Zero Trust Connectivity with adam:ONE®

White Paper
June 2022

David Redekop | Francois Driessen

//adamnet.works

Content

Background

The Missing Tenet

Zero Trust (ZT) is a set of philosophical tenets that was birthed out of security failures of the past. The core tenets of the ZT philosophy carries huge value, but it is important to note that there are substantial gaps between the raw philosophies and the practical implementation that is possible for most security administration teams today.

Having a great idea but falling short of the ability to execute renders it mostly useless. One of the primary tenets that we see as a requirement for effective security is the ability to execute that philosophy.

> *Zero trust presents a shift from a location-centric model to a more data-centric approach for fine-grained security controls between users, systems, data and assets that change over time; for these reasons, moving to a ZTA is non-trivial.* - CISA Zero Trust Maturing Model Pre-decisional Draft v1.0 Jun 2021

The Marketing Fog

Another important aspect to note is that the Term *Zero Trust* has become an industry buzzword that marketers are attaching on to their products to appear more relevant. However, just using the term does not make a security solution provide the answer as to the needs for why Zero Trust exists.

For the sake of distinction, we will refer to the term Zero Trust Connectivity (ZTc) to address the specific aspects of the Zero Trust philosophy that is executed through controlled networked connectivity of not only verified assets, but even more importantly, between verified assets and any other possible outgoing connections, including the public internet.

Design from First Principles

Zero Trust is a reaction to security failures of previous philosophies. However, it is about much more than just user identity and access.

In our experience, to arrive at practical solutions that solve these failures, one has to identify the design flaws that resulted in the failures. Then, once Identified, implement design from first principles to address the issues effectively.

To clarify design failures we address in the context of inter connected devices, it is helpful to break down the evolution of the internet and its technologies into 3 phases.

1. Build Application Functionality - “Get stuff to work”
2. Optimize Application Delivery/Usability/Utility - “Get stuff optimized”

3. Secure the Application, Protect the User and Data - “Get it secured”

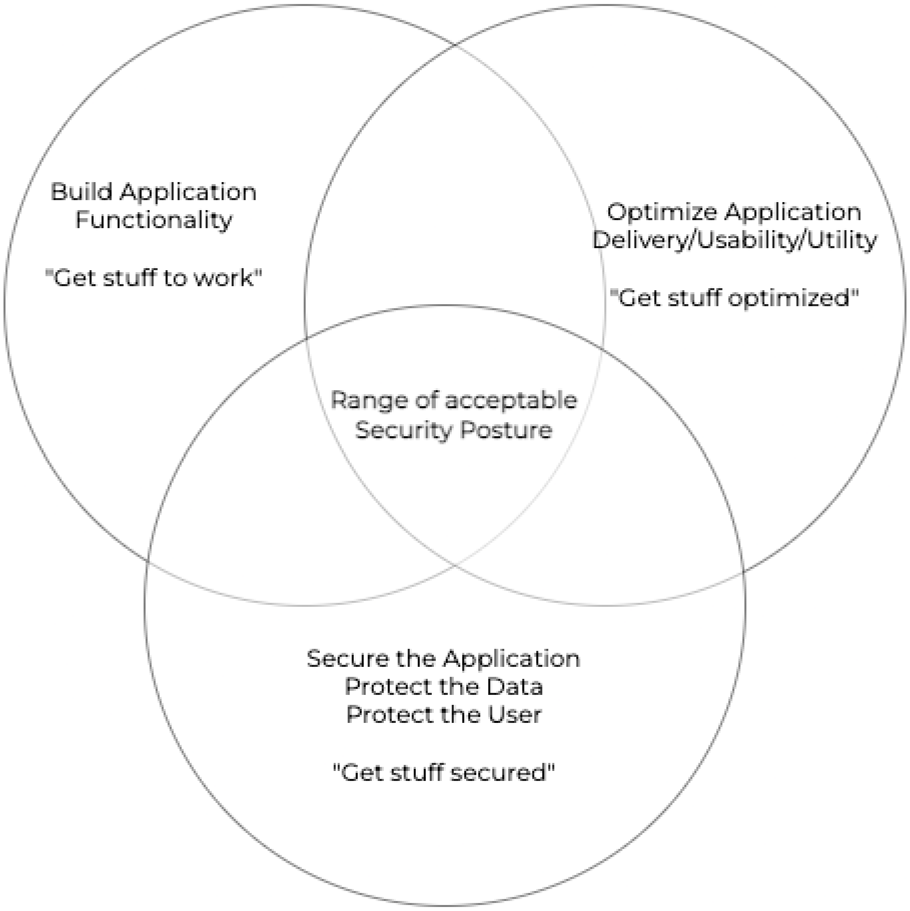

The building blocks of the Internet and how assets connect have matured well for phase 1 and 2. But the fundamental lack in phase 3, that of security, is quite evident today. The net result is the current state of the cyber threat landscape.

