

Performance testing of a Moodle system exclusively dedicated to massive online evaluations: determining its limits

José Fager

Emilio Penna

Servicio Central de Informática

Universidad de la República - Uruguay

moodlemoot
GLOBAL 2022

#mootglobal22

About Udelar

- Udelar: Universidad de la República (Udelar)
- Main university institution in Uruguay
- 100,000 students
- 11,500 teachers
- 450 careers
- 20 schools
- 16 centralized Moodle installations
- 9 decentralized Moodle installations
- Centralized (shibboleth) login



UNIVERSIDAD
DE LA REPÚBLICA
URUGUAY

During the pandemic

- Year 2020: Use of technologies for education intensified
- All university activities were 100% online
- New servers were installed to deploy moodle instances
- It is common practice at Udelar for thousands of users to participate simultaneously in evaluations



Requirements

Due to the large number of users, it was necessary to set up a Moodle-based system dedicated exclusively to carry out online evaluations.

Someone asked:

How many students does it support?... at the SAME time?

Looking for an answer, we decided to run a load test. We used Apache Jmeter – with this tool, we can develop a script that makes the same HTTP requests that the real users make from their web browsers.



Simulation design

- We configured a Moodle quiz with different type of questions (16) and some (small) images.
- Test data was generated, including 5000 test users enrolled in the test course.
- What does it mean that the users enter “at the same time”?
- Critical moment: quiz start
- It is not the same:
 - 1000 users start the quiz attempt in 10 seconds
 - 1000 users start the quiz attempt in 2 minutes

Timing

We analyzed real Apache logs, looking at the distribution of the time of the request to the start of the quiz attempts (startattempt.php). We observed, for example:

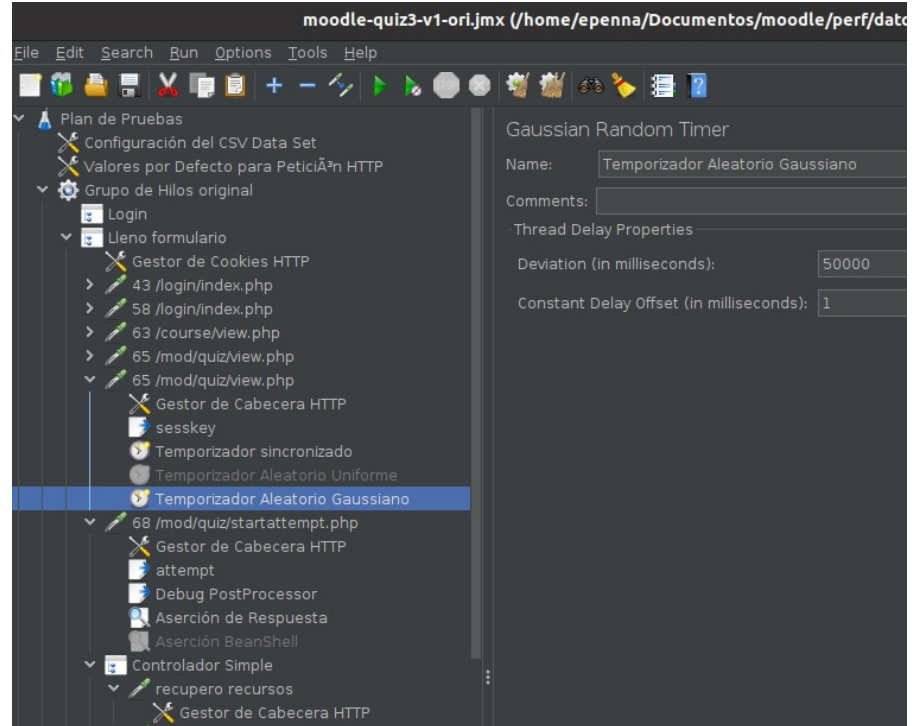
- In the first 10 seconds, 13% of users started the attempt
- In the first minute, 73% of users started the attempt

