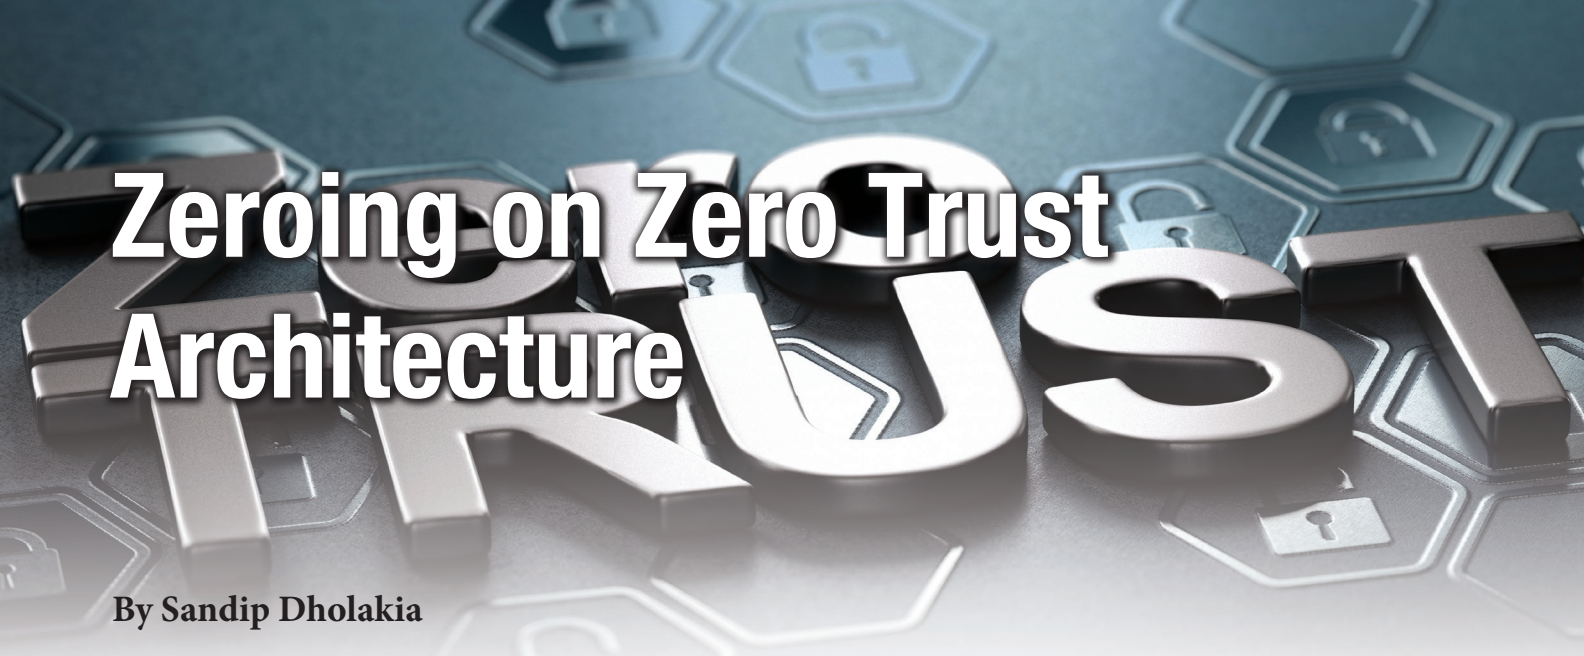


Zeroing on Zero Trust Architecture

By Sandip Dholakia

This article takes a look at Zero Trust Architecture and how it could be the way of the future for remote workers.

In the information security industry, people use some buzzwords to impress others. For the most part, using these buzzwords works because very few people in the room know enough about these buzzwords to ask you any challenging questions! Lately, Zero Trust has been one of those surefire buzzwords in the industry. In this article, I will explain why what and how of zero trust architecture so that you can intelligently impress the audience.

