

SafeDNS Shield

- [SafeDNS Shield](#)

SafeDNS Shield

1. Product Overview

Our on-premise solution for DNS traffic content filtering is a DNS Proxy which processes DNS queries to identify the user, it compares the target domain with the filtering policy of the client, and decides whether to block or allow the traffic.

As for the technical part, the blocking is implemented by substituting the target resource's IP address with the IP of the block page. This can either be a custom corporate page hosted outside of our solution, or a default block page embedded within our solution itself (which can also be customized).

An important limitation is the need to add our root certificate to the trusted list on every end-user device to display the block page over the https protocol. For block page display via http protocol this is not required. Without the certificate installed on the end device, when a domain is blocked over https, the user will not see the block page, but access to the resource will still be denied.

If blocking is not required, the DNS Proxy simply forwards the target domain resolution request to the next caching DNS server in the chain. This can be either a local corporate DNS server, an ISP's DNS, or any public DNS service.

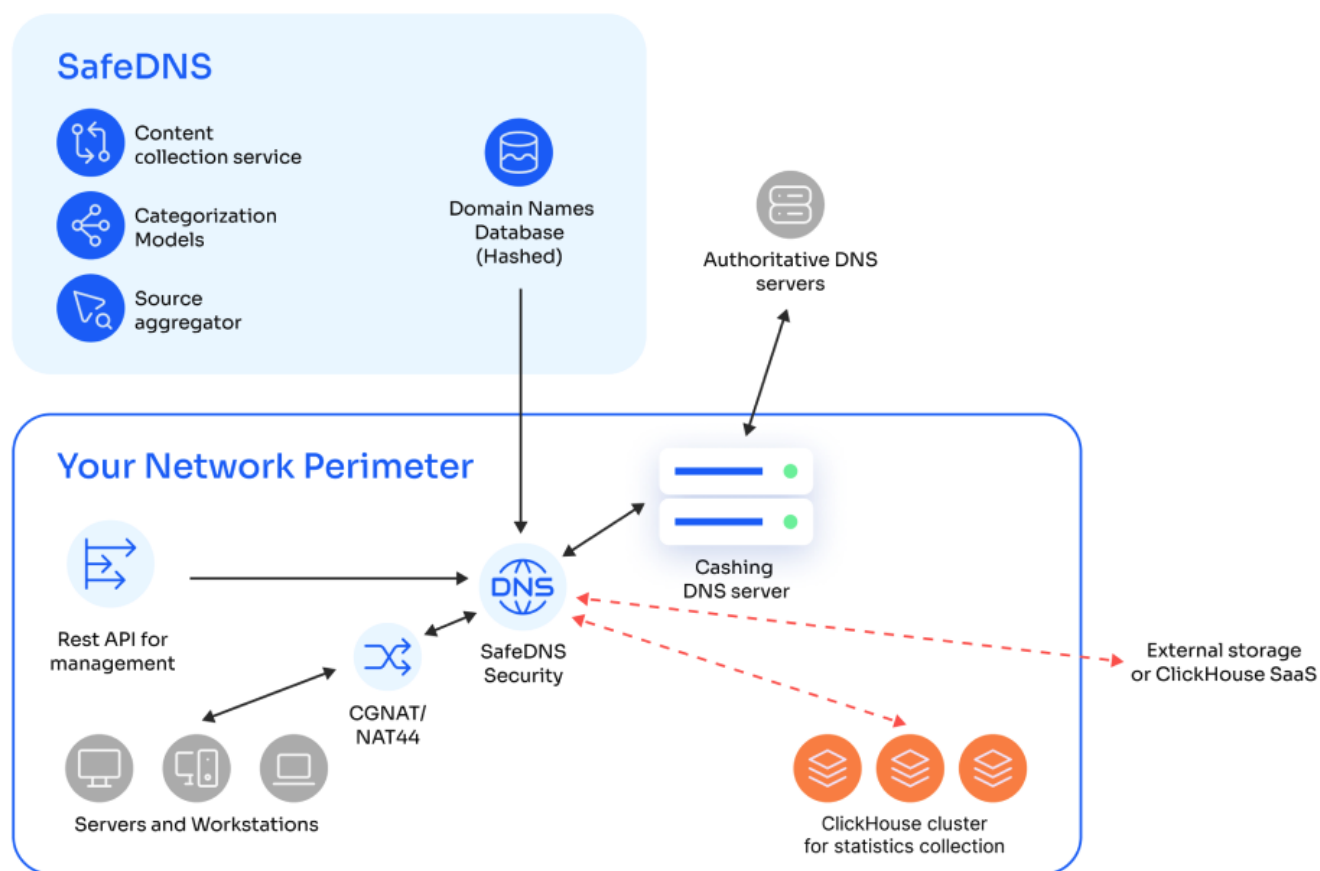
Furthermore, by processing all DNS traffic, this solution enables comprehensive traffic analysis on a per-user basis. Logs of all requests are compiled, and access to statistical information is provided.

