White Paper

Coalition Information Sharing in the Age of Great Power Competition

Introduction

Carl von Clausewitz, the 19th century military strategist, spoke on the importance of lines of communication (LOC) while observing Napoleon. That observation is as relevant today in all warfighting domains, including the cyber domain, as it was then.

As Clausewitz wrote, LOCs were necessary to move critical supplies to the front line, as well as provide an egress route for forces moving away from the front line.

Securing LOCs is paramount. Every military branch has incorporated LOCs into their own domain operational art for mission success. The Army establishes Ground LOCs, the Navy must secure Sea LOCs, and the Air Force identifies Air LOCs into and out of the area of operations (AOR). This concept is also relevant in the cyber domain. Zero trust—as a cybersecurity paradigm—enables cyber operators to securely move critical information around the battlefield, including to and from our coalition partners and everywhere it is needed for mission success.

Coalition Information Sharing Through Cyber Lines of Communication

The cyber domain, where we share information across the coalition and joint force, is unique. Much has been written about the cyber domain vis-à-vis the other domains. In 2011, General Larry Welch, former Air Force Chief of Staff and previous Institute for Defense Analyses President, made the point that every military domain must accomplish certain activities to be effective[1]. Those activities are:

- Passive Defense
- Active Defense
- Exploitation (Operational Prep of the Domain)
- Attack
- Defining and Developing Needed Capabilities

However, according to Welch, in the cyber domain, there is one other activity that makes it unique to the other warfighting domains: constructing. The combination of commercial information technology with unique government infrastructure makes up the cyber warfighting domain. As Welch further cautioned, there is a temptation to construct the cyber domain EVERYWHERE friendly forces will operate. Military-specific networks like Secret IP Router Network (SIPRNet) and uNclassifed IP Network (NIPRNet), military-only clouds for compute and store, and other unique hardware and software components like cross-domain solutions (CDS), are examples of how the cyber domain is constructed for military-specific requirements. The challenge for those that design and develop cyber capabilities is to determine the appropriate mix of general purpose commercial and legacy military-unique cyber infrastructure, all while ensuring critical information is protected.

As an example, the Navy cannot secure all the waterways on the globe. Therefore, the Navy must secure only that portion of the sea where the service must operate depending on the missions being executed. This operating environment is the Sea Line of Communication (SLOC). As the cyber domain is constructed, the mix of commercial vs. government (COTS vs GOTS) should strike the right balance to leverage private industry's innovation against government paramount need to protect national security information.

Zero trust is not only a modern security paradigm, but it also allows cyber operators to secure only that portion of the cyber operating environment needed to exchange coalition and joint information. Like the other warfighting domains, Cyber Lines of Communication (CLOC) creates the secure pathway to exchange information while fully exploiting existing commercial information technology purchased through Internet and Cloud Service Providers.

[1] Cyberspace — the fifth operational domain. (n.d.). https://www.ida.org/-/media/feature/publications/2/20/2011-cyberspace---the-fifth-operationaldomain/2011-cyberspace---the-fifth-operational-domain.ashx

Lessons Learned – Afghanistan Mission Network

The Afghanistan Mission Network (AMN)[2] leveraged a federated network architecture to connect coalition partners onto a common information sharing platform. The AMN architecture required each International Security Assistance Force (ISAF) country to connect to a NATO core network using their own infrastructure. The alternative approach—creating a network enclave dedicated to the mission requirements of the coalition—meant the coalition partners gave up control of their information and data to the country (presumably the US) responsible for the coalition enclave. Even with the benefits of federation, the evolution of the AMN required overcoming numerous technical and policy challenges, including data tagging, guard technology to aggregate the networks, and determining operational procedures for using the AMN to achieve desired mission objectives.

There is a realization by DoD leadership that future conflicts require a coalition network capable of rapidly onboarding coalition partners (both military and nonmilitary) with the flexibility of supporting operations that span disaster response to full-scale war. In a rare moment of unity across the entire DoD, the deputy commanders from each of the regional combatant commands and Special Operations Command signed out the "15-Star Memo" to accelerate the development of a coalition network and deliver it by the end of 2016.

To support Multi-Domain Operations, the US leverages coalition partners in a Mission Partner Environment, or MPE. In this coalition network environment, data is fluidly exchanged between interoperable systems so forces can contribute to a common operational picture from which commanders can make informed rapid decisions and task the appropriate military formations. Designing, developing, and deploying a MPE that enables the level of information sharing needed across the full range of military operations is the goal articulated in the 15-Star memo written in 2015, but is still not available to regional combatant commanders to this day.