Why do we need zero trust?

Before the inception of zero trust, Virtual Private Network (VPN) was the only secure way to access the remote network. The monopoly of VPN was hard to beat because of its ability to provide a dedicated connection and encrypted communication between a remote user and the corporate network. The VPN makes eavesdropping and unauthorized access very difficult. When VPN technology was first designed more than 25 years ago, the goal was to allow a handful of users connect to a single data center. However, the IT industry, with the popularity of the cloud, has come a long way in recent years. In today's IT landscape, most employees could be working off-site and could be connecting to multiple data centers, public cloud providers, or several SaaS applications. VPN was not designed for this type of use case. The exponential growth of cloud-based applications in the last few years has exposed limitations of VPN – ease and speed of access. Every device that connects to a remote network must have a specific VPN client. VPN clients can be process intensive and lead to slower communication. VPN server and client can be a single point of failure for users. The goal of the technology should be to provide easy and quick access to remote networks. This is where zero trust comes to our rescue!

What is zero trust?

As employees work remotely and teams collaborate across the globe, it is important to offer easy, fast, and secure access to enterprise applications without any additional overhead to end users or their devices. Zero Trust Architecture (ZTA) allows remote users to sign-in to the corporate network with their corporate single sign-on (SSO) credentials without establishing a VPN connection. Zero trust architecture offers accessibility and improved data sharing across corporate networks and partners with the utmost security. As shown in Figure 1, users can be working from home, from work, from a coffee shop, or from the airport and may want to access an app on the cloud or the company's data center using any device, zero trust can make it happen seamlessly! (Note: The figure is oversimplified for the illustration purpose.)

Zero trust connects the user to an application and not to the network! This translates into zero lateral movements.

Zero in zero trust implies the philosophy of "trust no one," inside or outside the network. There is no perimeter security in zero trust architecture. The foundation of zero trust architecture is based on the four components - strong authentication, segmented network, monitored devices/applications and granular authorization. The policy engine at the heart of the architecture controls the overall execution.

Zero trust in its present form was introduced in 2010 by Forrester Research.1 Later, Google announced that it had been using this philosophy internally. In 2014, Google introduced BeyondCorp,2 a model for zero trust implementation. Since then, zero trust has been implemented in other organizations' enterprise models.

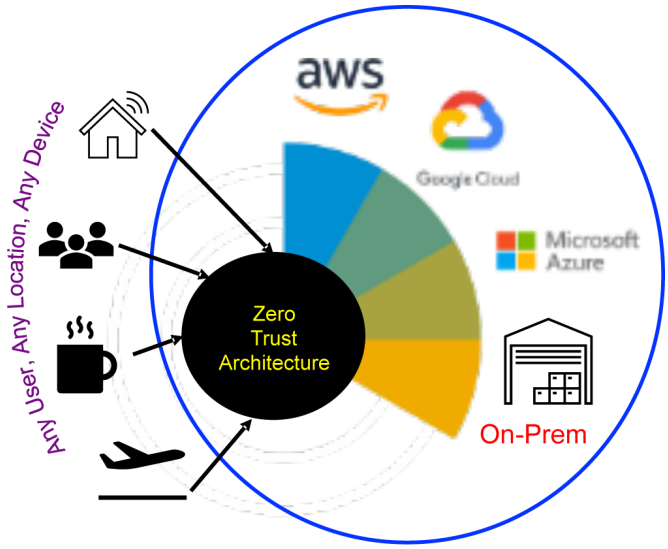


Figure 1. Any user, any location, any device connecting to any target using ZTA

Ask what zero trust can do for you!