Segundo	Cant Usuarios	Porcentaje
0	0	0 %
10	334	13 %
20	731	29 %
30	1100	44 %
40	1416	57 %
50	1685	67 %
60	1835	73 %
70	2020	81 %
80	2187	87 %
90	2295	92 %
100	2365	95 %
110	2418	97 %
120	2451	98 %

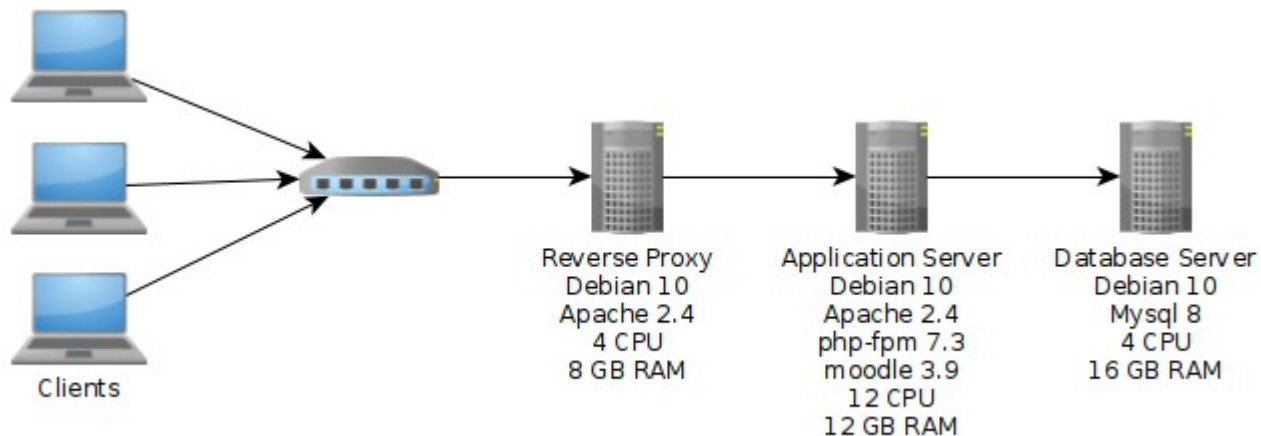


Simulation desing in jmeter

- Users logged in, in the 10 previous minutes to the quiz start time (10 min ramp up)
- Synchronized Timer, that waits until all Virtual Users are logged in, then the quiz starts
- A good approximation to the “start attempt” distribution was achieved using a “Gaussian Random Timer” (with a 50s deviation)



Infrastructure (test=prod)



Bottlenecks and adjustments

Bottlenecks detected in the load tests:

- Reached max DB connections
- Reached max Apache workers (using prefork mpm)
- Too high CPU load in app server

Main adjustments to enhance performance:

- Raised max DB connections
- Changed to Apache event mpm and php-fpm and raised apache workers
- Enabling and configuring opcache
- More CPU for application server (8 to 12 CPU)



Main results

- We ran tests with 2500 virtual users (VU) with no errors and good response times (560ms avg, 800ms std dev)
- With 3000 VU the response time got worse and the CPU load were too high (200%) during most of the test time.

