# Tencent Cloud EdgeOne

# Common Guidelines

# Product Documentation

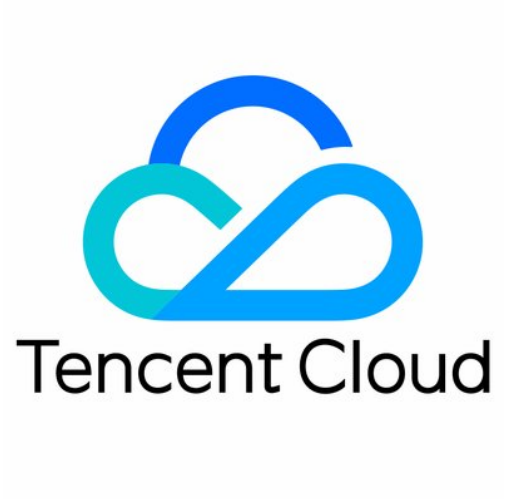

# Contents

# Common Guidelines
# FAQs
# Product Features FAQs

Last updated：2023-10-13 14:27:26

### How can I connect my site to EdgeOne?

EdgeOne supports NS and CNAME connection.

### What security capabilities does EdgeOne have?

It can prevent web application layer, DDoS, CC, bot, and crawler attacks and allows you to configure complicated custom access control rules based on your business needs.

### Does EdgeOne support cross-region accerlation?

EdgeOne deploys edge nodes in to fully meet your cross-region business needs. For specific available regions, contact us.

### Does EdgeOne support sites not deployed on Tencent Cloud?

Yes. For more details, please contact us.

### Does EdgeOne support API operations?

Yes. EdgeOne supports TencentCloud API and Terraform API.

### Does EdgeOne support dynamic acceleration?

Yes. It supports scenarios where requests for dynamic/static hybrid resources need to be accelerated. It can optimize the request response time and stability to deliver a high-quality and smooth access experience for websites.

### What site business security protection capabilities does EdgeOne offer?

EdgeOne provides web and bot protection for HTTP and HTTPS-based website businesses. Specific web protection rules include those for web security, OWASP rules, custom characteristics, and frequency control.

### What non-site business security protection capabilities does EdgeOne offer?

EdgeOne provides DDoS attack protection for TCP and UDP applications with specified ports, such as detection and protection against common types of DDoS attack, filtering rules by port, protocol, source IP region, and custom packet characteristic, and UDP watermark protection (coming soon).

# Domain Service and DNS FAQs

Last updated：2025-01-07 18:03:36

## Why do I get a CNAME and MX record conflict prompt when adding a DNS resolution record?

Take `example.com` as an example.

| Record type | Host | Value |
| --- | --- | --- |
| MX | www | mx.mail.com |
| CNAME | www | test.edgeone.com |

When performing a recursive resolution query, each record type has different priority, and CNAME has the highest priority. See RFC1034 and RFC2181. Therefore, during the resolution request process, the CNAME resolution record result will be returned first. When the host record value is the same, CNAME record and MX record cannot be configured at the same time, and you will get a prompt about the conflict.

If you do need to add CNAME and MX records at the same time when the host record is @, EdgeOne allows you to configure CNAME and  records at the same time:

| Record type | Host | Value |
| --- | --- | --- |
| MX | @ | mx.mail.com |
| CNAME | @ | test.edgeone.com |

**Reminder:**
This configuration will lead to unstable mailbox reception. If the Local DNS of the mailbox server prioritizes the resolution of the CNAME type of the @ record, the resolution of the MX type of the @ record will be affected, resulting in a resolution failure. If the host record is not @ , but the MX and CNAME records still indicate a conflict, please refer to the description of other record type conflicts below.

## Why do I get a CNAME and TXT record conflict prompt when adding a DNS resolution record?

Take `example.com` as an example.

| Record type | Host | Value |
| --- | --- | --- |
| TXT | www | edgeone-txt-flag |
| | | |

| | | |
|---|---|---|
| CNAME | www | test.edgeone.com |

The CNAME record has the highest priority, so if the host record is the same, configuring the CNAME record and the TXT record at the same time may cause the TXT record to fail to be parsed. **In this case, EdgeOne will prompt record conflict**.

If you do need to add CNAME and MX records at the same time when the host record is @, EdgeOne allows you to configure CNAME and TXT records at the same time:

| Record type | Host | Value |
|---|---|---|
| TXT | @ | edgeone-txt-flag |
| CNAME | @ | test.edgeone.com |

**Reminder:**
This configuration will cause the TXT verification to fail, you can remove the CNAME record to solve this problem. TXT and CNAME records will still conflict when the host record is not @.

### How do the record types conflict with one another?

See below for details:
✓: No conflict. When the HOST is the same, these two record types can both be configured. For example, after configuring the A record for `www.example.com`, you can still configure the MX record.
×: Conflict. When the HOST is the same, these two record types cannot be both configured. For example, after configuring the A record for `www.example.com`, you can not configure the CNAME record.

| Record type | A | AAAA | CNAME | MX | NS | TXT | SRV | CAA |
|---|---|---|---|---|---|---|---|---|
| A | ✓ | ✓ | × | ✓ | × | ✓ | ✓ | ✓ |
| AAAA | ✓ | ✓ | × | ✓ | × | ✓ | ✓ | ✓ |
| CNAME | × | × | × | × | × | × | × | × |
| MX | ✓ | ✓ | × | ✓ | × | ✓ | ✓ | ✓ |
| NS | × | × | × | × | ✓ | × | × | × |
| TXT | ✓ | ✓ | × | ✓ | × | ✓ | ✓ | ✓ |
| SRV | ✓ | ✓ | × | ✓ | × | ✓ | ✓ | ✓ |
| CAA | ✓ | ✓ | × | ✓ | × | ✓ | ✓ | ✓ |

**Note**

The table above shows the conflict relationship when the HOST is not @. If the HOST is @, a CNAME record does not conflict with an MX or TXT record.

### When the record type is A/AAAA/CNAME, can I configure both the resolution and acceleration when the HOST is the same?

Take the following configuration as an example:

| Record type | Host | Value |
| --- | --- | --- |
| A | www | 1.1.1.1 |
| A | www | 2.2.2.2 |

In this case, if you want to enable acceleration for one record, there will be a conflict. To enable acceleration for `1.1.1.1`, you need to delete `2.2.2.2` first.

**Note**

The above conflict happens on A/AAAA/CNAME records.

### Which access ports are supported by default?

EdgeOne supports the access ports 80, 8080, and 443 by default.

### What happens if there are multiple A/AAAA/CNAME records with the same host record name and no weights are set?

If there are multiple A/AAAA records with the same host record name and no weights are set, all the A/AAAA records will be returned. If there are multiple CNAME records with the same host record name and no weights are set, one of the CNAME records will be returned using a round-robin method.

### How long does domain configuration take effect?

Domain configurations generally take effect within 5 minutes after being issued. Some configurations may take 5-15 minutes to take effect due to the large number of configuration tasks. Please wait patiently.

# Site Acceleration FAQs

Last updated：2025-03-14 16:49:47

## How do I tell whether user access has hit the EdgeOne cache?

EdgeOne identifies whether a request hit the cache via EO-Cache-Status.

Open in browser

Curl command

Open the console in the browser and access the request URL (such as https://example.com/test.webp). Check the response header. If the value of `EO-Cache-Status` is `HIT`, the cache is hit.

