

CA Workload Control Center

Performance Tuning Guide r11.3



This documentation and any related computer software help programs (hereinafter referred to as the "Documentation") are for your informational purposes only and are subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be used or disclosed by you except as may be permitted in a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2010 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

Contents	3
General Recommendations	5
Requirements for Passive Users	6
Monitoring	7
How to Support More Concurrent Users.....	7
How to Support More Concurrent Requests (Load).....	7
Limitations for Flow Rendering.....	7
Limitations for the Monitoring Database.....	8
Optimize Settings for Better Performance.....	10
Job Status Console	11
General Recommendations.....	11
Limitations.....	11
Job Status Console Cache.....	13
Enable Global Session in Job Status Console.....	13
Application Editor	14
Optimize Commit for Better Performance.....	15
Quick Edit	16
CA Embedded Entitlements Manager (CA EEM)	17
General Recommendations.....	17
Event Policies.....	17
CA EEM Failover.....	17
CA EEM Updates.....	18

Logging Recommendations	19
Scheduling Engine Recommendations	20
Appendix A: How to Adjust Memory for Service	21
Appendix B: How to Set Up an External Database	23

General Recommendations

This document includes recommendations and best practices for performance tuning for CA Workload Control Center (CA WCC) r11.3.

General recommendations for performance tuning are as follows:

1. *Make sure that the CA WCC server machine meets the recommended specifications identified in the Release Notes.*

The CA WCC Release Notes include details of hardware requirements in the System Requirements section.

2. *Check the network as well as machine usage, especially if a subset of users have problems.*

Occasionally, there will be a problem with network bandwidth in an area of the enterprise. If there is a subset of users who are having problems, a simple test is to perform a ping to see what the transit time is between the user's machine and the CA WCC server.

3. *Resolve resource contention issues for virtual servers and shared servers.*


If CA WCC is installed on a shared server or virtual server, ensure that there are no resource contention problems. For example, if CA WCC is installed on a virtual server that hosts multiple environments and you experience periodic performance drops, verify that the other environments are not causing resource utilization spikes.

4. *On *NIX operating systems, use `ULIMIT -n = 4096` to increase the number of file descriptors per process.*

Requirements for Passive Users

The following apply to passive users (users who are logged in only):

- CA WCC with all services running and no users logged in occupies approximately 2 GB of RAM.
- When a user logs in to CA WCC with all CA WCC services running, roughly 8 MB of additional memory is consumed.



Memory consumption for passive users

Monitoring

How to Support More Concurrent Users

The number of possible active concurrent users logged in to the system depends on the memory available for the CA WCC Monitoring Server service. For more than 40 concurrent users, memory for the CA WCC Monitoring Server service should be increased. A rough estimate is 5 MB per concurrent user. Memory must be updated for the CA WCC Monitoring Server service. See **Appendix A: How to Adjust Memory for Service** for information about updating memory.

The CA WCC Monitoring Server service is a 64-bit service, so more than 2 GB can be allocated for heap space.

How to Support More Concurrent Requests (Load)

Several users opening the Monitoring page in a given time period is an example of concurrent requests. For a higher number of concurrent requests, you will experience longer response times. In most cases, response time is related to the number of concurrent users; however, response time also depends on the load generated by those users. Opening the View Details page takes the most time, so that page is used in the following example.

Suppose that 60 users are logged into CA WCC and 10 of those users request View Details each minute. Assuming the remaining users perform one request every 5 minutes (approximately 0.2 request per minute), the number of requests per minute can be expressed as follows:

