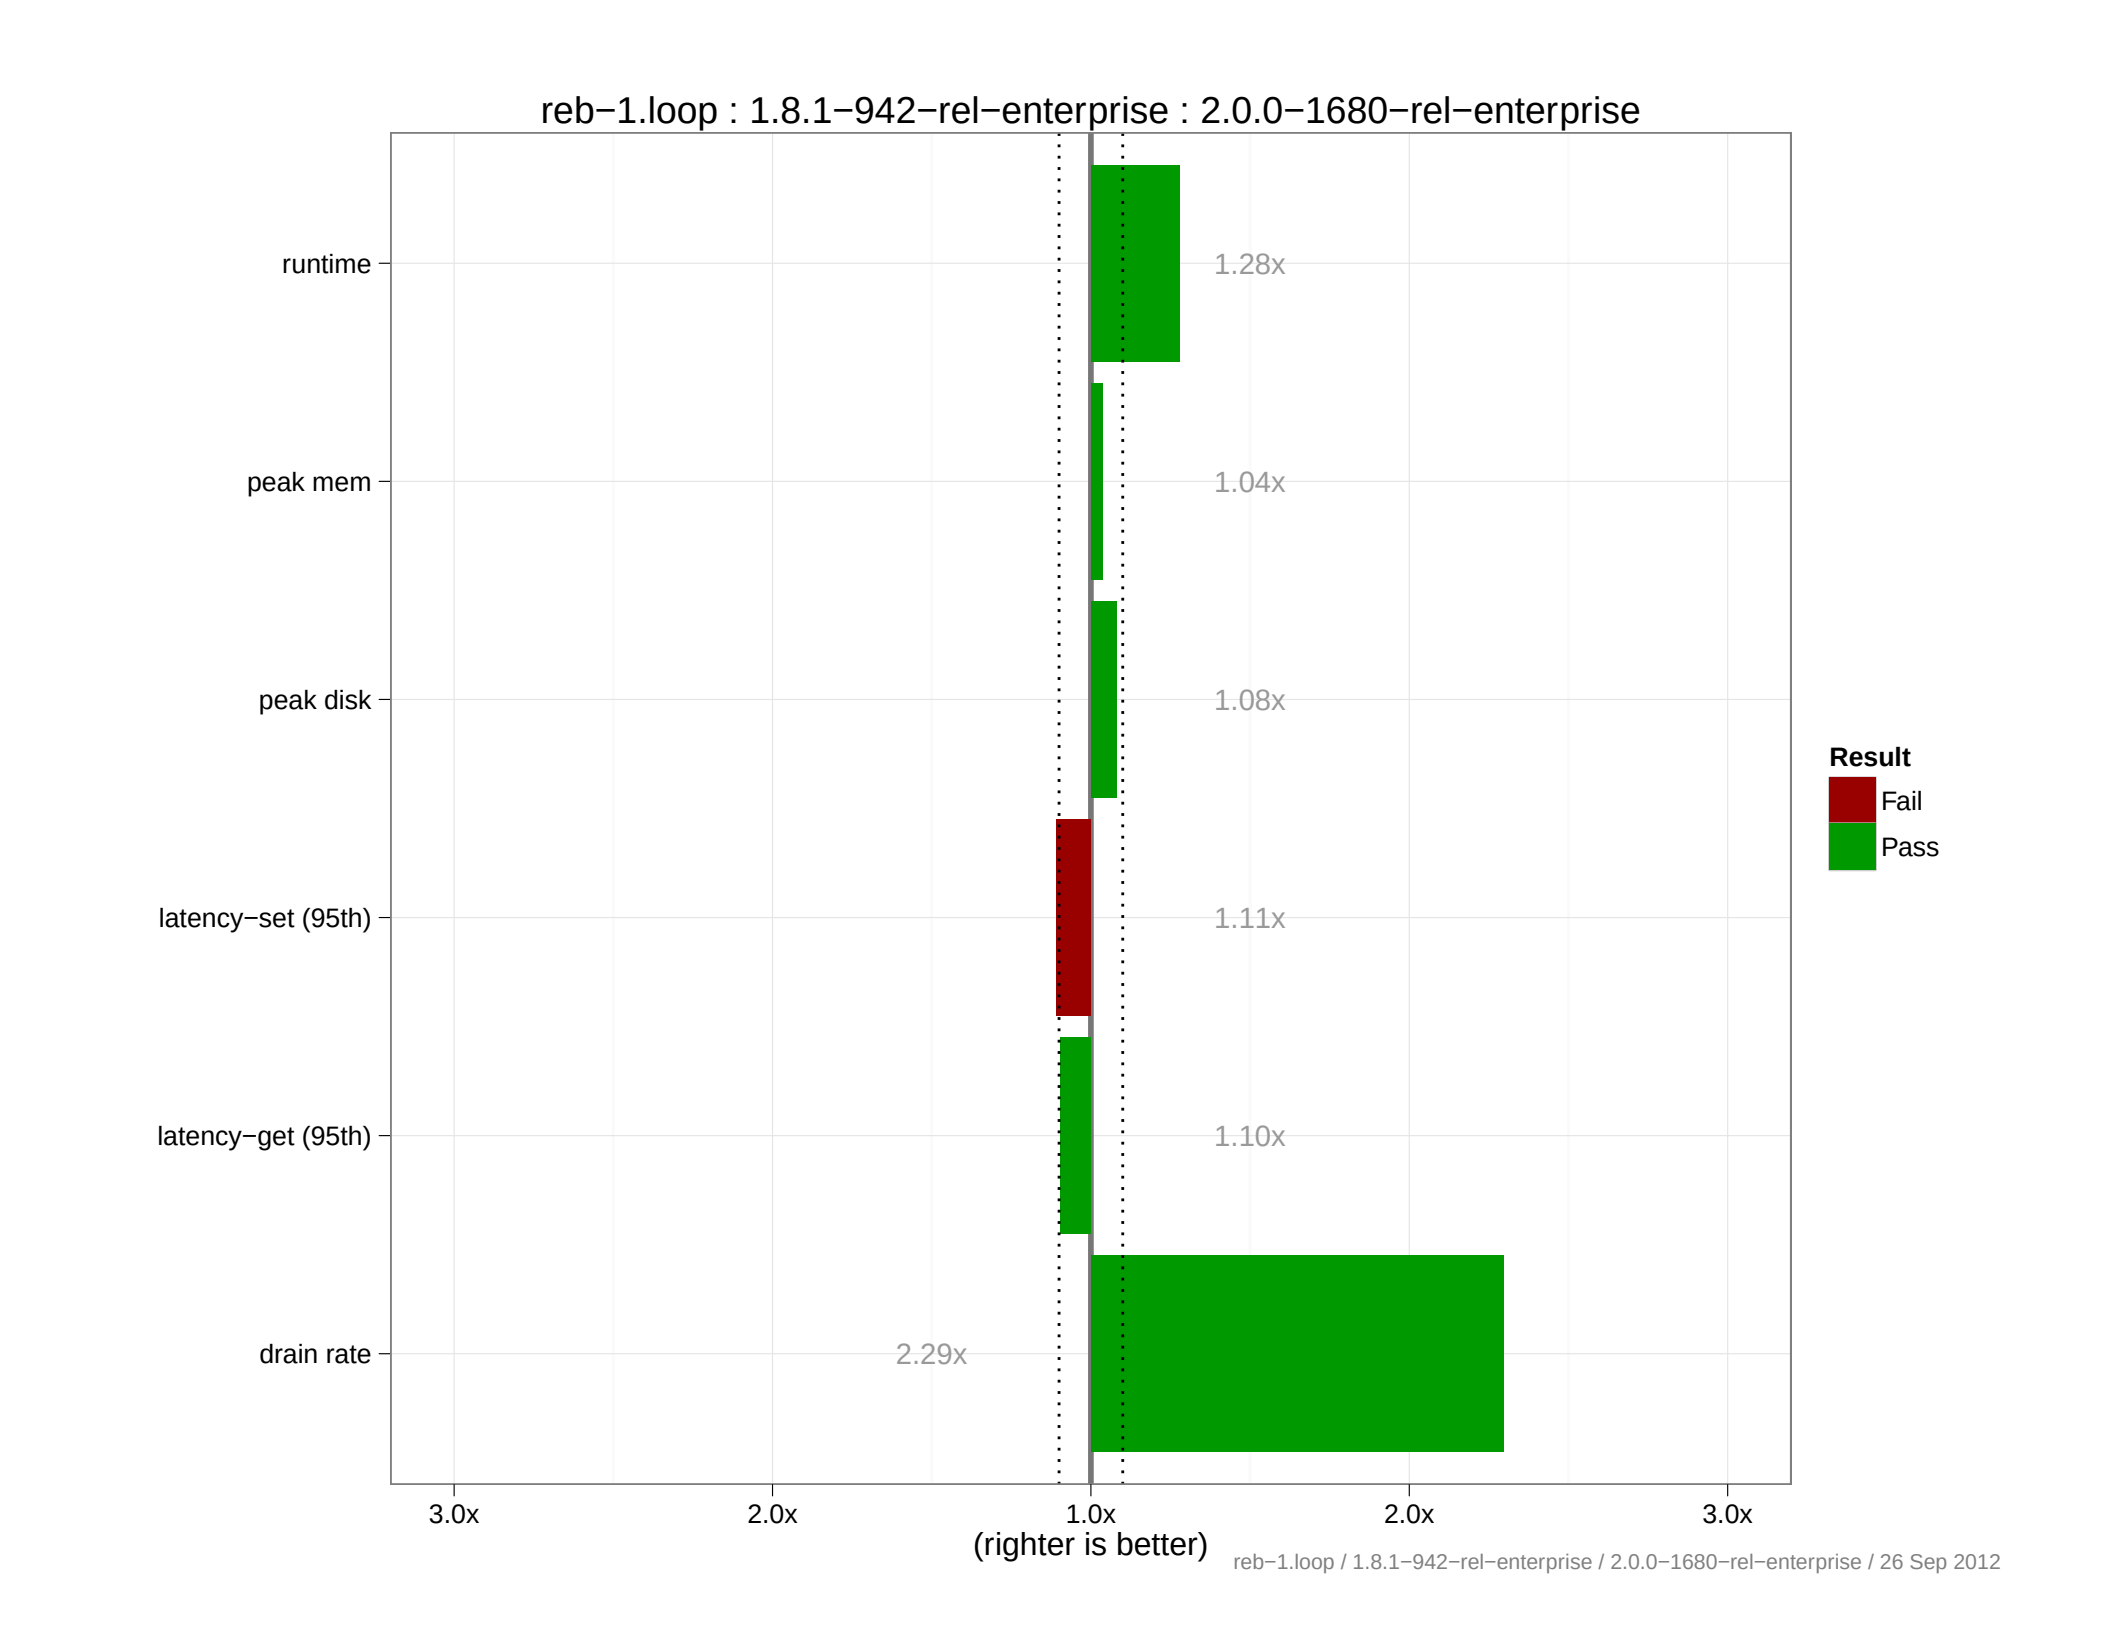
reb−1.loop : 1.8.1−942−rel−enterprise : 2.0.0−1680−rel−enterprise
runtime
peak mem
peak disk
latency−set (95th)
latency−get (95th)
drain rate
1.28x
1.04x
1.08x
1.11x
1.10x
2.29x
Result
Fail
Pass
3.0x
2.0x
1.0x
2.0x
3.0x
(righter is better)
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

