	2.0.0 − 1530	2.0.0 − 1530
Runtime (in hr)	3.27	NA
Avg. Drain Rate	423.33	NANA
Peak Disk (GB)	49.96	NA
Peak Memory (GB)	10.88	NA
Avg. OPS	8.73K	NANA
Avg. mem memcached (GB)	10.21	NA
Avg. mem beam.smp (MB)	554.92	NA
Latency−get (90th) (ms)	0.3	NA
Latency−get (95th) (ms)	0.53	NA
Latency−get (99th) (ms)	2.77	NA
Latency−set (90th) (ms)	0.47	NA
Latency−set (95th) (ms)	0.67	NA
Latency−set (99th) (ms)	2.99	NA
Latency−query (80th) (ms)	6.91	NA
Latency−query (90th) (ms)	9.95	NA
Latency−query (95th) (ms)	21.32	NA
Latency−query (99th) (ms)	81.62	NA
Latency−query (99.9th) (ms)	325.17	NA
Avg. QPS	670.57	NA
Rebalance Time (sec)	0	NA
Testrunner Version	799f985	NA


ops/sec
9.77K
7.81K
5.86K
3.91K
1.95K
0
ops/sec
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of ops per second
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

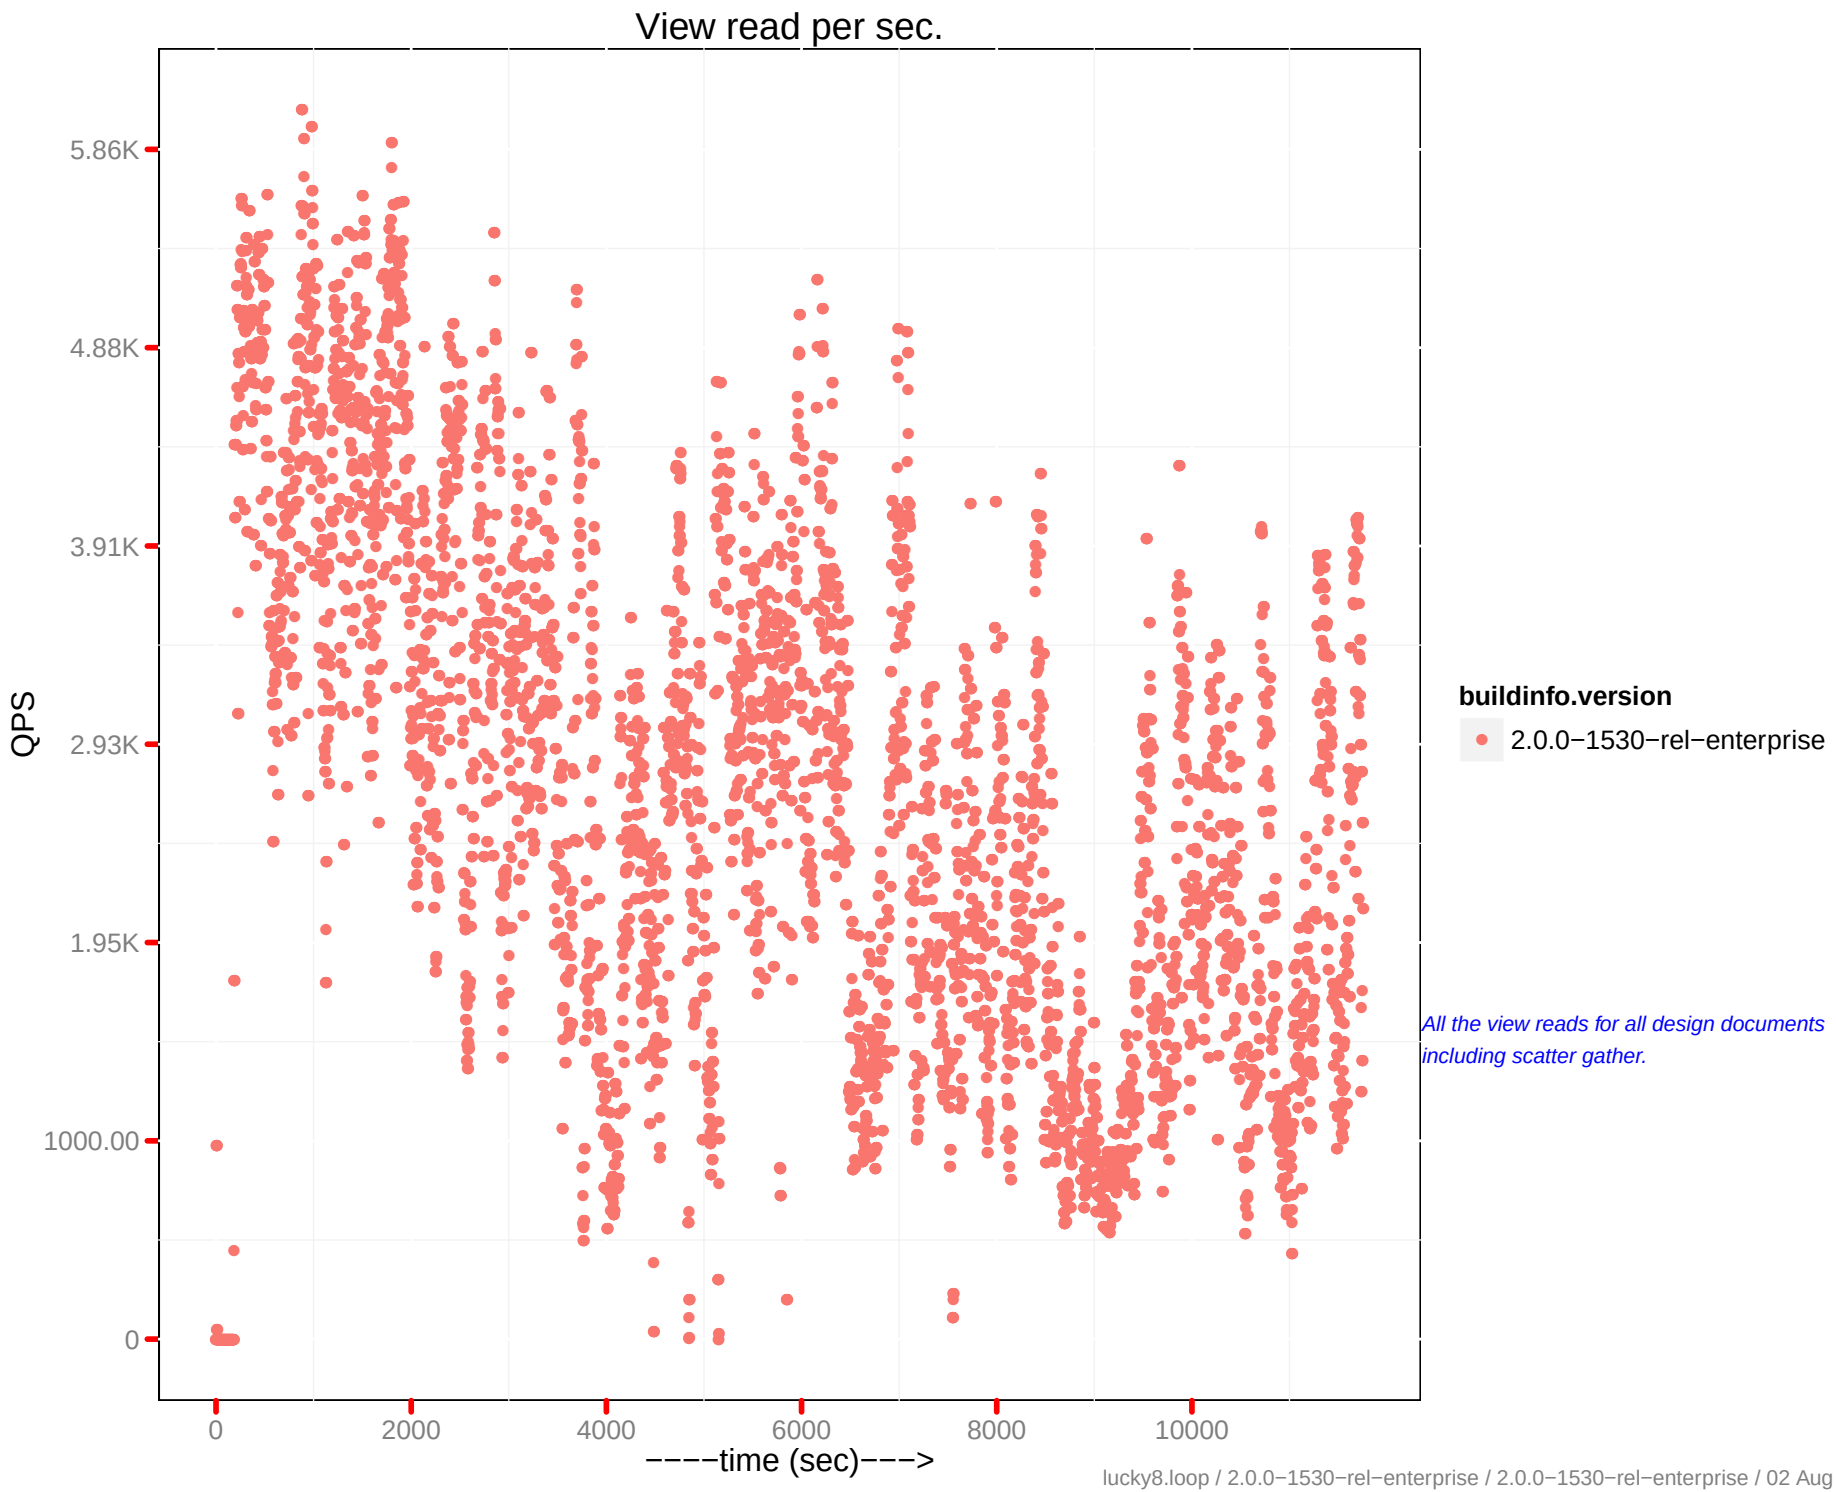
View read per sec.
5.86K
4.88K
3.91K
2.93K
1.95K
1000.00
0
QPS
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
All the view reads for all design documents including scatter gather.
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

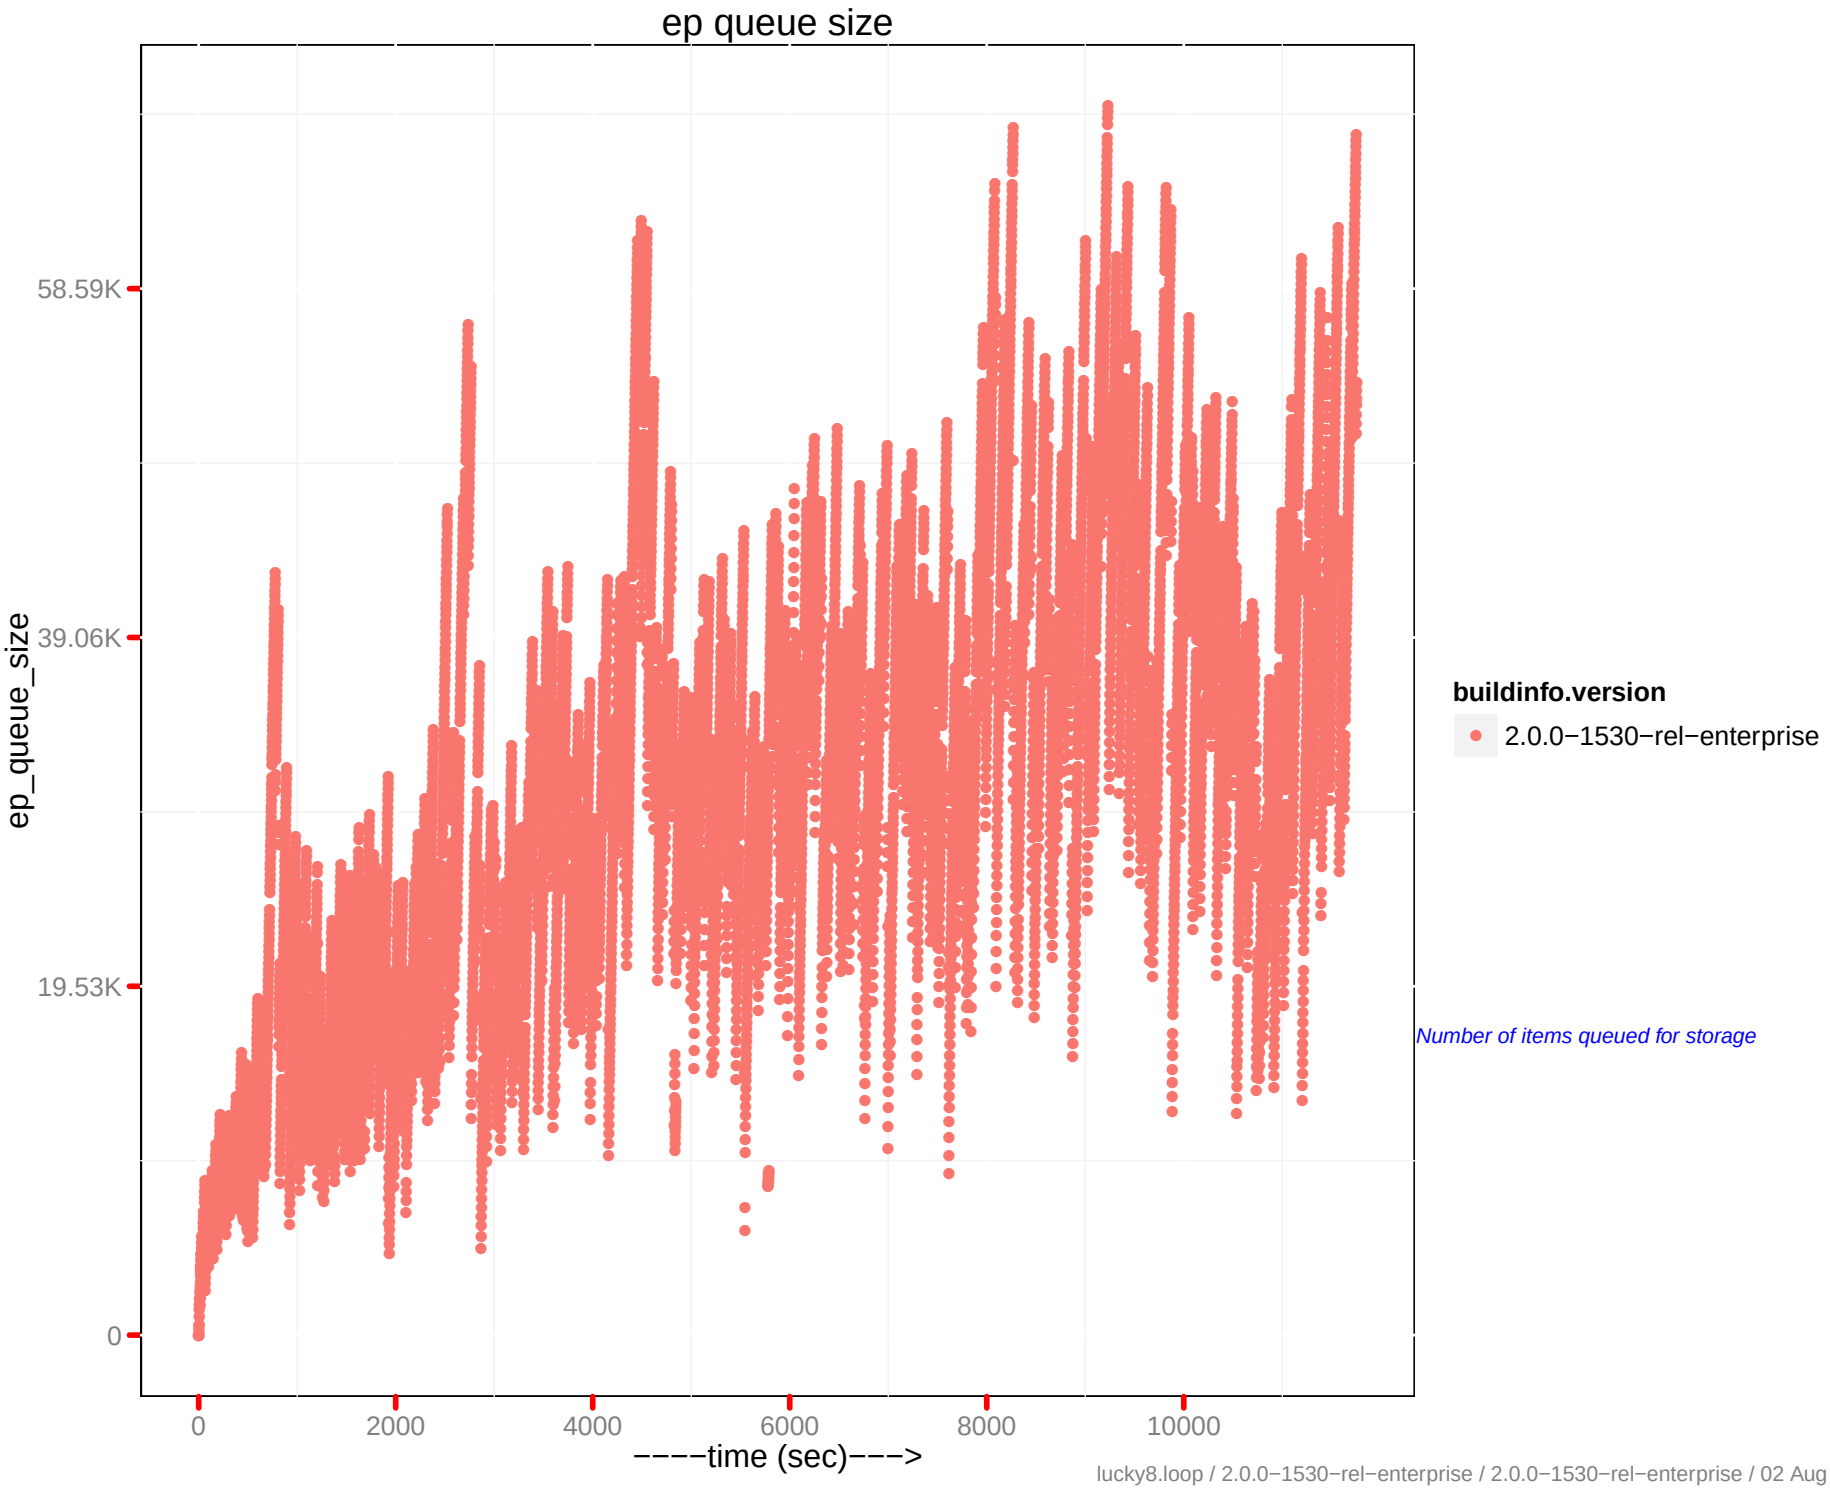
ep queue size
58.59K
39.06K
19.53K
0
ep_queue_size
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of items queued for storage
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

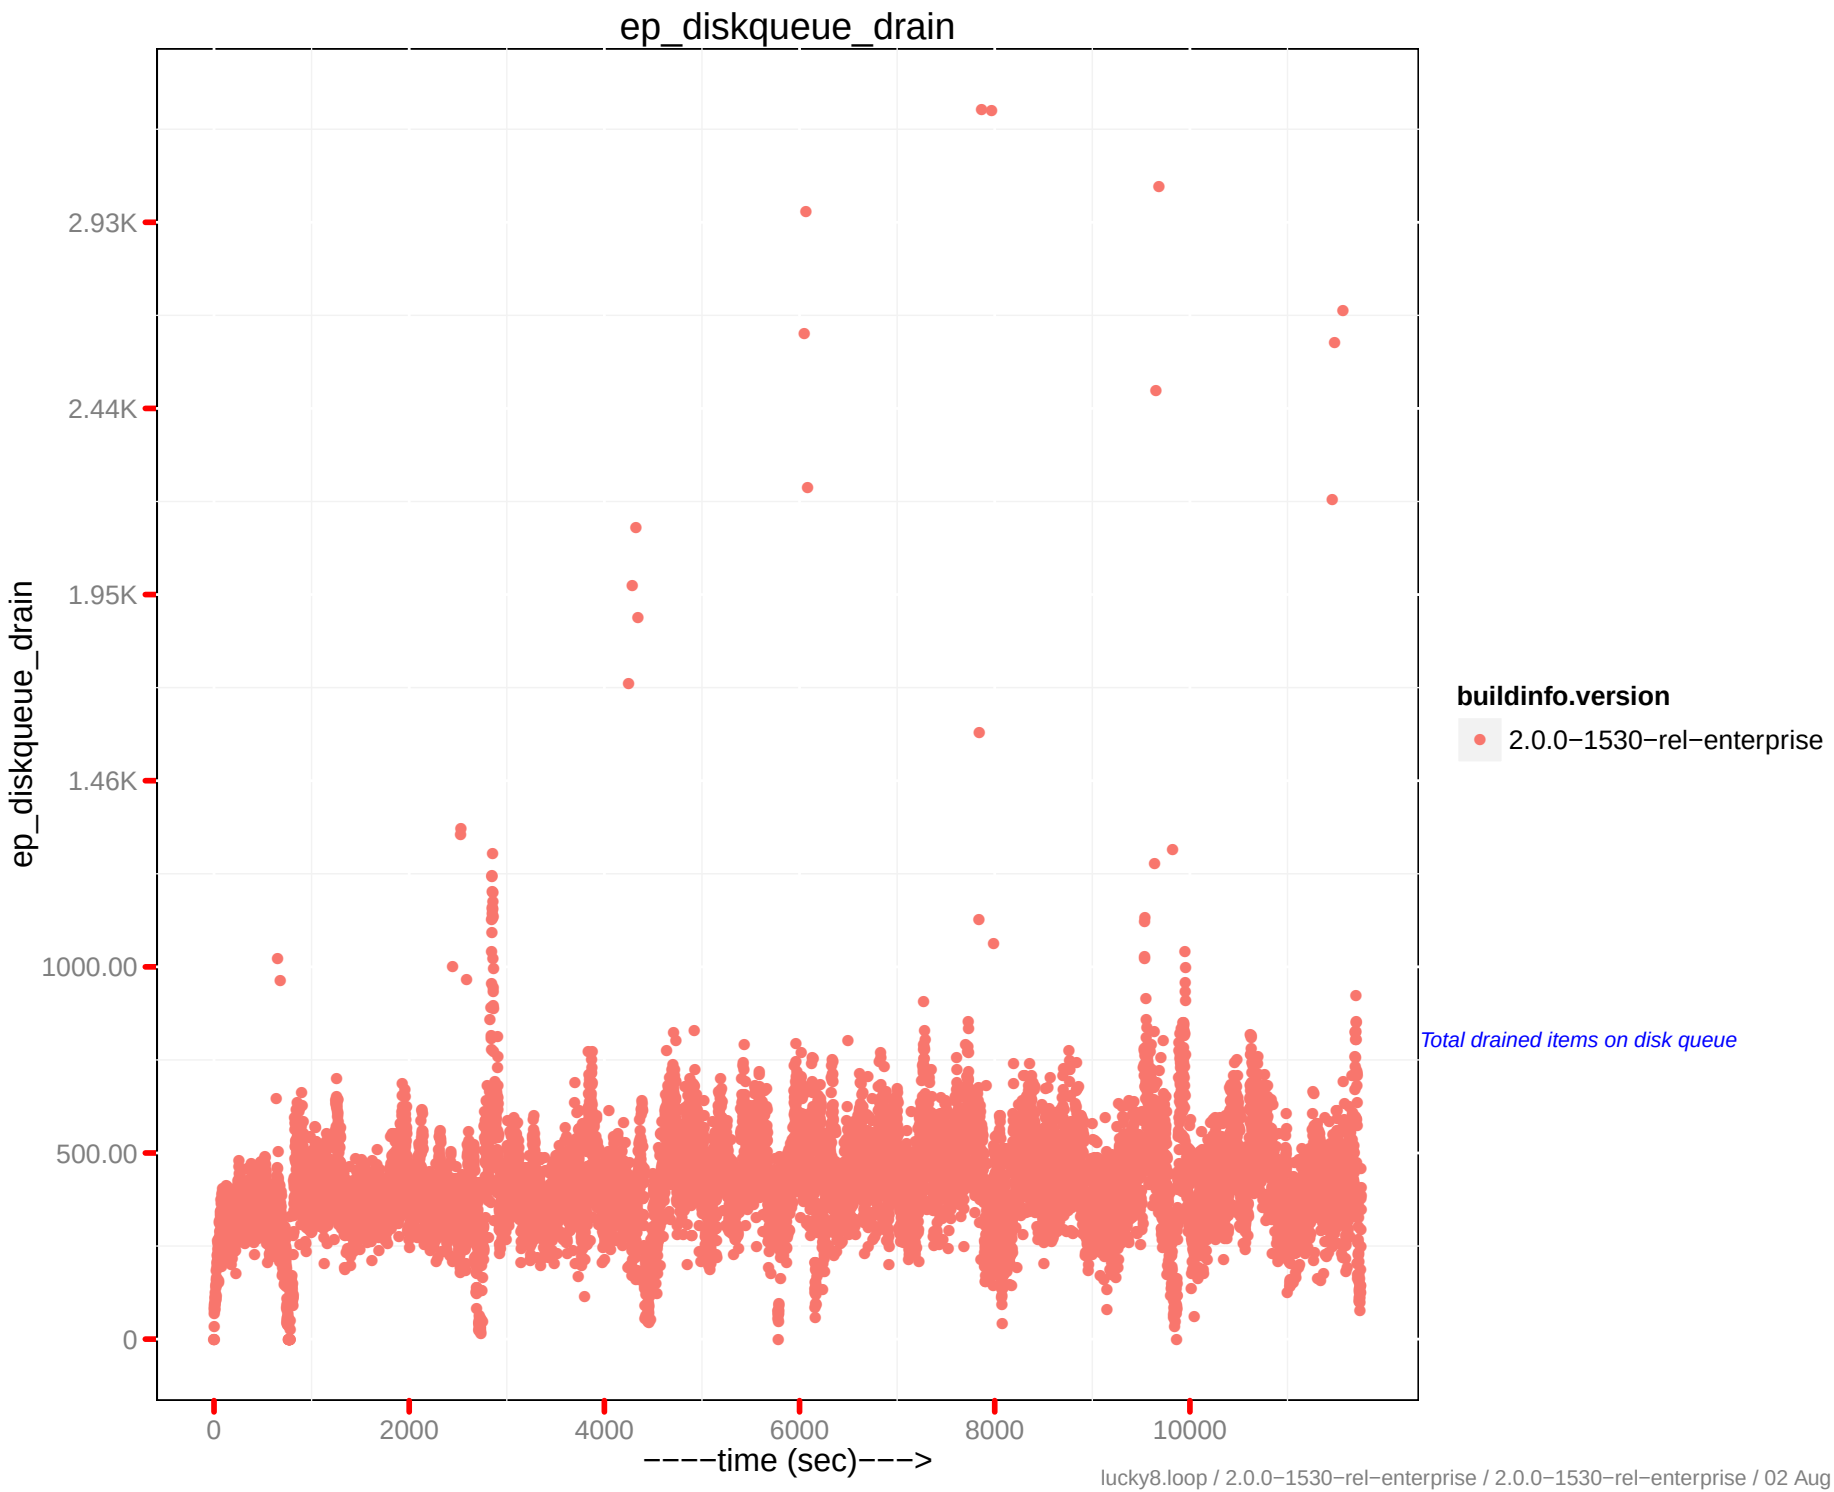
ep_diskqueue_drain
2.93K
2.44K
1.95K
1.46K
1000.00
500.00
0
ep_diskqueue_drain
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
Total drained items on disk queue
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

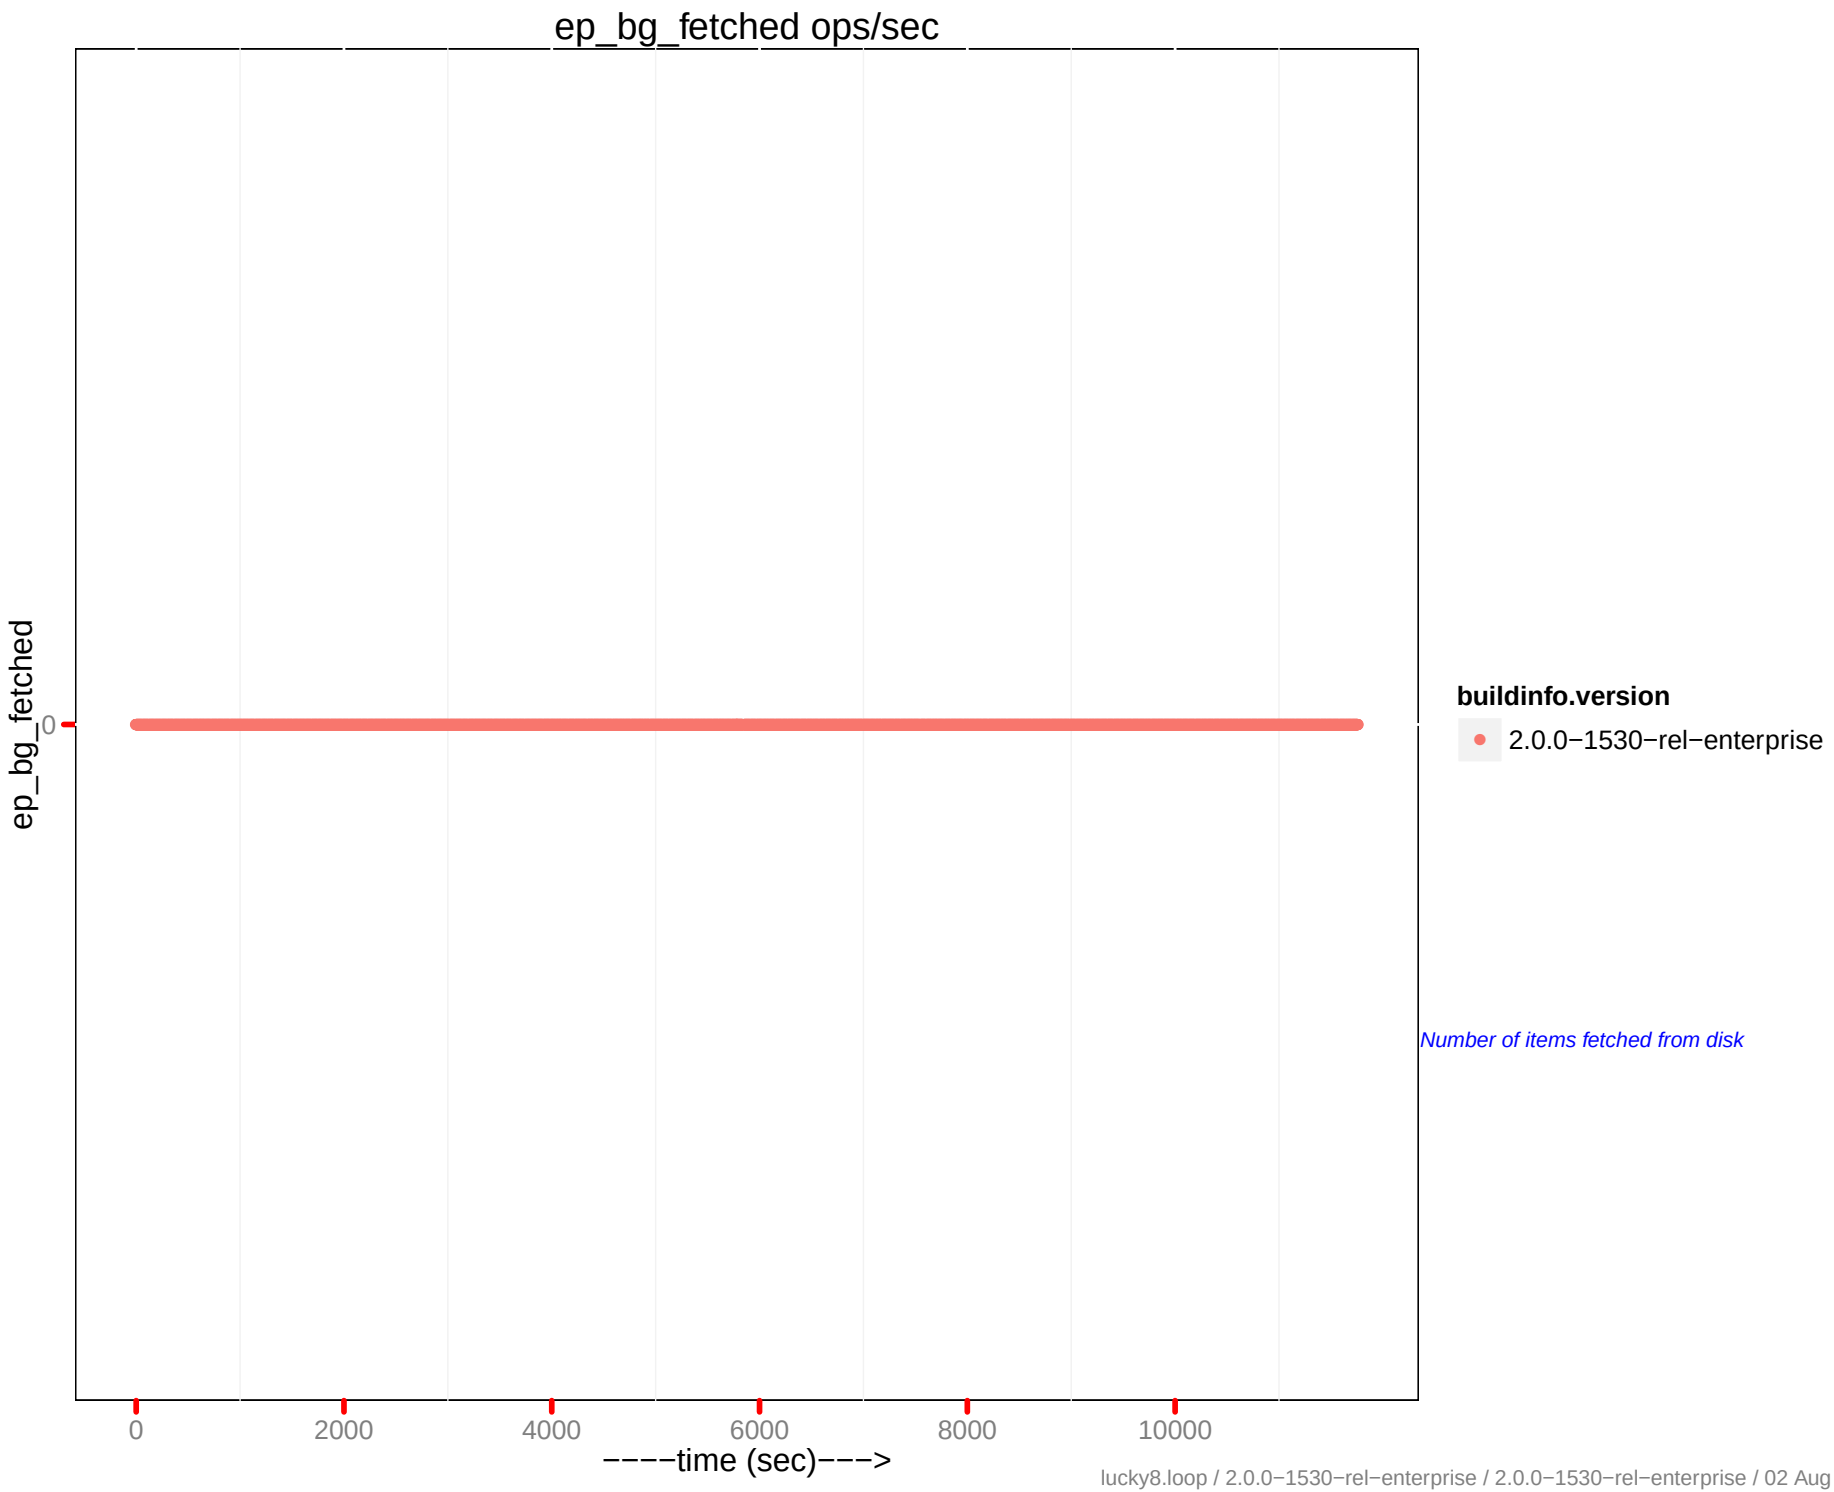
ep_bg_fetched ops/sec
ep_bg_fetched
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of items fetched from disk
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

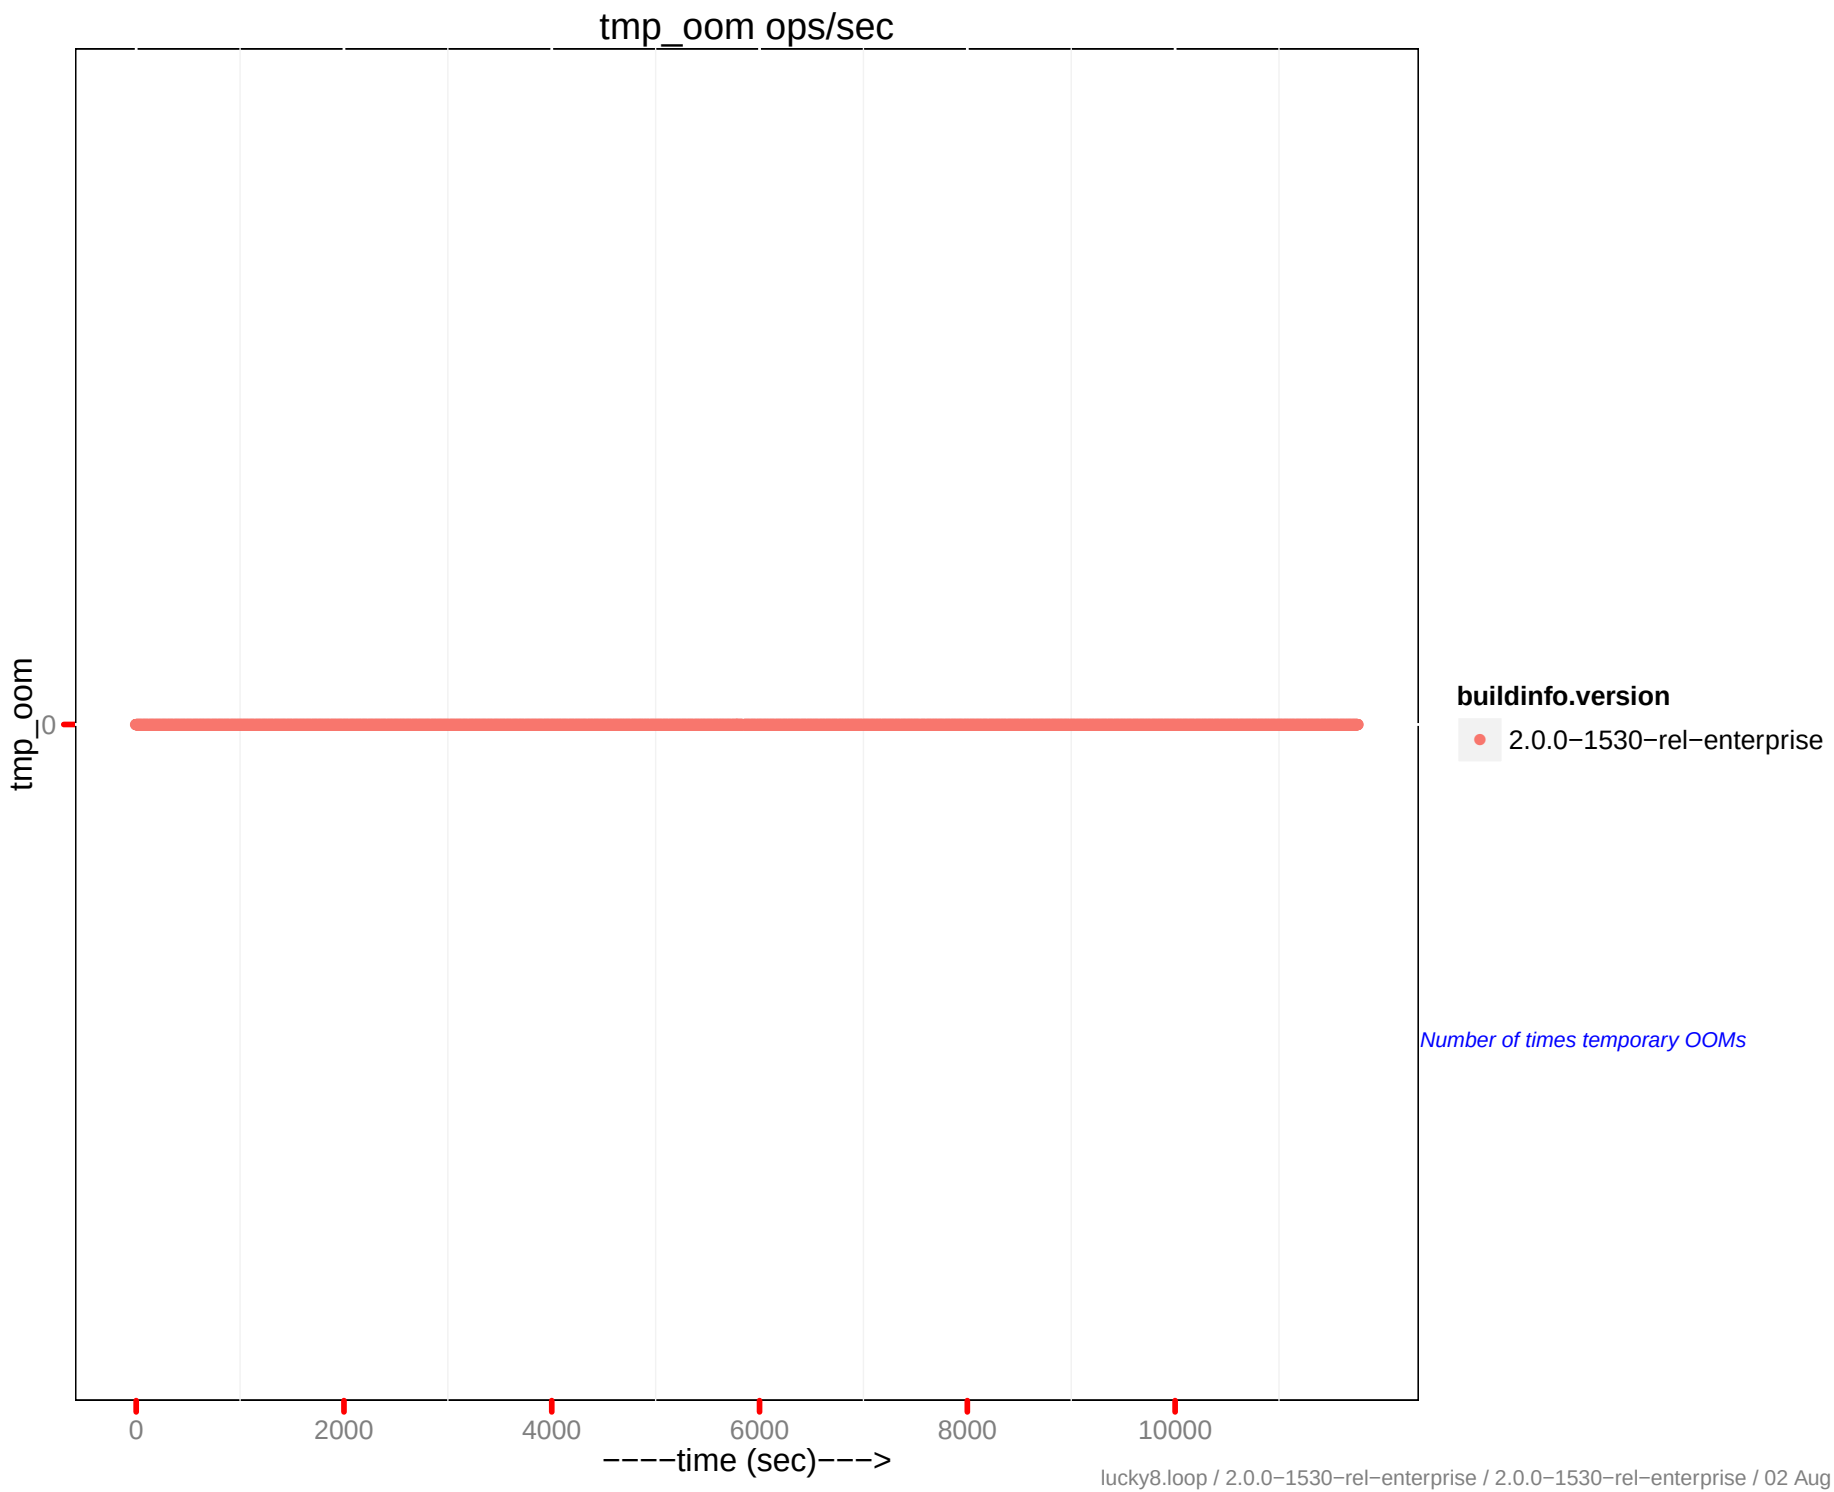
tmp_oom ops/sec
tmp_oom
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
Number of times temporary OOMs

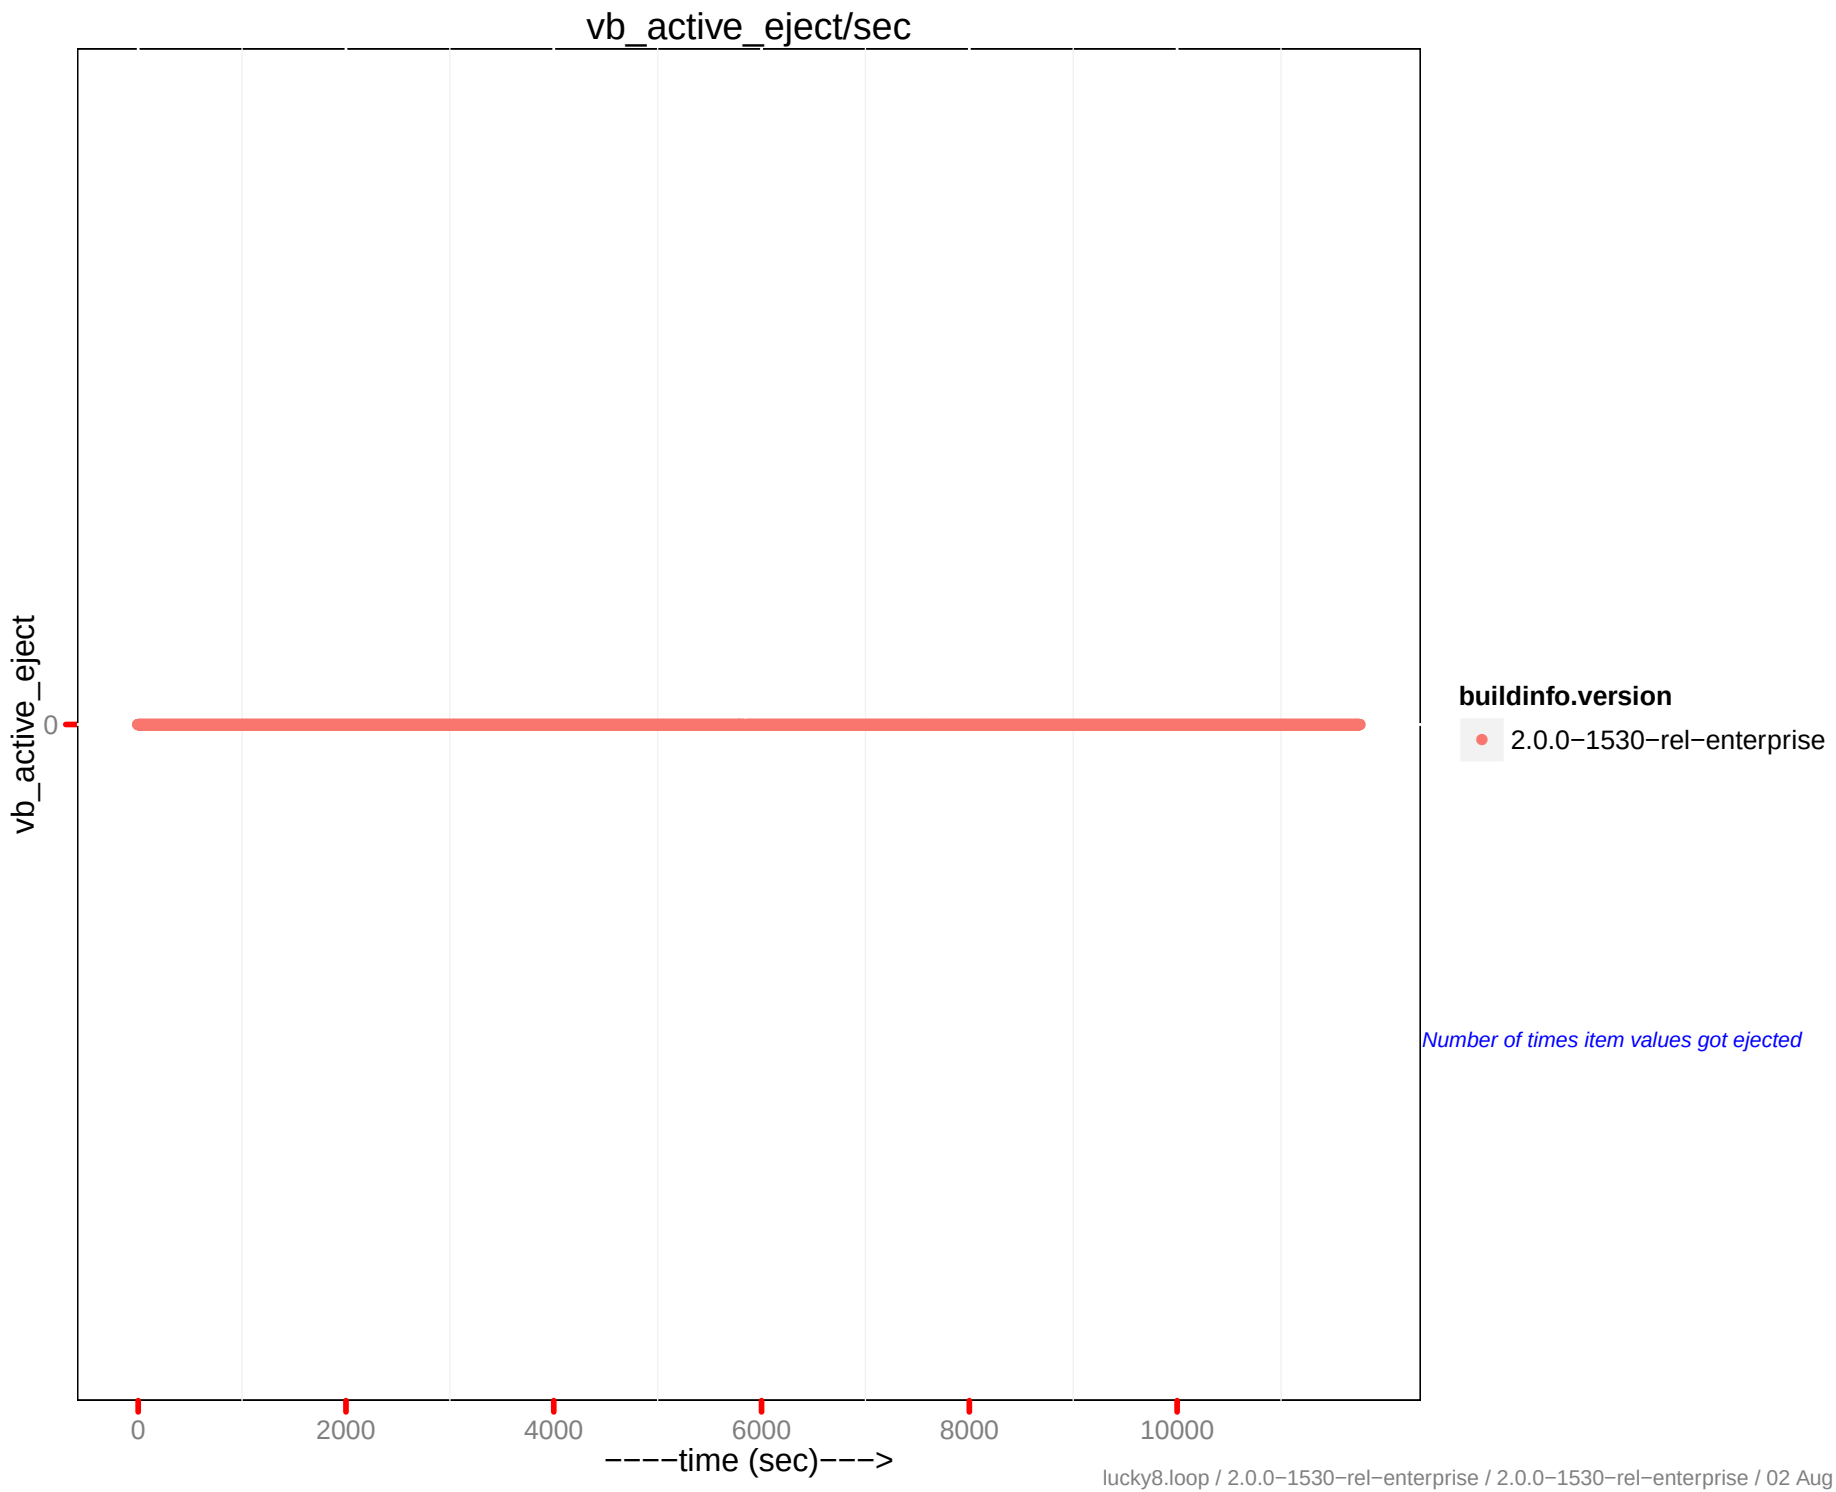
vb_active_eject/sec
vb_active_eject
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of times item values got ejected

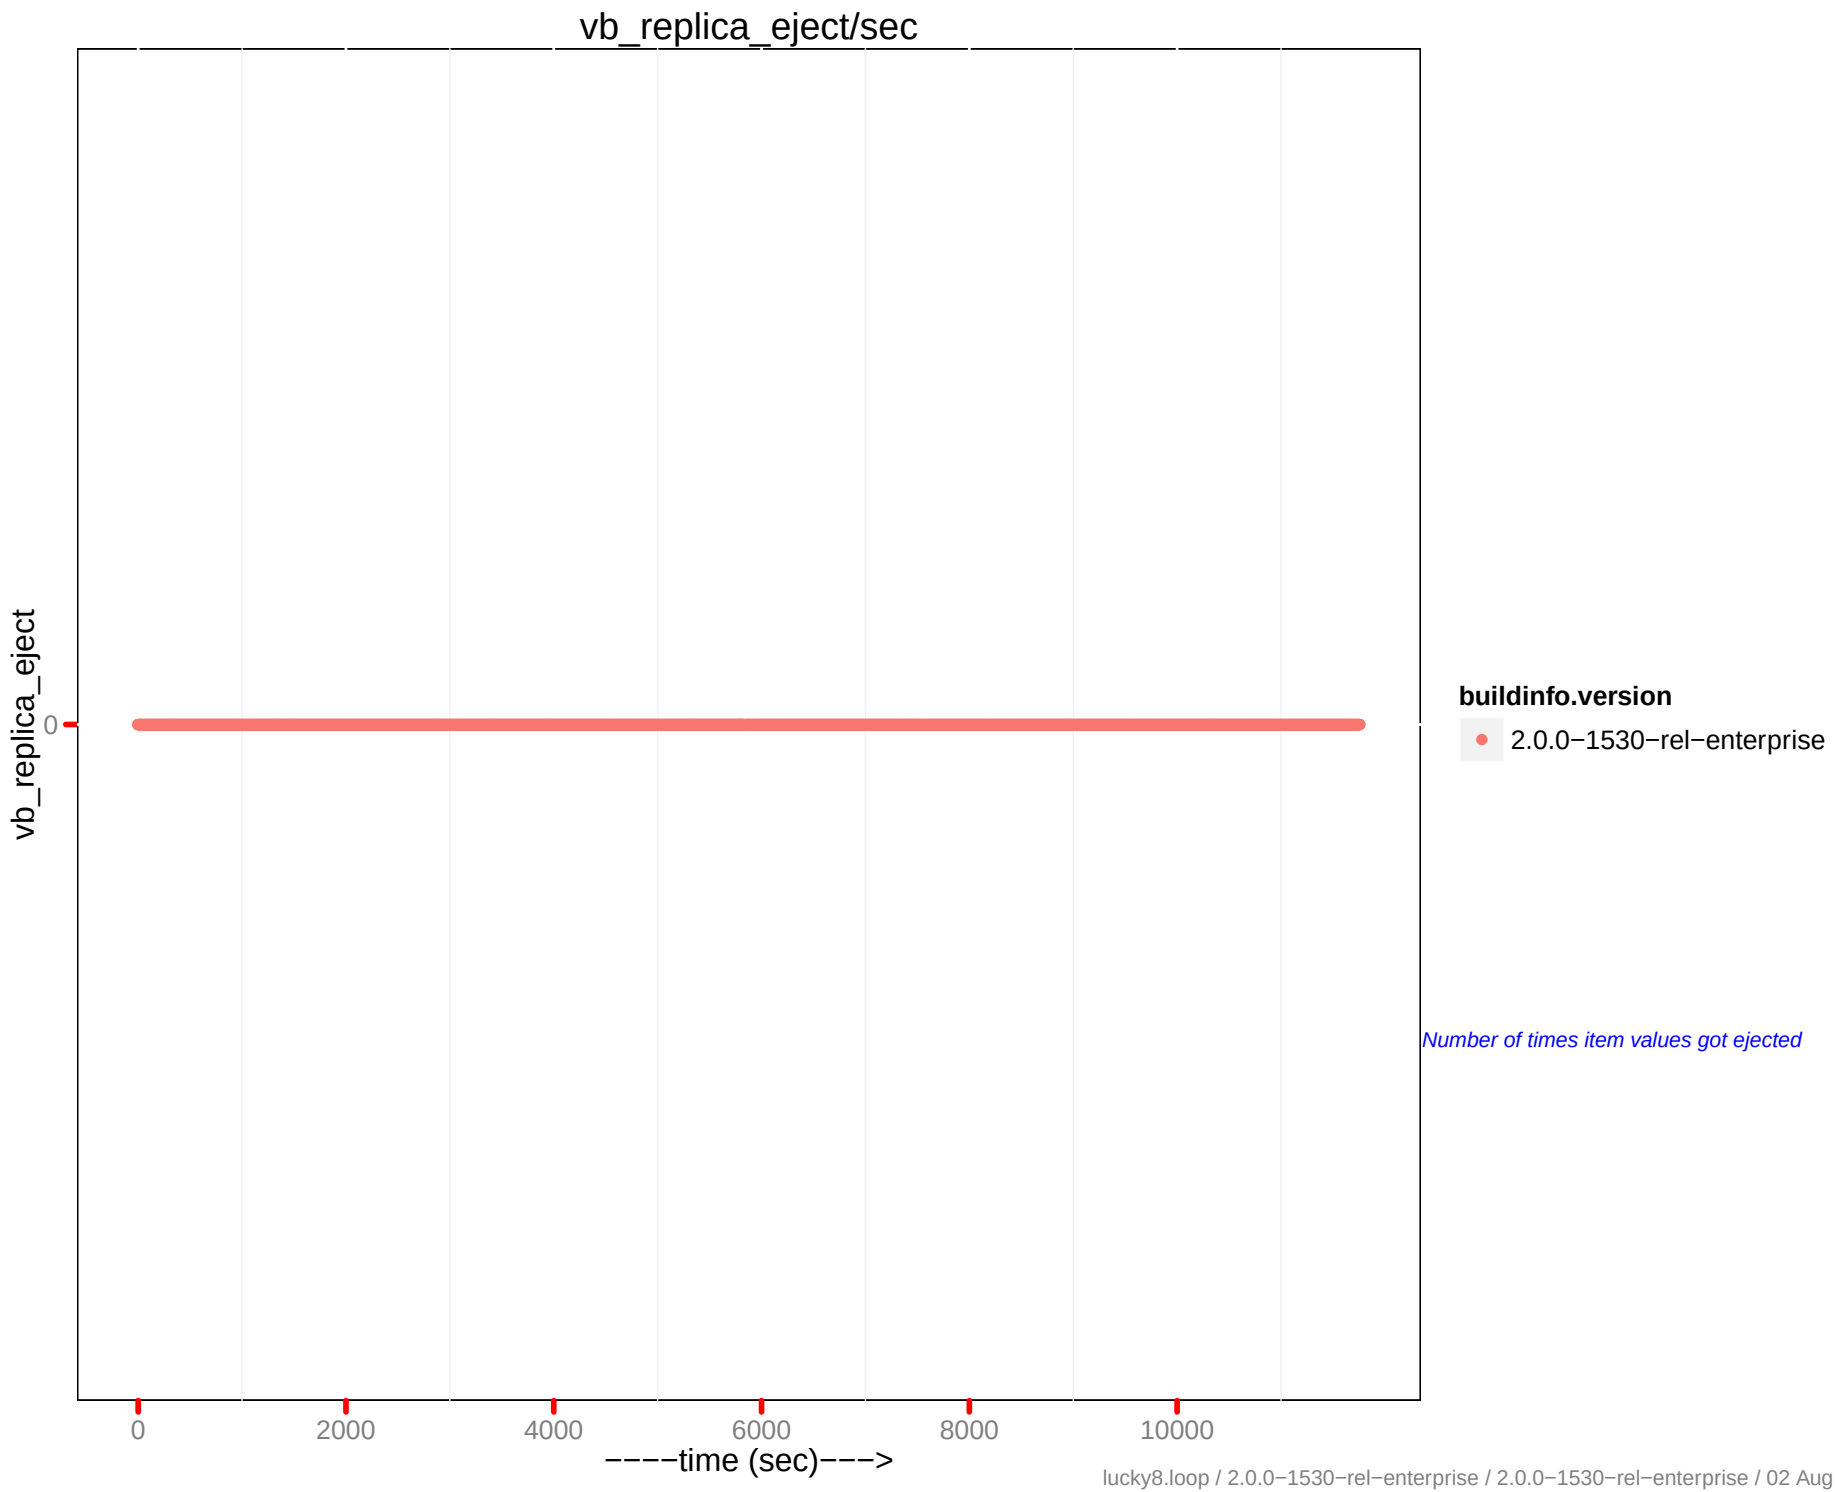
vb_replica_eject/sec
vb_replica_eject
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of times item values got ejected

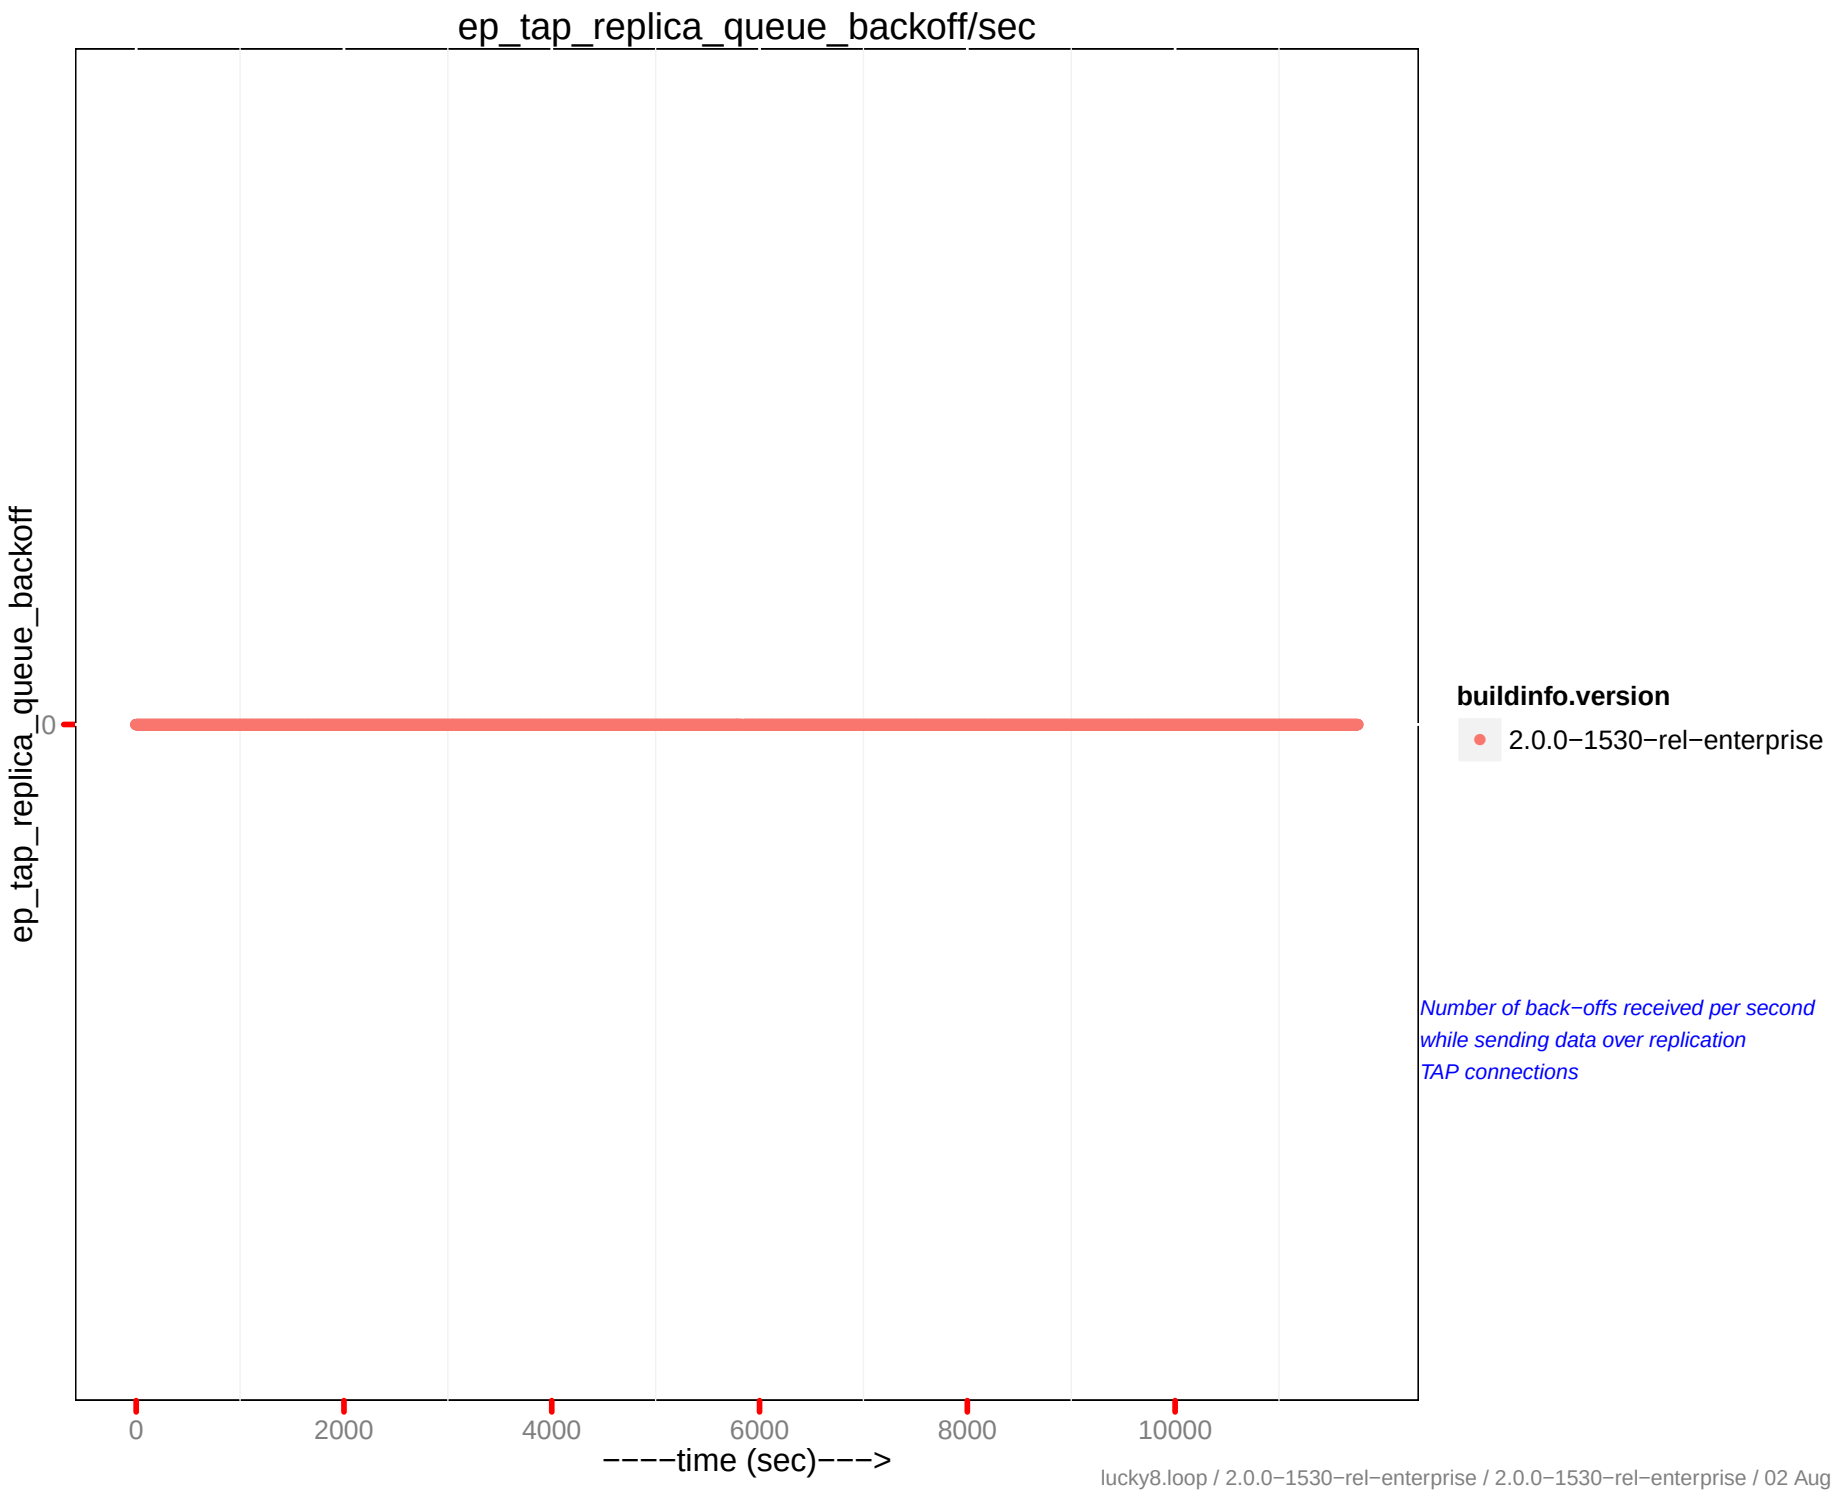
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of back-offs received per second
while sending data over replication
TAP connections

