

The background features a stylized globe with a network of glowing blue and green lines connecting various points, suggesting a global network or data flow. The globe is semi-transparent, showing the underlying network structure.

CS 419: Computer Security

# Week 12: Network Security


## Firewalls

Paul Krzyzanowski

© 2025 Paul Krzyzanowski. No part of this content may be reproduced or reposted in whole or in part in any manner without the permission of the copyright owner.

# Network Security Goals

- **Confidentiality**: sensitive data & systems not accessible
- **Integrity**: data not modified during transmission
- **Availability**: systems should remain accessible



Dragon artwork by Jim Nelson. © 2012 Paizo Publishing, LLC. Used with permission.

# Firewall

- **Separate your local network from the Internet**
  - Protect the border between trusted internal networks and the untrusted Internet
- **Approaches**
  - Network Address Translation (NAT)
  - Packet filters
  - Application proxies
  - Intrusion detection / intrusion protection systems

# Network Address Translation (NAT)

# NAT: Network Address Translation

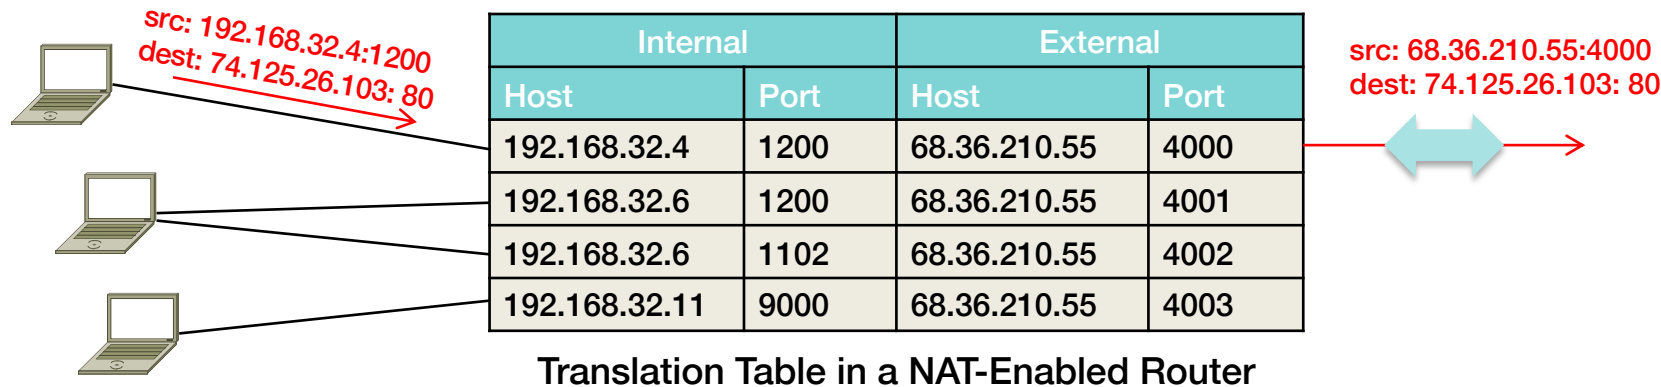
**NAT converts between  
private (internal) IP addresses and  
one or more public-facing (external) addresses**

# Running out of IP addresses

- **Every device on the Internet needs an IP address**
  - Every address must be unique  
... otherwise, how do you address a host?
- **IP addresses are not plentiful (with IPv4)**
  - Does an organization with 10,000 IP hosts really need 10,000 globally-unique addresses?
  - Prior to 1993, the answer was “YES!”

