

the
GORILLA
GUIDE[®] to...



Zero Trust: Using DNS as Your First Line of Defense

How To Leverage Your DNS to
Protect Your Apps, Users, and Data

ED TITTEL



POWERED BY  **ActualTech**
MEDIA

Zero Trust: Using DNS as Your First Line of Defense

By Ed Tittel

Copyright © 2022 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review. Printed in the United States of America.

ACTUALTECH MEDIA

6650 Rivers Ave Ste 105 #22489
North Charleston, SC 29406-4829
www.actualtechmedia.com

PUBLISHER'S ACKNOWLEDGEMENTS

EDITORIAL DIRECTOR

Keith Ward

DIRECTOR OF CONTENT DELIVERY

Wendy Hernandez

CREATIVE DIRECTOR

Olivia Thomson

SENIOR DIRECTOR OF CONTENT

Katie Mohr

WITH SPECIAL CONTRIBUTIONS FROM EFFICIENTIP®

Alexandre Chauvin-Hameau

Ronan David

Marc Gourvenec

Surinder Paul

Pascal Tangapregassam

ABOUT THE AUTHOR

Ed Tittel is a 30-plus year veteran of the IT industry who writes regularly about cloud computing, networking, security, and Windows topics. Perhaps best known as the creator of the Exam Cram series of certification prep books in the late 1990s, Ed writes and blogs regularly for GoCertify.com, Win10.Guru, ComputerWorld, and other sites. For more information about Ed, including a resume and list of publications, please visit EdTittel.com.

ENTERING THE JUNGLE

Introduction: The Importance of DNS	8
Chapter 1: What Is DNS? What Does It Do for You?	10
The Origins of DNS.....	10
Chapter 2: How DNS Can Enable Cyberattacks	13
The DNS Security Landscape.....	15
Types of Attacks.....	16
Chapter 3: Understanding ‘Zero Trust’	19
Perimeter Security Is Not Enough.....	19
The Principles of Zero Trust.....	20
Zero Trust Facts and Figures.....	21
The Zero Trust Perspective.....	22
Zero Trust Architecture Building Blocks.....	23
Chapter 4: How DNS Strengthens Zero Trust to Better Protect Apps, Users, and Data	25
Setting the Stage.....	25
Security Challenges.....	27
Chapter 5: The 5 Pillars of EfficientIP 360° DNS Security for Enabling DNS-Based Zero Trust	29
1. Client-Based Application Access Control Powered by DNS Client Query Filtering.....	29
2. Unequalled DNS Analytics for Behavioral Threat Detection at the Client Level.....	30

3. Proactive Security Powered by Threat Intelligence.....30

4. Adaptive Countermeasures to Ensure Service Continuity...31

5. Automated Sharing of Actionable Events and Data for Accelerated and Efficient Remediation.....32

Chapter 6: Secure DNS Use Cases and a Real-World

Example.....34

 DNS Security Use Cases.....35

 Case Study: STMicroelectronics.....38

 Protecting DNS Has Never Been More Important.....39

CALLOUTS USED IN THIS BOOK



SCHOOL HOUSE

The Gorilla is the professorial sort that enjoys helping people learn. In this callout, you'll gain insight into topics that may be outside the main subject but are still important.



FOOD FOR THOUGHT

This is a special place where you can learn a bit more about ancillary topics presented in the book.



BRIGHT IDEA

When we have a great thought, we express them through a series of grunts in the Bright Idea section.



DEEP DIVE

Takes you into the deep, dark depths of a particular topic.



EXECUTIVE CORNER

Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



KNOWLEDGE CHECK

Tests your knowledge of what you've read.



PAY ATTENTION

We want to make sure you see this!



GPS

We'll help you navigate your knowledge to the right place.



WATCH OUT!

Make sure you read this so you don't make a critical error!



TIP

A helpful piece of advice based on what you've read.

The Importance of DNS

Welcome to The Gorilla Guide® To... Zero Trust: Using DNS as Your First Line of Defense, Express Edition. This book is for everyone with an infrastructure that uses DNS and needs to make sure it's secure. In other words, this book is for everyone whose infrastructure uses DNS.

DNS has been fundamental since the launch of the Internet. It has its place in every nook and cranny of your network, and is of business-critical importance. In fact, it's well known that DNS misconfigurations are perennially a chief culprit when something breaks on the network. Indeed, sometimes the network doesn't work because DNS isn't properly configured (the wrong box gets checked, a default should be changed, and so on).

At other times, however, applications and services may become unavailable because DNS is attacked directly, and an attack succeeds. This renders the DNS Service unavailable, which in turn blocks access to applications and services. In other cases, DNS can also be used as a vector for attack, enabling the DNS Service to be abused through malware, command and control, or data exfiltration attacks. In both cases, such exploits are an order of magnitude scarier, because they involve longer and more difficult recovery scenarios.

This book will help you understand how to handle DNS attacks better and how to use DNS itself as part of your zero trust security strategy. While DNS security remains critical, ideas about how best to secure this service continue to evolve. Traditionally, local users or devices on the network weren't automatically viewed with extreme suspicion. Often, in a legacy view of perimeter security, they would be granted access without any serious checks performed to verify identity, access rights, and so forth.

Those days are long gone. There are more bad guys and even criminal organizations out there than ever before. Indeed, organizations have necessarily adopted a more paranoid view that everyone is potentially an attacker, including insiders. This is a good thing, as attacks become more automated, and even industrialized, and avenues of attack become much more common.

DNS sits at the front door to your crown jewels—your data, users, services, and applications. Thus, DNS itself not only needs protecting now more than ever, it must be leveraged as a key component in the modern security arsenal. How do you even start, though? Glad you asked! Let the Gorilla guide you toward a safer DNS posture by following him into Chapter 1, which provides an overview of the jungle you're about to enter.

CHAPTER 1

What Is DNS? What Does It Do for You?

DNS is an acronym for the Domain Name System (used on the Internet and on private networks), which allows symbolic names such as XYZ.com or ESU.edu to be used to identify points of Internet presence. The Domain Name System in turn relies on the Domain Name Service (also abbreviated as DNS) to resolve those symbolic names into current, valid Internet IP addresses. Thus, for example, addresses 142.250.189.174 and 2607:f8bo:4005:80e::200e point to google.com, while addresses 217.70.186.105 and 2604:3400:dc1:950::6 point to efficientip.com.

