	2.0.0 − 1916	2.0.0 − 1916
Runtime (in hr)	18.09	NA
Avg. Drain Rate	363.69	NANA
Peak Disk (GB)	140.88	NA
Peak Memory (GB)	20.14	NA
Avg. OPS	7.69K	NANA
Avg. mem memcached (GB)	19.66	NA
Avg. mem beam.smp (MB)	352.47	NA
Avg. CPU rate (%)	85.08	NA
Latency−get (90th) (ms)	1.47	NA
Latency−get (95th) (ms)	2.78	NA
Latency−get (99th) (ms)	7.53	NA
Latency−set (90th) (ms)	1.48	NA
Latency−set (95th) (ms)	2.73	NA
Latency−set (99th) (ms)	6.12	NA
Latency−query (80th) (ms)	49.9	NA
Latency−query (90th) (ms)	62.23	NA
Latency−query (95th) (ms)	80.43	NA
Latency−query (99th) (ms)	397.98	NA
Latency−query (99.9th) (ms)	1462.46	NA
Avg. QPS	435.72	NA
Avg. XDC ops/sec	NaN	NA
Avg. XDC queue	NaN	NA
Rebalance Time (sec)	0	NA
Testrunner Version	9acf95c	NA

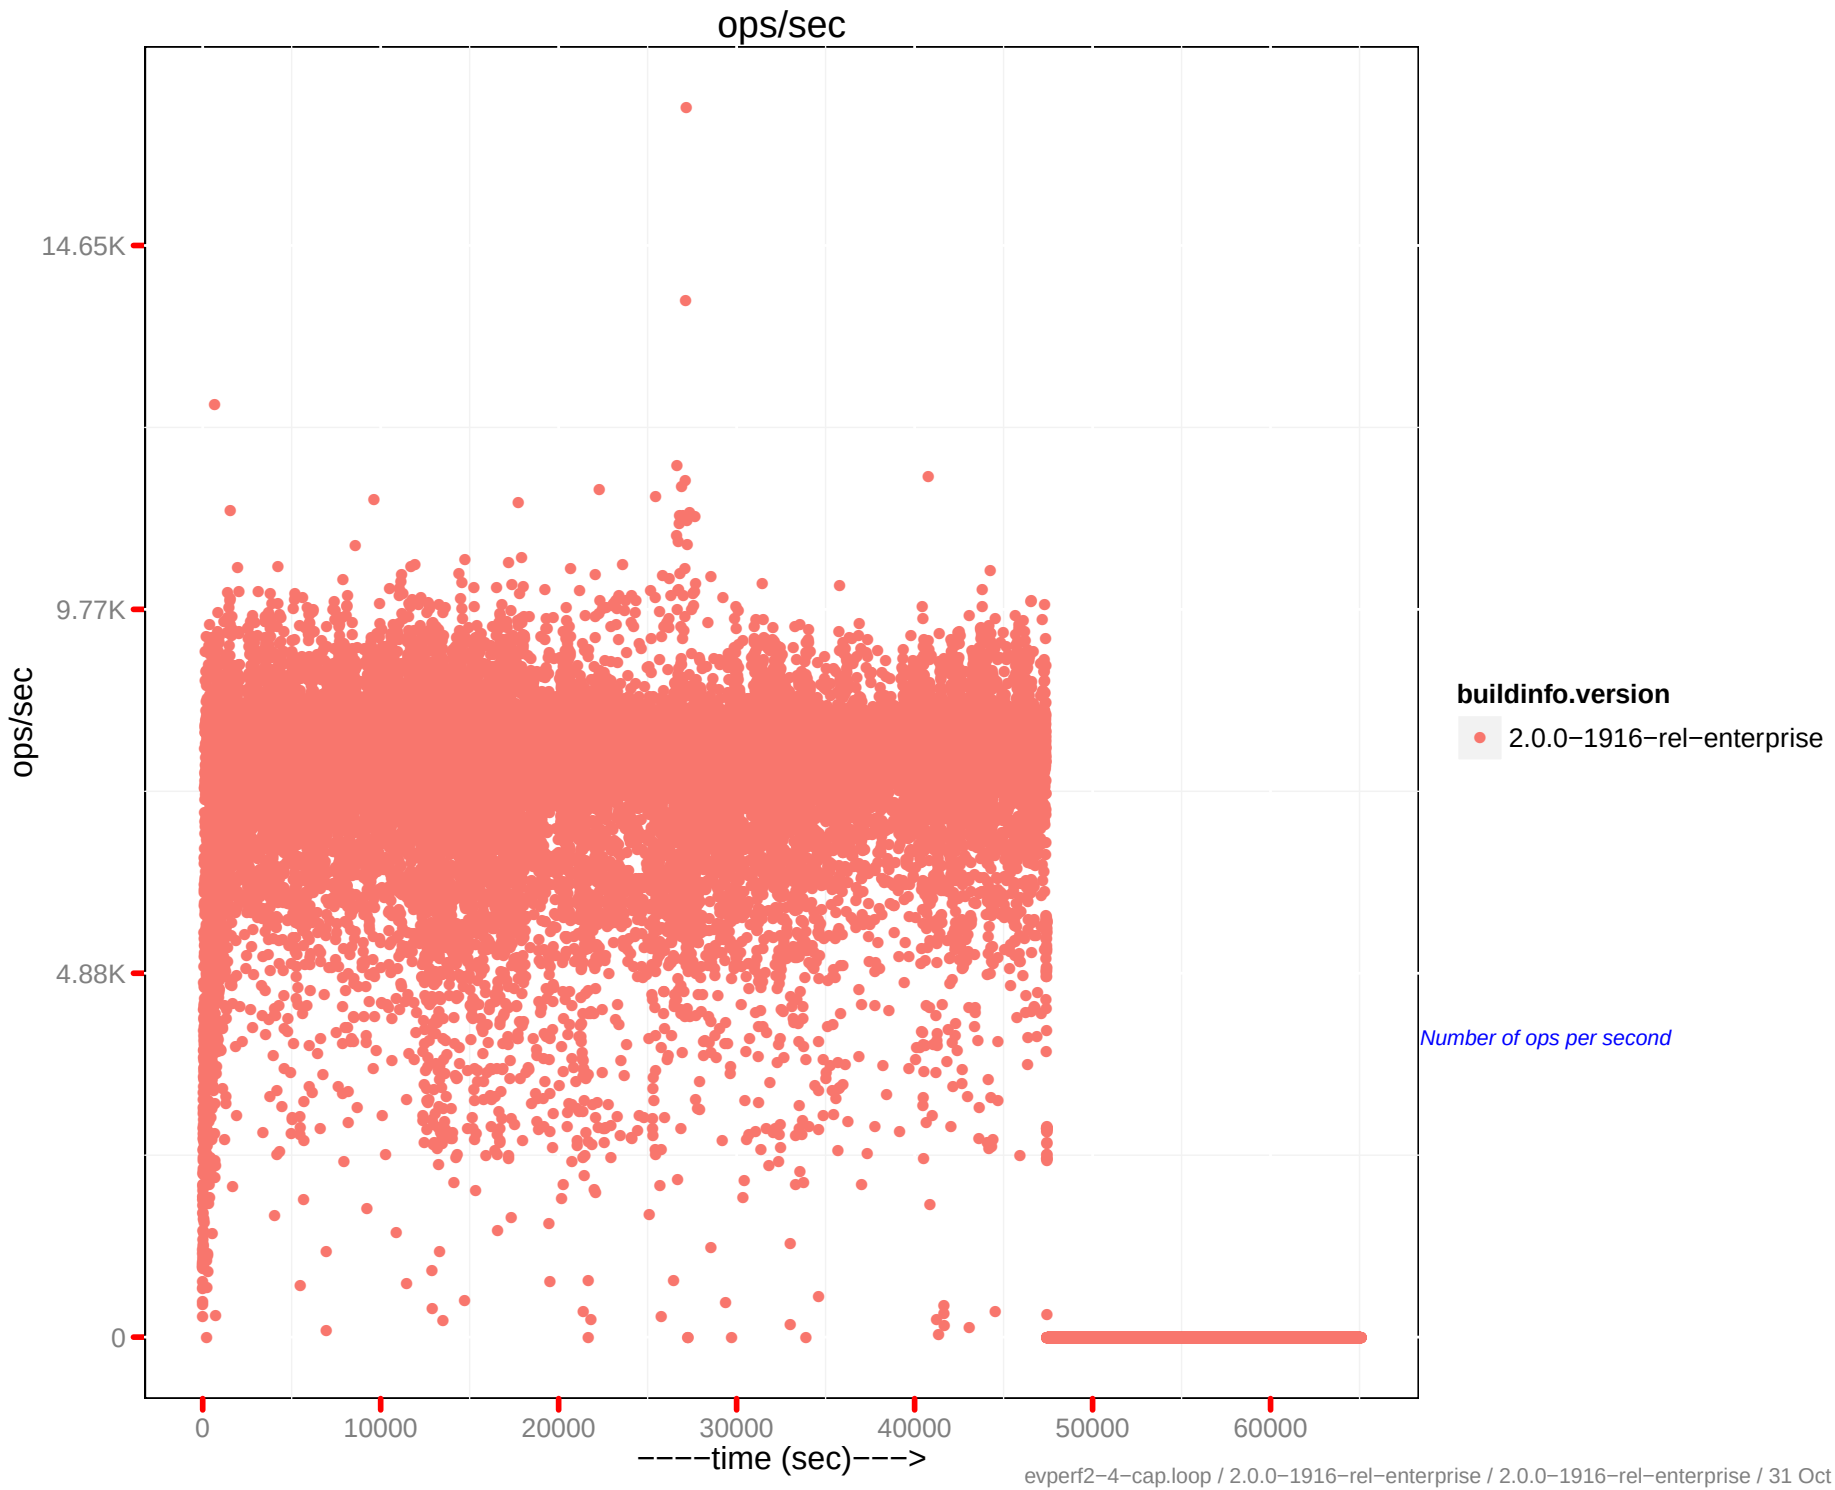
ops/sec
14.65K
9.77K
4.88K
0
ops/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of ops per second
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

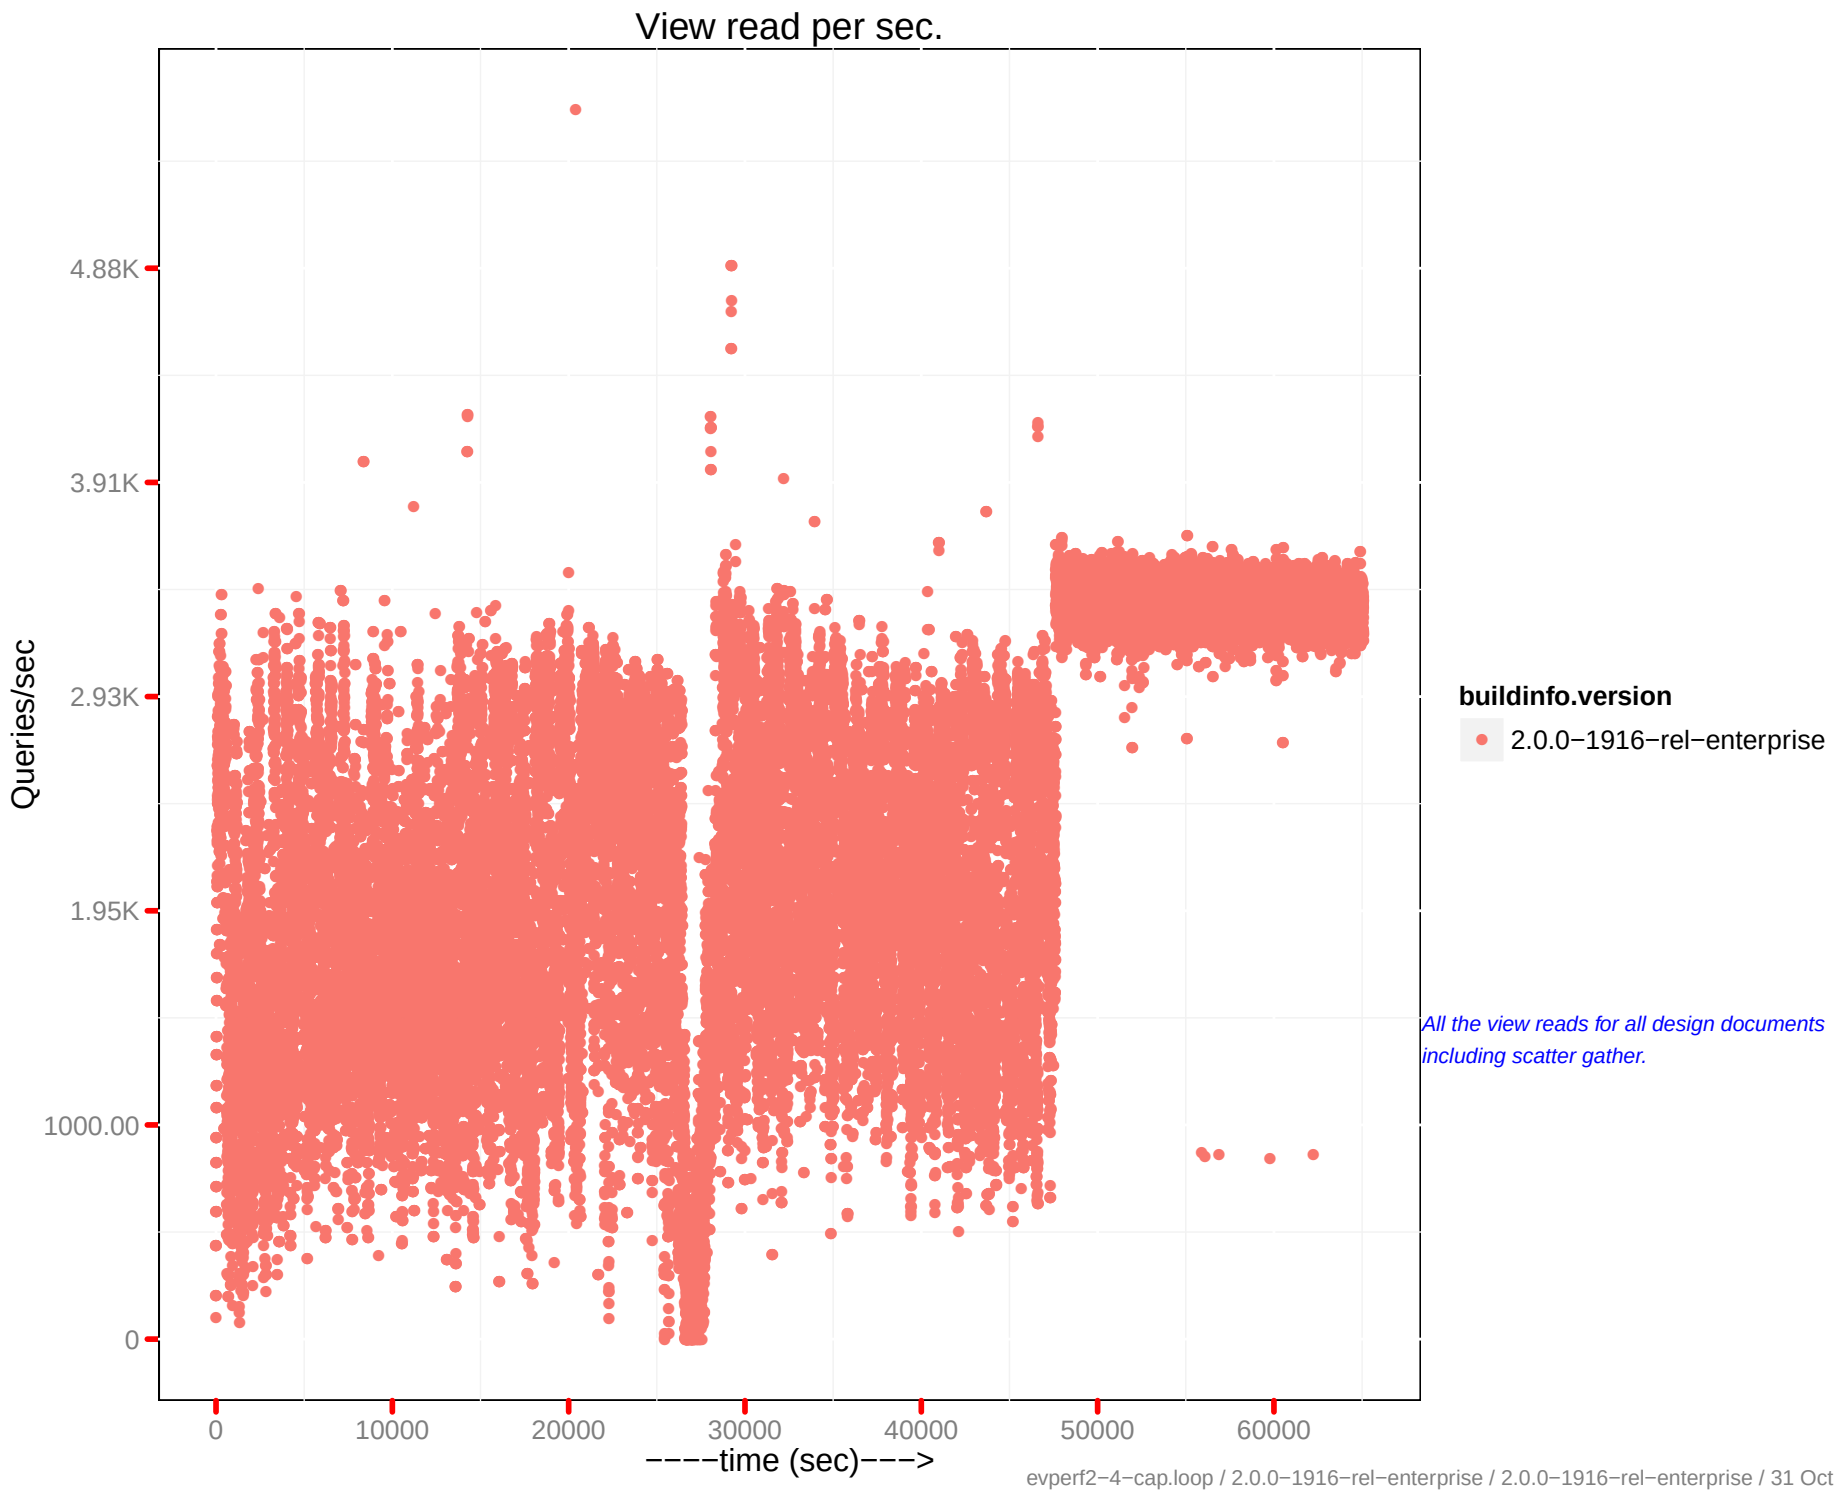
View read per sec.
Queries/sec
4.88K
3.91K
2.93K
1.95K
1000.00
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
All the view reads for all design documents including scatter gather.
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

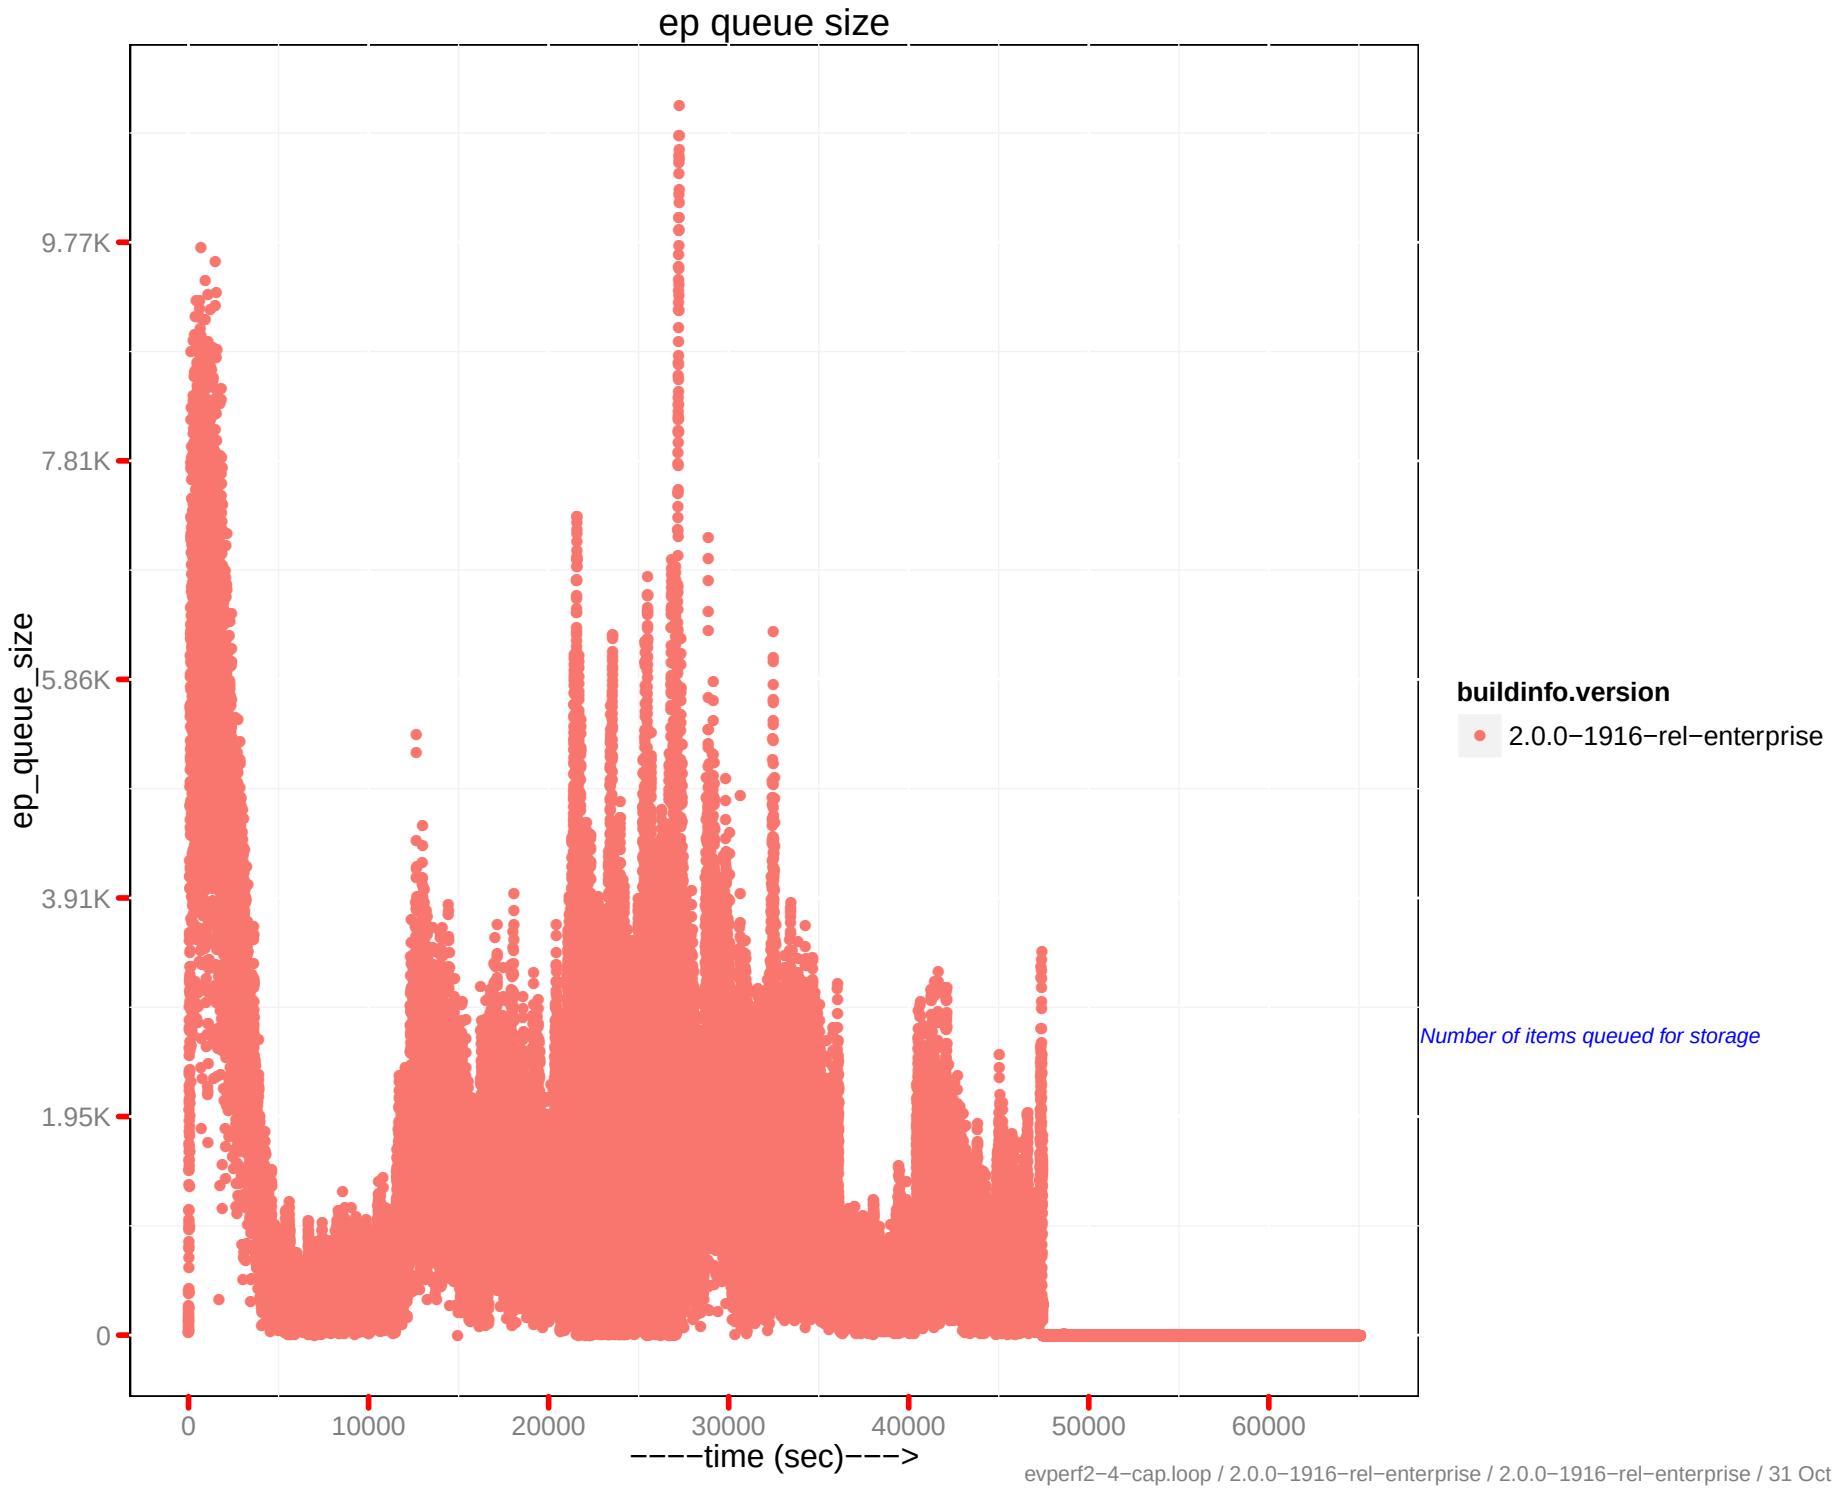
ep queue size
9.77K
7.81K
5.86K
3.91K
1.95K
0
ep_queue_size
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of items queued for storage
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

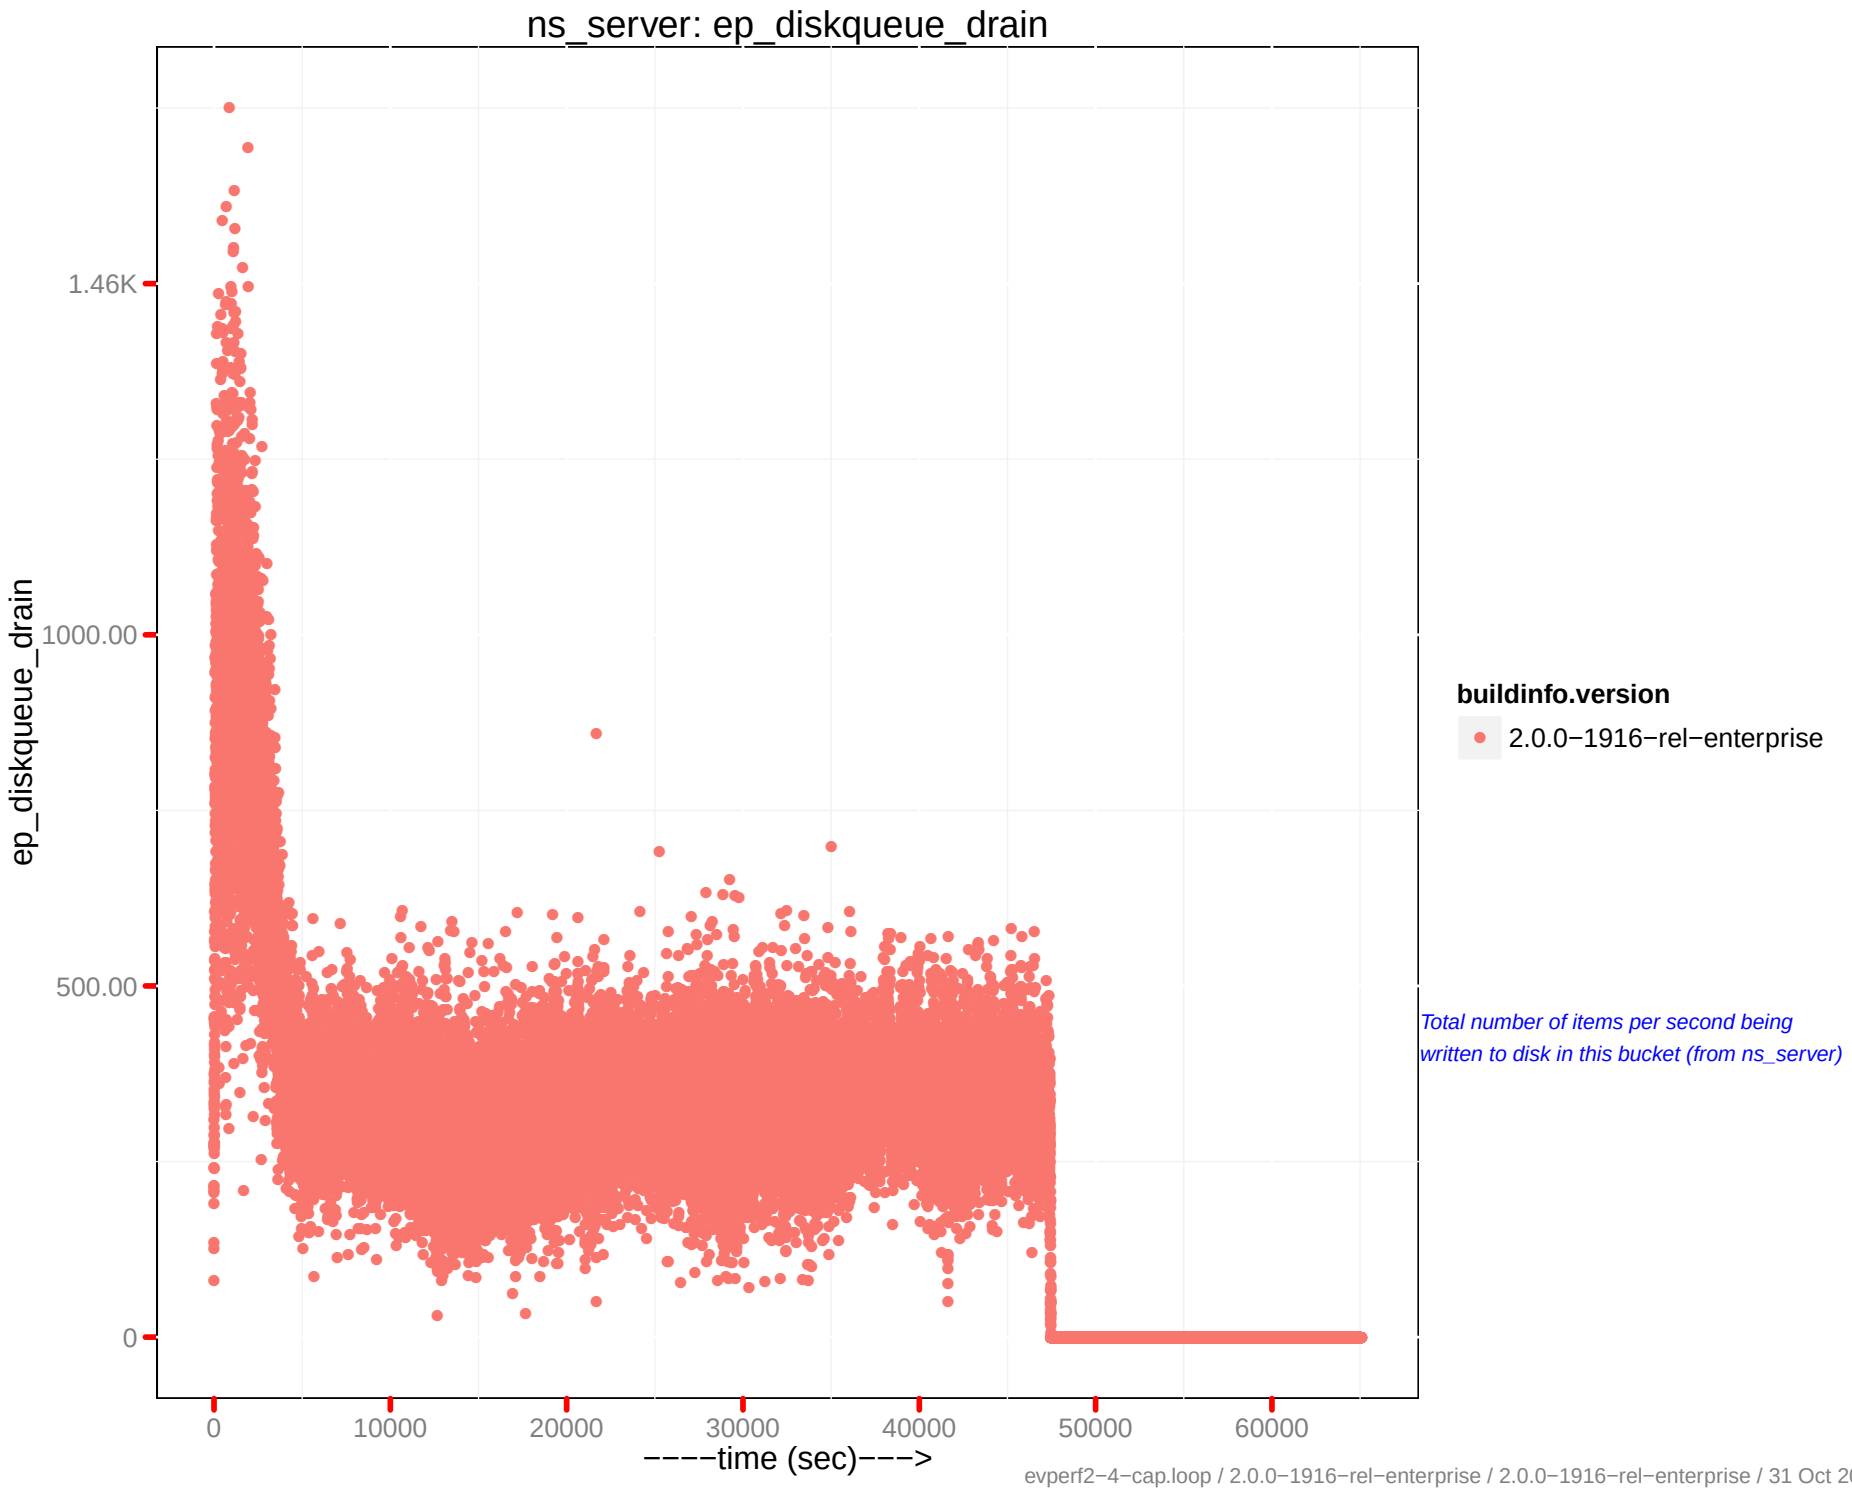
ns_server: ep_diskqueue_drain
ep_diskqueue_drain
1.46K
1000.00
500.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Total number of items per second being written to disk in this bucket (from ns_server)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

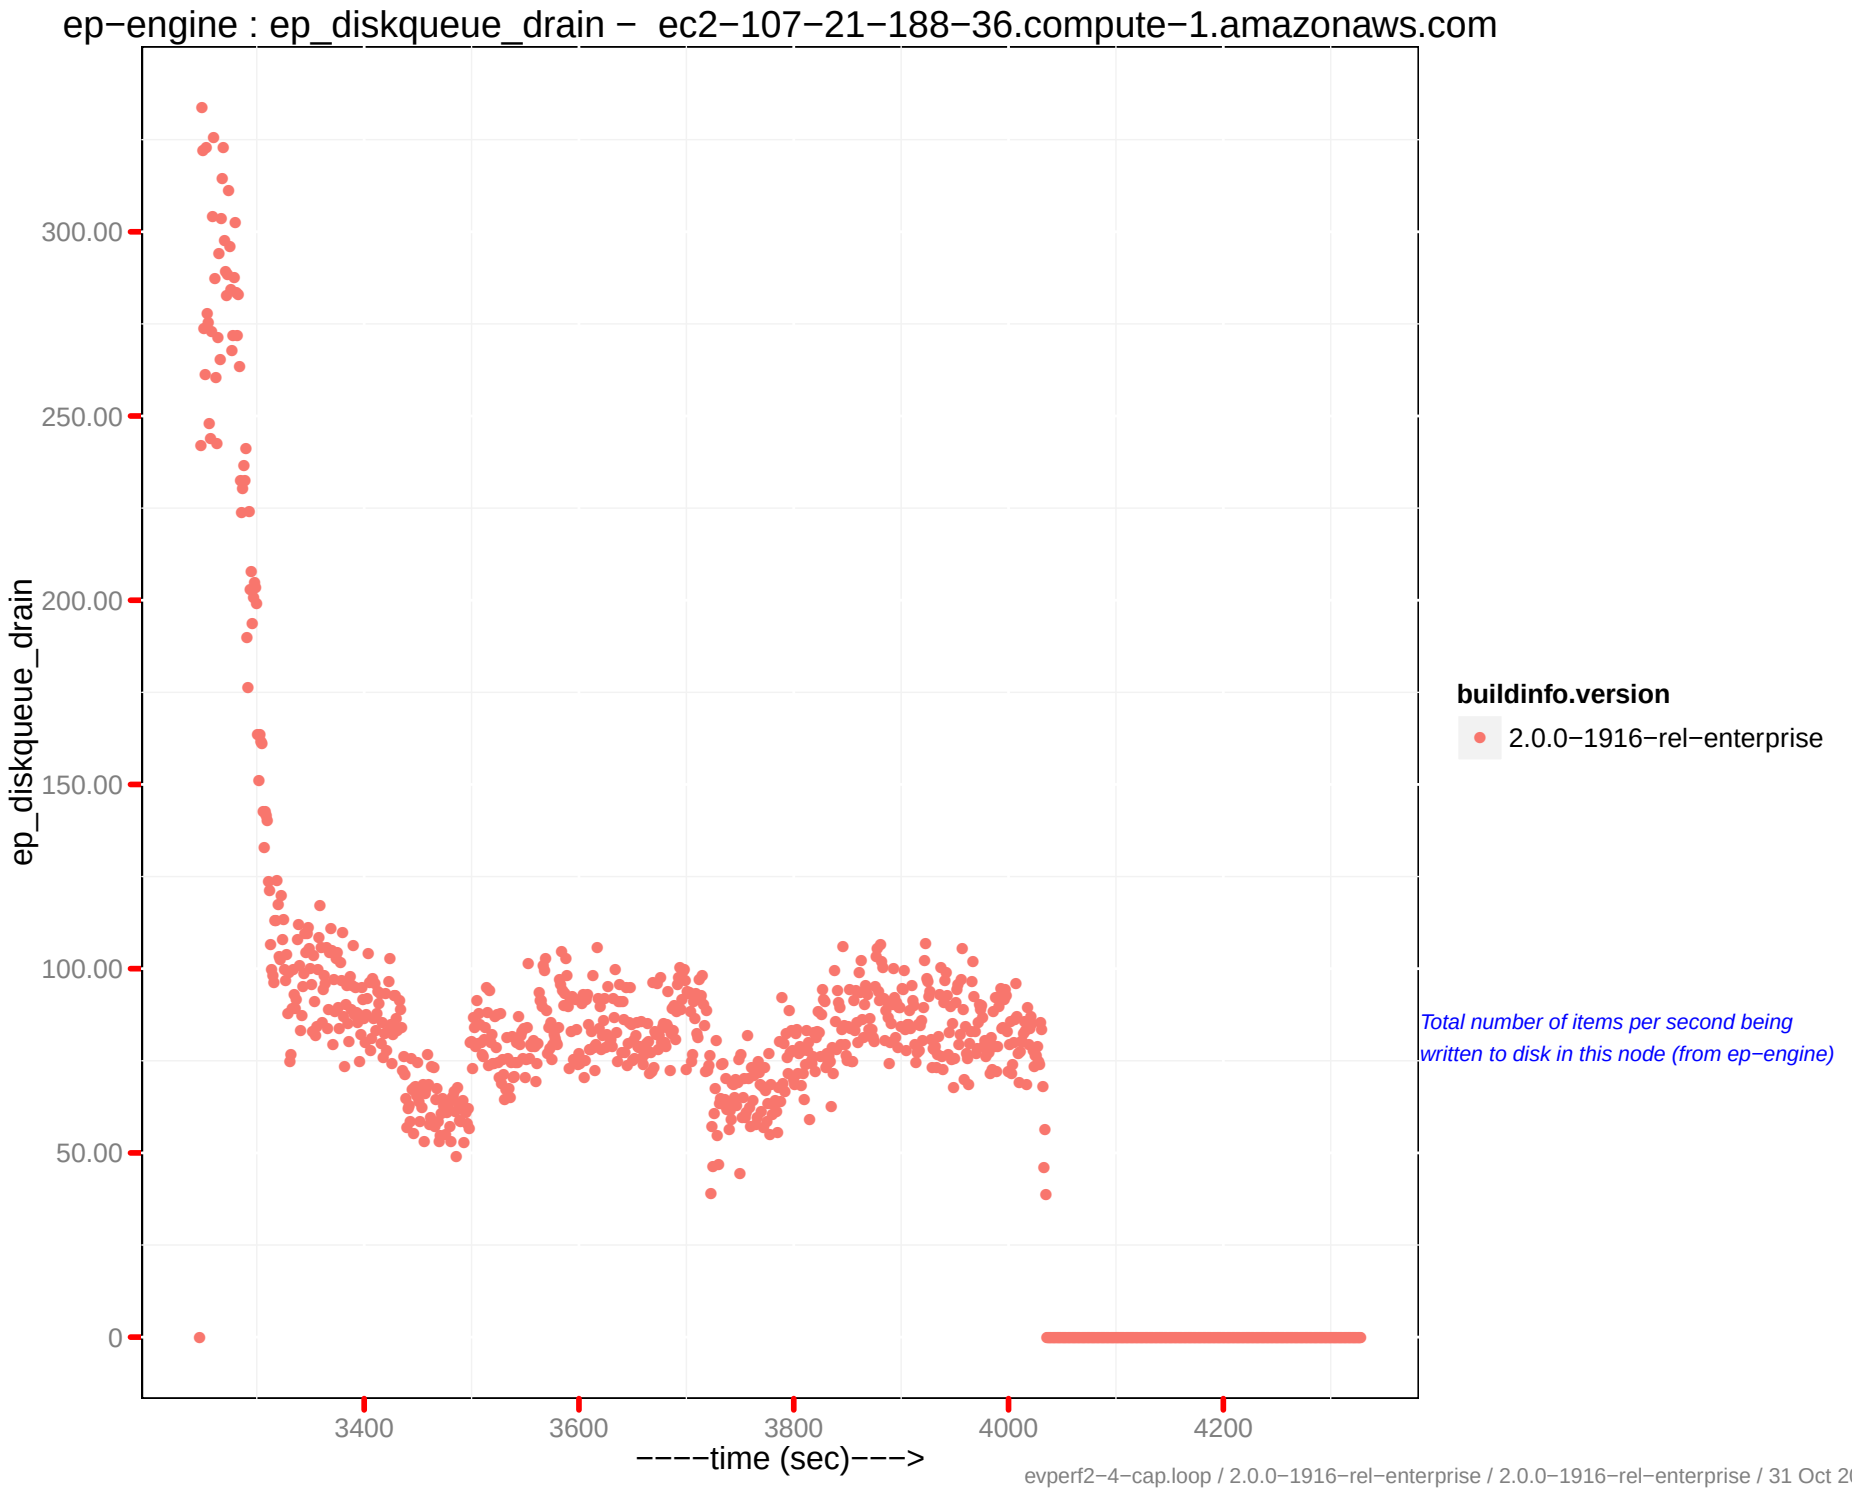
ep−engine : ep_diskqueue_drain − ec2−107−21−188−36.compute−1.amazonaws.com
ep_diskqueue_drain
300.00
250.00
200.00
150.00
100.00
50.00
0
3400
3600
3800
4000
4200
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

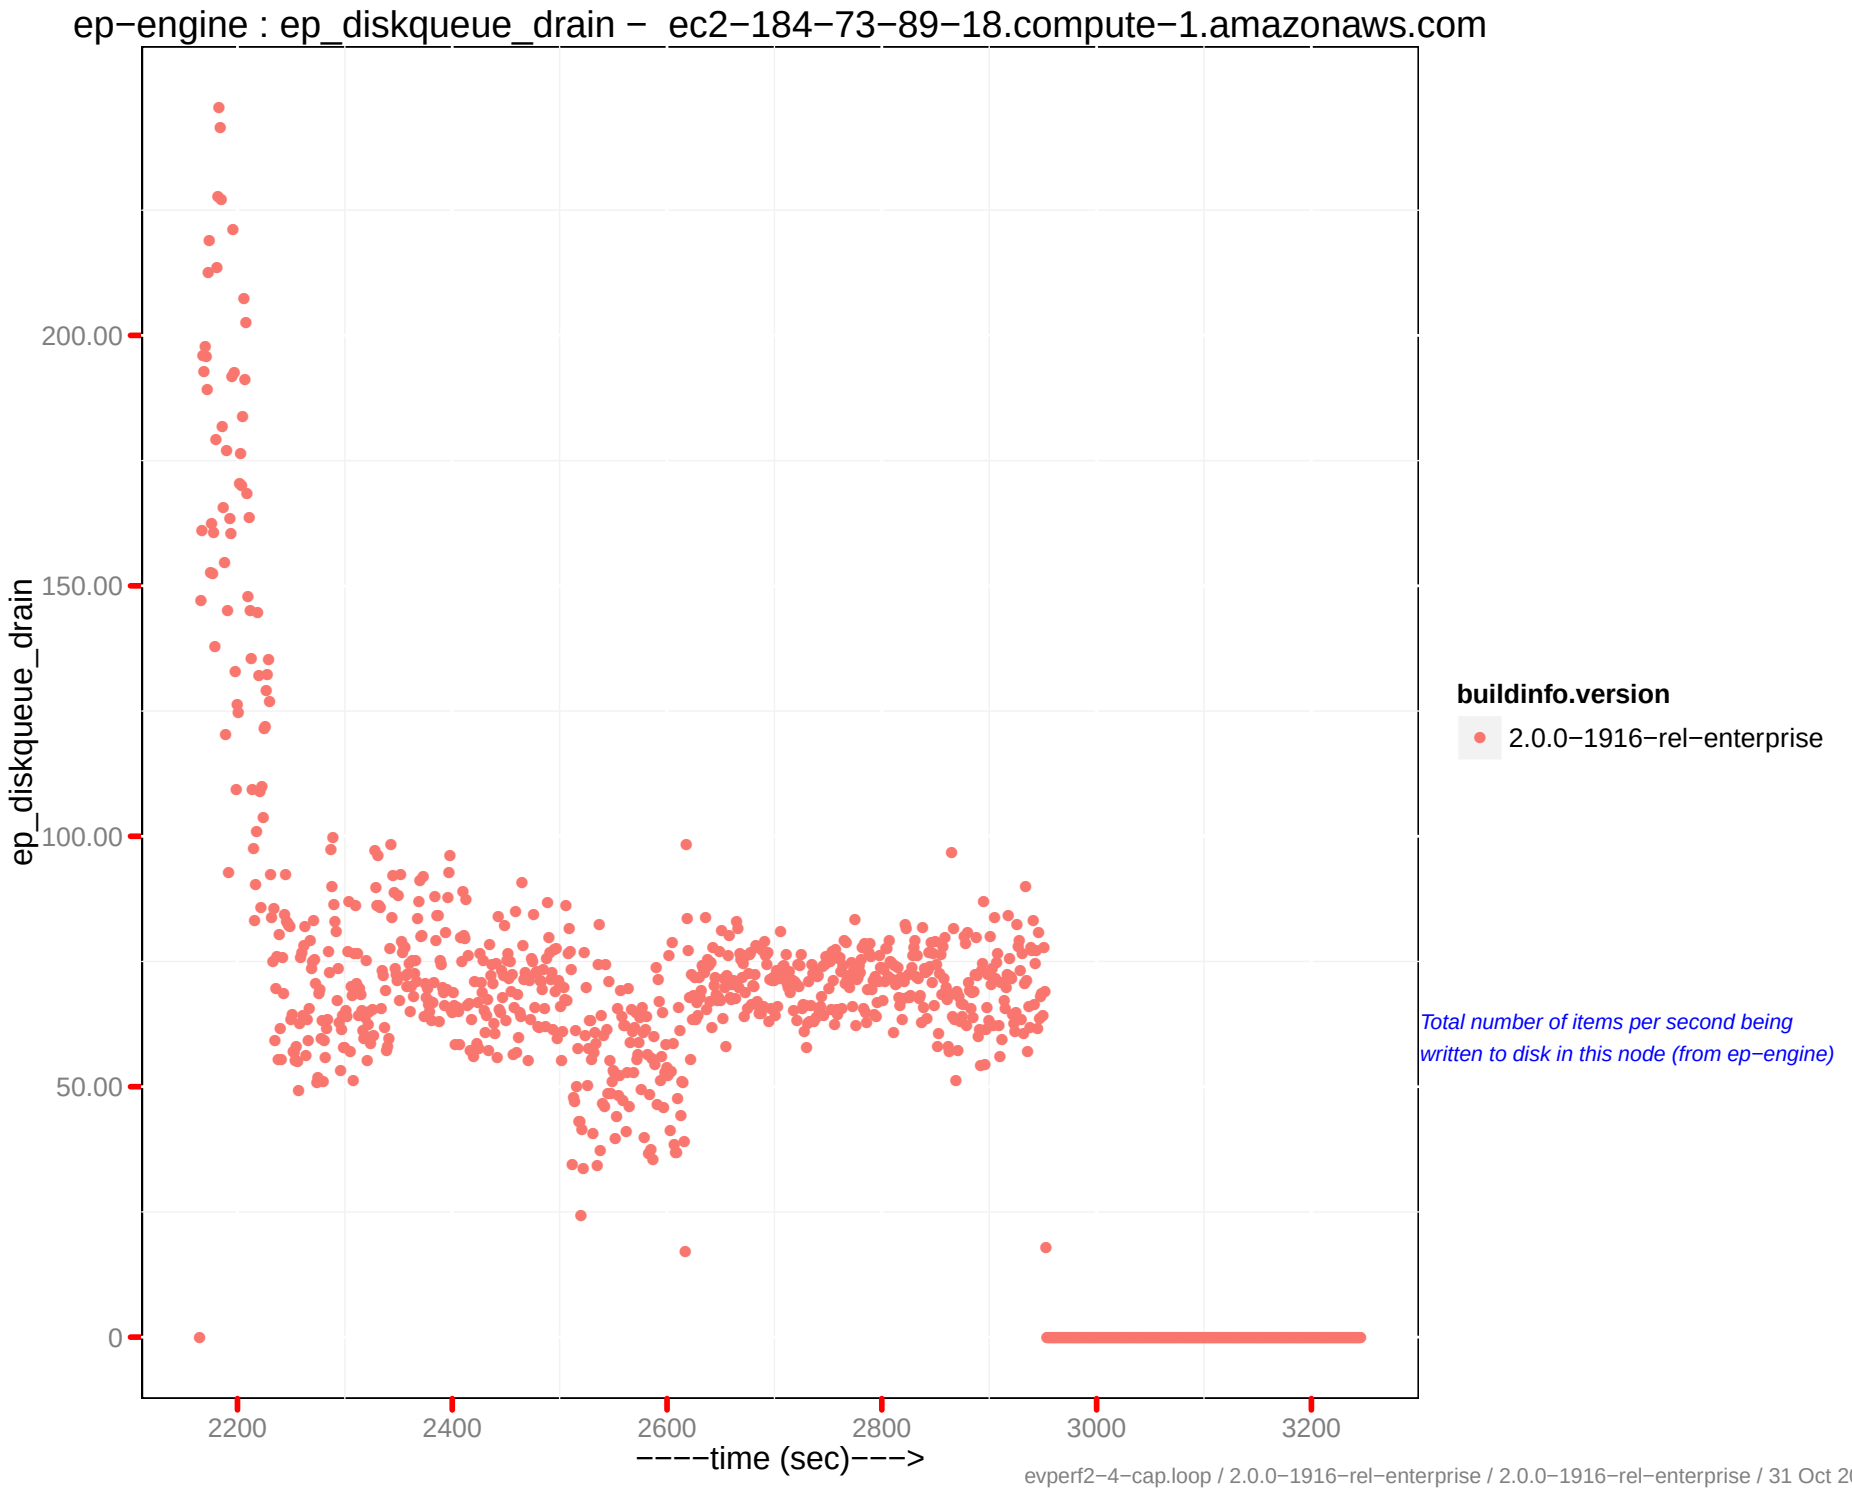
ep−engine : ep_diskqueue_drain − _ec2−184−73−89−18.compute−1.amazonaws.com
ep_diskqueue_drain
200.00
150.00
100.00
50.00
0
2200
2400
2600
2800
3000
3200
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

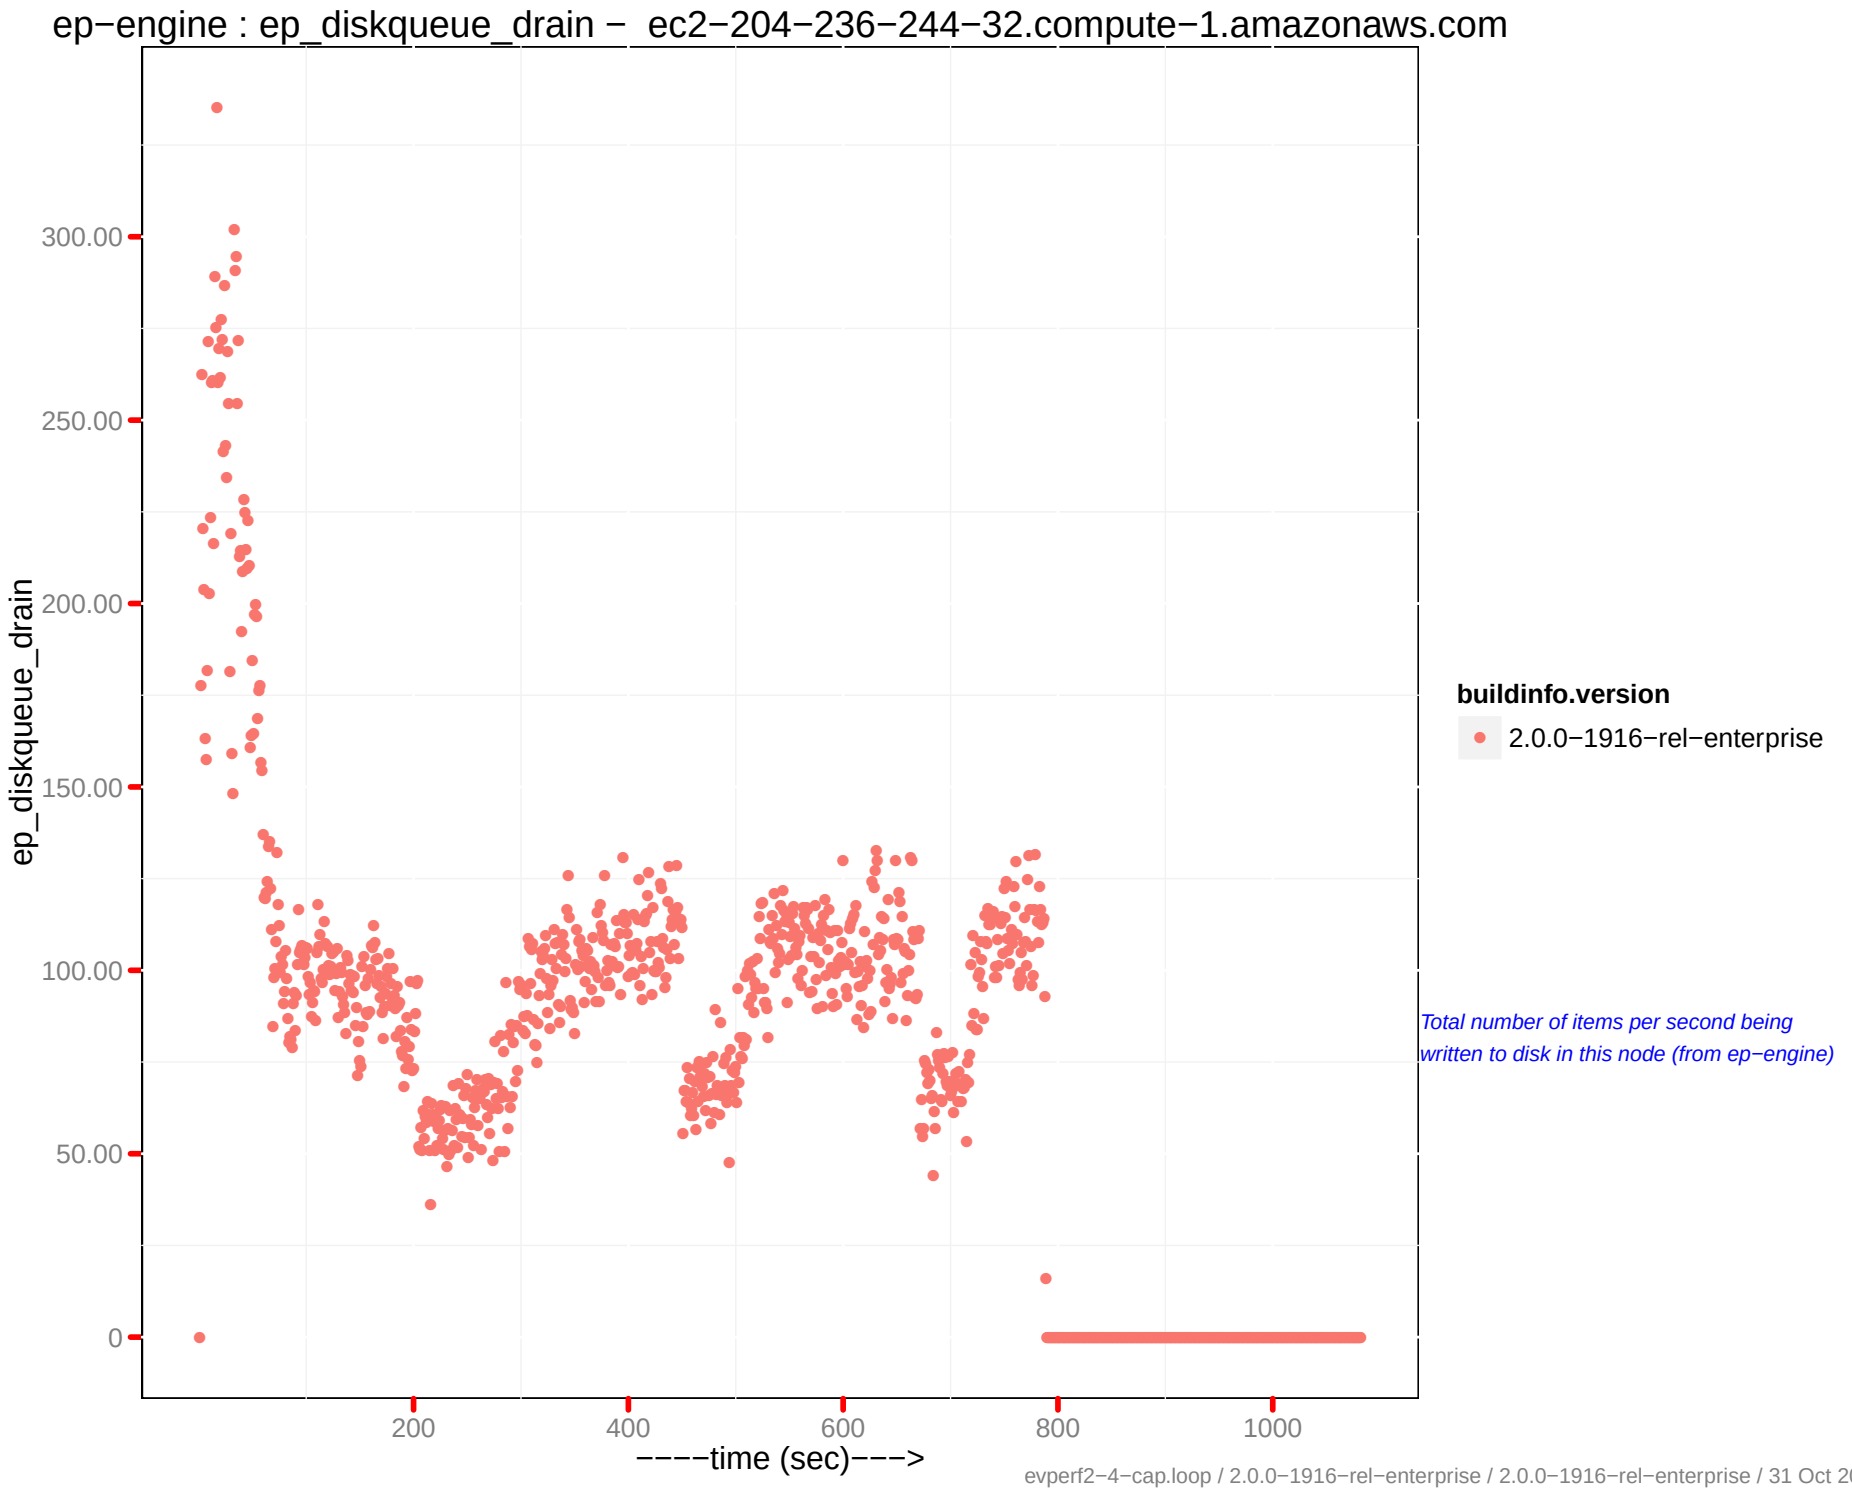
ep−engine : ep_diskqueue_drain − ec2−204−236−244−32.compute−1.amazonaws.com
ep_diskqueue_drain
300.00
250.00
200.00
150.00
100.00
50.00
0
200
400
600
800
1000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

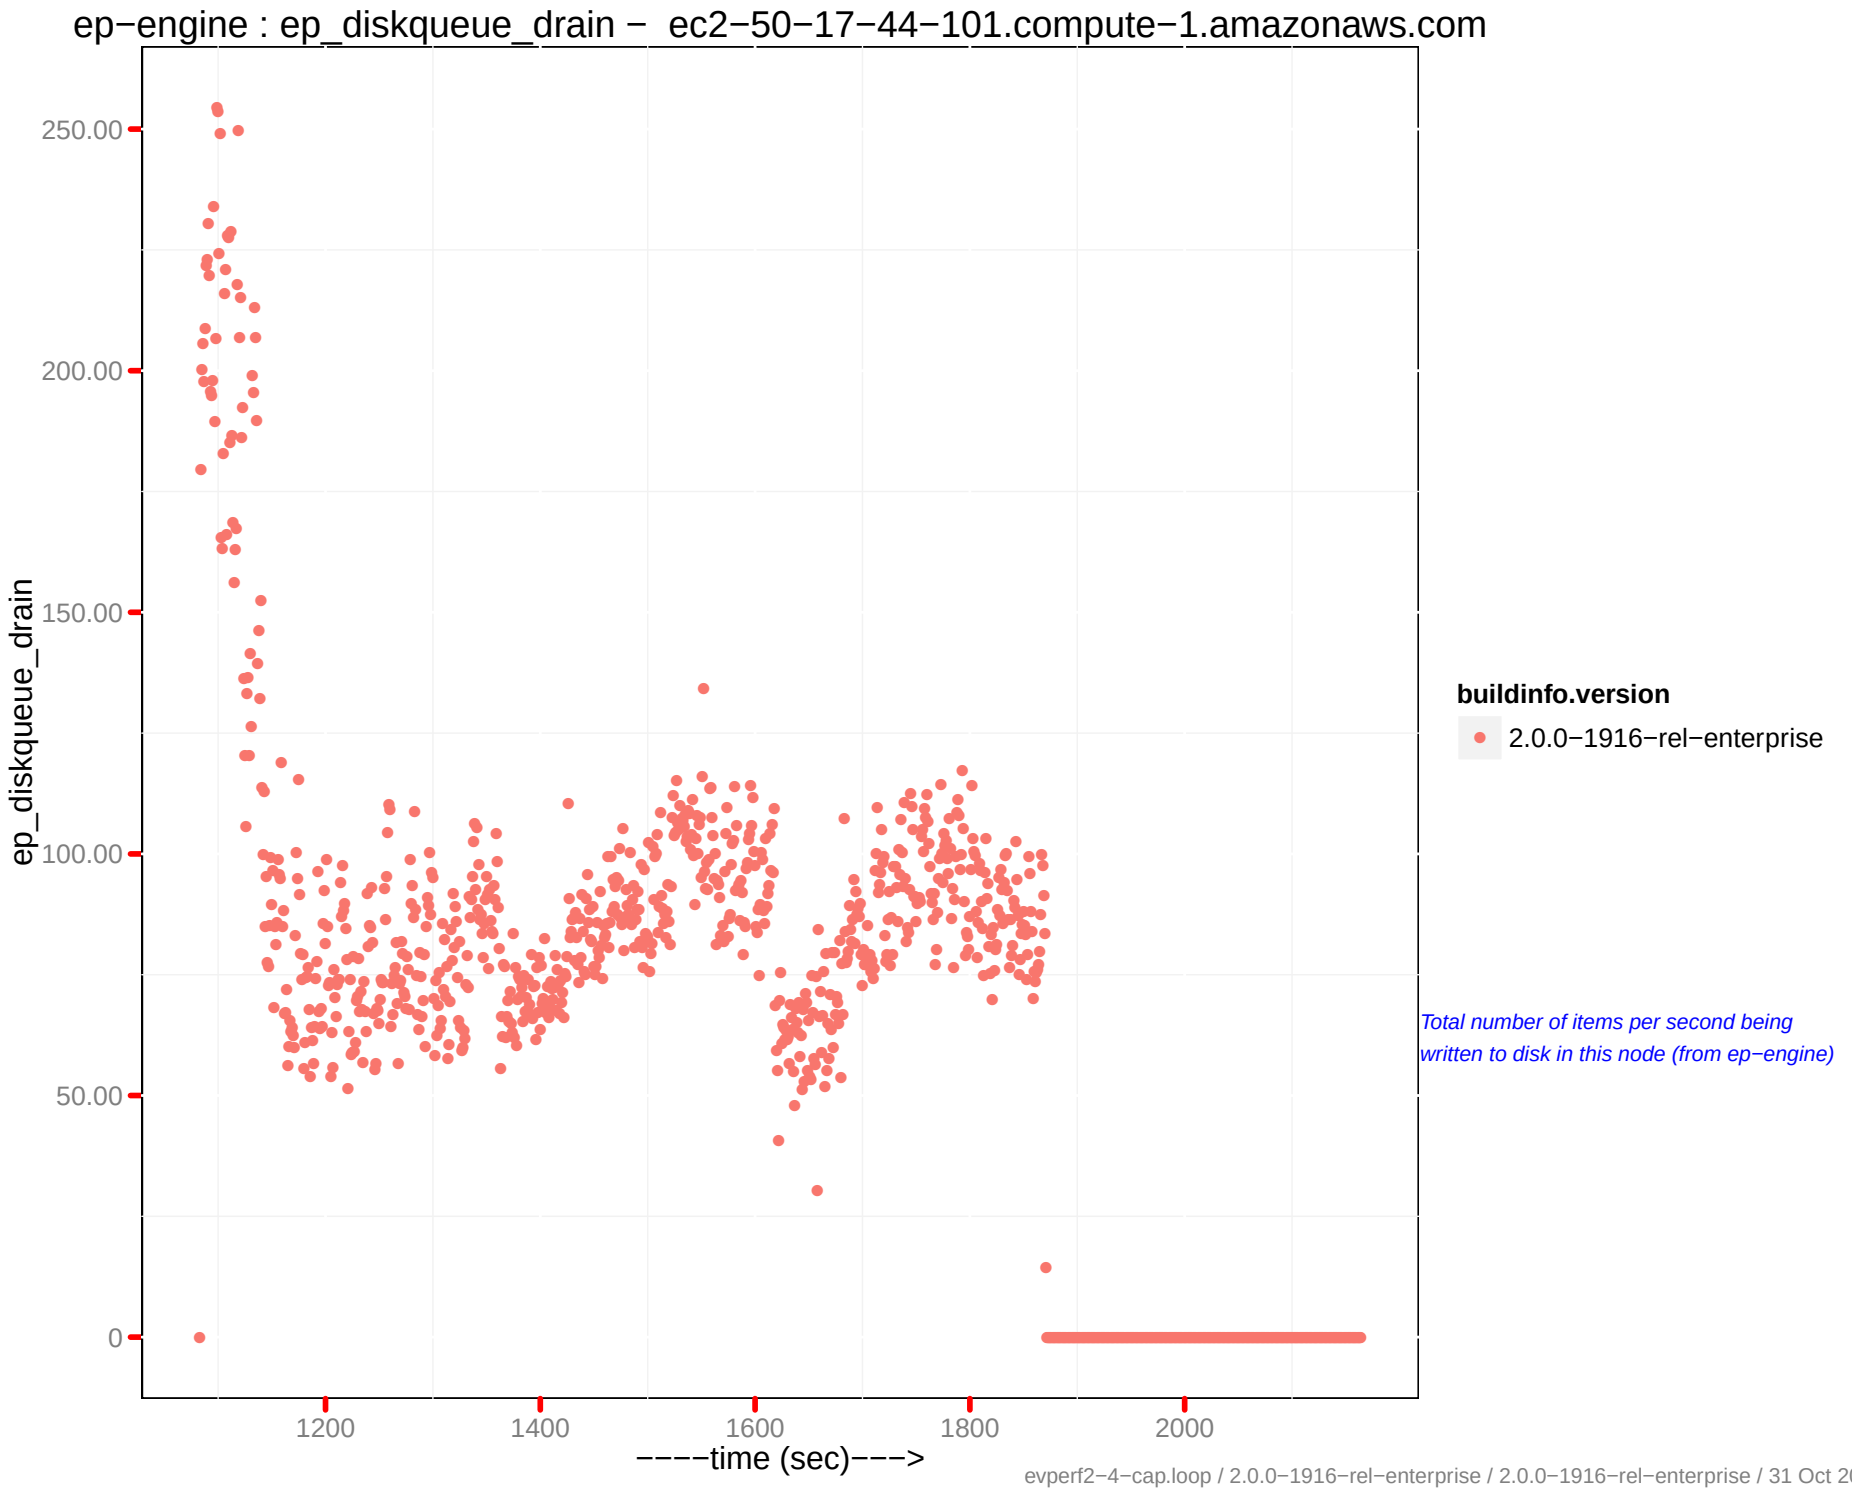
ep−engine : ep_diskqueue_drain − ec2−50−17−44−101.compute−1.amazonaws.com
ep_diskqueue_drain
250.00
200.00
150.00
100.00
50.00
0
1200
1400
1600
1800
2000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

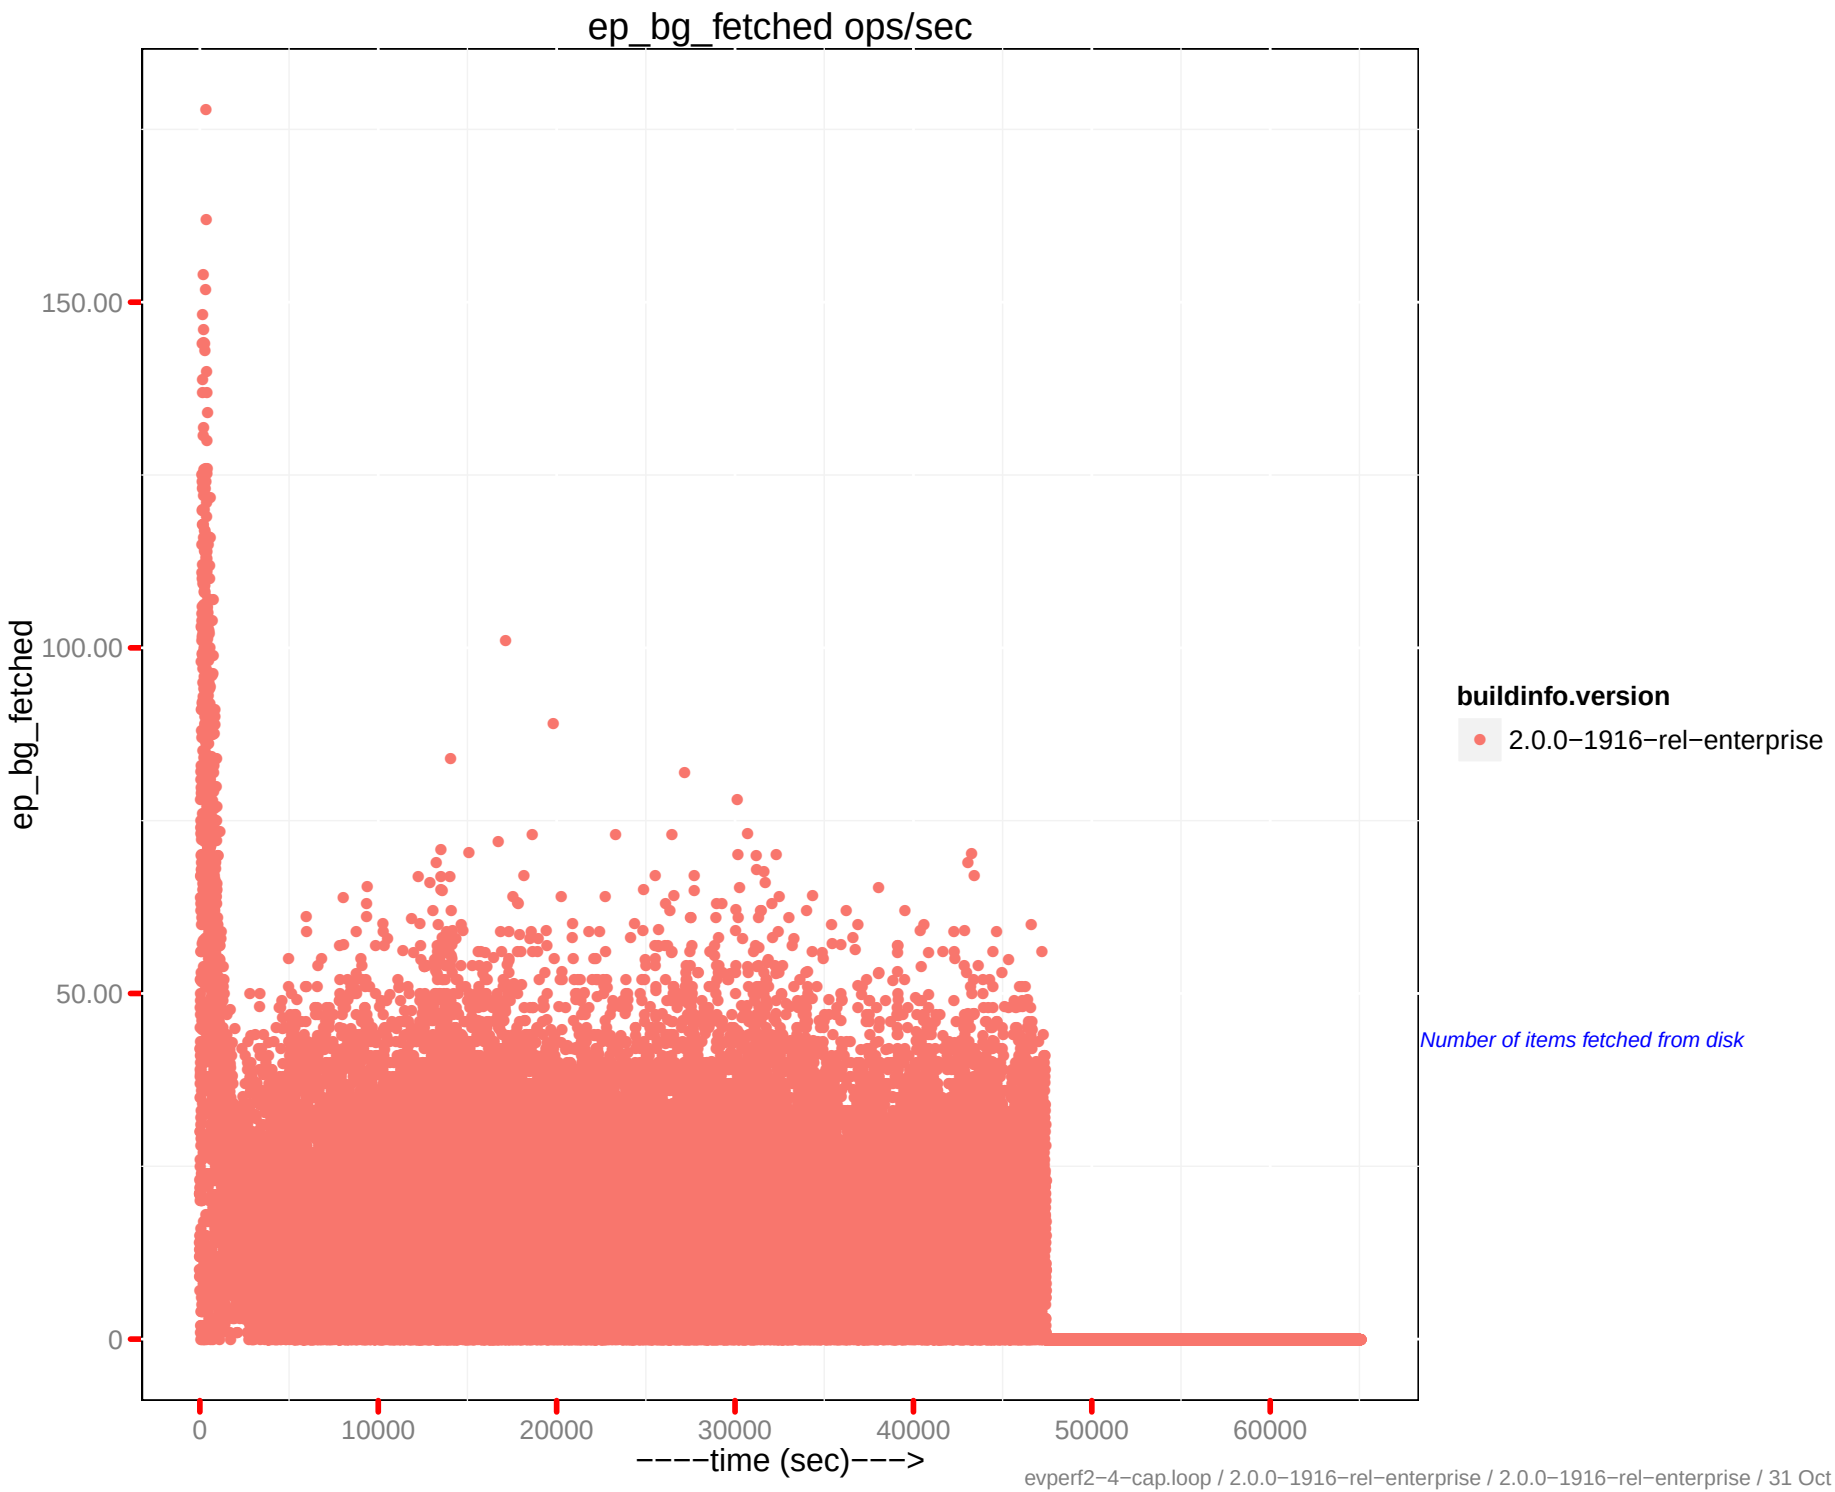
ep_bg_fetched ops/sec
ep_bg_fetched
150.00
100.00
50.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of items fetched from disk
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

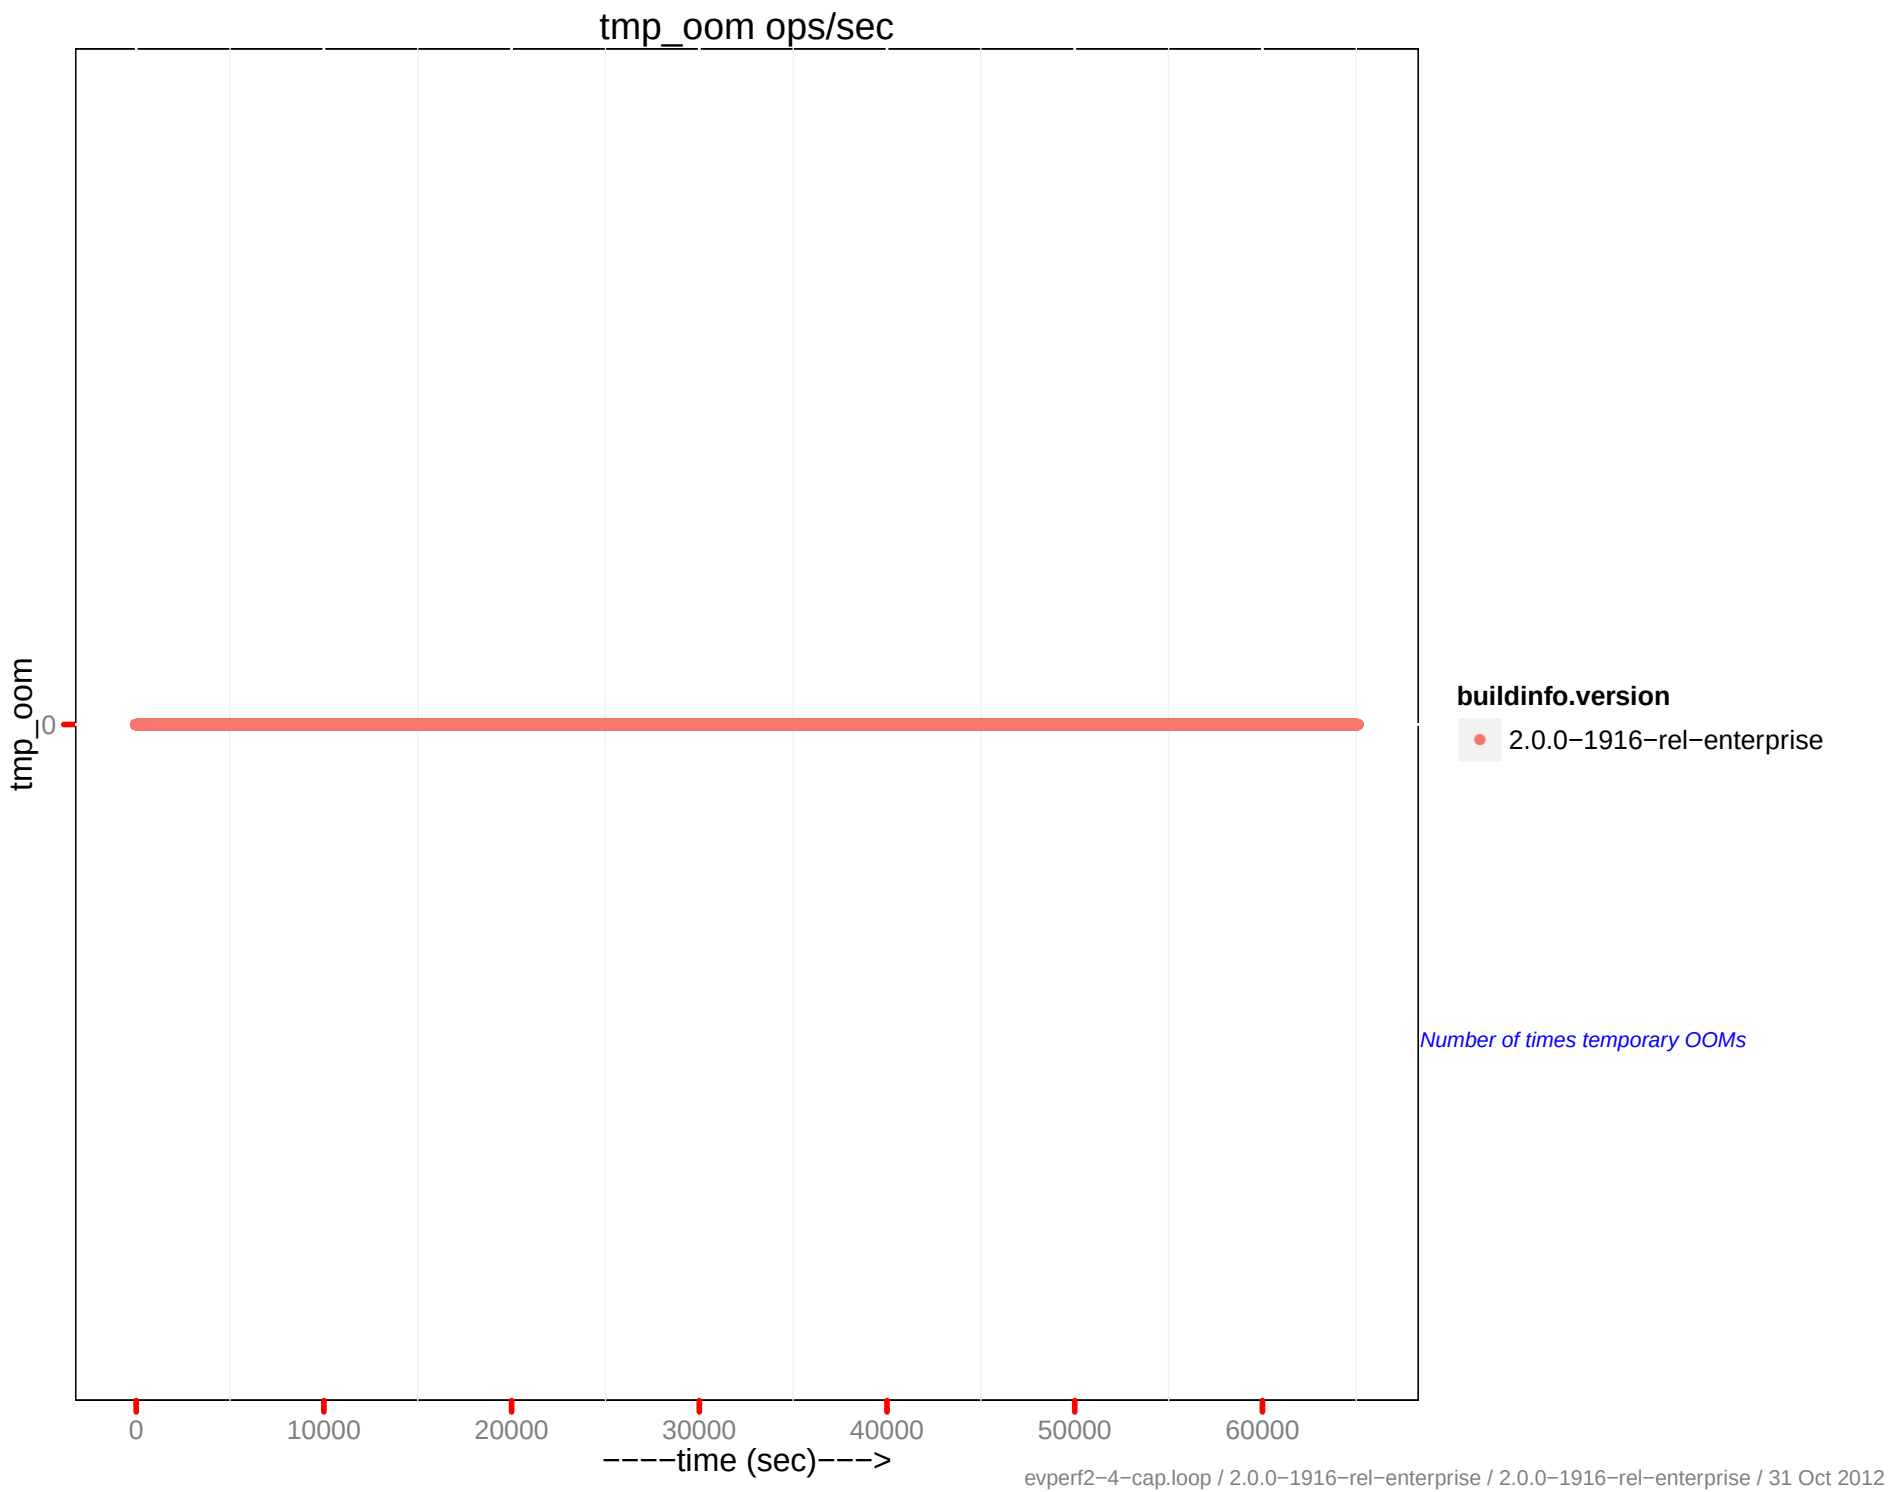
tmp_oom ops/sec
tmp_oom
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of times temporary OOMs
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

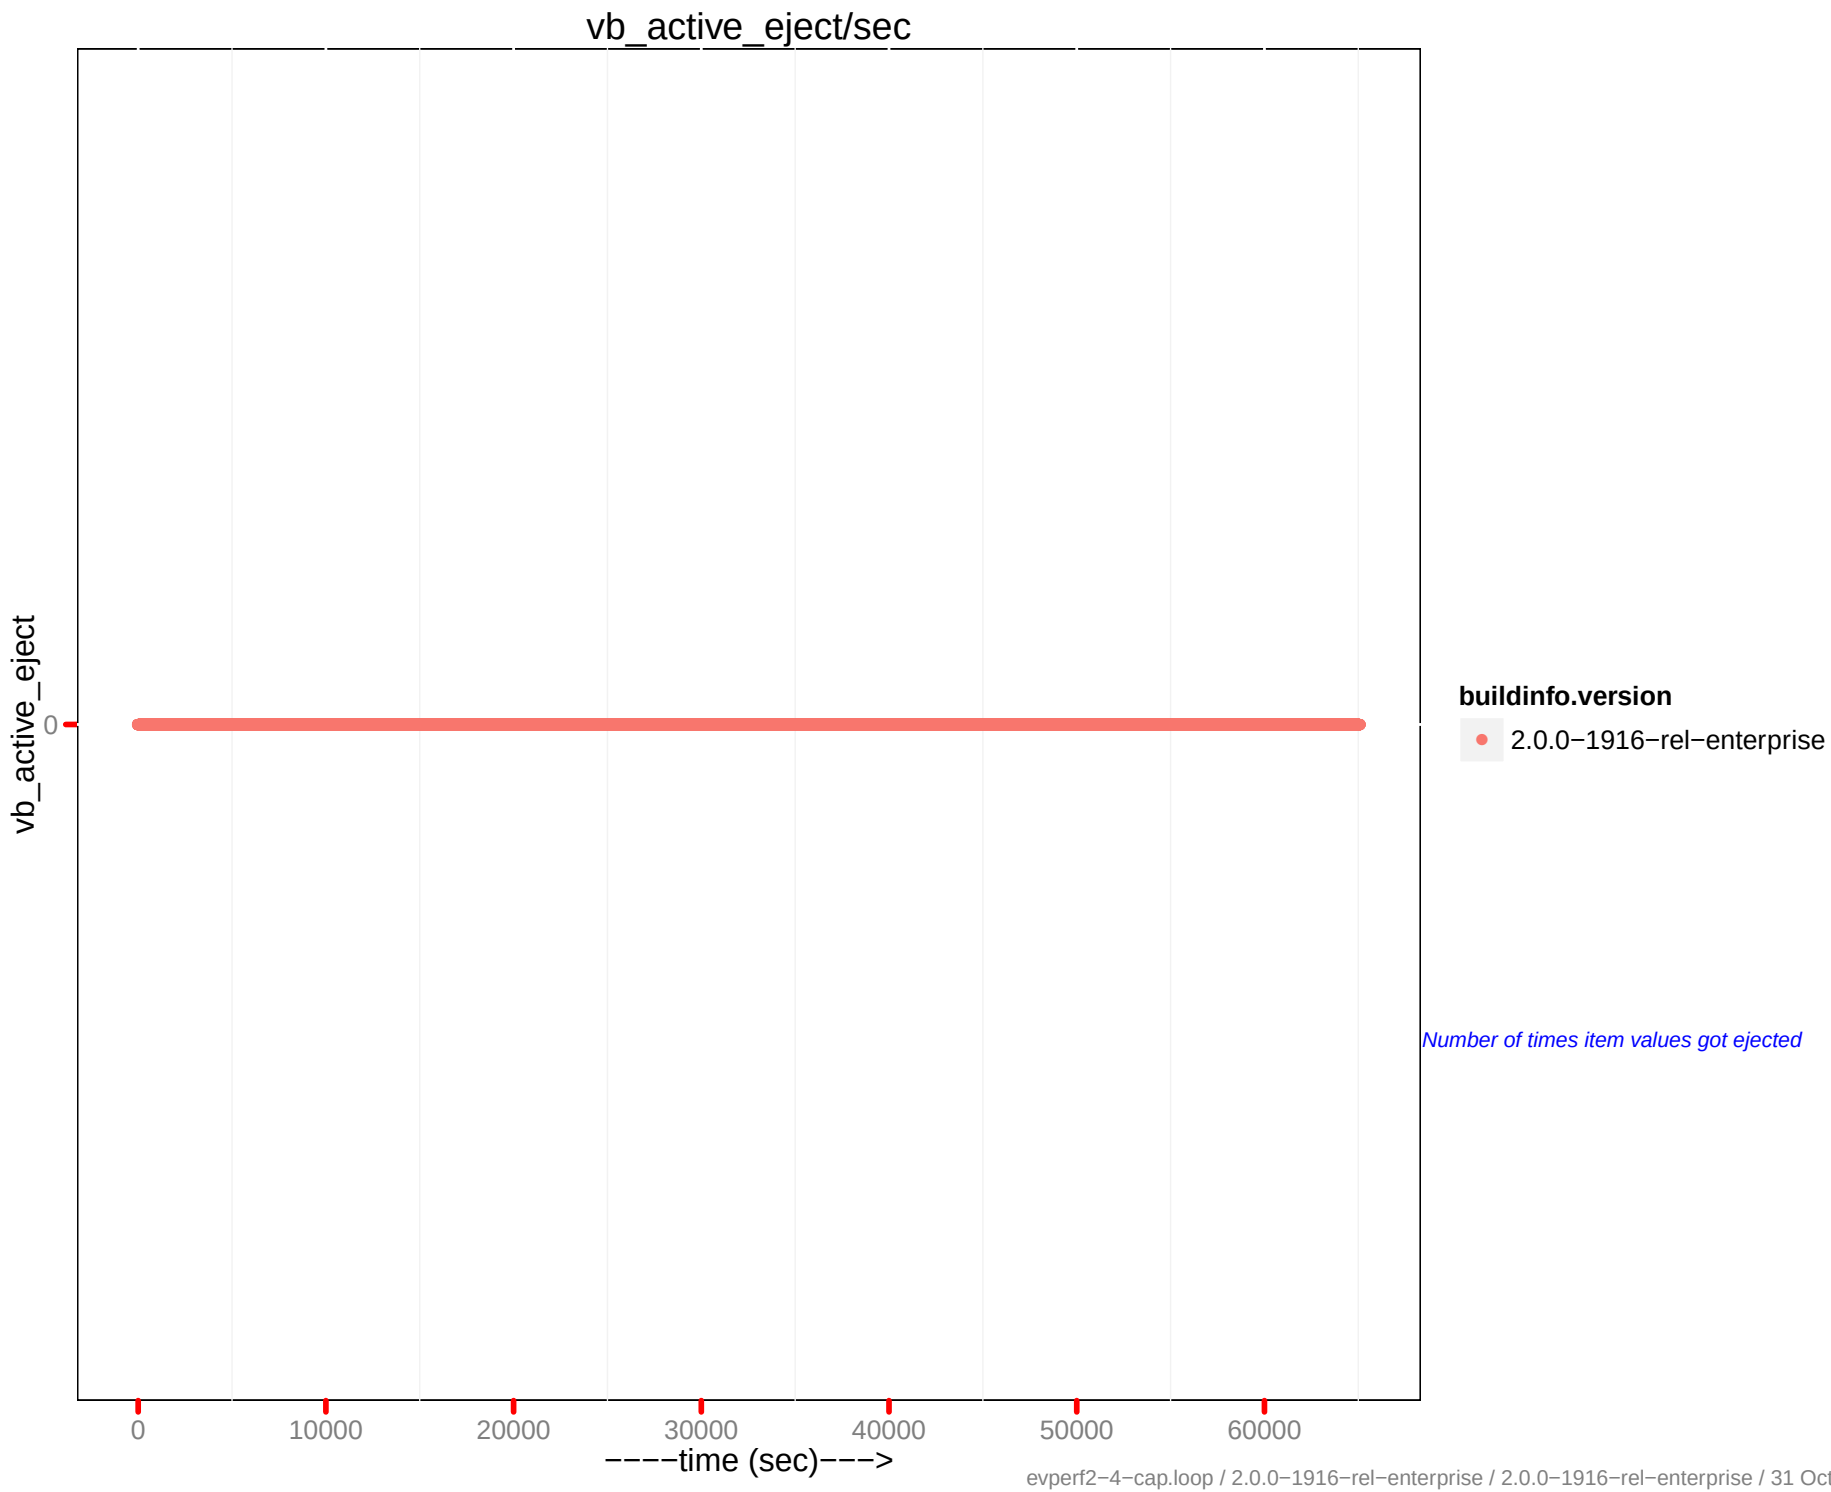
vb_active_eject/sec
vb_active_eject
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of times item values got ejected
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

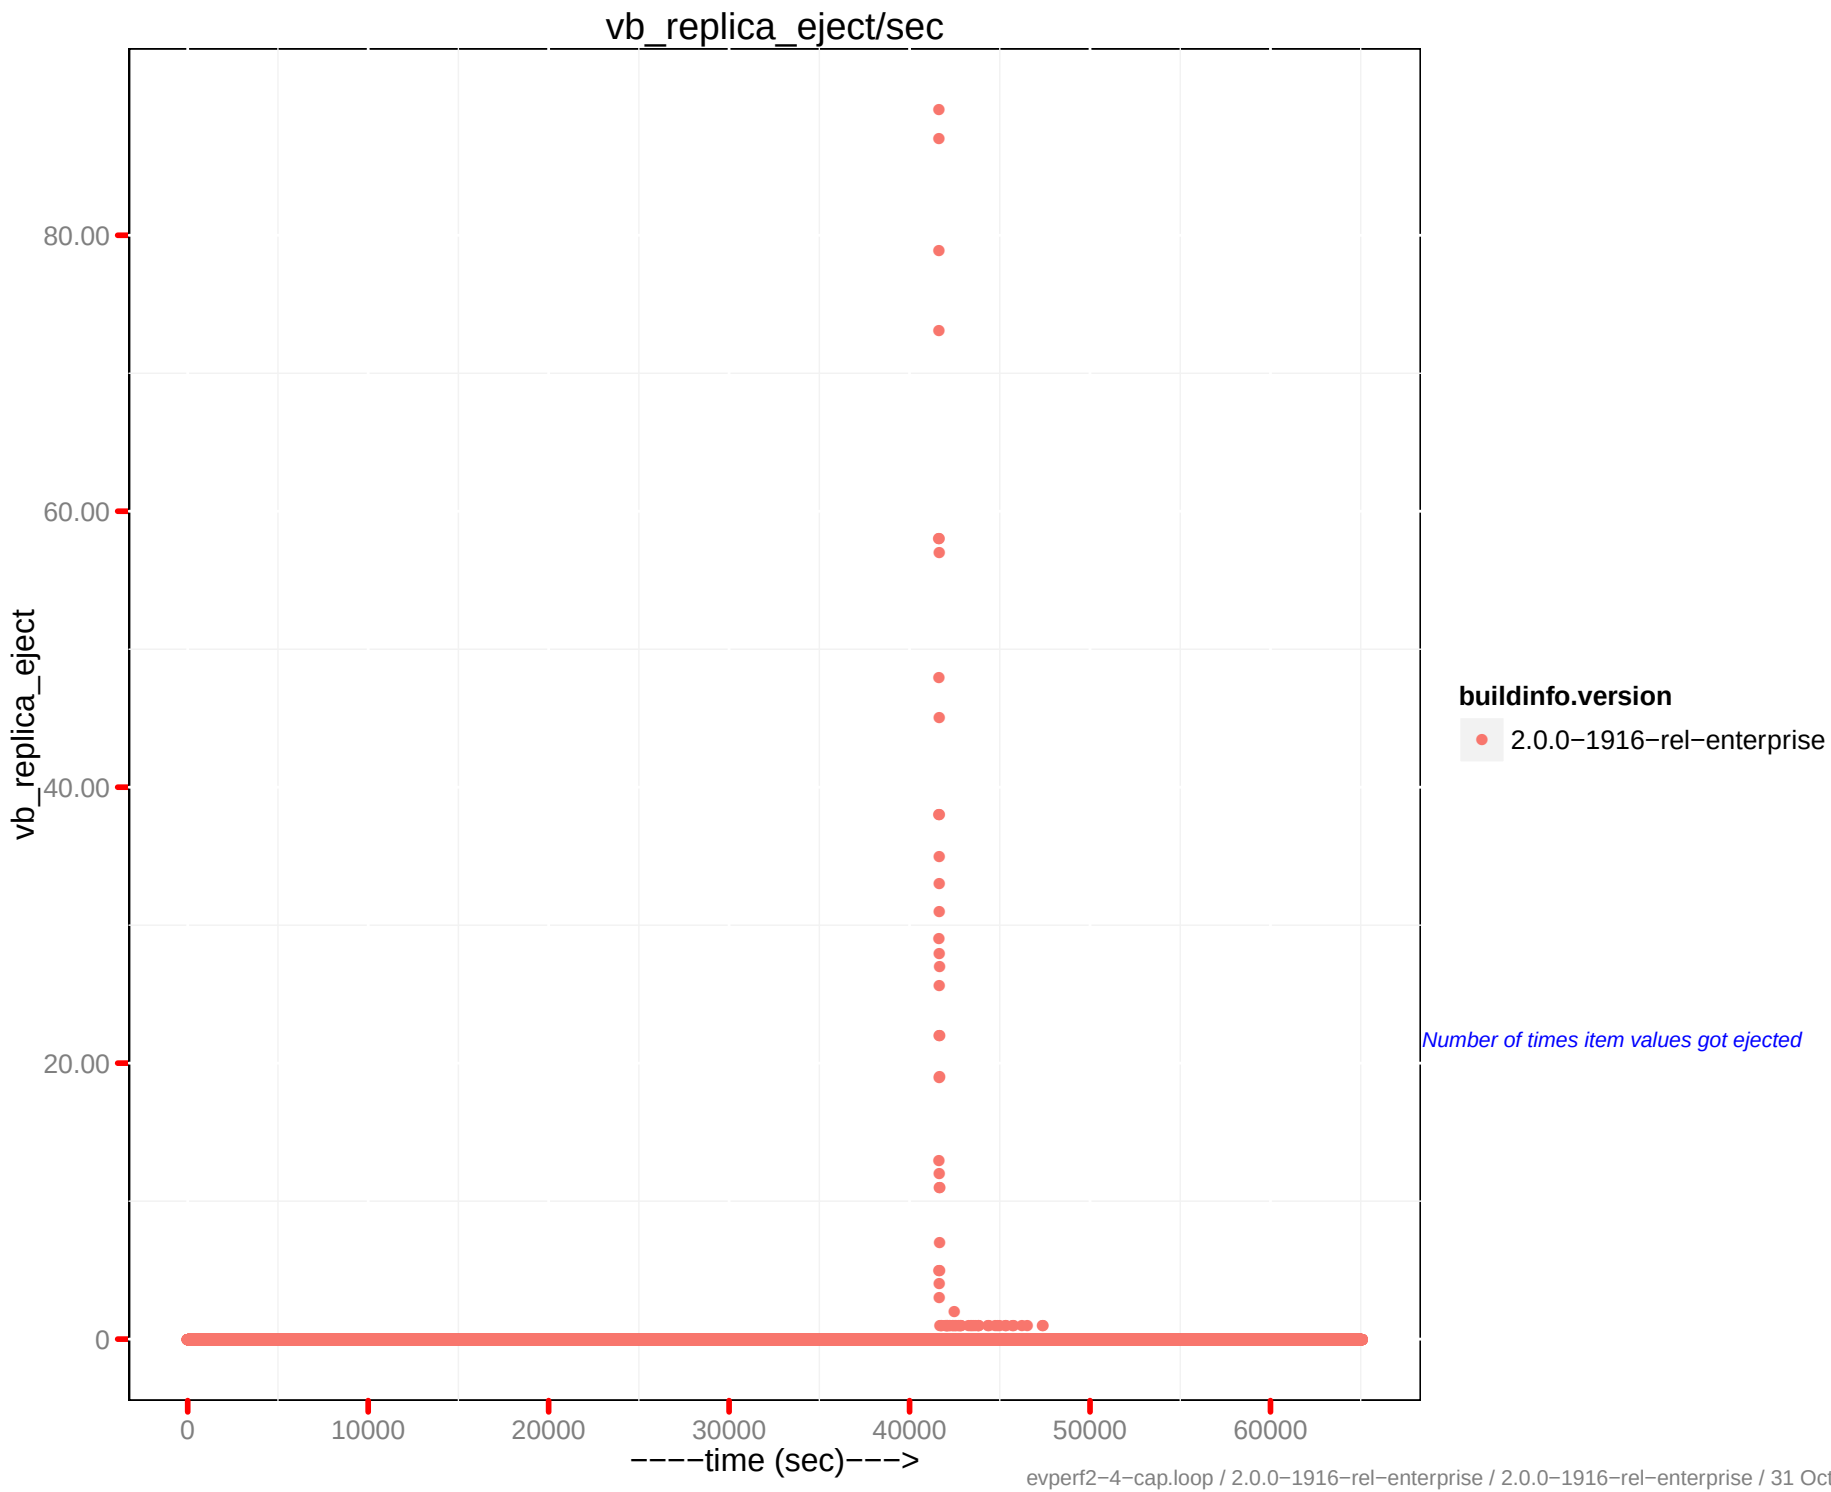
vb_replica_eject/sec
vb_replica_eject
80.00
60.00
40.00
20.00
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of times item values got ejected
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

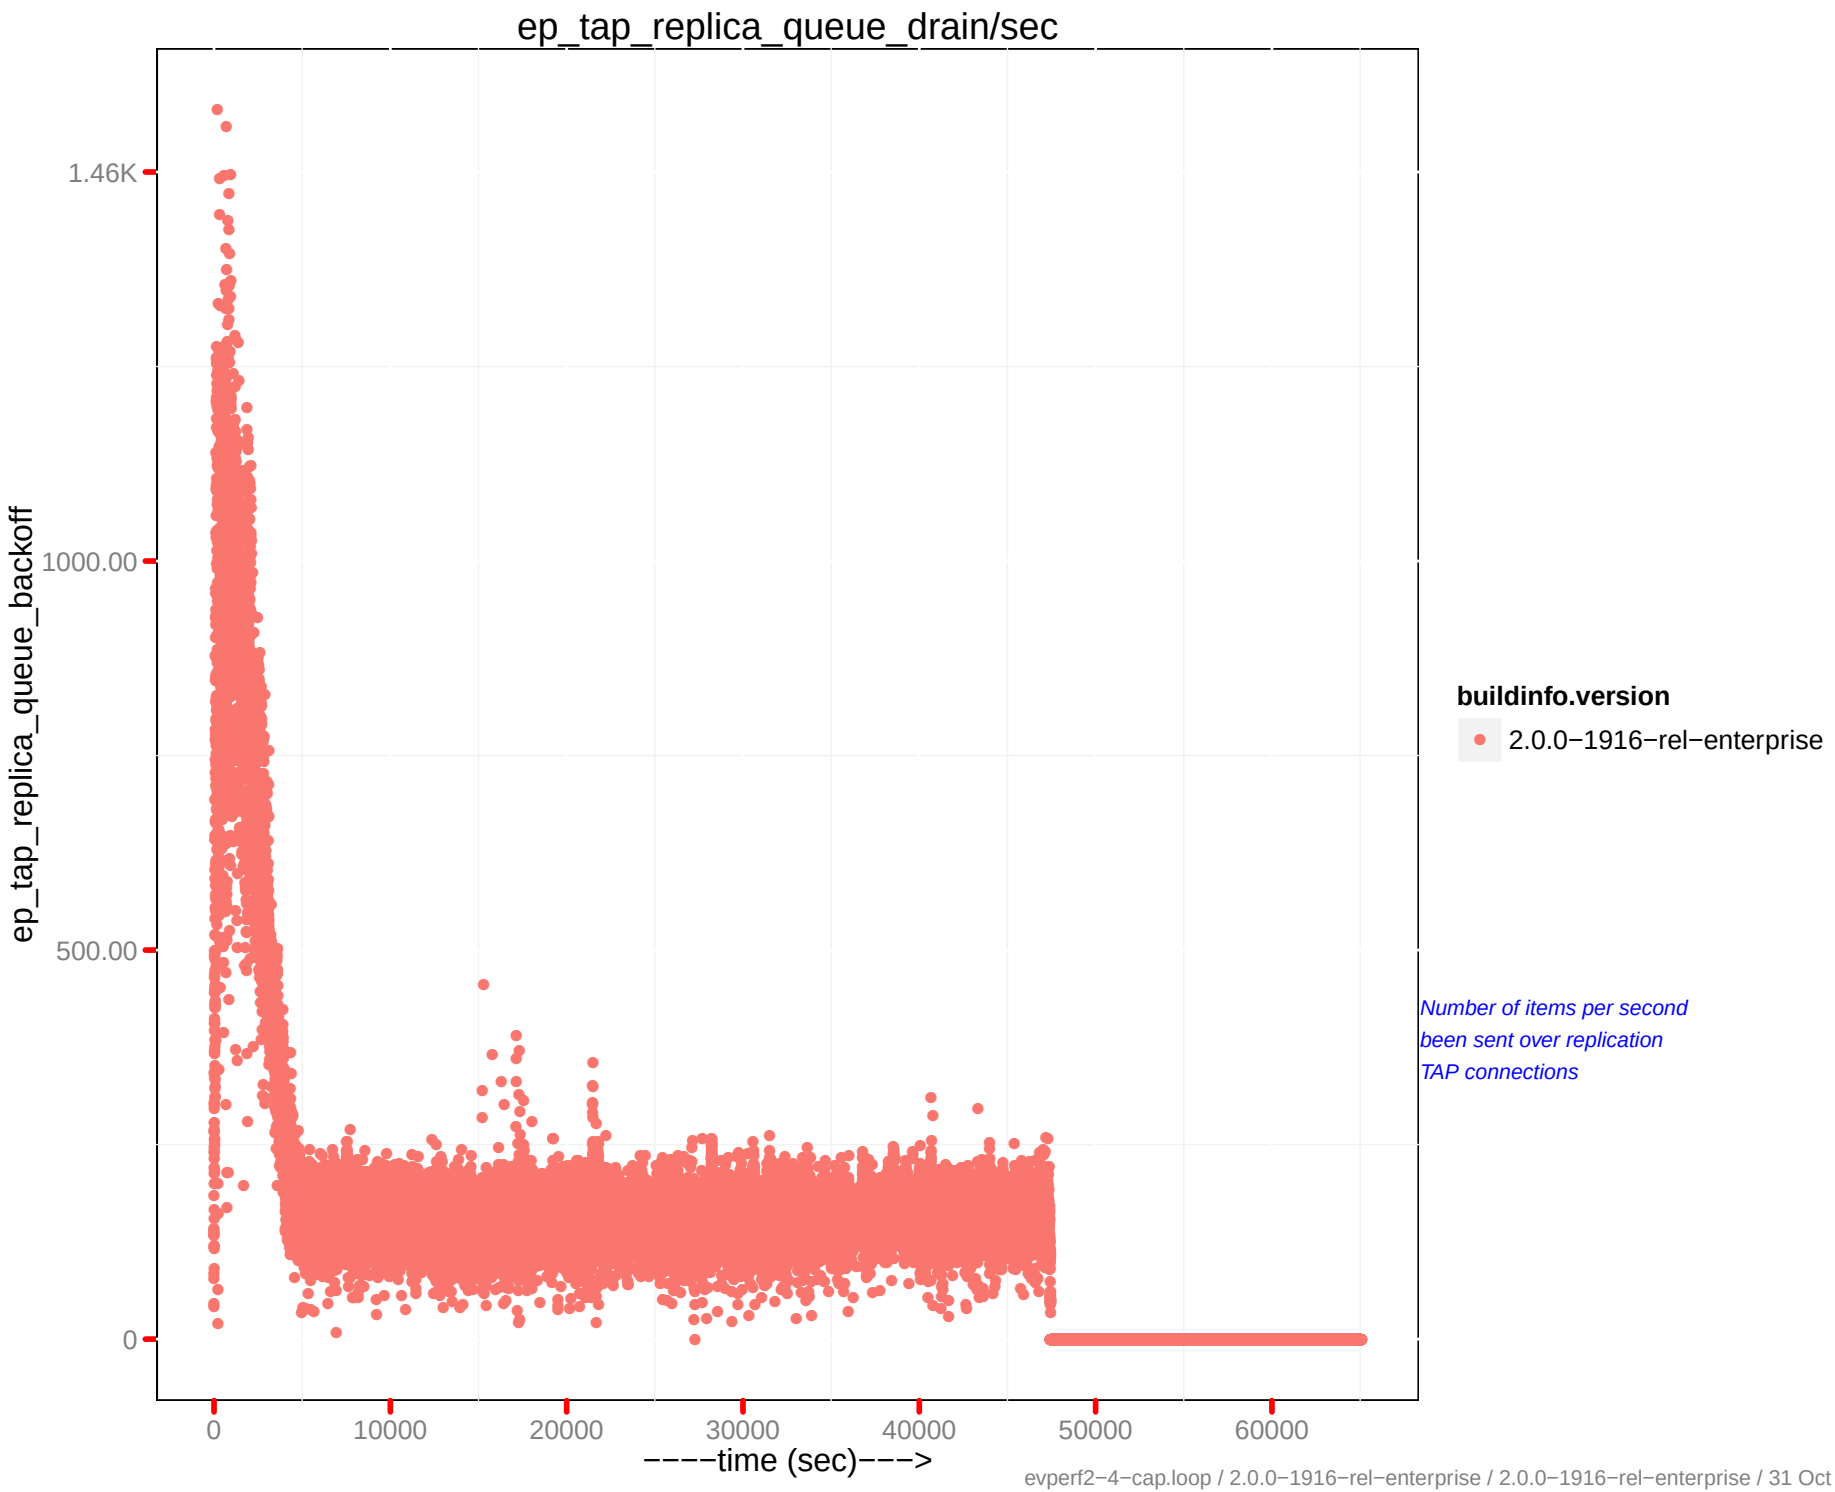
ep_tap_replica_queue_drain/sec
1.46K
1000.00
500.00
0
ep_tap_replica_queue_backoff
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of items per second
been sent over replication
TAP connections
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

