

SHARDSECURE WHITE PAPER

Microsharding

March 2021

Jesper Tohmo

CTO and Co-founder

Christer Roslund

VP Engineering and Co-founder

Robert A. Clyde[†]

*ISACA Board Member and
ShardSecure Senior Advisor*



SHARDSECURE[®]

Introduction

Encryption has long been the primary method of protecting data. Now, at a time when a convergence of factors is putting more strain on companies in safeguarding their data, a promising new technology, Microsharding, has been introduced, that should cause companies to examine whether there might be a better path forward.

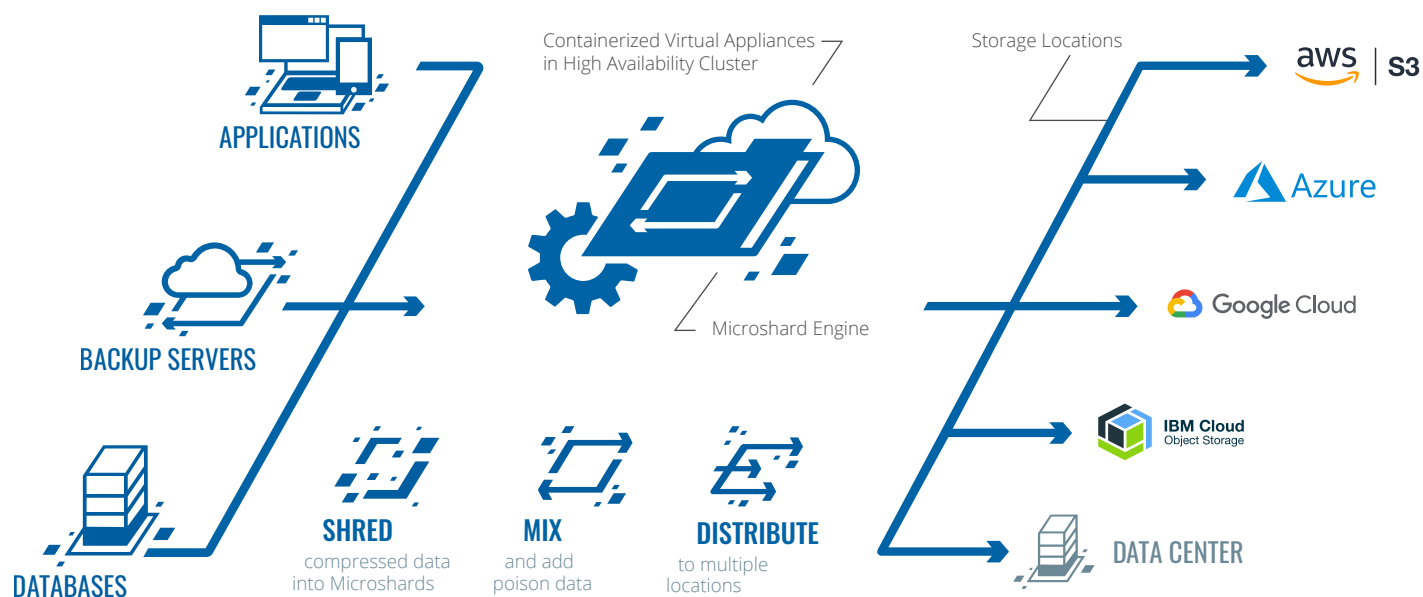
Most organizations are well along the journey of migrating their data and applications to the public cloud. The COVID-19 pandemic has prompted many organizations to accelerate their cloud usage, given the convenience of cloud access for their increasingly remote workforces.

This is all unfolding at a time when data protection and data privacy have never been more important, and privacy regulations such as the EU's GDPR and California's CCPA and CPRA have further elevated their prominence. The recent EU court ruling striking down Privacy Shield makes a challenging data protection and compliance environment even trickier.

