

One storage check a week, for a year

A calendar file with one storage, backup or recovery check each Monday for 52 weeks, each one doable in 15 minutes or less with what your team already runs.

A calendar file with one storage, backup or recovery check each Monday for 52 weeks, each one doable in 15 minutes or less with what your team already runs.

Work the checks in order. Each one takes 15 minutes or less of hands-on time; where a restore runs longer, start it, then come back to read the result. Week 1 sets up one notes file; every later week adds a few lines to it, so keep it where your team can find it.

The 52 checks

Know what you have and who owns it

Week 1 List your five biggest places data lives

Open a new notes file or spreadsheet and name it "Storage list". Put it where your team can find it, because every later week adds to it. Make three columns: system, where it sits, what it holds. Fill one row each for the five systems that hold the most company data (an array, a file server, a NAS, a cloud storage account), a few words per cell. Skip desktops and partner-hosted apps for now. Stop at 10 minutes. Five rows is a full week 1; add the rest a few at a time in later weeks.

Why: Every later check starts from this list. It is also the first thing an auditor or a new IT leader asks for, and five rows you can show beats a long list that never got finished.

Week 2 Name an owner for each system

Next to each system on your list, write one person who answers for it and one backup person. Where nobody comes to mind, write none and leave it there.

Why: A system with no named owner is the one that misses patches and tests. Who owns this is a plain question an executive can ask in any meeting.

Week 3 Match each system to its backup

For each system on the list, write which backup job covers it, or none found. Read it from the backup console, not from memory.

Why: Systems added since the backup plan was written are easy to miss. This shows what is covered as of today.

Week 4 Restore one file

Pick one file from last week's backup and restore it to a test folder, never over the original. Time it from the moment you start to the moment the file opens. Write down the time and any step that needed a password or a person you had to find.

Why: IBM's disaster recovery guide puts it plainly: an untested disaster recovery plan cannot be relied upon. The time you write down is a measured number to set beside the time people assume.

Week 5 Find the written steps

Look for the written steps for running storage and backup day to day: where they live, when they were last edited and who wrote them. If they live only in one person's head, note that.

Why: Written steps let someone else cover a sick day or a departure. Auditors ask for them by name.

Week 6 Before the holiday change freeze

List the changes planned before the year-end change freeze: firmware, patches, moves. For each one, write the way back if it goes wrong, and check a fresh backup of the affected systems is scheduled before the change.

Why: Changes squeezed in before a freeze often land when the team is thin. A written way back keeps a bad change small.

Week 7 Check that alerts reach a person

Find where storage and backup alerts are sent. Check that the address or channel still reaches someone who reads it, and send a test alert if the system offers one.

Why: An alert sent to an unwatched mailbox, or to someone who left, reaches nobody. A failed job then stays hidden until someone needs a restore.

Week 8 Holiday cover

Write down who covers storage and backup over the holidays and how to reach them. Have that person sign in to the backup console today to prove their access works.

Why: The person covering is often not the person who built it. A missing sign-in found today takes minutes to fix; found during an outage, it adds delay at the worst time.

Week 9 Quarter end: the backup history

Read the backup job history for the quarter. Count the failed and skipped jobs, and note any job that has shown a warning for weeks with nobody acting on it.

Why: A quarter of history shows patterns a daily glance misses. It also gives the executive one plain line for the quarter: what ran and what did not.

Week 10 New year: who has admin rights

List every account with admin rights on storage and backup systems. Mark any that belong to people who left or changed roles, and any login that several people share.

Why: Admin rights pile up over a year. CISA's ransomware guide notes that attackers hunt for stored credentials to reach backup systems, and advises separate admin accounts.

Week 11 Restore one whole system

Start a restore of one small virtual machine or server to an isolated test spot. Note every step you did by hand after it came back. The 15 minutes is your hands-on time; let the copy run and check on it later.

Why: A single file proves little about a full system. The hand steps you note are what make a real recovery slower than the plan says.

Week 12 Agree the recovery order

Write the three systems the business needs back first, in order, and what each one depends on, such as sign-in, network or a database. Ask one person outside IT whether they agree with the order.