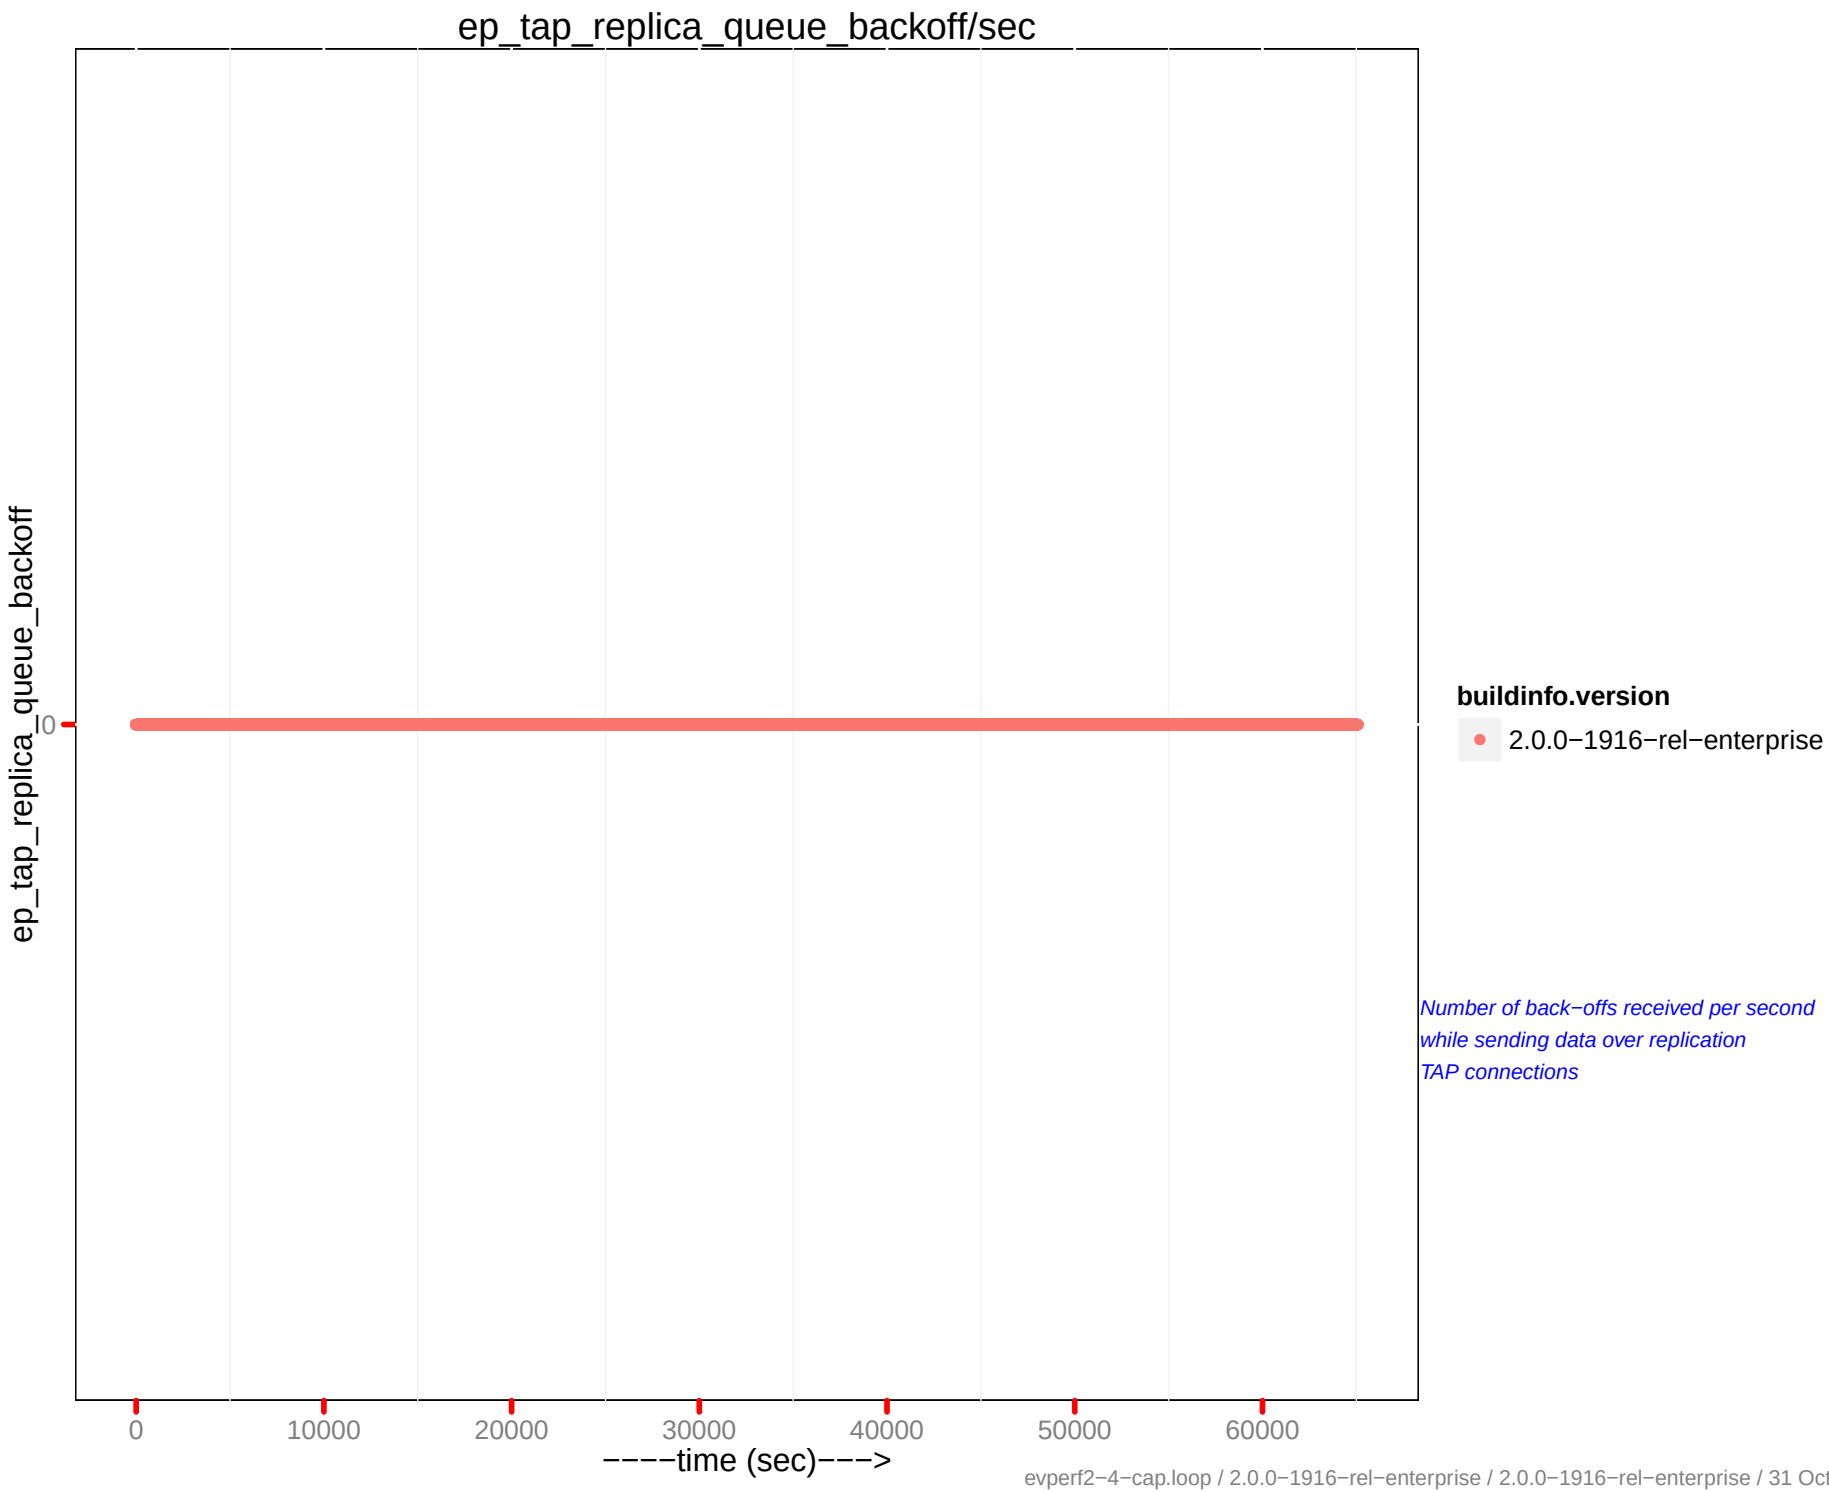
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of back-offs received per second
while sending data over replication
TAP connections
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

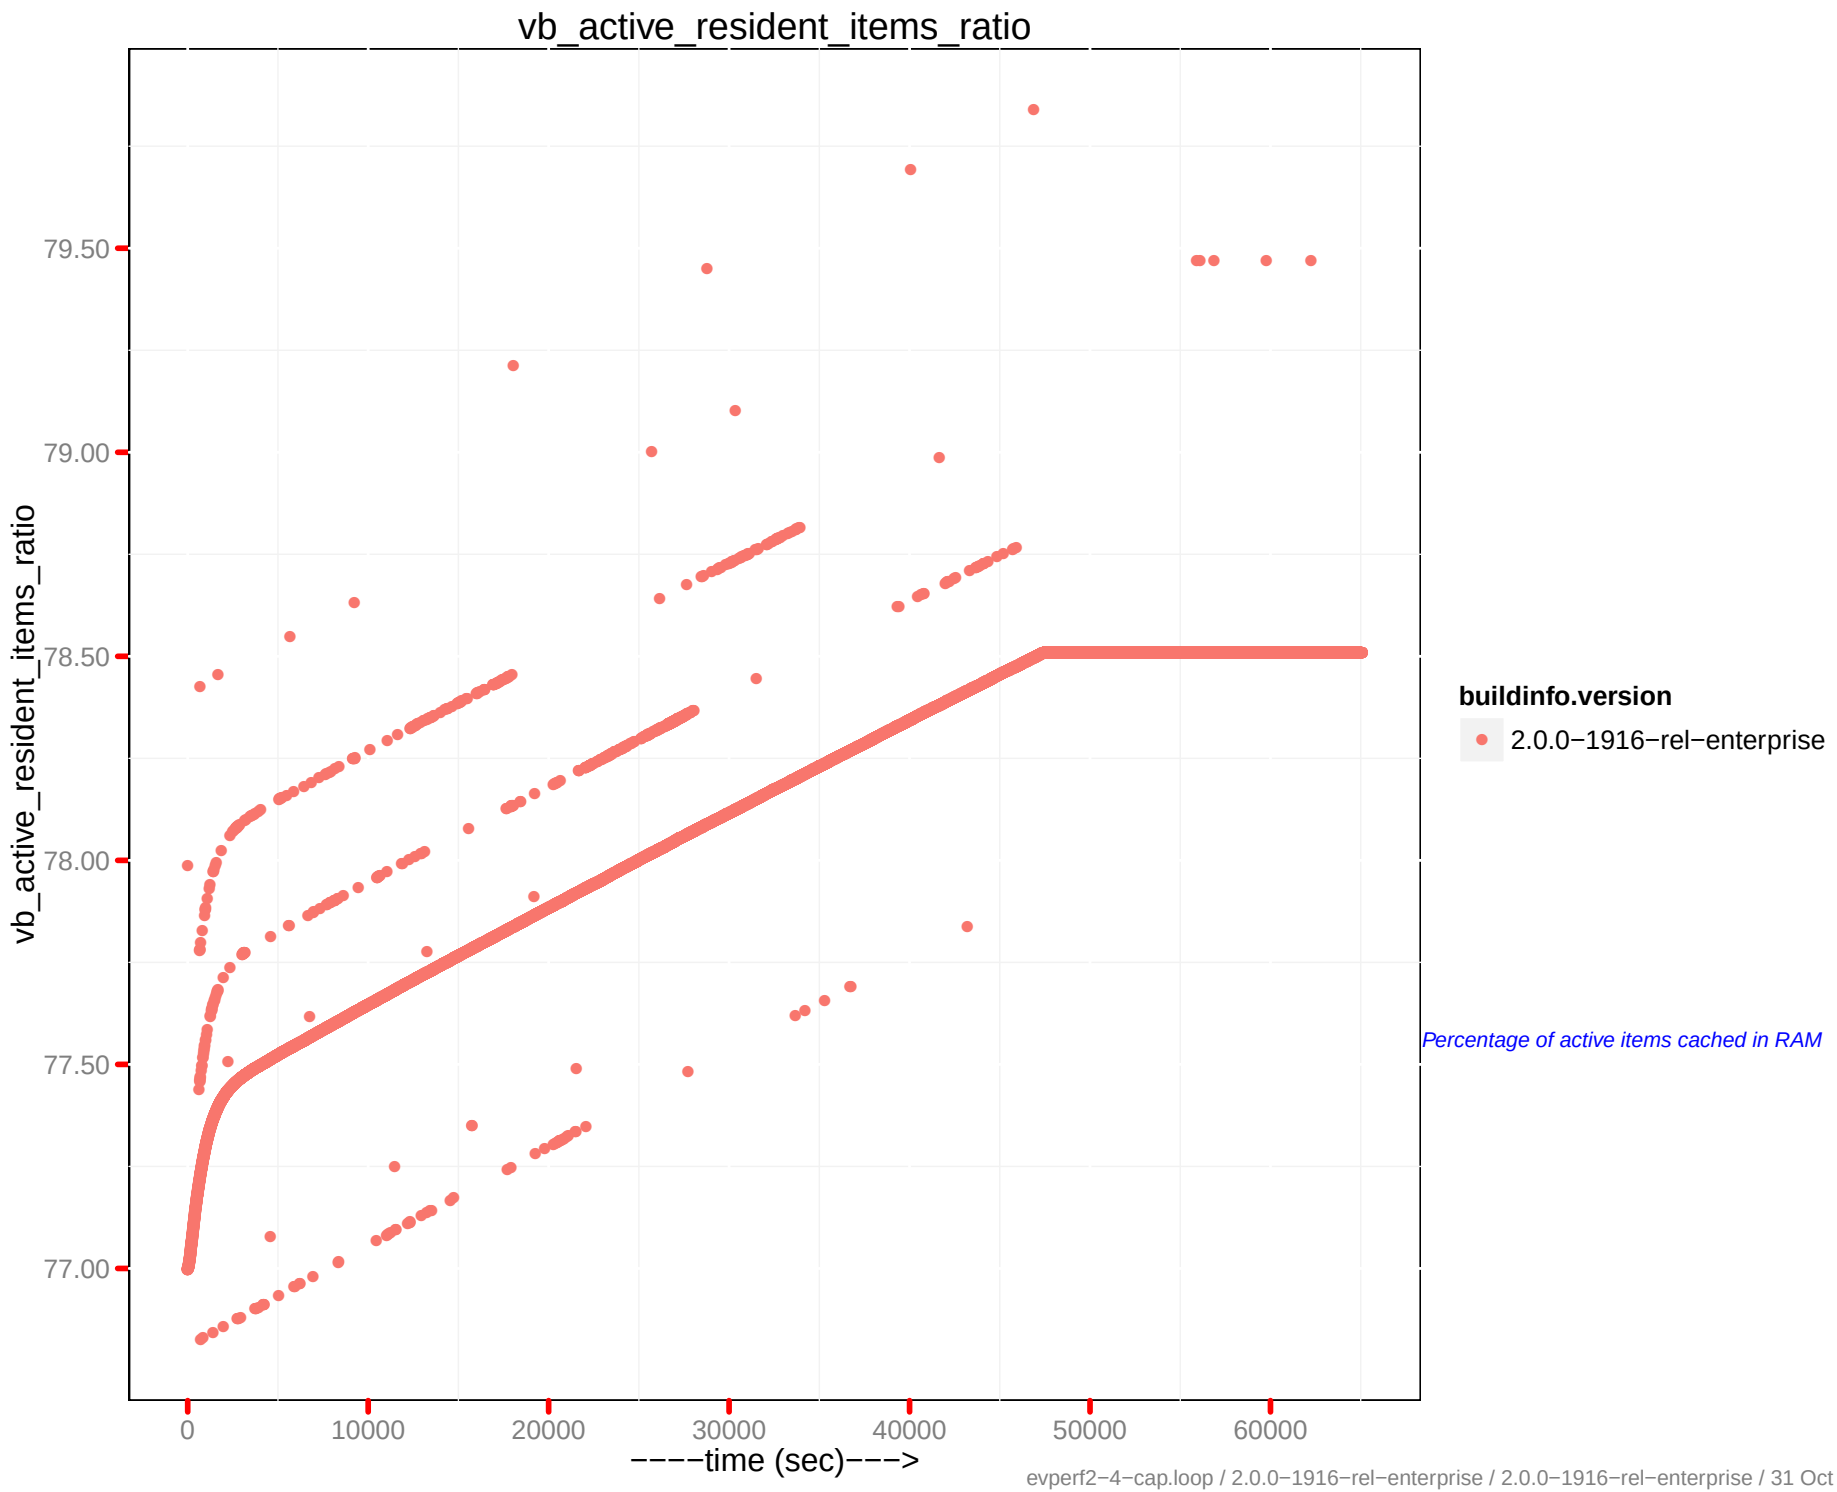
vb_active_resident_items_ratio
79.50
79.00
78.50
78.00
77.50
77.00
vb_active_resident_items_ratio
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Percentage of active items cached in RAM
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

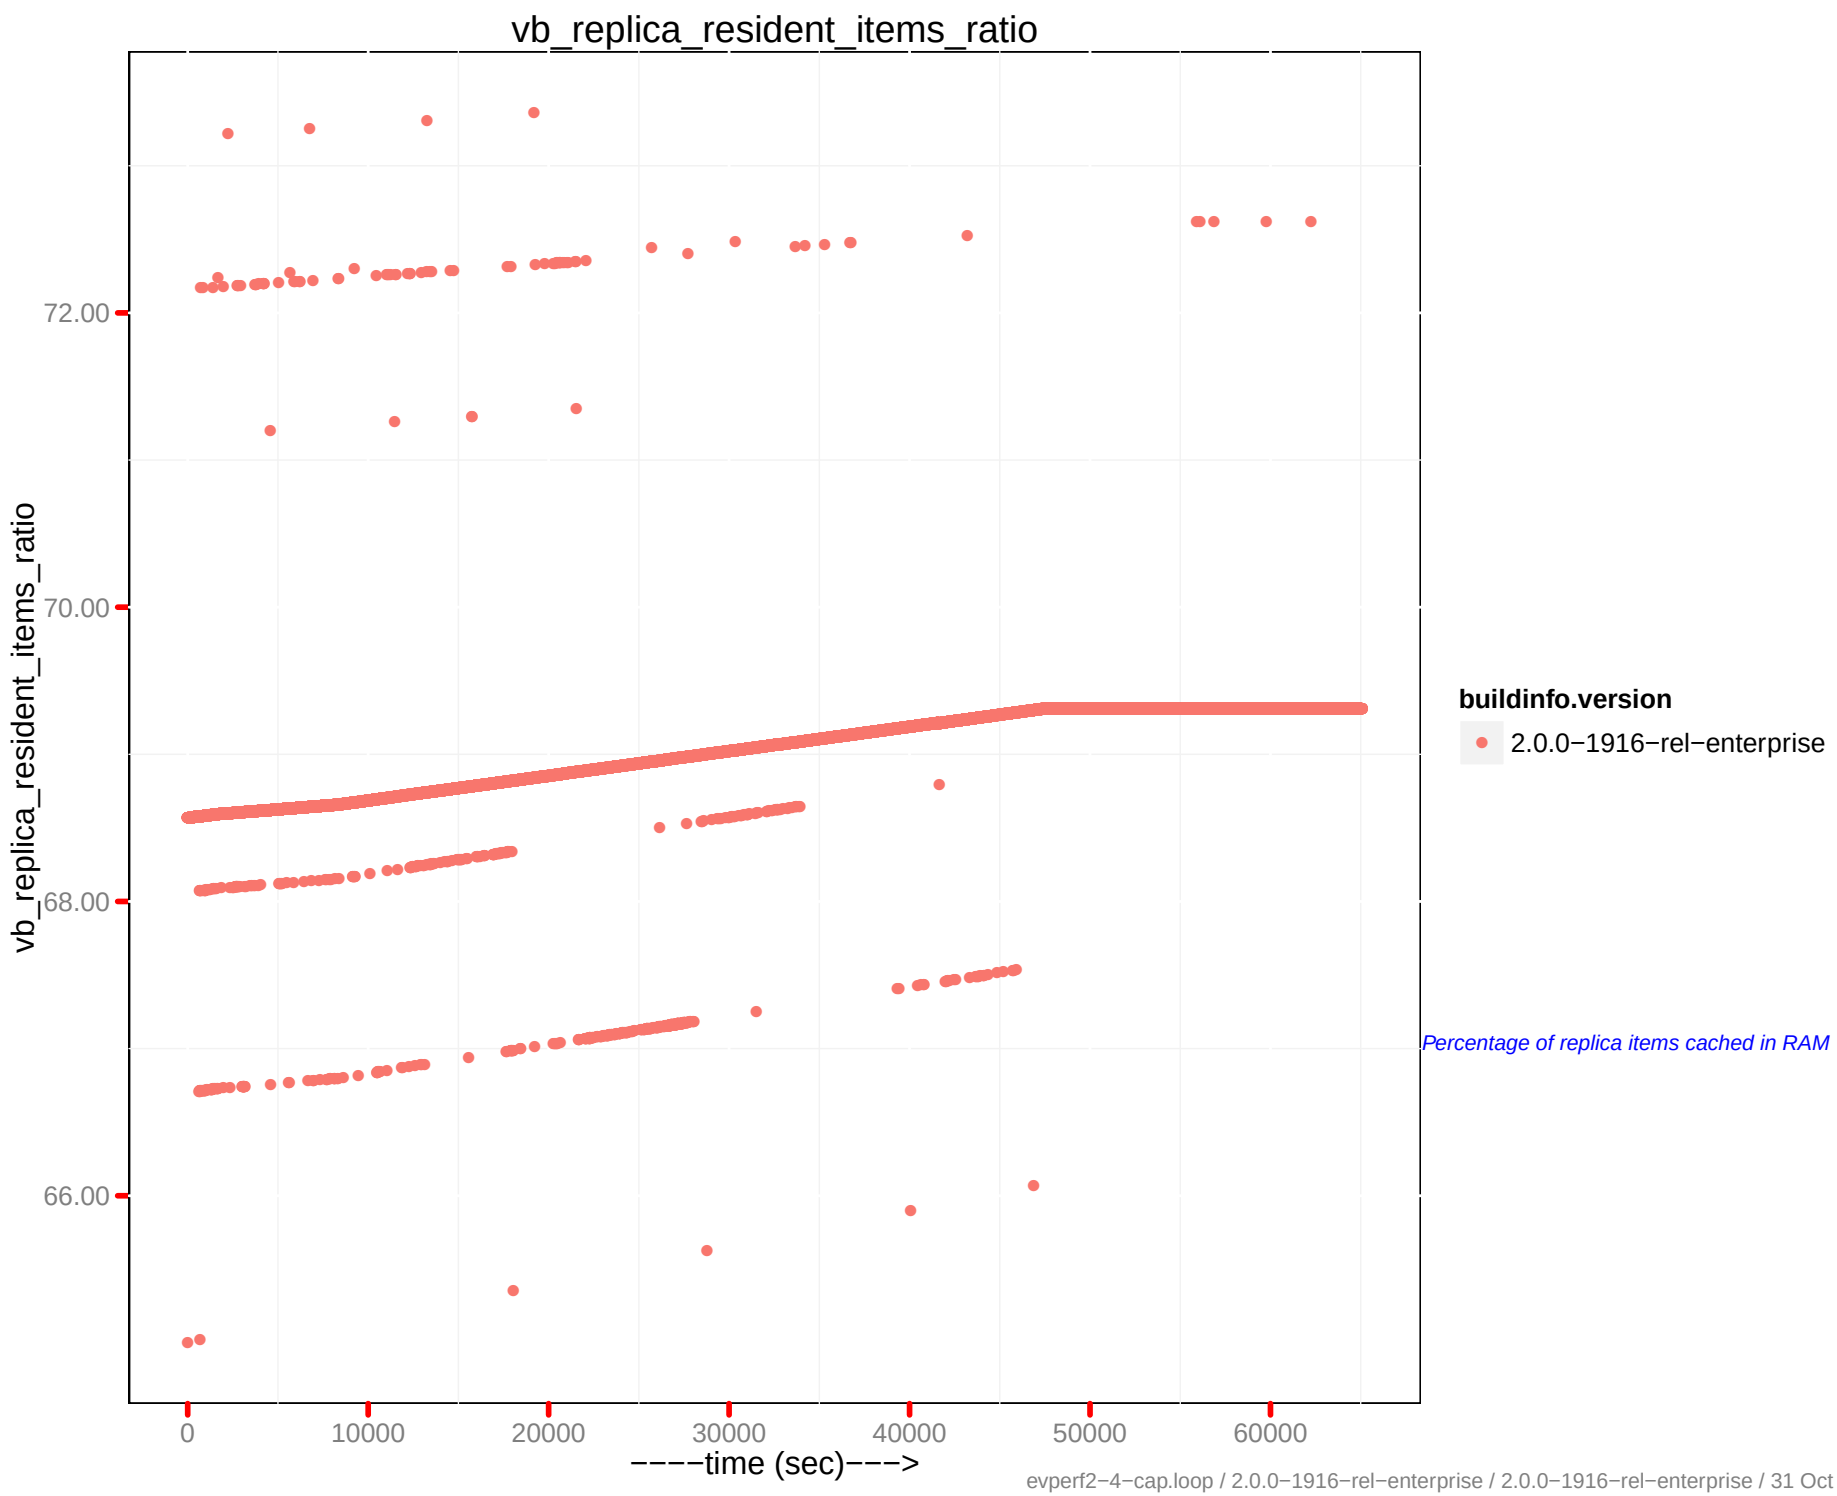
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
72.00
70.00
68.00
66.00
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Percentage of replica items cached in RAM
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

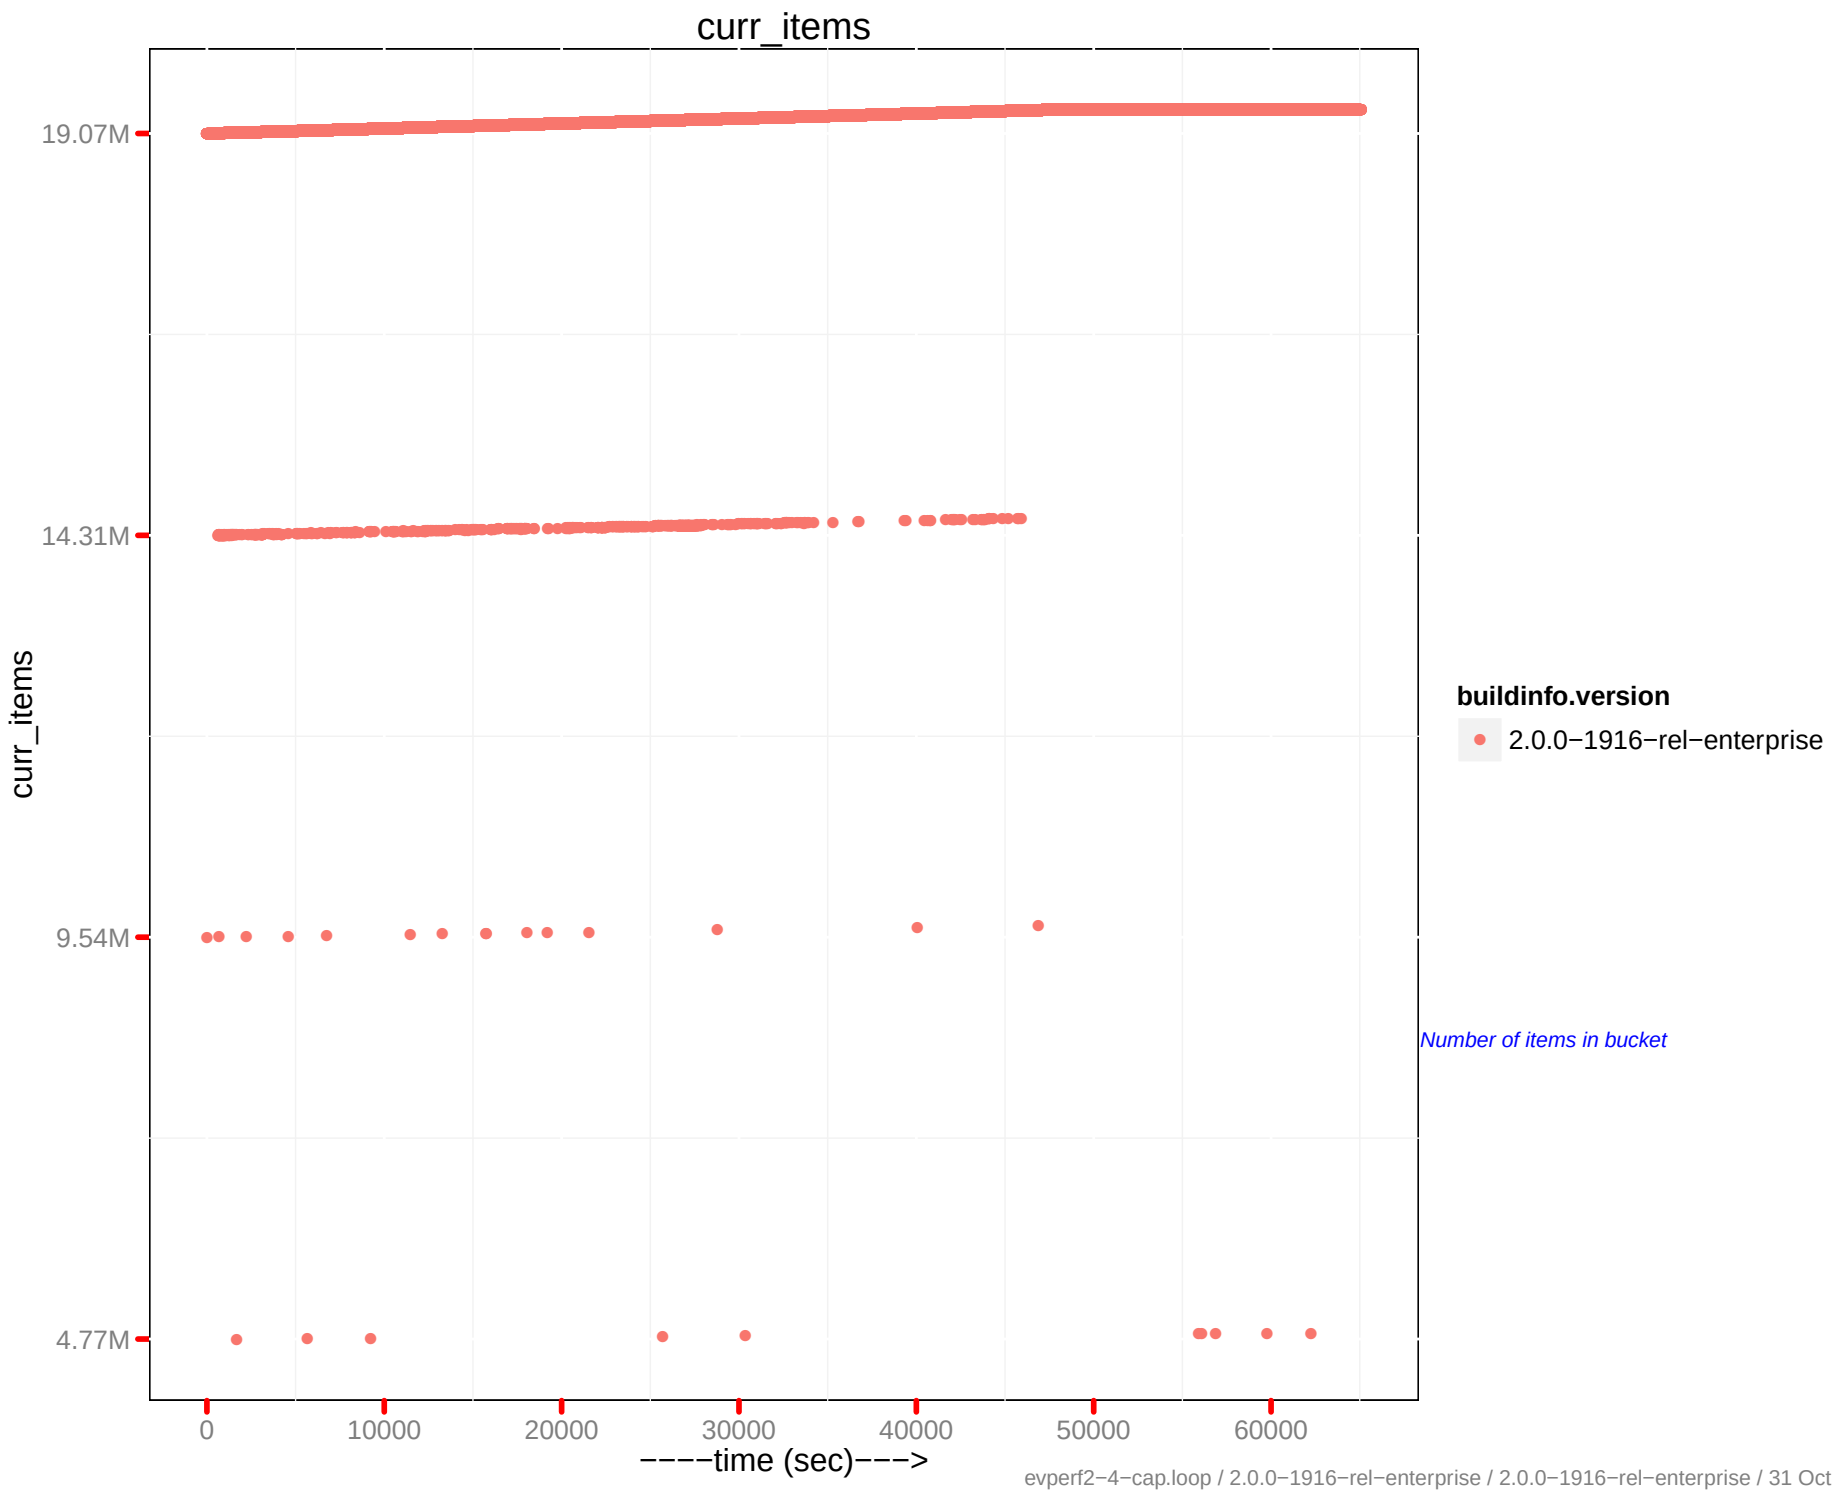
curr_items
19.07M
14.31M
9.54M
4.77M
curr_items
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of items in bucket
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

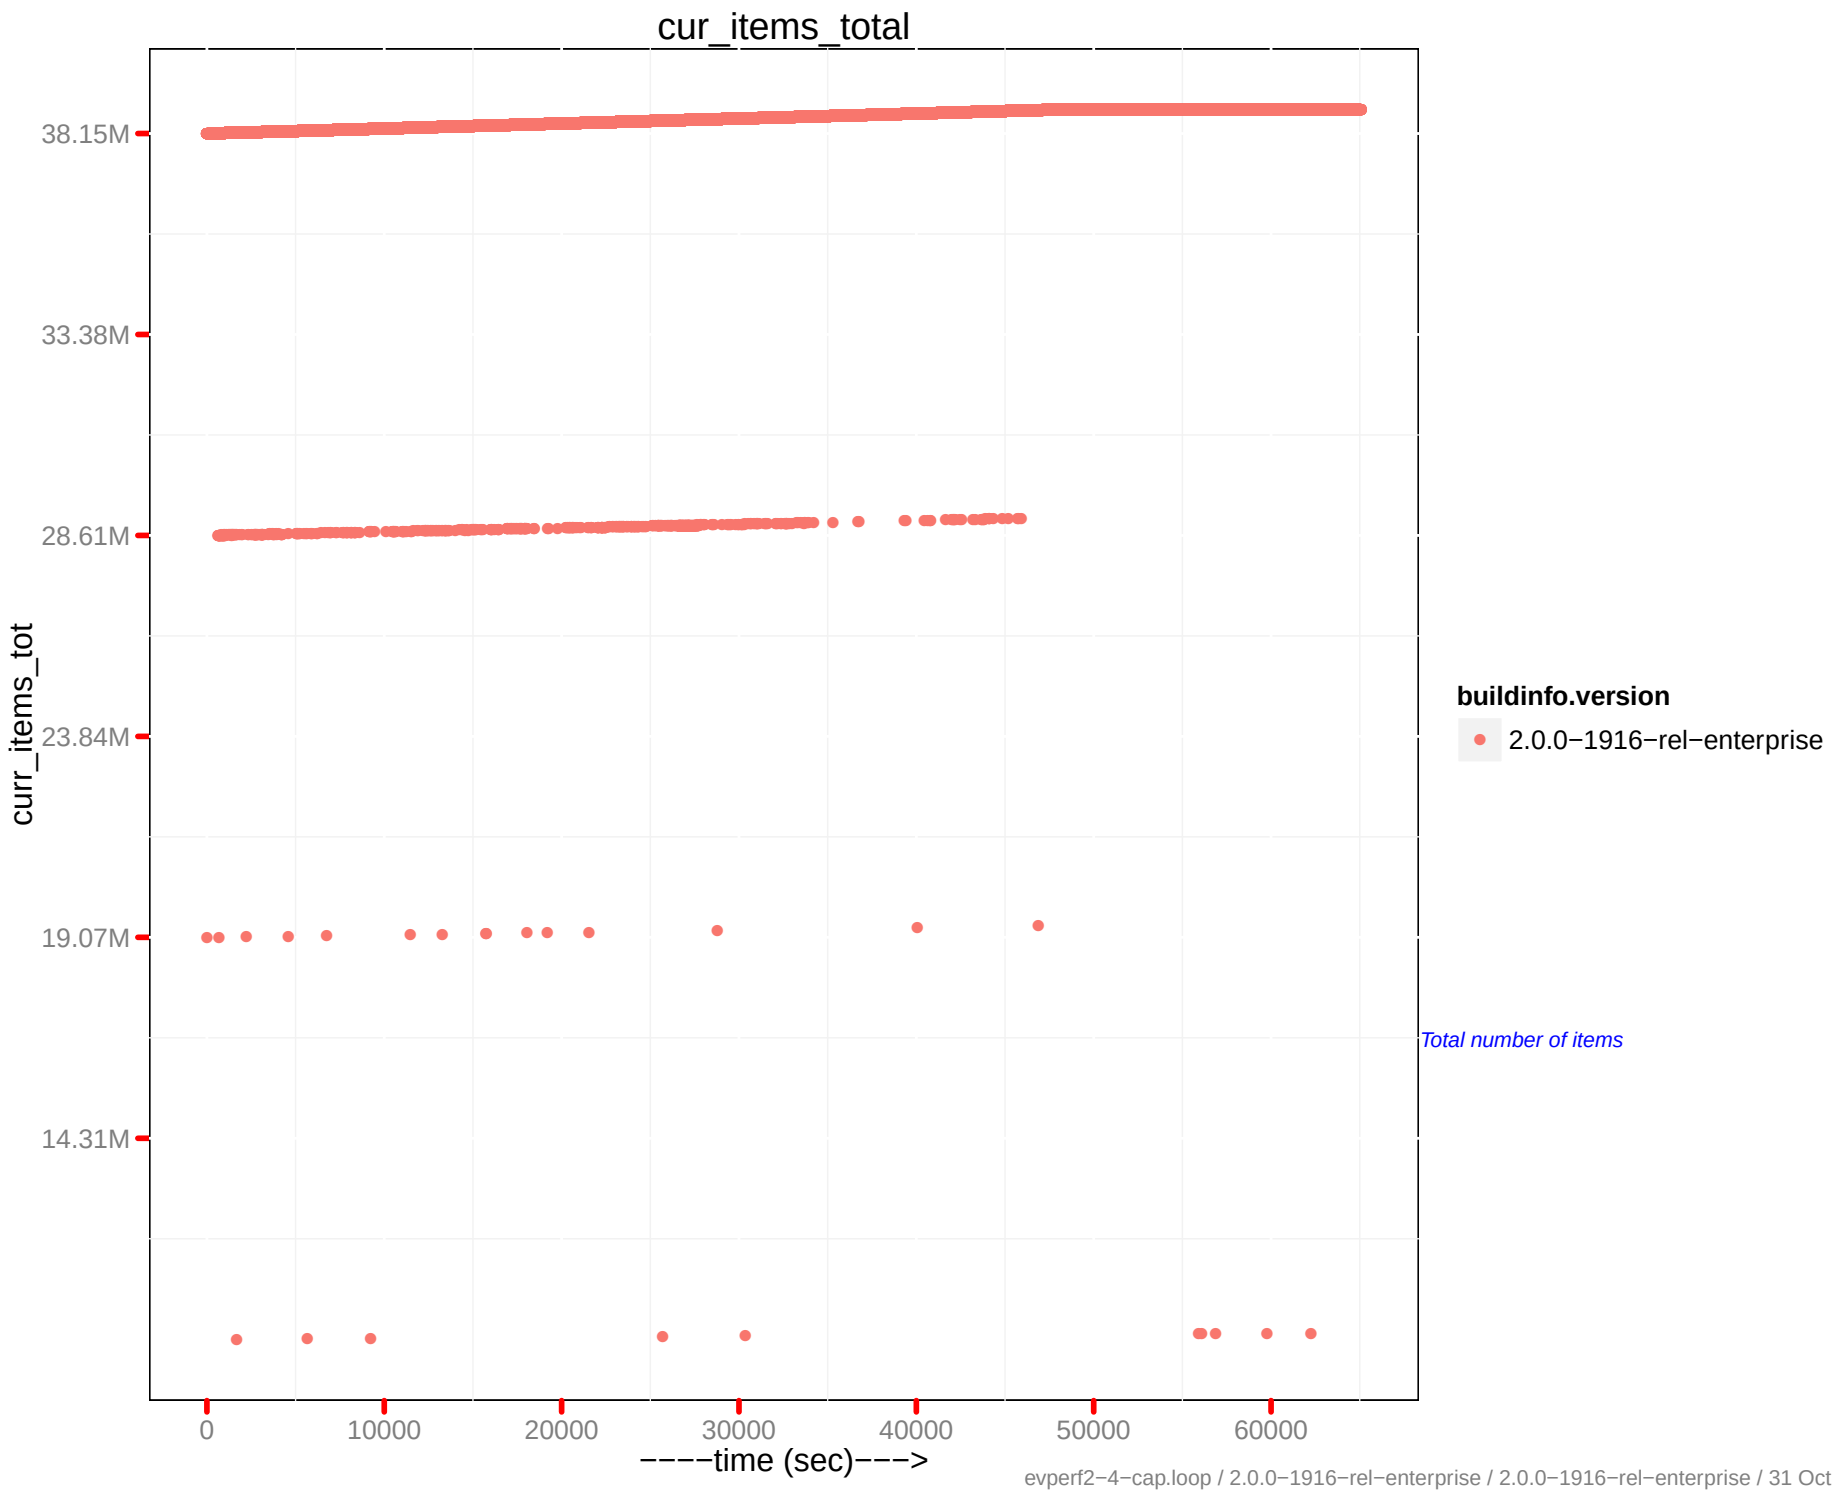
cur_items_total
38.15M
33.38M
28.61M
23.84M
19.07M
14.31M
curr_items_tot
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Total number of items
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

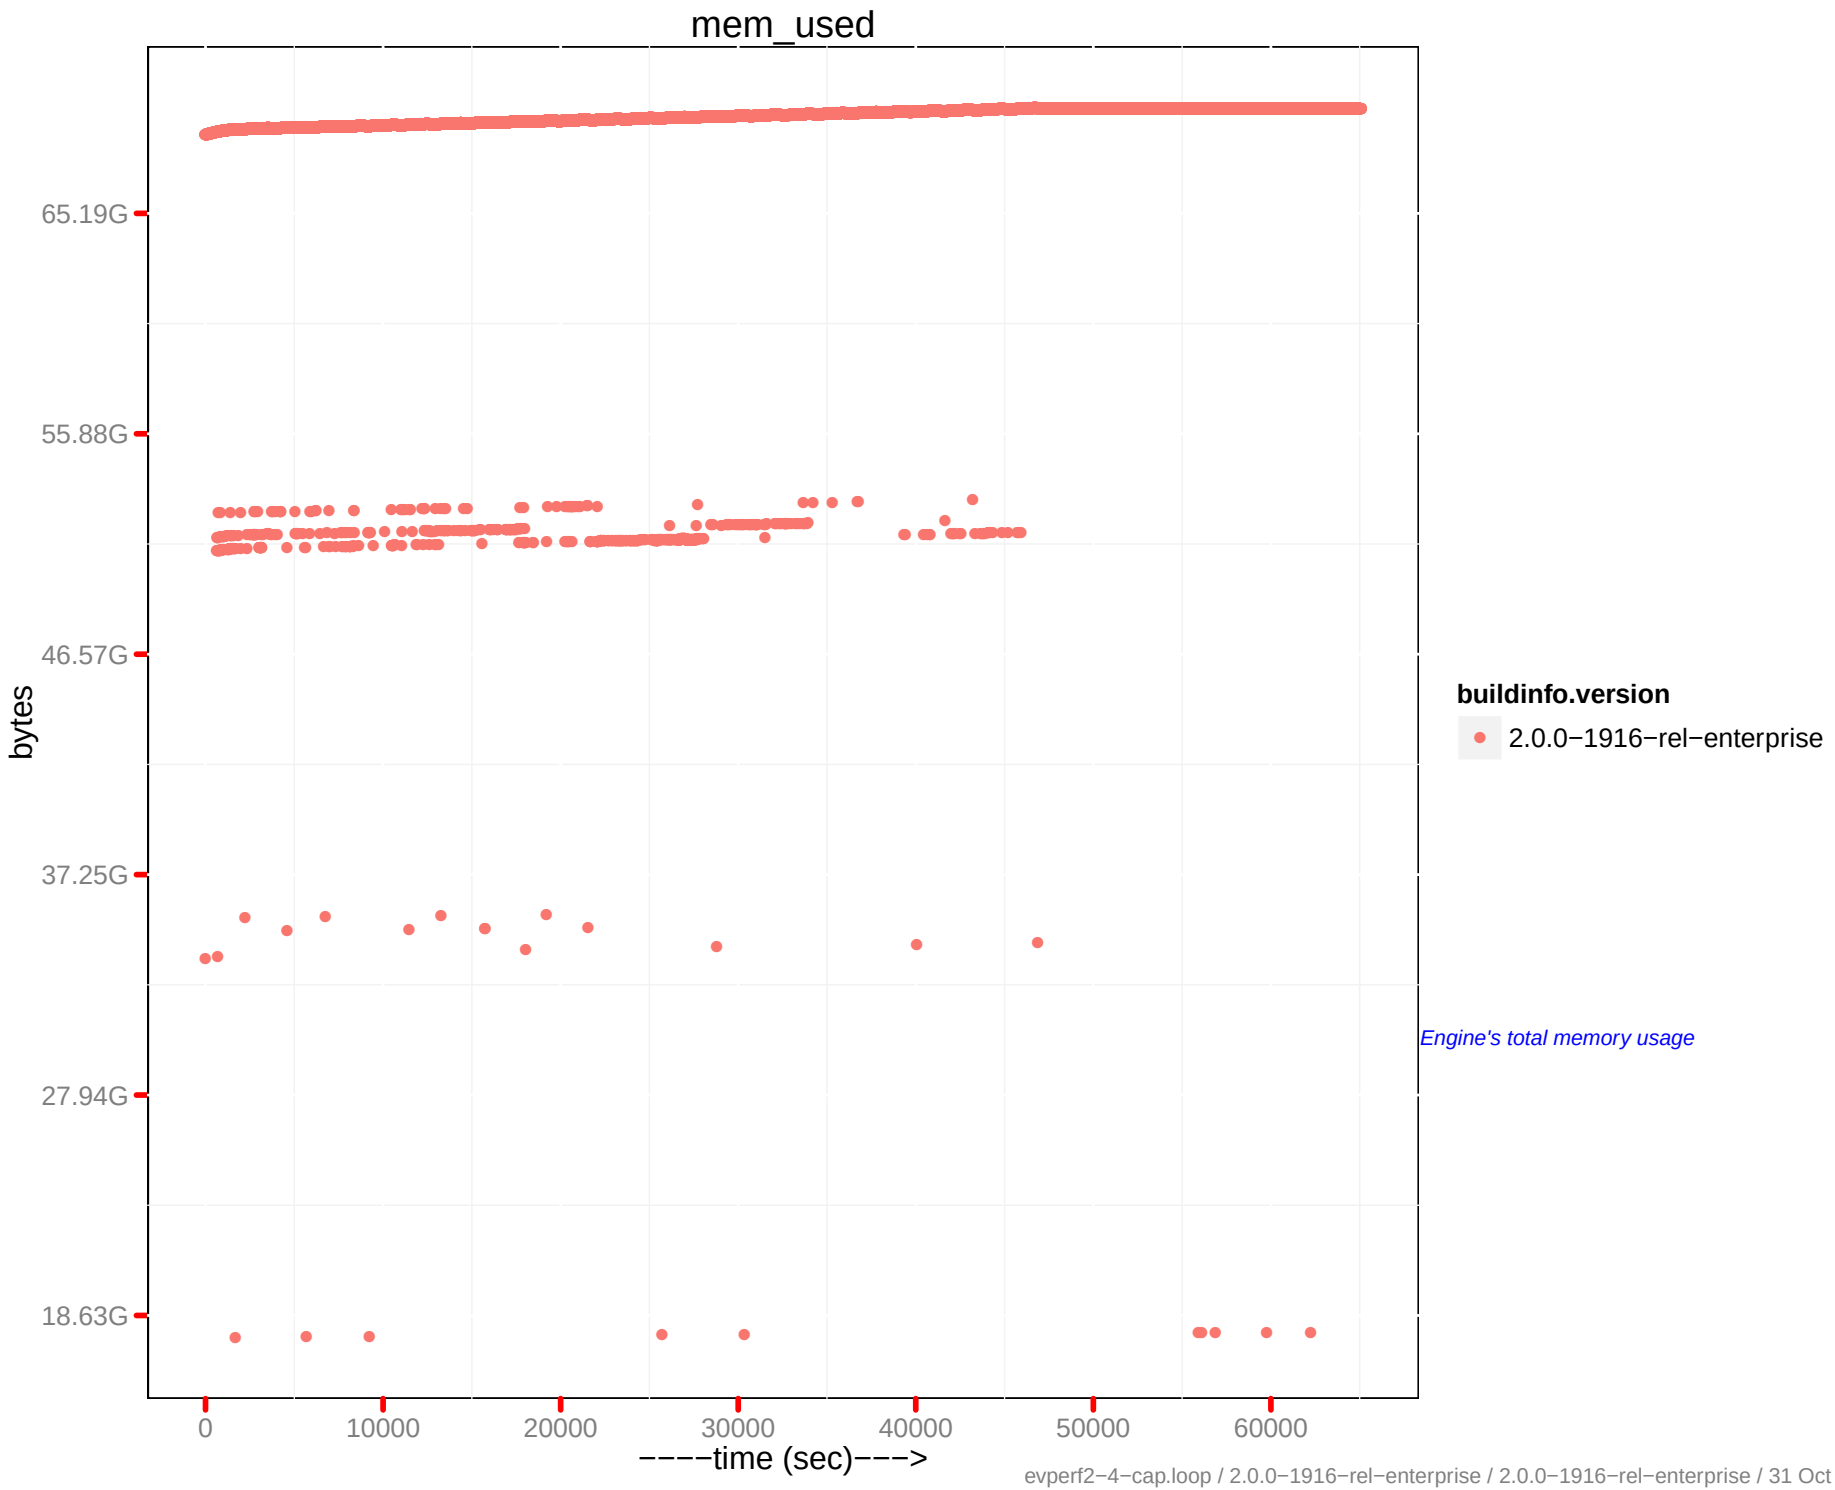
mem_used
65.19G
55.88G
46.57G
37.25G
27.94G
18.63G
bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Engine's total memory usage
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

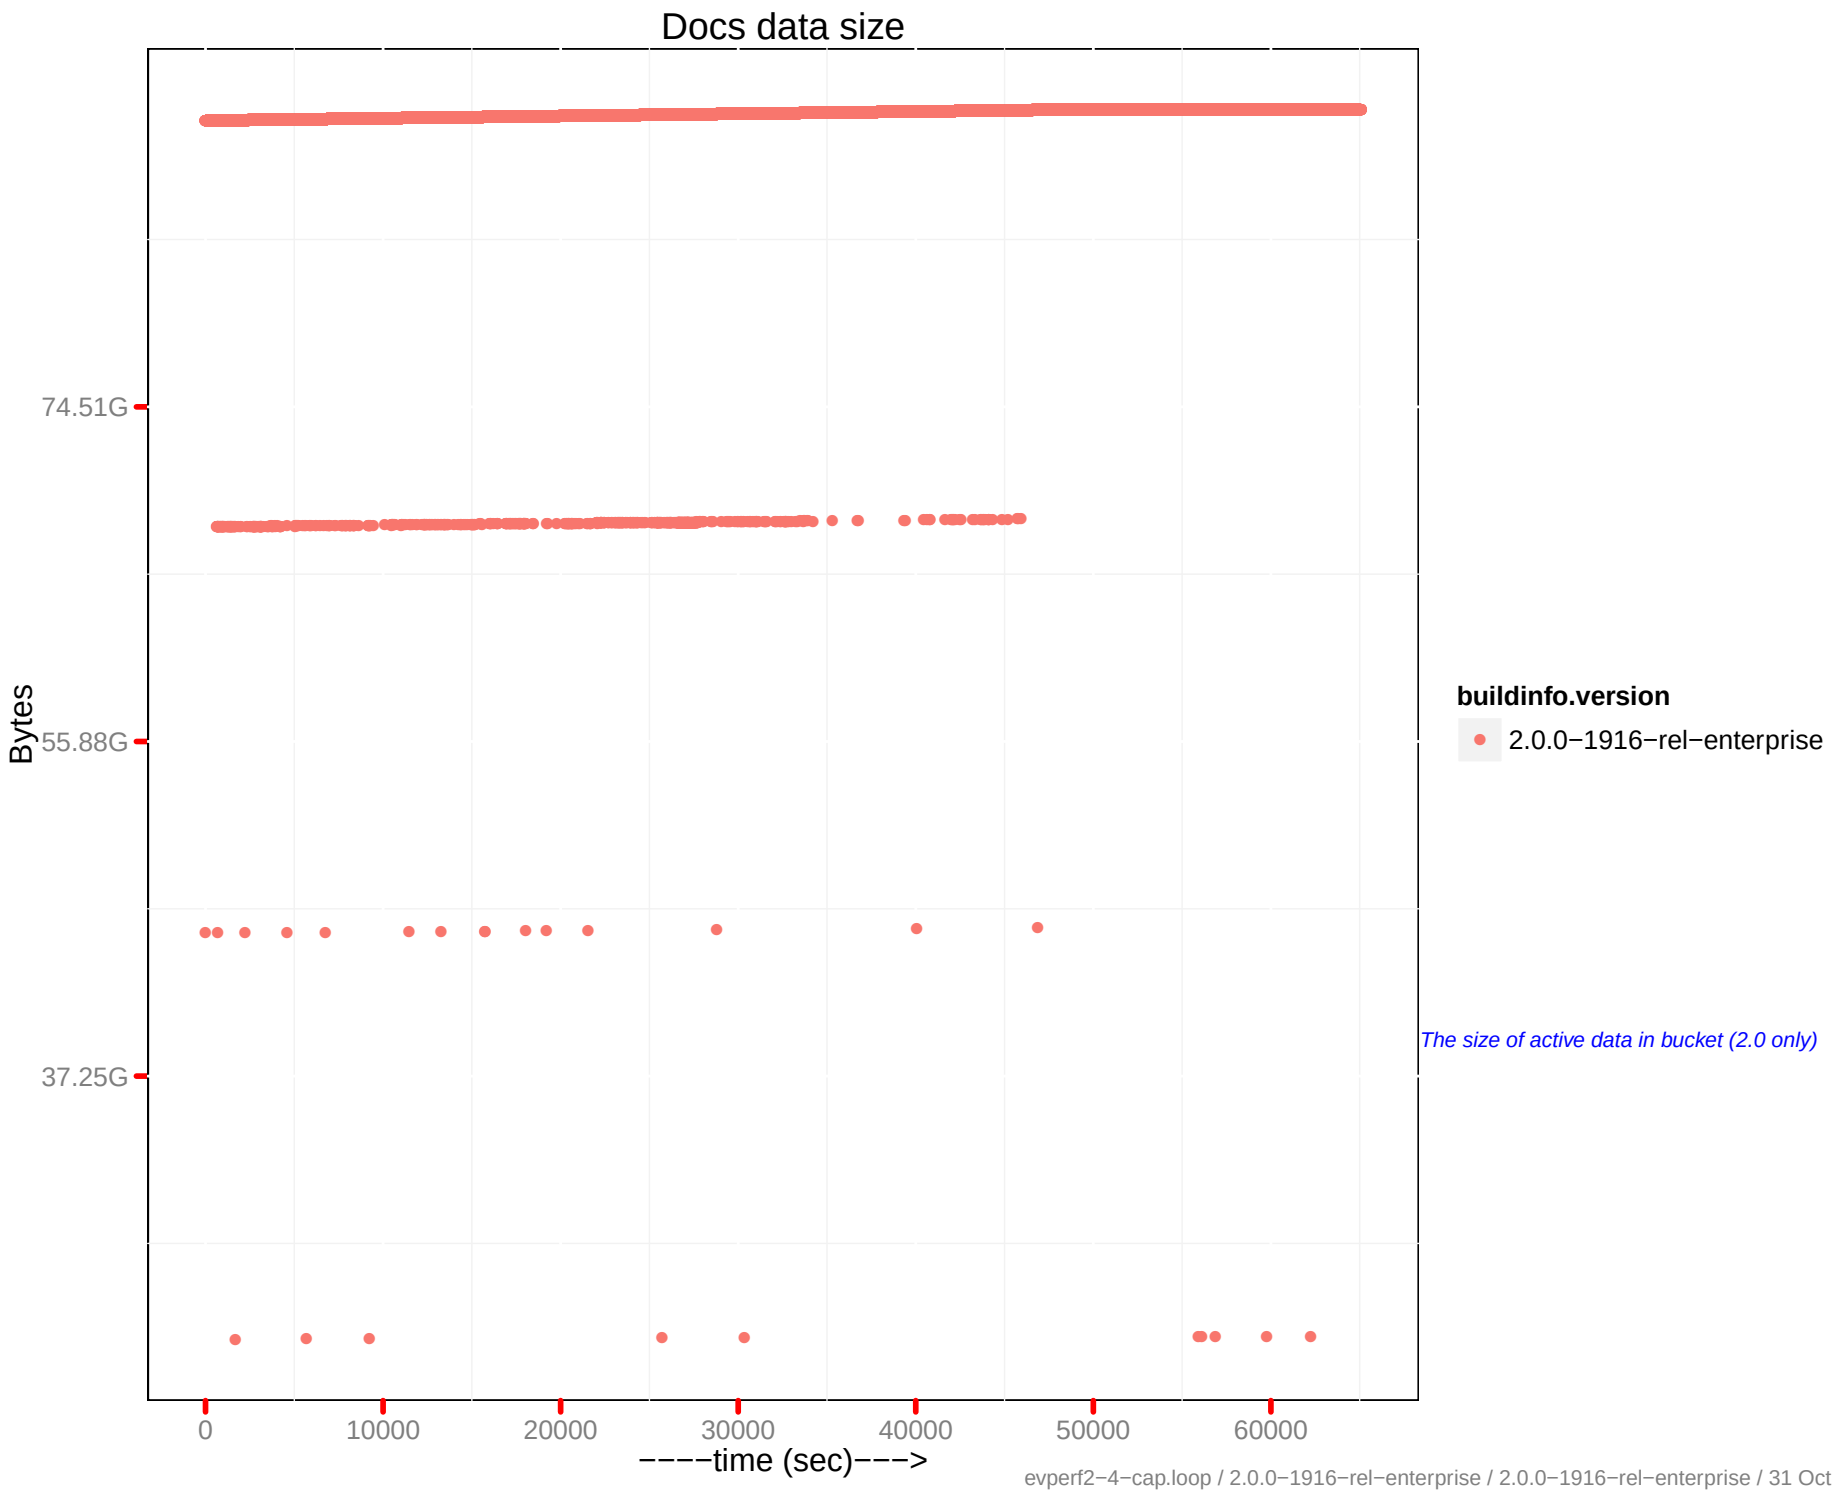
Docs data size
74.51G
55.88G
37.25G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
The size of active data in bucket (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

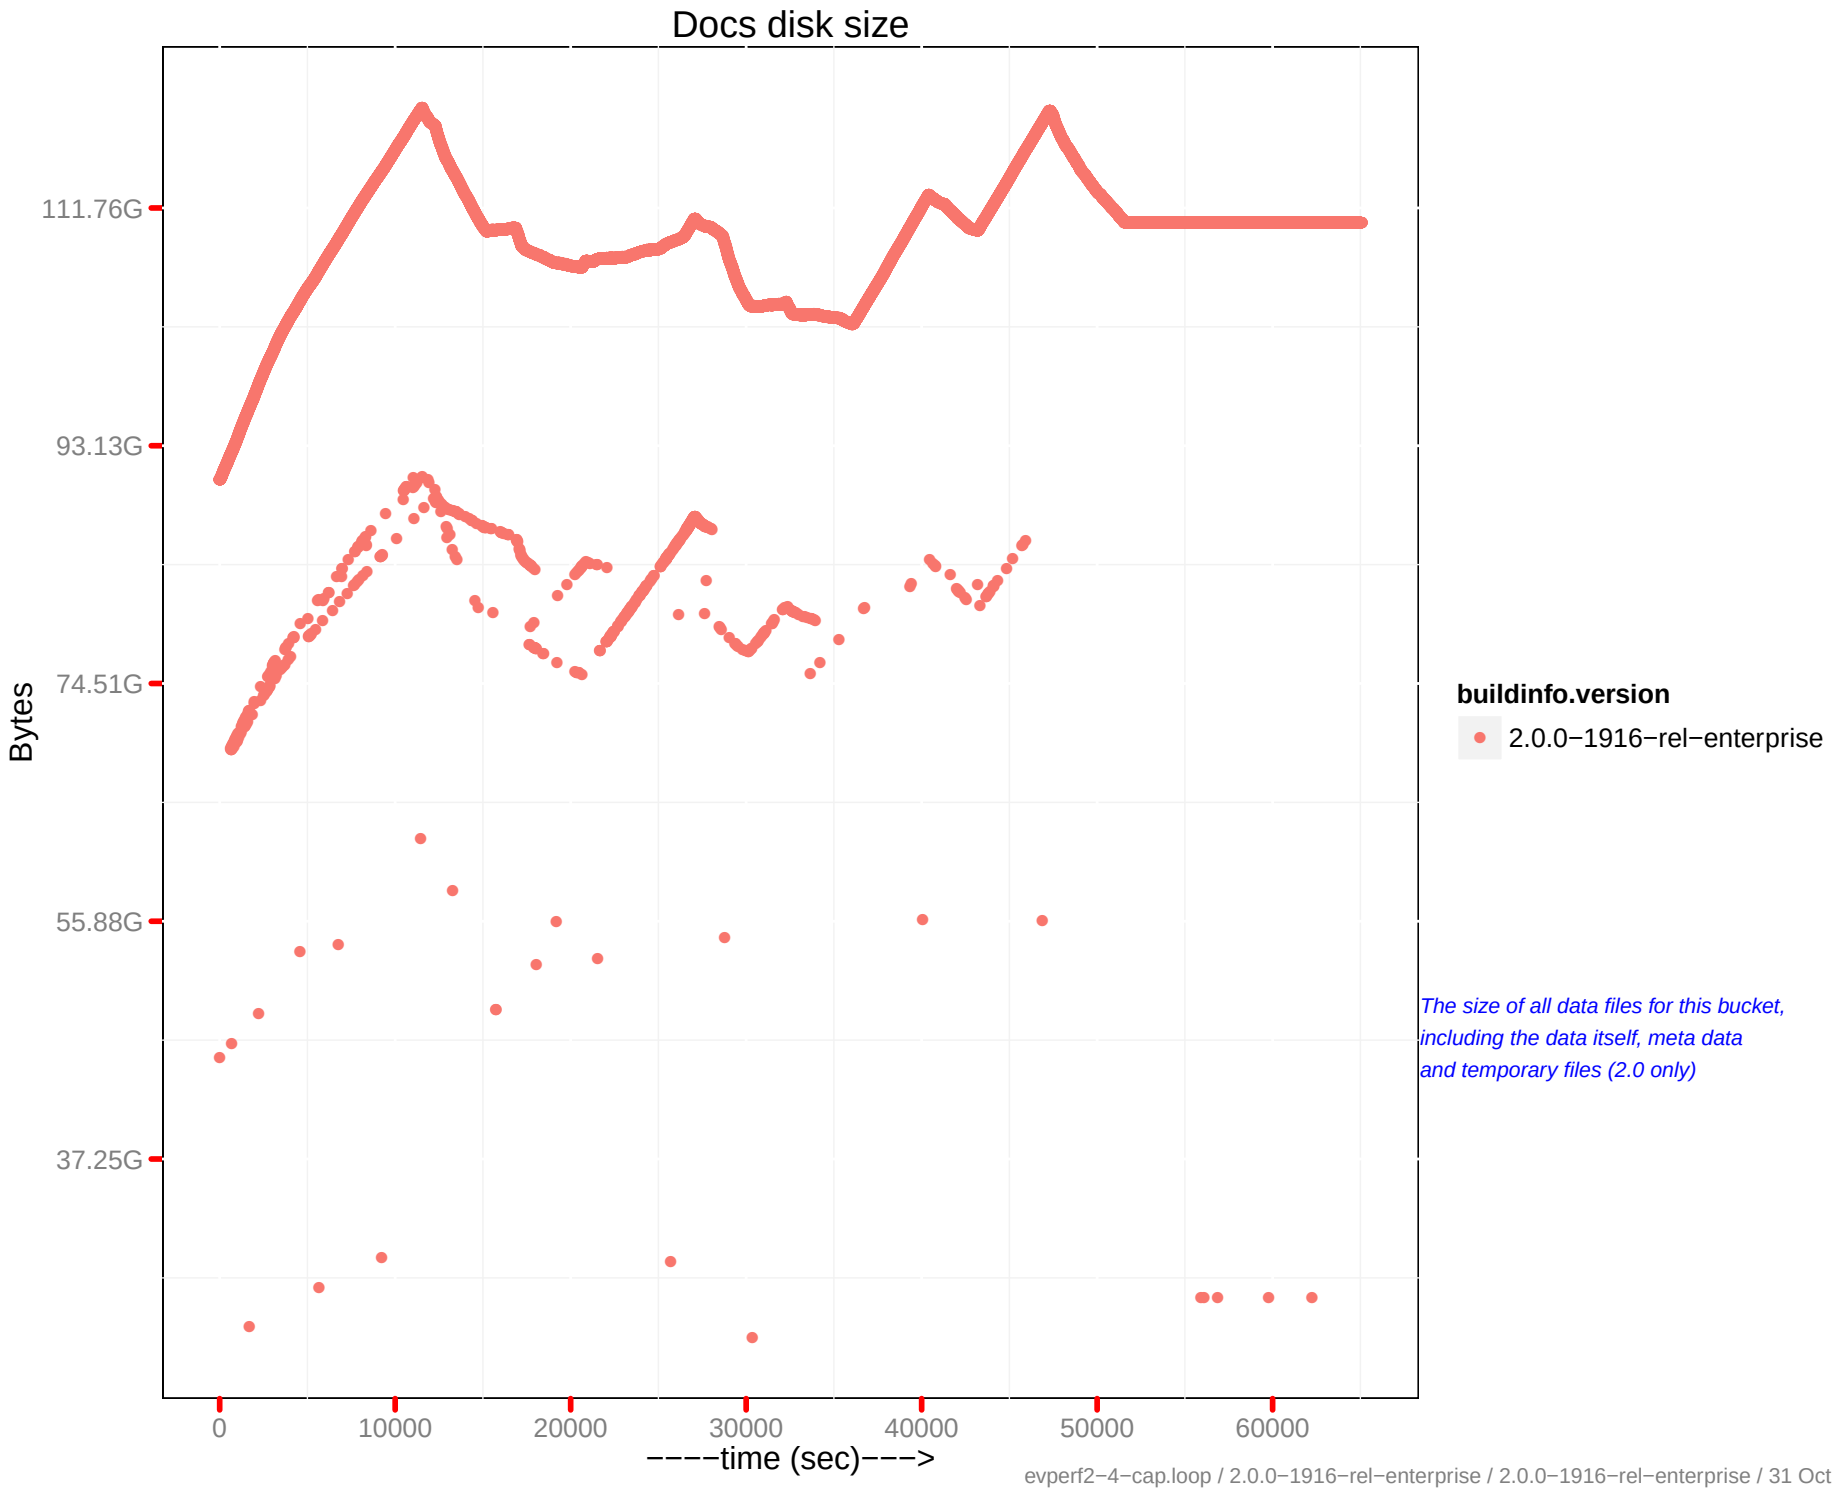
Docs disk size
111.76G
93.13G
74.51G
55.88G
37.25G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

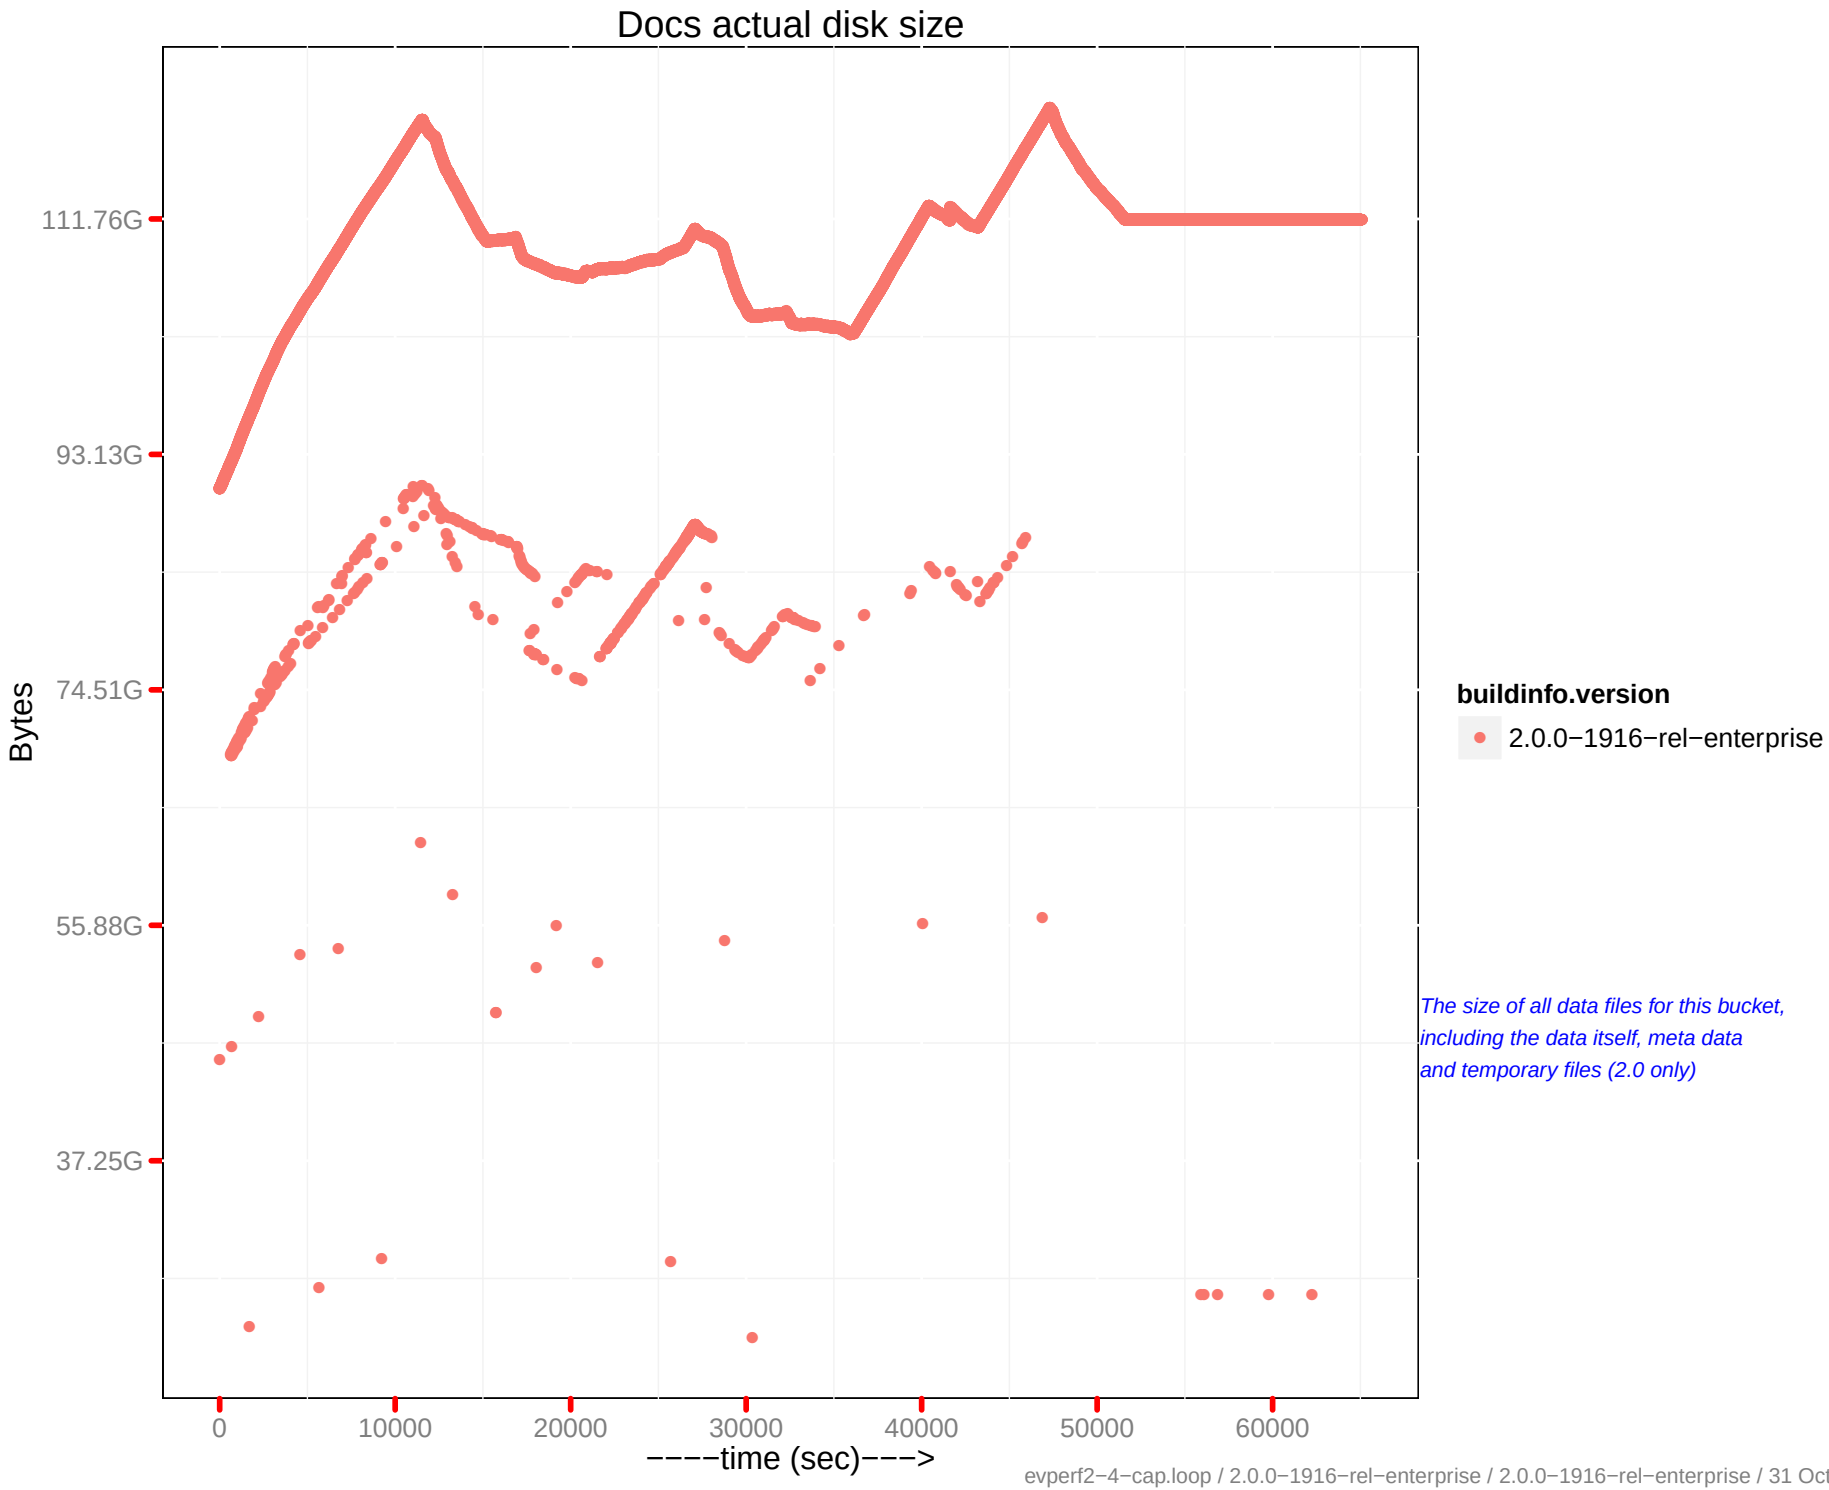
Docs actual disk size
111.76G
93.13G
74.51G
55.88G
37.25G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

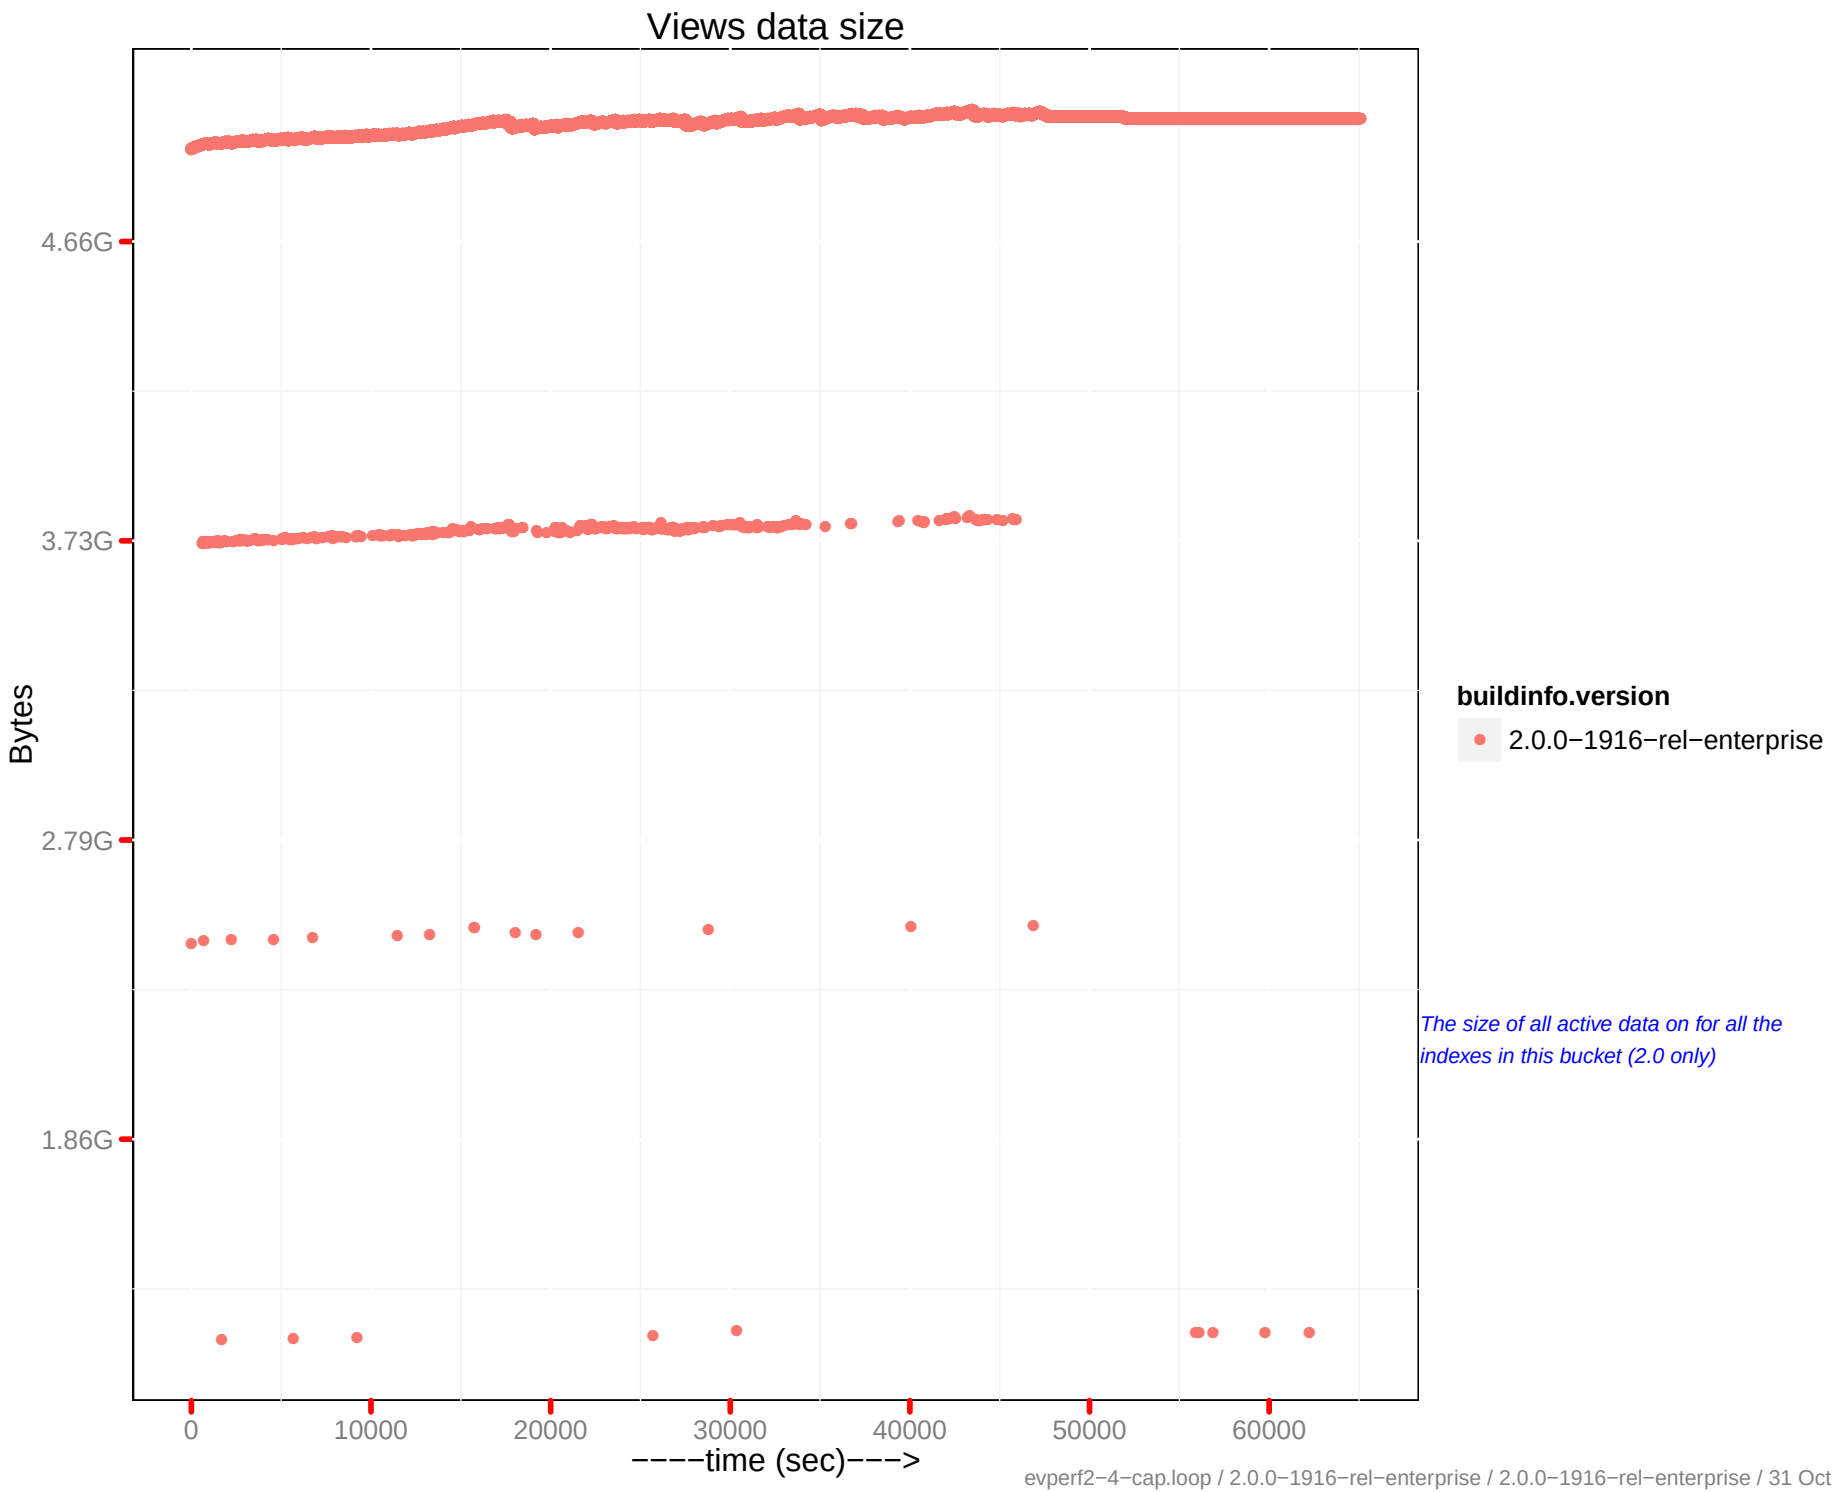
Views data size
4.66G
3.73G
2.79G
1.86G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

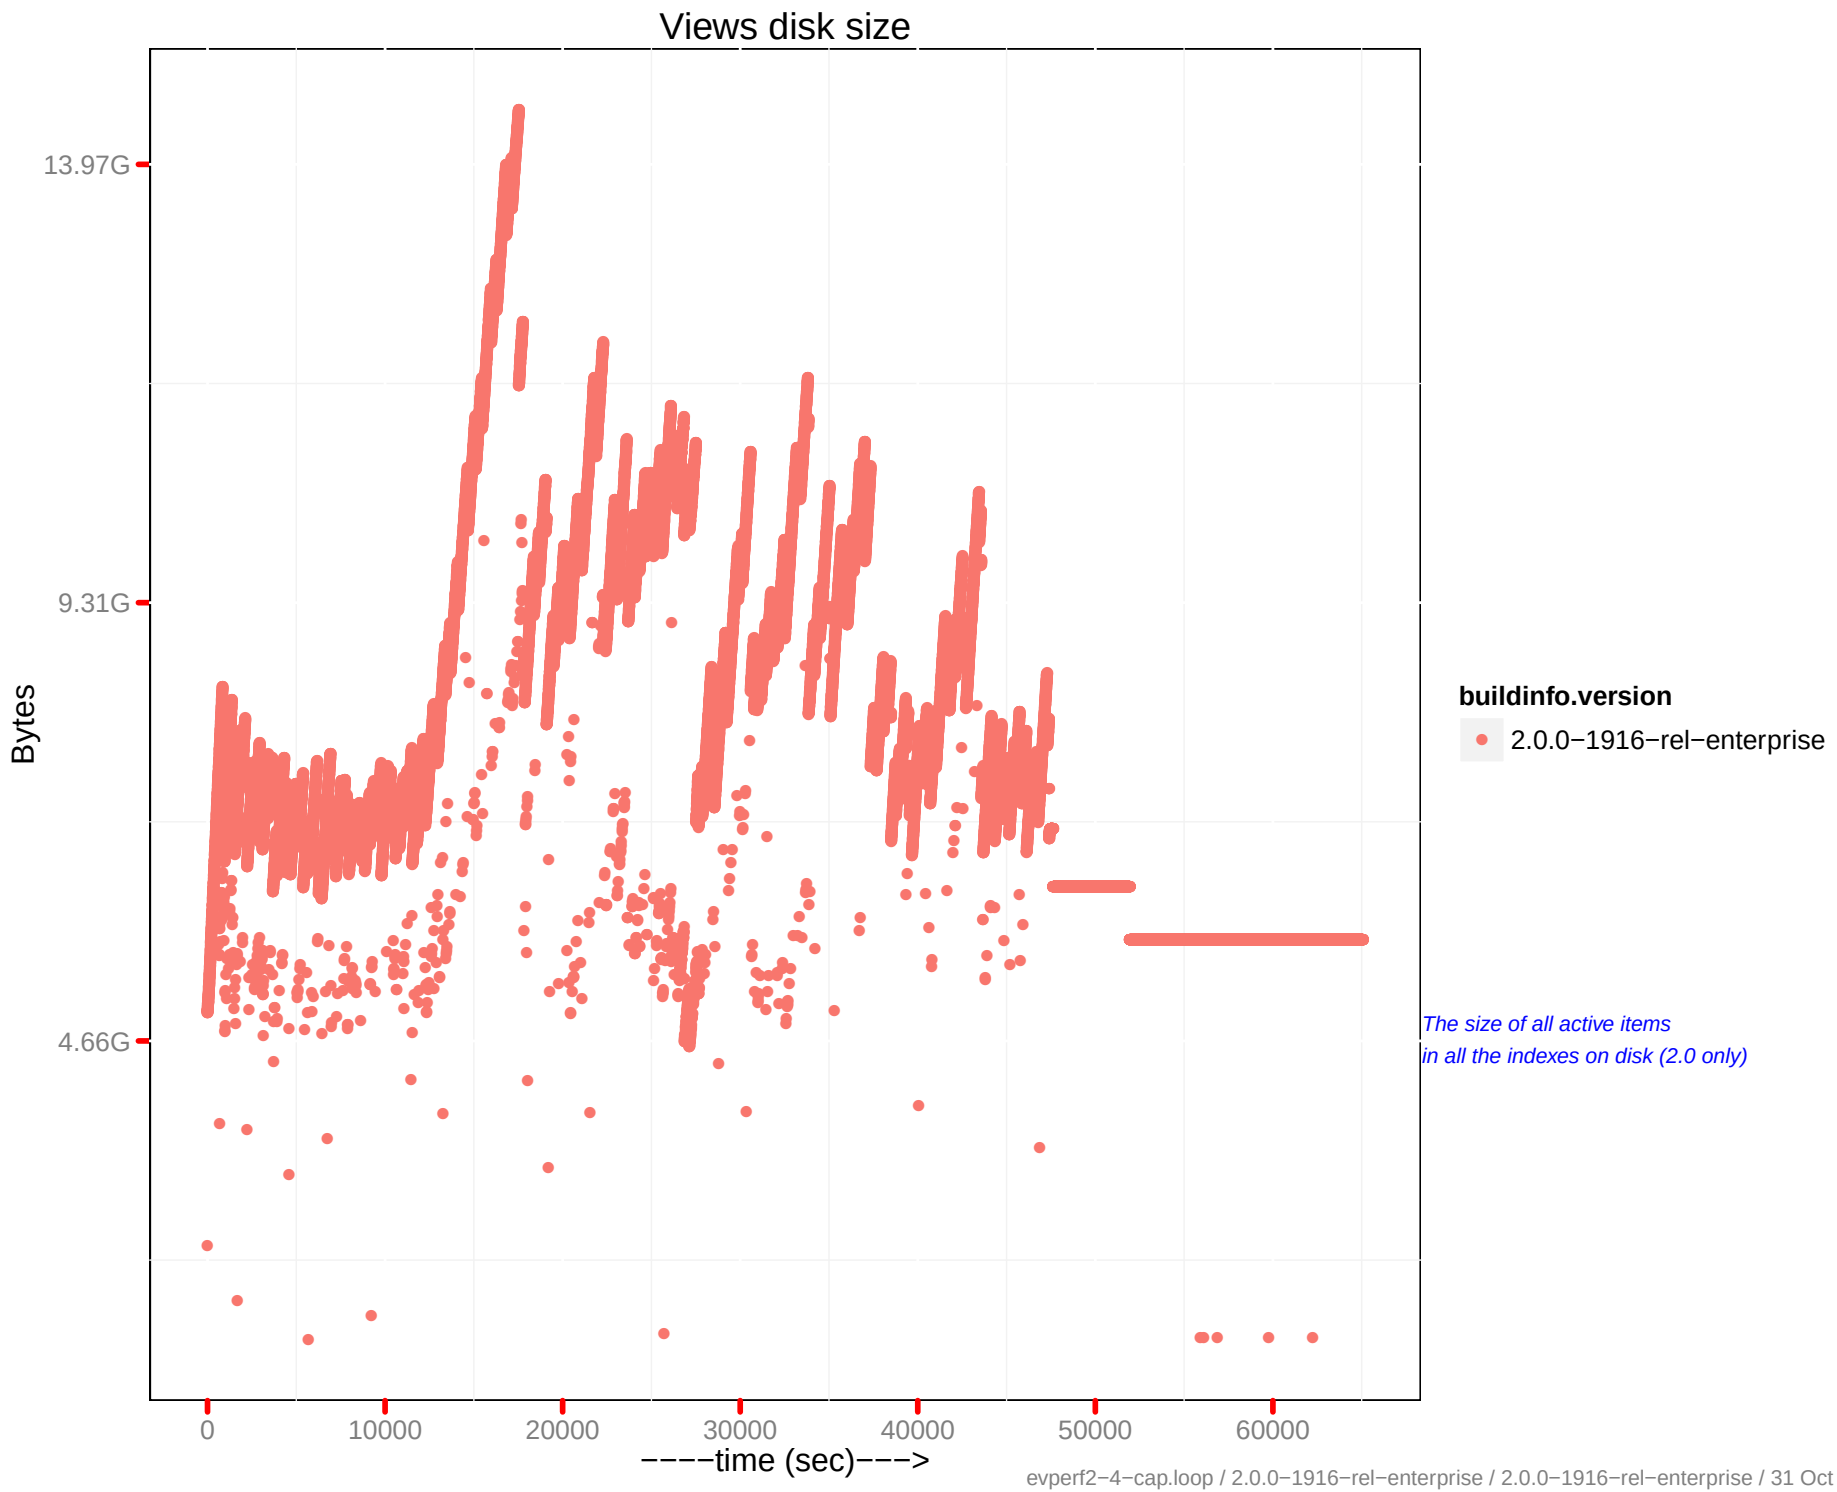
Views disk size
13.97G
9.31G
4.66G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

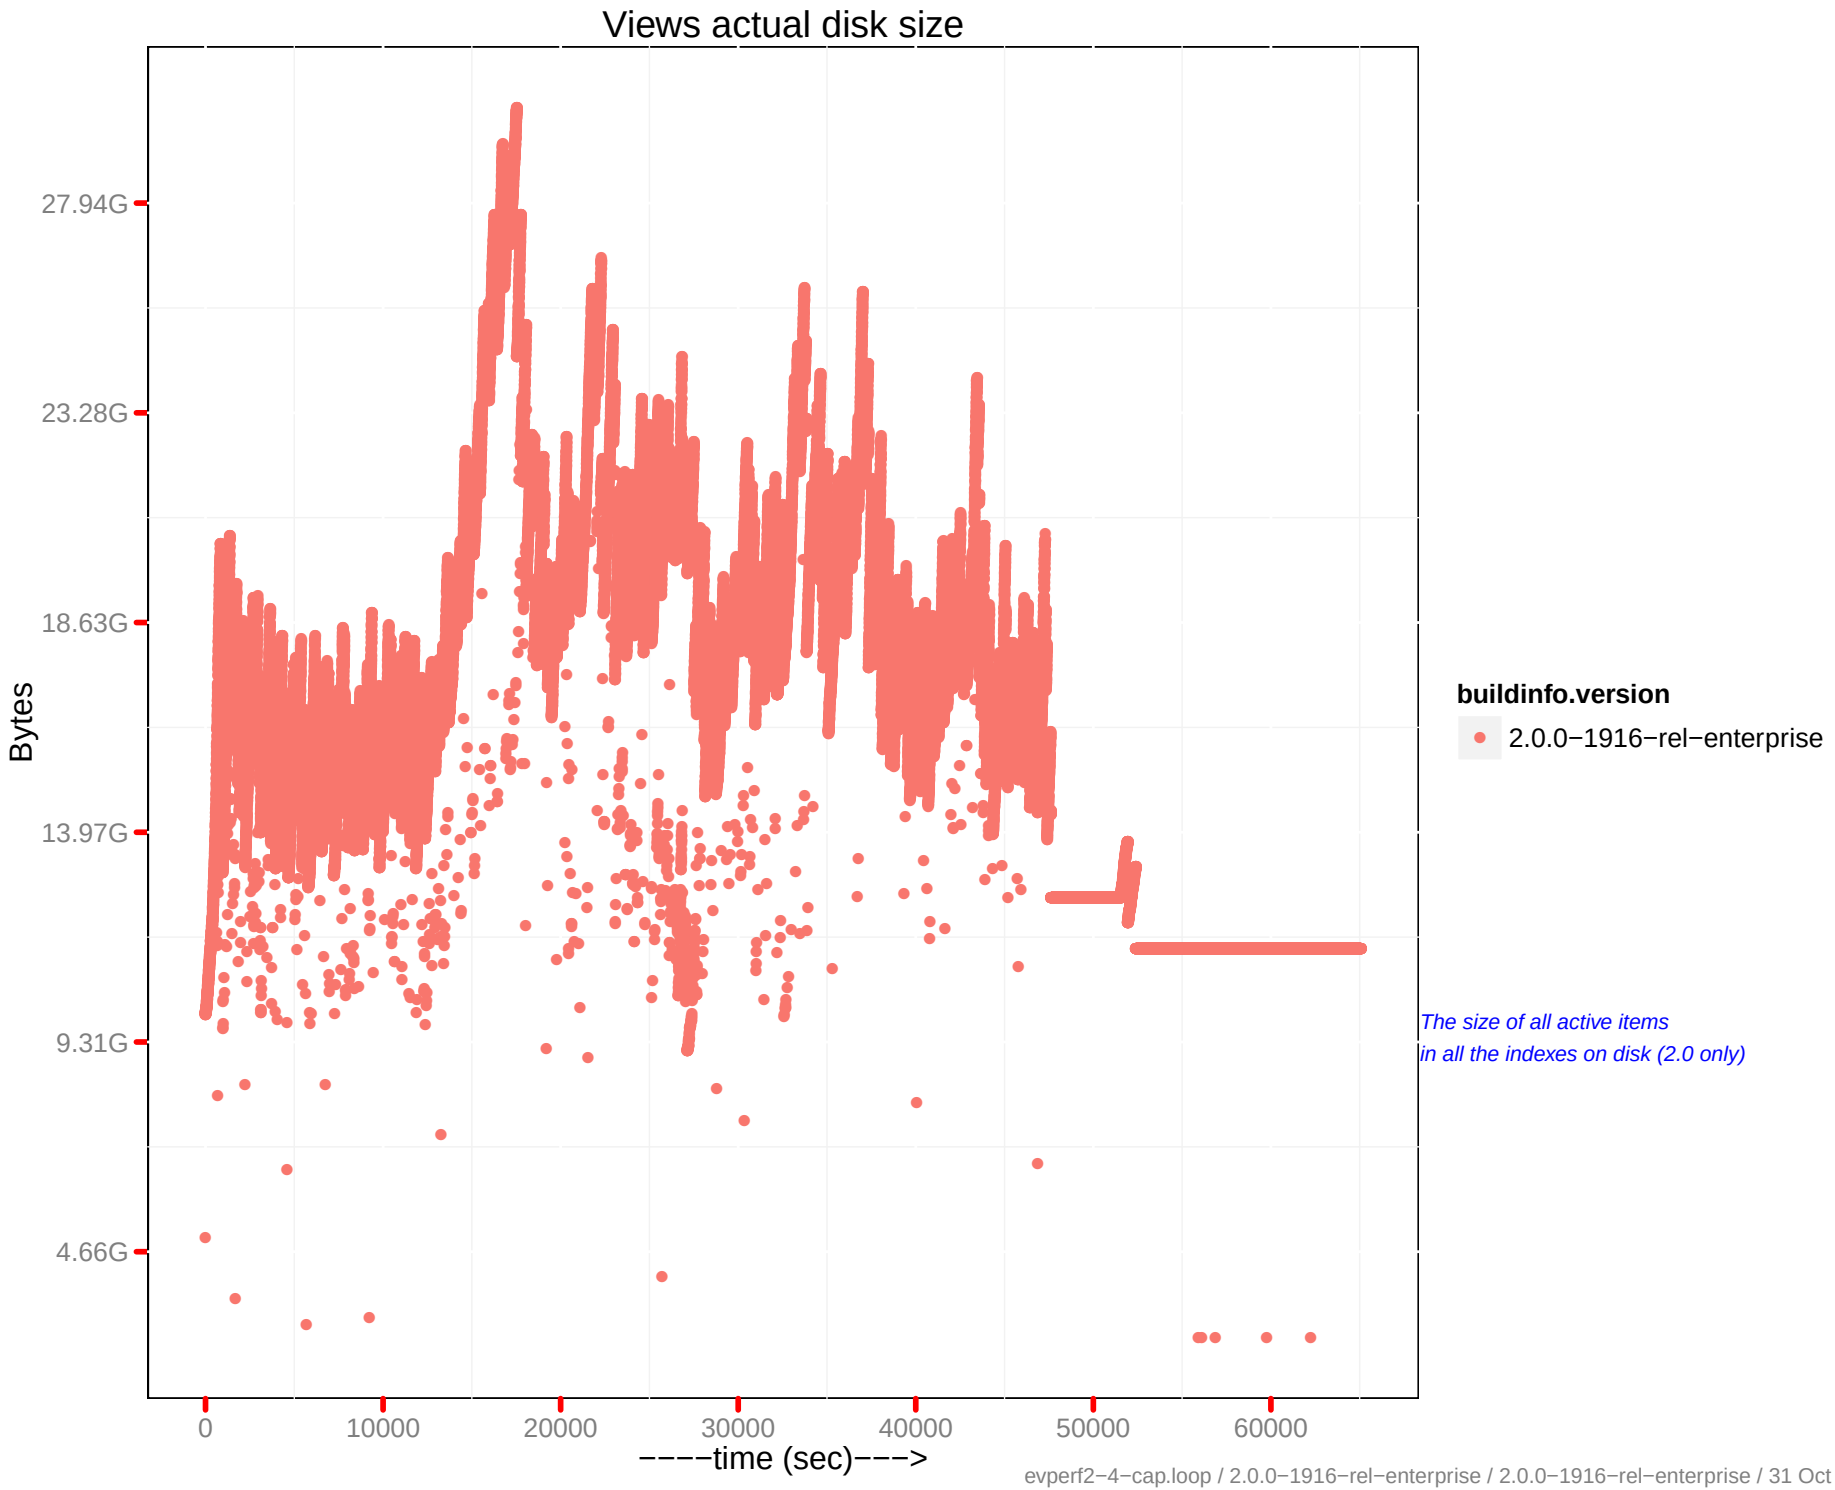
Views actual disk size
27.94G
23.28G
18.63G
13.97G
9.31G
4.66G
Bytes
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012


Total 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
2.0.0-1916-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

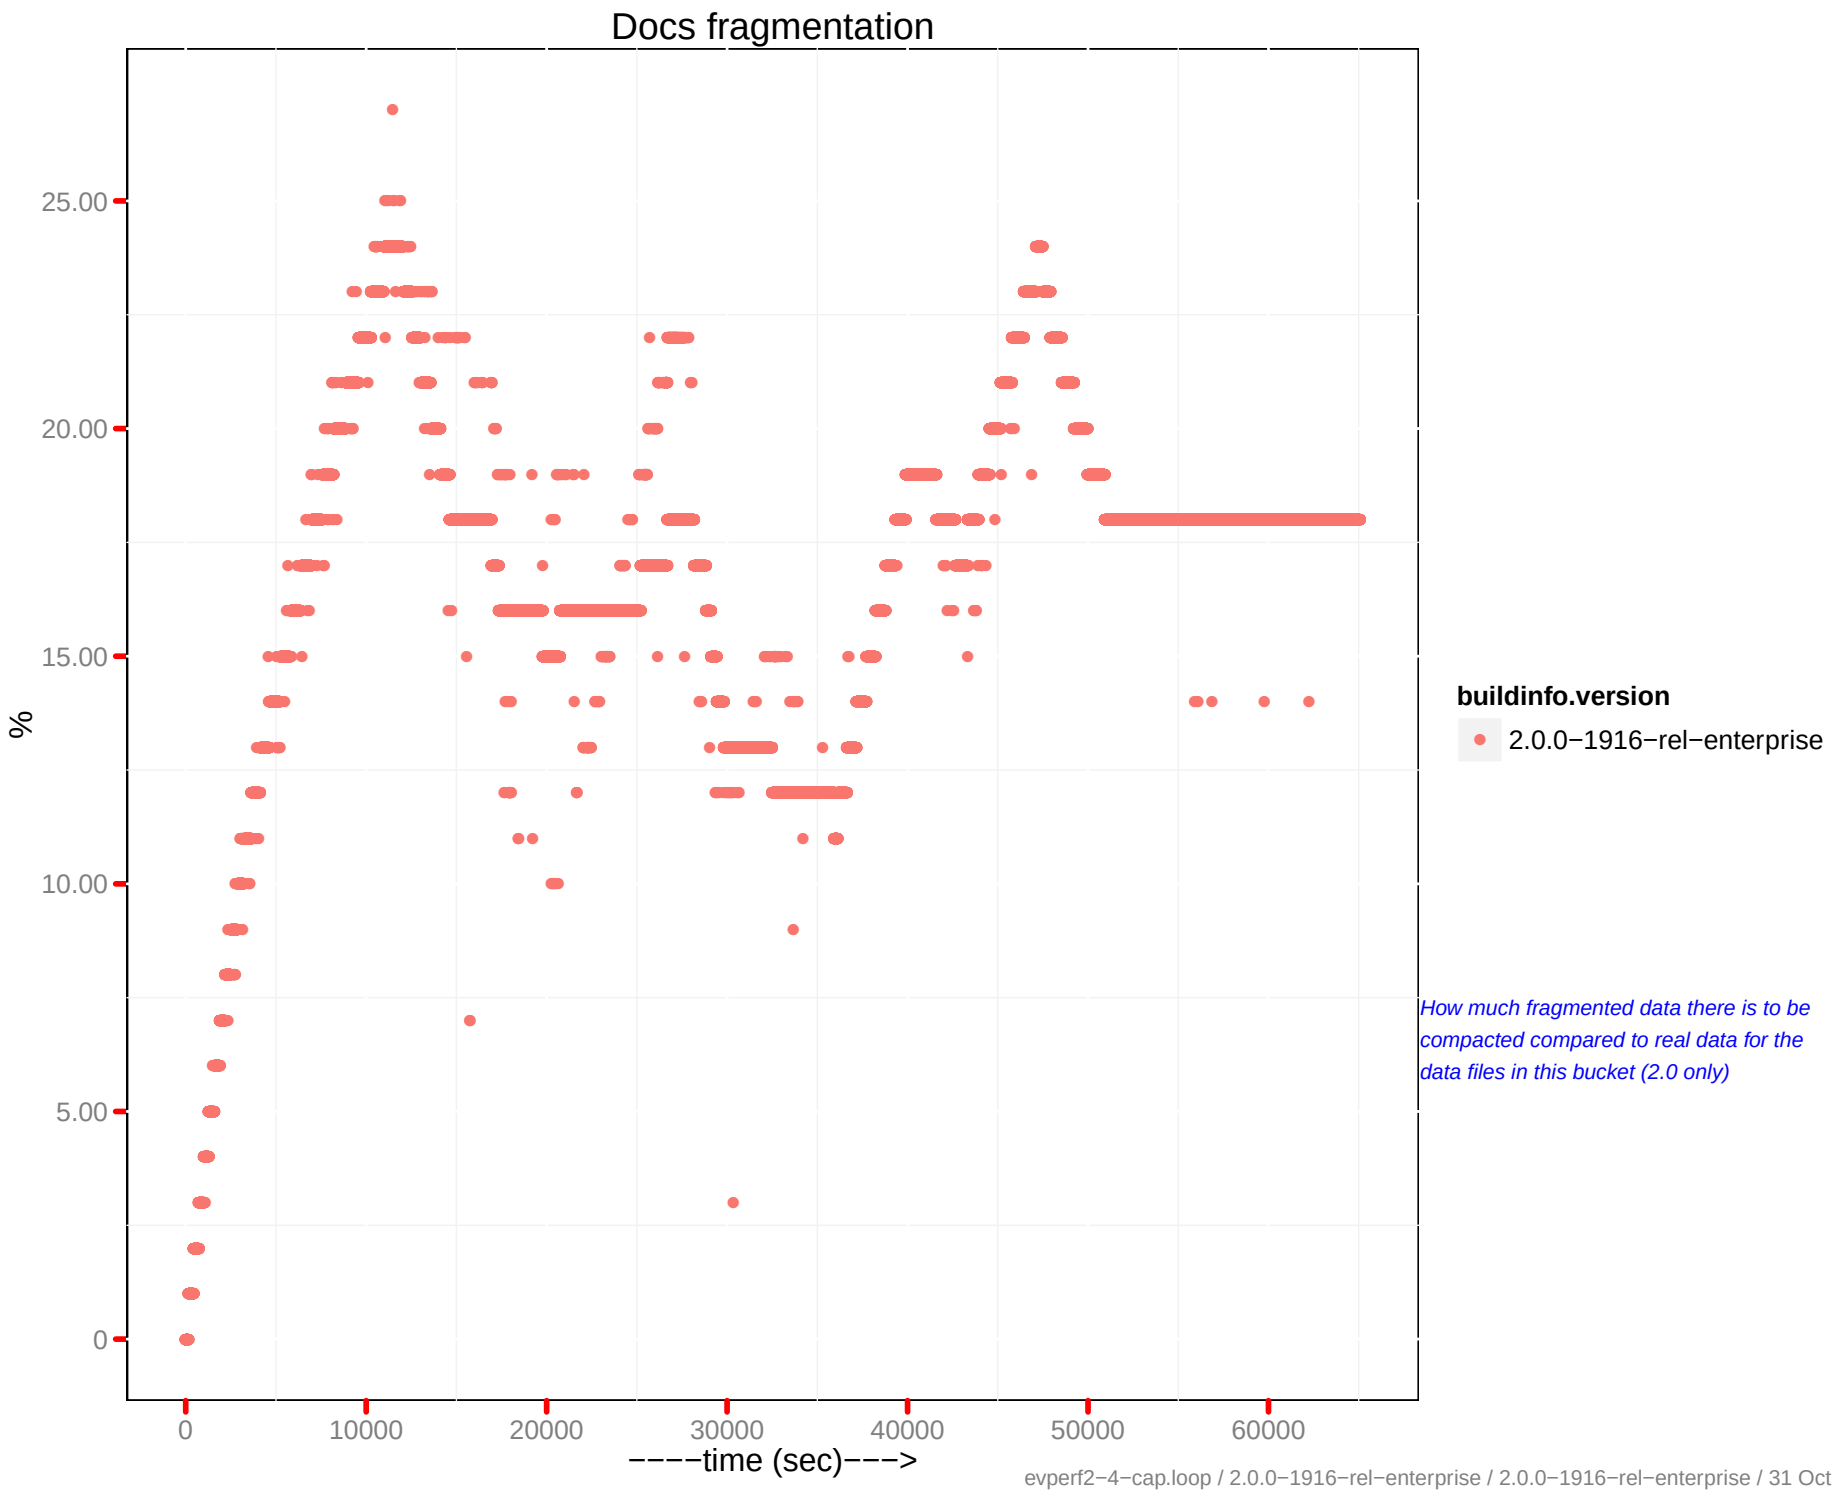
Docs fragmentation
25.00
20.00
15.00
10.00
5.00
0
%
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