2. Product Deployment options within the Company's Network

Depending on the client's network topology, there can be numerous deployment schemes. Below are examples of typical solutions.

For ISPs

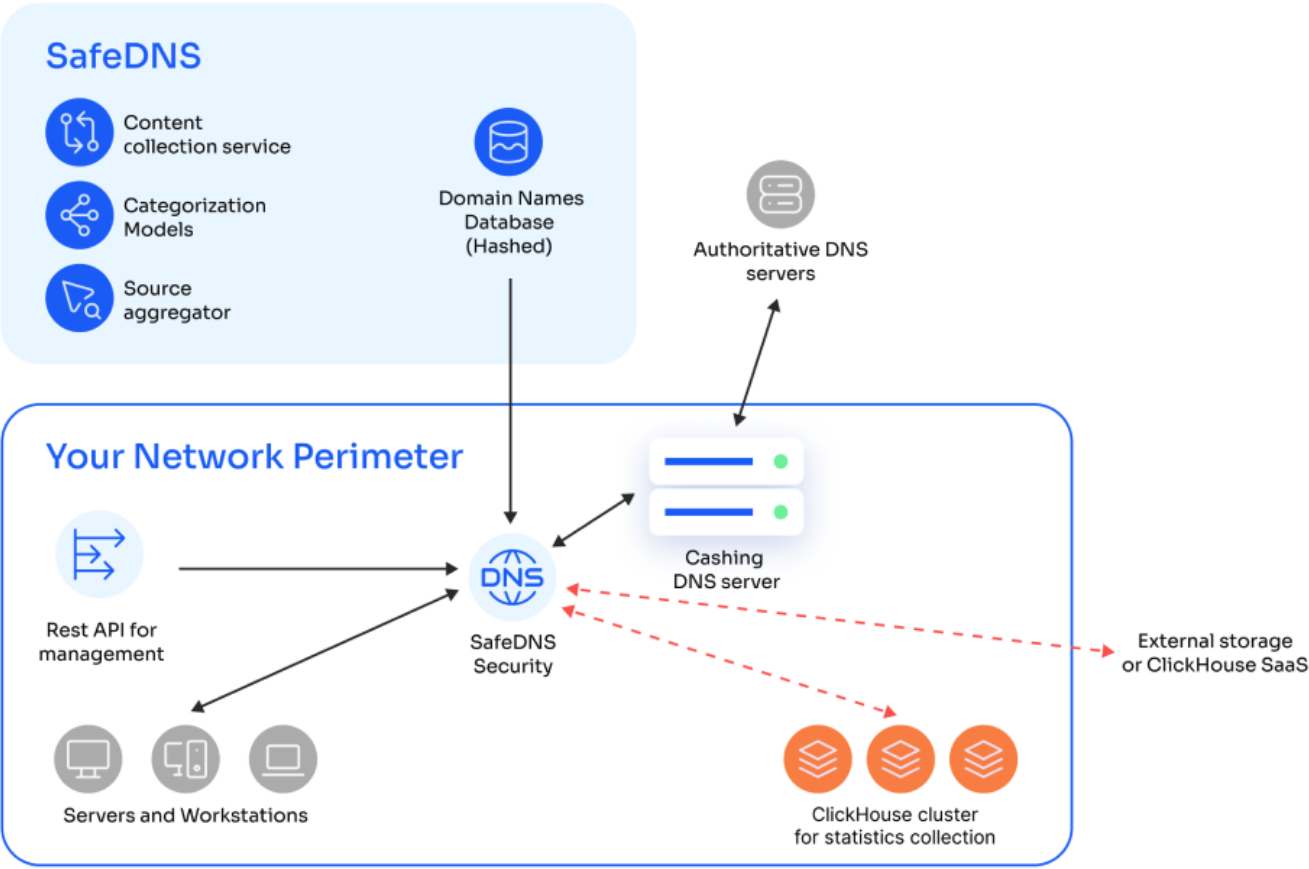
SAFEDNS



This deployment option is used in large networks where end users are behind NAT in relation to the client's main equipment, making it impossible to identify them by individual IP addresses.


