ownCloud Hardware Sizing

Introduction

One very common question as customers look to size ownCloud is "how much hardware do I need?" This is not a simple question to answer, and it depends a lot on ownCloud usage patterns and user behavior. The purpose of this document is to provide an overview of ownCloud sizing considerations, and provide, with appropriate assumptions, two examples of the hardware sufficient to handle 2 common ownCloud installations.

The first answer to the question of hardware sizing is yet another question: what do you expect your users to do with the system? This may sound trivial, but sizing is directly related to the size and number of files stored in the system, the number of application plug-ins enabled on the system, and the nature of the users, devices and bandwidth used to connect to the system. For example, higher bandwidth connections use less memory, but stress disk performance as more files are uploaded in a given time frame.

The key questions we look at are users, devices, behaviors, files, and geography. More specifically, these questions:

- How many users are expected to be active on the system each day?
- How many devices will they be connecting?
- What type and number of devices (e.g. mobile phone, tablet or desktop / laptop machine) will be used?
- How many files are they syncing and sharing?
- Is it a regional or global system?

There are many other questions, but these tend to be the most important to get right.

Example 1

For the first example, lets look at a typical proof of concept or departmental deployment scenario:

- 500 users active on the system
- Each user will have on average 1.25 desktops, 1 mobile device and also use a web browser to connect to the system
- The desktop clients are on 24 hours a day, mobile devices connect on average every 5 minutes during the work day, a web browser every hour spread around the clock
  - » *Mobile devices usually upload and download on average 4-8 MB files*
  - » *Desktop devices upload and download 16 MB files*
  - » *Web devices upload and download 16 MB files*
- Users are limited to 10GB of data total quota
- It is a regional office location, all in one time zone

This means that we will have up to 1,625 access points hitting the system, 3.25 for each user.

Recent performance tests of ownCloud suggest that this scenario is best handled by an n-tier web architecture, with 1 combined application server and database server running on an Intel Xeon E5-2640 based single socket server with 32GB of RAM – or equivalent. While this will handle the load, there are two questions that arise related to:

1. Fail over in case of a lost application server
2. Future expansion plans

If you plan on providing fail over capabilities, then a second identically configured application server / database server will be needed. The two databases will need to be synced, and a load balancer will be needed to redirect traffic to the backup server in case the primary is lost. If you are thinking about scaling in the future, it may be a lot simpler to separate the database server from the application server, use a clustered database implementation, and add application servers and database cluster nodes as needed to scale. While none of this is required for a proof of concept, it may be desired to make migration simpler in the future. For the purpose of the rest of this example, we will just assume a single server setup.

For memory sizing, this is highly dependent on the number of files being uploaded at any given time. The ownCloud architecture typically uses less memory and more processing power, so faster processors are more important than additional memory. Slower bandwidth connections also require more memory, as the server must keep the threads alive while the files are uploaded. We

recommend 32 GB, which will provide extra capacity if network connections ever slow and more web server sessions are needed.

For the public side of the server, a gigabit network is recommended. It is easy to saturate a 100MB network with requests to a single app server, so we suggest gigabit Ethernet out front. For server direct attached storage – where the OS and ownCloud software is installed – we recommend SSD drives because a potential bottleneck of the server is the writing of log files, but this is not required in this instance. The drives should be mirrored in case of failure, and large enough to handle the OS, database, Apache, php, and the ownCloud application, as well as the logs, cache files and sessions.

Back-end storage sizing is entirely dependent on the size and quota that is expected for each user. For example, 500 users with a 10 GB quota would have a total storage need of 5 TB, plus a little for new users and special exceptions, we would be looking at perhaps 10 TB in total. The storage speed is important to consider, as the more time it takes to write a file, the more processes may back up on the server if it saturates. It is also good to think of this storage in terms of growth over time, so it is good to pick a flexible storage solution on the back end. After all, very few users accumulate fewer files over time.

There are a number of options to tune the ownCloud installation. The database, for example, needs indexes in the most active tables. The number of live Apache connections needs to be turned up to 500 or more, and the number of allowed MySQL connections also has to be increased to the same. Turning on the Alternative PHP Cache (APC) will also increase performance on the app servers, and there are likely a number of environment and policy specific configurations needed as well in any given deployment.

Example 2

Now, for the second example, take the following requirements:

- 8,000 users active on the system
- Each user will have on average 1.25 desktops, 2 mobile device and also use a web browser to connect to the system
- The desktop clients are on 24 hours a day, mobile devices connect on average every 5 minutes during the work day, a web browser every hour spread around the clock
  - » *Mobile devices usually upload and download on average 4-8 MB files*
  - » *Desktop devices upload and download 16 MB files*
  - » *Web devices upload and download 16 MB files*
- Users are limited to 10GB of data total quota
- It is a regional office location, all in one time zone

This means that we will have up to 34,000 access points hitting the system, 4.25 for each user. As pointed out, every one of these users will connect on average 2 desktop clients to the server, and it is important to understand that they all expect the desktop sync client to respond quickly to changed files. Also note, most users user their laptops or desktops as the primary source of files in the enterprise, so it has to connect frequently and syncs often. This is why it is so important to understand the expected usage patterns, and estimate the anticipated desktop transaction volume.