When human interaction is involved, this creates a three-point paradox between elements that seem to be at odds with another. Security has an experience and convenience (optimization) cost and optimization or convenience of operation has a security cost. It is within the balance of all 3 that effective and practical security exists.

At the core a Zero Trust approach inverts the original design intent for the internet of being able to route around any obstruction. That legacy approach left us with the basic philosophy of:

“**Trust by default, deny threats**”

As threat actors take advantage of this fundamental and default philosophy, we have no choice but to invert it, to become:

“**Deny by default, trust verified connections**”

This has been the foundation of the development for adam:ONE®. Designing from First Principles has been the approach that piloted the development. The fact that the rising philosophies of Zero Trust align with the basis for the design tenets we developed for adam:ONE® is encouraging. It is for that reason that we use the terminology as it aligns with the path that we are already on.

We hope that the practical solutions that exist with adam:ONE to execute these philosophies are as valuable to others as it is to us in protecting people and critical infrastructure.

Main goals for the adam:ONE® design that intersect with Zero Trust philosophies

1. Implementation must be practical.
2. Remain operational while assuming breach.
3. Leakproof egress control.
4. Apply protection proactively, without needing a detection trigger.
5. Deny all connections and access unless verified to be necessary and good.
6. Achieve real-time visibility to each connection.
7. Achieve maximum visibility to each asset.
8. Automated real-time device inventory.
9. Protection for all asset types. Access devices, IoT, OT including legacy assets, regardless of OS.
10. Aggregate threat intelligence.
11. Multiple layers of protection. Defense in depth.
12. Distributed custodial protection.
13. Granular control to enforce policies of access and execution. Remain dynamic when appropriate.
14. Flexible Perimeters that remain dynamic.
15. Circumvention Protection.
16. Encryption friendly. Should not break any E2E encryption during protection.
17. Privacy friendly.

Predecessor technologies

In our brief internet security history, we went from firewalls to transparent packet inspection to next-generation Unified Threat Management (UTM) complimented by threat intelligence sources fed by constant threat hunting. This was all upended by our rapid adoption of encryption of traffic in transit.

One of the control planes that partially remained in place was DNS firewalling. Provided network devices behaved as they should, this prevented known threat actors from being reached by disabling DNS queries to known bad actors from resolving. The result was actually an acceleration of the cat-and-mouse game of threat intelligence feeds and malicious actors using new domain names, including widespread use of dynamically-generated algorithmic (DGA) domains. Advanced DNS firewalls subsequently complimented network control systems with IP address reputation systems.

Next, we observed vendors centralizing traffic redirection in order to make use of the increased capacity in the cloud to inspect traffic. Such vendors make varying use of proxies, issuing CAs and TLS certificates of their own to endpoints, in order to maintain visibility and manageability of all traffic as illustrated:

