

The Story of 3.0 File and Folder Excludes

David Wartell on the Story of File and Folder Excludes in 3.0 CDP

This summer of 2008 was exciting. A Hurricane ripped through our Houston Texas headquarters knocking out power for two weeks and during the same two weeks we completed new CDP device drivers for Windows and Linux. These new device drivers are the special low level "sauce" that makes it possible for us to include and exclude files and folders and still use block level backups and a near-Continuous method for computing deltas.

Everyone Said Block Level CDP + File and Folder Include/Excludes Was Impossible! - Including Me!

Doing backups at a level lower than files and folders is a real challenge. You don't look at the files and folders! This means incremental backups are very very fast and they scale regardless of file counts on your server. If you don't know what we mean by block level backups see these technical papers:

- [Block Based Backup Technology](#)
- [Computing Deltas - near-Continuous \(CDP\)](#)

Customers kept telling us that they REALLY LIKED our near-Continuous backup method. They all agreed that is was the future of backups not only in hosting but in the enterprise as well. They liked how our 2.0 CDP product shrunk backup windows down to a few minutes. They also liked the advantages of disk images and bare-metal restore.

But at the same time some of our customers were reluctant to give up the ability they currently had to choose what files and folders should be included in their backups, a feature that they had with their existing legacy file-based backup applications that often took hours or days just to complete an incremental backup.

Our customers were telling us they wanted:

1. Block level backups using our near-Continuous (CDP) backup method that drastically reduced backup windows down to minutes
2. The option to do a bare-metal restore with a disk image
3. AND The flexibility to select what files and folders were included in their backup sets just like they had always done with their [Legacy File-Based backup software](#).

Geez talk about wanting to eat your cake and have it too!

Everyone including myself said it was not possible or too hard and too prone to error. How could we exclude files and still have a consistent low level file system image? We always thought that to gain the advantages of CDP and disk imaging you had to back up a Disk or Volume all or nothing.

A Very Simple Idea

It was spring of 2008. Late one evening while sitting around with Mike B. and Brian V. (aka Brain) in our development room tossing around ideas Brain said, "David your ideas are all too complex. If we exclude files from the block level backups our method has to produce a perfect file system image. And it has to be easy so we don't mess it up. File systems are way too complicated for us to mess around with them excluding blocks. We need to let the O/S do the excluding for us."

But how would this be possible I thought? In Linux we take a point-in-time block level consistent snapshot of the disk while the server is running using the R1Soft CDP device driver. And in Windows we combine our [near-Continuous delta computation method](#) with Microsoft Volume Shadow Copies. Microsoft VSS give us consistent block level snapshots. We combine VSS with our own Volume filter driver in Windows to tell us what the deltas are in-between backup sets before the backup starts. This means the difference between a backup window that takes 8 hours vs. one that takes 8 minutes.

I knew that Windows standard Win32 APIs can tell us what blocks or sectors belong to a file. Windows made this API for file system defragment applications. We could exclude the blocks belonging to the files that the user wanted to exclude and the Win32 API could tell us what those blocks were. The challenge with this way of doing things is that it leaves the file system in a broken state. I didn't like this and neither did Brain.

Brain continued, he said, "You know. If we could some how just make a Volume Shadow Copy WRITABLE it would all be fairly simple. The same for our Linux snapshots", he continued. "If our snapshots were writable we could just use standard file API calls like delete or unlink to remove the files we didn't want. Then the O/S (Windows or Linux) would take care of modifying the file system to remove the unwanted files from the snapshot image in a way that was perfectly consistent and it would all be very fast. If the user wanted to exclude an entire directory for example we could just recursively delete that directory from the snapshot. This would be very fast since delete operations involve VERY FEW disk writes. Since the actual writes to disk are so few when deleting even large sets of files we could potentially store all of those Writes in RAM making the exclude operation fly."

Well great I thought. All we have to do now is re-write our Linux CDP snapshot device driver so it could some how make a snapshot writable. AND some how we had to trick Windows into making a Volume Shadow Copy writable! Oh My. Not trivial.

It Might be Possible in Windows

The next day I emailed our Windows device driver wizard. I told him we needed to somehow trick Windows (XP through 2008) into making its read-only Volume Shadow Copies writable! The email I got back said, "Wow that sounds like a challenge. It would be amazing if we could do that but I don't know if it's possible. Let me do some research and I'll get back to you."

Two months later he got back to me. He said there was still a lot of risk but that he thought he may have found a way to do it. It would require a new device driver capable of hooking-in below the file system (in the VSS snapshot) and above the Virtual Volume Shadow Copy Disk.

What About Linux?

In the spring we had already begun work on a new Linux Device driver. We decided that we would make it a requirement to incorporate writable snapshots into our new Linux CDP device driver.

Two new Drivers and a Hurricane

In September of 2008 a category-2 hurricane went ripping through downtown Houston. Power was out for a month in some areas and curfew was declared in the city of Houston with the downtown area looking like a war zone with all of the damage. R1Soft scrambled to set up an emergency base of operations at the Hotel Zaza and employees gathered at the homes of other employees with power to keep R1Soft open for business.

During all of the madness I got an email from our Windows device driver wizard. Writable VSS Snapshots were Alive! It took 2 months longer than planned but they were here and worked well. Not only could they handle the light disk I/O required for large scale deletes they could handle large scale writes almost as efficiently as the real disk. We now had writable Volume Shadow Copies!

Even better I learned that because of where we hooked into the Windows kernel this new Windows driver could be installed without a reboot. Even better I thought!

Two weeks after the hurricane I got more good news. Mike B. had completed our new Linux CDP device driver. Written from the ground up it had massive performance and scalability gains over and above our 2.0 CDP driver. We learned from all of our mistakes and did it all over again. Even better it delivered a virtual block device that was readable and writable! An instant point-in-time copy of hard disk while Linux was running! It was a virtual block device that looked just like the real disk and we can read and write to it. Not only that the new virtual hard disk was lightning fast. Twice as fast in some situations as our 2.0 Linux driver.

What's Left?

As of October 2008 we still have some work ahead of us before we are done with incorporating file and folder excluded in our 3.0 product. Now that we have the special device drivers we are working on incorporating them in our 3.0 CDP product. Thanks to the magic of our device drivers and writable snapshots it's almost as easy as find, then delete!