Results - 2500 VU

Etiqueta	# Muestras	Media	Mín	Máx	Desv. Estándar	% Error	Rendimiento
43 /login/index....	2500	105	82	686	49,20	0,00%	8,9/sec
58 /login/index....	2500	400	230	1677	136,67	0,00%	8,9/sec
63 /course/vie...	2500	145	110	650	42,20	0,00%	9,0/sec
65 /mod/quiz/vi...	5000	506	109	7031	698,58	0,00%	11,0/sec
68 /mod/quiz/st...	2500	1716	224	10051	1782,19	0,00%	14,2/sec
recupero recur...	15000	493	6	6894	639,17	0,00%	85,6/sec
Lleno primer pr...	2500	1837	139	11793	2219,19	0,00%	12,6/sec
Lleno segunda...	2500	918	135	8063	1450,42	0,00%	11,6/sec
paso de pagina	17500	635	121	15167	968,29	0,00%	28,8/sec
Lleno tercer pr...	2500	430	138	6898	562,52	0,00%	11,0/sec
Lleno cuarta pr...	2500	513	134	3841	560,09	0,00%	9,9/sec
Pregunta 5	2500	481	139	4727	543,74	0,00%	9,1/sec
Pregunta 6	2500	510	140	4576	559,08	0,00%	9,0/sec
Pregunta 7	2500	469	137	4122	524,95	0,00%	8,4/sec
Pregunta 8	2500	426	143	3486	448,65	0,00%	8,1/sec
Pregunta 9	2500	380	144	4089	378,03	0,00%	8,1/sec
Pregunta 10	2500	372	139	7591	382,06	0,00%	7,3/sec
Pregunta 11	2500	409	137	5200	467,71	0,00%	7,2/sec
Pregunta 12	2500	464	132	6170	601,96	0,00%	7,1/sec
Pregunta 13	2500	518	127	6135	707,93	0,00%	7,0/sec
Pregunta 14	2500	552	128	6061	757,94	0,00%	6,9/sec
Pregunta 15	2500	464	128	6053	673,44	0,00%	6,6/sec
Pregunta 16	2500	363	125	3347	543,40	0,00%	6,9/sec
finish	5000	443	147	6183	549,26	0,00%	13,8/sec
Total	92500	561	6	15167	885,85	0,00%	92,9/sec


App server performance (2500 vu)

- Max CPU usage: load average between 12 and 13 (110%) during 3 peak minutes (with 12 cores)
- Max Apache processes: 14 (150 threads each)
- Max php-fpm processes: 28
- Memory in peak minutes (no swap was used):


Total	used	free	shared	buff/ cache	available
11Gi	1,7Gi	2,4 Gi	111Mi	7,7Gi	9,9Gi



Monitoring with NMon



✓ Data T...	Metric	Minimum	Average	Weigh...	Maximum
✓ CPU_ALL	User%	0,000	10,058	51,978	76,100
✓ CPU_ALL	Sys%	0,000	2,088	12,207	21,200
✓ CPU_ALL	Wait%	0,000	0,007	0,440	1,400
✓ CPU_ALL	CPU%	0,000	12,146	63,750	93,100



✓	Data Type	Metric	Minimum	Average	Weighted Ave...	Maximum	Std Dev
✓	MEM	memtotal	12,009,200	12,009,200	12,009,200	12,009,200	0,000
✓	MEM	hightotal	-0,000	0,000	N/A	0,000	0,000
✓	MEM	lowtotal	-0,000	0,000	N/A	0,000	0,000
✓	MEM	swaptotal	3,904,000	3,904,000	3,904,000	3,904,000	0,000
✓	MEM	memfree	2,289,000	2,685,995	2,725,987	3,544,100	327,746
✓	MEM	highfree	-0,000	0,000	N/A	0,000	0,000

Other simulations using jmeter

- Cache comparison: file vs memcached (our file cache was better, using a dell-emc storage solution, unity400 FC multi-tier, with SSD and SAS disks)
- Useful: Moodle Site administration -> Plugins -> Caching -> Test performance
- Load test of Tasks with file uploads
- Load test of a quiz using Quiz Access Proctoring Plugin (takes photographs with webcam during exam). In our tests, the server supported 75% of the supported load without the plugin.



Conclusion

- Quite good results for only one application server
- Simple deployment, well known software
- No need for more RAM
- CPU intensive, maybe more CPU could give better results
- The system was used in production in more than 200 online evaluations, with up to 1250 students, without problems.
- Scripts were useful for testing other moodle systems in other universities.

Resources

- Jmeter script, test data, and documents:

<https://www.proyectos.udelar.edu.uy/redmine/projects/moodleperf/wiki/>

- The scripts and load tests were made by Valentina Parula (QA division of the university central IT service).





Questions?



Gracias





moodle moot

GLOBAL 2022

José Fager
Emilio Penna

soporte.eva@seciu.edu.uy

#mootglobal22

