

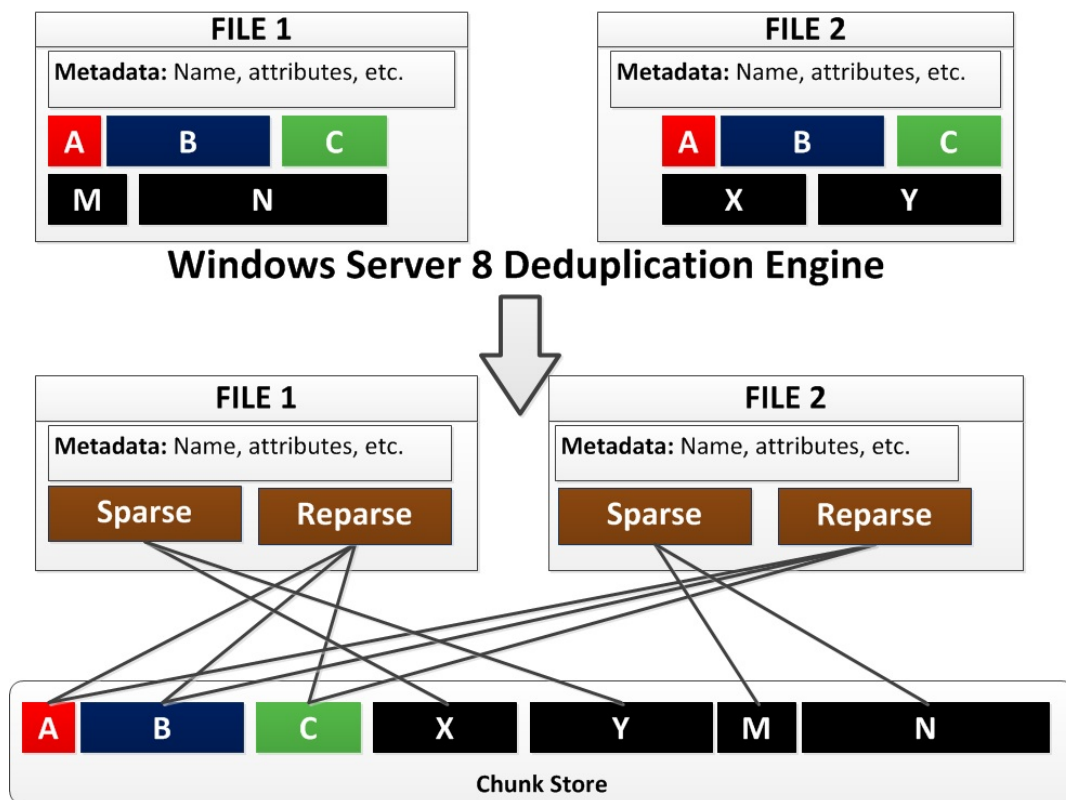
Windows Server 8 data deduplication: What you need to know

By Rick Vanover in The Enterprise Cloud, September 14, 2011, 5:05 AM PST // RickVanover

Data deduplication with Windows Server 8 provides storage space savings for Windows Servers running NTFS volumes. Read an overview of this feature.

Many of the new features in [Windows Server 8](http://www.microsoft.com/en-us/server-cloud/windows-server/v8.aspx) (<http://www.microsoft.com/en-us/server-cloud/windows-server/v8.aspx>) will deliver immediate benefits to server administrators. One of those features that I learned about at the Windows Server 8 Reviewers Workshop last week is a deduplication engine that can be enabled for NTFS volumes. [Deduplication can be done in various ways](http://www.techrepublic.com/blog/networking/three-modes-of-data-deduplication-how-do-you-decide/3869) (<http://www.techrepublic.com/blog/networking/three-modes-of-data-deduplication-how-do-you-decide/3869>) and, in the case of the Windows Servers implementation, it is a post-process deduplication that is enabled per volume and can be done through PowerShell.

