QAD Adaptive Sizing Guide

August 2026

Version 15.0

Contents

Introduction **4**
Sizing QAD Adaptive 4
Sizing Principles 4
Sizing for Peak Loads 5
Target Audience 5
How to Use This Document 5
A General Disclaimer 6
Quick/Ballpark Sizing **7**
Introduction 7
Single-Tier Quick Sizing 7
CPU Sizing 7
CPU Sizing AI Agents 8
Memory Sizing 8
Disk Sizing 9
Network Sizing 10
Client PC / Laptop Sizing 10
QAD Adaptive Client 10
User Count Sizing versus Throughput Sizing 10
Delta Sizing 10
Upgrading User Counts 11
Appendix A: Understanding the Sizing Process **12**
Understanding the Customer's Requirements 12
Requirements Gathering 12
The Questionnaire 12
The Sizing Workshop 12
Performance Monitoring 13
Security 13
On-Premise / Hosted / Cloud Considerations 13
High Availability / Disaster Tolerance 13
The Sizing Process 14
Using the Quick / Ballpark Sizing 14
Sizing to a Budget 14
Sizing by Reference Installations 14

Delta Sizing / Upgrade Sizing 14
Network 15
Appendix B : Sample Sizing Questionnaire 16
Existing QAD System 16
Provide details of existing software installations 16
Customizations 16
Current Workload and Performance 16
Provide details of new and future requirements (this sizing exercise) 16
Interfaces 16
Preferred Deployment 17
Hosting 17
High Availability and Disaster Tolerance 17
Users 17
Appendix C: Revision History 18

Introduction

Sizing recommendations here for QAD Adaptive 2608 components only and do not include or account for any other applications running that have their own resource demands such as:

- Logstash
- Performance Monitors
- Icinga
- Tomcat IO
- Anti-Virus
- Etc

Server resources should be increased to support non AUX components.

Sizing QAD Adaptive

This document is intended to provide background information to help you determine accurate sizing for customer systems running QAD Adaptive..

QAD has years of experience in sizing and supporting systems running its software. This experience is based on formal, high-user count load testing in hardware vendor labs, internal performance testing in the QAD product labs, and experience gained from auditing live customer environments.

Accurate sizing is difficult, mostly because of poorly defined requirements. A lack of clear facts introduces assumptions and guesses, which reduces accuracy. It is hard to predict the future if we do not know what is happening in the present, hence a large portion of this document contains guidance on how to increase the known facts about the proposed sizing.

Sizing Principles

System sizing for QAD software deployments is part of the broader performance engineering process.

1. Gather requirements
2. Define sizing inputs
3. Consider influencing factors
4. Select hardware models and processors
5. Define a balanced system around the base hardware selection
6. Purchase / implement
7. Monitor / plan capacity / tune
8. Resize

Sizing for Peak Loads

QAD Adaptive systems are sized for peak loads, not average loads. A characteristic of such systems is periods of high load during certain times of the day and certain days of the month. It is critical that during the periods when the use is highest (such as end of financial period, invoice posting, or shipping) that there are adequate system resources to cope.

A side effect of this can be observed under-utilization during non-peak periods. This is normal behavior.

Target Audience

The target audience for this document is anyone interested in sizing QAD Adaptive.

Typically the audience includes sales, technical pre-sales, technical services, and QAD performance subject matter experts (SMEs).

On the customer side, the audience may include CTOs, CIOs, IT Managers, IT Architects, Systems Administrators, Finance, and Project Managers.

How to Use This Document

This document is intended to help size QAD Adaptive for customers, using the sizing principles approach.

A quick / ballpark sizing section is intended to provide what should be a reasonably accurate sizing based on what the QAD performance team knows about QAD Adaptive.