For corporate clients

SAFEDNS



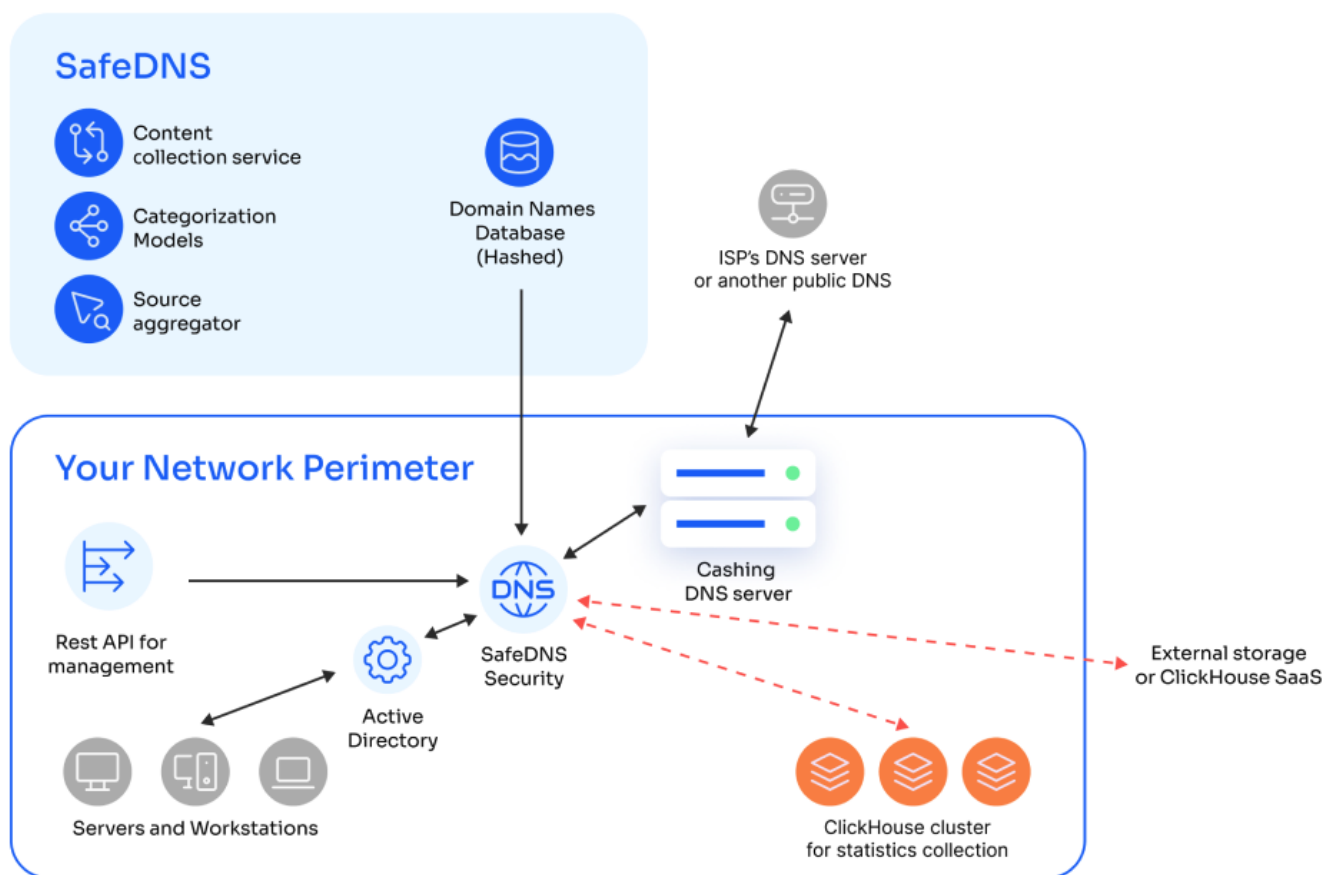
or

SAFEDNS



This setup is used in networks where the end user can be identified by their IP address at the deployment location of our solution. Depending on whether the company has its own caching DNS server, requests are either proxied to it or to external infrastructure (ISP's DNS or any public DNS, such as 1.1.1.1 or 8.8.8.8).

