	2.0.0 − 1692	2.0.0 − 1692
Runtime (in hr)	17.14	NA
Avg. Drain Rate	264.12	NANA
Peak Disk (GB)	225.28	NA
Peak Memory (GB)	20.16	NA
Avg. OPS	7.43K	NANA
Avg. mem memcached (GB)	19.69	NA
Avg. mem beam.smp (MB)	354.36	NA
Avg. CPU rate (%)	26.3	NA
Latency−get (90th) (ms)	1.29	NA
Latency−get (95th) (ms)	2.38	NA
Latency−get (99th) (ms)	7.68	NA
Latency−set (90th) (ms)	1.39	NA
Latency−set (95th) (ms)	2.41	NA
Latency−set (99th) (ms)	5.38	NA
Latency−query (80th) (ms)	183.09	NA
Latency−query (90th) (ms)	334.41	NA
Latency−query (95th) (ms)	478.58	NA
Latency−query (99th) (ms)	1059.87	NA
Latency−query (99.9th) (ms)	2101.25	NA
Avg. QPS	185.8	NA
XDC ops/sec	NaN	NA
Rebalance Time (sec)	0	NA
Testrunner Version	f3ae21f	NA

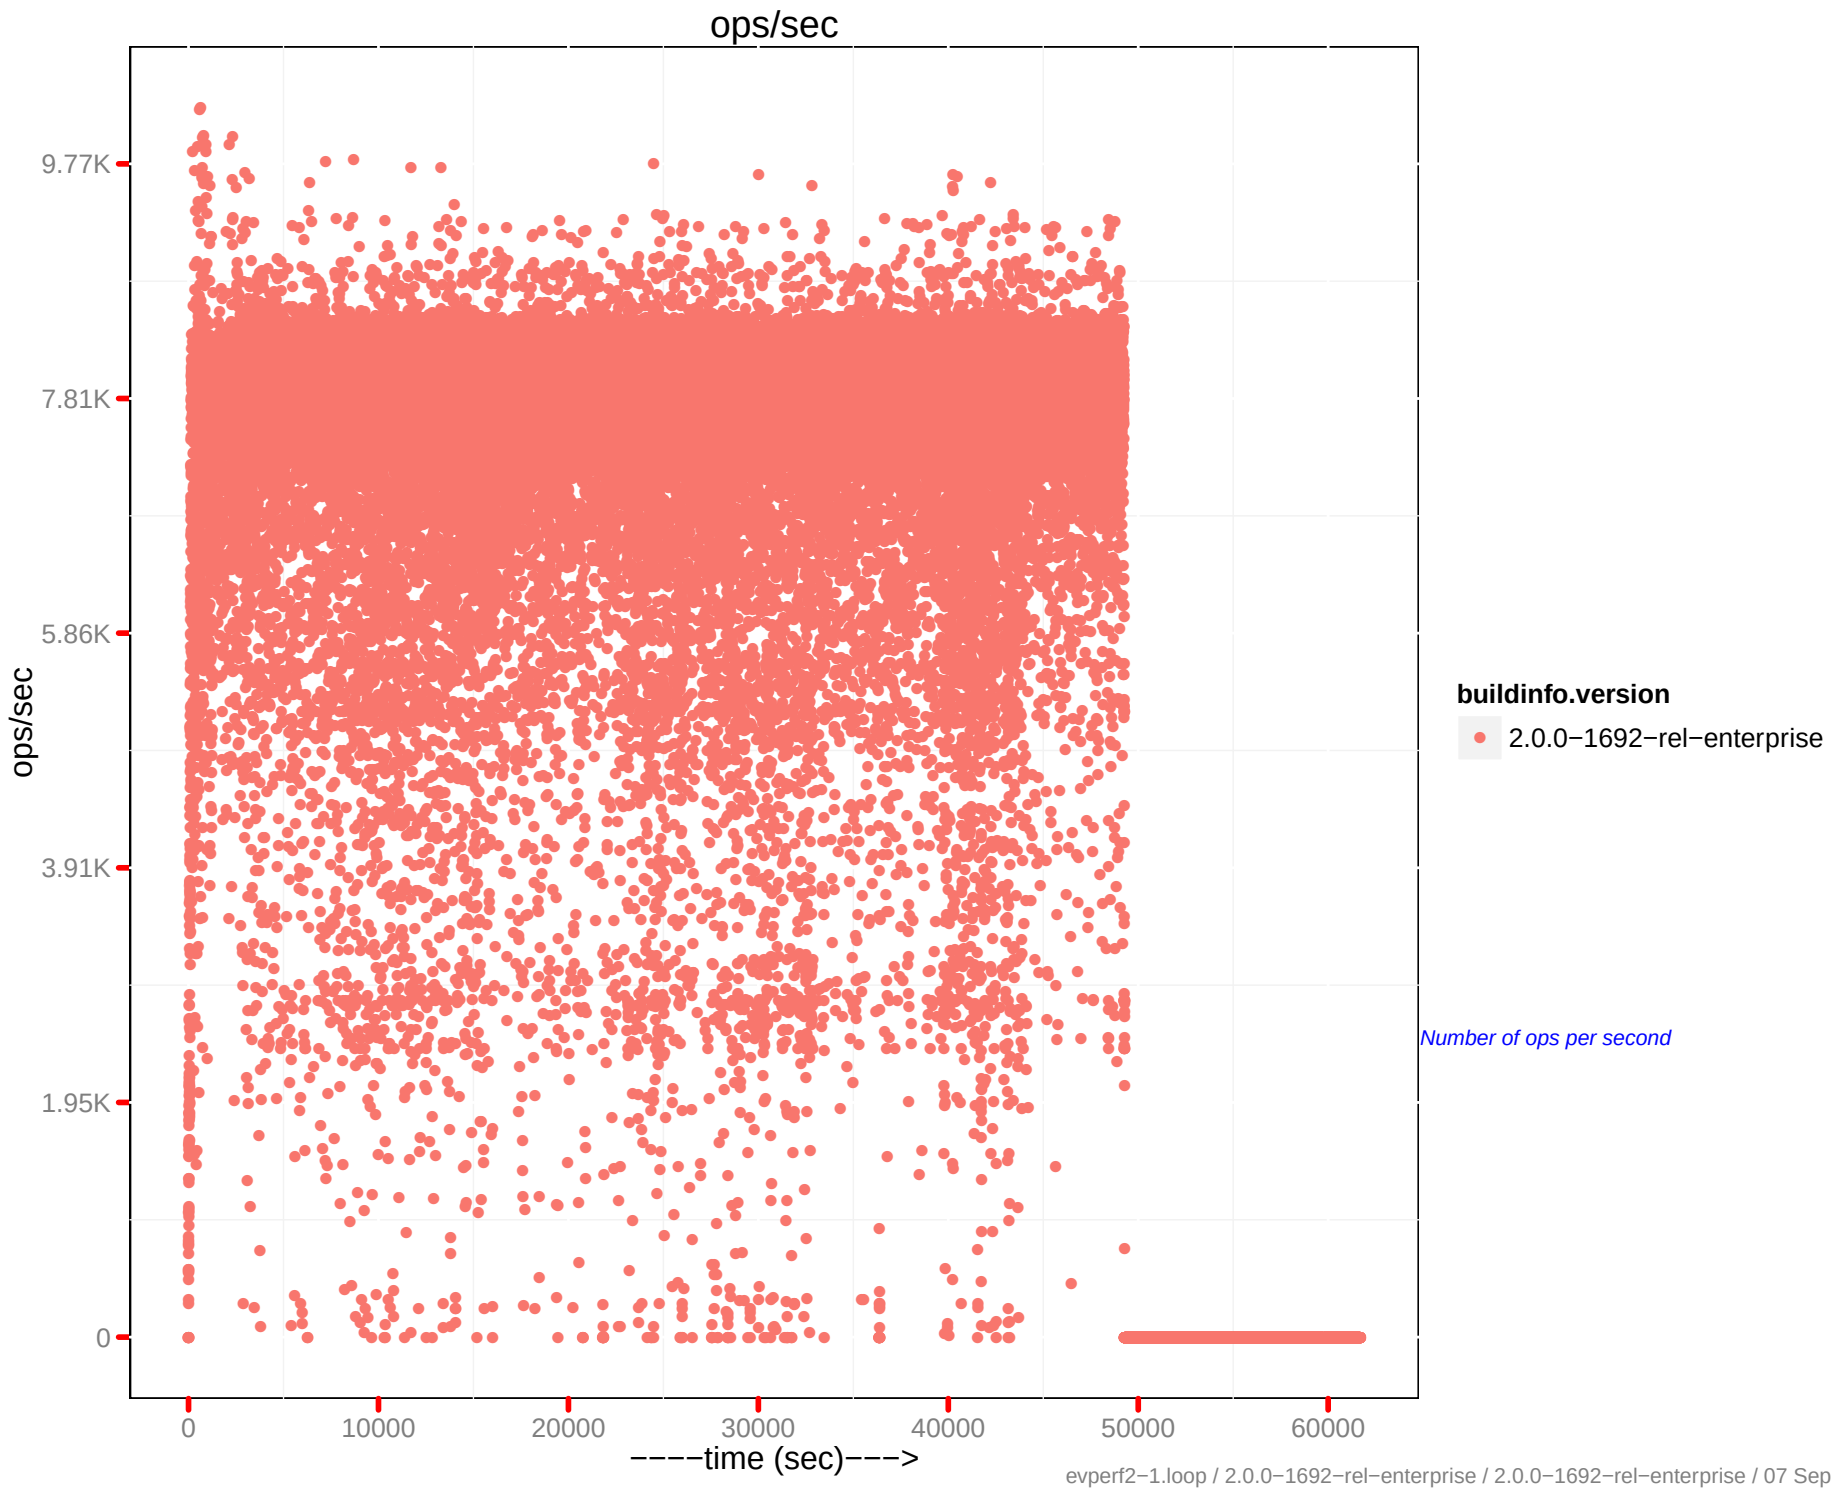
ops/sec
9.77K
7.81K
5.86K
3.91K
1.95K
0
ops/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of ops per second
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

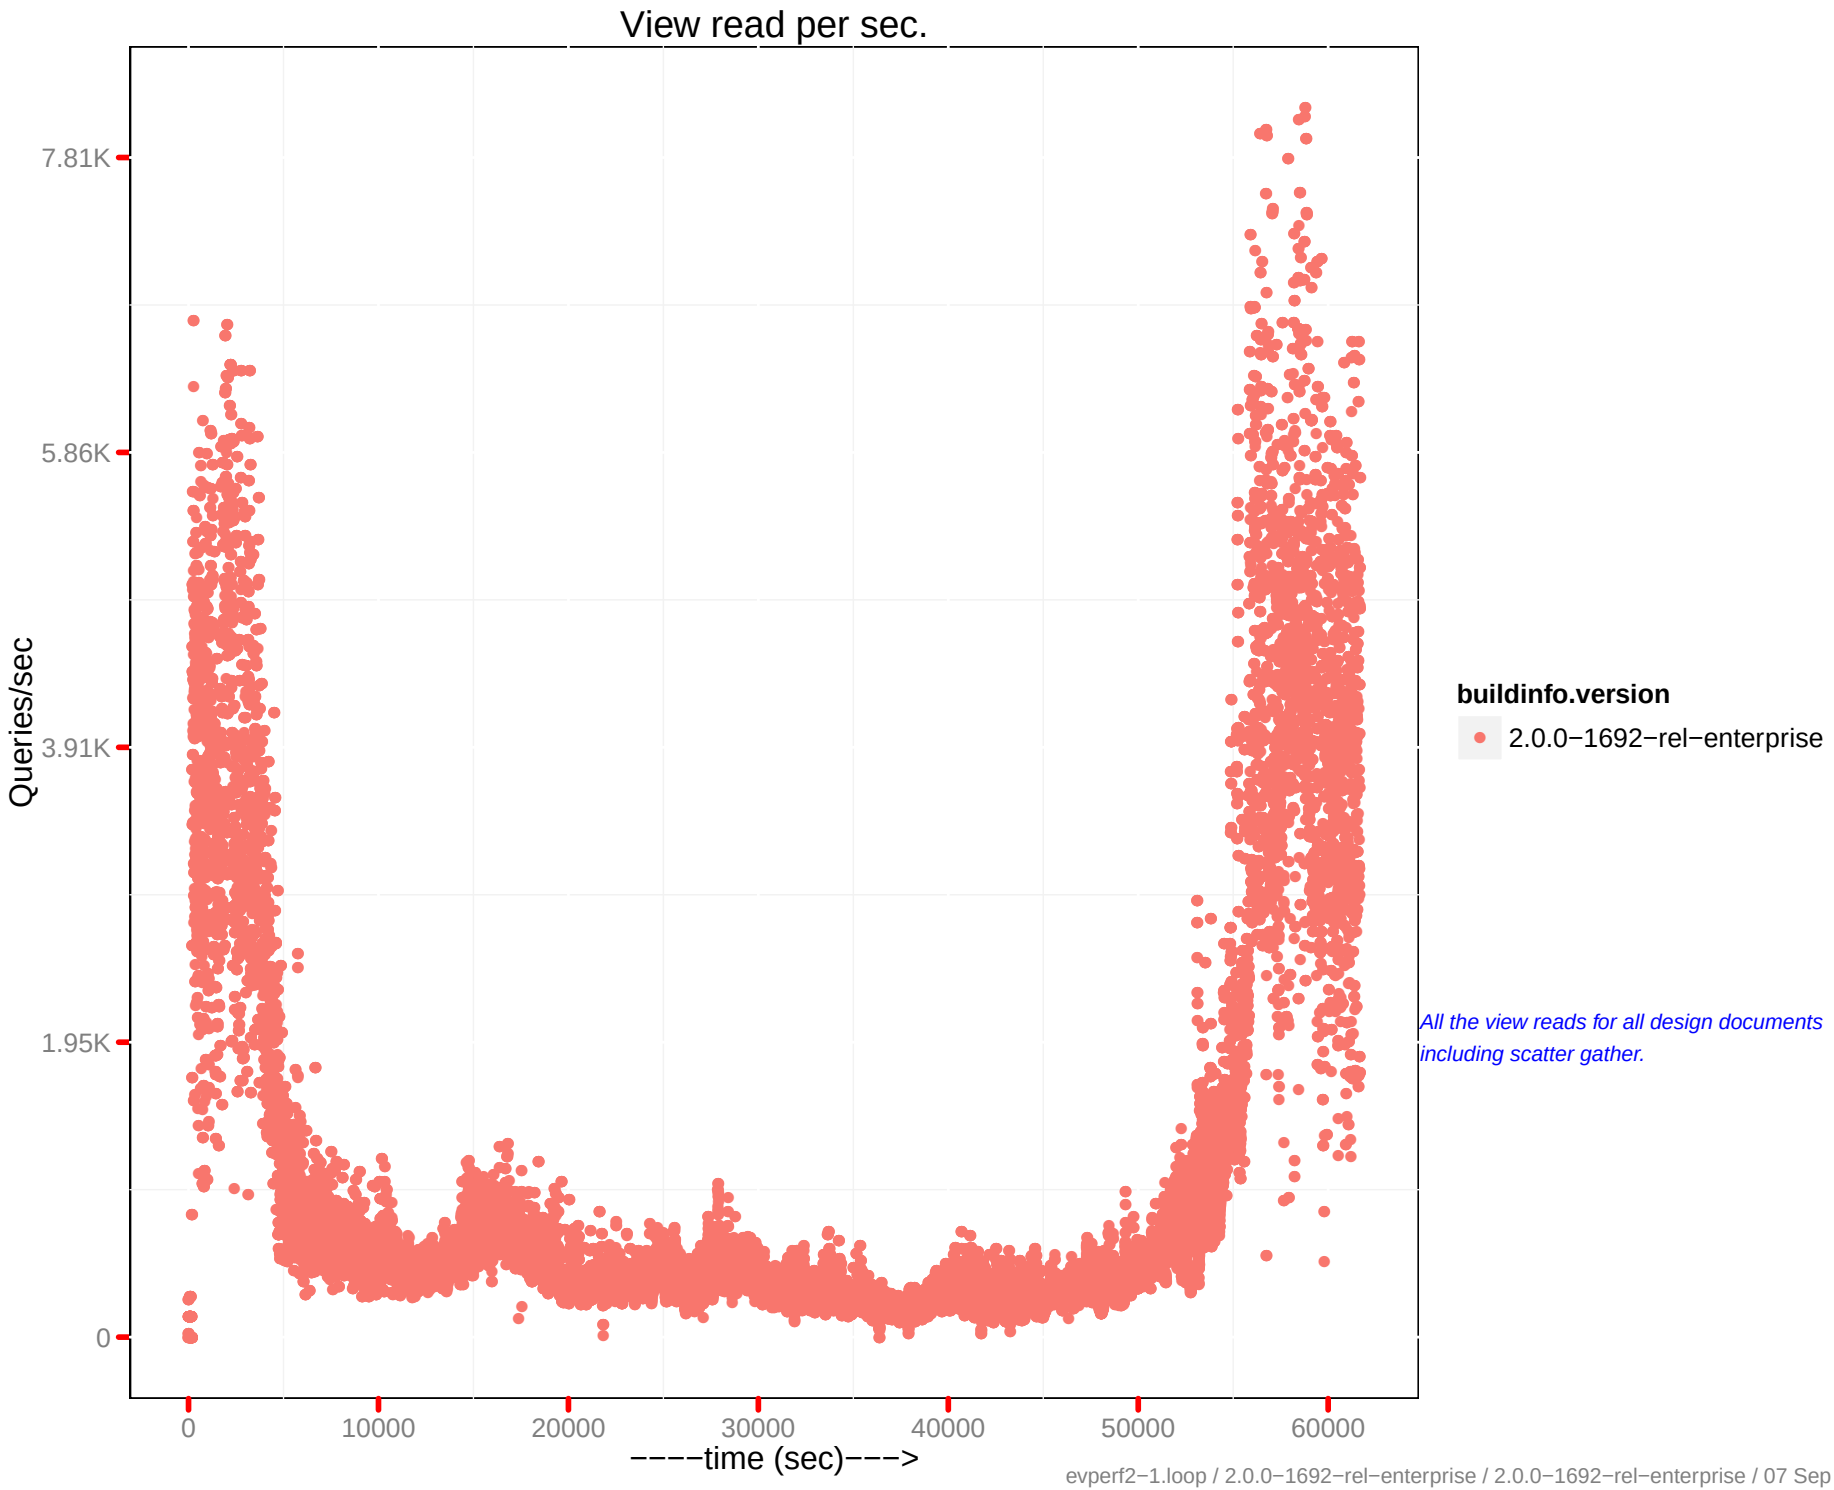
View read per sec.
7.81K
5.86K
3.91K
1.95K
0
Queries/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
All the view reads for all design documents including scatter gather.
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

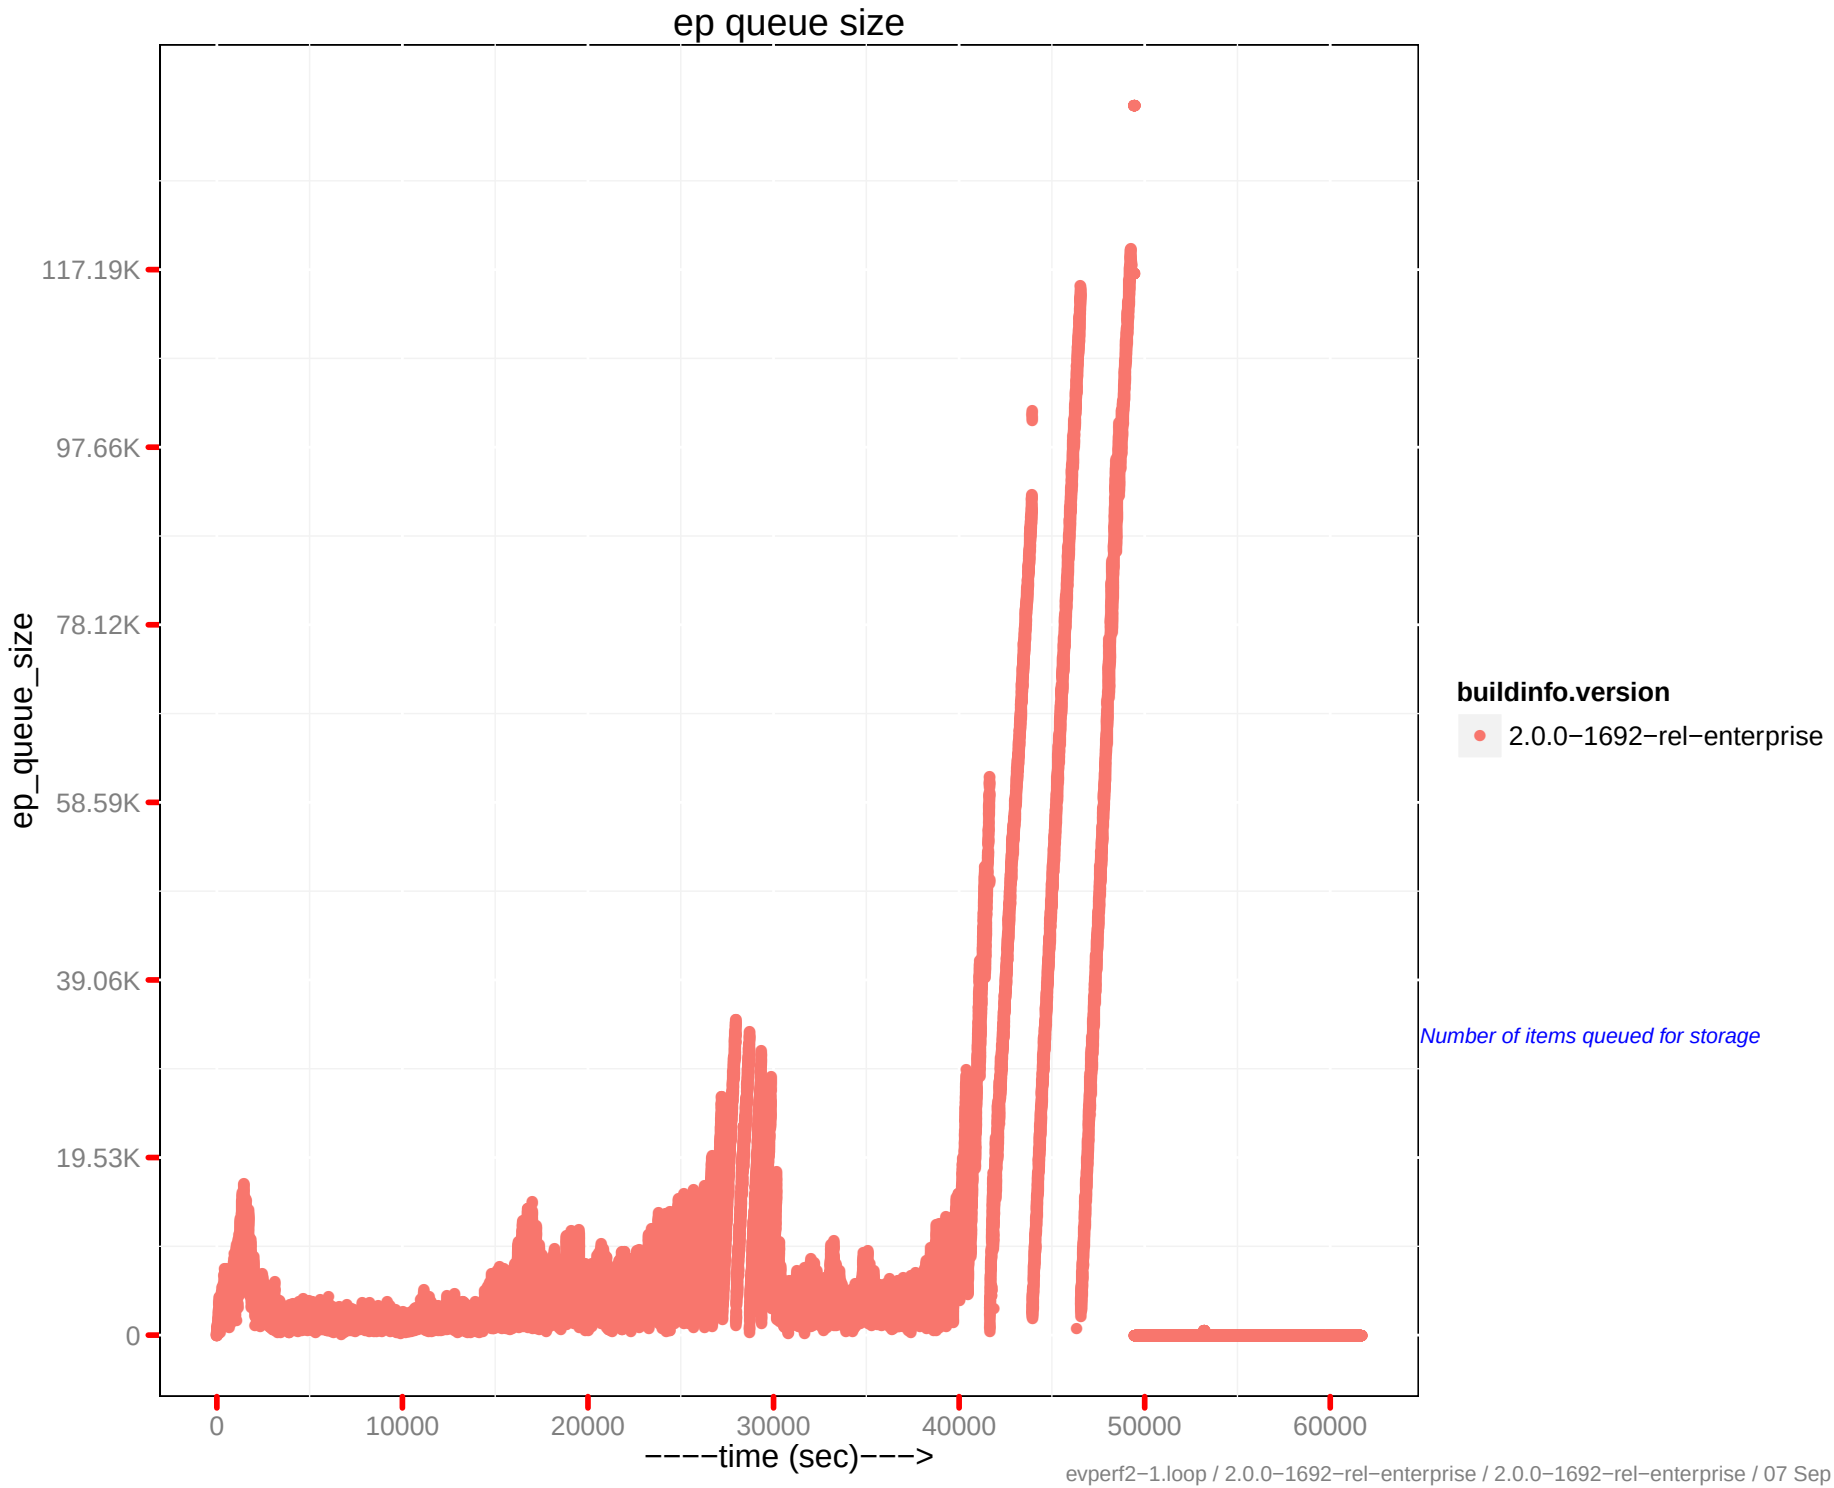
ep queue size
ep_queue_size
117.19K
97.66K
78.12K
58.59K
39.06K
19.53K
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of items queued for storage
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

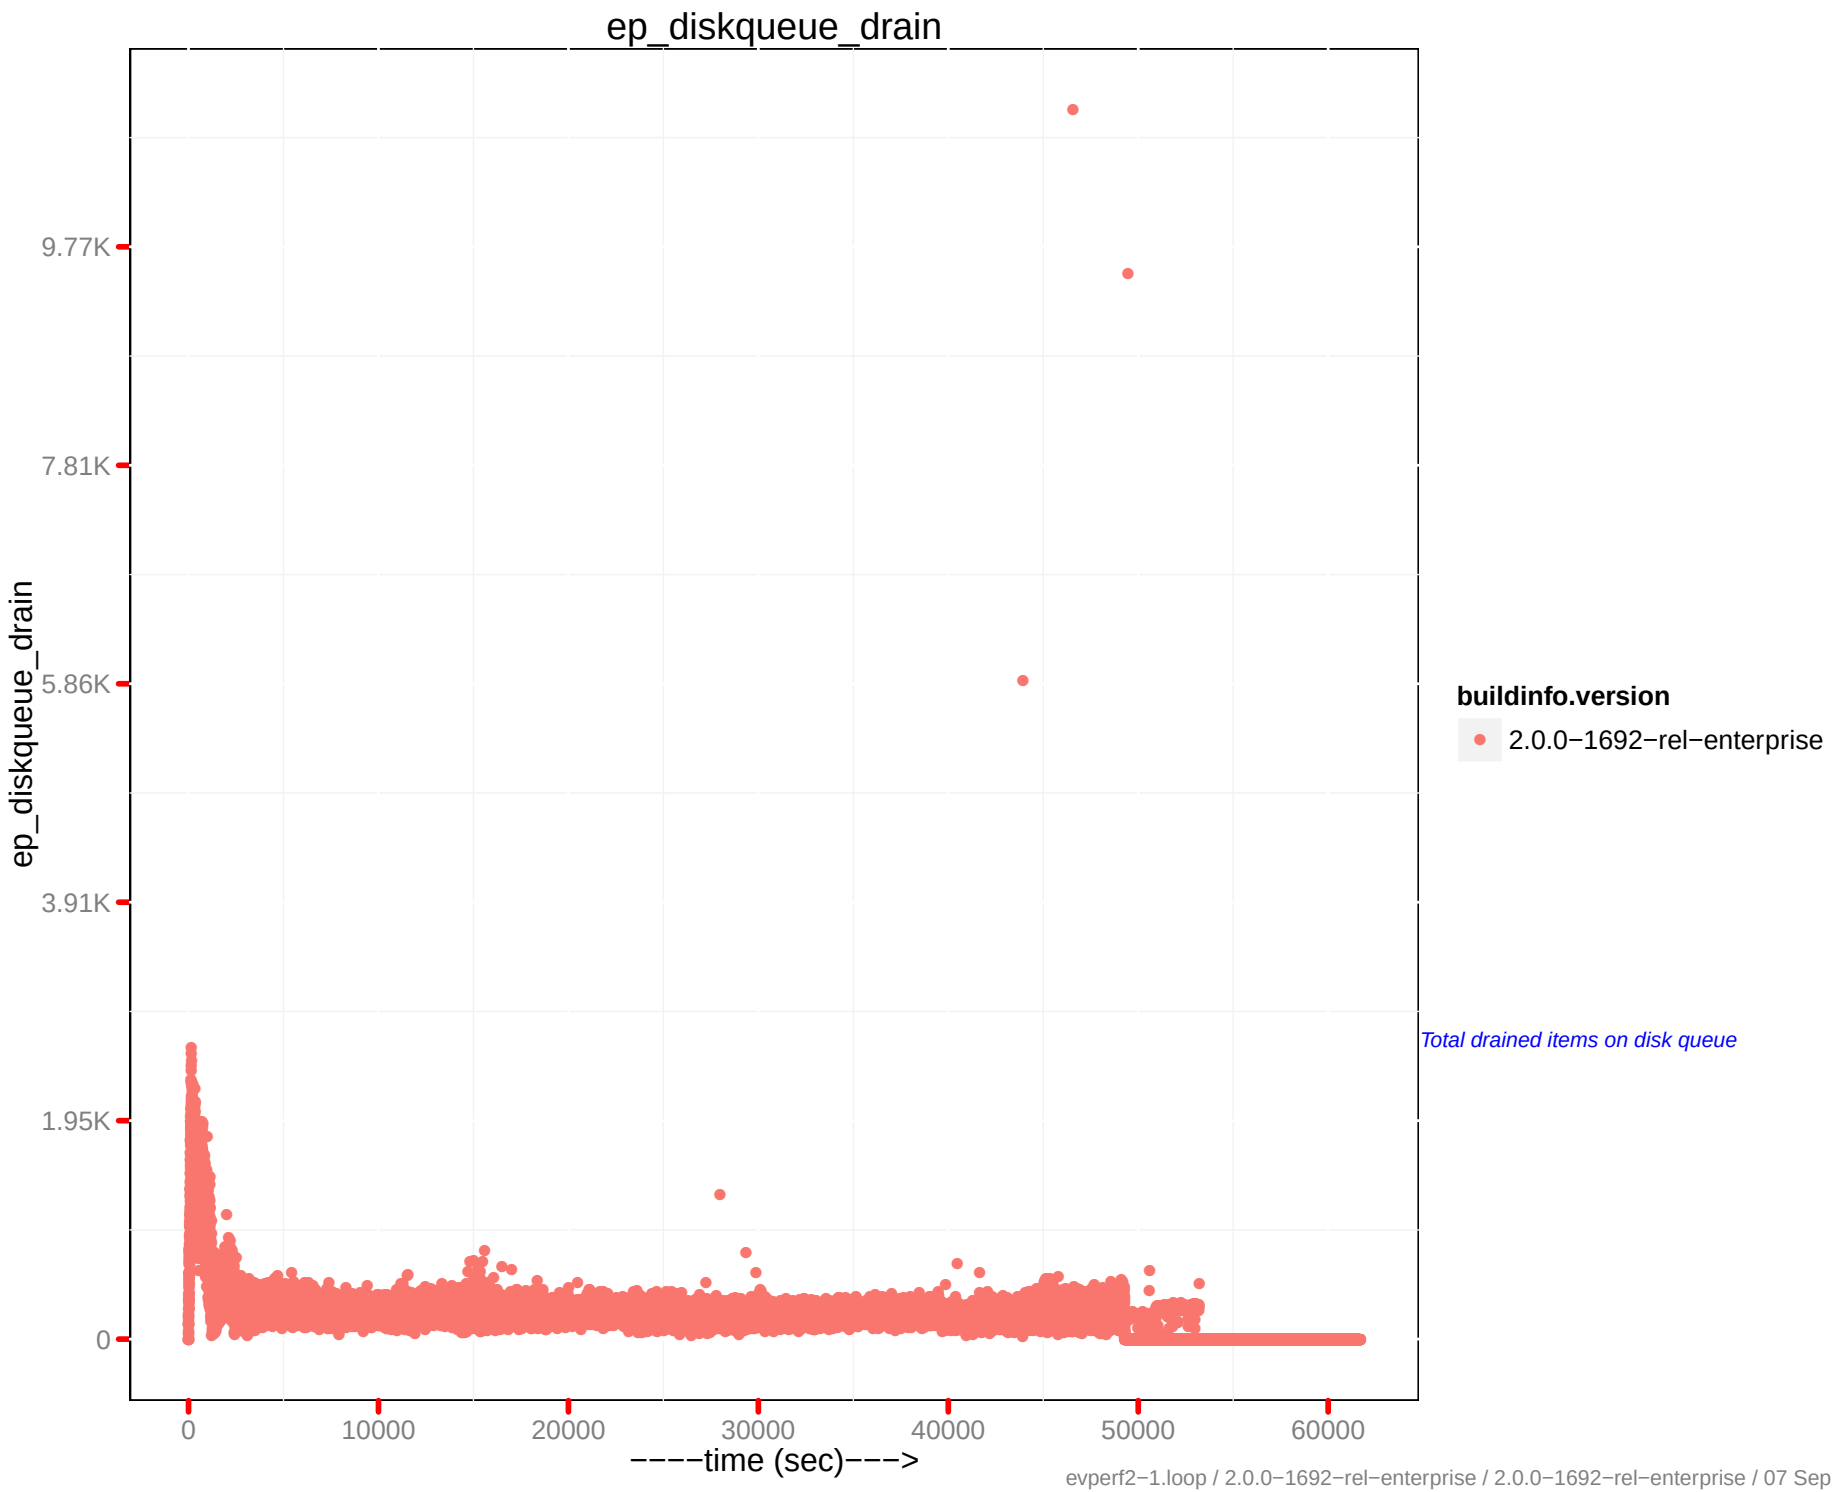
ep_diskqueue_drain
9.77K
7.81K
5.86K
3.91K
1.95K
0
ep_diskqueue_drain
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Total drained items on disk queue
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

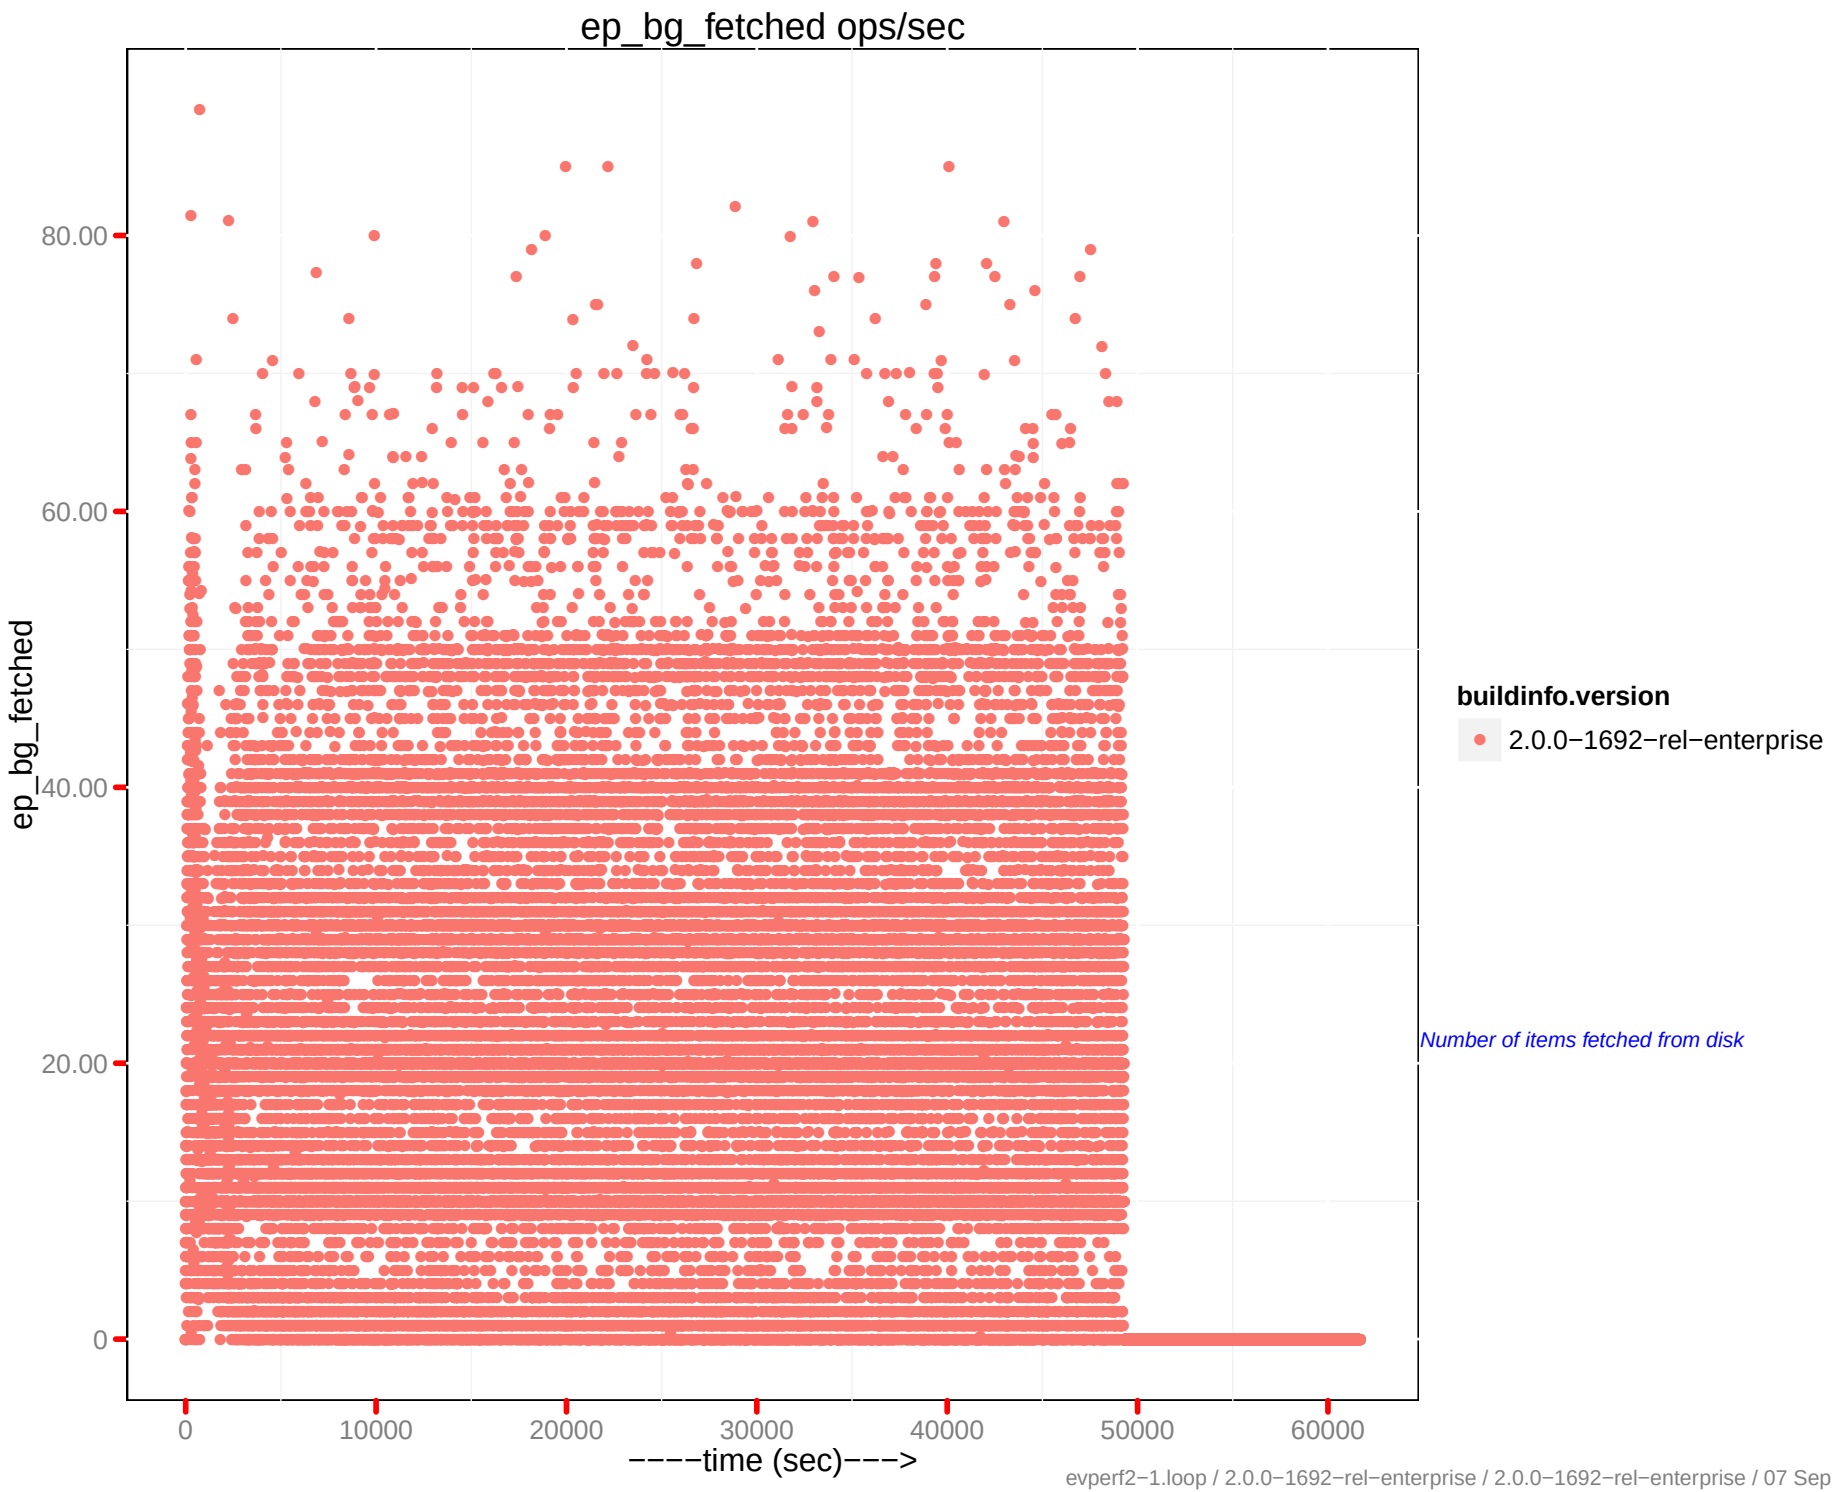
ep_bg_fetched ops/sec
ep_bg_fetched
80.00
60.00
40.00
20.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of items fetched from disk
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

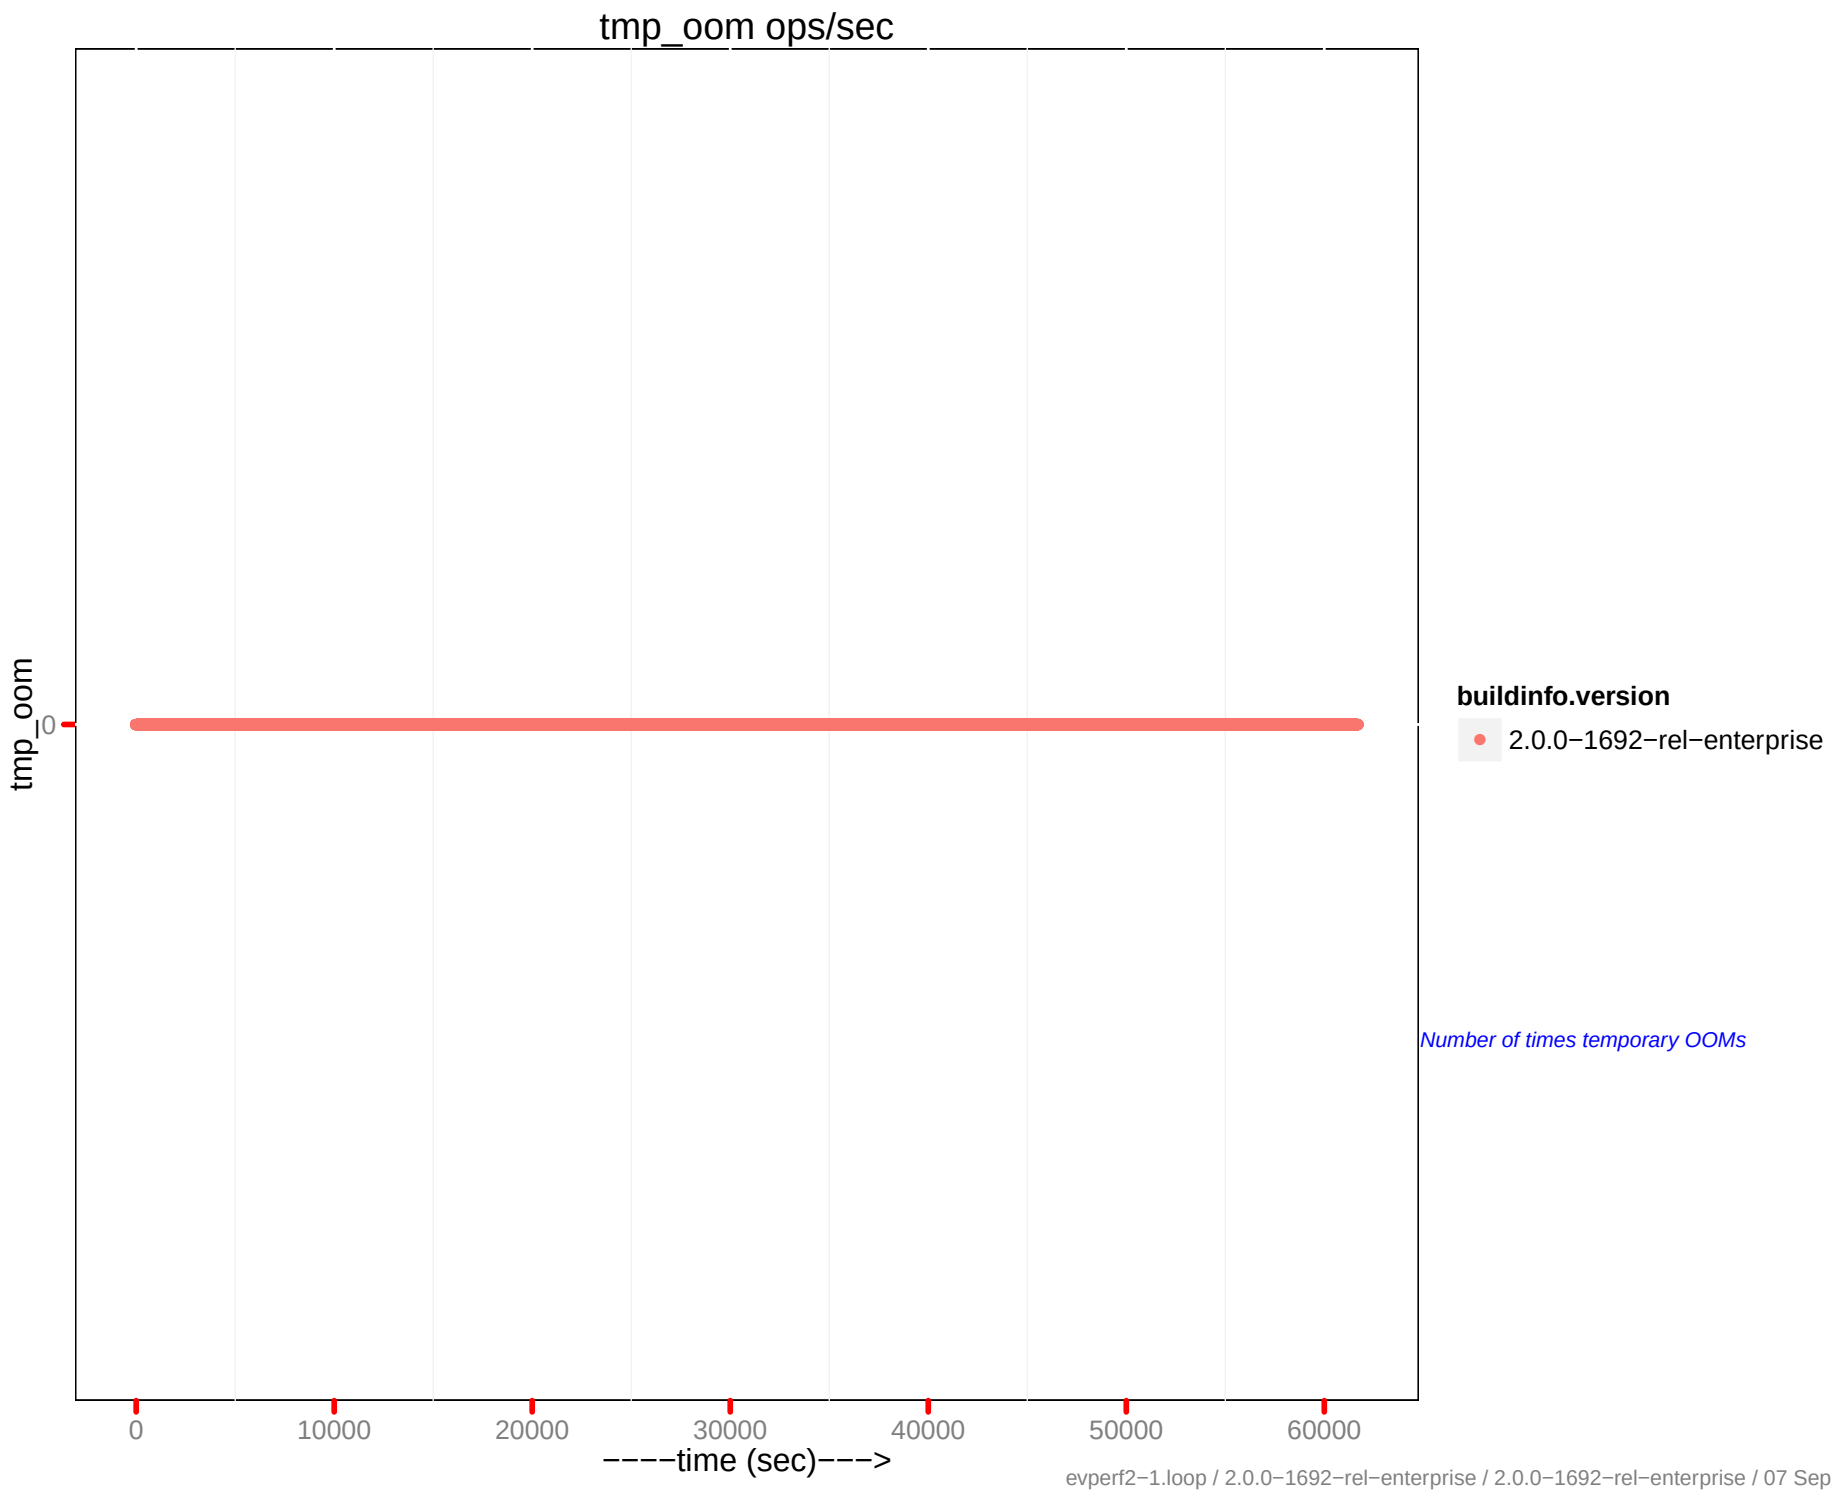
tmp_oom ops/sec
tmp_oom
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of times temporary OOMs
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

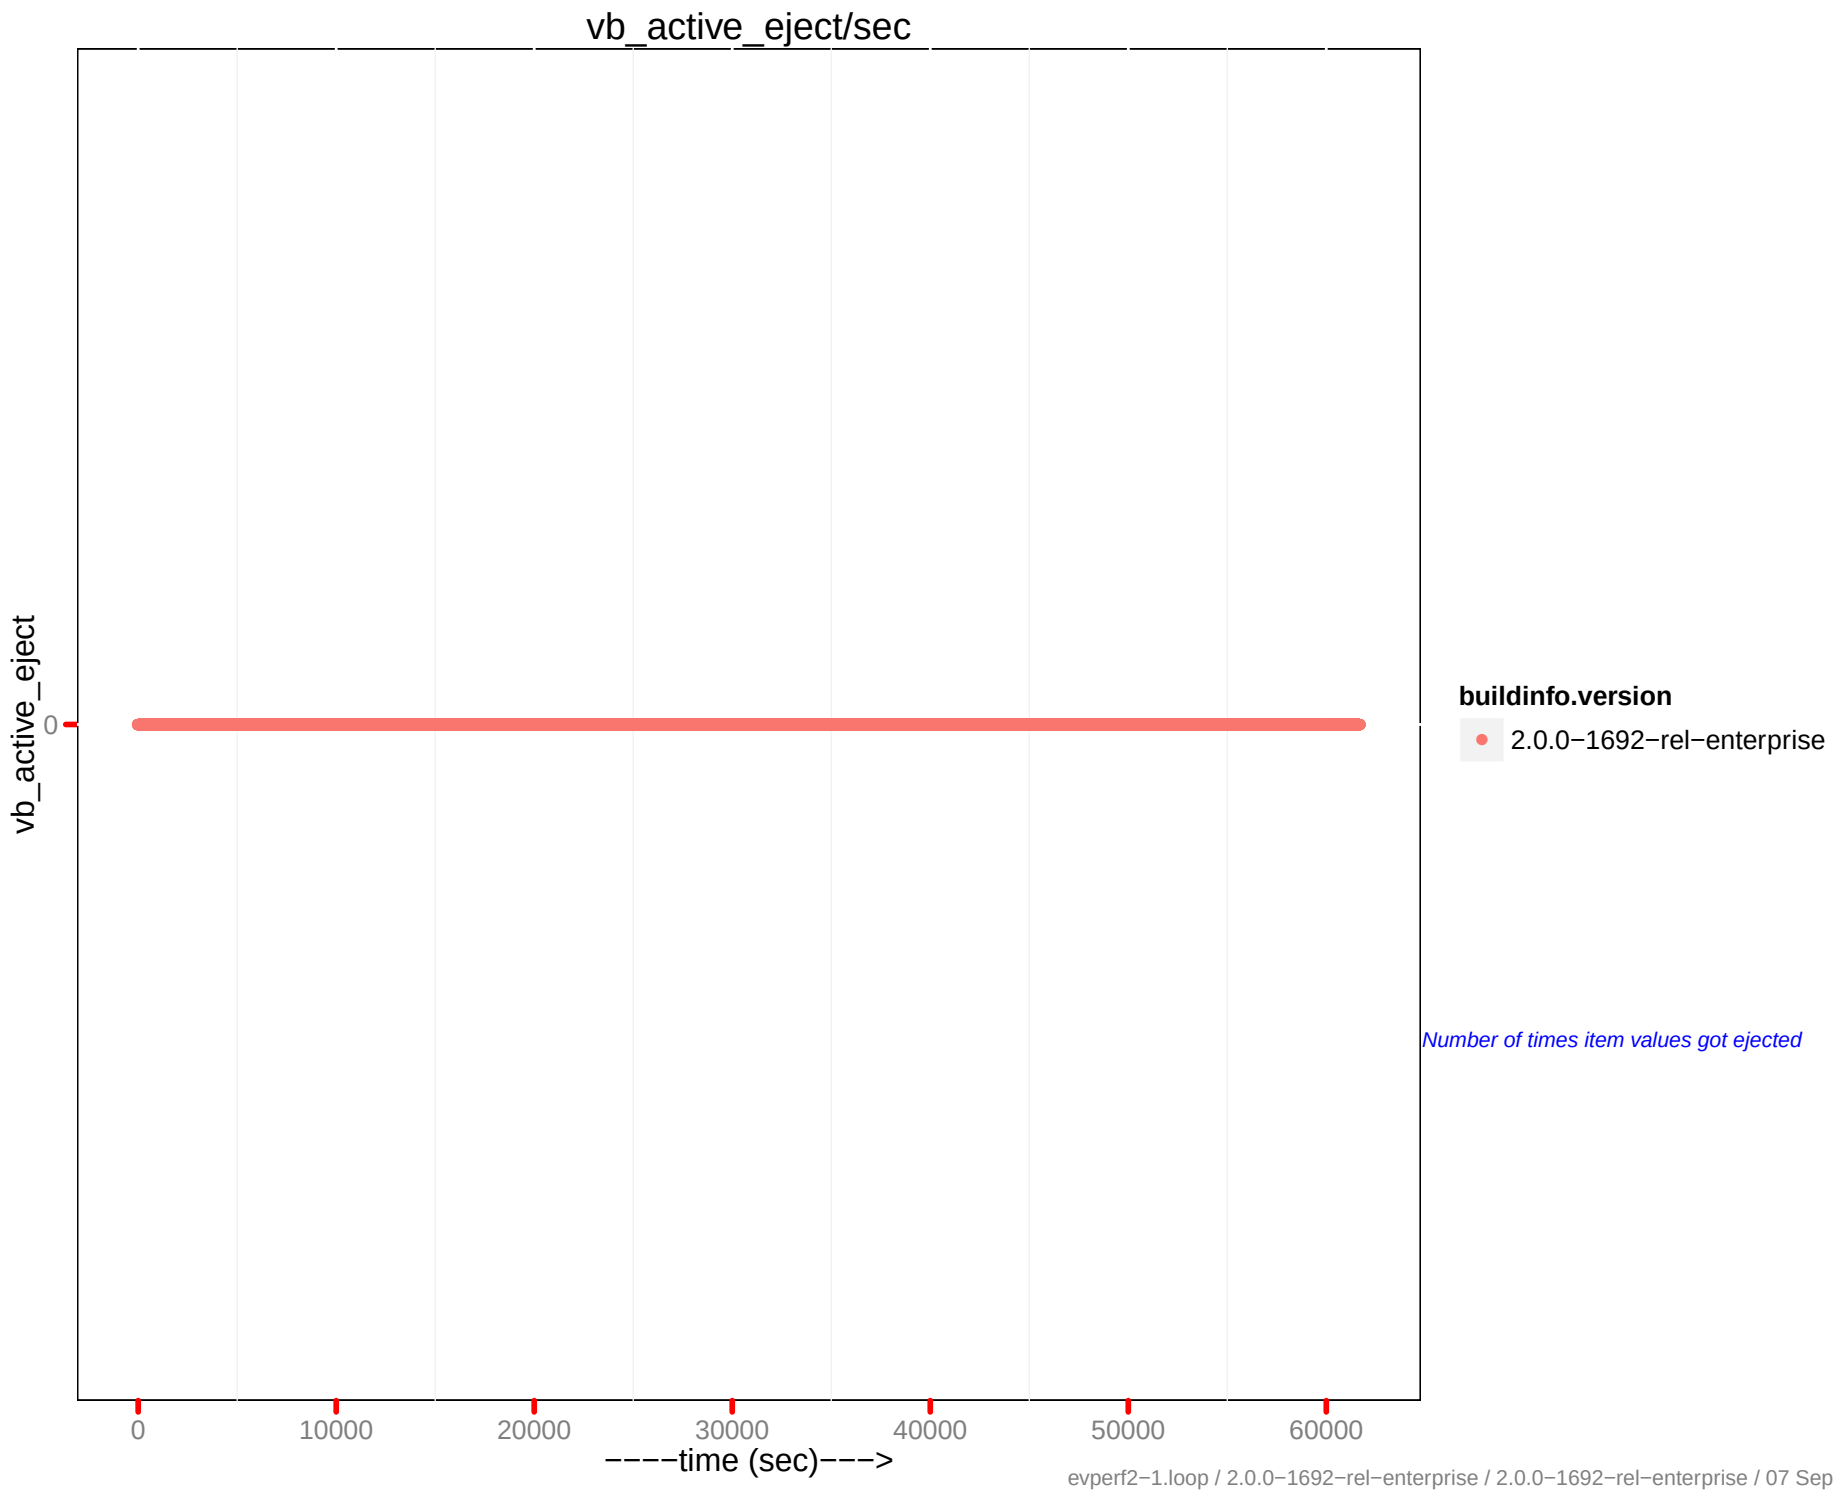
vb_active_eject/sec
vb_active_eject
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of times item values got ejected
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

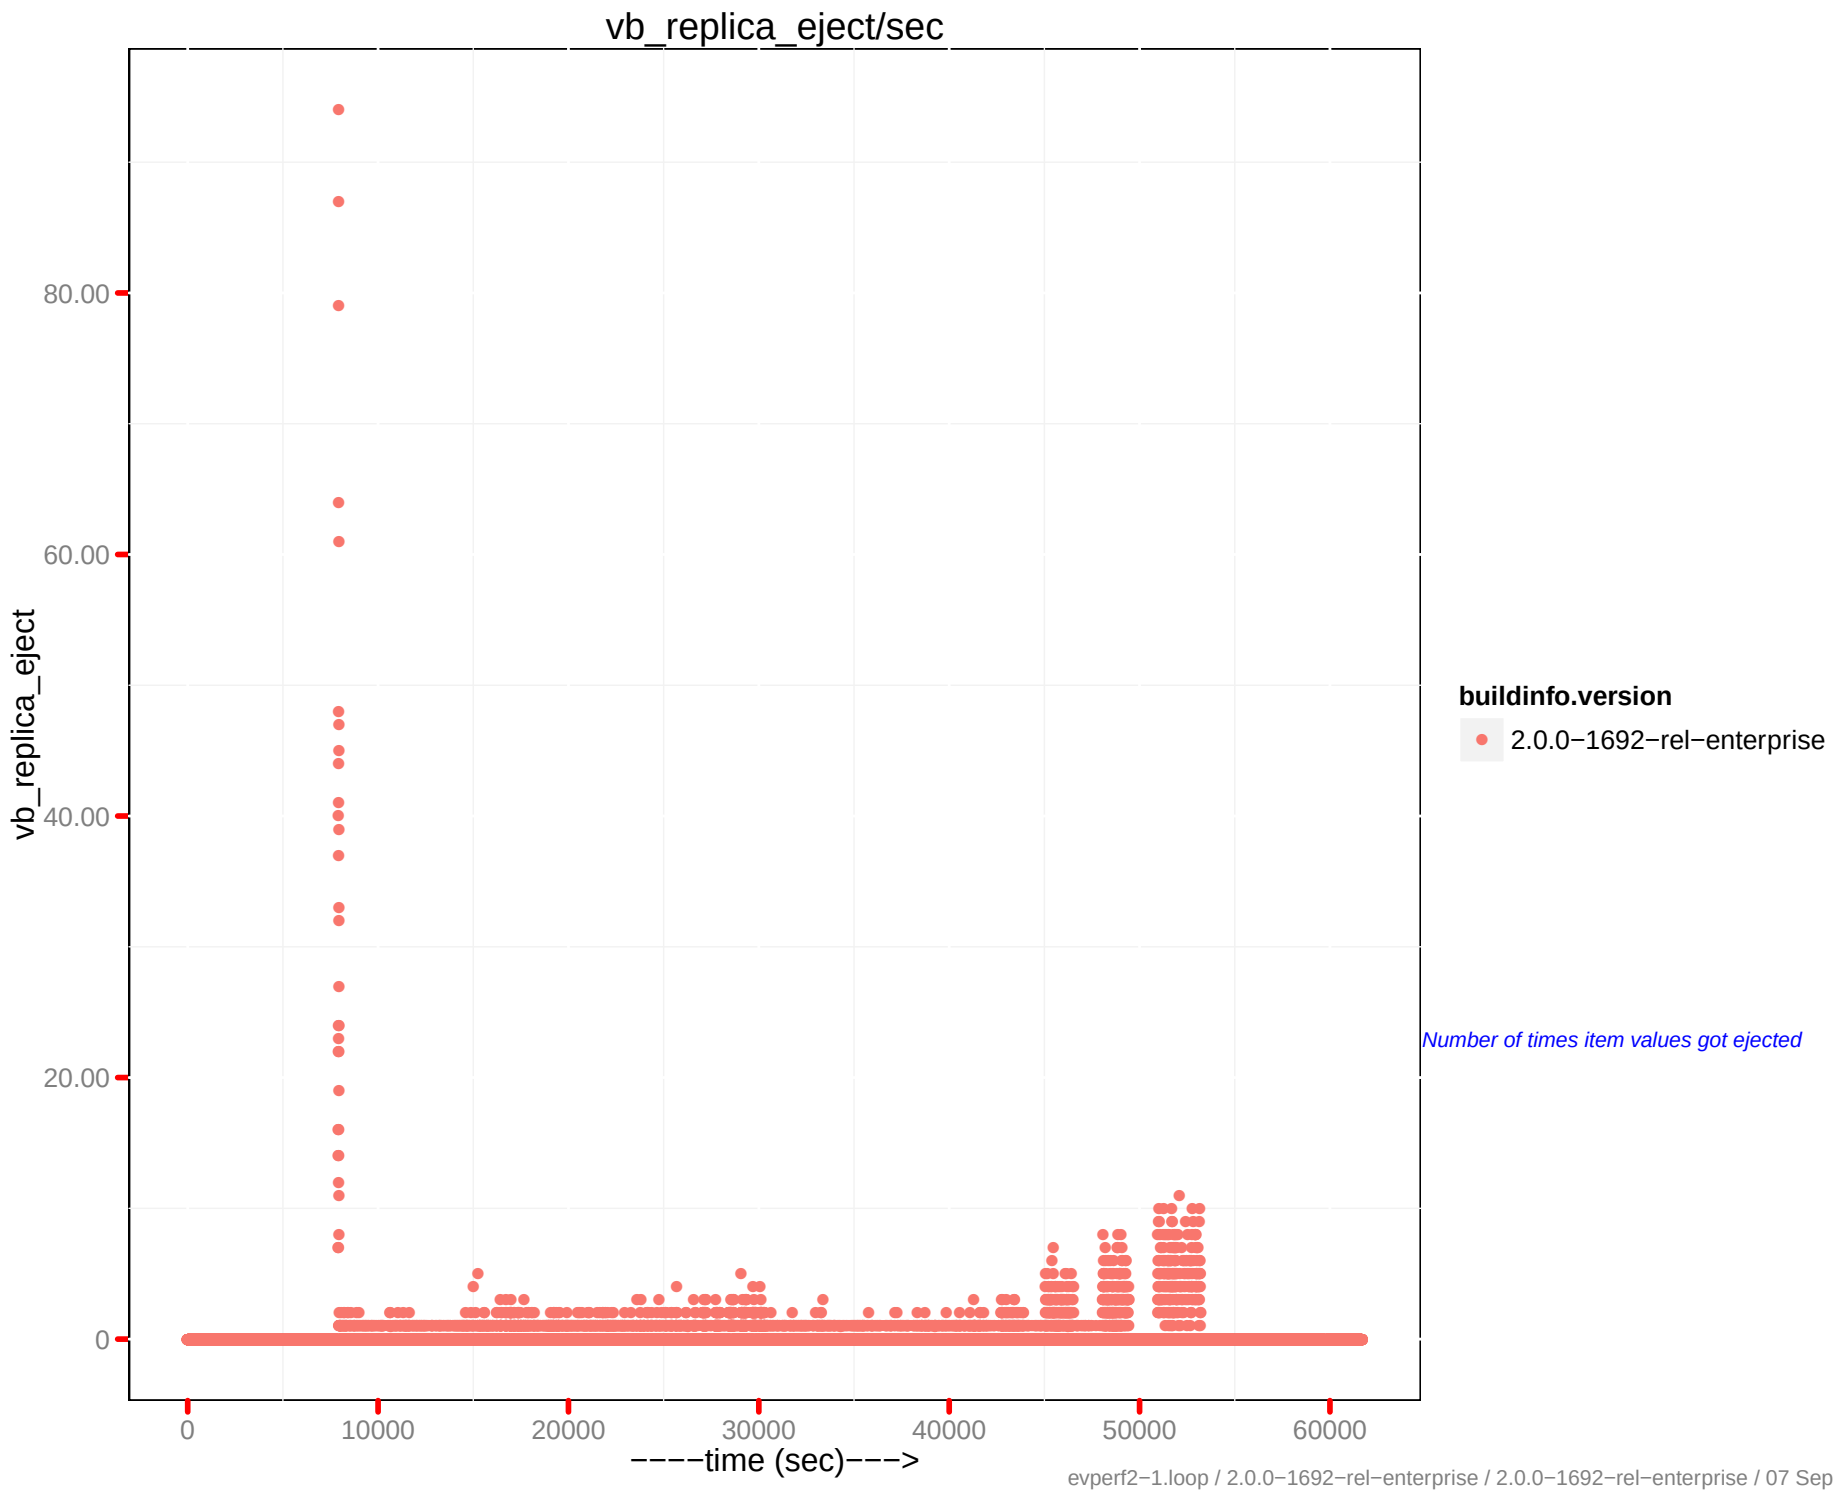
vb_replica_eject/sec
vb_replica_eject
80.00
60.00
40.00
20.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of times item values got ejected
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

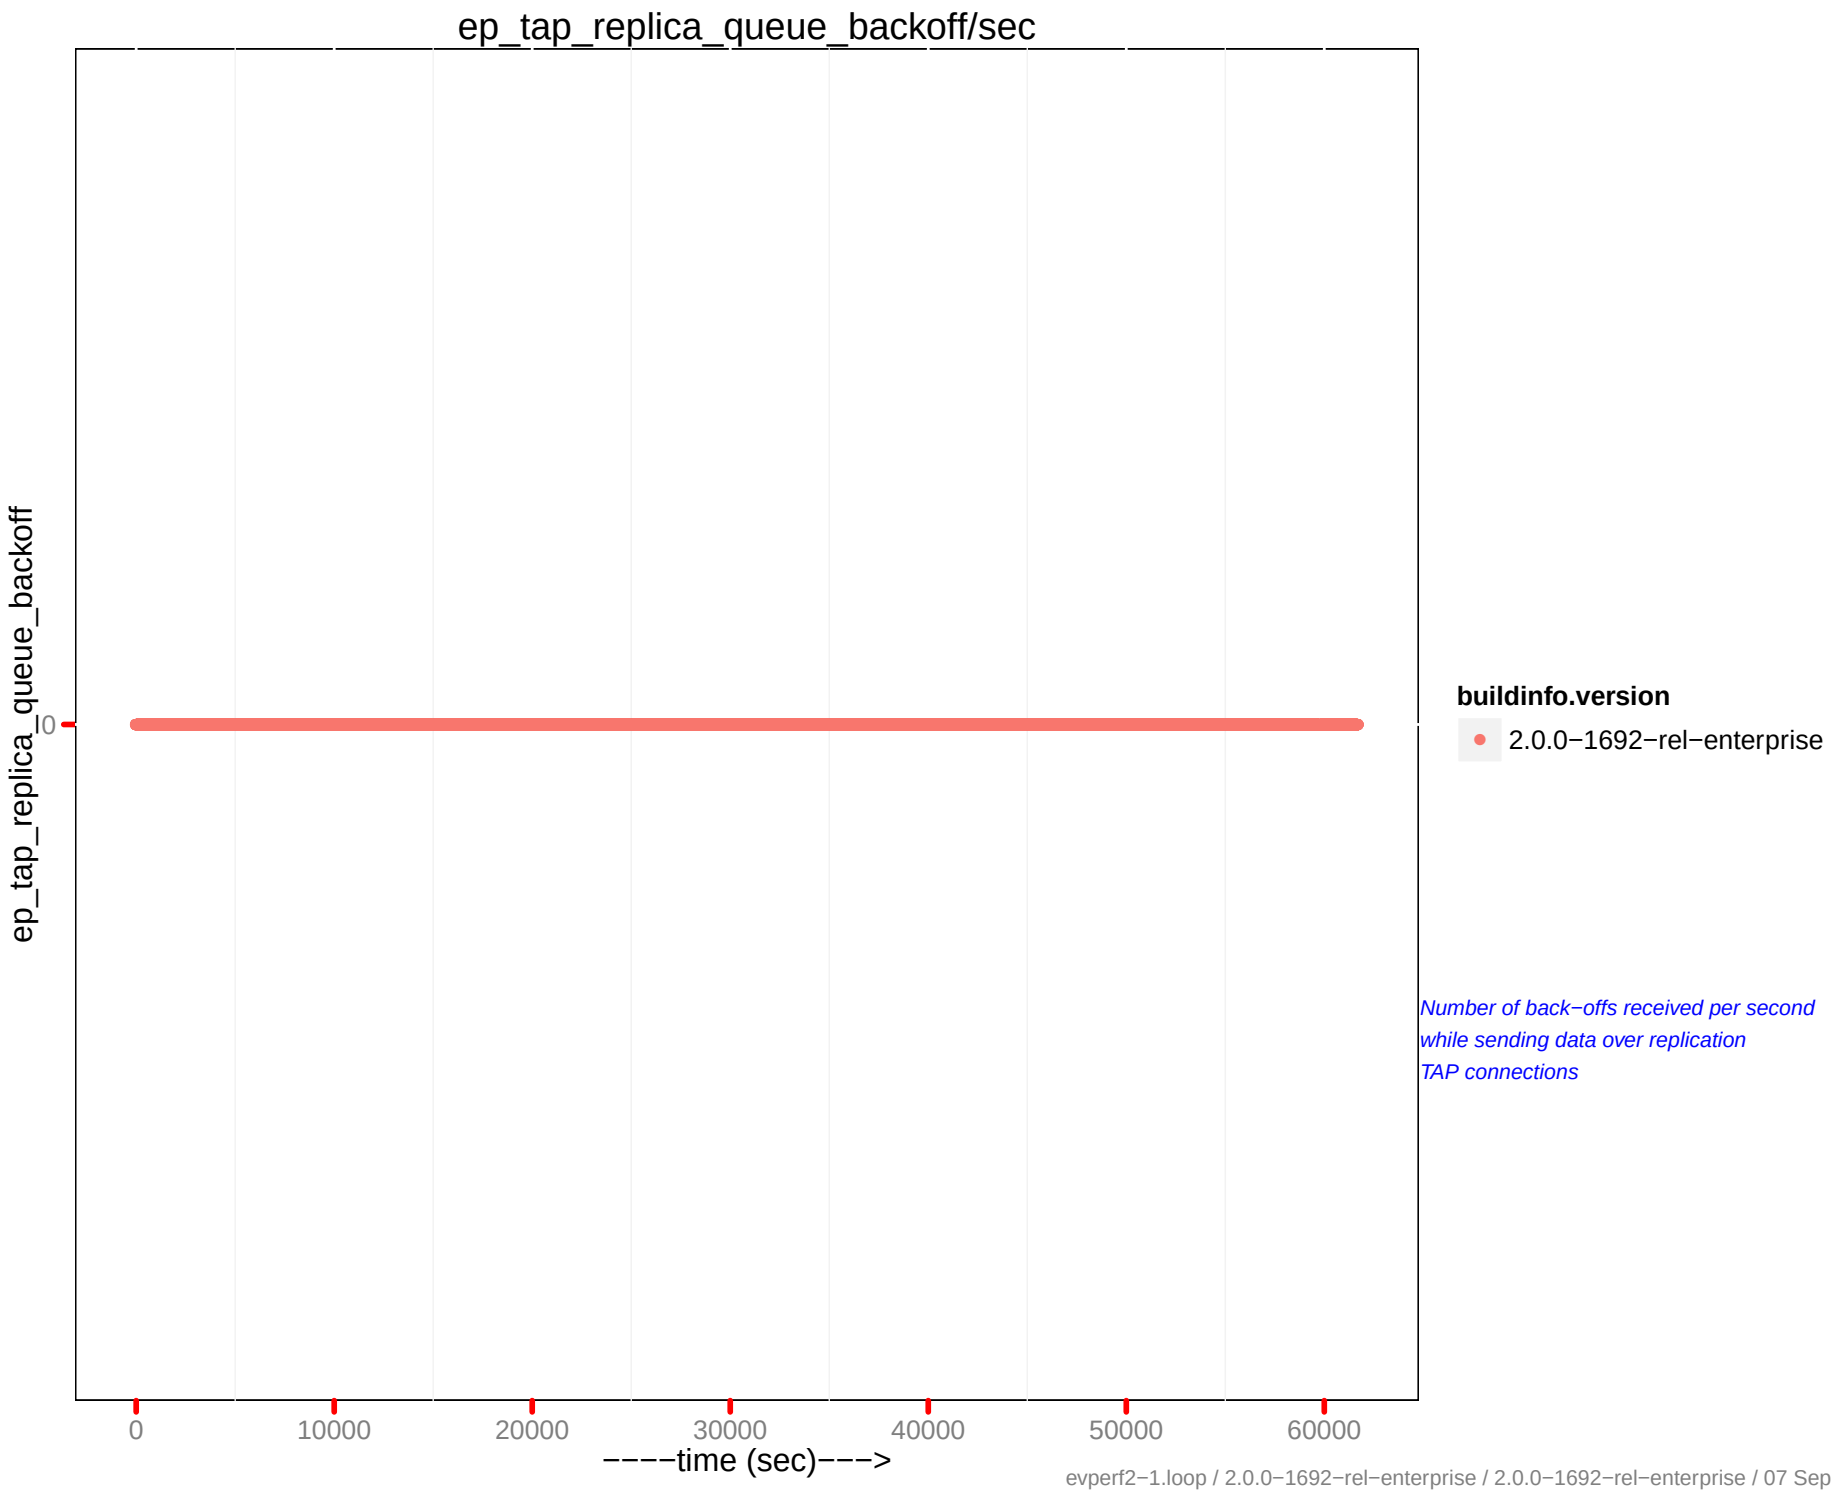
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of back-offs received per second
while sending data over replication
TAP connections
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

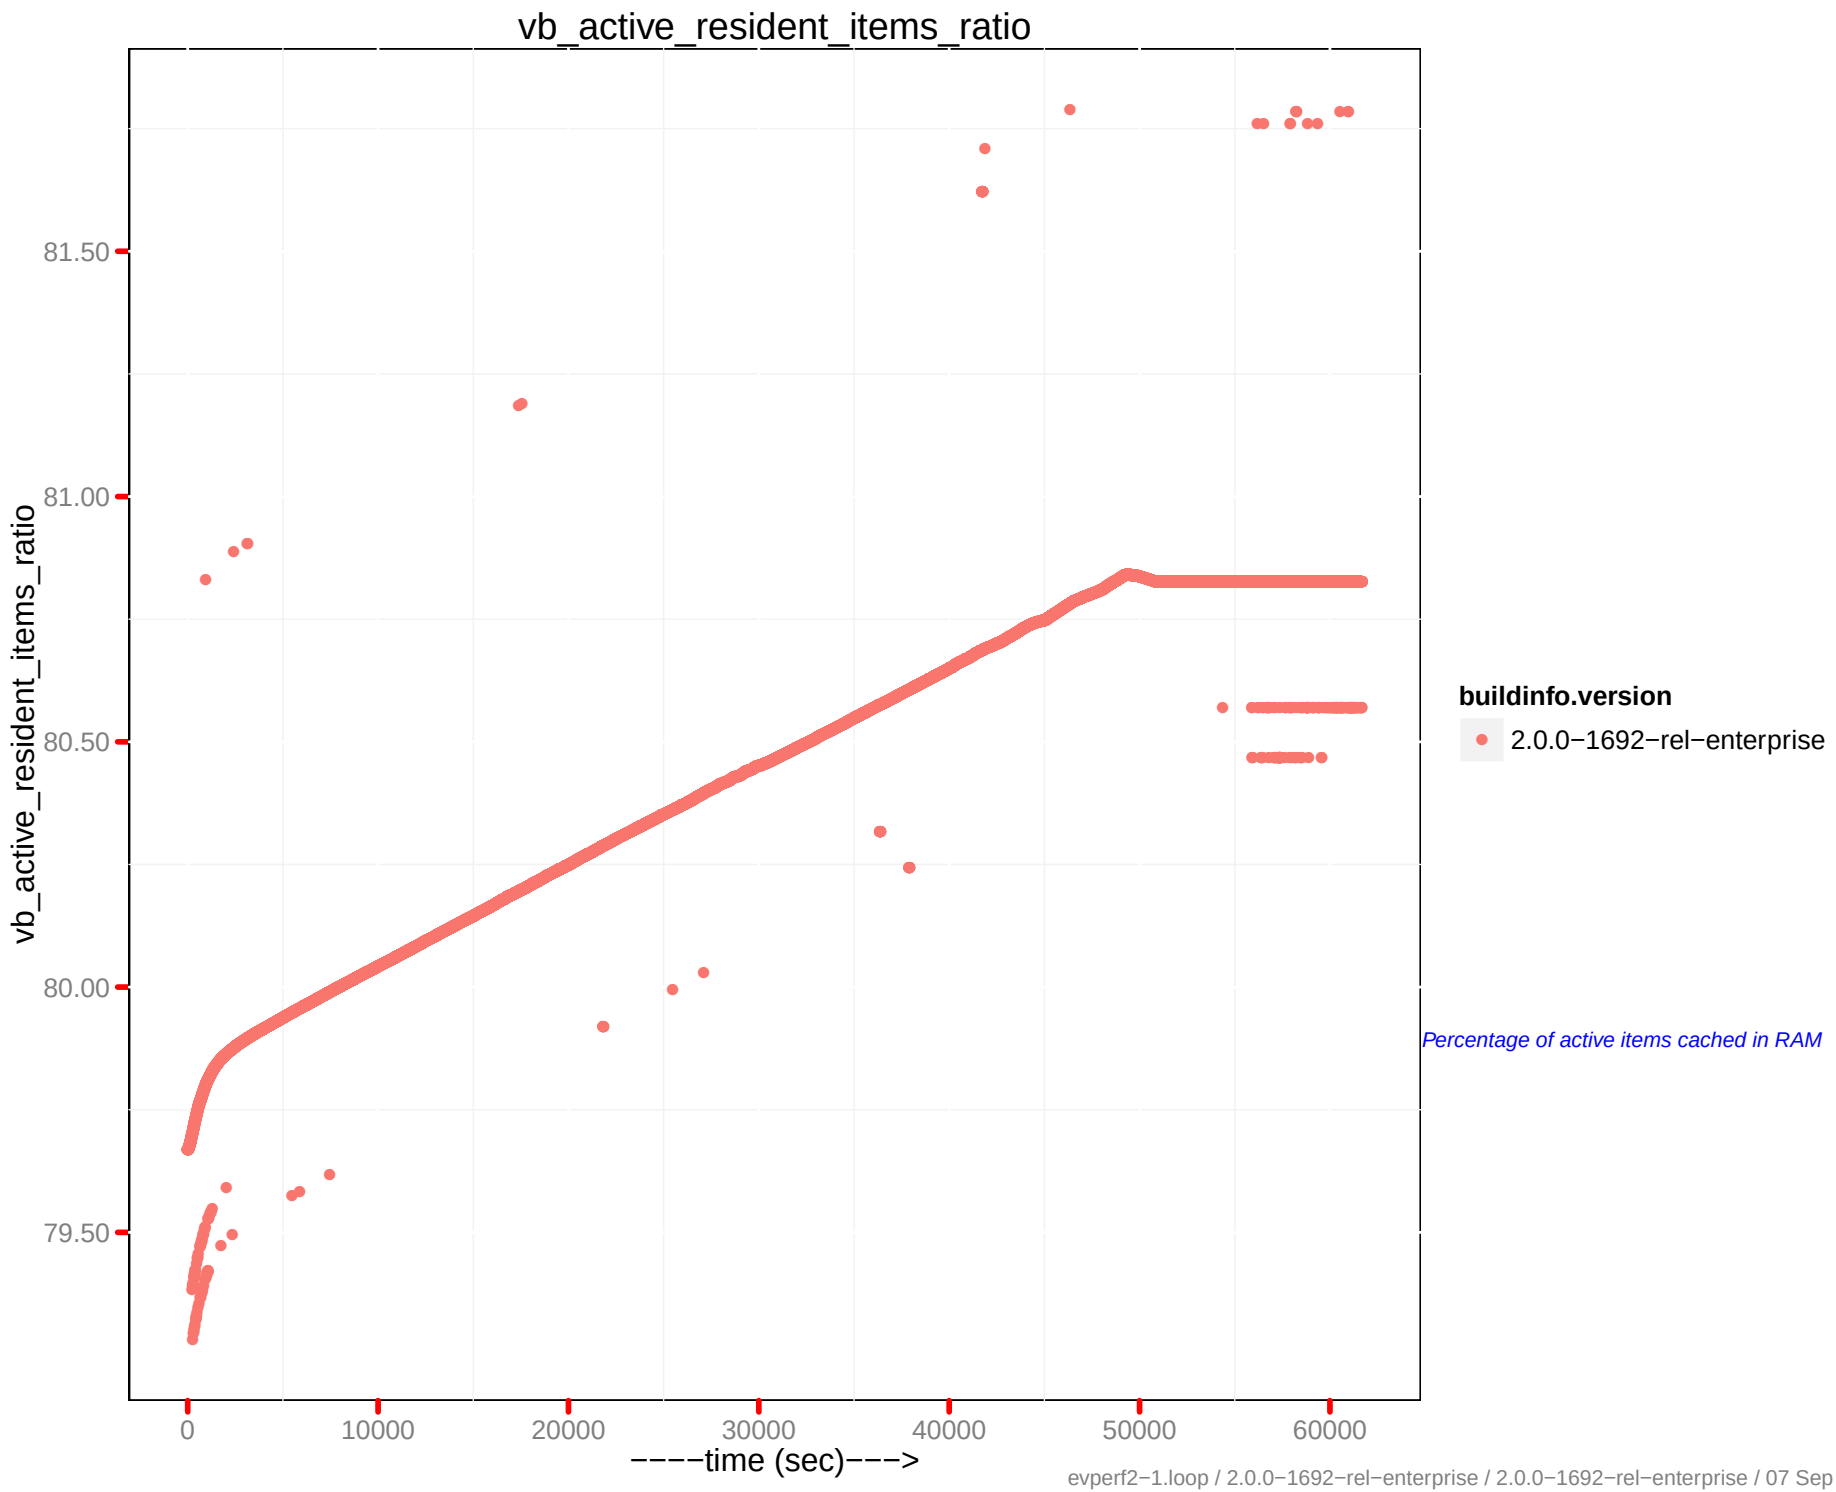
vb_active_resident_items_ratio
81.50
81.00
80.50
80.00
79.50
vb_active_resident_items_ratio
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Percentage of active items cached in RAM
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

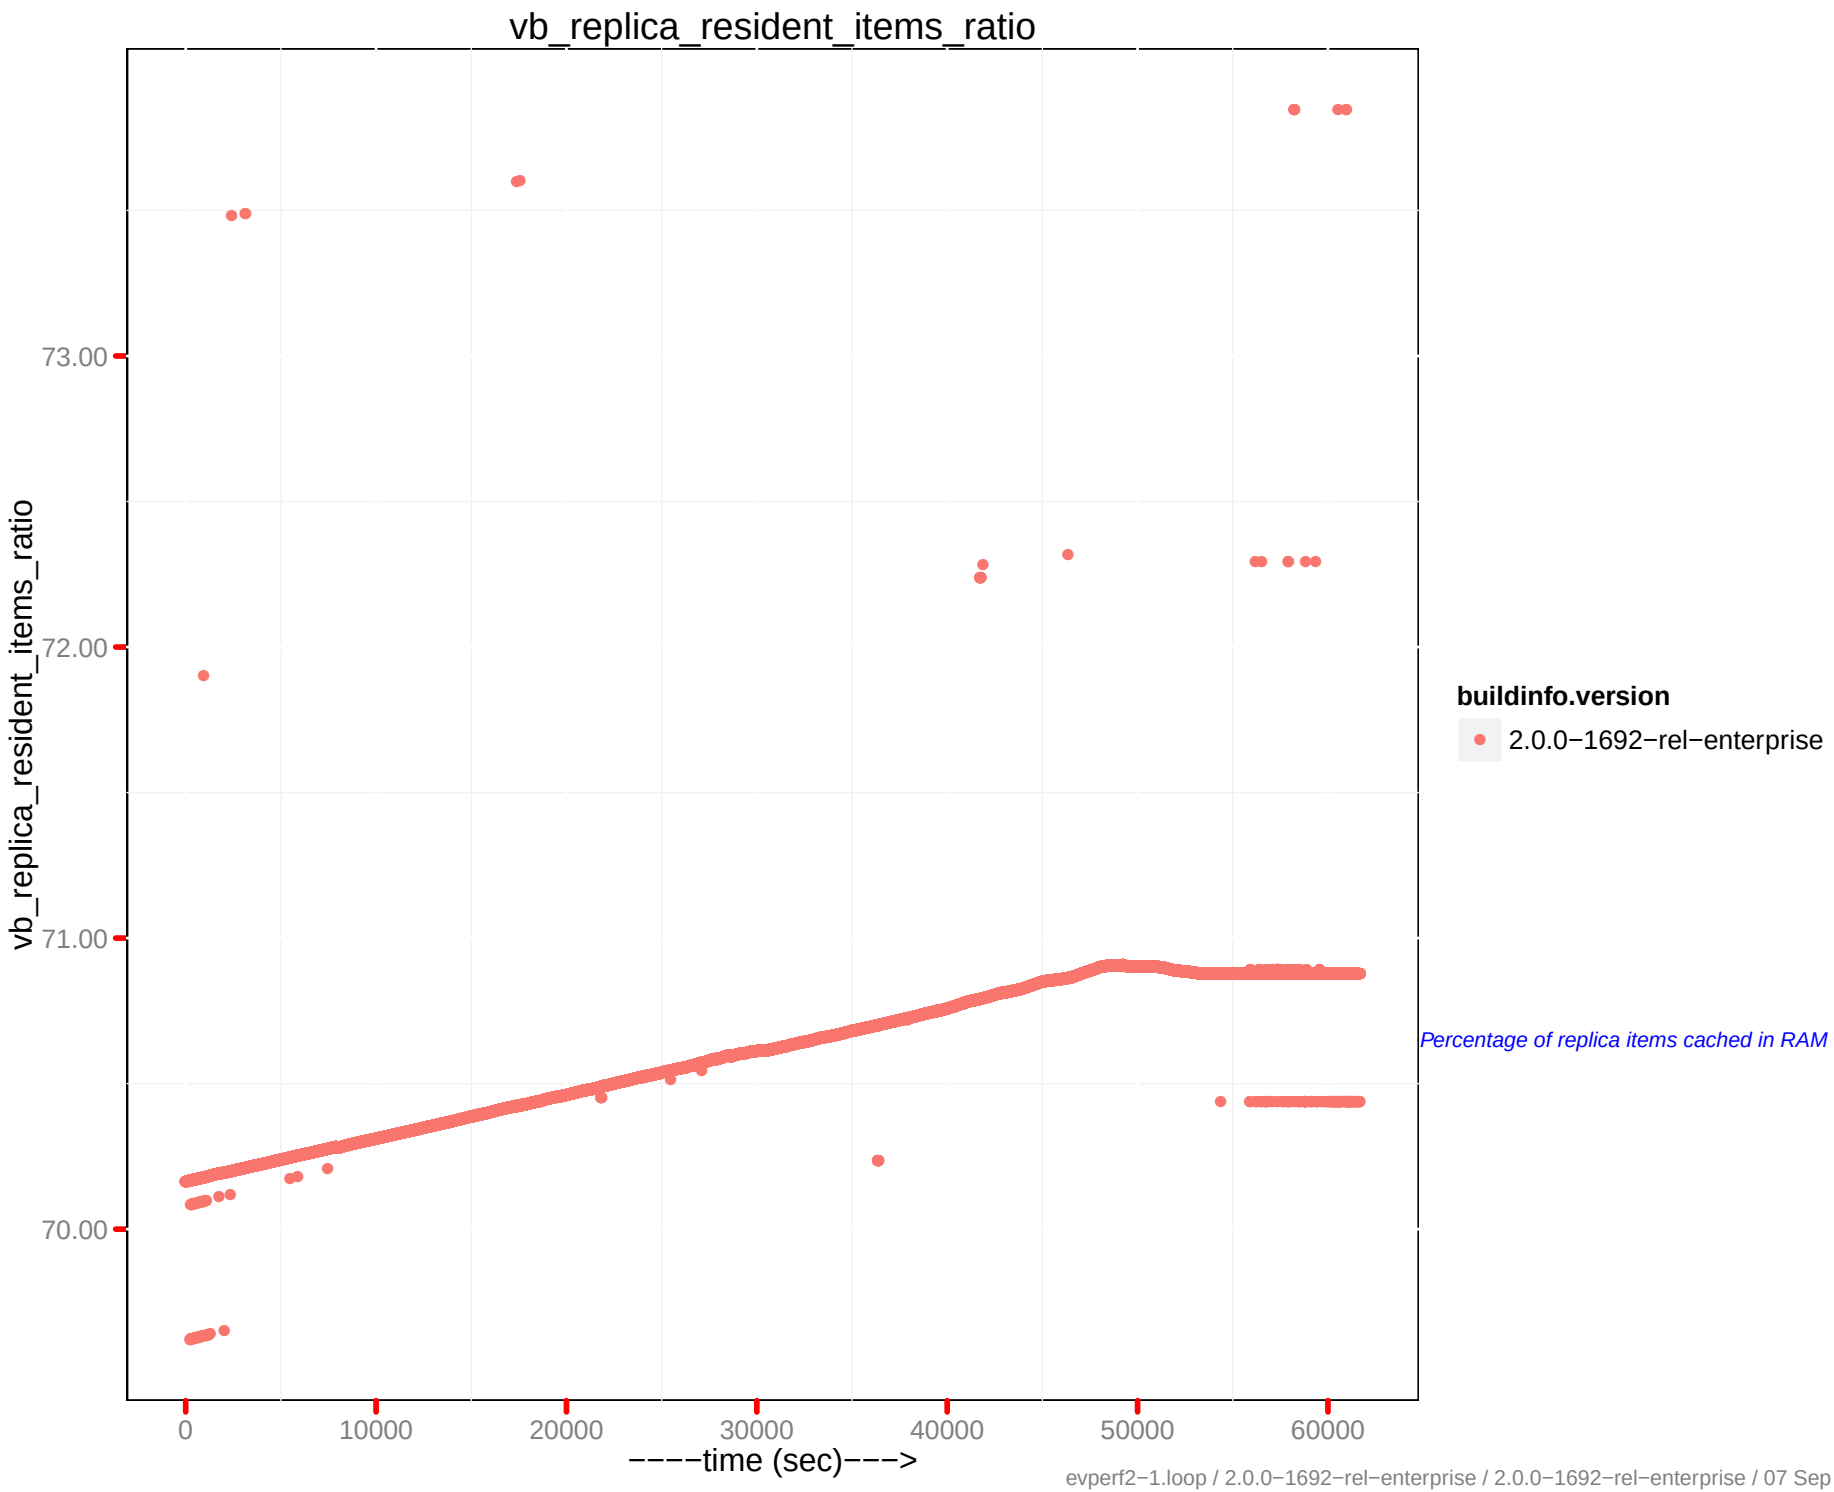
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
73.00
72.00
71.00
70.00
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Percentage of replica items cached in RAM
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