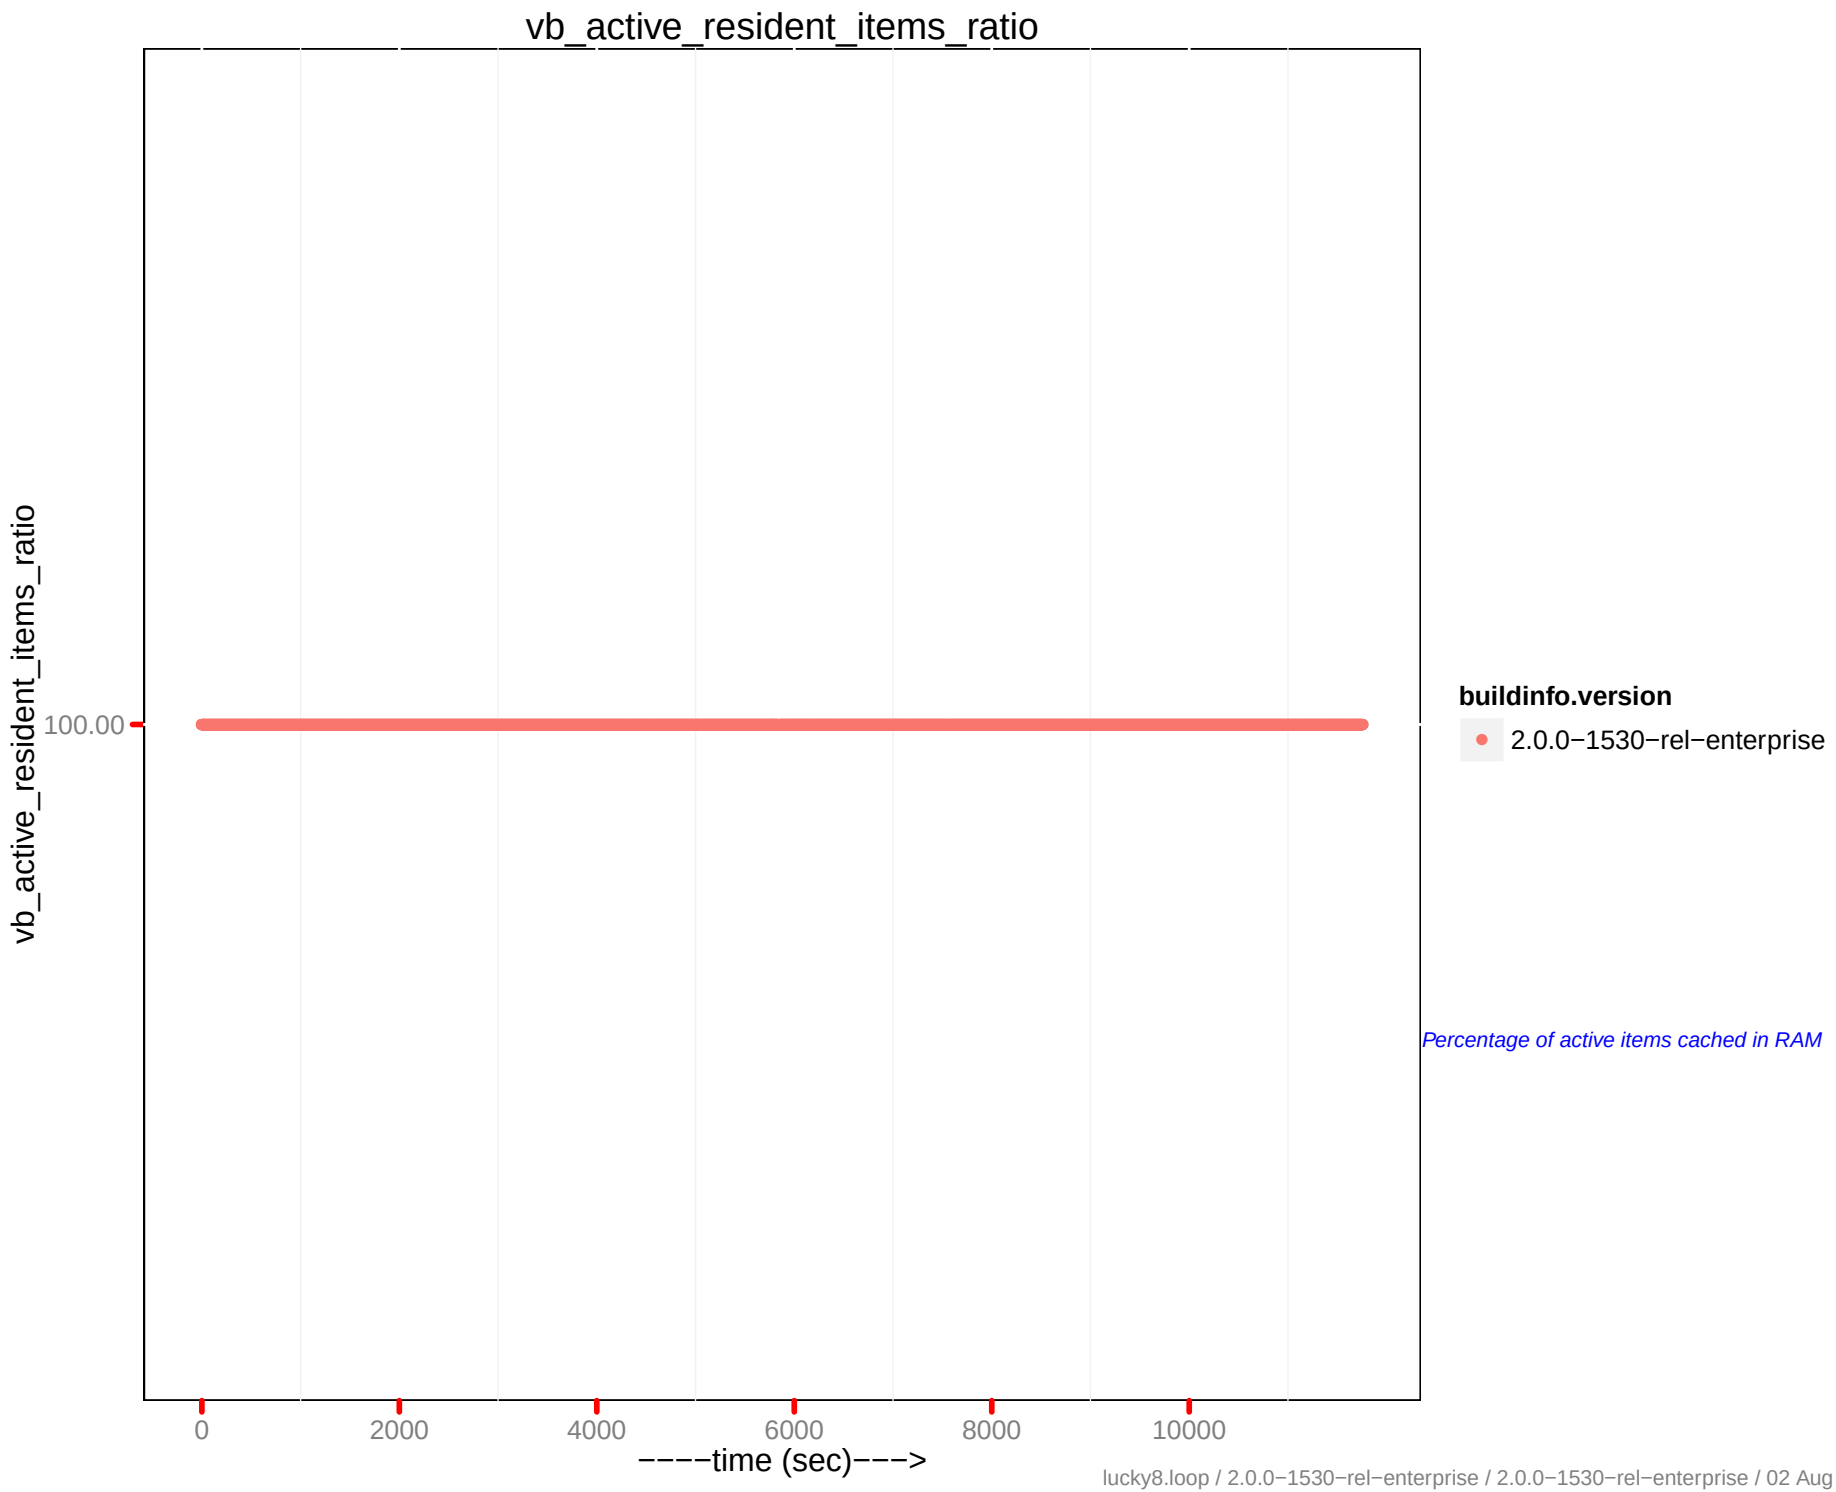
vb_active_resident_items_ratio
vb_active_resident_items_ratio
100.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Percentage of active items cached in RAM

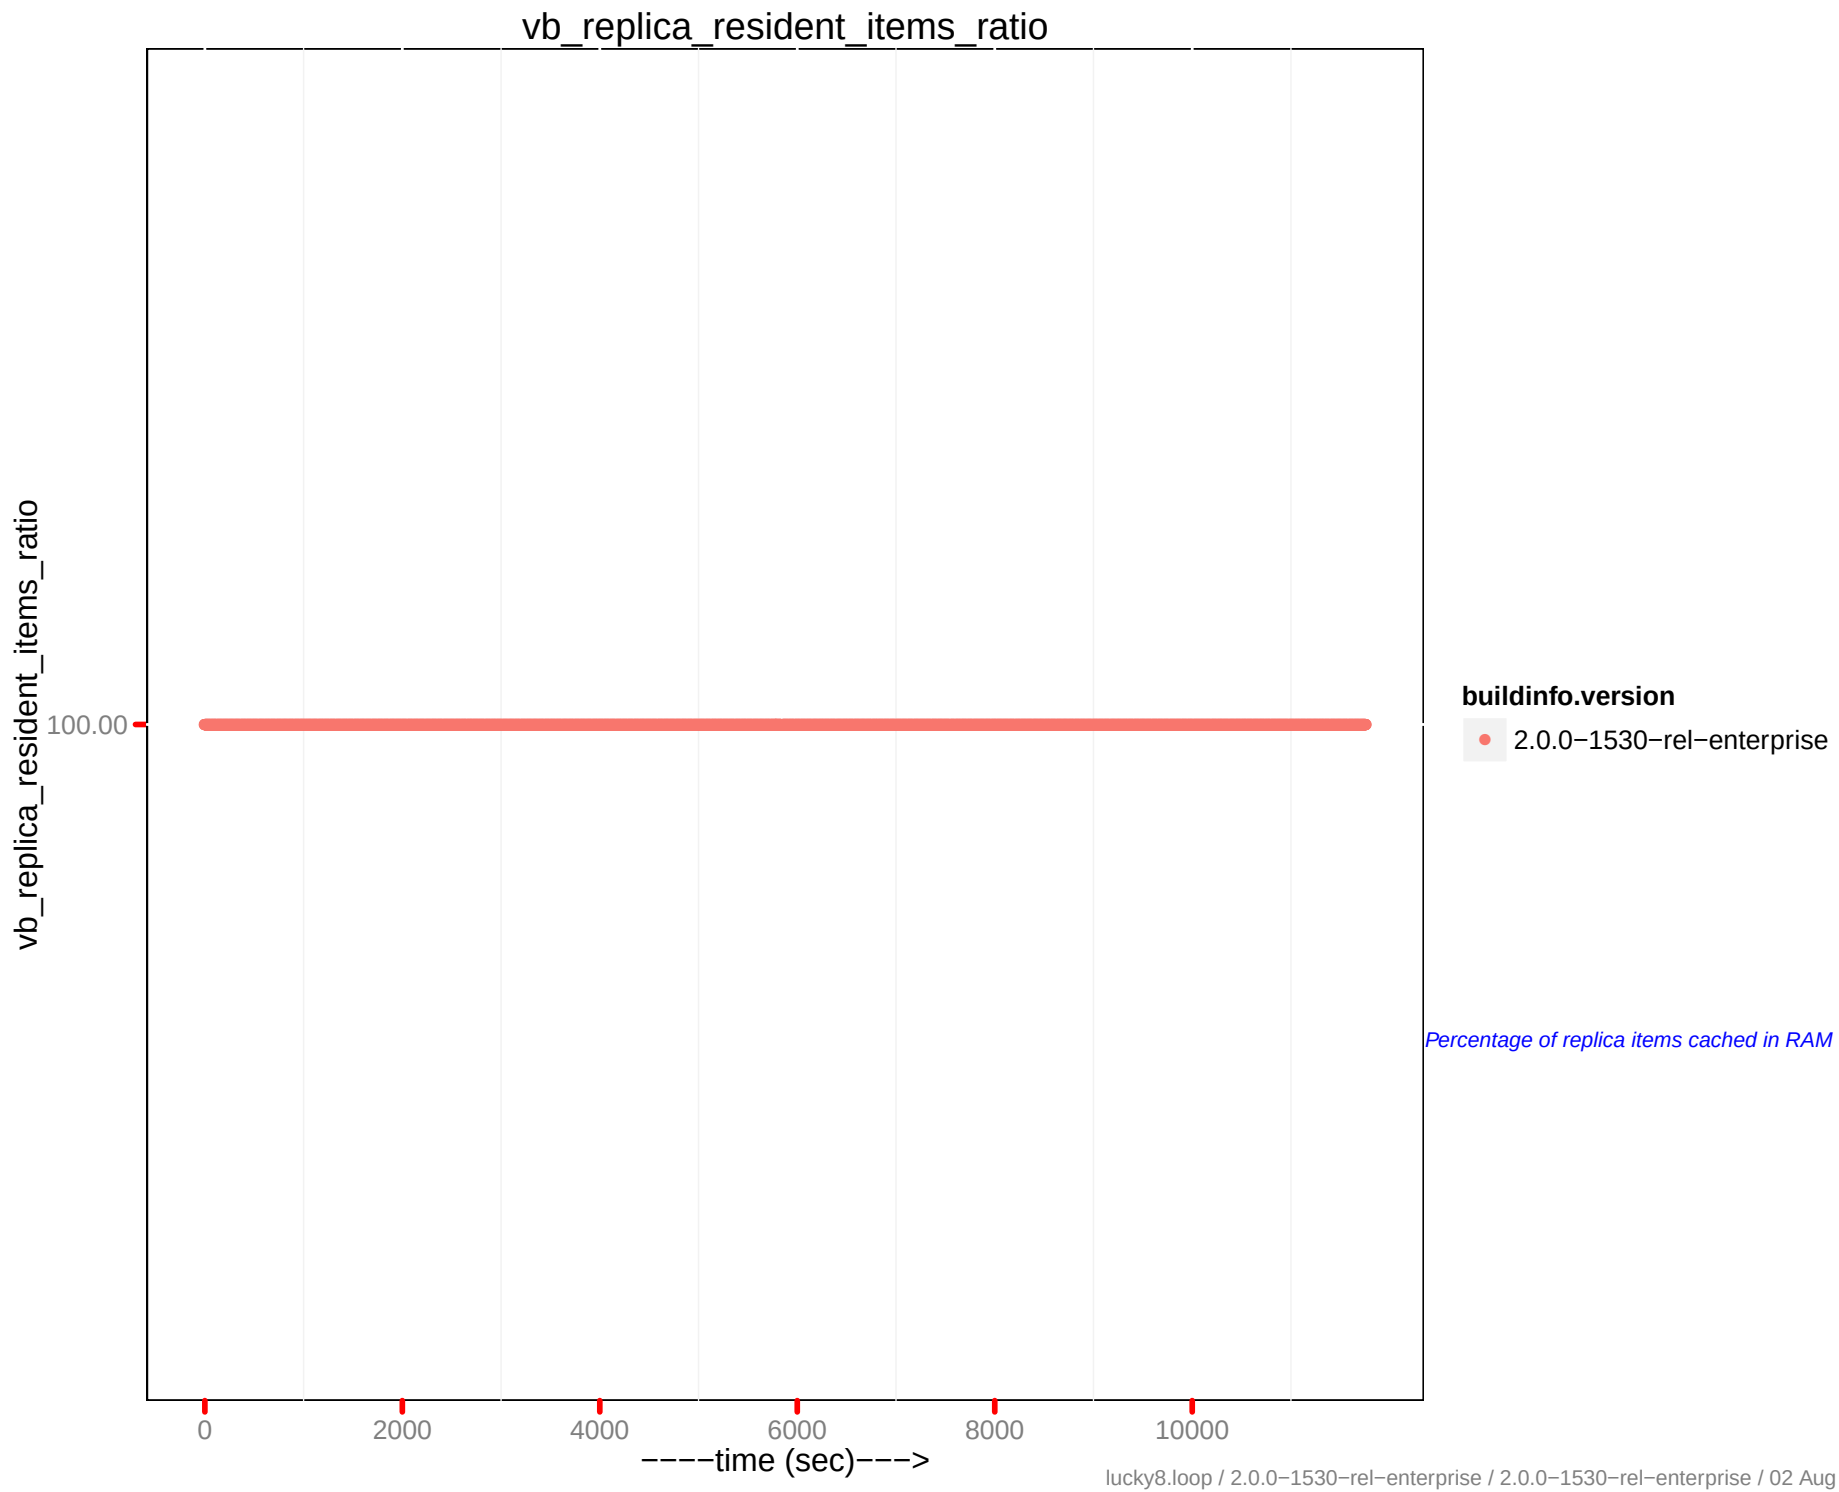
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
100.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Percentage of replica items cached in RAM
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

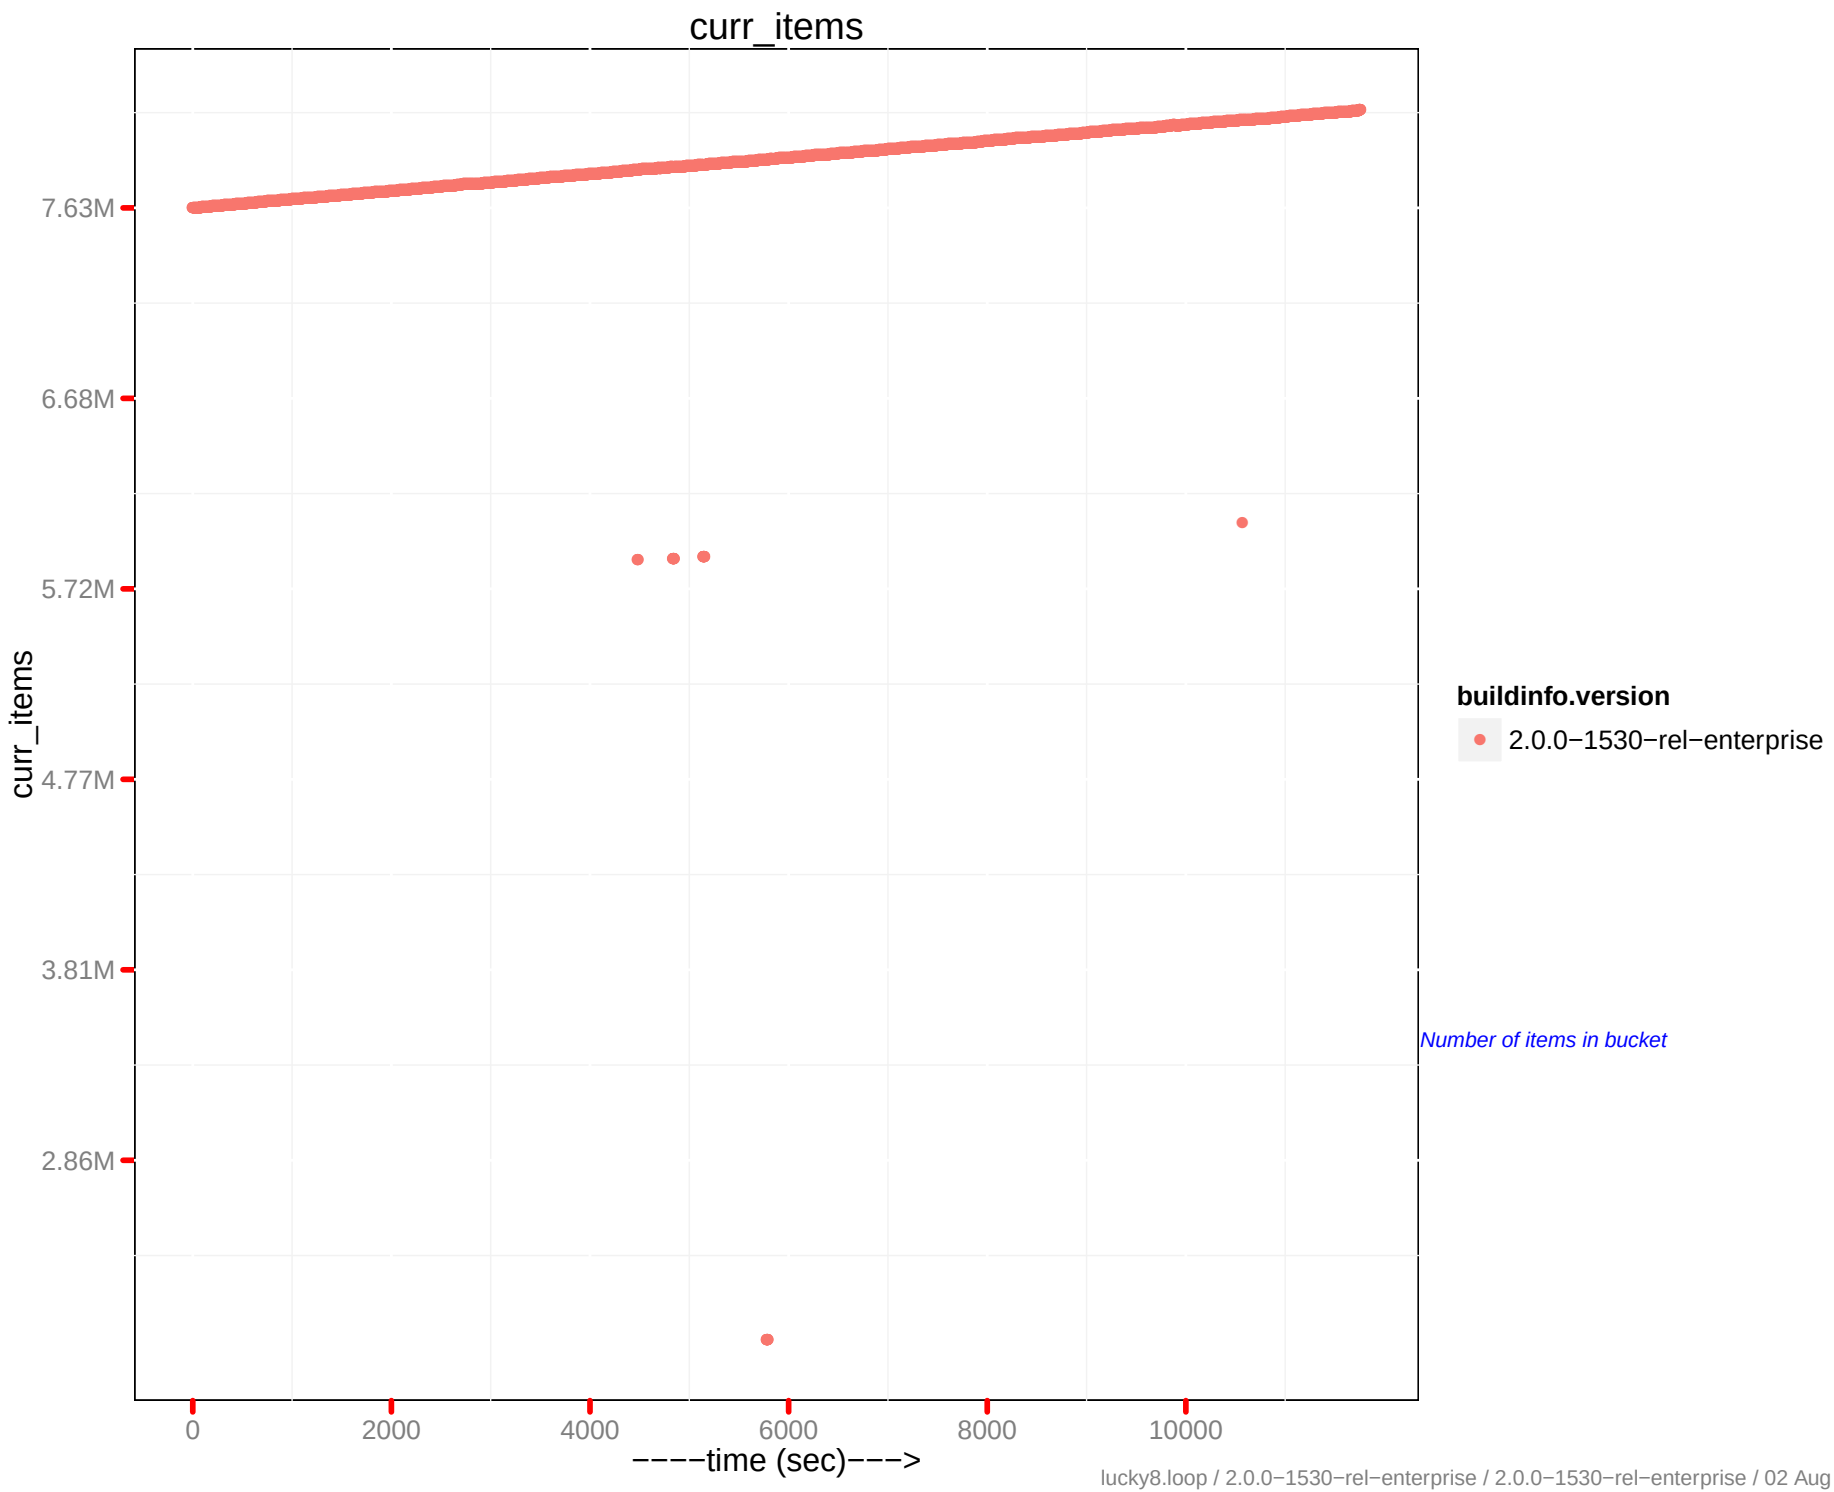
curr_items
7.63M
6.68M
5.72M
4.77M
3.81M
2.86M
curr_items
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of items in bucket
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

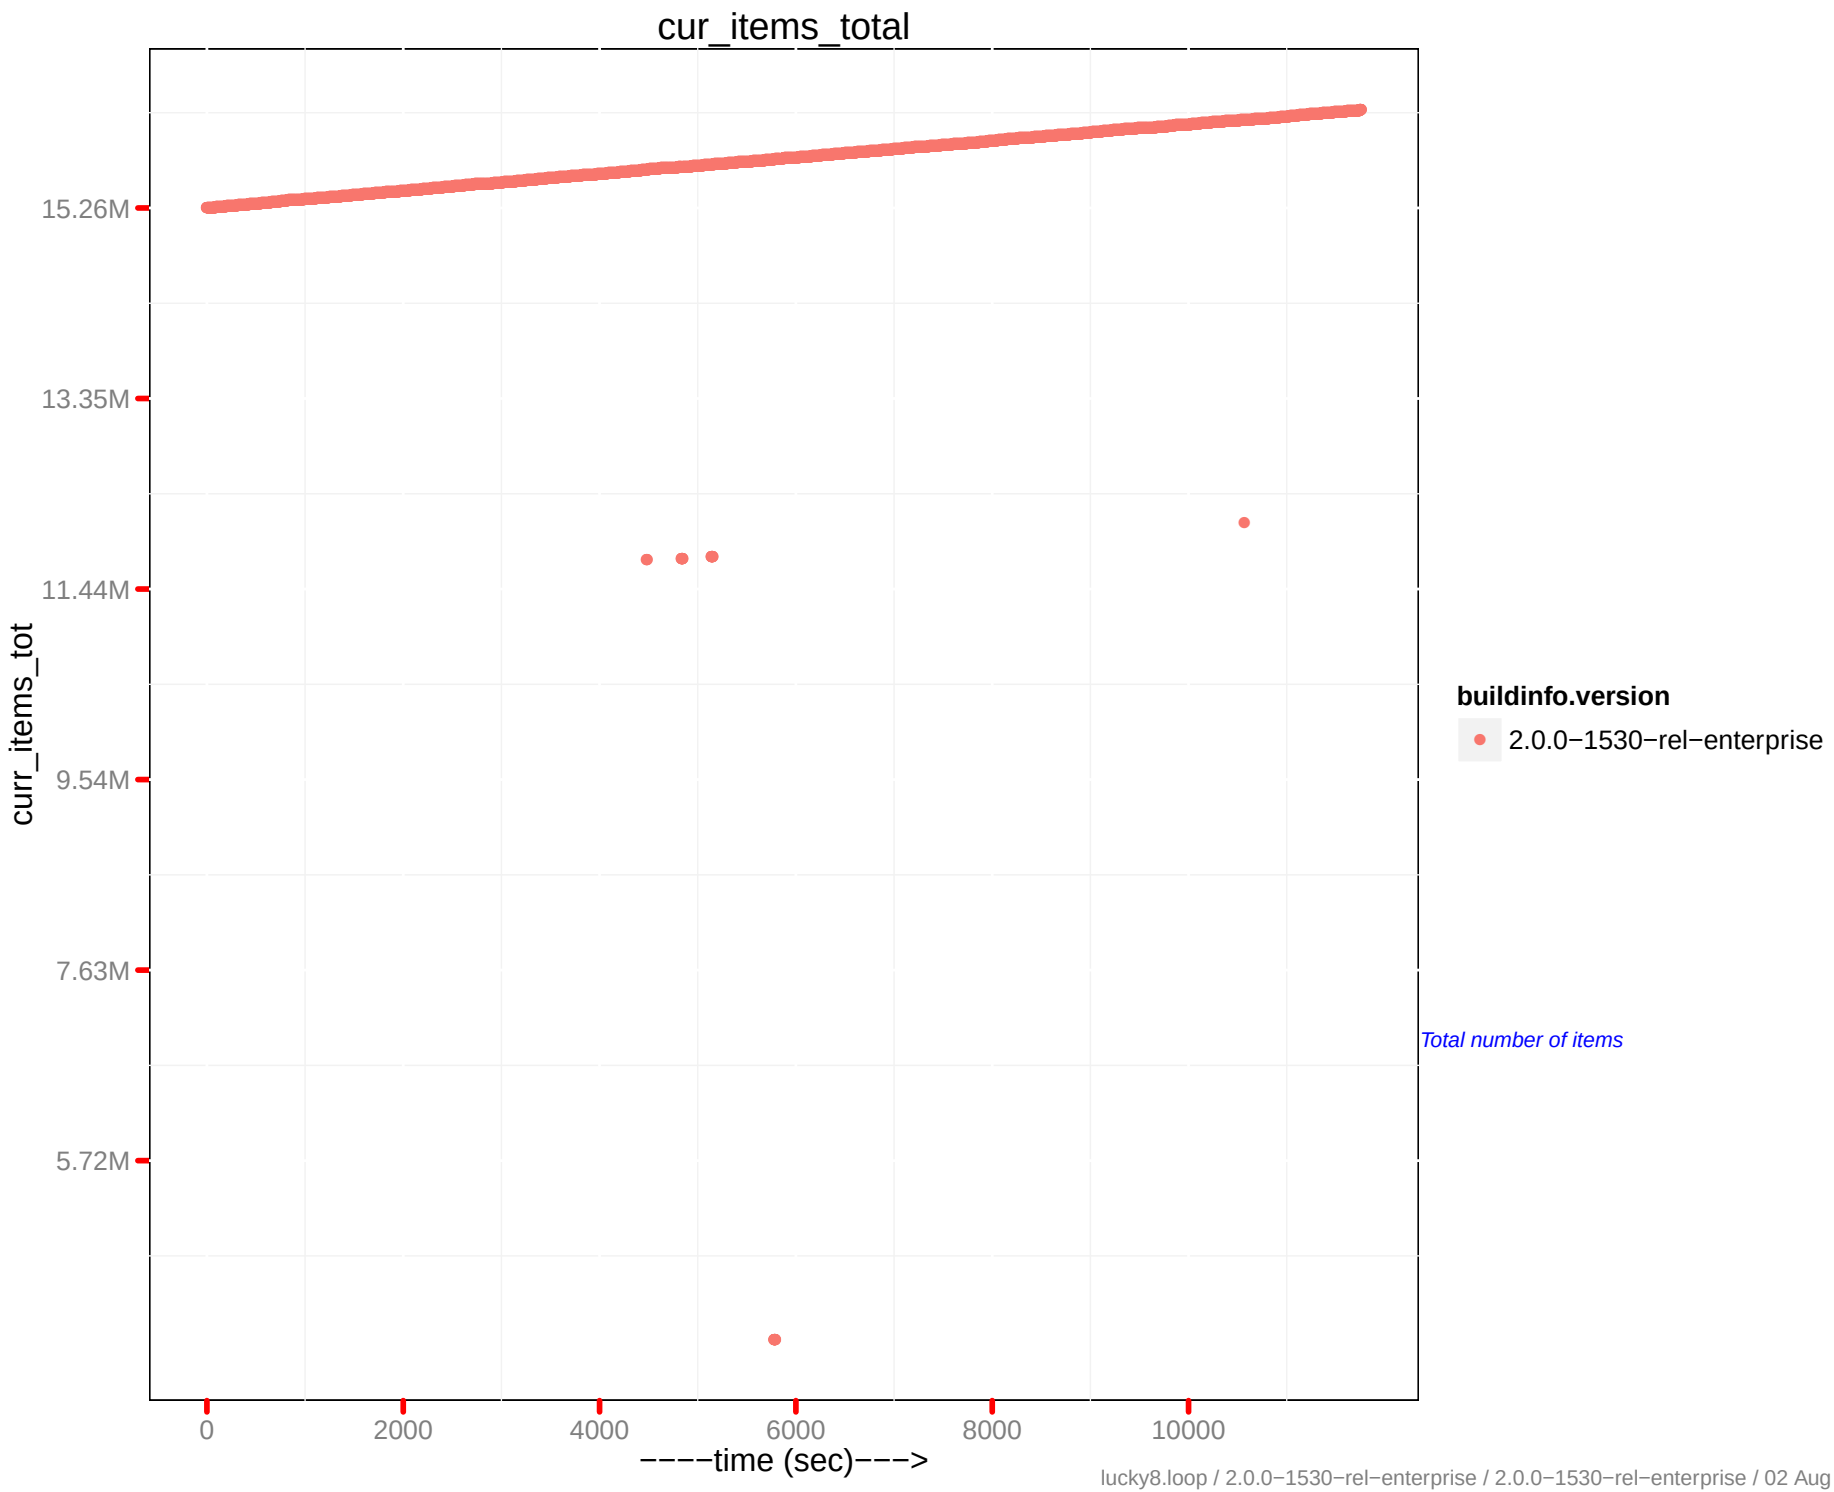
cur_items_total
15.26M
13.35M
11.44M
9.54M
7.63M
5.72M
curr_items_tot
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Total number of items
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

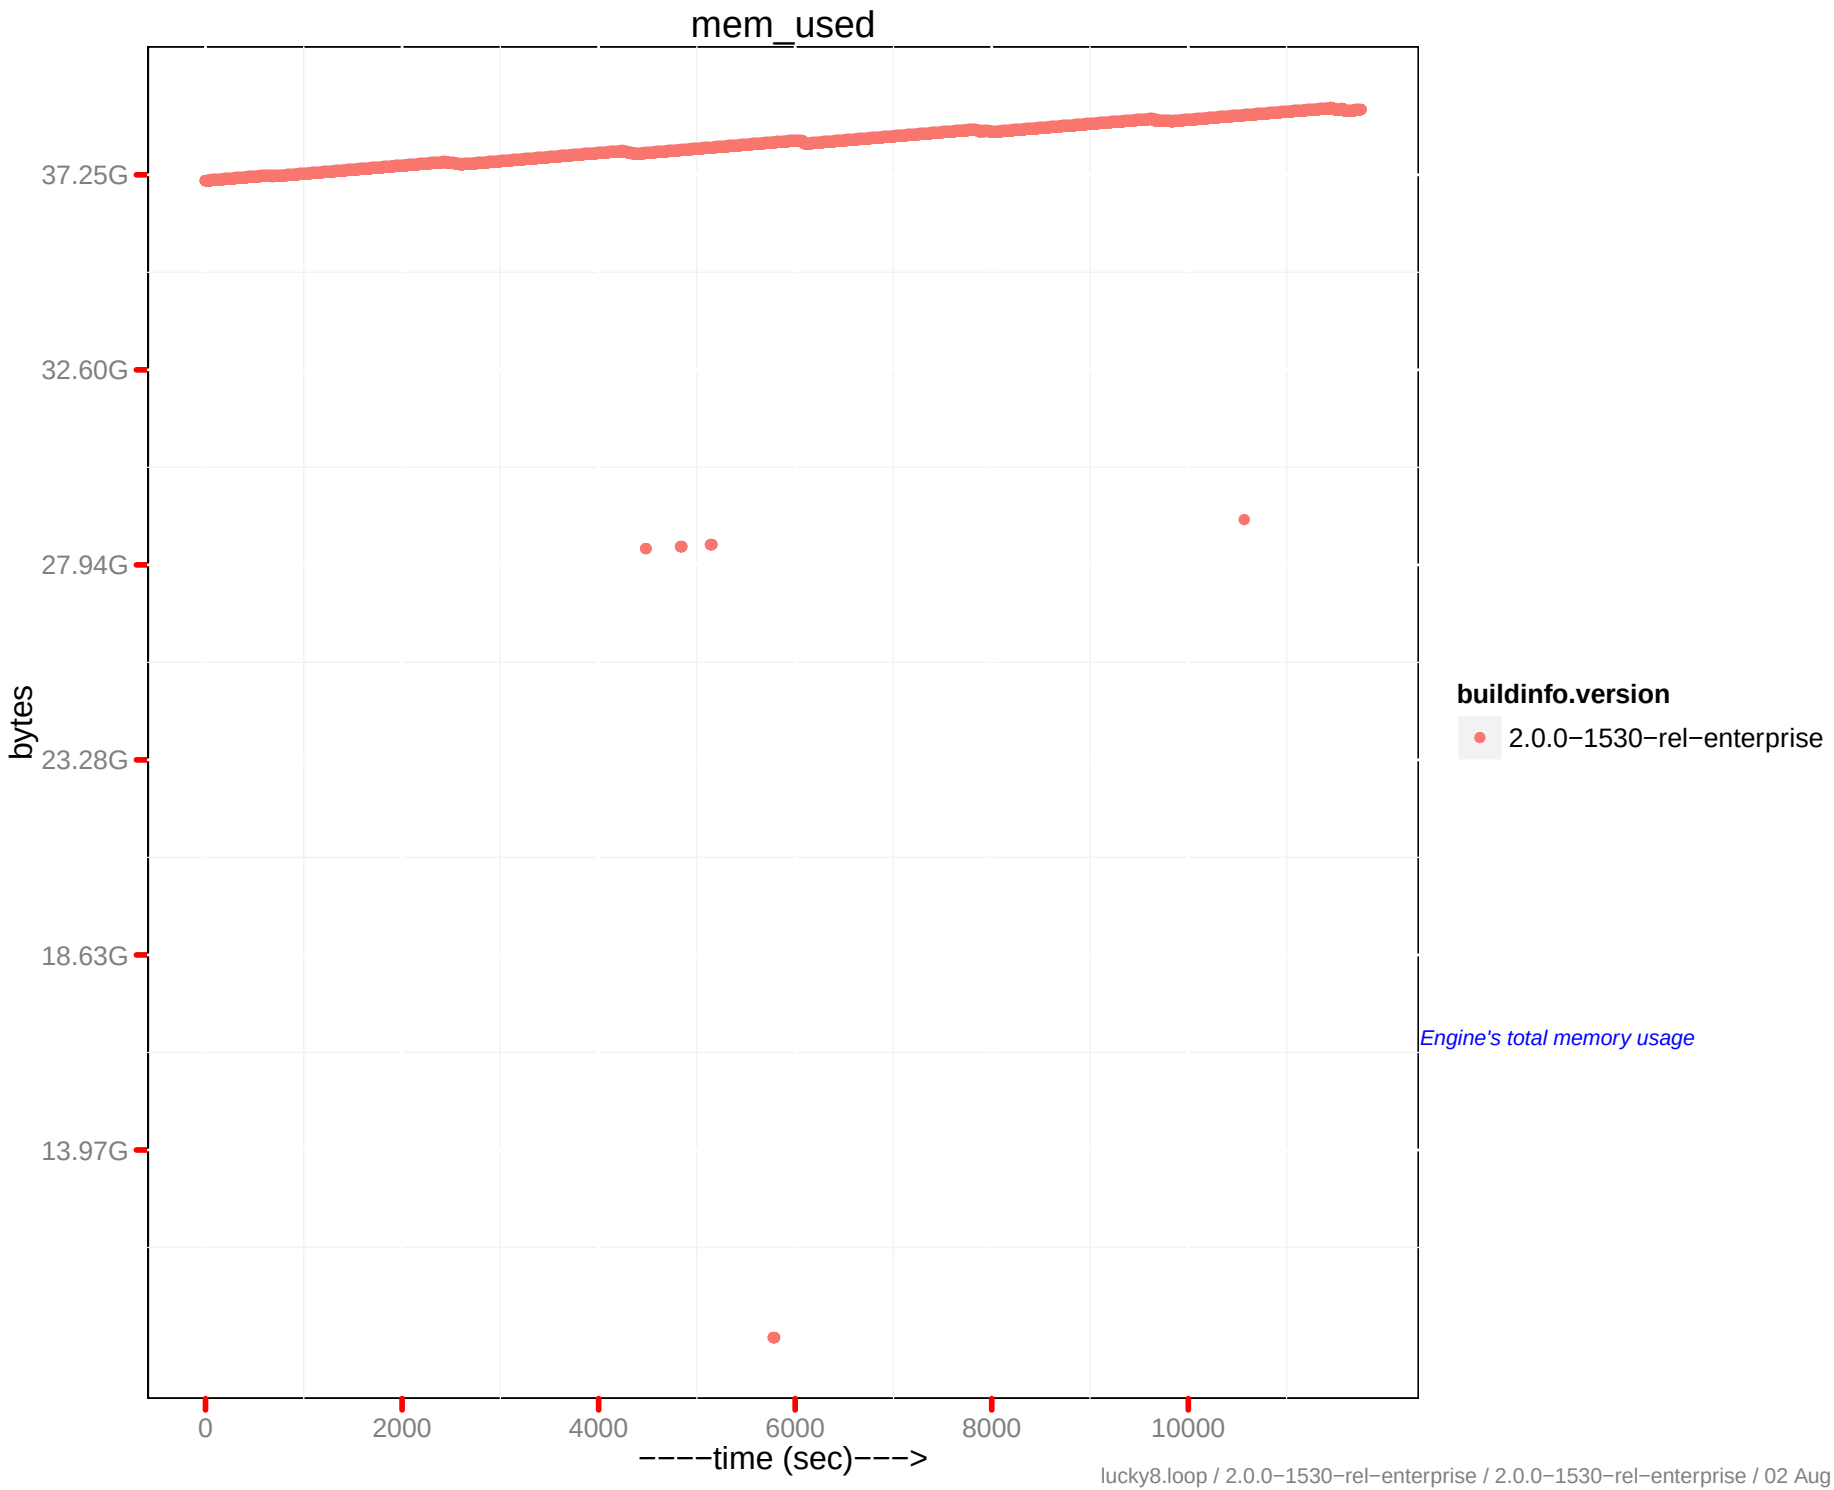
mem_used
37.25G
32.60G
27.94G
23.28G
18.63G
13.97G
bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Engine's total memory usage
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

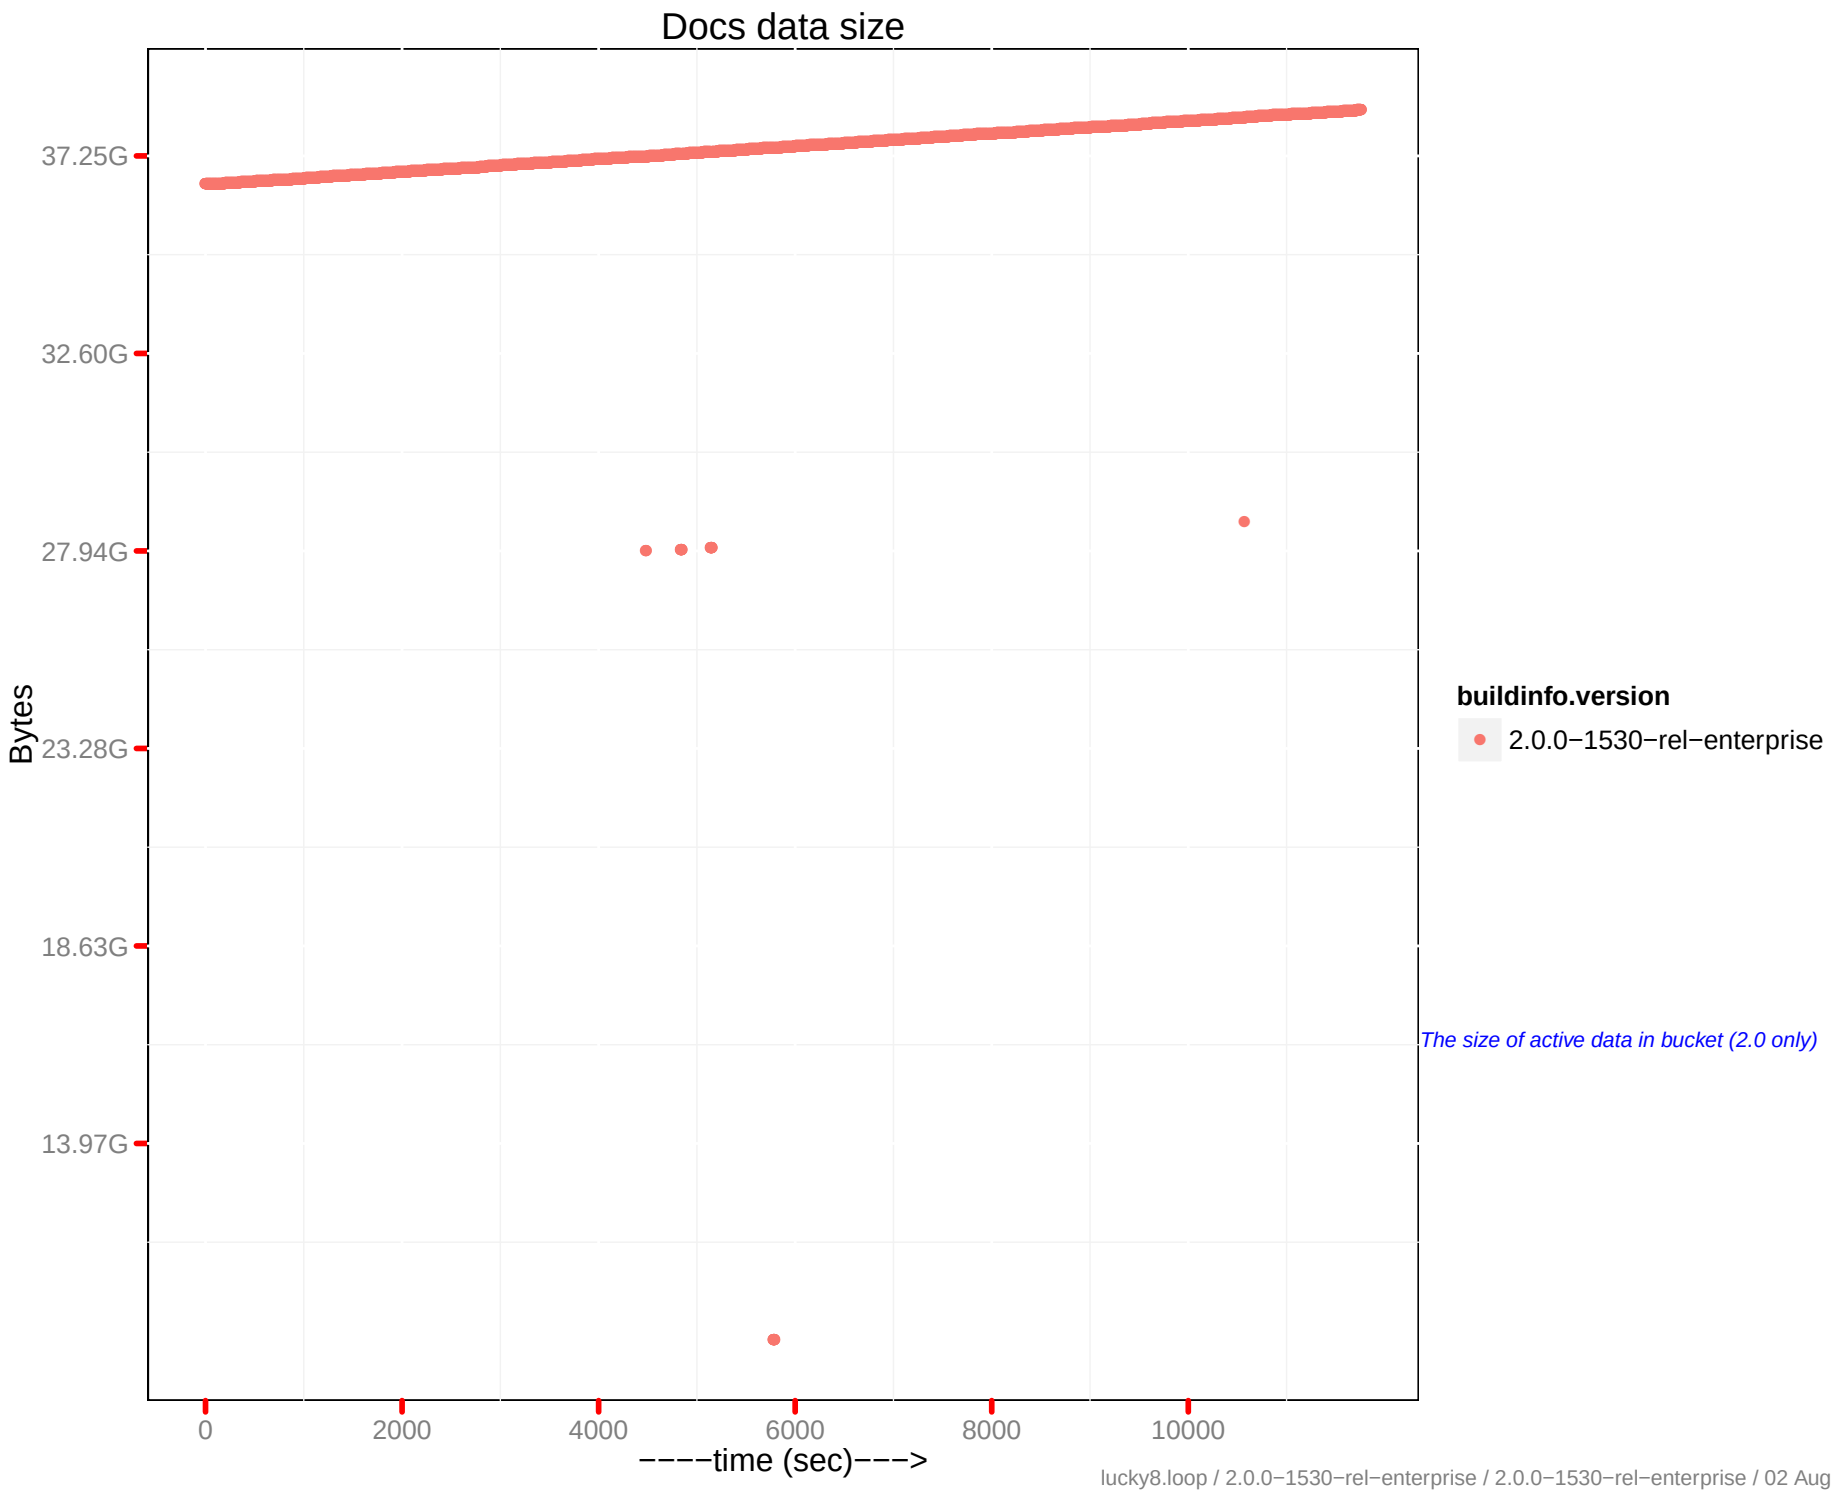
Docs data size
37.25G
32.60G
27.94G
23.28G
18.63G
13.97G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The size of active data in bucket (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

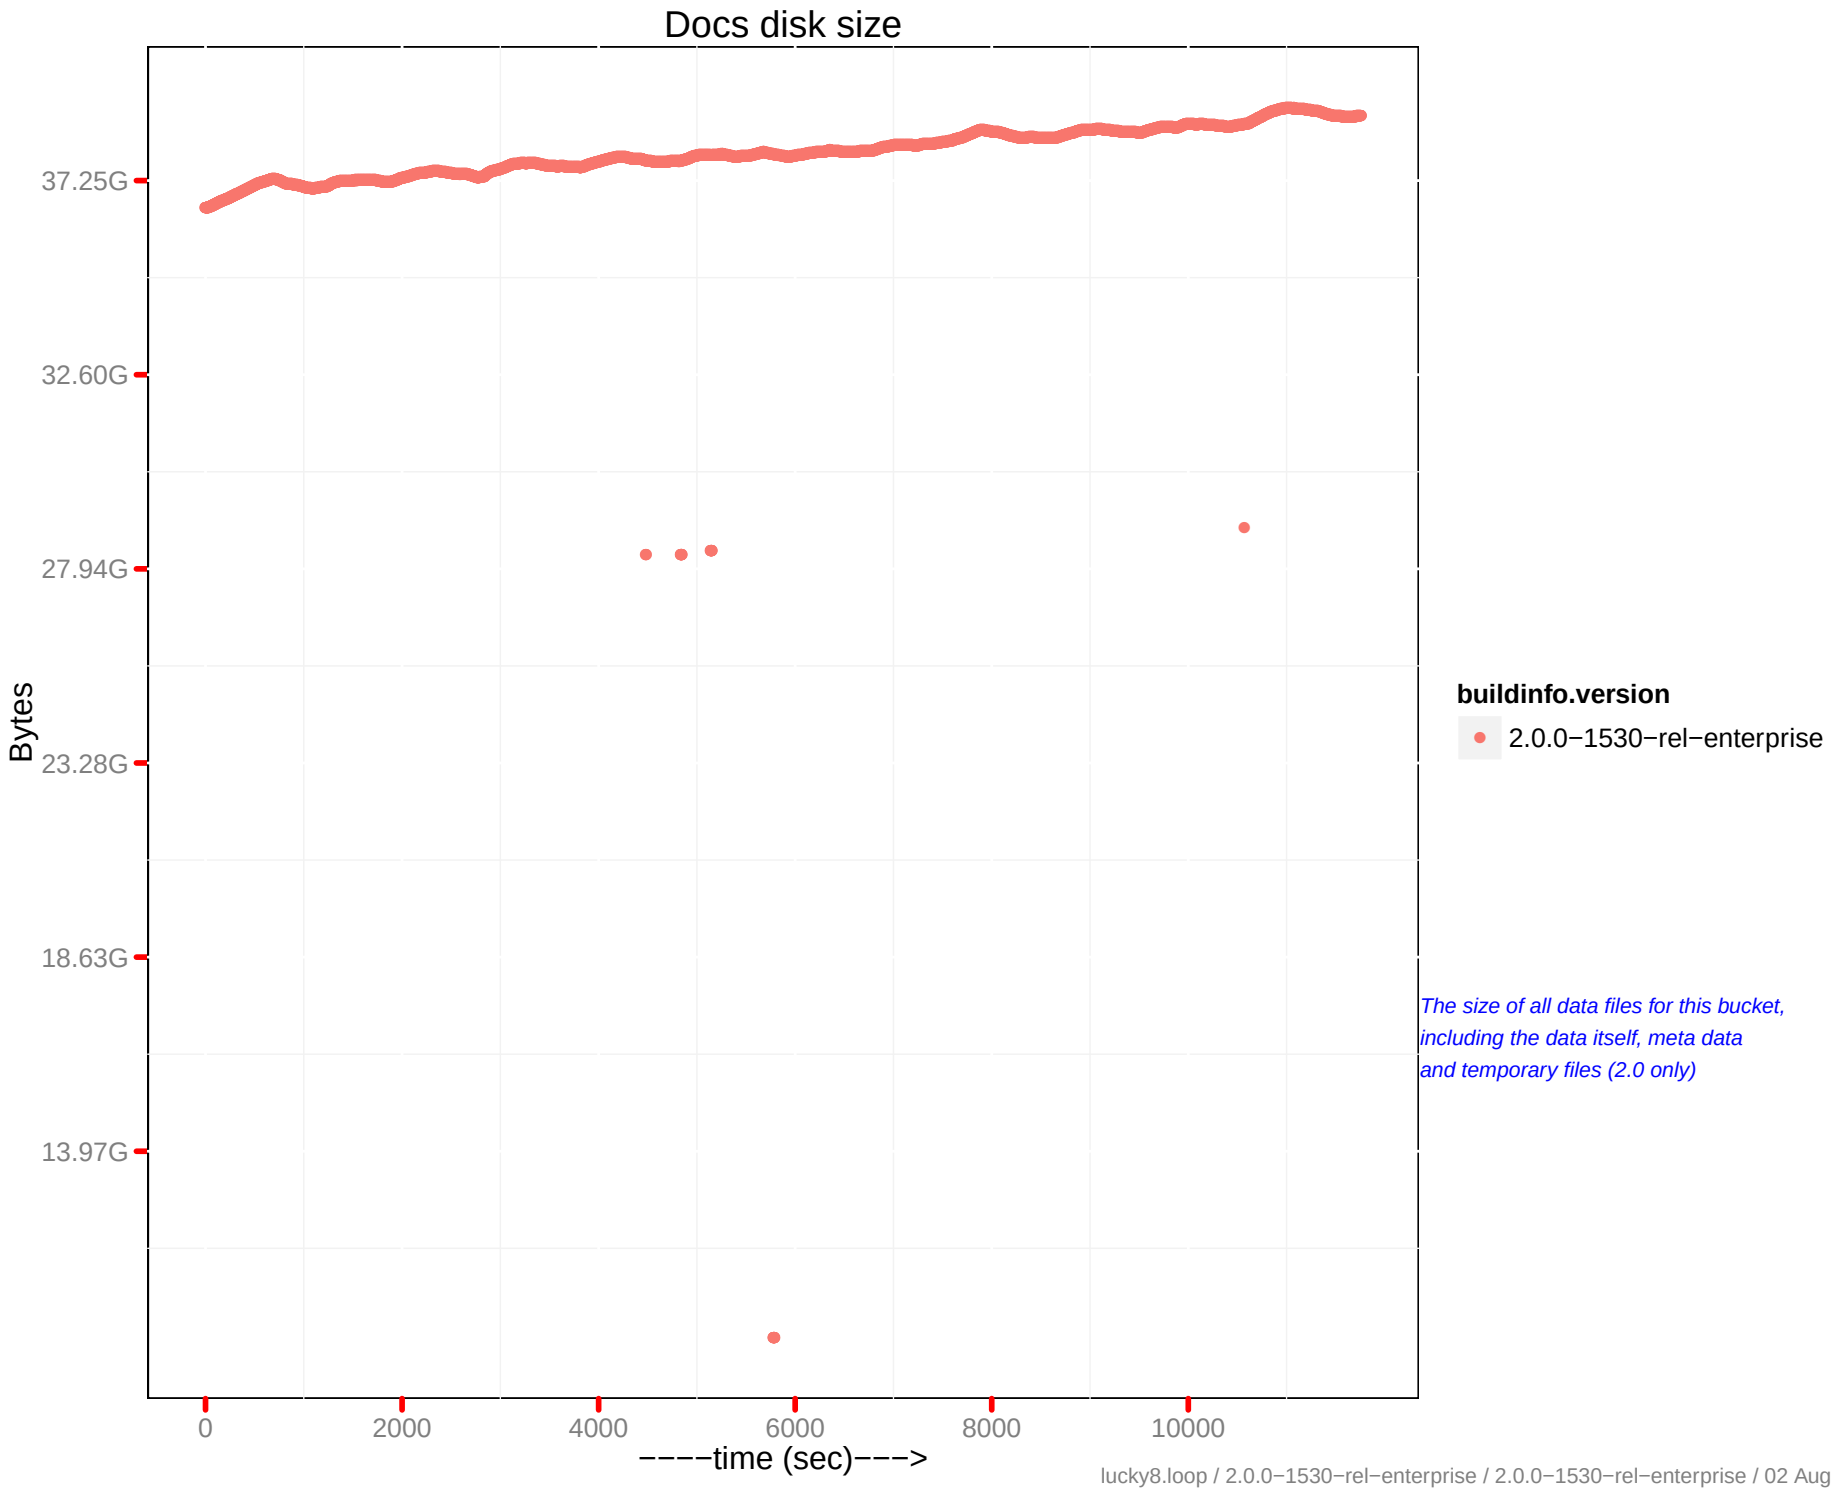
Docs disk size
37.25G
32.60G
27.94G
23.28G
18.63G
13.97G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

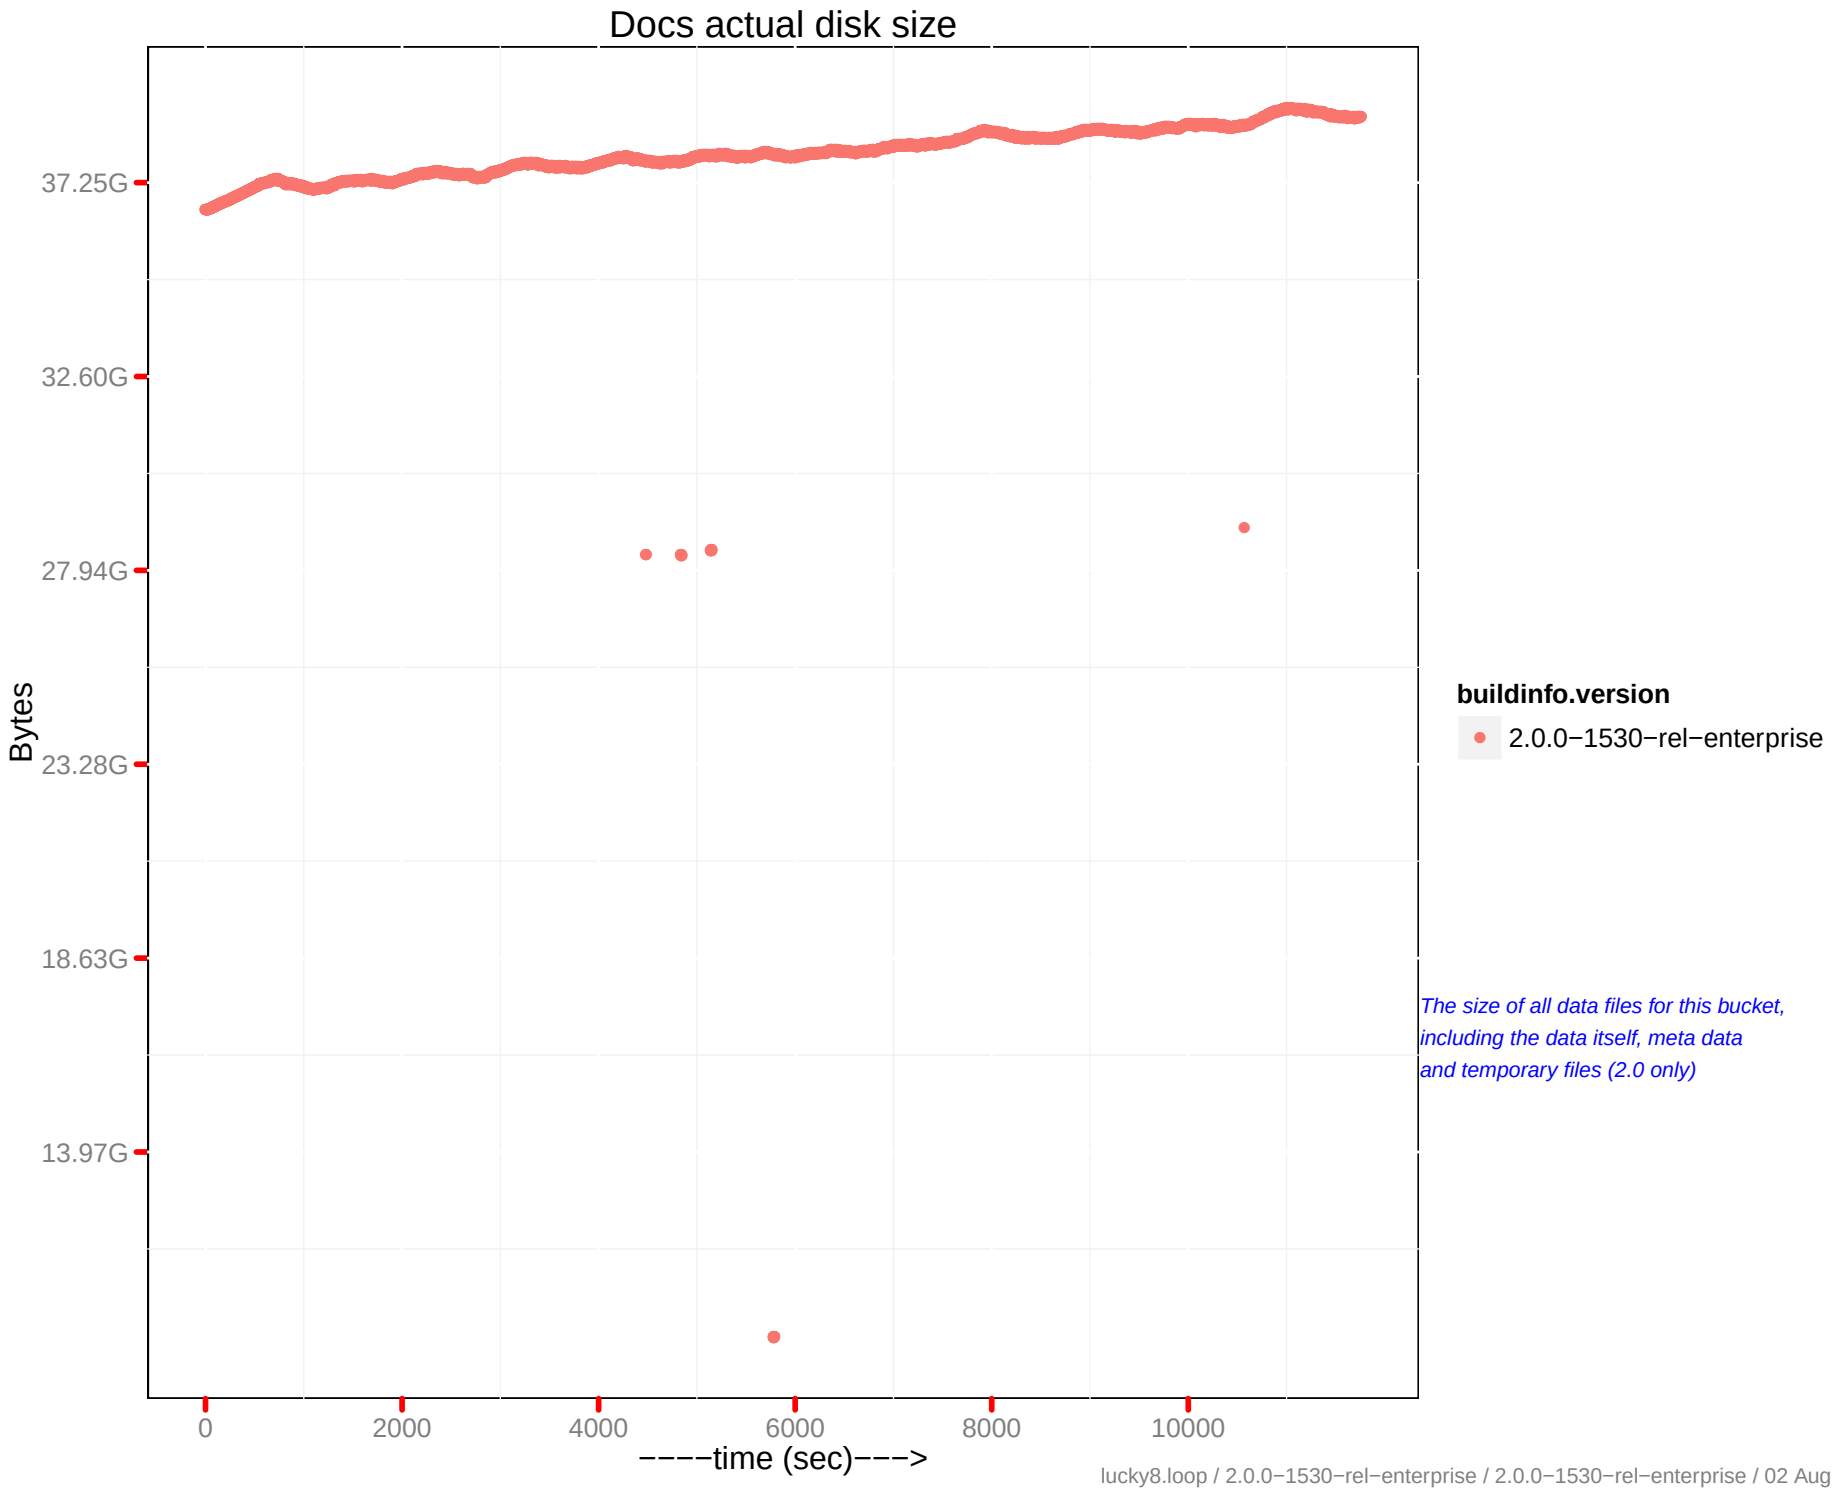
Docs actual disk size
37.25G
32.60G
27.94G
23.28G
18.63G
13.97G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

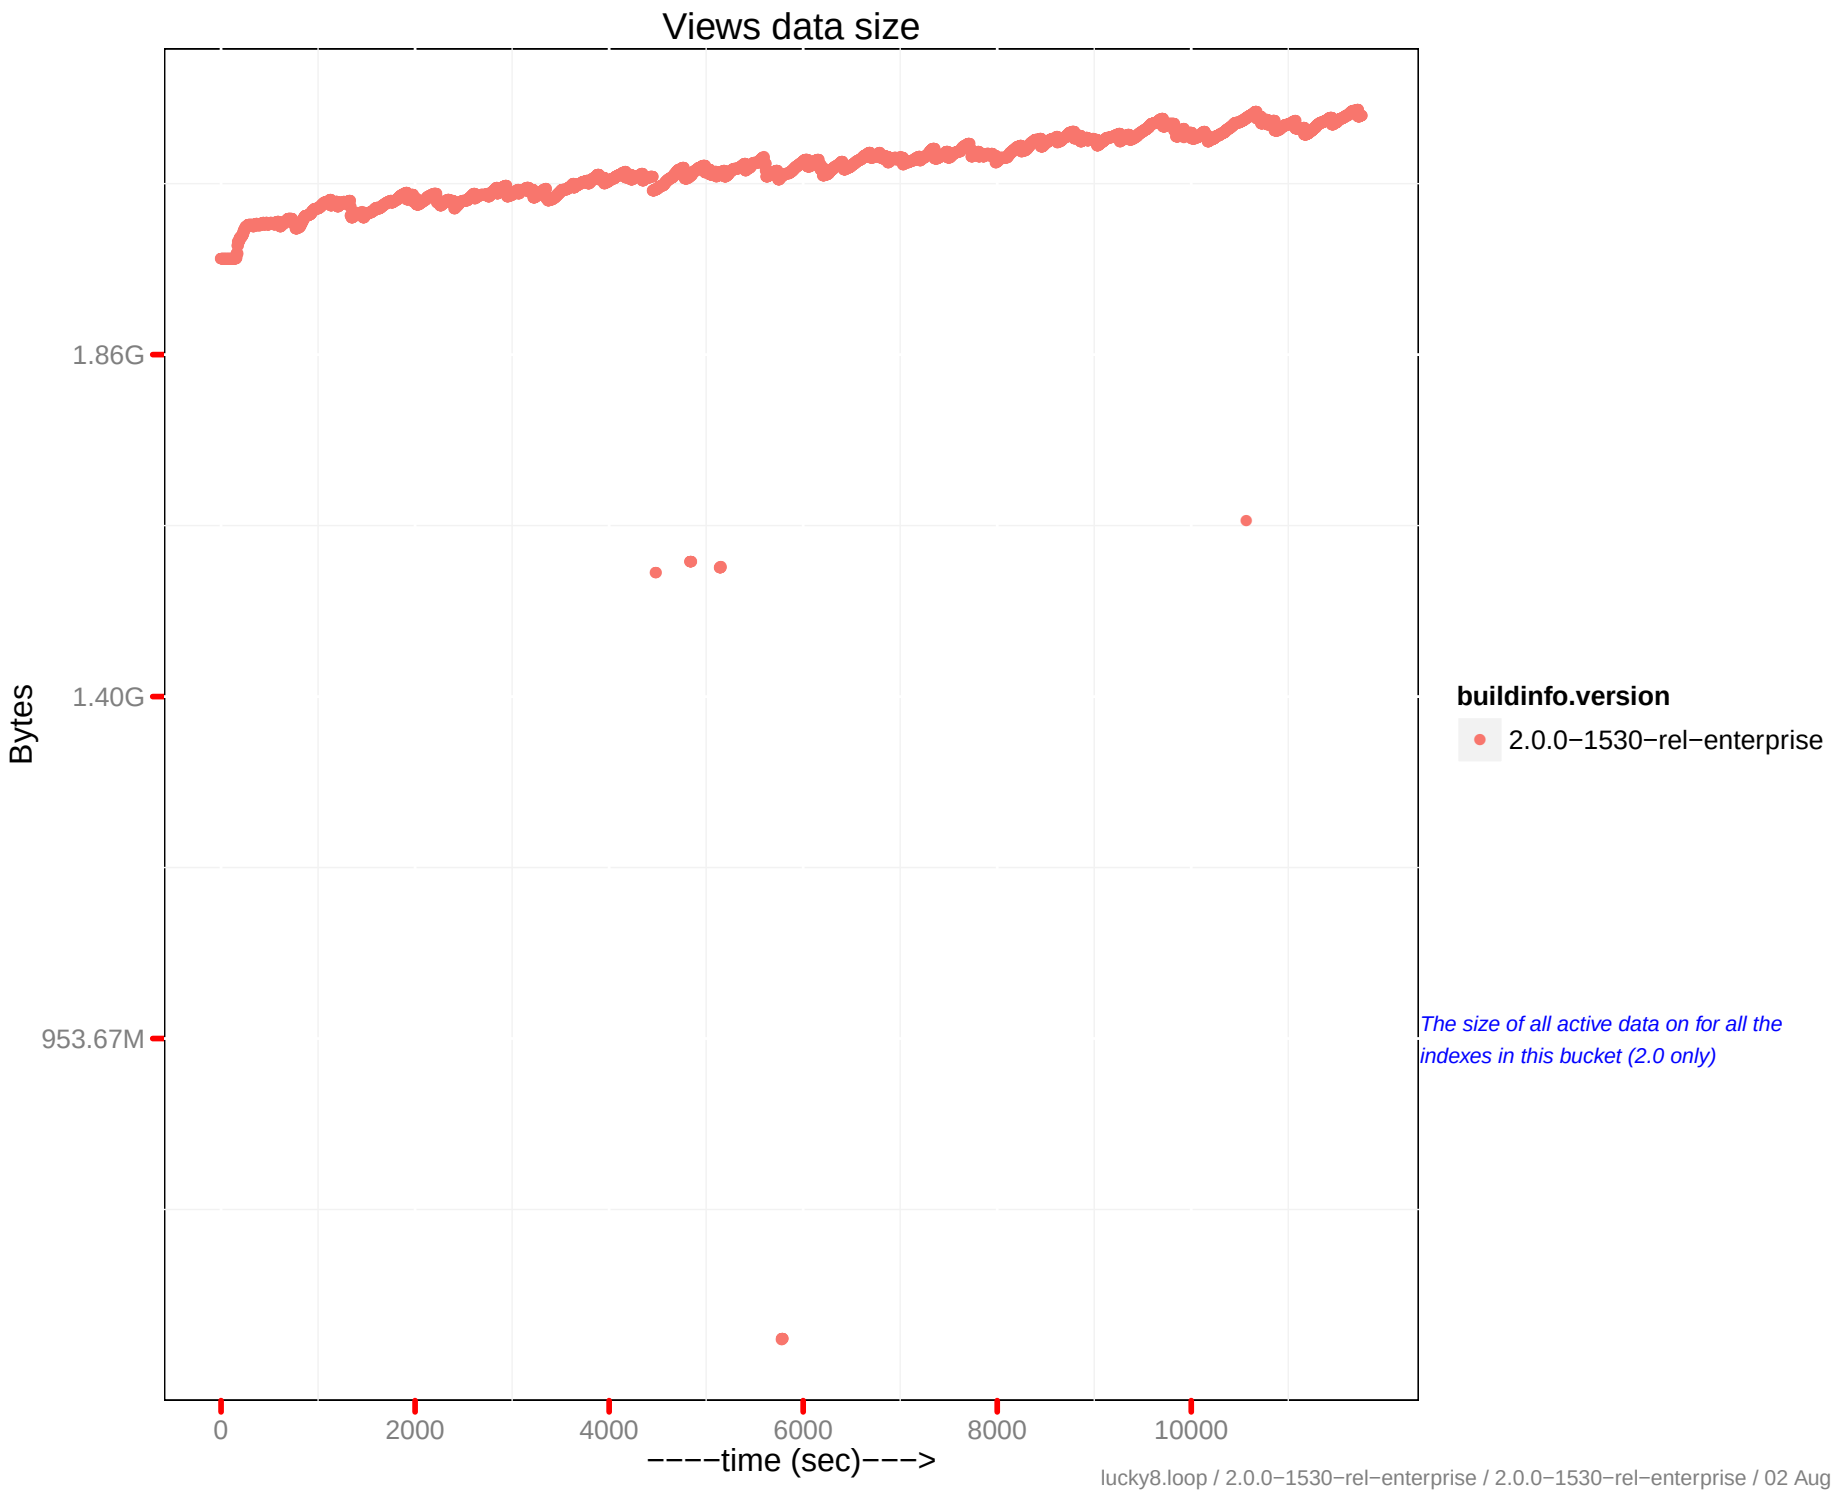
Views data size
1.86G
1.40G
953.67M
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

