Ransomware-resistant backups

Principles for making on-premises and cloud backups resistant to the effects of destructive ransomware.

Principles for ransomware-resistant cloud backups

Helping to make cloud backups resistant to the effects of destructive ransomware.

Principle 1: Backups should be resilient to destructive actions

Threat: Ransomware attacks may seek to frustrate or prevent effective recovery by destroying backups. The backup service should therefore be resilient to attempts to destroy backup data and should also protect that data from malicious editing, overwriting or deleting.

Suggested implementations

- **Blocking any deletion or alteration requests for a backup once it is created.**
  If there are no mechanisms in place for backups to be altered or deleted, a malicious actor has no way to delete backup data. In practice, backup data can't be stored forever, so system owners should be able to set policies in advance that specify how long a backup should be kept. This should align with the backup schedule, and may be different for different types of data.
- **Offering soft-delete by default.**
  A soft-delete mechanism flags data as inaccessible, meaning the wider system behaves as if the data has been permanently deleted, while in fact it is still recoverable for a certain period, perhaps in a separate storage area reserved for this purpose. This means that if a malicious actor deletes useful backup data, the customer should still be able to recover it. Customer accounts – and therefore also

attackers – shouldn't be able to delete or overwrite data from the storage area where soft-deleted data is held.

For soft-delete to be effective, the system owner needs to monitor the overall health of the backup system because if an attacker can soft-delete all backup data and the system owner doesn't notice until after the review period has ended, the attacker can subvert this.

- **Delaying implementation of any deletion or alteration requests.**
  This should be for a pre-agreed period of time, depending on the system owner's monitoring schedule. This will only be effective if an alert is also raised, and the system owner is confident that that alert will be successfully delivered even if their infrastructure is compromised.
- **Forbidding destructive requests from customer accounts.**
  This includes both user and machine identities, such as administrative users, backup agents, security scanning tools and protective-monitoring connections. In this scenario, all exceptional destructive requests must be authorised out-of-band using a pre-agreed mechanism between the customer and the backup service. This means that an attacker who has compromised a system can't make destructive requests without compromising another separate system.

Principle 2: A backup system should be configured so that it isn't possible to deny all customer access

Threat: If an attacker can stop a victim organisation from accessing its own backup data by disabling or deleting all customer accounts or corporate identities, they won't need to do anything as destructive as deleting the data itself.

The backup system should therefore be configured so that it isn't possible for an attacker to deny all customer access, either by deleting the individual accounts used to access backup data, or by deleting the entire customer account. This also includes setting identity and access management (IAM) policies to ensure this.

Suggested implementations

- **Allowing customer access to the backup service, even if all existing corporate IT systems and assets are unavailable,**
  by agreeing a separate out-of-band mechanism. This could be a separate authorised customer account and/or device, or a pre-shared passphrase that can be stored physically and authenticated by phone between the customer and the service.
- **Forbidding any IAM policy that restricts access to a single account within an attacker's control,**
  as this forces an attacker to undermine multiple accounts to achieve full control of the backup system.

Principle 3: The service allows a customer to restore from a backup version, even if later versions become corrupted

Threat: If an attacker can flood the backup store with corrupted backup data, they won't need to do anything as destructive as deleting the data itself.

The backup service should therefore allow customers to store backups for a retention period that aligns with their risk appetite, and system owners should monitor and test the state of their backups regularly.

Suggested implementations

- **Providing mechanisms so that system owners can test whether they can restore from the current backup state,**
  which can be on demand (in a way that isn't destructive to the company's existing infrastructure) and which should allow a system owner to detect if a backup has been corrupted. For this to be effective, the system owner should test this as part of a regular monitoring process.
- **Storing backup data according to a fixed time period,**
  rather than a fixed number of backups. This will prevent an attacker overwriting the backup store with a series of corrupted backups in quick succession.
- **Creating and retaining a version history**
  so that a system owner can restore from a version of their choice.

- **Offering flexible storage policies**
so that a system owner can decide how many backups to keep for different periods of time, according to their risk appetite. For example, a system owner might decide that they want to keep daily backups for a month, and monthly backups for a year.

Principle 4: Robust key management for data-at-rest protection is in use

Threat: If a stored backup is encrypted for data-at-rest protection, an attacker doesn't need to actually delete the data itself if they can simply delete or modify the encryption key.

Keys used to encrypt data at rest should therefore be protected, to make sure that backup data can be decrypted when necessary.

Suggested implementations

- **Following the NCSC's**
cloud key management guidance.
- **Offering an out-of-band key backup option,**
such as the option to commit a master key to paper in human-friendly text encoding or QR code form, so that it can be stored in a secure location, such as a safe.

Principle 5: Alerts are triggered if significant changes are made, or privileged actions are attempted

Threat: An attacker hopes that their attempts to compromise a backup won't be detected, since targeting a backup system can be a precursor to an attack on an organisation's main system.

If a significant change is attempted, a cloud backup service should raise alerts, and then initiate follow-on actions once those alerts are triggered. The service should offer different types of alert delivery mechanism, so that alerts can still be

received if the customer's infrastructure is compromised. Significant changes could include (but aren't limited to) mass deletion requests, backups stopping, altering global retention periods, changes to global encryption policies or changes to administrator account details. The alerts should be raised whether the attempts are successful or not.

Alerts are only effective if the customer initiates a follow-on incident management process once triggered.

Suggested implementations

- **The service offers a wide range of customisable alerts** for activity that affects the backup system that a system owner can ingest and monitor. For a larger organisation, this might be a SOC and for smaller organisations, it could be an automated email to a monitored group mailbox. Although fully customisable, alerts for significant changes should be switched on as default.
- **Significant changes to how the backup system behaves or is accessed require extra authorisation and should automatically initiate extra protective monitoring.** The NCSC guidance on logging and auditing administration activities may be helpful here.

PUBLISHED

22 November 2024

REVIEWED

22 November 2024

VERSION

1.0

WRITTEN FOR

Small & medium sized organisations

Public sector

Cyber security professionals

Large organisations