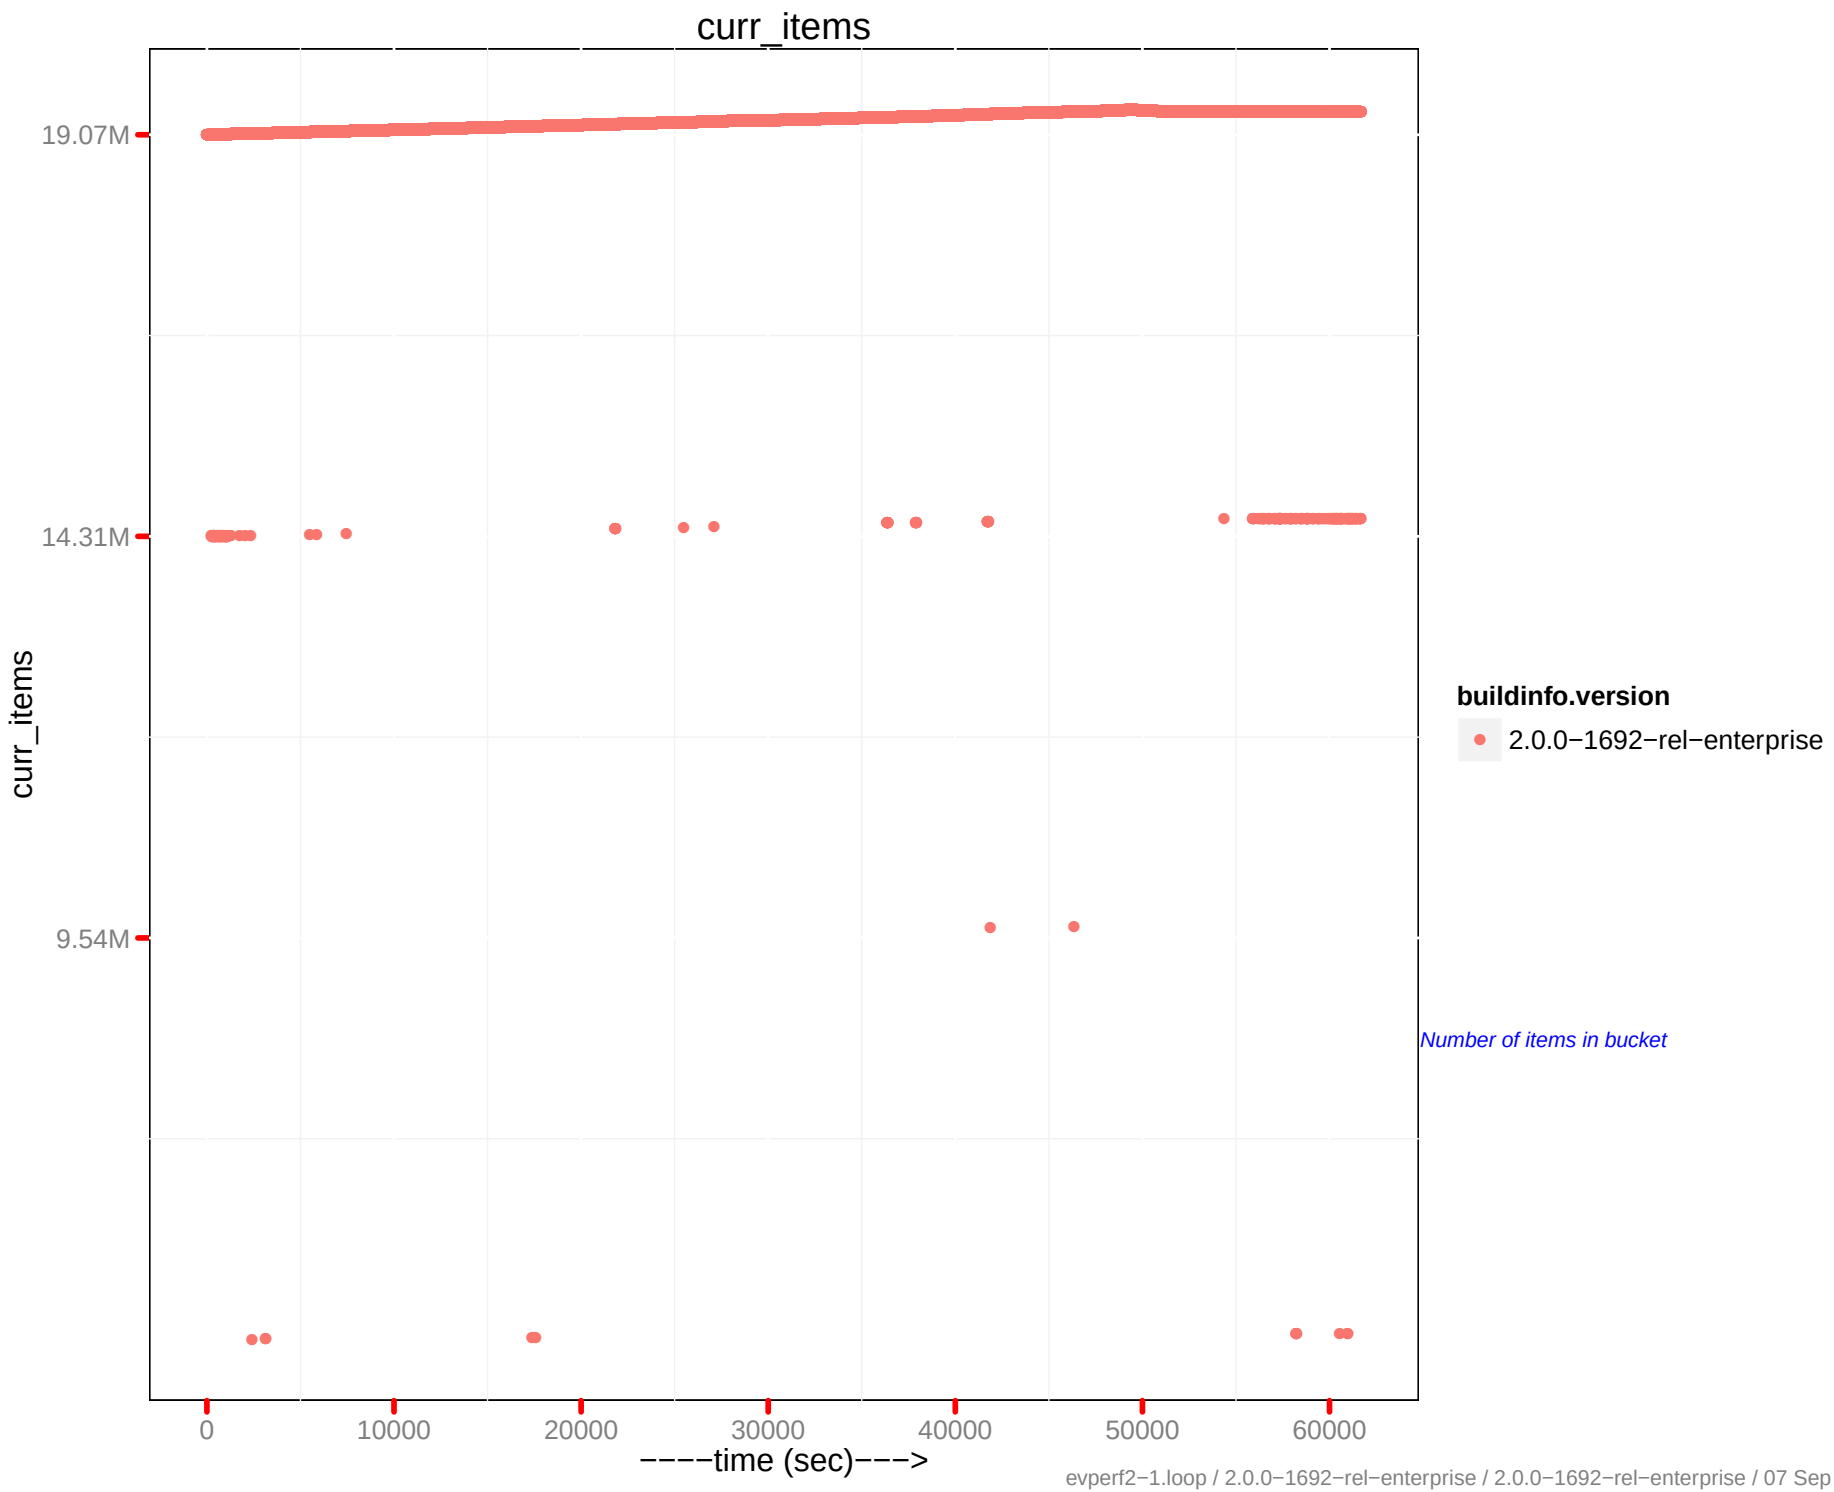
curr_items
19.07M
14.31M
9.54M
curr_items
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of items in bucket
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

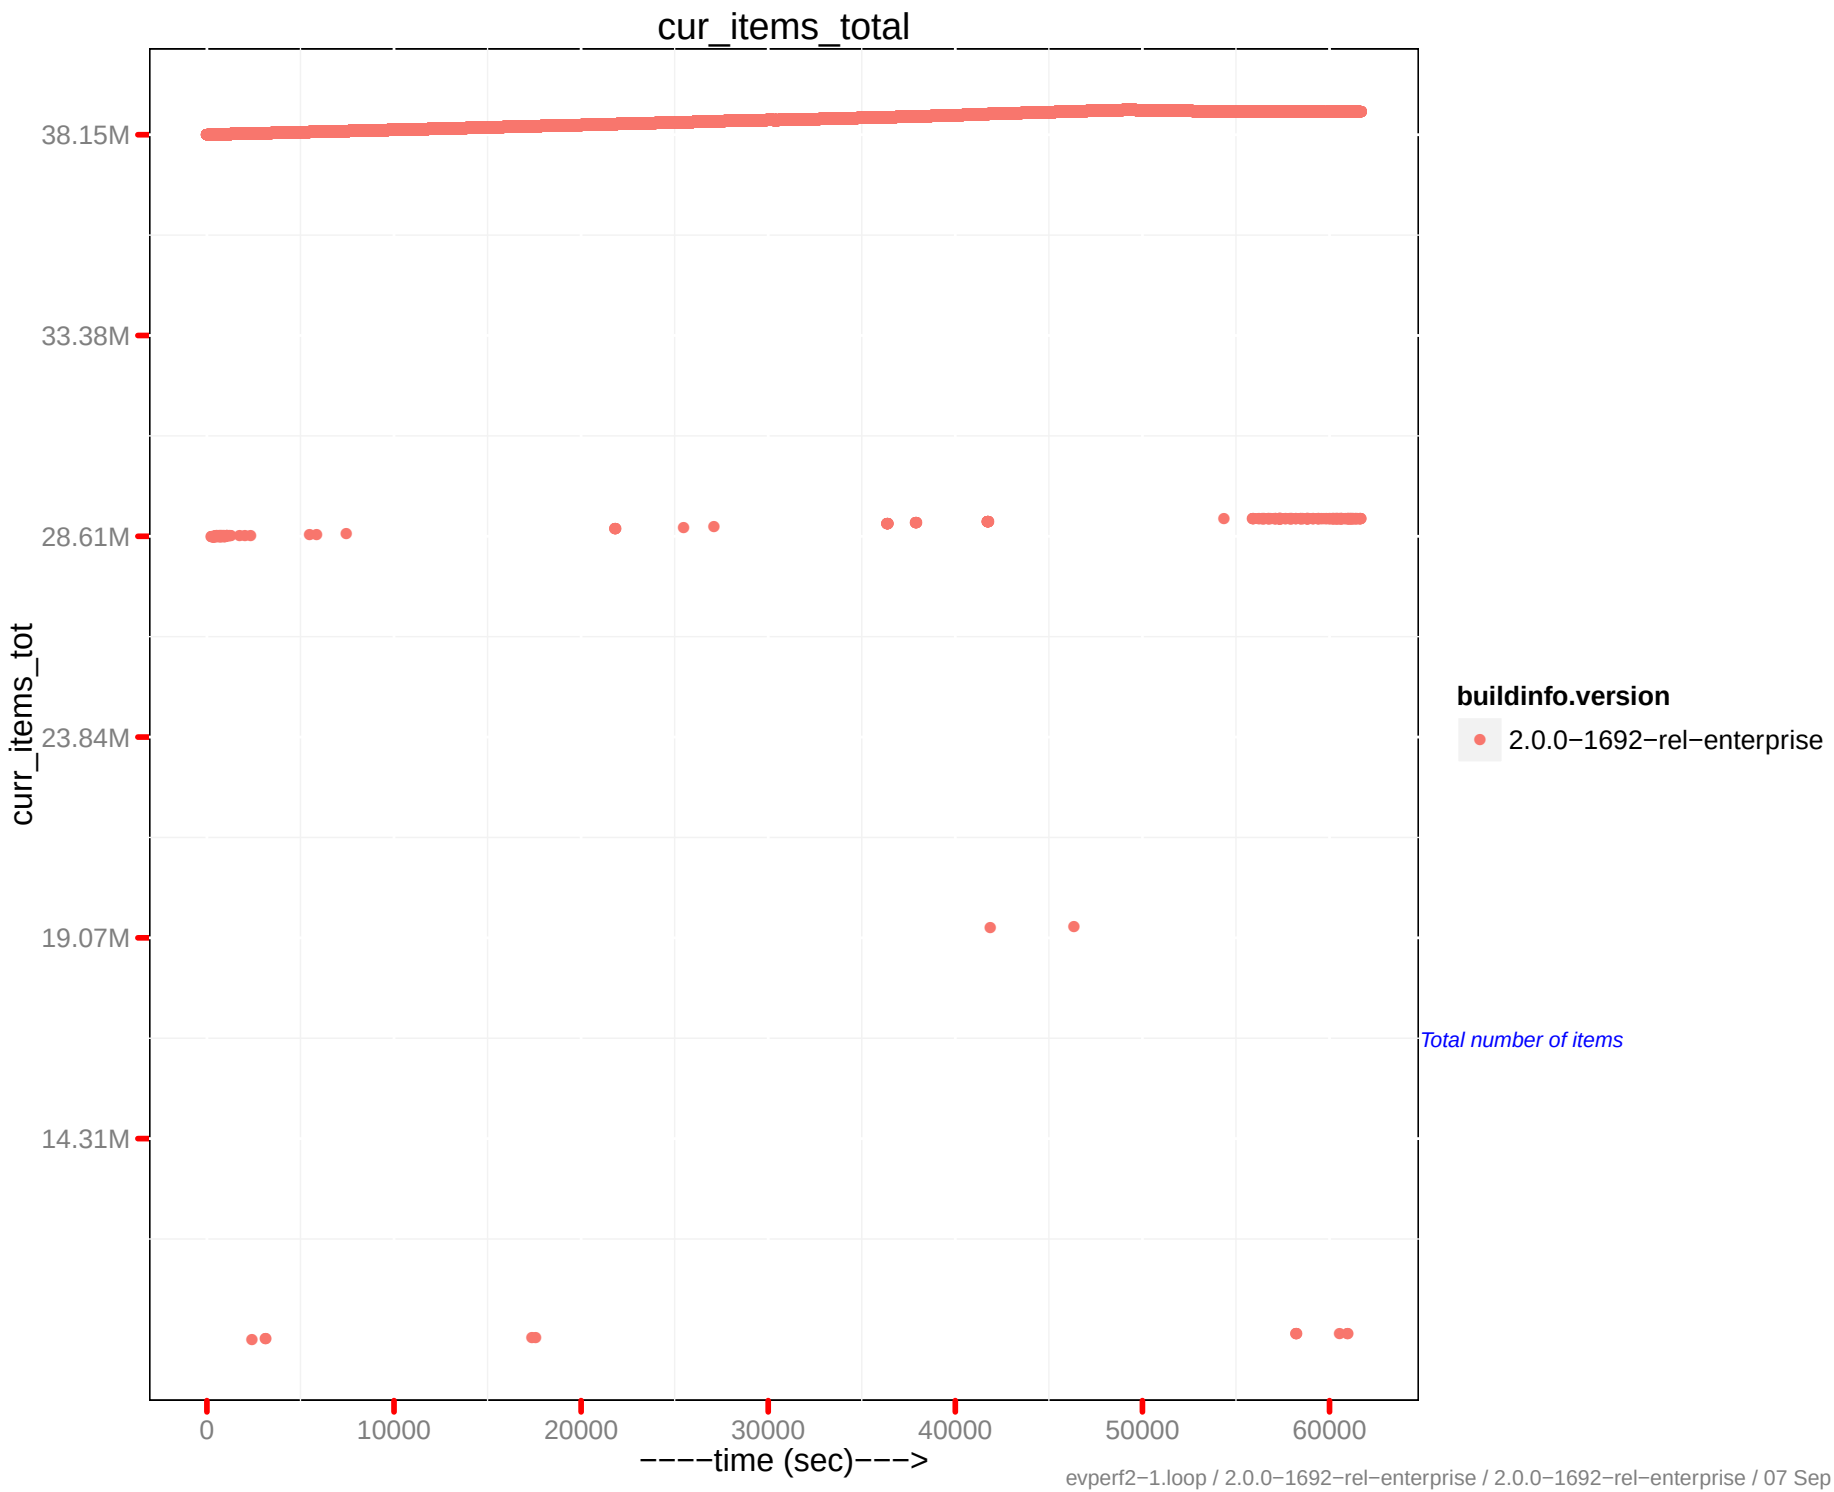
cur_items_total
38.15M
33.38M
28.61M
23.84M
19.07M
14.31M
curr_items_tot
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Total number of items
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

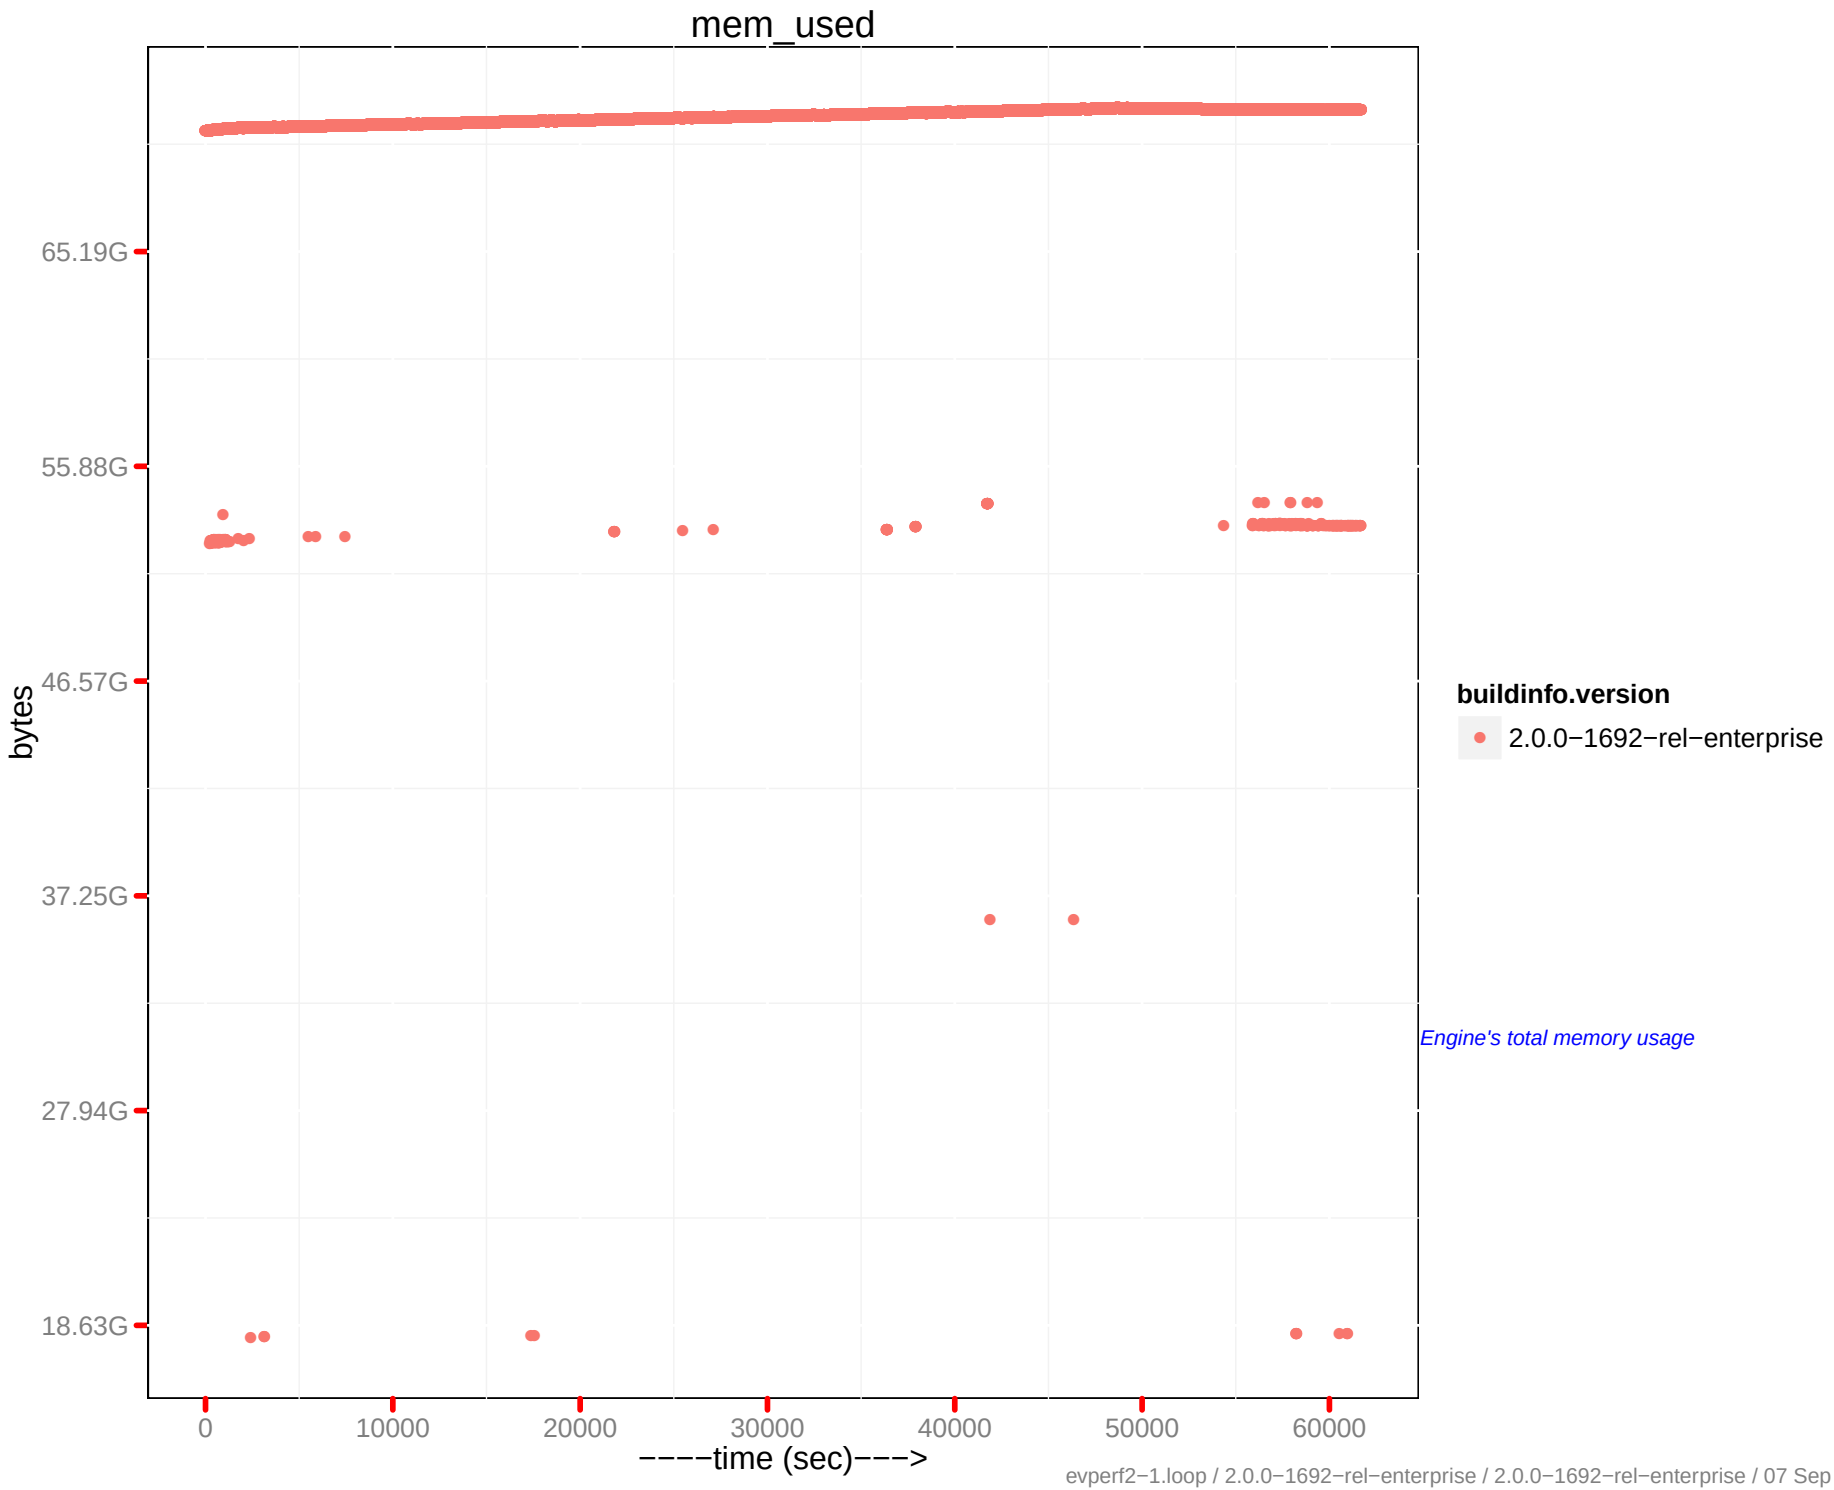
mem_used
65.19G
55.88G
46.57G
37.25G
27.94G
18.63G
bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Engine's total memory usage
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

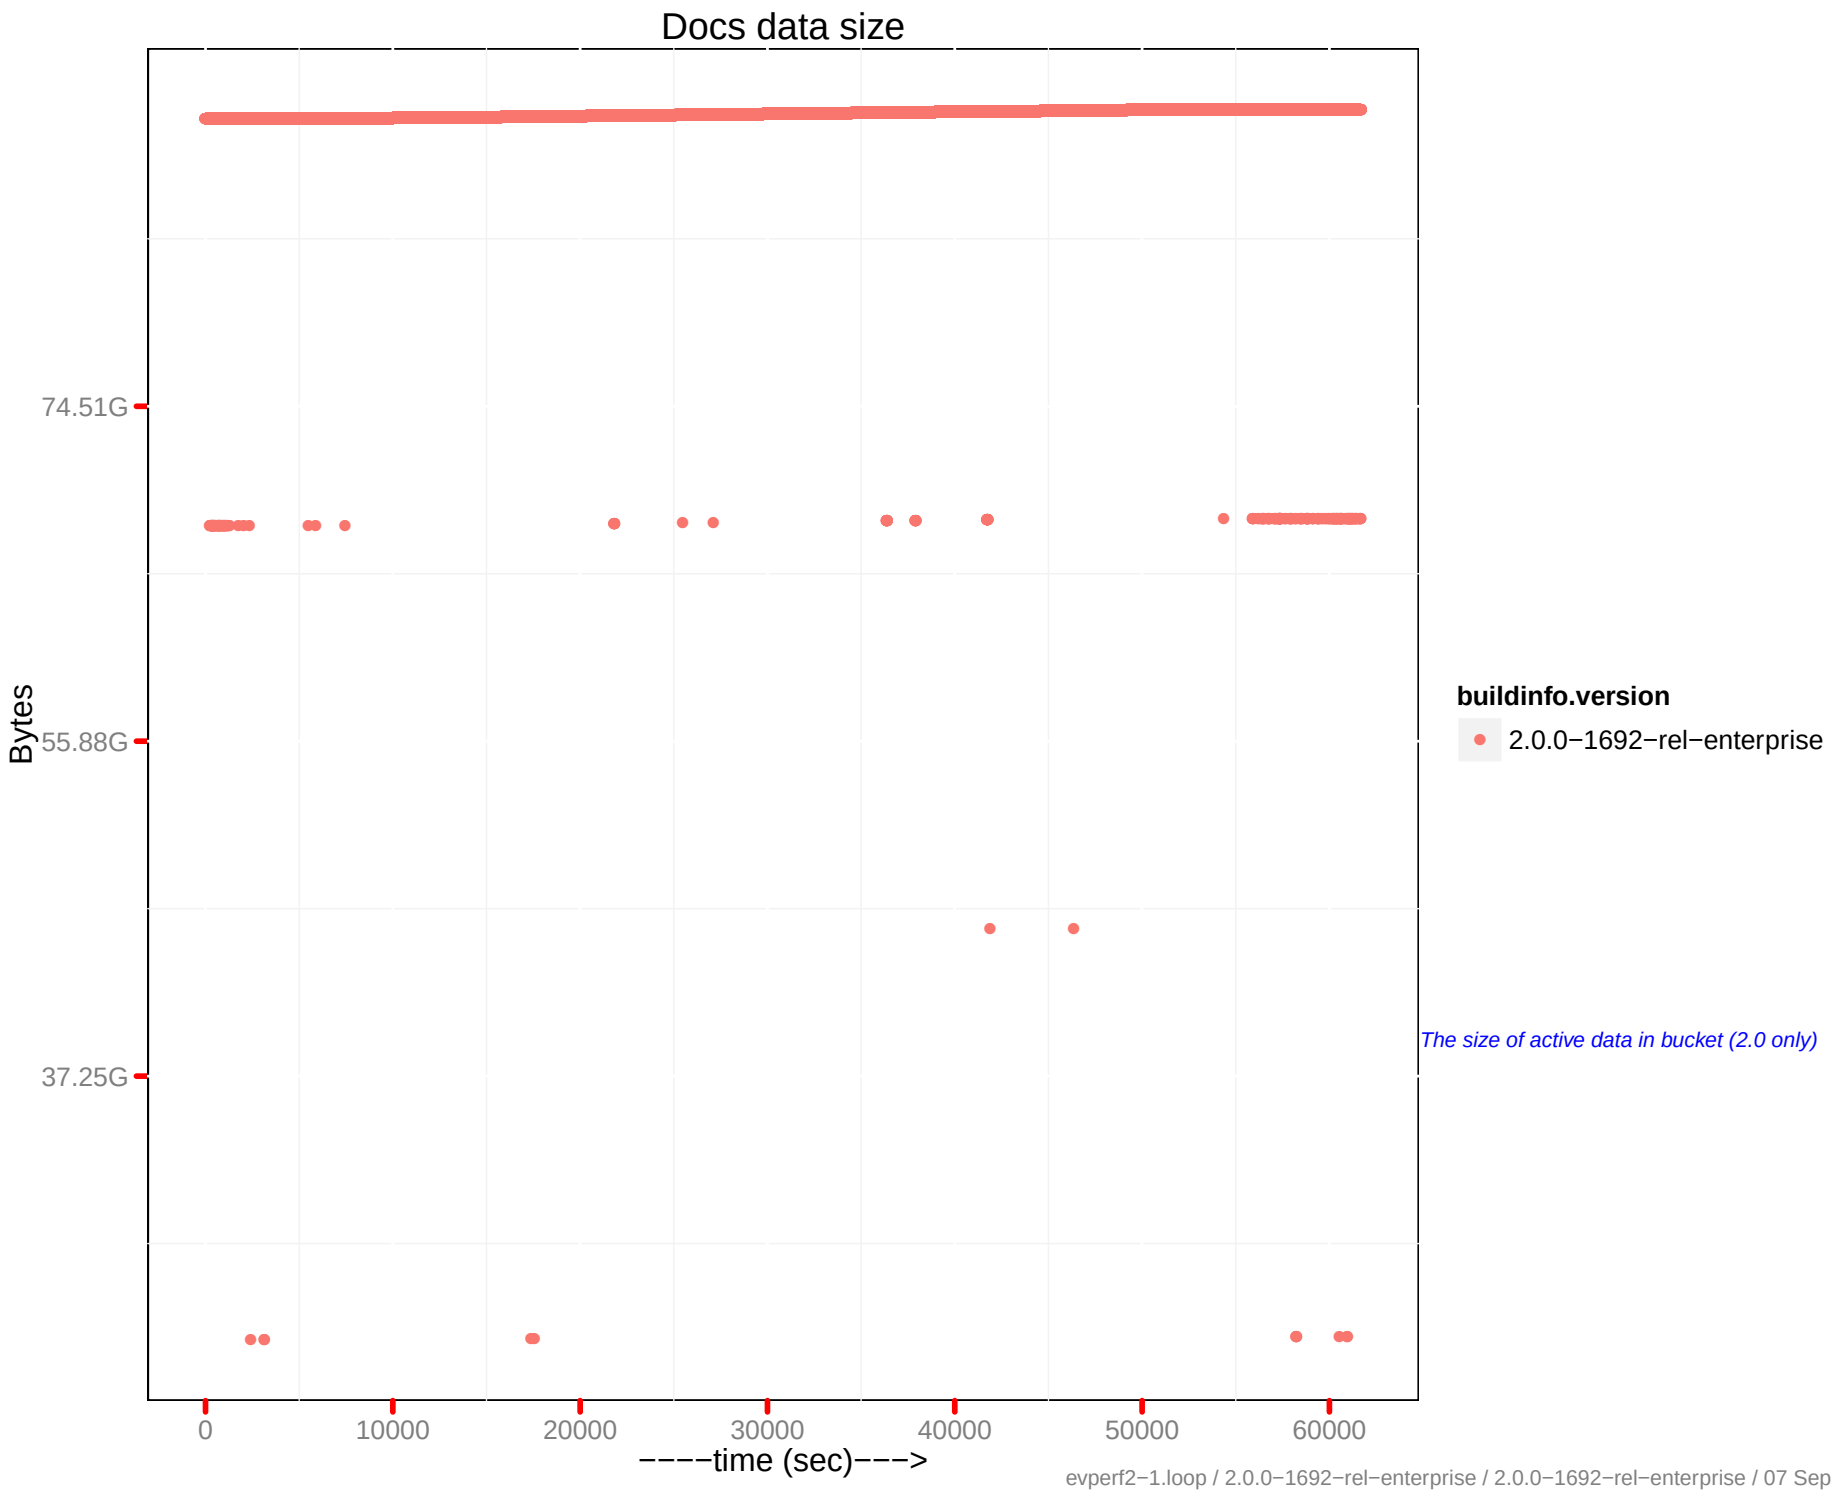
Docs data size
74.51G
55.88G
37.25G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The size of active data in bucket (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

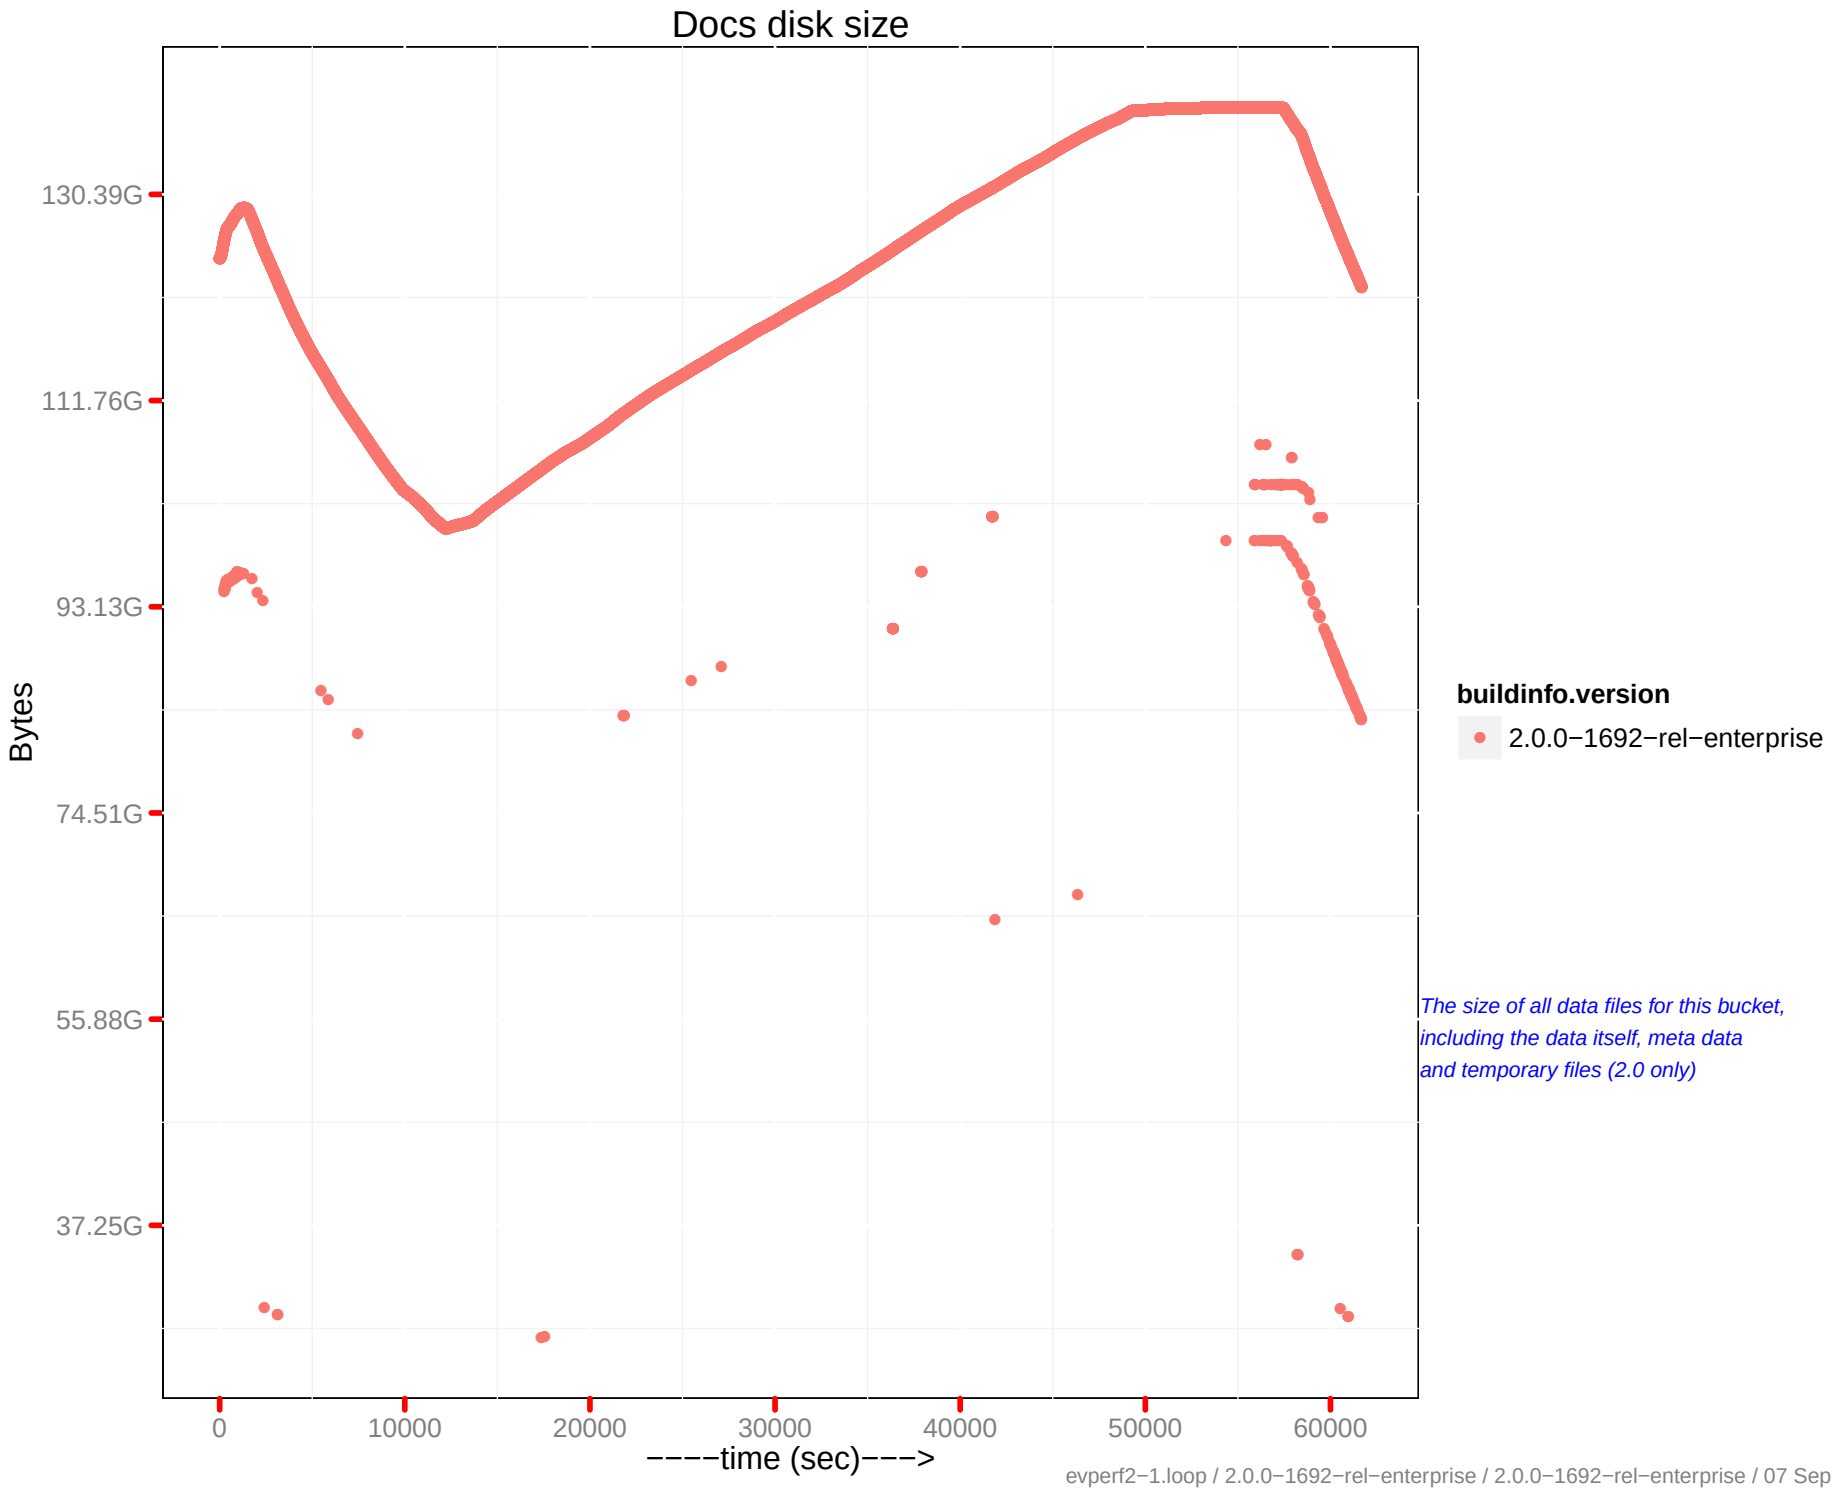
Docs disk size
130.39G
111.76G
93.13G
74.51G
55.88G
37.25G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

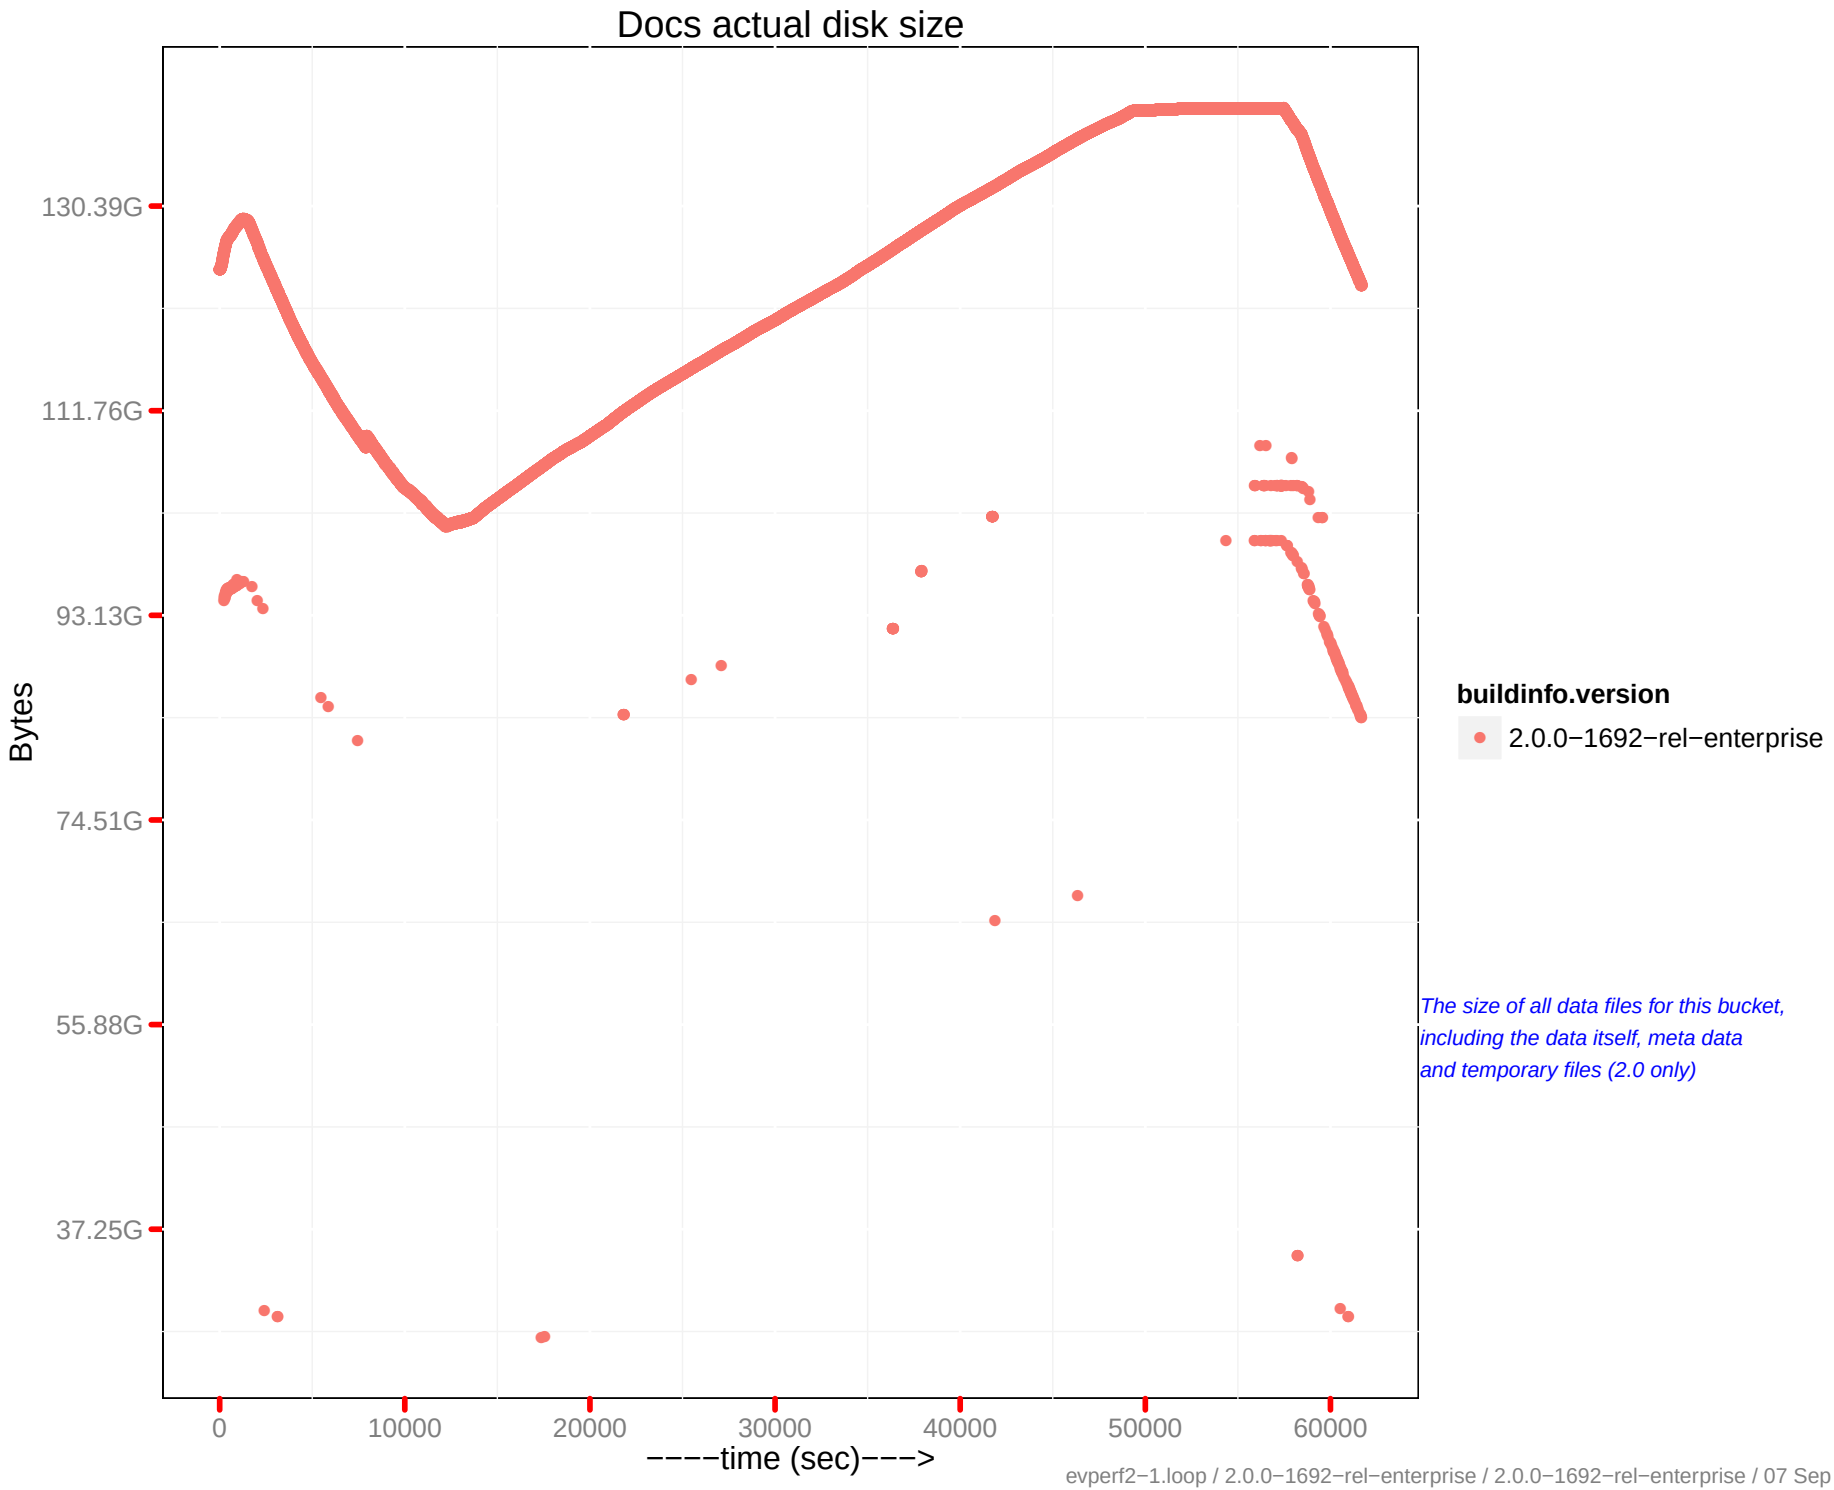
Docs actual disk size
130.39G
111.76G
93.13G
74.51G
55.88G
37.25G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

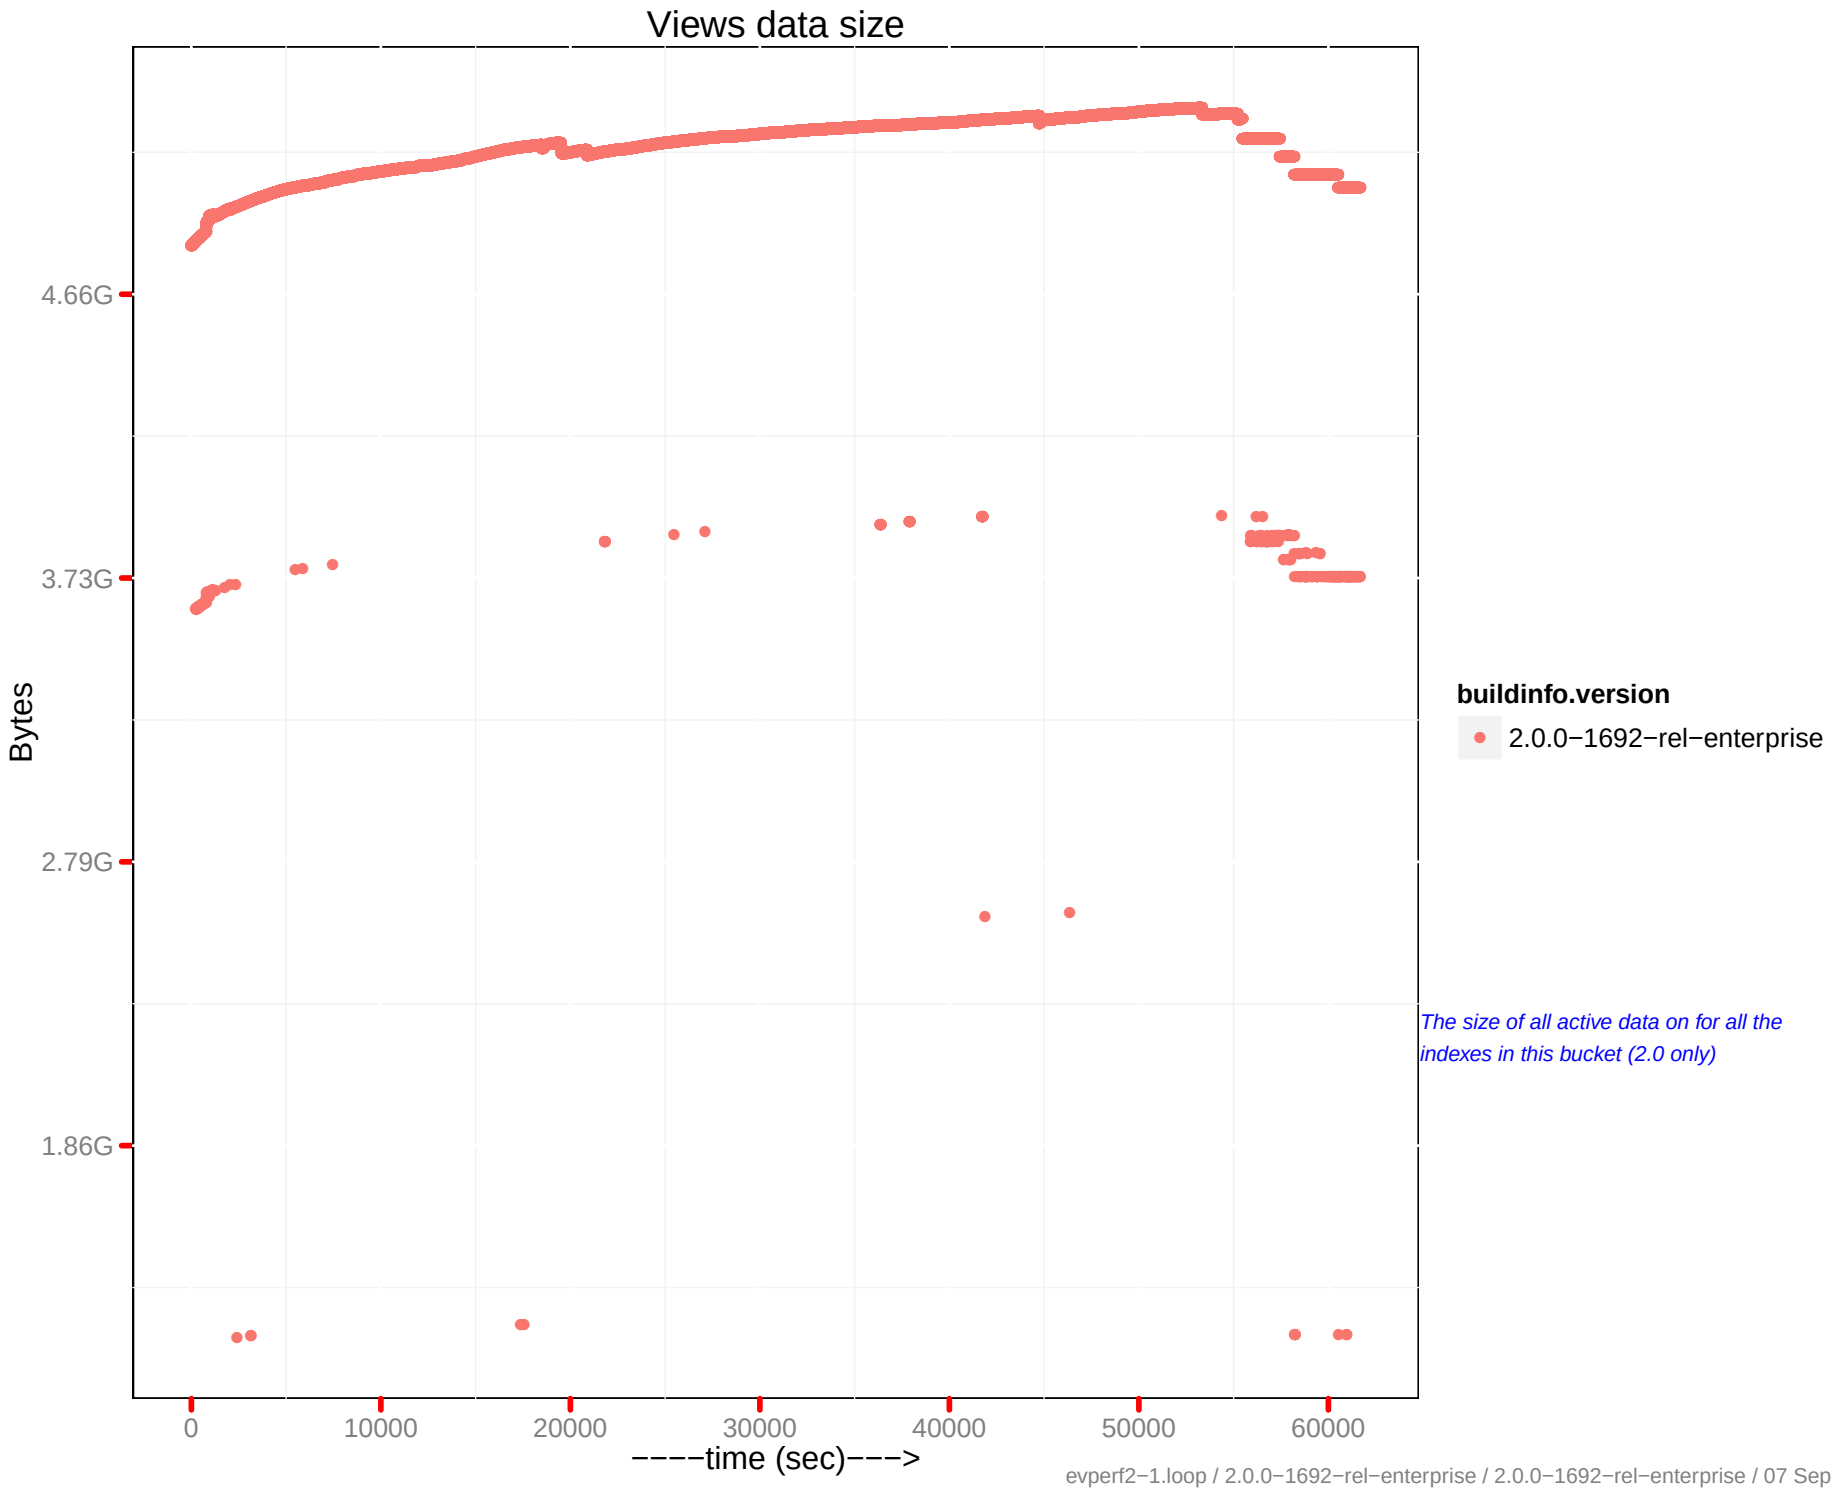
Views data size
4.66G
3.73G
2.79G
1.86G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

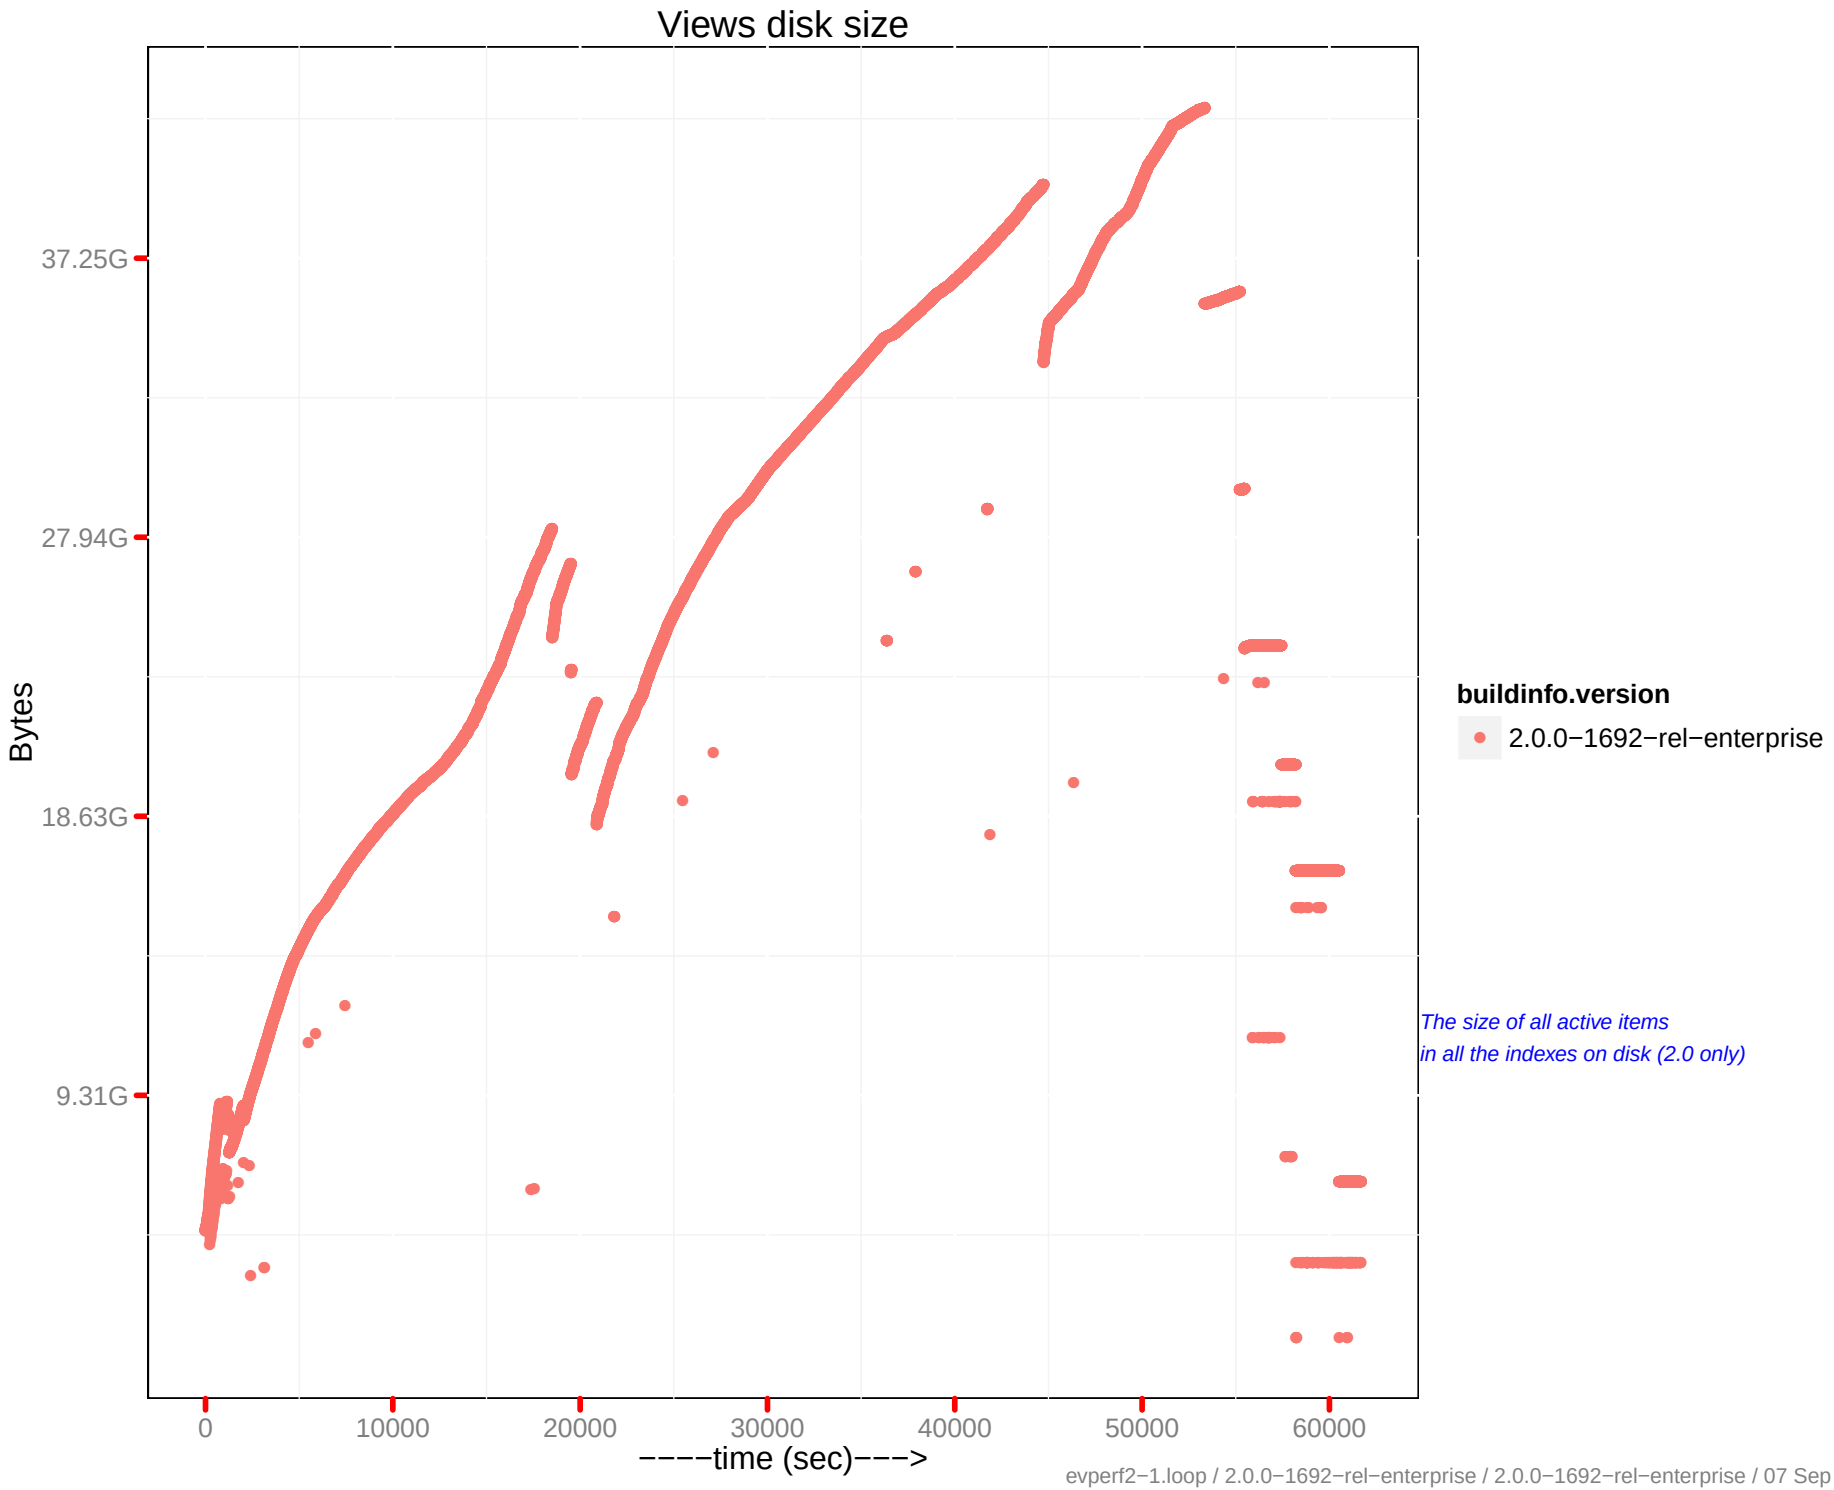
Views disk size
37.25G
27.94G
18.63G
9.31G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

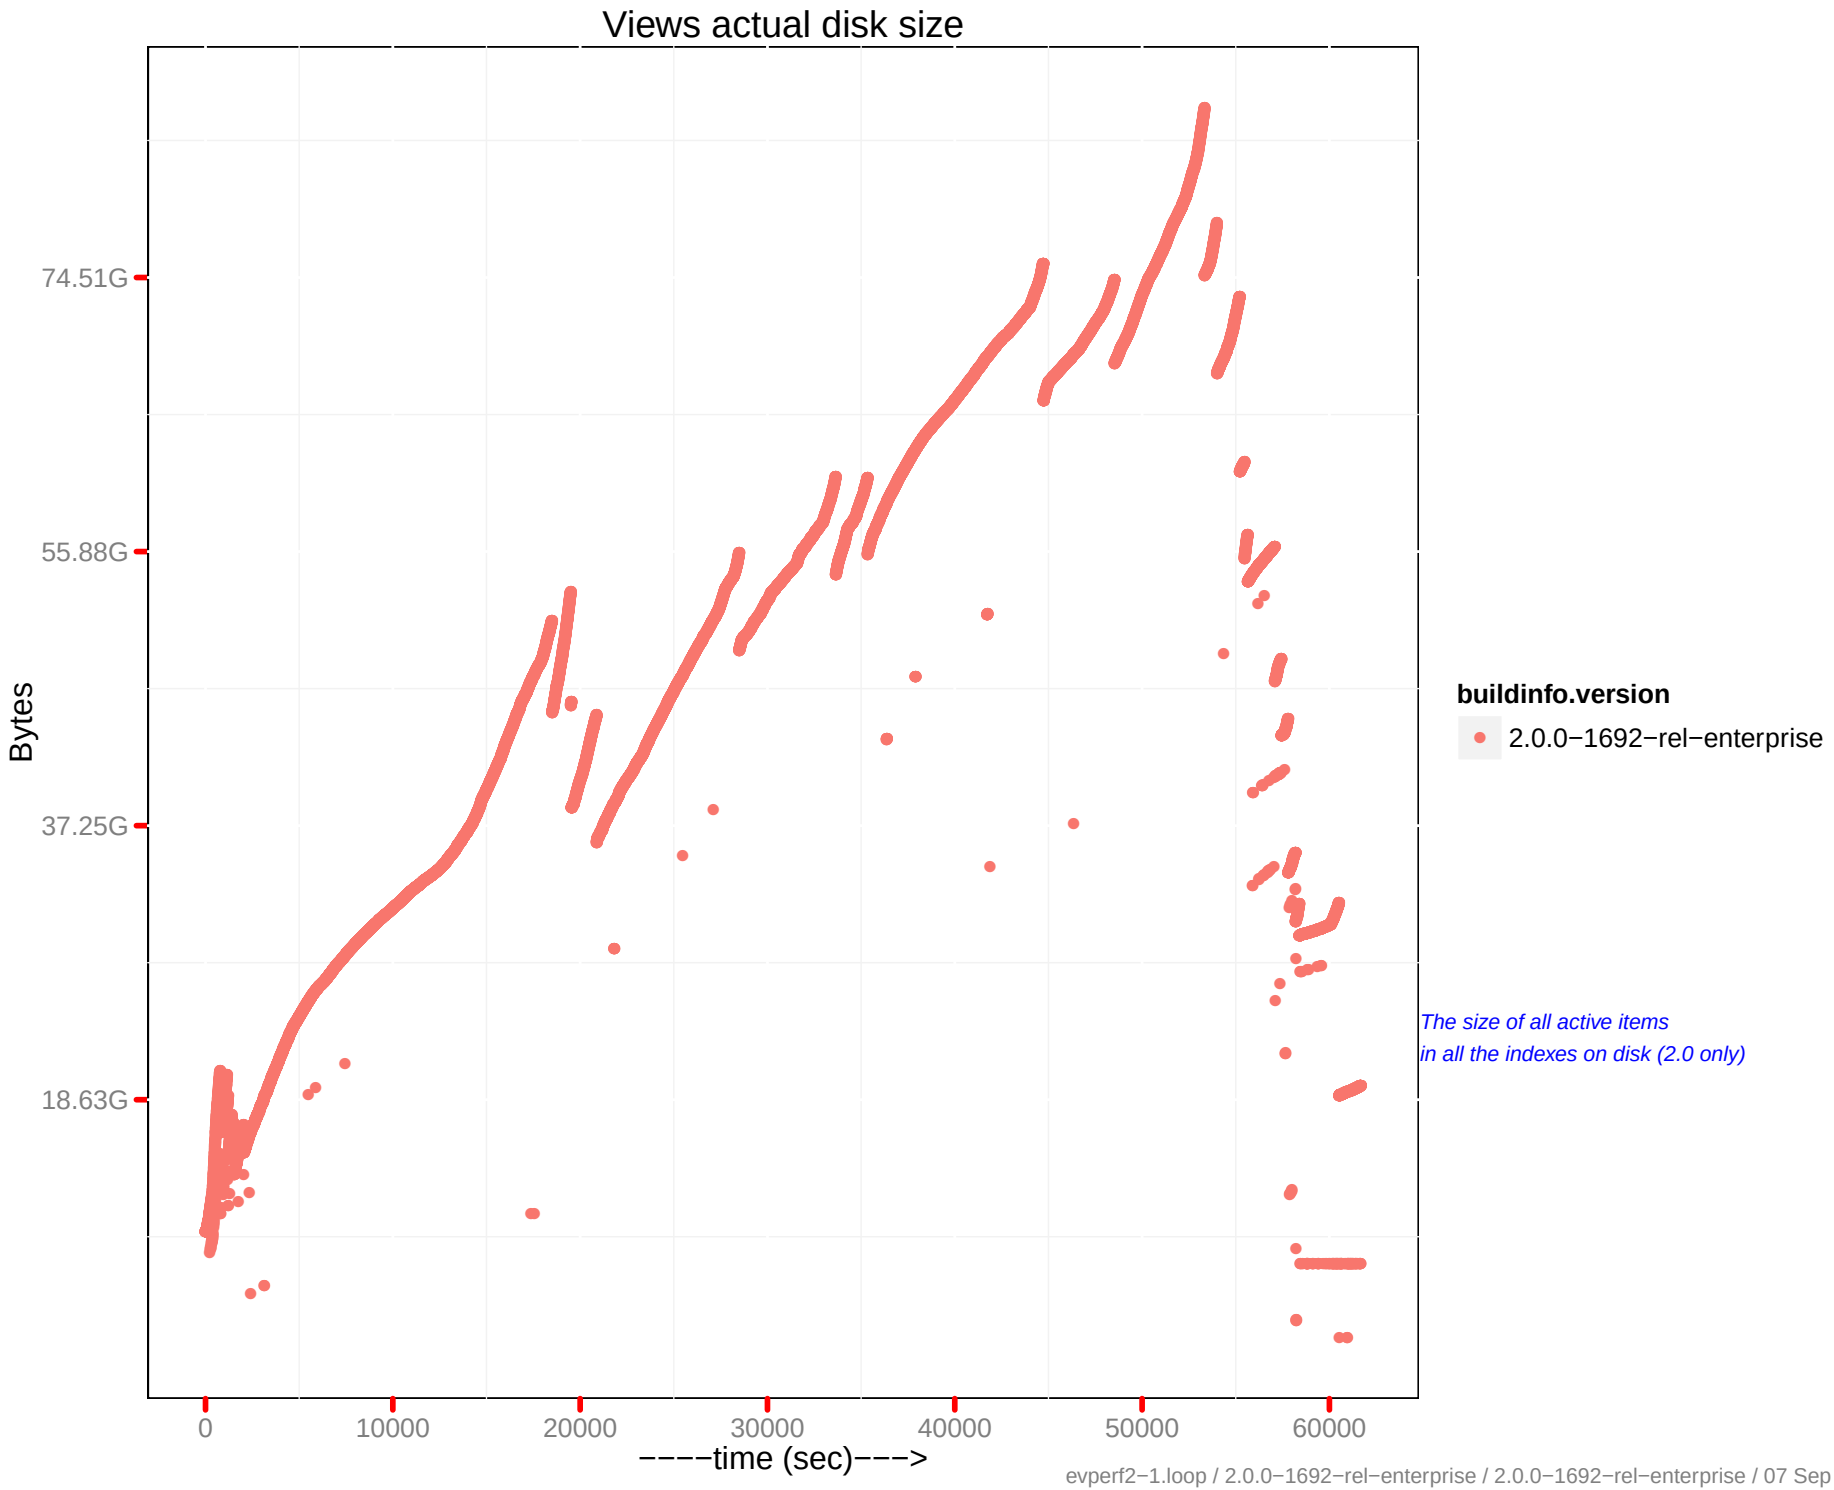
Views actual disk size
74.51G
55.88G
37.25G
18.63G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

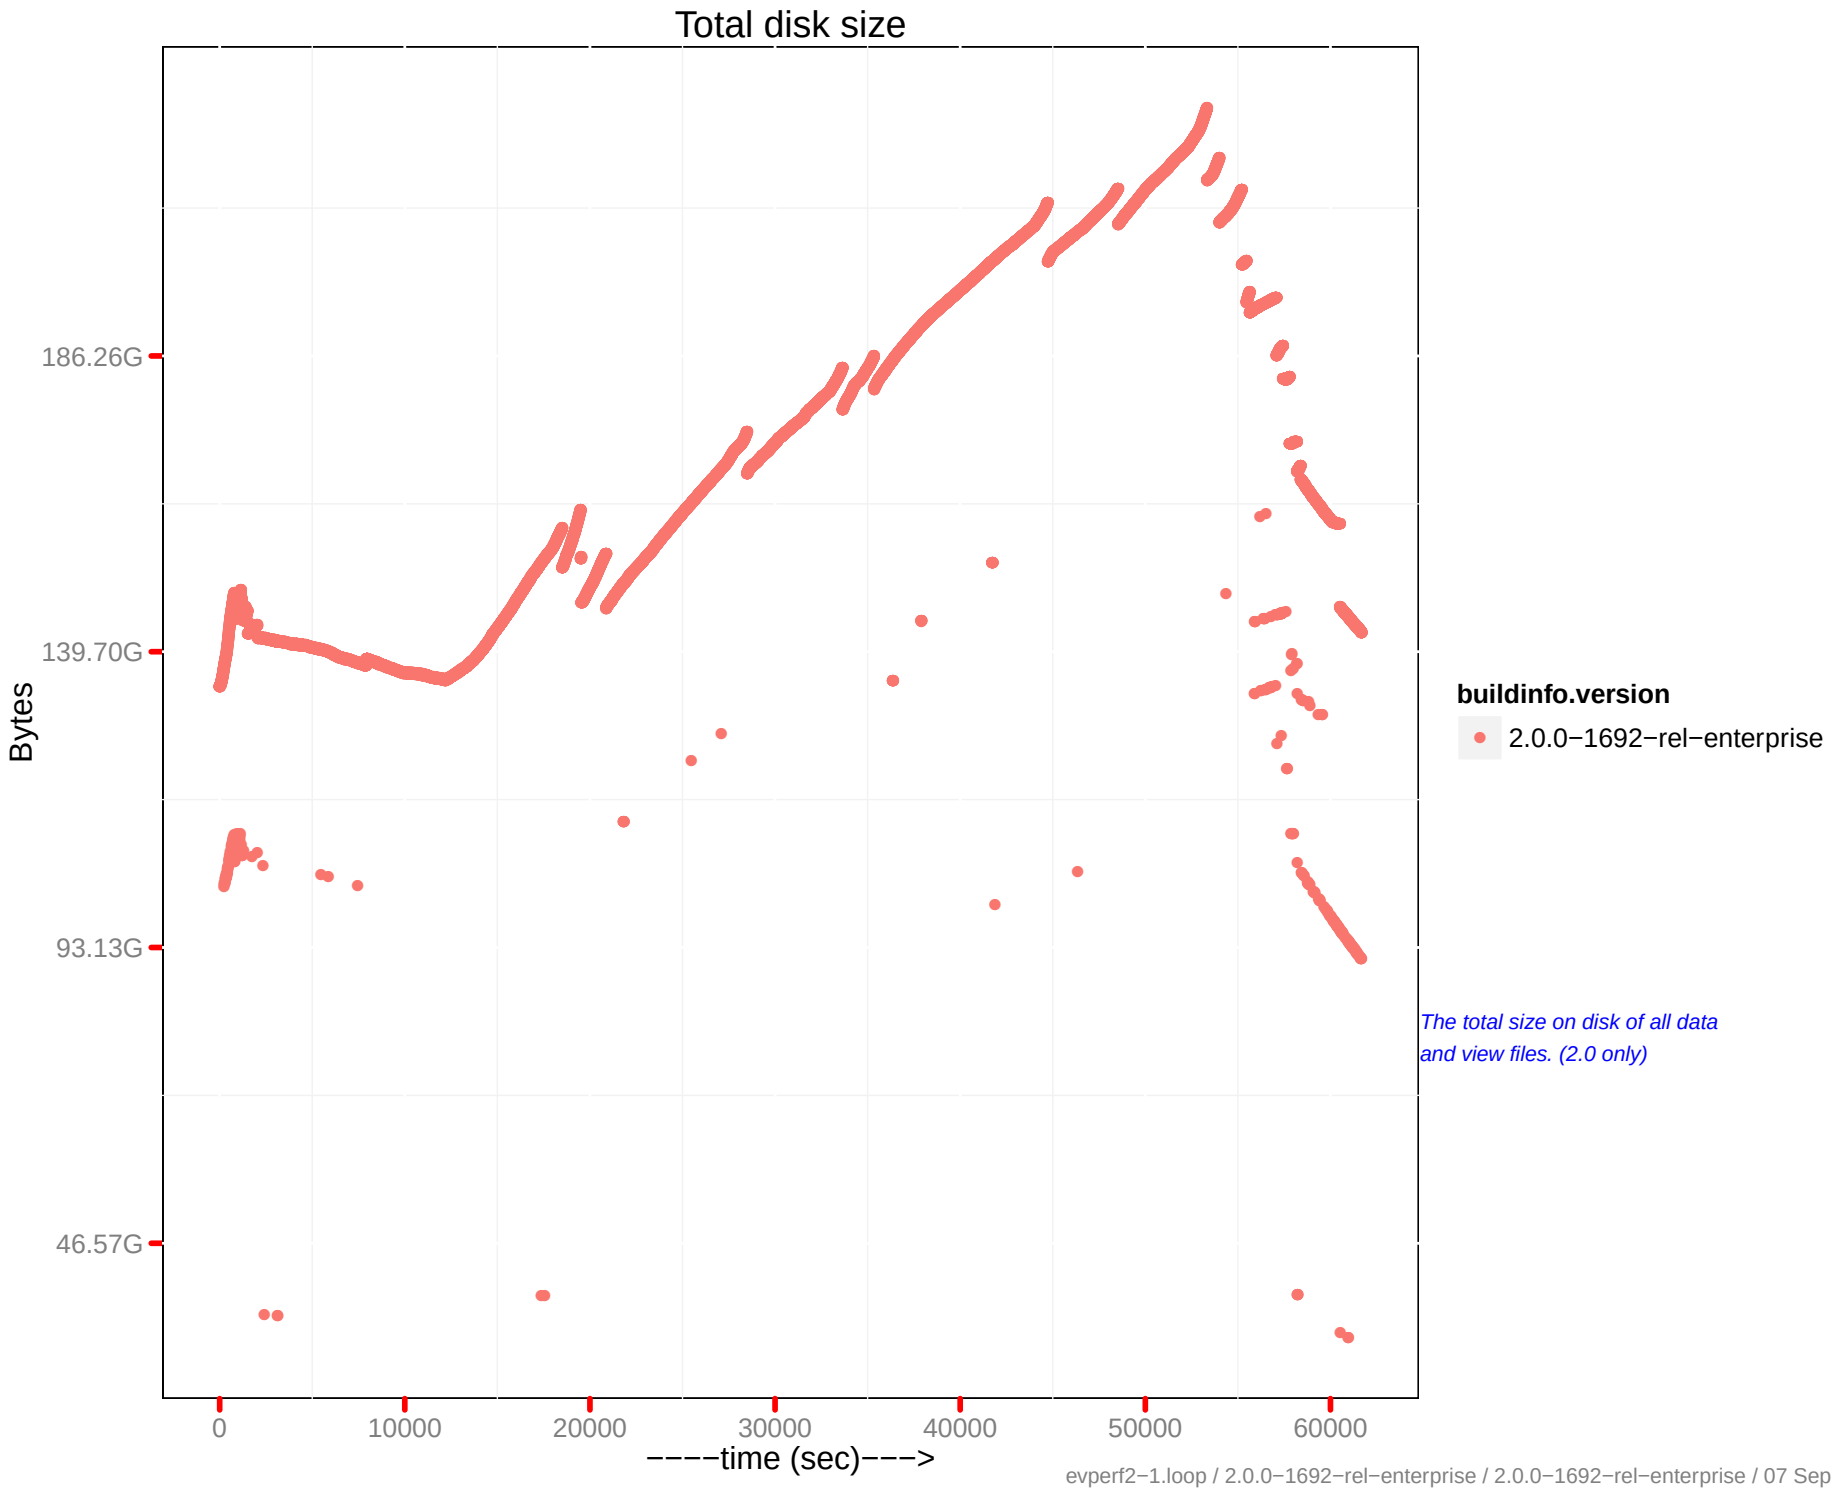
Total disk size
186.26G
139.70G
93.13G
46.57G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

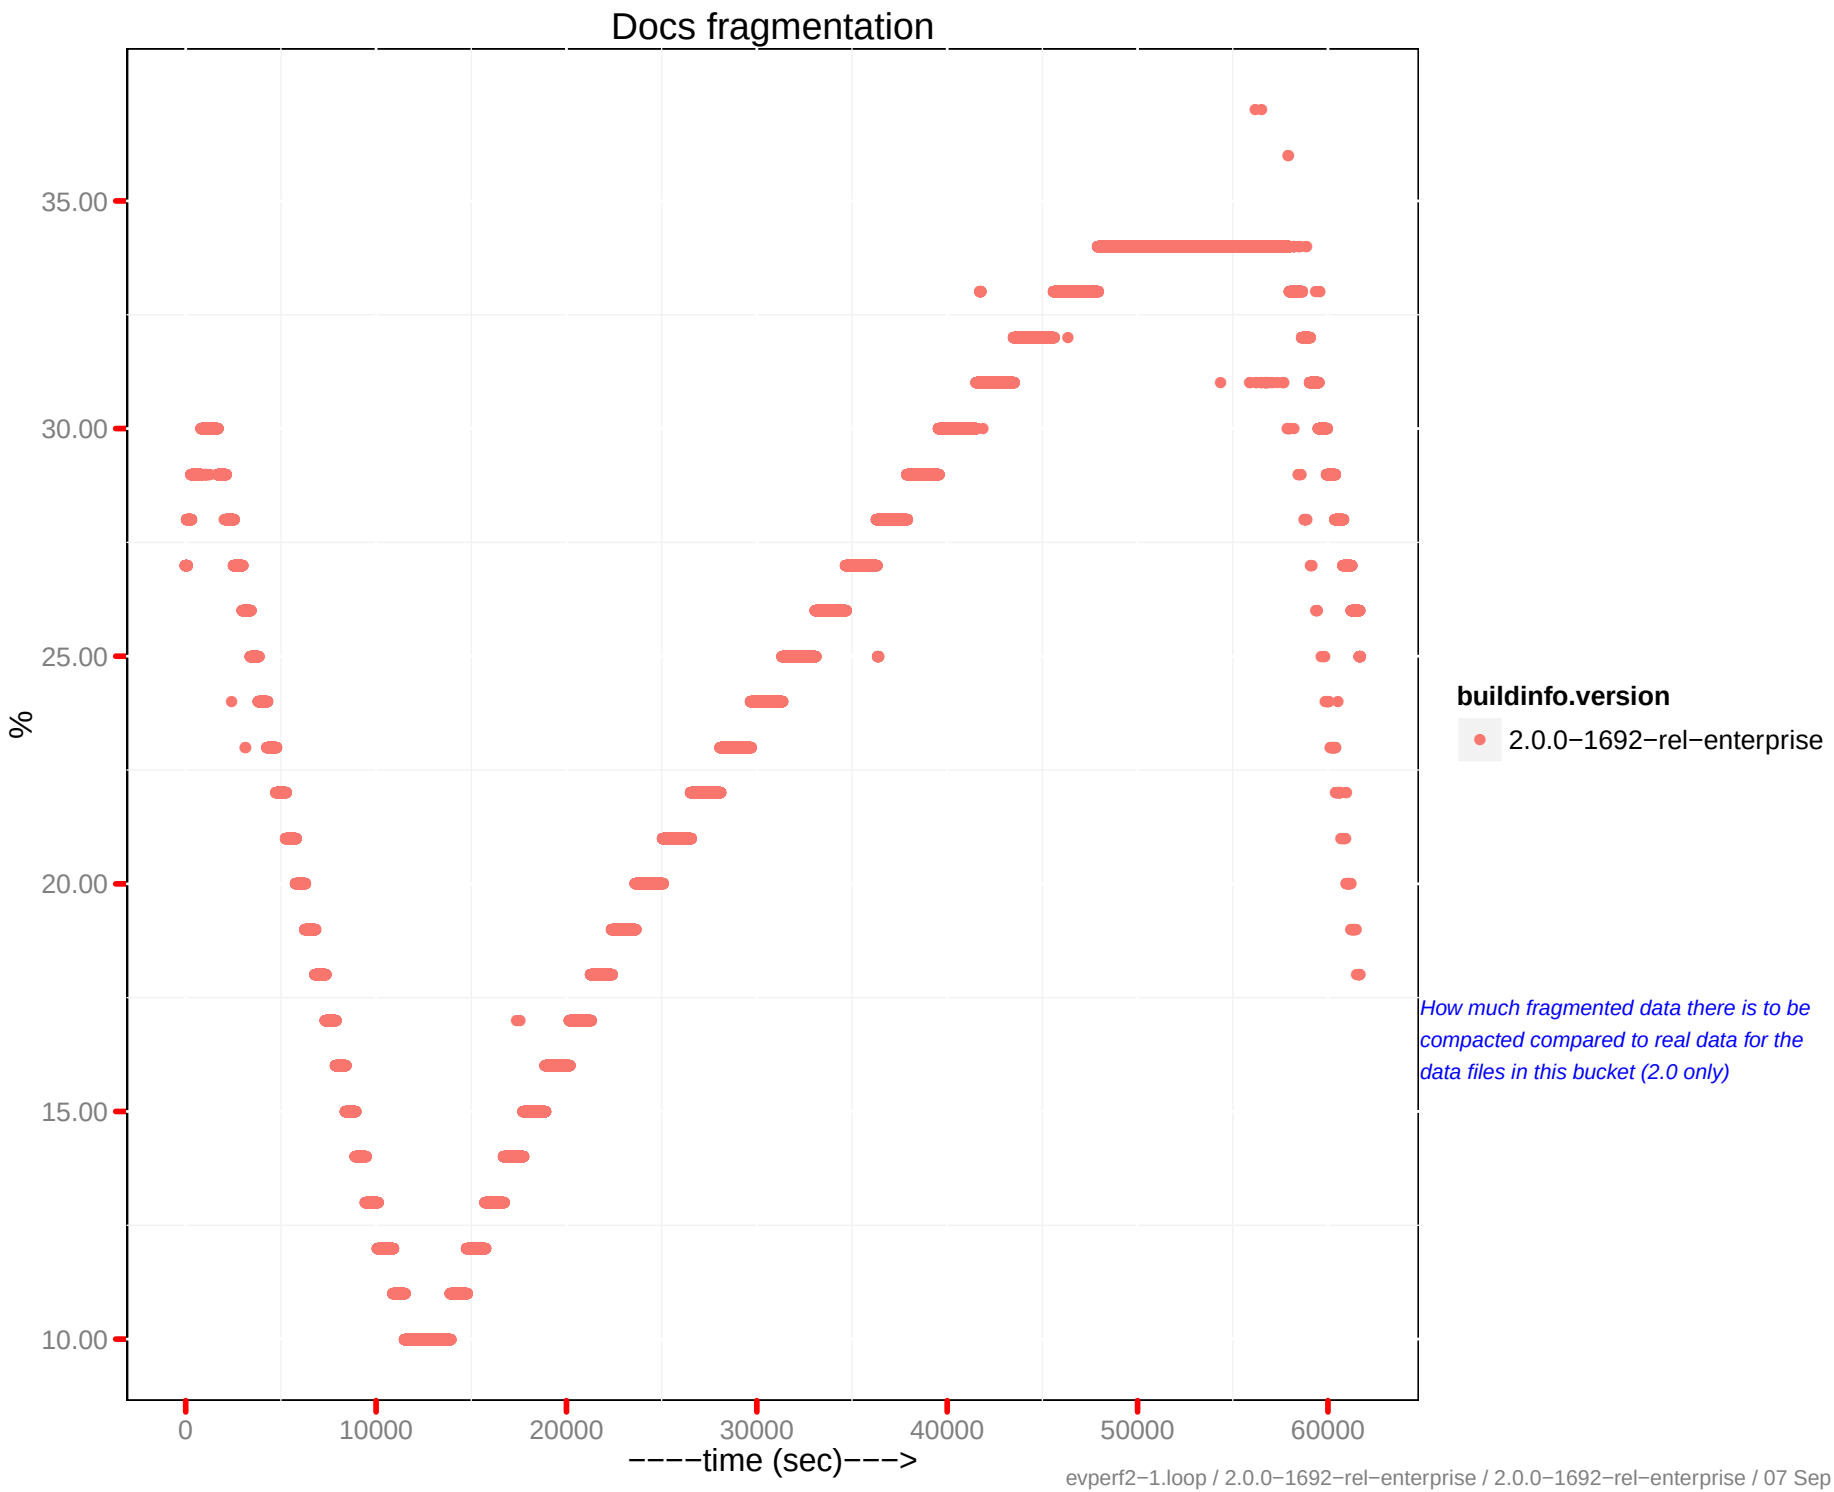
Docs fragmentation
35.00
30.00
25.00
20.00
15.00
10.00
%
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