To cybersecurity veterans, it is no surprise that data protection has emerged as a hot topic. With no shortage of examples of the steep consequences caused by not implementing adequate data security measures, professionals and organizations recognize its importance and are acting accordingly to bolster security efforts.

New benchmarks and technical privacy certifications such as the Certified Data Privacy Solutions Engineer (CDPSE) have been released by ISACA (Information Systems Audit and Control Association) in 2020. For enterprises around the world, determining how to best protect their organization's data while also ensuring data stored in the cloud is compliant with data protection and privacy regulations is a central business imperative.

Introducing Microsharding



Microsharding has emerged as a leading solution to meet this rising demand for ensuring data protection.

Many security practitioners are familiar with sharding, which has long been used by storage and database companies like Oracle, AltiBase and MongoDB, and tools like ElasticSearch and MySQL, to improve performance by splitting files into multiple pieces and storing them in different locations so that input/output (I/O) can be done in parallel to make the process faster. Normally these pieces are a few thousand to a few million bytes in size.

Microshard™ technology, a newer approach used by ShardSecure®, also splits a file up into multiple pieces, but the pieces are extremely small. Shards can be as tiny as a single byte, but practically speaking, each microshard tends to be a few bytes. Microshards are stored in different locations including across multiple cloud providers as well as on-premise locations.

The benefits of Microsharding are multi-faceted. Not only does the technology help organizations address a critical need to effectively protect

the data of customers and employees, it also provides significant value from a regulatory standpoint. Data privacy and security professionals also value the way Microsharding drastically reduces the attack surface by shredding and distributing data, another key advantage to the approach.

Significantly, industry experts accept Microsharding as a means to reduce what is in scope for data sensitivity; if the data has been shredded and scattered across multiple storage locations to the extent that a bad actor can't extract a single credit card number or Social Security number, is the data in a single storage location still sensitive data? Most contend the answer is no, which can dramatically reduce companies' data protection burden and the cost of compliance. As cyber audit and assurance firm, UHY Advisors asserted, "In our opinion, microsharded sensitive data is no longer sensitive. As a result, ShardSecure has the potential to lower cyber risks and compliance costs while maintaining compliance with the spirit of European and US data protection regulations."



Outlook: Microsharding and Legacy Security Solutions

With Microsharding, if the cloud storage administrator is hacked – a common concern in complying with regulations such as GDPR – the data is not in jeopardy.

In a bigger picture sense, in almost all cases with adding privacy and security, there is a performance cost. Implementing effective security measures tends to slow things down. Conversely, Microsharding represents a rare case in which layering on additional data protection also improves overall performance.

There are also distinct benefits of Microsharding compared to utilizing encryption. While encryption provides adequate protection in many cases, Microsharding allows companies to avoid worrying about the key management concerns that accompany encryption. For example, if an encryption key is lost, the data is also lost. Moving to a new key or encryption algorithm requires decrypting and re-encrypting

all the data, often a costly and time-consuming exercise.

Additionally, certain encryption algorithms might not be quantum-resistant, an increasingly important consideration with quantum computing looming as a probable game-changer for the industry. Imagine an organization with terabytes or petabytes of data that must be put through the process of decrypting and re-encrypting, at tremendous cost, due to broken encryption.

Any organization dealing with large amounts of structured or unstructured data – for example, using Microsoft 365 applications that include research and sensitive data – should strongly consider Microsharding, which can be particularly effective in sectors such as healthcare and for law firms.

Already, Microsharding is beginning to catch fire on a large scale. The timing is right for Microsharding to become a promising alternative or supplement to encryption.

Reducing the Attack Surface and Sensitivity of Data in the Storage Area with a Revolutionary Approach to Sharding

Microsharding provides unique data protection, privacy and security advantages as compared to traditional sharding methods.

When used to improve performance, legacy sharding methods involve splitting files or volumes into multiple pieces that are a few thousand to a few million bytes in size. For context, even a single kilobyte fragment is large enough to contain 111 social security numbers.

