Having Zero Trust in the cloud

The Zero Trust framework provides high-level principles to be applied to an enterprise architecture in order to protect its workloads and users. Leveraging cloud computing resources for public Cloud Service Providers (CSPs) is increasingly becoming a standard within an organization's IT architecture. With the advent of cloud computing and the current adopting pace of public cloud services within organizations, traditional security practices need to be tailored to this new way of developing and deploying business applications. This article sets out to provide cloud professionals and organizations, who are growing their (public) cloud footprint, with tangible approaches to boost their organizations security practice through Cloud Native solutions.

Christian Siegers
is a senior manager at KPMG.

Lucas van Biert
is a manager at KPMG.

INTRODUCTION

Zero trust is a security concept that assumes that any user, device, or system attempting to access a network or (cloud) resources should not be trusted by default and must be verified before being granted access ([Rose20]). While trust in everyday life has been studied extensively for centuries, the formalizing and application of what trust and zero trust entails from a computation perspective came from S.P. Marsh written in 1994. In the early 2000s, security forums and researchers understood the challenges and limitations of perimeter-only defense and generally referred to multi-layer perimeter defense to overcome these challenges. Around 2009, the true concept of zero trust architecture, advocating for the need for a stricter and, seemly contradictory, more open security and access approach within an organization's network, began to take shape; especially through research conducted by Forrester analysts. In 2014, Google implemented Zero Trust architecture as part of its BeyondCorp platform ([Goog]); enabling employees to access applications hosted in the cloud or on-premises from anywhere without leveraging traditional VPN services. Around the same timeframe, Microsoft increased focus on the "Assume Breach" element of the Zero Trust framework.

The Zero Trust framework is especially suitable in the context of (public) cloud computing and the continuous pivot from privately owned datacenters and infrastructure to rapidly deployable cloud services and abstracting application code away from underlying infrastructure. While traditional, (typically on-premises) security measures certainly apply to cloud workloads, a slight shift in the paradigm is required. In order to maximize the business value of cloud workloads and keep time-to-production low, application teams prefer to leverage the latest services that Cloud Service Providers (CSPs) have to offer. Depending on the maturity of an organizations' cloud practice and the high rate at which (new) cloud services are developed, misconfiguration of cloud service and level of cloud operator access to tenant data remains a thread. Principles introduced by the Zero Trust framework, like isolation and inherent distrust between workloads, can help to maintain a security baseline without significantly impacting innovation and further adoption of (new) cloud services.

> Any system or network, no matter how secure it is, can be breached

Where securing the cloud platform is one of the key aspects, it is also vital that end users in a company have a good understanding of security related risks. Security awareness, through initiatives like foxhunts and gamedays, are an important component of a Zero Trust strategy for cloud environments. These types of initiatives help to educate employees about the risks and threats they may encounter in the cloud, as well as the actions they can take to protect the organization's assets.

The remainder of the article will dive into several aspects of the Zero Trust that are often adopted and how these could be applied in the context of (public) cloud.

CLOUD SECURITY FIRST

A Cloud Security First (CSF) strategy is a holistic paradigm to securing cloud environments, which prioritizes security measures above all else. The goal of a CSF strategy is to establish a comprehensive and holistic security mentality that covers all aspects of the cloud environment and integrated into all software development practice. Key components of the CSF strategy are understanding the shared security responsibility model with CSPs, assessing your current security posture, understanding your cloud architecture, implementing security controls, encouraging a culture of security first. Especially the latter cannot be overlooked in the context of a fast-growing set of managed services from CSPs and features provided by new Infrastructure as Code (IaC) frameworks.

Where CSF focusses on adopting and cultivating a mindset of security first, Zero Trust provides a more practical application of enforcing security through trust mechanisms. CSF touches every aspect of security awareness in your organization; everything from detection and response to security awareness in business operations.

Security awareness tends to impact almost every aspect of organizations today: organizational structure, roles, resources, policies, strategy and operations. The Zero Trust framework provides a more tangible approach to tackling security in modern (cloud) architectures.