	1.8.1 − 942	2.0.0 − 1680
Runtime (in hr)	3	2.34
Avg. Drain Rate	1.65K	3.79K
Peak Disk (GB)	48.43	44.82
Peak Memory (GB)	17.01	16.39
Avg. OPS	6.79K	8.39K
Avg. mem memcached (GB)	16.87	15.98
Avg. mem beam.smp (MB)	90.22	349.03
Avg. CPU rate (%)	41.98	61.82
Latency−get (90th) (ms)	3.83	4.12
Latency−get (95th) (ms)	5.86	6.42
Latency−get (99th) (ms)	16.55	19.14
Latency−set (90th) (ms)	3.89	4.23
Latency−set (95th) (ms)	5.9	6.55
Latency−set (99th) (ms)	15.24	18.92
Latency−query (80th) (ms)	NA	NA
Latency−query (90th) (ms)	NA	NA
Latency−query (95th) (ms)	NA	NA
Latency−query (99th) (ms)	NA	NA
Latency−query (99.9th) (ms)	NA	NA
Avg. QPS	0	0
XDC ops/sec	NA	NaN
Rebalance Time (sec)	5035.43	1998.91
Testrunner Version	2543397	46a600d

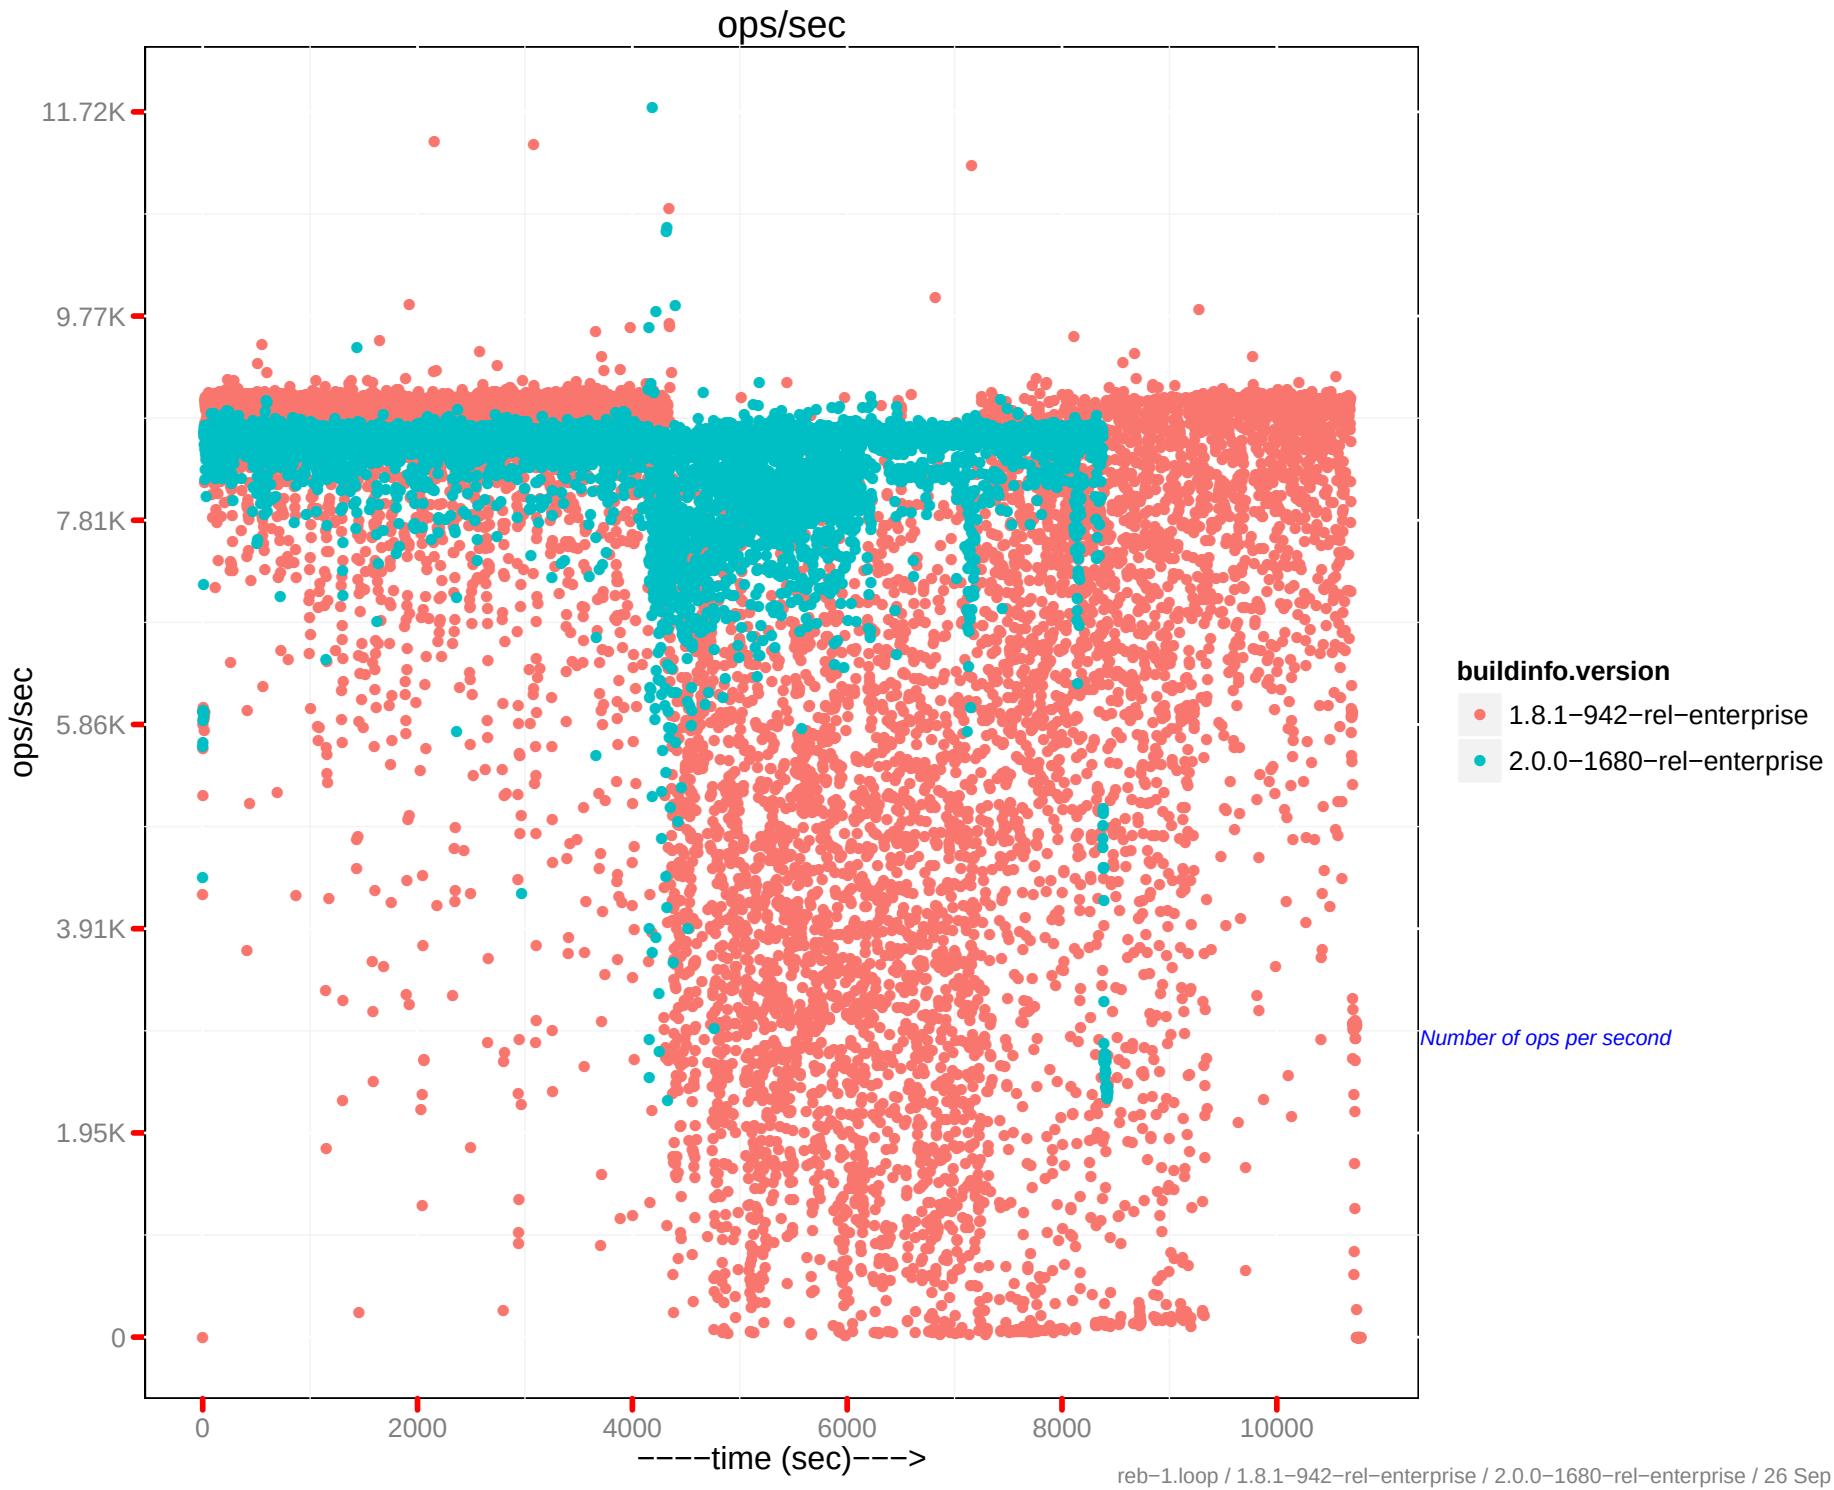
ops/sec
11.72K
9.77K
7.81K
5.86K
3.91K
1.95K
0
ops/sec
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of ops per second
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

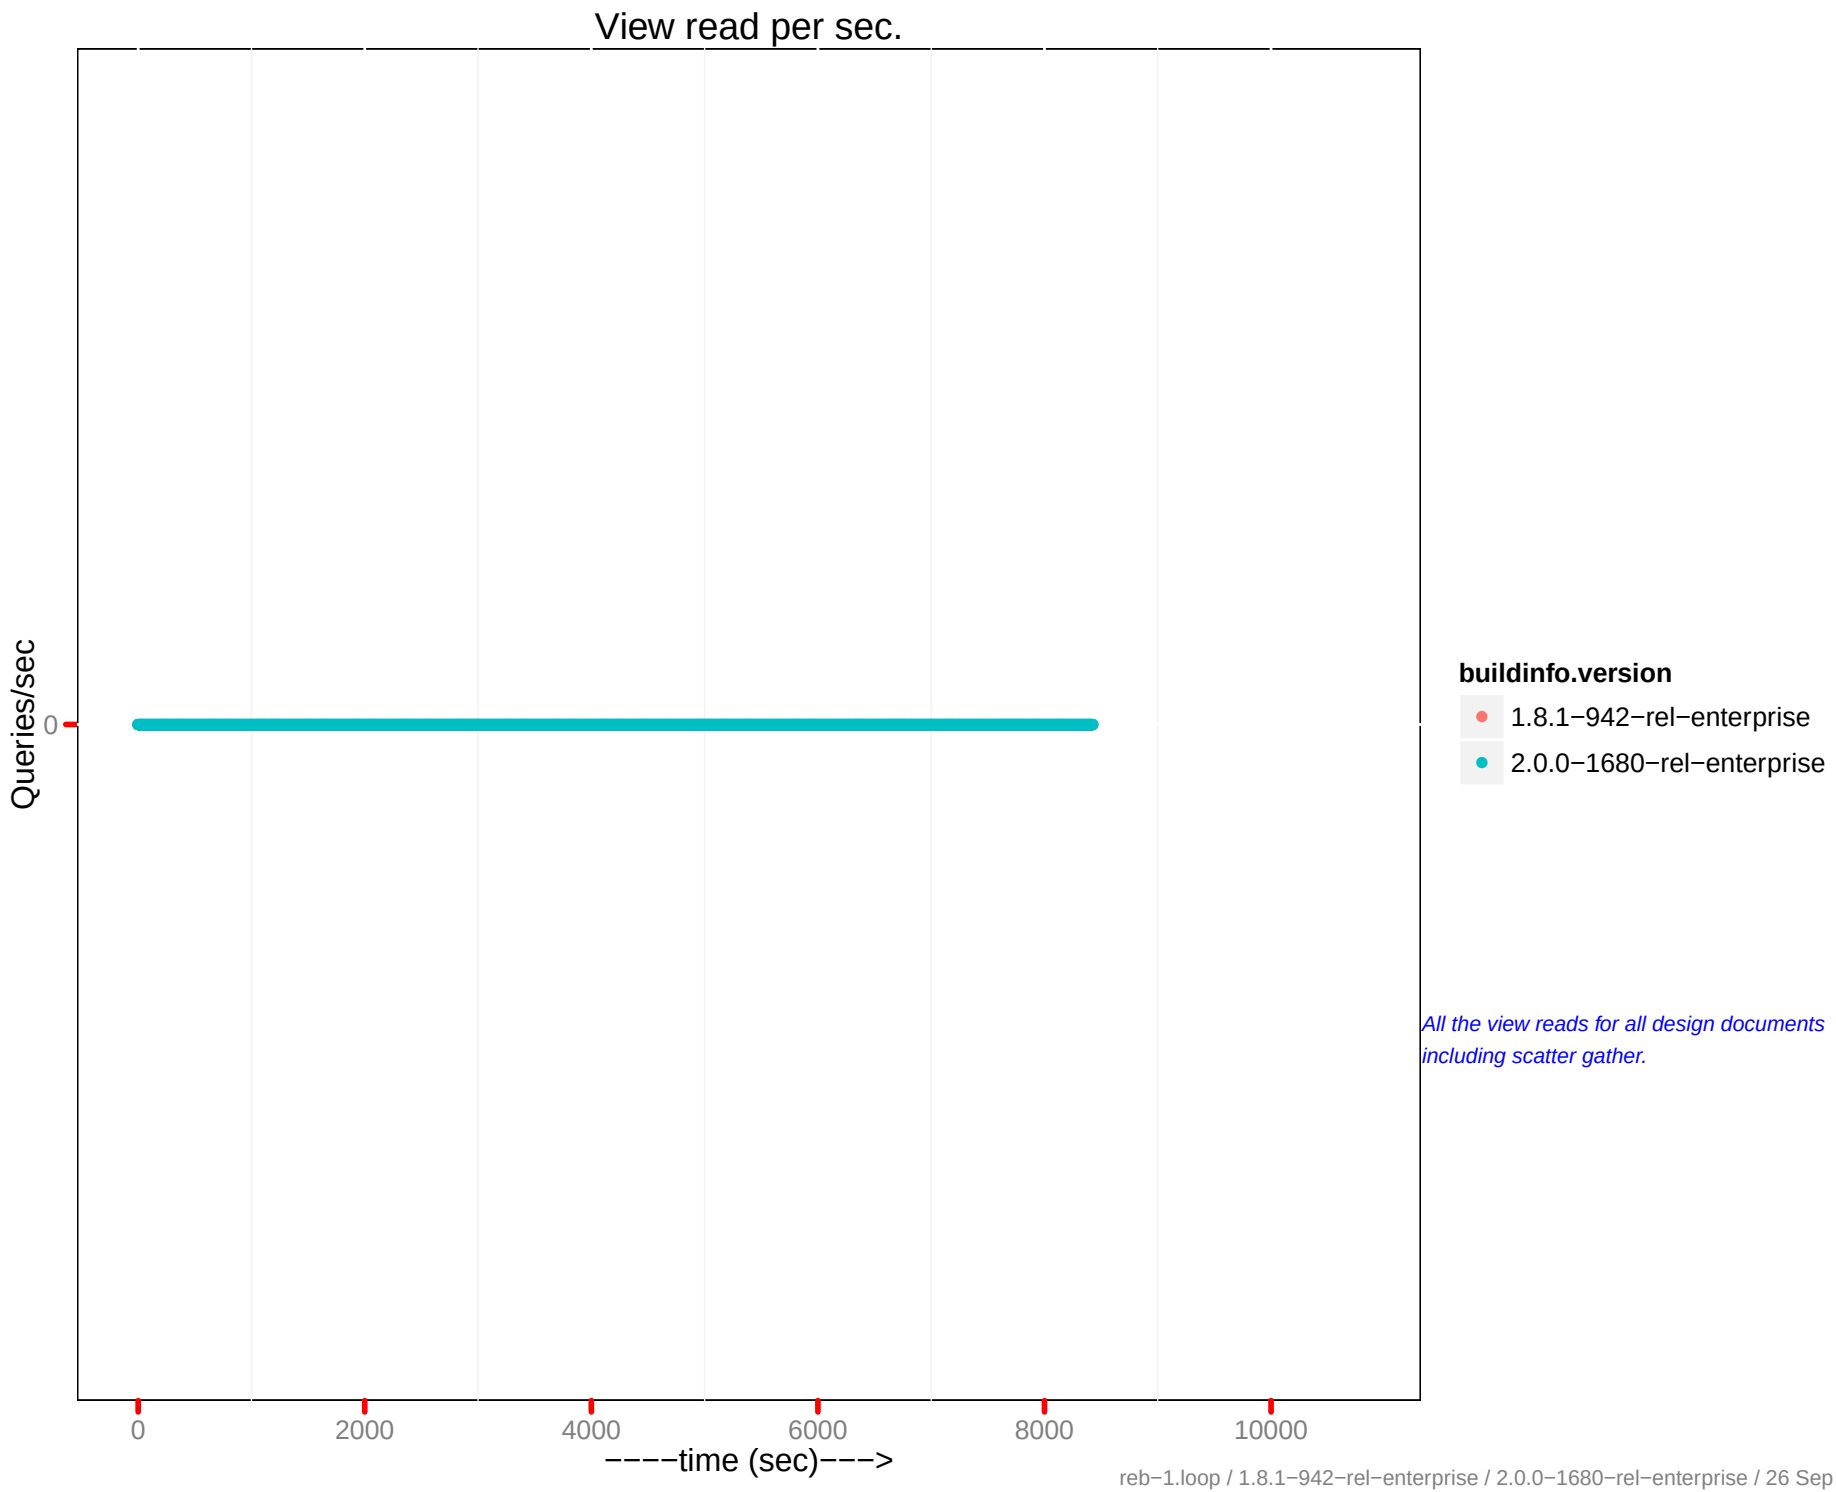
View read per sec.