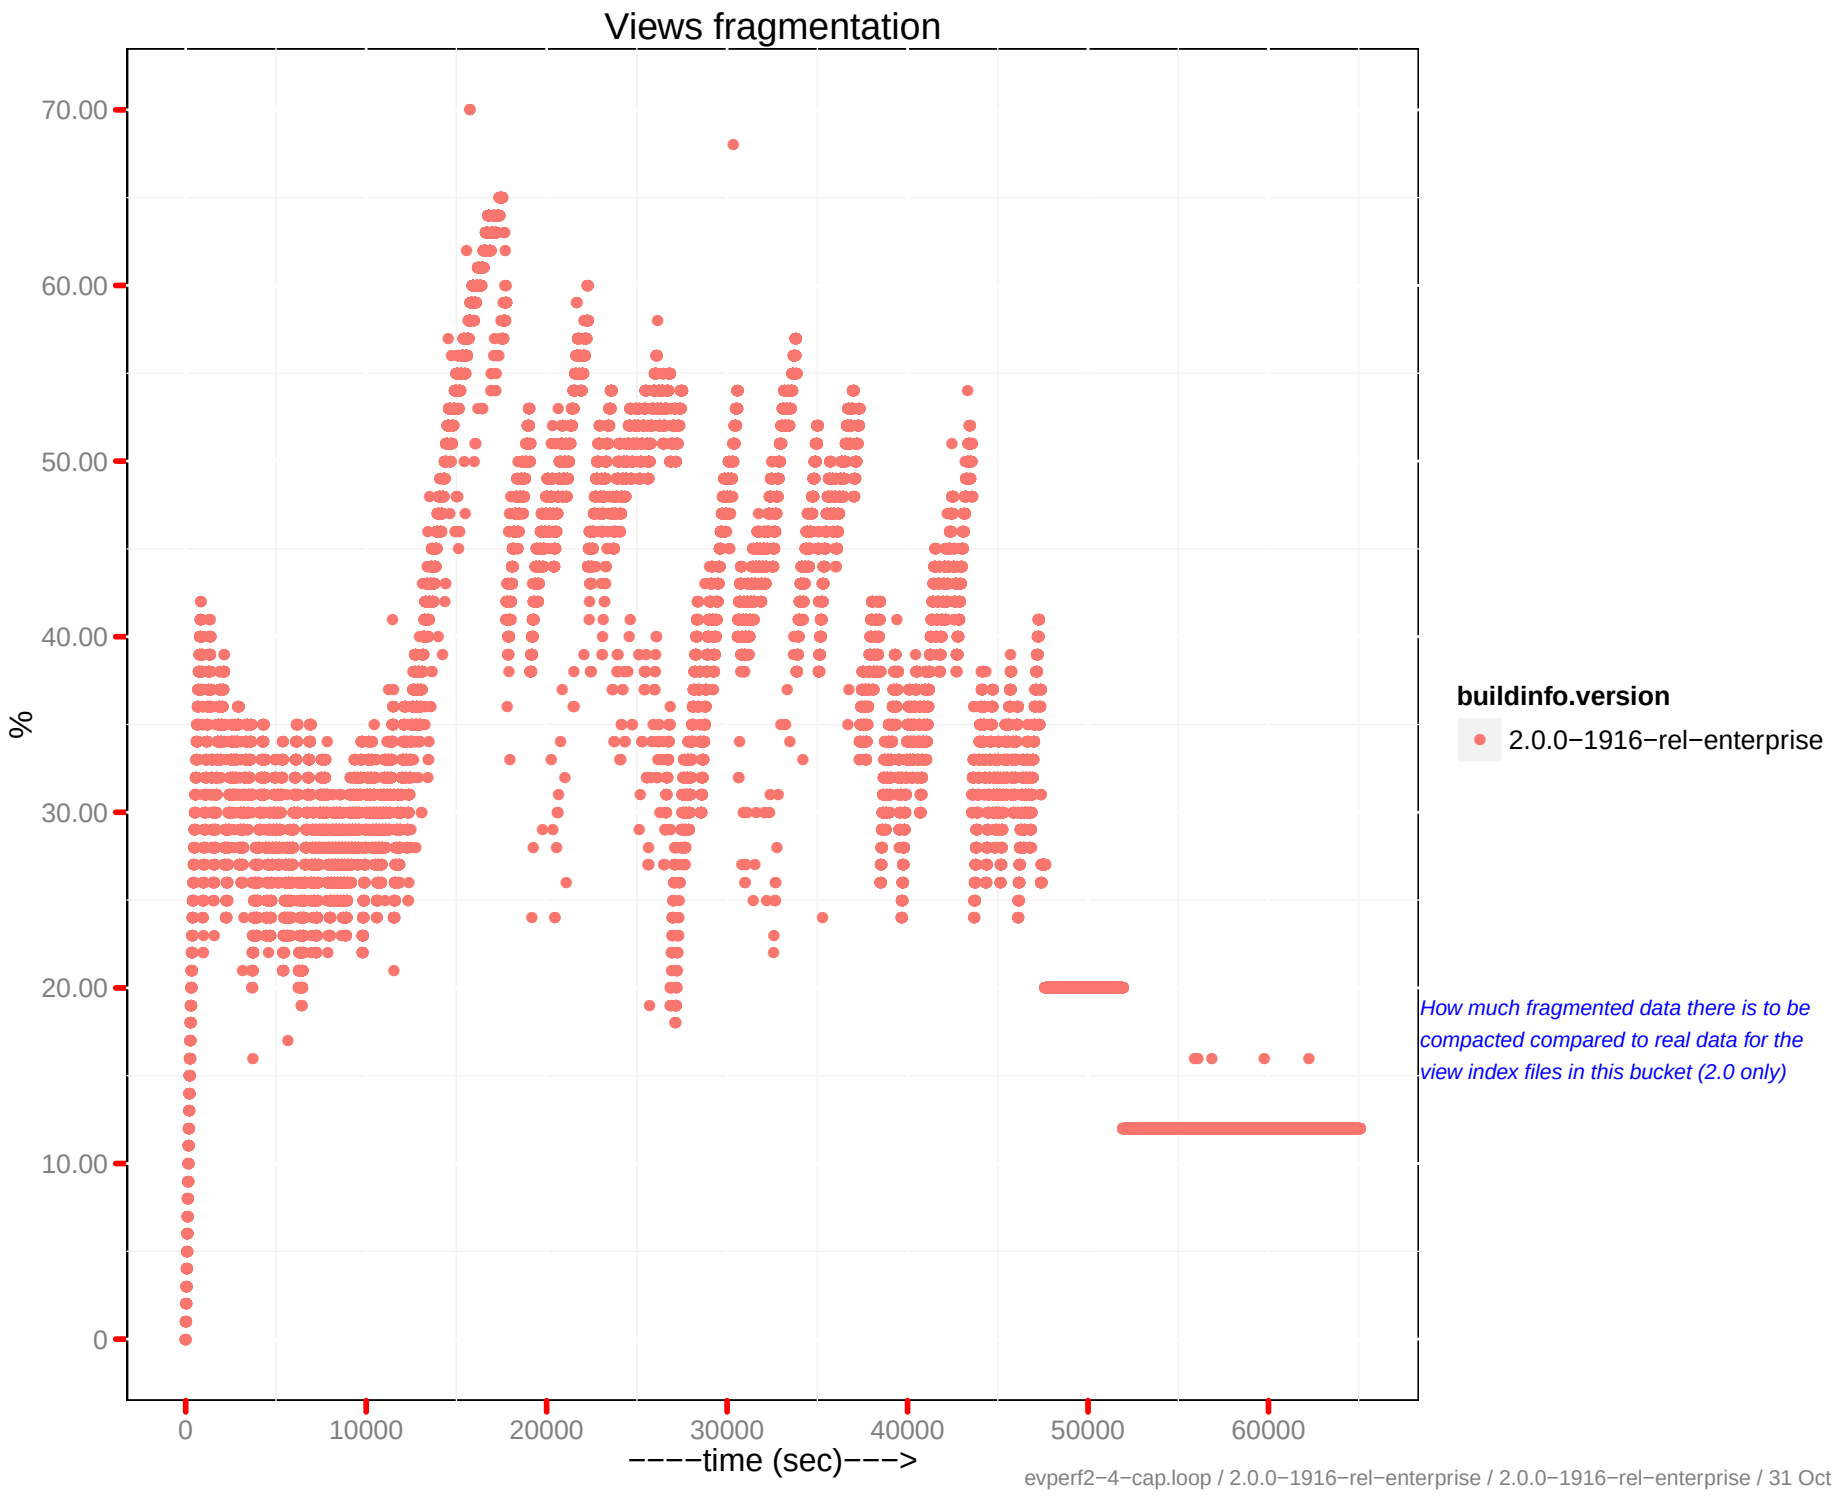
Views fragmentation
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0
%
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012


cmd_get ops/sec
11.72K
9.77K
7.81K
5.86K
3.91K
1.95K
0
ops/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Cumulative number of retrieval reqs
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012


cmd_set ops/sec
1.95K
1.46K
1000.00
500.00
0
ops/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Cumulative number of set reqs
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

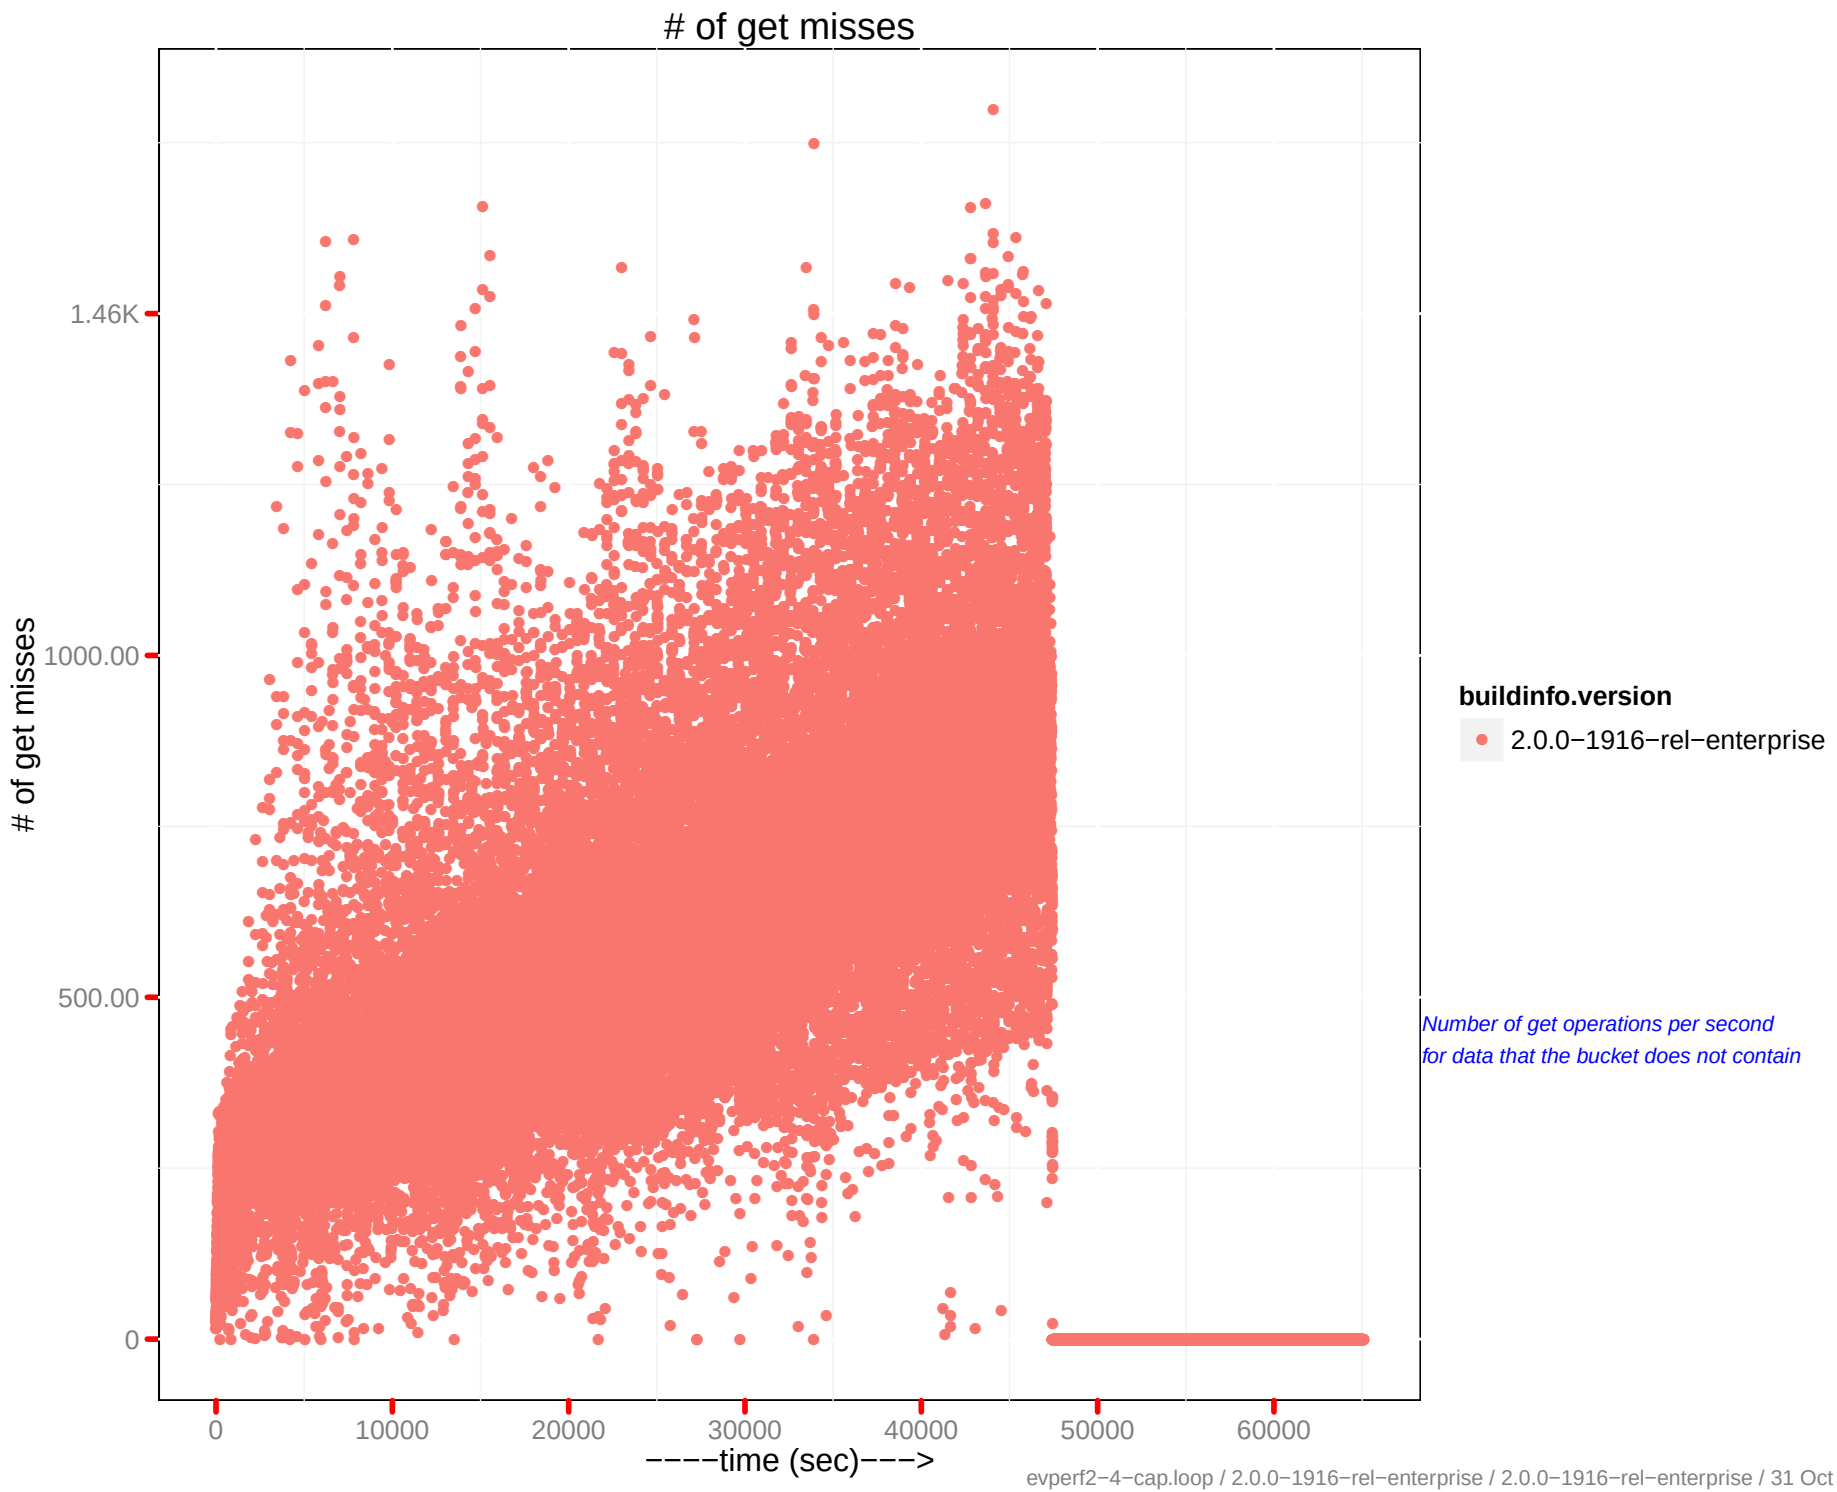
# of get misses
1.46K
1000.00
500.00
0
# of get misses
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of get operations per second
for data that the bucket does not contain
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

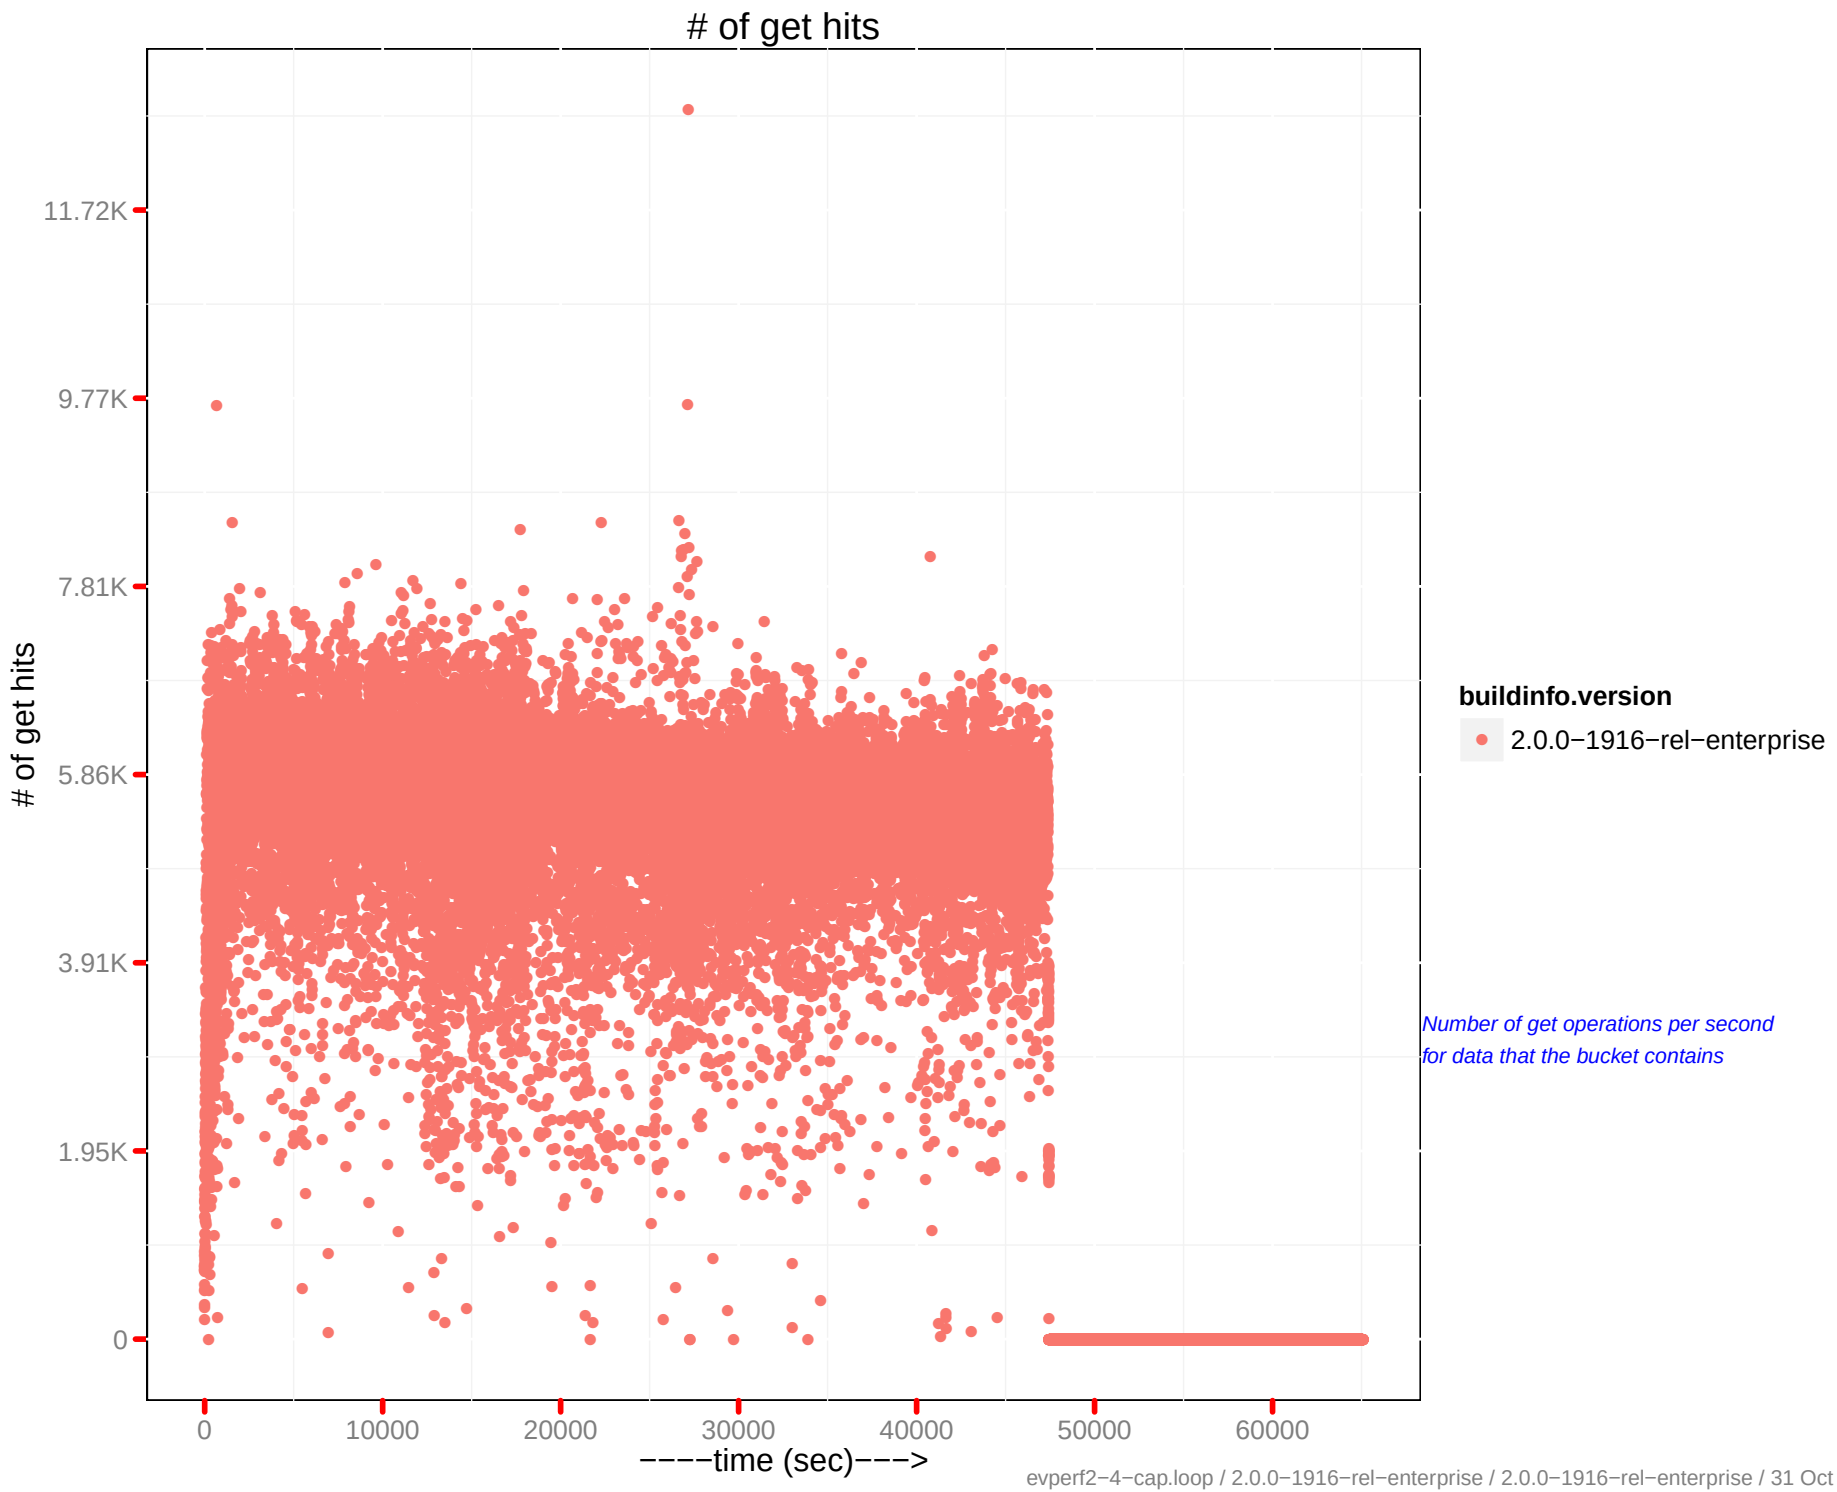
# of get hits
11.72K
9.77K
7.81K
5.86K
3.91K
1.95K
0
# of get hits
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of get operations per second
for data that the bucket contains
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

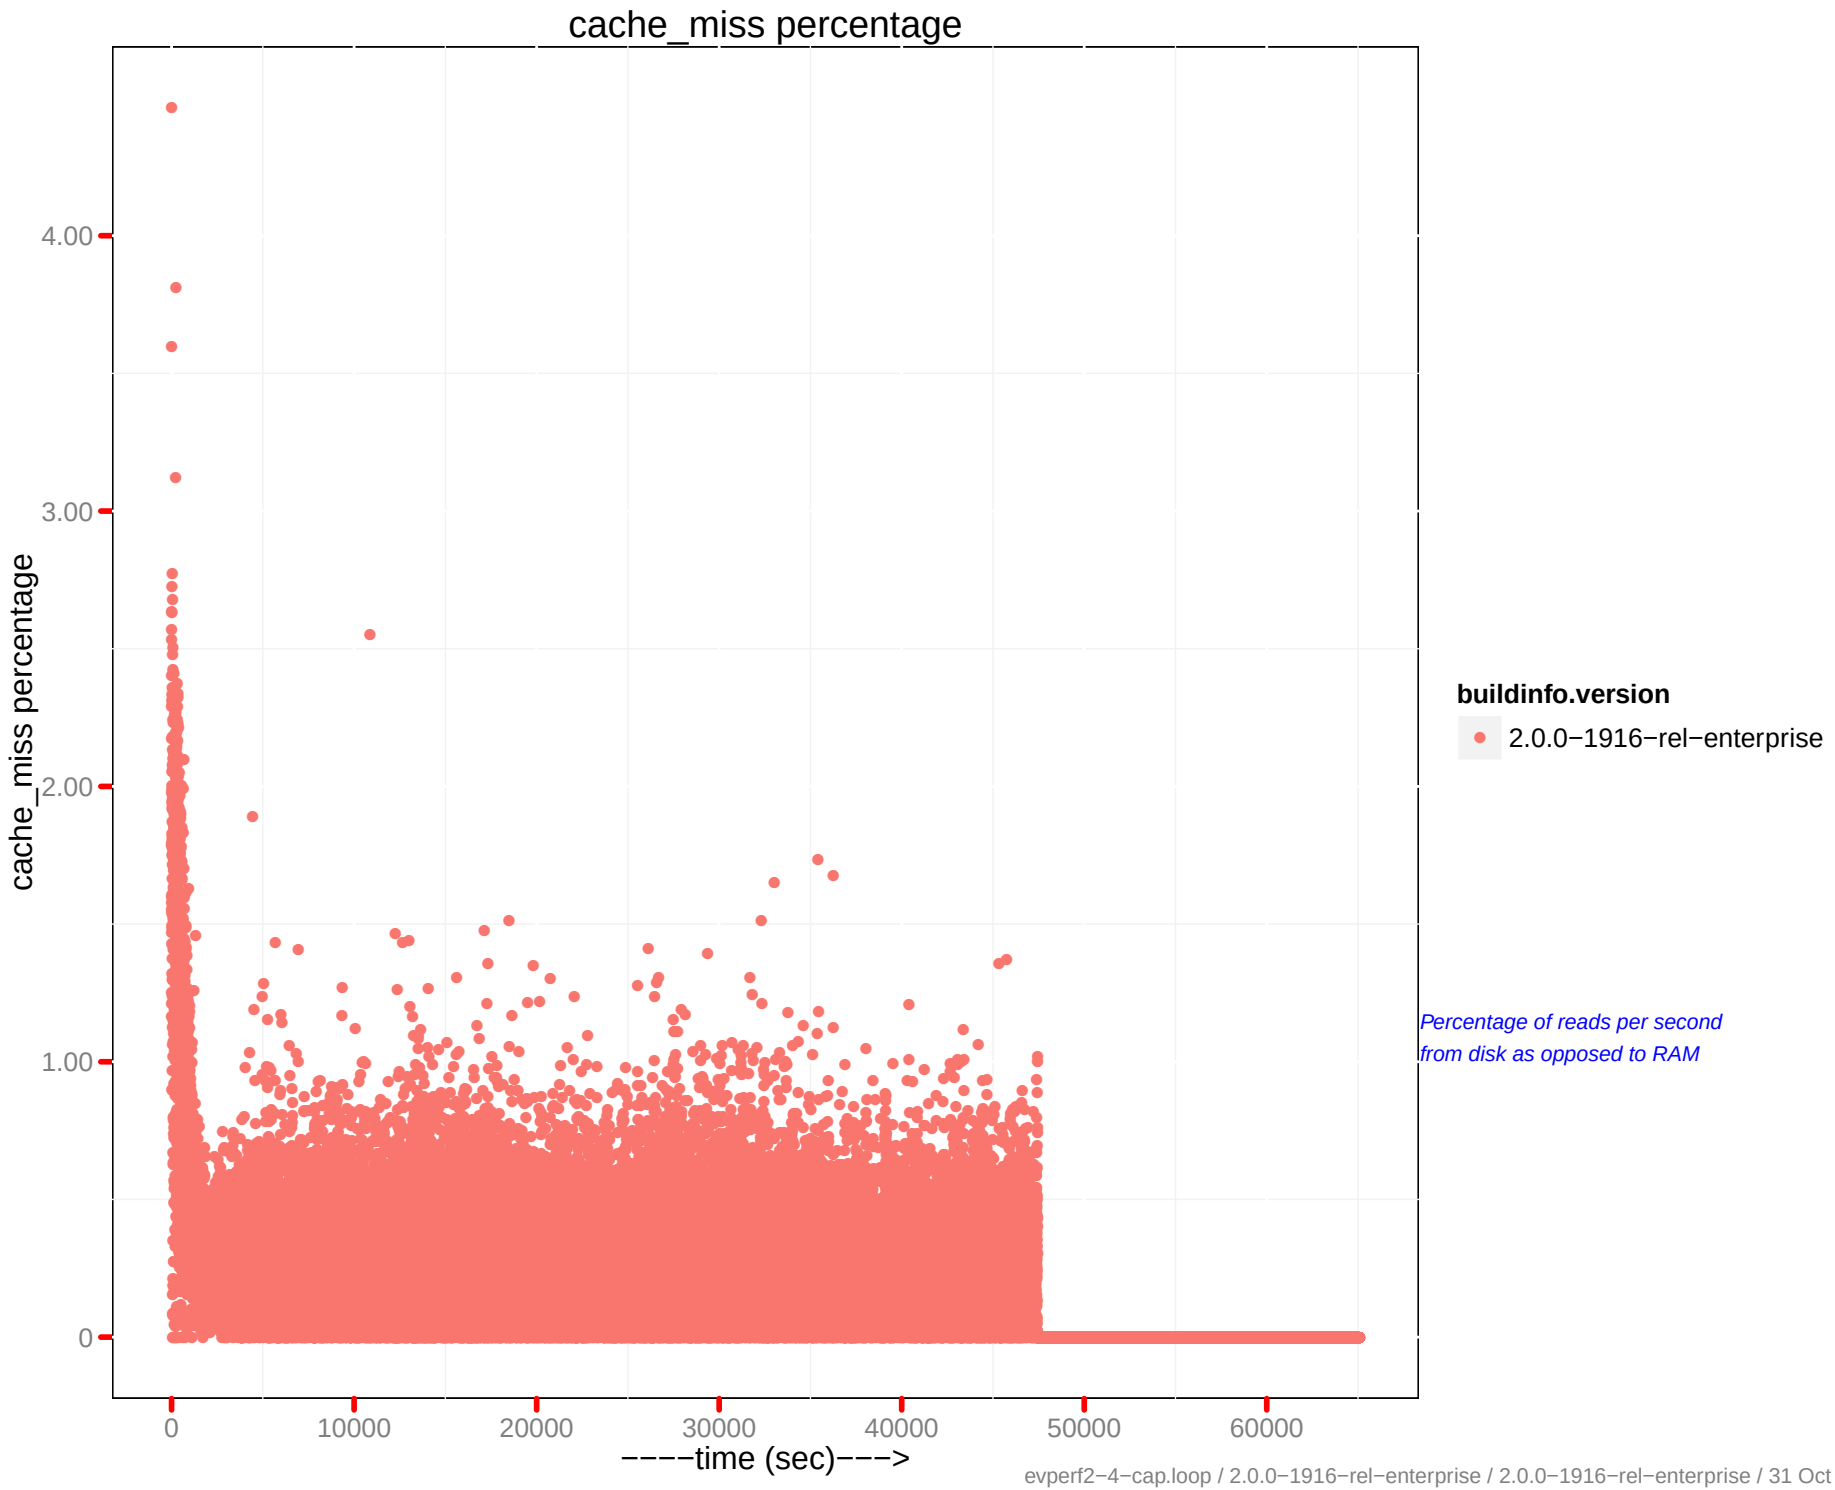
cache_miss percentage
cache_miss percentage
4.00
3.00
2.00
1.00
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
Percentage of reads per second
from disk as opposed to RAM
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

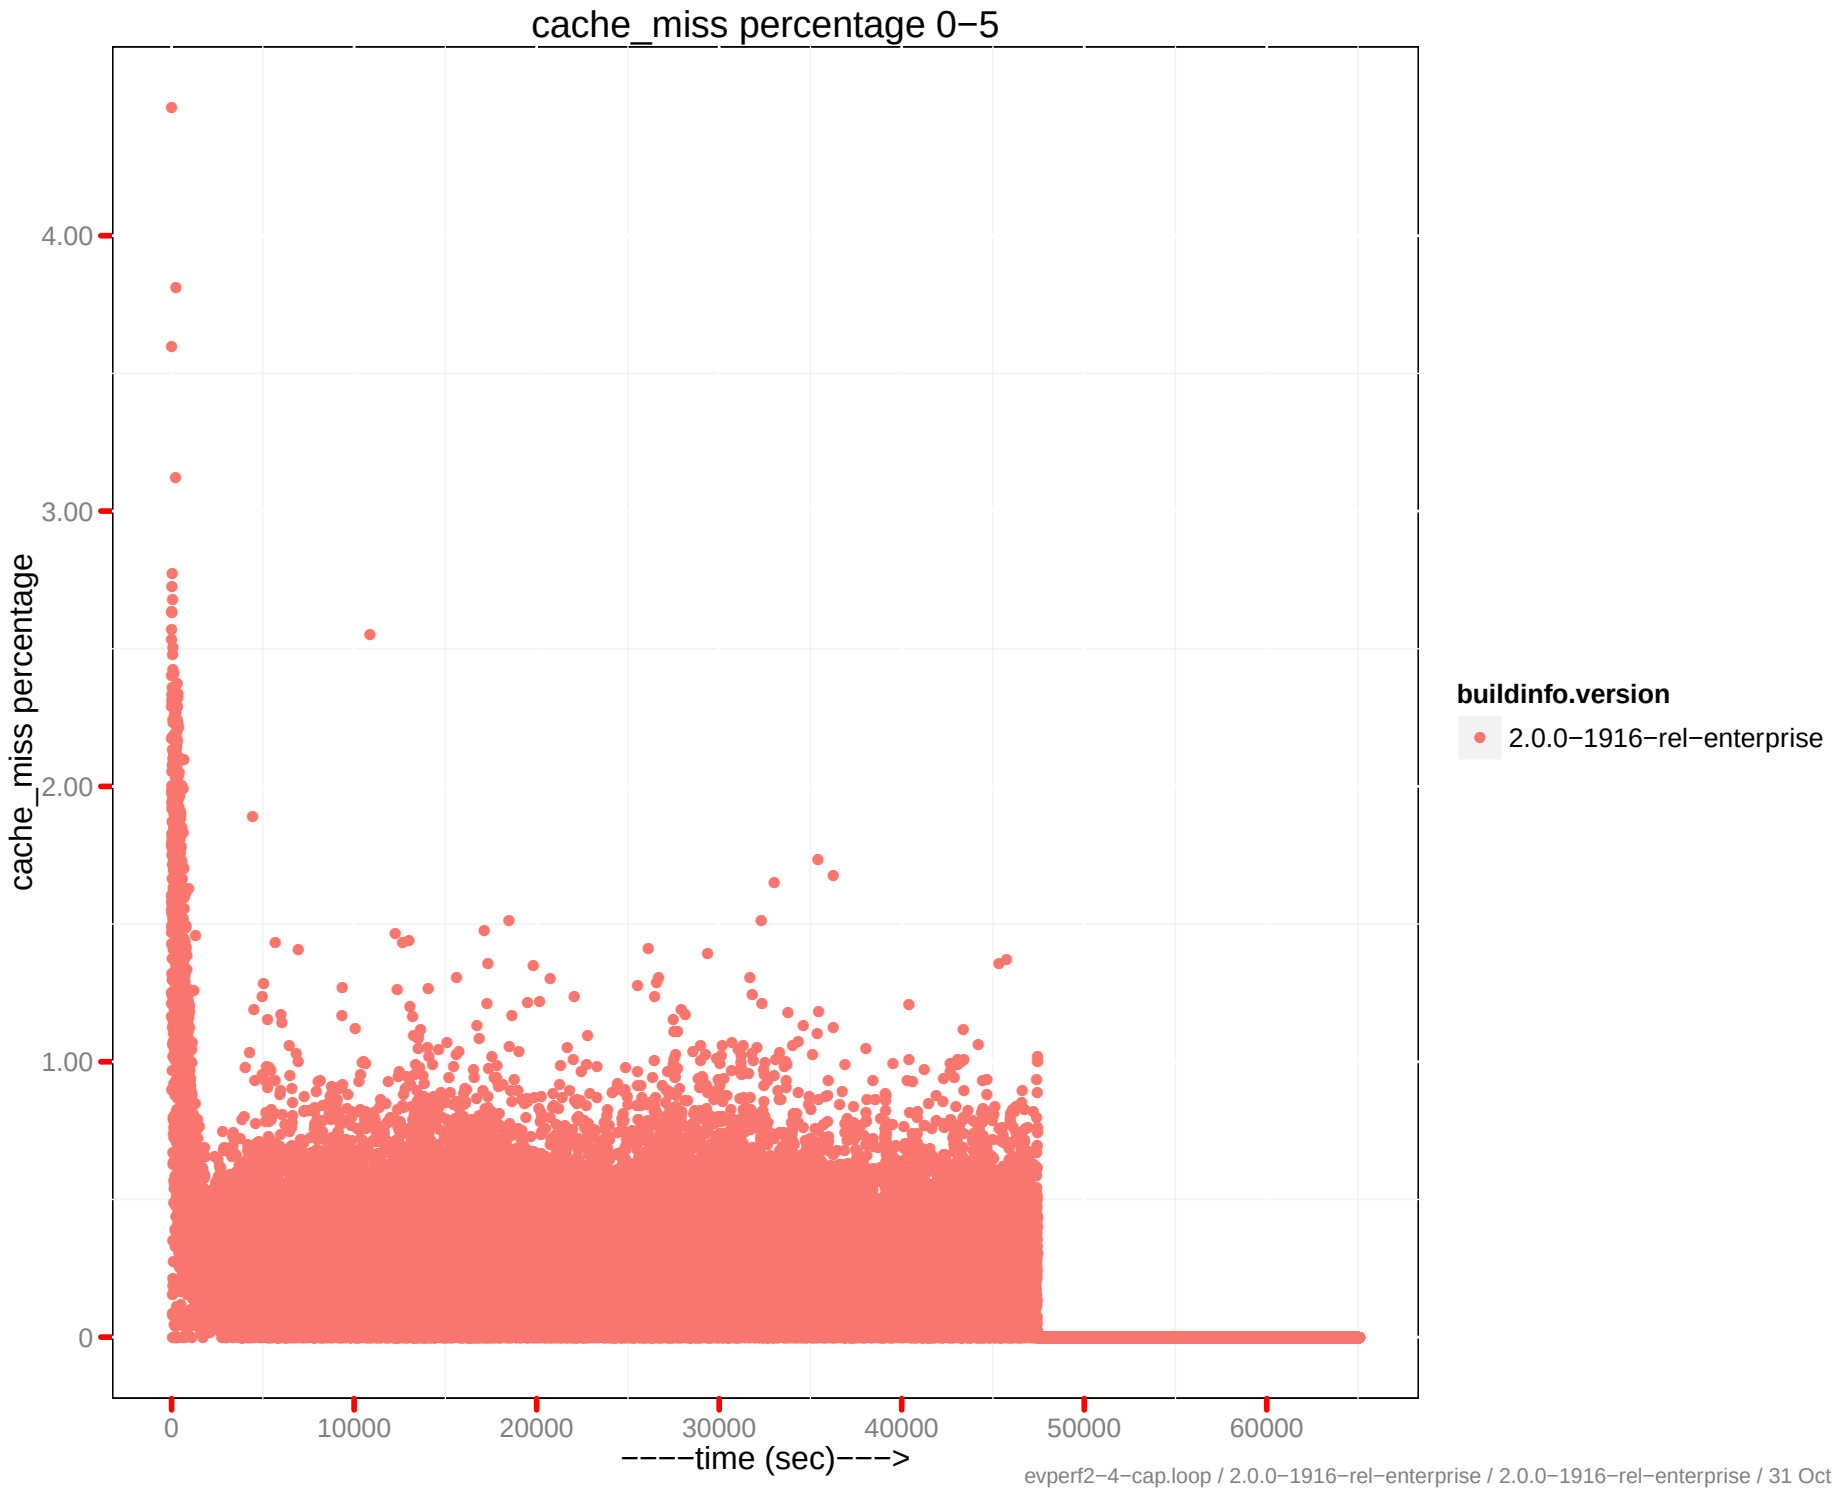
cache_miss percentage 0−5
cache_miss percentage
4.00
3.00
2.00
1.00
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

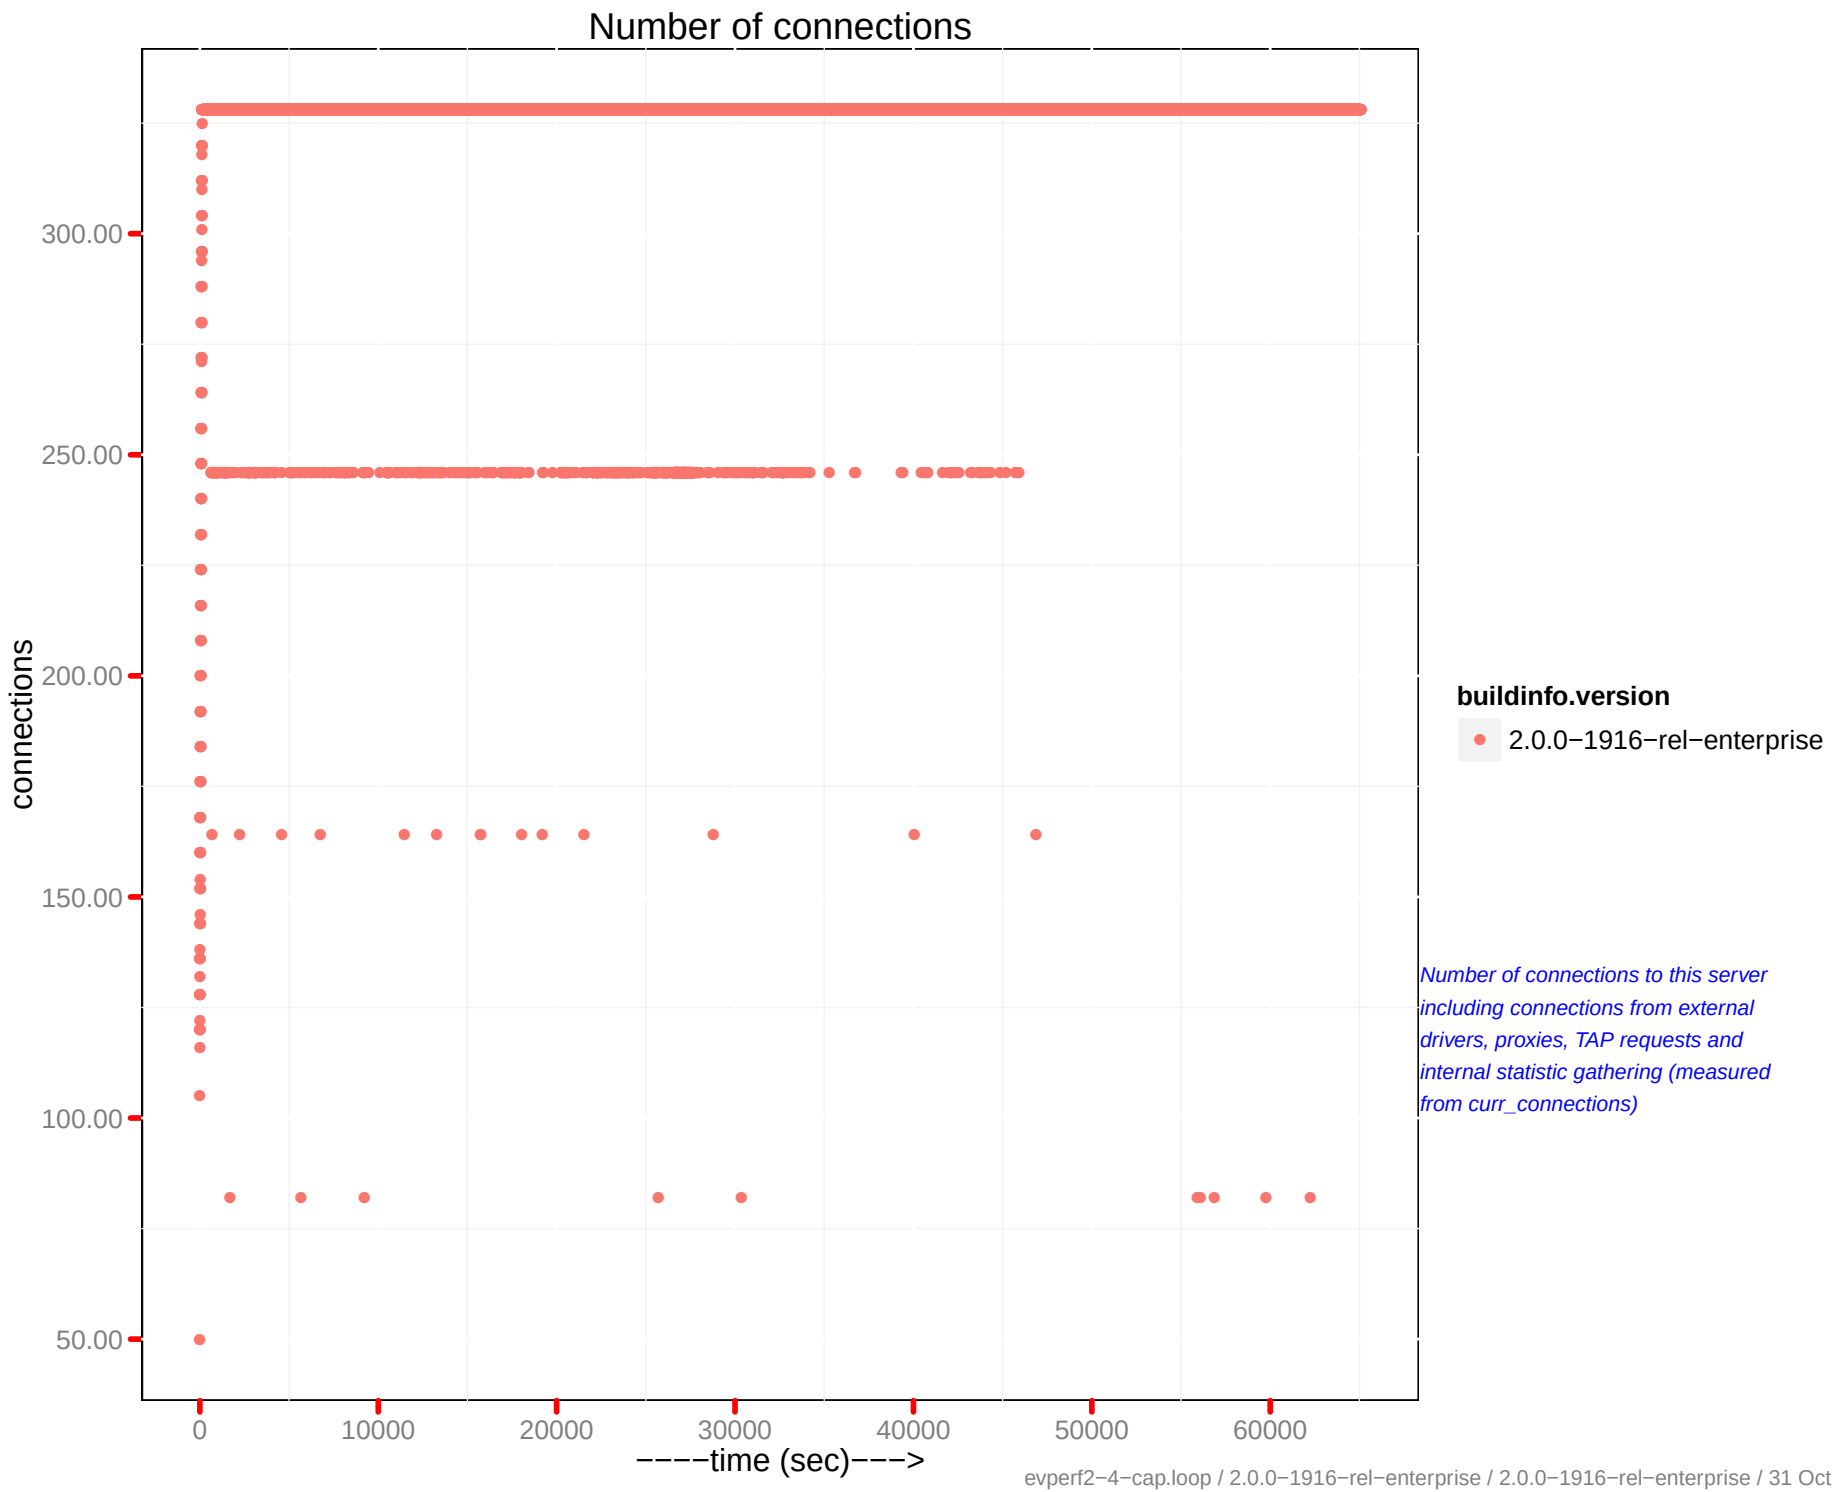
Number of connections
connections
300.00
250.00
200.00
150.00
100.00
50.00
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of connections to this server including connections from external drivers, proxies, TAP requests and internal statistic gathering (measured from curr_connections)
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

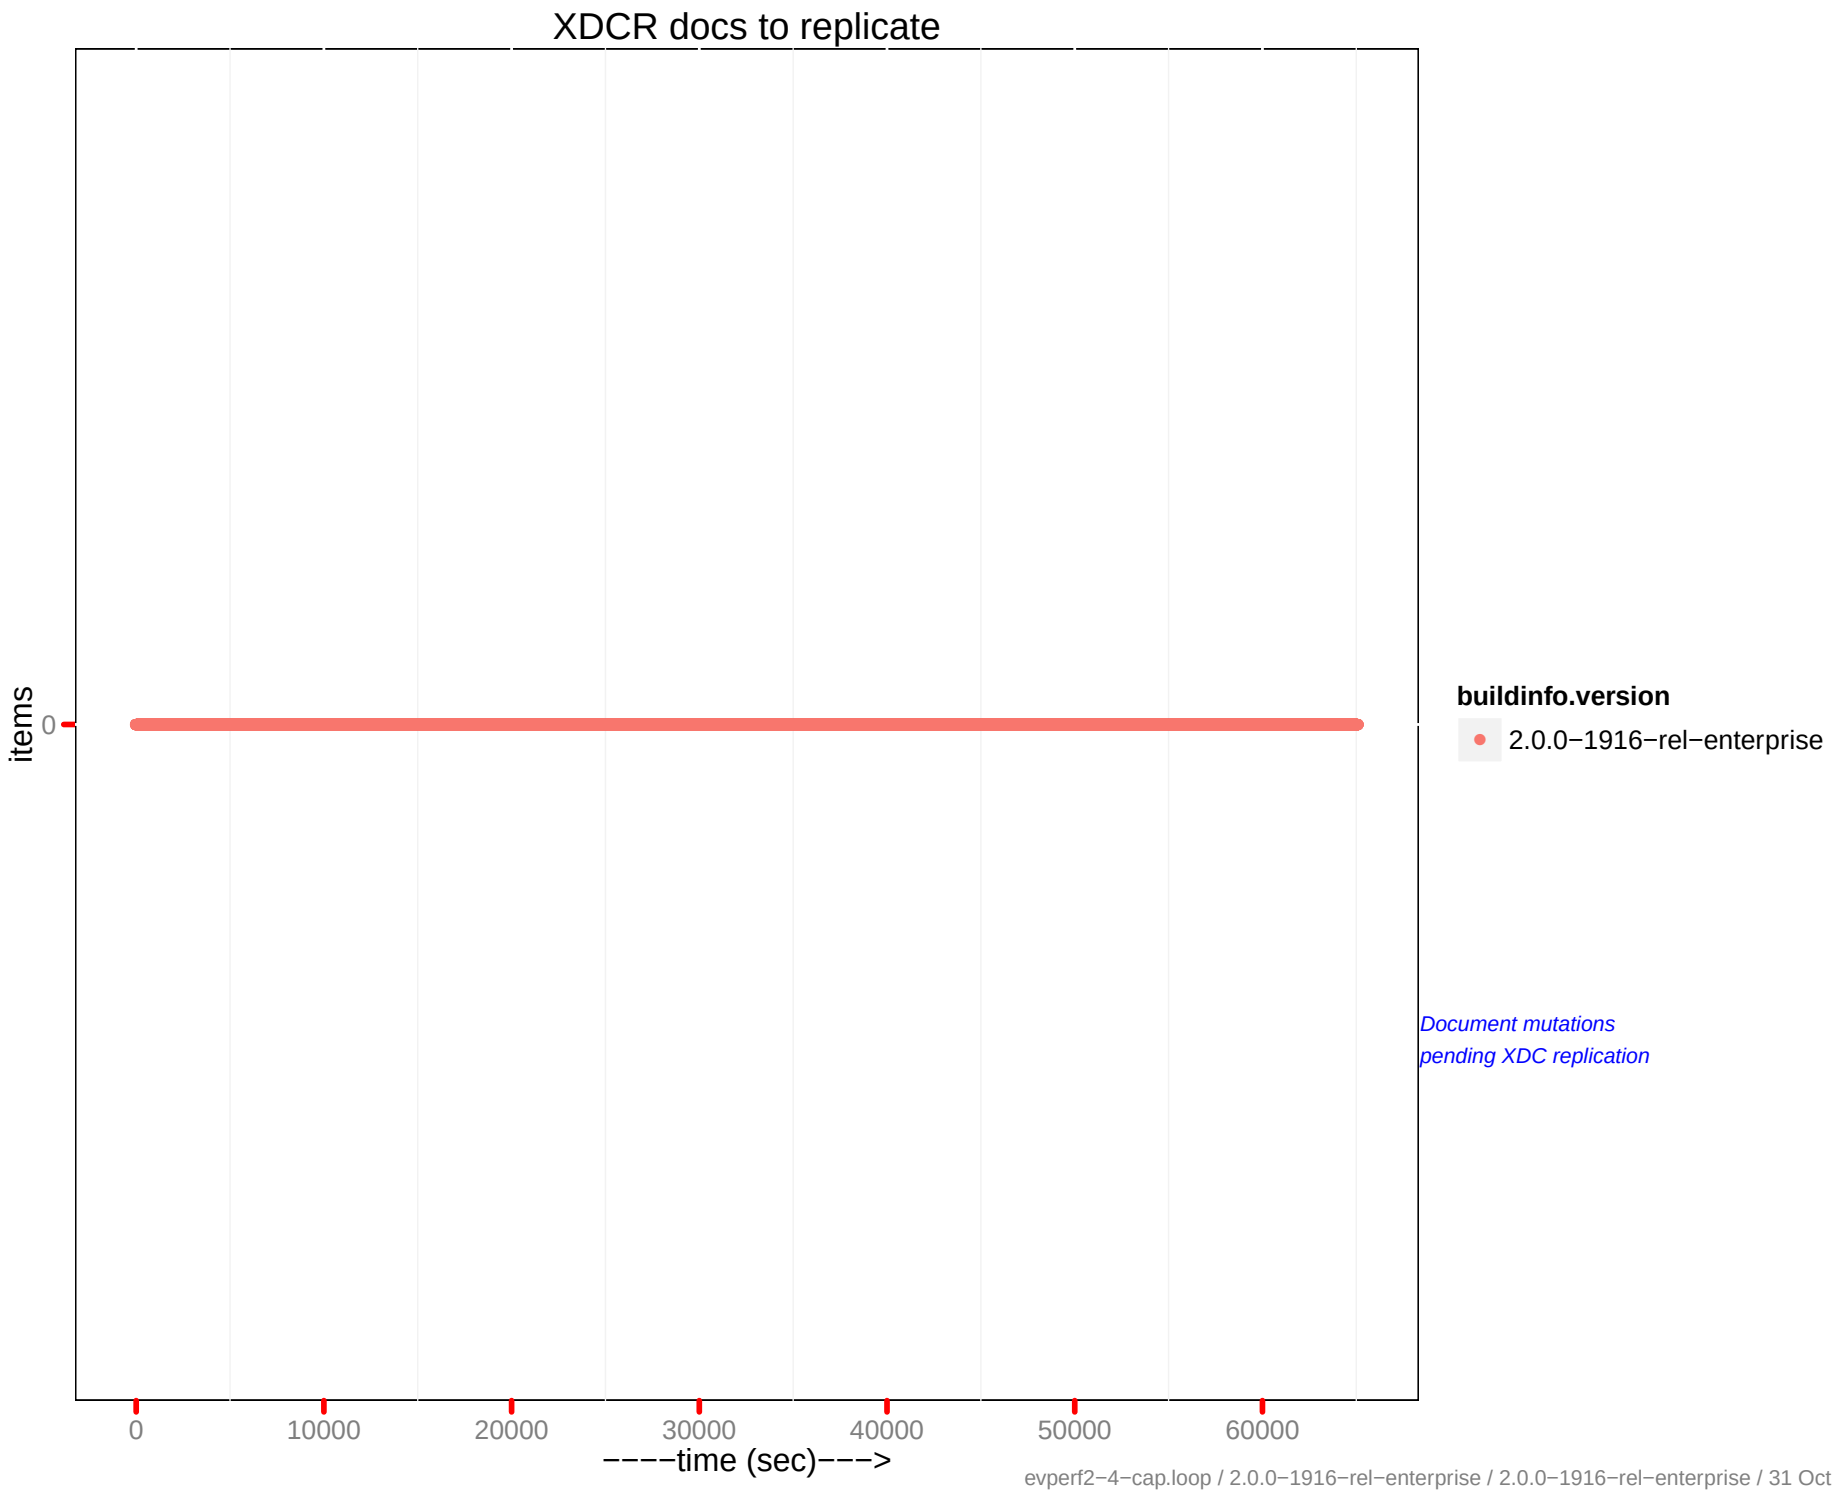
XDCR docs to replicate
items
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Document mutations
pending XDC replication
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

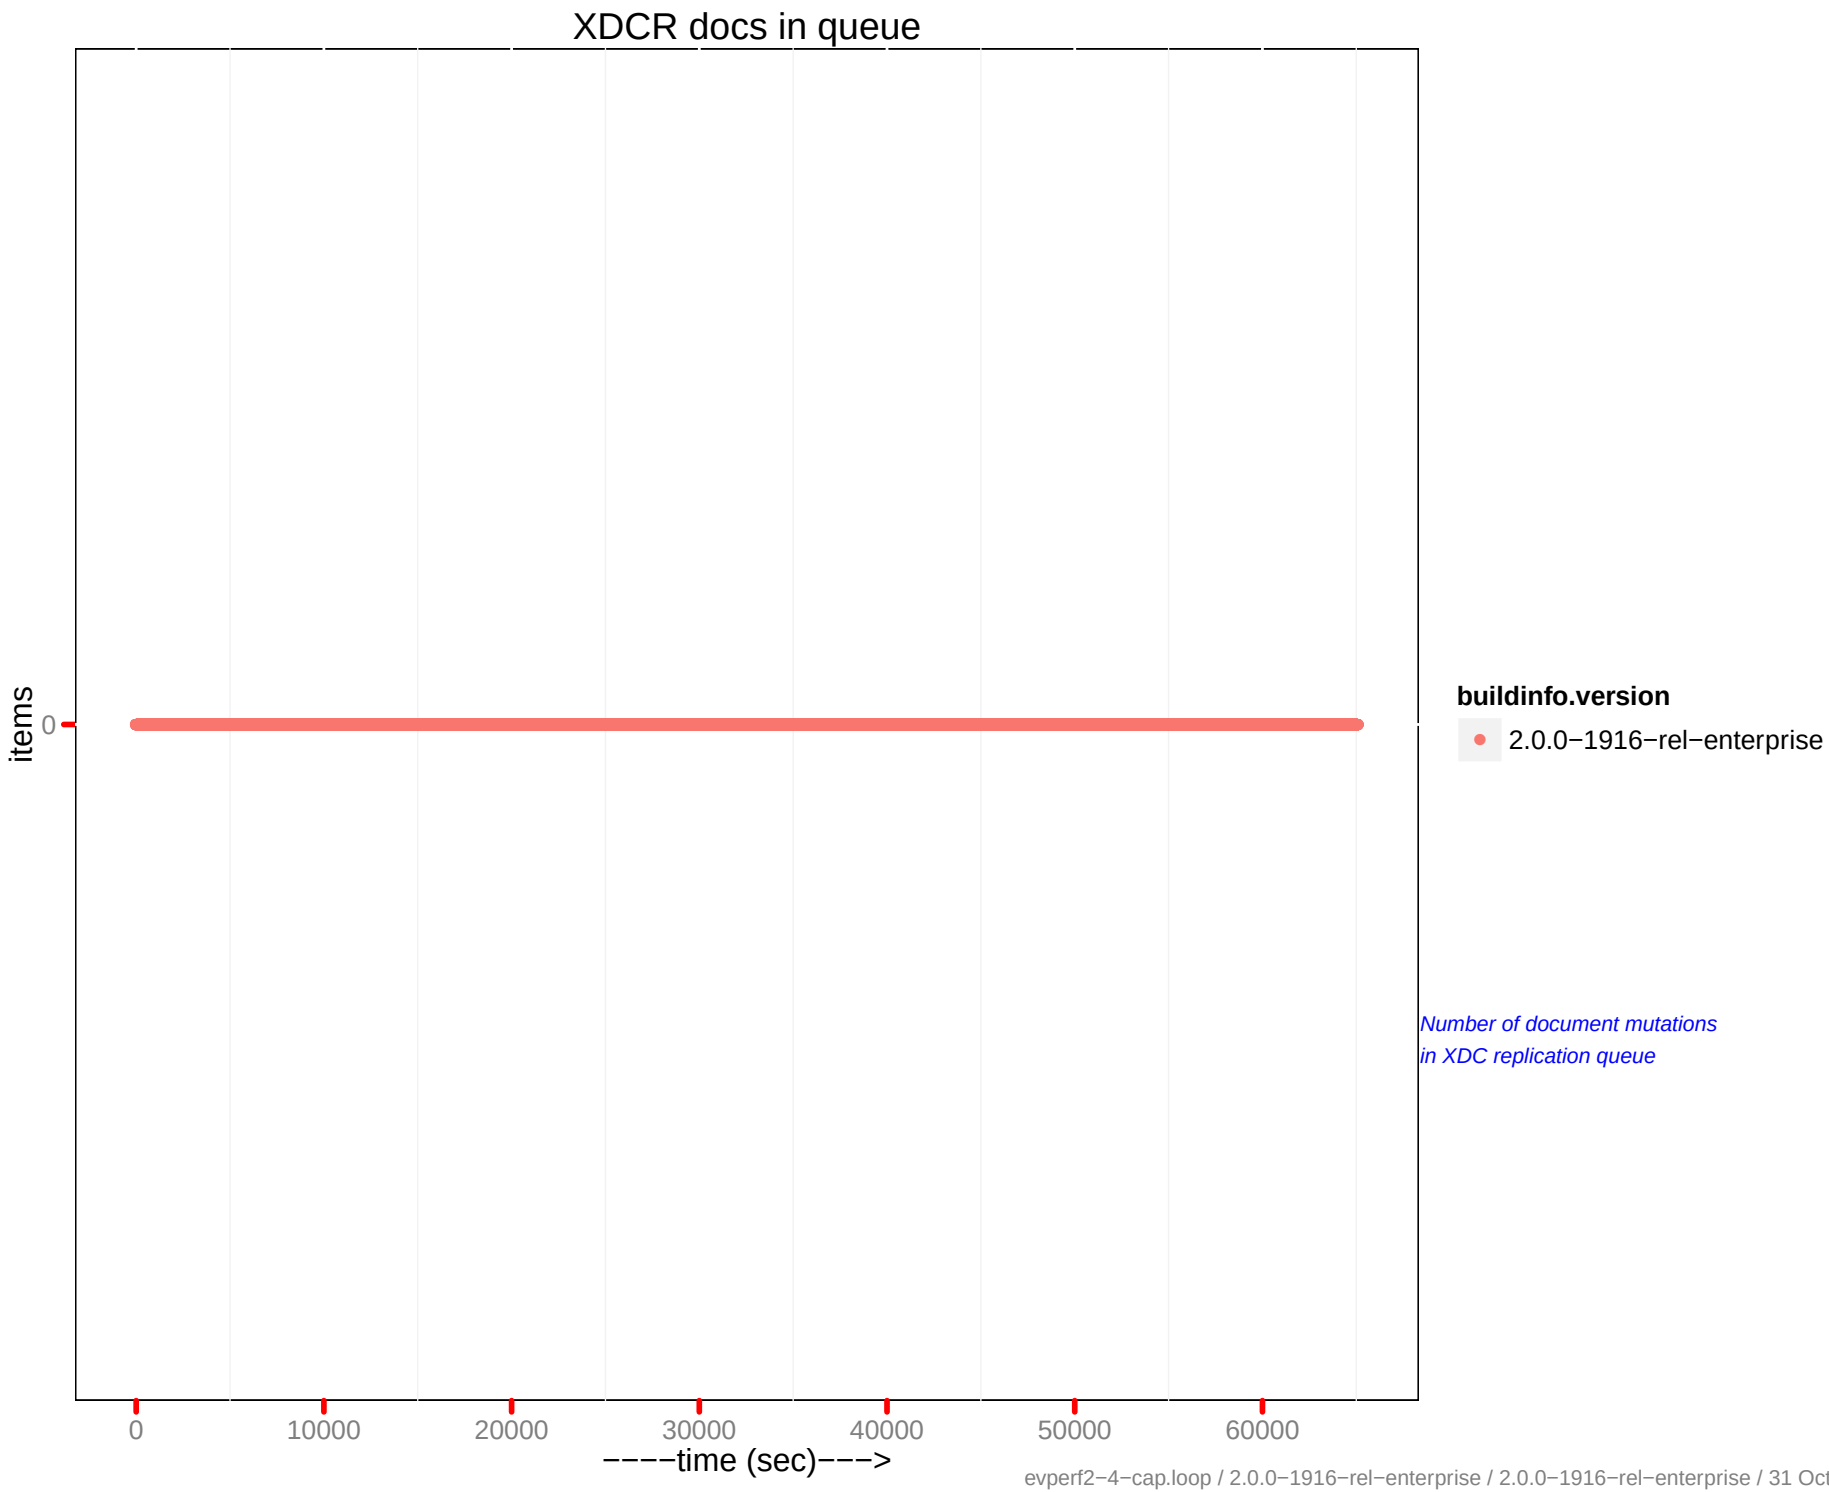
XDCR docs in queue
items
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of document mutations
in XDC replication queue
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

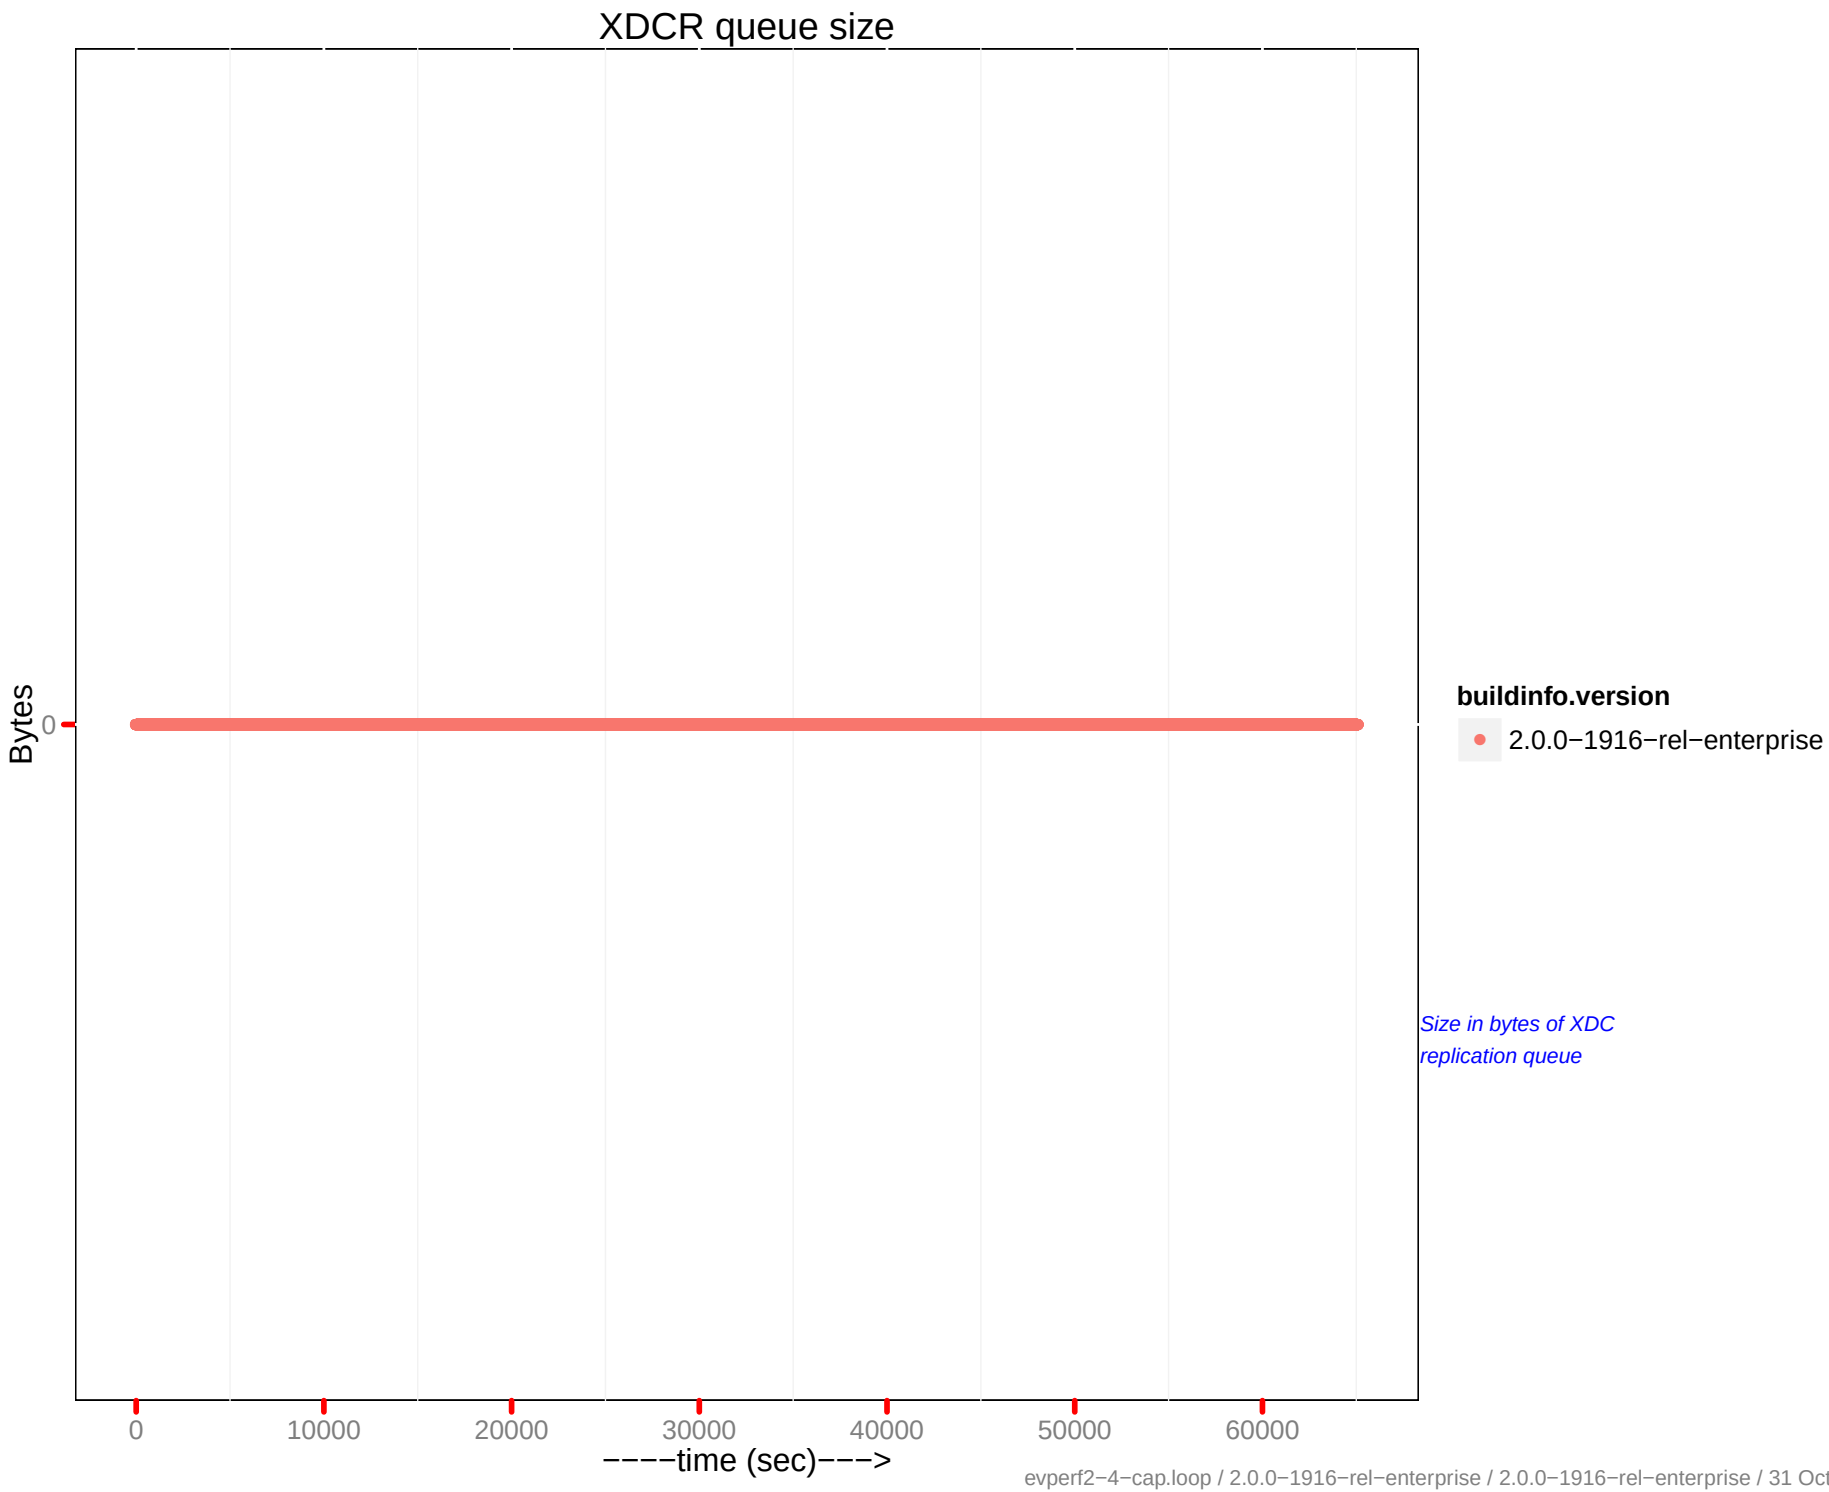
XDCR queue size
Bytes
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Size in bytes of XDC replication queue
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

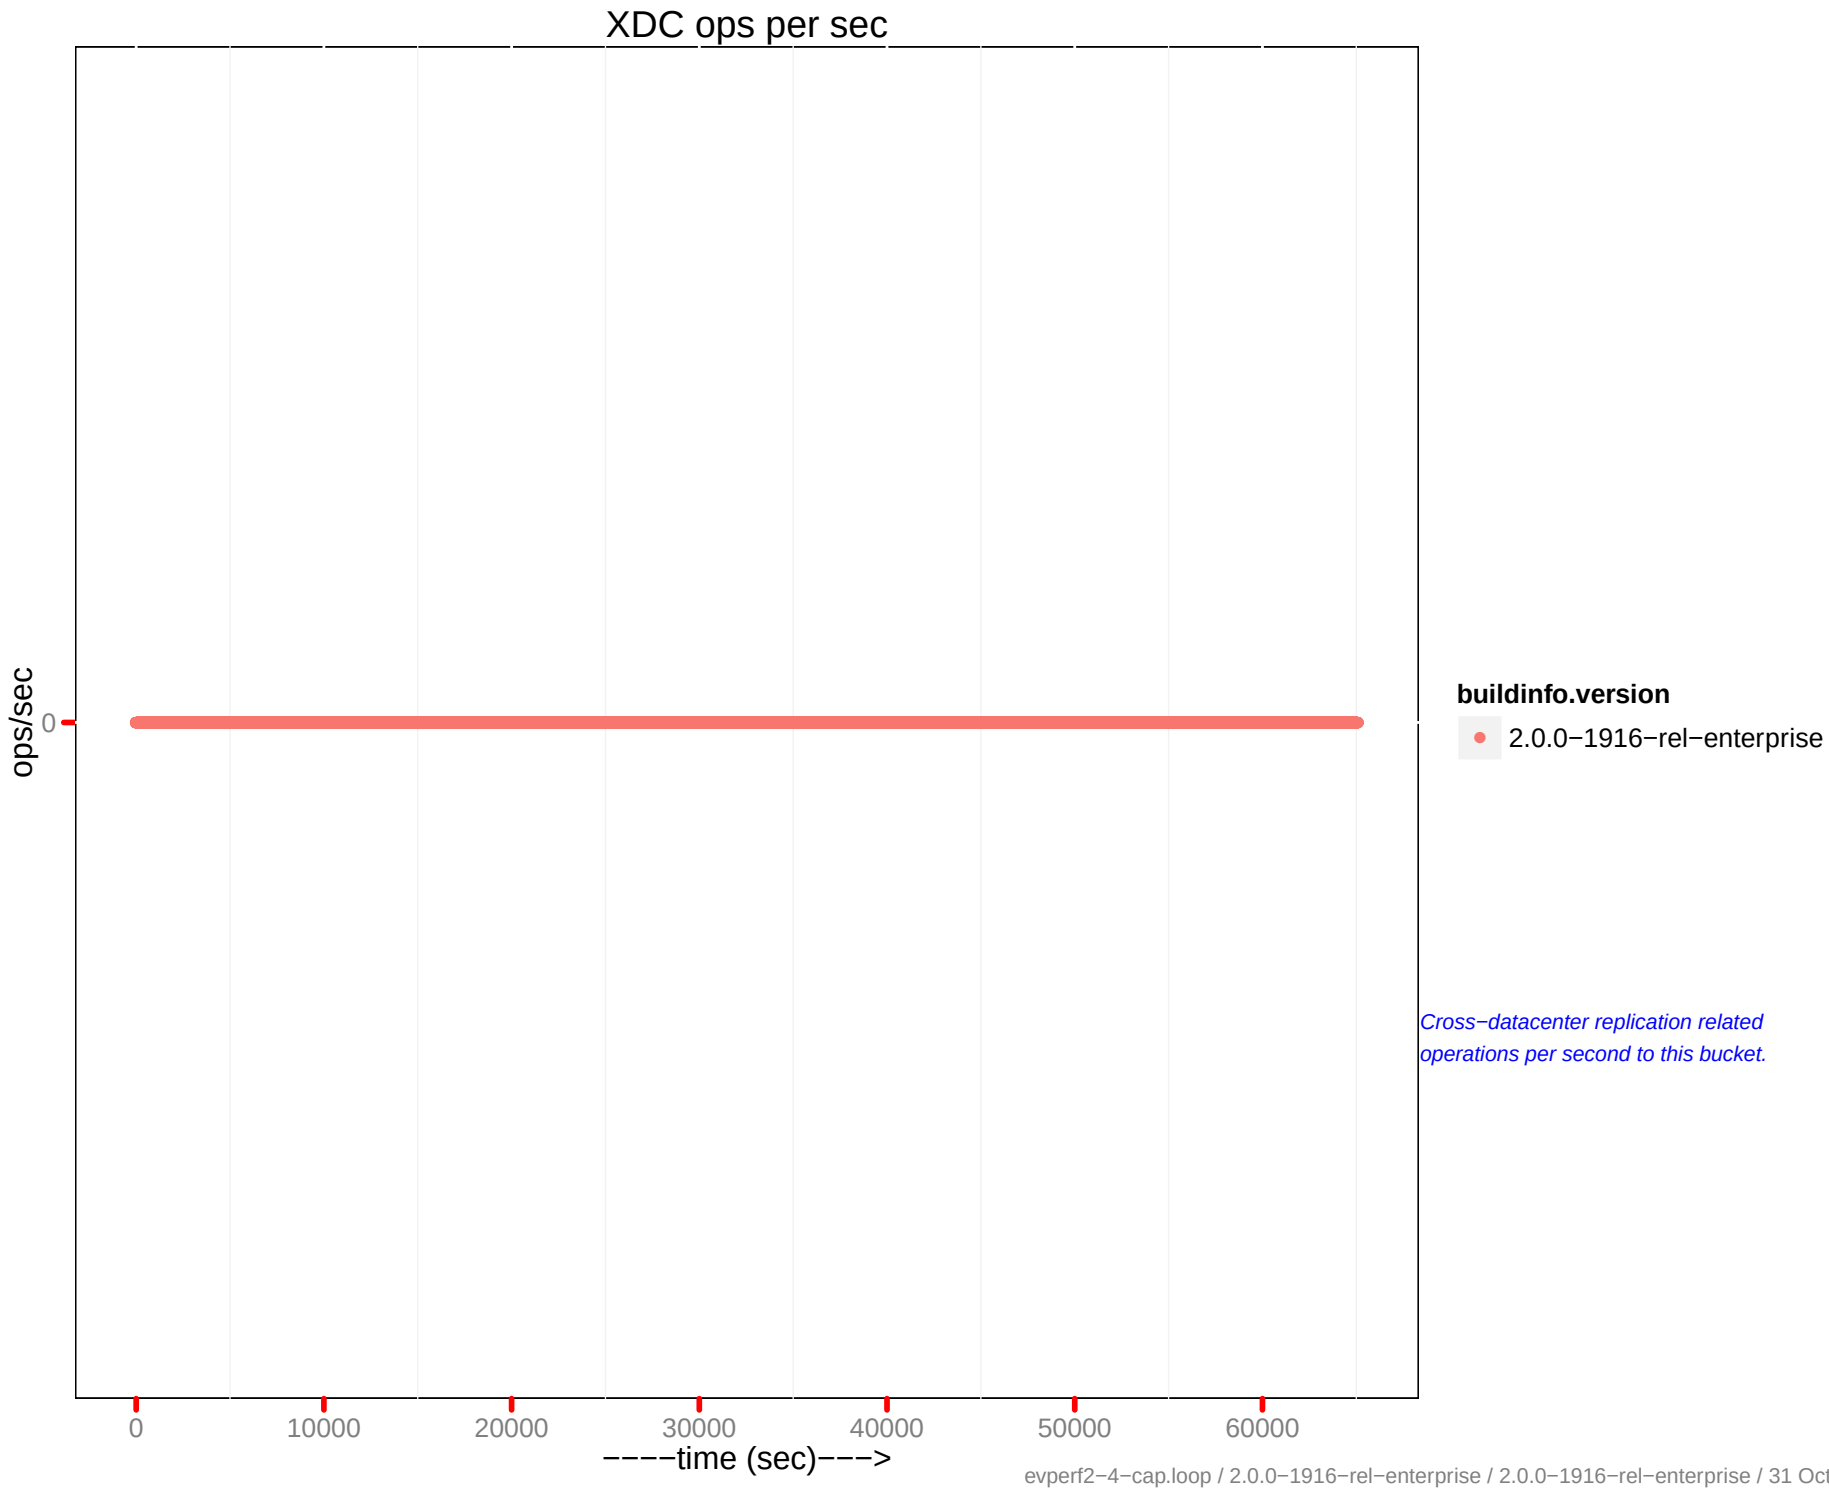
XDC ops per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Cross-datacenter replication related
operations per second to this bucket.
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

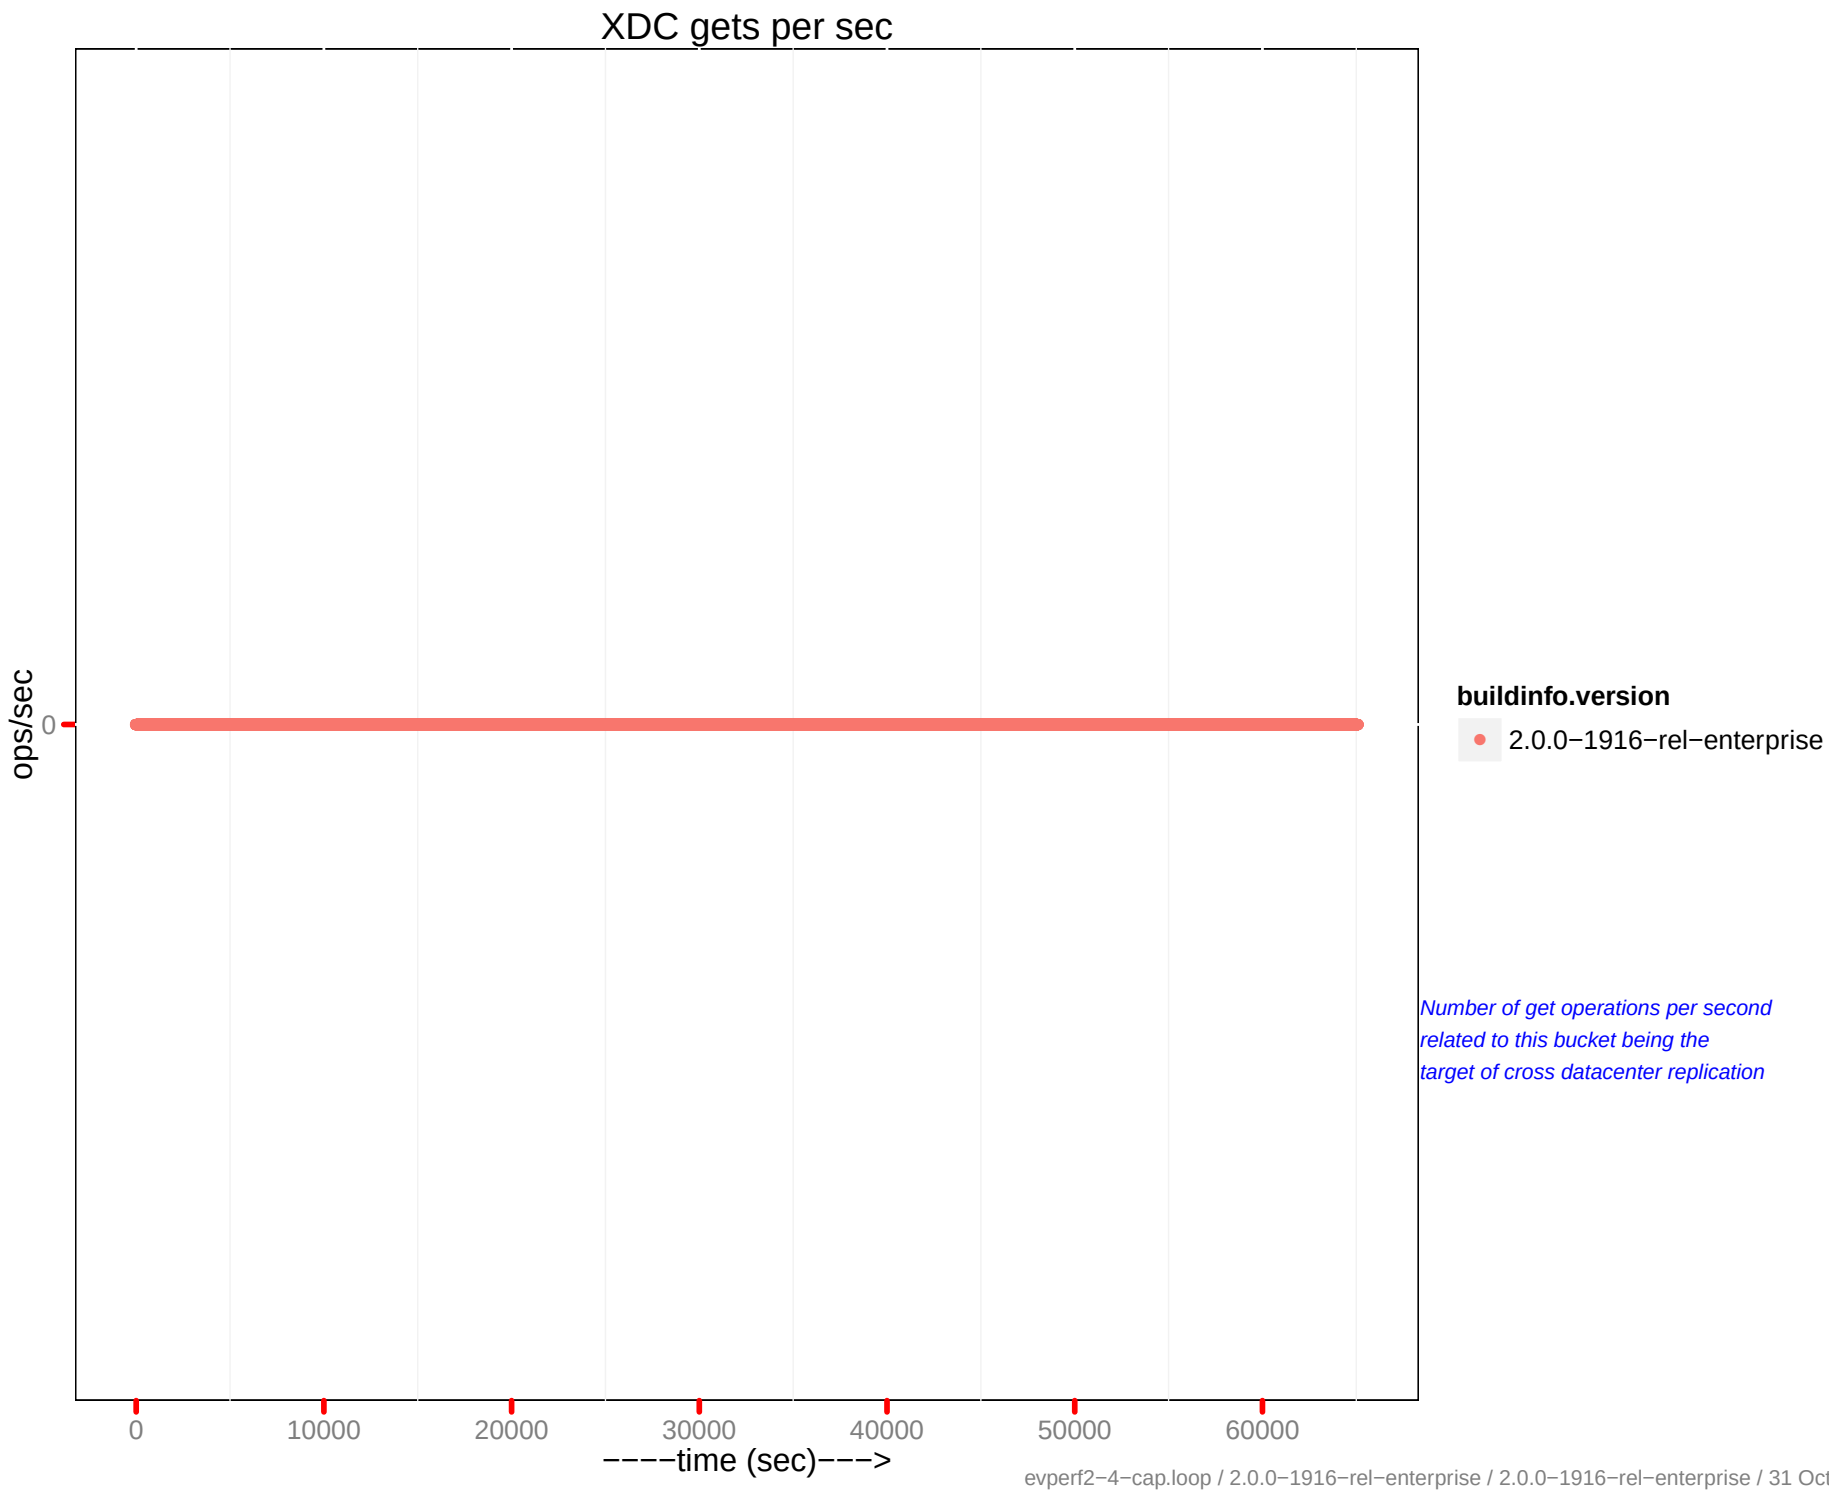
XDC gets per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of get operations per second related to this bucket being the target of cross datacenter replication
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

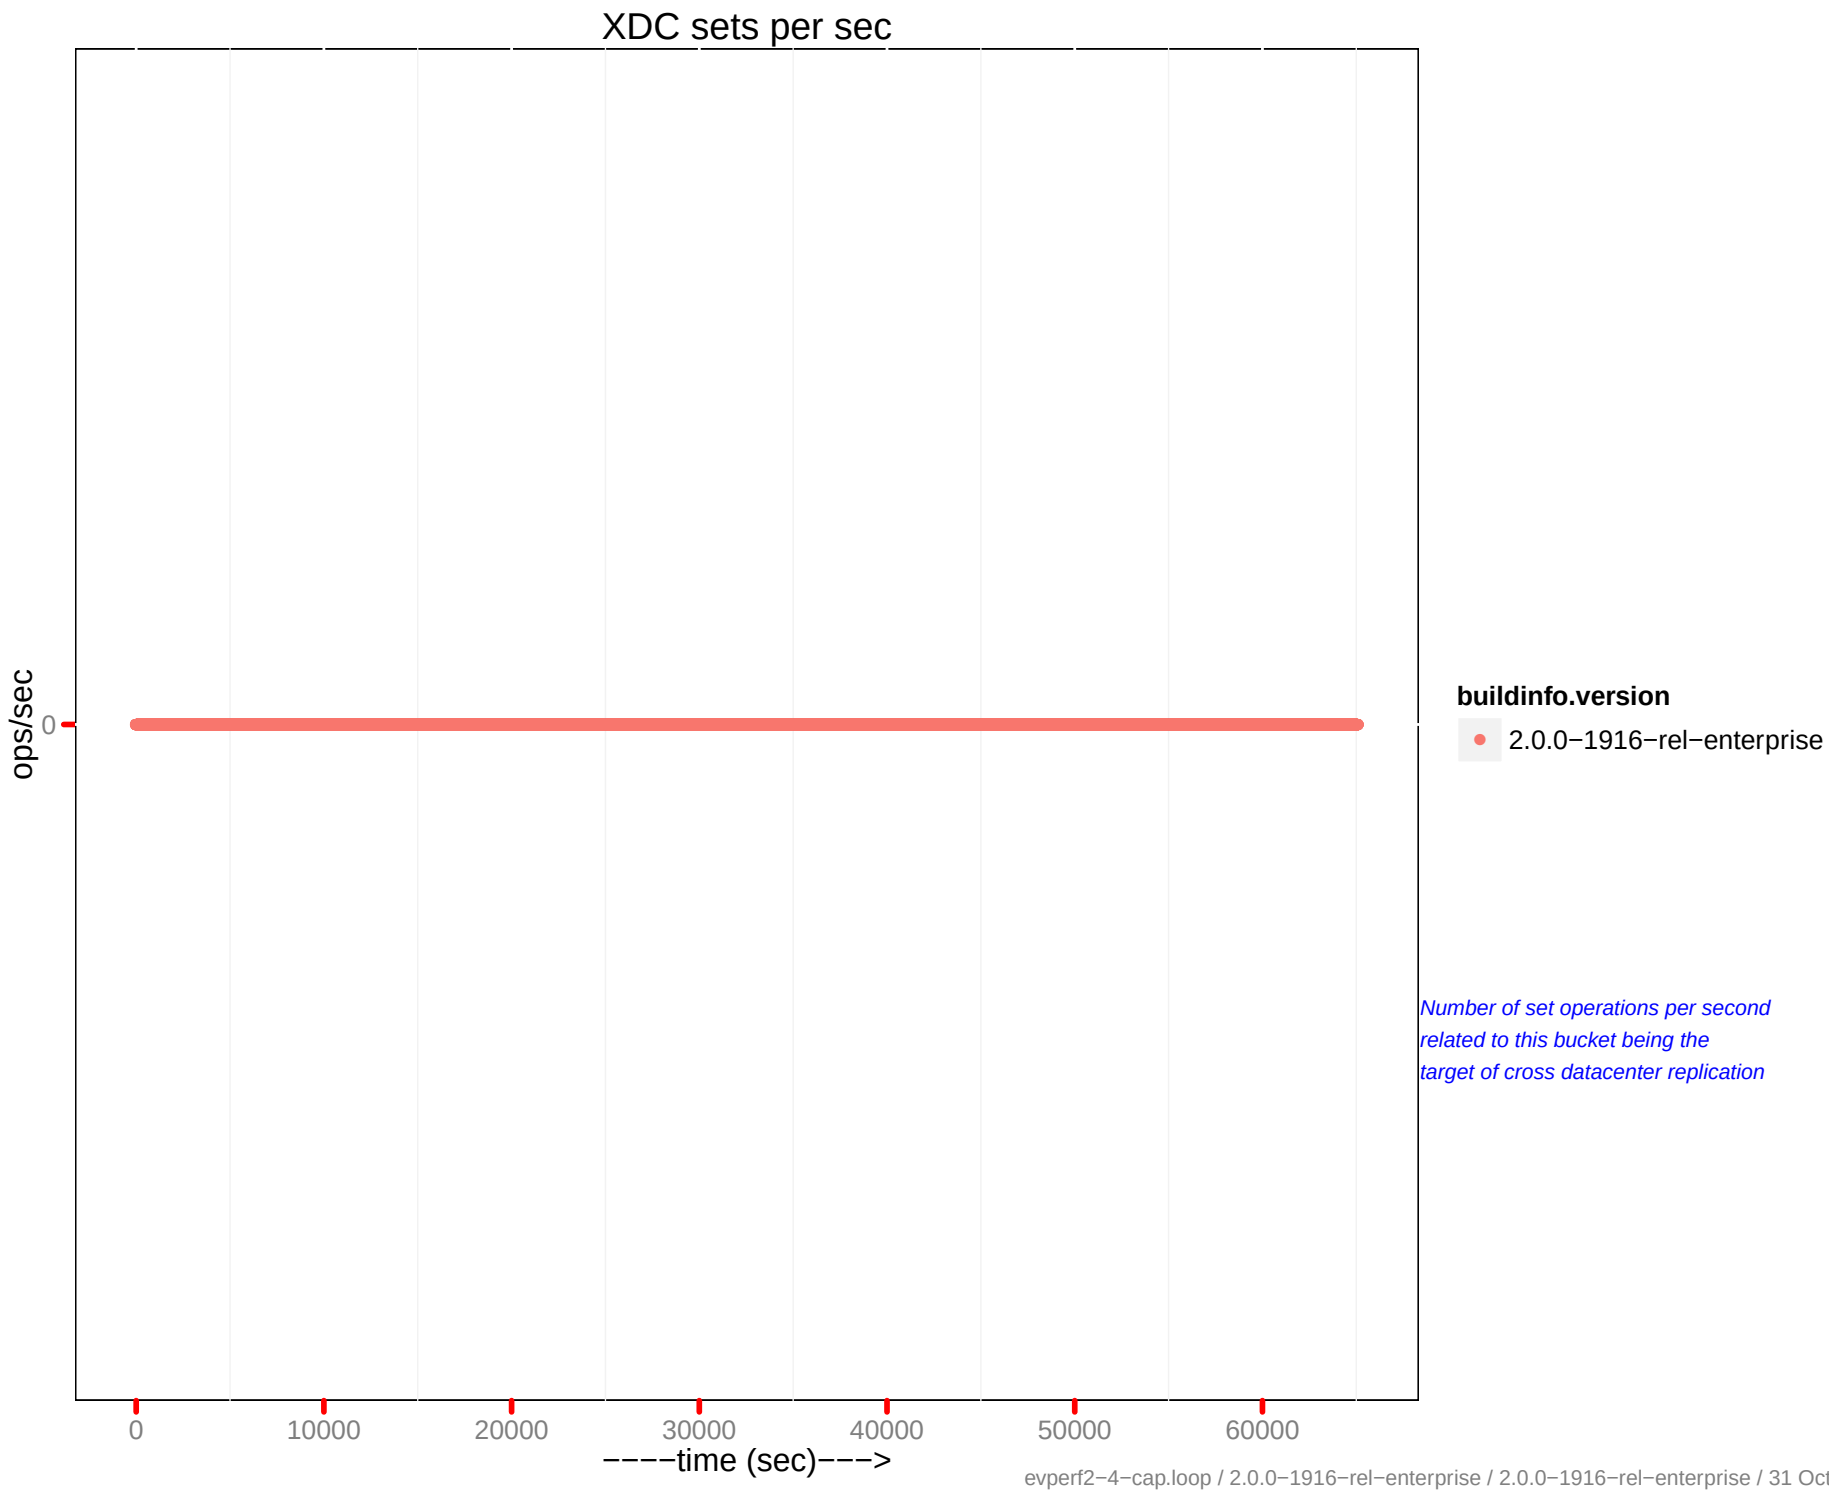
XDC sets per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of set operations per second
related to this bucket being the
target of cross datacenter replication
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

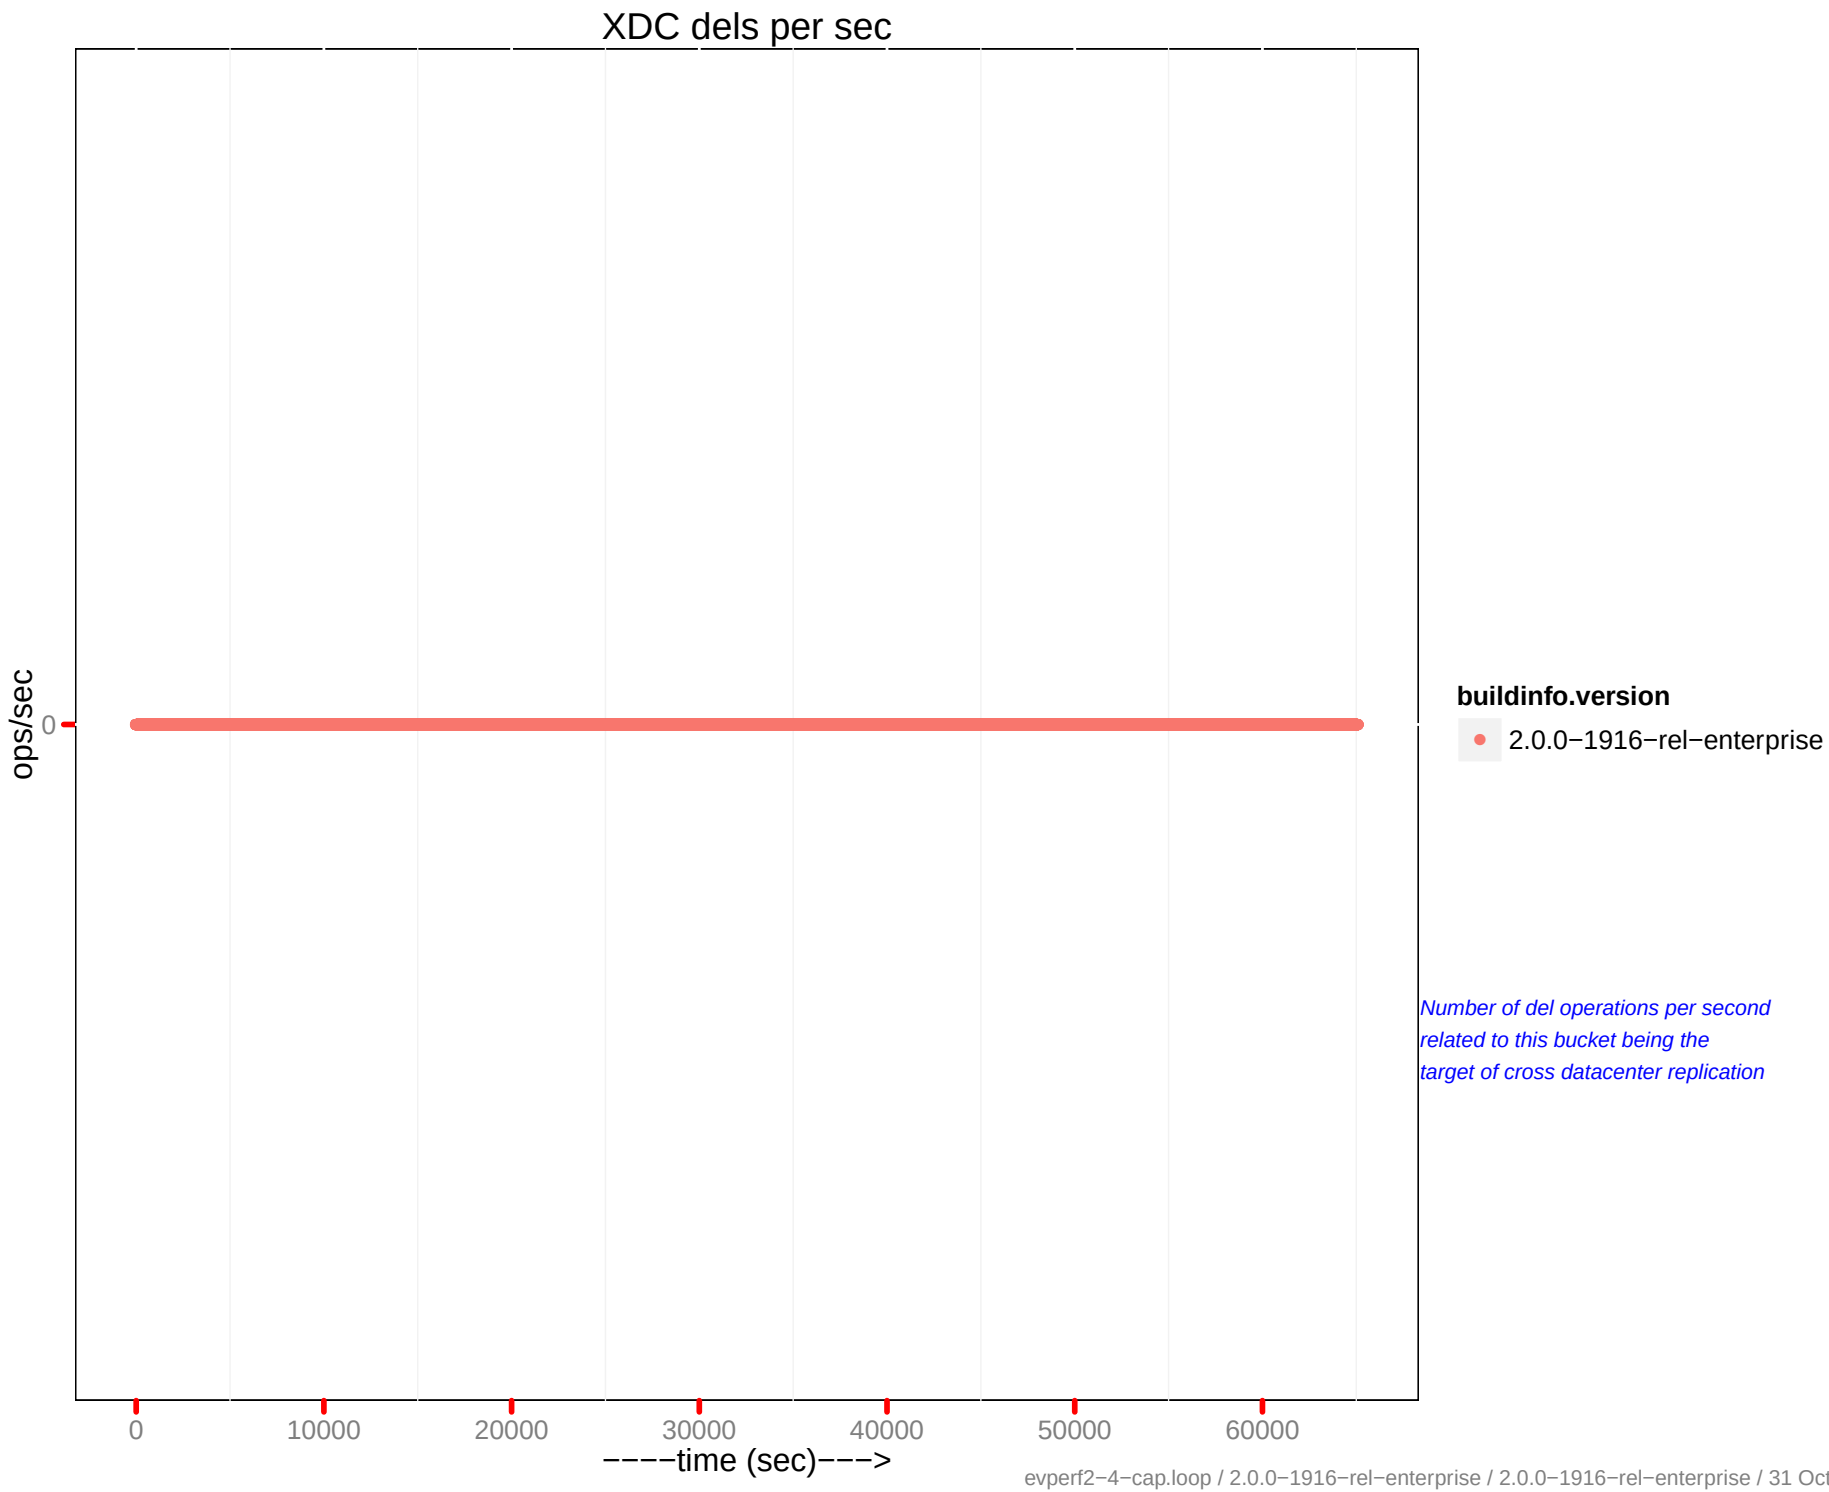
XDC dels per sec
ops/sec
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of del operations per second
related to this bucket being the
target of cross datacenter replication
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

