



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Network Monitoring and Threat Detection In-Depth (Security 503)"
at <http://www.giac.org/registration/gcia>

*Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters
to A Zero Trust Network Architecture*

Author: John Becker, jbecker42@gmail.com

Advisor: Tanya Baccam

Accepted:

2017-09-11

Abstract

Traditional IP-based access controls (e.g., firewall rules based on source and destination addresses) have defined the network perimeter for decades. Threats have evolved to evade and bypass these IP restrictions using techniques such as spear phishing, malware, credential theft, and lateral movement. As these threats evolve, so have the demands from end users for increased accessibility. Remote employees require secure access to internal resources. Cloud services have moved the perimeter outside of the enterprise network. The DevOps movement has emphasized speed and agility over up front network designs. This paper identifies gaps to implementation for organizations in the discovery phase of migrating to identity-based access controls as described by leading cloud companies.

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

1. Introduction

Traditional network access controls such as firewall rules or router Access Control Lists (ACLs) are ineffective at enforcing a network perimeter because the perimeter has moved to include the public cloud, DevOps automation tools, and remote workers. Connectivity is critical to productive workers but is a constant tradeoff with security. System access has evolved from physical requirements for workstations and mainframes to remote connections over Wide Area Networks (WANs), and then to remote employees and cloud services. Public cloud Infrastructure-as-a-Service (IaaS) adoption is growing (Gartner, 2017). While outsourcing most of the security management is possible with Platform as a Service (PaaS) providers, IaaS providers impose a shared responsibility model (Ullrich et al., 2016, pp. 2-3). This shared responsibility model forces changes to firewall and network management to support dynamic cloud resources. Public IP addresses change routinely between public cloud customers, and private network connections can add devices to workflows that are not automated. A Virtual Private Network (VPN) connection on a physical network device would have different configuration management than public cloud services leveraging infrastructure-as-code solutions such as Hashicorp Terraform or Amazon Web Services (AWS) CloudFormation. These scenarios create contention between feature-driven DevOps teams and firewall administrators.

Cloud DevOps and Site Reliability Engineer (SRE) teams need agility to provision environments using infrastructure-as-code. Often automated solutions connect systems that were previously-isolated. Examples include version control (e.g., GitHub, Perforce), project management (e.g., Jira) and chat (e.g., Slack, HipChat) coordinated for triggering changes and releases to cloud-based repositories (e.g., Amazon S3). FireEye summarizes these trends as “Networks that traditionally had clean borders and limited demarcation points are expanding” and that “network perimeter has dramatically shifted” (FireEye, 2017). This shift leads organizations to find better models to enforce and defend their network perimeters.

John Becker, jbecker42@gmail.com

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

As cloud services and DevOps team stretch the network model, so too do employee demands on accessibility. Up to 37% of employees telecommute regularly and require some level of remote access (Jones, 2015). This accessibility comes in the form of VPN connections to enterprise networks and mobile devices accessing services like Office 365. The remote devices establishing these VPN connections lack the physical protections. Workers can leave their laptops unguarded at coffee shops or disconnect from the VPN to check personal email. A compromised endpoint with a VPN connection becomes a gateway to poorly-protected internal systems.

Traditional IP-based perimeters enforce network isolation with ACLs and firewalls using IP addresses and ports as the key value to grant or deny access. While segmentation is necessary, it can hinder employee productivity if the access rules are not managed at the same rate end users operate. Identity-based perimeters disregard network addresses when assessing trust and access controls. Instead, identity-based systems use Multi-Factor Authentication (MFA) and short-lived access tokens or certificates whenever possible. These authentication solutions improve security by reducing risk from credential theft.

Determining the right conditions to employ identity-based perimeters requires an understanding of the connectivity risks along with compensating controls. Forrester (2013) describes the IP perimeter approach like M&M candy with a “hard crunchy outside” and a comparatively soft and chewy interior. This IP perimeter model “is no longer an effective way of enforcing security” (Forrester, 2013). Forrester (2013) goes on to state the current “trust, but verify” model is no longer valid. IP-based controls also provide a false sense of security for weak authentication methods of internal systems. Even without the increased pressures of cloud, remote workers, and DevOps access, the status quo is failing to prevent breaches of internal systems. The goal is not to weaken the perimeter, but rather to strengthen authentication and access controls.

Google, Facebook, Netflix, Uber, and Lyft have all published documentation and open source software around their identity solutions. Google has fully exposed their internal

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

services to the Internet while other cloud-based companies have retained IP controls to varying degrees. Organizations face several significant gaps when attempting to follow the designs provided by leading cloud companies. These gaps include security architecture, operations maturity, engineering talent, and executive support.

