VPN by Google One, explained

At Google, keeping our users safe online means continuously protecting the privacy and security of their personal information. As we design our products, we focus on three core principles: keeping our users' information secure, treating it responsibly, and putting our users in control. Over the years, we've put these principles into practice introducing tools that allow users to manage their passwords, view and delete their activity history, set auto deletion timelines, turn off web & app activity, and turn on private browsing sessions.

We've also long encouraged the use of Transport Layer Security (TLS), the widely adopted cryptographic protocol for securing communication over networks, and other protections across the wider web ecosystem. 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 can access their network traffic data. Even if security protections are properly utilized by online service providers, information such as your IP address and the sites you visit are not always encrypted and can be accessed by others.

A VPN provides both encrypted transit and IP address dissociation for packets between users' devices and the VPN servers. When securely implemented, this hides online activity from network nodes along the way that might have visibility into user traffic data, like public WiFi hotspots or other service providers. While this removes the ability for intermediaries to snoop on user traffic, this puts the VPN provider in a position to see all of a user's unencrypted browsing traffic, e.g., the domain of every website visited. Because the VPN provider occupies this privileged position, the user must be able to trust that the VPN provider has strong privacy and security guarantees.

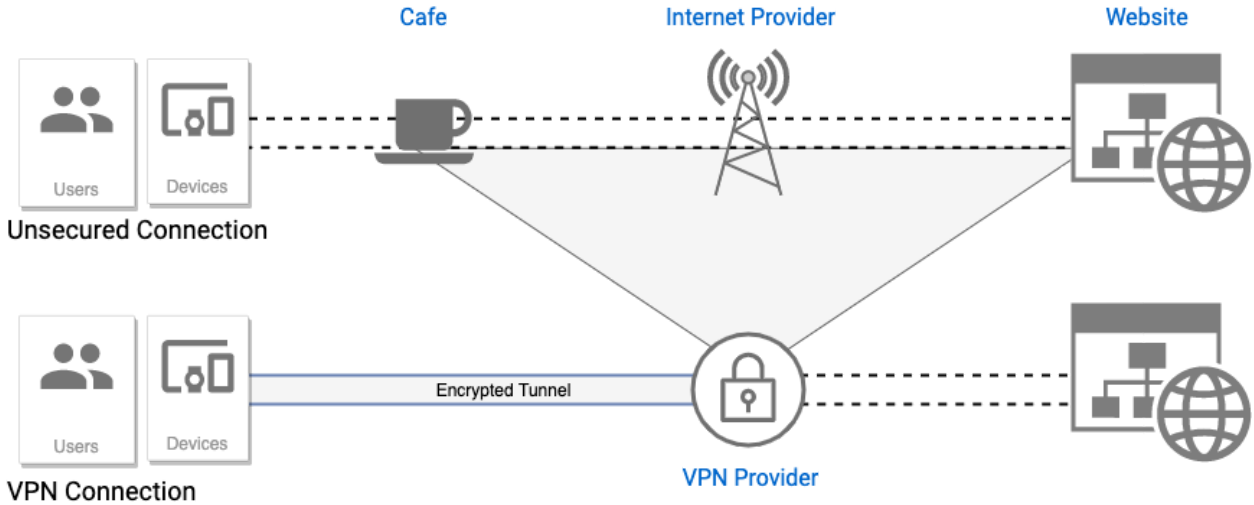

Figure 1: Unsecured connections v. VPN Connections

[1] https://www.ssllabs.com/ssl-pulse/

Demand for VPNs is growing, with evidence that it's becoming more mainstream -- up to 25% of all Internet users accessed a VPN within the last month of 2019.[2] Unfortunately, not all VPN providers have been proven to be trustworthy: some services are vulnerable[3], others request unnecessary access to their users' data or monetize the same data that users are utilizing the VPN to keep private and secure, while others fail to deliver on the promise of not logging their users' online activity.[4]

With growing demand for better privacy in a mixed landscape of solutions, we have used our expertise in privacy, cryptography, and infrastructure to build a Google-grade VPN that provides additional security and privacy to online connectivity without undue performance sacrifices. With VPN by Google One, users' online activity is not identifiable to the VPN and not logged by the VPN. We believe a VPN must be transparent, and robust. That's why we have open sourced our client and will provide a third party audit of the end-to-end solution to make them externally verifiable.