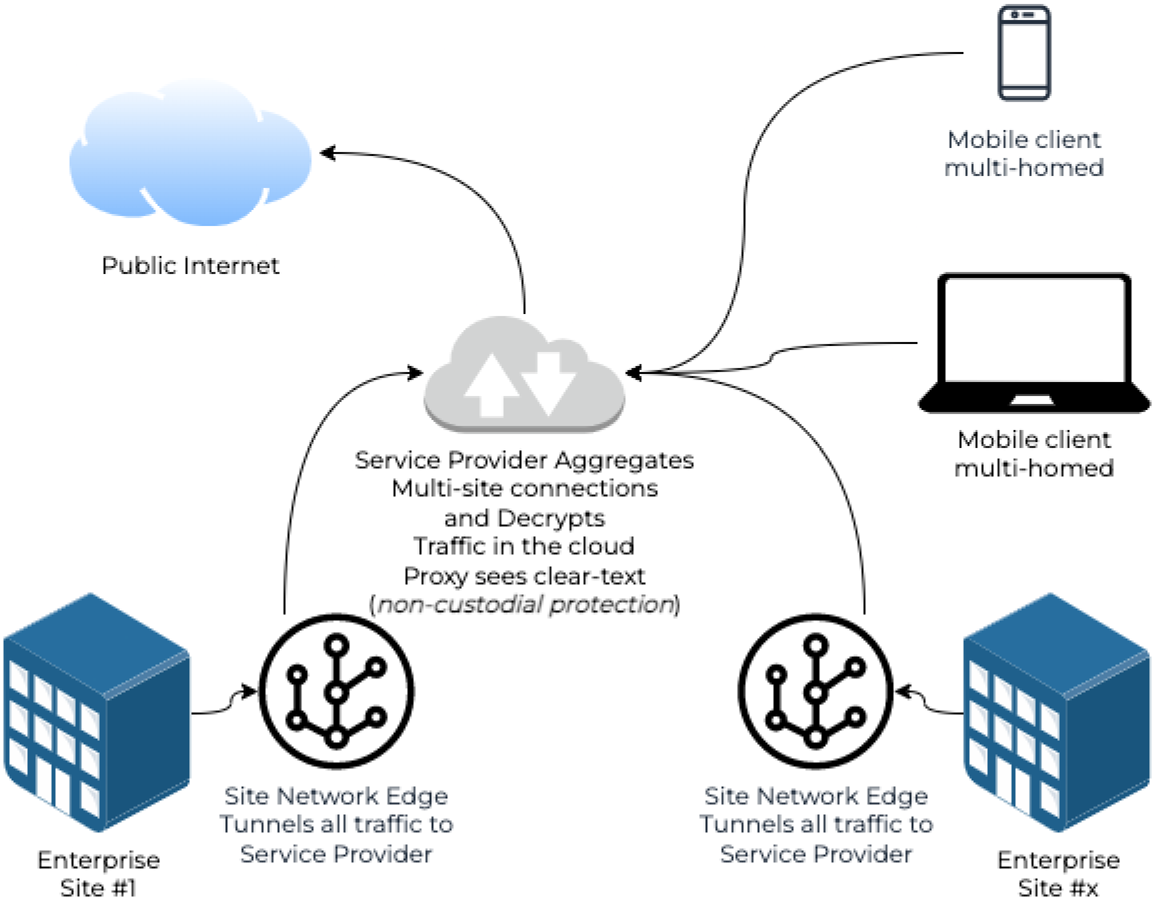

Well-intentioned service providers rely on the ability to decrypt SSL/TLS traffic via proxies or self-signed certificate-trusting methods.

Honouring **TLS** and **end-to-end encryption** without proxies

Endpoint users and applications rely on the trustworthy connection to a service provider. Network administrators are traditionally at odds with this approach since they carry the responsibility of keeping the endpoint safe from threats.

In order to maintain end-user privacy and honour end-to-end encryption, the security must necessarily be broken into two separate control systems:

1. Secure Access Service Edge (SASE) - responsible to allow or disallow all OSI layer 3+ connections
2. Endpoint security - responsible to inspect and protect from malicious intent post-decryption and pre-encryption

Comparatively speaking, distributed custodial protection looks like this:

