WHITE PAPER

4 Common Pitfalls of Zero Trust Solutions

The Zero Trust Security Solution

When designing a Zero trust architecture, it's not just about following best practice guidelines like: the NIST Cybersecurity Framework, the CISA Extensible Visibility Reference Framework, the CISA Zero Trust Maturity Model, or even US Executive Orders like EO 14028 – "Improving the Nation's Cybersecurity", from May 2021 and OMB memorandums like M-22-09 and M-21-31. There are several other factors that are important to keep mind.

For instance:

- Does your Zero Trust architecture actually cover all of your needs?
- Will it work as designed?
- How effective will it be?
- Are you missing anything?
- Does your security architecture restrict critical application or user performance?

These are the questions that keep government SecOps teams up at night because bad actors only have to be right one time for your security architecture to be labeled a failure. You and your government agency can't afford a mistake.

This whitepaper will provide several useful insights on how to overcome four common pitfalls for Zero Trust security architectures.

Four Common Pitfalls That You Can Avoid

When designing and implementing your Zero Trust security architecture, there are some common factors (pitfalls actually) that can get overlooked or dismissed.

Here are four of them:

1. Did you include a visibility architecture?
2. Do you have packet-level visibility?
3. How do you plan to validate your architecture and change management policies?
4. Did you include self-healing security opponents as well?

The following sections will discuss each pitfall in detail.

Pitfall #1 – Don’t Forget a Visibility Architecture

For any network security project, network visibility is the keystone. This is because the visibility architecture captures key pieces of data that help: secure the network, reduce operational costs, and improve performance. Without this information, IT security projects will struggle to achieve stated goals and cost targets.

Simply put, data is the lifeblood of modern networks. Terabytes, and possibly exabytes, of data traverse your network every day. This huge amount of data is forcing IT teams to acquire an even better insight and understanding of their network(s). Any lack of visibility will create blind spots and jeopardize network security. It’s unacceptable to not know what’s lurking within in your network — not when there’s technology that can expose those potential problems.

So, how do you expose these threats? The first place to start is by creating a visibility architecture that consists of taps (for data collection), a network packet broker (for data manipulation), and purpose-built security tools, like intrusion detection systems (IDS), to examine the data. By integrating visibility technology into your security architecture, you can clearly see what is happening on your network and implement the proper responses.

Taps are passive network data access points that let you copy data anywhere on your network. This gives you the opportunity to get the type of data you need. Taps make a complete copy of the data, not a summarized excerpt. This full copy of data allows you to get the whole story. You can see degradations as they start to occur — before an impairment happens. Pre-incident data can be extremely valuable in helping to diagnose the root cause of problems and save you a significant amount of troubleshooting time. You will also need a complete copy of data for threat hunting activities.

Taps are quick and easy to deploy. There are also several different versions. One version is for use with your physical on-premises environment. Another version, called cloud-based taps, are used to capture virtual data. There are also copper and optical versions available.

Network packet brokers are another core visibility component. Packet brokers allow you to aggregate data from taps anywhere across your network. You can also perform general Layer 2-4 packet data filtering to optimize data monitoring and management. Higher end models also allow you to filter data at the application layer (Layer 7). In addition, the packet broker can remove unnecessary duplicate data, remove unnecessary packet headers and content, and then replicate data to as many tools as needed.

The third component of a visibility architecture consists of purpose-built security tools, like an IDS, security information and event management (SIEMs), data loss prevention (DLP), or threat hunting tools. One huge benefit of the tap and packet broker combination is that once they are installed, you can essentially insert out-of-band security tools at will because there will be a limited need for Change Board Approvals (since these tools will not directly affect the flow of production data traffic). This is important if you want to insert new security tools for additional functionality or to replace older, lower speed tools.

For inline security tools and packet brokers, a Change Approval Board will typically still be required but the use of inline packet brokers and externa; bypass switches can significantly reduce risk and make the insertion of new security tools, or updates to existing tools, much easier than serial security tool deployments.