Simply put, DNS lets people use short, straightforward, human-readable names to access their networks. Behind the scenes, DNS manages the translation between such names and various corresponding Internet Protocol (IP) addresses that uniquely identify and locate websites, Internet-facing physical and virtual servers, and so on. **Figure 1** shows the so-called root-level domains .edu and .com, where both google.com and efficientip.com reside inside the .com root.

The Origins of DNS

DNS first appeared in 1983, via the work of Dr. Paul Mockapetris at USC, in a prototype version known as Jeeves. This led to a pair of Internet specifications, RFC 882 and 883, that laid the groundwork for DNS in November of that year. In 1984, a quartet of UC Berkeley graduate students created the Berkeley Internet Name Domain (aka BIND) system,

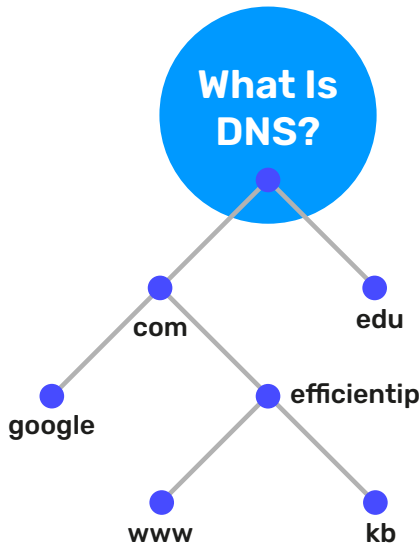


Figure 1: EfficientIP also supports `www.efficientip.com` (website) and `kb.efficientip.com` (knowledgebase: requires login credentials)

upon which all modern DNS implementations are based. By 1987, RFCs 1034 and 1035 defined a set of formal specifications for DNS, and still govern its formats and behaviors today (along with RFCs 1123, 2181, and 3403).

The domain name space consists of a tree data structure, like the one shown in **Figure 1**, where domain names are constructed by picking nodes from right to left, starting at least two nodes in from the left. As already noted, that makes `google.com` and `efficientip.com` the only valid domains on display, where `www.efficientip.com` and `kb.efficientip.com` represent logical partitions within the `efficientip.com` name space.

Behind the scenes, a hierarchy of DNS servers mirrors the top levels of the DNS name hierarchy. Across DNS servers, a distributed client/server database captures the fully fleshed-out tree of domain names around the globe. The nodes in this distributed database are “name servers,” where each domain has one or more authoritative DNS servers that publish information about that domain, plus any name servers for

delegated domains beneath it (e.g., for `efficientip.com`, Figure 1 shows that `www.efficientip.com` and `kb.efficientip.com` both qualify as such; both also inherit their name servers from the parent domain, namely, `ns-1000.awsdns-61.net`, based within Amazon Web Services, or AWS).



DNS offers incredible ease and convenience in using the Internet. The Domain Name System collectively defines a robust but distributed name hierarchy that ties the whole Internet together. It turns a loosely coupled and confederated collection of servers and information into a massive, single, logical, and global namespace. DNS scales well with hundreds of millions of domains and billions of IP addresses under its purview. The downside is that DNS's amazing ease of use and convenience also provides an enormous attack surface, because DNS must be everywhere to properly do its job.

In the next chapter, you'll learn how DNS can provide a focus for cyber-attacks of varying kinds and levels of severity. Please observe that as an old, well-established protocol, DNS is open by design and is not encrypted by default for backward-compatibility reasons. Thus, DNS is subject to easy attacks from cybercriminals who use it as both a target and a vector for attacks. These must be addressed and foiled, because doing without DNS is impossible.

CHAPTER 2

How DNS Can Enable Cyberattacks

There are several important reasons why DNS sits atop cybercriminals' lists of potential attack points. To begin with, DNS is a mission-critical network service: Almost all applications and services, users and customers, and so forth, must use the DNS service as a basic starting point for communication.

In fact, DNS plays a pivotal role in application traffic routing: It's what makes the links between users and internal or external applications possible. Further, without robust, secure, and agile DNS infrastructures, important IT services such as email, Internet access, or VOIP can't work correctly. In other words: No DNS, no business!

Next and perhaps more critically, DNS is easy to exploit. That's because DNS was designed from the start to be an open service. As a consequence of their fundamental role in IT infrastructures, DNS servers must be accessible to everyone.

Using its mandatory presence to enable address resolution and direct traffic, hackers make use of a dual DNS role in what's sometimes called the "cyber kill-chain." That is, DNS serves as both a threat vector, through malicious use of malware and the DNS protocol to communicate with attackers' remote command and control servers, and as a direct attack objective where, for example, a Denial of Service (DoS) attack on DNS servers can impact business continuity and viability.

Another issue is that the DNS protocol is connectionless, employing the User Datagram Protocol (UDP) by default. This makes it easier for hackers to launch attacks, because it doesn't require establishing a connection with a targeted device or address. Most firewalls cannot efficiently manage and maintain network security when UDP traffic handles DNS queries, replies, and other valid message types.

Eventually (and involuntarily), DNS is open to a wide variety of attacks, some quite sophisticated. These include DoS and distributed DoS (DDoS), phishing, zero-day exploit, data exfiltration, and many other well-documented attacks. Today's DNS threats are complex and multi-layered, where sophisticated threats combine multiple vectors and various types of communication in a single attack.

Finally, traditional security solutions don't offer effective DNS protection. Indeed, DNS isn't adequately protected by standard security systems such as next-generation firewalls, web proxies, IPS, and so forth because they're not purpose-built to handle and protect DNS. These systems don't understand the DNS protocol in any depth, nor do they implement advanced DNS security mechanisms. Three primary reasons explain this deficit:

1. These systems include no DNS analytics from which they can obtain real-time behavioral threat detection. In fact, their threat detection is based on limited security mechanisms, including signatures (which protect only against known threats) or basic analyses of DNS packet frequencies, request entropy, or payload contents.
2. These systems lack sufficient power to handle volumetric attacks. The majority of DNS DDoS attacks (88%) send more than 1 Gbps. Standard security solutions simply cannot handle that kind of traffic volume.
3. These systems implement only basic mitigation techniques, which are subject to a high risk of false positives. In fact, their countermeasures are limited to blocking all DNS traffic, not by filtering policies defined for individual clients.

