

Oracle® Fusion Middleware

Configuration File Reference for Oracle Traffic Director



12c (12.2.1.4.0)

E96375-02

November 2019

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Oracle Fusion Middleware Configuration File Reference for Oracle Traffic Director, 12c (12.2.1.4.0)

E96375-02

Copyright © 2016, 2019, Oracle and/or its affiliates. All rights reserved.

Primary Author: N. T. Thrupthi Contributors: Nanda Kishore

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	x
Documentation Accessibility	x
Related Documents	x
Conventions	xi

1 Overview of Configuration Files and Directories

Configuration Files	1-1
The server.xml File	1-1
The obj.conf File	1-2
NZ Wallet Files	1-2
Default Paths	1-2
Directory Structure	1-2
bin	1-3
Legal	1-3
lib	1-3
admin-server	1-3
net-server-id	1-4
Dynamic Reconfiguration	1-4

2 Syntax and Use of server.xml

Overview of server.xml	2-1
Editing the server.xml File	2-1
Editing Element Values	2-1
Adding Elements	2-2
Validating server.xml	2-2
Understanding server.xml	2-2
Server Pools	2-2
Health check	2-3
High Availability	2-3
HTTP Protocol	2-3

Logging and Monitoring	2-3
Performance Tuning	2-3
SSL, TLS, and PKCS #11	2-4
Variables	2-4
Virtual Servers	2-4
Sample server.xml File	2-4

3 Elements in server.xml

List of Elements	3-1
access-log	3-2
access-log-buffer	3-3
cert	3-3
crl	3-4
dns	3-4
dns-cache	3-5
event	3-5
event-subscription	3-6
file-cache	3-7
failover-group	3-8
failover-instance	3-9
ftp-filter	3-9
health-check	3-10
http	3-13
http-listener	3-15
keep-alive	3-17
localization	3-18
log	3-18
maintenance	3-19
origin-server-pool	3-20
origin-server	3-22
property	3-24
proxy-cache	3-24
proxy-server	3-25
qos-limits	3-25
server	3-26
snmp	3-28
ssl	3-28
ssl3-tls-ciphers	3-30
ssl-session-cache	3-33
stats	3-33

status-listener	3-34
tcp-access-log	3-34
tcp-listener	3-35
tcp-proxy	3-36
tcp-thread-pool	3-37
thread-pool	3-38
time	3-39
variable	3-39
virtual-server	3-40
webapp-firewall-ruleset	3-41

4 Syntax and Use of obj.conf

Request-Handling Process Overview	4-1
Steps in the Request-Handling Process	4-1
Directives in obj.conf	4-2
Objects in obj.conf	4-3
Objects That Use the name Attribute	4-3
Objects That Use the ppath Attribute	4-3
Using the Client, If, Elseif, and Else Tags	4-3
Client	4-3
If, Elseif, and Else	4-5
Flow of Control in obj.conf	4-5
AuthTrans	4-5
NameTrans	4-6
How and When Oracle Traffic Director Processes Other Objects	4-6
PathCheck	4-6
ObjectType	4-6
Input	4-7
Output	4-7
Route	4-7
Service	4-8
AddLog	4-8
Error	4-8
Changes in Function Flow	4-8
Restarted Requests	4-8
Internal Requests	4-8
URI Translation	4-9
Editing obj.conf	4-9
Order of Directives	4-9
Parameters	4-9

Case Sensitivity	4-9
Separators	4-9
Quotation Marks	4-9
Spaces	4-9
Line Continuation	4-10
Path Names	4-10
Comments	4-10

5 Predefined Server Application Functions and Filters in obj.conf

The bucket Parameter	5-1
AuthTrans	5-2
get-sslid	5-2
qos-handler	5-2
webapp-firewall	5-3
NameTrans	5-4
assign-name	5-5
block-request-cookie	5-6
map	5-6
reverse-map	5-7
rewrite	5-8
sed-request-header	5-8
strip-params	5-9
PathCheck	5-9
check-request-limits	5-10
deny-existence	5-12
get-client-cert	5-12
nt-uri-clean	5-13
ssl-logout	5-14
unix-uri-clean	5-14
ObjectType	5-14
block-auth-cert	5-16
block-cache-info	5-16
block-cipher	5-16
block-ip	5-17
block-issuer-dn	5-17
block-jroute	5-17
block-keysize	5-17
block-proxy-agent	5-18
block-secret-keysize	5-18
block-ssl	5-18

block-ssl-id	5-19
block-user-dn	5-19
block-via	5-19
block-xforwarded-for	5-19
forward-auth-cert	5-20
forward-cache-info	5-20
forward-cipher	5-20
forward-ip	5-21
forward-issuer-dn	5-21
forward-jroute	5-22
forward-keysize	5-22
forward-proxy-agent	5-22
forward-secret-keysize	5-23
forward-ssl	5-23
forward-ssl-id	5-24
forward-user-dn	5-24
forward-via	5-24
forward-xforwarded-for	5-25
http-client-config	5-25
proxy-cache-config	5-27
proxy-cache-override-http	5-28
proxy-websocket-config	5-30
reverse-block-date	5-30
reverse-block-server	5-31
reverse-forward-date	5-31
reverse-forward-server	5-31
set-basic-auth	5-31
set-cache-control	5-32
set-cookie	5-32
type-by-exp	5-33
type-by-extension	5-34
Input	5-34
Output	5-35
Route	5-35
set-origin-server	5-36
set-proxy-server	5-37
Service	5-38
proxy-retrieve	5-39
remove-filter	5-40
service-proxy-cache-dump	5-41
service-trace	5-42

stats-xml	5-42
AddLog	5-43
flex-log	5-44
Error	5-44
qos-error	5-45
send-error	5-45
Common SAFs	5-46
insert-filter	5-47
match-browser	5-48
redirect	5-49
remove-filter	5-50
Example	5-51
restart	5-51
rewrite-cookie	5-52
sed-param-name	5-52
sed-param-value	5-53
sed-response-header	5-53
set-priority	5-53
set-variable	5-54

A Using Variables, Expressions, Wildcards, and String Interpolation

If, Elself, and Else Tags	A-1
Variables	A-2
Predefined Variables	A-2
Custom Variables	A-4
Resolving Variables	A-4
Expressions	A-5
Expression Syntax	A-5
Expression Results as Boolean Values	A-6
Expression Literals	A-6
String Literals	A-6
Numeric Literals	A-7
Expression Variables	A-7
Expression Operators	A-8
Expression Functions	A-10
atime	A-10
choose	A-11
ctime	A-11
escape	A-12
external	A-12

httpdate	A-13
lc	A-13
length	A-14
lookup	A-14
lookupregex	A-15
mtime	A-16
owner	A-16
uc	A-17
unescape	A-17
uuid	A-18
Regular Expressions	A-18
String Interpolation	A-19
Using Variables in Interpolated Strings	A-19
Using Expressions in Interpolated Strings	A-19
Wildcard Patterns	A-20

B Using the Custom Access-Log File Format

C Using Time Formats

D Alphabetical List of Server Configuration Elements and Predefined SAFs

Preface

This document describes the purpose and use of the configuration files for Oracle® Fusion Middleware, including `server.xml`, and `obj.conf`. It provides a comprehensive list of the elements and directives in these configuration files.

Audience

The intended audience for this document is the person who administers and maintains Oracle® Fusion Middleware.

This document assumes you are familiar with:

- Working in a terminal window
- HTTP
- XML
- Executing operating system commands on UNIX-like platforms

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents, which are available on the Oracle Technology Network:

- *Installing Oracle Traffic Director*
- *WebLogic Scripting Tool for Oracle Traffic Director*
- *Administering Oracle Traffic Director*
- *About WebLogic Server MT*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Overview of Configuration Files and Directories

The configuration and behavior of Oracle Traffic Director is determined by a set of configuration files. You can use the Oracle Fusion Middleware Control and the WebLogic Scripting Tool (WLST) to change the configuration file settings. You can also manually edit these files.

Note:

As of 12.2.1.4.0, Oracle Traffic Director is deprecated. In the future, use Oracle HTTP Server, Microsoft IIS Web Server, or Apache HTTP Server plug-ins, or a native Kubernetes load balancer, such as Traefik for equivalent functionality.

This chapter includes the following topics:

- [Configuration Files](#)
- [Directory Structure](#)
- [Dynamic Reconfiguration](#)

Configuration Files

Each server instance has its own directory, referred to as `INSTANCE_HOME` in this document. The `INSTANCE_HOME/config` directory contains configuration files for Oracle Traffic Director components. The exact number and names of the configuration files depend on the components that were enabled or loaded into the server. For the default location of the `INSTANCE_HOME`, see [Default Paths](#).

The following sections describe the configuration files and related information pertaining to Oracle Traffic Director:

- [The server.xml File](#)
- [The obj.conf File](#)
- [NZ Wallet Files](#)
- [Default Paths](#)

The server.xml File

The `server.xml` file contains the Oracle Traffic Director configuration. For more information about the `server.xml` file, see [Syntax and Use of server.xml](#).

The obj.conf File

The `obj.conf` file contains directives for HTTP request processing. For more information about the `obj.conf` file, see [Syntax and Use of obj.conf](#).

NZ Wallet Files

Wallet files are stored in the `INSTANCE_HOME/config` directory:

- `ewallet.p12` - an editable wallet, with PIN
- `cwallet.sso` - a 'single sign-on' wallet, with obfuscated password (no PIN needed)

There are three NZ wallet operational modes. All three modes are supported by the OTD core server, but the admin may not support all of them. The modes are:

- only `ewallet.p12` - a PIN is always needed, whether OTD accesses the wallet in read-only mode, or the admin accesses the wallet in read-write mode. OTD prompts for the PIN on startup. The PIN is not stored in `server.xml`.
- only `cwallet.sso` - called "auto-login only" mode in `orapki`. No PIN is required, whether OTD accesses the wallet in read-only mode, or the admin accesses the wallet in read-write mode.
- both `ewallet.p12` and `cwallet.sso` - called "auto-login" mode in `orapki`. OTD doesn't require a PIN to access the wallet in read-only mode, but the admin requires a PIN to access the wallet in read-write mode. Edits are made to `ewallet.p12`. After each edit, `ewallet.p12` is converted to `cwallet.sso`.

Use the Oracle `orapki` command line utility to manage public key infrastructure (PKI) elements such as wallets. For more information, see the appendix *Using the orapki Utility to Manage PKI Elements* in the *Oracle Database Security Guide*.

Default Paths

The default paths used in Oracle Traffic Director are listed below. [Table 1-1](#) describes the default paths and file names.

Table 1-1 Default Paths

Placeholder	Description
<code>ORACLE_HOME</code>	The directory in which Oracle Traffic Director is installed.
<code>INSTANCE_HOME</code>	The directory that contains the files pertaining to the Oracle Traffic Director administration server and Oracle Traffic Director instances.

Directory Structure

This section describes the directory structures that are created when you first install Oracle Traffic Director. In a typical OTD installation, all directories are located in the `ORACLE_HOME` directory, except for the `admin-server` directory and `net-server-id` instance subdirectory. The `admin-server` directory and the `net-server-id` subdirectory are located in the `INSTANCE_HOME` directory. Creating an administration server or an

administration node requires you to execute the `configure-server` command. To create an administration node in OTD, you must make sure that the administration server is up and running. The newly created administration node will be registered with the administration server. Before creating any instances, create an administration server by running the `configure-server` command. Later, execute the `create-instance` command to create instances in OTD. For more information about the CLI commands, see *Oracle Traffic Director Command-line Reference*. For more information about the default locations for these directories, see [Default Paths](#). The following directories exist under the Oracle Traffic Director installation directory, `ORACLE_HOME`:

- [bin](#)
- [Legal](#)
- [lib](#)
- [admin-server](#)
- [net-server-id](#)

bin

The `bin` directory contains the commands to administer Oracle Traffic Director, including the WebLogic Scripting Tool Command interface (`WLST`).

Legal

The `Legal` directory contains the third party information pertaining to the Oracle Traffic Director software.

lib

The `lib` directory contains internal binaries, scripts, libraries, and bundled plug-ins. These are private files, for internal use only.

For information about the `admin-server` directory and `net-server-id` subdirectory see:

- [admin-server](#)
- [net-server-id](#)

admin-server

The `admin-server` directory contains the following subdirectories:

- `bin` - contains binary files used to start, stop, and restart Oracle® Fusion Middleware. On UNIX systems, this directory also contains the file required for rotating logs.
- `config` - contains private configuration files for the administration server. These files are for internal use.
- `config-store` - contains files used by the administration server to track server configuration information.

 **Note:**

The files in this directory are created by Oracle Traffic Director for internal use. Do not edit, run scripts on, or otherwise access any files in the `config-store` directory.

- `generated` - contains files generated by the instance, such as Java class files corresponding to JavaServer Pages (JSP).
- `logs` - contains any error messages or access log files generated by a server instance.

net-server-id

A `net-server-id` directory is created for every instance you create in OTD. This directory has the following subdirectories and files:

- `bin` - contains the commands for starting, stopping, restarting, and reconfiguring the server. It also contains the command for rotating the log files.
- `config` - contains the following instance-specific configuration files:
 - `<Cl>-obj.conf` - virtual server specific directory.
 - `obj.conf` -Oracle® Fusion Middleware directives for handling HTTP requests from clients.
 - `pkcs11.txt` - NSS PKCS #11 module database.
 - `server.xml` - most of the server configuration settings.
- `logs` - contains log files generated by this server instance.

Dynamic Reconfiguration

Dynamic reconfiguration enables you to make configuration changes to a running Oracle Traffic Director instance, without having to stop and restart it in order for the changes to take effect.

Dynamic configuration happens:

- When you deploy a configuration through the Administration Console or CLI
- When you run the `reconfig` script in the server instance's `bin` directory

You can dynamically change the configuration settings in the `obj.conf` file without restarting the server. In addition, most settings in the `server.xml` file can be changed without restarting the server. If you must restart the server, a warning message appears in the server log when you deploy the configuration or run the `reconfig` command.

You cannot dynamically reconfigure the following `server.xml` configuration parameters:

- `user`
- `temp-path`

- log (with the exception of log-level)
- thread-pool
- pkcs11
- stats
- dns
- dns-cache
- ssl-session-cache
- access-log-buffer

When you run the `reconfig` command, a new configuration object is created, and all new incoming requests are processed based on this new configuration object. The current configuration object is removed when no HTTP requests are using the object.

In case an erroneous configuration occurs during dynamic reconfiguration, the server displays an error message. The server logs the error message to a log file specified by the last configuration that worked.

Certain erroneous configurations result in warning messages but do not cause the server to reject the configuration. Other erroneous configurations trigger error messages and cause the server to reject the configuration. If the server rejects a configuration during startup, the server does not start. If the server rejects a configuration during a dynamic reconfiguration, the server reverts to the last configuration that worked.

2

Syntax and Use of server.xml

The `server.xml` file contains most of the server configuration. This chapter describes the basic syntax of the `server.xml` file and provides a high-level view of the elements that are used to configure features of the server. This chapter contains the following topics:

- [Overview of server.xml](#)
- [Understanding server.xml](#)
- [Sample server.xml File](#)

Overview of server.xml

The `server.xml` file contains the elements that define the configuration. The `server.xml` file is located in the `INSTANCE_HOME/net-server-id/config` directory.

The file encoding is UTF-8 to maintain compatibility with UNIX text editors.

Editing the server.xml File

The structure of the `server.xml` file is a hierarchy, with `server` as the topmost element. The `server` element has many subelements, many of which have subelements of their own.

In general, you do not need to edit `server.xml` directly. Instead, use the Administrator Console and the `tadm` command-line interface to change values in the `server.xml` file. Changes made using the Administrator Console and `tadm` command-line interface affect the `server.xml` file. Using `tadm` when creating scripts to change the `server.xml` file ensures forward compatibility. If you edit the `server.xml` file directly, ensure that the resulting `server.xml` file is valid.

Editing Element Values

To change the values in the `server.xml` file, change the value between the tags associated with the element you are editing. For example, to change the value of `<log-level>` from `NOTIFICATION:1` to `TRACE:1`, find the `log` child element of the `server` element. In this example, you see the following lines:

```
<log>
  <log-file>../logs/server.log</log-file>
  <log-level>NOTIFICATION:1</log-level>
</log>
```

For example:

Change the `log-level` from `NOTIFICATION:1`

```
<log-level>NOTIFICATION:1</log-level>
```

```
to TRACE:1  
  
<log-level>TRACE:1</log-level>
```

After altering the `server.xml` file, you must deploy your configuration for most changes to take effect. Use the command-line interface command `tadm pull-config` to pull the modified `server.xml` file, then use the Administrator Console or the `tadm deploy-config` command to deploy your changes. Some changes require a server restart in order to take effect. For information about changes that require a restart and which do not, see [Dynamic Reconfiguration](#).

Adding Elements

To add a new element to the file, add the element and any required subelements. Elements begin with a tag, for example `<virtual-server>`, and end with the closing tag, for example `</virtual-server>`. The tags are case-sensitive.

Validating server.xml

After editing the `server.xml` file, Oracle Traffic Director automatically validates the XML code when you start or dynamically reconfigure a server.

You can also use the `-configtest` option of the `startserv` script to validate your configuration. From the instance's `bin` directory, run:

```
startserv -configtest
```

Understanding server.xml

To change the `server.xml` file for your environment, you must know which elements contain the relevant settings. The following sections contain brief descriptions of the elements that configure the functional areas:

- [Server Pools](#)
- [Health check](#)
- [High Availability](#)
- [HTTP Protocol](#)
- [Logging and Monitoring](#)
- [Performance Tuning](#)
- [SSL, TLS, and PKCS #11](#)
- [Variables](#)
- [Virtual Servers](#)

For more information about all the `server.xml` elements and their subelements, see [Elements in server.xml](#).

Server Pools

The `origin-server` element defines a member of a server pool. The `origin-server-pool` element configures a pool of origin servers that are used to load-balance

requests. An origin server is a back-end server—such as an Oracle WebLogic Server instance or an Oracle iPlanet Web Server instance—to which Oracle Traffic Director should forward requests that it receives from clients, and from which it receives responses. A set of origin servers providing the same service constitute an origin server pool. For more information, see [origin-server-pool](#), [origin-server](#), [server](#).

Health check

The `health-check` element configures the parameters that are used to determine the status of each origin server in an origin-server pool. The `health-check` element is a subelement of the `origin-server-pool` element. See [health-check](#) and [origin-server-pool](#).

High Availability

The `failover-group` element is a grouping of a VIP (Virtual IP), an instance that is designated as the primary server and another instance designated as the backup server. The Active-Passive or Active-Active cluster failover configurations are represented as Failover Groups. The `failover-group` element defines a failover group. See [failover-group](#).

HTTP Protocol

The `http` element configures the general HTTP protocol options. The `keep-alive` element configures the HTTP keep-alive connection management. The `http-listener` element configures the ports and IP addresses on which the server listens for new HTTP connections. The `virtual-server` element configures a method by which the server processes the HTTP requests. See [http](#), [keep-alive](#), [http-listener](#), and [virtual-server](#).

Logging and Monitoring

The `access-log` element configures the file name and format of access logs. The `access-log-buffer` element configures the frequency of access log updates and ordering of the access log entries. See [access-log](#) and [access-log-buffer](#). For more information about the log file format, see [Using the Custom Access-Log File Format](#).

The `log` element configures the file name and contents of the server log. The `event` element configures the access log and server log rotation. See [log](#) and [event](#).

The `snmp` element configures Simple Network Management Protocol (SNMP), and the `stats` element configures statistics collection. See [snmp](#) and [stats](#).

Performance Tuning

The `thread-pool` element configures the number of threads used to process requests and the maximum number of HTTP connections that the server queues. See [thread-pool](#).

The `keep-alive` element configures the HTTP keep-alive connection management. See [keep-alive](#). The `dns-cache` element configures the DNS caching. See [dns-cache](#).

SSL, TLS, and PKCS #11

The `ssl` element configures Secure Sockets Layer (SSL) and Transport Layer Security (TLS). SSL and TLS can be configured separately for each HTTP listener. For more information, see [ssl](#) and [http-listener](#).

The `pkcs11` element configures the PKCS #11 subsystem, including certificate revocation lists (CRLs) and third-party cryptographic modules.

Variables

The `variable` element defines a variable for use in expressions, log formats, and `obj.conf` parameters. See [variable](#). For more information about variable and expression use, see [Using Variables, Expressions, Wildcards, and String Interpolation](#).

Virtual Servers

The `virtual-server` element configures the virtual servers. Each virtual server processes HTTP requests from one or more HTTP listeners. The `http-listener` element configures the HTTP listeners. See [virtual-server](#) and [http-listener](#).

You can define variables within a virtual server using the `variable` element, as described in [Variables](#).

Sample server.xml File

Example 2-1 shows an excerpt from a `server.xml` file.

Example 2-1 server.xml file

```
<?xml version="1.0" encoding="UTF-8" ?>

<!--
  Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved
-->

<server>
  <cluster>
    <local-host>www.example.com</local-host>
    <instance>
      <host>www.example.com</host>
    </instance>
  </cluster>
  <log>
    <log-file>../logs/server.log</log-file>
    <log-level>NOTIFICATION:1</log-level>
  </log>
  <platform>64</platform>
  <temp-path>tmp/net-test-8a4af444</temp-path>
  <user>myuser</user>
  <access-log>
    <file>../logs/access.log</file>
  </access-log>
  <http-listener>
    <name>http-listener-1</name>
    <port>1894</port>
```

```
<server-name>www.example.com</server-name>
  <default-virtual-server-name>test</default-virtual-server-name>
</http-listener>
<virtual-server>
  <name>test</name>
  <http-listener-name>http-listener-1</http-listener-name>
  <host>www.example.com</host>
  <object-file>test-obj.conf</object-file>
</virtual-server>
<origin-server-pool>
  <name>origin-server-pool-1</name>
  <type>http</type>
<origin-server>
  <host>www.example.com</host>
  <port>20005</port>
</origin-server>
</origin-server-pool>
</server>
```

3

Elements in server.xml

This chapter describes the elements in the `server.xml` file in alphabetical order.

List of Elements

This section describes the elements in the `server.xml` file in alphabetical order.

- [access-log](#)
- [access-log-buffer](#)
- [cert](#)
- [crl](#)
- [dns](#)
- [dns-cache](#)
- [file-cache](#)
- [event](#)
- [failover-group](#)
- [failover-instance](#)
- [ftp-filter](#)
- [health-check](#)
- [http](#)
- [http-listener](#)
- [keep-alive](#)
- [localization](#)
- [log](#)
- [maintenance](#)
- [origin-server-pool](#)
- [origin-server](#)
- [property](#)
- [proxy-cache](#)
- [proxy-server](#)
- [qos-limits](#)
- [server](#)
- [snmp](#)
- [ssl](#)

- [ssl3-tls-ciphers](#)
- [ssl-session-cache](#)
- [stats](#)
- [tcp-access-log](#)
- [tcp-listener](#)
- [tcp-proxy](#)
- [tcp-thread-pool](#)
- [thread-pool](#)
- [time](#)
- [variable](#)
- [virtual-server](#)
- [webapp-firewall-ruleset](#)

access-log

The `access-log` element configures the settings for the access log. This element can appear zero or more times within the `server` element and zero or more times within the `virtual-server` element. See [server](#) and [virtual-server](#).

[Table 3-1](#) describes the subelements of `access-log`.

Table 3-1 `access-log` Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies whether the server writes to this access log. Default Value: <code>true</code> .
<code>name</code>	0 or 1	The name that uniquely identifies the access log. If you specify a name, the server does not automatically write to this access log. Instead, you explicitly configure this access log in an <code>obj.conf</code> <code>AddLog</code> directive.
<code>file</code>	1	The file name of the access log. If a relative path is used, it is relative to the server's <code>config</code> directory, for example, <code>../logs/access.log</code> .
<code>format</code>	0 or 1	The format of the access log entries. The default format is an extended custom log format. For more information about access log format, see Using the Custom Access-Log File Format .

Related Topics

- [access-log-buffer](#)
- [event](#)
- [log](#)

access-log-buffer

The `access-log-buffer` element configures the settings for access log buffering subsystem. This element can appear zero or one time within the `server` element. For more information, see [server](#).

[Table 3-2](#) describes the subelements of `access-log-buffer`.

Table 3-2 `access-log-buffer` Subelements

Element	Occurrences	Description
<code>direct-io</code>	0 or 1	Specifies if the file system cache access log writes. Default value: <code>false</code> . It indicates that the file system write to a cache. Setting the value to <code>true</code> indicates that the file system should not to write to a cache. The setting is purely advisory; either the server or the operating system may choose to ignore it.
<code>enabled</code>	0 or 1	Specifies whether the server buffers the access log entries. Default value: <code>true</code> .
<code>buffer-size</code>	0 or 1	The size (in bytes) of individual access log buffers. The value can be from 4096 to 1048576. Default value: 8192.
<code>max-buffers</code>	1	Specifies the maximum number of access-log buffers per server. Values: 1 to 65536. Default value: 1000.
<code>max-buffers-per-file</code>	0 or 1	Specifies the maximum number of access-log buffers per access-log file. Values: 1 to 128.
<code>max-age</code>	0 or 1	The maximum time (in seconds) to buffer a given access log entry. The value can be from 0.001 to 3600. Default value: 1.

Related Topics

- [access-log](#)

Related Topics

- [event](#)
- [log](#)

cert

The `cert` element uniquely identifies a certificate. This element can appear zero or more times within the `ssl` element. See [ssl](#).

[Table 3-3](#) describes the subelements of `cert`.

Table 3-3 cert Subelements

Element	Occurrences	Description
subject	1	Required. Subject name of the certificate.
serial-number	0 or 1	This optional field can be specified in order to disambiguate between multiple certificates with the same subject name. This field is specified in hexadecimal and is not case-sensitive. The 0x prefix is optional.
issue	0 or 1	This optional field can be specified in order to disambiguate between multiple certificates with the same subject name.

crl

The `crl` element uniquely identifies a certificate revocation list (CRL). This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-4](#) describes the subelements of `crl`.

Table 3-4 crl Subelements

Element	Occurrences	Description
enabled	0 or 1	Defines whether CRLs should be applied while verifying SSL certificates. Default value: <code>true</code> .
crl-path	0 or 1	Specifies the path to a directory where CRLs are stored. Defaults to 'crl' sub-directory under the config directory of the instance.
crl-cache-size	0 or 1	Defines a cache size between 0 and 2,147,483,647 bytes inclusive. Default value: 52428800.
require-crls	0 or 1	Specifies whether a CRL is required when verifying peer certificates during SSL/TLS handshakes. This affects both <code>libproxy</code> (back-end) and client (front-end) authentication. Default value: <code>false</code> .

dns

The `dns` element configures how the server uses the domain name system (DNS). This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-5](#) describes the subelements of `dns`.

Table 3-5 dns Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the server does DNS lookups. Default value: <code>false</code> .

Table 3-5 (Cont.) dns Subelements

Element	Occurrences	Description
<code>async</code>	0 or 1	Specifies whether the server uses its own asynchronous DNS resolver, instead of the Operating System's synchronous resolver. Default value: <code>true</code> .
<code>timeout</code>	0 or 1	Specifies the duration (in seconds) after which the asynchronous DNS lookups should time out. The value can be from 0.001 to 3600.

Related Topics

- [dns-cache](#)

dns-cache

The `dns-cache` element configures the DNS cache. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-6](#) describes the subelements of `dns-cache`.

Table 3-6 dns-cache Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies whether the server writes to a cache for DNS lookup results. Default value: <code>true</code> .
<code>max-age</code>	0 or 1	Specifies the duration (in seconds) for which the entries must be kept in the cache. The value can be from 0.001 to 604800. Default value: 120.
<code>max-entries</code>	0 or 1	Specifies the maximum number of DNS lookup results to write to the cache. The value can be from 1 to 1048576. Default value: 1024.

Related Topics

- [dns](#)

event

The `event` element configures a recurring event. The element can appear zero or more times within the `server` element. See [server](#).

[Table 3-7](#) describes the subelements of `event`.

Table 3-7 event Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies whether the event is enabled at run time. Default value: <code>true</code> .

Table 3-7 (Cont.) event Subelements

Element	Occurrences	Description
time	0 or more	Configures a specific time when the event occurs. See time .
interval	0 or 1	Specifies the interval (in seconds) at which the event occurs. The value can be from 60 to 86400.
rotate-log	0 or 1	Rotates the log files. Default value: <code>false</code> .
rotate-access-log	0 or 1	Rotates the access log files. Default value: <code>false</code> .
command	0 or more	The command to execute to get an event to run.
reconfig	0 or 1	Dynamically reconfigures the server. Default value: <code>false</code> .
restart	0 or 1	Restarts the server. Default value: <code>false</code> .
description	0 or 1	The description of the event. The value of this element is in text format.

Related Topics

- [access-log](#)

Related Topics

- [access-log-buffer](#)

Related Topics

- [log](#)

event-subscription

The `event-subscription` element configures parameters for receiving notifications on HTTP endpoint URLs when an event occurs. The element may appear zero or once within the element. See [origin-server](#).

[Table 3-8](#) describes the subelements of `event`.

Table 3-8 event-subscription Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies the configuration in the event-subscription block is enabled or disabled. Default value: <code>true</code> (enabled).
name	1	Specifies the user defined name of the event subscription. Values: string.
URL	1	Specifies the subscription URL. If this is configured, Oracle® Fusion Middleware publishes the notification to this URL. Value: a valid HTTP URL.

 **Note:**

HTTPS endpoints are not supported.

file-cache

The `file-cache` element configures the file cache. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-9](#) describes the subelements of `file-cache`.