Why: IBM's disaster recovery guide says mapping dependencies shows engineers the necessary recovery order. An order agreed now saves an argument in the middle of an outage.

Week 13 Expected time against measured time

Ask two business owners how long they think their main system would take to come back. Write their answers next to the times from your week 4 and week 11 restores.

Why: The gap between what people expect and what you measured is the most useful line you can take upward. It turns a vague worry into one plain conversation.

Make sure the way back works

Week 14 The copy kept apart

Check whether at least one backup copy sits where your everyday admin accounts cannot change or delete it: offline media, a locked copy that cannot be changed until a set date, or a separate account. Write down which one, and who can reach it. On IBM FlashSystem this is called Safeguarded Copy; other makers use their own names.

Why: CISA's ransomware guide advises offline, encrypted backups that are tested regularly. A copy that the same login can delete is exposed to the same mistake or attack as the original.

Week 15 Where the backup keys live

Find where the backup encryption keys or passwords are stored. Check that two people can get to them, and that they are not kept only inside a system the same backup tool protects.

Why: An encrypted backup without its key cannot be restored. It takes minutes to check and is easy to skip.

Week 16 How long backups are kept

Look at how long each backup set is kept. Set that beside any rule your organization follows, such as legal holds, industry rules or your own policy, and note the gaps.

Why: Retention settings chosen years ago may no longer match what the business or the auditors expect. A gap found by your team is easier to close than one found in an audit.

Week 17 Restore from the copy kept apart

Restore one file from the copy you named in week 14, not from the everyday backup. Time it, and note every person and password the restore needed.

Why: This is the copy you would lean on after a security incident. Knowing its steps now means nobody learns them under pressure.

Week 18 Capacity: today’s starting point

Write today's used and total capacity for each main array and each backup target, with the date. If you have a figure from a few months ago, put it alongside; if not, this becomes your starting point.

Why: A full array or backup target stops jobs. A simple trend gives months of notice, which is the time an executive needs to plan.

Week 19 Snapshot schedules

List the snapshot schedules on your arrays: how often they run, how many are kept and which volumes are left out. Check that at least one recent snapshot exists for each important volume.

Why: Snapshots are often the fastest way back from a mistake. A volume left out of the schedule goes unnoticed until someone needs it.

Week 20 Support status of each system

Write down the firmware or software level on each array, storage switch and backup server, and whether its maker still supports it. The maker's support site lists end of support dates.

Why: A system past end of support stops getting fixes. That date is a plain fact an executive can plan around.

Week 21 Who can delete backups

In the backup tool, check which accounts can delete backups or shorten how long they are kept. Note whether a delete needs a second person or a waiting period.

Why: CISA's ransomware guide notes that attackers try to find and delete or encrypt any backups they can reach. Fewer accounts with delete rights is a small change with a large effect.

Week 22 Quarter end: the one-page summary

Write a short summary of the year so far: systems listed, owners named, restores tested with their times, failed jobs and one open problem. Keep it to one page.

Why: One page is what an executive will read and repeat. Updated each quarter, it shows progress without a meeting.

Week 23 Emergency admin access

Check there is an emergency admin login for storage and backup that does not depend on your normal sign-in system. Confirm its password is sealed somewhere safe and that two people know where.

Why: When the normal sign-in system is down, the tools to fix it can be locked too. Microsoft, which sells the Entra ID sign-in service, advises emergency access accounts for exactly this case.

Week 24 Restore to different hardware

Restore one file or small system to different hardware or a different location from where it came. Note any driver, network or license step that got in the way.

Why: In a real recovery the original hardware may be gone or busy. Steps that only work on the original box show up here, while there is time to fix them.

Week 25 Single points in the storage network

Sketch the path from your main servers to their storage: switches, links and adapters. Circle any single switch or link that everything passes through.

Why: One shared link can take several systems down at once. It is quick to spot on paper and slow to find during an outage.

Week 26 Features you already own

List the features your storage and backup tools already have that are switched off or unused, such as snapshots, replication, reports, locked copies or alerts. Pick one to look into next quarter.