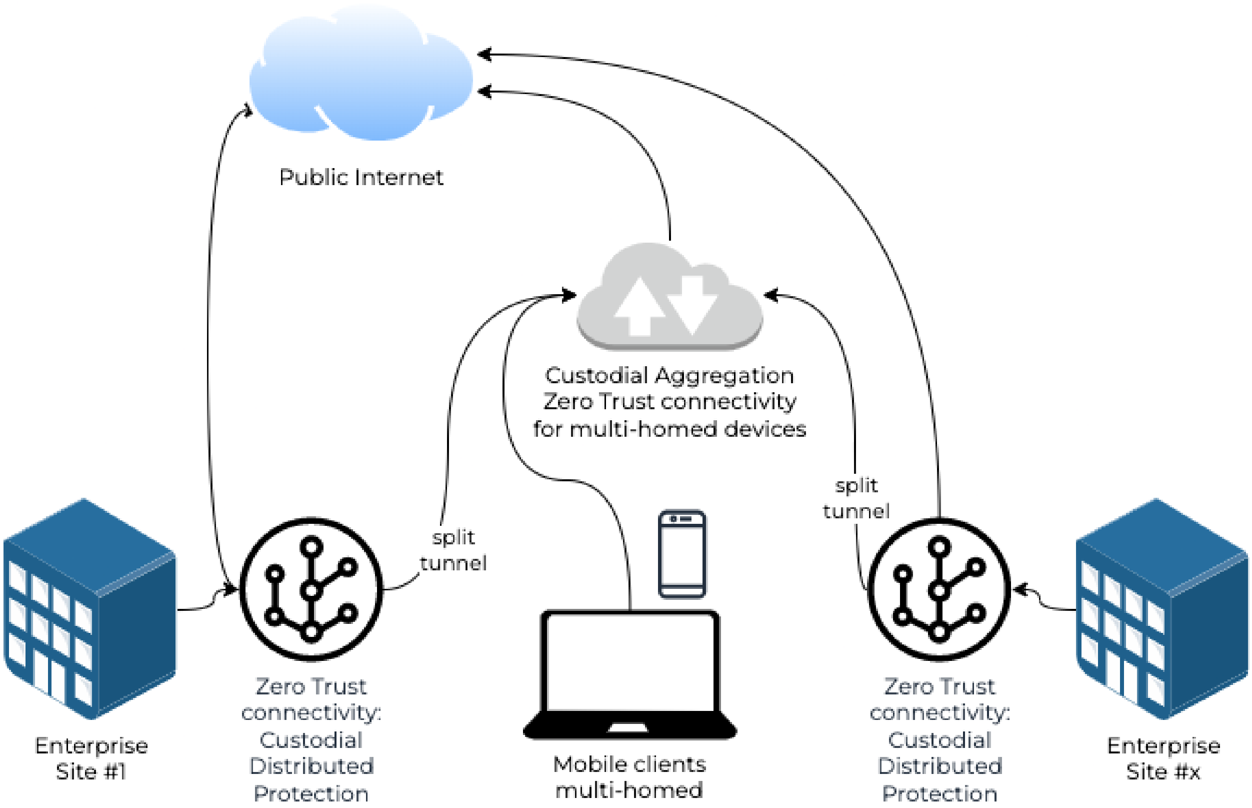

Only this approach honours the expected privacy and security with end-to-end encryption.

Distributed Custodial Network Protection

Muscle-Brain

Using an on-premise gateway-hosted client (the “muscle”) which receives its configuration continuously from a centralized cloud controller (the “brain”), resilience and local performance are simultaneously optimized. DNS queries and answers are channeled to be in geographic proximity to the endpoint and all network traffic decisions are made locally.

AI powered dynamic allow-listing

In a default-deny policy where we commit to the “Deny all, trust verified connections” approach, there must be flexibility in creating policies that match the right level of security posture, considering convenience and practicality.

Multiple different approaches in Allow-listing offers right-sizing a policy to match the role of a user or device. In summary:

Allow-list policy type	How it works	Typical use case
Fixed rule(s) policy	Using pre-made lists or by observation, a collection of required internet domains and/or fqdns are established as an allow-list	Active Directory Controllers File/Application Servers Web Servers IoT Devices Printers/Scanners
Adaptive Allowlisting policy	Per-site/per-network policy adds dynamically each time a domain is: ● Requested by the user ● Passes cloud-based sandbox testing and AI inspection ● Is in a safe-for-work category ● etc	End-user devices such as computers, laptops, devices where: ● User’s role is very defined such as Accounting ● Protecting user/device from all phishing is essential ● Circumvention protection is paramount ● High value assets
Reflex Allowlisting policy	Per-device allow-listing policy where 80 categories are sorted into 5 decision buckets: ● Always Block ● Bock ● Review ● Allow ● Always Allow Conflicts resolution can then be specified as Permissive, Protective, or Hold for Human	Any user requiring relative permissive policy while experiencing strong protection, specifically ones that do not fall into categories mentioned above.

