





# Performance Dashboard

 The Performance Dashboard was removed from XperienCentral in version R28. Use the [Administrative Pages](#) instead.

The XperienCentral Performance Dashboard provides detailed information about how the parts of XperienCentral's infrastructure are functioning. Because XperienCentral is a complex system, the interaction between many different sub-systems and configuration options has a direct impact on how it performs. Using the Performance Dashboard, system administrators, developers, and application managers can quickly identify problems and then make the necessary configuration changes to ensure that XperienCentral performs optimally. The Performance Dashboard measures the speed and response time of page requests and internal queries as well as other settings that affect how XperienCentral runs and rates the results according to the optimal expected results.

The status of how the various XperienCentral subsystems are performing is shown using an easily identifiable color-coded system:

| Performance Indicator   | Description   |
|---|---|
|  | The functionality/configuration setting is functioning within optimally expected parameters.  |
|  | The functionality/configuration setting is functioning below its optimal level.   |
|  | The functionality/configuration setting is having a pronounced, negative impact on XperienCentral's performance and should be immediately examined and addressed. |



- In a clustered environment, the statistics shown in the Performance Dashboard are only for the node/instance/server on which the Performance Dashboard is running, not the whole environment.

To open the Performance Dashboard, navigate to **Configuration > Performance Dashboard**.

## In This Topic

- [System Performance Indicators](#)
- [Page Response Time Indicators](#)
- [Page Cache Ratio Indicators](#)
- [Queries Indicators](#)
- [Improving System Performance](#)
- [Creating Custom Performance Indicator Tabs](#)

## System Performance Indicators

The System Performance Indicators tab shows performance information related to web page caching, web page response times, and the effect that important configuration settings are having on the overall performance of XperienCentral.

### Caching

For caching, the following indicators related to performance are provided. In the "Status" column, you can click the orange and red indicators for more information on the meaning of the indicator itself and what you can do in order to improve the performance for the related item.

| Indicator           | Description  |
|---------------------|--|
| Average cache ratio | Indicates the average percentage of all page requests that could be returned from the proxy cache, that is, the number of pages that were returned from the cache divided by the total number of page requests. The average page cache ratio is always a number between 0 and 1 (between 0% and 100%). |

|                   |   |
|-------------------|---|
| Timestamp caching | Specifies the timeout in milliseconds that timestamp caching is currently set to.   |
| Caching enabled   | <p>Specifies whether caching is enabled or disabled.</p> <p>The caching module of XperienCentral caches the web pages of the website. When a page is first requested, the web page is generated and stored in this cache. All subsequent requests for this page will be returned from the cache - this prevents the page from being generated each time it is requested. Web pages stay cached until the content is modified. The performance of a website is directly affected by the caching setting: If disabled, website visitors will experience much slower response times.</p> |

## Response Time

For response times, the following indicators related to performance are shown. In the "Status" column, you can click the orange and red indicators for more information on the meaning of the indicator itself and what you can do in order to improve the performance for the related item. The values listed are those that have been measured since the last time the list was cleared.

| Indicator                            | Description   |
|--------------------------------------|---|
| Largest response time                | The largest response time in milliseconds for all page requests across the entire web site. To view the response time for individual pages, see <a href="#">Page Response Time</a> .                        |
| Largest query execution time         | The largest query execution time is the execution time of the longest of all SQL queries performed on all databases. To view the query execution time for individual queries, see <a href="#">Queries</a> . |
| Largest response time Page Generator | The largest response time in milliseconds the Page Generator took to fulfill a request. To view the Page Generator response time for individual pages, see <a href="#">Page Response Time</a> .             |
| Average response time Page Generator | The average response time in milliseconds that the Page Generator took to fulfill requests. To view the Page Generator response time for individual pages, see <a href="#">Page Response Time</a> .         |
| Average response time                | The average response time in milliseconds for all page requests. To view the response time for individual pages, see <a href="#">Page Response Time</a> .   |
| Average query execution time         | The average query response time is the average for all SQL queries performed on all databases. To view the query execution time for individual queries, see <a href="#">Queries Indicators</a> .            |

## Settings

The settings tab shows the status of XperienCentral settings that can affect performance. In the "Status" column, you can click the orange and red indicators for more information on the meaning of the indicator itself and what you can do in order to improve the performance for the related item.