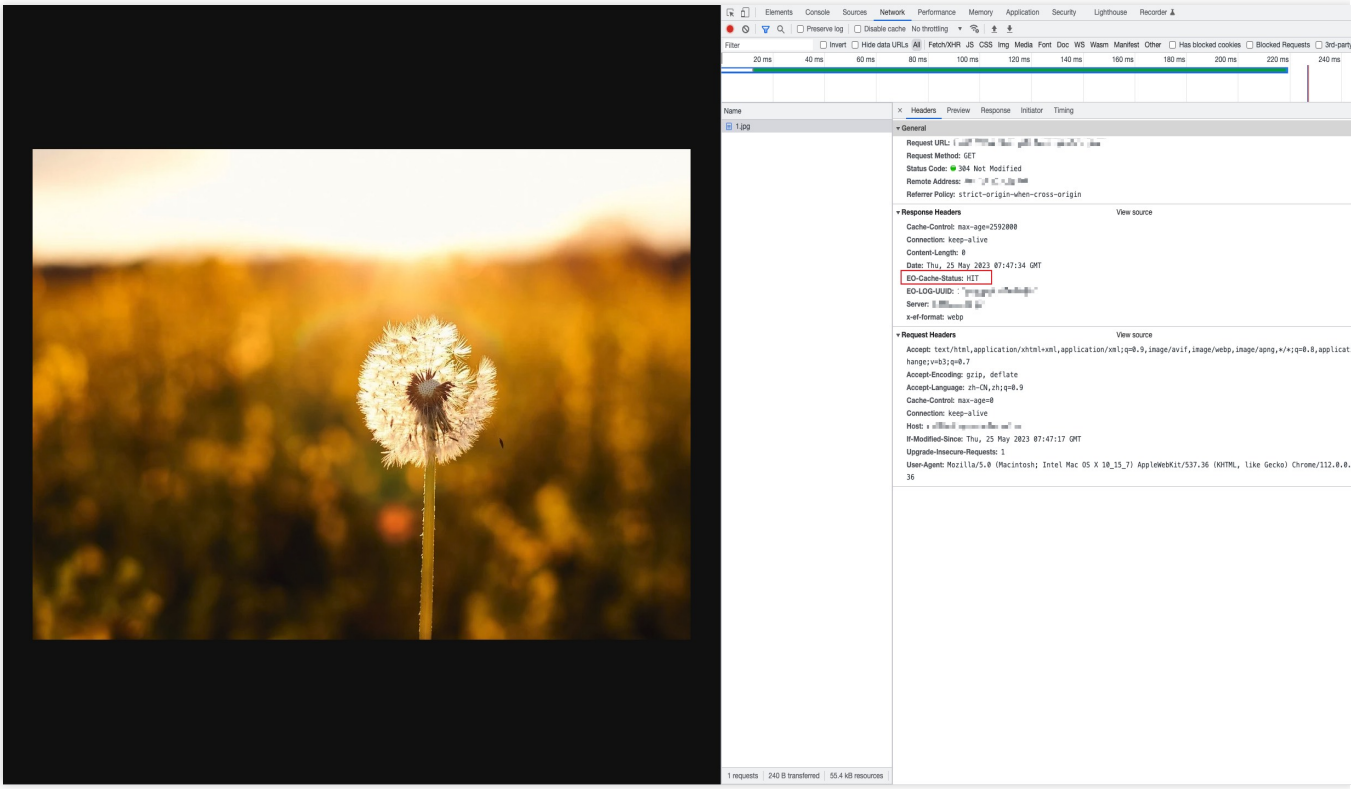


For Mac/Linux OS, you can use curl command to verify (such as `curl https://example.com/test.webp -i`). Check the response header. If the value of `EO-Cache-Status` is `HIT`, the cache is hit.


## How to handle cross-origin errors when prefetching?

Since resources are prefetched through URLs, cross-origin headers are not required. When a cross-origin request is initiated, the request fails because these headers are not present in the cache.

To enable cross-origin support for your resources when prefetching, you can customize the HTTP response header in EdgeOne.

## How long does it take for Cache Purge and Cache Pre-Warming to take effect after each content submission?

**Cache Purge：**

| Type | Single Submission quantity | Effective Time |
| --- | --- | --- |
| URL | ≤5000 URLs | Normally completed within 1 minute. If the number of submissions is large, it may take 3 - 5 minutes |
| Directory | ≤1000 directories | |
| Hostname | ≤1000 Hostname | |
| Cache-Tag | ≤100 Cache-Tag | |
| All Cache | - | |

**Cache Pre-Warming:**

| | | |
| --- | --- | --- |

| Type | Single Submission quantity | Effective Time |
| --- | --- | --- |
| URL | ≤5000 URLs | 5-30 minutes |

**Note：**

1. When the cache TTL configured for a file is less than 5 minutes, it is suggested not to use the purge tool, but to wait for the timeout update.
2. The actual total time for any type of cache purge mainly depends on the quantity of submitted content, the more content, the longer the waiting time.
3. The actual total time for cache pre-warming mainly depends on the file size, the larger the file, the longer the waiting time. The pre-warming effective time for more large files (≥100MB) may be extended, exceeding 30 minutes.

# Data and Log FAQs

Last updated：2024-07-30 15:25:56

## Why are the traffic data in the console and the traffic data derived from logs inconsistent?

The traffic data derived from the byte count recorded in the `EdgeResponseBytes` field of the site acceleration access logs may not match the traffic data displayed on the console and the billing traffic data. The reasons are as follows:

Access logs can only record application layer data. In actual network transmission, the network traffic generated is 5-15% more than the pure application layer traffic. It consists of two parts:

Consumption by TCP/IP packet headers. In TCP/IP-based HTTP requests, each packet has a maximum of 1,500 bytes, including TCP and IP headers of 40-60 bytes, which generate traffic that cannot be counted by the application layer. The overhead of this part is approximately 3-4%.

TCP retransmission. During normal network transmission, around 3-10% of packets are lost on the Internet, and the server retransmits the lost packets. This type of traffic cannot be counted by the application layer and constitutes approximately 3-7% of the total traffic.

## The monitoring data I see in Tencent Cloud Observability Platform and EdgeOne are not the same.

Data trends on Tencent Cloud Observability Platform and EdgeOne are generally consistent. However when it comes to the 1-minute granularity, the data can be slightly different. See below for details:

**Tencent Cloud Observability Platform:** Collect data from edge servers and aggregate the data with 1-minute granularity on the domain name level. This can guarantee the timeliness and stability. But it only provides data related to key metrics on the domain name level.

**EdgeOne**: Collect and analyze logs in real-time upon receiving the request, and then print out the result. It supports more metrics, such as traffic and requests by the device type and browser type. But the print-out time can be affected in case of request surges.

Assume that a user requests a 1 GB file. The download starts at 10:00:00 and ends at 10:01:40.

**Tencent Cloud Observability Platform:** Every edge server reports the metric data at a 1-minute interval. Data of this event is recorded at both 10:01 and 10:02.

**EdgeOne**: Every edge server prints a log when the download ends (10:01:40). The data is recorded at 10:01.

Therefore, data from Tencent Cloud Observability Platform and EdgeOne can differ at a 1-minute granularity due to the difference of sampling rules.

# Security Protection-related Queries

Last updated：2024-10-18 11:43:55

## What Security Features Does EdgeOne Have?

EdgeOne provides reverse proxy and protocol-specific security protection for Web application services and TCP/UDP application services.

| Access Service Type | L3/L4 DDoS Protection | HTTP DDoS Protection (L7 CC Attack Protection) | Web Protection | Bot Management |
|---|---|---|---|---|
| L4 Proxy (TCP/UDP Application Service) | ✓[1] | - | - | - |
| L7 Zone(Web Application Service) | ✓[1] | ✓ | ✓ | ✓[2] |

**Note:**

Note 1: Default platform-level protection is provided. If you have specific protection capacity requirements, please use Exclusive DDoS Protection Usage.

Note 2: Bot Management subscription is required; see Billing Overview (New Version).

## I've already configured a Web Application Firewall (WAF) on my origin server. Do I need to use EdgeOne security protection?

EdgeOne aims to provide integrated acceleration and security capabilities. Therefore, when you connect your application and services to EdgeOne, EdgeOne starts providing protection services. In addition to the protection already in place on your origin server, EdgeOne offers:

Distributed Security Protection: Provides protection resources distributed in multiple independent cleansing centers worldwide, offering efficient redundancy and disaster recovery through a distributed access architecture.

WAF and Web Site Protection: Provides application security protection features such as vulnerability attack detection, rate limiting, and Bot management[3].

