VPN by Google, explained

At Google, keeping our users safe online means continuously protecting the privacy and security of their personal information. We focus on three core principles: keeping data secure by default, building products that are private by design, and putting our users --you-- in control.

When it comes to networking privacy and security, we've long encouraged the use of Transport Layer Security (TLS) and other protections across the wider web and app ecosystems. Unfortunately, not every online service provider is committed to implementing rigorous data protection standards[1], leaving gaps in how well consumers are protected and in how much control they have over who accesses their network traffic. And even if security protections are properly implemented, sensitive data such as your IP address and the sites you visit can be visible to others[2].

When securely implemented, a VPN provides additional protection by:

- **Providing encrypted transit** that hides your data and network activity from hackers and network nodes along the way, such as public WiFi hotspot or other service providers
- **Masking your IP address** from trackers, sites and apps you visit, which could be used to track your location or your network activity

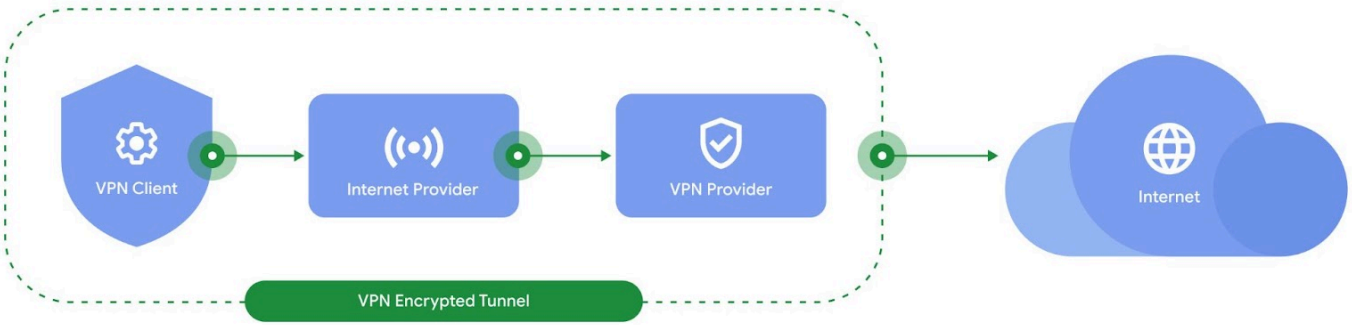

Figure 1: how a VPN connection works

While a VPN removes the ability for intermediaries to snoop on your traffic, it puts the VPN provider in a privileged position to potentially access your sensitive data. Therefore, it is important to choose a VPN provider who provides robust privacy and security guarantees. Unfortunately, not all VPN providers have been

[1] https://www.ssllabs.com/ssl-pulse/
[2] https://www.ietf.org/rfc/rfc8744.pdf

proven to be trustworthy: some services are vulnerable[3], others request unnecessary access or monetize their users network data, and others fail to deliver on the promise of not logging their users' online activity[4].

With growing demand for VPNs[5] in a mixed landscape of solutions, we have used our expertise in privacy, cryptography, and networking infrastructure to build a Google-grade VPN. With VPN by Google One, users' network traffic is not identifiable to the VPN and never logged by VPN. We will **never use the VPN connection to track, log, or sell your online activity**.

Verifiably private

We've built VPN by Google One to address some of the potential vulnerabilities of traditional architectures. A traditional VPN could compromise a user's sensitive data by linking their identity to their network traffic by means of a session ID. This ID could allow VPN operators, or attackers that compromise their infrastructure, to "eavesdrop" and identify users' and their network activity.

We wanted to eliminate that vulnerability by separating the authentication of the user from their use of the service. By employing a cryptographic blind signing step between user authentication and connecting to the VPN, we give users a stronger guarantee that their network activity can't be tied back to their identity.