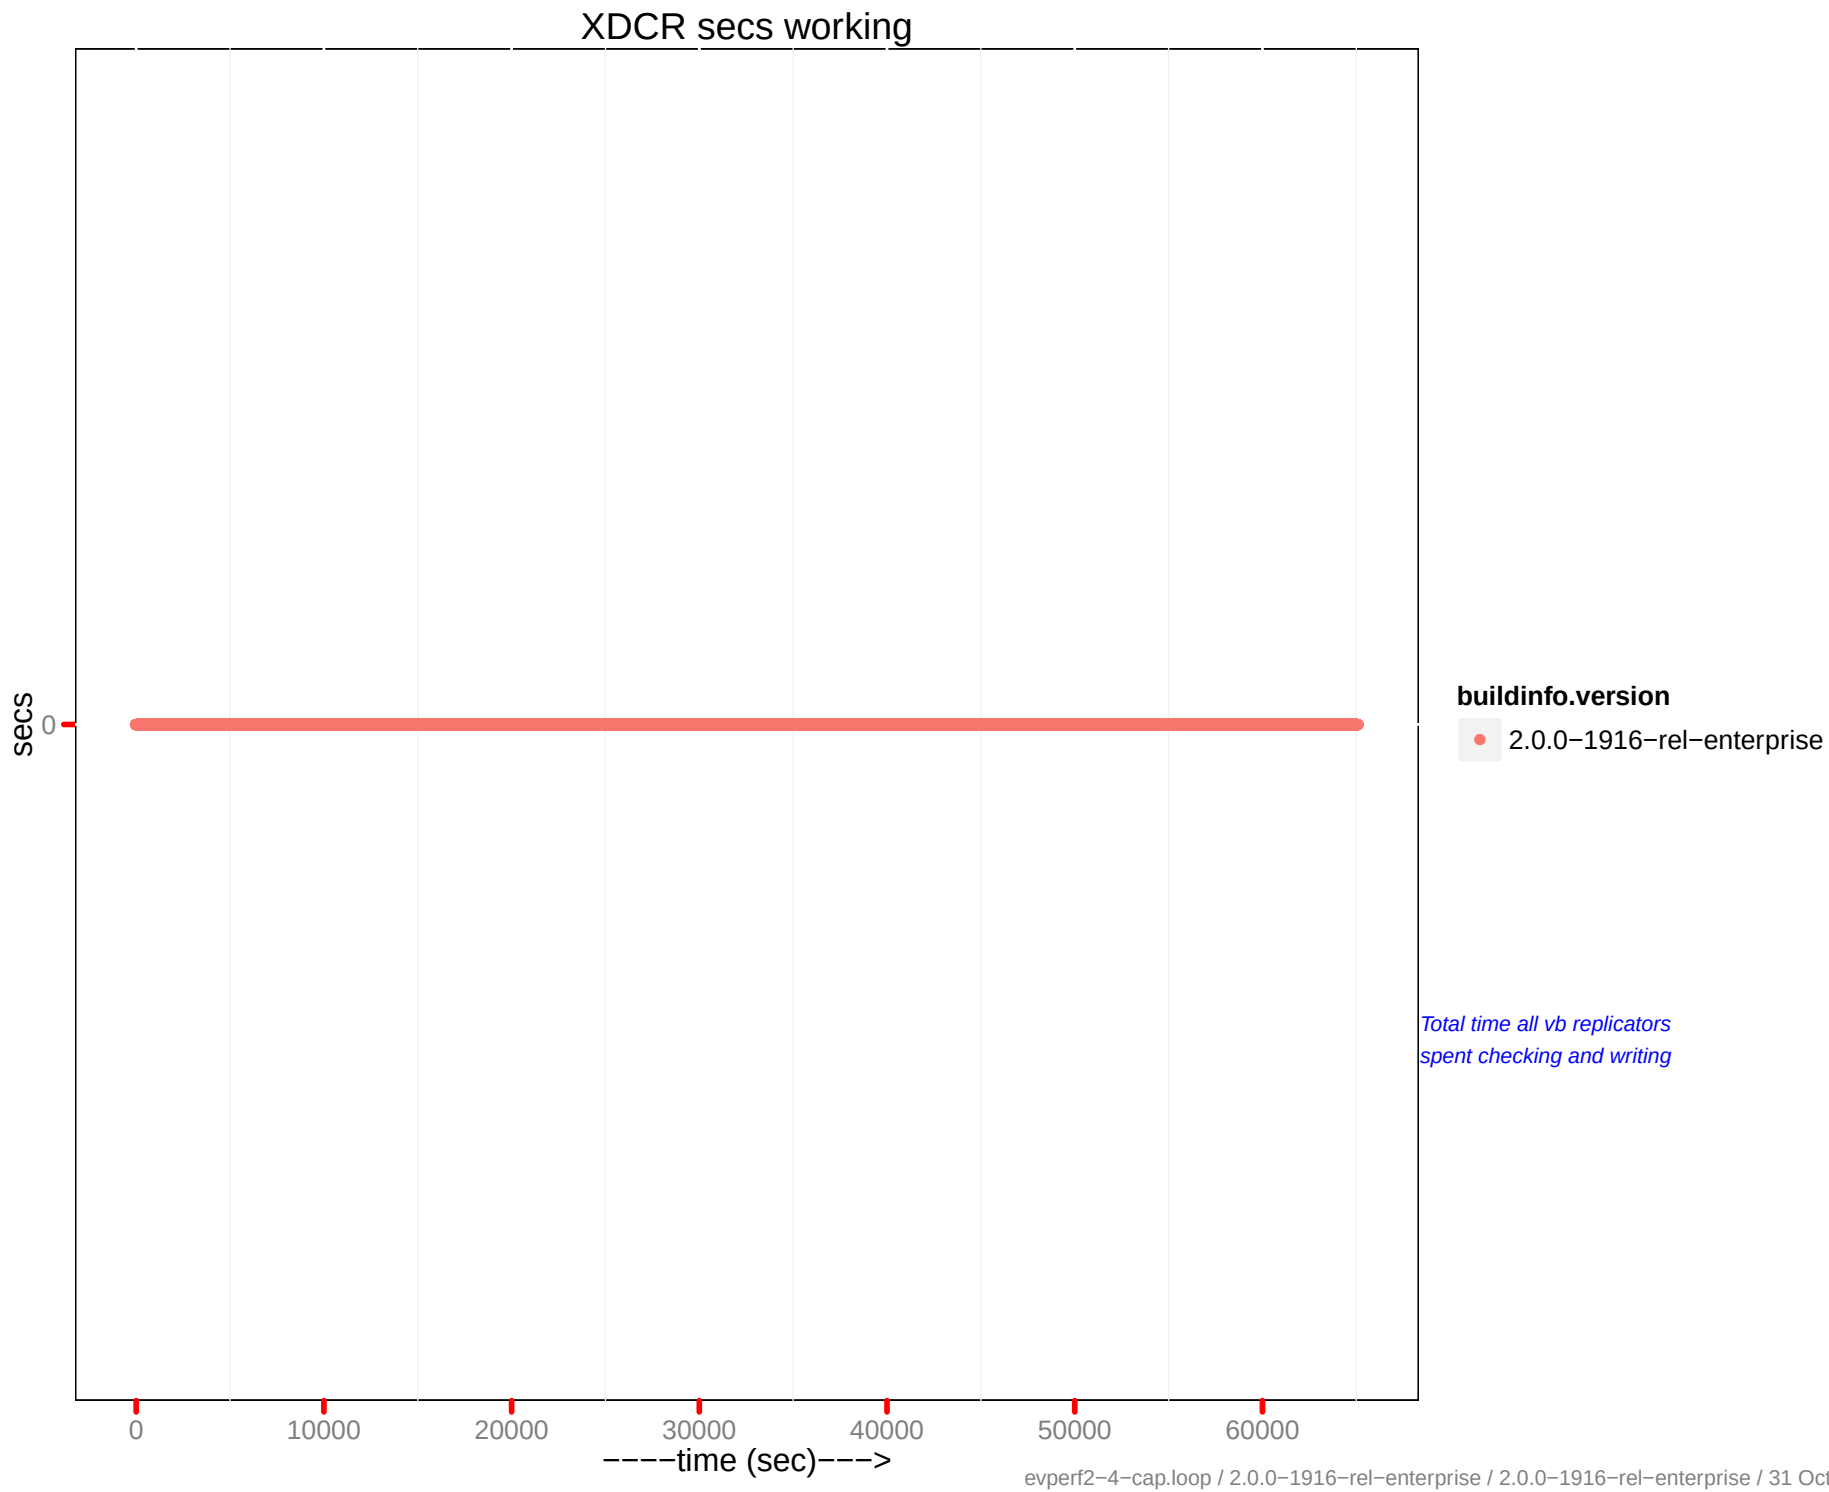
XDCR secs working
secs
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Total time all vb replicators
spent checking and writing
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

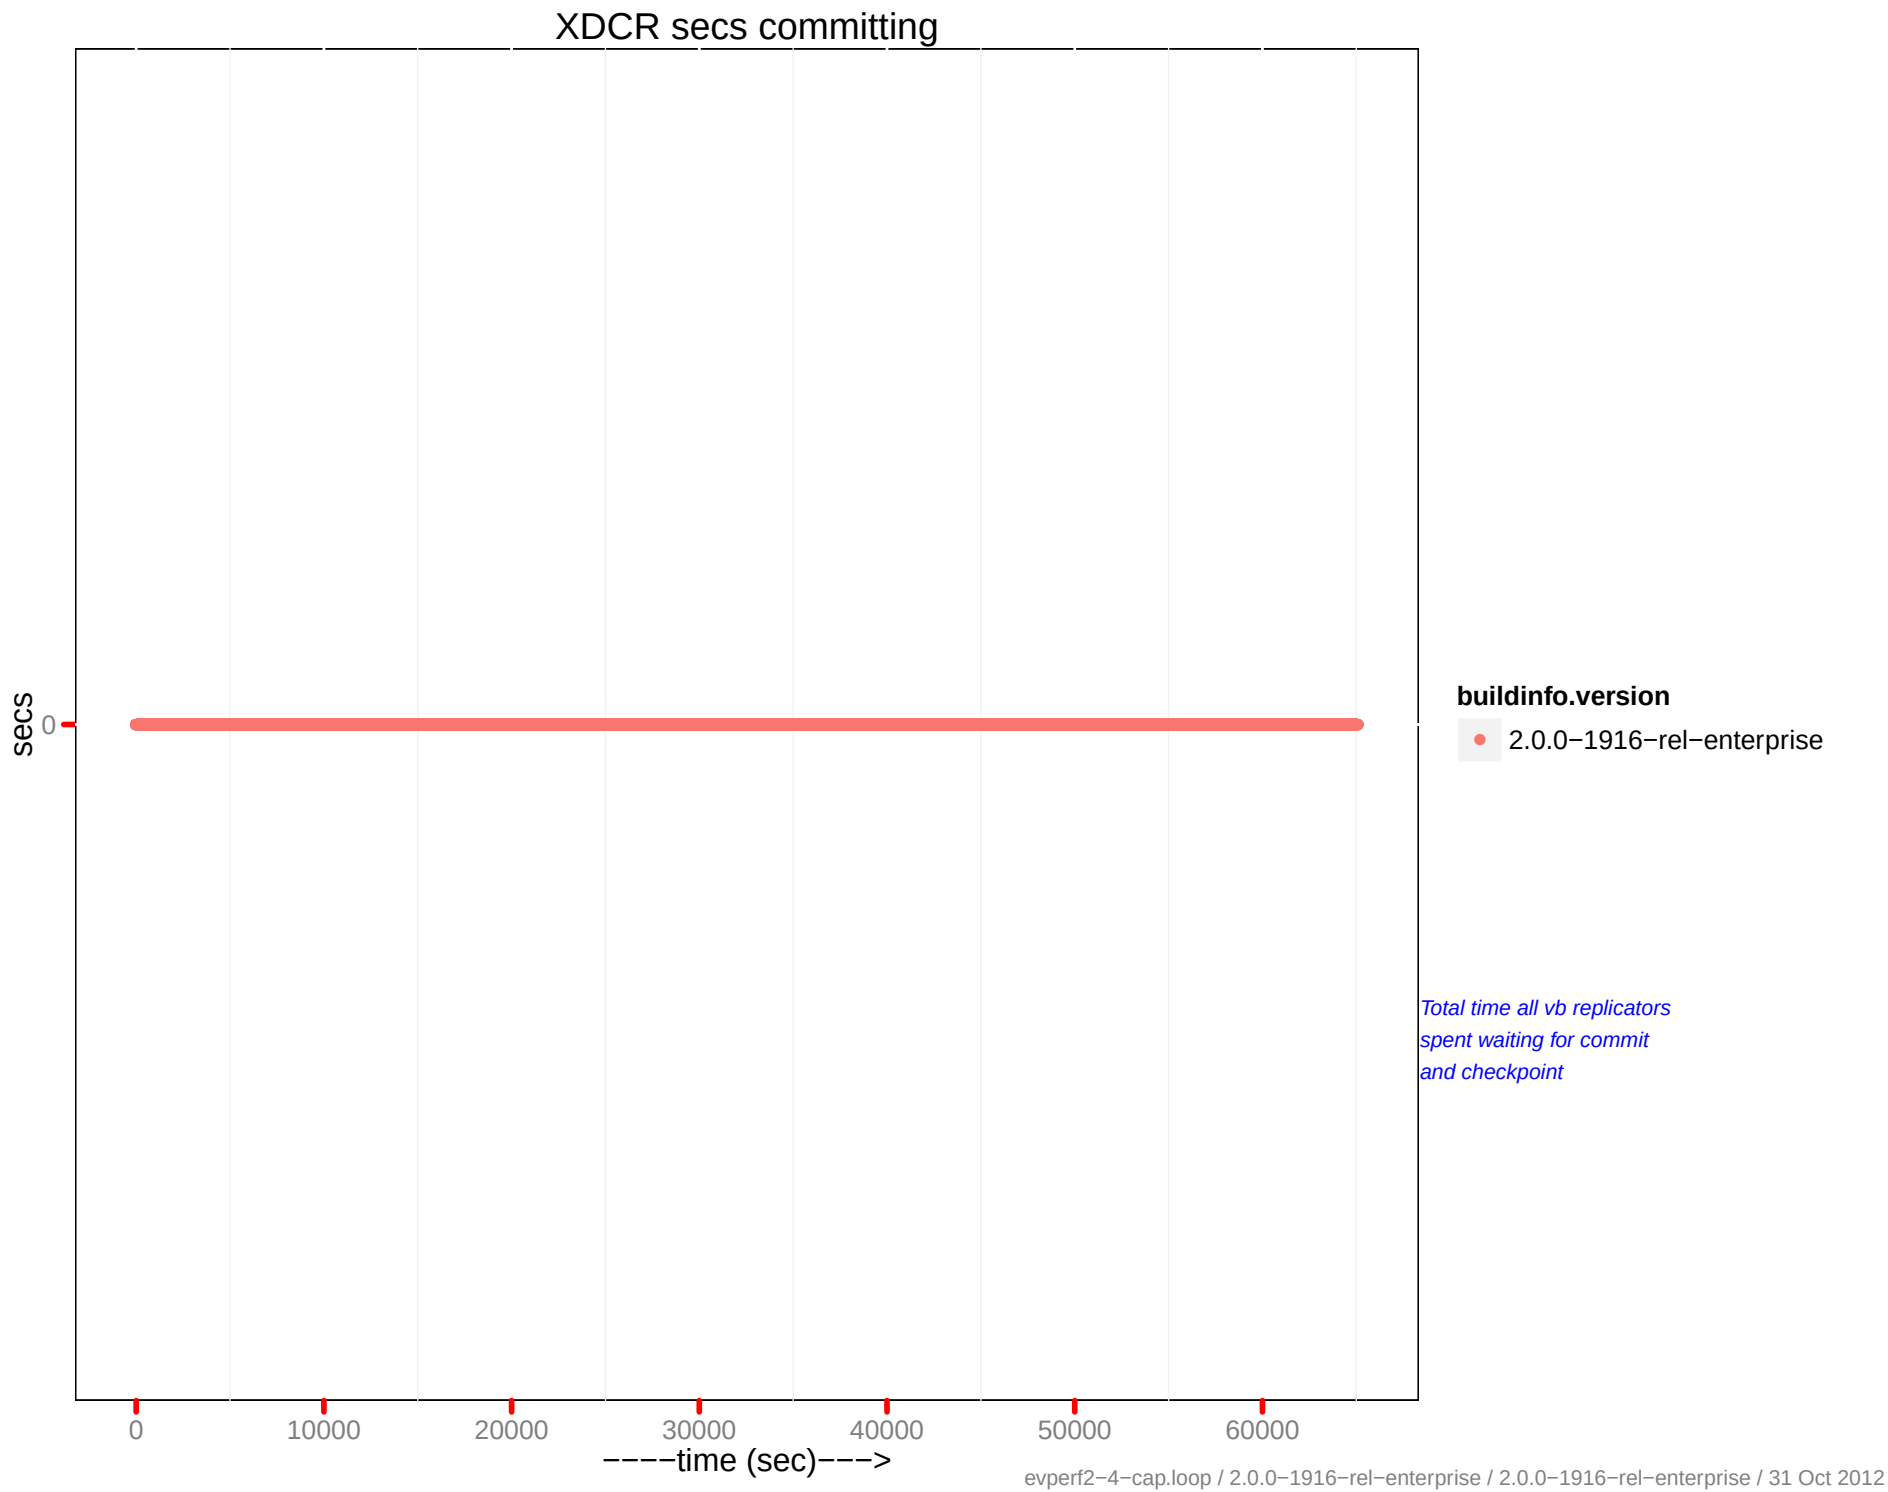
XDCR secs committing
secs
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Total time all vb replicators
spent waiting for commit
and checkpoint
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

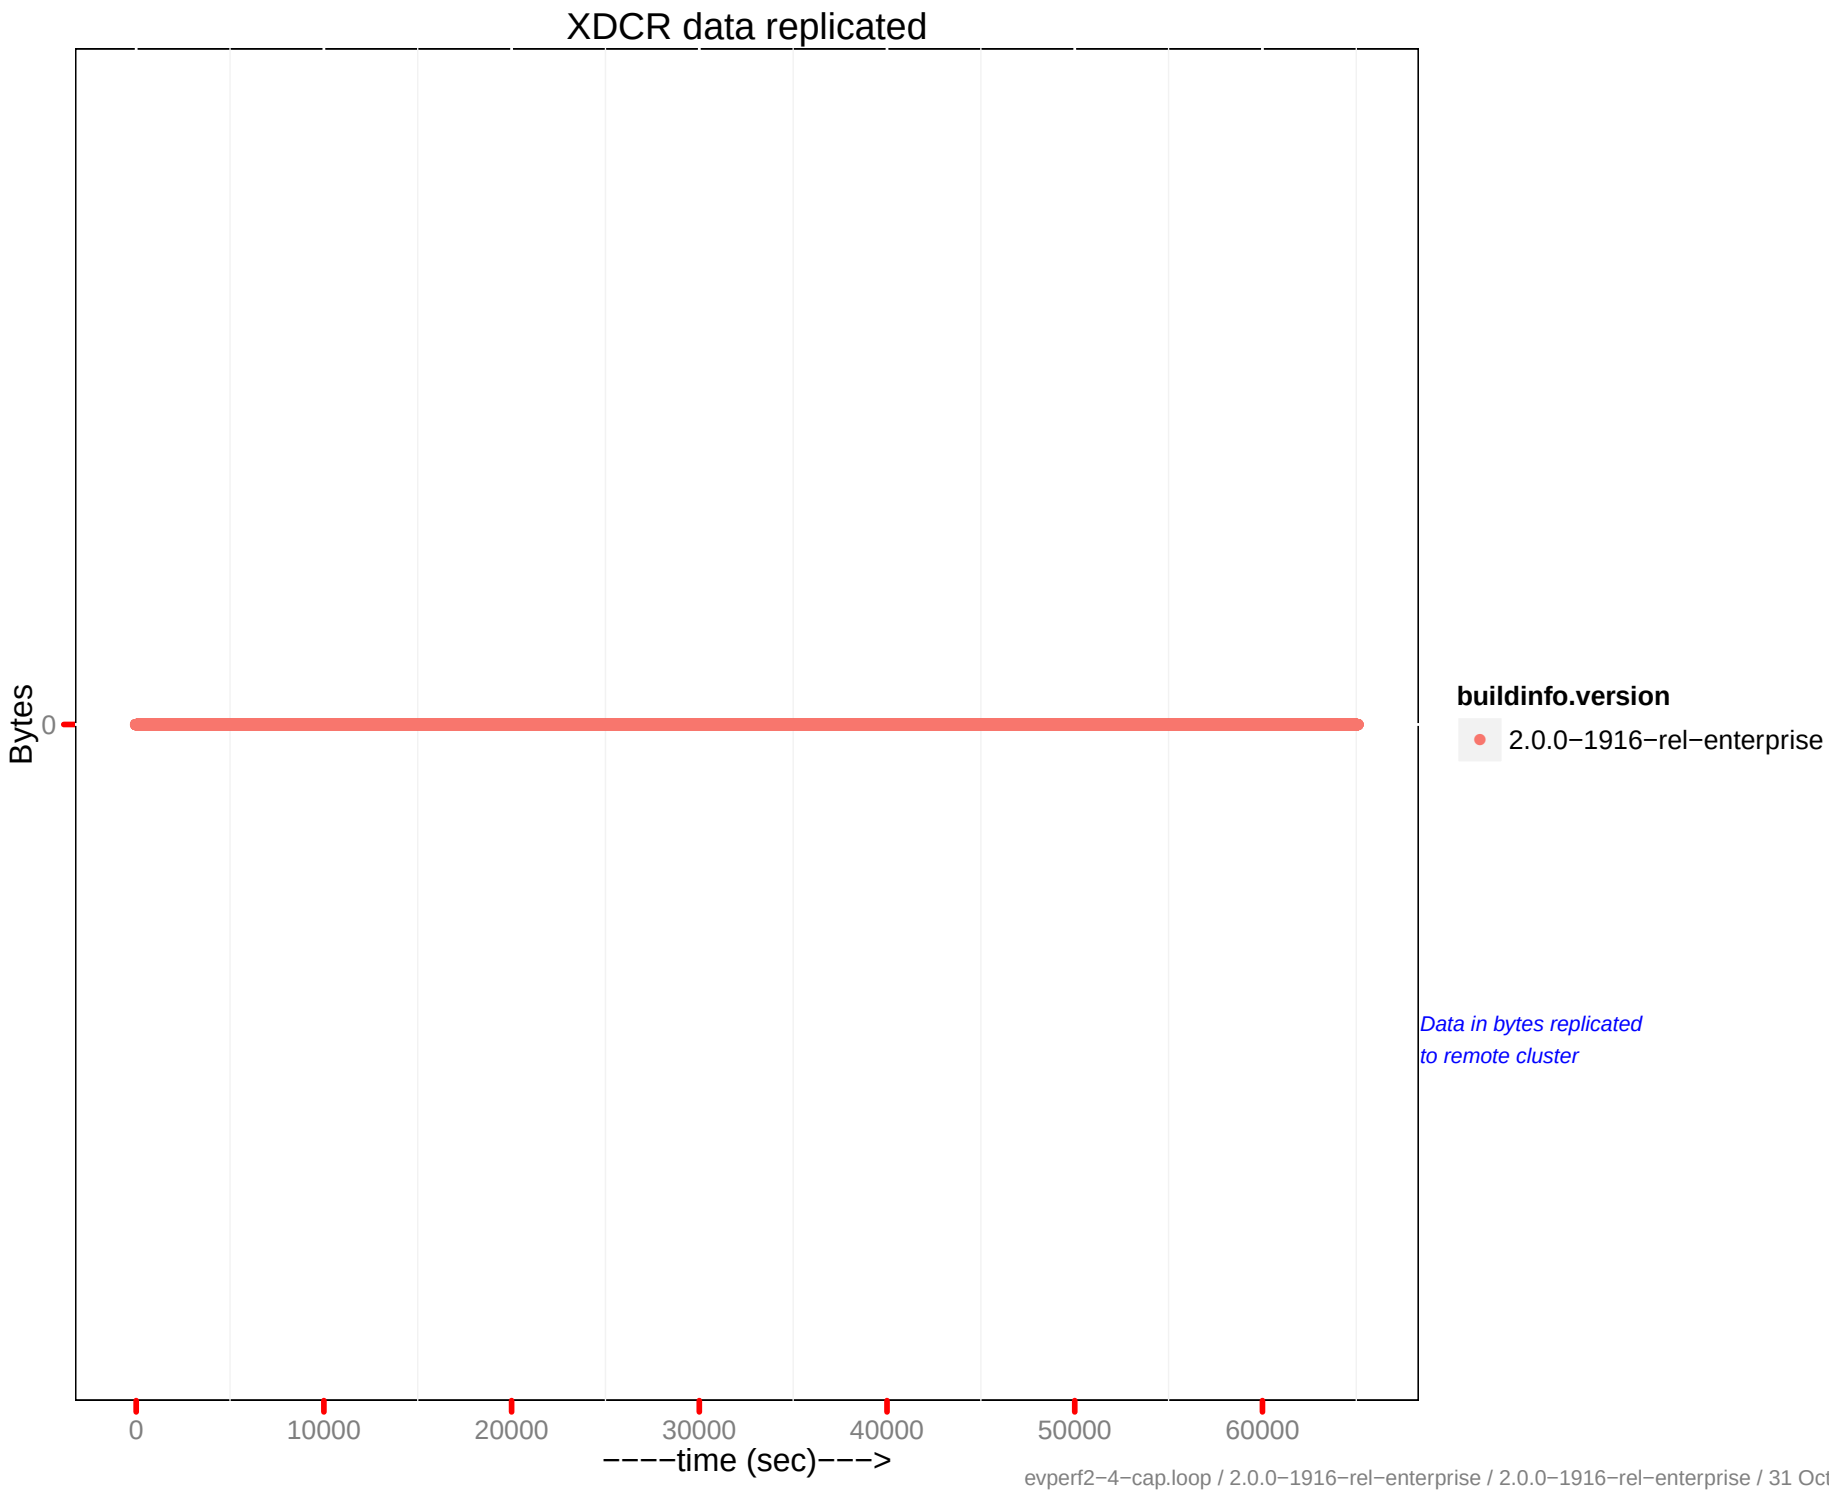
XDCR data replicated
Bytes
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Data in bytes replicated
to remote cluster
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

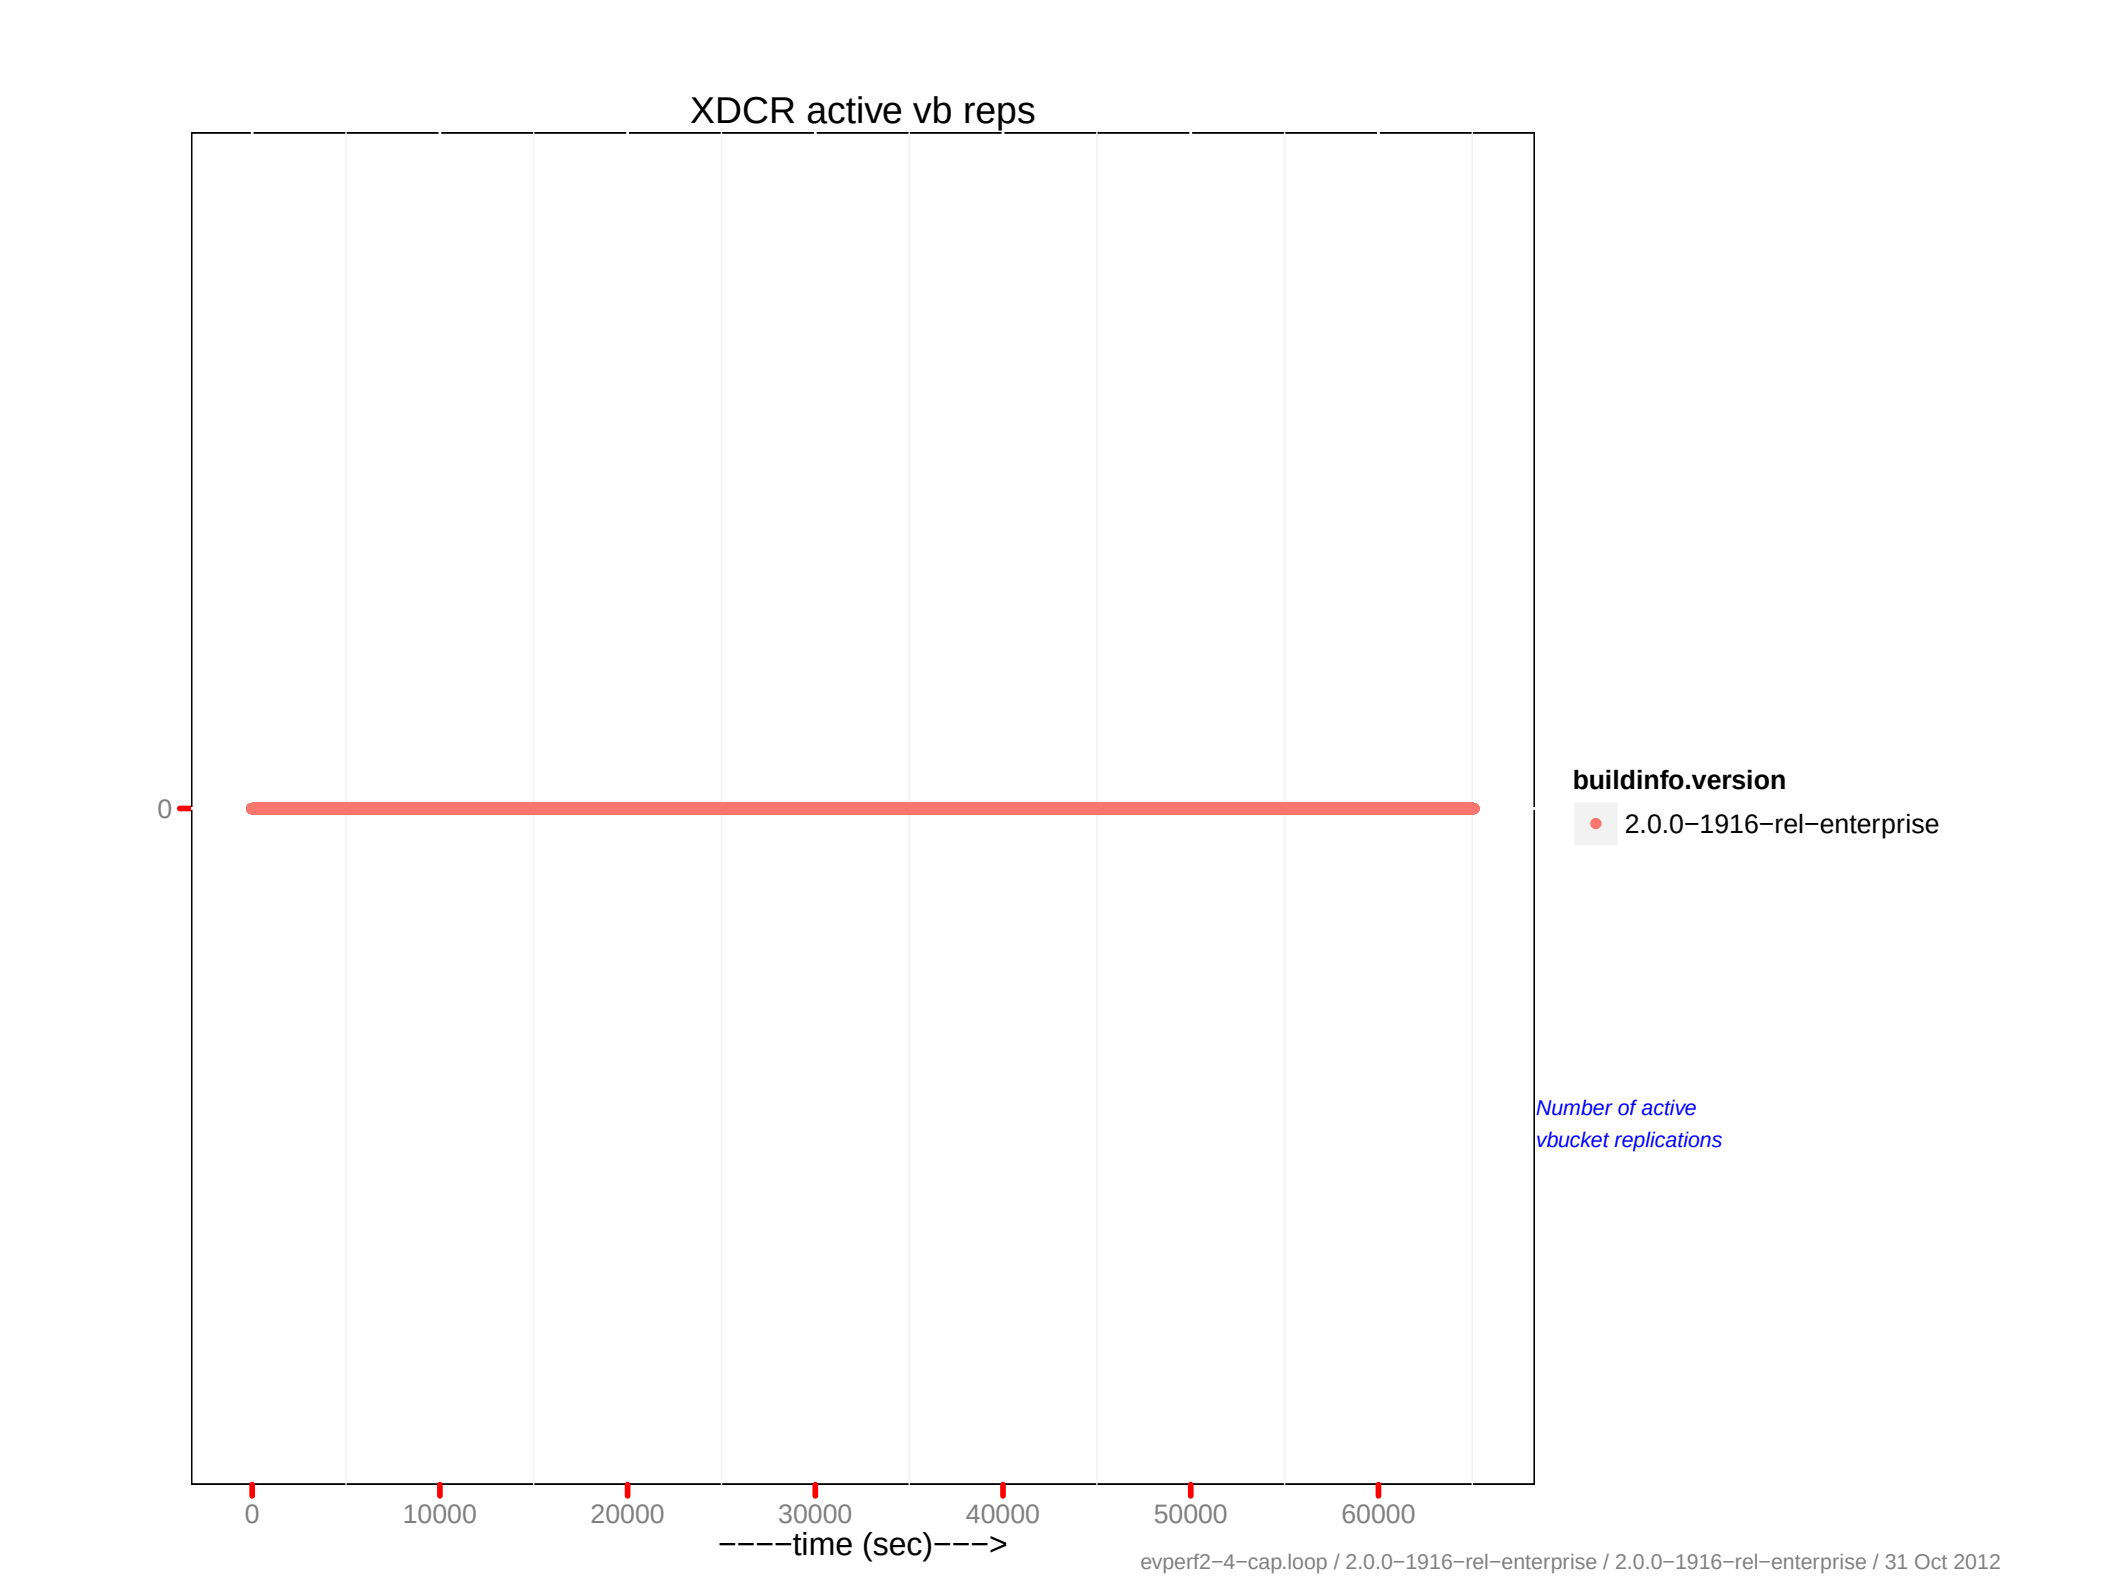
XDCR active vb reps
0
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Number of active
vbucket replications
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

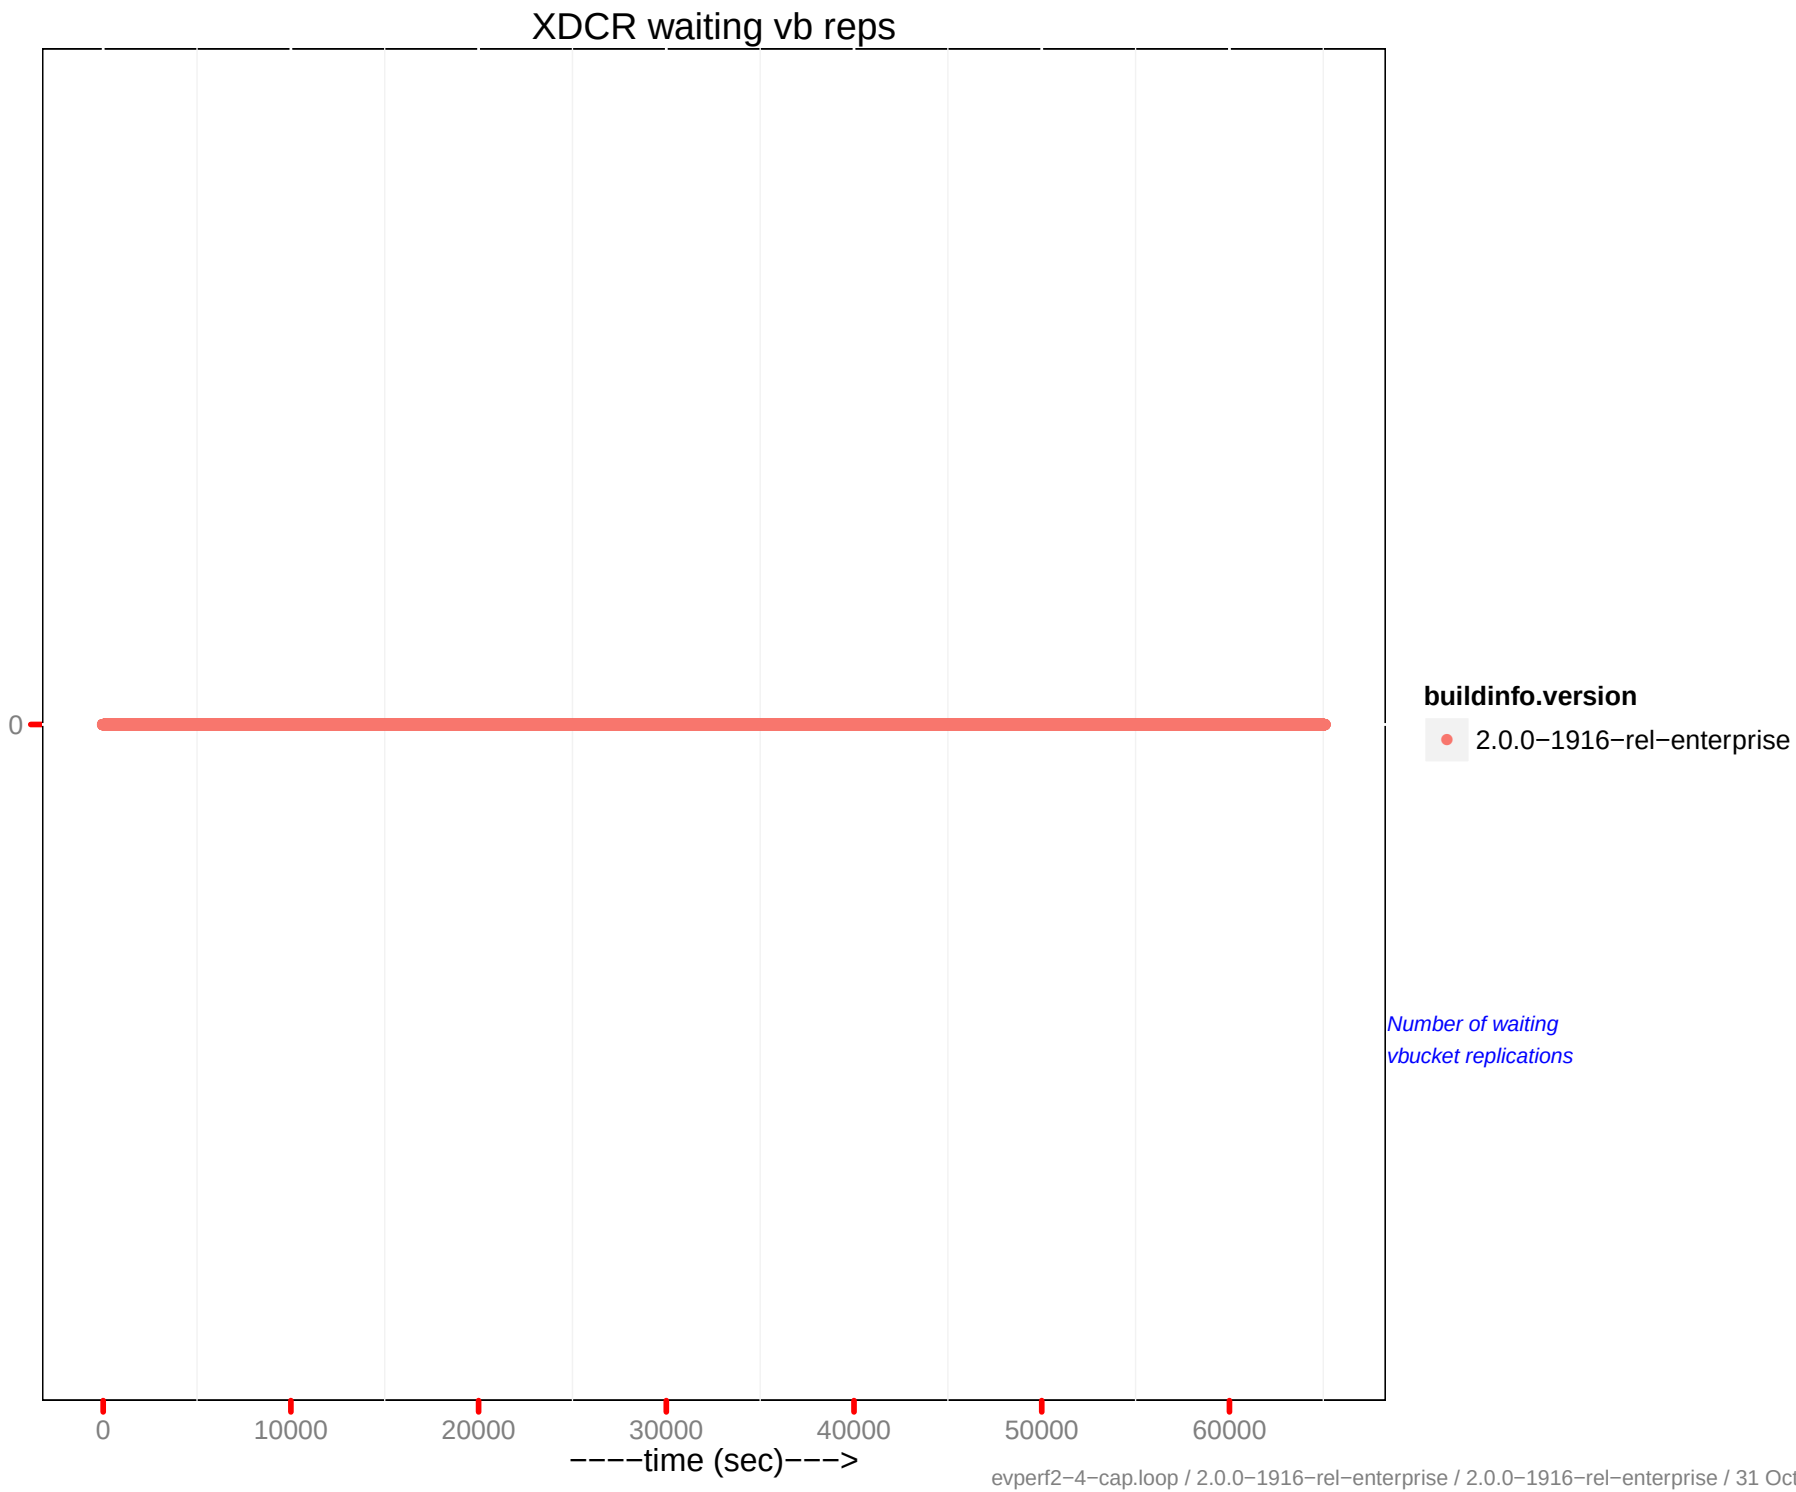
XDCR waiting vb reps
0
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
Number of waiting
vbucket replications
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012


CPU utilization − ec2−107−21−188−36.compute−1.amazonaws.com:8091
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

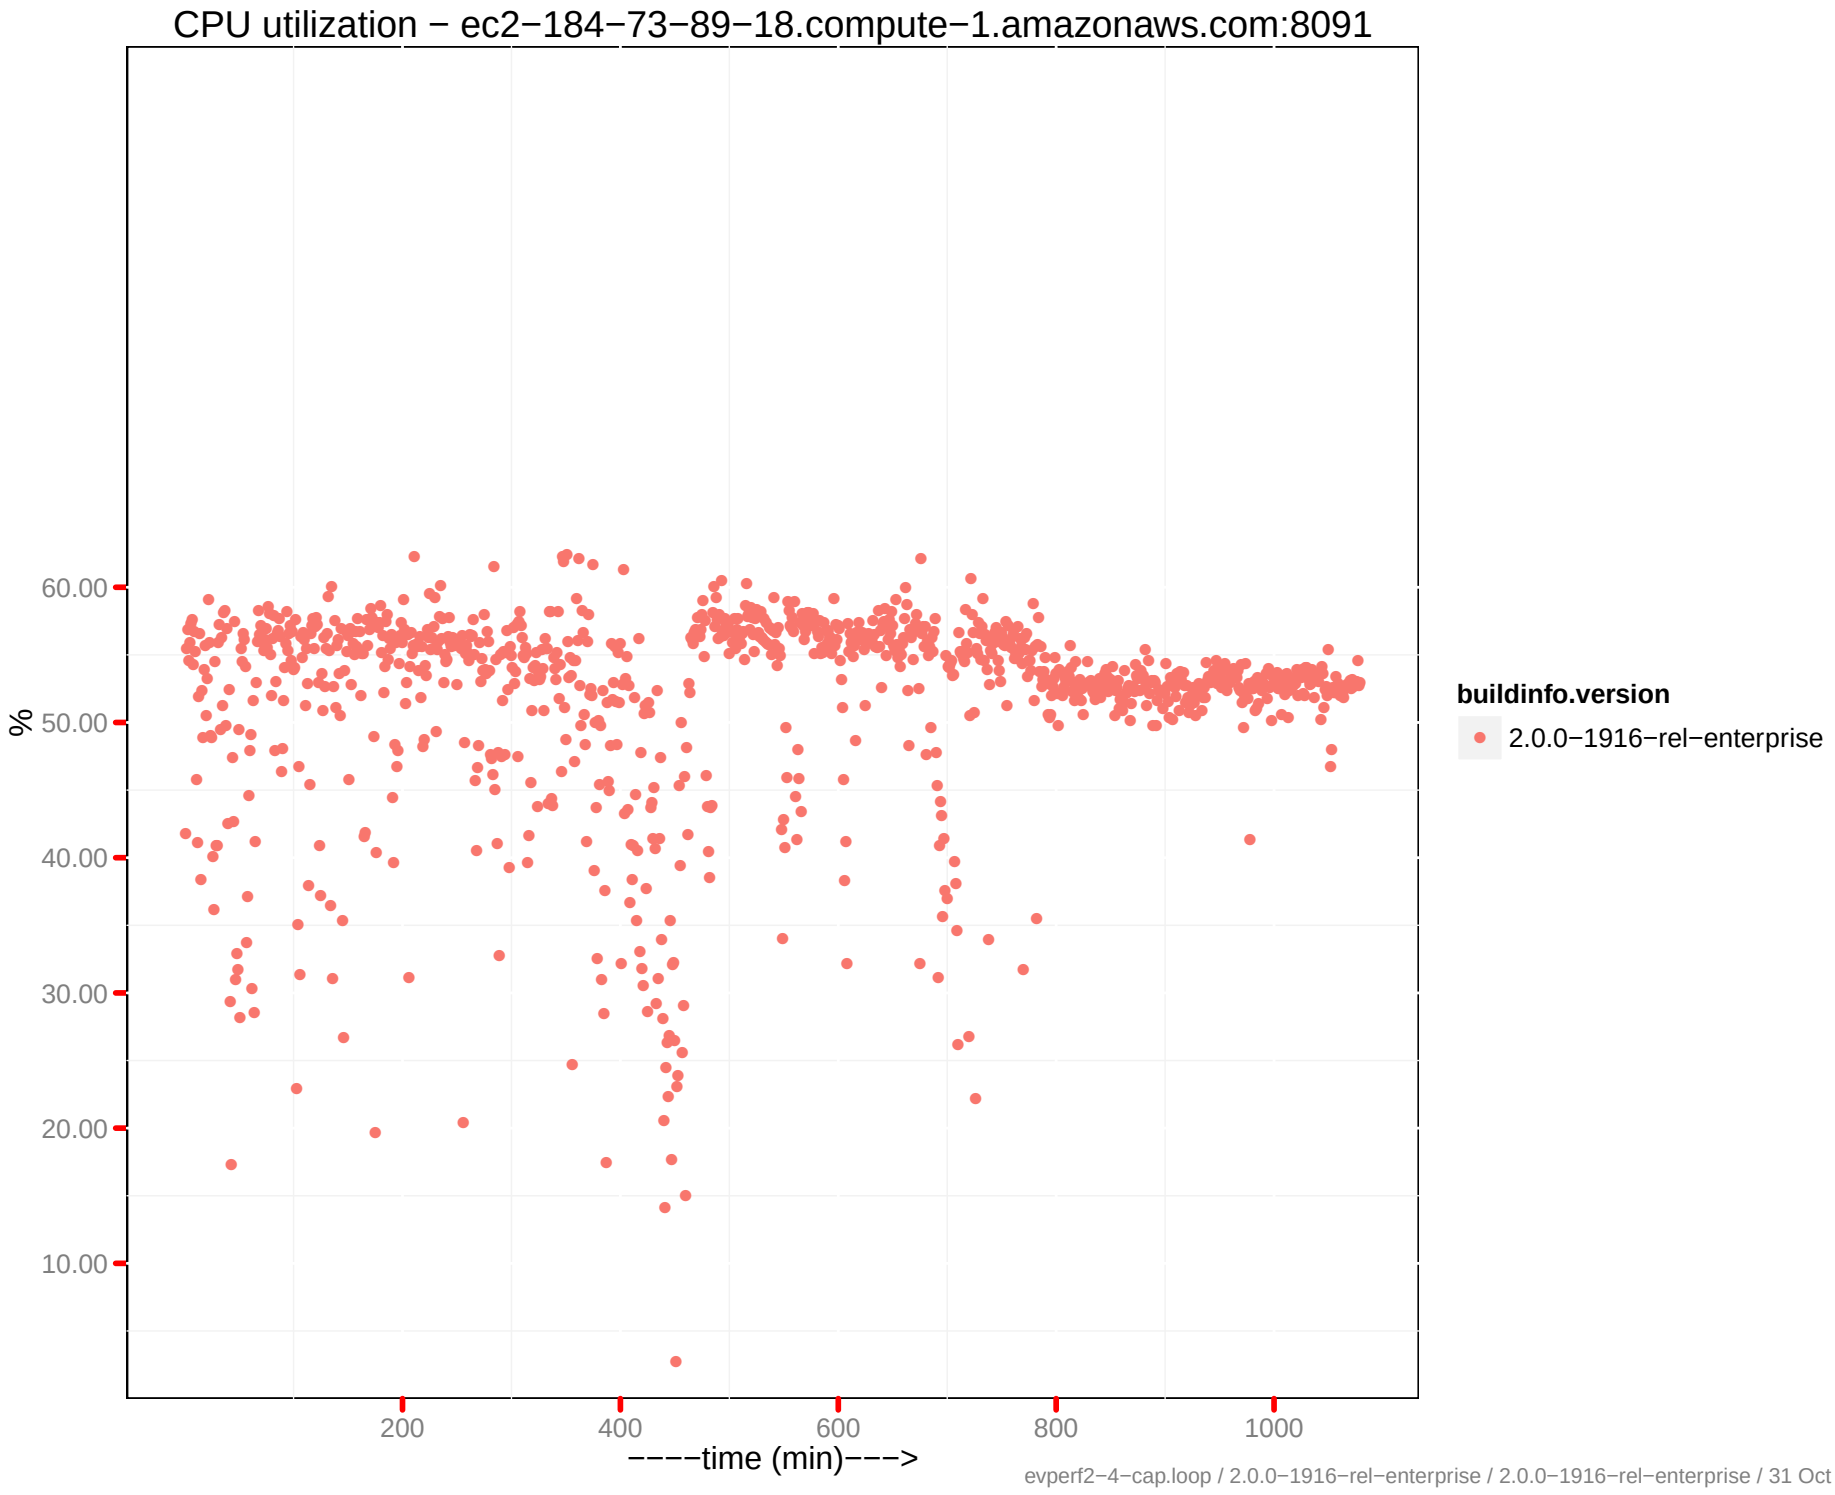
CPU utilization − ec2−184−73−89−18.compute−1.amazonaws.com:8091
60.00
50.00
40.00
30.00
20.00
10.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012


CPU utilization − ec2−204−236−244−32.compute−1.amazonaws.com:8091
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

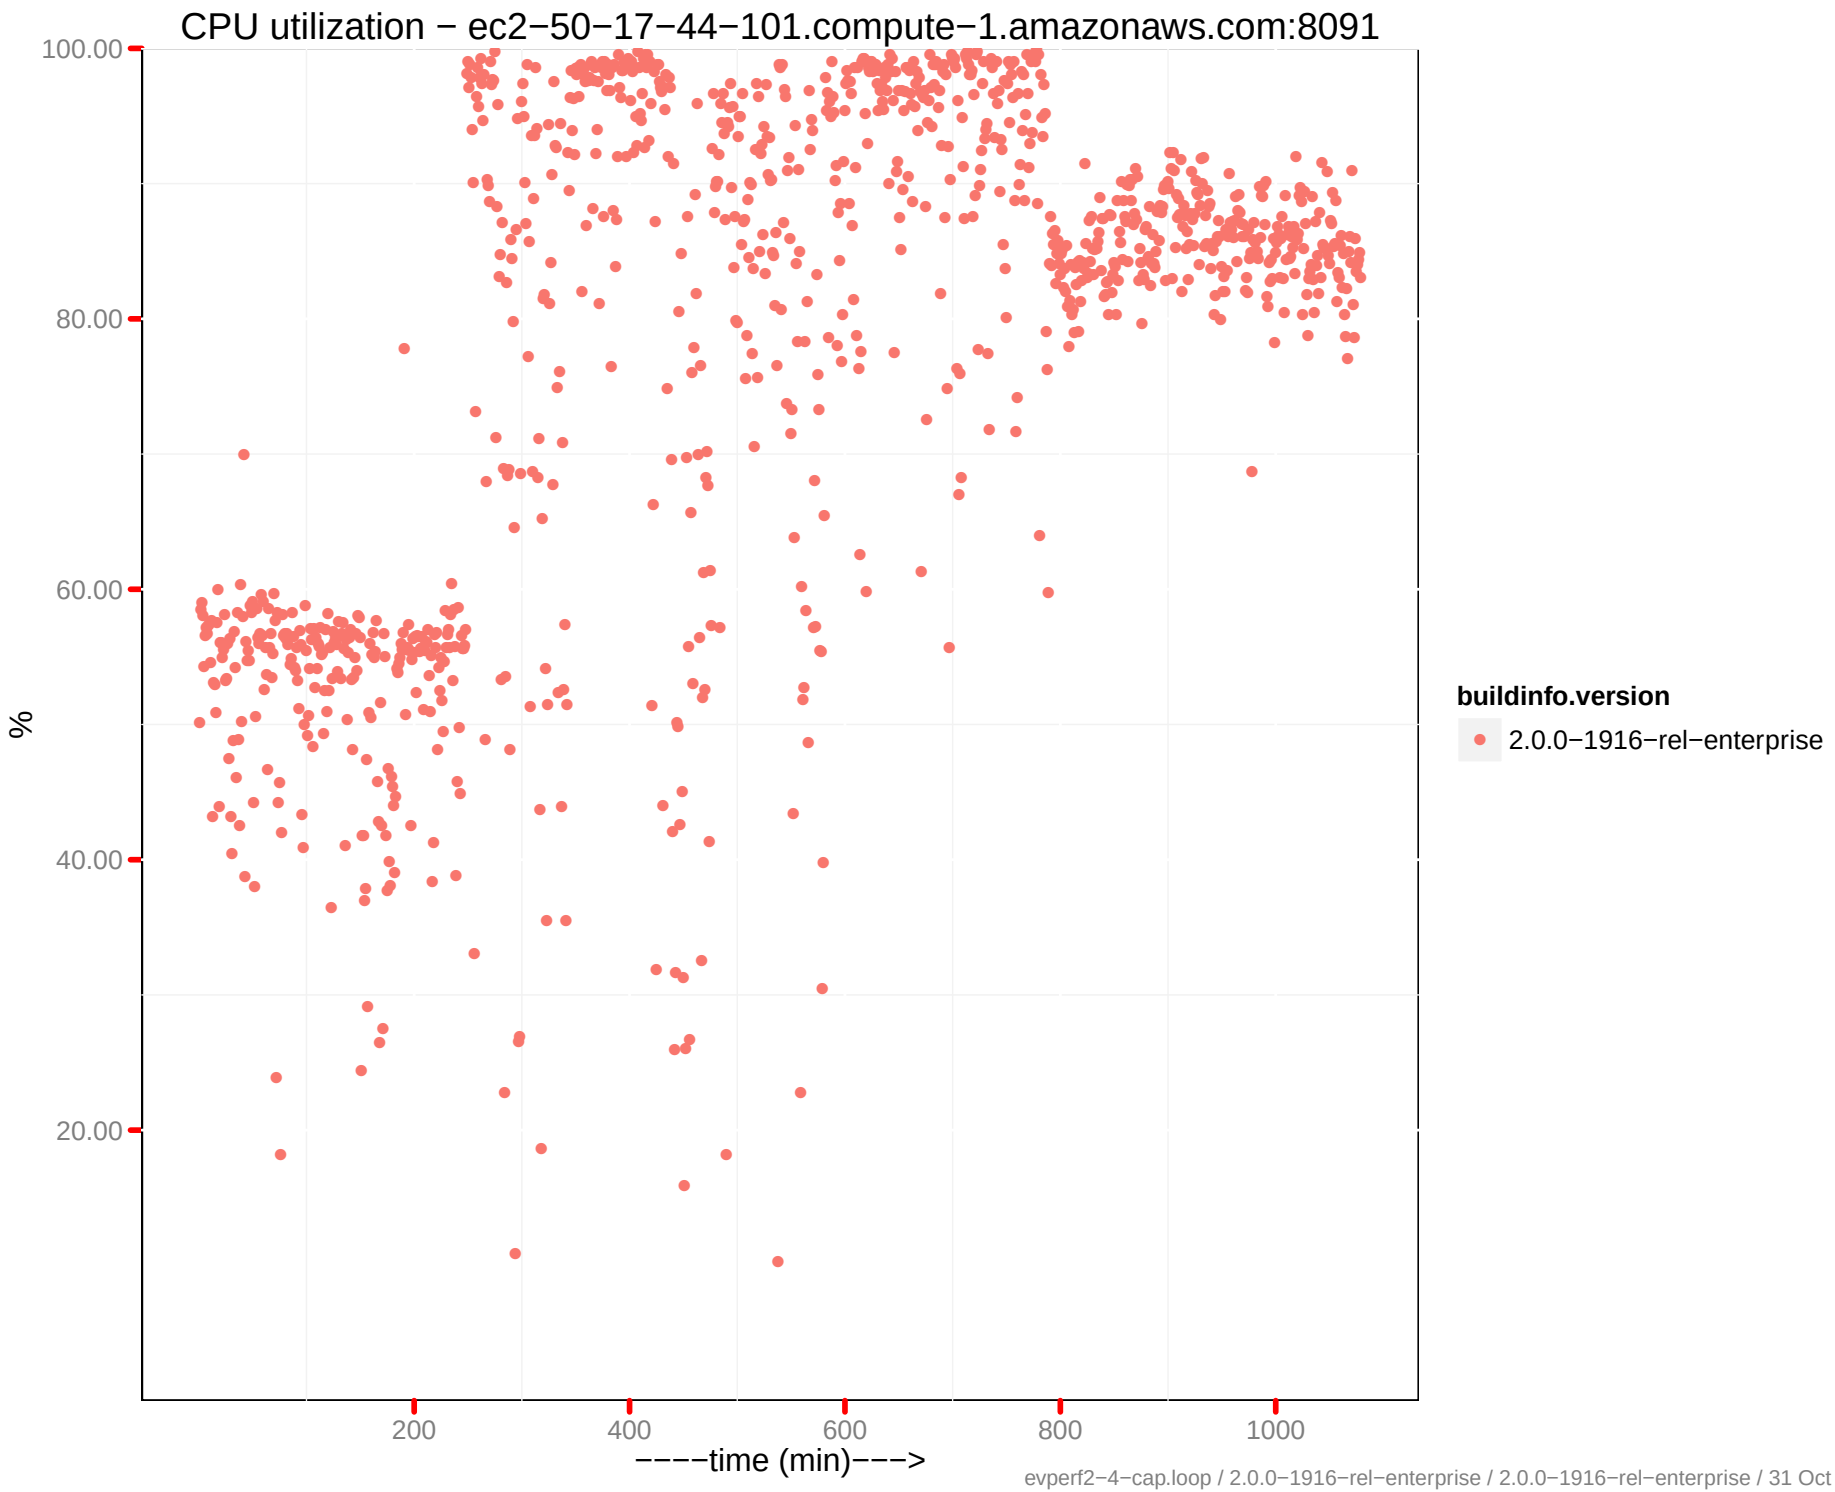
CPU utilization − ec2−50−17−44−101.compute−1.amazonaws.com:8091
100.00
80.00
60.00
40.00
20.00
%
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

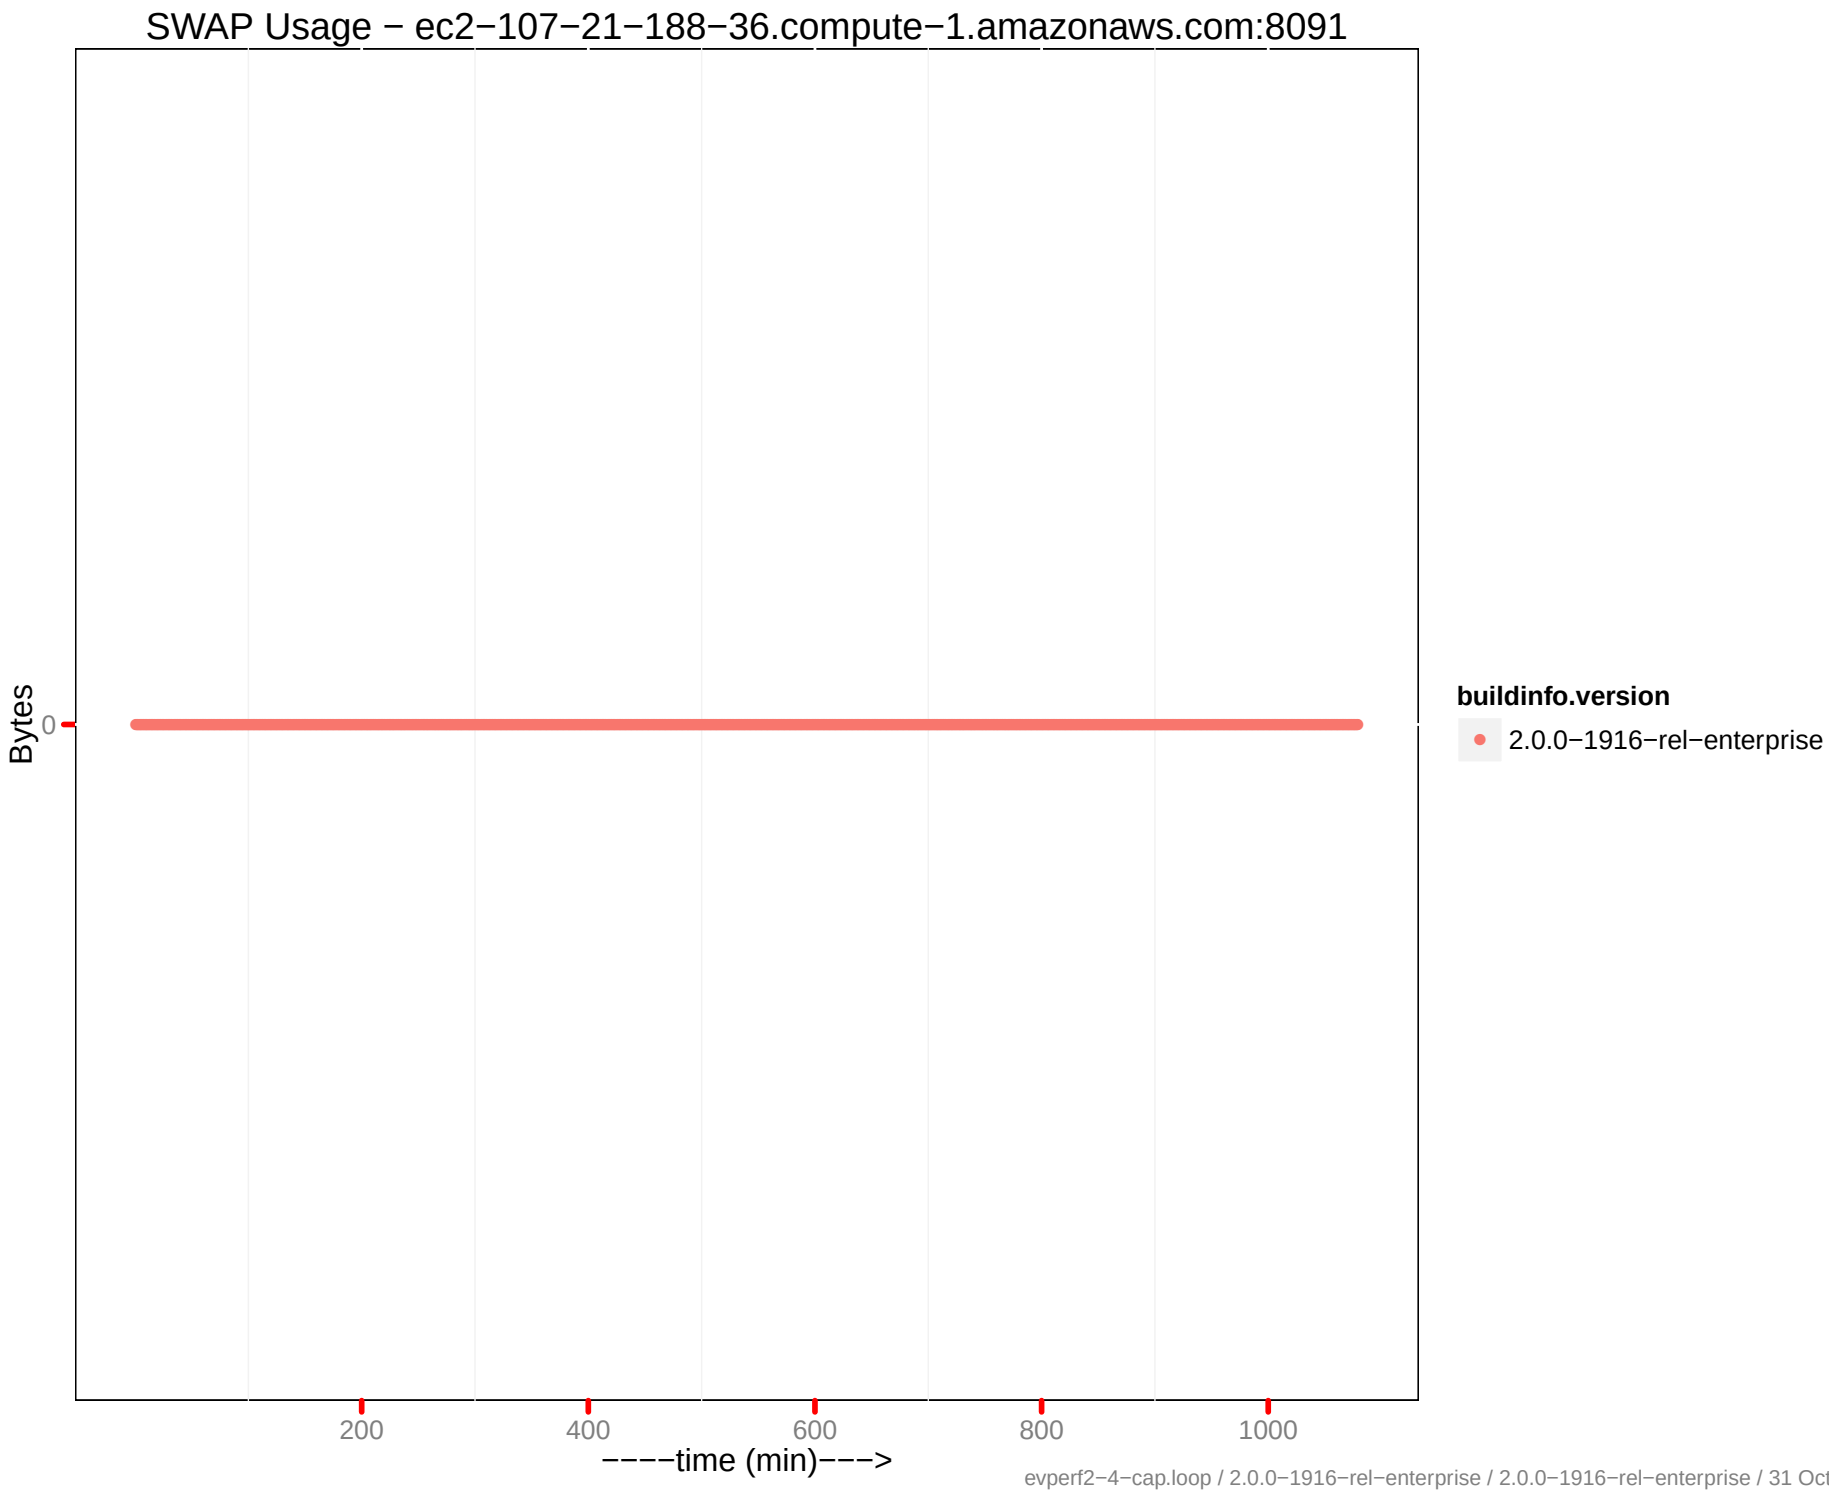
SWAP Usage − ec2−107−21−188−36.compute−1.amazonaws.com:8091
Bytes
0
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

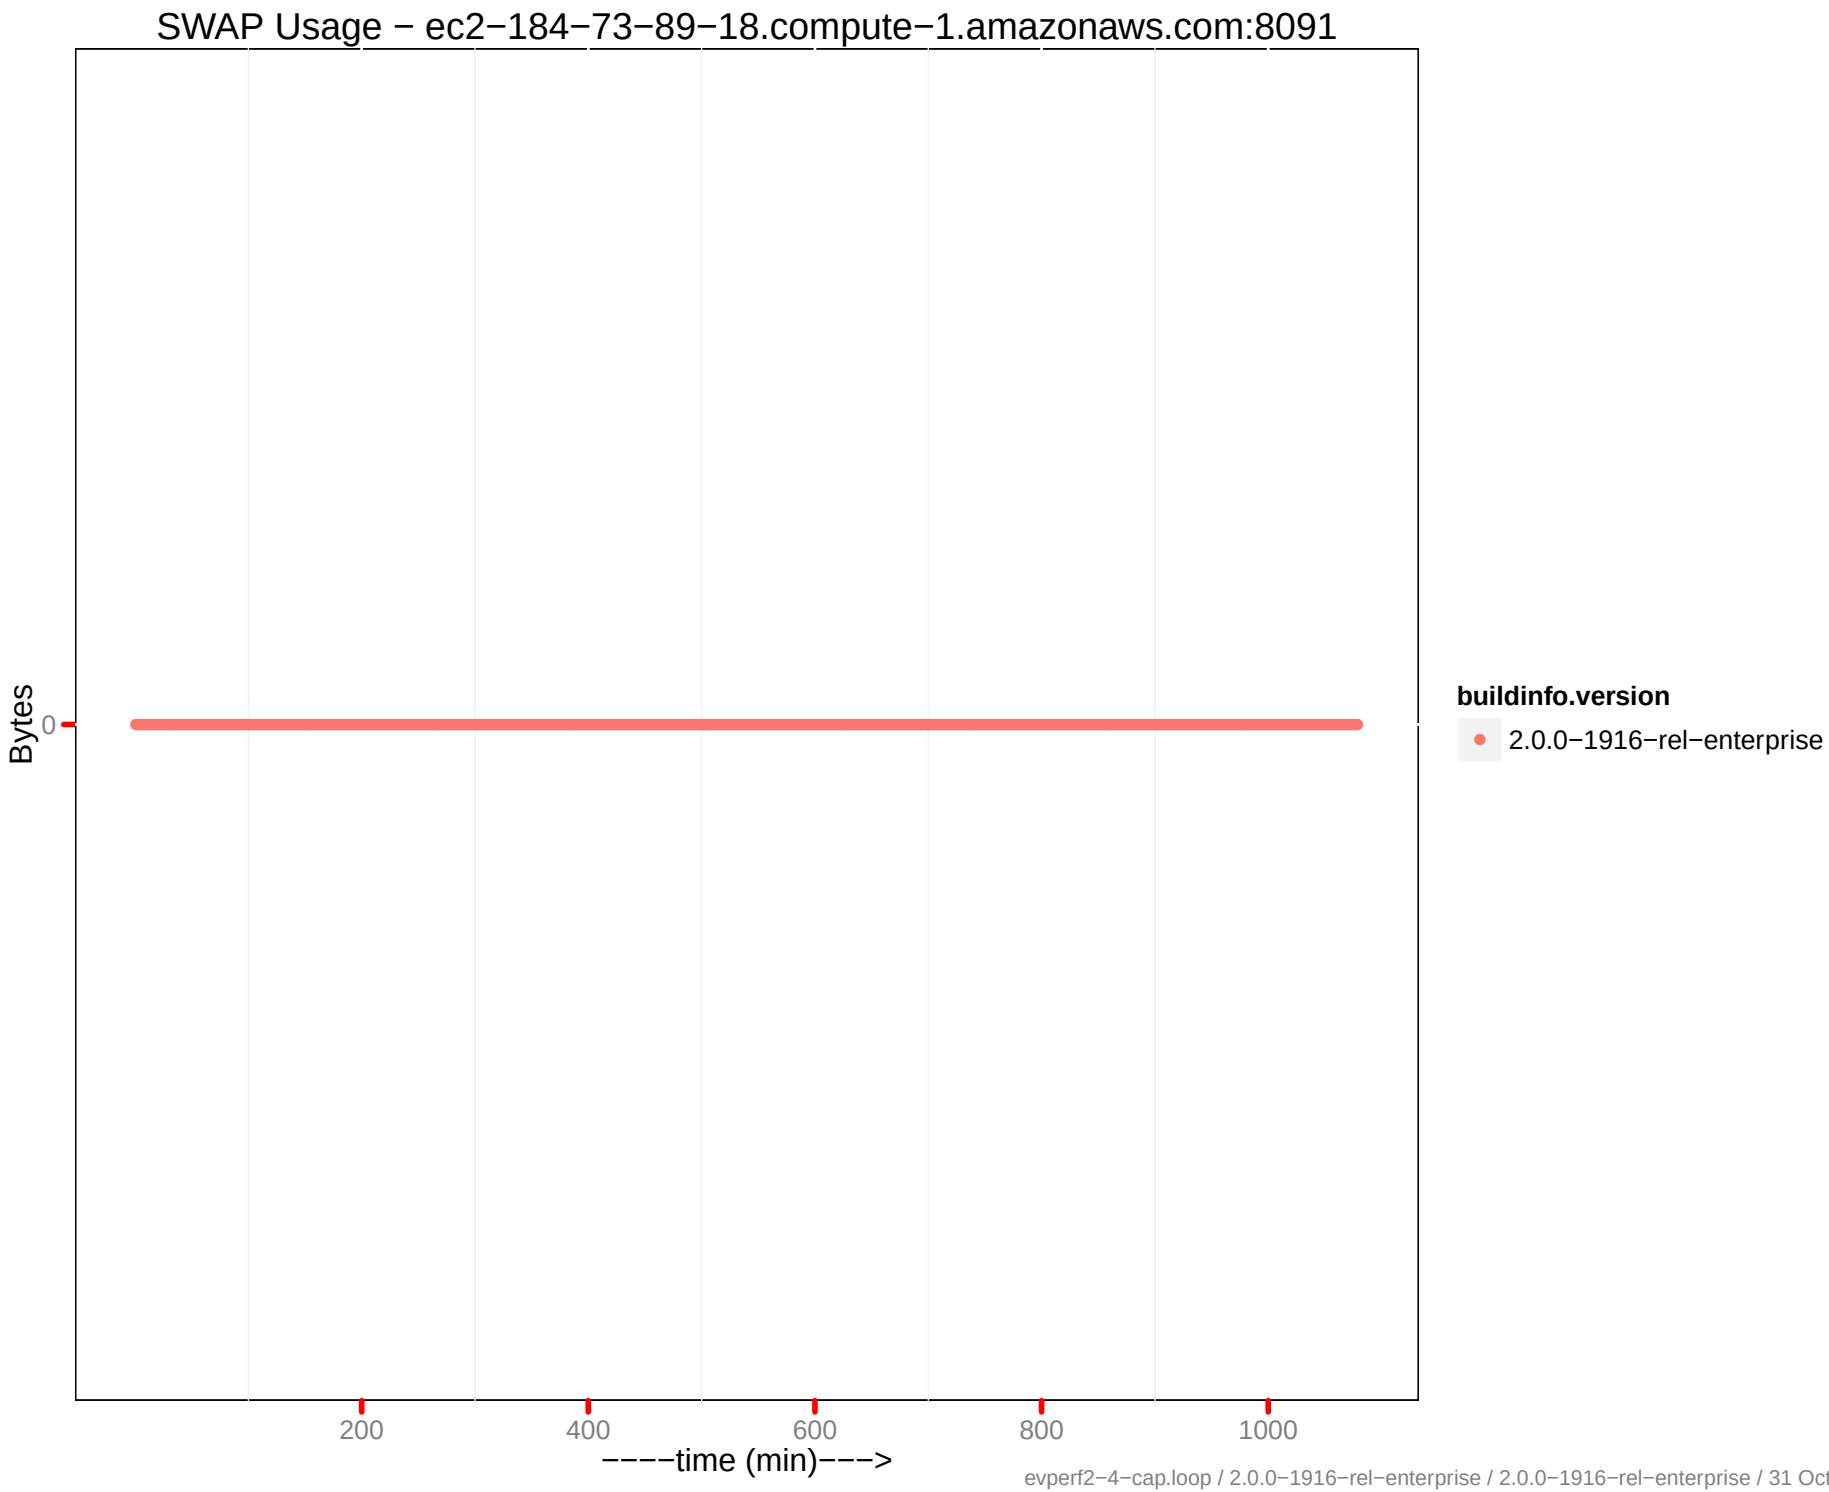
SWAP Usage − ec2−184−73−89−18.compute−1.amazonaws.com:8091
Bytes
0
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

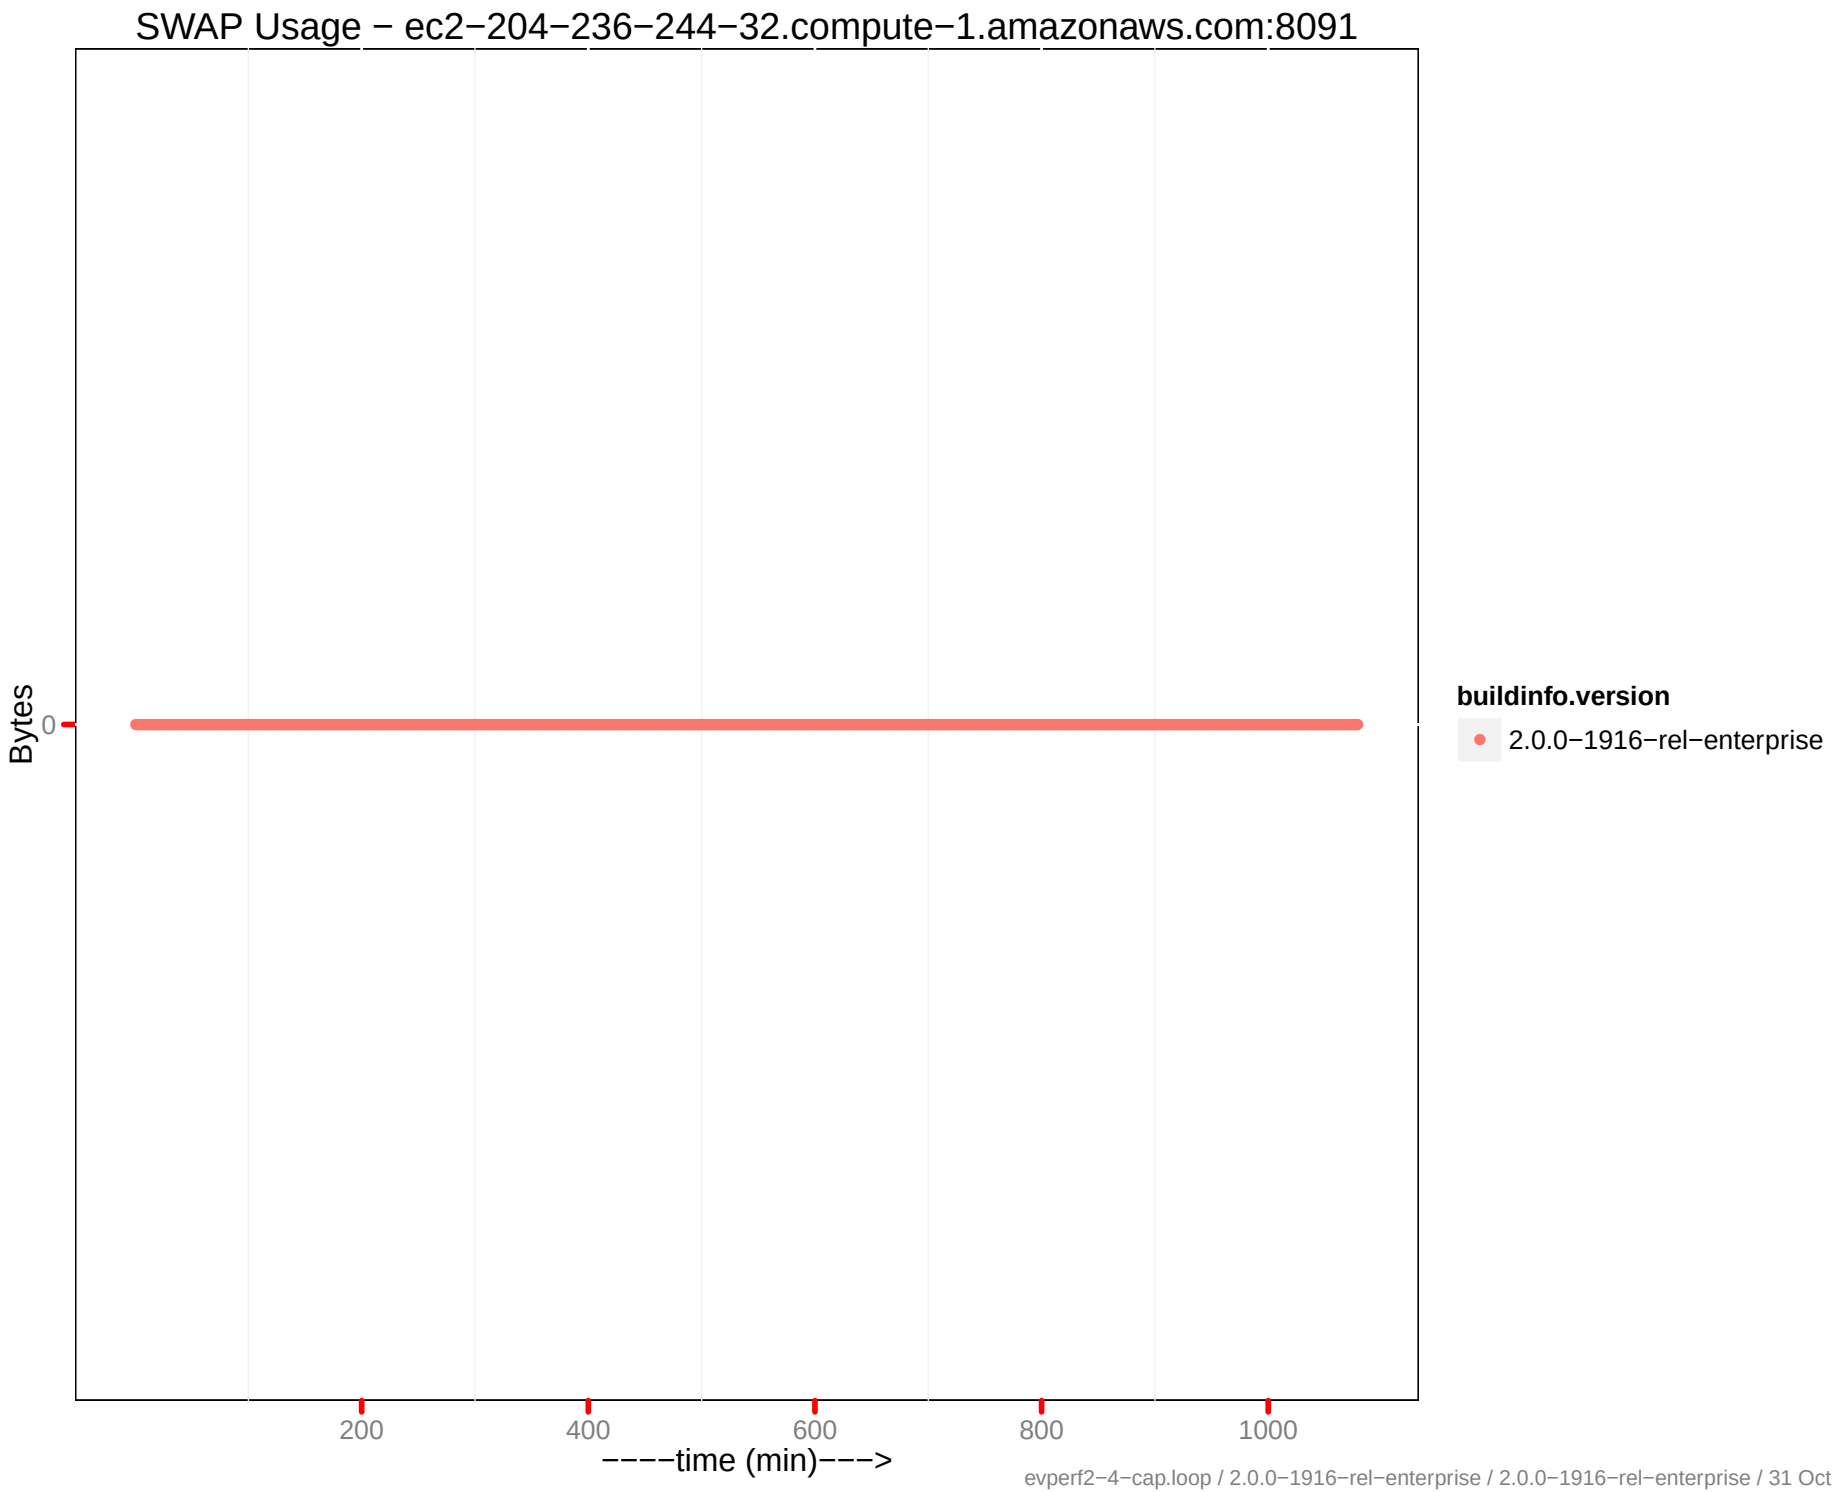
SWAP Usage − ec2−204−236−244−32.compute−1.amazonaws.com:8091
Bytes
0
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

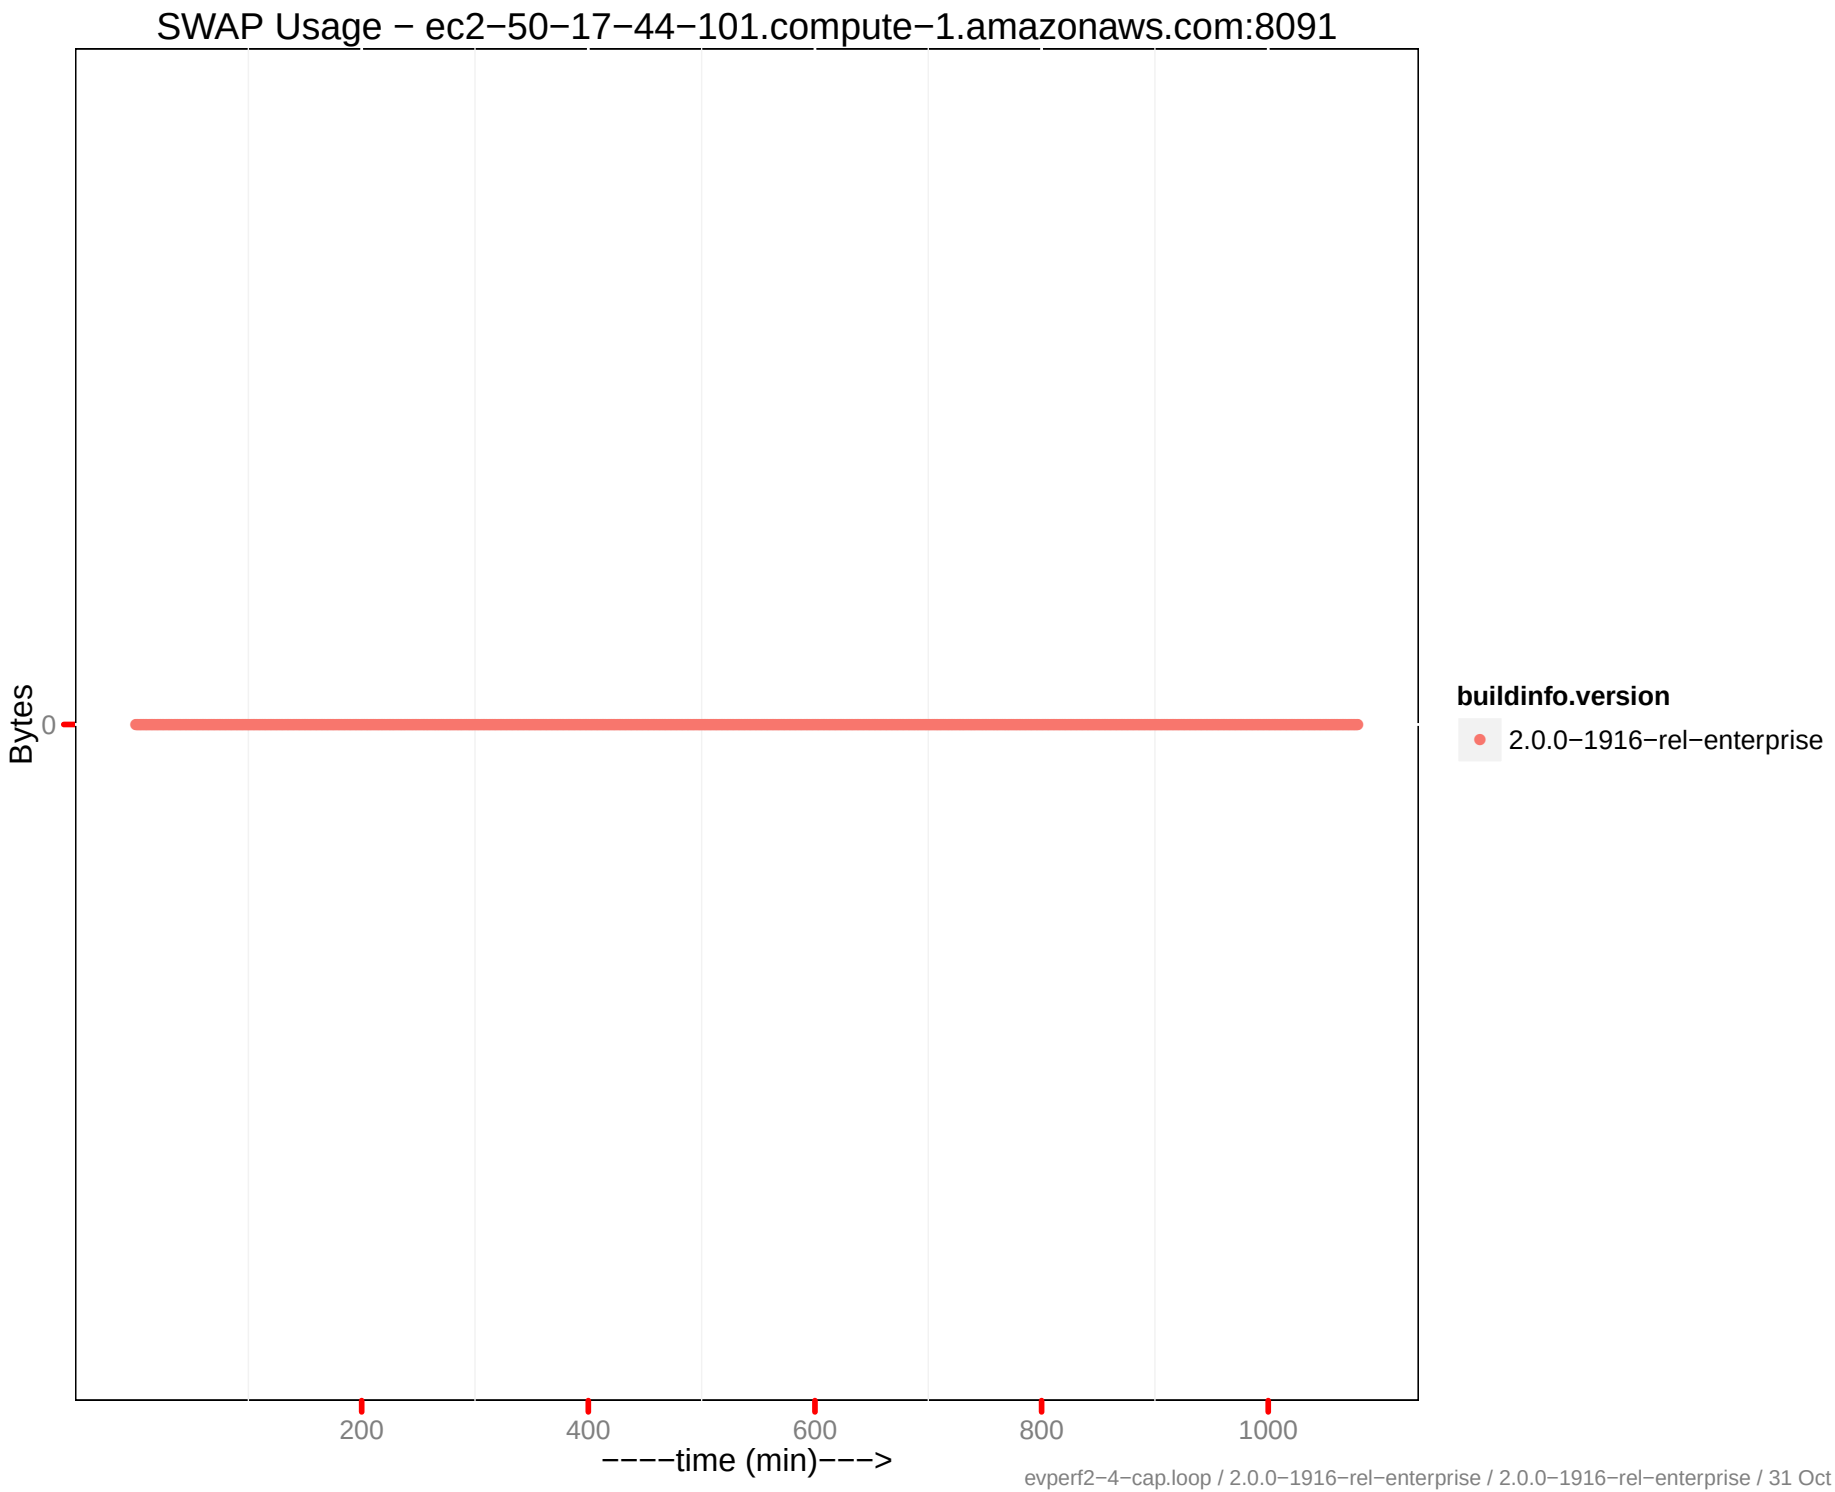
SWAP Usage − ec2−50−17−44−101.compute−1.amazonaws.com:8091
Bytes
0
200
400
600
800
1000
−−−−time (min)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

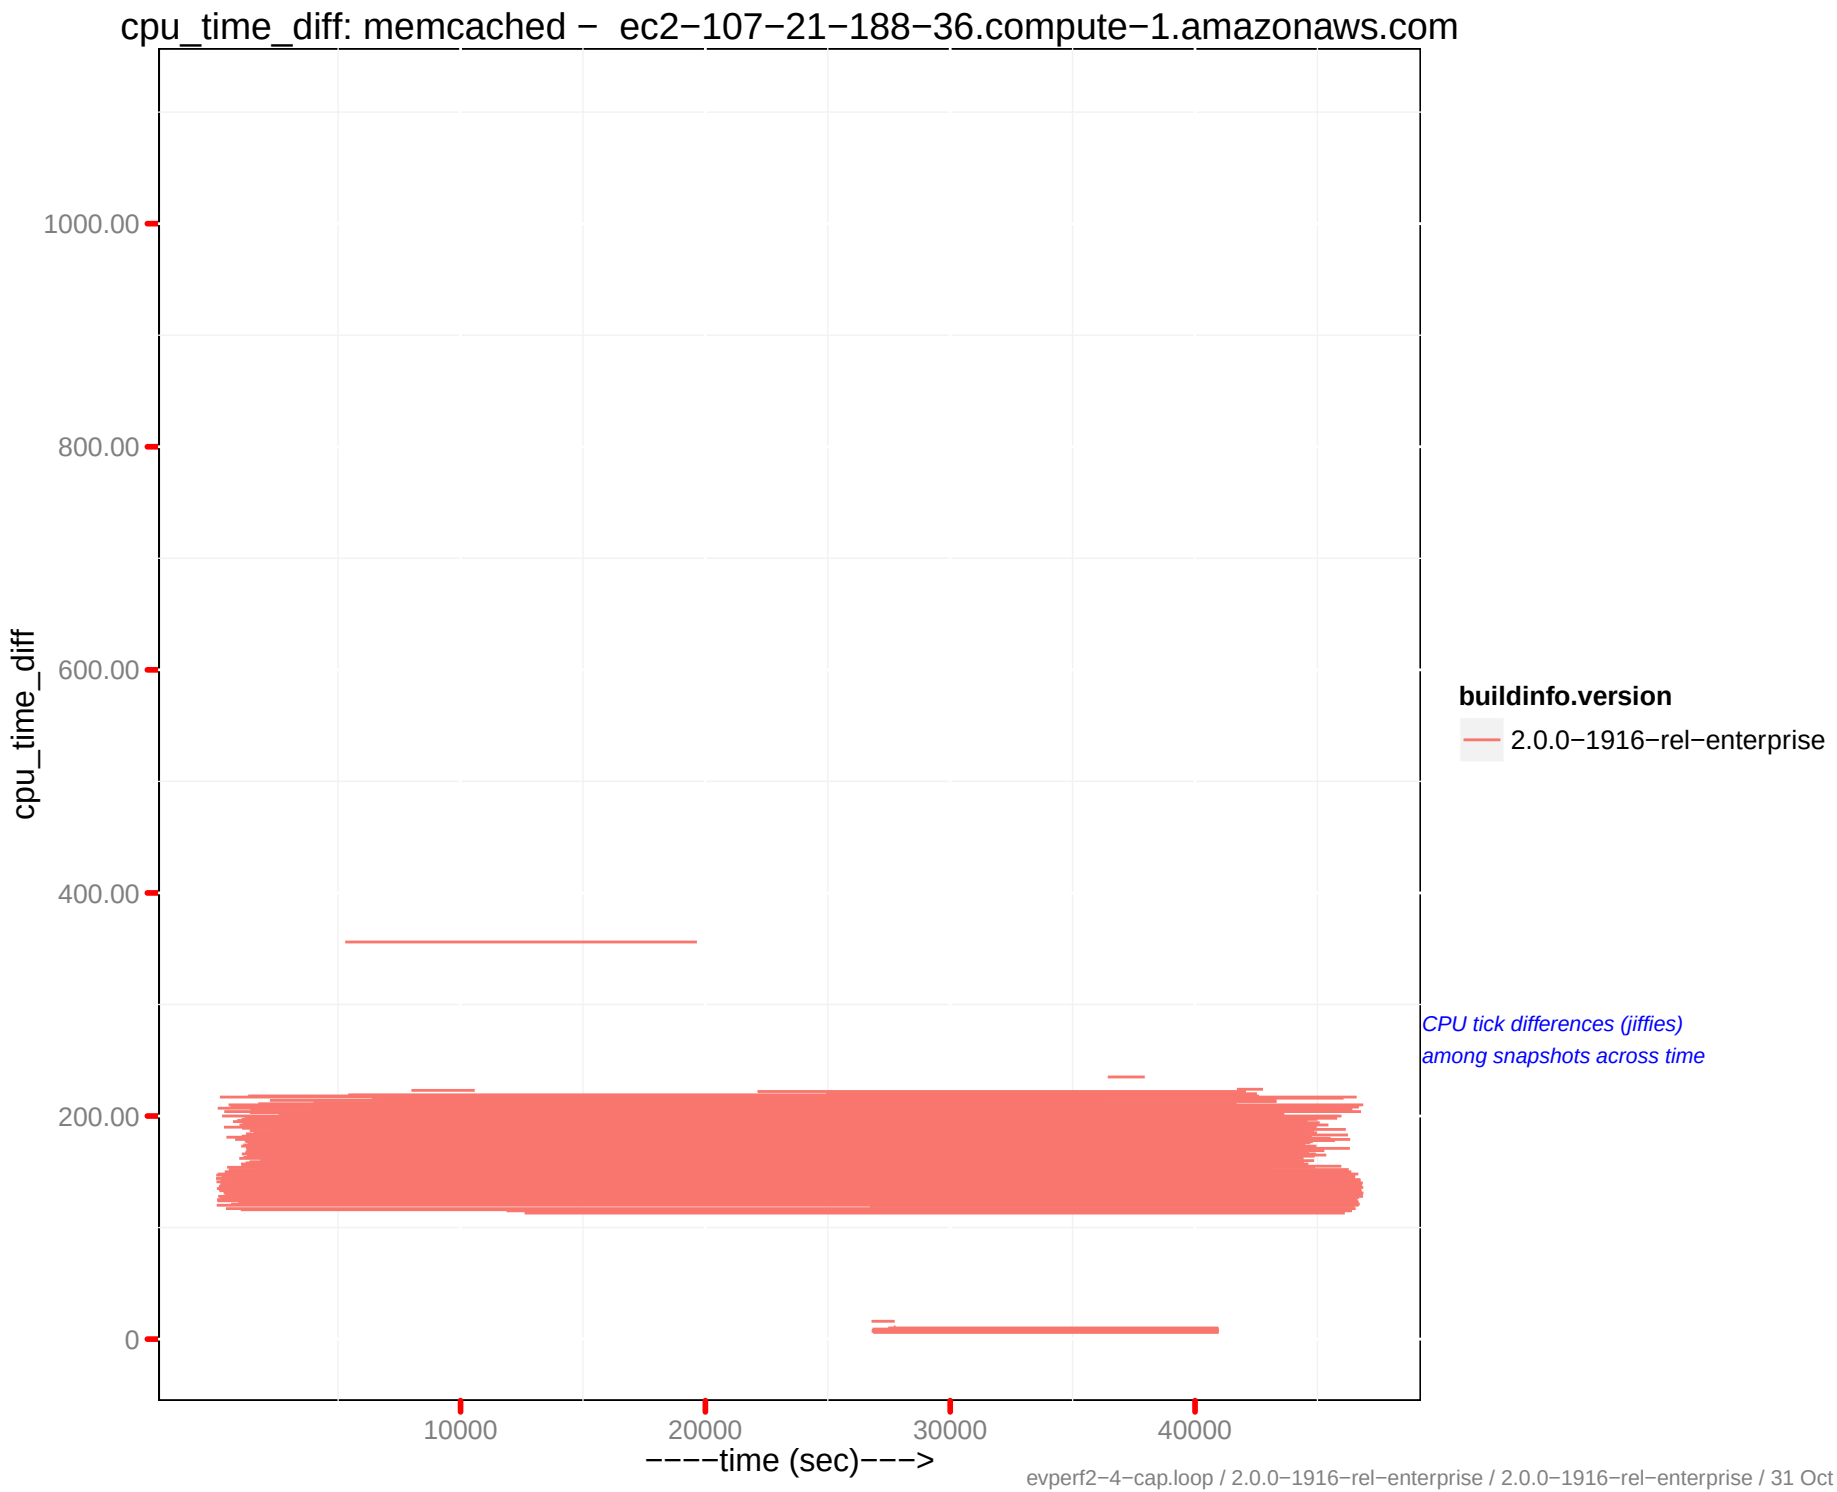
cpu_time_diff: memcached − ec2−107−21−188−36.compute−1.amazonaws.com
1000.00
800.00
600.00
400.00
200.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