Protection Capability for Cached Resources: Can simultaneously check requests accessing cached resources. The usage of security policies intercepted by EdgeOne is not billed, reducing unnecessary content delivery costs.

Identification of Threats Closest to the Client: Generally, an access request is directly initiated by a client. EdgeOne can collect and analyze L4 connection session characteristics and TLS fingerprint characteristics of the client, which

are used together with policies to identify malicious access.

Compatibility with Your Origin Server Security Policies: Supports marking of origin-pull requests [3]allowing further analysis of requests at the origin server.

**Note:**

Note 3: You need to subscribe to and enable Bot Management. Bot Management includes identification headers in origin-pull requests to assist in further analysis.

## Does EdgeOne Support IP Blocklists/Allowlists?

If you need to configure an IP blocklist (i.e., block specified client IPs), you can configure the **Basic Access Control** in Custom Rules, select **Client IP Control**, configure the list of IPs to be blocked, and choose the blocking method.

If you need to configure an IP allowlist (i.e., allow specified client IPs), you can use Exception Rules, select the **Client IP** matching condition, and choose the security modules to be skipped.

**Note:**

The application scenarios for an IP allowlist may vary:

(1) Allow specified client IPs to pass. In this scenario, configure Exception Rules to skip specified security modules.

(2) Only allow specified client IPs to access. In this scenario, configure Basic Access Control rules in Custom Rules to block client IPs not in the specified list.

## How to configure region blocking? How to block access from regions outside the Chinese mainland?

You can use **Basic Access Control** in Custom Rules, select **Regional Control**, configure the list of client regions to be blocked, and choose the blocking method. If you need to block access from regions outside the Chinese mainland, select the **Region Mismatch**, match the content to Chinese mainland region, and choose the blocking method.

## How to configure Hotlink Protection? How to allow access only from this domain and specified domains?

Hotlink protection is mainly used to prevent static resources from being loaded by external website pages.

### Common Hotlink Protection Techniques

The basic hotlink protection policy judges whether the request comes from page loading through the Referer header, intercepting requests for resources referenced by external sites and requests not accessed directly through page loading (example: directly accessing static resources by entering the URL in the browser). You can use **Basic Access Control** in Custom Rules to block requests with a Referer header not in the specified domain list.

### Further Validation of Data Access Security

Using HTTP header fields can address common hotlinking scenarios, but malicious requests can still generate legitimate HTTP requests through technical means to obtain site resources. To further improve the security of resource access, you can dynamically generate URLs with time-sensitive random signatures. Before providing access to resources, verify the legality and validity of the signature to identify whether the request has permission to access resources. EdgeOne's Rule Engine offers Token Authentication options, assisting in generating signed URLs and providing a signature verification mechanism. You can also use EDGE-FUNCTION to implement custom dynamic access authentication.

### What is "Monitor," and does the "Monitor" action involve interception?

The "Monitor" action only logs information and does not intercept requests. This is helpful for evaluating policies, as rules set to "Monitor" won't impact your business. Therefore, you can assess the impact on normal business and evaluate matching situations with malicious requests by checking the logs. This helps determine whether to enable interception. See Actions for more details.

### What is "JavaScript Challenge," and what impact does the "JavaScript Challenge" action have on business?

The "JavaScript Challenge" action responds with a page that verifies whether the requesting client supports Cookie and JavaScript runtime environments. Browsers that meet the verification conditions can proceed with access, while other tools (example, cURL) will be intercepted. This method helps identify some non-browser tools.

**Note:**

1. Most APIs cannot handle JavaScript responses, so they will be blocked by the "JavaScript Challenge" action.
2. Native APP and mini program requests are also API requests and cannot handle JavaScript responses, so they will be blocked by the "JavaScript Challenge" action. For compatibility with the JavaScript Challenge, clients can use a web-view or H5 framework to pass the JavaScript Challenge before accessing APIs, to avoid API access requests from being blocked by the "JavaScript Challenge".

### Can the Origin Server Be Accessed Through All Ports Opened by EdgeOne?

By default, ports opened by EdgeOne do not provide access to site business. After access requests are parsed based on the protocol and port configuration of the accessed site, EdgeOne will decide whether to handle requests from specified ports and then respond by blocking or origin-pull based on the security and acceleration configuration.
If no port is enabled or no access protocol is configured for your business, the domain name resource or origin server of the business cannot be accessed from a client through the port or protocol:
For site domain names that have not completed the access process, the domain name business cannot be accessed from EdgeOne through protocols such as HTTP, HTTP/2, and QUIC.
If HTTPS, HTTP/2, and QUIC services or corresponding certificates are not enabled or configured, the corresponding domain name resource or origin server cannot be accessed from a client through HTTPS, HTTP/2, and QUIC protocols.

If no L4 proxy forwarding rule is configured for a specified port, the corresponding port business of the origin server cannot be accessed through the port.

**Note:**

After domain name service is accessed, EdgeOne by default supports accessing HTTP services of sites through specified ports. For details, see Domain Service FAQs.

# Origin Configuration FAQs

Last updated：2025-01-15 10:44:22

### How to Fill in the Weight in the Origin Server Group?

When a weight is set for an origin server in the origin server group, all other origin servers must also have weights set. The weight can be an integer from 0 to 100. It is not supported to set weights for only some origin servers.

### How Does It Work When an Origin Server in the Origin Server Group Has a Weight of 0?

When you set the weight of an origin server in the origin server group to 0, no traffic will be directed to that origin server. EO will route traffic back to other origin servers based on their weighted proportion out of 100.

### How Does the Back-to-Origin Logic Work When Different Weights Are Configured for Multiple Origin Servers in the Origin Server Group and Intelligent Acceleration Is Enabled for the Corresponding Domain Name?

When using the combination of Smart Acceleration and origin server group weights, the following logic will apply:

| Scenario | Activation Logic |
| --- | --- |
| Configure Weights for Multiple Origin Servers in the Origin Server Group and Enable Intelligent Acceleration for the Corresponding Domain Name | Prioritize selecting the origin server based on weight, then Smart Acceleration will optimize the return link. |
| Configure Weights for Multiple Origin Servers in the Origin Server Group without Enabling Intelligent Acceleration for the Corresponding Domain Name | Origin-pull by weight ratio. |
| Multiple origin servers in the origin server group without configuring weights, enable intelligent acceleration for the corresponding domain name | Origin-pull from the optimal origin server selected by Smart Acceleration. |
| Multiple origin servers in the origin server group without configuring weights, do not enable intelligent acceleration for the corresponding domain name | Poll each origin server in the origin server group, proportional origin return. |

### What is the impact on existing connections when changing the origin server group configuration?

When you adjust the weight of an origin server in the origin server group from a value greater than 0 to 0 or directly delete an origin server, existing established return connections will not be affected. New connections will not be forwarded to the origin server with a weight of 0 or the deleted origin server.

# Troubleshooting
# Reference for Abnormal Status Codes

Last updated：2025-04-18 18:06:43

EdgeOne responds with the following types of exception status codes:

## Standard Status Codes

**You can refer to the HTTP Status Code Standard to see the specific meanings of these status codes.**
**These codes are typically responded to in the following situations:**
After the request is sent back to the origin, the status code information is responded by the origin server, and the node will pass transmit status code from the origin server to the client.
Direct responses from EdgeOne nodes, for example, Token authentication fails, responding with a 403 status code.
Possible status codes that may be directly responded by EdgeOne nodes include the following status ones:

| Status Code | Description |
| --- | --- |
| 400 | The client's request is deemed invalid because the requested method is not within the allowed range set by EdgeOne. For details, see HTTP Restrictions. |
| 403 | Fail to pass the link protection verification, such as the Token authentication in the rule engine. |
| 416 | Range exception, for example, rangeStart < 0,< 0, rangeStart > rangeStart > rangeEnd, or rangeStart > FileSize. |
| 418 | For each domain name connected to EdgeOne, the system automatically allocates service nodes for the domain. Each corresponding node receives the domain's configuration file, which includes settings such as the origin server, cache, headers, etc. When a request is sent to a node, it reads the domain's configuration file. If the file is not found, the node responds with a 418 status code. For example, if a client sends a request to `http://example.com/test.jpg` , the node reads the configuration file for the domain `example.com` . If, due to reasons such as the service node binding to a `non-example.com` domain name, CNAME configuration errors, or issues with the scheduling system, the configuration file is not found, the client will receive a 418 response. |
| 423 | Request loopback is triggered, that is, the Loops value in the CDN-Loop header is greater than or equal to 16. For details, see CDN-Loop. |

## EdgeOne Custom Status Codes

EdgeOne defines special status codes with unique meanings. Status codes within the range of 520-599 are reserved by EdgeOne for custom non-standard responses. It is recommended to avoid using status codes within this range in your business to prevent confusion with EdgeOne's status codes.