Queries/sec
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
All the view reads for all design documents including scatter gather.
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

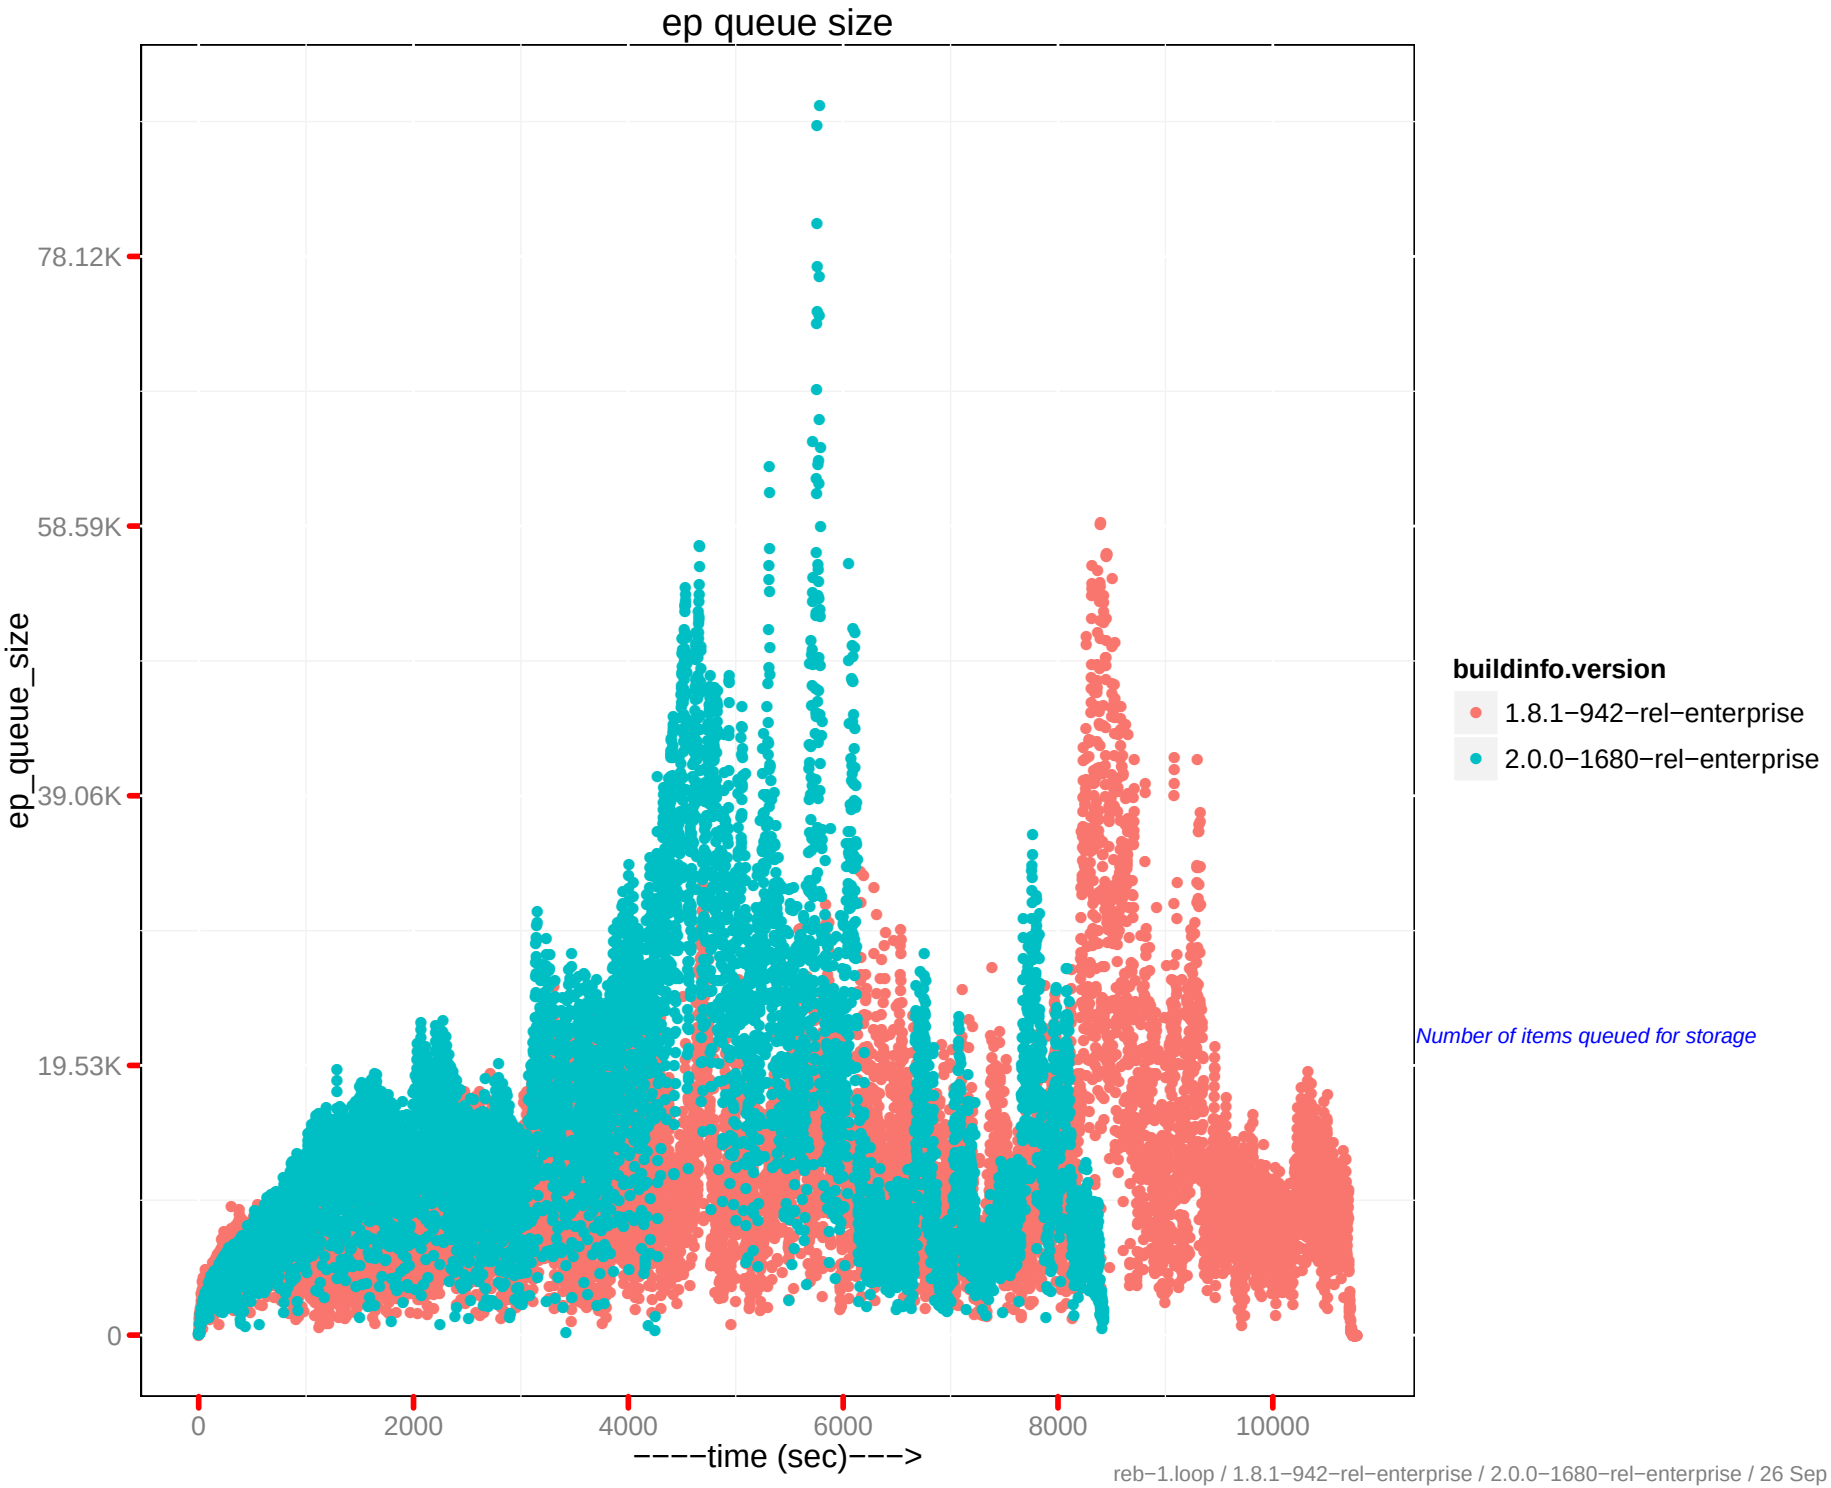