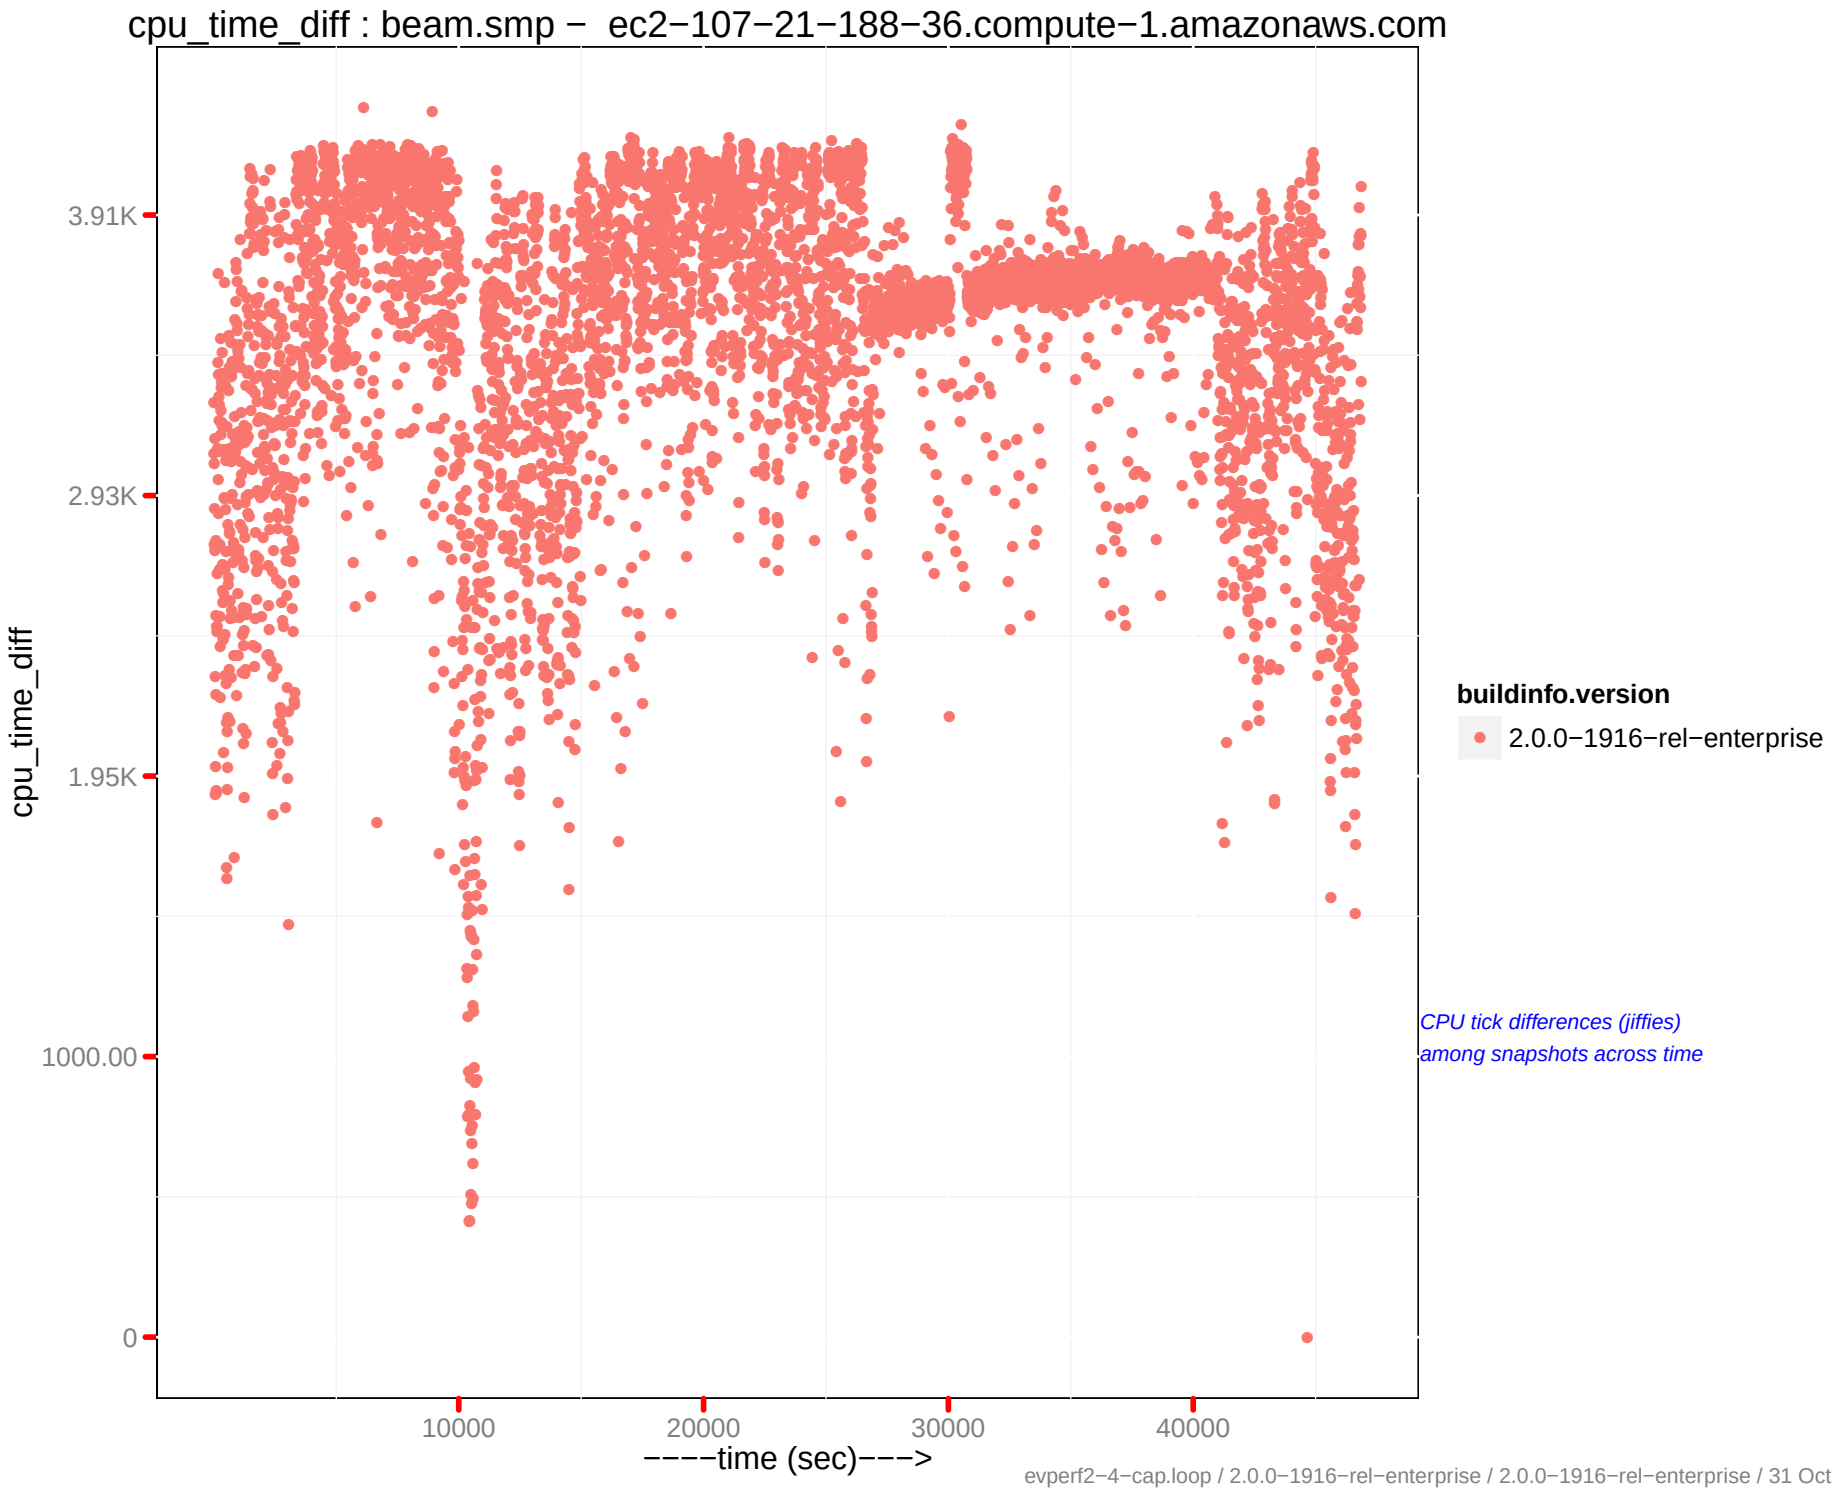
cpu_time_diff : beam.smp − ec2−107−21−188−36.compute−1.amazonaws.com
cpu_time_diff
3.91K
2.93K
1.95K
1000.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

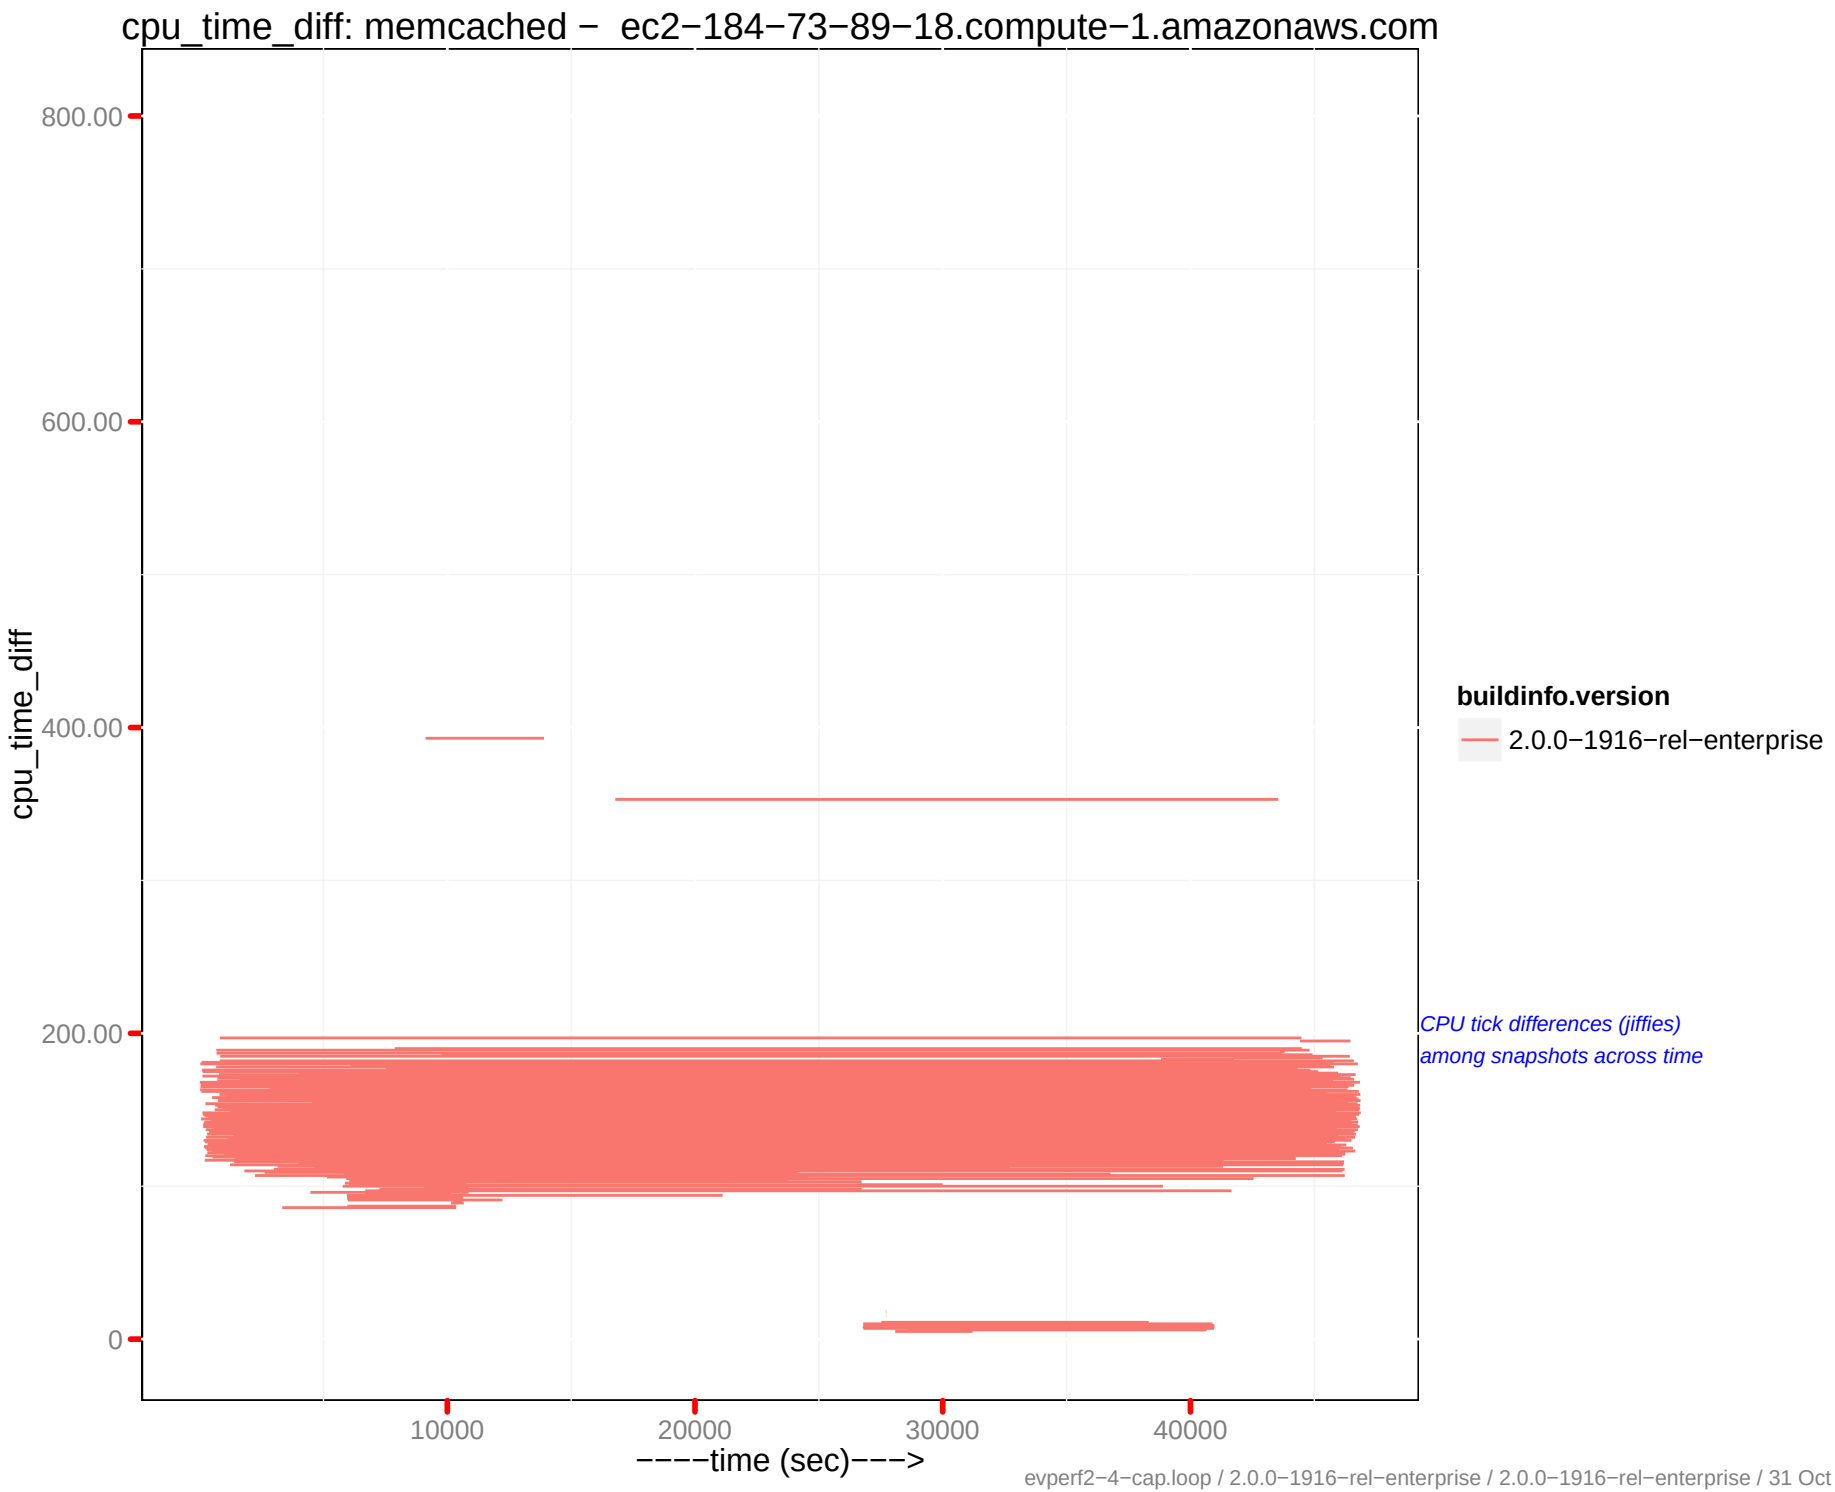
cpu_time_diff: memcached − ec2−184−73−89−18.compute−1.amazonaws.com
800.00
600.00
400.00
200.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

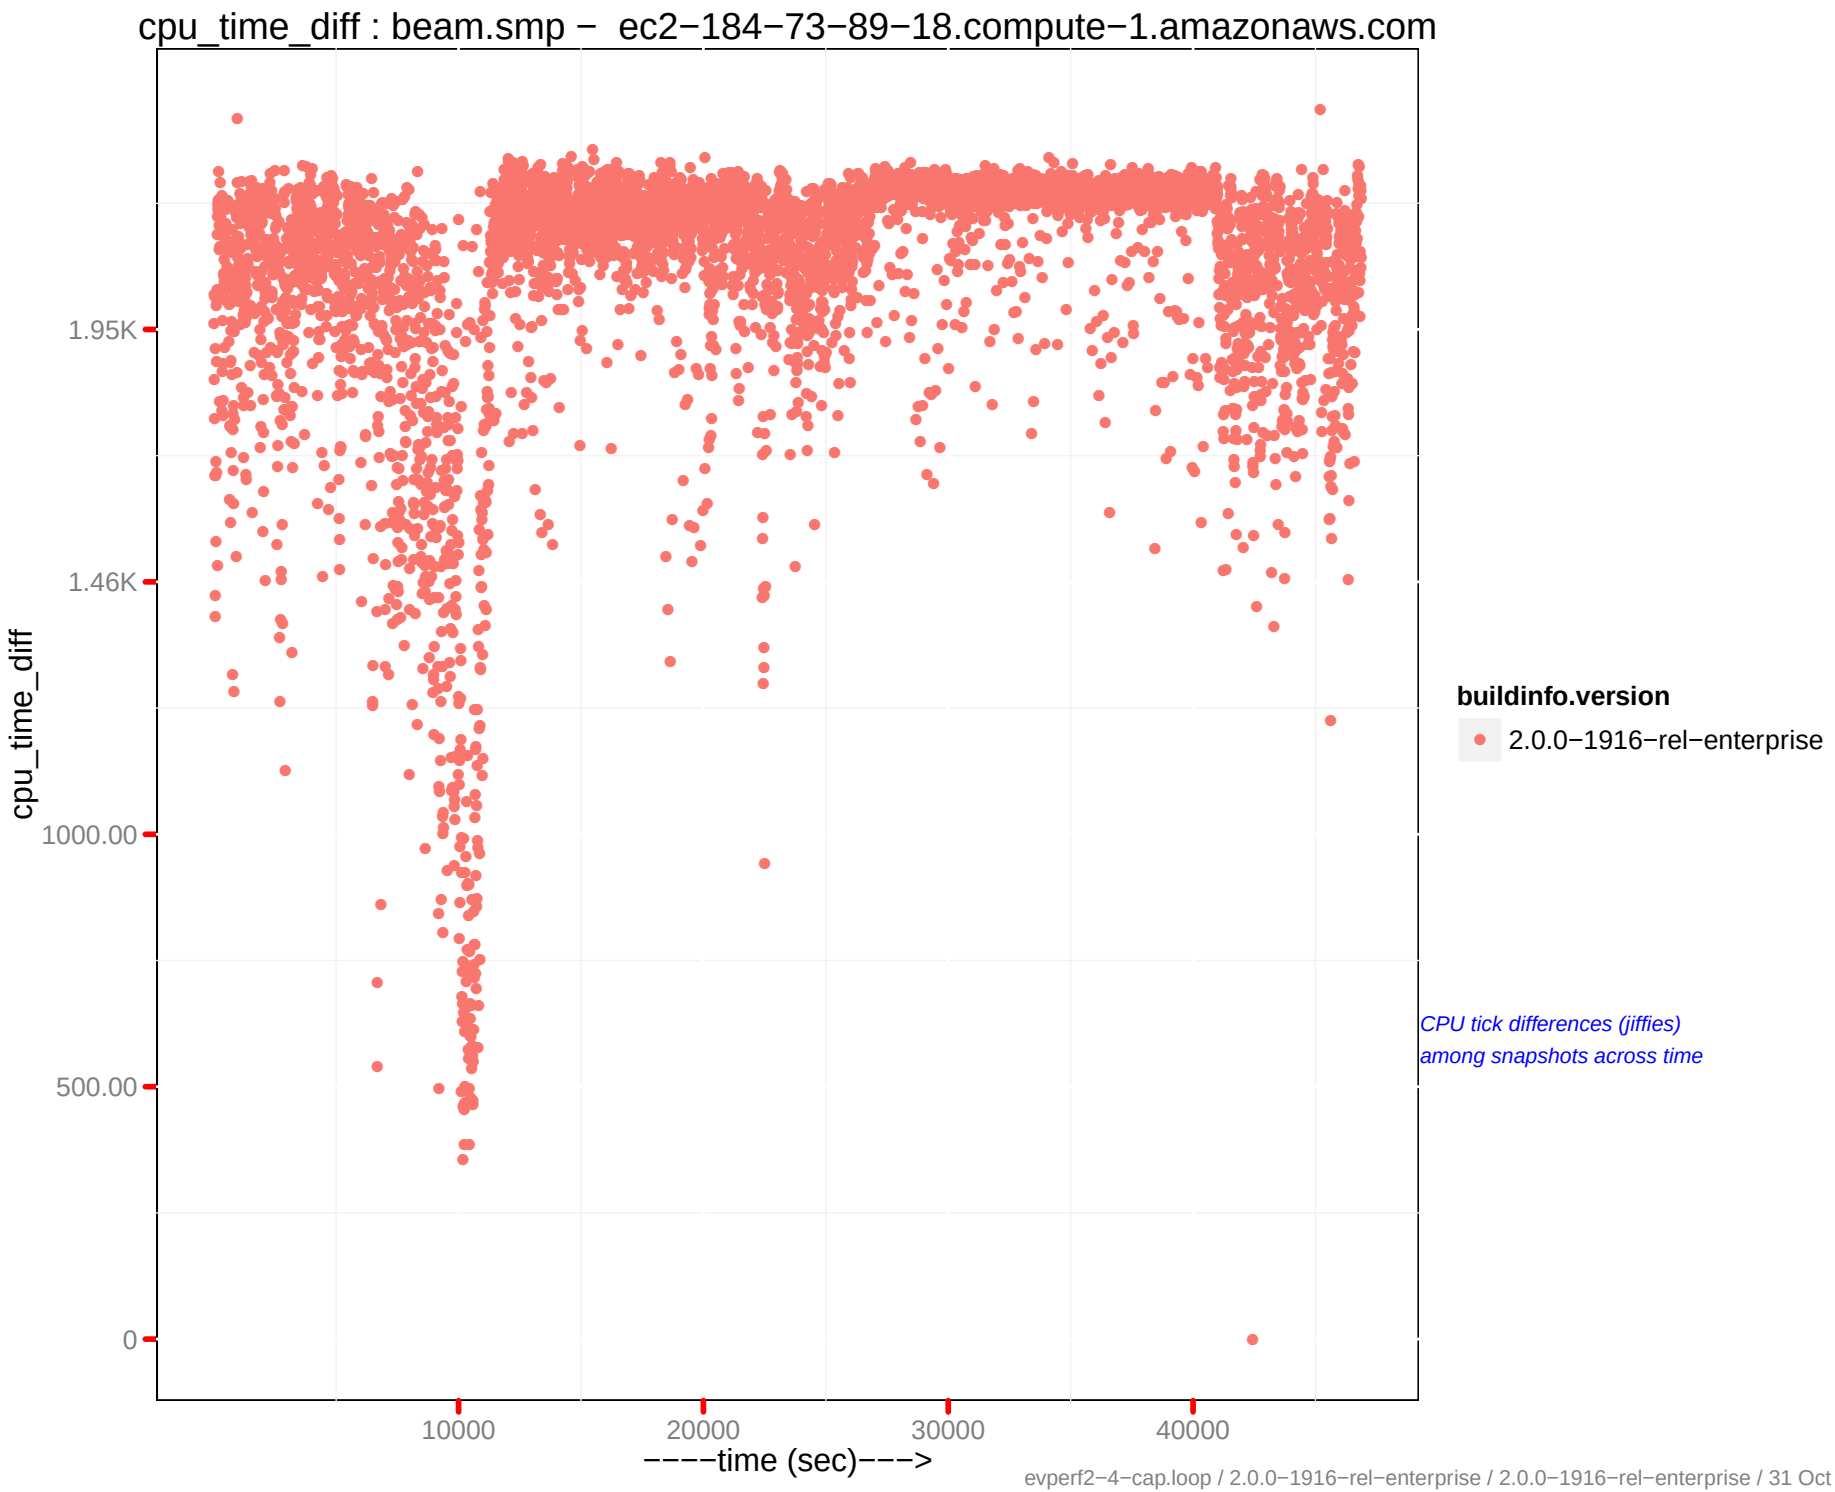
cpu_time_diff : beam.smp − ec2−184−73−89−18.compute−1.amazonaws.com
cpu_time_diff
1.95K
1.46K
1000.00
500.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

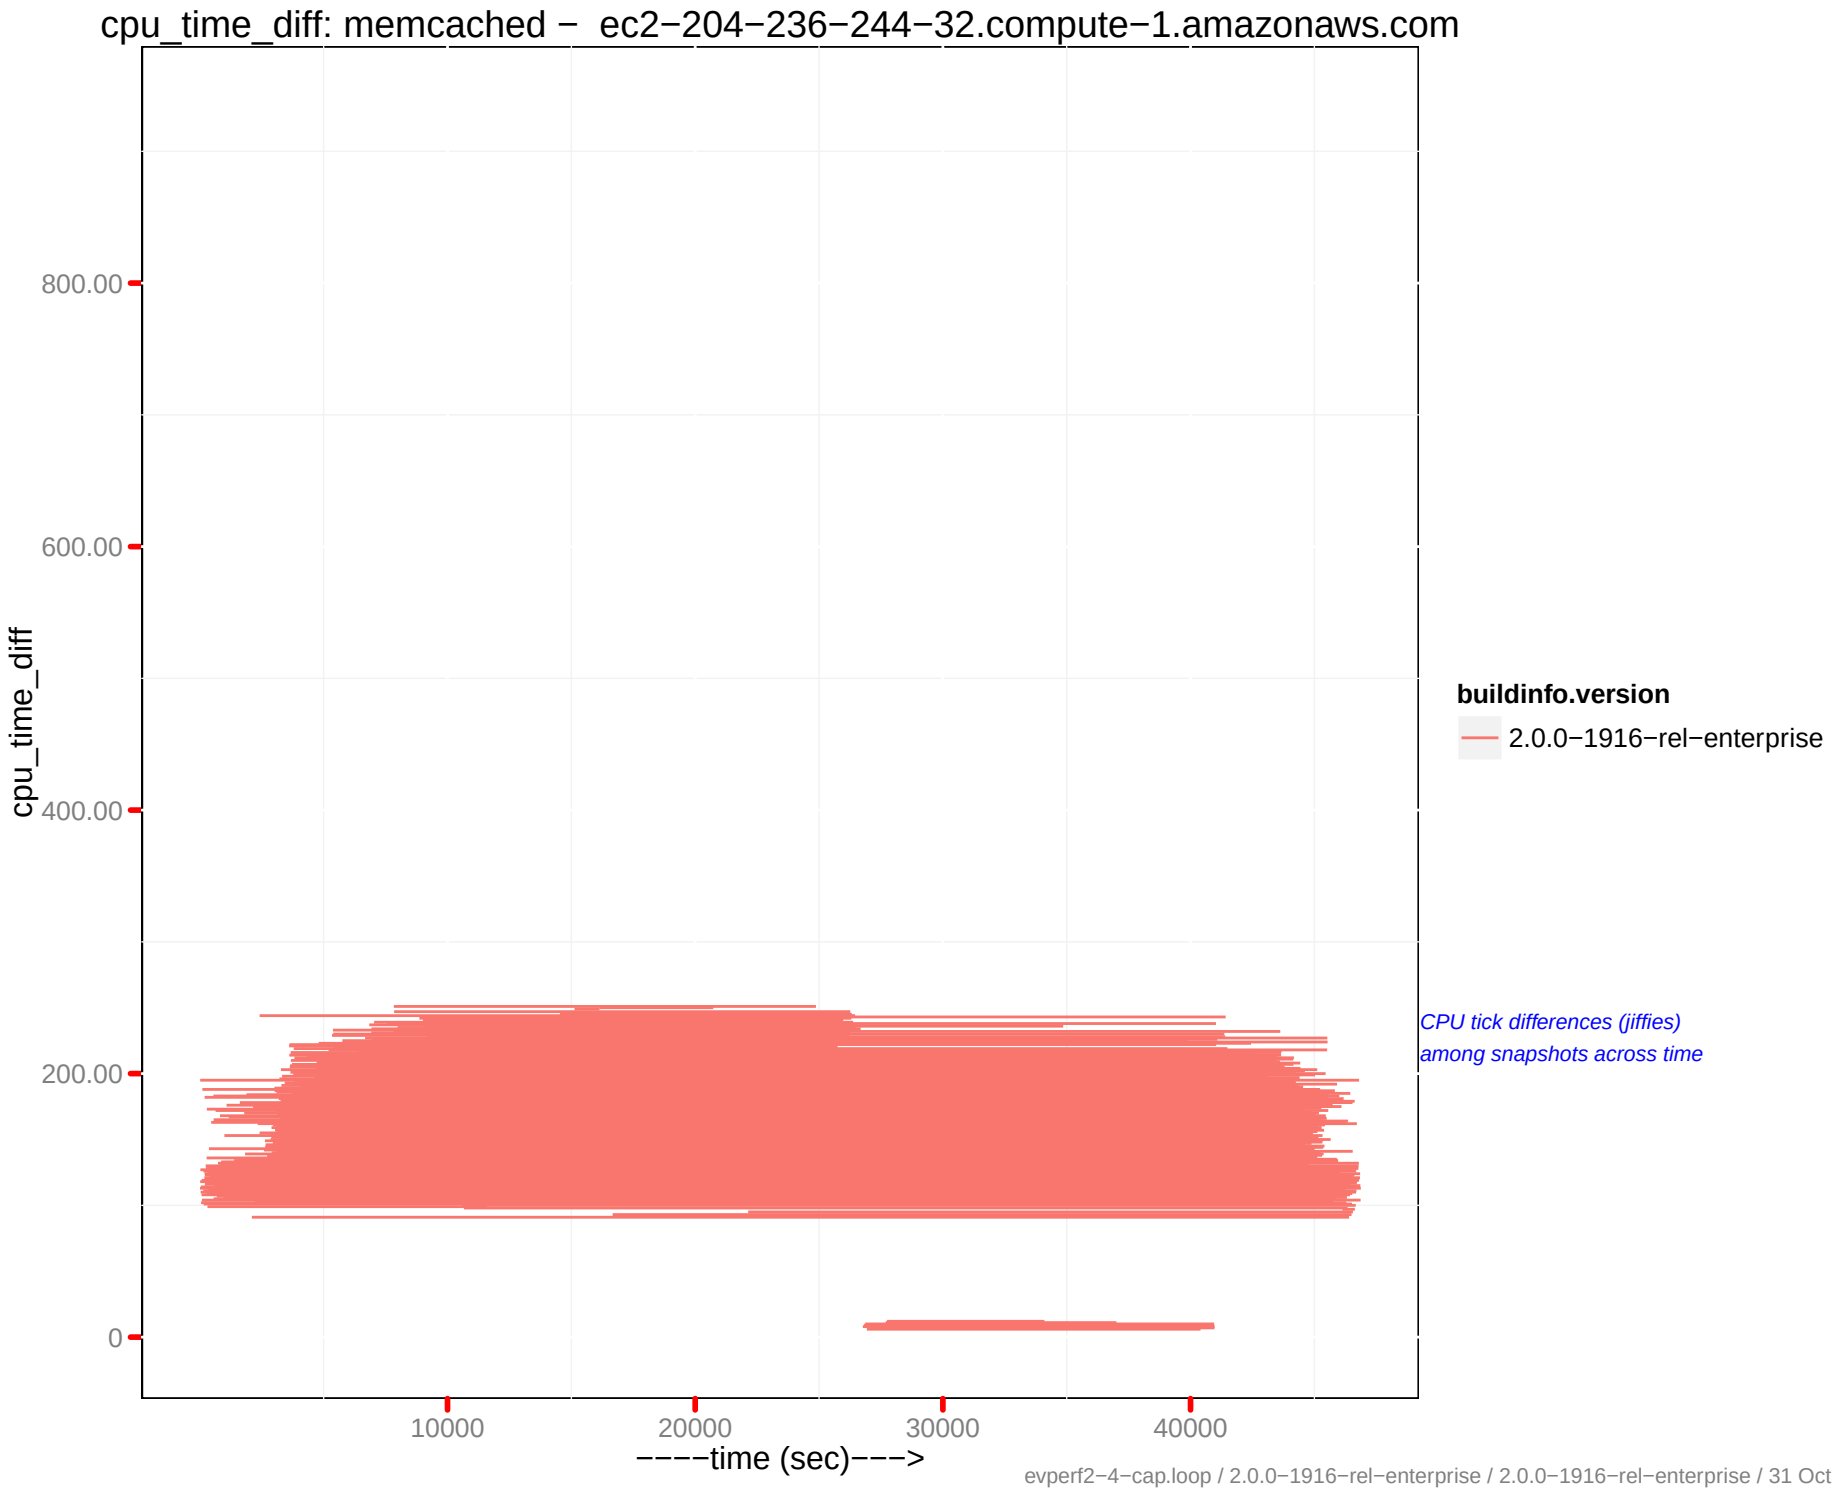
cpu_time_diff: memcached − ec2−204−236−244−32.compute−1.amazonaws.com
cpu_time_diff
800.00
600.00
400.00
200.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

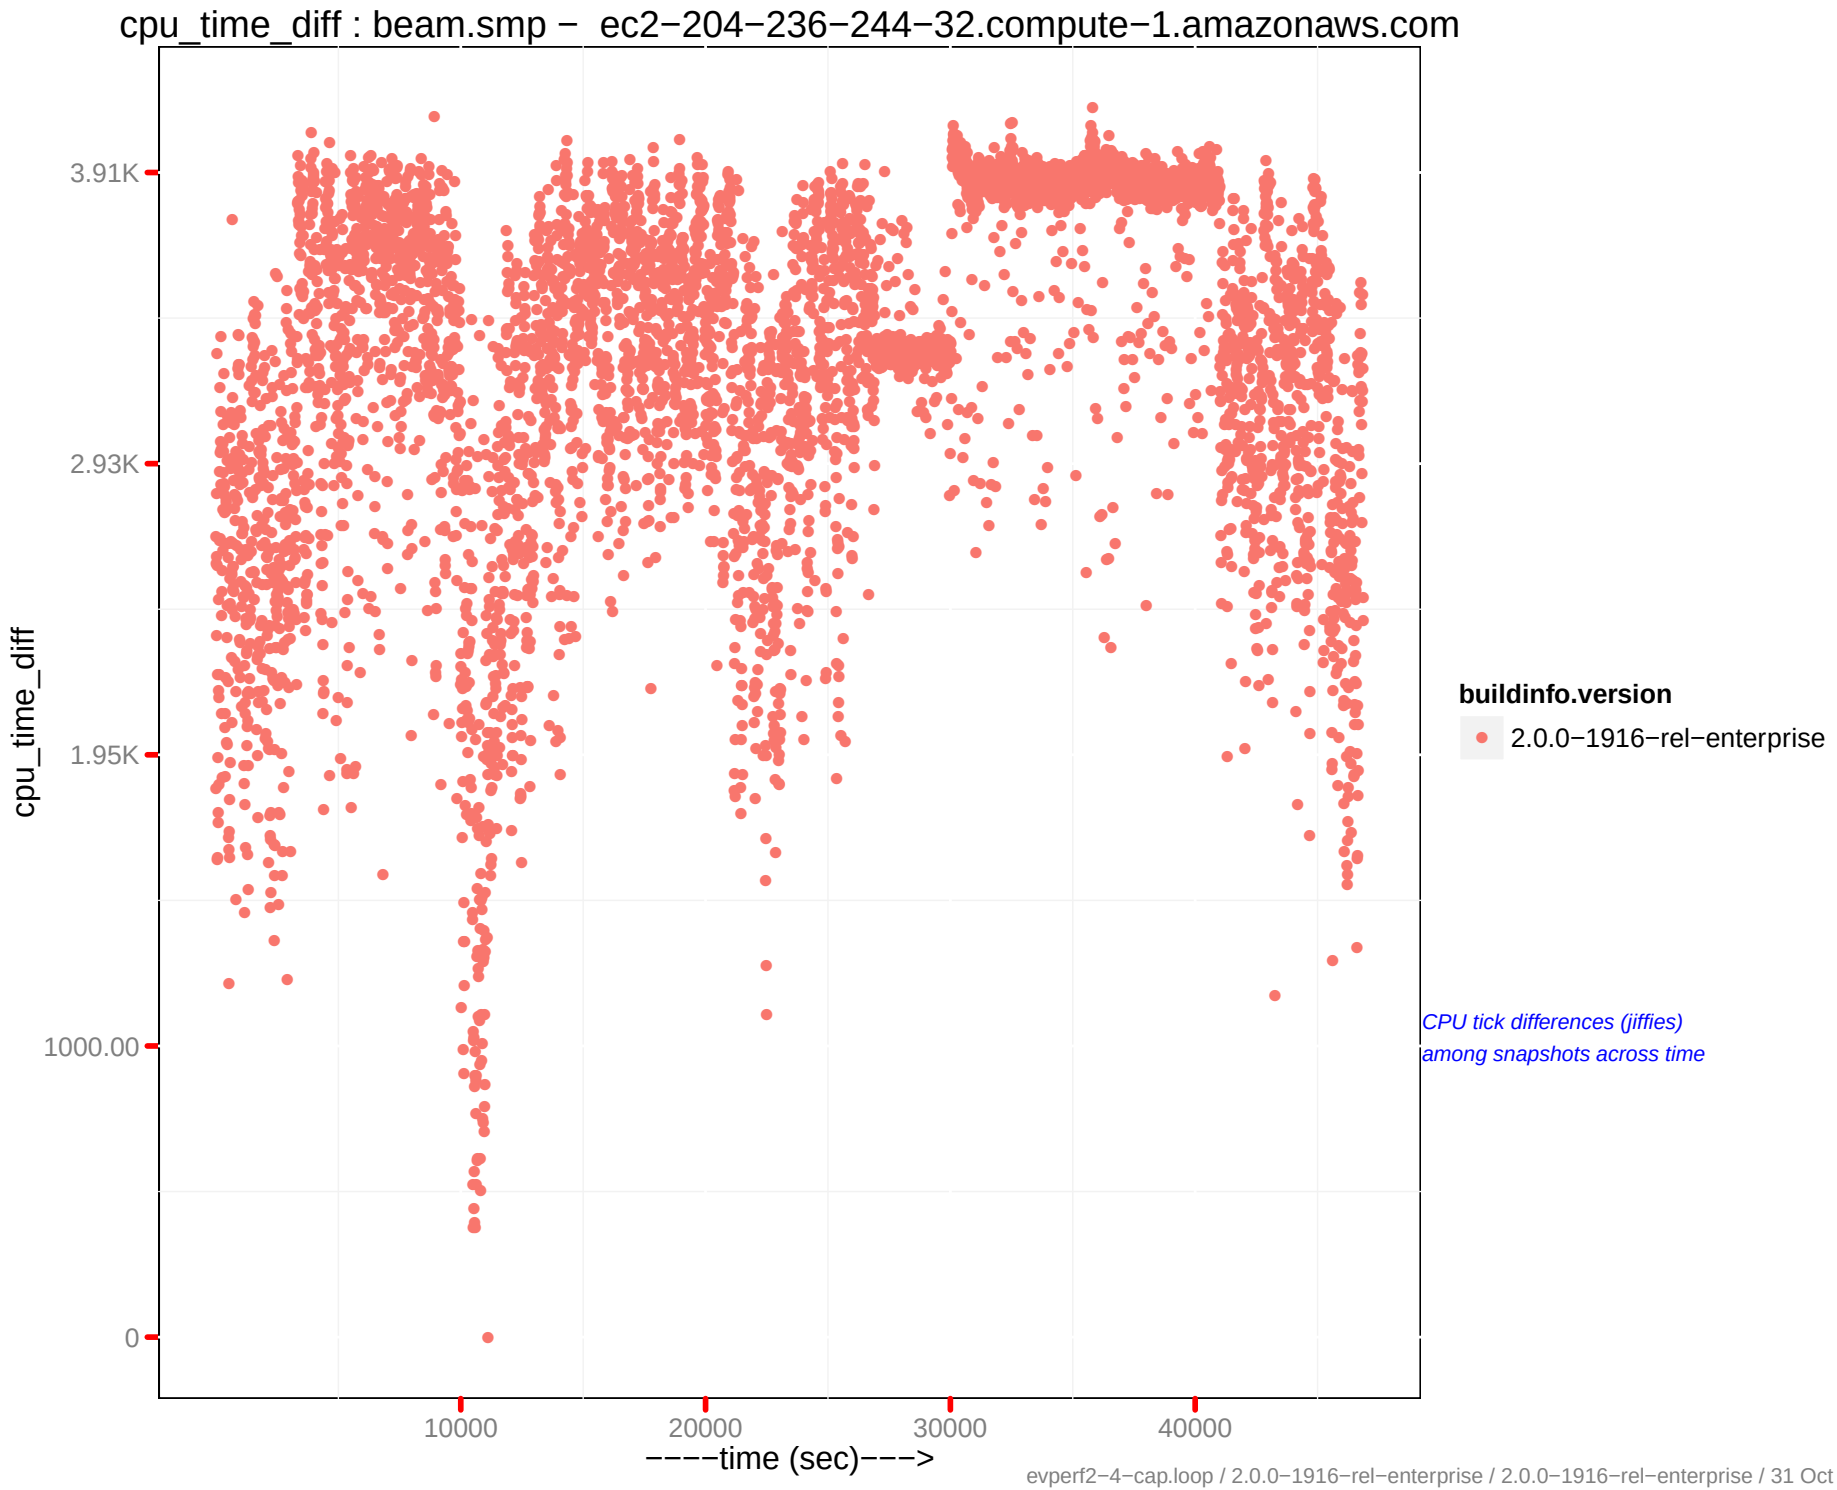
cpu_time_diff : beam.smp − ec2−204−236−244−32.compute−1.amazonaws.com
3.91K
2.93K
1.95K
1000.00
0
cpu_time_diff
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

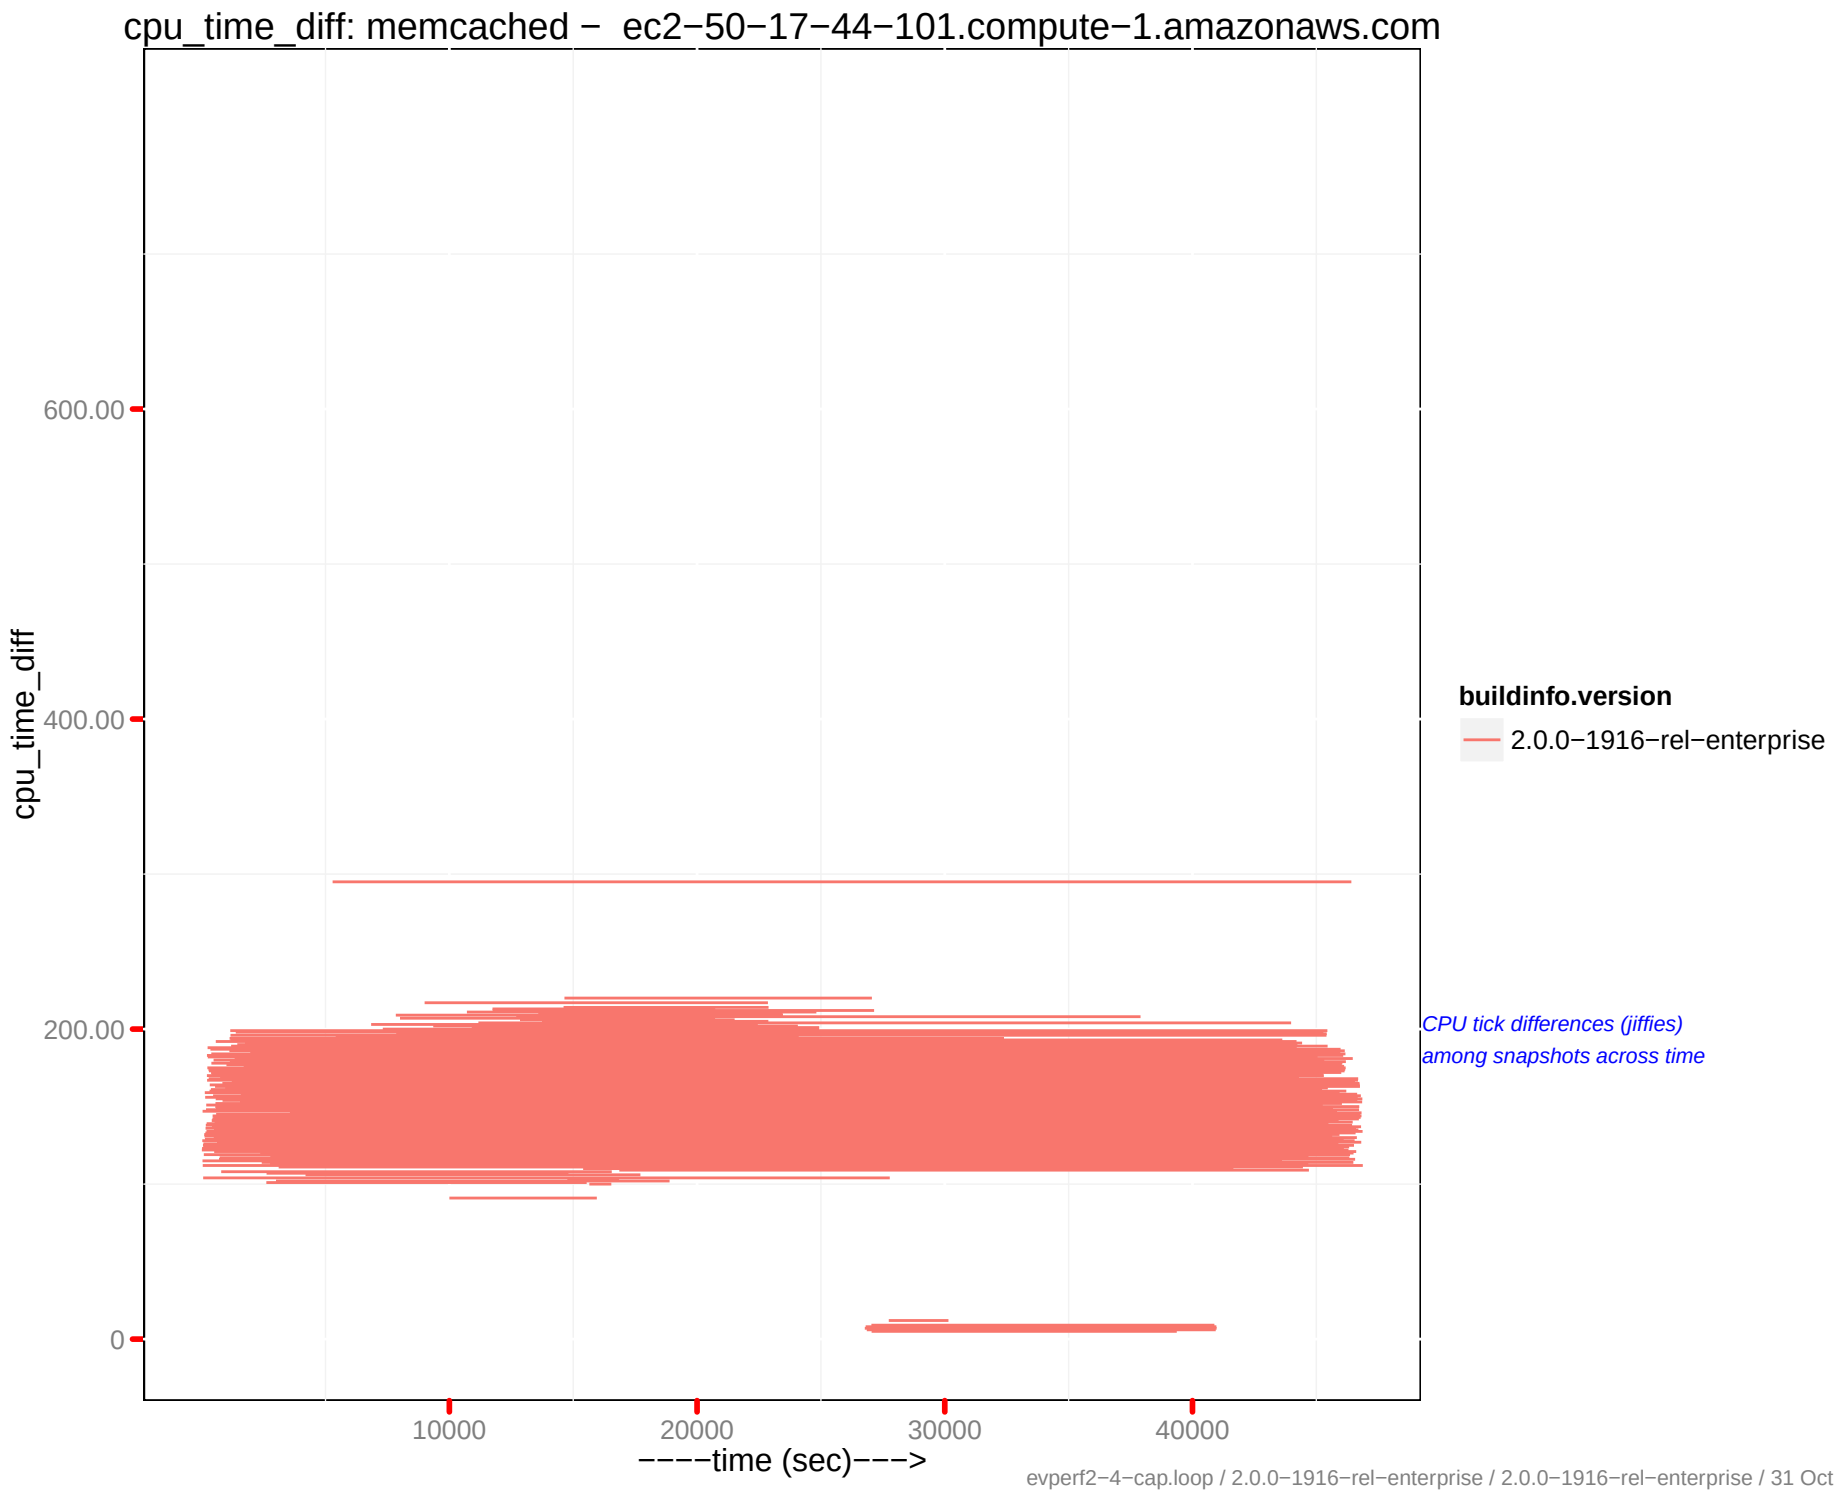
cpu_time_diff: memcached − ec2−50−17−44−101.compute−1.amazonaws.com
cpu_time_diff
600.00
400.00
200.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012


cpu_time_diff : beam.smp − ec2−50−17−44−101.compute−1.amazonaws.com
cpu_time_diff
5.86K
4.88K
3.91K
2.93K
1.95K
1000.00
0
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

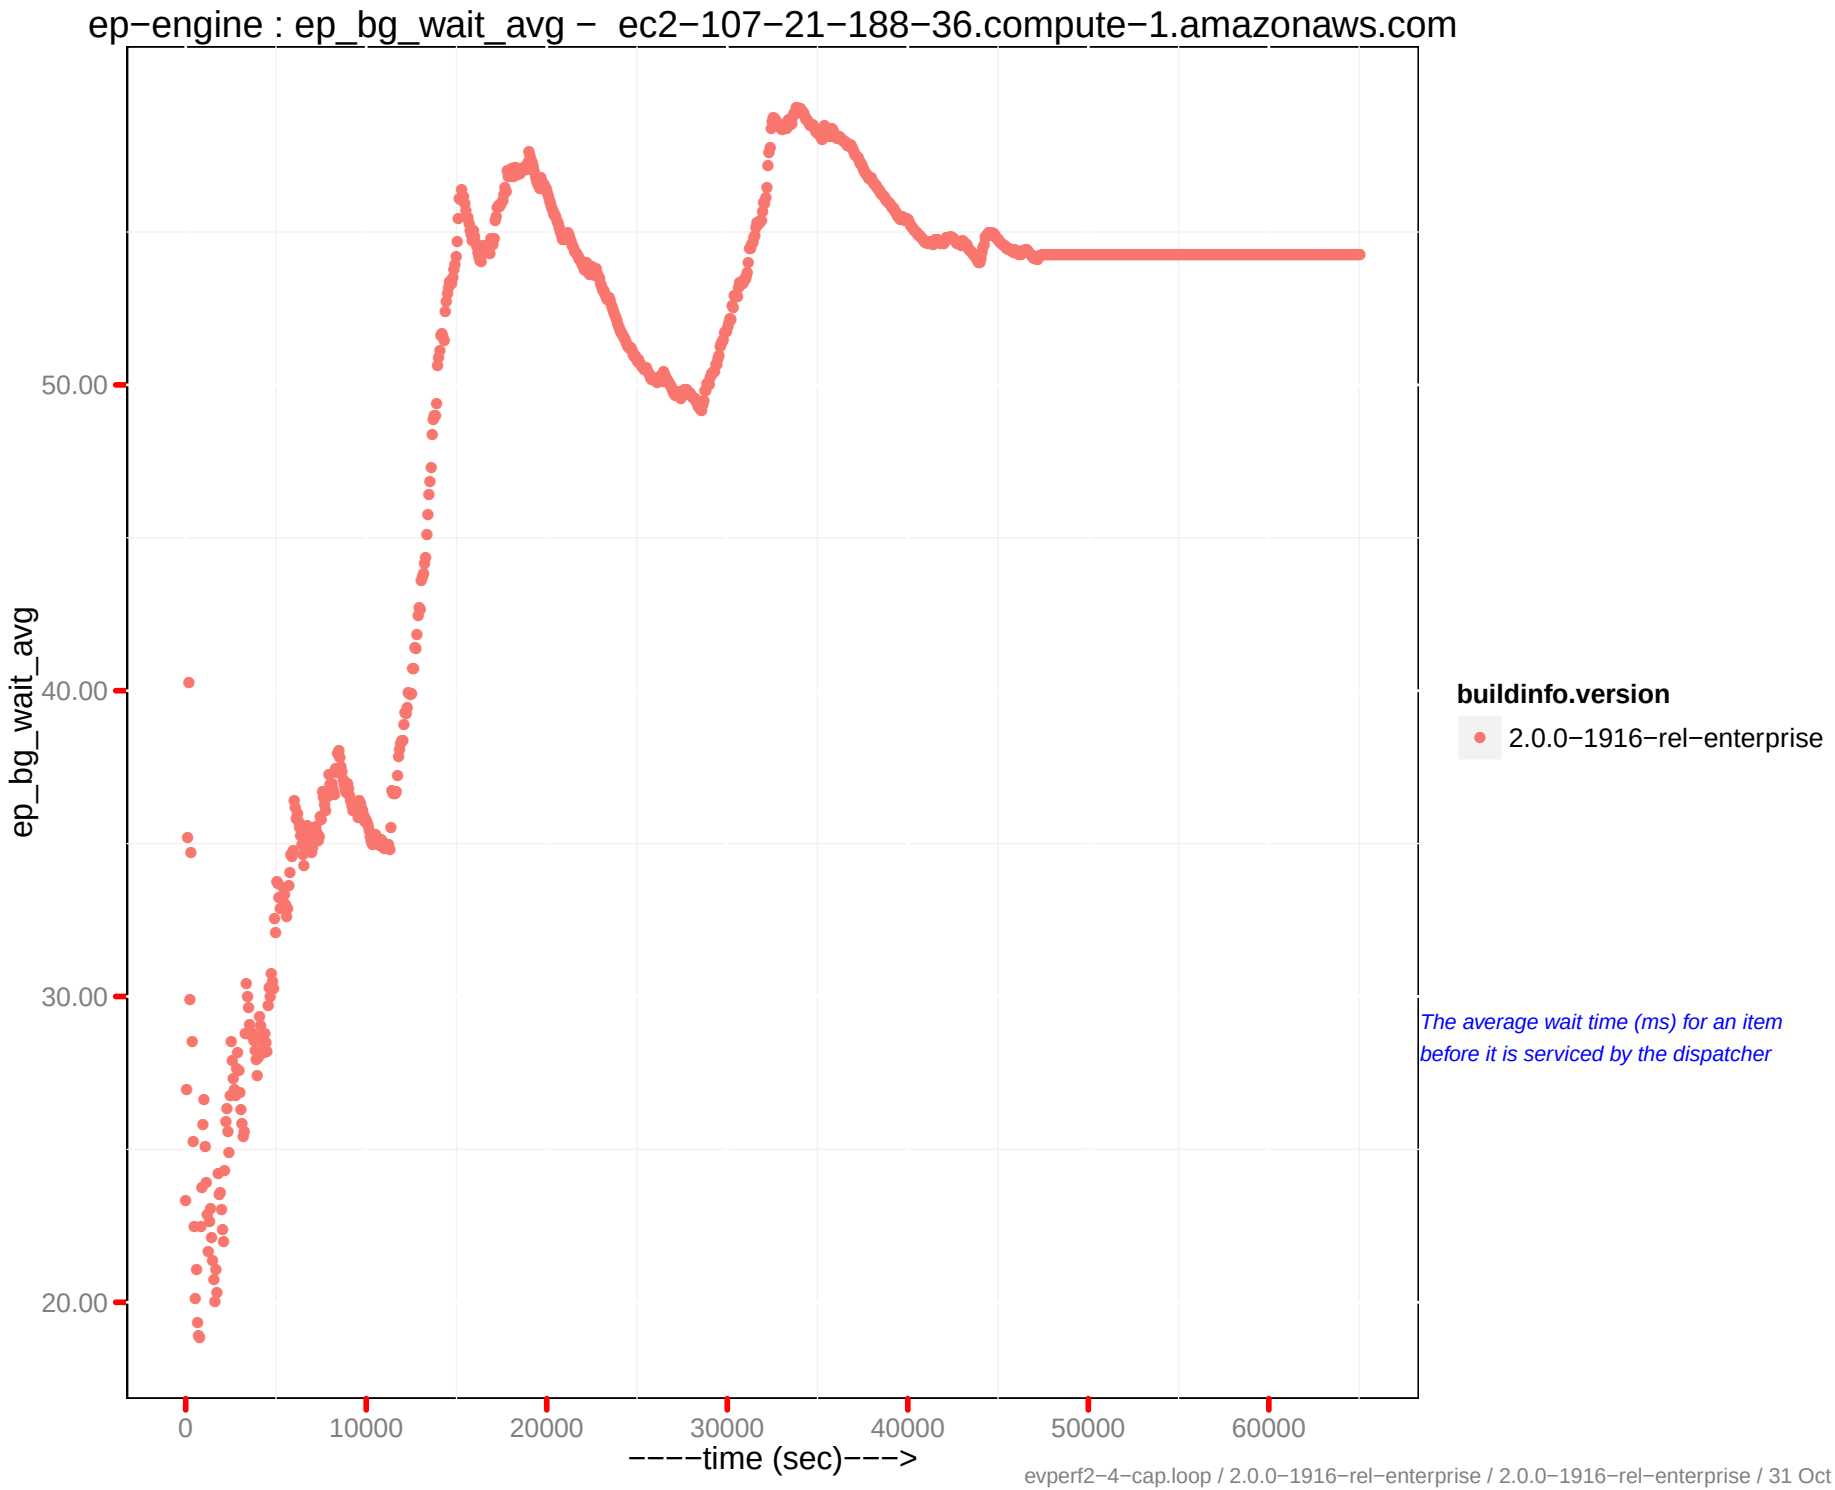
ep−engine : ep_bg_wait_avg − ec2−107−21−188−36.compute−1.amazonaws.com
ep_bg_wait_avg
50.00
40.00
30.00
20.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

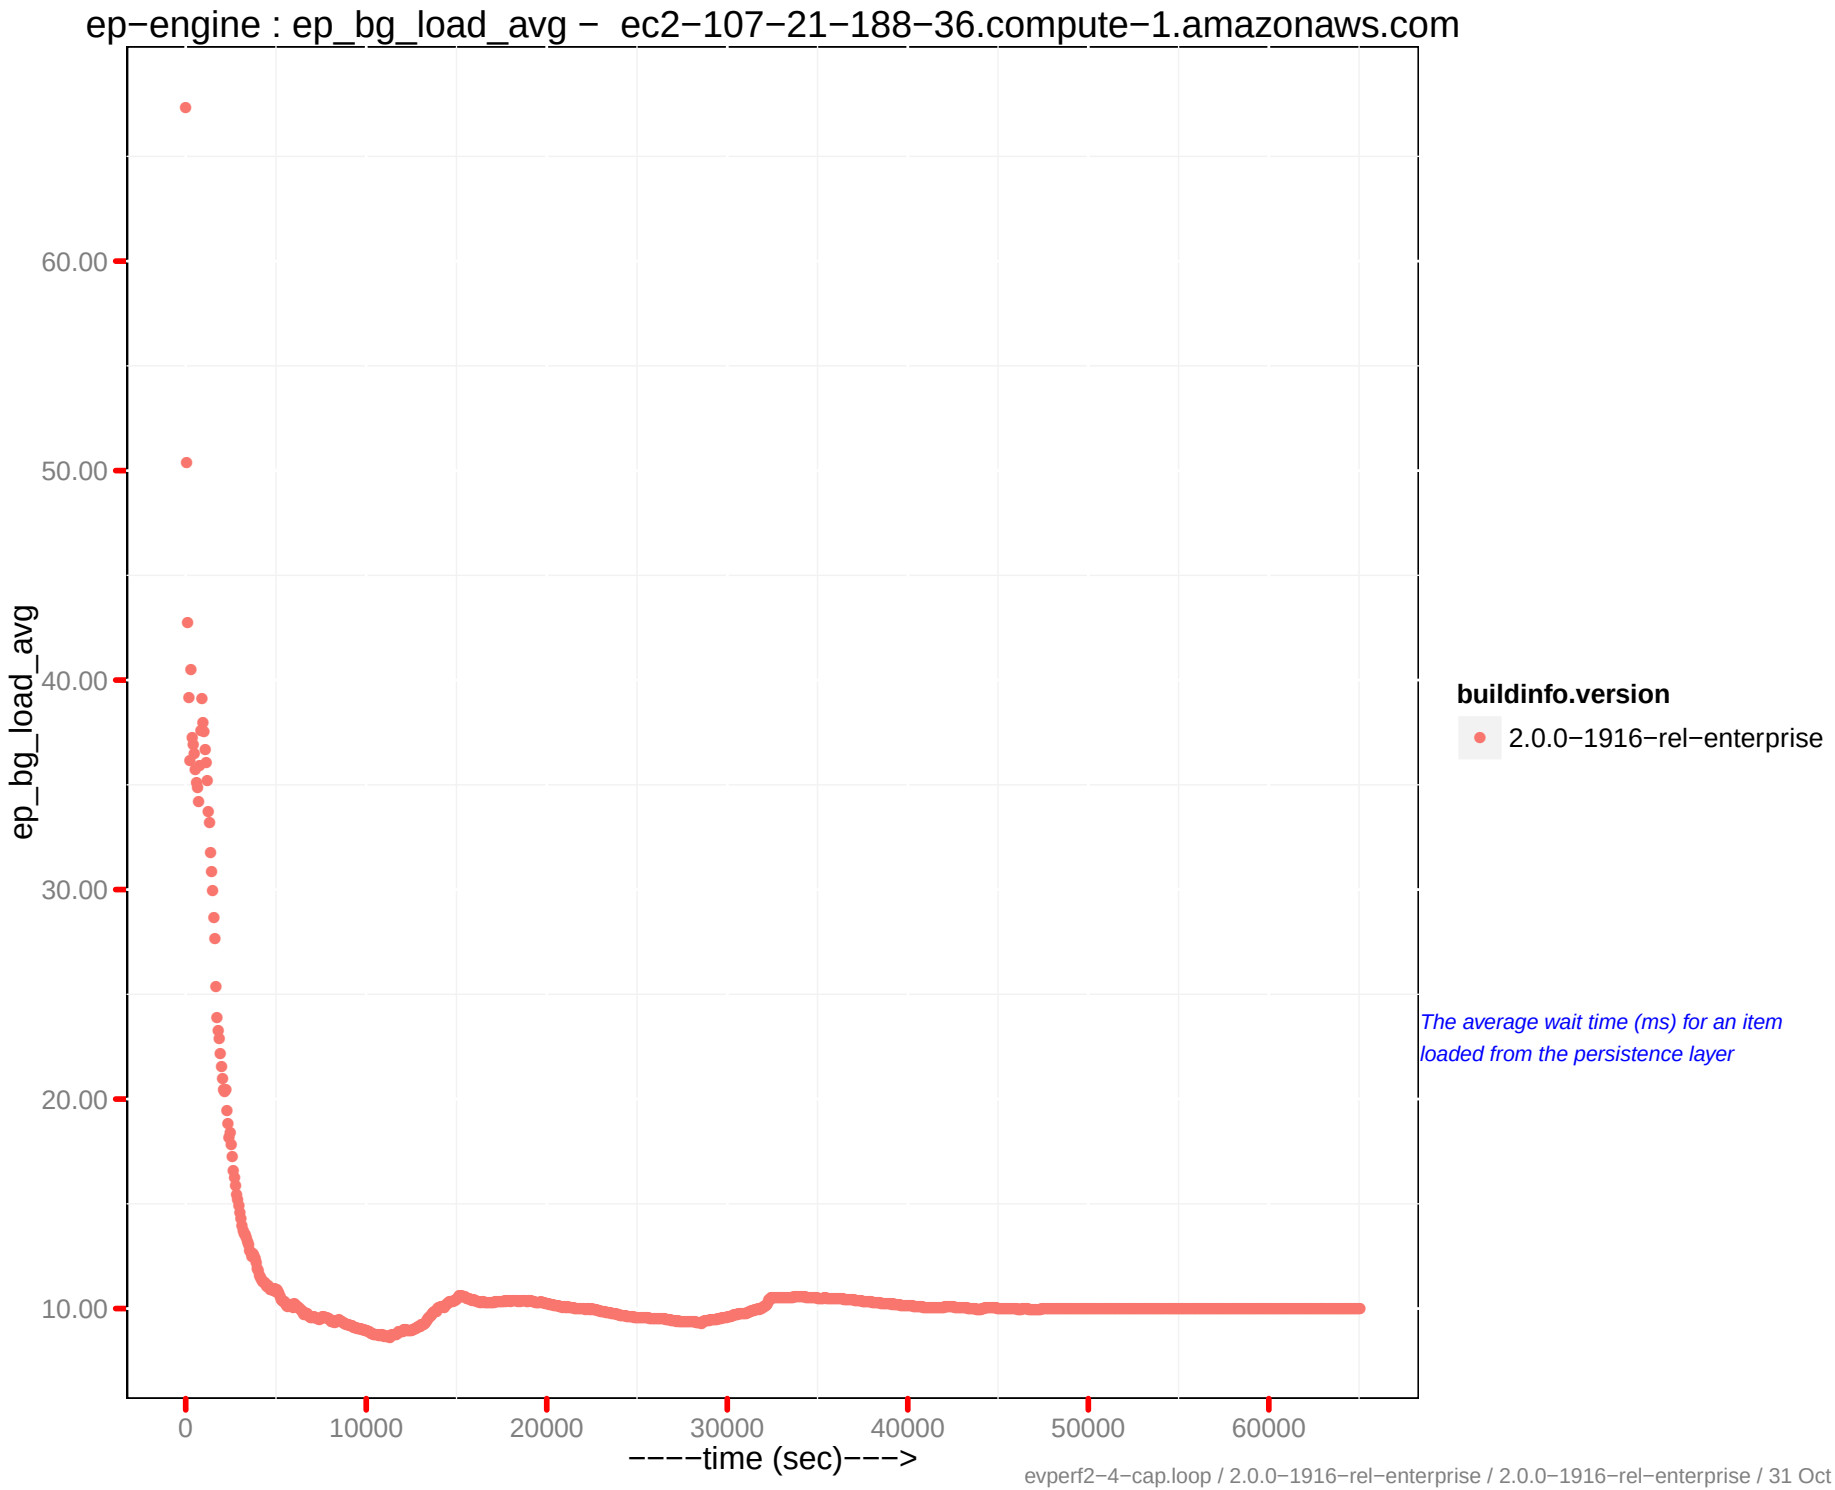
ep−engine : ep_bg_load_avg − ec2−107−21−188−36.compute−1.amazonaws.com
ep_bg_load_avg
60.00
50.00
40.00
30.00
20.00
10.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

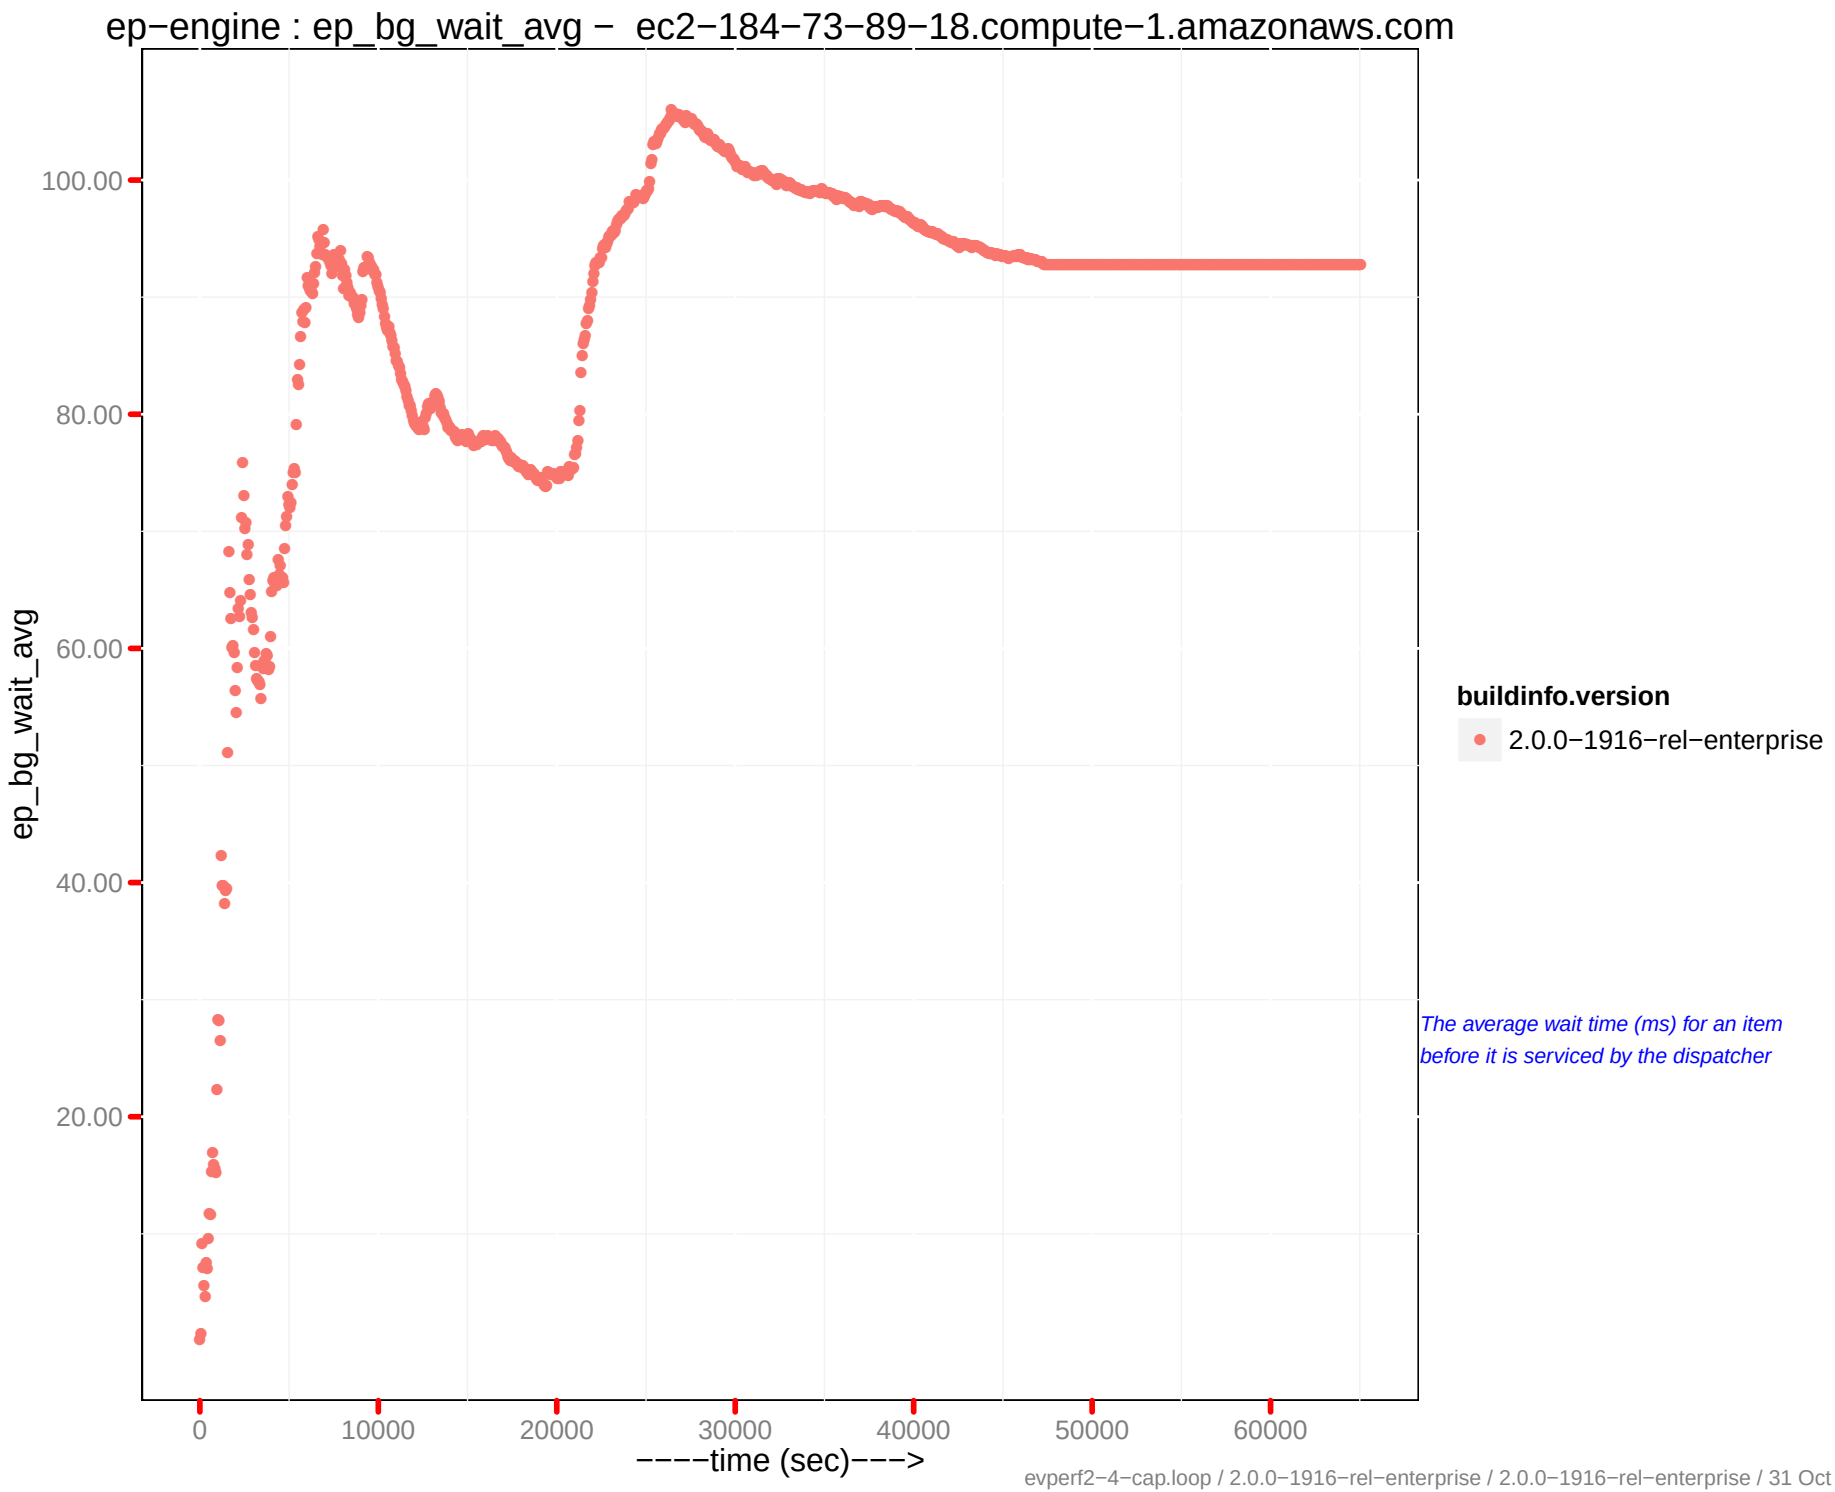
ep−engine : ep_bg_wait_avg − ec2−184−73−89−18.compute−1.amazonaws.com
ep_bg_wait_avg
100.00
80.00
60.00
40.00
20.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

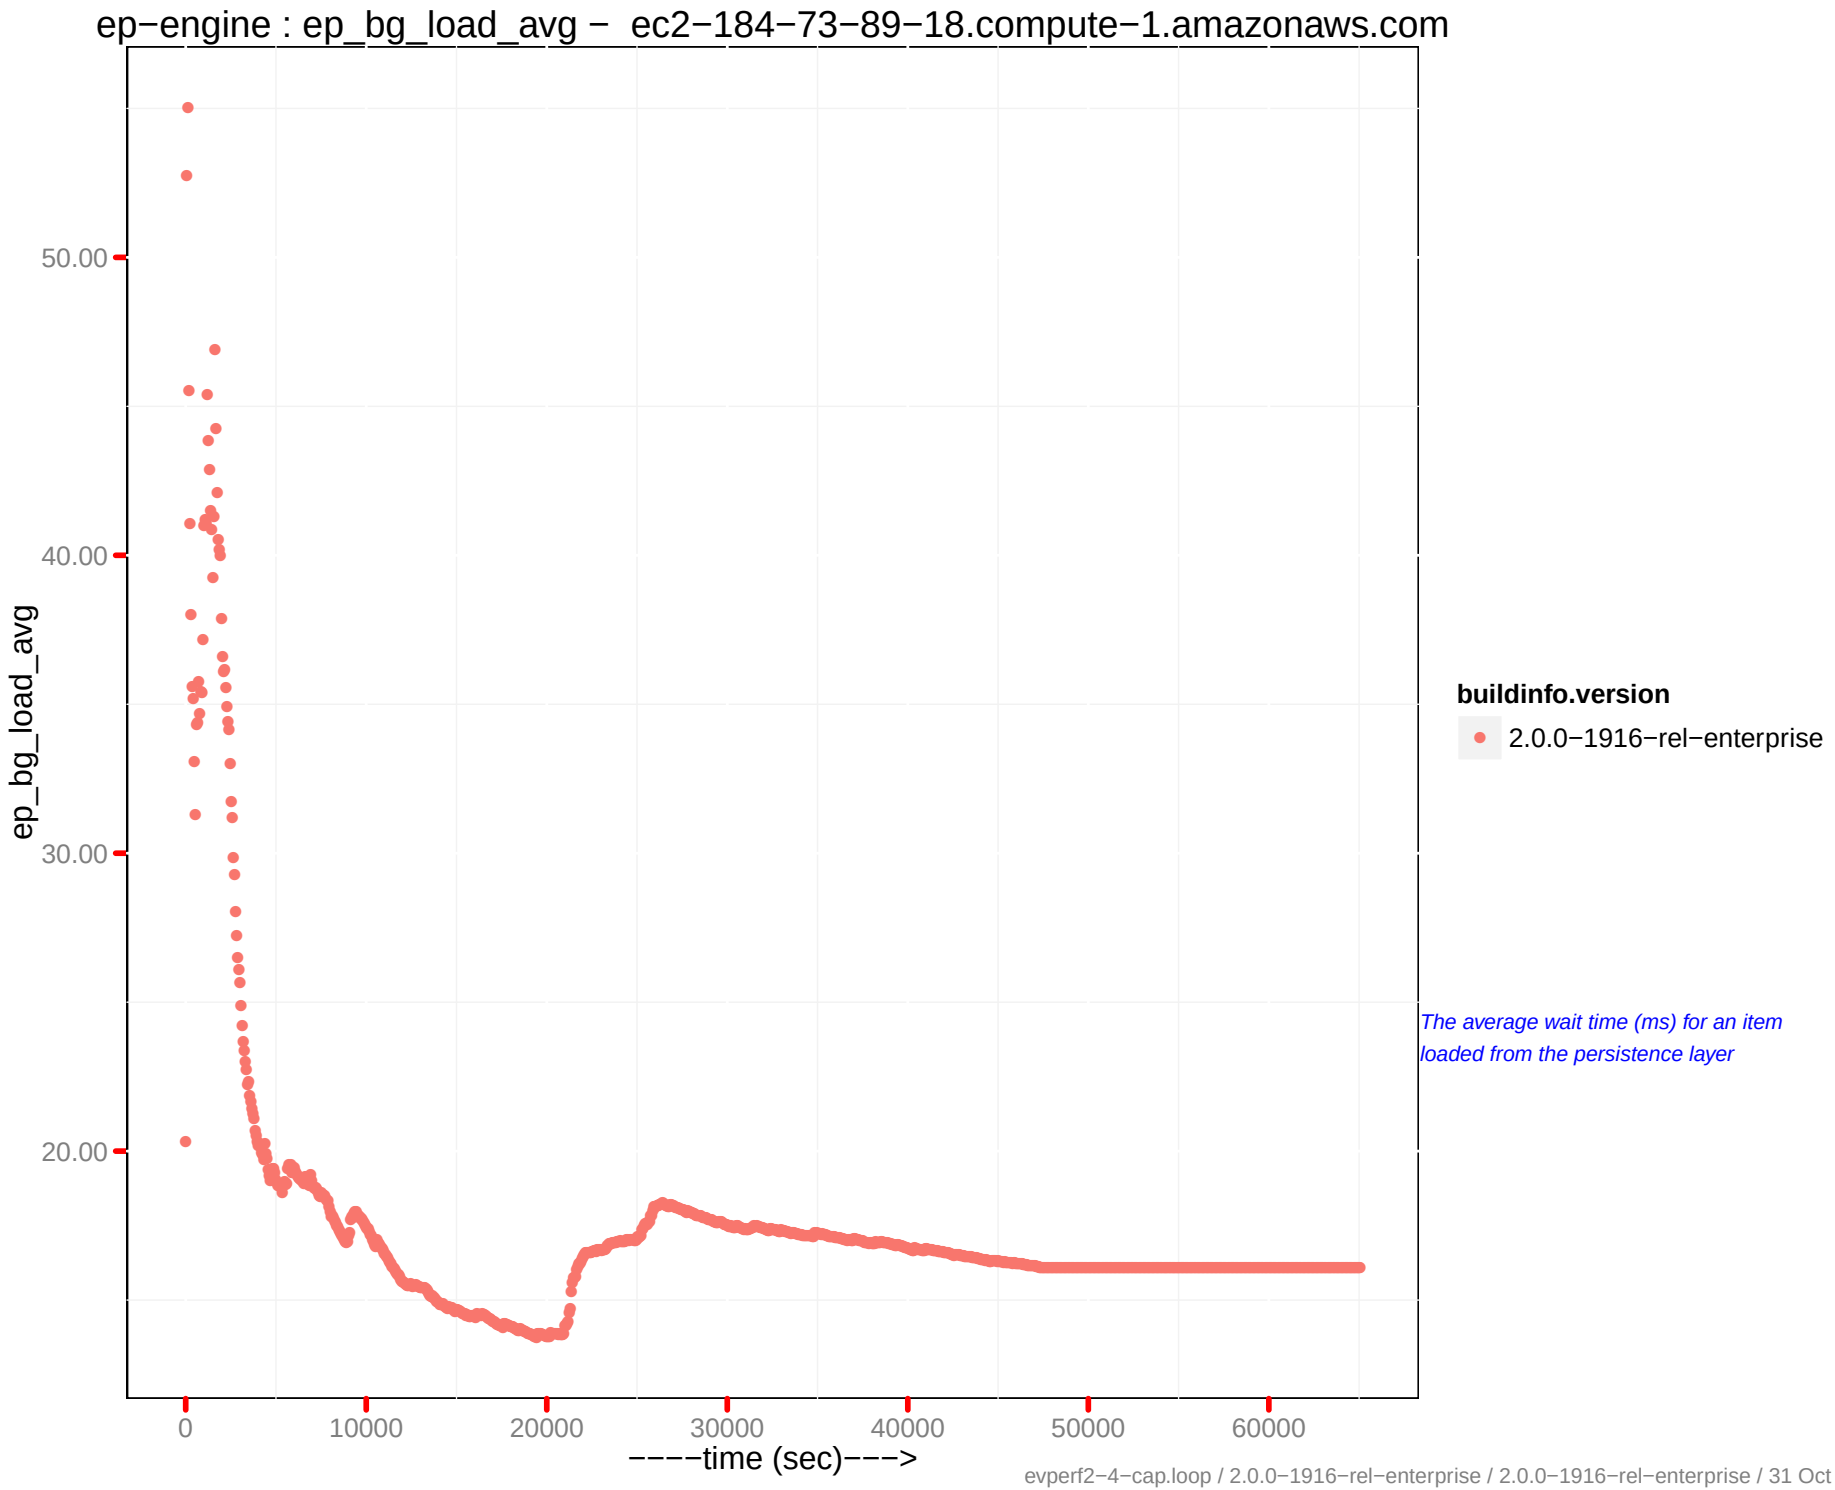
ep−engine : ep_bg_load_avg − ec2−184−73−89−18.compute−1.amazonaws.com
50.00
40.00
30.00
20.00
ep_bg_load_avg
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

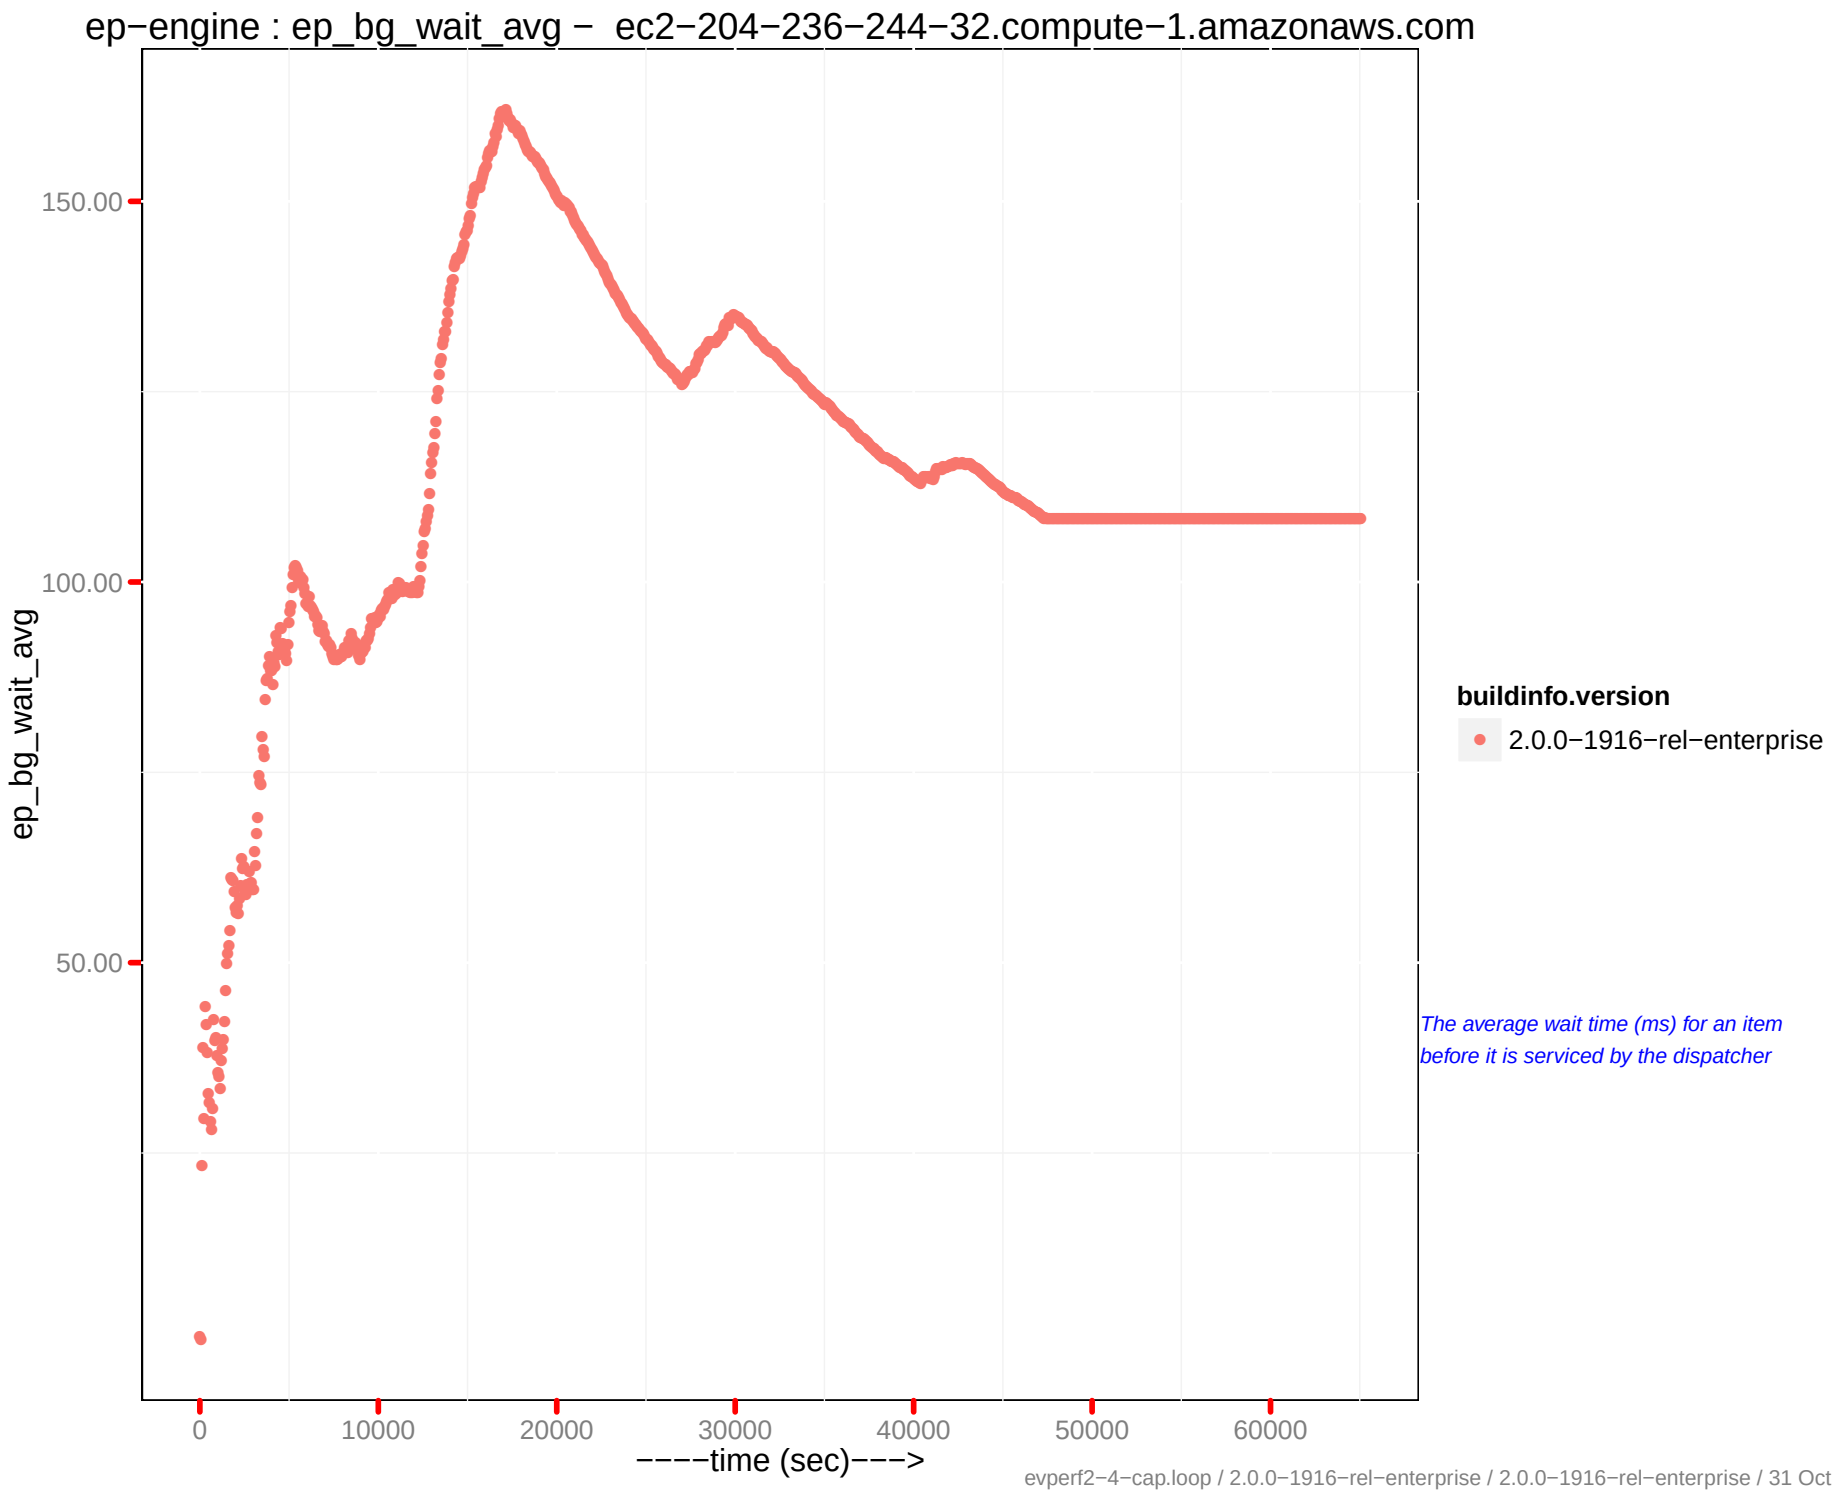
ep−engine : ep_bg_wait_avg − ec2−204−236−244−32.compute−1.amazonaws.com
150.00
100.00
50.00
ep_bg_wait_avg
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

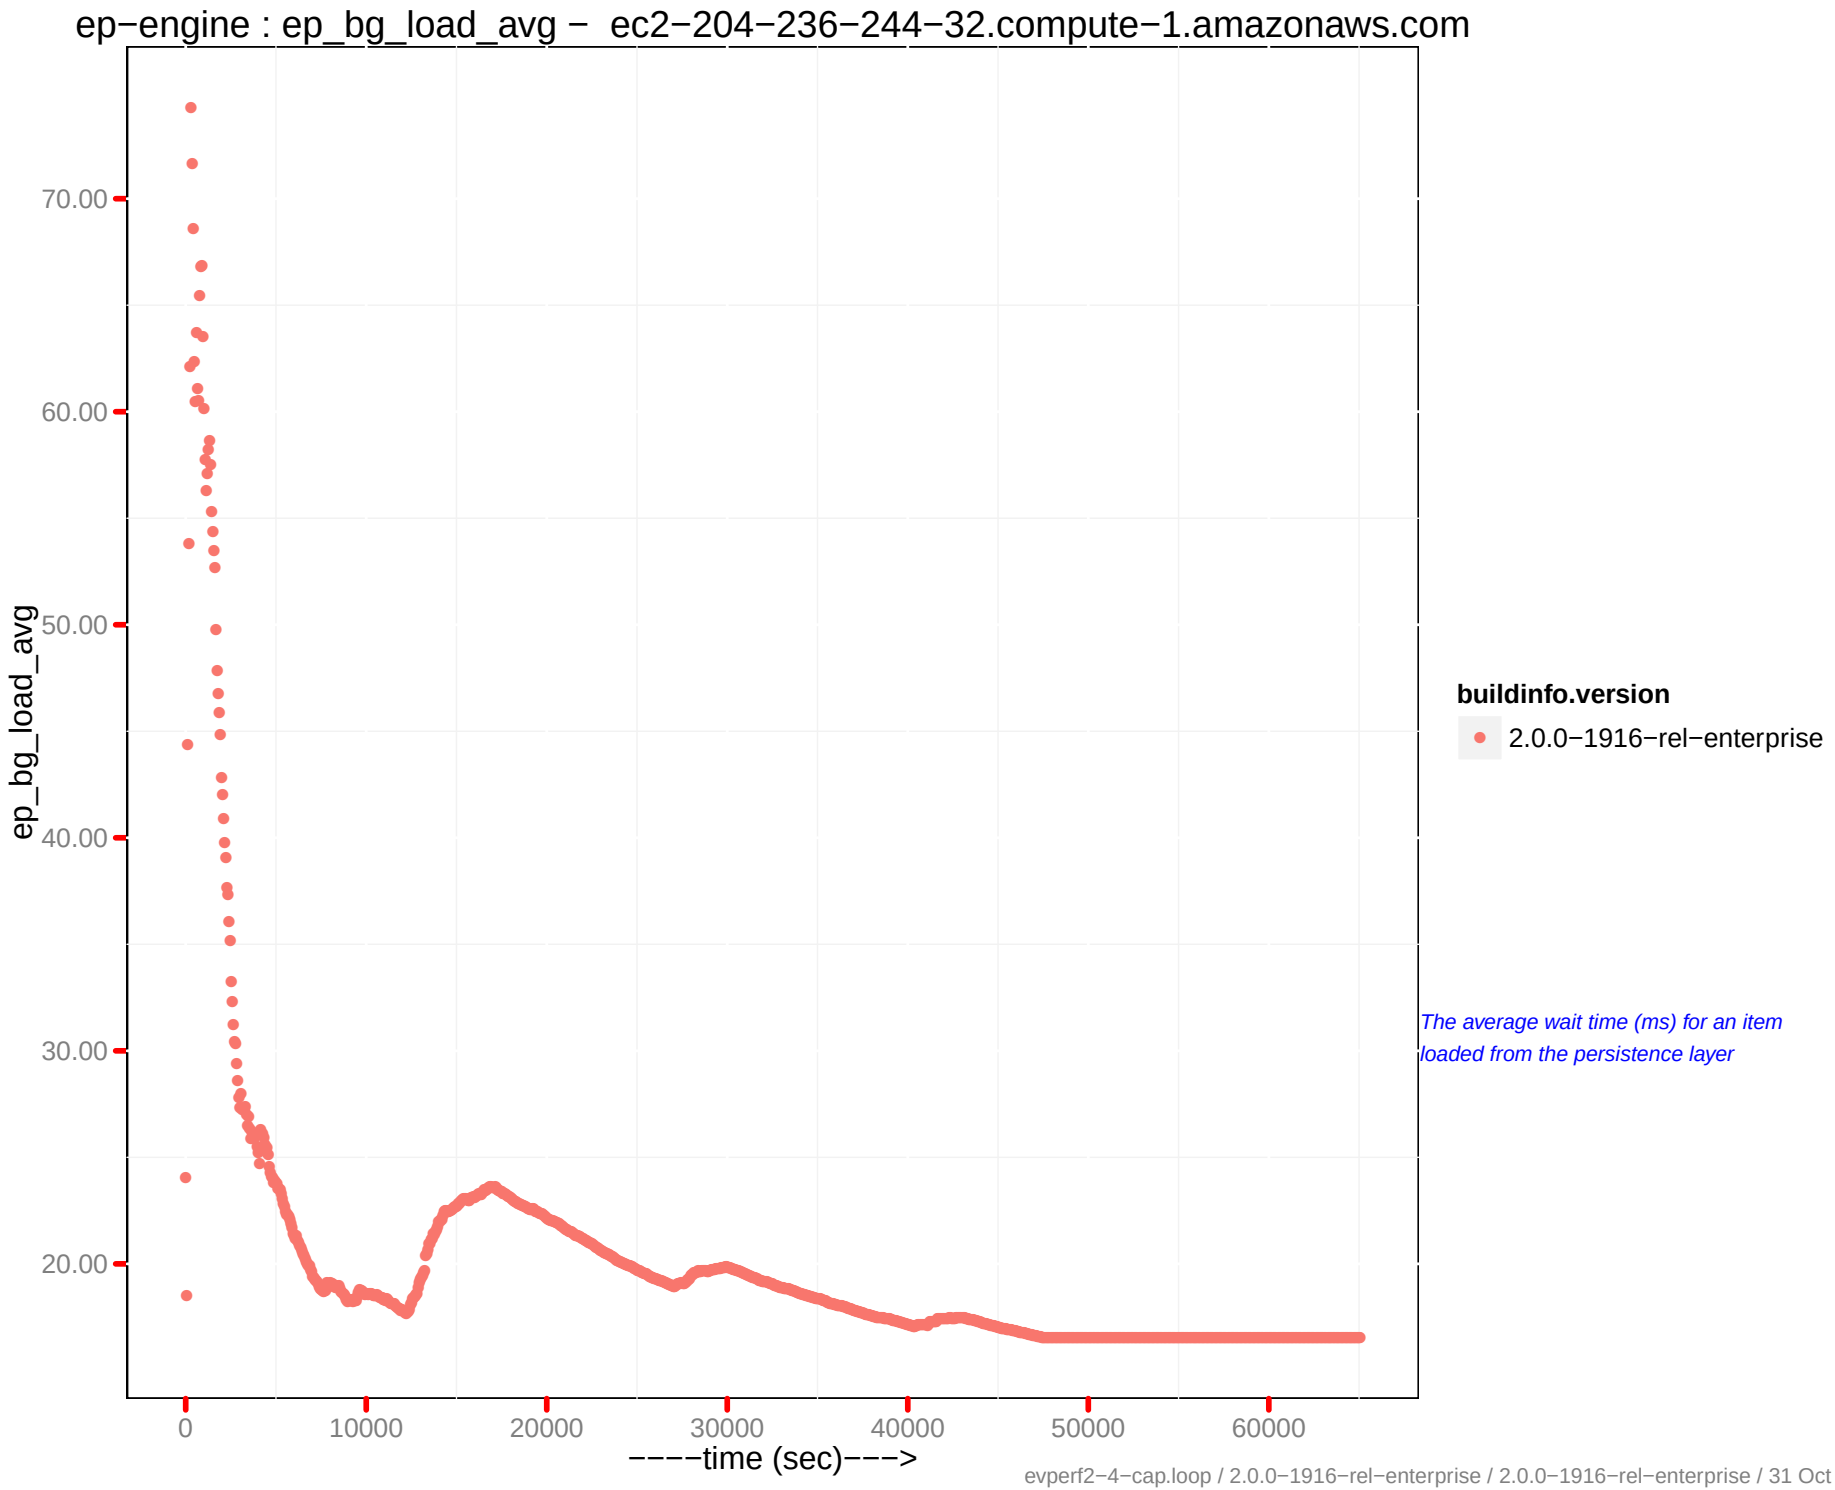
ep−engine : ep_bg_load_avg − ec2−204−236−244−32.compute−1.amazonaws.com
ep_bg_load_avg
70.00
60.00
50.00
40.00
30.00
20.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

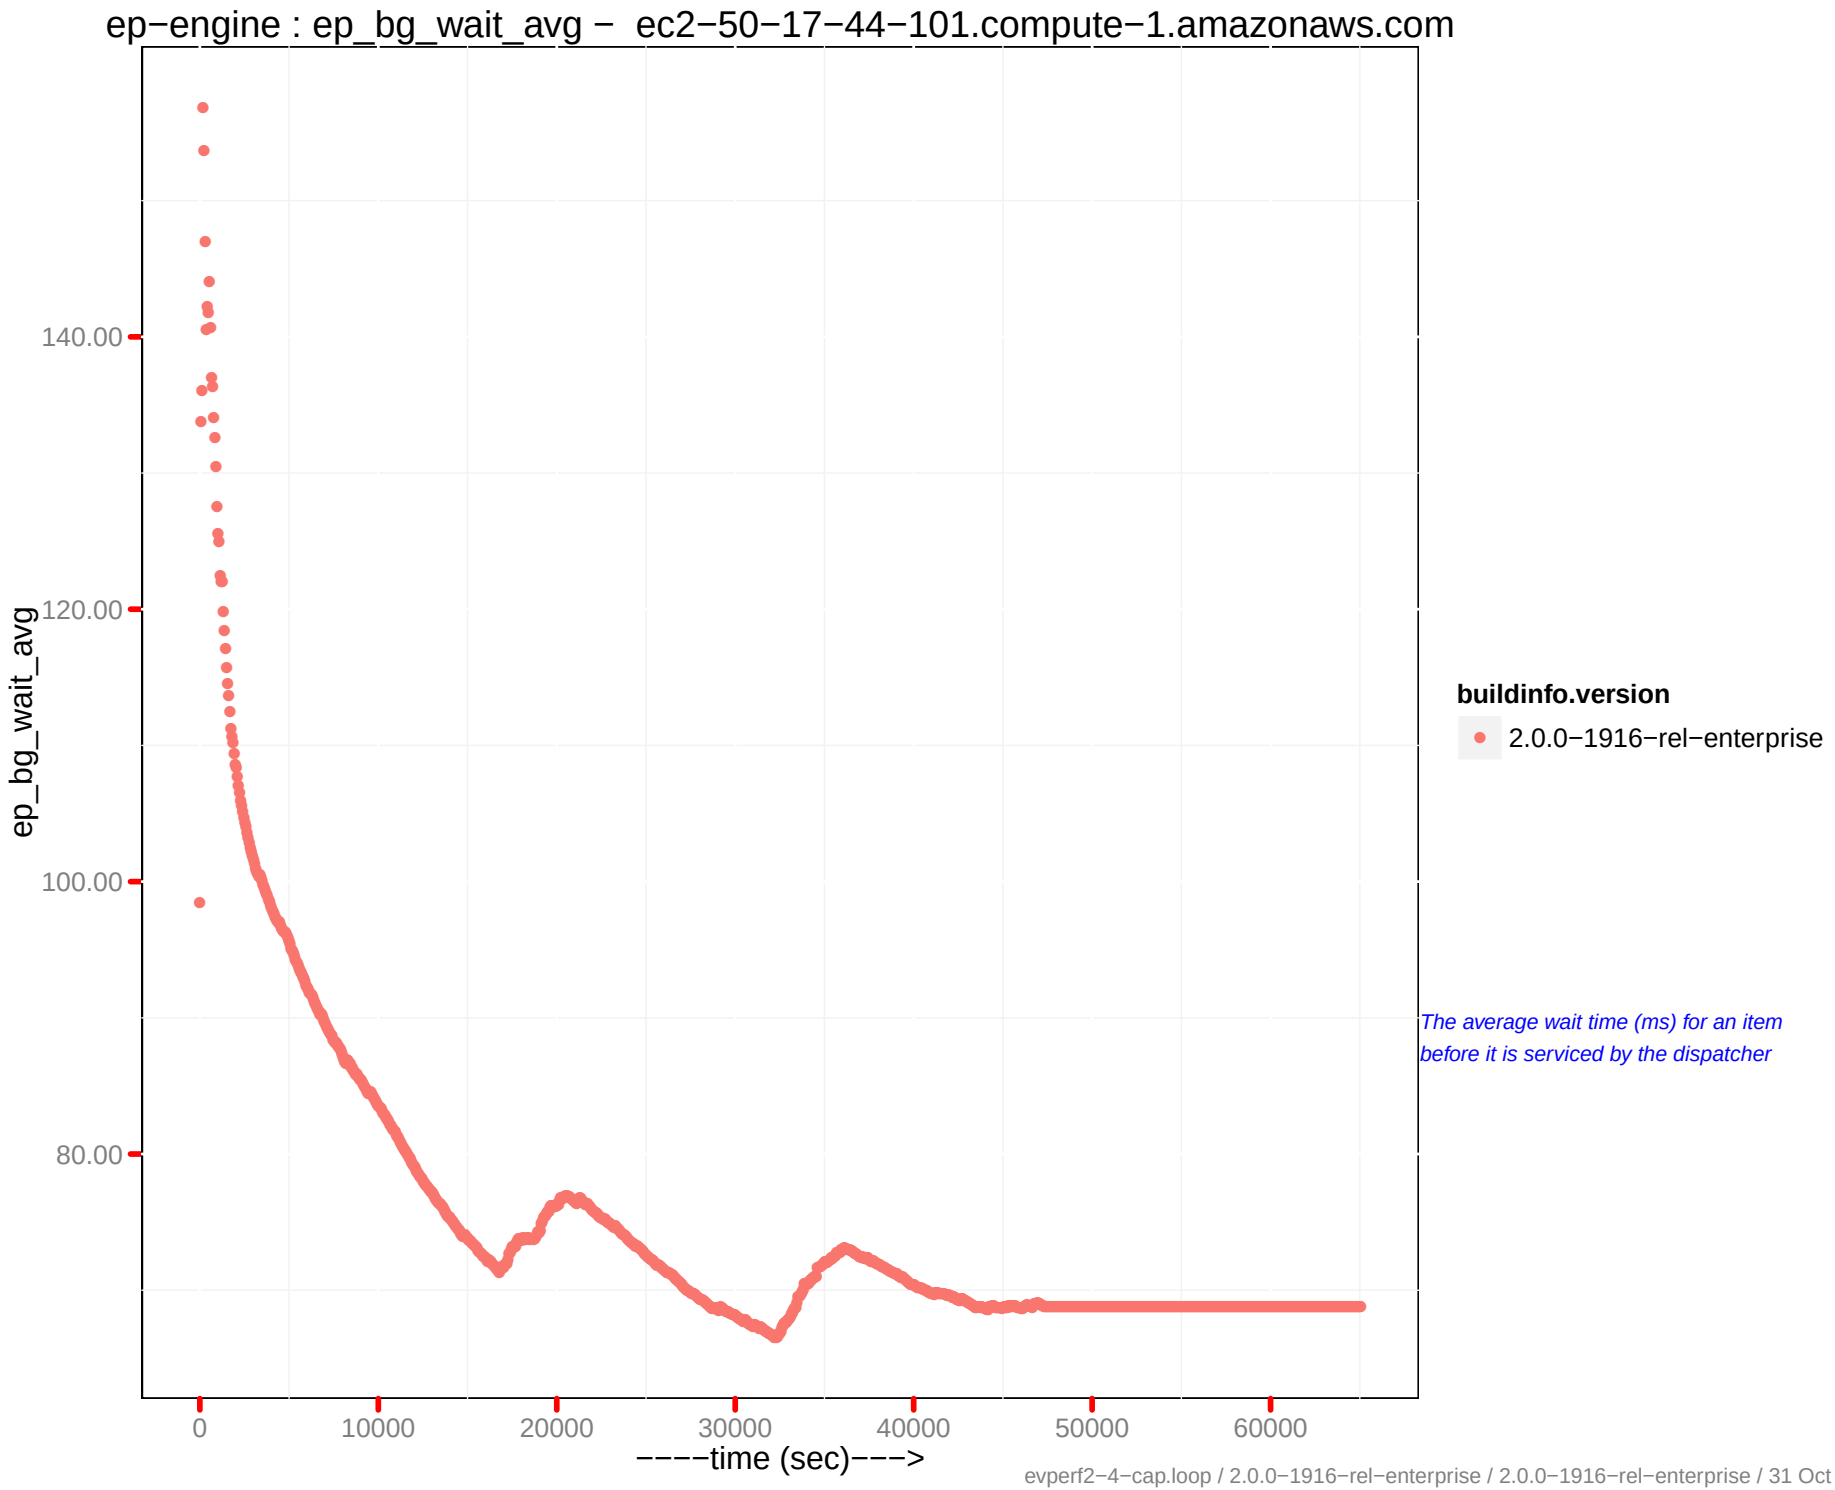
ep−engine : ep_bg_wait_avg − ec2−50−17−44−101.compute−1.amazonaws.com
ep_bg_wait_avg
140.00
120.00
100.00
80.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