The basic construct of Windows deduplication is that it is a variable chunk that ranges between 32 KB and 128 KB. Chunks that are a duplicate are copied to a chunk store that is managed by Windows and kept in the System Volume Information section of the disk; this means we can't see the inner workings of the deduplication engine. **Figure A** shows Windows being applied to two files. **Figure A**



(<http://b2b.cbsimg.net/blogs/sept2011-windows-server-tips-blog-tip1-figa.jpg>)

In Figure A, data blocks A, B, and C are deduplication candidates. When the deduplication engine runs, an eligible file has its deduplication blocks copied into the chunk store. From there, the file has two dimensions of its data: a spare and a reparse region. The reparse regions call to the chunk store to access the common chunks or deduplicated data.

The deduplication processes are run via Windows scheduled tasks or can be run interactively via PowerShell. The *Get-DedupStatus* command will quickly show a percentage of deduplication on a single volume. These scheduled tasks will scour the volume for the deduplication candidates and then coordinate movement into the chunk store. Windows Server 8 data deduplication is not allowed to run on C:\ drives. When I spoke with Microsoft product managers about the reasoning for this restriction, I was told that it was for experience of the system to be kept high. This is because any deduplication engine has overhead, no matter what any vendor says about it.

My observations on the deduplication process for Windows:

- If a file is fully deduplicated, it will only consume 4.00 KB on disk. This is only for the metadata sections of the files, and all of the reparse (and no parse) regions would be in the chunk store.
- When you click a file (even those not on a deduplicated volume) in Windows Explorer, the size and size on disk values may have some new behavior. These two values will show the fully hydrated size of the file and the deduplicated size consumption (outside of the chunk store).

Does the Windows Server 8 deduplication feature interest you? Share your comments.



About Rick Vanover

Rick Vanover is a software strategy specialist for Veeam Software, based in Columbus, Ohio. Rick has years of IT experience and focuses on virtualization, Windows-based server administration, and system hardware.

You May Also Like



[10 Celebrity Couples with a Large Age Gap](#)
(Celebrity Gossip Answers)



[The Gmail Trick They Don't Want You To Know About](#)
(The Good, The Bad & The Techie)



[These Are The 10 Sexiest CEOs](#)
(The Richest)



10 Online & Surprisingly Quick Degree Programs
(Education Portal)



Larry Ellison Doesn't Get the Cloud: The Dumbest Idea of 2013
(Forbes.com)



13 Amazing Uses for WD-40
(Reader's Digest)

Recommended by

Add your Comment

Free \$50 vnCloud Voucher

 vaultnetworks.com

\$50 to Test out our Cloud Server. Easy to Scale & Deploy in Minutes

Join IT's Best

Austin Data Center

Compare Visio & SmartDraw

Storage Optimization


Looking At Storage & SDN?

IT Consulting Services

IronBrick and NetApp

CrashPlan Backup

Top 10 Best Cloud Storage



Editor's Picks



Fine tune Mavericks' iCloud settings



AWS EC2 cloud automation tricks



The 15 worst startup names of all time




Cracking Open the Samsung Galaxy S5

White Papers, Webcasts, and Downloads

WEBCASTS // FROM DELL SOFTWARE

On-demand Webcast: The Next Generation Firewalls



Don't miss this interactive webcast about the changing nature of network threats and what your firewall absolutely must be able to do to keep you safe.

[VIEW THIS NOW](#)

WEBCASTS // FROM SAP

On-demand Webcast: Behind the Scenes in Mobile ROI

WATCH NOW

WEBCASTS // FROM UNITRENDS

Live Webcast: Smart Strategies for Data Backup and Recovery

WATCH NOW

WHITE PAPERS // FROM HEALTHTECH HOLDINGS

Three Ways Clinicians Can Leverage a Patient Portal to Craft a Healthcare Community

FIND OUT MORE

Don't Miss Our Latest Updates**Editor's Daily Picks**

Subscribe

Week in Review

Subscribe

**Mobile Device
Management** maas360.com/Free-GuideFree Smartphone Security Guide For Setting
Corporate IT Policies.**如果你在国外，孤单
一个人****How To Fix
"Missing.dll"****Resource Center**

Useful content from our premier sponsors

See what best practices for SMBs are emerging. Visit ZDNet today!

\$50 to test our vnCloud

 vaultnetworks.com

Scale & Deploy in Minutes! IaaS
Secure & Reliable. 99.999% Uptime