Don't Talk To Strangers (DTTS)®

A stranger in this context is an IP address that has not been resolved via a permitted DNS query. Disallowing such communications essentially follows this logic:

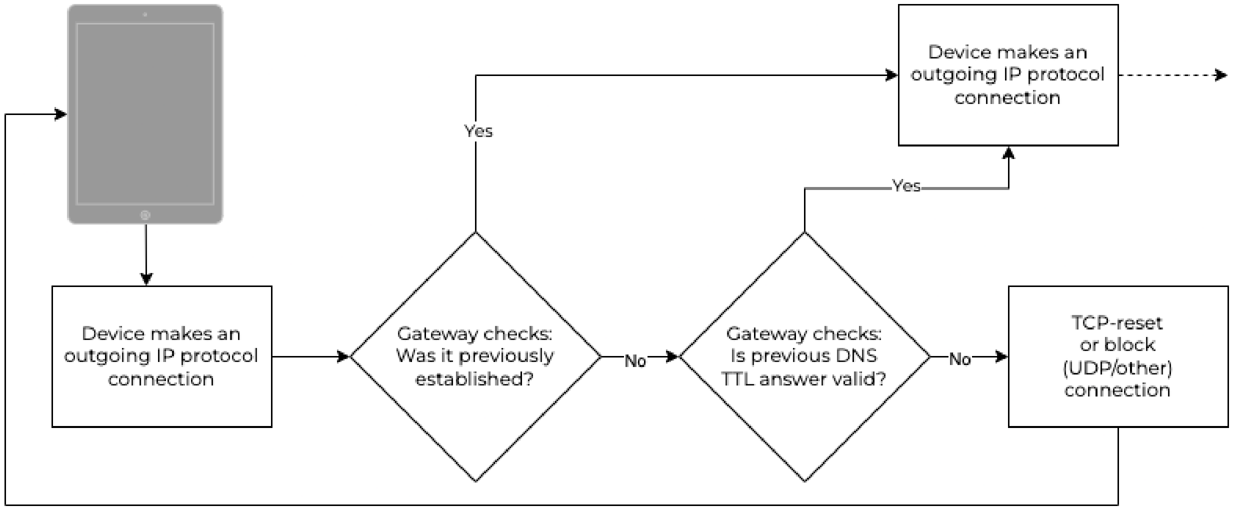
Device makes an outgoing IP protocol connection
Yes
Yes
Device makes an outgoing IP protocol connection
Gateway checks: Was it previously established?
No
Gateway checks: Is previous DNS TTL answer valid?
No
TCP-reset or block (UDP/other) connection

All modern Advanced Persistent Threats at some point in the attack chain rely on a direct-by-IP connection to a Command & Control (C2) operator waiting for "phone-home" connections. These are notoriously difficult to detect even by Security Operation Centers (SOCs) and all too often are identified only when it is too late. DTTS® stops the threat before it does damage.

It is important to note that this approach leads to many thousands of firewall rule changes per minute at each layer 3 connection on a typical network and yet it is entirely scalable as they can all be done in user-space.

DTTS® makes DNS firewalling **leakproof**.

Exceptions to allowing strangers are easily administered via **Enablers**. An **Enabler** is a simple definition of destination IPv4/IPv6 address spaces along with protocol and port numbers. These are typically needed for latency-sensitive services such as voice and video, which can easily be identified and enabled with supplied administrative tools.

DNSharmony®

Domain-based reputation systems and threat intelligence sources are readily available both commercially and free/open source. There is no need to rely on a single source. That's the concept behind DNSharmony®. Consider this scenario:

- End-user device wants to resolve `fqdn.example.com`
- For the most acceptably-permissive policy, consider aggregating the following threat intelligence sources:

    - Cloudflare 1.1.1.3/1.0.0.3 (block known malware and adult content)
    - CleanBrowsing Security Filter (185.228.168.9/185.228.169.9)
    - Quad9 (9.9.9.9/149.112.112.112) - itself containing commercial sources
    - Enterprise-sourced Threat intelligence via bind RPZ subscriptions
    - Pi-hole running on a local network edge
- A DNSharmony® powered policy would block access to `fqdn.example.com` if any of those providers answered with their respective block (NXDOMAIN, block page IP, etc)

This effective aggregation of any and all available sources allows for superior protection against known threats.

Additional Network Security Principles applied to ZTc

In order for devices and channels of all kinds to be protected with data integrity, additional concepts must be applied.

- Must not rely on proxies anywhere (whether local, or by a third party operating in the cloud)
- Do not break encryption (e.g. SSL decryption appliances)
- Offer flexible perimeters for multi-homed or mobile devices

Network Segmentation

This technique is a reasonably light administrative load when executed properly, and even internal resources are accessed via FQDNs vs IPs. Rather than offering a business unit access to `10.12.13.14` as an internal application, consider `https://app1.internal.example.com` which also offers the flexibility of moving the resource anywhere without the administrative task of future pointer changes.