Table 3-9 `file-cache` Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Indicates whether the server caches file content and meta information. Whether file content is cached in addition to meta information is controlled by the <code>cache-content</code> sub-element. Default value: <code>true</code> .
<code>max-age</code>	0 or 1	The maximum amount of time (in seconds) to cache file content and meta information. The value can be from 0.001 to 3600. -1 indicates no limit. Default value: 30.
<code>max-entries</code>	0 or 1	The maximum number of paths to cache content and/or meta information. The value can be from 1 to 1073741824. Default value: 1024
<code>max-open-files</code>	0 or 1	The maximum number of file descriptors the file cache will keep open. The value can be from 1 to 1073741824.
<code>sendfile</code>	0 or 1	Determines whether the server will attempt to use the operating system's <code>sendfile</code> , <code>sendfilev</code> , <code>send_file</code> , or <code>TransmitFile</code> system call. The default value is <code>true</code> on Windows and <code>false</code> on other platforms.
<code>copy-files</code>	0 or 1	Determines whether the server copies cached files to a temporary directory. Default value: <code>true</code> on Windows, <code>false</code> on other platforms.
<code>copy-path</code>	0 or 1	The temporary directory that is used when <code>copy-files</code> is <code>true</code> . If a relative path is used, it is relative to the server's config directory.
<code>replacement</code>	0 or 1	The cache entry replacement algorithm. The value can be <code>false</code> , <code>lru</code> , or <code>lfu</code> . Default value: <code>lru</code> .
<code>cache-content</code>	0 or 1	Determines whether the server caches file content in addition to the meta information. The default value is <code>true</code> .
<code>max-heap-file-size</code>	0 or 1	The maximum size (in bytes) of files to cache on the heap. The value can be from 0 to 2147483647. -1 indicates that there is no maximum size. Default value: 524288.
<code>max-heap-space</code>	0 or 1	The maximum amount (in bytes) of heap to use for caching files. The value can be from 0 to 1099511627776. Default value: 10485760.

Table 3-9 (Cont.) file-cache Subelements

Element	Occurrences	Description
max-mmap-file-size	0 or 1	The maximum size (in bytes) of files to mmap. The value can be from 0 to 2147483647. -1 indicates that there is no maximum size. Default value: 0.
max-mmap-space	0 or 1	The maximum amount (in bytes) of mmap address space to use for caching files. The value can be from 0 to 1099511627776. Default value: 0.
buffer-size	0 or 1	Specifies the size of the input/output buffer used on cache misses. The value can be from 512 to 1048576. Default value: 8192.
sendfile-size	0 or 1	sendfile-size is used only when sendfile is enabled and the file size is greater than max-heap-file-size. When sendfile-size is set to its default value of 0, the entire file is attempted to be sent out at once. Otherwise, the file-cache subsystem attempts to send the file in chunks of at most sendfile-size. The values can range from 0 and 2147483647. Default value: 0.

failover-group

The `failover-group` element defines a failover group. This element may appear zero or one time within the `cluster` element.

[Table 3-10](#) describes the subelements of `failover-group`.

Table 3-10 failover-group Subelements

Element	Occurrences	Description
virtual-ip	1	Specifies the virtual IP for the failover group. The virtual IP should belong to the same subnet as that of the instances in the failover group, and must be accessible to clients.
router-id	0 or 1	Specifies the router identity for the failover-group. The value must be unique across the failover-groups. It is used to identify the router group of all the participating routers for the same VIP. Values are positive integer. Range of values: 1 to 255. Default value: 255.
failover-instance	1 or more	Represents an Oracle Traffic Director instance that belongs to the failover group. See failover-instance .
failover-type	0 or 1	Specifies the failover type. Values: active-active or active-passive. Default value: active-passive.
property	0 or more	On Linux, these specify custom properties that will be passed to the <code>vrp_instance</code> within <code>keepalived.conf</code> . See property .

failover-instance

The `failover-instance` element defines a failover instance.

[Table 3-11](#) describes the subelements of `failover-group`.

Table 3-11 `failover-instance` Subelements

Element	Occurrences	Description
<code>instance-name</code>	1	Specifies the name of the Oracle Traffic Director instance which is part of the failover group.
<code>host</code>	1	Specifies the hostname of the node where the instance has been created.
<code>priority</code>	0 or 1	Specifies the priority value for the instance. This value identifies whether the instance is the primary or the backup for the <code>failover-group</code> . Values: positive integer. Range of values: 1 to 254. Default value: 250.
<code>network-interface</code>	1	Indicates the network interface on the node where this instance is created on which the VIP is moderated.

ftp-filter

The `ftp-filter` element configures parameters that are used to make a TCP proxy to front-end an FTP server. See [tcp-proxy](#).

[Table 3-12](#) describes the subelements of `tcp-proxy`.

Table 3-12 `ftp-filter` Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies if the FTP filter is enabled or disabled. Default value: <code>false</code>
<code>client-explicit-ftp</code>	0 or 1	Specifies if client-side SSL should be enabled explicitly. Default value: <code>true</code>
<code>origin-explicit-ftp</code>	0 or 1	Specifies if server-side SSL should be enabled explicitly. Default value: <code>true</code>
<code>ssl-termination</code>	0 or 1	Specifies if SSL should terminate at Oracle Traffic Director. Default value: <code>false</code>
<code>passive-port-min</code>	0 or 1	Specifies the lower limit of port range for FTP passive data connections. Range of values: 1025 to 65535. Default value: 1025

Table 3-12 (Cont.) ftp-filter Subelements

Element	Occurrences	Description
passive-port-max	0 or 1	Specifies the upper limit of port range for FTP passive data connections. Range of values: 1025 to 65535. Default value: 65535
active-port-min	0 or 1	Specifies the lower limit of port range for FTP active data connections. Range of values: 1025 to 65535. Default value: 1025
active-port-max	0 or 1	Specifies the upper limit of port range for FTP active data connections. Range of values: 1025 to 65535. Default value: 65535

Related Topics

- [tcp-access-log](#)
- [tcp-listener](#)
- [tcp-thread-pool](#)

health-check

The `health-check` element configures the parameters that are used to determine the status of each origin-server in an origin-server pool. This element may appear zero or one time within the `origin-server-pool` element. See [origin-server-pool](#).

[Table 3-13](#) describes the subelements of `health-check`.

Table 3-13 health-check Subelements

Elements	Occurrences	Description	TCP health check on HTTP servers	TCP health check on TCP servers
protocol	0 or 1	<p>Specifies the type of connection—HTTP or TCP, or an external executable—that Oracle Traffic Director should attempt with the origin server to determine its health.</p> <p>TCP: Oracle Traffic Director attempts to open a TCP connection to each origin server. The success or failure of this attempt determines whether Oracle Traffic Director considers the origin server to be online or offline.</p> <p>HTTP: Oracle Traffic Director sends an HTTP GET or OPTIONS request to each origin server in the pool, and checks the response to determine the availability and health of the origin server.</p> <p>COMMAND: Oracle Traffic Director invokes the executable specified in <command> for the health check.</p> <p>Default value: HTTP.</p>	Valid	Valid; HTTP is not a valid value for origin-server-pool elements that specify tcp in the type subelement.
interval	0 or 1	<p>Specifies the time interval (in seconds) between successive health check operations. Values: 0.001 to 3600. Default value: 30.</p>	Valid	Valid
failover-threshold	0 or 1	<p>Indicates the number of consecutive failures for marking a server down. Values: 1 to 256. Default value: 3.</p>	Valid	Valid

Table 3-13 (Cont.) health-check Subelements

Elements	Occurrences	Description	TCP health check on HTTP servers	TCP health check on TCP servers
timeout	0 or 1	Specifies the timeout value (in seconds) for a connection. Values: 0.001 to 3600. Default value: 5.	Valid	Valid
command	0 or 1	Specifies the full path of an external health check executable. You must configure this parameter if the protocol is COMMAND.	Ignored	Ignored
request-method	0 or 1	Specifies the method used during HTTP health check operations. Values: GET, HEAD and OPTIONS. Default value: OPTIONS.	Ignored	Ignored
request-uri	0 or 1	Specifies the URI that is used for HTTP health check operations. Default value: "/".	Ignored	Ignored
response-code-match	0 or 1	A modified regular expression used to specify the types of response status codes acceptable for a healthy origin server. The expression is a union of three character patterns that contain only digits or 'x'. 'x' represents any digit, for example, the following three expressions are valid: 200, 2xx 304, 1xx 2xx 3xx 4xx. If the parameter is not specified, all other codes except 5xx server error are considered acceptable. This is applicable only when protocol is HTTP.	Ignored	Ignored

Table 3-13 (Cont.) health-check Subelements

Elements	Occurrences	Description	TCP health check on HTTP servers	TCP health check on TCP servers
response-body-match	0 or 1	A regular expression that is used to match the HTTP response body to determine the origin server's health. This is applicable only when protocol is HTTP.	Ignored	Ignored
response-body-match-size	0 or 1	Specifies the maximum length of the response body that should match. Values: 0 to 2147483647. Default value: 2048.	Ignored	Ignored
dynamic-server-discovery	0 or 1	Specifies whether the server should dynamically discover Oracle WebLogic Server cluster nodes and add them to the pool. Default value: false.	Valid for HTTP Health Check	Ignored
use-object-file	0 or 1	Specifies whether the obj.conf file processing for health-check requests is enabled. Default value: True	Ignored	Ignored

http

The `http` element configures the settings for the miscellaneous HTTP protocol options. This element can appear zero or one time within the `server` element. For more information, see [server](#).

[Table 3-14](#) describes the subelements of `http`.

Table 3-14 http Subelements

Element	Occurrences	Description
version	0 or 1	Specifies the highest HTTP protocol version the server supports. The default HTTP version string is HTTP/1.1.
server-header	0 or 1	Specifies the server header information such as server software and version. The default server header is Oracle-Traffic-Director/<version>. Note that setting this to an empty string suppresses the server header.

Table 3-14 (Cont.) http Subelements

Element	Occurrences	Description
etag	0 or 1	Indicates if the server includes an Etag header field in its responses. Default value: true
request-header-buffer-size	0 or 1	Specifies the size (in bytes) of the buffer used to read HTTP request headers. The value can be from 0 to 2147483647. Default value: 8192.
strict-request-headers	0 or 1	Indicates whether the server rejects certain malformed HTTP request headers. Default value: false.
case-sensitive-headers	0 or 1	Indicates whether the server should preserve the case of HTTP request/response headers. Default value: false.
websocket-strict-upgrade	0 or 1	Enables/disables strict RFC 6455 adherence during the WebSocket upgrade request. Default value: false.
discard-misquoted-cookies	0 or 1	Indicate whether to discard misquoted cookies. Default value: true.
max-request-headers	0 or 1	Specifies the maximum number of header fields in an HTTP request header. The value can be from 1 to 512. Default value: 64.
body-buffer-size	0 or 1	Defines the maximum size of the request body content that OTD will expose via the \$body variable in obj.conf. The value can range from 0 to 2147483647. Default value: 1024.
output-buffer-size	0 or 1	Specifies the size (in bytes) of the buffer for HTTP responses. The value can be from 0 to 2147483647. Default value: 8192.
max-unchunk-size	0 or 1	Specifies the maximum size (in bytes) of a chunked HTTP request body that the server will unchunk. The value can be from 0 to 2147483647. Default value: 8192.
unchunk-timeout	0 or 1	Specifies the maximum time (in seconds) that the server waits for a chunked HTTP request body to arrive. The value can be from 0 to 3600, or -1 for no timeout. Default value: 60.
io-timeout	0 or 1	Specifies the maximum time (in seconds) that the server waits for an individual packet. The value can be from 0 to 3600. Default value: 30.
request-header-timeout	0 or 1	Specifies the maximum time (in seconds) that the server waits for a complete HTTP request header. The value can be from 0 to 3600. Default value: 30.
request-body-timeout	0 or 1	Specifies the maximum time (in seconds) that the server waits for a complete HTTP request body. The value can be from 0 to 604800, or -1 for no timeout. Default value: -1.

Table 3-14 (Cont.) http Subelements

Element	Occurrences	Description
favicon	0 or 1	Specifies whether the server replies to requests for <code>favicon.ico</code> with its own built-in icon file. Default value: <code>true</code> .
ecid	0 or 1	Specifies whether the server generates, propagates, and logs the execution context. The value of the ECID is a unique identifier that can be used to correlate individual events as being part of the same request execution flow. For example, events that are identified as being related to a particular request typically have the same ECID value. However, the format of the ECID string itself is determined by an internal mechanism that is subject to change; therefore, you should not have or place any dependencies on that format. ECID is defined as a part of the execution context. The execution context consists of ECID and RID. You may also refer to the whole execution context, which is the combination of ECID and RID, as just ECID. Default value: <code>true</code> .

Related Topics

- [http-listener](#)
- [keep-alive](#)
- [thread-pool](#)
- [virtual-server](#)

Related Topics

- [thread-pool](#)

http-listener

The `http-listener` element configures an HTTP listener. This element can appear zero or more times within the `server` element. See [server](#).

[Table 3-15](#) describes the subelements of `http-listener`.

Table 3-15 http-listener Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the HTTP listener is enabled to accept connection requests. Default value: <code>true</code> .
name	1	Specifies the name that uniquely identifies the HTTP listener. Whitespace is not permitted.

Table 3-15 (Cont.) http-listener Subelements

Element	Occurrences	Description
ip	0 or 1	Specifies an IP address to which to listen. The value of this element is a specific IP address or hostname or an asterisk * to listen on all IP addresses.
port	1	Specifies the port to which to listen. The value of this element is the port number. The value can be from 1 to 65535.
acceptor-threads	0 or 1	Specifies the number of threads dedicated to accept connections received by this listener. The value can be from 1 to 128.
server-name	1	Specifies the default server name. Tells the server what to put in the host name section of any URLs it sends to the client. This affects URLs the server automatically generates, say on a redirect; it does not affect the URLs for directories and files stored in the server. This name should be the alias name if your server uses an alias. Values: The value can be a hostname, fully qualified domain name, ip address, or a url prefix that contains one. The url prefix must not specify a path. It can include a scheme (for example, prefix http://) and port suffix (for example, :80)
blocking-io	0 or 1	Specifies whether the server uses blocking I/O. Default value: <code>false</code> .
family	0 or 1	Specifies the socket family that is used to connect to the origin server. Values: <code>inet</code> , <code>inet6</code> , <code>inet-sdp</code> , and <code>default</code> . <code>inet</code> and <code>inet6</code> represent IPV4 and IPV6 protocols respectively. <code>inet-sdp</code> is used for Sockets Direct Protocol (SDP). Default value: <code>default</code> .
handle-protocol-mismatch	0 or 1	Indicates whether the server responds to SSL or non-SSL protocol mismatches in client requests. Default value: <code>true</code> , meaning the server will attempt to detect SSL or non-SSL protocol mismatches and send an HTTP redirect or SSL alert when a mismatch is detected. Default value: <code>true</code> .
listen-queue-size	0 or 1	Specifies the size (in bytes) of the listen queue. The value of this element can be from 1 to 1048576. Default value: 128.
receive-buffer-size	0 or 1	Specifies the size (in bytes) of the operating system socket receive buffer. The value of this element can be from 0 to 2147483647.
send-buffer-size	0 or 1	Specifies the size (in bytes) of the operating system socket send buffer. The value of this element can be from 0 to 2147483647.

Table 3-15 (Cont.) http-listener Subelements

Element	Occurrences	Description
default-virtual-server-name	1	Specifies the name of the virtual server that processes request that do not match a host. The value of this element is the name value from a virtual-server element. See virtual-server .
ssl	0 or 1	Configures SSL/TLS. If the ssl element is omitted, SSL/TLS support is disabled. See ssl
description	0 or 1	Specifies the description of the HTTP listener. The value of this element must be in text format.
max-requests-per-connection	0 or 1	Specifies the maximum number of keep-alive requests that will be handled per HTTP connection after which the keep-alive connection will be closed. The value can be from 1 to 2147483647, or -1 for no limit. Default value: -1
blocking-accept	0 or 1	Enables/disables blocking of the server listen socket while retaining client end points as non blocking (useful when MaxProcs > 1). Default value: false.

Related Topics

- [http](#)

Related Topics

- [keep-alive](#)
- [virtual-server](#)

keep-alive

The `keep-alive` element configures the settings for the keep-alive subsystem. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-16](#) describes the subelements of `keep-alive`.

Table 3-16 keep-alive Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the keep-alive subsystem is enabled at runtime. Default value: true.
threads	0 or 1	Specifies the number of keep alive subsystem threads. The value can be from 1 to 256.
max-connections	0 or 1	Specifies the maximum number of concurrent keep alive connections that the server supports. The value can be from 1 to 1048576.
timeout	0 or 1	Specifies the timeout (in seconds) after which an inactive keep alive connection can be used. The value can be from 0.001 to 3600. -1 indicates no timeout. Default value: 30 seconds.

Table 3-16 (Cont.) keep-alive Subelements

Element	Occurrences	Description
poll-interval	0 or 1	Specifies the interval (in seconds) between polls. The value can be from 0.001 to 1. Default value: .001.

Related Topics

- [http](#)

Related Topics

- [http-listener](#)

Related Topics

- [virtual-server](#)
- [thread-pool](#)

Related Topics

- [thread-pool](#)

localization

The `localization` element defines a method by which the server chooses a language with which it presents information to the client. This element may appear zero or one time within the `server` element, and zero or one time within the `virtual-server` element. See [server](#) and [virtual-server](#).

[Table 3-17](#) describes the subelement of `localization`.

Table 3-17 localization Subelements

Element	Occurrences	Description
default-language	0 or 1	The default language with which the messages and content are displayed. The value is a language tag.
negotiate-client-language	0 or 1	Specifies whether the server uses the <code>accept-language</code> HTTP header to negotiate the content language with clients. Default value: <code>false</code> .

log

The `log` element configures the logging subsystem. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-18](#) describes the subelements of `log`.

Table 3-18 log Subelements

Element	Occurrences	Description
log-stdout	0 or 1	Specifies whether the server logs data that applications write to <code>stdout</code> . Default value: <code>true</code> .
log-stderr	0 or 1	Specifies whether the server logs data that applications write to <code>stderr</code> . Default value: <code>true</code> .
log-virtual-server-name	0 or 1	Specifies whether the server includes the virtual server name in log messages. Default value: <code>false</code> .
log-to-syslog	0 or 1	Specifies whether the server writes log messages to <code>syslog</code> . Default value: <code>false</code> .
archive-command	0 or 1	This is executed after the server rotates a log file. The program is passed the post-rotation file name of the log file as an argument. A program command line, for example: <code>gzip</code>
log-level	0 or 1	Specifies the log verbosity for the server as a whole. Values: <code>INCIDENT_ERROR:1</code> , <code>NOTIFICATION:1</code> , <code>ERROR:1</code> , <code>ERROR:16</code> , <code>ERROR:32</code> , <code>WARNING:1</code> , <code>TRACE:1</code> , <code>TRACE:16</code> , <code>TRACE:32</code> . Default value: <code>NOTIFICATION:1</code>
log-file	0 or 1	Specifies the name and location of the log file. Value: User defined name and location. Default value: <code>../logs/server.log</code>

Related Topics

- [access-log](#)

Related Topics

- [access-log-buffer](#)

Related Topics

- [event](#)

maintenance

The `maintenance` element enables maintenance for an `origin-server-pool`.

[Table 3-19](#) describes the subelements of `maintenance`.

Table 3-19 maintenance Subelements

Element	Occurrences	Description
enabled	0 or 1	<p>Indicates whether the origin server pool has been put under maintenance. Values: true, false.</p> <p>true means that maintenance is enabled and the server pool will not handle any sticky or non-sticky requests. Instead a response based on the response-code and response-file configurations.</p> <p>false means that maintenance is disabled and origin servers of this pool (depending on the state of the origin server) can handle sticky and non-sticky requests.</p>
response-code	0 or 1	<p>Specifies the response code to be sent for requests that land on this pool when it is under maintenance.</p> <p>If this is set to 200, response-file is expected to be configured. If response-file is not specified, it will be considered as a misconfiguration and server will throw an error at start-up.</p> <p>Values: 200 or HTTP response codes in the range 400 to 599.</p>
response-file	0 or 1	<p>Specifies the absolute path of an HTML file to send to the client when the request lands on a maintenance enabled origin server pool. The file is sent as text/html regardless of its name or actual type. If the file does not exist or is not accessible, the server returns the default 503 response code.</p>

origin-server-pool

The `origin-server-pool` element configures a pool of origin servers that are used for load balancing requests. This element may appear zero or more times within the `server` element. See [server](#).

[Table 3-20](#) describes the subelements of `origin-server-pool`.

Table 3-20 origin-server-pool Subelements

Element	Occurrences	Description
name	1	Specifies the name by which the server pool is identified.

Table 3-20 (Cont.) origin-server-pool Subelements

Element	Occurrences	Description
load-distribution	0 or 1	<p>The load-balancing method for distributing requests to the origin-server pool. Values: round-robin, least-connection-count, least-response-time, and ip-hash. Default value: round-robin.</p> <p>For more information about load-balancing methods, see the section <i>Modifying an Origin-Server Pool</i> in the <i>Oracle Traffic Director Administrator's Guide</i>.</p>
ssl	0 or 1	<p>Specifies outgoing proxy SSL connections.</p> <p>This allows you to select the client certificate as well as ciphers. With Oracle Traffic Director 12.2.1, it replaces the ssl-client-config SAF. See ssl,</p>
type	1	Indicates the kind of requests that are handled by every server in the server pool. Values: http and tcp.
family	0 or 1	Specifies the socket family that is used to connect to the origin server. Values: inet, inet6, inet-sdp, and default. inet and inet6 represent IPV4 and IPV6 protocols respectively. inet-sdp is used for Sockets Direct Protocol (SDP). Default value: default.
origin-server	0 or more	Represents an origin server that belongs to the server pool. See origin-server
health-check	0 or 1	Specifies the health check settings for the server pool. See health-check
proxy-server	0 or 1	Specifies a proxy server through which connections to origin servers can optionally go through. See proxy-server .
queue-timeout	0 or 1	Time (in seconds) for which the request waits in the queue for a connection to an origin-server. If a timeout occurs, the request is rejected. Range of values: 0.001 to 3600. Default value: 30.

Table 3-20 (Cont.) origin-server-pool Subelements

Element	Occurrences	Description
tls-sni-extension	0 or 1	This element is applicable only if ssl is enabled. If present, it represents the value of the TLS Server Name Indication extension to be sent for TLS connections to the origin-server during the TLS handshake. If omitted, the TLS SNI extension is not sent. Note that for HTTP origin server pools, the administrator needs to ensure that the Host header in HTTP requests matches the value of this extension, otherwise, there may be an error on the origin server.
maintenance	0 or 1	This element is used to enable maintenance on an origin server pool. See maintenance .

origin-server

The `origin-server` element defines a member of a server pool. This element may appear zero or more times within the `origin-server-pool` element. For more information, see [origin-server-pool](#).

[Table 3-21](#) describes the subelements of `origin-server`.

Table 3-21 origin-server Subelements

Element	Occurrences	Description
host	1	Specifies the host name or the IP address of the origin server.
port	0 or 1	Specifies the port number of the origin server. Value: 1 to 65536. The port defaults to 80 if ssl is disabled and to 443 when ssl is enabled.
weight	0 or 1	Specifies the load distribution weight for the origin server. Value: 1 to 1000. Default value: 1.

Table 3-21 (Cont.) origin-server Subelements

Element	Occurrences	Description
mode	0 or 1	<p>Specifies the state of the origin-server. Values: enabled, disabled or draining.</p> <p>enabled means origin server is active and handles both sticky and non-sticky requests.</p> <p>disabled means origin server does not handle any (sticky or non-sticky) requests and is not part of the pool (e.g. when it is under maintenance).</p> <p>draining means origin server handles only sticky requests. New non-sticky requests are handled by other origin servers in the pool. This mode is used mainly to bring down the server gracefully for maintenance (like patch rollout). This mode is not applicable for origin servers of tcp server pool.</p>
backup	0 or 1	<p>Specifies whether the origin server is a backup server. Requests are sent to the backup origin server only when none of the primary (non-backup) origin servers is available. Default value: false.</p>
max-connections	0 or 1	<p>Specifies the maximum number of active connections an origin-server can have at any given point of time. Values: 1 to 1048576. -1 indicates that there is no maximum. Default value: -1.</p>
ramp-up-time	0 or 1	<p>The time (in seconds) that Oracle Traffic Director should take to ramp up the request sending rate to the full capacity of this origin server. Values: 0.001 to 3600. Default value: 0.001.</p>
max-requests-per-connection	0 or 1	<p>Specifies the maximum number of keep-alive requests that will be handled per origin server connection after which the keep-alive connection will be closed. The value can be from 1 to 2147483647, or -1 for no limit. Default value: -1.</p>
max-request-bps	0 or 1	<p>Total bandwidth limit in byte/second enforced on request. Values: 0 to 1099511627776, or 0 for no limit. Default value: 0.</p>
max-response-bps	0 or 1	<p>Total bandwidth limit in byte/second enforced on response. Values: 0 to 1099511627776, or 0 for no limit. Default value: 0.</p>

Table 3-21 (Cont.) origin-server Subelements

Element	Occurrences	Description
bandwidth-queue-timeout	0 or 1	Time in seconds before a request waiting in the queue for bandwidth is aborted. Values: 0 to 86400. Default value: 60.

property

The `property` element defines a `name-value` pair. The `property` element can appear zero or more times within `failover-group`. On Linux, these properties can be used to specify custom properties to be passed to the `vrp_instance` within `keepalived.conf`.

[Table 3-22](#) describes the subelements of `property`.

Table 3-22 property Subelements

Element	Occurrences	Description
name	1	The name of the property.
value	1	The value of the property.

Related Topics

- [variable](#)

proxy-cache

The `proxy-cache` element configures the HTTP reverse proxy cache configuration. This element can appear zero or one time within the `server` element. For more information, see [server](#).

[Table 3-23](#) describes the subelements of `proxy-cache`.

Table 3-23 proxy-cache Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether response caching is enabled. Default value: <code>true</code> .
max-heap-space	0 or 1	Specifies the maximum amount (in bytes) of heap to use for caching response objects. It should not be more than available memory or process address space. Values: 0 to 1099511627776 (1024 GB). Default value: 10485760 (10 MB).
max-heap-object-size	0 to 1	Specifies the maximum size of objects that should be cached. Objects larger than the specified size are not cached. Values: 0 to 2147483647. Default value: 524288 (512 KB). -1 indicates that there is no maximum size.

Table 3-23 (Cont.) proxy-cache Subelements

Element	Occurrences	Description
replacement	0 to 1	Specifies the algorithm for cache replacement. Values: lru, lfu, and false. Default value: lru. <ul style="list-style-type: none"> lru (Least Recently Used): Oracle Traffic Director discards the least recently used entry first. lfu (Least Frequently Used): Oracle Traffic Director discards the least frequently used entry first. false: Cache replacement is disabled.
max-entries	0 to 1	Specifies the maximum number of entries in the cache. The range is 1 to 1073741824. Default value: 1024.

proxy-server

[Table 3-24](#) describes the subelements of `proxy-server`.

Table 3-24 proxy-server Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the proxy server is enabled at runtime. Default value: true.
host	1	Specifies the host name or the IP address of the proxy server.
port	0 or 1	Specifies the port number of the proxy-server. Default: 80.

qos-limits

The `qos-limits` element configures the Quality of Service (QoS) limits. This element may appear zero or one time within the `server` element and zero or one time within the `virtual-server` element. See [server](#) and [virtual-server](#).

[Table 3-25](#) describes the subelements of `qos-limits`.

Table 3-25 qos-limits Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the QoS limits are enforced at runtime. Default value: true.
max-bps	0 or 1	Specifies the maximum transfer rate (bytes/second). Range of value: 1 to 2147483647. -1 indicates that there is no maximum. Default : -1.

Table 3-25 (Cont.) qos-limits Subelements

Element	Occurrences	Description
max-connections	0 or 1	Specifies the maximum number of concurrent connections. Range of value: 1 to 1048576. -1 indicates that there is no maximum. Default : -1.

server

The `server` element defines a server. This is the root element. There can be only one `server` element in the `server.xml` file.

[Table 3-26](#) describes the subelements of `server`.

Table 3-26 server Subelements

Element	Occurrences	Description
cluster	0 or 1	The server cluster to which the server belongs.
fips	0 or 1	Enables the FIPS-140 mode of operation for the security library.
crl	0 or 1	Defines a certificate. See crl .
log	0 or 1	Configures the logging subsystem. See log .
max-fd	0 or 1	Specifies a configurable upper limit on the file descriptor usage of the Oracle Traffic Director server process. The default value of <code>max-fd</code> element is 2 million. This means that by default, Oracle Traffic Director does not assume more than 2 million available file descriptors even if the actual file descriptor availability is configured to be higher.
mime-file	0 or 1	Specifies the path to a file that contains the MIME type mappings for the server. If a relative path is used, it is relative to the server's config directory.
user	0 or 1	The account the server runs as (UNIX only). The value is the user account. If the server is started as <code>root</code> , any UNIX account can be specified. If the server is started by a <code>non-root</code> account, only that <code>non-root</code> account can be specified.
temp-path	0 or 1	The directory where the server stores its temporary files. If a relative path is used, it is relative to the server's <code>config</code> directory. The directory must be owned by the account that the server runs as.
variable	0 or more	Defines a variable for use in expressions, log formats, and <code>obj.conf</code> parameters. See variable .