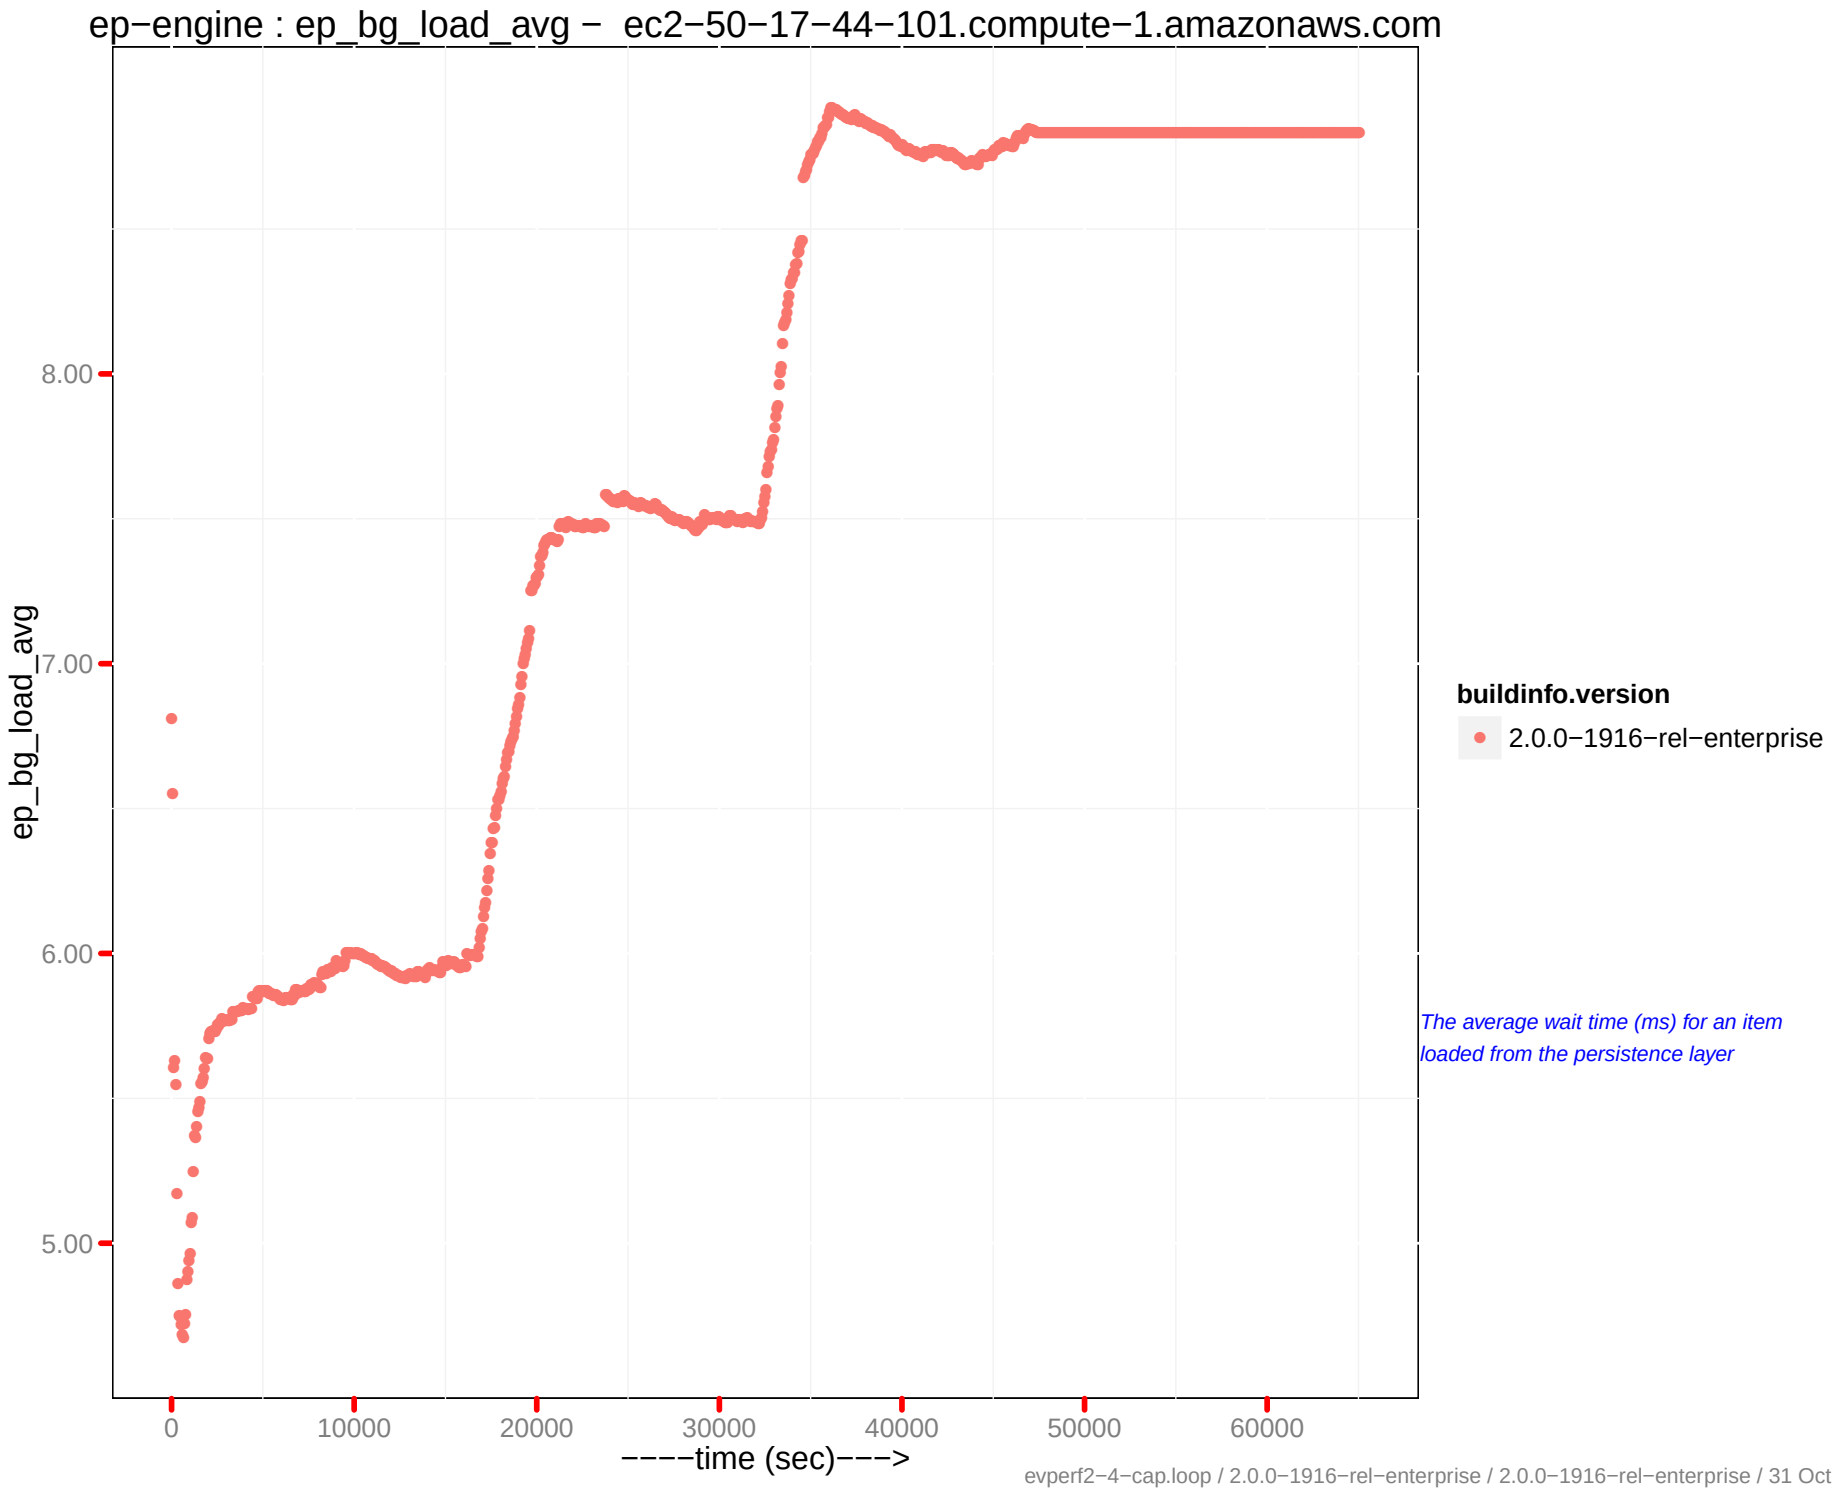
ep−engine : ep_bg_load_avg − ec2−50−17−44−101.compute−1.amazonaws.com
ep_bg_load_avg
8.00
7.00
6.00
5.00
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

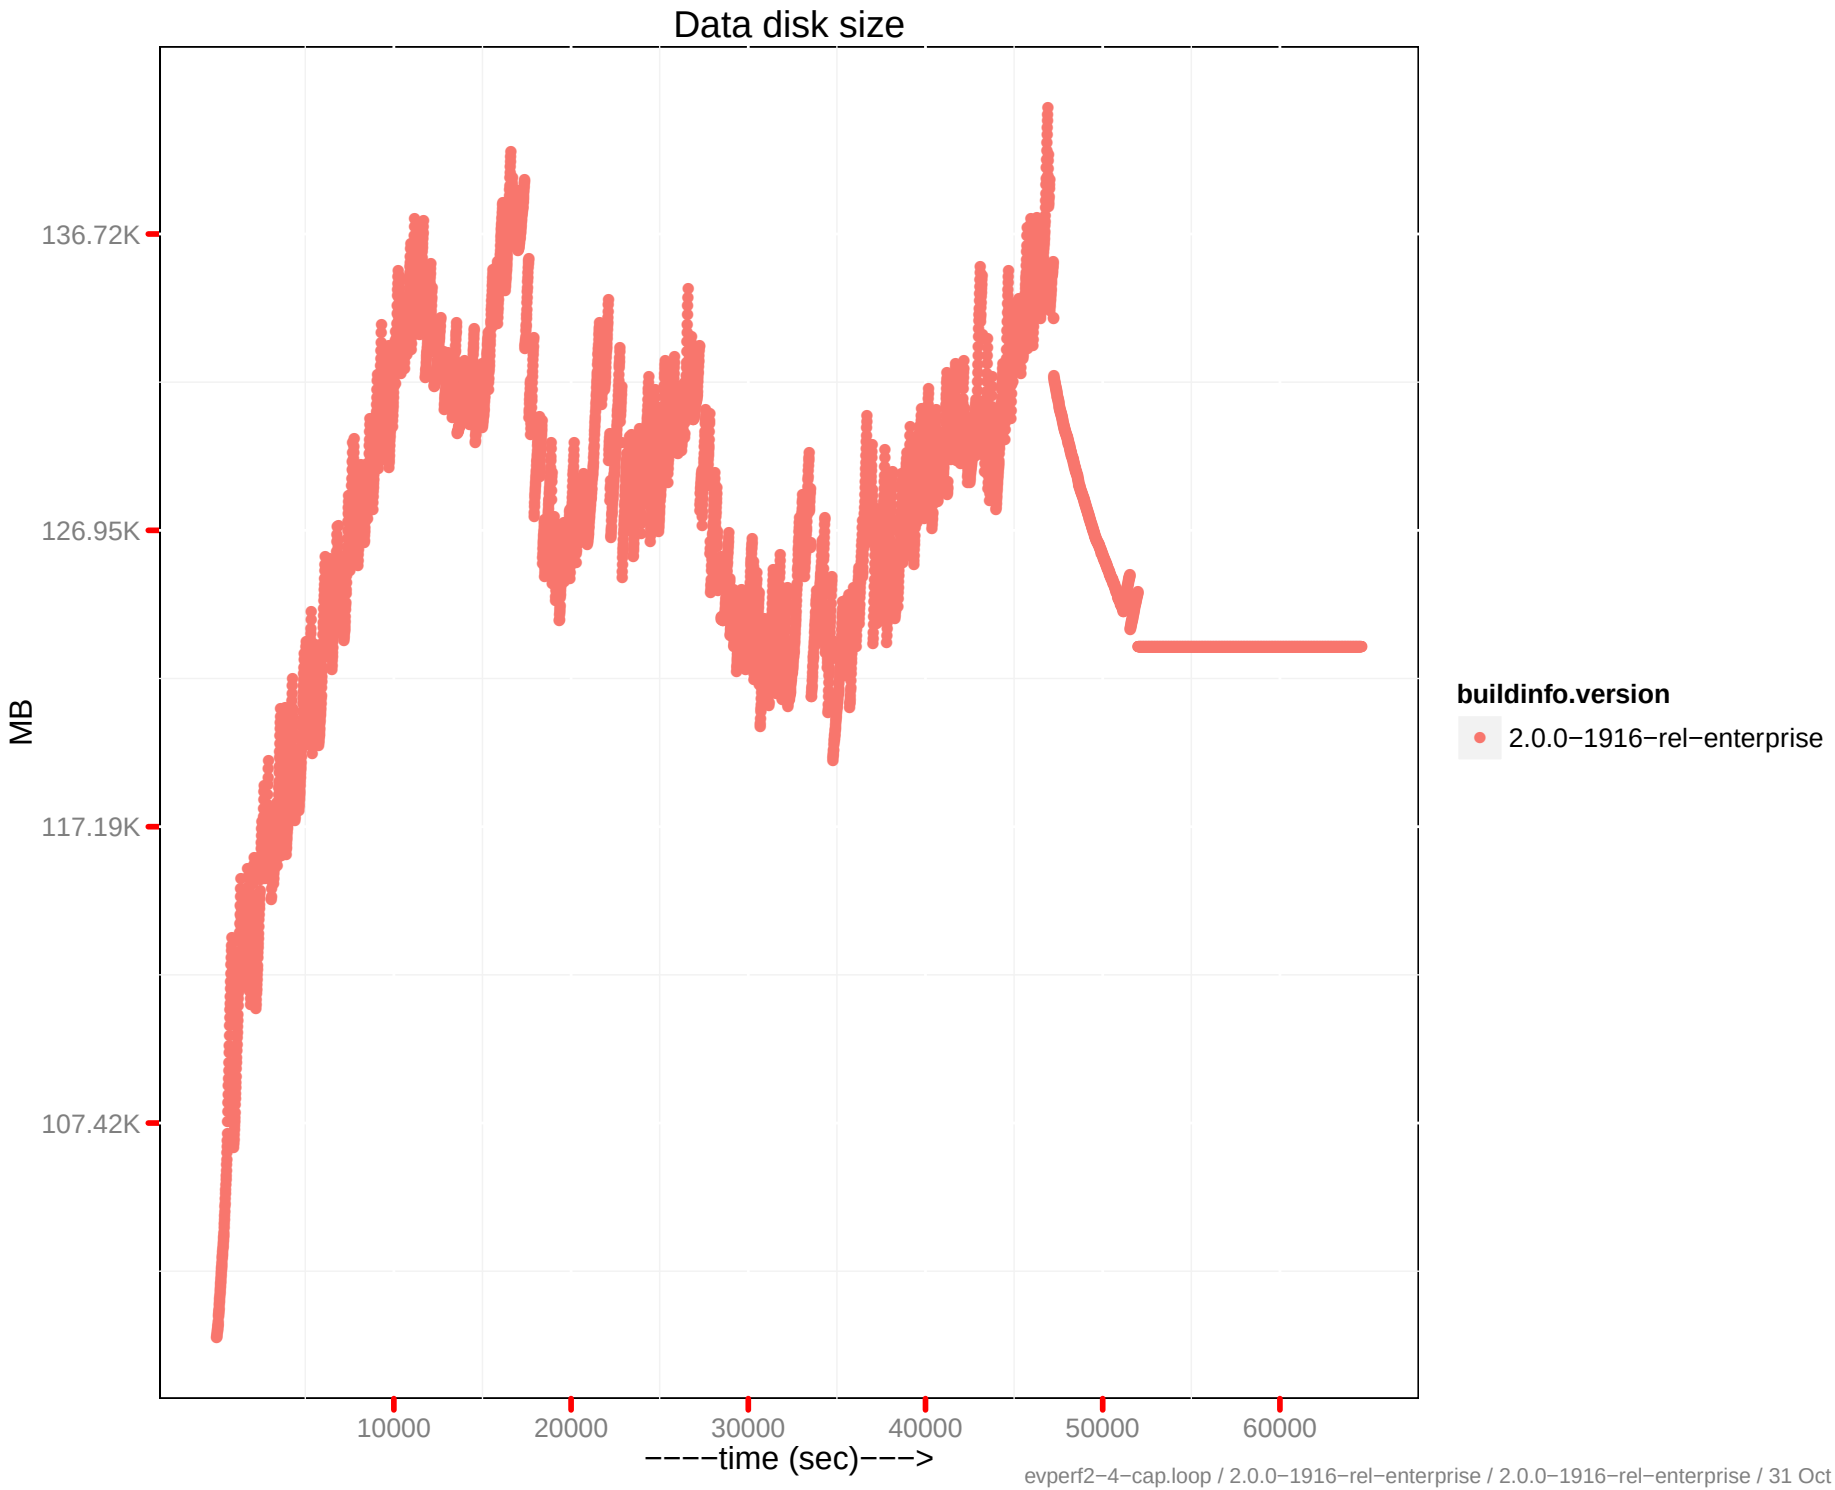
Data disk size
136.72K
126.95K
117.19K
107.42K
MB
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

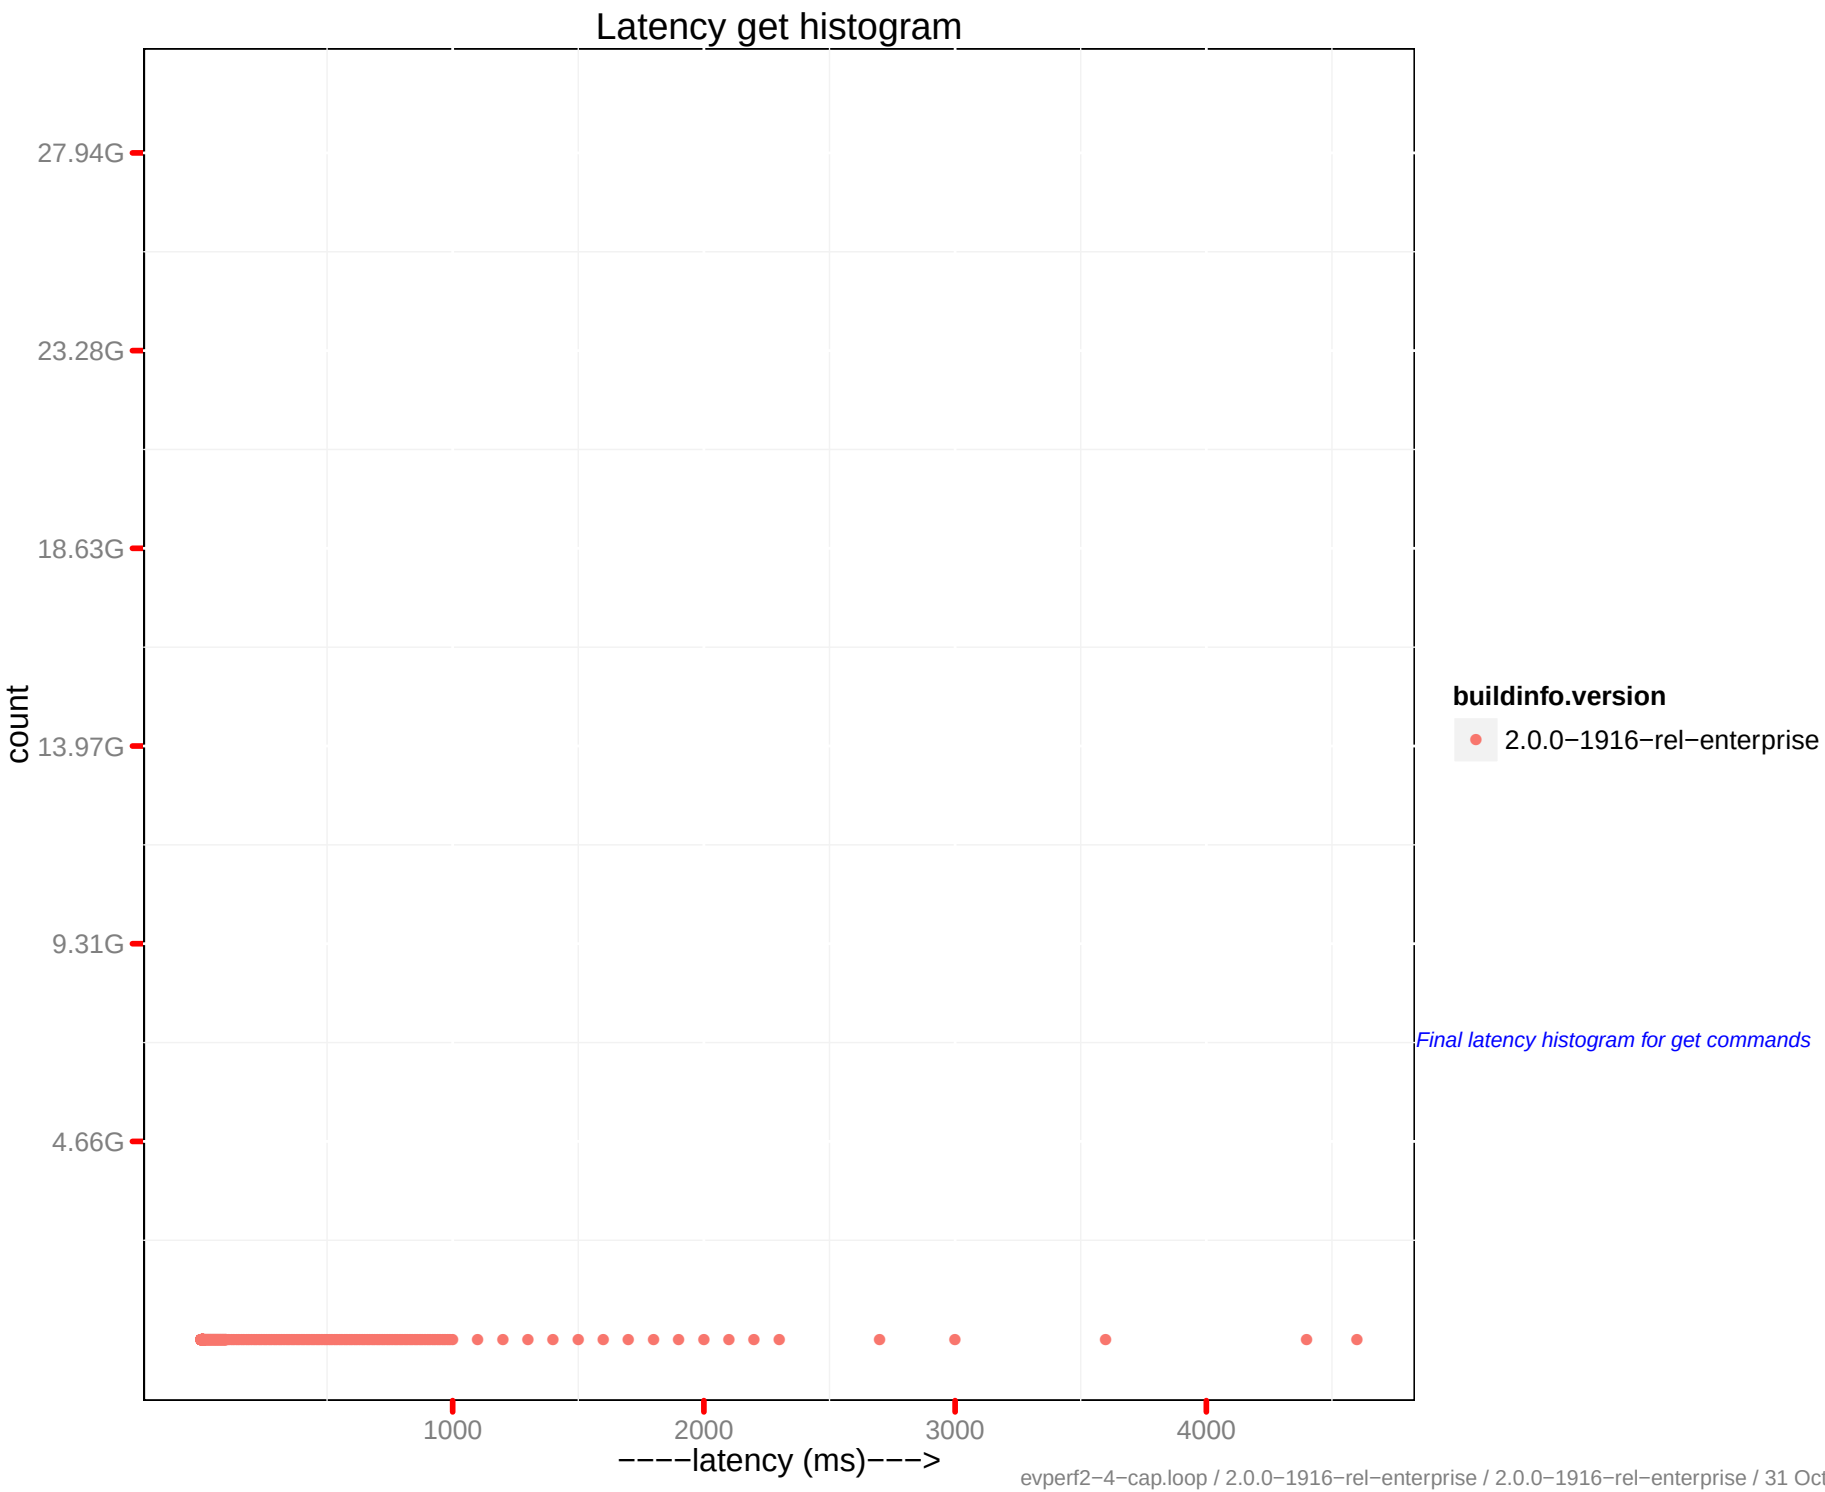
Latency get histogram
27.94G
23.28G
18.63G
13.97G
9.31G
4.66G
count
1000
2000
3000
4000
----latency (ms)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Final latency histogram for get commands
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

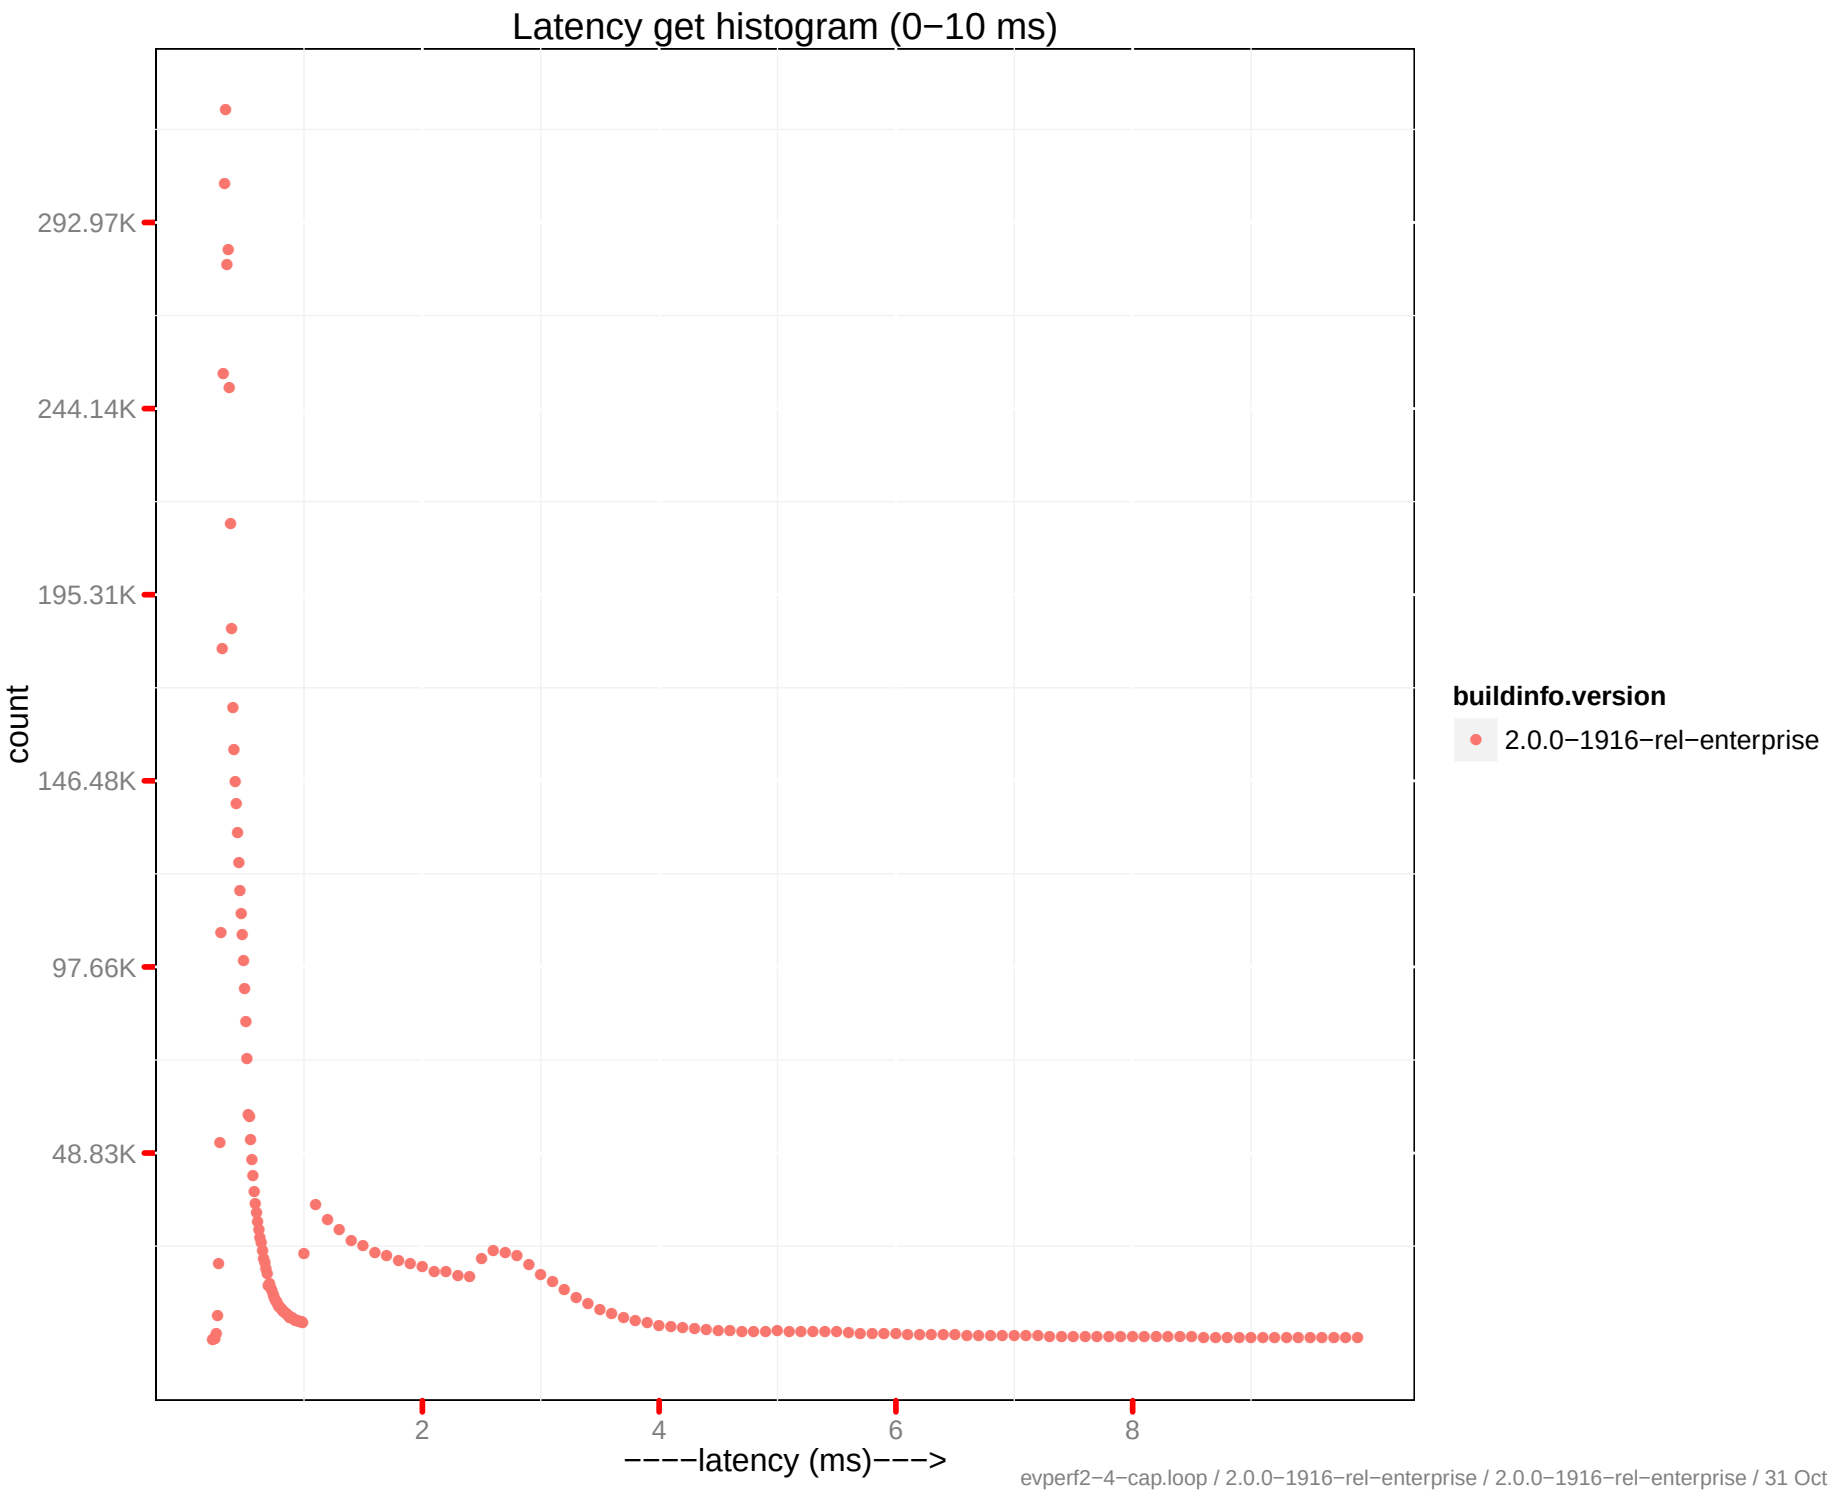
Latency get histogram (0−10 ms)
292.97K
244.14K
195.31K
146.48K
97.66K
48.83K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

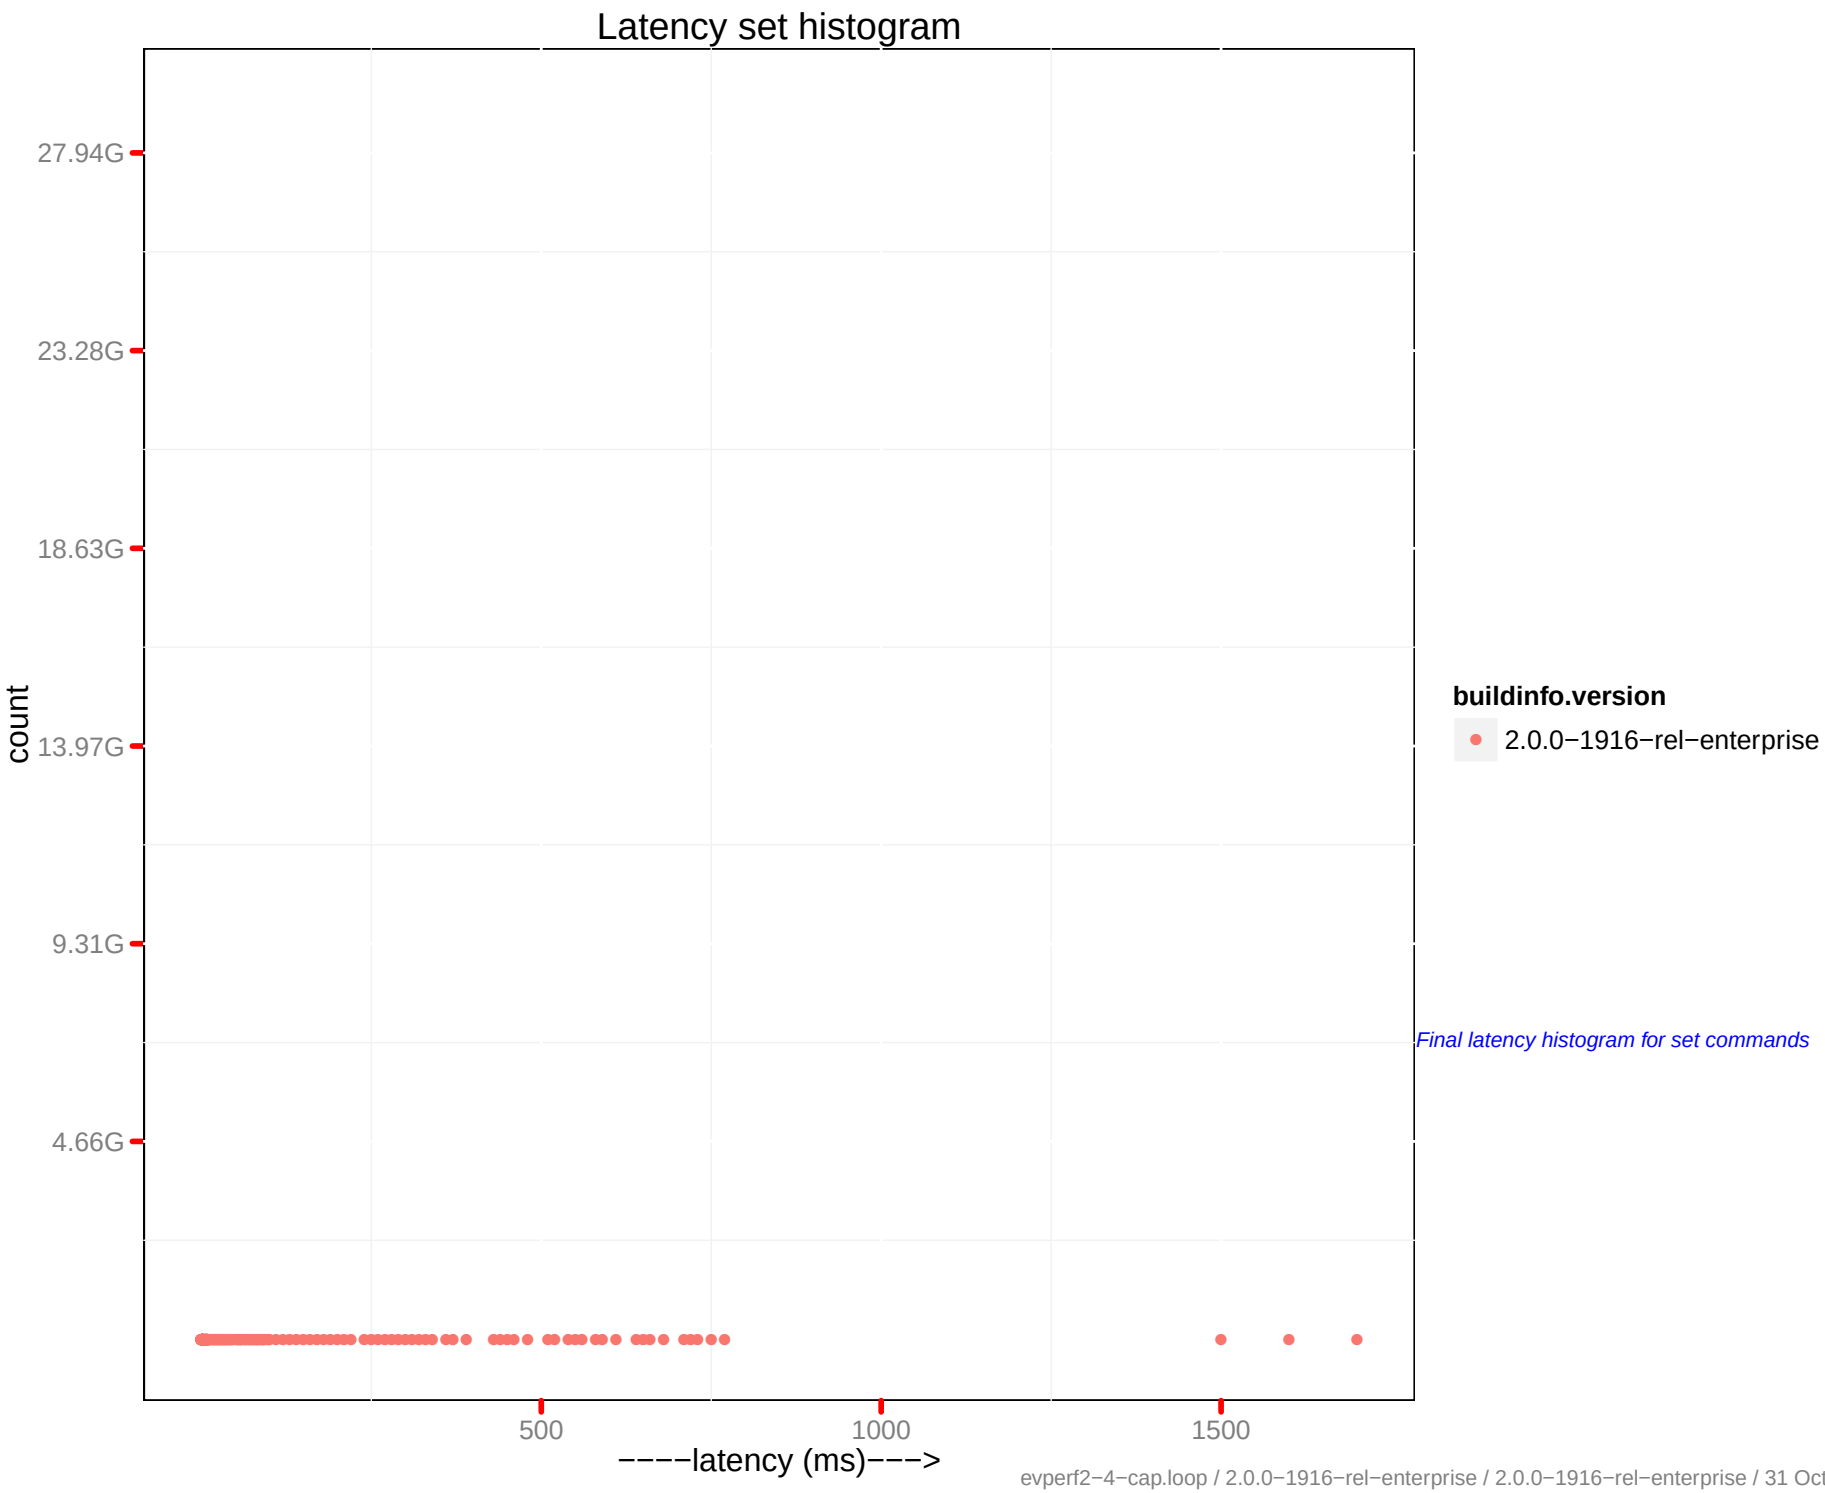
Latency set histogram
27.94G
23.28G
18.63G
13.97G
9.31G
4.66G
count
500
1000
1500
----latency (ms)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Final latency histogram for set commands
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

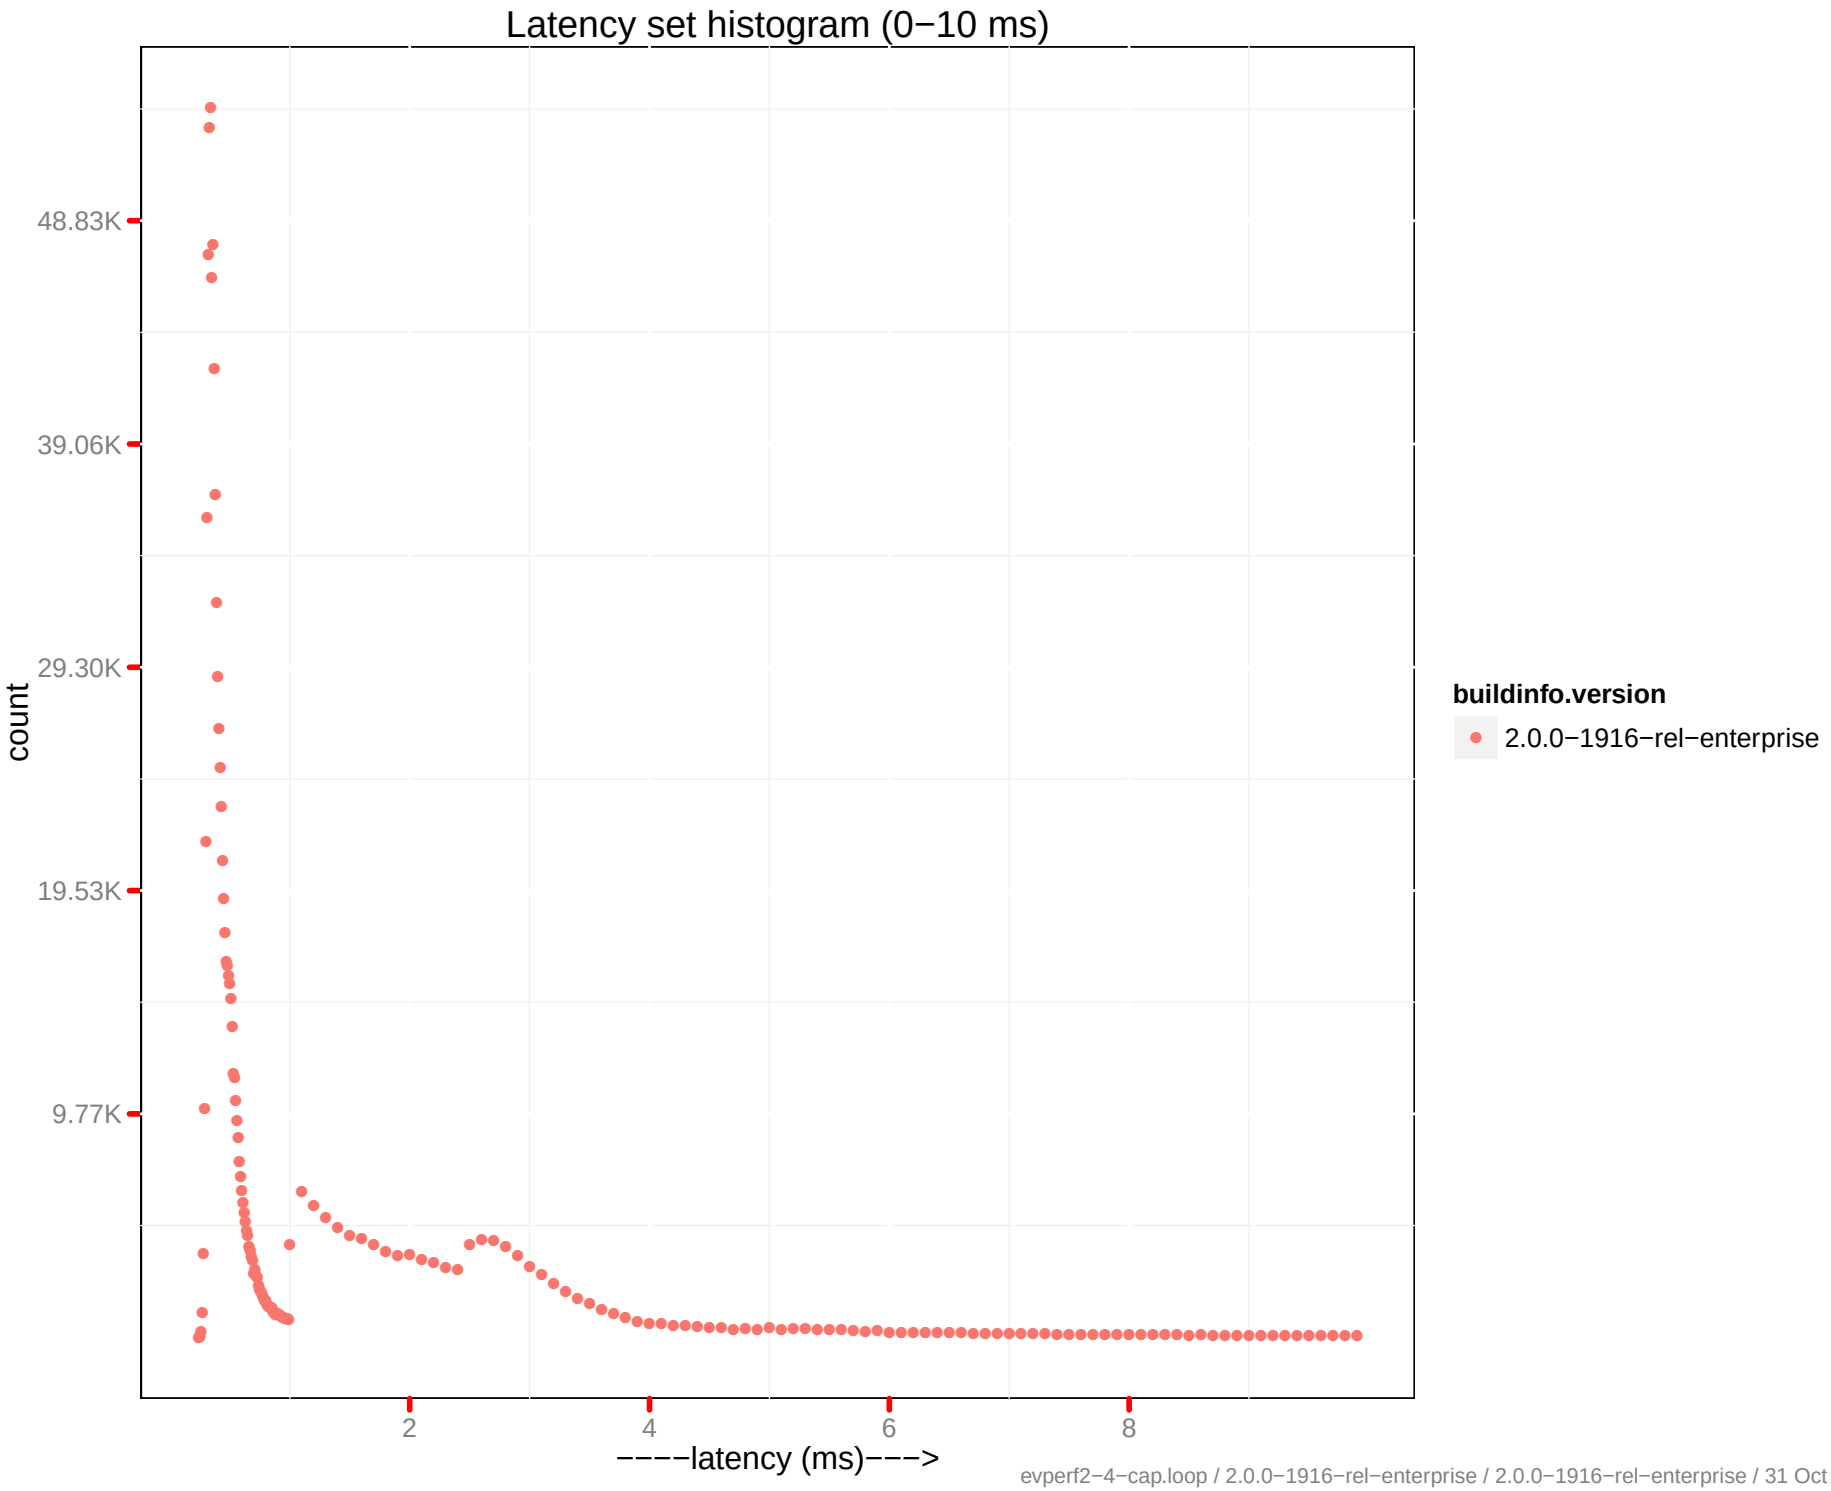
Latency set histogram (0−10 ms)
48.83K
39.06K
29.30K
19.53K
9.77K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

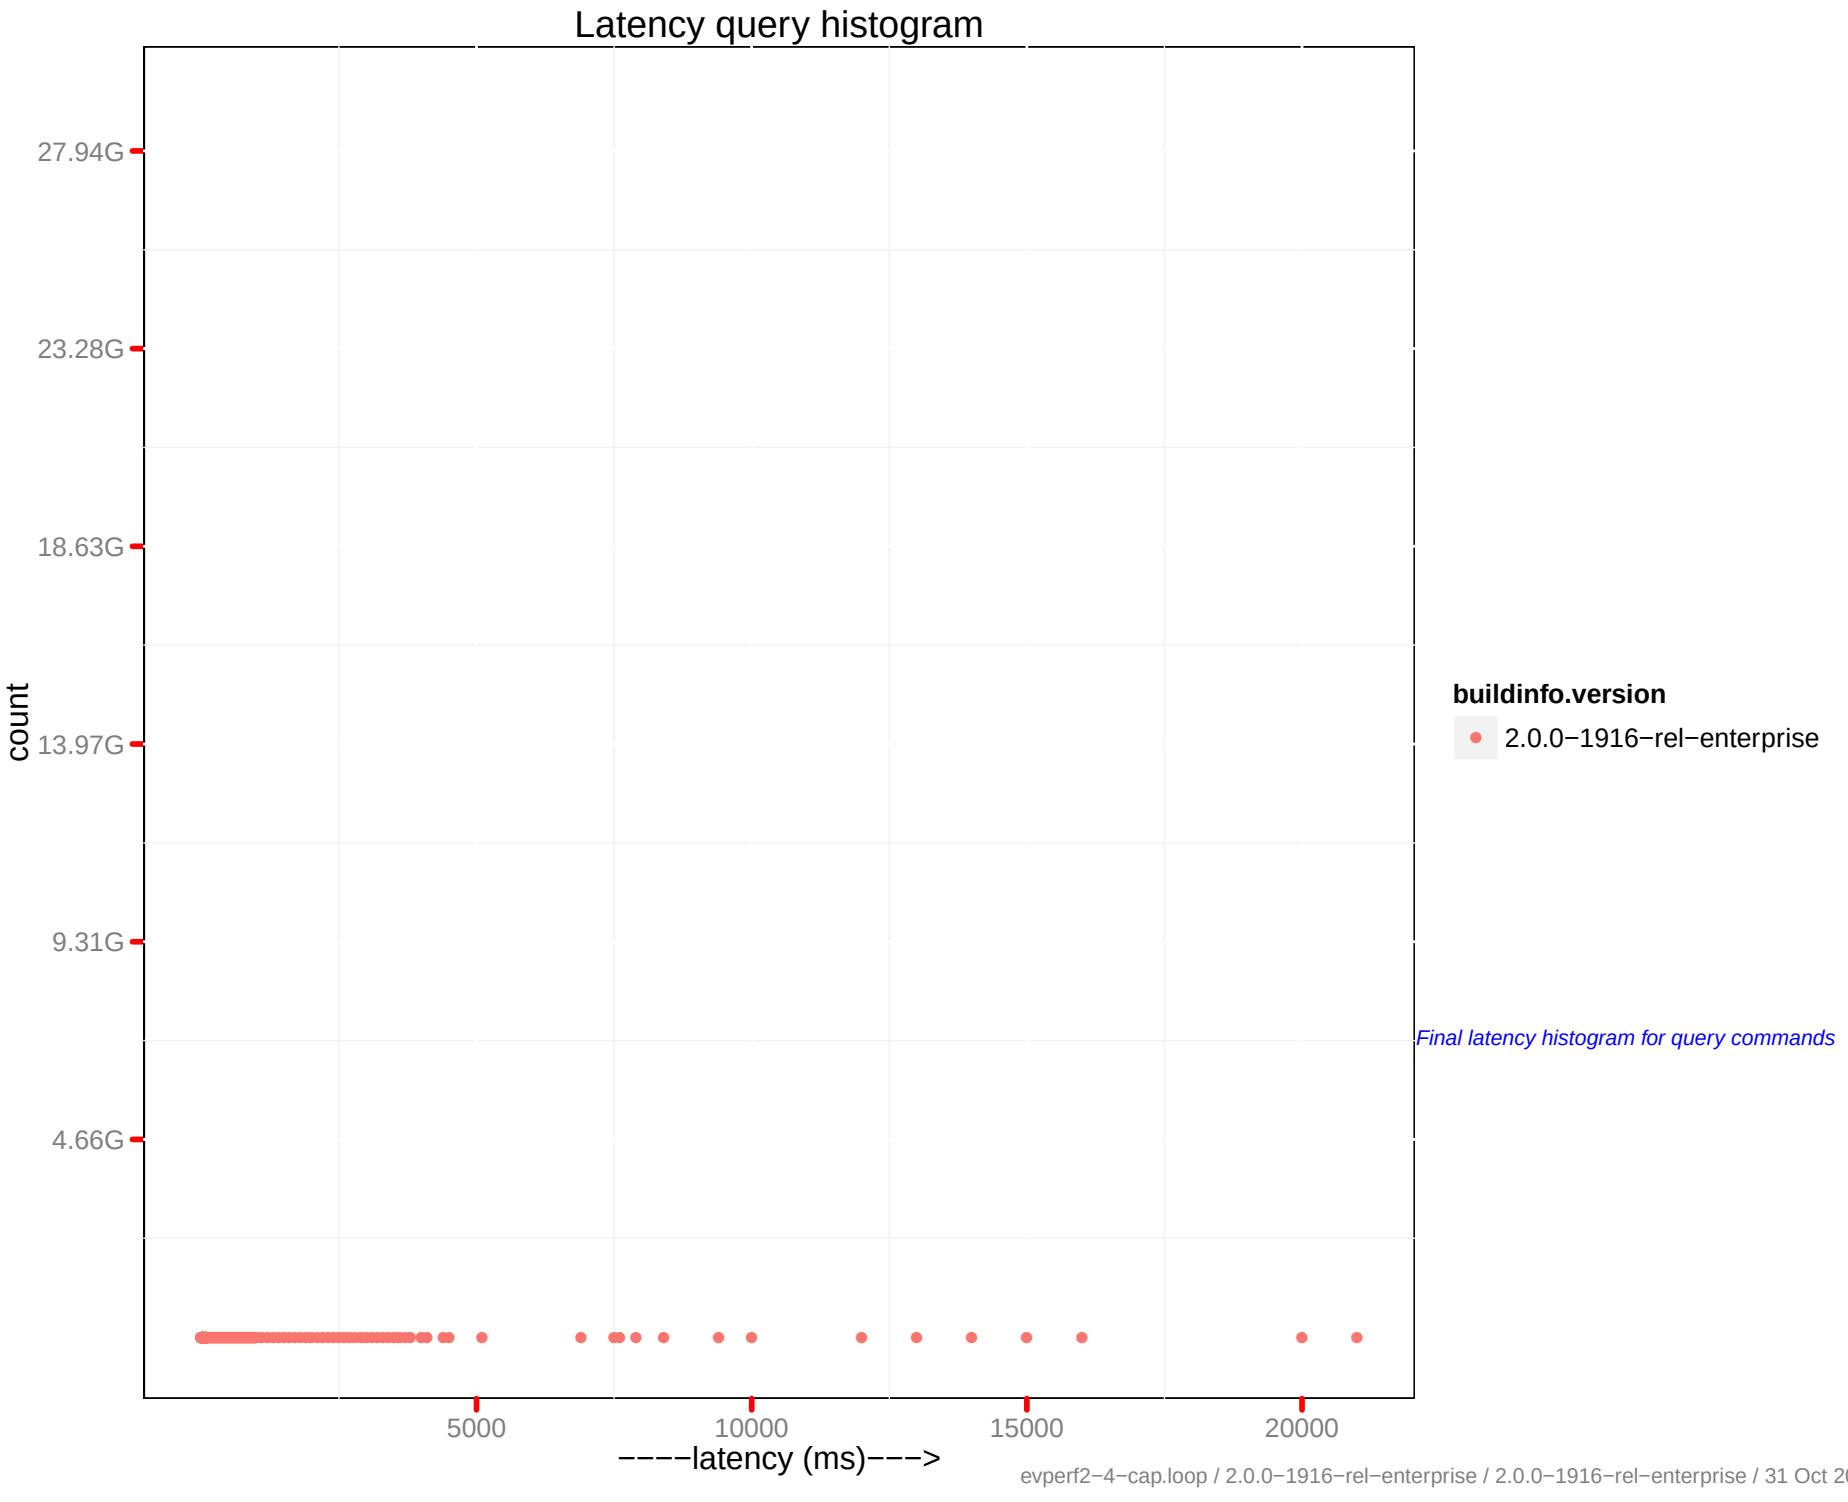
Latency query histogram
27.94G
23.28G
18.63G
13.97G
9.31G
4.66G
count
5000
10000
15000
20000
----latency (ms)--->
buildinfo.version
2.0.0-1916-rel-enterprise
Final latency histogram for query commands
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

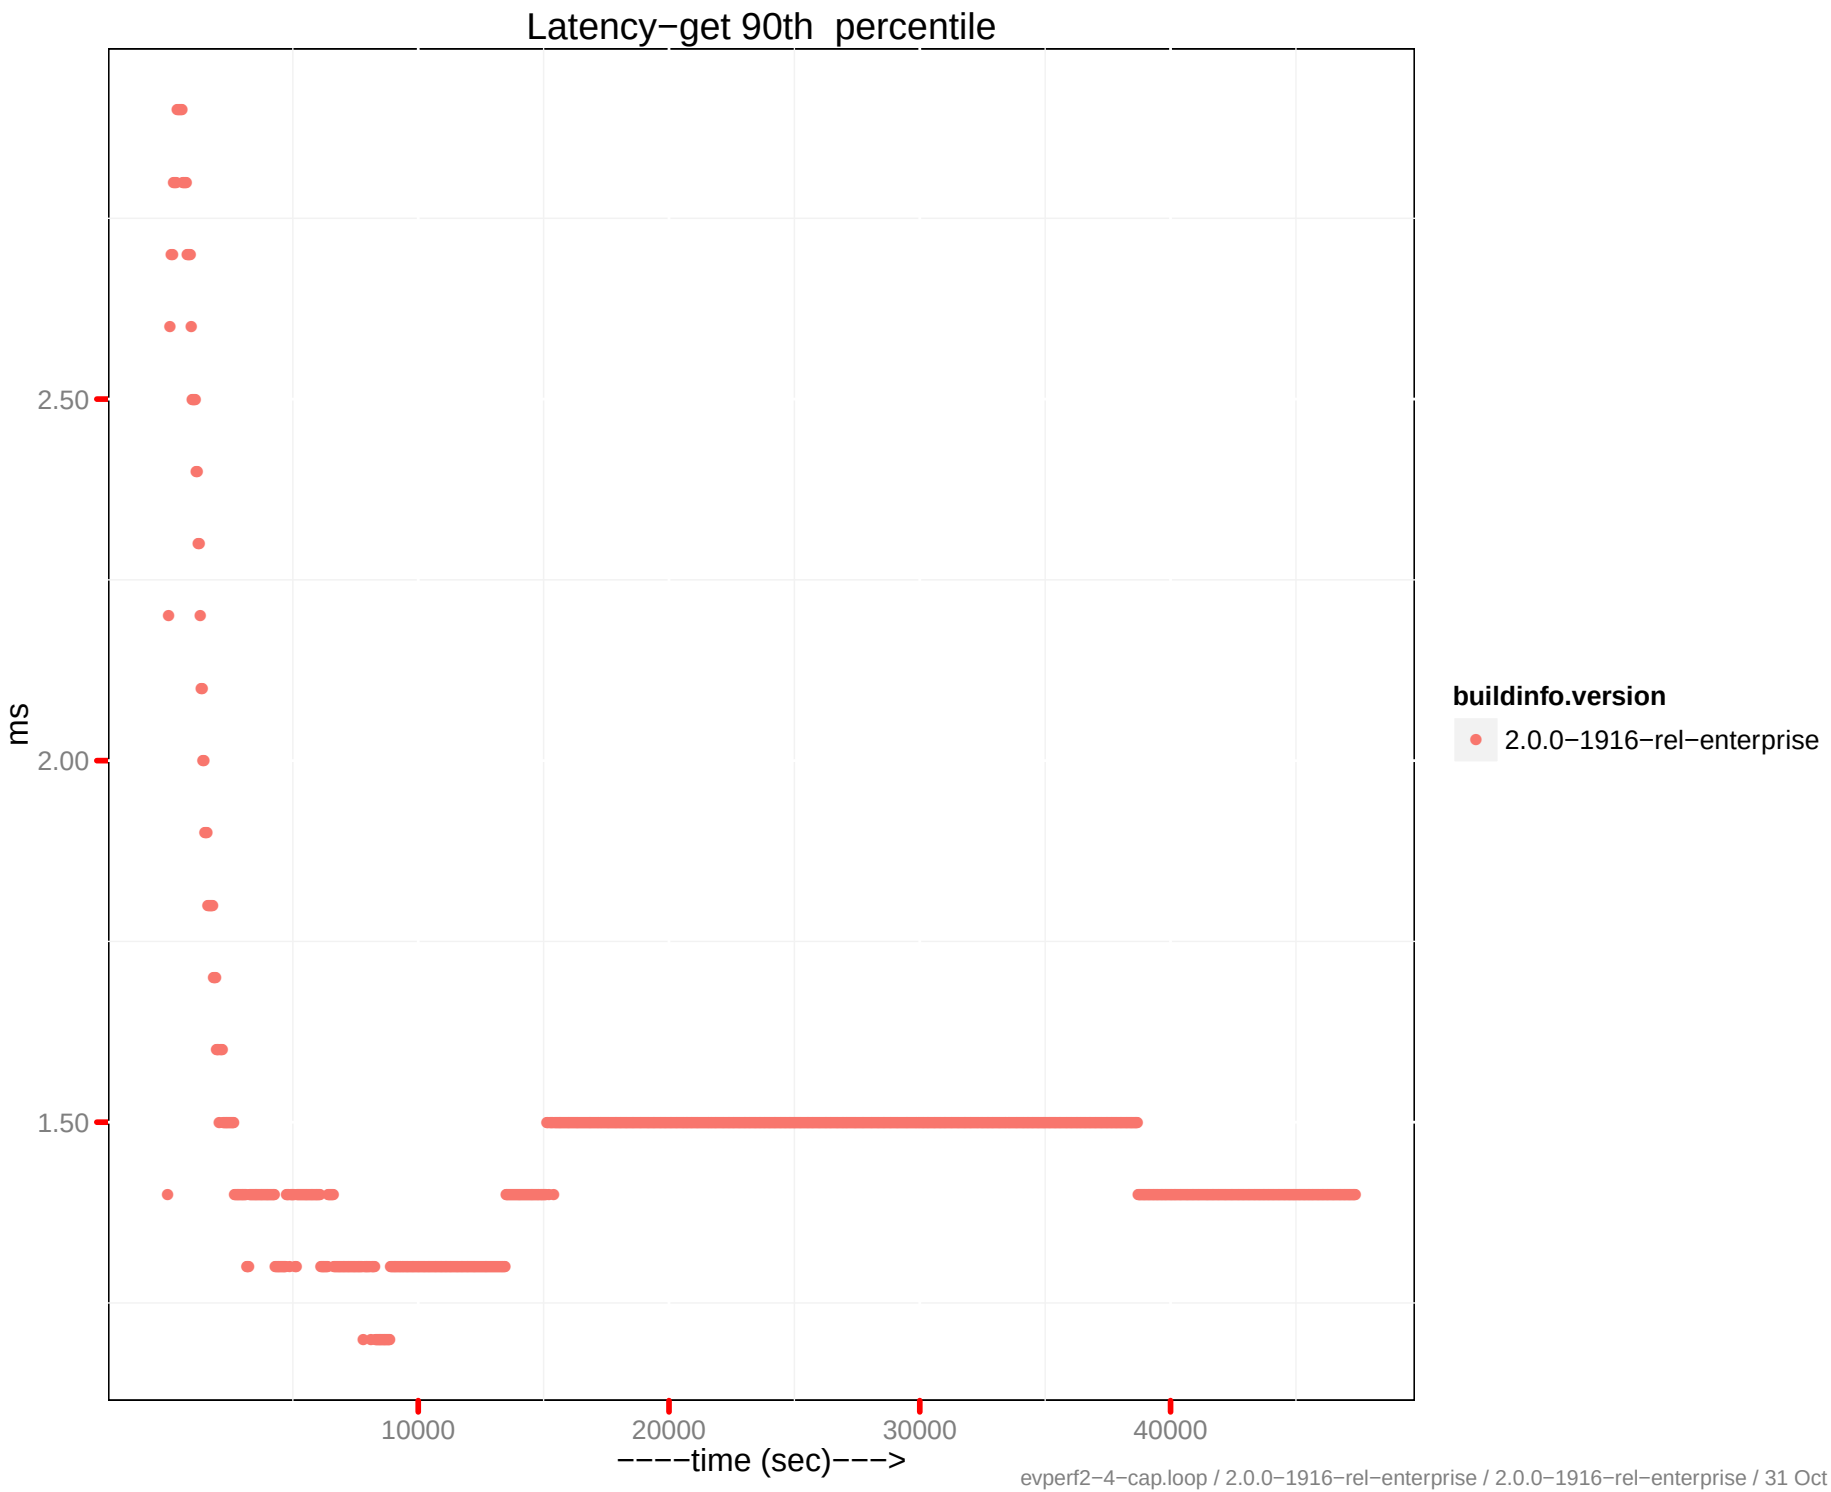
Latency−get 90th percentile
2.50
2.00
1.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

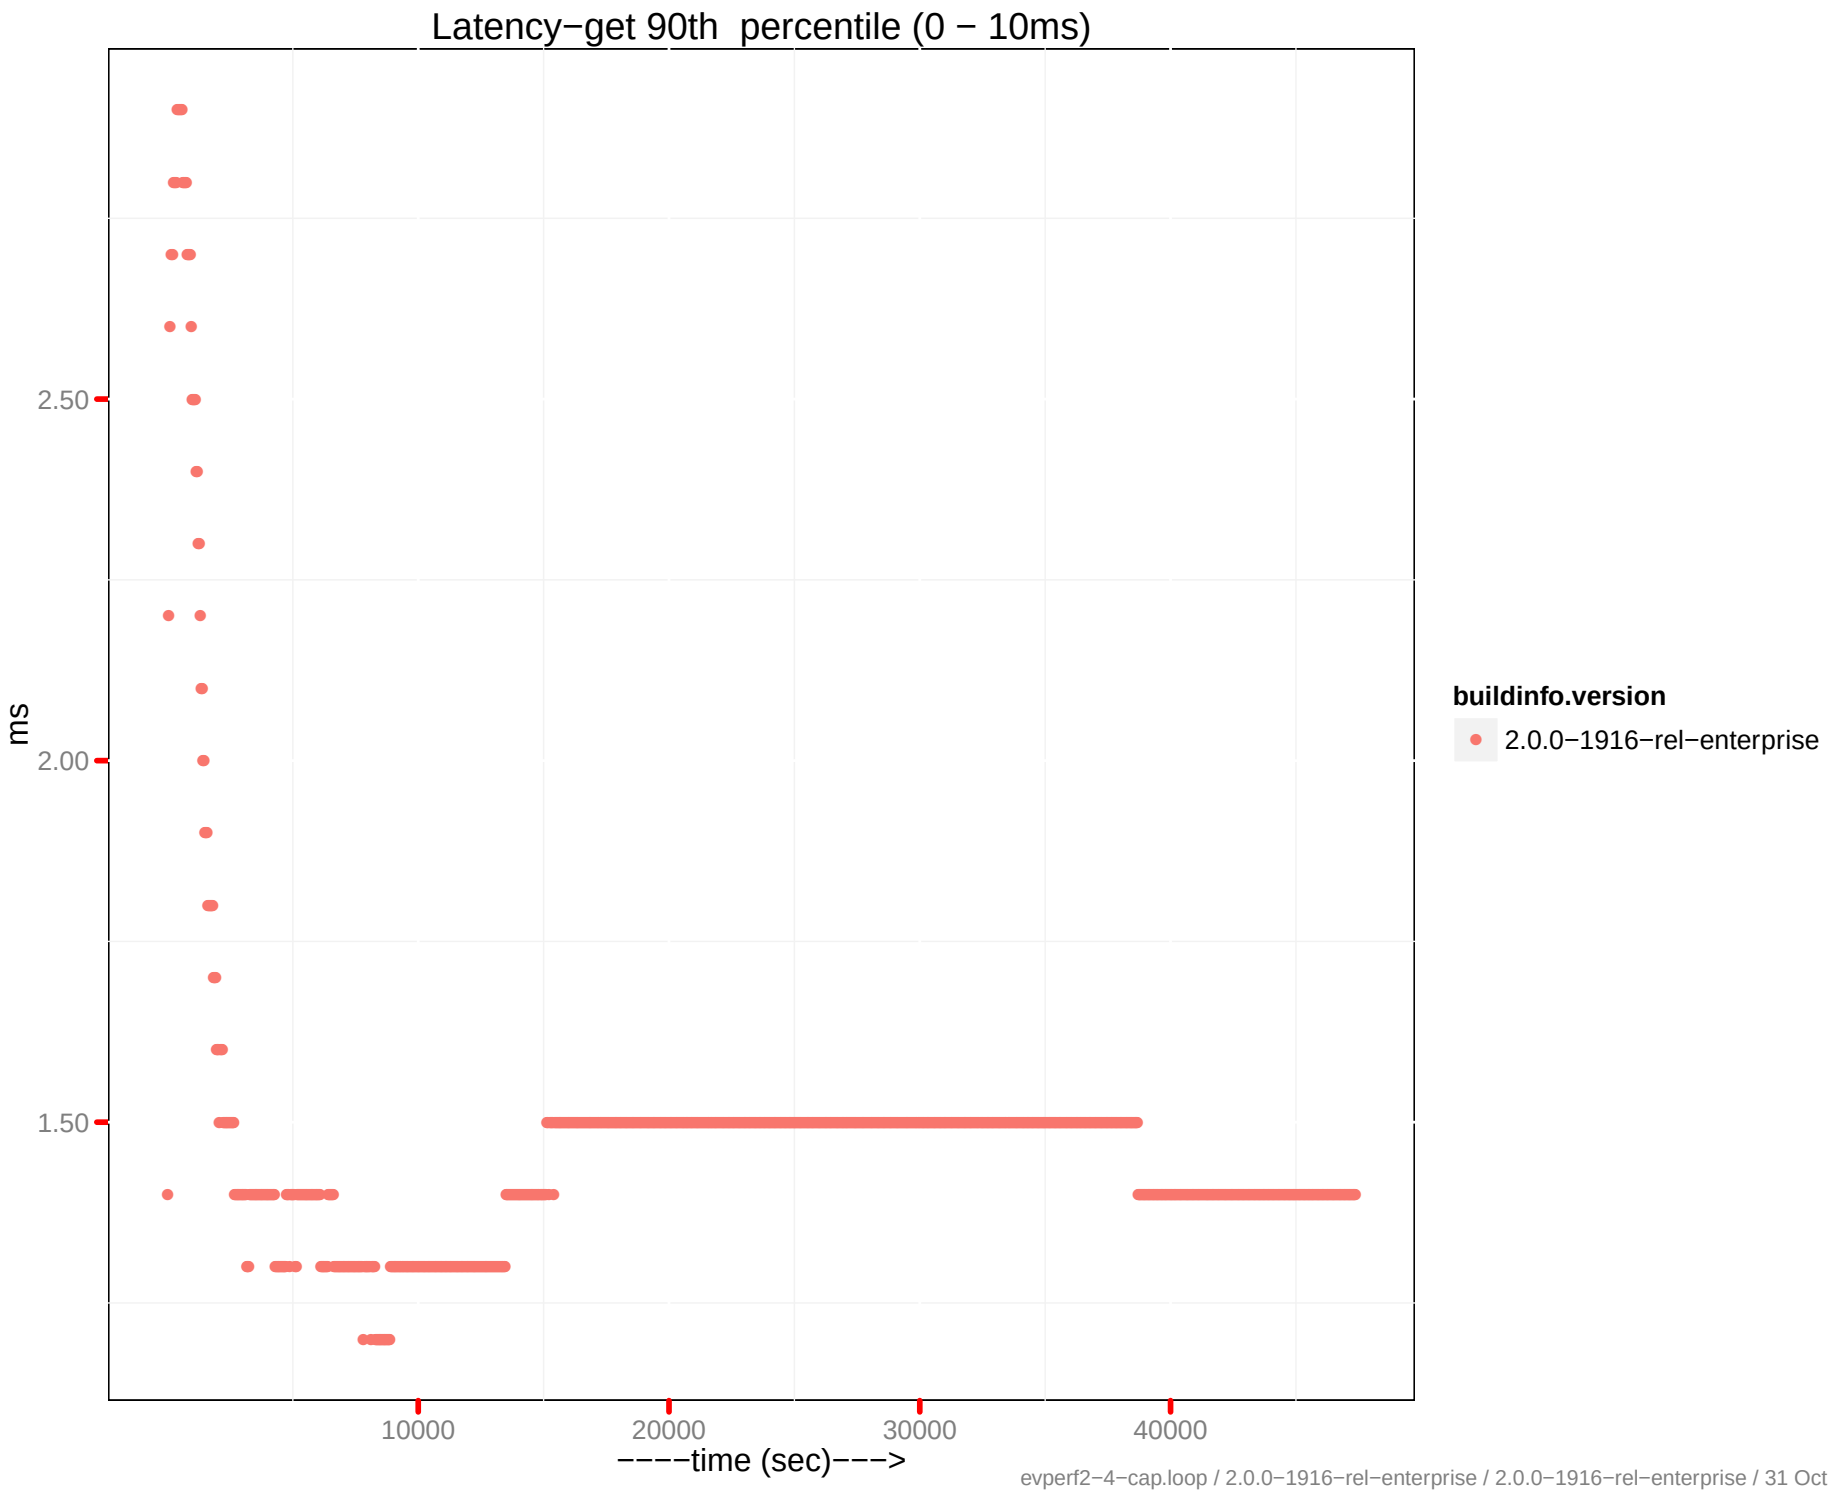
Latency−get 90th percentile (0 − 10ms)
2.50
2.00
1.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

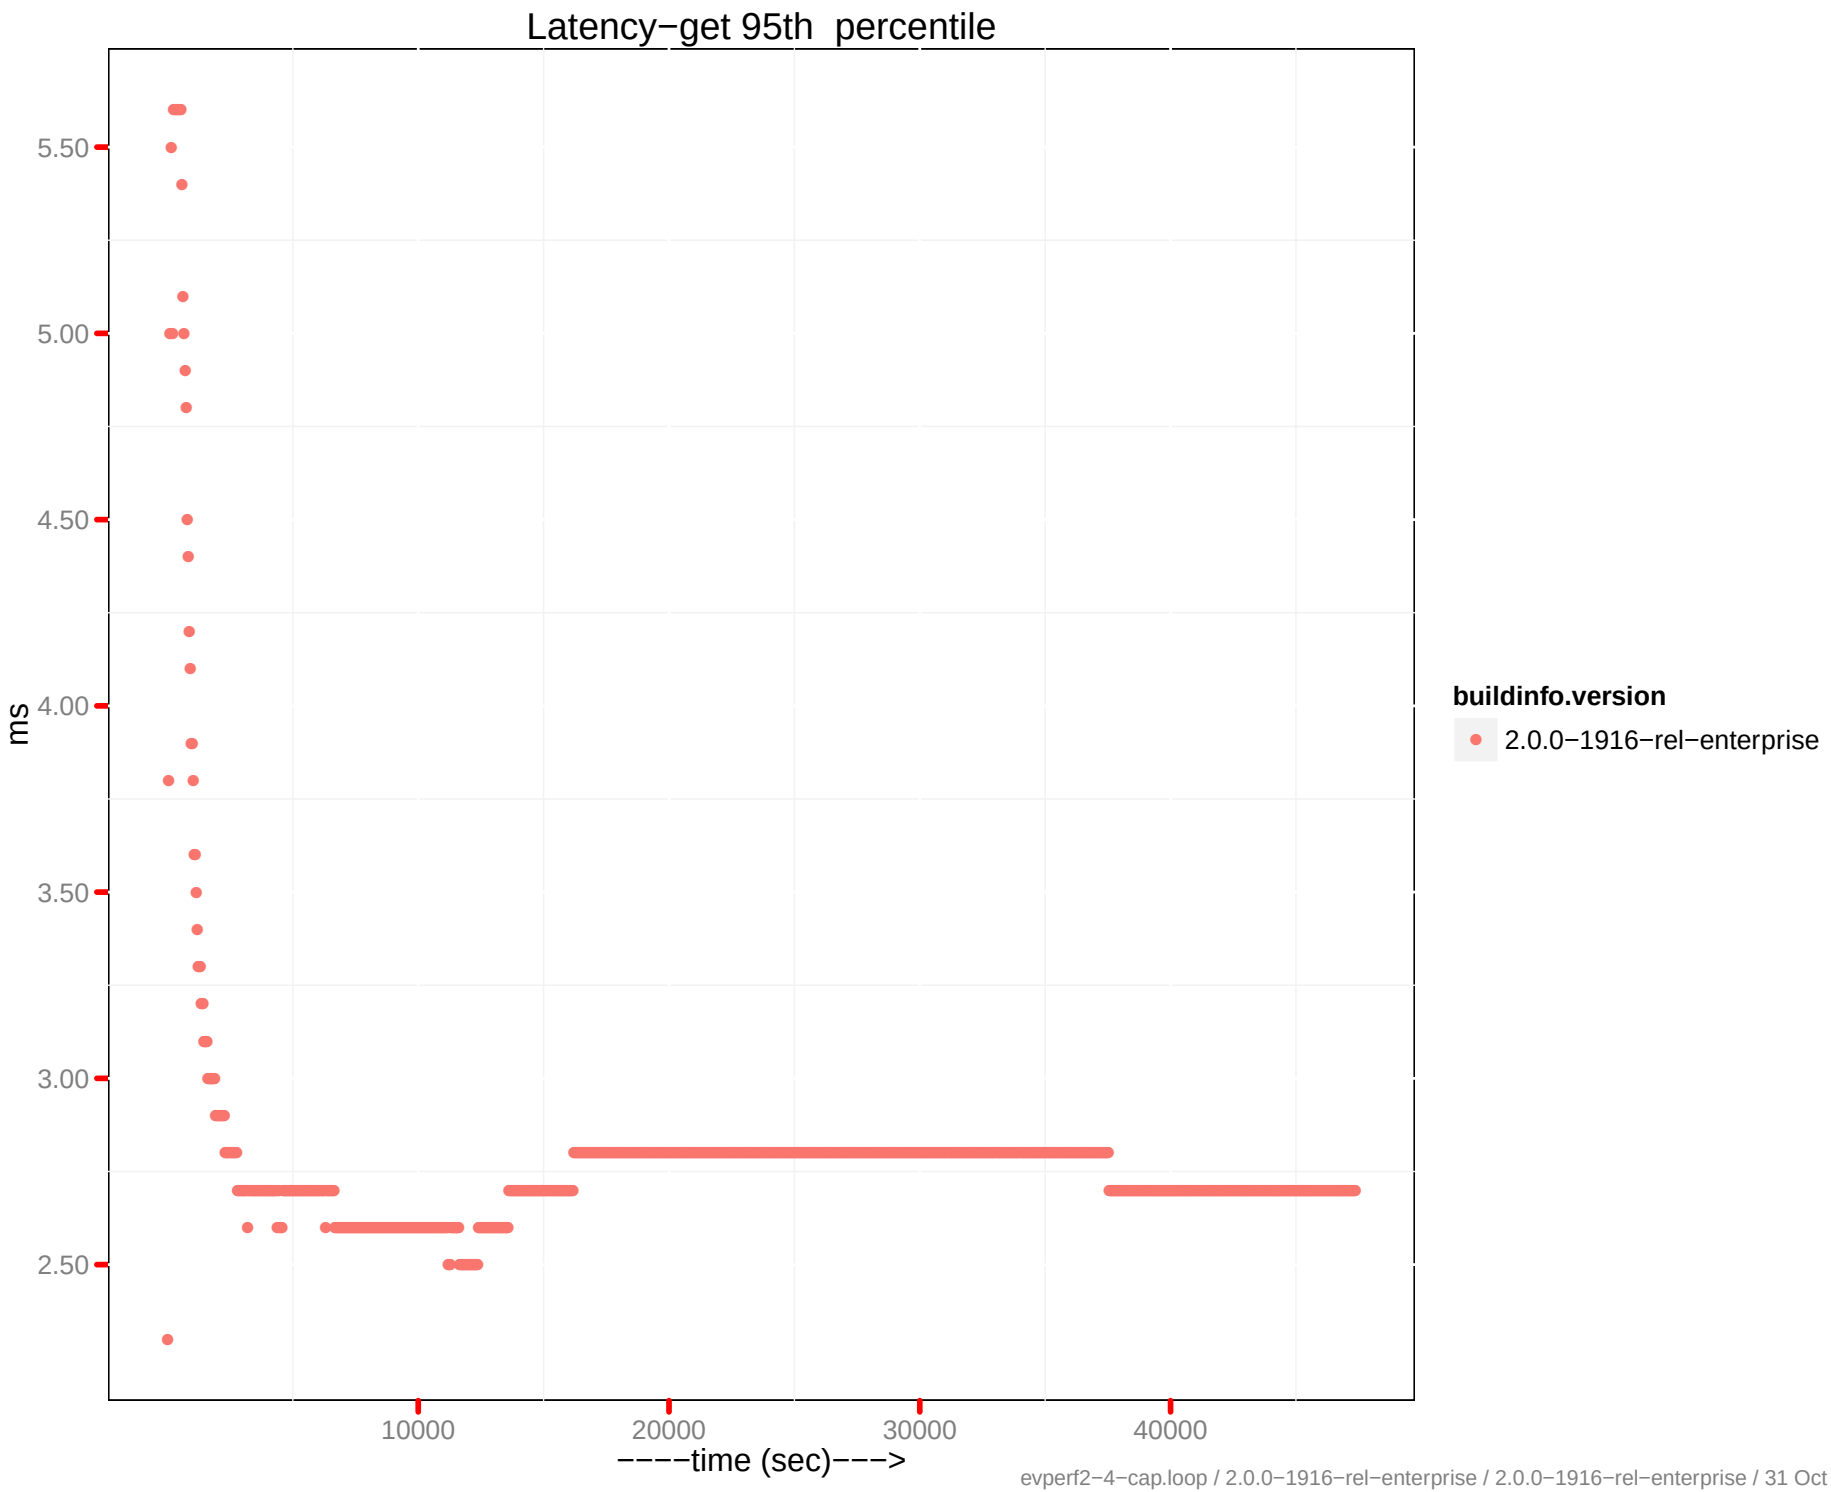
Latency−get 95th percentile
5.50
5.00
4.50
4.00
3.50
3.00
2.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

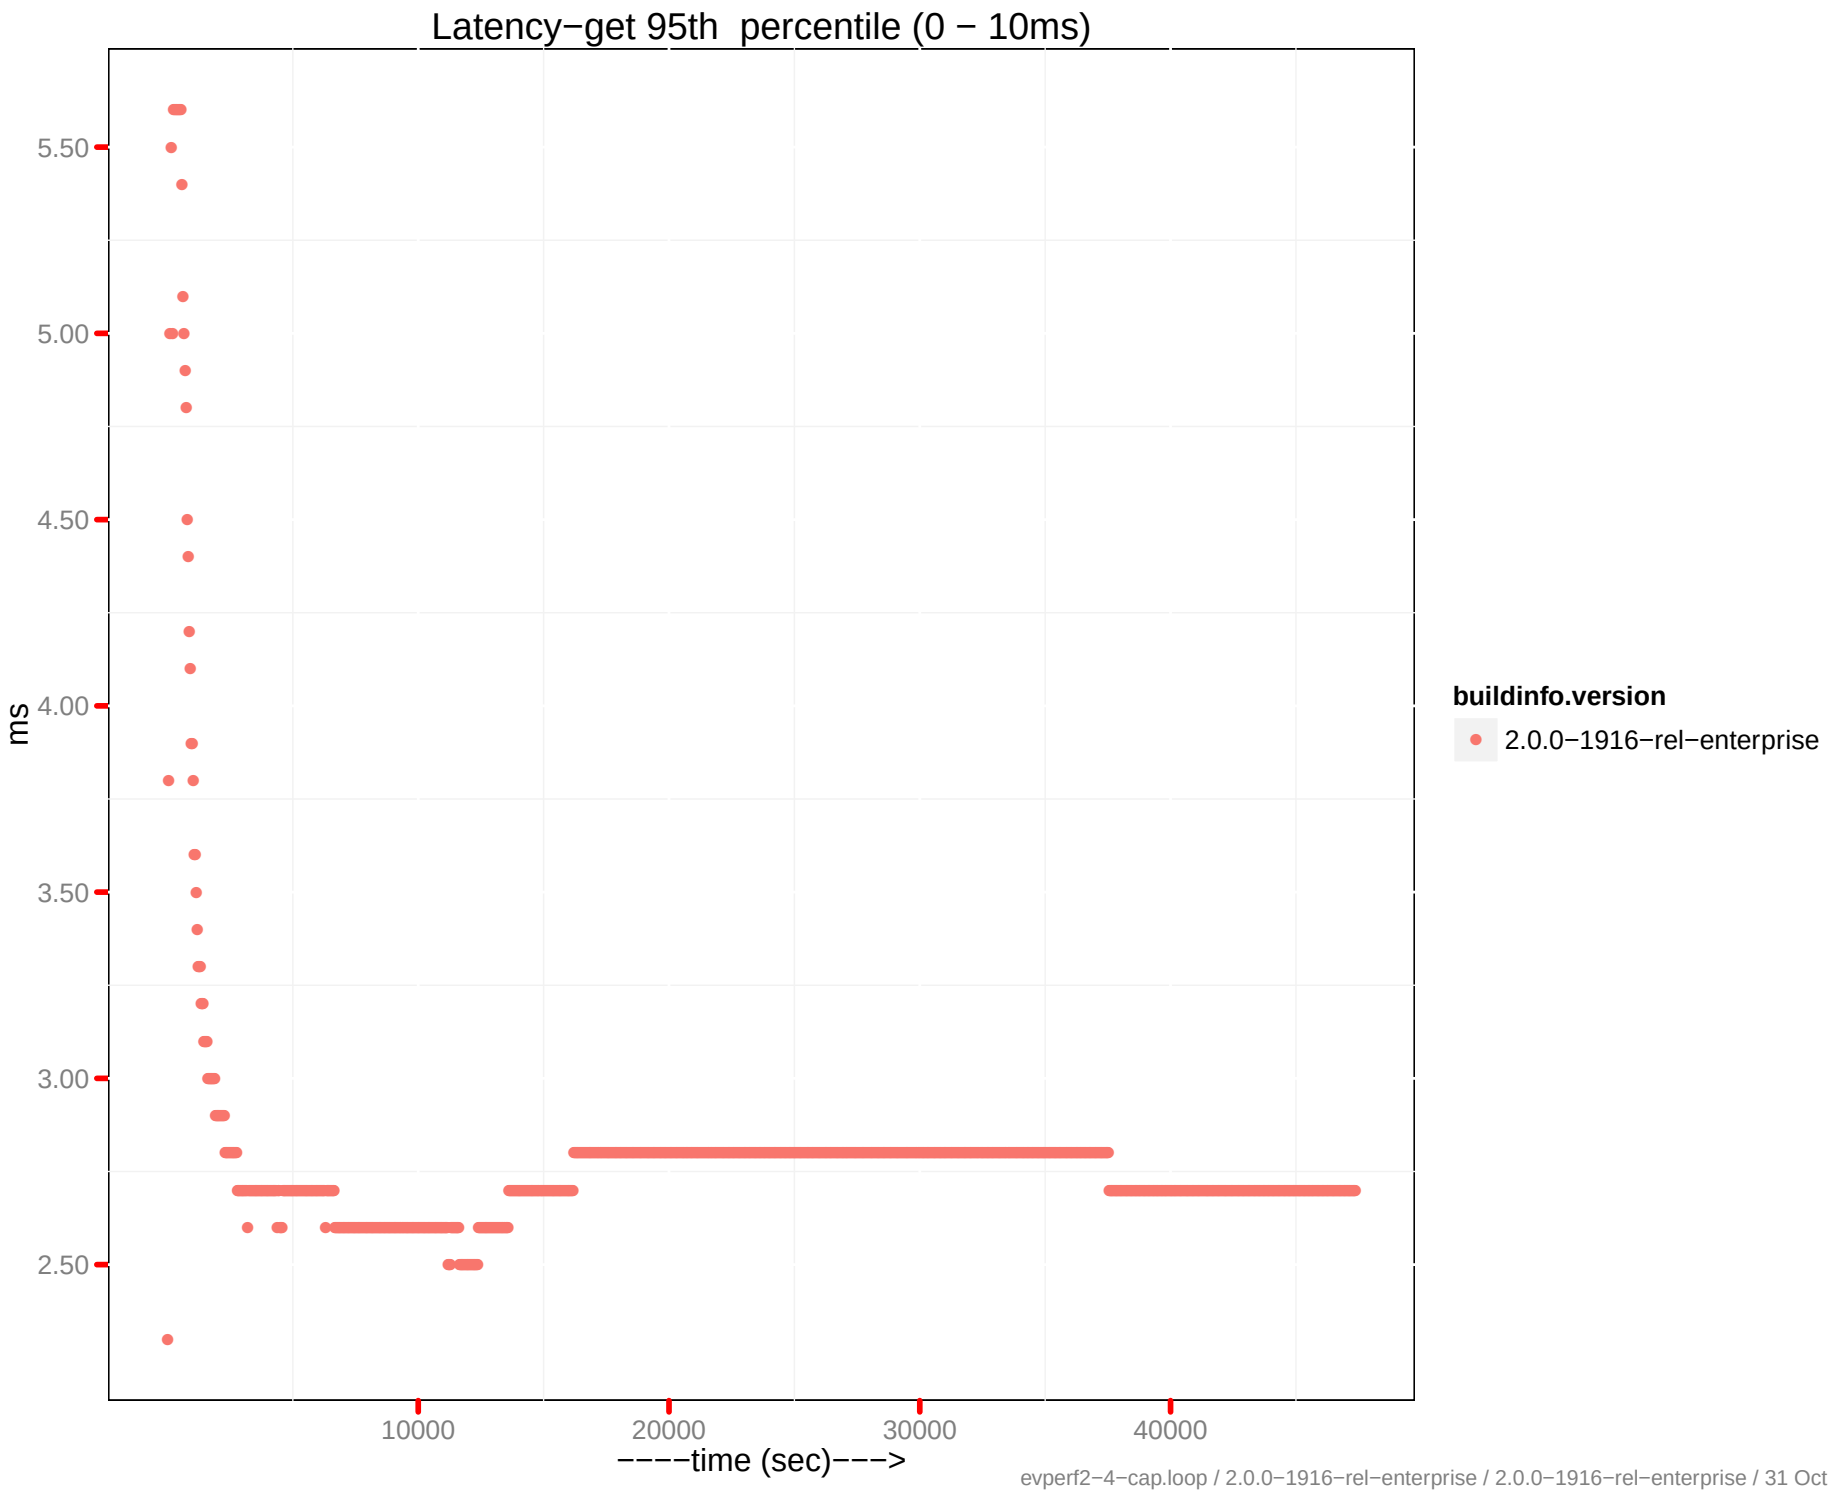
Latency−get 95th percentile (0 − 10ms)
5.50
5.00
4.50
4.00
3.50
3.00
2.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

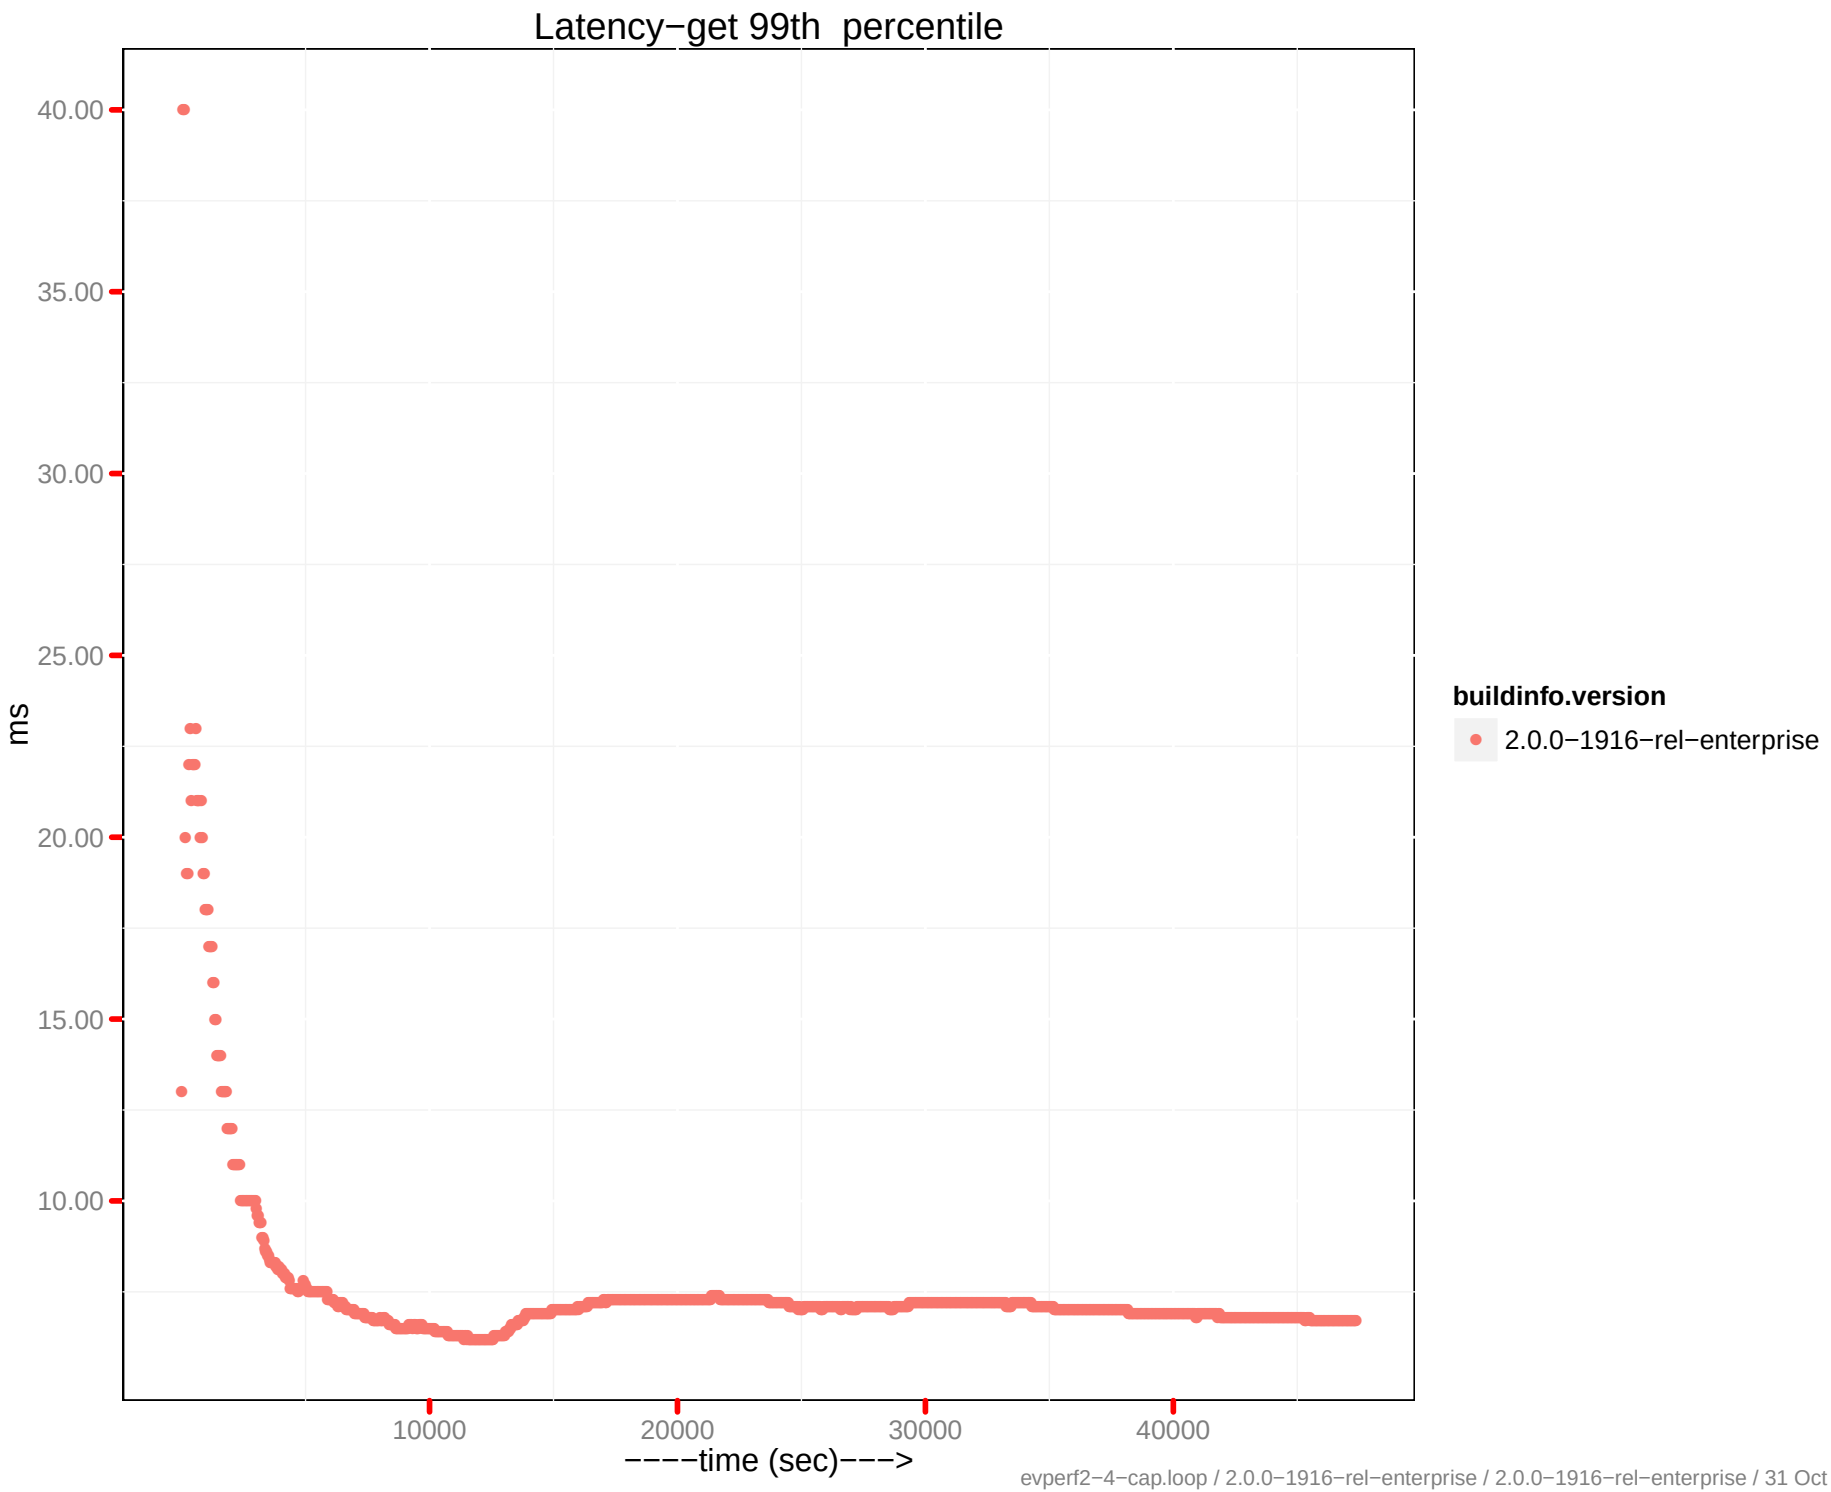
Latency−get 99th percentile
40.00
35.00
30.00
25.00
20.00
15.00
10.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

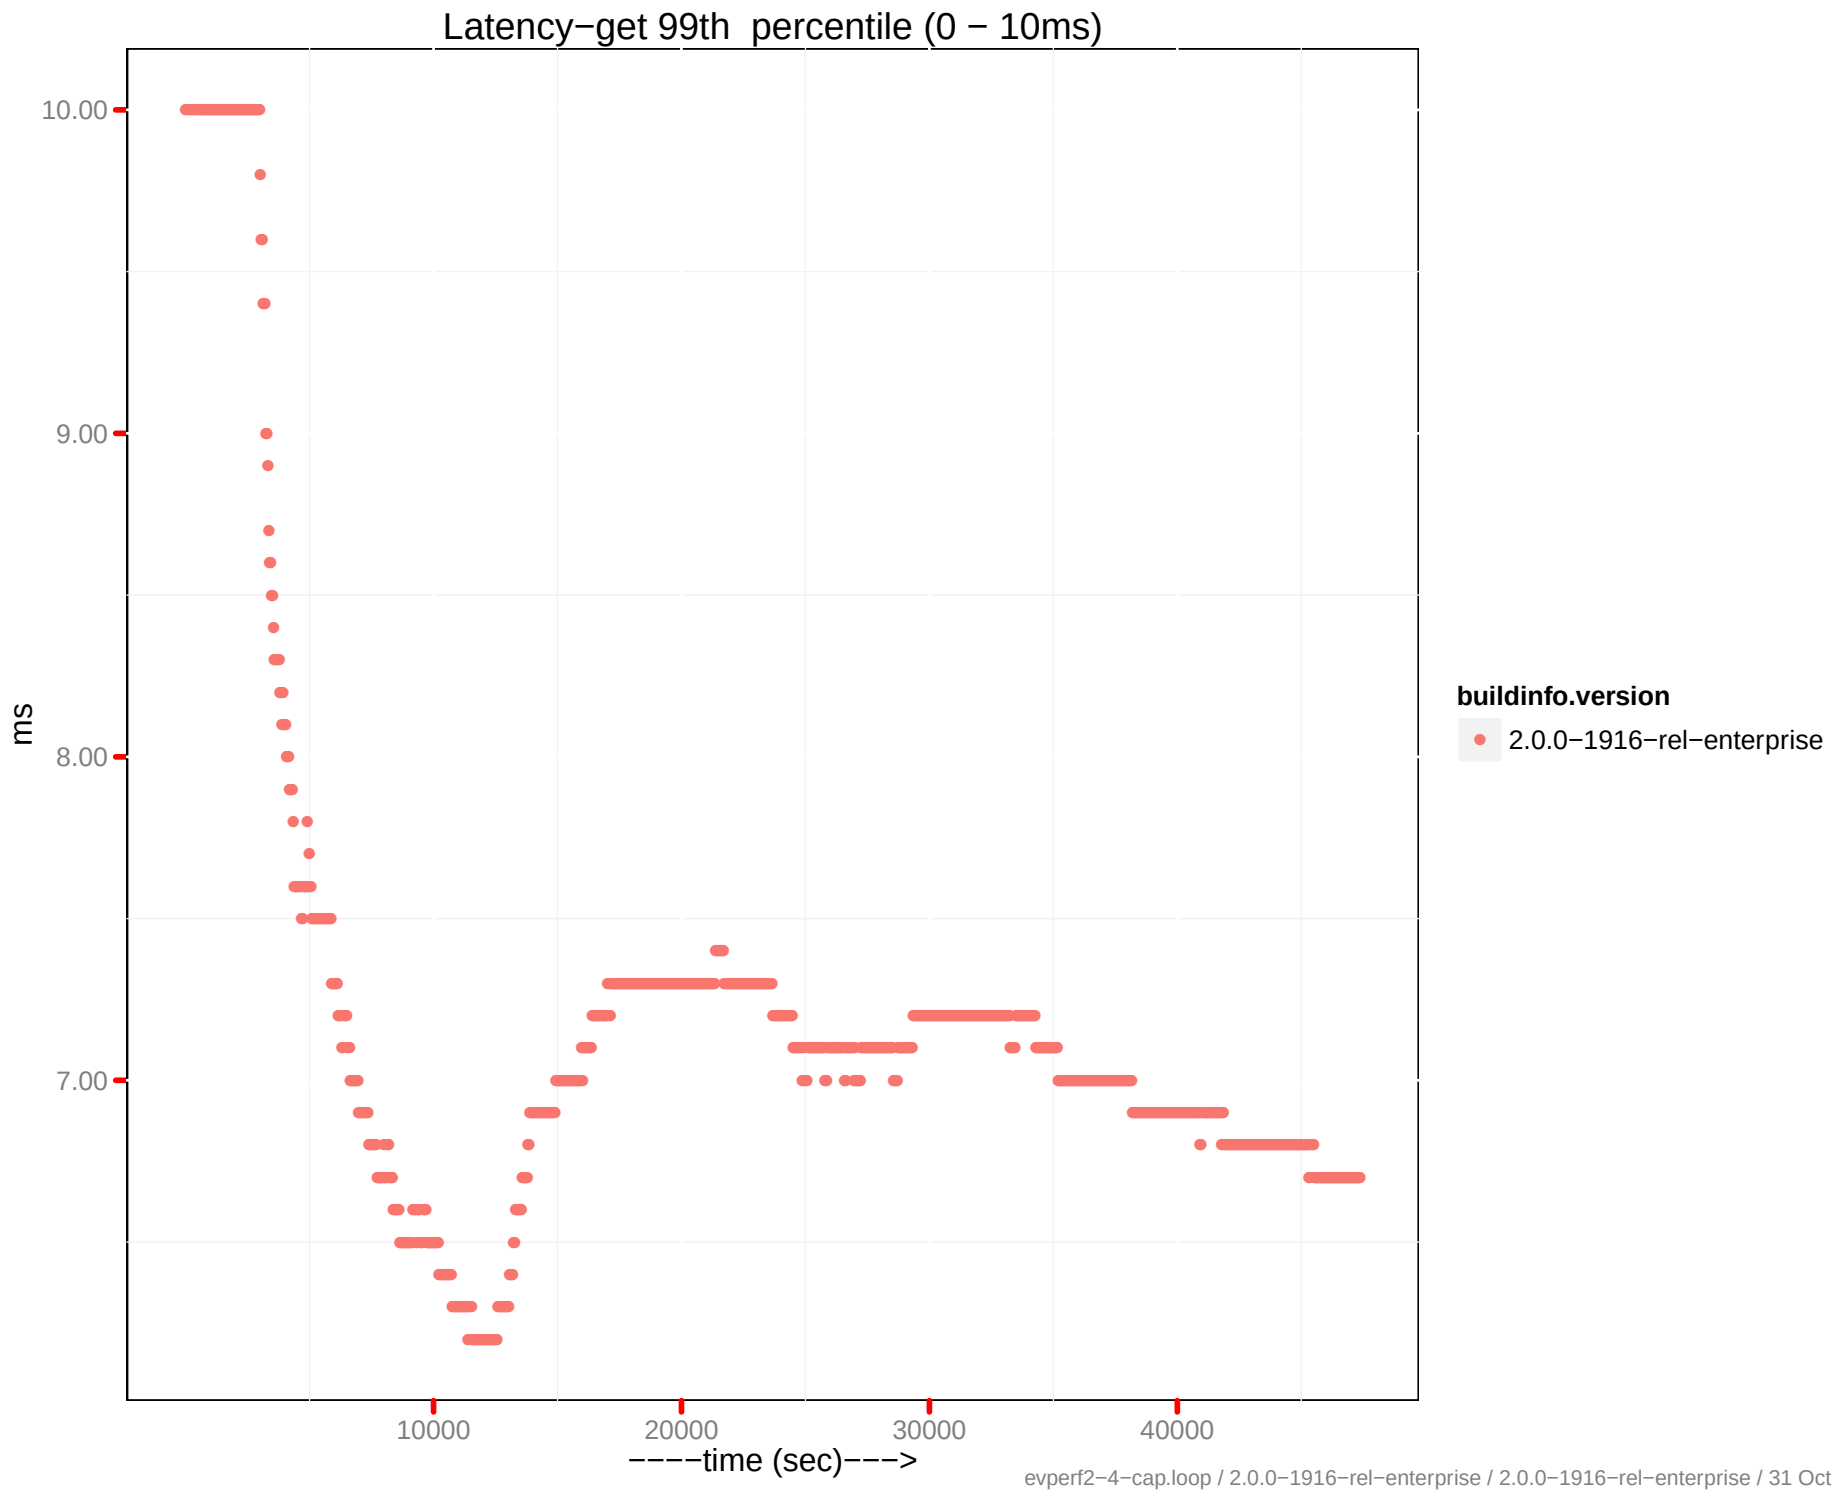
Latency−get 99th percentile (0 − 10ms)
10.00
9.00
8.00
7.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