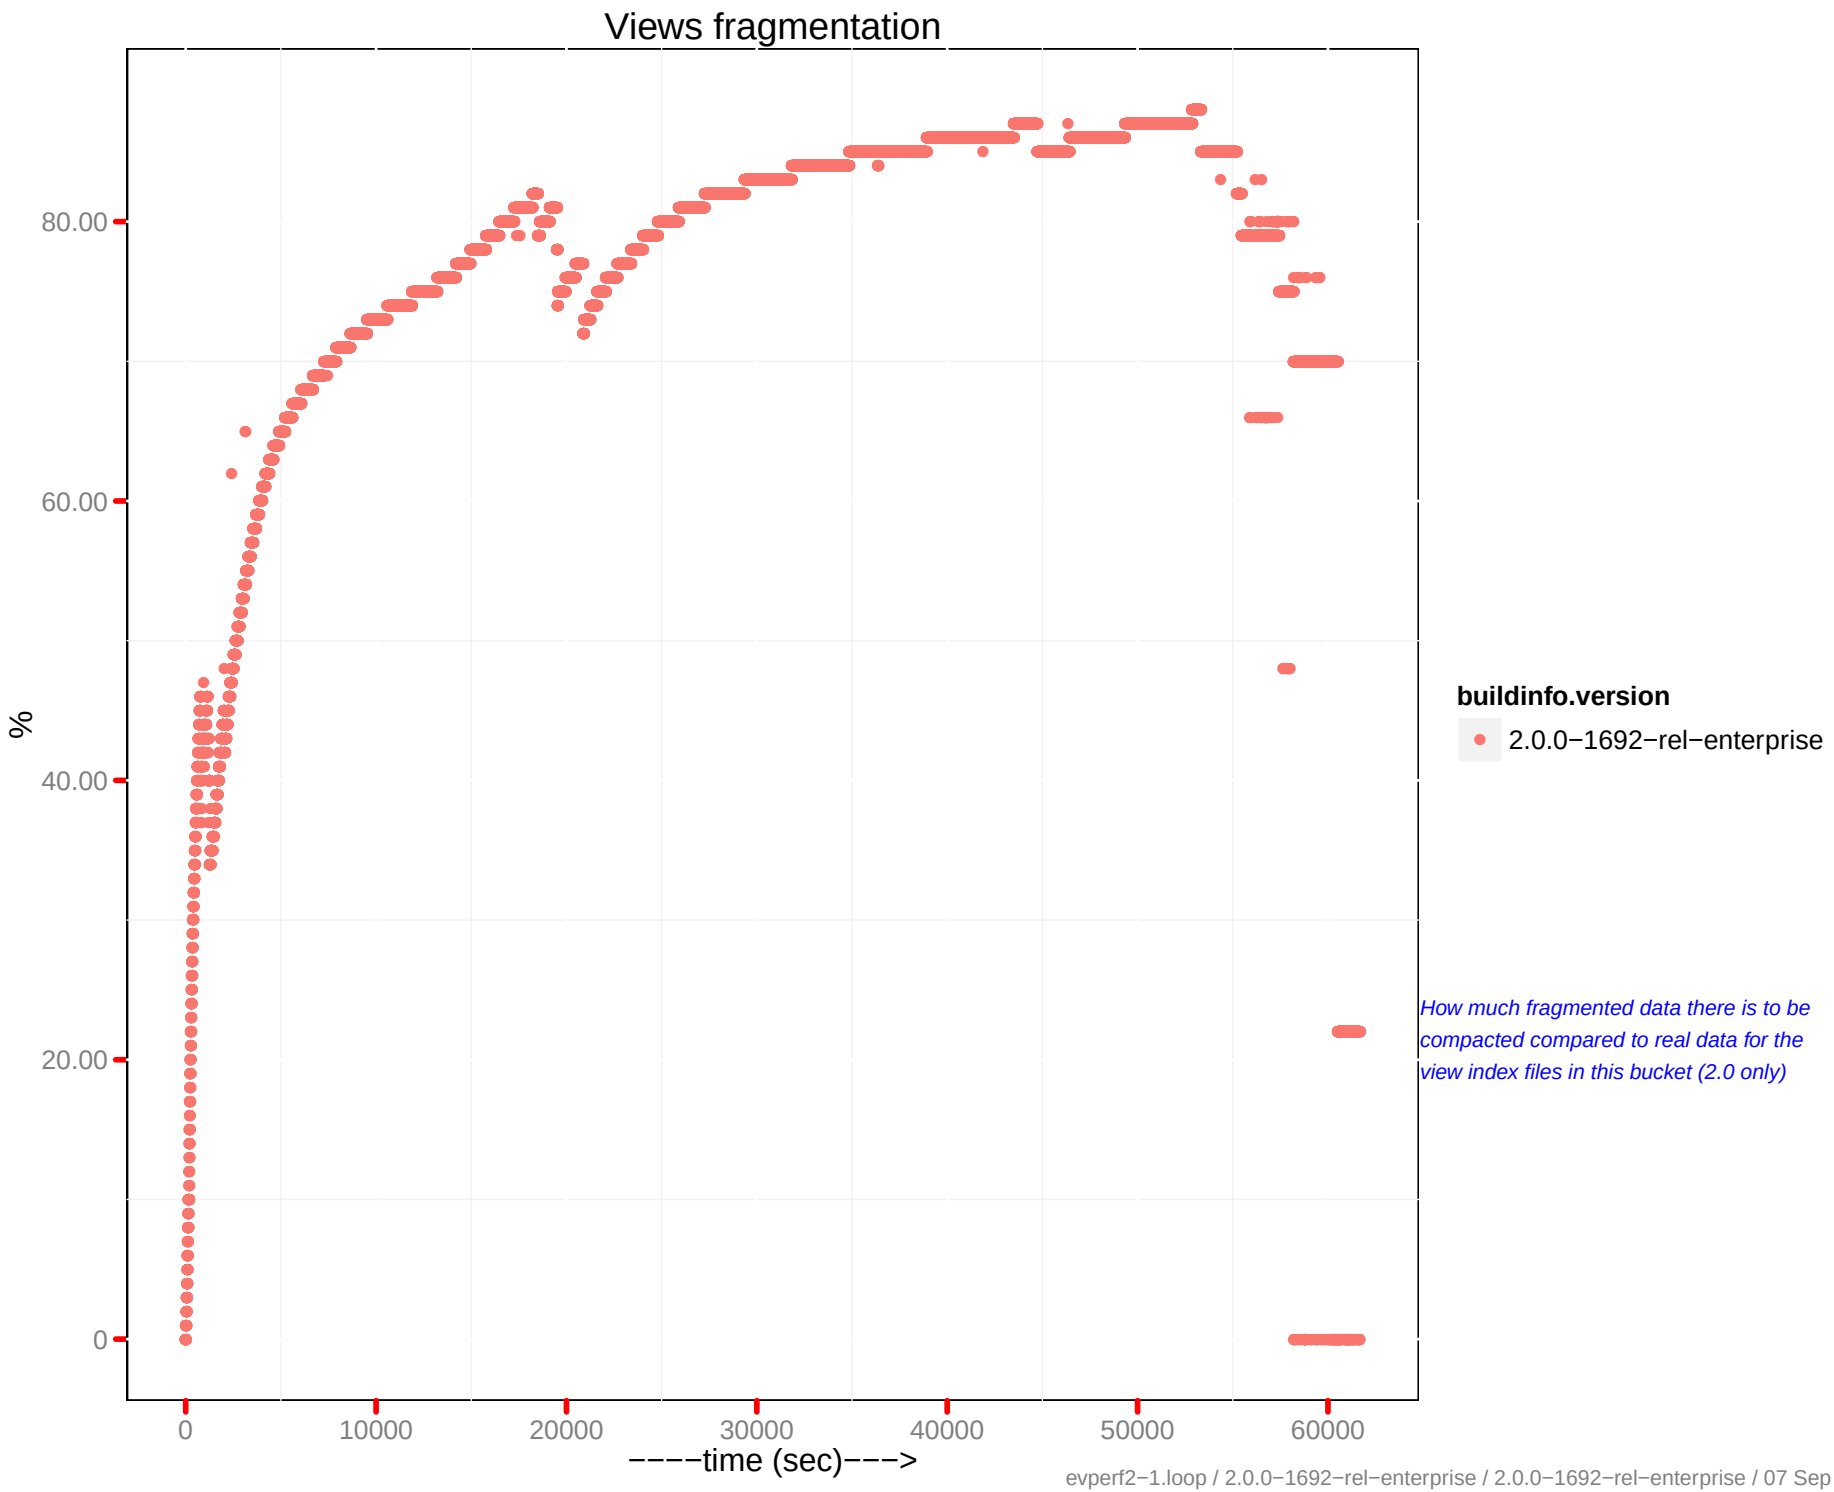
Views fragmentation
80.00
60.00
40.00
20.00
0
%
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

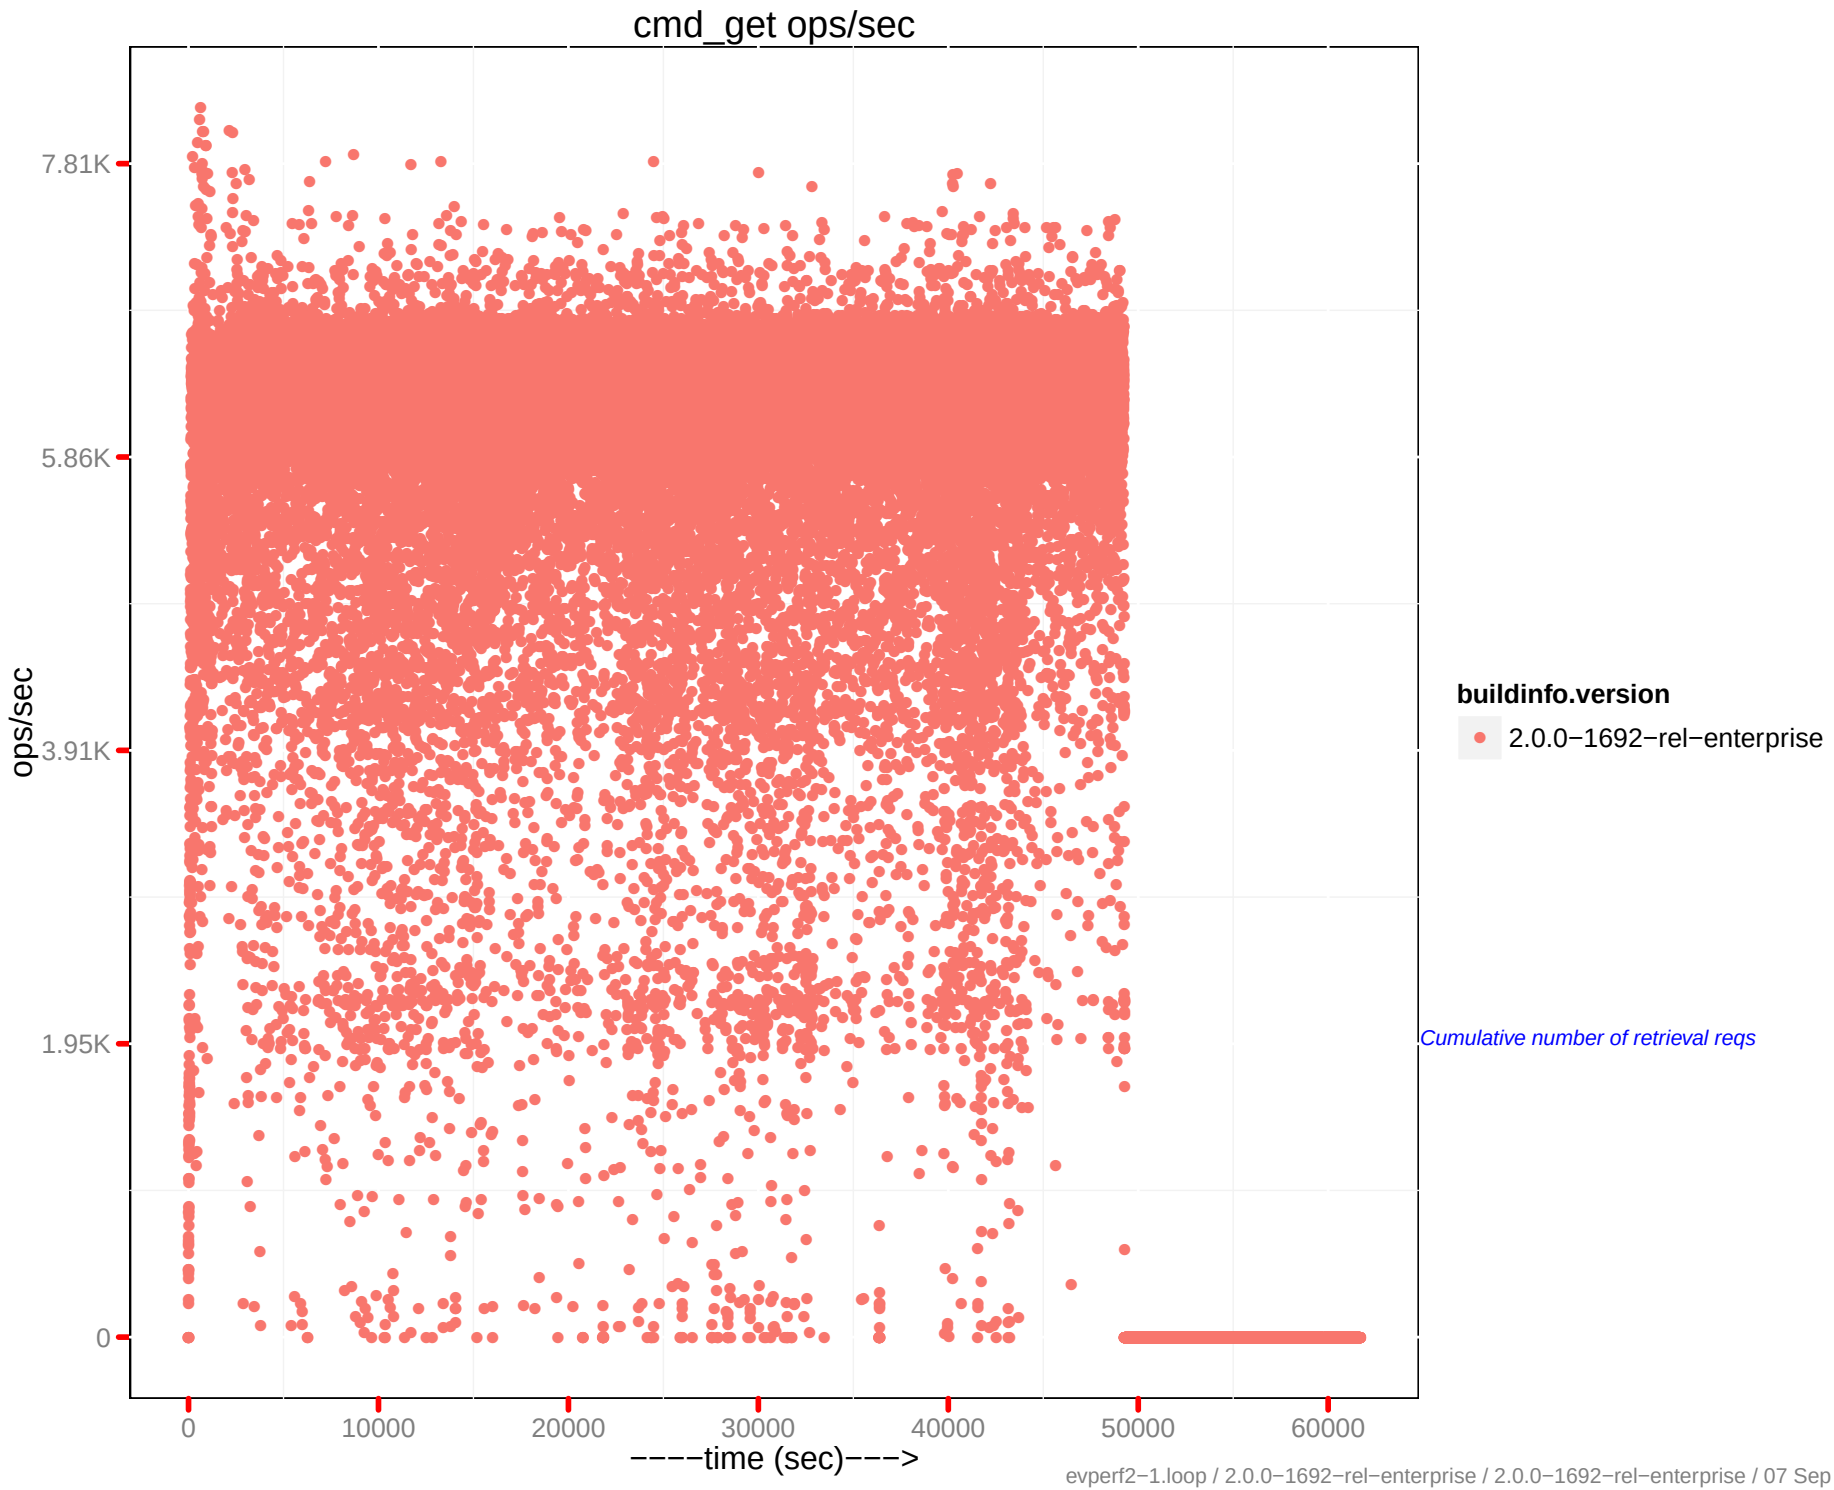
cmd_get ops/sec
7.81K
5.86K
3.91K
1.95K
0
ops/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Cumulative number of retrieval reqs
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

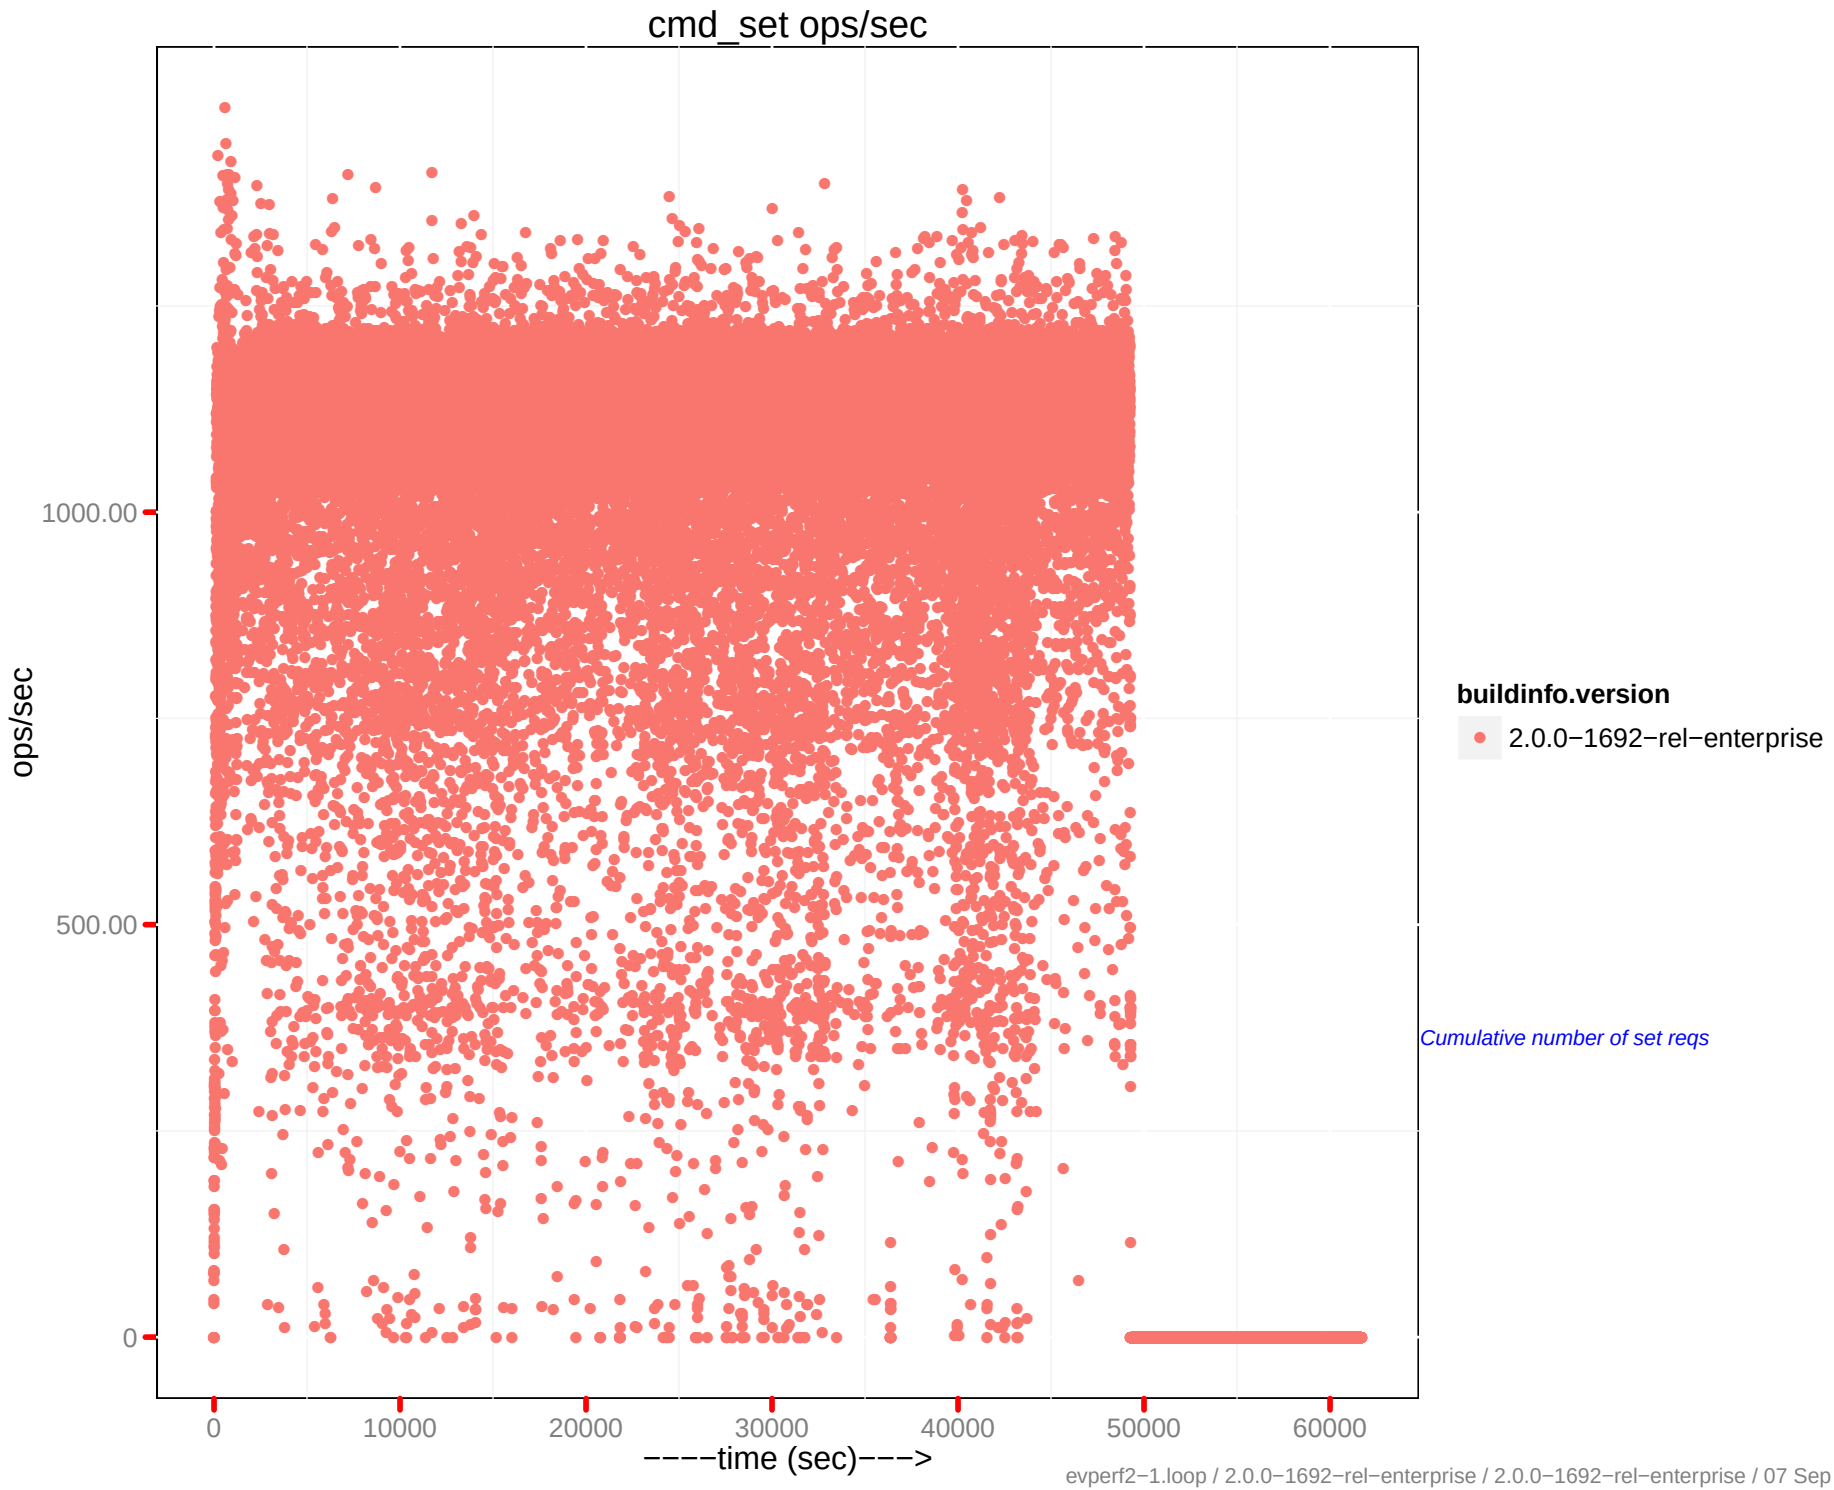
cmd_set ops/sec
ops/sec
1000.00
500.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Cumulative number of set reqs
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012


# of get misses
1.95K
1.46K
1000.00
500.00
0
# of get misses
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of get operations per second
for data that the bucket does not contain
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

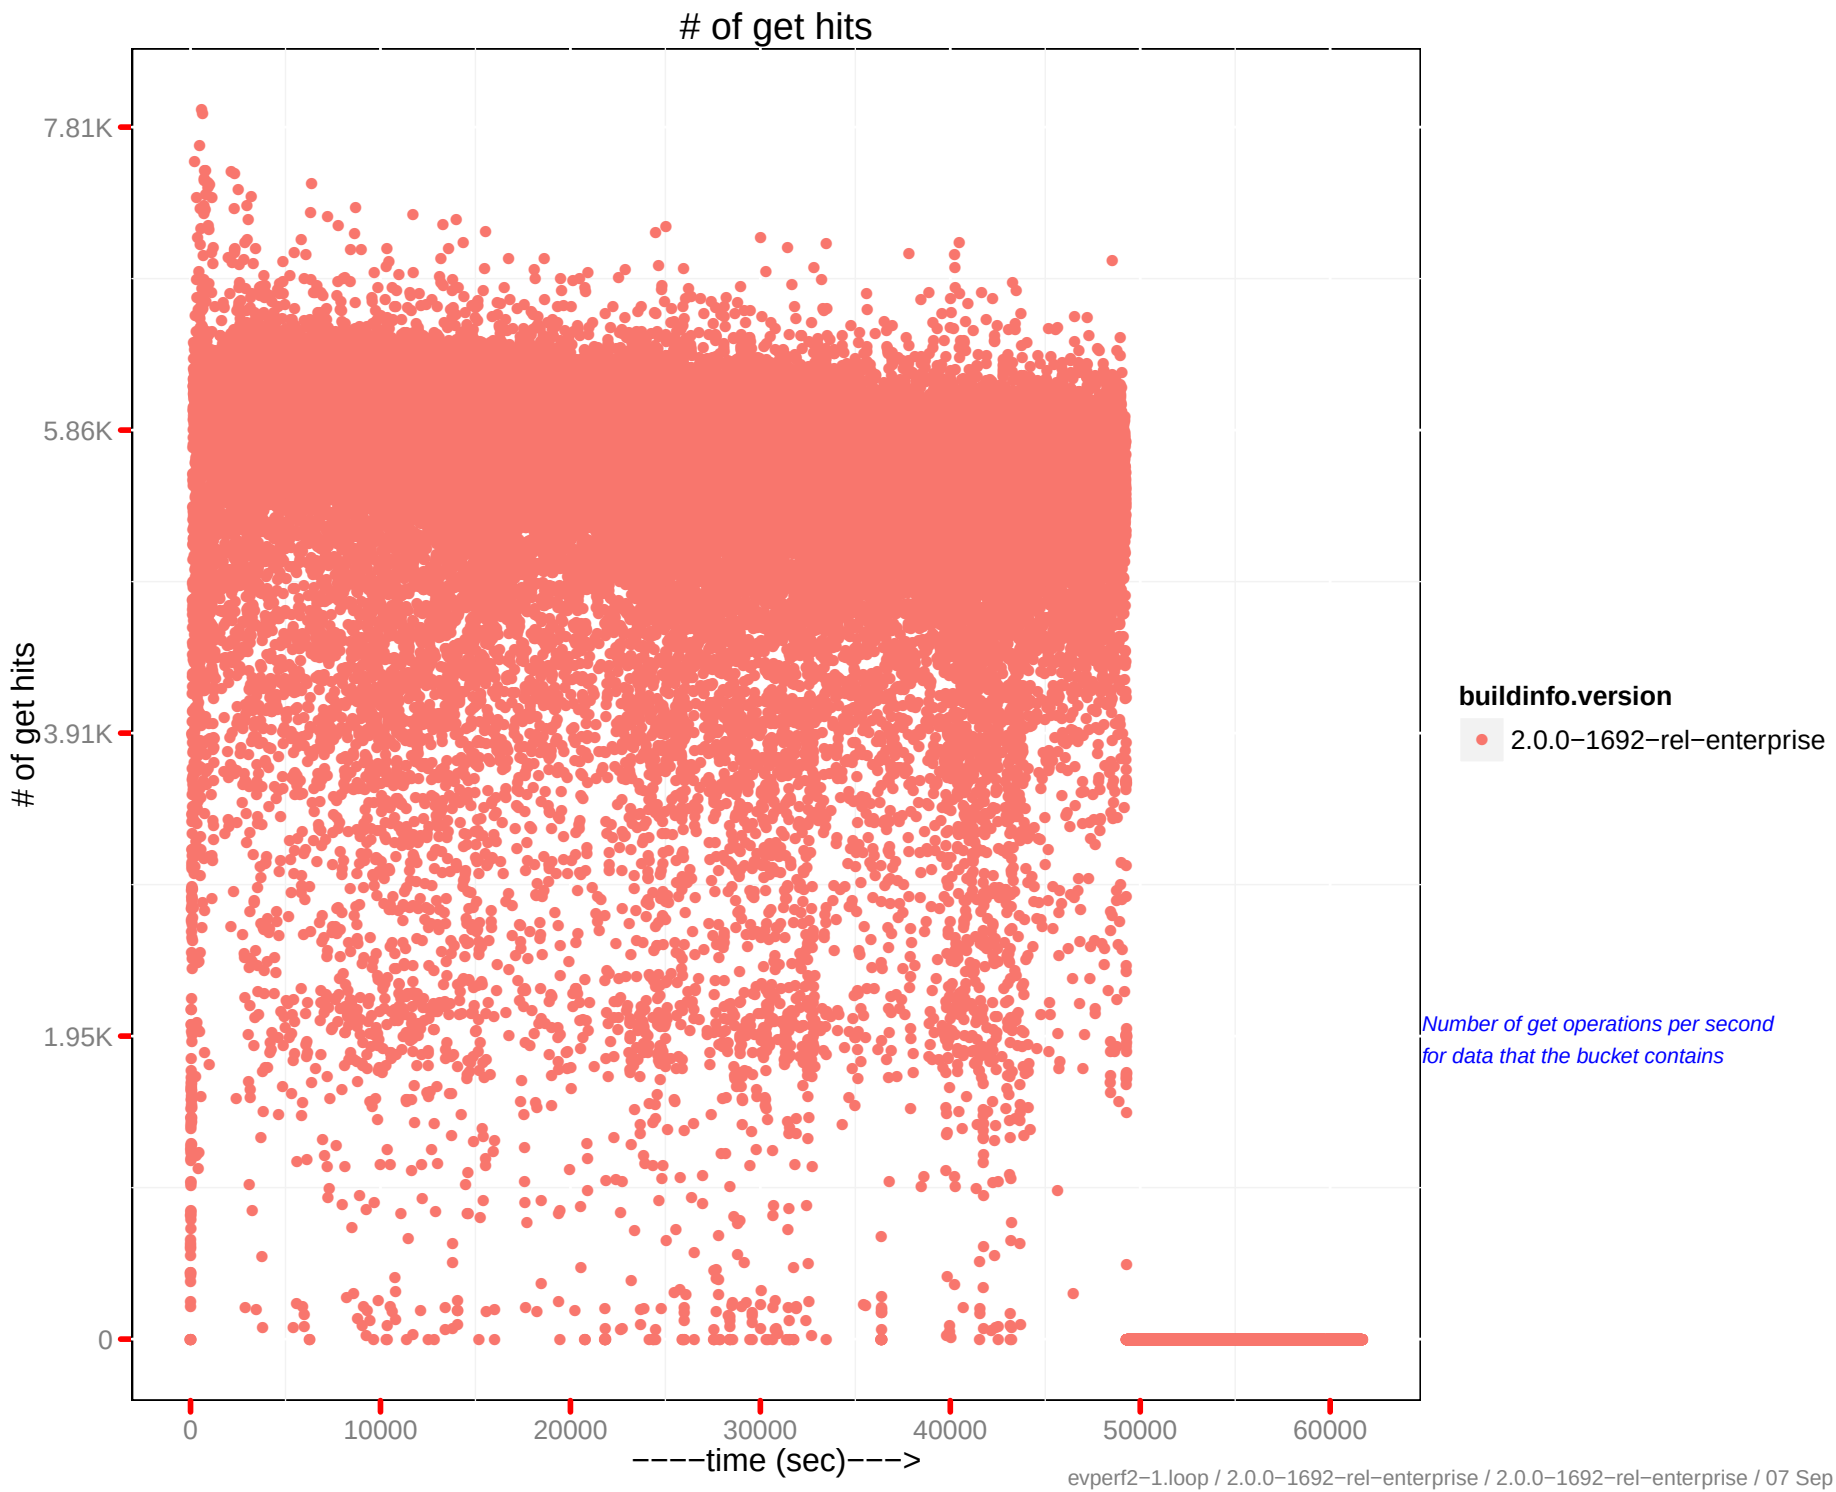
# of get hits
7.81K
5.86K
3.91K
1.95K
0
# of get hits
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of get operations per second
for data that the bucket contains
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

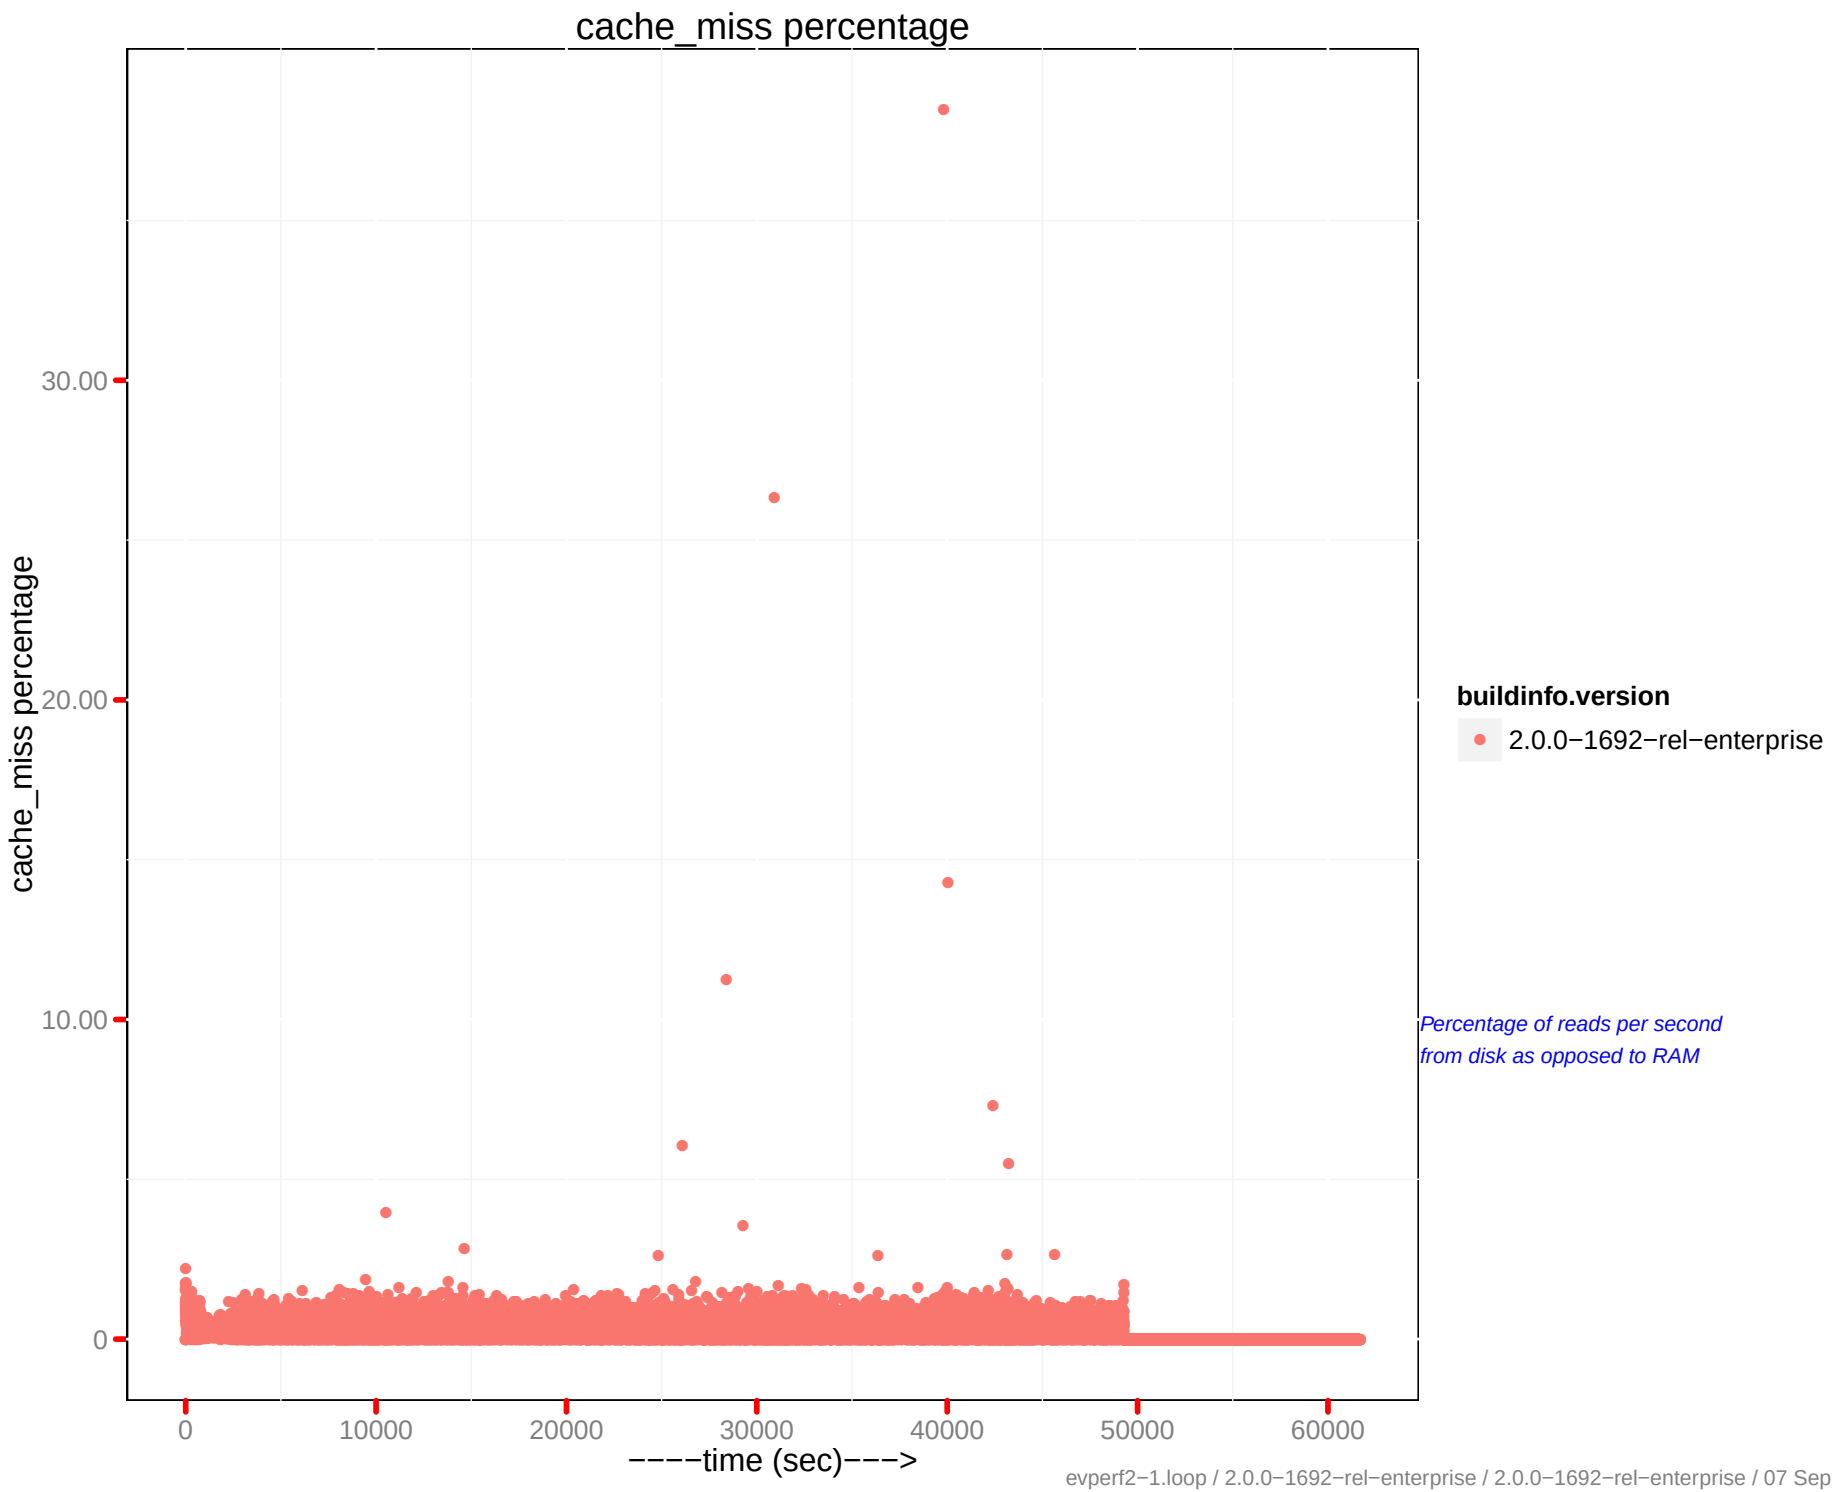
cache_miss percentage
cache_miss percentage
30.00
20.00
10.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Percentage of reads per second
from disk as opposed to RAM
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

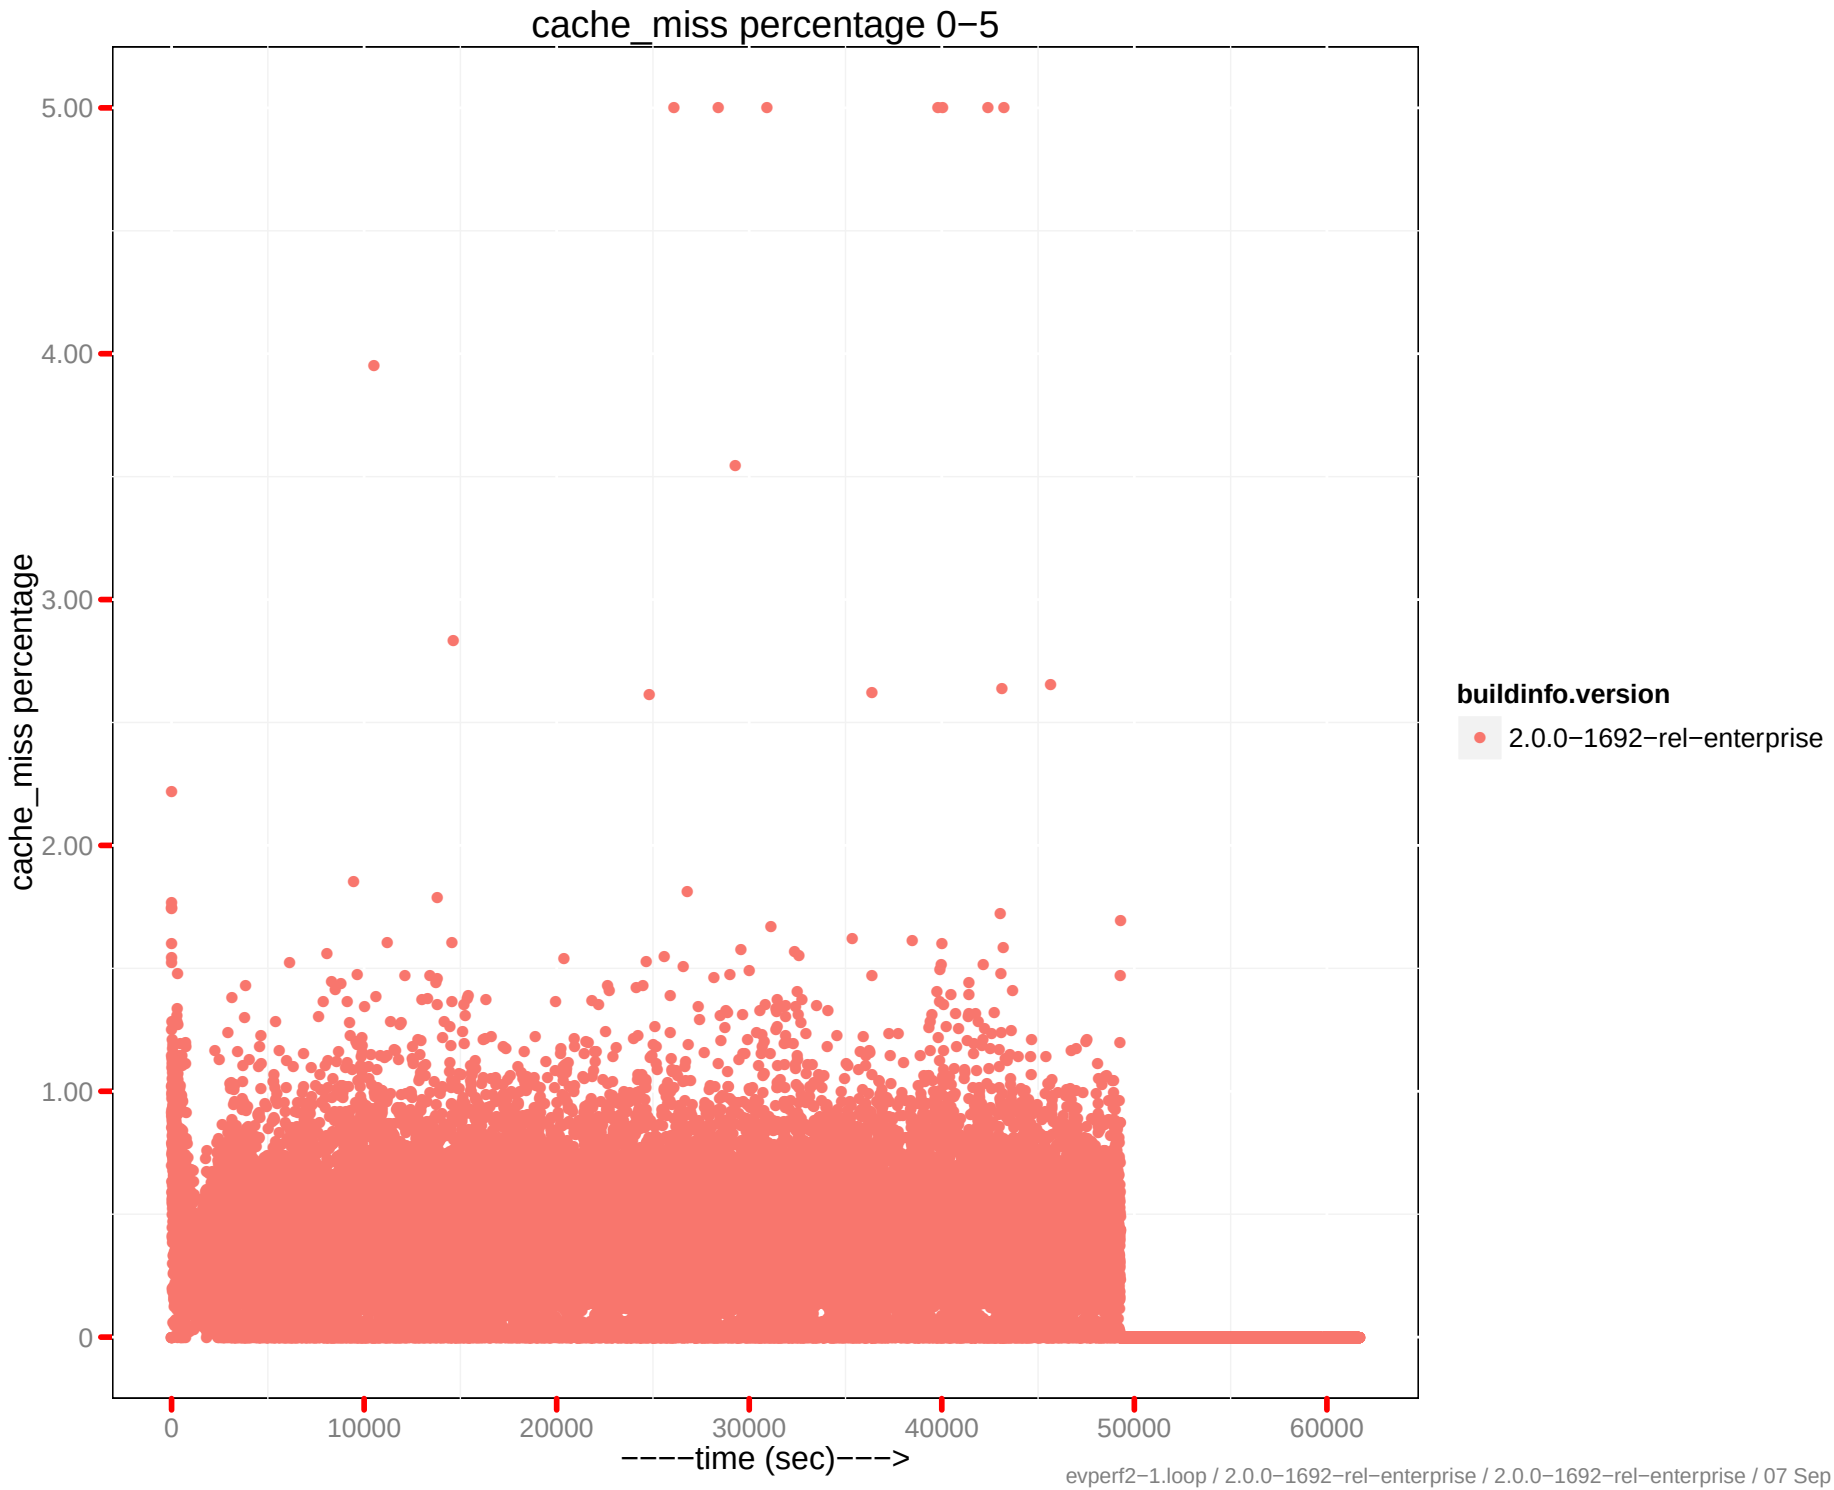
cache_miss percentage 0−5
cache_miss percentage
5.00
4.00
3.00
2.00
1.00
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

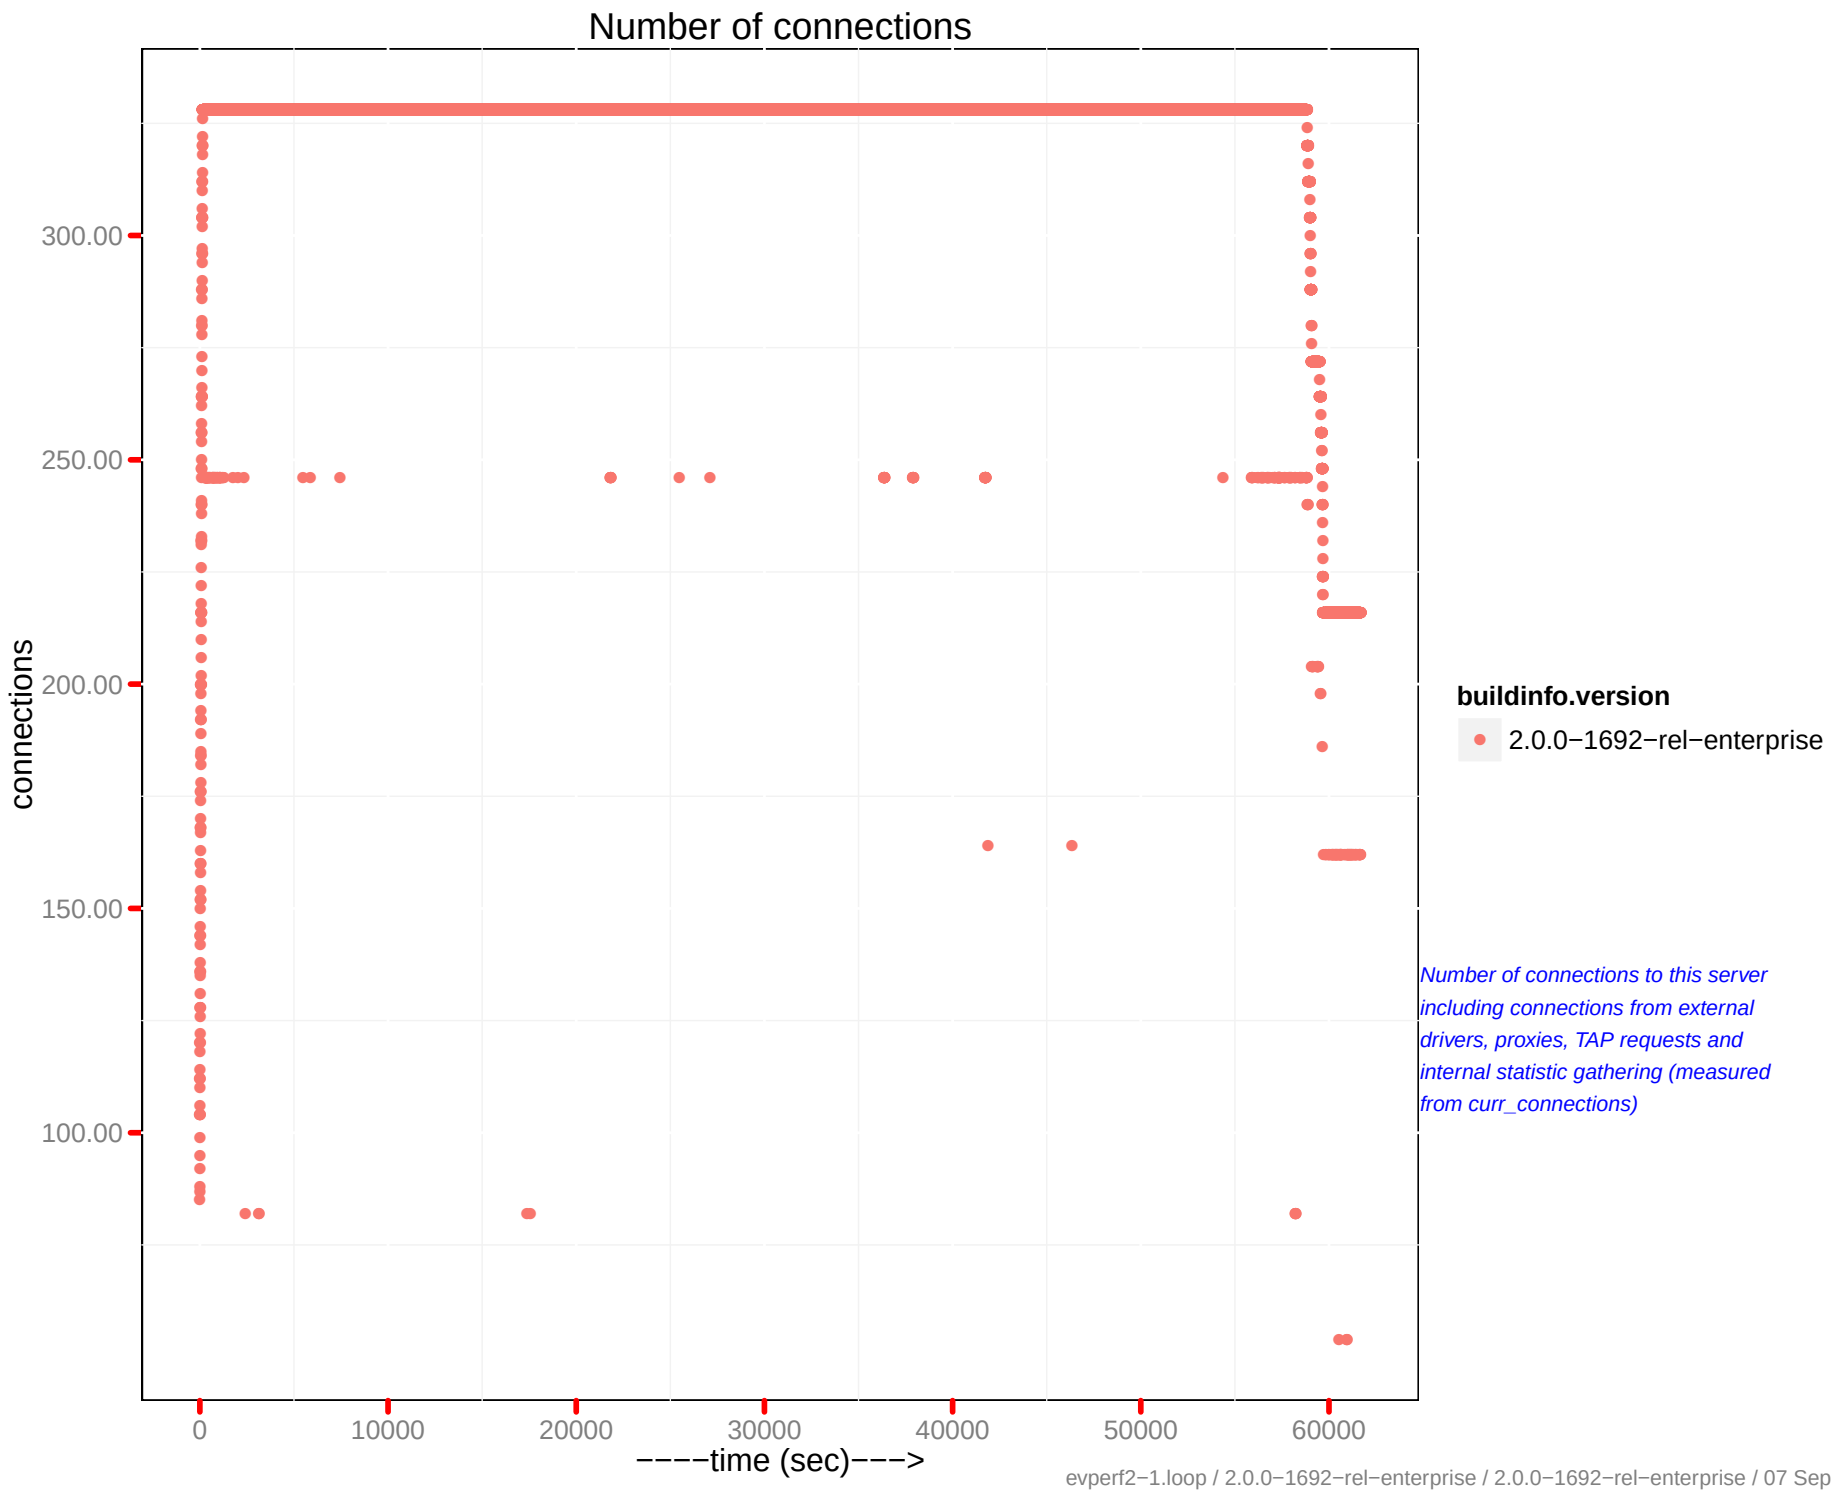
Number of connections
connections
300.00
250.00
200.00
150.00
100.00
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of connections to this server including connections from external drivers, proxies, TAP requests and internal statistic gathering (measured from curr_connections)
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

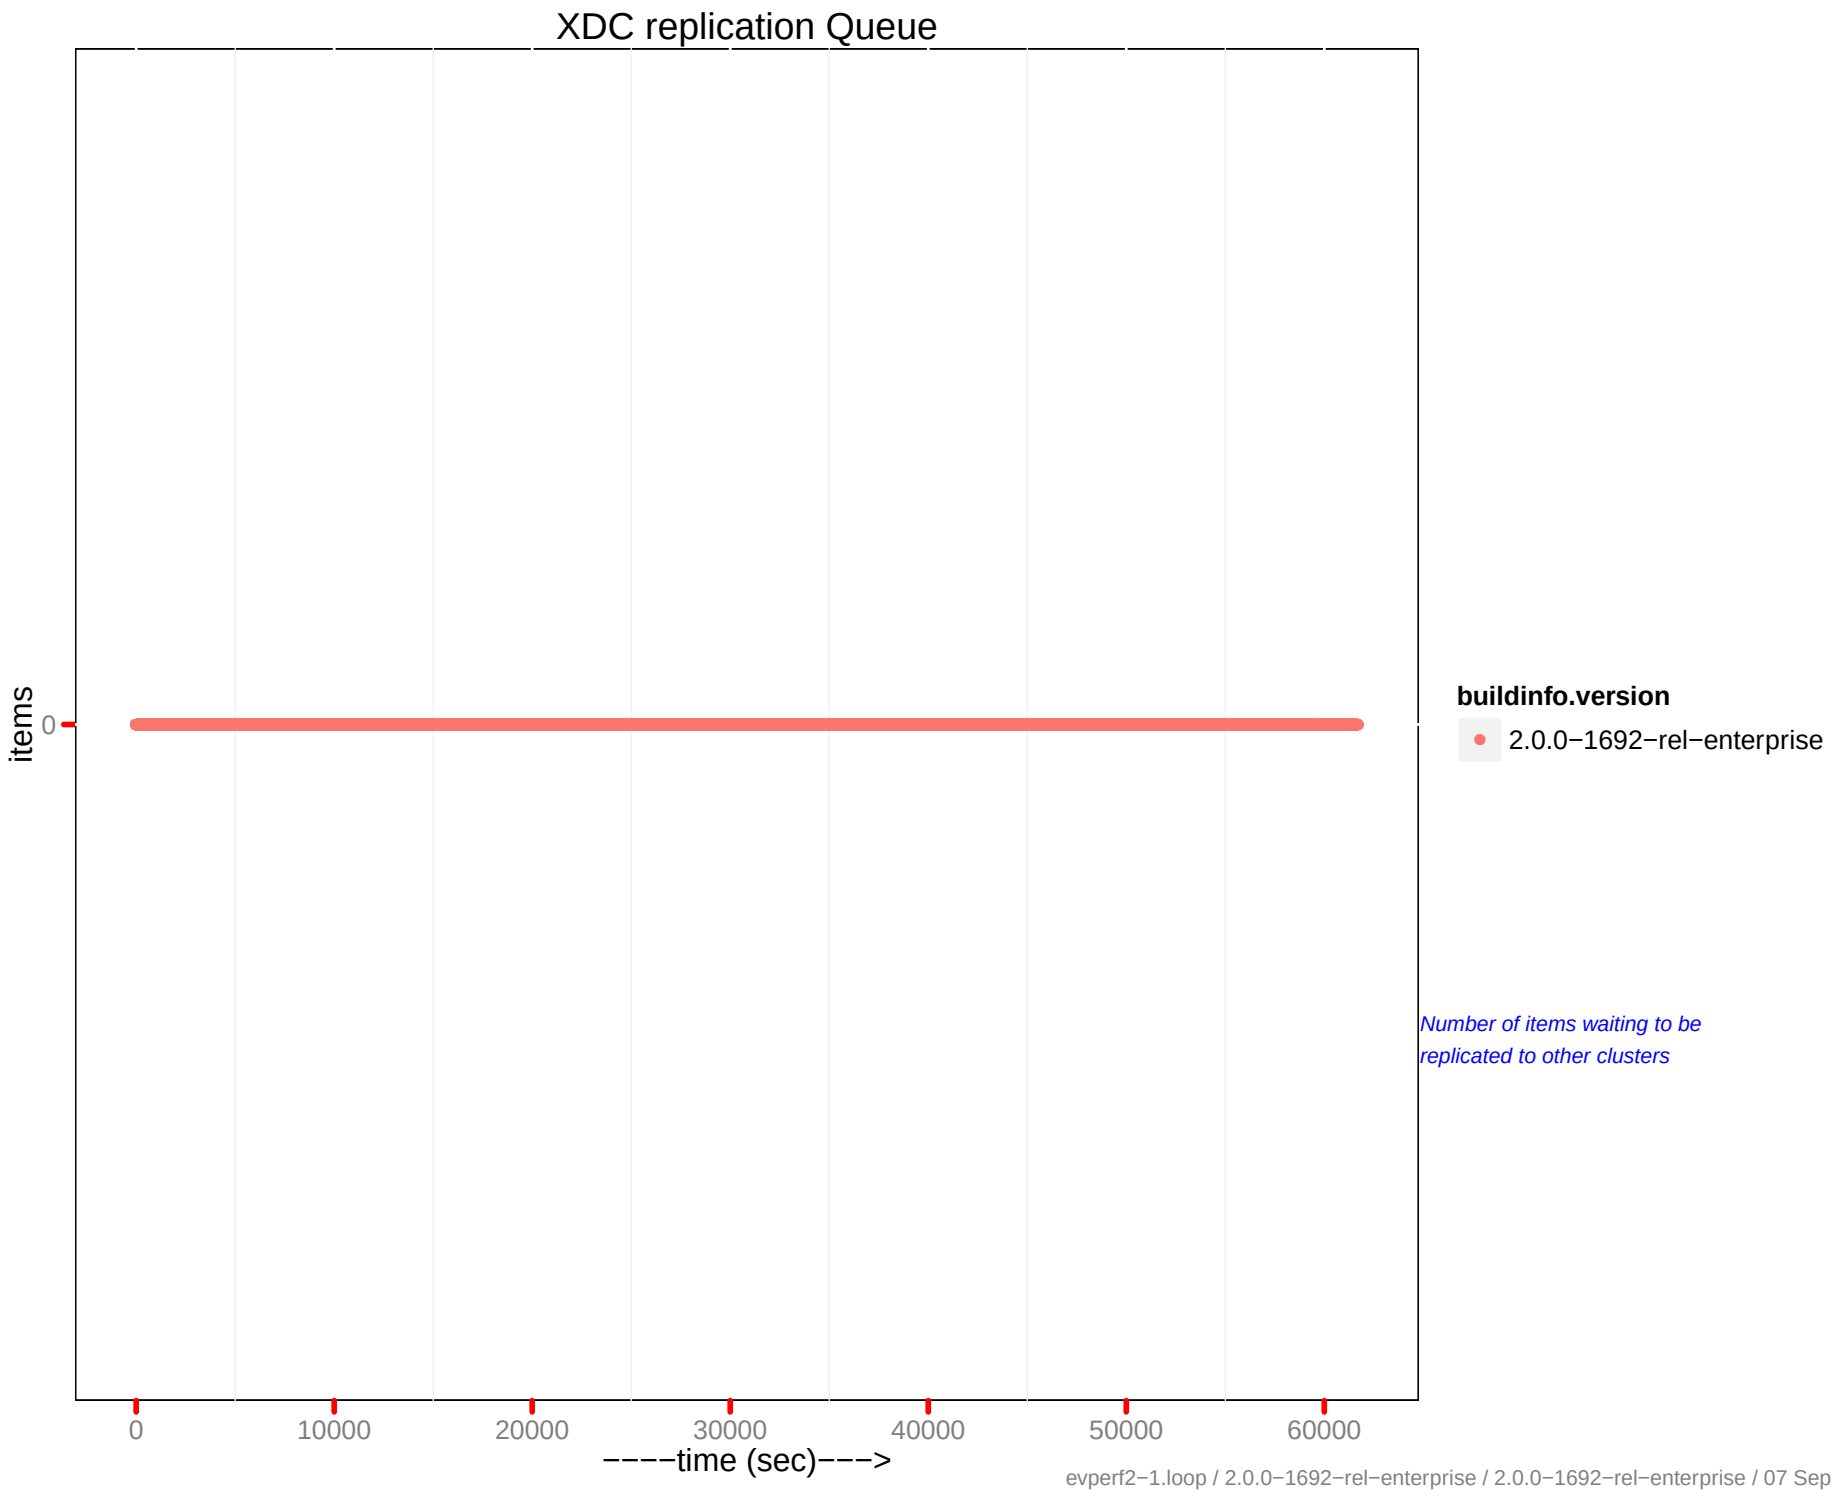
XDC replication Queue
items
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of items waiting to be replicated to other clusters
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

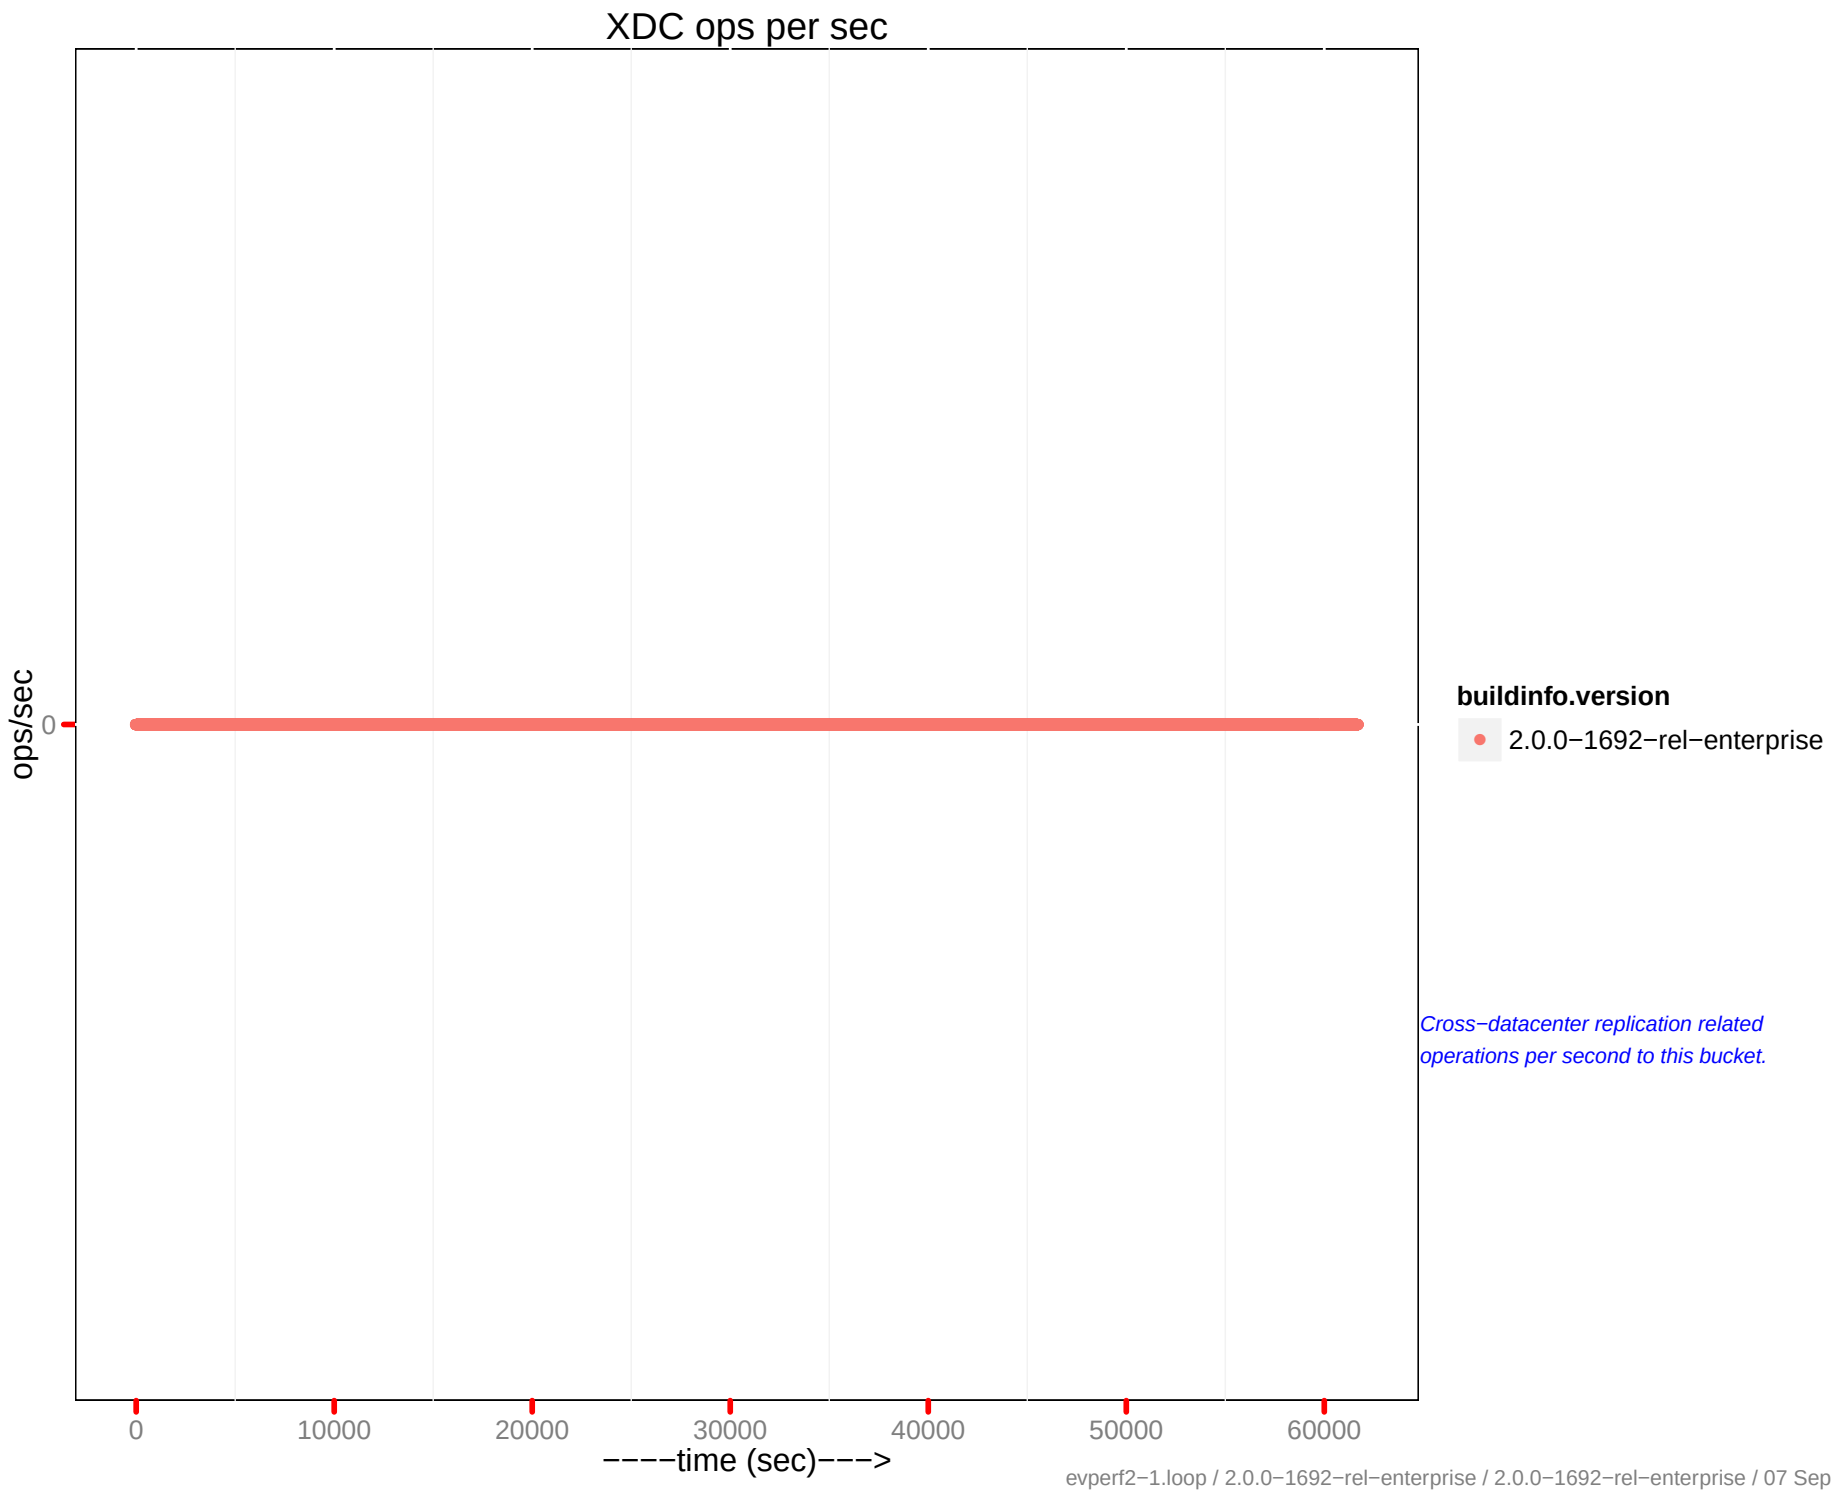
XDC ops per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
Cross−datacenter replication related
operations per second to this bucket.
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

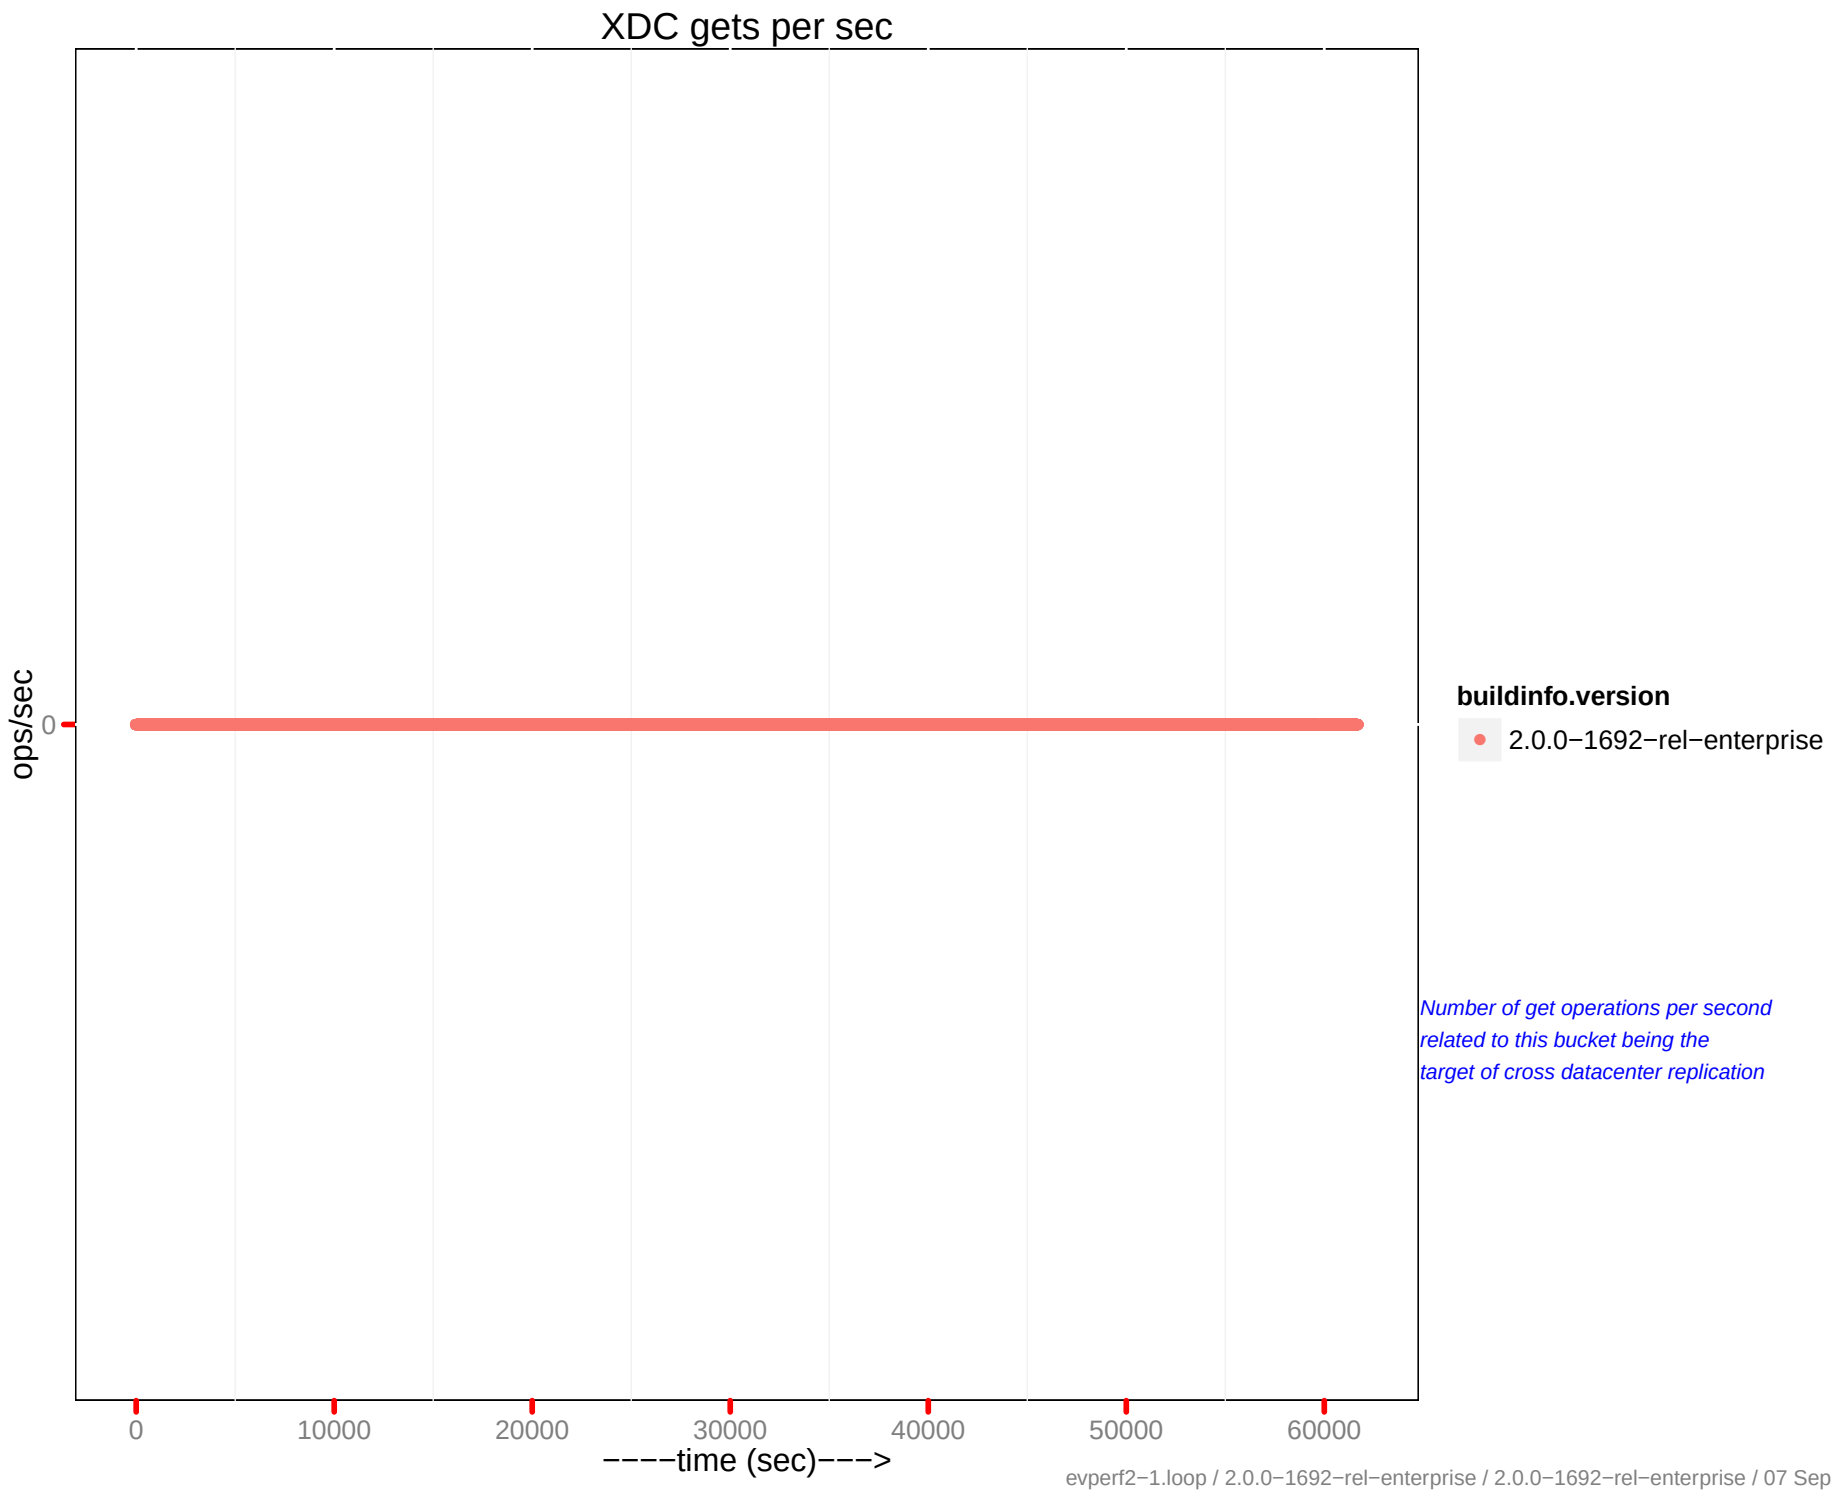
XDC gets per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of get operations per second related to this bucket being the target of cross datacenter replication
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

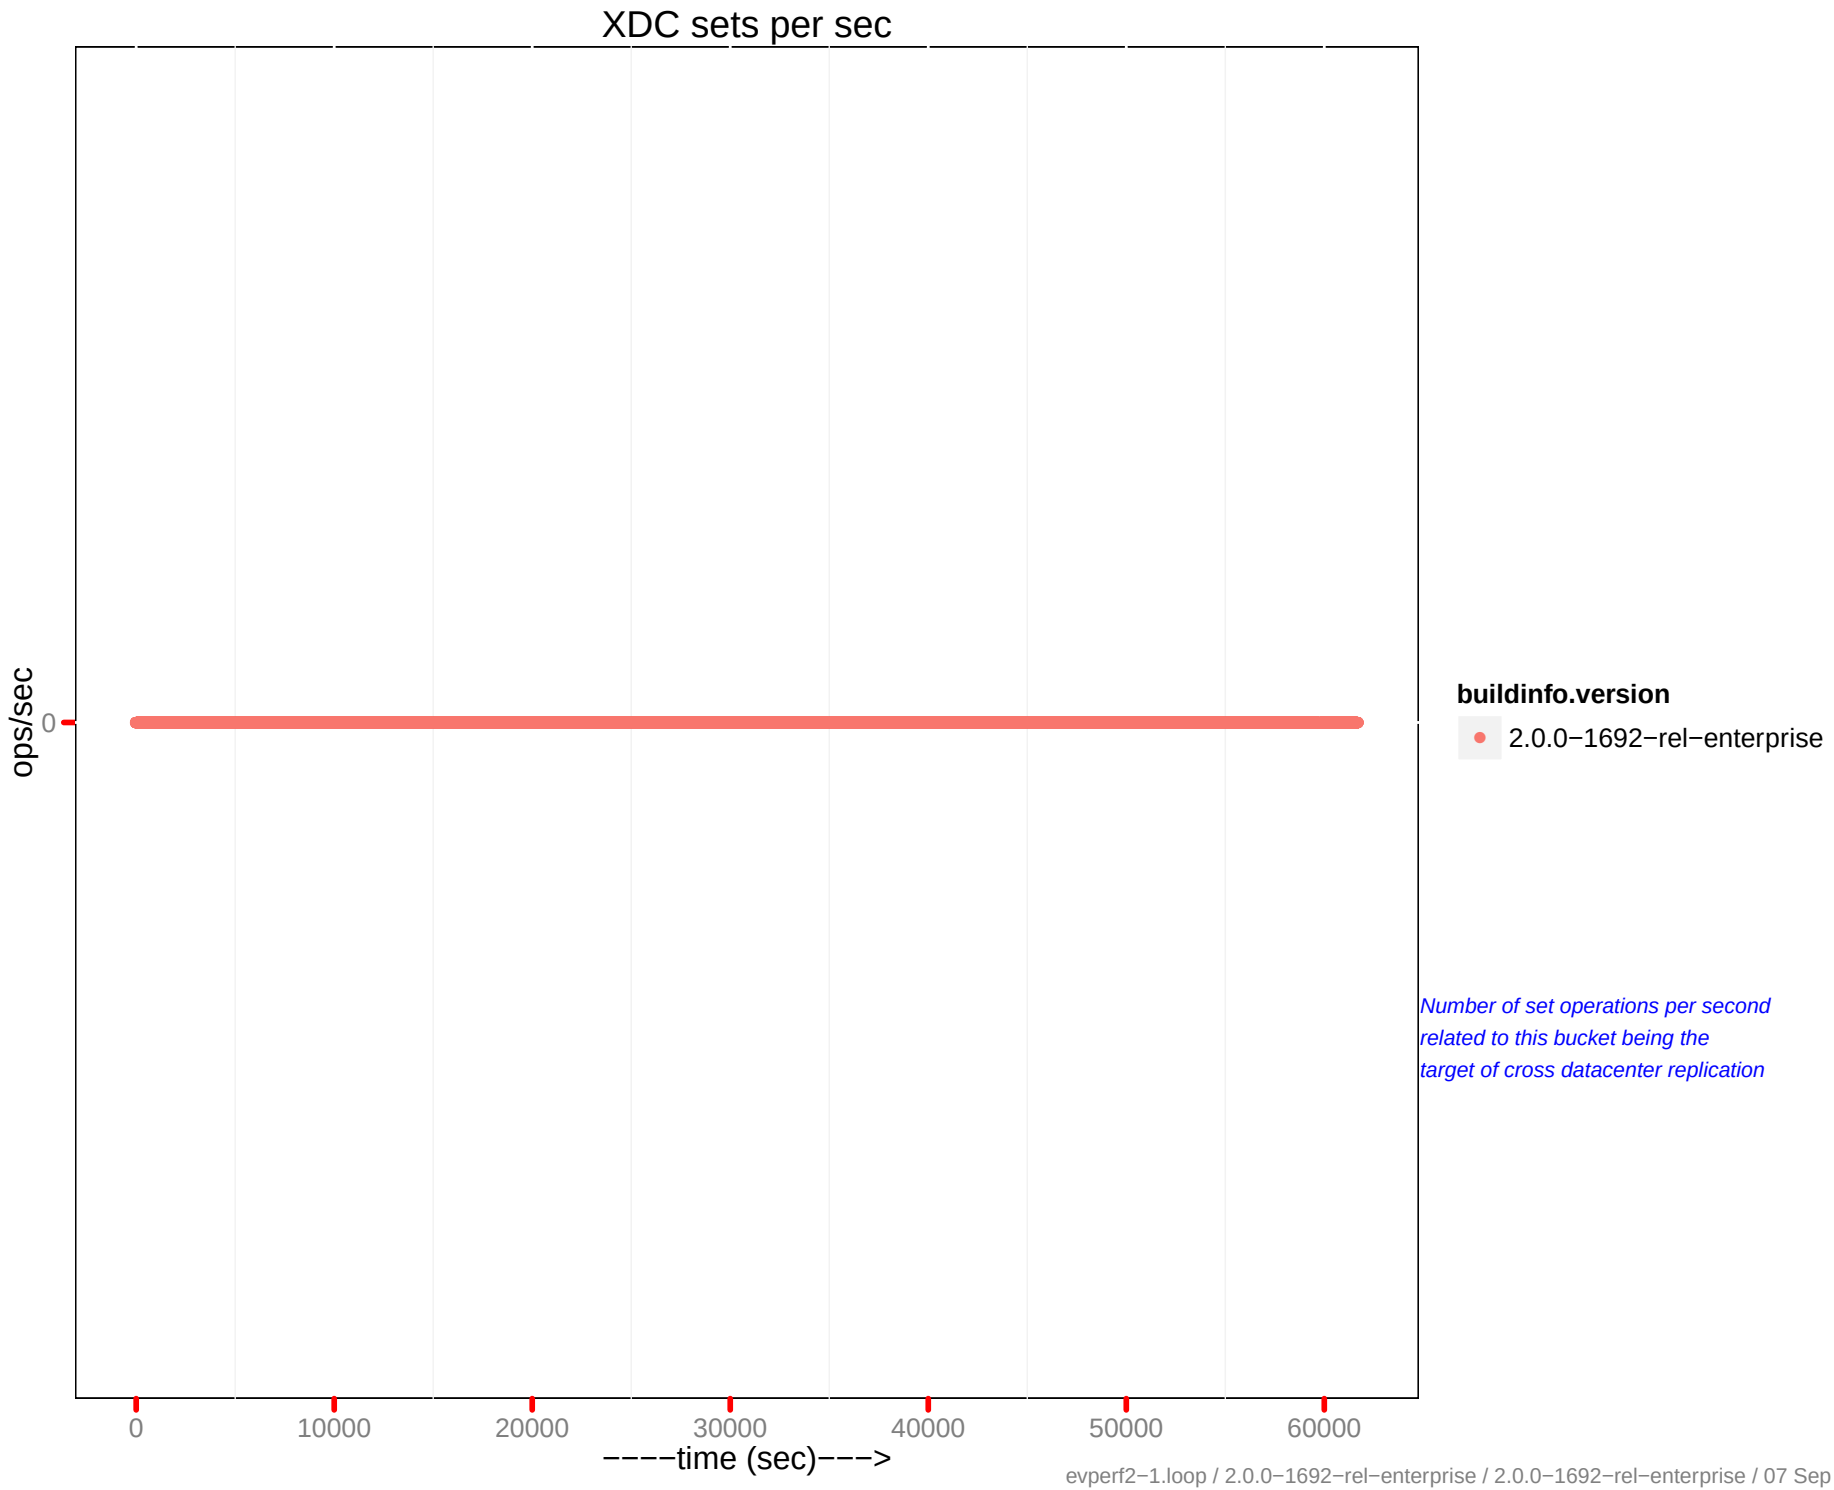
XDC sets per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Number of set operations per second
related to this bucket being the
target of cross datacenter replication
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