Possible Approach : Simple hardware budget sizing (small customers)

> For smaller customers (fewer than 200 named users), the simplest and most straightforward approach could be to:
>
> - determine the customer hardware budget and preferred hardware vendor
> - select a server from a hardware vendor website that falls within the budget
> - sanity check the selection against the ballpark sizing

Possible Approach : Estimate and refine (small-to-medium customers)

> It might be useful to initially consult the quick / ballpark sizing to determine the approximate level of hardware required, and then apply known facts and deployment influencing factors to

define the final sizing. This could be an iterative approach based on feedback from the customer or QAD experts.

Possible Approach : Sizing by experts (large / complex customers)

For large and complex deployments, it is strongly recommended to take a more formal approach to the sizing process. Appropriate data gathering should be carried out, using the sample questionnaire as a guideline.

For particularly large or sensitive projects, a formal round of load testing on acceptance testing hardware should be considered.

A General Disclaimer

A sizing project is not a performance guarantee or promise that the deployed system will work at acceptable levels.

When delivering a sizing specification to a customer, caveats and disclaimers should be used to highlight the fact that the recommendations given are not a contract nor can they be used as the basis for formal service level agreements.

It is important to note that even an oversized system can suffer poor performance if customized browses / inefficient code is allowed onto a production system without understanding the impact they can have on performance. Customers should also expect a large impact on system performance if they make heavy use of exporting browse data to flat files or have large numbers of users concurrently refreshing action centers.

Quick/Ballpark Sizing

Introduction

This section is intended to provide a quick / ballpark sizing based on internal QAD performance testing and the experience of existing customers running QAD Adaptive.

While it is preferred that a more formal process is followed for sizing, sometimes expediency or a lack of known facts, or just the need for a ballpark estimate, drives the need for such quick sizing information.

Even though this section provides a simplified sizing model, it has been historically shown to be reasonably accurate and consistent from version to version of core QAD Adaptive software.

Single-Tier Quick Sizing

The information in this section assumes a monolithic / single-tier deployment. All QAD Adaptive components are located on a single server. User counts are estimated **concurrent** users.

Assumptions: QXtend will be used at a moderate level relative to the proposed concurrent user count.

CPU Sizing

Benchmarking and customer validation has confirmed optimal performance for QAD Adaptive on AWS R8I instance types for **non AI Agent Users.**

QAD Version	Minimum CPUs	First 1000 Users	Additional Users
QAD Adaptive	4	25 users per core	30 users per core

- **Note 1:** It is recommended to use the fastest available server side CPU chipsets for optimal performance
- **Note 2**: If sizing is less than the minimum number of CPUs, we should go with the minimum
- **Note 3**: Figures do not include sizing for any other applications running (eg monitors, logstash, anti-virus, icinga etc)
- **Note 4**: Important CPU Considerations
    - CPU sizing guide assumes dedicated access to the required CPUs

- Heavy concurrent Action Center use may require additional CPU Resources
- Activity Feed Entity tracking for calculated fields results in additional API fetches, which may require additional CPU resources
- Bulk record changes for fields that have Activity Tracking enabled for non-database fields may require more CPU resources to avoid the server becoming overloaded or Activity Feed events being delayed.
- Heavy usage of background processing may require additional CPU resources
- Heavy usage of customer batch jobs may require additional resources.

Assumptions:

- Moderate Logi / Action Center Usage / background processing
- Moderate Entity Field Tracking on calculated fields (non-database fields)

CPU Sizing AI Agents

On top of real user sizing, additional CPU Cores are required for AI Agents with **1 additional CPU Core per concurrent agent requests** from AI Agents.

Memory Sizing

All values below are in gigabytes (GB)

QAD Version	< 50	< 100	< 200	< 400	< 800	< 1000	< 2000
QAD Adaptive	50	66	82	130	170	224	274

- **NOTE 1:** Increasing OE12 parameters related to ABL sessions per agent may require more RAM
- **NOTE 2:** Figures do not include sizing for any other applications running (eg monitors, logstash, anti-virus, icinga etc)
- **NOTE 3**: Important: Assumes the correct tuning has been applied to the correctly sized server. This is particularly important for PASOE appservers, as increasing the number of PASOE Multi-Sessions Agents can have a very large extra impact on RAM demands.
- **NOTE 4**: Larger RAM may be required to support large databases and buffer pool memory allocation

Disk Sizing

The following requirements are a minimum for designing the storage requirements for QAD Adaptive. The actual size of the disks (how many GB of storage) is not listed, because in general the number of disks required will provide more than enough storage.