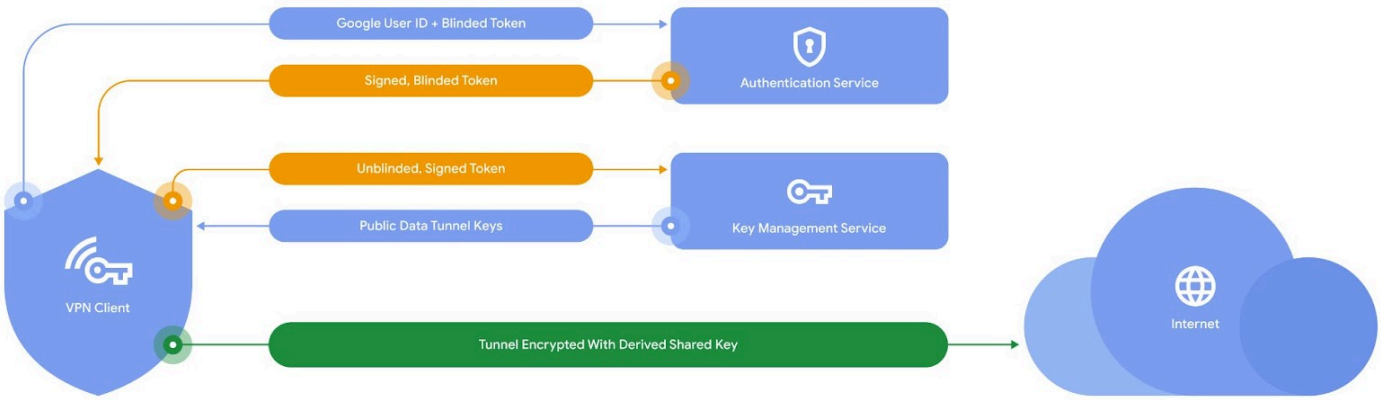

Figure 2: VPN by Google One's authentication with blind signatures

Architecturally, we've split authentication from the data tunnel setup, so that the connection to the VPN tunnel is anonymous and can't be associated with a user's account or identity. In order to achieve this, our authentication method employs a cryptographic technique called "blind signing". Let's go through this method step by step:

- First, the **client** generates an OAuth token and a blinded token (see below for definition) and sends it to the Authentication Service.

[3] https://dl.acm.org/doi/abs/10.1145/3407023.3407029
[4] https://dl.acm.org/doi/pdf/10.1145/3278532.3278570
[5] https://thebestvpn.com/vpn-usage-statistics/#vpnreasons

- The **Authentication Service** validates the user's eligibility to access the VPN. If successful, it exchanges the OAuth token for a signed blinded token

- The client can then 'unblind' this signed token using cryptographic blinding, and request the public data tunnel keys from the **Key Management Service** to connect to the VPN server.

- When the client connects to the **Data Tunnel**, it provides only the signed unblinded token. Thus, the only piece that links the Authentication Service to the Data Tunnel server is a single, public key, used to sign all blinded tokens presented during a limited period of time.

The blinding algorithm employed was first described by Chaum in 1982[6], and is commonly referred to as 'RSA Blind Signing'. The goal is to never use the same identifier in the Authentication server and the Key Management Service, which in this case, provides a separation between the user identity used to authenticate and the use of the VPN tunnel.

The servers are physically distinct and only share a cryptographic root-of-trust to validate the signed unblinded token. Due to this careful authentication architecture, it would be infeasible for an attacker to break the cryptographic protections of one of the services with enough time to break the second and thus be able to associate a user to their network activity. We've calculated that it would take years to break both services, even when using the equivalent of roughly Google's entire global computational capacity.

But, if the VPN server doesn't know anything about the client, how does it know what GeoIP to assign or other user customizable configurations of the data tunnel?

Encoded in the blind token, the authentication service and the VPN server exchange what we call *public metadata,* which is the only shared information between these two services. Public metadata does not include any user identifiers or other fields that could be used to tie a token back to a user. Most importantly, to provide assurance that this is true, public metadata is fully inspectable by the client, which checks that only the expected fields with their pre-established possible values are included. For example, public metadata contains an indicator of the user's coarsened location, which represents either the user's current broad region (such as their country) or a more local region with an online population of at least 1 million.

VPN logging practices

The authentication step has already separated the user's identity from the data tunnel that handles your network traffic. On top of that protection, the following data is **never logged**:

- Network traffic, including DNS

- IP addresses of the devices connecting to the VPN

- Bandwidth utilized by an individual user

- Connection timestamps by user

[6] https://www.hit.bme.hu/~buttyan/courses/BMEVIHIM219/2009/Chaum.BlindSigForPayment.1982.PDF

The VPN authentication and data plane services only record aggregate metrics —without any user identifiable information— for service reliability and performance optimization. These include aggregate throughput, uptime, latency, CPU/memory load and failure rates. They also include the coarsened user locations, aggregated and anonymized to prevent identifiability. Client applications running on the user's device may log additional metrics to understand product and feature adoption and engagement, prevent fraud, and to ensure VPN connection health. Client applications also provide the option to send feedback and errors to us, which include application and system logs, and are used for debugging purposes.

We believe an easy to use, highly private and performant VPN will significantly help improve user privacy online. So it should come as no surprise that we want to make VPN technology available to as many users as possible.