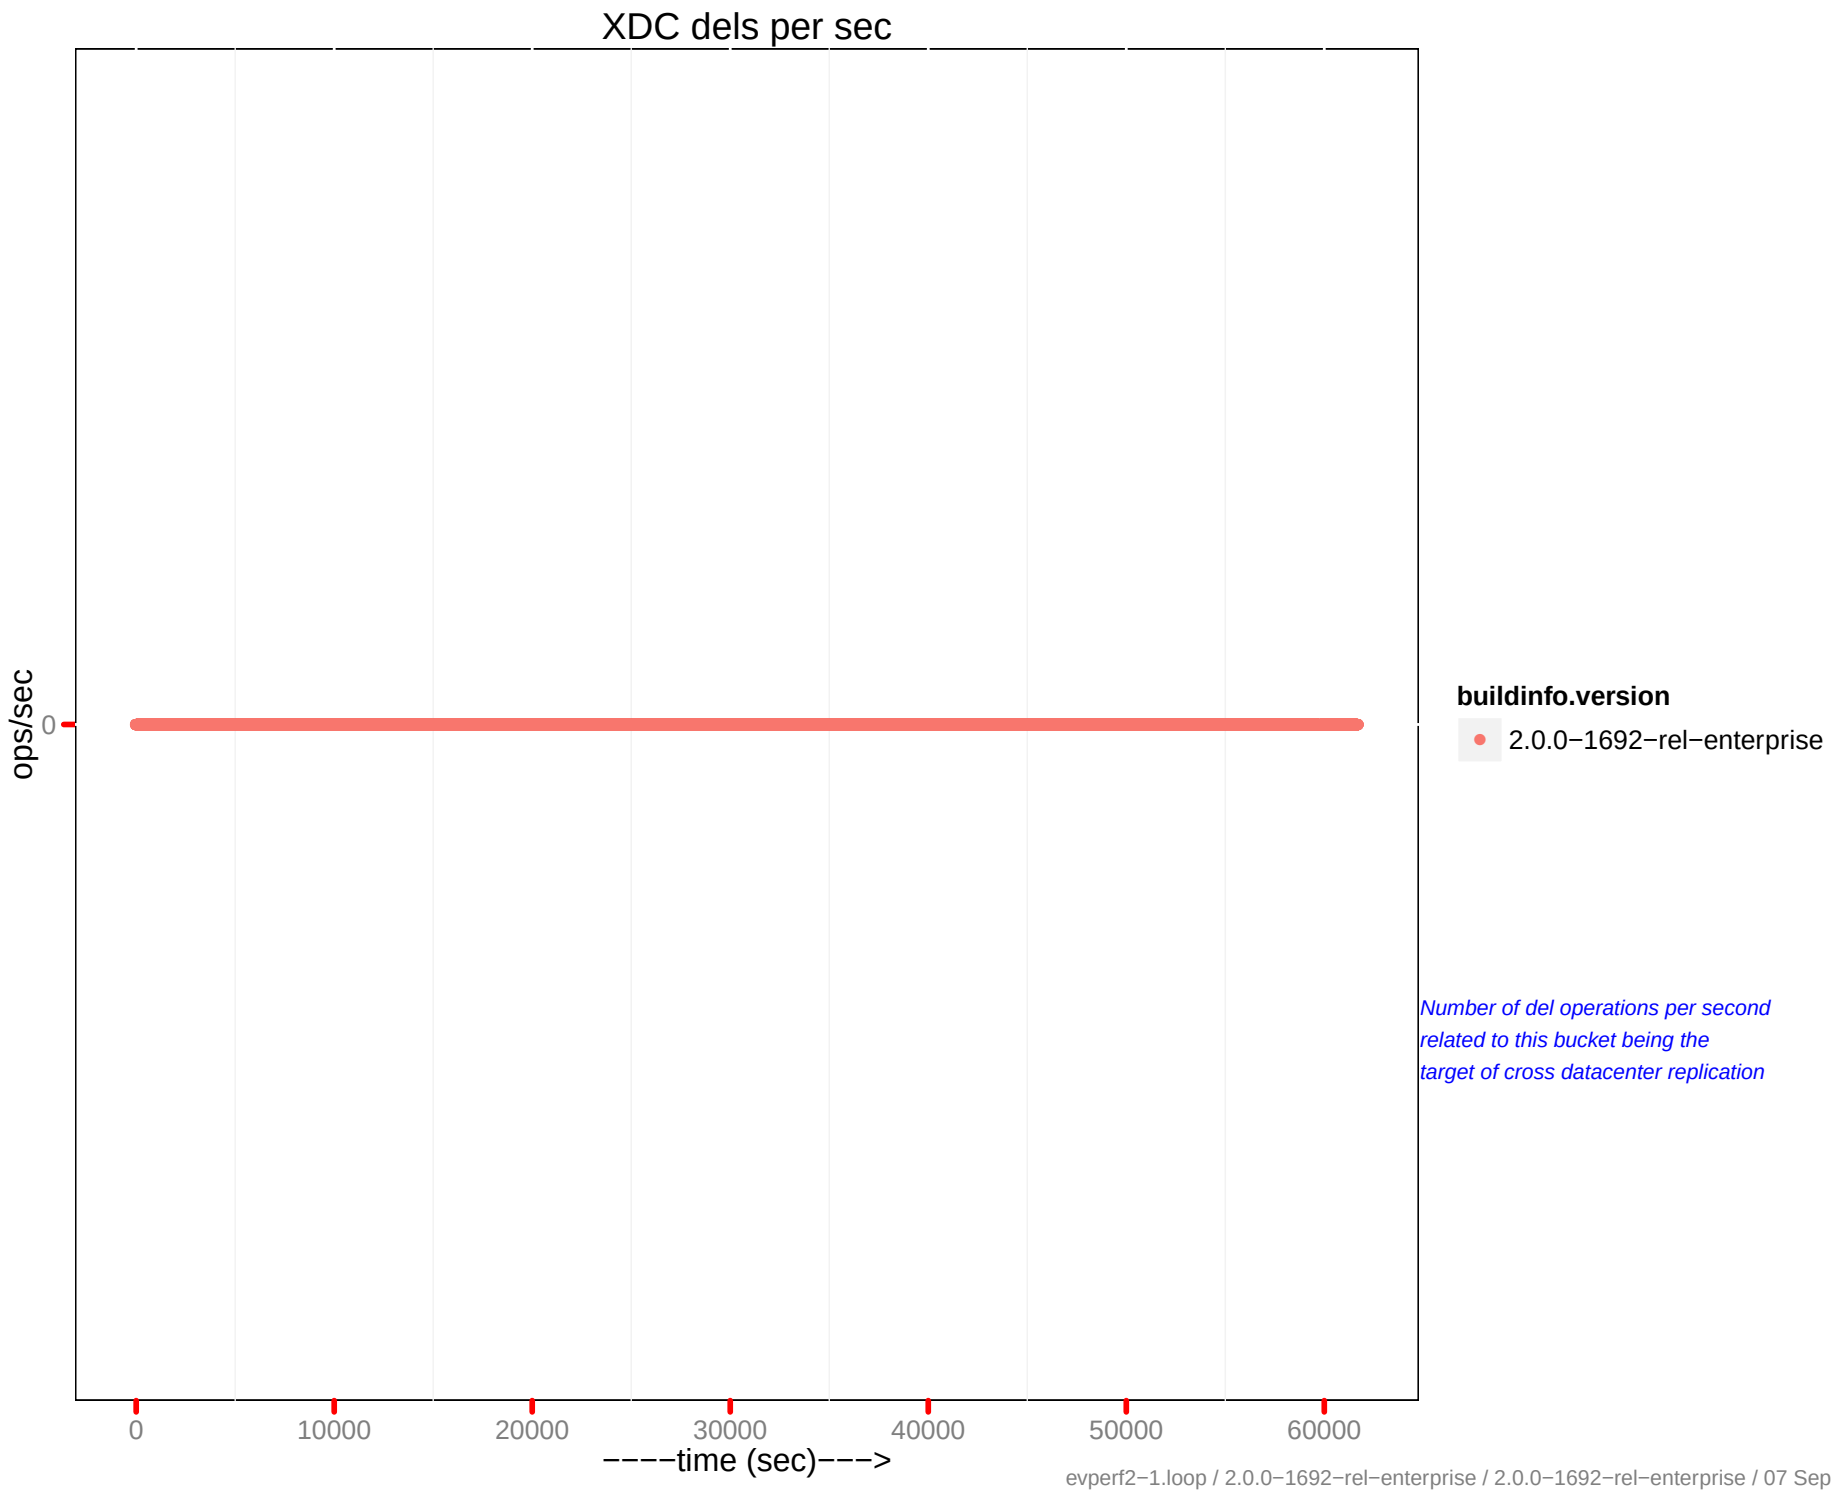
XDC dels per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
Number of del operations per second
related to this bucket being the
target of cross datacenter replication
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

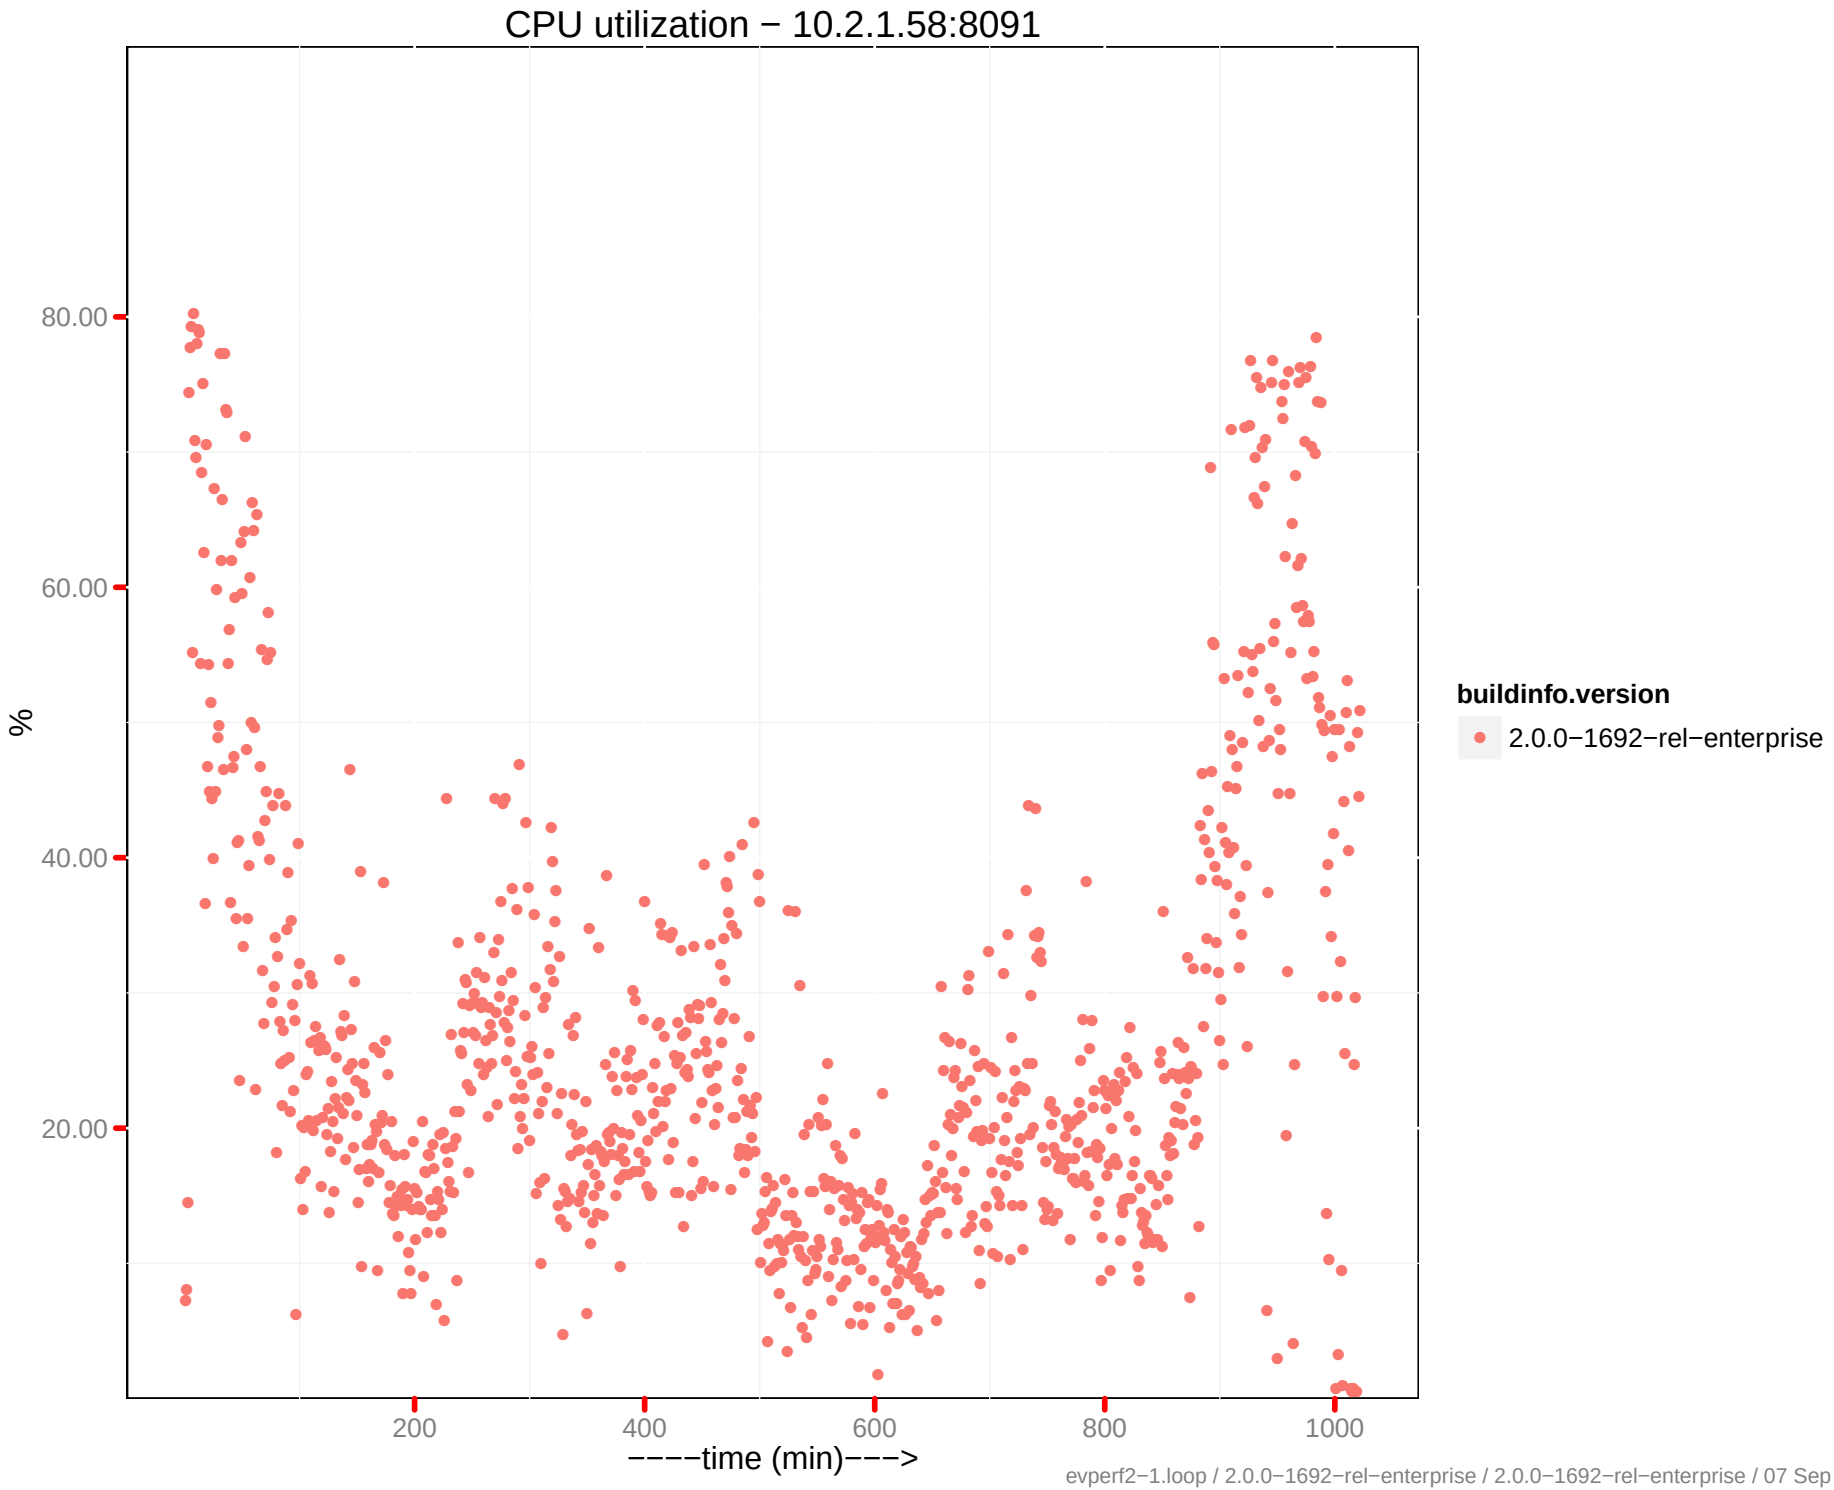
CPU utilization − 10.2.1.58:8091
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

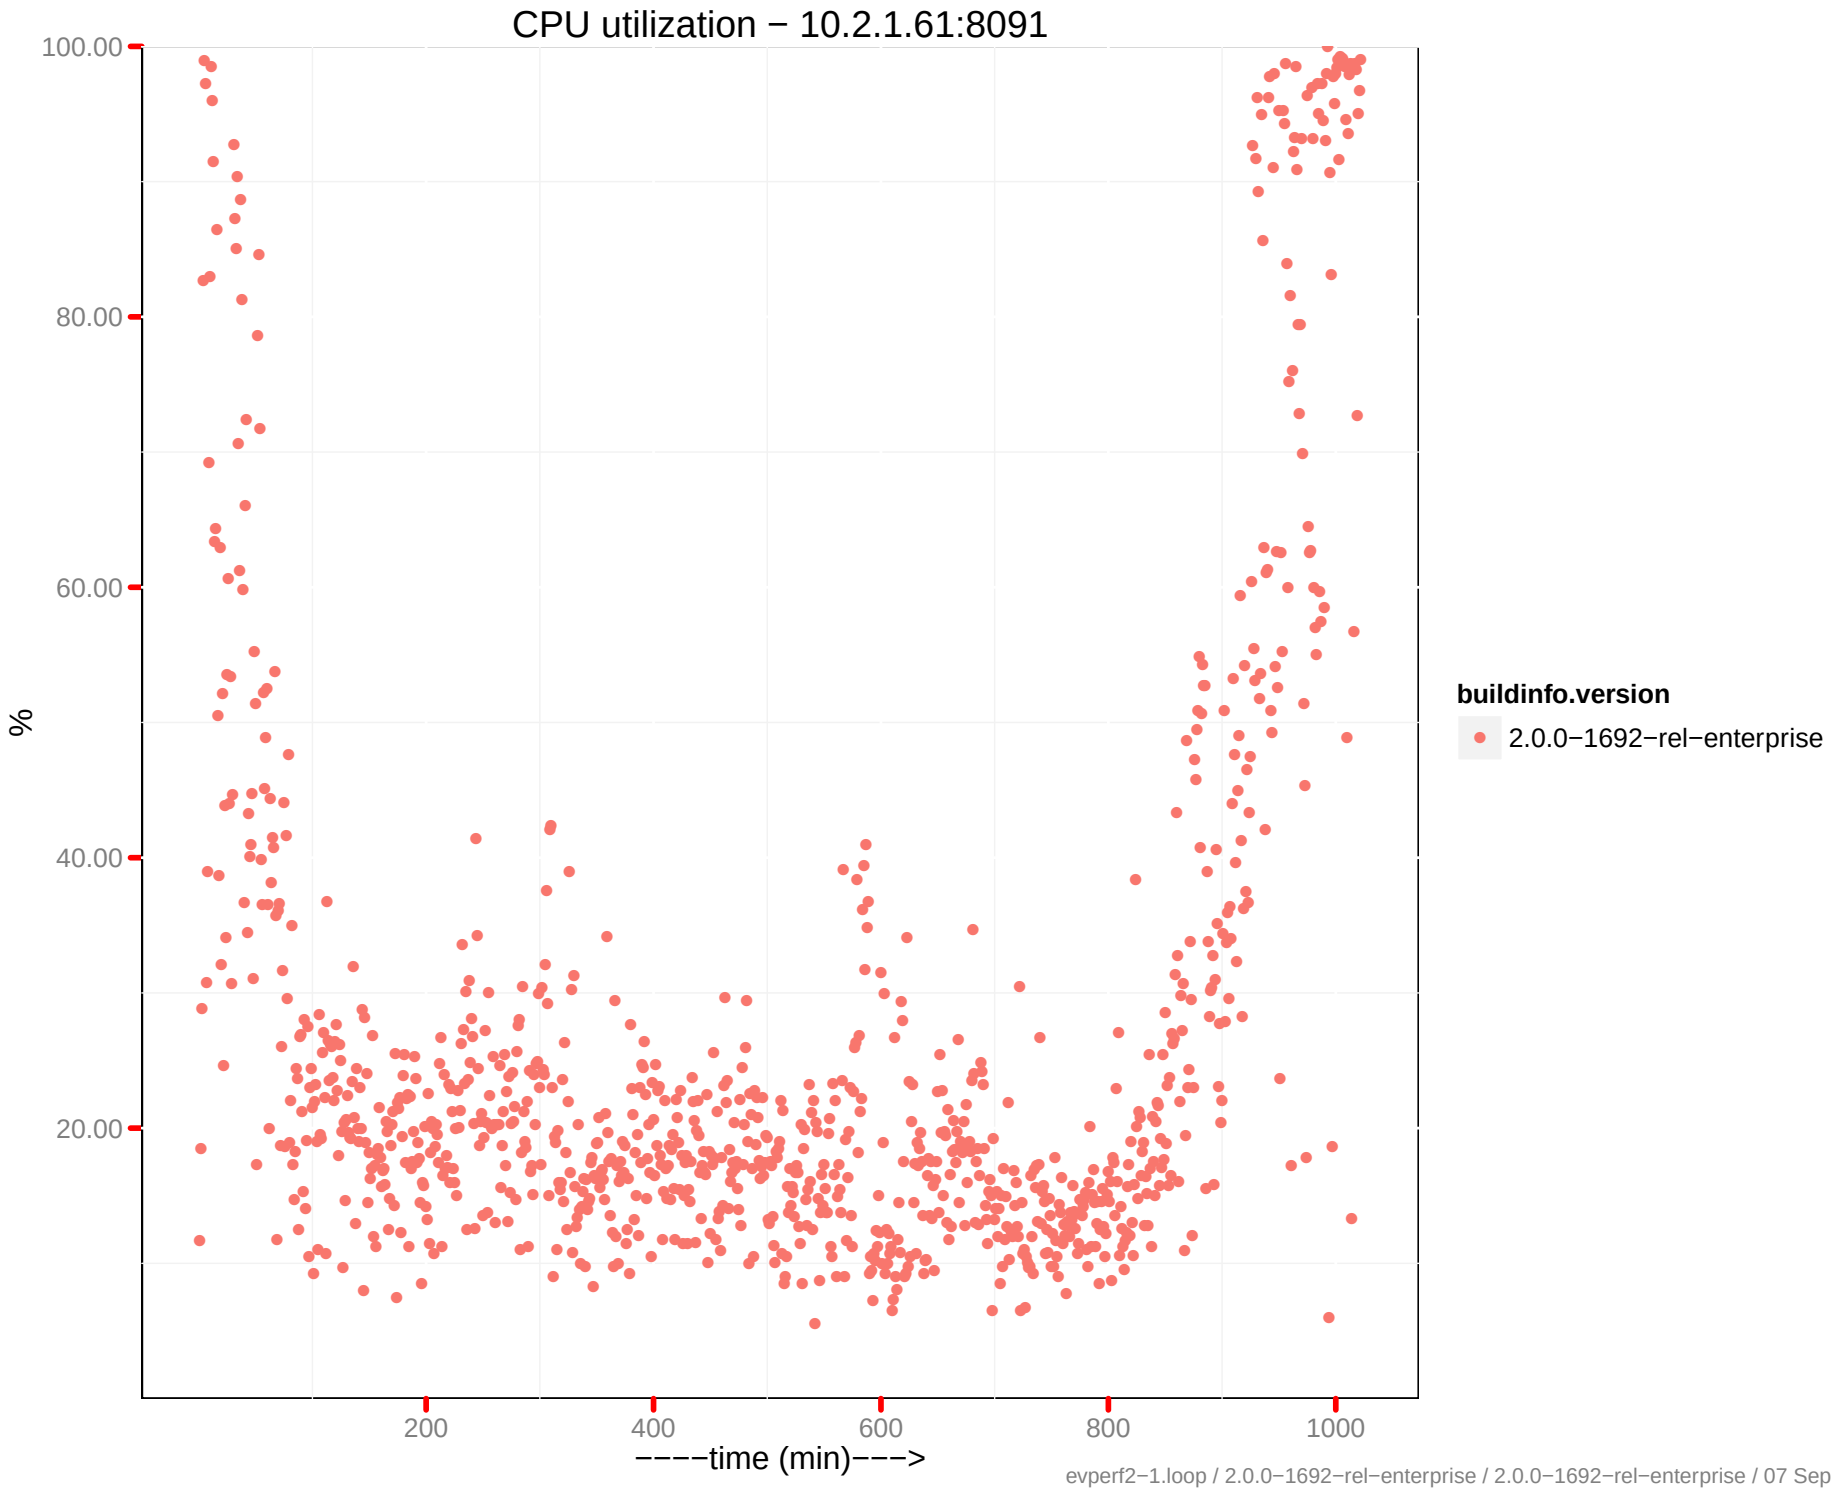
CPU utilization − 10.2.1.61:8091
100.00
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

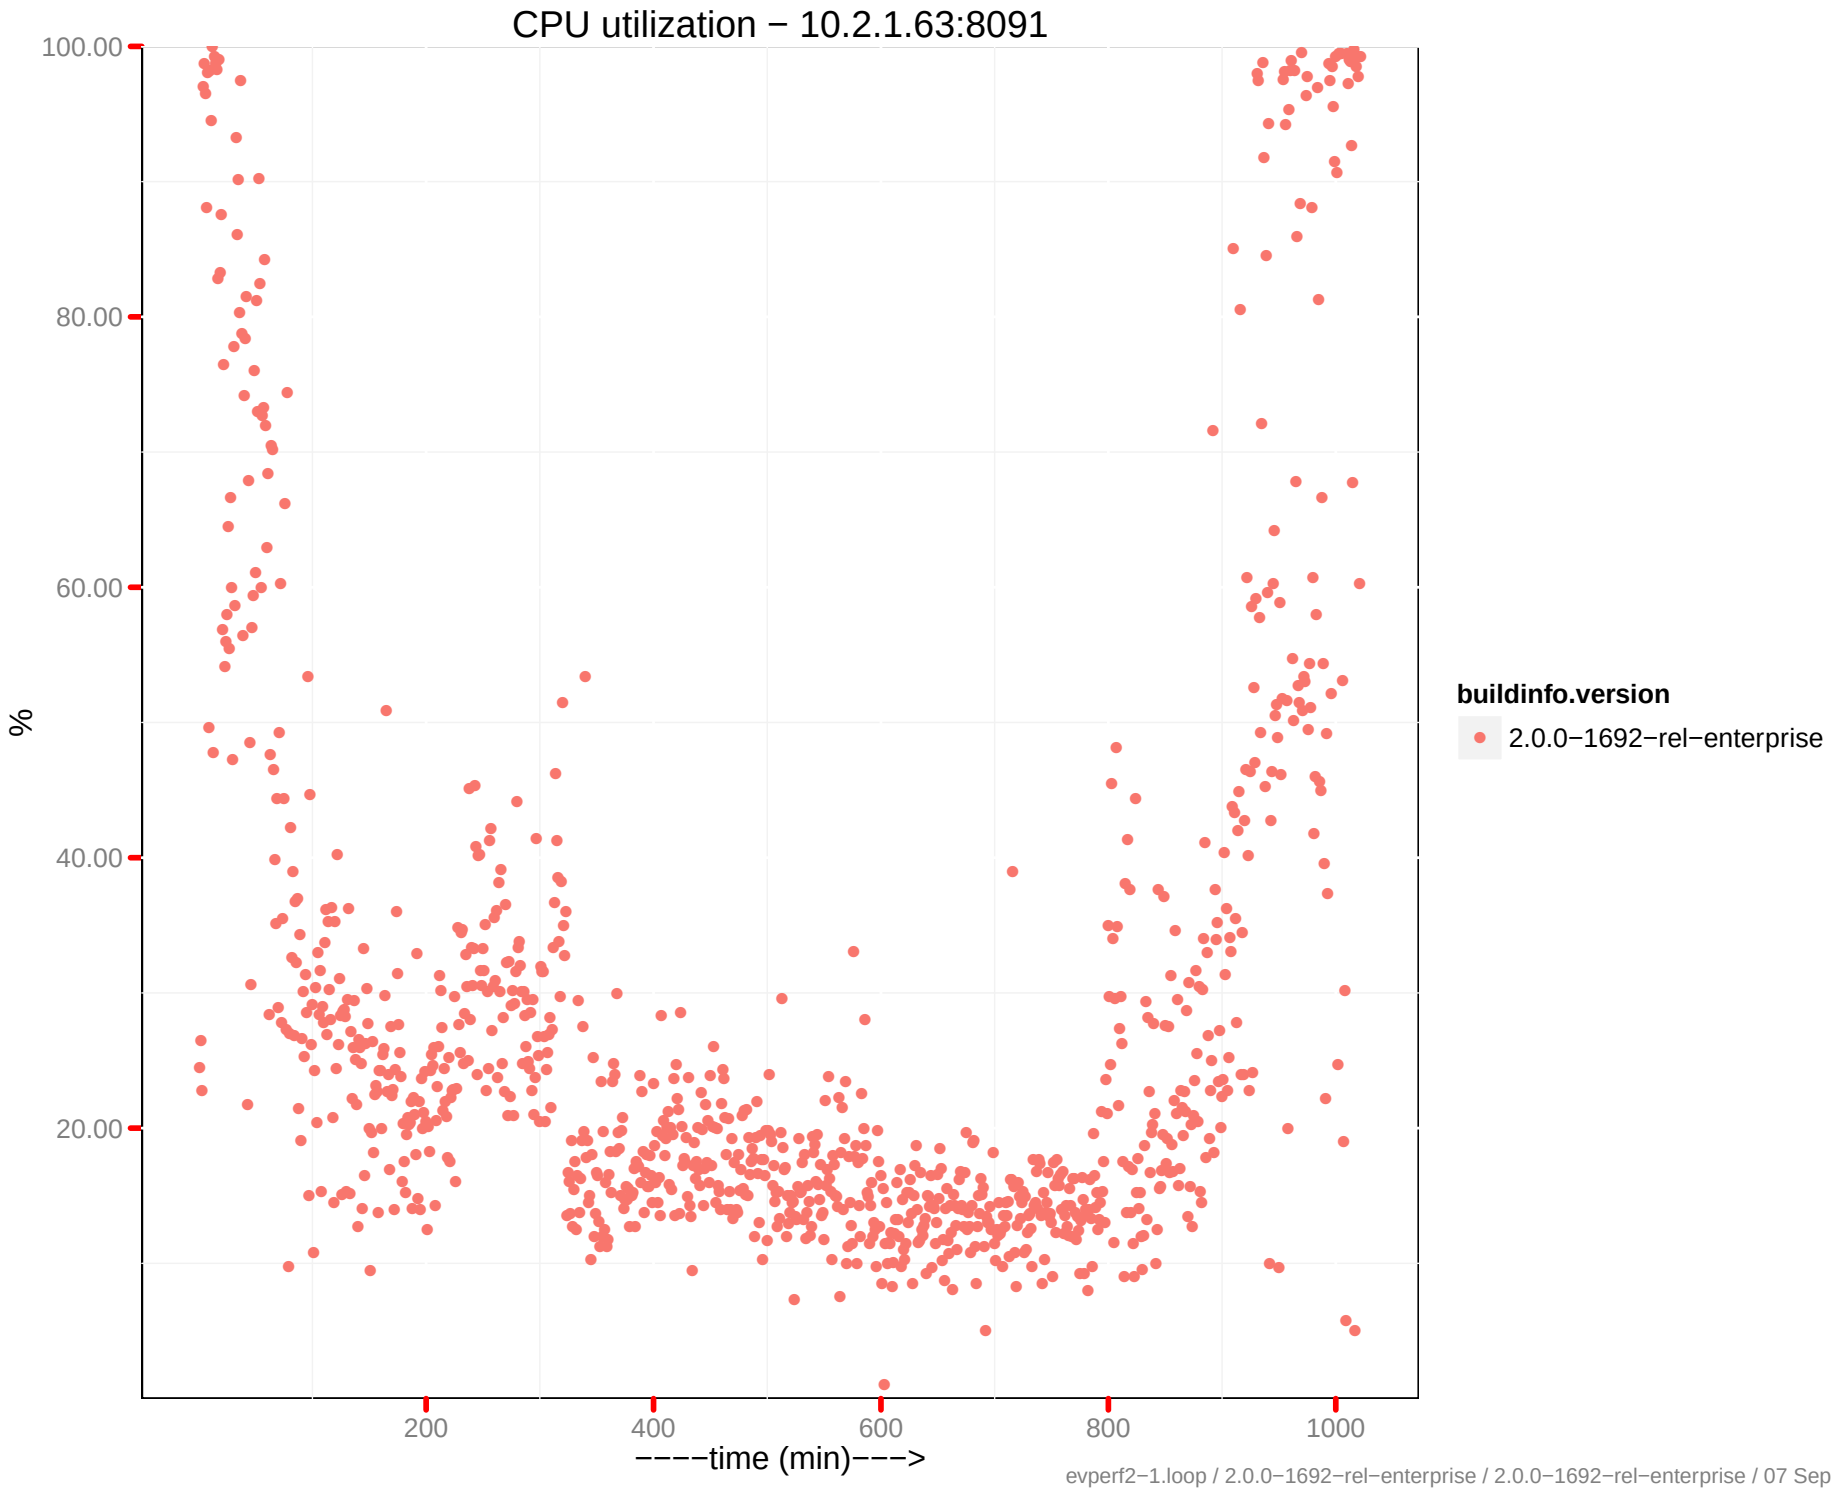
CPU utilization − 10.2.1.63:8091
100.00
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

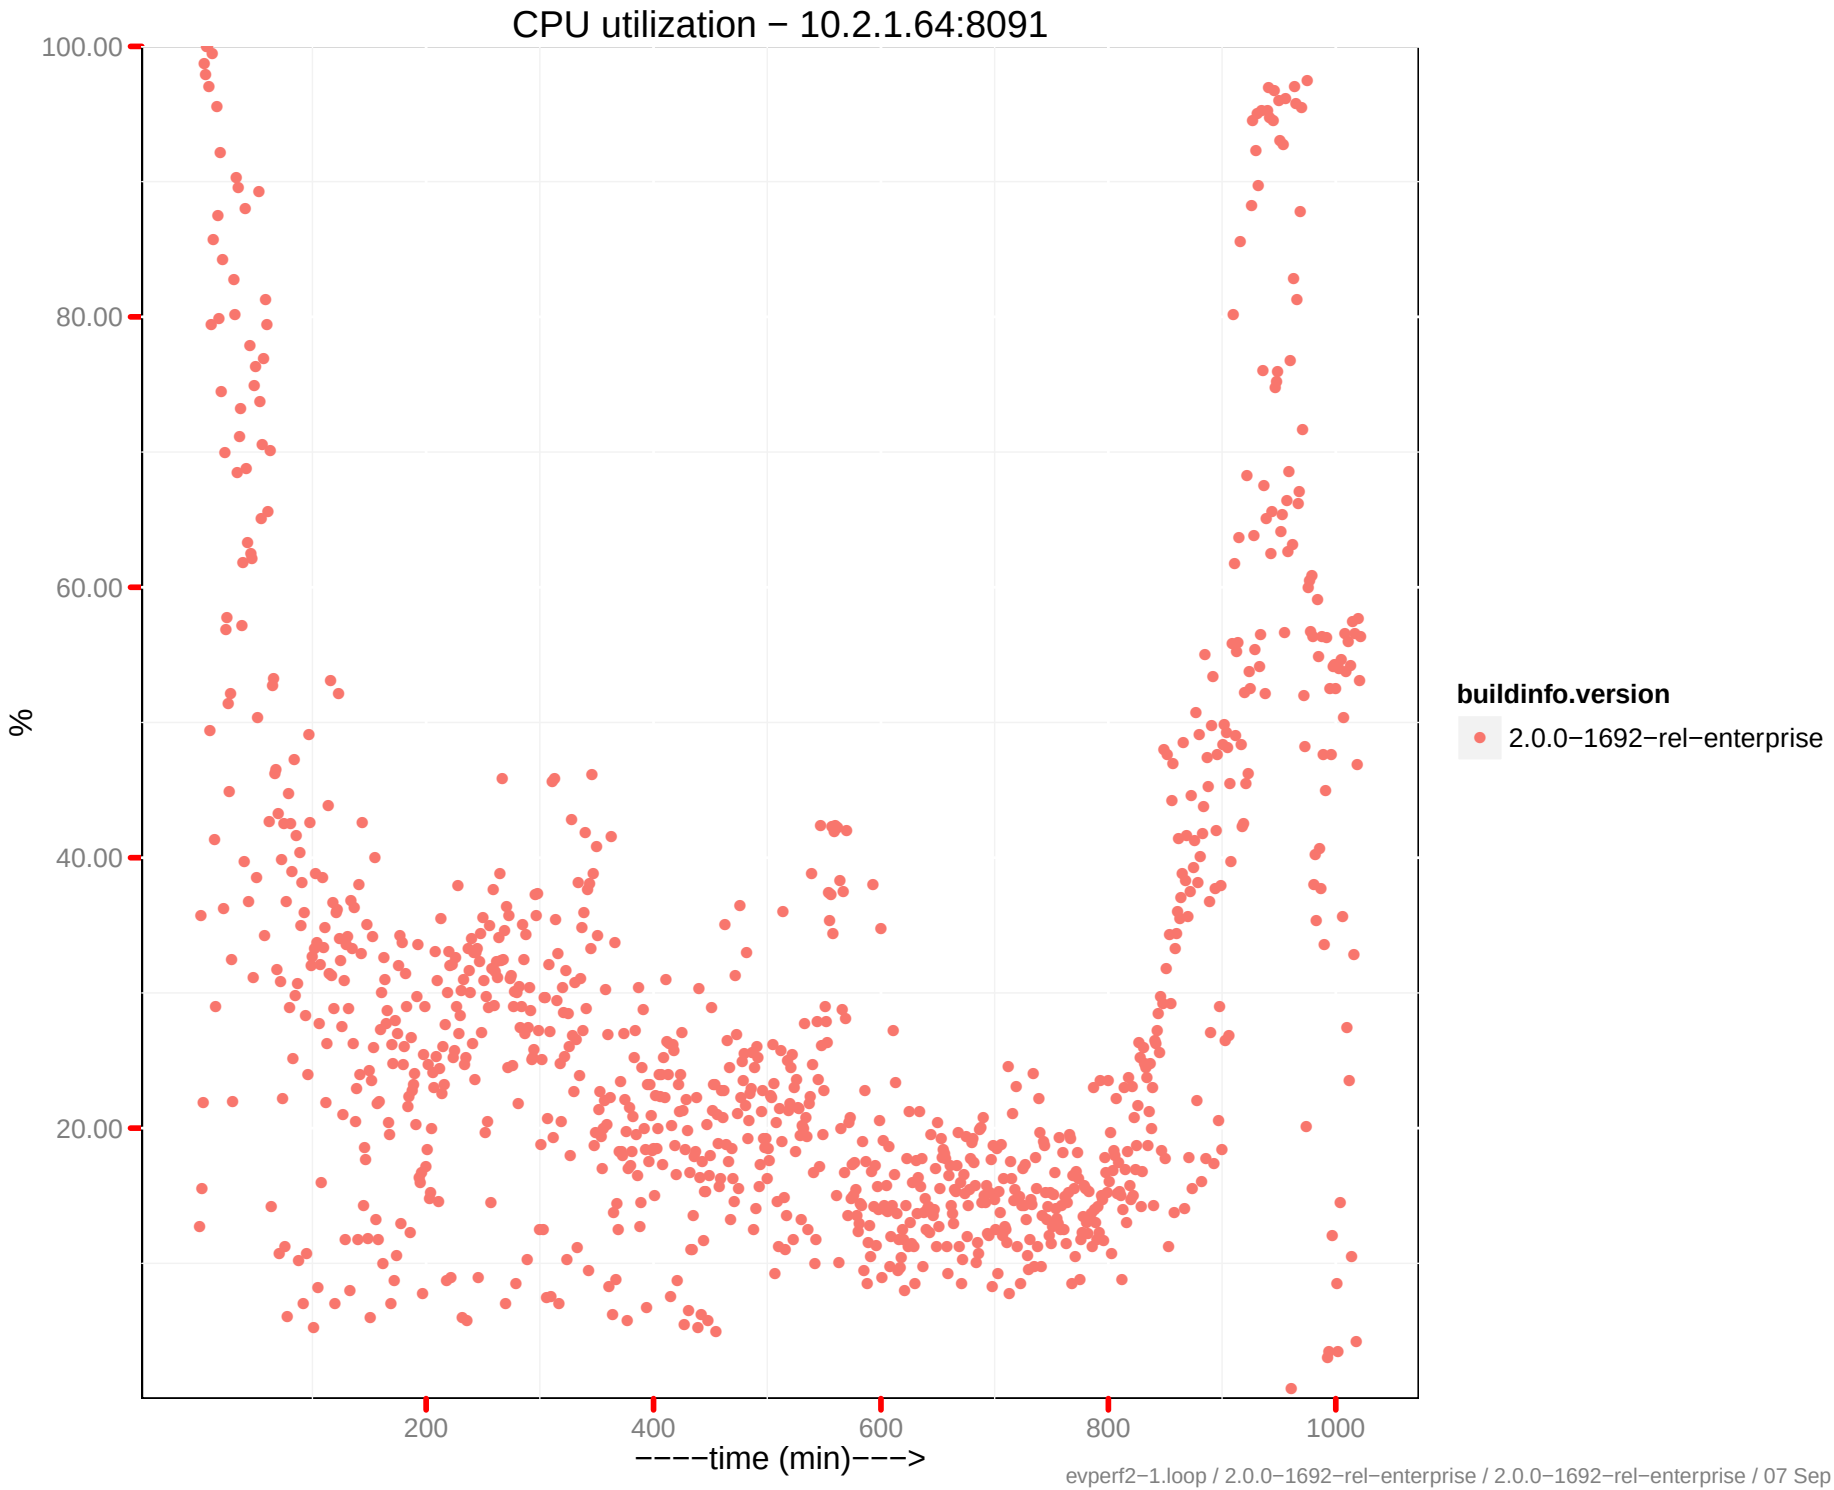
CPU utilization − 10.2.1.64:8091
100.00
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

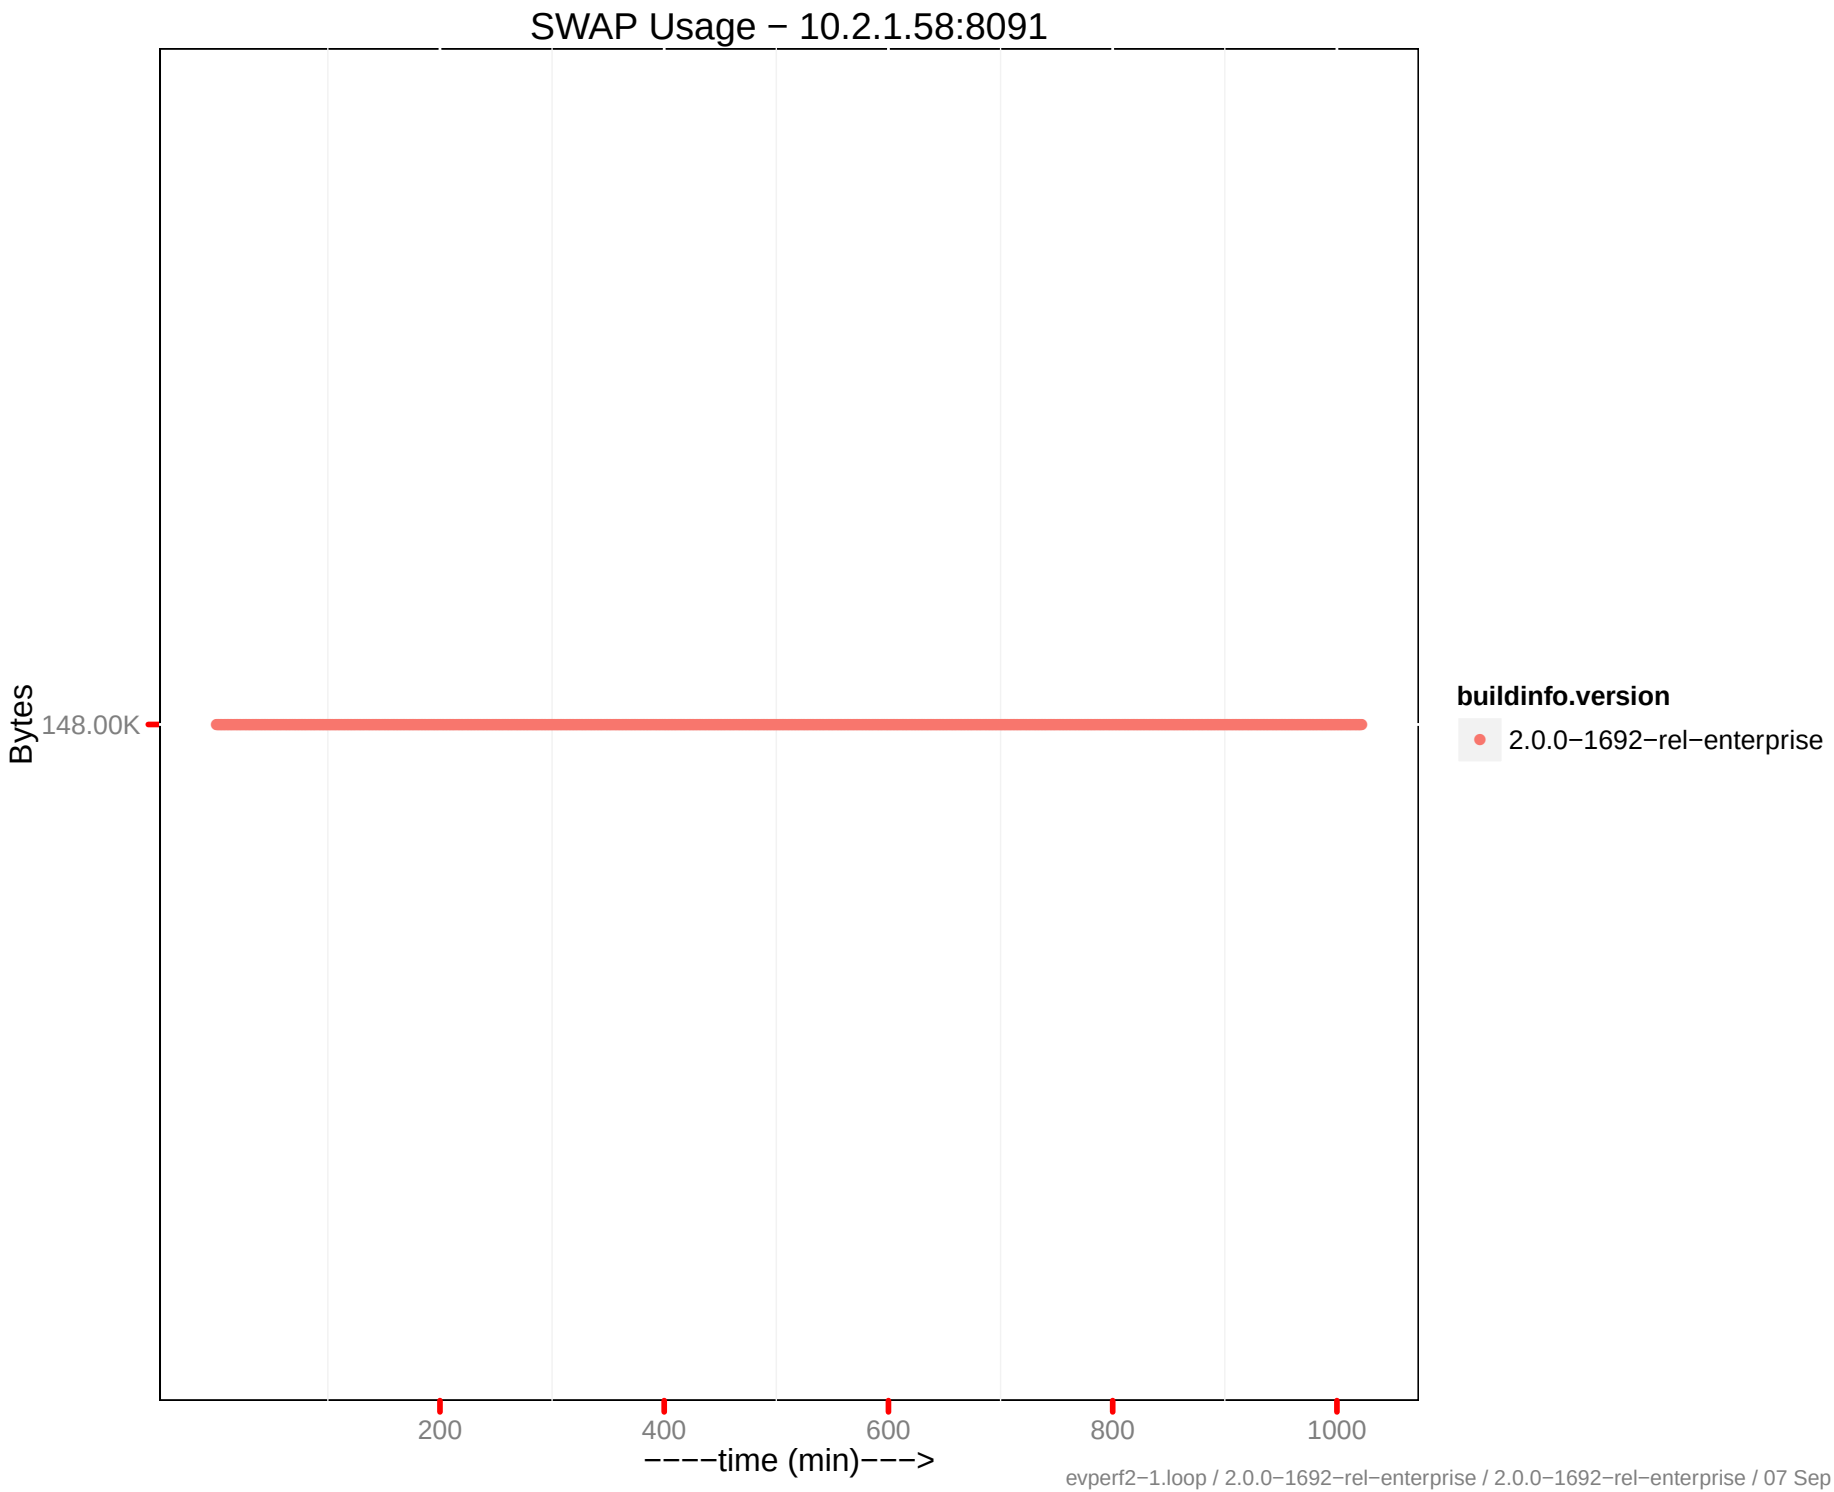
SWAP Usage − 10.2.1.58:8091
Bytes
148.00K
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

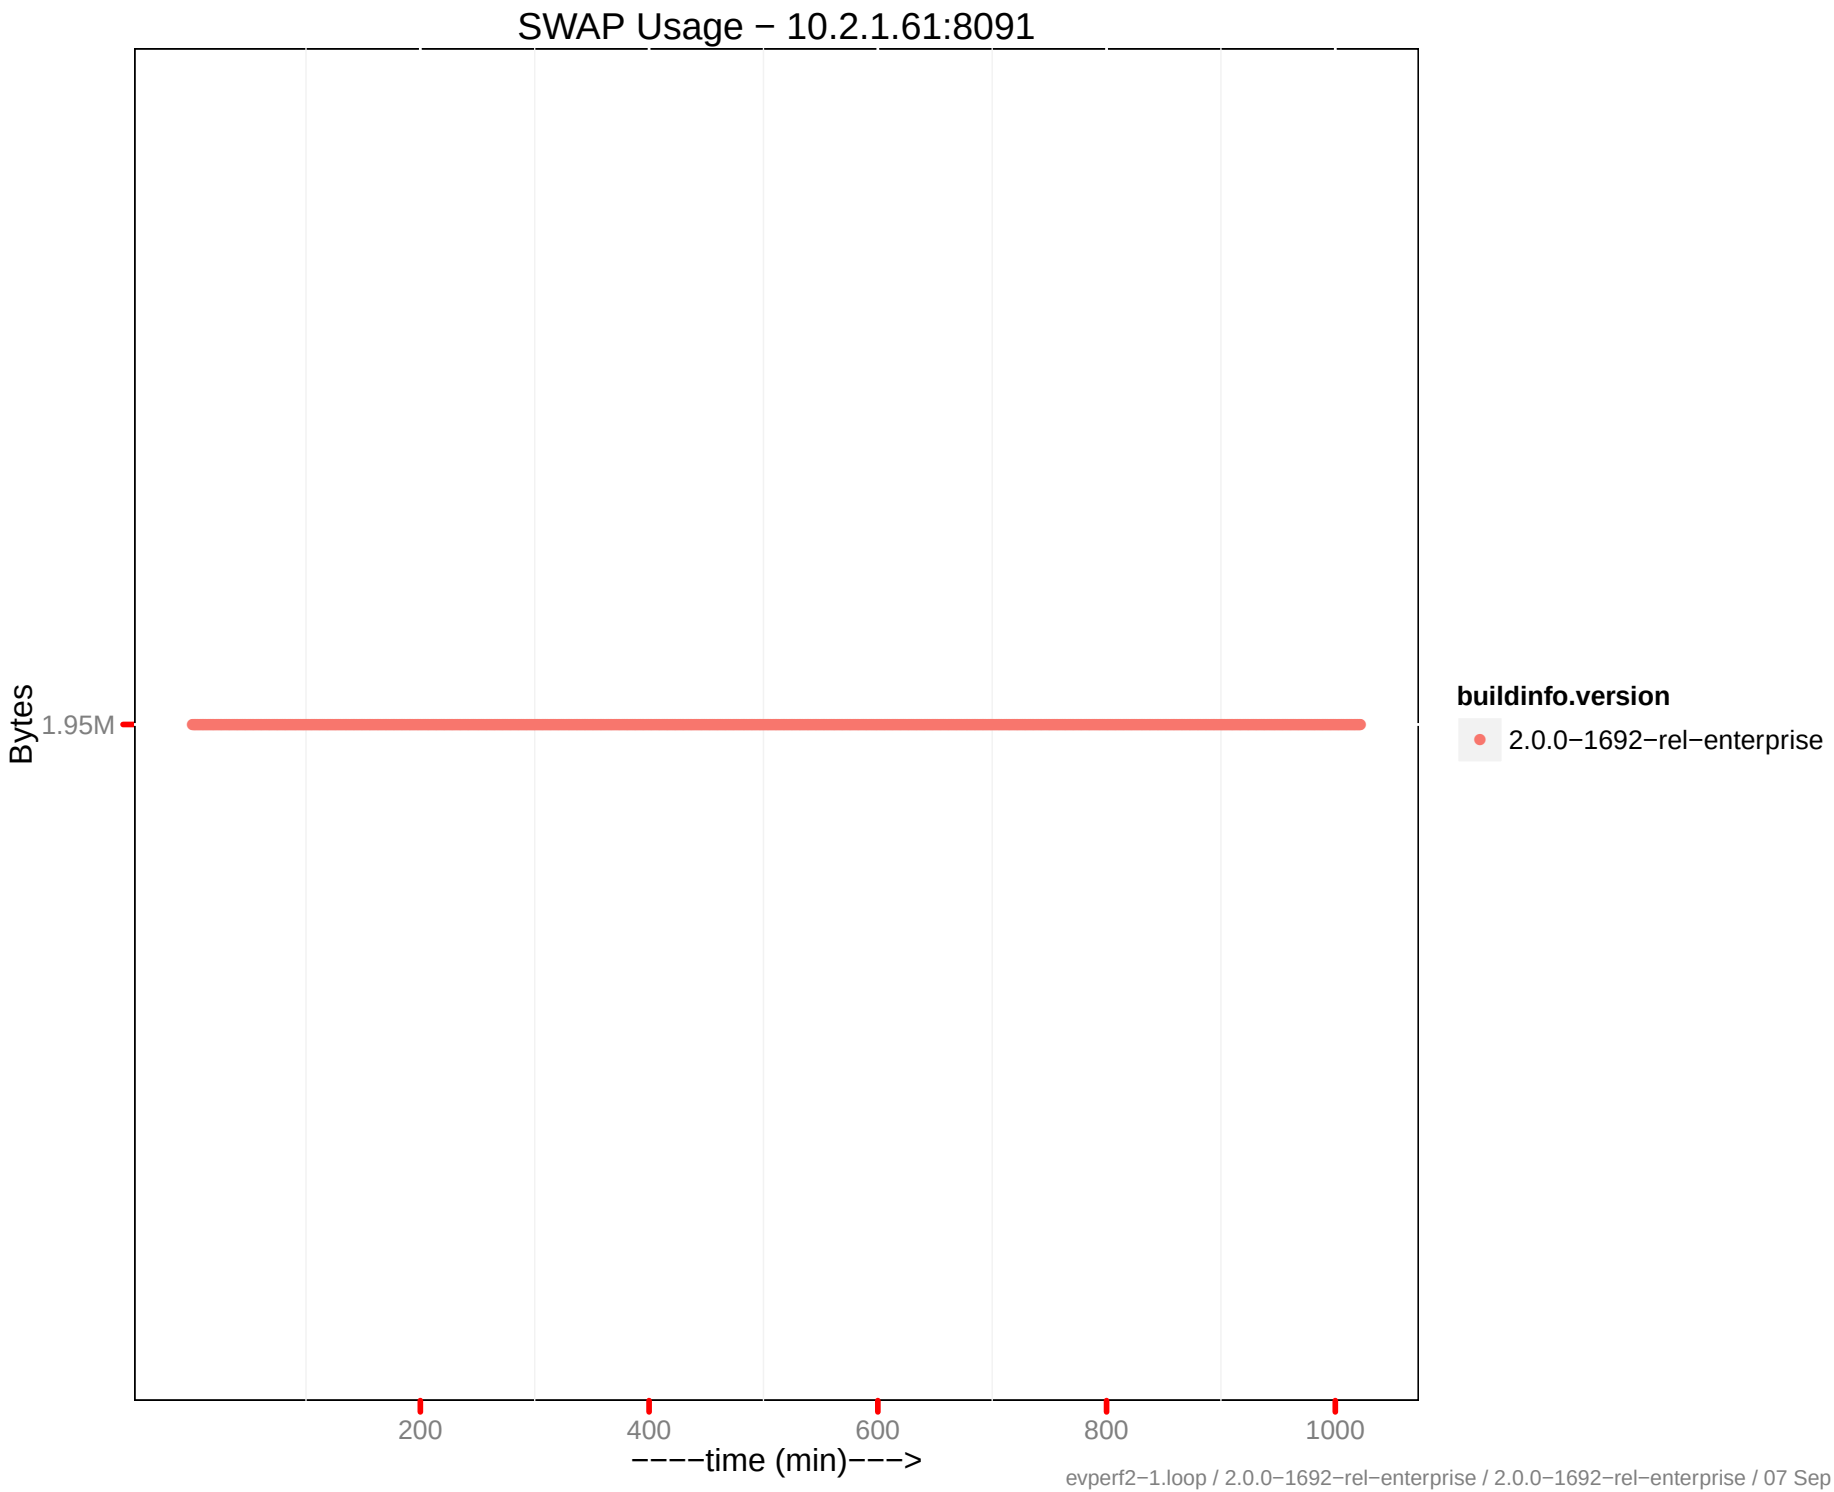
SWAP Usage − 10.2.1.61:8091
Bytes
1.95M
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

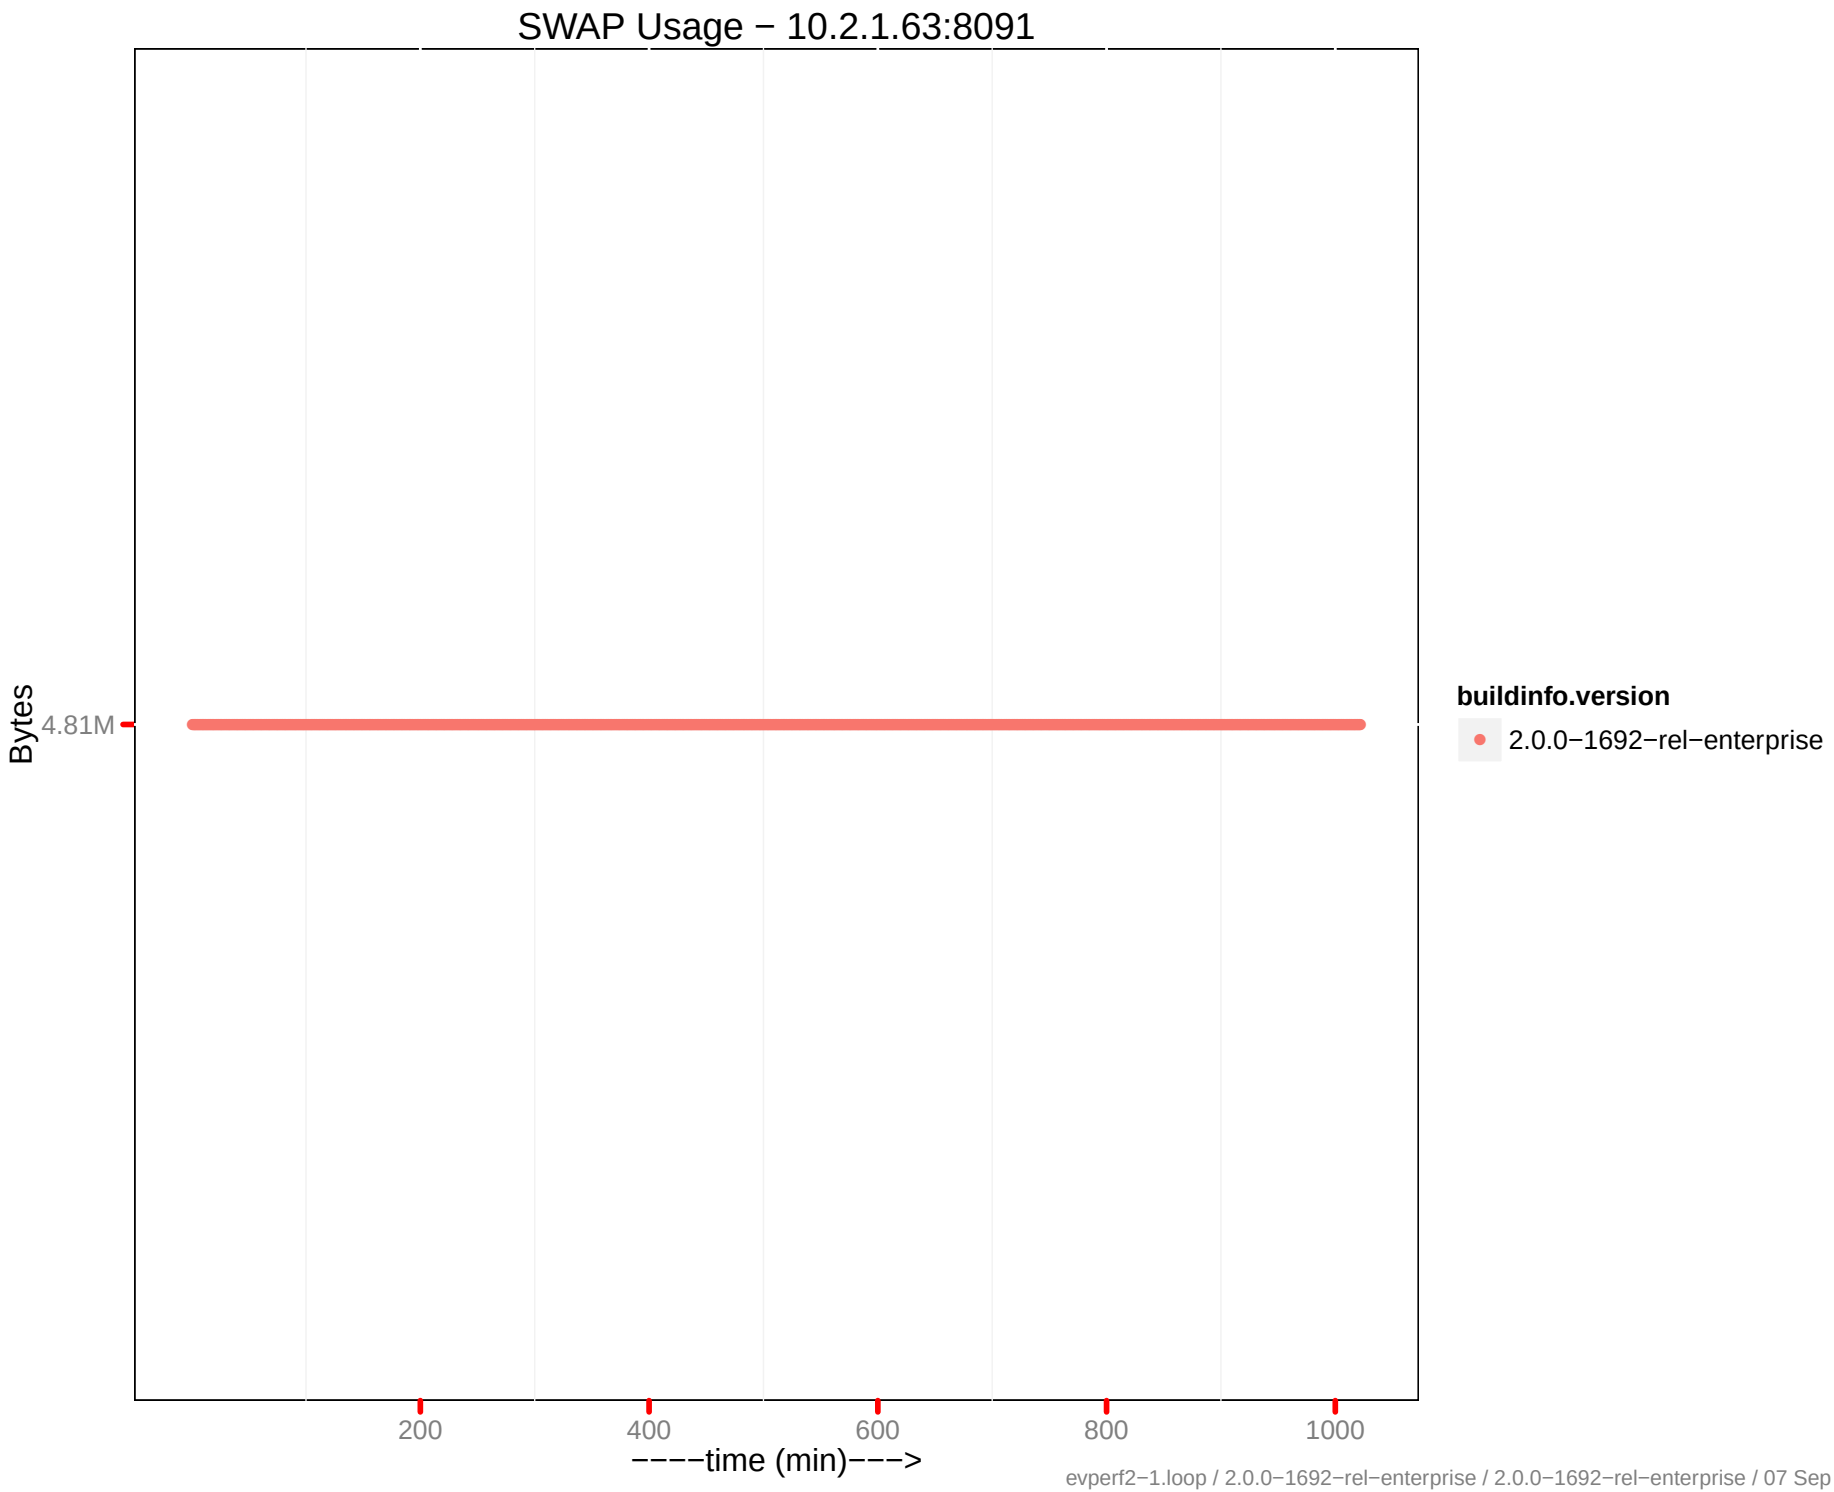
SWAP Usage − 10.2.1.63:8091
Bytes
4.81M
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

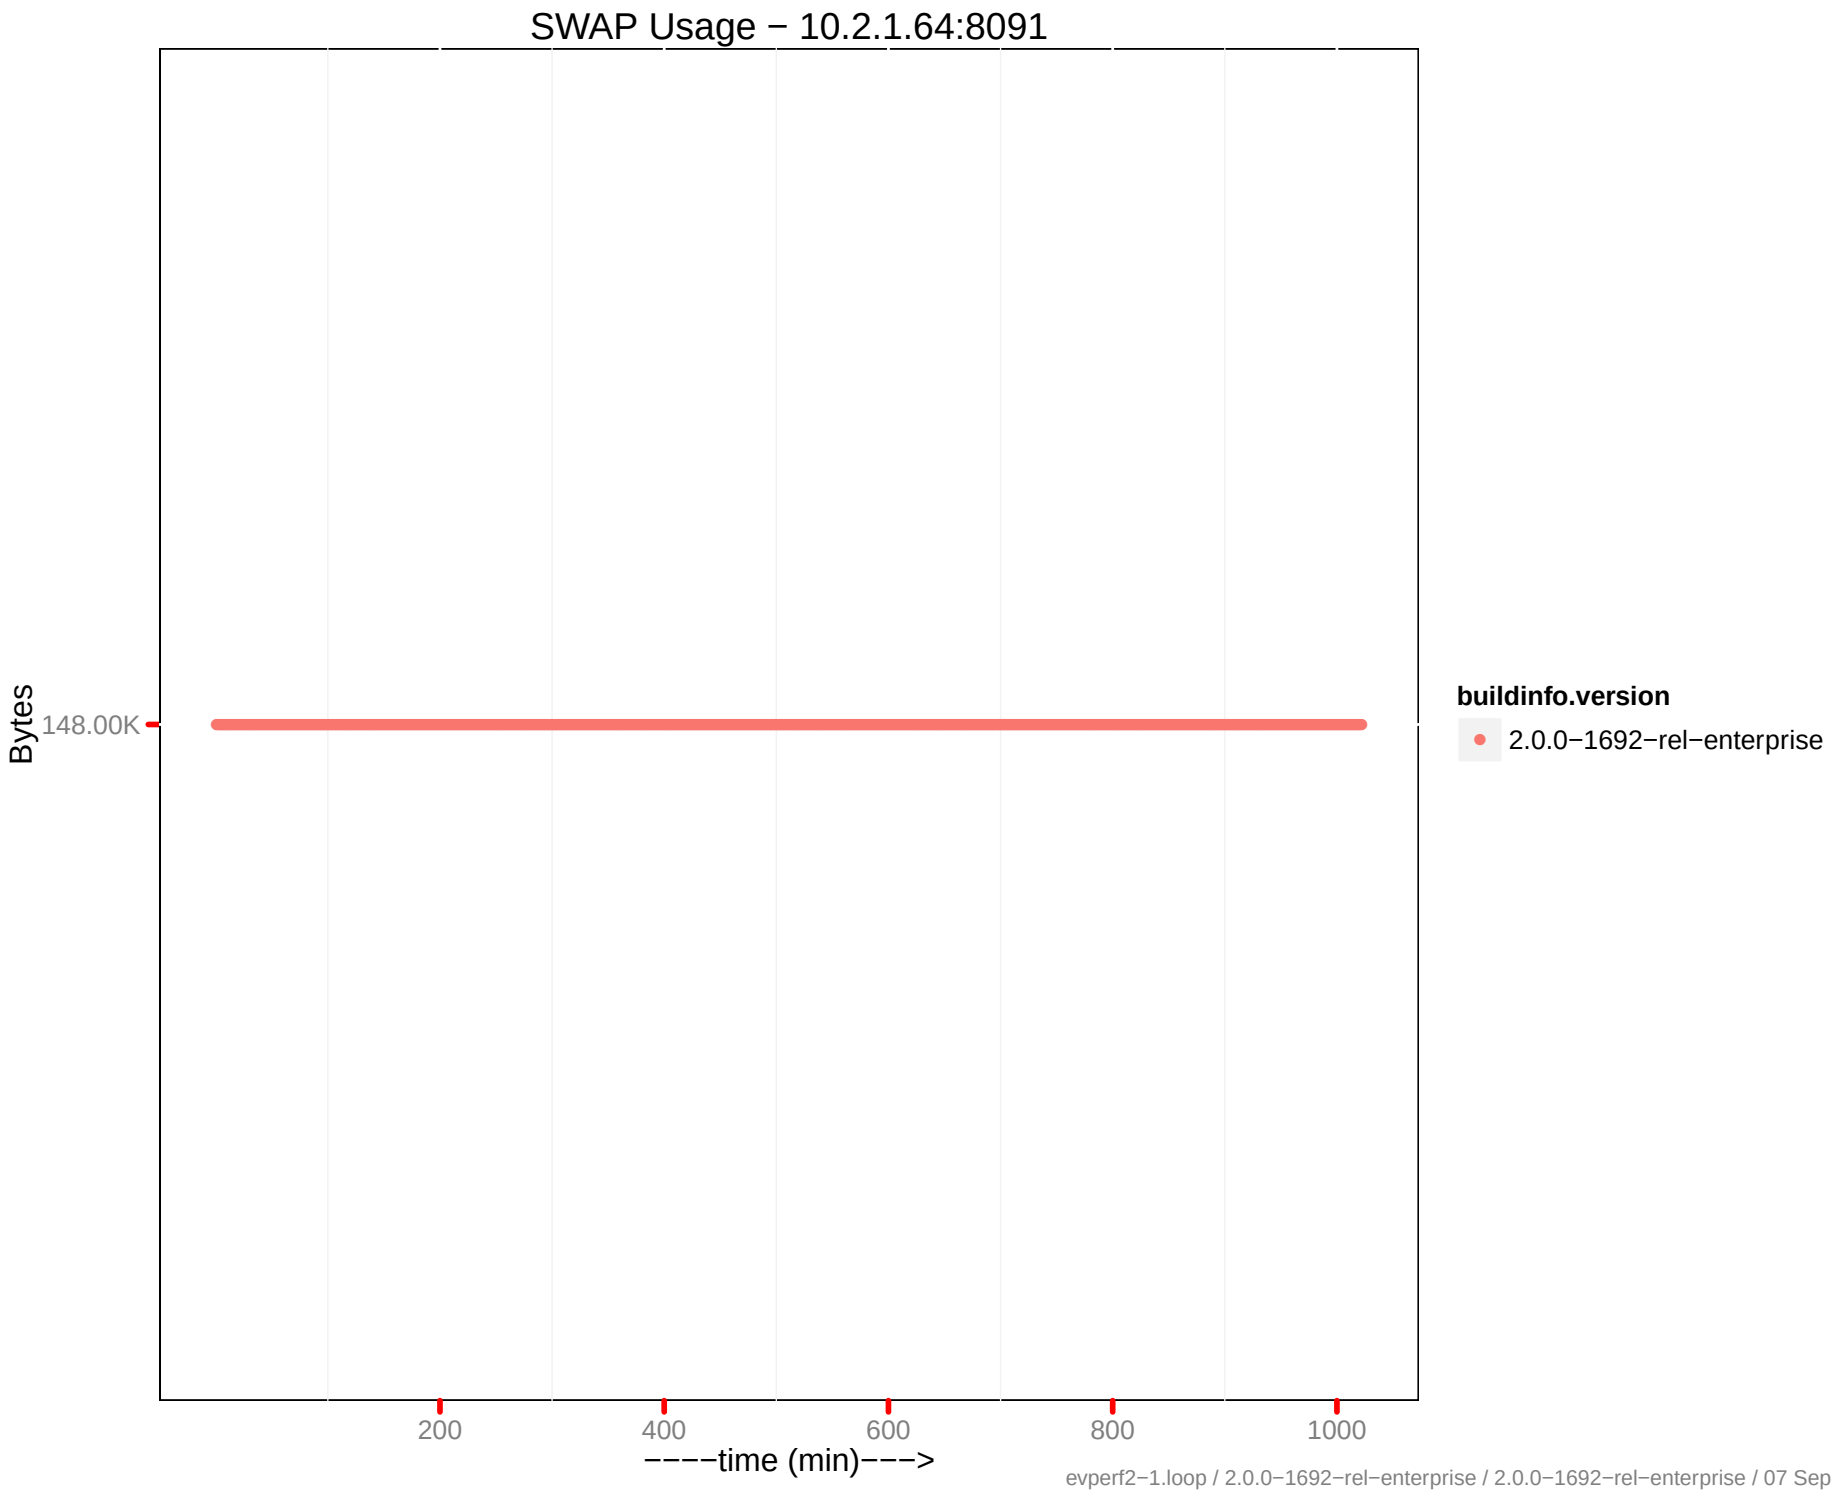
SWAP Usage − 10.2.1.64:8091
Bytes
148.00K
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

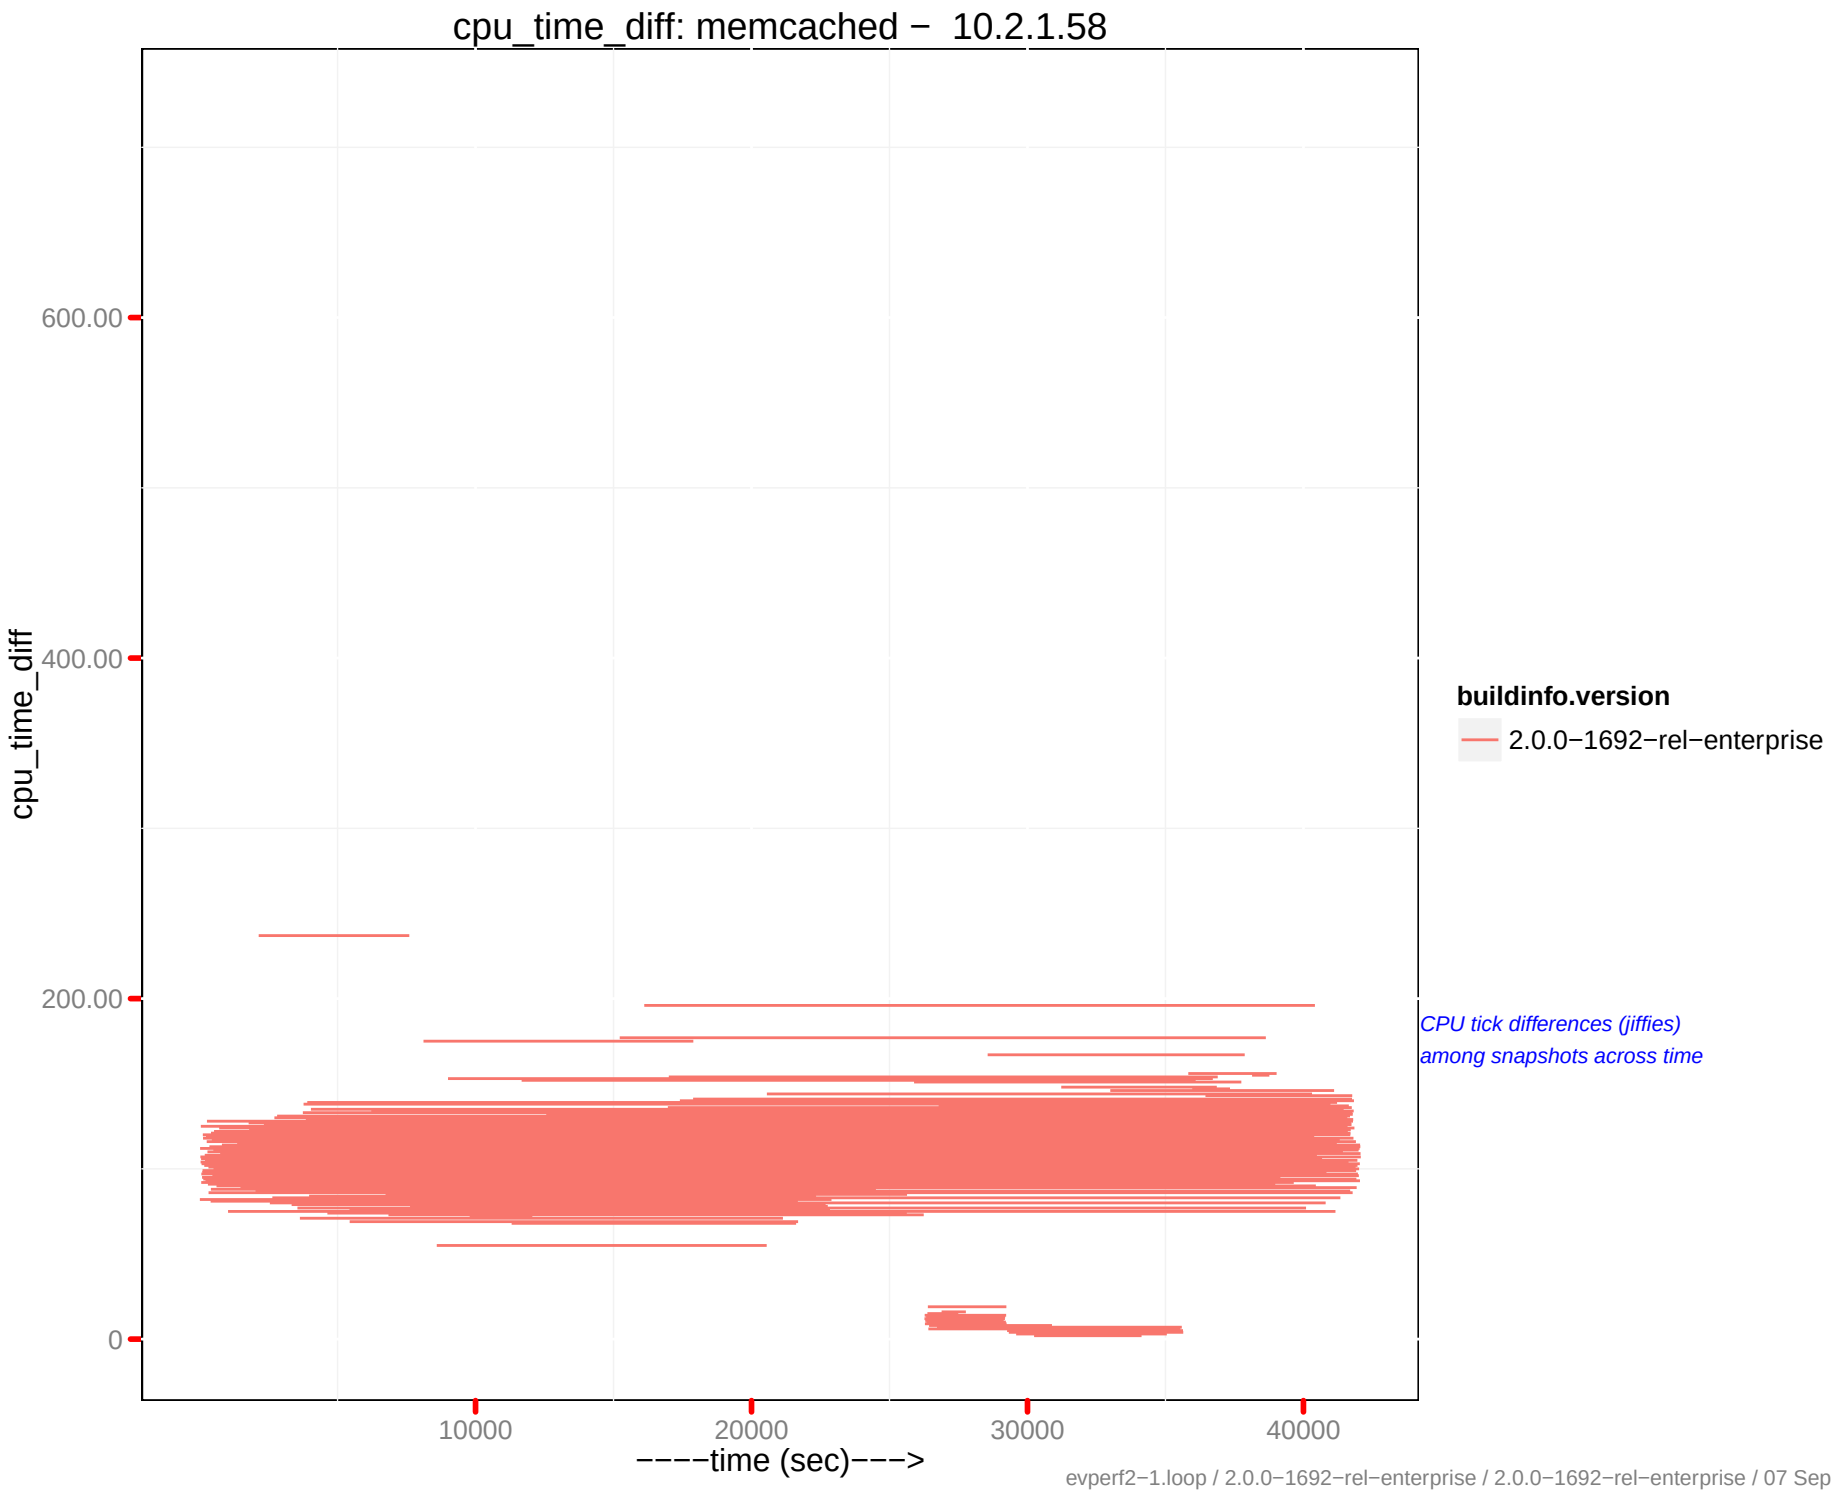
cpu_time_diff: memcached − 10.2.1.58
600.00
400.00
200.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