Why: A feature already licensed is the quickest improvement to make. The list also shows the executive what the current setup can do.

See how it runs day to day

Week 27 A normal day’s performance reading

At a busy hour, note the response time and throughput of your main array from its own dashboard. Save a screenshot with the date and time.

Why: Without a normal reading, it feels slow cannot be checked. A dated reading turns a complaint into a measurement.

Week 28 How long new storage takes

Take the last request for new storage and time it from the ask to the moment it was ready to use. Note each hand-off and each wait.

Why: Slow provisioning shows up as delays in other teams' projects. A measured time is easier to discuss than a feeling.

Week 29 Data nobody opens

Find the largest folders or shares nobody has opened in a long time. List the top few and ask each owner whether it is still needed.

Why: Old data still gets backed up and copied every night. Owners deciding is better than IT guessing.

Week 30 Restore a database and read it

Restore one database backup to a test server and run one simple query to confirm the data is there and readable. Time it.

Why: A database can restore without errors and still not open. One query shows the data came back, not only the files.

Week 31 Data held in cloud services

List the company data held in cloud services: file shares, business apps and cloud storage. For each one, write who backs it up, if anyone, and read the service's own terms on what it keeps and for how long.

Why: Cloud data is easy to leave off the list because it sits outside the server room. Who backs this up is a question an executive can ask of any service.

Week 32 Before summer leave

Check that someone besides the main admin has run a restore at least once. If nobody has, pair them up for the next restore test.

Why: A restore only one person can run is a risk the business may not know it carries. Holiday season is when it shows.

Week 33 Where the logs go

Check how long storage and backup logs are kept, and whether a copy lands somewhere an admin on that system cannot delete.

Why: Logs are how you find out what happened after an incident. An auditor will ask how long you keep them.

Week 34 What happens to old drives

Find out what happens to old drives and tapes when they leave service: who wipes them, how, and where the record is kept. NIST publishes a public guide on media sanitization you can work from.

Why: A written record of how media was wiped is what an auditor asks to see. It also answers the executive's question of where old data ends up.

Week 35 Quarter end: update the summary

Update the one-page summary from week 22: what changed, which restores you tested and their times, and the one open problem. Strike out anything that is now fixed.

Why: Two summaries side by side show direction, not only status. That is the line an executive repeats upward.

Week 36 Read one offline copy

If you keep tape or removable media, take one cartridge or drive from storage and check that it reads and that its label matches the catalog. With no offline media, read one file from the copy you named in week 14 instead.

Why: Offline media sits untouched for months. Reading one shows the drive, the media and the catalog still work together.

Week 37 Bring back a whole application

Restore one complete application, with all its parts, to a test area and ask one user to confirm it works. Count only your hands-on time; let the copy run.

Why: A system is only back when a user says it works. This test measures that, not only the copy time.

Week 38 What nobody monitors

List which storage and backup systems send health data to a monitoring tool and which do not. Mark the ones where a failing disk or a full volume would go unseen.

Why: A system nobody watches fails quietly. Knowing which ones are unwatched is the first step to fewer of them.

Week 39 Where tools do not connect

Note where storage, backup and security tools do not share information: separate consoles, separate alerts, separate accounts. Pick the one gap that costs your team the most time.

Why: Each separate console is another place to look during an incident. Naming the gaps is how a team decides which to close first.

Keep it current and close the year

Week 40 Support contacts that work offline

Check that the phone lines and support contract details for your storage and backup makers are current. Keep a copy where you can reach it if email and file shares are down.

Why: Opening a support case quickly depends on having those details at hand. CISA's ransomware guide advises keeping key documents in hard copy and offline.

Week 41 Power and clean shutdown

Check the battery status and last self-test of the UPS units behind your storage. Find out what the storage does when the power runs low, and whether it shuts down cleanly.

Why: An unclean power loss can damage data that backups then copy. A known shutdown path is a quick thing to confirm.

Week 42 Repeat the first restore

Run the same restore you ran in week 4, using the same kind of file. Time it and set it beside the week 4 time.

