Leading SaaS data security

Raising the bar for data protection in the cloud era

August 2024, Copenhagen

Foreword

Information security has never been more important—or challenging—than it is today. There are growing threats from external actors and insiders, new systems and software are being deployed faster and more frequently than ever before, the zero-trust model is leading to widespread changes in visibility and access control, and economic pressures demand that costs be reduced or avoided altogether.

At the same time, Chief Information Security Officers (CISOs) also need to ensure appropriate data privacy and protection throughout the company's entire tech stack.

These evolving demands exist in parallel to a paradigm shift in which a range of workloads—from customer-facing applications to business-critical operational systems—have migrated from on-premises appliances and servers to third-party cloud environments.

The cloud brings the agility and availability enterprises require, but it also introduces new challenges. Perhaps foremost among these issues, but one of the least-recognized, is a shared data responsibility model and its important implications for business continuity planning, in general, and disaster recovery specifically.

Keepit's mission is to protect data in the cloud. Since 2013, we have worked tirelessly to purpose-build an efficient and—above all—secure data protection solution that provides simple, reliable, cost-effective, and vendor-neutral backup and recovery for software-as-a-service (SaaS) workloads.

With deep roots and 20+ years of experience building best-in-class data protection and hosting services, we recognized that the only way to create such a solution was to architect it from the ground up—so that's exactly what we did. The output of our efforts is The Keepit Cloud.

When presented with a new solution, it is only natural—and prudent—to ask questions. We understand, and we would expect no less from security professionals. Fortunately, we believe we have created the industry's best SaaS data backup and recovery solution.

We believe in transparency, because it enables prospective buyers and technical evaluators to make informed decisions. Accordingly, we have prepared this document to proactively address the most common inquiries about our solution, to explain what we built and why we made certain decisions, and—frankly—to show how our solution is fundamentally different from, and meaningfully superior to, other products in the market.

We sincerely hope this guide answers your questions and we would greatly value the opportunity to respond to any that we may not have anticipated.

Happy reading,

Jakob Østergaard Hegelund

CTO at Keepit

Recommended by leading analysts and loved by customers globally,
Keepit is the world's only independent cloud dedicated to SaaS data protection.

2.

Table of contents

Executive summary

Backing up cloud SaaS data is the responsibility of the SaaS customer, not the vendor. However, this reality is not yet well understood within the marketplace, and with that lack of understanding comes significant risk: In 2023, 65% of data breaches affected data hosted in public cloud or hybrid environments.[1]

Similarly, many organizations have not yet recognized that backing up SaaS data within the same public cloud infrastructure that hosts their primary data fails to provide the needed assurances—since an event impacting that public cloud could render both the primary and backup data inaccessible, even with geographic distribution.

Only an infrastructure completely independent from public cloud environments can provide true backup for SaaS data. This premise led us to build, from the ground up, The Keepit Cloud—a vendor-neutral and dedicated SaaS data backup solution that is resilient, secure, and exceptionally easy to use.

We architected the solution with a security-first design, incorporating innovative features that meet the needs of even the most demanding enterprises, and we extend that same care to its physical and operational security. As an added benefit, constructing and managing our own infrastructure also allows us to offer our backup service with a predictable, all-inclusive pricing structure.

Introduction

As software-as-a-service (SaaS) applications enable modern businesses, critical services are migrating from the trusted confines of the on-premises computing facility to cloud providers, far away from traditional backup solutions.

This widespread and ongoing shift introduces a shared model in which the SaaS provider and the SaaS customer—you—each assume ownership of particular responsibilities with respect to data security (Figure 1).

SaaS Provider's Responsibility

Application	Hardware failure
Operation system	Software failure
Virtualisation	Natural disaster
Hardware	Power outage
Network	Physical intrusion

Your Data — Your Responsibility

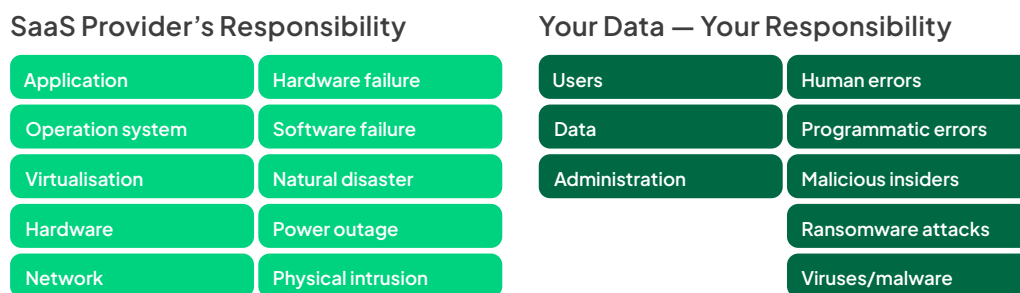|	
Users	Human errors
Data	Programmatic errors
Administration	Malicious insiders
	Ransomware attacks
	Viruses/malware

Figure 1—Division of responsibilities between the SaaS provider (left) and the SaaS customer (right)

While adoption of cloud services has soared—and accelerated by the COVID-19 pandemic's disruption[2] —enterprises are only now beginning to understand the implications of shared data responsibility. Despite the widespread dependence upon SaaS applications, in 2024 only 15% of enterprises prioritize backing up SaaS applications as a critical requirement, according to Gartner. However, they expect this number to increase to 75% by 2028.[3]

Unfortunately, many enterprises will learn the hard way that relying on their SaaS provider to keep their data safe is a risky proposition: in 2023, 65% of data breaches affected data stored in public cloud or hybrid environments.[1]

To help enterprises avoid disruption due to lost SaaS data, Keepit has architected a dedicated, vendor-neutral SaaS data backup solution that is resilient, secure, and easy to use—and that's provided with a predictable, all-inclusive pricing structure.

