Zero Trust for MSPs:
The Definitive Guide

So, you're really interested in Zero Trust, eh?

I'm glad you're here. By now, you've heard from me that vendors don't solve Zero Trust. They might help solve it. But just like Smokey Bear said with forest fires, the same holds true for Zero Trust: Only you can prevent… wait that doesn't work. But only you can create Zero Trust.

What is Zero Trust?

Zero Trust is a framework. It's a Zen. It's a thing we do. It gives some freedoms around how we approach things. As always risk appetites and mitigating controls should rule the day here.

Zero Trust is, at its core, a framework that requires every user and device (physical or virtual, permanent or ephemeral), inside your network or outside of it, to be authorized, authenticated, and validated (continuously) before it can operate.

It's Dangerous to Go Alone!

Here's what I mean: if you're starting with Zero Trust, as amazing as it is, but you haven't yet begun alignment to a security framework, you should really start there first. You see, security frameworks are more holistic. The creators of Zero Trust NEVER intended to supplant security frameworks like the Center for Internet Security's (CIS) 18 Controls.

Check out our podcast on the CIS Controls, hosted by Andrew Morgan. It's yours truly, Ryan Weeks and the literal Director of the CIS Controls Phyllis Lee. We cover every control at length. It's hours and hours of content.

https://www.thecybercast.com

Getting Started on Zero Trust

Alright, so NOW you're ready for Zero Trust, eh? Good. Because you get a huge advantage that I waited until just now to tell you. If you're already following the CIS Controls, especially though the beginning level one, you'll notice that a lot of what I outline below you're already doing.

There are a few areas where I think you can really focus on as an MSP to get the most bang for your buck, so to speak. I'm cognizant that you have to juggle scores of clients, each with their own risk tolerances and needs. I also know you don't operate with the complexity or budget of Bank of America. You don't have an army of devs. You likely don't have a single cybersecurity analyst. (And if you do, that's a luxury most folks don't have… but they're likely tied down with far too many other projects.)

Here four tenants that we should ascribe towards in an ideal Zero Trust environment (thank you Microsoft):

1

Identities are validated and secure with multifactor authentication (MFA) everywhere. Or as much as possible. Favor Authenticators over text based one-time passcodes (OTP). Leverage FIDO2 devices if possible. Seek to replace passwords with biometrics. (Ok that last one is a stretch, but those days are coming.)

Assess Yourself:

MFA: Identity Management Provider:

Are authenticators exclusively in use?

If not, where do you (or the client) still utilize OTP?

Have you assessed the use of FIDO 2 devices?

Where are you at when it comes to biometrics like Windows?

2

Devices are managed and validated as healthy. You need to establish a minimum standard of device “health” before it can access any company resources. There are multiple ways to do this. You have an RMM, so use it. But also, don’t forget about things like conditional access from Microsoft. And there are even tools like Network Access Control (NAC), or other configuration tools that can help do the job. Lastly, don’t forget about the cloud. Are you monitoring the cloud for the right configurations?

Assess Yourself:

Have you established a written minimum “health” standard for user and device configurations?

What tools are at your disposal to enforce device health standards?

Do you have a process in place to be sure all devices, even newly deployed, are enforced to a minimum healthy standard before being granted resource access?

3

Telemetry is pervasive. You must have standardized auditing as much as possible when it comes to network data, device data, and identities. That doesn't mean you log everything and keep it forever. But you need to pull in as much relevant data as you can. And you should ensure that those data points are monitored 24x7 by a qualified managed detection and response firm (MDR). Be sure you MDR understands their role in Zero Trust monitoring.

Assess Yourself:

What data sources are we logging today?

What data sources are we NOT logging but should?

How long are we keeping those data sources?

Do we have an MDR watching over our data 24x7?

What types of data does our MDR monitor?

What is our MDR provider's documented strategy on monitoring data?

4

Principle of least privilege is enforced and deployed everywhere. You should begin to limit access only to applications, services, and infrastructure (on-prem or cloud) required to perform the job function. It it has pre-requisites like data flow diagrams and software inventories (including inventory of privileged accounts.) Your scope should include network segmentation as well as user segmentation. This can (and should!) include cloud segmentation as well.

