WHITEPAPER

CLOUDFLARE

How the Cloudflare network maintains data privacy

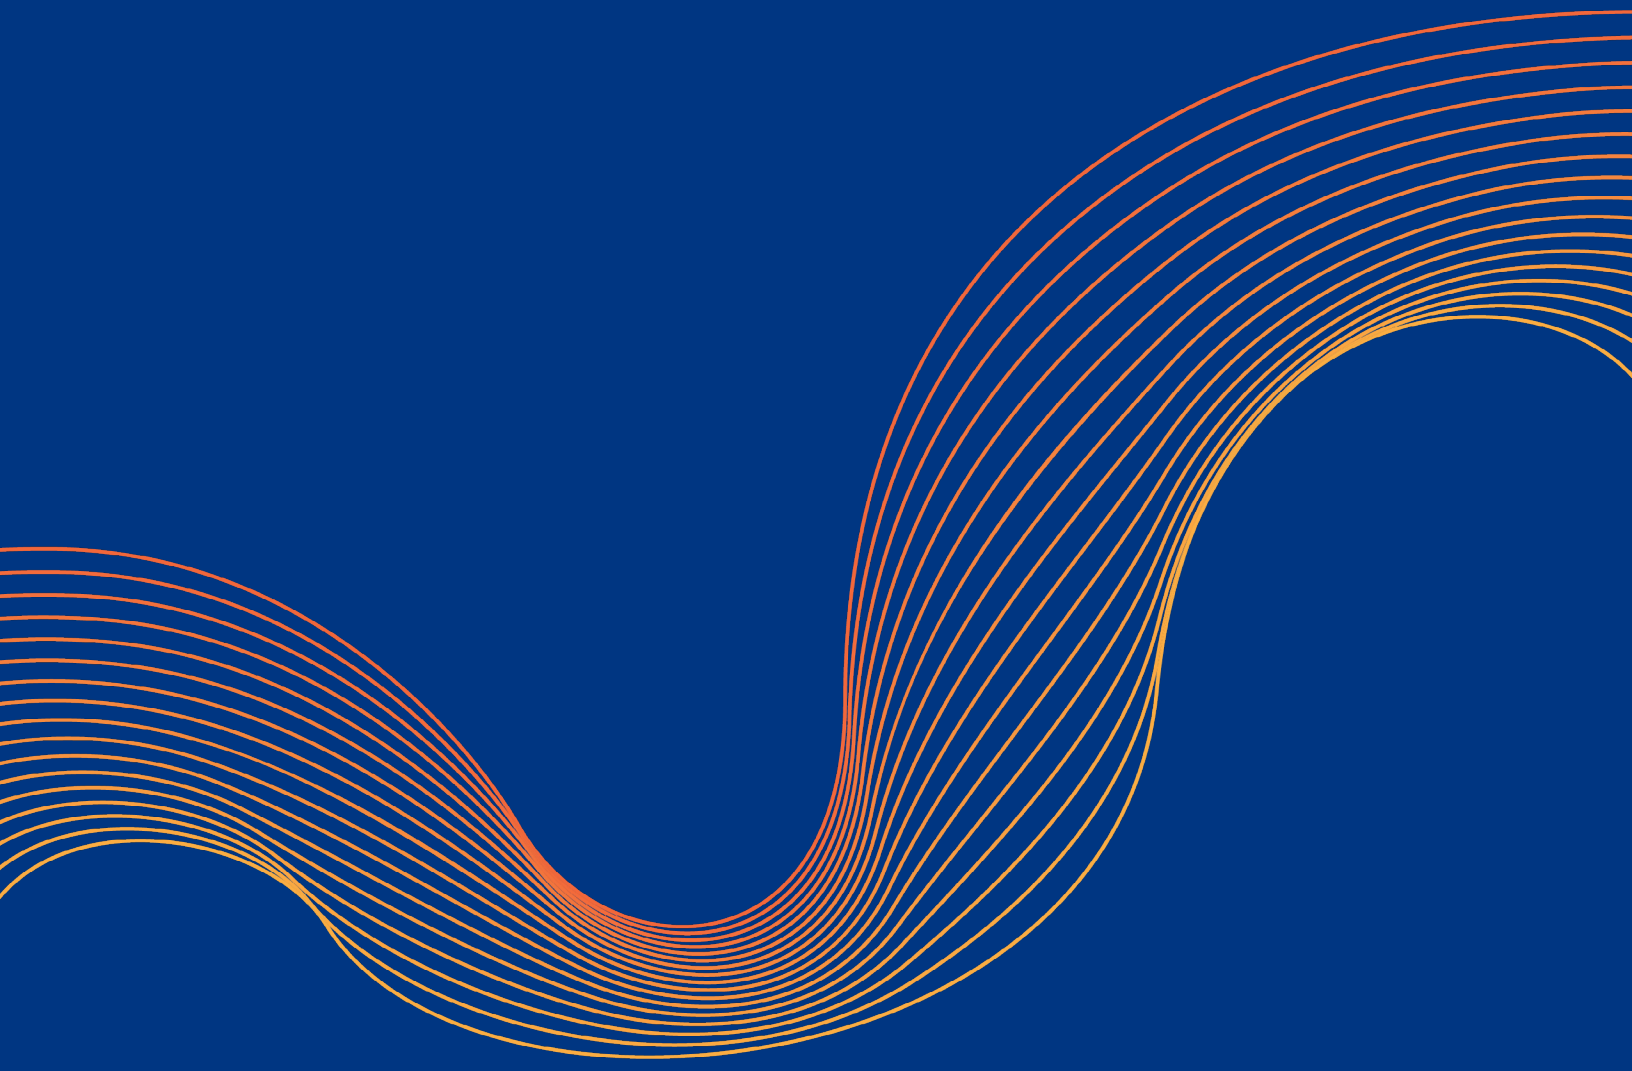

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

TABLE OF CONTENTS

INTRODUCTION

Cloudflare's network and business are all ultimately built on customer trust. We seek to continually earn and maintain that trust not only with privacy-first policies and procedures that guide how we manage customer and end-user data on our systems, but also by building privacy into our products and services. For this reason, we are constantly improving the security of our systems, we encrypt data at rest and in transit, and we allow our customers to determine how traffic is inspected across different locations around the world.

This paper breaks down how Cloudflare uses security measures to protect data as it crosses our network and as we analyze metadata from our network.

Part 1 explains how data traverses our global edge network of data centers, and how we incorporate encryption into that network to guarantee privacy.

Part 2 explains how we use encryption to protect the metadata we collect from that edge network.

Part 3 discusses our protection of encryption keys so that our encryption cannot be broken.

PART 1: HOW DATA TRAVELS ON THE CLOUDFLARE NETWORK

The Cloudflare network includes data centers in over 200 cities and over 100 countries. For resilience and for faster performance, it has been constructed as an Anycast network, which means that every location announces all Cloudflare IP addresses. Because of this construction, users who make a request for a website or application on the Cloudflare network are always directed to the closest data center to them.

Suppose Katerina in Cologne, Germany loads a website that is on the Cloudflare network. Her request for the website goes to the Cloudflare data center located in Düsseldorf, Germany, which is only 45 kilometers away (approximately). Now suppose Katerina drives to Frankfurt, Germany the next day and loads the same website: because the Cloudflare data center in Frankfurt announces the same IP addresses as the one in Düsseldorf, Katerina's request is now directed to the Frankfurt data center.

