

**TAG CYBER**

**ESTIMATING THE  
SECURITY RETURN  
ON INVESTMENT (ROI)  
FOR CLOUD  
MICROSHARDING**

EDWARD AMOROSO, CHRISTOPHER R. WILDER  
TAG CYBER

**SHARDSECURE**

# ESTIMATING THE SECURITY RETURN ON INVESTMENT (ROI) FOR CLOUD MICROSHARDING

EDWARD AMOROSO, CHRISTOPHER R. WILDER

---

Under reasonable assumptions, enterprise teams employing microsharding for cloud-hosted data will see a reduced likelihood of data compromise incidents in public cloud service infrastructure. The associated reduction in response costs is shown to be sufficient to justify purchasing a commercial platform license. This implies that enterprise teams storing critical data in the cloud would be wise to invest in a microsharding solution.

## 1. INTRODUCTION

The risk of hosting sensitive data in public cloud infrastructure has begun to diminish gradually as practitioners obtain access to new tools, platforms, services, and methods for improving security. Commercial vendors now provide excellent solutions to address cloud entitlement management, cloud access visibility, cloud security posture management, and cloud data encryption.

Each of these controls, however, focuses on the risks of unauthorized access to cloud data through front-end access channels. One can view these security methods as focus on protecting access through the normal user interface. While this is obviously important and necessary, such solutions have not typically addressed the uncomfortable issue of back-end access to data through cloud administrative access with insider privileges.

As a result, when critical data is hosted in a public cloud service, it is likely that back-end access by potentially compromised or disgruntled insiders might not include proper controls. Accidental data compromise by administrator mishandling is an additional use-case to include. Security teams can address these risks through service level agreements (SLAs) with their cloud providers, but functional controls are more desirable.

A promising technique being used to address this cloud hosted data compromise risk is *microsharding*. Derived from data management algorithms, the method is designed to address the cyber risks of back-end access to data hosted in cloud infrastructure. Conveniently, the approach is most well-suited to multi-cloud hosting, which has become common in most modern hybrid architectures.

This report provides a brief overview of microsharding and then uses a methodology to estimate the return on investment (ROI) of using the technique in a typical cloud setting. The analysis shows that the high preventive benefit of microsharding is sufficient to warrant investment under normal circumstances. This claim is supported by analyzing the impact of microsharding on the various cost variables to host data in a public cloud.

## 2. HOW CLOUD MICROSHARDING WORKS

The microsharding process is designed to break up important data into multiple components that are separated, obfuscated, and stored across disparate cloud infrastructure. The result is that back-end access to the data by administrators and other cloud hosting insiders cannot result in a data breach because the data has been sharded across multiple cloud storage entities (see Figure 1).<sup>1</sup>