$$10*1 + 50*0.2 = 20 \text{ requests per minute}$$

Limitations for Flow Rendering

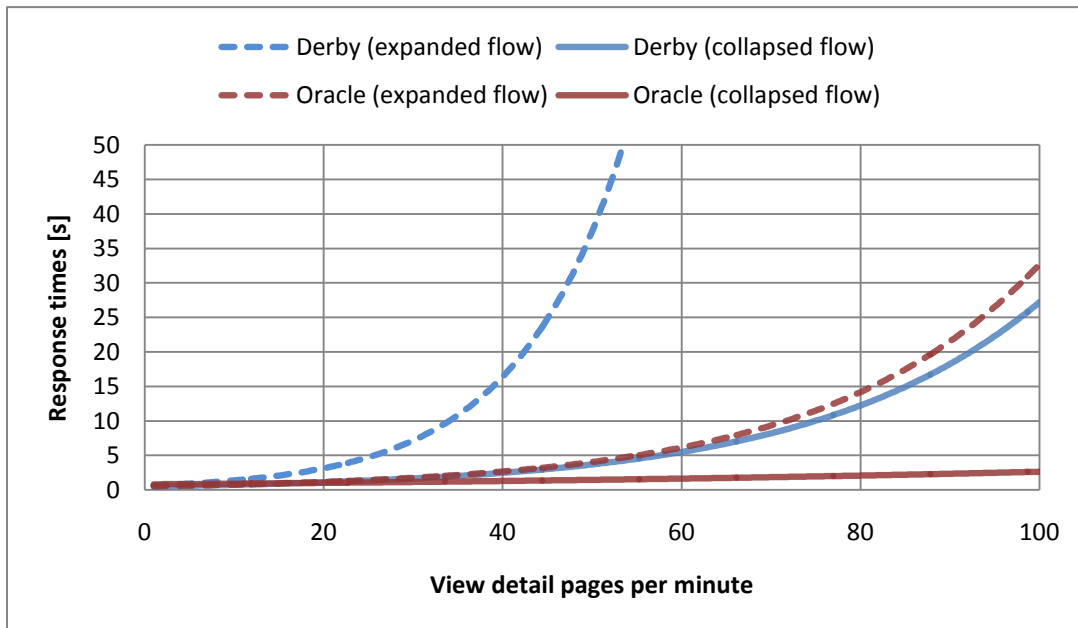
Performing calculations and rendering the Flow section in Monitoring increase the load on the CPU and the Monitoring database significantly. You can improve performance if you set up the Flow section so that it is not rendered or so that it is collapsed by default. Note that the collapsed Flow section can be expanded by the user when they want to view the flow.

Flow section defaults are set up in Configuration Manager as follows:

- The option to disable rendering can be set up in the view definition with the **render_flow** attribute, or in Monitoring Preferences in Configuration Manager with the **Flow Threshold Per Flow** or **Flow Threshold All Flows** attributes.
- The options to collapse the Flow section can be set up in Monitoring Preferences in Configuration Manager with the **Details Flow Section Collapsed** or **Flow Threshold Collapsed** attributes.

It is recommended to collapse the Flow section by default for the following use cases:

- When you have a large number of jobs in views (> 100 jobs per view)
- When you have a large number of job runs per day (> 15,000 job runs)



Response times for expanded and collapsed Flow section (52,000 job runs per day).

