rewind

The ROI of SaaS Data Resilience

How a routine Jira mistake becomes a six-figure recovery at a 2,300-person organization.

The problem

Under the **Shared Responsibility Model**, Atlassian keeps the platform running, but you protect the data your teams create inside it.[1] If an admin accidentally deletes issues or an automation rule fires on the wrong filter, Atlassian will not restore it. Deleted Jira issues are permanently removed. There is no recycle bin.[2]

The industry-standard **3-2-1 backup rule** requires three copies of your data, in two locations, with one copy independent of the platform it protects.[25] Atlassian's native Backup & Restore does not meet this standard:

- **No granular restore:** Full-instance only, rolling back all users' last 24 hours of work (24hr RPO, 12hr RTO)
- **30-day retention:** Insufficient for most compliance requirements (HIPAA, SOX, GDPR)
- **Not independent:** Backups live within Atlassian's infrastructure, with no export or customer storage option
- **Premium/Enterprise tier only:** Standard plan customers have no native backup at all
- **Data caps:** 300 GB for Jira, 32 GB for Confluence; no unlimited option

AI tools are compounding the risk. Agents like Claude Code now operate across Jira, Confluence, and GitHub simultaneously through MCP integrations. A single mistake propagates across every connected platform at machine speed, with no cross-platform undo.

How often does this happen?

68%
of data breaches involve a non-malicious human element (errors or social engineering)[4]

Verizon DBIR, 2024

Only 15%
of enterprises prioritize SaaS backup as a critical requirement[5]

Gartner, August 2024

16x
more data moved by one AI agent than all human users on the same platform[22]

Obsidian Security, 2025

The risk comes from three directions:

- **Human error:** The dominant cause of data incidents: bulk edits, bad JQL filters, automation misconfiguration. 87% of organizations experienced SaaS data loss of some severity in 2024.[3] (Kaseya, n=3,000+; vendor survey).
- **AI agents:** Tools like Claude Code and Atlassian Rovo now have write access to production data through MCP integrations. In one observed deployment, a single AI agent moved 16x more data than all human users on the same platform. Mistakes at that velocity propagate across platforms before anyone notices.
- **Ransomware:** 94% of attacks target backups before triggering encryption.[20] CVE-2023-22518 (CVSS 10.0) in Confluence was actively exploited by Cerber ransomware to destroy instance data.[23]

Existing controls, like permission restrictions, confirmation dialogs, audit logging, and user training, reduce the probability of incidents but cannot eliminate them. None of these controls help you recover after data is lost. Backup is the control that limits the blast radius when prevention fails—yet only 15% of enterprises treat it as critical.

The scenario

Company: 2,300 employees. Engineering team of 500 developers, supported by 125 PM, QA, and operations staff (625 total Jira users). Either of the following incidents occurs:

Human error

An engineer runs an API script to clean up resolved tickets. The JQL filter is misconfigured. Instead of targeting one project's completed backlog, the script triggers bulk deletion across three active projects, removing approximately 2,000 issues, including in-progress sprint work, customer-reported bugs, and active feature development tickets.

AI agent error

An engineering team uses an AI coding agent with MCP access to Jira and Confluence. The agent is instructed to update and close a batch of resolved tickets. It misinterprets the scope and bulk-modifies 2,000 issues across three projects, overwriting custom fields, clearing descriptions, and deleting linked attachments.

Regardless of how the deletion occurs, the outcome is the same: Jira issues have no recycle bin. Atlassian confirms they cannot restore customer-initiated changes. The only path forward is manual reconstruction.

This is not hypothetical. Atlassian Community forums document multiple cases of administrators permanently losing issues through bulk operations, automation errors, and script mistakes, all with no recovery path.

What recovery looks like without a backup

Day 1	**Triage.** Contact Atlassian Support. Confirm customer-initiated changes are not recoverable. Assess scope.
Day 2-7	**Scavenge and reconstruct.** Mine email notifications, Git commits, Confluence links, and Slack messages for fragments. Recreate issues manually via CSV import. Change history, timestamps, workflow transitions, and attachments are permanently lost.
Day 7-14+	**Validate.** Cross-reference with team members based on memory. No authoritative source exists to validate against.