Secret and Below Releasable Environment (SABRE)

To fully ealize the vision articulated in the 15-Star memo while incorporating the lessons learned from the AMN, the DoD has set out to develop SABRE. According to Danielle Metz, chief IT strategist for the office of the secretary of defense, SABRE, will "blend the intel aspects as well as the [command and control aspects] together in a cloud-based approach." During the AFCEA TechNet Cyber 2022 conference Metz said, "I think that we have struggled for a very long time on the mission partner environment." Metz went on to comment that, "The combatant commanders have been screaming for the need to be able to seamlessly... collaborate not only internally with themselves, but across the mission partners. And a mission partner can mean anything to anybody depending on where you are located and depending on what that actual mission is."

To fully realize the vision of an MPE, SABRE will need to look beyond networks and focus on the data that must be shared for operational and mission success. The temptation is to look at data sharing in terms of network access through a dedicated network coalition accessible enclave or federation of coalition networks. This network-centric approach will prevent advancing an MPE to full operating capability. The alternatives discussed in the lessons learned from the AMN come from two competing requirements.

[2] Serena, Chad C., Isaac R. Porche III, Joel B. Predd, Jan Osburg, and Brad Lossing, Lessons Learned from the Afghan Mission Network: Developing a Coalition Contingency Network. Santa Monica, CA: RAND Corporation, 2014. https://www.rand.org/pubs/research_reports/RR302.html.

One, is the need for an episodic information sharing capability that is effective for the specific requirements of the contingency, including the participating coalition partners and the type and scale of the information sharing needs from humanitarian operations to full-scale war. The second is the need for a persistent capability that enables regional combatant command planners and the Special Operations Command to begin sharing information quickly. Striking the balance between these two competing priorities is an MPE that is both tailored to the needs of the coalition partners (episodic) and available immediately (persistent).

Cross-domain solutions (CDS) are often discussed as the way forward to achieve the flexibility needed for episodic data sharing while also the availability necessary for persistent MPE data access for coalition partners. The value of a CDS is in enabling connection between different domains which could allow for exchanging information across security boundary levels, like secret to top-secret or from a partner network to a US network. According to Department of Defense Instruction 8540.01, a CDS is deployed to limit access between the two domains based on rules for either the transferring of data or user access to data. A CDS can be used to strike the balance between information security to only those with a need to know with fully enabling our coalition partners for mission success through granting data access because of a need to share. However, a CDS solution that strikes the right mission balance while not being overly complex has eluded the DoD.

The SABRE solution must tackle the competing requirements in an MPE as discussed above: episodic vs. persistent and need to know vs. need to share, while recognizing that the data that will inevitably be needed will originate on the US-only SIPRNet and have to be moved through CDS for coalition partners. However, recasting the challenge not just as granting data access while protecting network infrastructure, as a CDS is designed to do, but granting access to protected data independent of the network the data resides on is needed. This reframing of the MPE challenge to establishing Cyber Lines of Communication from a coalition user to the application access where the data resides greatly simplifies the SABRE solution and enables the MPE vision.