In this guide, we will explain how this dedicated cloud raises the bar for secure backup and we will present some of the security standards, practices, and features that protect your data.

But before we get to those topics, we'll start by explaining why we built our own dedicated data protection cloud in the first place.

2. For example, Microsoft Teams usage jumps 50 percent to 115 million daily active users [The Verge]

3. Source: Top Trends in Enterprise Backup and Recovery 2024 [Gartner]

Why Keepit built a dedicated cloud

Key Takeaways:

- Because backing up SaaS data in the public cloud comes with significant security compromises (and introduces hidden risks), Keepit operates our own data center regions, running on our own hardware, managed by our own people

- Once a customer chooses a Keepit region for backup, their data never leaves that region

- Building our own cloud allows Keepit to manage supply chain risks and provides tremendous control over costs, which enables us to offer simple and predictable "all-in" pricing

The 3–2–1 principle of backup (Figure 2) mandates that you must have one copy of your data off site. In the days of tape backup, where fire and theft were the only credible threats to your backup data, the off-site copy effectively ensured that your backup data would survive any calamity that could befall your primary data and your primary site.

3 Copies of data
(1 primary and 2 backups)

2 Copies stored
(At separate locations)

1 Off-site storage
(Secure private cloud)

Figure 2—The time-tested 3–2–1 principle of data backup

In the cloud age, however, backups have become much more complicated: geographic dispersal is insufficient to ensure your data is secure, and hidden risks are introduced by relying on clouds that may be taken offline to protect the providers' primary business interests.

True backup requires a separate logical infrastructure

While the requirement for separate backup infrastructure may seem self-evident to many IT and security professionals, most alternative cloud backup solutions will readily store your backup directly on the same cloud infrastructure that hosts your primary data—introducing unnecessary risk and really stretching the meaning of "backup."

You can choose to store your backup data in the public cloud, and you can even specify a data center location that is separate from your primary data—but in no way does this approach provide the assurances that you need, because all your data (now including your only backups!) is managed under the same administrative infrastructure.

Consequently, if that infrastructure is compromised, then both your primary data and your backups are at risk, no matter if they reside on servers in different parts of the country or even on different continents.

In our view, offering "true backup" means guaranteeing to our customers that their backup data doesn't reside on the same logical infrastructure as their primary data, regardless of which combination of primary workloads are being protected.

"True backup requires a logical infrastructure **separate from the primary data**"

To ensure that your backups are never stored on the same cloud infrastructure that hosts your primary data (e.g., Amazon Web Services, Microsoft Azure, etc.), Keepit does not run on public cloud infrastructure; instead, The Keepit Cloud:

- runs in our own data center regions

- operates on our own hardware

- is managed by our own people

To ensure data sovereignty, once a customer chooses a Keepit region for backup, their data never leaves that region. Not only does Keepit guarantee that backup data will always go to that region, but no processes exist that even could transfer backup data out of that region (aside from customer initiated restores and downloads, of course).
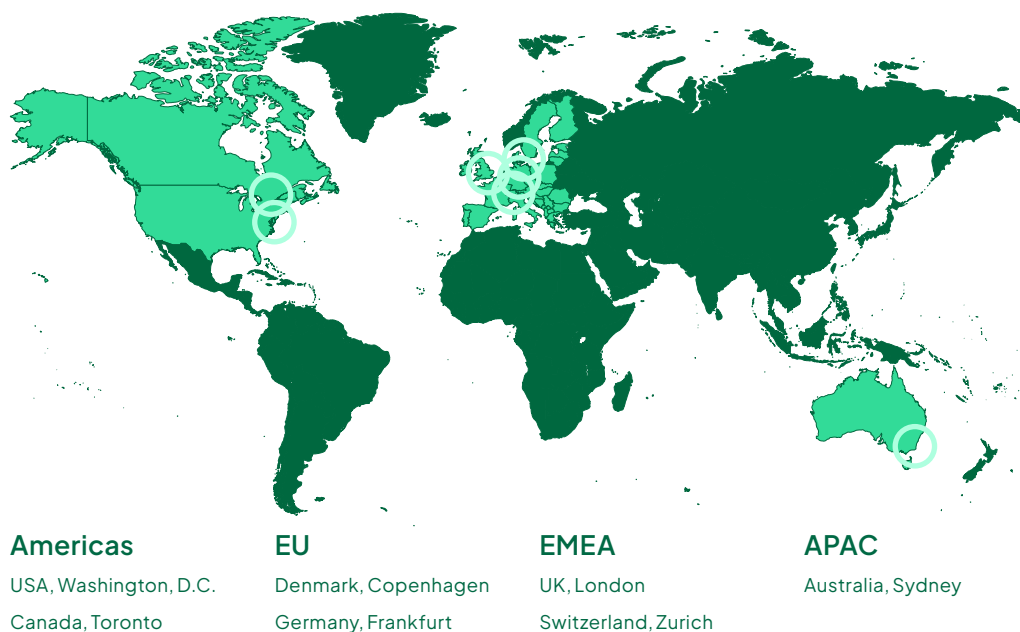

Americas
USA, Washington, D.C.
Canada, Toronto

EU
Denmark, Copenhagen
Germany, Frankfurt

EMEA
UK, London
Switzerland, Zurich

APAC
Australia, Sydney

Figure 3: The Keepit Cloud spans the globe with seven locations, each of which has two data centers

Gaining control and managing hidden risks

Aside from the fact that a dedicated private cloud is a fundamental requirement of any legitimate backup solution, building our own cloud gives Keepit a higher degree of control over both the supply chain and costs, with considerable benefits for our customers.

