



Apache Web Server Considerations (UNIX®/Linux/macOS)

Version 2024.1
2024-05-16

Apache Web Server Considerations (UNIX®/Linux/macOS)
InterSystems IRIS Data Platform Version 2024.1 2024-05-16
Copyright © 2024 InterSystems Corporation
All rights reserved.

InterSystems®, HealthShare Care Community®, HealthShare Unified Care Record®, IntegratedML®, InterSystems Caché®, InterSystems Ensemble®, InterSystems HealthShare®, InterSystems IRIS®, and TrakCare are registered trademarks of InterSystems Corporation. HealthShare® CMS Solution Pack™ HealthShare® Health Connect Cloud™, InterSystems IRIS for Health™, InterSystems Supply Chain Orchestrator™, and InterSystems TotalView™ For Asset Management are trademarks of InterSystems Corporation. TrakCare is a registered trademark in Australia and the European Union.

All other brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

This document contains trade secret and confidential information which is the property of InterSystems Corporation, One Memorial Drive, Cambridge, MA 02142, or its affiliates, and is furnished for the sole purpose of the operation and maintenance of the products of InterSystems Corporation. No part of this publication is to be used for any other purpose, and this publication is not to be reproduced, copied, disclosed, transmitted, stored in a retrieval system or translated into any human or computer language, in any form, by any means, in whole or in part, without the express prior written consent of InterSystems Corporation.

The copying, use and disposition of this document and the software programs described herein is prohibited except to the limited extent set forth in the standard software license agreement(s) of InterSystems Corporation covering such programs and related documentation. InterSystems Corporation makes no representations and warranties concerning such software programs other than those set forth in such standard software license agreement(s). In addition, the liability of InterSystems Corporation for any losses or damages relating to or arising out of the use of such software programs is limited in the manner set forth in such standard software license agreement(s).

THE FOREGOING IS A GENERAL SUMMARY OF THE RESTRICTIONS AND LIMITATIONS IMPOSED BY INTERSYSTEMS CORPORATION ON THE USE OF, AND LIABILITY ARISING FROM, ITS COMPUTER SOFTWARE. FOR COMPLETE INFORMATION REFERENCE SHOULD BE MADE TO THE STANDARD SOFTWARE LICENSE AGREEMENT(S) OF INTERSYSTEMS CORPORATION, COPIES OF WHICH WILL BE MADE AVAILABLE UPON REQUEST.

InterSystems Corporation disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

For Support questions about any InterSystems products, contact:

InterSystems Worldwide Response Center (WRC)
Tel: +1-617-621-0700
Tel: +44 (0) 844 854 2917
Email: support@InterSystems.com

Table of Contents

Apache Web Server Considerations (UNIX®/Linux/macOS)	1
1 Security	1
2 Apache Process Management and the Web Gateway	1
2.1 Maximum Server Connections	2
3 State-Aware Sessions (Preserve mode 1)	2
3.1 Prevent Apache Process Recycling	3

Apache Web Server Considerations (UNIX®/Linux/macOS)

This page contains information about the [recommended option](#) for UNIX®, Linux, and macOS and atypical option 1 ([Alternative Option 1: Apache API Module with NSD \(mod_csp24.so\)](#)).

1 Security

When an Apache web server starts, it initializes a parent process that usually runs with superuser privileges. This is necessary in order to bind to TCP port 80.

After that, Apache launches the child processes (*worker processes*) which do the work of serving web requests. These child processes run as a less-privileged user and group, which you specify using the `User` and `Group` Apache configuration directives.

The child processes must be able to read all the content that they are responsible for serving (and have read/write access to the Web Gateway's configuration and Event Log files). Beyond this, however, these child processes should be granted as few privileges as possible. Refer to the Apache documentation for further information.

2 Apache Process Management and the Web Gateway

InterSystems Web Gateway modules are directly bound to Apache *worker processes*. Therefore, the way Apache is configured to manage its process pool has a direct effect on the Web Gateway.

Apache provides three Multi-Processing Modules (MPMs) for process management: Prefork MPM (<http://httpd.apache.org/docs/current/mod/prefork.html>), Worker MPM (<http://httpd.apache.org/docs/current/mod/worker.html>), or Event MPM (<http://httpd.apache.org/docs/current/mod/event.html>). Because the Web Gateway [dynamic shared object modules \(DSOs\)](#) are thread-safe, they can be deployed alongside any MPM.