Here is one example of a visibility architecture for out-of-band security tools:

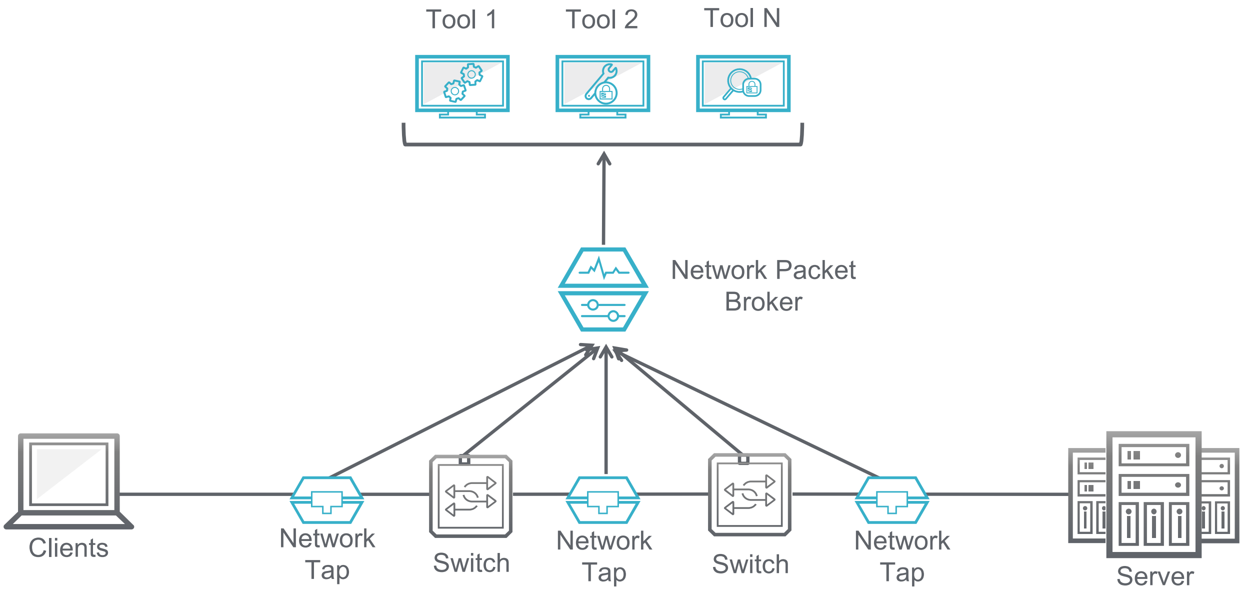

Figure 1: Out-of-band Visibility Architecture Example

Visibility architectures should also support public and private cloud environments. This is important as hybrid (cloud + physical on-premises) environments continue to grow. Virtual taps and packet brokers can be used to perform data capture and filtration in cloud environments and then backhaul critical forensic data to existing physical on-premises tools. This provides one central location for security data analysis and reduces the cost of deploying duplicate security tool functions in both on-premises and cloud locations.

Pitfall #2 – You Need Packet Visibility

Packet data will be critical to your security architecture as it can provide a single source of detailed truth. Don't misunderstand. While flow data is good, it only provides general trend information, not actionable details if you want to perform any type of threat hunting. Log data is also useful, but it can be corrupted or even erased by malware. Only packet data gives you all of the details that you need, like: who, what, where, when, and how. The devil truly is in the details. Metadata can never tell the whole story while packet data is the absolute truth.

One key point to keep in mind is that you need to be able to read the packet data. According to public research data, more than 70% of malware may now hide in encrypted packet data. Passive and active TLS decryption allows you to look into those packets and see what's hidden. After that, you can monitor and flag potentially malicious communications.

At the same time, decryption efforts should be strategic. Just "turning on" decryption on every inline security tool will cause you many headaches. For instance, decryption functions on security tools can be different, which can cause feature/function mishaps. In addition, most security tools suffer a significant performance degradation once decryption is enabled, i.e. the decryption feature consumes a lot of CPU resources which also results in an increase in data processing times. This in turn impacts your security tool planning, as you now have to buy more security tool processing capability since your security tools have less processing capability. ZK Research found in a study that 45% of security professionals had turned off decryption functions due to the processing overhead.

