Secure Network Connections

An evaluation of the U.S. Trusted Internet Connections program

AWS Government Handbook | Secure Network Connections

August 2017

Contents

AWS Government Handbook | Secure Network Connections

August 2017

Purpose

As a global first mover, the USG has invested considerable time in developing approaches to network perimeter security. As discussed below, however, while these approaches have been operating in the traditional IT space, additional innovation and iteration is necessary to better align with newer, non-traditional technologies, such as cloud.

This document discusses the following:

- A Summary of lessons learned from AWS's work with various U.S. Government (USG) agencies, including the Department of Homeland Security (DHS).
- The various USG federal-wide secure network connections programs, and focuses on one of the programs known as the "Trusted Internet Connections" (TIC) initiative.
- The AWS policy position and recommendations for how governments can consider establishing or enhancing their cloud-based network perimeter monitoring capabilities.

Introduction

With a vast number of government systems connecting to the public Internet, the USG is concerned about monitoring, controlling, and securing data that flows between its private networks and the outside world. Accordingly, the USG has established three initiatives to maintain visibility and control over data flows with the Internet and, in some cases, with cloud service providers (CSPs):

1) Civilian Trusted Internet Connections (TIC)- applies to all connections with external networks;
2) Department of Defense (DoD) Internet Access Point (IAP) - designed for general Internet traffic; and
3) DoD Cloud Access Point (CAP) - controls traffic flows to multi-tenant commercial clouds.

While these initiatives were originally designed to address traditional network protection paradigms, they are not well suited to address all technologies, such as cloud computing. The high-level goals of the initiatives — visibility and control over flows between networks—are reasonable, but there needs to be greater flexibility to account for how the initiatives are implemented based on different technologies. This will allow government customers to leverage technologies like the cloud to achieve much higher levels of security.

AWS has previously developed a whitepaper that addresses how the concept of a TIC can be used to accommodate the USG's network security initiatives in a cloud computing environment. Although TIC is not a perfect solution in light of the inflexibility presented by the initiatives[1], modernization is necessary to ensure government customers have access to the most agile, innovative, and cost-effective means of achieving their technology and information security goals. Today's TIC/IAP/CAP initiatives are limited in scope and prescriptive with respect to mechanisms and implementations. To achieve the desired security state, while accommodating the innovations of technologies like hyperscale commercial cloud, initiatives and policies should focus on higher-level strategies, principles, and goals that focus on an outcome-based approach, while providing agencies with greater flexibility in how to achieve those outcomes.

[1] See, for example, AWS's White Paper from February 2016 on how TIC requirements can be met: https://d0.awsstatic.com/whitepapers/compliance/Guidance_for_Trusted_Internet_Connection_TIC_Readiness_on_AWS.pdf

What is the history of TIC?

The White House's Office of Management and Budget (OMB) started the TIC initiative in November, 2007 in an attempt to improve the USG's security posture and incident response capabilities by reducing and consolidating individual external network connections used by federal agencies. TIC includes passive network monitoring via the EINSTEIN 2[2] intrusion detection capability to improve situational awareness of external network connections and agency network perimeter security.[3] EINSTEIN 2 detects specific custom signatures of known or suspected threats and alerts the US Computer Emergency Readiness Team (US-CERT), which then analyzes malicious activity occurring across the entire federal enterprise.

DHS's objectives for TIC are for agencies to know:

- Who is on my network?
- When is my network being accessed and why?
- What resources are being accessed?

This information allows DHS and other agencies to respond to inappropriate activity and ensure that only authorized individuals are performing authorized activities. TIC was designed to perform intrusive network analysis of all inbound and outbound traffic through agency networks to identify specific signatures or pattern-based data and uncover behavioral anomalies, such as botnet activity. A core technological objective to facilitate this near real-time monitoring is to route all traffic through EINSTEIN devices hosted at a limited number of network transit points with large bandwidth capacity. In so doing, the USG made a tradeoff in favor of perceived network security to the detriment of network efficiency. TIC was a long-standing Cross Agency Priority[4] for federal agencies and agencies report on their work under the TIC in annual OMB reports (Federal Information Security Modernization Act report) to Congress.