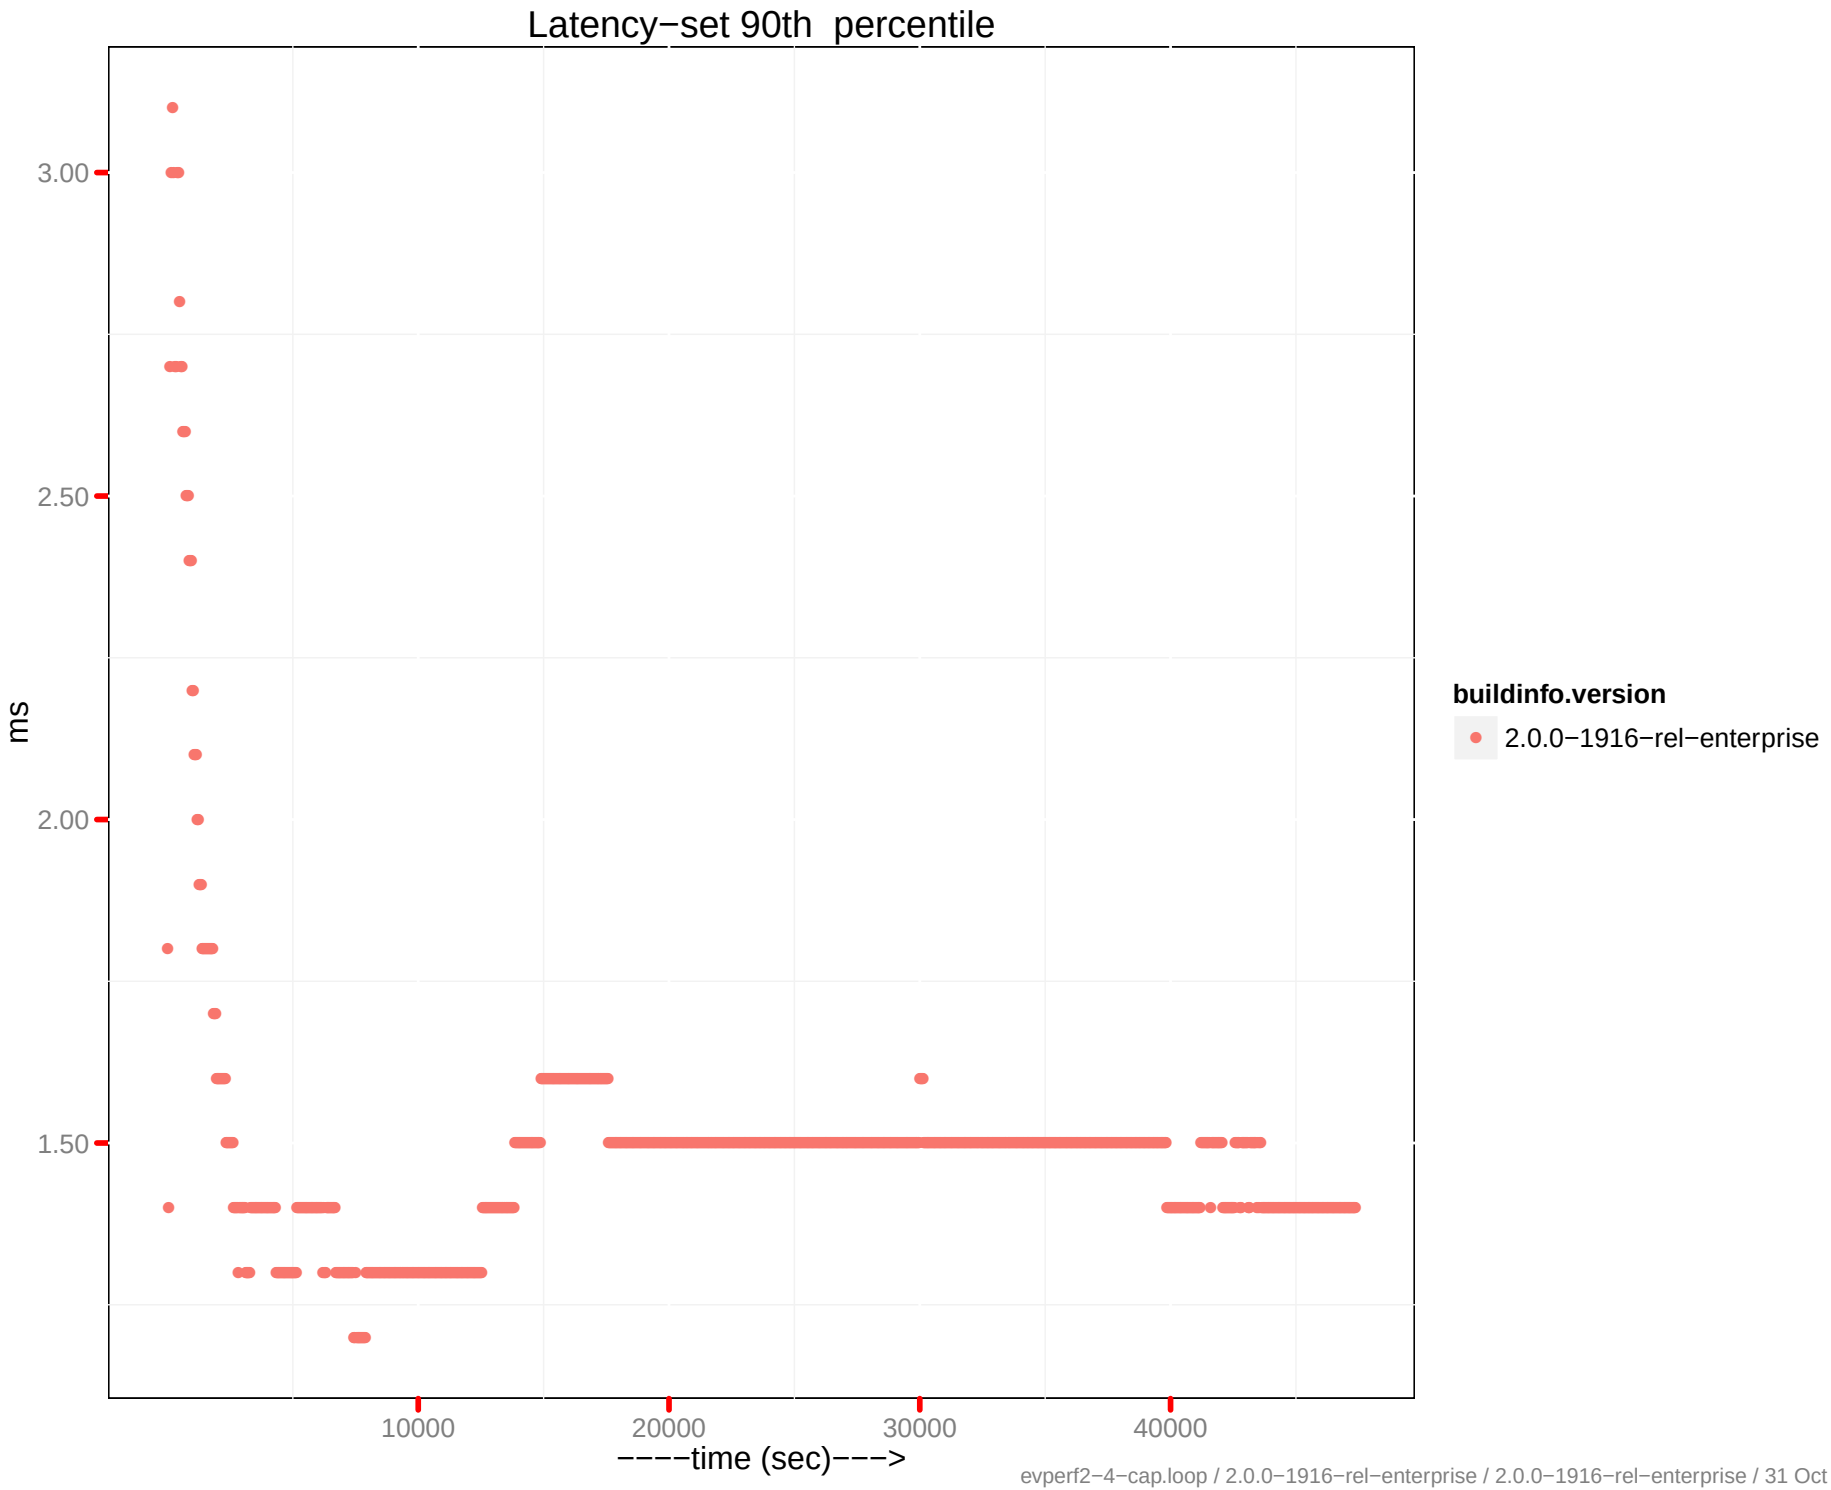
Latency−set 90th percentile
3.00
2.50
2.00
1.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

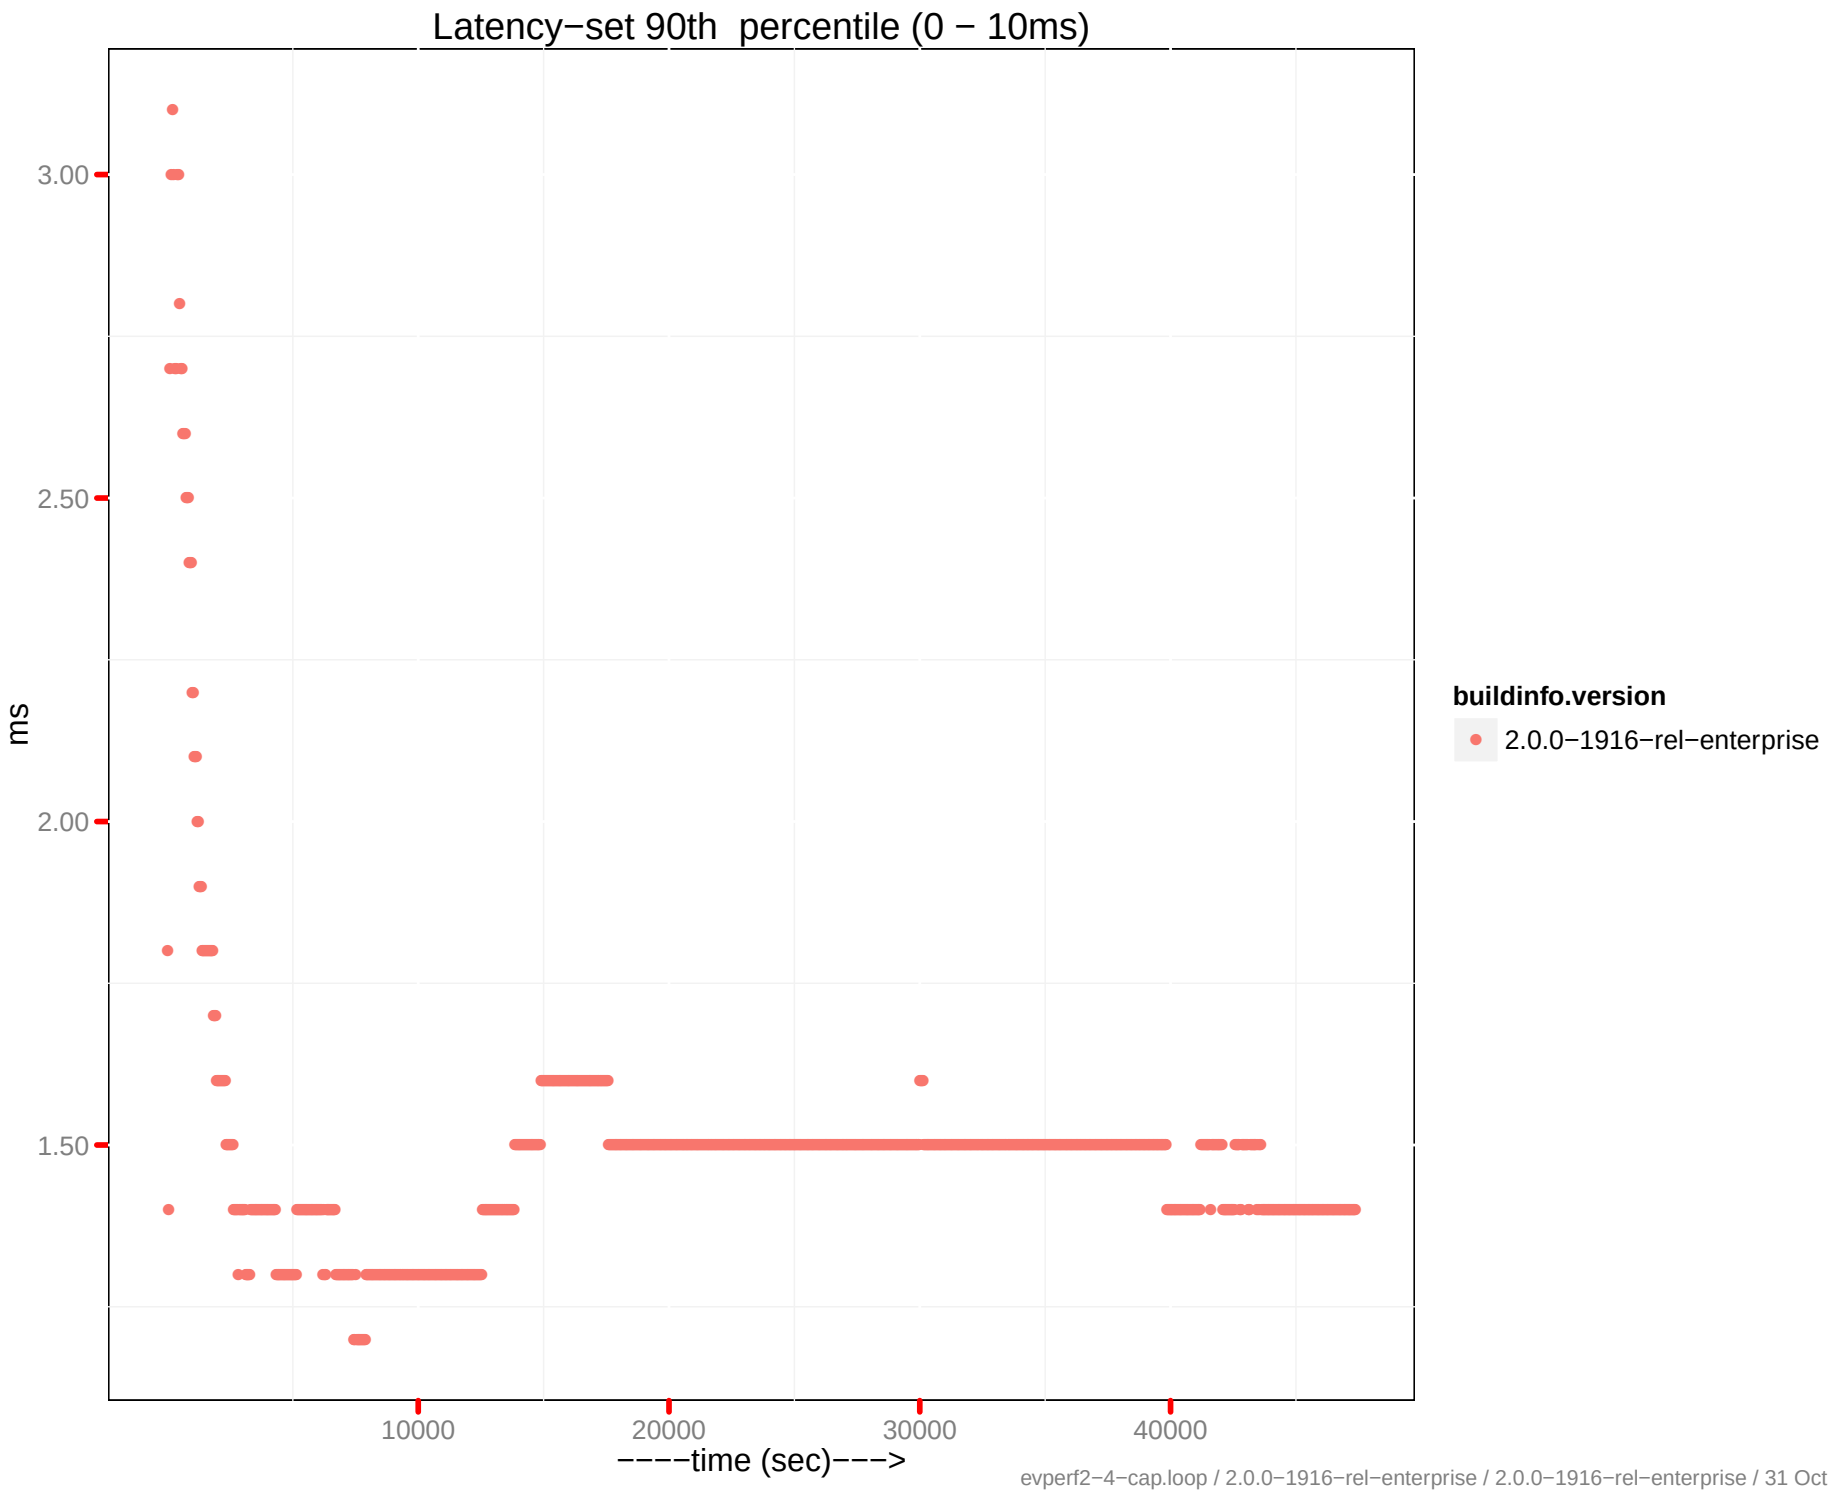
Latency−set 90th percentile (0 − 10ms)
3.00
2.50
2.00
1.50
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

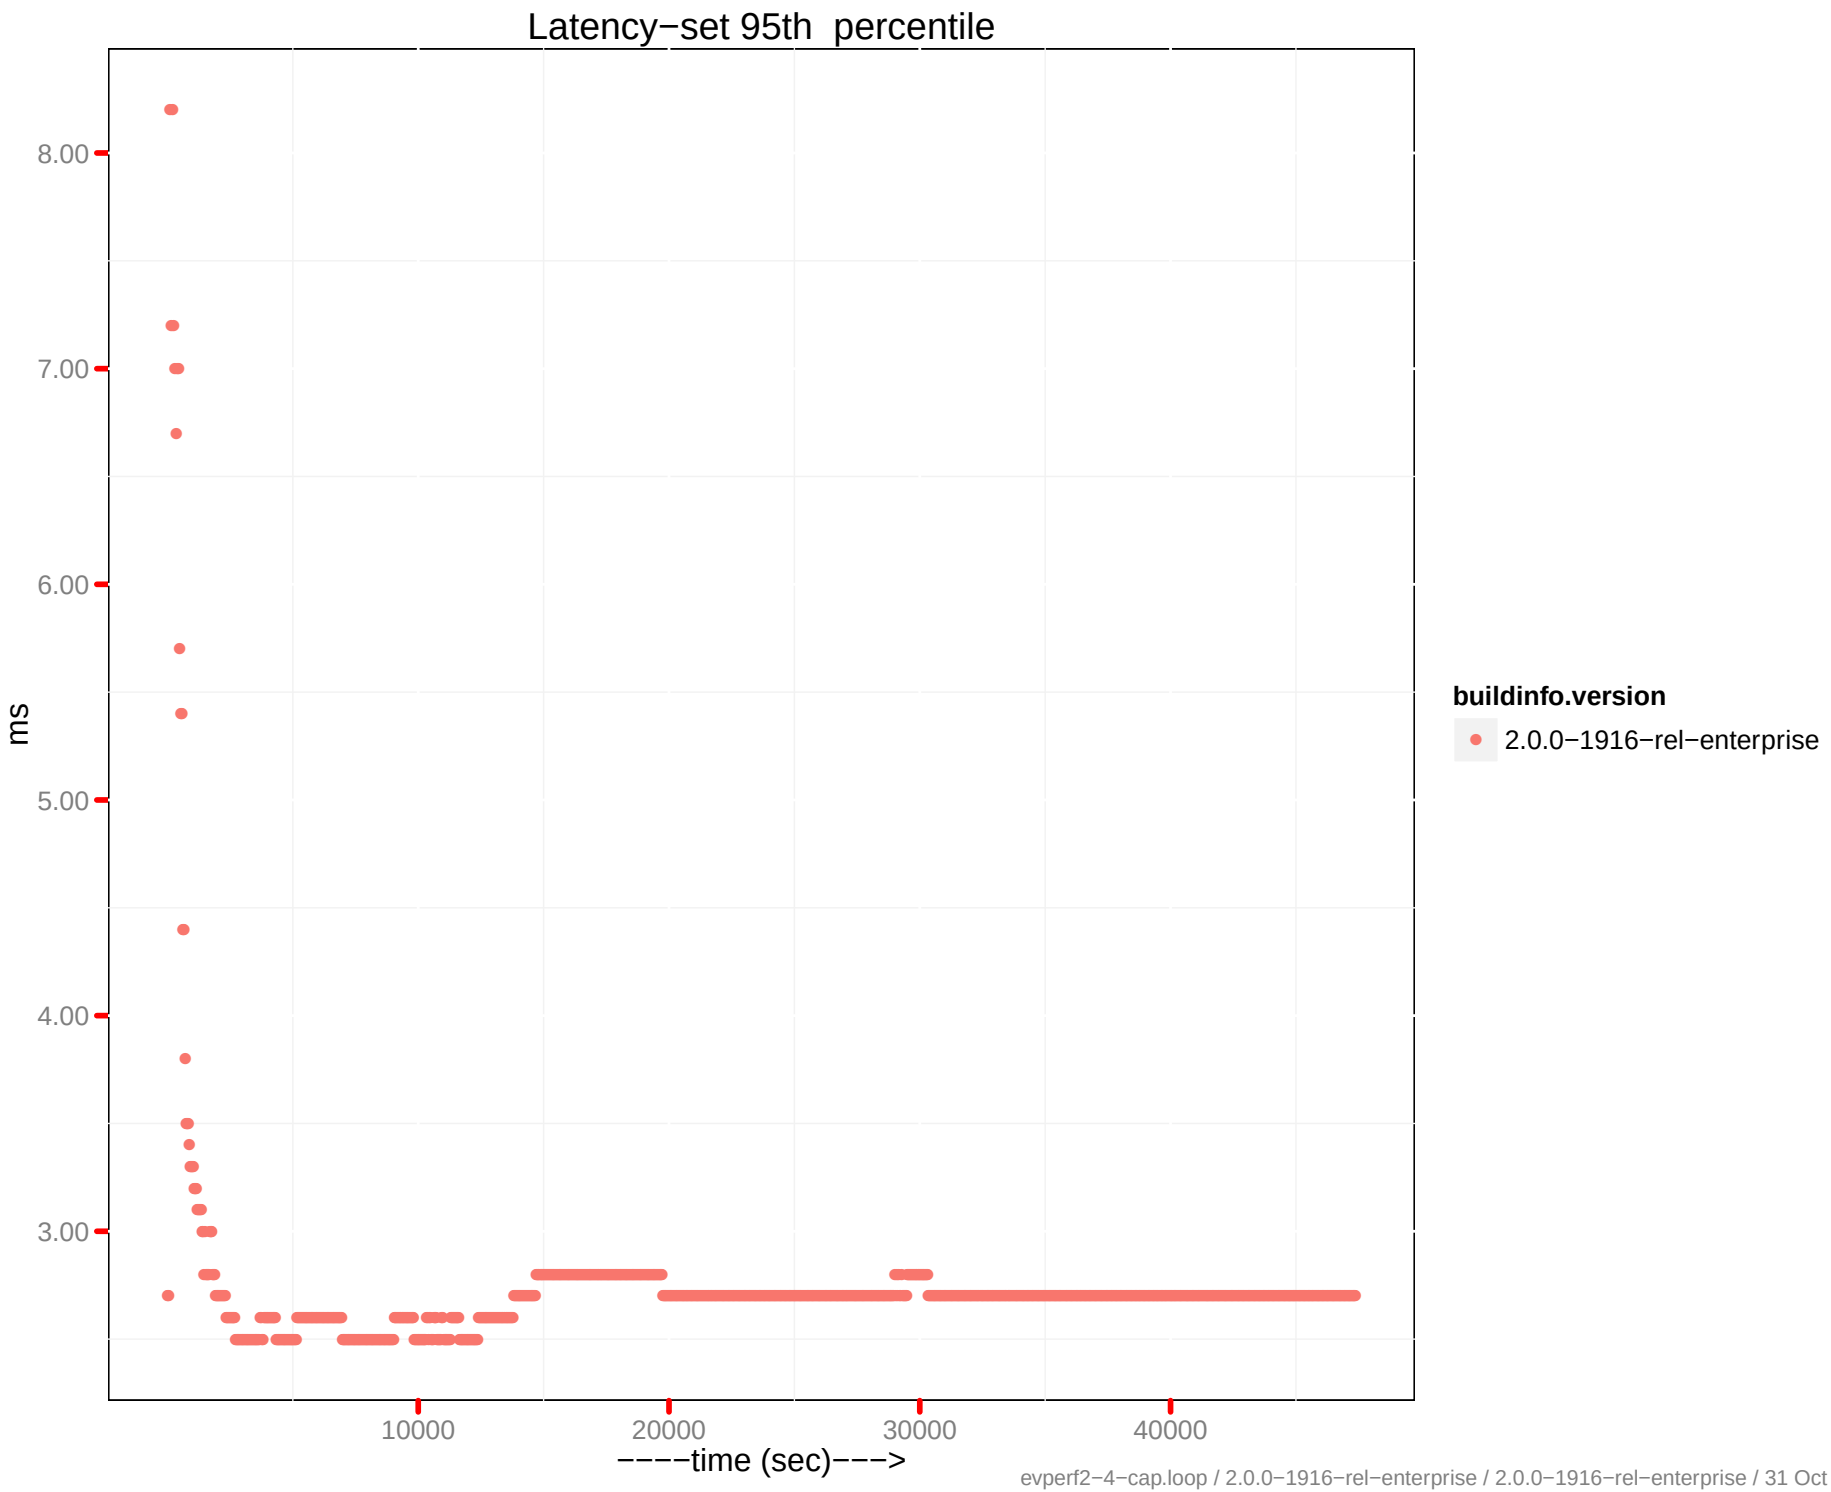
Latency−set 95th percentile
8.00
7.00
6.00
5.00
4.00
3.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

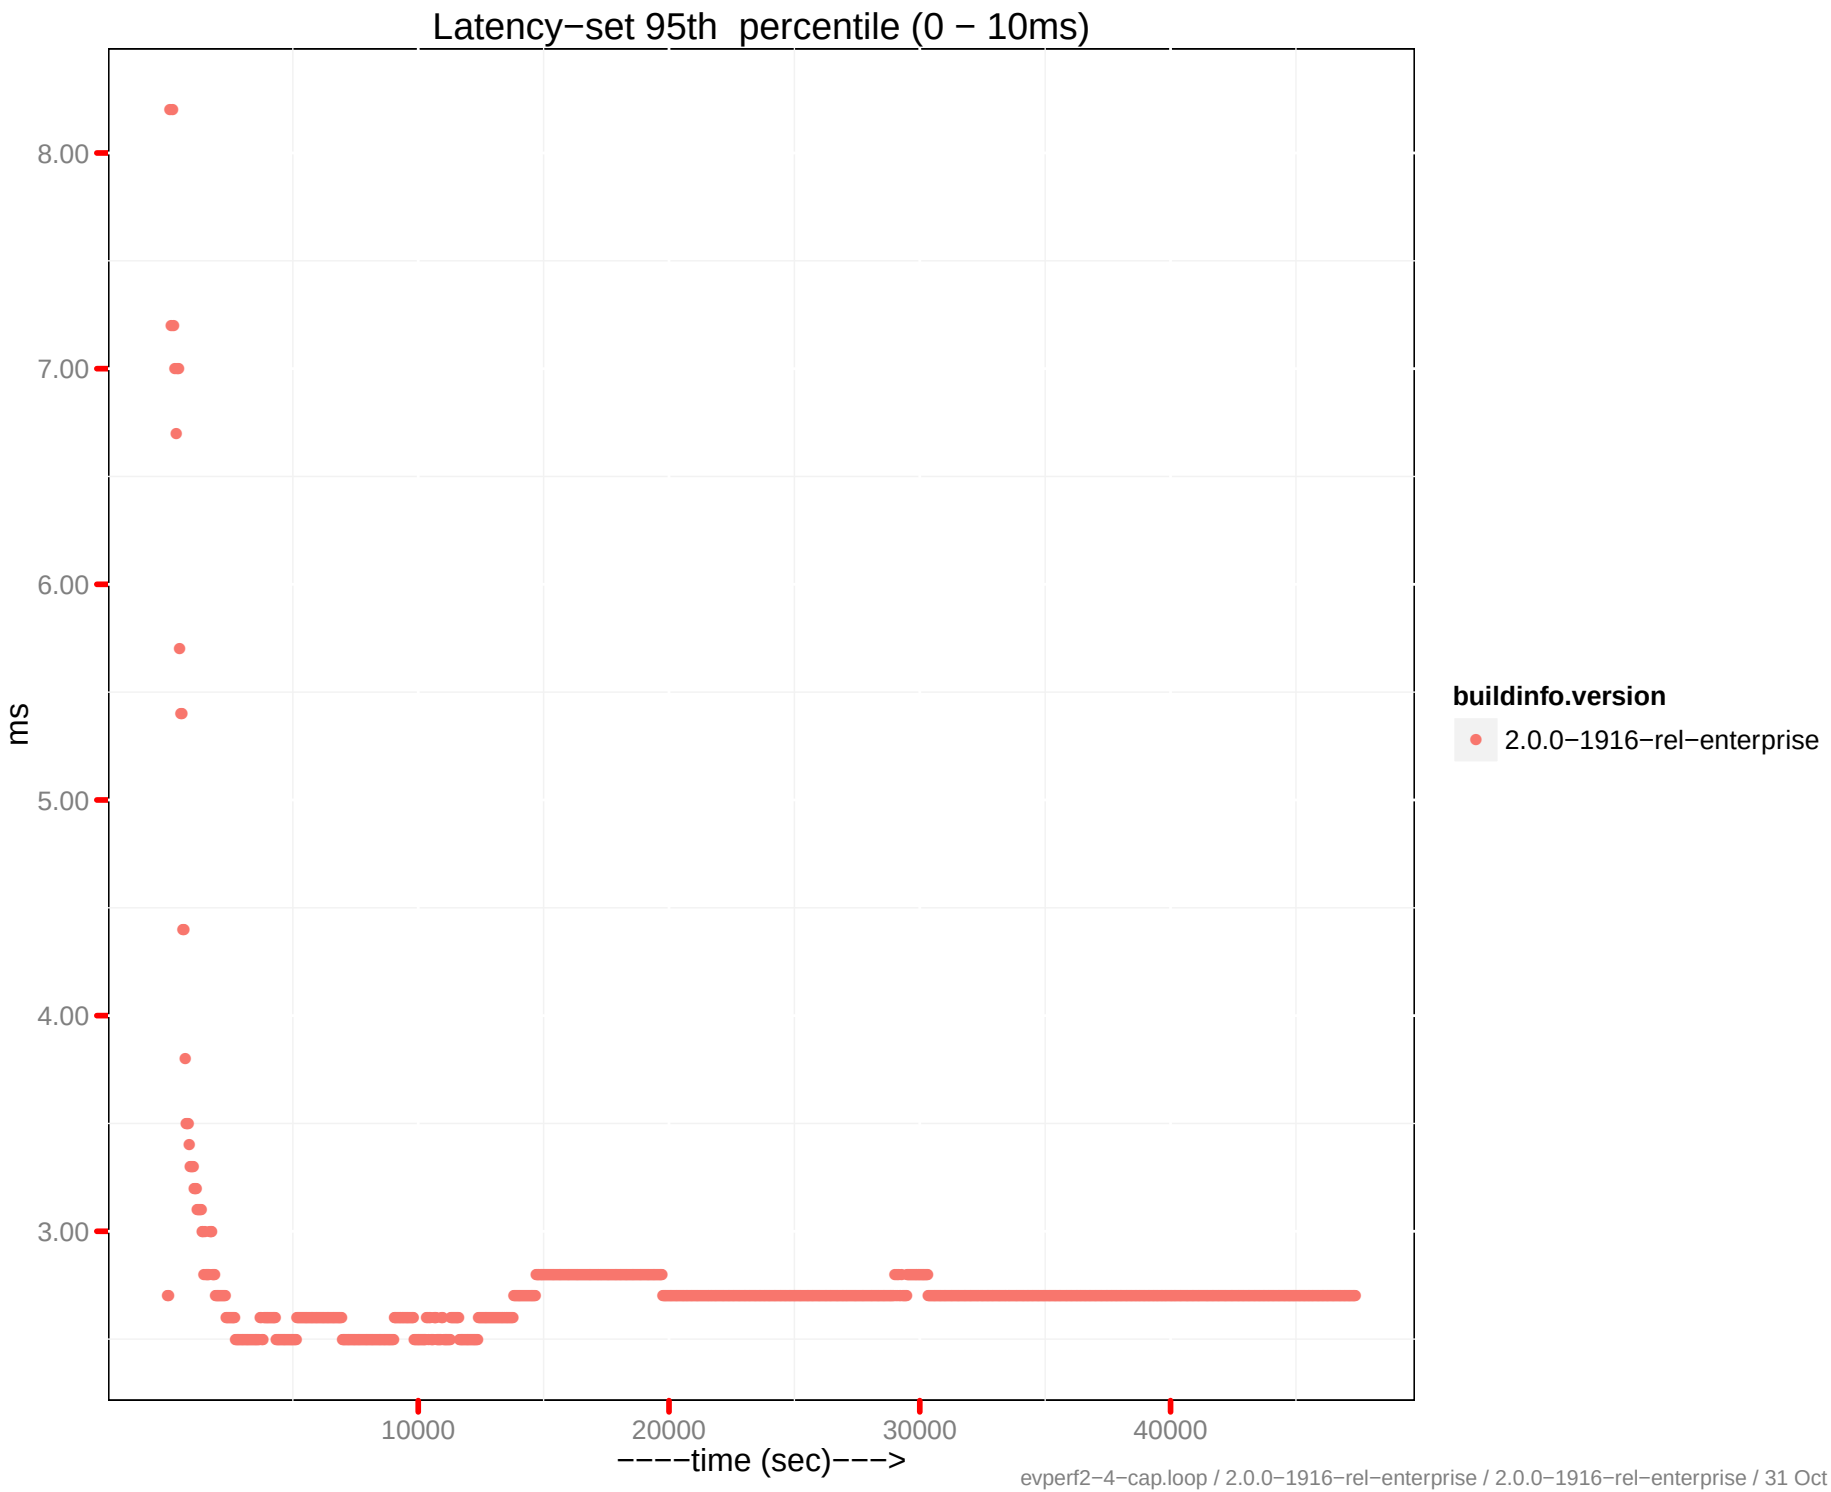
Latency−set 95th percentile (0 − 10ms)
8.00
7.00
6.00
5.00
4.00
3.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

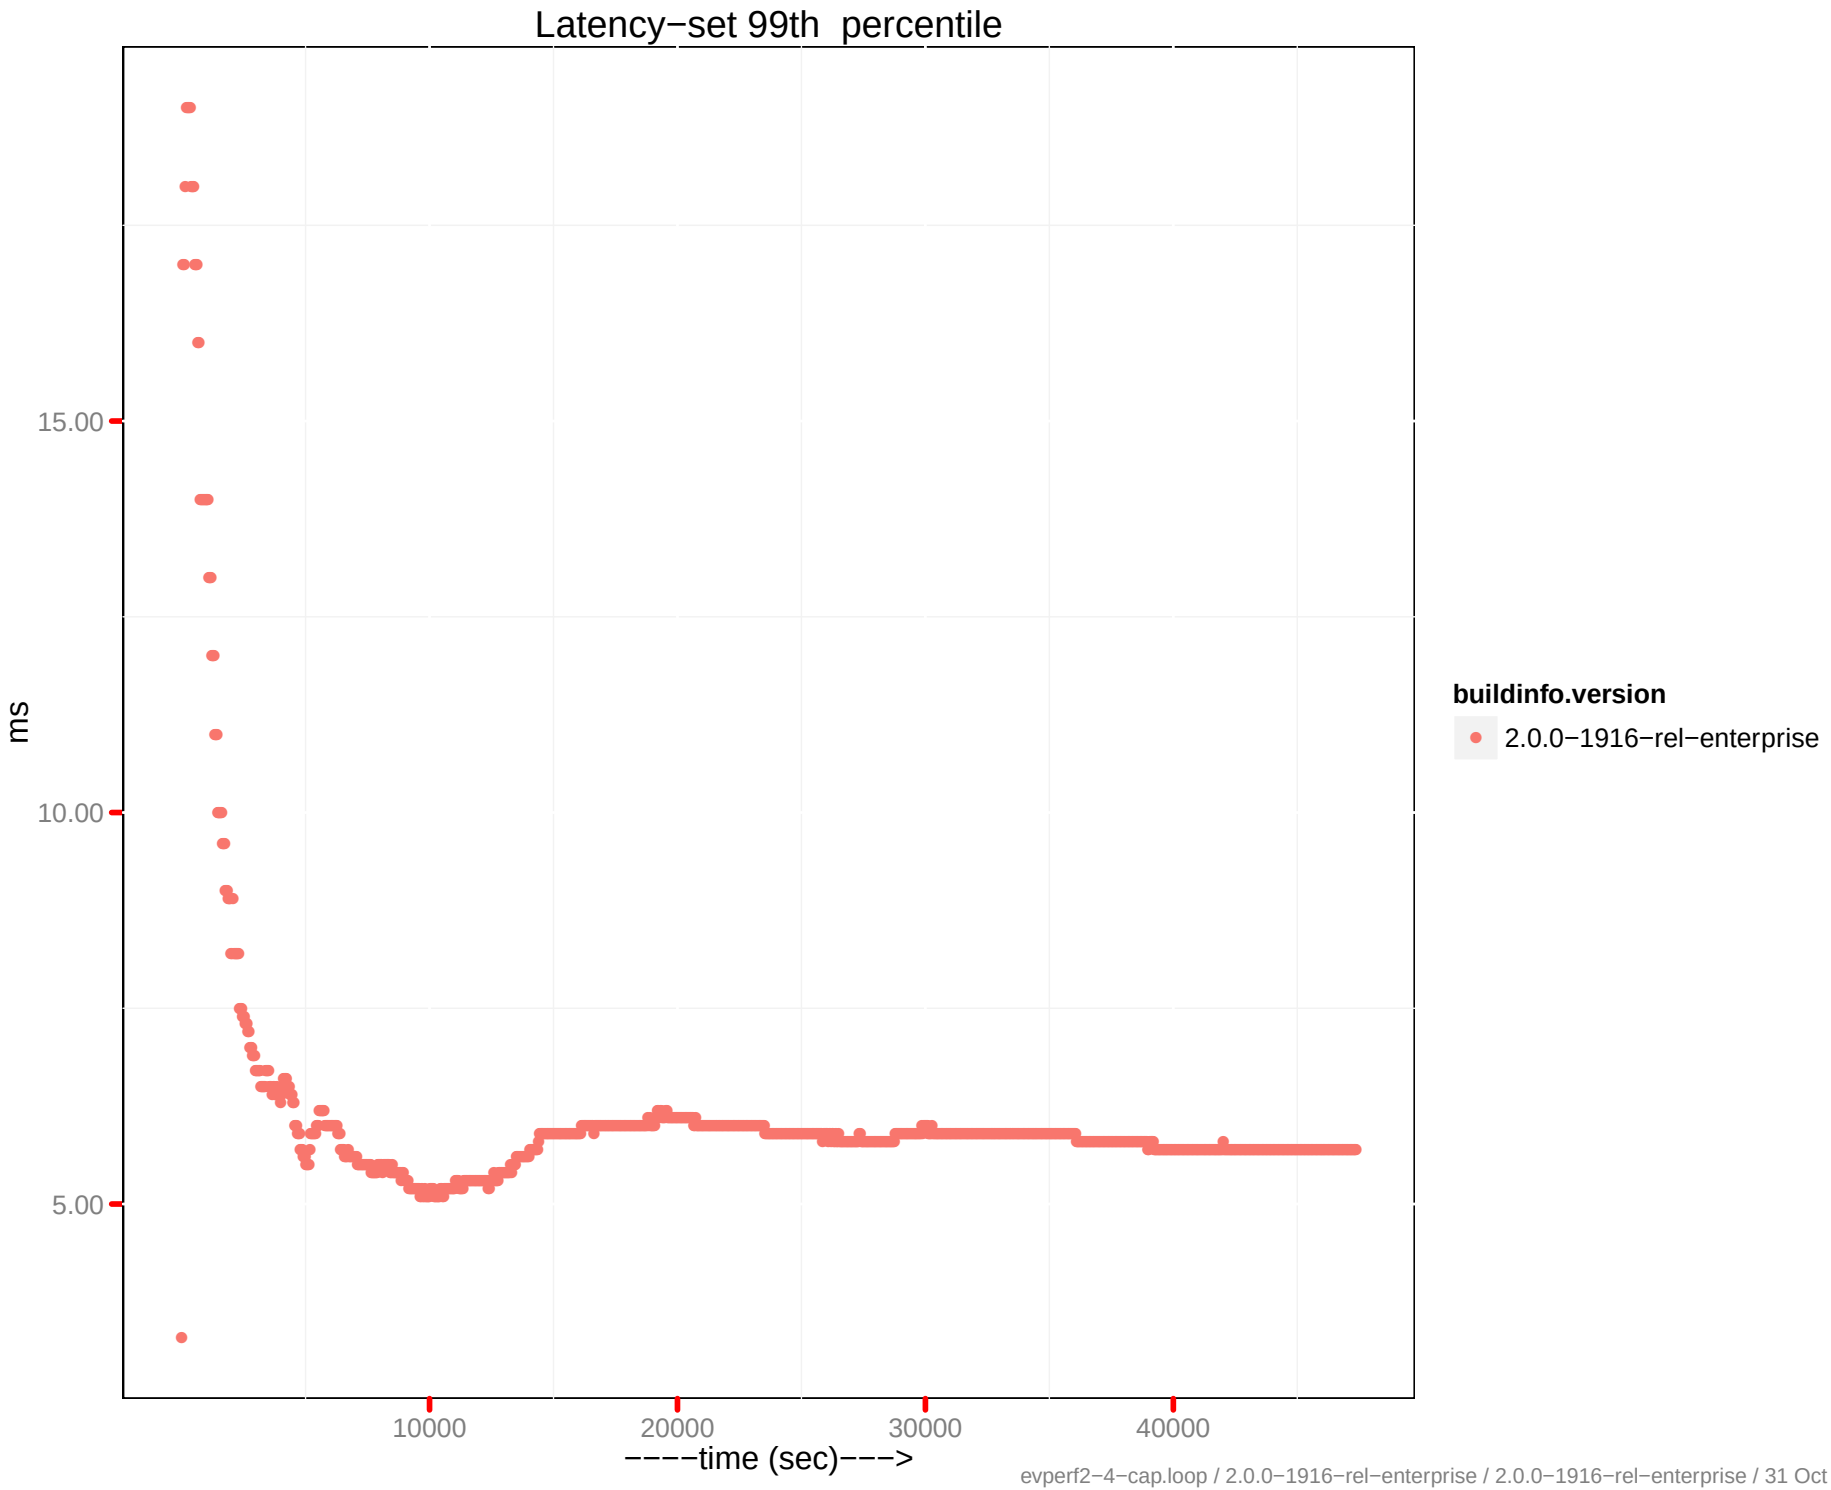
Latency−set 99th percentile
15.00
10.00
5.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

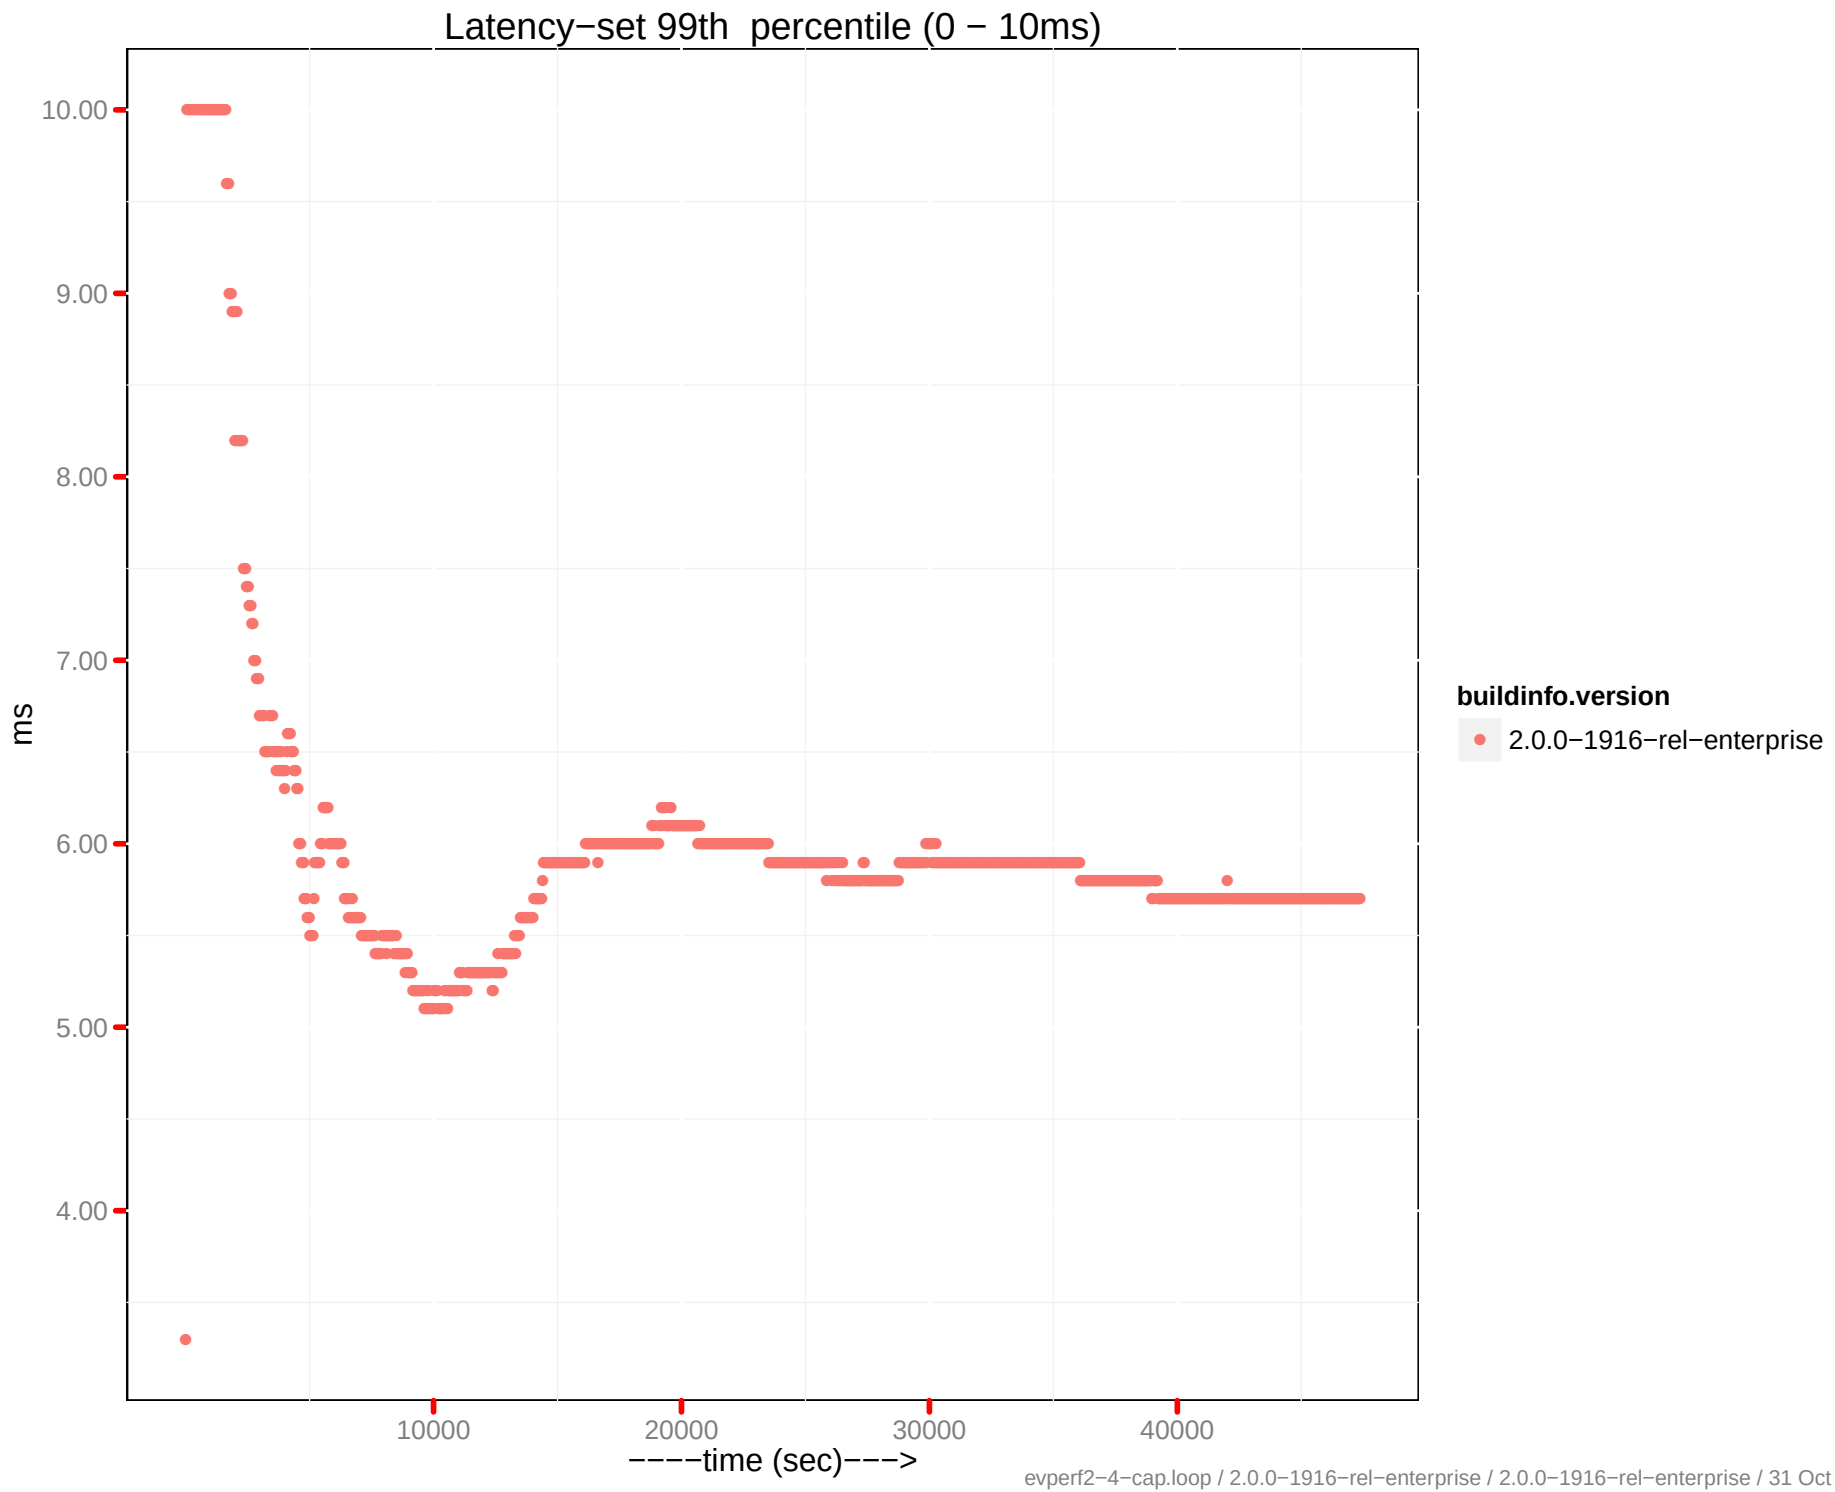
Latency−set 99th percentile (0 − 10ms)
10.00
9.00
8.00
7.00
6.00
5.00
4.00
ms
10000
20000
30000
40000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

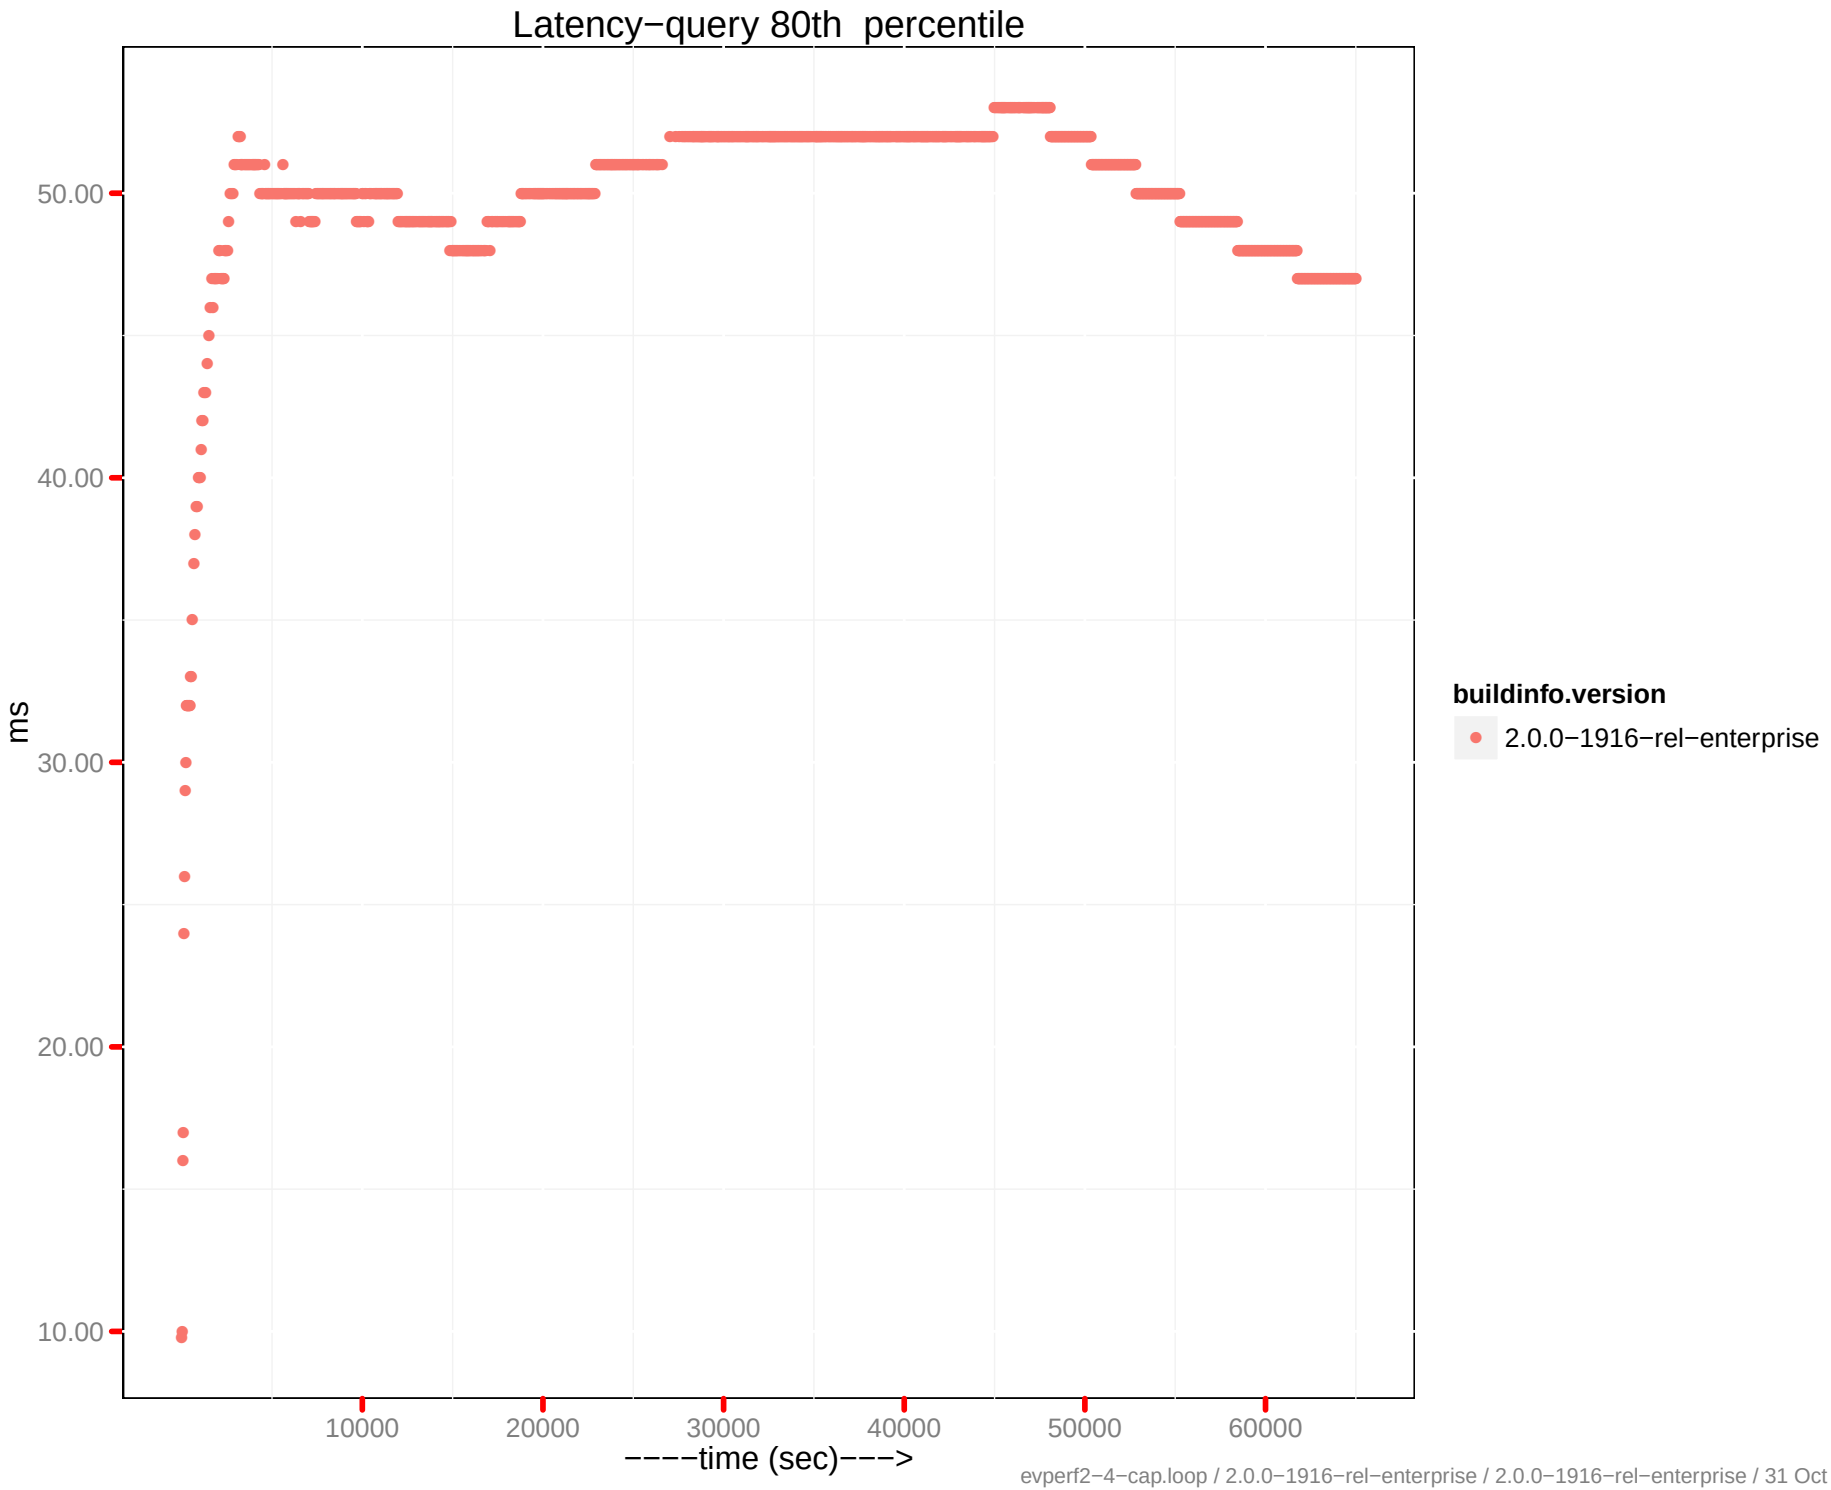
Latency−query 80th percentile
50.00
40.00
30.00
20.00
10.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

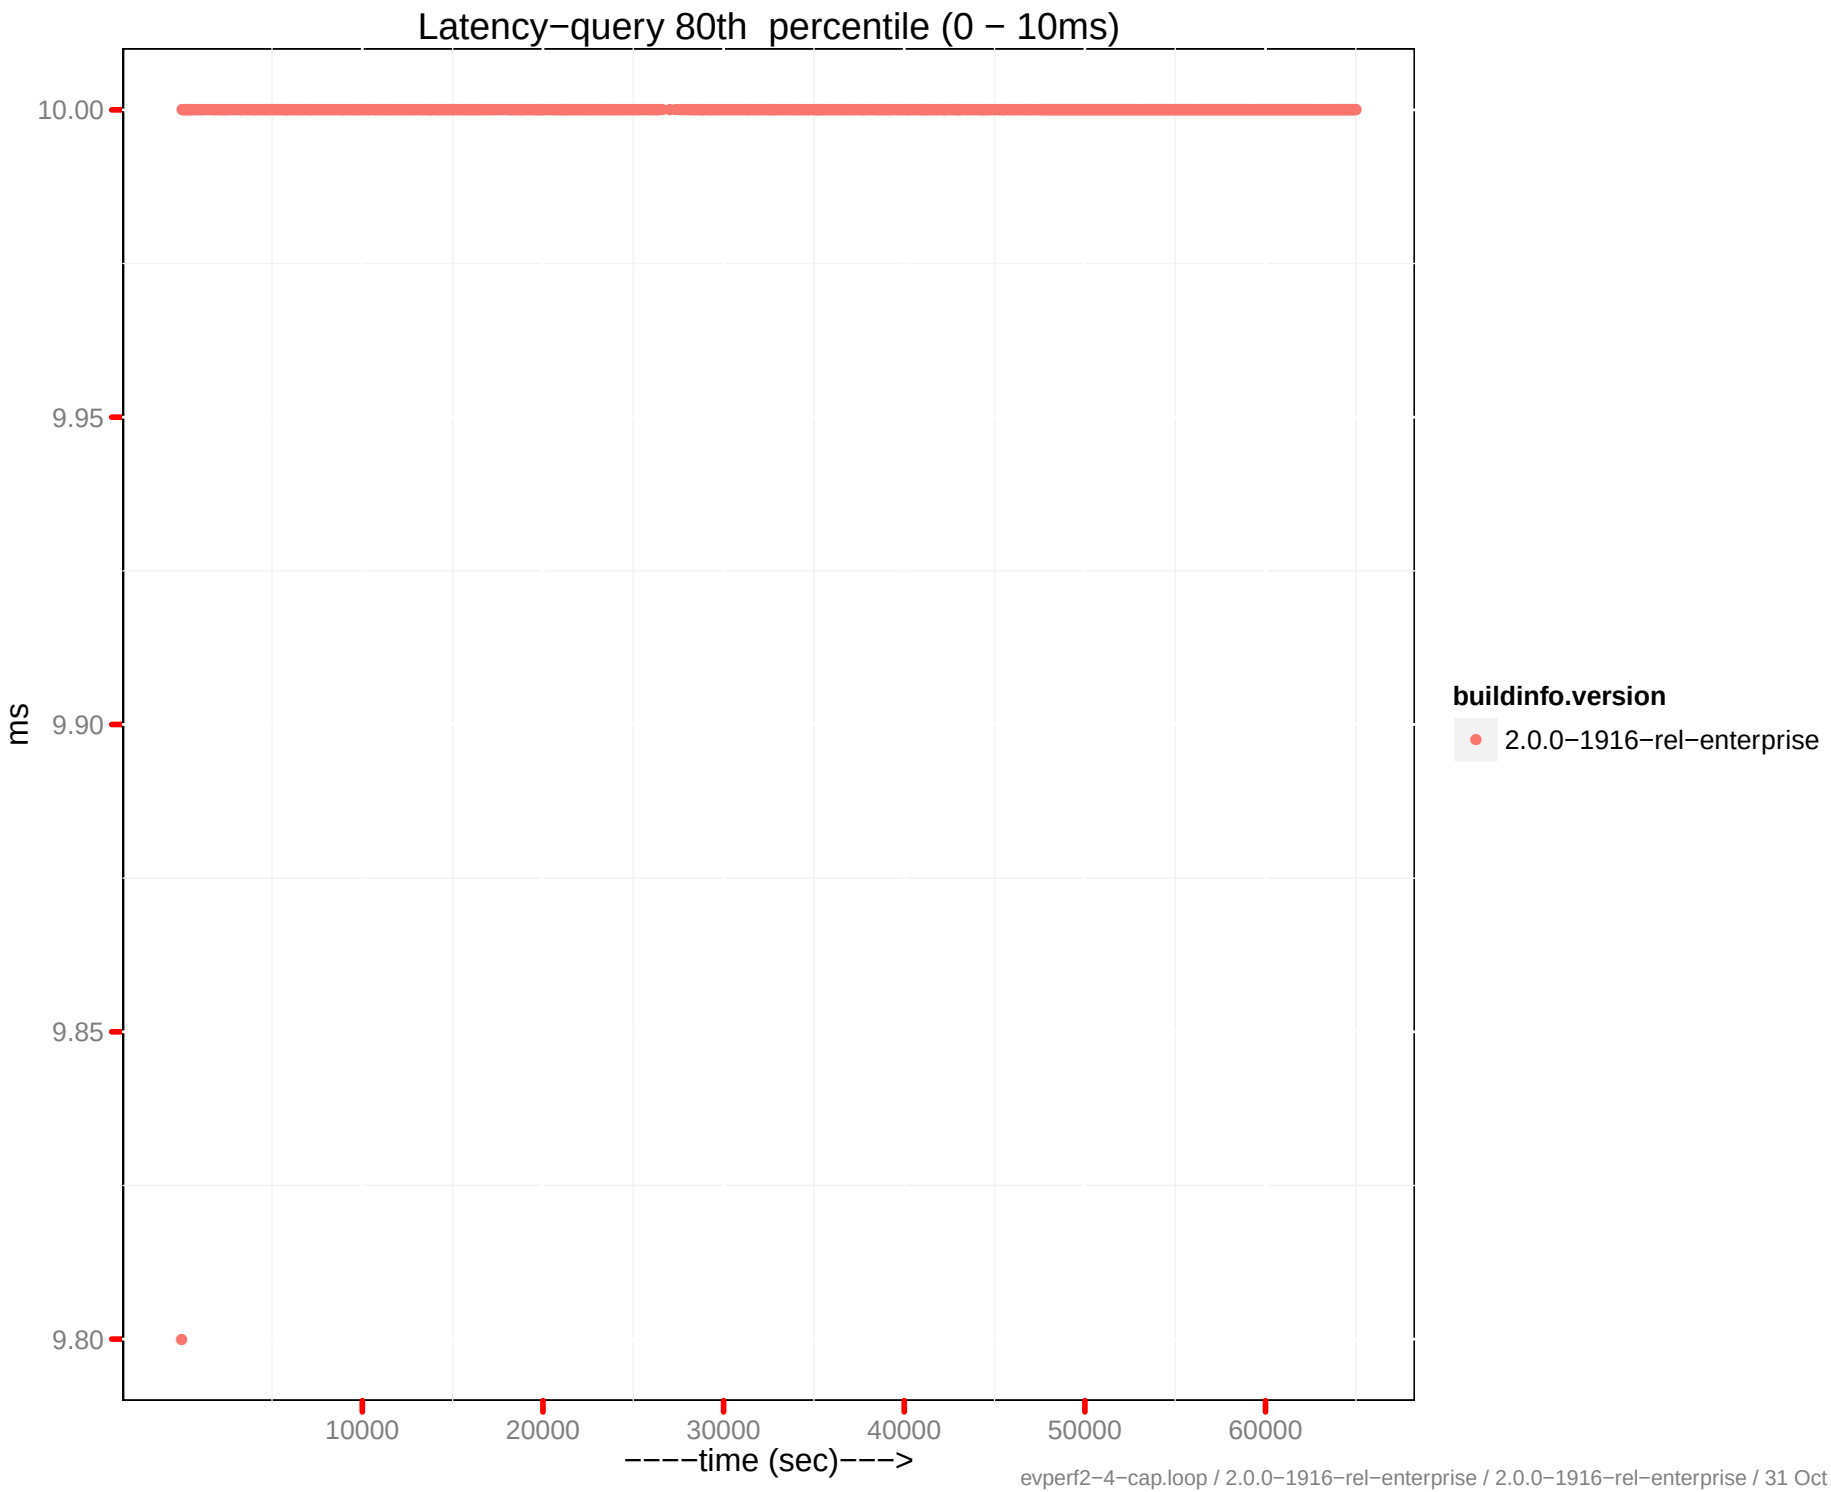
Latency−query 80th percentile (0 − 10ms)
10.00
9.95
9.90
9.85
9.80
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

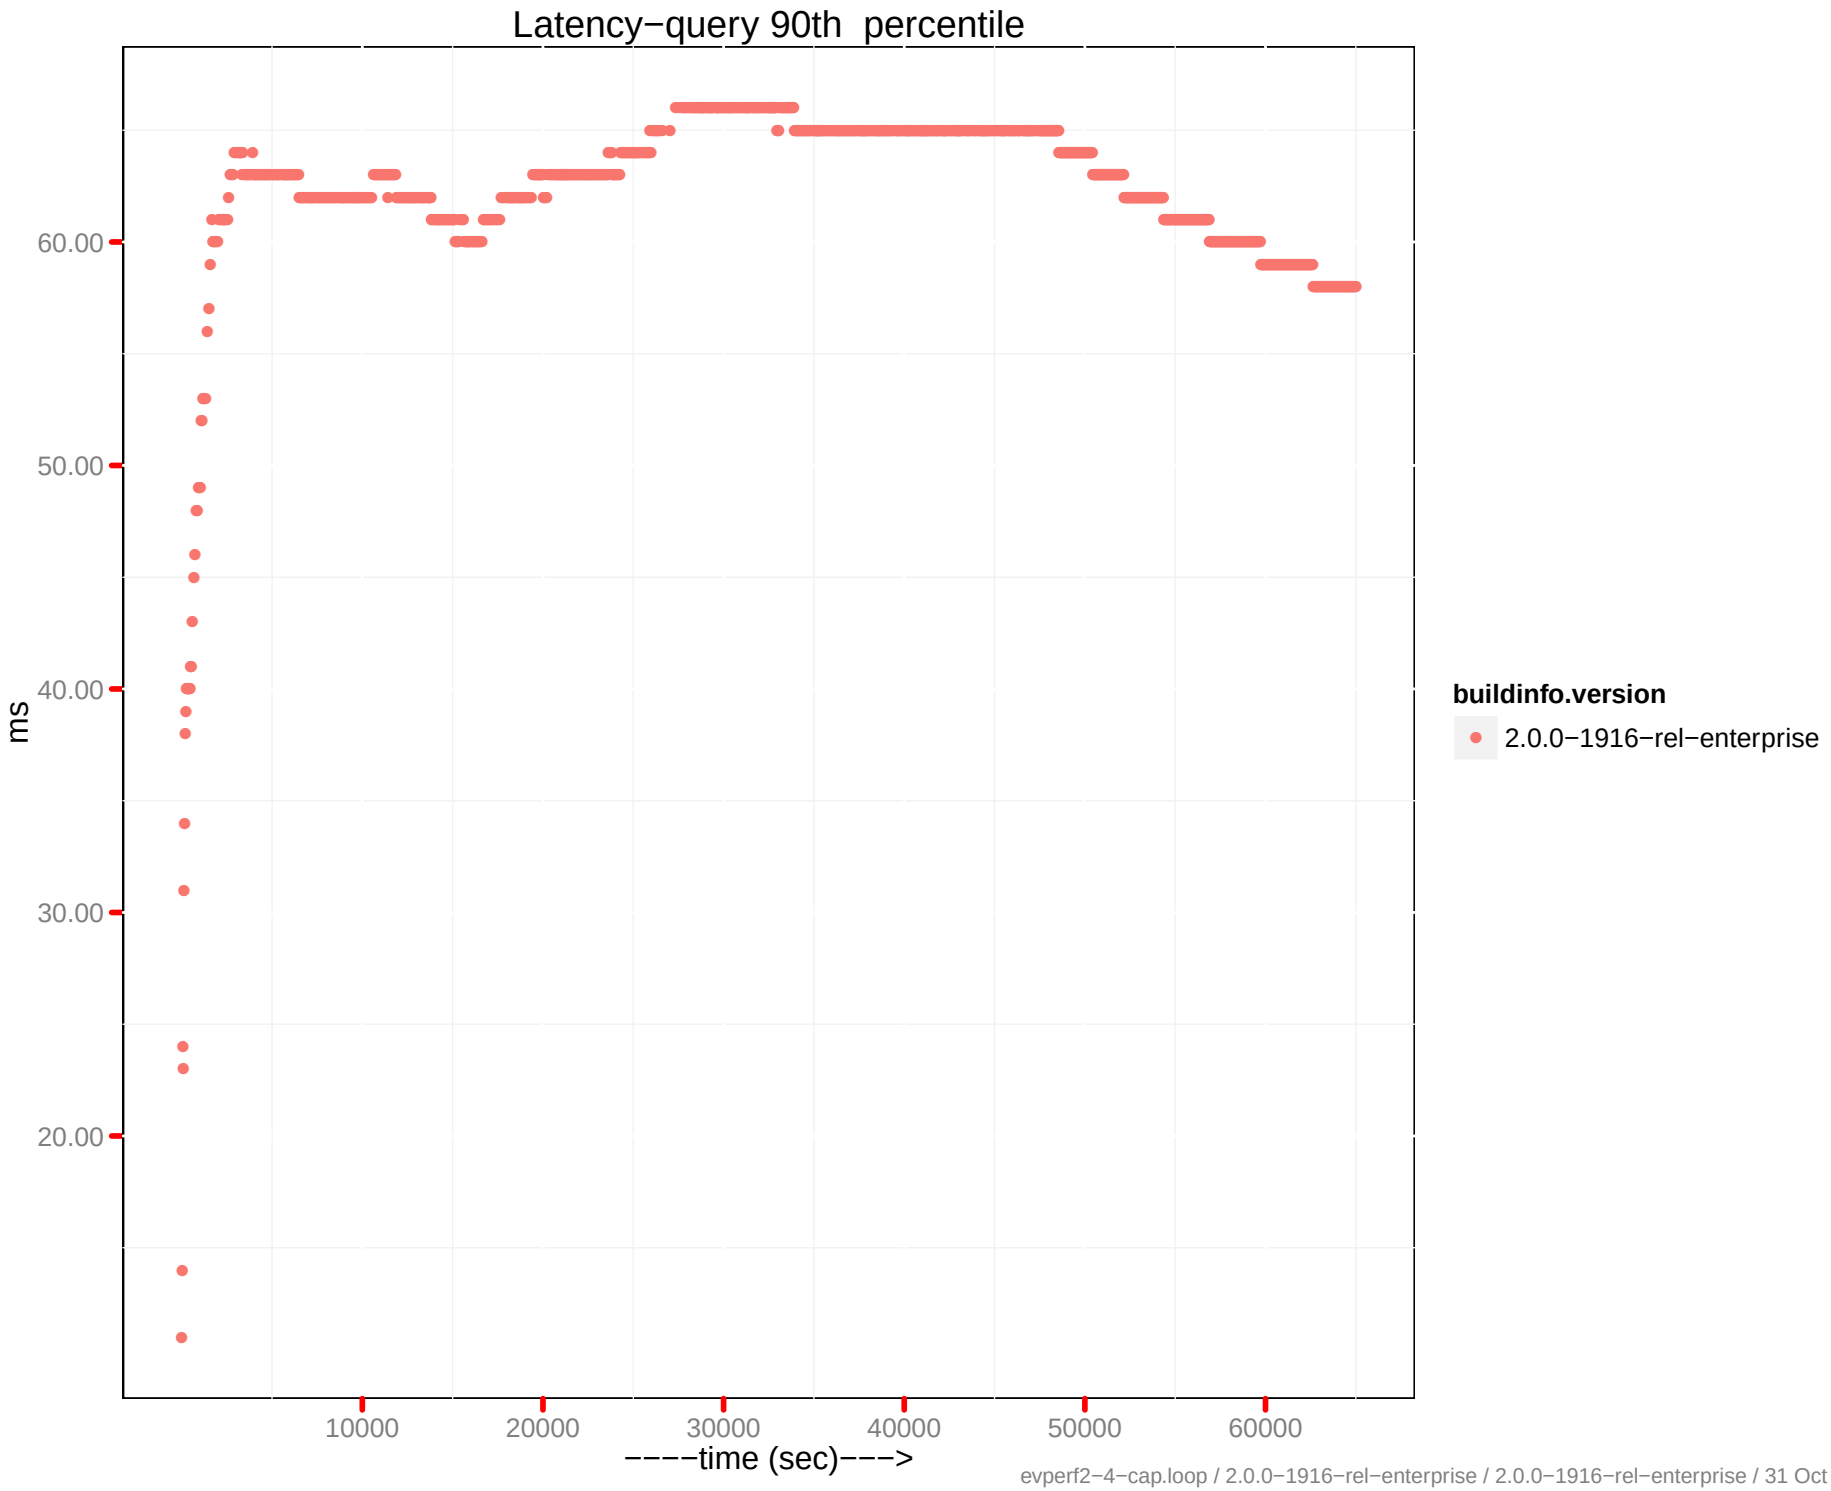
Latency−query 90th percentile
60.00
50.00
40.00
30.00
20.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

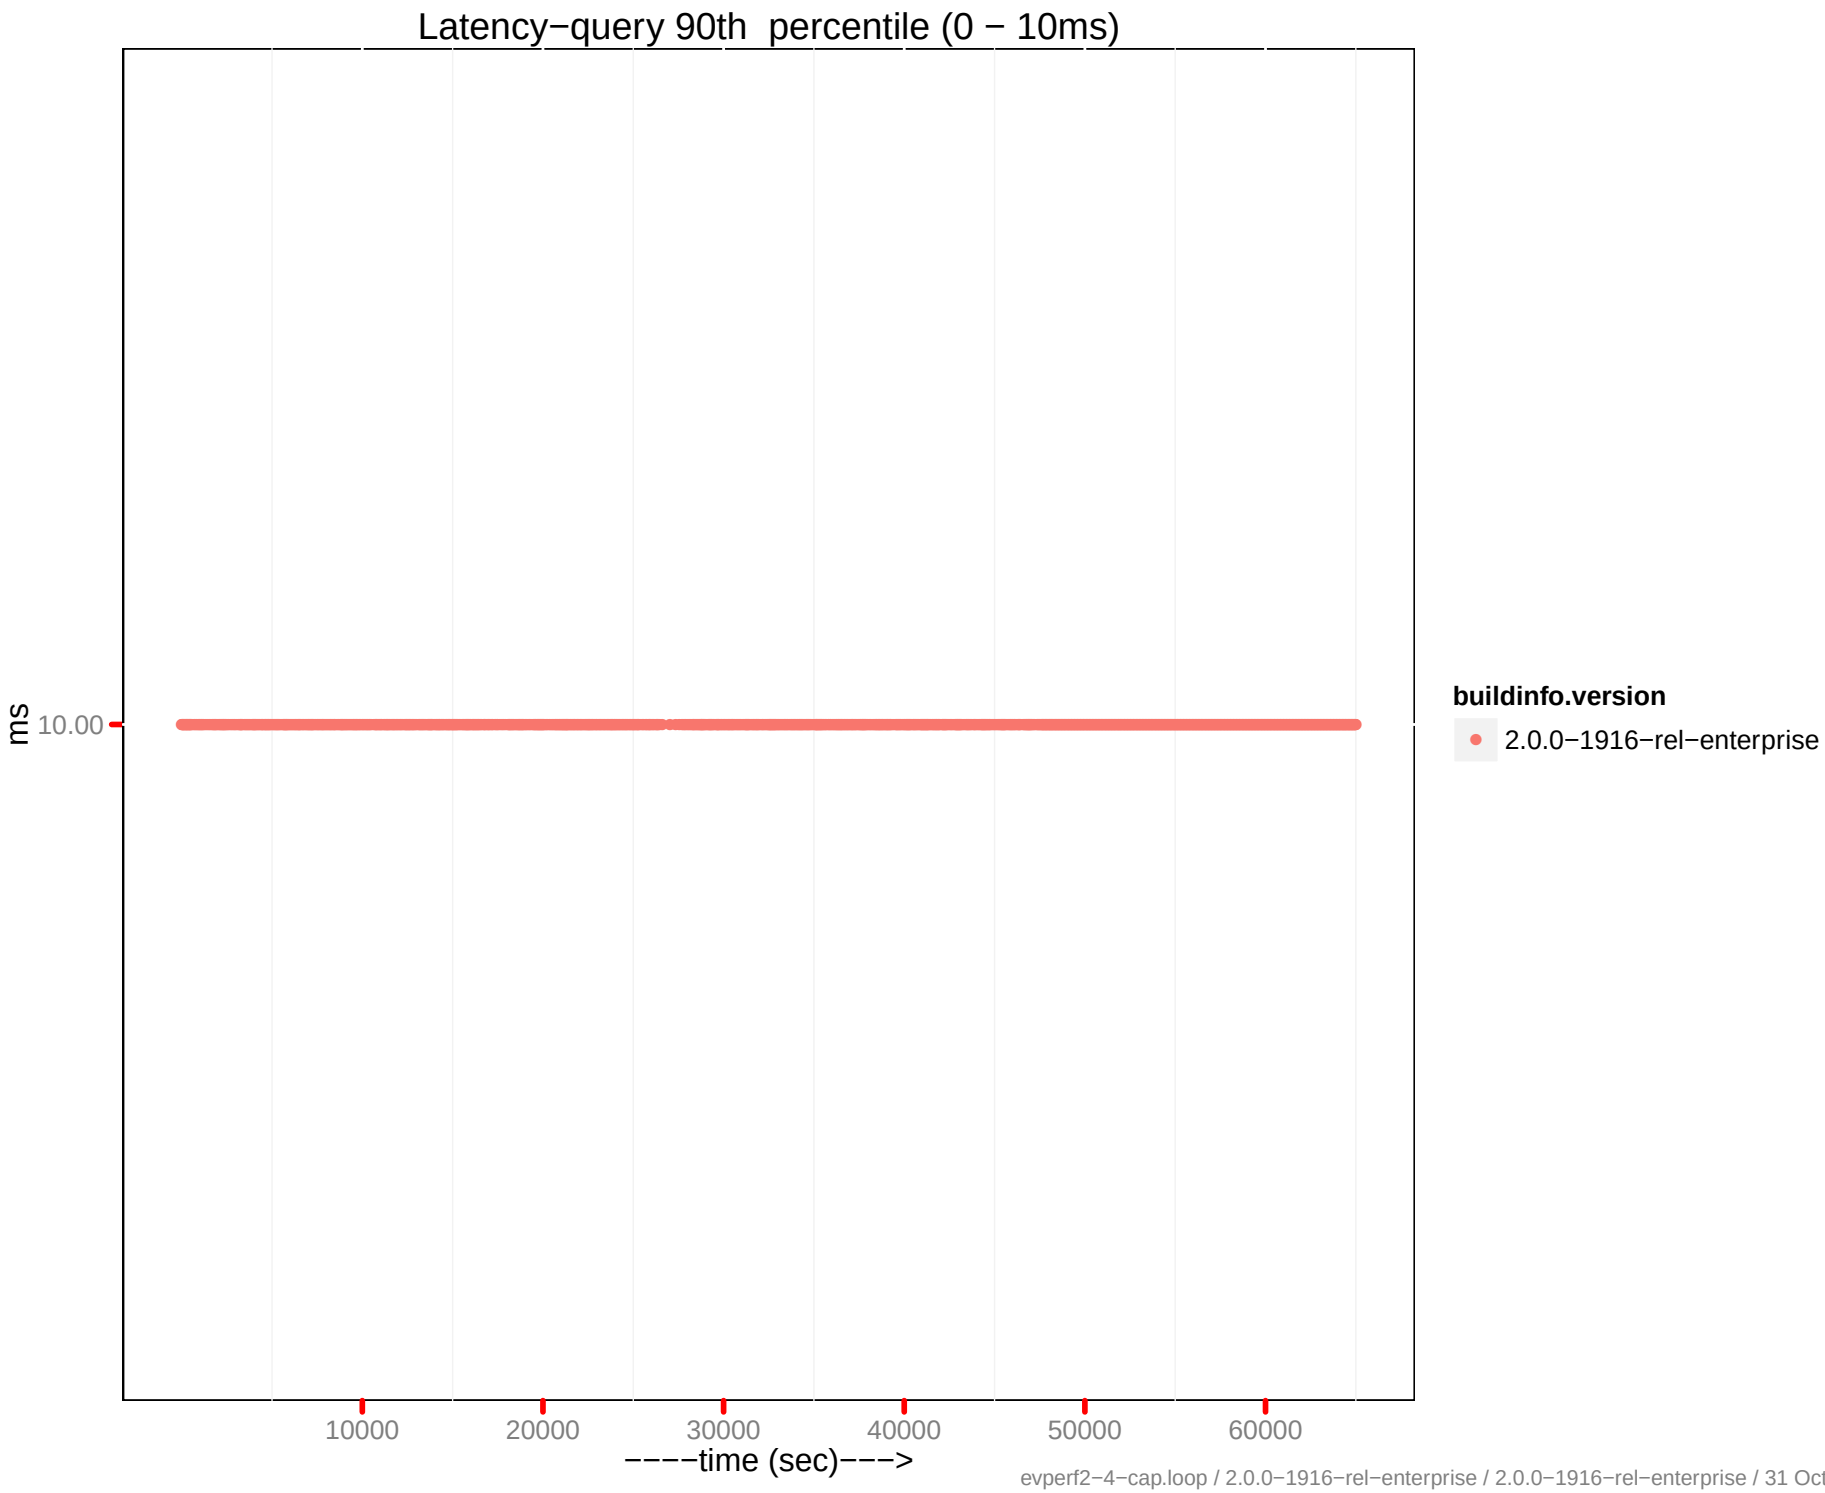
Latency−query 90th percentile (0 − 10ms)
ms
10.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

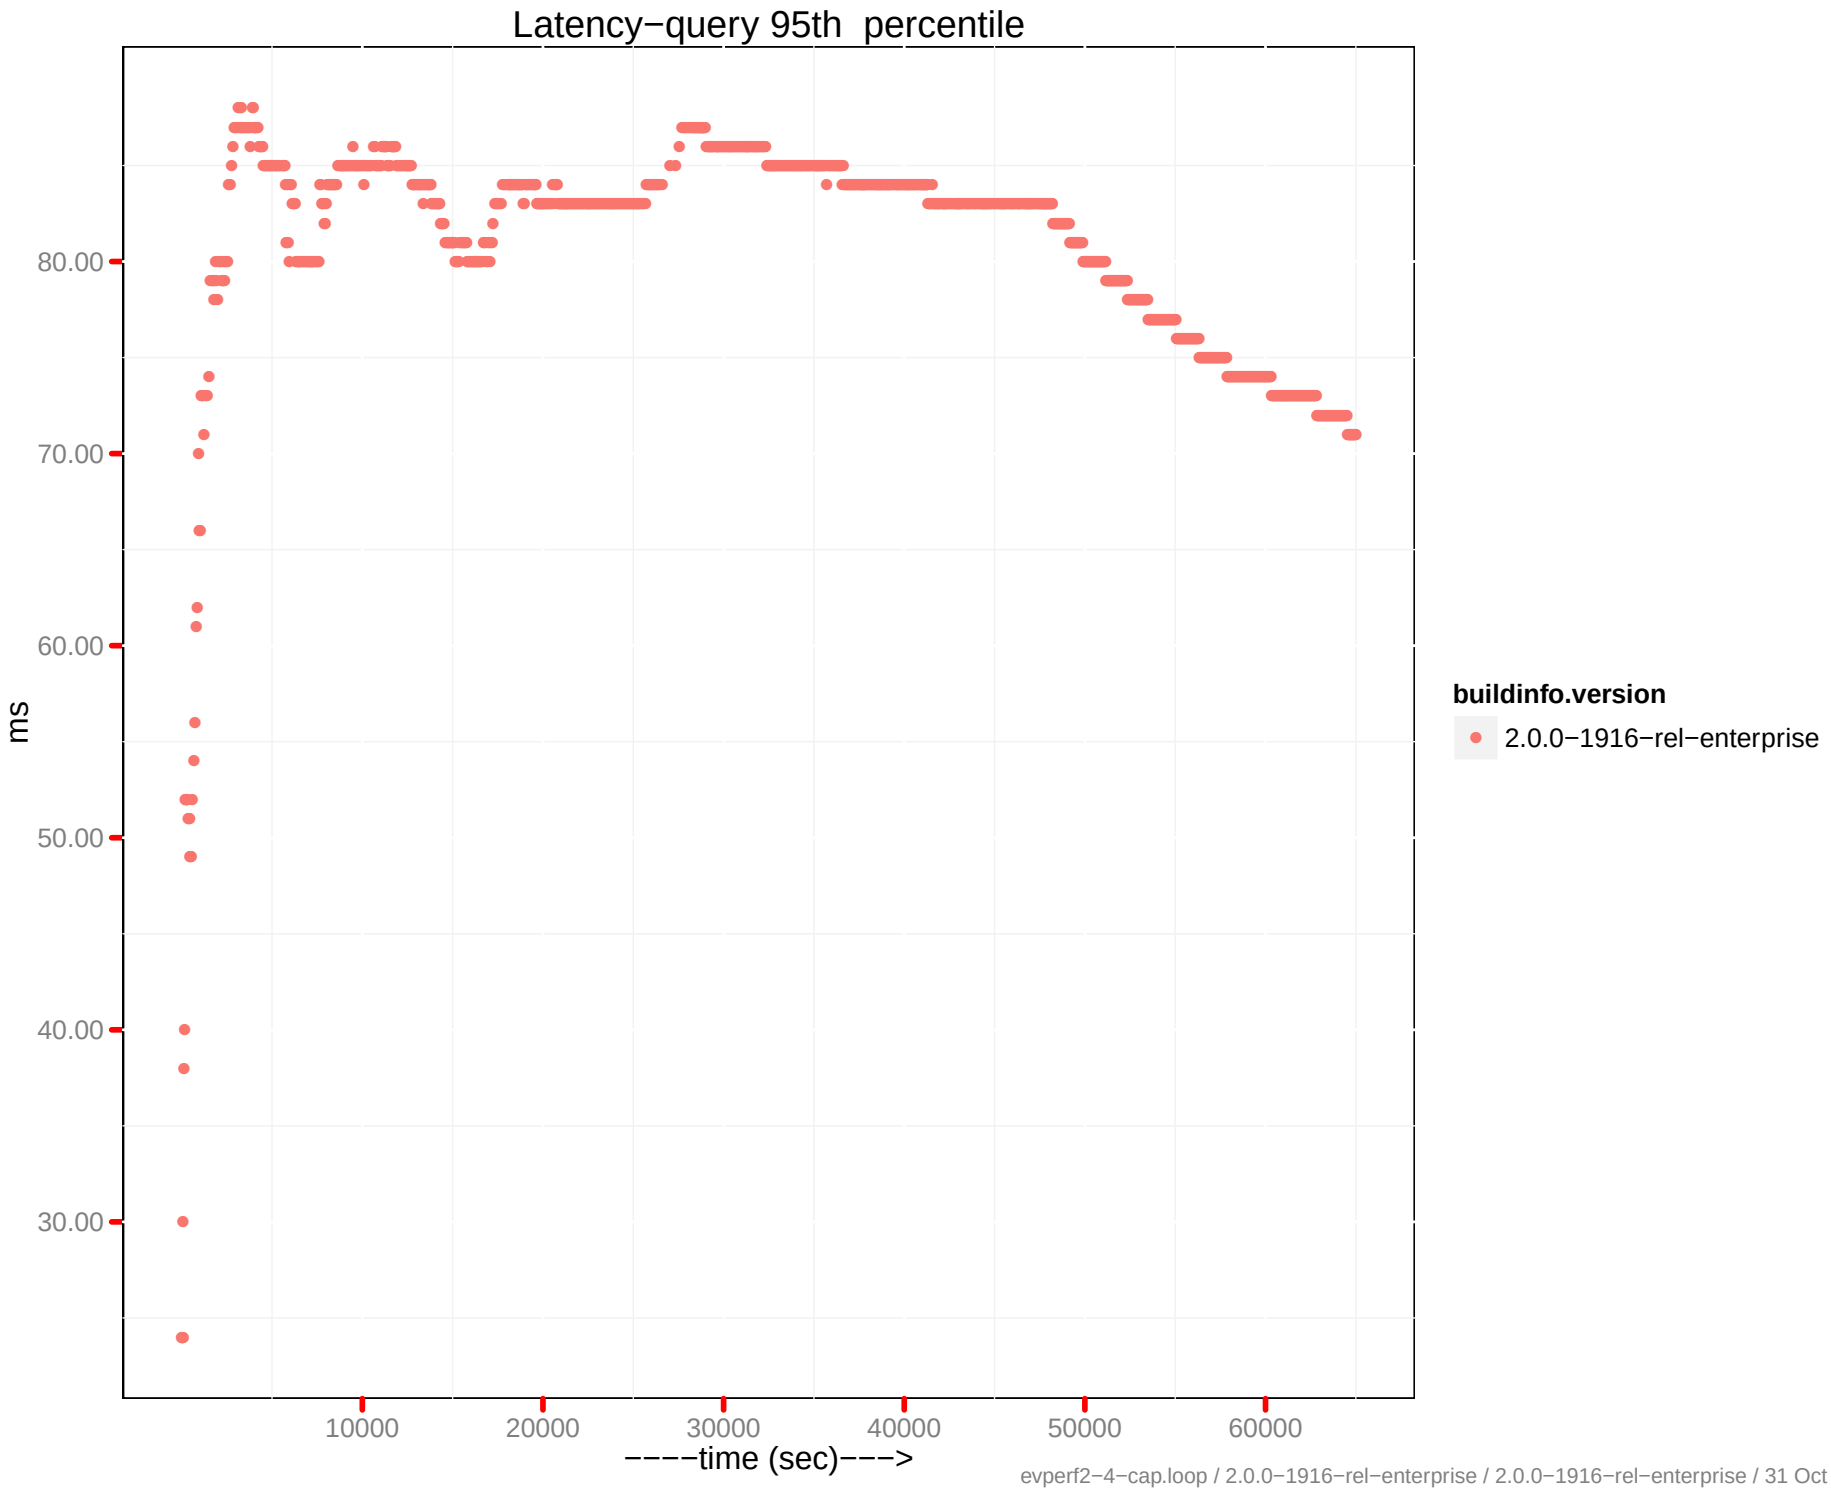
Latency−query 95th percentile
80.00
70.00
60.00
50.00
40.00
30.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

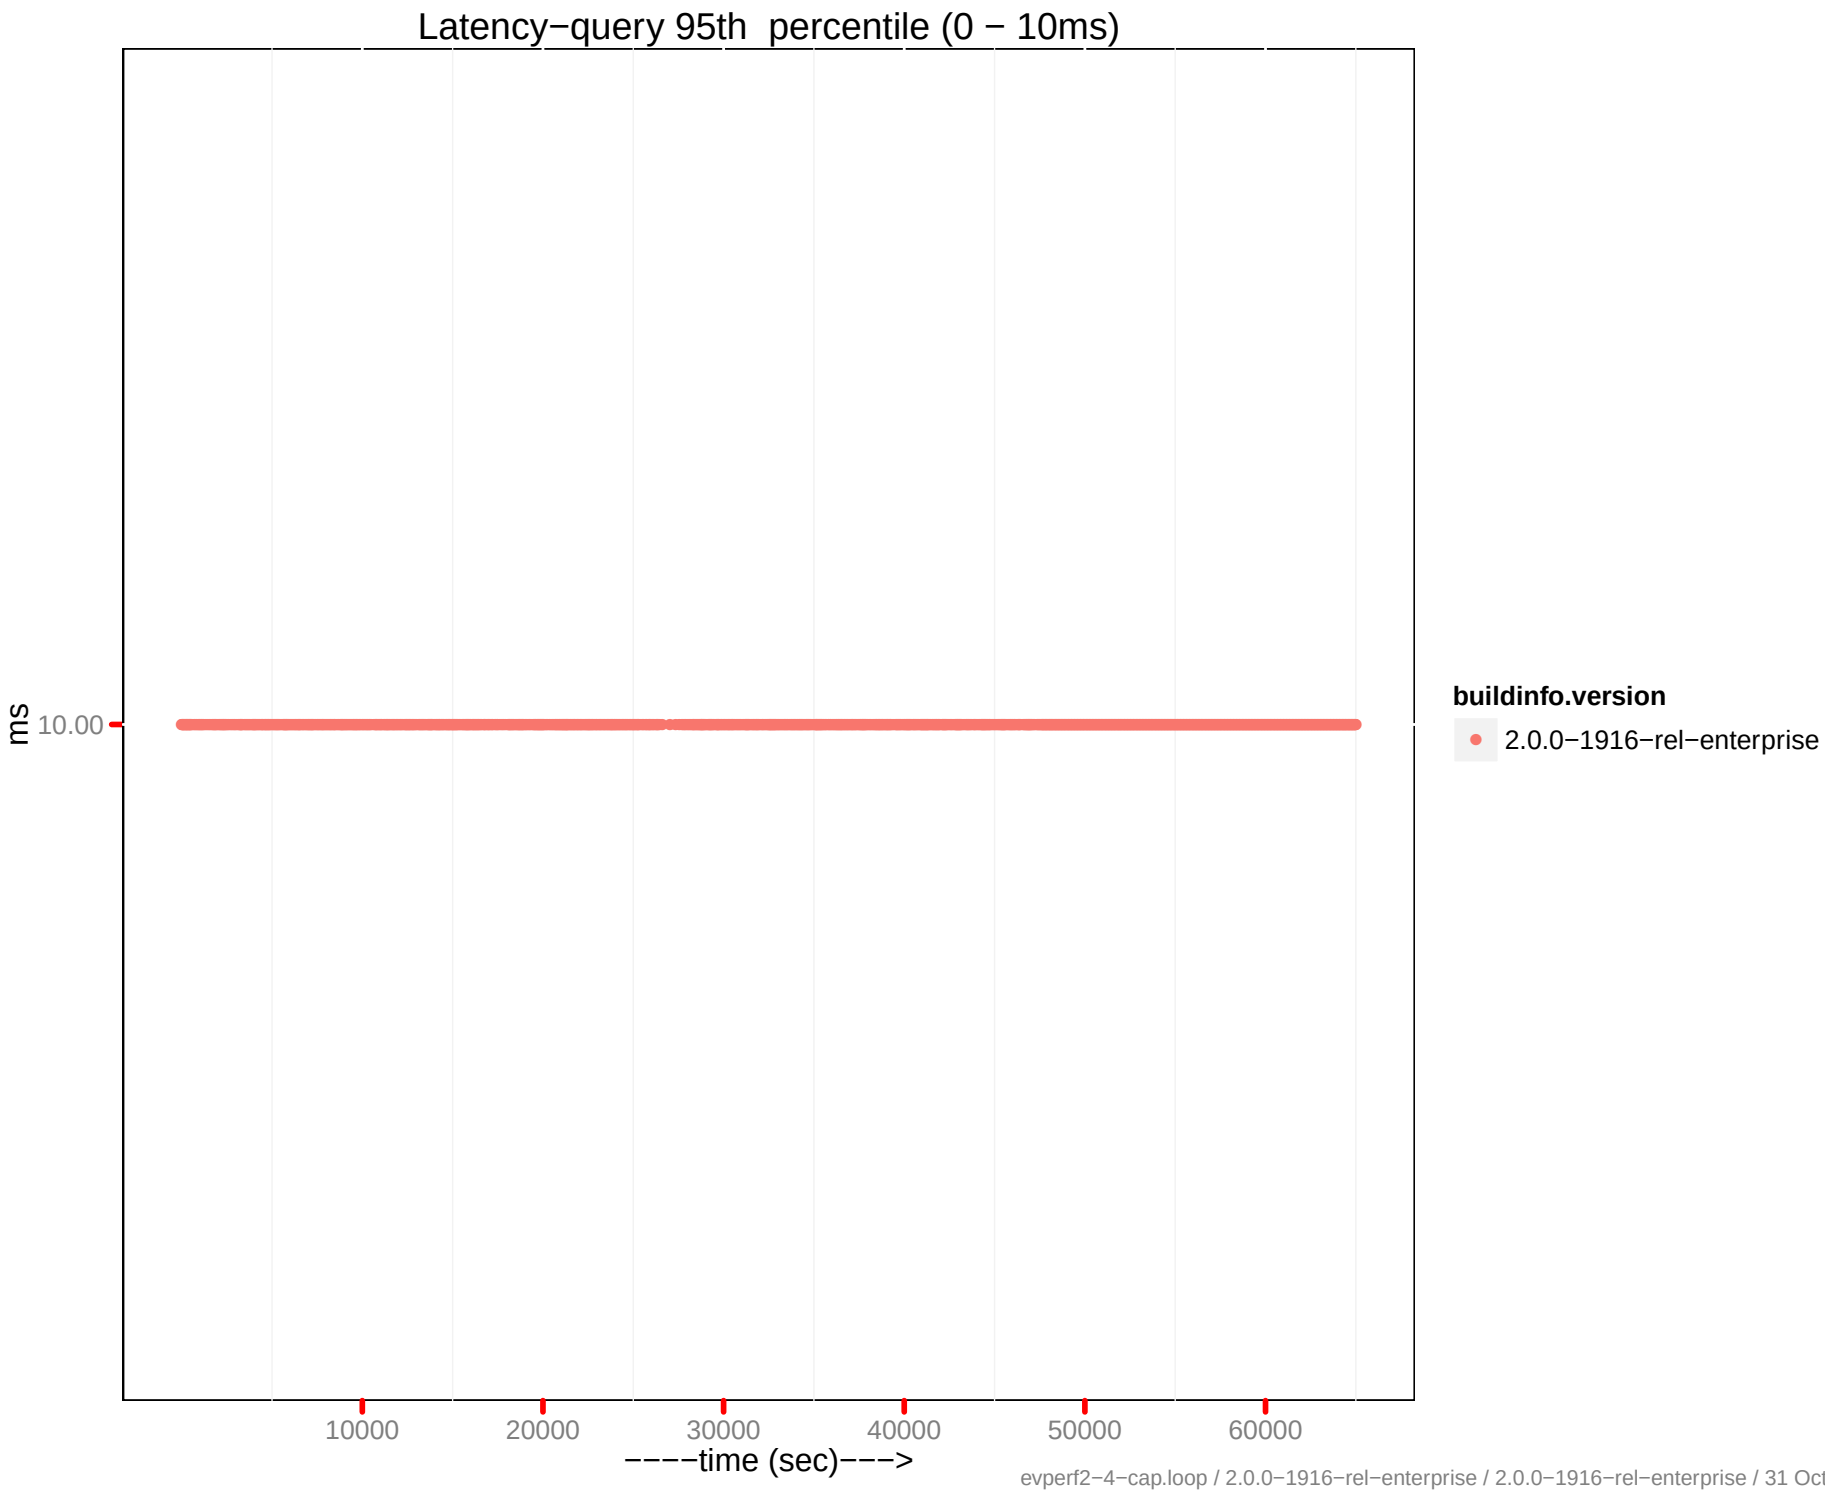
Latency−query 95th percentile (0 − 10ms)
ms
10.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

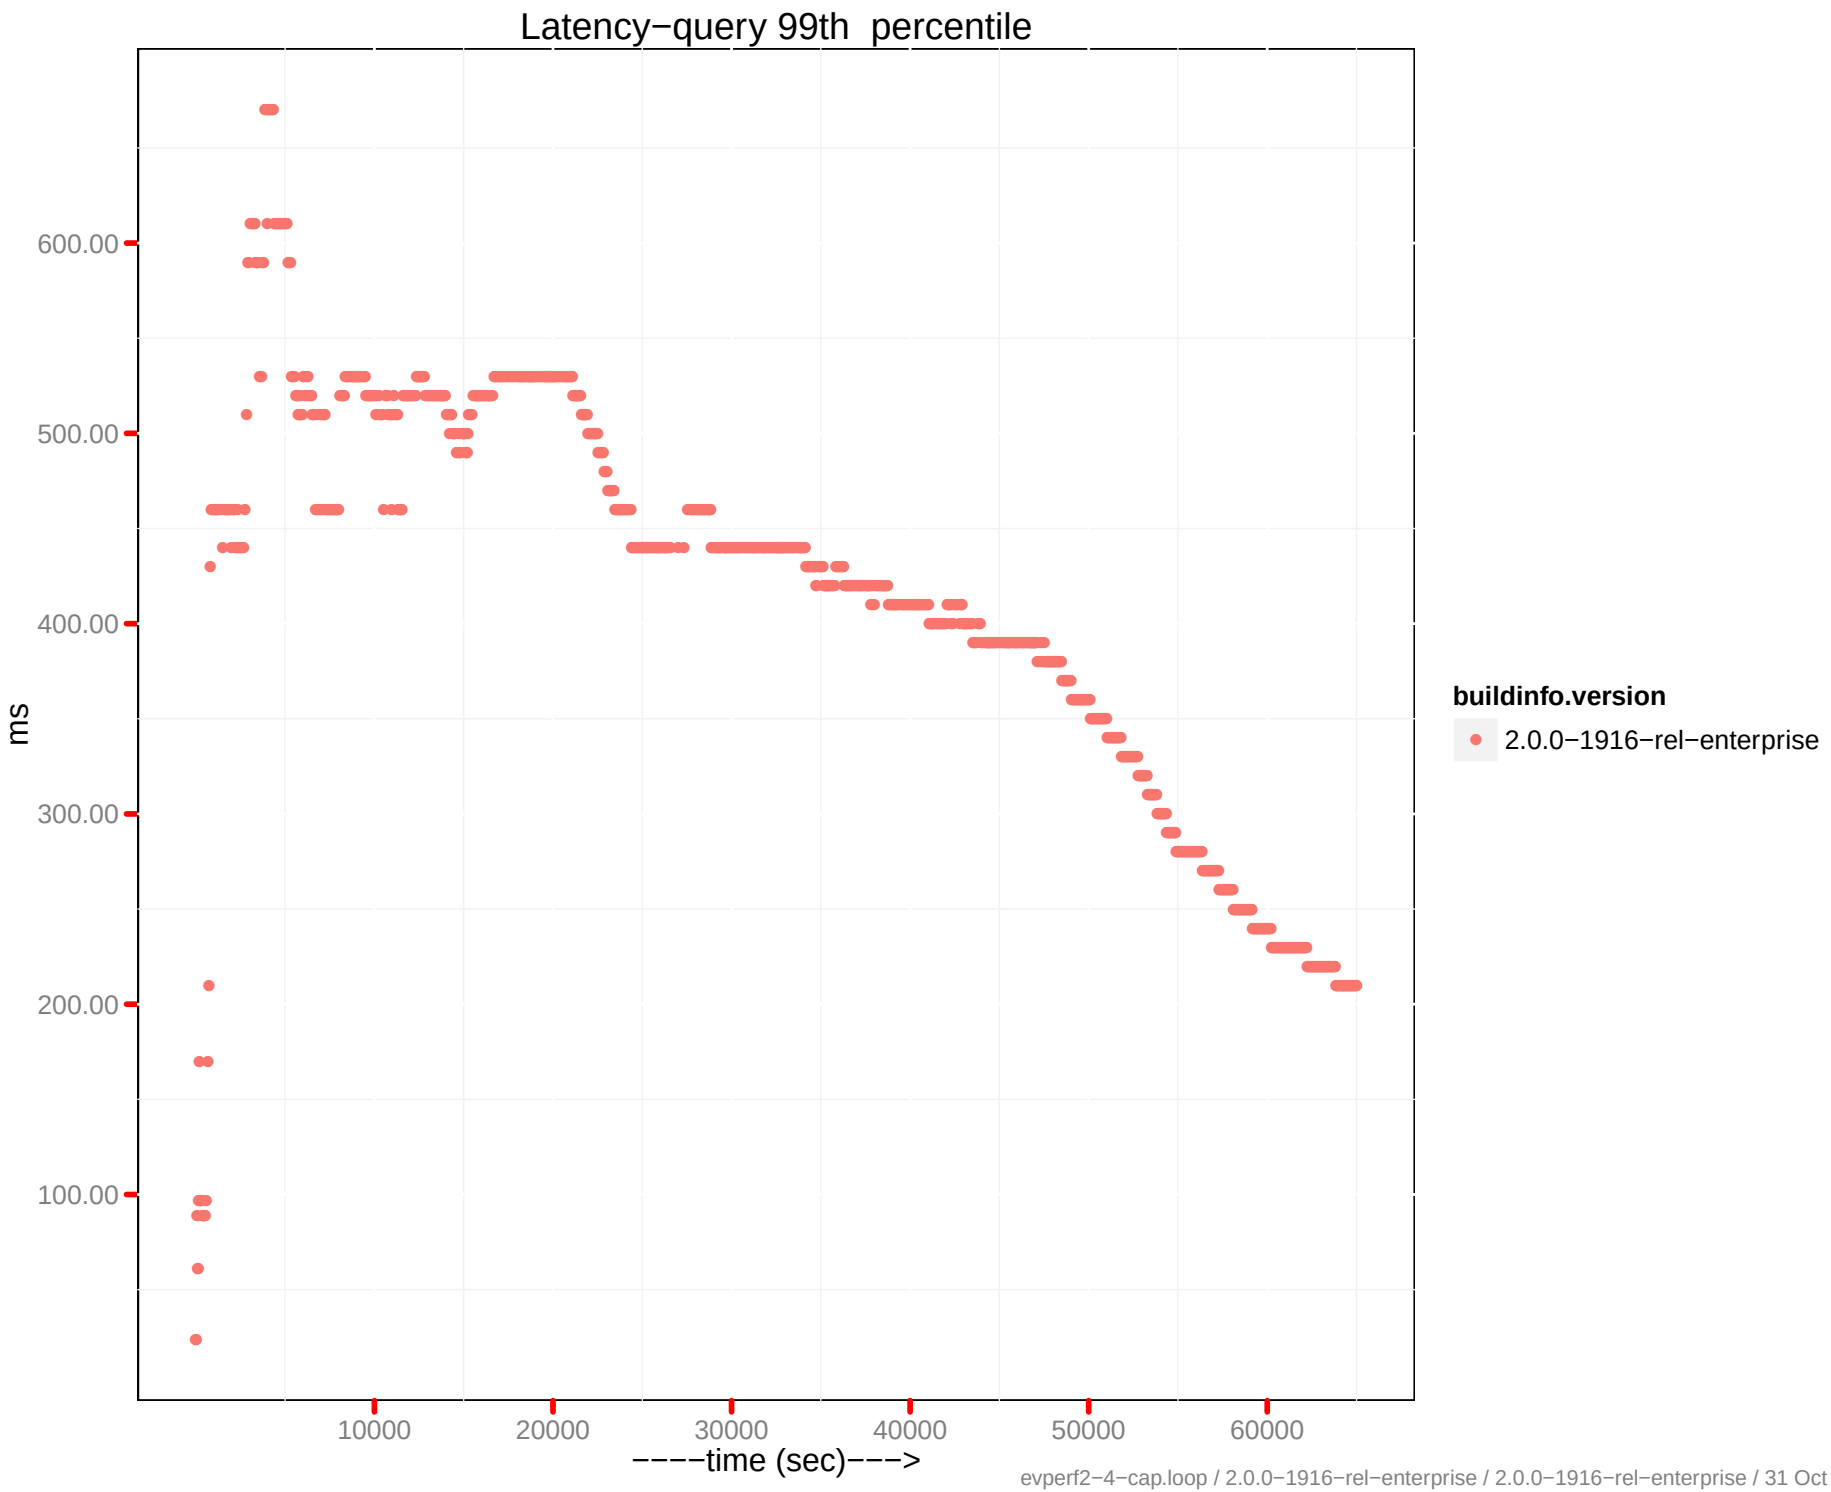
Latency−query 99th percentile
600.00
500.00
400.00
300.00
200.00
100.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