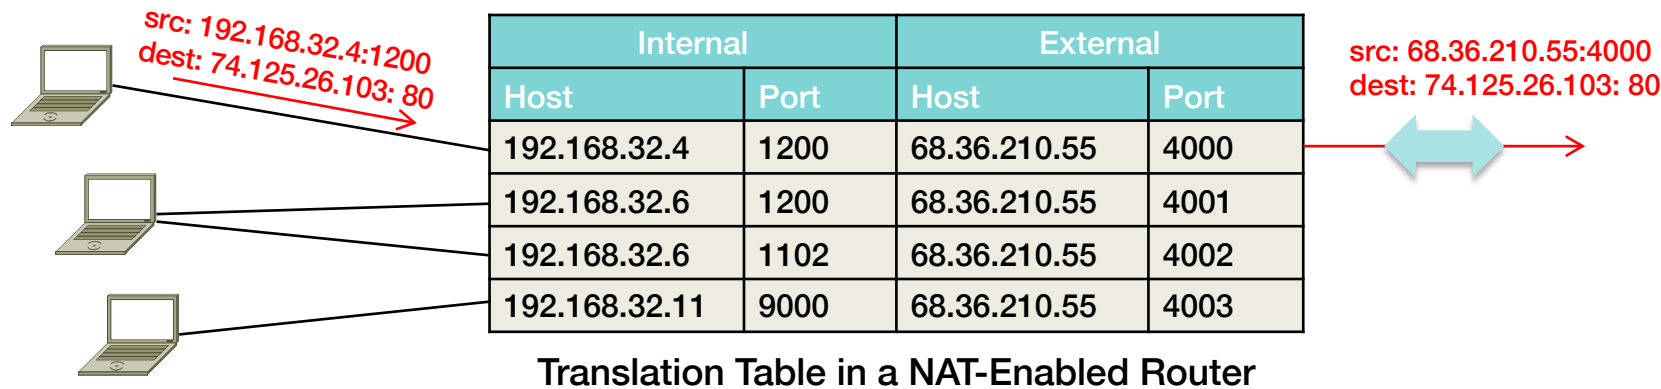
# IP Address Translation

- **Private IP address space in the organization**
  - One external IP address, multiple internal addresses
- **NAT – Translation Table**
  - Map source address:port in outgoing IP requests to a unique external address:port
  - Inverse mapping for incoming requests
- **A NAT-enabled router looks like a single device with one IP address**



# IP Address Translation

- **NAT requires a router to look at the transport layer**
  - Source port (outgoing) & destination port (incoming) changes
  - TCP/UDP checksum recomputed





# Private Addresses

- We cannot use IP addresses of valid external hosts locally  
... how will we distinguish local vs. external hosts?
- **RFC 1918: Address Allocation for Private Internets**
  - Defines unregistered, non-routable addresses for internal networks

Address Range	# addresses	IP address block
10.0.0.0 – 10.255.255.255	16,777,216	10.0.0.0/8
172.16.0.0 – 172.31.255.255	1,048,576	172.16.0.0/12
192.168.0.0 – 192.168.255.255	65,536	192.168.0.0/16

# Advantages of NAT

- **Internal address space can be much larger than the addresses allocated by the ISP**
- **No need to change internal addresses if ISP changes your address**
- **Enhanced security**
  - A computer on an external network cannot contact an internal computer ... unless the internal computer initiated the communication  
Even then – it can only contact the computer on that specific port

# First-Generation Firewalls: Packet Filters

# Screening router

## Border router (gateway router)

- Router between the internal network(s) and external network(s)
- Any traffic between internal & external networks passes through the border router

*Instead of just routing the packet, decide whether to route it*

## Screening router = Packet filter

**Allow or deny packets based on**

- Incoming & outgoing interfaces
- Source & destination IP addresses
- Protocol (e.g., TCP, UDP, ICMP, IGMP, RSVP, etc.)
- Source & destination TCP/UDP ports, ICMP command

# Filter chaining

An IP packet entering a router is matched against a set of rules:  
**access control list (ACL)** or **chain**

Each rule contains **criteria** and an **action**

- **Criteria**: packet screening rule
- **Actions**
  - **Accept** – and stop processing additional rules
  - **Drop** – discard the packet and stop processing additional rules
  - **Reject** – and send an error to the sender (ICMP Destination Unreachable)

Also

- **Route** – reroute packets
- **Nat** – perform network address translation
- **Log** – record the activity