3 Main Reasons DNS Is Targeted by Cybercriminals

1

Mission
Critical

2

Easy to
Exploit

3

Not
Efficiently
Protected

The DNS Security Landscape

Here are some key facts and figures about DNS in the context of cyber-security. They point to its continuing popularity as an attack vector and target.

- [87%](#) of companies were hit by DNS attacks in 2021. Thus, it remains a nearly constant inhabitant of the security landscape, and will remain a huge, juicy target that's constantly probed for weaknesses
- In fact, DNS remains a leading target for application layer attacks because organizations provide DNS servers that must respond to incoming queries, even malicious ones, unless proper screening and filtering is put in place
- In [2020](#), [85%](#) of malware used the DNS protocol in some way, typically during initial or early stages of reconnaissance and attack
- According to the CVE database, five meaningful BIND vulnerabilities appeared in 2021 (there were nine in 2020, and seven in 2019); as a primary software vehicle for providing DNS services, addressing BIND vulnerabilities is essential

Thus, protecting DNS, including managing access to its configuration and other meta-data, is an essential part of establishing and maintaining proper security.

Types of Attacks

[Methods of DNS attack](#) (**Figure 2**) might target DNS servers themselves, or they may use the Domain Name System (and its protocols and services). It's best to understand DNS server targeting as a form of direct attack, while DNS-based attacks are more indirect (and often, more subtle, and thus also harder to detect and fend off).

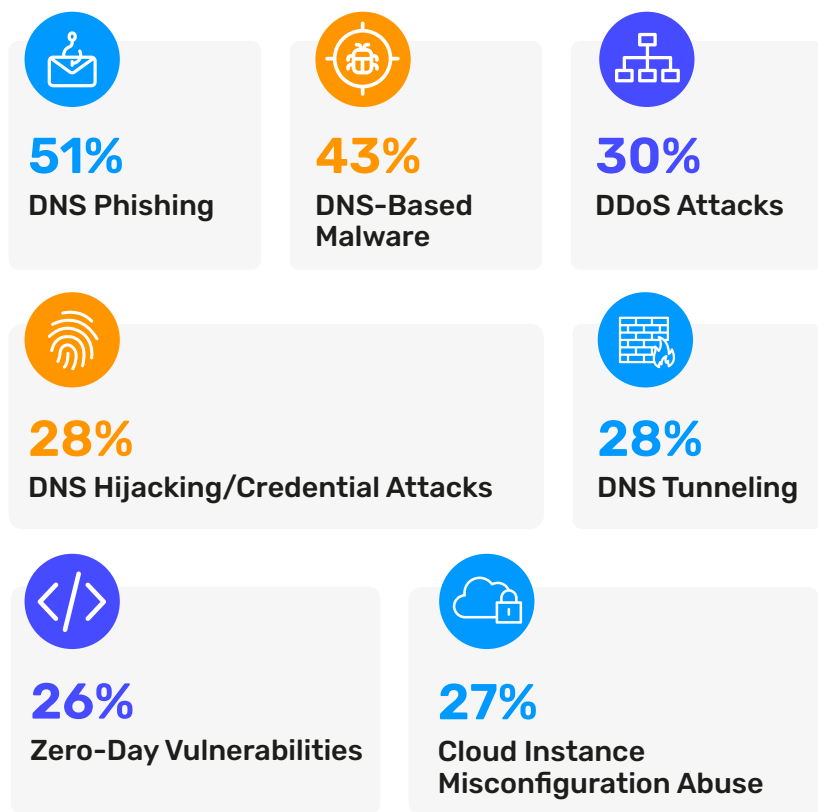


Figure 2: The IDC 2022 Global DNS Threat Report shows 7 different attacks as most common for the year (Source: IDC Global DNS Threat Report 2022)

Among other forms of direct attack, the following are well-known and documented in information security circles:

- **Volumetric attacks** attempt to overwhelm the DNS server with enormous numbers of requests from one source (DoS) or multiple sources (DDoS). The intent is to degrade the server's capabilities, or to outright deny access to its services.
- **Stealth attacks**, also called Slow Drip attacks, involve a constant, but low volume of specific, malicious DNS queries. These resemble a DoS or DDoS attack in that they consume query response capability in a DNS server, leading to reduced or no outgoing response volume. But they attempt to disguise their activity and intent through less obvious and intrusive use of malicious queries. Examples include Sloth Domain Attacks and Phantom Domain attacks.
- **Outright exploits** are simply attempts to use known bugs or flaws in DNS services and protocols, or the operating systems upon which DNS services may run. These are many and legion (a [search of the Common Vulnerabilities database with DNS as its focus](#) returns more than 10,000 results).

Where indirect attacks occur, DNS provides a method to inform or guide an attack, and plays a role in attempting to exploit DNS services and protocols as part of that attack, including:

- **DNS server/site spoofing** involves “man-in-the-middle” methods that intercept traffic and redirect it to malicious sites. They might also involve deliberate DNS configuration changes to explicitly redirect traffic to malicious sites behind the scenes.
- **Command and control** (see Figure 3) is a type of attack in which cybercriminals use a malicious external server to command and control already compromised machines over the internal network via DNS Tunneling. This command-and-control server also receives payloads sent from those compromised devices.

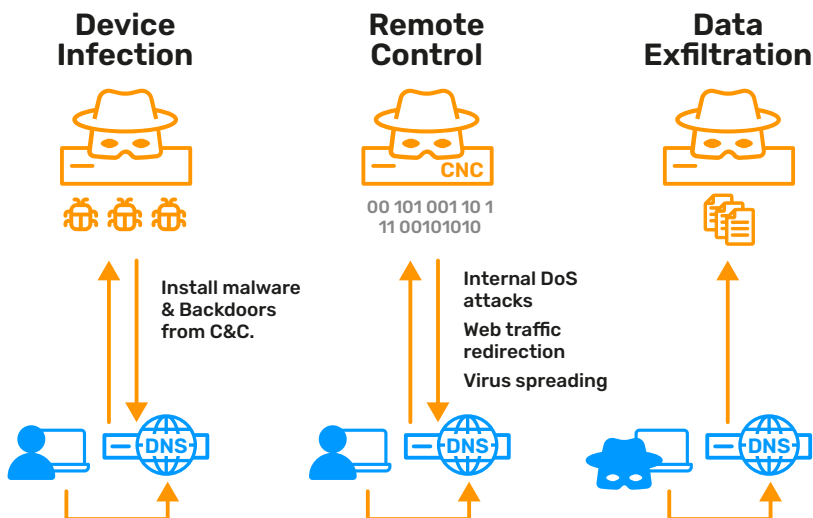


Figure 3: An illustration of command and control, as well as data exfiltration

- **Data exfiltration** attacks use DNS in malicious, unintended ways that expose an organization’s data or files to be accessed by unauthorized third parties (also depicted in **Figure 3**).