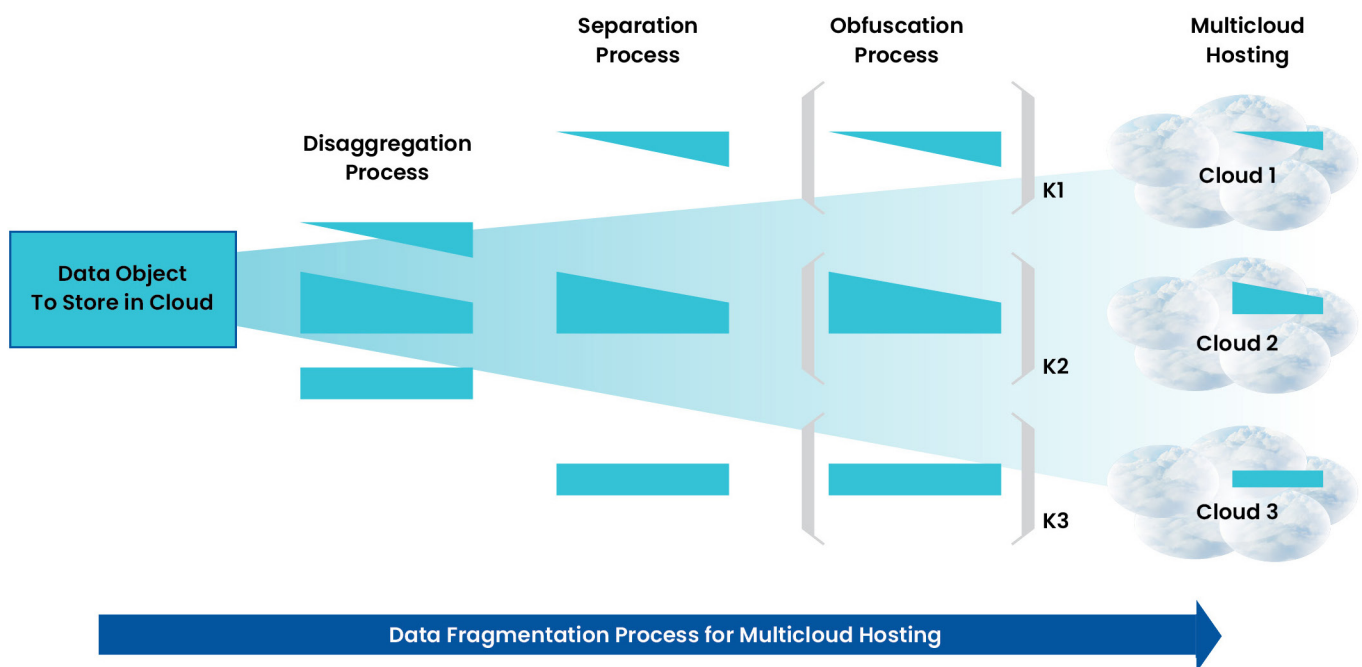


Figure 1. Microsharding Process

The components of the microsharding process are arranged into a pipeline that results in the cloud data protection. Each of these component processing tasks contributes to the overall security scheme. Insights into the algorithmic strategy for these processes are listed below:

- **Disaggregation** – Breaking up data to be cloud-stored into constituent parts is an important aspect of the microsharding process. By disassembling in this manner, the threat of direct, back-end access by an insider or an intruder is greatly reduced.
- **Separation** – The separation of the disaggregated components is a related task that further drives down the risk of unauthorized access using back-end channels with administrative access.
- **Obfuscation** – Obfuscation refers to the process of rendering each disaggregated data shard undiscernible on inspection. This can be done with encryption, blinding algorithms, and other practical means.

The most unique aspect of the microsharding process is that it complements the plethora of cloud security solutions that address front-end risk. That is, existing cloud security is all about ensuring that the overt path to hosted data is controlled by identity, access, encryption, and other policy enforcement mechanisms. Microsharding provides protection for back-end access to the data by administrators and other cloud hosting personnel.

### 3. MICROSHARDING ROI ANALYSIS

To measure the ROI for cloud microsharding requires definition of the use-cases being considered in the analysis. Two situations arise in which microsharding becomes important: An enterprise might be moving critical data into the cloud from their existing on-premises environment or an enterprise that is already hosting critical data in one or more public clouds might be considering microsharding to improve security.

Both situations are important, but we focus on the latter case where microsharding is being introduced for existing cloud-hosted data. This case is attractive for ROI estimation because it allows for straightforward *before* and *after* costs calculations. That is, by identifying the relevant cost variables for cloud hosting administration and security, we can estimate the differences in overall costs when microsharding is and is not being used.

Readers should keep in mind, however, that use of microsharding in the first case – where an enterprise is moving its critical data to cloud from premises – will have ROI savings over not including this step. The primary issue is that response costs for public cloud incidents, possibly involving some insider-initiated breach of the provider, can be so high that avoidance is imperative, from a purely financial perspective. By microsharding stored data in public cloud, direct access by administrators with privileged access will not reveal useful information.

#### 3.1 ROI Methodology

We will thus begin our analysis by first examining the existing costs of an enterprise hosting critical data in public cloud without microsharding, and by then examining the same set of costs with microsharding. The analysis should be applicable to all types of critical enterprise data, regardless of the domain (e.g., finance, health care, government). The costs and benefits associated with microsharding of important data in the cloud will be categorized as *small*, *medium*, or *large*.<sup>2</sup>

