UNIVERSITY OF TORONTO | Information Security

Guideline on resilient backups

Office of the Chief Information Security Officer

Effective: Oct. 3, 2025

Table of contents

Overview

Preamble

Data on systems can sometimes be lost unexpectedly, due to system failure, user error, or malicious action. Backups are copies of the data, kept in a suitably safe place to ensure that if the data were to be lost, it can be restored from a backup. A backup is resilient to a particular kind of data loss event when it is designed to survive that type of event. For instance, a backup is resilient to ransomware if ransomware cannot access it.

Purpose

This guideline is intended to inform about the need for suitable backups and provide guidance about how to achieve them. Simple, straightforward language is used, and technical terms are avoided as much as possible.

Audience

The entire University of Toronto community, including faculty, staff and students.

Scope and applicability

This guideline addresses any situation where data is held on a computer or information system, where that data could be lost due to unforeseen circumstances, and where the impact of that loss should be mitigated by a backup.

Guideline

Data on systems can sometimes be lost unexpectedly due to system failure, user error or malicious action. This could be disastrous. Copies of the data on the system are needed if this happens, so that the data is not lost. These copies are called “backups” and making them is called “backing up your system”.

How often you need to make a backup

If you lose your data, the only copy you have is your most recent backup. This means you need to do backups often enough so that you can live without the data that was created since the last one was made. If you can afford to lose a week, you will need to make backups every week. If you can only afford to lose a day, you will need to make backups each day. This assessment is sometimes called a *Recovery Point Objective*, or RPO.

How much of the data you need to backup

Prioritize backing up data that you cannot afford to lose and cannot recreate. If you can recreate your data, consider whether the cost of recreating it still warrants backing it up. If most of the data on your system doesn’t change very often, you may not need to back up all of it each time, but only what has changed since a previous backup (this is called an “incremental” backup). But you must make sure that you do have some backups of all your data, recently changed or not (this is called a “full” backup). If you choose to make incremental backups, be careful not to force yourself to have to go too far back in time to retrieve data from too many backups if you have to recover your data from backups, because this makes recovery take longer.

Data vital to digital assets that have been assigned to high availability levels (as defined by the University’s Digital Asset Classification Standard) must be prioritized for backup, to ensure that the availability of the asset can be preserved.

Where you store your backup

Your backups need to be kept in a safe place. A place is “safe” when it is not accessible to potential attackers, it cannot be destroyed by disasters such as fires or floods, and it preserves the privacy and integrity of the data. Consider also how long it will take you to retrieve your backup if you need to recover the data, because it will affect the time it takes to recover from a data loss event. Shipping your backup media to a secure warehouse or making an encrypted copy to a cloud service may keep it safe, but if it takes a long time to ship it back, or if you are trying to retrieve large amounts of data from a cloud service over a slow network connection, it will impact your recovery time.

How you store and protect your backup

The best way to protect backups from malicious attackers, such as criminals deploying ransomware, is to ensure they are off-line. As an alternative to off-line backups, some commercial backup vendors offer *immutable* backups. These enable clients to create backups over the Internet, but the backups cannot be changed after they are created. While immutable backups are still vulnerable to attack at the service provider site, they do protect clients from attacks at their own sites, and they may be easier to implement than off-line backups. Off-line or immutable backups are sometimes called *ransomware-resilient* backups because they are inaccessible to ransomware.

The best way to protect off-line backups from fire and flood is to ensure they are off-site, typically in a separate building. Some types of data loss events (e.g. tornado, riot) may affect multiple near-by buildings, so do consider the distance between buildings when assessing risk.

The best way to protect the integrity and privacy of off-line backups is to ensure backup media are appropriately encrypted at all times. In general, backups need to be protected according to the University's Information Security Standards, particularly the controls in the Media Protection and Access Control sections.

Making sure your backup is OK

After you make a backup, you need to read it back to make sure it has been properly written. If you write backups to the same piece of media at different times, be aware that a problem on the machine being backed up could damage previously created backups on that media, so try to avoid using a single piece of media if at all possible. It is a good practice to test your backups yearly to make sure they are still usable.

How long do you keep your backups

Make sure your backups are kept long enough so that if you need the data as of the day they were made, that you still have them. But having backups is not a substitute for archiving. You still need to follow appropriate data retention policies.

How to put backups to use

There are two main uses for backups. One is to recover an entire system, when things go badly wrong. The other is to recover some data from the system when needed (e.g. if someone accidentally damages or deletes something). Make sure your backups are good for both purposes.

What to do

If you have access to a file storage service that is equipped with resilient backups, then use it. Any data you store there will be properly backed up as part of the operation of the service.

If you are responsible for doing backups, here are some options in order of preference:

1. Use a resilient backup service endorsed by the University (e.g. UTORrecover)

2. Leverage a commercial backup service whose risks have been assessed by the University as being acceptable for the type of data you are backing up.

3. Use appropriate (preferably university-assessed and approved) backup software writing to local protected media and/or cloud storage. Media can be protected through secure encryption and/or lock & key. One prudent approach for local media is to have a 3-2-1 strategy, which is to have at least three copies, of which at least two are on separate media and at least one is offsite.

What to talk to

Start with your local IT help desk because they will know best the services accessible to you in your unit. When seeking to better understand the risks that backups are intended to mitigate, please do not hesitate to reach out to the University's information security team, who are experts in the sorts of risks that backups are designed to mitigate.

References

Information Security Standards

The University of Toronto Information Security Standards govern the design and operation of computer systems. Backups are generally a highly necessary component of the "Recovery" aspect of Incident Response. Backups themselves must be protected in accordance with the Standards, particularly Media Protection and Access Control.

Digital Asset Classification Standard

The University's Digital Asset Classification standard governs the classification of digital assets by availability level.