Under the hood

Privacy is at the core of the products and services we build. With VPN by Google One, we will never use the VPN connection to track, log, or sell your online activity. Some minimum logging is performed to ensure quality of service, but your network traffic or IP associated with the VPN is never logged. To demonstrate how our design works, we have open sourced the code that runs on a user's device and in the coming months we will be open sourcing the server side user authentication mechanism as well as providing the results of a third party audit, currently underway. These will provide further assurances of how user data is handled and how robust the VPN's security is.

Open sourcing our VPN and providing an audit are just some of the steps we are taking to ensure user privacy. While building VPN by Google One we realized it was important to strengthen some of the systems that are often attacked or compromised in order to access users' personal data. Traditional VPNs can sometimes compromise a users' identity or online activity by linking the usage of their service to the activity they conduct 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 activity.

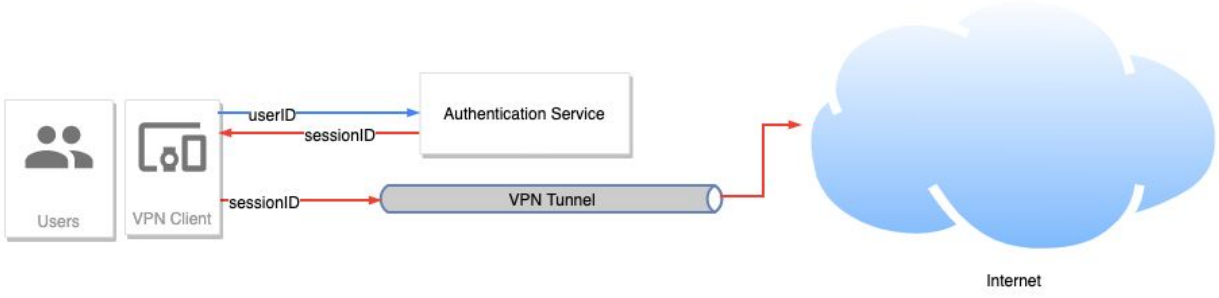

Figure 2: Traditional Authentication Architecture

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

We wanted to eliminate that vulnerability by separating the authentication of the subscriber from their use of the service. By employing a cryptographic blinding step between user subscription validation and connecting to the VPN, we give users a stronger guarantee that their online activity won't be tied back to their identity.

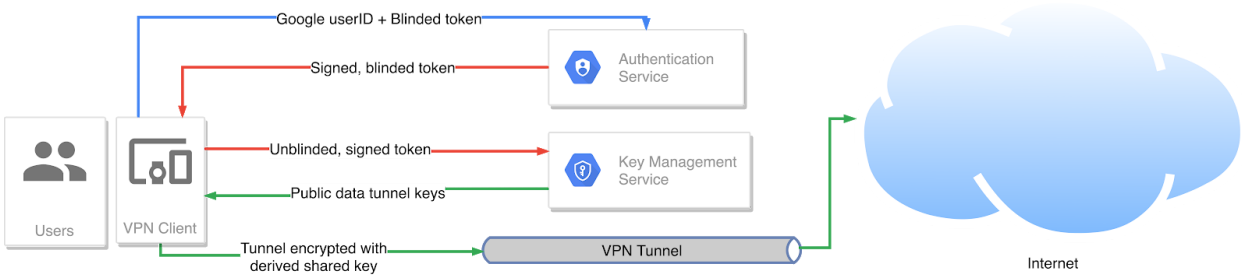

Figure 3: VPN by Google One's authentication with blind signatures

Architecturally, we've split authentication from the data tunnel setup into two separate services:

- **Authentication service:** This service validates users' access to VPN by Google One. The client first generates an OAuth token and a blinded token (see below for definition). Then, the authentication service validates and exchanges the OAuth token for a signed blinded token.
- **Key Management Service:** The client can then 'unblind' this signed blinded token using cryptographic blinding. When the client connects to the data tunnel server, it provides only this signed unblinded token to the data tunnel server. Thus, the only piece that links the authentication server 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[5], 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. To accomplish this, the client generates a token, hashes it using a Full Domain Hash, and combines it with a random value and the server's public signing key to produce a blinded token. That blinded token is then signed by our authentication server. When the client wants to connect to the VPN, it can unblind the blinded token and its signature using the random value only it knows. The unblinded token and the signature are then verifiable by our Key Management Server.

