Intrusion

# Shield OnPremise

## Case Study

Autonomous Reputation-Based Enforcement in a Defense Contractor Environment

# Executive Summary

This case study examines the deployment of Intrusion Shield OnPremise within a small defense contractor environment over a three-month operational period. The objective of the deployment was to strengthen network security posture through autonomous inline enforcement while minimizing operational disruption and administrative burden.

The findings demonstrate several significant outcomes:

- Extremely high churn rates among domains, hostnames, resolved IP addresses, and outbound server infrastructure
- Meaningful reductions in exposure to untrusted infrastructure through autonomous enforcement
- Strong effectiveness in identifying and blocking previously unknown infrastructure
- Exceptionally low operational false positive rates
- Minimal impact on latency and business operations

Most importantly, this deployment highlights a key differentiator of Shield OnPremise. Unlike traditional intrusion detection and prevention technologies that primarily rely on known indicators of compromise or community-confirmed malicious infrastructure, Shield proactively blocks communications associated with unknown or untrusted reputation states.

This distinction is operationally significant. In practical enterprise environments, previously unknown infrastructure is frequently later identified by the broader cybersecurity community as malicious. By enforcing trust decisions before consensus is established, Shield reduces the exposure window commonly exploited by modern threat actors.

# Deployment Overview

Shield OnPremise was deployed inline between the organization's core switch and firewall infrastructure in full Protect Mode operation.

The deployment architecture enabled Shield to:

- Inspect ingress and egress communications
- Maintain visibility into NAT clients
- Operate transparently within the existing environment
- Provide autonomous enforcement without requiring substantial architectural changes

In this deployment scenario, the Shield appliance was positioned inline between a Cisco Systems Meraki firewall and the Meraki core switch across a 10 Gbps network link.