# Filter structure is vendor specific

- **Windows: *Allow, Block***
  - Options such as
    - Discard all traffic except packets allowed by filters (*default deny*)
    - Pass through all traffic except packets prohibited by filters (*default allow*)
- **OpenBSD: *Pass (allow), Block***
- **Linux nftables (netfilter)**
  - Chain types: ***filter, route, nat***
  - Chain control
    - ***Return*** – stop traversing a chain
    - ***Jump*** – jump to another chain (***goto*** = same but no return)

# Network Ingress Filtering: incoming packets

## Basic firewalling principle

No direct inbound connections external systems (Internet) to any internal host – all traffic must flow through a firewall and be inspected

1. Determine which services you want to expose to the Internet
2. Allow only those inbound ports and protocols to the machines hosting the services
  - E.g., Web server: 10.0.0.10 TCP port 80, TCP port 443
  - Mail server: 10.0.0.12 TCP port 587

## Default Deny model – by default, *deny all*

- Anything not specifically permitted is dropped
- May want to log denials to identify who is attempting access

# Network Ingress Filtering (inbound)

- **Disallow IP source address spoofing**
  - Restrict forged traffic (RFC 2827)
- **Disallow incoming/outgoing traffic from private, non-routable IP addresses**
  - Helps with **DDoS attacks** such as SYN flooding from lots of invalid addresses
- **At the ISP**
  - Filter upstream traffic - prohibit an attacker from sending traffic from forged IP addresses
  - Attacker must use a valid, reachable source address

				address	mask	port	
access-list	199	deny	ip	192.168.0.0	0.0.255.255	any	log
access-list	199	deny	ip	224.0.0.0	0.0.0.255	any	log
				...			
access-list	199	permit	ip	any	any		



# Network Egress Filtering (outbound)

- **We don't usually worry about outbound traffic**
  - *Communication from a higher security network (internal) to a lower security network (Internet) is usually fine*
- **Why might we want to restrict it?**
  - Consider: if a computer is compromised & all outbound traffic is allowed, it can connect to an external server and download more malicious code  
... or launch a DoS attack on the internal network
  - Also, log which servers are trying to access external addresses

# Second-Generation Firewalls: Stateful Packet Inspection (SPI)

# Stateful Inspection – 2<sup>nd</sup> generation firewalls

**Retain state information about a stream of related packets**

## Examples

### – TCP connection tracking

- Disallow TCP data packets unless a connection is set up
- Allow return traffic

### – ICMP echo-reply

- Allow ICMP echo-reply only if a corresponding echo request was sent.


### – Related traffic

- Identify & allow traffic that is related to a request or connection
- Example: related ports in FTP
  - Client connects to server on port 21 to send commands
  - Server connects back to client on port 20 to send data

# Security Zones

- **Packet-filtering firewalls (almost always) live in routers**
- **Routers connect network zones together**
- **Firewalls allow you to control traffic between zones**

# Security Zones: DMZ




## Bastion hosts

Systems critical to network security. They will be carefully configured, run only essential software, have only the required user accounts, and be regularly audited.

Dragon artwork by Jim Nelson. © 2012 Paizo Publishing, LLC. Used with permission.

# Security Zones: DMZ




## Clients from the Internet:

- Can access allowed services in the DMZ
- Cannot access internal hosts

## The firewall:

- Blocks impersonated packets

# Security Zones: DMZ




## Clients in the internal subnet:

- Can access the Internet
- Can access allowed services in the DMZ
- May access extra services in the DMZ

Dragon artwork by Jim Nelson. © 2012 Paizo Publishing, LLC. Used with permission.

