Part One

Moving Beyond the Perimeter

The Theory Behind Google's BeyondCorp Security Model

AUTHOR
WENDY NATHER

DESIGNER
HAFSAH MIJINYAWA

Advancing Enterprise Security

Is the perimeter really dead?

The invisible line that we draw between what belongs to the enterprise and what doesn't – servers, desktops, network, applications, and logins – traditionally depends on firewalls to protect those assets, but the headlines are full of examples where firewalls fell down on the job. People have certainly been promoting the perimeter's demise for years now: the **Jericho Forum** was created to tackle "de-perimeterisation" as early as 2003. The idea really picked up steam as the cloud became more accepted as a common place to store and process data.

Now that Google has come out and described in detail how they have made it happen, and dubbed it "**BeyondCorp,**" it's within practical reach for many more organizations, with a concrete example to consider implementing.

The idea of getting rid of the perimeter is generally too scary for enterprises to contemplate, especially if they've only recently solidified one. So let's not think of it as getting rid of the perimeter, but rather as **tightening security on the inside so that the perimeter isn't the only thing keeping the attacker at bay.**

Enterprise Risks Addressed

The "BeyondCorp" model addresses several important risks for the enterprise:

- An attack that can bypass the firewall, or that starts on the internal network, can spread out to compromise critical systems and steal sensitive data.
- When an application or system is protected with different controls dependent on whether the user is "inside the perimeter" or not, an attacker can compromise the looser set of controls.
- External cloud-based applications and mobile users can face attacks that are outside of the enterprise perimeter protections.
- Users can make the organization vulnerable by using unmanaged and unpatched devices to connect to critical systems and data.

Moving Toward a BeyondCorp Model

To start implementing this new framework, organizations should consider taking the following steps:

- Enroll users and their endpoints into inventories
- Use digital certificates or other techniques to identify endpoints as "trusted" or "managed".
- Classify resources (such as applications) according to risk levels.
- Create access policies based on the authenticated combination of user and endpoint.

Other elements in the framework include single sign-on, device inspection, the trust inference engine, end-to-end encryption, and the reverse proxy that protects applications and enforces the enterprise access policies.

Enterprises often have many of these components already available and can make use of them; Duo also combines many of them in the new edition of its **Trusted Access** platform, **Duo Beyond**.

A New Enterprise Architecture

The kernel explodes in a tiny puff of steam, turning its insides out and expanding far beyond its original size. This describes popcorn, but also describes today's enterprise architecture.

With so many external services available, organizations can be partially or even fully "popped," storing their data outside of the traditional firewalled perimeter. To make things more complicated, a mobile workforce can take its laptops and smartphones to work anywhere, far outside the enterprise's walls and network. And finally, people are using the same software as a service (SaaS) applications for both personal and work purposes. This dynamic environment requires a new security model.

The Google BeyondCorp Vision

Google's vision is similar to **John Kindervag's "zero-trust model"** of information security: to assume that no traffic within an enterprise's network is any more trustworthy by default than traffic coming in from the outside. Of course, enterprises can't operate without any kind of trust; the trick is to set the conditions under which they will decide to trust something.

> **Google's implementation rests on the combination of validated users using validated endpoint devices. This combination is further locked down with end-to-end encryption between these devices and the resources they access.**

Finally, users are allowed only the bare minimum access needed for their roles (which is also known as "least privilege"). As long as the user is authenticated with the right number of factors, and is using an endpoint that has been enrolled and inspected for security vulnerabilities, they can access exactly those resources that they're allowed to by a centralized proxy.