2. Let the Right Ones In

Organizations are looking to enable employees, contractors, and customers access to appropriate resources while preventing malicious activity and abuse. A standard approach is to apply network segmentation based on trust levels. Segmentation limits lateral movement and reduces exfiltration. The challenge is IP-based network segmentation is simple to design but complex to enforce. Remote workers, public cloud services, and DevOps or SRE teams place increased pressure on IP firewall and ACL management.

2.1. Defining the Perimeter

The Critical Security Controls (CSC) describe the importance of perimeter - or boundary - defense because “attackers focus on exploiting systems that they can reach across the Internet, including not only DMZ systems but also workstation and laptop computers that pull content from the Internet through network boundaries” (CIS, 2016). The controls also describe the weakening of boundary lines between traditional internal and external resources in CSC 12 (CIS, 2016). This boundary erosion is impacted by public cloud capabilities as well. Allen (2016) identifies CSC sub-controls 12.2 – 12.10 as requiring a security appliance to implement in the cloud. AWS has since launched VPC Flow Logs which covers CSC 12.9 (Amazon, 2017e).

A modern network could include various internal and intranet systems isolated from DMZ and Internet-facing services. Many organizations are adopting PaaS cloud offerings such as email with Office 365, CRM solutions with Salesforce, or HR tools from Workday. These services are hosted outside of the physical and IP perimeter but are inside the logical perimeter. The pattern continues with IaaS solutions. AWS, Google Cloud, and

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

Microsoft Azure all offer a combination of IP and identity or key-based access controls. The management of the IP ACLs themselves uses identity-based access controls with Internet-facing APIs.

2.1.1. Remote Workforce

Remote workers often translate to VPN connections on laptops and mobile devices. User experiences with VPNs vary with different services and implementations. Some common causes for users rejecting VPN access include audio and video performance, multiple VPN connection requirements for accessing isolated resources (e.g., Engineering VPN vs. Standard Corporate), or basic connectivity issues due to port or protocol restrictions. Microsoft recommends against using VPNs with Skype for Business (Microsoft, 2017).

VPNs configured with split tunneling can avoid some of these performance and connectivity issues. However, there are potential risks of compromised endpoints which may allow malicious traffic through the tunnel. Furthermore, VPN access complicates logging and visibility with changing IP addresses on the network. To some, VPN termination for remote users is the most vulnerable. Some solutions emphasize Network Intrusion Detection System (NIDS) inspection of unencrypted traffic between remote users and internal resources (Obregon, 2015). There are options for decrypting, proxying and inspecting HTTPS traffic. These add additional expense and complexity as well as miss other protocols like Secure Shell (SSH).

2.1.2. Cloud

Public cloud use introduces new pressures on IP-based perimeters. According to Filkins, a traditional, perimeter-oriented model is not completely effective in cloud environments (2017). Auto-scaling and containers change IP addresses regularly. Services such AWS S3 and Azure Storage have options for restricting access by source IP address. However, these IP restrictions reside within AWS Identity and Access Management (IAM) policies or Azure Shared Access Signatures instead of network devices. Figure 1

John Becker, jbecker42@gmail.com

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

shows a typical use-case for AWS S3. A company creates an S3 bucket and restricts access to the egress IP address of the corporate firewall using a bucket policy. This approach has several downsides: all devices behind the firewall have equal access to the S3 bucket, egress filtering is complicated on the firewall as S3 objects do not use static IP addresses, and IP access controls reside in bucket policies rather than firewall rules. Auditing and managing this setup is complicated and error-prone.




Figure 1 – Internet-based Access to S3

Another approach is to extend the corporate network and IP space to the cloud using AWS Direct Connect or Azure ExpressRoutes as depicted in Figure 2. At first glance, this is a valid solution to maintaining IP filtering in the cloud. The Virtual Private Cloud (VPC) uses the same network space as the corporate network and S3 is available on a private IP address. Standard inline firewall or network devices manage access. This model has some significant issues to consider. First, AWS bucket policies cannot use the source IP address restrictions (Amazon, 2017c). Secondly, the bucket is still accessible from the Internet with appropriate access roles or keys. Services like Amazon S3 are inherently accessible from the public internet and rely on identity-based protections such as IAM profiles. Applying IP-based access controls often creates more complexity without a reduction in risk.

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture




Figure 2 - Direct Connect with S3 VPC Endpoint

The network access problem is worse when whitelisting traffic inbound from the cloud. A public cloud IP address, such as an Elastic IP (EIP) address within AWS, can be reallocated to different instances or accounts. For example, an Elastic Compute (EC2) instance with an EIP of 54.239.31.91/32 could be a bastion host with whitelisted access to corporate network resources on day 1. On day 2, the same 54.239.31.91/32 EIP could associate with a test database instance. By day 3, 54.239.31.91/32 could be released from the AWS account and assigned to a different AWS customer altogether.