Additionally, this strategy showed an unanticipated—but vital—benefit when the COVID-19 pandemic struck. To protect their own primary businesses, many public cloud vendors shut down customer workloads, forcing thousands of those customers offline. All the while, Keepit's dedicated private cloud continued normal operation.

Passing along the benefits: Why Keepit customers don't worry about deduplication, compression, and storage

We want to make it simple for you to become (and to happily remain) our customer. Fortunately, building our own dedicated private cloud provides advantages for Keepit that translate into benefits for our customers.

For example, the Keepit licensing model is simple: There are no complicated API transaction fees, network egress or ingress fees, or even storage consumption fees. That's right, storage is included—Keepit customers don't need to worry about deduplication capabilities or compression ratios!

We can offer simple and predictable pricing because we retain the freedom to innovate and to evolve the storage technology behind the scenes—something that would not be possible if we were using a public cloud. Since customers and partners are not sensitive to the cost structures around the underlying storage, we can continue to improve our systems regardless of how this will affect byte-for-byte consumption on whatever storage media the future brings.

There will continue to be some demand for BYOS (Bring Your Own Storage) business models in which the customer effectively pays twice: once for the backup service and once more for the public cloud storage onto which the backup data is subsequently copied. However, this approach is fundamentally at odds with our philosophy that true backups cannot be stored on the same cloud as the primary data; therefore, BYOS is not a model we have chosen to pursue.

Raising the Bar for Secure SaaS Backups

Key Takeaways:

- Keepit's backup service is vendor neutral and already supports major SaaS application providers

- Resilience and data immutability are achieved through immediate encryption of backup data, 30-day delete retention, data region sovereignty and dual data center redundancy, and a unique, tamper-proof infrastructure

- A blockchain-like Merkle tree architecture and innovative "incremental forever" data transfer ensure secure, reliable and change-based long-term data storage

- Restoring from a backup is quick and easy: simply browse and search through time and space, then click to restore once you find what you lost with our multiple restore options.

When designing and building The Keepit Cloud, we aspired to provide a number of important fundamental features and characteristics, including:

- Vendor-neutrality

- Resilience and data immutability

- Secure, reliable and change-based long-term data storage

- Easy and fast data recovery unparalleled in the industry

The following subsections describe these in detail.

Vendor-neutrality

Today's enterprises rely on a growing list of SaaS applications and providers, headlined by Microsoft Office 365 and Dynamics 365, Google Workspace (formerly G Suite), and Salesforce. It is imperative that any practical SaaS data backup solution support these services, otherwise clients will be burdened with the extra overhead, complexity, and potential risk that comes with maintaining multiple backup systems.