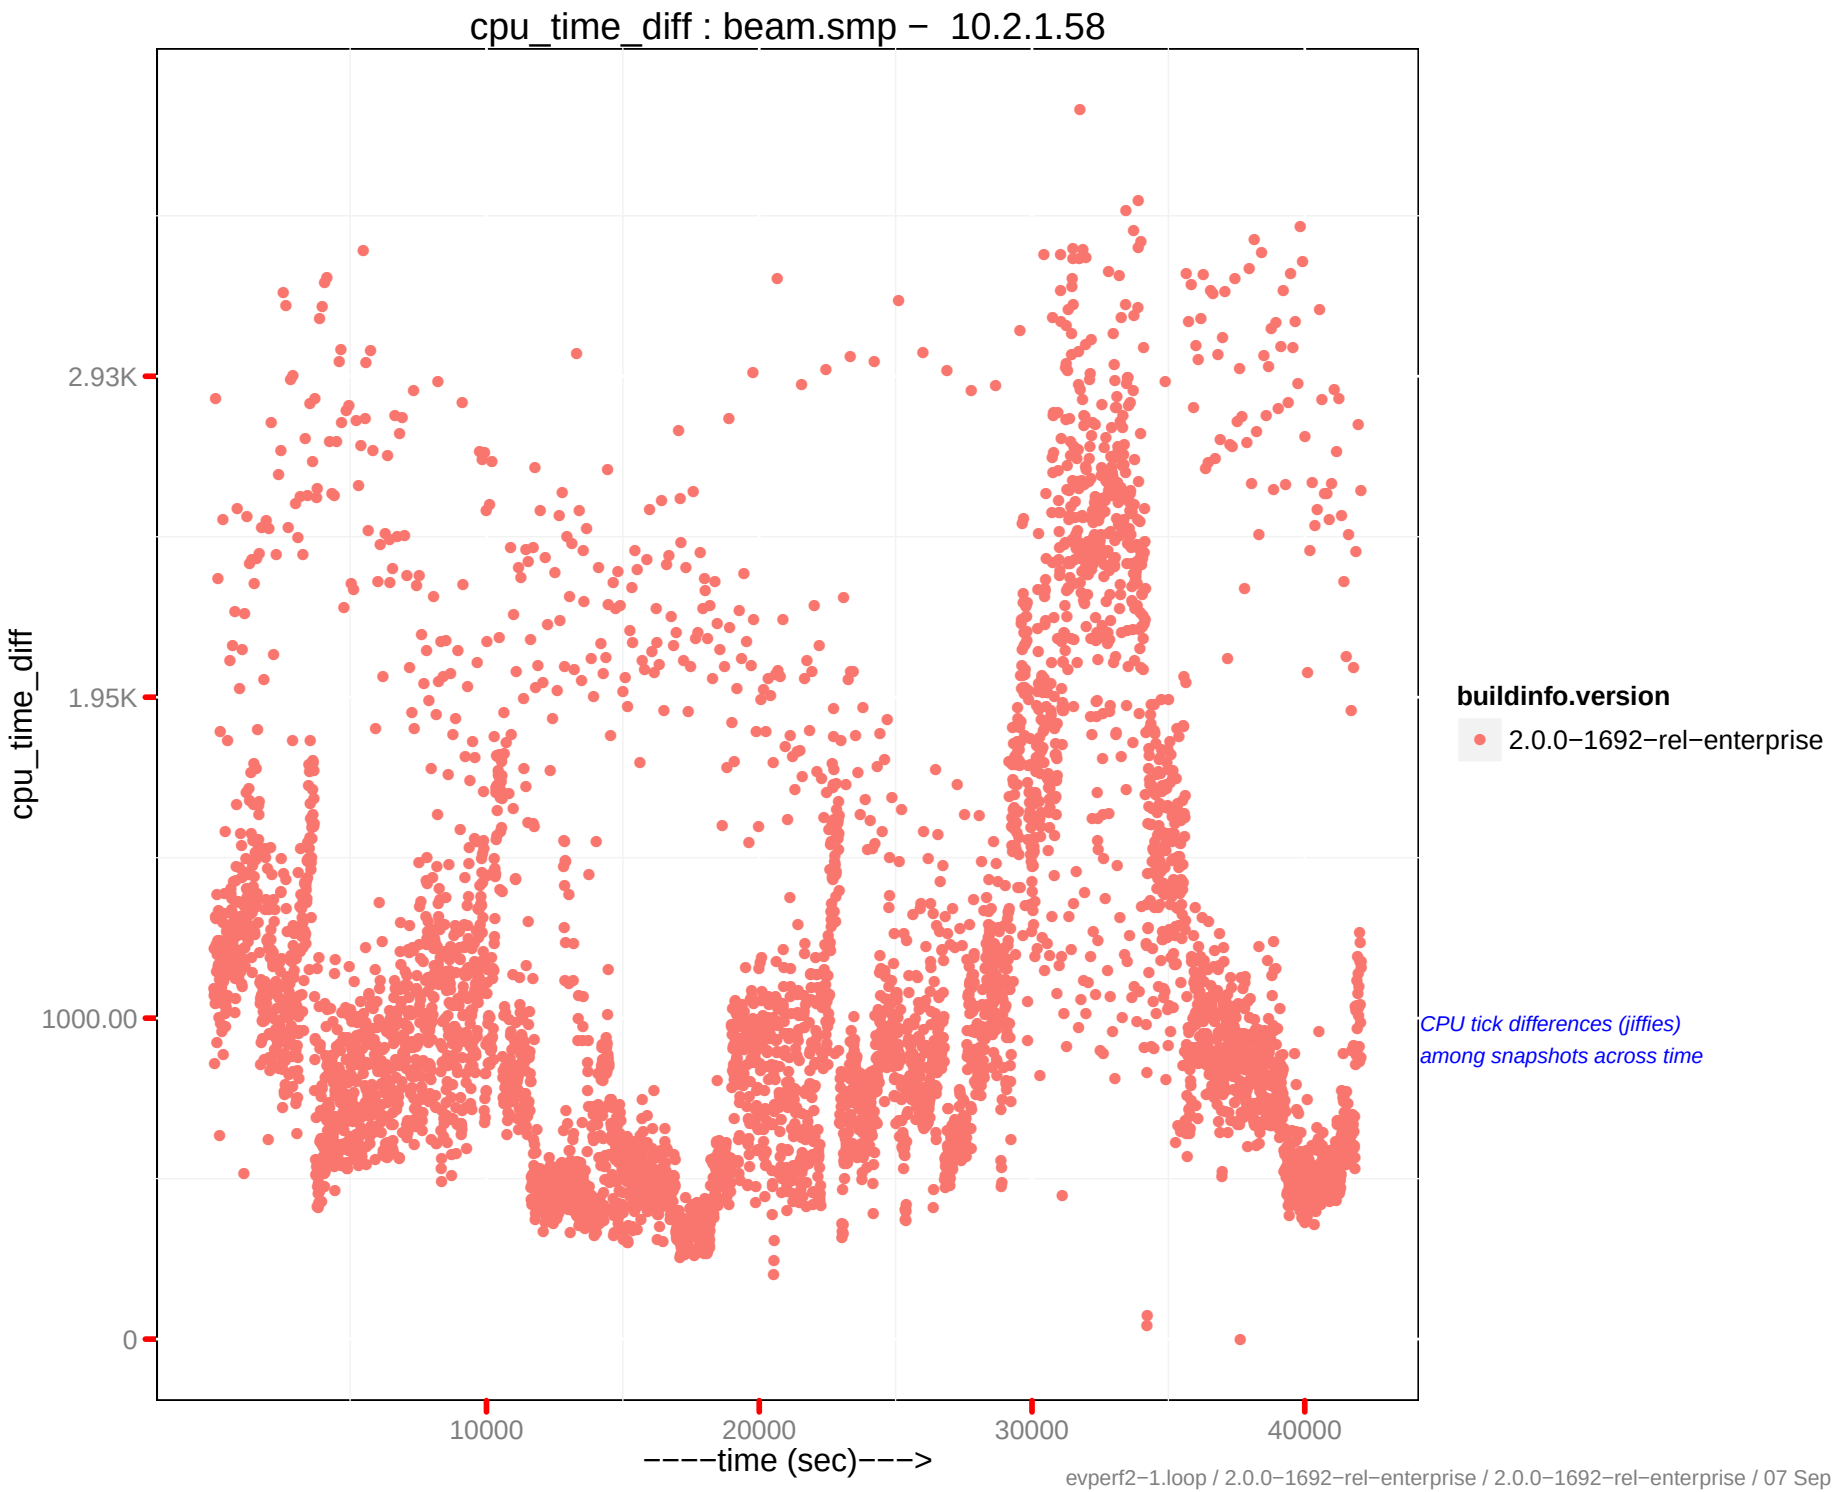
cpu_time_diff : beam.smp − 10.2.1.58
cpu_time_diff
2.93K
1.95K
1000.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

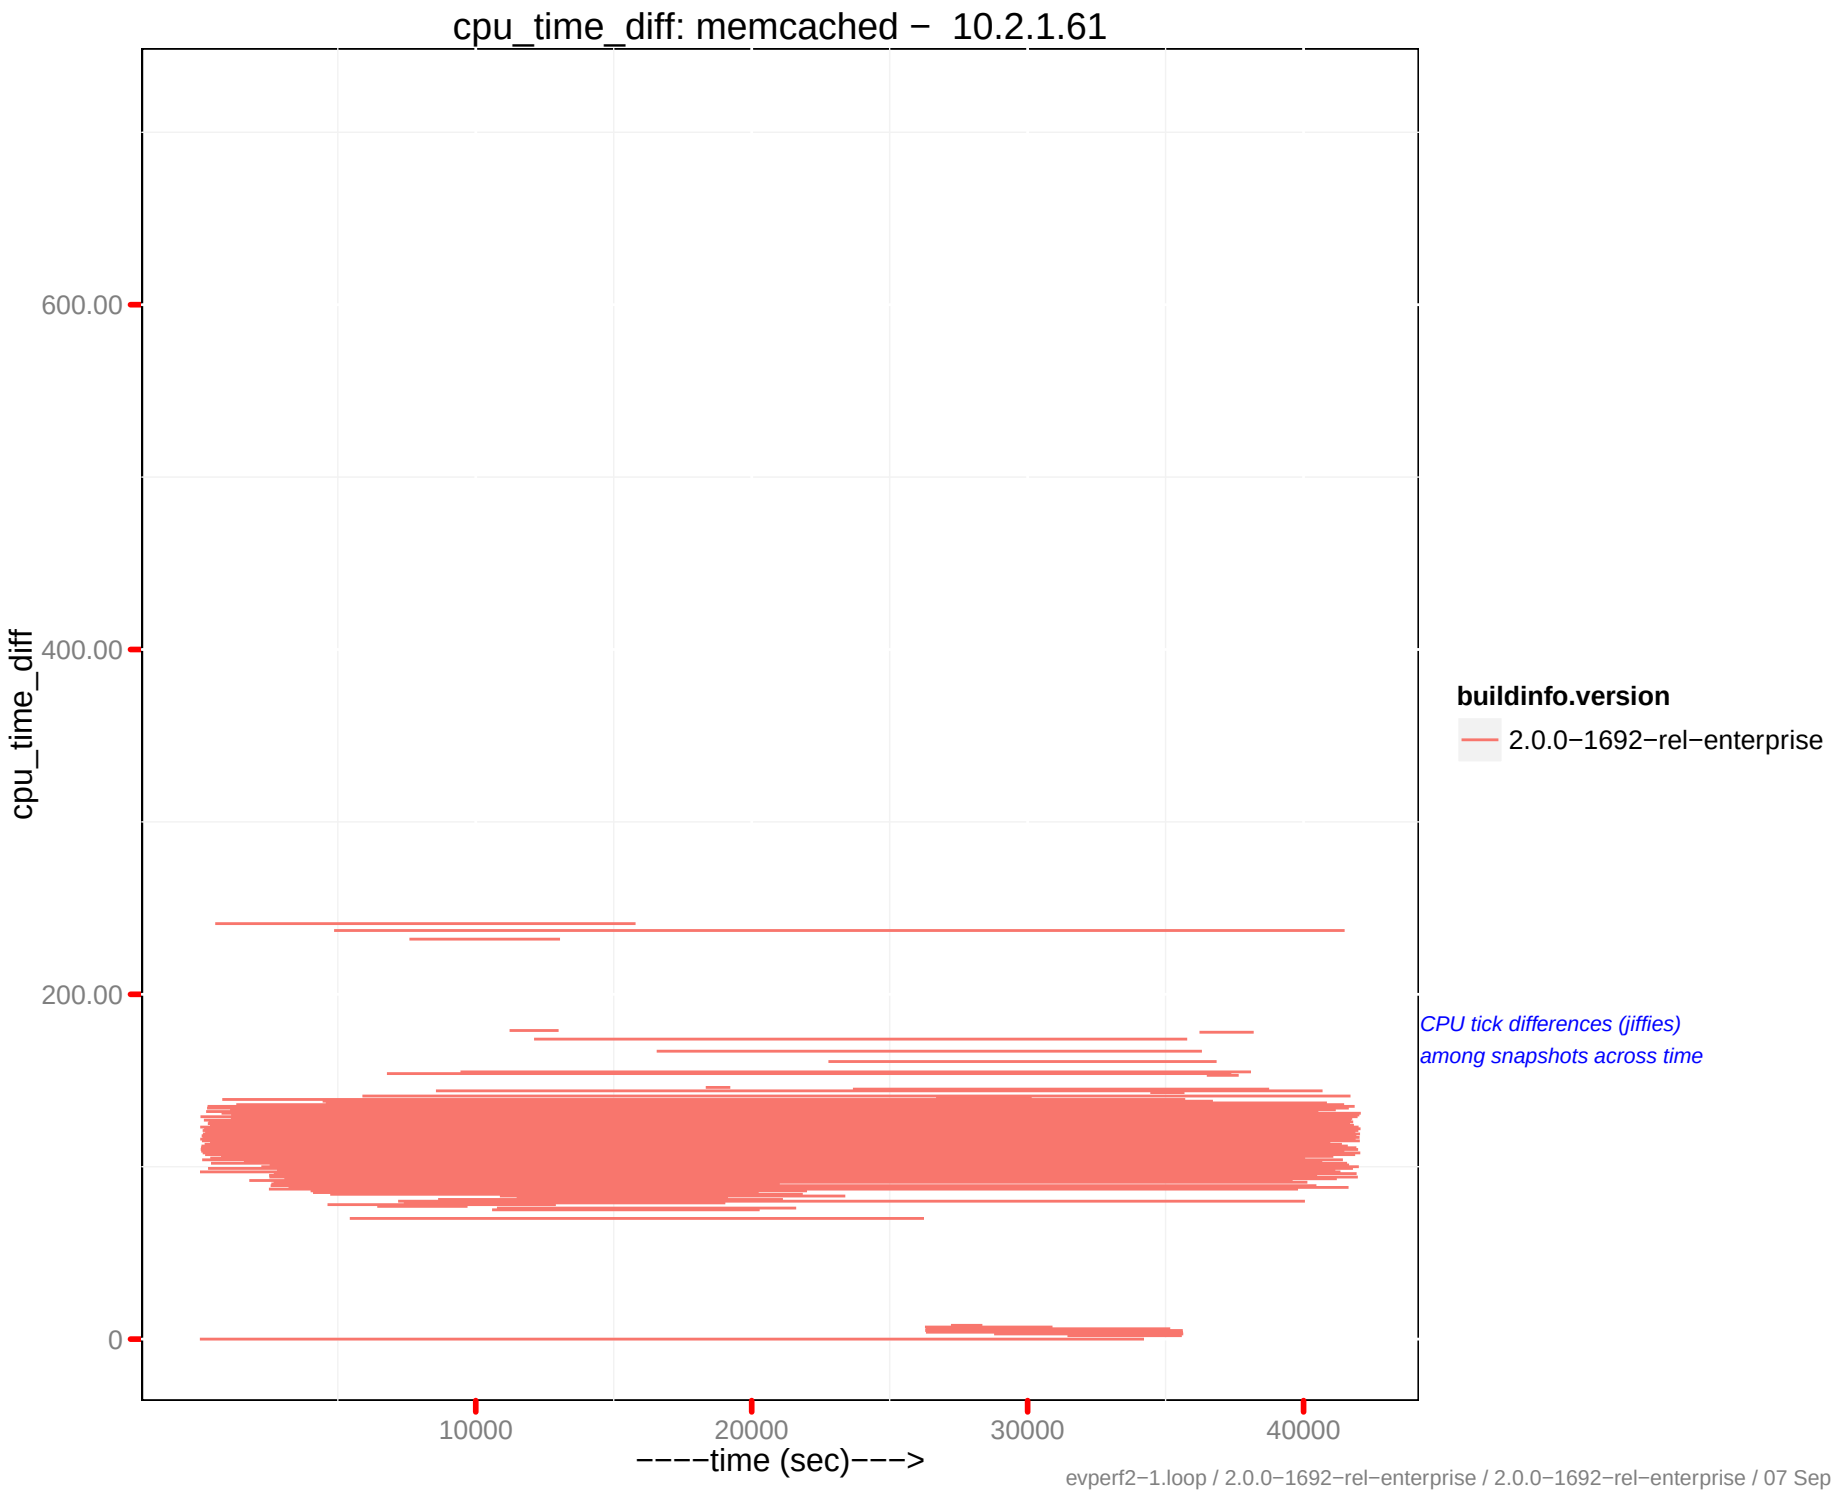
cpu_time_diff: memcached − 10.2.1.61
600.00
400.00
200.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

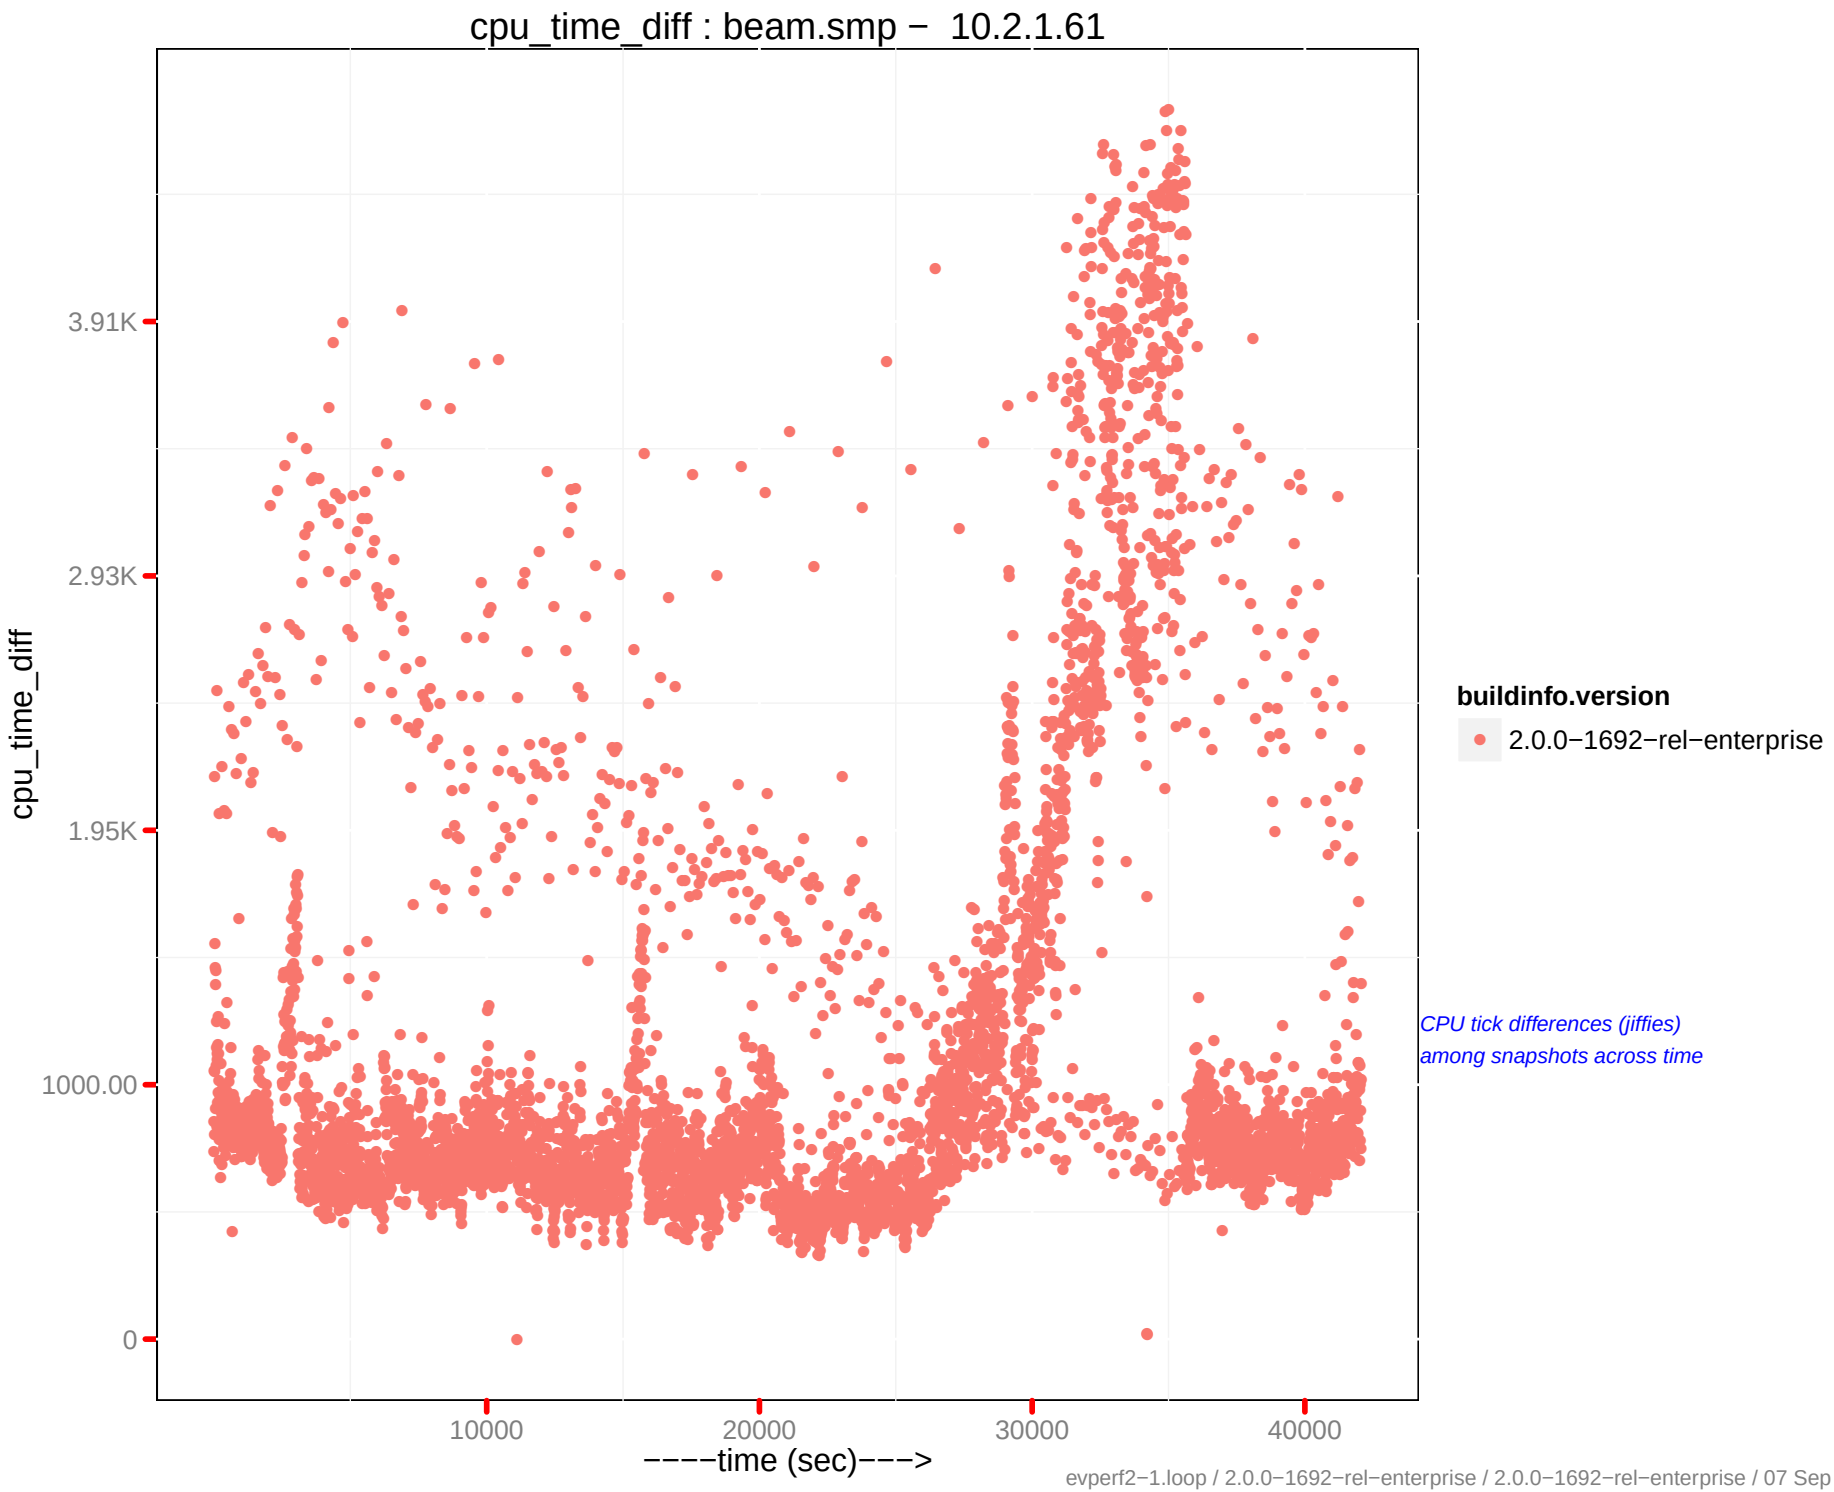
cpu_time_diff : beam.smp − 10.2.1.61
cpu_time_diff
3.91K
2.93K
1.95K
1000.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

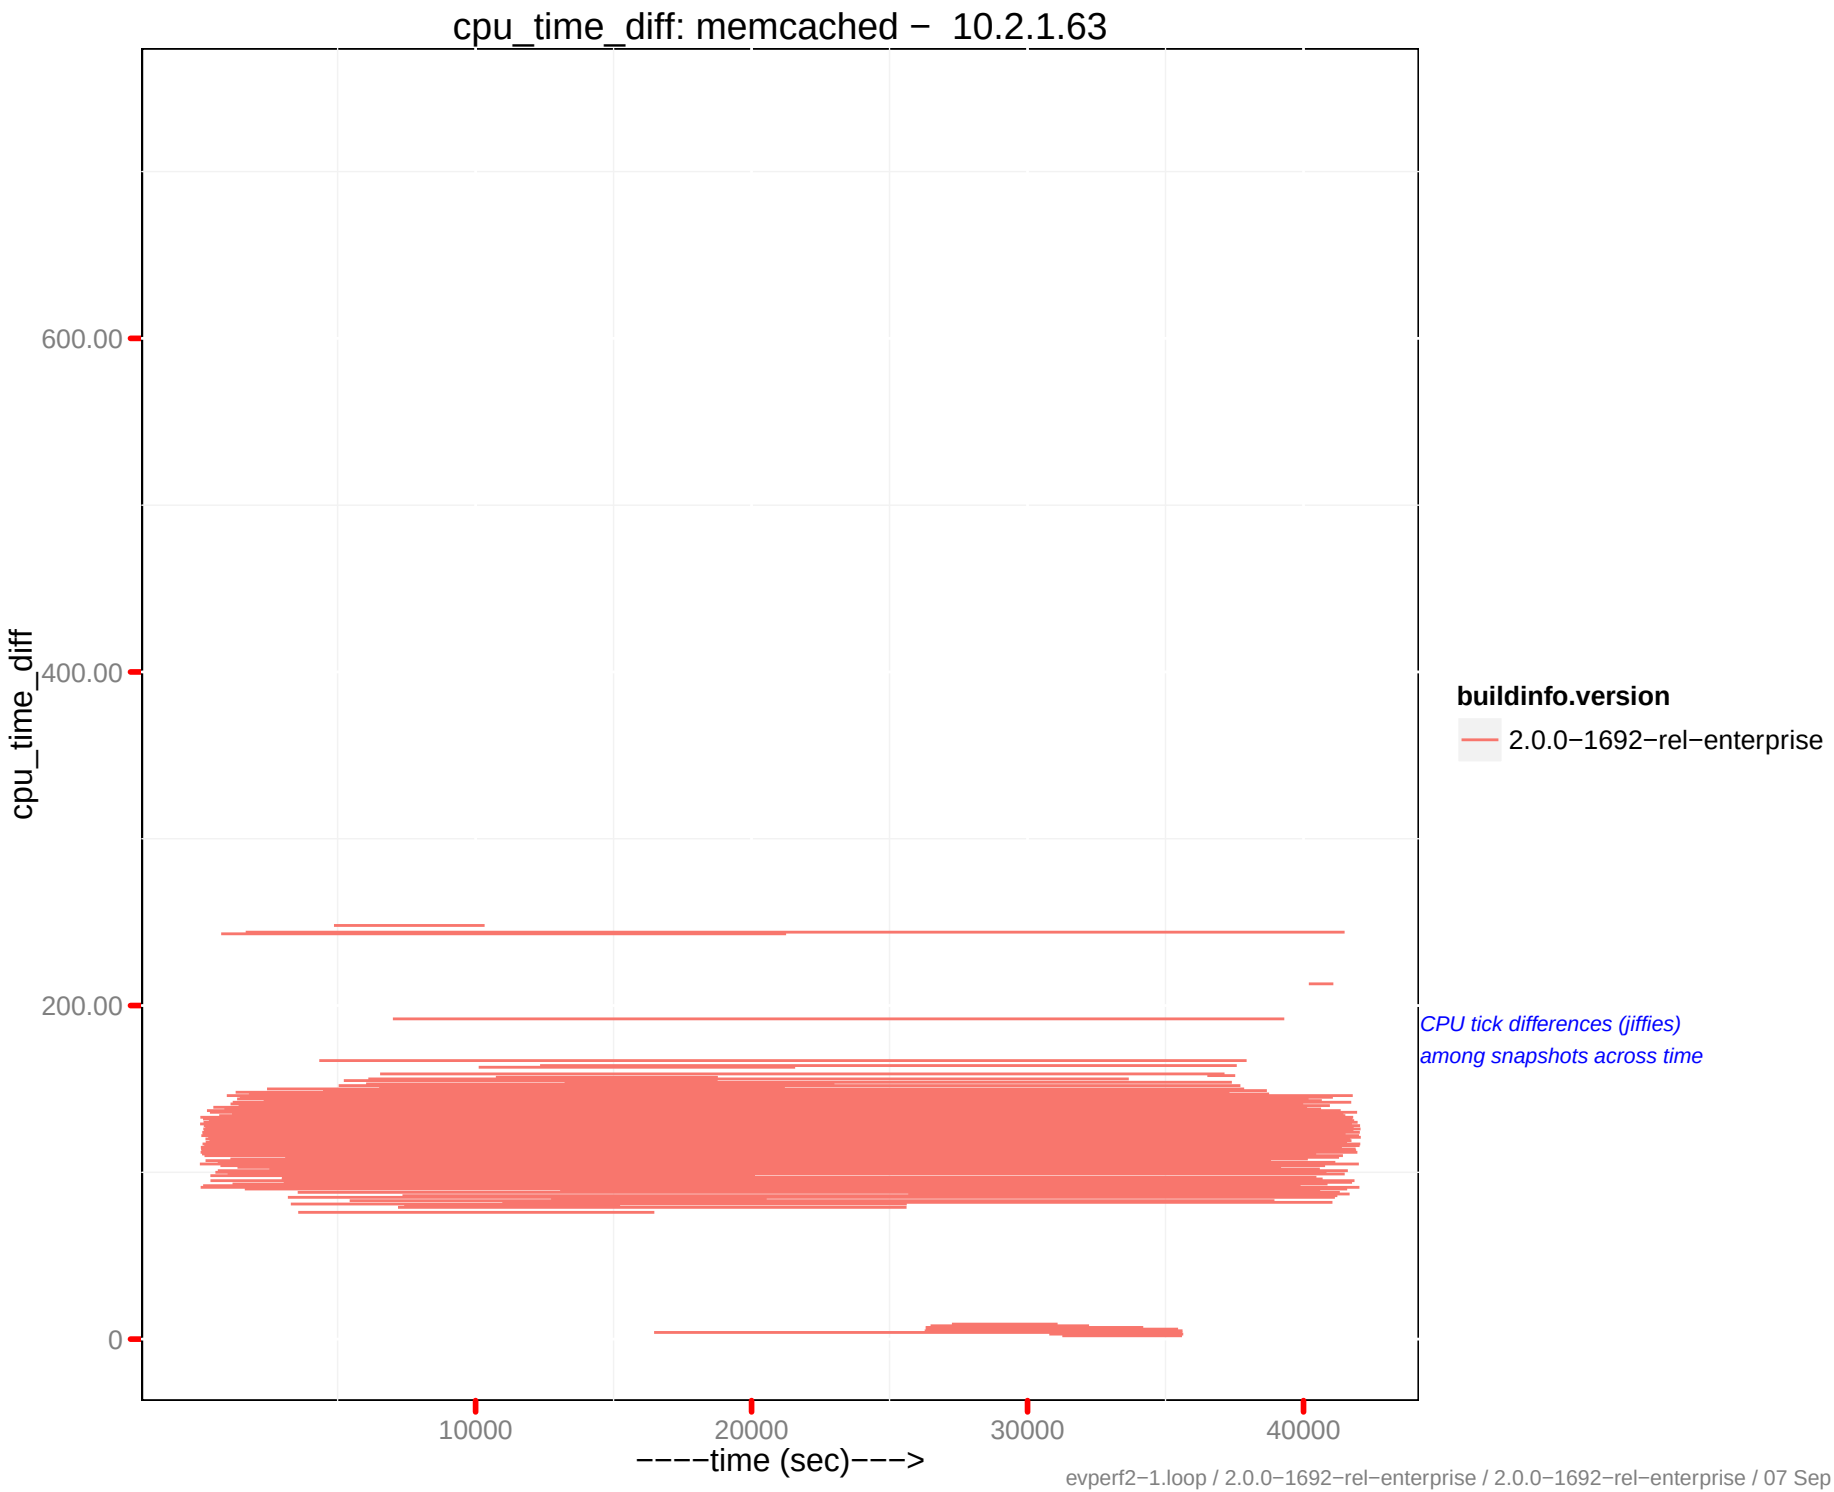
cpu_time_diff: memcached − 10.2.1.63
600.00
400.00
200.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

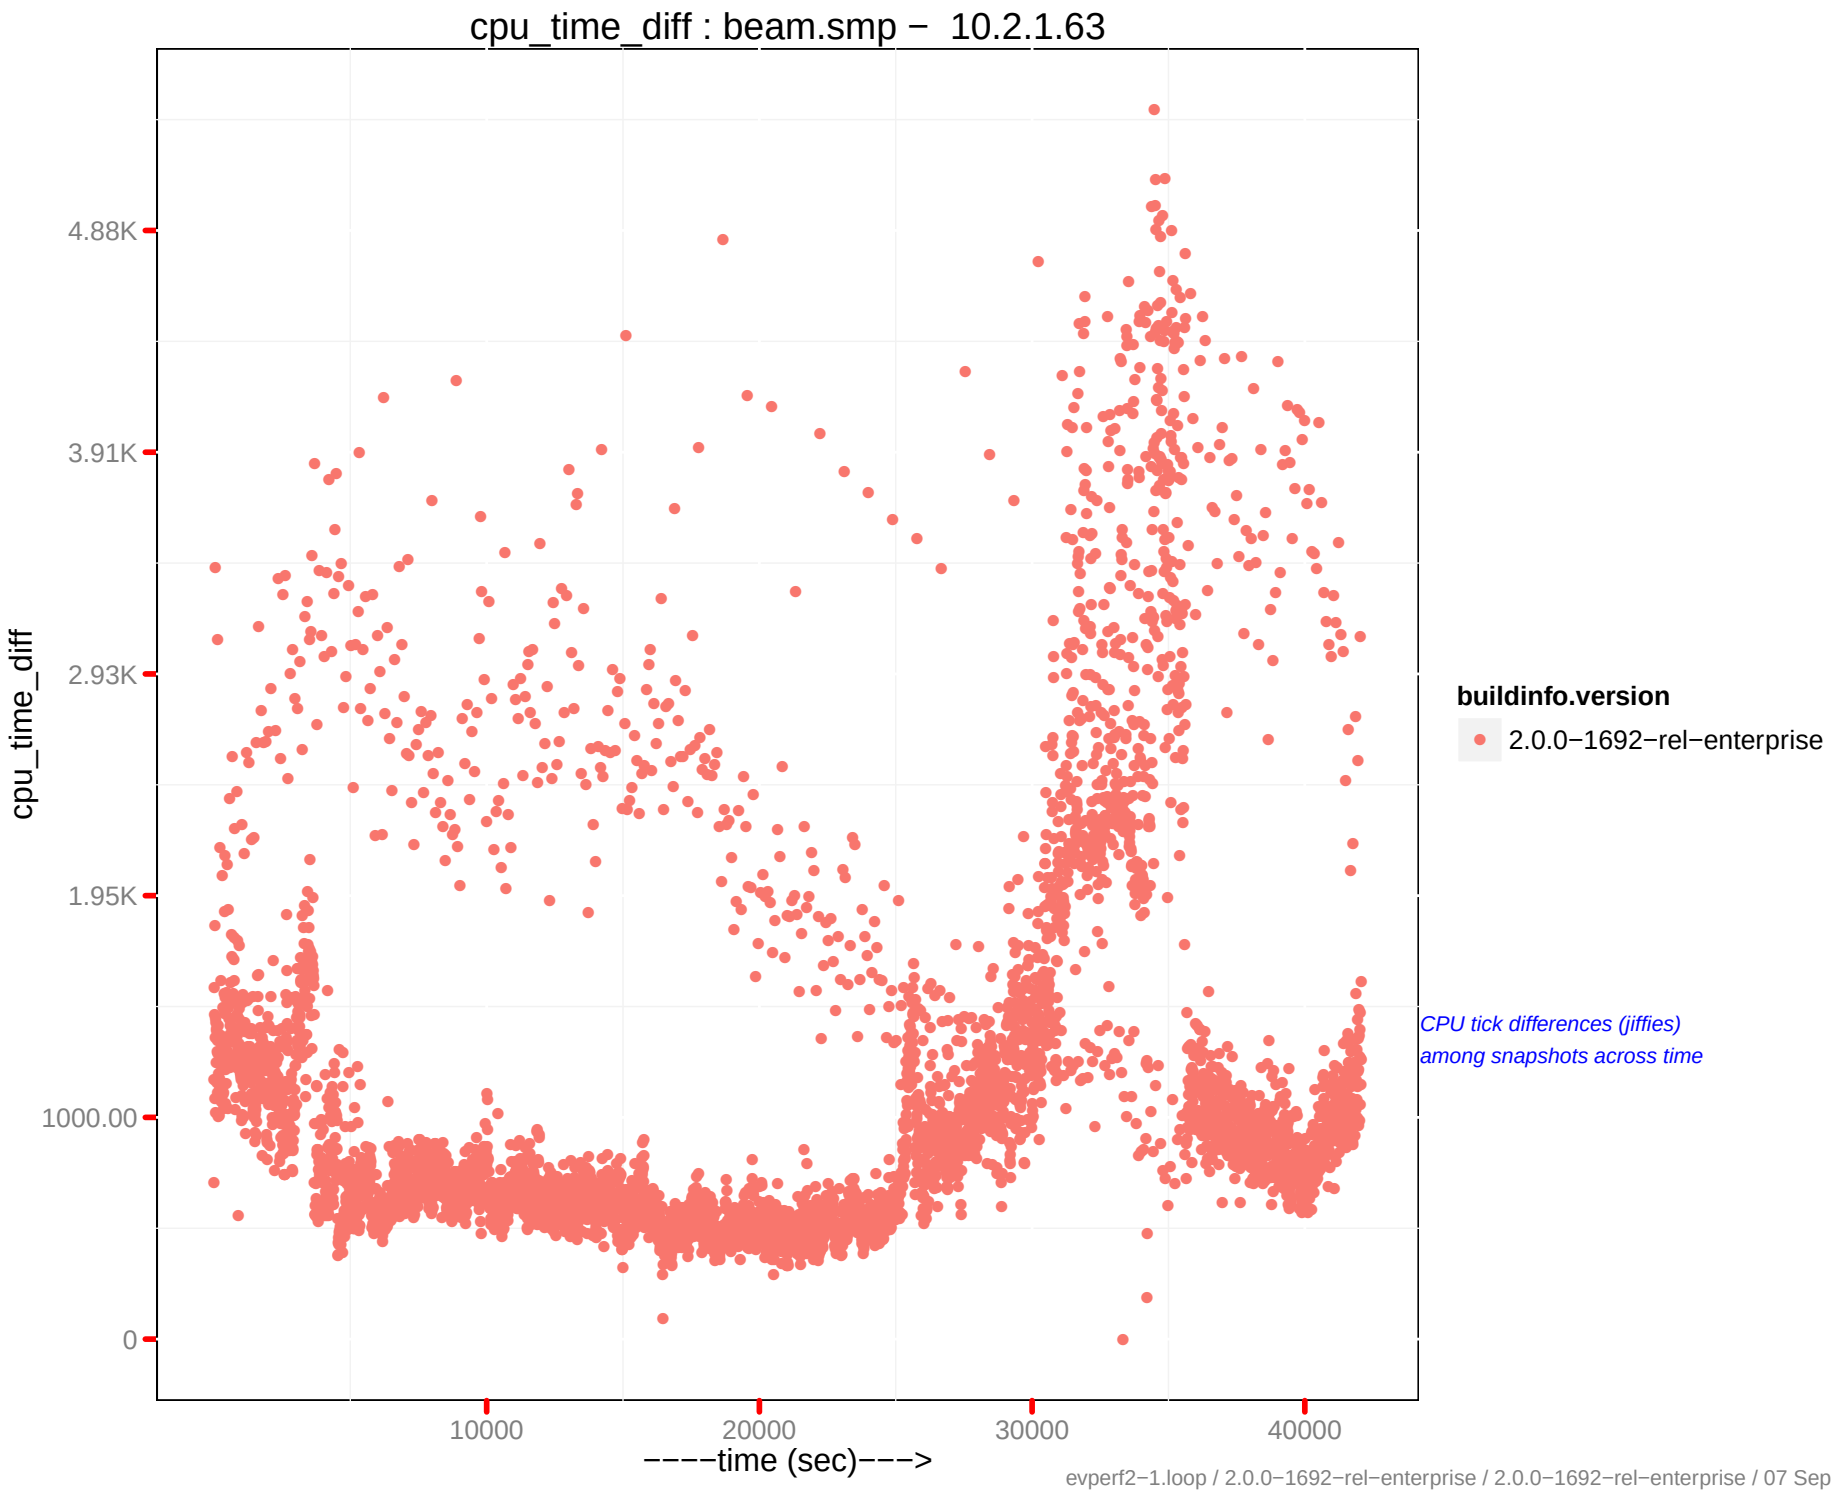
cpu_time_diff : beam.smp − 10.2.1.63
cpu_time_diff
4.88K
3.91K
2.93K
1.95K
1000.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

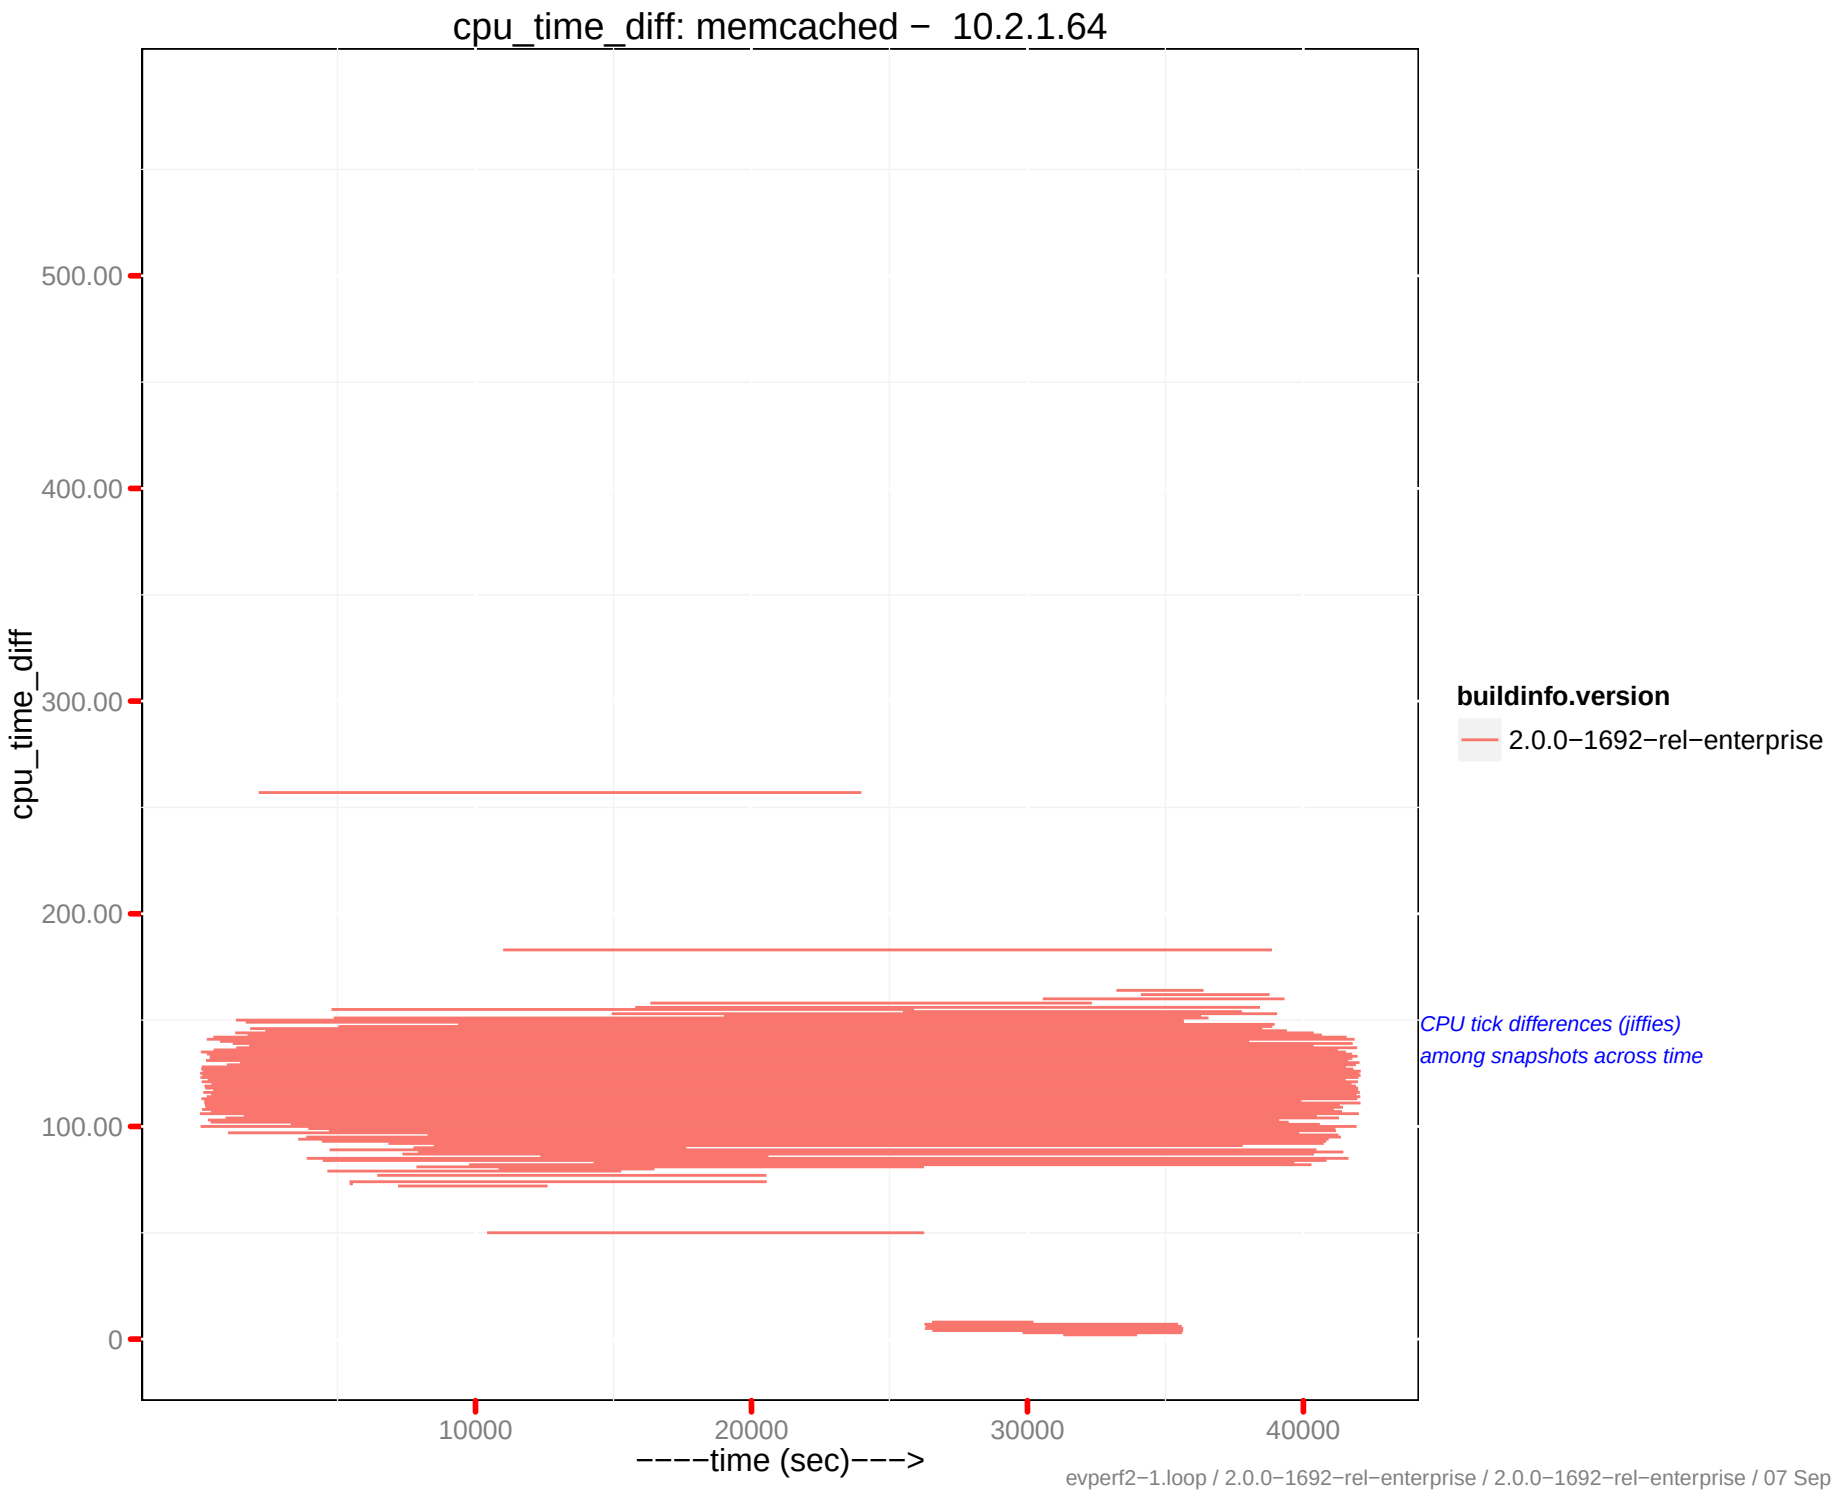
cpu_time_diff: memcached − 10.2.1.64
500.00
400.00
300.00
200.00
100.00
0
cpu_time_diff
10000
20000
30000
40000
----time (sec)--->
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

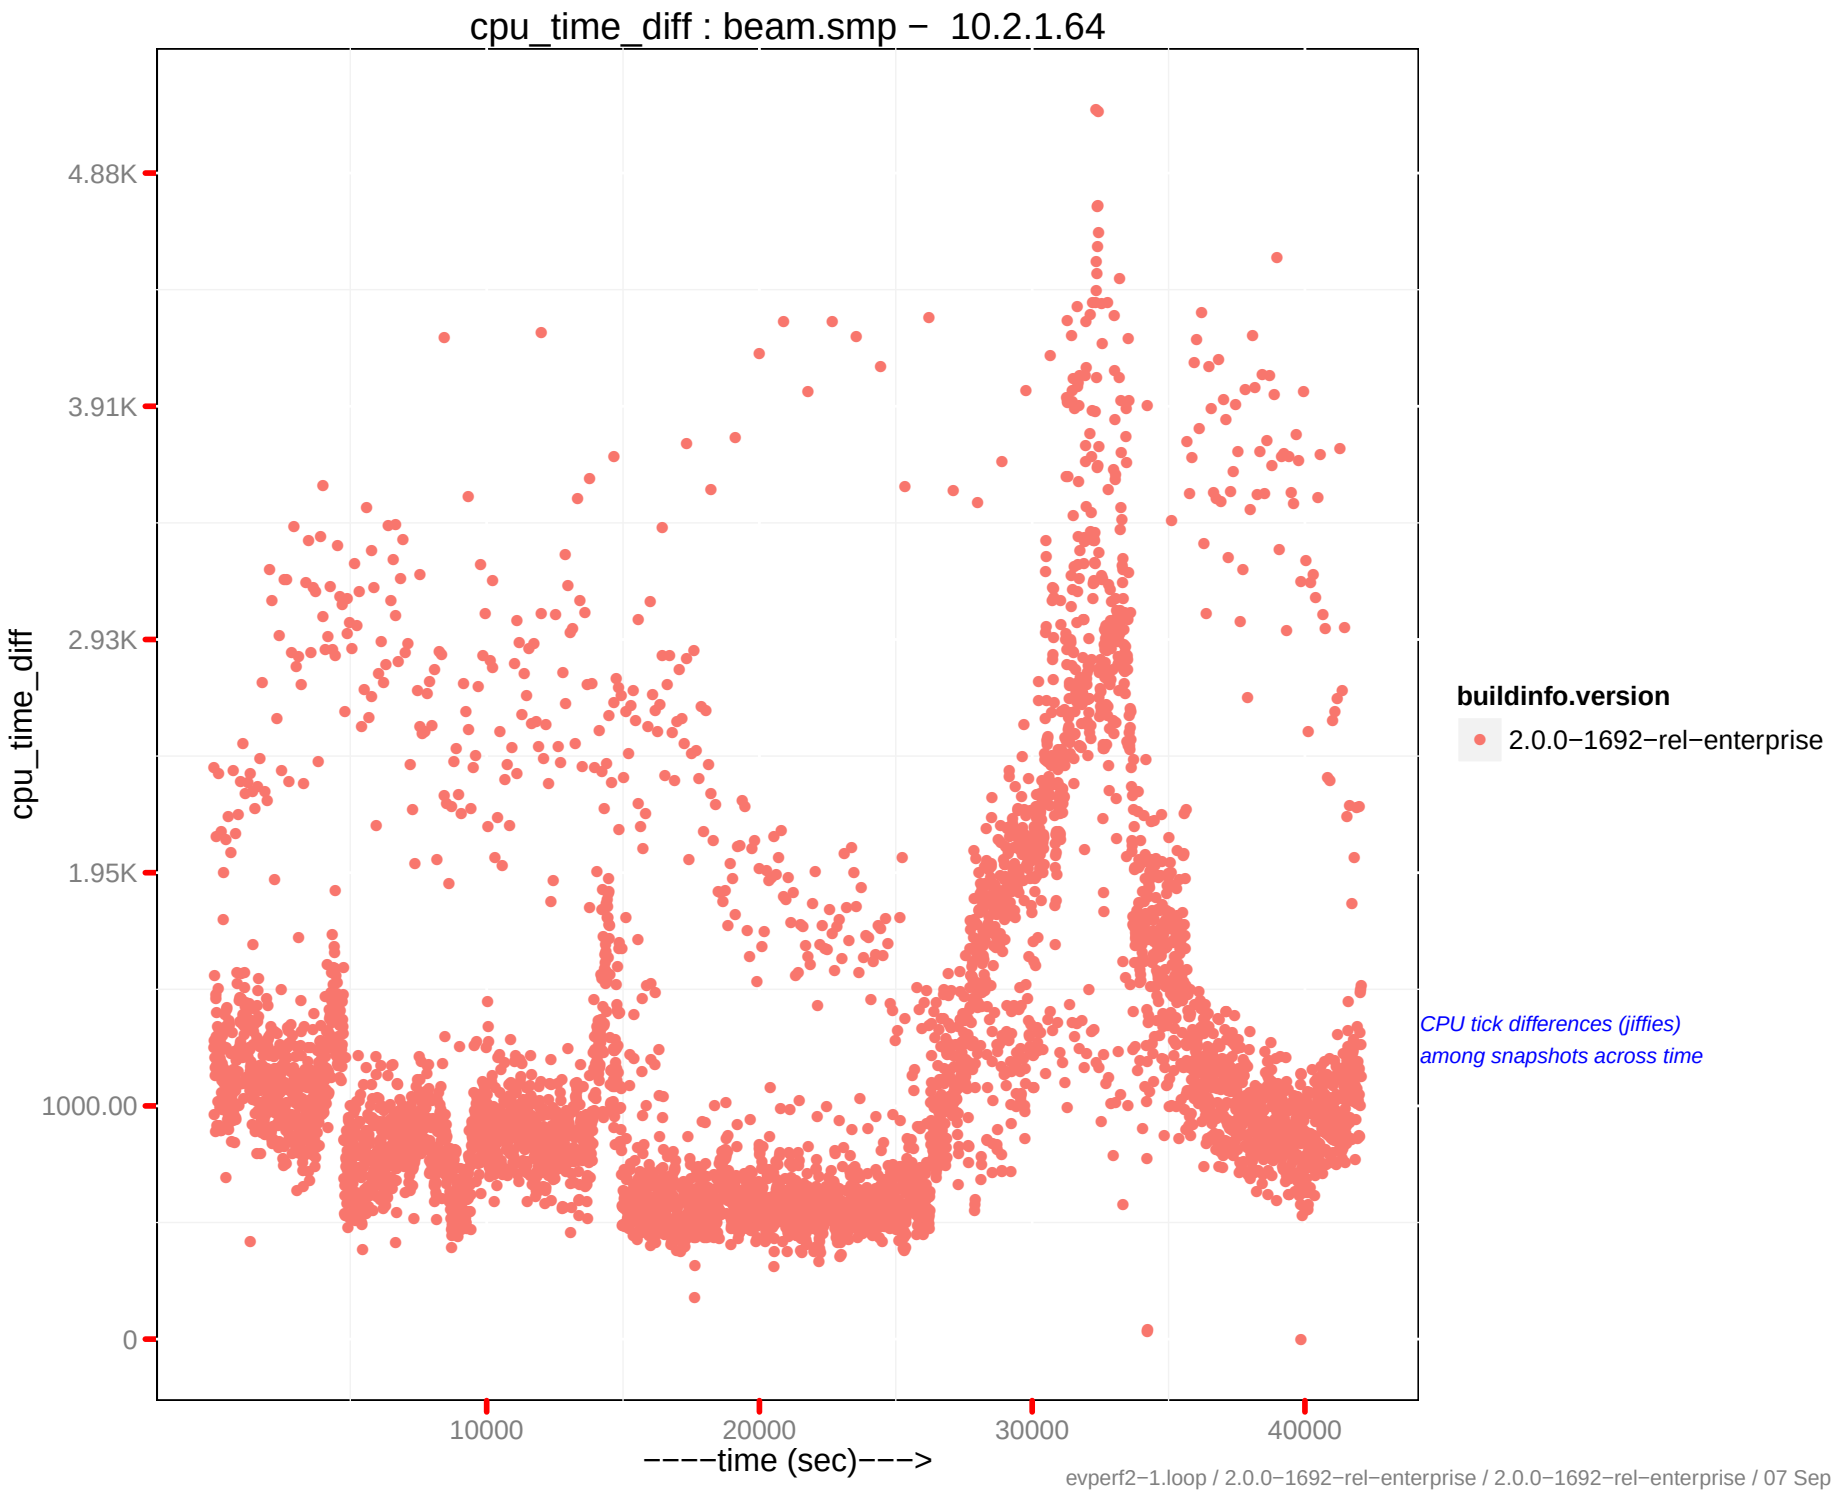
cpu_time_diff : beam.smp − 10.2.1.64
4.88K
3.91K
2.93K
1.95K
1000.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

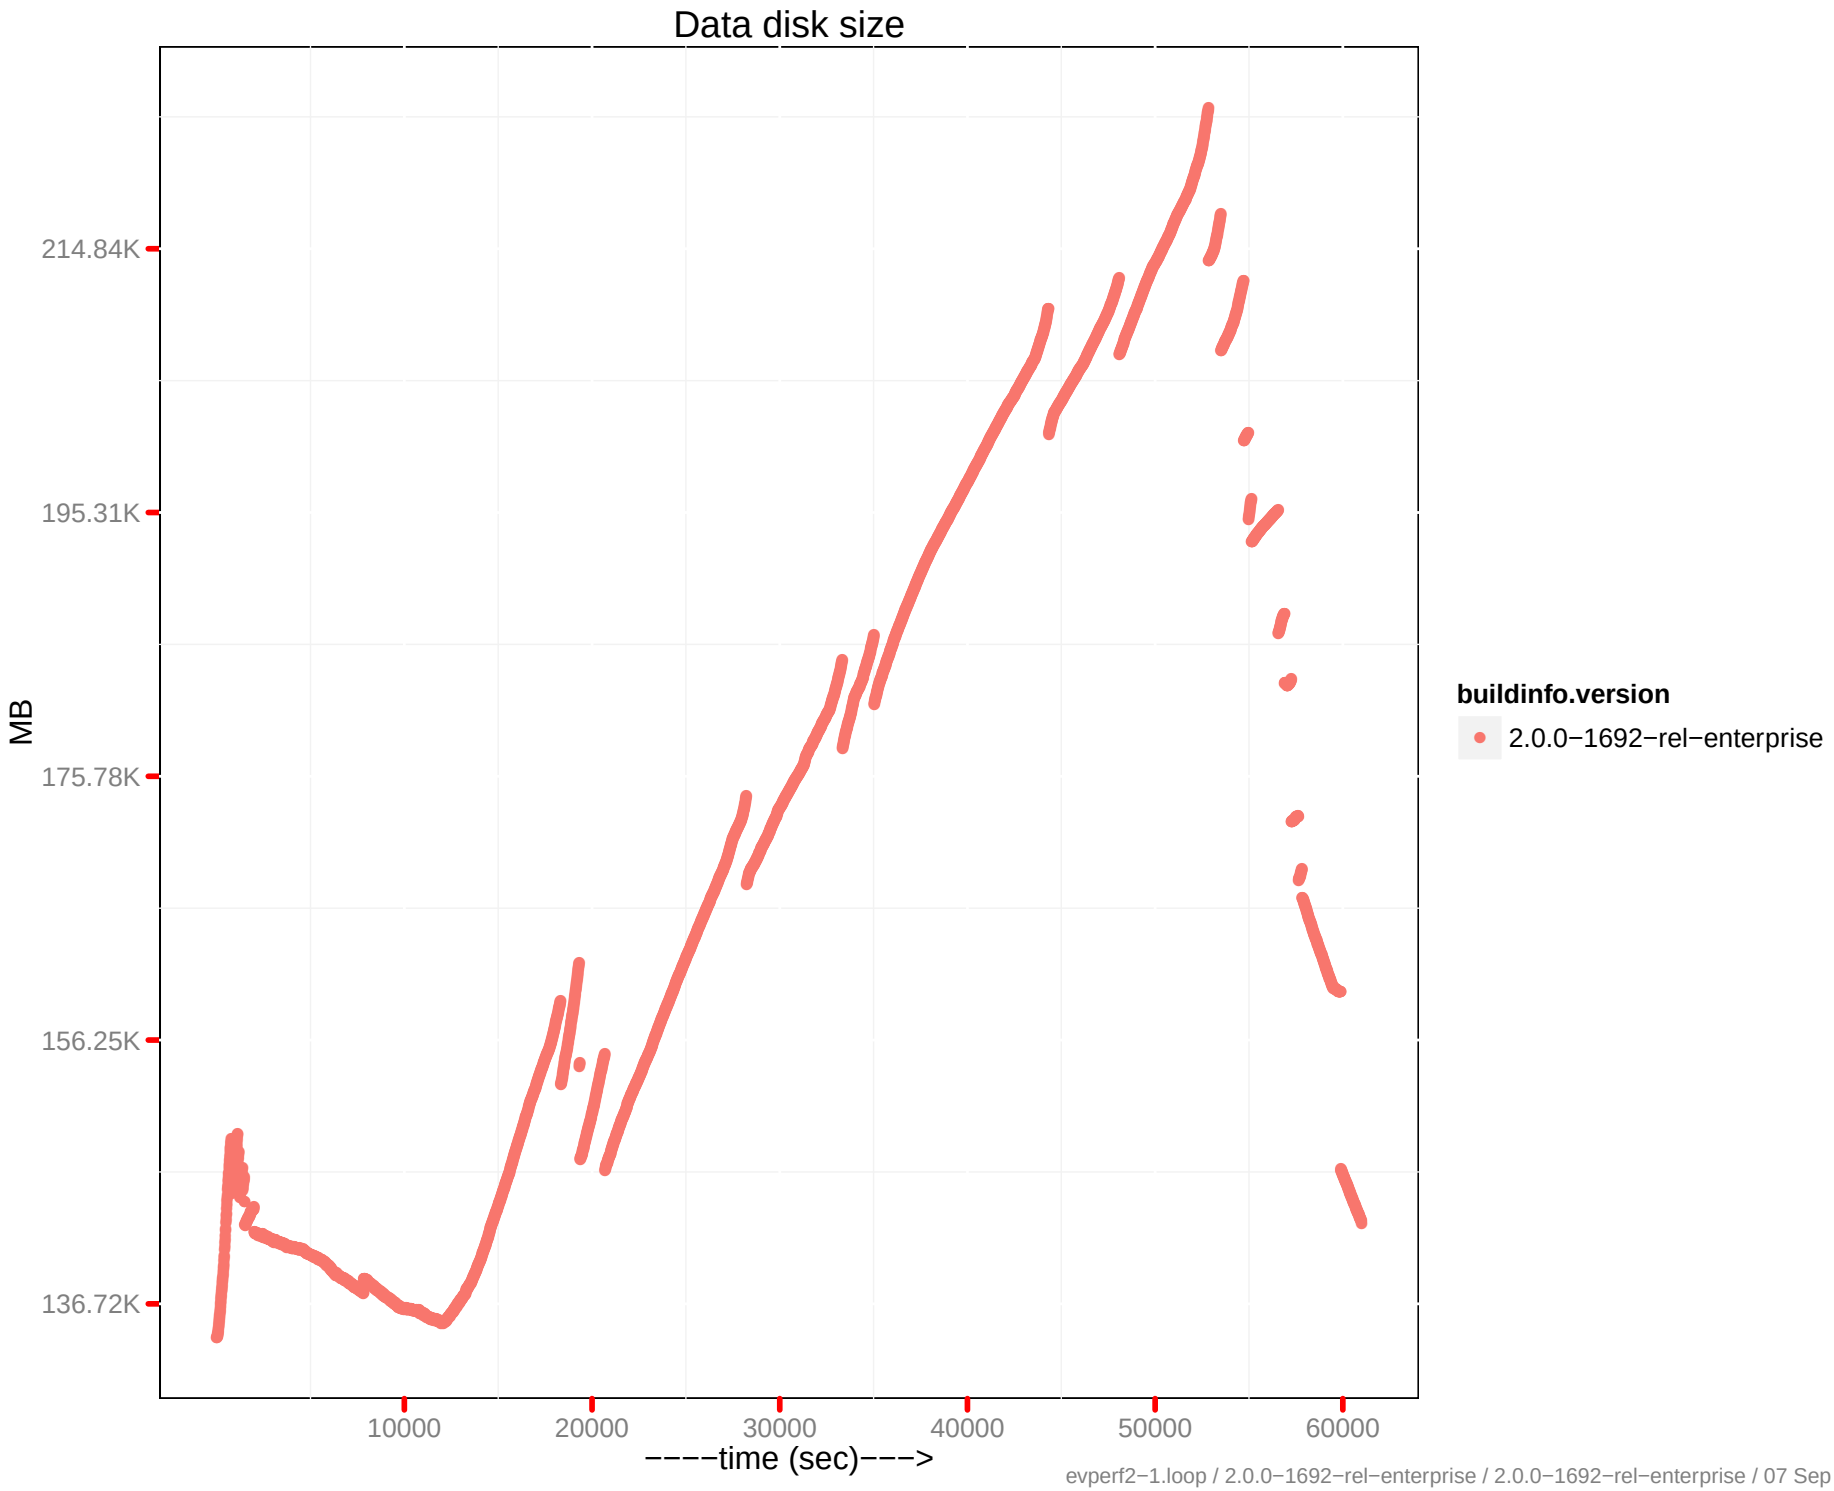
Data disk size
214.84K
195.31K
175.78K
156.25K
136.72K
MB
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

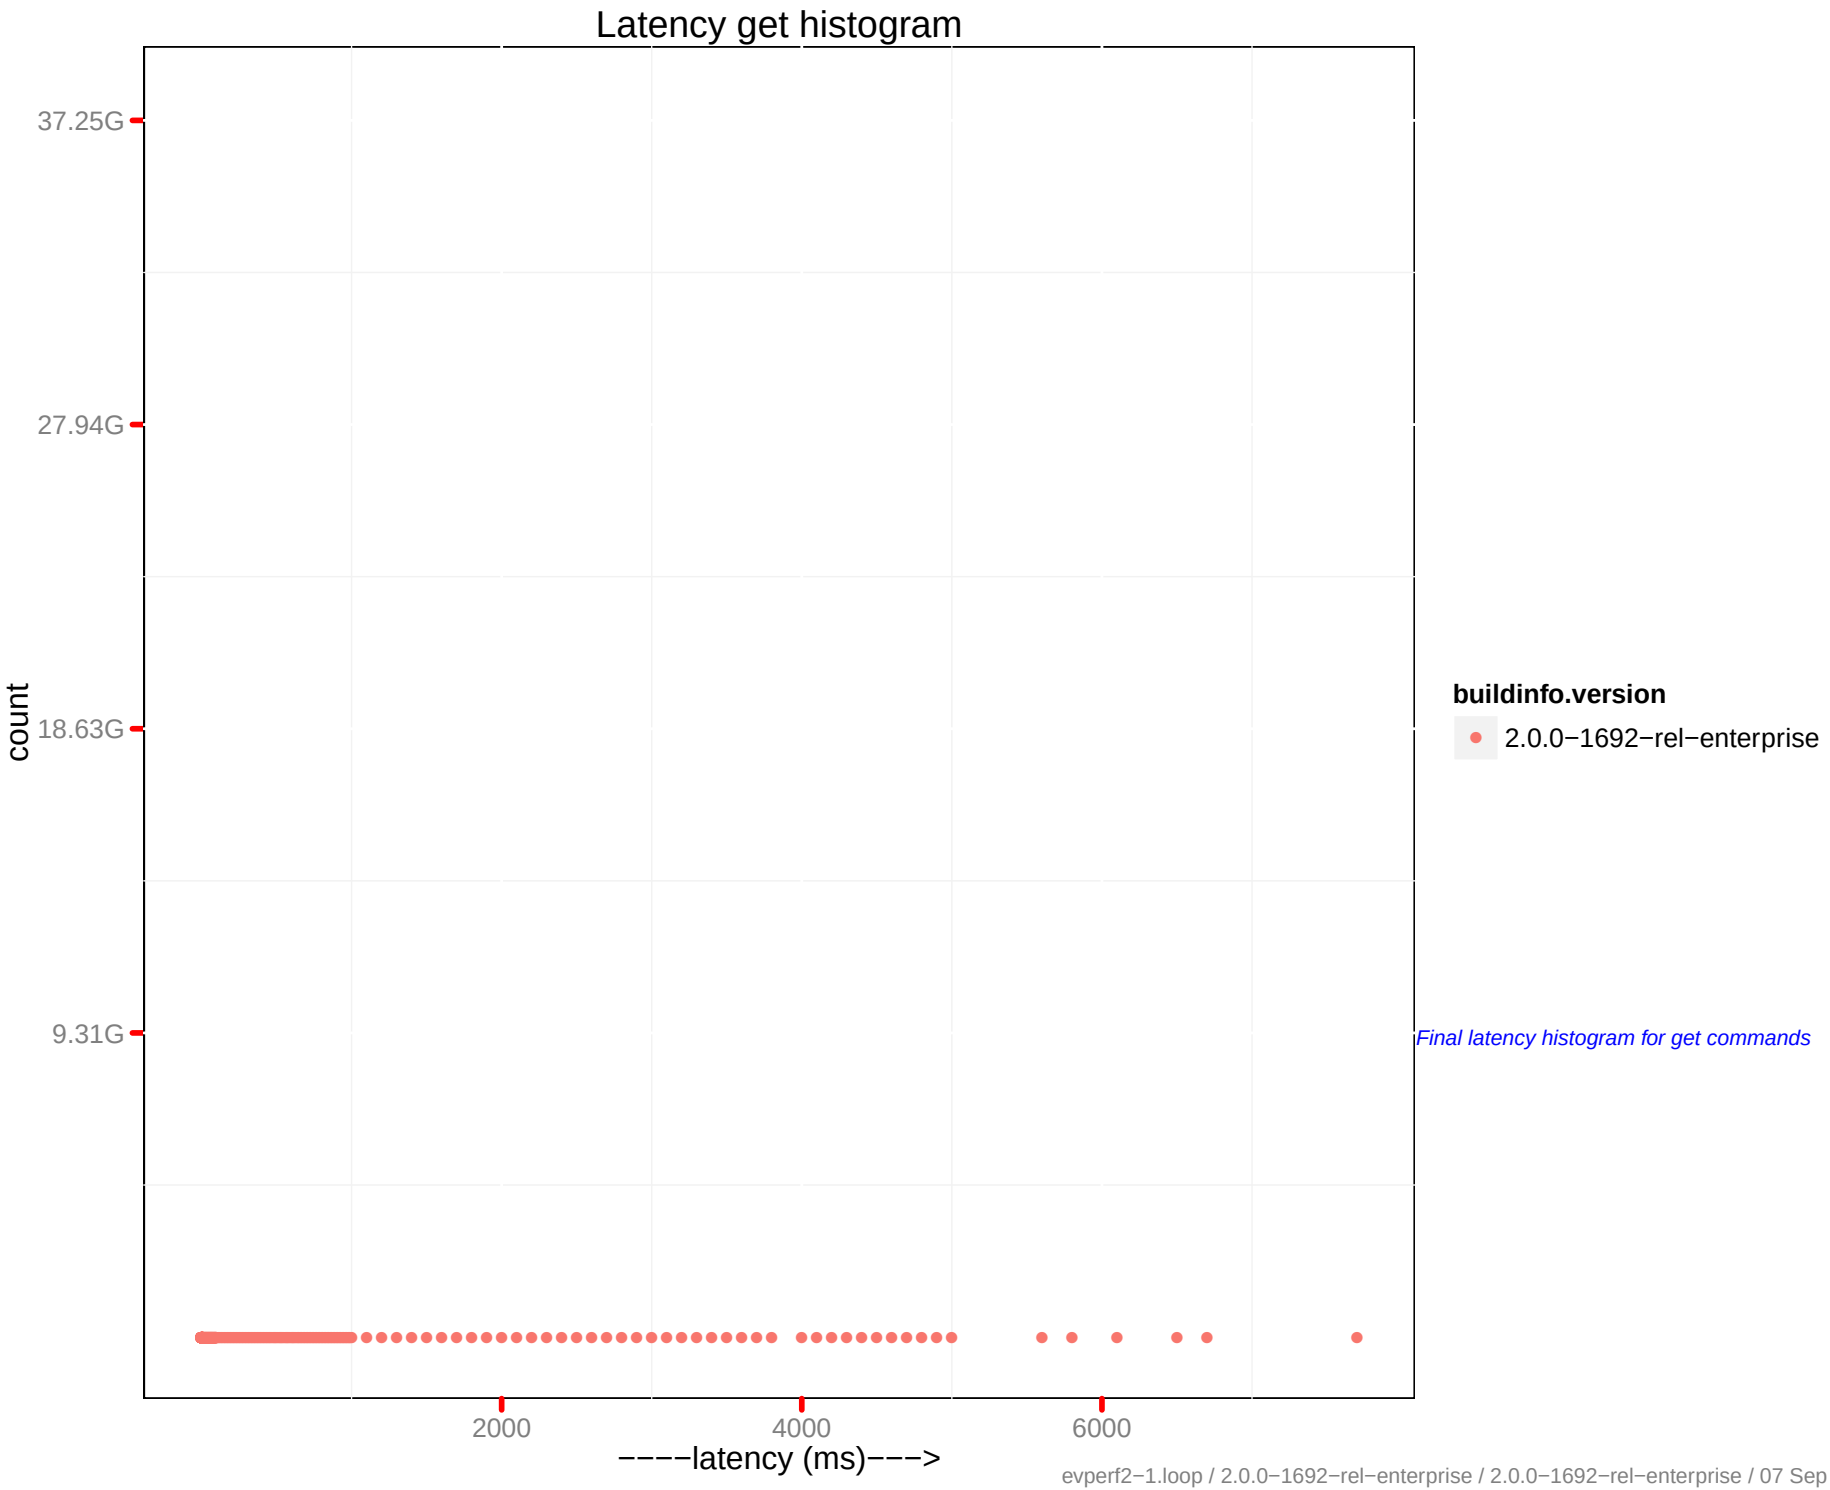
Latency get histogram
37.25G
27.94G
18.63G
9.31G
count
2000
4000
6000
----latency (ms)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Final latency histogram for get commands
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

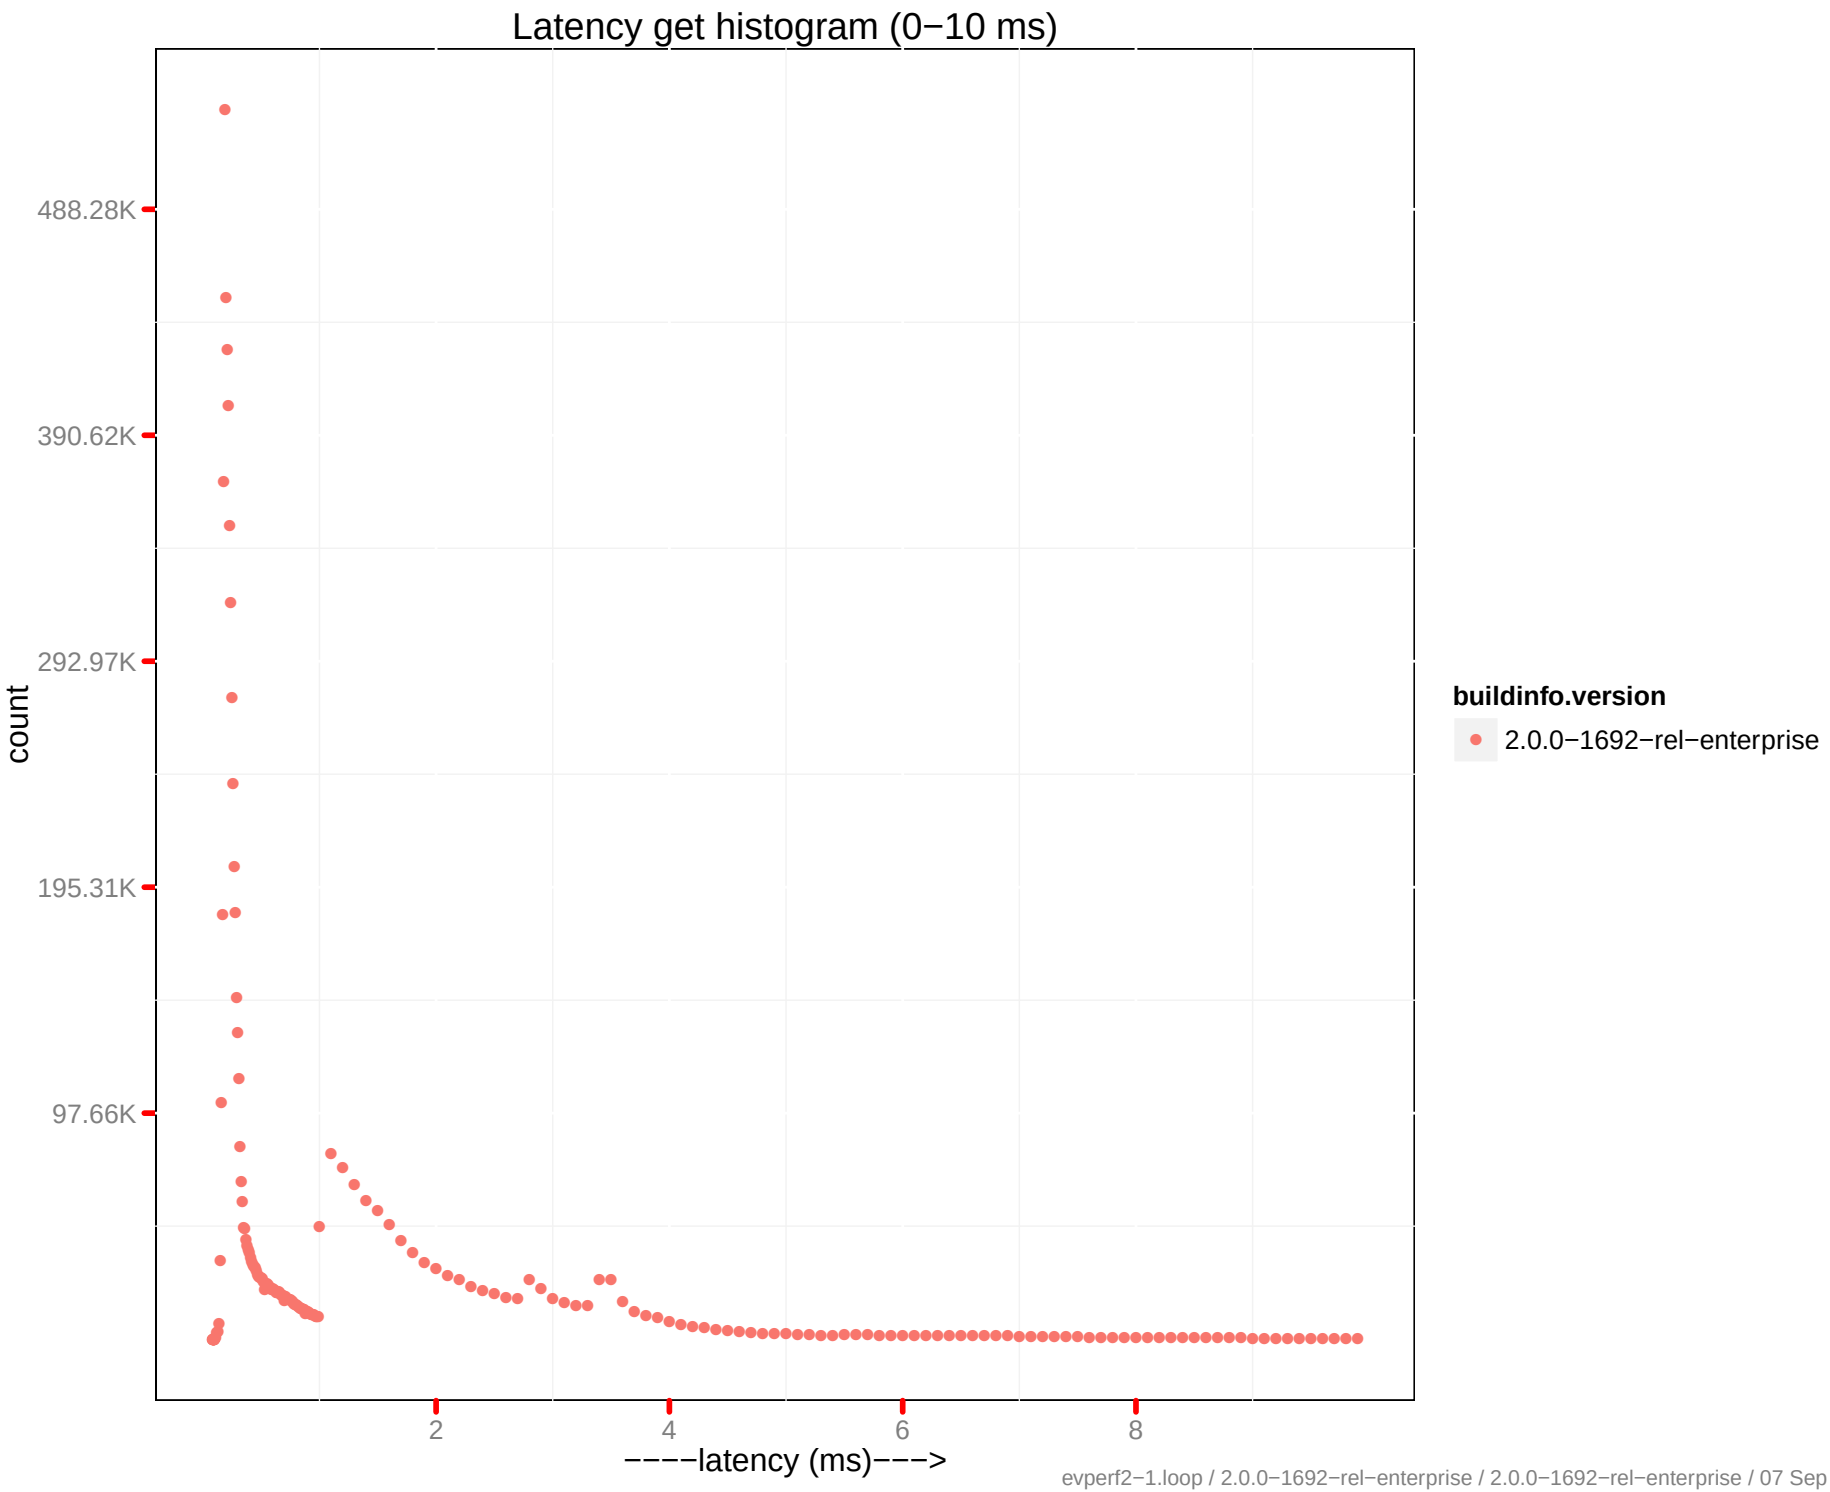
Latency get histogram (0−10 ms)
488.28K
390.62K
292.97K
195.31K
97.66K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

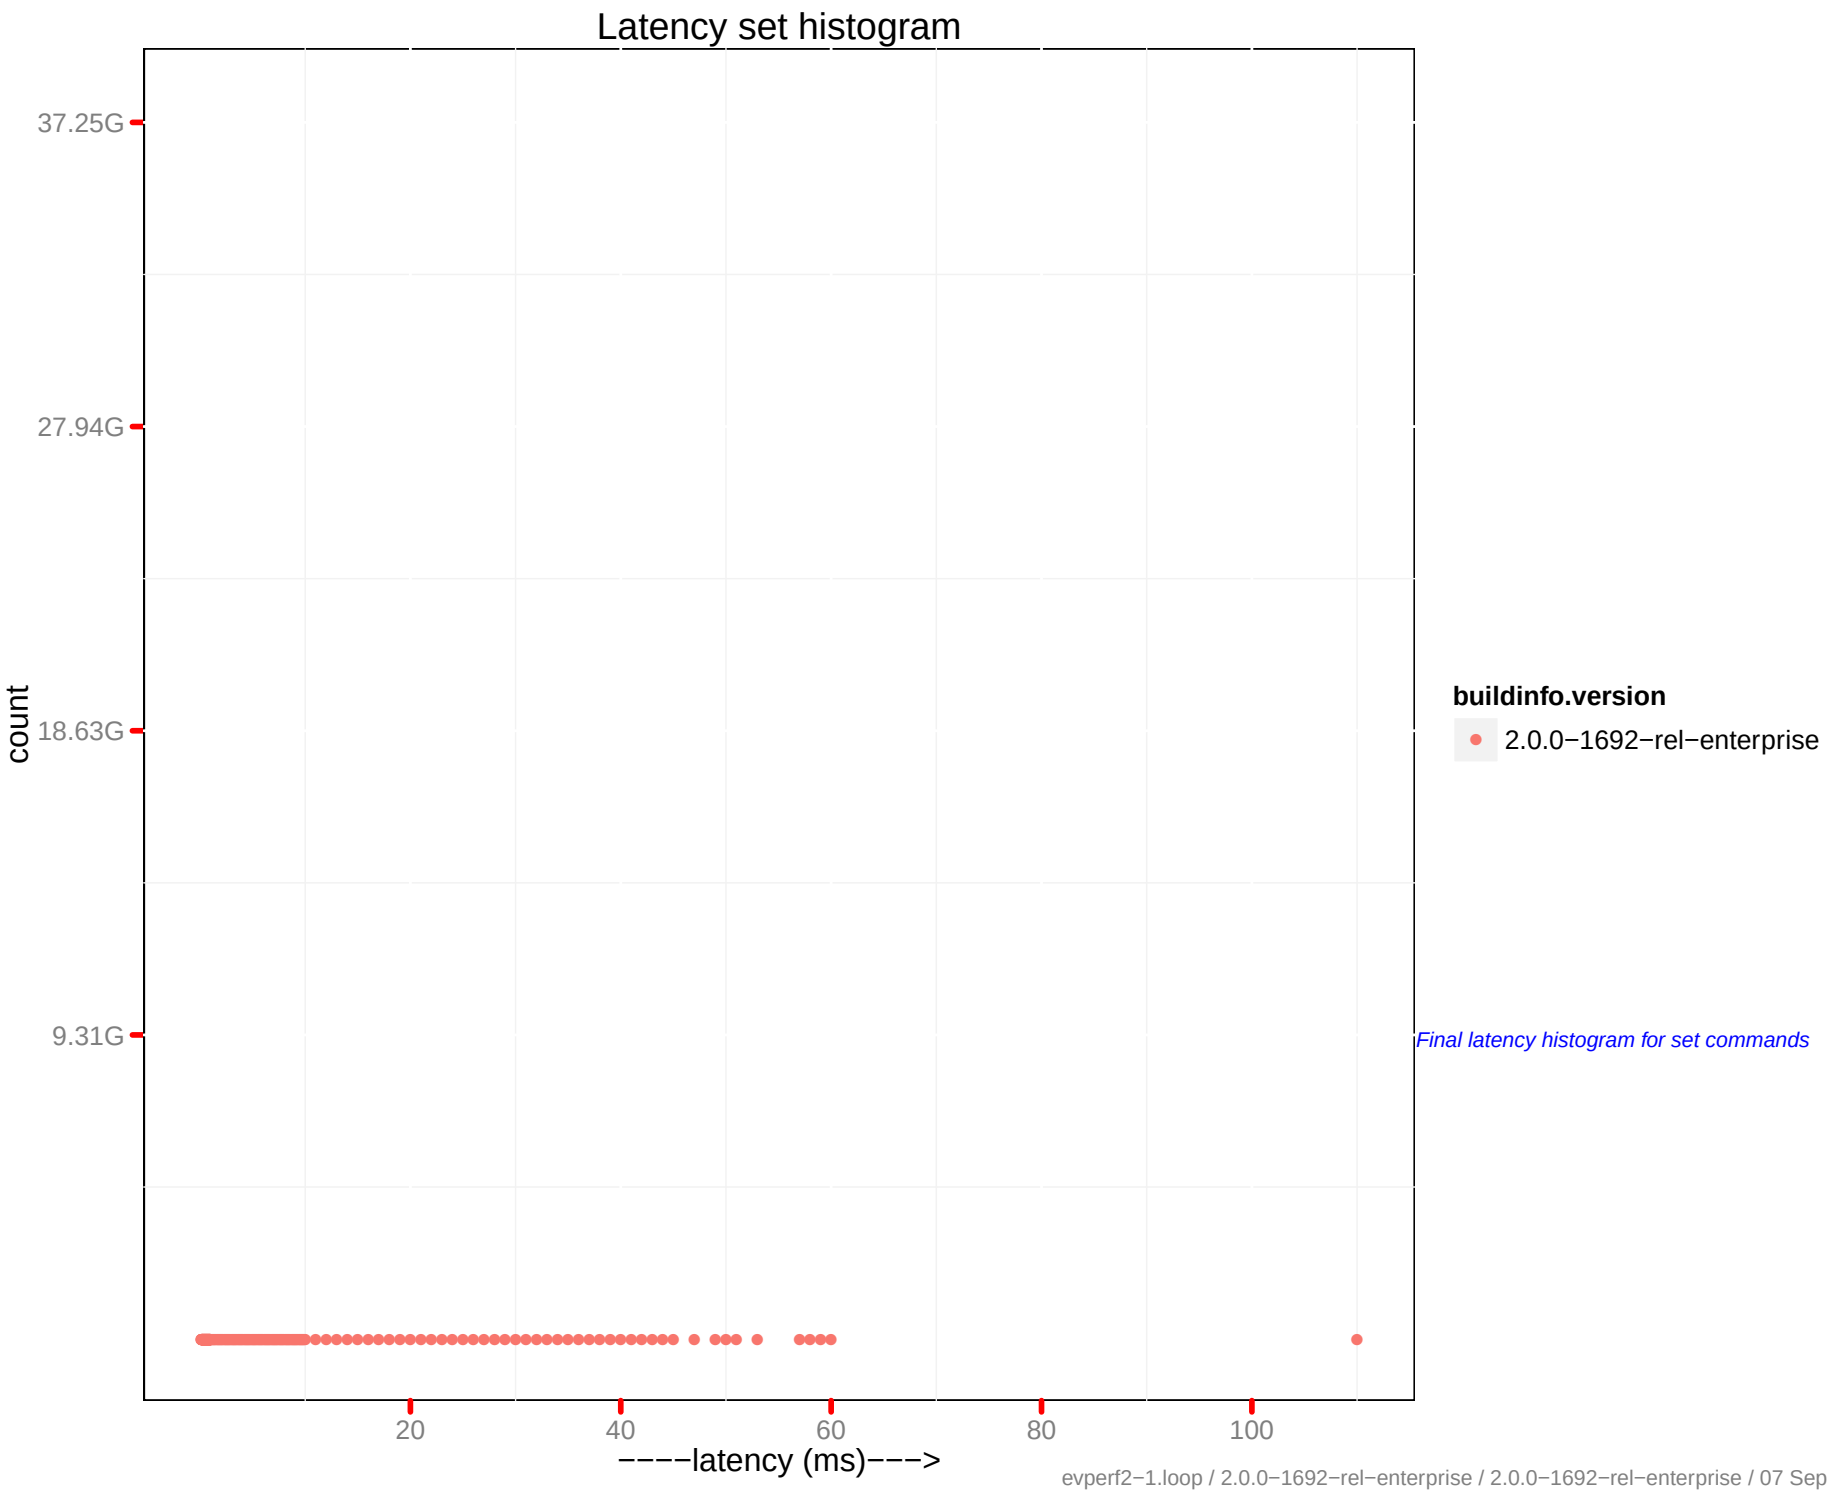
Latency set histogram
37.25G
27.94G
18.63G
9.31G
count
20
40
60
80
100
----latency (ms)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Final latency histogram for set commands
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

