

Disk Safe Best Practices

The all new CDP Disk Safe is built with industrial grade protection from crashes and power loss. The CDP Disk Safe is highly reliable and robust. It has automatic mechanisms that protect Terabytes of archived data from crashes and power failures not found in any other disk to disk backup software.

This is accomplished using an atomic write journal. Before any changes are made to a Disk Safe file, a new on-disk journal file is created. Before any Disk Safe pages are altered, the original content of the Disk Safe pages are written to the journal file. This allows any transaction to be completely rolled back even if interrupted by a crash or power failure. The Disk Safe assumes that the operating system will buffer writes and that a write request will return before data has actually been written to disk and that write operations will be reordered by the operating system. For this reason, the Disk Safe performs [FlushFileBuffers\(\)](#) on Windows or [fsync\(\)](#) operation on Linux at key points.

[Always Use Stable and Reliable Storage | Standard and Advanced Editions - Store Disk Safes on a Different Disk | Only use USB Drives to Transport Disk Safes | 32-Bit or 64-Bit for CDP Enterprise Servers? | Windows Best Practices | Linux Best Practices | Defragment File Systems | Avoid Disk Safe Vacuum | Corruption of Data Archived in the Disk Safe | Data Corruption in Your Environment](#)

Always Use Stable and Reliable Storage

Tips for selecting stable and reliable storage:

- If you are using hardware RAID, it is recommended to use a Battery Backup Unit (BBU) for the RAID card or RAID system.
- If you do not have a BBU for your RAID system, it is recommended to disable the write back cache. Consult the documentation for your RAID system.
- IDE drives are not recommended. It has been reported many IDE drives ignore signals to flush their on-disk cache.
- If using network attached block storage such as a: iSCSI or ATA-over-Ethernet, it is highly recommended you thoroughly test your storage environment for its ability to handle failures. In particular network level failures. It has been reported by our customers that these types of storage systems are the most prone to hardware and network faults.
- Avoid budget RAID cards. It has been reported that frequently entry level RAID cards can have serious performance and reliability issues. At the bottom end are some budget RAID cards that do not even perform hardware RAID and instead are simply a SATA card with a software RAID driver.

- USB and other portable media are great for temporary transport of a Disk Safe from one location to another. USB and other portable media should not be used for regular storage of Disk Safes. These types of inexpensive storage have been reported to ignore signals to flush their on-disk cache. In addition, their portable nature can cause disconnects or power failures simply by nudging a cable.
-

Standard and Advanced Editions - Store Disk Safes on a Different Disk

It is recommended that when using Standard or Advanced Editions that you store your Disk Safes on a different physical disk from the disk or disks you are replicating. While CDP does support writing the Disk Safe to the same disk and even the same partition and file system you are replicating it is NOT recommended. Replicating a drive to itself can cause severe disk saturation and performance issues.

Windows - Example Recommended Configurations

- Replicate C:\ a partition on a single SATA disk to E:\ a partition on a different physical SATA disk.
- Replicate C:\ to a network file system (CIFS).
- Replicate C:\ a partition or mapped LUN on a hardware RAID controller to E:\ a different physical SATA disk.
- Replicate C:\ a partition or mapped LUN on a hardware RAID controller to a network file system (CIFS).
- Replicate C:\ a mapped LUN on a hardware RAID controller to E:\ a mapped LUN on a different RAID controller.

Linux - Example Recommended Configurations

- Replicate `/dev/sda3` a partition on a single SATA disk to `/dev/sdb3` a partition on a different physical SATA disk.
- Replicate `/dev/sda3` to a network file system (NFS).
- Replicate `/dev/sda3` a partition or mapped LUN on a hardware RAID controller to `/dev/sdb3` a different physical SATA disk.
- Replicate `/dev/sda3` a partition or mapped LUN on a hardware RAID controller to a network file system (CIFS).
- Replicate `/dev/sda3` a mapped LUN on a hardware RAID controller to `/dev/sdb3` a mapped LUN on a different RAID controller.