This is where the selection of a well-designed packet broker can help out. The decryption capability can be offloaded to the packet broker, which can be more cost-effective than turning on decryption for every security tool. In addition, it's easier for the SecOps team. The packet broker decrypts the data once, passes it to the appropriate security tool(s), and then re-encrypts the data before it is sent into the rest of your network. Other security tools and sensors may not even have the ability to decrypt traffic. So, the packet broker concept makes sense here as well.

When it comes to packet visibility, it is important to note that you need a complete copy of the packet data. This means that you need data capture devices (like taps and network packet brokers) that operate at full line rate and don't drop packets. This is where you will want FPGA-based packet brokers. CPU-based packet broker designs often have bottlenecks that result in data being lost, often without your knowledge. FPGAs are the gold standard because they can operate up to line speed without dropping packets.

Incomplete/obscured visibility can lead to the wrong conclusions, unnecessary delays in fighting security threats, unnecessary security breaches, and unnecessary costs. This is because gaps in your data make it impossible to reassemble all of the session data. Data analysis then becomes impossible. Therefore, ensuring full, uninterrupted data collection is critical.

Pitfall #3 – Make Sure You Validate Your Security Architecture

Someone once said, "A great design is worthless if it fails a Litmus test." While it may seem obvious to thoroughly test your architecture, many shortcut the process because it takes too much time, costs too much money, or "just isn't needed because the design is perfect and rigorous testing isn't warranted." Unfortunately, security operations center (SOC) teams end up finding out this reasoning is flawed. The last thing you want is to discover a flaw in your design when your agency network is attacked.

The government has some excellent Zero Trust proof-of-concept environments, but it is impossible for them to test every type of architecture, vendor capability, software version, interoperability, performance, automation, attack complexity, application QoE, etc. Government agencies should consider using test and modeling tools to help them validate the completion of their goals.

In addition, validation isn't just required at initial deployment, it is needed all the time. With every change to your network (hardware updates, software updates, minor configuration changes to a firewall or intrusion prevention system (IPS), or SIEM, whatever), this can affect your network in hidden ways.

Breach and attack simulation (BAS) systems allow you to routinely test your network. You can set them up to automatically run once a month, once a week, once a day, or even once an hour. It's up to you, and how thorough you want to be with your validation strategy. BAS allows SOC engineers and administrators to measure their network to get operational insights into the effectiveness of their security posture and actionable intelligence to improve it. This type of testing ensures base level control system security using regular, comprehensive, and safe BAS assessments of the production network that provide innovative security solutions to enhance threat identification activities.

Penetration testing can be a useful tool, but this approach has some inherent problems. For instance, pen testing only tests one specific point in time. Again, the network changes and new threats are released into the wild at a fast clip nowadays. Hiring a pen tester to repeatedly test the network gets expensive fast. In addition, most pen testers just try to break into the network, they don't try to test every aspect of the security architecture. A BAS approach lets you pick and choose the exact aspects you want to test. If you want, this can be all of the aspects you are concerned about, not just a select few.

Here are some example use cases of a BAS approach:

- Perform continuous security monitoring and testing of live networks to check for security threats
- Routinely test the network after new configuration changes for configuration-created vulnerabilities
- Identify environment drifts between the current state and last week
- Test against newly released malware
- Reduce latency and risk by using computer analysis to formulate correct conclusions and recommendations to fix problems, display that information on a dashboard, and then transmit that information to a SIEM

If the right BAS solution is purchased, setup is very easy and cost-effective while being very effective at the same time. In addition, a solution that has a built-in "recommendation engine" makes it quick and easy to fix trouble spots within the network. That recommendation engine tells you what to fix and the exact location on the network.

Pitfall #4 – It's Not Just About Prevention, You Also Need to Respond Quickly

One final, but very important, note is that Zero Trust isn't just about defensive security. The sad fact is that you probably will encounter a successful attack against your network. There are two fundamental questions here: how painful do you want this experience to be, and do you want to try and prevent the attack from becoming a full-blown breach?

If your answer to those two questions is that you want to minimize the effects of the attack, then you need to include self-healing components within your architecture as well. With all of the emphasis on preventing an attack from being successful, this aspect cannot be forgotten about.