For a corporate client with AD (or another local DNS)



In this case, requests are first handled by the AD server, which manages local addresses. Then, the domain is passed to our service as an external source, where certain domains are filtered. And only if access to the target resource is necessary will the request be forwarded to an external recursive DNS. In this setup, it is impossible to identify the end user, as all requests originate from the domain controller or another local DNS server.

However, SafeDNS is currently developing a solution to address this issue, which may allow maintaining the connection between end users and the domain controller while also enabling their identification on our equipment. Additionally, it will become possible to manage filtering through AD group policies. This solution is expected to be available by the end of 2024.

3. User Identification

To apply various filtering policies and to separate statistics by requests, it is necessary to identify the end user. In the on-premises solution, user identification is achieved through one of the following methods:

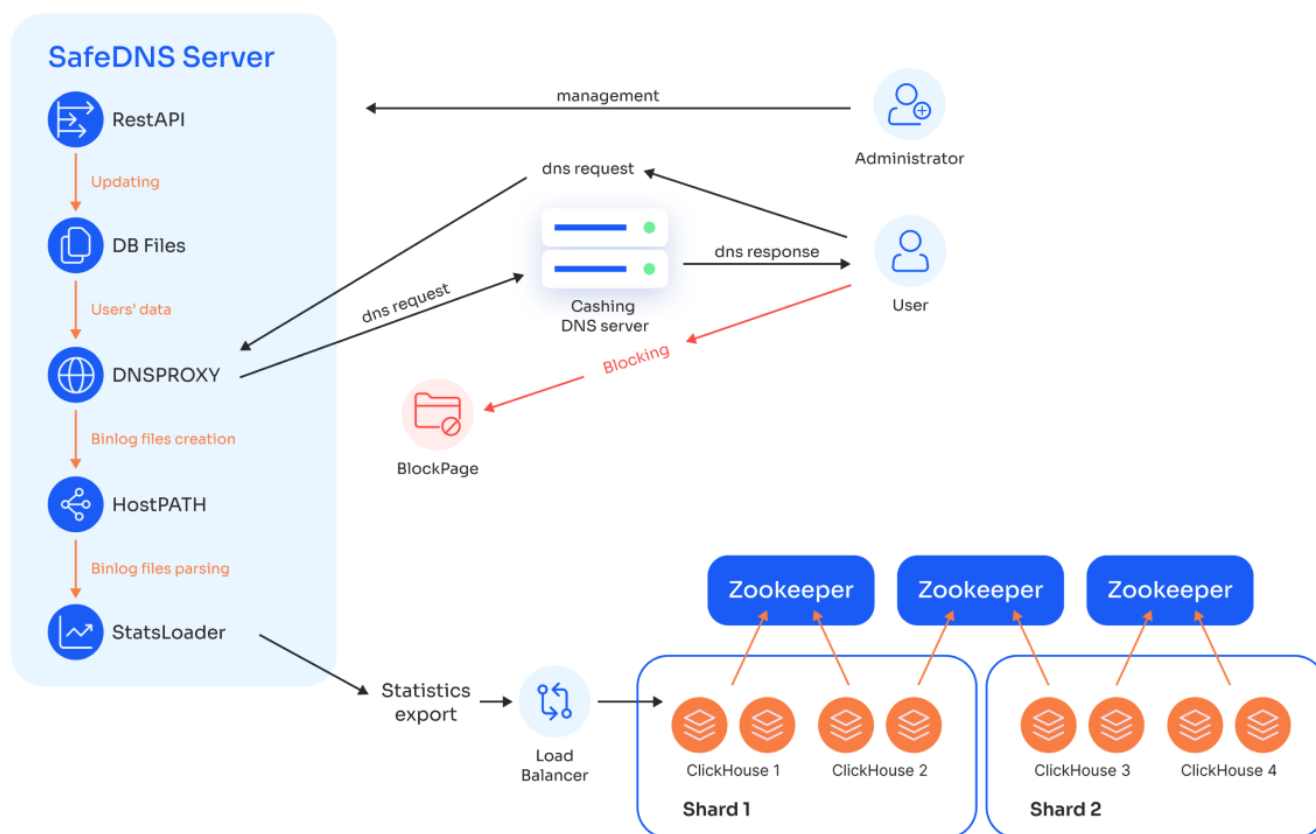
- IP
- IP/subnet
- IP:port
- IP:[port]-[port]

If each user has a unique IP address (from which DNS requests are made), we identify the client by this IP. Alternatively, if detailed separation is not required, we can identify them by the subnet.