The traditional IP view of the world is misleading in the cloud. Figure 3 represents an AWS VPC with an IP-restricted connection (Site-to-Site VPN or Direct Connect with firewall) to the corporate Enterprise Resource Planning (ERP) system. From a firewall perspective, the rule allows 10.10.10.0/24 inbound to port 443/tcp on the ERP host. A DevOps team manages the AWS resources and the Network team audits and approves Security Groups for the VPC. The full 10.10.10.0/24 is allowed inbound access to provide agility for the DevOps team to scale and rebuild instances without slowing down for firewall changes.

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

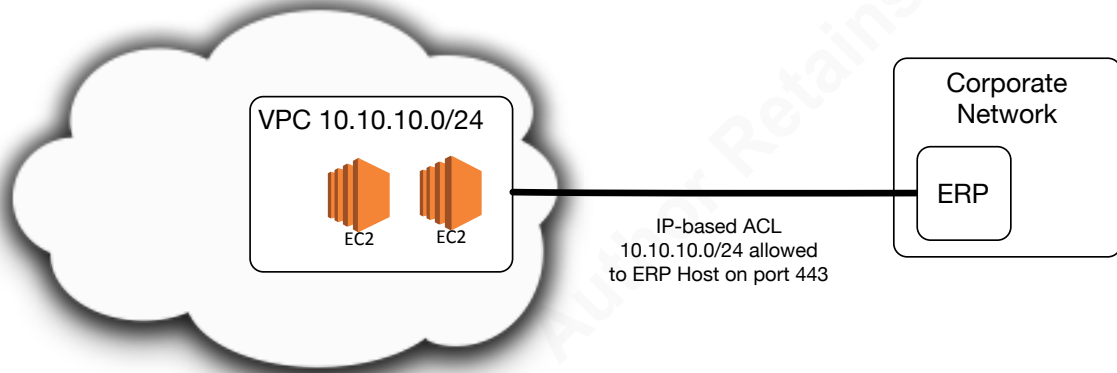


Figure 3 - Perceived Topology with IP-based Boundary

EC2 and RDS instances are not the only items that can use a private VPC address. AWS Lambdas are serverless, short-lived containers that can launch inside of a VPC (Amazon, 2017b). They will execute their request with a private IP address and terminate.

Application Programming Interface (API) Gateways are an AWS service to direct API requests to various services including Lambda (Amazon, 2017a). Figure 4 shows how traffic can completely bypass the firewall and security group for requests against the ERP system on the private network.

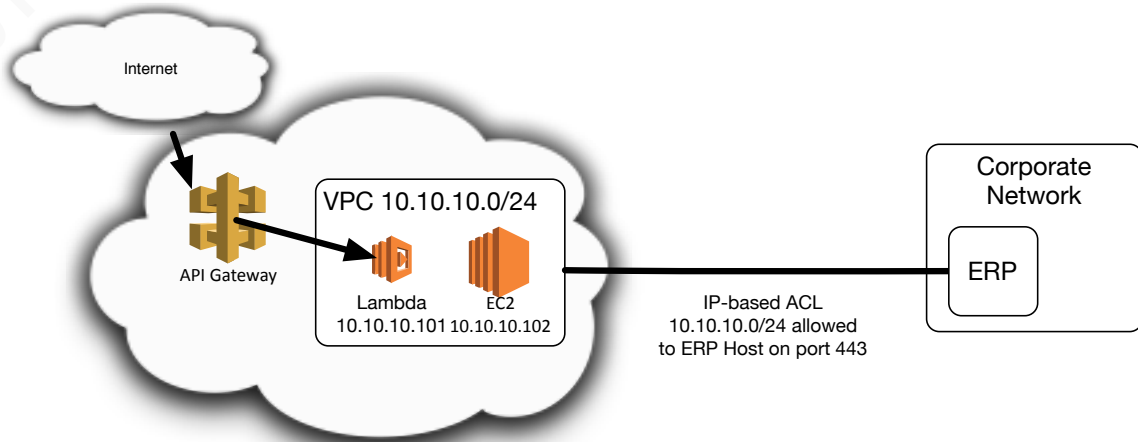


Figure 4 - API Gateway and Lambda

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

The API Gateway and Lambda combination essentially create an ingress NAT proxy from the untrusted internet to internal resources. Worse, tracking the traffic from an IP perspective is nearly impossible. The Lambda functions will dynamically use an IP address on the subnet and terminate in seconds or less. From an IP-perimeter perspective, this is ineffective and exposes internal systems that may have weak authentication. From an identity-based perimeter viewpoint, this model expands access to the internal resource while adding new, and more granular, access controls at the identity perimeter. API Gateways offer detailed logging along with auto-scaling, rate-limiting, and signed-request access controls. Additionally, these technologies are serverless, so there is no server audit and patching required. In this scenario, the security team review should include the Lambda functions, the API Gateway configuration, and the IAM profiles. These cloud services are as important as the EC2 instance settings, network access controls, and the ERP system itself.