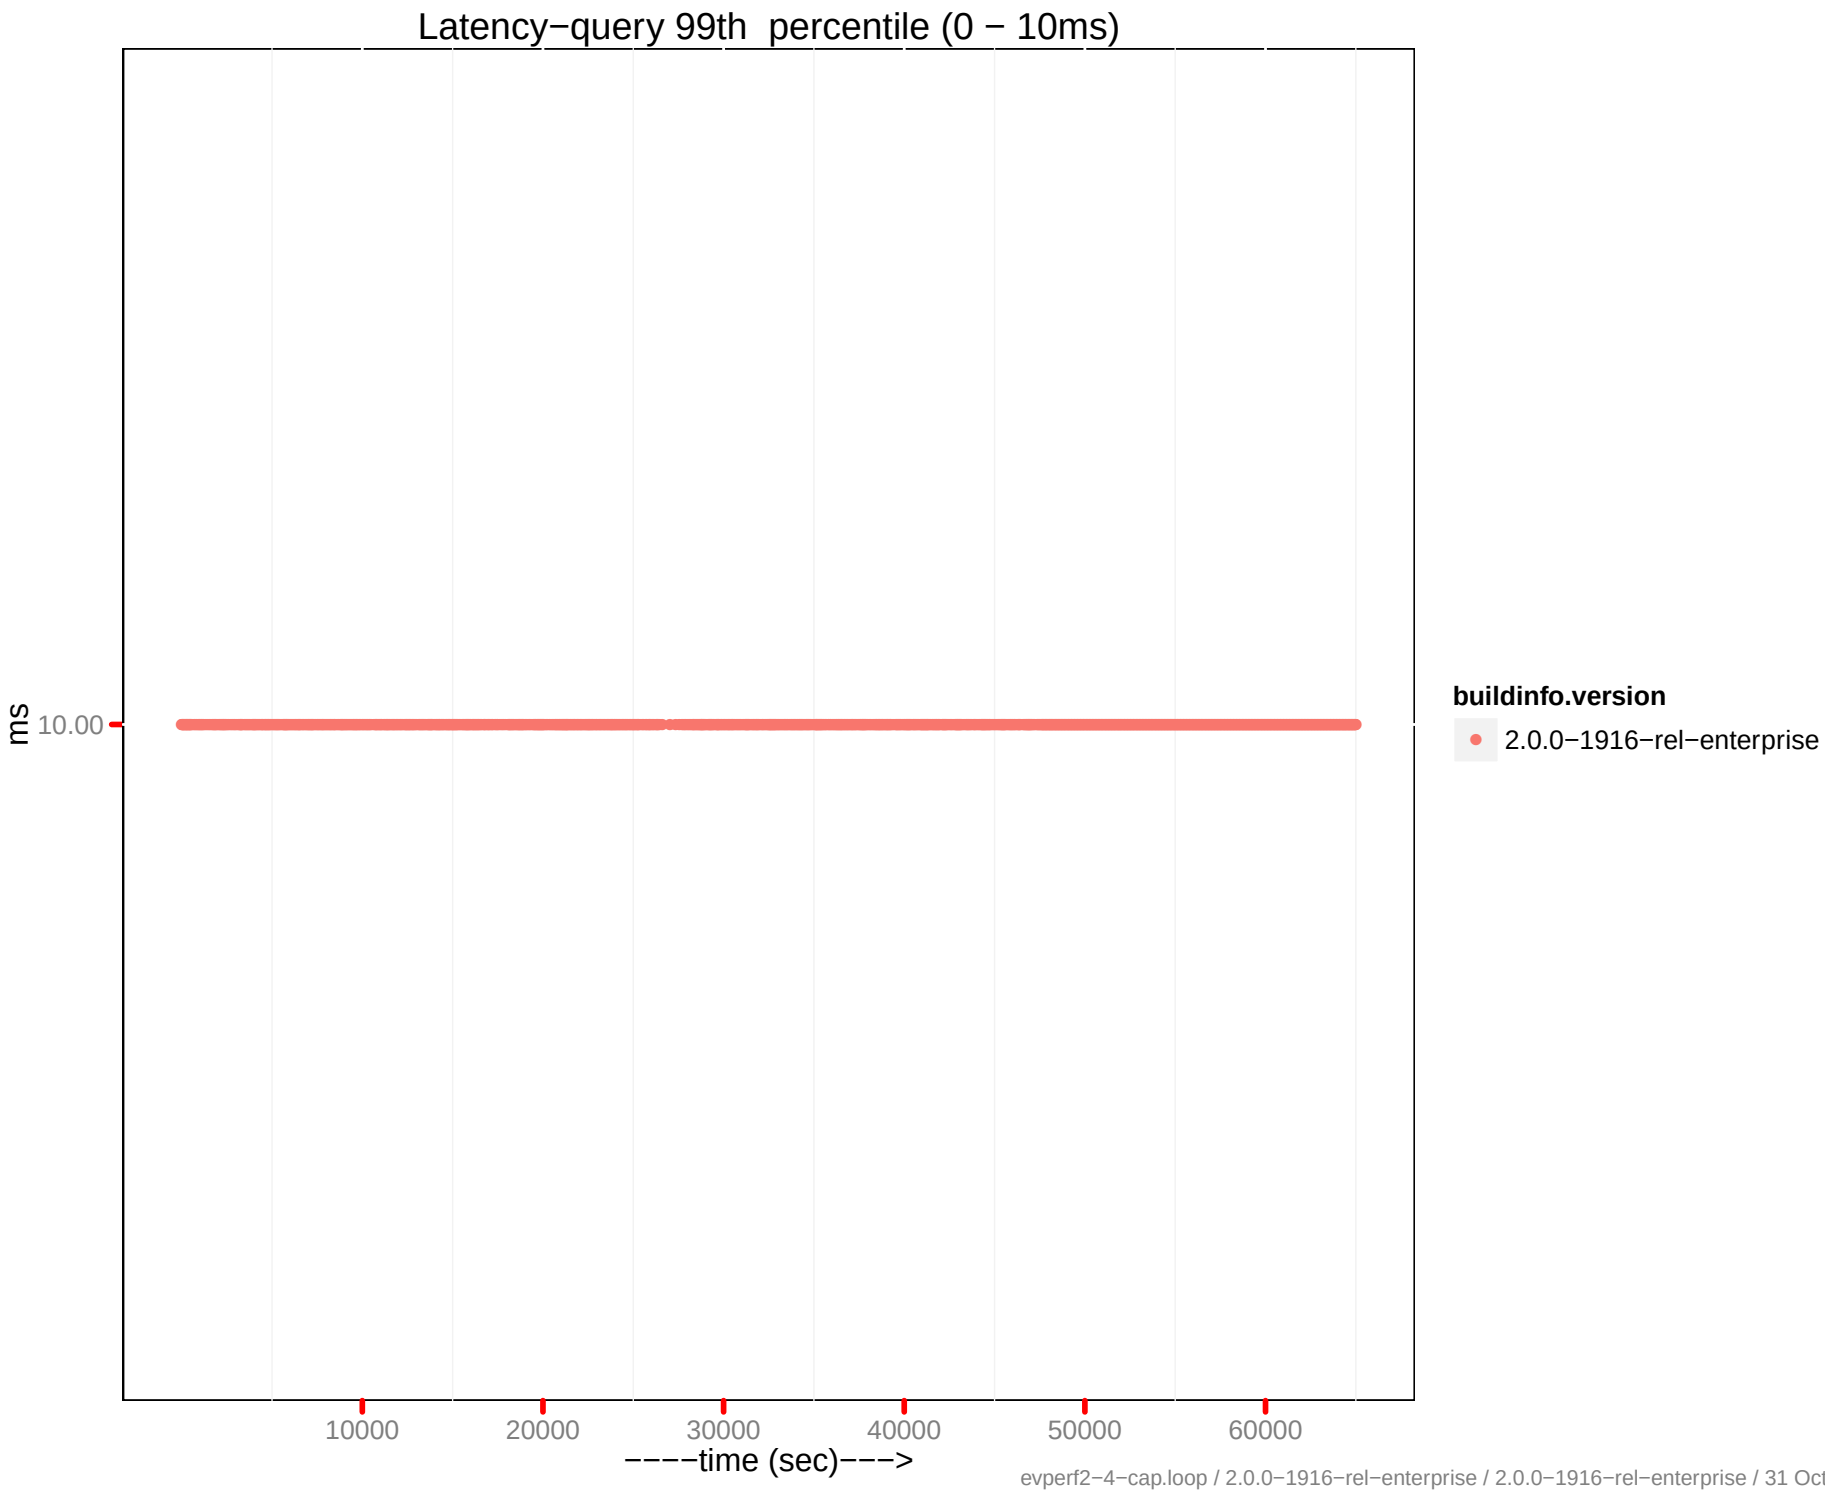
Latency−query 99th percentile (0 − 10ms)
ms
10.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

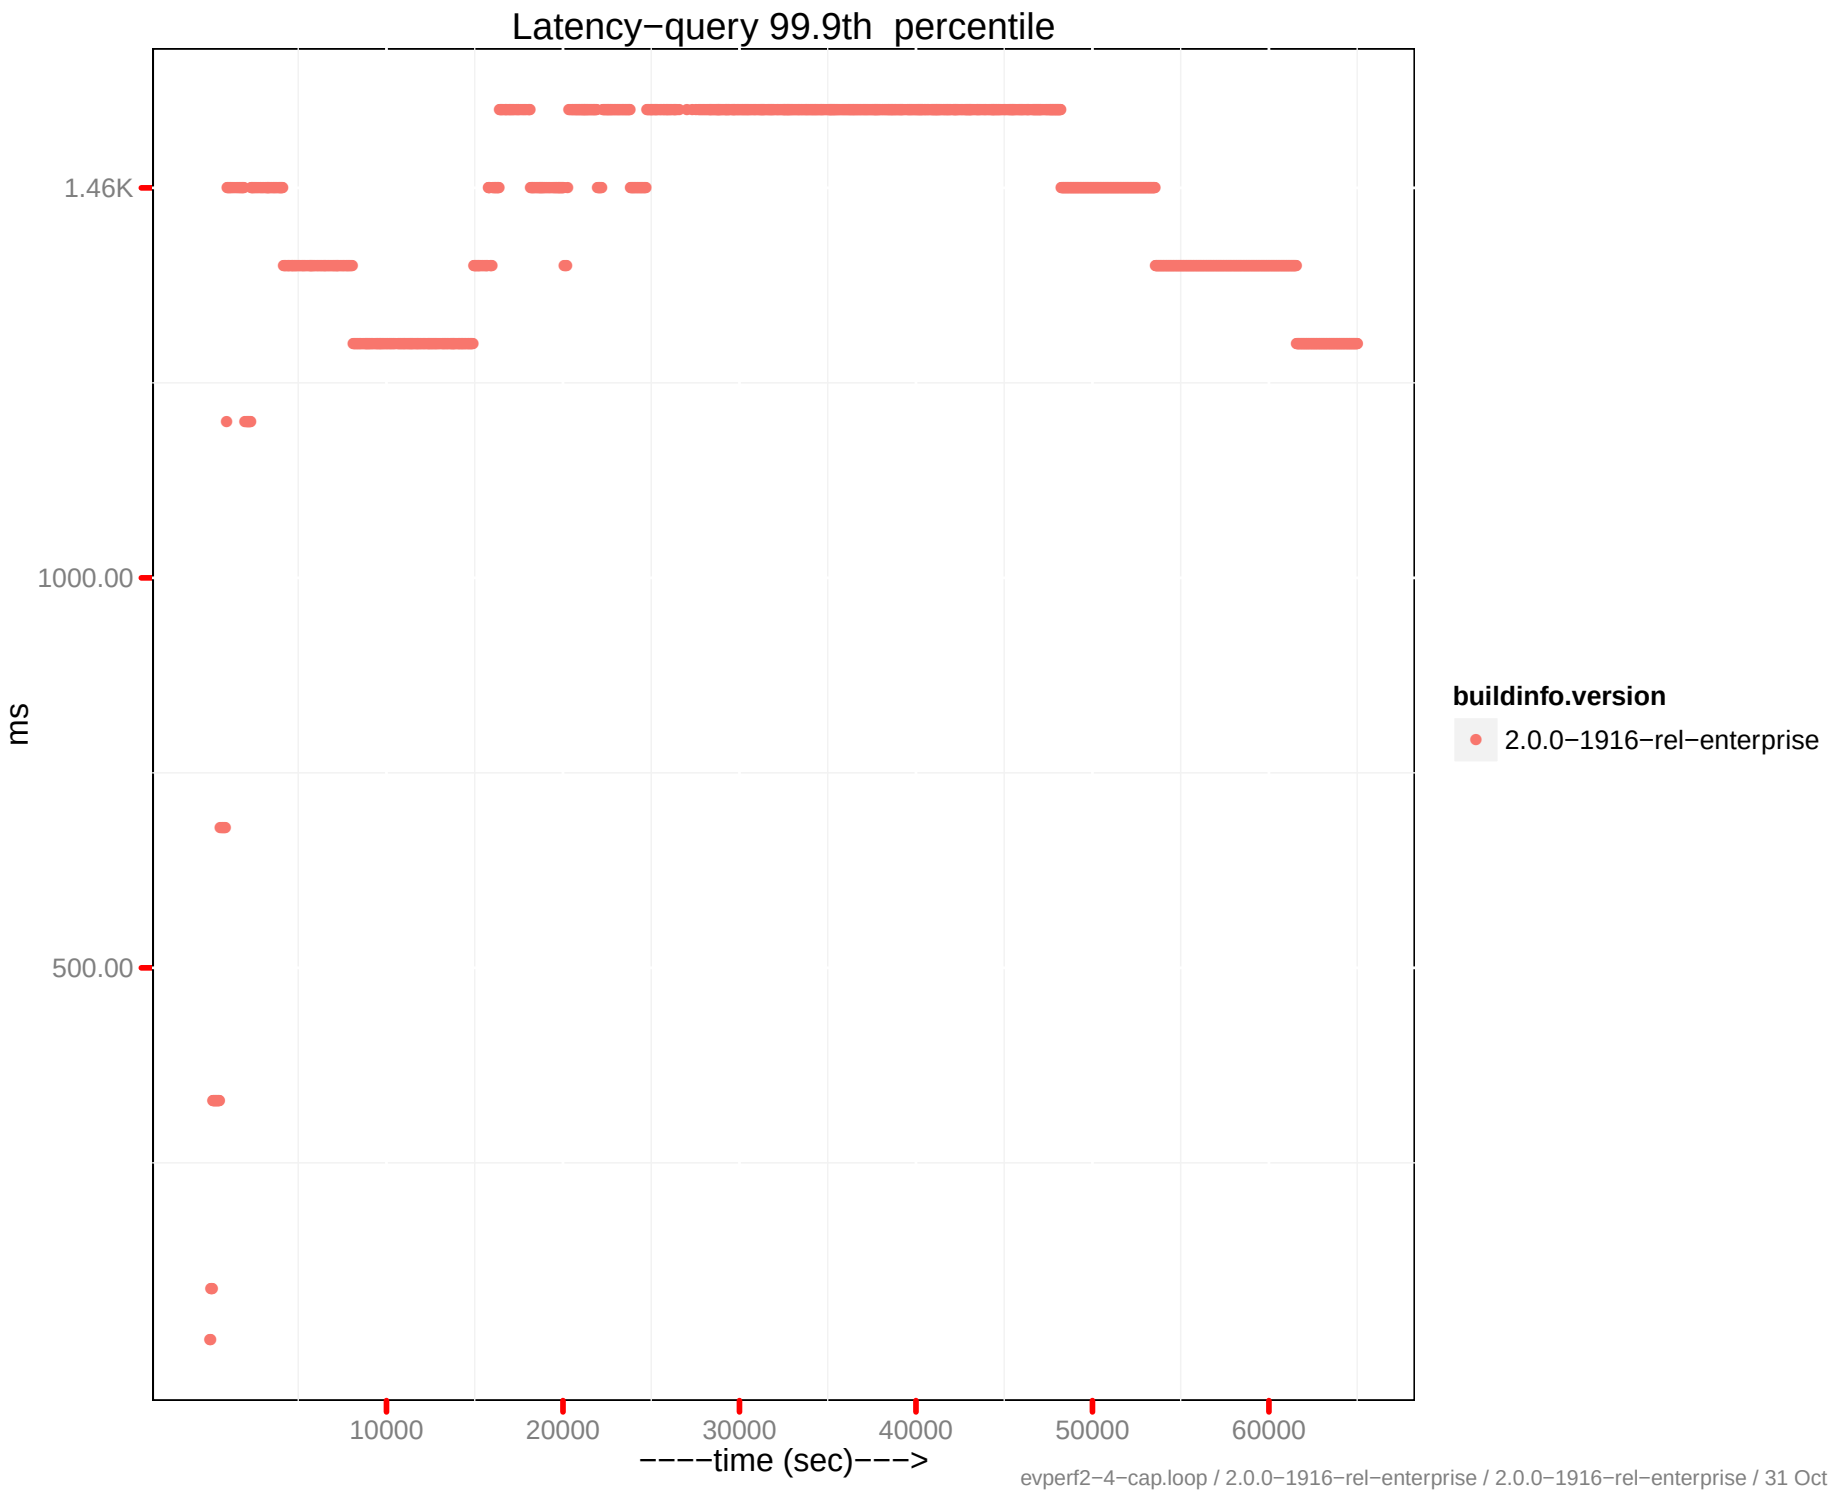
Latency−query 99.9th percentile
1.46K
1000.00
500.00
ms
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

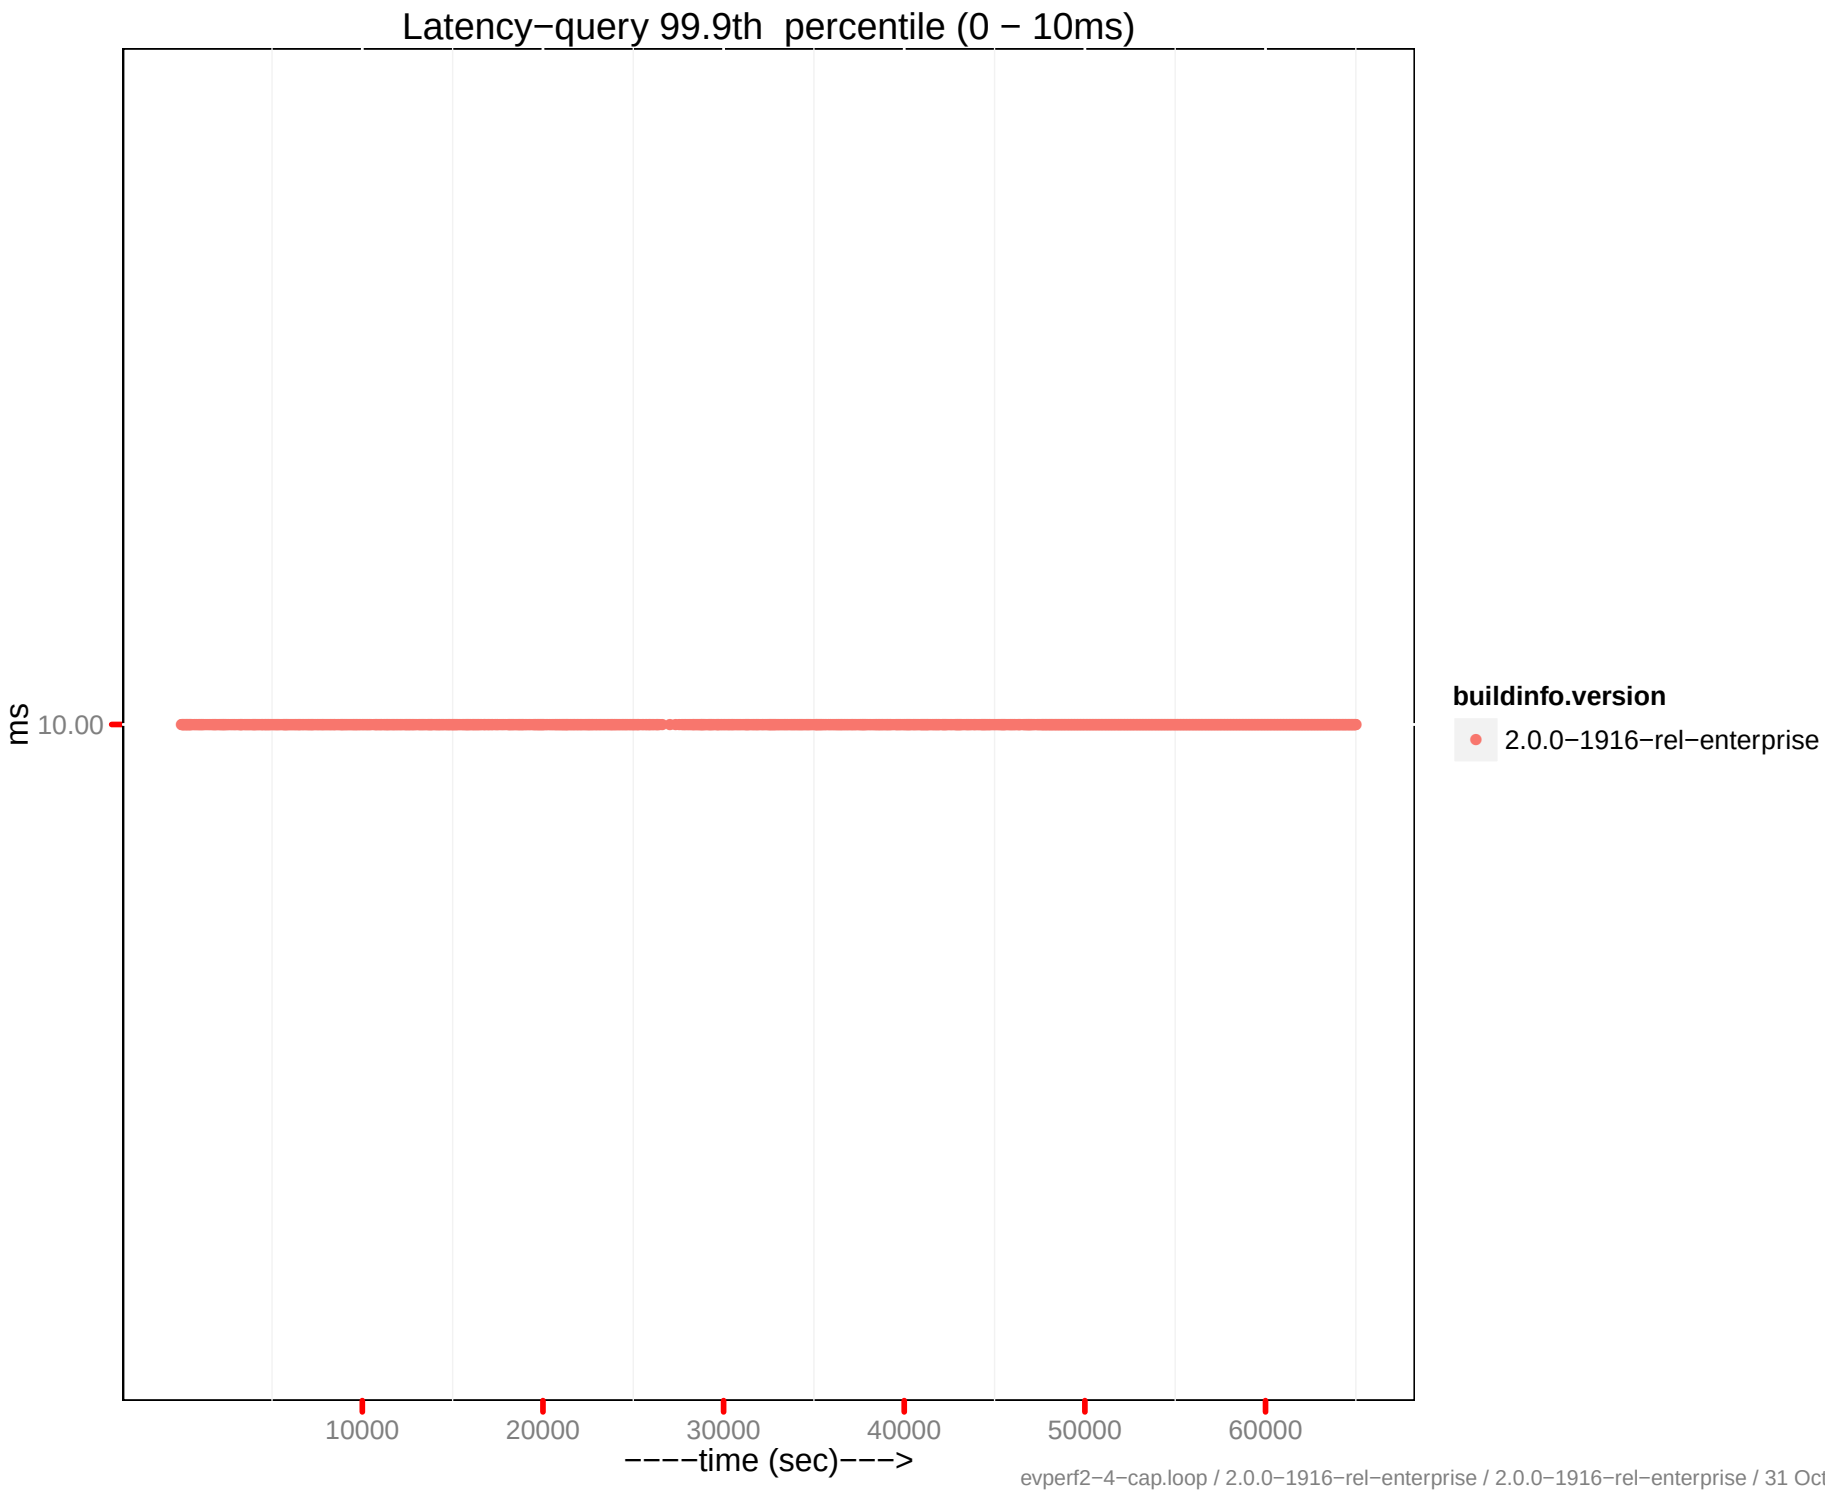
Latency−query 99.9th percentile (0 − 10ms)
ms
10.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

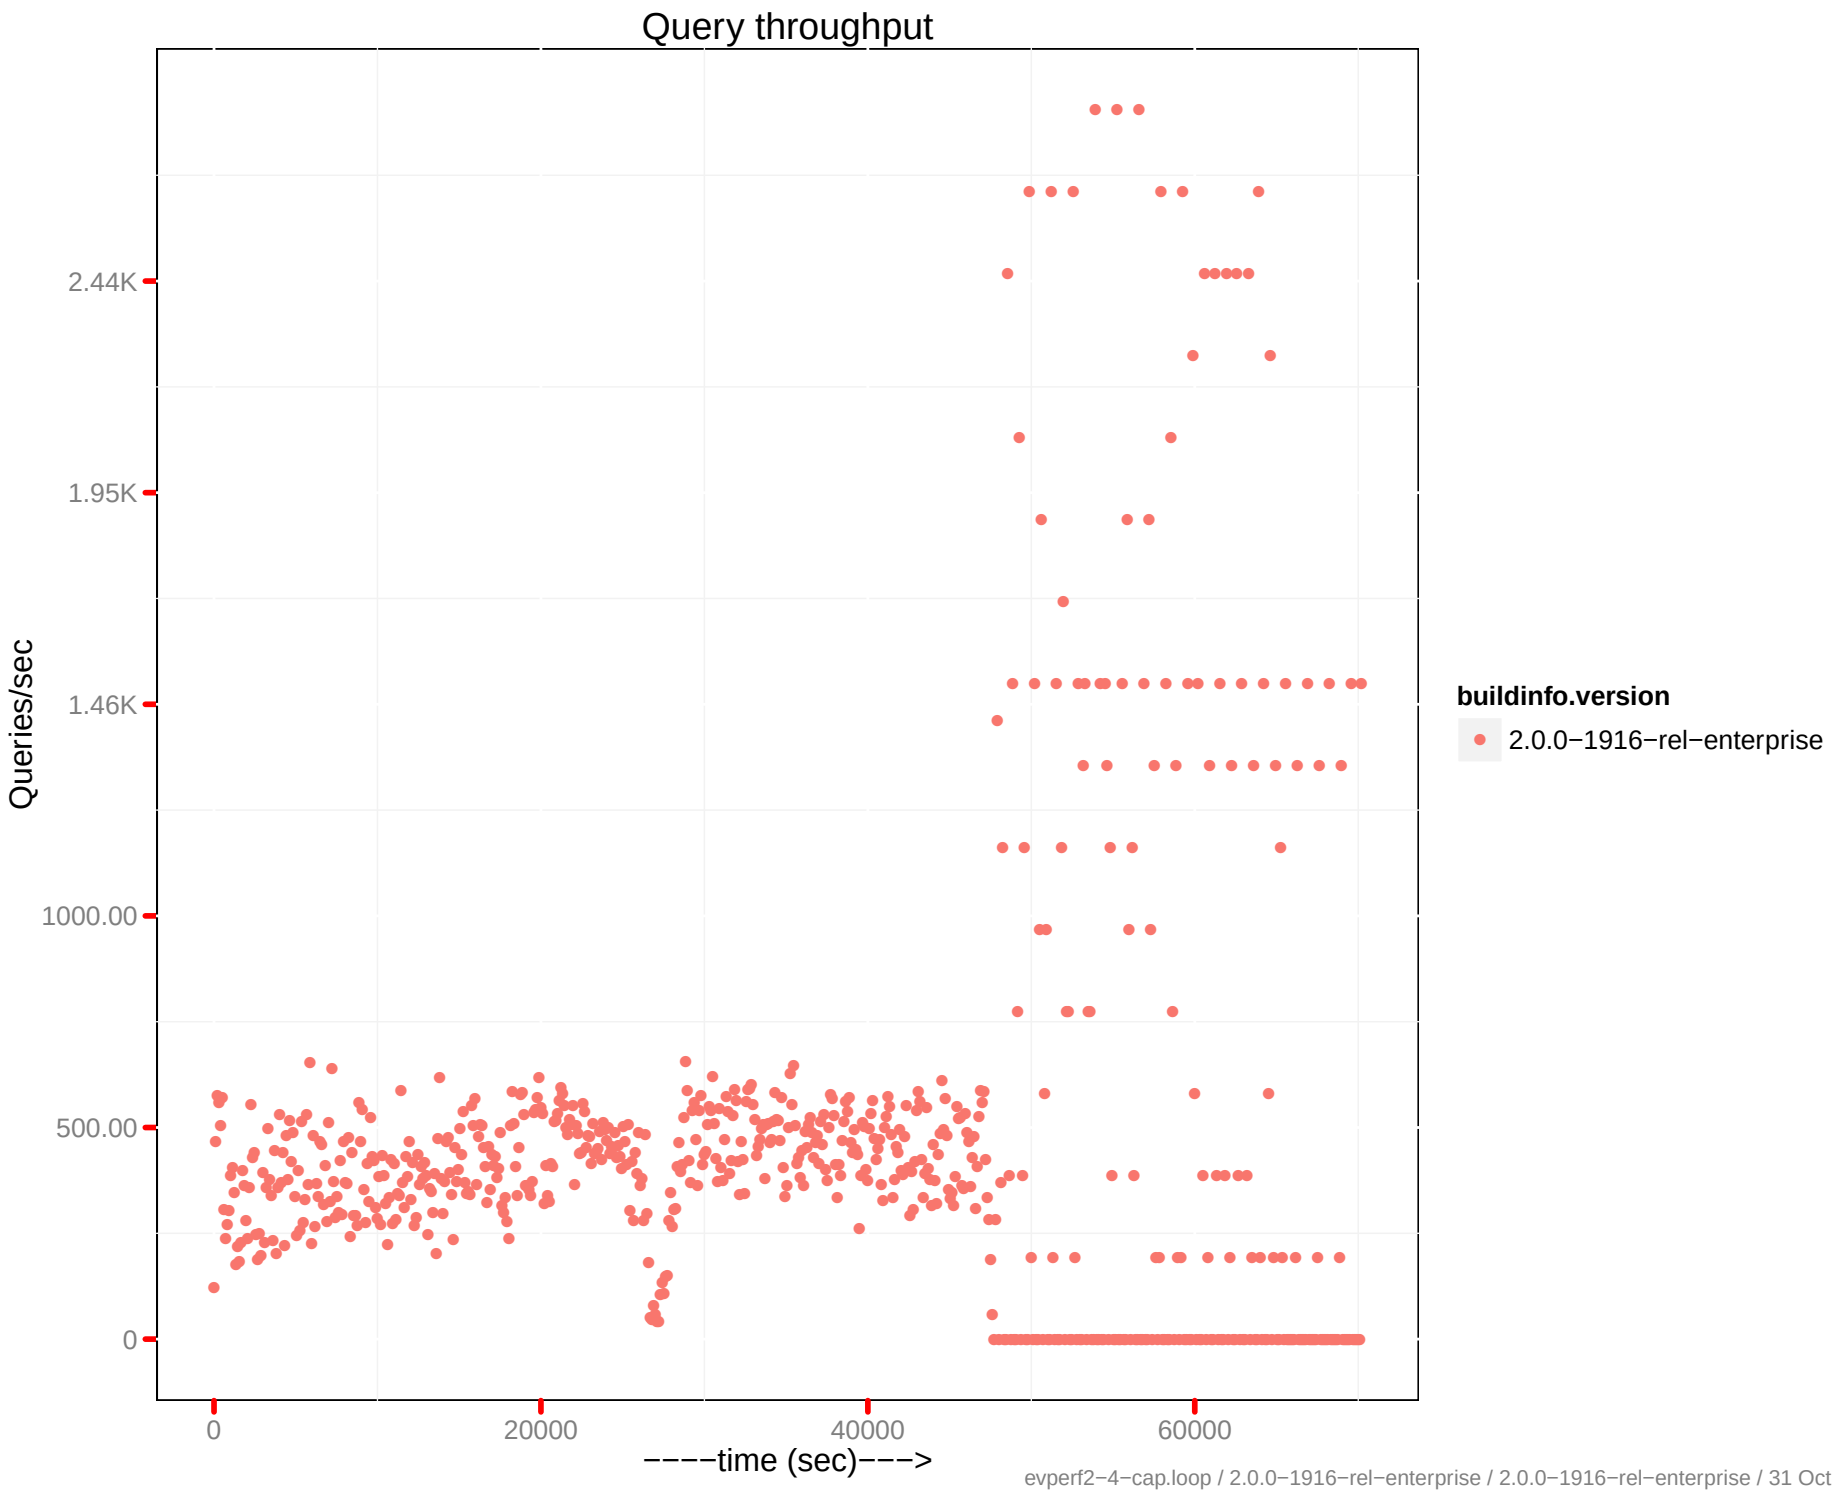
Query throughput
Queries/sec
2.44K
1.95K
1.46K
1000.00
500.00
0
0
20000
40000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

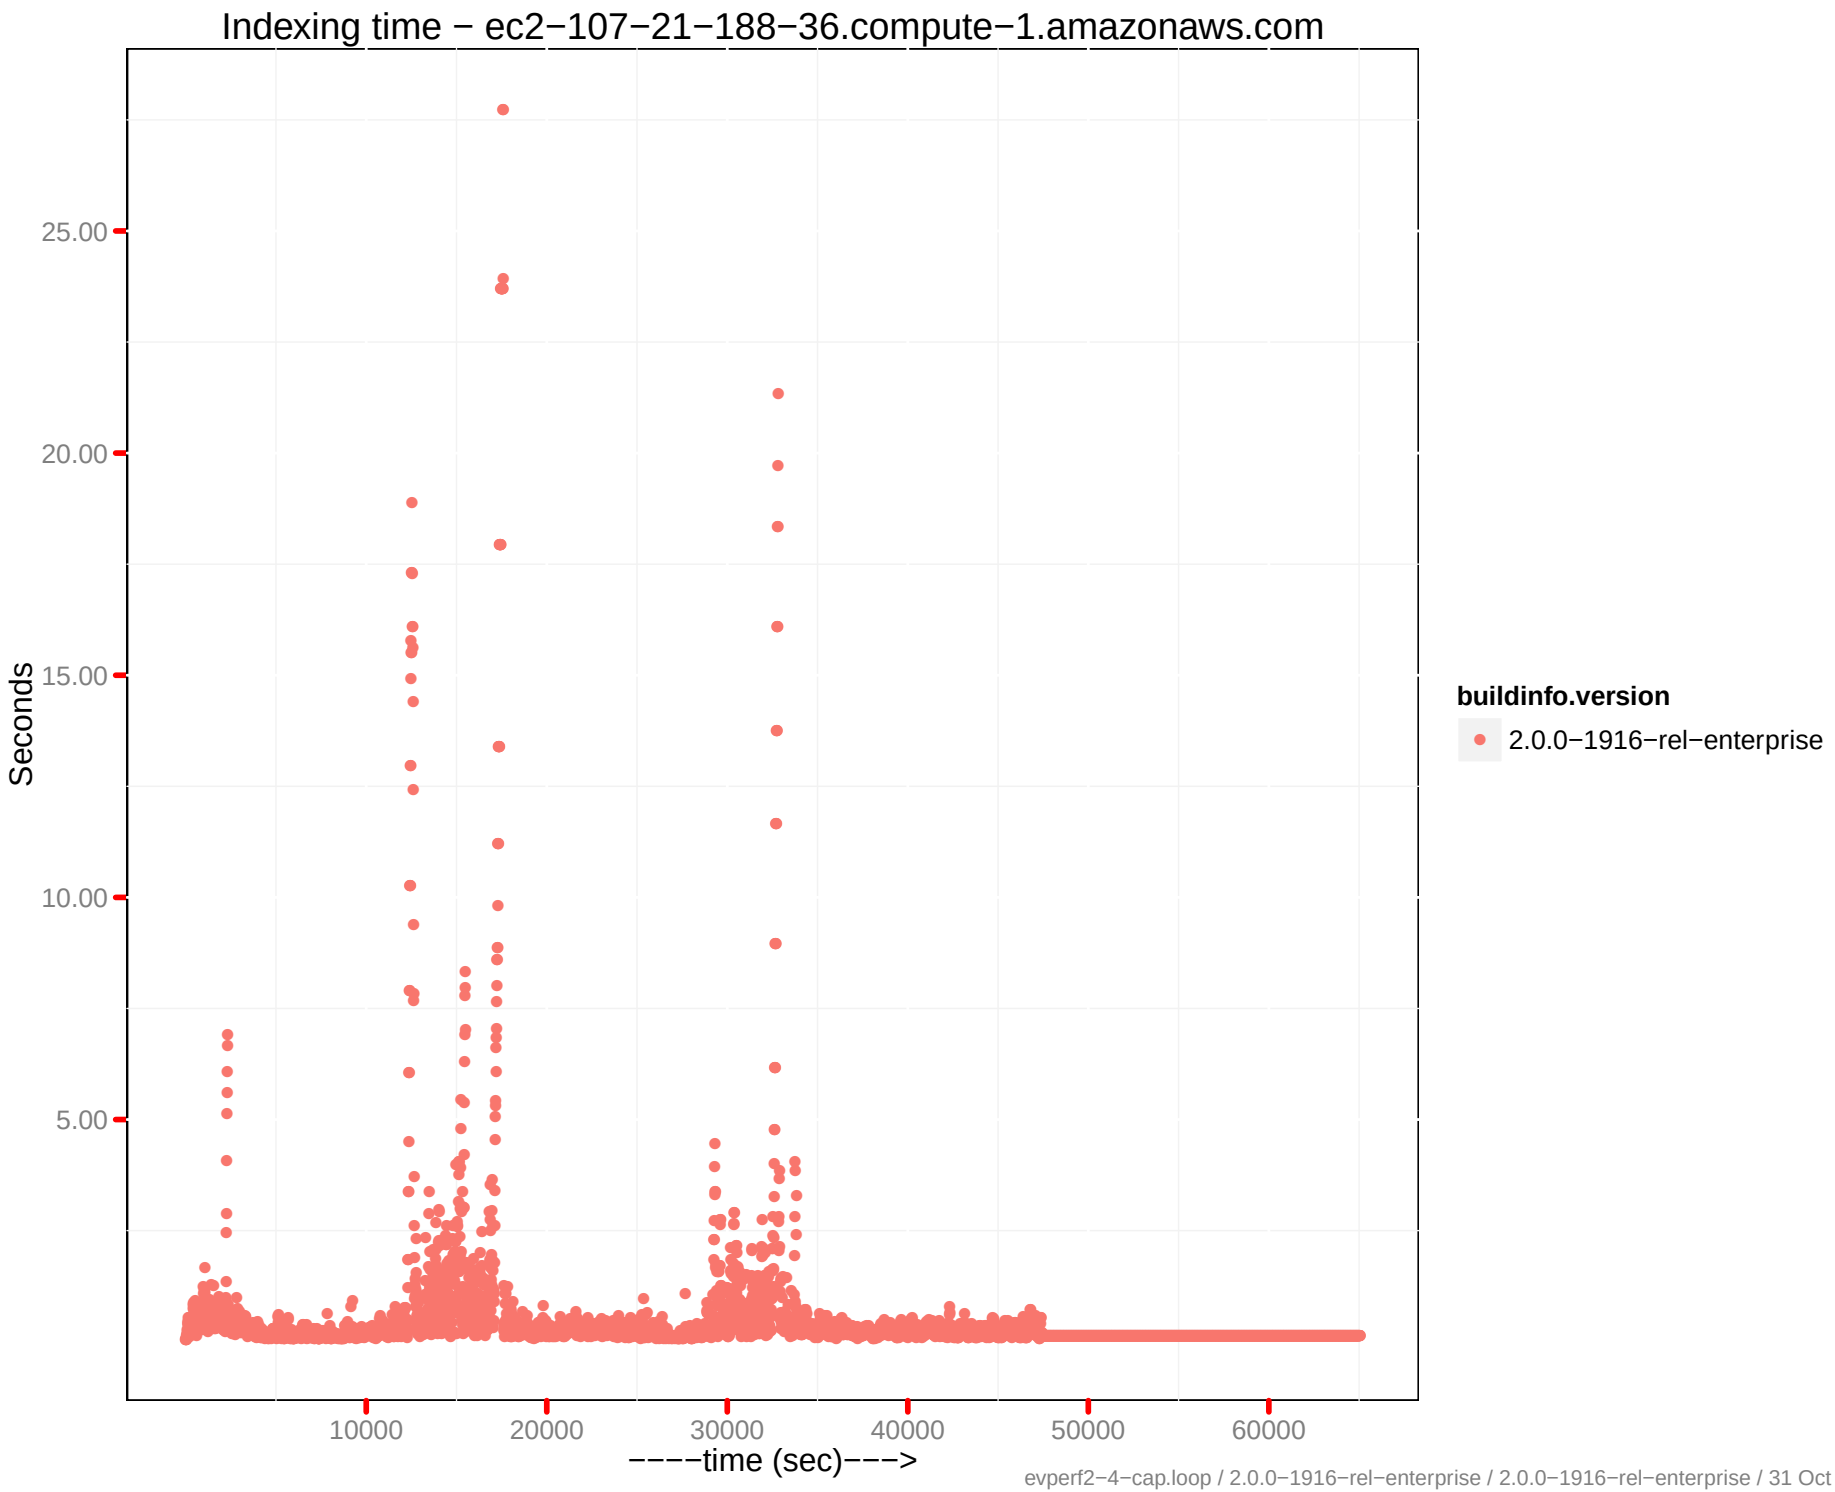
Indexing time − ec2−107−21−188−36.compute−1.amazonaws.com
Seconds
25.00
20.00
15.00
10.00
5.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

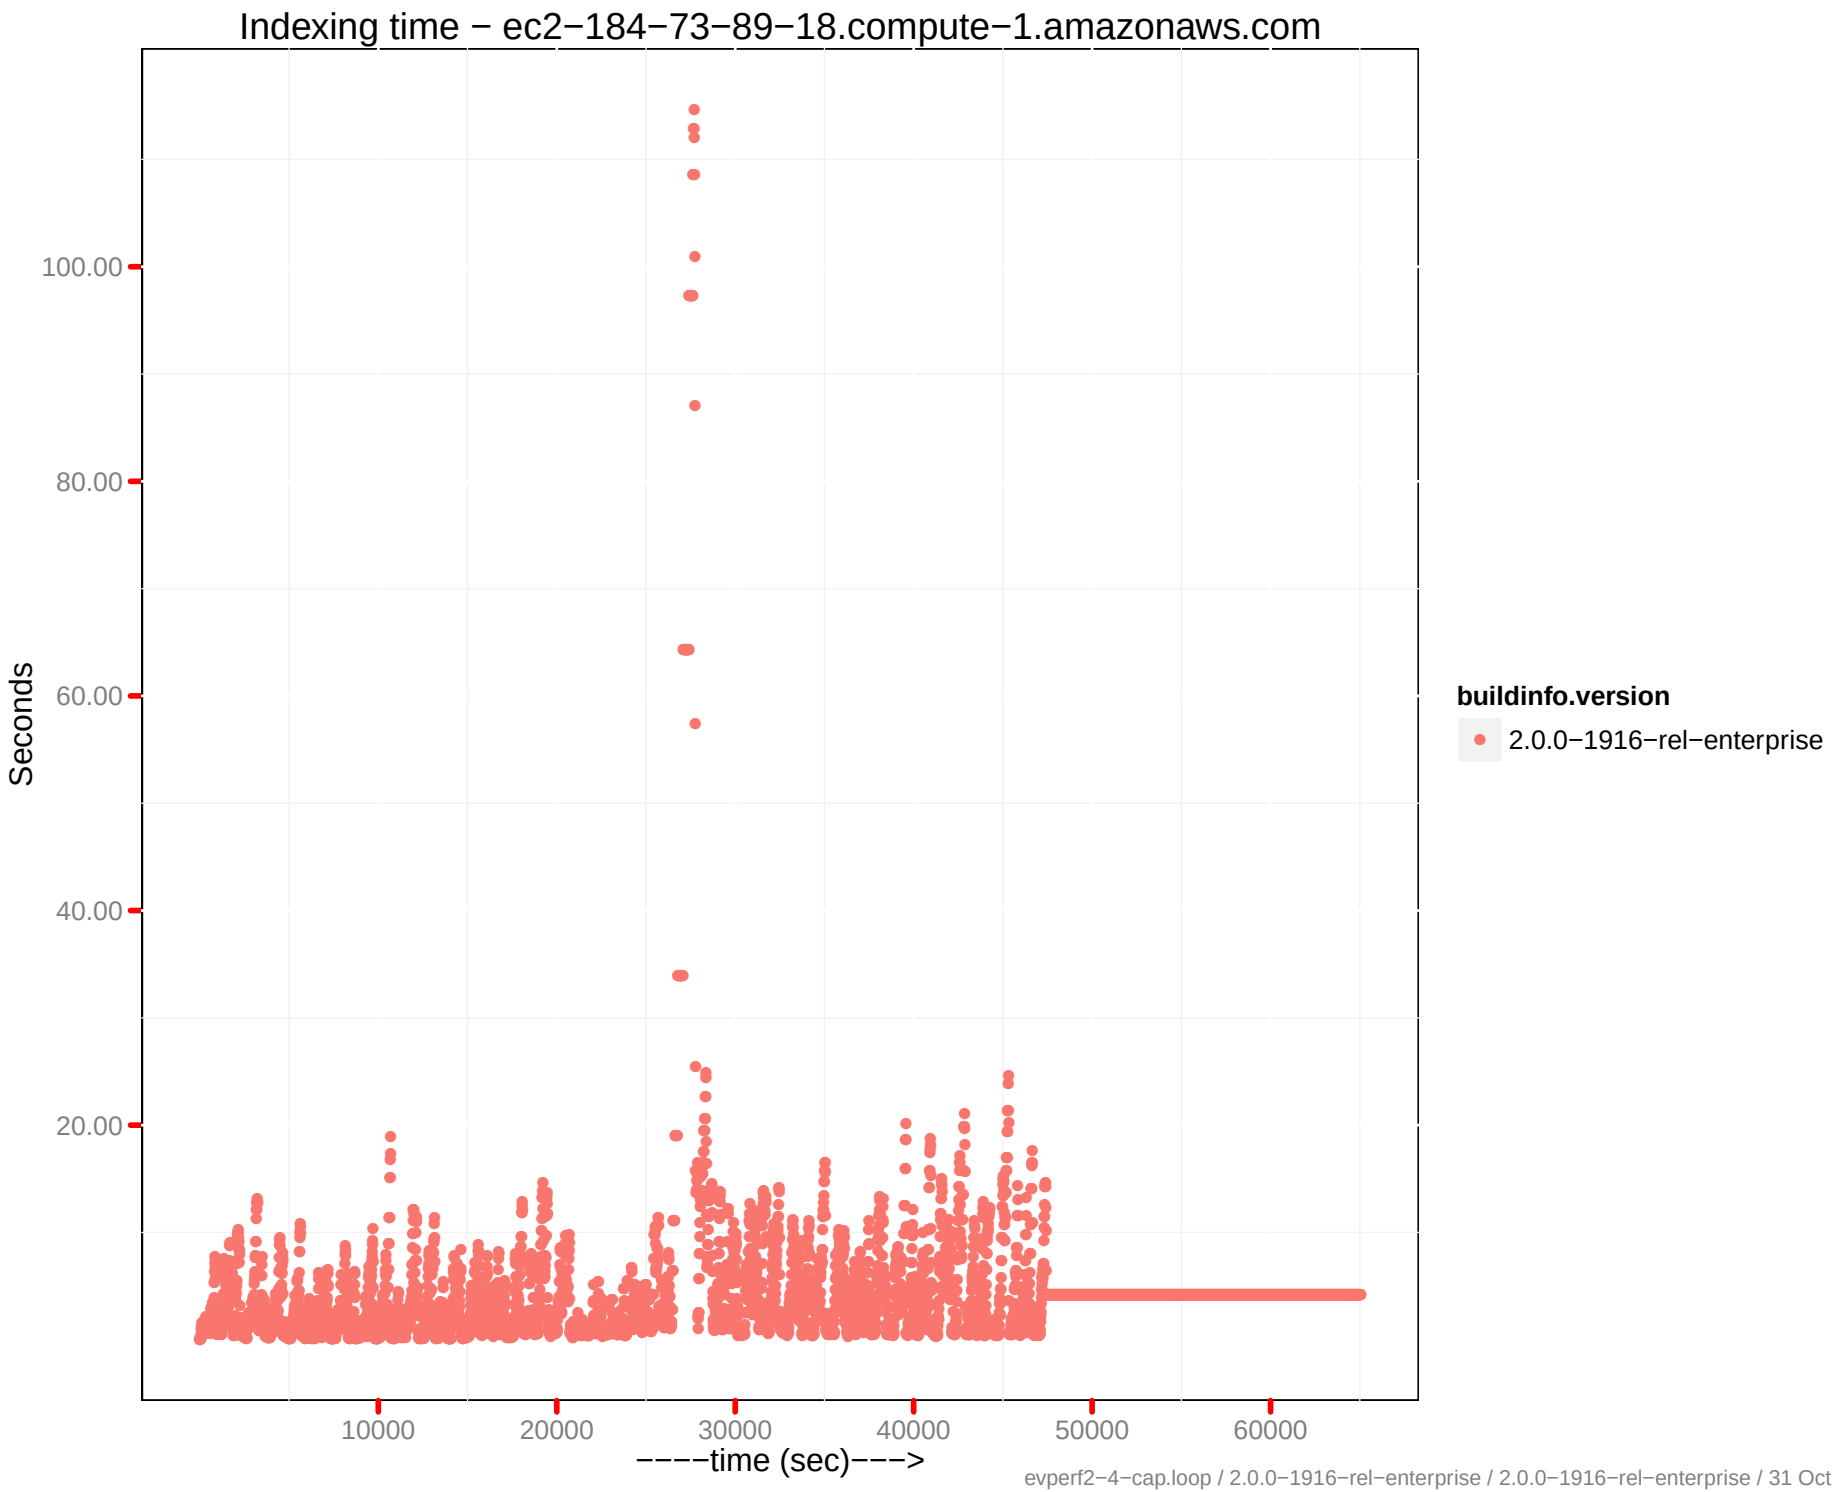
Indexing time − ec2−184−73−89−18.compute−1.amazonaws.com
Seconds
100.00
80.00
60.00
40.00
20.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

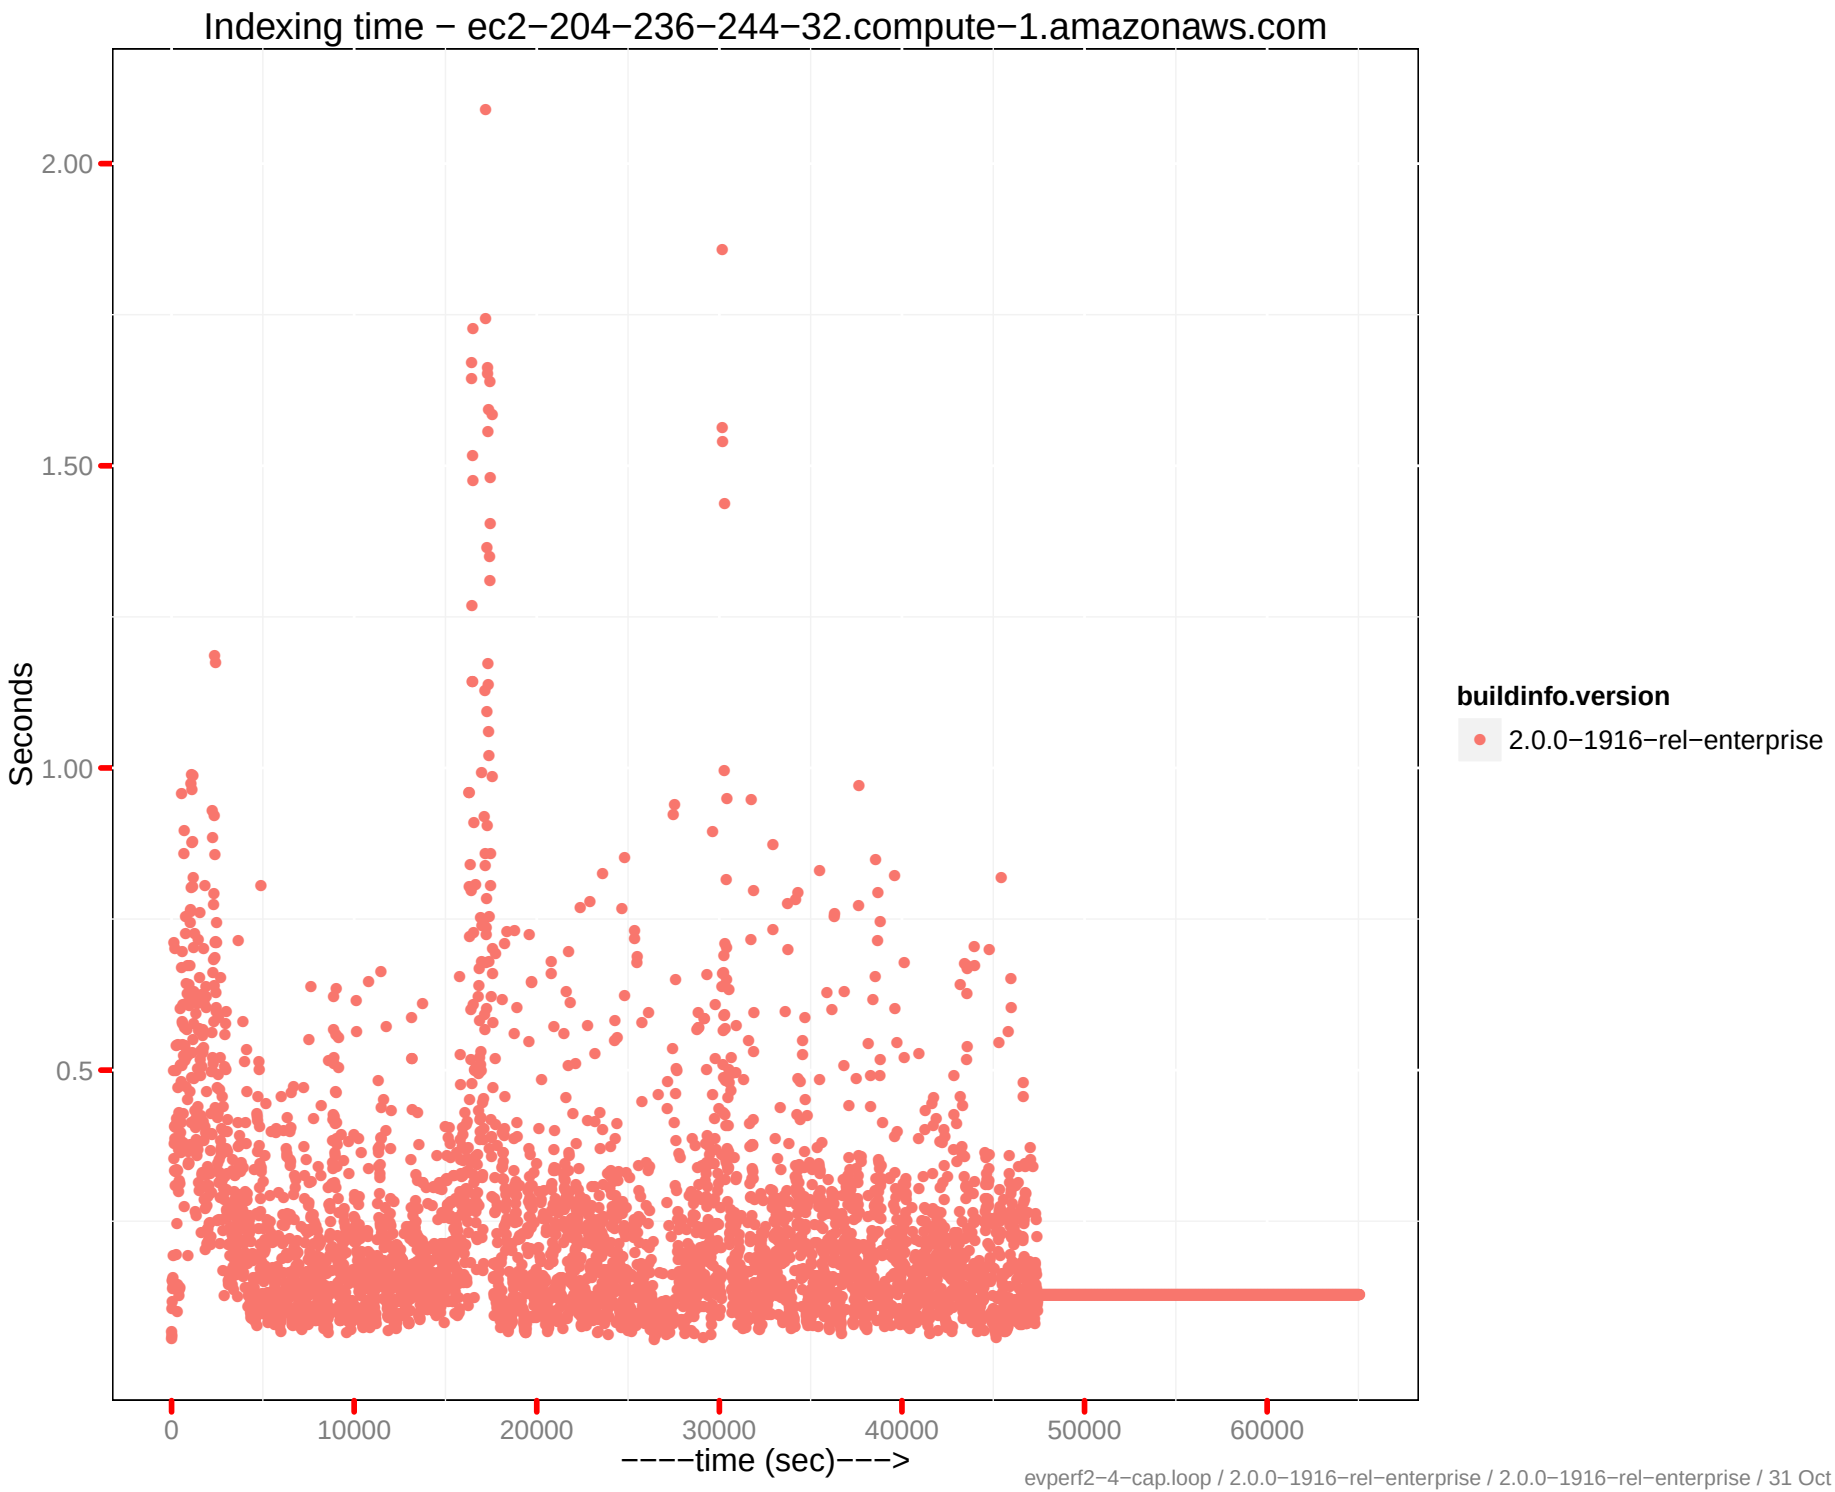
Indexing time − ec2−204−236−244−32.compute−1.amazonaws.com
Seconds
2.00
1.50
1.00
0.5
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

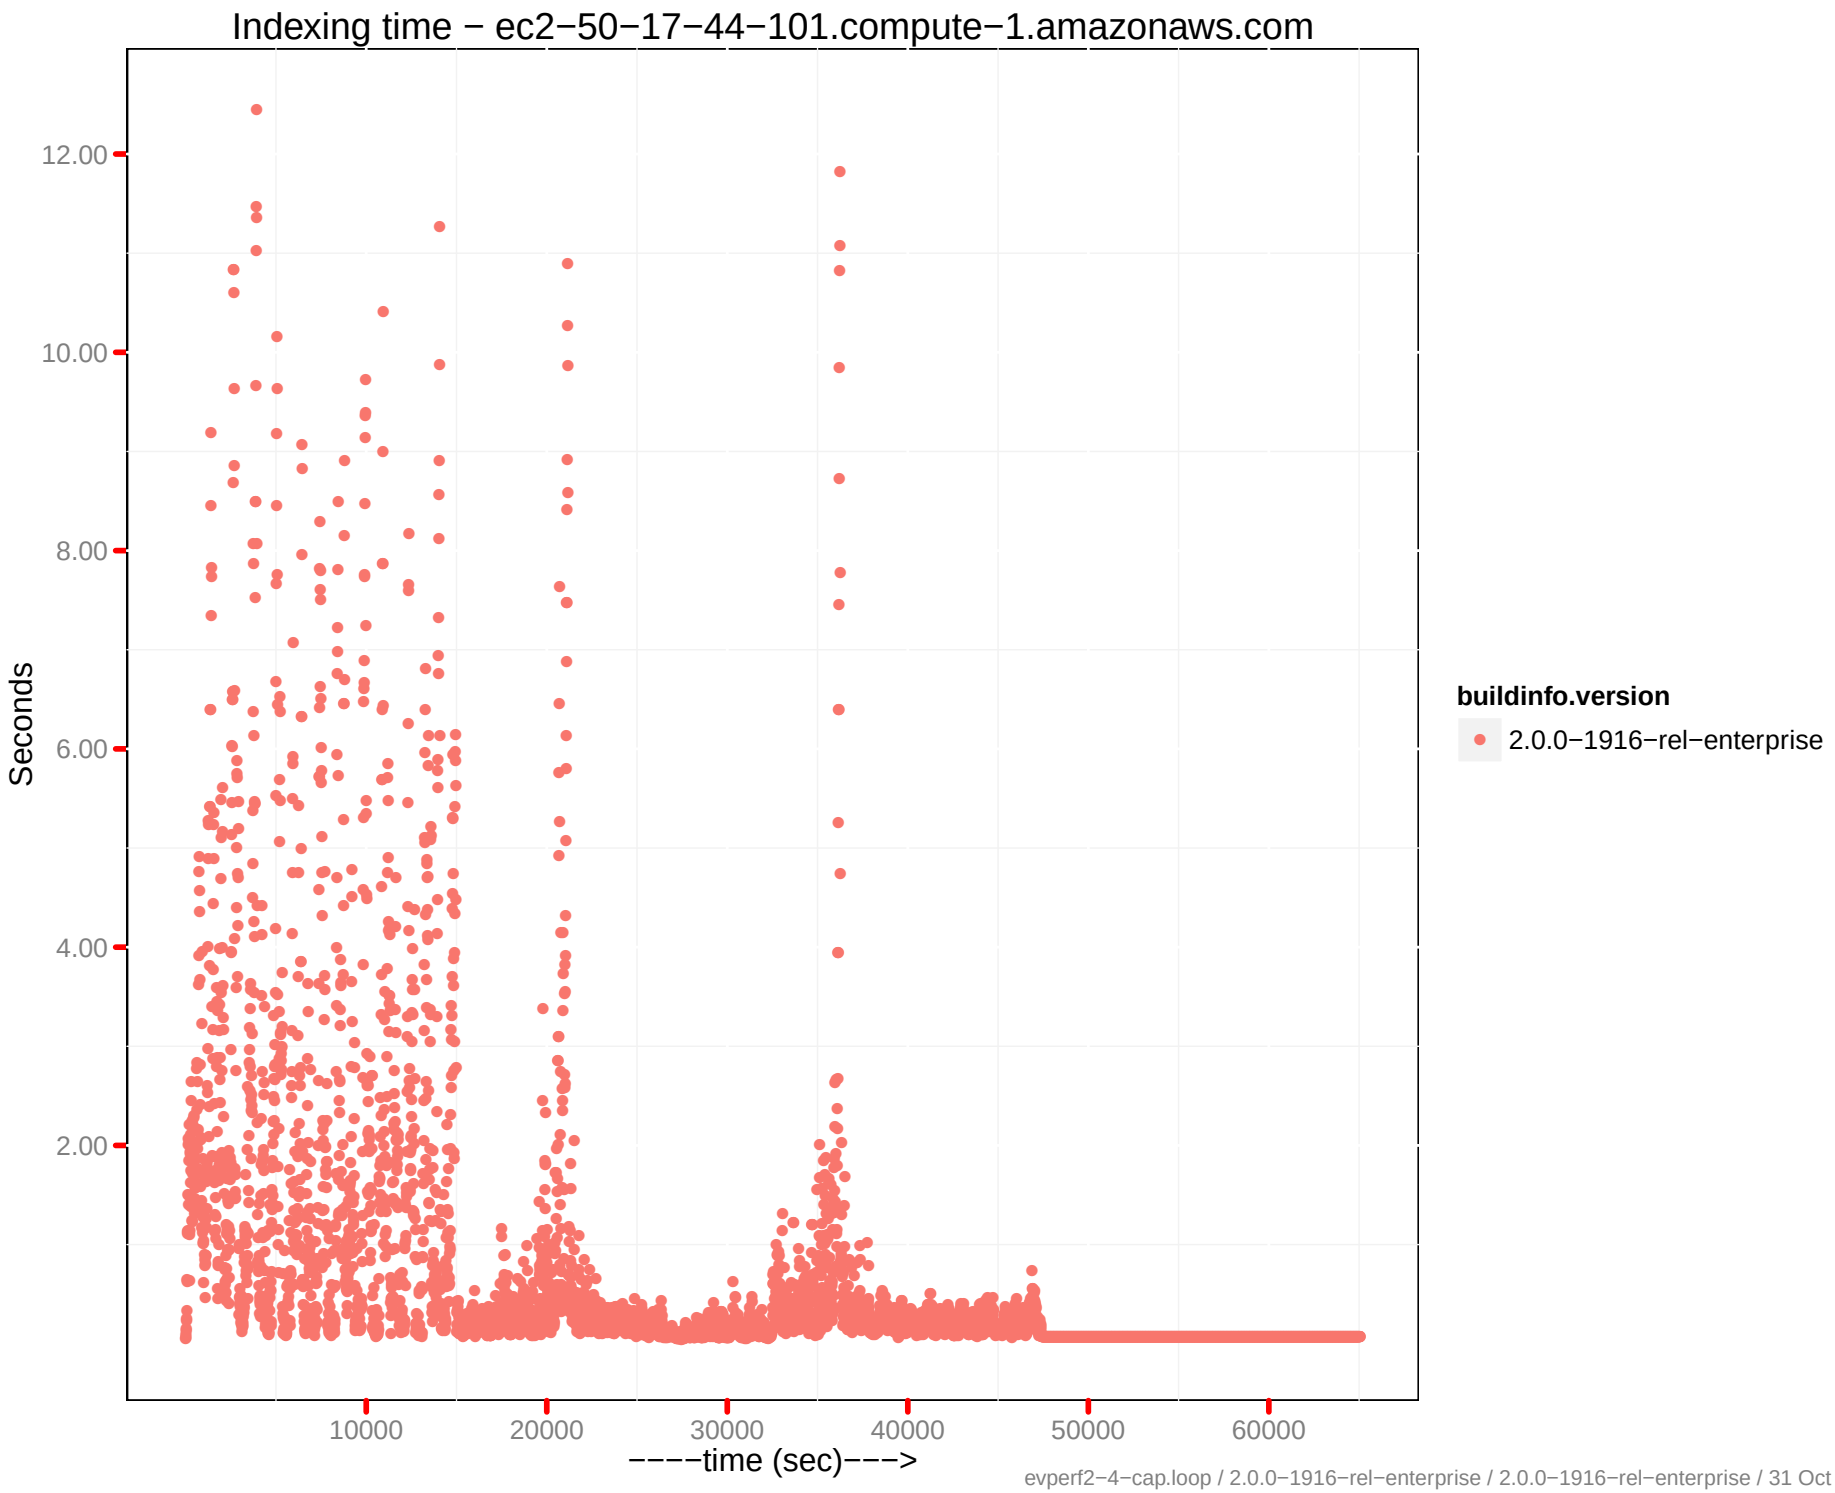
Indexing time − ec2−50−17−44−101.compute−1.amazonaws.com
12.00
10.00
8.00
6.00
4.00
2.00
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

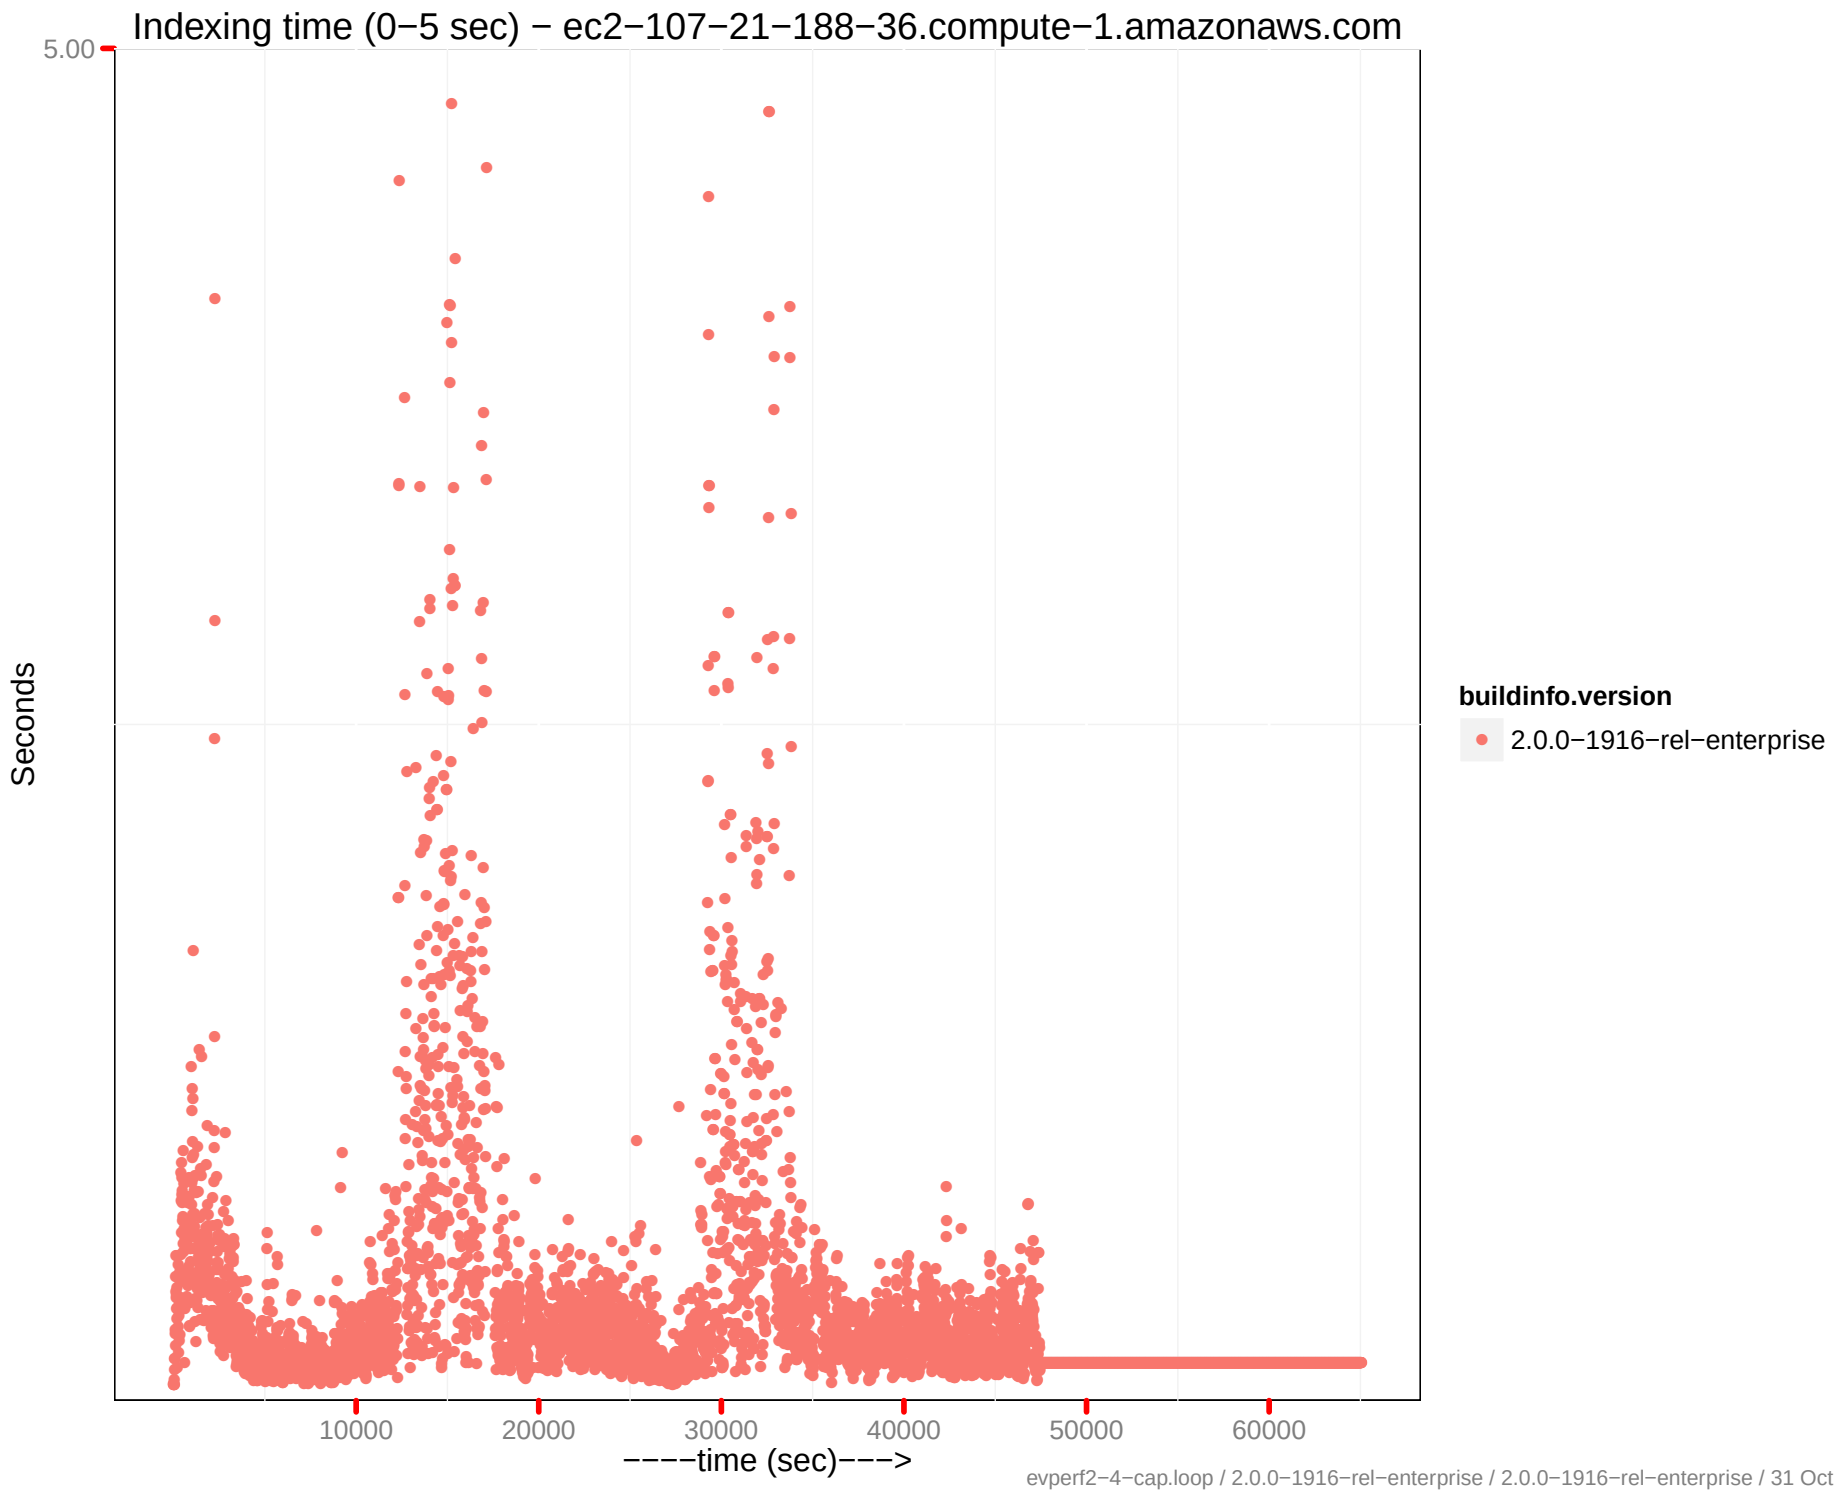
Indexing time (0−5 sec) − ec2−107−21−188−36.compute−1.amazonaws.com
5.00
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

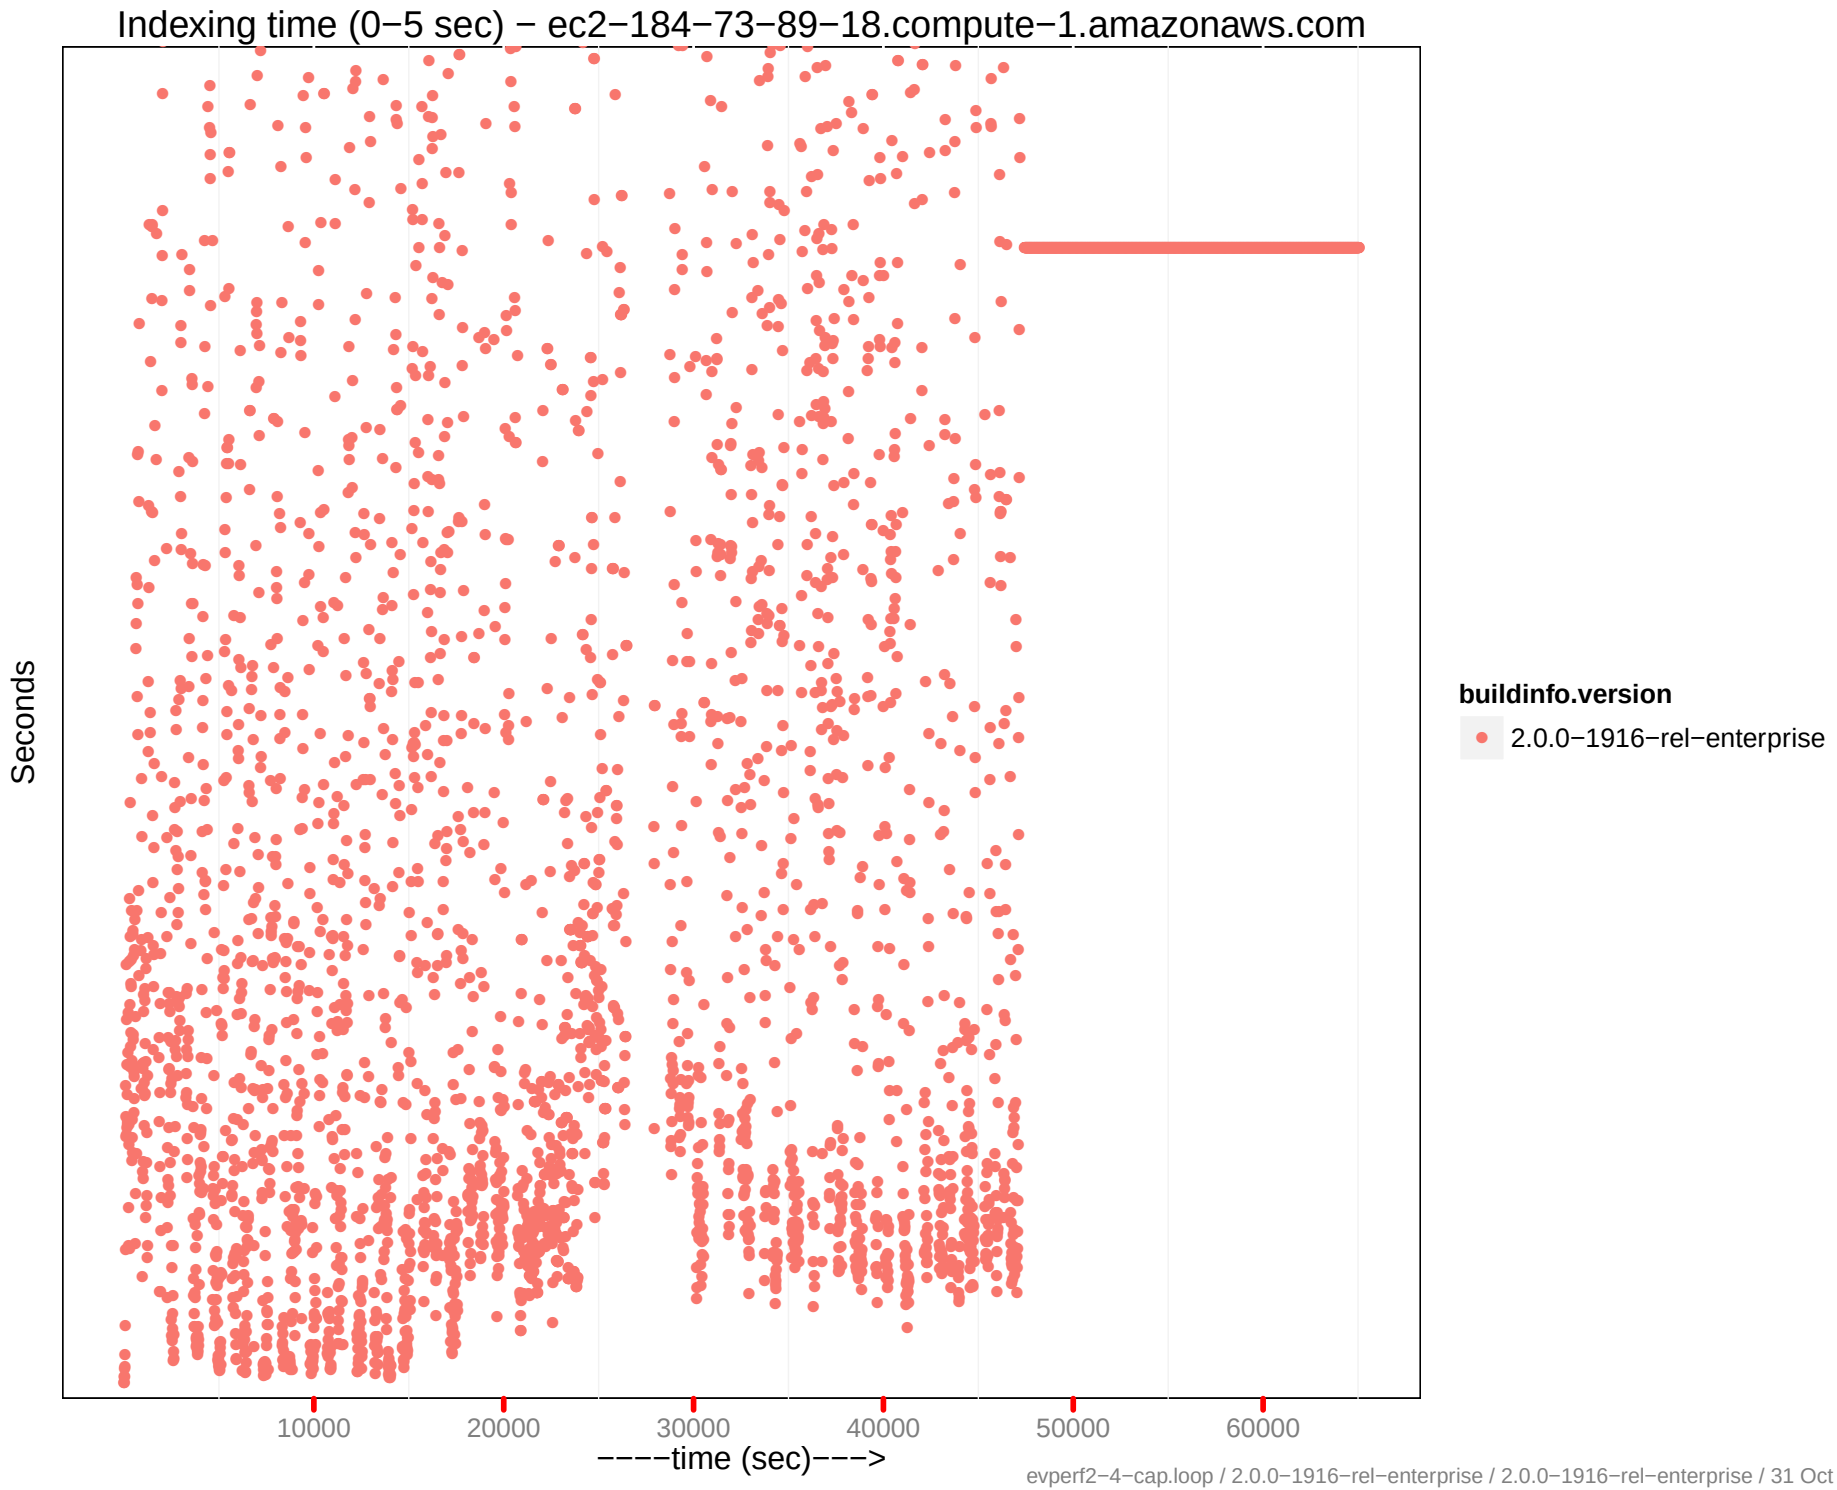
Indexing time (0−5 sec) − ec2−184−73−89−18.compute−1.amazonaws.com
Seconds
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

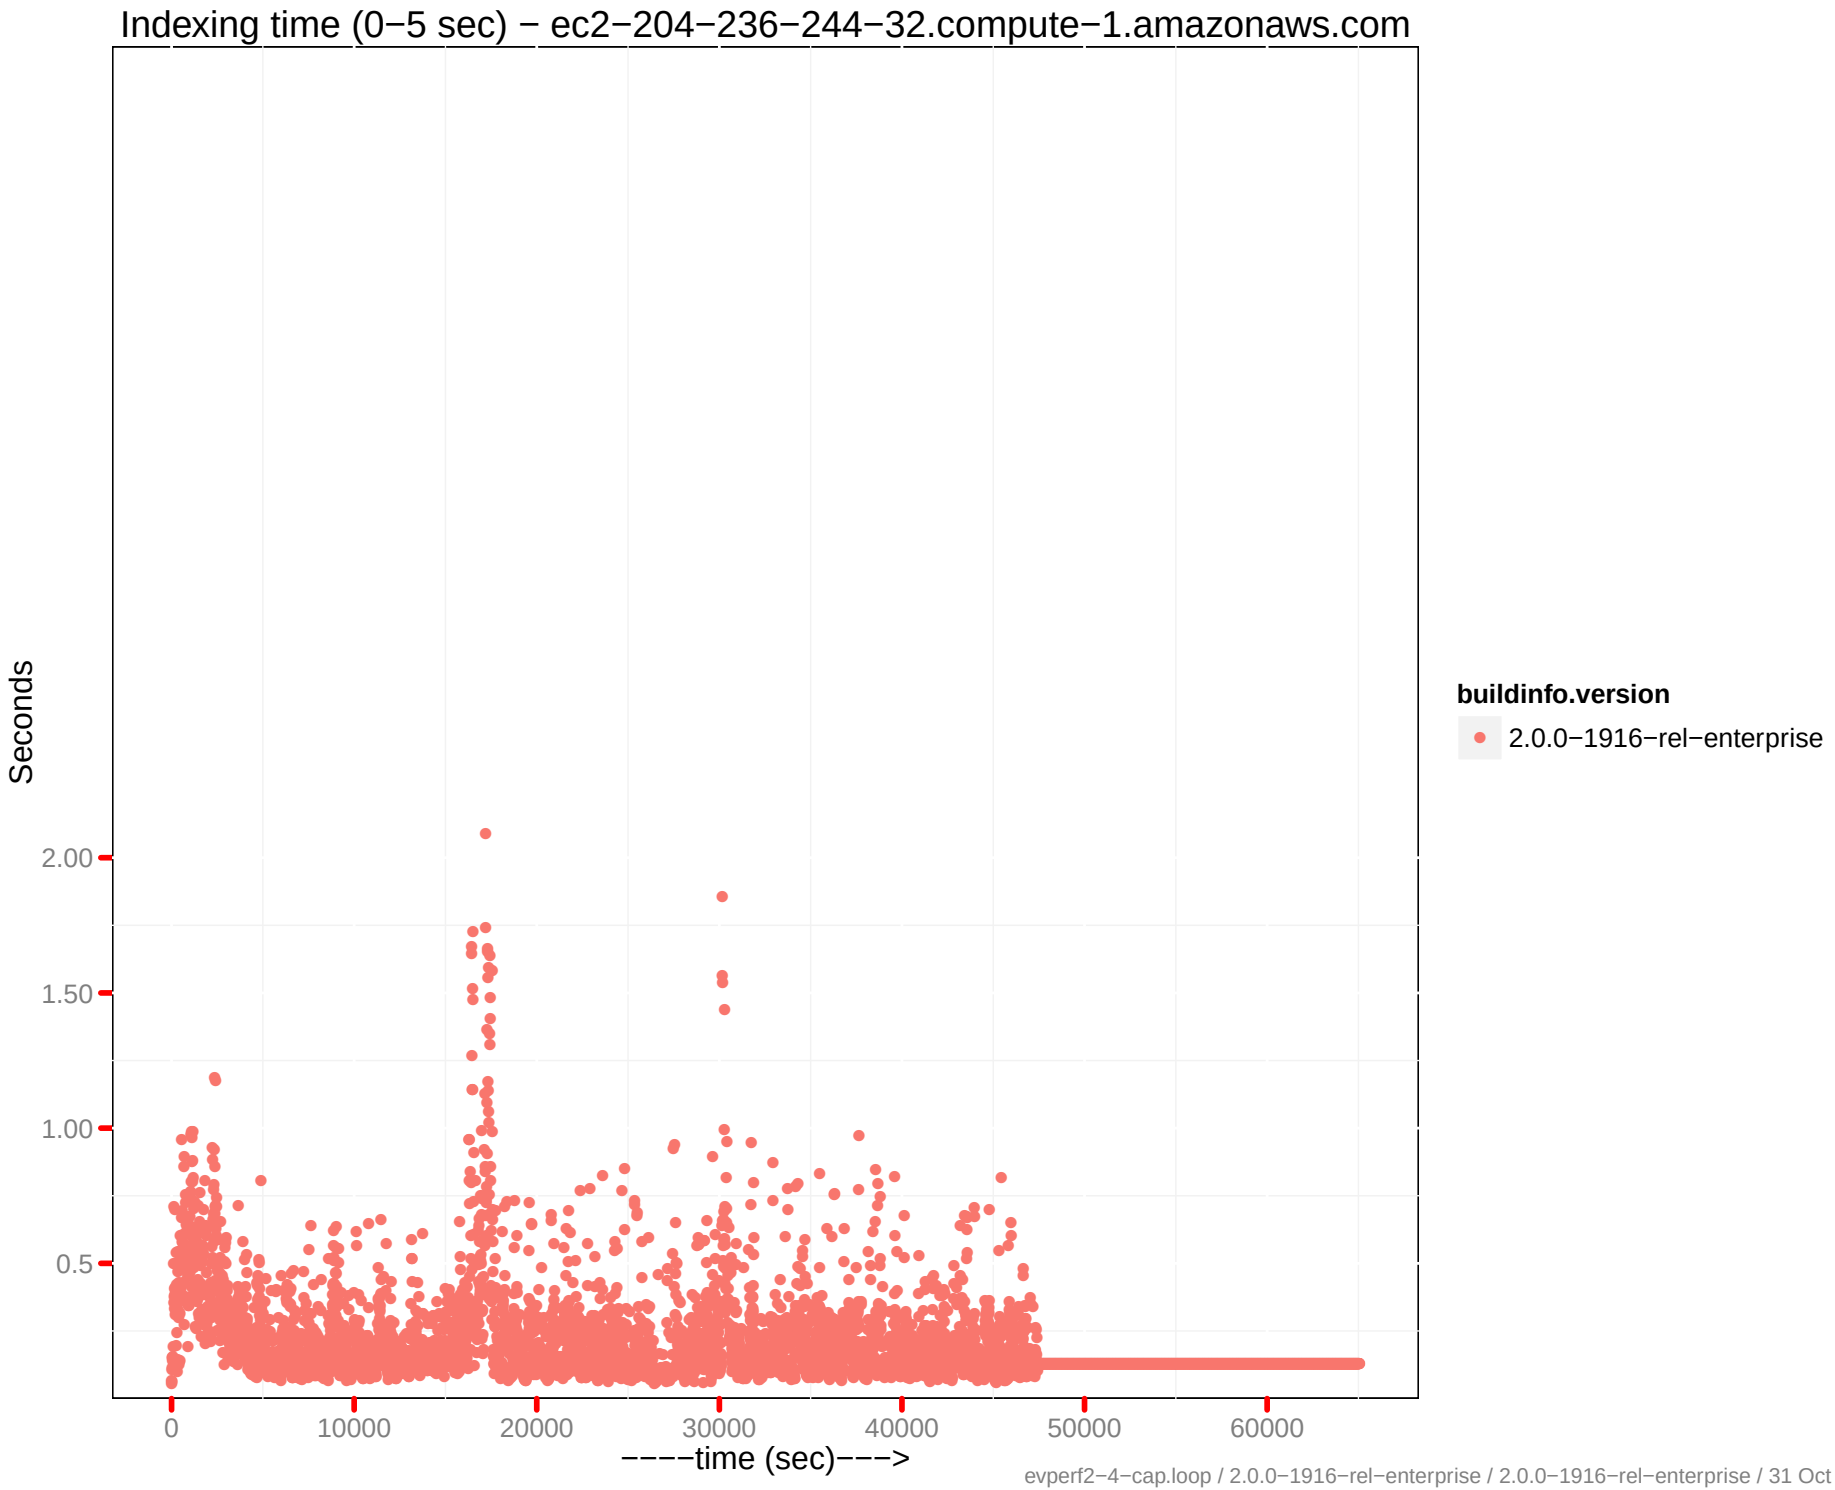
Indexing time (0−5 sec) − ec2−204−236−244−32.compute−1.amazonaws.com
Seconds
2.00
1.50
1.00
0.5
0
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

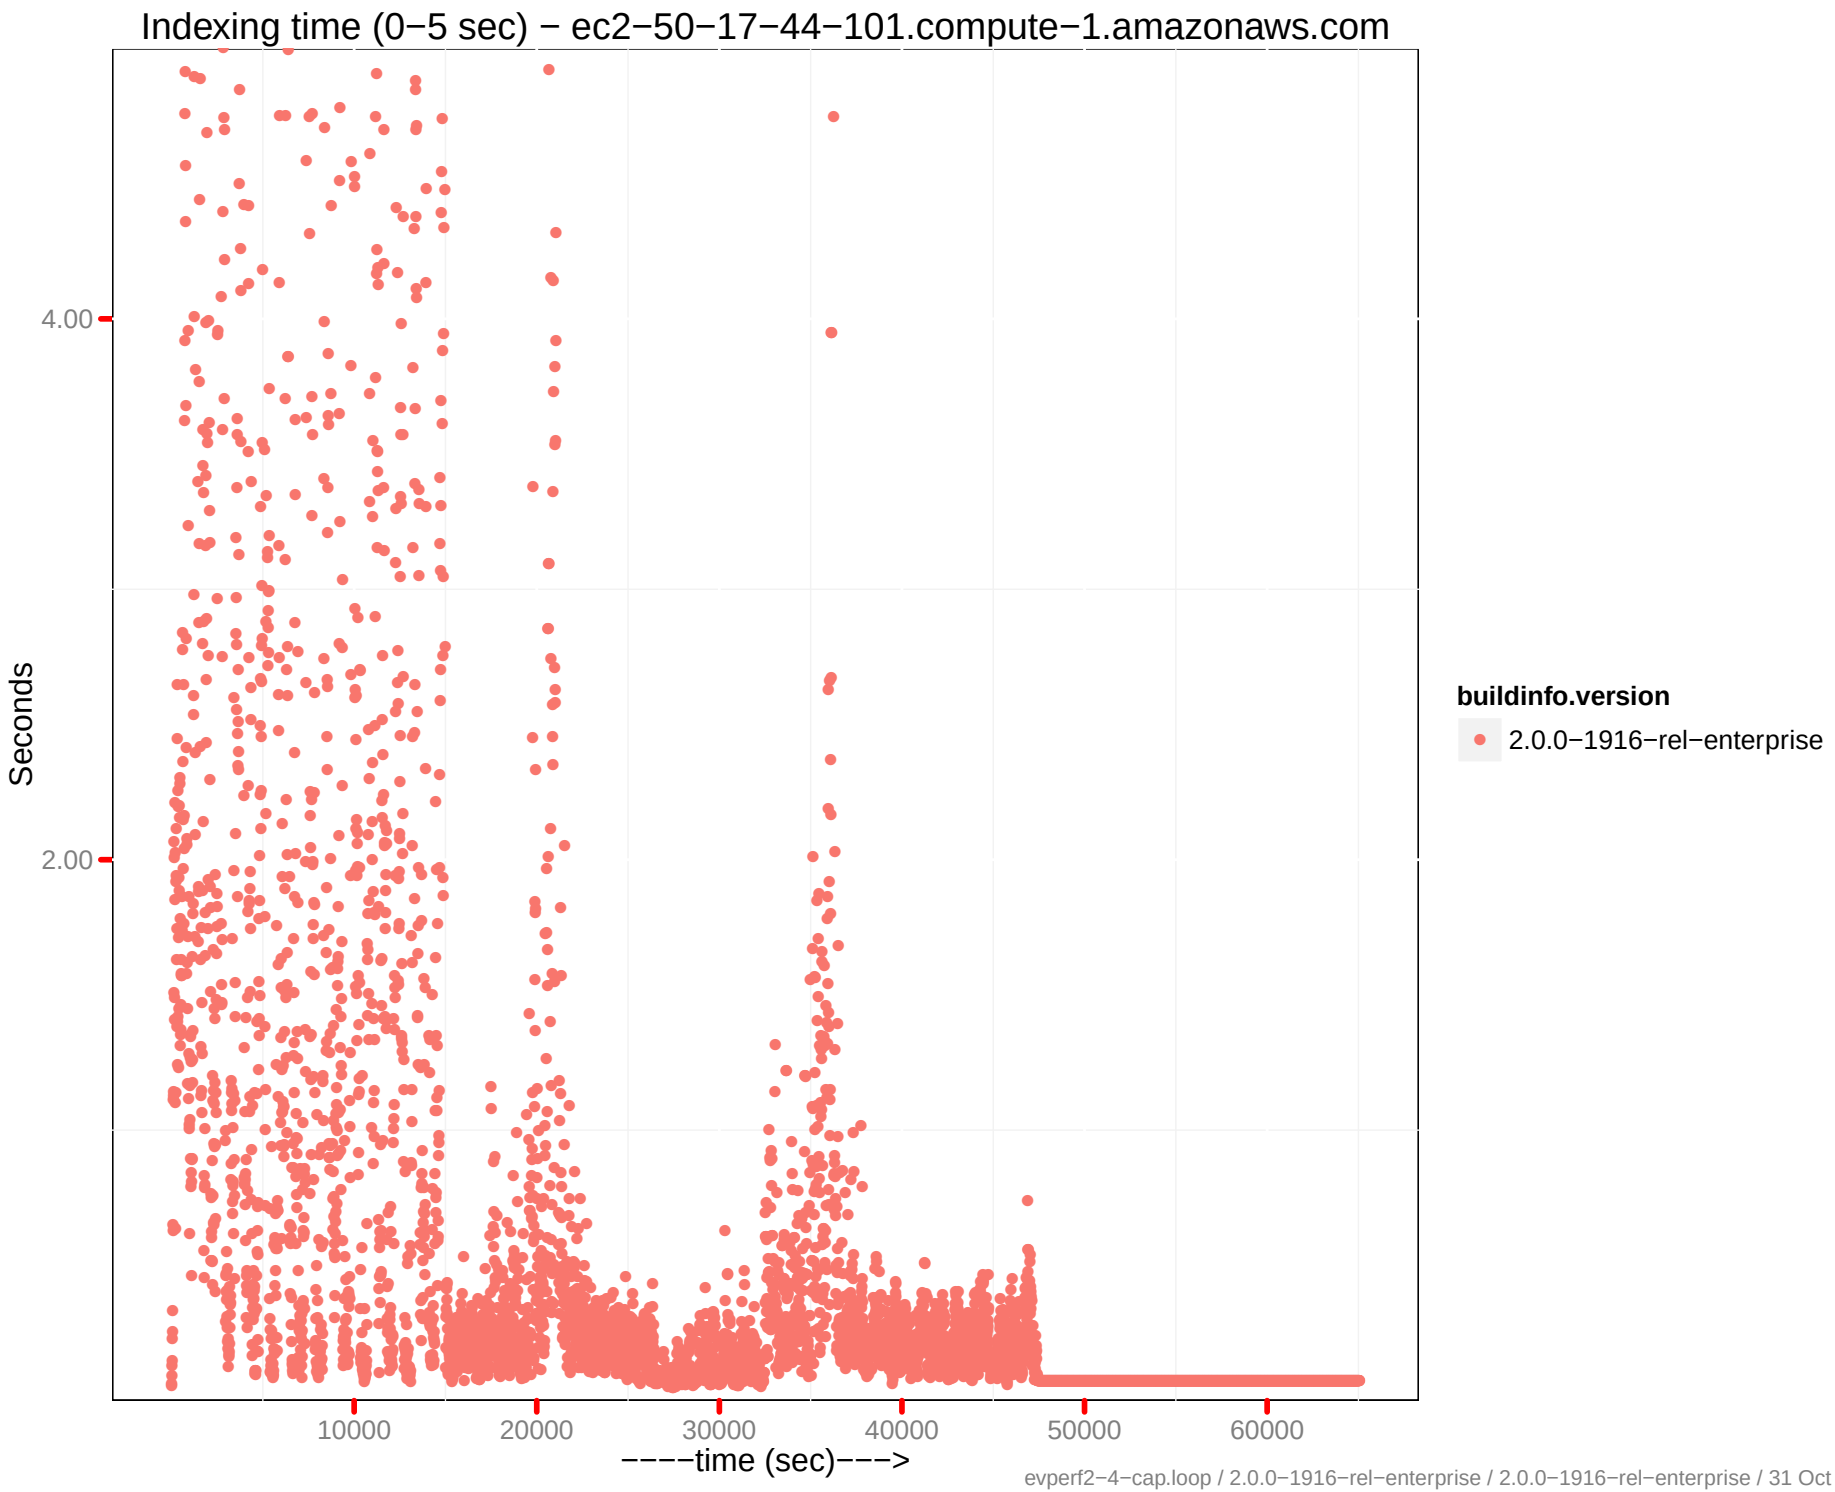
Indexing time (0−5 sec) − ec2−50−17−44−101.compute−1.amazonaws.com
Seconds
4.00
2.00
10000
20000
30000
40000
50000
60000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1916−rel−enterprise
evperf2−4−cap.loop / 2.0.0−1916−rel−enterprise / 2.0.0−1916−rel−enterprise / 31 Oct 2012

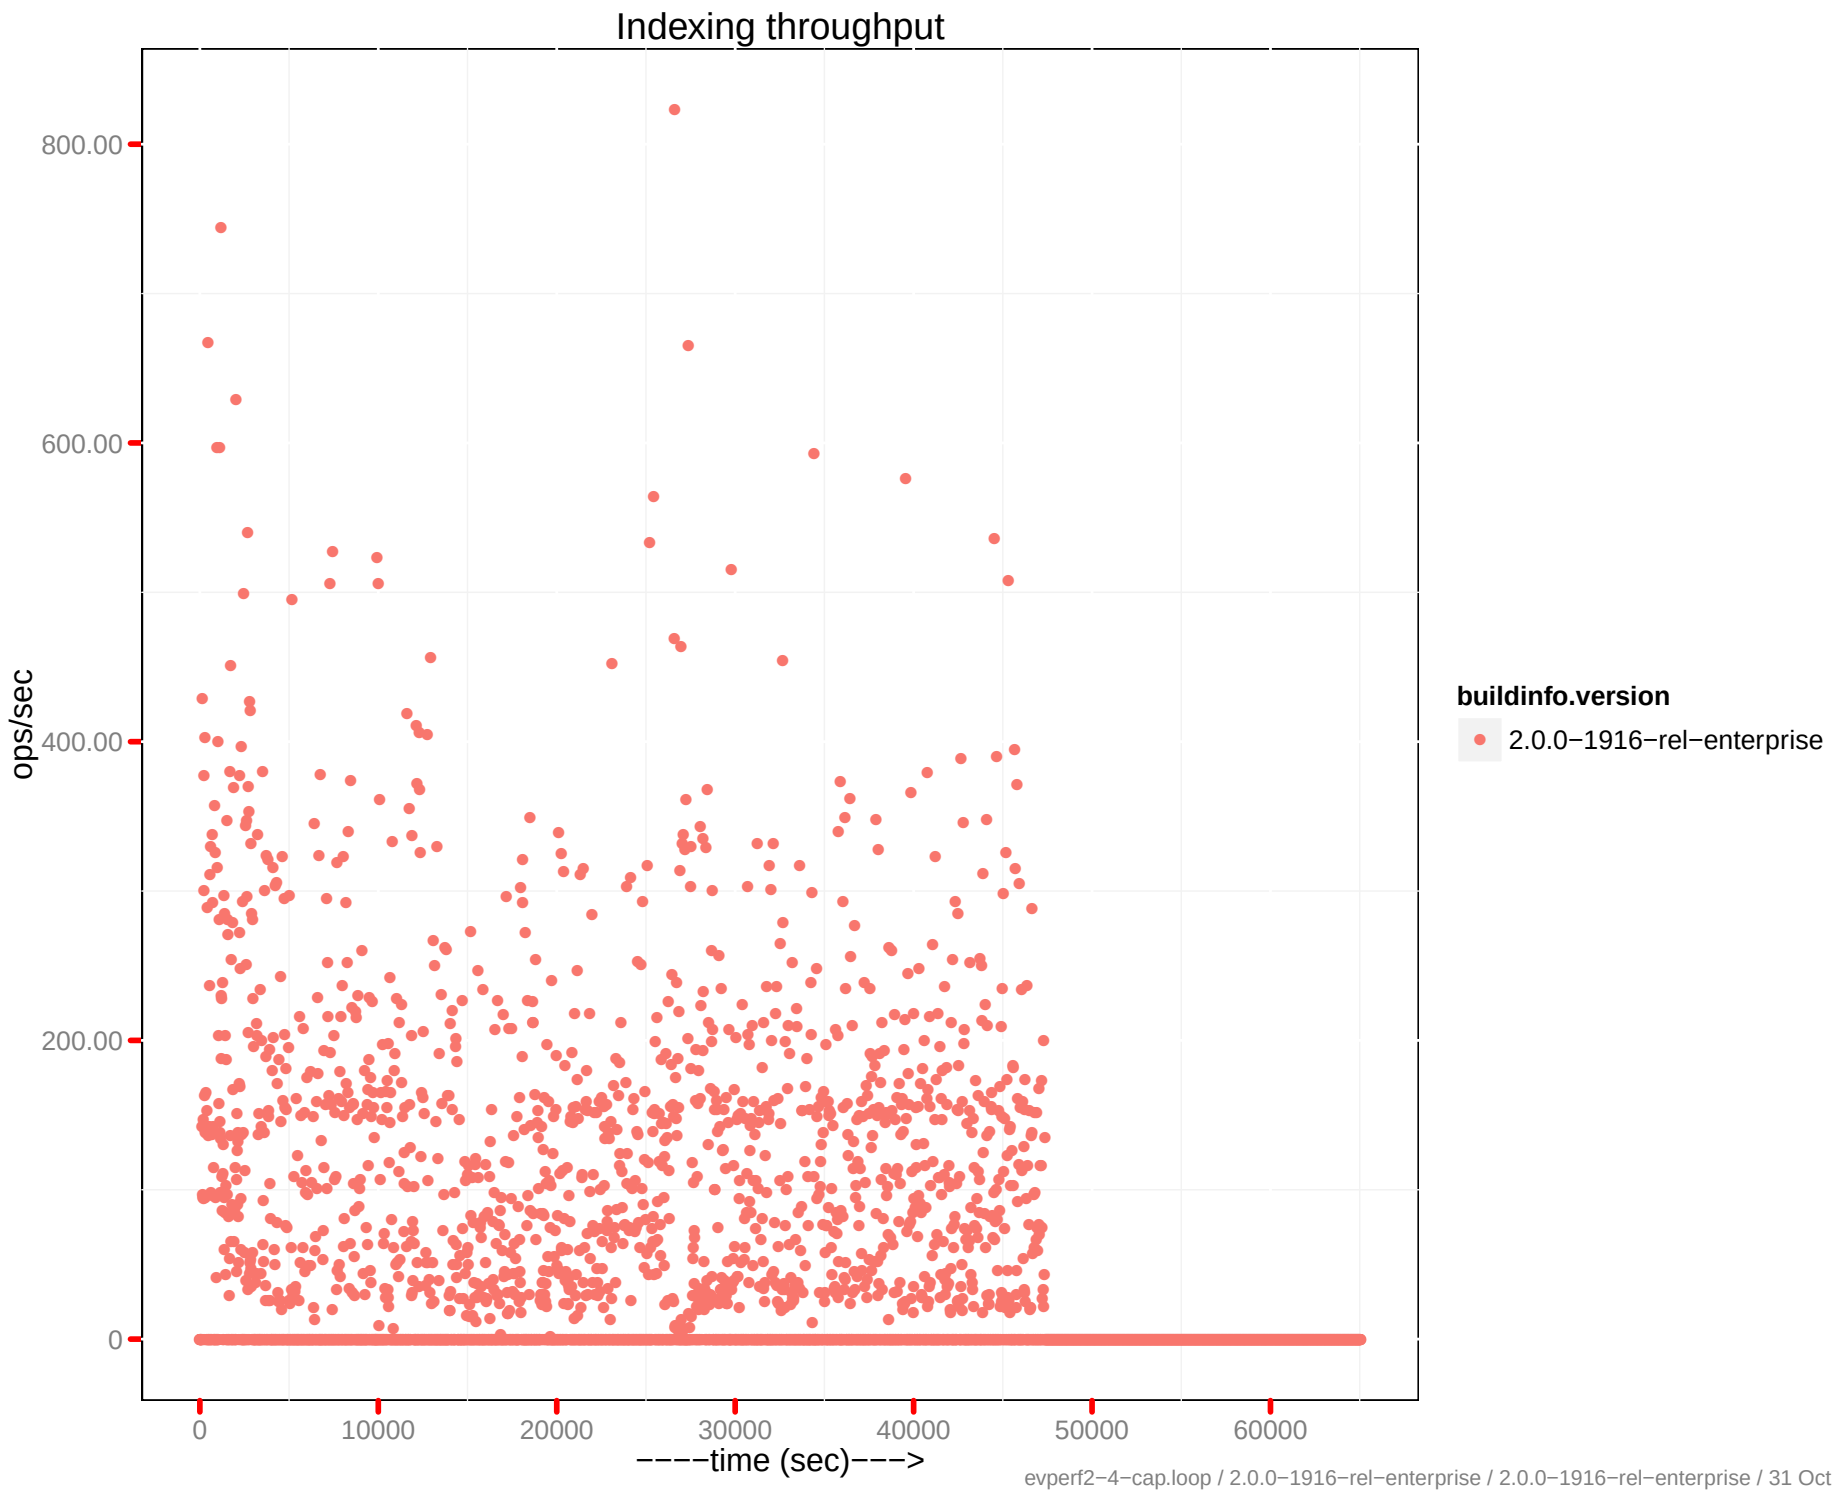
Indexing throughput
800.00
600.00
400.00
200.00
0
ops/sec
0
10000
20000
30000
40000
50000
60000
----time (sec)--->
buildinfo.version
2.0.0-1916-rel-enterprise
evperf2-4-cap.loop / 2.0.0-1916-rel-enterprise / 2.0.0-1916-rel-enterprise / 31 Oct 2012

```
evperf2-4-cap.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 45M total queries (tuned to be ~14 hours)

performance.iperf.MultiClientTests.test_evperf2

params:

# general
batch=50
kind=json
mem_quota=25000

# load phase
items=20000000
hot_init_items=5000000

# 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_per_sec=30
fg_max_ops=45000000
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
```

```
ec2.ini
[global]
username:ec2-user
ssh_key:/home/jenkins/qa-key.pem
port:8091
data_path:/data

[servers]
1:10.104.101.200
2:10.217.55.200
3:10.71.83.95
4:10.104.53.243

[membase]
rest_username:Administrator
rest_password:password
```