Therefore, in the remainder of this article it will be used as a guide to highlight several different security practices, which are essential for deploying Zero Trust in your cloud environment. The most important aspect is operating from the assumption that your cloud environment is already breached.

ASSUME YOU ARE BREACHED

One of the foundational principles of Zero Trust is to assume that any system or network, no matter how secure it is, can be breached. It is based on the idea that organizations should not trust any user, device, or system on the network until they have been properly authenticated and authorized. In such an environment, all network traffic is treated as untrusted and is subject to strict security controls, such as multi-factor authentication, network segmentation, and micro-segmentation. This helps to prevent attackers from moving laterally across the network once they have breached a single system.

The model also emphasizes the need for continuous monitoring and real-time threat detection to quickly identify and respond to security incidents. Security Automation is the essential keyword in this area. While traditional Security Incident and Event Management (SIEM) solutions are well versed in collecting and consolidating information for security specialist, traditional human incident detection and resolution is not scalable in an ever growing and complex enterprise cloud ecosystem. Therefore, to be prepared for breaches and respond within seconds, the development of security automation solutions should be a priority, either commercially off the shelve or self-developed based on cloud-native services. To allow automation of security responses, many cloud security (native) services already provide a level of integration with cloud application - and data services. It is therefore important that during design time engineers and architects (from different disciplines) collaborate to consider safeguards and security controls upfront – i.e. shift the level of security concerns and reuse central security solutions. A starting point for this could be to define your infrastructure security posture in the context of an "outside-in" and "inside-out" defense approach.

OUTSIDE-IN VS INSIDE-OUT

An *outside-in* approach, also known as perimeter-based security, focuses on securing the boundaries of the cloud environment. This includes implementing security controls, such as firewalls and intrusion detection/prevention systems (IDS/IPS), which will be discussed further on, to protect against external threats. The goal of an outside-in approach is to prevent unauthorized access to the cloud environment.

An *inside-out* approach, on the other hand, focuses on securing the resources within the cloud environment. This includes implementing security controls such as access controls (IAM), data control access and data encryption to protect against internal threats. The goal of an inside-out approach is to prevent data breaches (e.g. exfiltration) and unauthorized access to cloud resources.

A Zero Trust strategy typically employs both outside-in and inside-out approaches, as they complement each other. Namely, implementing security controls, such as firewalls and IDS/IPS, ensures perimeter protection of the cloud environment, while implementing access controls and data encryption security controls aim to protect data and resources within the cloud environment.

Before looking at the IDS/IPS aspect, a suitable next step is to understand data flows and network connectivity requirements for your cloud environment, which are commonly categorized into two types of traffic: ingress and egress.

THE FLOW OF NETWORK TRAFFIC

Ingress and egress refer to the flow of (network) traffic into and out of a network or system. The Zero Trust model assumes that any system or network, no matter how secure it is, can be breached. As such, it emphasizes the need for strict security controls to protect against unauthorized access, regardless of whether the traffic is incoming or outgoing.

For ingress traffic, a Zero Trust model would involve verifying the identity of the source of network traffic and ensuring that it is authorized to access the network or system. This can be achieved by implementing multi-factor authentication (MFA) and using security protocols such as Transport Layer Security (TLS) (or also known from its predecessor Secure Sockets Layer (SSL)) to encrypt the traffic.

Additionally, traffic coming in could be inspected and analyzed by security devices such as firewalls, Intrusion Prevention Systems (IPS) or Next-generation firewalls (NGFWs), which would be able to identify and block malicious traffic. Also, the use of a demilitarized zone (DMZ) could strengthen the security posture for incoming traffic. DMZs will be further discussed in the next paragraph.

For egress traffic, a Zero Trust model would involve ensuring that traffic is being sent to authorized destinations, and that the data is protected from exfiltration. This can be achieved by implementing encryption, monitoring for data leakage, and using security protocols, such as TLS or SSL, to encrypt the traffic.