What is permanently lost

Change history and timestamps: Who changed what and when. Creation dates, resolution dates, workflow transitions. Audit trails depend on these.	**Sprint and velocity data:** Burndown charts, velocity trends, capacity baselines, and SLA metrics.
Relationships and context: Attachments, linked issues, parent-child hierarchies, and embedded Confluence references.	**Data completeness:** Recovery without backup yields partial data at best, with zero metadata fidelity.

The cost

Company profile:

2,300 employees | 500 engineers + 125 PM/QA/ops = 625 Jira users | 3 projects affected | ~2,000 issues lost

For this moderate data loss scenario, affecting 150 engineers and 45 PM/QA/ops staff directly across three projects:

Cost component (estimate USD)	
A. Direct recovery costs	
Incident response team (6 staff, 10 business days) 6 dedicated staff: Jira admin, 2 DevOps[7], engineering manager[8], IT staff, PM. Blended rate: $77/hr.[6]	$37,000
External Atlassian consultant (Solution Partner, 1 week) Atlassian Solution Partner for 1 week at standard engagement rate.	$20,000
B. Productivity loss	
Developer productivity loss (150 engineers, 2 weeks declining) 150 of 500 engineers. Week 1: 30% loss. Week 2: 15%. Rate: $81/hr.	$218,700
PM/QA/ops productivity loss (45 staff, 2 weeks declining) 45 staff on affected projects. Week 1: 25% loss. Week 2: 10%. Blended rate: $62/hr.[9]	$39,060
C. Business impact	
Sprint disruption and rework (3 teams, 1 sprint cycle) 3 teams x 6 developers x 32 hours rework x $81/hr.	$47,000
Management overhead (3 directors, 30 hours each) 2 engineering directors + 1 VP, $125/hr avg, 30 hrs each.[10]	$11,250
Estimated total cost of incident	**$373,010**

All salary figures: BLS May 2024 median wages with standard benefits multiplier. All estimates use conservative assumptions.[11]

Regulatory exposure: When your Jira data falls under regulation

Whether regulatory penalties apply depends on what data your Jira instance contains and what industry you operate in. Where triggering conditions are met, the compliance exposure is additive to the $373K direct recovery cost above.

Regulation	Triggered when	Data loss requirement	Penalty range	Enforcement example
GDPR	Jira processes personal data of EU/EEA individuals	Art. 32(1)(b-c): ensure availability; Art. 4(12): "accidental or unlawful destruction, loss"[12]	Up to EUR €10M or 2% of global turnover	EUR 14.5M: Deutsche Wohnen (2019) - data management failure[13]
HIPAA	Covered entity/business associate; Jira contains ePHI	Security Rule requires availability of ePHI, data backup plan, disaster recovery procedures[14]	$145 - $2.19M per provision, per year	$1.3M: Cignet Health (2011) - denying patient record access[15]
SEC / FINRA	Broker-dealer; Jira contains records subject to retention	Rule 17a-4: preserve in non-rewritable format. Strict liability.[16]	$5K - no cap (FINRA: $5K-$310K/violation)	$125M: JP Morgan (2021); $289M: 11 firms (2023)[17]
NIS2	Essential/important entity under NIS2; Jira supports critical ops	Art. 21(2)(c): "backup management and disaster recovery." Personal liability for management.[18]	Up to EUR €10M or 2% of turnover	Enforcement nascent; ~20/27 EU states transposed

Does this apply to your organization?

Whether regulatory exposure applies depends on your specific circumstances. Ask: Does your Jira instance contain personal data of EU individuals? Protected health information? Records subject to SEC/FINRA retention rules? If yes, the compliance dimension is additive to the $373K direct recovery cost.

The insurance angle: Cyber insurance is shifting from reimbursement to prevention requirements. Coalition lists encrypted offline backups as one of five essential controls for coverage eligibility.[19] 94% of ransomware attacks target backups before triggering encryption. Rewind operates outside your SaaS provider's infrastructure, meeting the independence requirement that insurers increasingly demand.