Here are two suggestions to consider in this area:

- Make sure your security architecture includes visibility components to optimize threat hunting activities
- Include cyber resilience functionality within your architecture

Once you find a threat, or suspect one, you will want to be able to rapidly implement threat detection (i.e. threat hunting) techniques. This means deploying a threat hunting solution that uses deep packet analysis to specifically look for potential threats early in the cyber kill chain processes of delivery, exploitation, and installation.

The best source of data for threat hunting is packet data. Some flow data adds value, but packets contain the critical details that threat hunting solutions need to perform deep packet analysis. This data can come from anywhere across your network. Once collected, it requires aggregation and filtering for irrelevant material before analysis for security threats.

Before you start a threat hunting practice, you must have the right information. This is true for any security architecture. The data must be complete and reliable. The wrong data will lead to wrong conclusions and missed security threats.

While this may sound rudimentary, even trivial, there are hidden traps everywhere. First, many enterprises use SPAN ports on their network switches to capture monitoring data. SPAN ports can drop packets when switches become overloaded. Also, SPAN ports automatically drop malformed and improper packets that could contain key pieces of information about when, where, and how a threat started. Worst of all, there is no indication that any of the data is being dropped and definitely not which pieces of data were dropped. You'll have to try to figure out all of that on your own.

SPAN ports are not the only concern. Surprisingly, some network packet broker solutions use a CPU-based design which drops packets under load or when multiple features and filters are employed. What's worse is that the CPU-based packet broker doesn't tell you it dropped the packets, so you are literally left blindsided.

A well-built NPB delivers wire speed data to threat hunting appliances for analysis. However, not all NPBs and taps are created equal. It is important to select taps and packet brokers that can process data at wire speed and not drop packets. While you want to filter out unnecessary data, you do not want your data collection architecture to randomly drop data, since that lost data might have been critical to your analysis process.

Concurrent to threat hunting activity, you should be able to implement cyber resilience capabilities that were built-in to your Zero Trust architecture. Cyber resilience capabilities allow you to control how painful an attack will be. Basically, cyber resilience is about mitigating breach risk and returning the network to a "normal" state as fast as possible.

So, what do we really mean by resilience? Traditional resilience refers to the ability of an entity to return to its original form after being bent, stretched, or compressed. From our perspective, we are specifically talking about the ability of an IT network to recover to a normal, steady state operations after a security attack or breach has occurred.

Cyber (or network) security resilience then is the set of activities that can be conducted to help the network after an attack happens. By adopting a resilient security architecture approach, the "time to observance" and "time to remediation" can be reduced. This is where cyber resilience comes into play. Cyber resilience helps you reduce the cost and risk associated with a data breach which is why important security best practice standards recommend it. For instance, the RECOVER section of the NIST Cyber Security Architecture Framework specifically calls out the need for cyber resilience functionality.

Here are just a few visibility-focused security resilience techniques that can be implemented:

- Optimize network continuity with external bypass switches and heartbeat messaging. These devices can be set to Fail Open or Fail Closed, as you choose. The reason for an external bypass is that if you have to completely replace a security tool (and you are relying upon an internal bypass), then your network goes down during the changeout.
- Inline and out-of-band network packet brokers using load balancing and n+1 survivability allow you to maintain operations during "impaired" network situations. The right packet broker choice also provides reversion capability which means that they can automatically sense when out of service security tools become operational again (i.e., if a security tool does a reboot and comes back online). This provides a "self-healing" component to your security architecture.
- Inline packet brokers with Active-Active processors provide enhanced business continuity without loss of data. Active-Standby solutions will lose data while the standby processor comes online.
- The ability to completely simulate the attack in your labs to validate any fixes is especially important. This is where you need a security threat generator to faithfully reproduce the security attack in your lab so that you can determine whether your security fix actually works. The last thing you want is to shoot yourself in the foot by rolling out a security fix that doesn't work. This could lead to another successful attack/breach and become a career limiting event for yourself.
- Something else to consider is network packet brokers that support integration to SIEMs. This allows your network to support automation to collect data faster and thwart security attacks as fast as possible.
- Start conducting cyber range training exercises so that you can recognize and respond to attacks faster. It's one thing to suspect that a certain type of attack has happened, or is happening, and another to be able to "see" the indicators of different types of attacks in real-time. Practice seeing these attacks in a cyber range is critically important. While you may not be able to tell a Petya attack from Ryuk, you can at least narrow down your search to the fact that it is probably a ransomware attack and proceed forward with that information.