ShardSecure's technology is a revolutionary approach that breaks data into extremely tiny fragments too small to be valuable to malicious actors. ShardSecure is the only solution capable of breaking data into single-digit bytes without sacrificing performance. ShardSecure then distributes these Microshards to multiple locations that could include cloud providers like Amazon, Microsoft or Google, as well as on-premises locations. Importantly, shard locations are unrelated and not known to each other.

Thus, an attacker intercepting Microshard data has no way to put the pieces back together because they will always have an incomplete set.

The ShardSecure approach is in contrast with encryption, in which the full set of data is compromised and needs to be unscrambled. Unscrambling data requires time and compute power. Reconstituting data fragments requires most or all of the data fragments, something the attacker cannot obtain without compromising all possible storage locations everywhere. ShardSecure's Microshard technology changes the attacker's challenge from a time and compute power problem to a time, compute power, and spatial problem.

Encryption may slow an attacker down, but Microshard data protection persists over time. Faster computers won't help an attacker against ShardSecure, not even quantum computers, as they simply do not have the data localized to unscramble.

SHARDING




Sharding is the process of fragmenting data down into small shards, then distributing the shards to multiple storage locations.

TYPICAL SHARD



Typical shards are between a few thousand to few million bytes. A single kilobyte fragment is large enough to contain 111 social security numbers.

MICROSHARD



ShardSecure's unique microshards can be as small as single-digit bytes, too small to reveal a social security number, credit card number or even email address.



Securing the ShardSecure Software Appliance

ShardSecure software uses industry-standard commercial compression and tokenization technologies to secure our pointer database.

Should attackers get a copy of the ShardSecure pointer database, and know the algorithms for the computed portions of the pointers, they still would not be able to find the remote data fragments. A key element of each pointer (the remote host's address) has been replaced by a token. An attacker might learn that a data fragment was in location 7, but without the mapping file to tell them where location 7 is (perhaps a file in Google Drive or a bucket in Amazon S3), they cannot reassemble the data. The ShardSecure token mapping file persists in a different location from our database for added security.

In addition to tokenizing our pointers, access to the ShardSecure appliance (whether physical, virtual, or hosted) can be limited and secured from "front door" attacks while the appliance secures data from "side door" attacks. Protecting a single entry point is far easier than protecting access to all data everywhere. This is particularly true when the data to be protected from unauthorized access is created dynamically by different groups in your organization. Rather than tracking and protecting each file, you simply limit access to the fixed "front door" of the ShardSecure appliance.

As outlined in this white paper, ShardSecure software dramatically shrinks the attack surface from all data everywhere to a single front door of an appliance. In doing so, the solution can focus controls on this one point without having to worry about where large data sets may sit remotely (such as cloud environments).



Data Security for the Cloud

ShardSecure can be deployed to secure on-premises and cloud data, giving organizations the freedom to embrace the infrastructure that best suits their needs, knowing data sensitivity is eliminated in the back-end cloud infrastructure.

Microsharding addresses key vulnerabilities introduced by public cloud adoption, for example misconfigurations that can leave data exposed in storage buckets.

ShardSecure provides a unique focus on securing data on back-end cloud infrastructure, where privileged cloud administrators perform important daily activities including patch management, software updates and other critical tasks that bear serious consequences in the event of data breaches. Even if encryption is used, the cloud admins often have access to the keys. Whereas legacy solutions have provided little in the way of back-end cloud

data security, Microsharding separates sensitive data from privileged administrators, who could be compromised, disgruntled or simply make mistakes unintentionally that cause data breaches, and is an excellent way to achieve zero trust in data security.