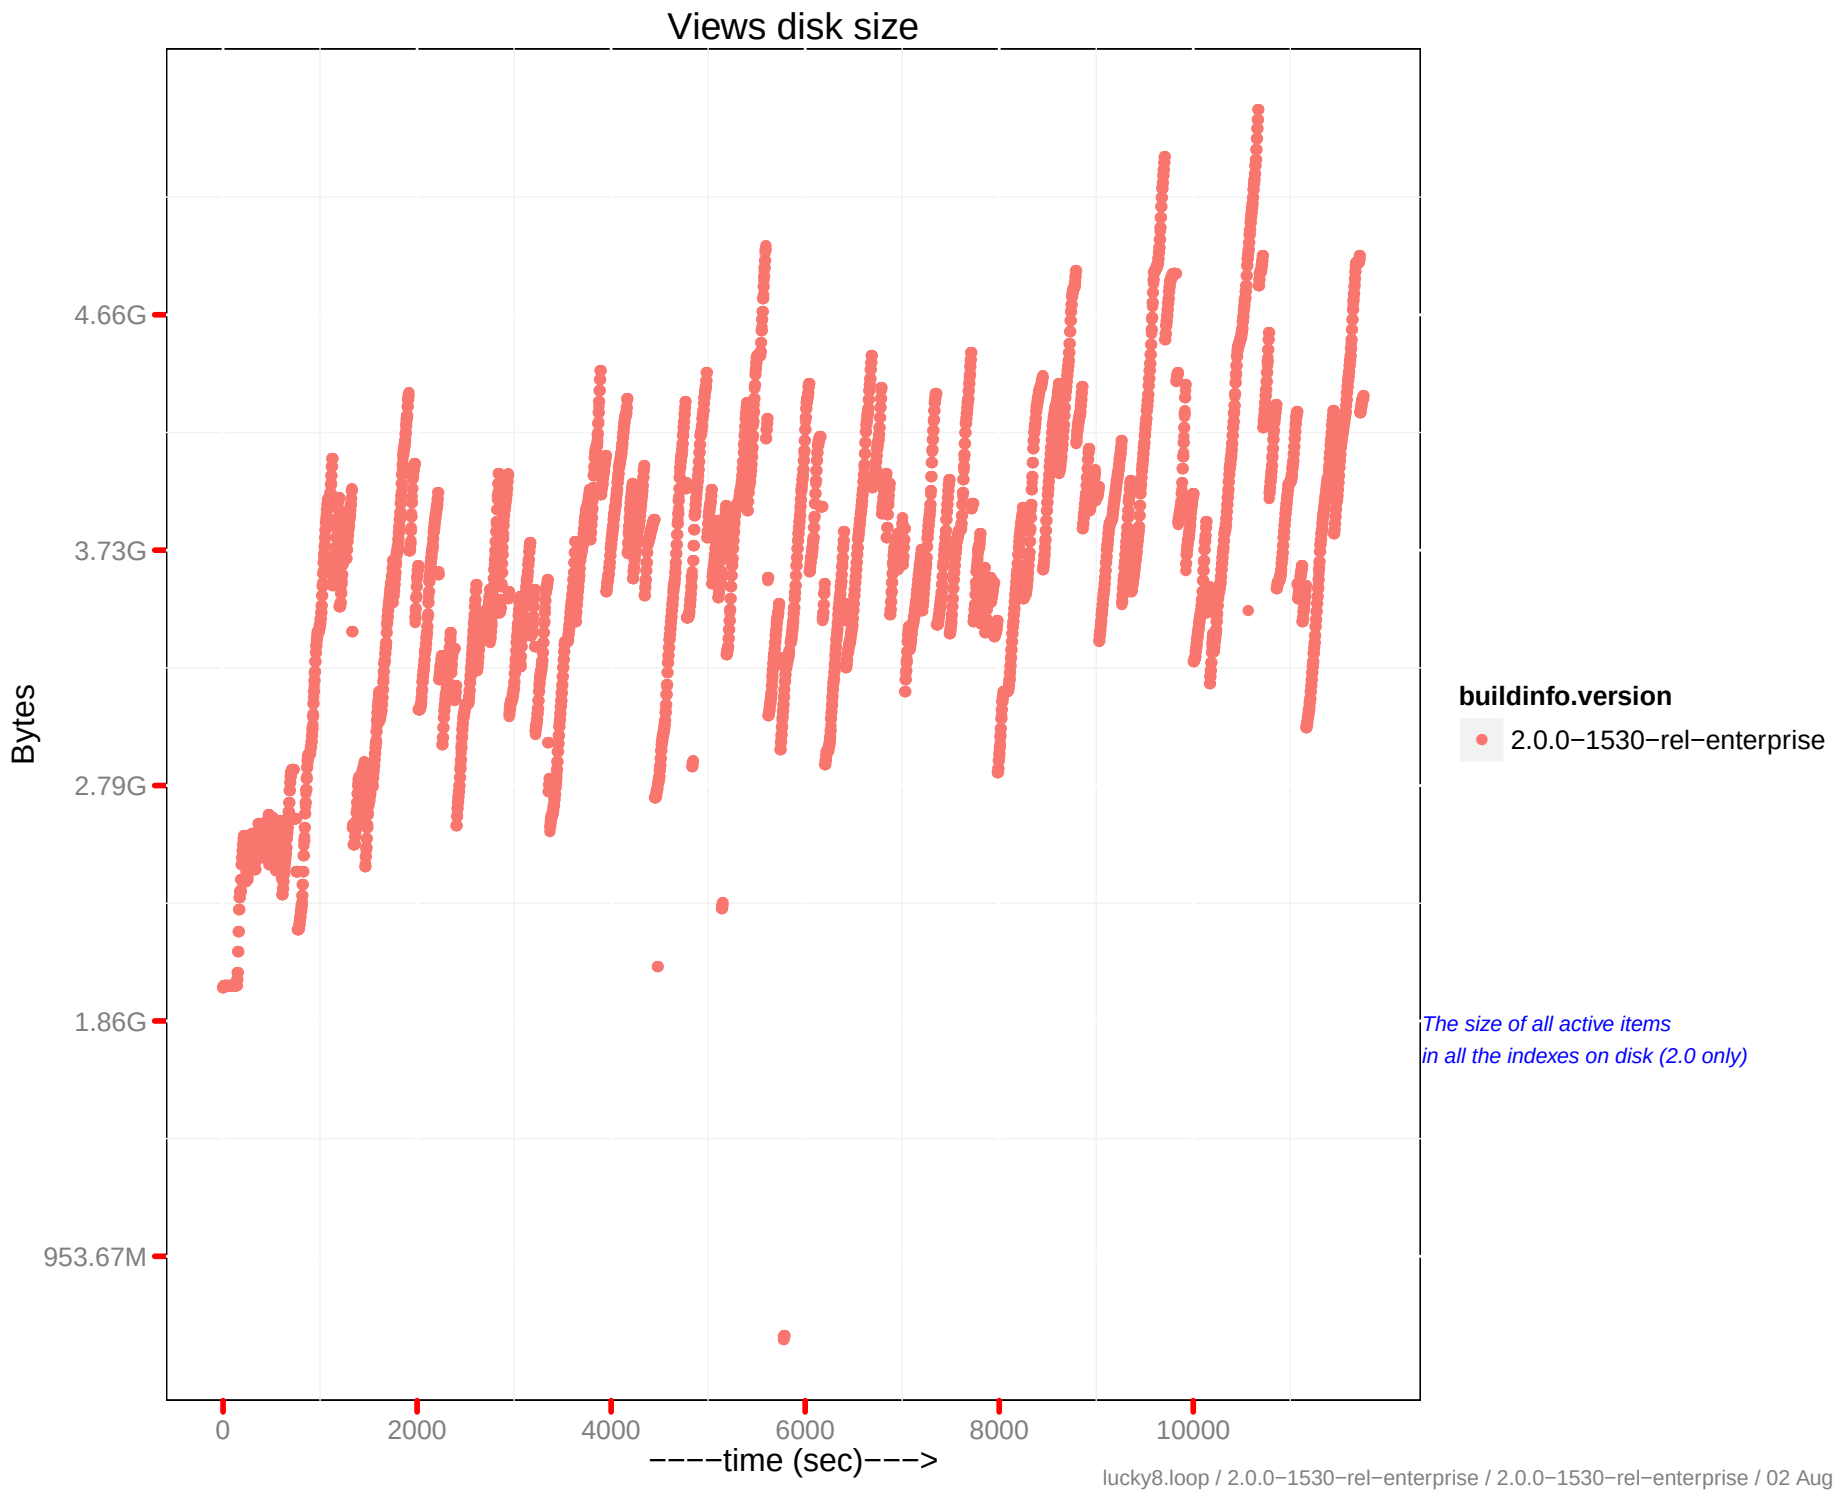
Views disk size
Bytes
4.66G
3.73G
2.79G
1.86G
953.67M
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

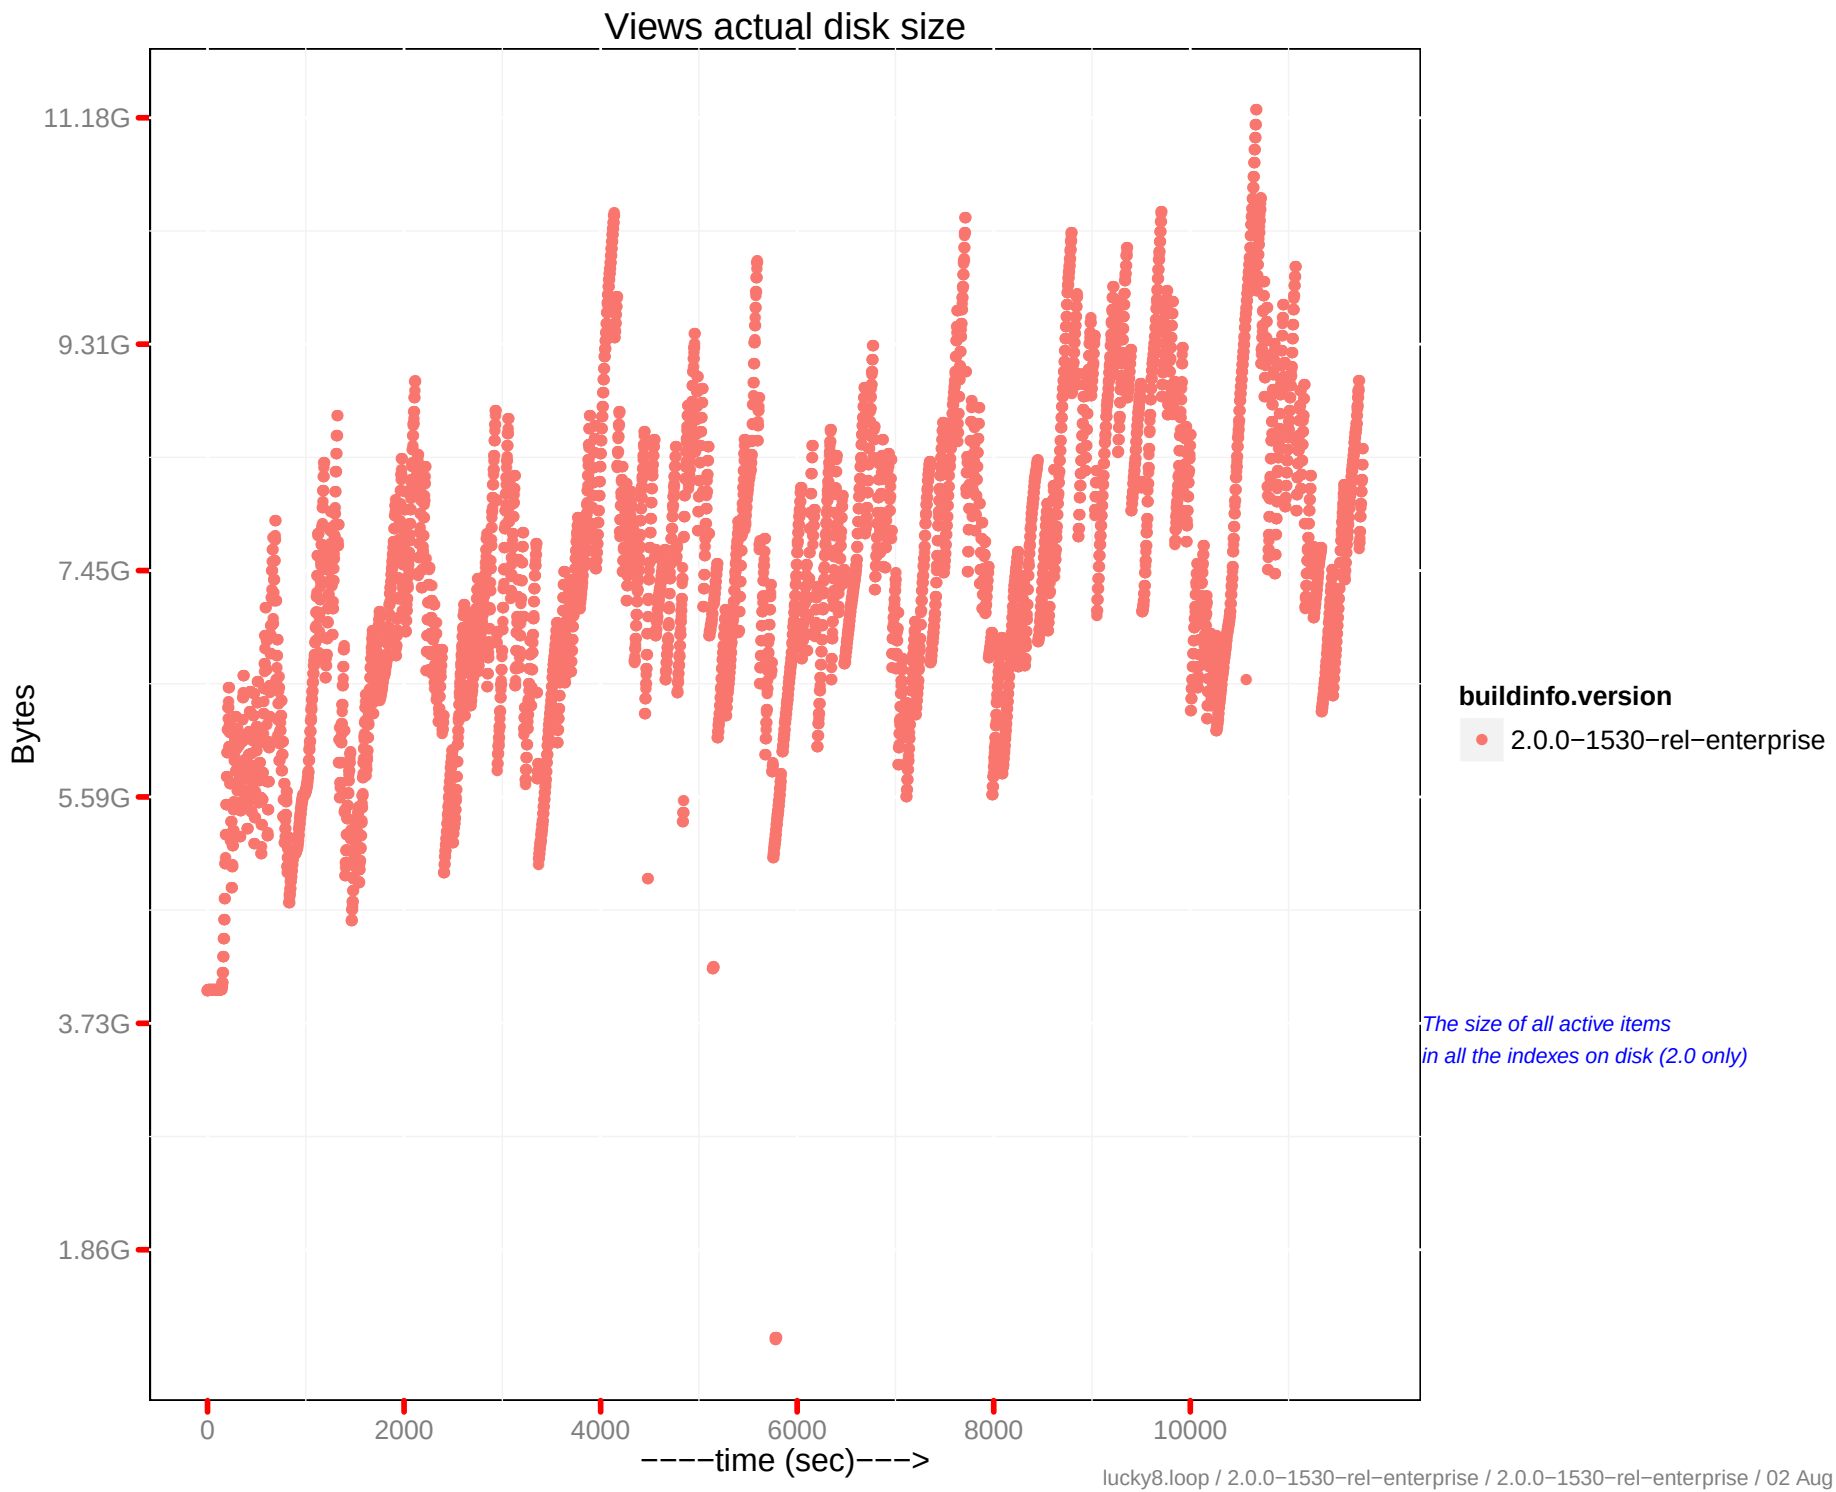
Views actual disk size
11.18G
9.31G
7.45G
5.59G
3.73G
1.86G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

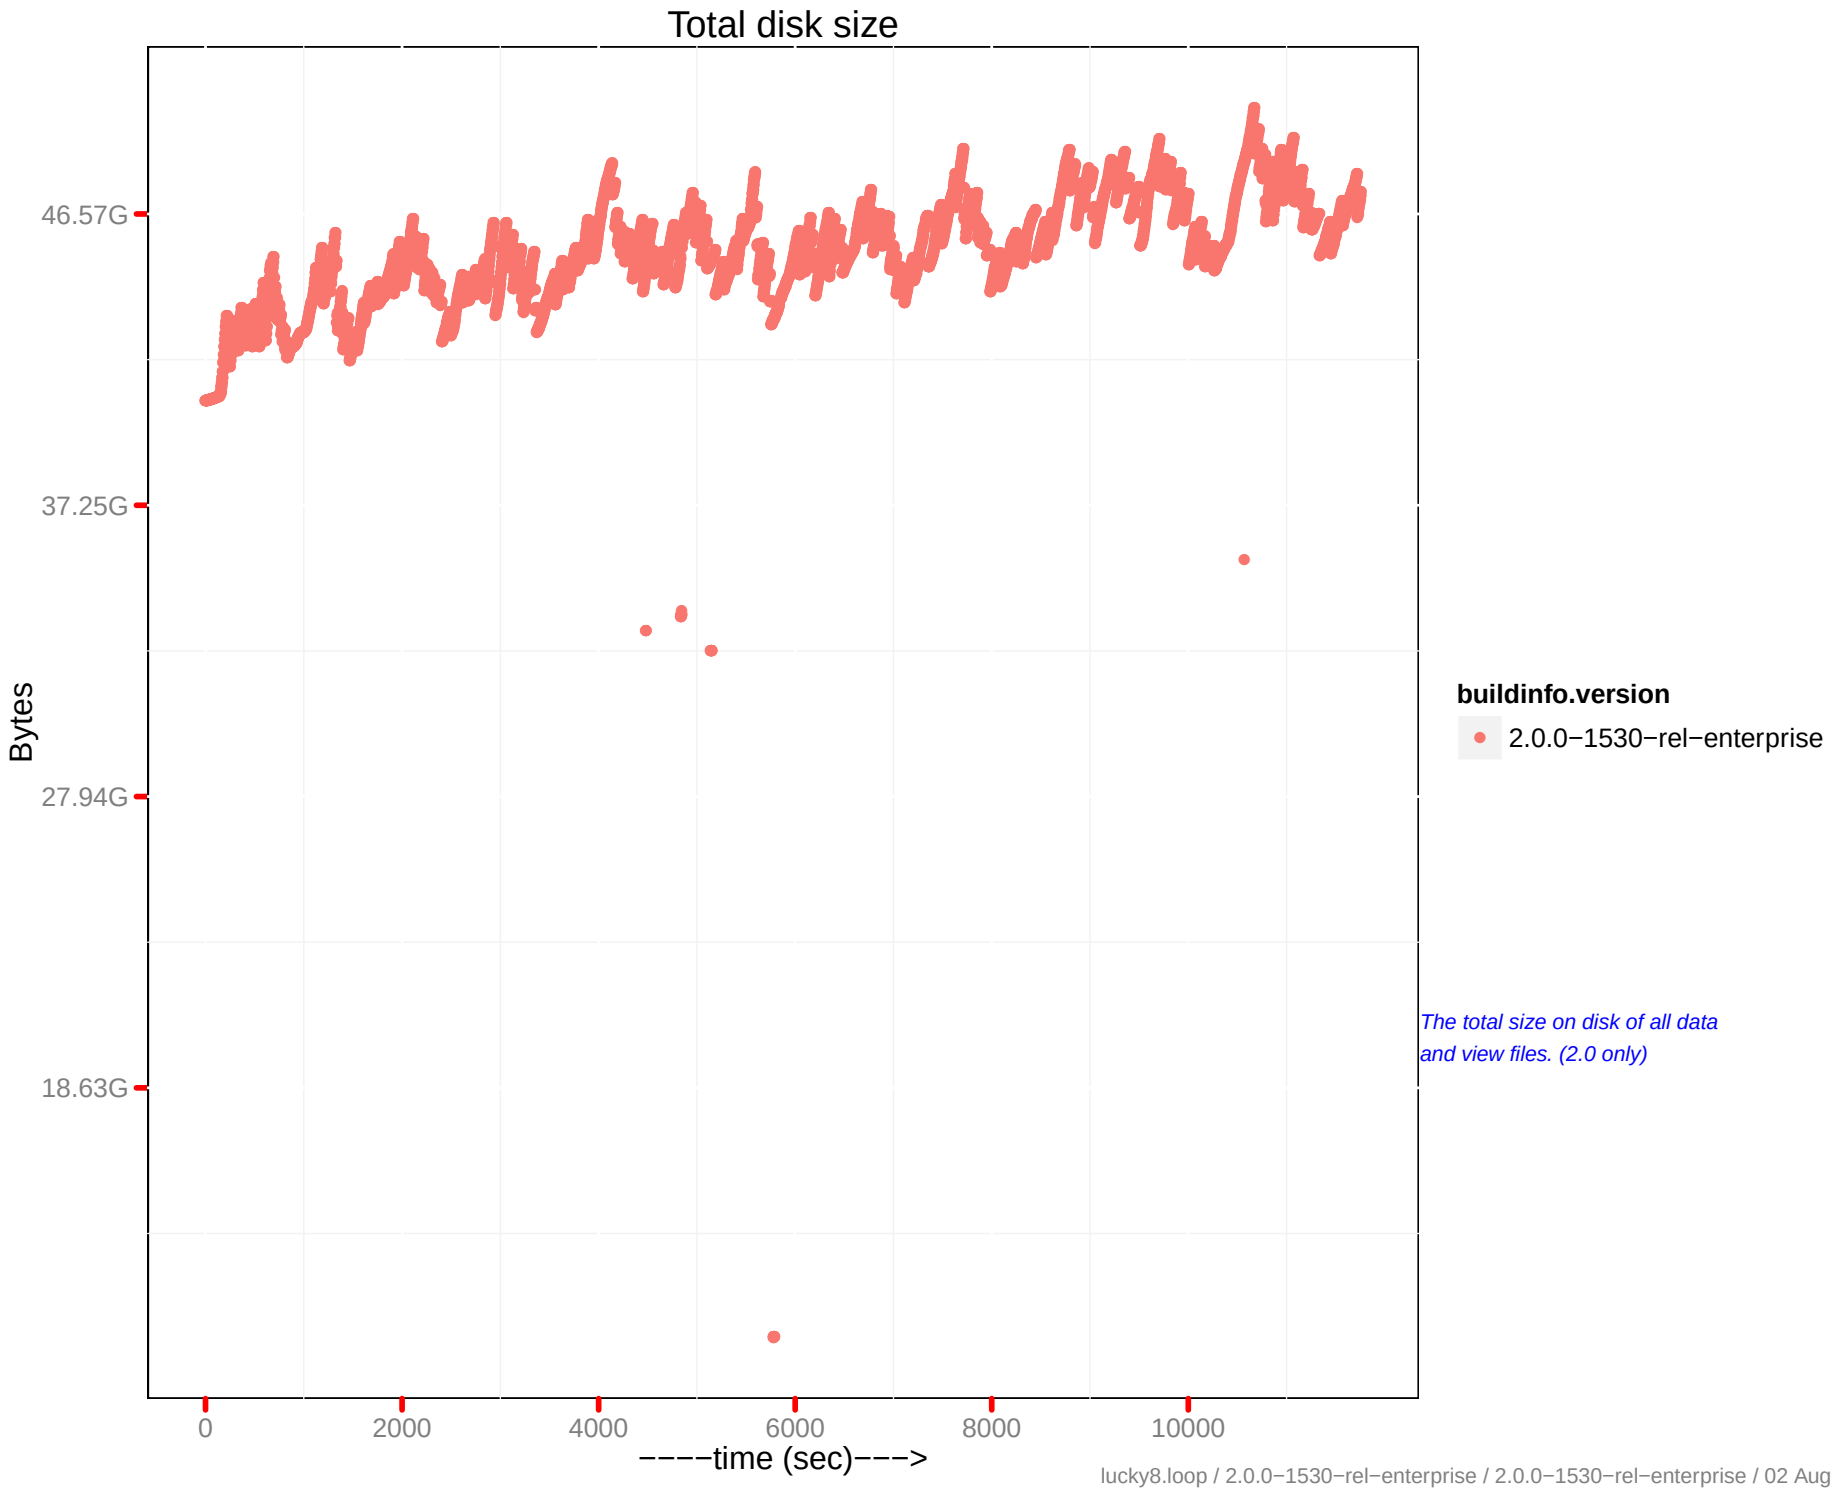
Total disk size
46.57G
37.25G
27.94G
18.63G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

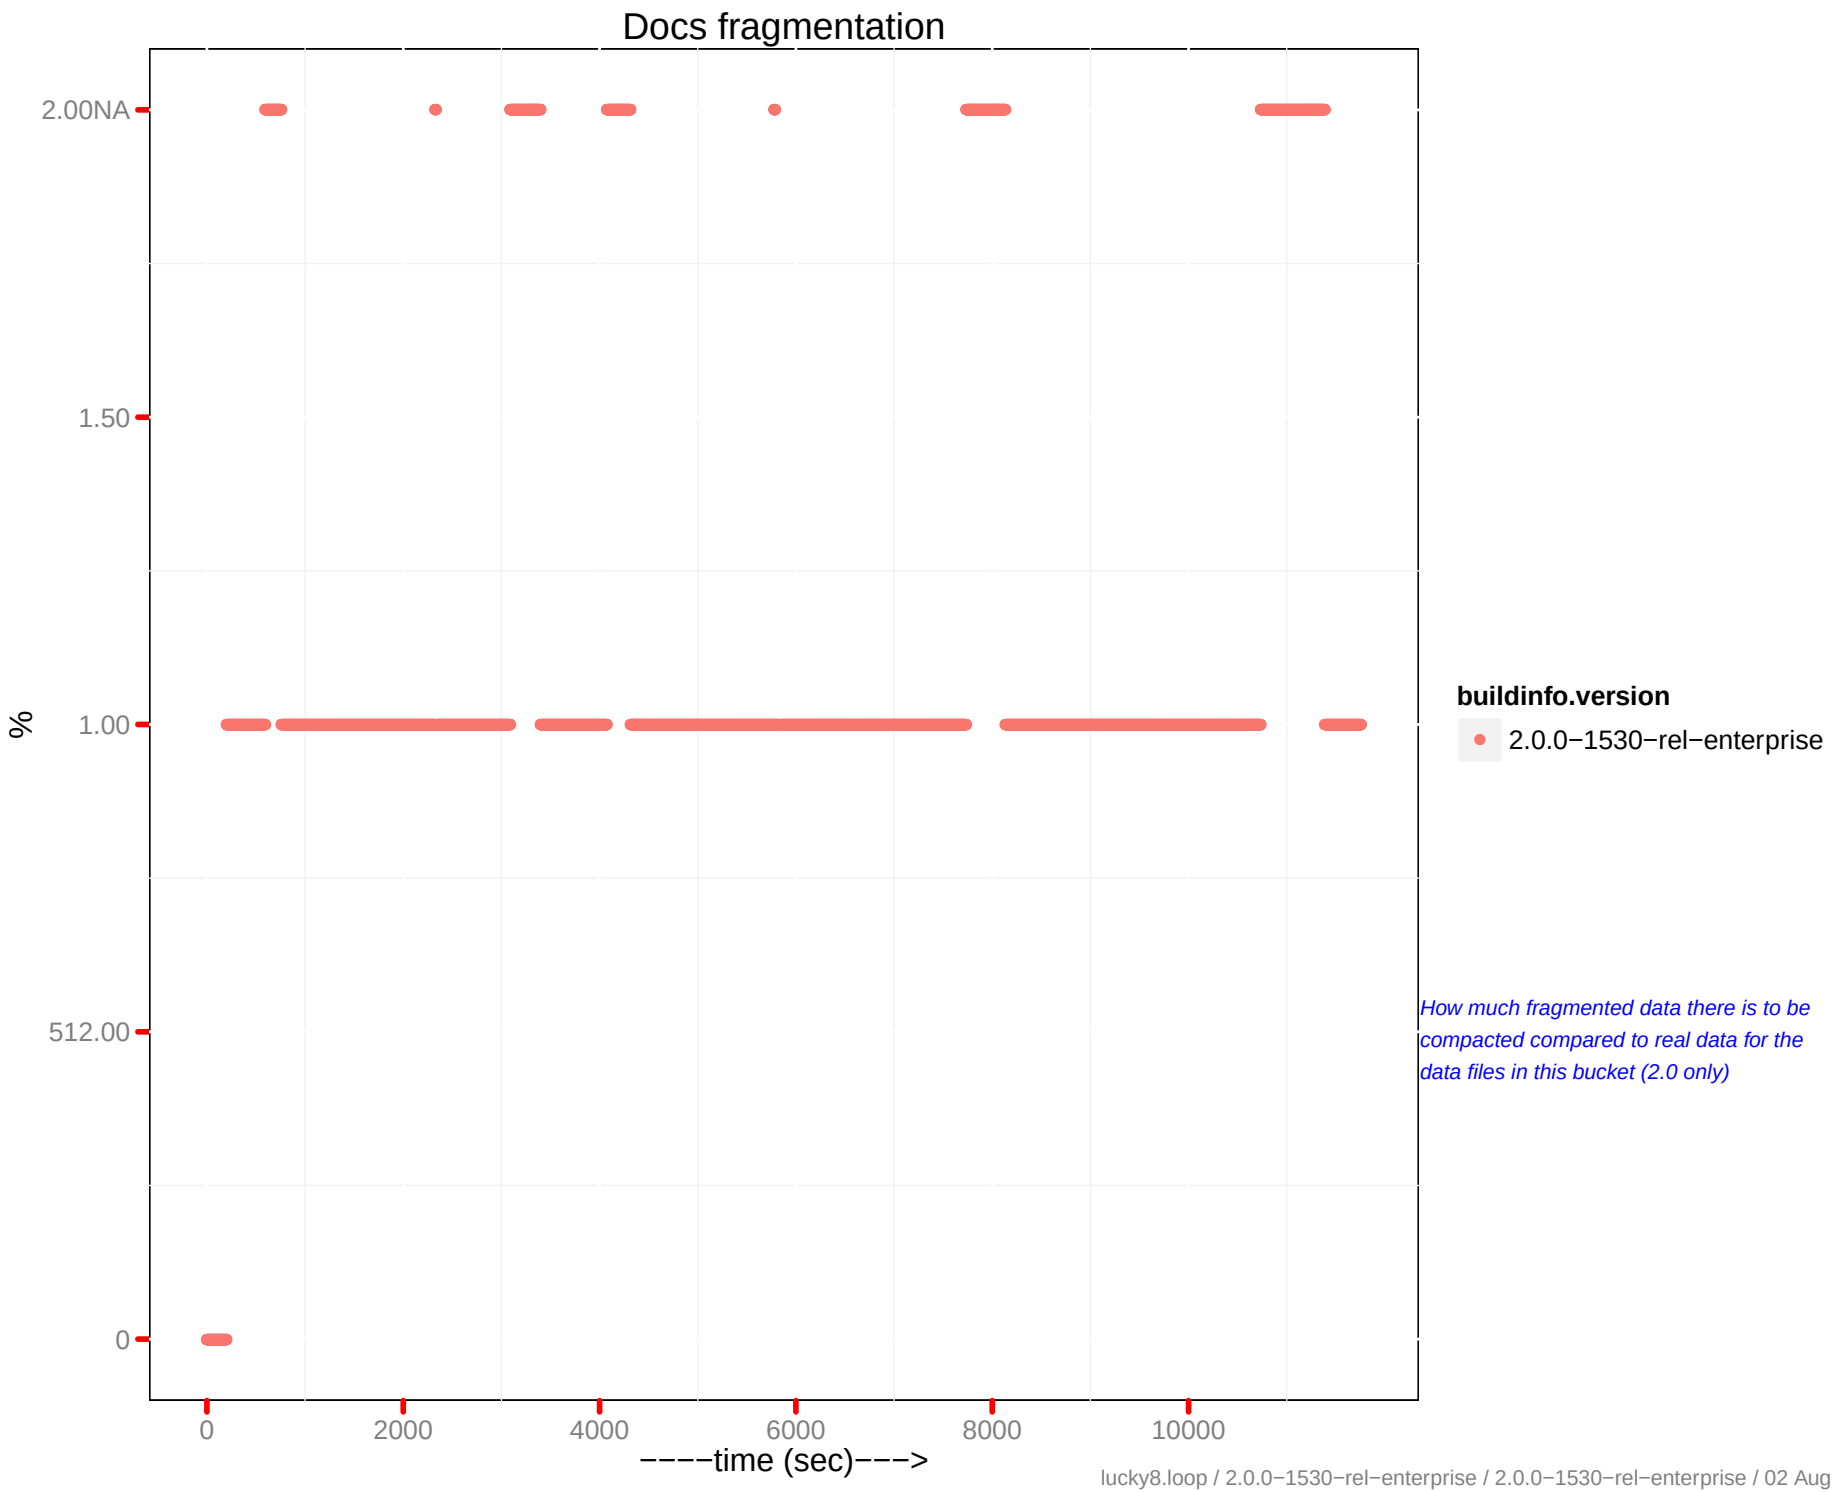
Docs fragmentation
2.00NA
1.50
1.00
512.00
0
%
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

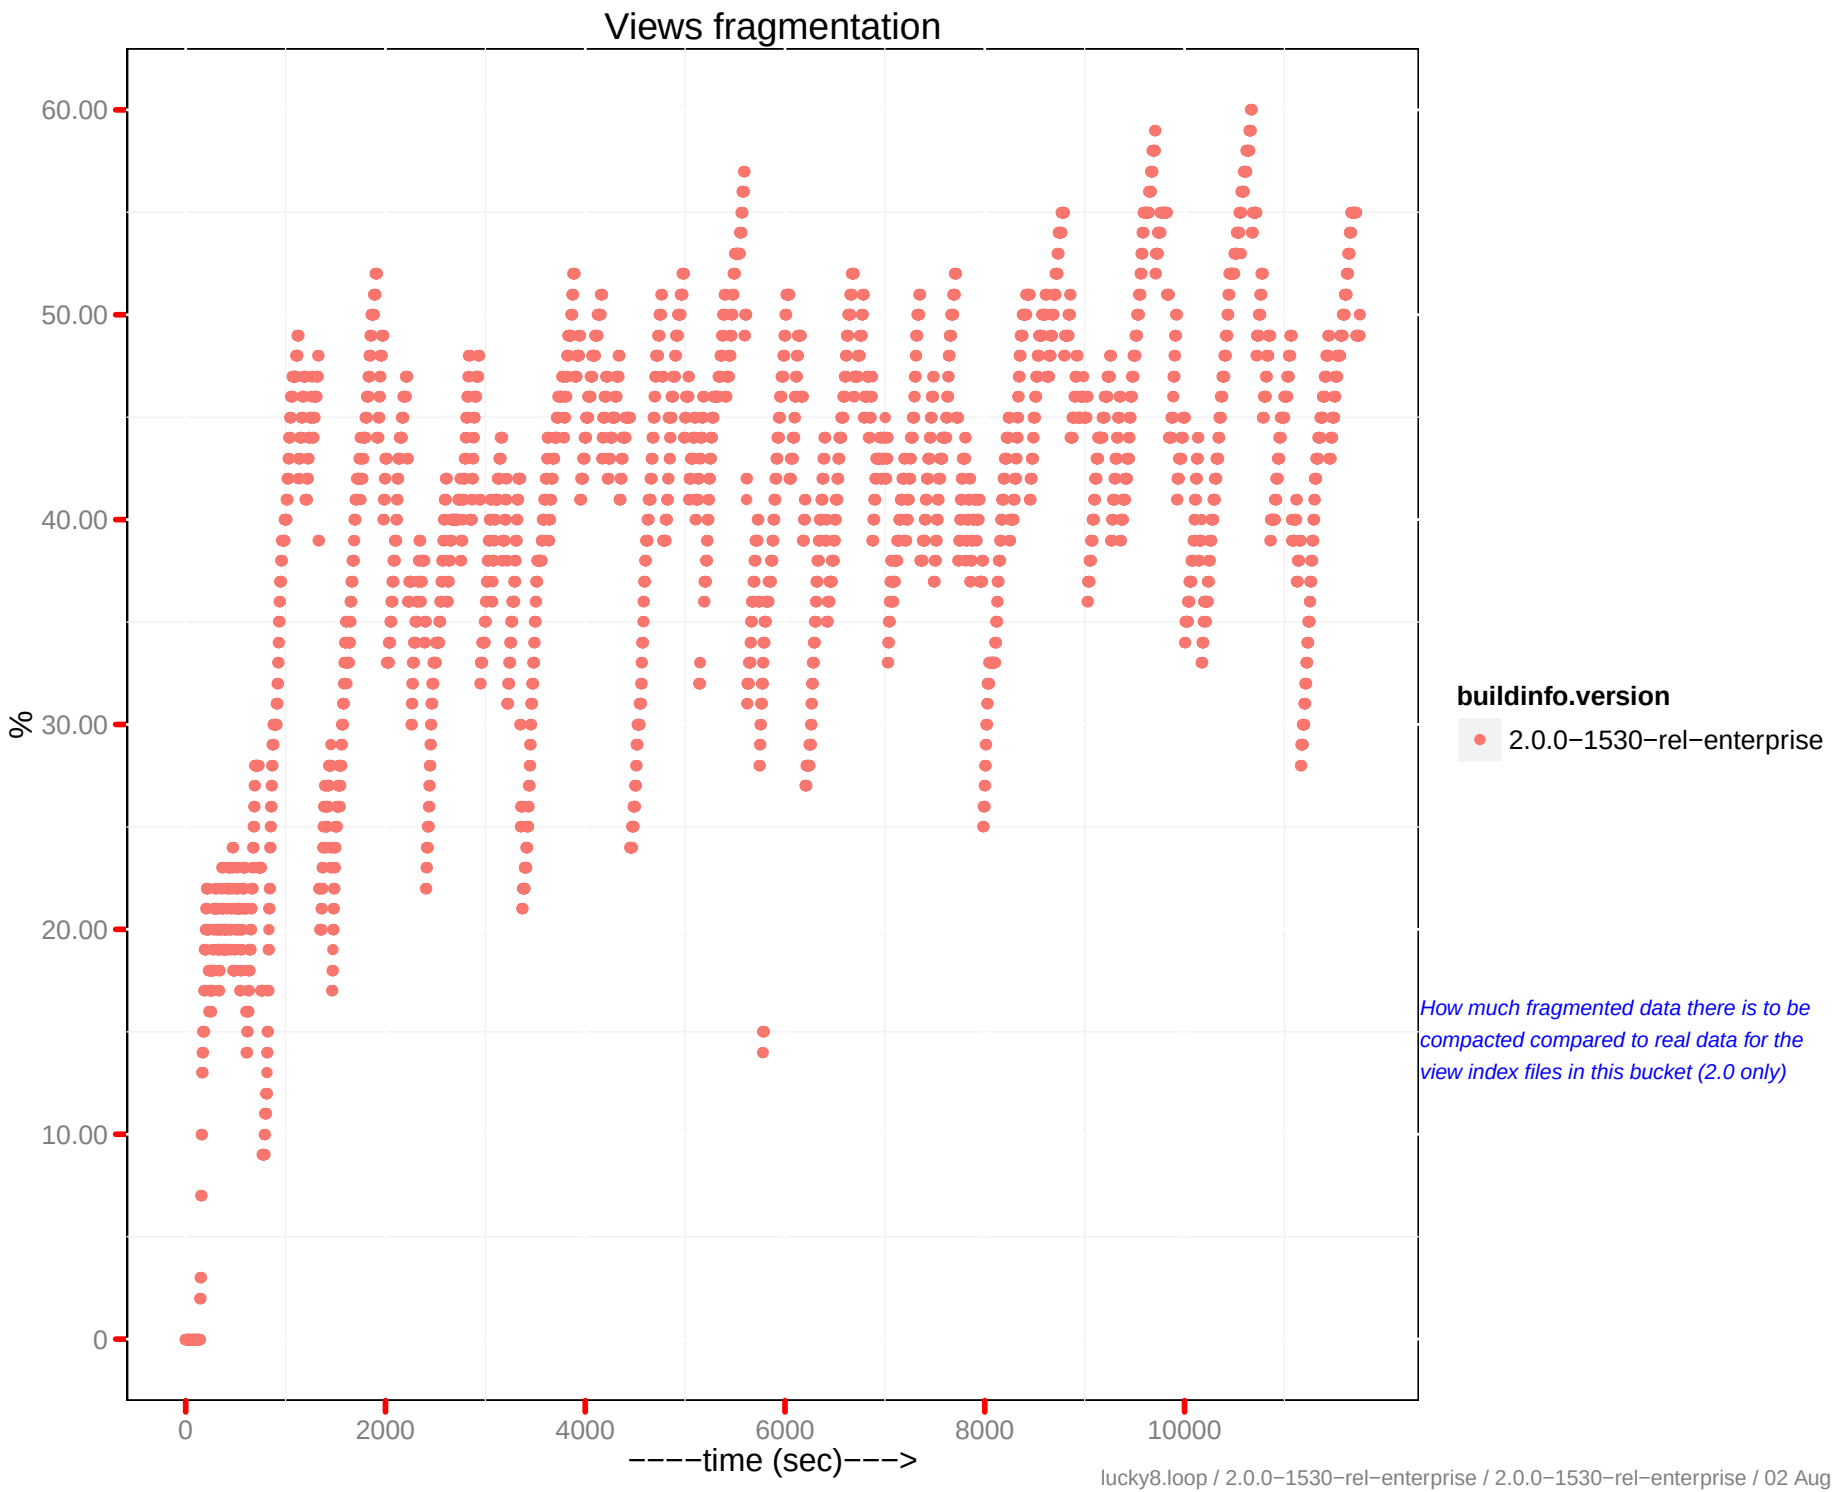
Views fragmentation
60.00
50.00
40.00
30.00
20.00
10.00
0
%
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012


cmd_get ops/sec
ops/sec
5.86K
3.91K
1.95K
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Cumulative number of retrieval reqs
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012


cmd_set ops/sec
1.17K
1000.00
800.00
600.00
400.00
200.00
0
ops/sec
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Cumulative number of set reqs
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

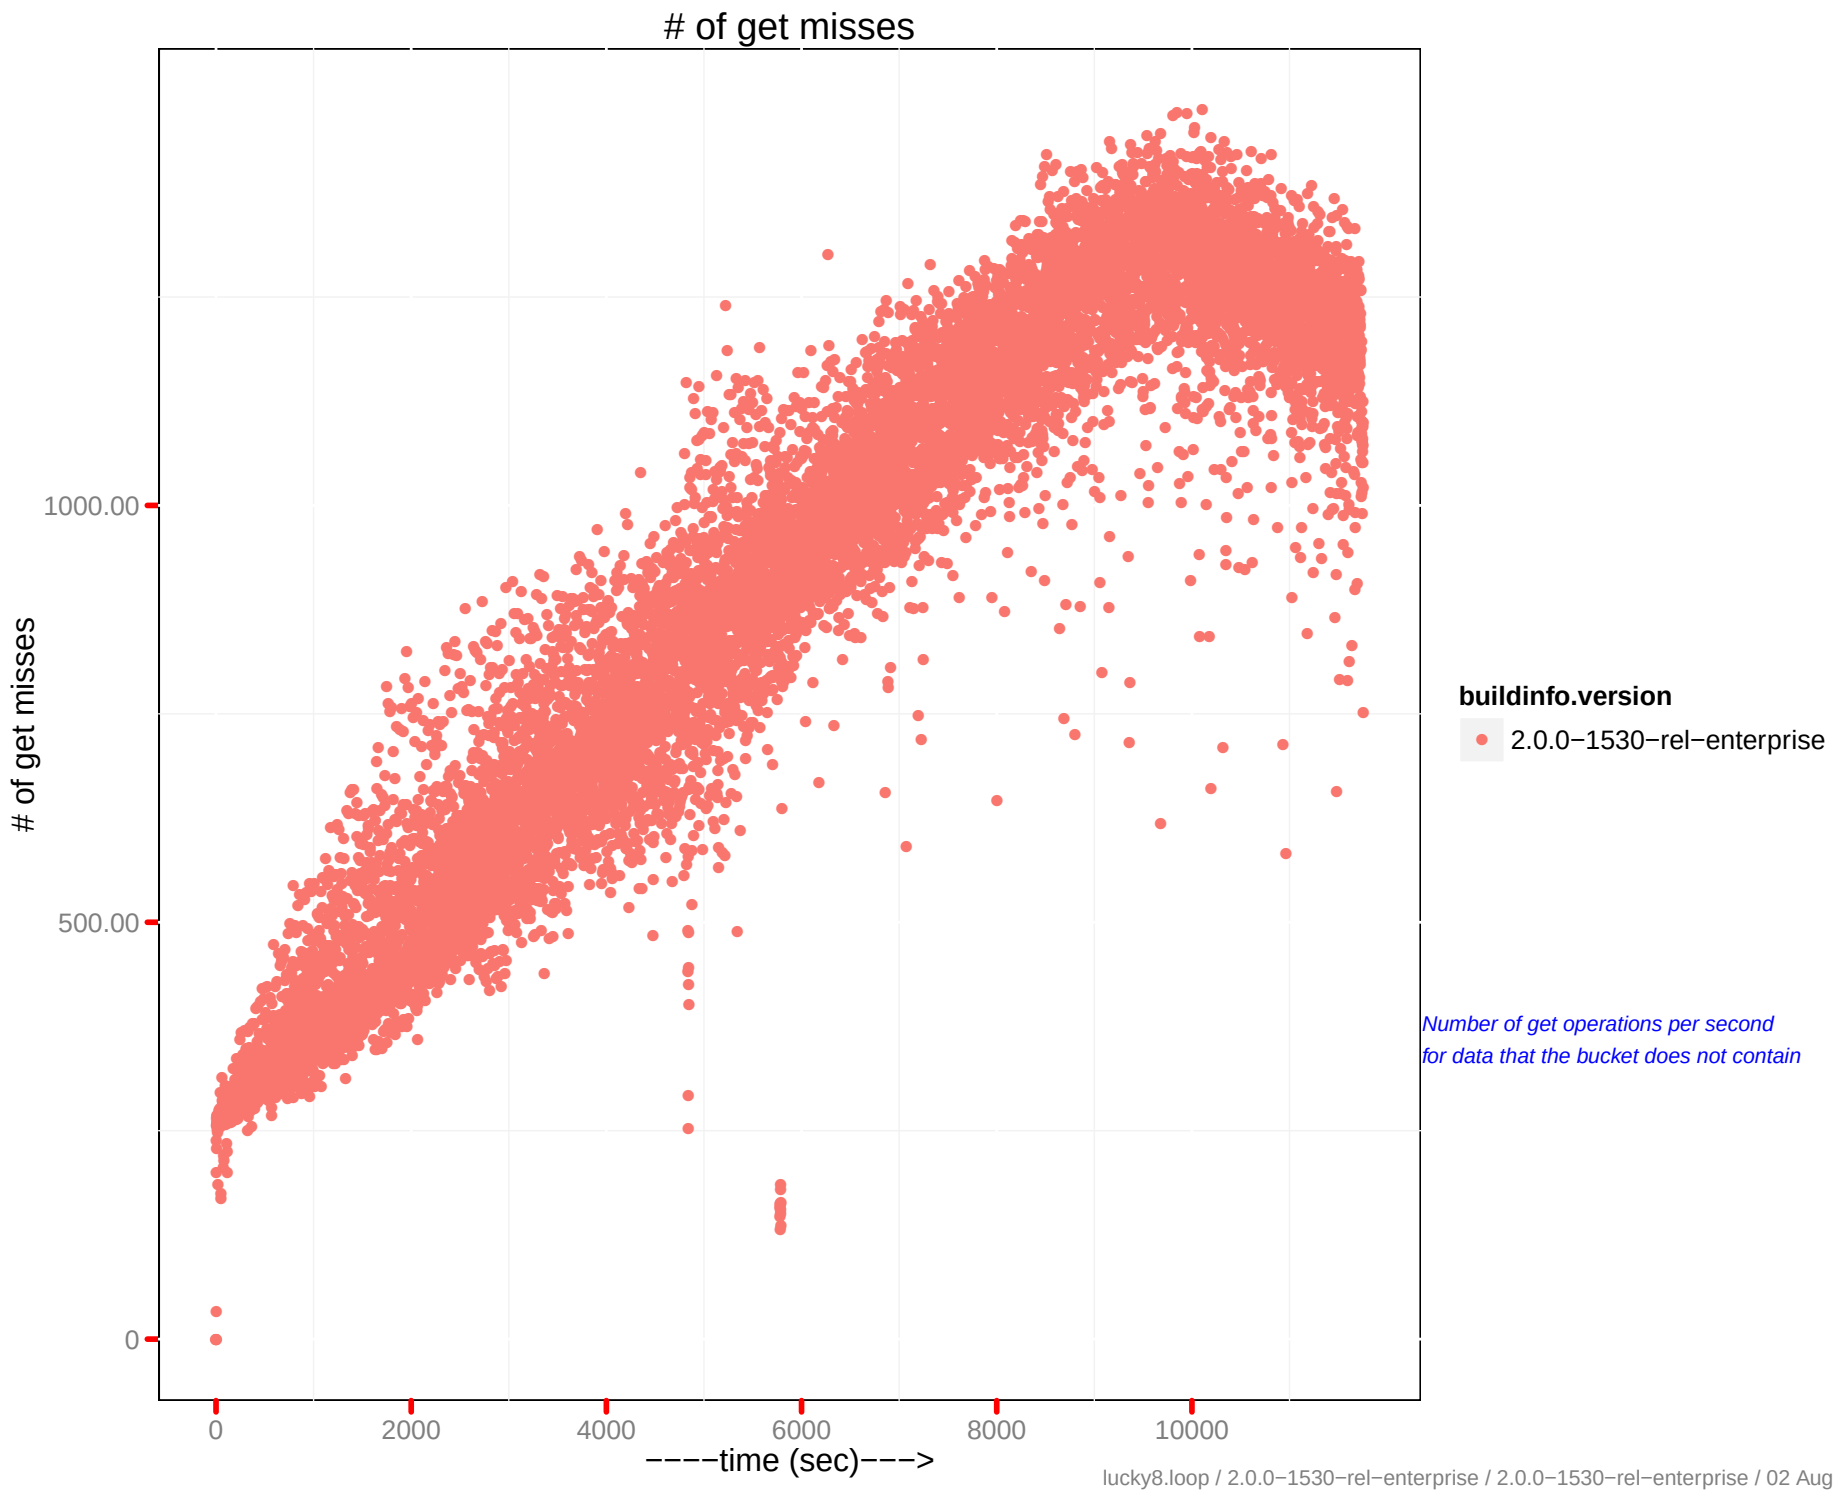
# of get misses
1000.00
500.00
0
# of get misses
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of get operations per second
for data that the bucket does not contain
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

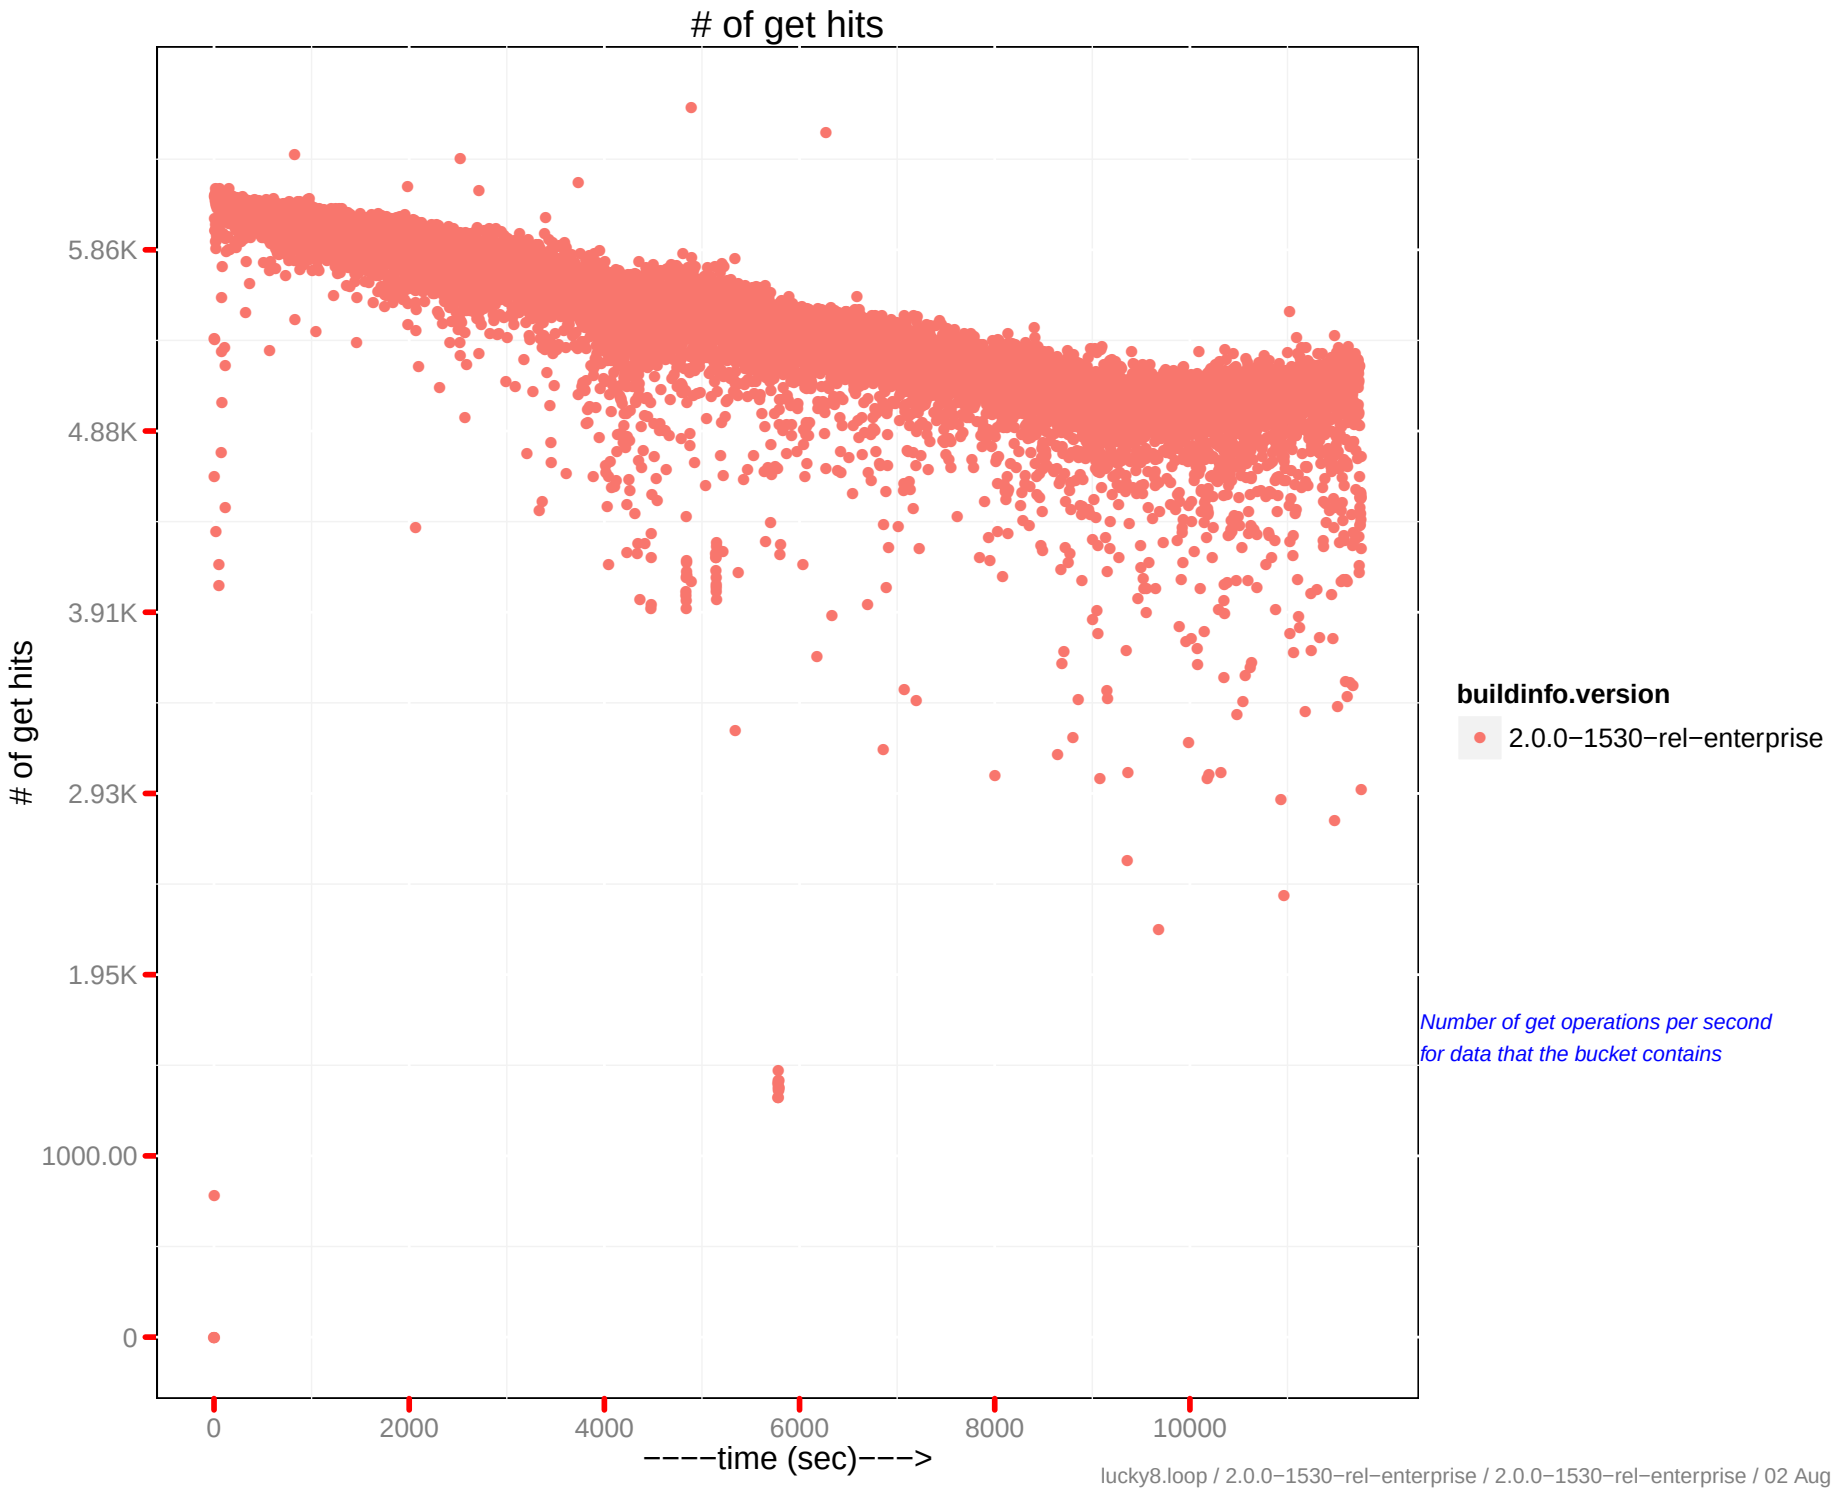
# of get hits
5.86K
4.88K
3.91K
2.93K
1.95K
1000.00
0
# of get hits
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of get operations per second
for data that the bucket contains
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

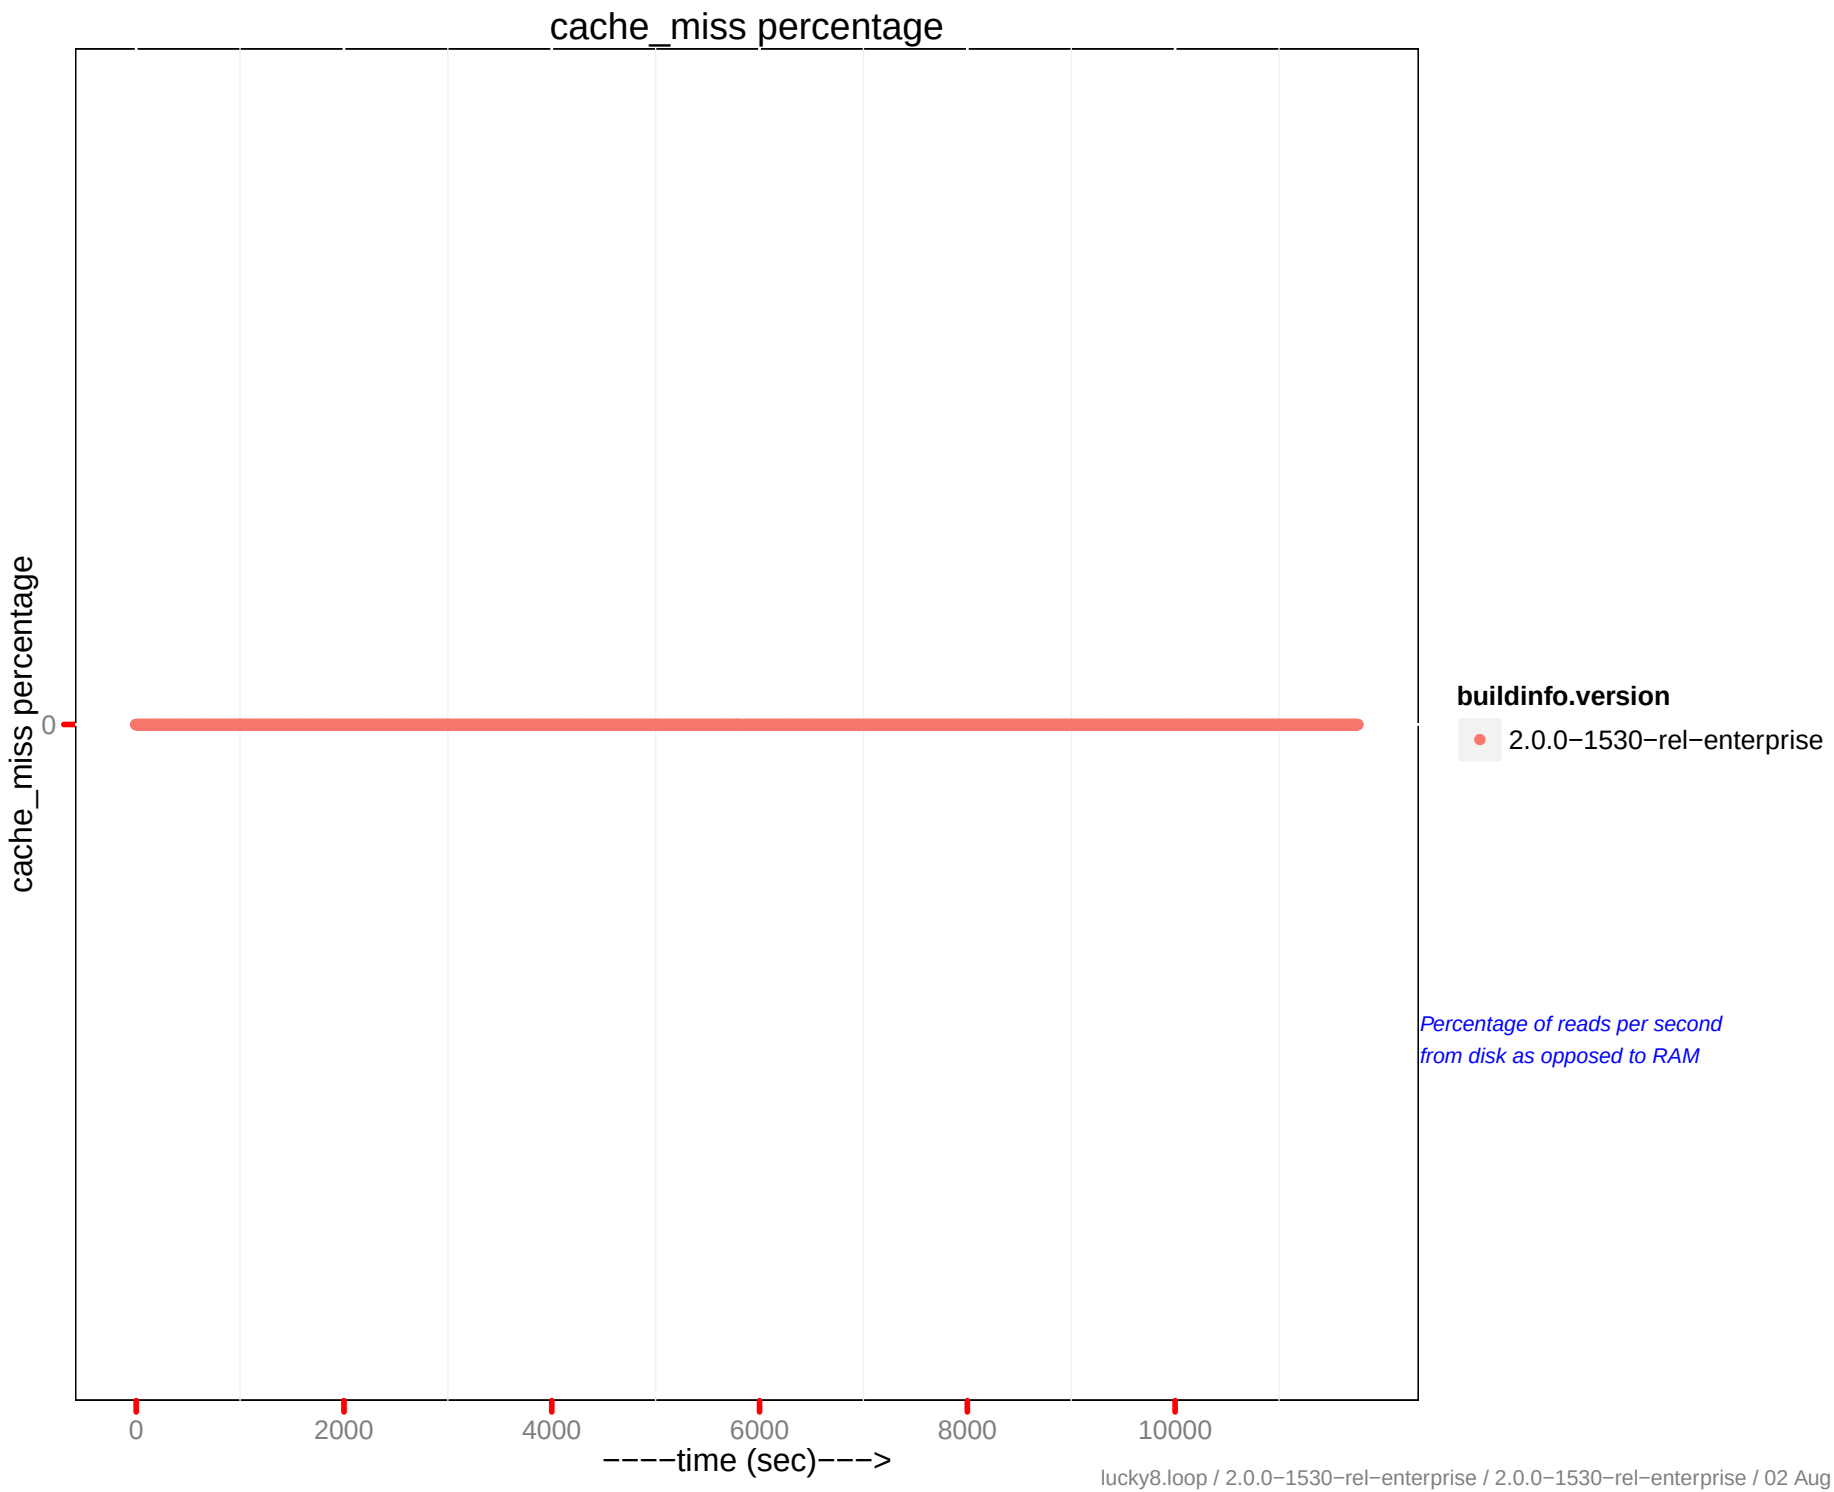
cache_miss percentage
cache_miss percentage
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
Percentage of reads per second
from disk as opposed to RAM
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

