# Achieving Zero-Trust Security in Federal Agencies

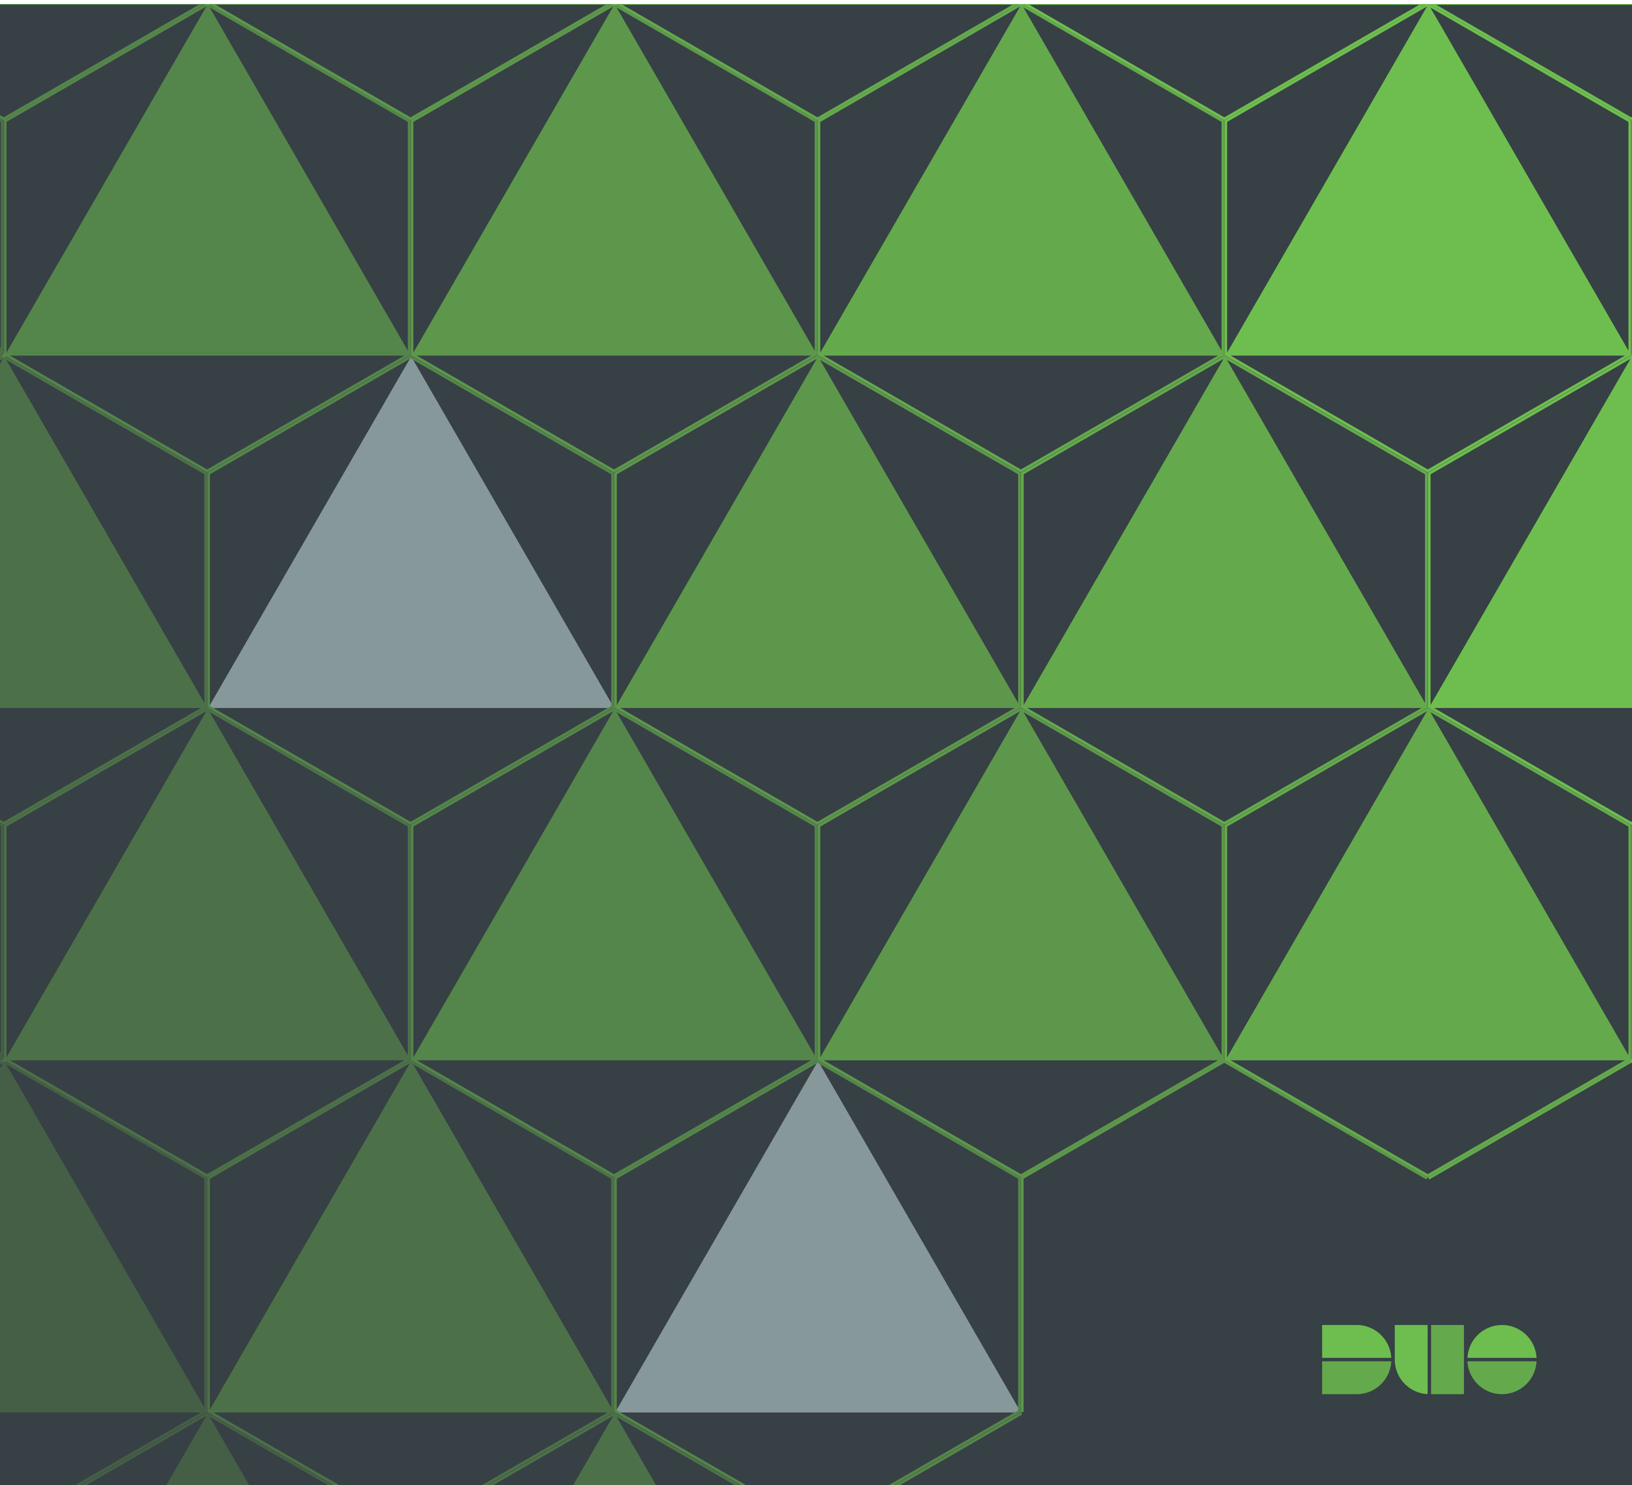

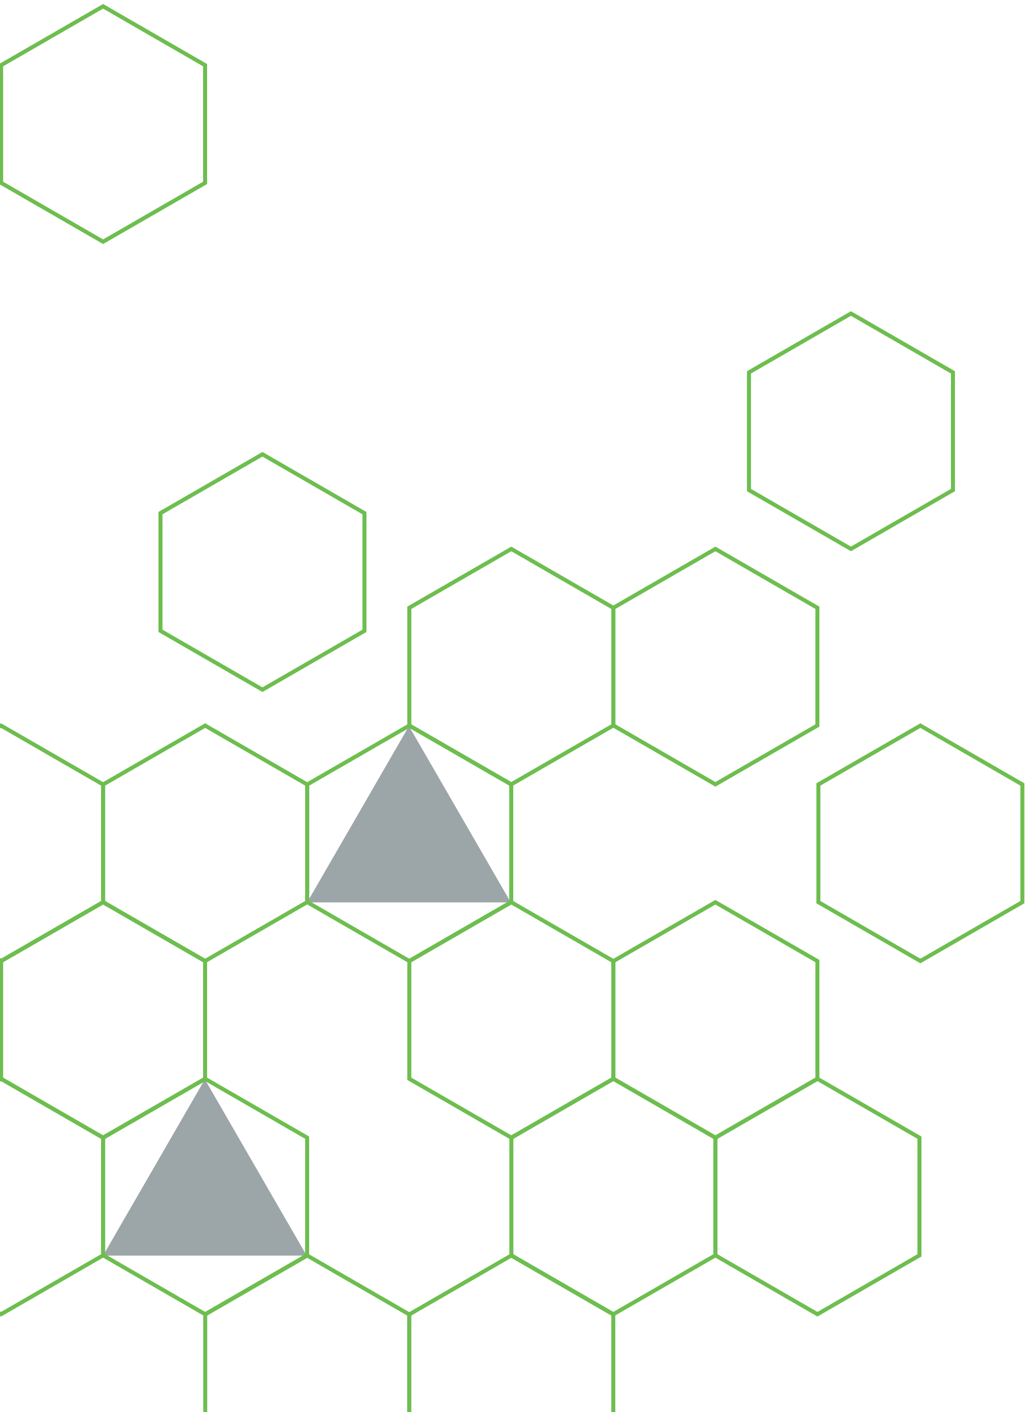

In most federal agencies, terms like “agile” and “responsive” aren’t used very often; like many large organizations, change happens incrementally over long stretches of time. So how can we expect federal agencies to stay up to speed with the constantly evolving digital landscape and keep secure?

Current IT modernization initiatives are challenging federal agencies to implement big changes to their infrastructure at an uncomfortable pace, as they look to accommodate the shift to cloud and mobile. An ideal security solution needs to account for both protecting users and rolling out on a realistic but still workable timeline. That’s where zero trust comes in.

Zero trust is a security model that shifts the access conversation from traditional perimeter-based security and instead focuses on secure access to applications based on user identity, the trustworthiness of their device and the security policies you set, as opposed to the network from where access originates.