Windows - Not Recommended Except for Testing purposes

- Replicate C:\ to a Disk Safe stored on C:\.
- Replicate C:\ a partition on a single SATA disk to D:\ a different partition on the same physical SATA disk.
- Replicate C:\ a partition or mapped LUN on a hardware RAID controller to E:\ a different

partition or mapped LUN on the same physical RAID controller.

Linux - Not Recommended Except for Testing purposes

- Replicate [/dev/sda3](#) to a Disk Safe stored on [/dev/sda3](#).
 - Replicate [/dev/sda3](#) a partition on a single SATA disk to [/dev/sda5](#) a different partition on the same physical SATA disk.
 - Replicate [/dev/sda3](#) a partition or mapped LUN on a hardware RAID controller to [/dev/sdb3](#) a different partition or mapped LUN on the same physical RAID controller.
-

Only use USB Drives to Transport Disk Safes

Use USB drives to transport Disk Safes. For example, a USB drive is excellent for moving a Disk Safe from your office to your data center. Or between two CDP Servers connected by a slow network connection. When using a USB drive, follow the steps to safely remove the device from your computer.

R1Soft does not recommend using a USB drive as the every day, primary storage for Disk Safes. USB drives tend to be unreliable and it has been reported that many USB drives ignore [FlushFileBuffers\(\)](#) and [fsync\(\)](#) requests to clear hard disk cache.

32-Bit or 64-Bit for CDP Enterprise Servers?

Always use a 64-Bit CPU(s) and 64-Bit operating system for CDP 3 Enterprise whenever possible. The main benefit of a 64-Bit environment is the CDP Server process will not be limited to a 32-bit address space.

If you are using a 32-Bit server, each processes address space is limited to less than 2 GB, even when Physical Address Extensions ([PAE](#)) are enabled. The CDP 3 Server is a single multi-threaded process Each thread will readily execute on a different CPUs/cores at the same time however all the threads share the same virtual address space limiting the total memory a CDP Server on 32-bit to less than 2 GB. Why not 4 GB? Windows and Linux virtual memory space is actually 31-bit (2 GB) because the last bit is reserved. Also some portions of the 31-bit virtual address space is further reserved by the operating system. Typical usable virtual address space on 32-bit is somewhere between 1.5 - 1.8 GB depending on the environment.

Windows Best Practices

- Schedule a weekly defrag of the file system where Disk Safes are stored.
- Install CDP 3 Enterprise Edition server on Windows 2008 R2 if possible. Windows 2008 R2 has highly optimized disk I/O performance.
- Store Disk Safes on a NTFS file system formatted with a 32 KB cluster size. 32 KB

matches the internal "page" size used by the Disk Safe when reading and writing from disk

- To store Disk Safes on a network file system, use a Windows Storage Server or Windows Server O/S for your file server. A reliable Network Attached Storage Device supporting the CIFS protocol is also recommended. Using Linux SAMBA software for network file storage of Disk Safes is not recommended.
-

Linux Best Practices

- If your kernel version has stable support for Ext4, then Ext4 is highly recommended for performance over Ext3 for Disk Safe storage.
- If using Ext4, make sure you are using Ext4 extents. Extents help reduce file system fragmentation.
- Add the mount option `barrier=1` to `/etc/fstab` on Ext3 file systems where Disk Safes are stored OR disable your storage controller's write cache. Ext3 does not check sum the journal. If `barrier=1` is not enabled as a mount option (in `/etc/fstab`), and if the hardware is doing out-of-order write caching, you run the risk of severe file system corruption during a crash.

