How much does downtime really cost your organization?

A practical business continuity guide for CIOs and IT managers

YOU FOCUS, WE MANAGE.

EPACT

Introduction

Almost every organization has backups today. But how many organizations can say with certainty that, when it actually matters, they can also recover effectively?

That difference, between having backups and being able to recover, is the subject of this ebook. It's a difference that often only becomes visible during an incident. And by then, it's too late to fix it.

The numbers on cyberattacks are well known by now. Ransomware attacks hit more organizations every year. But recovery after such an attack, according to recent research by Veeam, takes an average of several weeks rather than hours. And attackers deliberately target backup environments, because they know it's the last line of defense.

Yet in many organizations, backup remains something approached in purely technical terms: a task on the IT department's checklist, rather than a strategic pillar of the business.

With this ebook, we want to open up that discussion. We approach backup not as an IT topic, but as a business question. With **one very concrete question: what does downtime actually cost your organization?**

Why do backups fail exactly when you need them? How do you translate technical parameters like RTO and RPO into a decision that leadership and IT need to make together? And how do you use backup as a demonstrable compliance measure in an age of NIS2 and GDPR?

This ebook isn't meant to alarm anyone. It's an invitation to ask the right questions, before an incident forces you to answer them.

CHAPTER 1

The downtime bill: what an incident really costs you

This chapter is essentially a simple calculation model. A way to turn an abstract risk ("our systems could fail someday") into a concrete, quantified risk you can present to your leadership team so they can decide what steps to take.

Most organizations underestimate the cost of downtime. Not because they don't know the numbers, but because they've never actually done the full math.

"Our systems are never down for long" is not a risk analysis. It's an assumption that has, so far, held true. Anyone who actually does the downtime math usually discovers three things: the cost is higher than expected, it's largely predictable, and it consists of layers that are rarely all mapped out together.

Three cost layers that together form the full picture

1

Direct costs

These are the costs that become visible first and are easiest to calculate: the time of your own IT team, external consultants, or forensic experts brought in. Think of replacing compromised hardware, overtime and weekend work, and any costs for specialized incident response...

In a ransomware incident, this often comes with the question of whether or not to pay the ransom. That's a decision that's rarely wise, by the way, since paying offers no guarantee of recovery and only keeps attackers' business model alive.

2

Operational costs

This is the layer that grows fastest and is most easily underestimated: lost productivity from employees who can't work, missed orders and invoicing, delayed deliveries, SLA penalties toward your own customers, and the cost of manual workarounds (working with pen and paper is rarely efficient, and certainly doesn't scale). For some sectors, such as manufacturing, logistics, or retail, this is often the single largest cost of all.

3

Indirect and strategic costs

Think of lost customer trust, contractual consequences (such as penalty clauses or losing out on a tender), possible GDPR or NIS2 sanctions after a data breach, higher insurance premiums afterward, and management and communication time that goes into crisis management instead of the business. At publicly listed companies, a serious incident often even has a measurable effect on the share price.

A note on international benchmarks

Many published figures on the cost of downtime are based on large, often American companies, and run into the hundreds of thousands of euros per hour. Those figures are useful as a frame of reference, but not very useful as a benchmark for a Belgian SME or mid-sized organization.

So always calculate your own figure, based on your own revenue, headcount, and operational context, rather than treating a published average as gospel.

A simple calculation method

You don't need to hire an expensive consultant to make a first, realistic estimate. A basic calculation consists of four building blocks:

1

Lost revenue per hour

your annual revenue divided by the number of operating hours per year, optionally weighted for the timing of the incident (an outage during peak season weighs more heavily than one during a quiet month).

2

Lost productivity

the number of employees affected, multiplied by their average hourly labor cost and the percentage of their time that becomes genuinely unproductive.

3

Recovery costs

the estimated cost of internal and external hours, hardware, and any specialized support needed to get back up and running.