Because DNS is in such widespread and universal use, its presence—and in this context, its security—are too often taken for granted. Within an organization, security operations (SecOps) and network operations (NetOps) teams and their members must work together to protect, strengthen, and, ultimately, make proper use of DNS to boost their security posture. They must collaborate to foil attacks directed at DNS as their targets as well as those that use DNS as a vector.

In the next chapter, you’ll learn how the twin mantras of “trust no one” and “trust, but verify” combine to create the right level of constructive paranoia needed to make DNS security take its proper place in establishing and maintaining a correct security posture.

CHAPTER 3

Understanding ‘Zero Trust’

Zero trust means exactly what it says: By default, no one who requests access to systems and resources is trusted. All requests must be checked and verified, and will go forward only if explicitly allowed to do so. Let’s explain why such a tight policy needs to be in effect.

Perimeter Security Is Not Enough

The foundations of information security rest on the notion of a secure perimeter around an organization’s networks. This is a model that goes back to the mid-1990s (and earlier) as described in Cheswick and Bellovin’s classic 1994 book, “Firewalls and Internet Security: Repelling the Wily Hacker.” This approach introduces the concepts of internal and external networks, a “demilitarized zone” (DMZ) between them, and places firewalls as a crucial element between insiders and outsiders (especially those with bad intent and hacking skills).

These days, perimeter security no longer suffices to establish or maintain a proper security posture. Simple interposition between internal and external traffic, addresses, and flows offers no protection against internal threats. Even stateful inspection of packets between communicating parties offer insufficient protection to screen out bad actors and bad actions.

Today, behavior in which clients engage establishes what’s safe and acceptable, or what’s unsafe and unwanted. It also provides the basis for establishing user-based notions of normal and expected activity and access, as a baseline against which current requests and activities can

be assessed and evaluated. Simply put, tracking ports and protocols no longer provides enough data to define, establish, and maintain security.

The Principles of Zero Trust

Zero trust, or Zero Trust Network Access (aka ZTNA), is emerging as a cornerstone of modern information security best practices. Zero trust is based on two fundamental principles:

1. Deny all access by default.
2. Only allow access when trust can be safely and properly established.

Thus, the starting assumption for ZTNA may be succinctly stated as: “[Never Trust, Always Verify.](#)”

In its purest form, zero trust is a network security model based on a strict and always-enforced identity verification process. This means that only authenticated users and devices can access applications and data.

It’s important to note that ZTNA means that ordinary users must be authenticated to access the network, along with those requesting higher levels of access (such as administrators, operators, technicians, and so on). As such, zero trust is a key ingredient for modern digital transformation. ZTNA is also the cornerstone upon which safe and secure business network architectures and approaches can rest. **Figure 4** depicts the elements that should be involved in any sound and secure zero trust approach.

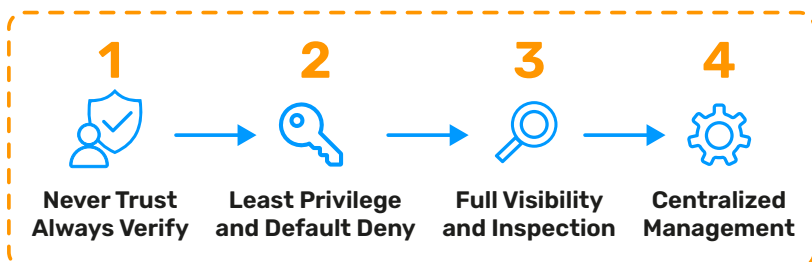


Figure 4: Beyond the principle of “never trust, always verify” Zero Trust/ ZTNA goes considerably further

Figure 4 shows all the basic principles of zero trust in their proper order, read from left to right:

1. **Never Trust, Always Verify.** This is the foundation on which ZTNA rests.
2. **Least Privilege and Default Deny.** Also known as the “Principle of Least Privilege,” or PLP, this approach never extends more rights or privileges than needed. By default, all requests for access are denied. They’ll only be granted if and when valid authorization applies to a validated identity.
3. **Full Visibility and Inspection.** This means that access is audited and monitored, especially when high levels of privilege are invoked. Visibility provides the basis for accountability.
4. **Centralized Management.** All the other zero trust principles operate in the open, under explicit management and control as per organizational policy, governance, and compliance requirements.

Zero Trust Facts and Figures

Zero trust is more than just a good idea—it’s becoming the cornerstone for best security practices, digital transformation, and safe and secure business network architectures and approaches, which is why [76%](#) of organizations have implemented or plan to implement ZTNA regimes. They’re moving this way for several reasons, including:

- Internal threats are powerful and dangerous; zero trust protects equally against all threats, both internal and external. According to TechJury, 60% of security breaches involve an insider when data theft or exfiltration occurs.
- Distributed network topologies and multi-cloud use make security more complex; ZTNA provides a sound basis for securing complex environments.

- Macro-segmentation of networks—e.g., intranet, extranet, and DMZ—do not protect against internal threats. Zero trust reduces the attack surface, and provides added and vital protection against unauthorized or unwanted internal activities.
- Security and network operations and the data that govern them are too often separate and disjointed. ZTNA provides a consistent approach across the whole organization and beyond.
- Zero trust helps overcome typical functional disconnects within organizations, where NetOps, CloudOps, and DevOps teams don't communicate enough (or at all) with SecOps.

The Zero Trust Perspective

Zero trust's increasing pervasiveness and adoption speak well of its importance and value in establishing and maintaining security, especially in today's complex, cloud-forward IT environments. It's safest to assume nothing, to trust no one, and to authenticate and verify all requests for access, irrespective of origination or location.