Notional MPE with Cross-Domain Solutions

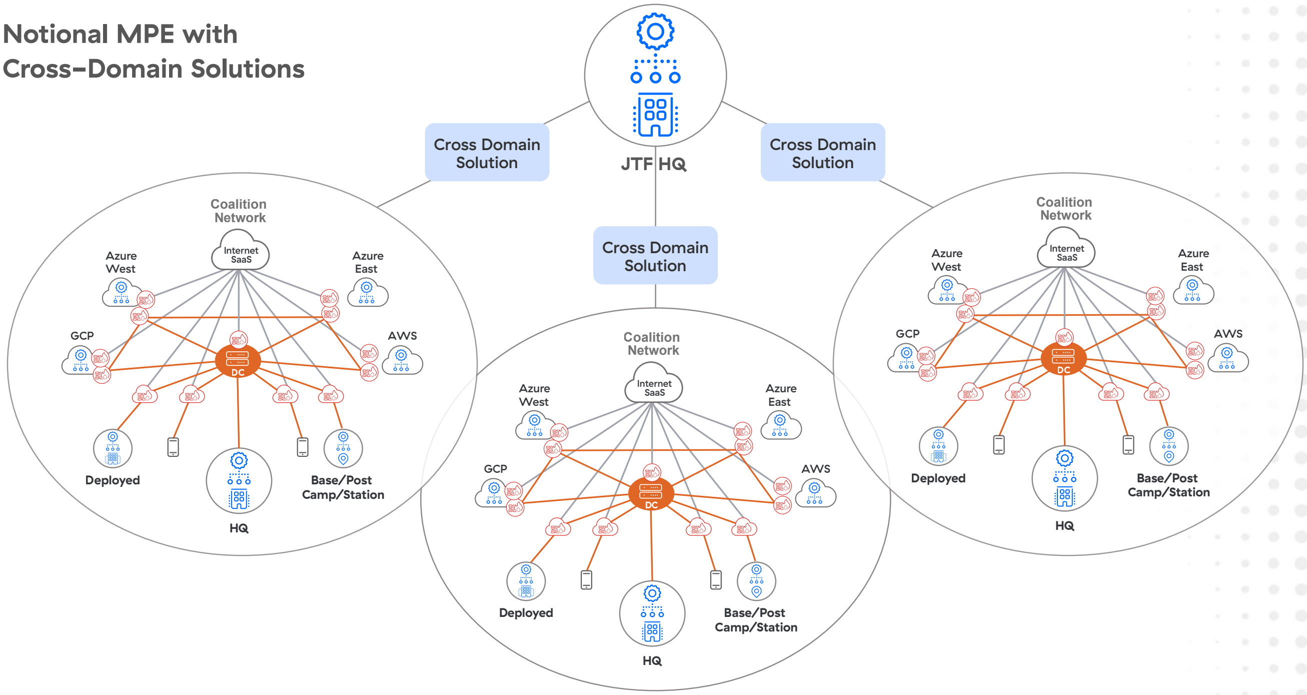

Zero trust is a game-changer for building out the MPE. Implemented as a network independent, device independent, and location independent solution, a zero trust architecture enabling SABRE will allow combatant commanders the flexibility to quickly grant access to critical mission applications regardless of where they reside, from an Impact Level 6 SIPRNet cloud to a coalition partner's unclassified data center.

Zero Trust to Enable SABRE

Striking a balance between need to know and need to share is not a unique challenge to combatant commands integrating coalition partners into regional operations. Over the last 30 years, private industry and the federal government have built "hub-and-spoke networks" where connectivity to the data center requires access through a virtual private network hardware and software to access enterprise resources. They then deployed numerous and disparate security appliances or networking appliances, like routers, switches, and firewalls, to protect the data center where all the applications reside. This was called a "castle-and-moat security" model. The combination of cross-domain security appliances, meant to exchange data between security classification levels as well from US-only to partner networks, within today's operate from anywhere VPN-based security infrastructure is a tremendous and complex challenge.

Legacy coalition interoperable architectures leveraging a CDS to exchange data do not provide an optimal user experience because they introduce unneeded latency and complicated routing to reach wherever the data are hosted. SABRE should focus on granting direct access to coalition data securely regardless of location of the user or where the data resides, regardless of network. Zero trust enables security transformation – a move away from the network-based castle-and-moat model, which is built on the installation, operation, and maintenance of firewalls, VPNs and CDSs. A zero trust architecture (ZTA) transforms both data security and information access – striking the right balance between need to know and need to share.

The challenges caused by legacy network and security architectures are pervasive and long-standing and require rethinking the way connectivity is granted in today's current nation state threats characterized by Great Power Competition. Warfighter information and data are everywhere, from commercial cloud service providers to unclassified/secret/top secret data centers to coalition partner hosted data repositories. This is where ZTA is leveraged—an architecture where no user or application is trusted by default. Zero trust is based on least-privileged access, which ensures that trust is only granted once identity and context are verified, and policy checks are enforced.

NIST SP 800-207, Zero Trust Architecture

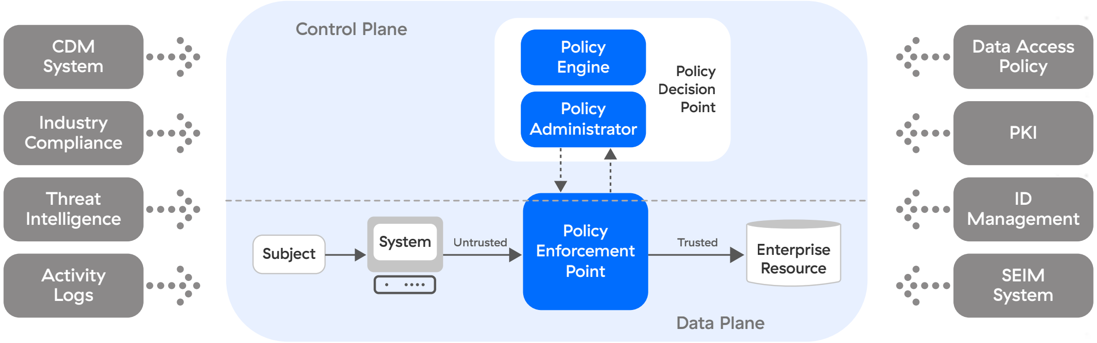

The National Institute of Standards and Technology (NIST) Special Publication 800-207 defines the underlying principle of a zero trust architecture as “no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the internet) or based on asset ownership (enterprise or personally owned)”. The Department of Defense describes seven principles that guide the development of the zero trust reference architecture, with #1 being “assume no implicit or explicit trusted zone in networks.”

This approach treats all network communications as hostile, where communications between users and applications or data are blocked until validated by identity-based policies. It ensures that inappropriate access and lateral movement are prevented. This validation carries across any network environment, where the network location of an entity is no longer a factor and not reliant on rigid network segmentation.

Zscaler Private Access

Today, the security perimeter extends well beyond garrison and tactical military networks to anywhere warfighters connect globally to wherever applications run and data resides. Traditional network security architectures, anchored in on-premises data centers that rely on appliances, have become less effective for modern warfare.