If users are behind NAT (such as CGNAT, NAT44, etc.), and requests from different end users come from the same IP address, we can identify the end user by the combination of IP address and port, or by a range of ports.

4. Components

(((SAFEDNS



Our solution consists of the following components:

- DNS Proxy Module: Receives DNS requests and returns responses in the form of IP addresses.
- Internal Database: Stores information about all users and filtering policies.
- RestAPI: Used to update the Database with:

- Creation/modification of filtering profiles (specifying which categories should be blocked)
- Creation/modification of block pages
- Creation/updating of users or user groups, including:
 - Identifier (subnet/IP/port)
 - Filtering profile
 - Block page

- Block Pages
- Binary Log Parsing Module for DNS proxy (StatsLoader)
- Statistics Export Module to Clickhouse DBMS.

Explanation of other elements of the set up:

Client: Represents the end user who will use the DNS filtering service. The client sends DNS requests to SafeDNS Shield.

DB Files: This is the database that stores information about clients and filtering rules. RestAPI retrieves information from this database.

DNS Proxy: This is the core component of the DNS filtering process. It receives DNS requests from the client, applies filtering rules based on the data from DB Files, and then forwards the requests to the caching DNS.

Caching DNS Server: This is a cache server deployed on the client's side. It performs the actual resolution of DNS requests after they have been filtered through the DNS Proxy.

HostPATH: After processing DNS requests, the DNS Proxy generates binlog files (log files containing information about DNS queries). HostPATH is responsible for storing these binlog files.

StatsLoader: Once the binlog files are stored, StatsLoader retrieves and analyzes the binary data from HostPATH, processes the statistics, and sends them to the ClickHouse cluster for storage and analysis.

Load Balancer (LB): StatsLoader sends processed statistics to the load balancer, which distributes them across ClickHouse nodes.

ClickHouse Cluster: This is a distributed database used for storing and analyzing DNS request statistics. The cluster is divided into shards, with each shard containing multiple ClickHouse nodes for fault tolerance (mirroring). Shard 1 and Shard 2 are designed for high performance, as reading and writing are done in parallel.

Zookeeper: These instances manage the configuration and coordination of the ClickHouse cluster, ensuring data consistency and the reliability of the distributed system.

5. Product Setup and Interaction

All setup, maintenance, and support are handled exclusively by SafeDNS specialists, which clearly defines the areas of responsibility between SafeDNS and the client. The client provides the hardware (a server with x64 architecture) according to the requirements, along with remote access to the server, after what we handle the full setup "turnkey".

Our specialists, in collaboration with the client, carry out the initial configuration of filtering rules and provide training for the client's staff for future adjustments.

If needed, a dedicated support line is made available for clients using On-Prem solutions.

6. Working with Statistics

Binary logs are stored on the SafeDNS Security server. Our solution includes a module that parses these logs and exports them to an external DBMS for further analysis and report generation. Currently, there is a connector for exporting logs to the ClickHouse DBMS, which provides the best performance for this type of data. However, if needed, logs can be exported to other databases as well. The statistics module is not essential for the core functionality of SafeDNS Security, but without it, it is impossible to assess performance or investigate incidents effectively.

7. System Requirements for the Product

The minimum system requirements for a server running the SafeDNS Security module are as follows:

- CPU: 4 cores (x64 architecture)
- RAM: 8 GB
- 200 GB NVMe
- OS: Debian 11

These specifications are sufficient to support all system modules and handle up to 1,000 requests per second.

With 64 CPU cores and 128 GB RAM, the system is capable of processing up to 2 million requests per second, provided the network interface bandwidth is sufficient.

The minimum system requirements for a single server for the ClickHouse DBMS cluster:

- CPU: 8 cores (x64 architecture)
- RAM: 16 GB
- 500 GB NVMe (up to 2 TB depending on client traffic volume)

We recommend a minimum cluster configuration of 4 servers arranged in a 2x2 setup: two shards for parallel read/write operations, with two servers in each shard for redundancy. Optionally, for lower traffic volumes, a standalone ClickHouse server can be deployed, eliminating the need for the Zookeeper module.