

Architecture and Server Sizing for IBM Cognos Controller 8.5

Nature of Document: **Guideline** Product(s): **IBM Cognos Controller 8.5** Area of Interest: **Infrastructure**





Copyright and Trademarks

Licensed Materials - Property of IBM.

© Copyright IBM Corp. 2010

IBM, the IBM logo, and Cognos are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

While every attempt has been made to ensure that the information in this document is accurate and complete, some typographical errors or technical inaccuracies may exist. IBM does not accept responsibility for any kind of loss resulting from the use of information contained in this document. The information contained in this document is subject to change without notice.

This document is maintained by the Best Practices, Product and Technology team. You can send comments, suggestions, and additions to <u>cscogpp@ca.ibm.com</u>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Microsoft, Windows, Windows NT, SQL Server, Terminal Services and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Citrix and all citrix-based trademarks and logos are trademarks of Citrix Systems, Inc.



Table of Contents

1	Introduction5		
	1.1	Purpose5	
	1.2	Applicability5	
	1.3	Exclusions and Exceptions5	
2	Contro	oller Architecture Overview6	
	2.1	Overview of the main architecture components	
	2.2	Description of the sub-components	
	2.3	Benefits and Drawbacks of a Multi-Server design9	
3	Contro	oller Architecture Overview11	
	3.1	VMWare (and other virtual platforms, such as ESX / Virtual PC11	
	3.2	64 Bit Operating System11	
4	Softwa	are Requirement13	
	4.1	Purpose	
	4.2	Microsoft Components	
	4.3	Other Third Party Components	
	4.4	Database Server14	
5	Softwa	are Requirement15	
	5.1	Server ⇔ Server Backbone15	
	5.2	Client to Gateway (application server) connection15	
6	Datab	ase Storage Sizing16	
	6.1	Database Storage (drive/volume area) Sizing16	
	6.2	Database driver arrays16	
7	Hardw	vare recommendations for individual boxes17	
	7.1	Overview	
	7.2	Hyperthreading17	
	7.3	Multi-core CPU Chips	
	7.4	Windows 2003 Server CPU capabilities/licensing18	
	7.5	Typical Database Server Specifications	
	7.6	Typical individual Application Server Specification21	
	7.7	Typical individual Citrix/Terminal Server Specifications22	
	7.8	Typical end user client PC specifications	
8	Muli S	erver (number) recommendations24	
	8.1	Overview24	
	8.2	Example #1 - up to and including 15 Concurrent Users25	
	8.3	Example #2 - up to and including 40 Concurrent Users25	



Architecture and Server Sizing for IBM Cognos Controller 8.5

	8.4	Example #3 - up to and including 75 Concurrent Users	.26
	8.5	Example #4 – Approximately 75 to 150 Concurrent Users	.27
	8.6	Example #5 – Large Implementations 150+ Concurrent users	.28
9	Using	optional IBM Cognos Controller features	.28
	9.1	Overview	.28
	9.2	Data Marts (RDBMS / OLAP)	.29
	9.3	Financial Analytics Publisher (FAP)	.30
	9.4	Integration with IBM Cognos Products	.33
10	Real-L	ife Example	.33
	10.1	Customer Example #1 – 100 to 120 Concurrent Users	.33
	10.2	Customer Example #2 – 100 to 125 Concurrent Users	.34
	10.3	Customer Example #3 – 30 Concurrent Users	.34



1 Introduction

1.1 Purpose

This document is written by an IBM Cognos Support analyst. It is intended to give an overview on the technical architecture of the Controller system, and then to describe best practices on server sizing (based on customer feedback) for IBM Cognos Controller 8.5.

By following these "best practices" the intention is to make customers as successful as possible, by ensuring that they have the correct server architecture/ infrastructure in place for their system.

1.2 Applicability

IBM Cognos Controller 8.5

1.3 Exclusions and Exceptions

Here are the following exceptions and exclusions to this document

- 1. This document is based on feedback to IBM Cognos Support.
- 2. Official documentation that comes with the product (for example ctrl_arch.pdf 'IBM Cognos 8 Controller ARCHITECTURE AND DEPLOYMENT GUIDE') takes precedence over this Proven Practice document
- 3. Although this document demonstrates proven practices suitable for *most* environments, it is not necessarily perfect for *all* environments.
- 4. There are an infinite variety of possible customer I.T. environments, many different ways to install/ configure Controller 8.5, and therefore the advice in this document may have to be modified by the reader to fit in with their needs/environment.

This document is only intended to be a basic guide. Customers are recommended to always **employ an experienced Cognos technical consultant** to analyze their exact requirements. Your consultant will produce an architecture which best fits for your exact requirements



2 Controller Architecture Overview

2.1 Overview of the main architecture components

The entire Controller 8.5 architecture is explained in detail inside the file 'ctrl_arch.pdf' ('IBM Cognos 8 Controller - ARCHITECTURE AND DEPLOYMENT GUIDE') which is found on the installation CD-ROM.