Assess Yourself:

Do we have a data flow diagram?

Do we have a software inventory of everything running in our networks and cloud?

Do we have a list of all privileged accounts and system accounts for all systems?

Have we disabled local administrator at the endpoint for all users?

Have we adopted a privilege access management (PAM) tool?

Are we use it a Zero Trust Network Access tool to control traffic flows including cloud?

The Zero Trust Litmus Test

Whew! That was a lot of work. No doubt you have many things here to start on. Just remember, Rome wasn't built in a day. And I would venture zero honest CISOs would ever say their organization is at a Zero Trust level they're ultimately finished with. It's an ongoing process.

At this point, it might be helpful for you to have an objective standard to aspire towards when you're building out Zero Trust. Here's four scenarios Microsoft provides us that are extremely helpful to judge ourselves against. Use these to determine any gaps you may have in accomplishing a minimum standard:

Litmus Test 1: Applications and services are capable of utilizing MFA and ensuring device health where possible and necessary.

Litmus Test 2: Employees can enroll new devices into your systems and guarantee the health of the device before being granted resource access.

Litmus Test 3: Employees and non-employees (where applicable) have the ability to access corporate resources when not using a managed device.

Litmus Test 4: Access to resources is limited to the minimum amount required to perform a function. Both the user running with minimum privileges and accessing the minimum amount of data necessary.

What Systems, Data Types, and Resources Do We Start With First?

Remember that pesky data flow diagram we talked about before? It's really critical here. You should start deployed Zero Trust around your Crown Jewels.

For example, a CPA client, you might want to focus on the tools they use most, the data repositories that store the most critical data, and the line of business (LoB) applications they use most.

Make an **inventory** like:

- User Accounts: Local, Azure AD, LoB Accounts, oAuth to a sketchy AI app
- Devices in use: MacOS, iOS, Azure Servers
- Tools: Microsoft 365, QuickBooks, Office, and Canva
- Cloud storage: Dropbox Cloud, SharePoint, OneDrive, Canva (see how it shows up twice?) and some sketchy AI app you discovered.
- Data flows (how the client operates and how that data flows for them to accomplish their work)

Notice that with this inventory, you can now begin to design and implement things like:

- Network segmentation for on-prem and cloud access
- Privileged account controls locally and in the cloud
- Conditional Access in Microsoft
- Deployment of MFA for every account and every network

While this example is simplistic, do you see how easy this can become? Once you start with the Crown Jewels, you can then expand Zero Trust to other less sensitive areas of the client network. In fact, this is exactly how the creators of Zero Trust intended us to complete our work.

The Microsoft Approach to Zero Trust

Here is how Microsoft puts all of this together as an approach to Zero Trust across their entire ecosystem. I believe this is a helpful tool MSPs can use as they begin their designs towards Zero Trust.

Assess Yourself:

Which of the major columns above are you strongest in?

Which of the major columns above are you weakest in?

Which individual areas in the graph above will you focus on first?

Which individual areas in the graph above are the most challenging for you to solve?

Finishing the Zero Trust Puzzle

As you come to the end of your planning and strategy for Zero Trust, you should document and diagram how the Zero Trust journey should flow for the client.

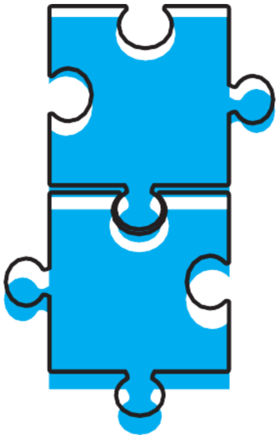

Just start with what you can control. And you can then begin to document for yourself (and the client) certain areas that simply cannot be addressed without making architecture (or even vendor) changes. And also understand sometimes you'll end up with areas that feel like an epic and colossal dumpster fire (looking at you Mac… as I write this on a Mac, oh this irony) (obligatory: there are tools that can help build towards Zero Trust with Mac of course.)

Remember the mantra here: don't let perfection be the enemy of good. Zero Trust can take years to get to a standard level of health that you'll be happy with. The key is building a strategy, and getting towards incremental progress as you go along.