Limitations for the Monitoring Database

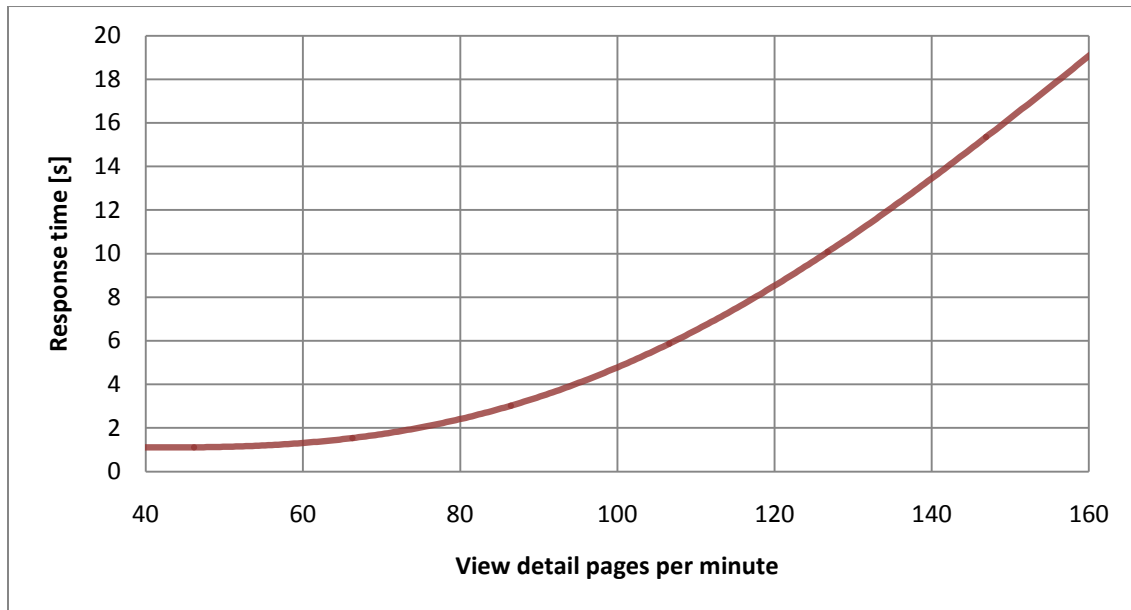
The default Derby database for Monitoring can be used in the following scenarios:

- When you have a small load.
- If your environment processes 50 or fewer requests for view detail with the Flow section collapsed.
- If you have up to 25 requests for view detail with the Flow section expanded. (The threshold is around 25 requests.)

For medium and large loads, one of the following supported external databases should be used instead of the Derby database:

- MS SQL Server 2008
- MySQL Enterprise 5.1
- Oracle 11g

All three supported external database engines (MSSQL, MySQL Enterprise and Oracle) have similar response times. The graph that follows shows the limitations for those databases.



Response times for external Oracle database and collapsed flow (52,000 job runs per day)

Test environment:

- Linux VM
- Two cores
- 8 GB RAM

For the test environment, the threshold is around 100-120 users. In this case, the threshold also depends on the available computing power of the CA WCC server. CA WCC will perform better with more CPUs available or on physical machines (the exact threshold depends on specific hardware and cannot be generalized in a graph).

For information about setting up an external database, see **Appendix B: How to Set Up an External Database.**

Note the following:

- The external database does **not** need to reside on the same machine as CA WCC; however, it is recommended that both the CA WCC server and the database server are located close to each other.
- If an external database is used, the CA WCC Monitoring Repository service should be stopped.
- There must be sufficient network bandwidth and latency between the CA WCC server and the database server.

Optimize Settings for Better Performance

Only settings related to performance tuning are mentioned in this section. Settings are located in Monitoring Preferences in Configuration Manager. To configure settings, click the CONFIG tab and select Monitoring from the drop-down list in the Preferences section.

Server section:

- **Collector Sentry Mode** – Set to the *filter* option when you want to collect data from servers that are defined in active Monitoring view filters only; otherwise, it should remain at the default value, *all*.
Example: CA Workload Automation AE servers defined in Configuration Manager are: A1, A2, A3, A4, A5. Only A1, A2 are defined in active monitoring view filters.
The following show the effects of the settings:
 - *all* – collect data from servers A1, A2, A3, A4, A5.
 - *filter* – collect data only from servers A1, A2.
- **Collector Sentry Thread Pool Size** - For the default value (-1), the pool size equals the number of CPU cores; however, it is useful to set it to the exact value (the number of CPU cores minus 1) if the collector service generates a high load (> 25% on average). Update times will be increased in this case.
- **Security Collector Run Interval** – This value can be increased if only rare changes in CA EEM access policies for Monitoring Views are expected between the time a user logs in and logs out. CA EEM checks increase the load on the database and CPU significantly.
- **Job Collector Run Interval** – This value can be increased if the collector service generates a high load on the server (10 or 15 seconds is sufficient). This can happen if CA WCC is attached to many CA Workload Automation AE servers that generate a high load (there are many updates during a 5-second interval).