Moreover, the SaaS vendor ecosystem is always expanding. Accordingly, The Keepit Cloud is completely vendor agnostic, meaning there is no technical reason why support for a particular SaaS application provider cannot be introduced in addition to the major vendors already integrated into the platform (Figure 4).
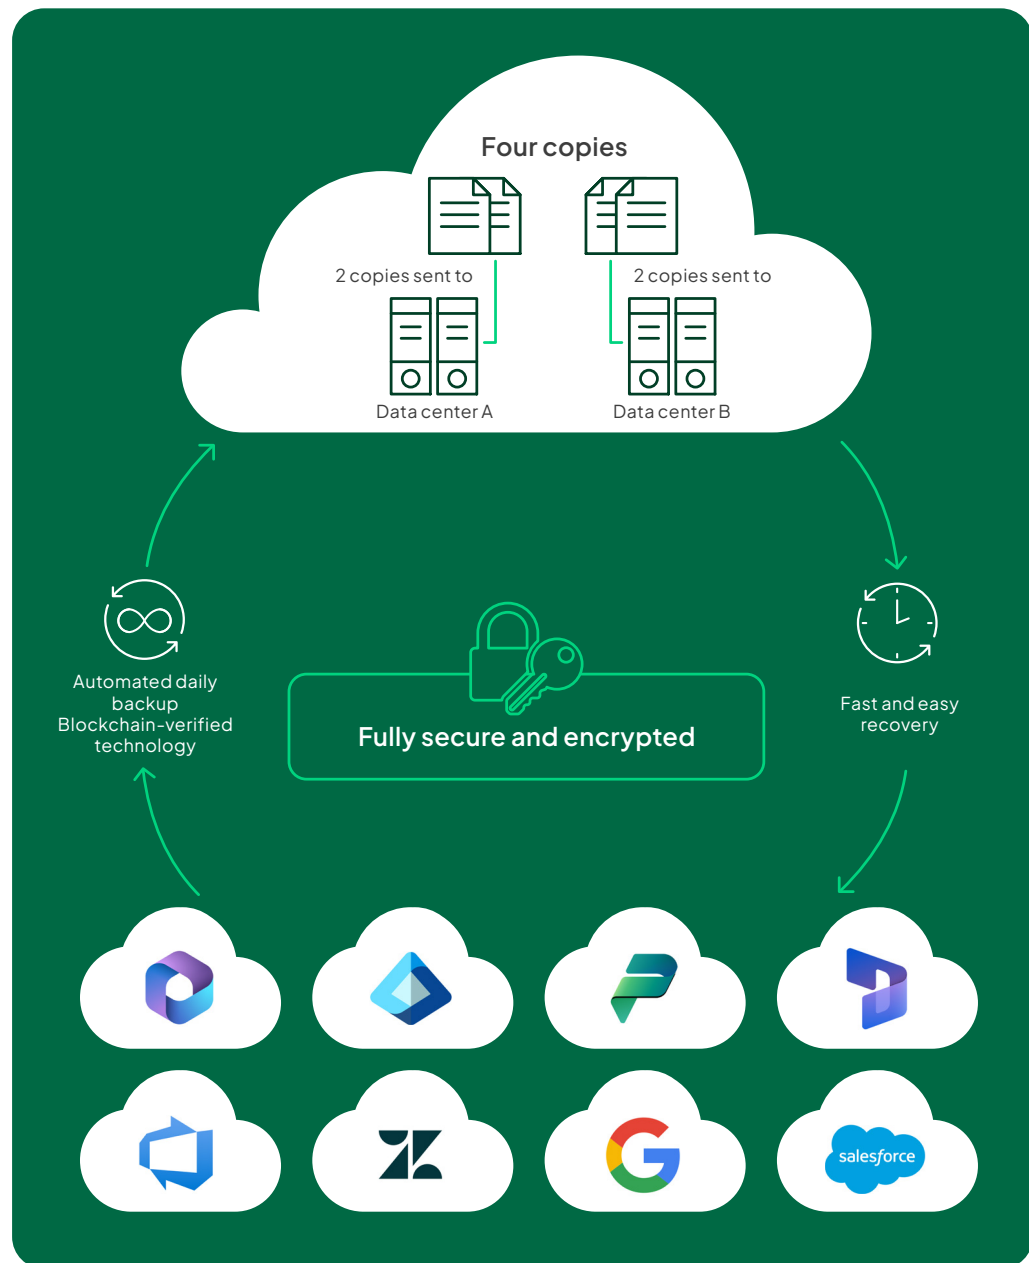
Four copies

2 copies sent to   2 copies sent to

Data center A   Data center B

Automated daily
backup
Blockchain-verified
technology

Fully secure and encrypted

Fast and easy
recovery

Figure 4—The Keepit Cloud is completely vendor agnostic

Resilience and data immutability

Assuming physical security is adequate, no threat actor can erase the data stored on an offline tape. In the cloud, however, everything is always online and expectations for ease of use dictate a high degree of accessibility.

To ensure security while maintaining convenience, Keepit combines several protective measures (Table 1).

Protective measure	Benefits
Immediate encryption	Upon ingestion by the Keepit platform, backup data is immediately encrypted before it is written to storage media. This process employs industry best practice algorithms believed to offer unbreakable security for at least the next 30 years.
30–day delete lock	There is no simple way to immediately delete your data in Keepit: you, or the attacker who has successfully taken over your identity, will have to wait for 30 days for datasets to be deleted or for accounts to be closed. This deliberate delay is the first line of defense against user error, insider threats, and ransomware attacks that target backups.[4]
Data region sovereignty and dual data center redundancy	To comply with data sovereignty requirements and desires for geographical dispersal of information assets, Keepit offers multiple regions around the world. Once your data enters a given Keepit region, your data stays in that region - forever. To provide protection from the most common disasters such as fire, power or cooling system outages—or other events that can render a building or complex of buildings unusable—all data is copied into two separate data center facilities within the chosen region.
Unique, tamper–proof infrastructure	The Keepit infrastructure runs on software that was designed and built from the ground up by our team. This proprietary storage infrastructure allows us to provide fast and reliable backup at a reasonable cost, and supports the data immutability that protects backup datasets—like the backup tapes of yesteryear, our disk-based storage systems do not offer a mechanism for modifying backup data. This core characteristic means that even in a nightmare scenario—e.g., administrative accounts are compromised, primary data is corrupted or encrypted, attackers gain access to your backups—there is no way for the system to comply with an attacker's efforts to tamper with your backup.

Table 1—These protective measures ensure security without limiting convenience

4. Many ransomware families are adept at targeting backups. A recent study found that 94% of organizations hit by ransomware said the cybercriminals attempted to compromise their backup, showing just how much backup and security products are under attack. [Sophos]

Secure, reliable and change-based long-term data storage

We designed the Keepit platform for the specific purpose of providing secure and reliable long-term storage for SaaS backup data. In pursuit of this objective, we created a uniquely powerful information architecture (Table 2) that employs a Merkle tree structure and an innovative "incremental forever" approach to deliver performance and cost benefits that would simply not be possible on general-purpose storage services.

Architectural element	Benefits
Merkle Tree architecture	The Keepit information architecture is inspired by the Merkle tree structure that gained renown from blockchain systems such as Bitcoin and the 'git' version control system. In addition to providing benefits including strong consistency checks, this design also allows us to represent and access backup sets over time—completely removing the antiquated model of 'full' and 'incremental' backups, along with all the overhead associated with that legacy approach.
"Incremental forever" data transfer	In the Keepit data model, when a backup executes, we effectively only transfer the differences in your dataset; however, when you view your data in the Keepit platform, all access is instant across time and space, and every single backup—no matter how old—appears as if it was a standalone full copy of your dataset at the time it was taken. We have even gone a step further: by designing, completely from scratch, an object storage architecture for this information model, we are able to store backup data very efficiently on simple high-density hard drive-based storage systems, which simplifies our supply chain and lowers both the risk and the price point for our customers.