ep queue size
78.12K
58.59K
39.06K
19.53K
0
ep_queue_size
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of items queued for storage
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

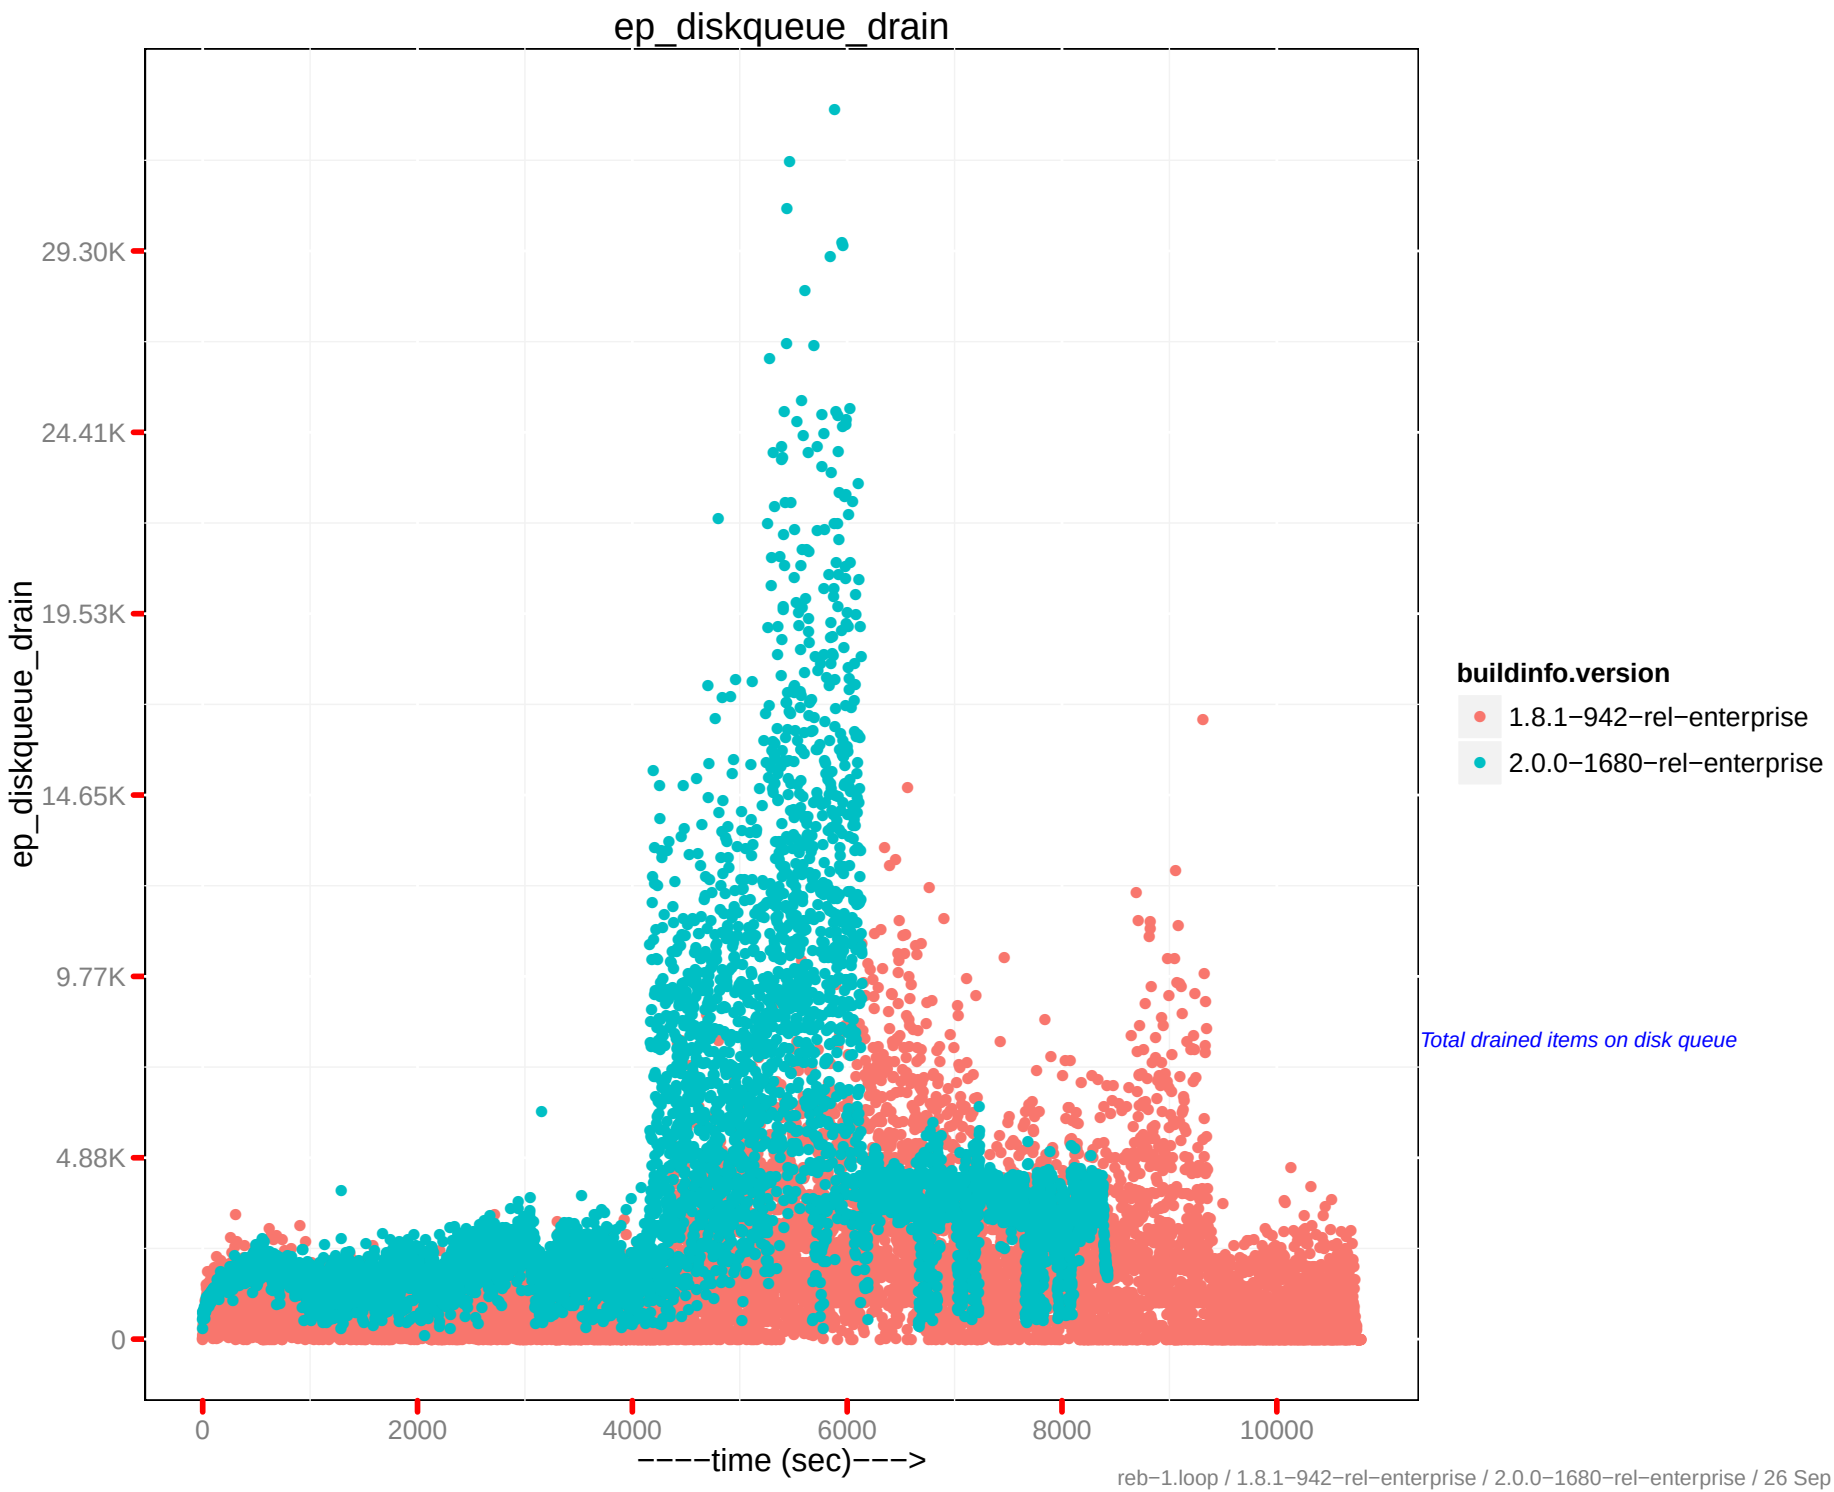