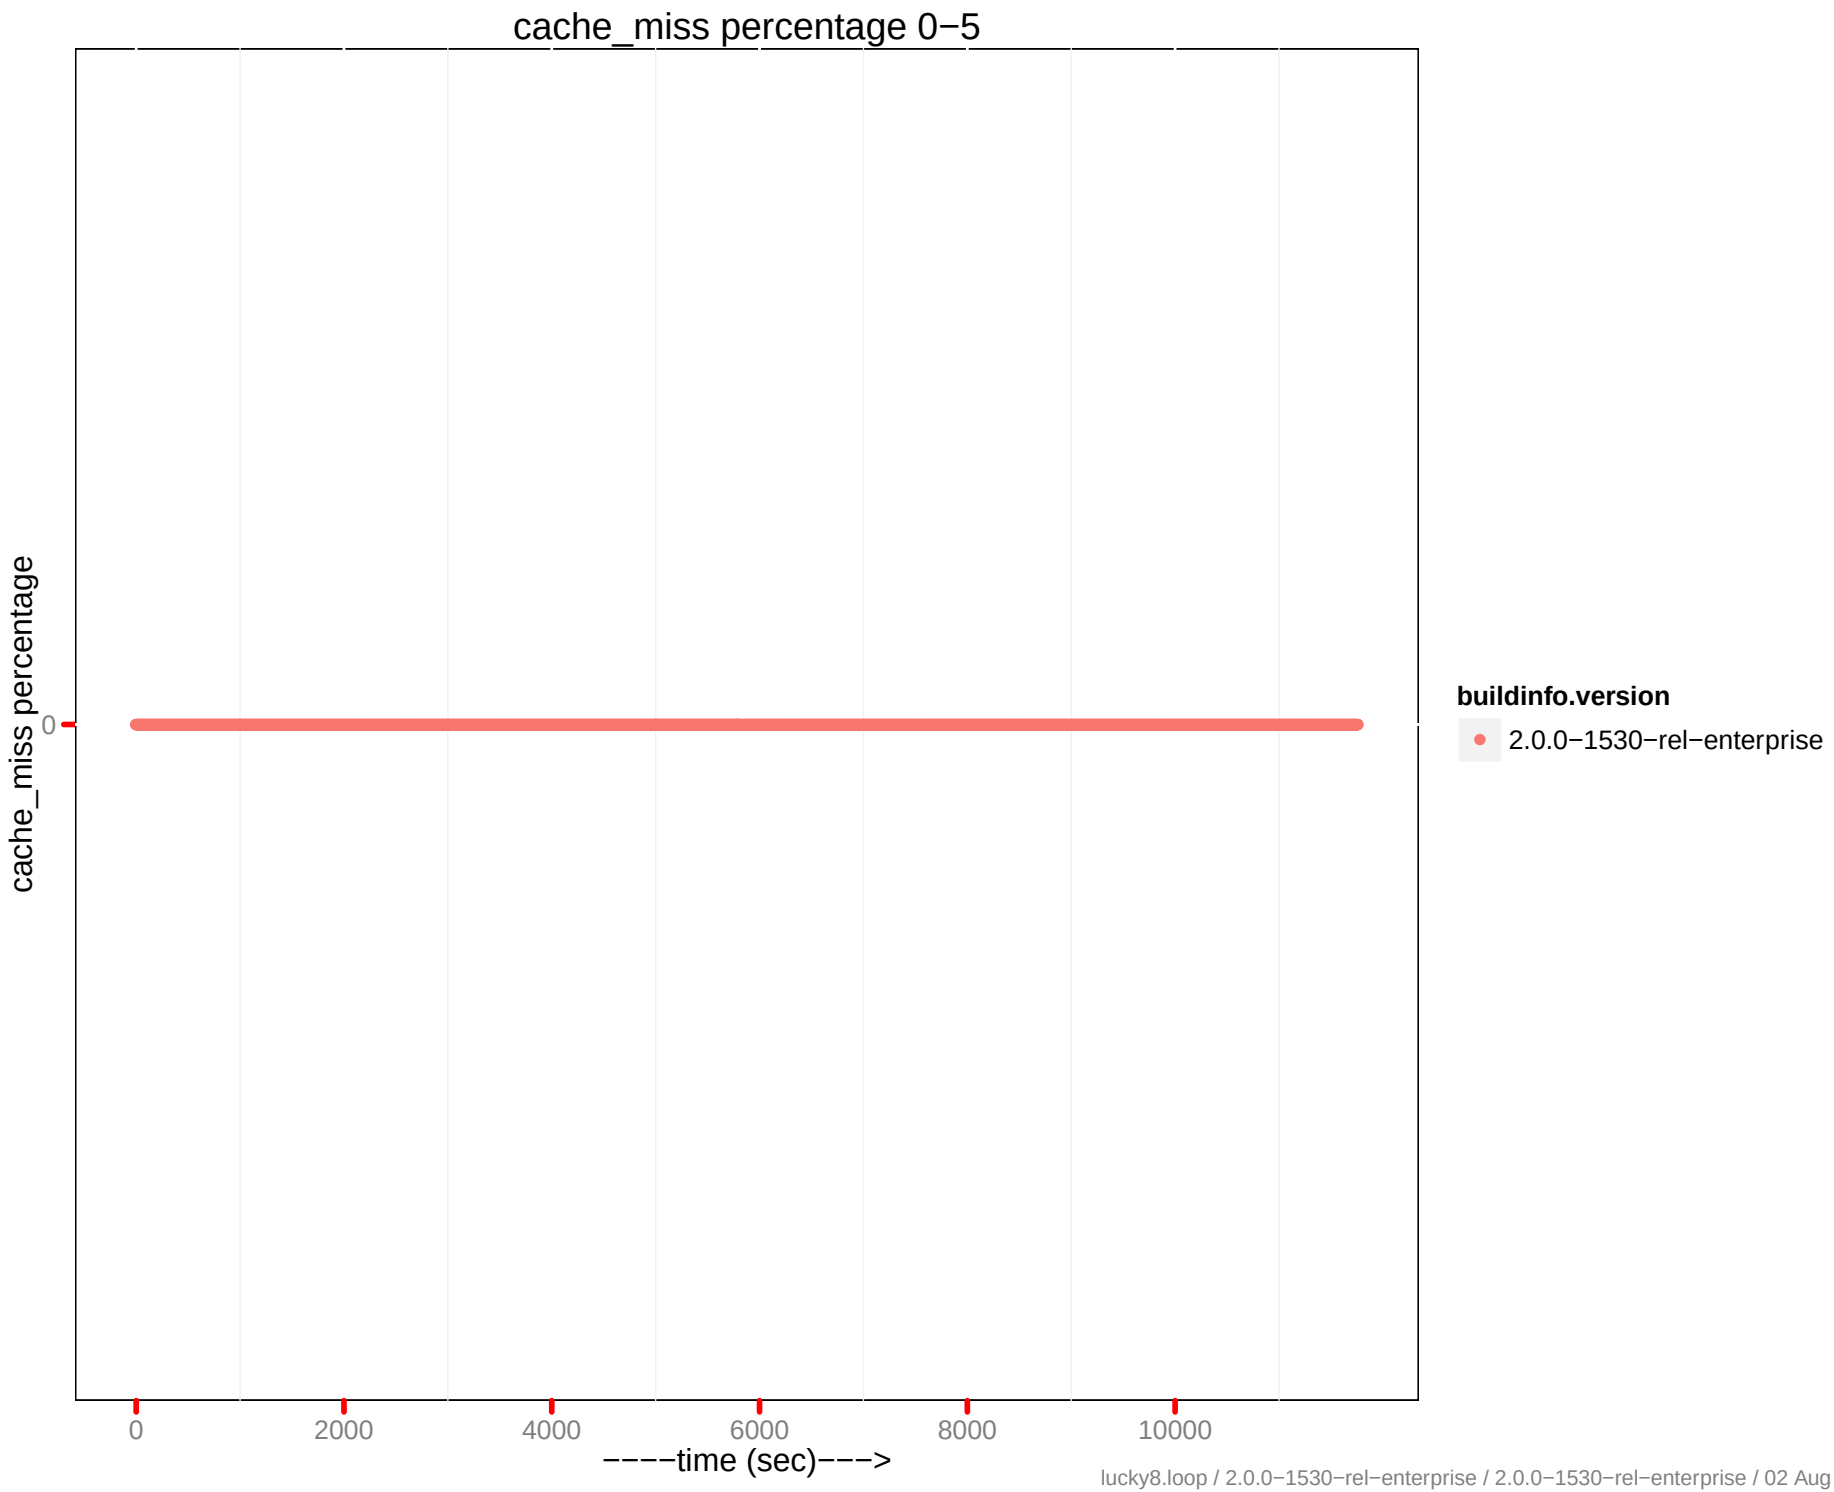
cache_miss percentage 0−5
cache_miss percentage
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

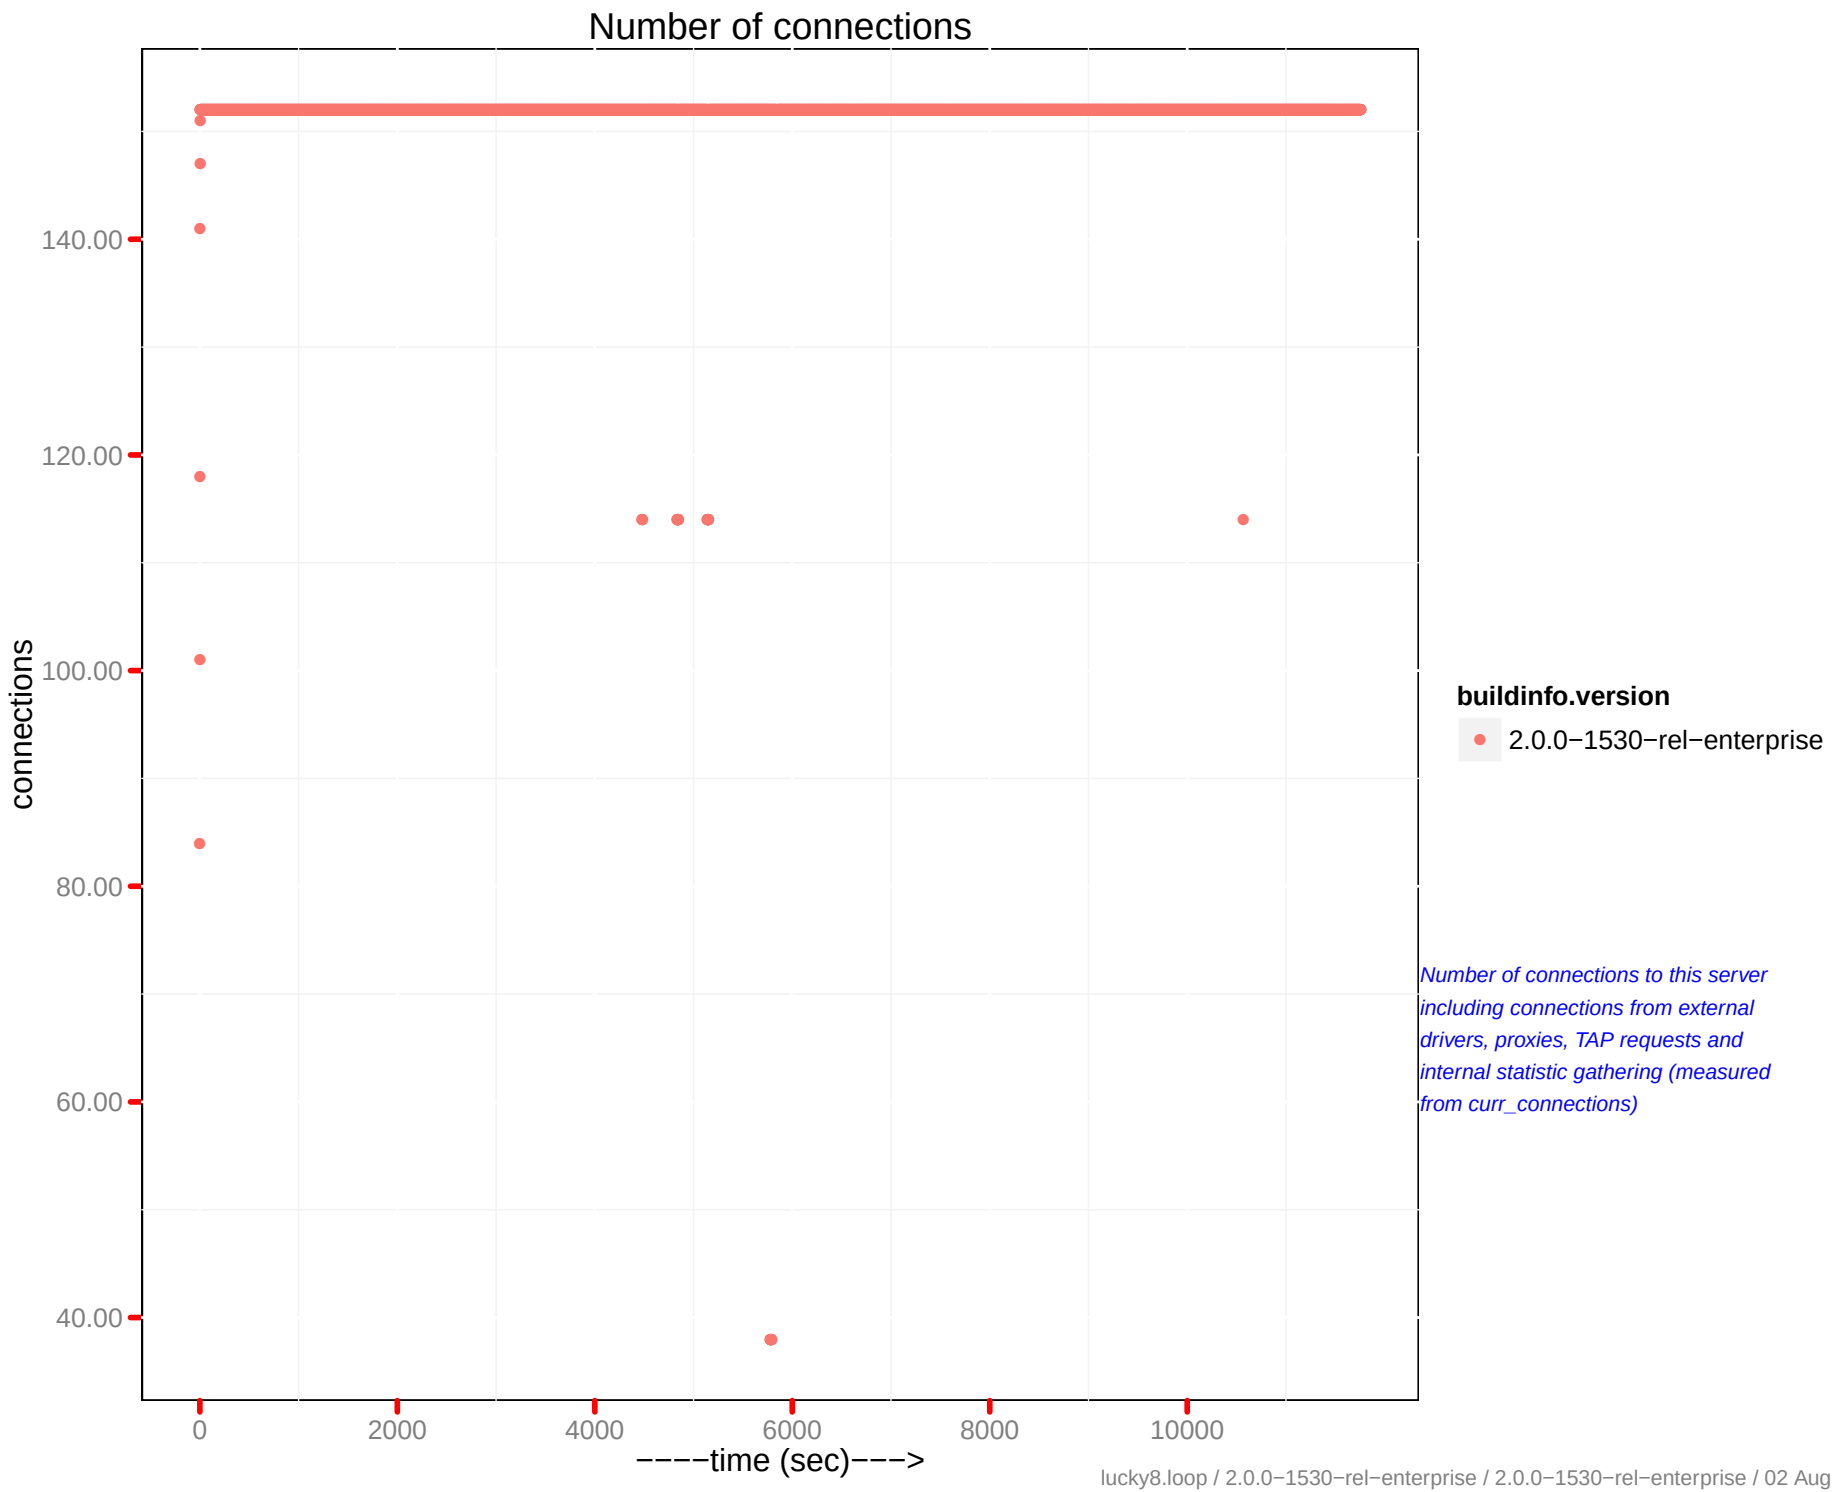
Number of connections
connections
140.00
120.00
100.00
80.00
60.00
40.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Number of connections to this server including connections from external drivers, proxies, TAP requests and internal statistic gathering (measured from curr_connections)
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

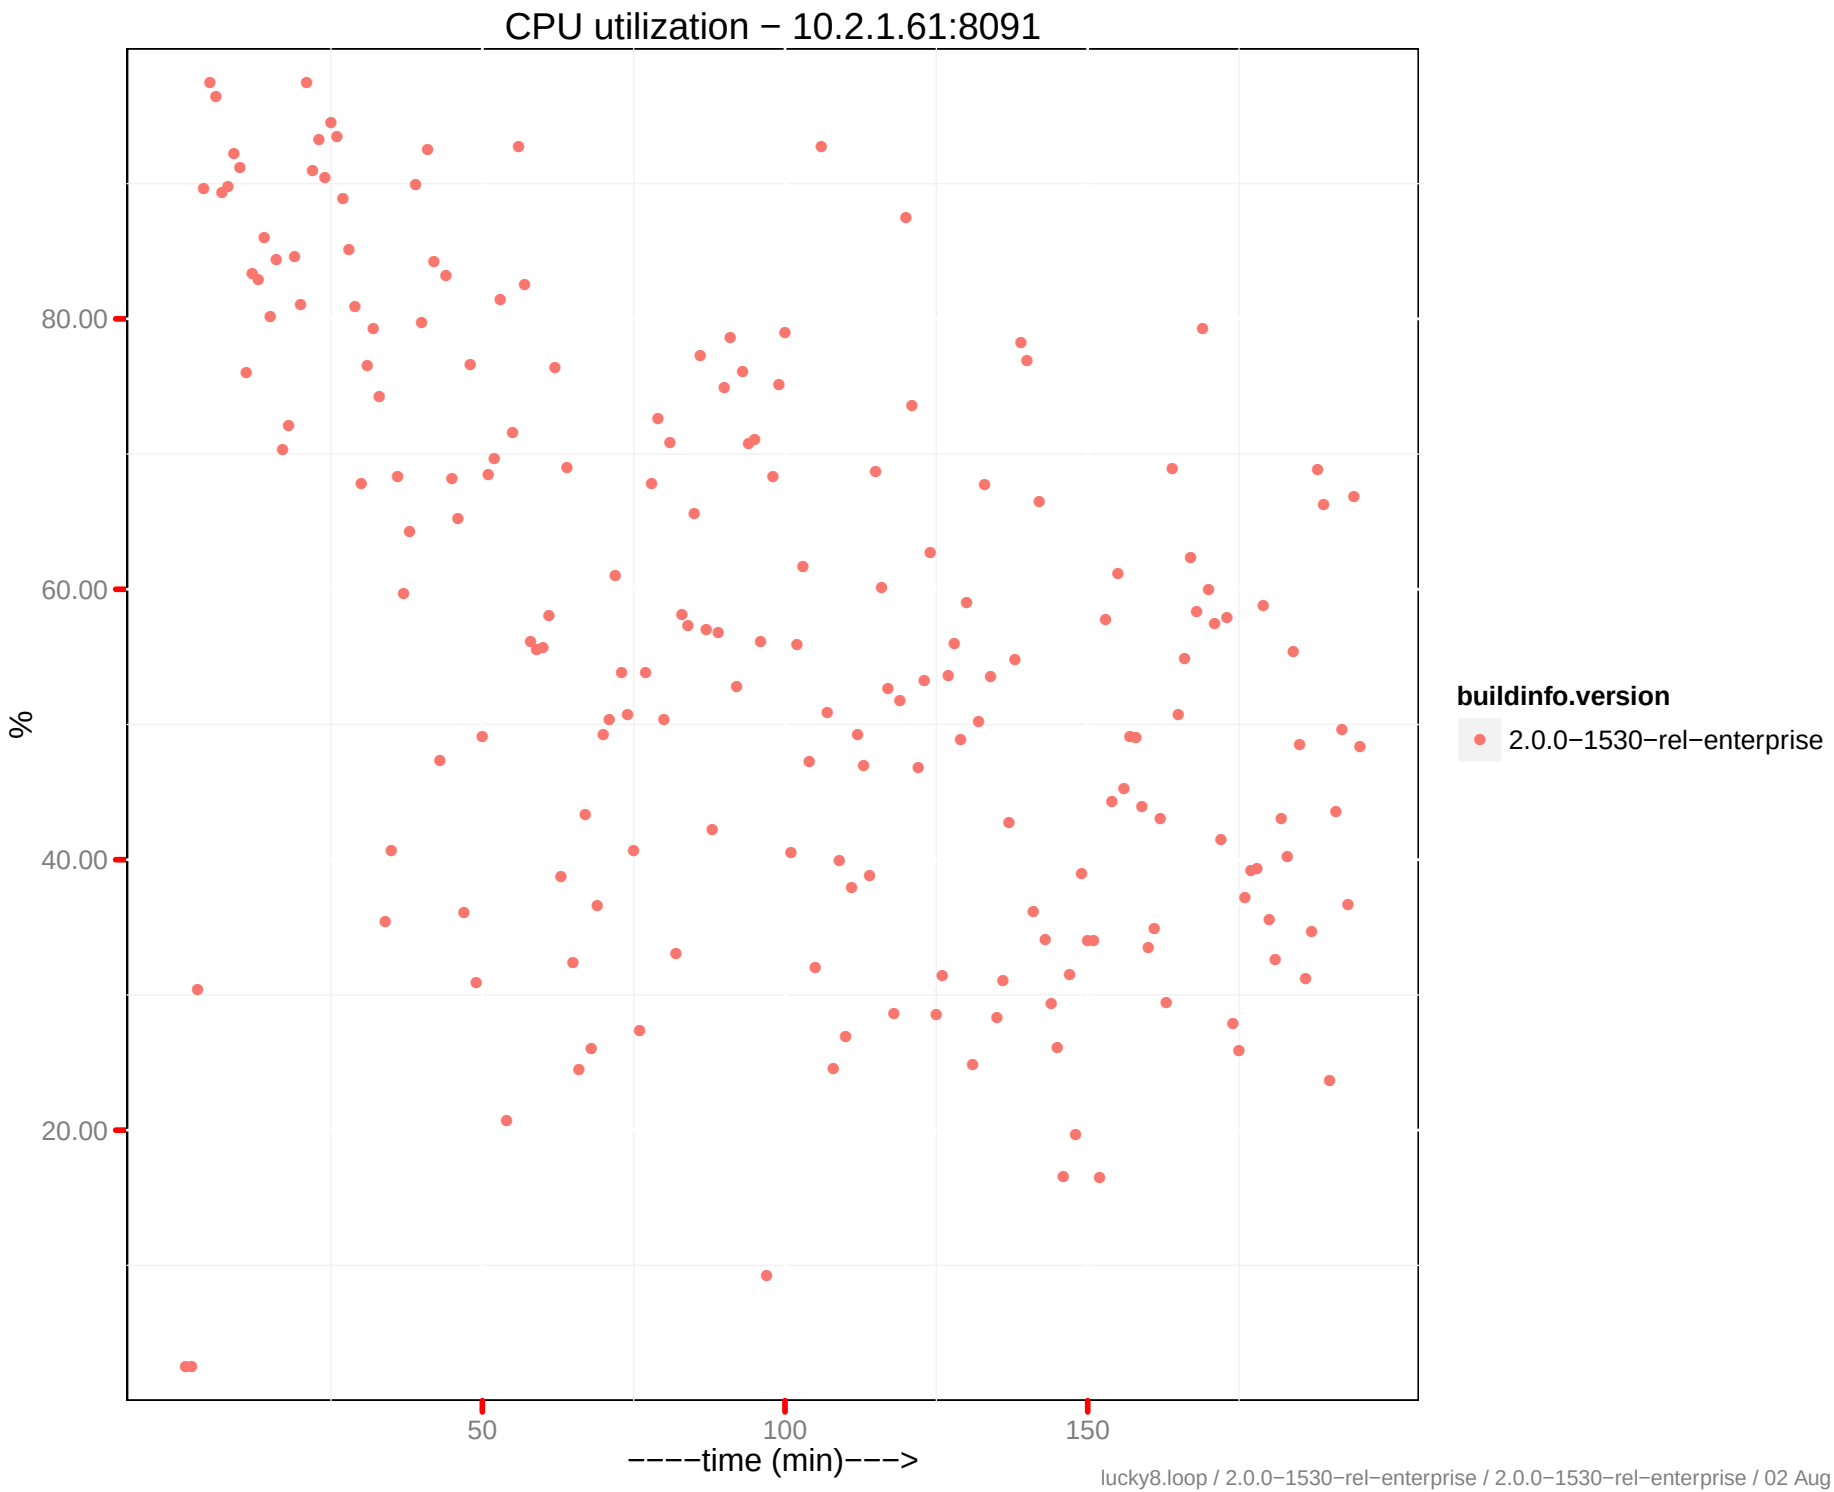
CPU utilization − 10.2.1.61:8091
80.00
60.00
40.00
20.00
%
50
100
150
−−−−time (min)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise

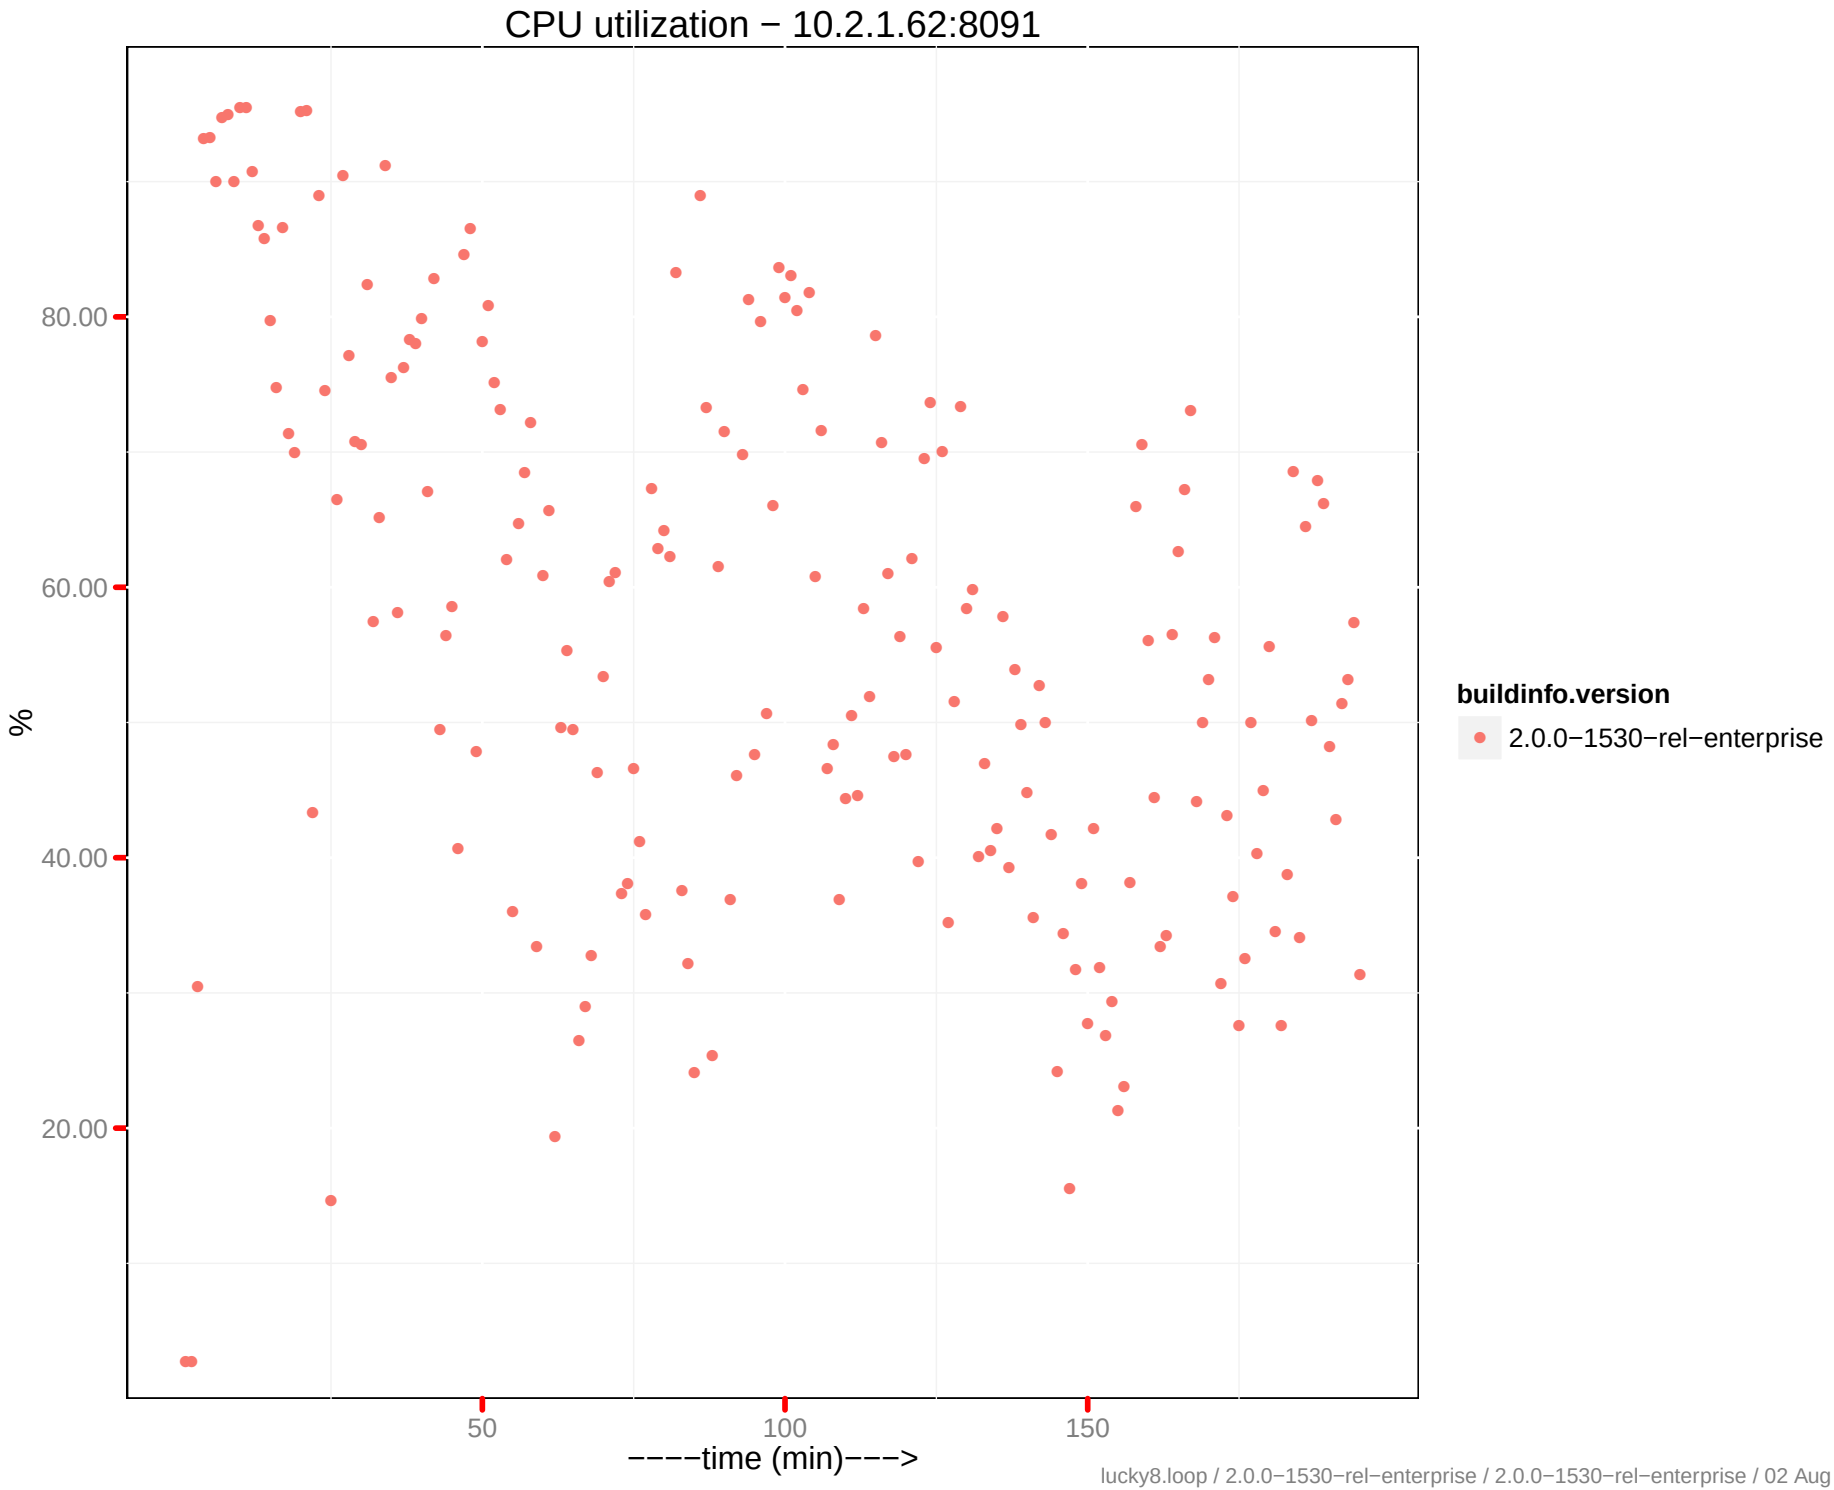
CPU utilization − 10.2.1.62:8091
80.00
60.00
40.00
20.00
%
50
100
150
−−−−time (min)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

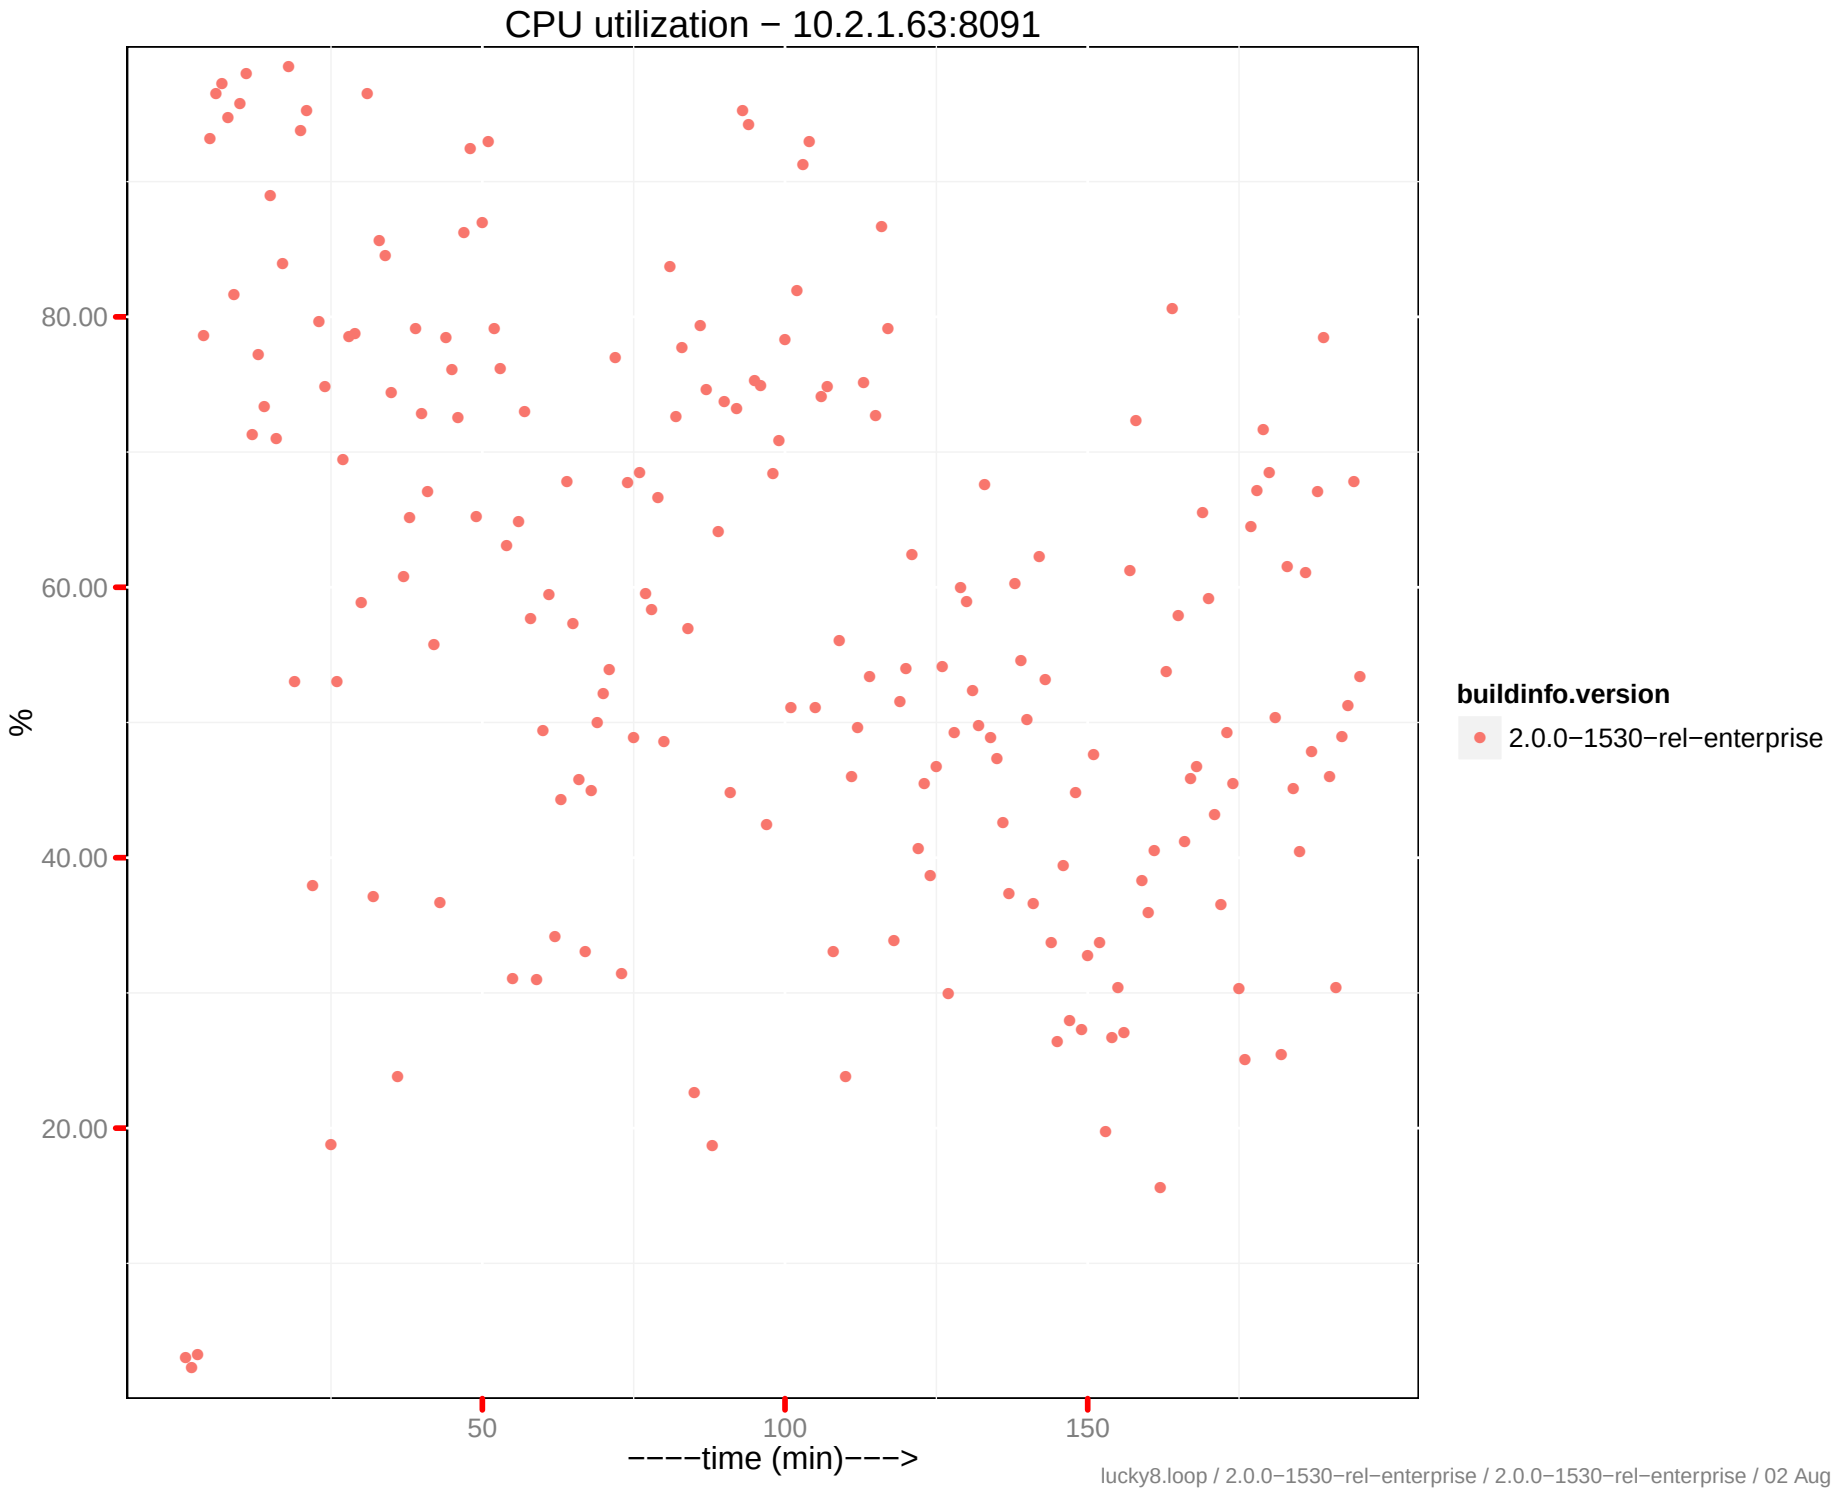
CPU utilization − 10.2.1.63:8091
80.00
60.00
40.00
20.00
%
50
100
150
−−−−time (min)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

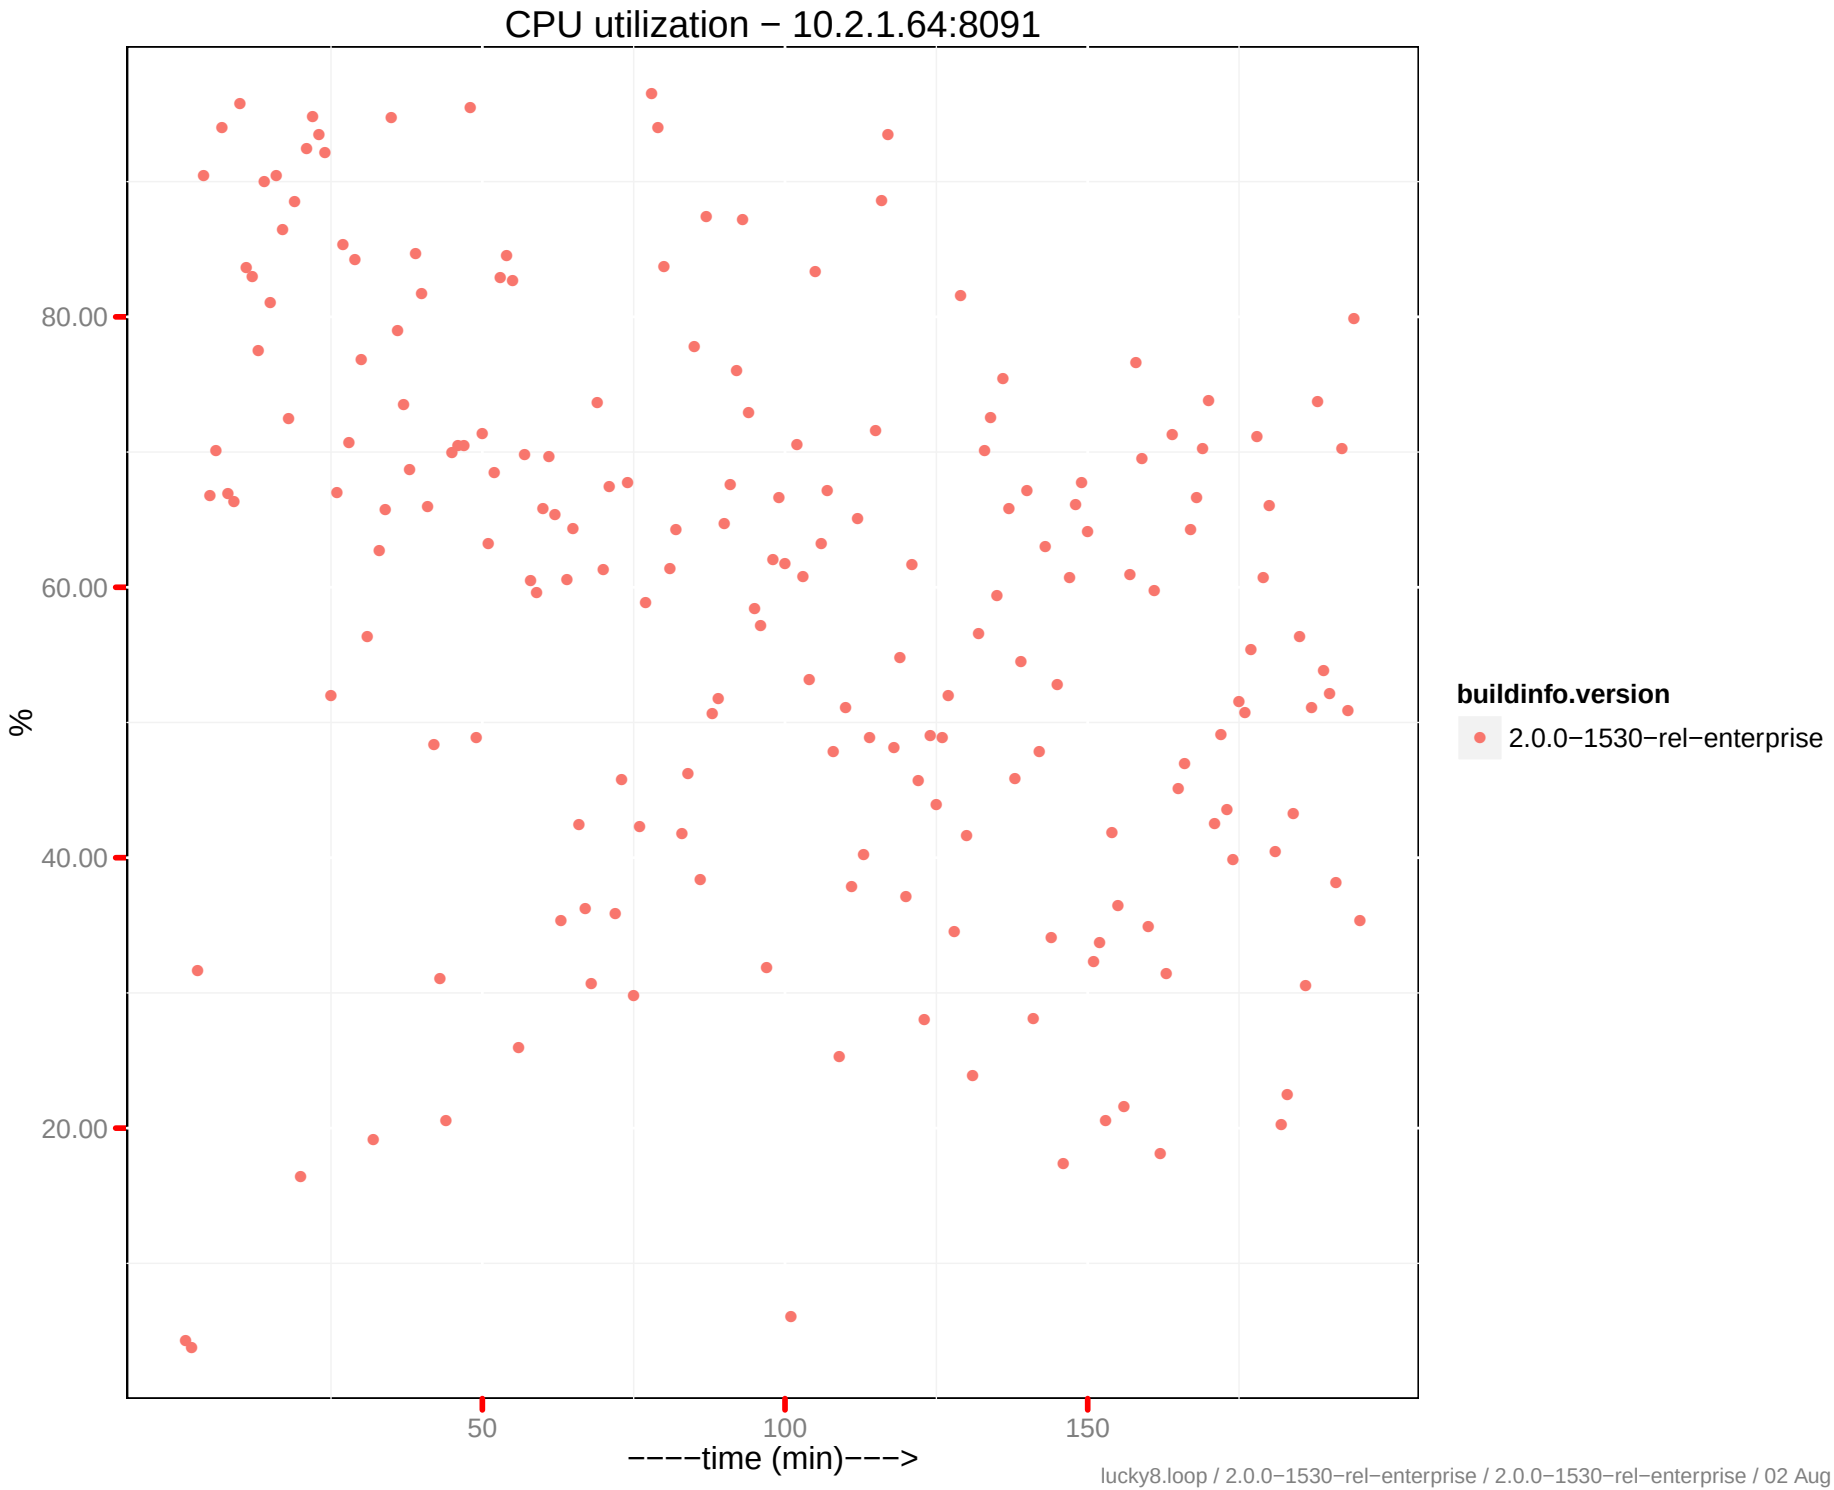
CPU utilization − 10.2.1.64:8091
80.00
60.00
40.00
20.00
%
50
100
150
−−−−time (min)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

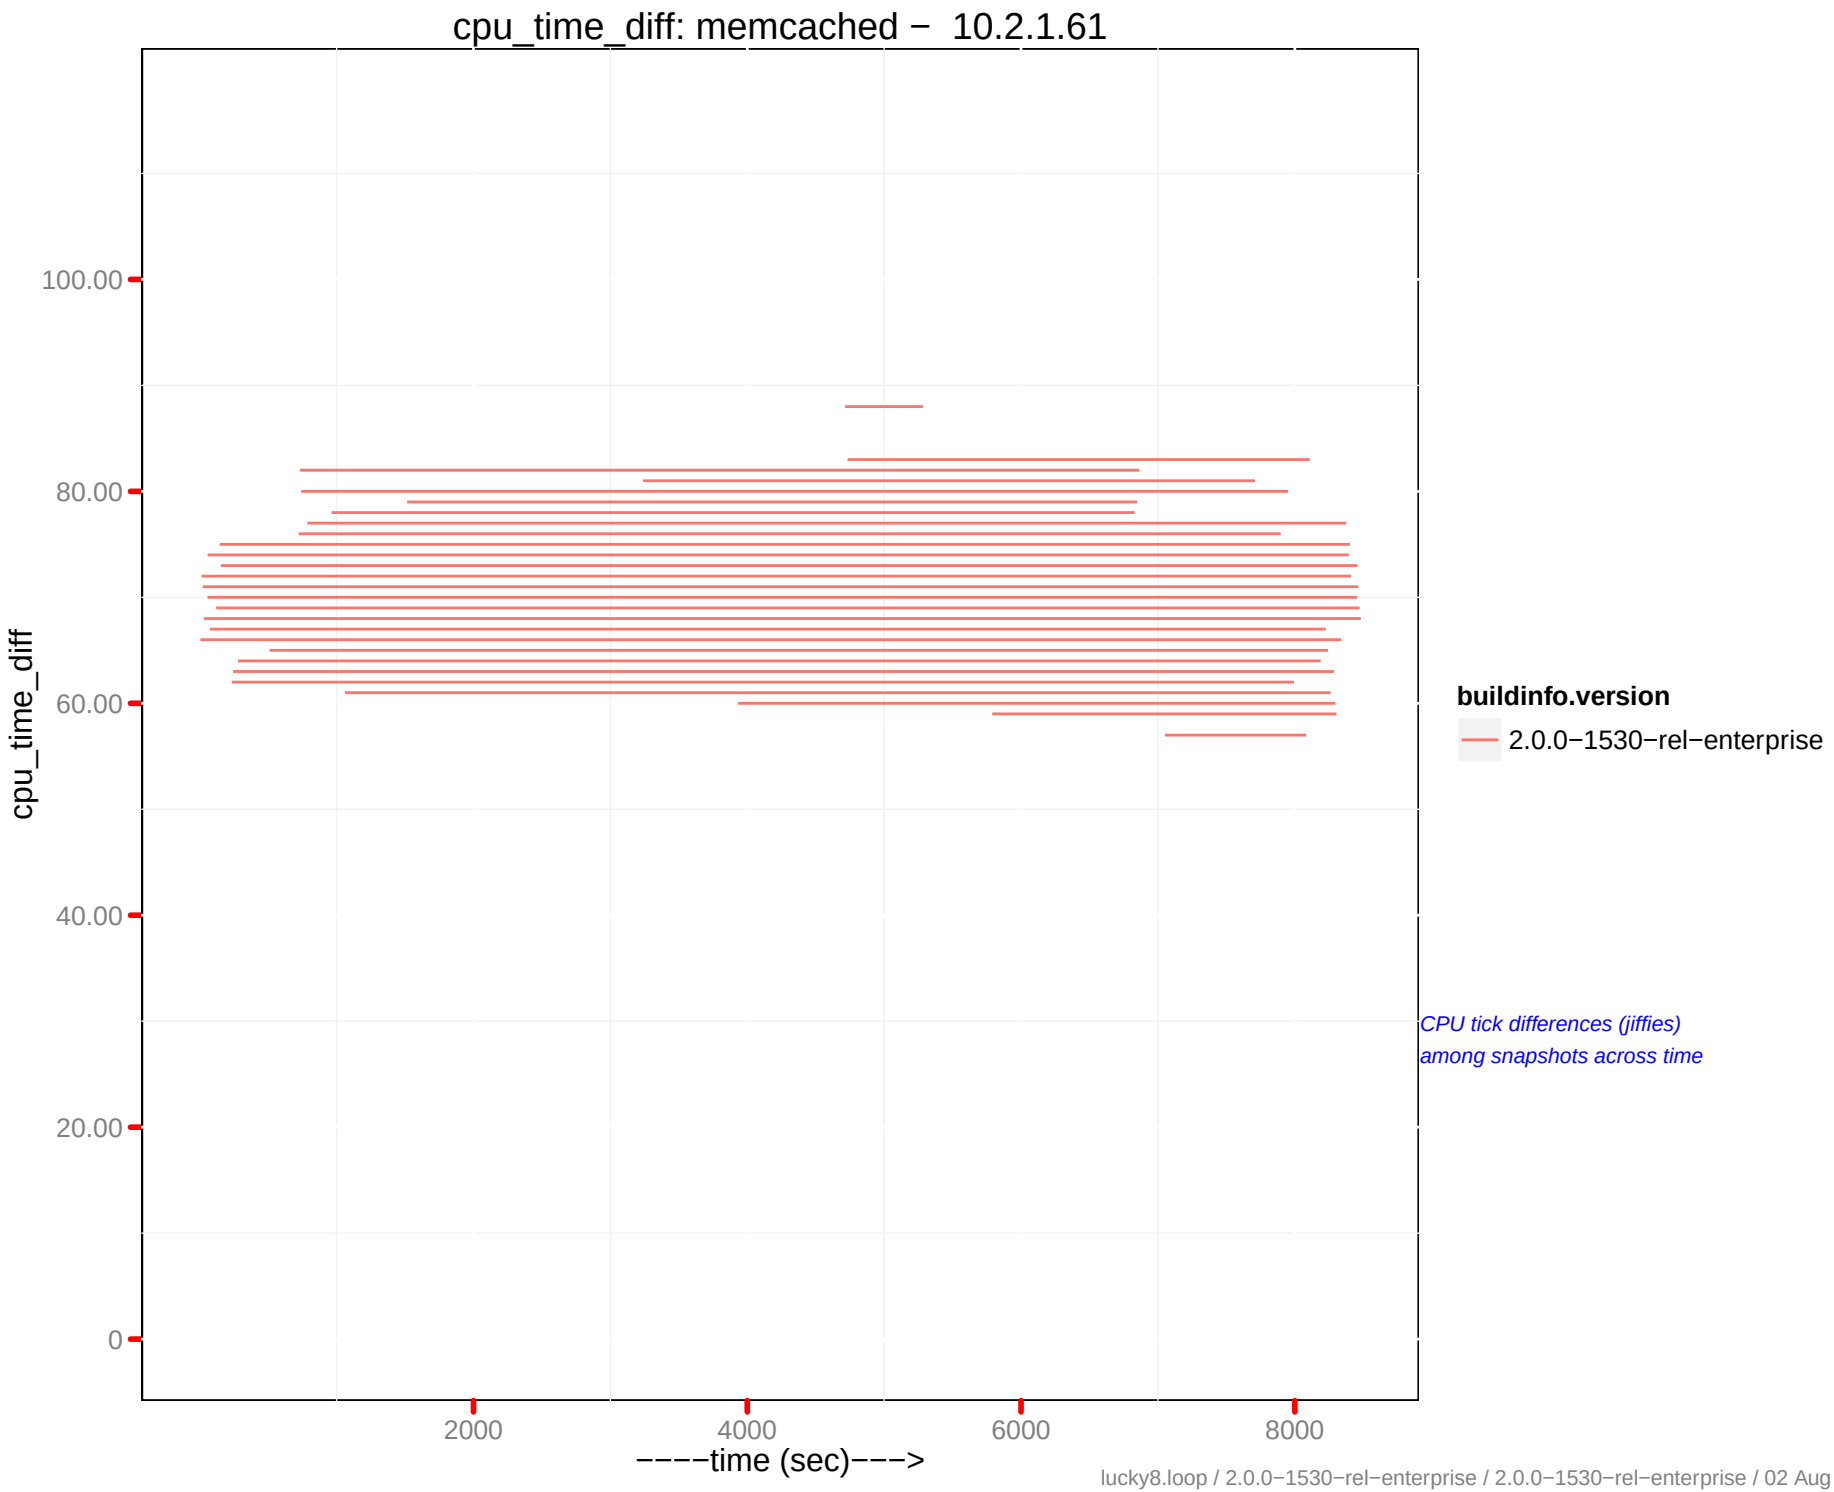
cpu_time_diff: memcached − 10.2.1.61
100.00
80.00
60.00
40.00
20.00
0
cpu_time_diff
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

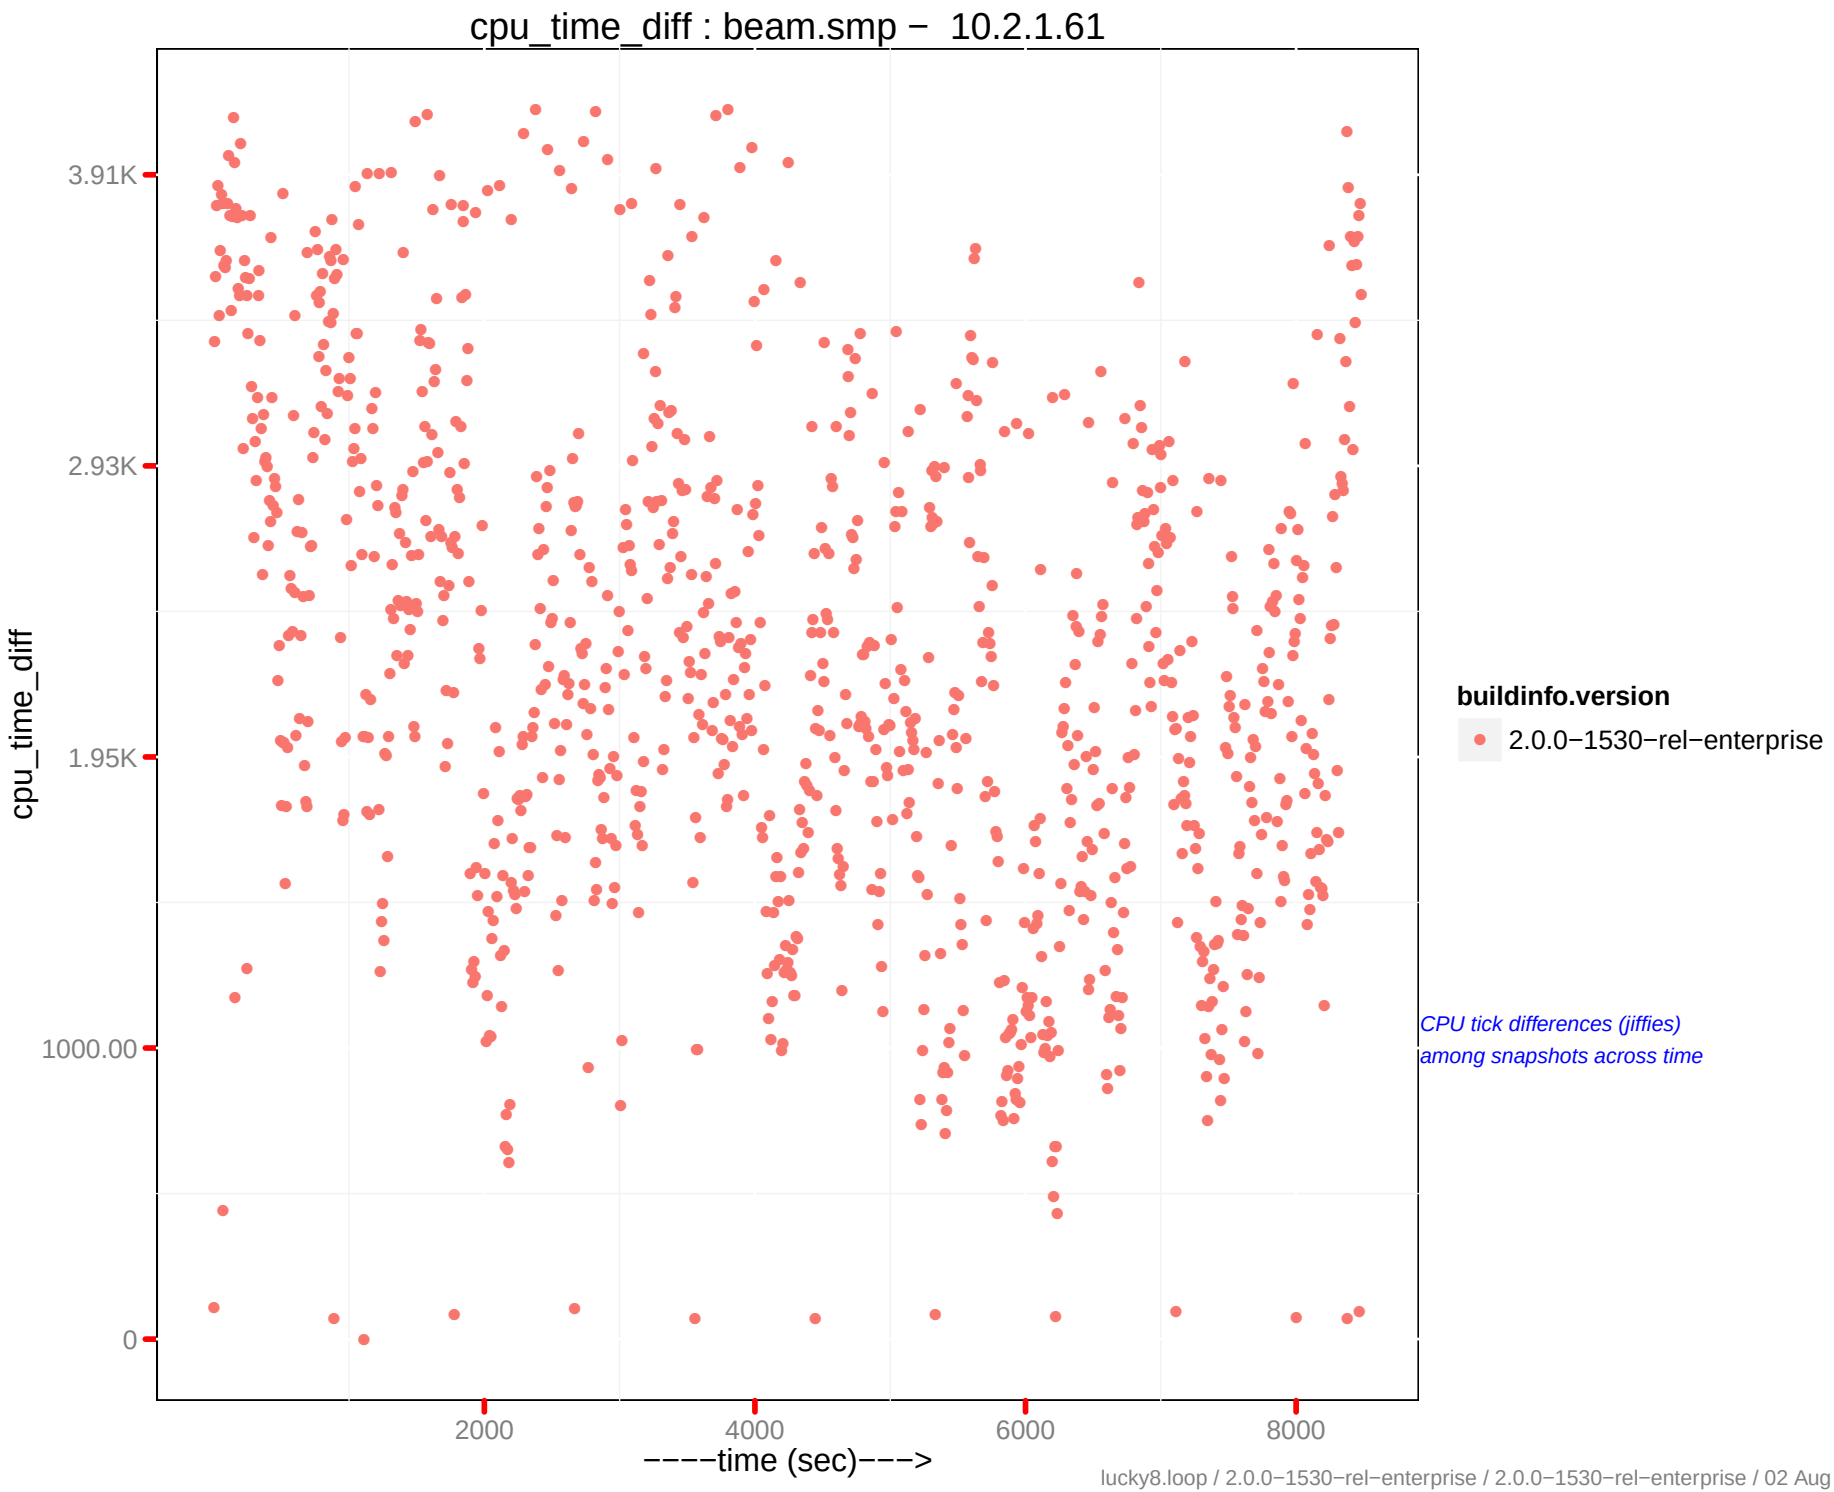
cpu_time_diff : beam.smp − 10.2.1.61
3.91K
2.93K
1.95K
1000.00
0
cpu_time_diff
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

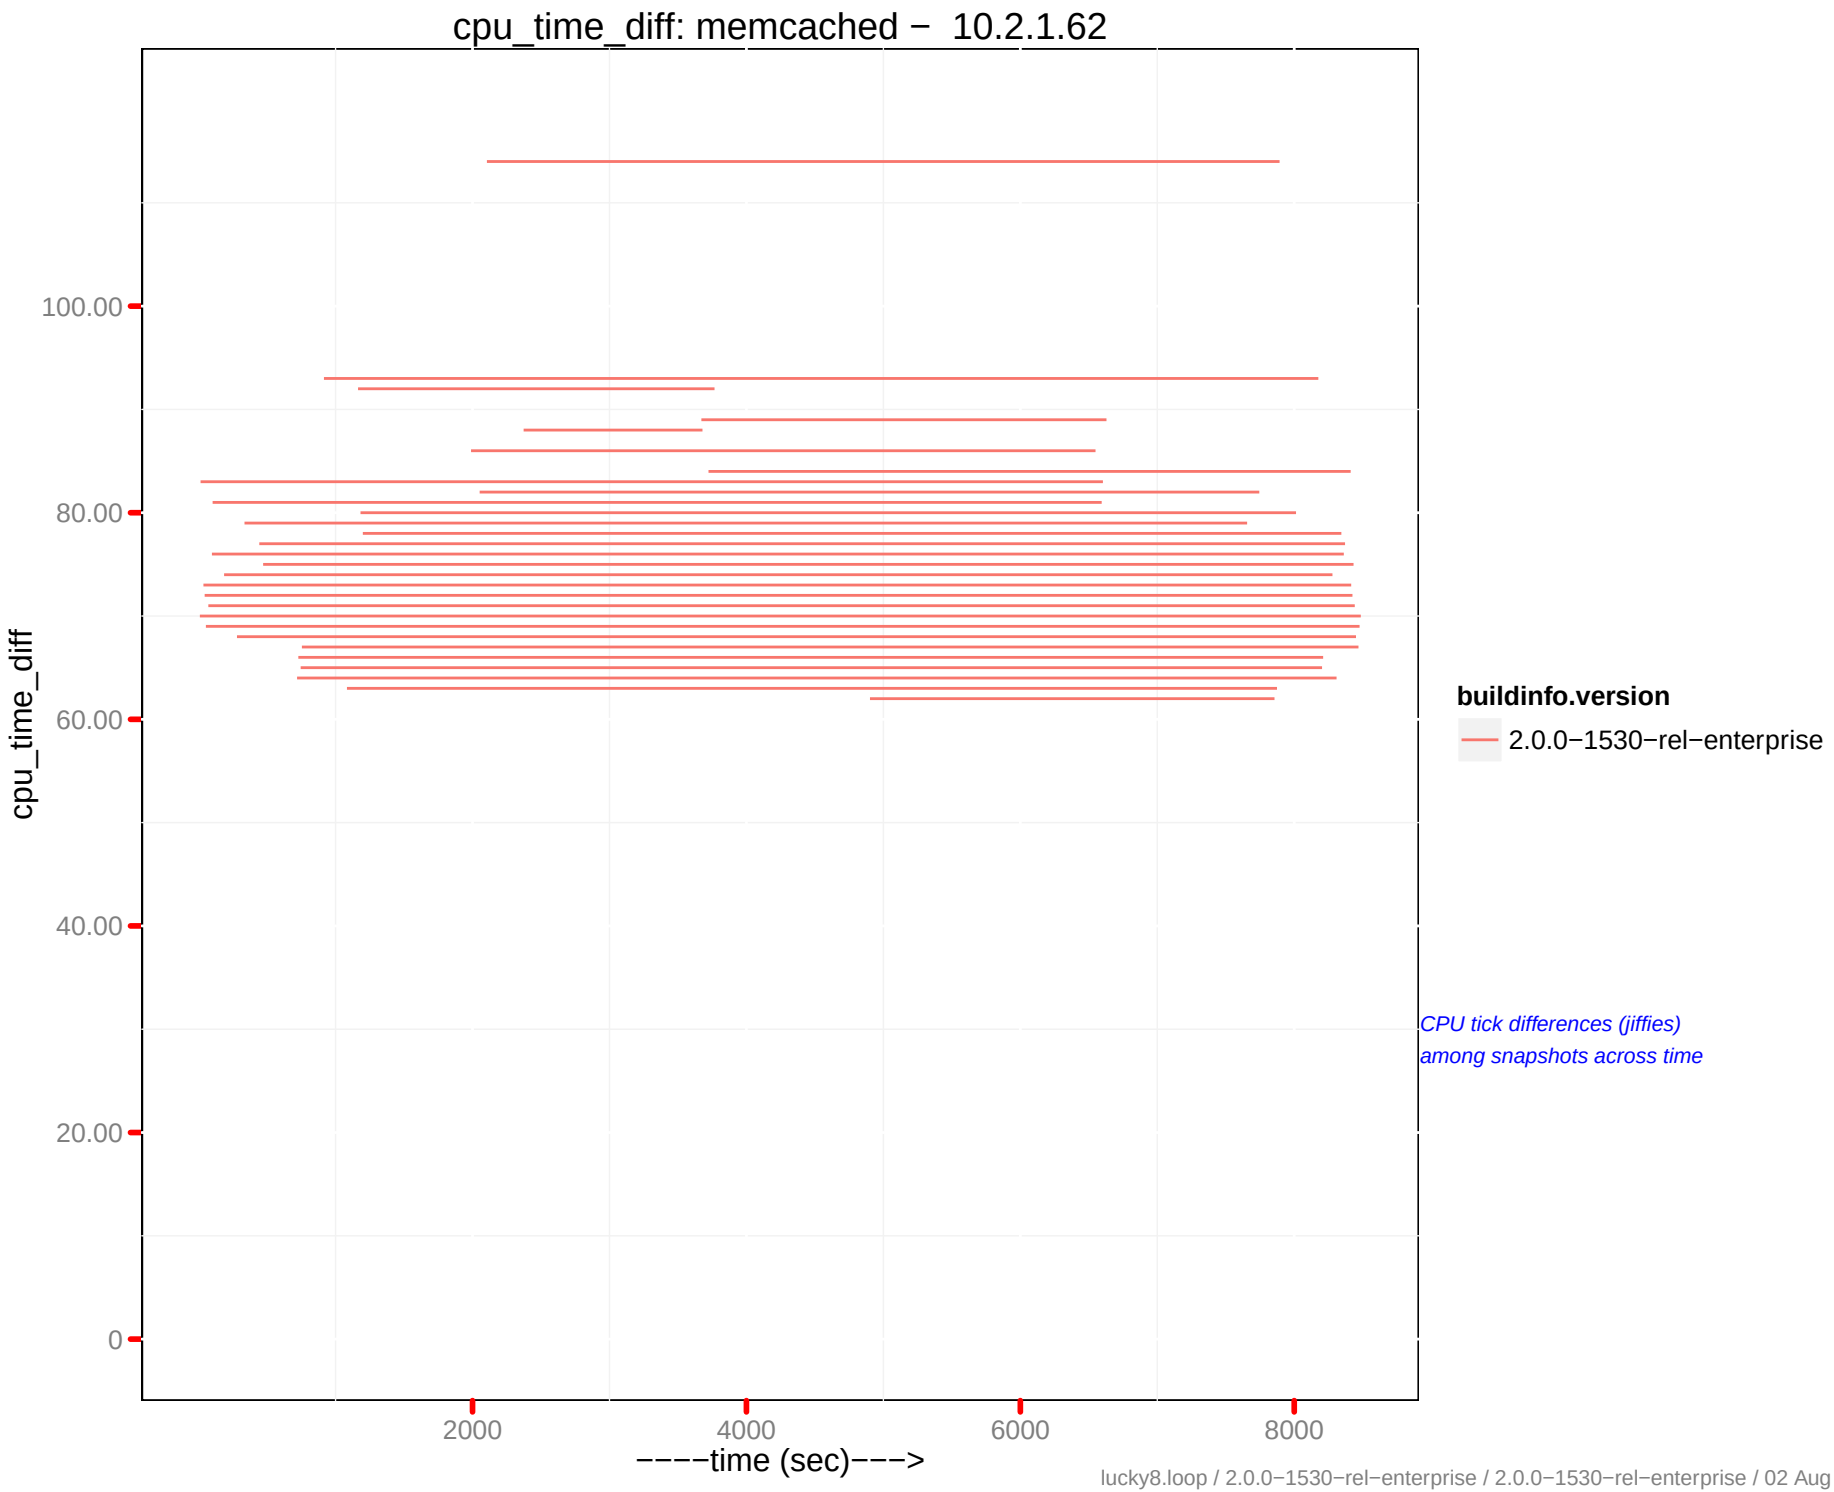
cpu_time_diff: memcached − 10.2.1.62
cpu_time_diff
100.00
80.00
60.00
40.00
20.00
0
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