To be clear, we're not suggesting that you stop trying to prevent a breach. You always want to do that. What we are saying is that you will want to add another set of capabilities to lower your agency risk and the cost of a breach.

How Keysight Solutions Connect to the CISA Zero Trust Maturity Model

The Zero Trust Maturity Model, created by the US CISA agency, represents a Zero trust Model that uses five distinct pillars of activity. There are also three, cross-pillar functions, as shown in Figure 2.

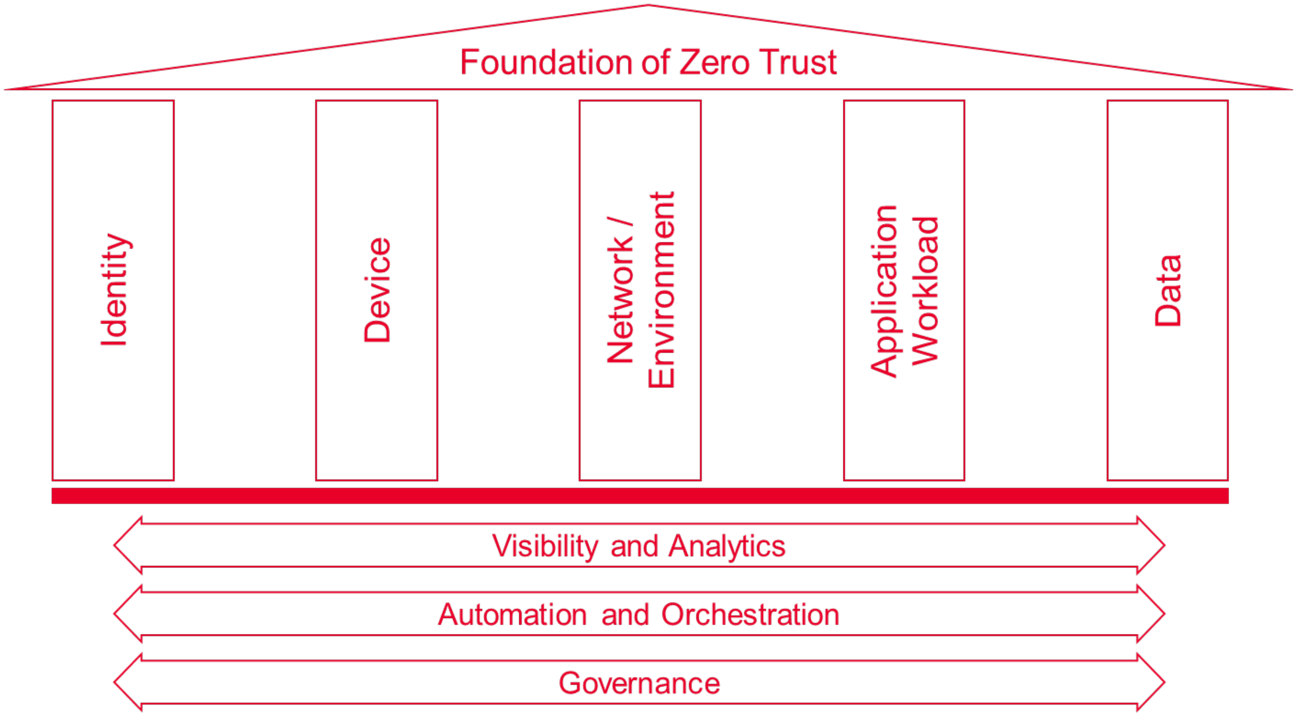

Figure 2: Illustration of CISA Zero Trust Maturity Model

NetworkDataPedia investigated the Keysight (formerly Ixia) visibility solution and found this to be the best visibility solution on the market. Specifically, Keysight Technologies can help you overcome the four pitfalls mentioned earlier, as well as help you with your overall visibility and analytics across the five pillars of the CISA Zero Trust model. Their taps, bypass switches, and packet brokers provide the visibility and confidence you need that you are seeing EVERYTHING in your network.