UI-View/Collection Details section:

- **Flow Threshold Collapsed** – The recommended threshold depends on database usage and on flow usability with a large number of jobs (see the **Limitations for Flow Rendering** section for details).
- **Details Flow Section Collapsed** – With this setting, the Flow section is collapsed for every view. Same as threshold = 0.

Job Status Console

General Recommendations

Job Status Console (JSC) performance is largely dependent upon the number of open views and the size of the views (the number of jobs the view contains). For best performance, the size of the views should not exceed 500 jobs.

To reduce the load on CA Workload Automation AE servers, the JSC Cache should always be used (see the Job Status Console Cache section for details).

When performance is important, you should also minimize Alerts usage. Since Alerts generation is a relatively costly operation in terms of resources, they should be used only if necessary. CA Workload Automation AE Alarms should be utilized instead if they meet all requirements of the enterprise.

When Alerts are needed, alert policies must be created carefully; only a small set of jobs should be targeted by the alert policy settings. If possible, alert and alarm filters in a view should also be restricted to the jobs of that particular view only. If you have issues with JSC performance and there are too many alerts being generated, you can decrease the **Alert Deletion Threshold** parameter value. This value specifies number of days that alerts remain in the Alerts database before they are deleted.

To improve JSC performance, the **AS_THREAD_STACK** variable should be set as follows in the

`%CA_WCC_INSTALL_LOCATION%/JobStatusConsoleServer/conf/wrapper.conf` file:

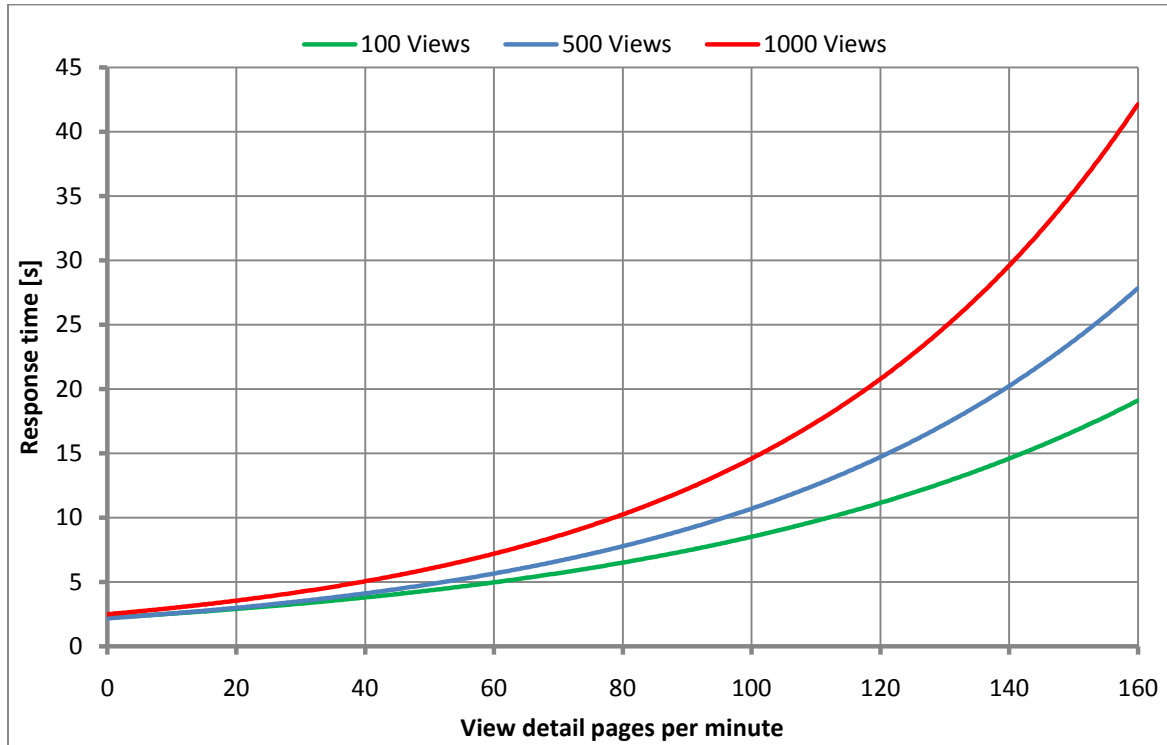
```
set.AS_THREAD_STACK=65536
```

Limitations

The threshold for concurrent requests in JSC depends on number of views and the size of the views. Generally, the threshold is around 100 concurrent requests. We strongly recommend not exceeding this limit significantly; otherwise, you can encounter a high number of failed or timed out requests.