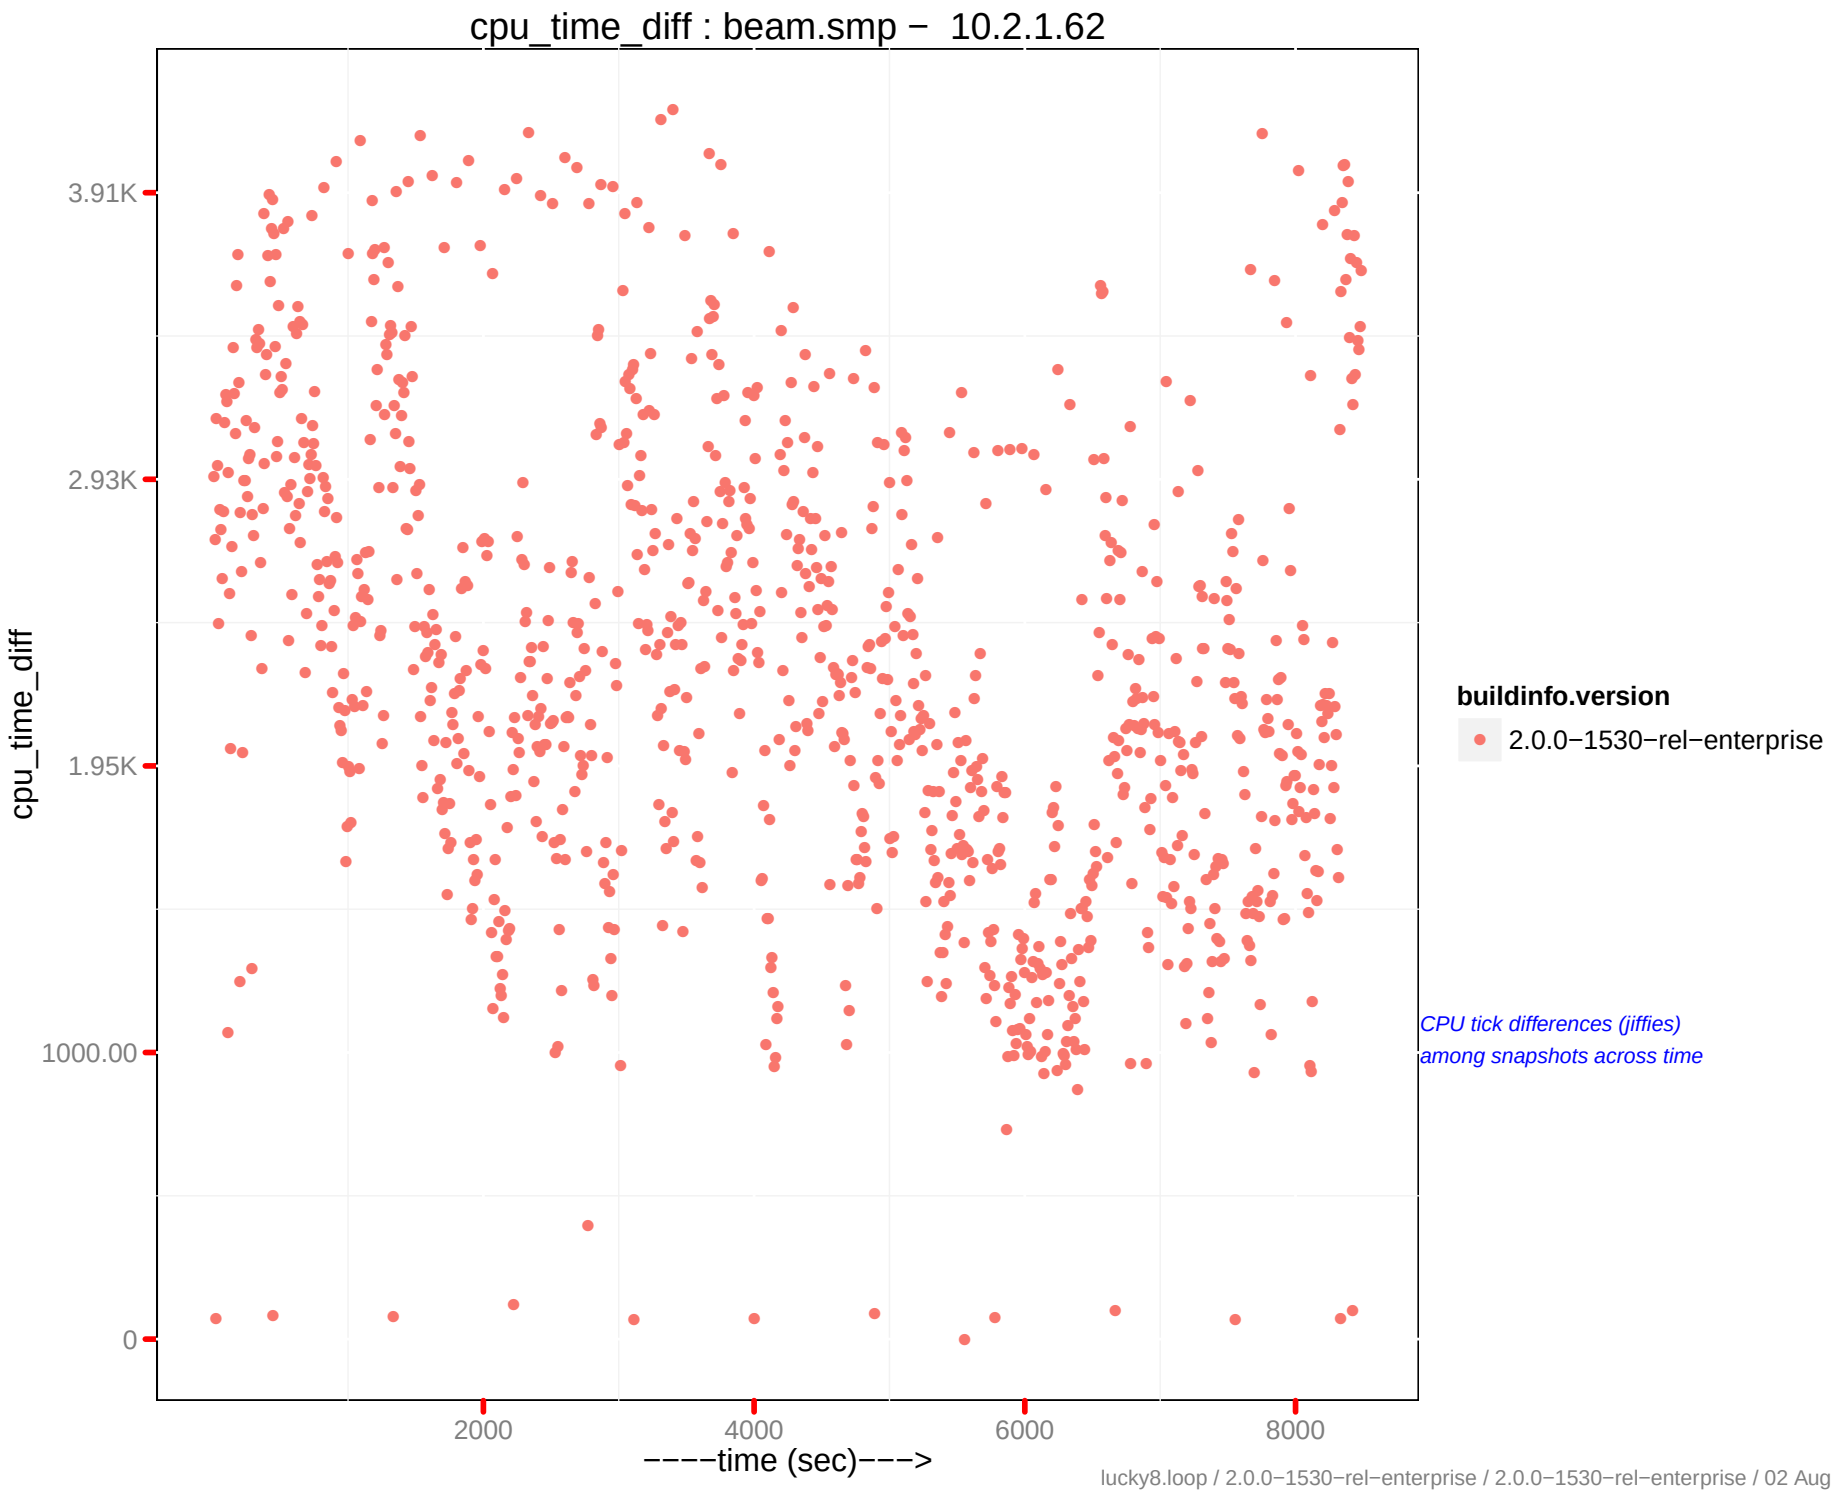
cpu_time_diff : beam.smp − 10.2.1.62
cpu_time_diff
3.91K
2.93K
1.95K
1000.00
0
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

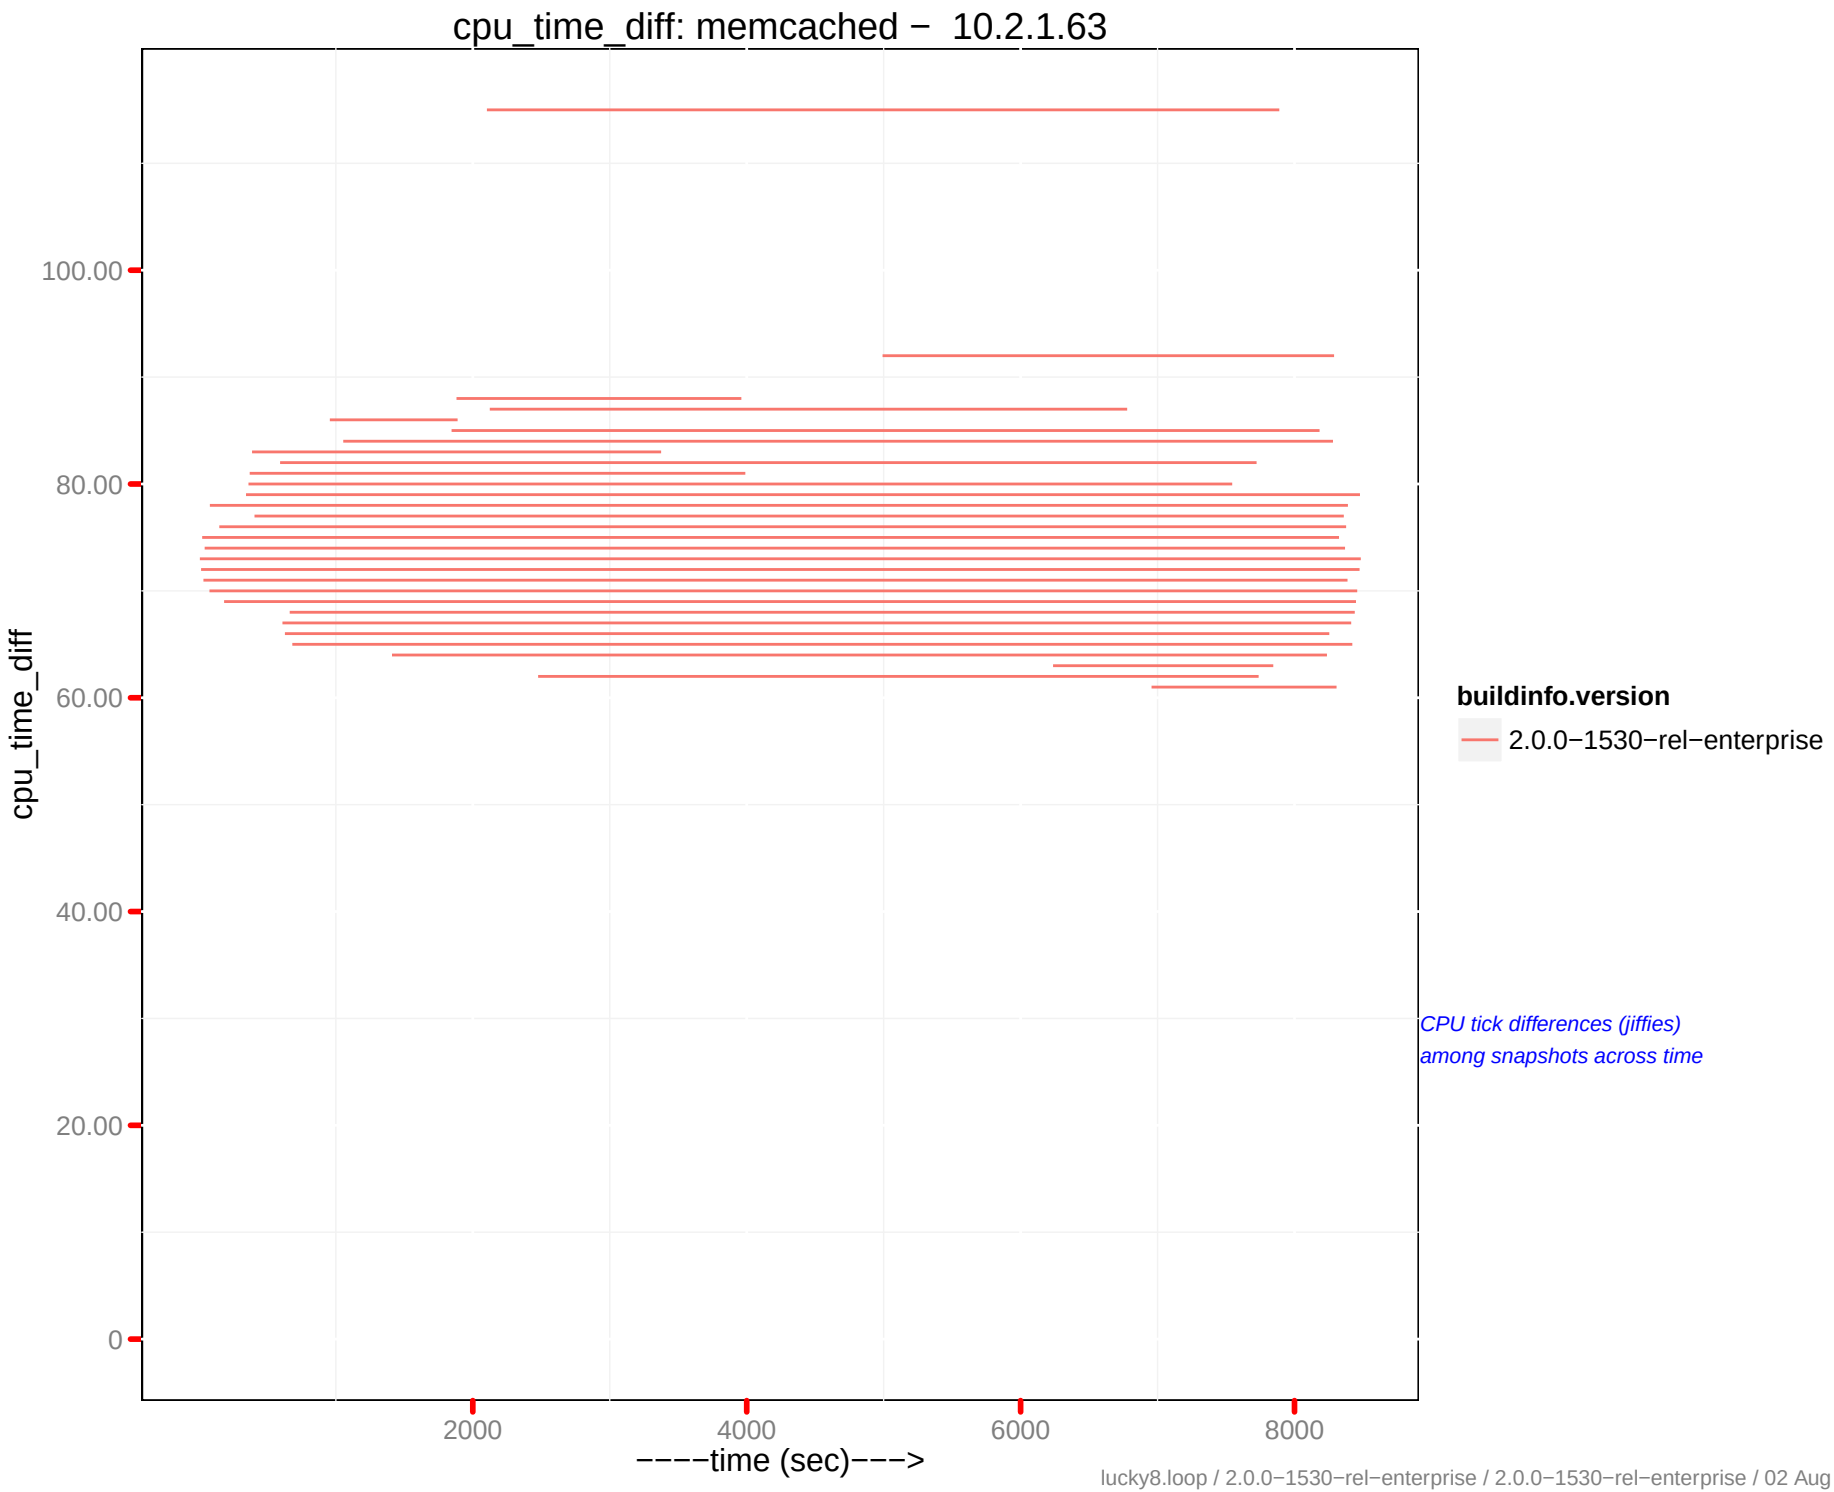
cpu_time_diff: memcached − 10.2.1.63
100.00
80.00
60.00
40.00
20.00
0
cpu_time_diff
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

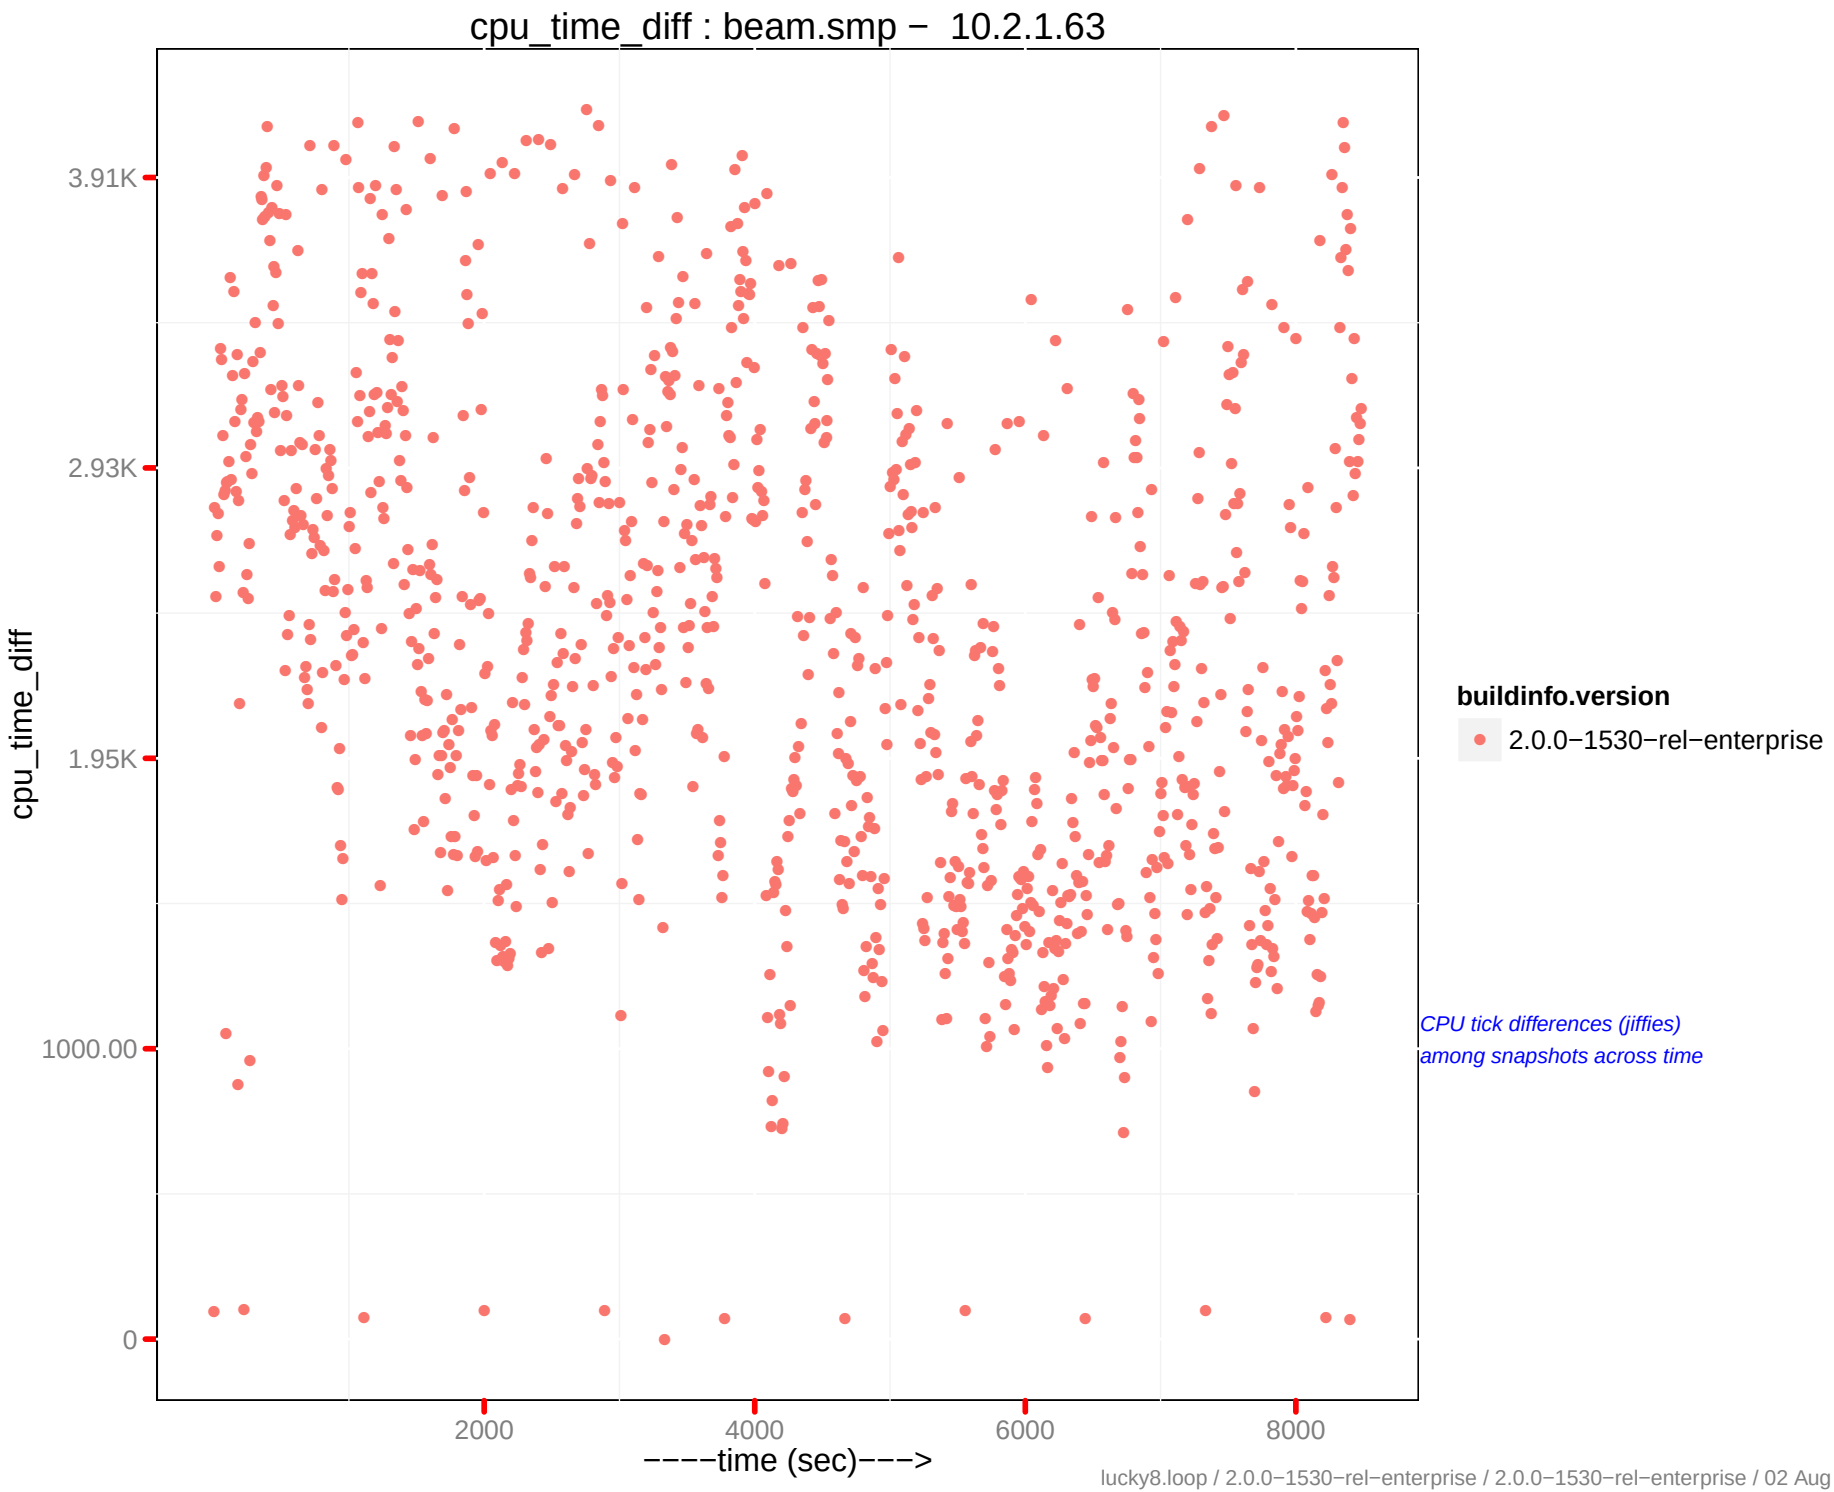
cpu_time_diff : beam.smp − 10.2.1.63
cpu_time_diff
3.91K
2.93K
1.95K
1000.00
0
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

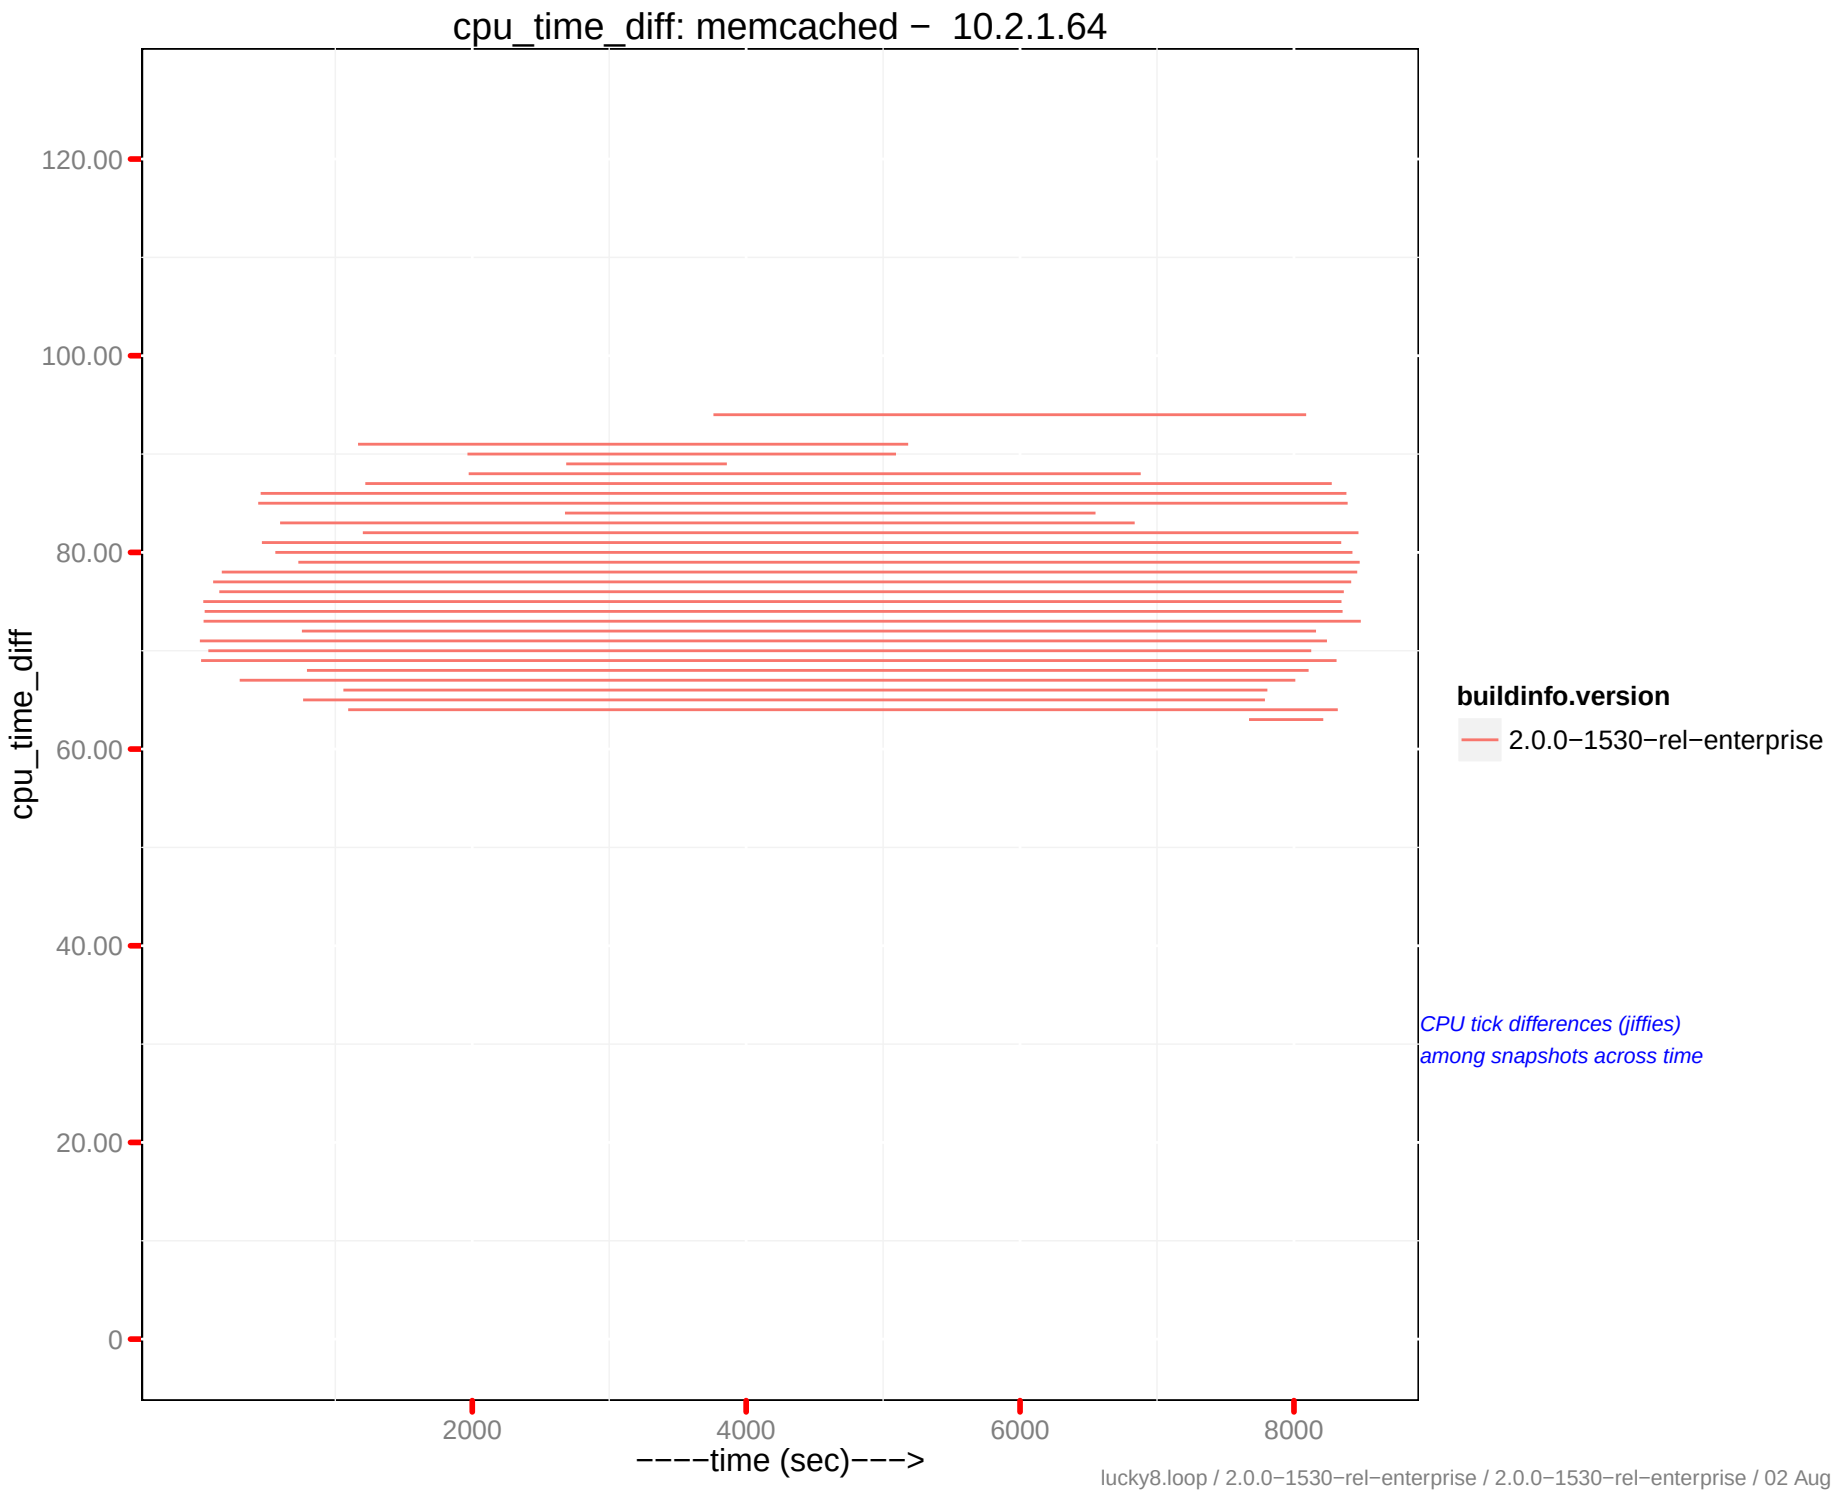
cpu_time_diff: memcached − 10.2.1.64
120.00
100.00
80.00
60.00
40.00
20.00
0
cpu_time_diff
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

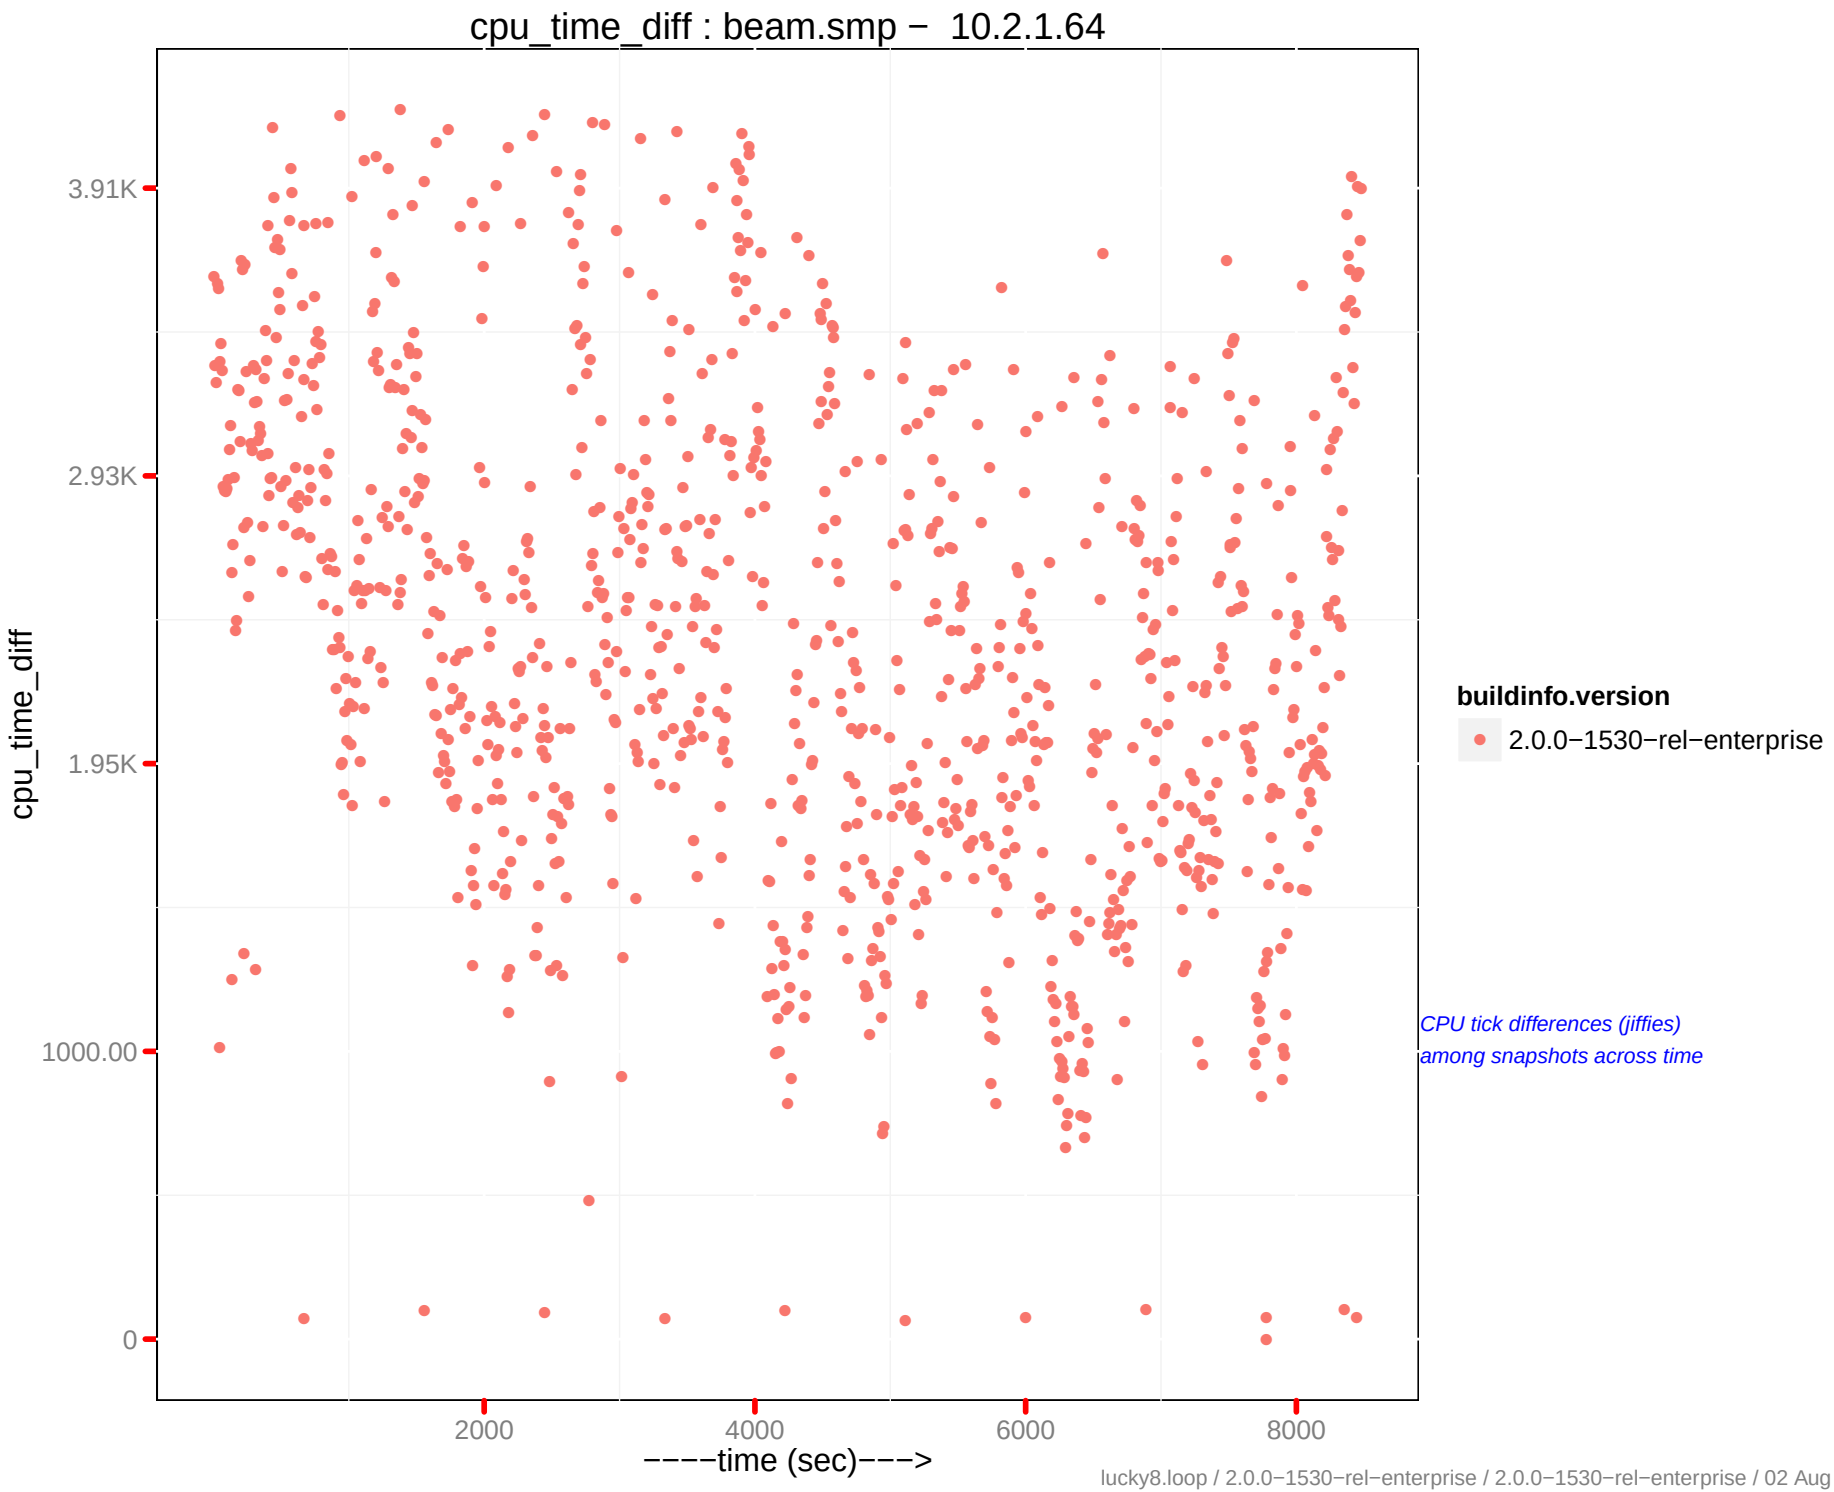
cpu_time_diff : beam.smp − 10.2.1.64
cpu_time_diff
3.91K
2.93K
1.95K
1000.00
0
2000
4000
6000
8000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

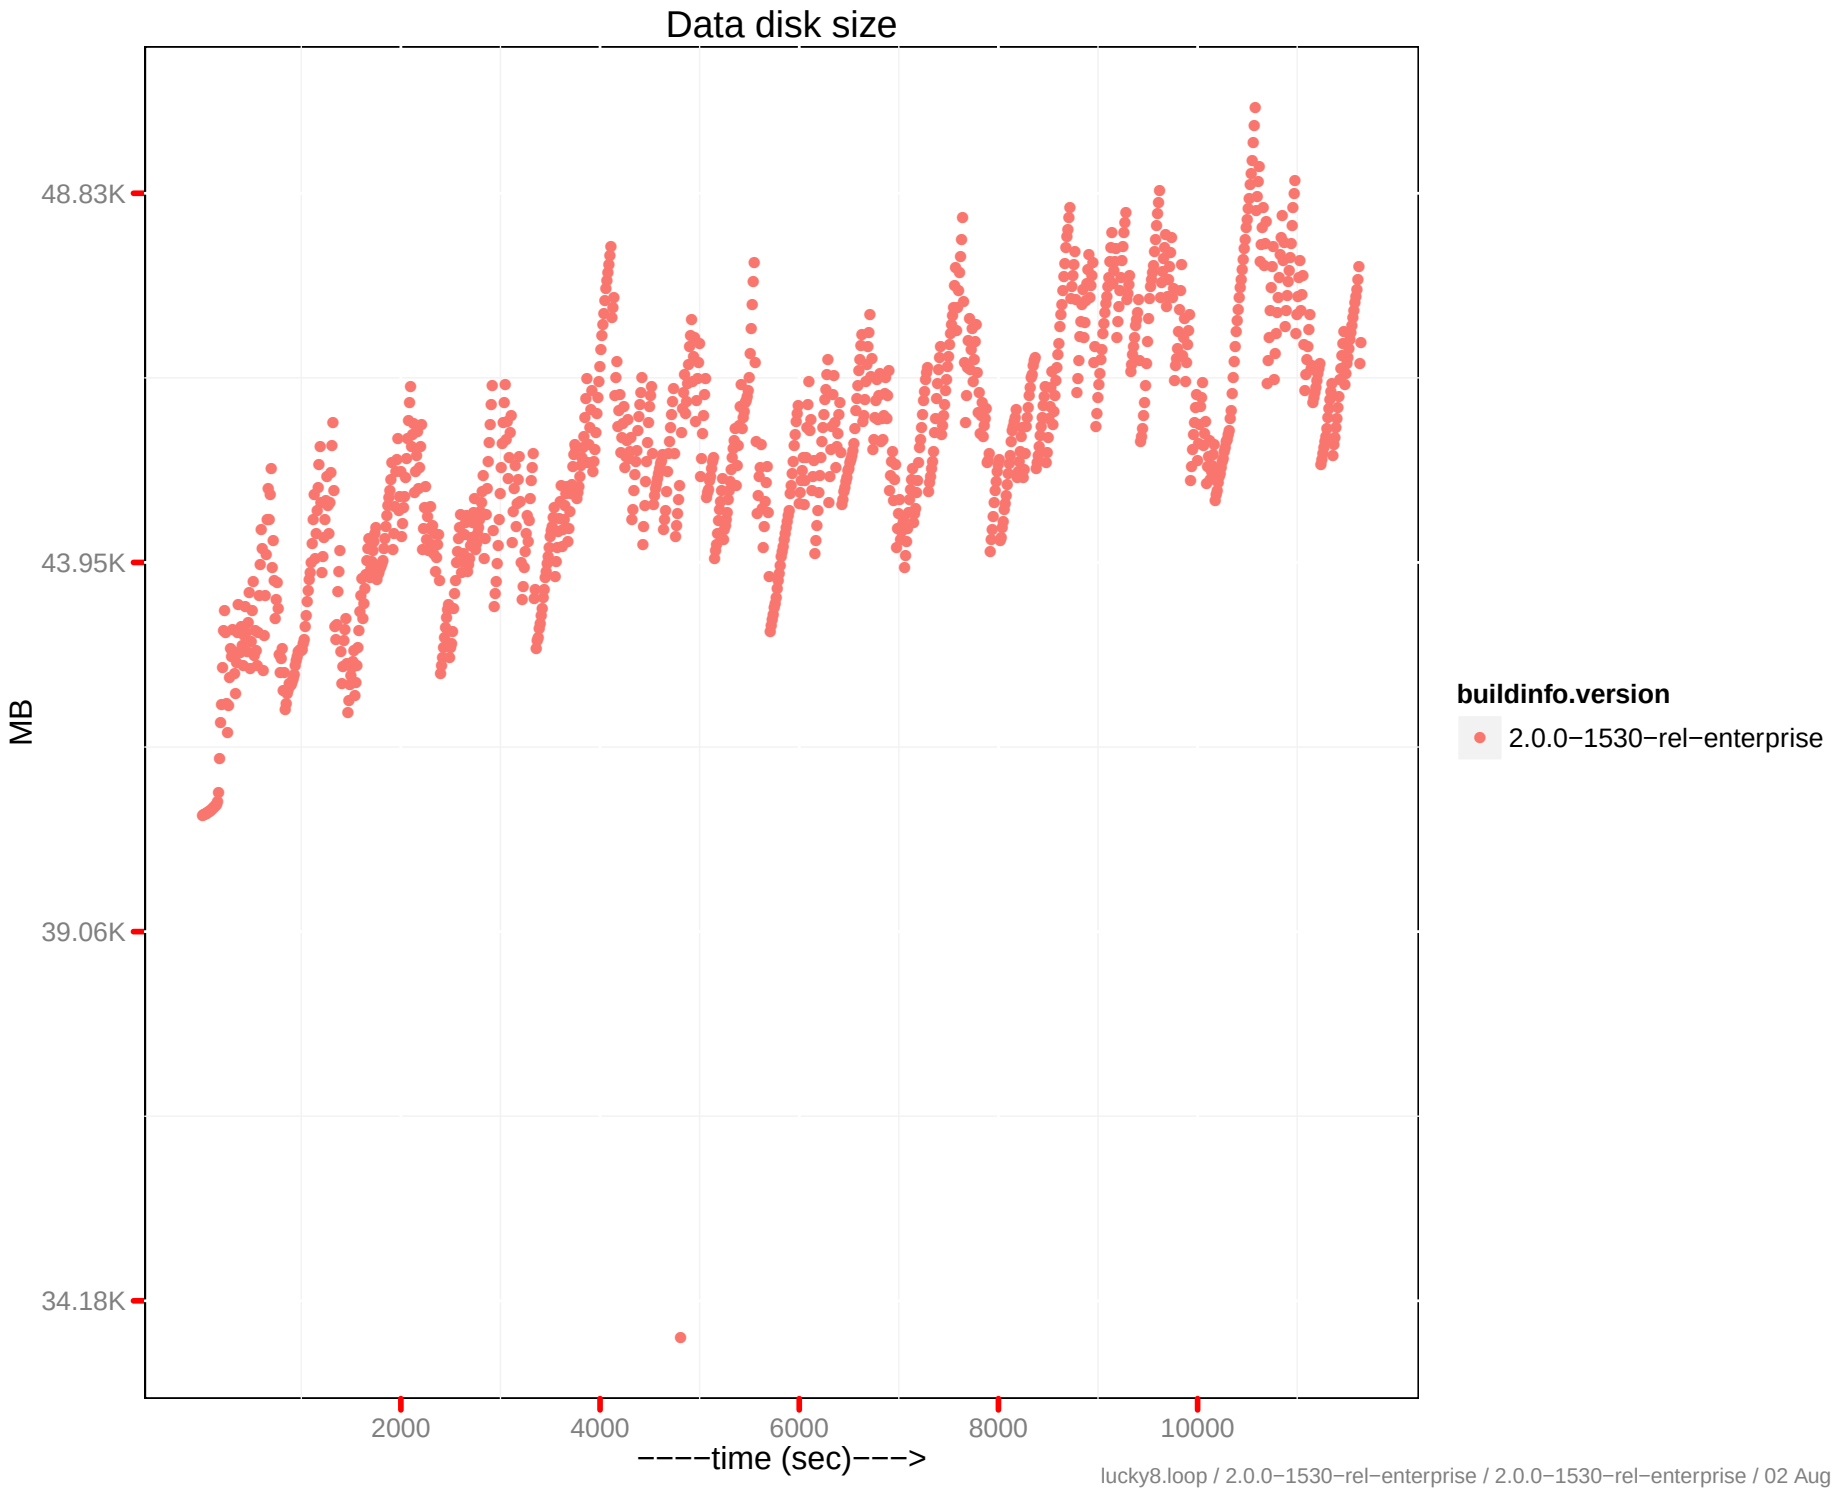
Data disk size
48.83K
43.95K
39.06K
34.18K
MB
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1530-rel-enterprise
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

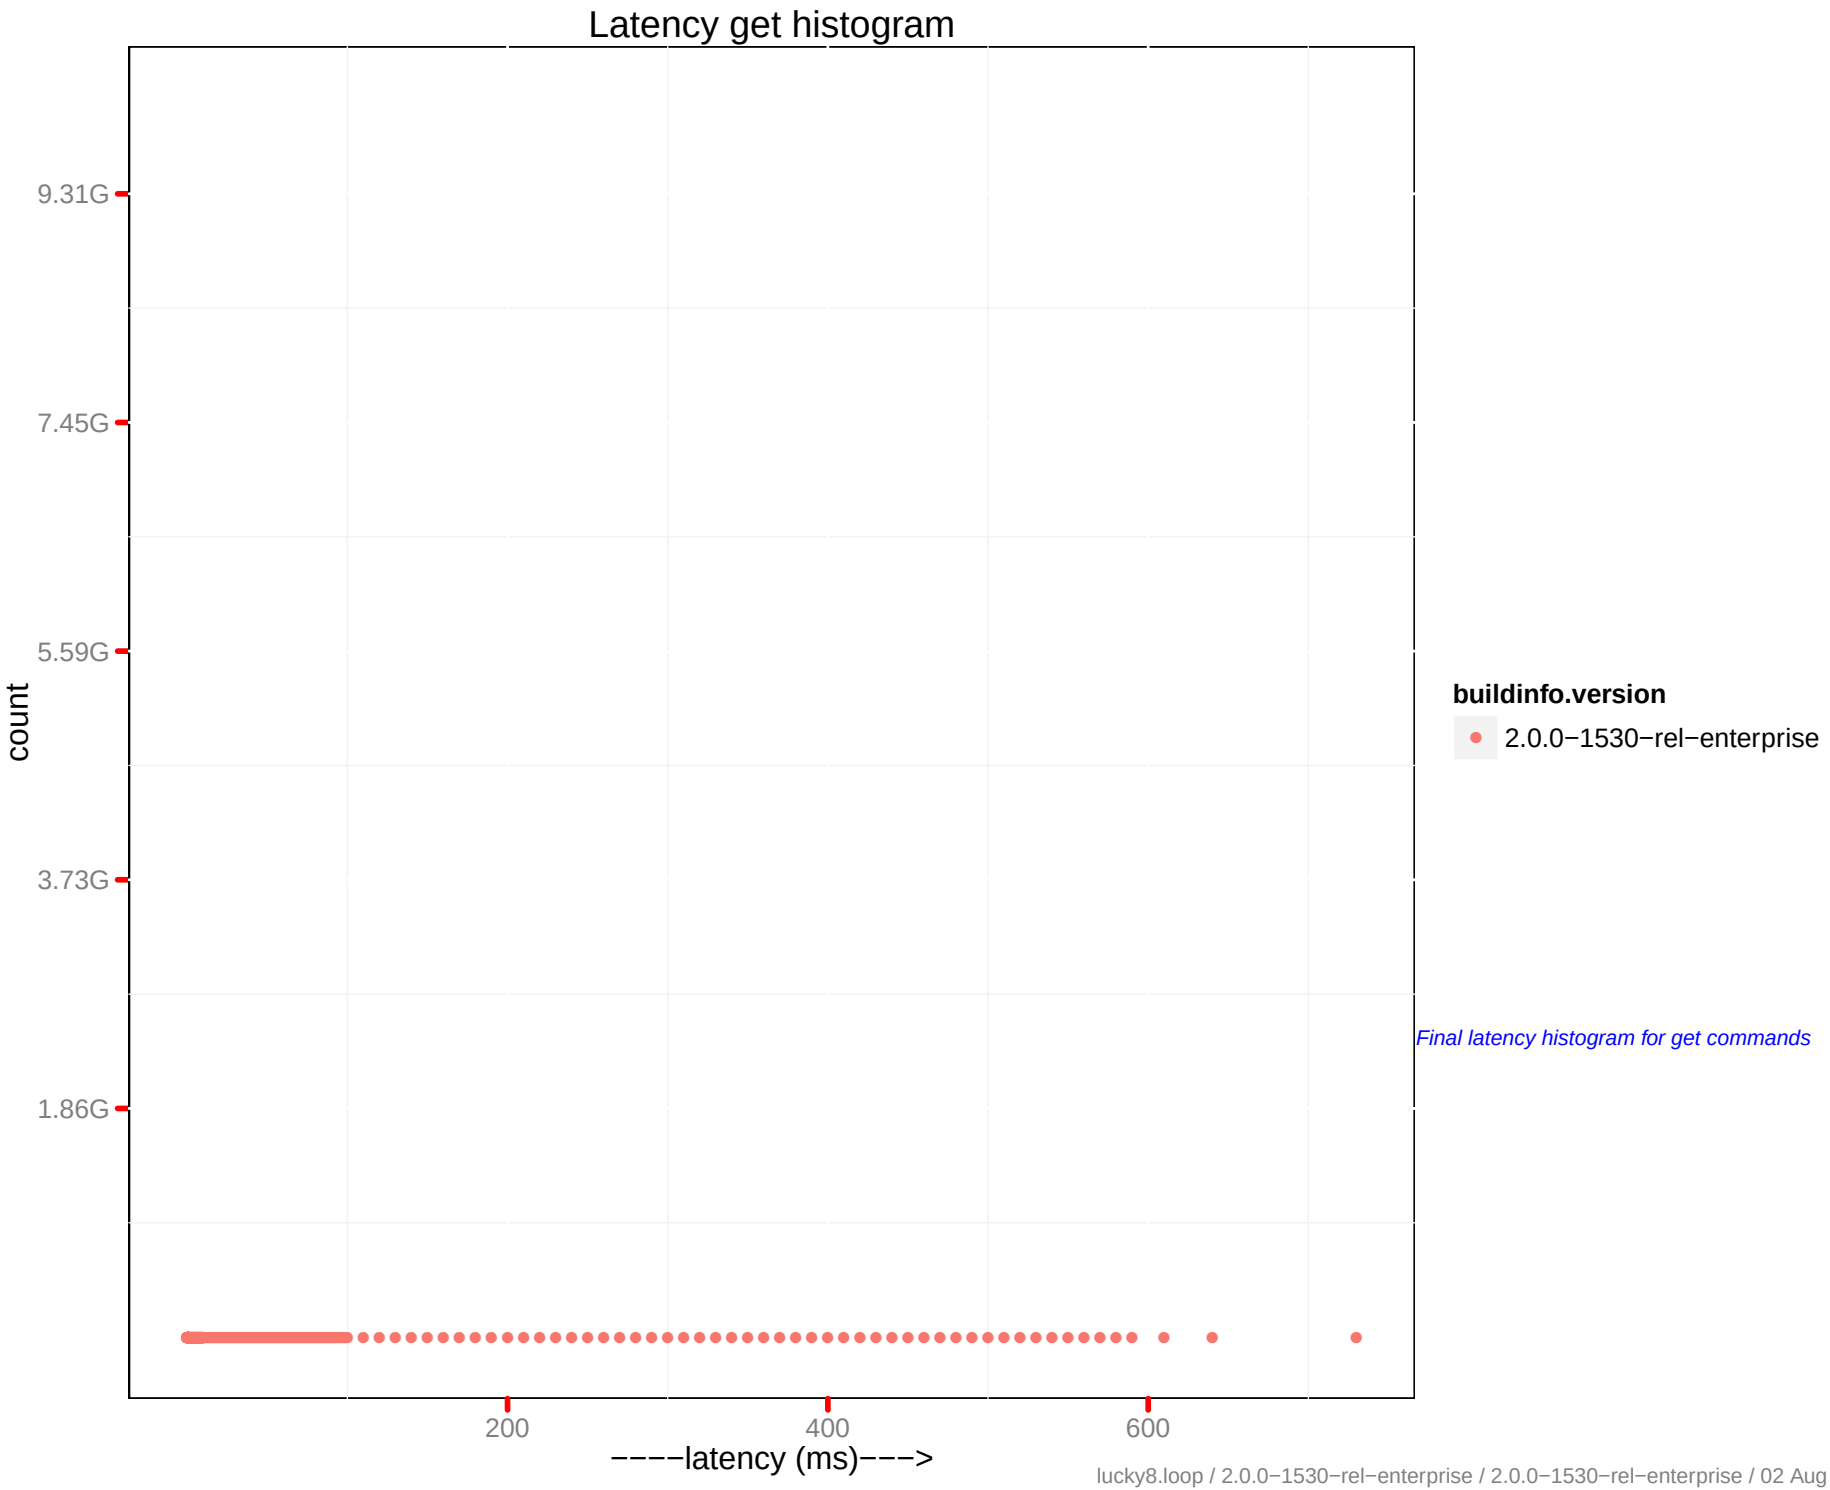
Latency get histogram
9.31G
7.45G
5.59G
3.73G
1.86G
count
200
400
600
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
Final latency histogram for get commands
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

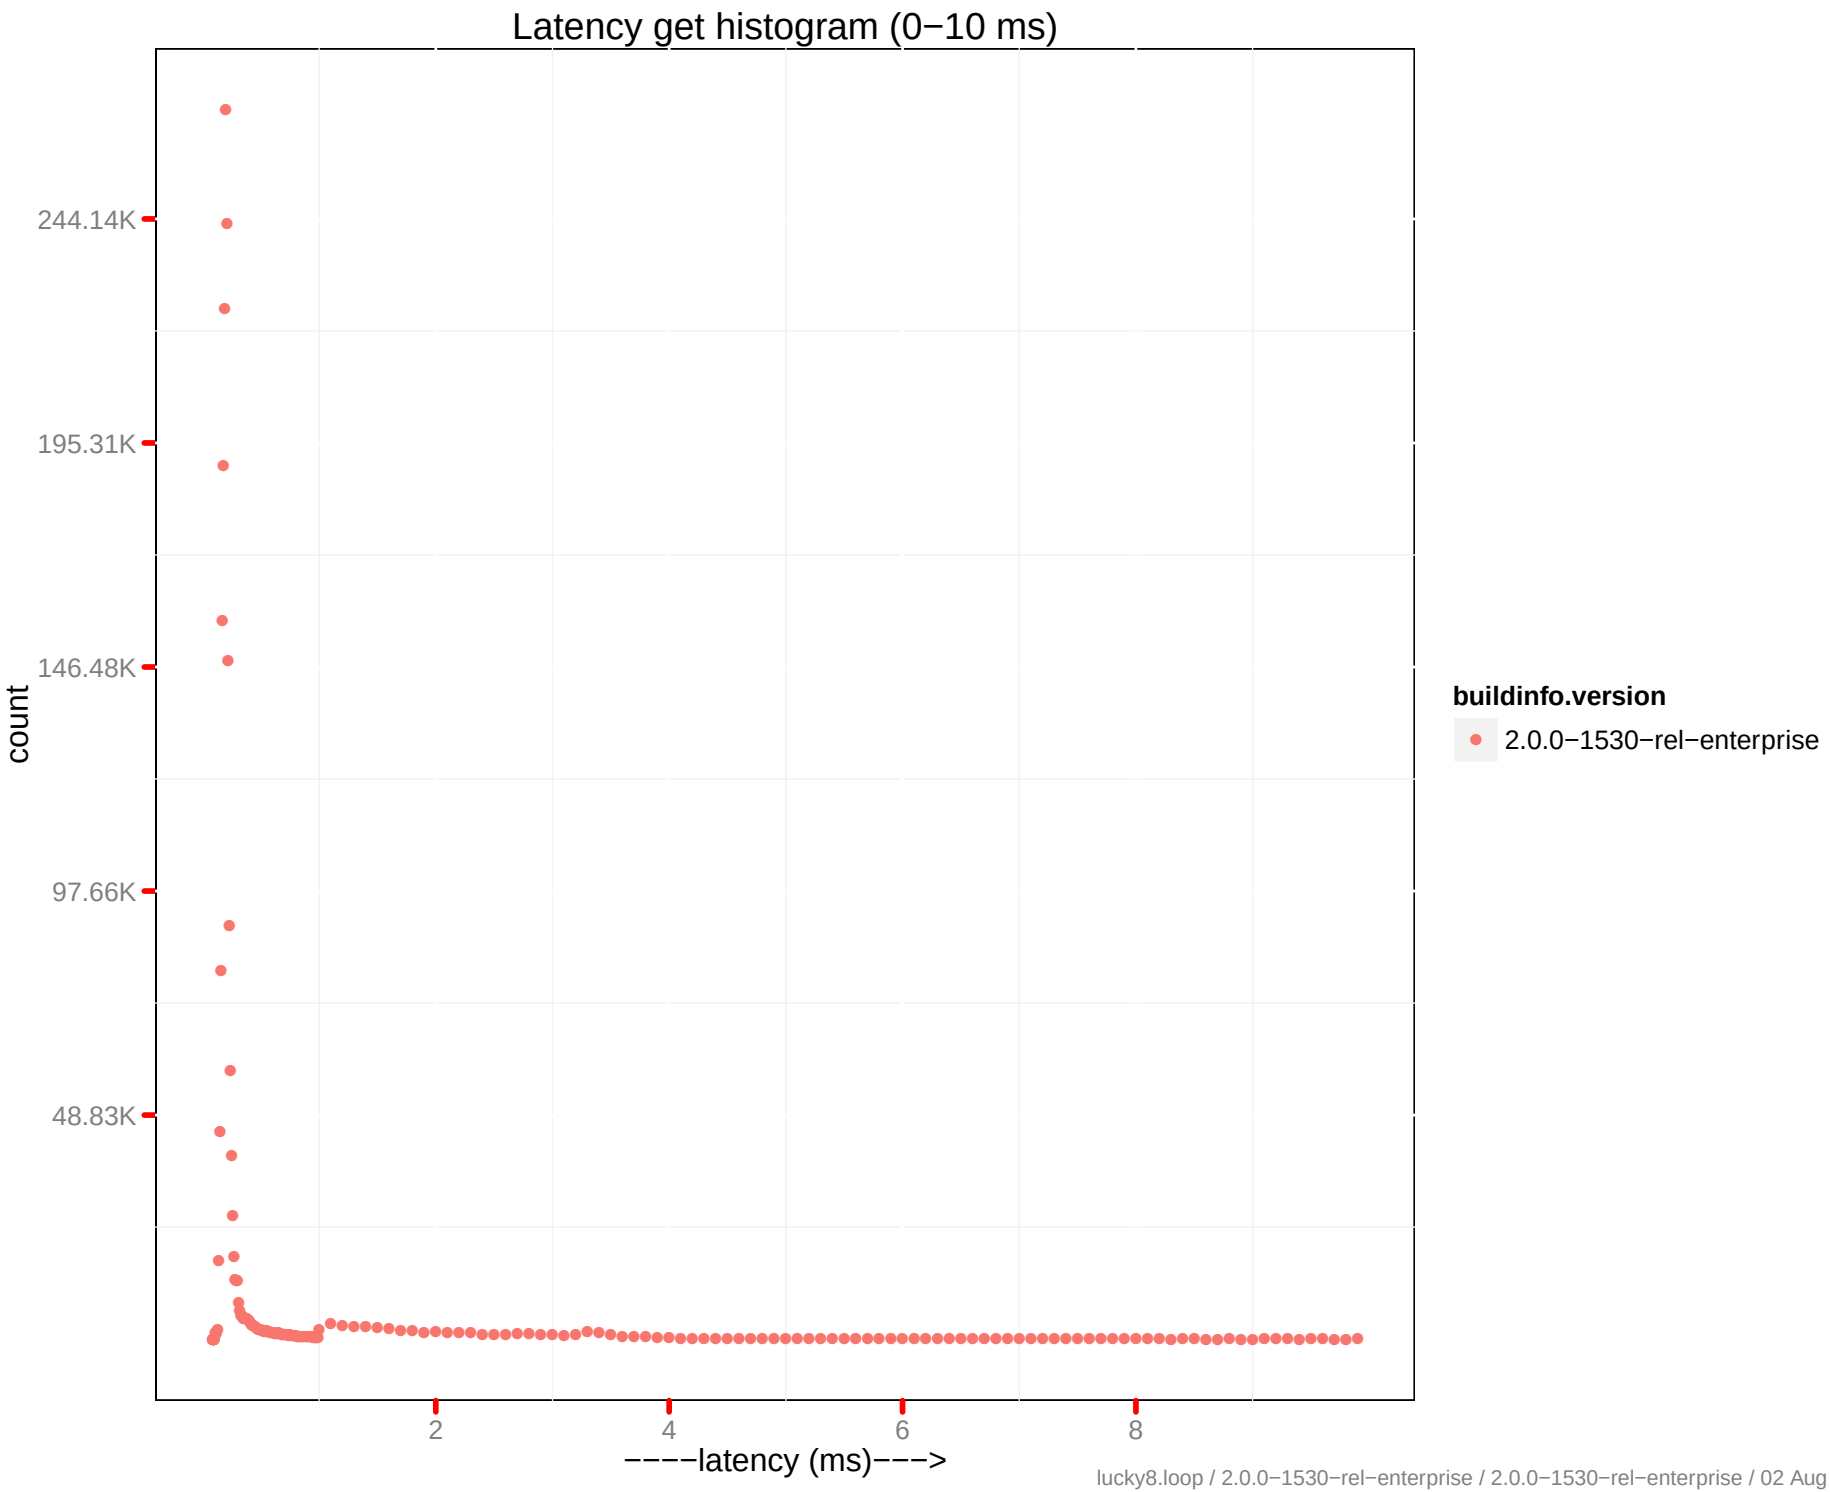
Latency get histogram (0−10 ms)
244.14K
195.31K
146.48K
97.66K
48.83K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