In order to determine which of the server models is in use for an existing installation, issue the following command from a command prompt:

RHEL

```
httpd -v
```

Ubuntu/SUSE

```
apache2 -v
```

All MPMs involve spreading the load over multiple child (*worker*) processes. For all MPMs, the `StartServers` directive specifies the number of worker processes to start. Because Web Gateway modules are directly bound to worker processes, this directive also determines the number of Web Gateway instances in operation.

However, the running configuration, connection table and form cache is held in a shared memory sector. This allows the contents of the Web Gateway System Status form to consistently reflect the Web Gateway activity across the all worker

processes. For Apache servers, the connection table (and connection numbers) includes an additional column indicating the web server process ID to which each InterSystems IRIS connection is associated is included.

2.1 Maximum Server Connections

While the Web Gateway load is spread over multiple web server processes, the Maximum Server Connections configuration parameter sets a single overall limit on the number of connections the Web Gateway can make to a particular InterSystems IRIS server. This means that the number of worker processes started by Apache does not affect the maximum number of connections the Web Gateway can create.

For installations where most of the Apache workload is devoted to serving InterSystems web application requests, it is better to not assign a value to the Web Gateway's Maximum Server Connections directive and control the amount of concurrent work that can be done with the corresponding Apache configuration parameters instead. Setting an independent value for the Web Gateway's Maximum Server Connections directive would, however, make sense in installations where InterSystems file types represent only part of the workload for the Apache installation as a whole.

3 State-Aware Sessions (Preserve mode 1)

To provide support for state-aware sessions across multiple Apache worker processes, the Web Gateway routes requests between worker processes using UNIX® domain sockets.

As an example, consider a web server installation that distributes its load over 3 worker processes: P1, P2 and P3. Each worker process can potentially start any number of threads (T1, T2 ... Tn) according to the web server MPM and configuration in use.

Suppose an application makes a request to mark its session as state-aware (preserve mode 1) and the Web Gateway acknowledges this instruction in process P2. The connection and (security context) to the now private InterSystems IRIS process is hosted by web server worker process P2. All further requests for that user/session must now be processed by worker process P2. However, the Web Gateway has no control over which worker process the web server routes subsequent requests to, so the Web Gateway must establish an IPC channel between P2 and (potentially) any other worker process in the set.

When the Web Gateway marks the connection as state-aware in P2, it starts a listening service in a separate, detached, thread. For [log level v2](#), the Event Log would include a message similar to the following:

```
IPC Server
Process ID: 28457 Listening on Domain Socket: /tmp/csp28457.str
```

If another request for the same session is processed by worker process P3, the Web Gateway forwards that request to process P2 via the IPC channel previously established and then waits for the response. For [log level v2](#), the Event Log would include a message similar to the following:

```
Route request over IPC to another web server process
PrivateSession=2; pid_self=28456; ipc_to_pid=28457;
```

Of course, if the web server routes a request for this session to P2, then no further routing is necessary in the Web Gateway environment.

If the Web Gateway is unable to connect and forward a request to a previously created IPC channel, one of the following messages is recorded to the Event Log (depending on the context):

```
IPC CLIENT: Error
Cannot connect
```

Or:

```
IPC CLIENT: Error  
Cannot send request
```

This error may occur if Apache has closed or [recycled](#) a worker process. A worker process may also have crashed; in this case, the Apache error log would also record an error.

3.1 Prevent Apache Process Recycling

By default, Apache periodically recycles worker processes. Therefore, if you use state-aware sessions, configure Apache such that it doesn't recycle worker processes by configuring the installation as follows:

- Set the value of `MaxConnectionsPerChild` to zero
- Set the value of `MaxSpareThreads` to the same value as `MaxRequestWorkers`

If it is not possible to prevent Apache periodically recycling processes (perhaps as a result of a malfunctioning module) and state-aware sessions must be used, then an NSD based Gateway configuration can be used. An NSD-based architecture avoids the problems discussed above because it effectively separates the process management of the Web Gateway from the web server. Options for using the Web Gateway's network service daemon (NSD) are covered in [Using the NSD on Microsoft Windows](#) and [Using the NSD on UNIX®, Linux, and macOS](#).