As a point of background on the other USG initiatives, the DoD also uses the IAP program to provide general Internet access from the Department of Defense Information Network (DODIN) and the CAP program to monitor and control traffic flows between multi-tenant commercial clouds and the DODIN. Both of these programs are based on traditional network protection paradigms and also suffer from being too inflexible and narrowly defined to support newer technologies like cloud. The high-level goals—visibility and control over network flows between networks—are still

[2] EINSTEIN 1 analyzed network flow information from participating agencies to observe potentially malicious activity. EINSTEIN 2, the second iteration, is a passive, automated system that incorporates intrusion detection based on predefined attack signatures. It relies primarily on commercial tools and is able to alert US-CERT of malicious activity.
[3] TIC was identified as part of the President's Comprehensive National Cybersecurity Initiatives (NSPD-54/HSPD-23)-
Initiative #1 Manage the Federal Enterprise Network as a single network enterprise with Trusted Internet Connections.
[4] TIC, identity management (i.e. HSPD-12 PIV smartcard credentials), and continuous monitoring comprise the top three cross agency cyber security priorities.

reasonable, but the particular means chosen to accomplish those goals should be less prescriptive to provide agencies with flexibility in achieving the goals.

How does TIC work?

A TIC provides three basic functions:

- Network connectivity between an agency network and a CSP
- Network traffic monitoring
- Network security

All federal agency external connections must be routed through an OMB-approved TIC. Federal agencies are required to participate in the TIC initiative either as a TIC Access Provider (TICAP) or by contracting services with an approved Managed Trusted Internet Protocol Service (MTIPS) provider.[5] TIC includes mandatory critical capabilities that are performed by the agency and MTIPS provider. In the current version of TIC, the EINSTEIN 2 intrusion detection devices are deployed at each TICAP and MTIPS and the agency establishes a Memorandum of Understanding (MOU) with DHS to deploy EINSTEIN capabilities to federal systems.

What is the TIC Overlay?

DHS's primary objectives with TIC are incident response and ongoing visibility into anomalous network activity. Today, network flow[6] data (also referred to as "netflow" data) gives DHS insight into the majority of the .gov enclaves. Netflow data is "non-content data" and contains information on network flow, such as source IP, destination IP, logon, logoff, and provides information around content (e.g. who's connected to who and for how long). With commercial cloud, DHS is concerned about losing insight into network flows and security where federal employees remotely access agency resources via the Internet. The goal of the netflow requirement, a TIC Overlay focal point, is to regain some of that insight where .gov traffic is going straight to the cloud service provider. DHS considers netflow and similar monitoring to be a foundational capability similar to how it is used in the traditional network environment today.

AWS was among a small group of CSPs that engaged directly with the White House, DHS, and the FedRAMP Program Management Office on the challenges with TIC implementation in the cloud. One of the main concerns was related to the separate requirements and security assurance vehicles for FedRAMP, TIC, and agency specific requirements such as the DoD Security Requirements Guide. These sometimes competing (and overlapping) requirements present duplicative and

[5] MTIPS allow agencies to physically and logically connect to the public Internet and other external connections in compliance with TIC. TICAPs and MTIPS have baseline security capabilities including firewalls, malware policies, and network/security operation centers.
[6] A network flow is identified as a unidirectional stream of packets between a given source and destination-both are defined by a network-layer IP address and transport-layer source and destination port numbers. Specifically, a flow is identified as the combination of the following key fields: Source IP address, Destination IP address, Source port number, Destination port number, Layer 3 protocol type, Type of service (ToS), and Input logical interface.

unreasonably burdensome barriers for contractors and agencies trying to implement new technologies.

Following industry consultation, DHS created a TIC Overlay to evaluate TIC security capabilities for CSPs as part of the FedRAMP assessment. DHS mapped the TIC critical capabilities to the FedRAMP baseline and determined that compliance with FedRAMP satisfied approximately 90% of the TIC capabilities. DHS also conducted TIC Overlay pilots with a select group of CSPs, including AWS. The results of the pilots were shared with OMB and the CIO Council, which concluded that *"the focus on network-level security also misses important modern security data inputs such as end-host and application-level logs. As demonstrated by the TIC Overlay pilot program implemented by DHS in coordination with OMB in 2016, there are other methods of implementing security controls that can improve security without significant additional costs."*[7]

What is the challenge with TIC in the cloud?

The current design of these network and data protection mechanisms rely on flowing all traffic bound for "the Internet" (even if not literally on the Internet, see below) or through commercial cloud infrastructure using a relatively small number of hardware devices designed to centralize the USG's management, monitoring, and control. It is precisely this centralized approach, however, that is creating barriers to scalability and elasticity to accommodate the demand for secure access. The current approach results in USG systems being either over-provisioned (have more capacity than needed) or, what is more common, under-provisioned, which results in system bottlenecks and choke points. This slows down the USG's ability to access and deliver services. Moreover, these traditional perimeter systems are largely limited to the observation of outcomes (e.g., network flows), while CSPs offer the opportunity to centrally monitor and proactively address security vulnerabilities (e.g., instance launches, changes in security configurations, API execution data,). By using software-defined and software-driven cloud technology, traditional visibility and control objectives can be enhanced in a distributed, scalable fashion, without increasing management complexity as would be required by deploying larger numbers of physical systems.