4

An estimate of indirect costs

based on comparable incidents in your sector, or on concrete contractual risks you're already aware of today (SLAs, penalty clauses, insurance terms).

A simplified example

A Belgian SME with 60 employees and €12 million in annual revenue that's offline for a full working day (8 hours) quickly loses between €15,000 and €40,000 in revenue and productivity for that day alone, excluding recovery costs and excluding any reputational damage or contractual consequences.
For organizations with critical, time-sensitive processes (healthcare, logistics, financial services), that figure is typically even higher.

What this calculation model doesn't account for

Even a careful calculation remains an underestimate.

It rarely captures the opportunity cost of projects that get postponed because the team is busy managing the crisis, the impact on employee morale and turnover, or the long-term damage to a client relationship that isn't terminated outright, but is quietly scaled back. Those costs are real, even if they don't appear in our calculation model.

CONCLUSION

Downtime is a management issue, not an IT detail

As long as the cost of downtime remains an abstract, unquantified risk, it stays easy to ignore or postpone.

The moment you put a number on it, the conversation changes. It becomes an investment decision a CFO or business owner can weigh: what does a robust backup and recovery strategy cost, versus what a single serious incident would cost you?

In the vast majority of cases, that comparison favors the investment. The downtime bill isn't meant to scare anyone. It's meant to support a decision.

CHAPTER 2

Why most backups fail exactly when it matters

“We run backups.” It’s one of the most misleading sentences in IT. It sounds reassuring, but it says nothing about the question that actually matters: when you need to, can you recover, effectively and on time?

Most organizations hit by data loss did have a backup solution in place. The problem rarely lies in the absence of backups, but in three classic failure scenarios that keep recurring, across virtually every sector.

Three scenarios that keep coming back

1

The backup exists, but has never been tested

A backup job reported as “successful” doesn’t mean the data is actually recoverable. Configuration errors, incomplete backup sets, corrupted files, or forgotten applications (think of a database that falls just outside the backup schedule) often go unnoticed until the moment you try to restore. A backup that’s never been tested is, in practice, an assumption, not a guarantee.

2

The restore takes far longer than expected

Even when a backup is technically sound, recovery in practice is often slower than planned. Network bandwidth, the order in which systems need to be restored (an application that depends on a database server that isn't back online yet, for example), and the lack of a documented recovery plan mean that an estimated recovery time of a few hours can, in practice, stretch into a full working day or longer.

Recent research by Veeam shows that the average recovery time after a ransomware attack stretches to more than three weeks, and that barely one in ten affected organizations can actually recover more than 90% of their data.

3

The backup itself is infected or encrypted

Modern ransomware attacks don't target production data alone. Attackers deliberately hunt for backup infrastructure, to prevent the victim from recovering without paying the ransom.

That same Veeam research shows that in 89% of organizations that suffered a ransomware attack, attackers also actively tried to compromise the backup environment. Yet only a minority of organizations use immutable storage to structurally remove that risk.