Keysight's advantage is due to its packet broker architecture — which uses FPGA's to process important data features, instead of a CPU running software. As mentioned earlier, the CPU/software combination has inherent issues. This can result in a feature blocking architecture (i.e. certain important features cannot run at the same time) and also has the distinct possibility of introducing data packet loss. This is a known problem with some packet broker solutions on the market. Keysight's FPGA design is superior because the FPGA works at line rate — which allows for faster data processing and eliminates delays that result in feature blocking and lost data.

Figure 3 provides an overview of Keysight's integrated Zero Trust solution that interweaves security products, forensics, and performance functionality to create a successful solution.

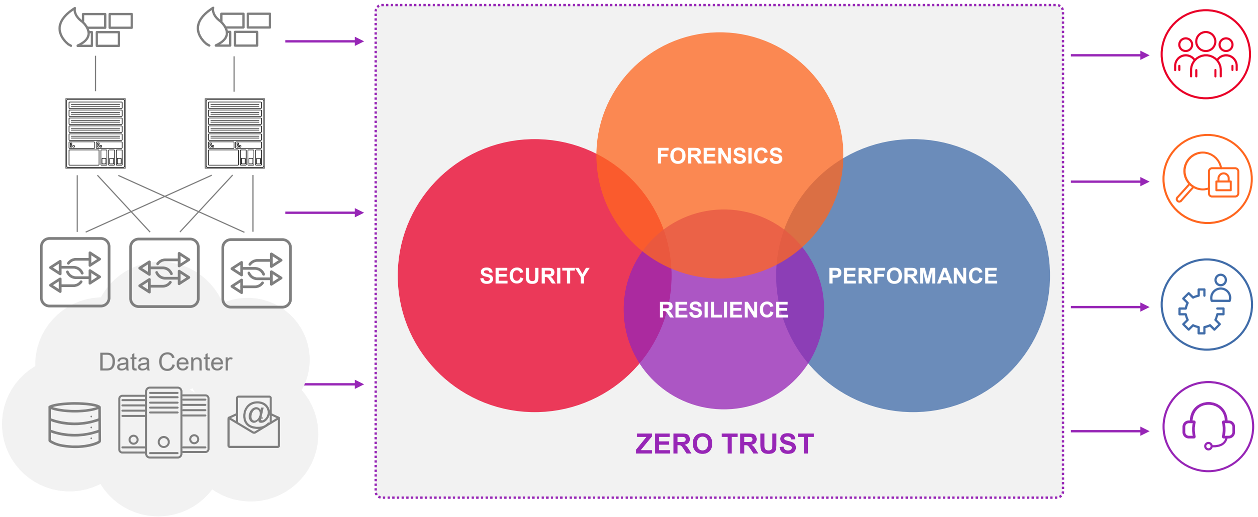

Figure 3: Generic Keysight Zero Trust Solution

Here is a quick summary of available Keysight solutions:

- Taps – Includes a vast array of interfaces including copper (10/100/1000 MB) and optical (1/10/40/100/400 GE). Keysight also has a large portfolio of tap split ratios including 50/50, 60/40, 70/30, 80/20, and 90/10 splits.
- Vision series packet brokers – Supports zero packet loss for full featured, non-blocking monitoring up to 400GE. Keysight's patented GUI interface is intuitive and easy to use which saves significant programming time and cost.
- Inline Vision series packet brokers – Supports internal high availability as well dynamic load balancing to create cyber resilience with security appliance survivability and self-healing architectures
- SecureStack application – Integrated SSL/TLS decryption for the Vision series packet broker that exposes hidden security threats while removes the inefficient and heavy decryption burden from your security tools
- AppStack application – Provides high-value intelligence features for the Vision series packet broker that delivers empirical data to identify bandwidth usage by application type, flow data, geolocation, and various pieces of user data to look for indicators of compromise and an early warning of potential problems.
- iBypass external bypass switches – Provides superior fail-over and fail-back techniques to increase network reliability and mission continuity for cyber resilience
- Threat Simulator – A BAS solution that performs continuous tests of your live network cyber security defenses, WAF, and web policy engines to identify any vulnerabilities. Once identified, a patented Recommendation Engine provides detailed easy-to-follow instructions on how to optimally configure your security products to close those gaps

and improving your security. These recommendations can also be integrated directly to your SIEM.