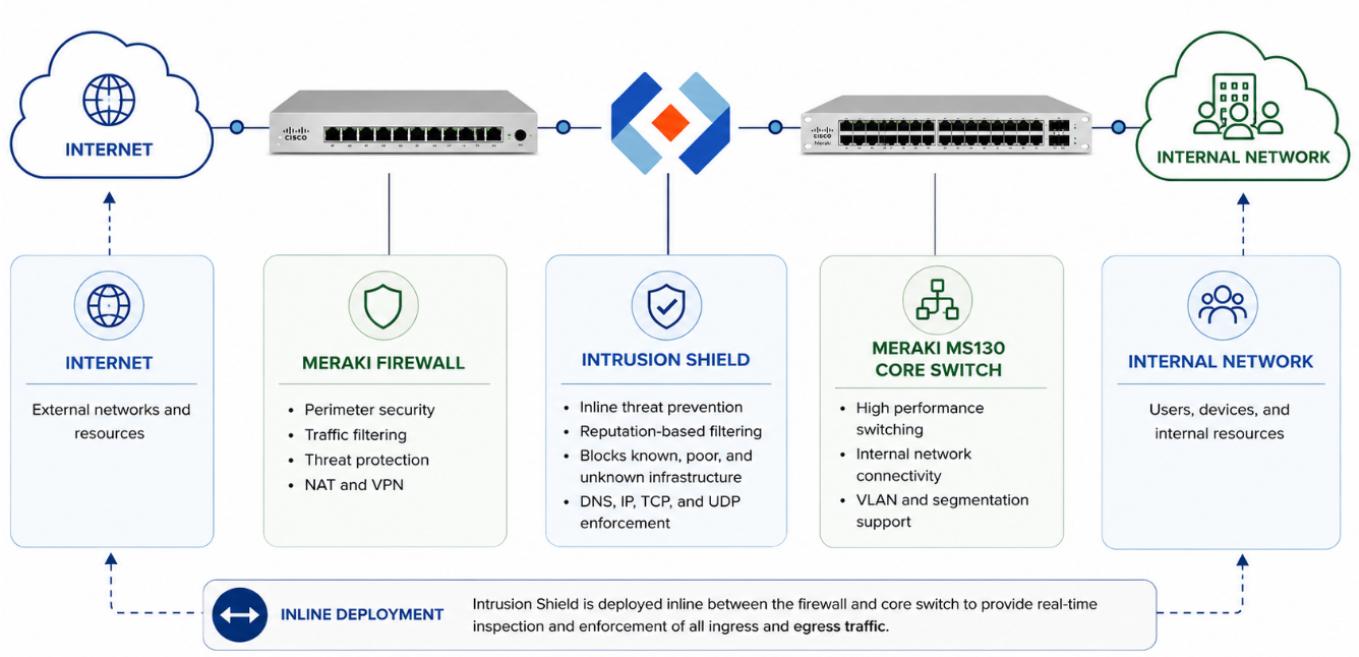


This network is also protected by a Cisco Meraki firewall and uses Akamai's GovShield for Protective DNS. The deployment of Shield compliments these solutions, adding additional protection for blocking high-risk and no-reputation domains and IPs.

Throughout the observation period:

- Average bandwidth utilization remained below 750 Mbps
- Typical inbound utilization ranged between 400 and 600 Mbps
- Outbound bandwidth generally remained below 50 Mbps
- Peak daily throughput occasionally approached 750 Mbps

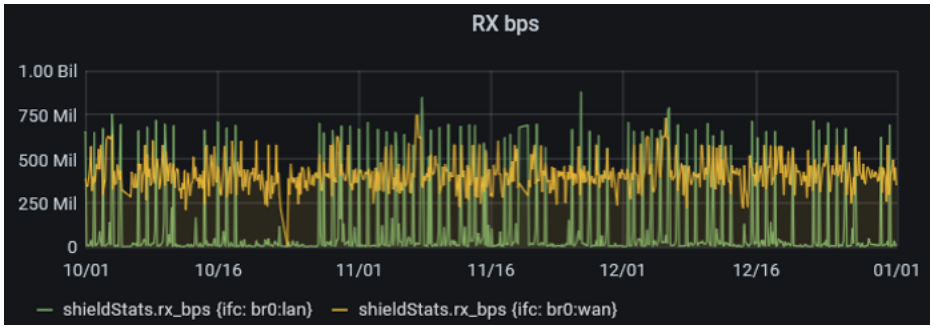


Operational impact to users and applications was negligible.

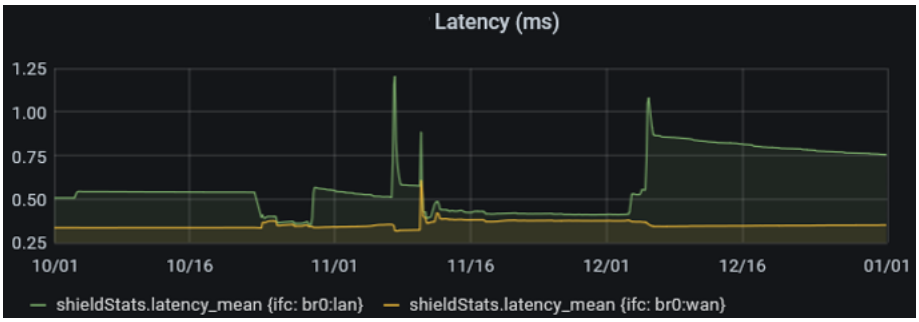


**Inline latency introduced by Shield remained consistently below 1 millisecond.**

This deployment model allowed the organization to rapidly operationalize advanced reputation-based enforcement capabilities without disrupting business operations.

# Understanding Internet Infrastructure Churn

Modern internet infrastructure is highly dynamic. Domains, hostnames, cloud instances, and IP addresses constantly evolve as organizations adopt distributed architectures, cloud-native services, and content delivery platforms. The set of domains and IPs with which a typical mid-size business communicates is not a stable set.

Threat actors exploit this same dynamism.

Malicious infrastructure increasingly relies on:

- Newly registered domains
- Ephemeral cloud environments
- Disposable command-and-control servers
- Rapidly rotated phishing infrastructure
- Fast-flux hosting techniques
- Short-lived IP allocations

Traditional security controls often struggle to respond quickly enough because many defenses rely on community validation or historical reputation before enforcement decisions are made.

Other approaches, such as building explicit allow lists for north-south traffic are highly effective for special-purpose networks but are not scalable for general Internet usage, as the effort of maintaining domain and IP allow lists creates a large operational overhead with the constant churn of Internet infrastructure.

This case study demonstrates how rapidly infrastructure changes even within a relatively stable enterprise environment.