The number one advantage of zero trust is the "zero attack surface." In ZTA implementation, no ports, firewalls, VPCs, or VPNs are exposed to the internet – you can't attack something if you can't find it! This not only increases the security, but it also saves money to your organization as there are no firewalls or VPNs to install and manage. Another factor to consider is the implementation workload. Implementation of zero trust does not require any additional work – assuming evaluating security posture and updating, if needed, is an ongoing process in your organization. Security professionals should always implement defense in depth as well as controls for stronger authentication, granular authorization, monitoring, etc., to improve the security posture. For example, SSO is a required component of zero trust implementation. However, if your organization is already using SSO, there is no additional work or expense for your organization to implement the SSO. The next big advantage is that the zero trust approach allows you to choose the level of security based on the data classification. For example, sensitive data, such as PII, can have more security controls, the database can be in an isolated Virtual Private Cloud (VPC) and access is allowed only to a handful of users. On the other hand, webservers with non-sensitive information can be in a VPC with other servers and should have fewer controls. Unlike VPN, every application doesn't need the highest level of security. This translates into cash savings for most organizations. Lastly, zero trust makes it easy to scale – users can use any device and can log in from anywhere. Users are not restricted to devices with VPN clients installed or to corporate-issued devices.

How zero trust works?

People are the weakest link in managing security. Most breaches happen because users simply fail to keep their credentials safe. Implementation of proper Identity Lifecycle Management (ILM) policy is very crucial. Users should be onboarded and offboarded right away – preferably these processes should be automated. Some guidance for strong authentication includes password management policy, multi-factor authentication and user behavior. User behaviors, such as recognition of the type of device being used or the user's location is required to log in. For example, if a user logs in from the United States every day, and one day the same user logs log in from another part of the world, then the login attempt should be put on hold, and more identity confirmation must be requested from the user before access is granted. User behavior is not yet commonly used for authentication but can add a layer of control in stronger authentication.

Unlike VPN, flexibility and ease of use is at the heart of the zero trust philosophy. This means a user can change a device at any time or even use a personal device to access the target application. This leads us to another component of the zero trust architecture - continuous monitoring of user devices and target applications. The security posture of user devices must be checked before allowing access. Parameters, such as the latest versions of OS, updated patches, native encryption mechanisms and recent virus scan signatures should be checked. When users are connecting with a corporate issued device these checks and verifications are easy since the devices are managed by the corporation. However, the issue becomes a bit tricky and complicated to manage when a user is connecting with a personal device. Device posture must meet the minimum requirements before allowing access. Targets or applications being connected must be hardened along with the infrastructure, OS, and middleware it is running on. Application, middleware, OS, and infrastructure must go through the proper DevSecOps processes, various vulnerability assessments, security testing, and other standard application hardening processes. Necessary logs from devices and applications must be aggregated and analyzed using automated System Information and Event Management (SIEM) tools.

Networks form the third component of the zero trust architecture. Networks, broadly speaking, can be divided into two types: the corporate network, where applications, data, and infrastructure are hosted, and the network used to connect to the corporate network, such as a public Internet connection. Needless to say, all communication between the endpoint and corporate network must be encrypted using a secure protocol like TLS/HTTPS.

Zeroing on Zero Trust Architecture

continued from page 13

In zero trust architecture, there is no single network perimeter. All applications should be isolated and hosted on small individual VPCs. Databases must be protected further by hosting them in the private subnets. Only application servers and jump hosts for administration (using IP address–allowed listings) should be able to connect to the database in a private subnet. This process of creating small trust zones for host applications is known as microsegmentation. Microsegmentation offers two advantages. First, it allows just the right amount of security based on the value of the asset in the network. Second, if a hacker can gain access, the damage is minimized; the impact will be limited only to that one application, as the hacker will not be able to reach other applications –no lateral movement. Modern technologies, such as Containerization and Microservices make microsegmentation and related security very easy using VPCs, software-defined networks, firewalls, security groups, and the like.