However, to summarise, here is an overview description of the common components that are involved:





The three main components are:

- 1. Gateway (front-end 'portal' / proxy)
- 2. Web Services Server (main COM+ system)
- 3. Report Server (Cognos 8 BI reporting engine)

The diagram gives an example of what happens when a user runs a typical request (in this example a report) in IBM Cognos Controller:

- 1. The user clicks a report to run it, and the request goes to the gateway, which forwards the request to Controller Web Services Server. If a gateway is not part of your installation, the request is sent directly to Controller Web Services Server.
- 2. Controller Web Services Server forwards the request to the Controller COM+ application for processing.
- 3. The COM+ application prepares the data in the Controller database for Report Server to retrieve later. To prepare the data, the COM+ application inserts data for the report into the dedicated tables created during installation in the Controller database. When Report Server generates the report, SQL queries are run against these tables.
- 4. The COM+ application retrieves the report template, which is an XML file stored on the Controller Web Services Server computer. The COM+ application updates the report template based on selections made by the user when requesting the report. Updates include modifications to the data source, formatting, and SQL queries.
- 5. The modified report template is sent to IBM Cognos Controller (client).
- 6. IBM Cognos Controller (client) makes the request to the Report Server dispatcher. The request also specifies whether a PDF or HTML report is required.
- 7. The presentation service sends the request to the report service.
- 8. The report service uses the SQL queries in the report template to retrieve data from the dedicated tables for the report in the Controller database.
- 9. The report service returns one of the following results to the presentation service, and then to IBM Cognos Controller (client):
 - an error page
 - a not ready page
 - a page of an HTML or PDF report, depending on which format was requested, for display in Cognos Viewer

2.2 Description of the sub-components

As described earlier, the three main architectural components are:

- 1. Gateway (front-end 'portal' / proxy)
- 2. Web Services Server (main COM+ system)
- 3. Report Server (Cognos 8 BI engine)

However, each of these contain separate 'sub-components'. Below is a more detailed picture of the **sub-components** that form the architecture:



Information Management



2.3 Benefits and Drawbacks of a Multi-Server design

Of course, it is *possible* to place all the Controller system components onto one single physical server, or instead place each individual sub-component on its own dedicated/individual physical server.

Naturally, each customer must (in conjunction with their Cognos Technical Consultant) decide for themselves what is the best server architecture. For example, they should consider the following factors:

Fewer servers	More servers
Lower hardware cost	Higher hardware cost
Easier to install	More complex to install
Easier to document/understand/maintain	Harder to document/understand/maintain
Less scalable / slower performance	More scalable / higher performance
Less flexible	More flexible (for example can tailor to
	individual requirements, such as DMZ)

The customer should balance all the factors (such as the ones listed above) to choose their favoured architecture design.

Here are the following exceptions and exclusions to this document

- I) This document is based on feedback to IBM Cognos Support.
- II) Official documentation that comes with the product (for example ctrl_arch.pdf 'IBM Cognos 8 Controller - ARCHITECTURE AND DEPLOYMENT GUIDE') takes precedence over this Proven Practice document
- III) Although this document demonstrates proven practices suitable for *most* environments, it is not necessarily perfect for *all* environments.
- IV) There are an infinite variety of possible customer I.T. environments, many different ways to install/configure Controller 8.5, and therefore the advice in this document may have to be modified by the reader to fit in with their needs/ environment.

This document is only intended to be a basic guide. Customers are recommended to always employ an experienced Cognos technical consultant to analyze their exact requirements. Your consultant will produce an architecture which best fits for your exact requirements



TIP: Although splitting the Controller application components over several servers can give higher scalability, please do not make the mistake of using underpowered servers.

For example, imagine a customer who decided that they wanted to split the Controller system over 4 servers:

Server 1 = Gateway Server 2 = WSS

Server 3 = Report Server

Server 4 = Consolidation Server

However, each server was a VMWare server, with only 1 CPU and 1Gb RAM allocated. Therefore (in total) there was only 4 CPU cores and 4Gb RAM available.

- ⇒ In total, therefore, the solution gave no more performance benefit than hosting ALL the components on a SINGLE server, with 4 cores and 4 Gb RAM
- ⇒ In this case, all the customer had achieved was to massively increase the complexity of the system, therefore increasing time/effort required to (a) install (b) upgrade (c) support (troubleshoot) the system. There would be no performance benefit (and quite probably plenty of performance drawbacks) to this design.



3 Controller Architecture Overview

TIP: You can also find useful links (on subjects such as 64-bit and VMWare compatibility statements) here: <u>http://www-01.ibm.com/support/docview.wss?rs=3528&uid=swg27014782</u>

3.1 VMWare (and other virtual platforms, such as ESX / Virtual PC

Virtual platforms (such as VMWare) are designed to be transparent to the software that runs on them. In other words, just like any other piece of software, naturally Controller will run OK on virtual platforms.