Typical Network Segments include a segregation of role-based functions per site, for example:

- Active Directory Controllers
- Network Management (Switches, controllers, etc)
- File Server(s)
- Printers/Scanners
- IoT Ethernet (including security cameras, door control systems, etc)
- Staff WiFi (802.1X)
- IoT WiFi
- Guest WiFi

Microsegmentation is also an effective strategy, especially for multi-homed devices. The distributed custodial protection shown above includes microsegmentation by default.

adam:GO™ can be distributed by way of an MDM-pushed always-on VPN to the aggregation site(s).

The default firewall rules between segments permit no traffic. Traffic is allowed only as defined by the device/user policy.

Network segmentation is sometimes discounted in modern Zero Trust philosophies, which unfortunately undermines a significant source of threat vectors. Network devices which aren't end-user devices have no other protection mechanism available. Consider the high level of vulnerabilities present almost universally in printers, scanners, cameras, control systems. Vulnerabilities are slow to be fixed at the source, and even slower to be deployed. ZTc provides an important protection against lateral movement from such devices.

Layer 2 visibility

Every device in a network must have layer 2 visibility to the SASE.

Inventory Enforcement via Quarantine State Default Policy

This approach ensures that any time a new device is discovered at the SASE, it is automatically assigned a policy which passes typical connectivity checks and nothing else. This limited connectivity state causes the device to reveal its next-stage functionality and thus makes it very practical to passively investigate and apply an appropriate policy.

Per-user policy assignment leveraging 802.1X

In an age where most modern mobile devices use random MAC addresses, 802.1X facilitates the immediate device:user mapping for policy assignment purposes.

Making TLS/SSL blocking/redirection user-friendly without breaking trust

The challenge of offering a user-friendly page to a browser when a secure site is blocked is solved. It isn't solved by trusting a self-signed certificate but rather by a browser extension that reads the TCP reset and redirects to a friendly block page at `mytools.management` (see `adamnet.io/extension`). The result is transparent to the user without breaking SSL/TLS.

DoH and DoT - proper placement of secure DNS

In cases where the DNS channel needs to be protected, it must be done internally. DoH-based C2 channels alone are a justification for disallowing enterprise user devices to use publicly-available DoH servers.

Disabling UDP reflection attack participation

Two most commonly-used standard protocols that make it easy for an adversary to launch UDP-based reflection and amplification attacks are DNS (operating on UDP port 53 by default) and NTP (operating on UDP port 123). This attack vector is completely mitigated by:

a) Offering DNS and NTP services on all internal segments and
b) Forcing internet-bound requests to be answered locally

The result of Zero Trust connectivity

- Strong mitigation against typical adversarial network-based attacks (for example, Active Directory is not reachable by an attacker attempting to move laterally)
- Protected against unknown threats from tomorrow
- Attack surface reduction to near Zero – effectively 7000:1
- C2 communication is prevented before it can occur
- Data exfiltration protected
- Human Factors considered, phishing protection
- Spear phishing mitigated. IPv6 has enough unique addresses to attack every person 47 Octillion times
- Shadow IT cleanup, no unauthorized remote connections remain possible, including on BYOD endpoints
- Circumvention protection
- Privacy
- Productivity
- Distributed Custodial Protection
- Decentralized Performance, Centralized Control

Platform availability

pfSense® (available with optional HA and SLA)
ClearOS® (via marketplace)
Debian (available on managed platform)
Asuswrt-Merlin (for evaluation, POC, SOHO purposes)
Docker Container
VyOS® (coming soon)

Binaries are available for multiple architectures including x86/ARM/MIPS.

Glossary

ZTc	Zero Trust connectivity
SASE	Secure Access Service Edge
DoH	DNS over HTTPS
DoT	DNS over TLS
CA	Certificate Authority
FQDN	Fully Qualified Domain Name
DNS	Domain Name Service
NTP	Network Time Protocol
E2E	End-to-End Encryption
BYOD	Bring Your Own Devices
IoT	Internet of Things
OT	Operational Technologies
MITM	Man in the Middle
HA	High Availability
SLA	Service Level Agreement

Contact information

For further contact with ADAMnetworks, please feel free to reach us here:

https://adamnet.works/contact/

David Redekop
Francois Driessen