The range of agency requirements governing network traffic requirements pose problems across the federal government, as they create layers of complex compliance requirements, which burden service providers without providing the USG with any real security protections. For example, a FedRAMP/High CSP that has gone through rigorous assessments before receiving its FedRAMP accreditation has been assessed and determined to meet stringent guidelines backstopped by NIST, including guidelines surrounding the CSP's network security capabilities. Nevertheless, government agencies often will require the use of mandated network security requirements like the TIC or CAP, which effectively then prevent the agencies from leveraging the CSP's accredited

[7] https://www.cio.gov/2017/03/06/new-cio-council-report-on-developer-platforms-and-common-apis-and-services/

services absent additional steps by DHS, the agency, and/or CSP. This is true even where data flows only between private government networks and virtual private networks provided by the CSP. This presents inefficient, duplicative, and impractical barriers that provide little or no improvement in the agency's security posture.

The following table illustrates some of the challenges with the current system and the opportunities afforded by a cloud-centric modernization effort.

Issue	Current Models	Proposed Models
Level of policy and specification relative to valid goals	Too specific, outdated, and tends to force inflexible technologies and methodologies to achieve desired goals.	Commercial cloud services can provide agencies flexibility to leverage new innovative and emerging technologies to keep pace with rapidly evolving threats and improve security and mission support while meeting or exceeding government requirements Regulations or initiatives should focus on a model that is outcome-oriented instead of prescriptive on static technologies They should also allow for a wider variety of technical solutions that provide the 'best fit' for the desired outcome and the ability to adapt to evolving threats faster.
When is a TIC/IAP/CAP required?	For data moving to and from "the Internet," often interpreted to include private connections to private cloud networks.	Only when exiting / entering from trusted to untrusted networks.
Can TIC/IAP/CAP requirements match different levels of data sensitive for the relevant systems?	No; the current model has no way of distinguishing between low, moderate, and high levels of data sensitivity at the network level.	Yes; workload segregation and network micro-segmentation allow applications with different levels of data sensitivity to be treated differently (more efficiently and appropriately) at the network monitoring and control level. AWS services such as Amazon Virtual Private Cloud, Security Groups, AWS Identify and Access Management, AWS Key Management Service and Access Control Lists allow workloads to be categorized, encrypted, and segmented from one another even within the same virtual datacenter and network. This means data can be processed at the appropriate level for that workload reducing the overall capacity burden for processing highly sensitive data.

Issue	Current Models	Proposed Models
Use of classified signatures for detection/prevention such as the EINSTEIN program	Can be run in-band but at high cost, performance bottle-necking, and lack of scalability and flexibility.	Appropriate signals can be sent off-cloud to a classified enclave for final determination, and flow control responses implemented in-cloud within milliseconds of a threat detection and decision to block.
Architectural pattern (general)	Send all network flows from all clients, servers, and services through a limited number of inflexible deployments of expensive equipment with expensive maintenance contracts.	Allow network information gathering and management as appropriate to the service offering, without losing insight and control of critical data. This can be accomplished using the elastic infrastructure of the cloud to scale with traffic instead of producing choke points that hinder processing capacity. Also, policy-based security controls and the innate API architecture of cloud services allow for deep visibility, command, and control over data transit and processing without the high level of overhead required for onboarding processing cloud based traffic to alternate infrastructure.
Architectural pattern (virtual machine services)	Same as above.	Allow horizontal scaling of network flow intelligence, monitoring, and control via auto-scaling fleets of packet-scrubbing cloud nodes (semi-centralized), or agent-based network monitoring on modern, fully validated operating systems (fully decentralized). Using AWS services such as AWS Elastic Load Balancing, AWS Auto Scaling and AWS CloudFormation, a packet inspection fleet can be decentralized and independently scaled to keep pace with specific demand hot spots on the network.
Architectural pattern (other network-hardened, cloud-scale APIs and managed services).	No options.	Use the built-in identity management and logging/ auditing features of perimeter-less, "Internet-hardened" services such as object storage (*e.g.*, Amazon S3), NoSQL databases (*e.g.*, Amazon DynamoDB), and messaging services (e.g., Simple Queuing Service). Eventually, USG APIs and services will also be capable of perimeter-less operations.

