The Unprotected Workload Gap

You backup system is protecting what you used to run, but not what you're running now

Contents

Executive Summary

Most organizations still define “backup” as what they protected a decade ago: On-premises servers, virtual machines (VMs), and databases. However, the term “backup” has grown to encompass Microsoft 365, Salesforce, cloud-native apps, and data pipelines as well.

This shift unfortunately creates a protection gap that stays invisible until you need to restore something. Dashboards will look green because they only report on what they can see. The result of this is a growing gap between what organizations believe is protected and what they can actually recover when data is lost.

Traditional Backup Coverage vs. Modern Workloads

✓ **Protected**	✗ **Often Protected**
On-premises servers	Microsoft 365 (Exchange, Teams, SharePoint)
Virtual machines	Salesforce CRM data
SQL databases	Cloud-native application
File servers	AI training datasets

Backup Assumptions Haven’t Kept Pace with Modern Workloads

Traditional backup strategies assumed that everything important ran inside your data center. Servers had predictable lifecycles and known boundaries and VMs lived in defined storage locations. Databases belonged to specific teams with clear backup windows, and protection was pretty straightforward.

These assumptions worked when everything that mattered ran inside your four walls. But today’s IT environments no longer align with those assumptions, and many critical workloads now operate outside the scope of traditional backup systems.

What Traditional Backup Was Built to Protect

Legacy backup tools excel at protecting physical and virtual infrastructures, such as:

- Servers with known IP addresses
- VMware environments with defined storage
- SQL databases that sit on infrastructure you control
- Anything with a predictable file system and a clear backup window

For these workloads, traditional backup still works exactly as designed. These tools assume that your data is reachable via file systems, agents, or snapshots.

How “Protected” Became a False Sense of Security

Here’s where the blind spots form: Backup dashboards report only on what they can see, but that might not include everything that needs protection. If your backup system can’t reach Microsoft 365 or Salesforce, for example, it won’t report an issue. The infrastructure it *can* protect looks fine. However, the platforms that run your revenue operations, customer communications, and collaboration sit entirely *outside* that coverage.

IT teams often assume that SaaS platforms include built-in protection. Cloud providers handle the infrastructure, so surely they handle backup, too, right? Not quite. While cloud providers may offer native backup services, these typically function as point solutions that protect specific workloads without the coverage or unified visibility that enterprise compliance requirements demand. The result? Retention policies that don’t guarantee point-in-time recovery, short recycle-bin windows, and fragmented protection across multiple platforms.

Evolution from On-Prem to SaaS-First Environments

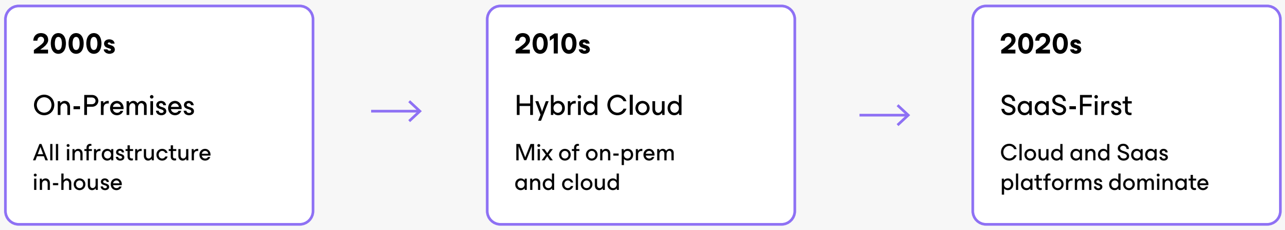

2000s

On-Premises

All infrastructure in-house

2010s

Hybrid Cloud

Mix of on-prem and cloud

2020s

SaaS-First

Cloud and Saas platforms dominate

Whitepaper

The Modern Workloads Most Backup Systems Miss

The platforms most critical to your daily operations are often the least protected. That's not because anyone deliberately left them exposed, but because traditional backup tools weren't designed to reach them. They're outside the infrastructure boundary and only accessible via APIs.

Microsoft 365: Email, Files, and Collaboration Data

Exchange Online hosts your business communications, while OneDrive stores individual work products. SharePoint holds project files and institutional knowledge, and Teams contains meeting recordings, chat history, and shared documents. However, most traditional backup infrastructure doesn't back up these important files.

Microsoft provides retention policies but make no mistake: Retention is *not the same as* backup. Retention keeps data for compliance purposes. Backup provides point-in-time recovery, version history, and the ability to restore data after accidental deletion or ransomware encryption.