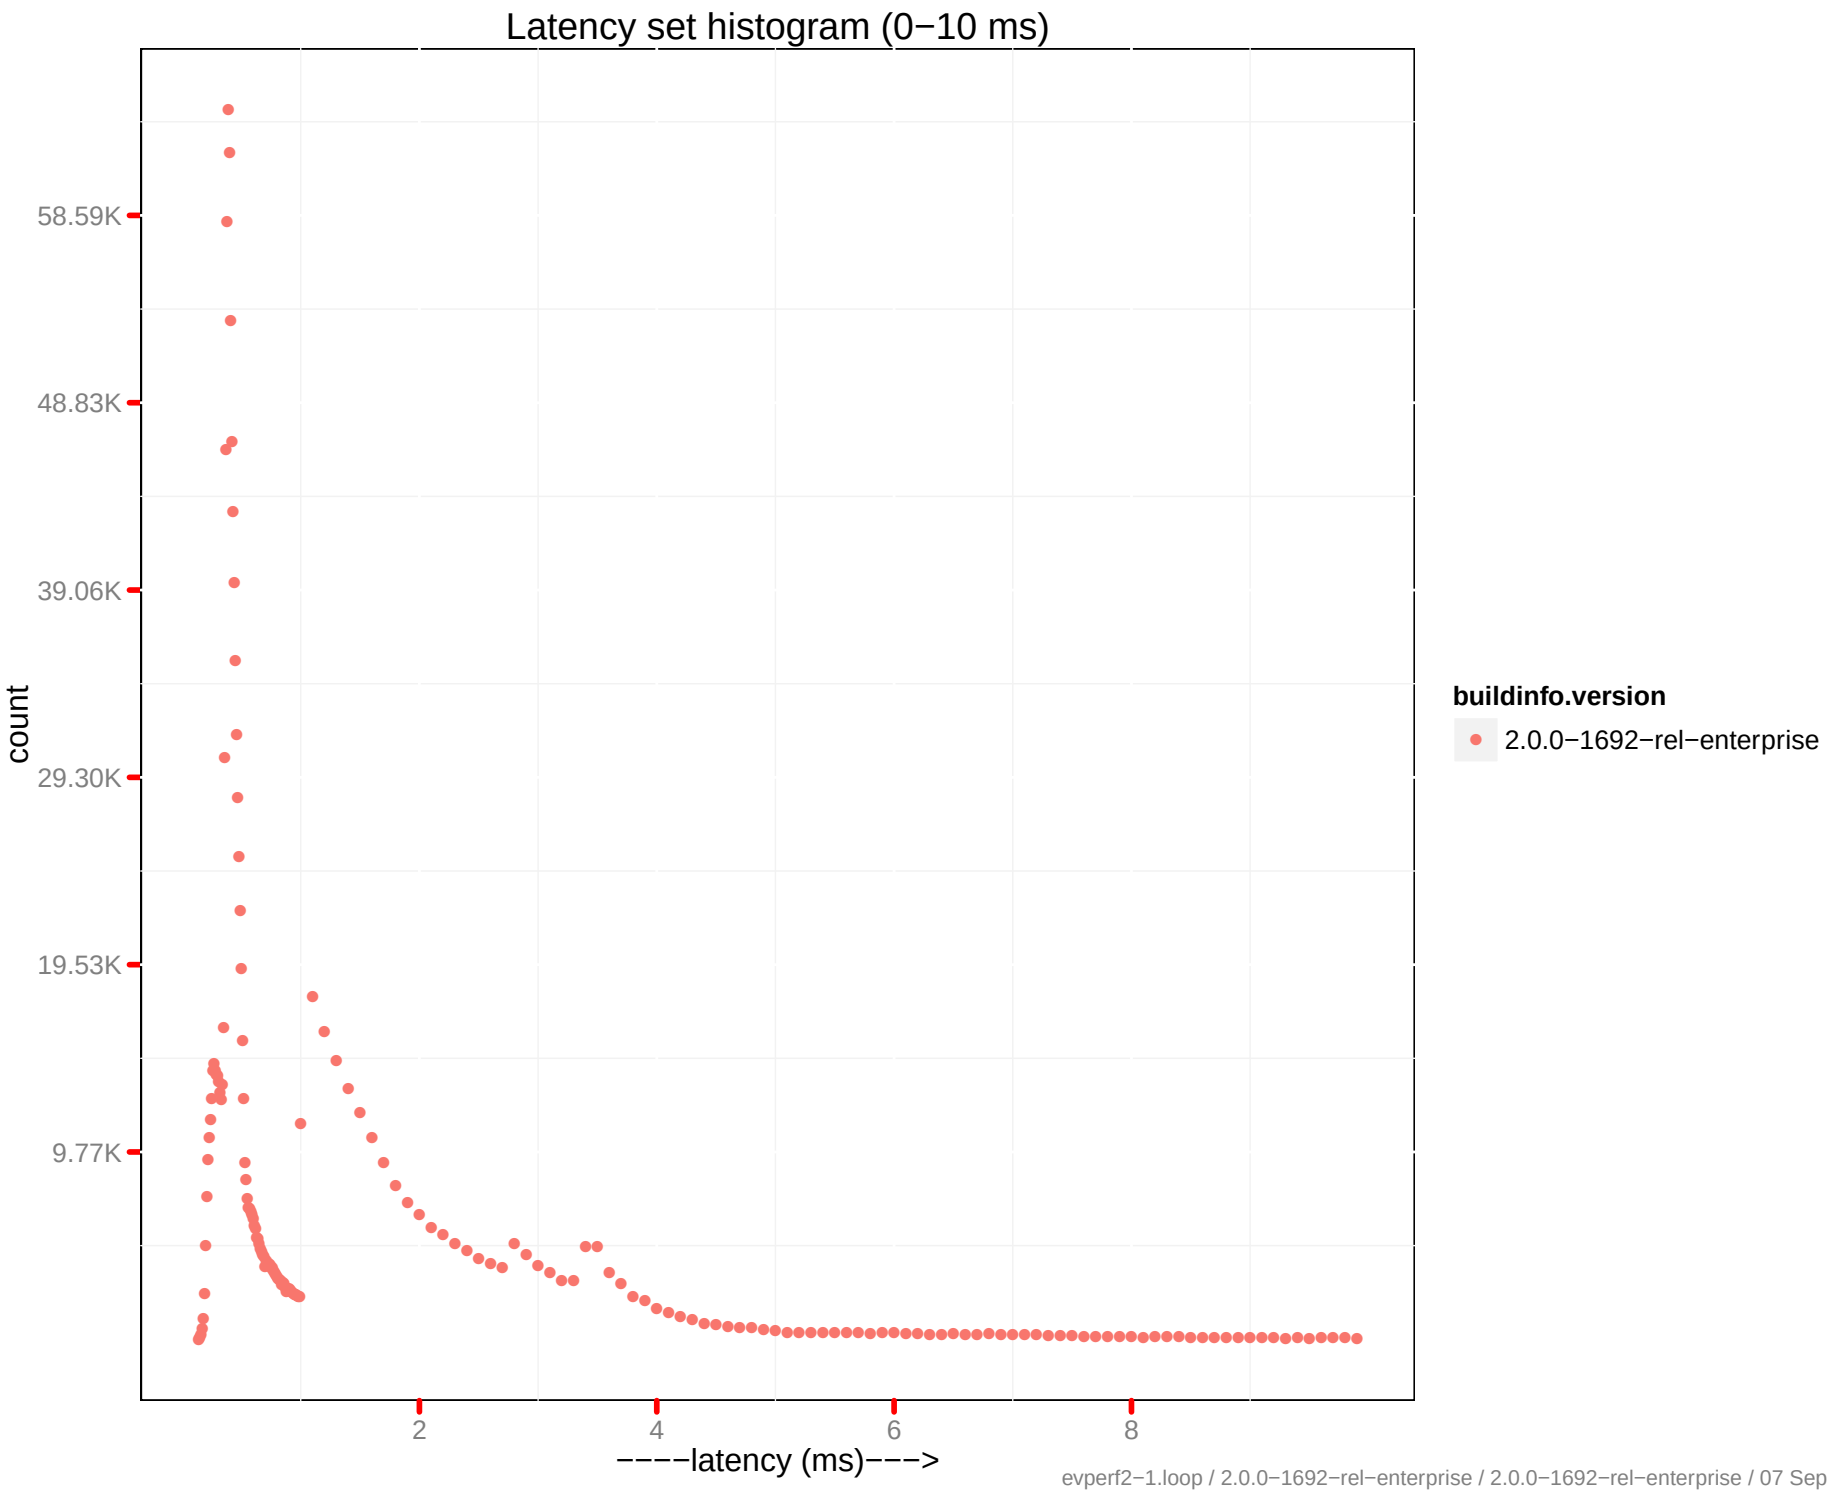
Latency set histogram (0−10 ms)
58.59K
48.83K
39.06K
29.30K
19.53K
9.77K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

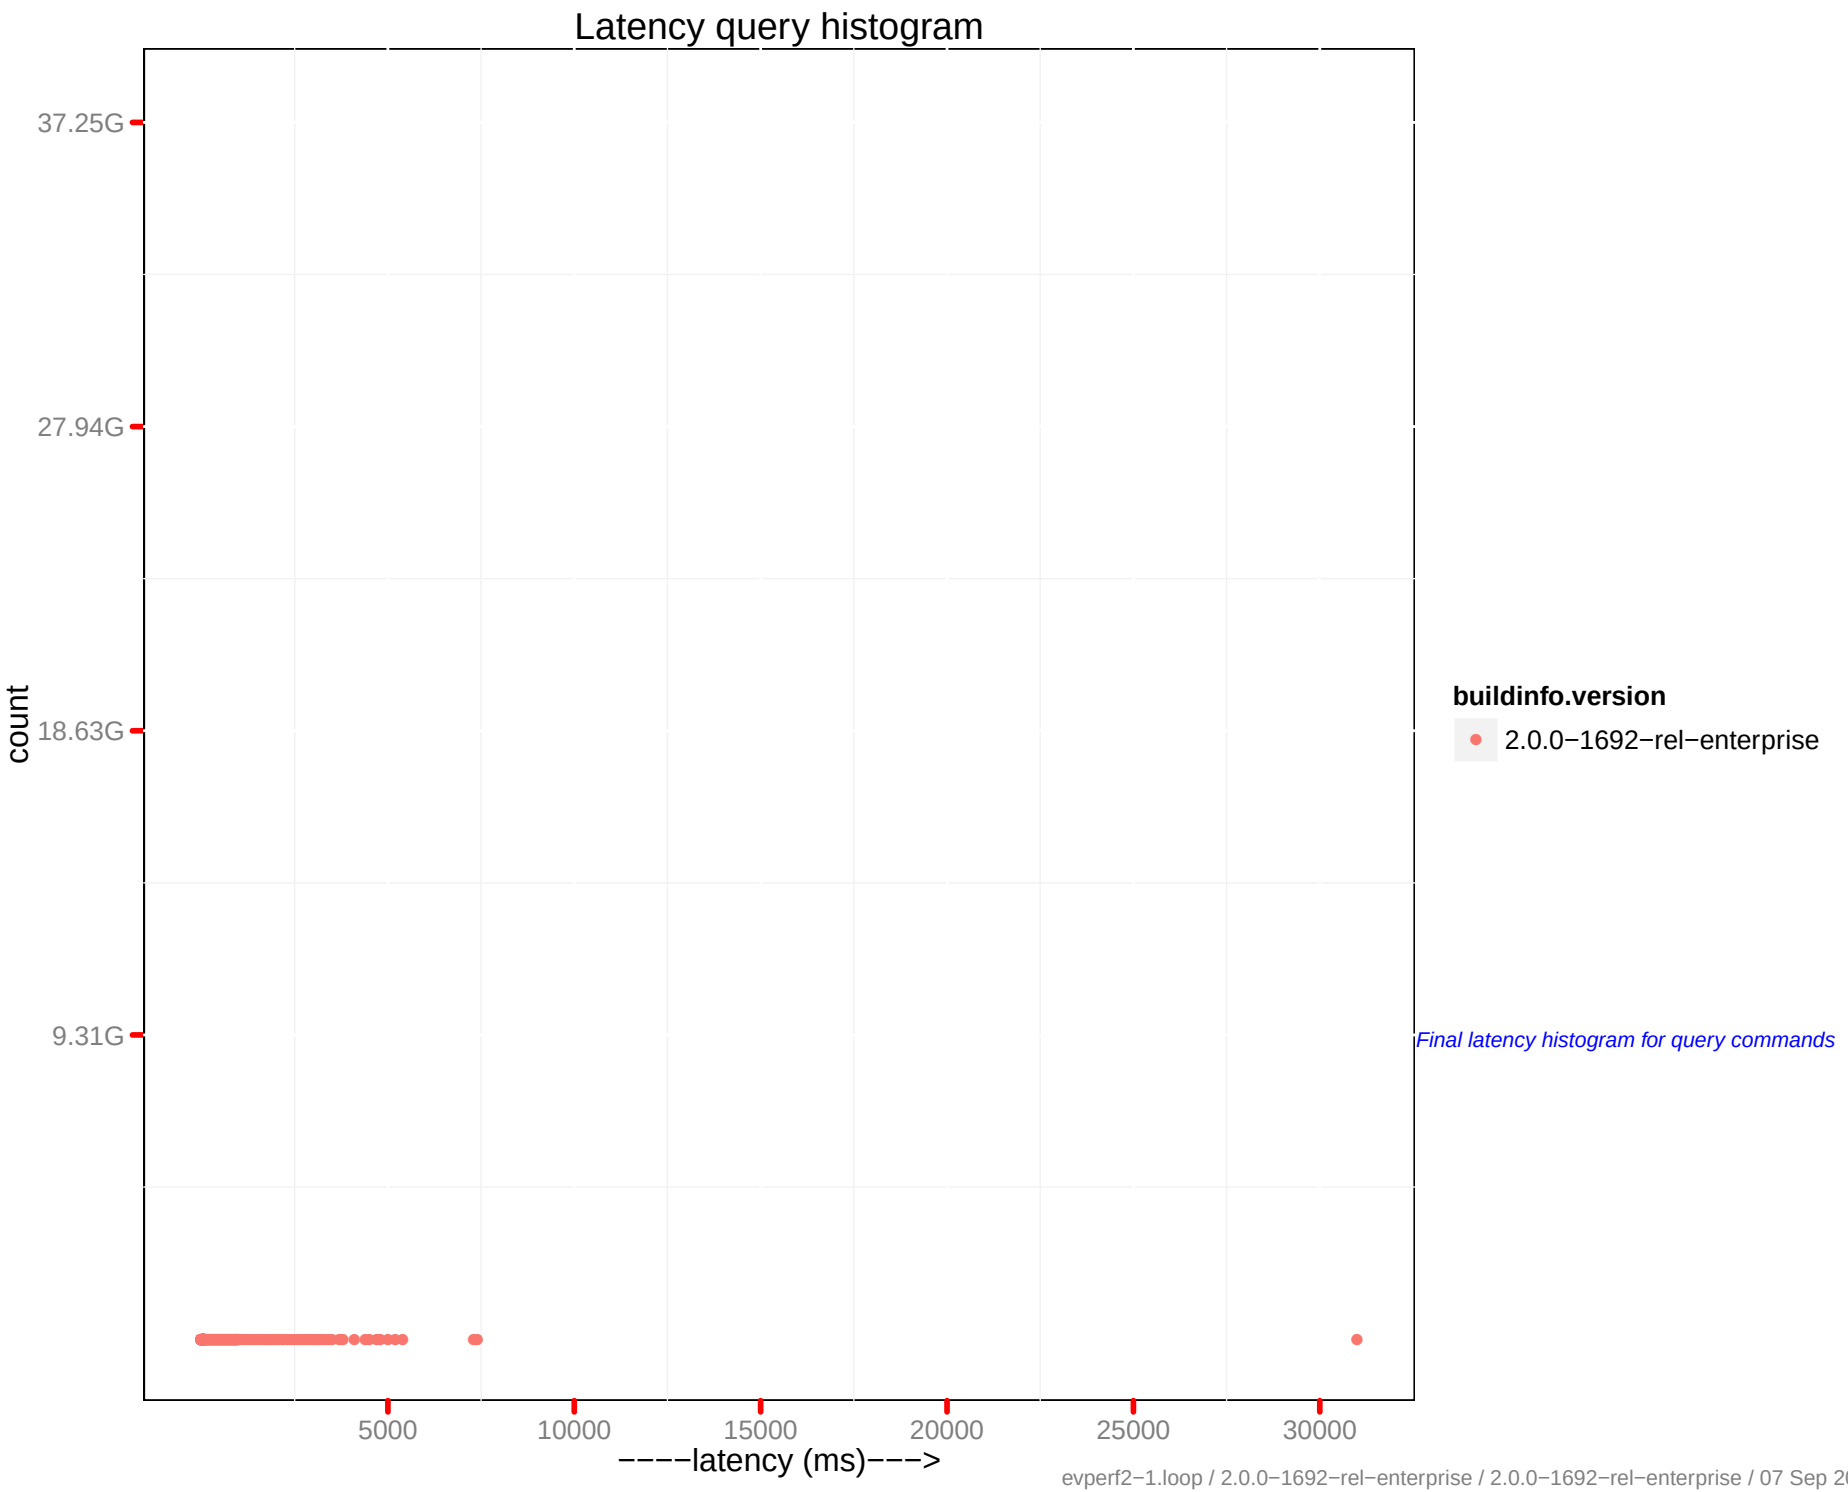
Latency query histogram
37.25G
27.94G
18.63G
9.31G
count
5000
10000
15000
20000
25000
30000
----latency (ms)--->
buildinfo.version
2.0.0-1692-rel-enterprise
Final latency histogram for query commands
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

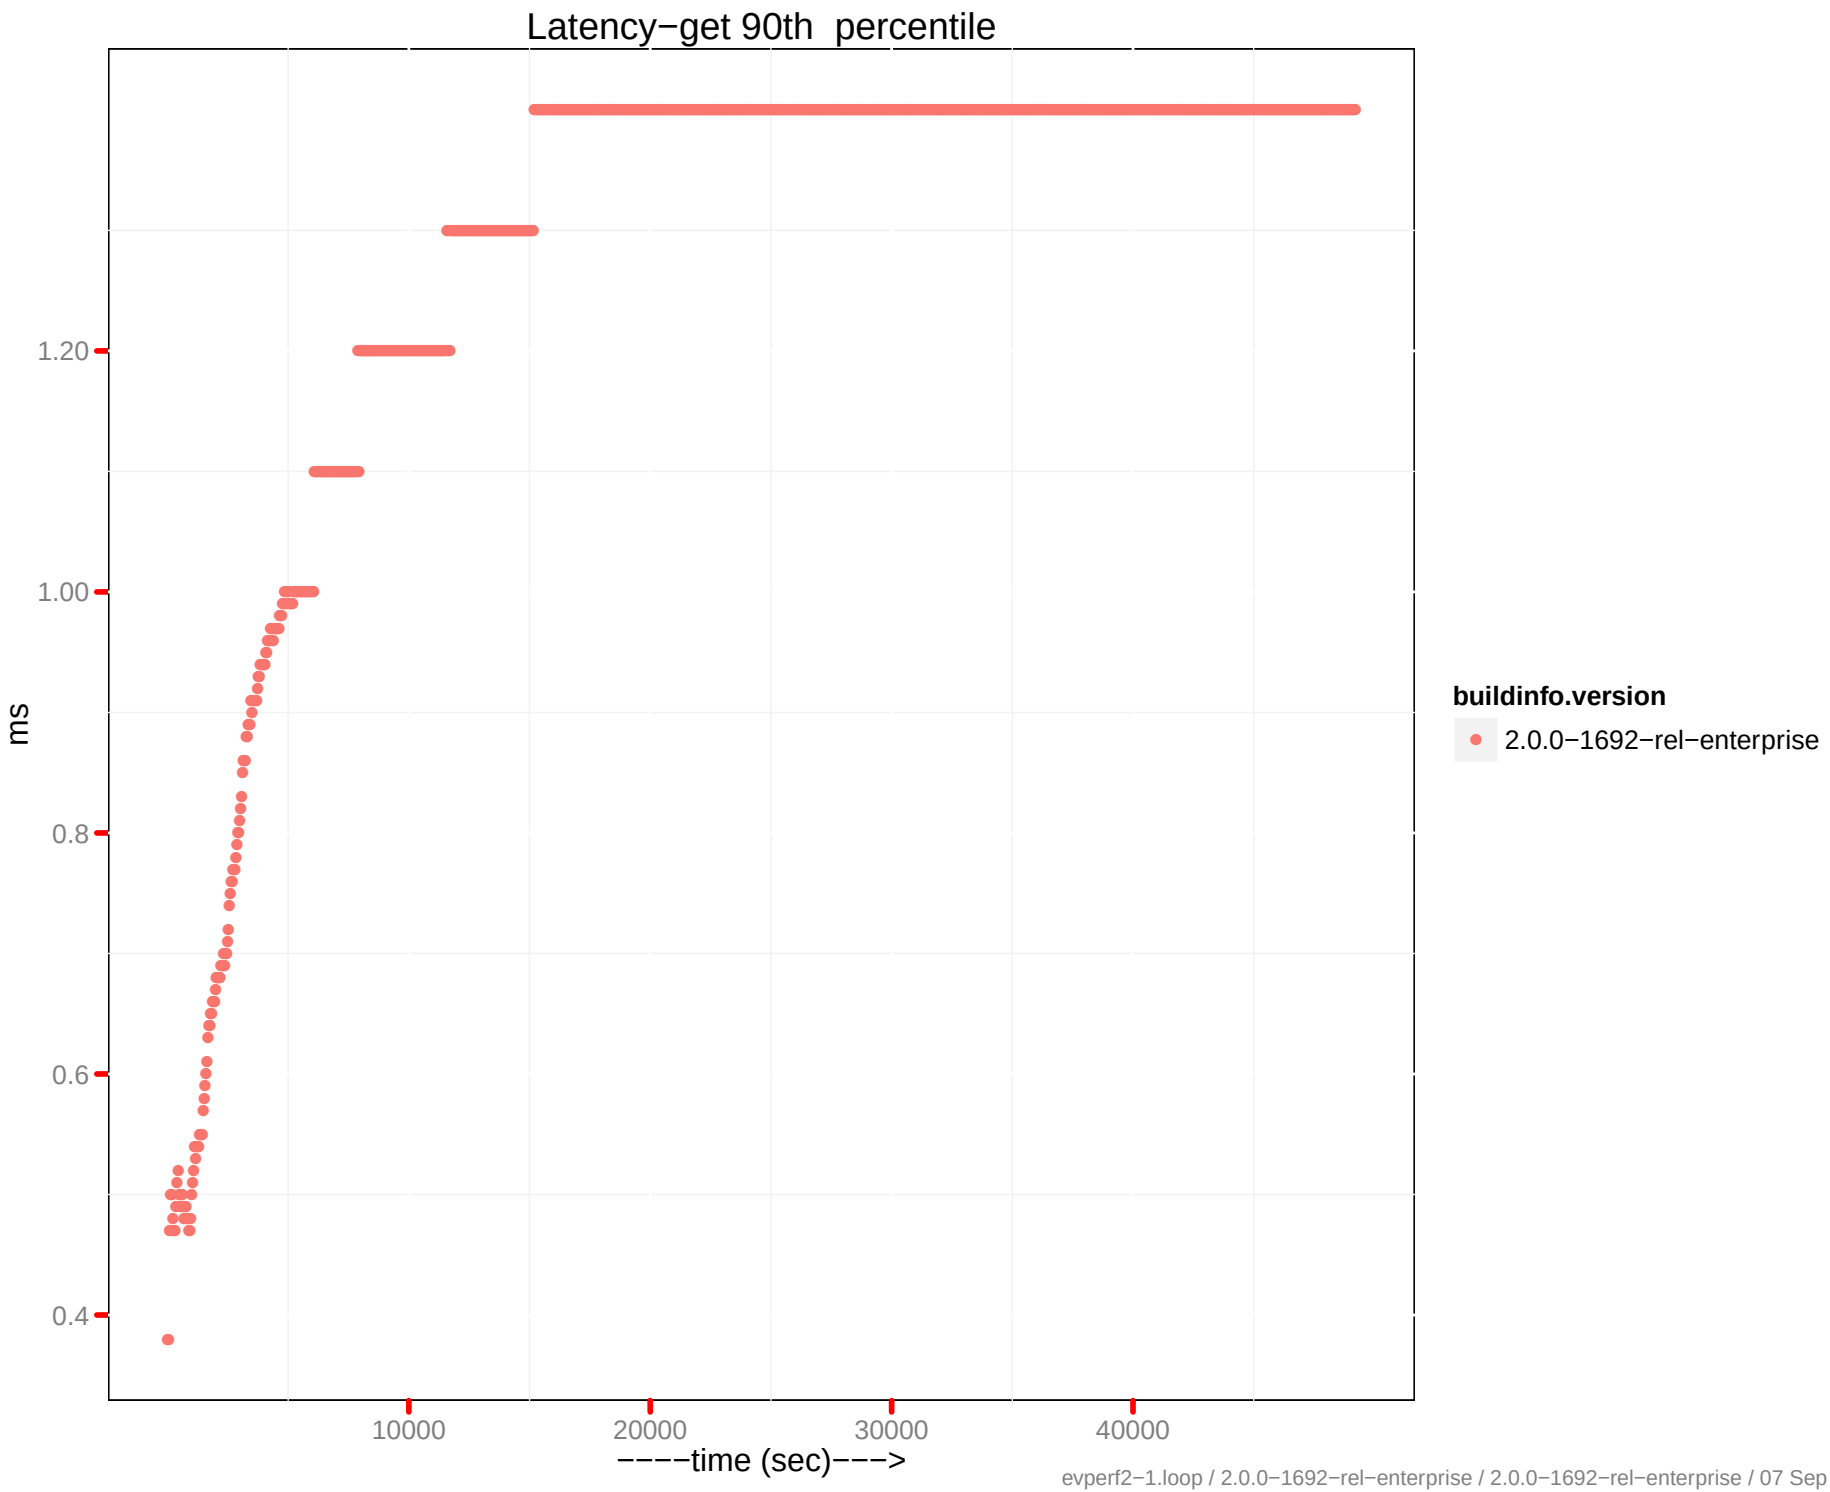
Latency−get 90th percentile
1.20
1.00
0.8
0.6
0.4
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

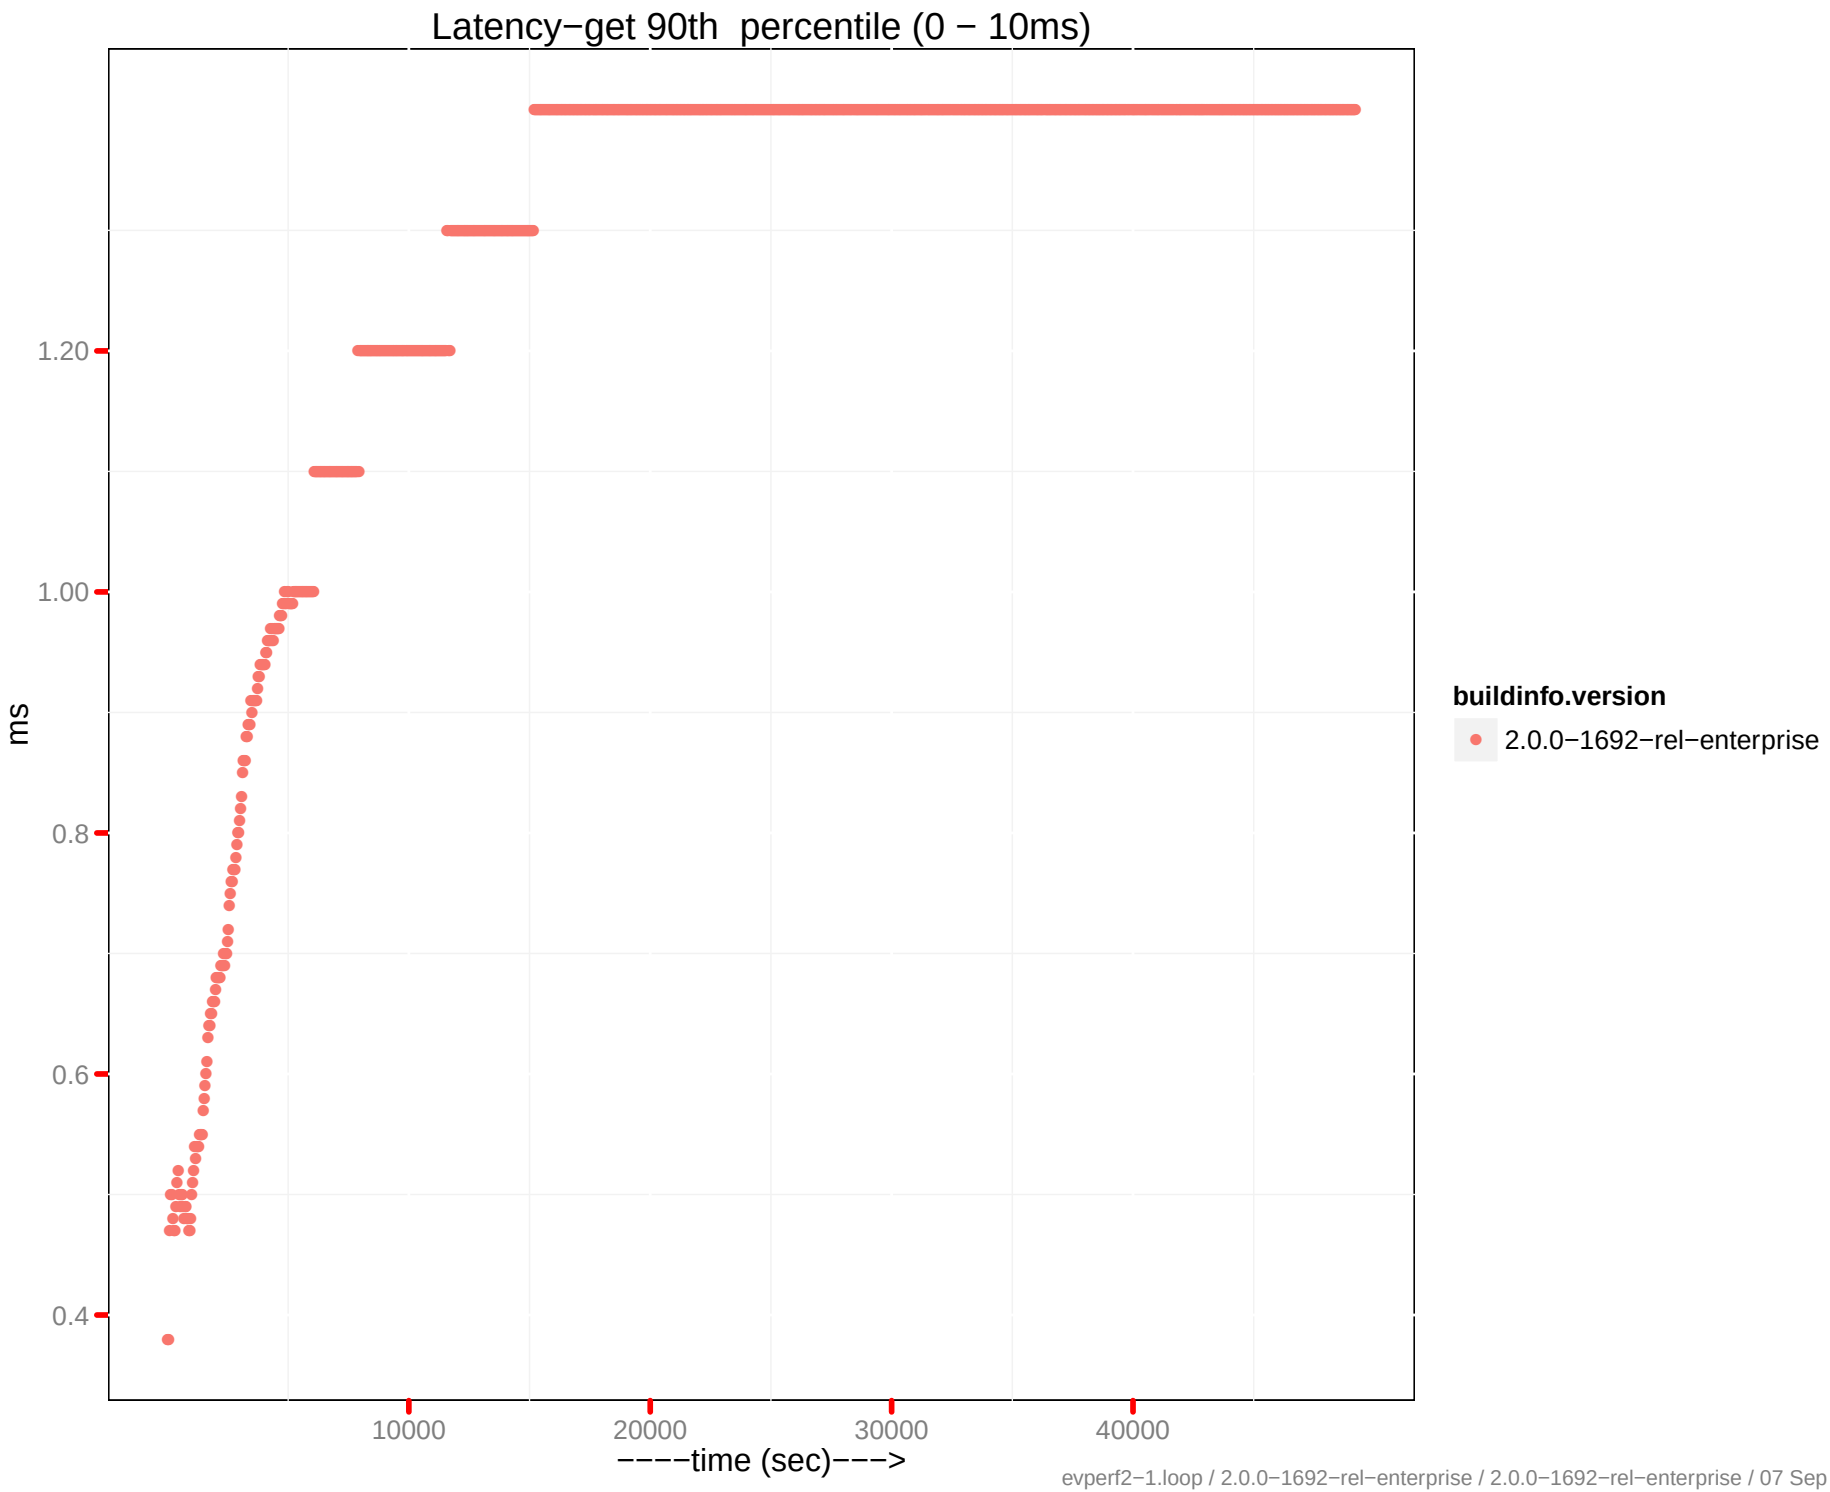
Latency−get 90th percentile (0 − 10ms)
1.20
1.00
0.8
0.6
0.4
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

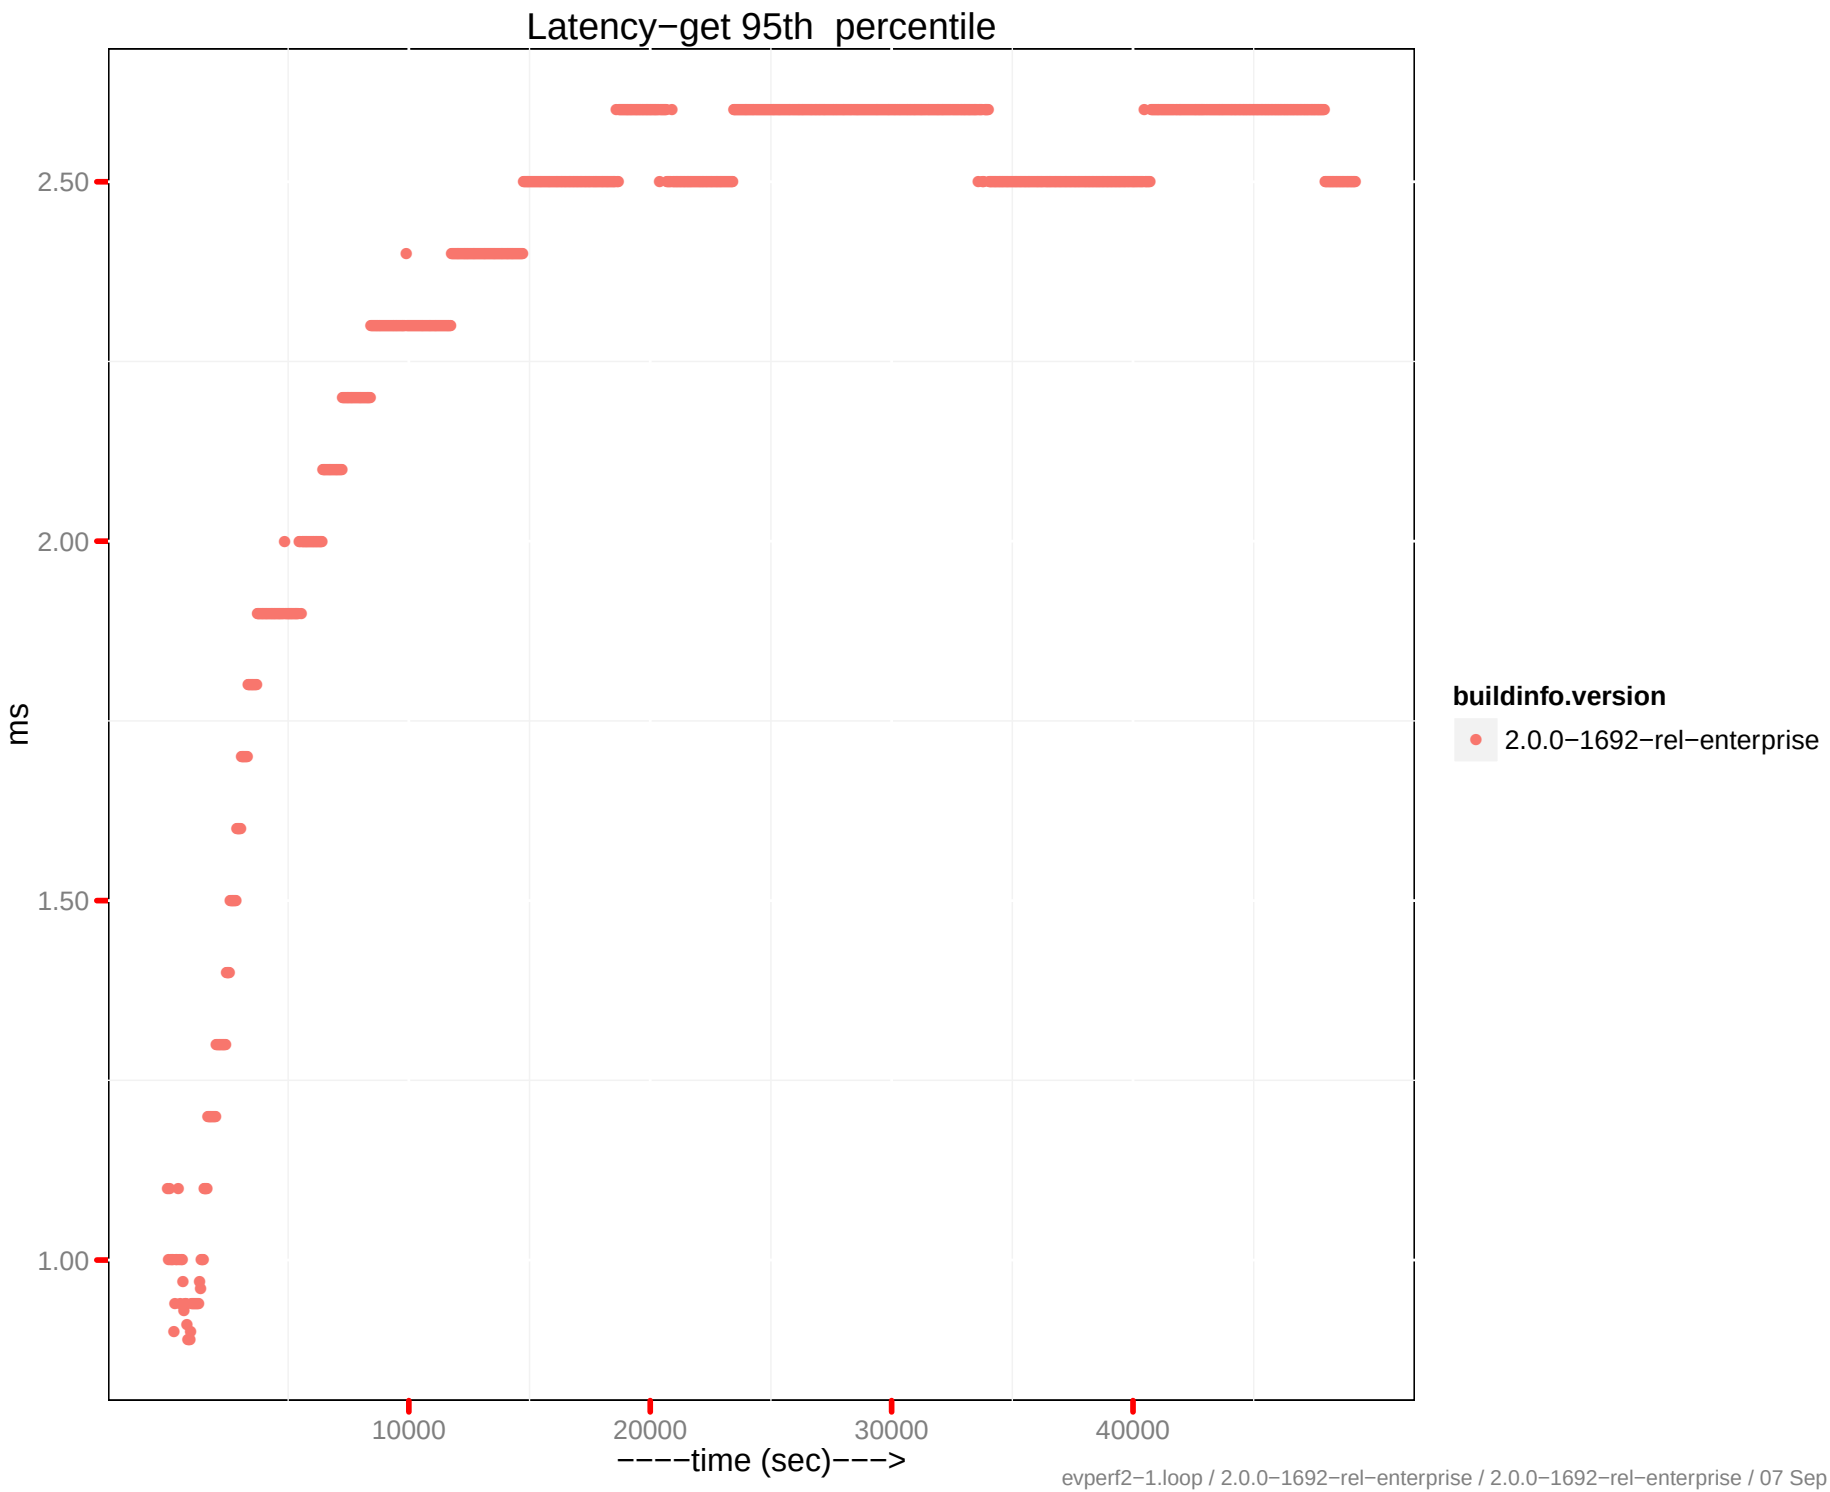
Latency−get 95th percentile
2.50
2.00
1.50
1.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

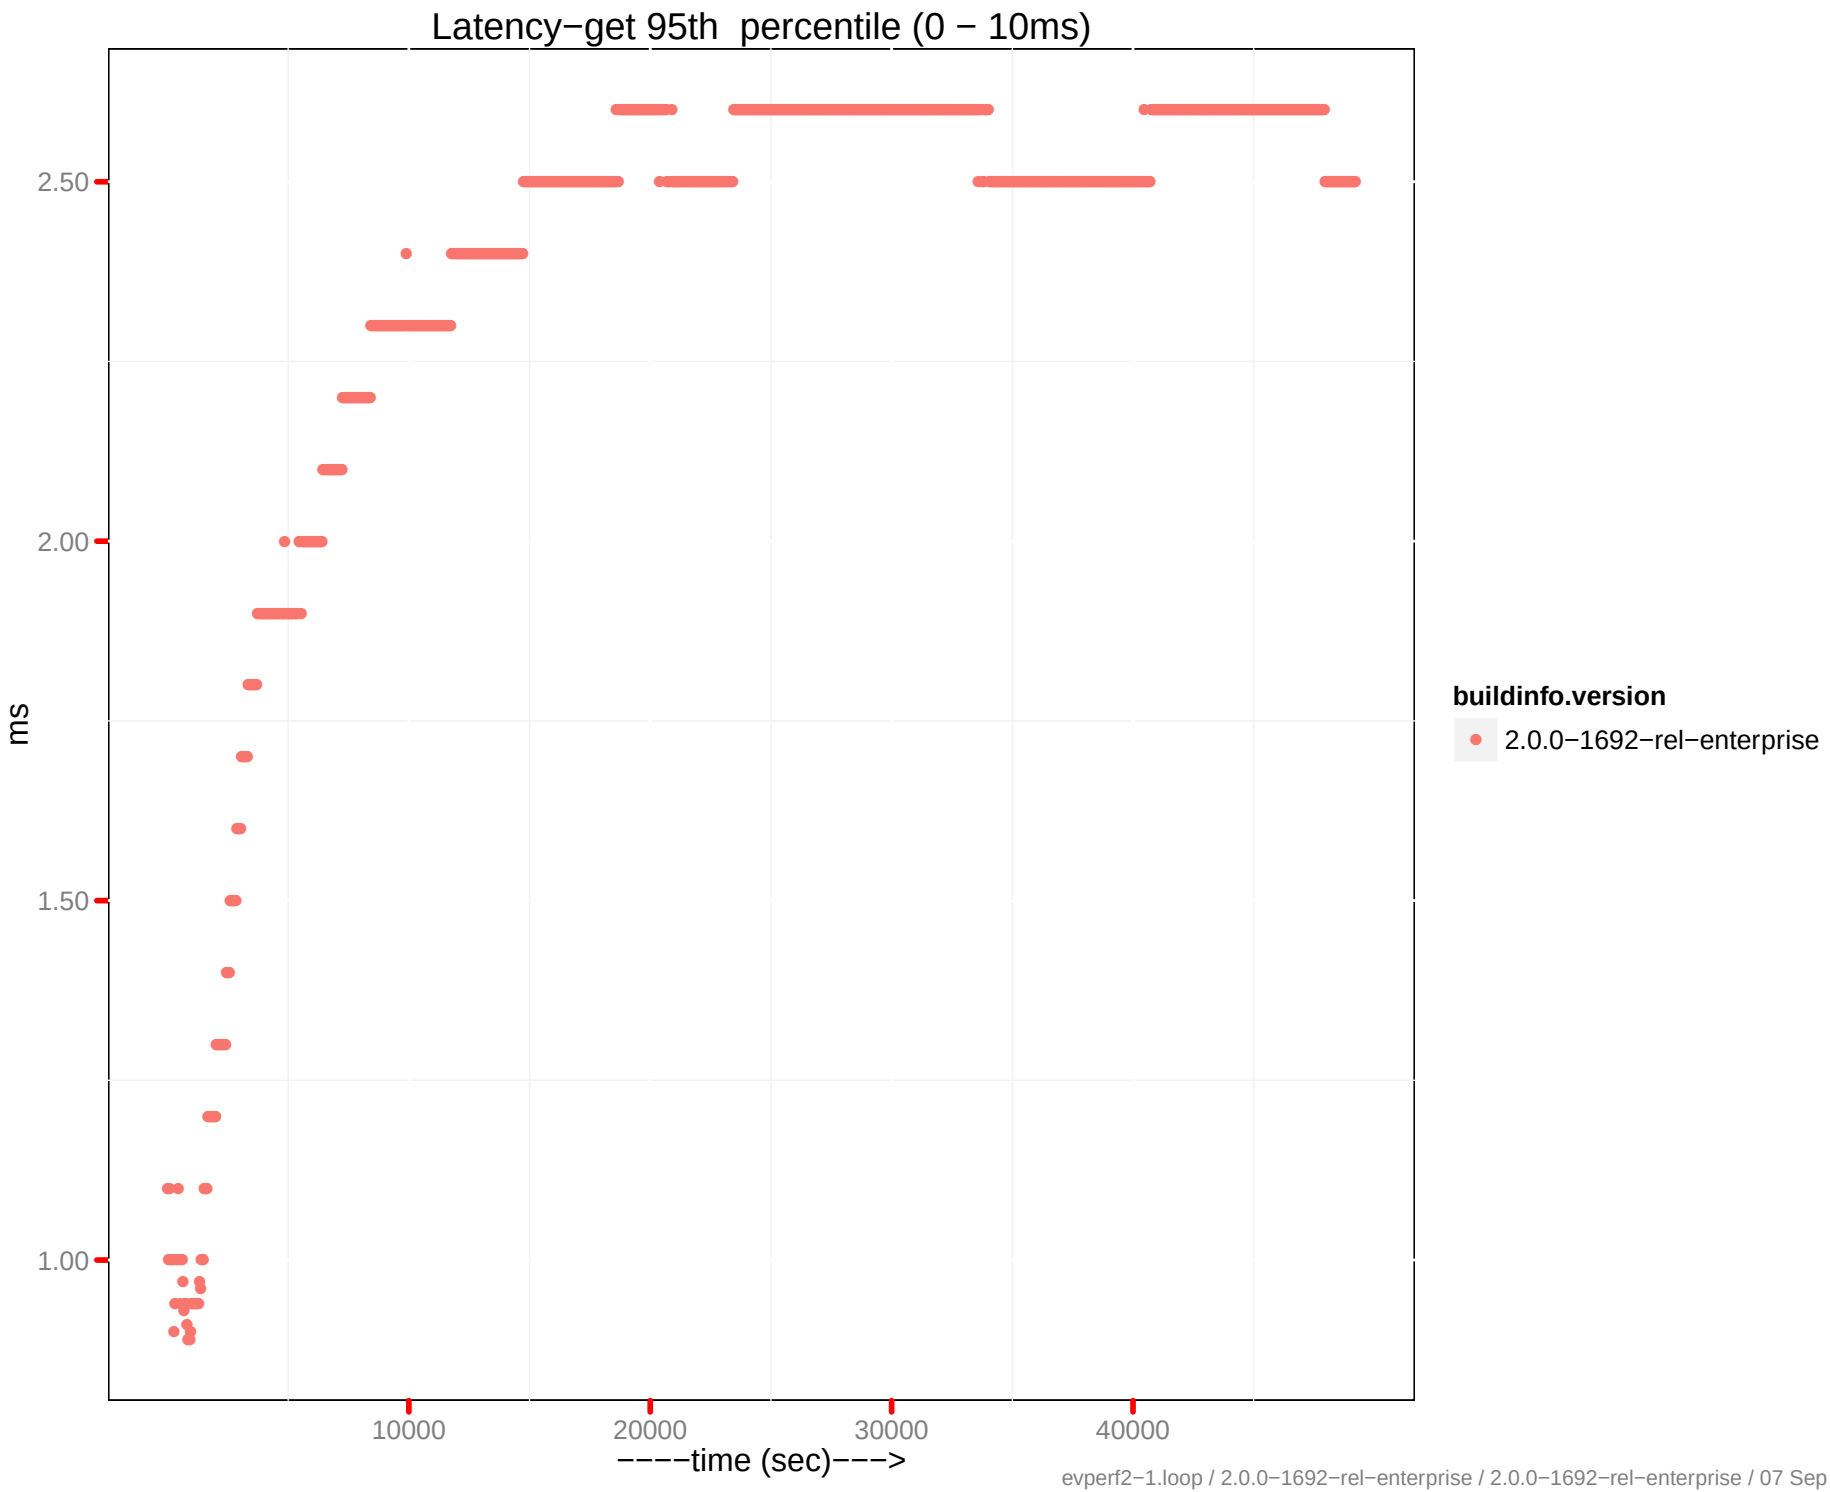
Latency−get 95th percentile (0 − 10ms)
2.50
2.00
1.50
1.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

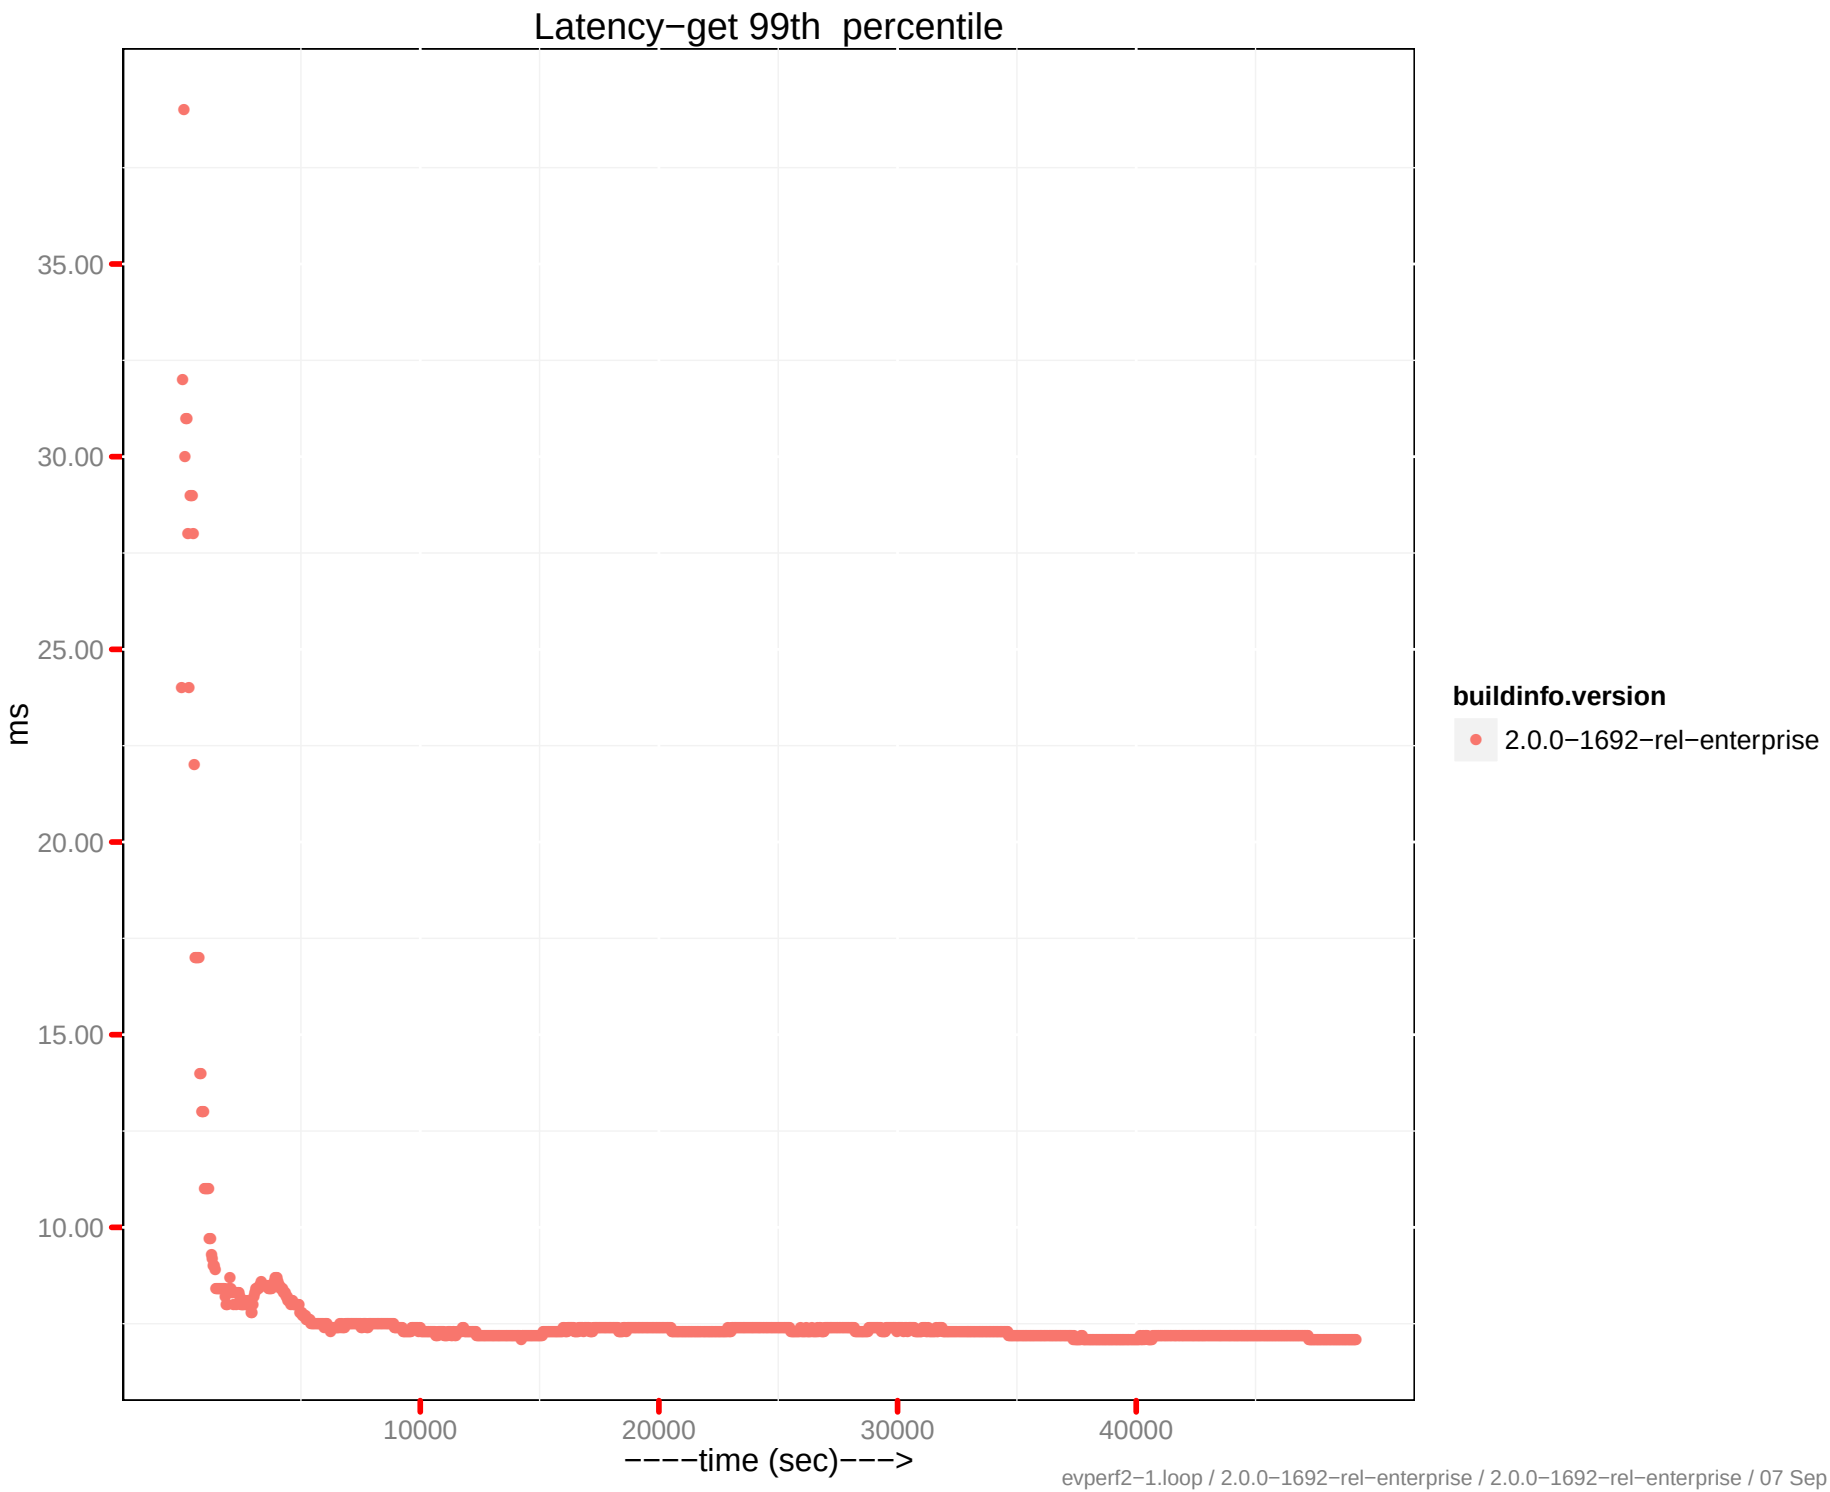
Latency−get 99th percentile
35.00
30.00
25.00
20.00
15.00
10.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

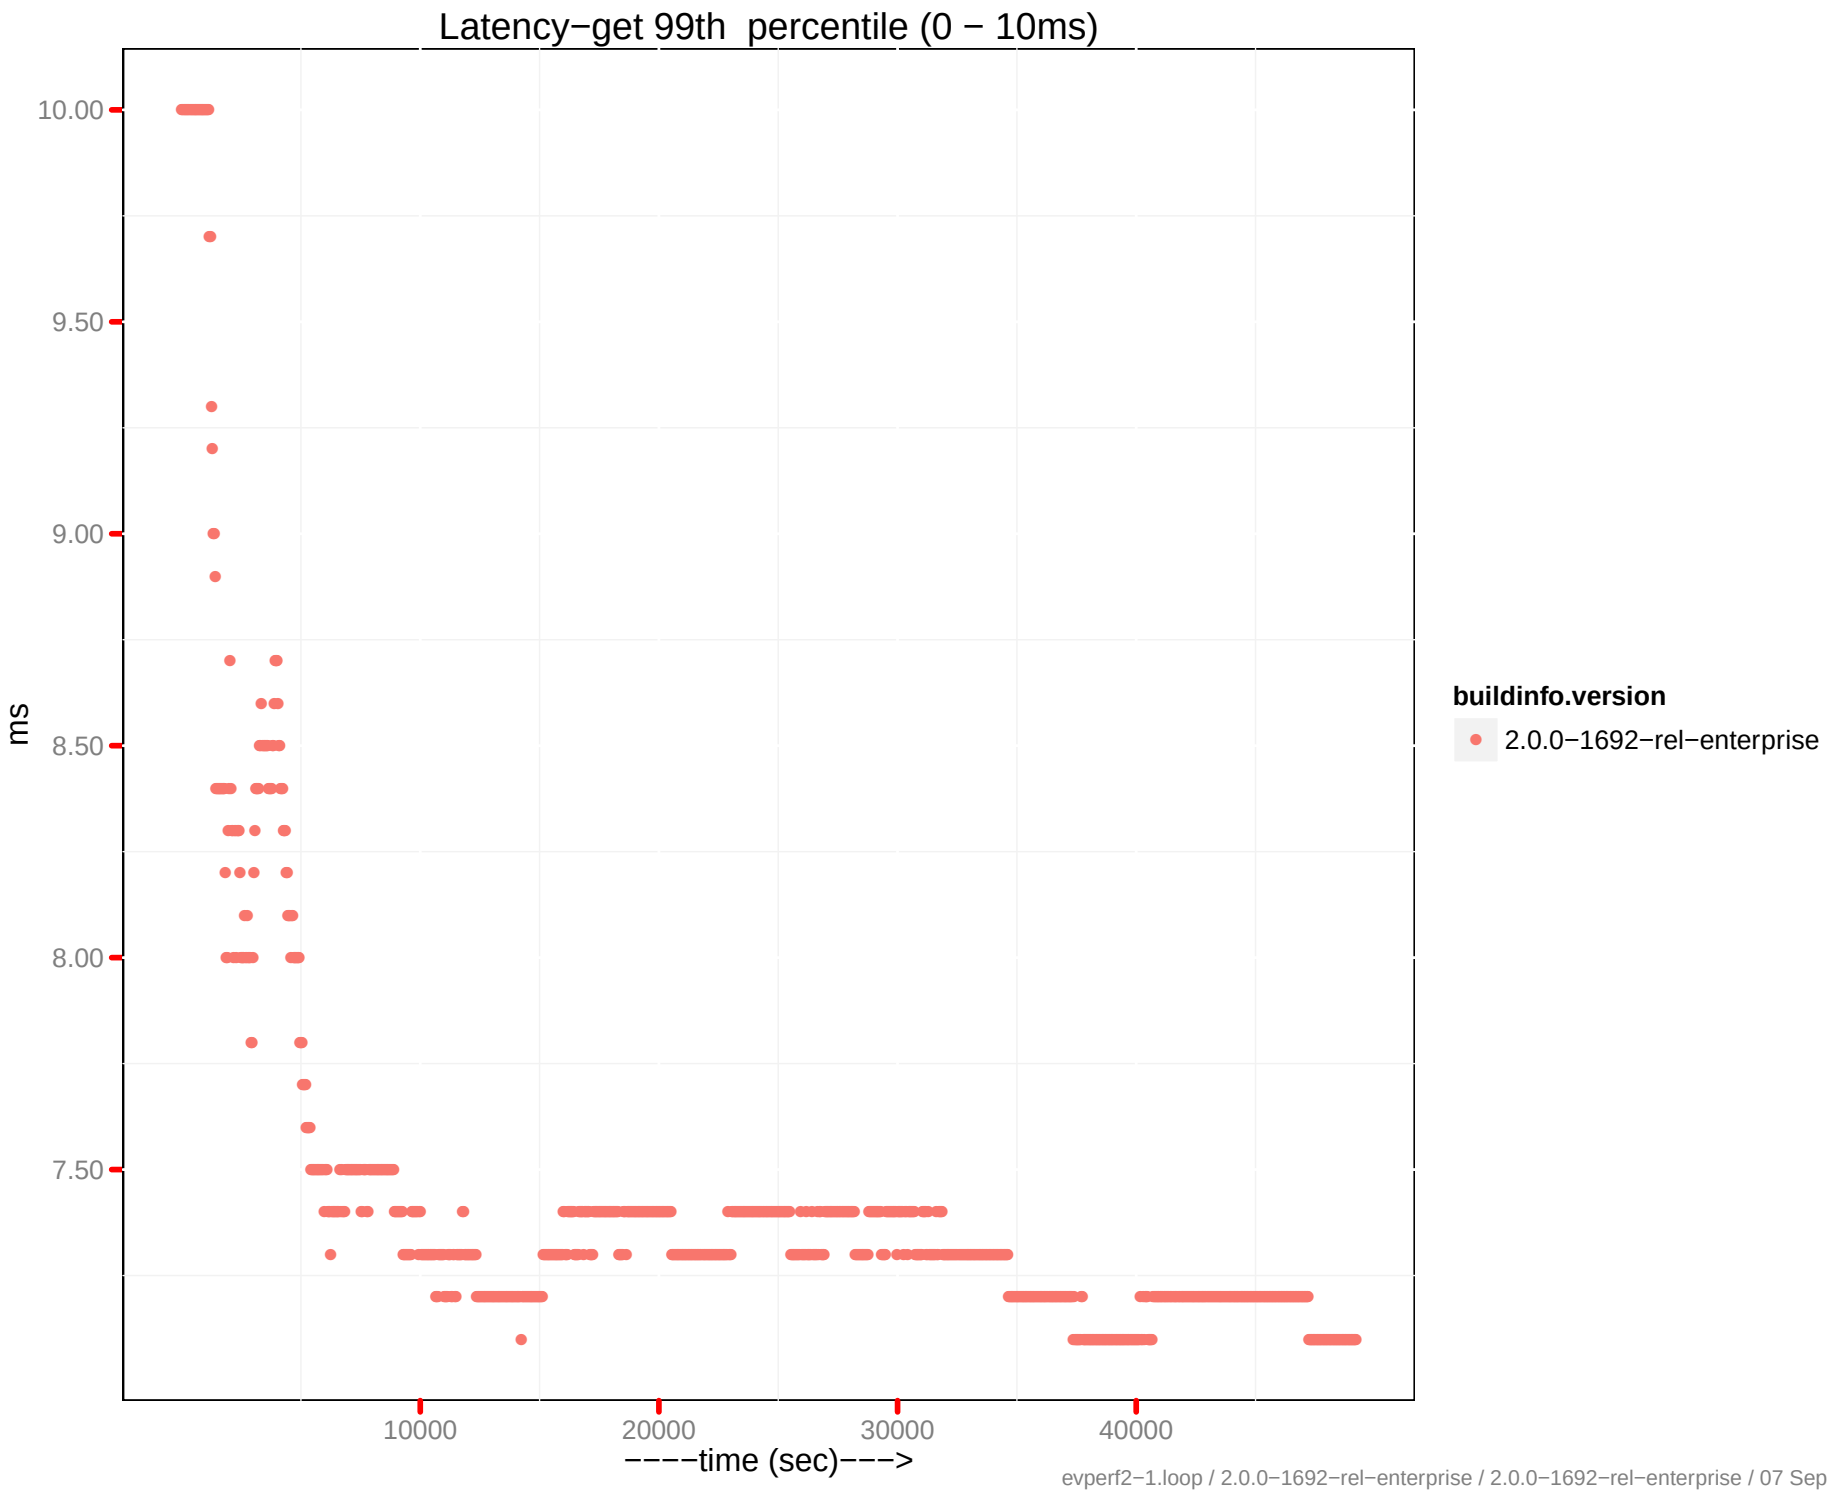
Latency−get 99th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
8.00
7.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

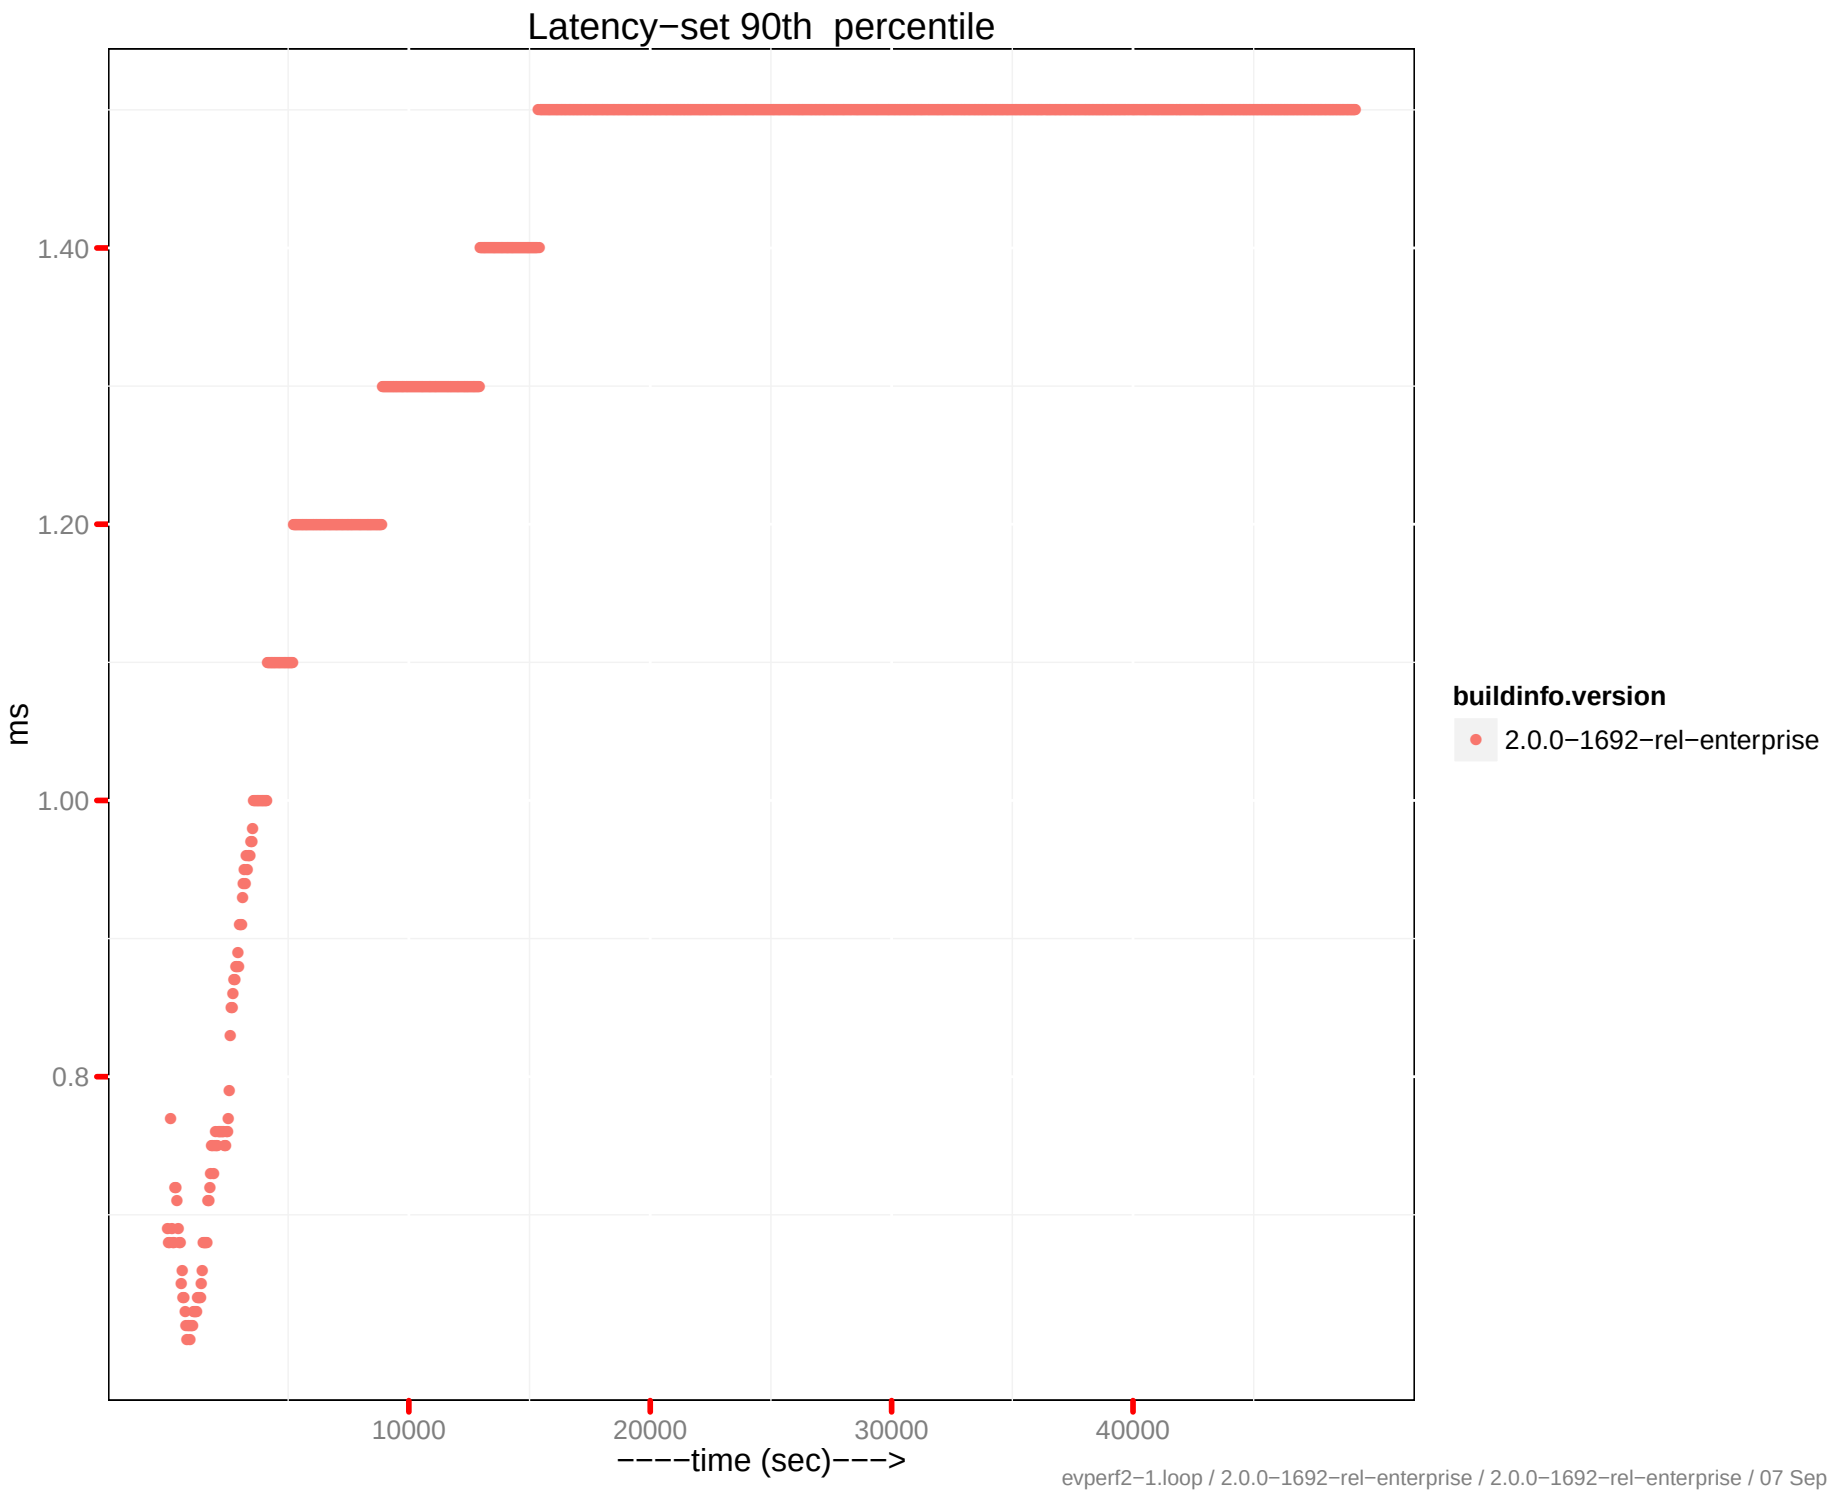
Latency−set 90th percentile
1.40
1.20
1.00
0.8
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

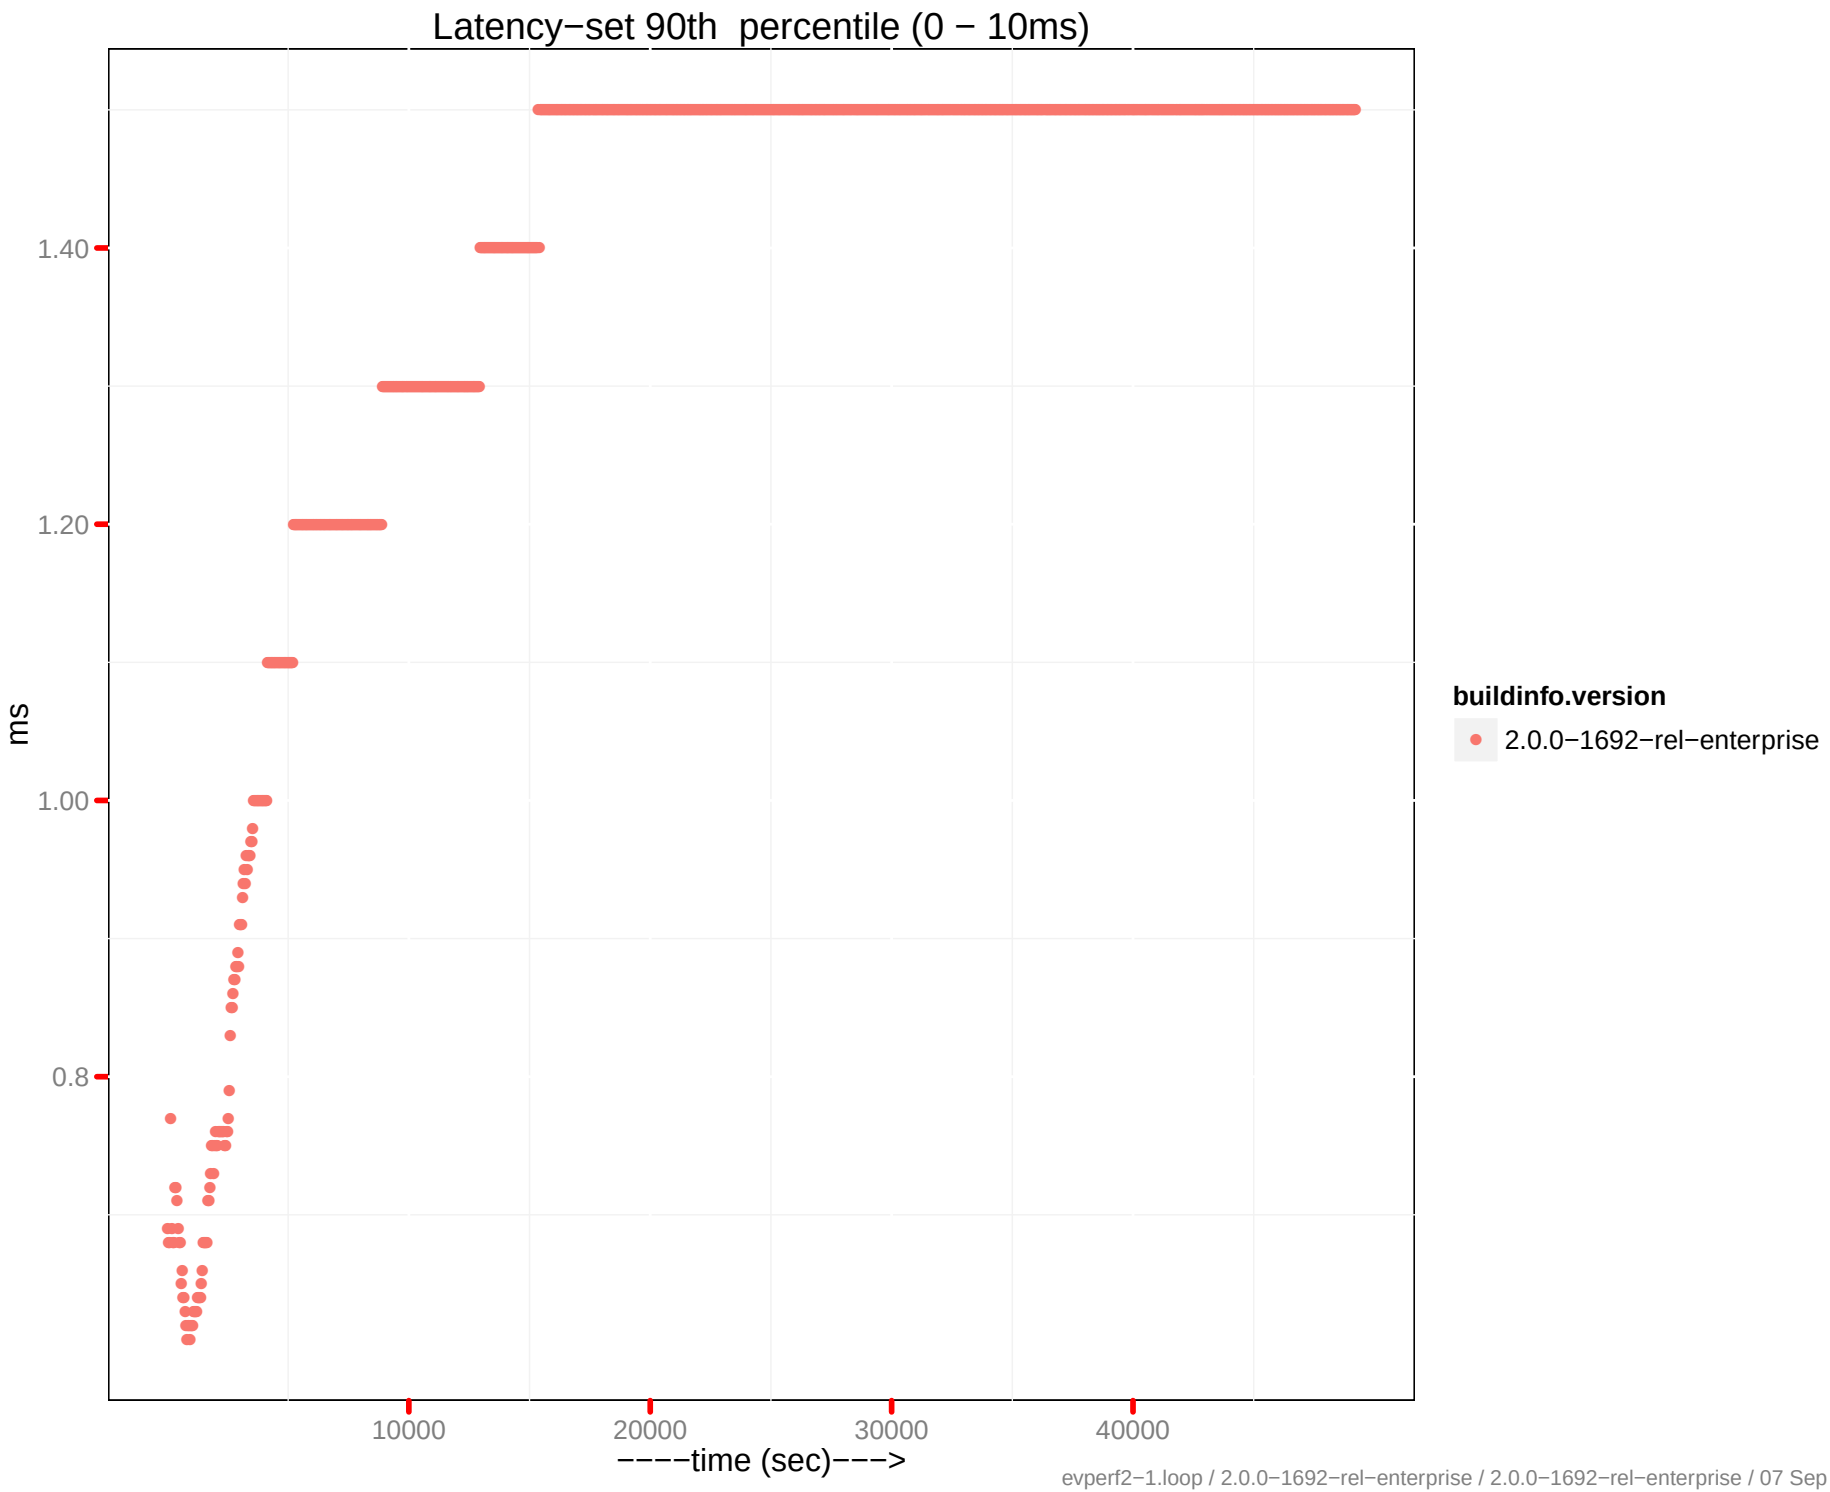
Latency−set 90th percentile (0 − 10ms)
1.40
1.20
1.00
0.8
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

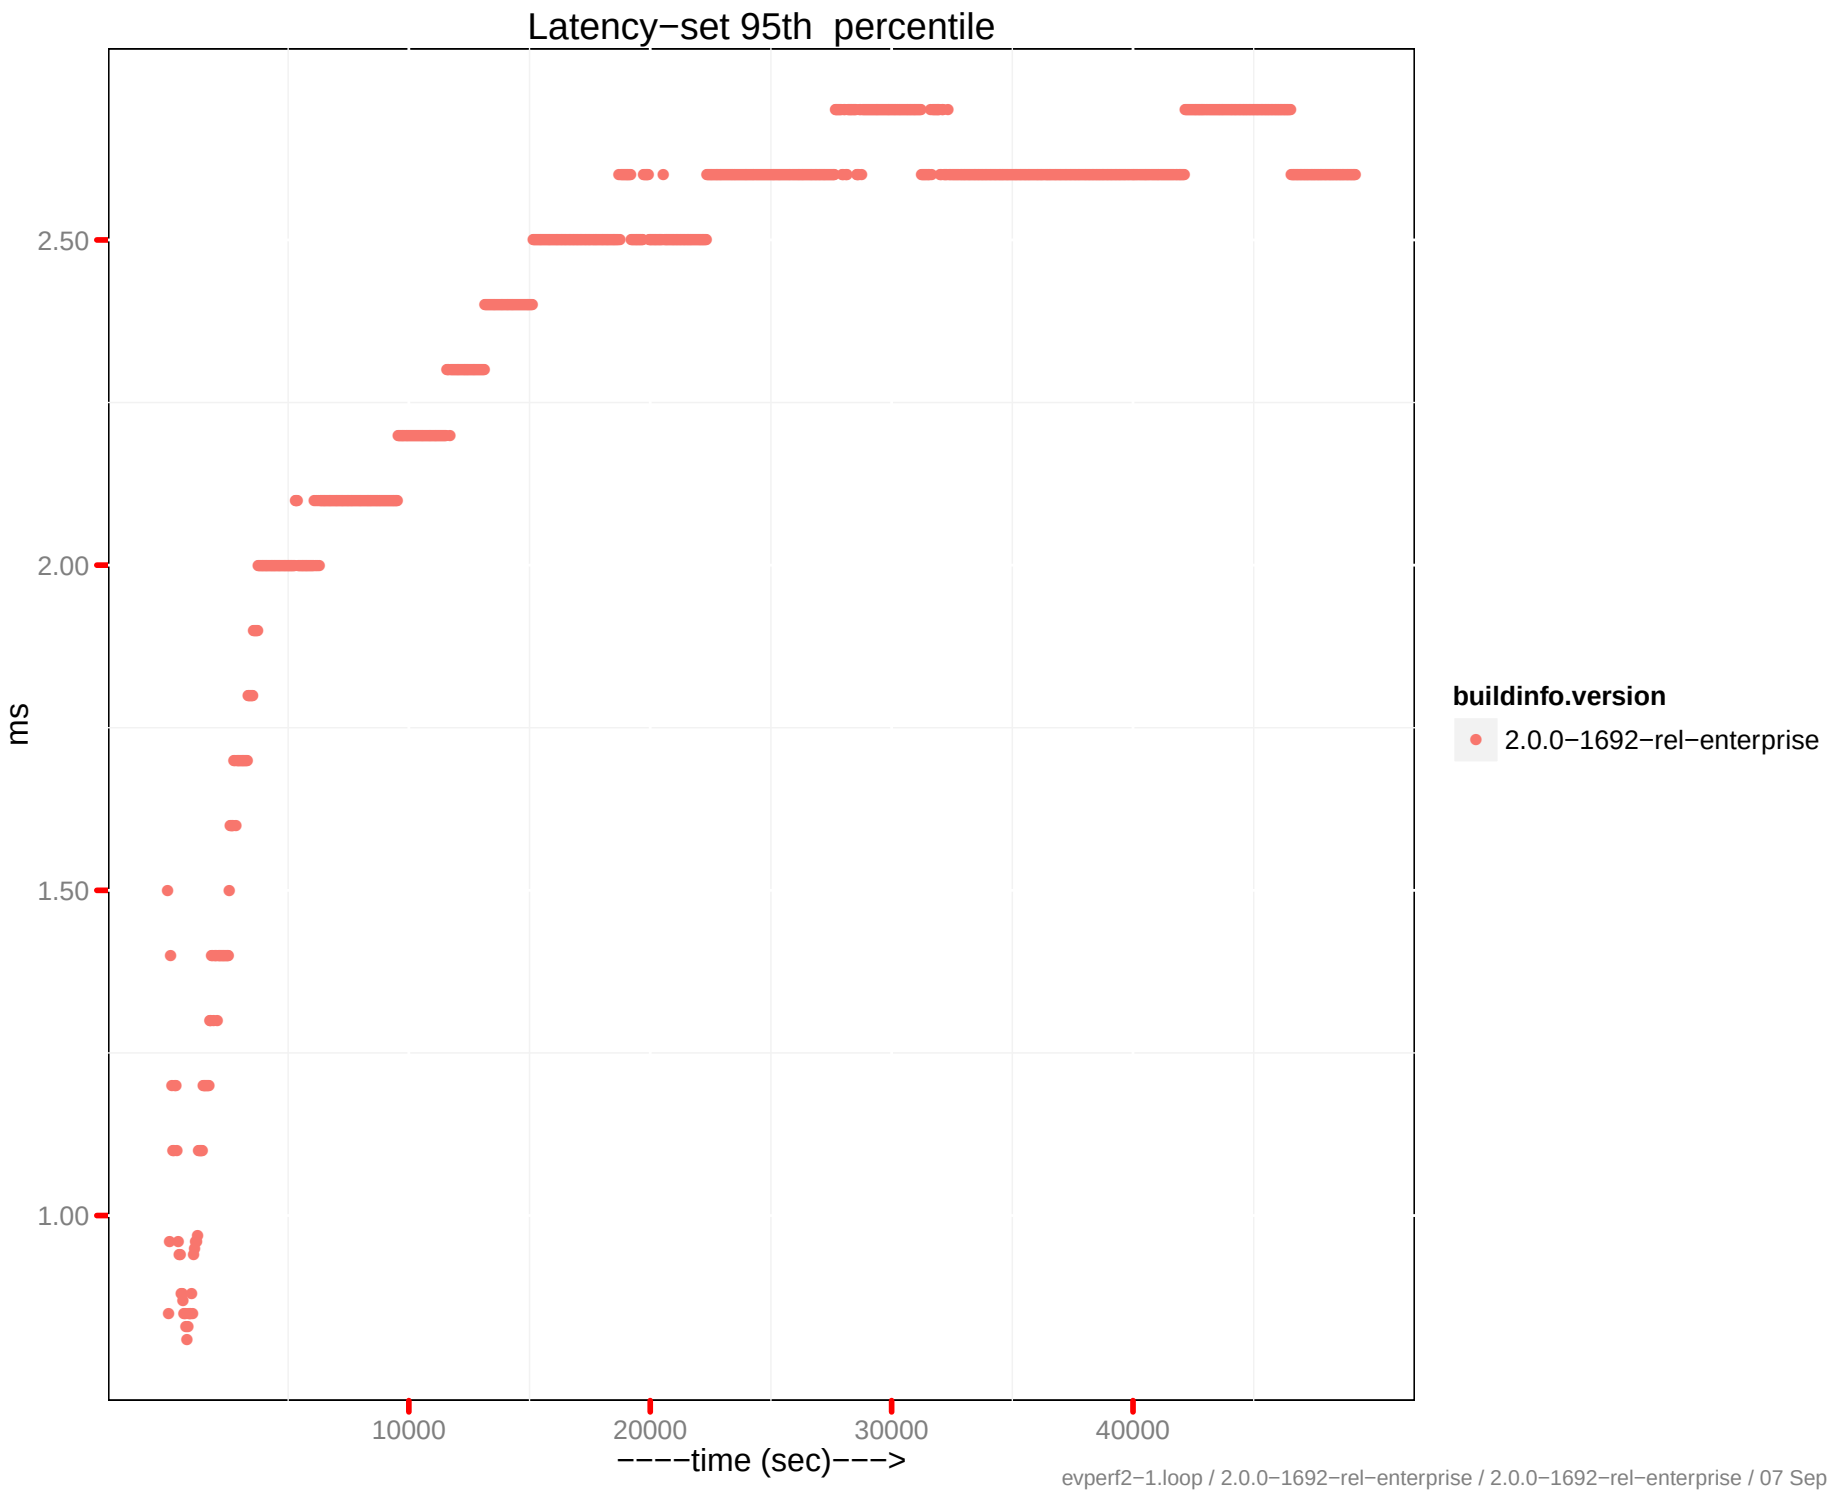
Latency−set 95th percentile
2.50
2.00
1.50
1.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