| Indicator         | Description   |
|-------------------|---|
| Debugging options | Specifies the status of the XperienCentral debug options. If any debug option is enabled, the status is red.  |
| Log level         | <p>Specifies the level of logging that is currently set for XperienCentral. Logging can be enabled for the following system functionality:</p> <ul style="list-style-type: none"> <li>• XperienCentral</li> <li>• Jackrabbit</li> <li>• Spring framework</li> <li>• Quartz</li> </ul> <p>What you log for each and how many of these functionalities you enable logging for can directly affect XperienCentral performance.</p> |

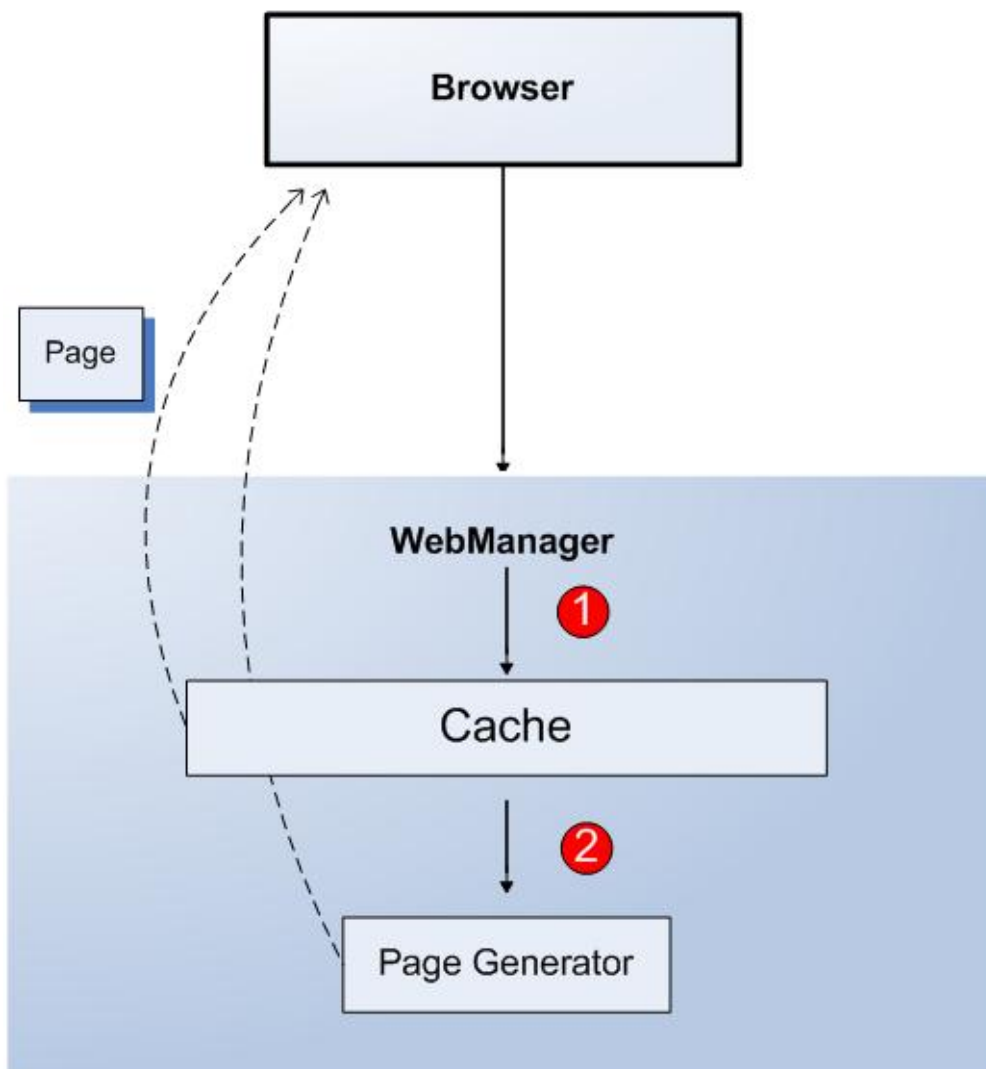
[Back to top](#)

## Page Response Time Indicators

The Page response tab shows a list of the slowest page response times in milliseconds for pages requested from the website. The page response tab lists two different types of pages:

- The slowest pages requested by website visitors that were not retrieved from the cache but had to be generated by the Page Generator
- The overall slowest pages requested by website visitors, both those returned from the cache as well as those generated by the Page Generator

When a page is requested and caching is enabled, XperienCentral determines whether the page can be returned from the cache (see 1 in the figure below) and returns it; if the page cannot be returned from the cache, it is generated by the Page Generator and returned (see 2 in the figure below):



To enable monitoring for web pages and/or pages generated by the Page Generator pages, select "Monitoring enabled".