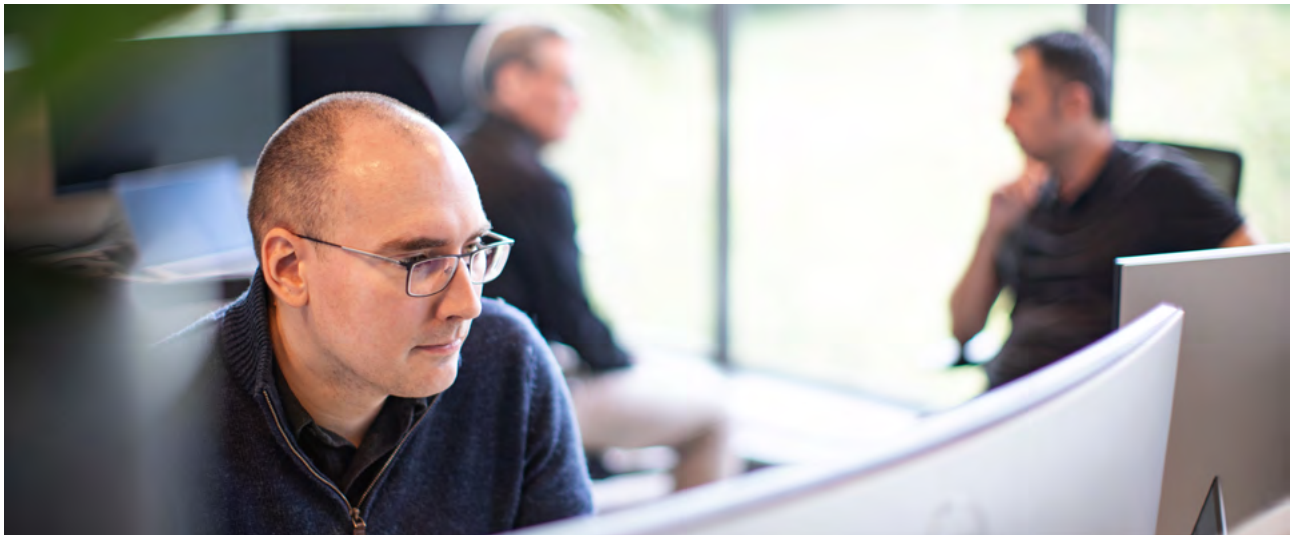

The 3-2-1-1-0 rule as a practical answer

There's no silver bullet against these three scenarios, but there is a proven architectural rule of thumb that structurally limits the risk: the 3-2-1-1-0 rule.

3 **3 copies of your data**
the production data plus at least two backup copies.

2 **2 different types of storage media**
for example, local disk storage and cloud or object storage, so that a single type of storage failure doesn't wipe out all your copies at once.

1 **1 offsite copy**
physically separated from your production environment, so a local incident (fire, theft, a full site outage) doesn't destroy your backups too.

1 **1 offline or immutable copy**
inaccessible or unmodifiable to an attacker who has access to your production environment, even with administrator rights.

0 **0 errors**
regular, documented, and verified restore tests. Not just a green checkmark in a dashboard.

These last two elements are the main evolution compared to the classic 3-2-1 rule, and precisely the elements most often missing in practice.

What this looks like in practice

At Epact, depending on the environment, we combine Veeam Backup & Replication or Veeam Data Cloud with immutable object storage (including via Wasabi, using Object Lock functionality) for the offsite and unchangeable copy.

That means that once a backup is written, no one can modify or delete it during the configured retention period, not even an attacker who has gained full access to your production environment, or even to the backup management itself.

Testing isn’t a one-off snapshot, it’s a habit

A restore test that happens once a year, at a quiet moment, on a non-critical application, says very little about how your environment will behave during a real incident.

A mature backup strategy includes periodic, documented restore tests across different types of workloads (file servers, databases, full virtual machines), with a corresponding measurement of the actual recovery time. Only then do you know whether your RTO target (see chapter 3) is genuinely achievable, rather than just a number on paper.

CONCLUSION

A backup that’s never been tested is not a backup

The three failure scenarios in this chapter have one thing in common: they only become visible at the moment you can least afford them.

A backup strategy structurally built on multiple, separated, and unchangeable copies, combined with regular, documented restore tests, is the difference between an organization that survives an incident and one that takes weeks to get back on its feet, if it manages to at all.

CHAPTER 3

From backup to recovery strategy

As a CIO or IT manager, you obviously know the terms RTO and RPO. But have you ever translated those terms into a concrete business decision? That's a missed opportunity, because the answer to that question is rarely a technical matter. It's a decision leadership and IT really need to make together.

RTO and RPO in brief

Recovery Time Objective (RTO)

answers the question:
how long can a system stay offline before the impact becomes unacceptable?