Proper zero trust architecture is segmented, parallelized, and centralized. Network domains and services should be segmented for logical separation and improved security. Centralized management via a single console provides a global overview of all components and domains, and is more easily managed for security, compliance, and proper governance.

Such architectures are best built using multiple, parallelized switching cores to support highly distributed and virtualized operations around software-defined architectures and networking. They can make use of extensive automation to handle setting up, provisioning, rebalancing and repositioning, and rapid response to emerging security threats or attacks.



According to Forrester Research, “[Zero Trust Is Not a Security Solution; It’s a Strategy](#).” This means that zero trust is neither a product nor a platform. Rather, ZTNA is best understood as a security framework built around the default concepts of “never trust, always verify,” and also, “assume a breach has occurred.” To that end, Forrester has defined a zero trust eXtended (Zero TrustX) ecosystem that involves assessing zero trust maturity across an organization’s people, skills, technology, capabilities, and so on. Next, involved teams evaluate maturity levels to identify areas of relative strength and weakness, especially when certain capabilities need improvement. Such an exercise also involves considering tools and technologies to address weaknesses and consistent implementation of ZTNA principles across the organization’s business, IT, and security projects.

Zero Trust Architecture Building Blocks

Micro-segmentation is a major component of the ZTNA architecture. It requires some adjustments to the classical port and VLAN-based segment, allowed by the separation of the data plane and the control plane. Micro segments can allow or deny certain communications even within the same VLAN, helping devices to be separated from each other and enforcing communications to go through security devices like firewalls, proxy, and WAF. Combined with Network Access Control (NAC) through Identity Management, it provides security at the network level that helps other zero trust blocks to be deployed.

This approach benefits from the deep knowledge of users and applications that comes thanks to integrated, across-the-board Identity and Access Management (IAM). In fact, IAM enables filtering of network traffic flows at a functional level. This supports a standard, network-wide approach to security. It operates not just at the network boundary, or between DNS zones, but everywhere on the network.

Within the zero trust architecture, data acquisition provides statistics, logs (subject to automated analysis), and telemetry. These combine to create an ongoing, continuous view of network status and conditions in real time.

Events may be linked to Security Information and Event Management (SIEM) correlation tools, and include User Behavioral Analysis (UBA), so that automation can drive rapid response to changing conditions, security threats, or abnormal behaviors. This explains the emphasis on provisioning access controls for all users based on automation, tools, and software-defined networking.

In the next chapter, you'll learn how DNS can (and most definitely should) play an important role in establishing and maintaining ZTNA. Because DNS is involved in almost every Internet communication, securing its use provides unique opportunities to boost security, observe usage behavior, and filter or block unwanted actors and actions.

CHAPTER 4

How DNS Strengthens Zero Trust to Better Protect Apps, Users, and Data

Given that DNS is bound to be part and parcel of any organization's IT infrastructure, it's essential to protect it from attack. And be assured that those attacks are coming—its protocols and queries offer a common attack vector. But at the same time, if DNS is implemented correctly and based on a purpose-built technology, it can provide useful and important information, and act as an active security component as part of a zero trust framework.

Setting the Stage

DNS is well aligned with the basic principles of the zero trust framework previously described. First, as an entry point into an organization's networks, DNS is ideally positioned to become the first checkpoint to apply the “never trust, always verify” principle. Therefore, all DNS client requests can be denied by default, and DNS resolution is carried out only if a validated identity is authorized to proceed. Note that user identities include trusted services, OS processes, applications, and other identities that can request and consume computing resources and facilities. Unfortunately, the DNS protocol has not been designed to make identity checks, nor to validate that a specific client is authorized to reach specific domains. Such capabilities aren't built into DNS.

However, a proper DNS solution can filter on domain names and IP addresses, providing granular protection and control using a feature

called Domain Name Service Response Policy Zones (DNS RPZ). DNS RPZ is a mechanism implemented in all modern recursive DNS engines that allows dynamic modification of DNS answers and provides alternate answers to any DNS query. This mechanism permits DNS administrators to define precise filtering and redirection of rules applied to DNS traffic according to the queried domain name, NameServer (NS) or IP address.



Secure DNS works from the level of network micro-segmentation tied to a single, unique IP address and associated user ID. What may seem complex at the overall network level becomes simple and straightforward through the DNS lens regarding per-user ID behavior and activity.

DNS RPZ can be leveraged to control DNS traffic to prevent any connection to known malicious services used to steal credentials or deliver initial infectious malware payloads. However, maintaining appropriate filtering rules related to known malicious domains is difficult because this type of threat is dynamic and unpredictable. Attackers use several, often randomly generated domains (using DGAs, aka “Domain Generation Algorithms”) to control their botnets and leverage huge numbers of poorly secured servers to run their activities. Consequently, using a dynamically updated filtering rule repository that can be extended through a customized filtering policy provides the most sustainable solution.

On the basic principle of “full visibility and inspection,” incoming DNS requests also offer key insights and intelligence at the individual client level. If DNS is suitably monitored and its data properly collected, it can provide valuable information about resources requested by which clients. DNS has the data to see normal traffic patterns, and the locations, applications, and services clients normally visit. Any anomalies that deviate from this baseline must start with out-of-range or unusual

DNS requests. Thus, DNS provides perfect visibility into behavior and access via network traffic for each client, resource, and server. That's why the right DNS solution can offer early anomaly detection and recognition, and spur speedier, better-informed security responses.

Security Challenges

The real challenge is to perform real-time analysis of overall traffic at the client level. Indeed, standard DNS servers are simply not designed to deliver such capability. Thus, traditional DNS security solutions usually base their detection algorithms on DNS packet frequency, payload, data encoding, or request entropy. They have no insight into DNS traffic

Understanding DDI

DDI stands for DNS, DHCP and IPAM (DDI). Here's what it means:

- **DNS:** the Domain Name System (and Domain Name Service), upon which access to enterprise apps and the Internet is based, via human-readable names.
- **DHCP:** the Dynamic Host Configuration Protocol, an automated service that assigns addresses and other data (including DNS server addresses) to network devices from a pool of managed addresses. Essentially, DHCP lets users log onto and use the network without having to manage all the name and address details for themselves.
- **IPAM:** IP address management is a set of services and a centralized IP data repository that contains important metadata used to manage IP addresses, support SD-WAN, avoid misconfigurations, and support extensive automation.