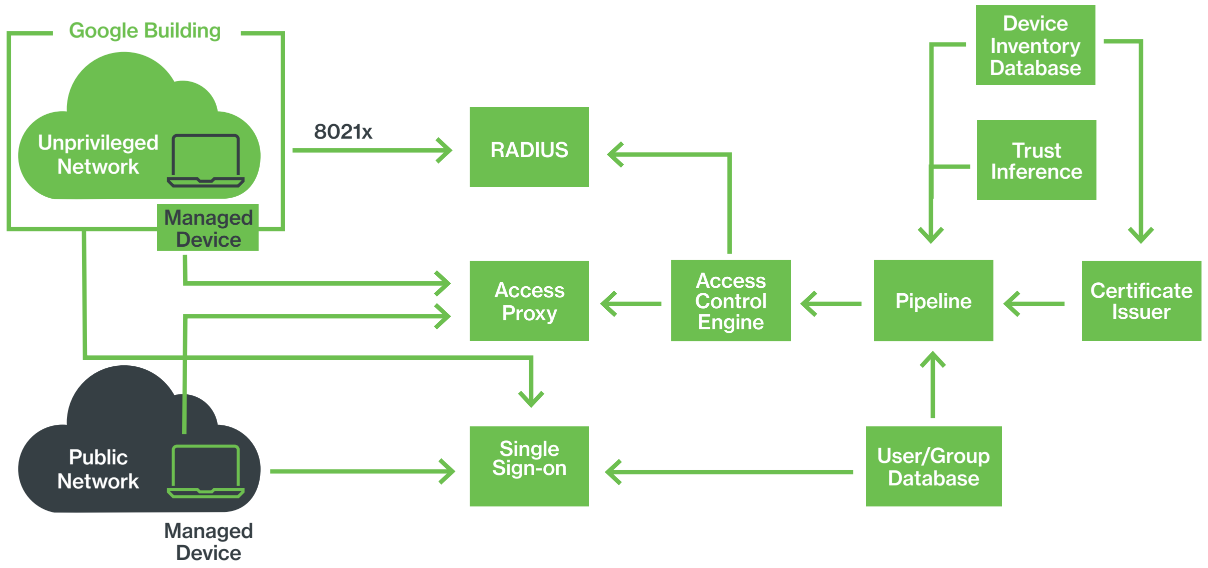

As Google illustrates above, it relies on a device inventory database, a user/group database, and client-side certificates for strong identification and control. To migrate a huge and complex infrastructure to this model, Google had to map and simulate workflows, using transition measures such as split DNS to make sure nothing broke while it was being gradually moved out of the unrestricted internal network (also known as the “soft and chewy center”).

What Risks Does This Approach Address?

The biggest risk, of course, is that an attacker breaks through the perimeter and then has free rein within the trusted internal network. Google specifically referred to the the **"Aurora" attacks** as an example of what prompted BeyondCorp.

THE INSIDER RISK

The other risk with a fully trusted internal network is that you don't have to start by breaking through the perimeter; if you're an insider planning malfeasance, you're already there. The traditional way to deal with this risk is to segment the network. But creating segmentation after the fact can be a major project, disrupting traffic and application tiers, and in many organizations, it never gets done. **And let's face it: a sufficiently successful outsider looks exactly like an insider.** An external attacker will use the same means to get in that work for the legitimate user, so you have to make sure to limit what everyone can do.

POLICY GAPS

Another risk is that the attacker exploits the gaps between different policies or enforcement that apply to the same asset. If the same confidential data is available in two different systems using different types of authentication, the attacker will go after the one that's easier to reach – either because it trusts something else that you can leverage, or because that one authentication method has a flaw in it.

ATTACK SCENARIO: 2FA WORKAROUND

For example, let's say one database requires two-factor authentication (2FA), but the same data is available in another application that doesn't require 2FA – and it has weaker passwords that are shared with a third system. An attacker would try to break into the third system, grab someone's username and password, and use it to get into the non-2FA application. You can prevent this kind of arbitrage by trusting nothing by default and making everyone pass the same tests each time.

VULNERABLE ENDPOINTS

A common risk that every organization faces is the vulnerable endpoint, where out-of-date software contains security flaws that can be exploited by attackers. At the very least, endpoints should be up to date on the operating system (OS) and plugins that they need to use. This isn't always practical due to legacy software that is dependent on older versions of other software, or that is only certified by the vendor for a particular set of infrastructure. But users who simply don't get around to upgrading – especially on their personal devices – are a security headache for the enterprise.

ADDRESSING RISKS WITH APPLICATION POLICIES

With a centralized access proxy, you can have one set of policies for each application, regardless of where the system or user is located. A third-party SaaS could have the same trust requirements for access as an internal web application. This is important because attackers try to come from the "most trusted" location, whether that's a known IP address, an "internal" system, or a favored geographic area.

> **With the BeyondCorp model, it's the combination of validated user and endpoint that earns the trust, not the network.**

Note: you can have different requirements based on whether it's an internal or external app, but once you start making that distinction, you're back on the road to destroying that security model you just tried to implement. Make sure your policies are based on business criticality and confidentiality, not on "inside" versus "outside."

Getting Started with BeyondCorp

If you're already in the hybrid environment – with some of your infrastructure located on-premises and some hosted in the cloud – it's time to think about how you could potentially use the BeyondCorp model to rebalance your security policies to extend to cover assets that aren't within your perimeter. For those with a large network who haven't been able to segment it as much as they'd like, or for tighter control, the BeyondCorp model offers a chance to focus on combining user and endpoint verification with encryption.