One of the disadvantages of VPN technology is the same level of access for everyone. Once the hacker breaks in, he has the key to the kingdom. Granular authorization in zero trust addresses this issue. Users should have access to data and applications that they are expected to use. Need to know is one of the concepts used for the granular authorization. Not every user needs to know or access everything that the enterprise has to offer. Access to data and applications should be granted based on business needs. This will minimize the damage if a user's credentials are compromised. For example, a software developer should not be able to read payroll data. The concept of Least privilege is another concept used in granular authorization. Users should be granted the access required to perform their work. This means that users should have the minimum (or least) amount of access. For example, an auditor should have "read-only" access to data and should not be able to modify the data, because this is not necessary to perform their work. Implementing SSO and reviewing user privileges and permissions regularly are key measures any enterprise can take for secure access management.

Figure 2 shows how each of these components works together. As shown in the figure, the policy engine is at the heart of the architecture. The zero trust policy determines who can transit the microperimeter at any point in time, preventing unauthorized users from gaining access to your protected surface and preventing the exfiltration of sensitive data. The engine evaluates all the parameters and decides if the connection should be established and access should be granted. The policy engine is also responsible for closing the connection when the task is complete, when access is not granted, or authorization fails.

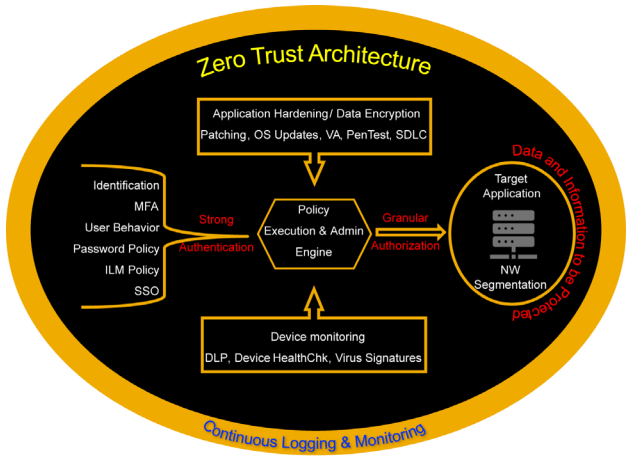


Figure 2. Zero Trust Architecture

The primary goal of information security is to protect data and information, and zero trust is designed to just do that. All components of the zero trust architecture discussed above help protect data. However, to make the architecture completely robust, additional protection to the data itself should also be considered. Depending on its classification and sensitivity, data must be encrypted in transit as well as at rest. Wherever possible, data at rest must be encrypted using data volume encryption and disk encryption. Disk encryption is useful in protecting the data in case the disk, where the data is stored, is lost. When data is stored on the cloud, customers should try and manage their own encryption key for data-at-rest encryption. If that is not an option, the key management process must be understood and included in the contract.

Rome wasn't built in a day, but they were laying bricks every hour!

Implementing zero trust is a journey. It will take some time for any organization to adapt to zero trust. The organization must devise a detailed plan and work towards it one step at a time – just like Romans! The first step of the plan should be to provide security training to all employees. One of the major challenges in the implementation of zero trust architecture is a user mindset. Zero trust architecture provides a robust defense against stolen credentials, but attackers still could take advantage of a user's temptation to respond to phishing attacks, social engineering, or a combination of different strategies.

It is impossible to go "Cold-Turkey" with VPN. Organizations will use both methods, VPN and zero trust, in parallel for a while. A migration plan could include grabbing the low hanging fruits first to show quick progress but migrating everything to zero trust in a short span can add a significant cost to the company. Even while using VPN, the organization should continue to improve security posture, processes and policies. Improving security should be an ongoing process.

Logging & Monitoring and Identity & Access Management (IAM) are the two low hanging fruits, and most organizations can start their zero trust journey in these two areas. Logging & Monitoring is a good start as it is easy to turn on logging with a few simple configuration changes. Analyzing the logs using a SIEM provides good visibility into the system usage. Security professionals should use these data to create a baseline for the organization. This is the first and very important step of the zero trust journey. IAM is a core component of the security and controls, such as SSO, MFA, granular authorization, etc. must be implemented regardless of zero trust adaptation.