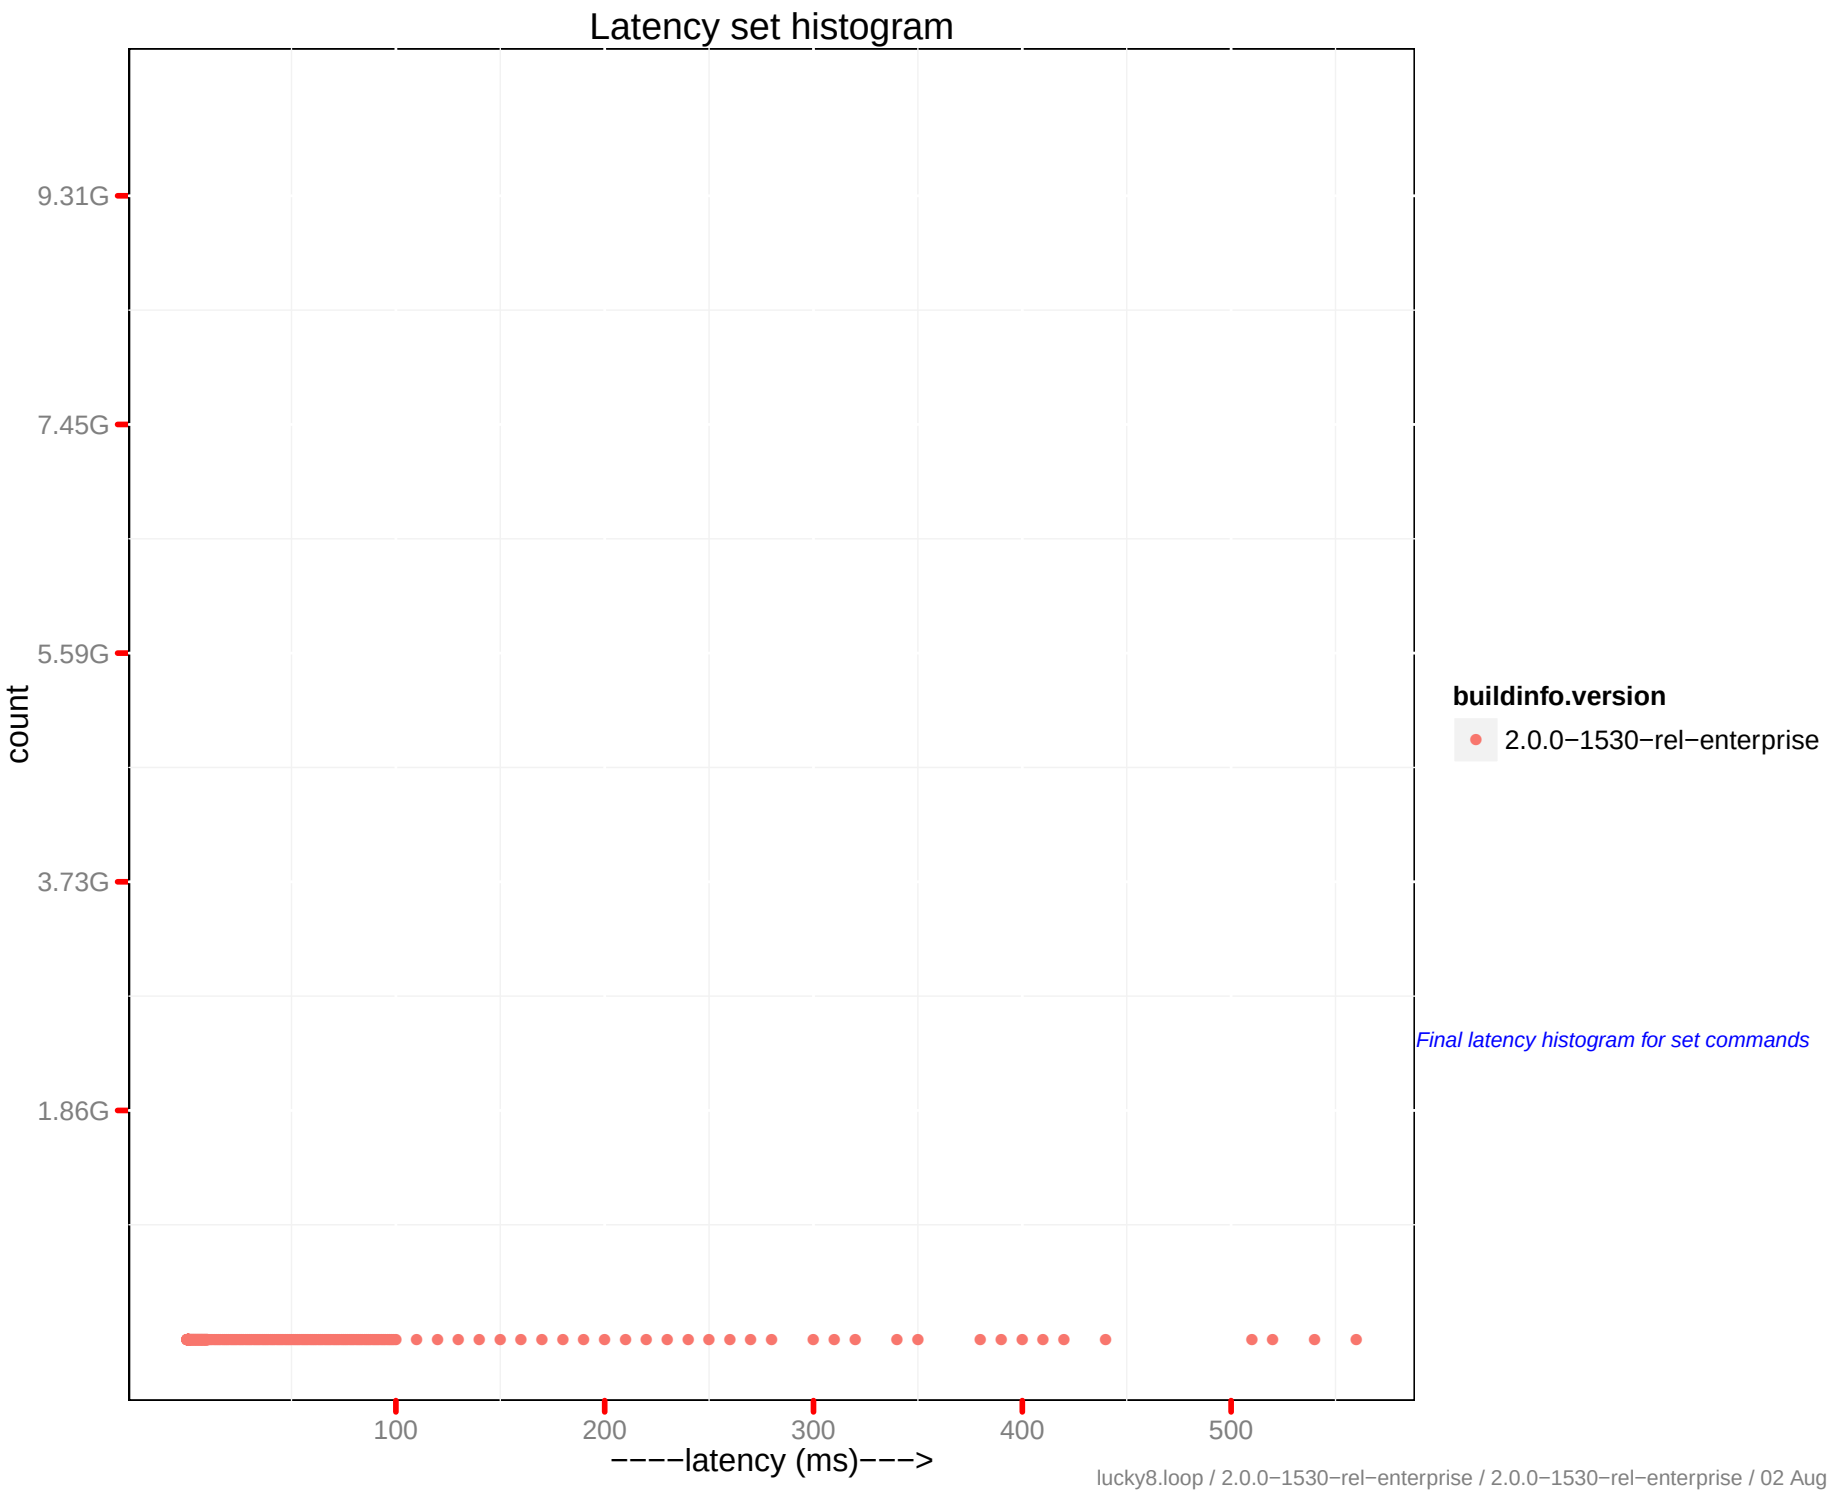
Latency set histogram
9.31G
7.45G
5.59G
3.73G
1.86G
count
100
200
300
400
500
----latency (ms)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Final latency histogram for set commands
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

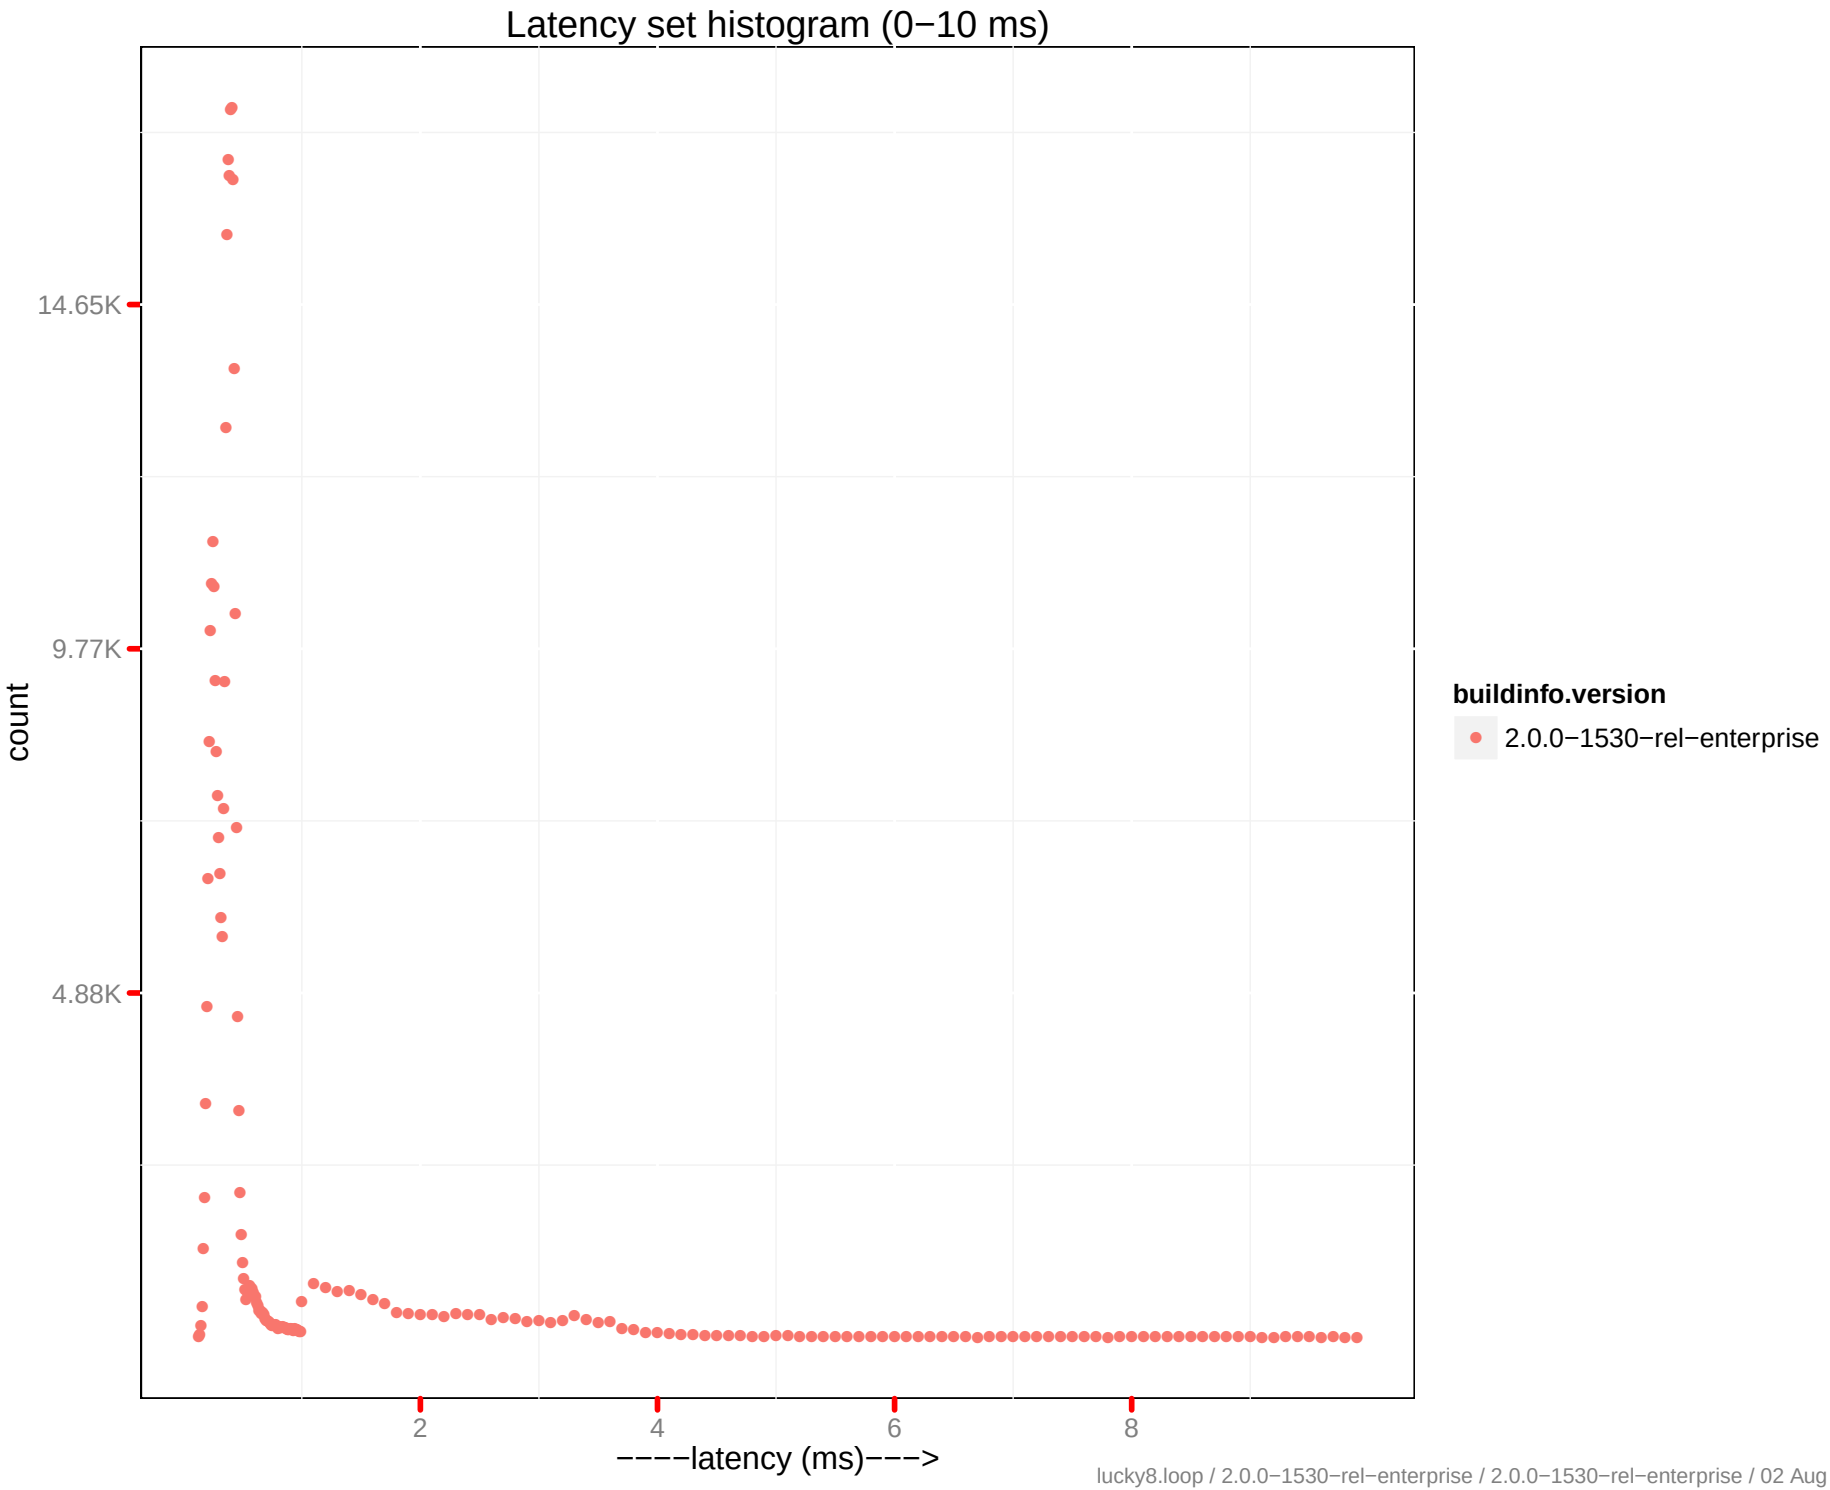
Latency set histogram (0−10 ms)
14.65K
9.77K
4.88K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

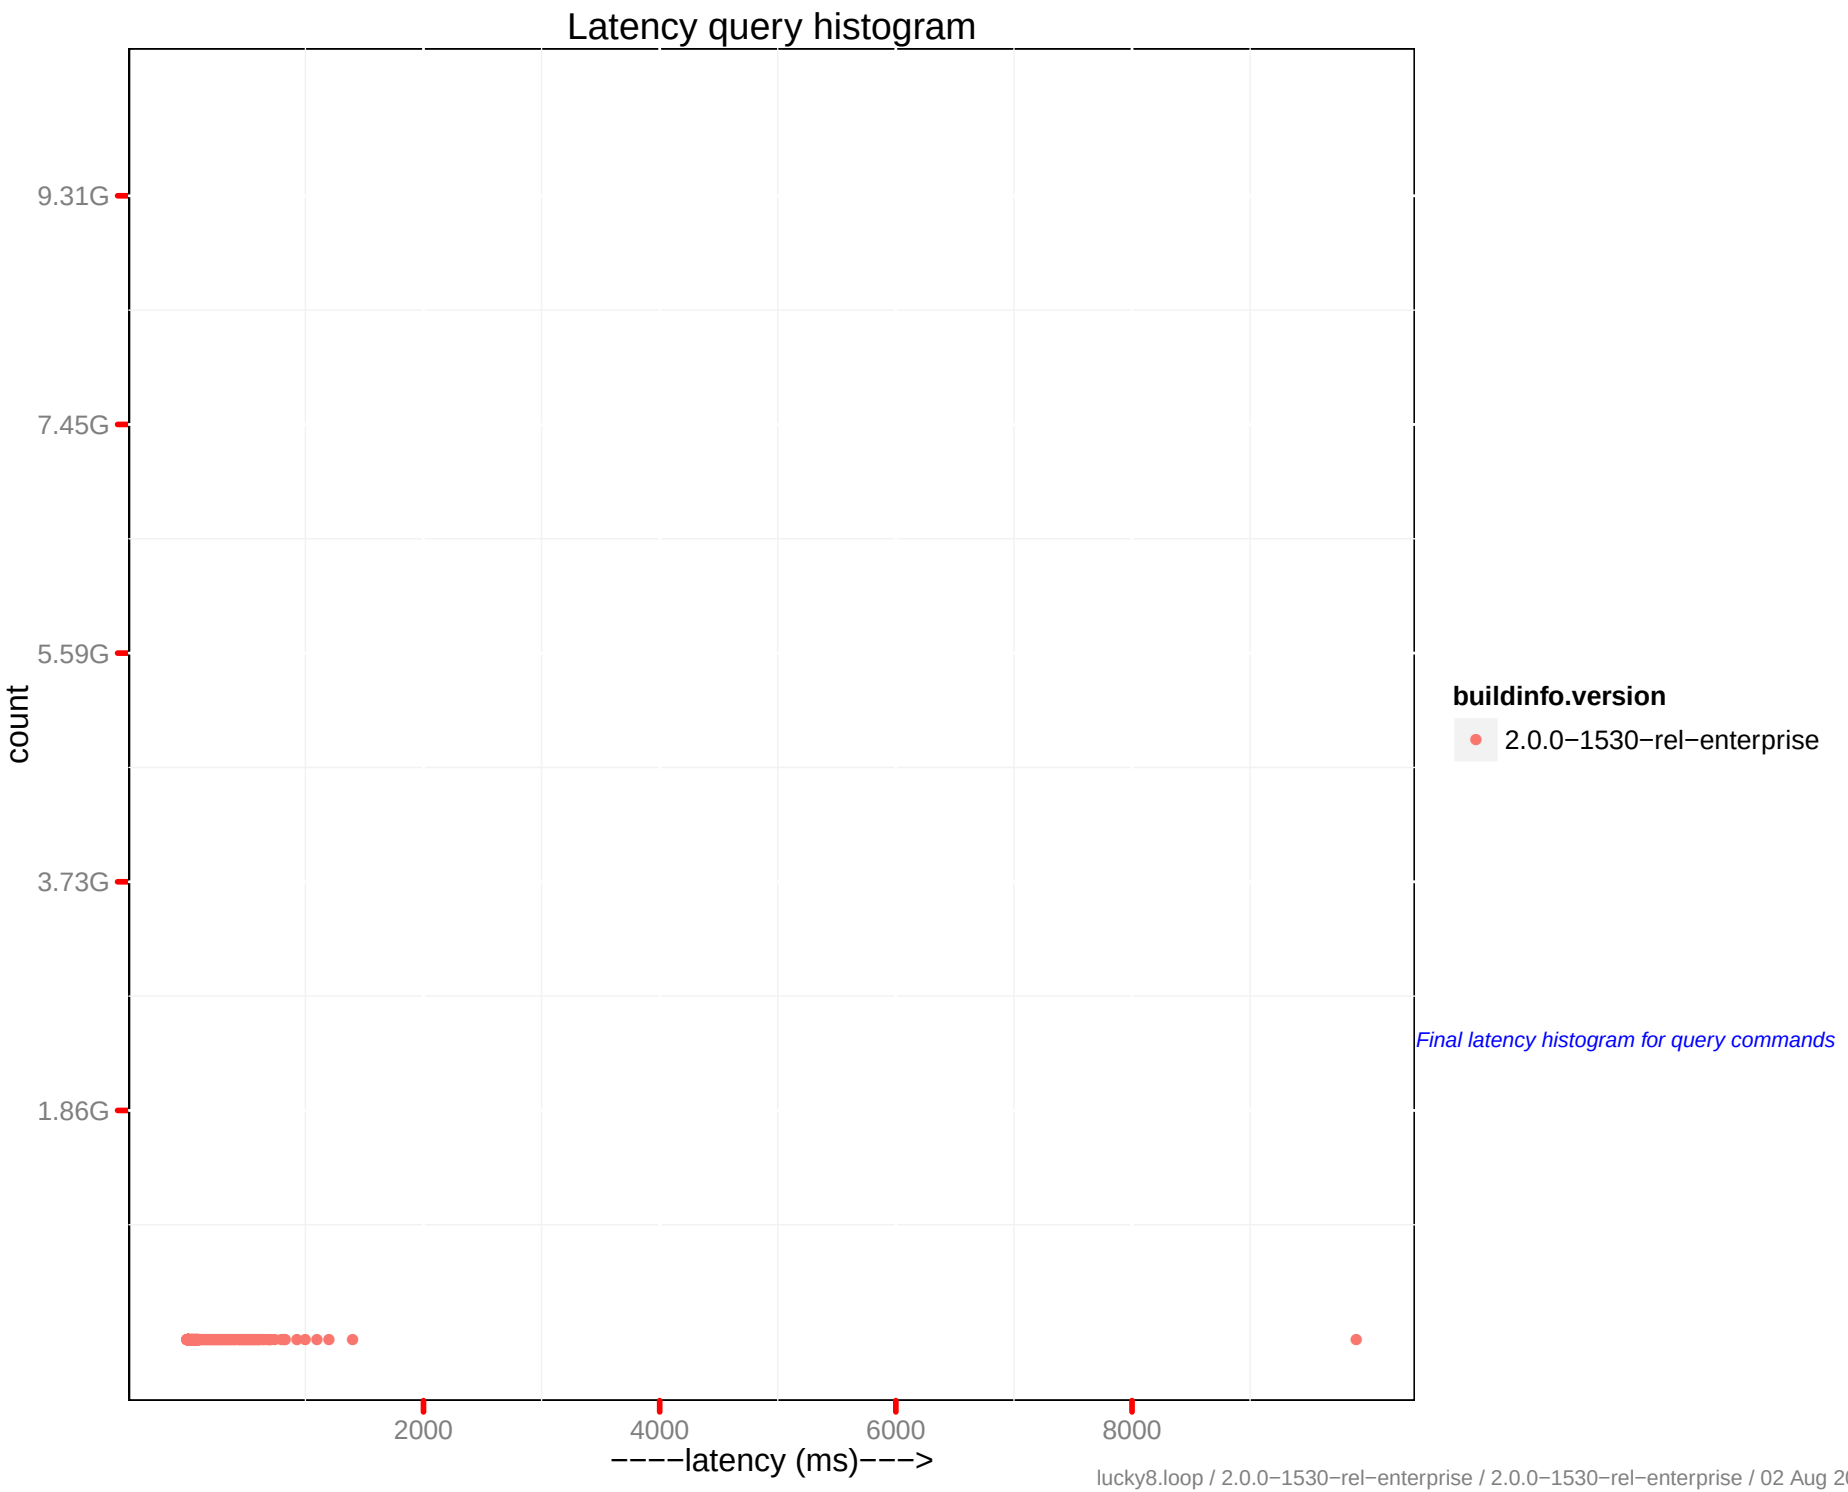
Latency query histogram
9.31G
7.45G
5.59G
3.73G
1.86G
count
2000
4000
6000
8000
----latency (ms)--->
buildinfo.version
2.0.0-1530-rel-enterprise
Final latency histogram for query commands
lucky8.loop / 2.0.0-1530-rel-enterprise / 2.0.0-1530-rel-enterprise / 02 Aug 2012

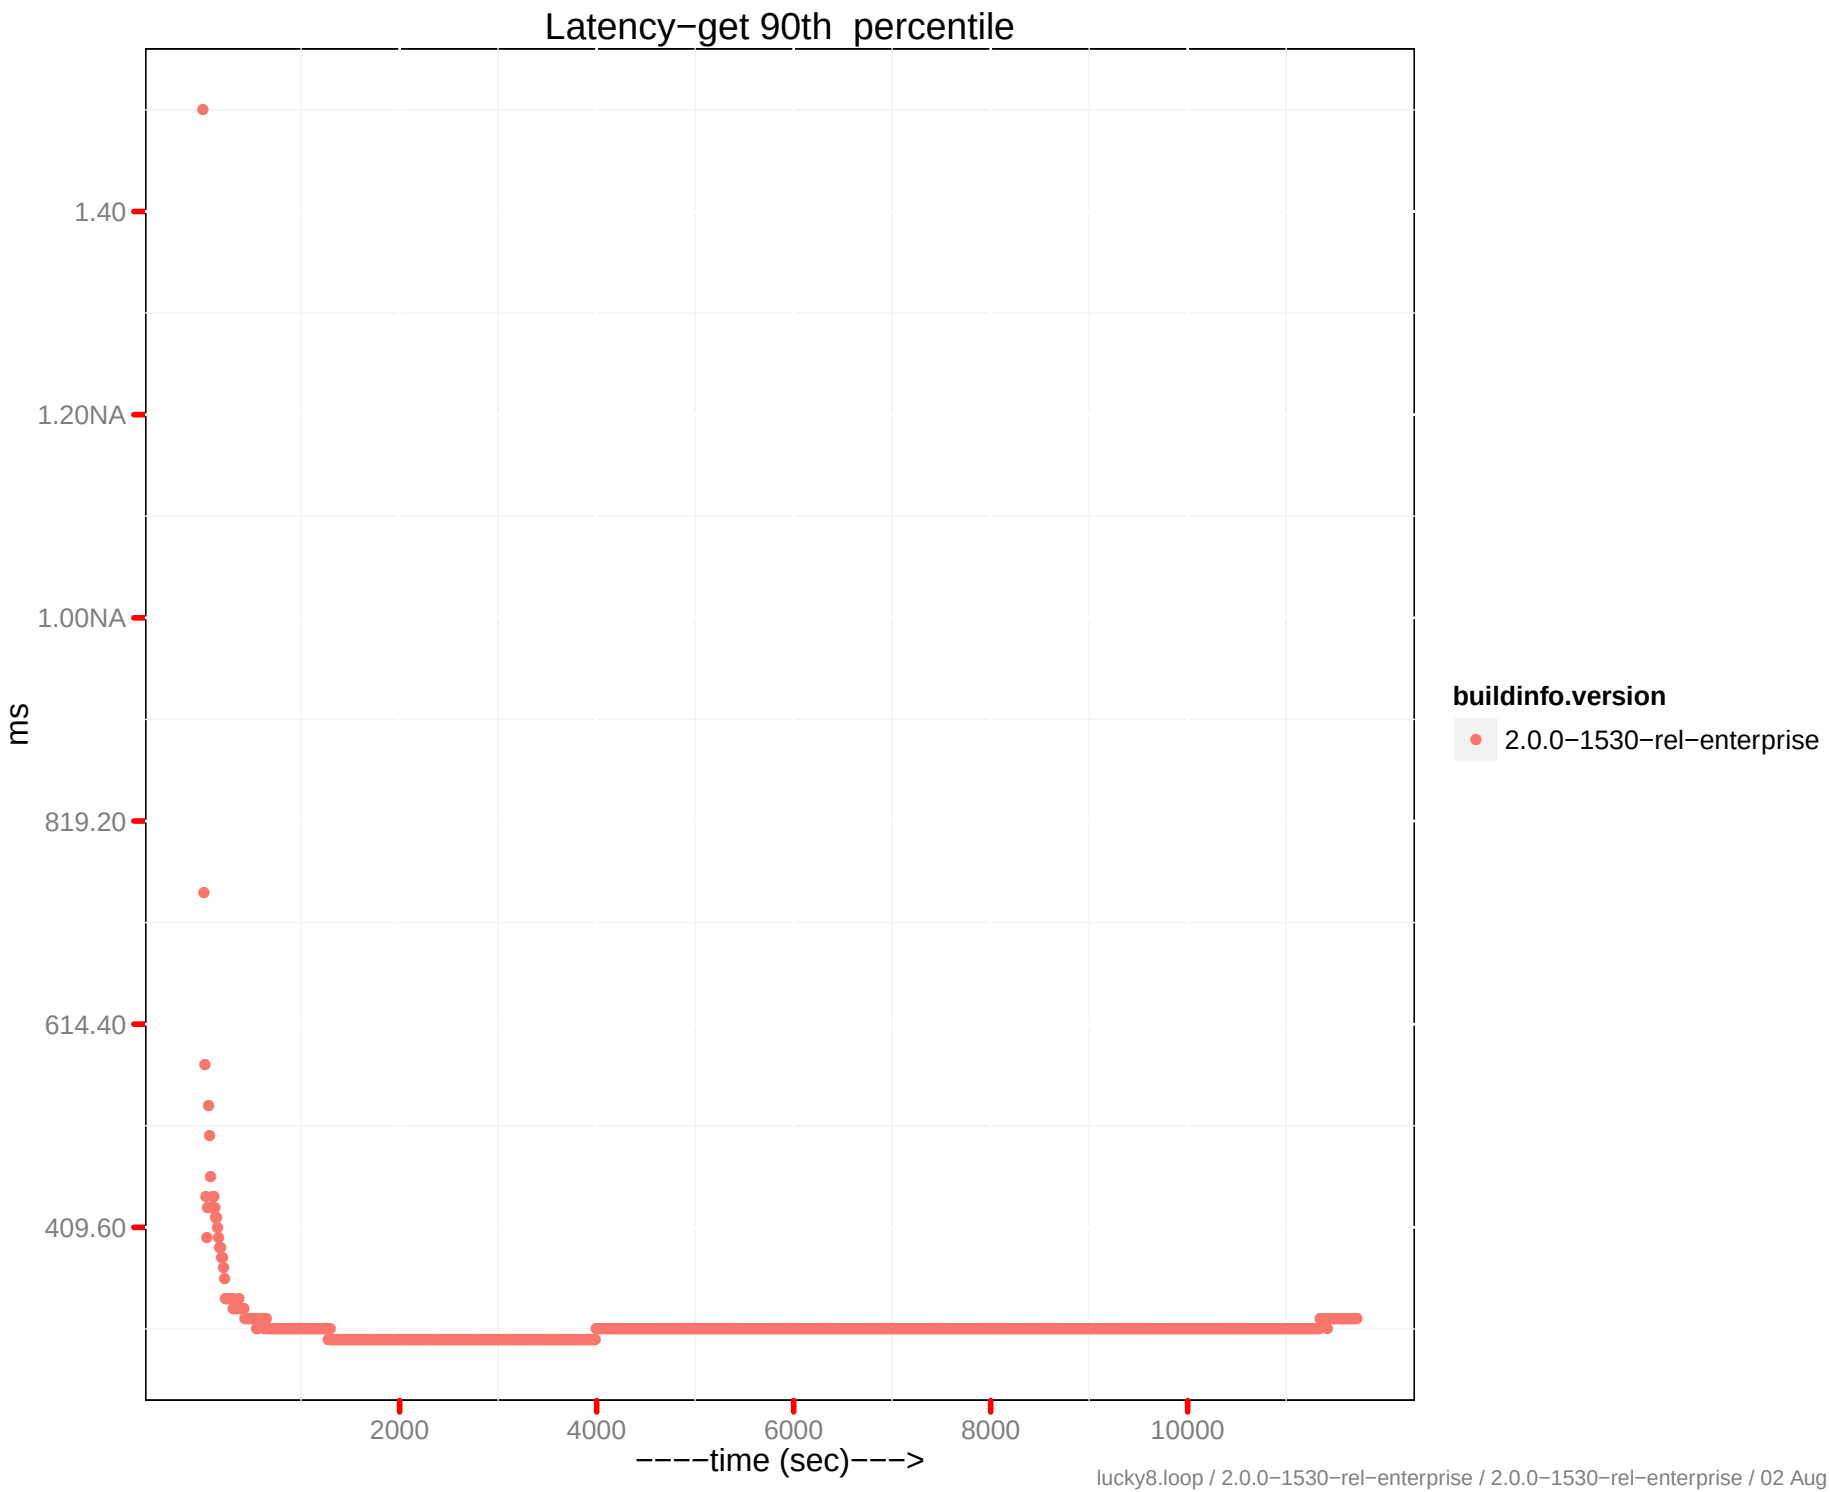
Latency−get 90th percentile
1.40
1.20NA
1.00NA
819.20
614.40
409.60
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

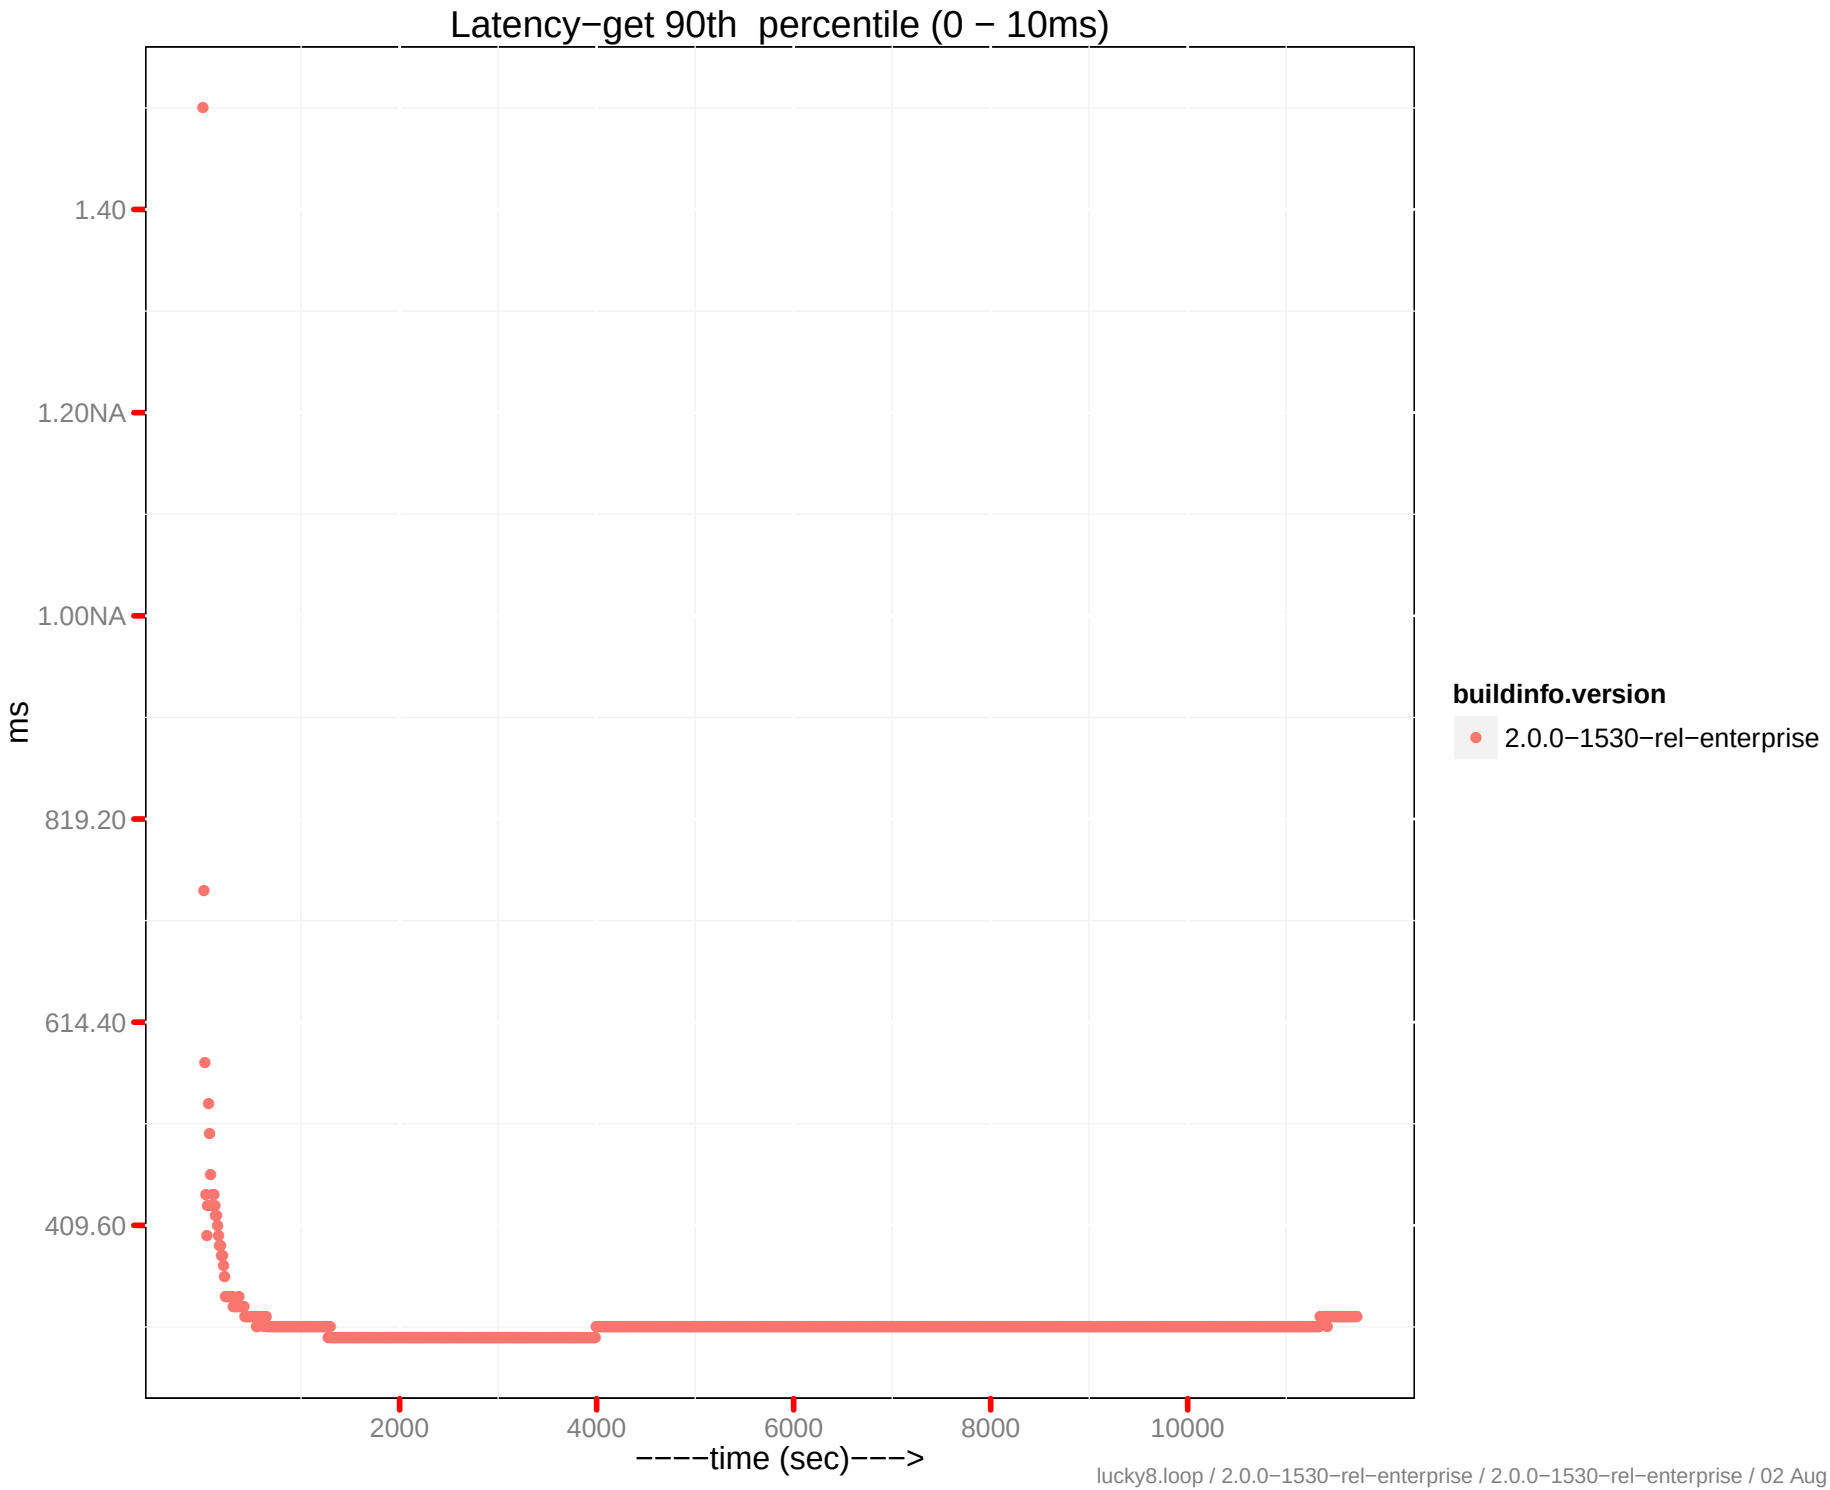
Latency−get 90th percentile (0 − 10ms)
1.40
1.20NA
1.00NA
819.20
614.40
409.60
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

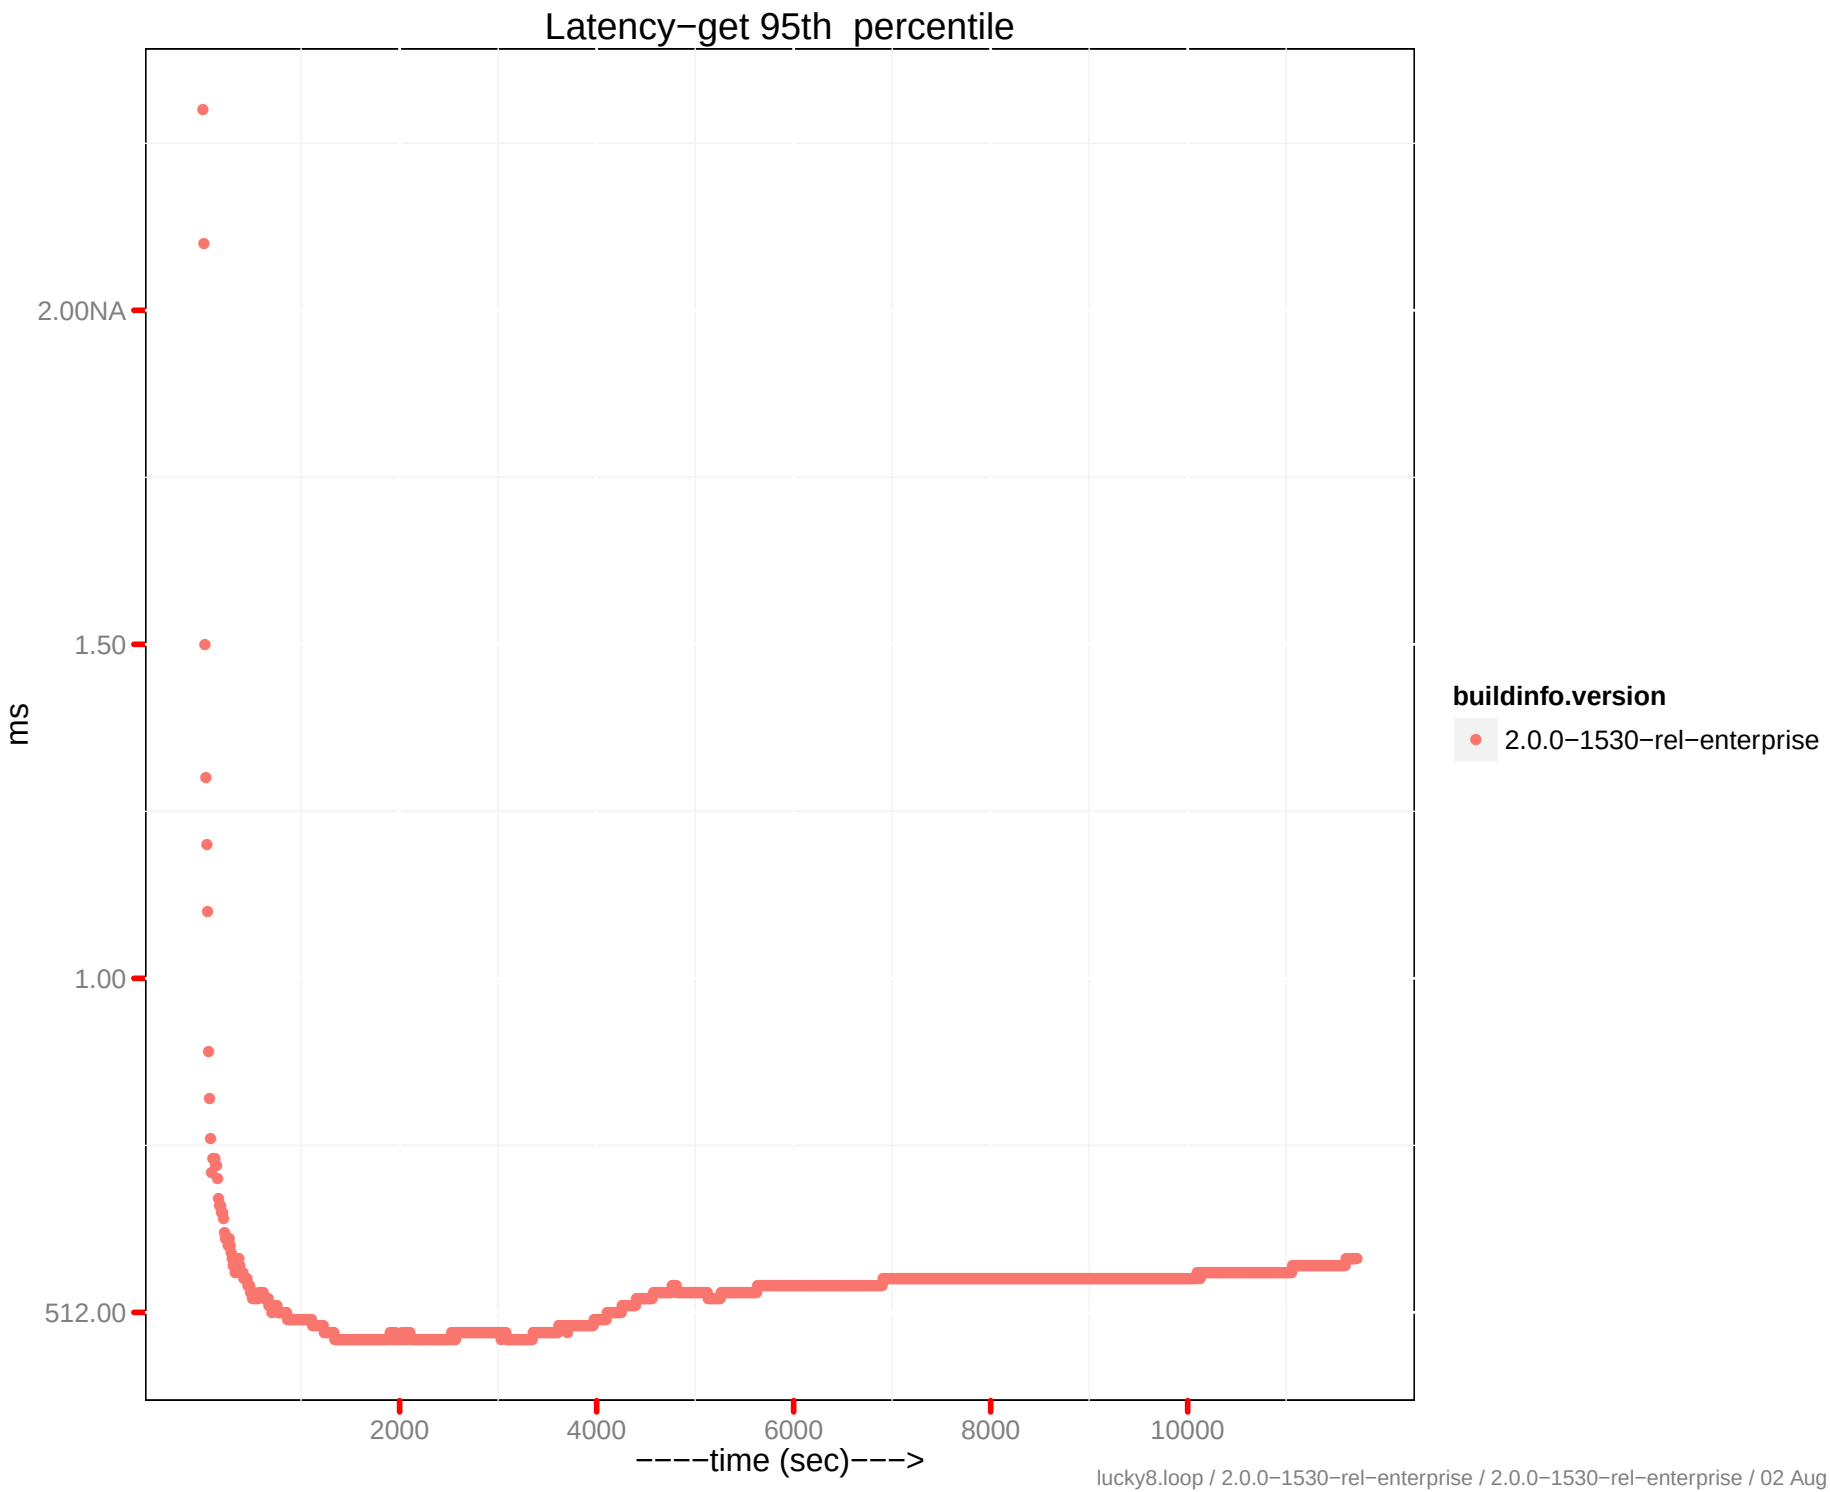
Latency−get 95th percentile
2.00NA
1.50
1.00
512.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

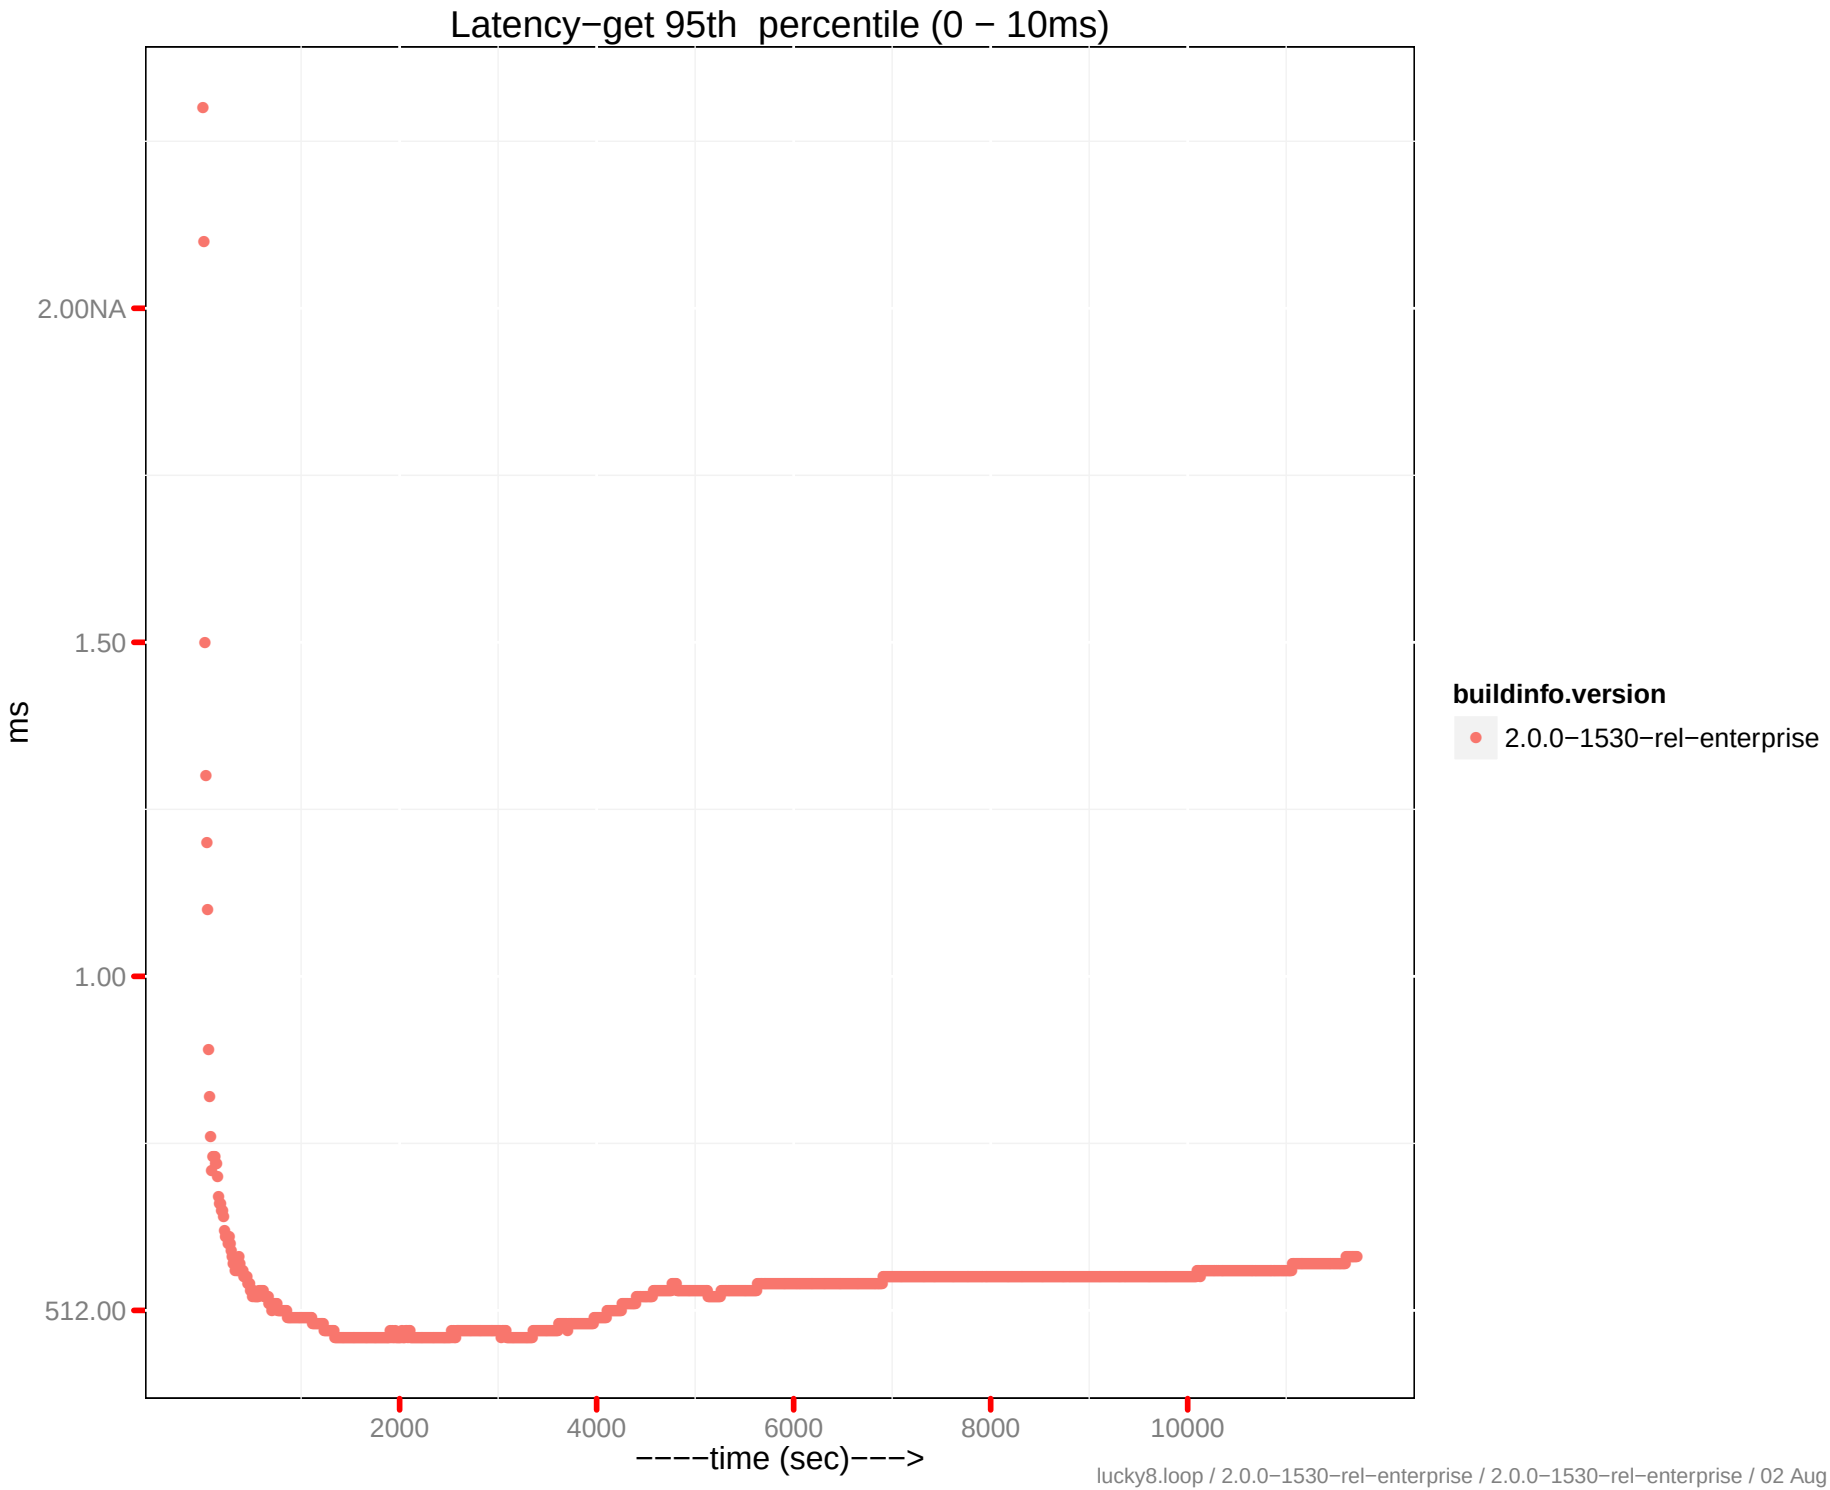
Latency−get 95th percentile (0 − 10ms)
2.00NA
1.50
1.00
512.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

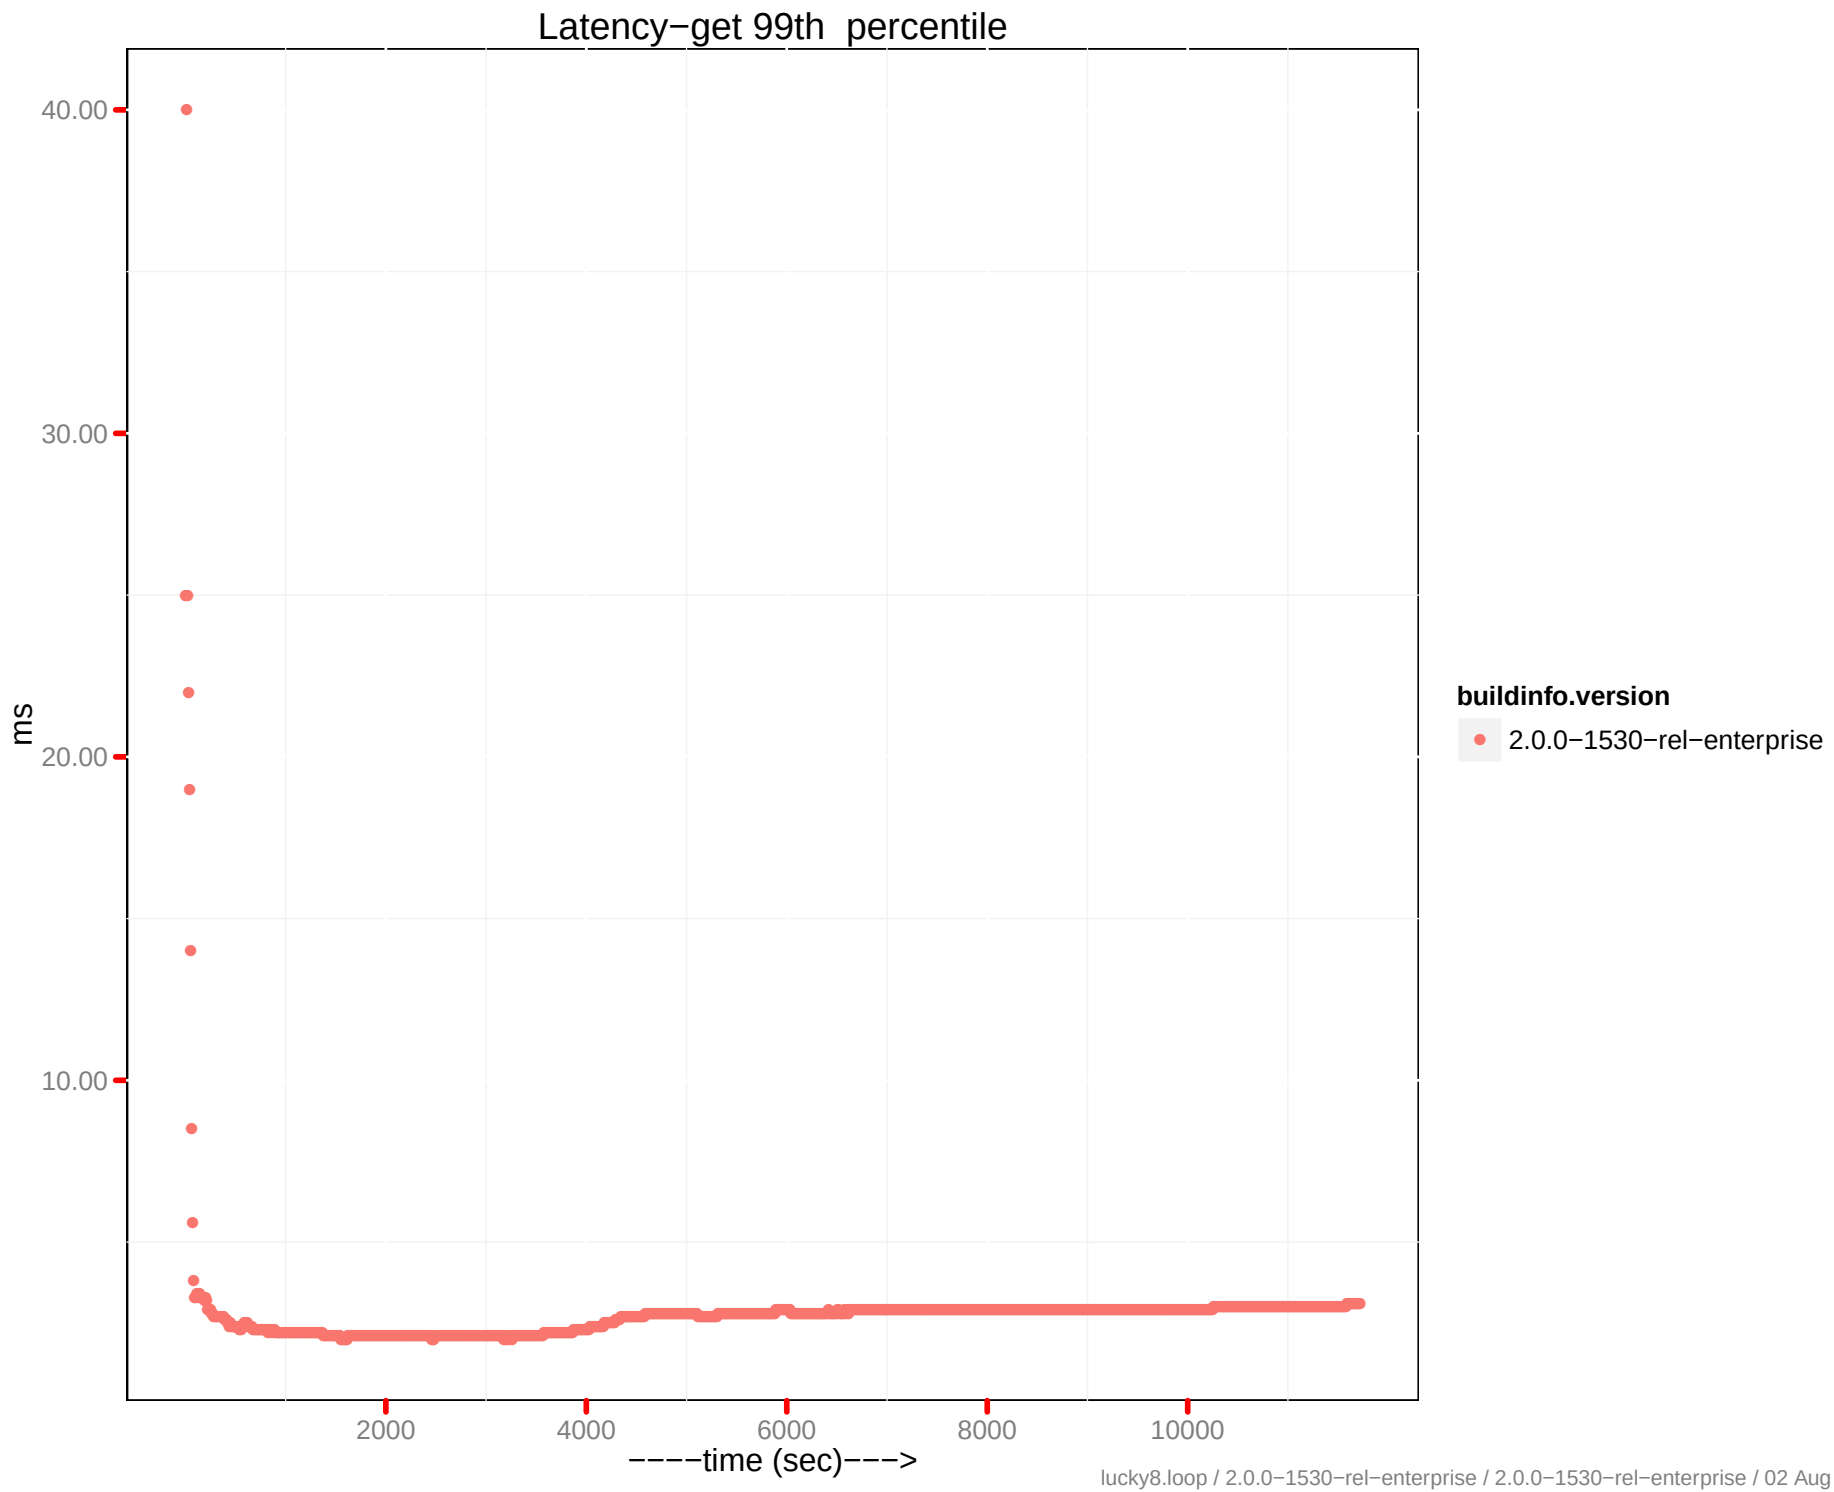
Latency−get 99th percentile
40.00
30.00
20.00
10.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