Table 3-26 (Cont.) server Subelements

Element	Occurrences	Description
localization	0 or 1	Configures localization. See localization .
http	0 or 1	Configures the HTTP protocol options. See http .
keep-alive	0 or 1	Configures the HTTP keep-alive subsystem. See keep-alive .
thread-pool	0 or 1	Configures the HTTP request processing threads. See thread-pool .
stats	0 or 1	Configures the statistics collection subsystem. See stats .
dns	0 or 1	Configures the server's use of DNS. See dns .
dns-cache	0 or 1	Configures the DNS cache. See dns-cache .
ssl-session-cache	0 or 1	Configures the SSL/TLS session cache. See ssl-session-cache .
access-log-buffer	0 or 1	Configures the access log buffering subsystem. See access-log-buffer .
snmp	0 or 1	Configures SNMP. See snmp .
access-log	0 or more	Configures an HTTP access log for the server. See access-log .
http-listener	0 or more	Configures an HTTP listener. See http-listener .
virtual-server	0 or more	Configures a virtual server. See virtual-server .
event	0 or more	Configures a recurring event. See event-subscription .
event-subscription	0 or more	Configures notifications for origin server status change. See event .
origin-server-pool	0 or more	Configures a pool of origin servers that are used for handling load balancing requests. See origin-server-pool .
proxy-cache	0 or 1	Defines the HTTP reverse proxy caching configuration mechanism. See proxy-cache .
qos-limits	0 or 1	Specifies information related to QoS settings. See qos-limits .
status-listener	0 or 1	Configures a Status Listener. See status-listener .
tcp-thread-pool	0 or 1	Configures the TCP request processing threads. See tcp-thread-pool .
tcp-access-log	0 or 1	Configures TCP access log for the server. See tcp-access-log .
tcp-listener	0 or more	Configures a TCP listener. See tcp-listener .
tcp-proxy	0 or more	Configures a TCP service. See tcp-proxy .
webapp-firewall-ruleset	0 or more	Specifies the path to a file containing the Web Application Firewall (WAF) module rules. See webapp-firewall-ruleset .

snmp

The `snmp` element configures the server's SNMP subagent. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-27](#) describes the subelements of `snmp`.

Table 3-27 `snmp` Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies whether the SNMP agent is enabled. If enabled, the SNMP subagent gathers information about the server and passes the information to the master agent. Default value: <code>true</code> .
<code>description</code>	0 or 1	(Optional) Specifies the description of the server. The value must be in text format.
<code>organization</code>	0 or 1	(Optional) Specifies the name of the organization responsible for the server. The value must be in text format.
<code>location</code>	0 or 1	(Optional) Specifies the location of the server. The value must be in text format.
<code>contact</code>	0 or 1	(Optional) Specifies the contact information of the person responsible for the server. The value must be in text format.

Related Topics

- [stats](#)

ssl

The `ssl` element configures the SSL/TLS settings. This element can appear zero or one time within the `http-listener` element. See [http-listener](#).

To configure outgoing proxy SSL connections, this element can appear zero or one time within the `origin-server-pool` element. See [origin-server-pool](#).

[Table 3-28](#) describes the subelements of `ssl`.

Table 3-28 `ssl` Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies whether SSL support is enabled for the listener. Disabled by default for listeners when no <code>cert</code> element is specified, otherwise enabled.
<code>cert</code>	0 or more	Specifies the certificate that the server presents to clients. See cert . You can specify zero or one RSA certificate, and zero or one ECC certificate.

Table 3-28 (Cont.) *ssl* Subelements

Element	Occurrences	Description
numCtx	0 or 1	Allows the creation of more than one NZ global context to get around NZ lock contention in high load situations. Default value: 1.
context-pool-size	0 or 1	Creates a pool to support the re-use of NZ <i>ssl</i> contexts, boosting performance. At creation, the pool is empty, so there is no additional startup time. Once the pool is full, new contexts are created but not reused, and performance drops. This setting should be tuned to the maximum number of expected concurrent SSL connections. Note that increasing the pool size will increase memory usage, as the contexts saved in the pool will not be freed until re-configuration or shutdown. Range of values: -1 to 1000000. If size is 0 or -1, we do not create a pool of reusable contexts. Default value: 4096.
tls10	0 or 1	Specifies whether TLS 1.0 connections are accepted. Default value: <code>false</code> .
tls11	0 or 1	Specifies whether TLS 1.1 connections are accepted. Default value: <code>true</code> .
tls12	0 or 1	Specifies whether TLS 1.2 connections are accepted. Default value: <code>true</code> .
ssl3-tls-ciphers	0 or 1	Configures the TLS cipher suites. See ssl3-tls-ciphers .
client-auth	0 or 1	Specifies the method of client certificate authentication. The value can be <code>required</code> , <code>optional</code> , or <code>false</code> . When you choose <code>required</code> option, the server requests the client for a certificate; if the client does not provide a certificate, the connection is closed. When you choose <code>optional</code> option, the server requests the client for a certificate, but does not require it. The connection is established even if the client does not provide a certificate. Default value: <code>false</code> . The client authentication is disabled by default.

Table 3-28 (Cont.) ssl Subelements

Element	Occurrences	Description
<code>client-auth-timeout</code>	0 or 1	Indicates the duration (in seconds) after which a client authentication handshake fails. The value can be from 0.001 to 3600. Default: 60.
<code>max-client-auth-data</code>	0 or 1	Specifies the number of characters of authentication data that the server can buffer. The value can be from 0 to 2147483647. Default: 1048576.
<code>validate-server-cert-hostname</code>	0 or 1	Specifies whether validate SSL certificate hostname is on or off. Applies only to outgoing connections. The remote certificate or CA must still be trusted locally in the wallet. NZ does not provide a programmatic override if the remote certificate is completely untrusted (for example, self-signed). Default value: <code>true</code>
<code>wallet-location</code>	0 or 1	Allows selection of an alternate wallet for a virtual server, listener, or origin server group. If this is omitted, the wallet from the instance's config directory is omitted. This is primarily to support SNI for multi-tenant, so that each virtual server can use a different wallet.

Related Topics

- [ssl3-tls-ciphers](#)
- [ssl-session-cache](#)

ssl3-tls-ciphers

The `ssl3-tls-ciphers` element configures SSL3 and TLS cipher suites. This element can appear zero or one time within the `ssl` element. See [ssl](#).

Note that if `ssl3-tls-ciphers` is not present in the configuration, the default enablement value for each of the ciphers is used. If `ssl3-tls-ciphers` is present, you must include a `cipher` element for each cipher that you want enabled.

[Table 3-29](#) describes the subelements of `ssl3-tls-ciphers`.

Table 3-29 ssl3-tls-ciphers Subelements

Element	Occurrences	Description
cipher	0 or more	<p>Specifies a TLS cipher to be enabled. This element can appear zero or one time within the <code>ssl3-tls-ciphers</code> element.</p> <p>The following ciphers are supported:</p> <ul style="list-style-type: none"> • TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 • TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 • TLS_RSA_WITH_AES_128_GCM_SHA256 • TLS_RSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 • TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 • TLS_RSA_WITH_AES_128_CBC_SHA256 • TLS_RSA_WITH_AES_256_CBC_SHA256 • TLS_RSA_WITH_AES_128_CBC_SHA <p>If the <code>ssl3-tls-ciphers</code> element is omitted, the above set of ciphers is configured implicitly and in the order specified above. If <code>ssl3-tls-ciphers</code> is present, you must include a cipher element for each cipher that you want enabled.</p>

 **Note:**

Apart from the above, the following ciphers are also supported, but are deprecated and disabled by default:

- TLS_EC
DHE_RS

Table 3-29 (Cont.) ssl3-tls-ciphers Subelements


Element	Occurrences	Description
		<div style="border-left: 2px solid #0070C0; border-right: 2px solid #0070C0; border-bottom: 2px solid #0070C0; padding: 10px;">  <ul style="list-style-type: none"> A_WITH _AES_12 8_CBC_ SHA • TLS_EC DHE_EC DSA_WI TH_AES _128_CB C_SHA • TLS_EC DHE_RS A_WITH _AES_25 6_CBC_ SHA • TLS_EC DHE_EC DSA_WI TH_AES _256_CB C_SHA • TLS_RS A_WITH _AES_25 6_CBC_ SHA • TLS_EC DHE_RS A_WITH _3DES_ EDE_CB C_SHA • TLS_EC DHE_EC DSA_WI TH_3DE S_EDE_ CBC_SH A • SSL_RS A_WITH _3DES_ EDE_CB C_SHA </div>

Table 3-29 (Cont.) ssl3-tls-ciphers Subelements

Element	Occurrences	Description
override-cipher-order	0 or 1	This setting applies only to server-side listeners (ie. HTTPS and TCP listeners), and SNI virtual servers. If set, the server chooses a cipher in the order specified in <ssl3-tls-ciphers>. The first cipher from this list supported by the client is selected. If not set, the first cipher from the ClientHello message supported by the server is selected. Default value: false.

Related Topics

- [ssl](#)

ssl-session-cache

The `ssl-session-cache` element configures the SSL/TLS session cache. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-30](#) describes the subelements of `ssl-session-cache`.

Table 3-30 ssl-session-cache Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the server writes SSL/TLS sessions to the cache. Default value: true.
max-entries	0 or 1	Specifies the maximum number of SSL/TLS sessions that are written to the cache by the server. The value can be from 1 to 524288. Default: 10000.
max-ssl3-tls-session-age	0 or 1	Specifies the maximum amount of time (in seconds) a SSL/TLS session is written to the cache. The value can be from 1 to 86400. Default: 86400.

Related Topics

- [ssl](#)

stats

The `stats` element configures the statistics collection subsystem. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-31](#) describes the subelements of `stats`.

Table 3-31 stats Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the server collects the statistics. Default value: <code>true</code> .
interval	0 or 1	Specifies the interval (in seconds) at which statistics are updated. The value can be from 0.001 to 3600.
profiling	0 or 1	Specifies whether the performance buckets used to track NSAPI function execution time are enabled at runtime. Default value: <code>true</code> .

Related Topics

- [snmp](#)

status-listener

The `stats-listener` element configures dedicated Status Listeners to check the status of Oracle® Fusion Middleware instances. This element can appear zero or one time within the server element. See [server](#).

[Table 3-32](#) describes the subelements of `stats`.

Table 3-32 stats-listener Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the Status Listener is enabled to accept connection requests. Default value: <code>true</code> .
ip	0 or 1	Specifies the IP address to listen.
port	1	Specifies the port to listen.
family	0 or 1	Specifies the protocol family that is used to connect to the origin server. Values: <code>inet</code> , <code>inet6</code> or <code>default</code> . Default value: <code>default</code> .
ssl	0 or 1	Configures SSL/TLS. See ssl .
description	0 or 1	Specifies the description of the Status Listener. The value of this element must be in text format.
blocking-accept	0 or 1	Enables/disables blocking of the server listen socket, while retaining client end points as non-blocking (useful when <code>Maxprocs > 1</code>). Default value: <code>false</code> .

tcp-access-log

The `tcp-access-log` element configures the settings for the TCP access log. If the `tcp-access-log` element is missing TCP access logging is disabled. See [server](#).

[Table 3-33](#) describes the subelements of `tcp-access-log`.

Table 3-33 tcp-access-log Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether TCP access logging is enabled. If the element is enabled, the server writes a log entry for every request received by TCP listeners. Default value: <code>true</code> .
file	1	Specifies the filename of the access log file (absolute path or path relative to the server's config directory).

Related Topics

- [ftp-filter](#)

Related Topics

- [tcp-listener](#)
- [tcp-proxy](#)
- [tcp-thread-pool](#)

tcp-listener

The `tcp-listener` element configures a TCP listener. See [server](#).

[Table 3-34](#) describes the subelements of `tcp-listener`.

Table 3-34 tcp-listener Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the TCP listener is enabled to accept connection requests. Default value: <code>true</code> .
name	1	Specifies the name that uniquely identifies the TCP listener.
ip	0 or 1	Specifies the IP address to listen. The value of this element is a specific IP address or hostname or an asterisk <code>*</code> to listen on all IP addresses.
port	1	Specifies the port to listen. The value of this element is the port number. The value can be from 1 to 65535.
family	0 or 1	Specifies the socket family that is used to connect to the origin server. Values: <code>inet</code> , <code>inet6</code> , <code>inet-sdp</code> and <code>default</code> . <code>inet</code> and <code>inet6</code> represent IPV4 and IPV6 protocols respectively. <code>inet-sdp</code> is used for Sockets Direct Protocol (SDP). Default value: <code>default</code> .
acceptor-threads	0 or 1	Specifies the number of threads dedicated to accept connections received by this listener. The value can be from 1 to 128.
tcp-proxy-name	1	Specifies the name of the TCP proxy that processes requests received by the listener.

Table 3-34 (Cont.) tcp-listener Subelements

Element	Occurrences	Description
listen-queue-size	0 or 1	Specifies the size (in bytes) of the listen queue. Value: 1 to 1048576. Default value: 128.
receive-buffer-size	0 or 1	Specifies the size (in bytes) of the operating system socket receive buffer. Value: 0 to 2147483647.
send-buffer-size	0 or 1	Specifies the size (in bytes) of the operating system socket send buffer. Value: 0 to 2147483647.
ssl	0 or 1	Configures SSL/TLS. If the ssl element is omitted, SSL/TLS support is disabled. See ssl .
description	0 or 1	Specifies the description of the TCP listener. The value of this element must be in text format.
blocking-accept	0 or 1	Enables/disables blocking of the server listen socket, while retaining client end points as non-blocking (useful when MaxProcs > 1). Default value: false.

Related Topics

- [ftp-filter](#)

Related Topics

- [tcp-access-log](#)

Related Topics

- [tcp-proxy](#)
- [tcp-thread-pool](#)

tcp-proxy

The `tcp-proxy` element is used to support LDAP/T3 listeners. See [server](#).

[Table 3-35](#) describes the subelements of `tcp-proxy`.

Table 3-35 tcp-proxy Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the TCP service is enabled. Default value: true.
name	1	A name that uniquely identifies the TCP proxy.
session-idle-timeout	0 or 1	Specifies the maximum timeout (in seconds) that the server waits while receiving/sending data. Range of values: 0.001 and 604800. It defaults to the timeout value under <code>tcp-thread-pool</code> .

Table 3-35 (Cont.) tcp-proxy Subelements

Element	Occurrences	Description
origin-server-pool-name	0 or 1	Specifies the name of a server pool that provides the TCP service. The value must be a name value from an origin-server-pool element.
ftp-filter	0 or 1	Specifies if the FTP protocol is enabled for the TCP proxy. See ftp-filter .

Related Topics

- [tcp-access-log](#)

Related Topics

- [tcp-listener](#)

Related Topics

- [tcp-thread-pool](#)

tcp-thread-pool

The `tcp-thread-pool` element configures the threads used to process WebSocket requests and requests received by TCP listeners. See [server](#).

[Table 3-36](#) describes the subelements of `tcp-thread-pool`.

Table 3-36 tcp-thread-pool Subelements

Element	Occurrences	Description
enabled	0 or 1	Specifies whether the pool is enabled. Default value: true.
threads	0 or 1	Specifies the number of TCP/WebSocket request processing threads. The value can be from 1 to 512. Default value: 1 per CPU.
max-connections	0 or 1	Specifies the maximum number of connection pairs that the server will support. The value can be from 1 to 1048576. Default value: the default value is the value of the keep-alive <code>max-connections</code> value.
timeout	0 or 1	Specifies the idle timeout (in seconds), after which connection pairs will be closed. The value will be overridden by the <code>tcp</code> or <code>WebSocket</code> subsystem. The value can be from 0.001 to 3600. -1 indicates no timeout. Default value: 300 seconds.
stack-size	0 or 1	The value can be from 8192 to 268435456. 0 indicates that the platform-specific default stack size should be used. Default value: 65536.
poll-interval	0 or 1	Specifies the interval (in seconds) between polls. The value can be from 0.001 to 1. Default value: 0.010 seconds.

Table 3-36 (Cont.) tcp-thread-pool Subelements

Element	Occurrences	Description
buffer-size	0 or 1	Specifies the size of the buffer (in bytes), used by each connection for transferring data. The value can be from 512 to 1048576. Default value: 16384.

Related Topics

- [ftp-filter](#)

Related Topics

- [tcp-access-log](#)

Related Topics

- [tcp-listener](#)

Related Topics

- [tcp-proxy](#)

thread-pool

The `thread-pool` element configures the threads used to process HTTP requests. This element can appear zero or one time within the `server` element. See [server](#).

[Table 3-37](#) describes the subelements of `thread-pool`.

Table 3-37 thread-pool Subelements

Element	Occurrences	Description
min-threads	0 or 1	Specifies the minimum number of HTTP request processing threads. The value can be from 1 to 20480. Default value: 1.
max-threads	0 or 1	Specifies the maximum number of HTTP request processing threads. The value can be from 1 to 20480. Default value: If there are 1, 2 CPUs, then the default value is 256 and if there are 3, 4 CPUs, the default value is 512. If there are more than 4 CPUs, the default value is 1024. Note that the the default value will never be more than quarter of the maximum number of file descriptors available to the process.
stack-size	0 or 1	Specifies the stack size (in bytes) for HTTP request processing threads. The value can be from 8192 to 268435456. 0 indicates that the platform-specific default stack size should be used. Default value: 262144.
queue-size	0 or 1	Specifies the maximum number of concurrent HTTP connections that can be queued for processing. The value can be from 1 to 1048576.

Related Topics

- [http](#)

Related Topics

- [keep-alive](#)

Related Topics

- [http](#)
- [keep-alive](#)

time

The `time` element schedules when an event occurs. This element can appear zero or more times within the `event` element. See [event](#).

[Table 3-38](#) describes the subelement of `time`.

Table 3-38 `time` Subelements

Element	Occurrences	Description
<code>time-of-day</code>	1	Specifies the time when the event occurs. The value must be in the <code>hh:mm</code> format.
<code>day-of-week</code>	0 or 1	Specifies the day of the week. The value can be Sun, Mon, Tue, Wed, Thu, Fri, or Sat.
<code>day-of-month</code>	0 or 1	Specifies the day of month. The value can be from 1 to 31.
<code>month</code>	0 or 1	Specifies the name of the month. The value can be Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, or Dec.

variable

The `variable` element defines a variable for use in expressions, log formats, and `obj.conf` parameters. This element can appear zero or more times within the `server` element, and zero or more times within the `virtual-server` element. For more information, see [server](#) and [virtual-server](#).

[Table 3-39](#) describes the subelements of `variable`.

Table 3-39 List of `variable` Subelements

Element	Occurrences	Description
<code>name</code>	1	Specifies the name of the variable. The value must be in text format.
<code>value</code>	1	Specifies the value of the variable. The value must be in text format.
<code>description</code>	0 or 1	The description of the variable. The value must be in text format.

Related Topics

- [property](#)

Related Topics

- [Using Variables, Expressions, Wildcards, and String Interpolation](#)

virtual-server

The `virtual-server` element configures an HTTP virtual server. Each server typically has at least one virtual server. This element can appear zero or more times within the `server` element. See [server](#).

[Table 3-40](#) describes the subelements of `virtual-server`.

Table 3-40 `virtual-server` Subelements

Element	Occurrences	Description
<code>enabled</code>	0 or 1	Specifies whether the virtual server is enabled at runtime. Default value: <code>true</code> .
<code>name</code>	1	A name that uniquely identifies the virtual server.
<code>ssl</code>	0 or 1	Specifies SSL for a <code>virtual-server</code> . SSL is configurable for each virtual server for SNI. You can select a different certificate and cipher for each <code>virtual-server</code> . <code>ssl</code> can be set only on a <code>virtual-server</code> explicitly bound to an <code>http-listener</code> . The default <code>virtual-server</code> cannot contain an <code>ssl</code> element as the SSL parameters would come from the listener's <code>ssl</code> settings. A <code>virtual-server</code> with <code>ssl</code> is accessible only by SNI-capable SSL clients that send an SNI extension and HTTP host header, both of which must match one of the <code>host</code> elements for the <code>virtual-server</code> . To support non-SNI capable clients, configure OTD without including <code>ssl</code> in a <code>virtual-server</code> . For example, configure multiple listeners on different IP addresses with separate certificates, or configure one listener and one certificate with multiple <code>subjectAltNames</code> .
<code>http-listener-name</code>	0 or more	The name of a HTTP listener associated with one or more of the virtual server's host name. The value is the name from an <code>http-listener</code> element. See http-listener .
<code>host</code>	0 or more	Indicates the host name that the virtual-server services. The value can be a specific hostname or a wildcard pattern that matches one or more hostnames. See Wildcard Patterns .

Table 3-40 (Cont.) virtual-server Subelements

Element	Occurrences	Description
canonical-server-name	0 or 1	The canonical name of the virtual server. Requests using a different hostname are redirected to this hostname. The value can be a hostname, fully qualified domain name, IP address, or a URL prefix that contains one (the URL prefix must not contain a path).
object-file	1	The obj.conf file that controls request processing for virtual server. Default value: default-virtual-server-name-obj.conf, and the user can specify any valid file.
default-object-name	0 or 1	The name of the root obj.conf object. Default value: default.
localization	0 or 1	Configures localization. See localization .
access-log	0 or more	Configures an HTTP access log for the virtual server. See access-log .
log-file	0 or 1	Specifies the path to the log file for the virtual server. The value is the log file name, for example, ../logs/errors.
variable	0 or more	Defines an obj.conf variable for the virtual server. See variable .
description	0 or 1	The description of the virtual server.
qos-limits	0 or 1	Specifies information related to QoS settings. See qos-limits .
webapp-firewall-ruleset	0 or more	Specifies the path to a file containing Web Application Firewall (WAF) rules or configuration. See webapp-firewall-ruleset .

Related Topics

- [http](#)

Related Topics

- [http-listener](#)

Related Topics

- [keep-alive](#)

webapp-firewall-ruleset

The `webapp-firewall-ruleset` element configures the path to a web application firewall configuration file, which contains ModSecurity rules/configuration directives. The path may be an absolute path or a relative path. If a relative path is used, it is relative to the server's `config` directory. The file name component may contain wildcard characters to specify multiple files within the given directory.

The `webapp-firewall-ruleset` element may be present at the `virtual-server` level as well as at the `server` level and can appear zero or more times within the `server` and `virtual-server` elements. Configuration settings at the `virtual-server` level

take precedence over the `server` level. However some configuration directives can only be specified at the `server` level. The scope of these directives is considered to be `Main`. Similarly, scope of directives that can be specified at either `server` level or `virtual-server` level is considered to be `Any`. Note that if a directive with `Main` scope is specified within the `virtual-server` level configuration file, then an error will be logged and the server will fail to start. For information about the scope of different directives, see the Web Application Firewall section in the *Oracle Traffic Director Administrator's Guide*.



Note:

For information about various web application firewall use cases, see the appendix, Web Application Firewall Examples and Use Cases in the *Oracle Traffic Director Administrator's Guide*.

4

Syntax and Use of obj.conf

The `obj.conf` file contains directives for HTTP request processing.

If you configure multiple virtual servers using the Oracle Traffic Director Administrator Console or Command-Line Interface (CLI), separate `obj.conf` files will be created for each virtual server. These files are named `virtual-server-name-obj.conf`, where `virtual-server-name` is the name of the virtual server.

In this document, the `obj.conf` file refers to the `obj.conf` file specified by the `object-file` sub element of the `virtual-server` element in the `server.xml` file.

This chapter describes the `obj.conf` directives; the use of `Object`, `Client`, `If`, `ElseIf`, and `Else` tags; the flow of control in `obj.conf`; and the syntax rules for editing `obj.conf`.

This chapter includes the following topics:

- [Request-Handling Process Overview](#)
- [Directives in obj.conf](#)
- [Objects in obj.conf](#)
- [Flow of Control in obj.conf](#)
- [Changes in Function Flow](#)
- [Editing obj.conf](#)

Request-Handling Process Overview

When you first start Oracle Traffic Director, it performs some initialization tasks and then waits for an HTTP request from a client (such as a browser). When Oracle Traffic Director receives a request, it first selects a virtual server. The `obj.conf` file of the selected virtual server determines how Oracle Traffic Director handles a request.

The `obj.conf` file contains a series of instructions known as directives that tell Oracle Traffic Director what to do at each stage in the request-handling process. These directives are grouped inside `Object` tags. Each directive invokes a function with one or more arguments.

Each directive is applied to a specific stage in the request-handling process. For example, a directive that is applied during the authorization stage in the request-handling process is an `AuthTrans` directive.

Steps in the Request-Handling Process

1. `AuthTrans` (authorization translation)
Verify the authorization information (such as name and password) sent in the request.
2. `NameTrans` (name translation)

Translate the logical URI into a local file system path.

3. PathCheck (path checking)

Check the local file system path for validity and check if the requestor has access privileges to the requested resource on the file system.

4. ObjectType (object typing)

Controls the flow of information from Oracle Traffic Director to the origin server and also configures Oracle Traffic Director to origin server connection attributes.

5. Input (prepare to read input)

Select filters that will process incoming request data read by the `Service` step.

6. Output (prepare to send output)

Select filters that will process outgoing response data generated by the `Service` step.

7. Route (request routing)

Select where to route the request.

8. Service (generate the response)

Generate and return the response to the client.

9. AddLog (adding log entries)

Add entries to log files.

10. Error (error handling)

Send an error message to the client and exit processing. This step is executed only if an error occurs in the previous steps.

Directives in obj.conf

The directives in the `obj.conf` file invoke functions known as server application functions (SAFs). Each directive calls a function, indicating when to call it and specifying parameters for it.

The syntax of each directive is:

```
Directive fn="function" name1="value1"...nameN="valueN"
```

The value of the function (`fn`) parameter is the name of the SAF to execute. All directives must supply a value for the `fn` parameter; if there is no function, the instruction does nothing. The remaining parameters are the arguments needed by the function, and they vary from function to function.

Parameters can contain references to variables and expressions. The variables can be predefined variables, variables defined at request time using the `set-variable` SAF, or variables defined in the `server.xml` file. See [set-variable](#). For more information about defining variables in the `server.xml` file, see [variable](#). For more information about expressions and variables, see [Using Variables, Expressions, Wildcards, and String Interpolation](#).

Oracle Traffic Director has a set of built-in SAFs that you can use to create and modify directives in the `obj.conf` file. [Predefined Server Application Functions and Filters in obj.conf](#) describes these SAFs in detail.

The `magnus.conf` file contains `Init` directive SAFs that initialize NASPI plug-ins.

Objects in `obj.conf`

Directives in the `obj.conf` file are grouped into `Object` tags. The default `object` contains instructions to Oracle Traffic Director about how to process requests by default. Each new object modifies the default behavior of the object.

An `Object` tag can contain a `name` or `ppath` attribute. Either parameter can be a wildcard pattern.

Oracle Traffic Director starts handling a request by processing the directives in the default object. However, Oracle Traffic Director switches to processing directives in another object after the `NameTrans` stage of the default object if either of the following conditions is true:

- The successful `NameTrans` directive specifies a `name` argument.
- The physical path name that results from the `NameTrans` stage matches the `ppath` attribute of another object.

When Oracle Traffic Director is alerted to use an object other than the default object, it processes the directives in the other object before processing the directives in the default object. For some steps in the process, Oracle Traffic Director stops processing directives in that particular stage (such as the `Service` stage) as soon as one is successfully executed. Whereas for other stages, Oracle Traffic Director processes all directives in that stage, including the ones in the default object and those in the additional object. See [Flow of Control in `obj.conf`](#).

Objects That Use the `name` Attribute

If a `NameTrans` directive in the default object specifies a `name` argument, Oracle Traffic Director switches to processing the directives in the object of that name before processing the remaining directives in the default object.

Objects That Use the `ppath` Attribute

When Oracle Traffic Director completes processing the `NameTrans` directives in the default object, the logical URL of the request is converted to a physical path name. If this physical path name matches the `ppath` attribute of another object in the `obj.conf` file, Oracle Traffic Director switches to processing the directives in that object before processing the remaining ones in the default object.

Using the `Client`, `If`, `Elseif`, and `Else` Tags

Additional tags are available to use within the `Object` tag. These tags give you greater flexibility when invoking directives within an object.

Client

The `Client` tag enables you to limit the execution of a set of directives to requests received from specific clients. Directives listed within the `Client` tag are executed only when information in the client request matches the parameter values specified.

Table 4-1 describes the parameters for the Client tag.

Table 4-1 Client Tag Parameters

Parameter	Description
browser	The User-Agent string sent by a browser to Oracle Traffic Director.
chunked	A Boolean value set by a client requesting chunked encoding.
code	The HTTP response code.
dns	The DNS name of the client.
internal	The Boolean value indicating internally generated request.
ip	The IP address of the client.
keep-alive	The Boolean value indicating whether the client has requested a keep-alive connection.
keysize	The key size used in an SSL transaction.
match	The match mode for the Client tag. The valid values are all, any, and none.
method	The HTTP method used by the browser.
name	The name of an object as specified in a previous NameTrans statement.
odds	A random value for evaluating the enclosed directive. The value can be a percentage or a ratio (for example, 20% or 1/5).
path	The physical path to the requested resource.
ppath	The physical path of the requested resource.
query	The query string sent in the request.
reason	The text version of the HTTP response code.
restarted	A Boolean value indicating that a request was restarted.
secret-keysize	The secret key size used in an SSL transaction.
security	An encrypted request.
type	The type of document requested (such as text/html or image/gif).
uri	The URI section of the request from the browser.
urlhost	The DNS name of the virtual server requested by the client. (The value is provided in the Host header of the client request).
variable-headers	Prevents access to a specific site, based on the request by the client, for example: <pre>Client variable-headers="Weferer:SKVFWVRKJVZCMHVIBGDA Service type="image/*" fn="deny-existence" </Client></pre>

The Client tag parameter provides greater control when the If directive is executed. In the following example, the odds parameter gives the request a 25% chance of being redirected:

```
<Client odds="25%">
NameTrans fn="redirect"
          from="/Pogues"
          url-prefix="http://pogues.example.com"
</Client>
```

One or more wildcard patterns can be used to specify the `Client` tag parameter values. Wildcards can also be used to exclude clients that match the parameter value specified in the `Client` tag. In the following example, the `Client` tag and the `AddLog` directive are combined to direct Oracle Traffic Director to log access requests from all clients except those from the specified subnet:

```
<Client ip="*~192.85.250.*">
AddLog fn="flex-log" name="access"
</Client>
```

You can also create a negative match by setting the `match` parameter of the `Client` tag to `none`. In the following example, access requests from the specified subnet are excluded as are all requests to the virtual server `example.com`:

```
<Client match="none" ip="192.85.250.*" urlhost="www.example.com">
AddLog fn="flex-log" name="access"
</Client>
```

See [Wildcard Patterns](#).