IBM Cognos' official position on VMWare is stated here: <u>http://download.boulder.ibm.com/</u> ibmdl/pub/software/data/cognos/support/en/products/vmware_statement.pdf

To summarise IBM Cognos' position:

- 1. IBM Cognos supports customers who deploy Controller on a virtual platform
- 2. However, in the unlikely event that the virtual platform is the specific cause of a problem (it cannot be reproduced on a non-virtual server), we reserve the right to ask the customer to contact their virtual platform vendor (for example VMWare) to explain/solve the problem
- 3. There is always going to be a performance overheard when using a virtual platform (compared with using equivalent 'real' server hardware).
 - 1. **If configured correctly**, this overhead **may** be so small that it is barely noticeable to the end user
 - 2. However, a **badly-configured** virtual server may give very **poor performance**. In this case, the customer must work with their virtual platform vendor to solve the performance problem.

3.2 64 Bit Operating System

IBM Cognos' official position on 64-bit compatibility (for Controller 8.4 FP1 onwards) is stated here:

http://download.boulder.ibm.com/ibmdl/pub/software/data/cognos/support/en/products/Cognos84_64-bit_Jan2009.pdf

From version 8.4 FP1 onwards, the Controller software was changed to a "true 32 bit" application (instead of an "AnyCPU" application).

NOTE:

- 1. All of the Controller code is still 32-bit
- 2. In other words, Controller is entirely a 32 bit application.

Cognos.

software

This enables the Controller software to run in a 64 bit server environment. However, since Controller is a 32-bit application, please be aware that there is no performance advantage in running Controller 8.5 on a 64-bit platform.

IBM Cognos recommends that the only x64 platform (where customers should consider running Controller client code) is a x64 Citrix (or Terminal Services) server.

1. This has been tested by IBM Cognos R&D and is supported

The reason why customers may wish to use x64 Citrix servers, is to increase capacity more cheaply. The maximum amount of RAM that Windows 2003 Server supports (<u>http://msdn.microsoft.com/en-us/library/aa366778%28VS.85%29.aspx#physical memory limits windows server 2003</u>) is:

- 1. 32-bit STANDARD edition 4Gb RAM
- 2. 64-bit STANDARD edition 16Gb RAM
- 3. 32-bit ENTERPRISE edition 32Gb RAM
- 4. 64-bit ENTERPRISE edition 64Gb RAM

Therefore, for customers who prefer to have fewer (but more powerful) Citrix servers, they may prefer to use fewer 64-bit 'Standard' Edition servers, since it is cheaper (in Microsoft server license terms) than (a) multiple 32-bit 'standard' edition servers and (b) using expensive 32-bit 'Enterprise' edition servers.

- 1. In other words, customers may receive a cost benefit (cheaper scalability) when delivering the Controller client using 64-bit Citrix/Terminal servers
- 2. However, please remember that they will NOT receive any *performance* benefits (Controller does NOT run quicker on a 64-bit server).

To summarise IBM Cognos' position:

- 1. IBM Cognos actively supports customers who use a 64-bit version of a supported database server, to host their Controller database
 - 1. For example, IBM Cognos supports the use of 64-bit Microsoft SQL 2005 and 2008, plus Oracle 10G and Oracle 11.
- 2. IBM Cognos actively supports customers who use a 64-bit version of a supported client operating system, to run the Controller client
 - 1. For example, IBM Cognos supports the using a Citrix server which is based on Windows 2003 Server 64-bit edition
- 1. Based on the claims of compatibility that Microsoft give regarding their 64-bit operating systems, IBM Cognos R&D believe that a Controller 8.5 application server can successfully be used when hosted on a 64-bit Microsoft operating system.
 - 1. IBM Cognos therefore supports 64-bit operating systems on a 'compatible' (non-active) level
 - 2. For example, IBM Cognos R&D believes that Controller 8.5 application server will work when using a Windows 2003 Server 64-bit edition server.

- 2. However, because there will be no performance benefit of running the 32-bit Controller server code on a 64-bit Windows server, IBM Cognos do not recommend that customers should use a 64-bit version of a supported server operating system, to run the Controller application server
 - 1. For example, IBM Cognos do not recommend that customers install the Controller 8.5 application server components on Windows 2003 Server 64-bit edition

4 Software Requirement

TIP: The official Controller 8.5 current supported environments are listed here: <u>http://www-01.ibm.com/support/docview.wss?rs=3450&uid=swq27017475</u>

4.1 Purpose

To summarise the official web page, the most important/relevant sections are:

- 1. your server should be running Windows 2003 Service Pack 2
 - 2. Other operating systems and service packs are supported, but the most 'actively' tested (therefore recommended) server environment is Windows 2003 Server Enterprise edition SP2.
- 3. your client device should be running Windows XP Service pack 3 or Vista SP1
 - Other operating systems and service packs are supported, but the most 'actively' tested (therefore recommended) client environments are WinXP SP3 and Vista SP1.