Below are the custom EdgeOne status codes along with their meanings, facilitating self-troubleshooting when encountering abnormal business access,please refer to: Troubleshooting Guide for EdgeOne 4XX/5XX Status Codes.

| Status Code | Meaning Explanation |
| --- | --- |
| 520/550 | After the node successfully establishes a connection with the origin server and sends a request, if the origin server directly sends a RST packet, the node responds to the client with a 520/550 status code. |
| 521/551 | When the node requests the origin server, during the TCP connection establishment phase, if the origin server directly sends a RST packet, the node responds to the client with a 521/551 status code. |
| 522/552 | When the node requests the origin server, during the TCP connection establishment phase, if the origin server does not respond, causing the node to time out, the node responds to the client with a 522/552 status code. |
| 523/553 | If the domain is configured with a domain name as the origin server, when the node goes to the origin, it needs to resolve the domain to obtain the origin server's IP. If the resolution fails, the node cannot go to the origin, and it responds to the client with a 523/553 status code. |
| 524/554 | After successfully establishing a connection with the origin server, if the node initiates a request to the origin server and there is no response from the origin server, causing a timeout at the node, the node responds to the client with a 524/554 status code. |
| 525/555 | If the origin protocol is HTTPS, the node needs to perform an SSL handshake with the origin server when going back to the source. If the handshake fails, the node responds to the client with a 525/555 status code. |
| 566 | When a request is intercepted by Web Protection - Managed Rules, the default response is a 566 status code along with the default interception page. If the user has configured a custom interception status code, the configured status code will be used. |
| 567 | When a request is intercepted by Web Protection - Custom Rules, Web Protection - Rate Limiting or Bot Management rules, the default response is a 567 status code along with the default interception page. If the user has configured a custom interception status code, the configured status code will be used. |

# Troubleshooting Guide for EdgeOne 4XX/5XX Status Codes

Last updated：2025-01-14 15:55:44

After connection to EdgeOne, if your business request encounters a 4XX/5XX status code, you can troubleshoot by referring to the following common causes and solutions.

**Note:**

If you are still unable to resolve the issue by referring to the following troubleshooting guide, you can contact us. To facilitate troubleshooting, please provide EO-LOG-UUID information. If it is unavailable, please provide the following details:

User IP and EdgeOne node IP;

Specific 4XX/5XX error code and message;

Time and timezone for the 4XX/5XX error;

The URL causing the HTTP 4XX/5XX error (for example, https://www.example.com/images/icons/image1.png:).

## HTTP 400

1. Meaning: The server cannot or will not process the request due to certain causes that are considered client errors (such as request syntax errors, invalid request message formats, or deceptive request routing).

2. Possible causes and solutions:

Origin server response: Test direct access to the origin server. If the origin server responds with 400, then modify the origin server configuration or adjust the client request behavior to obtain a correct response.

EdgeOne node response:

Check whether the request method is within the scope supported by EdgeOne. If it is not one of the following request methods, EdgeOne nodes will directly respond with a 400 status code.

GET, HEAD, POST, PUT, DELETE, TRACE, CONNECT, OPTIONS, PATCH, COPY, LOCK, MKCOL, MOVE, PROPFIND, PROPPATCH, and UNLOCK.

Protocol format errors, such as accessing HTTP via Port 443.

The HTTP request does not comply with RFC standards.

## HTTP 403

1. Meaning: The server cannot or will not process the request due to certain causes that are considered client errors (such as request syntax errors, invalid request message formats, or deceptive request routing).

2. Possible causes and solutions:

Origin server response: Test direct access to the origin server. If the origin server responds with 403, then modify the origin server configuration or adjust the client request behavior to obtain a correct response.

EdgeOne node response: A common cause is failure to pass the hotlink protection verification.

Check whether HTTP response has been configured in the rule engine. If so, verify whether it hits basic access controls such as referer blocklist/allowlist, IP blocklist/allowlist, and user-agent blocklist/allowlist.
Check whether token authentication has been configured in the rule engine. If so, verify whether the URL in the client request has an expired timestamp or an incorrect MD5 encrypted string.
Check whether remote authentication has been configured through the edge function. If so, verify whether the URL in the client request does not comply with the authentication and release rules of the authentication server.

### HTTP 416

1. Meaning: The server cannot process the requested data range. The most common situation is that the requested data range is not within the file range,
2. Possible causes and solutions:
Origin server response: Test direct access to the origin server. If the origin server responds with 416, then modify the origin server configuration or adjust the client request behavior to obtain a correct response.
EdgeOne node response:
Abnormal range values, such as: rangeStart < 0, rangeStart > rangeEnd, and rangeStart > FileSize.
The range header of the client request is not standard. For example, the request is Range: Bytes=0-1023 instead of Range: bytes=0-1023.

### HTTP 418

1. Meaning: This response is typically used by servers to handle requests they do not want to deal with.
2. Possible causes and solutions:
Origin server response: Test direct access to the origin server. If the origin server responds with 418, then modify the origin server configuration or adjust the client request behavior to obtain a correct response.
EdgeOne node response: For a domain name connected to EdgeOne, the system automatically assigns service nodes to the domain name, and the corresponding nodes issue the configuration file for the domain name. The file content depends on the domain name's configuration, such as the origin server, cache, and header. When a request reaches an EdgeOne node, the node will read the configuration file for the domain name. If the configuration file is not found, a 418 status code will be returned.
Example: A client requests the domain name `http://example.com/test.jpg` that has been connected to EdgeOne, and the relevant request configuration for the domain name exists on nodes `1.1.1.1` and `1.1.1.2`. When a domain name request reaches these two nodes, they will read the configuration file for the domain name `example.com` and respond successfully. However, if a request reaches node `2.2.2.2`, a 418 status code will be returned. Check whether the domain name is bound to a service node without the domain name `example.com` and whether the CNAME configuration is correct. If there are no issues, please contact us.

### HTTP 423

1. Meaning: It indicates "locked". In the scenario of connection to EdgeOne, it generally means request loopback is triggered.

2. Possible causes and solutions:
Origin server response: Test direct access to the origin server. If the origin server responds with 423, then modify the origin server configuration or adjust the client request behavior to obtain a correct response.
EdgeOne node response: Request loopback is triggered, that is, the CDN-Loop header's Loops value is ≥ 16. For details, see CDN-Loop. This is usually seen in accelerated domain names connected to EO, and accelerated domain names connected to EO or CDN are also set on the origin server.

### HTTP 520

1. Meaning: EdgeOne's custom status code. After the node successfully establishes a TCP connection with the origin server, it initiates an HTTP request to the origin server. However, the origin server directly sends an RST packet, and the node responds to the client with a 520 status code.