## DNS Registered Domain Churn

The first analysis examined distinct registered domains observed during DNS resolution activity over the three-month period.

| Month | Unique Registered Domains Requested |
|---|---|
| October | 5,906 |
| November | 4,805 |
| December | 4,310 |
| Three-Month Total | 9,196 |
| **Common Across All Three Months** | **2,167 (24%)** |

The results were substantial.

Between October and November, approximately 50% of observed domains were never seen again. Between November and December, approximately 49% of domains disappeared from subsequent observations.

Across the entire three-month period, only 24% of domains remained consistent throughout all months analyzed.

**In practical terms, 76% of observed domains appeared during only a single month.**

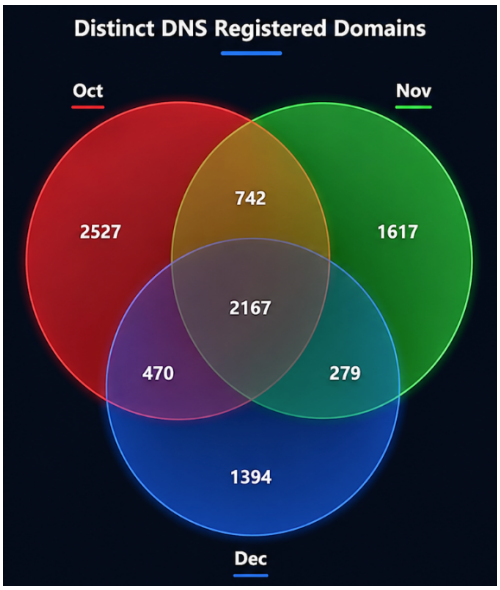


These findings reinforce a critical operational reality. Enterprise traffic patterns are constantly changing, and trust decisions based solely on previously categorized infrastructure create significant defensive gaps.

## DNS QName Churn

The analysis next examined DNS QNames, or fully qualified domain names (FQDNs), observed within DNS requests.

| Month | Unique FQDNs Requested |
|---|---|
| October | 28,744 |
| November | 23,867 |
| December | 21,445 |
| Three-Month Total | 54,762 |
| **Common Across All Three Months** | **7,193 (13%)** |

Hostname churn proved even more significant than registered domain churn.

**Only 13% of observed QNames persisted across all three months. Approximately 87% appeared during only a single month of operation.**

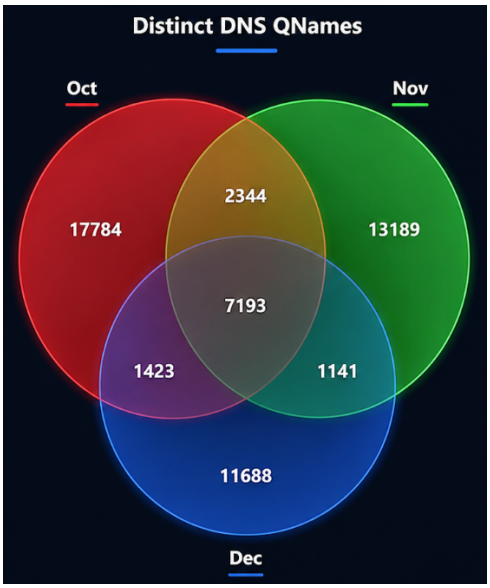


These findings further illustrate why static allow-list approaches and purely reputation-confirmed models often struggle operationally in modern enterprise environments.

## DNS Resolved IP Churn

Shield also analyzed resolved IPv4 addresses appearing within DNS responses.

| Month | Unique IPv4 Resolved IPs |
|---|---|
| October | 48,813 |
| November | 41,649 |
| December | 36,360 |
| Three-Month Total | 84,842 |
| **Common Across All Three Months** | **14,873 (18%)** |

Only 18% of resolved IP addresses remained consistent across all three months.