Table 2—By rethinking information architecture, the Keepit platform provides secure and reliable long-term backup with unmatched performance
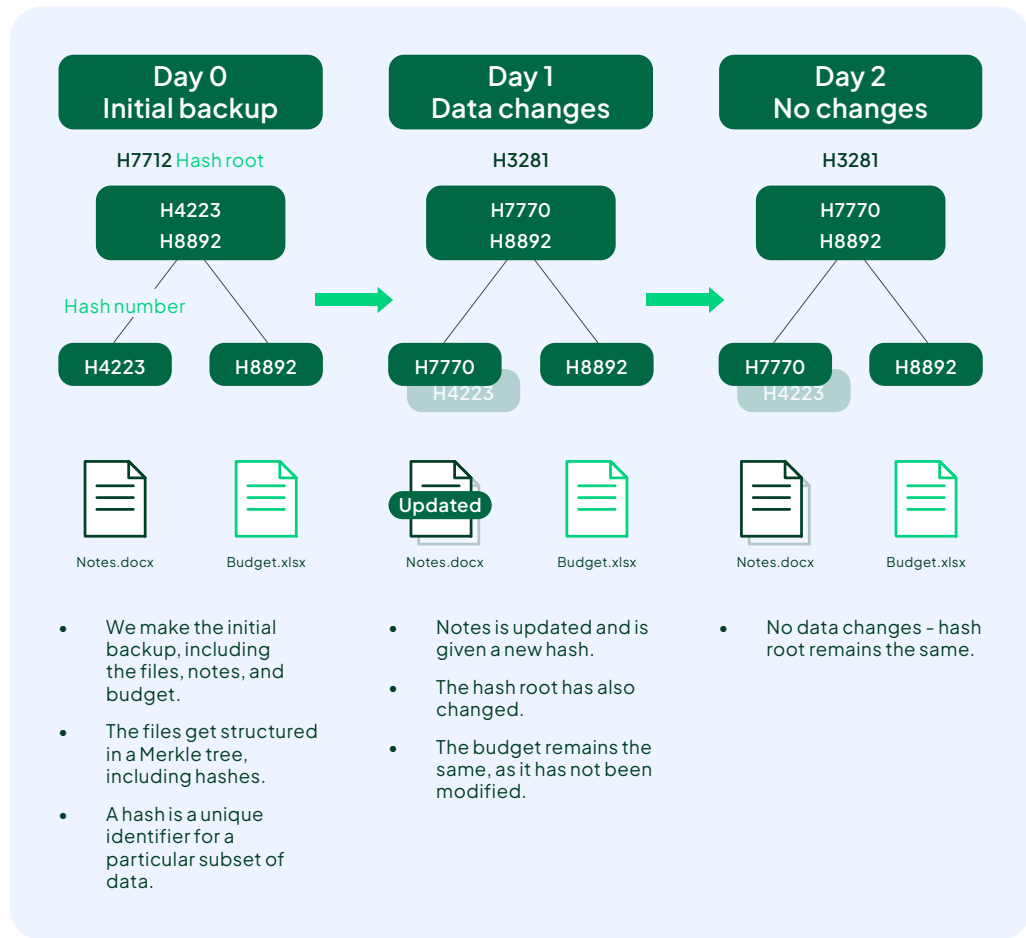
Day 0 Initial backup

H7712 Hash root

H4223
H8892

Hash number

H4223   H8892

Notes.docx   Budget.xlsx

- We make the initial backup, including the files, notes, and budget.
- The files get structured in a Merkle tree, including hashes.
- A hash is a unique identifier for a particular subset of data.

Day 1 Data changes

H3281

H7770
H8892

H7770   H8892
H4223

Updated
Notes.docx   Budget.xlsx

- Notes is updated and is given a new hash.
- The hash root has also changed.
- The budget remains the same, as it has not been modified.

Day 2 No changes

H3281

H7770
H8892

H7770   H8892
H4223

Notes.docx   Budget.xlsx

- No data changes – hash root remains the same.

Figure 5 — The Merkle Tree architecture

Unparalleled easy and fast data recovery

What good is reliable backup if you can't find what you're looking for, or if it takes days or weeks to recover your data?

With Keepit, all data is instantly accessible. To restore from a backup, simply search for or browse to the data you need and click "Restore." All your data with all your history is readily available for you in a modern web-based user interface that offers not only live browsing but also provides previews, downloads, and restores of your data elements.