If, Elself, and Else

Similar to the `Client` tag, these tags can only appear inside an `Object` tag. In addition, these tags can evaluate an expression, then conditionally execute one or more contained directives. See [If, Elself, and Else Tags](#).

Flow of Control in obj.conf

Before Oracle Traffic Director can process a request, it must direct the request to the correct virtual server. After the virtual server is determined, Oracle Traffic Director executes the `obj.conf` file of the specified virtual server. This section describes how Oracle Traffic Director determines the directives to execute in `obj.conf`.

AuthTrans

When Oracle Traffic Director receives a request, it executes the `AuthTrans` directives in the `default` object to check if the client is authorized to access Oracle Traffic Director. If there is more than one `AuthTrans` directive, Oracle Traffic Director executes them in sequence until one succeeds in authorizing the user, unless one of them results in an error. If an error occurs, Oracle Traffic Director skips all other directives except for the `Error` directive.

`AuthTrans` directives work in conjunction with the `PathCheck` directives. The `AuthTrans` directive checks if the user name and password associated with the request are acceptable, but it does not allow or deny access to the request; that is done by the `PathCheck` directive.

The authorization process is split into two steps to incorporate multiple authorization schemes easily and provide the flexibility to have resources that record authorization information.

When a client initially makes a request, the user name and password are unknown. The `AuthTrans` directive gets the user name and password from the headers associated with the request. The `AuthTrans` and `PathCheck` directives work together to reject the request if they cannot validate the user name and password. When a request is rejected, Oracle Traffic Director displays a dialog box. The client includes the user name and password in the headers and resubmits the request.

NameTrans

Oracle Traffic Director executes `NameTrans` directives in the `default` object to associate a named object (for example, an object that specifies the routing rules) with the URL of the requested resource.

Oracle Traffic Director evaluates each `NameTrans` directive in the `default` object in turn, until it finds one that can be applied.

Because Oracle Traffic Director might not execute all `NameTrans` directives, the order in which the directives appear is important.

How and When Oracle Traffic Director Processes Other Objects

As a result of executing a `NameTrans` directive, Oracle Traffic Director might start processing directives in another object. This happens if the `NameTrans` directive that was successfully executed specifies a name or generates a partial path that matches the `name` or `ppath` attribute of another object.

If the successful `NameTrans` directive assigns a name by specifying a `name` argument, Oracle Traffic Director starts processing directives in the named object (defined with the `object` tag) before processing directives in the `default` object for the rest of the request-handling process.

PathCheck

After converting the logical URL of the requested resource to a physical path name in the `NameTrans` step, Oracle Traffic Director executes `PathCheck` directives to verify that the client is allowed to access the requested resource.

If there is more than one `PathCheck` directive, Oracle Traffic Director executes all directives in the order in which they appear, unless one of the directives denies access. If access is denied, Oracle Traffic Director switches to executing directives in the `Error` section.

If the `NameTrans` directive assigned a name or generated a physical path name that matches the `name` or `ppath` attribute of another object, Oracle Traffic Director first applies the `PathCheck` directives in the matching object before applying the directives in the `default` object.

ObjectType

Assuming that the `PathCheck` directives approve access, Oracle Traffic Director next executes the `ObjectType` directives to determine the MIME type of the request. The MIME type has three attributes: `type`, `encoding`, and `language`. When Oracle Traffic Director sends the response to the client, the `type`, `language`, and `encoding` values are transmitted in the headers of the response. The `type` also frequently helps Oracle

Traffic Director to determine which `Service` directive to execute to generate the response to the client.

If there is more than one `ObjectType` directive, Oracle Traffic Director applies all directives in the order in which they appear. However, once a directive sets an attribute of the MIME type, further attempts to set the same attribute are ignored. The reason why all `ObjectType` directives are applied is that one directive can set one attribute, for example `type`, while another directive sets a different attribute, such as `language`.

As with the `PathCheck` directives, if another object was matched to the request as a result of the `NameTrans` step, Oracle Traffic Director executes the `ObjectType` directives in the matching object before executing the `ObjectType` directives in the `default` object.

Input

The `Input` directive selects filters that process incoming request data read by the `Service` step. `Input` directives are invoked when Oracle Traffic Director or plug-in first attempts to read entity body data from the client. You can add the NSAPI filters that process incoming data by invoking the `insert-filter` SAF in the `Input` stage of the request-handling process. NSAPI filters enable a function to intercept and potentially modify the content presented to or generated by another function. The `Input` directives are executed once per request.

The order of `Input fn="insert-filter"` and `Output fn="insert-filter"` directives in `obj.conf` is important if two or more filters are defined to be in the same location in the filter stack. Filters that were inserted later appear higher than filters that were inserted earlier.

Output

The `Output` directive selects filters that process outgoing response data generated by the `Service` step. The `Output` directive allows you to invoke the `insert-filter` SAF to install NSAPI filters that process outgoing data. NSAPI filters enable a function to intercept and potentially modify the content presented to or generated by another function. `Output` directives are executed when Oracle Traffic Director or a plug-in first attempts to write entity body data from the client. The `Output` directives are executed once per request.

The order of `Input fn="insert-filter"` and `Output fn="insert-filter"` directives in `obj.conf` is important if two or more filters are defined to be in the same location in the filter stack. Filters that were inserted later appear higher than filters that were inserted earlier.

Route

If a `Service` directive requires that the HTTP request be sent to another server, Oracle Traffic Director executes `Route` directives to determine how the request be routed. Routing a request can involve selecting Oracle Traffic Director that will service the request and selecting a proxy through which the request is sent.

Service

Oracle Traffic Director executes a `Service` directive to generate the response to send to the client. Oracle Traffic Director looks at each `Service` directive to find the first one that matches the type, method, and query string. If a `Service` directive does not specify type, method, or query string, then the unspecified attribute matches anything.

If there is more than one `Service` directive, Oracle Traffic Director applies the first one that matches the conditions of the request and ignores all remaining `Service` directives.

For the `PathCheck` and `ObjectType` directives, if another object was matched to the request as a result of the `NameTrans` step, Oracle Traffic Director checks the `Service` directives in the matching object before considering the ones in the `default` object. If Oracle Traffic Director successfully executes a `Service` directive in the matching object, it does not execute the `Service` directives in the `default` object, because it only executes one `Service` directive.

AddLog

After Oracle Traffic Director generates the response and sends it to the client, it executes `AddLog` directives to add entries to the log files. All `AddLog` directives are executed. Oracle Traffic Director can add entries to multiple log files.

Error

If an error occurs during the request-handling process, for example, if a `PathCheck` or `AuthTrans` directive denies access to the requested resource or the requested resource does not exist, the SAF sets the HTTP response status code and returns the value `REQ_ABORTED`. When this happens, Oracle Traffic Director stops processing the request. Instead, it searches for an `Error` directive matching the HTTP response status code or its associated reason phrase and executes the directive's function. If Oracle Traffic Director does not find a matching `Error` directive, it returns the response status code to the client.

Changes in Function Flow

There are times when the function flow changes from the normal request-handling process. This happens during internal redirection, restarts, and URI translation functions.

Restarted Requests

Requests can be restarted, for example, a `PathCheck` directive might restart a request for `http://server_name/` as a request for `http://server_name/index.html`.

Internal Requests

Oracle Traffic Director can generate internal requests, for example, an SHTML file or servlet might include a file. While processing the original request, Oracle Traffic Director makes an internal request to retrieve this file.

URI Translation

Oracle Traffic Director can execute `AuthTrans` and `NameTrans` directives to translate a URI to a physical path name without starting a new request.

Editing `obj.conf`

Use caution when making changes to this file. Simple mistakes can make Oracle Traffic Director fail to start or operate incorrectly.

Order of Directives

The order of directives is important, because Oracle Traffic Director executes them in the order in which they appear in `obj.conf`. The outcome of some directives affects the execution of other directives.

For `PathCheck` directives, the order within the `PathCheck` section is not important because Oracle Traffic Director executes all `PathCheck` directives. However, the order within the `ObjectType` section is very important, because if an `ObjectType` directive sets an attribute value, no other `ObjectType` directive can change that value.

Similarly, the order of directives in the `Service` section is very important. Oracle Traffic Director executes the first `Service` directive that matches the current request and does not execute the others.

Parameters

The number and names of parameters depend on the function. The order of parameters on the line is not important.

Case Sensitivity

Items in the `obj.conf` file are case-sensitive including function names, parameter names, parameter values, and path names.

Separators

Function names in the C language can be composed of letters, digits, and underscores. You can use the hyphen (-) in the configuration file in place of underscore (_) for your C code function names. This is only true for function names.

Quotation Marks

Quotation marks (") are only required around value strings when there is a space in the string. Otherwise, they are optional. Each open quotation mark must be matched by a closed quotation mark.

Spaces

- Spaces are not allowed at the beginning of a line except when continuing from the previous line.

- Spaces are not allowed before or after the equal sign (=) that separates the name and value.
- Spaces are not allowed at the end of a line or on a blank line.

Line Continuation

A long line may be continued on the next line by beginning the next line with a space or tab.

Path Names

Always use forward slashes (/), in path names. A backslash escapes the next character.

Comments

Comments begin with a pound sign (#). If you manually add comments to the `obj.conf` file, then use the Administration Console or CLI to make changes to your server: your comments are overwritten when the `obj.conf` file is updated.

5

Predefined Server Application Functions and Filters in `obj.conf`

This chapter describes the predefined server application functions (SAFs) and filters that are used in the `obj.conf` file. For details about the syntax and use of the `obj.conf` file, see [Syntax and Use of `obj.conf`](#).

Each SAF has its own parameters that are passed to it by an `obj.conf` directive. SAFs can examine, modify, or create server variables. Each SAF returns a result code that indicates whether it has succeeded, did nothing, or has failed.

The SAFs in this chapter are grouped by the type of directive that calls them. For an alphabetical list of predefined SAFs and server configuration elements, see [Alphabetical List of Server Configuration Elements and Predefined SAFs](#).

This chapter includes the following topics:

- [The `bucket` Parameter](#)
- [AuthTrans](#)
- [NameTrans](#)
- [PathCheck](#)
- [ObjectType](#)
- [Input](#)
- [Output](#)
- [Route](#)
- [Service](#)
- [AddLog](#)
- [Error](#)
- [Common SAFs](#)

The `bucket` Parameter

The `bucket` parameter is common to all SAFs. You can measure the performance of any SAF in the `obj.conf` file by adding a `bucket=bucket-name` parameter to the function, for example, `bucket="cache-bucket"`. The bucket statistics are displayed by the `perfdump` utility, which can be set up through the Administrator Console, CLI, or through the `service-dump` SAF.

The following performance buckets are predefined:

- The `default-bucket` records statistics for the functions not associated with any user-defined or built-in bucket.

- The `all-requests bucket records perfdump` statistics for all NSAPI SAFs, including those in the `default-bucket`.

AuthTrans

The `Authtrans` directive instructs Oracle Traffic Director to check for authorization before allowing a client to access resources. See [AuthTrans](#).

The following `AuthTrans`-class functions are described in detail in this section:

- [get-sslid](#)
- [qos-handler](#)

In addition, the following common SAFs are valid for the `AuthTrans` directive:

- [match-browser](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-priority](#)
- [set-variable](#)

get-sslid

The `get-sslid` function retrieves a string that is unique to the current SSL session and stores it as the `ssl-id` variable in the `Session->client` parameter block.

Note:

This function is provided for backward compatibility. The functionality of `get-sslid` was incorporated into the standard processing of an SSL connection.

If the variable `ssl-id` is present when a CGI is invoked, it is passed to the CGI as the `HTTPS_SESSIONID` environment variable. The `get-sslid` function has no parameters and always returns `REQ_NOACTION`. It has no effect if SSL is not enabled.

qos-handler

The `qos-handler` function examines the current quality of service (QoS) statistics for a virtual server, logs the statistics, and enforces the QoS parameters by returning an error. This function must be the first `AuthTrans` function configured in the `default` object.

Example

```
AuthTrans fn= "qos-handler"
```

Related Topics

- [qos-error](#)

webapp-firewall

The `webapp-firewall` function controls the enabling and disabling of the rule engine. If this function is present in a virtual server specific `obj.conf`, it indicates that the rule engine is enabled for that particular virtual server.

The `webapp-firewall` function is not configured by default and hence, the rule engine is not enabled. If the rule engine is not enabled, neither the directives nor the rules within the configuration files, specified by `webapp-firewall-ruleset` element, are applied.

Note:

- If the directive `SecRuleEngine` is specified within the configuration file(s) specified by the `webapp-firewall-ruleset` element, then it will be ignored. However, this condition is not applicable if `SecRuleEngine` is set to `DetectionOnly` mode.
- If there are other SAFs that could return `REQ_PROCEED`, then the SAF `webapp-firewall` must be on top of the list. If this is not the case, the execution of `webapp-firewall` might get skipped.
- For information about various web application firewall use cases, see the appendix, *Web Application Firewall Examples and Use Cases* in the *Oracle Traffic Director Administrator's Guide*.

[Table 5-1](#) describes parameters for the `webapp-firewall` function. These parameters take precedence over the equivalent settings specified within the `webapp-firewall-ruleset` element.

Table 5-1 webapp-firewall Parameters

Parameter	Equivalent setting within <code>webapp-firewall-ruleset</code>	Description
<code>detect-only</code>	<code>DetectionOnly</code>	<p>(optional) Indicates whether the rule engine should enforce the rules or not.</p> <p>The value <code>true</code> indicates that the directives should be evaluated but the result of the evaluation should not be enforced. This is equivalent to setting the <code>SecRuleEngine</code> directive to <code>DetectionOnly</code>. The value <code>false</code> indicates that the rules should be enforced.</p> <p>If this parameter is not specified and if <code>SecRuleEngine</code> is set to <code>DetectionOnly</code> mode (in the configuration file specified by <code>webapp-firewall-ruleset</code>), then the behavior is the same as setting <code>detect-only</code> to <code>true</code>.</p>

Table 5-1 (Cont.) webapp-firewall Parameters

Parameter	Equivalent setting within webapp-firewall-ruleset	Description
process-request-body	SecRequestBodyAccess	(Optional) Indicates whether request bodies are processed by web application firewall. When the <code>body-buffer-size</code> parameter in <code>server.xml</code> is configured to be a positive value, Oracle Traffic Director buffers the request body in memory, up to the limit defined by <code>body-buffer-size</code> parameter. This parameter dictates whether web application firewall accesses the buffered request body. The value <code>on</code> indicates that request bodies will be processed. The value <code>off</code> indicates that response bodies will not be processed. The default value is set by the <code>SecRequestBodyAccess</code> directive (if any) in the configuration files specified by the <code>webapp-firewall-ruleset</code> element. If the <code>SecRequestBodyAccess</code> directive is not present, the value is <code>off</code> .
process-response-body	SecResponseBodyAccess	(Optional) Indicates whether response bodies are buffered and processed by web application firewall. When response body processing is enabled, the server buffers the entire response body in memory, up to the limit defined by the <code>SecResponseBodyLimit</code> directive (if any) in configuration files specified by the <code>webapp-firewall-ruleset</code> element. If the <code>SecResponseBodyLimit</code> directive is not present, the value is 524288 (512 KB). The value <code>on</code> indicates that response bodies will be processed. The value <code>off</code> indicates that response bodies will not be processed. The default value is set by the <code>SecResponseBodyAccess</code> directive (if any) in the configuration files specified by the <code>webapp-firewall-ruleset</code> directive. If the <code>SecResponseBodyAccess</code> directive is not present, the value is <code>off</code> .

NameTrans

The `NameTrans` directive translates virtual URLs to physical directories on your server. The `NameTrans` directive must appear in the `default` object. See [NameTrans](#).

The following `NameTrans`-class functions are described in detail in this section:

- [assign-name](#)
- [block-request-cookie](#)
- [map](#)
- [reverse-map](#)

- [rewrite](#)
- [sed-request-header](#)
- [strip-params](#)

In addition, the following common SAFs are also valid for the `NameTrans` directive:

- [match-browser](#)
- [redirect](#)
- [restart](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

assign-name

The `assign-name` function specifies the name of an object in the `obj.conf` file that matches the current request. Oracle Traffic Director processes the directives in the named object in preference to those in the default object.

For example, given the following directive in the `default` object:

```
NameTrans fn="assign-name" name="personnel" from="/personnel"
```

Assume that Oracle Traffic Director receives a request for `http://server-name/personnel`. After processing this `NameTrans` directive, Oracle Traffic Director searches for an object named `personnel` in the `obj.conf` file and continues by processing the directives in the `personnel` object.

The `assign-name` function returns `REQ_NOACTION`.

[Table 5-2](#) describes parameters for the `assign-name` function.

Table 5-2 assign-name Parameters

Parameter	Description
<code>from</code>	(Optional) Wildcard pattern that specifies the path to be affected. If you do not specify the <code>from</code> parameter, all paths are affected.
<code>name</code>	Specifies an additional named object in the <code>obj.conf</code> file whose directives are applied to this request.
<code>find-pathinfo-forward</code>	<p>(Optional) Instructs Oracle Traffic Director to look for the <code>PATHINFO</code> forward in the path right after the <code>ntrans-base</code>, instead of backward from the end of path as Oracle Traffic Director function <code>assign-name</code> does by default.</p> <p>The <code>find-pathinfo-forward</code> parameter is ignored if the <code>ntrans-base</code> parameter is not set in <code>rq->vars</code>. By default, <code>ntrans-base</code> is set.</p> <p>This feature can improve performance for certain URLs by reducing the number of statistics performed.</p>

Table 5-2 (Cont.) assign-name Parameters

Parameter	Description
nostat	<p>(Optional) Prevents Oracle Traffic Director from performing a stat on a specified URL.</p> <p>The effect of <code>nostat="virtual-path"</code> in the NameTrans function <code>assign-name</code> is that Oracle Traffic Director assumes that a stat on the specified <code>virtual-path</code> will fail. Therefore, use <code>nostat</code> only when the path of the <code>virtual-path</code> does not exist on the system. For example, use <code>nostat</code> for NSAPI plug-in URLs to improve performance by avoiding unnecessary stats on those URLs.</p> <p>When the default PathCheck server functions are used, Oracle Traffic Director does not stat for the paths <code>/ntrans-base/virtual-path</code> and <code>/ntrans-base/virtual-path/*</code> if <code>ntrans-base</code> is set (the default condition). It does not stat for the URLs <code>/virtual-path</code> and <code>/virtual-path/*</code> if <code>ntrans-base</code> is not set.</p>

Example

```
# This NameTrans directive is in the default object.
NameTrans fn="assign-name" name="proxy-cache" from="/.proxycache"
...
<Object name="proxy-cache">
...additional directives..
</Object>
```

block-request-cookie

The `block-request-cookie` needs one parameter, "name". This new SAF is introduced to block request cookies by its name.

[Table 5-4](#) describes the parameters for the `block-request-cookie` function.

Table 5-3 block-request-cookie Parameter

Parameter	Description
name	(Optional) Specifies an additional named object in the <code>obj.conf</code> file. The directives of the named object are applied to this request.

map

The `map` function maps a request URI to a URL on another server, enabling you to specify that a request should be serviced by another server. To load balance a given URI across multiple servers, use the `map` function in conjunction with the `set-origin-server` function. The `map` function looks for a certain prefix in the URI that the client is requesting. If `map` finds the prefix, it replaces the prefix with the mirror site prefix.

[Table 5-4](#) describes the parameters for the `map` function.

Table 5-4 map Parameters

Parameter	Description
from	The URI prefix to map. The prefix must not contain trailing slashes.
to	The URL prefix to which the request should be mapped. The prefix must not contain trailing slashes. The <code>Host</code> and <code>Port</code> values specified in this parameter are silently ignored.
name	(Optional) Specifies an additional named object in the <code>obj.conf</code> file. The directives of the named object are applied to this request.
rewrite-host	(Optional) Indicates whether the <code>Host</code> HTTP request header is rewritten to match the host specified by the <code>to</code> parameter. In a reverse proxy configuration where the proxy server and origin server service the same set of virtual servers, you can specify <code>rewrite-host="false"</code> . Default value: <code>true</code> . It indicates that the <code>Host</code> HTTP request header is rewritten.

Example

```
NameTrans fn="map" from="/" name="reverse-proxy" to="/"
```

Related Topics

- [set-origin-server](#)

Related Topics

- [reverse-map](#)

reverse-map

The `reverse-map` function rewrites the HTTP response headers when Oracle Traffic Director is functioning as a reverse proxy. `reverse-map` looks for the URL prefix specified by the `from` parameter in certain response headers. If the `from` prefix matches the beginning of the response header value, `reverse-map` replaces the matching portion with the `to` prefix.

[Table 5-5](#) describes the parameters for the `reverse-map` function.

Table 5-5 reverse-map Parameters

Parameter	Description
from	URL prefix to be rewritten.
to	URL prefix that will be substituted in place of the <code>from</code> prefix.
rewrite-location	(Optional) Indicates whether the <code>location</code> HTTP response header is rewritten. Default value: <code>true</code> (the <code>location</code> header is rewritten).
rewrite-content-location	(Optional) Indicates whether the <code>Content-Location</code> HTTP response header is rewritten. Default value: <code>true</code> (the <code>Content-Location</code> header is rewritten).

Table 5-5 (Cont.) reverse-map Parameters

Parameter	Description
<code>rewrite-<i>headername</i></code>	(Optional) Indicates whether the <i>headername</i> HTTP response header is rewritten, where <i>headername</i> is a user-defined header name. With the exception of the Location and Content-Location headers. Default value: <code>false</code> (the <i>headername</i> header is not rewritten).

Example

```
NameTrans fn="reverse-map" from="http://download.oracle.com/app/docs" to="/docs"
```

Related Topics

- [map](#)

rewrite

The `rewrite` function allows flexible mappings between URIs and file system paths.

The following table describes parameters for the `rewrite` function.

Table 5-6 rewrite Parameters

Parameter	Description
<code>from</code>	(Optional) Wildcard pattern that specifies the path of requests that should be rewritten. The default is to match all paths.
<code>root</code>	(Optional) File system path to the effective root document directory.
<code>name</code>	(Optional) Name of an object in <code>obj.conf</code> whose directives will be applied to this request.
<code>path</code>	(Optional) Rewritten partial path. If non-empty, the path must begin with a slash (<code>/</code>).

Example

The following `obj.conf` code maps requests for the URI `~/user/index.html` to the file system path `/home/user/public_html/index.html`:

```
<If $path =~ "^/~([^/]+)(|/.*)$" >
NameTrans fn="rewrite"
    root="/home/$1/public_html"
    path="$2"
</If>
```

Related Topics

- [restart](#)

sed-request-header

The `sed-request-header` rewrites a request header and it needs two parameters, "name" and "sed".

The following table describes parameters for the `sed-request-header` function.

Table 5-7 sed-request-header Parameters

Parameter	Description
<code>name</code>	Header name for the specified variable.
<code>sed</code>	sed expression value for the specified variable.

Example

Rewrite the value of a specified request header `NameTrans fn="sed-request=header" name="x-someheader" sed="s/abcd/123/g"`

strip-params

The `strip-params` function removes the embedded semicolon-delimited parameters from the path. For example, a URI of `/dir1;param1/dir2` would become a path of `/dir1/dir2`. When used, the `strip-params` function should be the first `NameTrans` directive listed.

Example

```
NameTrans fn="strip-params"
```

PathCheck

The `PathCheck` directive checks the URL that is returned after the `NameTrans` step to verify that the client is allowed to access the specified origin server. For more information, see [PathCheck](#).

The following `PathCheck`-class functions are described in detail in this section:

- [check-request-limits](#)
- [deny-existence](#)
- [get-client-cert](#)
- [nt-uri-clean](#)
- [ssl-logout](#)
- [unix-uri-clean](#)

In addition, the following common SAFs are valid for the `PathCheck` directive:

- [match-browser](#)
- [restart](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

check-request-limits

The `check-request-limits` function monitors incoming requests that match a given attribute (for example, client IP address) and computes an average requests per second on a configurable time interval. When requests that match the monitored attribute exceed a threshold that you configure, subsequent matching requests are not serviced until the request rate drops. Use this function to detect possible denial-of-service attacks.

You must specify either `max-rps` or `max-connections`, otherwise `check-request-limits` does nothing. If you do not enter an attribute or attributes to monitor, the function monitors all requests.

By default, the function keeps entries on requests for 300 seconds (5 minutes) before purging them. To adjust this time, use the `init-request-limits` SAF in the `magnus.conf` file.

Table 5-8 describes the parameters for the `check-request-limits` function.

Table 5-8 check-request-limits Parameters

Parameter	Description
<code>max-rps</code>	(Optional) Threshold for matching requests per second. If this threshold is exceeded subsequent connections matching the criteria are not serviced. Because an acceptable threshold value can vary widely between sites, there is no default value for this parameter.
<code>max-connections</code>	(Optional) Maximum number of concurrent matching connections. If Oracle Traffic Director receives a request that matches the criteria while the number of matching requests currently being processed meets or exceeds this number, the request is denied. Note that this number is the current requests at any time, and is independent of the <code>interval</code> parameter. As soon as the number of concurrent requests falls below this limit, new matching requests are processed. Because an acceptable value can vary widely between sites, there is no default value for this parameter.
<code>interval</code>	(Optional) Time interval in seconds during which average requests per second is computed. The <code>max-rps</code> limit is not applied until the next request rate computation. Because potential attackers can have unlimited requests serviced during this interval, balance the length of this interval against the performance cost of recomputing the maximum requests per second. Default value: 30 seconds.
<code>continue</code>	(Optional) The condition that must be met in order for a blocked request type to become available again for servicing. Valid values are: <ul style="list-style-type: none"> <code>silence</code> - Refused requests must fall to zero in a subsequent interval for service to resume. <code>threshold</code> - Refused requests must fall below the <code>max-rps</code> value for service to resume. Default value: <code>threshold</code> .

Table 5-8 (Cont.) check-request-limits Parameters

Parameter	Description
error	(Optional) The HTTP status code to use for blocked requests. Default value: 503 (the Service Unavailable error).
monitor	(Optional) A request attribute to monitor. Request rates are tracked in a bucket named by the value of this parameter. If the <code>monitor</code> parameter is not specified, the matching requests are tracked in an unnamed (anonymous) bucket. Note that these buckets differ from the buckets you specify with the standard <code>obj.conf</code> <code>bucket</code> parameter. Although the value of the <code>monitor</code> parameter can be a fixed string, it is most useful when you use predefined variables, for example, <code>monitor="\$ip"</code> . You can also specify multiple variables, separated by a colon, for example, <code>monitor="\$ip:\$uri"</code> . For a list of predefined variables, see Predefined Variables .

Example

The following example limits a client IP to a maximum request rate of 10 requests per second in the default interval of 30 seconds:

```
PathCheck fn="check-request-limit" monitor="$ip" max-rps="10"
```

The following example limits a client IP to a maximum request rate of 10 requests per second when accessing any Perl CGIs. Other types of requests are unlimited:

```
<If path = "*.pl">
PathCheck fn="check-request-limits" monitor="$ip" max-rps="10"
</If>
```

For more information on using the `If` tag, see [If, Elseif, and Else](#).

The following example limits requests globally for Perl CGIs to 10 requests per second. No specific `monitor` parameter is specified:

```
<If path = "*.pl">
PathCheck fn="check-request-limits" max-rps="10"
</If>
```

The following example limits a client IP from generating more than 10 Perl CGI requests per second, or 5 JSP requests per second. To track the Perl and JSP totals separately, the specified `monitor` parameters contain both a fixed string identifier and the client IP variable:

```
<If path = "*.pl">
PathCheck fn="check-request-limits" max-rps="10" monitor="perl:$ip"
</If>
<If path = "*.jsp">
PathCheck fn="check-request-limits" max-rps="5" monitor="jsp:$ip"
</If>
```

The following example limits any one client IP to no more than 5 connections at a given time:

```
PathCheck fn="check-request-limits" max-connections="2" monitor="$ip"
```

deny-existence

The `deny-existence` function sends a 404 Not Found message when a client tries to access a specified path.

[Table 5-9](#) describes parameters for the `deny-existence` function.

Table 5-9 deny-existence Parameters

Parameter	Description
<code>path</code>	(Optional) Wildcard pattern of the file system path to hide. If the path does not match, the function does nothing and returns <code>REQ_NOACTION</code> . If the path is not provided, it is assumed to match.
<code>bong-file</code>	(Optional) Specifies a file to send rather than responding with the 404 Not Found message. The value is a full file system path.

