AKAMAI SOLUTION BRIEF

Confidential Computing: Protecting Data in Use

As threats continue to grow in scope, scale, and sophistication, security teams are generally managing to rise to the challenge — particularly when encrypting data as it is being moved and limiting access when it is stored. But it is becoming increasingly apparent that teams need to also protect data while it is actively being edited, read, or processed, commonly referred to as *data in use*.

This gap in protection for data in use is becoming more important amid the evolution of computing and the rise of AI. The prevalence of hybrid and multicloud computing has expanded the ways in which organizations collect and store data. Meanwhile, as organizations seek to take advantage of AI, they are putting huge datasets — often their most valuable and sensitive data — in use where it is unencrypted and unprotected.

These risks are fueling the interest in confidential computing, an approach to security that ensures any sensitive data being used by applications, processes, or services remains encrypted and protected.

APIs add complexity

APIs are proliferating because they serve critical functions in two areas that companies are consistently putting resources into: cloud environments and services, and AI models. In the cloud, APIs are essential for allowing technologies to communicate and share data. With AI, large language models (LLMs) use APIs for accessing and combining data to carry out intricate tasks like language comprehension and text generation.

Unfortunately, APIs do not receive the same attention from security teams as applications and infrastructure do. Attackers are taking advantage of this security gap, with 84% of organizations having experienced API security incidents in the past 12 months.[1] To protect the sensitive data that every one of their cloud- and AI-related APIs touches, companies need comprehensive API security capabilities running in their confidential computing environments.

Locking all three doors

Locking up your data in transit and storage can still leave one door — data in use — wide open, exposing companies to risk.

In confidential computing, that data is processed in an environment that is deemed trustworthy at the hardware level. With APIs, organizations can deploy their own private machine learning instances that are purpose-built for securing API traffic rather than utilizing a public cloud API service, drastically reducing their attack surface. Running an API security solution in a confidential computing environment creates an extra layer of security. Even if part of the system is compromised, the data within the protected environment remains secure. Running API analysis on this data in a trusted environment is more secure and eliminates the risk present in traditional environments.

This combination of AI, API security, and confidential computing helps prevent unauthorized entities — like the hypervisor, host operator system infrastructure owner, or anyone with physical access — from viewing or changing code or data during execution, protecting

Business benefits

Enhanced data security
Limit access to data in use with strong controls, reducing the attack surface and protecting sensitive API-driven processes from unauthorized access

Protection for APIs
Execute deep API traffic analysis while keeping sensitive data encrypted in use, reducing the risk of exposure during monitoring

Stronger compliance
Meet evolving and stringent global data protection regulations, ensuring compliance with industry and government standards

1. Akamai, API Security Impact Study, 2024

against both internal threats (e.g., rogue system administrators or workloads running on untrusted infrastructure) and external threats (e.g., attackers who seize upon vulnerabilities).

The benefits

As API threats proliferate, and with data in use presenting itself as an attractive target, attackers won't be far behind. Forward-thinking organizations are beginning to embrace confidential computing for a number of reasons:

- Limiting access to data in use in the first place, through strong controls
- Securely analyzing the growing number of APIs
- Meeting new and stringent data protection compliance requirements around the world with the controls that confidential computing makes available

Confidential computing has the greatest benefit for highly regulated businesses, whether it's a financial services company looking to protect online transactions, or a life sciences company protecting patient data. This approach can also help an independent software vendor safeguard an AI model it distributes to customers across multiple locations from the edge to the cloud. Indeed, any enterprise IT organization running real-time analytical processing on its vital data needs to be thinking about confidential computing.

How we and our partners can help

Effective confidential computing requires an integrated set of solutions working closely together to provide comprehensive control and protection. Akamai, together with our partners Intel and IBM, delivers security for data in use from the hardware level to the cloud to APIs.

First, Intel® Trust Domain Extensions (TDX) provide trusted execution environments that:

- Protect against outside intrusion originating from threat actors and/or non-malicious entities that should not have access
- Improve security for the software that controls the technology used for creating virtual resources in the cloud, such as networks, servers, and storage
- Add a much-needed layer of security around the people administering any of these distributed systems — reducing the risk of honest mistakes and potential instances of insider malicious activity

Additionally, the Intel Tiber™ Trust Authority verification and tokens let organizations limit and control access to unencrypted data in use.

The Akamai API Security solution provides an inventory of APIs used in the enterprise, then monitors and detects how those APIs are used. It automatically detects and prevents malicious API requests by analyzing traffic patterns and behaviors, effectively blocking threats at the network edge without manual intervention. This allows for real-time protection against API attacks like data breaches, unauthorized access, and logic abuse.

Together, Akamai's remote machine learning engines combined with Intel Xeon® processors on IBM Cloud Virtual Servers — which are in turn secured with Intel TDX and attested with the Intel Tiber Trust Authority — deliver a private, hyperscalable environment designed to prevent any outside threat from accessing data when unencrypted in the final stage of data in use, whether it's from AI-powered bot attacks or human attackers.

The time to protect your in-use data has arrived

Organizations need a highly trusted environment to secure their most valuable corporate data — not only when it is being stored or accessed, but when it is actually in use. They are turning to Akamai and our partners to provide comprehensive security. Together, these trusted names in computing are ensuring security as it passes through every layer of the data lifecycle.