For a use case with approximately 500 views, where each view contains an average of 500 jobs and all of the jobs are defined on a single CA Workload Automation AE server, the most time-consuming page - view detail - is displayed within approximately 10 seconds for up to 100 concurrent requests.

The approximate dependency of response time on the number of concurrent requests in the JSC view detail pages for three use cases (100, 500, and 1000 JSC open views) is demonstrated on the following chart. All of these views had an average of 500 jobs.



Dependency of response time on the number of concurrent requests

Note that these measurements were performed while using all of the recommended optimizations as described in the following sections and with CA WCC residing on a server with the following configuration:

- Linux VM
- Two cores
- 8 GB RAM

Also note that the actual response times may vary significantly from these results based on number of factors, for example:

- The configuration of the server on which CA WCC resides
- The performance of the connected CA Workload Automation AE servers
- Current load on other CA WCC components (the number of concurrent users working with CA WCC)
- The number and complexity of Alert policies defined in JSC

Job Status Console Cache

When a user opens a view or displays job statuses for a CA Workload Automation AE server, the data is automatically stored in the Job Status Console Cache. To get fresh data, the user must click the Refresh button in the view.

To change the Job Status Console cache properties, click the CONFIG tab, select Job Status Console from the drop-down list in the Preferences section, and select the following properties:

- **Caching for AE** – Enable (this option should never should be disabled)
- **AE Cache Size** – Consider the following when setting the value:
 - The number of views that are cached
 - Set the value to number of views + expected number of concurrent users of JSC
 - Setting it to a high rather than low value will not degrade performance

Enable Global Session in Job Status Console

To enable the Global Session option in JSC:

1. Click the CONFIG tab to open CA WCC Configuration Manager.
2. Select Job Status Console from the drop-down list in the Preferences section.
3. Check the Use Global Session check box only if you plan to use the Global Session.
Important! If you do not plan to use the Global Session, make sure that you uncheck the Use Global Session check box.
4. In the AE Server Configuration, make sure the Monitor ID and Monitor password are specified correctly.
5. If you made a change in the server configuration, deploy the change in the Deploy section.
6. If you made a change to any Job Status Console or server property, restart the services selected in the Services section in Configuration Manager.