The concept of zero-trust security isn’t wholly new, it was posited by the Jericho Group (a group of industry leading CISOs) way back in 2005 as they had thoughts and conversations on how to implement security in this new Internet-centric world we find ourselves in. Since then there have been several companies who have moved to implement zero-trust principles, the most famous one being Google. Google’s BeyondCorp architecture journey started in 2014 and continues to this day. Google did a great job of documenting this journey to provide other organizations guidance and data on “lessons learned.” And while everyone’s journey will be unique, this is great place to start and a great architecture to use a reference.

**In the world of federal agencies, there are four underlying methods that can be coordinated to help achieve zero trust: continuous authentication, device assessments, user controls and application access.**

# Continuous Authentication

One of the key premises behind zero trust is that the perimeter is now anywhere you make an access control decision, and a critical component is verifying user trust. Authentication or user verification needs to be continuous and should happen as a result of multiple coordinated data points instead of a single password or point of failure. Continuous authentication takes a novel approach to that idea by creating a “fingerprint” for each user based on how they use a device – and not any personally identifiable information – and then continually manages access based on the factors that make up that fingerprint.

That isn’t as invasive as it sounds. Continuous authentication simply takes note of how you typically use your device, and then denies access when it senses off-pattern behavior that indicates an unauthorized user. For example, continuous authentication will track things like how a user types or how they use a laptop trackpad, and will allow access when those factors positively identify the right user; if another user happens to be able to log in, but their general usage patterns don’t match the logged in individual, access will be denied.

Continuous authentication is user-specific, but doesn’t rely on privacy-protected information, and it creates a security environment that protects from threats that other approaches don’t account for. If a user has left their workstation while still logged in, continuous authentication prevents anyone else from taking advantage of the open session. Similarly, if a malicious virus or application attempts to spoof a user and execute its own commands, it would be disabled for not meeting all of the criteria established by continuous authentication.

Continuous authentication isn’t perfect, but it can provide a balance between strong security and the usability that current users have become accustomed to, *and* it’s valuable because it creates data points that are difficult to simulate. IT can also work in conjunction with other factors to moderate access, like the state of the device in question – so let’s explore why that’s so important.