DDI provides a powerful combination of capabilities that inherently (and explicitly) supports a zero trust approach to DNS name and address management and networking.



at the heart of the DNS resolver engine, which requires monitoring requests at both cache and recursive levels. There's no deep understanding of overall DNS traffic between each specific client within a network and internal or external DNS servers.

This external analysis of DNS protocol can never provide enough information to accurately detect malicious traffic. It also allows DNS to easily serve as a threat vector. Passing DNS logs to a SIEM solution is impractical for such a workload, too. That's because it needs to store, index, and analyze large volumes of data, ideally in-memory. Additionally, analysis of such logs takes hours, and only provides post-mortem analyses. As a result, DNS services are not protected as quickly or completely as they should be. That's why a comprehensive solution to DNS security goes beyond the limitations of SIEM and works within the DNS runtime context itself, through secure, purpose-built DDI environments.

In the next chapter, we explain how DDI can provide signals and guidance to better protect applications, as well as their users and data.

The 5 Pillars of EfficientIP 360° DNS Security for Enabling DNS-Based Zero Trust

EfficientIP provides a holistic solution that protects public and private DNS from both internal and external threats and exposures. The EfficientIP answer rests on five key capabilities within its offerings, collectively known as the “5 Pillars.” Taken together, they deliver unique security capabilities to support your zero trust strategy.

1. Client-Based Application Access Control Powered by DNS Client Query Filtering

Within an organization, application access works at various hierarchical levels—org, OU, group, role, and so on—as dictated by organizational security policy. Certain applications (usually mission-critical ones) require special access and run on a dedicated infrastructure, with no sharing for main components. Filtering at the network level often hinges on Access Control Lists (ACLs) and firewall rules. Adding filtering at the DNS level helps organizations raise their security bar still higher. This eliminates the possibility of resolving an application’s IP address, except for clients explicitly allowed to do so. That makes DNS security a great approach for zero trust environments.

EfficientIP brings granular network segmentation functionality into the mix through its DNS Client Query Filtering (CQF) feature. NetOps and SecOps teams can dynamically update DNS Client Query Filtering lists with application or client entries (Allow and deny policy rules). This ensures security is raised dynamically and automatically as needed, and application zoning restricts exposure and data visibility for unknown or unauthorized users.

2. Unequaled DNS Analytics for Behavioral Threat Detection at the Client Level

EfficientIP's protective DNS solution delivers built-in security to cache, recursive and authoritative DNS servers. It offers complete and real-time DNS Transaction Inspection (DTI), enabling in-depth understanding of the context for client requests. By analyzing transactions at the heart of the DNS server (queries, responses, fragments, recursions), threat visibility is enhanced well beyond known attack patterns and overcomes the limitations of signature-based protection. By requiring all DNS requests to be valid and authentic, and filtering out known and probable bad actors and actions, behavioral threat detection prevents unauthorized and unwanted access to systems and resources. It works for both internal and external clients, and enables predictive security to protect against the most advanced attacks such as zero-day malicious domains, data exfiltration, and DNS tunneling.

3. Proactive Security Powered by Threat Intelligence

DNS Threat Intelligence delivers dynamic feed data containing a list of malicious domain names built from various distributed sources (internal or external). It aggregates reports of suspicious activity from identified IP addresses or domains (such as abuse and spam, phishing, malware,

cracked websites), enabling it to proactively block any attempts to connect to known malicious destinations. It prevents initial malware infection and communications with known command and control servers to avoid malware updates, remote control, and most data exfiltration attempts.

New DNS-based threats identified from real-time behavioral client traffic inspection enable building of internal threat intelligence services (specific to the context of your organization) that can be immediately shared across the entire DNS infrastructure, strengthening your security posture.

4. Adaptive Countermeasures to Ensure Service Continuity

EfficientIP security protections go beyond traditional security systems and are not only limited to blocking DNS traffic altogether for everyone, but delivers graduated and adaptive countermeasures according to the threat analysis and defined policies. It provides intelligent DNS protection to prevent, contain and block attacks while mitigating the risks of false positives. Malicious requests are blocked or abusive client's traffic is rate limited while suspicious compromised clients can be isolated and put in quarantine mode. A client in quarantine has only access to DNS cache service and not to the recursive function, eliminating the risk of any communication with an external accomplice command and control server. Adaptive countermeasures can be applied on a per client basis to ensure highest security response.

The Rescue Mode ensures Cache Service continuity even under unidentified attack sources. In rescue mode all cache entries are frozen, and the recursive function is delivered in "best-effort" mode to ensure service continuity.

5. Automated Sharing of Actionable Events and Data for Accelerated and Efficient Remediation

DDI services deliver unique contextual information about network activity and offer a “Network Source of Truth” (NSoT) that can be leveraged by security systems to deliver advanced automated protection capacity.

EfficientIP DNS, DHCP and IPAM solutions can provide Security Orchestration Automation Response solution (SOAR) and SIEM with key network and security event information in real-time such as IP addresses

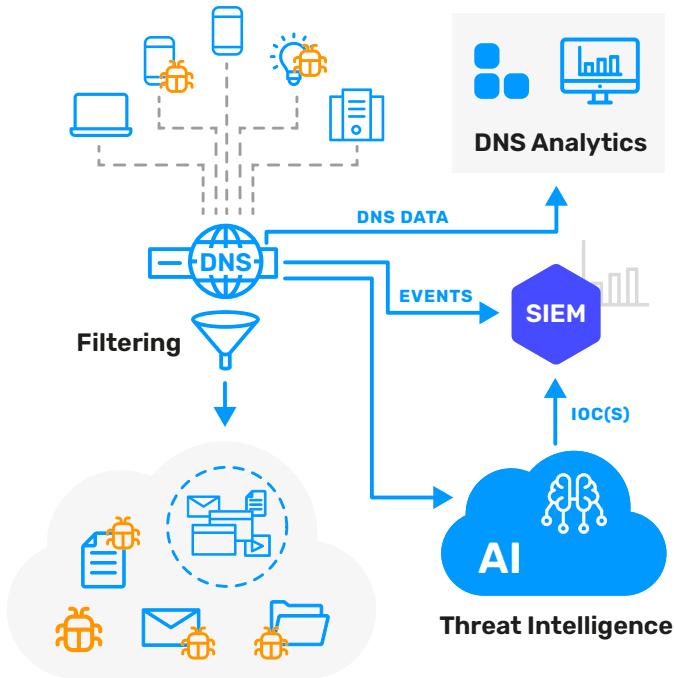


Figure 5: By bringing DNS data and events into SIEM, along with AI-driven threat intelligence, companies gain greater security and improved security and data protection

of an infected device by a DNS-based malware or a newly detected malicious domain name. Consequently, the SIEM can better correlate the qualified security event for advanced threat impact analysis, while the SOAR platform can use that information to facilitate and accelerate the security alert remediation at the user level (see **Figure 5**).