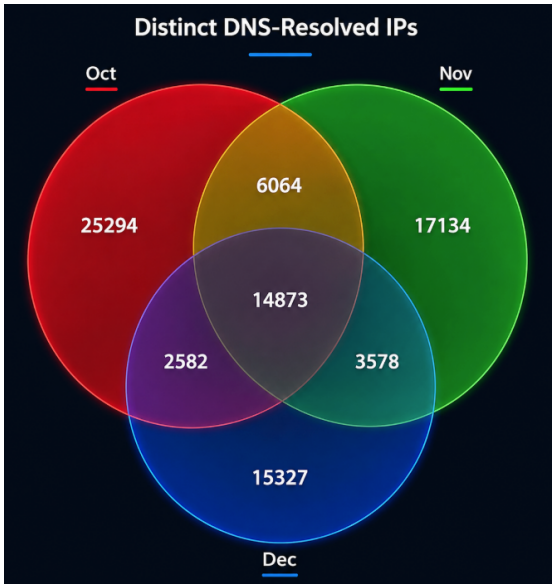


This behavior is expected in modern internet architectures utilizing:

- Content delivery networks
- Cloud load balancing
- Dynamic infrastructure scaling
- Distributed service delivery

However, these same characteristics are heavily leveraged by adversaries attempting to evade detection and rapidly rotate malicious infrastructure.

## Outbound TCP Server IP Churn

Shield additionally analyzed outbound TCP communications initiated from internal systems to external internet infrastructure.

During the observation period, Shield observed approximately 22 million outbound TCP connections.

| Month | Unique Server IPs |
|---|---|
| October | 24,882 |
| November | 20,989 |
| December | 18,906 |
| Three-Month Total | 42,069 |
| **Common Across All Three Months** | **8,077 (19%)** |

Only 19% of outbound server IP addresses persisted across all three months.

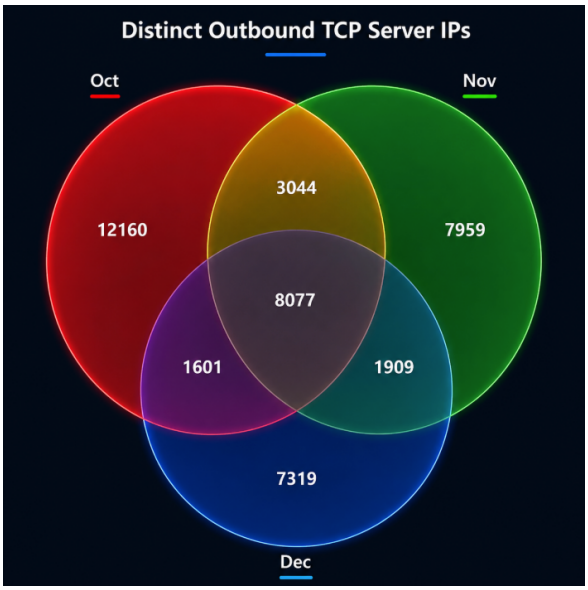


**Infrastructure volatility is no longer exceptional. It is normal.**

Security architectures that assume unknown infrastructure is inherently trustworthy create exploitable windows for attackers operating within this constantly shifting ecosystem.

# Reputation-Based Blocking and Autonomous Enforcement

Shield performs both DNS-based and connection-based enforcement.

Each DNS request is evaluated at multiple levels:

- The queried hostname itself
- Associated CNAME records
- IP addresses returned within DNS answers

Each element is evaluated against Shield's reputation intelligence framework.

Infrastructure may be:

- Explicitly allowed
- Explicitly blocked due to known malicious or high-risk reputation
- Blocked due to unknown or insufficient reputation

This final category represents one of Shield's most important differentiators.

**Rather than assuming unknown infrastructure is trustworthy until proven otherwise, Shield applies a verification-first enforcement philosophy aligned with zero trust principles.**

## DNS Domain Reputation Blocking

| Type | Unique Registered Domains |
|---|---|
| Total DNS Registered Domains | 9,196 |
| Allowed | 8,012 (87%) |
| QNAME Blocked Due to Poor Reputation | 186 (2%) |
| QNAME Blocked Due to Unknown/No Reputation | 266 (3%) |
| CNAME Blocked Due to Poor Reputation | 12 (0.1%) |
| CNAME Blocked Due to Unknown/No Reputation | 20 (0.2%) |
| Resolved IP Blocked | 700 (8%) |

Approximately 13% of domains encountered during the study period resulted in blocking actions.

Importantly, a significant portion of these enforcement actions involved previously unknown infrastructure rather than already-confirmed malicious entities.

This is operationally meaningful because newly established infrastructure is frequently leveraged during early-stage attack campaigns before broader industry recognition occurs.

