Leonovus

Blockchain hardened ultra-secure
software defined object storage

A LEONOVUS WHITE PAPER

Solving the enterprise data storage problem

How compliant, secure and cost-effective storage of petabyte-scale data is possible

_ ABOUT THIS PAPER

The volume of data organizations have to store is already massive — and growing at a rapid pace. The need to store that data securely, cost-effectively and in compliance with internal and external regulations is straining many businesses and agencies around the world. This white paper examines the problem and explains how Leonovus addresses it with a single, flexible, blockchain-hardened software solution.

Note: This paper may contain forward-looking statements. These are not meant to imply any guarantee of product features, timing or release.

Contents

An unacceptable compromise

The specifics may vary, but observers agree the digital universe is going to keep expanding at a jaw-dropping rate well past 2020. According to insideBIGDATA, digital information is growing 10 times faster than traditional business data, and machine-generated data is growing five times faster than that.[1]

Companies already scrambling to manage data growth at double-digit compound annual rates won't be surprised by these stats. But the implications are bound to keep executives up at night, caught in a thought loop that goes something like this:

"The cloud would be so much more cost-effective for storing our data. But we can't choose the cloud because it won't give us the control to meet our compliance requirements. And it's not secure. But is our on-premises stora ge even secure enough? We can't afford to spend any more capex on IT. The cloud would be so much more cost-effective…."

And on it goes.

The reality is, security, compliance and ROI have historically been mutually exclusive. At best, enterprises could pick any two. That kind of compromise won't cut it today, not with regulations like Europe's General Data Protection Regulation (GDPR) carrying penalties of €10 million or more, cases like Equifax underscoring the high financial and reputational cost of a major data breach, and the average petabyte-scale business spending somewhere in the neighborhood of $18 million per petabyte a year on storage infrastructure and operations alone.

Secure, compliant data storage is straining organizations. Leonovus 3.0 eliminates the security–compliance–ROI problem with a flexible, easy-to-use blockchain-hardened software solution.

Pressure points

Corporate data volumes are skyrocketing in part because their sources and formats are diverse. Information is being generated by internal and external users, by people and machines, as documents, emails, images, video and more. The Internet of Things and data analytics will only drive those volumes up further.

1 insideBIGDATA. *The Exponential Growth of Data.* February 2017.

Laws and regulations like GDPR, HIPAA, PCI-DSS, Sarbanes-Oxley and Canada's Bill C-198 impose demands on organizations to store more of their data for longer periods, with high expectations of auditability and accountability and steep penalties for compliance failures.

This would be a non-issue in a world of limitless server capacity, free electricity and unbounded space, but when many companies' IT budgets are already constrained and shrinking, it poses a real and mounting problem. While virtualization had promised to ease the pain of IT infrastructure sprawl, it has often worked against compliance by edging data farther and farther out of the enterprise's direct control.

Then there's the small matter of security.

A breach is inevitable

Companies specialized in network and IT security routinely issue warnings about the world's constantly evolving threat matrix, which encompasses everything from corporate espionage, black hat hacking and government-backed intrusions to internal malfeasance and plain old human error.

A 2018 study published by the U.S. Identity Theft Resource Center (ITFC) found deliberate breaches have gone up by 45% since 2005, with businesses and healthcare organizations the two most affected types of entities.[2] These breaches come with hard costs and even harder consequences, affecting customer trust, public image, brand reputation and potentially even an organization's viability to continue doing business.

Arguably, the strongest (or at least, most realistic) position is to assume a breach will happen, and to have strategies in place to minimize the damage when it does. One thing is clear: today's threats demand an intrinsic, data-centric approach to security. The days of overlays and afterthoughts being enough are over.

Why the problem is hard to solve

If beauty is in the eye of the beholder, so are potential solutions to the enterprise data storage problem.

Looked at through a cost lens, going to the cloud is an obvious move since on-premises data volumes are growing faster than rack density. The sheer variety of on-premises solutions — storage area networks (SANs), flash

2 Identity Theft Resource Center. *2017 Annual Data Breach Year-End Review.* 2018.

storage, hyperconverged equipment — has introduced head-spinning complexity for IT teams. And even as environments are becoming more varied, IT departments are not necessarily growing their headcounts. They're often so absorbed in dealing with complex hardware they don't have the time or resources to manage and classify their data. In many organizations, the unwritten rule is to buy the biggest box with the most features and throw data at it — tantamount to buying the dealership just because you need a new car.

In light of all this, the appeal of the cloud, especially the public cloud, is clear. Public cloud storage is flexible, affordable and scalable. Sure, it might require a little more replication, and maybe it's a little less efficient as a storage mechanism, but the tradeoffs are worth it.

At least they seem to be until you switch perspectives and look at the situation through a compliance lens. Suddenly, the fact data is less reachable and traceable in a virtualized environment make the cloud seem risky and reckless.

