keep in touch

OPEN-XCHANGE

Open-Xchange Hardware Needs

OPEN-XCHANGE: Open-Xchange Hardware Needs

Publication date Wednesday, 18 January 2012 version 1.2

OX – Hardware Sizing

1. Hardware Needs with Open-Xchange

1.1. Overview

The purpose of this document is to get an idea about the hardware needed when rolling out Open-Xchange, depending on the amount of concurrent users.

The picture below shows how an Open-Xchange cluster could look like. The Open-Xchange architecture is designed to scale vertically as well as horizontally. This means all services can run on one physical server as well as on different machines. If more CPU or RAM is added to a single server, it is efficiently used. If more resources are required, either the cluster can be extended, services can be outsourced to separate machines or hardware can be extended later on. The architecture currently has no known physical limits. The application also runs in virtual environments e.g., XEN and VMware.

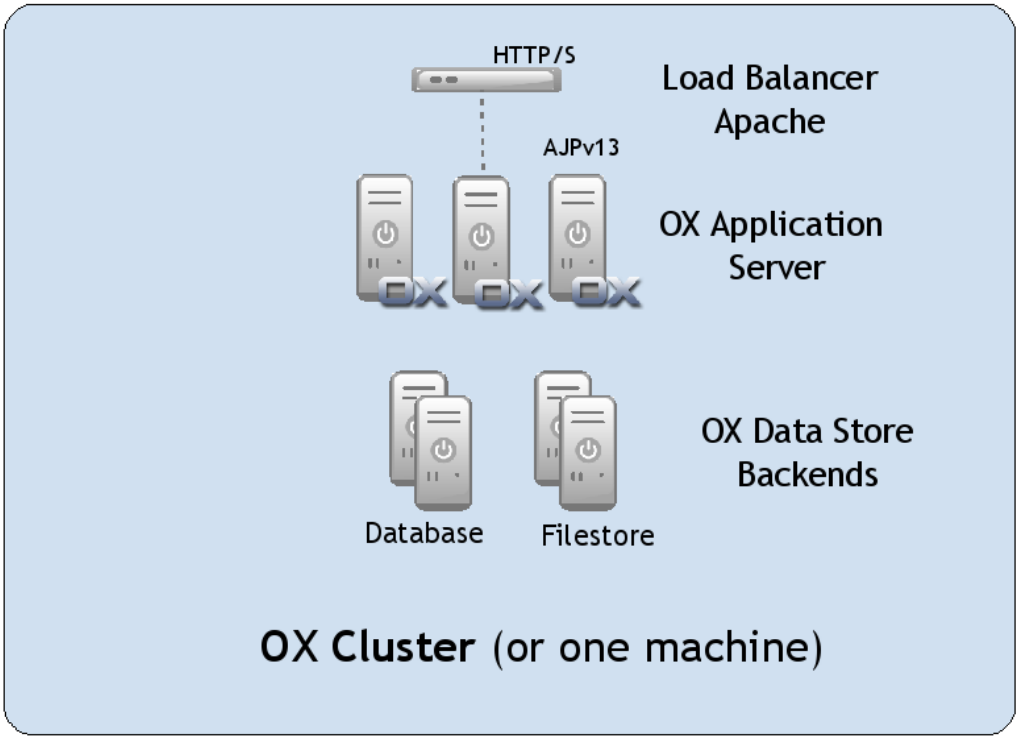

Figure 1. OX Cluster

1.2. How to calculate concurrent sessions

If several users are connected to the Open-Xchange application at the same time we are speaking of concurrent sessions. Based on Open-Xchange's experience, the number of concurrent sessions can be calculated from existing mailboxes. For this calculation we need to distinguish between Webmail accounts which are given to the end users for free, and fully featured Groupware accounts which are paid by the end user. The main reason, why this accounts into the relation, is because free webmail accounts have a much shorter session

lifetime as they are mostly used to check emails for example once a day. The paid groupware accounts normally have a much longer session lifetime, as the accounts are used during the whole working day. The relation between groupware accounts (existing in the system) and concurrent users at peak time is between 1:5 and 1:10. For cost free Webmail mailboxes the relation varies from 1:50 to 1:100 – in this case it is safe to calculate the amount from already existing services which should get replaced by Open-Xchange.

1.3. Planning the Open-Xchange application servers

On the Open-Xchange application servers, the main factors are RAM and CPU. If the Filestore resides on a separate server, IO is almost negligible. The main factor is RAM. For each concurrent user, 4MB of RAM needs to be available in the Cluster to allow efficient caching. It is also recommended to have 64Bit CPUs and to assign one CPU core to each 2GB of RAM.

1.4. Planning the Apache Instances

In small environments the Apache instances can run on the same physical hosts as the application servers. Starting with 4 to 5 application servers it is recommended to have own Apache servers on separate machines. The number of Apache instances should not be higher than the number of application servers available. The Apache Server acts as a proxy between clients and groupware servers and transfers the application data to the users' web browsers.

1.5. Planning the MySQL Database

In Open-Xchange clusters, the database write-connections in proportion to the read-connections is about 1:3. Open-Xchange has a pool for database connections where 10 active users share one connection in average. The memory for database servers should be at least double of the data which is used for MySQL files. This makes sure that read connections can be served out of RAM and there is room for a growing database. It is recommended to have at least one MySQL server that serves two application nodes. For high-availble setups, the minimum for production use is one MySQL master and one MySQL slave server (MySQL replication is supported by Open-Xchange). If more slaves are needed, a load balancer is required to distribute the read requests.