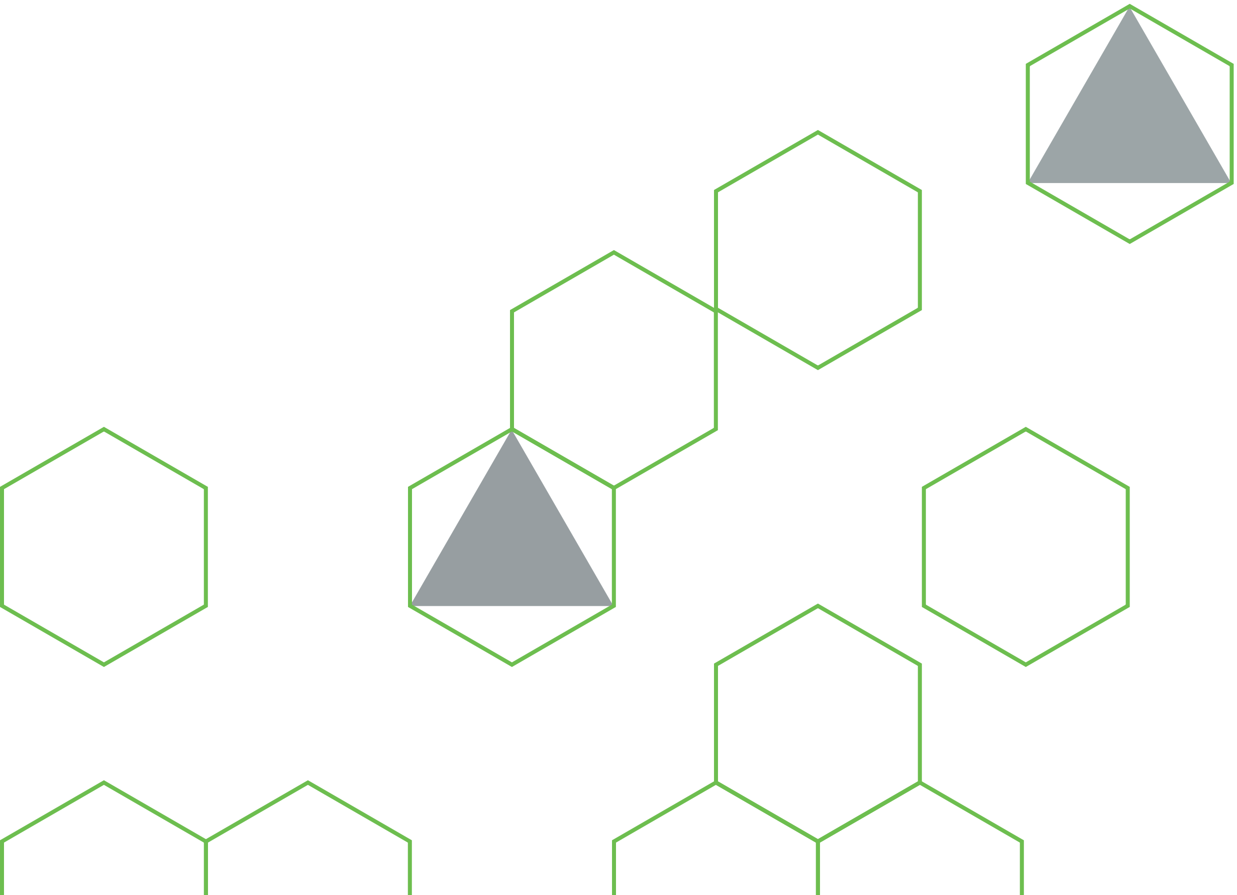

# Device Assessments

The bring your own device (BYOD) revolution can represent a big challenge for large organizations like federal agencies, especially those used to dealing mostly with government furnished equipment (GFE). As full-time employees and contractors introduce their own hardware into a network, administrators in typical perimeter-based environments have to be prepared for personal phones, tablets and laptops that may not have the same standardized protections in place as government-owned hardware or GFE.

In the zero-trust model, that's not a problem, because administrators are able to establish access policies that only grant access to devices that meet specific criteria, these devices are then considered "trusted."

Device assessment is the first step toward creating a coordinated security policy at the device level. Remember that continuous authentication monitors for the signs of the right authenticated users on a device, and device assessments in a zero-trust framework monitor the devices themselves to make sure they meet a predetermined set of rules before they're granted access.