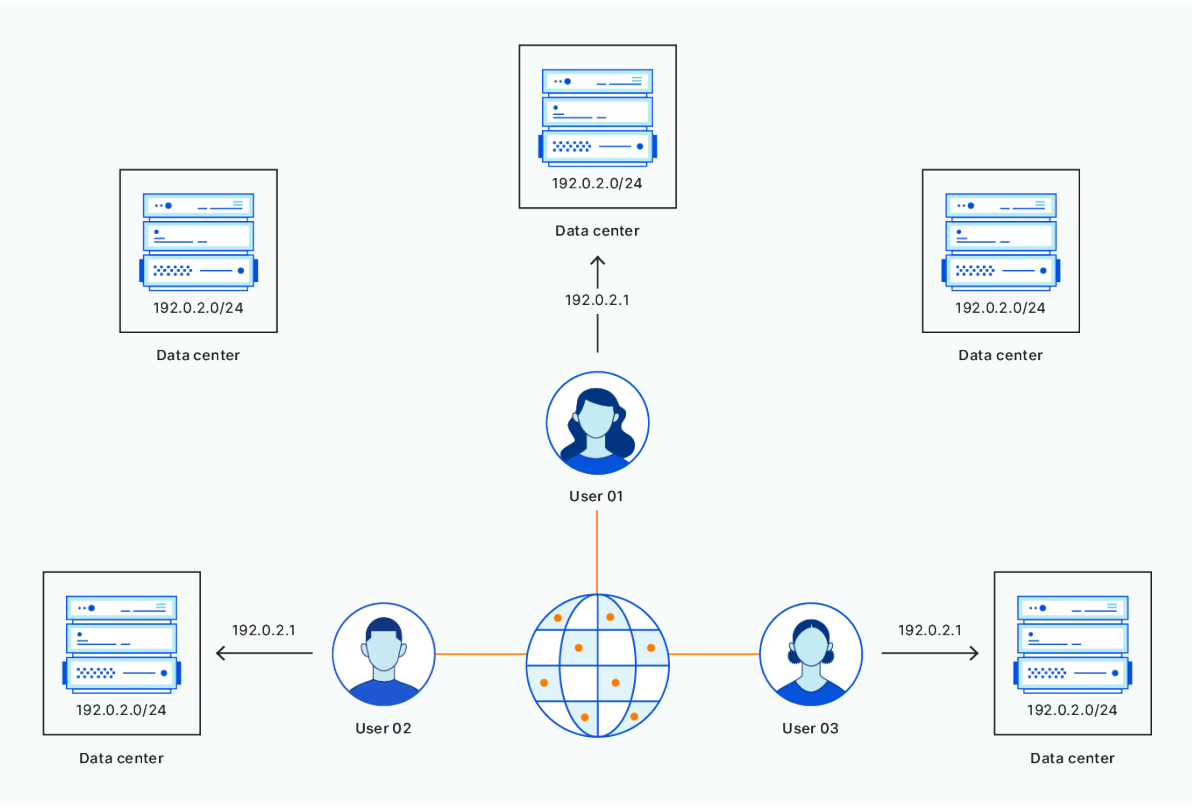

PART 1: HOW DATA TRAVELS ON THE CLOUDFLARE NETWORK

Once the device of a user like Katerina connects with the Cloudflare data center (no matter which one), the request is processed in the following way:

- Certain types of requests for data that can be used for cyber attacks are immediately dropped based on the addressing information.
- Next, the request is inspected using the business logic requested by the customer: page rules, firewall rules, rate limiting, and so on. This enables the detection and prevention of a variety of different types of cyber attacks and malicious traffic, including layer 7 DDoS attacks, automated bot traffic, credential stuffing, and SQL injection.
- The request is then passed to the cache. If the cache can fulfill the request with a cached copy of the content, it does so; if not, it forwards the request to the customer's origin server over the Internet. Traffic between the Cloudflare server and the origin server is encrypted, if enabled by the customer.
- When the response comes from the customer's origin server, any static content is cached onto encrypted disks. The response then goes back through the chain to the user across the Internet.