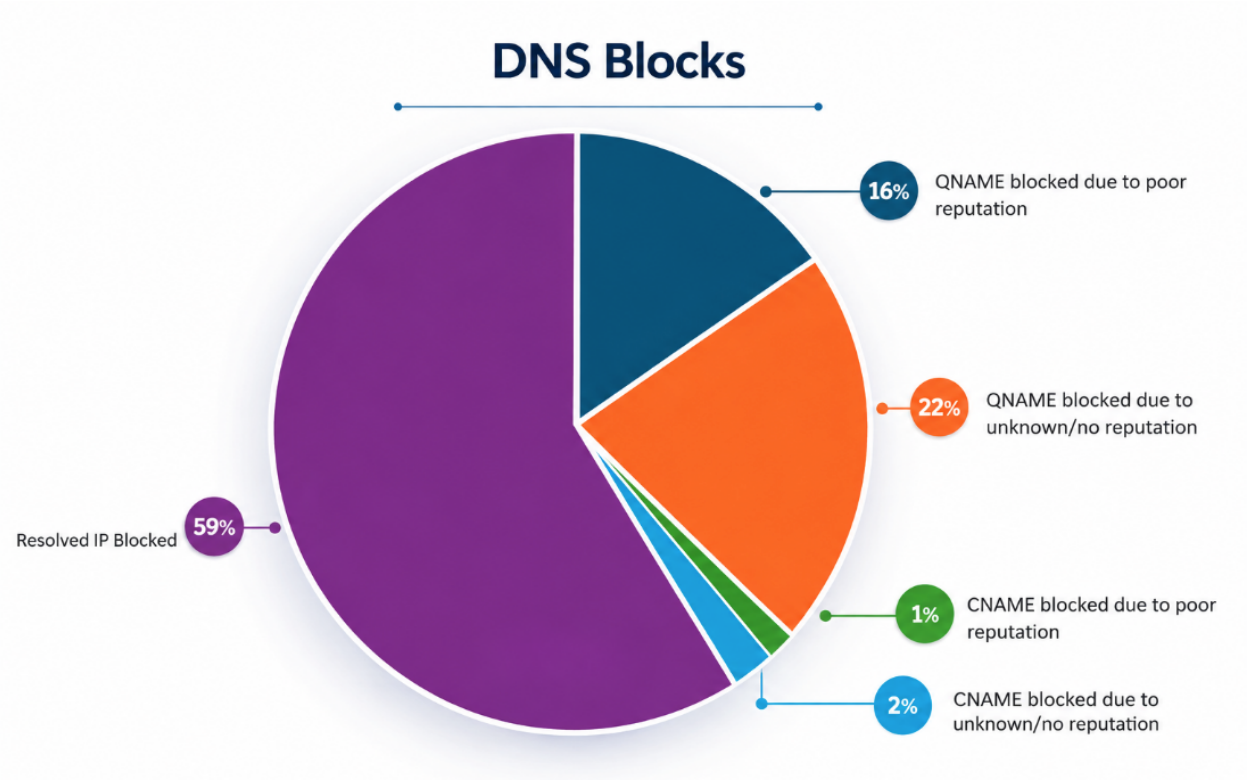


## DNS QName Reputation Blocking

| Type | Unique FQDNs |
|---|---|
| Total DNS Registered Domains | 54,762 |
| Allowed | 54,074 (99%) |
| Blocked Due to Poor Reputation | 308 (0.6%) |
| Blocked Due to Unknown/No Reputation | 380 (0.7%) |

Although the percentage values appear relatively small, the operational impact remains substantial given the large volume of communications occurring within enterprise environments.

These statistics also demonstrate that Shield's enforcement model remains highly selective and controlled despite aggressively evaluating unknown infrastructure.

Note: This dataset overlaps the registered domain dataset, but the numbers are diluted because the legitimate registered domain names typically have more QNames (subdomains) as high traffic sites use Content Delivery Networks (CDNs) and other legitimate subdomains (such as mail.yahoo.com, finance.yahoo.com, answers.yahoo.com) than the less reputable sites (such as a phishing domain that only has one host).