#### 3.2 Baseline Cost Equation

We will perform the ROI analysis using a fictitious medium-sized enterprise called ACME that hosts a critical business application called AppX in the public Amazon Web Services (AWS) infrastructure. The

critical data associated with AppX is stored as S3 Objects in Amazon S3 Storage Buckets. ACME's high-level cost equation for this AWS-hosted data can thus be represented as follows:

$$\text{Costs(AppX)} = \text{DevOps Costs(AppX)}^3 + \text{Admin Costs(AppX)}^4 + \text{License Costs(AppX)}^5 + \text{Security Costs(AppX)}^6$$

The size of DevOps and Admin costs will depend on the local environment and where it resides in the maturity lifecycle. Organizations that have more experience and expertise in these areas will have lower costs than ones just getting started. We will assume that no special licensing situations exist, and that ACME is paying industry-average license fees for any data hosting-related services.

### 3.3 Security Costs

The protection costs for cloud hosted applications can be partitioned into prevention costs and detection/response costs. Prevention costs are designed to avoid cyber threats (often referred to as shifting left), whereas detection/response costs are intended to deal with on-going attacks (referred to as shifting right). We can therefore adjust ACME's cost equation for AppX to include the following variables:

$$\text{Costs(AppX)} = \text{DevOps Costs(AppX)} + \text{Admin Costs(AppX)} + \text{License Costs(AppX)} + \text{Prevention Costs(AppX)} + \text{Detection/Response Costs(AppX)}$$

Cyber security experts agree that detection/response represents the highest security cost. They also agree that whenever possible, front-end prevention investments will reduce back-end detection and response costs. This is a not a controversial concept, but it only works if the prevention is effective. Where this equation becomes questionable is when the prevention has no impact (e.g., early generations of software antivirus).

### 3.4 Estimating Cost Values

We can start by associating high, medium, and low values with the cost equation for ACME. Before microsharding is introduced, we can assume medium costs in all cases to reflect the generic, industry-consistent aspect of the use-case. This results in the cost equation for ACME cloud data hosting looking as follows:

$$\text{Costs(AppX)} = \text{DevOps Costs(AppX)} + \text{Admin Costs(AppX)} + \text{License Costs(AppX)} + \text{Prevention Costs(AppX)} + \text{Response Costs(AppX)}$$

$$\text{DevOps Costs(AppX)} = \text{Medium}, \text{Admin Costs(AppX)} = \text{Medium}, \text{License Costs(AppX)} = \text{Medium}, \text{Prevention Costs(AppX)} = \text{Medium}, \text{Response Costs(AppX)} = \text{Medium}$$

$$\text{Costs(AppX)} = \text{Medium} + \text{Medium} + \text{Medium} + \text{Medium} + \text{Medium}$$

We can now associate simple numeric cost values (1, 2, 3, 4, 5) in an ordinal scale (where respective values are ordered, but not associated with an arithmetic) with (*very low, low, medium, high, very high*) cost estimate categories. While ordinal values should not be used for detailed calculations, they can help to illustrate comparative aggregate cost estimations.

$$\text{Costs(AppX)} = 3 + 3 + 3 + 3 + 3 = 15$$

Obviously, the use of ordinal values in the above cost equation just provides a simple basis for comparison, because no absolute meaning can be associated with each value other than their respective order (e.g., 15 is greater than 3). What we can do, however, is show how changes in different cost variables might add to or offset corresponding changes in other cost variables.

### 3.5 Estimating Microsharding ROI

The primary cost assumptions being made with respect to the use of data microsharding in the cloud can be stated as follows: When microsharding is introduced using a second cloud (e.g., Microsoft Azure), the following cost impacts should be expected:

**DevOps: No change (3)**

**Admin: No change (3)**

**License: Increase for microsharding tool plus second AWS cloud (3 to 4)<sup>7</sup>**

**Prevention: No change (3)**

**Response: Decrease to reflect reduced incident likelihood (3 to 1)**

Using these estimates allows for a simple use-case representation of introducing microsharding to a typical environment.

**Pre-Microsharding:**

**Costs (Appx) = 3 + 3 + 3 + 3 + 3 = 15**

**Post-Microsharding:**

**Costs (Appx) = 3 + 3 + 4 + 3 + 1 = 14**

Again, ordinal values cannot be compared accurately, so the two equations should be interpreted to show that as license costs increase, corresponding response costs decrease. In the case shown, the response costs decrease exceeds the license cost increase, so the ROI will be positive. Our analysis at TAG Cyber suggests that the positive case is likely under reasonable assumptions (see below).