CRM and Revenue Systems Like Salesforce

Customer records, opportunity pipelines, workflow automations, custom objects, and metadata all live in Salesforce. This means that a data loader error, a bad integration that pushes null values into required fields, or an accidental mass deletion can wipe out months of work across thousands of records. Without independent backup, your recovery options will depend entirely on the platform's native tools. This often leads to incomplete restores, data loss, and compliance gaps when regulations require full recoverability.

Cloud-Native Applications and AI-Driven Data

SaaS applications built without traditional file systems can only be accessed through APIs, not the file-level methods legacy backup tools rely on. Many of these platforms maintain short retention windows — sometimes just days or weeks — before data ages out permanently.

AI training datasets and analytics pipelines represent particularly high-value operational assets. Recreating this data often requires significant compute resources, specialized expertise, and time, which delays decisions and disrupts downstream systems. Because ownership frequently spans development teams, data science groups, and IT, protection responsibilities are often unclear, leaving these datasets exposed by default.

Assumed Protected vs. Actually Protected

Workload	**Assumed Protected**	**Actually Protected**
Microsoft 365 Email	✓	✖
SharePoint Files	✓	✖
Teams Chat History	✓	✖
Cloud Databases	✓	✖
AI Training Data	✓	✖

Why Legacy Backup Tools Can't Cover These Systems

This gap exists for technical reasons, not just organizational ones. Traditional backup architecture doesn't translate to modern platforms.

Architectural and Visibility Mismatch

Legacy backup systems work by mounting file systems, accessing storage arrays, or installing agents on servers. They expect data to sit still long enough to be copied. However, SaaS platforms work differently. Data lives behind APIs, not file systems. Content constantly changes, and versioning happens at the application layer. Traditional backup tools have no way to see or access this data, much less protect it. Without consistent visibility into application-level data and versioning, organizations cannot reliably validate restores, leaving them unsure whether recovery will actually succeed when it matters.

Shared Responsibility Creates Blind Spots

SaaS and cloud providers manage the infrastructure, but protecting the data itself still falls to the customer. Salesforce keeps the platform running and Microsoft ensures your tenant stays available. However, if you delete a folder, corrupt a database, fall victim to a business email compromise that wipes your mailboxes, or suffer a breach through a compromised admin, recovery is still your problem. Many organizations discover this boundary only when they need to restore something and find out their provider's tools can't help.

Shared Responsibility Model

SaaS Provider Responsibility

- Infrastructure availability
- Platform uptime and performance
- Physical security

Customer Responsibility (Often Overlooked)

- Data backup and recovery
- Protection from accidental deletion
- Ransomware recovery
- Long-term retention for compliance
- Point-in-time recovery

When the Gap Becomes a Business Problem

This protection gap remains theoretical until it isn't. When unprotected workloads fail, the impact cascades quickly from a technical issue to business disruption.

Operational Disruption

When recovery is incomplete or impossible, technical data loss quickly becomes a business-level disruption. Lost email threads stall customer conversations, missing SharePoint files delay project deliverables, and corrupted Salesforce data breaks revenue reporting. Teams then scramble to recreate work from memory or partial records and what should take hours can stretch into days or weeks.

Compliance and Legal Risk

Regulatory requirements often specify the need for retention *and* recoverability, not just retention alone. When auditors ask you to demonstrate recovery capabilities or produce specific historical records, "We don't have that anymore" isn't an acceptable answer. eDiscovery requests become impossible to fulfill in their entirety and compliance gaps turn into audit findings.

Financial and Reputational Impact

Reconstructing lost data costs real money: Staff time, consultant fees, delayed decisions, and missed deadlines. Customer trust erodes when you can't access their account history or communications and disappears entirely when they realize you left their personal data unprotected. Revenue suffers when your CRM data disappears.