Network-based architectures are also vulnerable due to excessive trust. Remote users connecting from an approved list of IP addresses (via VPN) are assumed to be trusted and are granted access to network resources through a gateway which is often exposed to the internet. On-premises users on the network can move laterally across it. Ultimately, this inherent trust, violating the DoD’s Zero Trust Reference Architecture Principle #1, leads to risk from an over-privileged network access methodology. The security paradigm needs to shift from a static network perimeter and, instead, focus on the entity, resource, and user device, as described in NIST SP 800-207.

Zscaler Private Access (ZPA) is a cloud-delivered zero trust service that uses a distributed architecture to provide fast and secure access to private applications running on-prem or in the public cloud. The service provides access based on four key principles:

- Application access is based on context and should not require network access
- Outbound-only connections make applications invisible to unauthorized users
- Application segmentation connects users to a specific app and limits lateral movement
- The internet becomes the enterprise’s new transport network

When a joint or coalition user attempts to access a MPE application, the user’s identity and device posture are verified via the Zscaler Client Connector software installed on the user device. Policy is checked, and a ZPA Service Edge, either on-prem (or at a tactical or deployed location) or in the cloud, determines where the closest MPE application instance exists. ZPA uses the location of the client and determines the closest application to the user based on reachability to the ZPA App Connector (lightweight VM in the app environment). Lastly, two outbound tunnels, one from the Client Connector on the device and the other from the App Connector, are stitched together by a ZPA Service Edge. All of this takes place automatically and in real time.

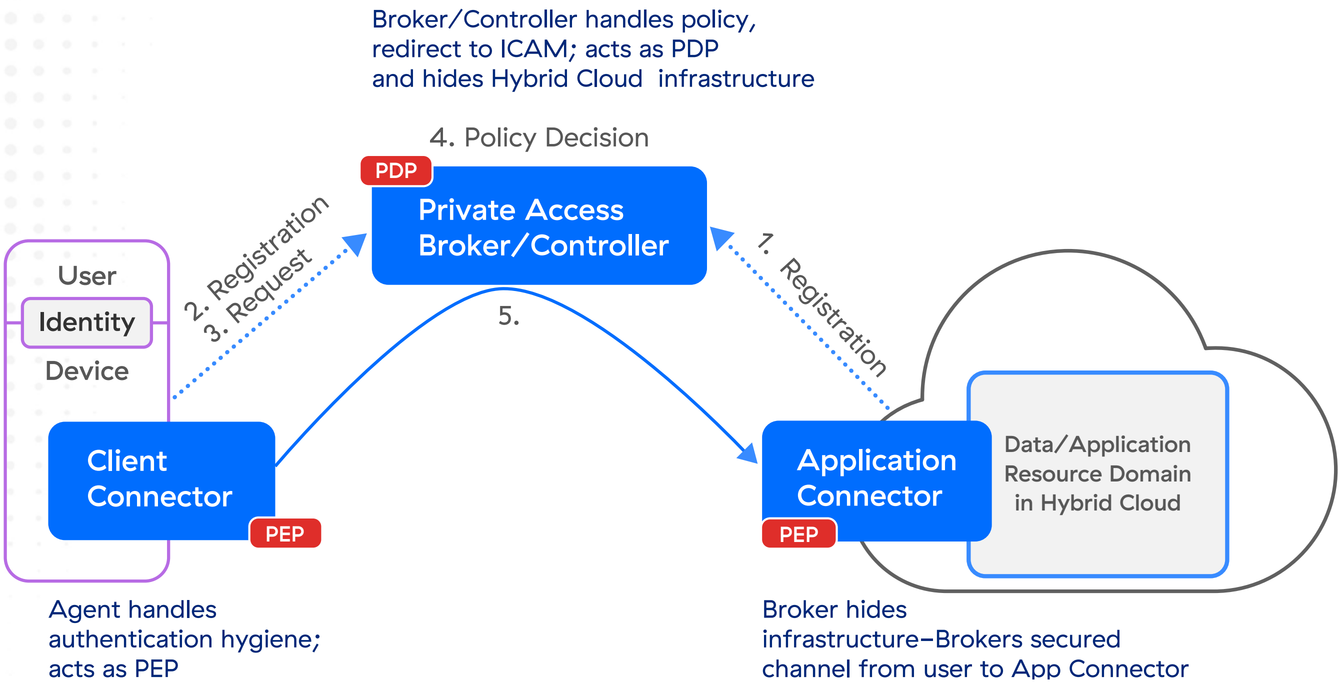
Broker/Controller handles policy, redirect to ICAM; acts as PDP and hides Hybrid Cloud infrastructure
4. Policy Decision
PDP
Private Access Broker/Controller
2. Registration
3. Request
1. Registration
5.
User
Identity
Device
Client Connector
PEP
Application Connector
PEP
Data/Application Resource Domain in Hybrid Cloud
Agent handles authentication hygiene; acts as PEP
Broker hides infrastructure–Brokers secured channel from user to App Connector