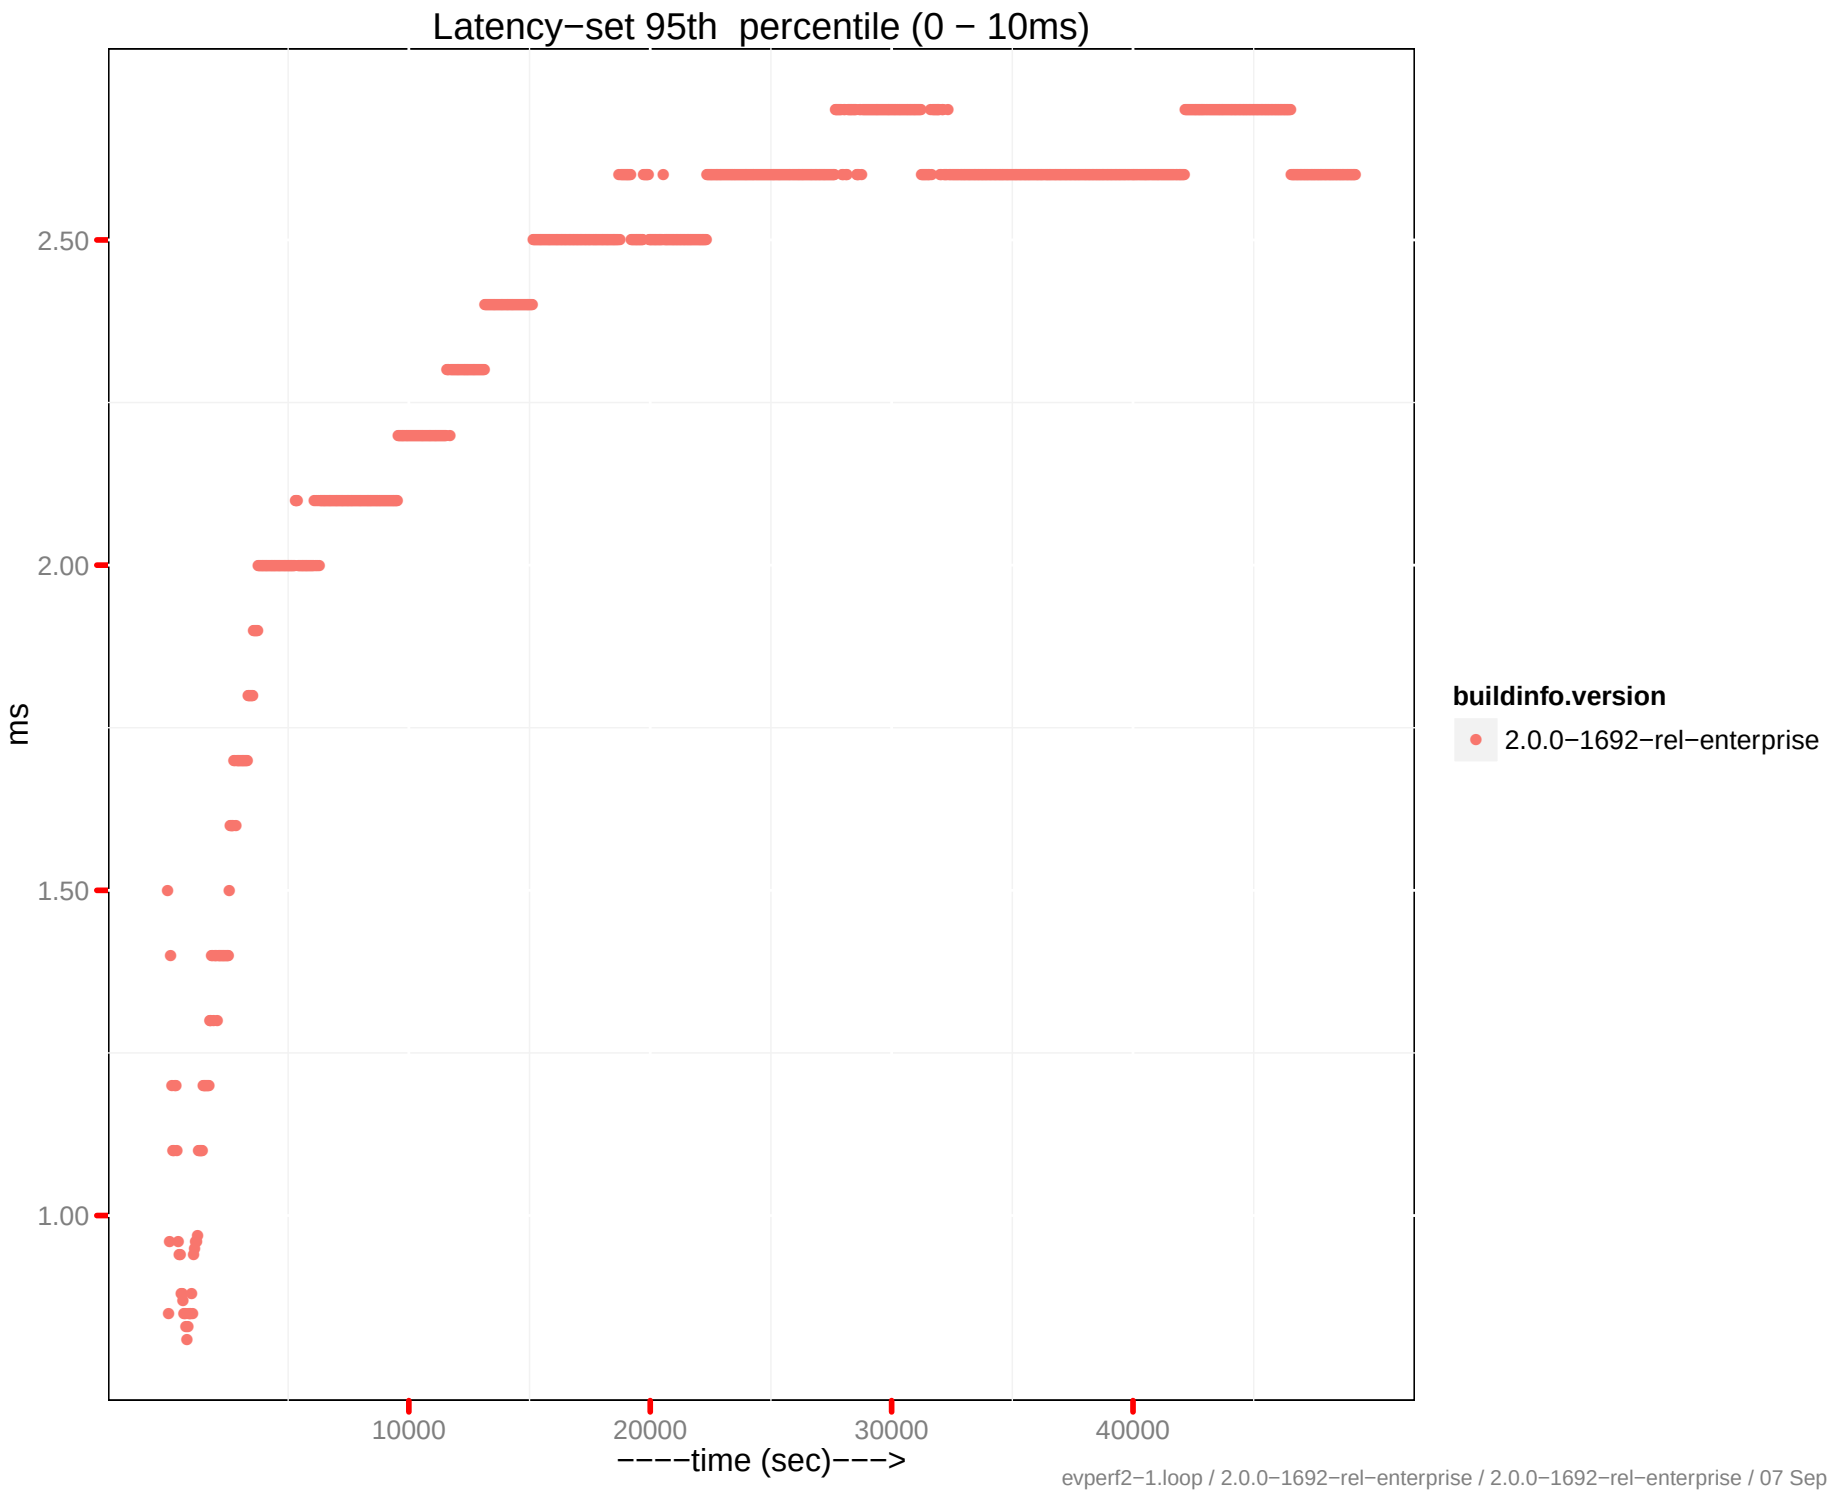
Latency−set 95th percentile (0 − 10ms)
2.50
2.00
1.50
1.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

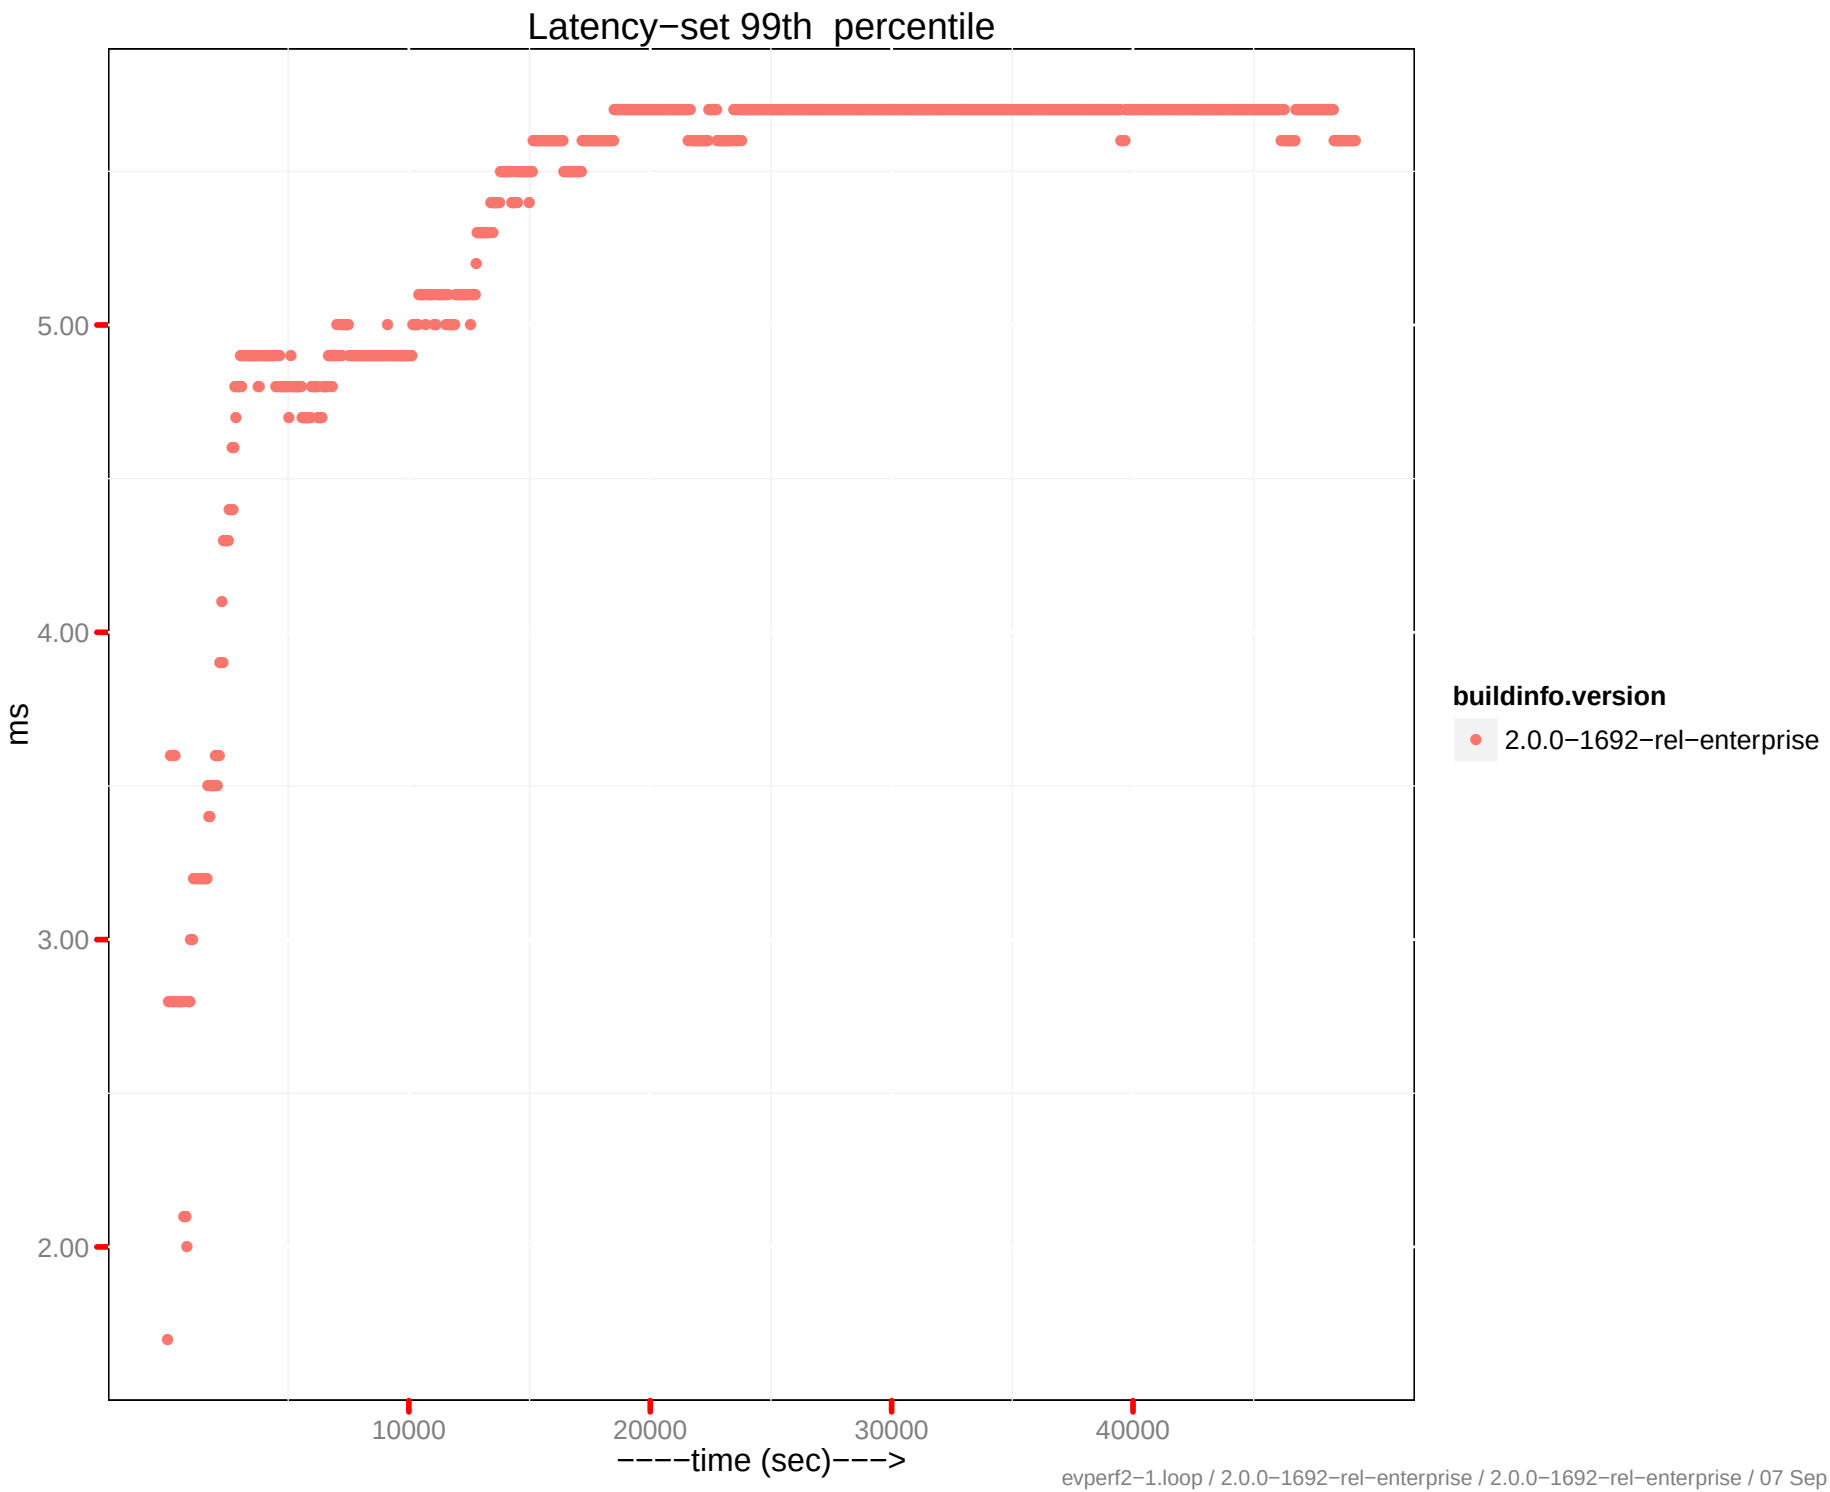
Latency−set 99th percentile
5.00
4.00
3.00
2.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

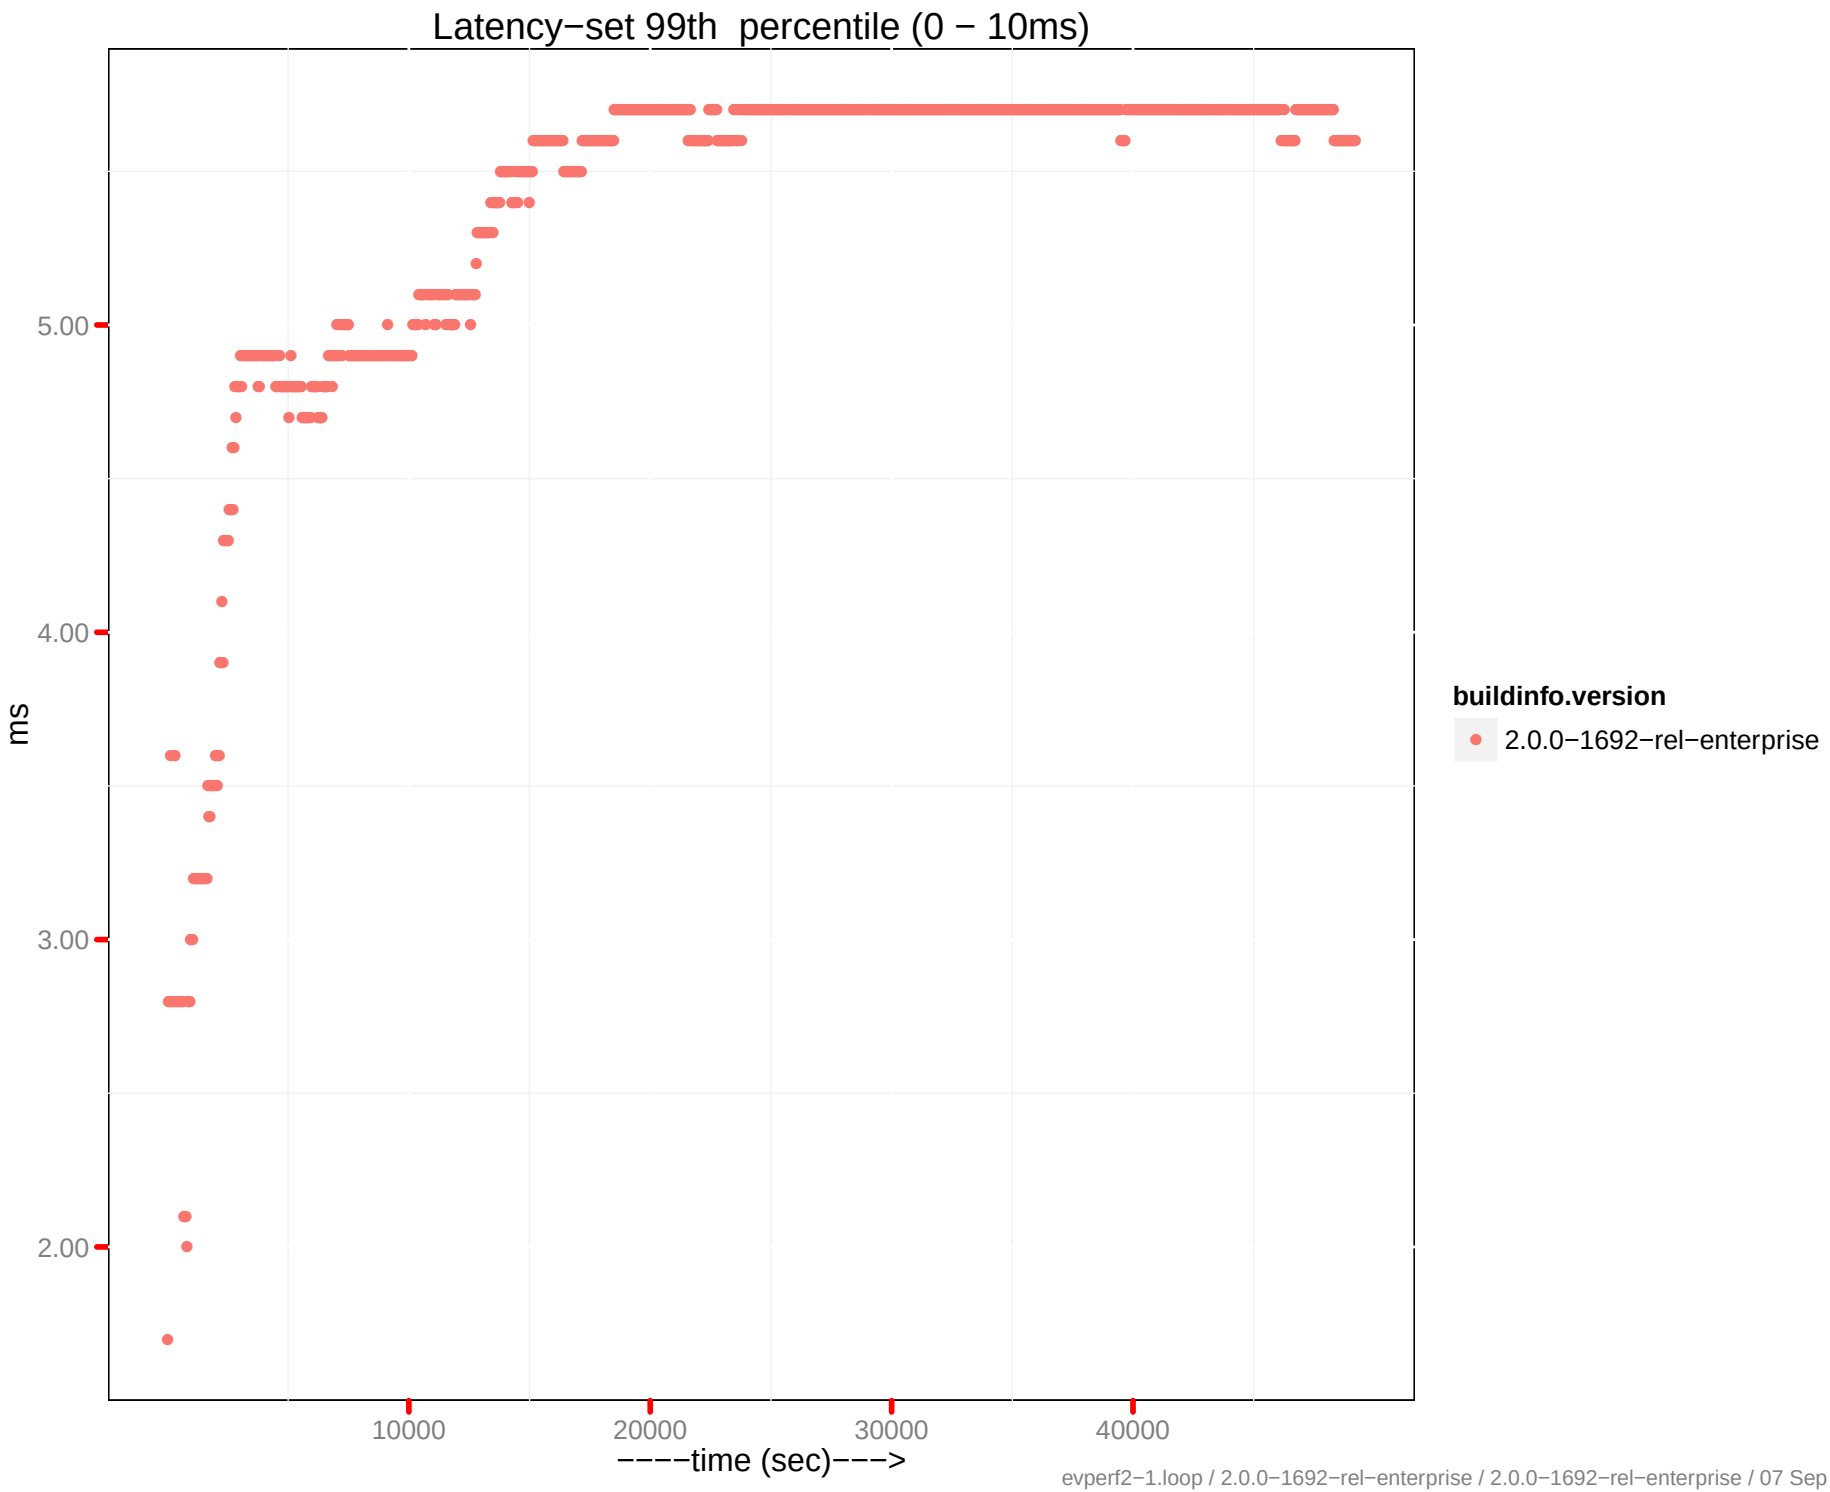
Latency−set 99th percentile (0 − 10ms)
5.00
4.00
3.00
2.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

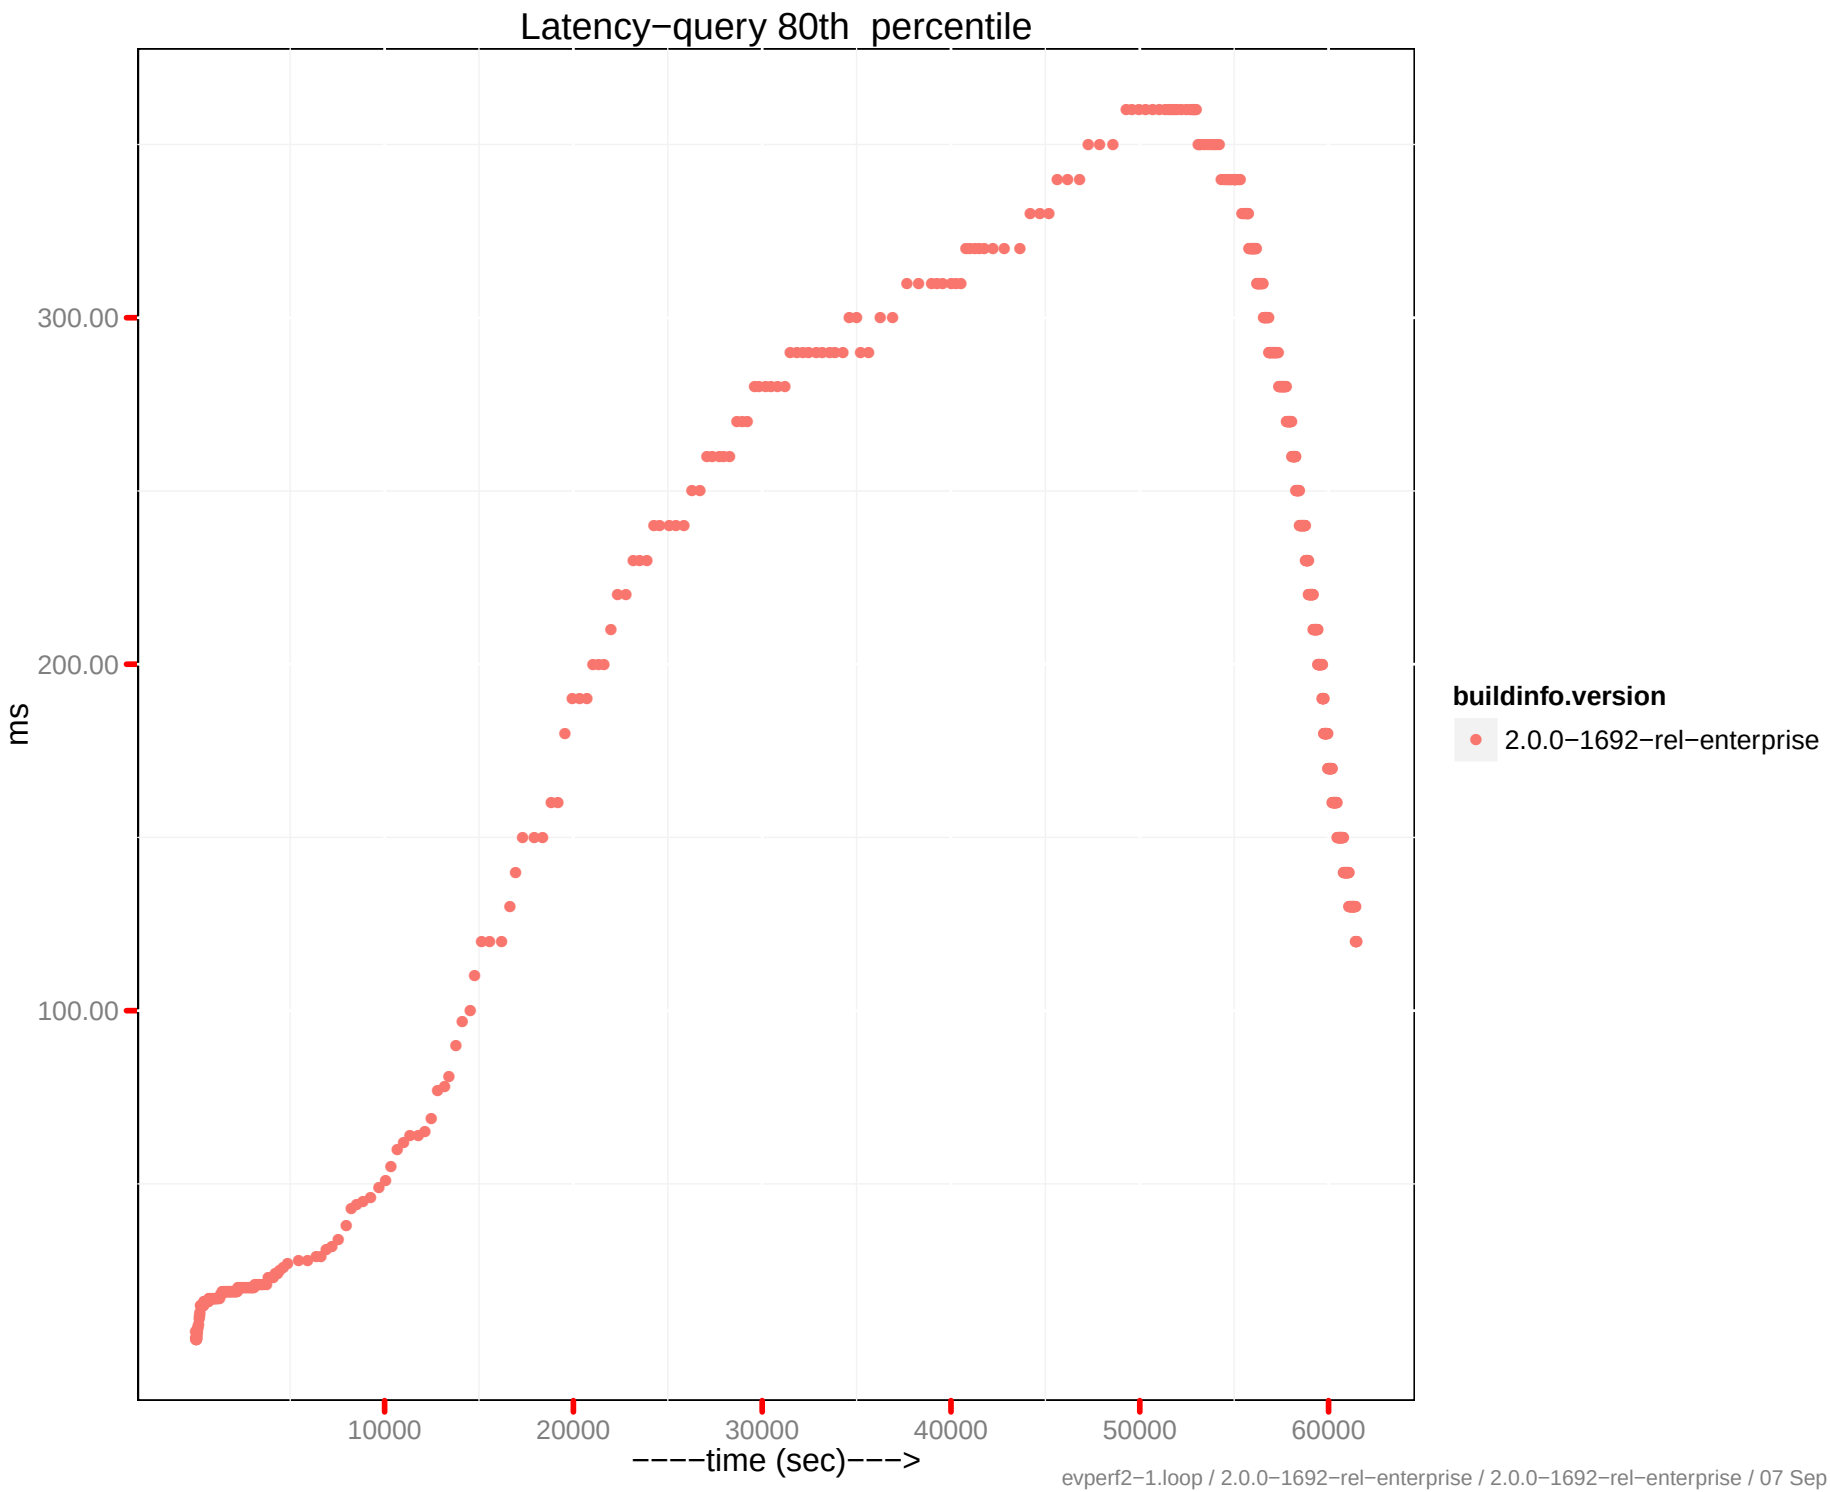
Latency−query 80th percentile
300.00
200.00
100.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

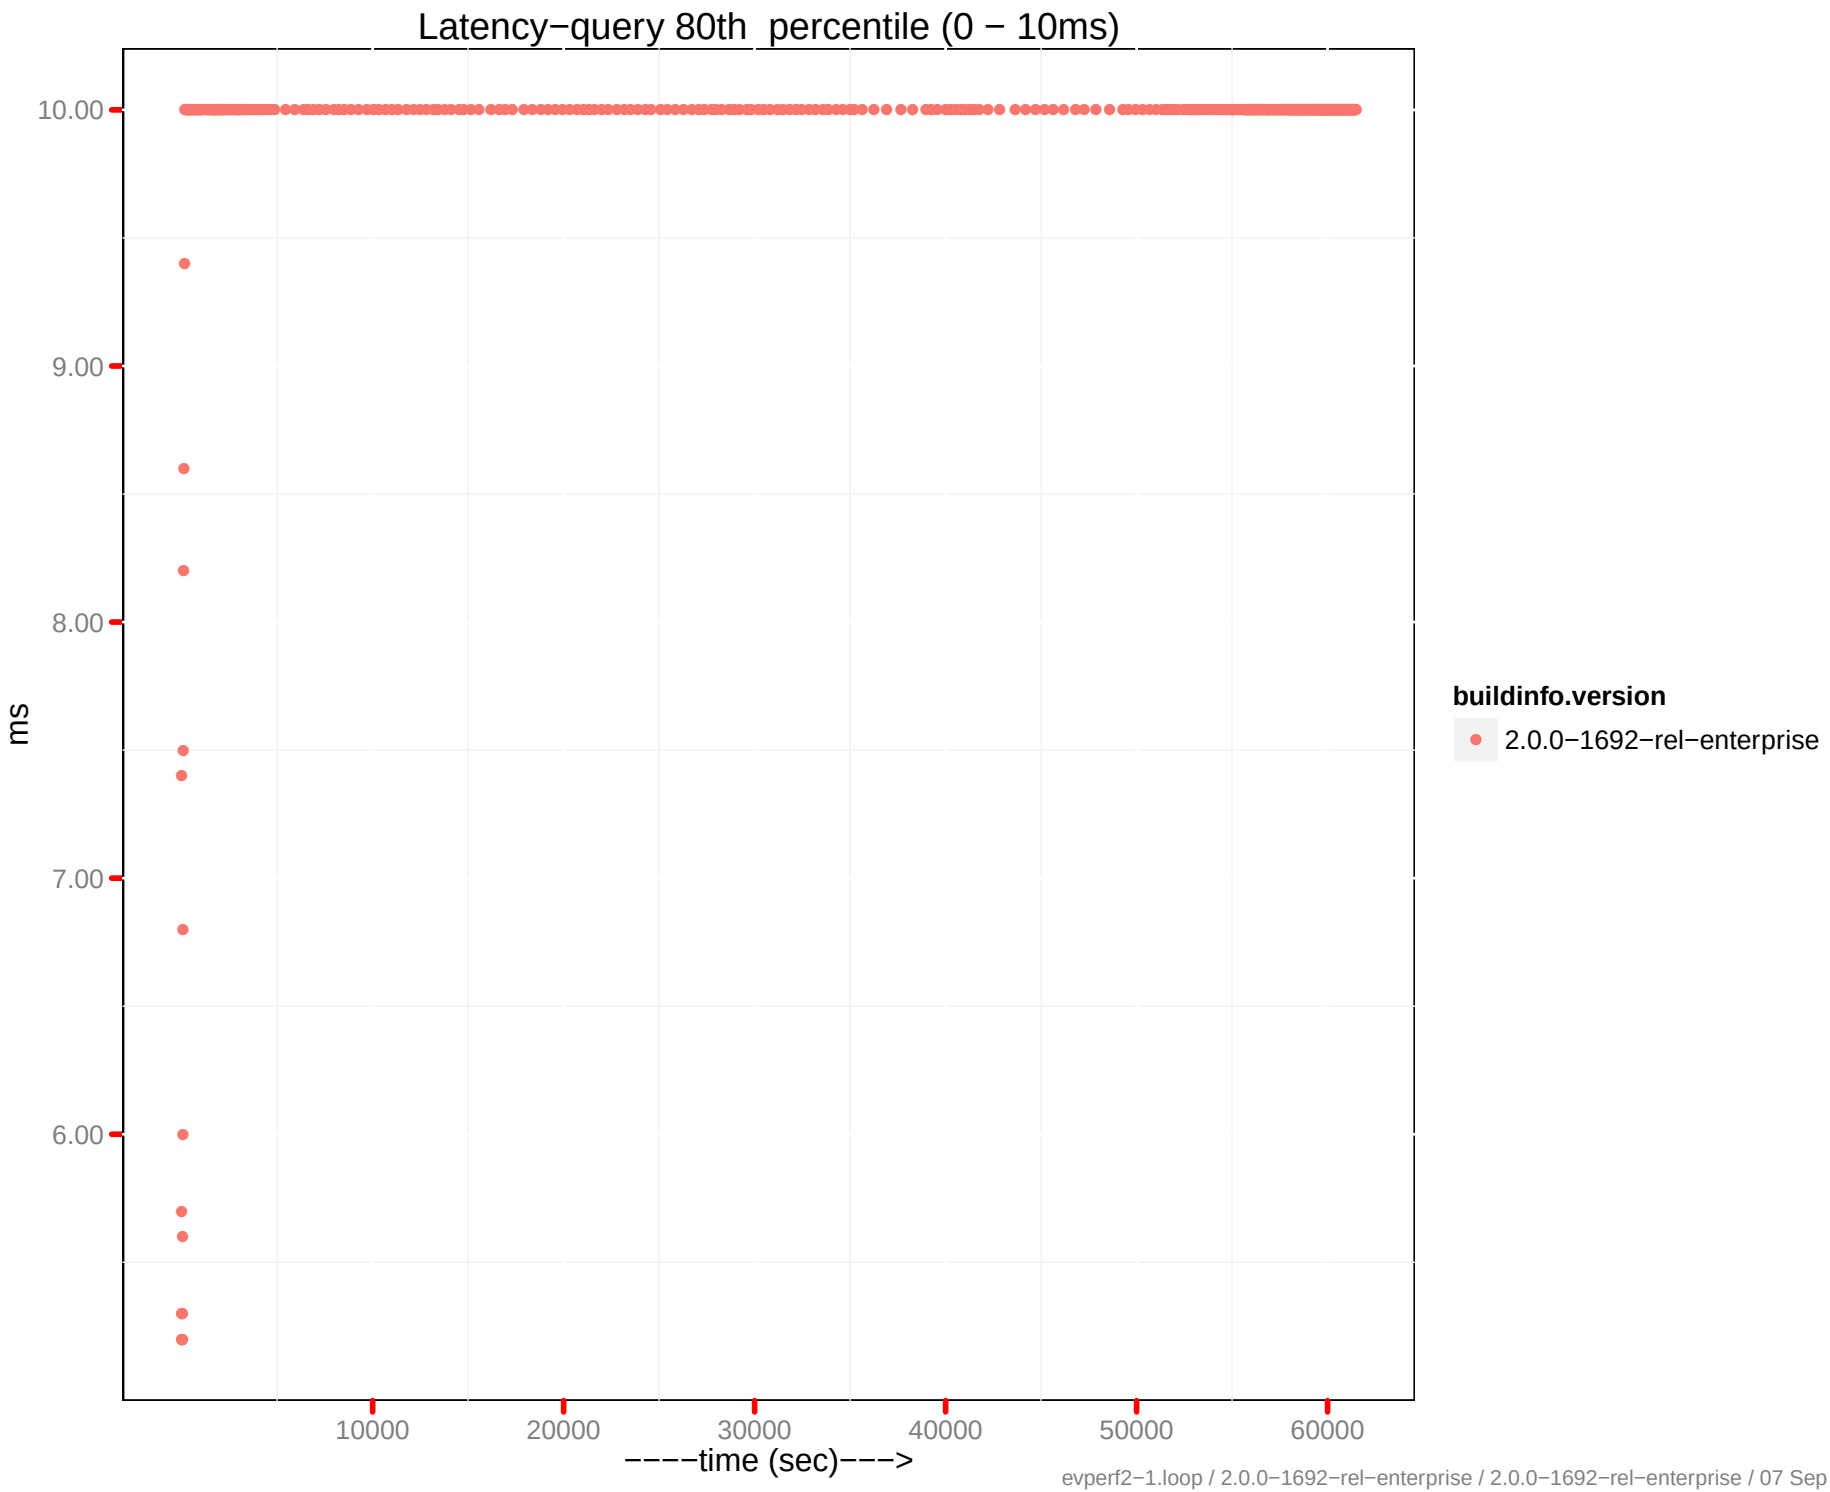
Latency−query 80th percentile (0 − 10ms)
10.00
9.00
8.00
7.00
6.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

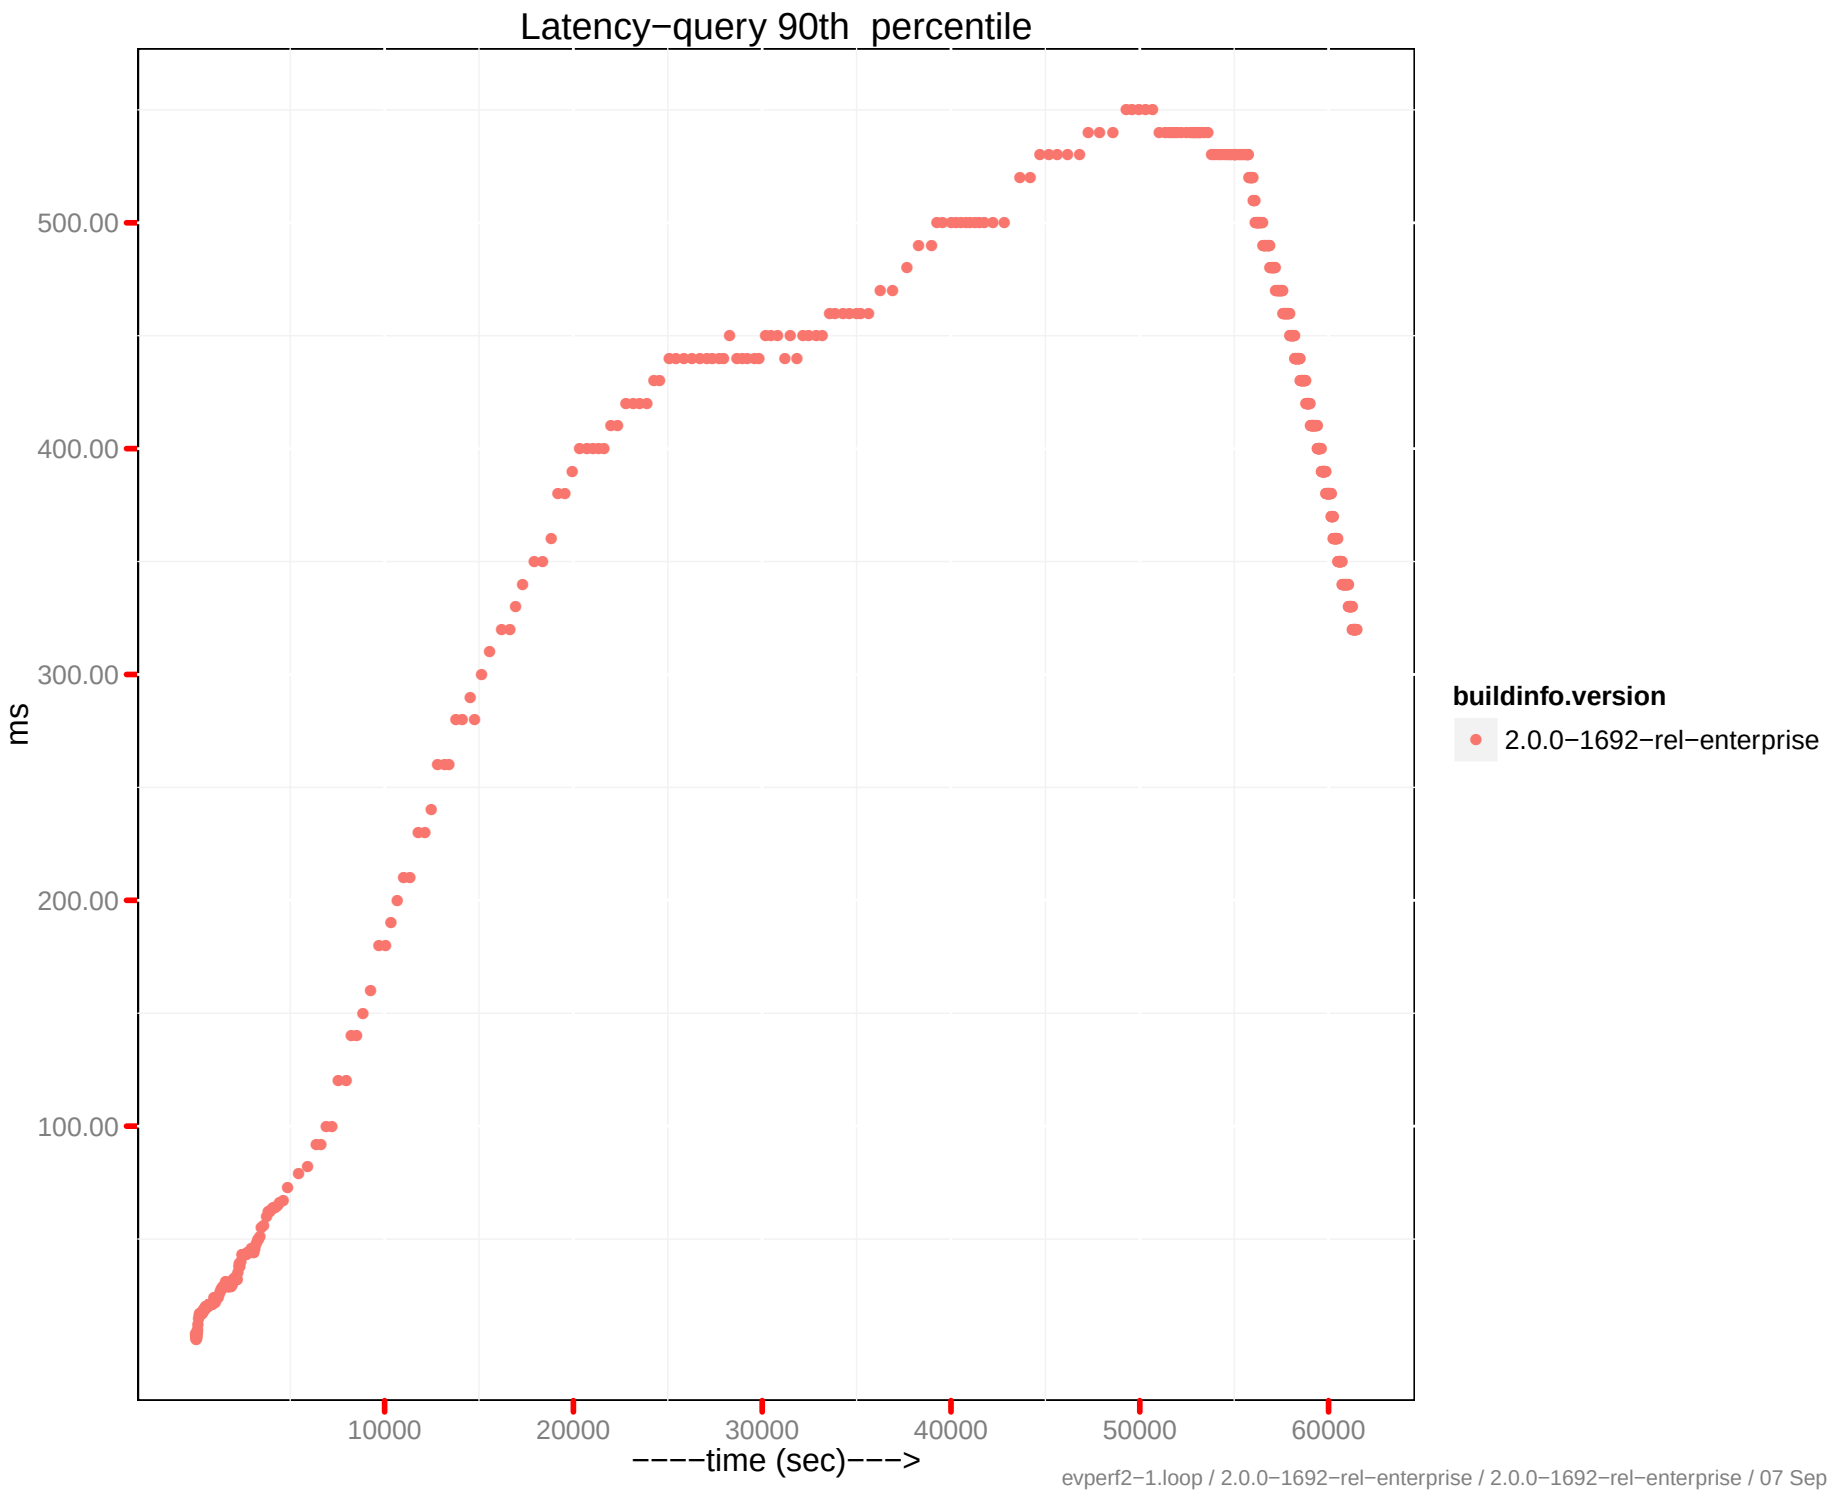
Latency−query 90th percentile
500.00
400.00
300.00
200.00
100.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

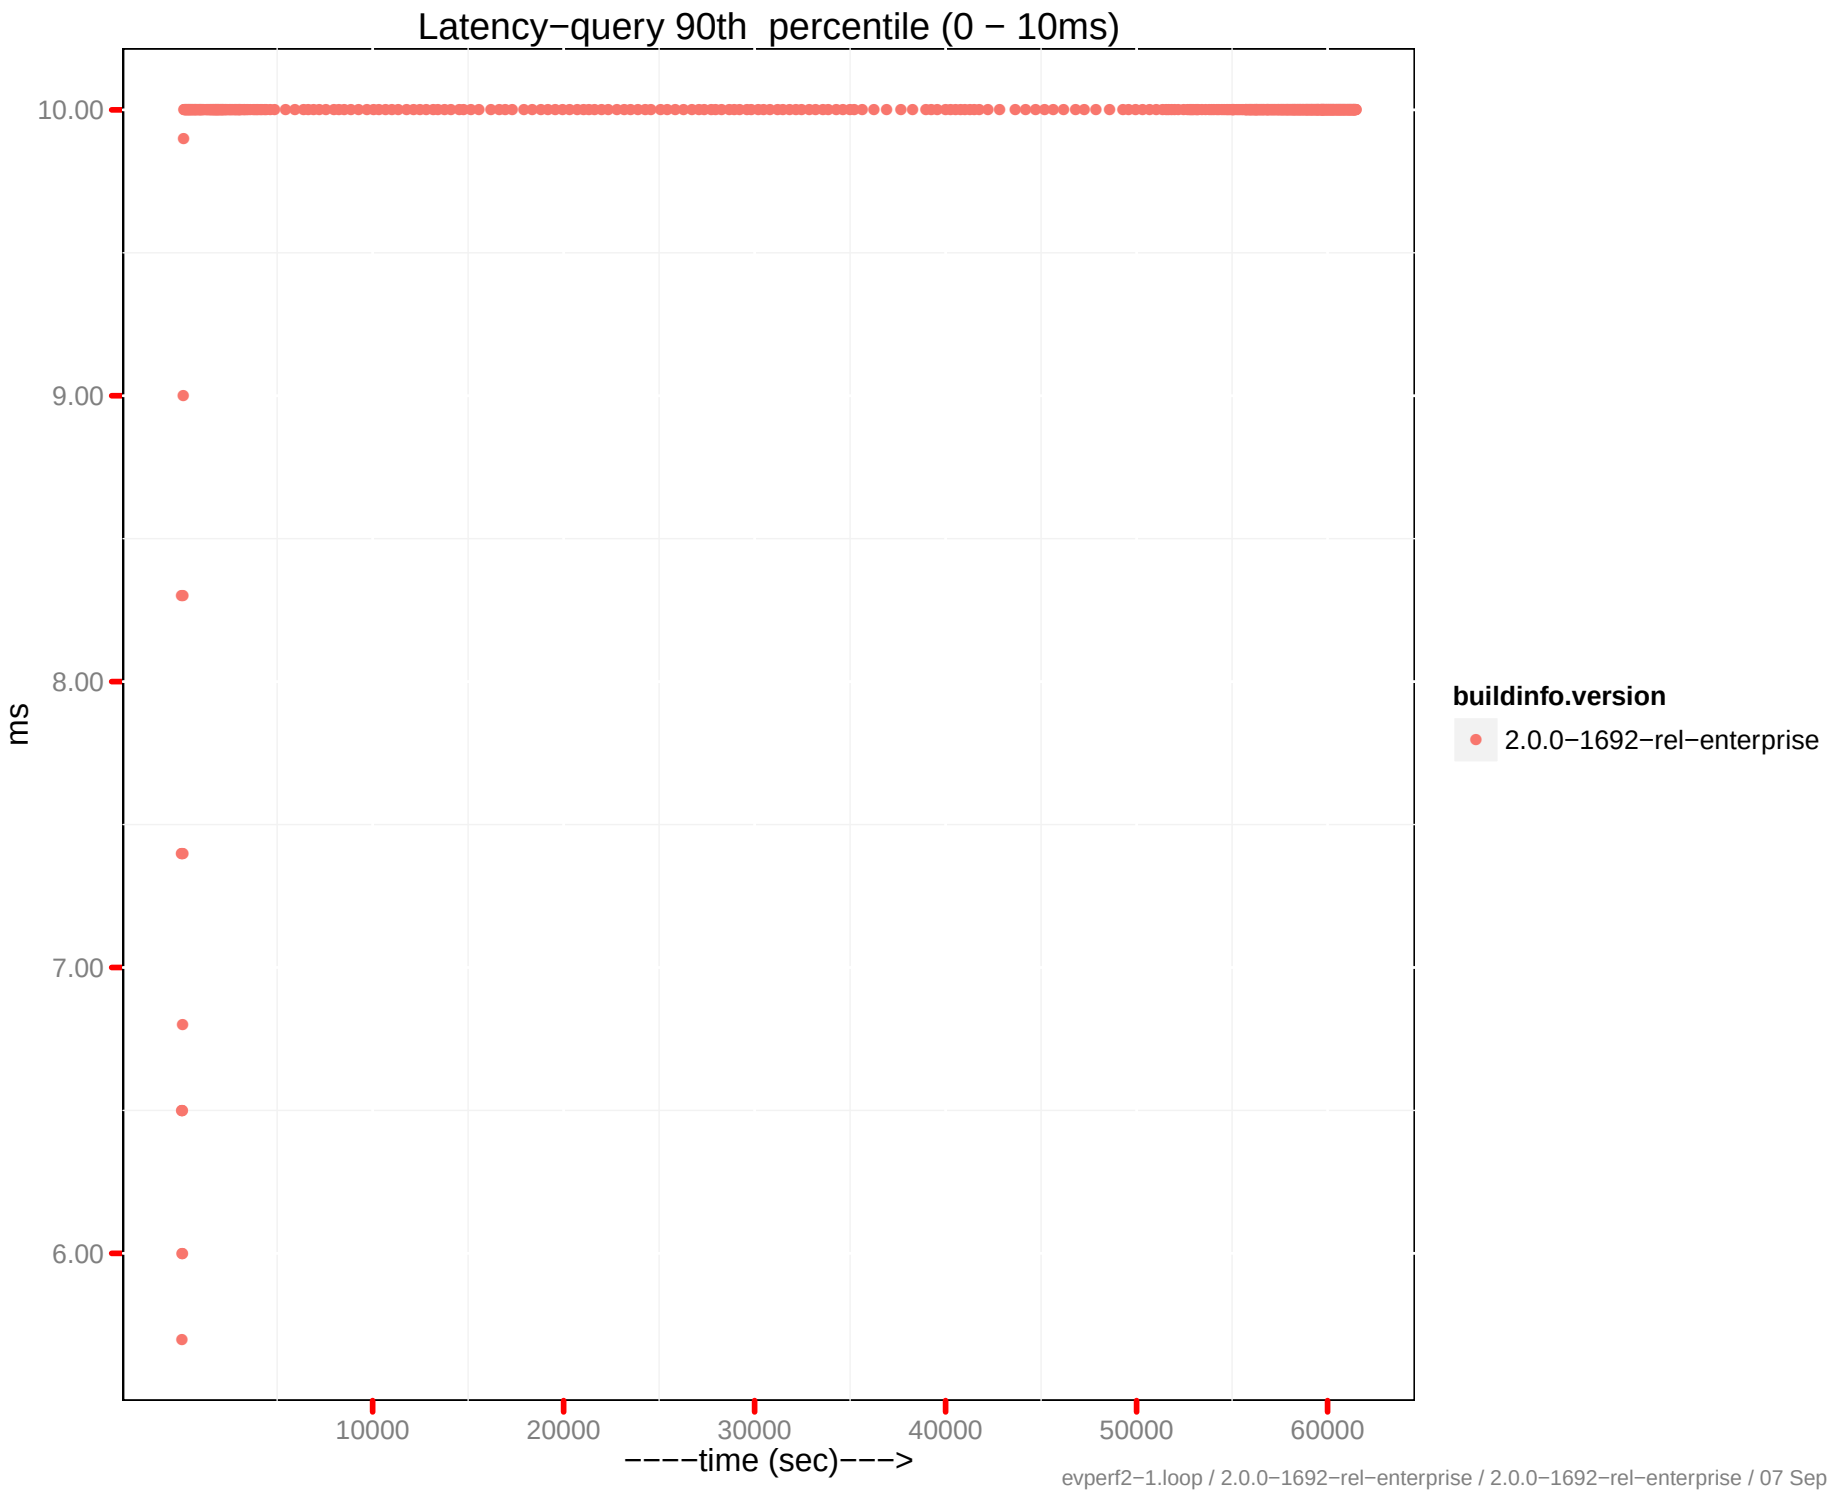
Latency−query 90th percentile (0 − 10ms)
10.00
9.00
8.00
7.00
6.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

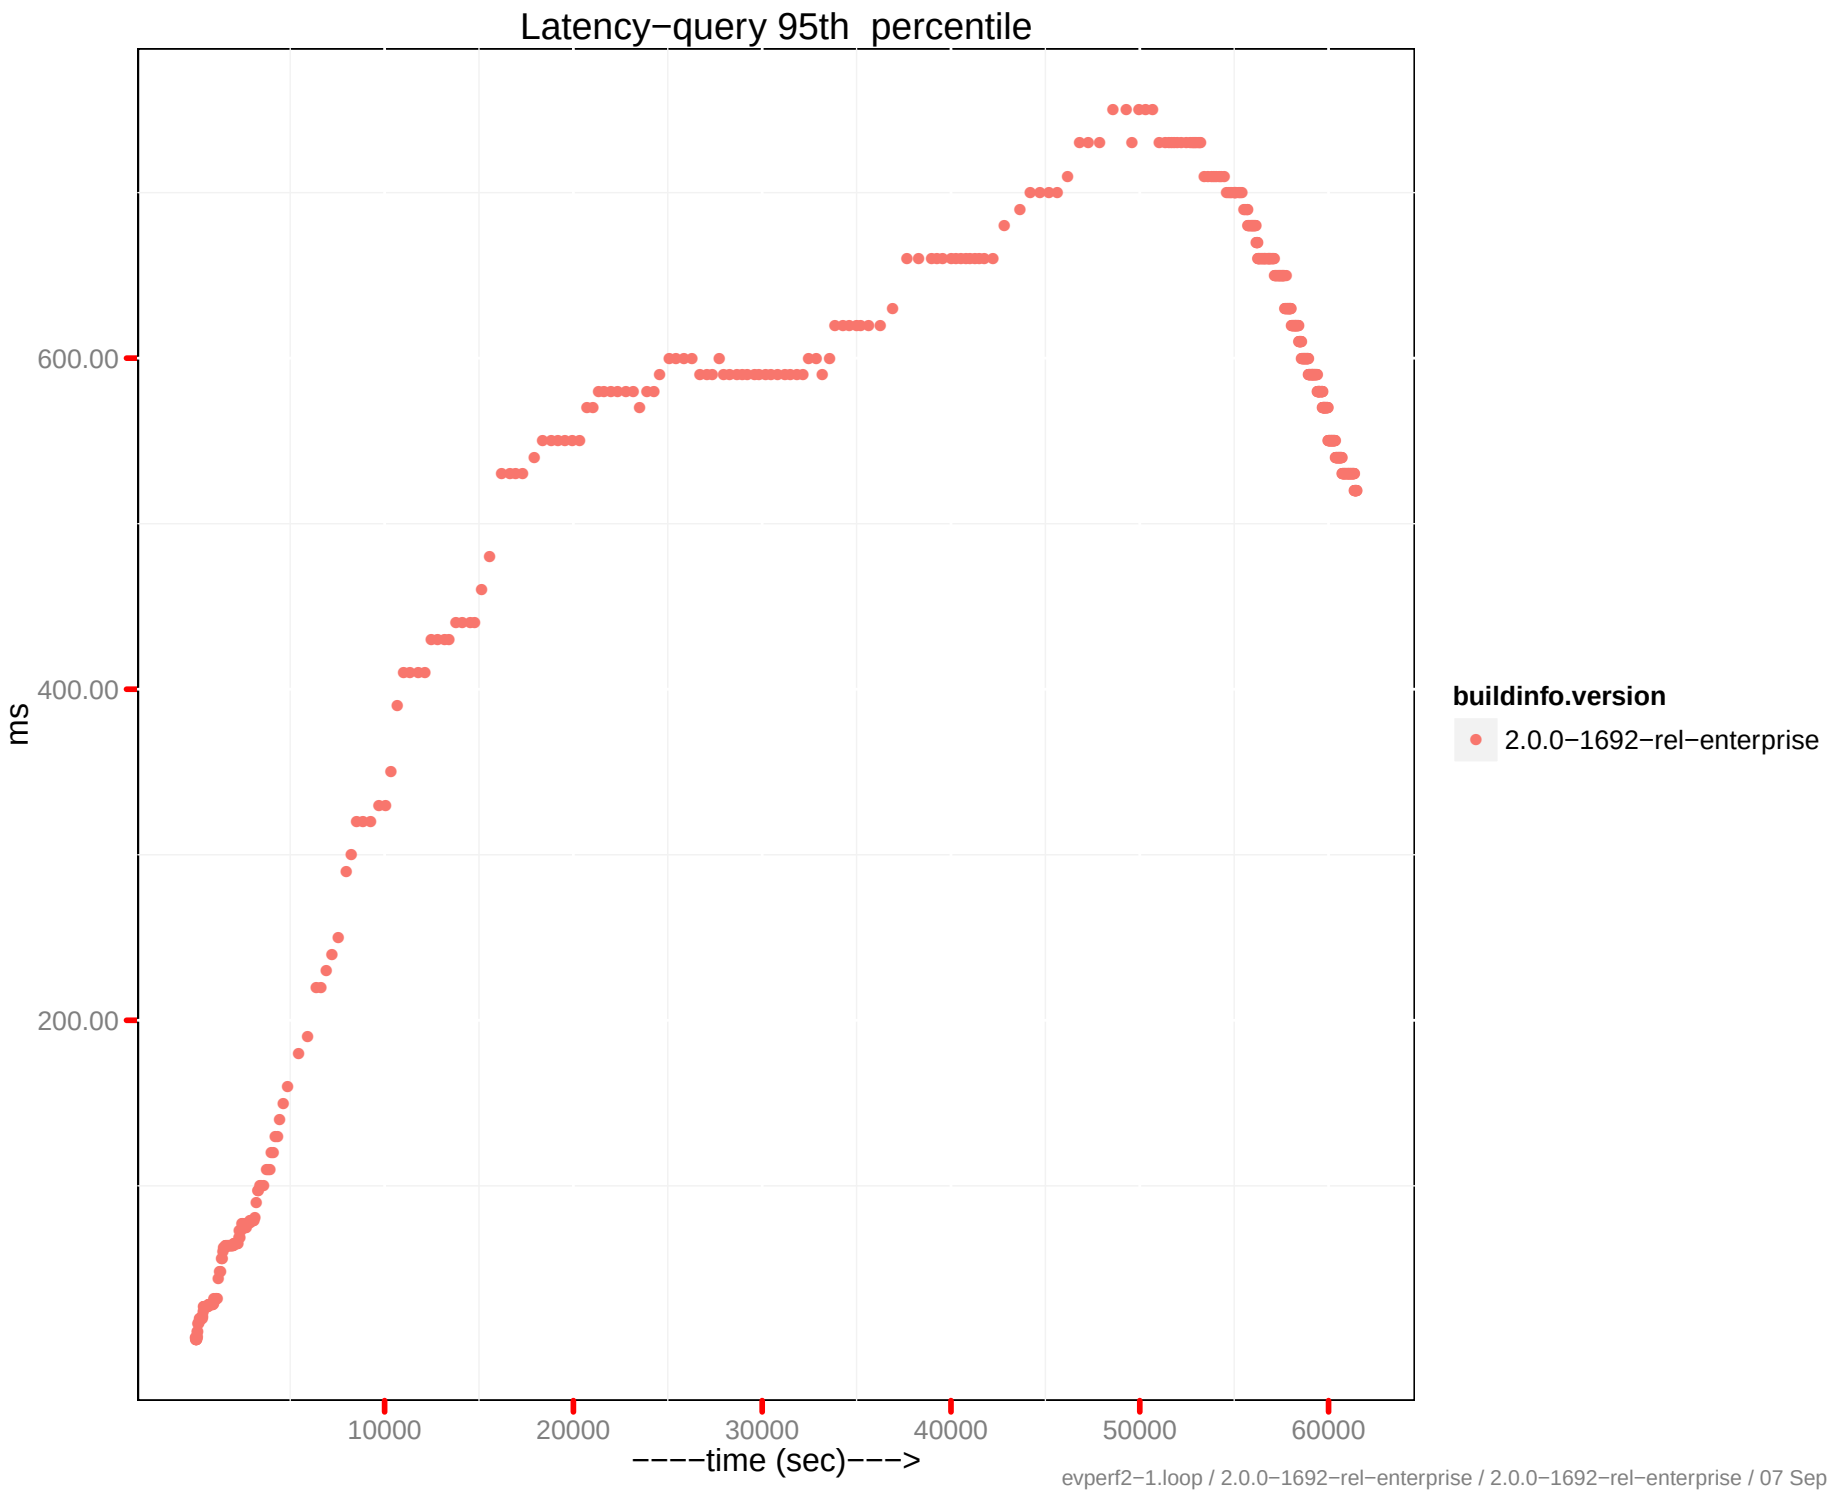
Latency−query 95th percentile
600.00
400.00
200.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

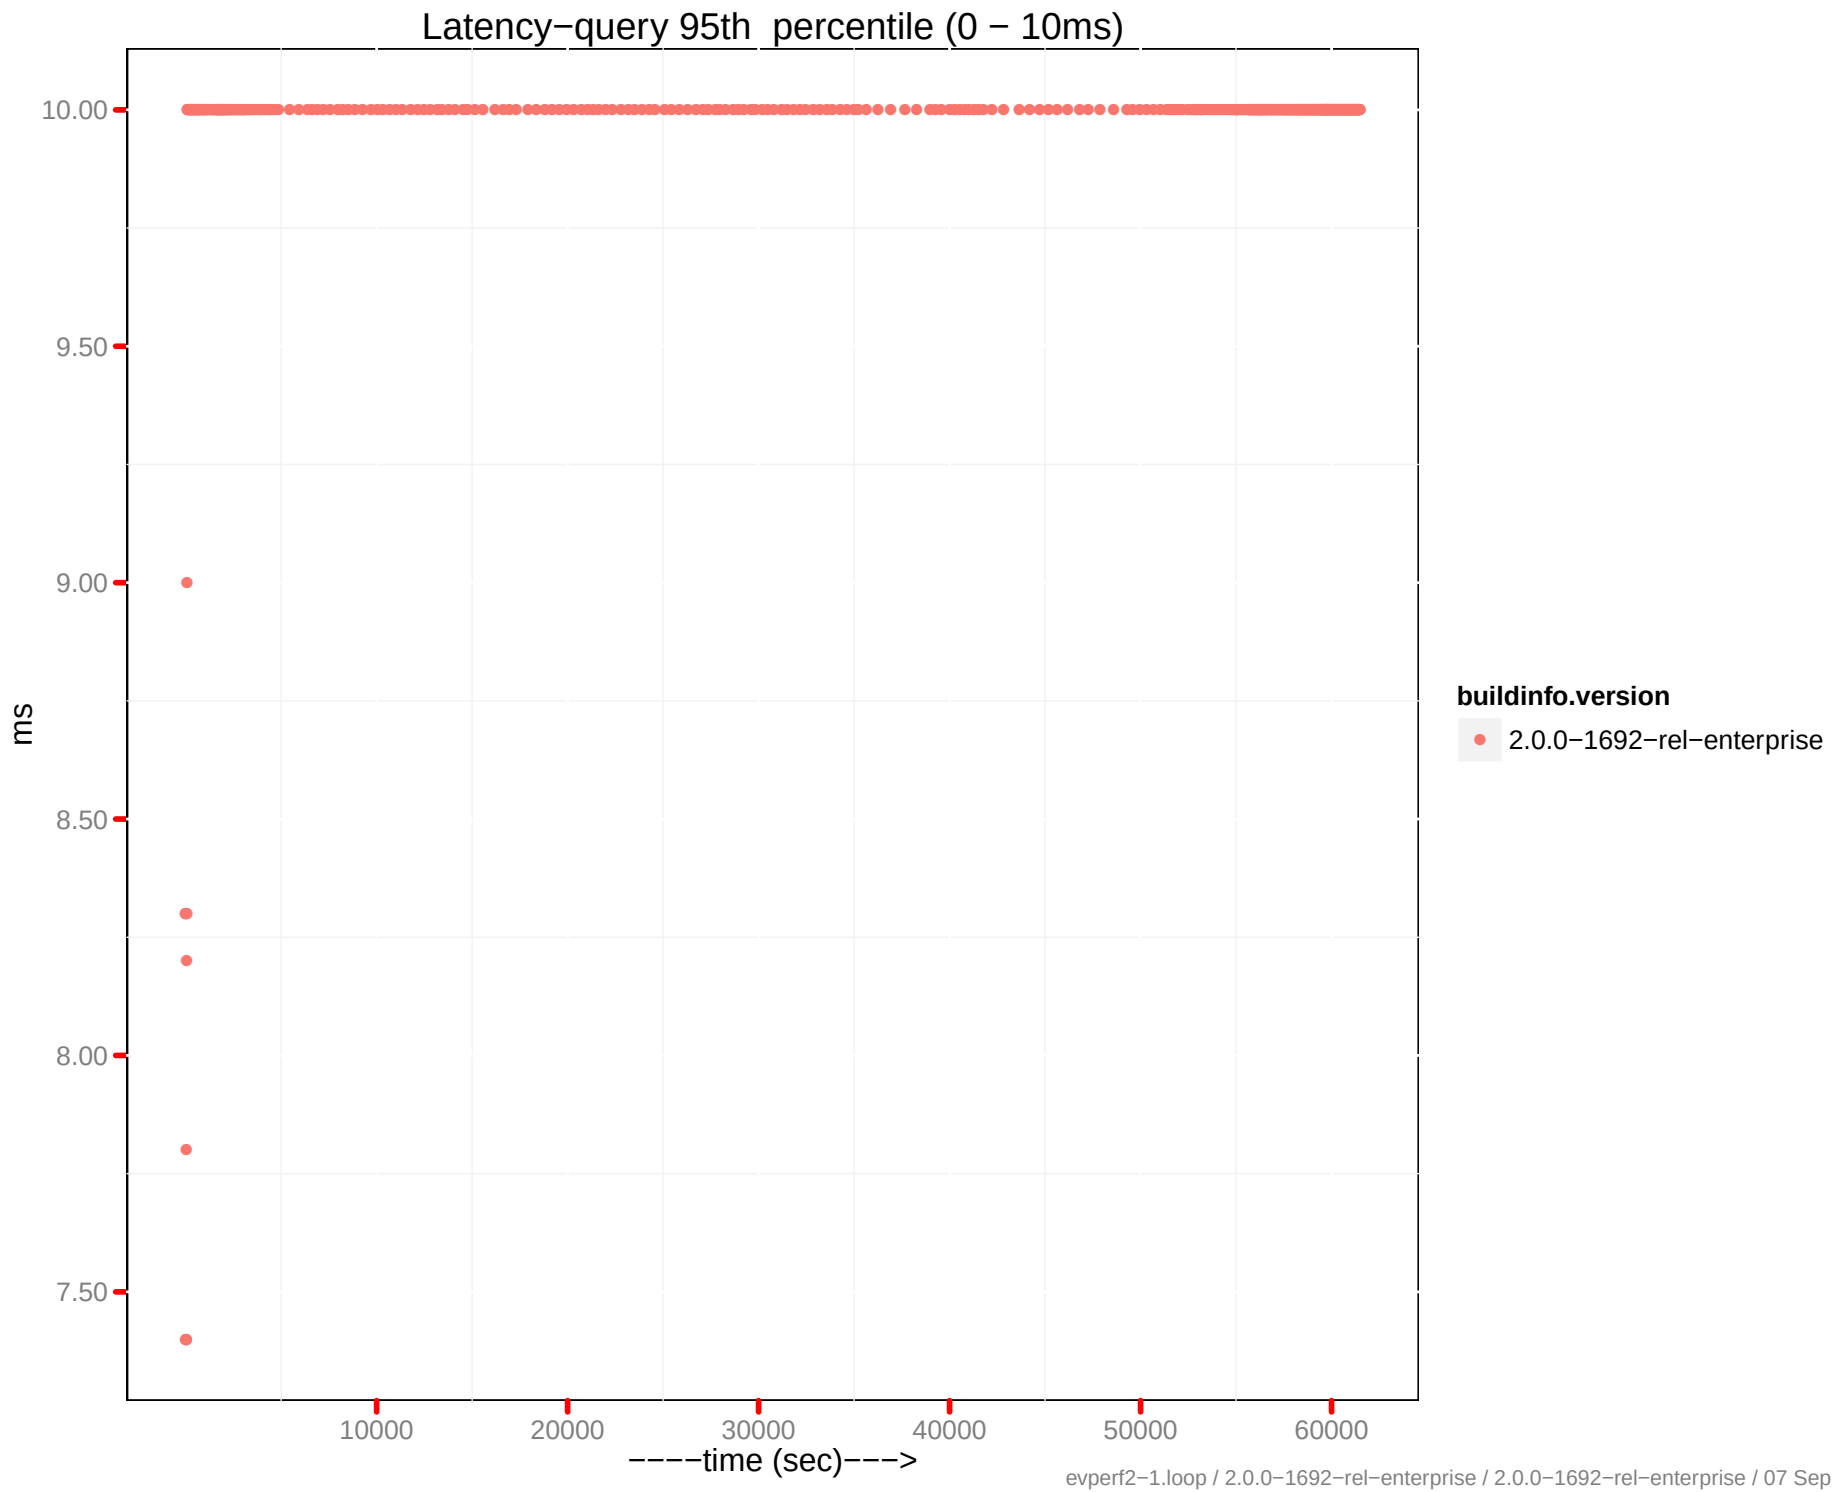
Latency−query 95th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
8.00
7.50
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

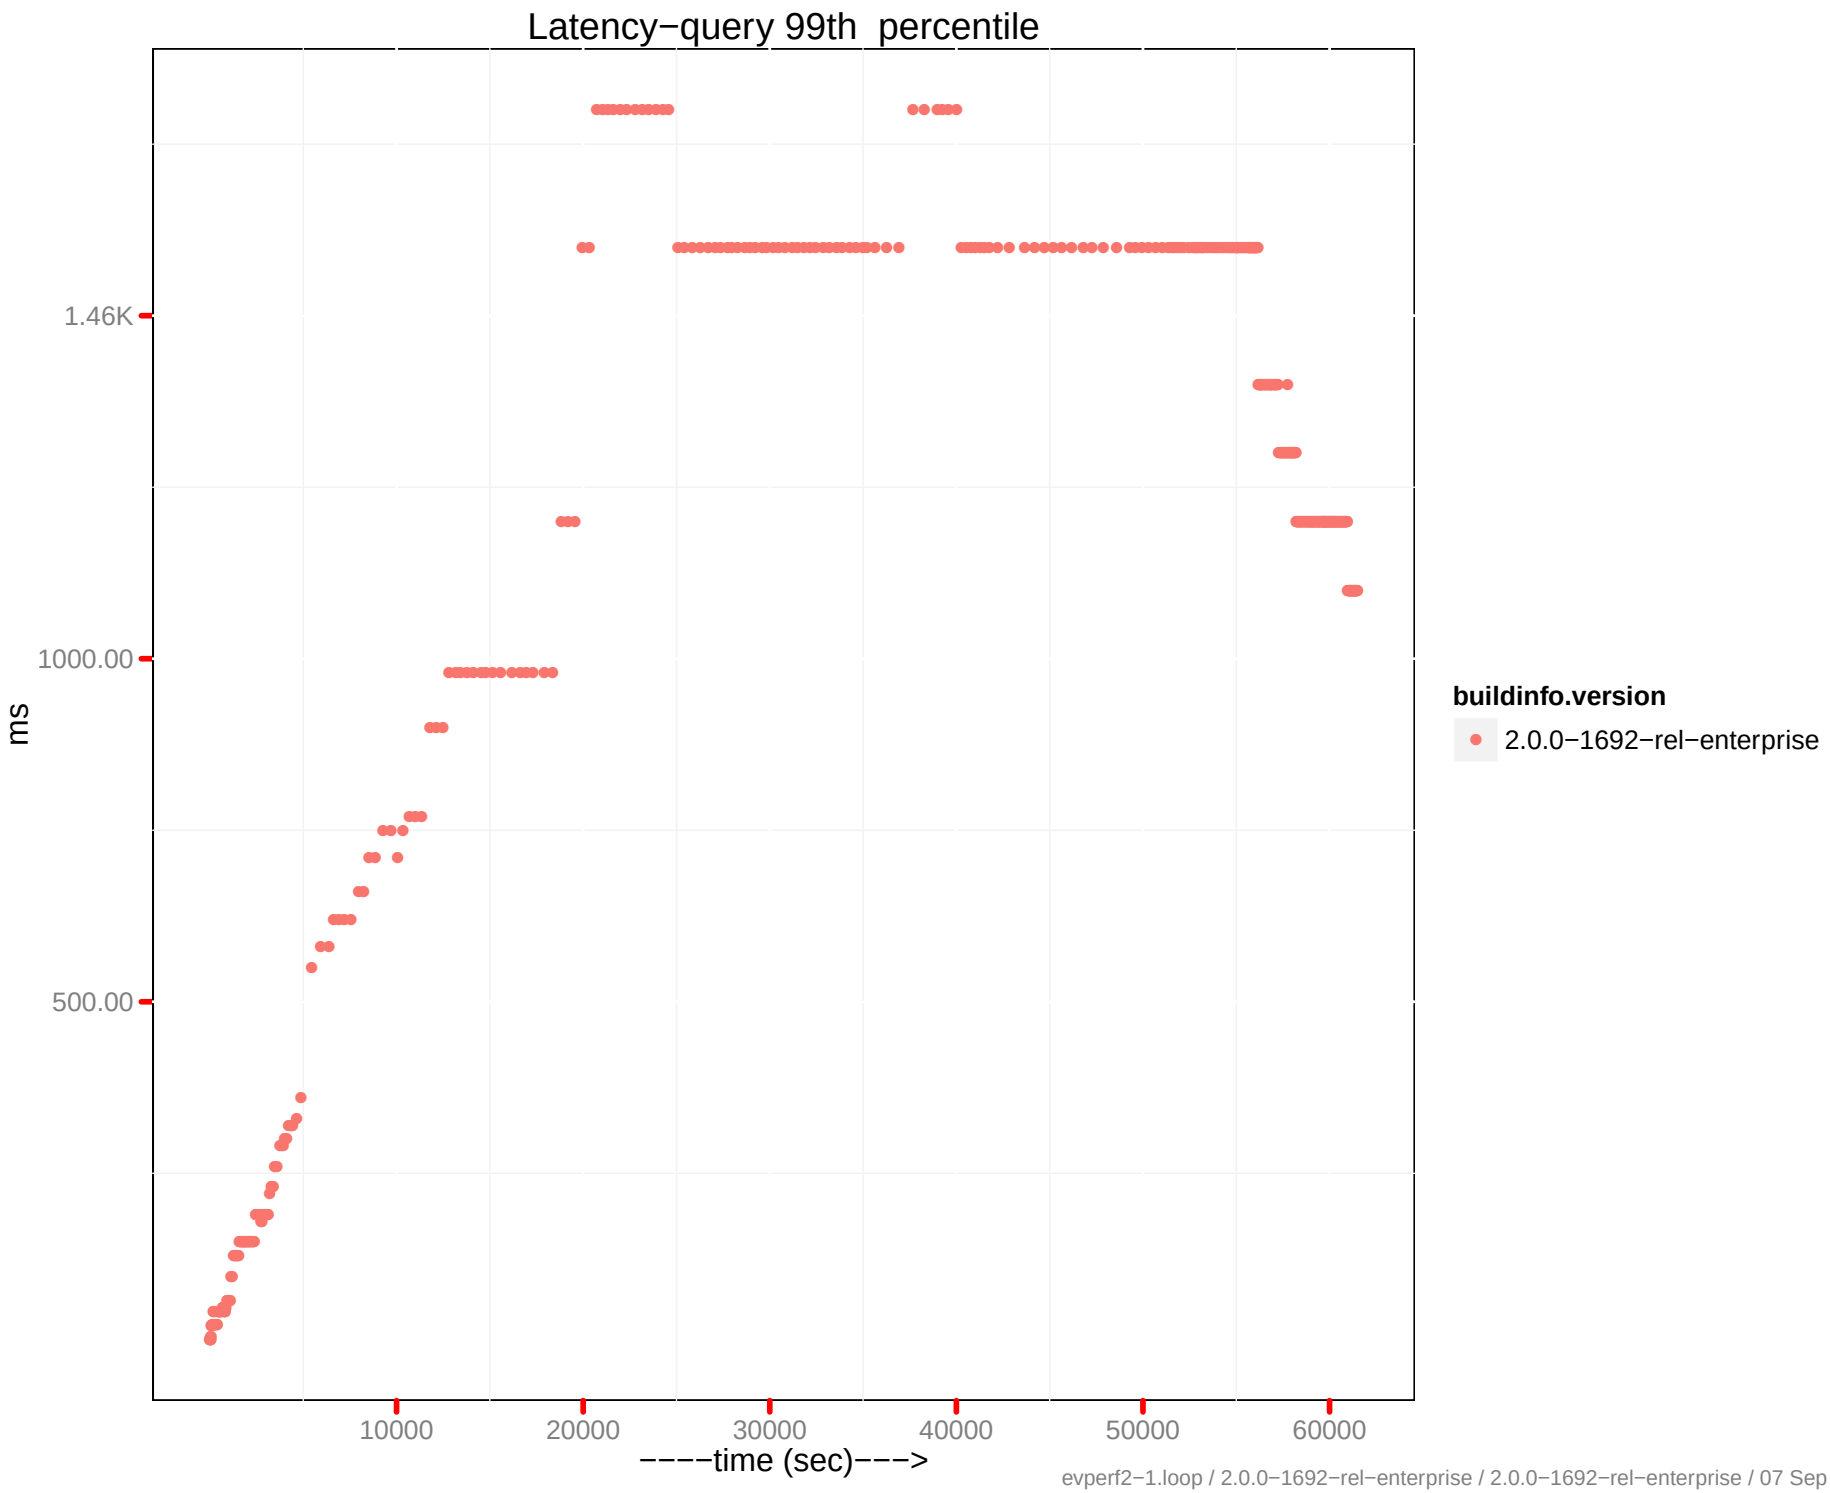
Latency−query 99th percentile
1.46K
1000.00
500.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

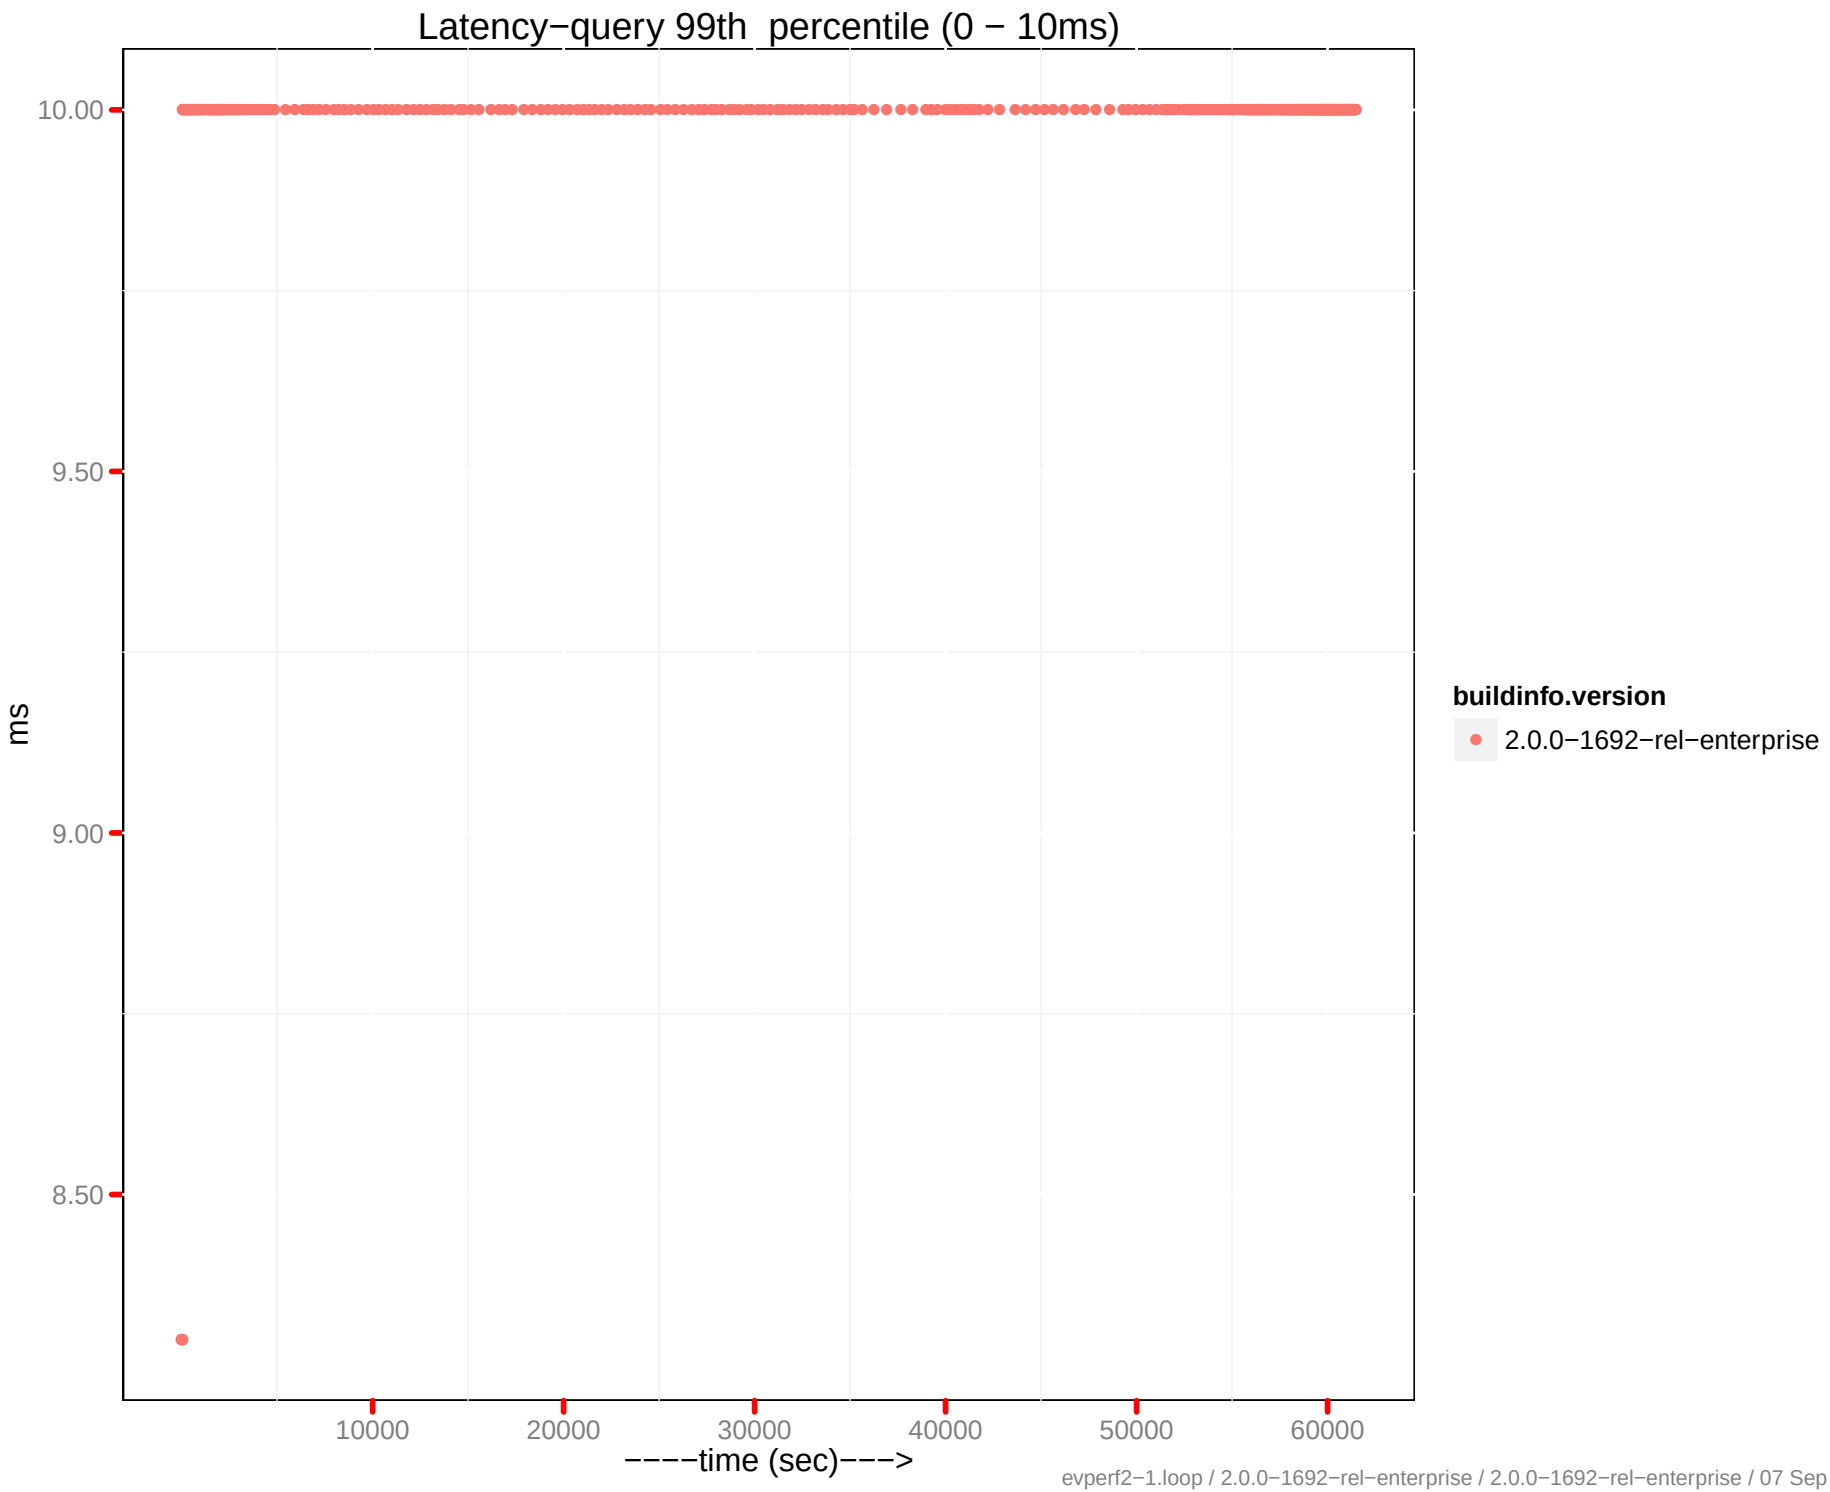
Latency−query 99th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