Why: Data grows and setups change. A repeat test shows whether recovery is getting faster or slower.

Week 43 The first calls in an incident

Write the first five calls the team makes if storage or backups are hit: staff, leadership, insurer and outside help. Check that each number still works.

Why: CISA's ransomware guide advises a regularly exercised incident response plan with a communications plan. A short call list is the part people reach for first.

Week 44 When things were last patched

Check when your storage, backup servers and storage switches were last patched, and whether the patch steps, including the way back, are written down.

Why: NIST's patch management guide treats patching as routine preventive maintenance and notes a patch may need to be reversed or restored from backup. Written steps make the next one routine.

Week 45 How far behind the second copy runs

If you copy data to a second site or to the cloud, check when the last copy finished and how far behind it runs. Read the time of the last good copy, not the setting that says it is turned on.

Why: Replication that lags or has stopped can still show as enabled. The lag is the data you would lose.

Week 46 A new person runs a restore

Ask the newest person on the team to restore one file using only the written steps. Note where they got stuck and fix the steps.

Why: Steps that only work for their author are not yet written steps. This tests the procedure, not the person.

Week 47 Systems with no owner or use

From your week 1 and week 2 lists, find systems or volumes with no owner or no recent use. List them for a decision; delete nothing this week.

Why: Unused systems still need patching, backup and licenses. A decision list keeps the cleanup safe.

Week 48 Quarter end: capacity and summary

Write today's used capacity next to the week 18 figures, then update the one-page summary. Note anything that will fill up before this time next year.

Why: Growth seen early gives the business time to plan. It turns a full array from a surprise into a scheduled task.

Week 49 Gather the year's evidence

Put this year's notes in one folder: restore times, access reviews, summaries and lists. Date each item.

Why: Audit questions get easier when the answers already sit in one place. It also shows the year's work at a glance.

Week 50 Recovery steps that work offline

Print the recovery steps for your top system, or save them somewhere that does not depend on your network. Check they can be read with the file shares down.

Why: Steps kept only on the file server go down with it. CISA's ransomware guide advises keeping a hard copy and an offline version of the response plan.

Week 51 Pick next year's three fixes

Read back through every check you marked none or gap this year. Pick the three to fix first and name an owner for each.

Why: Three fixes, each with an owner, turn a year of checks into a plan. It is also the clearest thing to take to the executive.

Week 52 Update the list and start again

Repeat week 1: update the list of every place data lives. Set it beside last year's list and note what changed.

Why: The list is only as good as its last update. Next year's checks start from this version.

What this does not prove

A finished check shows what your team looked at and when. It does not prove your systems would come back in a real incident, at full size and under pressure, and it does not find problems nobody thought to look for.

Words used here

locked copy
a backup copy that cannot be changed or deleted until a set date, sometimes called immutable

snapshot
a point-in-time copy an array keeps of a volume

UPS
the battery unit that keeps equipment running through a short power cut

Sources

- IBM, What is disaster recovery (DR)?, IBM Think, published 2023 (IBM sells storage and recovery products)
- IBM, What is disaster recovery (DR)?, IBM Think, published 2023 (IBM sells storage and recovery products)
- IBM Documentation, Managing the Safeguarded Copy function, Storage Virtualize 9.1.0 (IBM sells FlashSystem)
- CISA, FBI, NSA and MS-ISAC, #StopRansomware Guide, published 2023
- CISA, FBI, NSA and MS-ISAC, #StopRansomware Guide, published 2023
- CISA, FBI, NSA and MS-ISAC, #StopRansomware Guide, published 2023
- Microsoft Learn, Manage emergency access accounts in Microsoft Entra ID, published 2026 (Microsoft sells Entra ID)
- NIST, Guide to Enterprise Patch Management Planning, published 2022
- NIST, Guidelines for Media Sanitization, published 2025

Made by StepKnock, a small team in San Diego, paid through our partners when a customer buys. This page sells nothing.

Nothing to sign up for. No cookies, no analytics, and nothing you type here leaves this page.