Network segmentation is a powerful tool for ensuring that both ingress and egress traffic is properly controlled in a Zero Trust environment. This involves dividing the network into smaller, more manageable segments and implementing security controls on each segment. This can help to prevent attackers from moving laterally across your cloud network once attackers have identified a vulnerable entry-point into your cloud environment.

People are still fallible, and that fallibility can cause data breaches

DMZ, AN EXTRA LAYER OF SECURITY?

A demilitarized zone (DMZ) is a security architecture commonly used to provide an additional layer of protection in cloud environments. This pattern originates from on-premises solutions and is still commonly used to add an extra layer of security to cloud environments. However, from a pure Zero Trust perspective, a DMZ is not required, as Zero Trust does not solely rely on network-level security measures and perimeter defense.

A DMZ is typically implemented as a separate network segment that sits between an organization's internal network and the Internet. The DMZ serves as a buffer zone, where incoming traffic is screened and filtered before it is allowed to access the internal network.

In a cloud environment, a DMZ can be implemented to protect cloud-based resources such as servers, databases and applications. For example, public-facing (web) servers can be placed in the DMZ, while sensitive data is stored on servers located in the internal network.

The DMZ can also include security controls such as firewalls, intrusion detection/prevention systems (IDS/IPS) and load balancers to provide protection against various types of cyber-attacks.

By implementing a DMZ in a cloud environment, organizations can:

- reduce the attack surface of their cloud infrastructure by limiting access to only necessary resources;
- improve security by placing security devices in the DMZ;
- segment the internal and external network;
- isolate public-facing resources to prevent unauthorized access of sensitive data;
- enable better incident response and forensic analysis.

Implementing a DMZ in a cloud environment can help organizations to better protect their cloud resources and ensure that only authorized access is allowed. CSPs often provide several cloud-native services to help you provision security solutions in your DMZ.

IDS/IPS: ESSENTIAL SECURITY CONTROLS

An Intrusion Detection System (IDS) and an Intrusion Prevention System (IPS) are key components of a Zero Trust strategy for cloud environments. While newer concepts, such as Endpoint Detection and Response (EDR), Extended Detection and Response (XDR) and Network Detection and Response (NDR) are being introduced by 3rd party solutions, cloud-native services for IDS/IPS are still widely relied on.

An IDS is a security solution that monitors network traffic and identifies potential security threats, for example, unauthorized access attempts or malicious activity. An IDS can be configured to alert security teams of potential threats and may also be configured to take automated actions such as blocking the traffic.

An IPS, on the other hand, is a security solution that monitors network traffic and actively blocks or mitigates security threats. It can be considered as an advanced version of IDS. It examines network traffic in real-time and takes action to prevent malicious packets from reaching their intended targets.

When implemented in a cloud environment, an IDS/IPS can provide several key benefits, such as:

- monitoring and protecting cloud resources against various types of cyberattacks;
- real-time threat detection, alerting and blocking;
- analyzing the traffic and identifying malicious traffic;
- providing detailed information on security incidents to the Security Operations Centre (SOC) team;
- helping to prevent data breaches and unauthorized access to cloud resources.

These solutions are the vital part of a secure cloud environment. However, large segments of traditional security solutions rely on human operation and control. In some cases, extra layers of protection can be implemented as automation. The Assume Breach mindset already touched on the need to consider security controls from the outset and throughout the design and development of your cloud environment. In practice this means that automated security controls are essential for a secure cloud environment.

AUTOMATION PLAYS A CRITICAL ROLE

In a cloud-based environment, automation plays a vital role in deployment of solutions, auto-healing and other tasks where traditionally human intervention was required. People are still fallible, and that fallibility can cause data breaches. Misconfiguration is responsible for 10% of security breaches ([Veri22]).

