

Calculating disk and memory

Calculate memory and disk space requirements based on the questionnaire responses. There is no "formula" that would give the answer because there is no such thing as a typical installation. The amount of content and the number of concurrent authors are the main factors that decide how much load the Magnolia CMS instance needs to handle. Do the calculation separately for author and public instances as the requirements are different.

Author instance

Basic memory

- Operating system needs 1 GB of memory
- Each Tomcat instance running Magnolia CMS needs a minimum of 512MB, but optimally at least 1 GB of memory (RAM).
- A local database needs a minimum of 512 MB, but optimally at least 1 GB as well. If the database is not on the same server machine it doesn't need the 512/1024 MB but you need another server for the database.
- Total server memory is thus 2 GB at minimum, optimally at least 3GB
- When you buy a physical server you typically get 16 GB of RAM. (early 2016)
- When you buy a VM, those with get 1GB are too small.

Authors and memory

- How many authors do you have?
- A typical site has 100-200 author.
- If you have more than 1000 authors, add more memory for Tomcat.
- Number of authors as such does not matter. What matters is *concurrent* authors who edit at the same time.
- 20 concurrent editors, each editing 5 pages per day, generate 100 page edits (one edit every 6 minutes on average). They can work with 512 MB of Tomcat memory, assuming they are not manipulating hi-res images. In case retina-ready images are used on the site and manipulated via image variations, server needs at least 1 GB RAM.
- Uploading binary content is considered a heavy edit. Add memory.
- If you make heavy use of Imaging (resizing, generating variants), add even more memory.
- If you use lots of images, flash and videos, add memory.
- A multisite configuration needs one cache per site. Each cache uses 30 MB of RAM. If you have 10 sites (10 caches) you need 300 MB of memory just for EHcache.
- If you are using advanced cache strategies, those keep two copies of cache (live and one that is being pre-heated) so double the amount of memory calculated previously for cache.
- In total, if you end up using more than 8GB RAM however, something might be incorrect in your calculation or setup. Double check with someone at Magnolia.

Disk space

- Magnolia CMS uses the disk to store data, to create an index of content and to cache content. Caching is not heavily used on author instance, mostly on public.
- Data storage is heavier on author because of versioning. Versioning is done by copying. This is expensive for binary content.
- Is the data stored in a local database?
- Assign 70GB of disk space for the database, wherever it resides.
- Assign 30GB for Magnolia CMS and Tomcat. Used for indexes, temp files and cache.
- Do you store large binary files in the database or on the local file system? If you have file based data store you need more disk space.

Public instance

Requests

- The number of users does not matter. Number of *requests* matters.
- Typically each page request generates 10-20 requests on the server for dependent assets and resources.
- 1 page requests = 10 server requests
- 1 page request = 200 threads on Tomcat
- 1 page request = 200 ms processing time on Tomcat
- In 1 second you can handle 1000 requests.
- Tomcat takes 200 ms to process one page request if the page is not cached. Caching brings response time down to 20 ms.
- Tomcat cannot handle more than 200 requests simultaneously. Furthermore you can handle 1000 requests.
- WebSphere and WebLogic can handle more simultaneous requests.

Type of content

- If you use heavy imaging, add more time.
- If you get data from some remote source, add time.
- Aim at no more than 100 requests served by any given instance at the same time.
- Streaming video? Add 1GB of memory on the public.
- UCG? Add 512 MB of memory on the public.
- If you are not using a local database, take network throughput into account. The data has to be transferred twice: first from the database to Tomcat through the network interface, then from Tomcat to the user.
- Always use a local disk storage if possible. This means a physical disk on the same machine.
- MySQL needs 512 MB as minimum, optimally 1 GB
- Oracle and Microsoft SQL Server need 1 GB
- The bigger your pages, the more memory is needed.
- If you have over 1GB data on the website, double the memory.
- Imported data from other sources such as integrated systems may flush the cache on the public instance

Example of Hardware / Software

Ref	Name	Description	Software	Virtual Resources		
				Ram (GB)	vC PU	Disk (GB)
Production Front-End						
A1	Web Application Firewall	Optional Fortinet FortiWeb Appliance: Secures web applications to help customers meet compliance requirements	FortiWeb	4	2	50
A2	Web Application Firewall	Optional	FortiWeb	4	2	50
B1	JBoss EWS	Front-End web server. Reverse proxy the requests to the application server; manage the http sessions through the Jboss cluster. Also used for rewrite rules; access logs ...	Windows 2012 JBoss EWS Mod_Cluster	4	2	50
B2	JBoss EWS	"	Windows 2012 JBoss EWS Mod_Cluster	4	2	50
Production Servers						
C1	Magnolia Public	Public Web sites are served by the Magnolia Public instance. Magnolia Public application is installed in the Jboss cluster.	Windows 2012 JBoss EAP 6.3.x Magnolia CMS EE 5.x JDK 7.x	8	4	120
C2	Magnolia Public	"	Windows 2012 JBoss EAP 6.3.x	8	4	120

			Magnolia CMS EE 5.x JDK 7.x			
D	Magnolia Public Stage	Web content is first published in the Magnolia Public Stage server and then, when approved, to the Magnolia Public Cluster (c1,c2)	Windows 2012 JBoss EAP 6.3.x Magnolia CMS EE 5.x JDK 7.x	6	2	12 0
E	Magnolia Author	Magnolia Authoring server	Windows 2012 JBoss EAP 6.3.x Magnolia CMS EE 5.x JDK 7.x	10	4	12 0
F 1	Database	The MS-SQL cluster will contain all the databases needed by the different magnolia instances.	Windows 2012 MS-SQL 2014	8	4	20 0
F 2	Database	"	Windows 2012 MS-SQL 2014	8	4	20 0
Total				64	30	10 80