Circling back to square one

Security experts would be inclined to agree. While public cloud providers do secure their environments and the majors make a point of promoting the ISO/IEC 27001 alignment of their platforms and infrastructure, it's impractical for them to assume liability for their customers' data. The responsibility for data protection remains with the enterprise, which in the end has the regulatory accountability. This creates a conundrum, since (as the compliance-minded people have already pointed out), the enterprise has less visibility into and control over its data in the cloud.

One option is for enterprises to go for a hybrid architecture that combines on-premises and cloud storage. Security-wise, this provides the reassurance that sensitive data is under the company roof. Compliance and data governance can be managed directly for anything stored in-house. And using the cloud for other, less sensitive data relieves some of the cost burden.

But even on-premises data is not perfectly secure: breaches happen wherever data lives. When on-premises environments become excessively complex, it's hard to be sure policies and compliance measures are truly being enforced. And the cloud introduces unavoidable latency that can affect enterprise applications.

It's easy for the situation to seem hopeless, or at least without a solution that doesn't involve a less-than-ideal compromise. There is, however, an alternative.

Secure, object-based cloud storage

What enterprises really need is some way of storing data that addresses compliance, security and cost together: a solution that provides a traceable, auditable way of ensuring regulations are adhered to, protecting data even if a breach occurs, and doing both while permitting flexible use of the cloud in whatever way best meets the needs of the business.

Some principles emerge fairly quickly from this. A winning solution must:

1. **Be data centric** — built around the data itself and reaching out from every conceivable endpoint across the storage architecture, on-premises and in the cloud.
2. **Provide full, end-to-end control** — allowing the organization's IT department to see, control, and be able to manage all data, wherever it is stored, which (looking ahead into the near future should include any number and variety of clouds.
3. **Be highly efficient** — which means both squeezing every byte of available storage out of all existing assets before acquiring new ones and handling the data intelligently to keep overall volumes down to the absolute minimum.
4. **Provide an irrevocable chain of evidence** — so that the provenance of data is unquestionable, and that if (or when) there is a breach or other form of data loss, the original can be retraced and recovered.
5. **Be easy to use** – allowing for transparent integration with existing applications and the enterprise operating environment, usable by non-technical employees, and eliminating restrictive vendor lock-in into the future.

The Leonovus solution

Leonovus 3.0 satisfies all five of these requirements through a platform that decouples structured and unstructured enterprise data from the corporate infrastructure, pulling it into a secure plane: the Leonovus Web Integrated Services Environment (WISE). The WISE Network™ reaches all of an organization's endpoints, down to the device and desktop.

Decoupling is the foundation of Leonovus' data-centric approach because it allows data and metadata to be treated on their own, irrespective of users, devices and corporate applications. All data is encrypted in flight and at rest,

and only authorized, authenticated users and services can access it in line with company policies and permissions.

Because the WISE Network™ is software defined framework, it can extend across any architecture: on-premises, hybrid cloud and multi-cloud. This not only gives organizations cost-effective options when deciding what storage to utilize but also provides scalability. When more storage is needed, a pricey build-out is no longer the only choice.

Through WISE and a unified dashboard, corporate IT has a complete view of all data. Controls and data governance policies and procedures are extended into the cloud and across multiple clouds as need be.

Figure 1. Leonovus storage architecture

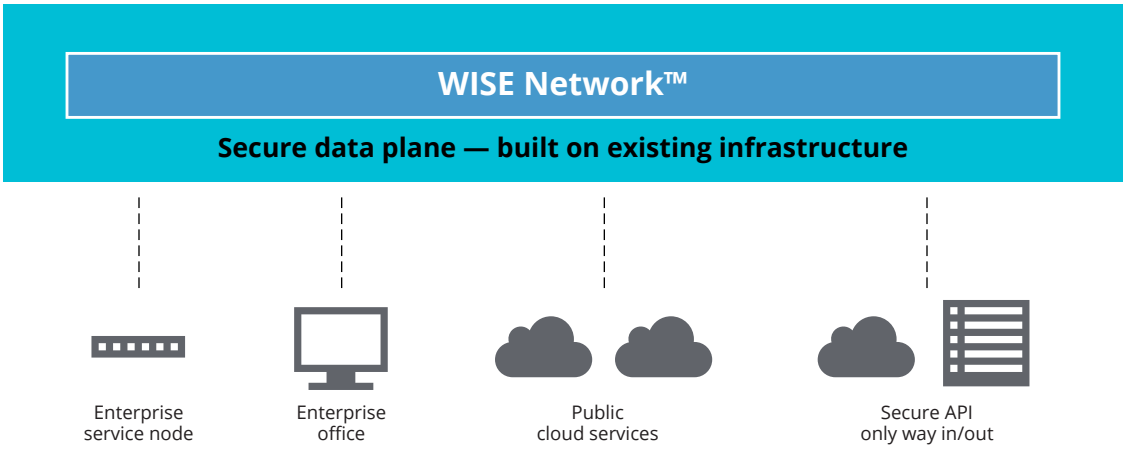

Military-grade security

No matter where data is saved — whether in nearline resources for active use, or cool or cold storage — Leonovus 3.0 ensures it is secured, logged and digitally signed. Mechanisms are built in to enable sharing and collaboration within the organization and with partners and suppliers outside.