2.1.3. DevOps and SRE

DevOps and SRE teams move quickly to bring new features and services to market. These teams release often utilizing infrastructure-as-code tools such as HashiCorp Terraform or AWS CloudFormation. Infrastructure-as-code can provision entire stacks including network, instances, policies, keys, security groups, and related services. The capability to turn on services in public clouds can significantly expand services faster than traditional controls can accommodate.

For example, a DevOps team could provision an entire application stack in a new AWS region in minutes. The longest portion of the setup is network access to traditional networks using VPNs or Direct Connect. These steps often require manual configuration and approval on the organization's edge routers and firewalls. Organizations have increased pressure to deliver services faster, and managing firewall rules and network connections can be a source of friction between teams. Removing IP filters and focusing on identity-based access reduces contention and time-to-market.

3. Think Different

Large-scale cloud-based companies like Google, Facebook, and Netflix routinely open source projects and share solutions to industry problems. Google has published numerous papers on their internal systems and processes. The BeyondCorp white papers explain Google's design, rollout, and details of their application proxy. At the center of Google's decision to remove IP restrictions to internal resources is their belief that both internal and external systems are equally untrusted (Saltonstall et al., 2016). Facebook, Netflix, Lyft, and Uber have presented various presentations and open source projects detailing similar patterns for implementing identity-based access models. Many of these solutions begin with the concept of Zero Trust.

3.1. Zero Trust Model

A Zero Trust Model asserts all network traffic is untrusted and follows three key concepts: all resources are accessed securely regardless of location, use least-privilege with strong access control, as well as log and inspect all traffic (Forrester, 2013). More succinctly, Forrester states "cybersecurity professionals must stop trusting packets as if they were people" (Forrester, 2013). Many organizations invest in a robust IP perimeter but then allow for weak internal access controls. These IP-based restrictions are inadequate as evidenced by 81% of breaches involving stolen or weak passwords (Verizon, 2017). IP-based filters and firewall rules are insufficient compensation for weak authentication. Attackers will find a way around the filters and exploit the weak authentication.

Google bases its access model on trust levels with strong enforcement of device and user identities (Beyer et al., 2014). Facebook restricts direct SSH access to production systems and enforces access through trusted networks and bastion hosts (Dutra, 2016). While Netflix is removing their perimeter (Amplify Partners, 2014), they rely on bastion access for SSH administration using their Bless SSH Certificate Authority (CA) (Lewis, 2016). These cloud companies share a common approach to mapping trust and access using strong identity management. Simultaneously, these companies reduce or eliminate

IP addresses as a means of access control. They differ their trust of devices with some favoring controlled bastion hosts over direct endpoint connections.

Google created a custom system with its Access Proxy working in tandem with device and user identity systems. To access the most-trusted resources, a user must be active and a member of the correct group while using a device that conforms to various levels of patching and secure configurations. Devices are all Google-owned (no Bring Your Own Device) and identified using certificates allocated from their Device Inventory Database (Beyer et al., 2014). Google trusts a combination of users and devices over IP addresses because they have invested heavily in writing custom management solutions. Netflix and Facebook do not advocate the same commitment to endpoint trust. Instead, both organizations reference bastion hosts for accessing systems via SSH.

3.2. Short but Sweet

A common theme for strong identity management is short-lived access tokens or certificates with secure application proxies and bastion hosts. In contrast, more traditional IP access controls allow “internal-only” addresses (located on physical or via VPN connections) to connect using password or SSH key-based authentication. This pattern is particularly well-documented regarding the use of SSH certificate authentication instead of passwords or keys. Passwords and keys can be lost and are difficult to rotate. Conversely, SSH certificates can be set to automatically expire after a set time (Moody, 2017).

This short-lived access is common in many modern applications to combat credential theft and replay scenarios. AWS has the Security Token Service (STS) which enables roles and identity federation (Amazon, 2017d). Kerberos has relied on short-lived tickets for decades (Neuman et al., 2005). These solutions are successful as key rotation and expiration are part of the standard. There is no requirement to force a user to pick a new password or generate a new SSH key every 90 days. Rather, the expiration can easily be set to reflect the level of trust assigned to the service. The result is a low-friction and scalable solution for authorizing access to services at a large scale.