And that's it—there are no extra steps. You browse and search through time and space, and you click the button once you find what you've lost. The efficiency of the restore process is unparalleled in the market.
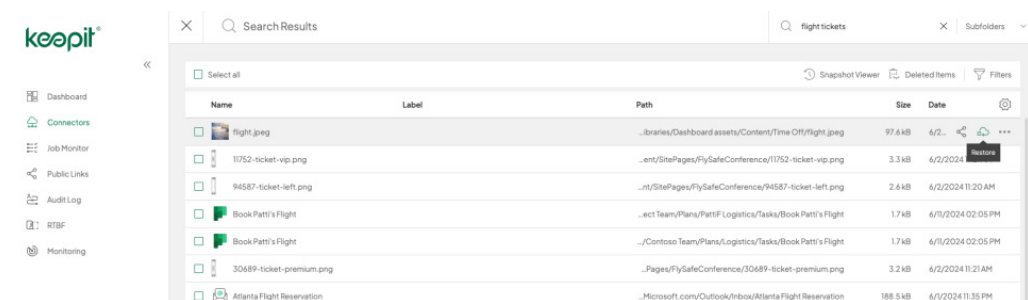

Figure 6 — Easily find the right data with our Smart Search and restore across time with a few clicks

Security standards, practices, and features

Key takeaways:

- All data center facilities employed by Keepit conform to a high physical security baseline and hold ISO 27001 certification plus complementary certifications (e.g., SOC-2, ISAE 3402, PCI/DSS, HIPAA)

- Technical and organizational measures are in place to ensure operational security. Keepit demonstrates its dedication and ability to deliver best-in-class security technology to its customers by holding the ISO/IEC 27001:2013 certification and the ISAE 3402–II certification. These certifications are documentation that our entire organization has been audited and is certified: Not just our data centers, but the services, the technology, business continuity, operations, disaster recovery, sales, and legal. And in all of our locations.

- Mature software development lifecycle processes ensure that security is incorporated from the inception of a new project and continued throughout the entire life of the system

- To enhance protection of customer data, the Keepit platform incorporates TLS transport security and secure, logged access controls; the recommended best practice for deployment uses the customer's existing authentication infrastructure by means of SAML integration

Providing secure backup demands more than just a secure platform design. Security extends to the physical and operational aspects of the platform and must be ingrained in the software development lifecycle.

Physical security

As noted in Table 1, once backup data reaches the Keepit platform, it is immediately copied into systems in two separate data centers within the designated region. This practice provides resilience in case a facility is permanently lost (e.g., to fire or a natural disaster) and also ensures continuous availability of data in case of a more benign facility problem (e.g., temporary failure of power or cooling systems).

In every region, the Keepit platform operates in active-active mode between the two data centers. This model:

- Allows the platform to continue virtually uninterrupted in the unlikely scenario that a full building is incapacitated

- Provides system-level redundancy, allowing the platform to continue servicing customers in case any single system is lost or needs to be taken offline for service, upgrades, maintenance, or repairs

This system-level redundancy is complemented by significant component-level redundancy on the individual servers that host the Keepit platform. Components that most commonly fail on modern servers (e.g., cooling fans, power supplies and spinning disks) can typically be hot swapped without even powering down the system holding the failed component.

Figure 7 - Each region has dual data centers which typically holds the above certifications.

While there are some local differences between the regions [5], all data center facilities employed by Keepit must conform to a physical security baseline that:

- Ensures high availability

- Restricts access to the physical systems that constitute the Keepit platform

Physical security at a facility can feature a combination of multi-factor authentication including biometric identification for facility access, controlled access with man traps, CCTV monitoring of the facility and facility perimeter, plus onsite security staff. Typically, the processing facilities will hold an ISO 27001 certification and several complementary certifications such as SOC-2, ISAE 3402, PCI/DSS, HIPAA etc.

When providing large-scale storage, failure of storage media is a regular and undramatic event. Failed storage media is never sent back to the media vendor for analysis; instead, it is kept onsite for physical disintegration (mechanical shredding) before being disposed of, upholding Keepit's commitment that customer data never leaves its designated region.

Additionally, and as noted previously, customer backup data is encrypted before it is written to storage media. Therefore, even in the most unlikely event that a storage device is stolen from a processing facility, any data on the media remains inaccessible.

Operational security

When it comes to backup and recovery, businesses seeking solutions need to be incredibly thorough in their due diligence processes. The Keepit service holds the ISO 27001:2013 certification and the ISAE 3402–II certification.

End-to-end ISO/IEC 27001:2013 certification:

By achieving the ISO/IEC 27001:2013 certification, Keepit continues to demonstrate its dedication and ability to deliver best-in-class security technology to its customers. The certificate states that:

"[Keepit] operates an Information Security Management system which complies with the requirements of the ISO/IEC 27001:2013 for the following scope:

The ISMS scope includes development, operations, and maintenance of services that support the company's business & B2B SaaS Backup Solutions System, in accordance with the ISMS Statement of Applicability V.20220120."

The scope of what we have done is very ambitious: It is not standard to include the entire organization. With the breadth of this certificate, our customers can rest assured that all processes within our organization live up to the highest international security standards.

Benefits of Keepit's ISO/IEC 27001:2013 certification include:

- A systematic, verified approach to information security that results in superior customer data protection

- Ongoing performance evaluations and internal audits that ensure Keepit continues to meet the requirements of the ISO/IEC 27001 standard

- Continued improvement of business continuity management and disaster recovery plans

- Risk, vulnerability, and security incident management practices that enhance overall information technology (IT) operations security

- Compliance with current and future legal and regulatory requirements

ISAE 3402–II certification:

Keepit holds the ISAE 3402–II certification and is audited by Deloitte on Service Organizational Controls annually. In addition to the conditions required by ISAE 3402–II certification, Keepit has implemented a number of operational security measures (Table 3).