- ThreatARMOR – Threat intelligence gateway that stops incoming malware and outgoing security threats by leveraging known bad IP address information. It can also be used to isolate IP links to create an air gapping solution to thwart security threats.
- BreakingPoint – Simulates real-world legitimate traffic, distributed denial of service (DDoS), exploits, malware, and fuzzing, to validate an organization's security infrastructure and reduce the risk of network degradation by almost 80%.
- CloudLens – Allows you to capture and filter packet data in public and private cloud networks.
- Cyperf – Simultaneously generates both legitimate traffic mixes and malicious activities across a complex network of proxies, software-defined wide area networks, TLS inspection, elastic load balancers, and web application firewalls for cloud networks.

The chart in Figure 4 shows a mapping of the Keysight solutions to the CISA Zero Trust Maturity Model.

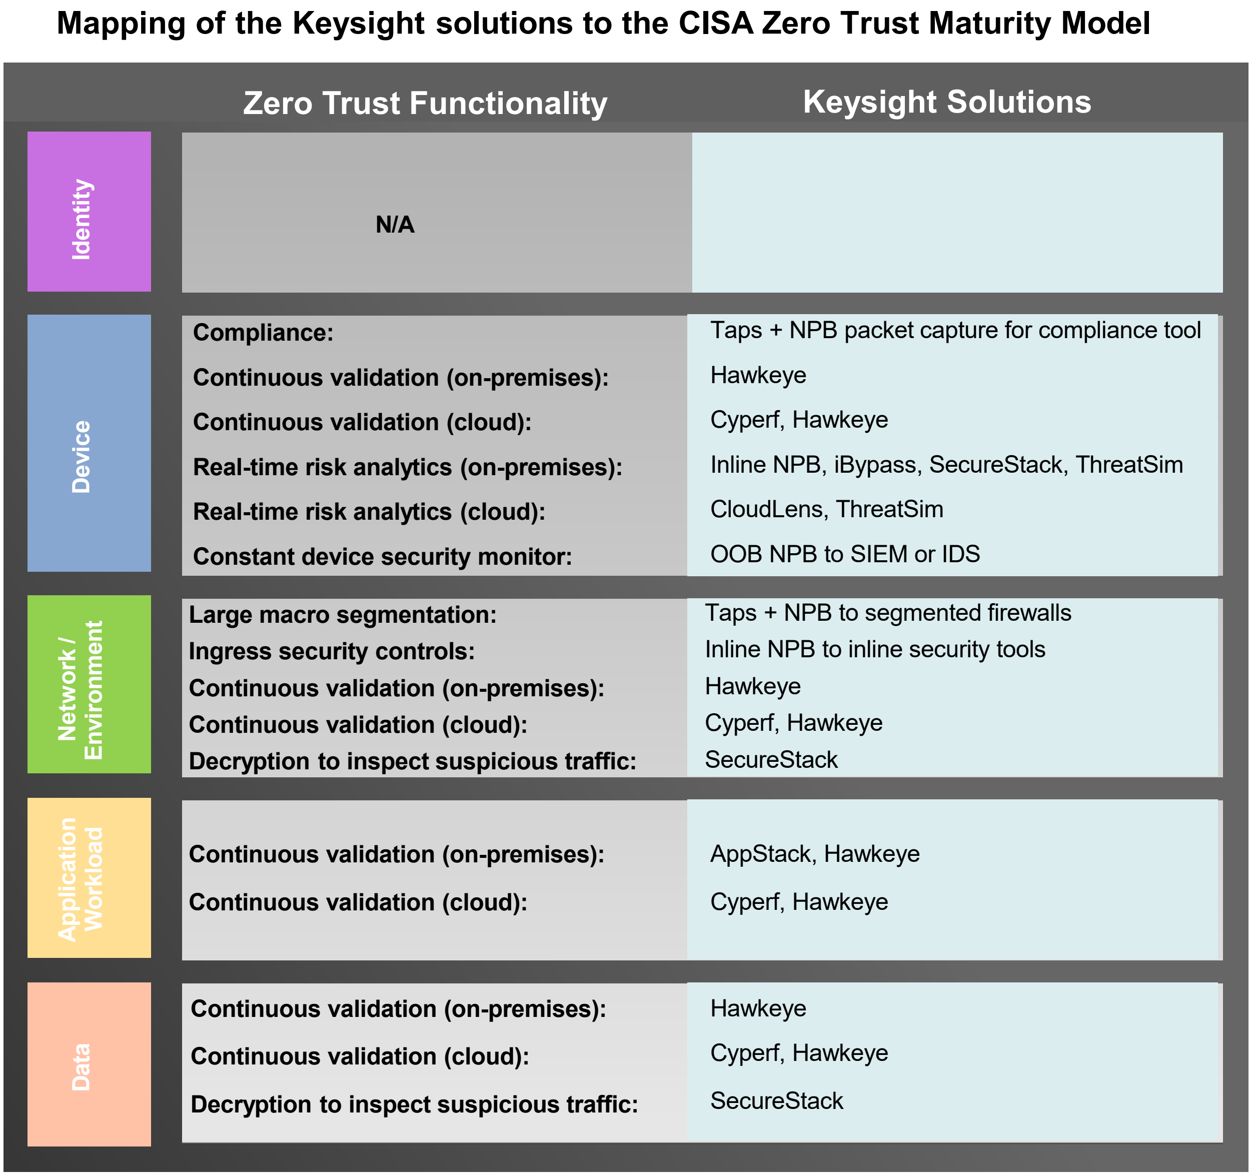

Mapping of the Keysight solutions to the CISA Zero Trust Maturity Model

	Zero Trust Functionality	Keysight Solutions
Identity	N/A	
Device	**Compliance:**	Taps + NPB packet capture for compliance tool
	Continuous validation (on-premises):	Hawkeye
	Continuous validation (cloud):	Cyperf, Hawkeye
	Real-time risk analytics (on-premises):	Inline NPB, iBypass, SecureStack, ThreatSim
	Real-time risk analytics (cloud):	CloudLens, ThreatSim
	Constant device security monitor:	OOB NPB to SIEM or IDS
Network / Environment	**Large macro segmentation:**	Taps + NPB to segmented firewalls
	Ingress security controls:	Inline NPB to inline security tools
	Continuous validation (on-premises):	Hawkeye
	Continuous validation (cloud):	Cyperf, Hawkeye