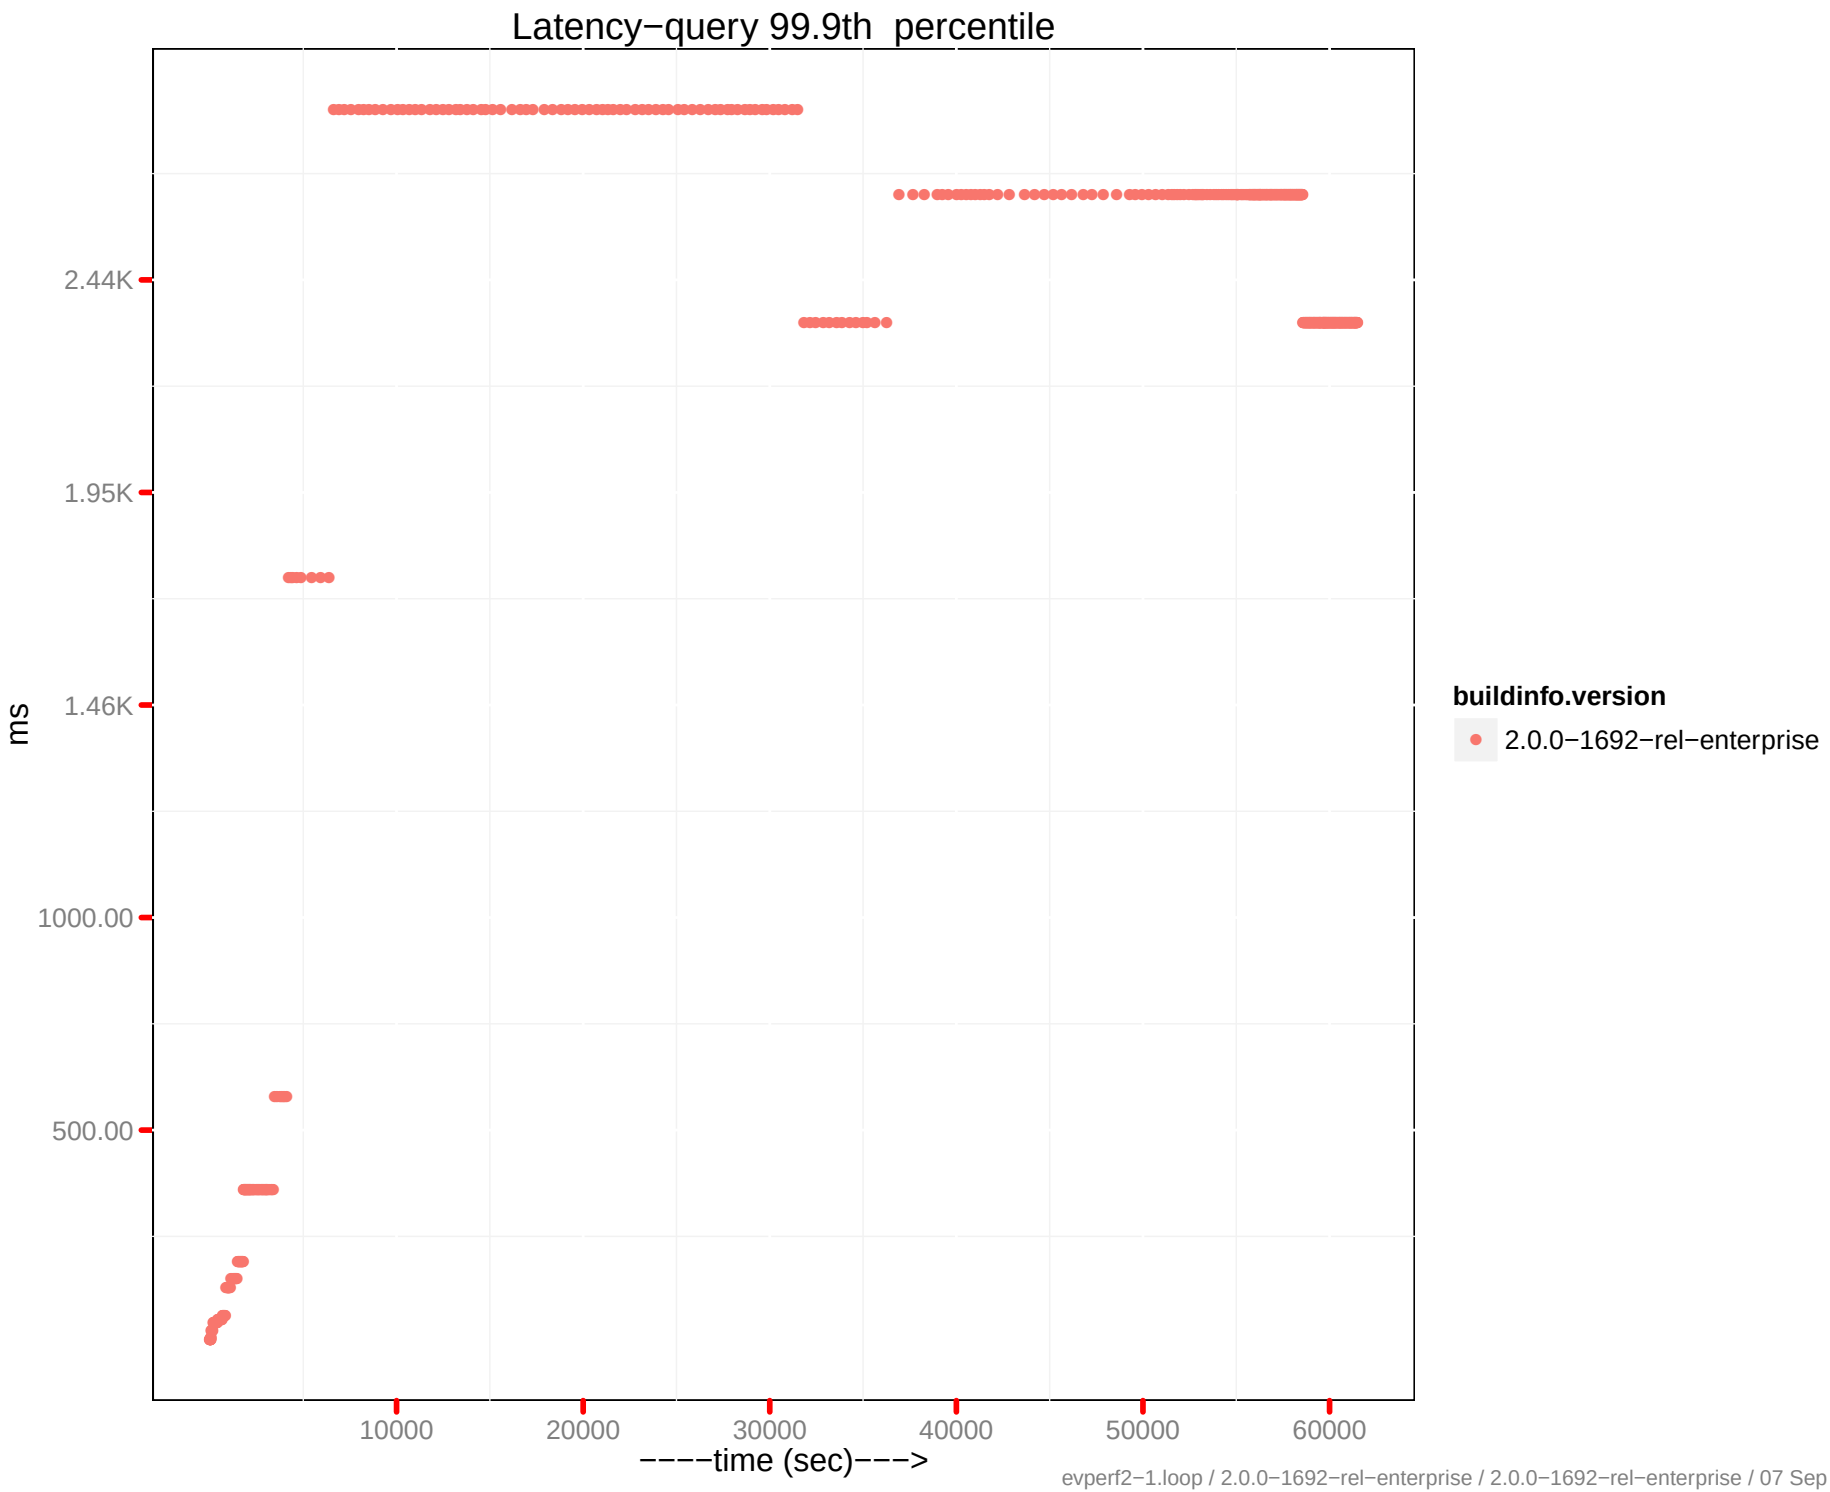
Latency−query 99.9th percentile
2.44K
1.95K
1.46K
1000.00
500.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

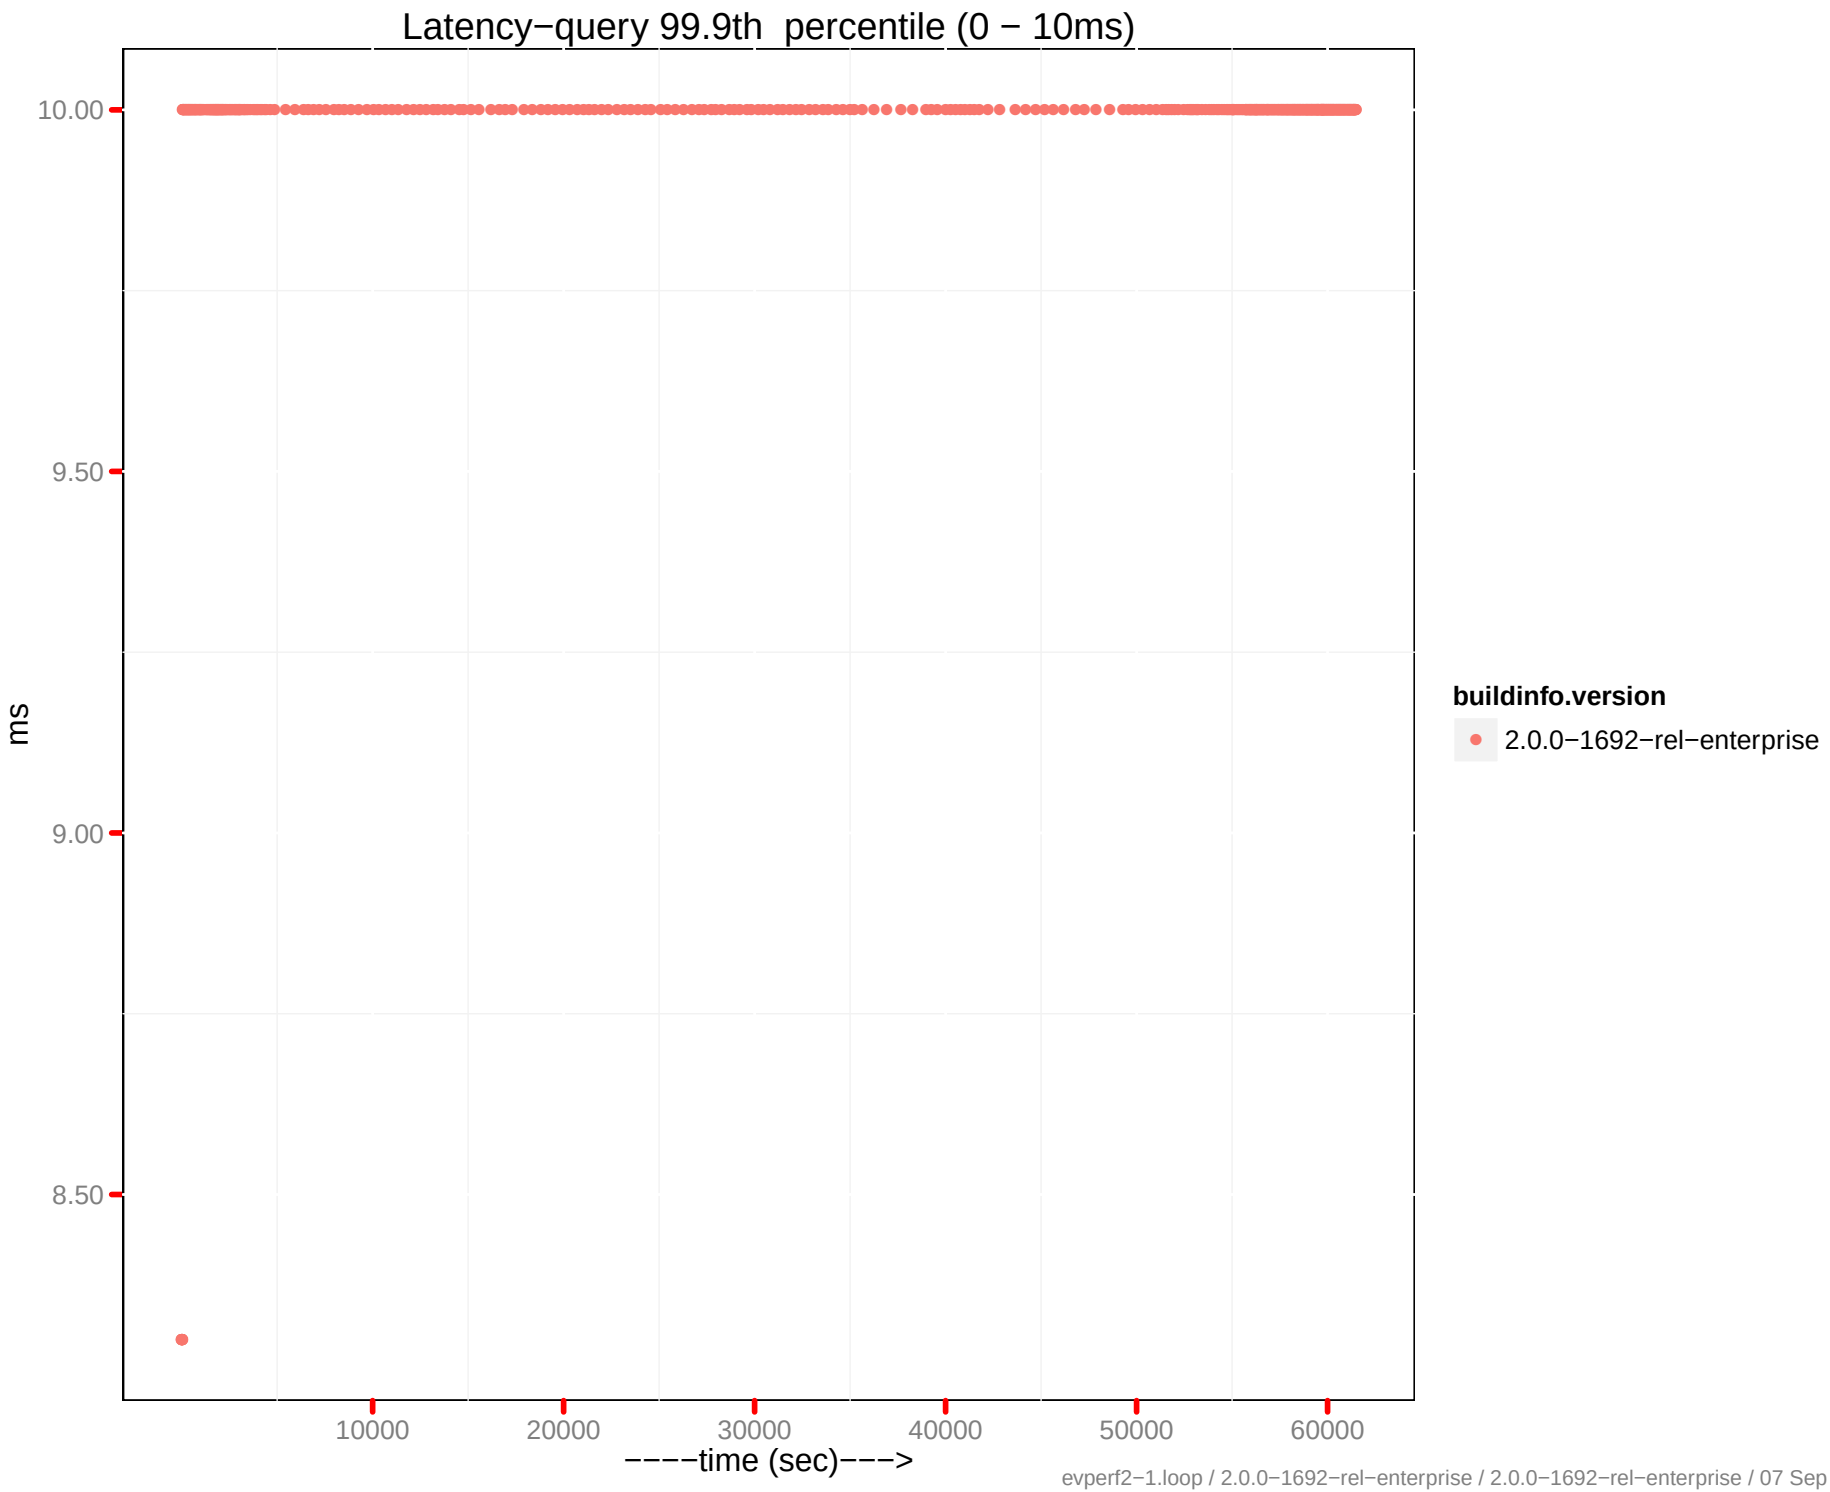
Latency−query 99.9th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

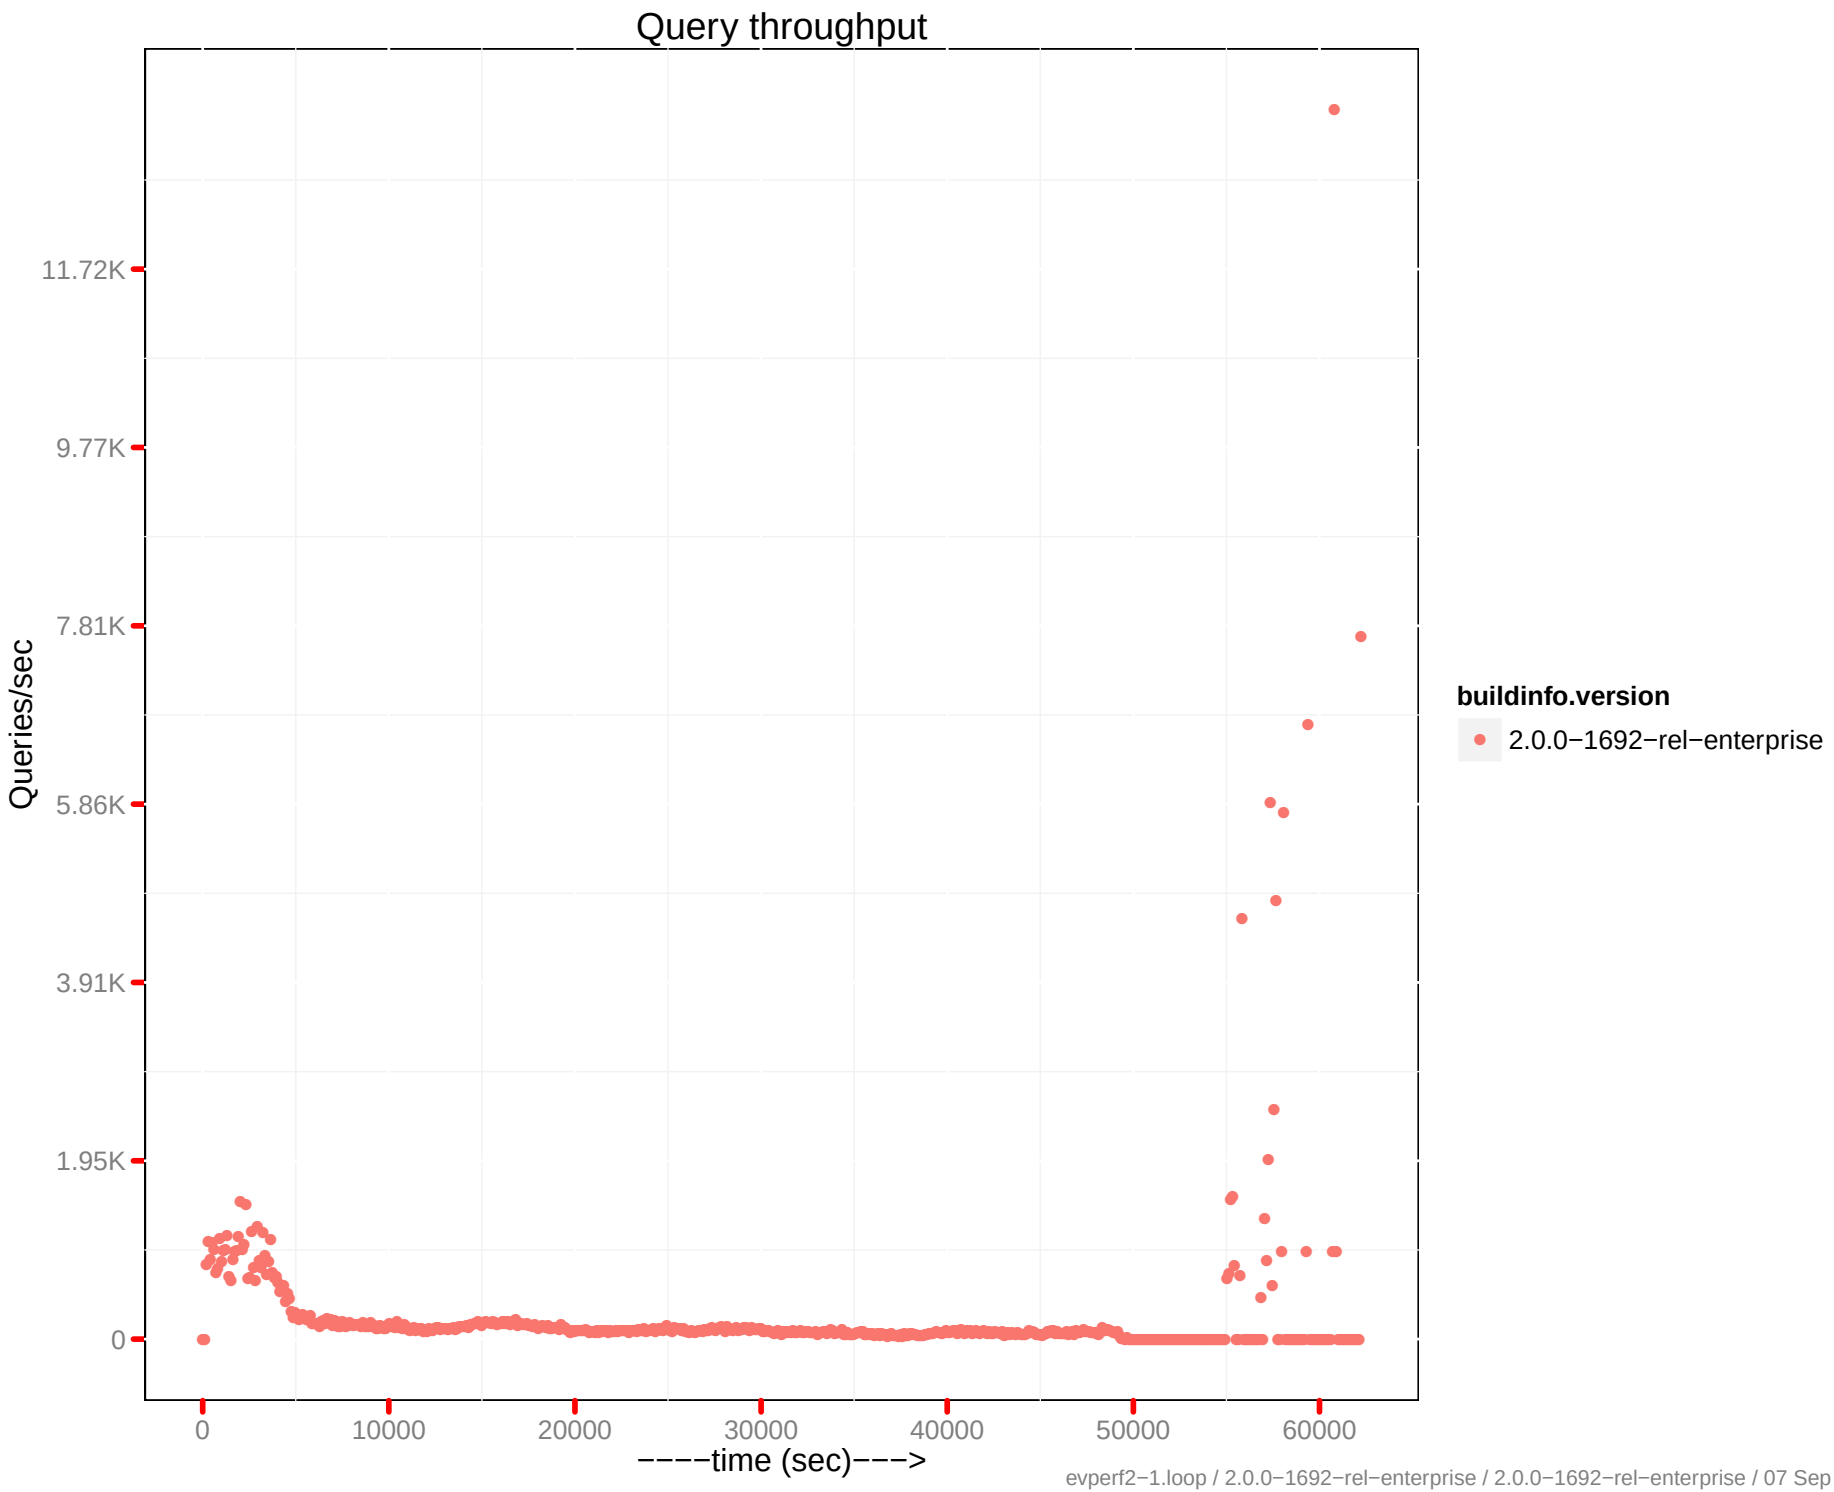
Query throughput
Queries/sec
11.72K
9.77K
7.81K
5.86K
3.91K
1.95K
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1692-rel-enterprise
evperf2-1.loop / 2.0.0-1692-rel-enterprise / 2.0.0-1692-rel-enterprise / 07 Sep 2012

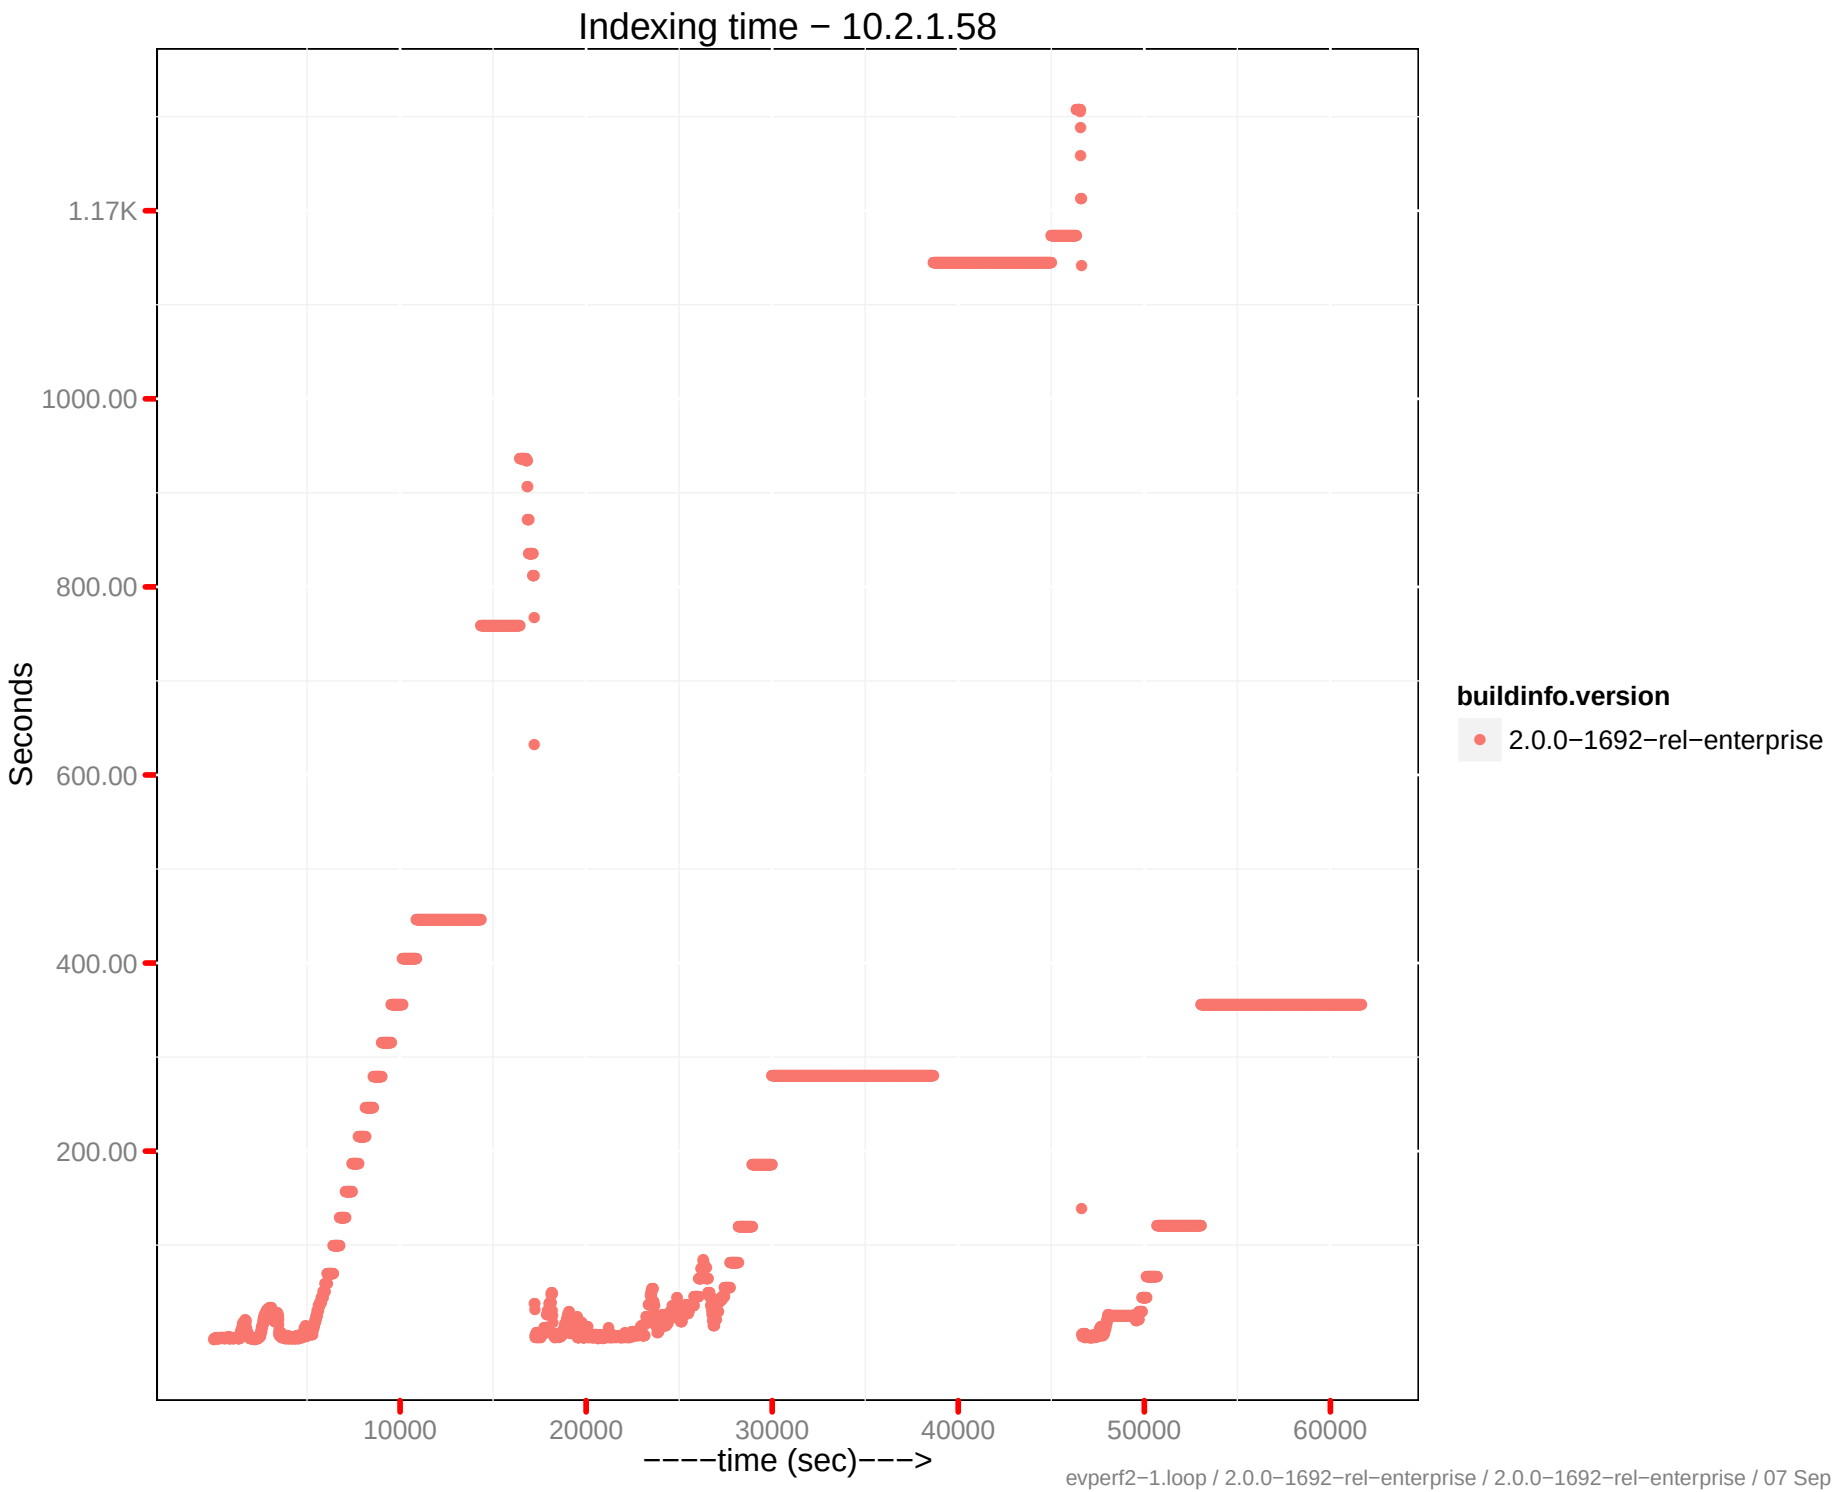
Indexing time − 10.2.1.58
1.17K
1000.00
800.00
600.00
400.00
200.00
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

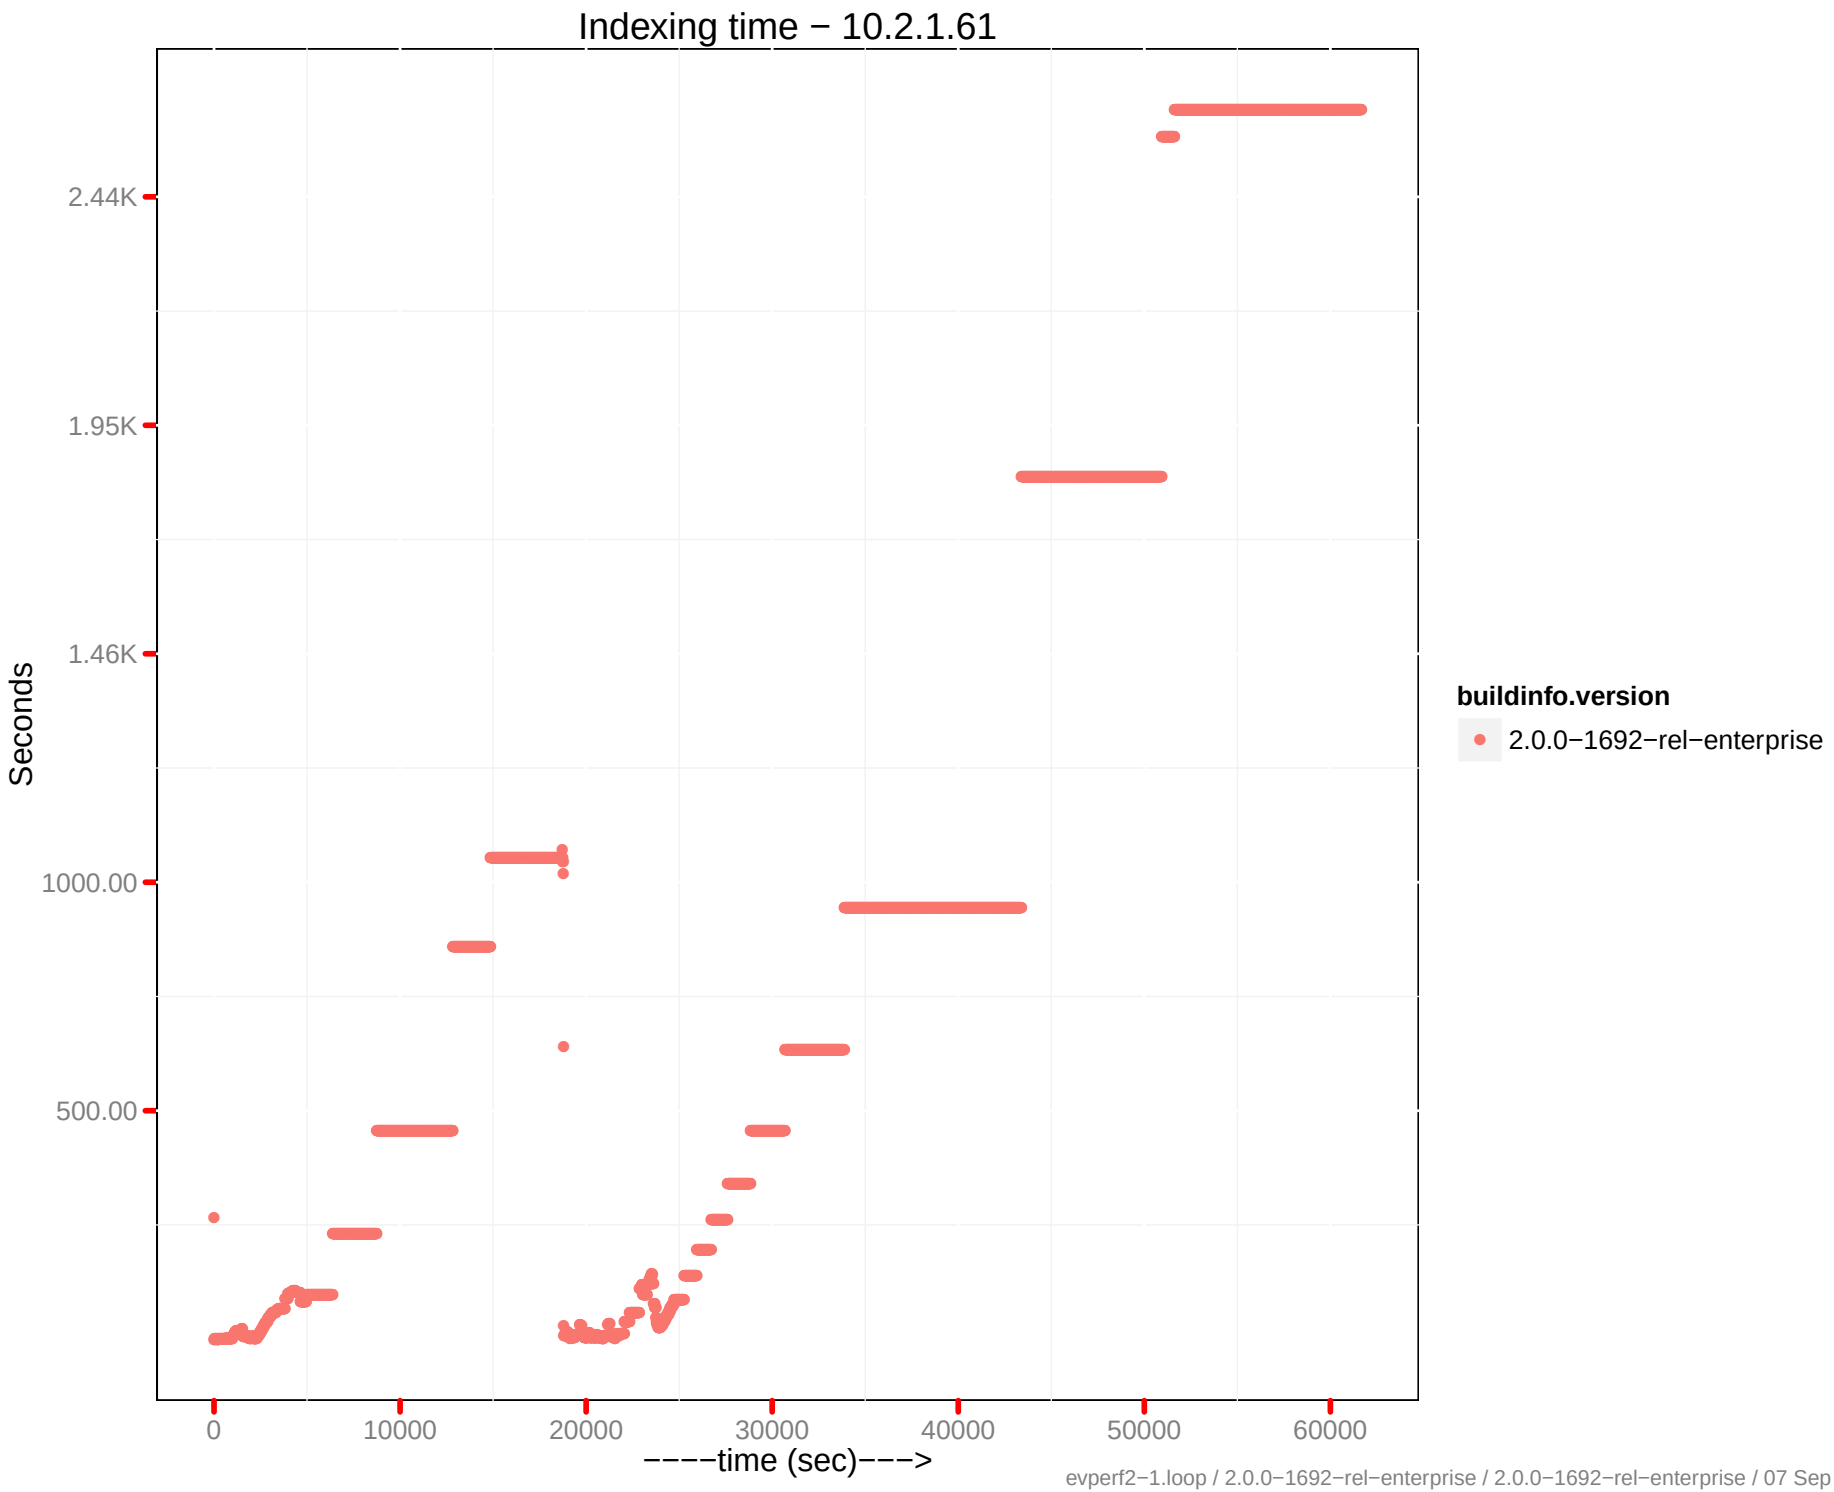
Indexing time − 10.2.1.61
Seconds
2.44K
1.95K
1.46K
1000.00
500.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

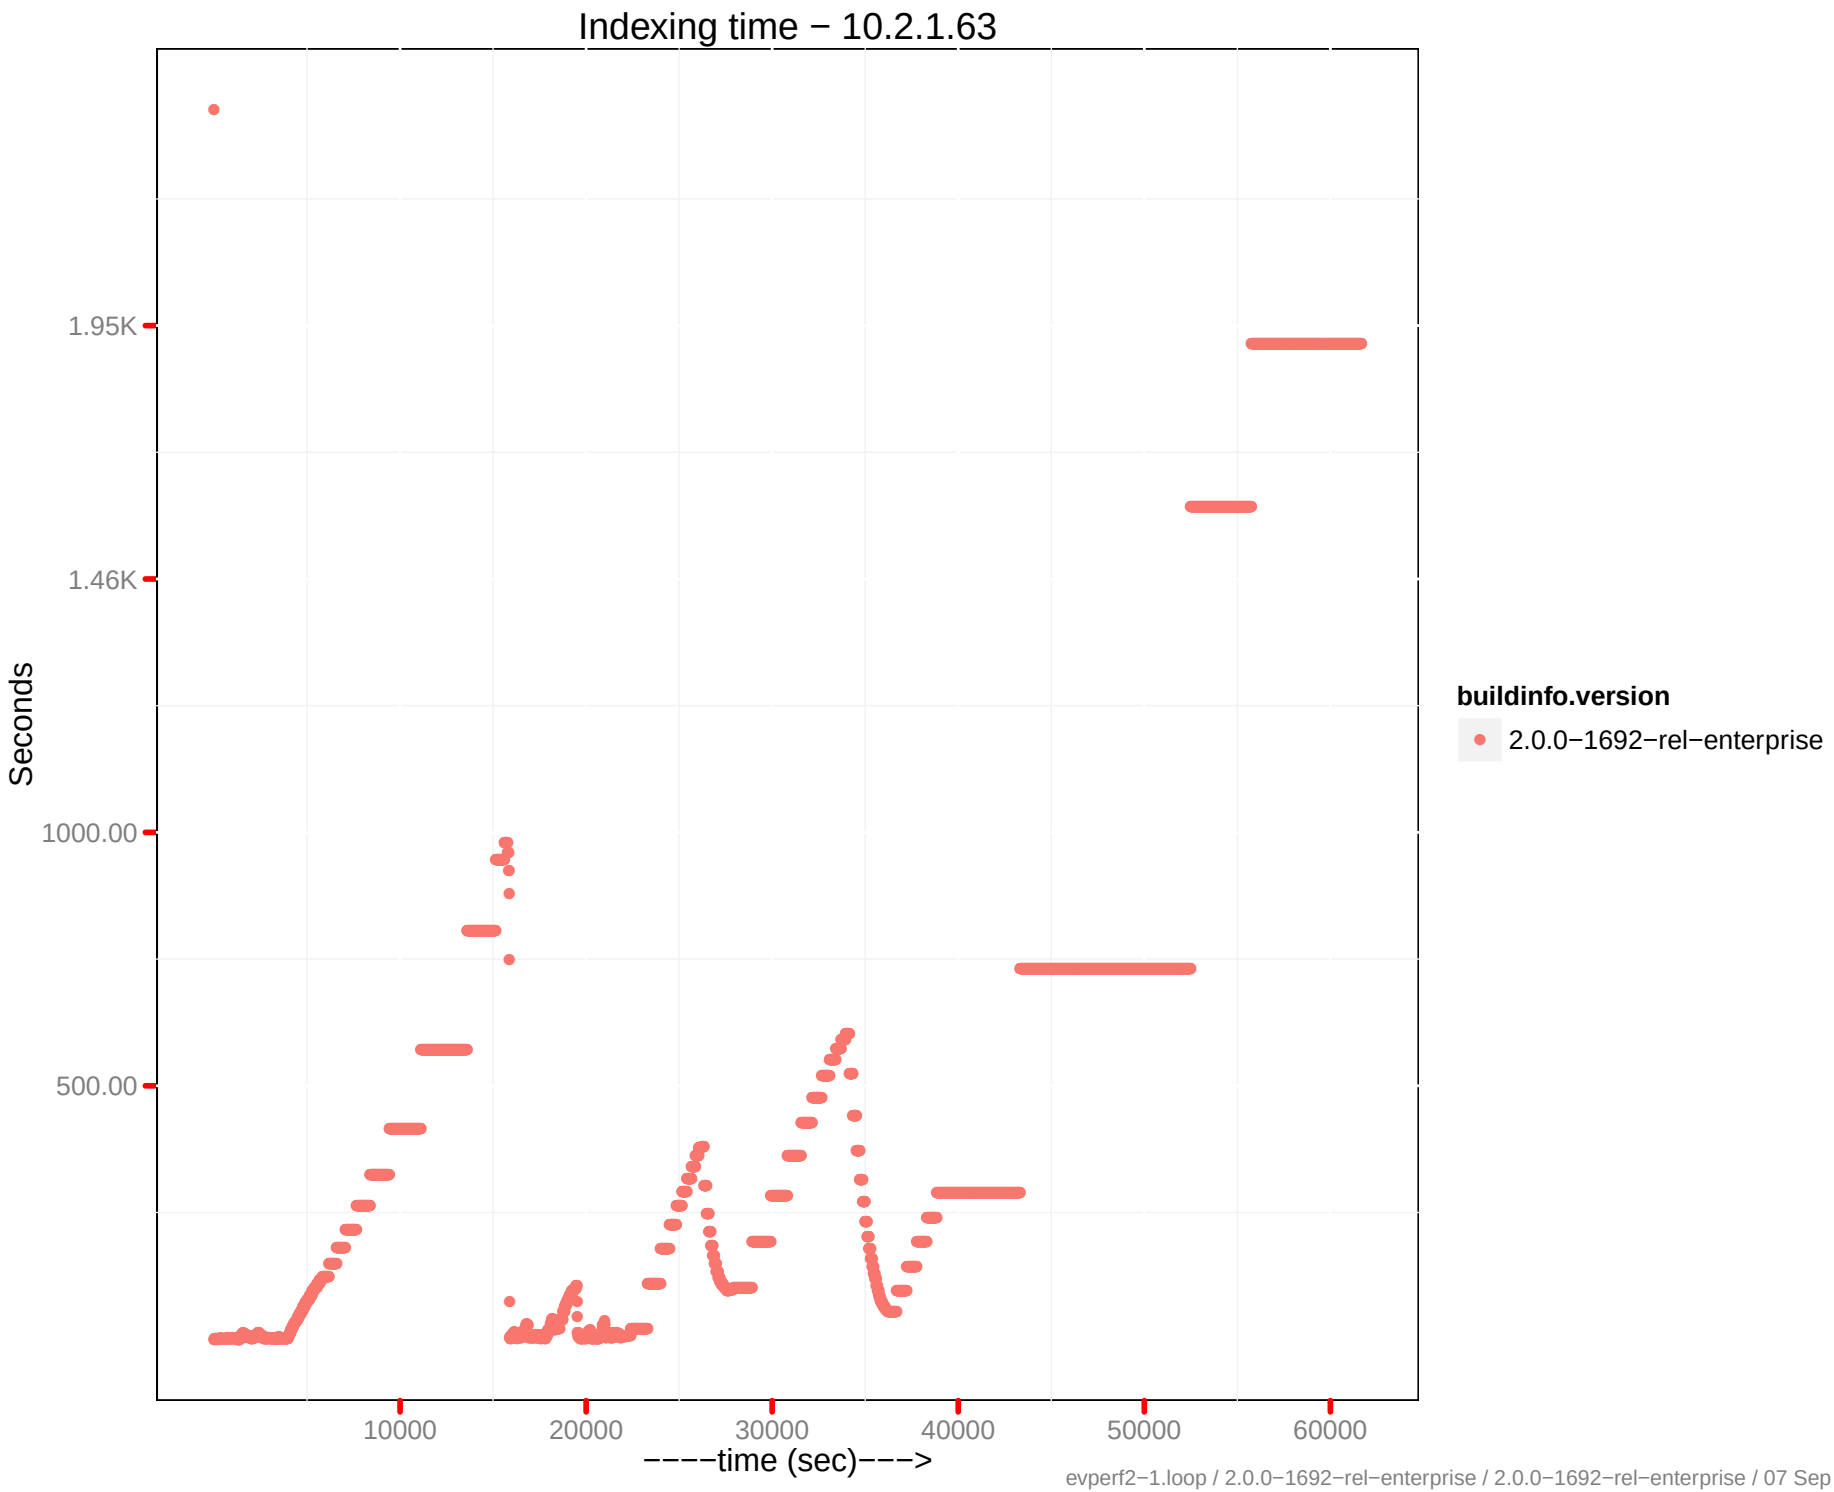
Indexing time − 10.2.1.63
1.95K
1.46K
1000.00
500.00
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

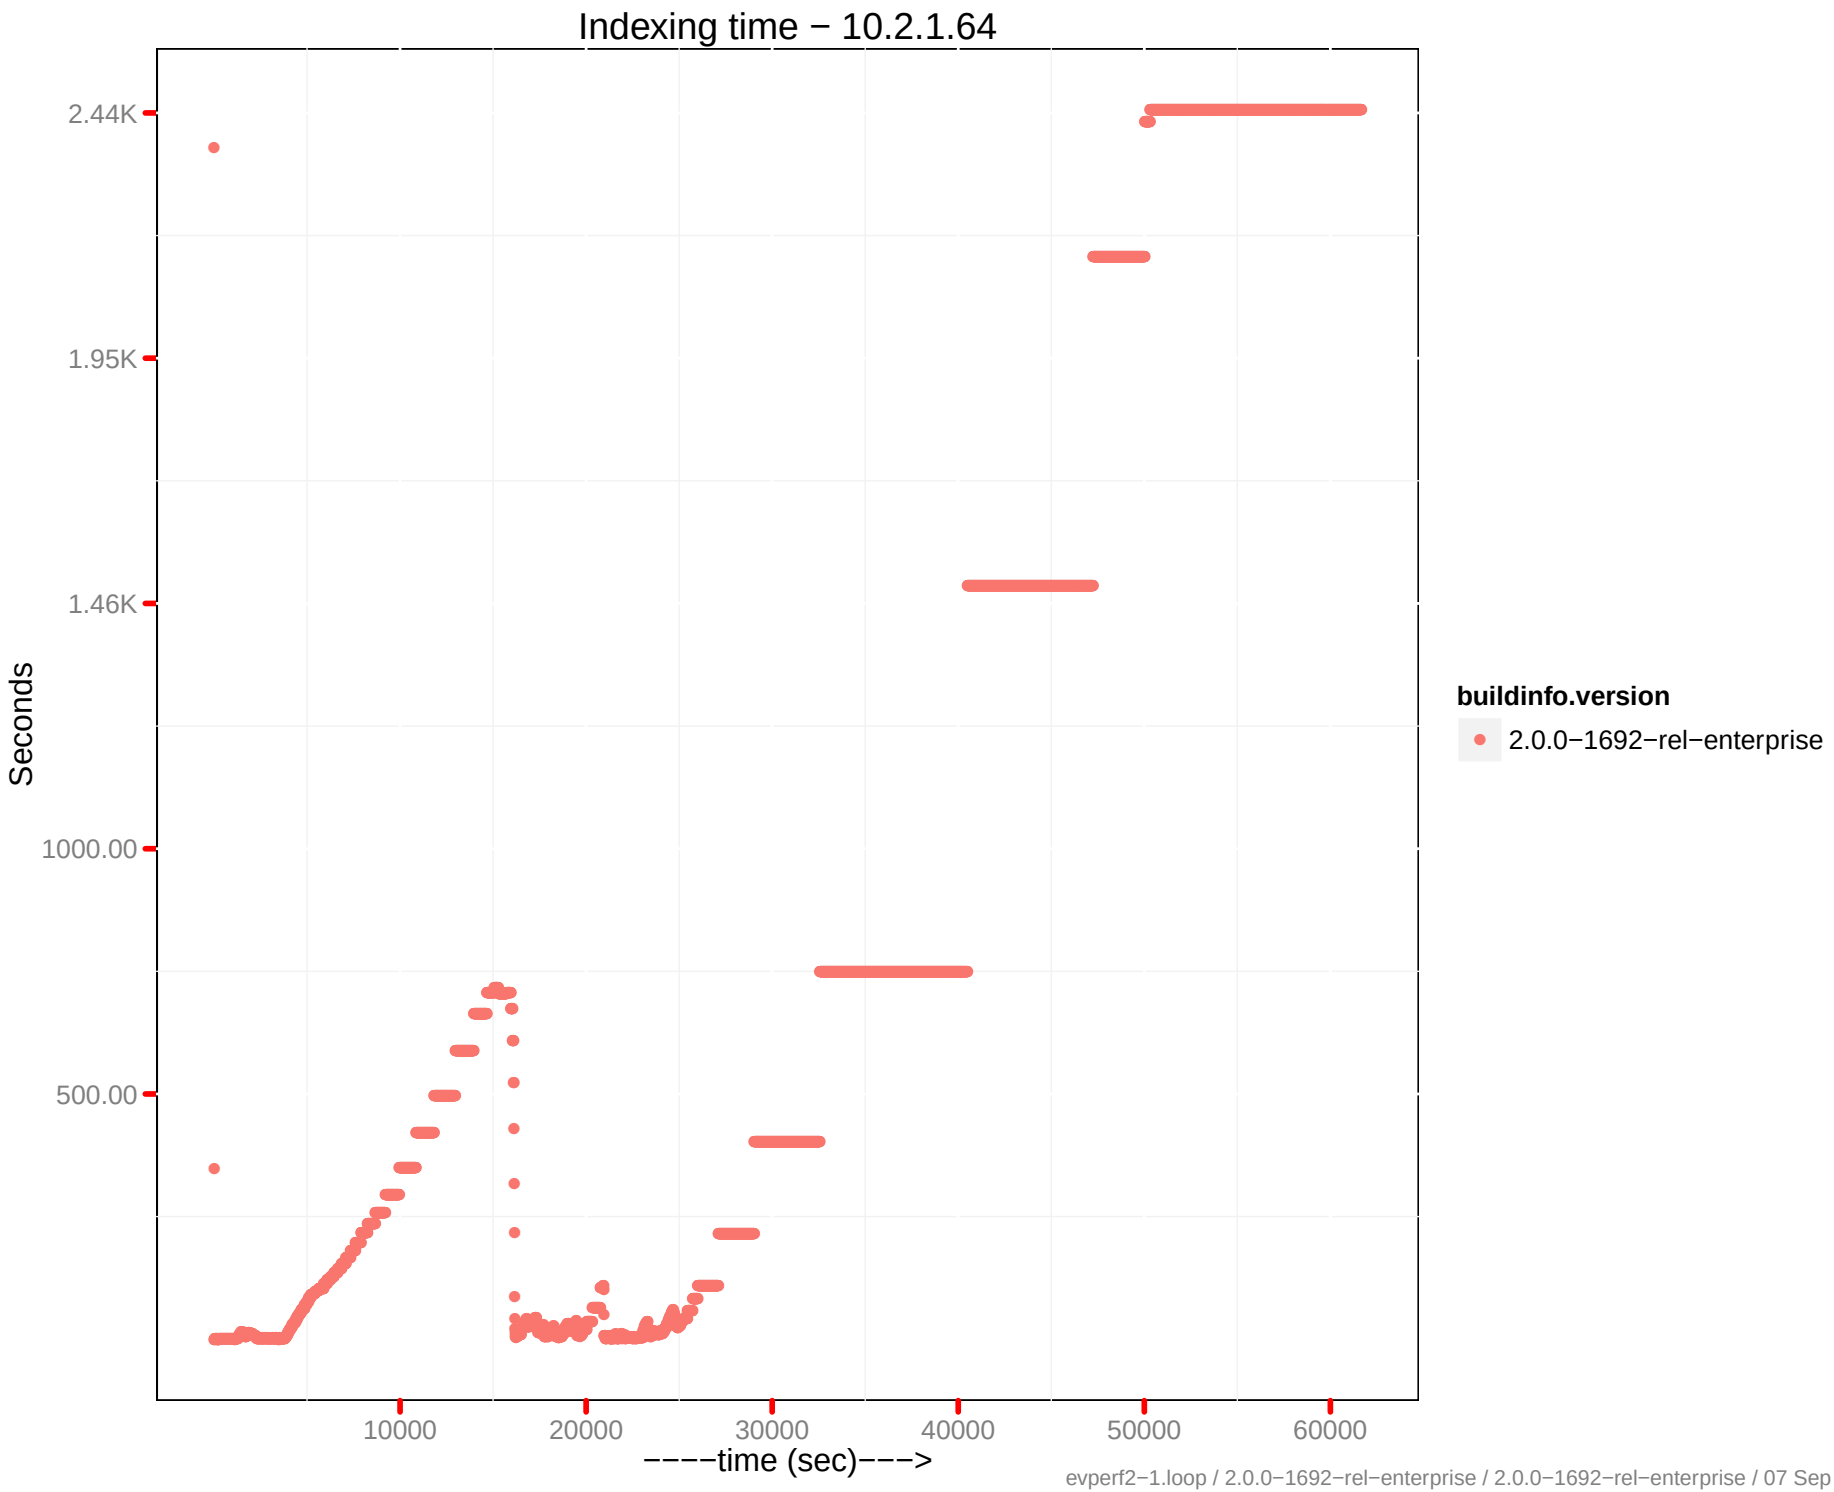
Indexing time − 10.2.1.64
2.44K
1.95K
1.46K
1000.00
500.00
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

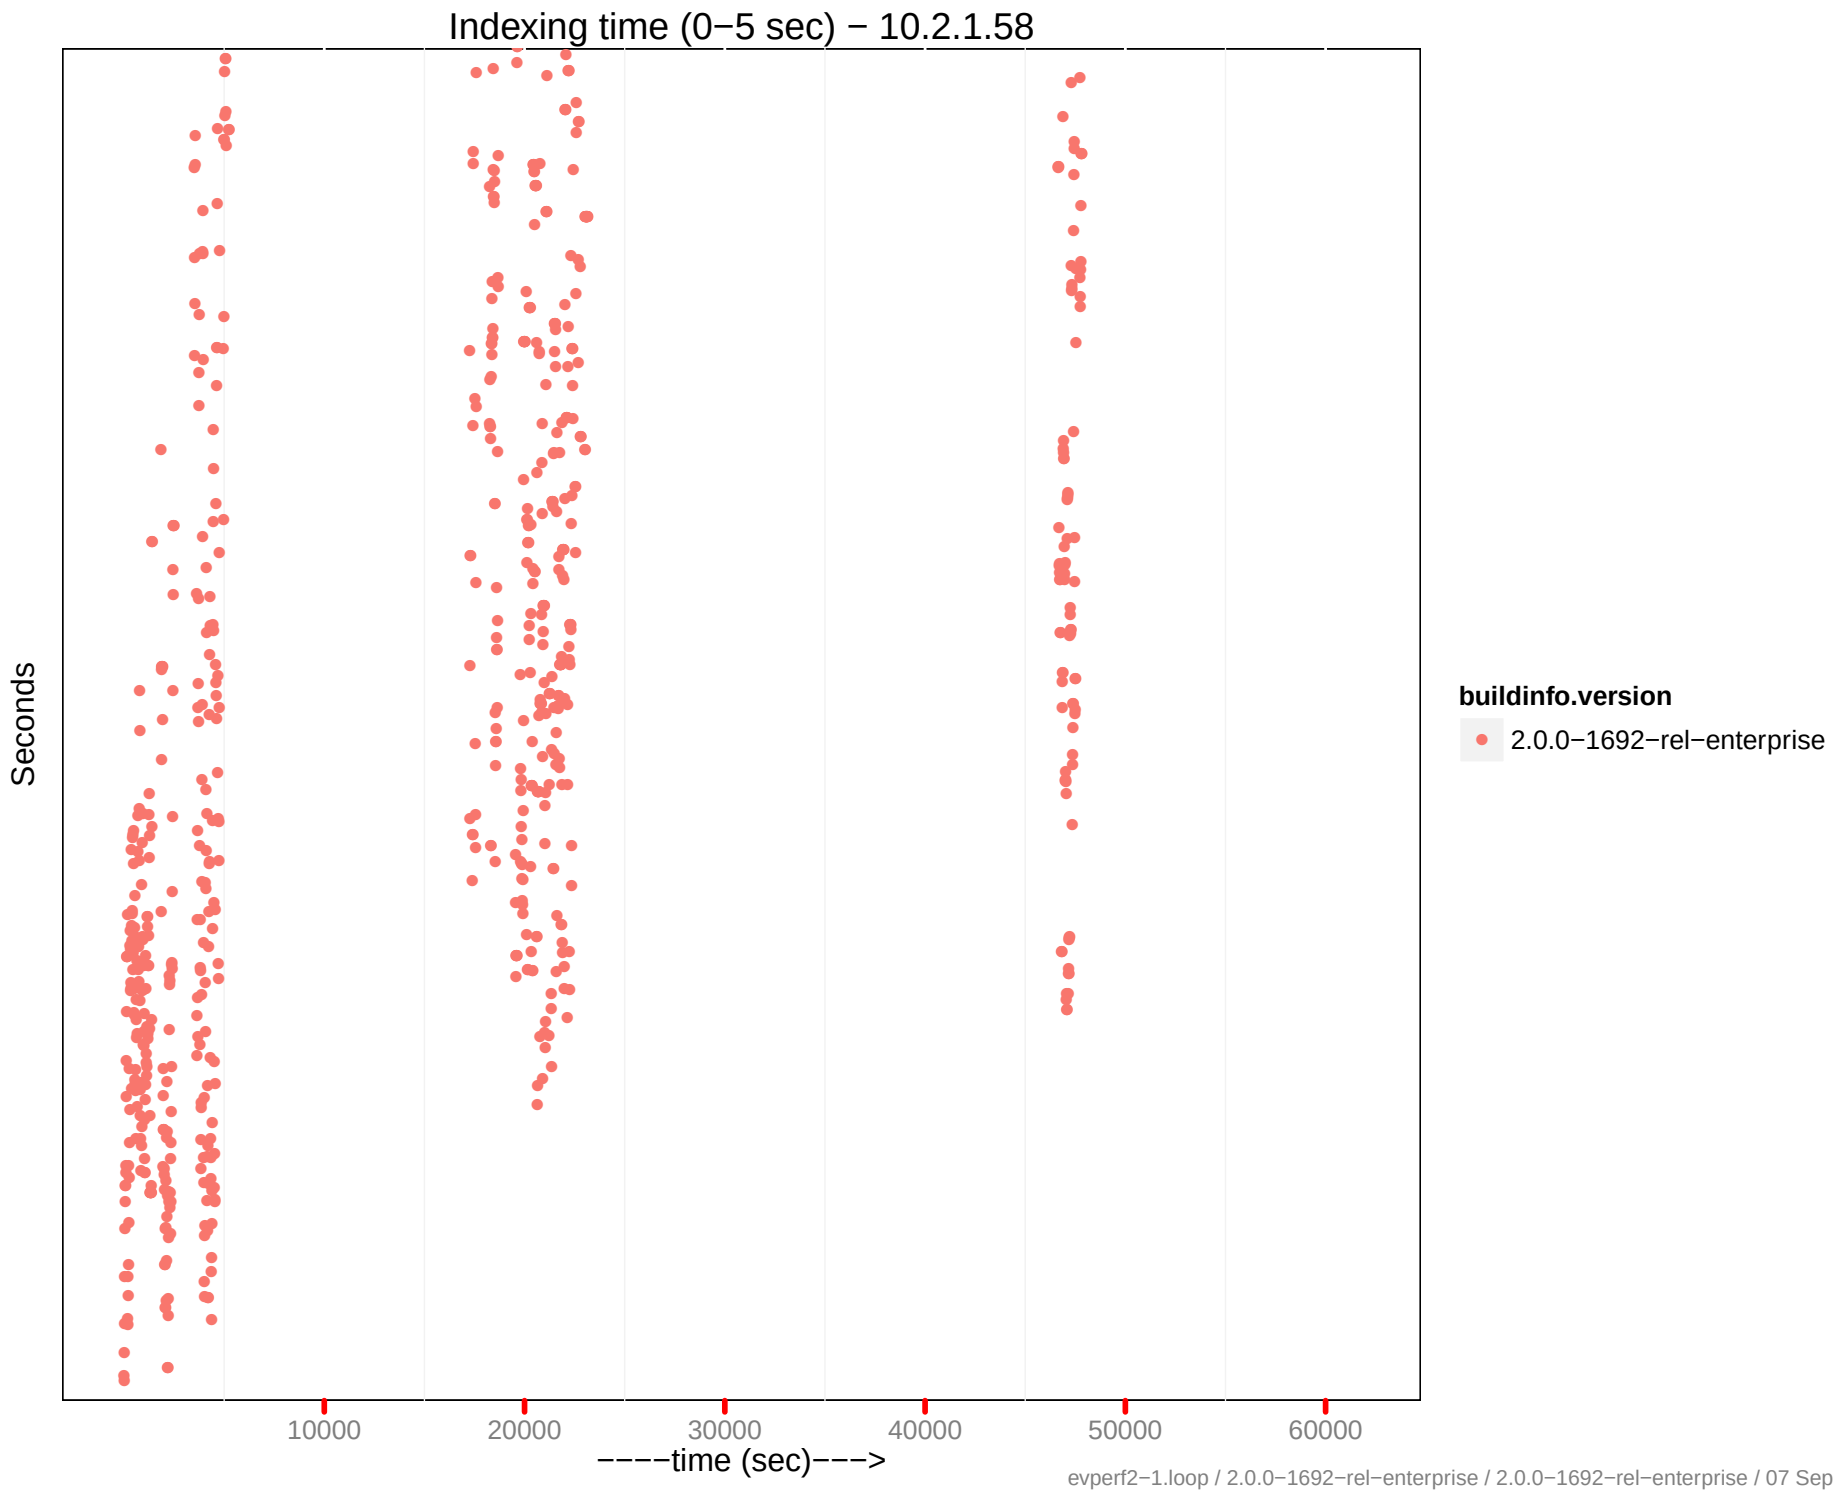
Indexing time (0−5 sec) − 10.2.1.58
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

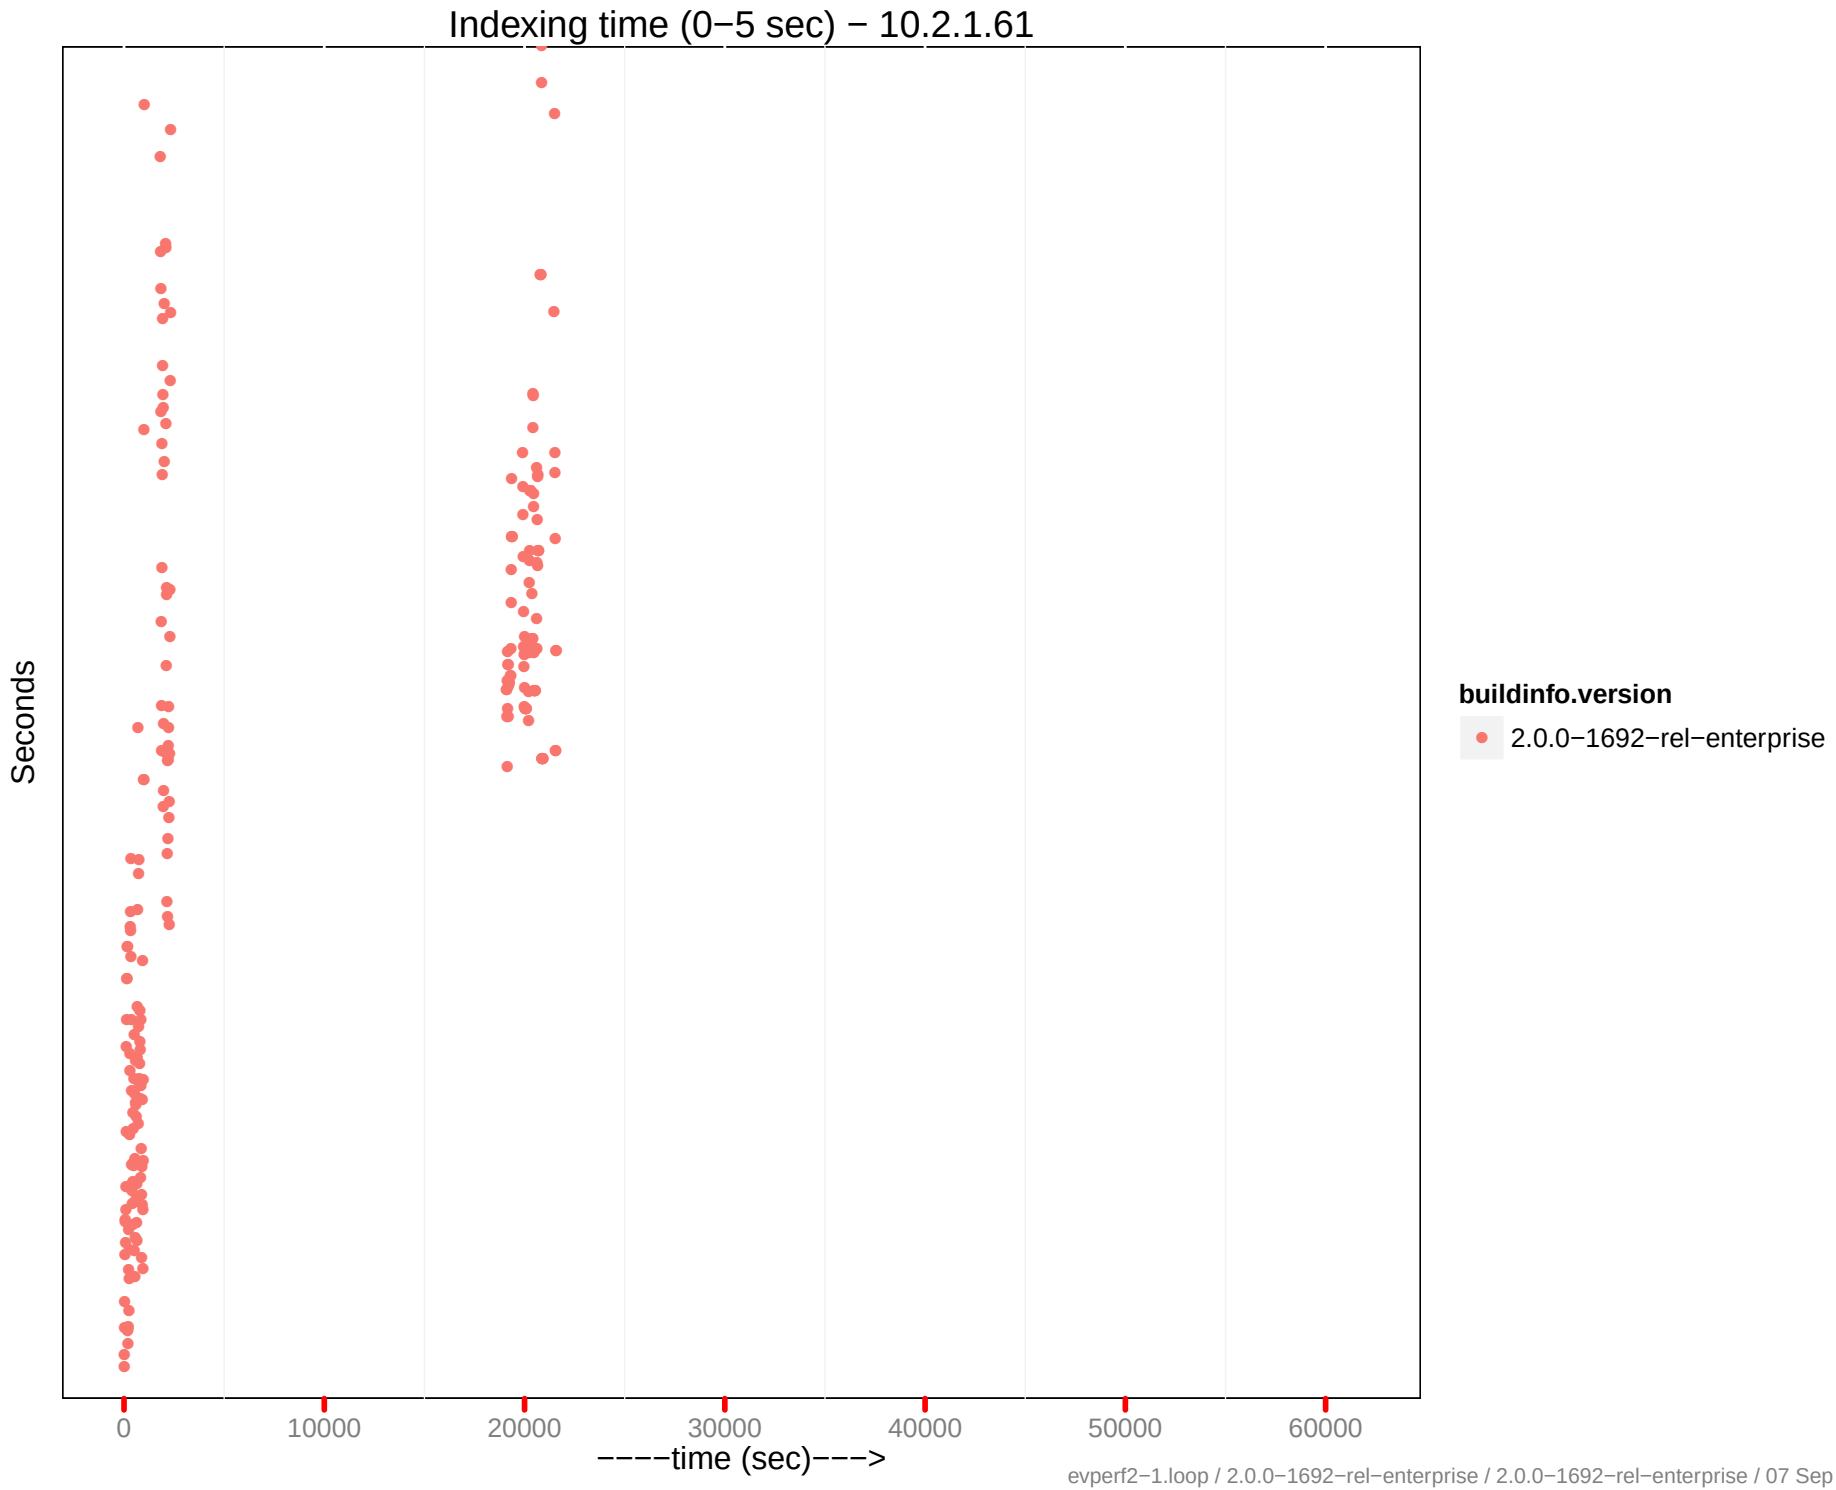
Indexing time (0−5 sec) − 10.2.1.61
Seconds
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

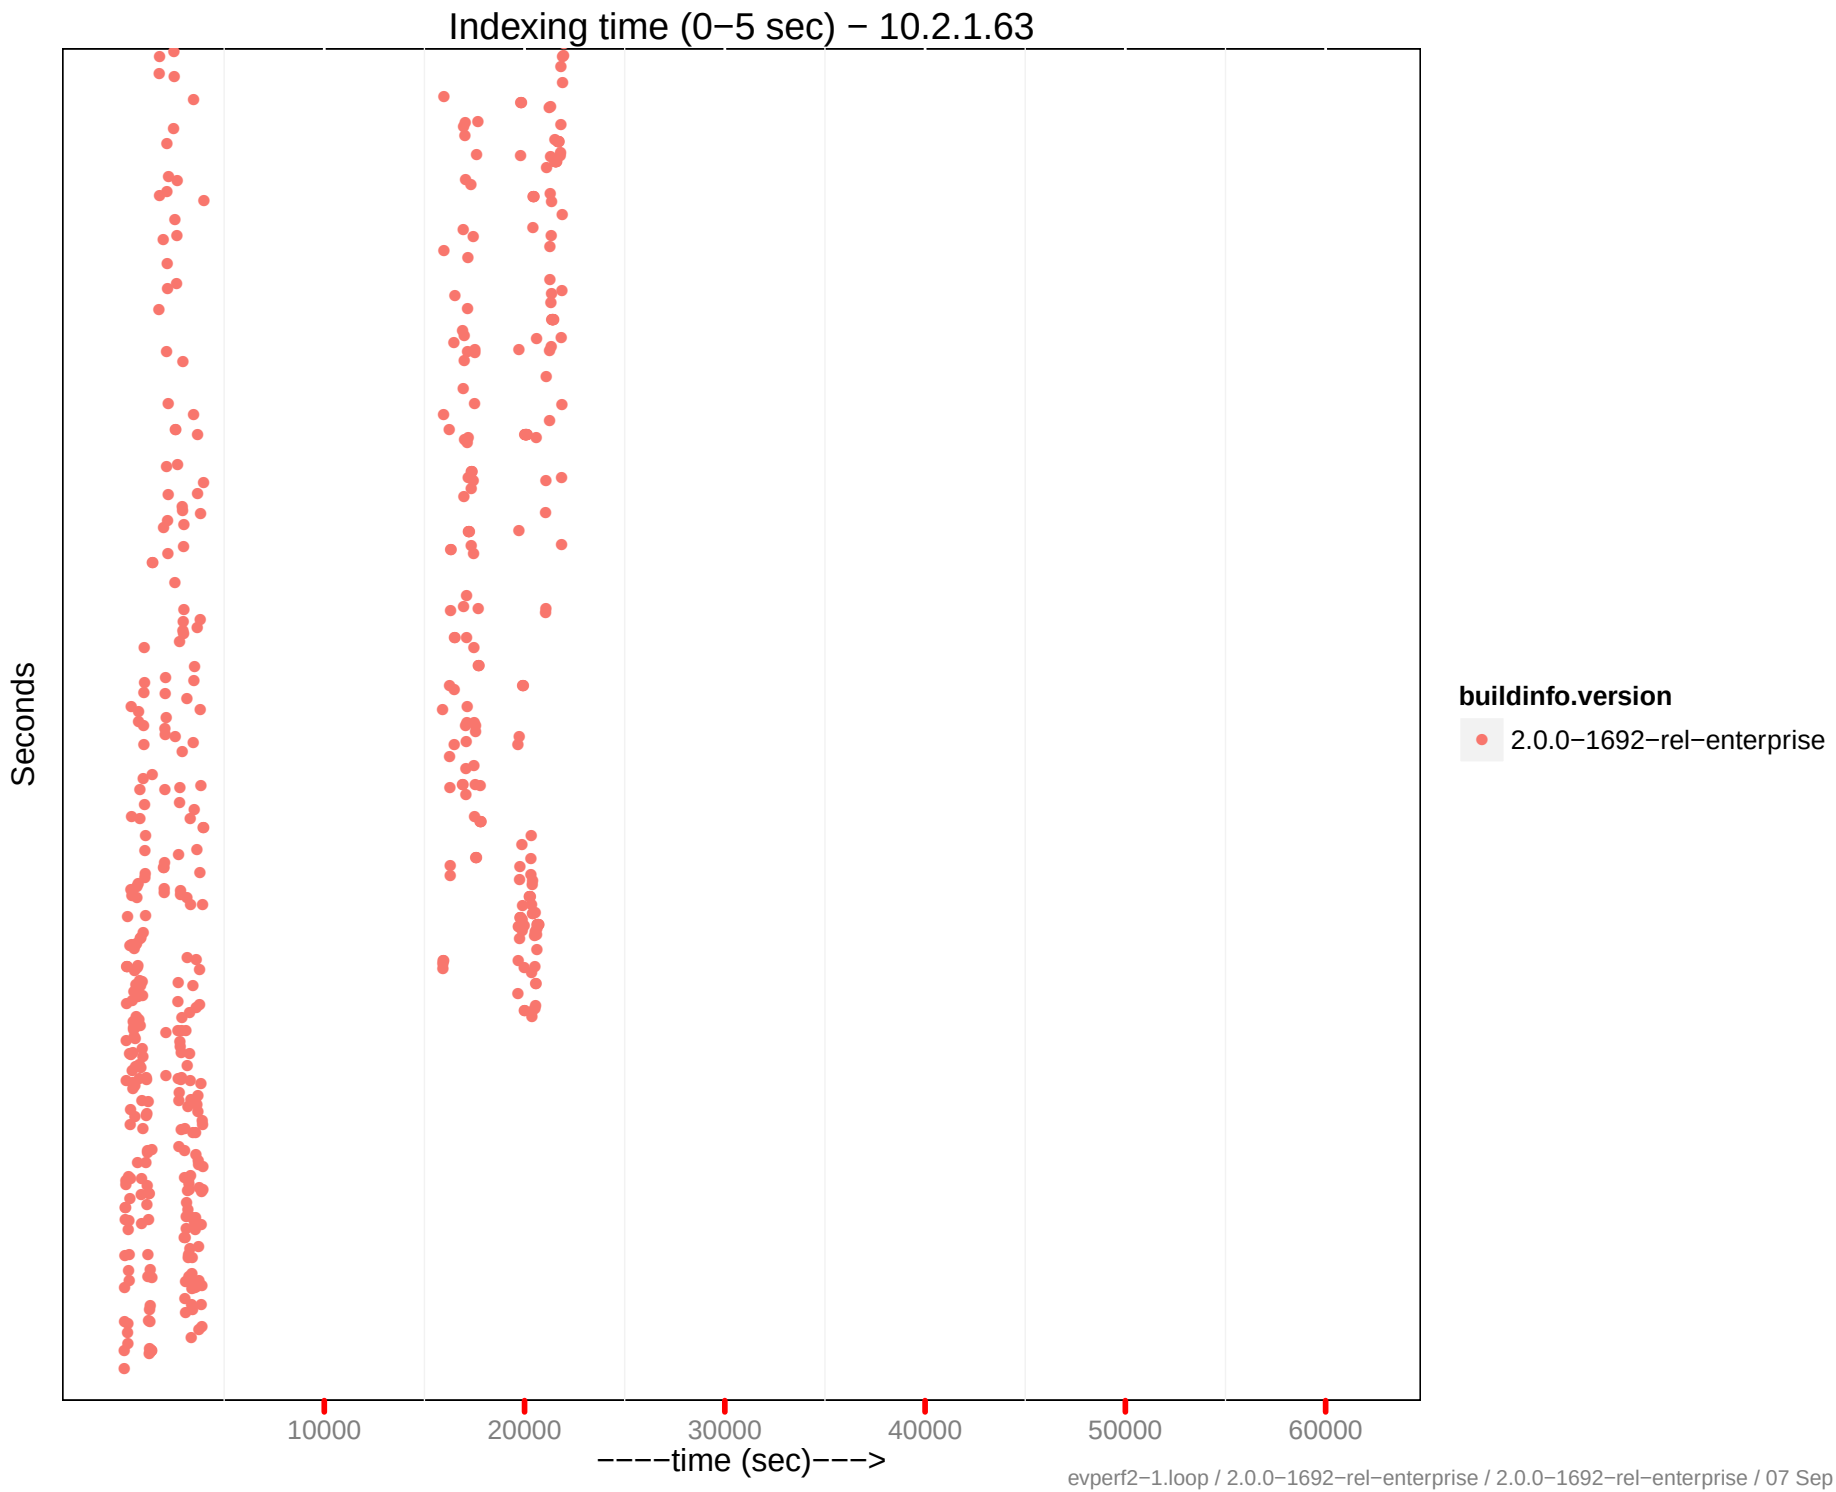
Indexing time (0−5 sec) − 10.2.1.63
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

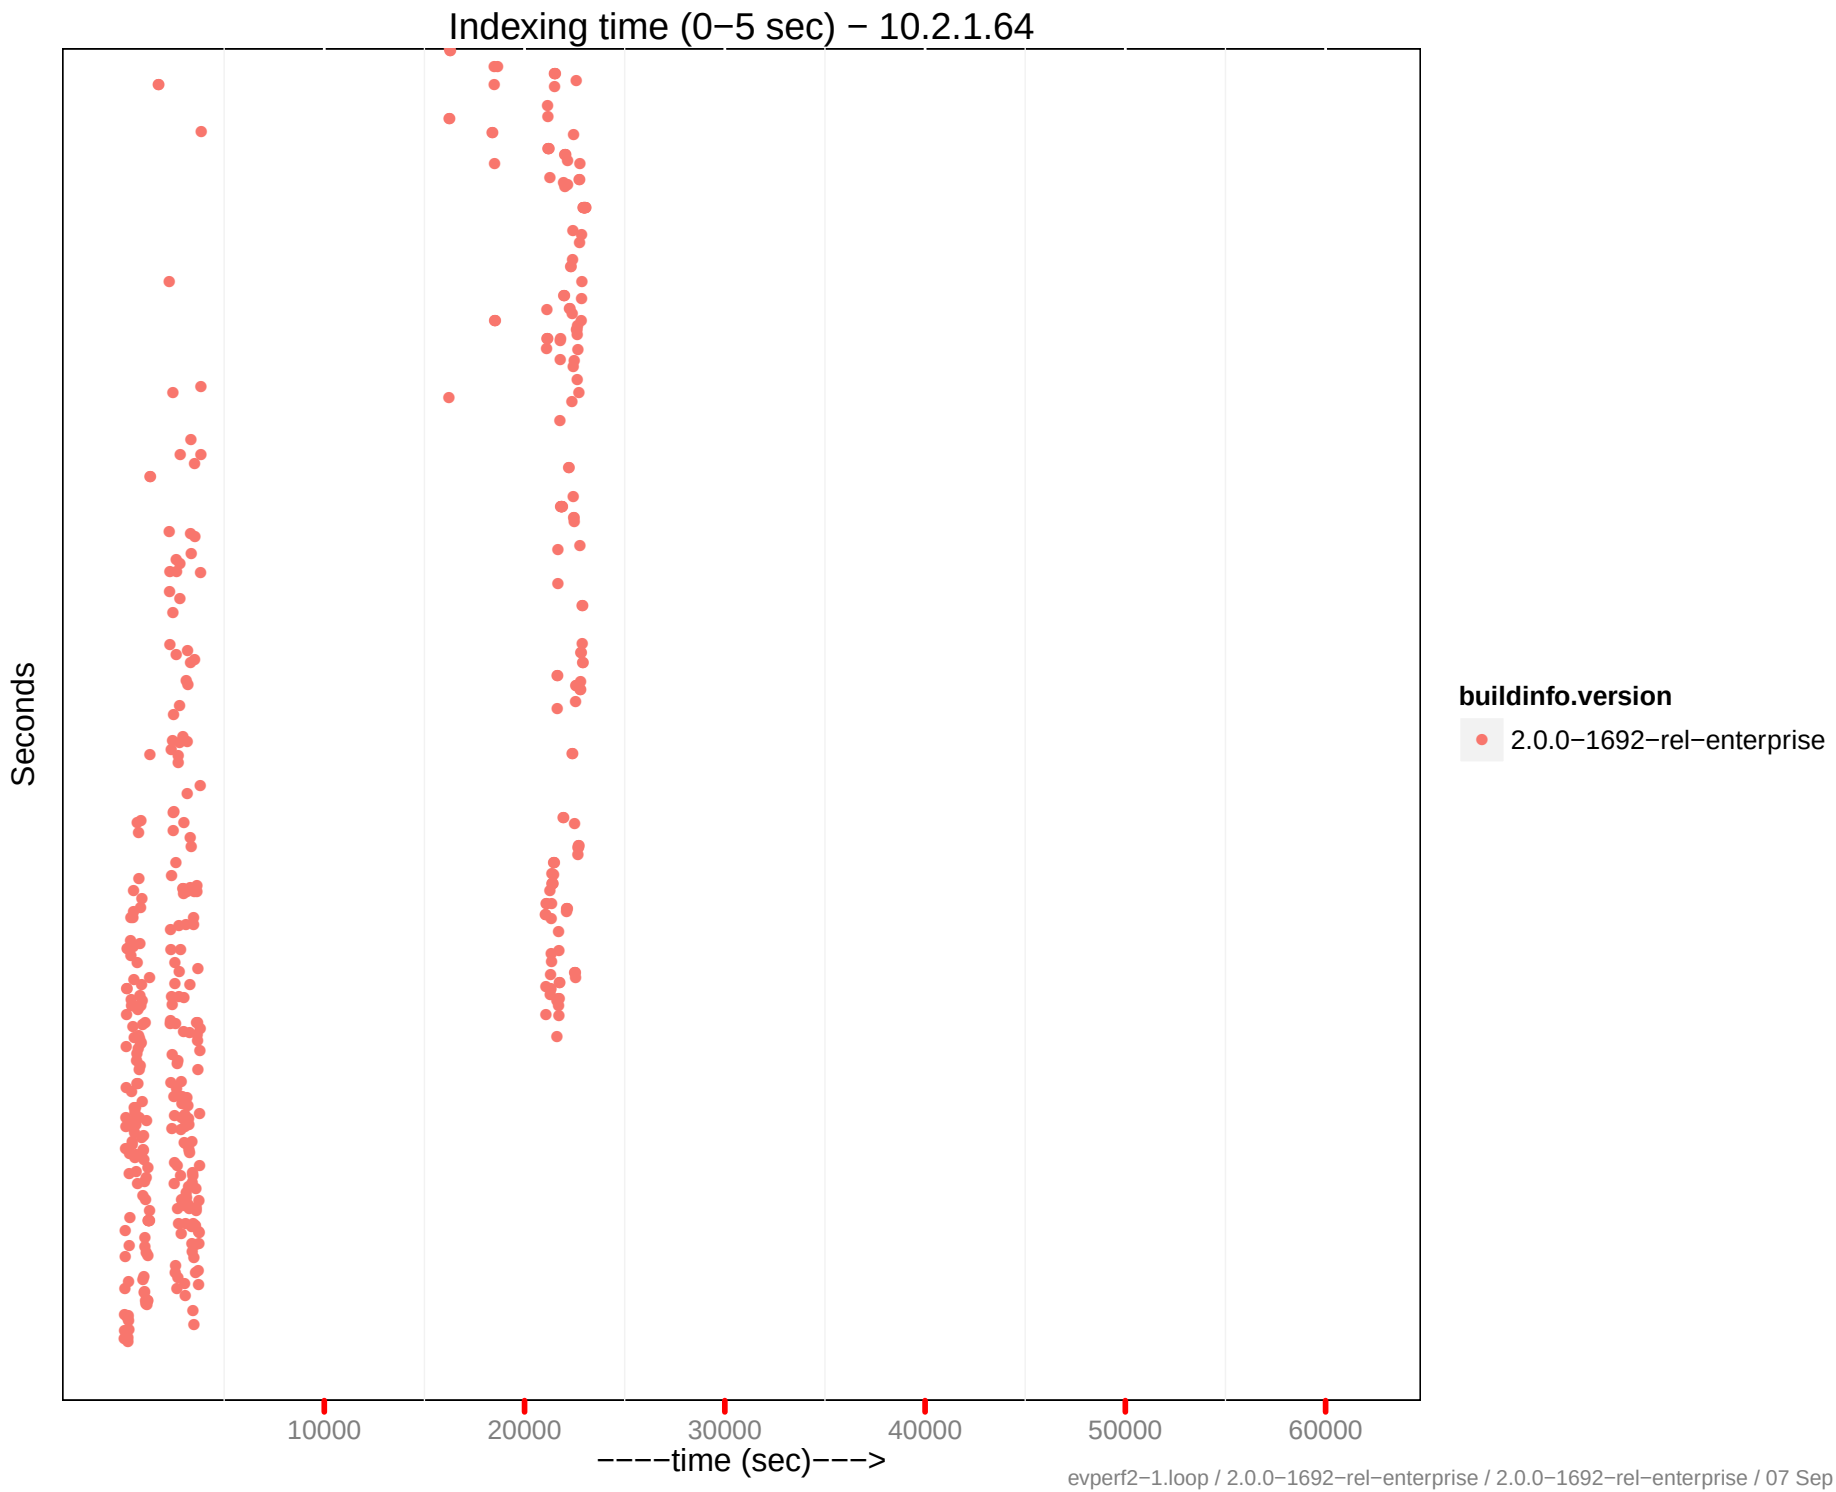
Indexing time (0−5 sec) − 10.2.1.64
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1692−rel−enterprise
evperf2−1.loop / 2.0.0−1692−rel−enterprise / 2.0.0−1692−rel−enterprise / 07 Sep 2012

evperf2-1.conf

```
# "EVPERF'2" view performance test:
# 3 ddocs with 8 views per ddoc
# 20M initial items
# 25GB RAM quota (32GB total RAM)
# DGM
# 30 clients
# 8K ops/sec total background load (memcached commands)
# 80% reads, 20% write (12% updates/deletes, 8% inserts)
# Cache miss ratio < 1%
# Stop on 18M total queries (tuned to be 6 hours)

performance.iperf.MultiClientTests.test_evperf2

params:

# general
batch=50
kind=json
mem_quota=25000
loglevel=error

# load phase
items=20000000
hot_init_items=16000000

# 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=0.975
ratio_hot_sets=0.975
ratio_expirations=0.0
bg_max_ops_per_sec=265
fg_max_ops=18000000
total_clients=30
start_delay=5

# 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
2:10.2.1.58
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
```