From a Zero Trust perspective, automation will play an even more critical role in enforcing security policies and maintaining the integrity of a cloud network. Automation can be applied in virtually all security control scenarios in a cloud environment:

- Verify the identity of users and devices. Automated systems can verify the identity of users and devices through multi-factor authentication. Multi-factor authentication refers to aspects of a user to be tested to verify their identity. In a multi-factor this includes multiple combinations of something the user knows (e.g., a password), something the user has (e.g., a hardware token), and something the user is (e.g., a fingerprint).
- Monitor network activity. Automated systems can continuously monitor network activity for unusual or suspicious behavior. For example, machine learning algorithms can be trained to detect anomalies in network traffic, such as excessive data exfiltration or unauthorized access attempts.
- Enforce security policies. Automated systems can be used to enforce security policies, such as firewalls and access controls, in real-time. For example, an automated system could automatically block traffic from a known malicious IP address.
- Continuous discovery and monitoring on infrastructure, if using Infrastructure as a service, automation could work on discovering resources, and continuous monitoring to ensure only authorized access is granted and to detect any misconfigurations in the infrastructure.
- Scanning solutions minimizes your security risks and accelerates the remediation process by comparing cloud application configurations to compliance policies to identify gaps quickly.

CLOUD SCANNING SOLUTIONS

There are different types of security scanning solutions available, but when it comes to cloud native solution, the two main types of solutions focus on the Network (Layer 3) and Application (Layer 7) layers, as defined in the Open Systems Interconnection model. Layer 3 scans focus on the underlying infrastructure of a cloud environment, such as the network, servers, and other devices. These scans typically include checks for vulnerabilities in the operating system, software, and network configuration. They can also include checks for misconfigurations, such as open ports and weak passwords. The different cloud platforms have built in Level 3 security as (Web Application) Firewalls, DDOS and Vulnerability scanners.

Layer 7 scans, focus on the upper layers of the cloud environment, such as the web applications and APIs. These scans typically include checks for vulnerabilities in the web application code, such as SQL injection, cross-site scripting (XSS), and cross-site request forgery (CSRF). They can also include checks for misconfigurations and weak access controls.

Both layer 3 and layer 7 scans are important for identifying vulnerabilities and misconfigurations in a cloud environment, as they provide different perspectives and a comprehensive coverage.

When implemented as part of the cloud environments, these scans can help organizations to:

- identify vulnerabilities and misconfigurations in the cloud environment;
- prioritize vulnerabilities and misconfigurations based on their level of risk;
- provide detailed information on vulnerabilities and misconfigurations that need to be fixed;
- help prevent data breaches and unauthorized access to cloud resources;
- enable incident response and forensic analysis.

These scanning solutions aim to detect common vulnerability and known breaching tactics. One of the most used awareness lists is OWASP ([OWAS22]). On a yearly basis, OWASP creates a list of the top 10 biggest security risks. OWASP's (annual) publications could be a useful trigger for organizations to actively test their cloud environments security. One common approach for this is called Red Teaming.

RED TEAMING

Red Teaming is a security testing method that simulates real-world attacks to assess an organization's security posture and identify vulnerabilities. In the context of a Zero Trust strategy for cloud environments, Red Teaming can be used to test the effectiveness of security controls and identify any gaps in protection.

The red team simulates a variety of attacks, such as phishing, social engineering, and advanced persistent threats, to test an organization's ability to detect and respond to cyber-attacks. The red team also evaluates the security of cloud infrastructure, including network segmentation, access controls, and incident response processes.

The goal of Red Teaming is to identify and exploit vulnerabilities before they can be used by malicious actors. By simulating real-world attacks, the red team can provide valuable insight into an organization's security posture and help to identify areas that need improvement.

When implemented as part of a Zero Trust strategy for cloud environments, Red Teaming can help organizations to:

- identify and prioritize vulnerabilities;
- improve incident response capabilities;
- validate the effectiveness of security controls;
- enhance security awareness and employee training;
- provide a comprehensive security testing and validation.

By conducting regular Red Teaming exercises, organizations can ensure that their cloud environments are as secure as possible, and that they are prepared to detect and respond to cyber-attacks.