In the final chapter, we'll look at several secure DNS use cases and a case study, showing its effectiveness in action.

CHAPTER 6

Secure DNS Use Cases and a Real-World Example

A key value from using DNS to inspect and apply security comes in the arena of application access control. This approach takes network segmentation to the individual user level, so that each user's activity, behavior, and access stands out on its own (see **Figure 6**).

This approach confers numerous benefits and capabilities to organizations that adopt the right kinds of DNS security solutions, including:

- More granular filtering
- Better application access control
- Earlier security barrier
- New business opportunities
- Stronger security ecosystem

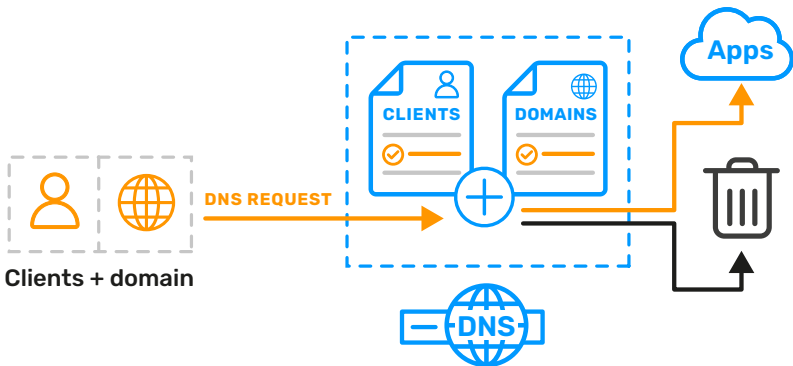


Figure 6: The filtering process in EfficientIP's DNS Client Query Filtering (CQF)

- Nothing to deploy on the network
- Lightweight and flexible implementation
- Low cost without performance bottlenecks

DNS Security Use Cases

Client-Based Application Access Control

Client-Based Application Access Control confers the ability to use data beyond domain names and associated IP addresses (such as the extended DNS Client Subnet field, CPE ID, MAC address, and more) to provide more detailed and descriptive client information and intelligence.

As a technical enabler of Client-Based Application Access Control, DNS Client Query Filtering brings greater power to DNS filtering. Security can be based on source client information mapped to the requested domain, rather than filtering based solely on the domain, as is typical for other DNS security solutions.

A specific filtering policy can therefore be applied just to specific clients requesting access to specific applications. This filtering policy can be an explicit Allow or Deny, establishing specific, fine-grained access policies between clients and applications. This strengthens DNS security by combining client and destination information with Allow and Deny lists, enabling application security enhancement at the earliest point in the security chain.

This may be used effectively along with cascaded DNS servers, DNS over HTTP external engines, and even an ISP's DNS relays. Thus, it supports complex telecom and service provider networks, including those in the cloud. It provides extensive DNS filtering, blocking, baseline data, and more.

IoT Security

Another use case concerns the Internet of Things (IoT). With input and output targets both restricted and well-defined on IoT networks, DNS security provides a mechanism well suited to lock down traffic to and from IoT devices (see **Figure 7**). Outsiders can neither see nor access IoT nodes, and node communication is prohibited outside the scope of allowed server, data collection, and other valid addresses.

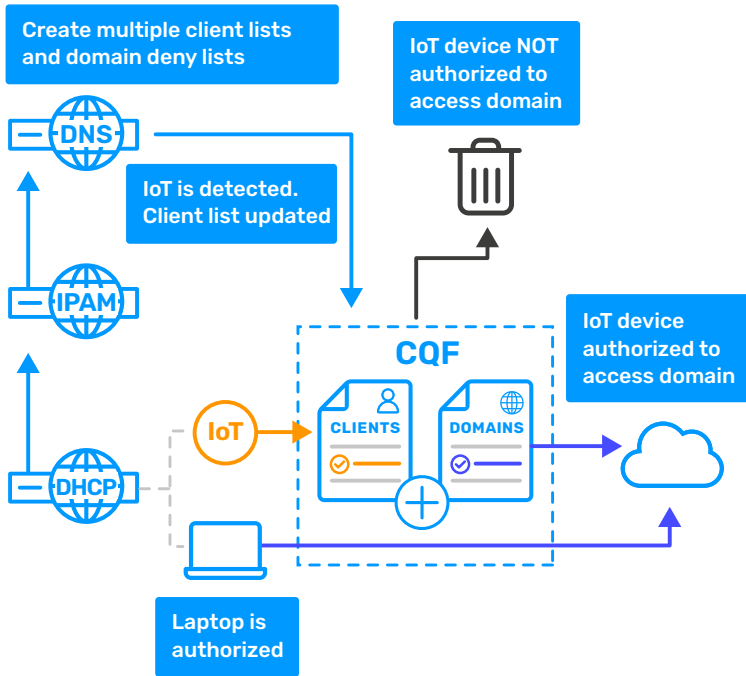


Figure 7: EfficientIP's DNS Client Query Filtering (CQF) applied to IoT

Cloud Services Access Control

The cloud provides yet one more excellent use case for DNS security. In the cloud, DNS Client Query Filtering serves to limit the users and services (and through them, the APIs) that can access cloud-based applications (see **Figure 8**).