# Security Zones: DMZ



## Clients in the DMZ:

- Can access Internet services only to the extent required
- Can access internal services only to the extent required

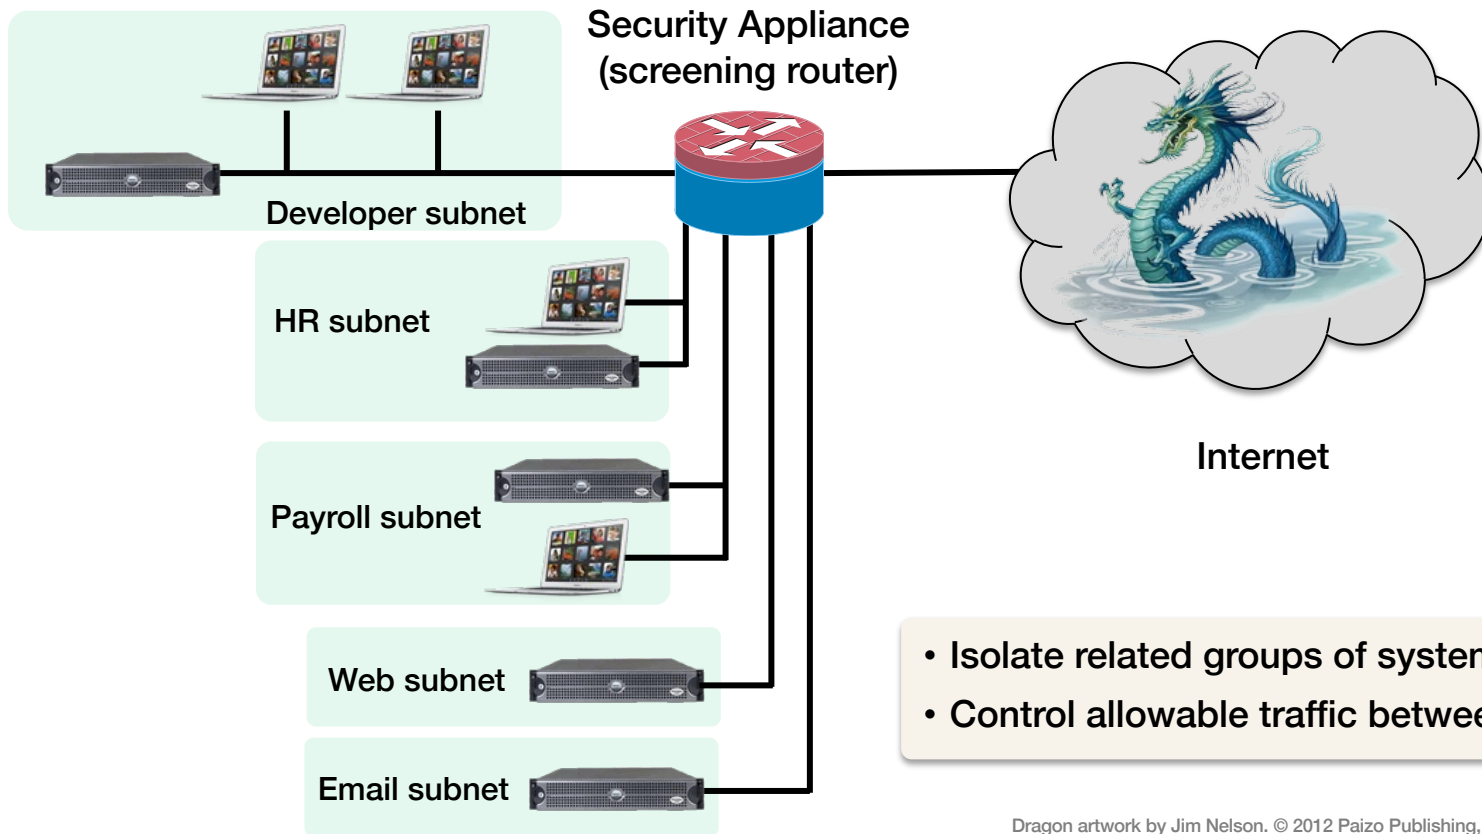
## Goal:

*Limit possible damage if DMZ machines are compromised*

Dragon artwork by Jim Nelson. © 2012 Paizo Publishing, LLC. Used with permission.



# Security Zones: Segmentation



Dragon artwork by Jim Nelson. © 2012 Paizo Publishing, LLC. Used with permission.