4.2 Microsoft Components

To summarise the official web page:

- 1. your server should be running:
 - 1. Microsoft .NET Framework 2.0 Service Pack 1 or later (SP2 recommended)
- 2. your client device should be running:
 - 1. Microsoft .NET Framework 2.0 Service Pack 1 or later (SP2 recommended)
 - 2. Microsoft Excel (2002) or ideally 2003 or 2007, each with the latest service pack
 - 3. Microsoft Internet Explorer (6) or 7 (recommended)

4.3 Other Third Party Components

Summarise the official web page:

- 1. your server and client device should be running:
 - 1. Adobe Acrobat Reader 8.1.2



4.4 Database Server

To summarise the official web page:

- 1. your **database server** should be running
 - 1. MS SQL 2005 or SQL 2008 (latest service pack recommended)
 - 2. or Oracle 10G (Release 2 recommended) or Oracle 11G (Release 1)
 - 1. If using Oracle, then please be aware that:
 - 2. Cognos Controller depends on components that are only packaged with Oracle *Enterprise* Edition.
 - 3. Using Oracle is not supported for Traditional and Simplified Chinese locales.
 - 4. If using Oracle 10G, then Oracle patch 5473334 is required.
- 2. If using Microsoft Analysis Services (OLA) then your MSAS server should be running
 - 1. MSAS (SQL) 2000 or MSAS (SQL) 2005 (latest service pack recommended)

5 Software Requirement

5.1 Server \Leftrightarrow Server Backbone

The network connections between your application server(s) and database server should be 100MB Full Duplex as minimum. However, some processes will benefit from large performance increases if you use Gigabit (1Gb) connections between the servers (especially the Web Services Server \Leftrightarrow database server connection).

5.2 Client to Gateway (application server) connection

Some of the performance (of the end user's experience) will be affected by the speed of the network between his/her client device (for example, their desktop PC) and the application server (gateway). Naturally, the faster the network connection, the faster the experience.

Bandwidth and network latency:

The bandwidth requirement (client ó application server) for Controller v8 is approximately 256kbps per concurrent user, for "standard" (basic) usage of Controller.

NOTE:

- 1. this (256kbps) is not an 'absolute' / 'fixed' value
- 2. there will be performance benefits if you have a larger available bandwidth (e.g. 512kbps)
- 3. however, from the tests that we have run, there is not a huge difference between 256kbps and 512kbps
 - 1. i.e. it's not twice as fast when running with twice the available bandwidth!)
 - 2. instead (depending on what you are doing of course) you may see something like (for example) a 20% performance increase with twice the bandwidth
- 4. more complex tasks (in Controller) would certainly benefit from there being bandwidth available >512kbps

Also, it is not just bandwidth size that is important – network latency is just as important. There should be a quick (e.g. <80ms) round-trip ("ping") network latency between client and server, for best performance.

Even though remote network links (a.k.a. "WAN" links) are getting faster and cheaper every year, many customers will want to try to make as efficient use of their remote network as possible. Since Citrix gives extremely good (e.g. 10x) network bandwidth savings (over the traditional client⇔server deployment of Controller 8), they may choose to deploy via Citrix (or its 'little-brother' – Microsoft Terminal Services).

To summarise, as a general rule of thumb:

- 1. Local (same LAN as servers) users will not notice much benefit between 100Mb and 1Gb network connections
- 2. Remote (WAN) users will generally find their Controller performance is acceptable if the (per concurrent user) bandwidth (client ó application server) is at least:
 - 1. 256kbps for 'standard' (for example data-entry) users
 - 2. 512kbps for 'advanced' (for example super-user) users



and also that their round-trip ("ping") network latency is less than approximately 80ms

If such suitable network connections are not available, then the typical solution is to utilize Citrix and/or Microsoft Terminal Services to deploy Controller

6 Database Storage Sizing

6.1 Database Storage (drive/volume area) Sizing

The following information is based on Microsoft SQL, from customer feedback to IBM Cognos support :

A typical customer's Controller database will be approximately 2Gb -> 10Gb in size

• The maximum size seen is approximately 30Gb or so

The finance department would almost certainly benefit from having several databases (for example "live", "test", "training" etc.)

• In addition, you will want to scope for "room to breath"

=> this suggests you should have a total of at least 144Gb storage for your data (MDF) partition.

6.2 Database driver arrays

- 1. Some customers prefer a single RAID5 array
- 2. However, the author suggests that ideally 3x RAID1 (144Gb+ each array) would be faster/more flexible

1. [One array for each of "system/tempdb", "data" and "transaction logs"]

TIP: For optimal speed, ensure that each hard drive spins at 15k rpm (not cheaper 10k rpm drives)



7 Hardware recommendations for individual boxes

7.1 Overview

In general, this document shall assume that all servers (including the database server) are dedicated to Controller-only use. This is to give the Controller system maximum performance/flexibility

Also, we shall often talk about the maximum number of "concurrent" users. This is different from the 'total' number of 'named' users (which is the entire number of users that could ever *possibly* use the system)

