TAG CYBER

Data Sharding for Back-End Cloud Security

Data sharding for back-end cloud security addresses the threat of compromised insiders with privileged access. The method disaggregates, separates, and obfuscates data so that insiders within cloud service infrastructure cannot make sense of stored assets.

Prepared by
Dr. Edward G. Amoroso
Chief Executive Officer, TAG Cyber LLC
eamoroso@tag-cyber.com

Version 1.0
April 27, 2020

Introduction

The challenge of protecting enterprise assets in public cloud infrastructure has emerged as one of the top concerns for the modern Chief Information Security Officer (CISO). This contrasts with earlier debates about whether to allow for such external hosting of internal resources. Those discussions have completed – and almost uniform agreement now exists that cloud infrastructure will be used for at least some applications, often in hybrid mode.

The cloud security challenge is typically represented as follows: An employee needs access to an enterprise resource hosted on a public cloud. Since the cloud is Internet accessible, the access path from the employee is publicly visible to attackers. End-to-end controls such as cloud access security brokers (CASB), multi-factor authenticated login, and microsegmented workload architectures are thus placed in-line to reduce cloud access risk.

The primary threat being addressed with this approach involves non-authorized individuals somehow gaining access to the cloud-hosted resource. This might involve use of credential theft, identity spoofing, or some other malicious means to trick the cloud host into believing that the request is valid. It's a tough risk to manage, because many attacks, such as credential theft, are performed outside the purview of the cloud.

In this report, we examine a related threat, but one that could be more serious, and that is almost certainly being paid less attention by enterprise cyber security practitioners. The risk involves *back-end access* to cloud-resident data by individuals, groups, and even automation without proper authorization or purpose. Many of the use-cases for back-end threats involve compromised administrators, but unintended or erroneous actions are also common.

Front-End Cloud Security Protection

As alluded to above, the most familiar security approach used to protect cloud infrastructure involves *front-end* protection, because it addresses the normal interface used by employees and other authorized individuals for access to cloud-hosted resources. The protections sit in the end-to-end access path to authenticate users, detect anomalies, generate telemetry, support mitigation, and provide a means for performing incident response and forensics.

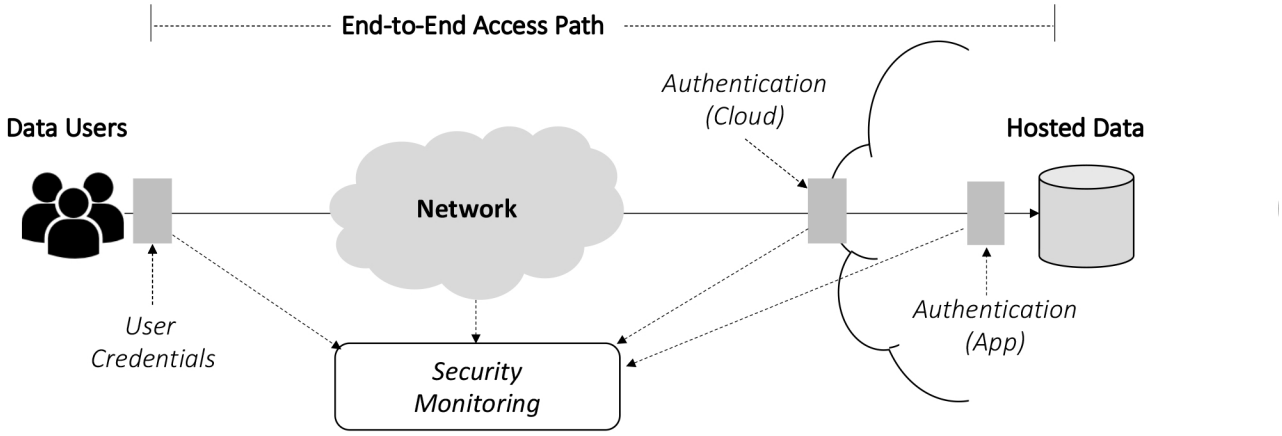

Figure 1. Front-End Security Protection for Cloud

It is worth mentioning that this front-end view also includes automated systems accessing cloud assets. Machine-to-machine transactions increasingly use cloud systems through an application programming interface (API). The security protections for APIs often include comparable controls for human users. These include gateways, authentication, and monitoring, and can be used for both IT and operational technology (OT) resources in cloud.