Example

```
PathCheck fn="deny-existence" path="/opt/oracle/webserver7/docs/private"
```

```
PathCheck fn="deny-existence" bong-file="/svr/msg/go-away.html"
```

get-client-cert

The `get-client-cert` function gets the authenticated client certificate from the SSL3 session. It can apply to all HTTP methods, or only to those that match a specified pattern. It only works when SSL is enabled on Oracle Traffic Director.

If the certificate is present or obtained from the SSL3 session, the function returns `REQ_NOACTION` and allows the request to proceed. Otherwise, it returns `REQ_ABORTED` and sets the protocol status to 403 forbidden, causing the request to fail.

The following table describes parameters for the `get-client-cert` function.

Table 5-10 get-client-cert Parameters

Parameter	Description
dorequest	<p>(Optional) Controls whether to get the certificate, or just test for its presence.</p> <ul style="list-style-type: none"> • 1 tells the function to redo the SSL3 handshake to get a client certificate, if Oracle Traffic Director does not already have the client certificate. This typically causes the client to present a dialog box to the user to select a client certificate. Oracle Traffic Director might already have the client certificate if it was requested on the initial handshake, or if a cached SSL session has been resumed. • 0 tells the function not to redo the SSL3 handshake if Oracle Traffic Director does not already have the client certificate. <p>If a certificate is obtained from the client and verified successfully by Oracle Traffic Director, the ASCII base 64 encoding of the DER-encoded X.509 certificate is placed in the parameter <code>auth-cert</code> in the <code>Request->vars</code> pblock, and the function returns <code>REQ_PROCEED</code>, allowing the request to proceed.</p> <p>Default value: 0.</p>
require	<p>(Optional) Controls whether failure to get a client certificate aborts the HTTP request.</p> <ul style="list-style-type: none"> • 1 tells the function to abort the HTTP request if the client certificate is not present after <code>dorequest</code> is handled. In this case, the HTTP status is set to <code>PROTOCOL_FORBIDDEN</code>, and the function returns <code>REQ_ABORTED</code>. • 0 tells the function to return <code>REQ_NOACTION</code> if the client certificate is not present after <code>dorequest</code> is handled. <p>Default value: 1.</p>
method	<p>(Optional) Specifies a wildcard pattern for the HTTP methods for which the function will be applied. If <code>method</code> is absent, the function is applied to all requests.</p>

Example

```
# Get the client certificate from the session.
# If a certificate is not already associated with the session, request one.
# The request fails if the client does not present a
#valid certificate.
PathCheck fn="get-client-cert" dorequest="1"
```

nt-uri-clean

(Windows only) The `nt-uri-clean` function denies access to any resource whose physical path contains `\.`, `\.\` or `\\` (these are potential security problems).

[Table 5-11](#) describes the parameters for the `nt-uri-clean` function.

Table 5-11 nt-uri-clean Parameters

Parameter	Description
tildeok	(Optional) If present, allows tilde (~) characters in URIs. This is a potential security risk on the Windows platform, where, for example, longfi~1.htm might reference longfilename.htm but does not go through the proper ACL checking. If present, "/" sequences are allowed.
dotdirok	(Optional) If present, /. / sequences are allowed.

Example

```
PathCheck fn="nt-uri-clean"
```

Related Topics

- [unix-uri-clean](#)

ssl-logout

The `ssl-logout` function invalidates the current SSL session in Oracle Traffic Director's SSL session cache. This does not affect the current request, but the next time that the client connects, a new SSL session is created. If SSL is enabled, this function returns `REQ_PROCEED` after invalidating the session cache entry. If SSL is not enabled, it returns `REQ_NOACTION`.

unix-uri-clean

(UNIX only) The `unix-uri-clean` function denies access to any resource whose physical path contains `./` or `../` or `//` (these are potential security problems).

The following table describes parameters for the `unix-uri-clean` function.

Table 5-12 unix-uri-clean Parameters

Parameter	Description
dotdirok	If present, /. / sequences are allowed.

Example

```
PathCheck fn="unix-uri-clean"
```

Related Topics

- [nt-uri-clean](#)

ObjectType

The `ObjectType` directives determine the MIME type of the file that has to be sent to the client in response to a request. See [ObjectType](#).

The following `ObjectType`-class functions are described in detail in this section:

- block-auth-cert
- block-cache-info
- block-cipher
- block-ip
- block-issuer-dn
- block-jroute
- block-keysize
- block-proxy-agent
- block-secret-keysize
- block-ssl
- block-ssl-id
- block-user-dn
- block-via
- block-xforwarded-for
- forward-auth-cert
- forward-cache-info
- forward-cipher
- forward-ip
- forward-issuer-dn
- forward-jroute
- forward-keysize
- forward-proxy-agent
- forward-secret-keysize
- forward-ssl
- forward-ssl-id
- forward-user-dn
- forward-via
- forward-xforwarded-for
- http-client-config
- proxy-cache-config
- proxy-cache-override-http
- proxy-websocket-config
- reverse-block-date
- reverse-block-server
- reverse-forward-date
- reverse-forward-server
- set-basic-auth

- [set-cache-control](#)
- [set-cookie](#)
- [type-by-exp](#)
- [type-by-extension](#)

In addition, the following common SAFs are valid for the `ObjectType` directive:

- [match-browser](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

block-auth-cert

The `block-auth-cert` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-auth-cert` header to the origin server. In addition, if the incoming request contains this header, then the SAF will allow Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-auth-cert"
```

Related Topics

- [forward-auth-cert](#)

block-cache-info

The `block-cache-info` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-cache-info` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-cache-info"
```

Related Topics

- [forward-cache-info](#)

block-cipher

The `block-cipher` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-cipher` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-cipher"
```

Related Topics

- [forward-cipher](#)

block-ip

The `block-ip` function instructs Oracle Traffic Director to **not** generate and forward its own `Client-ip` header (or `Wl-proxy-client-ip` header for WebLogic Server) to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-ip"
```

Related Topics

- [forward-ip](#)

block-issuer-dn

The `block-issuer-dn` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-issuer-dn` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-issuer-dn"
```

Related Topics

- [forward-issuer-dn](#)

block-jroute

The `block-jroute` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-jroute` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-jroute"
```

Related Topics

- [forward-jroute](#)

block-keysize

The `block-keysize` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-keysize` header (or `Wl-proxy-client-keysize` header for WebLogic Server) to the origin server. In addition, if the incoming request contains this header, then the SAF will allow Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-keysize"
```

Related Topics

- [forward-keysize](#)

block-proxy-agent

The `block-proxy-agent` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-agent` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-proxy-agent"
```

Related Topics

- [forward-proxy-agent](#)

block-secret-keysize

The `block-secret-keysize` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-secret-keysize` header (or `Wl-proxy-client-secretkeysize` header for WebLogic Server) to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-secret-keysize"
```

Related Topics

- [forward-secret-keysize](#)

block-ssl

The `block-ssl` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-ssl` header (or `Wl-proxy-ssl` header for WebLogic Server) to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-ssl"
```

Related Topics

- [forward-ssl](#)

block-ssl-id

The `block-ssl-id` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-ssl-id` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-ssl-id"
```

Related Topics

- [forward-ssl-id](#)

block-user-dn

The `block-user-dn` function instructs Oracle Traffic Director to **not** generate and forward its own `Proxy-user-dn` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-user-dn"
```

Related Topics

- [forward-user-dn](#)

block-via

The `block-via` function instructs Oracle Traffic Director to **not** generate and forward its own `Via` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-via"
```

Related Topics

- [forward-via](#)

block-xforwarded-for

The `block-xforwarded-for` function instructs Oracle Traffic Director to **not** generate and forward its own `X-forwarded-for` header to the origin server. In addition, if the incoming request contains this header, the SAF allows Oracle Traffic Director to pass-through the incoming request containing this header to the origin server.

Example

```
ObjectType fn="block-xforwarded-for"
```

Related Topics

- [forward-xforwarded-for](#)

forward-auth-cert

The `forward-auth-cert` function instructs Oracle Traffic Director to generate information about client's SSL/TLS certificate within the header `Proxy-auth-cert` and forward it to origin server. If an incoming request includes the header `Proxy-auth-cert`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-13](#) describes the parameters for the `forward-auth-cert` function.

Table 5-13 forward-auth-cert Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate the client's DER-encoded SSL/TLS certificate in Base 64 encoding. Default value: <code>Proxy-auth-cert</code> .

Related Topics

- [block-auth-cert](#)

forward-cache-info

The `forward-cache-info` function instructs Oracle Traffic Director to generate information about local hits within the header `Cache-info` and forward it to the origin server. If an incoming request includes the header `Cache-info`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-14](#) describes the parameters for the `forward-cache-info` function.

Table 5-14 forward-cache-info Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate information about local cache hits. Default value: <code>Cache-info</code> .

Related Topics

- [block-cache-info](#)

forward-cipher

The `forward-cipher` function instructs Oracle Traffic Director to generate information about the client's SSL/TLS cipher suite within the header `Proxy-cipher` and forward it to origin server. If an incoming request includes the header `Proxy-cipher`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-15](#) describes the parameters for the `forward-cipher` function.

Table 5-15 forward-cipher Parameters

Parameter	Description
hdr	(Optional) Name of the HTTP request header used to communicate the name of the client's SSL/TLS cipher suite. Default value: <code>Proxy-cipher</code> .

Related Topics

- [block-cipher](#)

forward-ip

The `forward-ip` function instructs Oracle Traffic Director to generate the client's IP address within the header `Client-ip` (or `WI-proxy-client-ip` for WebLogic Server) and forward it to origin server. If an incoming request includes the header `Client-ip` (or `WI-proxy-client-ip` for WebLogic Server), this SAF causes Oracle Traffic Director to remove the header from the request that is forwarded to the origin server. Subsequently, Oracle Traffic Director generates and inserts this header with the appropriate value before forwarding the request to the origin server.

[Table 5-16](#) describes parameters for the `forward-ip` function.

Table 5-16 forward-ip Parameters

Parameter	Description
hdr	(Optional) Name of the HTTP request header used to communicate the client's IP address. Default value: <code>Client-ip</code> , when the origin server is non-WLS and <code>WI-proxy-client-ip</code> -> when origin server is WLS.

Related Topics

- [block-ip](#)

forward-issuer-dn

The `forward-issuer-dn` function instructs Oracle Traffic Director to generate information about the client's SSL/TLS certificate within the header `Proxy-issuer-dn` and forward it to origin server. If an incoming request includes the header `Proxy-issuer-dn`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-17](#) describes the parameters for the `forward-issuer-dn` function.

Table 5-17 forward-issuer-dn Parameters

Parameter	Description
hdr	(Optional) Name of the HTTP request header used to communicate the distinguished name of the issuer of the client's SSL/TLS certificate. Default value: <code>Proxy-issuer-dn</code> .

Related Topics

- [block-issuer-dn](#)

forward-jroute

The `forward-jroute` function instructs Oracle Traffic Director to generate information about request routing within the header `Proxy-jroute` and forward it to origin server. The `Proxy-jroute` header field is used by the `set-origin-server` function and some Servlet containers to implement session stickiness. If an incoming request includes the header `Proxy-jroute`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-18](#) describes the parameters for the `forward-jroute` function.

Table 5-18 forward-jroute Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate the request routing information. Default value: <code>Proxy-jroute</code> .

Related Topics

- [block-jroute](#)

forward-keysize

The `forward-keysize` function instructs Oracle Traffic Director to generate information about the size of the client's SSL/TLS key within the header `Proxy-keysize` and forward it to origin server. If an incoming request includes the header `Proxy-keysize`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-19](#) describes the parameters for the `forward-keysize` function.

Table 5-19 forward-keysize Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate the size of the client's SSL/TLS key. Default value: <code>Proxy-keysize</code> .

Related Topics

- [block-keysize](#)

forward-proxy-agent

The `forward-proxy-agent` function instructs Oracle Traffic Director to generate its version information within the header `Proxy-agent` and forward it to origin server. If an incoming request includes the header `Proxy-agent`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-20](#) describes the parameters for the `forward-proxy-agent` function.

Table 5-20 forward-proxy-agent Parameters

Parameter	Description
hdr	(Optional) Name of the HTTP request header used to communicate server version. Default value: <code>Proxy-agent</code> .

Related Topics

- [block-proxy-agent](#)

forward-secret-keysize

The `forward-secret-keysize` function instructs Oracle Traffic Director to generate information about the size of the client's SSL/TLS secret key within the header `Proxy-secret-keysize` (or `Wl-proxy-client-secretkeysize` for WebLogic Server) and forward it to origin server. If an incoming request includes the header `Proxy-secret-keysize` (or `Wl-proxy-client-secretkeysize` for WebLogic Server), this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

Table 5-21 forward-secret-keysize Parameters

Parameter	Description
hdr	(Optional) Name of the HTTP request header used to communicate the client's SSL/TLS secret key. Default value: <code>Proxy-secret-keysize</code> , when the origin server is non-WLS and <code>Wl-proxy-client-secretkeysize</code> -> when origin server is WLS.

Example**Related Topics**

- [block-secret-keysize](#)

forward-ssl

The `forward-ssl` function instructs the server to forward information to remote (origin) servers to check if the client sent the request to Oracle Traffic Director over an SSL connection. Accordingly, if the client connects to OTD using a non-SSL connection, this header is set with the value `False`. Similarly, if the client connects to OTD using an SSL connection, this header is set with the value `True`. If an incoming request includes the header `Proxy-ssl` (or `WI-proxy-ssl` for WebLogic Server), this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

Table 5-22 forward-ssl

Parameter	Description
hdr	Name of the HTTP request header used to communicate that the connection between the client and OTD was over SSL. Default value: <code>Proxy-ssl</code> , when the origin server is non-WLS and <code>WI-proxy-ssl</code> -> when origin server is WLS.

Example

```
ObjectType fn="forward-ssl"
```

Related Topics

- [block-ssl](#)

forward-ssl-id

The `forward-ssl-id` function instructs Oracle Traffic Director to generate information about the client's SSL/TLS session ID within the header `Proxy-ssl-id` and forward it to origin server. If an incoming request includes the header `Proxy-ssl-id`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-23](#) describes the parameters for the `forward-ssl-id` function.

Table 5-23 forward-ssl-id Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate the client's SSL/TLS session ID. Default value: <code>Proxy-ssl-id</code> .

Related Topics

- [block-ssl-id](#)

forward-user-dn

The `forward-user-dn` function instructs Oracle Traffic Director to generate information about the distinguished name of the subject of the client's SSL/TLS certificate within the header `Proxy-user-dn` and forward it to origin server. If an incoming request includes the header `Proxy-user-dn`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-24](#) describes the parameters for the `forward-user-dn` function.

Table 5-24 forward-user-dn Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate the distinguished name of the subject of the client's SSL/TLS certificate. Default value: <code>Proxy-user-dn</code> .

Related Topics

- [block-user-dn](#)

forward-via

The `forward-via` function instructs Oracle Traffic Director to generate information about request routing within the header `Via` and forward it to origin server using the

HTTP/1.1 `Via` format. The HTTP/1.1 `Via` header field records the proxy servers and protocol versions that were involved in routing a request. If an incoming request includes the header `Via`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

[Table 5-25](#) describes parameters for the `forward-via` function.

Table 5-25 forward-via Parameters

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate routing information. Default value: <code>Via</code> .

Related Topics

- [block-via](#)

forward-xforwarded-for

The `forward-xforwarded-for` function instructs Oracle Traffic Director to generate information about user-specified `X-Forwarded-For` header values within the header `X-Forwarded-For` and forward it to origin server. If the function is enabled, Oracle Traffic Director sends the `X-Forwarded-For` header value to the origin server, where the value is a comma-separated list of IP addresses. This SAF is enabled by default. If an incoming request includes the header `X-forwarded-for`, this SAF causes OTD to remove the header from the request that is forwarded to the origin server.

Table 5-26 forward-xforwarded-for

Parameter	Description
<code>hdr</code>	(Optional) Name of the HTTP request header used to communicate routing information. Default value: <code>X-forwarded-for</code> .

Example

```
ObjectType fn="forward-xforwarded-for"
```

Related Topics

- [block-xforwarded-for](#)

http-client-config

The `http-client-config` function configures Oracle Traffic Director's HTTP client.

[Table 5-27](#) describes the parameters for the `http-client-config` function.

Table 5-27 http-client-config Parameters

Parameter	Description
always-use-keep-alive	(Optional) Indicates whether the HTTP client can reuse existing persistent connections for all types of requests. Default value: true for WLS origin servers. It indicates that GET/HEAD/OPTIONS requests uses keep-alive by default.
exclude-escape-chars	(Optional) Specifies the list of characters that Oracle Traffic Director should not escape. Various applications deployed in the application server require certain characters not to be escaped. If you do not specify this parameter, Oracle Traffic Director might escape those characters. For example: ObjectType fn="http-client-config" exclude-escape-chars="%&"
keep-alive	(Optional) Indicates whether the HTTP client uses persistent connections. Default value: true.
keep-alive-timeout	(Optional) The maximum number of seconds to keep a persistent connection open. Default value: 29.
log-headers	(Optional) Specifies whether to log request or response headers in server log that Oracle Traffic Director sends and receives from the origin server. This parameter is useful for diagnostic purposes. For example: ObjectType fn="http-client-config" log-headers="true"
protocol	(Optional) HTTP protocol version string. By default, the HTTP client uses either HTTP/1.0 or HTTP/1.1 based on the contents of the HTTP request. In general, do not use the protocol parameter unless you encounter specific protocol interoperability problems.
proxy-agent	(Optional) Value of the proxy-agent HTTP request header. The default is a string that contains the web server product name and version.
proxy-buffer-size	(Optional) Specifies the size of the buffer used by Oracle Traffic Director to store data before it is sent to the client. A larger buffer size results in a lower number of write system calls. By default, the value of the proxy buffer size is 16 KB. To change the value to 32 KB, use the parameter as follows: ObjectType fn="http-client-config" proxy-buffer-size="32768"
retries	(Optional) The number of times to retry when getting content from the origin web server before sending an error to the client. Only GET/HEAD/OPTIONS requests are retried by default. If GET requests have a body associated, those requests are also not retried. If unsafe requests such as POST/PUT/DELETE need to be retried on receiving a retryable error response from the origin-server, set the retry-unsafe-requests to true. Acceptable values are 0 through 100. The value 0 indicates that no retries should be attempted. Default value: 3.

Table 5-27 (Cont.) http-client-config Parameters

Parameter	Description
retry-unsafe-requests	If true, unsafe requests (POST/PUT/DELETE) that receive a retryable error response (e.g. 503 error) from the origin server are retried. The number of times a request is retried before an error is sent to client is determined by the property retries.

Example

```
ObjectType fn="http-client-config" keep-alive="false"
```

proxy-cache-config

The `proxy-cache-config` function configures reverse proxy cache settings.

[Table 5-28](#) describes the parameters for the `proxy-cache-config` function.

Table 5-28 proxy-cache-config Parameters

Parameter	Type	Description	Default Value
enabled	Boolean	Enable or disable caching of reverse proxy content.	false
max-reload-interval	Integer	(Optional) Specifies the maximum time in seconds allowed between consecutive up-to-date checks. If set to 0 (the default), a check is made every time the document is accessed, and the <code>last-modified-factor</code> value has no effect. Note: Setting a value for this element will cause the server to behave in a manner different from what the HTTP specification mandates.	3600
min-reload-interval	Integer	Specifies the minimum time in seconds allowed between consecutive up-to-date checks of a cached document.	0
last-modified-factor	Float	(Optional) Represents the factor used in estimating the expiry time, which defines how long a document will be up-to-date based on the time it was last modified. The time elapsed since the last modification is multiplied by this factor. The result gives the estimated time the document is likely to remain unchanged. Specifying a value of 0 turns off this function. The caching system then uses only explicit expiry information, which is rarely available. Only explicit <code>Cache-Control</code> or <code>Expires</code> HTTP headers are used. This value has no effect if <code>max-reload-interval</code> is set to 0.	0
min-object-size	Integer	The minimum size, in bytes, of any document to be cached. This is useful if you prefer to cache only larger documents.	0

Table 5-28 (Cont.) proxy-cache-config Parameters

Parameter	Type	Description	Default Value
max-object-size	Integer	Limits the maximum size of cached documents to the specified value. This value cannot exceed the maximum value specified in <code>server.xml</code> in <code>proxy-cache->max-heap-object-size</code> .	<code>proxy-cache->max-heap-object-size</code>
query-maxlen	Integer	Specifies the number of characters in the query string (the "?string" part at the end of the URL that are still cacheable). The same queries are rarely repeated exactly in the same form by more than one user, and so caching them is often not desirable.	0
compression	Boolean	If this parameter is set to <code>true</code> , the proxy compresses the document before storing. This consumes less cache space, allowing the cache to accommodate more entries for a given cache size. For clients that accept compression, compressed content is sent. For other clients, the server dynamically uncompresses the content. The cache will not store two copies of the same content.	<code>false</code>
cache-https	Boolean	If this parameter is set to <code>true</code> , responses from HTTPS connections are cached. If <code>cache-https</code> is <code>false</code> , responses from HTTPS origin server connections are not cached irrespective of HTTP headers.	<code>false</code>

Example

```
ObjectType fn="proxy-cache-config" enable="1" max-reload-interval=300 min-reload-interval=60
```

Related Topics

- [proxy-cache-override-http](#)
- [service-proxy-cache-dump](#)

proxy-cache-override-http

The `proxy-cache-override-http` function configures reverse proxy cache parameters that override certain HTTP caching rules.

[Table 5-29](#) describes the parameters for the `proxy-cache-override-http` function.

Table 5-29 proxy-cache-override-http Parameters

Parameter	Type	Description	Default Value
ignore-server-no-cache	Boolean	If this parameter is set to true, a pragma: no-cache or cache-control: no-cache header from the origin server is ignored and the response is cached. This behavior violates the HTTP standard.	false
ignore-server-no-store	Boolean	If this parameter is set to true, a cache-control: no-store header from the origin server is ignored and the response is cached. This behavior violates the HTTP standard.	false
ignore-private	Boolean	If this parameter is set to true, a cache-control: private header from the origin server is ignored and the response is cached. This behavior violates the HTTP standard.	false
ignore-client-no-cache	Boolean	If set to true, a pragma: no-cache or cache-control: no-cache header from the client is ignored and the request is served from the cache. This behavior violates the HTTP standard.	false
override-expire	Boolean	If this parameter is set to true, min-reload-interval is enforced over the value of an Expires header and Cache-Control: max-age value. This behavior violates the HTTP standard.	false
override-lastmod	Boolean	If this parameter is set to true, min-reload-interval is enforced over the value of a Last-modified header. This behavior violates the HTTP standard.	false
reload-into-ims	Boolean	If this parameter is set to true, reload request from clients are converted into conditional GET requests with an If-modified-since header. This behavior violates the HTTP standard.	true
require-expires	Boolean	If this parameter is set to true, a response without an Expires header will not be cached. This behavior violates the HTTP standard.	false
without-lastmod	Boolean	If this parameter is set to true, the absence of a Last-modified header is ignored and the response cached. This behavior violates the HTTP standard.	false

Example

```
<If uri =~ '^/images/'>   ObjectType fn="proxy-cache-config" enable="1" max-reload-
interval=600
  ObjectType fn="proxy-cache-override-http" ignore-client-no-cache="true"
</If>
<Else uri =~ '^/myapp/'>
  ObjectType fn="proxy-cache-config" enable="1" max-reload-interval=120
</Else>
```

Related Topics

- [proxy-cache-config](#)

Related Topics

- [service-proxy-cache-dump](#)

proxy-websocket-config

The `proxy-websocket-config` SAF disables WebSocket upgrade and modifies the `idle-timeout` for WebSocket connections. WebSocket upgrade is enabled by default. If WebSocket upgrade must be disabled, `proxy-websocket-config` can be used with `enabled` set to `off`. The `proxy-websocket-config` directive may be present in the route object for a route or the default object for the whole virtual server. This enables administrators to disable WebSocket traffic or to set a different `idle-timeout` value for certain routes or for the whole virtual server.

[Table 5-30](#) describes the parameters for the `proxy-cache-override-http` function.

Table 5-30 proxy-websocket-config Parameters

Parameter	Default Value
<code>enabled</code>	'on' or 'off'
<code>idle-timeout</code>	Default is the timeout value mentioned in the <code>tcp-thread-pool</code> element

Example

```
ObjectType fn="proxy-websocket-config"
```

reverse-block-date

The `reverse-block-date` SAF blocks the `Date` header sent from the origin server and causes Oracle Traffic Director to generate and insert its own `Date` header in the response.

Example

```
ObjectType fn="reverse-block-date"
```

Related Topics

- [reverse-forward-date](#)

reverse-block-server

The `reverse-block-server` SAF blocks the Server header sent from the origin server and causes Oracle Traffic Director to insert its own Server header in the response.

Example

```
ObjectType fn="reverse-block-server"
```

Related Topics

- [reverse-forward-server](#)

reverse-forward-date

The `reverse-forward-date` SAF forwards the Date header sent from the origin server. In Oracle Traffic Director, this is the default behavior.

Example

```
ObjectType fn="reverse-forward-date"
```

Related Topics

- [reverse-block-date](#)

reverse-forward-server

The `reverse-forward-server` SAF forwards the Server header sent from the origin server. If origin server does not generate any Server header then Oracle Traffic Director generates and uses its own Server header. This is the default behavior.

Example

```
ObjectType fn="reverse-forward-server"
```

Related Topics

- [reverse-block-server](#)

set-basic-auth

The `set-basic-auth` function allows you to set the HTTP basic authentication credentials used by the server when it sends an HTTP request. Use `set-basic-auth` to authenticate to a remote origin server or proxy server.

The following table describes parameters for the `set-basic-auth` function.

Table 5-31 set-basic-auth Parameters

Parameter	Description
<code>user</code>	Name of the user to authenticate.
<code>password</code>	Password of the user to authenticate.

Table 5-31 (Cont.) set-basic-auth Parameters

Parameter	Description
hdr	(Optional) Name of the HTTP request header used to communicate the credentials.
bucket	(Optional) Common to all <code>obj.conf</code> functions. Adds a bucket to monitor performance. See The bucket Parameter .

Example

```
ObjectType fn="set-basic-auth" user="admin" password="secret" hdr="proxy-authorization"
```

set-cache-control

The `set-cache-control` function allows you to specify the HTTP caching policy for the response being sent back to the client.

The following table describes parameters for the `set-cache-control` function.

Table 5-32 set-cache-control Parameters

Parameter	Description
control	HTTP cache control directives. Separate multiple directives by commas.

The following table describes some of the useful cache control directives defined by the HTTP/1.1 protocol.

Table 5-33 Cache Control Directives

Directive	Description
public	The response may be cached by any cache.
private	The response must not be cached by a shared cache (for example, a proxy server).
no-cache	Clients must ask Oracle Traffic Director for updated content on each access.
max-age= <i>n</i>	The response should not be cached for more than <i>n</i> seconds.

Example

```
ObjectType fn="set-cache-control" control="private,max-age=60"
```

set-cookie

The `set-cookie` function allows you to set a cookie in the response being sent back to the client.

The following table describes parameters for the `set-cookie` function.

Table 5-34 set-cookie Parameters

Parameter	Description
name	Name of the cookie.
value	(Optional) Value of the cookie. Default value: null.
path	(Optional) Base URI to which the cookie applies. Default value: / (slash).
domain	(Optional) The domain name of servers to which the cookie must be sent. If no domain is specified, web browsers send the cookie only to Oracle Traffic Director that sets the cookie.
max-age	(Optional) Maximum time (in seconds) after which the cookie expires. If max-age is not specified, web browsers delete the cookie when the user closes the web browser.

Example

```
<If not defined $cookie{'FIRSTVISITTIME'}>
ObjectType fn="set-cookie"
    name="FIRSTVISITTIME"
    value="$time"
    max-age="31536000"
</If>
```

type-by-exp

The `type-by-exp` function matches the current path with a wildcard expression. If they match, the `type` parameter information is applied to the file. This is the same as `type-by-extension`, except that you use wildcard patterns for the files or directories specified in the URLs.

[Table 5-35](#) describes the parameters for the `type-by-exp` function.

Table 5-35 type-by-exp Parameters

Parameter	Description
exp	Wildcard pattern of paths for which this function is applied.
type	(Optional) Type assigned to a matching request (the <code>Content-Type</code> header).
enc	(Optional) Encoding assigned to a matching request (the <code>Content-Encoding</code> header).
lang	(Optional) Language assigned to a matching request (the <code>Content-Language</code> header).
charset	(Optional) The character set for the <code>magnus-charset</code> parameter in <code>rq->srvhdrs</code> . If a browser sends the <code>Accept-Charset</code> header or its User-Agent is Mozilla/1.1 or newer, then append <code>"; charset=charset"</code> to <code>Content-Type</code> , where <code>charset</code> is the value of the <code>magnus-charset</code> parameter in <code>rq->srvhdrs</code> .

Example

```
ObjectType fn="type-by-exp" exp="*.test" type="application/html"
```

Related Topics

- [type-by-extension](#)

type-by-extension

The `type-by-extension` function instructs Oracle Traffic Director to look in a table of MIME type mappings to find the MIME type of the requested resource. The MIME type is added to the `Content-Type` header that is sent back to the client.

The table of MIME type mappings is created by a `mime-file` element in the `server.xml` file, which loads a MIME types file or list and creates the mappings.

For example, the following two lines are part of a MIME types file:

```
type=text/html exts=htm,html  
type=text/plain exts=txt
```

If the extension of the requested resource is `htm` or `html`, the `type-by-extension` file sets the type to `text/html`. If the extension is `.txt`, the function sets the type to `text/plain`.

Example