- 1. This maximum "concurrent" user number will be a *proportion* of the 'named' users
 - 1. this proportion will vary between customers, because of where its divisions are located (time-zones etc.)
 - 2. typically, however, we often see that the maximum concurrent users is approximately 40% of the total 'named' users
 - 3. e.g. if had 100 'named' users, then at most we could expect to see 40 users actively using the system at any one time

7.2 Hyperthreading

Older generations (approximately 2003 -> 2006) of servers (based on older Intel Xeon CPUs) were single-core, but with HYPER THREADING

- 1. HYPER THREADING makes it "appear" that there are two CPUs in each physical chip
- 2. *However* hyper threading is known to give poor performance (and stability problems) with applications that give very high CPU load (such as Controller)
- 3. Therefore IBM Cognos have always recommending disabling HYPER THREADING (if using older Intel Xeon chips)
- 1. Some customers prefer a single RAID5 array
- 2. However, the author suggests that ideally 3x RAID1 (144Gb+ each array) would be faster/more flexible

1. [One array for each of "system/tempdb", "data" and "transaction logs"]

2.

TIP: For optimal speed, ensure that each hard drive spins at 15k rpm (not cheaper 10k rpm drives)

Cognos.

software

7.3 Multi-core CPU Chips

However, in recent years, both Intel and AMD have moved their CPU technology to multi-core (initially dual-core, but now quad-core and beyond!) technology

- 1. Although no official tests have been performed, feedback from customers suggests that multicore CPUs (e.g. modern AMD and new Intel chips) can significantly outperform the older Intel Xeon chips.
- 2. Therefore, we 100% recommend that customers use the latest/fastest multi-core chips on the market when they come to purchase new Controller servers

IMPORTANT NOTE: There are multiple different brands of CPU chips (for example Intel/AMD etc). Each brand has a huge number of different chip models/families, with different features (for example single-core, dual-core and quad-core) and different CPU clock speeds

- 1. It would be impossible for the author to make a detailed comparison of all the different brands/models of CPUs
- 2. Therefore (for the sake of simplicity) when reading this document, please be aware that the author shall generally servers based on the number of "physically separate" chips (as opposed to the number of cores)
- 3. For example, I shall typically recommend a server containing 2 x physically-separate CPUs (separate chips). Assuming that each of these are 'dual-core' chips/CPUs, this will give you a server with 4 separate cores

Note:

General customer feedback suggests that customers would be <u>unwise</u> to assume that 1x quadcore CPU will perform as fast as 2x dual-core CPUs or 4x single-core CPUs. Despite this, multi-core CPUs are generally (of course) recommended over single-core CPUs.

7.4 Windows 2003 Server CPU capabilities/licensing

IBM Cognos cannot provide any official statements on third-party (Microsoft) products. However, third-party websites (for example <u>http://www.itnewsgroups.net/group/</u> <u>microsoft.public.windows.server.general/topic31536.aspx</u>) state Microsoft's position as:

Windows Server is licensed per physical processor, irrespective of how many cores are on the processor.

Standard Edition Windows support 4 processors. These processors can be single, dual or quad core today. Windows will use all cores on all processors that it supports.

For example, if a server has 4 processors, and each is dual core, the customer can install Windows Server Standard Edition and it will use all 8 cores. Task manager will show 8 CPUS and applications can leverage all 8.

Another example, if the server had 4 Quad Core CPU's installed, then the customer could install Windows Server Standard edition and use all 16 cores

18

7.5 Typical Database Server Specifications

However, in recent years, both Intel and AMD have moved their CPU technology to multi-core (initially dual-core, but now quad-core and beyond!) technology

In the author's experience, customers generally deploy Controller onto multiple identically (hardware) specified application servers. In other words, imagine a customer who has decided to have 2 application servers & 4 Citrix servers. It is very typical that all 6 of these servers have the exact same hardware specification.

Typically, the only server which has a different (from the application server, for example) recommended specification is the database server. This is because (inside a Controller server environment) we often find that it is the database server which is under the most load compared to all the other servers.

Controller 8.5 supports Microsoft SQL 2005 and SQL 2008.

- 1. Both SQL 2005 and SQL 2008 have
- 2. a maximum number of CPUs of 4, for the 'standard edition'
- 3. and a memory limit the exact same as the operating system's memory limit (for example 4Gb for 32-bit standard edition)
- 4. Typically, the CPU is more of a performance bottleneck (compared to memory) on the database server
- 5. Therefore, for many customers, SQL 2005/2008 Standard edition is sufficient for their needs, despite the fact that the 32-bit edition is limited to 4Gb RAM
- 6. For some advanced customers, Enterprise editions and/or 64-bit editions (which allow the use of more CPUs and more memory) may be useful/necessary.
- 7. Naturally, it is recommended that the latest Microsoft patches (e.g. SQL 2005 SP3 and SQL 2008 SP1) are applied.