3.2.1. Blessed SSH

Netflix open-sourced their AWS Lambda-based SSH Certificate Authority on GitHub (Netflix, 2017a). For organizations using AWS, this is a straightforward means to adopting short-lived SSH access without the use of IP restrictions. Figure 5 shows the separation of SSH signing to a dedicated AWS account hosting the Lambda function and KMS for securely storing the CA password. The signing of the SSH key on the bastion instance happens without any IP restrictions – the trust boundary resides within the AWS IAM roles and policies. Lyft took the lessons and technology from Netflix Bless and extended it to include SSH Certification authentication from the endpoint to the bastion servers (Steipp, 2017).

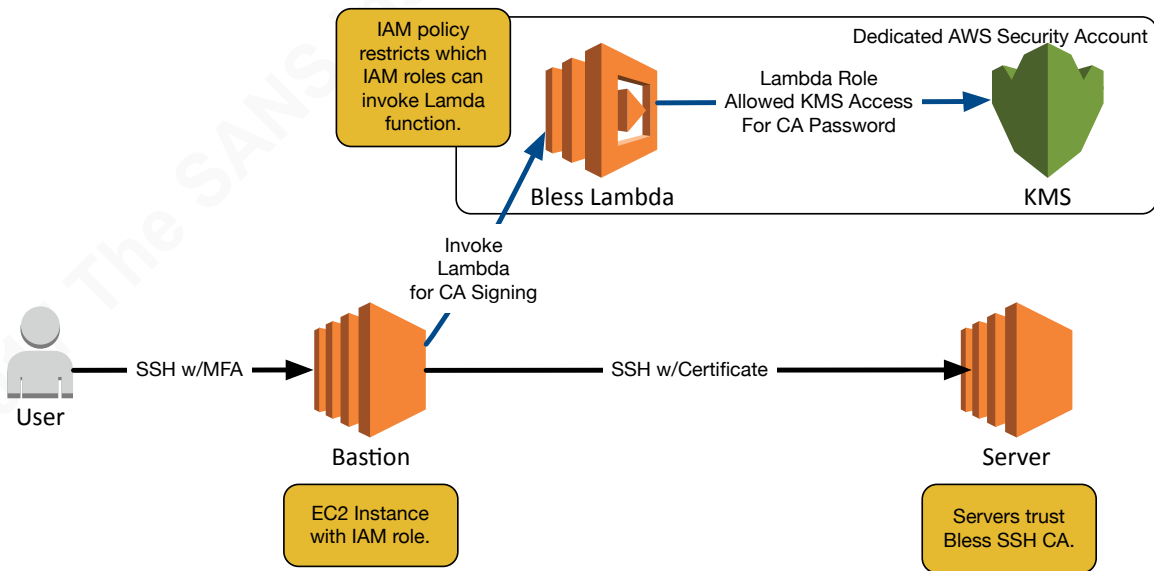


Figure 5 – Bless CA with Bastion Host

4. Mind the Gap!

Organizations wanting to move to identity-based perimeters face three challenge areas. First, inventory solutions will need to change with dynamic cloud resources and DevOps release cycles. The first four Critical Security Controls (device inventory, software inventory, secure configurations, and vulnerability assessment and remediation) are

foundational. Identity begins with CSC 5 and relies on controls 1-4. Secondly, a strong engineering culture is needed to coordinate the applications with identity access systems. Organizations not familiar with coding solutions for open source and commercial tools may find it difficult to deploy the solutions advocated by leading cloud organizations. Lastly, organizations will need to expand their trust zone definitions to focus on identity boundaries that permeate through internal systems, cloud services, and DevOps automation.

4.1. Basic Instincts

The authors of the Critical Security Controls prioritized them for a reason. Advanced controls with access management and segmentation cannot happen if the components are unknown. Scanning finds vulnerabilities (CSC 4) on the devices (CSC 1) and software (CSC 2) in inventory as related to their configuration (CSC 3). Google talks a length about their device and user inventories as well as privilege managed and malware defenses (Beyer et al., 2014). With these basics established, the identity-based access controls can be created (CSC 4). These systems may supplant IP-based requirements in CSC 9 or CSC 12. Removing IP controls is the outcome of a foundation of controls that makes IP restrictions sufficiently irrelevant.

4.2. Culture Club

Organizations successfully using identity-based perimeters have mature engineering cultures and often focus on open source software. This atmosphere is critical due to the interaction between identity management and applications. It is comparatively simple to setup a VPN with MFA to enable access to an internal web service. Identity-based access using an access proxy with TLS termination and device, configuration, and user validation is difficult. Google went so far as to create a tunneling proxy for SSH as well as rewrote cURL to work through their Access Proxy (Cittadini et al., 2016). This level of access management requires a significant commitment to building and supporting custom tooling in-house.

Google does not inherently rely on private DNS and IP addresses for system management. Google describes their container-based model as follows:

Because tasks are fluidly allocated over machines, we can't simply rely on IP addresses and port numbers to refer to the tasks. We solve this problem with an extra level of indirection: when starting a job, Borg allocates a name and index number to each task using the *Borg Naming Service* (BNS). (Beyer et al., 2016). This mindset eases a transition to an identity-based framework because IP addresses are not central to daily operations. In contrast, many organizations have inventories based on hostnames and IP addresses.