Microsharding reduces the huge attack surface of applications in the cloud and the entire data storage area (gigabytes to petabytes) to the small attack surface of the microshard engine, pointers, host map file and applications (mere megabytes). If attackers breach a cloud administrator's account for one cloud provider, the microsharded data on that cloud provider cannot be used to reconstruct any files or even a small amount of sensitive information

Even if attackers breached all the enterprise's cloud storage, the microsharded data could not be put back together without access to the microshard engine, pointers, and host map file. As mentioned in the previous section, cloud admins and nearly all other admins should never have access to any of these items.

Accelerating Cloud Migration

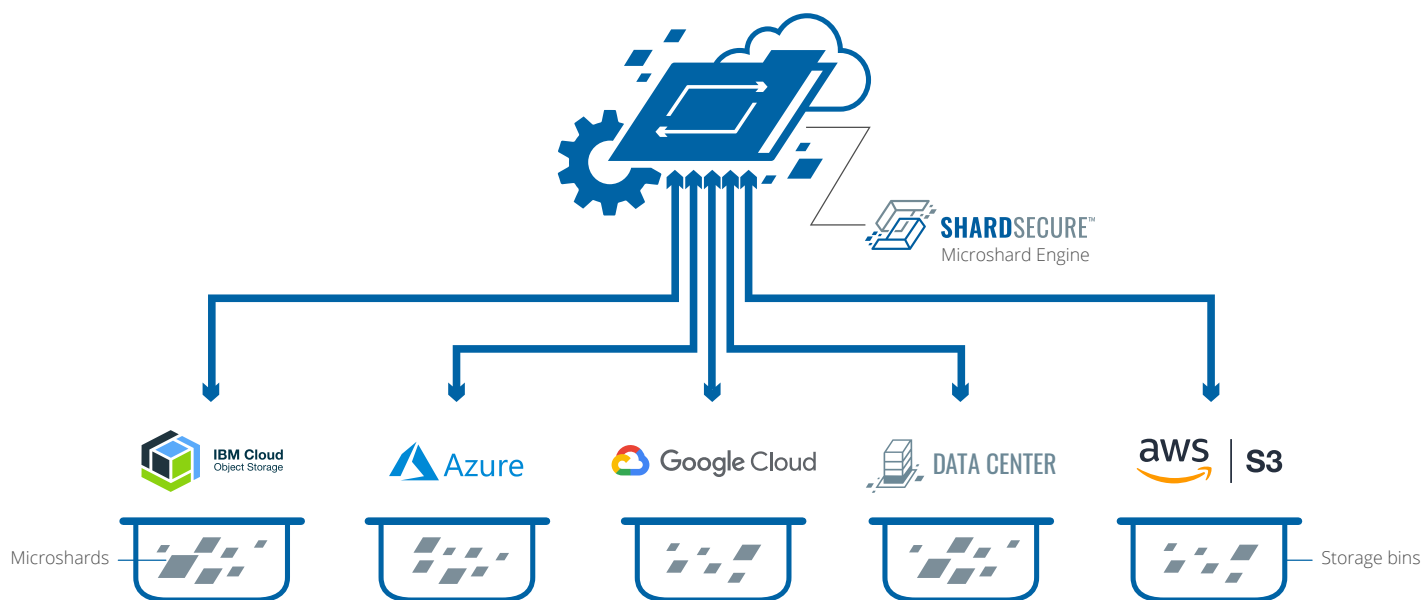
One of the biggest hurdles to embracing cloud is the challenge it presents to IT stakeholders.

Cloud providers each present their own landscape to navigate, requiring proprietary configuration and best practice insight, and making it difficult to move data to, between and out of different cloud providers. It is no surprise that misconfigurations remain the leading cause of cloud data breaches, with Gartner estimating that as much as 95% of cloud security failures can be attributed to customer faults.

ShardSecure addresses this issue head on by eliminating the sensitivity of data before it is sent to cloud storage to mitigate the risks caused by exposure, distributing Microshards to drastically reduce the attack surface, and by providing a layer of abstraction between applications and the storage locations.