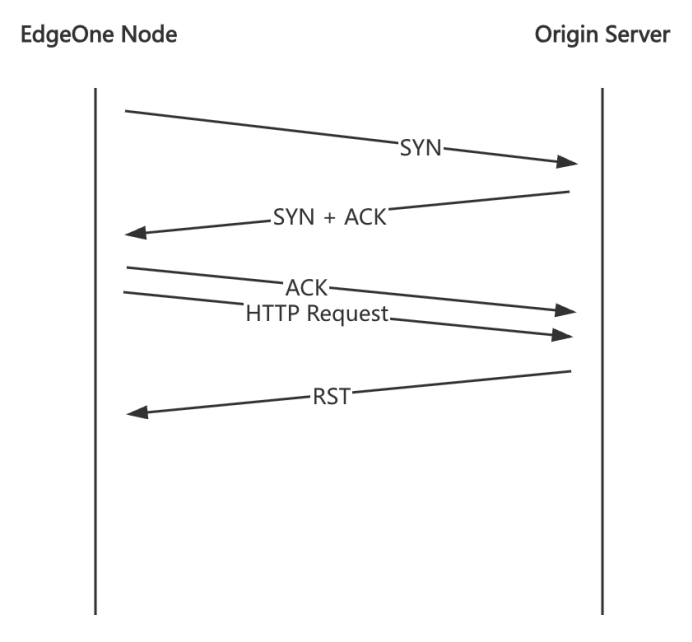


2. Possible causes and solutions:
Origin server service exception: Test direct access to the origin server and capture packets to check whether the origin server responds to the HTTP request with an RST packet. This may be caused by the origin server's firewall or service exception.

### HTTP 521

1. Meaning: EdgeOne's custom status code. When the node requests the origin server, if the origin server directly sends an RST packet during the TCP connection establishment phase, the node responds to the client with a 521 status code.

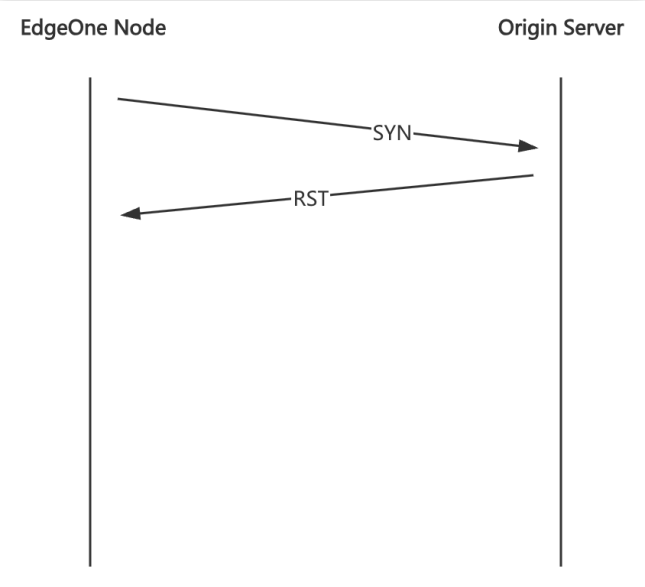


2. Possible causes and solutions:
Origin server service exception: Test direct access to the origin server. You can use a command tool such as curl or telnet to check whether the TCP connection can be established. This is usually caused by certain ports of the origin server not being open to the public network or origin server node network issues.

### HTTP 522

1. Meaning: EdgeOne's custom status code. When the node requests the origin server, if the origin server does not respond during the TCP connection establishment phase, causing the node to time out, the node responds to the client with a 522 status code.

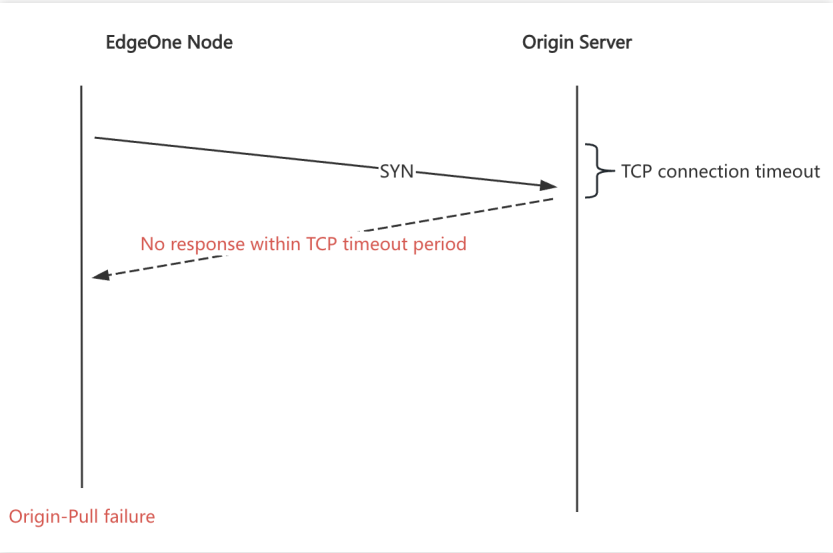


2. Possible causes and solutions:

Origin server service exception: Test direct access to the origin server. You can use a command tool such as curl or telnet to check whether the TCP connection can be established. This is usually caused by certain ports of the origin server not being open to the public network or origin server node network issues.

### HTTP 523

1. Meaning: EdgeOne's custom status code. If the origin server configured for the domain name is a domain name, then when the node requests the origin server, it needs to resolve the domain name to obtain the IP of the origin server. If the resolution fails, the node cannot request the origin and responds to the client with a 523 status code.
2. Possible causes and solutions:
Try using dig to resolve the origin server domain name to confirm whether it can be resolved normally.
If the origin server domain name can be resolved normally, please contact us.

### HTTP 524

1. Meaning: EdgeOne's custom status code. After the node successfully establishes a connection with the origin server and sends a request to the origin server, if the origin server does not respond, causing the node to time out, the node responds to the client with a 524 status code.

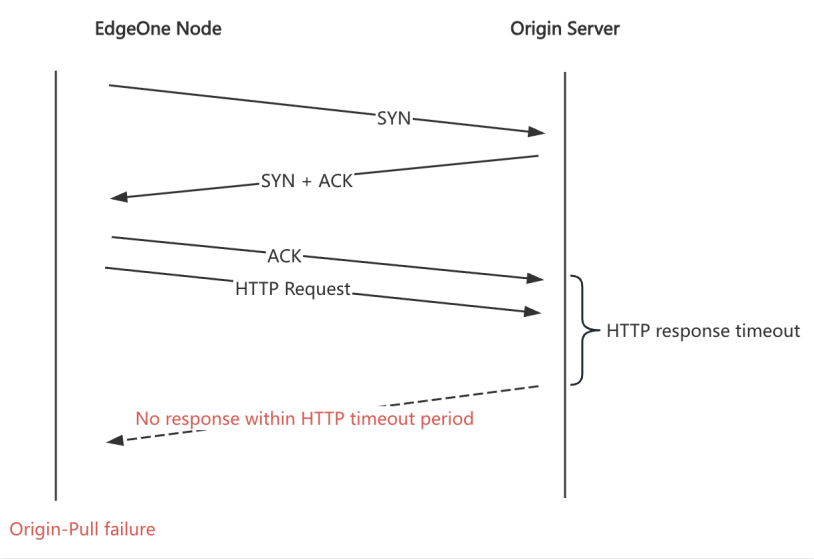


2. Possible causes and solutions:
Origin server service exception: Test direct access to the origin server to check for a response to the HTTP request. If the origin server responds normally, you can try adjusting the origin-pull timeout. If the issue persists, please contact us.

### HTTP 525

1. Meaning: EdgeOne's custom status code. If the origin-pull protocol is HTTPS, the node needs to perform an SSL handshake with the origin server when it requests the origin server. If the handshake fails, the node responds to the

client with a 525 status code.