#### Notes

- These options are not persisted — Monitoring is enabled for the current session only. If the Tomcat server is restarted, these checkboxes are cleared.
- Enabling page response time monitoring can affect XperienCentral performance. Monitoring should not be constantly enabled.

When these monitoring settings are enabled, the pages returned by the Page Generator that took the longest to generate as well as the pages that took the longest time to load are listed. The pages that appear in the list are those that have been read since monitoring was enabled or since the last time that the list was cleared using the [Clear list] button.


### Slowest Pages from the Page Generator

The pages listed in the "Slowest Pages from the Page Generator" list are pages requested from the website that were generated by the Page Generator and not taken from the cache. The pages in this list are a subset of the pages in the "Slowest pages" list and will always appear in the "Slowest Pages" list as well. See [Lowering Page Response Times](#) for troubleshooting information.

## Slowest Pages

The pages listed in the "Slowest Pages" list are the pages requested from the website that took the longest time to load from the cache and/or be generated by the Page Generator. The pages that took the longest time to be generated by the Page Generator are also listed separately in the "Slowest Pages Page Generator" list. See [Lowering Page Response Times](#) for troubleshooting information.

To view the contents of a page in the list in a separate window, click its URL in the list. This makes it easy for you to quickly view the contents of the slowest pages. The page will display in a separate window.

 Not all pages in the list are viewable due to context-related factors. If the page cannot be displayed, a JavaScript error will be returned.

To refresh the list, click [Refresh this page]. To clear the list, click [Clear list]. For information about steps you can take to improve the performance of slow pages, See [Lowering Page Response Times](#).

[Back to top](#)


---

## Page Cache Ratio Indicators

The Page cache ratio tab shows the pages with the smallest cache ratio, that is pages that are more often being regenerated instead of being returned from the cache. The higher the number, the better this is for performance. The page cache ratio is always a number between 0 and 1 (between 0% and 100%).

To enable page cache ratio monitoring, select "Monitoring enabled". The average cache ratio for the pages with the lowest cache ratio are listed. Ideally this number should always be 1 (100% of requests could be returned from the cache). The pages that appear in the list are those that have been read since monitoring was enabled or since the last time that the list was cleared using the [Clear list] button.

To view the contents of a the pages with the smallest cache ratios, click its URL in the list. This makes it easy for you to quickly view the contents of the pages. The page will display in a separate window.

 Not all pages in the list are viewable due to context-related factors. If the page cannot be displayed, a JavaScript error will be returned.

To refresh the list, click [Refresh this page]. To clear the list, click [Clear list]. For information on improving cache ratio, see [Lowering the Cache Ratio](#).

[Back to top](#)

---

## Queries Indicators

Use the Queries tab to monitor and examine the queries that have taken the most time (the slowest). To enable query monitoring, select "Monitoring enabled":

SLOWEST QUERIES

Monitoring enabled ☒



#### Notes

- The enabling of SQL query monitoring is not a persisted setting — Monitoring is enabled for the current session only. If the Tomcat server is restarted, the "Enable monitoring" checkbox is cleared.
- Enabling SQL query monitoring directly affects XperienCentral performance. Monitoring should not be constantly enabled.

## Query List

After query monitoring is enabled, the speed of all queries performed by XperienCentral will be monitored. In the "Slowest Queries" list, the slowest query monitored appears at the top of the list. The queries that appear in the list are those that have run since monitoring was enabled or since the last time that the list was cleared using the [Clear list] button.