Cascade from Data Loss to Business Impact

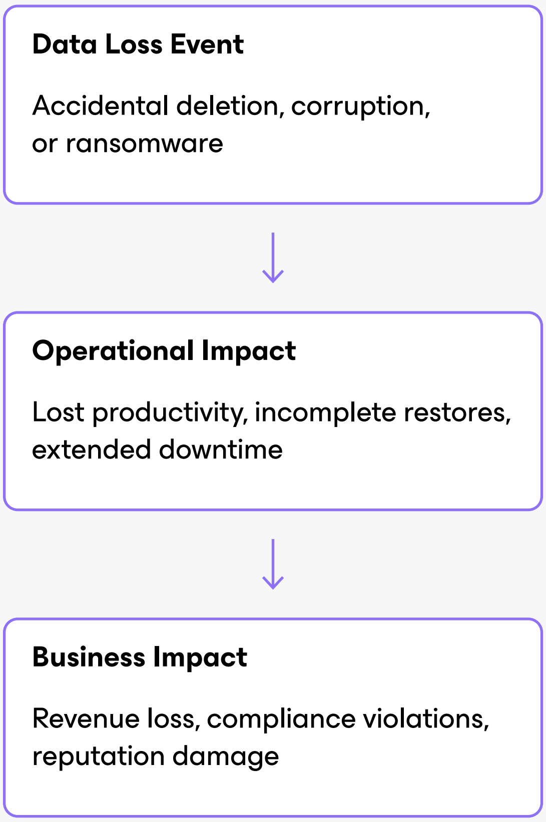

Data Loss Event

Accidental deletion, corruption, or ransomware

Operational Impact

Lost productivity, incomplete restores, extended downtime

Business Impact

Revenue loss, compliance violations, reputation damage

How to Identify Your Own Unprotected Workload Gap

To find the gap, you must honestly assess your environment.

Inventory: What Actually Runs the Business

Start by making a list of the platforms your teams use every day. Look beyond what your backup system reports and focus on the applications that actually run the business, such as Microsoft 365, Salesforce, Workday, cloud databases, container platforms, and other SaaS tools. Be thorough and don't forget to include data pipelines and AI workflows. This often surfaces overlooked dependencies such as identity platforms, automation workflows, or analytics pipelines that never appear in traditional backup reports.

Map Protection Methods to Each Workload

Next, document how you're protecting each platform:

- Backup
- Replication
- Native retention
- Nothing at all/no backup

Be specific about your recovery capabilities. Can you restore individual items? Roll back to a specific point in time? Recover after ransomware encryption?

Validate Recovery, Not Just Coverage

If you can't remember the last time you dusted off your data recovery run book, you aren't alone. But your backup is only as good as the ability to restore your data. Prioritize testing whether your backup actually works. Try restoring a single email, a SharePoint site, or a Salesforce custom object to validate recovery across dependencies. Testing lets you identify any potential issues now rather than during an emergency.

Self Assessment Worksheet

Use this framework to audit your current data protection coverage

Platform/Workload	Protection Method	Recovery Tested?
Microsoft 365		
Salesforce		
Cloud Databases		
SaaS Applications		
AI/ML Datasets		
Container Platforms		

Closing the Gap: What Modern Data Protection Requires

Closing the protection gap requires a shift in how data protection is defined and applied across modern workloads.

SaaS-Aware Backup and Recovery

Protecting SaaS platforms requires approaches that operate at the application layer rather than at the file system level. Data is accessed through APIs, governed by platform-specific structures, and subject to continuous change. Effective protection must account for application metadata, version history, and granular recovery scenarios that native retention alone cannot support.

Unified Visibility Across Environments

As data spreads across on-premises infrastructure, cloud services, and SaaS platforms, visibility becomes fragmented. Without a consolidated view of what is protected, what is retained, and what is not recoverable, organizations struggle to identify gaps before recovery is required. Consistent reporting and documentation are necessary to support audit, compliance, and recovery planning too.

Protection for Emerging Workloads

Emerging workloads such as cloud-native applications, container platforms, and AI data pipelines do not conform to traditional backup models. These environments are dynamic, distributed, and often ephemeral, which complicates ownership and recovery planning. Protecting these workloads requires strategies that can scale, adapt, and handle non-persistent data structures.

Unified Data Protection Platform

On-Premises
Servers, VMs, Databases

API-Level Protection

Cloud
AWS, Azure, GC

Granular Recovery

SaaS
M365, Salesforce, etc

Audit-Ready Reporting

Whitepaper

Building a Future-Ready Backup Strategy

Closing the gap isn't just about filling gaps in your current coverage; it's about building resilience that aligns with how your business operates.

Align Backup with How Work Actually Happens

Collaboration happens in Teams and SharePoint, customer relationships live in Salesforce, development runs in containers, and AI models train on cloud platforms. With this level of continuous change, you need to ensure your protection covers where your work is, not just where the infrastructure is.

Make Recoverability a Measurable Outcome

Regular testing, clear ownership, documented recovery procedures, and alignment with risk planning turns backup from a checkbox into actual resilience. The question isn't "Are we backing up?" but "Can we recover what matters, when it matters?"

Conclusion

Protection gaps stay invisible until they matter. The difference between “we’re protected” and “we can recover” becomes painfully clear only when you need your data back and discover it isn’t there.

The workloads that drive your business have moved beyond what traditional backup was designed to handle. Modern data protection isn’t about replacing what works; it’s about covering what’s exposed. Organizations that identify and close these protection gaps now will be better positioned to recover quickly, maintain trust, and demonstrate resilience when it matters most.

Learn more: veeam.com