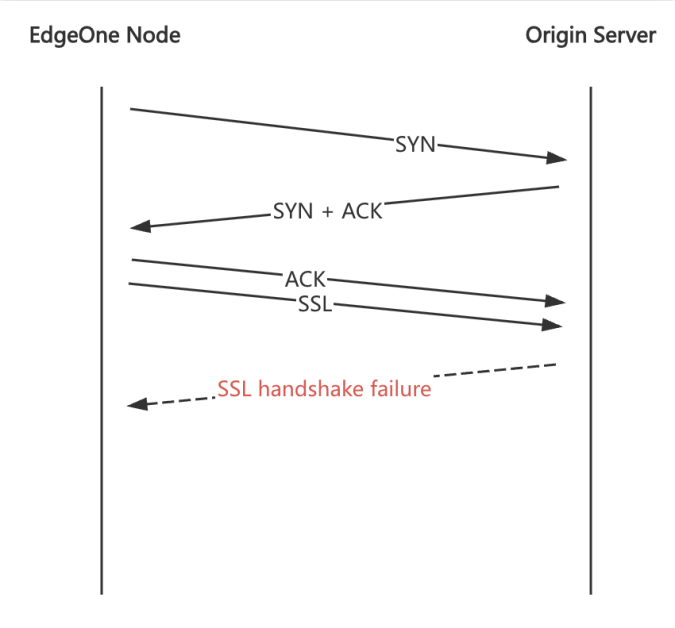


2. Possible causes and solutions:

Check whether the domain name's origin-pull protocol is configured as HTTPS, but the origin server has no certificate deployed. If so, you can change the origin-pull protocol to HTTP in the domain name management section or deploy the corresponding domain name certificate on the origin server.

Packet loss occurs due to network issues during the SSL handshake between the node and the origin server.

Failure to pass the origin-pull certificate validity verification. Currently, EdgeOne verifies the certificate's validity period by default.

# Tool Guide
# Speed Test Tools
# Real User Monitoring

Last updated：2025-01-14 15:50:55

**Note:**
The EdgeOne Performance Monitoring page was discontinued on November 21, 2024. Note that this change does not affect your existing performance monitoring services and data. You can continue to view and use related services on the Tencent Cloud Observability Platform - Real User Monitoring - Application Management page without worrying about service interruption or data loss.

## Overview

Real User Monitoring is a feature interconnected with EdgeOne. It provides one-stop frontend monitoring solutions. You only need to install its SDK to your project and complete simple configuration, and then it will take care of the user page quality in an all-around manner by monitoring the page performance and frontend quality in real time, truly enabling cost-effective usage and non-intrusive monitoring.
**Note:**
RUM provides a free tier of 500,000 reports per day for each application. Reports exceeding the free tier (500,000) will be billed. The fees are not part of your EdgeOne plan but are charged by RUM. For billing details, see Billing Overview.

## Use Cases

Page performance analysis: RUM offers metrics such as firstScreenTime, TCP connection establishment duration, time to first byte (TTFB), and SSL handshake duration. In addition, it supports latest Web Vitals standards, Google's webpage loading speed and experience metrics, helping you optimize the user experience in an all-around manner.
User access analysis: RUM displays the business PV/UV and top access metrics of each page. It analyzes the user access data in various dimensions including network, browser, and region, so that you can stay on top of and analyze the user access information.
Static resource speed test: RUM supports different types of resource speed tests on image loading, CDN resource operation, etc., so you can view diverse information such as resources used on a page and loading duration of each resource.

# Directions

1. Log in to the EdgeOne console and click **Speed Test Tools** > **Real User Monitoring** on the left sidebar.
2. If you enter the **Real User Monitoring** page for the first time, as this feature is based on EdgeOne and RUM, you need to click **Enable** to grant the relevant permissions.

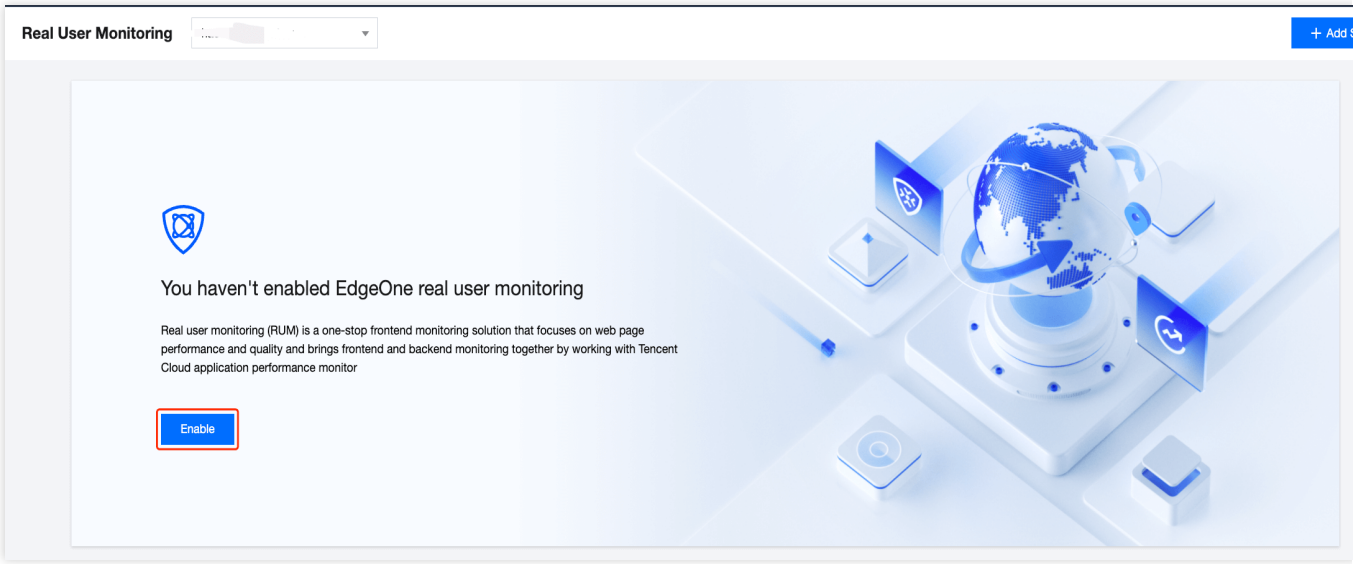


3. On the **Real User Monitoring** page, click **Application connection**.
4. In the **Application connection** window, enter the application name and description, select **I have understood the billing details**, and click **Next**.