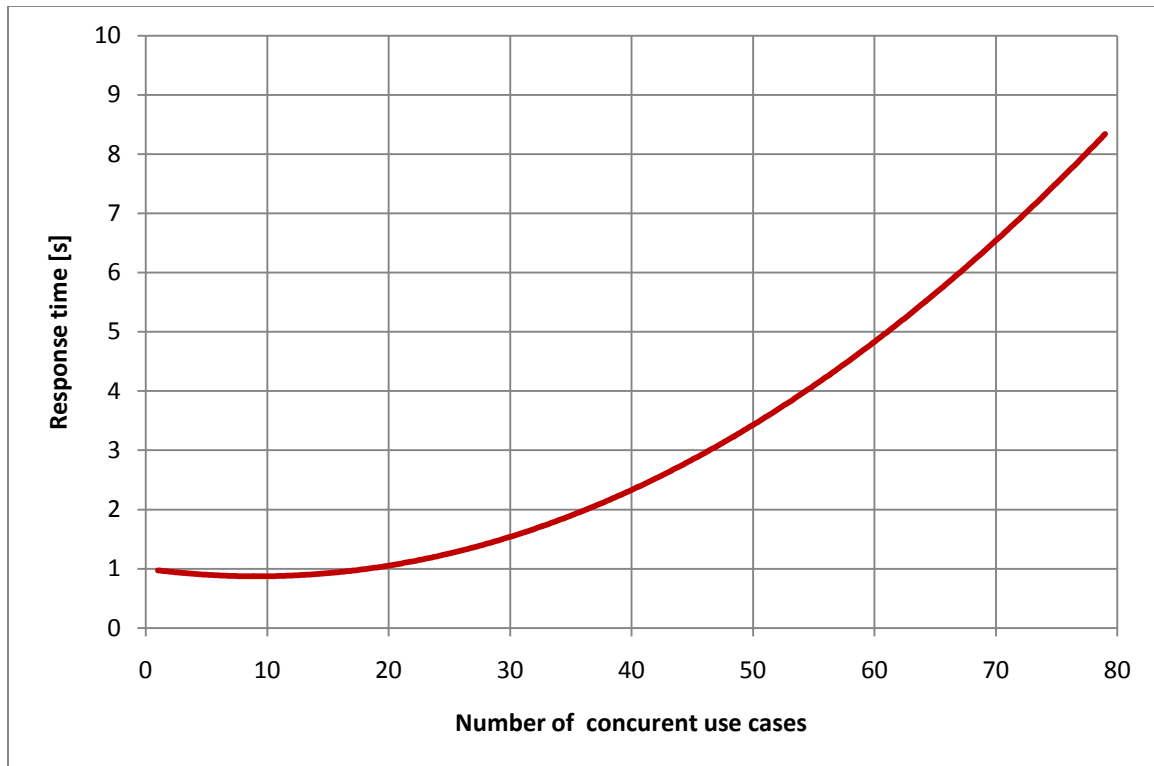
Guidelines for using the JSC Global Session

To help you determine whether to use the JSC Global session, note that the break-even point is two watchers for a view. Watchers are users who simply watch Job Status Views on the Main page and do not actively use a view. Active view users should not influence your decision because they work the same in both methods. Note also that if you have access to five views but actively work in one, you are a watcher of four views.

If you have 800 views and each view belongs to one unique user only, it is recommended to uncheck the Global Session option. If you have 800 views and three watchers for each view, it is recommended to check the Global Session option. Considering those two scenarios, you need to know exact usage patterns to determine which method makes the fewest calls to CA Workload Automation AE and uses the least amount of memory.

Application Editor

The number of active concurrent users logged into system depends on the memory available for the CA WCC Application Editor Server service. Using the default settings, Application Editor supports up to 50 users. If you increase the memory setting to 1 GB, up to 80 concurrent Application Editor users are supported.



Calculations in the graph above were based on periodic import/clean of a flow using 400 jobs, 100 logical operators, and 300 links. The memory setting for the CA WCC Application Editor Server service was 1 GB.

The CA WCC Application Editor Server service is a 32-bit process, therefore a maximum of 1 GB of memory can be allocated. For details, see **Appendix A: How to Adjust Memory for Service**.

Note: In order to achieve the scalability of Application Editor described in this section, fix T5TN001 (258-APPEDIT PERFORMANCE) must be applied first.

Optimize Commit for Better Performance

In Application Editor, you can specify how many threads you want to use in parallel for a flow commit operation. Valid values are 1–50, and the value can be changed in Configuration Manager. A higher value will allow a faster commit speed.

Because CA Workload Automation AE box jobs are always committed using a single thread before other jobs, you may notice a difference in the commit speed only in the case of an increased number of non-box jobs present in the flow.

Quick Edit

The default settings for Quick Edit support up to 50 concurrent Quick Edit users. If you increase the memory setting to 1 GB, up to 80 concurrent Quick Edit users are supported.

The CA WCC Quick Edit Server service is a 32-bit process, therefore a maximum of 1 GB memory can be allocated. For details, see **Appendix A: How to Adjust Memory for Service**.