# Third-Generation Firewalls: Deep-Packet Inspection (DPI)

# Application-Layer Filtering

**Firewalls don't work well when everything is a web service**

## Deep packet inspection (DPI)

- Look beyond layer 3 & 4 headers
- Need to know something about application protocols & formats

## Examples

- **URL filtering**
  - Normal source/destination host/port filtering + URL pattern/keywords, rewrite/truncate rules, protocol content filters
  - Detect ActiveX and Java applets, media types; configure specific content as trusted
    - Remove others from the HTML code
- **Keyword detection**
  - Prevent classified material from leaving the organization
  - Prevent banned content from leaving or entering an organization

# Design Challenges With DPI

- **DPI matches IP packet data against specified patterns**
- **This must be done at network speeds**
  - DPI hardware can only hold a limited number of packets for matching
  - DPI hardware can only store a limited amount of malware patterns

# Deep Content Inspection (DCI)

## Deep Packet Inspection evolves to Deep Content Inspection

- **Deep Packet Inspection** systems
  - Examines packets, including the data in the packet
  - Rely on pattern matching and reputation lookup
- **Deep Content Inspection** systems
  - Examines content, even if it spans multiple packets
  - Unpacks encoded data
    - Example: base64-encoded MIME data in web and email content
  - Signature matching, compliance analysis (including data loss prevention)
  - Behavior analysis via correlation with previous sessions

# Intrusion Detection/Prevention Systems: IDS/IPS

# Intrusion Detection/Prevention Systems

**IDS/IPS systems are part of Application-layer firewalls**

**Identify threats and attacks**

**IDS:** *Intrusion Detection System*

- Monitor traffic at various points of the network and report problems

**IPS:** *Intrusion Prevention System*

- Sit in between two networks & control traffic between them (like a firewall)
- Enforce admin-specified policy on detection of problems

**Types of Systems**

- Protocol-based
- Signature-based    *We know what is bad; anything else is good*
- Anomaly-based    *We know what is good; anything else is bad*

# Protocol-Based IDS

**Reject packets that do not follow a prescribed protocol**

- Permit return traffic as a function of incoming traffic
- Define traffic of interest (filter), filter on traffic-specific protocol/patterns

## Examples

- **HTTP inspection**: prevent malicious HTTP attacks:  
    validate headers, cookies, URL string, content types
- **DNS inspection**: prevent spoofing DNS replies:  
    make sure they match IDs of sent DNS requests
- **SMTP inspection**: restrict SMTP command set  
    ... and command count, arguments, addresses
- **FTP inspection**: restrict FTP command set  
    ... and file sizes and file names



# Signature-based IDS

**Don't search for protocol violations but for possible data attacks**

## **Match patterns of known “bad” behavior**

- Viruses
- Malformed URLs
- Buffer overflows

## **Need a database of known protocol attacks & malware**

- Signature = data segments & order of packets that make up the attack
- Only detects known attacks

# Anomaly-based IDS

**Search for statistical deviations from normal behavior**

**Establish baseline behavior first**

## **Examples:**

- Port scanning
- Imbalance in protocol distribution
- Imbalance in service access

## **Challenge**

- Distinguish anomalies from legitimate traffic

# Next-Generation Firewalls (NGFW)

**Term for a firewall that combines**

**Stateful packet inspection +**

**Deep packet inspection +**

**Intrusion prevention**

- **Decrypt & re-encrypt TLS & ssh traffic**
  - Breaks end-to-end encryption; firewall is a man-in-the-middle
  - Clients will need to get & validate the firewall's certificate
- **Application awareness**
  - Classify types of apps; assign risk levels & define app-specific policies

# Host-based (personal) firewalls


- **Run on the user's systems, not as dedicated firewalls**
- **Application awareness**
  - Manage network-facing effects of malware
    - Allow only approved applications to send or receive data over the network
- ***Important for defense in depth***
- **Problem**
  - If malware gets elevated privileges, it can reconfigure or disable the firewall
- **Personal IDS**
  - E.g., **fail2ban** on Linux
    - Scan log files to detect & ban suspicious IP addresses
    - High number of failed logins, probes, URLs that try to target exploits