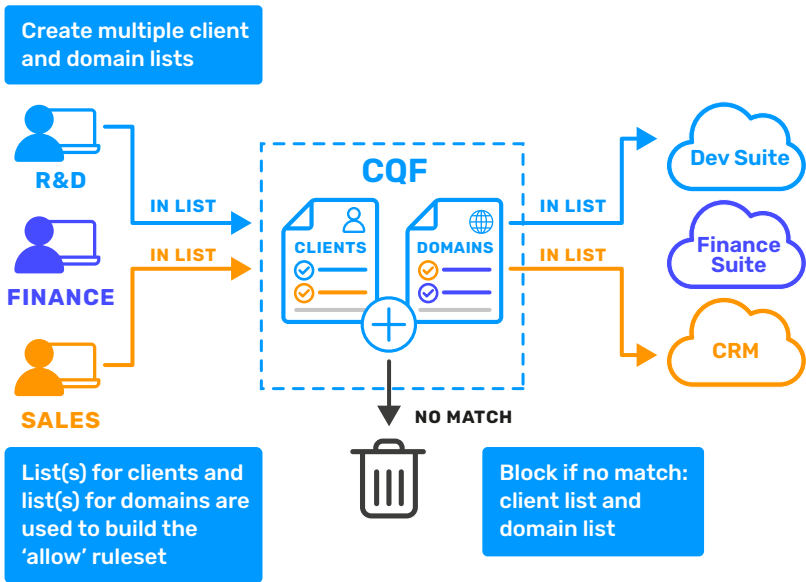


Figure 8: EfficientIP's DNS Client Query Filtering (CQF) for Cloud Services Access Control

In the same way, DNS Client Query Filtering can also limit outgoing communication from applications to only users and services authorized to receive such data. That means that DNS Client Query Filtering helps restrict traffic to cloud-based applications, both inbound and outbound, to only authorized clients. Dynamic Application Access Control over related Clients and Domains lists means that this pool of clients and domains can change over time, yet remain current and accurate.

Parental Control for Telcos

Parental control is another case that calls for DNS security. Service providers and telecom companies can grant parents the ability to assign a variety of restrictions to younger family members, including access hours, off-limits content designations, allowed applications, and more. This provides a way for families to manage access to technology and services, based on the identity of (and restrictions associated with) individual family members.

Case Study: STMicroelectronics

STMicroelectronics (ST) is a large, French-Italian multinational electronics and semiconductor maker headquartered near Geneva, Switzerland.¹ Recently the company used EfficientIP's DDI management solutions to improve its collaboration capabilities along with its security posture. The company adopted EfficientIP's SOLIDserver DDI to unify its core network services around its key constituent elements—DNS, DHCP, and IPAM. The company's DHCP managed hundreds of thousands of IP addresses across both its many manufacturing networks and globally distributed enterprise networks.

ST's networking manager for its manufacturing networks said that EfficientIP's solution “was most scalable for DDI.” In late 2019, ST launched a global initiative to address security risk. To that end, it eliminated manual rules management in the zone-based firewalls it had used to manage traffic between subnets inside the company network.

The company wanted zone-based security with its firewalls, which required classifying every subnet in its manufacturing network and establishing rules for zone communication.

According to the report, integration and automation of rules permitted network administrators to classify each new subnet using EfficientIP's IPAM capabilities. Such classifications automatically push new rules to the company's firewalls, and make sure all zone-based rulesets remain current and accurate. The integration, completed in 2021, eased the burden of knowledge transfer from the network team to the security team.

DNS security integration had additional uses as well. ST says it will enable micro-segmentation projects, like network access control. This integration also gives its teams a foundation for working together in the future, not just on security, but for network automation.

¹ Based on an EMA research report study from Shamus McGillicuddy: “[NetSecOps: Aligning Networking and Security Teams to Ensure Digital Transformation](#),” October 2021.

Recently, ST has also implemented EfficientIP's DNS Guardian to help secure cloud migration. As you've seen, DNS serves as a common attack vector for malicious attacks of all kinds, but cloud-based enterprise services face constant attack via DNS spoofing, and other paths as described in Chapter 2. ST is moving part of its IT to the cloud, and needs extra security threat protection and data leak prevention. ST praises EfficientIP's high level of quality processes, responsiveness to requests, and desire to continuously improve.

Protecting DNS Has Never Been More Important

In this guide, you've learned about the central role that DNS plays in every infrastructure. As any IT admin can tell you, when DNS doesn't work, pretty much nothing in the environment works.

This importance is also, of course as seen in Chapter 2, the reason that hackers and cybercriminals are constantly trying to compromise your DNS system. It's a gold mine of information that can bring them great rewards at little expense for enterprises not properly protected, and therefore is their primary target.

That's why the zero trust mindset is crucial for organizations to adopt. You must start with the idea that *everyone* is a potential enemy of your infrastructure, and only let in those who have proven that they can be trusted. In the modern era, to do anything else is to invite disaster.

If this has convinced you to take your DNS security more seriously, consider EfficientIP and what its solutions can do for you. The company offers peace of mind that's hard to put a price on.

Your next step, then, should be to contact [EfficientIP](#), and [request a free trial](#) and a [DNS Risk Assessment](#) that will identify Vulnerabilities: Expert Assessment of Your DNS Traffic. You may be surprised at what you find, and how vulnerable your network really is.

ABOUT EFFICIENTIP



EfficientIP is a network security and automation company, specializing in DNS-DHCP-IPAM (DDI). We promote business continuity by making your IP infrastructure foundation reliable, agile, and secure.

Since 2004, we have continued to expand our reach, providing solutions, professional services, and support all over the world with the help of select business partners. Our passionate teams have delivered successful projects to over 1,000 customers globally, and ensured operational efficiency through dedicated customer care.

Our goal is to enable secure and dynamic IP communication between users and apps/services. We achieve this by:

- Securing DNS services to protect users, apps, and data and ensure service continuity
- Simplifying lifecycle management of DDI resources, via smart automation, cross-platform visibility, and policy control through a single pane of glass

Companies rely on us to help control the risks and reduce the complexity of challenges they face. This applies particularly to modern key IT initiatives such as cloud applications, virtualization, mobility, digital transformation, and SDN. For more information, visit www.efficientip.com and follow [@efficientip](https://twitter.com/efficientip) on Twitter.


ABOUT ACTUALTECH MEDIA



ActualTech Media is a B2B tech marketing company that connects enterprise IT vendors with IT buyers through innovative lead generation programs and compelling custom content services.

ActualTech Media's team speaks to the enterprise IT audience because we've been the enterprise IT audience.

Our leadership team is stacked with former CIOs, IT managers, architects, subject matter experts and marketing professionals that help our clients spend less time explaining what their technology does and more time creating strategies that drive results.



If you're an IT marketer and you'd like your own custom Gorilla Guide® title for your company, please visit <https://www.gorilla.guide/custom-solutions/>