Recovery Point Objective (RPO)

answers a different question:
how much data can you afford to lose, expressed as time since the last usable backup?

Both figures sound technical, but they're essentially **business decisions**.
Two quick examples to illustrate that:

RTO of 5 minutes

- \+ technically feasible
- \- the price tag isn't justifiable for most workloads

RPO of 24 hours

- \+ perfectly acceptable for an archive system
- \- untenable for an order-processing system

Not every workload deserves the same protection

The most common mistake, then, is a one-size-fits-all approach: the same backup frequency and recovery time for every system, regardless of how important it is.

That creates two problems at once: critical systems end up under-protected, while less important systems soak up unnecessary budget. Classifying systems by business criticality is the first step toward a strategy that's both cost-efficient and robust.

Workload type	Typical RTO / RPO expectation
Business-critical (ERP, order processing, production control)	**RTO:** minutes to a few hours **RPO:** minutes
Important (internal applications, file servers, email)	**RTO:** a few hours to one working day **RPO:** a few hours
Less critical (test environments, archives, historical data)	**RTO:** several days **RPO:** 24 hours or more

These figures are obviously indicative. The right values for your organization depend on your sector, your contractual obligations toward customers, and the real financial impact you calculated in chapter 1.

The conversation leadership and IT need to have

RTO and RPO aren't technical parameters you can set on your own as a CIO or IT manager. They're trade-offs between cost and risk that belong at the leadership level.

A few questions that can make that conversation with leadership concrete:

Which systems are genuinely business-critical, and which ones just feel that way because everyone uses them daily?

What's the financial and contractual impact if a specific system is unavailable for one hour, one day, or one week?

Who decides, during an incident, which systems get restored first?

Which budget is tied to which RTO/RPO target, and is that a deliberate choice or something that just happened over time?

The right recovery technique for the right objective

Once the targets are set, they help determine which technical approach is needed:

For business-critical solutions and systems (near-zero downtime)

replication to a standby environment or Disaster Recovery as a Service (DRaaS), where a critical system can be taken over by a secondary environment within minutes.

For widely used solutions and systems (limited downtime)

fast restore techniques such as instant recovery, where a virtual machine can be started directly from the backup repository while the full data set is restored in the background.

For less critical solutions and systems (e.g. archives)

classic, cost-efficient backups with longer retention, where a longer recovery time is acceptable.

Disaster Recovery as a Service (DRaaS)

For organizations with workloads that can tolerate almost no downtime, a DRaaS solution provides an answer for the moment an entire environment, or an entire site, goes down.

Instead of recovering from backups, an up-to-date copy of your environment is kept continuously running in a second environment, ready to take over the role of the primary environment within an agreed timeframe.

CONCLUSION

RTO and RPO are business decisions, not technical specifications

The question “how much downtime and data loss can we afford?” is too important to leave to IT alone, and too technical to leave to leadership alone. It’s a conversation both sides need to have together, using the downtime bill from chapter 1 as their shared language.

CHAPTER 4

Back-up as a compliance tool

For many organizations, backup isn't just a technical line of defense. It's also one of the most concrete, auditable pieces of evidence that an organization takes its legal obligations around data protection and cybersecurity seriously.

GDPR, NIS2, and sector-specific regulations rarely spell out exact technical requirements. What they do spell out are clear expectations: can you demonstrate that data is available, recoverable, and protected? This chapter translates those expectations into concrete, auditable measures.

GDPR: availability as a security requirement

Article 32 of the GDPR requires organizations, taking into account the state of the art, to implement appropriate technical and organizational measures to ensure, among other things, the availability and resilience of processing systems, and to be able to restore timely access to personal data in the event of an incident.

In other words, a backup strategy isn't a side condition under GDPR, but an explicitly named security measure. A data breach where personal data can't be restored in time is therefore both an operational and a legal problem.