The good news is that you don't have to do everything all at once. While Google's description of a comprehensive migration sounds daunting, moving to this different concept of security also works when you do it incrementally.

> **Remember, you're not actually getting rid of the perimeter controls; you're raising the level of security on the inside so that it looks more like the outside. Any progress is a significant improvement.**

Here are some of the high-level steps to plan for:

1. **Enroll your users and their endpoints.** This may require a discovery process, since users might not always be using the corporate assets you assigned them. By routing those users to a popular application through an **authentication gateway** such as the ones that Duo provides, you can get an inventory on the fly, and discover which devices are actually connecting to your corporate systems.
2. **Deploy certificates** to the user endpoints that you want to identify as "managed" or "trusted." The level of trust is up to you, but for some organizations, it means that these endpoints are officially supported and maintained by the enterprise; for others who embrace Bring Your Own Device (BYOD), it means that they've done the initial hygiene check during enrollment and validated that the device belongs to an authorized user.
3. **Classify applications according to risk levels** so that you can enforce different access policies for each. Some resources require global access and contain less sensitive data, such as internal announcement pages, employee directories, and cafeteria menus. Other resources, such as financial systems, HR, customer or patient data, or intellectual property, would have more restricted access.
4. **Create access policies** based on the requirements for each application or system you want to protect. These policies can include how often you want users to re-authenticate; whether they can use personal devices; and which level of hygiene you want to enforce. These policies can be adjusted dynamically based on security events. For example, if a new vulnerability is being actively exploited in a particular endpoint OS or plugin, you can block affected users until they update it. This drives users to **update on their own rather** than waiting for IT to organize a scheduled maintenance window (no more Terror Tuesday, when patches are released!).

> **As a result, you will get better visibility and a tighter set of controls over what your users and endpoints are accessing, regardless of where they are. By adapting to the new reality — that applications, users and devices can change locations at the drop of a hat — you'll be able to maintain a more consistent level of security and user experience.**

Mapping BeyondCorp Components

If you were to build your own BeyondCorp, what components would it entail?

USER/GROUP DATABASE

To keep the information and attributes about your users, and to group them where necessary according to organizational, geographical or other aspects.

DEVICE INVENTORY DATABASE

An up-to-date repository for information on all devices you allow to access the network, including type, purpose, network addresses, asset tags, components, configuration, and responsible user or maintainer.

MANAGED DEVICES

If you do not allow BYOD, this will be the whitelist of corporate-owned devices you will allow to access your resources. If you are using an enterprise asset manager such as LANDESK, Jamf, or Active Directory, you probably have this list already.

CERTIFICATE ISSUER

This is used to mark your managed or otherwise approved devices with a client-side certificate. Depending on which types of certificates you plan to use, the public key infrastructure (PKI) for this may already be part of another security product.

ACCESS CONTROL ENGINE

The repository of all your access policies, such as "only this group of users, together with their up-to-date, assigned and managed devices, may use this sensitive application."

ACCESS PROXY

The part that carries out the connections and policy enforcement. Google's description of its own access proxy can be found **here**. It is much more complex and handles traffic load balancing, Transport Layer Security (TLS), authentication, access control list (ACL) evaluation, authorization, and self-service for users.

TRUST INFERENCE

Deciding what conditions will cause you to place or lose trust in a given device (such as hardware changes). The trust inferrer will rely on a steady input of data from the sources you choose. **Google's description** includes checking to see whether the device is encrypted; whether it has all management agents working; whether its software is up to date; and whether all of the information about that device is current.

SINGLE SIGN-ON (SSO)

Make it much easier on your users by providing **one portal for access** to all of their applications and systems.

Other components you will need

Google's BeyondCorp architecture doesn't explicitly mention **two-factor authentication (2FA)** or **multi-factor authentication (MFA)**, since in Google's case, it's integrated with its own **identity provider (IdP)** service, but it's vital to the strategy of making it harder to compromise an account. In addition to MFA, if you don't have a centralized system for identity management, this will likely make the BeyondCorp implementation more complex.

Summary

Can you ever make the security state of the network irrelevant? That's the end goal of BeyondCorp, and although that theoretically means that enterprises can ditch their traditional firewalls, practically speaking, it's not likely to happen. Any enterprise that is still hosting any connected infrastructure will be responsible for protecting it against many other sorts of network-based attacks (such as denial-of-service), not just authentication-level ones.