Cognos.

software

Controller		Total			
System	Separate	Number	Clock		SQL
Size	CPUs	CPU cores	Speed	RAM	Edition
small	2	4 (or 8)	2.4GHz+	4Gb (32bit) or 8Gb+ (64-bit)	Standard
medium	4	8 (or 16)	2.4GHz+	4Gb to 16Gb	Std/Ent
large	4 to 8	16+	2.4GHz+	16Gb	Enterprise

To summarise the typical database server recommendations:

TIP: The above suggestions are based on a balance between sensibly-priced and reasonably powerful server hardware, as of January 2010. More RAM, more CPU cores and higher clock speeds will always provide more performance, but feedback suggests that the above specifications are a sensible balance.

Here is how systems are defined for IBM Cognos Controller

- 1. Small Controller systems
 - 1. Relatively low (approximately 40 or fewer) concurrent users
 - 2. Limited reporting functionality
- 2. Medium Controller systems
 - 1. Between approximately 50 to 100 concurrent users, with limited reporting functionality
 - 2. Or alternatively, fewer users but advanced reporting functionality.
 - 1. For example, multiple concurrent users using Controller reports to perform budgeting, which often means that 12+ periods of data are used (instead of 'typical' reports with 1-3 months of data are used) in each report, causing high SQL CPU% usage)
- 3. Large Controller systems
 - 1. For example 100+ concurrent users

Controller also supports Oracle 10G and 11G. These products can be run on non-Microsoft operating systems with non-Intel/AMD x86 hardware. Due to these added complexities, this document shall not refer specifically to non-x86 hardware. Instead, suffice to say that customers (who want to host their Controller system via an Oracle database) should ensure that their non-x86 hardware is equivalent to the x86 hardware that is described above.

7.6 Typical individual Application Server Specification

It is perfectly possible to have multiple application servers, each with a completely different hardware specification. For example, server#1 may have 1x dual-core CPU, server#2 may have 2x dual-core CPUs and server#3 may have 1x quad-core CPU.

However, in general, most customers typically purchase multiple 'identical' application servers (i.e. each based on the same hardware specification). Therefore, this document will assume that the customer wishes each of their Controller application servers to be identical hardware.

=> In other words, the only decision (that the customer has to make) is not what specification their application servers should be, but (instead) how many application servers they need.

The recommended hardware specification (for each separate application server) would be:

Physically					
Separate	Total number		Operating		Hard
CPUs	CPU cores	CPU speed	System	RAM	Drive
2	4 or 8	approx 2.4GHz+	Win2003 Server	4Gb	1x RAID1
					72Gb



7.7 Typical individual Citrix/Terminal Server Specifications

In general, most customers typically purchase multiple 'identical' Citrix/TS servers (i.e. each based on the same hardware specification).

Naturally, each and every customer uses Controller in a slightly different way. As a ballpark estimate, customer feedback suggests that Citrix servers can host approximately 20 concurrent Controller sessions per 4 (or ideally 8) CPU cores and 4Gb RAM. 64-bit servers can make use of more RAM than 32-bit versions, and are therefore sometimes more cost-efficient for some customers.

Therefore, the recommended hardware specification (for each separate Citrix server) should be

Physically						Approx
Separate	Total number		Operating		Hard	Users
CPUs	CPU cores	CPU speed	System	RAM	Drive	per server
2	4 (or ideally 8)	approx 2.4GHz+	Win2003 Server	4Gb	1x RAID1	20
			32-bit		72Gb	
4	8 (or ideally 16)	approx 2.4GHz+	Win2003 Server	8Gb	1x RAID1	40
			64-bit	(or 16Gb)	72Gb	

=> If you had a maximum of 50 concurrent Controller users, and you followed the above guidelines, you should have the following number of Citrix servers:

1. 32-bit: approximately 3 Citrix/Terminal servers

2. 64-bit: approximately 2 Citrix/Terminal servers



7.8 Typical end user client PC specifications

Software Requirements:

- 1. Windows XP Pro SP3 or Vista SP1 (recommended)
- 2. Office XP, 2003 or 2007 (with latest Microsoft service pack, plus hotfix if using Office 2007)
- 3. Adobe Acrobat Reader (v8.12 recommended)
- 4. Internet Explorer 6.0+(IE 7 recommended)
- 5. MDAC 2.8+
- 6. .NET Framework 2.0 (SP1 minimum, SP2 recommended)

Hardware Recommendations Naturally:

- 1. software will physically run on under powered client PCs
- 2. However, the performance that end-users receive may be unacceptable on such slow PCs.

Feedback to IBM Cognos Support, based on customer's experience, gives the following suggested guidelines:

	minimum	recommended	ideal
CPU	P4 CPU	2.4GHz+	3GHz+ dual-core
RAM	512 Mb	1Gb or more	2Gb or more

NOTE: Although it is possible to run Controller 8 using the "minimum" specifications, unless there are exceptional circumstances all customers should use at least the <u>recommended</u> hardware specifications for their client PCs.