### Performance Dashboard

[System performance indicators](#)[Page response time](#)[Page cache ratio](#)[Queries](#)

#### SLOWEST QUERIES

Monitoring enabled ☒

| Avg. exec. time (ms) | Query sql string  |
|----------------------|---|
| 52                   | call tssetobjecttimestamp(26111)  |
| 43                   | call wjaddwebcontent(?,?,?,?)   |
| 7                    | select * from wmdiscussion where for_contentid is null  |
| 5                    | insert into wmjellycontentactivityinstance(contentid,activity,state)values(?,?,?)                         |
| 3                    | update wmlocalrevision set revisionid=? where node=?  |
| 3                    | call tssetjellycontenttimestamp(?,?)  |
| 3                    | select wmjellywebconfigmodel.discussionlayout as forumsettings,wmdiscussion.id as forum from wmjelly...   |
| 2                    | insert into wmhhistory(timestamp,contentid,action,username,representation,usercomment,webid,contenttyp... |
| 1                    | delete from wmglobalindexeventqueue where revisionid <=?  |
| 1                    | select wmjellywebterm.id from wmjellywebcontentterm,wmjellywebterm where wmjellywebcontentterm.webte...   |
| 1                    | select wmjellywebreference.typeofpage,wmjellywebreference.dbid from wmjellywebcontentreference,wmjel...   |
| 0                    | select display_on from wmjellywebcontentdisplayon where content=?   |
| 0                    | select distinct max(revisionid)as revisionid,entityid,type from wmglobalindexeventqueue where revisi...   |
| 0                    | update clusterlock set heartbeat=? where id=?   |
| 0                    | insert into clusterlock(name,owner,id,created,heartbeat)values(?,?,?,?/?)                                 |
| 0                    | delete from clusterlock where id=?  |
| 0                    | select id from wmjellywebconfigmodel where webid=? order by name  |
| 0                    | select id,url_title,querystring,iscurrent,current_identifier from wmfriendlyurl where querystring li...   |
| 0                    | select id,url_title,url_identifier,entity_identifier,sequence,new_alt_path,iscurrent,querystring,cur...   |

[Help](#)[Close](#)

To see the full query, click it in the list. In the "Details" box the full query appears. From here you can also copy the query and paste it in another application:

## Performance Dashboard

| System performance indicators | Page response time | Page cache ratio | Queries  |
|-------------------------------|--------------------|------------------|--|
| 5                             |                    |                  | insert into wmjellycontentactivityinstance(contentid,activity,state)values(?,?,?)                          |
| 3                             |                    |                  | update wmlocalrevision set revisionid=? where node=?   |
| 3                             |                    |                  | call tssetjellycontenttimestamp(?,?)   |
| 3                             |                    |                  | select wmjellywebconfigmodel.discussionlayout as forumsettings,wmdiscussion.id as forum from wmjelly...    |
| 2                             |                    |                  | insert into wmhistory(timestamp,contentid,action,username,representation,usercomment,webid,contenttyp...   |
| 1                             |                    |                  | delete from wmglobalindexeventqueue where revisionid <=?   |
| 1                             |                    |                  | select wmjellywebterm.id from wmjellywebcontentterm,wmjellywebterm where<br>wmjellywebcontentterm.webte... |
| 1                             |                    |                  | select wmjellywebreference.typeofpage,wmjellywebreference.dbid from wmjellywebcontentreference,wmjel...    |
| 0                             |                    |                  | select display_on from wmjellywebcontentdisplayon where content=?  |
| 0                             |                    |                  | select distinct max(revisionid)as revisionid,entityid,type from wmglobalindexeventqueue where revisi...    |
| 0                             |                    |                  | update clusterlock set heartbeat=? where id=?  |
| 0                             |                    |                  | insert into clusterlock(name,owner,id,created,heartbeat)values(?,?,?,?/?)                                  |
| 0                             |                    |                  | delete from clusterlock where id=?   |
| 0                             |                    |                  | select id from wmjellywebconfigmodel where webid=? order by name   |
| 0                             |                    |                  | select id,url_title,querystring,iscurrent,current_identifier from wmfriendlyurl where querystring li...    |
| 0                             |                    |                  | select id,url_title,url_identifier,entity_identifier,sequence,new.alt_path,iscurrent,querystring,cur...    |
| 0                             |                    |                  | select obj_timestamp from tsoobjecttimestamp where sw_object=?   |

### DETAILS

call tssetjellycontenttimestamp(?,?)

Help

Close

To refresh the list to see the latest queries, click [Refresh this page]. To clear the list, click [Clear list].

[Back to top](#)

## Improving System Performance

This section describes the reason(s) that a system indicator may display a status of orange or red and the steps you can take to investigate the matter and change your XperienCentral settings in order to improve your website's performance.

### Caching

If [caching](#) is enabled in XperienCentral, a green indicator is shown and if it is disabled, a red indicator is shown. Caching allows XperienCentral to handle the load created by many page requests by using an intelligent mechanism that returns pages without having to regenerate them completely every time they are requested. In all cases, caching improves website performance and it is highly recommended that you enable it for your website.