Leonovus renders backed up, archived and bulk data hackerproof through a unique approach to object-based storage. Rather than store whole files on any one server, Leonovus 3.0 uses patented algorithms to "shred" data into discrete, encrypted objects and distribute it across multiple storage devices. The typical configuration is to create 25 objects out of a segment of a single file, but that threshold — and the number of objects a user needs to reconstitute the whole — are fully configurable based on an organization's security needs.

Because the data is broken up in this way and distributed across the storage architecture, even when a breach occurs, hackers who break into a single storage node do not possess enough objects to reconstitute a single file. The data is useless to them.

Figure 2. What hackers see with Leonovus

The same process used to reconstitute saved files is used to restore data if a node were to fail and where there would be risk of data is loss. Through its handling of metadata, Leonovus software automatically re-creates any missing objects to restore the original file. In other words, data can't be lost.

Blockchain-hardened compliance

Leonovus 3.0 uses enterprise-grade private blockchain technology as a platform to track and secure the metadata associated with all distributed data. For security, the metadata is stored separately from the data itself, within the same overall storage architecture and fully under the enterprise's control.

The blockchain protects the compass and the map that allow archived data to be reconstituted. It validates and verifies all metadata creation and distribution by establishing agreement across multiple nodes. It also generates an immutable record that begins the moment metadata is created — immutable because multiple nodes are keeping the same record and checking for concurrence.

This handling of metadata and immutability provides fully auditable records of all data, who has interacted with it, and where it's saved throughout the entire data lifecycle.

High-efficiency data storage

Leonovus 3.0 is hardware-, software-, and cloud-agnostic, which allows it to bring virtually any storage device or service within the secure enterprise infrastructure. This includes partitions on currently underused assets and paid, application-specific public cloud storage such as Microsoft OneDrive or Google Drive, which are bundled with software subscriptions and in many enterprises go up to 95% unutilized.

Leonovus software also allows for data classification and rules-based data management with aggregation and deduplication to manage storage efficiently and reduce the overall volume of data that needs to be stored. Because the WISE Network™ spans the entire enterprise, this deduplication is exponentially more effective than when deployed on a departmental or group basis.

Leonovus is well suited to any organization with strong need to protect valuable data in a controllable way including:

- **Large enterprises/institutions**
- **Financial organizations**
- **Health-sector organizations**
- **Law enforcement agencies**
- **Heavily regulated sectors**

How Leonovus can be used

In addition to simple storage applications, Leonovus 3.0 has any number of potential uses for organizations with petabyte-scale data, whether nearline, bulk, departmental, backup or archival. Four key scenarios associated with common business functions include:

Secure managed file transfers	Create internal and external transfer zones ('swim lanes') to exchange files securely within your organization and with partners and customers outside.
Mobile/desktop collaboration	Extend data protection to all users, including corporate and personal storage drives and devices.
Multi-cloud archiving	Aggregate and deduplicate data to manage the total volume stored. Classify data and set rules for automatic routing to the right storage tier based on its classification. Support Write-Once-Read-Many (WORM) zones across multiple storage nodes for long-term retention of critical data without any loss of integrity. Because blockchain provides an immutable record of provenance, even if data is ransom-wared, Leonovus 3.0 can retrace back to the last "clean" version and restore it.
Chain of evidence	Ensure data integrity and immutability through blockchain and digital security best practices.

Working with WORM data

Many organizations use WORM architectures to produce permanent, singular instances of data. In law enforcement, for example, eyewitness video from an incident must be kept in its unaltered, original form. When that video is produced as evidence, however, redactions may be needed to abide by privacy laws. With Leonovus3.0, the chain of evidence is fully intact, so the original version is always identifiable.

The simple way to compliant, secure, cost-effective storage

Leonovus 3.0 allows organizations to address all three requirements for compliance, security and storage ROI, enabling true, highly flexible hybrid and multi-cloud storage with operational efficiency, ease of use, and flexibility.

Because it is hardware-, software-, and cloud- agnostic, Leonovus 3.0 eliminates vendor and cloud lock-in, freeing organizations to choose whatever mix of storage solutions offers the best functionality for the lowest cost. At the same time, Leonovus 3.0 interfaces with all standard data management platforms, is AWS S3 compliant, and provides full data lifecycle support. It extends security and IT control from the cloud to the user device or desktop, easily integrating applications and requiring little or no end-user training.

For enterprises with petabyte-scale data storage needs that are unwilling to sacrifice compliance, security or cost, Leonovus 3.0 integrates seamlessly with the existing — and future — data storage environment.

About Leonovus

Leonovus is a global innovation company with offices in the United States and Canada. With more than 40 patents in process, Leonovus is a leader in the application of enterprise private blockchain and in architecting software-defined object storage for on-premises, hybrid, cloud and multi-cloud environments. www.leonovus.com

www.leonovus.com

Follow us on:

LinkedIn | Twitter | Facebook