8 Muli Server (number) recommendations

8.1 Overview

As explained previously in this document, the most typical/recommended approach to server sizing is to split the Controller components onto multiple servers. This approach normally gives:

- 1. Maximum scalability
- 2. Minimum cost (as a ratio to performance)
- 3. Maximum flexibility

Therefore, this section shall assume that all application servers are identically specified (see section 5)

IMPORTANT: The following examples are purely for illustrative purposes.

- 1. They are based on 'basic' implementations where the customer is only using 'core' functionality (such as consolidations and data entry)
- 2. Later in this document, the author discusses 'optional' functionality such as Data Marts and FAP
- 3. Therefore, it is possible that customer-specific needs mean that they require higherperformance servers for smaller-numbers of concurrent users.

For example, customer X may have 30 concurrent users and therefore (see 'Example 2' later) you may imagine 2 servers to be the preferred architecture. However, customer X wants to use all the optional features (see next chapter) of Controller to an advanced level and yet still have top performance. In this customer's case, 3 (or more!) separate servers would be sensible.

Therefore, please use the following guide *only* as a <u>basis</u> for your discussions, <u>not</u> as an absolute/definitive recommendations.

8.2 Example #1 - up to and including 15 Concurrent Users

In this scenario it is assumed that there is 'standard' usage of Controller (for example noncomplex consolidation requirements) for a limited number of end users. Therefore, all Controller server components can be installed on a single "application" server

8.3 Example #2 - up to and including 40 Concurrent Users

In this scenario it is assumed that there is 'standard' usage of Controller (for example noncomplex consolidation requirements) for a limited number of end users.

Therefore, all Controller server components can be installed on a single "application" server:



8.4 Example #3 - up to and including 75 Concurrent Users

In this scenario, there is too much workload for a single server to run all the Controller



'application' services. Therefore, we shall use two separate physical servers for our application layer.

The most resource-intensive tasks are running Reports and Consolidations. Therefore, we shall split one of these roles off. Different customers may have different priorities – some may run many consolidations, whilst others may have more of a reporting requirement.

In the example below, we have decided that the customer's reporting requirement is more intensive (than their consolidation requirement), so we have separated the Cognos 8 BI 'run-time' reporting system onto a separate server:



However, in a different customer's case, it may be a better idea to have a separate (new) consolidating application server (and keep the reporting) services on the first shared web/app server.



8.5 Example #4 – Approximately 75 to 150 Concurrent Users

In order to support the anticipated load of up to 150 'standard' users it is necessary to install both a separate Cognos 8 reporting server and consolidation server:





8.6 Example #5 – Large Implementations 150+ Concurrent users

Large implementations of Controller 8.x will benefit from the distribution of the consolidation functions of Controller to a separate server. This reduces CPU and I/O conflicts for Web Service and Reporting users during consolidation tasks. Large Cognos 8 implementations can also be deployed in a fully redundant topology. Please refer to the Cognos 8 Planning and Architecture Guide for more information regarding the advanced deployment options for Cognos 8.

9 Using optional IBM Cognos Controller features

9.1 Overview

IBM Cognos Controller is an incredibly flexible product, which gives customers access to a huge number of features. Each customer has different needs and requirements.

- 1. Some customers are simply looking for a 'standard' consolidation system, where their needs are simply data entry, consolidation and basic reporting
- 2. Other customers would like to integrate Controller tightly into many of their other systems, such as their Data Warehouse and/or have complex reporting requirements.

The features described below are therefore considered 'optional', even though they are readily available for customers to use on any Controller 8.5 system.

Therefore, any server hardware recommendations (inside each section of this chapter) should be thought of as in addition to the 'core' Controller server requirements listed in the previous chapter (Chapter 8).

9.2 Data Marts (RDBMS / OLAP)

IBM Cognos Controller can publish data to an RDBMS ("Cognos 8 Framework Manager Model" method) or an OLAP (Microsoft Analysis Service / MSAS method). This allows advanced reporting through many IBM Cognos and third-party products (for example Query Studio and PowerPlay).

RDBMS Data Mart:

1. Depending on the customer's specific requirements, this may require the use of a extra server (s), for example a 'dedicated' Cognos 8 BI Report server

OLAP Data Mart:

1. Depending on the customer's specific requirements, this may require the use of a extra server (s), for example a 'dedicated' MSAS (OLAP) server



9.3 Financial Analytics Publisher (FAP)

FAP is a new feature, available from Controller 8.5 onwards. It allows near real-time reporting from Controller, by trickle-publishing the Controller data to a 'staging' data mart area (FAP Database), and from then into a TM1 cube from where it can be reported from.

Logically, the FAP services can be represented by the following picture:



Financial Analytics Publisher client

This is the 'admin console' for FAP. This is typically installed on the Controller application server.

Financial Analytics Publisher database

This may be located on a separate/dedicated database server.