Figure 3: DoD ZT Software Defined Perimeter: Zscaler Architecture to create secure Cyber Lines of Communication

1. Joint or Coalition Data Resource, through a Zscaler Application Connector, registers with the broker anywhere in the world—as an example: a coalition on–prem data center, US–operated IL–6 private, or an unclassified commercial cloud.
2. Joint or Coalition warfighter, via their lightweight agent, registers with broker over available commercial or military transport (SATCOM, Cellular 5G, or fiber).
3. Warfighter via their agent requests access to resource
4. Broker checks security controls and access policy
5. Cyber Line of Communication is securely created from warfighter to application

The future join and coalition information sharing platform must strike the right balance between security, agility, and accessibility.

Going Forward: Securing an Agile Coalition and Joint Force

Project managers and acquisition professionals are familiar with the Iron Triangle of cost, schedule, and performance. The theory recognizes that to deliver any system it's quite unlikely that there are unlimited dollars, time, and scope to address every requirement a customer can articulate. One way to manage the three axes is to focus on the impact of how constraining any two elements have on the third. In other words, if keeping costs down and schedule tight is a firm requirement, the functional requirements (scope) of what is delivered will have to be severely limited. If scope and time are hard requirements, the price will not be cheap. Finally, if scope and price are hard requirements, the time to deliver will likely be much longer than desired.

Iron Triangle of Project Management

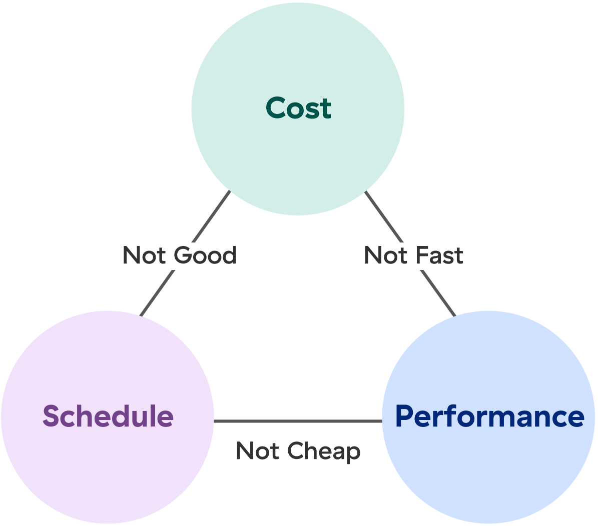

These same trade-offs are found in developing the future Mission Partner Environment and Joint All-Domain C2 (C/JADC2) platform. The three axes for Coalition and Joint Information Sharing are: Security, Accessibility, and Agility. Focusing on security and accessibility prevent SABRE from being agile. This is a good example of the Afghanistan Mission Network: while it achieves the requirements for security and accessibility for NATO coalition partners, it was not agile enough to be used outside of the Afghanistan theater of operations. Focusing on security and agility will require a trade-off with accessibility.

The Agile Combat Employment (ACE) vision requires a SABRE platform that is rapidly deployable and secure, but leveraging legacy US only infrastructure for connectivity and information resource access will introduce challenges integrating coalition partners. In the future, ACE must balance security with agility and accessibility to rapidly onboard coalition partners. Finally, maximizing agility and accessibility will result in an unsecure solution. An existential fight tonight with service members in harm's way or humanitarian assistance/disaster response mission requires maximizing accessibility and agility at the expense of security. Allowing any one axis to dominate the requirements will result in a less than optimal solution. The good news is technology is in place to enable the deployment of a capability that allows rapid sharing of joint and coalition data and to do it in a way that is secure and agile.

Iron Triangle of Coalition Information Sharing

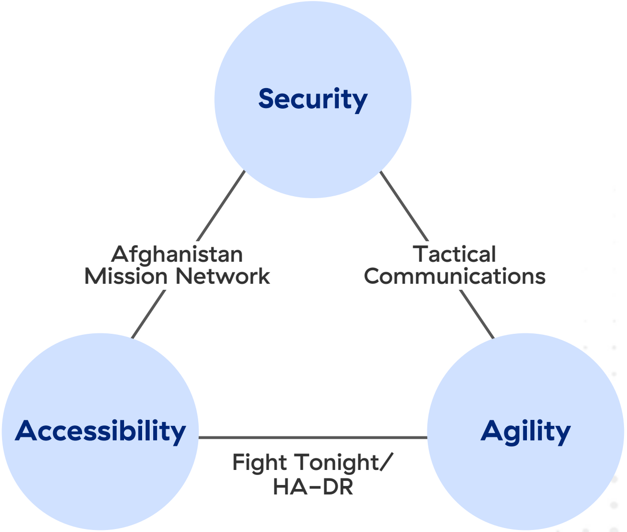

Digging into each of the components of the future SABRE or C/JADC2 platform helps articulate the requirements necessary for the acquisition community to best interact with industry to arrive at the optimal solution.