Note that for some larger customers who have corporate storage solutions, the customer will typically request "an amount of storage in gigabytes" with the backend hardware details abstracted. This is known as virtualized storage. When sizing for customers with virtualized storage, you can ignore this disk requirements section altogether.

- Disk Technology : SAS, Ultra-320 SCSI, Ultra-640 SCSI, or Fiber Channel only (earlier SCSI standards / ATA / SATA are not recommended)
- Disk speeds : 15000 (15k) RPM minimum
- Hardware RAID level : RAID 10 preferred

Note: RAID 5 and RAID 6 are strongly discouraged.

In the following sizing table, it is assumed that RAID 10 will be used.

User Count	Storage Controllers	Disks Required	Approx IOPS Required
< 50	1	6	200
< 100	1	8	500
< 200	2	10	800
< 400	2	12	1500
< 800	2	16	2500
< 1000	3	20	2800
< 1500	4	24	3400
< 2000	4	30	5000

Network Sizing

The following requirements are averaged across the entire user community. This means that at any given time some users are using trivial amounts of bandwidth, while others may be using very high amounts. For example, a user running a browse can easily consume 500 kbps for the duration.

- QAD Adaptive: 30 kbps per user

QAD Adaptive should use a **nginx with compression enabled.**

Client PC / Laptop Sizing

Client sizing is considerably easier than it was in the past, due to the rapid growth in the power of modern desktop and laptop computers.

QAD Adaptive Client

- The general recommendation is that client machines should be no more than four years old for running QAD Adaptive.
- The faster the CPU the better for QAD Adaptive due to client side JavaScript processing.

User Count Sizing versus Throughput Sizing

When trying to determine the sizing for a QAD Adaptive deployment, the expected user counts may not be known. Another method of determining sizing is based on the document throughput (creation) rate.

You can use the following conversion rates to translate throughput into equivalent user counts :
100 concurrent users = 700 documents created per hour, where a document is a sales order, purchase order, invoice, shipper, work order, journal, or similar.

Delta Sizing

Delta sizing refers to sizing from technical data from an existing customer deployment.

Delta sizing is most often used when:

1. Upgrading from one version of QAD to QAD Adaptive
2. Adding additional user counts
3. Migrating to a new server

Delta sizing is the most accurate sizing method and should be used where possible.

Upgrading User Counts

When upgrading user counts, use a simple mathematical equation:

- Additional resources required = current resources + (current resources * % new users)

For example, if you want to increase users by 35%:

- Additional resources required = current resources * 1.35

Appendix A: Understanding the Sizing Process

Understanding the Customer's Requirements

While the sizing process ultimately produces a technical specification of hardware requirements and their deployment topology, sizing is all about meeting the customer's requirements.

To this end, much of the sizing process involves human interaction, discovery, workshopping, and meeting the requirements of the business.

This section contains information about some of the business considerations that influence and guide the sizing process.

Requirements Gathering

In order to understand the customer's requirements, you must engage in a requirements gathering process.

The Questionnaire

A common way to determine customer requirements is through a questionnaire. See the sample questionnaire for details. This is probably the simplest way to gather requirements, but also tends to have the longest lead time from project start to final sizing specification, due to necessary follow-up questions.

The Sizing Workshop

A sizing workshop is the most effective and accurate way to generate the final sizing for a QAD deployment. Sizing workshops tend to be used with very large, complex deployments with a budget for formal sizing.

