

# **Comparison of Lucene mechanism and database engine in Asset publisher as a test of short term resolution for LPS-22332**

2011-11-07



NETWORKS

## SUMMARY

In Liferay there was change in behavior of Asset Publisher between version 6.0.5 and 6.1. This change involved more permission checks than which caused huge negative impact on system performance.

After little brain storm we had in [LPS-22332](#) we choose to test solution based on Lucene indexes.

The goal of this study was to prepare comparison of performance of default Asset Publisher using two different strategies and see if this is a good direction to go.

There were 2 strategies that we tested:

- Search for web content using database query (default asset publisher behavior),
- Search for web content using Lucene indexes (functionality added by [eo Networks](#) as a proof of concept for [LPS-22332](#)).

## TEST ENVIRONMENT

- Liferay version: 6.1 CE (revision 87489)
- processor: AMD Athlon(tm) 64 X2 Dual Core Processor 5000+
- RAM: 3.5GB
- Database: MySql 5.0.51a (database was deployed on the same machine as Liferay Portal)
- HDD: Seagate Barracuda 7200.10 ST3250820AS
- JVM parameters: Xmx: 1024m, MaxPermSize: 256m

## RESULTS

This study was commissioned with combination of following options:

1. Web content entries in system:
  - a) 500
  - b) 5000
  - c) 25000
2. Cache state:
  - a) System cold boot , no cached entries
  - b) Warmed system, cache filled with data
3. Different number of Asset Publisher on a test page:
  - a) 1 AP on a test page without any categories or tags set as filters

- b) 5 AP on a test page where one AP had no categories and tags set and the others had different sets of those filters including “contains” and “does not contain”.
  - c) 8 AP on a test page where one AP had no categories and tags set and the others had different sets of those filters including “contains” and “does not contain”.
4. User state:
- a) Logged
  - b) Not logged
5. Strategy:
- a) DB
  - b) Lucene

Results in tables below present server response time for requests made using Jmeter.

### Not logged user

Entries	Number of AP on test page	Response time with cold boot using DB strategy [in seconds]	Response time using DB strategy [in seconds]	Response time with cold boot using Lucene [in seconds]	Response time using Lucene strategy [in seconds]
<b>500</b>	1	9 - 10	0,25 - 0,3	11 – 13	0,1 – 0,2
	5	12 – 15	0,9 – 1	8 – 10	0,3 – 0,4
	8	16 - 18	1,5 – 1,75	8 - 9	0,5 - 0,6
<b>5000</b>	1	52 - 63	1,4 – 1,5	7	0,18 – 0,2
	5	60 - 67	5 – 5,9	10	0,4 – 0,5
	8	65 – 72	11 – 12	9,3	0,7 – 0,9
<b>25000</b>	1	62 – 66	10 – 12	8 – 11	0,3 – 0,5
	5	134 – 152	40 – 47	9 – 12	0,9 – 1,4
	8	140 – 160	60 – 70	9 – 13	1,2 - 2

### Logged user

Entries	Number of AP on test page	Response time with cold boot using DB strategy [in seconds]	Response time using DB strategy [in seconds]	Response time with cold boot using Lucene [in seconds]	Response time using Lucene strategy [in seconds]
<b>500</b>	1	11	0,2 – 0,3	4 – 5	0,1 – 0,2
	5	12 – 16	1	5 – 6	0,4 – 0,5
	8	12 – 17	2 – 2,4	8 – 9	0,55 – 0,62
<b>5000</b>	1	57 - 59	1,3 – 1,7	5 - 6	0,17 – 0,22
	5	63 - 66	5,5 – 8	7 - 9	0,4 – 0,6
	8	68 - 72	11 – 12	7 - 9	0,8 – 0,9
<b>25000</b>	1	65 – 72	9 – 12	7 – 8	0,3 – 0,4

	5	120 – 123	40 - 44	8 - 10	0,9 – 1,6
	8	140	60 - 70	10 – 12	1,2 – 1,9

### Offset change test

In this test we checked how AP behaved after we changed offset of results it was showing. We changed offset from page 2 to page 3 and next from page 3 to page 500.

Entries	Number of APs on test page	Change from page 2 to page 3 using DB strategy [avg. in seconds]	Change from page 3 to page 500 using DB strategy [avg. in seconds]	Change from page 2 to page 3 using Lucene strategy [avg. in seconds]	Change from page 3 to page 500 using Lucene strategy [avg. in seconds]
<b>25000</b>	1	10	375	0,3	0,3

### Conclusion

If web content number in system stays below 5000 this problem can be avoided by making warming up requests to application so Liferay can generate cache. In all other cases time in which page loads is unacceptable for users. The more AP we use the worse it gets.

Permission checks using Lucene are way faster than using DB and implementation effort (in terms of time) of this approach is low enough to become permanent fix for [LPS-22332](#).