While it is key (and fun!) to regularly apply Red Teaming and penetration test your cloud environment, security strategy should never overlook human system interactions. Human actors are part and parcel of every application, workload and cloud resource. Therefore, continuous focus is required on security awareness to maintain knowledge and create "muscle memory" to recognized and respond to possible security threats. The education of stakeholders in your organization can be done in many ways. Two common and proven approaches are covered in the next section.

FOXHUNTS AND GAMEDAYS TO INCREASE SECURITY AWARENESS

Security awareness training is an important component, as it helps to educate employees about the risks and threats they may encounter in the cloud, as well as the actions they can take to protect the organization's assets. 82% of breaches involved a human element, including social engineering, plain human errors and misuse ([Veri22]).

One way to increase security awareness is to use gamified training methods such as foxhunt and gamedays. A foxhunt is a simulated security incident where a team of security experts, known as the red team, simulates an attack on the organization's network. The goal of the foxhunt is to identify and exploit vulnerabilities before they can be used by malicious actors. Employees are trained to detect and respond to a simulated attack, which helps to improve their ability to detect and respond to real-world attacks.

Gamedays are similar to foxhunts, but instead of simulating a security incident, the red team simulates a business continuity event, such as a natural disaster or power outage. The goal of the gameday is to test the organization's ability to respond to and recover from an unexpected event.

Foxhunts and gamedays can be adapted to the cloud environment and can include simulations of different types of cloud-specific attacks, such as a cloud account compromise or a misconfiguration in a cloud infrastructure.

Security awareness training with foxhunts and gamedays can help organizations to:

- improve employee knowledge and understanding of cyber threats and risks;
- improve incident response capabilities;
- validate the effectiveness of security controls;
- enhance security awareness and employee training;
- provide a comprehensive security testing and validation.

By conducting regular foxhunt and gameday exercises, organizations can ensure that their employees are prepared to detect and respond to cyber-attacks, and that they are aware of the best practices to protect the organization's assets in the cloud.

CONCLUSION

In conclusion, a cloud-based Zero Trust security model is a proactive approach to securing access, data and network resources that can help organizations prevent data breaches and unauthorized access. It works by continuously verifying the identity and trustworthiness of users and devices, and by enforcing security policies through automation. Zero Trust in the cloud, provides a framework to protect against threats, both inside and outside the organization, by implementing multiple layers of security and validation, and by continuously monitoring for suspicious activity. This helps organizations to reduce their attack surface, making it more difficult for malicious actors to gain access to sensitive information.

References

[Goog] Google (n.d.). *BeyondCorp*. Retrieved from: https://cloud.google.com/beyondcorp#:~:text=BeyondCorp%20Enterprise%20is%20a%20modern,a%20traditional%20remote%2Daccess%20VPN

[OWAS22] OWASP (2022). *OWASP Top Ten*. Retrieved from: https://owasp.org/www-project-top-ten/

[Rose20] Rose, W, Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture*. [NIST Special Publication 800-207]. Retrieved from: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[Veri22] Verizon (2022). *Data Breach Investigations Report*. Retrieved from: https://www.verizon.com/business/resources/reports/2022/dbir/2022-data-breach-investigations-report-dbir.pdf

About the authors

Christian Siegers is a senior manager at KPMG NL. He is a highly skilled solution architect with years of experience in digital transformation and technology. He helps clients in defining solution architectures which are future proof. With a broad experience in cloud, solution and integration architecture, he looks into the (emerging) technology and how this can enable and solve business challenges and/or problems and create solutions and innovations.

Lucas van Biert is a manager at KPMG NL and has experience as cloud (solution) architect. He has supported several organizations in their cloud journey, predominantly in the financial services industry. Using his (data and cloud) engineering background, he provides valuable technical insight and architectural recommendation on cloud-based solutions, enabling clients to design, build and deploy business-oriented solutions, leveraging cloud-native services and proven architectural patterns.