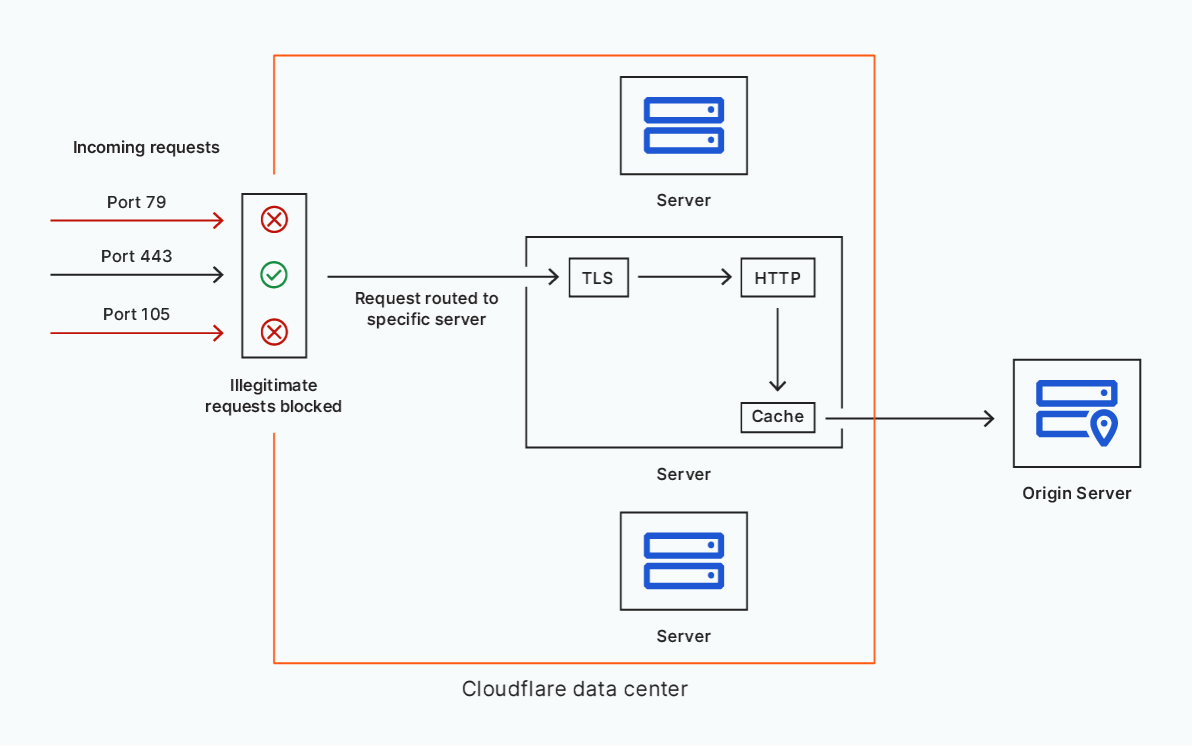

PART 1: HOW DATA TRAVELS ON THE CLOUDFLARE NETWORK

Privacy of HTTP requests

Cloudflare enforces the usage of TLS encryption for all data in transit between an end user and any Cloudflare data center as directed by our customers. No intermediary third party, whether it is the in-between networks that provide transit or a malicious attacker, can view the encrypted data, keeping it private and secure.

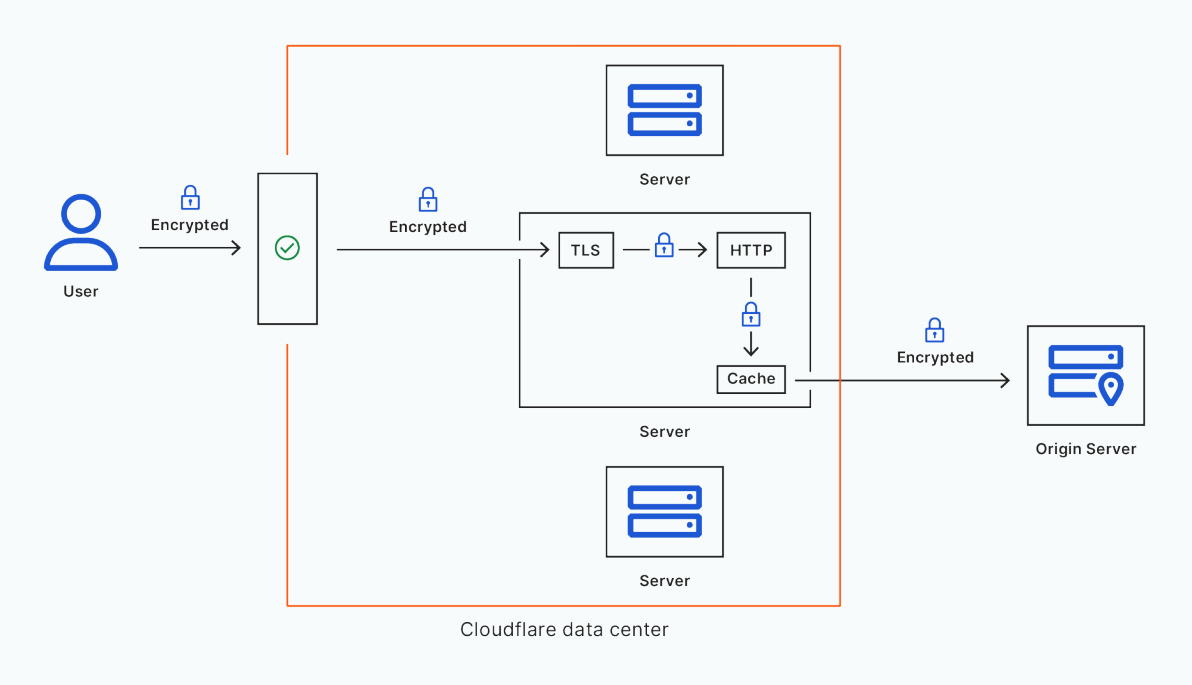

Traffic within the data center, between data centers, and between the Cloudflare network and a customer's origin is also encrypted. In addition, all request and response processing within a Cloudflare data center takes place in memory; nothing is written to disk except for cached static content, and all cache disks are encrypted.

How data travels using Cloudflare Regional Services

Many Cloudflare customers have expressed an interest in having granular control over where and how their data is handled.