## DNS Resolved IP Blocking

Shield uniquely evaluates resolved IP addresses contained within DNS responses.

| Type | Unique DNS Answer IPs |
|---|---|
| Total Resolved IPs | 84,852 |
| Allowed | 82,507 (97%) |
| Blocked Due to Poor Reputation | 2,345 (3%) |

This capability allows Shield to identify situations such as:

- Legitimate-looking domains resolving to malicious infrastructure
- DNS hijacking scenarios
- Fast-flux infrastructure rotation
- Infrastructure changes occurring after prior evaluations

Over 2,300 resolved IP addresses were blocked due to poor reputation during the study period.

## Outbound Server IP Blocking

Shield extends enforcement beyond DNS analysis and evaluates IP, TCP, UDP, and ICMP communications directly.

For outbound TCP communications:

| Type | Unique Server IPs |
|---|---|
| Total Outbound Server IPs | 42,069 |
| Allowed Server IPs | 40,084 (95%) |
| Blocked Server IPs | 1,985 (5%) |
| Server IPs Without Preceding DNS Requests | 7 |

Approximately 5% of observed outbound server IPs were blocked.

Notably, several blocked server IPs had no preceding DNS requests at all. This behavior is commonly associated with hardcoded infrastructure frequently used within malware and botnet operations.

These findings demonstrate why DNS filtering alone is insufficient and why layered communication enforcement remains critical.

# False Positive Analysis

One of the primary concerns surrounding aggressive enforcement models is the potential for operational disruption caused by false positives. However, the absence of third-party threat intelligence associated with a domain or IP address does not inherently indicate that the infrastructure is trustworthy or legitimate. Modern threat actors frequently leverage newly established or previously unseen infrastructure before it becomes widely recognized by the broader security community.

Shield evaluates infrastructure based on reputation and trustworthiness rather than relying solely on confirmed malware classifications. If Shield blocks a connection associated with poor, unknown, or insufficient reputation, and that enforcement action does not negatively impact business operations, the event is not considered a false positive. Instead, it is treated as the successful prevention of an unnecessary or potentially risky communication.

Within this deployment, false positives were measured operationally based on whether a block materially affected users or required administrative override. The customer's IT team actively managed permit requests and selectively approved specific domains or IP addresses when legitimate business requirements justified access.

**This deployment demonstrated exceptionally low operational impact.**

During the three-month period:

- Only 7 domain or hostname permits were added by administrators
- Only 21 IP permits were added

| Type | Count |
|---|---|
| Blocked Domains/Hosts | 9,196 |
| User-Permitted Domains | 7 |
| **Effective Domain False Positive Rate** | **0.07%** |
| Blocked IPs | 2,345 |
| User-Permitted IPs | 21 |
| **Effective IP False Positive Rate** | **0.89%** |

These findings are particularly significant because Shield aggressively evaluates unknown infrastructure while still maintaining extremely low operational friction.

**The data demonstrates that proactive reputation enforcement and business continuity are not mutually exclusive objectives.**

# Conclusion

This deployment demonstrated that Shield OnPremise can materially strengthen enterprise security posture while maintaining operational stability and minimizing administrative burden.

Several findings stand out:

- Internet infrastructure churn is extraordinarily high
- Unknown infrastructure constitutes a substantial percentage of enterprise communications
- Traditional trust assumptions create exploitable defensive gaps
- Autonomous reputation-based enforcement meaningfully reduces exposure
- Aggressive reputation filtering can be achieved with very low false positive rates

Most importantly, this case study validates the operational value of Shield's core philosophy:

**Verify before trust.**

Modern attackers increasingly rely on infrastructure that exists only briefly before disappearing or becoming widely recognized as malicious. By proactively controlling communications associated with unknown or untrusted infrastructure, Shield significantly reduces the exposure window commonly exploited in modern cyberattacks.

For organizations operating sensitive, regulated, or mission-critical environments, this capability provides a meaningful advancement beyond traditional reactive detection models.

# Appendix A

## Frequently Used Terms and Definitions

### Autonomous Enforcement

A security model in which enforcement decisions are made automatically and in real time without requiring manual analyst intervention. Shield OnPremise operates autonomously inline, allowing malicious or untrusted communications to be blocked immediately rather than generating alerts for later review.

### Botnet

A collection of compromised systems controlled remotely by an attacker. Botnets are frequently used for distributed denial-of-service attacks, credential theft, spam campaigns, malware distribution, and command-and-control operations.