The sizing workshop typically gathers all relevant experts from QAD (or a QAD partner) and the company deploying the QAD software. This group of experts can include systems and database administrators, network and security architects, project managers and sponsors, hardware vendors, and QAD performance SMEs.

The workshop is an intense, interactive experience tasked with coming up with a sizing based on requirements and agreement from all stakeholders.

Performance Monitoring

Forrester Research has reported that among companies with revenue greater than 100 billion USD, nearly 85% report significant application performance degradation ("best practices in problem management"). **Nearly 85% of corporate applications are failing to meet and sustain their performance requirements**.

Poor application performance leads to lost productivity, lost confidence, lost credibility in the QAD brand, potentially reduced revenue, and possible financial penalties.

In order to sustain required performance levels, performance monitoring is critical. Unless performance is actively monitored and trended, proactive problem solving and performance management is extremely difficult.

Security

Some customers, especially larger ones, may have a security architecture team, or a security policy in place that dictates how the QAD Adaptive can be deployed. Hosting Data Centers also may have such policies.

On-Premise / Hosted / Cloud Considerations

The possible sizing scenarios can depend on whether the customer wishes to host on-premise, via a data center, or through a cloud-based arrangement (QAD Cloud).

The customer should be asked where the hardware will be hosted.

In the context of where the hardware will be hosted, there should be follow-up questions that ask what restrictions this choice entails.

High Availability / Disaster Tolerance

An open discussion should be had with the customer emphasizing the importance of highly available and disaster tolerant QAD deployments.

QAD recommends that all customers run with high availability at a minimum, and strongly recommend the implementation of a disaster-tolerant deployment.

The Sizing Process

The sizing process is not a single “this is the way we size” methodology. There are many paths to a final sizing.

Sizing could be as quick and simple as sizing to a budget (“buy the fastest server for xyz dollars”) to an involved workshop involving QAD sizing experts.

This section details the common approaches, and explains when they might be appropriate.

Using the Quick / Ballpark Sizing

If the sizing request has a short turnaround time, or there are few facts or data to support more in-depth sizing exercises, then the quick / ballpark sizing should give a reasonably accurate sizing with a minimal amount of effort and invested time.

Sizing to a Budget

If there is a limited budget, then sizing to a budget can make things simple. Determine the minimal set of hard requirements for the deployment, and then choose the most powerful and acceptable hardware to meet these requirements, within the notional budget.

Sizing by Reference Installations

Another common method of sizing is “sizing by reference installations.” Chances are that a customer that wishes a properly sized system has similar characteristics to one or more recently deployed QAD systems.

Contacting the wider QAD technical consulting audience with questions about reference customers may provide examples that can be used to validate or lay the foundation for a proposed sizing.

Delta Sizing / Upgrade Sizing

Delta sizing is useful when the customer is already running QAD software.

Delta sizing uses the notion that the most accurate sizing is gained from using the most accurate data. If the customer is already running QAD software and the proposed upgrade is an extension of this, then data from the running system is the most accurate.

To carry out effective delta sizing, tools such as the QAD Monitoring Framework are used to analyze existing resource usage patterns, and the new sizing is some multiple of the current system.

For example, if the current system is running at 70% average CPU Utilization and there is an expected user count growth of 50%, then the new server must have 50% greater CPU capacity to provide the same level of performance with the higher user count. If the current system has 8 cores, then the new system will require 12 cores (50% more).

Network

An often overlooked part of the sizing process is the network. While it is unlikely that the modern gigabit networking cards on modern servers will be overloaded, it is possible that geographically remote users could be affected by undersized network links.

Appendix B : Sample Sizing Questionnaire

Existing QAD System

Provide details of existing software installations

Provide details of the current QAD deployment, including all non-QAD components on QAD servers.