Other cloud organizations have similar backgrounds. Facebook designs open source hardware and has open-sourced over 170 projects (Facebook, 2017). Netflix has open-sourced over 125 projects on GitHub with a focus on their AWS hosting environment (Netflix, 2017b). Uber released their ussh PAM module to map SSH certificate principals to sudo access (Uber, 2017). Management supports this culture as shown by the range of company-sponsored repositories and presentations. Google explicitly states a prerequisite is for top-level management to drive the effort at all levels (Peck et al., 2017). Netflix's VP of IT Operations describes a move to more open source across all their technology areas while maintaining a balance between "build vs. buy" solutions (Amplify Partners, 2014).

Another pattern is the scale at which companies like Netflix, Google, and Facebook operate. Facebook discusses the benefits of central authentication with LDAP and Kerberos but mentions these technologies become a significant point of failure at large scale (Dutra, 2016). Most organizations will not approach the scale of Facebook and should be careful with abandoning existing LDAP and Kerberos solutions. In fact, LDAP and Kerberos are the authentication source to bastion hosts at Facebook. In turn, these bastion hosts use highly-scalable SSH Certificate Authorities to grant access to the fleet of servers (Dutra, 2016).

SSH has multiple open source and commercial solutions for certificate and key-based access. Examples include native SSH CA signing built into OpenSSH or open source offerings like Vault (Hashicorp, 2017). Application access proxies as described by Google will be more difficult. Open source proxies like Nginx and Apache are options for mutual authentication or various Single Sign On solutions. The security plugins ModSecurity (Trustwave, 2017) and Fail2ban (2017) can further improve security. However, web proxies are only a start and do not offer the full range of access described by Google.

4.3. Trust Fund

Security is rooted in assigning trust to people and objects. Segmentation of trust zones is foundational for any organization to reduce risk and lateral movement. Unfortunately, Mandiant reports many organizations still lack basic network and data segmentation (FireEye, 2017). The absence of segmentation policy leads to inconsistent enforcement that frustrates users or exposes systems to compromise.

Access between these segments is often described with either IP-based controls (e.g., allow subnet 10.10.10.0/24 to 192.168.0.0/16 on port TCP/80) or through identity-based controls (e.g., allow john.becker@sans.org to portal.sans.org). The real world can be a combination of IP and identity controls with certain users allowed access from specific IP addresses. Figure 6 illustrates an AWS IAM policy using the identity of “john.becker” and an IP restriction of “10.10.10.0/24”.