That means that in the context of zero trust, administrators can create an access policy, for example, that requires devices to be running specific OS versions, have certain security features turned on, like encryption and jailbreak/root detection, or have updated versions of Flash or Java.

**The most important part about device assessments is that different policies can be used in concert with one another.**

An administrator can require a phone to have the most recent version of Android or iOS as well as a lock screen that's protected by a PIN code, and deny access to any device that doesn't meet all of those criteria. With device assessment, administrators don't have to worry about unapproved devices gaining access, and users are given a clear understanding of why they were denied access, and what they'll need to do to resolve the issue.

# User Controls

In addition to creating access policies based on typical user behavior and device hygiene, a zero-trust model also relies on appropriate controls at the user level. That's typically achieved via access control on a per-application basis with role-based controls.

Here's what that means. Managing users individually requires a lot of care and feeding, and it increases the risk of human error inadvertently creating security flaws. For example, if a system relies on humans to maintain an accurate running list of which individuals should have access to which applications or files, that security is dependent upon the list always being accurate. In these cases, all it takes is one instance of human error like forgetting to update a user account to create a potential security hole.

With zero trust, administrators establish access roles across applications that govern which job roles should be allowed access – and not which individuals. For example, in federal agencies it's common to see access roles like "Contractor," where assigned users are able to see relevant materials like project-specific documents, but unable to access internal files or applications intended only for full-time employees.

Similarly, role-based access policies can differentiate between different internal users, so anyone with a junior-level role isn't able to access content that's only appropriate for senior-level roles.

# Application Access

At its core, zero trust is about access to applications. All applications. Establishing user identity and device trust set the table for seamless and secure access to applications – on-premises or in the cloud – with no discernible difference.

Secure application access from any device is the destination of the zero-trust journey. Providing granular, role-based access to specific applications based on your customized access policies reduces the attack surface should both the user and their device be compromised.

Zero trust doesn't mean that no trust exists. It means that trust is determined on a narrow path to a successfully authenticated user from a successfully interrogated device to a specific application. It also means that network location – inside vs. outside the perimeter, for example – isn't part of the recipe.

# Conclusion

Zero trust is absolutely achievable in federal agencies; to make it a reality, you'll need to keep three ideas in focus:

## The perimeter is now anywhere you make an access control decision.

Continuous authentication, device assessments, user controls and application access are each good security methods on their own, but for optimal security – and to count as a true zero-trust model – they need to be used in coordination with one another.

Zero trust is designed to let administrators gradually migrate from a perimeter-based framework, so federal agencies can create plans that align with their IT modernization initiatives and allow them to take action without requiring massive changes all at once. This means you can apply these principles today, now, and they can grow with your agency on your mobile and cloud journey; effectively providing a seamless security architecture for the future.

Duo works with many federal agencies to advise and assist with migration to a zero-trust framework using Duo Beyond, which was built on the core tenets and principles of zero trust.

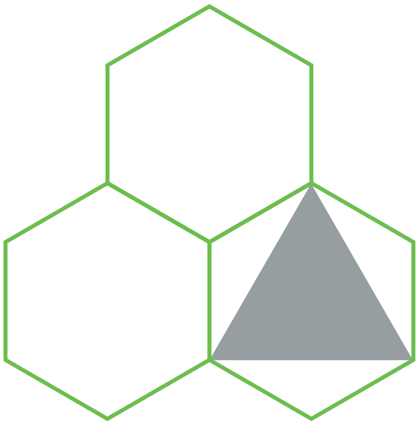

# Duo makes security painless so you can focus on what’s important.

## Trusted Users

Verify user identities with advanced two-factor authentication, and enforce user access policies to limit access to applications.

## Trusted Devices

Duo checks the security health of your users’ devices. Block, warn or notify users of risky devices with our device access policies.

## Every Application

Protect both cloud and on-premise applications, and simplify access with our secure single sign-on.

**The Most Loved Company in Security**

**duo.com**

Email

**sales@duo.com**

North America

**866.760.4247**

Europe, Middle East & Africa

**+44 8003 585 309**