CA Embedded Entitlements Manager (CA EEM)

General Recommendations

General recommendations for CA EEM are as follows:

- Ensure that the CAELM application is installed and operational.
- Review event policies in both the CA Workload Automation AE and CA WCC applications to ensure that they are correctly defined. (See the Event Policies section below.)
- If you are using Active Directory as your CA EEM user repository, ensure that you set appropriate filters in the EEM/AD connection to limit the allowed users to groups and individuals that actually use CA WCC, rather than retrieving the entire Active Directory.

Event Policies

Auditing of CA EEM Events can affect the performance of CA WCC applications. If you are experiencing performance problems, check event policies settings for both CA WCC and CA Workload Automation AE applications and consider either limiting or disabling CA EEM auditing.

In the CA EEM Web Admin user interface, go to Manage Access Policies, Event Policies. The DefaultSubmitEvent policy should not contain the resource authorizeWith* (in the form of this pattern exactly), as well as all matching resources like authorizeWithSession, etc. If this resource is present, remove it from the policy. The same approach applies to custom, non-default policies that are created and whose filter includes any CA WCC application that operates with CA Workload Automation AE job list or views (for example, Monitoring, JSC, and so on). If the removal of authorizeWith* resource from Event Policies does not help, consider disabling those policies.

CA EEM Failover

When CA EEM servers are connected in a failover configuration, you may experience performance issues during the initial load of the CA WCC login page. The initial load is the first login into CA WCC that occurs after the Launcher Service starts or restarts. Performance can be affected when the list of CA EEM servers in failover mode is provided to CA WCC. In this case, CA WCC attempts to connect to each CA EEM server in turn until the first available CA EEM server is found. All attempts to connect to CA EEM servers that are down are performed within a timeout, which results in the CA WCC login page being successfully loaded after all unsuccessful attempts to connect to CA EEM servers have occurred. The total connection time is equal to the sum of all CA EEM server connection timeouts.

The same delay occurs in CA WCC Batch Interface utilities, which also attempt to connect to CA EEM servers in the order specified in the list. The startup time of Batch Interface utilities is affected when the first servers in the CA EEM server list are unavailable. To provide better performance, it is recommended to specify the CA EEM servers with the highest availability (the lowest probability of failure) at the beginning of the list and servers with a higher rate of failure at the end of the list.

For example, suppose you have three servers named A, B and C. Server A is completely dedicated to CA EEM and has good performance characteristics. Server B is shared by many services and frequent reboots are expected, due to the other services. Server C is down most of the time and is turned on when servers A or B are down. In this scenario, specifying the list of servers in the order of C,B,A is not recommended because server C (which is currently turned-off) will cause a delay in establishing the initial connection. Instead, the recommended order is A,B,C. Dedicated server A is listed first so that the first attempt to connect to CA EEM will be successful most of the time.

CA EEM Updates

CA WCC r11.3 is delivered with the CA Common Components DVD. The DVD includes an installer for CA EEM, which is used for CA WCC identity management. If performance problems are discovered in CA EEM, a patch will be released for the CA EEM server as well as for the CA EEM safe.jar library that is embedded in CA WCC. Check support.ca.com or contact CA Support for the latest patches.

Important! Refer to [CA WCC Certification Information on the latest supported versions of CA EEM \(Server\) for various CA WCC versions.](#)

Logging Recommendations

CA WCC includes extensive logging for all components. By default, logging levels are set to INFO or ERROR. To optimize performance, ensure that you do not set logging values to DEBUG.

Logs are located in `%CA_WCC_INSTALL_LOCATION%\Logs\Application (Windows)` or `$CA_WCC_INSTALL_LOCATION/Logs/Application (*NIX)` and there is a `log4j.properties` file in each subfolder that controls these logging settings. We recommend that if you change any logging settings, you first make a backup copy of the original `log4j.properties` file. When you have made changes to the `log4j.properties` file, save the file then stop and start the associated service(s).

To reduce the amount of storage used for logging, the number of rolled log files being kept can be specified by setting **`wrapper.logfile.maxfiles`** parameter in `wrapper.conf`. The log is rolled based on size and the value 0 indicates no limit on the number of logs.