	Decryption to inspect suspicious traffic:	SecureStack
Application Workload	**Continuous validation (on-premises):**	AppStack, Hawkeye
	Continuous validation (cloud):	Cyperf, Hawkeye
Data	**Continuous validation (on-premises):**	Hawkeye
	Continuous validation (cloud):	Cyperf, Hawkeye
	Decryption to inspect suspicious traffic:	SecureStack

Figure 4: Mapping of Keysight Solutions to the CISA Zero Trust Maturity Model

Conclusion

Government agency SecOps teams are under ever-increasing pressure to maintain a high level of network security and performance. To meet these challenges, they must implement security architectures that they know are both effective and maintainable.

Network visibility is a crucial component to network security for the simple reason that you cannot defend against what you cannot see. At the same time, you don't want to implement "a" visibility solution, you want to implement the "right" visibility solution. This solution will consist of physical taps and cloud taps, FPGA-based packet brokers that can perform layer 2-4 as well as layer 7 data filtering in a non-blocking architecture at line rate, and the right capabilities to optimize data for your security solutions so that your security tools make the right analysis as fast as possible.

Whether you are looking to achieve M-21-31 or M-22-09 compliance, or to enhance your Zero Trust architecture, Keysight is available to help. They have various network visibility and network security solutions that help with both NIST and CISA compliance. Reach out to Keysight Technologies and they will show you how to optimize your security solutions.

Learn more at: www.getnetworkvisibility.com/ZeroTrust