NIS2: backup management as a mandatory measure category

NIS2 goes a step further. Article 21 of the directive lists ten categories of mandatory risk management measures, and “backup management, disaster recovery, and crisis management” is one of them, named explicitly. This isn’t an implicit expectation. It’s a literal, stated obligation.

In Belgium, NIS2 is enforced by the Centre for Cybersecurity Belgium (CCB), through the CyberFundamentals framework (CyFun). Organizations in the designated sectors, generally from 50 employees or more than €10 million in revenue, fall under the obligation. Essential entities (energy, healthcare, digital infrastructure, government, and more) face stricter, proactive oversight. Important to know: NIS2 also introduces a form of personal liability for management when cybersecurity measures aren’t approved or followed up on.

What auditors specifically expect

1. **A documented backup policy**, with explicitly defined RTO and RPO targets per system.
2. **Application of the 3-2-1-1-0 rule:** multiple copies, different media, an offsite copy, an immutable or offline copy, and periodic, error-free restore tests.
3. **Encryption of backup data,** both in transit and at rest.
4. **Full scope:** not just on-premise servers, but also cloud environments and SaaS applications such as Microsoft 365 (email, SharePoint, OneDrive, Teams), which are often wrongly assumed to be “already covered by the provider.”
5. **Documented restore tests** with actual measured recovery times, not just automated backup reports.
6. **Certifications and audit rights with your hosting or backup partner,** such as ISO 27001, and clear agreements for handling a security incident.

The often-forgotten dimension: data sovereignty

For organizations with sensitive data, or in sectors such as healthcare, financial services, legal services, or government, the relevant question isn't just whether data can be restored, but also where that data and the associated backups actually reside.

The Schrems II ruling and the extraterritorial reach of US legislation such as the CLOUD Act mean that a backup physically located in Europe isn't automatically beyond the reach of foreign authorities if the underlying provider is American.

From technical measure to demonstrable control

The difference between "we have backups" and "we can prove our backups meet the legal requirements" comes down to documentation and traceability.

Anyone who can show that every restore test is documented, that retention periods were deliberately set based on legal and sector-specific requirements, and that access to the backup environment is restricted and logged, is in a much stronger position facing an auditor, a regulator, or, worst case, a judge after an incident.

CONCLUSION

Compliance is the result of good architecture, not a separate checklist

When your backup strategy is architecturally sound, built on immutability, separated copies, and periodic testing, you already meet, in practice, most of what NIS2 and GDPR require. Compliance then stops being a separate project alongside the technical reality, and becomes a logical result of it.

CHAPTER 5

The human side of an IT incident

This chapter isn't about technology. It's about what an IT incident does to people: to the IT manager who gets a call at 6 a.m., to employees who can't work, and to a leadership team that has to make decisions within hours it never prepared for.

The scenario below is a composite, anonymized picture based on experiences we've had with multiple organizations over the years. It's representative of how an incident tends to unfold in practice, rather than a literal account of any one specific client.

How it begins...

Monday morning, 6:34 a.m. An IT manager at a mid-sized organization gets an automated alert: unusual activity on the servers. By the time he logs in, it's already too late. Files are encrypted, with a ransom note in every folder. The first reaction is disbelief: "this happens to other companies, not to us."

The next hour is chaos. Which systems exactly were hit? Are the backups still intact? Should leadership be informed now, or wait until more is known? Will employees be able to work today? Do customers need to be notified?

What stress does to decision-making

Research by Veeam among organizations that went through a ransomware attack confirms what most IT managers already know from experience: 45% of those involved report a sharply increased workload after an incident, and 40% experience a clear rise in stress.

That same research reveals a second, equally important pattern: organizations that were confident they were well prepared before an attack saw that confidence drop by more than 20% afterward.

Almost all surveyed organizations (98%) had a ransomware playbook, but fewer than half had built the essential technical elements into it, such as verified backup frequencies or a pre-defined chain of command.