# Application proxies

# Application proxies

## Proxy servers

- Intermediaries between clients and servers
- Stateful inspection and protocol validation



# Summary

<b>Firewall (screening router)</b>	1 <sup>st</sup> generation packet filter that filters packets between networks. Blocks/accepts traffic based on IP addresses, ports, protocols
<b>Stateful inspection firewall</b>	2 <sup>nd</sup> generation packet filter – like a screening router but also considers TCP connection state and information from previous connections (e.g., related ports for services)
<b>Deep Packet Inspection firewall</b>	3 <sup>rd</sup> generation packet filter – examines application-layer protocols
<b>Application proxy</b>	Gateway between two networks for a specific application. Prevents direct connections to the application from outside the network. Responsible for validating the protocol.
<b>IDS/IPS</b>	Can usually do what a stateful inspection firewall does + examine application-layer data for protocol attacks or malicious content. Usually a part of Deep Packet Inspection firewalls
<b>Host-based firewall</b>	Typically screening router with per-application awareness. Sometimes includes anti-virus software for application-layer signature checking
<b>Host-based IPS</b>	Typically allows real-time blocking of remote hosts performing suspicious operations (port scanning, ssh logins)

# Firewall Challenges



# Intrusion detection & prevention problems

- **There's a lot of stuff going on**
  - People visit random websites with varying frequencies
  - Software accesses varying services
  - Buggy software may create bad packets
  - How do you detect what is hostile?
- **Traffic volume from attacks may be miniscule compared to legitimate traffic**
  - Even a small % of false positives can be annoying and hide true threats
  - Exceptions would be compromised systems launching a DDoS attack or exfiltrating data
- **Environments are dynamic**
  - Content from CDNs or other large server farms has a broad range of IP addresses
  - Malicious actors can coexist with legitimate ones

# Intrusion detection & prevention problems

- **Encrypted traffic cannot be easily inspected**
  - Just because you visit a web site using HTTPS doesn't mean the site is secure ... or hasn't been compromised
  - Encrypted  $\neq$  trustworthy
- **Packet inspection provides a limited view into activity**
  - You may need to extract data from multiple packets
  - You may need to reconstruct sessions
  - Both of these are time consuming and can affect performance
- **Threats & services change over time**
  - Rules must be updated

# Deperimeterization

**Boundaries between internal & external systems are harder to identify and may be fluid**

**Systems in a trusted network cannot implicitly be assumed to be trustworthy**

- Mobile devices
- Cloud-based computing
- USB flash drives
- Web applications, web services
- Internal systems may get compromised
  - Accidental downloads of malware
  - Attacks to exposed services or via outbound connections
  - Malicious insiders

# Zero-Trust Architecture (ZTA)

*Don't assume something is secure because it's within your network!*

**Don't allow access to a service until the user & service are mutually authenticated and the user is authorized to access the service**

- Enforce the **Principle of Least Privilege**
  - Enable access only when policies allow it
- No device is implicitly trusted
- **Rely on multifactor authentication, access control, and encryption**
  - Authentication to one resource doesn't mean you have access to others

**Core principle: “Never Trust, Always Verify”**

# 7 Tenets of Zero Trust (NIST SP 800-207)

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications and uses it to improve its security posture.

# Challenges with Zero-Trust Access

- **Ideally, every connection will be authenticated, authorized, and encrypted**
  - True end-to-end connectivity requires application awareness (e.g., link with a library providing the services or have operating systems enforce end-to-end security)
  - Many services do not support centrally-managed access control – and there isn't a standard
    - Manage user authentication credentials, entitlements (permissions)
- **Fallback: Zero Trust Network Access (ZTNA)**
  - Security managed at the transport layer rather than all the way to the application
  - Fallback where zero trust is provided between hosts rather than within apps
  - E.g., create host-host VPNs with user authentication and packet filtering
- **Networks may need micro-segmentation as a safeguard**
  - Move users or groups of resources into separate network segments
  - Defense in depth strategy: limit damage even if a system gets compromised
- **Insider threat is still a problem**
  - So are stolen credentials and compromised devices