This view does, however, ignore a significant aspect of cyber risk that is arguably more intense in its consequence for cloud-hosted resources. This aspect involves the administration of cloud infrastructure, including the management of resources, assets, and data. The challenge of protecting administrative access can be described accurately as back-end protection, because it does not sit in-line with normal user or even machine-based access over an API.

Back-End Cloud Security Protection

The functions typically involved in back-end access to cloud-hosted resources include the day-to-day tasks that are necessary to ensure proper hosting and a good user experience. This can include performance tunings, patch management, software updates, new feature introduction, and so on. These tasks generally require the highest level of privilege (e.g., Unix root) so the side-effects of an administrator becoming disgruntled or just making an error can be significant.

The types of security controls that are common to reduce risk in back-end cloud infrastructure settings often parallel the ones used to protect front-end access, but that are simply re-positioned for administrative access. This includes authentication, privilege management, and activity monitoring. These are excellent controls to detect improper access to administrative accounts, but they are weak controls to deal with a compromised or hacked administrator.

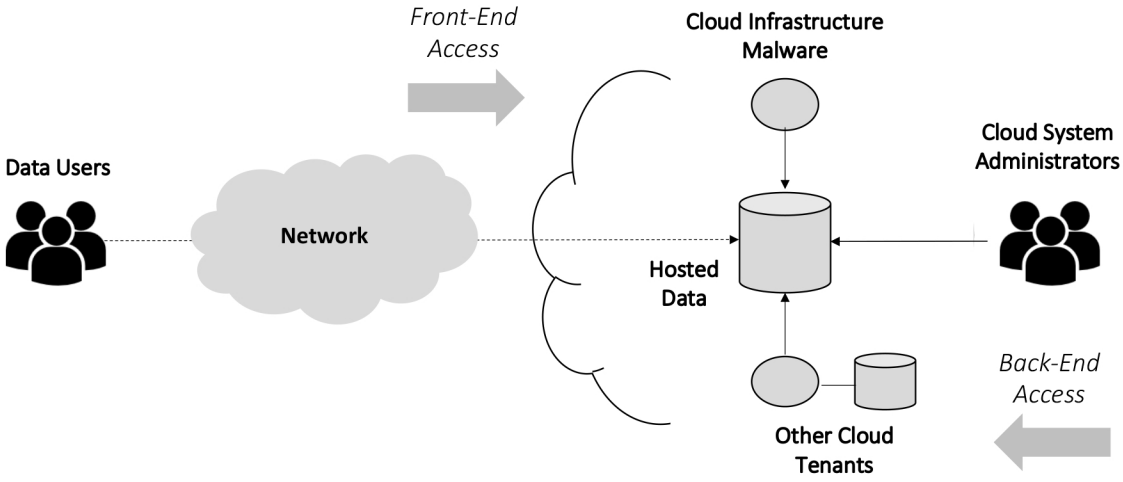

Figure 2. Back-End Security Threats for Cloud

To address the problem of rogue, disgruntled, or compromised cloud administrators, several options emerge. One involves strict contract management with public cloud services, with clearly defined penalties if evidence of insider attacks emerges. Another control involves the use of logs to uncover evidence of insider threats. Both of these approaches, however, might be covered up or suppressed by a competent malicious insider with high privilege.

The more common problem, however, involves non-malicious staff erroneously accessing, handling, or using cloud-resident resources. Errors during normal administration, maintenance update, patching, rehosting, and other tasks can easily result in a serious data compromise. The security controls to address such risk must be designed to deal with these common, but inadvertent problems.

What is needed, however, is a stricter control, one that cannot be subverted by a rogue or sloppy insider. This is where the emerging protection technique known as *data sharding* (or just sharding) has become increasingly important. The algorithms associated with sharding are relatively mature, but only recently have they been applied to modern protection of cloud-hosted resources.

Data Sharding to Reduce Back-End Cloud Risk

The general method of data sharding can be easily described schematically in terms of three algorithmic components: First, sharding involves disaggregating a given file or other asset into a set of smaller pieces. Second, the method involves separating these assets in manner that makes it difficult for any single hosting source to put the pieces together. Finally, the method involves obfuscating the disaggregated, separated asset pieces.

This three-step process offers an excellent means for securely storing sensitive data into the cloud. Furthermore, it supports the hosting of such data across disparate multi-cloud environments, thus dramatically reducing any back-end data access risk from unauthorized or unintended cloud actions. Tools are required obviously to retrieve and re-assemble the stored data for use by authorized individuals or data owners.