Security

A zero trust solution that proxies connections and applies security inline resolves the challenge of building out perimeter based security controls around critical data and resources. Zero trust security enables joint and coalition warfighters the ability to enact a risk informed go/no-go decision for any user accessing data and applications in real-time. The decision making available with zero trust creates a dynamic need to know environment that can rapidly adjust to battlefield conditions. Additionally, adding to Zscaler's policy enforcement are the inline technologies that were once part of a sophisticated hardware and software stack of appliances protecting the enterprise boundary. Applying everything from data loss prevention to intrusion detection and prevention, to filtering out dangerous websites from being accessed. Ultimately, security must follow the user, regardless of location and what data and applications they need to access.

The below graphic represents the in-line security process that shows the combination of protecting coalition data in applications both managed by others (Internet and SaaS applications in the cloud) and those applications managed by joint and coalition partners (military data centers, software factories, and apps hosted on commercial IaaS/PaaS). The overarching three phase approach of verifying, controlling, and enforcing is happening at scale, with over 350 billion transactions being secured everyday within the Zscaler cloud.

Realizing the promise of Zero Trust
Per Session Policy Enforcement

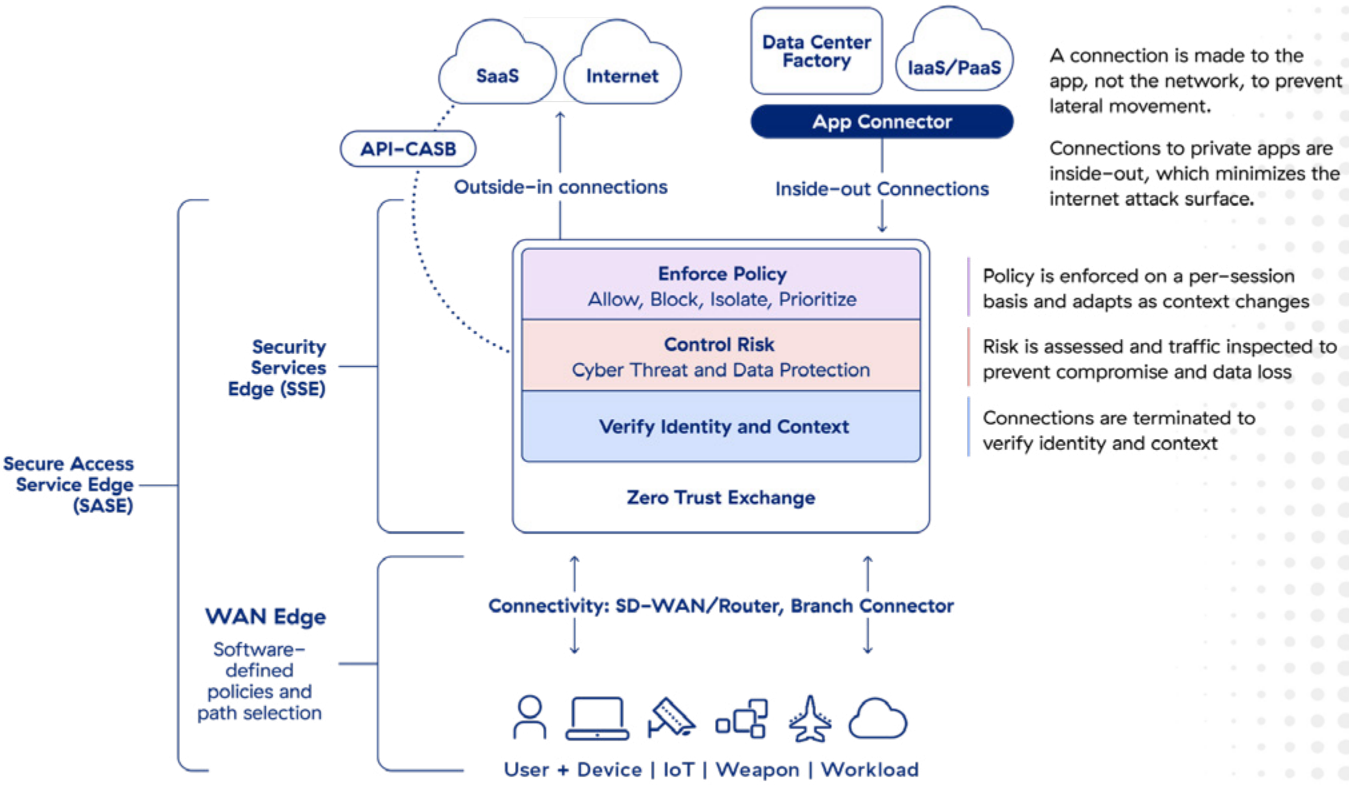
SaaS
Internet
Data Center Factory
IaaS/PaaS
App Connector
API-CASB
Outside-in connections
Inside-out Connections
Enforce Policy
Allow, Block, Isolate, Prioritize
Control Risk
Cyber Threat and Data Protection
Verify Identity and Context
Zero Trust Exchange
Security Services Edge (SSE)
Secure Access Service Edge (SASE)
WAN Edge
Software-defined policies and path selection
Connectivity: SD-WAN/Router, Branch Connector
User + Device | IoT | Weapon | Workload

A connection is made to the app, not the network, to prevent lateral movement.