Recent performance tests of ownCloud suggest that this scenario is best handled by an n-tier web architecture, with 2 ownCloud app servers and a 2 node MySQL database cluster if we use a moderate performance Intel Xeon E5-2640 based dual socket server with 32 GB of RAM, mirrored SSD OS drives, and appropriate external storage to meet quota. This provides a peak to mean ratio of about 2.5x, which should provide enough processing power in a peak scenario to meet demands without creating disastrous backlogs for the users. More powerful servers could be applied, with quad sockets or more expensive processors, and thus reduce the number of servers required.

The reason we ask about geographic distribution is that system usage tends to be bi-modal, with peaks at about 10:30AM and again at 2:30 PM local time. This coincides with the most productive times of the work day for most users. A geographically distributed system will have a lower peak to mean requirement for their server capacity, because their usage will be spread out over a large window of time across time zones. A local instance, on the other hand, will have 2-3x peak to mean ratios that coincide roughly with these times during the day.

Sample ownCloud Deployment Architecture

For memory sizing, this is highly dependent on the number of files being uploaded at any given time. The ownCloud architecture typically uses less memory and more processing power, so faster processors are more important than additional memory.

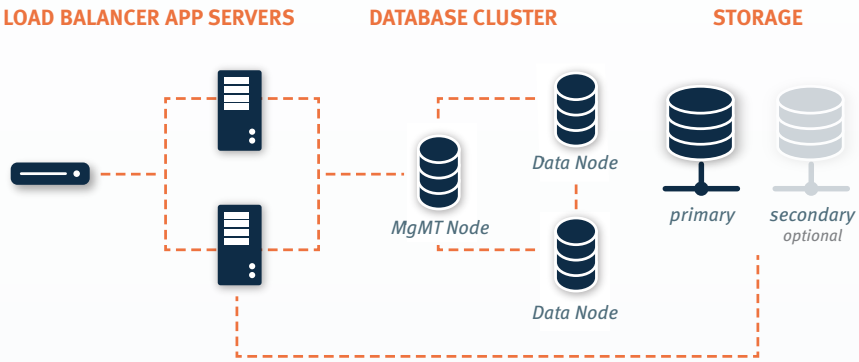

For the front-end network, a gigabit or 10GigE network is recommended leading to a load balancer, and from the load balancer to the three app servers. It is easy to saturate a 100MB network with requests to a single app server, so we suggest at least gigabit Ethernet. For internal networking between the servers, we also recommend at least a Gigabit Ethernet to reduce the potential for bottlenecks on the MySQL connection and back-end storage. We also recommend that each server have a dedicated connection to the database cluster, and also that each server have a dedicated connection to the back-end storage.

For server direct attached storage, we recommend the SSD drives because a potential bottleneck of the server is the writing of log files. The SSD drives should be mirrored in case of failure, and large enough to handle the OS, Apache, php, and the ownCloud application, as well as the logs, cache files and sessions. 128 GB SSDs are a minimum recommended size to leave enough room for all of this activity.

Back-end storage sizing is entirely dependent on the size and quota that is expected for each user. For example, 8,000 users with a 10 GB quota would have a total storage need of 80 TB, plus a little for new users and special exceptions, we would be looking at 100 TB in total. The storage speed is important to consider, as the more time it takes to write a file, the more processes may back up on the server if it saturates. It is also good to think of this storage in terms of growth over time, so it is good to pick a flexible storage solution on the back end. After all, very few users accumulate fewer files over time.

There are, of course, a number of options to tune the ownCloud installation and enable this level of performance. The database, for example, needs indexes in the most active tables. The number of live Apache connections needs to be turned up to 1000 or more, and the number of allowed MySQL connections also has to be increased to the same. Turning on the Alternative PHP Cache (APC) will also increase performance on the app servers, and there are likely a number of environment and policy specific configurations needed as well in any given deployment.

Conclusion

This reference architecture is provided as an example to assist customers in basic cost of ownership calculations. For 500 or 8,000 users, these architectures and setups should provide more than enough capacity to meet ownCloud's needs. If more detail is needed to size a specific project, we highly recommend contacting your account manager and scheduling a workshop to discuss this with you live, and to enable us to provide you with the best estimates customized for your user profile, environment and business need. We would be happy to work with you to determine how to best leverage your existing infrastructure to deliver the right solution for you. For more information, drop us a line at **info@owncloud.com**, or contact your account manager directly.

Corporate Headquarters
ownCloud, Inc.
10 Foster Road
Lexington, MA 02421
United States
info@owncloud.com

https://www.owncloud.com