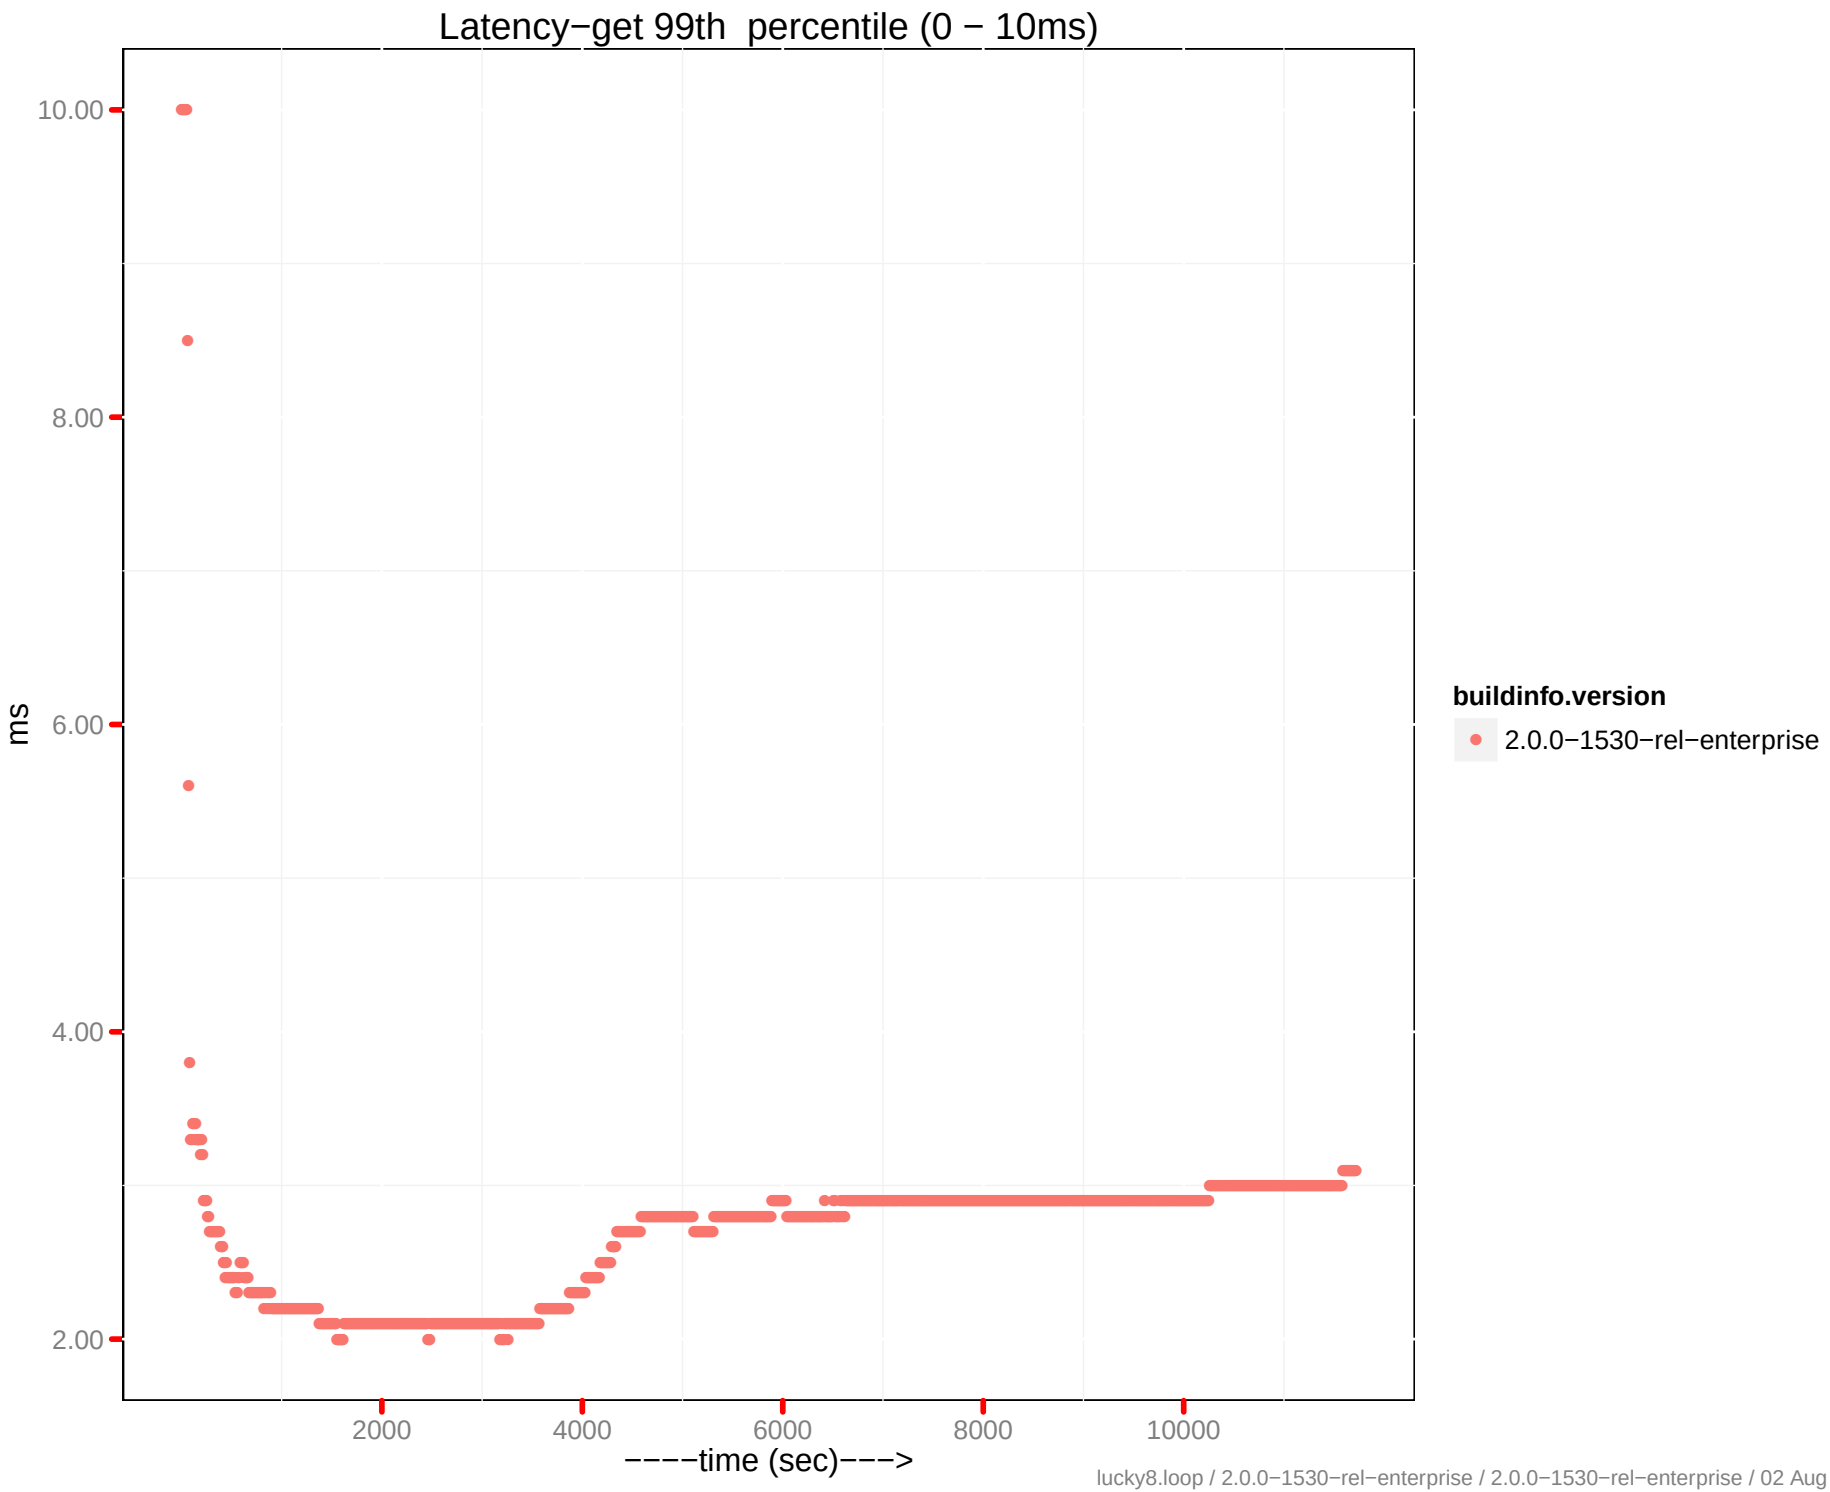
Latency−get 99th percentile (0 − 10ms)
10.00
8.00
6.00
4.00
2.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

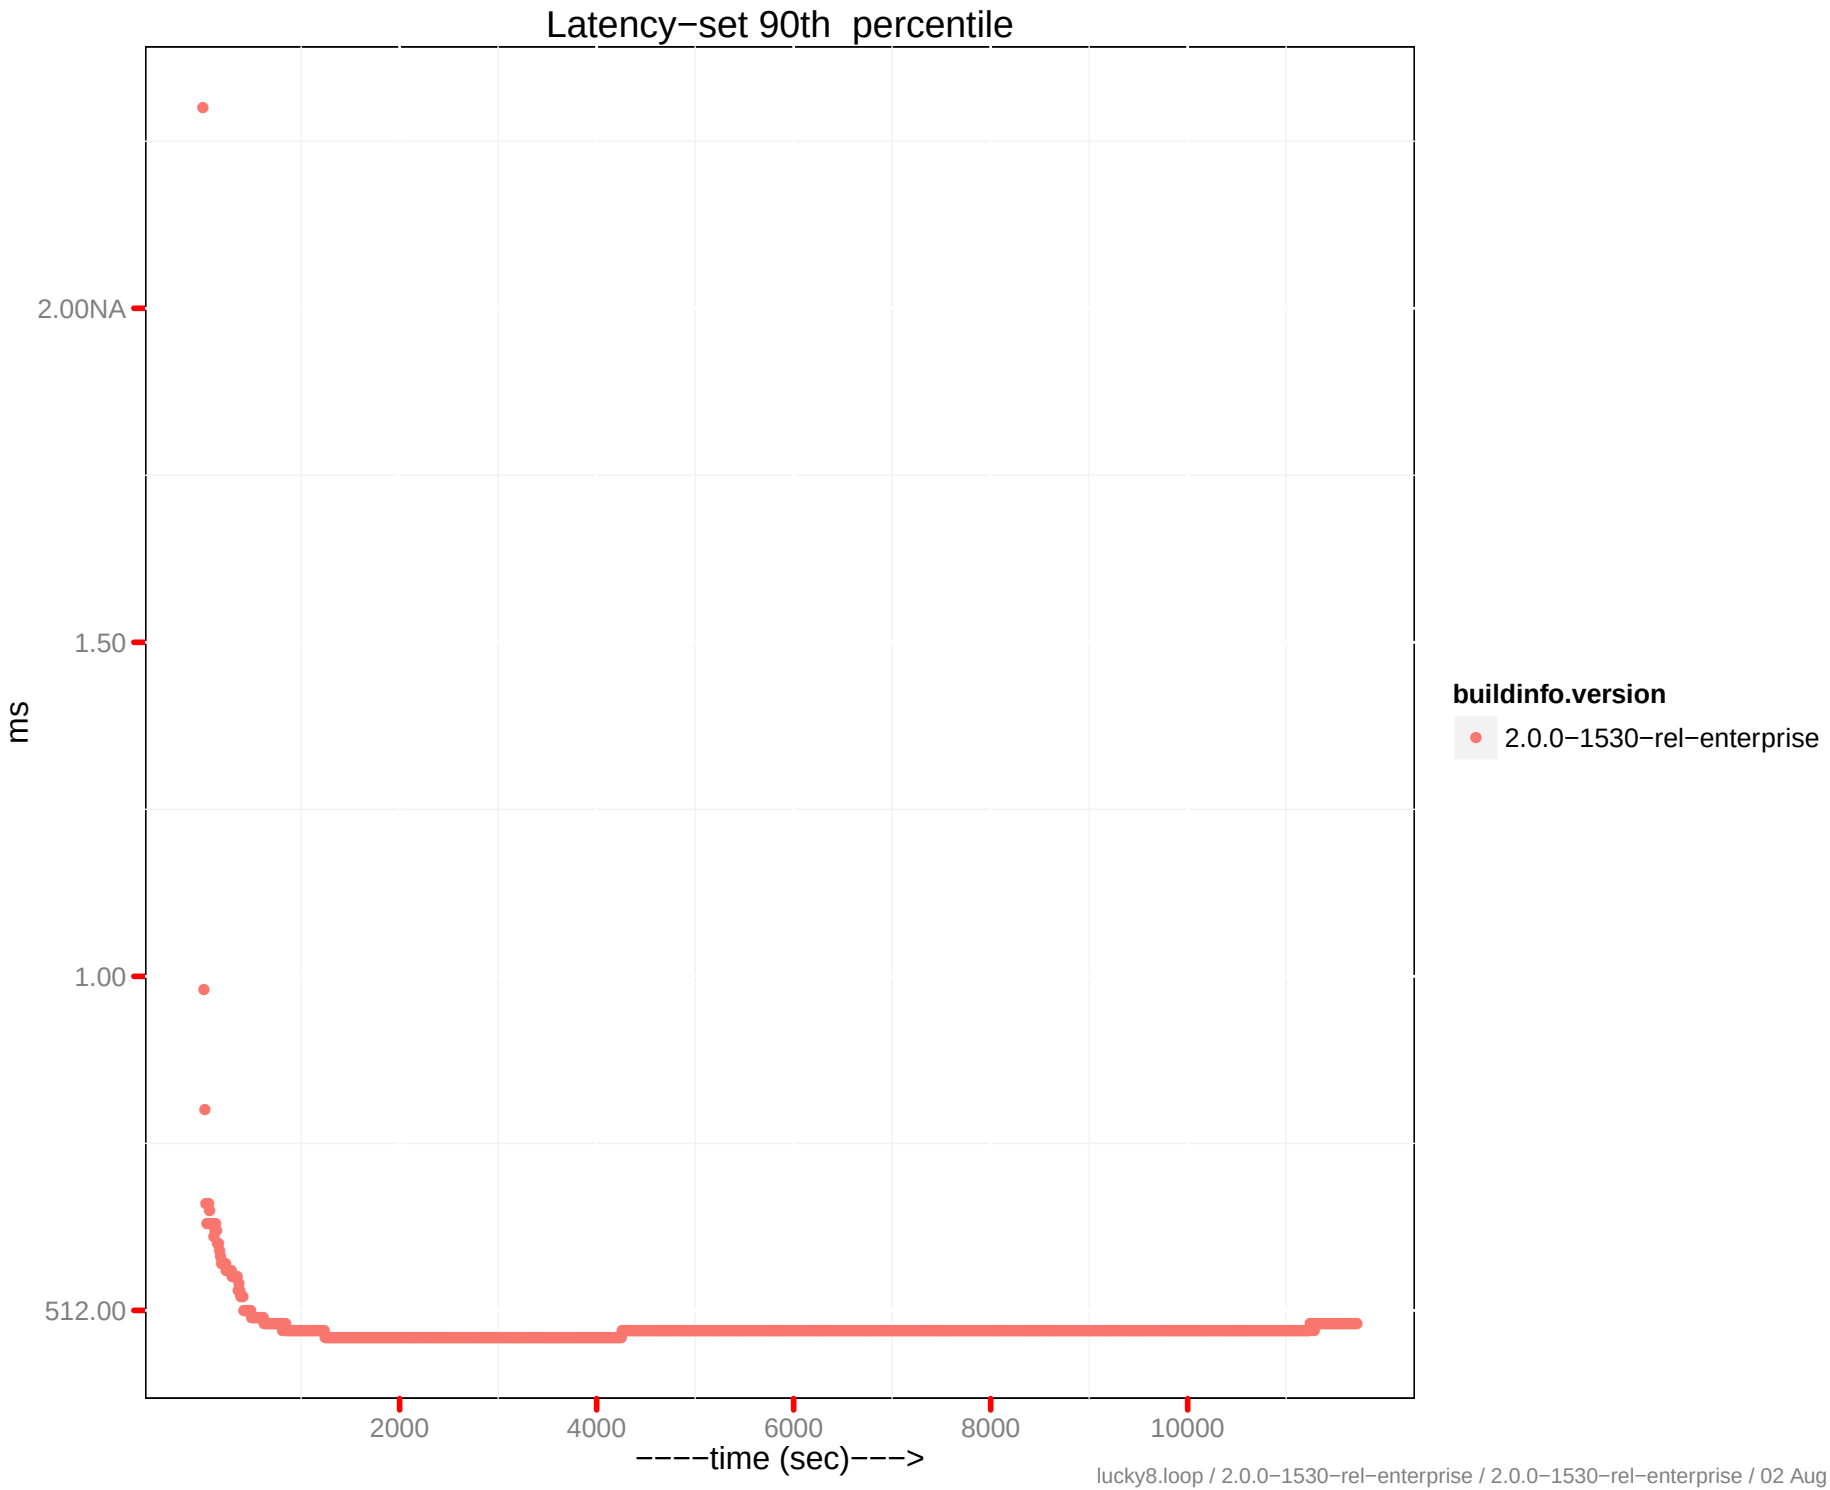
Latency−set 90th percentile
2.00NA
1.50
1.00
512.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

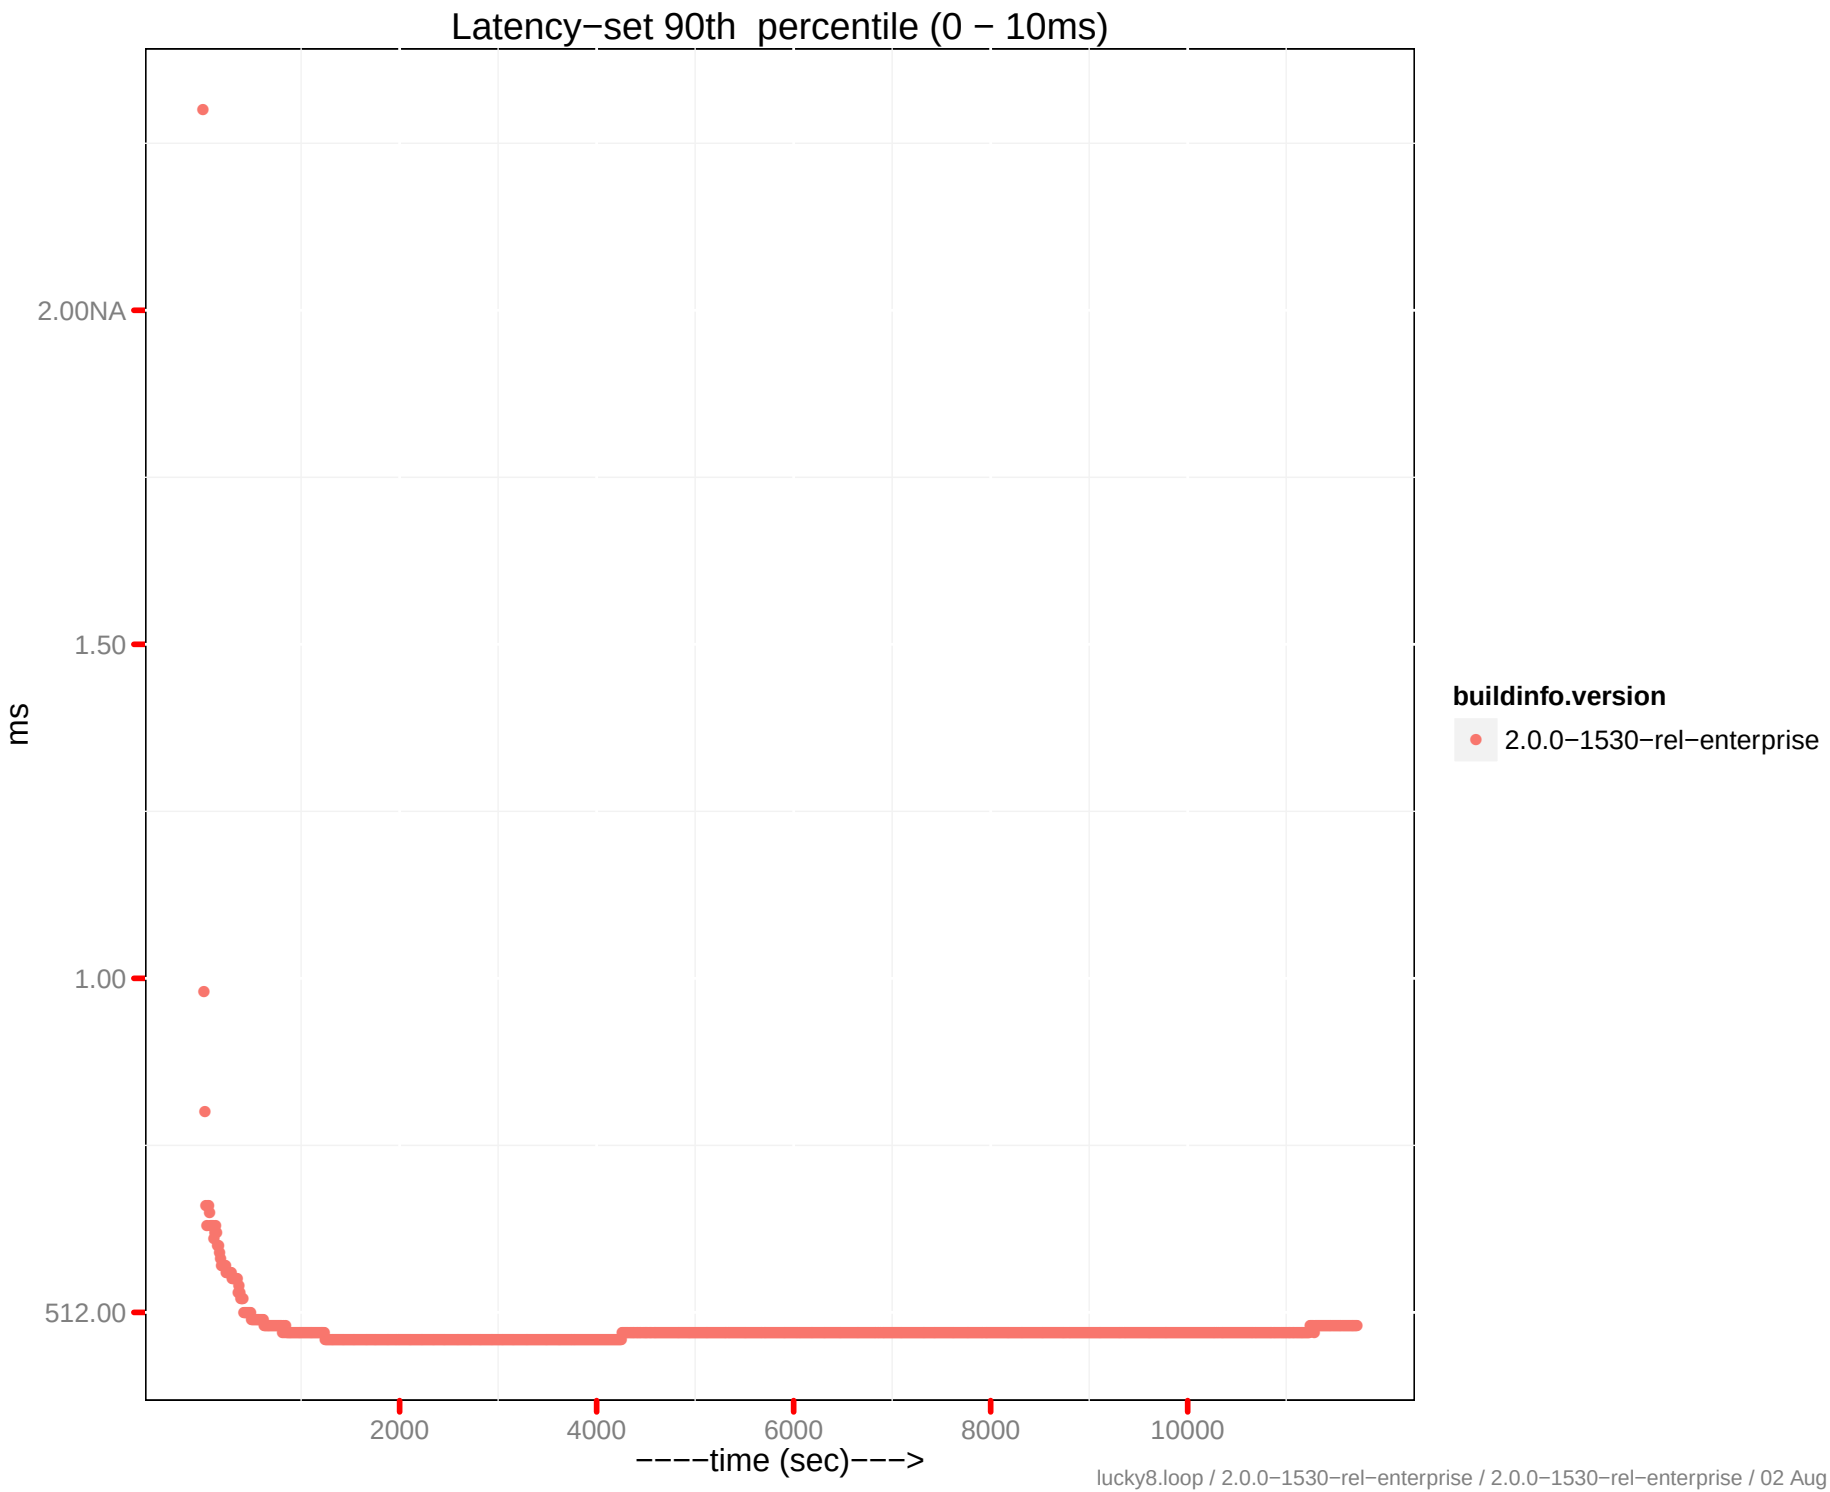
Latency−set 90th percentile (0 − 10ms)
2.00NA
1.50
1.00
512.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

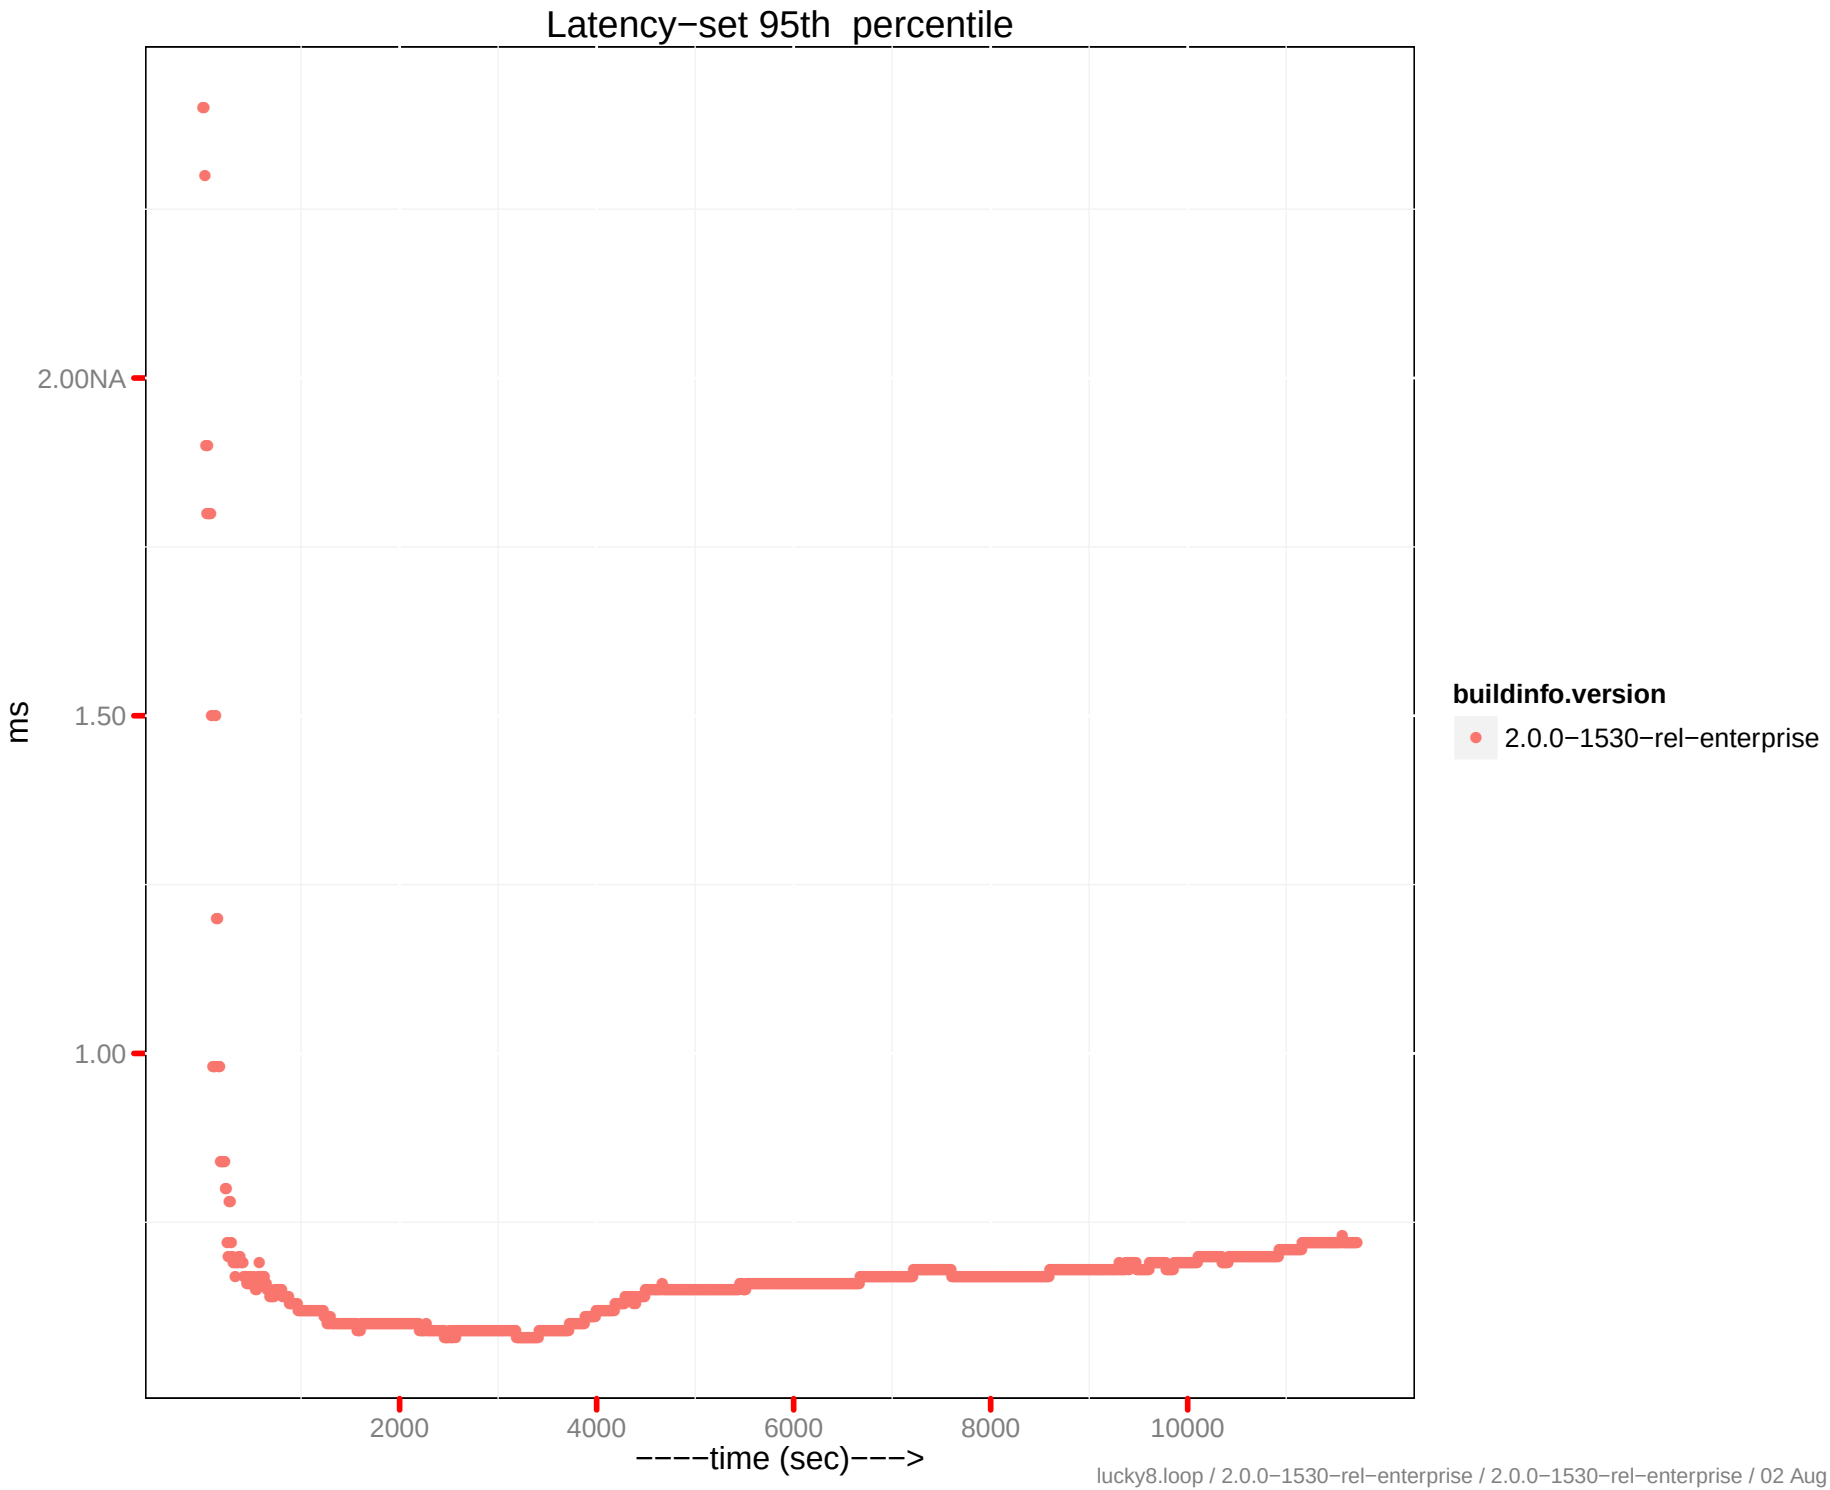
Latency−set 95th percentile
2.00NA
1.50
1.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

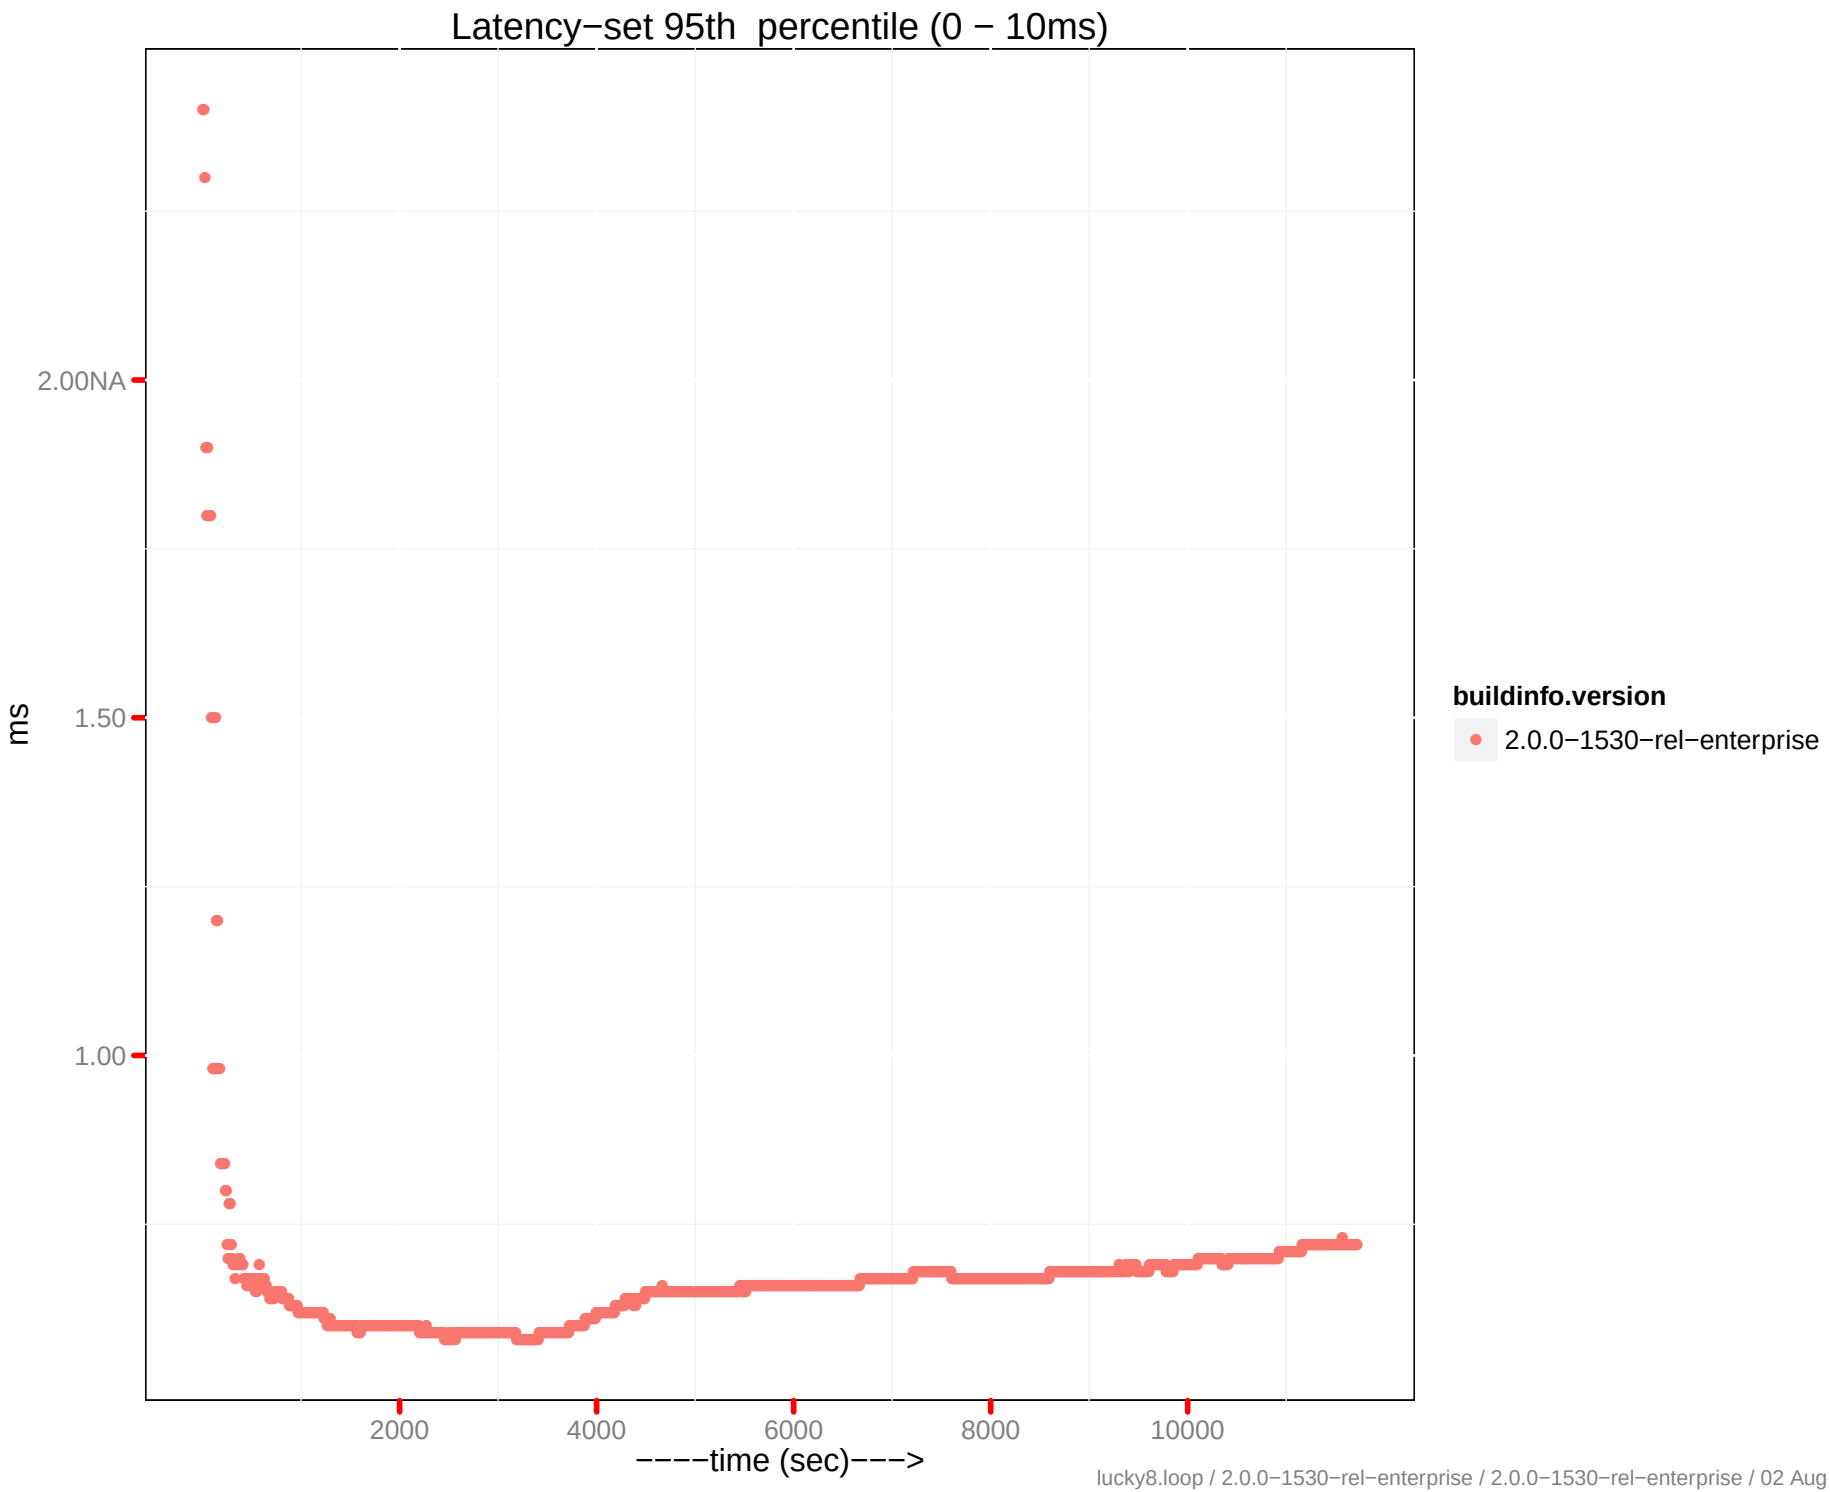
Latency−set 95th percentile (0 − 10ms)
2.00NA
1.50
1.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

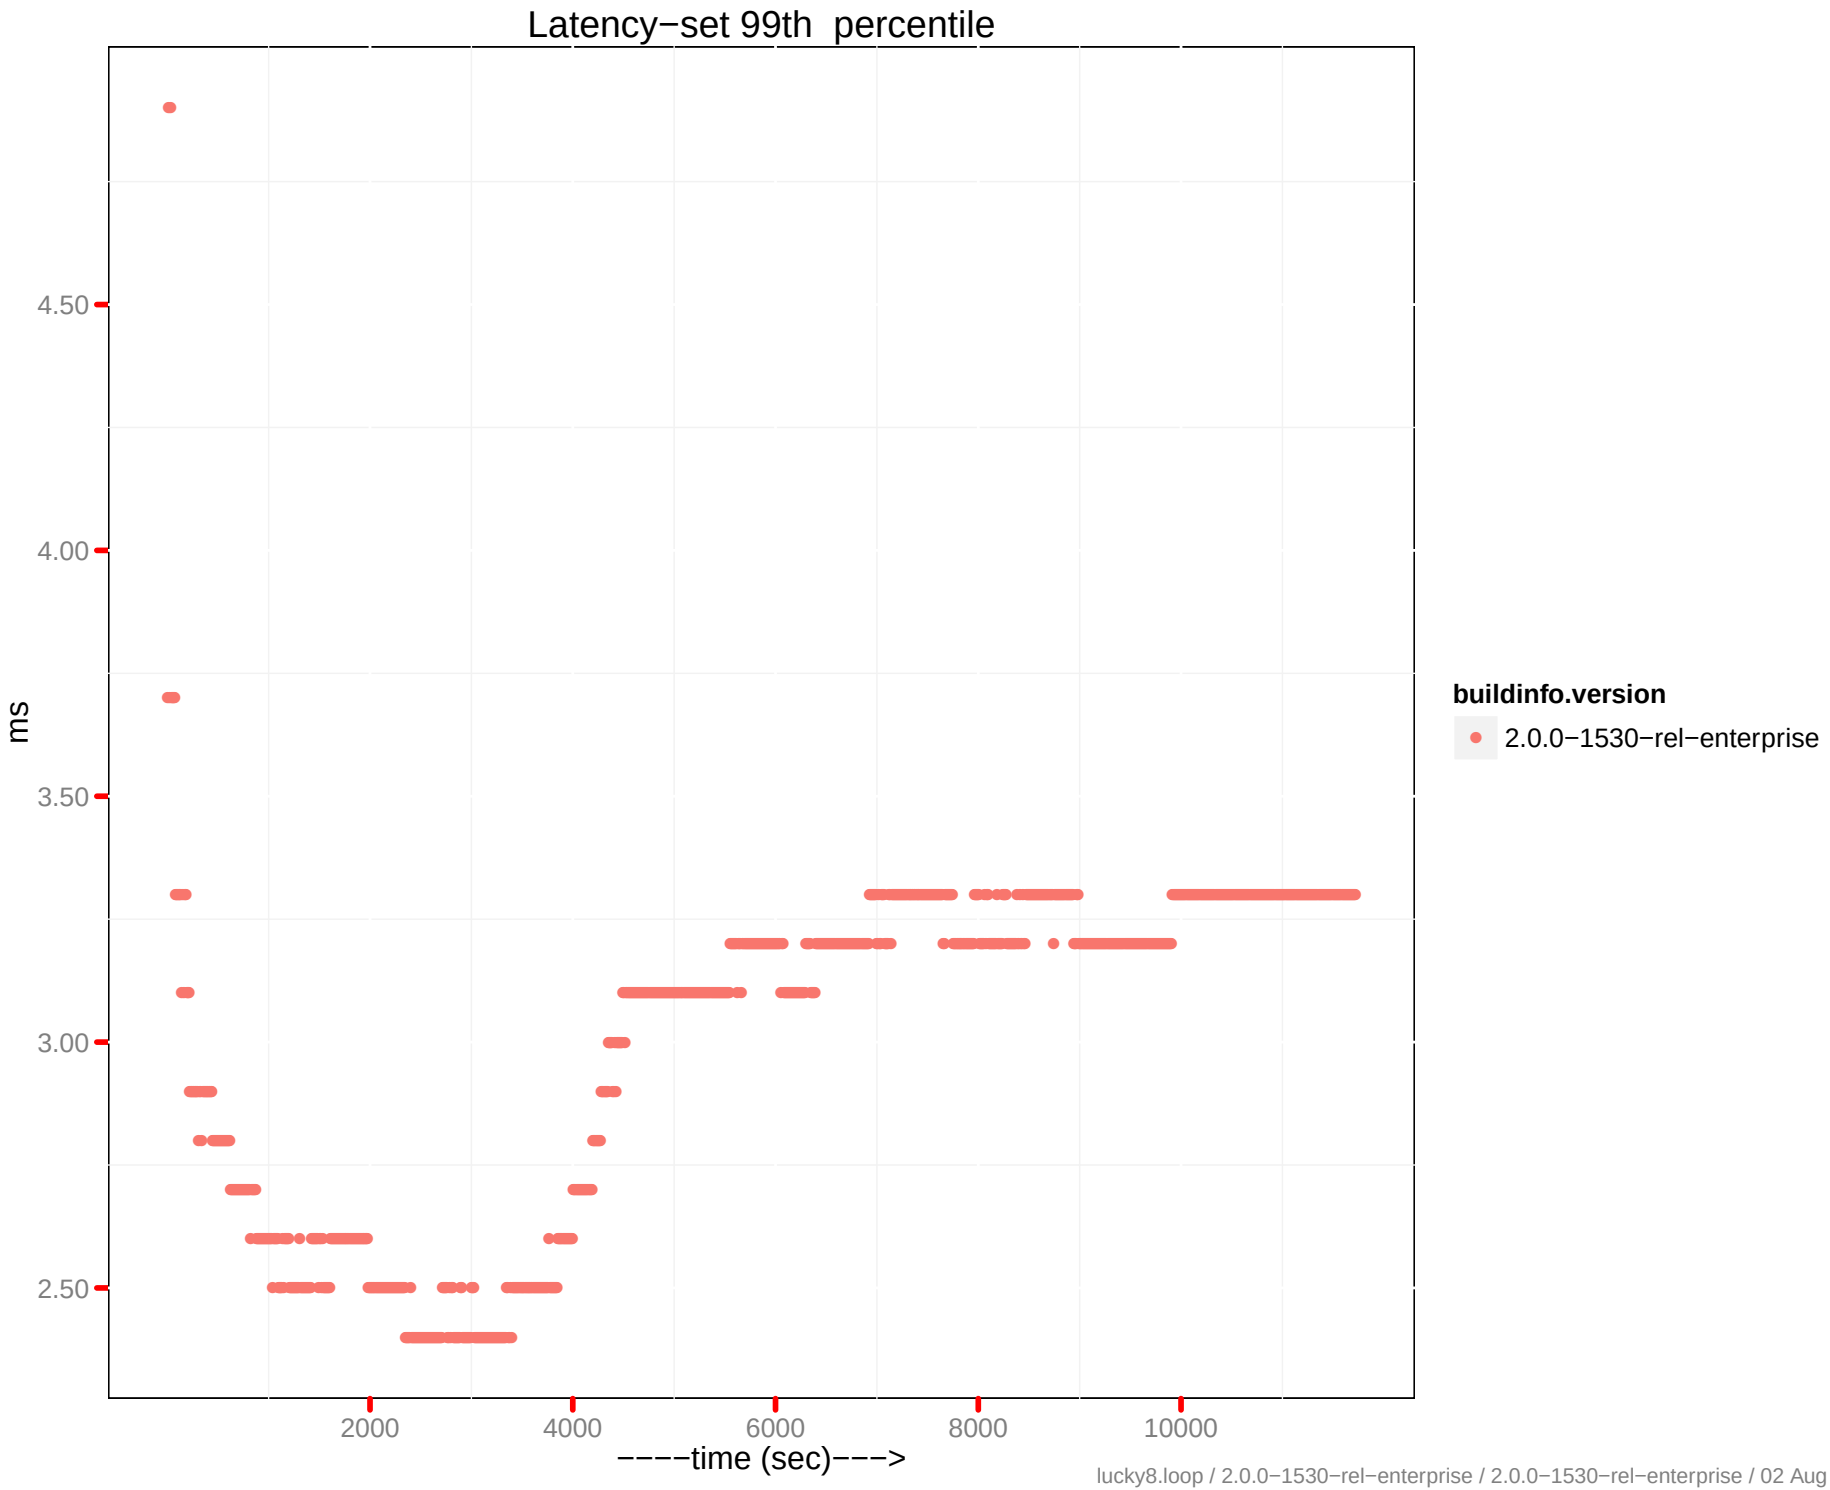
Latency−set 99th percentile
4.50
4.00
3.50
3.00
2.50
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

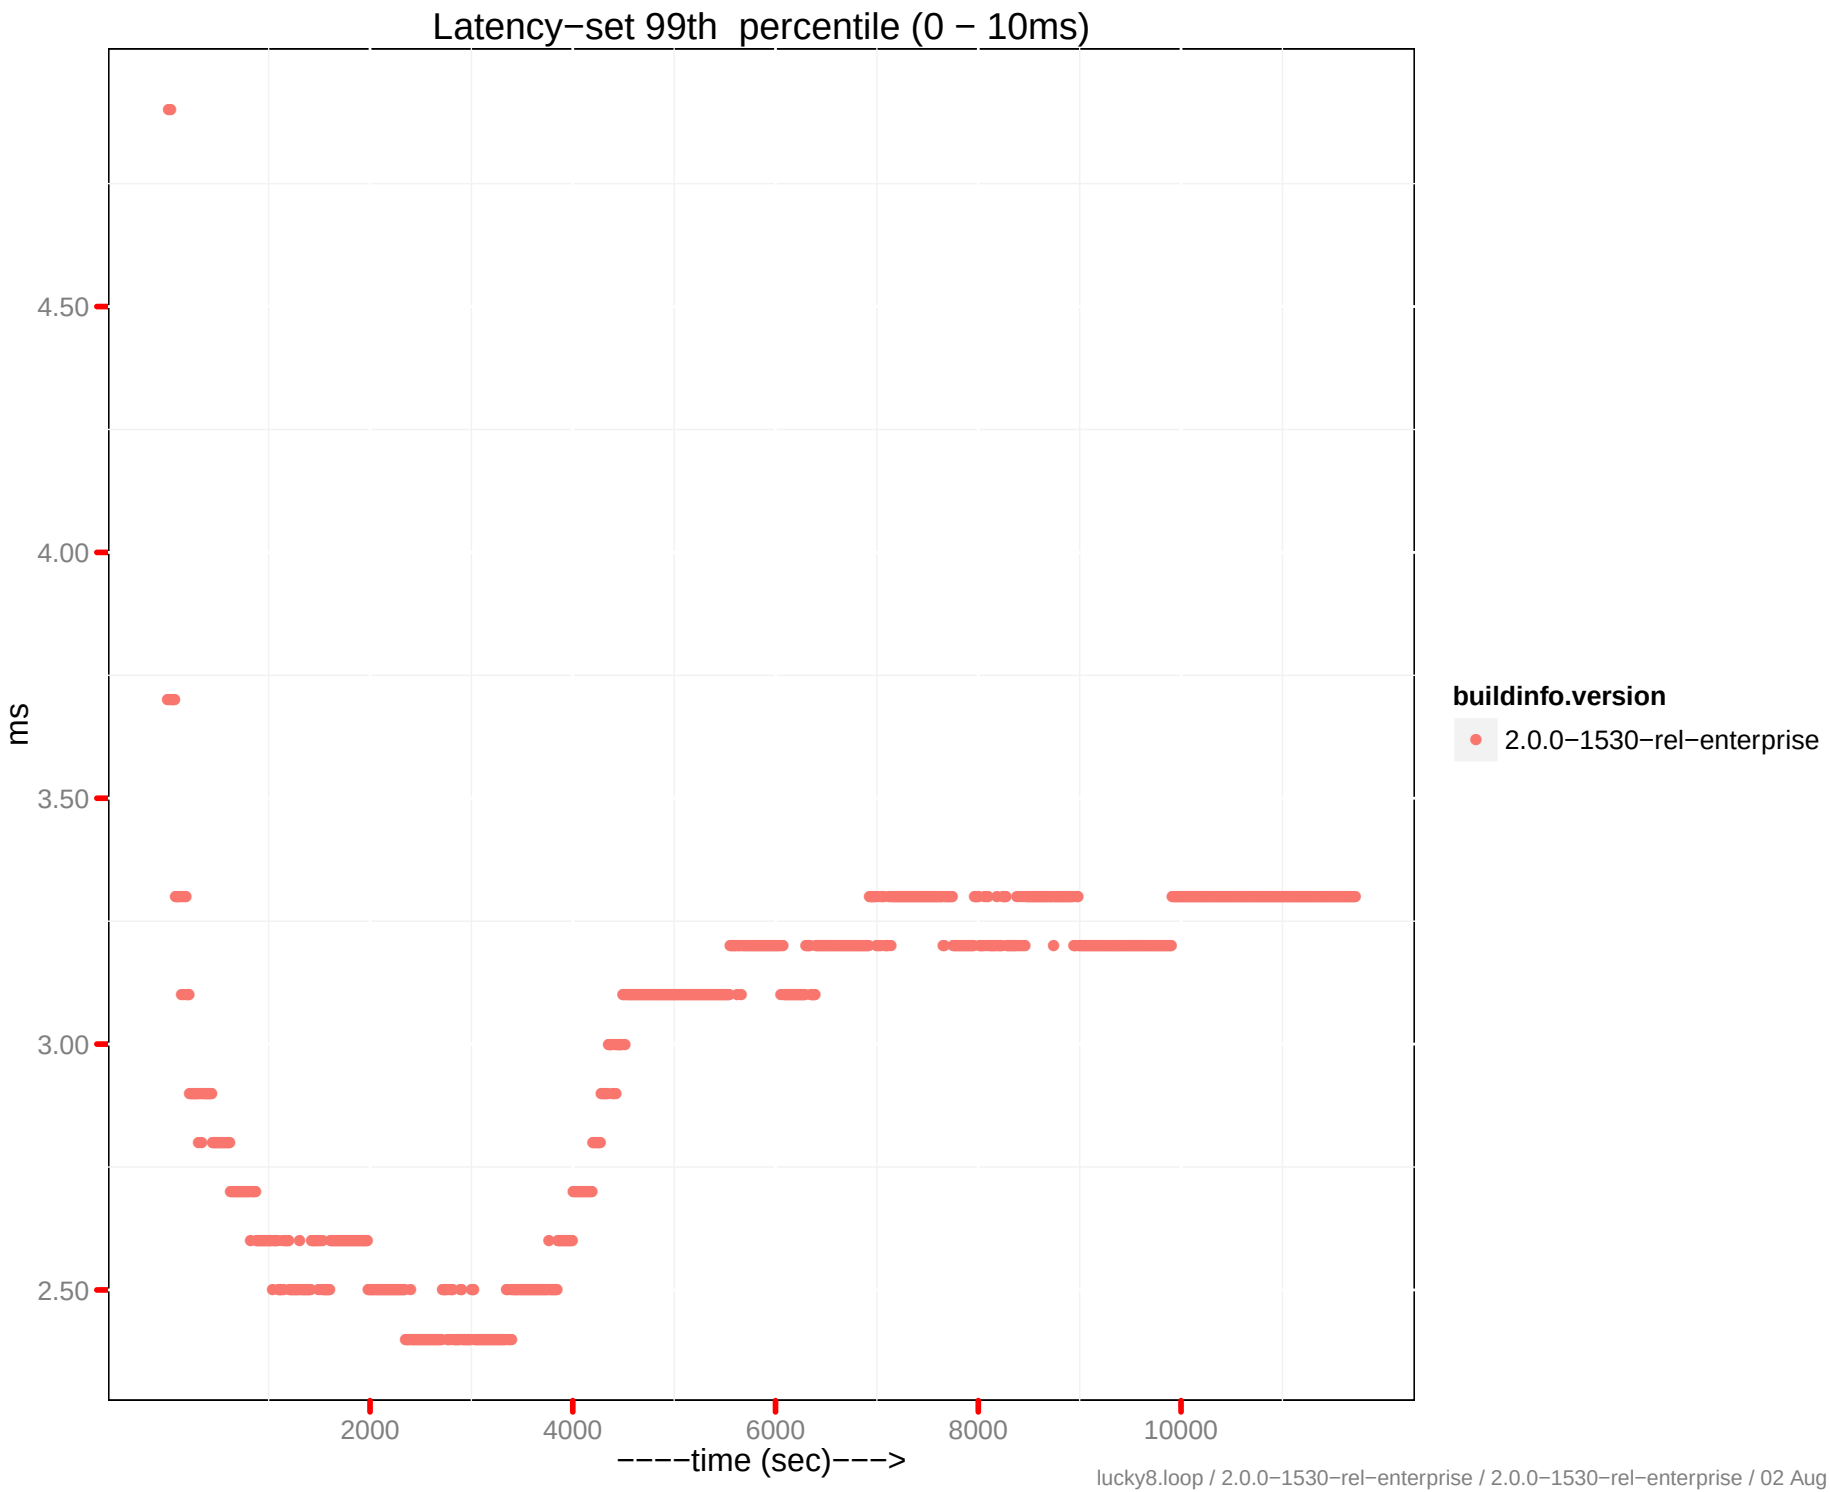
Latency−set 99th percentile (0 − 10ms)
4.50
4.00
3.50
3.00
2.50
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

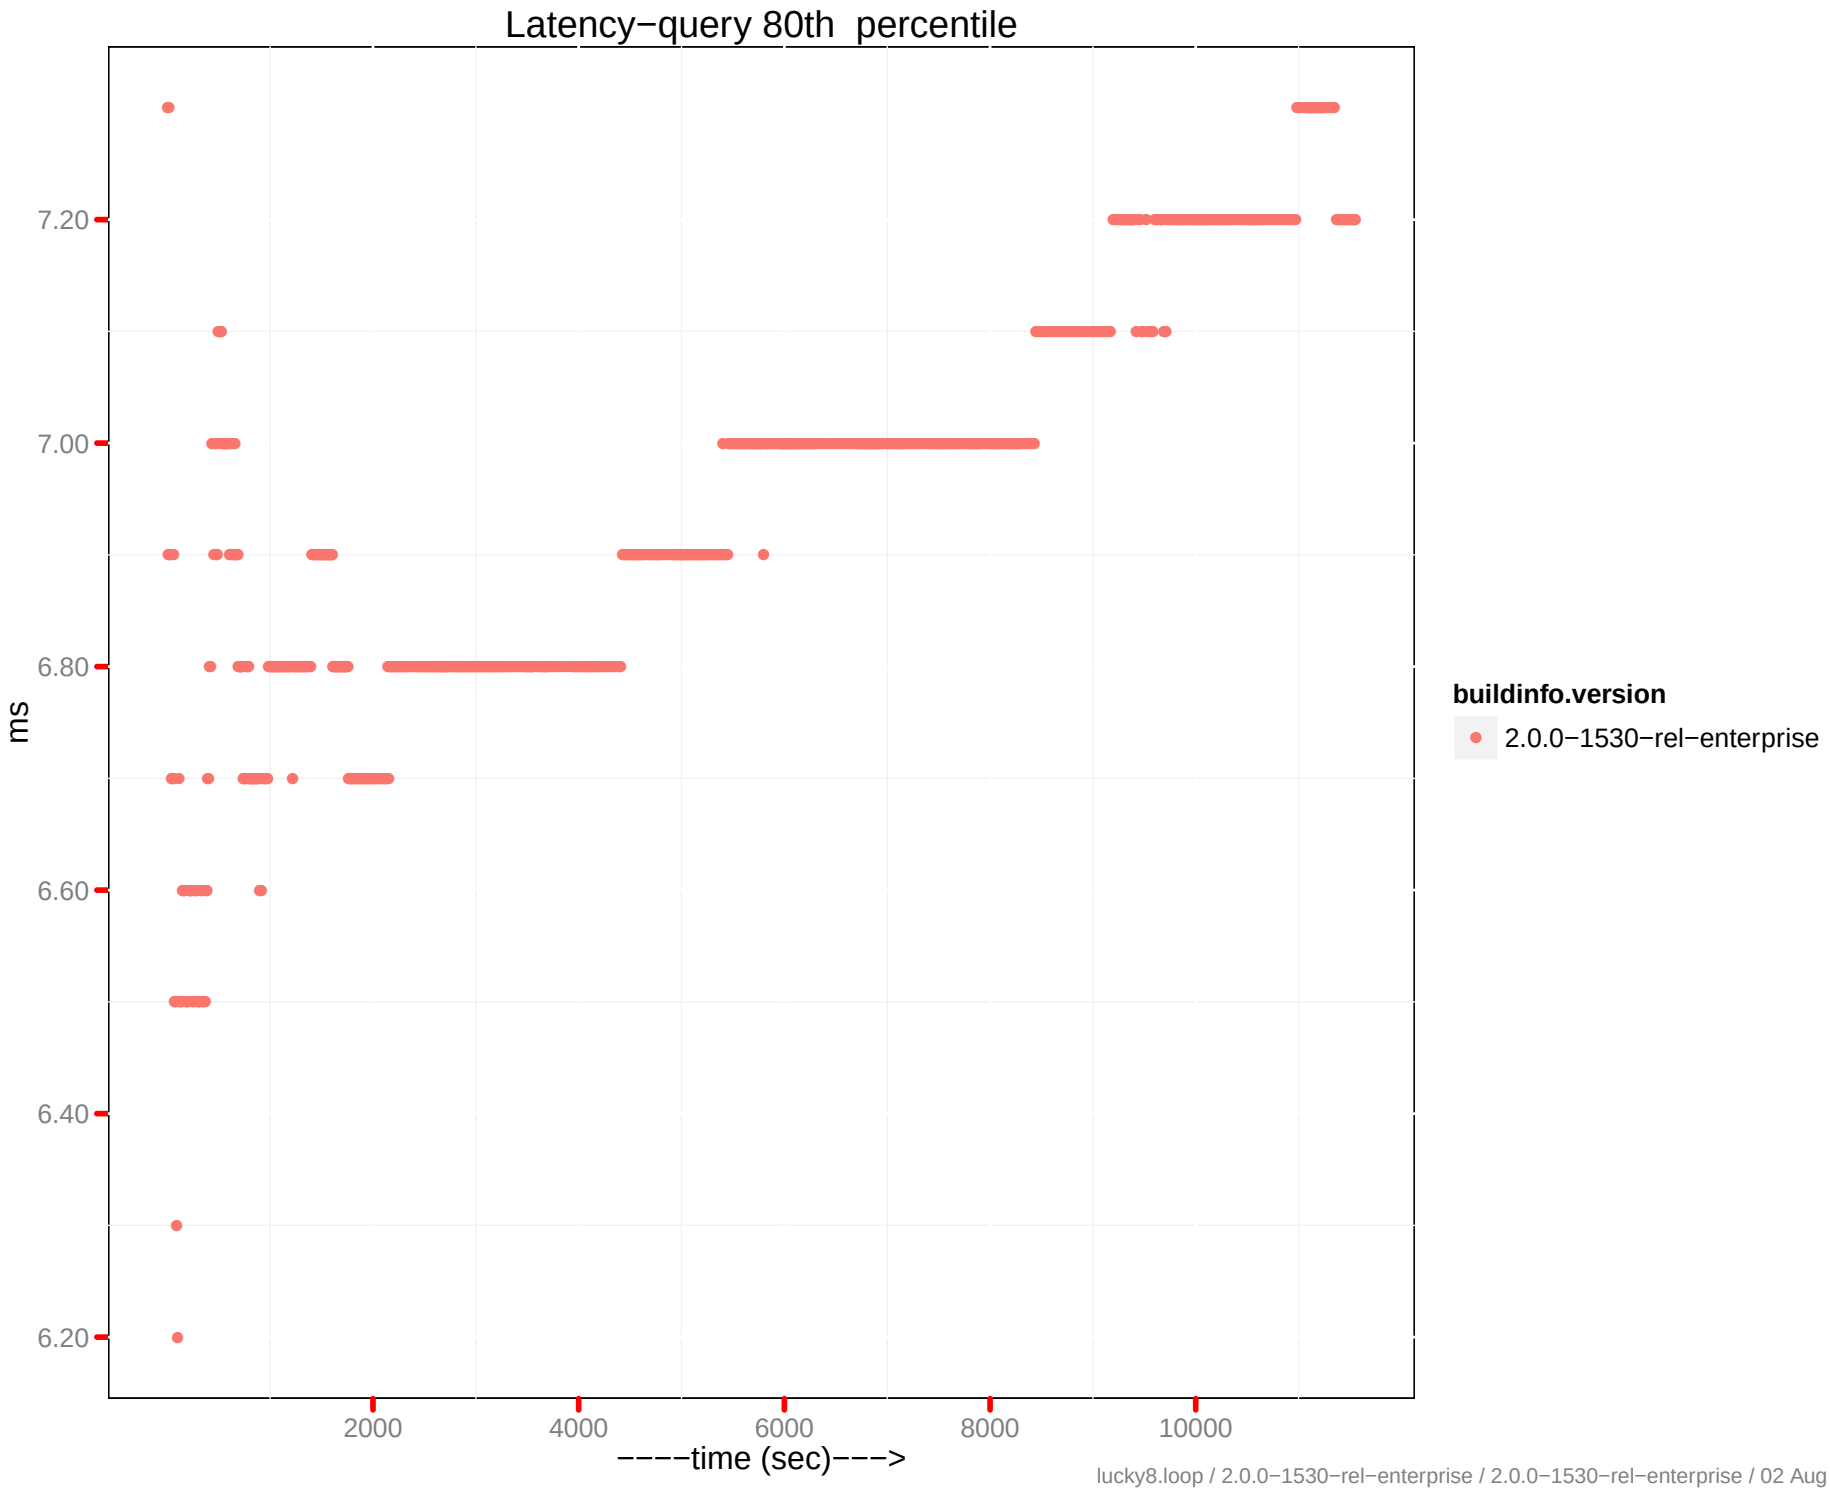
Latency−query 80th percentile
7.20
7.00
6.80
6.60
6.40
6.20
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

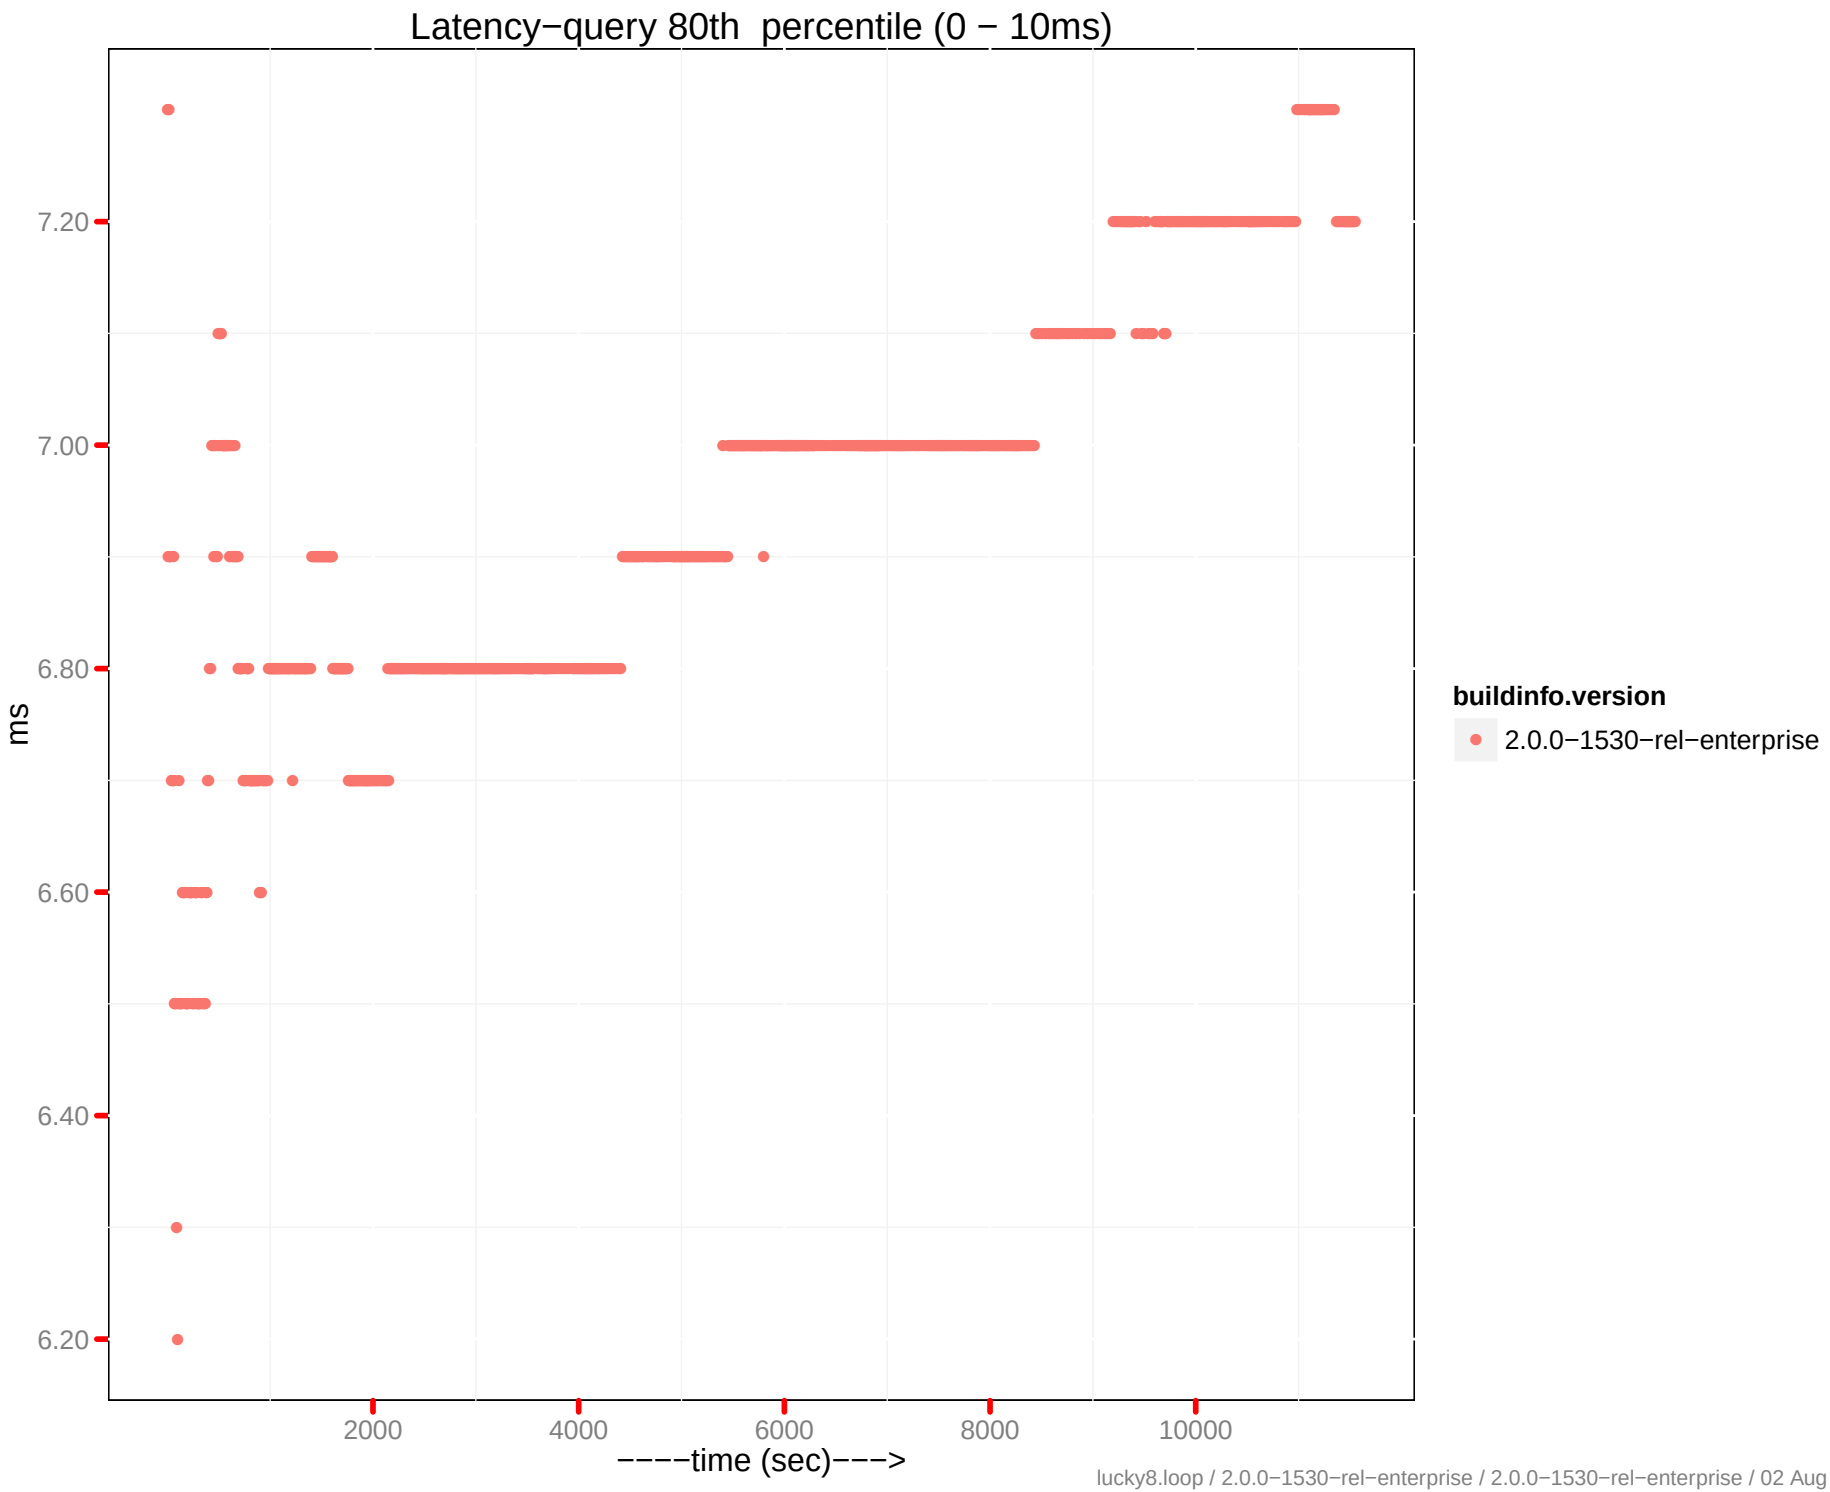
Latency−query 80th percentile (0 − 10ms)
7.20
7.00
6.80
6.60
6.40
6.20
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