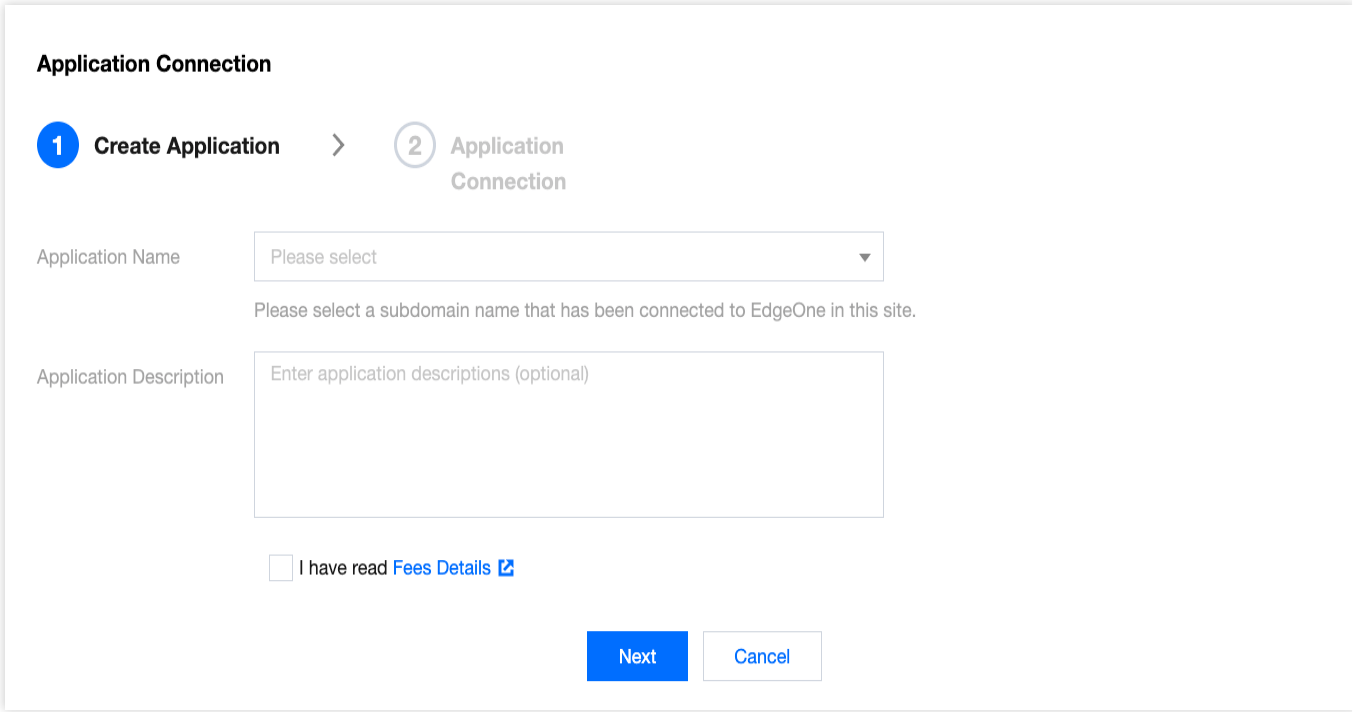

5. Install the SDK based on the connection type.

Install the SDK by importing the `<script>` tag

5.1.1 On the connection guide page, copy the provided `<script>` tag code.

5.1.2 Import the code below `<script> tag import` into the `<head></head>` tags of the site to be monitored.

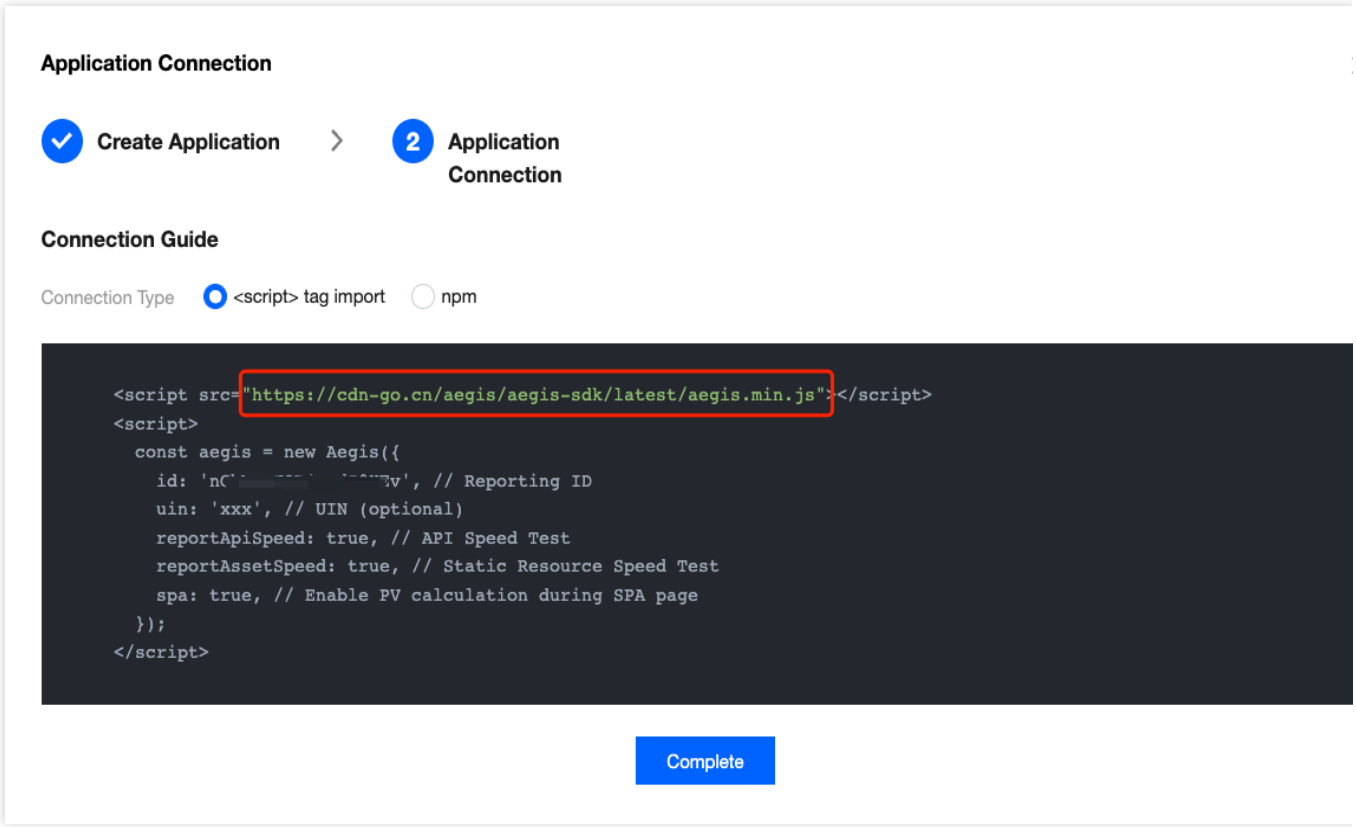


**Note:**

This connection method uses the “h3-Q050” protocol, where `cache-control` is `max-age=666` by default. To modify `cache-control`, you can add the `max_age` parameter, such as `<script src="https://cdn-go.cn/aegis/aegis-sdk/latest/aegis.min.js?max_age=3600"></script>`.