Cloudflare Regional Services enables Cloudflare customers to specify where they want data to be inspected. For customers with Regional Services activated, the Cloudflare network processes requests in the same way described above, with one crucial difference: requests are only inspected for cyber security risks and cached when they reach a data center within the specified region.

An end user's request still travels to the nearest data center ("nearest" measured by number of network hops). If the data center is outside of the specified region, then the request is forwarded to a data center within the designated region before business logic is applied. Regional Services customers are assigned dedicated IP addresses to enable this region-specific processing, while still getting the benefit of global DDoS protection.

PART 2: DATA COLLECTION AND PRIVACY

What data does Cloudflare collect from the edge?

Cloudflare's business has never been built around tracking users or selling advertising, and we minimize our collection of personal data. The metadata that Cloudflare collects from our global edge network of data centers in order to keep the network running smoothly and to optimize our service is largely technical in nature; the metadata contains extremely limited personal data, most often in the form of IP addresses. Data that is collected is sent to our core data center in the US for processing.

The metadata that Cloudflare collects falls into three main categories:

1. System metrics and debugging data
2. Data collection for Cloudflare's data analysis products
3. HTTP request metadata for operating Cloudflare's network

More details about each below.

1. System metrics and debugging data

Cloudflare collects and processes various metrics about server and network performance, such as the number of requests received per minute by a given data center or the round trip time (RTT) between data centers. These metrics are collected as aggregated statistics, stripped of their context. These aggregated statistics, which do not contain personal data, are sent to our core data center for processing, and all data sent from our edge to our core is encrypted with TLS in transit.

Cloudflare also collects debugging information for normal maintenance from all software that runs on the edge network.

All this data collection is necessary only to keep the Cloudflare network and Cloudflare services up and running. In addition to being encrypted in transit with TLS, all of this data is encrypted at rest in our core data center with hard disk encryption.

2. Cloudflare data analysis products

Cloudflare offers customers data products to understand Cloudflare's services and how they can make better use of Cloudflare. Our data products also help customers better protect their own origin servers, configure our services, and understand the behavior of their own business and systems.

All of our data products are derived from metadata about software running at our edge data centers. For HTTP requests, for example, Cloudflare incorporates URLs, information about what Cloudflare features were used, timing information, cache information, and select HTTP request and response headers. Similarly, Cloudflare collects metadata about DNS requests, TCP flows, Access Logins, Stream video views, and all of our products. We also provide a privacy-first web analytics service that provides customers insight into how users browse their websites.

If enabled by the customer, Cloudflare Logs provide detailed information about every event to customers. This data, encrypted in transit, is sent from the edge network back to the core and then pushed to customers. Customers may choose to store encrypted log data in the core for up to 7 days.

Both Analytics and Logs encrypt data in transit to and from the edge network and to and from our customers.

PART 2: DATA COLLECTION AND PRIVACY

3. Cloudflare's use of random-sampled HTTP traffic data

Cloudflare stores a simple random sample of all HTTP traffic served across our network. This data is pseudonymized and is not used to track individual users across Internet properties. Cloudflare has no means of connecting any one IP address to any one natural person. All collected sample data is transferred over encrypted connections, stored using secure hard disk encryption, and retained for a limited time (up to 12 months). We use sample data to diagnose malicious activity, help investigate customer incidents, and improve the overall effectiveness of our security products. Every Cloudflare customer benefits from the cumulative intelligence of the sampled data across our entire network.

PART 3: OUR PROTECTION OF ENCRYPTION KEYS

Encryption forms the backbone of the technical measures we use to help protect the privacy of our customers' content and metadata as it flows across our network. Encryption on our network cannot be broken by third parties. This has crucial implications for privacy: if the encryption keys remain secure, then the content protected by that encryption remains private.

As a security company, Cloudflare views the protection of encryption keys as paramount. To ensure the security of encryption keys, Cloudflare maintains strict physical security standards and access controls. We also have had a longstanding public commitment that we would fight any governmental attempt to access our encryption or authentication keys or our customers' encryption or authentication keys, as described in the Cloudflare Transparency Report.

Given how critical it is to keep encryption keys secret and protected, we believe not only in building robust access control systems, but in making sure our customers have options for granular control of their own encryption keys. Customers who want to limit the geographic regions where keys are stored with data center granularity can use Geo Key Manager. Customers who want to retain on-premise custody of their private keys while still using our SSL services can use Keyless SSL.

Conclusion

At Cloudflare, our mission is to help build a better Internet, and we believe data privacy is core to that mission. We will keep working, just as we have for over a decade, to find new ways to ensure privacy and security for our customers and for the Internet as a whole.