Operational security measure	Explanation
Automation	The Keepit service is fully automated, operating based on the instructions (i.e., the configuration) provided by the customer. Therefore, no human operators outside of the customer's own organization are involved in accessing, reviewing, or otherwise processing actual customer data.
Restricted access	The Keepit platform itself is regularly updated as part of the ongoing maintenance and evolution of our service. However, there are no recurring or routine tasks that require human operators to access customer data. Furthermore, there are technical measures in place to ensure that no such access is required and that it is not easily or accidentally invoked. Additionally, there are organizational measures in place to educate and train staff on information security practices and policies.
Regional support	If a customer needs assistance—and in the rare event that this assistance involves access to actual customer backup data—the Keepit support organization will work with the customer to ensure that such requests are served by support staff in an appropriate region, considering data sovereignty requirements, support availability, and time constraints.
Third-party security assessments	Not only do we employ external auditors for security assessments, we also subject our platform to regular penetration testing conducted by third-party security and risk management specialists. This process replicates the tactics threat actors employ to gain Initial Access[6], and extends beyond simple automated security scanners (which we also employ, of course).
No subprocessors	Keepit has no subprocessors in connection with the provision of the Keepit Services. We adhere to the strictest data protection standards and ensure that no one else will be accessing or handling your data. This guarantees compliance with local regulations and that your data will always be handled with the highest level of sensitivity.

Table 3—Keepit employs a number of operational security measures

6. See Initial Access within the MITRE ATT&CK Framework [MITRE]

Operational insights

To ensure reliable around-the-clock operation, we closely monitor our production systems' physical operational health and software stack health (Table 4).

Operational concern	Explanation
Physical operational health	We measure tens of thousands of datapoints every minute, aggregating them and making them instantly available to our operational staff, thereby ensuring efficient investigation and mitigation to meet our high uptime guarantees. System and data center performance is continuously compared to baseline thresholds, while our health monitoring system monitors the physical equipment (e.g., environmental, network, hardware, operating systems and services) with 30–second granularity, alerting on a range of unwanted situations (e.g., high temperatures in our data centers, failing disks in storage systems, congested network connections that threaten to impede ingress and egress, etc.)
Software stack health	Our Real-Time High-Frequency Event Monitoring (RTHFEM) system provides platform-specific deep stack insights on software components and their operational health. Our operations team has different dashboards and areas of interest they follow and monitor, and live metrics can even be shared across organizational groups to provide developers with real-world insights into how their code performs in production. The RTHFEM system provides insights into data ingress rates from different SaaS vendors and quickly identifies problems, should they arise within the software stack. Our backend engineers also use this system to troubleshoot specific customer-related situations based on anonymized views of the interoperability of each customer's connectors, data estate, and index. Additionally, detailed insights into software processes can assist with determining if an SLA is at risk or if an anomaly might be an indicator of malicious activity.

Table 4—Operational insights provide our team with timely feedback, visibility, and alerts

Looking for the current status and change history?
Visit **status.keepit.com** to view and subscribe to updates.

Software development life cycle

Keepit is dedicated to providing best-in-class solutions for cloud-to-cloud backup and data management. For us to make good on our promises, we are actively developing our platform in the pursuit of continuous improvement. Our software development lifecycle processes (Figure 8) ensure that security is incorporated from the inception of a new project and continued throughout the entire life of the system; these mature processes have been refined over the lifespan of the company, but we continue to invest in their ongoing evolution.
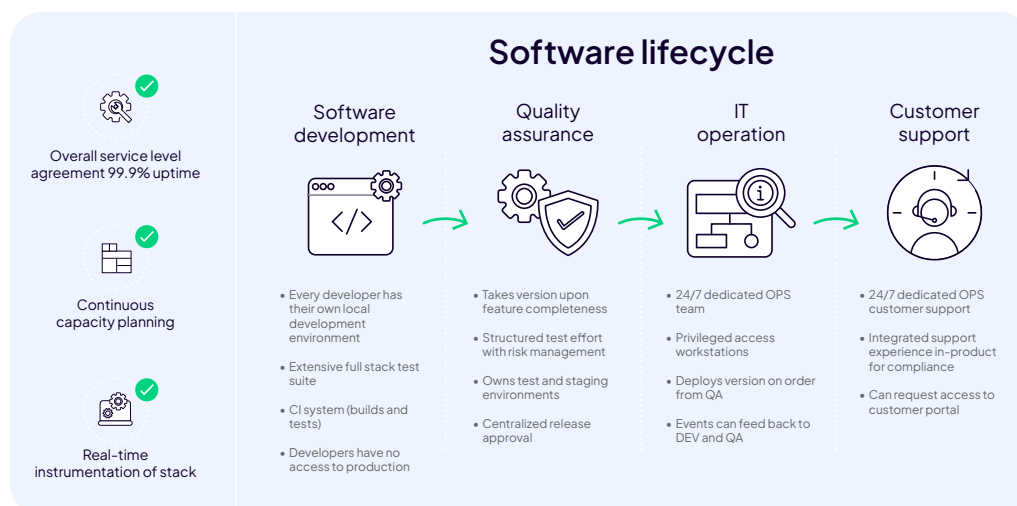

Figure 8—Keepit executes a mature and agile security-oriented software delivery methodology

Strong processes are a necessity for the secure development of new functionality on the Keepit platform. While information sharing is vital in any growing organization, segregation of duties is also an essential part of the actual execution of development, qualification, and—ultimately—deployment of software. We employ a number of technical and organizational measures to ensure that only the software approved by our Quality Assurance team goes to production. By enforcing significant workflow around "simple things" like software deployment, we ensure that we have a high degree of confidence in any and all software that is deployed to production and, in case a problem is discovered, that we are able to roll back a deployment or deploy hot fixes efficiently (once qualified).