Install the SDK through npm

5.1 On the connection guide page, copy the first command line to import `aegis sdk` into your development environment.

5.2 Then, copy the provided code to initialize the SDK in your JavaScript code.

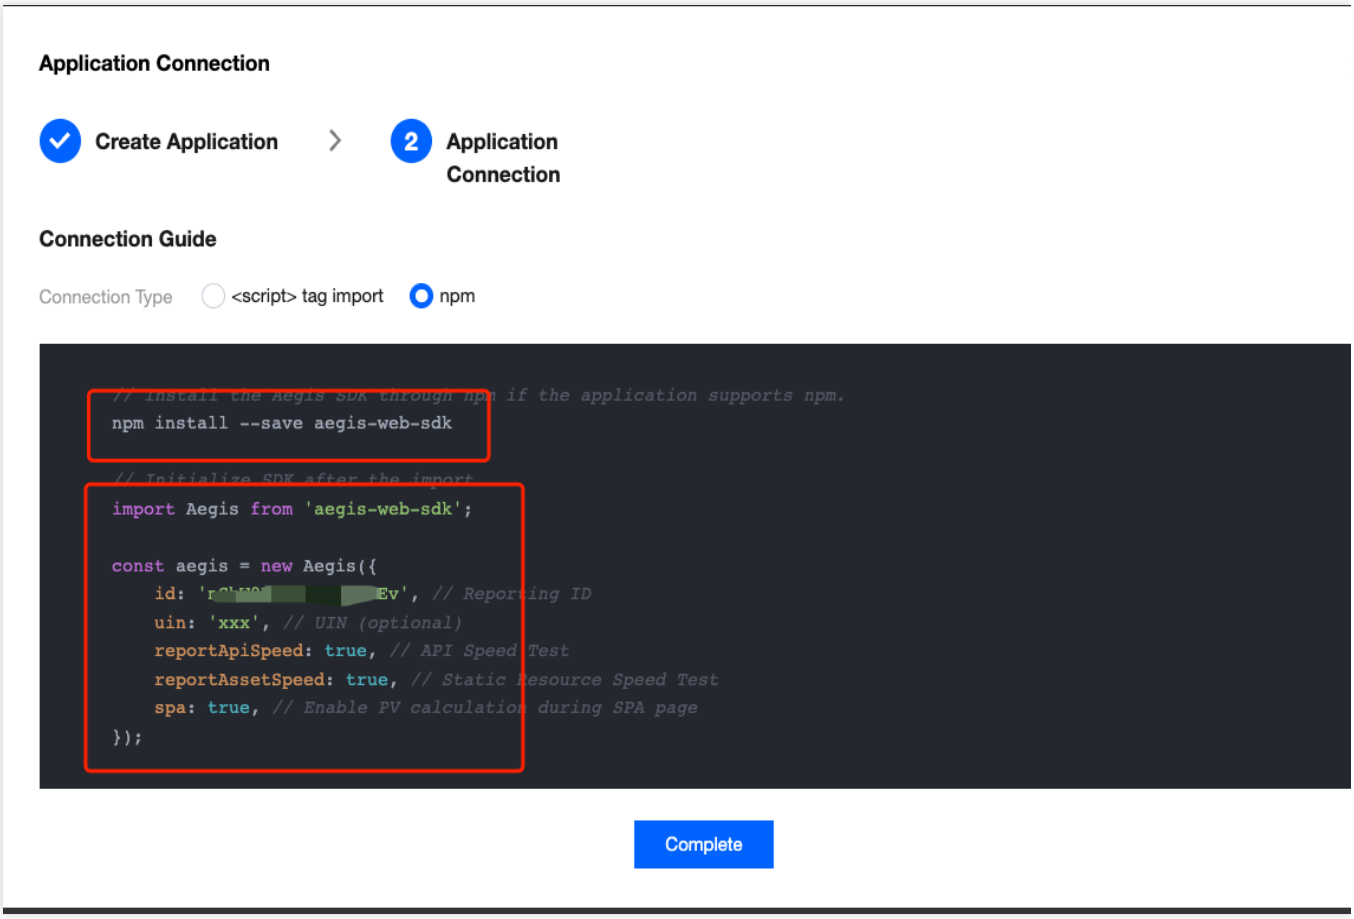


# Data Monitoring

After performing the above connection steps, go to the **Page performance**, **Page view**, and **Static resource** pages to view the relevant data.

### Page performance

The **Page performance** module supports multidimensional page performance analysis. You can analyze key page performance metrics such as firstScreenTime and request response through various views including performance change trend chart, page loading waterfall plot, and regional view. For more information, see Page Performance.

### Page view

The **Page view** module displays the page view information such as UV, PV, WAU and MAU, and supports multidimensional page access analysis. For more information, see Page View.

### Static resource

Frontend HTML pages mainly contain the following static resources: JavaScript, CSS, and image files. If such files fail to load, or it takes a long time to load them, the page will be affected or even crash. To address these problems, static

resource monitoring helps you analyze the frontend static resource status. For more information, see Static Resource.

# Diagnostic Tool
# Self-service debugging

Last updated：2025-03-26 15:53:34

## Feature Introduction

If you need to confirm whether the node cache rules, custom Cache Key, and other configurations currently configured in EdgeOne have taken effect for your resources, EdgeOne provides a self-service debugging tool to help you obtain node cache TTL, whether the resource is cacheable, Cache key, and other information, making it easy for you to debug your business configuration. After enabling self-service debugging, you can initiate a URL request from a specified client IP, carrying the `EO-Debug-Headers: all` header in the request, and view whether the resource is cached in the node, the corresponding Cache Key value, and cache time based on the returned response headers.

## Usage Scenarios

If you have configured more complex cache rules and custom cache keys in the rule engine of the console, and need to verify whether the configuration is effective, you can use this feature for verification.

## Directions

For example, the domain name `www.example.com` under the current site `example.com` has been configured to cache `.jpg` suffix files in EdgeOne nodes for 600 seconds; the cache Cache Key is configured to retain the specified parameter a as the cache key. After the configuration is completed, you need to verify whether the current configuration has taken effect, and you can follow the steps below to verify:
1. Log in to the EdgeOne console, click on the site list in the left menu bar, and click on the site to be configured in the site list.
2. On the site details page, click **Site Acceleration** to enter the global configuration page for the site. In the right-hand navigation bar, click **Self-service debugging**.
3. Locate the **Self-service debugging** configuration card and click the "Switch" to enable the Self-Diagnosis feature.

4. After enabling the debugging mode, you need to set the validity period and the allowed client source. The time range is 1-365 days, with a default of 7 days. The client IP allows for the input of 100 entries, accommodating both IPv4 and IPv6 IP/IP segments. The notation 0.0.0.0/0 signifies the permission for all IPv4 clients to execute debugging, while ::/0 indicates the allowance for all IPv6 clients to carry out debugging.

5. Click **Save**, and the allowed client IPs can debug within the effective time.
6. Initiate a curl request for verification from the specified client IP source in a Mac/Linux environment, for example: `curl -voa 'http://www.example.com/test.jpg?a=1' -H 'EO-Debug-Headers: all'` 。The request result is as follows：

In the response headers, you can see the corresponding Cache Key, cache status, and cache time for this request, which is consistent with the configuration in the example, indicating that the current configuration has taken effect.

## Related References

When the self-service debugging mode is enabled, the debug header explanations in the response are as follows:

| Header Name | Configurations | Meaning of the returned value |
| --- | --- | --- |
| EO-Debug-Status | Indicates whether the self-service debugging mode is enabled. | on: activated, and the request client IP is within the allowlist & the request time is within the validity period;<br>off: Off, or activated but the request time is beyond the validity period;<br>forbidden: activated, but the request client IP is not in the allowlist. |
| EO-Debug-Cacheable | The Request URL of this request, according to the configured EdgeOne node cache TTL, the final cacheable status of the Request URL resource in EdgeOne nodes. | yes：cacheable content<br>no: non-cacheable content |
| EO-Debug-CacheKey | The Request URL of this request, according to the custom Cache key, the final Cache key generated for the Request URL resource in EdgeOne nodes. | For example: `www.example.com/test.jpg a=1,` indicating the Cache Key generated for the Request URL resource in EdgeOne |
| EO-Debug-CacheTTL | The Request URL of this request, according to the configured EdgeOne node cache TTL, the final cache TTL | List values, including numbers and time units. d stands for days, h stands for hours, m |

| | | |
|---|---|---|
| | duration of the Request URL resource in EdgeOne nodes. | stands for minutes, and s stands for seconds, for example:<br>3d0h0m0s means the cache TTL is 3 days;<br>0d0h5m0s means the cache is 5 minutes;<br>0d0h0m5s means the cache is 5 seconds. |

# IP Location Query

Last updated：2023-04-10 18:14:58

This document describes how to verify whether an IP is owned by EdgeOne and query the IP geolocation.

## Directions

1. Log in to the EdgeOne console and click **IP Location Query** in the left sidebar.

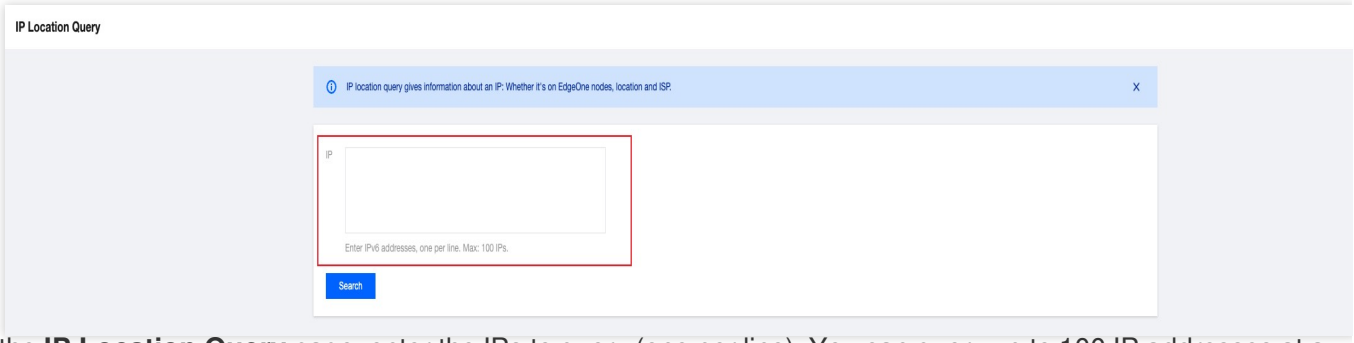


2. On the **IP Location Query** page, enter the IPs to query (one per line). You can query up to 100 IP addresses at a time. IPv6 addresses are supported.
3. Click **Search**. The **Query results** table shows the IP geolocation and whether they are owned by EdgeOne nodes. To export the query results, click the download icon

in the top-right corner of the table. The query results are exported to a CSV file.