This last point is important to emphasize: Data sharding is not designed for storage in a passive manner, but rather for active hosting with query, access, and normal usage by users through cloud interfaces and processes through APIs. The algorithms must be designed with a retrieval capability that can combine and aggregate the data that has gone through the data sharding process.

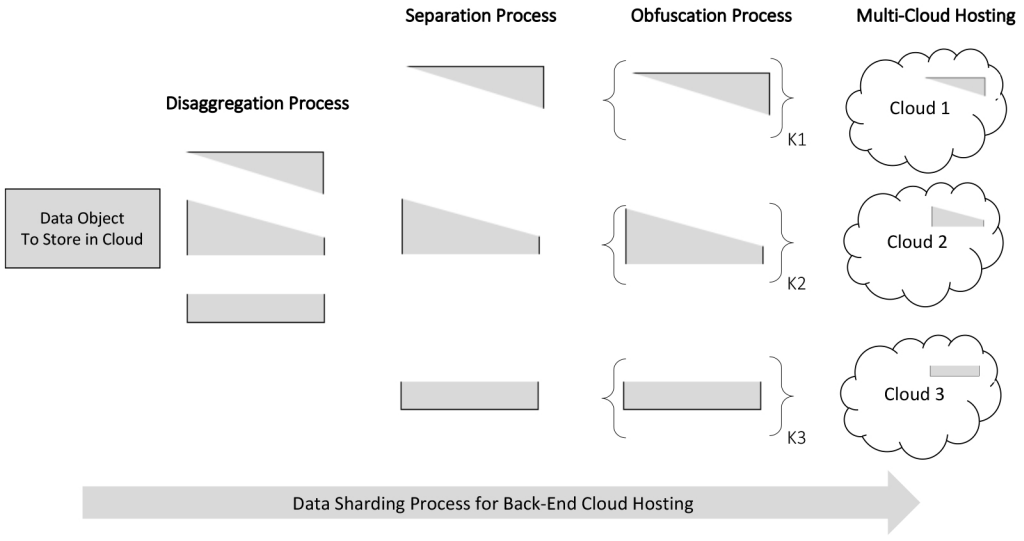

Figure 3. General Schematic Description of Data Sharding

The cyber security advantages of data sharding in the context of back-end cloud access include the following:

- *Supports Multi-Cloud Security* – When data is sharded across multiple target hosting locations, a powerful level of multi-cloud security can be achieved. This is helpful in many compliance settings, especially where the insider threat for externally hosted data is viewed as particularly intense. One might imagine financial services companies finding multi-cloud security controls important.

- *Complements Encryption* – Encryption of data before hosting into cloud is a common application-level control for hosted resources. This method is compatible with data sharding and can be used to provide layered security. The obfuscation involved in sharding can be over-laid onto any existing encryption without causing any functional issues.

- *Consistent with Multi-Person Controls* – The use of multi-person controls could be used for back-end access for hosted data sharding. That is, if administrators needed access to sharded data, then this could be accomplished through the design and development of tools that would aggregate and re-assemble the multi-cloud hosted data, but that could easily be only allowed with multiple administrators.

The motivation behind data sharding to reduce back-end cloud access threat is that hosted data is only useful if it can be retrieved and interpreted. If individual shards are neither complete nor meaningful when stored, then the threat is reduced. An additional consideration is that some vendors offer granular algorithms to manage the size and scope of the sharding – including, for example, micro-sharding options. Buyers should review options with their selected vendor.

Enterprise Recommendations

Any enterprise security team with responsibility to reduce risk in cloud-hosted infrastructure would be wise to begin examining commercial solutions in this emerging area. The TAG Cyber team provides regular guidance on excellent options for data sharding vendors, but since this is a new area, vendor mixes change frequently. Some of the commercial data sharding solutions examined for this work include the following:

Altibase (provides sharding for client and server-side database operations), *MongoDB* (includes sharding in its database solutions), and *ShardSecure* (start-up that provides a general data sharding solution for multi-cloud).

About TAG Cyber

Founded in 2016 by Dr. Edward Amoroso, retired Chief Security Officer for AT&T, Manhattan-based *TAG Cyber* democratizes cyber security industry research and advisory through high-quality reports, guidance, and information accessible to a wide expert audience. TAG Cyber helps close the communications gap between enterprise practitioners and security vendors. TAG Cyber also offers consultation and research subscriptions for enterprise practitioners.