```
ObjectType fn="type-by-extension"
```

Related Topics

- [type-by-exp](#)

Input

The `Input` directives allow you to select filters that process incoming request data read by the `Service` stage. See [Input](#).

The following common SAFs are valid for the `Input` directive:

- [insert-filter](#)
- [match-browser](#)
- [remove-filter](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

Every `Input` directive has the following optional parameters.

Table 5-36 Input Directive's Optional Parameters

Parameters	Description
type	(Optional) Specifies a wildcard pattern of MIME types for which this function is executed.
method	(Optional) Specifies a wildcard pattern of HTTP methods for which this function is executed. Common HTTP methods are GET, HEAD, and POST.
query	(Optional) Specifies a wildcard pattern of query strings for which this function is executed.

Output

The `Output` stage allows you to select filters that will process outgoing data. For more information, see [Output](#).

Every `Output` directive has the following optional parameters:

Table 5-37 Output Directive's Optional Parameters

Parameters	Description
type	(Optional) Specifies a wildcard pattern of MIME types for which this function is executed.
method	(Optional) Specifies a wildcard pattern of HTTP methods for which this function is executed. Common HTTP methods are GET, HEAD, and POST.
query	(Optional) Specifies a wildcard pattern of query strings for which this function is executed.

The following common SAFs are valid for the `Output` directive:

- [insert-filter](#)
- [match-browser](#)
- [redirect](#)
- [remove-filter](#)
- [restart](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [sed-response-header](#)
- [set-variable](#)

Route

The `Route` directive specifies information as to where the Web Server should route requests. See [Route](#).

The following `Route`-class functions are described in detail in this section:

- [set-origin-server](#)
- [set-proxy-server](#)

In addition, the following common SAFs are valid for the Route directive:

- [match-browser](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

set-origin-server

The `set-origin-server` function distributes the load across a set of homogeneous HTTP origin servers. This SAF chooses the origin server from a given origin server pool for this request. The `set-origin-server` SAF requires `origin-server-pool` as mandatory parameter.

[Table 5-38](#) describes the parameters for the `set-origin-server` function.

Table 5-38 set-origin-server Parameters

Parameter	Description
<code>origin-server-pool</code>	(Mandatory) Name of the configured origin server pool. From this pool, one of the origin servers will be chosen based on the load balancing properties defined within the <code>origin-server-pool</code> element in <code>server.xml</code> .
<code>sticky-cookie</code>	(Optional) Name of a cookie that, when present in a response, causes subsequent requests to stick to that origin server. Accordingly, subsequent requests with this cookie are sent to the same origin server. This parameter accepts <code>*</code> as value, which means that any cookie received from the origin server will be considered as sticky. Default value: <code>*</code> .
<code>sticky-param</code>	(Optional) Name of a URI parameter to inspect for route information. When the URI parameter is present in a request URI and its value contains a colon (<code>:</code>) followed by a route ID, the request will stick to the origin server identified by that route ID. Default value: <code>jsessionid</code> .
<code>route-hdr</code>	(Optional) Name of the HTTP request header used to communicate route IDs to origin servers. <code>set-origin-server</code> associates each origin server named by a server parameter with a unique route ID. Origin servers may encode this route ID in the URI parameter named by the <code>sticky-param</code> parameter to cause subsequent requests to stick to them. Default value: <code>Proxy-jroute</code> .
<code>route-cookie</code>	(Optional) Name of the cookie generated by Oracle Traffic Director when it encounters a <code>sticky-cookie</code> in a response. The <code>route-cookie</code> parameter stores a route ID that enables Oracle Traffic Director to direct subsequent requests back to the same origin server. Default value: <code>JROUTE</code> .

Table 5-38 (Cont.) set-origin-server Parameters

Parameter	Description
<code>rewrite-host</code>	(Optional) Indicates whether the host HTTP request header is rewritten to match the host specified by the <code>server</code> parameter. Default value: <code>false</code> . It indicates that the host header is not rewritten.
<code>rewrite-location</code>	(Optional) Indicates whether the Location HTTP response header that matches the <code>server</code> parameter should be rewritten. Default value: <code>true</code> . It indicates that the matching Location headers are rewritten.
<code>rewrite-content-location</code>	(Optional) Indicates whether the Content-Location HTTP response header that matches Oracle Traffic Director parameter should be rewritten. Default value: <code>true</code> . It indicates that the matching Content-Location headers are rewritten.
<code>rewrite-headername</code>	(Optional) Indicates whether the <code>headername</code> HTTP response headers that match Oracle Traffic Director parameter should be rewritten, where <code>headername</code> is a user-defined header name. <code>headername</code> is in lowercase. With the exception of the Location and Content-Location headers, Default value: <code>false</code> . It indicates that the <code>headername</code> header is not rewritten.

Example

```
Route fn="set-origin-server" origin-server-pool="origin-server-pool-1"
```

Related Topics

- [map](#)

Related Topics

- [set-proxy-server](#)

Related Topics

- [proxy-retrieve](#)

set-proxy-server

The `set-proxy-server` function directs Oracle Traffic Director to retrieve the current resource from a particular proxy server.

[Table 5-39](#) describes the parameters for the `set-proxy-server` function.

Table 5-39 set-proxy-server Parameters

Parameter	Description
<code>server</code>	URL of the remote proxy server. If multiple server parameters are given, Oracle Traffic Director distributes load among the specified remote servers.

Example

```
Route fn="set-proxy-server"
      server="http://webcache1.eng.sun.com:8080"
      server="http://webcache2.eng.sun.com:8080"
```

Related Topics

- [set-origin-server](#)

Related Topics

- [proxy-retrieve](#)

Service

The `Service` directives send the response data to the client. See [Service](#).

Every `Service` directive has the following optional parameters to determine whether the function is executed. All optional parameters must match the current request for the function to be executed.

Table 5-40 Service Directive's Optional Parameters

Optional Parameters	Description
<code>type</code>	Specifies a wildcard pattern of MIME types for which this function is executed. The <code>magnus-internal/*</code> MIME types are used only to select a <code>Service</code> function to execute.
<code>method</code>	Specifies a wildcard pattern of HTTP methods for which this function is executed. Common HTTP methods are <code>GET</code> , <code>HEAD</code> , and <code>POST</code> .
<code>query</code>	Specifies a wildcard pattern of query strings for which this function is executed.
<code>UseOutputStreamSize</code>	Determines the default output stream buffer size (in bytes), for data sent to the client. If this parameter is not specified, the default is 8192 bytes. Note: Set this parameter to zero (0) to disable output stream buffering.
<code>flushTimer</code>	Determines the maximum number of milliseconds between write operations in which buffering is enabled. If the interval between subsequent write operations is greater than the <code>flushTimer</code> value for an application, further buffering is disabled. This is necessary for monitoring the status of CGI applications that run continuously and generate periodic status update reports. If this parameter is not specified, the default is 3000 milliseconds.
<code>ChunkedRequestBufferSize</code>	Determines the default buffer size, in bytes, for unchunking request data. If this parameter is not specified, the default is 8192 bytes.
<code>ChunkedRequestTimeout</code>	Determines the default timeout, in seconds, for unchunking request data. If this parameter is not specified, the default is 60 seconds.

If there is more than one `Service`-class function, the first one matching the optional wildcard parameters (`type`, `method`, and `query`) are executed.

The `UseOutputStreamSize`, `ChunkedRequestBufferSize`, and `ChunkedRequestTimeout` parameters also have equivalent `magnus.conf` directives. The `obj.conf` parameters override the `magnus.conf` directives.

By default, Oracle Traffic Director sends the requested file to the client by calling the `send-file` function. The directive that sets the default is:

```
Service method="(GET|HEAD)" type="*~magnus-internal/*" fn="send-file"
```

This directive usually comes last in the set of `Service`-class directives to give all other `Service` directives a chance to be invoked. This directive is invoked if the method of the request is `GET`, `HEAD`, or `POST`, and the type does not start with `magnus-internal/`. Note here that the pattern `*~` means "does not match." For a list of characters that can be used in patterns, see [Wildcard Patterns](#).

The functions used in the `Service` directive are described in the following sections:

- [proxy-retrieve](#)
- [remove-filter](#)
- [service-proxy-cache-dump](#)
- [service-trace](#)
- [stats-xml](#)

In addition, the following common SAFs are valid for the `Service` directive:

- [match-browser](#)
- [remove-filter](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

proxy-retrieve

The `proxy-retrieve` function retrieves a document from a remote server and returns it to the client. This function also enables you to configure Oracle Traffic Director to allow or block arbitrary methods. This function only works on the HTTP protocol.

[Table 5-41](#) describes the parameters for the `proxy-retrieve` function.

Table 5-41 proxy-retrieve Parameters

Parameter	Description
<code>type</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of MIME types for which this function is executed. See Service .
<code>method</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. See Service .

Table 5-41 (Cont.) proxy-retrieve Parameters

Parameter	Description
query	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of query strings for which this function is executed. See Service .
UseOutputStreamSize	(Optional) Common to all <code>Service</code> -class functions. Determines the default output stream buffer size (in bytes), for data sent to the client. See Service .
flushTimer	(Optional) Common to all <code>Service</code> -class functions. Determines the maximum number of milliseconds between write operations in which buffering is enabled. See Service .
ChunkedRequestBufferSize	(Optional) Common to all <code>Service</code> -class functions. Determines the default buffer size, in bytes, for unchunking request data. See Service .
ChunkedRequestTimeout	(Optional) Common to all <code>Service</code> -class functions. Determines the default timeout, in seconds, for unchunking request data. See Service .

Example

```
# Normal proxy retrieve
Service fn="proxy-retrieve"
# Proxy retrieve with POST method disabled
Service fn="proxy-retrieve"
    method="(POST)"
```

Related Topics

- [set-origin-server](#)
- [set-proxy-server](#)

remove-filter

The `remove-filter` function is used to remove a filter from the filter stack. If the filter is inserted multiple times, only the topmost instance is removed. In general, it is not necessary to remove filters with `remove-filter`, as they are removed automatically at the end of a request.

The following table describes parameters for the `remove-filter` function.

Table 5-42 remove-filter Parameters

Parameter	Description
type	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of MIME types for which this function is executed. See Service .
method	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. See Service .

Table 5-42 (Cont.) remove-filter Parameters

Parameter	Description
query	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of query strings for which this function is executed. See Service .
UseOutputStreamSize	(Optional) Common to all <code>Service</code> -class functions. Determines the default output stream buffer size (in bytes), for data sent to the client. See Service .
flushTimer	(Optional) Common to all <code>Service</code> -class functions. Determines the maximum number of milliseconds between write operations in which buffering is enabled. See Service .
ChunkedRequestBufferSize	(Optional) Common to all <code>Service</code> -class functions. Determines the default buffer size, in bytes, for un-chunking request data. See Service .
ChunkedRequestTimeout	(Optional) Common to all <code>Service</code> -class functions. Determines the default timeout, in seconds, for un-chunking request data. See Service .

Example

```
Service fn="remove-filter"
```

service-proxy-cache-dump

The `service-proxy-cache-dump` function dumps the current reverse proxy caching statistics.

[Table 5-43](#) describes the parameters for the `service-proxy-cache-dump` function.

Table 5-43 service-proxy-cache-dump Parameters

Parameter	Description
list	Lists the objects in the cache. The cache listing includes the URI, a set of flags, the current number of references to the cache entry and the size of the entry.
refresh=n	Setting this parameter to a value <code>n</code> causes the client to reload the page every <code>n</code> seconds.
restart	Stops and restarts the cache.
start	Starts the cache.
stop	Stops the cache.

Example

```
<Object name="default"
NameTrans fn=assign-name name="proxy-cache" from="/.proxycache"
/>Object>
<Object name="proxy-cache">
  Service fn="service-proxy-cache-dump"
/>Object>
```

Related Topics

- [proxy-cache-config](#)

Related Topics

- [proxy-cache-override-http](#)

service-trace

The `service-trace` function services `TRACE` requests. `TRACE` requests are used to diagnose problems with web proxy servers located between a web client and web server.

[Table 5-44](#) describes the parameters for the `service-trace` function.

Table 5-44 `service-trace` Parameters

Parameter	Description
<code>type</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of MIME types for which this function is executed. See Service .
<code>method</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. See Service .
<code>query</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of query strings for which this function is executed. See Service .
<code>UseOutputStreamSize</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the default output stream buffer size (in bytes), for data sent to the client. See Service .
<code>flushTimer</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the maximum number of milliseconds between write operations in which buffering is enabled. See Service .
<code>ChunkedRequestBufferSize</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the default buffer size, in bytes, for unchunking request data. See Service .
<code>ChunkedRequestTimeout</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the default timeout, in seconds, for unchunking request data. See Service .

Example

```
<Object name="default">
...
Service method="TRACE" fn="service-trace"
...
</Object>
```

stats-xml

The `stats-xml` function creates a performance report in XML format. If performance buckets are defined, this performance report includes them.

The report is generated at:

```
http://server_id:portURI
```

For example:

```
http://example.com:80/stats-xml
```

The following table describes parameters for the `stats-xml` function.

Table 5-45 stats-xml Parameters

Parameter	Description
<code>type</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of MIME types for which this function is executed. See Service .
<code>method</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. See Service .
<code>query</code>	(Optional) Common to all <code>Service</code> -class functions. Specifies a wildcard pattern of query strings for which this function is executed. See Service .
<code>UseOutputStreamSize</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the default output stream buffer size (in bytes), for data sent to the client. See Service .
<code>flushTimer</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the maximum number of milliseconds between write operations in which buffering is enabled. See Service .
<code>ChunkedRequestBufferSize</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the default buffer size, in bytes, for unchunking request data. See Service .
<code>ChunkedRequestTimeout</code>	(Optional) Common to all <code>Service</code> -class functions. Determines the default timeout, in seconds, for unchunking request data. See Service .

Example

```
<Object name="default">
<If uri = "/stats-xml/*">
Service fn="stats-xml"
</If>
...
</Object>
```

AddLog

The `AddLog` directives are executed to record information about the transaction. For more information, see [AddLog](#).

The following `AddLog`-class function is described in detail in this section:

- [flex-log](#)

In addition, the following common SAFs are valid for the `AddLog` directive:

- [match-browser](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

flex-log

The `flex-log` function records request-specific data in a flexible log format. It can also record requests in the common log format. A log analyzer, `flexanlg`, exists in the `/bin` directory for Web Server. There are also a number of free statistics generators for the common log format.

Specify the log format using the `format` subelement of the `access-log` element in `server.xml`. See [access-log](#). For more information about the log format, see [Using the Custom Access-Log File Format](#).

[Table 5-46](#) describes the parameters for the `flex-log` function.

Table 5-46 flex-log Parameters

Parameter	Description
<code>name</code>	(Optional) Specifies the name of a log file. The name must previously been defined by an <code>access-log</code> element in <code>server.xml</code> . If no name is given, the entry is recorded in the default log file.
<code>iponly</code>	(Optional) Instructs Oracle Traffic Director to log the IP address of the remote client rather than looking up and logging the DNS name. This improves performance if DNS is turned off. The value of <code>iponly</code> has no significance, as long as it exists; you can use <code>iponly=1</code> .

Example

```
# Log all accesses to the default log file
AddLog fn="flex-log"
# Log accesses from outside our subnet (198.93.5.*) to
# nonlocallog
<Client ip="*~198.93.5.*">
AddLog fn="flex-log" name="nonlocallog"
</Client>
```

Error

If an SAF results in an error, Oracle Traffic Director stops executing all other directives and immediately starts executing the `Error` directives. See [Error](#).

The following `Error`-class functions are described in detail in this section:

- [qos-error](#)
- [send-error](#)

In addition, the following common SAFs are valid for the `Error` directive:

- [match-browser](#)
- [redirect](#)

- [remove-filter](#)
- [restart](#)
- [sed-param-name](#)
- [sed-param-value](#)
- [set-variable](#)

qos-error

The `qos-error` function returns an error page stating the quality of service that caused the error and the value of the QoS statistic.

[Table 5-47](#) describes the parameters for the `qos-error` function.

Table 5-47 qos-error Parameters

Parameter	Description
<code>code</code>	<p>(Optional) Three-digit number representing the HTTP response status code, such as 401 or 407. The recommended value is 503.</p> <p>This can be any HTTP response status code or reason phrase according to the HTTP specification.</p> <p>A list of common HTTP response status codes and reason strings is as follows:</p> <ul style="list-style-type: none"> • 401 Unauthorized • 403 Forbidden • 404 Not Found • 500 Server Error

Example

```
Error fn="qos-error" code="503"
```

Related Topics

- [qos-handler](#)

send-error

The `send-error` function sends an HTML file to the client in place of a specific HTTP response status. This allows the server to present a message describing the problem. The HTML page may contain images and links to the server's home page or other pages.

[Table 5-48](#) describes the parameters for the `send-error` function.

Table 5-48 send-error Parameters

Parameter	Description
<code>path</code>	Specifies the absolute path of an HTML file to send to the client. If the file does not exist or is not accessible, the server returns a 404 or 403 error page. The file is sent as text/html regardless of its name or actual type.

Table 5-48 (Cont.) send-error Parameters

Parameter	Description
code	(Optional) Three-digit number representing the HTTP response status code, such as 401 or 407. This can be any HTTP response status code or reason phrase according to the HTTP specification. A list of common HTTP response status codes and reason strings is as follows: <ul style="list-style-type: none"> • 401 Unauthorized • 403 Forbidden • 404 Not Found • 500 Server Error
type	(Optional) Common to all <i>Service</i> -class functions. Specifies a wildcard pattern of MIME types for which this function will be executed. See Service .
method	(Optional) Common to all <i>Service</i> -class functions. Specifies a wildcard pattern of HTTP methods for which this function will be executed. See Service .
query	(Optional) Common to all <i>Service</i> -class functions. Specifies a wildcard pattern of query strings for which this function will be executed. See Service .
bucket	(Optional) Common to all <i>obj.conf</i> functions. Adds a bucket to monitor performance. See The bucket Parameter .

Example

```
Error fn="send-error" code="401" path="/opt/oracle/webserver7/docs/errors/401.html"
```

Common SAFs

This section lists SAFs that are common to multiple directives.

Table 5-49 Common SAFs

Server Application Functions	Directives
insert-filter	<ul style="list-style-type: none"> • Input • Output
match-browser	<ul style="list-style-type: none"> • AuthTrans • NameTrans • PathCheck • ObjectType • Input • Output • Route • Service • AddLog • Error
redirect	<ul style="list-style-type: none"> • NameTrans • Output • Error

Table 5-49 (Cont.) Common SAFs

Server Application Functions	Directives
<code>remove-filter</code>	<ul style="list-style-type: none"> • <code>Input</code> • <code>Output</code> • <code>Service</code> • <code>Error</code>
<code>restart</code>	<ul style="list-style-type: none"> • <code>NameTrans</code>
<code>sed-param-name</code>	<ul style="list-style-type: none"> • <code>AuthTrans</code> • <code>NameTrans</code> • <code>PathCheck</code> • <code>ObjectType</code> • <code>Input</code> • <code>Output</code> • <code>Route</code> • <code>Service</code> • <code>AddLog</code> • <code>Error</code>
<code>sed-param-value</code>	<ul style="list-style-type: none"> • <code>AuthTrans</code> • <code>NameTrans</code> • <code>PathCheck</code> • <code>ObjectType</code> • <code>Input</code> • <code>Output</code> • <code>Route</code> • <code>Service</code> • <code>AddLog</code> • <code>Error</code>
<code>set-variable</code>	<ul style="list-style-type: none"> • <code>AuthTrans</code> • <code>NameTrans</code> • <code>PathCheck</code> • <code>ObjectType</code> • <code>Input</code> • <code>Output</code> • <code>Route</code> • <code>Service</code> • <code>AddLog</code> • <code>Error</code>

insert-filter

The `insert-filter` SAF is used to add a filter to the filter stack to process incoming (client to server) data. The order of `Input fn="insert-filter"` and `Output fn="insert-filter"` directives is important.

Returns

Returns `REQ_PROCEED` if the specified filter was inserted successfully or `REQ_NOACTION` if the specified filter was not inserted because it was not required. Any other return value indicates an error.

Parameters

The following table describes parameters for the `insert-filter` function.

Table 5-50 `insert-filter` Parameters

Parameter	Description
<code>filter</code>	Specifies the name of the filter to insert. For more information about predefined filters, see Input and Output .
<code>type</code>	(Optional) Common to all <code>Input-class</code> and <code>Output-class</code> functions. Specifies a wildcard pattern of MIME types for which this function is executed.
<code>method</code>	(Optional) Common to all <code>Input-class</code> and <code>Output-class</code> functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. Common HTTP methods are GET, HEAD, and POST.
<code>query</code>	(Optional) Common to all <code>Input-class</code> and <code>Output-class</code> functions. Specifies a wildcard pattern of query strings for which this function is executed.

Example

```
Input fn="insert-filter" filter="http-decompression"
```

This directive instructs the `insert-filter` function to add a custom filter, that is, `http-decompression` to the filter stack. The `http-decompression` filter will decompress the incoming HTTP request data before it goes to the service stage. See [Input](#) and [Output](#).

match-browser

The `match-browser` function matches specific strings in the `User-Agent` string supplied by the browser. It then modifies the behavior of Oracle Traffic Director based on the results by setting values for specified variables. This function is applicable in all directives.

Syntax

```
stage fn="match-browser" browser="string" name="value" [name="value" ...]
```

Parameters

The following table describes parameter values for the `match-browser` function.

Table 5-51 `match-browser` Parameters

Value	Description
<code>stage</code>	Stage directive used in <code>obj.conf</code> processing. The <code>match-browser</code> function is applicable in all stage directives.
<code>string</code>	Wildcard pattern to compare with the <code>User-Agent</code> header (for example, <code>*Mozilla*</code>).
<code>name</code>	Variable to be changed. The <code>match-browser</code> function indirectly invokes the <code>set-variable</code> function.
<code>value</code>	New value for the specified variable.

Example

```
AuthTrans fn="match-browser"
          browser="*[Bb]roken*"
          ssl-unclean-shutdown="true"
          keep-alive="disabled"
          http-downgrade="1.0"
```

If a browser's `User-Agent` header contains the string `Broken` or `broken`, the above `AuthTrans` directive instructs Oracle Traffic Director to do the following:

- Not send the SSL3 and TLS `close_notify` packet
- Not honor requests for HTTP Keep-Alive
- Use the HTTP/1.0 protocol rather than HTTP/1.1

For more information on the variables used in this example, such as `ssl-unclean-shutdown`, see [set-variable](#).

Related Topics

- [set-variable](#)

redirect

The `redirect` function lets you change URLs and send the updated URL to the client. When a client accesses your server with an old path, Oracle Traffic Director treats the request as a request for the new URL.

The `redirect` function inspects the URL to which the client will be redirected. If the URL matches the URL the client has requested (same scheme, hostname, port, and path), this function does not perform the redirect and instead returns `REQ_NOACTION`.

[Table 5-52](#) describes the parameters for the `redirect` function.

Table 5-52 redirect Parameters

Parameter	Description
<code>from</code>	(Optional) Specifies the prefix of the requested URI to match. If <code>from</code> is not specified, it defaults to " ".
<code>url</code>	(Optional) Specifies a complete URL to return to the client. If you use this parameter, do not use <code>url-prefix</code> .
<code>url-prefix</code>	(Optional) The new URL prefix to return to the client. The <code>from</code> prefix is replaced by this URL prefix. If you use this parameter, do not use <code>url</code> .
<code>escape</code>	(Optional) Indicates whether the value of the <code>url</code> or <code>url-prefix</code> parameter must be escaped. The default is <code>yes</code> , indicating that Oracle Traffic Director will escape the value. The value <code>no</code> indicates that the URL or URL prefix value has already been escaped. An example of an escaped value is one where any <code>%</code> characters have been replaced with <code>%25</code> and any spaces have been replaced with <code>%20</code> .
<code>status</code>	(Optional) Customizes the HTTP status code. If <code>status</code> is not specified, it defaults to 302.

Table 5-52 (Cont.) redirect Parameters

Parameter	Description
type	(Optional) Common to all Output-class functions. Specifies a wildcard pattern of MIME types for which this function is executed.
method	(Optional) Common to all Output-class functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. Common HTTP methods are GET, HEAD, and POST.
query	(Optional) Common to all Output-class functions. Specifies a wildcard pattern of query strings for which this function is executed.

Example

In the first example, any request for `http://server-name/whatever` is translated to a request for `http://tmpserver/whatever`.

```
NameTrans fn="redirect" from="/" url-prefix="http://tmpserver/"
```

In the second example, any request for `http://server-name/toopopular/whatever` is translated to a request for `http://bigger/better/stronger/morepopular/`.

```
NameTrans fn="redirect" from="/toopopular"
url="http://bigger/better/stronger/morepopular"
```

remove-filter

The `remove-filter` SAF is used to remove a filter from the filter stack. If the filter is inserted multiple times, only the topmost instance is removed. In general, it is not necessary to remove filters with `remove-filter`, as they are removed automatically at the end of a request.

Returns

Returns `REQ_PROCEED` if the specified filter was removed successfully, or `REQ_NOACTION` if the specified filter was not part of the filter stack. Any other return value indicates an error.

Parameters

The following table describes parameters for the `remove-filter` function.

Table 5-53 remove-filter Parameters

Parameter	Description
filter	Specifies the name of the filter to remove.
type	(Optional) Common to all Input-class, Output-class, and Service-class functions. Specifies a wildcard pattern of MIME types for which this function is executed. The <code>magnus-internal/*</code> MIME types are used only to select a Service function to execute.

Table 5-53 (Cont.) remove-filter Parameters

Parameter	Description
method	(Optional) Common to all <code>Input-class</code> , <code>Output-class</code> , and <code>Service-class</code> functions. Specifies a wildcard pattern of HTTP methods for which this function is executed. Common HTTP methods are GET, HEAD, and POST.
query	(Optional) Common to all <code>Input-class</code> , <code>Output-class</code> , and <code>Service-class</code> functions. Specifies a wildcard pattern of query strings for which this function is executed.
UseOutputStreamSize	(Optional) Common to all <code>Service-class</code> functions. Determines the default output stream buffer size (in bytes), for data sent to the client. See Service .
flushTimer	(Optional) Common to all <code>Service-class</code> functions. Determines the maximum number of milliseconds between write operations in which buffering is enabled. See Service .
ChunkedRequestBufferSize	(Optional) Common to all <code>Service-class</code> functions. Determines the default buffer size, in bytes, for unchunking request data. See Service .
ChunkedRequestTimeout	(Optional) Common to all <code>Service-class</code> functions. Determines the default timeout, in seconds, for unchunking request data. See Service .

Example

```
Input fn="remove-filter" filter="http-compression"
```

restart

The `restart` function allows URL rewriting within Oracle Traffic Director without sending an HTTP redirect to the client. The `restart` function replaces the `uri` and `query` values in `rq->reqpb` with the URI and query string specified by the `uri` parameter and restarts the request by returning `REQ_RESTART`.

If the `uri` parameter contains a `?` character, the value following `?` is used as the query string. Otherwise, the restarted request will not have a query string. Because the new request URI will be passed through the `AuthTrans` and `NameTrans` stages again, avoid creating infinite loops.

The following table describes parameters for the `restart` function.

Table 5-54 restart Parameters

Parameter	Description
from	(Optional) Wildcard pattern that specifies the path of requests that should be restarted. The default is to match all paths.
uri	URI and query string to use for the restarted request.

Example

The following `obj.conf` code directs Oracle Traffic Director to service requests for `/index.html` as though they were requests for `/index.jsp`:

```
NameTrans fn="restart" from="/index.html" uri="/index.jsp"
```

Related Topics

- [rewrite](#)

rewrite-cookie

The `rewrite-cookie` function allows modification of cookie/s returned by the back end server.

The following table describes parameters for the `rewrite-cookie` function.

Table 5-55 restart Parameters

Parameter	Type	Description
<code>name</code>	Boolean	It specifies a response cookie name
<code>secure</code>	Boolean	It specifies, if a "secure" attribute needs to be added/dropped from the response cookie specified by "name".
<code>httponly</code>	Boolean	It specifies, if a "HttpOnly" parameter need to be added/dropped from the response cookie specified by "name".
<code>passthroug h</code>	Boolean	It specifies, if the "path" attribute of the response cookie specified by "name" need to be rewritten or not.

sed-param-name

The `sed-param-name` function modifies a specified pblock parameter name based on a given sed expression.

The following table describes parameters for the `sed-param-name` function.

Table 5-56 sed-param-name Parameters

Parameter	Description
<code>name</code>	It specifies a parameter name
<code>sed</code>	It specifies sed expression value for the specified variable.
<code>pblock</code>	It specifies the pblock name.

Example

To convert "Content-length: 123" to "Content-Length: 123" `NameTrans fn="sed-param-name" pblock="headers" name="content-length" sed="s/-length/-Length/g"`

sed-param-value

The `sed-param-value` function modifies a specified pblock parameter value based on a given sed expression.

The following table describes parameters for the `sed-param-value` function.

Table 5-57 sed-param-value Parameters

Parameter	Description
<code>name</code>	It specifies the parameter name whose value will be modified
<code>sed</code>	It specifies sed expression value for the specified variable.
<code>pblock</code>	It specifies the pblock name.

Example

Rewrite an arbitrary pblock parameter value (corresponding to a specified name)
NameTrans fn="sed-param-value" pblock="reqpb" name="uri" sed="s/test/plan/g"

sed-response-header

The `sed-response-header` rewrites a response header and it needs two parameters, "name" and "sed".

Returns

Returns `REQ_ABORTED` if there is an error or `REQ_NOACTION` if the specified filter was not inserted because it was not required.

Parameters

The following table describes parameter values for the `sed-response-header` function.

Table 5-58 sed-response-header Parameters

Value	Description
<code>name</code>	Header name for the specified variable.
<code>sed</code>	sed expression value for the specified variable.

Example

Rewrite the value of a specified response header Output fn="sed-response-header" name="server" sed="s/backend/frontend/g".

set-priority

The `set-priority` function sets the priority (HIGH, NORMAL, or LOW) for other SAFs and filters (such as `check-req-limit`, `set-origin-server`, and bandwidth filters) that are called after it.

If `set-priority` is not used, requests are executed with the default NORMAL priority.

The priority value is stored as an ordinal in `rq->vars` with HIGH stored as 2, NORMAL as 1, and LOW as 0.

The `access-log` format can be modified to include the priority - for example, `%Req->vars.priority%`.

The following table describes parameters for the `set-priority` function.

Table 5-59 restart Parameters

Value	Description
<code>priority</code>	Supported priority values are: <ul style="list-style-type: none"> "HIGH" - the request has higher chance of getting processed before LOW and NORMAL priority requests "NORMAL" - the request has higher chance of getting processed before LOW priority requests "LOW" - the request has lower chance of getting processed before HIGH and NORMAL priority requests The priority values are case sensitive.
<code>uri</code>	URI and query string to use for the restarted request.

Example

The following `obj.conf` code sets the priority for requests based on the value of a URI - setting it to HIGH for `'/index.html'` and LOW for `'foo.html'`:

```
<Object name="default">AuthTrans fn="match-browser" browser="*MSIE*" ssl-unclean-shutdown="true"

<If $uri eq '/index.html'>AuthTrans fn="set-priority" priority="HIGH"</If>

<If $uri eq '/foo.html'>AuthTrans fn="set-priority" priority="LOW"</If>
.....Service fn="proxy-retrieve" method="*"AddLog fn="flex-log"</Object>
```

set-variable

The `set-variable` function enables you to change Oracle Traffic Director settings based upon conditional information in a request. This function is applicable in all directives.

It can also be used to manipulate variables in parameter blocks with the following commands:

- `insert-pblock="name=value"`
Adds a new value to the specified *pblock*.
- `set-pblock="name=value"`
Sets a new value in the specified *pblock*, replacing any existing values with the same name.
- `remove-pblock="name"`
Removes all values with the given name from the specified *pblock*.

The `set-variable` function recognizes many predefined variables as parameters. Additionally, when a `set-variable` parameter name begins with `$` but is not the name

of a predefined variable, the parameter and its value are stored in the `rq->vars` pblock. This functionality allows you to define or override the `$variable` values at the request time.

`set-variable` accepts both the `$variable` and `${variable}` forms, but the name of the parameter stored in the `rq->vars` pblock is always in the `$variable` form.

Syntax

```
stage fn="set-variable" [{insert|set|remove}-pblock="name=value" ...]
[name="value" ...]
```

Parameters

The following table describes parameter values for the `set-variable` function.

Table 5-60 set-variable Parameters

Value	Description
<i>pblock</i>	Specifies one of the following session or request parameter block names: <ul style="list-style-type: none"> <code>client</code>: Contains the IP address of the client machine and the DNS name of the remote machine. <code>vars</code>: Contains the server's working variables, which includes anything not specifically found in the <code>reqpb</code>, <code>headers</code>, or <code>srvhdrs</code> pblocks. The contents of this pblock differ, depending on the specific request and the type of SAF. <code>reqpb</code>: Contains elements of the HTTP request, which includes the HTTP method such as GET or POST, the URI, the protocol (generally HTTP/1.0), and the query string. This pblock does not change during the request-response process. <code>headers</code>: Contains all the request headers (such as <code>User-Agent</code>, <code>If-Modified-Since</code>, and so on) received from the client in the HTTP request. This pblock does not change during the request-response process. <code>srvhdrs</code>: Contains the response headers (such as <code>Server</code>, <code>Date</code>, <code>Content-Type</code>, <code>Content-length</code>, and so on) that are to be sent to the client in the HTTP response.
<i>name</i>	The variable to set.
<i>value</i>	The string assigned to the variable specified by <i>name</i> .

Variables

The following tables lists variables supported by the `set-variable` SAF.

Table 5-61 Supported Variables

Variable	Description
<code>abort</code>	A value of <code>true</code> indicates that the result code should be set to <code>REQ_ABORTED</code> . Setting the result code to <code>REQ_ABORTED</code> will abort the current request and send an error to the browser. For information about result codes, see <i>Administering Oracle Traffic Director</i> .

Table 5-61 (Cont.) Supported Variables

Variable	Description
error	<p>Sets the HTTP status code and exits the request by returning REQ_ABORTED. To set the HTTP status code without exiting the request, use the set-variable error parameter along with the noaction parameter. To rewrite an HTTP status code, use a Client tag to match the original status code and an Output directive to set the new status code.</p> <p>For example, the following code will rewrite all 302 Moved Temporarily responses to 301 Moved Permanently responses:</p> <pre><Client code="302"> Output fn="set-variable" error="301 Moved Permanently" noaction="true" </Client></pre> <p>Sets the error code to be returned in the event of an aborted browser request.</p>
escape	<p>A Boolean value signifying whether a URL should be escaped using util_uri_escape.</p> <p>See <i>Administering Oracle Traffic Director</i>.</p>
find-pathinfo-forward	Path information after the file name in a URI.
http-downgrade	HTTP version number (for example, 1.0).
http-upgrade	HTTP version number (for example, 1.0).
keep-alive	A Boolean value that establishes whether a keep-alive request from a browser will be honored.
name	Specifies an additional named object in the obj.conf file whose directives will be applied to this request. See also assign-name .
noaction	A value of true indicates the result code should be set to REQ_NOACTION. For AuthTrans, NameTrans, Service, and Error stage SAFs, setting the result code to REQ_NOACTION indicates that subsequent SAFs in that stage should be allowed to execute. For information about result codes, see <i>Administering Oracle Traffic Director</i> .
nostat	Causes the server <i>not</i> to perform the stat() function for a URL when possible. See also assign-name .
senthdrs	A Boolean value that indicates whether HTTP response headers have been sent to the client.
ssl-unclean-shutdown	<p>A Boolean value that can be used to alter the way SSL3 connections are closed.</p> <p>Caution: As this violates the SSL3 RFCs, you should only use this with great caution if you know that you are experiencing problems with SSL3 shutdowns.</p>
stop	A value of true indicates the result code should be set to REQ_PROCEED. For AuthTrans, NameTrans, Service, and Error stage SAFs, setting the result code to REQ_PROCEED indicates that no further SAFs in that stage should be allowed to execute.
url	Redirect requests to a specified URL.

Examples

- To deny HTTP keep-alive requests for a specific server class (while still honoring keep-alive requests for the other classes), add this `AuthTrans` directive to the `obj.conf` for the server class, and set the variable `keep-alive` to `disabled`:

```
AuthTrans fn="set-variable" keep-alive="disabled"
```

- To set the same server class to use HTTP/1.0 while the rest of Oracle Traffic Director classes use HTTP/1.1, the `AuthTrans` directive is:

```
AuthTrans fn="set-variable" keep-alive="disabled" http-downgrade="1.0"
```

- To insert an HTTP header into each response, add a `NameTrans` directive to `obj.conf` using the `insert-pblock` command and specify `srvhdrs` as your `Session` or `Request` parameter block.

For example, to insert the HTTP header `P3P`, add the following line to each request:

```
NameTrans fn="set-variable" insert-srvhdrs="P3P"
```

- To terminate processing a request based on certain URIs, use a `Client` tag to specify the URIs and an `AuthTrans` directive that sets the variable `abort` to `true` when there is a match. Your `Client` tag would be as follows:

```
<Client uri="*(system32|root.exe)*">
AuthTrans fn="set-variable" abort="true"
</Client>
```

- To use predefined variables so that Oracle Traffic Director rewrites redirects to host *badname* as redirects to host *goodname*:

```
<If $srvhdrs{'location'} =~ "^(http|https)://badname/(.*)$"
Output fn="set-variable" $srvhdrs{'location'}="$1://goodname/$2"
</If>
```

- To set a `$variable` value at request time:

```
<If "$time_hour:$time_min" < "8:30" || "$time_hour:$time_min" > "17:00">
AuthTrans fn="set-variable" $docroot="/var/www/docs/closed"
</If>
...
NameTrans fn="document-root" root="$docroot"
```

Regardless of whether the `$docroot` variable has been defined in `server.xml`, its value is set to `/var/www/docs/closed` when Oracle Traffic Director is accessed after 5:00 p.m. and before 8:00 a.m. local time.

Related Topics

- [match-browser](#)

A

Using Variables, Expressions, Wildcards, and String Interpolation

The regular expressions that you specify in this field define various conditions under which the server processes a client request. These regular expressions define the `If` conditions in the `obj.conf` file of the Oracle Traffic Director configuration and determine how a set of directives will be processed.

For example, if you specify the condition `$uri = "/images/*"` in the **Route Settings** of `route_1`, the following `If` tag will be defined in the `obj.conf` file of the configuration.

```
<If $uri = "/images/*">  
NameTrans fn="assign-name" id="route_1" name="route_1"  
</If>
```

This appendix contains the following topics:

- [If, Elseif, and Else Tags](#)
- [Variables](#)
- [Expressions](#)
- [String Interpolation](#)
- [Wildcard Patterns](#)

Related Topics

- [variable](#)

If, Elseif, and Else Tags

The `If`, `Elseif`, and `Else` tags enable you to define the conditions under which a set of directives will be executed. These tags can only appear inside an `Object` tag. In addition, these tags can evaluate an expression, then conditionally execute one or more contained directives. The usage of these tags is summarized below:

- `If` and `Elseif` tags offer a richer expression syntax, including support for regular expressions. For more information about the `If` and `Elseif` expression syntax, see [Expressions](#).
- `If`, `Elseif`, and `Else` tags can contain other tags.
- `If` and `Elseif` expressions are evaluated once per request, not once per contained directive.
- `If`, `Elseif`, and `Else` tags cannot contain multiple types of directives.
- Directives within the `If` and `Elseif` tags can contain regular expression backreferences.

When used, an `ElseIf` or `Else` tag must immediately follow an `If` or `ElseIf` tag. `ElseIf` and `Else` tags are skipped if the preceding `If` or `ElseIf` expression evaluates to logical true.

The following example shows `If`, `ElseIf`, and `Else` tag syntax:

```
<If $path eq "/">
<If $browser =~ "MSIE">
NameTrans fn="rewrite" path="/msie.html"
</If>
<ElseIf $browser =~ "Mozilla">
NameTrans fn="rewrite" path="/mozilla.html"
</ElseIf>
<Else>
NameTrans fn="rewrite" path="/unknown.html"
</Else>
</If>
```

This example displays a different page based on whether the browser is Microsoft Internet Explorer, Mozilla Firefox, or another browser.

Variables

The Oracle® Fusion Middleware includes a set of variables predefined by the server, and the capability for you to define custom variables. This section includes the following topics:

- [Predefined Variables](#)
- [Custom Variables](#)
- [Resolving Variables](#)

Predefined Variables

Predefined variables are implicitly defined by the server. The following table lists the predefined variables and their descriptions:

Variable	Description
<code>\$n</code>	Regular expression backreference (value of the <i>n</i> th capturing subpattern, <i>n</i> = 1...9), for example, <code>\$1</code> . Regular expression backreferences are only available within the body of <code>If</code> and <code>ElseIf</code> containers, and only if the container expressions includes one or more regular expressions.
<code>\$&</code>	Value that matched a regular expression. Regular expression backreferences are only available within the body of <code>If</code> and <code>ElseIf</code> containers, and only if the container expressions includes one or more regular expressions.
<code>\$n</code>	Regular expression backreference (value of the <i>n</i> th capturing subpattern, <i>n</i> = 1...9), for example, <code>\$1</code> . Regular expression backreferences are only available within the body of <code>If</code> and <code>ElseIf</code> containers, and only if the container expressions includes one or more regular expressions.

Variable	Description
<code>\$&</code>	Value that matched a regular expression. Regular expression backreferences are only available within the body of <code>If</code> and <code>ElseIf</code> containers, and only if the container expressions includes one or more regular expressions.
<code>\$body</code>	Contains the entity-body (if any) of the current HTTP request that OTD receives from the client (browser). The size of the data stored in this variable is configured using the <code>body-buffer-size</code> sub-element of the <code>http</code> element in <code>server.xml</code> .
<code>\$browser</code>	Web browser version (alias for <code>\$headers{ 'user-agent' }</code> if the client sent a <code>User-Agent</code> header or an empty string).
<code>\$chunked</code>	Boolean variable that indicates whether the request body was sent using chunked encoding.
<code>\$code</code>	Response status code.
<code>\$cookie{ 'name' }</code>	Value of the cookie <i>name</i> from the request.
<code>\$dns</code>	Alias for <code>\$client{ 'dns' }</code> .
<code>\$env{ 'name' }</code>	Value of the environment variable <i>name</i> (includes CGI/SHHTML environment variables).
<code>\$headers{ 'name' }</code>	Value of <i>name</i> from <code>rq->headers</code> , that is, value of the request header name where <i>name</i> is a lowercase string.
<code>\$id</code>	Virtual server ID as specified by the <i>name</i> subelement of the <code>virtual-server</code> element in the <code>server.xml</code> file.
<code>\$internal</code>	Boolean value that indicates whether request was internally generated.
<code>\$ip</code>	Alias for <code>\$client{ 'ip' }</code> .
<code>\$keep_alive</code>	Boolean value that indicates whether the connection is kept open.
<code>\$keysize</code>	Alias for <code>\$client{ 'keysize' }</code> .
<code>\$method</code>	Request method (alias for <code>\$reqpb{ 'method' }</code>).
<code>\$path</code>	Requested path (either URI, partial path, or file system path depending on stage). The predefined variable <code>path</code> is the value of <code>path</code> from <code>rq->vars</code> . If <code>path</code> is not set in <code>rq->vars</code> (for example, if <code>NameTrans</code> has not completed), <code>path</code> gets the value of <code>ppath</code> from <code>rq->vars</code> .
<code>\$path_info</code>	Alias for <code>\$vars{ 'path-info' }</code> .
<code>\$ppath</code>	Alias for <code>\$vars{ 'ppath' }</code> .
<code>\$protocol</code>	Request protocol (alias for <code>\$reqpb{ 'protocol' }</code>).
<code>\$query</code>	Request query string (alias for <code>\$reqpb{ 'query' }</code>).
<code>\$reason</code>	Response reason phrase.
<code>\$referer</code>	Alias for <code>\$headers{ 'referer' }</code> .
<code>\$reqpb{ 'name' }</code>	Value of <i>name</i> from <code>rq->reqpb</code> .
<code>\$restarted</code>	Boolean value that indicates whether the request was restarted.
<code>\$secret_keysize</code>	Alias for <code>\$client{ 'secret-keysize' }</code> .
<code>\$server_url</code>	Prefix for self-referencing URLs.

Variable	Description
<code>\$time</code>	Time the request was received as the number of seconds since 00:00:00 UTC, January 1, 2006.
<code>\$time_day</code>	Day of the month when the request was received. Value can be from 01 to 31.
<code>\$time_hour</code>	Hours since midnight when the request was received. Value can be from 00 to 23.
<code>\$time_min</code>	Minutes after the hour when the request was received. Value can be from 00 to 59.
<code>\$time_mon</code>	Month of the year when the request was received. Value can be from 01 to 12.
<code>\$time_sec</code>	Seconds after the minute when the request was received. Value can be from 00 to 61.
<code>\$time_wday</code>	Day of the week when the request was received. Value can be from 0 to 6, where 0 corresponds to Sunday.
<code>\$time_year</code>	Four-digit year when the request was received.
<code>\$type</code>	Alias for <code>\$srvhdrs{'content-type'}</code> .
<code>\$uri</code>	URI of the requested resource (alias for <code>\$reqpb{'uri'}</code>).
<code>\$url</code>	URL of the requested resource.
<code>\$urlhost</code>	Host name to which the client connected.
<code>\$vars{'name'}</code>	Value of <i>name</i> from <code>rq->vars</code> .
<code>\$security</code>	Boolean value that indicates if a secure transport was used.
<code>\$senthdrs</code>	Boolean value that indicates if a response headers were been sent.
<code>\$srvhdrs{'name'}</code>	Value of <i>name</i> from <code>rq->srvhdrs</code> , that is, value of response header name where <i>name</i> is a lowercase string.

Custom Variables

You can define custom variables in the `server.xml` file using the `variables` element. These variables can then be used in function parameters in `obj.conf` functions. You can also define variables at request time using the `set-variables` function in the `obj.conf` file.

Note:

The predefined variables take precedence over custom variables. It is a best practice to use uppercase names for custom variables. Using uppercase avoids conflicts with the lowercase predefined variables, if the list of predefined variables is extended in the future.

Resolving Variables

The server uses the following order when resolving a `$variable`:

1. Predefined variables

2. Variables defined at request time using `set-variable` in `obj.conf`
3. Variables defined by the `virtual-server` element's `variable` subelement in `server.xml`
4. Variables defined by the `server` element's `variable` subelement in `server.xml`

When you define a `$variable` at request time, it is stored as a name-value pair in the `rq->vars` pblock. These variables are given a higher precedence than `server.xml` variables so that `server.xml` variables can be overridden at request time.

Expressions

Expressions enable you to dynamically construct server application function (SAF) parameters and to select which SAFs to execute on a request-by-request basis. Expressions are constructed from literals, variables, functions, and operators. Use expressions in `If` and `ElseIf` tags, in log format strings, and SAF parameters.

This section contains the following topics:

- [Expression Syntax](#)
- [Expression Results as Boolean Values](#)
- [Expression Literals](#)
- [Expression Variables](#)
- [Expression Operators](#)
- [Expression Functions](#)
- [Regular Expressions](#)

Expression Syntax

The expression syntax is similar to the syntax used in Perl. Expressions are constructed from literals, variables, functions, and operators.

The following example illustrates how to use expressions in an `If` tag:

```
<If not $internal
  and $uri =~ "^/private/(.*)$"
  and $referer !~ "^https://example.com/">
NameTrans fn="redirect"
  url="http://example.com/denied.jsp?file=$1"
</If>
```

This example expression checks to see if a request meets certain criteria, for example if it is an internal request. If it does not meet the criteria, the server redirects the request to a request denied URL.

The expression contains the following components:

- **Literals** - `"/private/(.*)"` and `"^https://example.com/"`
- **Variables** - `$internal`, `$uri`, and `$referer`
- **Operators** - `not`, `and`, `=~`, and `!~`

For more information about `If` and `ElseIf` tags, see [If, Elseif, and Else Tags](#).

Expression Results as Boolean Values

In some circumstances, for example, after evaluating an `If` or `ElseIf` expression, the server must treat the result of an expression as a Boolean value. The server uses the following rules when converting a numeric value to a Boolean value:

- Zero evaluates to false.
- All other numeric values evaluate to true.

The server uses the following rules when converting a string to a Boolean value:

- Zero-length strings evaluate to false.
- The string 0 (zero) evaluates to false.
- All other strings evaluate to true.

Expression Literals

Expression literals are divided into string and numeric literals.

String Literals

A string literal is enclosed by either single quotation marks (`'`) or double quotation marks (`"`). When single quotation marks enclose a string literal, the value of the literal is the value within the quotation marks. When double quotation marks are used, any references to variables or expressions within the quotation marks are interpolated. For more information, see [String Interpolation](#).

The following expression examples show the use of single and double quotation marks.

```
# This expression evaluates to true.
('foo' eq "foo")

# This expression evaluates to false.
('foo' eq "bar")

# This expression evaluates to true.
('foo' eq "f${lc('O')}o")

# This expression may evaluate to true or false,
# depending on the value of the variable $foo
('$foo' eq "$foo")
```

To include an uninterpolated dollar sign `$` in a string enclosed in a double quotation marks, use the two dollar sign or a backslash dollar sign `$$` or `\$` escape sequences.

When a double quotation marks appears within a literal enclosed by double quotation marks, it must be prefixed with a backslash. When a single backslash (`\`) appears within a literal bracketed by double quotes, it must be prefixed with a backslash. When a single quote character appears within a literal bracketed by single quotes, it must be prefixed with a backslash.

The following examples show valid and invalid literals:

```
# The following are examples of valid literals
'this string literal is bracketed by single quotes'
```



```
"this string literal is bracketed by double quotes"  
"the backslash, \\, escapes characters in double quote string literals"  
'it\'s easy to use strings literals'  
  
# The following are examples of invalid literals  
'it's important to escape quote characters'  
"any \ characters in double quote string literals must be escaped"
```

Numeric Literals

A numeric literal can consist of decimal digits and an optional decimal point, a leading zero followed by octal digits, or a leading `0x` prefix followed by hexadecimal digits. Hexadecimal and octal numbers are automatically converted to decimal form.

The following examples show expressions that use numeric literals:

```
# The following expressions evaluate to true  
(1 < 2)  
(0x10 == "16")  
(1 == 1.00)  
  
# The following expressions evaluate to false  
(1 > 2)  
("0x10" == 16)  
(1 != 1.00)
```

Expression Variables

Any `$variable` can be used as a variable in an expression. To mirror the `Client` tag syntax, the dollar sign `$` prefix is optional for predefined variable names in expressions. For example, the following three portions of the `obj.conf` file are semantically equivalent:

```
<If $uri = "*.html">  
...  
</If>  
  
<If uri = "*.html">  
...  
</If>  
  
<Client uri = "*.html">  
...  
</Client>
```

Any variable names you define must use the `$` prefix. For example, the following expression is invalid even if `somecustomvariable` is defined in a `server.xml` variable element:

```
<If somecustomvariable = "foo">  
...  
</If>
```

To make this expression valid, add the dollar sign prefix:

```
<If $somecustomvariable = "foo">  
...  
</If>
```

Expression Operators

The following table lists the operators that are used in expressions.

Operator Symbol	Operator Name
!	C-style logical not
=	Wildcard pattern match
=~	Regular expression match
!~	Regular expression mismatch
+	Addition or unary plus
-	Subtraction or unary minus
.	String concatenation
defined	Value is defined
-d	Directory exists
-e	File or directory exists
-f	File exists
-l	Symbolic link exists
-r	File is readable
-s	File size
-U	URI maps to accessible file or directory
<	Numeric less than
<=	Numeric less than or equal to
>	Numeric greater than
>=	Numeric greater than or equal to
lt	String less than
le	String less than or equal to
gt	String greater than
ge	String greater than or equal to
==	Numeric equal
!=	Numeric not equal
eq	String equal
ne	String not equal
^	C-style exclusive or
&&	C-style logical and
	C-style logical or
not	Logical not
and	Logical and
or	Logical or

Operator Symbol	Operator Name
xor	Logical exclusive or

The following table lists the precedence of operators within expressions from highest to lowest precedence.

Symbol	Operands	Associativity	Description
(), []	0	Left to right	Parentheses
!, unary +, unary -	1	Right to left	Sign operators
=, =~, !~	2	Non-associative	Pattern matching operators
+, -, .	2	Non-associative	Additive operators
defined, -d, -f, -l, -r, -s, -U	1	Right to left	Named operators
<, lt, <=, le, >, gt, >=, ge	2	Non-associative	Relational operators
==, eq, !=, ne	2	Non-associative	Equality operators
^	2	Left to right	C-style exclusive or operator
&&	2	Left to right	C-style logical and operator
	2	Left to right	C-style logical or operator
not	1	Right to left	Logical not operator
and	2	Left to right	Logical and operator
or, xor	2	Left to right	Logical or operators

The numeric operators (<, <=, >, >=, ==, and !=) are intended to operate on numbers and not strings. To facilitate comparing numbers, dates, and timestamps, the numeric operators ignore any white space, colons, slashes, and commas in their arguments. Dashes after the first digit are also ignored.

 **Note:**

It is generally incorrect to use the numeric operators on non-numeric values.

For example, the following expression evaluates to true:

```
# The following expression evaluates to true because both
# "foo" and "bar" are numerically equivalent to 0
("foo" == "bar")
```

Expression Functions

Expression functions manipulate data for use in expressions. Expression functions are different from SAFs. While SAFs perform the actual work associated with an HTTP request, expression functions are used to select which SAFs run and what parameters to pass to the SAFs.

Some expression functions require one or more arguments. An expression function's argument list is enclosed in parentheses () and the individual arguments are separated by commas (,).

The individual expression functions are listed in the following sections:

- [atime](#)
- [choose](#)
- [ctime](#)
- [escape](#)
- [external](#)
- [httpdate](#)
- [lc](#)
- [length](#)
- [lookup](#)
- [lookupregex](#)
- [mtime](#)
- [owner](#)
- [uc](#)
- [unescape](#)
- [uuid](#)

atime

The `atime` function returns the time of the last access for the specified file or directory.

Syntax

```
atime (path)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>path</i>	The absolute path to the directory or file name for which you are requesting the last access.

 **Note:**`ctime, mtime`

choose

The `choose` function parses pipe-separated values from *values* and returns one at random.

Syntax

```
choose (values)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>values</i>	The list of values to choose from, separated by the pipe character ()

Example

The following `obj.conf` code demonstrates the use of `choose` to randomly select one of three images:

```
NameTrans fn="rewrite"
          from="/images/random"
          path="/images/${choose('iwsvi.jpg|0061.jpg|webservervii.jpg')}")"
```

ctime

The `ctime` function returns the time of the last status change for the specified file or directory.

Syntax

```
ctime (path)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>path</i>	The absolute path to the directory or file name for which you are requesting the last status change

 **Note:**`atime, mtime`

escape

The `escape` function encodes the URI using `util_uri_escape`, converting special octets to their percentage encoded equivalent and returns the result.

Syntax

```
escape(uri)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>uri</i>	The URI that the expression function converts



Note:

`unescape`

external

The `external` function passes a value to an external rewriting program and returns the result.

Each invocation of `external` results in a single newline-terminated line being written to the external rewriting program's `stdin`. For each line of input, the program must produce a single line of output. When developing an external rewriting program, it is important to avoid buffering `stdout`. In Perl, for example, `$| = 1`; is used to disable buffering. Because the server expects the external rewriting program to produce one line of output for each line of input, the server can stop responding if the external rewriting program buffers its output.

Syntax

```
external (program, value)
```

Arguments

The following table shows the arguments for the `external` function.

Argument	Description
<i>program</i>	The program argument is the file name of an external rewriting program. Because <code>program</code> is executed using the operating system's default shell (<code>/bin/sh</code> on Unix/Linux) or the command interpreter (<code>CMD.EXE</code> on Windows), <code>program</code> should be an absolute path or the name of a program in the operating system's <code>PATH</code> . The server starts the external rewriting program on demand. A given server process never executes more than one instance of the program at a time.
<i>value</i>	The value passed to the rewriting program.

Example

The following is an example of an external rewriting program `rewrite.pl`, used to change the prefix `/home/` to `/u/`:

```
#!/usr/bin/perl

$| = 1;

while (<STDIN>) {
    s|^/home/|/u/|;
    print $_;
}
```

In this example, the external expression function used to invoke `rewrite.pl` is as follows:

```
NameTrans fn="rewrite" path="$(external('rewrite.pl', $path))"
```

httpdate

The `httpdate` function returns an RFC 1123 date time stamp for use in HTTP header fields such as `Expires`.

Syntax

```
httpdate (time)
```

Argument

The following table describes the argument for the `httpdate` function.

Argument	Description
<i>time</i>	The time value.

Example

The following `obj.conf` code could be used to set an `Expires` header that indicates a response is not cached for more than one day:

```
ObjectType fn="set-variable"
    insert-srvhdrs="$(httpdate($time + 86400))"
```

lc

The `lc` function converts all the US ASCII characters in the string to lowercase and returns the result.

Syntax

```
lc(string)
```

Argument

The following table describes the argument for the `lc` function.

Argument	Description
<i>string</i>	The string the expression function converts to lowercase.

Example

The following `obj.conf` code can be used to redirect clients, who erroneously used uppercase characters in the request URI to the equivalent lowercase URI:

```
<If code == 404 and not -e path and -e lc(path)>
```

```
Error fn="redirect" uri="$(lc($uri))"
```

```
</If>e($time + 86400))"
```



Note:

UC

length

The `length` function returns the length of its argument, that is, a number representing the length of the string.

Syntax

```
length (string)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>string</i>	The string for which the expression function computes the length.

Example

The following `obj.conf` code can be used to send a **404 Not found** error to clients that request URIs longer than 255 bytes:

```
<If length($uri) > 255>
```

```
PathCheck fn="deny-existence"
```

```
</If>
```

lookup

The `lookup` function inspects a text file for a name-value pair with name *name* and returns the corresponding value. The name-value pairs in the file are separated by white space.

If the file does not contain a name-value pair with the specified name, this function returns the value of *defaultvalue*, if specified, or returns an empty string.

Syntax

```
lookup(filename, name, defaultvalue)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>filename</i>	Indicates the name of a text file that contains one name-value pair per line. <i>filename</i> can be an absolute path or a path relative to the server's <code>config</code> directory. Names and values are separated by white space. Lines beginning with a pound sign (#) are ignored.
<i>name</i>	The name of the name-value pair for which the function looks in the text file.
<i>defaultvalue</i>	The value returned by the function if <i>filename</i> exists but does not contain a name-value pair with a name matching the value of <i>name</i> . If <i>defaultvalue</i> is not specified, it defaults to an empty string.

Example

The following example shows a text file called `urimap.conf` that could be used with the `lookup` function to map shortcut URIs to URIs:

```
# This file contains URI mappings for Oracle Traffic Director.