For more information see:

<http://blog.nirkabel.org/2008/12/07/ext3-write-barriers-and-write-caching/>

- Storing Disk Safes on Ext3 can cause performance problems. The Disk Safe periodically (some times frequently) use `fsync()` to force changes to the file through to disk hardware to ensure data is protected from a power loss or crash. Ext3 has a well known broken implementation of `fsync()` causing a flush file buffers on one Disk Safe file to sync the entire file system to disk. This can have a severe performance penalty on Disk Safe I/O and writes to one Disk Safe can degrade reads of other Disk Safes due to Ext3's well known `fsync()` issues.
 - If storing your Disk Safes on Ext3, it is recommended that only other Disk Safes are stored on the file system due to Ext3's problematic `fsync()` implementation. Do not store applications and daemons on a Ext3 file system with your Disk Safes if you can avoid it.
 - When using Ext3 or Ext4 make sure you accept the default block size (4 KB). Forcing smaller block sizes will decrease performance for the Disk Safe's large files and can limit max file size to as little as 16 GB which will likely cause 'File too Large' errors during initial replicas.
 - Use NFS to store Disk Safe on a network file system. If using NFS, it is recommended to use the latest stable NFS versions and latest stable Linux kernels.
 - Never export NFS file systems with asynchronous writes enabled (`async` option). Exporting NFS with the `async` option can cause data corruption in the event of a failure.
-

Defragment File Systems

The CDP Disk Safe is a storage system for long-term archiving of unique block-level (units of data below the file system) deltas (small differences in data). CDP extends the Disk Safe files

and makes write in predictable 32 KB increments. This helps the modern Windows and Linux file systems pre-allocate space for the Disk Safe files and naturally reduce file fragmentation. As the Disk Safe block stores (.db files) for each device, you are protecting age some data inside those files remains forever and some is marked deleted and recycled for new deltas as old data is merged out over time. If there is no recycled free space inside of the file for new deltas, the file is extended on disk. This can happen anytime a new recovery point is created.

Generally, any files that remain on Disk for a long period of time and continue to be written to and extended in size are subject to [file system fragmentation](#). This is true for the R1Soft Disk Safe and other long-term storage mechanisms, for example, a relational database like MS SQL Server and MySQL.

R1Soft recommends for optimal performance that you periodically [defrag](#) your file systems where Disk Safes are stored. If it is feasible in your environment a weekly file system defrag would be optimal.

Windows

We recommend [contig](#) by MS Sysinternals to analyze the fragmentation of specific files and folders and [Auslogics Defrag](#) for regular scheduled file system defrags


1. Determine how fragmented your disk safe files are.

- Install [contig](#).
- Run [contig](#) with `-s` and `-a` options and give contig the path to the folder where your disk safes are stored like:

```
contig -a -s D:\PATH\TO\YOUR\DISK\SAFES
```

This may take a while.

2. Use Auslogics to schedule weekly file system defragment tasks on your CDP Server.



Linux

Unfortunately, the Linux operating system is lacking stable online file system defrag capability.

Here are your options:

1. XFS file system has defrag capability.

How to use XFS defrag <http://www.linux.com/archive/feature/141404>.

Tutorial for XFS on CentOS <http://blogwords.neologix.net/neils/?p=1>.



Note on Linux XFS

It is not possible to use Multi-Point Replication with XFS.

2. Ext4 promises Defrag in the future.

Ext4 has an experimental online defrag capability. Eventually, we all expect this to become stable for production servers.

Ext4 has extents which help reduce file fragmentation. Extents do not eliminate file system fragmentation and Ext4 can still benefit from defrag just like NTFS on Windows.

Be aware the consensus is ext4 online defrag is NOT ready for production systems.

Here is a thread in ubuntu bugs about the topic of ext4 online defrag:

<https://bugs.launchpad.net/ubuntu/+bug/321528> (look for I_KNOW_E4DEFRAG_MAY_DESTROY_MY_DATA_AND_WILL_DO_BACKUPS_FIRST).

2. Offline Defrag for Ext3 and Ext4.

Is it possible to perform an offline defrag of any Linux file system. Due to the work involved, it would be recommended to perform this task only once or twice a year.