Connections to private apps are inside-out, which minimizes the internet attack surface.

Policy is enforced on a per-session basis and adapts as context changes

Risk is assessed and traffic inspected to prevent compromise and data loss

Connections are terminated to verify identity and context

Agility

Cloud-Based Command and Control (CBC2) is the right vision for the future SABRE platform. Removing the requirement to palletize and deploy multiple pallets of command and control technology saves deploying military forces, both US and coalition, time and expense when moving into position to execute missions. As mentioned earlier, there is a propensity to "own" the IT infrastructure used to operate in the cyber domain. Through zero trust, military capability developers can allow commercial cloud providers to deliver the necessary infrastructure while security specialists focus on placing "trust zones" close to the data and applications needed for mission execution. Replicating US-owned and operated cloud infrastructure, to add additional layers of security, drives up costs and limits the agility for coalition forces.

Accessibility

The world is more and more connected every day being driven by commercial terrestrial fiber, 5G cellular, and satellite connectivity from vendors like Starlink. This world of ubiquitous internet access, or ubiquitous network transport, coupled with zero trust security that establishes secure Cyber Lines of Communication in any region of the world, greatly enhances the coalition force's accessibility to critical mission data and applications. The future SABRE should harness this reality to take full advantage of available bandwidth, and to do so securely with a zero trust inline security provider like Zscaler. This transport-agnostic approach, that does not rely solely on US-owned and operated networks and network appliances for security, would allow for seamless data exchanges where security follows the user. Gone would be the need to build complex cross-domain solutions to make SABRE "always on" and ready for the ad-hoc nature of today's national security environment—an era of Great Power Competition.

Future SARBE Architectural Components

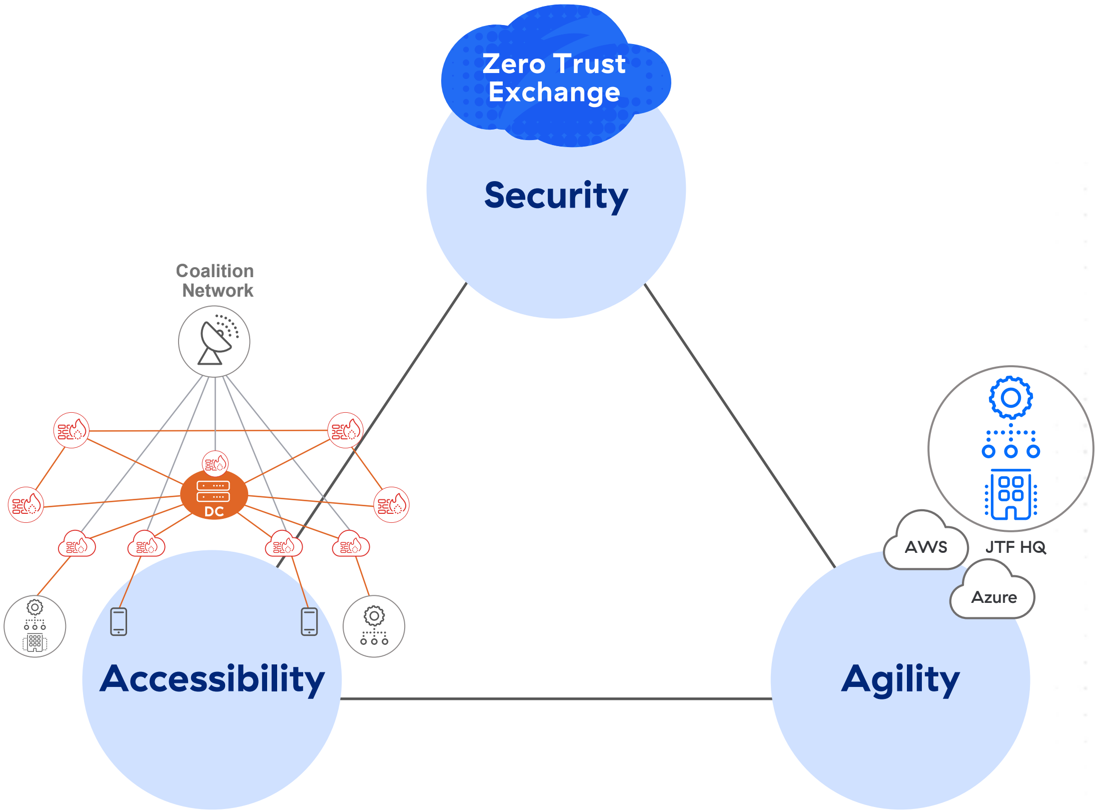

Deployed Mission Partners Accessing Coalition Information Inside the Zero Trust Exchange