Security features

In addition to the measures already presented (e.g., regional lock-in, dual data center redundancy), the Keepit platform also incorporates a number of additional security features to enhance protection of customer data. Two of the most important to govern are transport security and access control (Table 5).

Security feature	Explanation
TLS transport security	**TLS transport security** All primary workloads supported by Keepit employ modern Transport Layer Security (TLS) encryption on the endpoints used by Keepit and customers to access data. Both the Keepit Web UI and mobile app require the use of TLS. These communication protocols are designed to provide safe transport of data over insecure networks such as open Wi-Fi or broadband internet connections. **Network security** Additionally, to ensure a short path and the fastest possible data exchange with the primary customer data clouds for backup and restoration, Keepit connects directly to major internet exchanges.
Secure and logged access controls	**Single sign-on** Keepit's recommended best practice for deployment uses the customer's existing authentication infrastructure (such as AD/ADFS, Okta, etc.) by means of SAML integration. This approach allows customers to leverage the identity security measures they already have in place (e.g., multi-factor authentication). **Fine-grained role-based access control** Inside Keepit, an elaborate access control list (ACL) and role-based access control (RBAC) system allow fine-grained control over which identities can perform which operations on the customer tenant. For example, pre-defined policies permit an administrative role to manage the backup and a support role that can only restore data in place (as part of an internal IT support role) and cannot download or otherwise exfiltrate data from the platform. Additionally, separate groups of administrators can manage separate backup configurations; for example, one group can manage the Salesforce or Dynamics 365 backup while another manages the Microsoft Office 365 or Google Workspace backup. **Immutable system audit log** Finally, an audit trail is maintained for the duration of the customer engagement. The audit log can be viewed directly in the web application by authorized administrators and it can be accessed via the API for integration into third-party log analysis solutions. Data can also be pushed to Splunk, Sentinel or other third-party SIEM systems.

Table 5—Security features provide additional protection for data in transit and at rest

Conclusions

With increasing numbers of business-critical SaaS applications migrating to the cloud, and in a world in which new threats arise daily, continuity planning demands a secure SaaS data backup solution—and the SaaS vendors are clear that backups are your responsibility, not theirs.

The question then becomes: where to back up this critical data?

> "SaaS vendors are clear that backups are **your responsibility**, not theirs."

Keepit believes that real backup requires a separate logical infrastructure from the cloud in which the primary data is hosted. The most practical solution, then, is a dedicated cloud platform.

But simply having a dedicated platform is not enough: the platform must be secure. However, security is not something that can be bolted on as an afterthought—it needs to be designed into the very essence of a system. This secure-by-design philosophy guided the creation of The Keepit Cloud, which was architected from the ground up by information security and hosting experts to provide secure backup of SaaS data:

- Resilience and data immutability are achieved through immediate encryption of backup data, 30-day delete retention, data region sovereignty and dual data center redundancy, and a unique, tamper-proof infrastructure with immutability by default

- A blockchain-like Merkle tree architecture and innovative "incremental forever" data transfer ensure secure, reliable, and change-based long-term data storage

- All data center facilities employed by Keepit conform to a high physical security baseline and hold ISO 27001 certification plus complementary certifications (e.g., SOC-2, ISAE 3402, PCI/DSS, HIPAA)

- Technical and organizational measures are in place to ensure operational security, including a high degree of automation, comprehensive staff training, third-party security audits, proactive penetration testing, and detailed operational monitoring of the health of physical infrastructure and the software stack

- Mature software development lifecycle processes ensure that security is incorporated from the inception of a new project and continued throughout the entire life of the system

- To enhance protection of customer data, the Keepit platform incorporates TLS transport security and secure, logged access controls

These characteristics and others combine to form a SaaS data backup solution that is fundamentally different from and—we sincerely believe—superior to other products in the market, and that provides you with the benefits of offsite storage and the comforting knowledge that if your primary cloud is compromised, your backup is safe and secure elsewhere.

If you would like to know more about our solution, please visit our website at **keepit.com** or reach out to us at **sales@keepit.com**.

Keepit future proof

Take the next step toward protecting your SaaS data

Request demo

keepit

Keepit provides next-level SaaS data protection purpose-built for the cloud, by securing data in a vendor-independent cloud to safeguard essential business applications, boost cyber resilience and future-proof data protection.

Headquartered in Copenhagen with offices and data centers worldwide, over ten thousand companies trust Keepit for its ease of use and effortless backup and recovery of cloud data.

HQ – Copenhagen
Denmark
Keepit A/S
Per Henrik Lings Allé 4, 7.
2100 København, Denmark
CVR: 30806883
+45 8987 7792
sales@keepit.com

Dallas, Texas
United States
Keepit USA Inc.
3232 McKinney Avenue Suite 820
Dallas TX 75204, USA
TIN: 85–358 6335
+1 469–461–4892
sales@keepit.com

London
United Kingdom
Keepit Technologies UK, LTD.
Warnford Court
29, Throgmorton Street
London EC2N 2AT, UK
Company reg. no.: 13785045
sales@keepit.com

Munich
Germany
Keepit Germany GmbH
Maximiliansplatz 22
D–80333 München
Amtsgericht München
HRB 270094
Geschäftsführer Morten Felsvang
sales@keepit.com

www.keepit.com