Open-Xchange recommends to not exceed following limits:

250 database schemas per database instance
1000 contexts per database schema
5000 users per context

1.6. Planning the filestore server space

The filestore depends on the customer usage. If Open-Xchange acts only as webmail replacement without the store, no space is needed at all but required later if upsell to full groupware functionality is planned.

In case the filestore is in use, the space must be at least as big as the physical space of all customer documents plus versions of them. Since the application has the ability to overbook the physical space, it is required to monitor free space during operation and to extend the space as required.

1.7. Network Traffic Throughput Profile

The following tables show the network traffic demands of webmail end users. The figures are calculated from huge installations serving multiple thousands of concurrent sessions. The figures can be used for a base calculation.

	MB/day
in	3.2
out	13.7
sum	16.9

Table 1. Internet <-> user traffic per user

	MB/day
in	1.0
out	2.2
sum	3.2

Table 2. Database <-> Open-Xchange traffic per user

1.8. Load balancer

There is no special requirement for the load balancer. The load balancer can be a hardware or software solution. Session stickiness is provided by the proxy ajp functionality of the Apache server. www.keepalived.org is known to work reliable.

2. Different scenarios and recommendations

In the following sections you will find some recommendations for standard setups. These are to be regarded as general recommendations to have a basic guideline. Many other deployments are possible depending on the real infrastructure and existing services.

2.1. Minimum Setup On-Premises

All services can be installed on one machine. With this setup high availability is not given. As all services are running on this machine, it scales to approximately 1,000 to 2,000 concurrently logged in users.

This scenario is only recommended for companies which do not need high availability.

Service	**#**	**RAM**	**CPU Cores**
OX Application server	1	8GB	2
Mysql servers	0		
Apache servers	0		
File store Servers	0		
Load Balancer	0		

Table 3.

2.2. Minimum High Availability Setup On-Premises / Hosting Environments

From a scalability perspective, this deployment is equivalent to the one described above. The main difference is, that all services are running on two machines. OX Application servers are running loadbalanced and the database is separated to one master and one slave.

Failover can be implemented with all known Linux mechanisms.

For this setup it may be recommended to have the filestore mounted from an external system.

Service	#	RAM	CPU Cores
OX Application server	2	8GB	4
Mysql servers	0		
Apache servers	0		
File store Servers	0pt.		
Load Balancer	1		

Table 4.

2.3. Hosting Environment up to 50,000 / 500,000 users

The following setup is the minimum recommendation to allow a high quality service without significant interruption in case of hardware failures.

This setup gives flexible scalability for future requirements, so it is a safe starting point even for larger deployments, additional machines can be added during operation with growing user numbers.

Given the ratios from above, the system scales up to approximately 50,000 paid groupware accounts or equivalent to 500,000 free webmail accounts.

Service	#	RAM	CPU Cores
OX Application server	2	8GB	4
Mysql servers	2	4GB	2
Apache servers	0		
File store Servers	1		
Load Balancer	1		

Table 5.

2.4. Example deployment for millions of users

This is one example (!) for a real huge deployment using all available distribution mechanisms. Of course environments like this need to be defined together based on the real infrastructure.

In this example all services are running on multiple separate machines to gain maximum flexible transparent scalability.

The given example would allow to serve approximately half a million paid groupware accounts or equivalent 5Mio of free webmail users.

Service	#	RAM	CPU Cores
OX Application server	10	16GB	8

Service	#	RAM	CPU Cores
Mysql servers	6	4GB	2
Apache servers	6	2GB	2
File store Servers	1		
Load Balancer	2		

Table 6.

3. Test Results

Fujitsu, SUSE, and OX cooperated to provide sizing information for Open-Xchange Server Hosting Edition on a given setup of hardware and software.

The goal is to provide guidance for product managers at service providers.

On the setup used (Fujitsu CX1000, SLES11) the amount of users that can be served per blade (Fujitsu CX122) varies between 3.500 and 30.000 concurrent users. Resource requirements are highly dependent on specific usage scenarios. What this paper provides are upper and lower bounds of what can be expected from a given hardware / software scenarios in extreme cases.

3500 users would represent a call center where every user schedules a new appointment, sends an email, or adds a new contact every second. 30.000 users represent a realistic scenario on the hardware described below for users who pause between actions like adding calendar entries and deleting contacts from a list.

3.1. Results

- 150MBit/s traffic,
- 50% CPU load,
- 20% RAM used (ca. 15GB on OX servers / 3.75MB per user),
- 4000 connections to webserver,
- 6000 queries per second through 140 connections to the database

3.2. Scenario

- 1000 virtual OX domains with 5 users each
- 4000 emulated concurrent browser users
- pre-filled accounts, Inbox with 200 emails, new / reply / delete mail, new / accept / delete appointments, new / edit / delete contacts

3.3. Hardware Used For Tests

- 5x Fujitsu CX122 blades (Dual Xeon 5620, 48GB RAM),
- deployed in CX1000S1 Rack,
- containing Brocade 648S switches

3.4. Software Used For Tests

- 1x Webserver+OX (SLES11, Apache, OX 6.20-rev29)
- 1x OX (SLES11, OX 6.20-rev29)
- 1x DB Master (SLES11, MySQL 5.5, InnoDB with multi core setting)
- 1x DB Slave (SLES11, MySQL 5.5, InnoDB with multi core setting)
- 1x IMAP Server (SLES11, Cyrus, Postfix)
- Ramdisk for IMAP and DBMS