### CNAME (Canonical Name Record)

A type of DNS record that maps one domain name to another domain name. CNAMEs are commonly used for content delivery networks, cloud services, and infrastructure abstraction. Example: mail.company.com → hostedmail.provider.com

### Command-and-Control (C2)

Infrastructure used by attackers to communicate with compromised systems. Command-and-control servers allow threat actors to issue instructions, exfiltrate data, deploy malware, or maintain persistence within a compromised environment.

### Content Delivery Network (CDN)

A distributed infrastructure platform designed to improve performance and availability by serving content from geographically distributed servers. CDNs frequently introduce high levels of IP and hostname churn due to dynamic routing and load balancing.

### DNS (Domain Name System)

The distributed system responsible for translating human-readable domain names into IP addresses used by computers and network devices. Example: www.example.com → 192.0.2.10

### DNS Answer

The response returned by a DNS server containing one or more resource records, such as resolved IP addresses or CNAME records.

### Egress Traffic

Outbound network communications originating from systems inside an organization's network and destined for external infrastructure on the internet.

### Ephemeral Infrastructure

Short-lived cloud or internet infrastructure created temporarily and often discarded rapidly. Threat actors frequently leverage ephemeral infrastructure to evade detection and reputation tracking.

### False Positive

A legitimate communication or activity incorrectly identified as malicious or blocked by a security control. Within this case study, false positives are measured operationally based on whether a block created business disruption requiring administrator override.

### Fast-Flux Infrastructure

A technique commonly used by attackers in which domains rapidly rotate through large numbers of IP addresses to evade detection, increase resiliency, and complicate tracking efforts.

### FQDN (Fully Qualified Domain Name)

The complete domain name specifying the exact location of a host within the DNS hierarchy. Example: mail.research.organization.gov

### ICMP (Internet Control Message Protocol)

A network protocol used for diagnostics and operational communications, including functions such as ping and error reporting.

### Inline Deployment

A deployment architecture in which a security device sits directly in the traffic path and can actively inspect, permit, or block communications in real time.

### Intrusion Detection System (IDS)

A security technology designed to identify potentially malicious activity and generate alerts. Traditional IDS platforms are often passive and do not actively block traffic.

### Intrusion Prevention System (IPS)

A security technology capable of both detecting and actively blocking malicious communications.

### IP Address

A numerical identifier assigned to devices and systems communicating across a network using Internet Protocol. Example: 203.0.113.25

### Latency

The amount of delay introduced during network communications. Inline security controls are often evaluated based on their operational latency impact.

### NAT (Network Address Translation)

A networking technique that allows multiple internal systems to share one or more public IP addresses when communicating externally.

### North-South Traffic

Communications traveling between internal enterprise systems and external internet resources.

### QName

The exact hostname queried during a DNS lookup. Example: login.example.com

### Registered Domain

The primary domain registered with a domain registrar, excluding subdomains. Example: For mail.finance.example.com, the registered domain is example.com.

### Reputation Intelligence

Threat intelligence derived from historical observations, behavioral analysis, infrastructure tracking, malicious activity association, and long-term operational analysis of internet entities such as domains, IP addresses, and hostnames.

### Resolved IP

An IP address returned within a DNS response after a hostname lookup is performed.

### Resource Record (RR)

A structured DNS data entry contained within a DNS response. Common resource record types include A records, AAAA records, CNAME records, MX records, and TXT records.

### Server IP

The external IP address receiving a client connection during outbound communications.

### SIEM (Security Information and Event Management)

A centralized platform used to aggregate, correlate, analyze, and retain security logs and operational telemetry from multiple security tools and infrastructure sources.

### TCP (Transmission Control Protocol)

A stateful network protocol commonly used for reliable communications across the internet, including web browsing, email, and file transfers.

### Threat Intelligence

Information collected and analyzed regarding malicious infrastructure, attacker behavior, vulnerabilities, malware activity, and indicators of compromise used to improve defensive security operations.

### Zero Trust

A cybersecurity model based on the principle that no user, device, application, or communication should be implicitly trusted. Access and communications must be continuously validated before being permitted.

### Zero Trust DNS (ZTDNS)

A security architecture concept in which DNS resolution is tightly controlled using trust validation mechanisms to limit communications to explicitly approved infrastructure.