# Lines beginning with # are treated as comments.

# All other lines consist of a shortcut URI, whitespace, and canonical URI.

/webserver /software/products/web_srvr/home_web_srvr.html

/solaris   /software/solaris/

/java      /software/java/
```

Using the previous sample text file, you could use the following lookup expression to implement shortcut URIs for commonly accessed resources:

```
<If lookup('urimap.conf', uri)>
NameTrans fn="redirect" url="$(lookup('urimap.conf', uri))"
</If>
```

lookupregex

The `lookupregex` function inspects a text file for a regular expression-value pair. This function takes a string as an input and matches it with the regular expression in each line. It returns the corresponding value only if it matches. If the file does not contain a match, this function returns the default value, if specified, or returns an empty string.

Syntax

```
lookupregex (filename,string,defaultvalue)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>filename</i>	The filename is the name of a text file that contains one regular expression-value pair per line. filename can be an absolute path or a path relative to the server's configuration directory.
<i>string</i>	The string to match with every regular expression in the text file.
<i>defaultvalue</i>	The value returned by the function if the filename exists but does not contain a matching regular expression-value pair. If defaultvalue is not specified, it defaults to an empty string.

mtime

The `mtime` function returns the time of the last data modification for the specified file or directory.


Syntax

```
mtime(path)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>path</i>	The absolute path to the directory or file name for which you are requesting the last data modification.

 **Note:**
`atime`, `ctime`

owner

The `owner` function returns the owner of a file.

Syntax

```
owner(path)
```

Arguments

The following table describes the argument for the expression function.

Argument	Description
<i>path</i>	The absolute path to the directory or file name for which you are requesting the last data modification.

UC

The `uc` function converts all the US ASCII characters in string to uppercase and returns the result.

Syntax

```
uc(string)
```

Arguments

The following table describes the argument for the expression function.

Table A-1 `uc` Argument

Argument	Description
string	The string that the expression function converts to uppercase.

 **Note:**

`lc`

unescape

The `unescape` function decodes the URI using `util_uri_unescape`, converting percent-encoded octets to their unencoded form, and returns the result.

Syntax

```
unescape(uri)
```

Arguments

The following table describes the argument for the expression function.

Table A-2 `unescape` Argument

Argument	Description
<i>uri</i>	The URI that the function converts.

 **Note:**

`escape`

uuid

The `uuid` function returns a UUID as a string. No two calls to `uuid` return the same UUID. Because they are guaranteed to be unique, UUIDs are useful for constructing client-specific cookie values.

Syntax

```
uuid()
```

Regular Expressions

The `If` and `ElseIf` expressions can evaluate regular expressions using the equal sign and tilde (`=~`) and exclamation point and tilde (`!~`) regular expression matching operators. These regular expressions use the Perl-compatible syntax implemented by Perl-compatible Regular Expressions (PCRE).

By default, regular expressions are case sensitive. The `(?i)` option flag can be added to the beginning of a regular expression to request case insensitivity, for example:

```
$uri =~ '^[Ff][Ii][Ll][Ee][Nn][Aa][Mm][Ee]$'
```

```
$uri =~ '(?i)^/filename$'
```

When an `If` or `ElseIf` expression contains a regular expression, regular expression backreferences can appear within arguments in the container body. Regular expression backreferences are of the form `$n` where `n` is an integer between 1 and 9 corresponding to the capturing subpattern of the regular expression, for example:

```
<If $path =~ '^(.*)\.html|\.htm)$'>
NameTrans fn="rewrite" path="$1.shtml"
</If>
```

In the preview example, two subpatterns are used in the `If` expression, so `$1` and `$2` can be used as backreferences. In the example, the value of the first capturing subpattern is used within a `NameTrans fn="rewrite"` parameter. The value of the second capturing subpattern is ignored.

An `If` or `ElseIf` expression can contain backreferences to earlier regular expressions in that same `If` or `ElseIf` expression, for example:

```
<If "foo" =~ "(.*)" and $1 eq "foo">
# Any contained directives will be executed
# since $1 will evaluate to "foo"
...
</If>
```

The contents of the preview `If` expression are executed, because the given `If` expression always evaluates to true.

However, `If` and `Elseif` expressions, and contained directives, cannot contain backreferences to regular expressions in parent containers. For example, the following `obj.conf` entry is invalid:

```
<If $path =~ '(.*)\.css'>
<If $browser = "*MSIE*">
# This example is invalid as $1 is not defined
AuthTrans fn="rewrite" path="$1-msie.css"
```

```
</If>  
</If>
```

You can use the dollar sign and ampersand `$&` to obtain the value that last successfully matched a regular expression. Use the following `obj.conf` entry to redirect requests for HTML files to another server:

```
<If $path =~ '\.html$' or $path =~ '\.htm$' >  
NameTrans fn="redirect" url="http://docs.example.com$&"  
</If>
```

String Interpolation

Strings that contain references to variables or expressions are called interpolated strings. When you use interpolated strings, the embedded expressions and variables are evaluated, and the result is inserted into the string. The act of inserting data into a string is called string interpolation.

Use interpolated strings in expressions, log formats, and `obj.conf` parameters. In expressions, only string literals enclosed in double quotation marks are interpolated. See [Expression Literals](#).

Using Variables in Interpolated Strings

To include the value of a variable in a string, prefix the name of the variable with the dollar-sign (`$`). For example, the following `format` element in `server.xml` logs the client IP address, requested URI, and corresponding file system path for each HTTP request:

```
<access-log>  
  <file>access</file>  
  <format>$ip "$uri" $path</format>  
</access-log>
```

In this example, `$ip`, `$uri`, and `$path` are predefined variables. See [Variables](#).

For more information about access logs and log format, see [Using the Custom Access-Log File Format](#).

If the name of the variable is ambiguous, add braces `{ }` to the name. For example, the following string contains a reference to the predefined `$path` variable:

```
"${path}html"
```

Without the braces, the string contains a reference to a hypothetical variable named `pathhtml`.

Using Expressions in Interpolated Strings

To include the result of an expression in a string, prefix the expression with a dollar sign and a left parenthesis (`$(` and follow it with a right parenthesis `)`. For example, the following two strings are identical after interpolation:

```
"$(2 + 2)"  
"4"
```

When an interpolated string is used as an `obj.conf` parameter, the string is interpolated each time the corresponding instruction is executed. For example, the following lines could be used in the `obj.conf` file to redirect clients based on the requested URI and the contents of the file `redirect.conf`:

```
<Object ppath="/redirect/*">
NameTrans fn="redirect" url="$(lookup('redirect.conf', $uri, '/'))"
</Object>
```

In this example, the expression `lookup('redirect.conf', $uri, '/')` is evaluated each time the `NameTrans` directive is invoked, and the result is passed to the `redirect` SAF in its `url` parameter.

Wildcard Patterns

Oracle Traffic Director supports wildcard pattern matching in expressions. To use a wildcard without any special meaning, precede it with a backslash (`\`).

The following table describes various wildcard patterns and their uses.

Wildcard	Use
*	Matches zero or more characters.
?	Matches one occurrence of any character.
	An <code>or</code> expression. The substrings used with this operator can contain other special characters such as an asterisk <code>*</code> or dollar sign <code>\$</code> . The substrings must be enclosed in parentheses, for example, <code>(a b c)</code> , but the parentheses cannot be nested.
\$	Matches the end of the string. This is useful in <code>or</code> expressions.
[abc]	Matches one occurrence of the characters <code>a</code> , <code>b</code> , or <code>c</code> . Within these expressions, the only character that must be treated as a special character is the right bracket <code>]</code> ; all others are not special.
[a-z]	Matches one occurrence of a character between <code>a</code> and <code>z</code> .
[^az]	Matches any character except <code>a</code> or <code>z</code> .
*~	This expression, followed by another expression, removes any pattern matching the second expression.

The following table lists wildcard examples with pattern and the result.

Wildcard	Result
*.example.com	Matches any string ending with the characters <code>.example.com</code> .
(quark energy).example.com	Matches either <code>quark.example.com</code> or <code>energy.example.com</code> .
198.93.9[23].???	Matches a numeric string starting with either <code>198.93.92</code> or <code>198.93.93</code> and ending with any 3 characters.
.	Matches any string with a period in it.
~example-	Matches any string except those starting with <code>example-</code> .

Wildcard	Result
*.example.com~quark.example.com	Matches any host from domain example.com except for a single host quark.example.com.
*.example.com~(quark energy neutrino).example.com	Matches any host from domain .example.com except for hosts quark.example.com, energy.example.com, and neutrino.example.com.
.com~.example.com	Matches any host from domain .com except for hosts from sub domain example.com.
~.gif*	Matches any string except those including gif.

B

Using the Custom Access-Log File Format

This appendix contains information about the log format used by Oracle® Fusion Middleware. Use these format options to customize the format of your log files. You can enter them through the Admin Console, or edit the `format` subelement of the `access-log` element in the `server.xml` file.

You can use variables and expressions in log formats with the syntax `$(variable)` and `$(expression)`.

When creating a custom log format, anything between percent signs (%) is recognized as the name portion of a name-value pair stored in a parameter block in the server. Any additional text is treated as literal text, so you can add to the line to make it more readable. The one exception to the percent sign rule is the `%SYSDATE%` component, which delivers the current system date. `%SYSDATE%` is formatted using the time format `%d/%b/%Y:%H:%M:%S` and the offset from GMT.

If no format parameter is specified for a log file, the common log format is used:

```
"%Ses->client.ip% - %Req->vars.auth-user% [%SYSDATE%]  
\ "%Req->reqpb.clf-request%" %Req->srvhdrs.clf-status%  
%Req->srvhdrs.content-length%"
```

Typical components of log file format are listed in the following table. Because certain components could resolve to values that contain spaces, they are enclosed in escape quotes (\").

Option	Component
Client host name (unless <code>iponly</code> is specified in <code>flex-log</code> or DNS name is not available) or IP address	<code>%Ses->client.ip%</code>
Client DNS name	<code>%Ses->client.dns%</code>
System date	<code>%SYSDATE%</code>
Full HTTP request line	<code>\ "%Req->reqpb.clf-request%"</code>
Status	<code>%Req->srvhdrs.clf-status%</code>
Response content length	<code>%Req->srvhdrs.content-length%</code>
Response content type	<code>%Req->srvhdrs.content-type%</code>
Referer header	<code>\ "%Req->headers.referer%"</code>
User-agent header	<code>\ "%Req->headers.user-agent%"</code>
HTTP method	<code>%Req->reqpb.method%</code>
HTTP URI	<code>%Req->reqpb.uri%</code>
HTTP query string	<code>%Req->reqpb.query%</code>
HTTP protocol version	<code>%Req->reqpb.protocol%</code>
Accept header	<code>%Req->headers.accept%</code>

Option	Component
Date header	%Req->headers.date%
If-Modified-Since header	%Req->headers.if-modified-since%
Authorization header	%Req->headers.authorization%
Any header value	%Req->headers.headername%
Name of authorized user	%Req->vars.auth-user%
Value of a cookie	%Req->headers.cookie.name%
Value of any variable in Req->vars	%Req->vars.varname%
Virtual server ID	%vsid%
ECID	%Req->vars.ecid%. For more information, refer the elements of http .
RID	%Req->vars.rid%. For more information, refer the elements of http .
Duration	%duration% Records the time in microseconds the server spent handling the request. Statistics must be enabled before %duration% can be used.
System time	%Time% System time in seconds since 00:00:00 UTC, January 1, 1970.
Relative time	%RELATIVETIME% System time in seconds since logging started.
Method number	%Req->method_num% A number representing the HTTP method as used in NSAPI.
HTTP Protocol Version	%Req->protv_num% A number representing the HTTP protocol version as used in NSAPI.
HTTP request line	%Req->reqpb.clf-request.method% The method from the HTTP request line.
HTTP URI	%Req->reqpb.clf-request.uri% The URI from the HTTP request line.
URI path	%Req->reqpb.clf-request.uri.abs_path% The absolute path component of the URI
URI query	%Req->reqpb.clf-request.uri.query% The query component of the URI.
user_dn	%Ses->client.user_dn% The SSL client certificate authentication for web security.
HTTP protocol	%Req->reqpb.clf-request.protocol% The protocol from the HTTP request line.
Protocol name	%Req->reqpb.clf-request.protocol.name% The name of the protocol.

Option	Component
Protocol version	<code>%Req->reqpb.clf-request.protocol.version%</code> The version of the protocol.
Origin server	<code>%Req->vars.origin-server%</code> The origin server that served the request.
Cache Finish Statuses in Access Logs	<code>%Req->vars.proxy-cache-finish-status%</code> Cache status of the origin servers response in the proxy server

Additional log file parameters that can be configured are listed in the table below.

Option	Component
Cipher name	<code>%Ses->client.cipher%</code>
Size in bits of cipher key	<code>%Ses->client.keysize%</code>
Size in bits of private key	<code>%Ses->client.secret-keysize%</code>
DN for certificate issuer	<code>%Ses->client.issuer_dn%</code>
DN for certificate user	<code>%Ses->client.user_dn%</code>
SSL session identifier	<code>%Ses->client.ssl-id%</code>

C

Using Time Formats

This appendix describes the format strings used for dates and times in the server log. These formats are used by the NSAPI function `util_strftime`, by some built-in SAFs. The formats are similar to those used by the `strftime` C library routine, but not identical.

Table C-1 Format Strings for Date and Time

Attribute	Allowed Values
%a	Abbreviated day of the week (3 chars)
%d	Day of month as a decimal number (01-31)
%S	Second as a decimal number (00-59)
%M	Minute as a decimal number (00-59)
%H	Hour in 24-hour format (00-23)
%Y	Year with century, as a decimal number, up to 2099
%b	Abbreviated month name (3 chars)
%h	Abbreviated month name (3 chars)
%T	Time in HH:MM:SS format
%X	Time in HH:MM:SS format
%A	Day of the week, full name
%B	Month, full name
%C	"%a %b %e %H:%M:%S %Y"
%c	Date and time in "%m/%d/%y %H:%M:%S" format
%D	Date in "%m/%d/%y" format
%e	Day of month as decimal number (1-31) without leading zeros
%l	Hour in 12-hour format (01-12)
%j	Day of year as a decimal number (001-366)
%k	Hour in 24-hour format (0-23) without leading zeros
%l	Hour in 12-hour format (1-12) without leading zeros
%m	Month as a decimal number (01-12)
%n	Line feed
%p	a.m./p.m. indicator for 12-hour clock
%R	Time in "%H:%M" format
%r	Time in "%l:%M:%S %p" format
%t	Tab
%U	Week of year as a decimal number, with Sunday as first day of week (00-51)
%w	Weekday as a decimal number (0-6; Sunday is 0).

Table C-1 (Cont.) Format Strings for Date and Time

Attribute	Allowed Values
%W	Week of year as decimal number, with Monday as first day of week (00-51)
%x	Date in "%m/%d/%y" format
%y	Year without century, as decimal number (00-99)
%%	Percent sign

D

Alphabetical List of Server Configuration Elements and Predefined SAFs

This appendix provides an alphabetical list of server configuration elements, including `server.xml` elements and predefined SAFs in `obj.conf` file.

A

- [access-log](#)
- [access-log-buffer](#)
- [assign-name](#)

B

- [block-auth-cert](#)
- [block-cache-info](#)
- [block-cipher](#)
- [block-ip](#)
- [block-issuer-dn](#)
- [block-jroute](#)
- [block-keysize](#)
- [block-proxy-agent](#)
- [block-secret-keysize](#)
- [block-ssl](#)
- [block-ssl-id](#)
- [block-user-dn](#)
- [block-via](#)
- [block-xforwarded-for](#)

C

- [check-request-limits](#)

D

- [deny-existence](#)
- [dns](#)
- [dns-cache](#)

E

- [event](#)

F

- [failover-group](#)
- [flex-log](#)
- [forward-auth-cert](#)
- [forward-cache-info](#)
- [forward-cipher](#)
- [forward-ip](#)
- [forward-issuer-dn](#)
- [forward-jroute](#)
- [forward-keysize](#)

forward-proxy-agent
forward-secret-keysize
forward-ssl-id
forward-user-dn
forward-via
forward-xforwarded-for

G

get-client-cert
get-sslid

H

health-check
http
http-client-config
http-listener

K

keep-alive

L

localization
log

M

map

N

nt-uri-clean

O

origin-server-pool
origin-server

P

property
proxy-cache
proxy-cache-config
proxy-cache-override-http
proxy-retrieve

Q

qos-error
qos-handler

R

remove-filter
restart
reverse-map
rewrite

S

server
service-proxy-cache-dump
service-trace
set-cache-control

set-cookie
set-origin-server
set-priority
set-proxy-server
set-variable
snmp
ssl
ssl-logout
ssl-session-cache
stats
stats-xml
strip-params

T

thread-pool
time
type-by-exp
type-by-extension

U

unix-uri-clean

V

variable
virtual-server

W

webapp-firewall-ruleset

Index

A

access-log element, [3-2](#)
access-log-buffer element, [3-3](#)
AddLog, [4-2](#)

- flow of control, [4-8](#)
- function descriptions, [5-43](#)

admin-server directory, [1-3](#)
all-requests bucket, [5-2](#)
always-use-keep-alive parameter, [5-26](#)
assign-name function, [5-5](#)
AuthTrans, [4-1](#)

- flow of control, [4-5](#)
- function descriptions, [5-2](#)

B

backreferences, [A-18](#)
bin directory, [1-3](#)
block-auth-cert function, [5-16](#)
block-cache-info function, [5-16](#)
block-cipher function, [5-16](#)
block-ip function, [5-17](#)
block-issuer-dn function, [5-17](#)
block-jroute function, [5-17](#)
block-keysize function, [5-17](#)
block-proxy-agent function, [5-18](#)
block-secret-keysize function, [5-18](#)
block-ssl-id function, [5-18](#), [5-19](#)
block-user-dn function, [5-19](#)
block-via function, [5-19](#)
bong-file parameter, [5-12](#)
boolean

- expression results, [A-6](#)

bucket

- all request, [5-2](#)
- default, [5-1](#)

built-in SAFs in obj.conf, [5-1](#)

C

cache

- DNS, [3-5](#)

cache control directives, [5-32](#)
case sensitivity in obj.conf, [4-9](#)

charset parameter, [5-33](#)
check-request-limits function, [5-10](#)
choose function, [A-11](#)
ChunkedRequestBufferSize parameter, [5-38](#)
ChunkedRequestTimeout parameter, [5-38](#)
Client tag, [4-3](#)
cluster element, [3-3](#), [3-4](#)
code parameter, [5-45](#), [5-46](#)
comments

- in obj.conf, [4-10](#)

config directory, [1-4](#)
control parameter, [5-27](#), [5-29–5-32](#)
core SAFs in obj.conf, [5-1](#)
ctime function, [A-11](#)
custom log file format, [B-1](#)
custom variables, [A-4](#)

D

day of month, [C-1](#)
default object, [4-3](#)
default-bucket, [5-1](#)
deny-existence function, [5-12](#)
directives

- obj.conf, [5-1](#)
- order of in obj.conf, [4-9](#)
- syntax in obj.conf, [4-2](#)

directory structure, [1-2](#)
dns element, [3-4](#)
dns-cache element, [3-5](#)
domain parameter, [5-33](#)
dorequest parameter, [5-13](#)
dotdirok parameter, [5-14](#)
dynamic reconfiguration, [1-4](#)

E

elements in the server.xml file, [3-1](#)
Else tag, [4-5](#)
Elsif tag, [4-5](#)

- with regular expressions, [A-18](#)

enc parameter, [5-33](#)
Error directive, [4-2](#)

- flow of control, [4-8](#)
- function descriptions, [5-44](#)

error parameter, [5-11](#)
 errors
 sending customized messages, [5-45](#), [5-46](#)
 escape function, [A-12](#)
 escape parameter, [5-49](#)
 event element, [3-5](#)
 exp parameter, [5-33](#)
 expressions, [A-5](#)
 functions, [A-10](#)
 in interpolated strings, [A-19](#)
 literals, [A-6](#)
 operators, [A-8](#)
 regular, [A-18](#)
 results as Booleans, [A-6](#)
 syntax, [A-5](#)
 variables, [A-7](#)
 external function, [A-12](#)

F

failover-group element, [3-8](#)
 failover-instance element, [3-9](#)
 file-cache element, [3-7](#)
 filter parameter, [5-48](#), [5-50](#)
 find-pathinfo-forward parameter, [5-5](#)
 flex-log function, [5-44](#)
 flow of control in obj.conf, [4-5](#)
 FlushTimer parameter, [5-38](#)
 fn parameter
 in directives in obj.conf, [4-2](#)
 forward slashes, [4-10](#)
 forward-auth-cert function, [5-20](#)
 forward-cache-info function, [5-20](#)
 forward-cipher function, [5-20](#)
 forward-ip function, [5-21](#)
 forward-issue-dn function, [5-21](#)
 forward-jroute function, [5-22](#)
 forward-keysize function, [5-22](#)
 forward-proxy-agent function, [5-22](#)
 forward-secret-keysize function, [5-23](#)
 forward-ssl-id function, [5-23](#), [5-24](#)
 forward-user-dn function, [5-24](#)
 forward-via function, [5-24](#), [5-25](#)
 from parameter, [5-3](#), [5-5](#), [5-8](#), [5-49](#), [5-51](#)–[5-53](#)
 functions
 common, [5-46](#)
 expression, [A-10](#)

G

get-client-cert function, [5-12](#)
 get-sslid function, [5-2](#)

H

health-check element, [3-10](#)
 http element, [3-13](#)
 http-client-config function, [5-25](#)
 http-listener element, [3-15](#)
 httpdate function, [A-13](#)

I

If tag, [4-5](#)
 with regular expressions, [A-18](#)
 Input, [4-2](#)
 flow of control, [4-7](#)
 function descriptions, [5-34](#)
 optional parameters, [5-34](#)
 insert-filter function, [5-47](#)
 with Input directive, [4-7](#)
 with Output directive, [4-7](#)
 internal requests, [4-8](#)
 interpolated strings, [A-19](#)
 interval parameter, [5-10](#)
 iponly parameter, [5-44](#)

K

keep-alive element, [3-17](#)
 keep-alive parameter, [5-26](#)
 keep-alive-timeout parameter, [5-26](#)

L

lang parameter, [5-33](#)
 lc function, [A-13](#)
 length function, [A-14](#)
 lib directory, [1-3](#)
 line continuation
 in obj.conf, [4-10](#)
 literals
 expression, [A-6](#)
 numeric, [A-7](#)
 string, [A-6](#)
 localization element, [3-18](#)
 log analyzer, [5-44](#)
 log element, [3-18](#)
 log file
 analyzer for, [5-44](#)
 log file format, [B-1](#)
 lookup function, [A-14](#)

M

maintenance element, [3-19](#)
 map function, [5-6](#)

match-browser function, [5-48](#), [5-53](#)
 max-age parameter, [5-33](#)
 max-connections parameter, [5-10](#)
 max-rps parameter, [5-10](#)
 method parameter, [5-13](#), [5-35](#), [5-38](#)
 monitor parameter, [5-11](#)
 month name, [C-1](#)
 mtime function, [A-16](#)

N

name attribute
 in obj.conf objects, [4-3](#)
 in objects, [4-3](#)
 name parameter, [5-4](#), [5-5](#), [5-44](#)
 NameTrans, [4-1](#)
 flow of control, [4-6](#)
 function descriptions, [5-4](#)
 nondefault objects
 processing, [4-6](#)
 nostat parameter, [5-6](#)
 nt-uri-clean function, [5-13](#)
 ntrans-base, [5-5](#), [5-6](#)
 numeric literals, [A-7](#)

O

obj.conf
 case sensitivity, [4-9](#)
 Client tag, [4-3](#)
 comments, [4-10](#)
 directive syntax, [4-2](#)
 directives, [4-2](#), [5-1](#)
 Else tag, [4-5](#)
 Elseif tag, [4-5](#)
 flow of control, [4-5](#)
 function flow changes, [4-8](#)
 If tag, [4-5](#)
 Object tag, [4-3](#)
 order of directives, [4-9](#)
 overview, [4-1](#)
 parameters for directives, [4-9](#)
 processing other objects, [4-6](#)
 syntax rules, [4-9](#)
 Object tag, [4-3](#)
 name attribute, [4-3](#)
 ppath attribute, [4-3](#)
 objects
 processing nondefault objects, [4-6](#)
 ObjectType, [4-2](#)
 flow of control, [4-6](#)
 function descriptions, [5-14](#)
 operators
 expression, [A-8](#)

order
 of directives in obj.conf, [4-9](#)
 origin-server element, [3-22](#)
 origin-server-pool element, [3-20](#)
 Output, [4-2](#)
 flow of control, [4-7](#)
 function descriptions, [5-35](#)
 optional parameters, [5-35](#)
 owner function, [A-16](#)

P

parameters
 for obj.conf directives, [4-9](#)
 path names
 in obj.conf, [4-10](#)
 path parameter, [5-10](#), [5-12](#)
 PathCheck, [4-2](#)
 flow of control, [4-6](#)
 function descriptions, [5-9](#)
 ppath attribute
 in obj.conf objects, [4-3](#)
 in objects, [4-3](#)
 predefined SAFs in obj.conf, [5-1](#)
 predefined variables, [A-2](#)
 processing nondefault objects, [4-6](#)
 property element, [3-24](#)
 protocol parameter, [5-26](#)
 proxy-agent parameter, [5-26](#)
 proxy-retrieve function, [5-39](#)
 proxy-server element, [3-25](#)

Q

qos-error function, [5-45](#)
 qos-handler function, [5-2](#), [5-3](#)
 query parameter, [5-35](#), [5-38](#)
 quotes
 in obj.conf, [4-9](#)

R

reconfig, [1-4](#)
 redirect function, [5-49](#)
 regular expressions, [A-18](#)
 remove-dir function, [5-40](#)
 remove-filter function, [5-50](#)
 request-handling process, [4-1](#)
 flow of control, [4-5](#)
 requests
 internal, [4-8](#)
 restarted, [4-8](#)
 require parameter, [5-13](#)
 restart function, [5-51](#)–[5-53](#)
 restarted requests, [4-8](#)

retries parameter, [5-26](#)
 retry-unsafe-requests parameter, [5-27](#)
 reverse-map function, [5-7](#)
 rewrite function, [5-8](#)
 rewrite-content-location parameter, [5-7](#), [5-37](#)
 rewrite-headername parameter, [5-8](#), [5-37](#)
 rewrite-host parameter, [5-7](#), [5-37](#)
 rewrite-location parameter, [5-7](#), [5-37](#)
 root element, [3-26](#)
 Route, [4-2](#)
 flow of control, [4-7](#)
 function descriptions, [5-35](#)
 route-cookie parameter, [5-36](#)
 route-hdr parameter, [5-36](#)
 rules
 for editing obj.conf, [4-9](#)

S

SAFs
 predefined in obj.conf, [5-1](#)
 separators
 in obj.conf, [4-9](#)
 server
 flow of control, [4-5](#)
 instructions in obj.conf, [4-2](#)
 processing nondefault objects, [4-6](#)
 server element, [3-26](#)
 server instance directory, [1-4](#)
 server.xml, [2-1](#)
 editing, [2-1](#)
 elements, [3-1](#)
 overview, [2-1](#)
 sample, [2-4](#)
 validating, [2-2](#)
 Service, [4-2](#)
 flow of control, [4-8](#)
 function descriptions, [5-38](#)
 optional parameters, [5-38](#)
 service-trace function, [5-42](#)
 set-cache-control function, [5-32](#)
 set-cookie function, [5-32](#)
 set-origin-server function, [5-36](#)
 set-proxy-server function, [5-37](#)
 set-variable function, [5-54](#)
 snmp element, [3-28](#)
 spaces
 in obj.conf, [4-9](#)
 ssl element, [3-28](#)
 ssl-logout function, [5-14](#)
 ssl-session-cache element, [3-33](#)
 ssl3-tls-ciphers element, [3-30](#)
 stats element, [3-33](#)
 stats-xml function, [5-42](#)

sticky-cookie parameter, [5-36](#)
 sticky-param parameter, [5-36](#)
 string interpolation, [A-19](#)
 string literals, [A-6](#)
 strip-params function, [5-8](#), [5-9](#)
 syntax
 directives in obj.conf, [4-2](#)
 expressions, [A-5](#)
 for editing obj.conf, [4-9](#)

T

tags
 Client, [4-3](#)
 Else, [4-5](#)
 Elseif, [4-5](#)
 If, [4-5](#)
 Object, [4-3](#)
 thread-pool element, [3-34–3-38](#)
 tildeok parameter, [5-14](#)
 time element, [3-39](#)
 time format strings, [C-1](#)
 type parameter, [5-33](#), [5-35](#), [5-38](#)
 type-by-exp function, [5-33](#)
 type-by-extension function, [5-34](#)

U

uc function, [A-17](#)
 unix-uri-clean function, [5-14](#)
 uri parameter, [5-51–5-53](#)
 URI translation, [4-9](#)
 url parameter, [5-49](#)
 url-prefix parameter, [5-49](#)
 UseOutputStreamSize parameter, [5-38](#)
 util_strftime, [C-1](#)
 uuid function, [A-18](#)

V

variable element, [3-39](#)
 variables, [A-2](#)
 custom, [A-4](#)
 expression, [A-7](#)
 in interpolated strings, [A-19](#)
 predefined, [A-2](#)
 resolving, [A-4](#)
 supported by set-variable, [5-55](#)
 virtual-server element, [3-40](#), [3-41](#)

W

weekday, [C-1](#)