If zero trust is so obvious and secure, why everyone is not implementing it? I hear you asked.

Networking is one of the difficult areas when it comes to security. Microsegmentation usually takes a lot of planning and work. In on prem deployments, apps are usually not isolated very well since everything is behind the perimeter, while in the case of cloud deployment, it could be worse with ease of spinning the new VMs in seconds (virtual sprawl). To isolate every app in the organization and create microsegments can take a while for any organization. Networking coupled with a mindset of difficulty in centrally managing decentralized architecture demotivates many organizations from the early adaptation of zero trust architecture.

One option is to purchase a part of the solution or outsource the implementation. However, proper research is warranted. Many vendors in the market want to capitalize on the buzzword of zero trust. Most vendors provide only part of the solution. For example, some vendors have very robust identity management solutions, some vendors offer advanced DLP solutions while a small segment is specialized in a device health check solution. Almost all solutions that I came across seem to be working seamlessly with solutions offered by other vendors. This makes it easy for an organization to use multiple vendors, if needed, to implement zero trust architecture. Finally, the "trust no one" principle should be applied to your vendor selection process as well – every vendor should be carefully evaluated for their marketing claims!

The journey of zero trust started almost a decade ago and will continue for many years. The process and technology will only get better over time - enjoy the ride!

To learn more about Zero trust architecture, please check out the references provided below.

References

1. J. Kindervag, Forrester Research, "Build Security Into Your Network's DNA: The Zero Trust Network Architecture," 2012.
2. R. Ward, Google Research, "Beyond Corp: A New Approach to Enterprise Security," 2014.
3. Whitepaper, Identity Defined Security Alliance: The Path to Zero Trust Starts with Identity.
4. Scott, Rose et. Al., "NIST Special Publication, 800-207, Zero Trust Architecture", 2020
5. Cybersecurity and Infrastructure Security Agency, "Zero Trust Maturity Model", 2021

About the Author

Sandip is a passionate Information Security professional with experience in various security domains and industries. Sandip is working as a Principal Security Architect, Global Security Group, SAP America.

Rethink How You Hire

continued from page 9

Rebuilding company culture to highlight personal identity means adopting flexible environments and promoting a sense of inclusion. This will help employees contribute and collaborate, even advance. It's time to take the next step to rethink how and who you hire as well as building a flexible, transformational culture to retain talent. Leaders who account for the lack of daily in-person interaction and focus on how the organization can fulfill its goals and support employees wherever they may be working are best positioned to remain leaders.

Coming Next Month: *Alternatives to Traditional Education: Part 3 of Inspiring and Preparing the Next Generation of Cybersecurity Professionals*

About the Author

Dr. Curtis C Campbell, C|CISO, VP of Atlantic Capital Bank in Atlanta, GA, serves as Director on the ISSA International Board, and is President of the ISSA Chattanooga Chapter. Curtis holds a Ph.D. in Organizational Leadership in Information Systems Technology, and serves on the advisory board of University of TN-Chattanooga, a national Center for Academic Excellence for Cyberdefense (CAE-CD) studies. She was named ISSA Fellow in 2020. Connect with Curtis via curtis@mprotechnologies.com.

References

1. "A Resilient Cybersecurity Profession Charts the Path Forward", (ISC)2 Cybersecurity Workforce Study, 2021 [Online] Available: https://www.isc2.org/Research/Workforce-Study
2. A. Elder, 4 Questions Carla Harris Says Candidates Should Ask in an Interview –and How You Can Respond, Dec. 28, 2021, [Online], Available: https://www.linkedin.com/business/talent/blog/talent-acquisition/carla-harris-interview-questions
3. H. Tariq, "Four Tips for Retaining Your Best Employees", Nov. 13, 2021, [Online], Available: https://www.forbes.com/sites/forbescommunicationscouncil/2021/11/11/four-tips-for-retaining-your-best-employees/?sh=589e56a87f40