1. Shutdown your CDP Server:

```
/etc/init.d/cdp-server stop
```

2. Add an intermediate disk available capable of holding all of your Disk Safes (could be network storage).
3. Copy all of your Disk Safes to the intermediate storage:

```
cp -af /disk/safes/* to /mnt/storage/
```

4. Re-format the file system that has the primary copy of your Disk Safe:

```
mkfs /dev/YOUR_DEVICE
```

5. Copy all of the Disk Safes files back:

```
cp -af /mnt/storage/* /disk/safes/
```

6. Start the CDP Server:

```
/etc/init.d/cdp-server start
```



Note on Linux Defrag

There appears to be a widespread fallacy that somehow Linux file systems (ext2, ext3, and ext4) are magically immune to fragmentation. This could not be further from the truth. It's true that the kernel does what's called pre-allocation as you are writing to a file. It will attempt to notice that an application keeps extending a file and allocate contiguous blocks ahead of what it is doing.

Windows also does this for NTFS. No matter how fantastical pre-allocation or the file system may be files get fragmented and you need a way to pack them. Linux file systems are no exception.

Avoid Disk Safe Vacuum

For best performance only vacuum your Disk Safes when you absolutely must reclaim unused storage space. For more about Disk Safe vacuum, see: [Vacuuming Disk Safes](#).

If you are a service provider using CDP 3 Enterprise edition as a multi-tenant system use [Volume quotas](#) based on "Size of Deltas in Disk Safe" instead of "On Disk Size". This way your customers only pay for what is stored in the Disk Safe instead of the on-disk foot print which includes the unused parts of the Disk Safe file being recycled for future deltas.

Corruption of Data Archived in the Disk Safe

The CDP 3 Disk Safe is highly reliable and robust. Even with industrial grade protection, there are still ways for your data to become lost or damaged beyond repair. If any of the following events occur, you may corrupt your Disk Safe. If your Disk Safe becomes corrupted, you may lose all or some of your archived data.

If a CDP 3 Disk Safe is corrupted by any one of the events below, it can not be repaired.

- Delete a file in the Disk Safe Folder.
- Make an incomplete copy of the Disk Safe folder thereby corrupting the copy.
- Making a copy of the Disk Safe files when the Disk Safe is open and being written to by the CDP Server thereby corrupting the copy.
- A hardware or O/S fault causing incorrect data to be written to Disk Safe files.
- A faulty hard disk or storage controller failing to flush volatile cache when requested can break protection from unclean shutdowns and power failures.
- Rogue process writing to the Disk Safe files.
- Soft Linking any files Inside of the Disk Safe Folder. If the block deltas store and its associated write journal (created at run time) end up on different file systems data loss can occur if there is a crash or power failure.
- Failure to store the Disk Safe on a journaling file system can cause the write journal to be lost or moved to lost+found. If this happens the Disk Safe may likely become damaged beyond repair.



Note

Windows NTFS and Linux Ext3, Ext4, XFS, and ReiserFS ARE Journaling File Systems.

- NFS (Linux / Unix Network File System) faults or bugs.
If using NFS on Linux, R1Soft recommends you use the latest available NFS versions and latest stable Linux kernels, and do not export NFS file systems with asynchronous writes (async option).
-

Data Corruption in Your Environment

- If your file system is corrupt at the time you replicated it, this does not mean your Disk Safe

is corrupt. Instead it is corruption that was on your primary storage when CDP replicated it to the Disk Safe. If this happens the corrupt files will likely be unrecoverable.

- Use caution if you are replicating a server that has file system or disk subsystem warnings or errors in the event or system logs. Be aware it is possible for a working server to have file system corruption. Key file system data structures may be damaged on disk and loaded in memory by the operating system causing a corrupt file system to work until it is rebooted. If this happens your files were corrupt when replicated and you should understand you may not be able to restore them from the CDP replica.
- It is possible for hardware or O/S faults to cause incorrect data to be read from the Disk Safe at the time of restore even though the data on disk media is correct. In these cases, some or all of the archived data may appear damaged beyond repair until the fault is corrected. Determining if the fault is occurring only on read or if the data is damaged on media may be difficult or impossible. Furthermore, no amount of checking or validating of data in the Disk Safe can prevent or pre-warn these kinds of faults.