Our recommendations for TIC modernization to support operational visibility

Emerging technologies necessitate a reconsideration of how to modernize network protection rules. Any new initiatives should define high-level goals and requirements for network and data protection without being prescriptive on specific implementation details or technologies. This allows for technology and implementations to evolve and use state-of-the-art and innovative solutions as long as they achieve the stated goals. Goals should also be developed in conjunction with industry experts who can bring to the government "the art of the possible" in this fast-developing area of IT security. The changes will enable faster and more comprehensive adoption of cloud and other state-of-the-art technologies, allowing agencies to achieve their goals of modernization and efficiency while improving security operations.

Based on our understanding of the objectives for network monitoring, we offer the following recommendations on how to modernize these capabilities to better position government customers to innovate using new technologies while maintaining important security goals:

1) Much of the netflow data (e.g., IP address, logon, logoff) is available to agencies today through on-demand or near real-time audit logs based on the customer's implementation of a commercial cloud solution. Remote access traffic can be captured in application logs, and agency mobile device management deployments allow further insight into mobile traffic.

   The starting point is for agencies to determine what information they can obtain through a certain CSP or cloud-based solution. Different cloud service models will provide customers with different levels of access to data. For example, SaaS/PaaS logs may contain application and platform (e.g., database)-related log content; while IaaS logs may contain more infrastructure and network level data. The most valuable information is generally available through application-level, cloud-based audit logs that should be available to customers on-demand. No matter the implementation (commercial cloud or otherwise), the agency customer should always own, and have access to, their own audit logs. If customers need to monitor certain .gov customer actions, the customer can gain access to the logs and perform the monitoring without CSP interaction. Audit log solutions such as Amazon VPC Flow Logs, Amazon CloudWatch Logs, and AWS CloudTrail along with other third party open source or commercial solutions can support this level of visibility today. Amazon Macie also supports insight into abnormal session access and activity as it continuously monitors for data access activity anomalies, and delivers alerts when it detects risk of unauthorized access or inadvertent data leaks.

2) To the extent that network-edge protection is required, customers can consider the use of cloud-native capabilities in conjunction with open source or commercial products, services, and modern protection techniques. The elasticity of the cloud enables "right sizing" of edge network protection capabilities. That is, cloud-native architectures add and remove capacity in response to demand.

   A good approach may be to use a horizontally scaled, fleet-based edge protection strategy that works well in cloud environments. In this system, TICAP and MTIPS operators could offer managed TIC services to Federal workloads in the cloud. These providers would use commercial or open source products with unclassified indicators to deliver equivalent or superior outcomes to that of the EINSTEIN program. For classified data, network traffic could be mirrored to an off-site, secure facility for out-of-band evaluation. As necessary, this off-site infrastructure would deliver control messages back to the cloud-based inline fleet to impact data flow.

3) With the rise of Zero Trust networks and Software Defined Perimeters, industry is increasingly moving towards "boundary-less" operation that assumes no trust in the underlying infrastructure. The security boundary is reduced to a minimal footprint around the data and data processing. In this approach, the ability to establish network communication between two nodes in the network is driven by identity-based network security. That is, any given network connection is allowed based on the identity of the initiating principal and its permissions to access a resource. This type of capability is available in AWS (e.g., Amazon EC2 security groups and instance roles). There are also an increasing number of commercial, agent-based offerings. This approach shifts the burden for detection and monitoring from the edge of the infrastructure to each individual node of the infrastructure, unlocking tremendous visibility.

   In all three of the above recommended approaches, cloud offers users 360 degree in-depth visibility into both effect (e.g., network traffic flow) and cause (e.g., customer activity). When cloud audit data is fused with network flow data, insights become available that would otherwise be nearly impossible to achieve in a physical data center environment. For example, creation of a new user, a change in network access controls, the creation of a new node in the infrastructure, and a new outflow of data can all be correlated and traced back to the originating event with relative ease.

We are encouraged by the USG's evolution towards accepting innovative, cloud-adaptive solutions to achieve network perimeter monitoring objectives in the cloud. We are committed to ongoing collaboration with the USG and governments worldwide that are evaluating the merits, best practices, and lessons learned from the TIC program.

August 2017

Contributors

Mark Ryland, Director, Solutions Architect
Tim Anderson, Program Manager, Security Business Acceleration
Alan Halachmi, Sr. Manager, Specialists Team
Arash Heidarian, Sr, Corporate Counsel
Min Hyun, Cloud Security Policy Strategist
Jim Jennis, Sr. Solutions Architect