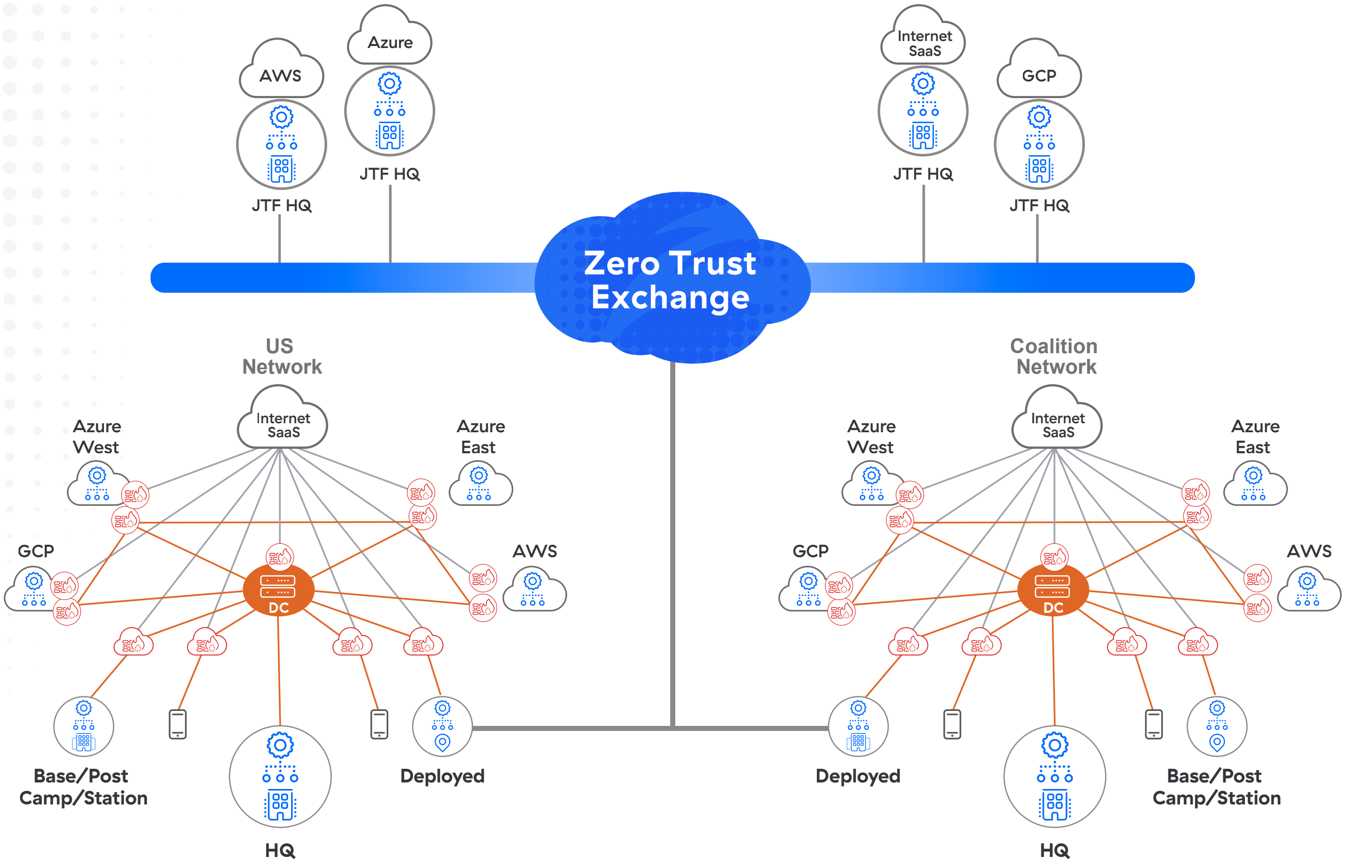

Conclusion

It is undeniable that the Department of Defense recognizes the imperative of joint and coalition warfare, and has so for a long time. Taking advantage of today's modern information technology capabilities like zero trust, cloud computing, and ubiquitous internet access, military planners responsible for command and control and information sharing can focus on building Cyber Lines of Communications. Managing the temptation to control the entire portion of the cyber domain where friendly forces will operate can introduce excessive risk and remove critical agility needed in an era of Great Power Competition. We have learned much from past multinational contingencies, but perhaps the most critical lesson is how important agility is to respond to threats across the entire spectrum of military operations. Zero trust is more than a security solution, it greatly enables connectivity on the battlefield and supports the challenge of building a SABRE program that is secure, agile, and accessible.

+1 408.533.0288 Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134 zscaler.com