To enable/disable caching, navigate to the [Setup Tool](#). On the [General](#) tab, navigate to the "frontend\_system\_settings" section and select allow\_cache:

#### frontend\_system\_settings (default)

| Name                           | Value                                     |
|--------------------------------|---|
| allow_cache                    | <input checked="" type="checkbox"/>       |
| cache_cacheable_response_codes | 200                                       |
| cache_directory                | \${startup_settings.base_directory}/cache |
| cache_directory_depth          | 3   |
| cache_max_stale_time           | 0   |

For more information on caching in XperienCentral, see [Caching](#).

### Adjusting Timestamp Caching

The caching module of XperienCentral stores a timestamp in the timestamp database for each cached web page. This timestamp is needed in order to determine whether the web page needs to be regenerated or whether it can be retrieved from the cache. For high volume websites, the number of web page requests per second can be so high that even the relatively simple task of performing one SQL query to retrieve this timestamp from the database can cause a bottleneck. If this exceeds 500 milliseconds, it may indicate a performance issue.

To prevent timestamp retrieval from becoming a bottleneck, timestamps can be cached in memory for <n> milliseconds. Note that this also causes a delay of <n> milliseconds; during the <n> milliseconds, modified web pages will not be regenerated. It is recommended that you set timestamp caching to at least 5000 milliseconds (5 seconds). If you set timestamp caching to fewer than 5000 milliseconds, this can cause database bottlenecks especially for websites that experience a large number of page requests.

Timestamp caching is configured in the XperienCentral [Setup Tool](#). Navigate to the "frontend\_system\_settings" section of the [General \(R30 and older\)](#) tab:

#### frontend\_system\_settings (default)

| Name                           | Value                                     |
|--------------------------------|---|
| allow_cache                    | <input checked="" type="checkbox"/>       |
| cache_cacheable_response_codes | 200                                       |
| cache_directory                | \${startup_settings.base_directory}/cache |
| cache_directory_depth          | 3   |
| cache_max_stale_time           | 0   |
| cache_number_of_directories    | 30  |
| cache_statement_timeout        | 2   |
| cache_timestamp_expire         | 30000                                     |

### Lowering Page Response Times

Page response time performance for pages returned from the cache as well as pages generated by the Page Generator is measured in the following way:

| Indicator | Value |  |  |  |
|-----------|-------|---|---|---|
|-----------|-------|---|---|---|

|                            |              |                            |                       |       |
|----------------------------|--------------|----------------------------|-----------------------|-------|
| Largest page response time | milliseconds | Less than or equal to 1000 | Between 1000 and 4999 | 5000+ |
| Average response time      | milliseconds | 0-99                       | 101-499               | 500+  |

## Large Response Times for Individual Pages

A large page response time can be caused by many factors, for example for pages that contain several complicated components (page sections, and so forth) that take time to generate from the JSP. If a particular page has an unusually large response time compared to other pages, you can be fairly certain that there is an issue with the page itself. Check the contents of the page to see what could be causing it to take longer to generate and/or load. You should also compare the structure and content of a page with a large response time with other similar pages on your website that do not have a large response time. This can help to identify what the cause or causes are that make the particular page take more time to generate or load.

Caching settings for individual page parts can be configured independently of each other. If the caching timeout for one or more page parts is set to a low threshold, this can cause the page response time for a page with several page parts to increase. Check the caching timeout for the page parts that make up the page experiencing large response times in order to ensure that the page section is not generated more often than needed.

## Large Average Response Time




If the average response time is high, it could indicate a more fundamental problem with the website itself, the webserver(s) in the production environment and/or the database(s) being used for the website. If other processes are competing with the XperienCentral components on the production web server, this can cause the average page response times to increase. Check the configuration of your web server to see whether there are too many processes running on it. Is the XperienCentral deployment running on a single server or in a clustered environment? If XperienCentral is running on a single server, both the Workspace environment as well as the front-end have to share the resources. It is highly recommended that you run XperienCentral in a [clustered environment](#) for the following reasons:

- **Performance:** more servers means more capacity to generate pages for visitors of the website and for the website editors to make changes to the website.
- **Security:** The edit environment can be separated from the website frontend.
- **Fail-over:** Prevents the website from becoming unreachable if a server goes offline.