The servers are physically distinct and only share a cryptographic root-of-trust to validate the signed unblinded token; they strictly share no other information - specifically, no user identifiable information is available to the data tunnel servers. As a result of 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 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.

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

Users only benefit from a VPN when it is used. If the VPN is slow or unreliable, users may turn it off, putting their traffic at risk. Speed and reliability are at the core of our VPN. At launch time, the VPN protocol will be a Google proprietary protocol; however, to ensure a high throughput while minimizing battery consumption, we will soon adopt IPsec as the data tunnel protocol due to its native support in Android. We may eventually use other protocols, such as Wireguard, as their native support improves or on platforms where no specific protocols have a specific advantage.

Before the client initiates the data tunnel setup, it does a DNS query to resolve a pool of exit nodes topologically nearby in one of Google's Points of Presence (PoP) locations. The client asks to be associated with this pool during the data tunnel setup, ensuring a low-latency connection on a node with reasonable load.

The VPN service does not limit the users throughput speed and will allow as much as the infrastructure is capable of delivering at any given time, often resulting in available user speeds above 300 Mbps.

While the user ID and IP data have already been protected by the authentication steps, the data tunnel path does not log any user data traffic; only recording aggregate metrics, without user identifiable information, where needed for reliability. This delivers a service with higher reliability while maintaining the VPN's privacy guarantees. We combine these approaches with strict standards on logging and auditing of our open-source code to ensure the VPN remains trustworthy.

Using a VPN shouldn't require that you completely turn over your trust to the VPN provider. A VPN provider should be able to transparently demonstrate how their service keeps your data private. Our VPN client-side code is open sourced so that users and privacy experts alike can verify how user data is handled. We will have external security experts audit VPN by Google One end to end, including the server-side implementation, and publish a report on our VPN privacy protections.

What does the VPN log?

In order to provide peace of mind for our users that their activity is private from the VPN operator and from potential attackers, VPN by Google One does not log user activity on the network or other information that could reveal personally identifiable information about them.

The following data is NOT logged by the VPN for a given user:

- Network traffic, including DNS
- IP addresses of the devices connecting to the VPN
- Bandwidth utilized by an individual user
- Connection timestamps by user

Because it is necessary to ensure a healthy and performant VPN, the following kinds of data are logged at an aggregate level (individual users' data cannot be identified) for performance monitoring and debugging purposes:

- Aggregate throughput
- Aggregate VPN tunnel uptime
- Aggregate VPN tunnel setup latency
- Aggregate Total bandwidth rate
- Aggregate Packet loss rate
- Aggregate VPN tunnel failure rates
- Aggregate VPN tunnel retries
- Aggregate Service/Server CPU and memory load
- Aggregate VPN tunnel setup error rates

In order to measure overall service experience, debug the service, and prevent fraud without compromising user privacy, the following data may be collected for a user:

- Use of the service in the last 28 days. This metric collects how often the service was used in the last 28 days but does not track the specific times they used the service nor the duration of the usage nor the amount of data used.
- Number of recent attempts by a user to set up a VPN session to ensure that the user does not exceed the maximum number of allowed concurrent sessions. User IDs are encrypted and therefore cannot be personally identified by the concurrent session check.
- Server error logs without request or response data.

Finally, for those users that wish to share feedback or errors with Google, the client application provides users the option to send Google application and system logs from their device which contain personally identifiable information (such as email) and is used for debugging purposes. This is an optional capability and requires user permission every time they wish to submit such information.

Looking forward

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

For starters, we will provide the service on Android. Over time, we plan to scale it across more platforms like iOS, ChromeOS, Windows, and Mac.

Because a VPN sits at the interface between the device and the network, there are many interesting opportunities for it to provide additional security to users by blocking threats before they get to their devices. We will continue to explore these for future additions.