A playbook that exists on paper but has never actually been lived through offers little to hold onto when the moment arrives.

The impact on a single employee

Not every incident is an attack from outside. Sometimes the cause of data loss is a personal mistake or a flaw in an internal process, and the impact is limited to a single individual, which doesn't make it any less significant.

An employee at an IT company told us how he discovered one morning that he could no longer access his Microsoft 365 account. Due to an error in an offboarding process, he had been swept up in his (large) employer's offboarding workflow, while still actively employed.

The result: his entire account, including his mailbox, was deleted. His organization had no separate backup of its Microsoft 365 data, an assumption many organizations still make, believing Microsoft already covers this.

Despite repeated attempts, both internally and directly with Microsoft, it took a full week before anything could be recovered. What eventually came back was only a very limited part of his mailbox. His notes in OneNote, and especially the links between those notes, related documents, and the corresponding emails, were gone for good.

This example illustrates a dimension that's often underexposed in the business case for backup: the impact of data loss isn't just a matter of company-wide downtime

It can just as easily be limited to a single employee, whose productivity, project knowledge, and working context are suddenly set back by weeks, without there ever having been an attack, a ransomware incident, or even a technical failure.

What makes the difference

In organizations that are well prepared, that same Monday morning looks fundamentally different.

There's a pre-agreed escalation procedure: everyone knows who decides what, and when. There's a partner at the table within the hour, working through it in a structured way: which systems were hit, which backups are intact and untouched, what's the fastest, safest path to recovery?

The difference isn't the absence of stress, because an incident remains disruptive, even with the best preparation. The difference is the presence of something to hold on to: a clear step-by-step plan, a partner who owns the technical side so the internal team can focus on communication and decision-making, and the certainty that a reliable, untouched backup exists to recover from.

What good preparation actually changes

An organization with a tested recovery plan and immutable backups no longer has to ask, during an incident, "will we ever get our data back?" The question shifts to "how quickly can we be operational again, according to plan?"

That's a fundamentally different conversation, with a fundamentally different effect on the people having it.

CONCLUSION

Preparation protects people, not just data

A backup and recovery strategy is often discussed purely in technical and financial terms.

But the most tangible impact of good preparation may be the most underrated one of all: less panic, fewer rushed decisions, and a team that knows that, even at the worst possible moment, a way out has already been planned.

CLOSING THOUGHTS

Downtime isn't a question of if, but of when

No organization is immune to an incident, regardless of size, sector, or how much budget goes into IT security.

What is within your control is how much that incident costs you, how long it lasts, and how well your team and your organization are positioned when it happens.

With this ebook, we wanted to give you a framework for that conversation: a calculation model to quantify the cost of downtime, insight into why backups fail exactly at the decisive moment, a framework for discussing RTO and RPO as a business decision, a guide to using backup as a demonstrable compliance measure, and an honest look at what an incident does to an organization on a human level.

YOU FOCUS, WE MANAGE.

How can Epact help?

At Epact, we combine Veeam technology with immutable, offsite storage and, where needed, fully managed backup and recovery planning, so backup is no longer an unanswered question when it matters most.

From Zero Effort Server and M365 backups, to Veeam Data Cloud, to Disaster Recovery as a Service: we help build a recovery strategy that matches your organization's RTO and RPO targets, and the compliance requirements that come with them.

Want to know whether your current backup strategy holds up to that standard? We're happy to think it through with you, with no strings attached.

Get in touch

For a no-obligation assessment or a technical deep dive. That way, you're not taking a leap in the dark, but making a well-informed decision.

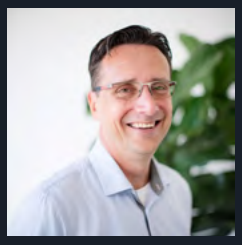

Christof Ugau
Managing Partner

christof.ugau@epact.be

Leo Chang
Business Developer

leo.chang@epact.be

EPACT