A large average response time can also be caused by XperienCentral caching settings. If caching is disabled, every time a page is requested it must be generated. This can lead to large page response times, especially if your website has many visitors simultaneously requesting pages. See [Caching](#) for more information.

## Lowering Query Execution Times




Query execution time is measured in the following way:

| Indicator       | Value        |  |  |  |
|-----------------|--------------|---|---|---|
| Slowest Queries | milliseconds | 500   | Between 500 and 2000  | 2000+   |

Query execution time that exceeds 500 milliseconds may indicate a performance issue. The list of the slowest queries contains both internal XperienCentral queries as well as queries that have been developed especially for your web site. If a query that you have developed is especially slow compared to other queries or is consistently slow, you might want to examine it carefully to see whether it is possible to restructure it for greater efficiency.

## Lowering the Average Query Execution Time

Average query response time is measured in the following way:




| Indicator                    | Value        |  |  |  |
|------------------------------|--------------|---|---|---|
| Average query execution time | milliseconds | 0-500   | 501-1999  | 2000+   |

If the average query execution time is higher than 501 milliseconds, it could indicate a fundamental problem with the performance of the database that XperienCentral is using. If other applications are using the same database as XperienCentral, this can cause the average query execution time to increase. If you consistently experience high average query execution times, you should investigate how the database you are using for XperienCentral is configured and whether it is functioning correctly/optimally.

## Lowering the Cache Ratio



The average cache ratio is measured in the following way:

| Indicator           | Value   |  |  |  |
|---------------------|---------|---|---|---|
| Average cache ratio | percent | 10+   | between 1 and 9   | less than 1   |



A cache ratio that is less than 10 indicates that XperienCentral is having to generate a large number of pages instead of retrieving them from the cache. This can be because:

- Caching is not enabled.
- Some page parts have too low of a cache threshold. The page parts that are used to construct a XperienCentral page can each have an individual cache timeout setting. The caching timeout threshold determines when the page part should be regenerated from the JSP. If one or more page parts, especially those that appear on many pages is regenerated more often than necessary, this can cause the cache ratio to decrease because more pages have to be generated instead of being retrieved from the cache.

When investigating why the average cache ratio is low, pay particular attention to pages that have particularly small average cache ratios in the [Page cache ratio tab](#). Examine the structure of these pages in the Workspace environment in order determine why they have to be regenerated more often than other pages. It is possible that the caching timeout for these page parts is unnecessarily short.

## Setting Debug Options




The XperienCentral debug options are useful during the website development phase but should not be enabled on a live website because they can slow down performance. The possible statuses of the debug option in the dashboard are:

| Indicator     |  |  |
|---------------|---|---|
| Debug options | All debug options are disabled.   | One or more debug options is enabled.   |

To enable/disable the XperienCentral debug options that affect performance, navigate to the [Setup Tool](#). On the [General \(R30 and older\)](#) tab, navigate to "development\_settings". Clear the checkbox for `enable_debug_messages`. In a production environment, the setting `enable_view_resolve_cache` should be selected.

## Setting the Log Level

The level at which you log messages has a direct effect on system performance. Each time a system event is logged, a small performance hit occurs. Logging messages of type `SEVERE` affects the system the least whereas logging at the `INFO` level can cause the system to slow down due to the sheer volume of messages that are written to the log. The effect that the logging level has on XperienCentral performance is indicated as follows.

| Indicator |  |  |  |
|-----------|---|---|---|
| Log level | SEVERE  | WARNING   | INFO  |

To change the log level, follow these steps:

1. Open the `<tomcat root>\conf\logging.properties` file.
2. At the end of the file, the following lines should appear. If they do not, add them.
  - `org.apache.jackrabbit.core.level = <level>`
  - `org.apache.jackrabbit.core.state.level = <level>`
  - `nl.gx.XperienCentral.level = <level>`
  - `org.apache.jackrabbit.level = <level>`
  - `org.quartz.level = <level>`
  - `org.springframework.level = <level>`
3. Substitute `<level>` with the level at which you want to log (`SEVERE`, `WARNING`, or `INFO`). Logging at the level `SEVERE` is recommended.

For more information on logging events in XperienCentral, see the [Log](#) administrative page.

[Back to top](#)

---

## Creating Custom Performance Indicator Tabs

In addition to the four default tabs, you can develop and extend the XperienCentral Performance Dashboard using custom plugins. See [Adding Custom Indicators to the Performance Dashboard](#).

[Back to top](#)