Scheduling Engine Recommendations

CA Workload Automation AE r11.3 utilizes an application server to manage secure communication with the CA Workload Automation AE database. The application server is typically installed on the CA Workload Automation AE Scheduler machine. If the machine where the application server is installed is heavily used, the machine may become CPU-bound. To maximize performance, install the instance of the application server that is used by CA WCC to a lower-load machine. There are three possibilities for achieving this:

- If CA WCC is on a physical machine with free memory and CPU resources (at least five cores and more than 10 GB of RAM), the application server for CA WCC can be installed on the CA WCC server.
- If CA Workload Automation AE has shadow scheduler installed, CA WCC can point to the shadow server.
- A new CA Workload Automation AE application server can be installed on different machine than the CA Workload Automation AE Scheduler.

Application servers are taken in order as defined in the Server Properties section of Configuration Manager in the **Application Server Host Address** attribute, so application server(s) dedicated to CA WCC must be first in the list.

Appendix A: How to Adjust Memory for Service

You must ensure that there is adequate memory on the machine for the settings; the Java Virtual Machine (JVM) will not load if maxmemory cannot be allocated.

Note: The **maxmemory** parameter should be set to a maximum of 1 GB for 32-bit services (see the table that follows).

1. Open `%CA_WCC_INSTALL_LOCATION%/ServerDirectory/conf/wrapper.conf`.
2. Change **wrapper.java.maxmemory** parameter. The value is in MB.
3. Save wrapper.conf file.
4. Restart the service related to the particular server according to the table that follows.

Service name	Service Directory	32- or 64-bit	CA WCC Applications
CA WCC Application Editor Server	AppEditorServer	32-bit	APPEDIT
CA WCC Command App Server	CmdAppServer	64-bit	RPT (Forecast),ECLI
CA WCC Configuration Server	ConfigServer	64-bit	CONFIG
CA WCC Event CPM Server	EventCPMServer	32-bit	EVTCON
CA WCC Reports Repository	ReportsDB	64-bit	Used by RPT
CA WCC High Availability Server	HAServer	64-bit	HA
CA WCC Job Status Console Server	JobStatusConsoleServer	32-bit	JSC
CA WCC Job Status Console Servant	JobStatusConsoleServant	32-bit	JSC
CA WCC Launcher Server	LauncherServer	64-bit	MAIN UI, CRED
CA WCC Monitoring Repository	MonitorDB	64-bit	Used by MONIT

Service name	Service Directory	32- or 64-bit	CA WCC Applications
CA WCC Monitoring Server	MonitoringServer	64-bit	MONIT
CA WCC Monitoring Collector	MonitoringCollector	32-bit	Used by MONIT
CA WCC Quick Edit Server	QuickEditServer	32-bit	QEDIT
CA WCC Quick View Server	QuickViewServer	32-bit	QVIEW
CA WCC Resources Server	ResourcesServer	64-bit	RES
CA WCC Remote Services	RemoteServices	32-bit	Used by MONIT,RES

Appendix B: How to Set Up an External Database

This procedure shows how to set up an external database for Monitoring.

1. Click the CONFIG tab to open Configuration Manager.
2. In the Preferences section, select Monitoring from the drop-down list.
3. Fill in the following mandatory fields in the Database section of the table:
 - Database Type
 - Host Name
 - Port
 - Database Name
 - User Name
 - User Password
4. Save the preferences changes.
5. Restart CA WCC Monitoring Server and CA WCC Monitoring Collector services.
6. (Optional) Stop the CA WCC Monitoring Repository service (the external database is used instead of this service).

Note the following:

- The external database must exist before the first CA WCC startup.
- CA WCC installs its own database drivers. It is not required to install the any database client software on the CA WCC server.
- The database user must have permissions to create tables, to drop and create indexes, and to select, insert, update, and delete rows.
- During the first CA WCC startup, the database tables are created. During each restart of the Monitoring service, the indexes are recreated.
- For MySQL Enterprise only: The database user must have granted access to the database from remote hosts if the database is not installed on the same host as CA WCC.