Financial Analytics Publisher Service and TM1

This is typically located on a separate/dedicated 'TM1' application



TIP: In addition to the server components (shown above) you will also need client tools to access the cube to consume and analyse the data. There are a number of reporting client tools you can use to do this, including:



TM1 server sizing is influenced by a number of factors, such as the amount Controller data being pushed over to TM1. Naturally this will vary greatly between each customer, for example:

- 1. the amount of fact data (in the customer's relational database)
- 2. business decisions such as "how many periods to trickle-publish/keep in the cube

Therefore, depending on the customer's scenario, you may need 1 or 2 extra 'FAP' servers (in addition to the initial 'Controller' servers).

The peak load (on the FAP hardware) will occur during:

- 1. Initial publish
- 2. The closing period (when the majority of changes are processed in Controller) 'trickle publish'
- 3. TM1 reporting peaks (reporting from data in the cube)



TIP: Detailed in-depth analysis of TM1 sizing is beyond the scope of this document. For more detailed information, either:

1. Use the followingweblink:

http://www.ibm.com/developerworks/data/library/cognos/page205.html

- 1. This is where you can download a compilation of recommended practices on the topic of TM1 server planning, sizing, and hardware performance (called 'tm1_server_administration.pdf').
- 2. Contact your TM1 consultant

Instead, this document shall simply give "general rule of thumb" advice. It assumes that TM1 usage is restricted to *only* using FAP - in other words, TM1 usage is restricted to only reporting from Controller data.

In this case:

- 1. Most/standard Controller customers probably only need a single (<u>one</u>) <u>extra</u> application server for their TM1 processes
- 2. this acts as a 'combined' TM1 'application server' and TM1 'data storage' server
- 3. Advanced Controller customers (typically those customer whose use of Controller is sufficiently advanced for them to require a separate/dedicated Consolidation server) may benefit from a <u>2</u> separate extra TM1 servers (TM1 application and TM1 data servers)

For each TM1 server, you can apply 'standard TM1 sizing' principles (liaison with your TM1 consultant for more details). However, in general:

- 1. Most customers benefit from their TM1 server(s) being based on:
 - 1. TM1 64-bit edition (running on Windows 64-bit OS)
 - 2. 8 CPU cores (for example dual-CPU with each CPU quad-core)
 - 3. 8G RAM
- 2. If the customer's requirements are more basic (and the user count is low) then often customers will again choose Windows x64, but based on a 4 CPU core and 4Gb RAM server

9.4 Integration with IBM Cognos Products

As a basic start for exploring this topic, you may find the following link useful: <u>http://www-01.ibm.com/support/docview.wss?uid=swg21343279.</u>

Cognos.

software

10 Real-Life Example

10.1 Customer Example #1 – 100 to 120 Concurrent Users

This customer has 400 named users, but normally sees a maximum of approximately 100-120 concurrent users actually logged in at any one time.



To summarise the above picture:



Information Management

Description	Role	Specification
SQL	SQL 2005 database server	8 x physically-separate Intel CPUs, 8Gb RAM
APP #1	Main Controller application server	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
APP #2	Dedicated consolidation-only app server	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
APP #3	Dedicated Reports(BI)-only app server	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
CITRIX #1	Citrix server, from where Controller is launched	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
CITRIX #2	"	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
CITRIX #3	"	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
CITRIX #4	"	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
CITRIX #5	"	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM
CITRIX #6	n	2x physical dual-core AMD @ 2.4GHz (4 cores in total), 4Gb RAM

They are all very similar because they use Citrix and have about 100-130 concurrent users.

10.2 Customer Example #2 – 100 to 125 Concurrent Users

This customer has 300 named users, with approximately 100 - 125 concurrent users as a maximum.

Description	Role	Specification
SQL	SQL 2005 database server	Intel E7340 @ 2.4GHz (8 cores in total), 16Gb RAM
APP #1	Gateway & main Controller application server	1x quad-core Intel E5420 @ 2.5GHz (4 cores in total), 3.25Gb RAM
APP #2	Dedicated consolidation-only app server	1x quad-core Intel E5420 @ 2.5GHz (4 cores in total), 3.25Gb RAM
APP #3	Dedicated Reports(BI)-only app server	1x quad-core Intel E5420 @ 2.5GHz (4 cores in total), 3.25Gb RAM
CITRIX x	Citrix server, from where Controller is launched	1x quad-core Intel E5420 @ 2.5GHz (4 cores in total), 8Gb RAM
NOTE:	There are multiple Citrix servers in use	

10.3 Customer Example #3 – 30 Concurrent Users

This customer has 400 <u>named</u> users but not more then 30 <u>concurrent</u> users.

Description	Role	Specification
ORAČLE	Oracle database server	Oracle 10g Database server on AIX 64bit, 2 CPU
APP #1	Main Controller application server	2 X INTEL XEON CPU 1,86 Ghz E5320, 4 GB RAM
APP #2	Dedicated consolidation-only app server	2 X INTEL XEON CPU 1,86 Ghz E5320, 4 GB RAM
CITRIX #1	Citrix server, from where Controller is launched	
CITRIX #2	"	
CITRIX #3	"	
CITRIX #4	"	
CITRIX #5	"	
CITRIX #6	"	



35

Information Management