Once ShardSecure is implemented, an organization can easily migrate data to a new location, such as from on-premises to the cloud or to a new cloud provider. This can be done without having to modify the applications. Once the data is fully migrated the old storage location can be decommissioned.

ShardSecure can be deployed easily and quickly, without having to turn off applications or otherwise interrupting IT operations. The performance benefits are not limited to deployment, as the use of parallel reads and writes can improve performance in the cloud. Those familiar with traditional data security methods such as encryption are aware of how unique it is for a security solution to augment performance, rather than create additional lag. As a result, ShardSecure provides a compelling solution for cloud migration, backup and DevOps security, ensuring both data privacy and fast restore.





Reducing Compliance Burden

ShardSecure so effectively eliminates the sensitivity of data that it also provides key benefits from a regulatory standpoint.



Regulations such as GDPR, HIPAA, CCPA, CPRA, and PCI DSS, can be costly to comply with and expensive if their standards are not upheld. However if data that has been shredded to the extent that a bad actor is unable to extract even a credit card number or Social Security number, it can be argued that the data is no longer classed as 'sensitive', dramatically reducing companies' data protection burden and the cost of compliance. For example, if the cloud storage administrator is hacked – a common concern in complying with regulations such as GDPR – the data that has been Microsharded is not in jeopardy.

Microsharding can reduce the scope of storage locations that must comply and scope of

regulatory and audit reviews. For instance, the cloud storage locations that store the disparate microshards might no longer have to comply with GDPR – none of them contain any data that would be considered sensitive. Of course, the applications that use the data reconstructed by ShardSecure from the microshards would need to comply.


By virtue of eliminating the sensitivity of data, Microsharding makes it easier for organizations to store larger quantities of data, for longer, knowing they are not multiplying the attack surface or exposure risk along with the data quantities. This includes the long-term security of storage backups, an often neglected category of data as compared to production data and a company's "crown jewels". Organizations can confidently store large volumes of Microsharded data on-premises, or affordably in the cloud, long-term thereby improving audit outcomes and helping ensure compliance.

Conclusion

	BEFORE SHARDSECURE 	AFTER SHARDSECURE 
ATTACK SURFACE FOR DATA BREACH	<ul style="list-style-type: none"> ■ All data storage areas ■ All admins with access to data storage area (likely many) ■ Cloud provider ■ Applications using the data ■ Users with access to the applications 	<ul style="list-style-type: none"> ■ Applications using the data ■ Users with access to the application ■ ShardSecure admin
DATA SENSITIVITY SURFACE	<ul style="list-style-type: none"> ■ Terabytes to exabytes ■ All data storage area ■ Applications using the data 	<ul style="list-style-type: none"> ■ Megabytes to gigabytes ■ Applications using the data
COMPLIANCE AND IT AUDIT SCOPE	<ul style="list-style-type: none"> ■ Many admins ■ All data storage areas ■ All admins with access to data storage area (likely many) ■ Cloud provider ■ Applications using the data ■ Users with access to the applications 	<ul style="list-style-type: none"> ■ Few admins ■ Applications using the data ■ Users with access to the applications ■ ShardSecure admin
RESISTANCE TO QUANTUM COMPUTING ATTACK	<ul style="list-style-type: none"> ■ Low 	<ul style="list-style-type: none"> ■ High
DATA ACCESS BY CLOUD OR STORAGE ADMINS	<ul style="list-style-type: none"> ■ Complete access to data 	<ul style="list-style-type: none"> ■ No access to data (zero trust)

† Any opinions or contributions by Robert A. Clyde made in conjunction with producing this paper are his own and do not represent an official position or endorsement by ISACA.

Learn more about Microsharding at
shardsecure.com/get-started



SHARDSECURE®

☎ +1 (800) 760 9445

✉ info@shardsecure.com

🐦 @ShardSecure

101 Avenue of the
Americas, 9th Floor
New York, NY 10013
United States of America