### 3.6 Using More Realistic Estimates to Calculate ROI

If we introduce more realistic cost assumptions for a typical enterprise hosting data in a public cloud, then we can use the ROI methodology introduced above to illustrate more meaningful sample values. Included are the specific costs changing for ACME doing the microsharding across AWS and Microsoft, the license cost increases (for the additional cloud and the microsharding tool), and the corresponding response costs decrease:

**Pre-Microsharding:**

**License: Assume 500TB at \$120K/year to AWS**

**Response: Assume team of 5 FTE at \$1M/year loaded response costs**

**Total License + Response: \$1.12M**

**Post-Microsharding:**

**License: Assume 250 TB at \$65K/year to AWS and 250 TB at \$65K/year to MS**

**License (Microsharding): \$100K/year (estimate based on ShardSecure pricing)**

**Response: Reduce two FTE @ \$200K loaded salary in response team, resulting in three FTE**

**@ \$200K loaded salary for response team = \$600K/year**

**Total License + Response: \$830K**

The implication of the analysis is that in this sample hypothetical case, ACME can reduce two full-time equivalent staff members on its incident response team due to a significant reduction in assumed data compromise exploits in public cloud. Despite minor increases in public cloud licenses for AWS, and the license costs for a microsharding tool such as ShardSecure, the overall costs for ACME will go down because of the decision to implement the security control.

## 4. ACTION PLAN FOR ENTERPRISE

An action plan for enterprise based on this ROI assessment should involve the following management steps to take full advantage of the microsharding benefit illustrated above:

- **Step 1: Inventory.** The enterprise team should take inventory of any critical data that is either already hosted or planned for hosting in a public cloud service.
- **Step 2: ROI Instantiation.** The enterprise team should use the sample generic ROI analysis presented above as basis for a more realistic analysis using actual costs being incurred.
- **Step 3: Platform Review.** The enterprise team should scan the commercial landscape for suitable microsharding tools to be reviewed and tested (such as ShardSecure, referenced throughout this report).
- **Step 4: Implementation Plan.** Assuming the ROI and vendor reviews proceed successfully, the enterprise team should begin an implementation plan to integrate microsharding into the data hosting infrastructure.

---

<sup>1</sup> The analysis of microsharding benefited from cooperation with the ShardSecure team which offers a commercial platform supporting this technique. The analysis was completed with their expert assistance to ensure that no technical assumptions were made that were either incorrect or unrealistic. The final ROI conclusion was developed independently by the TAG Cyber team, however, and was not pre-determined by ShardSecure or any other commercial participant in this industry.

<sup>2</sup> Enterprise teams using this ROI report can and should translate these broad groupings into more specific quantitative values. Also, any cost reductions that can be made without commensurate negative impact on cyber risk will be interpreted as cost benefits.

<sup>3</sup> DevOps Costs are presumed to include all development, test, deployment, update, and bug fixes.

<sup>4</sup> Admin Costs are presumed to include day-to-day care, maintenance, monitoring, and support.

<sup>5</sup> License Costs are presumed to include any fees paid to third parties for software, hosting, or other services.

<sup>6</sup> Security Costs are presumed to include any tasks consistent with prevention, detection, and response to cyber threats.

<sup>7</sup> This is an aggressive estimate that assumes the cost increase. Actual negotiation might produce a better result.

## ABOUT TAG CYBER

TAG Cyber is a trusted cyber security research analyst firm, providing unbiased industry insights and recommendations to security solution providers and Fortune 100 enterprises. Founded in 2016 by Dr. Edward Amoroso, former SVP/CSO of AT&T, the company bucks the trend of pay-to-play research by offering in-depth research, market analysis, consulting, and personalized content based on hundreds of engagements with clients and non-clients alike—all from a former practitioner perspective.

Copyright © 2022 TAG Cyber LLC. This report may not be reproduced, distributed, or shared without TAG Cyber's written permission. The material in this report is composed of the opinions of the TAG Cyber analysts and is not to be interpreted as consisting of factual assertions. All warranties regarding the correctness, usefulness, accuracy, or completeness of this report are disclaimed herein.