# Government of Zero Trust Initiatives


**US federal agencies**  
**Deadline of Sept 30, 2024 to**  
**move from perimeter-based**  
**defenses to zero trust**

**European Union**  
**Network and Information**  
**Security Directive (NIS2)**  
**Implementation deadline of**  
**October 2024**

**UK: No mandate but strongly**  
**promotes zero trust and**  
**provides guidelines**

**Canada: Core part of the**  
**Government's Cyber Security**  
**Strategy**

THE WHITE HOUSE



JANUARY 26, 2022

## Office of Management and Budget Releases Federal Strategy to Move the U.S. Government Towards a Zero Trust Architecture

Today, the Office of Management and Budget (OMB) released a Federal strategy to move the U.S. Government toward a "zero trust" approach to cybersecurity. The strategy represents a key step forward in delivering on President Biden's Executive Order on Improving the Nation's Cybersecurity, which focuses on advancing security measures that dramatically reduce the risk of successful cyber attacks against the Federal Government's digital infrastructure.

The growing threat of sophisticated cyber attacks has underscored that the Federal Government can no longer depend on conventional perimeter-based defenses to protect critical systems and data. The Log4j vulnerability is the latest evidence that adversaries will continue to find new opportunities to get their foot in the door. The zero trust strategy will enable agencies to more rapidly detect, isolate, and respond to these types of threats. By detailing a series of specific security goals for agencies, the new strategy will serve as a comprehensive roadmap for shifting the Federal Government to a new cybersecurity paradigm that will help protect our nation. These goals are directly aligned with and support existing zero trust models.

CIO

## Zero Trust Security for NIS2 compliance: What you need to know

BrandPost • By Eve-Marie Lanza, Senior Security Solutions Marketing Manager, Aruba  
Sep 12, 2023 • 6 mins

The NIS2 requirement to adopt a Zero Trust architecture reflects the limitations of models based on implicit trust—this security approach has exposed organizations to great risk.



CREDIT: ISTOCK

National Cyber  
Security Centre

GUIDANCE

## Zero trust architecture design principles

Eight principles to help you to implement your own zero trust network architecture in an enterprise environment.

### Zero trust architecture design principles

Introduction to Zero Trust

1. Know your architecture including users, devices, services and data
2. Know your user, service and device identities
3. Assess user behaviour, service and device health
4. Use policies to authorise requests
5. Authenticate and authorise everywhere
6. Focus your monitoring on users, devices and services
7. Don't trust any network, including your own
8. Choose services which have been designed for zero trust

Help implementing zero trust architecture

PAGE 1 OF 15

Government of Canada  
Gouvernement du Canada

Canada.ca • Canadian Centre for Cyber Security • Cyber security guidance

## Zero Trust security model - ITSAP.10.008

From: [Canadian Centre for Cyber Security](#)

November 2022 | Awareness series

Alternate format: [ITSAP.10.008 Zero Trust Security Model \(PDF, 294 KB\)](#)

The traditional security model used by organizations to protect information systems focused on perimeter defense and implicitly trusted anyone inside the corporate network, therefore granting them access to resources. As more governments and enterprises undergo digital transformation, adopt cloud-based technologies and embrace remote/hybrid work, the traditional perimeter-focused defenses are no longer sufficient to protect internal networks and data. This document provides information on Zero Trust (ZT) as a model to address the modern challenges of securing remote workers, protecting hybrid cloud environments and defending against cyber security threats.

### What is Zero Trust?

The term "Zero Trust" (ZT) does not apply to a single product, technology, or architecture layer. Rather, it represents a security framework for protecting infrastructure and data. ZT's central tenet is that no subject (application, user, or device) in an information system is trusted by default. Trust must be re-assessed and verified every time a subject requests access to a new resource. The degree of access provided is dynamically adjusted based on

# The End