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "arn:aws:iam::00123456789:user/john.becker"
    "Condition": {
      "IpAddress": {
        "aws:SourceIp": [
          "10.10.10.0/24"
        ]
      }
    }
  }
}
```

Figure 6 – AWS IAM Policy with IP Restriction

Without a consistent policy, access controls will frustrate users. Access that works at the office can fail on a VPN connection that leaves a user without any clear explanation what went wrong. The resulting confusion and troubleshooting prevent users from performing their tasks. Furthermore, attackers can find ways to bypass MFA by targeting OAuth to link malicious applications with API access. These attacks occur inside encrypted HTTPS channels and are difficult to identify on endpoints (FireEye, 2017). Managing and auditing identity controls (e.g., Google G Suite, AWS IAM, OKTA, etc.) is critical, and Trust Zone policies should include these scenarios.

Additionally, an anti-pattern emerges for some of the cloud services that provide an IP whitelist option to bypass MFA. For example, Salesforce (2017) offers a “Trusted IP Range” option to avoid login challenges such as one-time codes. In this case, an attacker with a presence on a network listed as a “Trusted IP” can attack the Salesforce data with a weak or compromised password. Enabling features like a “Trusted IP Range” exchange the stronger identity-based access for weaker IP-based access. A comprehensive Trust Zone policy should include details on the proper use of IP restrictions in all forms.

Adopting a Zero Trust model simplifies policies as the emphasis is the identity component and removal of IP restrictions.

5. Conclusion

In a defense-in-depth mindset, IP controls will always reduce risk. A bastion host on the Internet with SSH open to 0.0.0.0/0 is less secure than one restricted to a CIDR range. That same SSH bastion host is most secure using SSH Certificates, disabling password and key-based authentication, and routinely patching for vulnerabilities. The IP restrictions represent a diminished return after strong identity-based controls and vulnerability management are present. Trends in remote workers, cloud services, and DevOps have put unprecedented pressure on traditional IP perimeters and boundary access controls. Security organizations can be overwhelmed with many requests for firewall rule changes. All the while, threats increase from “internal” sources as IP-perimeters fail with compromised endpoints or malicious insiders.

Engineering-savvy cloud companies such as Google, Facebook, Netflix, Lyft, and Uber have built strong identity-based systems to enable access to all resources – internal or external. These efforts rely on a Zero Trust model that espouses all IP addresses are inherently untrustworthy. Zero Trust forces a mindset of securing the systems and identities regardless of their network address. All organizations should look closely at the benefits of a Zero Trust model. An IP-based perimeter is valid, but ultimately porous without sufficient internal identity-based protections.

Most organizations will not have the scale issues that large cloud organizations like Google, Facebook, and Netflix have reached. LDAP and Kerberos with MFA are sufficient for organizations with thousands of hosts and traditional data centers with some PaaS and IaaS cloud usage. These scenarios require security teams understand both IP and identity boundaries to avoid exposing systems and services. No organization or white paper reviewed advocates for the complete removal of all IP address restrictions. Instead,

John Becker, jbecker42@gmail.com

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

traffic passes through access control points with robust logging and identity management. The removal of IP-based access controls is the result of a journey that begins with a Zero Trust philosophy.

When implemented correctly, identity-based perimeters simplify access to services for users while simultaneously improving the overall security posture. The first step is to create identity trust zones for accessing systems and data. The design should include device and user identity criteria for accessing the various trust levels. The trust model should align with the inventories and configurations documented in Critical Security Controls 1-3. The identity perimeter emerges at CSC 5 with administrative privileges after vulnerability remediation and patching in CSC 4. Organizations should only remove an IP-Based perimeter when required for business use-cases; and then only after establishing robust inventory and vulnerability management systems in conjunction with identity-based access systems.

References

Allen, J. M. (2016, January 31). *Implementing the Critical Security Controls in the Cloud*.

Retrieved June 9, 2017, from <https://www.sans.org/reading-room/whitepapers/cloud/implementing-critical-security-controls-cloud-36725>

Amazon. (2017a). Amazon API Gateway. Retrieved August 5, 2017, from

<https://aws.amazon.com/api-gateway>

Amazon. (2017b). AWS Lambda - Serverless Compute. Retrieved August 5, 2017, from

<https://aws.amazon.com/lambda>

Amazon. (2017c). Endpoints for Amazon S3 - Amazon Virtual Private Cloud. (2017).

Retrieved August 8, 2017, from

<https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-endpoints-s3.html>

Amazon. (2017d). Temporary Security Credentials - AWS Identity and Access

Management. Retrieved August 11, 2017, from

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html

Amazon. (2017e). VPC Flow Logs - Amazon Virtual Private Cloud. Retrieved July 27,

2017, from <https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html>

Amplify Partners. (2014, March 14). *Netflix VP of It on the Future of Infrastructure*.

Retrieved June 17, 2017 from <http://www.amplifypartners.com/netflix-vp-of-it-on-the-future-of-infrastructure/>