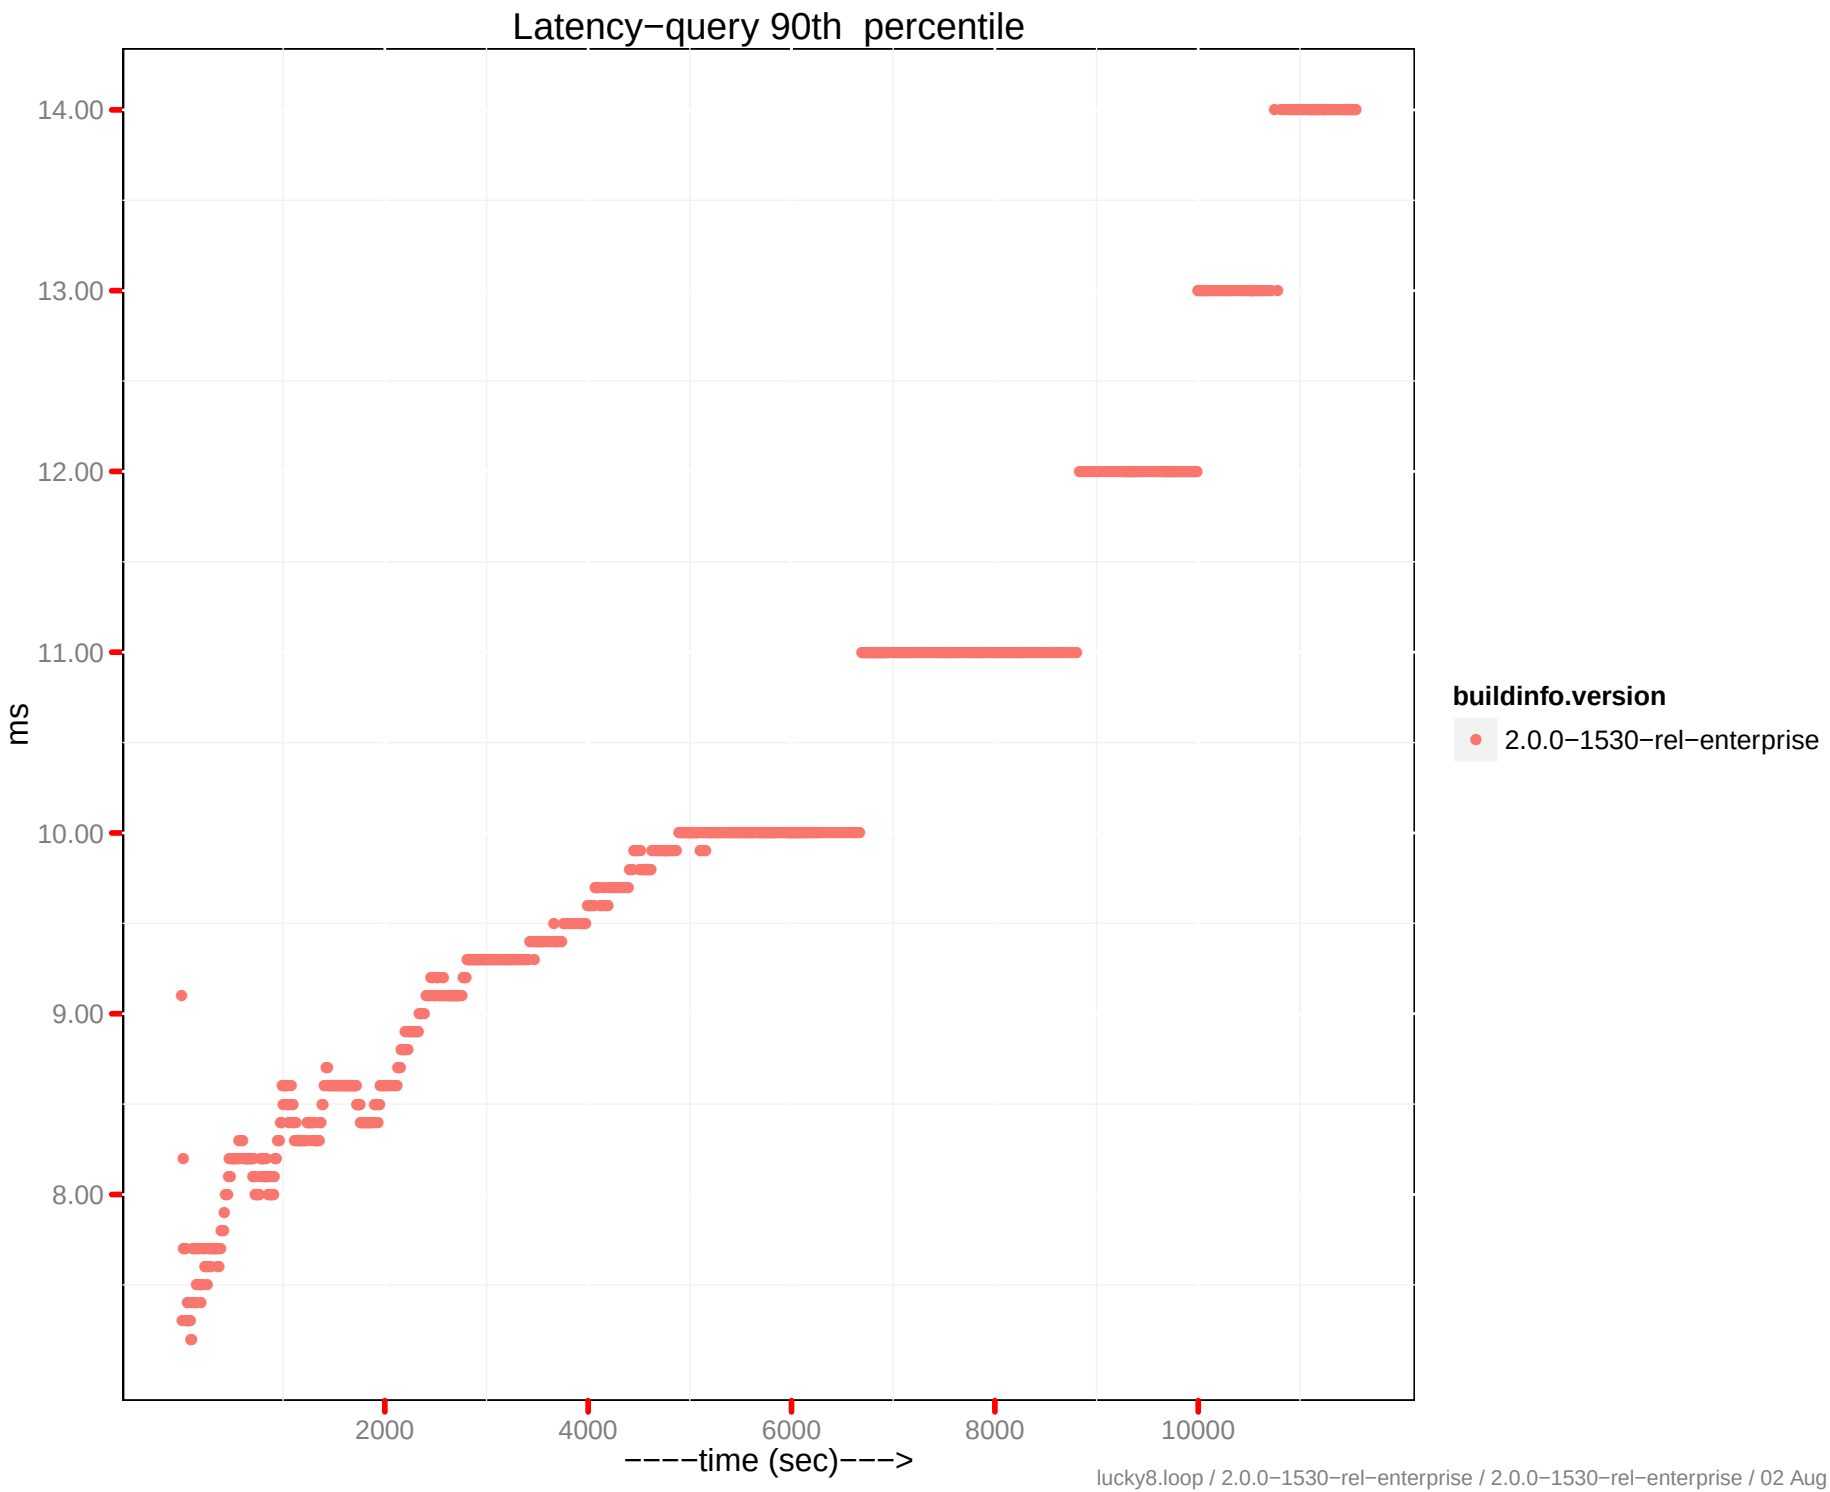
Latency−query 90th percentile
14.00
13.00
12.00
11.00
10.00
9.00
8.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

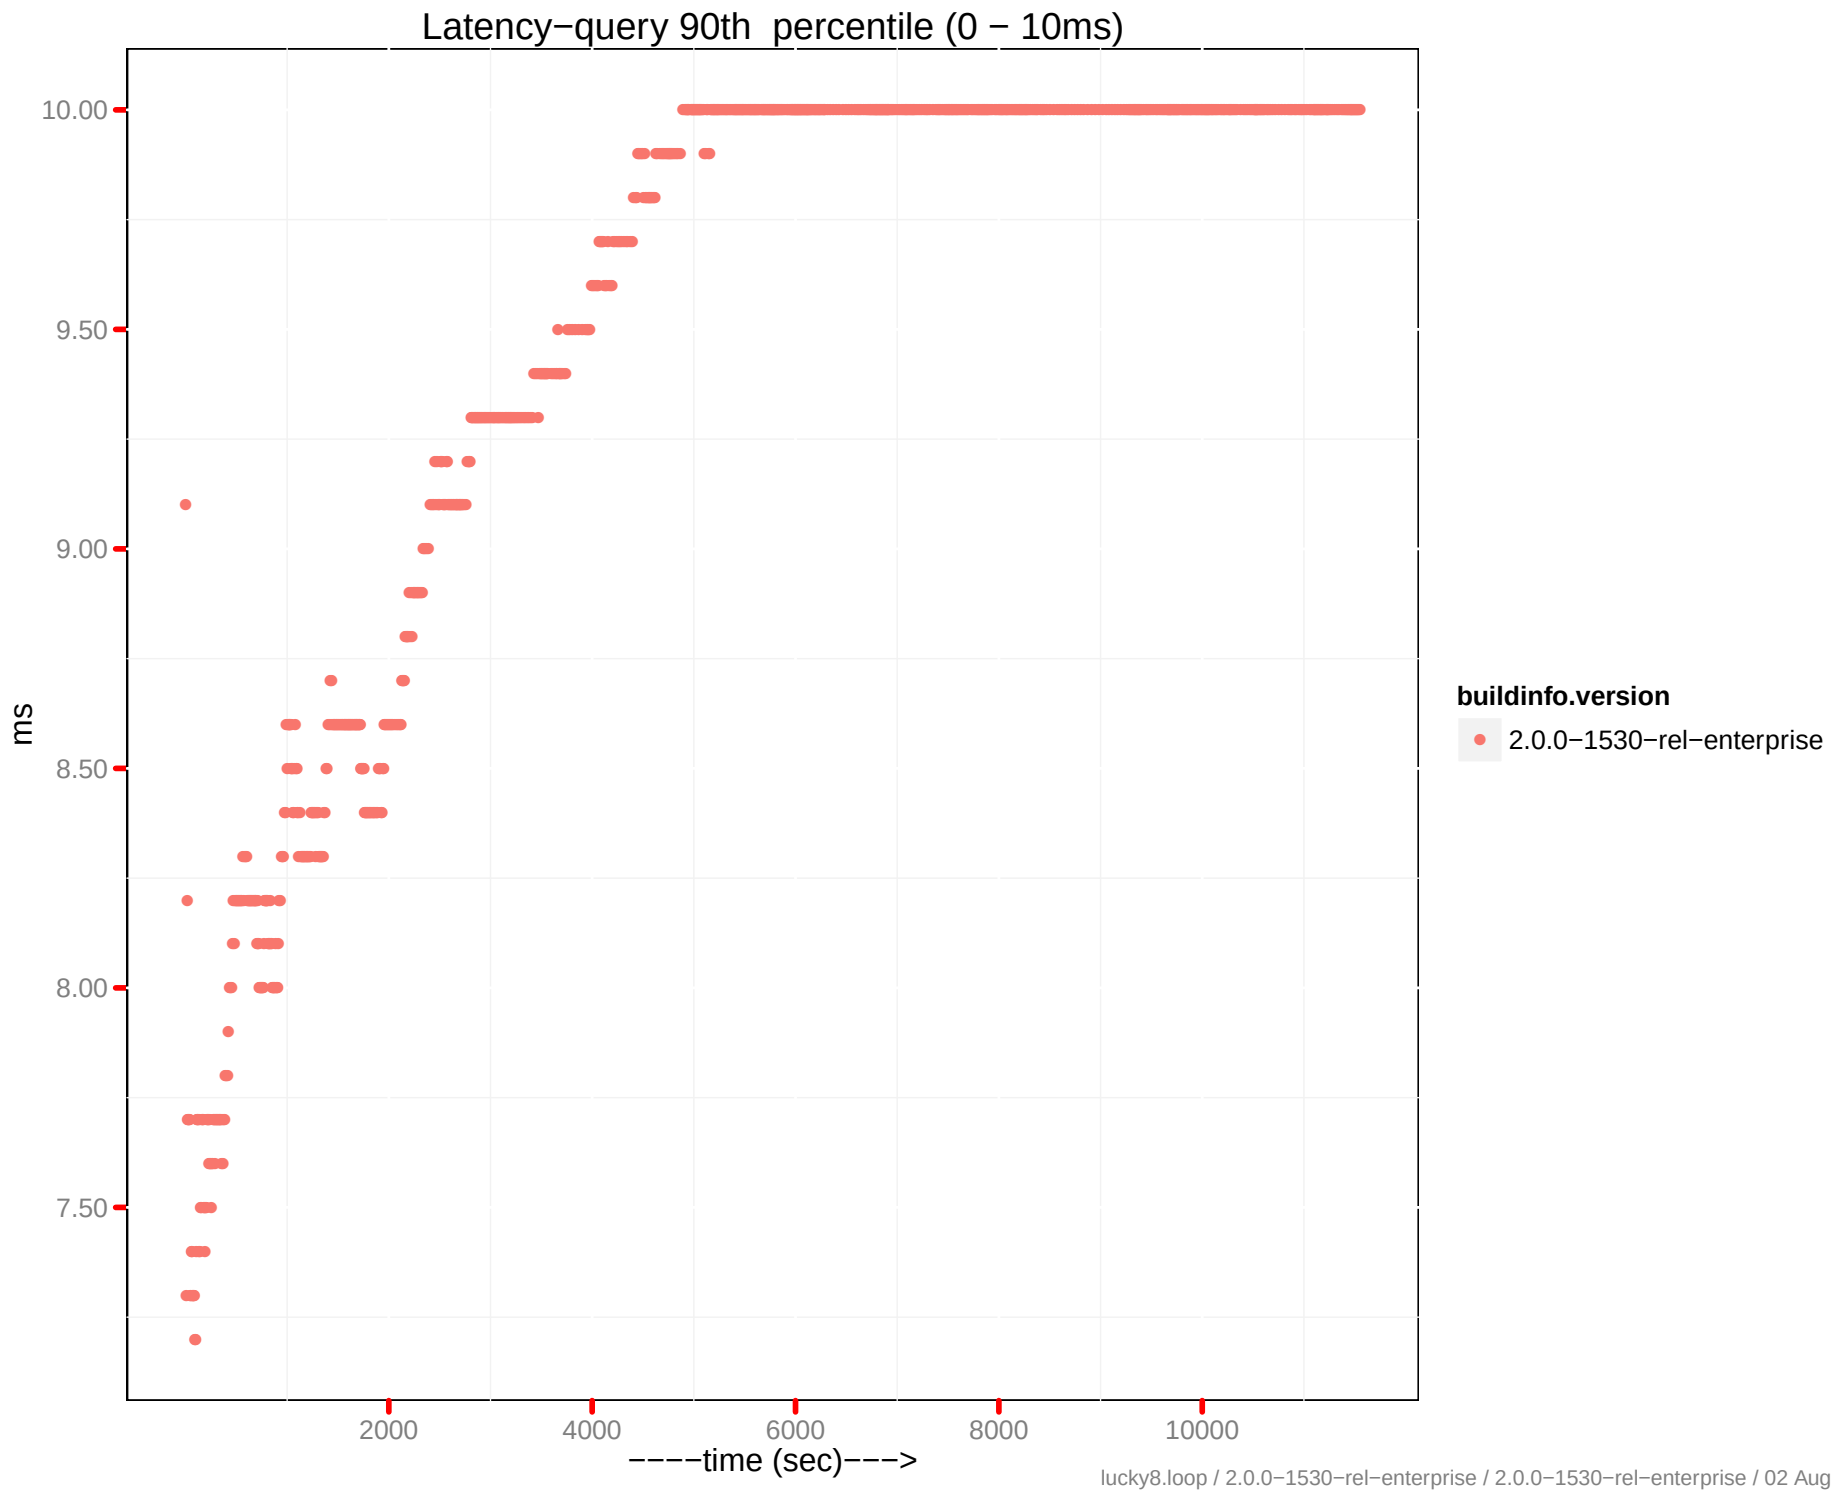
Latency−query 90th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
8.00
7.50
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

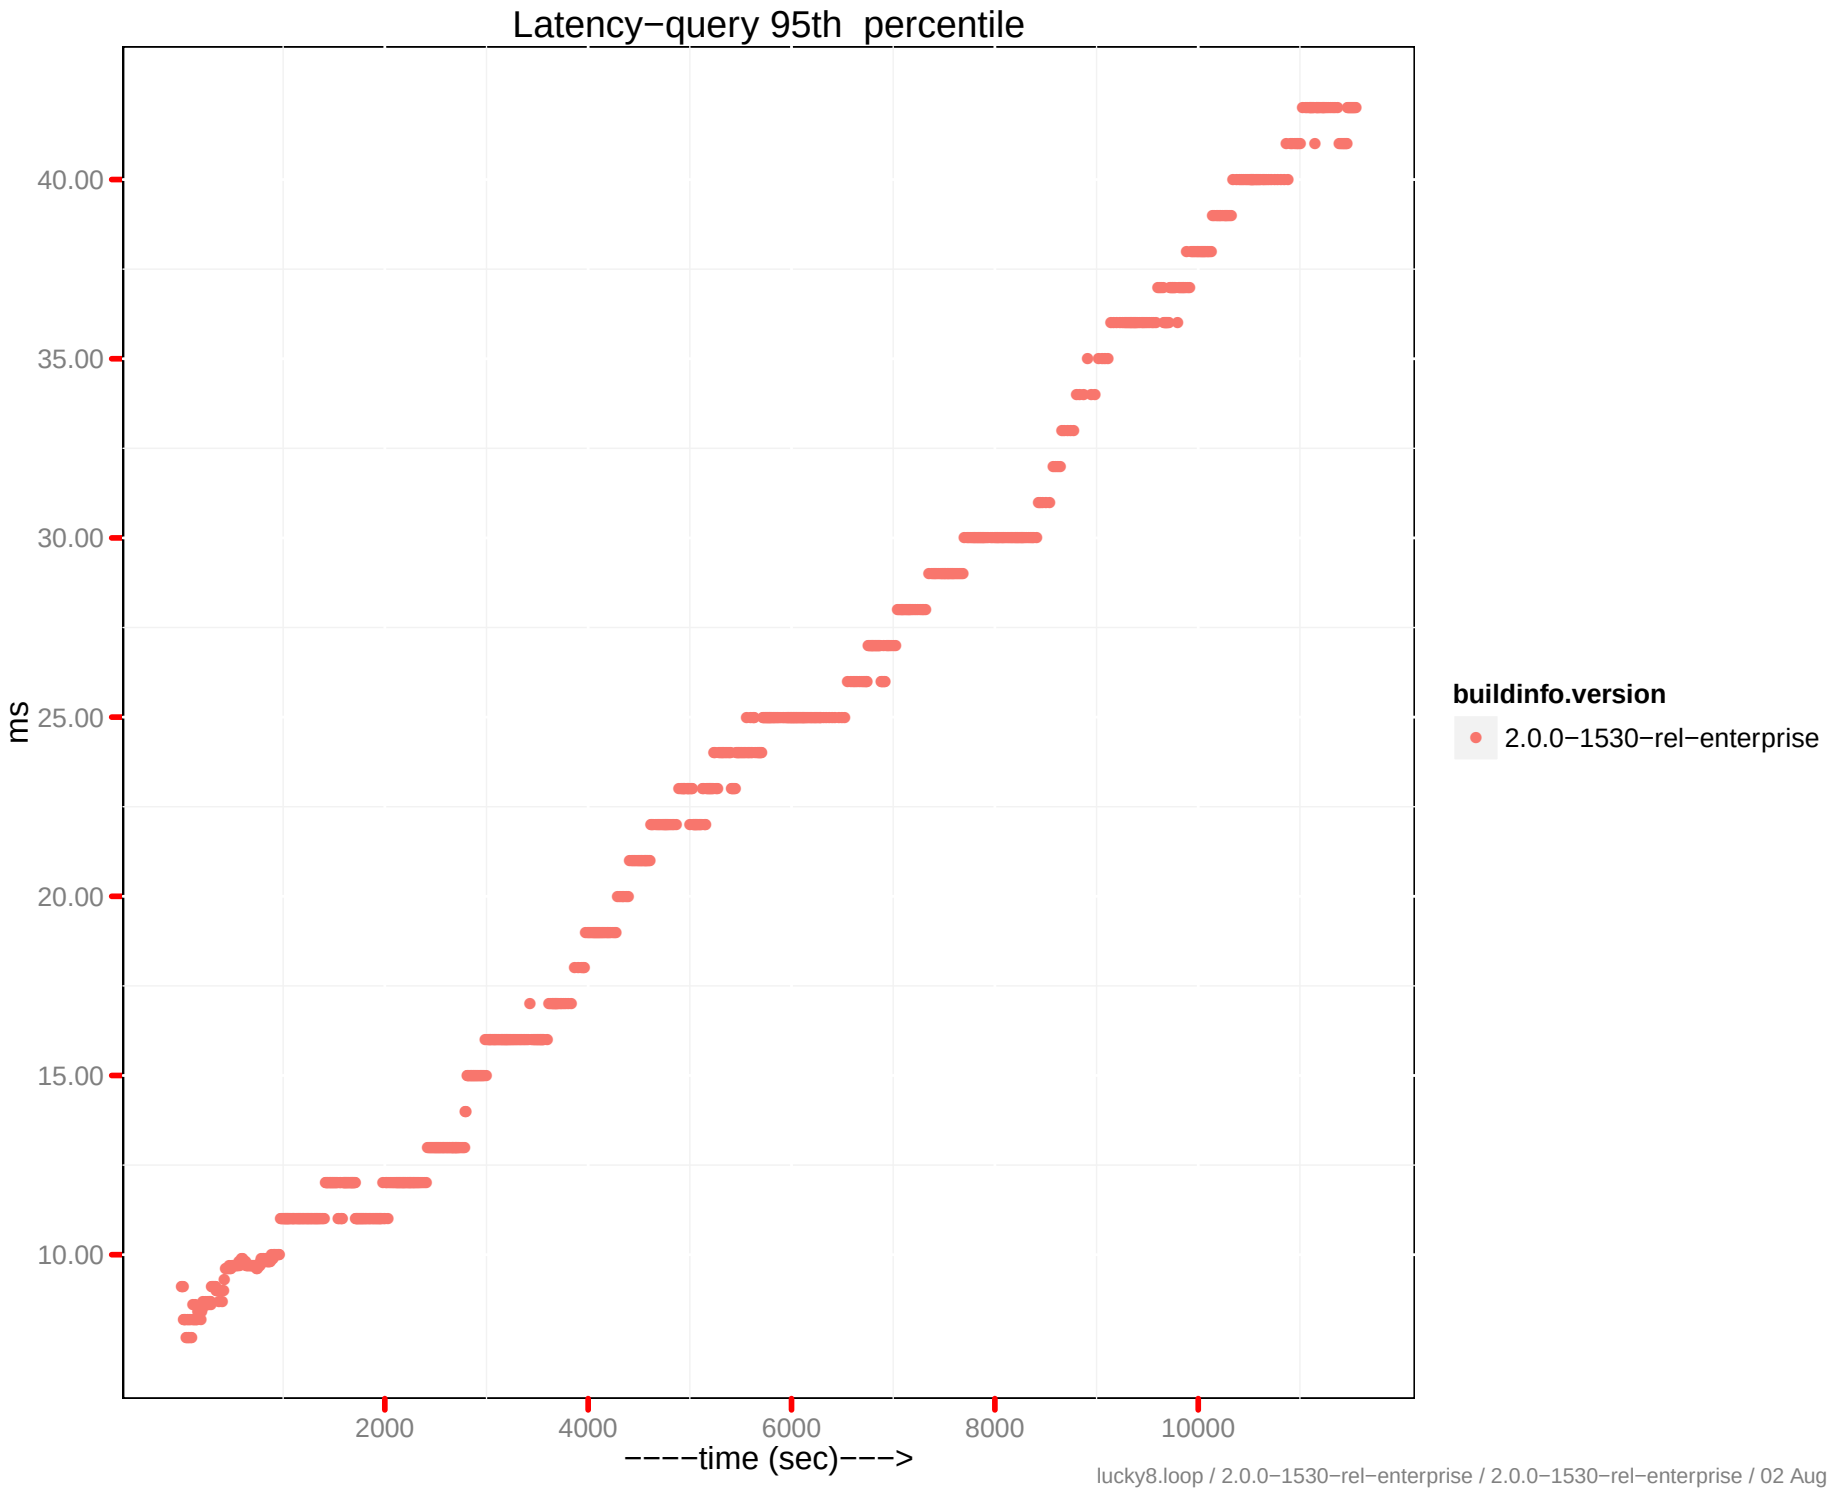
Latency−query 95th percentile
40.00
35.00
30.00
25.00
20.00
15.00
10.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

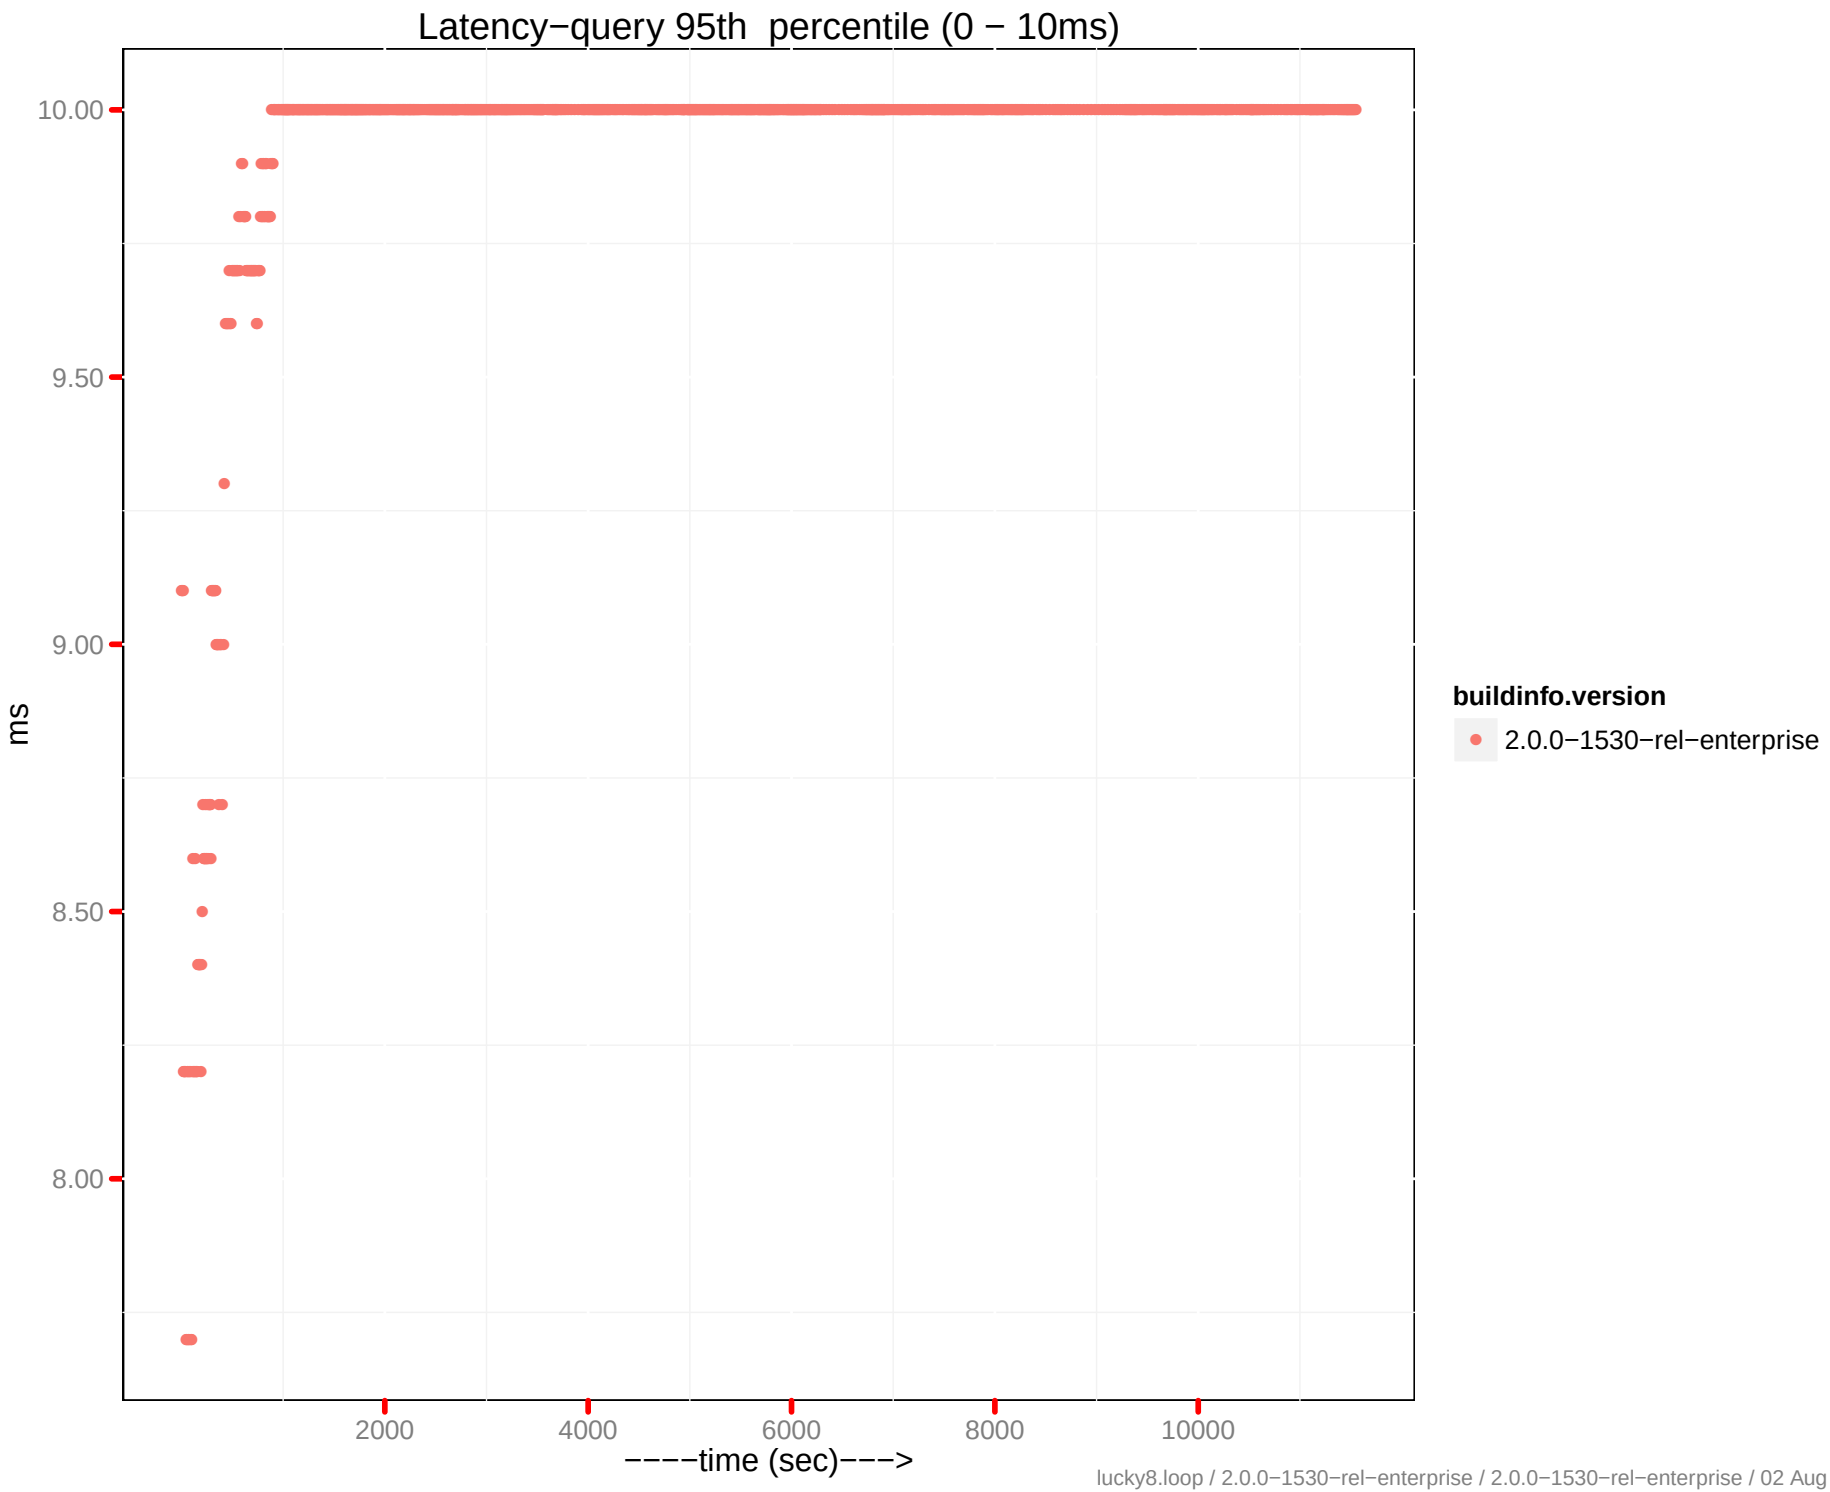
Latency−query 95th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
8.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

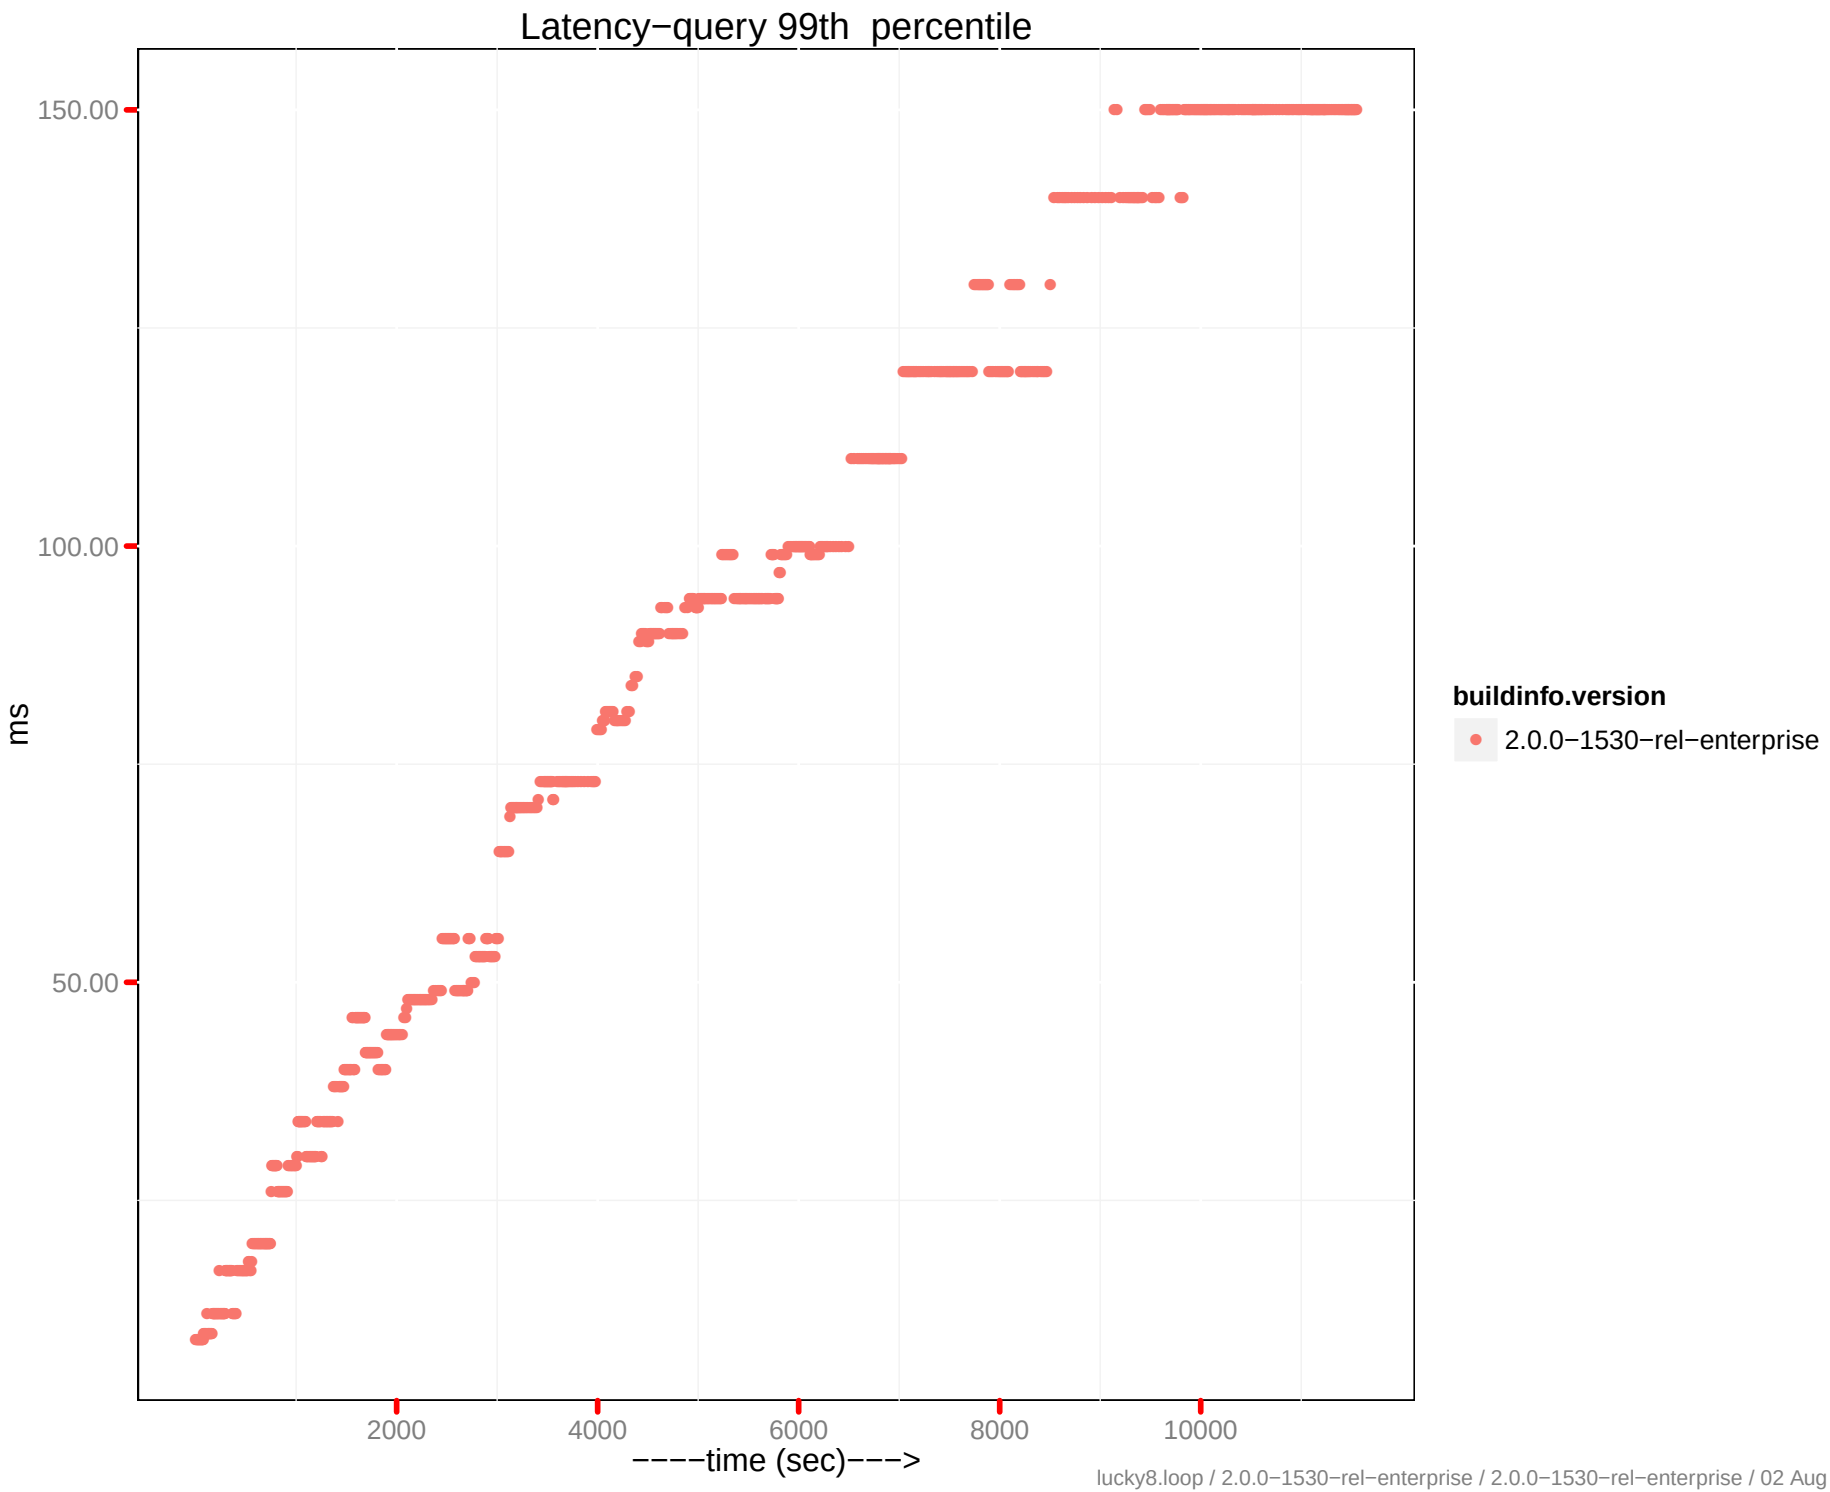
Latency−query 99th percentile
150.00
100.00
50.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

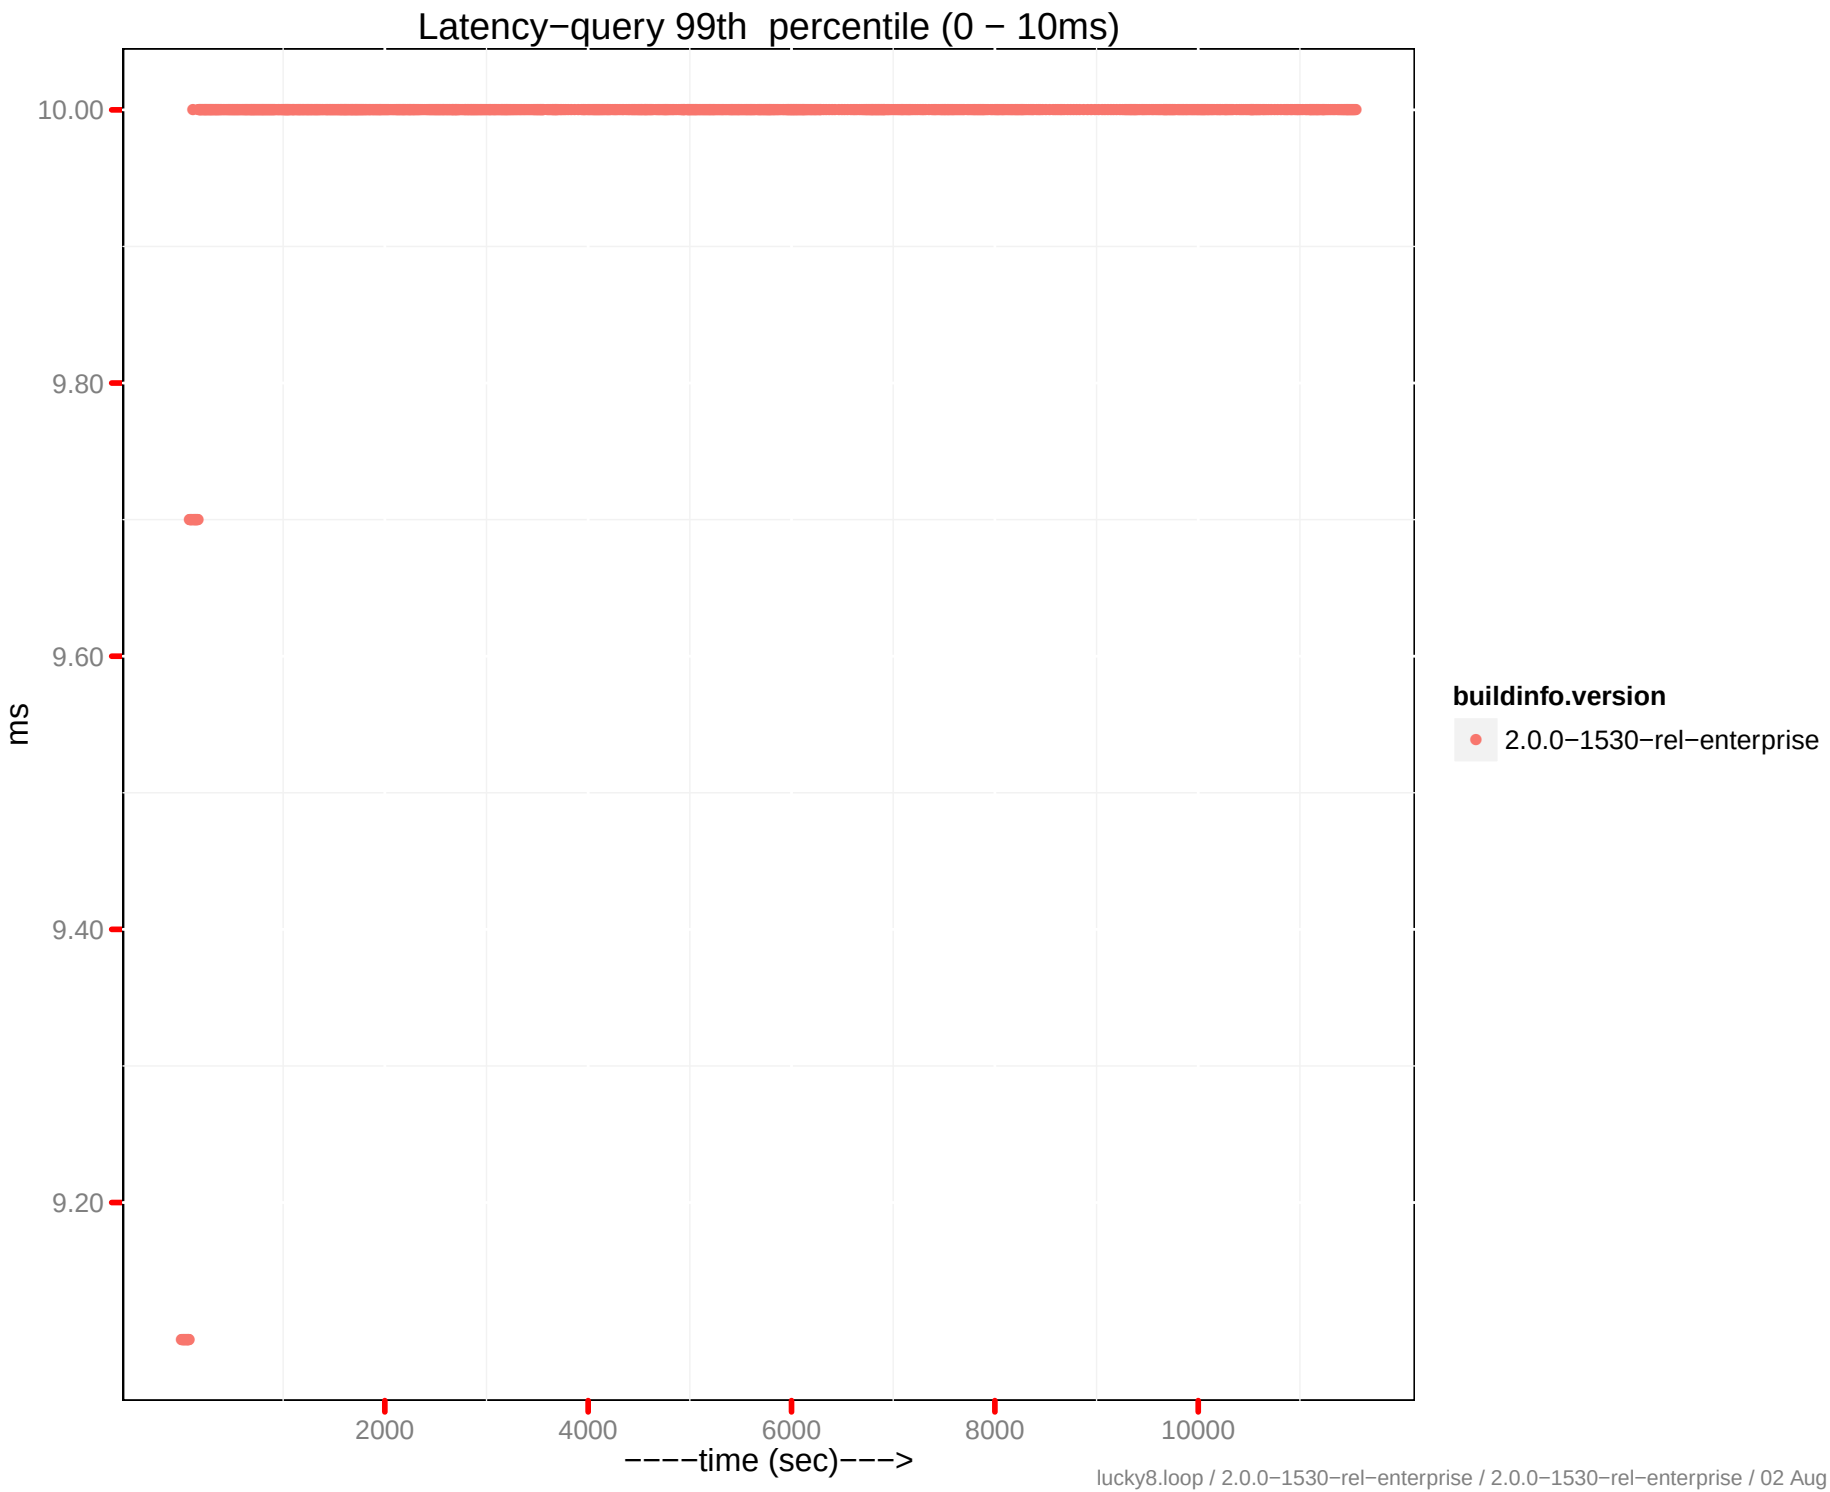
Latency−query 99th percentile (0 − 10ms)
10.00
9.80
9.60
9.40
9.20
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

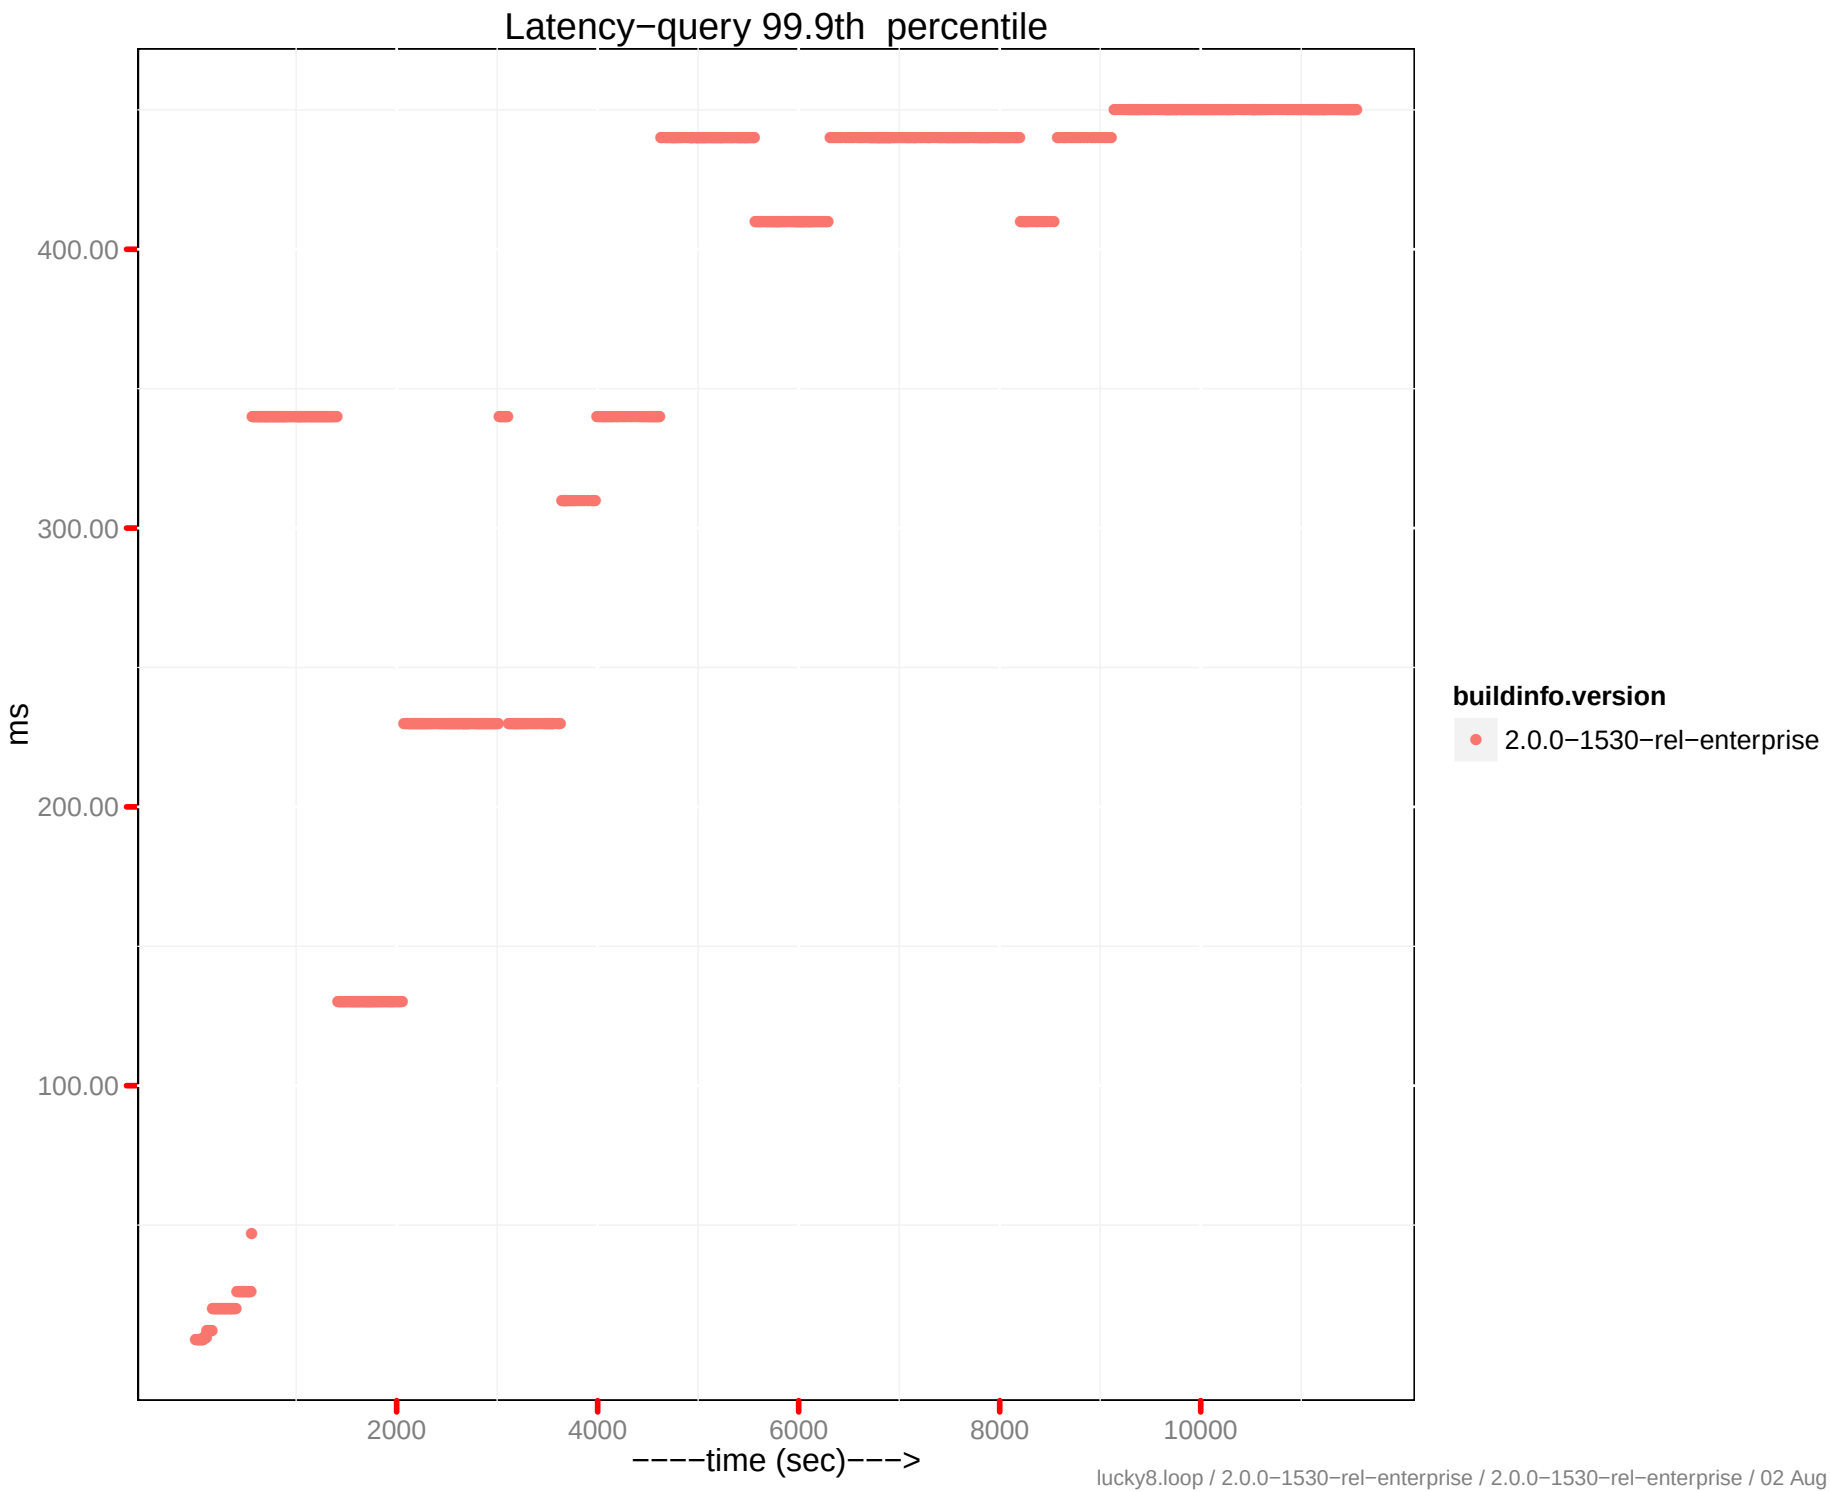
Latency−query 99.9th percentile
400.00
300.00
200.00
100.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

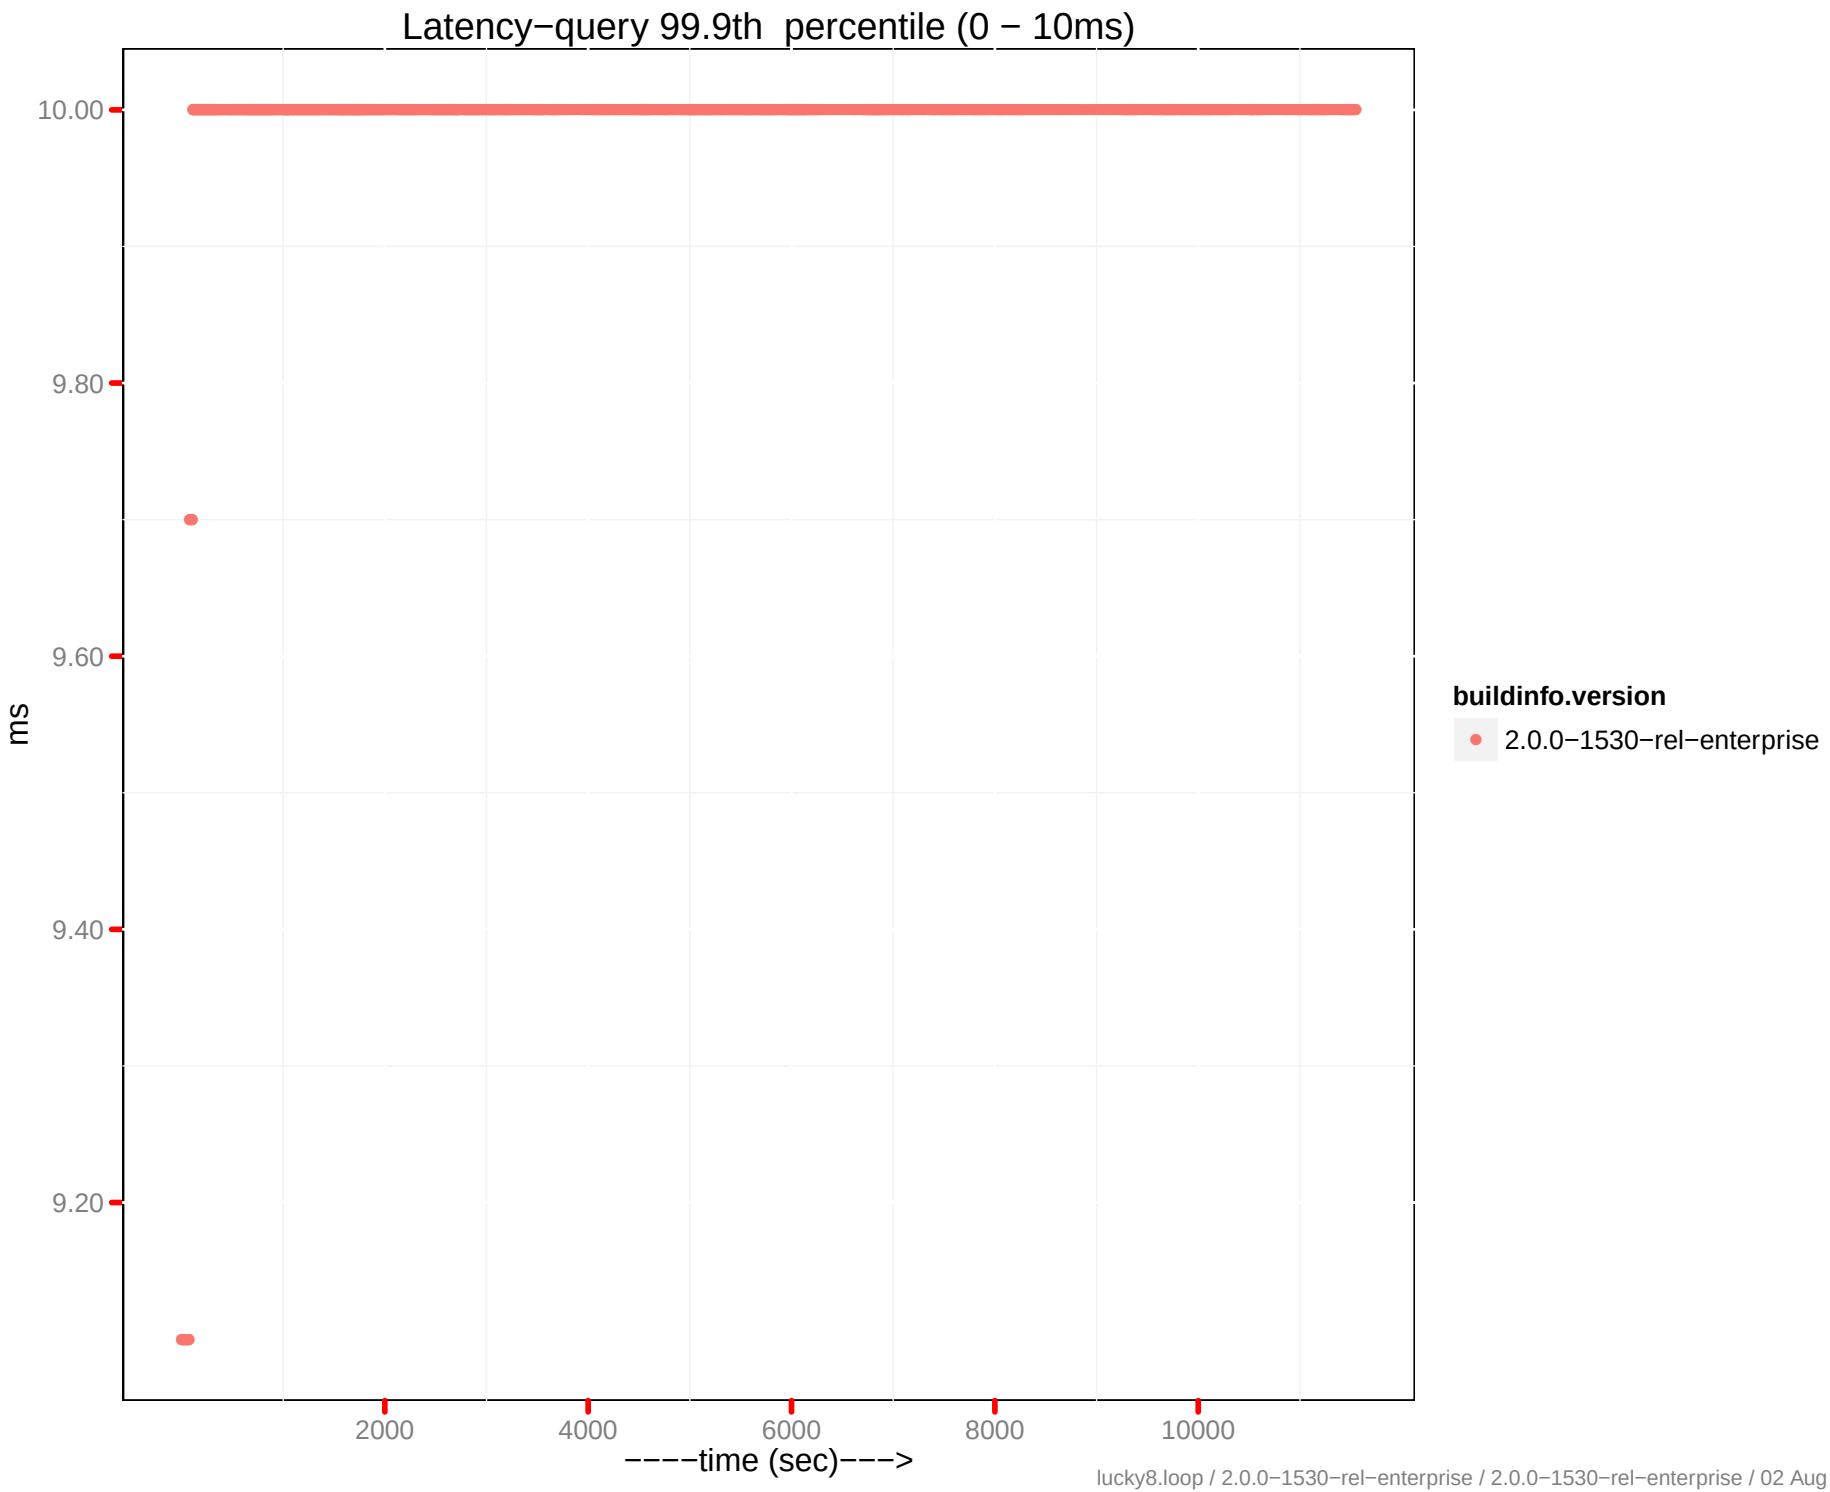
Latency−query 99.9th percentile (0 − 10ms)
10.00
9.80
9.60
9.40
9.20
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012


Query throughput
Queries/sec
1.46K
1000.00
500.00
0
0
2000
4000
6000
8000
10000
12000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1530−rel−enterprise
lucky8.loop / 2.0.0−1530−rel−enterprise / 2.0.0−1530−rel−enterprise / 02 Aug 2012

```
lucky8.conf
# "lucky 8" view performance test:
# 8K ops/sec (background, cluster-wide)
# 80% reads, 20% write (12% updates/deletes, 8% inserts)
# 8M dataset
# Stop after 8M total queries
# 1 ddoc with 8 views

performance.iperf.MultiClientTests.test_vperf2

params:

# general
batch=50
kind=json
mem_quota=16000
loglevel=error

# load phase
items=8000000

# access phase
ratio_sets=0.2
ratio_misses=0.04
ratio_creates=0.40
ratio_deletes=0.50
ratio_hot=0.2
ratio_hot_gets=0.95
ratio_hot_sets=0.95
ratio_expirations=0.0
bg_max_ops_per_sec=1000
fg_max_ops=8000000
total_clients=8

# control (defaults: pytests/performance/perf_defaults.py)
load_wait_until_drained=1
loop_wait_until_drained=0
mcsoda_heartbeat=3
tear_down=1
tear_down_proxy=1
tear_down_bucket=0
tear_down_cluster=1
tear_down_on_setup=0
```

terra.ini

```
[global]
username:root
password:couchbase
port:8091
data_path:/data

[servers]
1:10.2.1.61
2:10.2.1.62
3:10.2.1.63
4:10.2.1.64

[clients]
1:10.2.1.59

[membase]
rest_username:Administrator
rest_password:password

[dashboard]
1:dashboard.hq.couchbase.com:80
```