BeyondCorp is not a silver bullet that will take care of all risks; it's a way of increasing the security level of what used to be viewed as a "safe" environment. Until you remove the complexity of legacy systems, software and protocols, or move all the hosting off-premises, you'll need your traditional perimeter to continue standing its watch.

By contrast, cloud-first organizations can use the BeyondCorp model to increase the control that they have today over access to third-party SaaS applications. By allowing only known, validated user-device combinations to authenticate, they can filter out a whole realm of daily attempts from attackers trying to take over accounts. Anyone can try to log into a public SaaS application today with a set of stolen credentials - but it'll be much harder to do so if they have to use both that user's fully-patched endpoint and 2FA on yet another device.

> **Ultimately, BeyondCorp is a new way of thinking about security and trust. Applying the "zero-trust" attitude to every enterprise design and process is the real peak maturity on this curve.**

In part 1 of our Moving Beyond the Perimeter series, we discussed the **theory** behind the BeyondCorp security model. In part 2, we'll detail how you can **implement** the BeyondCorp security model in your organization.

Sign up for our **newsletter** to get notified when it's available.

How Duo Beyond Can Help

Building a whole infrastructure to accommodate a new way of thinking takes a long time. At Duo, we've shortened the path by building a platform, **Duo Beyond**, that contains most of the components already: the **device inventory**, **identification of trusted devices**, access control engine, access proxy, **single sign-on**, and **MFA** are all included.

Just bring your own identity provider, your list of users and corporate-owned endpoints, and of course, your strategy for building tiers of trust.

Duo has made the BeyondCorp journey viable for companies such as **KAYAK**, allowing them to tighten their security controls both inside and outside the perimeter, and saving them months or years of effort piecing together their own solutions.

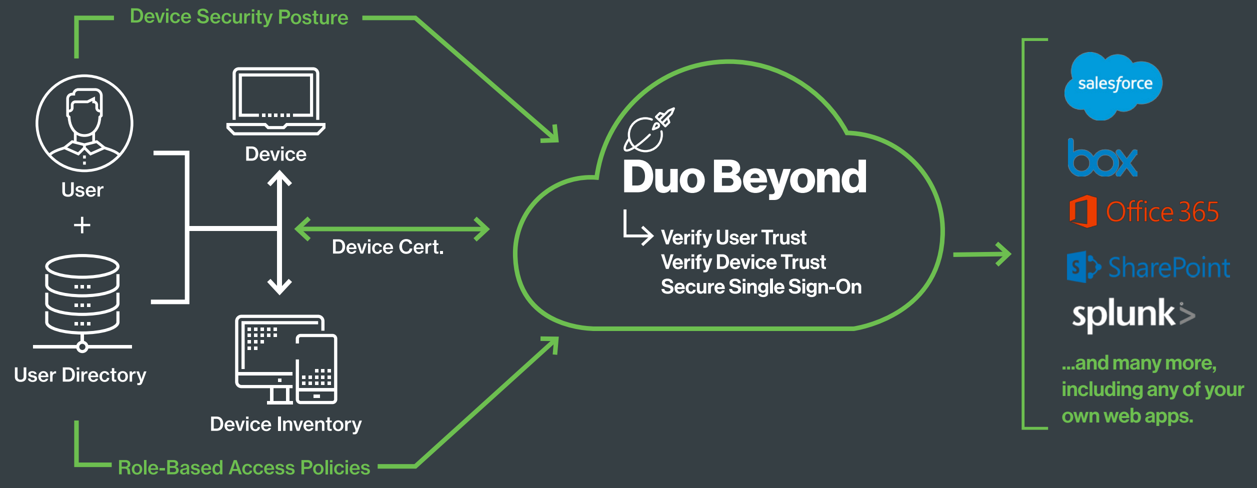

Duo Beyond

With Duo Beyond, organizations can easily implement a BeyondCorp security model within their own organization based on the identity of users and security of their devices. Give your users a consistent user experience while securely accessing cloud or on-premises applications.

Regain trust of your endpoints with Duo Beyond:

- Easy-to-use two-factor authentication
- A secure single sign-on experience
- Complete device visibility
- Identify corporate vs. personal devices
- Easy device certificate deployment
- Block untrusted endpoints
- Secure access to internal apps, without a VPN
- Phishing simulations to assess risk

Try it today at **duo.sc/beyond**

The Trusted Access Company

Follow **@duosec** on Twitter and Instagram.