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

- AWS. (2016, February). *Architecting for the Cloud*. Retrieved June 10, 2017, from https://d0.awsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf
- Beyer, B., & Ward, R. (2014). *BeyondCorp: A New Approach to Enterprise Security*. ;login:, 39(6), 6-11. Retrieved June 8, 2017, from <https://research.google.com/pubs/pub43231.html>.
- Beyer, B., Jones, C., Petoff, J., & Murphy, N. (2016). Chapter 2. The Production Environment at Google, from the Viewpoint of an SRE. In *Site Reliability Engineering: How Google Runs Production Systems*. Sebastopol, CA: O'Reilly.
- CIS. (2016, August 31). *The Center for Internet Security Critical Security Controls for Effective Cyber Defense v6.1*. Retrieved June 10, 2017, from <https://www.cisecurity.org/controls/>
- Cittadini, L., Spear, B., Beyer, B., & Saltonstall, M. (2016). *BeyondCorp Part III: The Access Proxy*. ;login:, 41(4), 28-33. Retrieved July 12, 2017, from <https://research.google.com/pubs/pub45728.html>
- Dutra, M. (2016, September 12). *Scalable and secure access with SSH*. Retrieved June 17, 2017 from <https://code.facebook.com/posts/365787980419535/scalable-and-secure-access-with-ssh>.
- Facebook. (2017). Facebook Code | Facebook. Retrieved July 16, 2017, from <https://code.facebook.com/hardware>
- Fail2ban. (2017). Fail2ban. Retrieved August 10, 2017, from https://www.fail2ban.org/wiki/index.php/Main_Page

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

2
1

- Filkins, B. (2017, May). *Network Security Infrastructure and Best Practices*. Retrieved June 11, 2017, from <https://www.sans.org/reading-room/whitepapers/analyst/network-security-infrastructure-practices-survey-37795>
- FireEye. (2017). *M-Trends 2017 Cyber Security Trends*. Retrieved June 13, 2017, from <https://www.fireeye.com/current-threats/annual-threat-report/mtrends.html>
- Forrester. (2013, April 13). *Developing a Framework to Improve Critical Infrastructure Cybersecurity*. Retrieved from http://csrc.nist.gov/cyberframework/rfi_comments/040813_forrester_research.pdf
- Gartner. (2017, February 22). *Gartner Says Worldwide Public Cloud Services Market to Grow 18 Percent in 2017* Retrieved June 15, 2017, from <https://www.gartner.com/newsroom/id/3616417>
- Hashicorp. (2017). *Signed SSH Certificates - SSH Secret Backend - Vault by HashiCorp*. Retrieved August 4, 2017, from <https://www.vaultproject.io/docs/secrets/ssh/signed-ssh-certificates.html>
- Jones, J. (2015, August 19). *Gallup Poll - In U.S., Telecommuting for Work Climbs to 37%*. Retrieved July 12, 2017, from <http://www.gallup.com/poll/184649/telecommuting-work-climbs.aspx>
- Lewis, R. (2016, June 30). *OSCON 2016: Netflix BLESS* [Video file]. Retrieved from <https://www.youtube.com/watch?v=JwLGsWYVjqU>
- Microsoft. (n.d.). *Optimizing your network for Skype for Business Online - Skype for Business*. Retrieved July 17, 2017 from <https://support.office.com/en->

John Becker, jbecker42@gmail.com

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

2
2

- [us/article/Optimizing-your-network-for-Skype-for-Business-Online-b363bdca-b00d-4150-96c3-ec7eab5a8a43](https://medium.com/uber-security-privacy/introducing-the-uber-ssh-certificate-authority-4f840839c5cc)
- Moody, P. (2017, February 8). *Introducing the Uber SSH Certificate Authority – Uber Security + Privacy – Medium*. Retrieved from <https://medium.com/uber-security-privacy/introducing-the-uber-ssh-certificate-authority-4f840839c5cc>
- Netflix. (2017a). GitHub - Netflix/bless: Repository for BLESS, an SSH Certificate Authority that runs as a AWS Lambda function. Retrieved July 19, 2017, from <https://github.com/Netflix/bless>
- Netflix. (2017b). Netflix, Inc. · GitHub. Retrieved July 19, 2017, from <https://github.com/Netflix>
- Neuman, C., Yu, T., Hartman, S., & Raeburn, K. (2005, July). *The Kerberos Network Authentication Service (V5)*. Retrieved from <https://www.ietf.org/rfc/rfc4120.txt>
- Obregon, L. (2015, December 2). *Infrastructure Security Architecture for Effective Security Monitoring*. Retrieved July 18, 2017 from <https://www.sans.org/reading-room/whitepapers/bestprac/infrastructure-security-architecture-effective-security-monitoring-36512>
- Peck, J., Beyer, B., Beske, C., & Saltonstall, M. (2017). *Migrating to BeyondCorp: Maintaining Productivity While Improving Security*. ;login:, 42(2), 49-55.
Retrieved from <https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/46134.pdf>

Trust No One: A Gap Analysis of Moving IP-Based Network Perimeters to A Zero Trust Network Architecture

2
3

Salesforce. (n.d.). Set Trusted IP Ranges for Your Organization. Retrieved July 17, 2017

from

https://help.salesforce.com/articleView?id=security_networkaccess.htm&type=0

Saltonstall, M., Beyer, B., McWilliams, J., & Osborn, B. (2016). *BeyondCorp: Design to Deployment at Google*. ;login:, 41, 28-34. Retrieved June 8, 2017, from

<https://research.google.com/pubs/pub44860.html>.

Steipp, C. (2017, April 11). *Blessing your SSH at Lyft – Lyft Engineering*. Retrieved from

<https://eng.lyft.com/blessing-your-ssh-at-lyft-a1b38f81629d>

Trustwave. (2017). ModSecurity: Open Source Web Application Firewall.

Retrieved August 11, 2017, from <https://modsecurity.org>

Uber. (2017). GitHub - uber/pam-ussd: uber's ssh certificate pam module. Retrieved July

20, 2017, from <https://github.com/uber/pam-ussd>

Ullrich, J., Cropper, J., Frühwirt, P., & Weippl, E. (2016). The role and security of

firewalls in cyber-physical cloud computing. *EURASIP Journal On Information Security*, 2016(1), 1-20. doi:10.1186/s13635-016-0042-3

Verizon. (2017, May 02) *DBIR: Understand Your Cybersecurity Threats*. Retrieved June

13, 2017, from <http://www.verizonenterprise.com/verizon-insights-lab/dbir/2017/>