What customers experience

Enterprise electronics manufacturer	**"A mass deletion hit multiple teams. The teams protected by Rewind were minimally impacted. It's hands off, no maintenance."**
Mid-size eCommerce brand	**"Instead of being down for a full week, we restored from backup and were back up in about an hour and a half."**

One moderate incident	Annual Rewind investment
$373K before compliance penalties	$25,220 advanced tier, 625 users

~15x return on investment | ~25 days payback period | Minutes-hours recovery | High data fidelity

What a full-instance loss looks like

The scenario above is moderate: 3 projects, 195 affected staff, 2-week recovery. A full Jira instance loss, triggered by a compromised admin account, a ransomware attack, or a catastrophic automation error, affects all 625 users across every project, with a 4-week recovery timeline. Applying the same model:

- **Recovery team:** $99,000 (8 staff, 20 business days)
- **External consultants:** $40,000 (2-week engagement)
- **Developer productivity:** $1,620,000 (all 500 engineers, 4 weeks at declining rates)
- **PM/QA/ops productivity:** $264,000 (all 125 staff, 4 weeks declining)
- **Sprint disruption:** $311,000 (10 teams, 2 full sprint cycles of rework)
- **Management overhead:** $35,000 (full incident command, 7 leaders)

Estimated cost: $2.37M+ before compliance penalties. ~6x the moderate scenario.

What this analysis does not quantify

- **Recurring risk:** Kaseya's 2025 survey found 87% of organizations experienced SaaS data loss of some severity in a single year. The probability of multiple incidents over a 3-5 year period is high.
- **Permanent data degradation:** Manual reconstruction recovers a fraction of original data, and teams live with that gap permanently.
- **Revenue impact:** No customer churn, delayed product launches, or reputational damage is included in either scenario.
- **Scale:** Costs grow with headcount. A 10,000-person organization faces proportionally larger productivity losses.

Assumptions and methodology

- **Conservative design:** Every estimate uses the lower bound where ranges exist. Developer productivity loss is modeled at 30% for week 1, declining to 15% in week 2. The 30% figure benchmarks against published data for high-impact tool disruptions (New Relic, 2024).[21] Only 30% of the engineering team is counted as directly impacted. Sprint disruption covers one cycle. No revenue impact, customer churn, or reputational damage is included.

- **Compliance scope:** Only regulations with provisions covering data loss and destruction are included. Each regulation cited has been verified against the actual legislation text.

- **Scaling:** All labor costs scale with headcount. A larger incident would multiply productivity losses proportionally.

- **Company profile:** The 625-user Jira deployment (~27% of a 2,300-person organization) reflects the median from Rewind's customer base analysis.

- **Cost of Rewind:** $25,220/year is the published list price for the 601-800 user tier on the Advanced plan (1.3x Standard). Actual pricing varies.