For each server in the QAD deployment, provide the following information :

- Server Model, number of CPUs, RAM, attached disk drives and technology, RAID setup
- Whether the server is virtualized or not. If virtualized, provide details including version number of the hypervisor
- List of QAD software products installed, including:
    - The number of instances of Tomcat
    - Related databases and database size
    - Approximate concurrent connections and rate of growth, if known
    - If the server is highly available
    - If the server is disaster tolerant
    - Where the server is geographically hosted
- Include details on network or SAN storage

Customizations

- List all major areas of customization to the QAD software.

Current Workload and Performance

- At a high level, list the current user counts and/or transaction rates.
- If known, list the approximate rates during peak and non-peak periods.
- List any areas of concern or performance problems with the current QAD deployment. Objective measurements are more useful than qualitative data.

Provide details of new and future requirements (this sizing exercise)

- Number of peak concurrent users and expected growth rate
- Number of non-production systems required (development, test, QA, etc) and whether or not it is expected these will be on the primary production system.

Interfaces

- List all interfaces to third-party systems.

  - Type of interface (REST, QXtend, EDI, FTP, HTTP, etc)
  - Approximate documents exchanged per hour during peak periods of interface use
  - Whether peak periods of interface use occur during normal working hours for interactive users

Preferred Deployment

- Provide as much information as possible about the preferred deployment scenario.

Hosting

- Is the system to be hosted on-premise, at a hosting center, or via QAD On Demand?
- What operating system, if any, is preferred?
- What hardware vendor, if any, is preferred?
- Is there a preference for blade or rack servers, if relevant?
- Do you have your own blade enclosures? If so, list the make and model, and number of spare slots.
- Is there a preference for virtualized environments versus physical hardware?

High Availability and Disaster Tolerance

- Is the system to be highly available (servers are redundant with shared storage and a failover mechanism)?
- Is the system to be disaster tolerant (servers are near-realtime replicated to a remote hosting site)?

Users

- What are the expected user counts (average and peak), and their geographical deployment?
- Is there a preference to use thin client technologies such as Citrix for remote users?
- If so, is there an existing thin client server farm that can be used?

Appendix C: Revision History

Version 3.0

- Initial, November 2012
- Derek Bradley (dbb@qad.com)

Version 3.1

- Last Updated April 2013 - Fixed Citrix sizing table
- Derek Bradley (dbb@qad.com)

Version 3.2

- Last Updated April 2016 - Updated for 2016 Release
- David Hughes (djh@qad.com)

Version 3.3

- Updated April 2018 to include Channel Islands, remove some old material and a general restructure of the document
- Feedback from David Hughes Included
- Eddie O'Sullivan (eos@qad.com)

Version 4.0

- Updated October 2018 to include Channel Islands, remove some old material and a general restructure of the document.
- Adjusted Channel Islands sizing for Query Service and additional information on Activity Feed
- Adjustments from David Hughes Included
- Eddie O'Sullivan (eos@qad.com)

Version 4.1

- Updated to March-2019 Release
- Minor changes included to mention cpu score for workbenches
- Adjust activity feed cpu comments based on non database fields based on new coding

Version 4.2

- Updated to March-2020 Release

Version 5.0

- Updated to March-2020 Release
- Introduced QAD AUX to replace Channel Islands

Version 6.0

- Updated based on review from Bridget Schmidt

Version 7.0

- Final version for March-2020

Version 8.0

- Migrated from March-2020 to Sept-2020 (Increase memory sizing from Nifi, Kafka and Zookeeper)

Version 9.0

- Added important note to memory sizing
- Minor adjustments
- Final version for March-2021

Version 10.0

- Minor adjustments for Sept-2021

Version 11.0

- Minor adjustments for March-2022

Version 12.0

- Adjustments to include PASOE
- Minor adjustments

Version 13.0

- Adjustments to memory sizing for Java17

- Applied new QAD template

- Minor adjustments in some areas

Version 14.0

- Removed a number of Appendix
- Fixed naming conventions
- Removed specific references to the impact of OE12
- Removed a number of sections to simplify this document

Version 15.0

- Initial version for Adaptive 2608
- To Be Added AI Agent Sizing