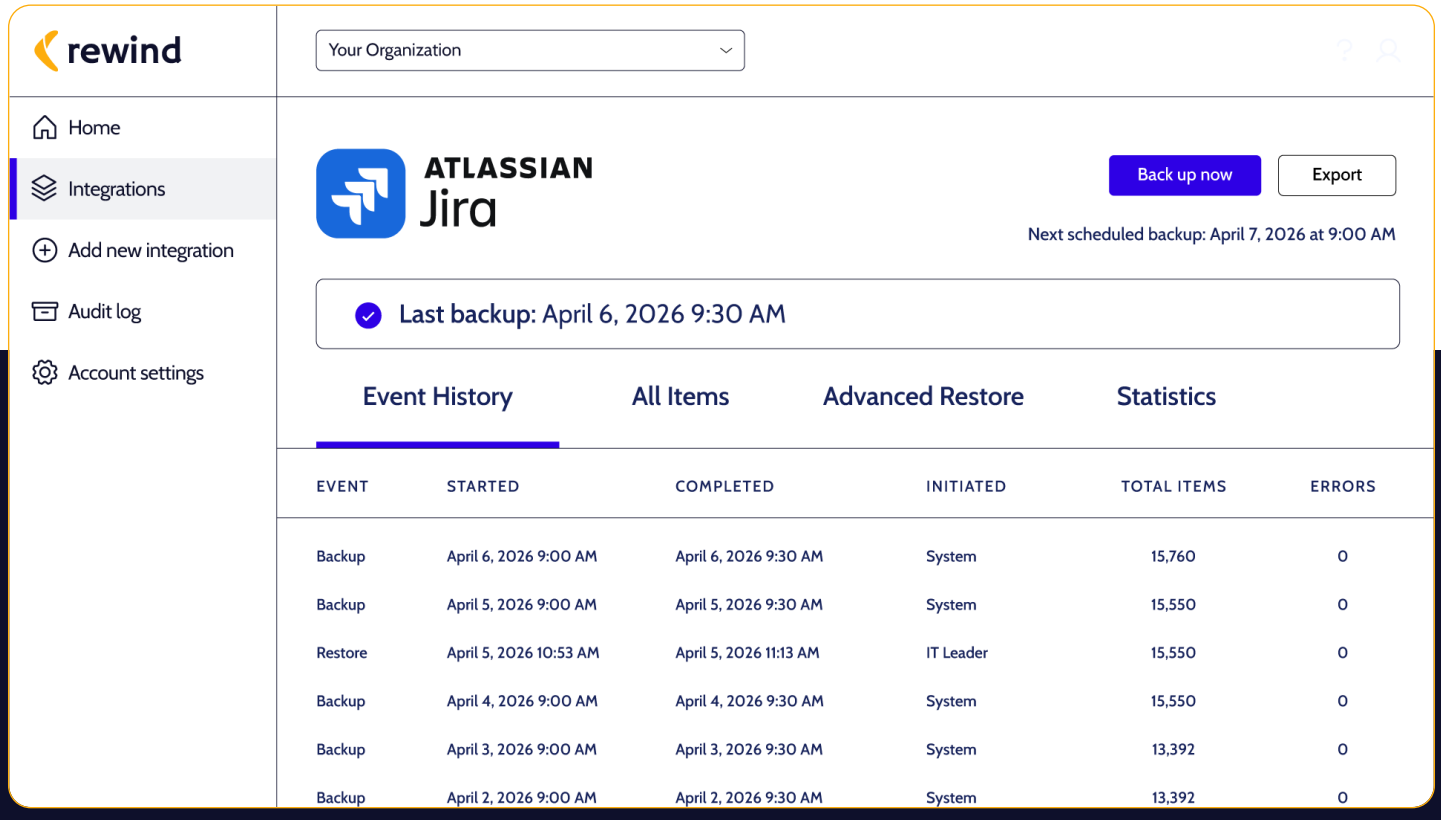

Appendix: Sources

1	Atlassian Trust Center, Data Management Policy.	Shared responsibility, recovery limitations
2	Atlassian Support Documentation and Community Forums.	Deletion mechanics, recovery limitations
3	Kaseya, State of SaaS Backup and Recovery Report, 2025.	Data loss frequency
4	Verizon, 2024 Data Breach Investigations Report (DBIR).	Human error as primary cause
5	Gartner, Strategic Planning Assumption, August 2024.	SaaS backup adoption gap
6	U.S. Bureau of Labor Statistics, May 2024.	Base salary: developers, IT admins, DevOps
7	Indeed, Glassdoor, Built In salary surveys, 2024-2025.	DevOps engineer compensation
8	Built In / Salary.com, Engineering Manager, 2025-2026.	Engineering manager compensation
9	BLS, May 2024. Project Management Specialists.	Project manager compensation
10	Salary.com, Director of Engineering, 2026.	Director-level compensation
11	BLS, Employer Costs for Employee Compensation (ECEC), 2024.	Benefits multiplier
12	GDPR, Articles 4(12), 32, and 83.	GDPR applicability to data loss
13	Berlin Commissioner (BlnBDI), Deutsche Wohnen SE, 2019. EUR 14.5M.	GDPR enforcement for data loss
14	HIPAA Security Rule, 45 CFR 164.306, 164.308(a)(7), 164.312.	HIPAA backup/recovery requirements
15	HHS OCR, Cignet Health, February 2011. $4.3M.	HIPAA enforcement for data unavailability
16	SEC Rule 17a-4, FINRA Rule 4511.	SEC/FINRA recordkeeping requirements
17	SEC Enforcement Actions, 2021-2025. $2B+ aggregate.	SEC enforcement precedent
18	NIS2 Directive (EU) 2022/2555, Article 21(2)(c).	NIS2 backup requirements
19	Coalition, Essential Cyber Insurance Requirements.	Cyber insurance backup requirements
20	Sophos, The State of Ransomware 2024. 94% target backups.	Ransomware targeting of backups
21	New Relic (via Level.io), Financial Impact of Inefficient IT, 2024.	Developer productivity loss benchmark
22	Obsidian Security, Govern AI Agents, 2025.	AI agent data volume, risk amplification
23	Atlassian Security Advisory, CVE-2023-22518, October 2023.	Ransomware targeting Atlassian products
24	Rewind, “The 3-2-1 Backup Rule for SaaS,” 2024.	3-2-1 backup rule, SaaS independence