ep_diskqueue_drain
29.30K
24.41K
19.53K
14.65K
9.77K
4.88K
0
ep_diskqueue_drain
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Total drained items on disk queue
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

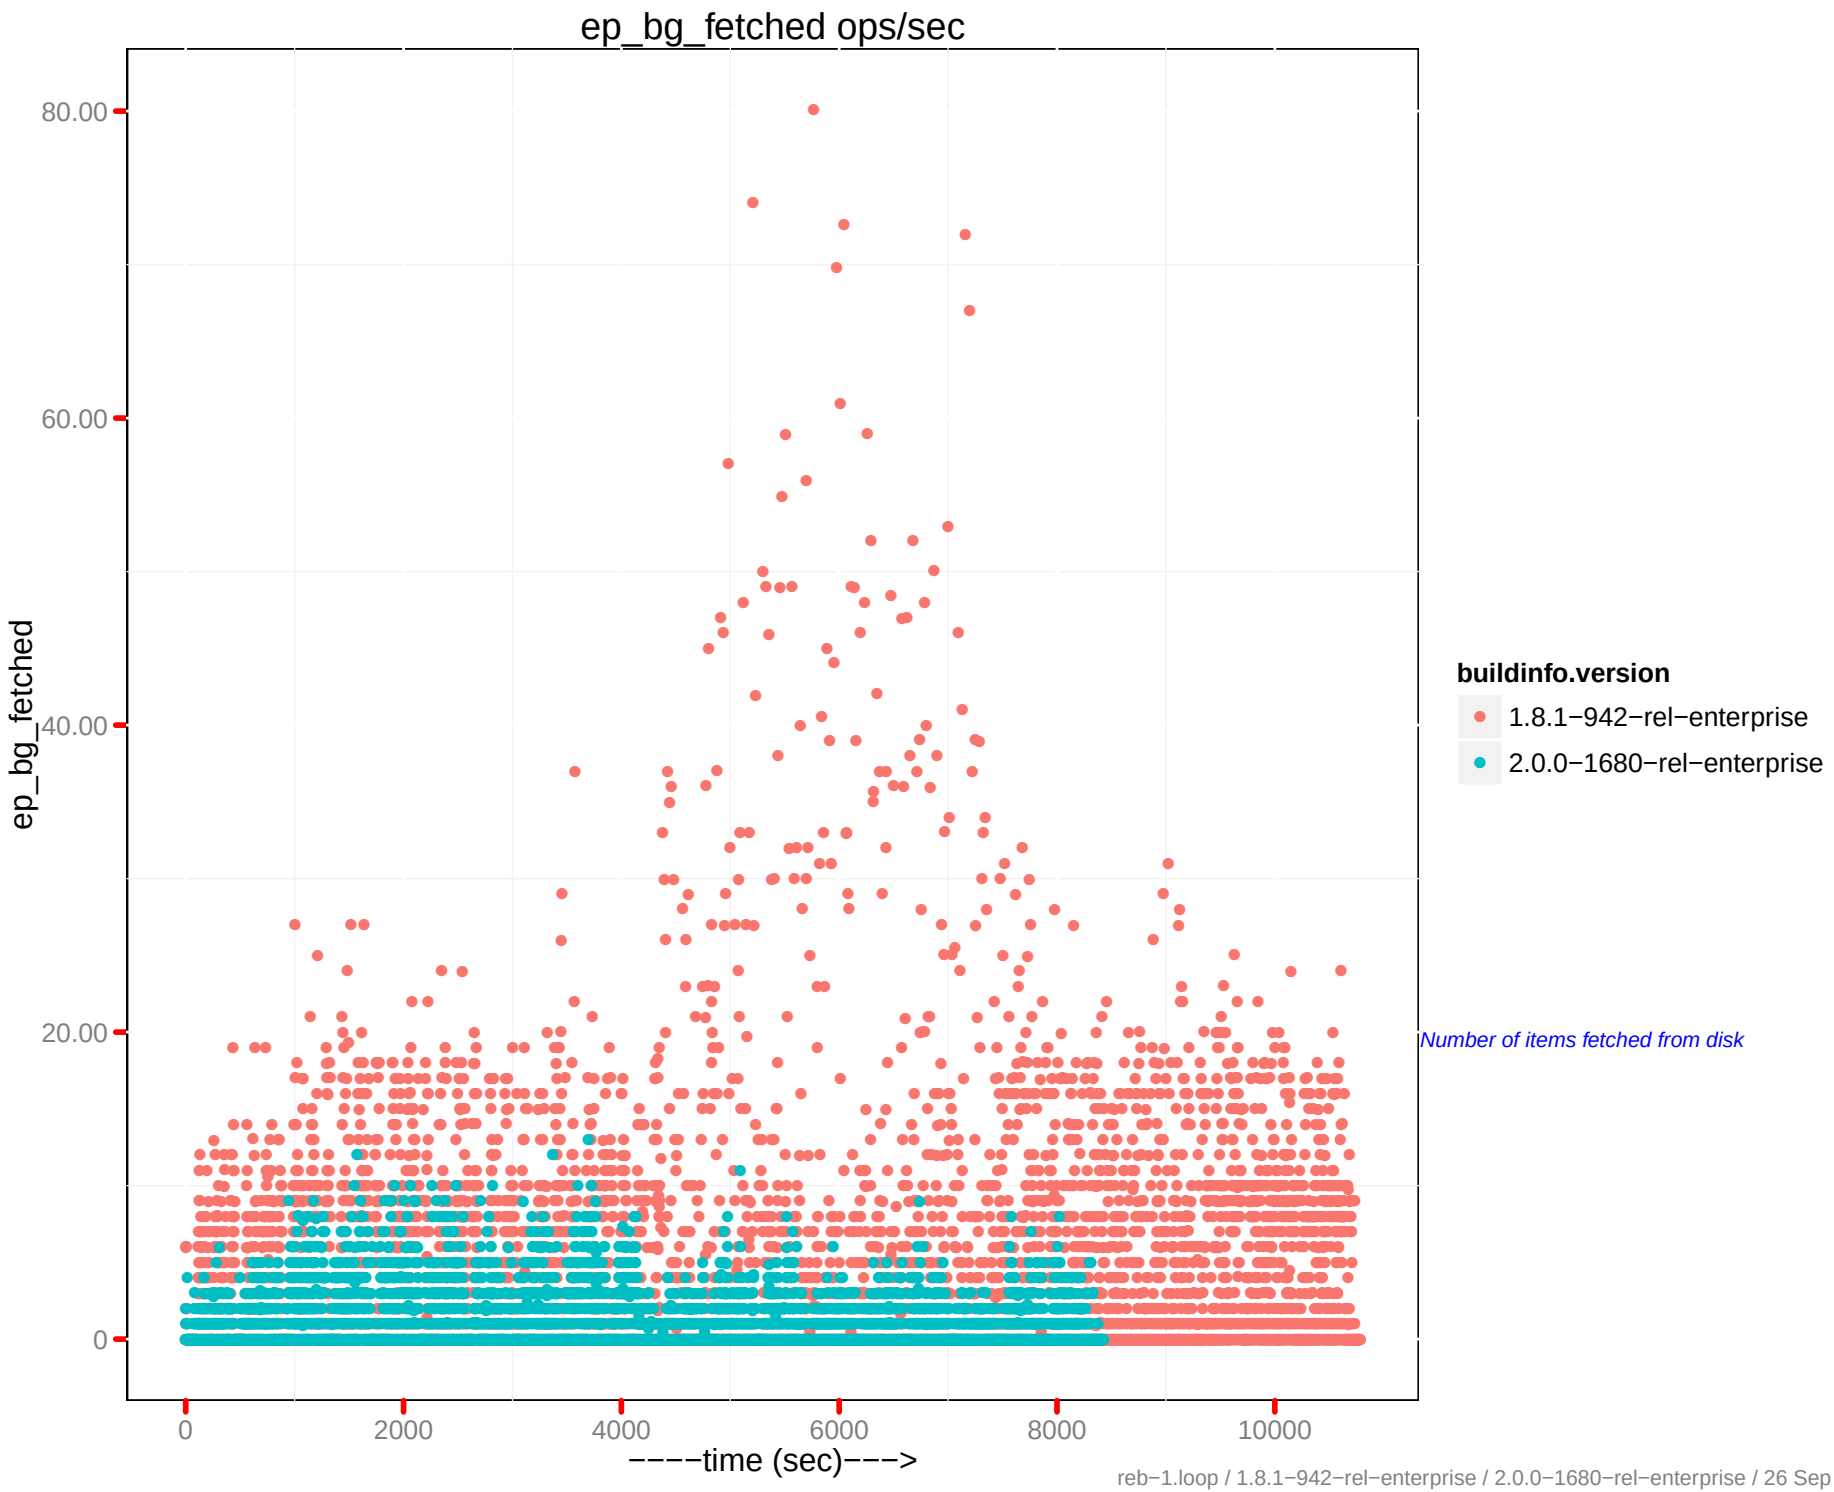
ep_bg_fetched ops/sec
80.00
60.00
40.00
20.00
0
ep_bg_fetched
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of items fetched from disk
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

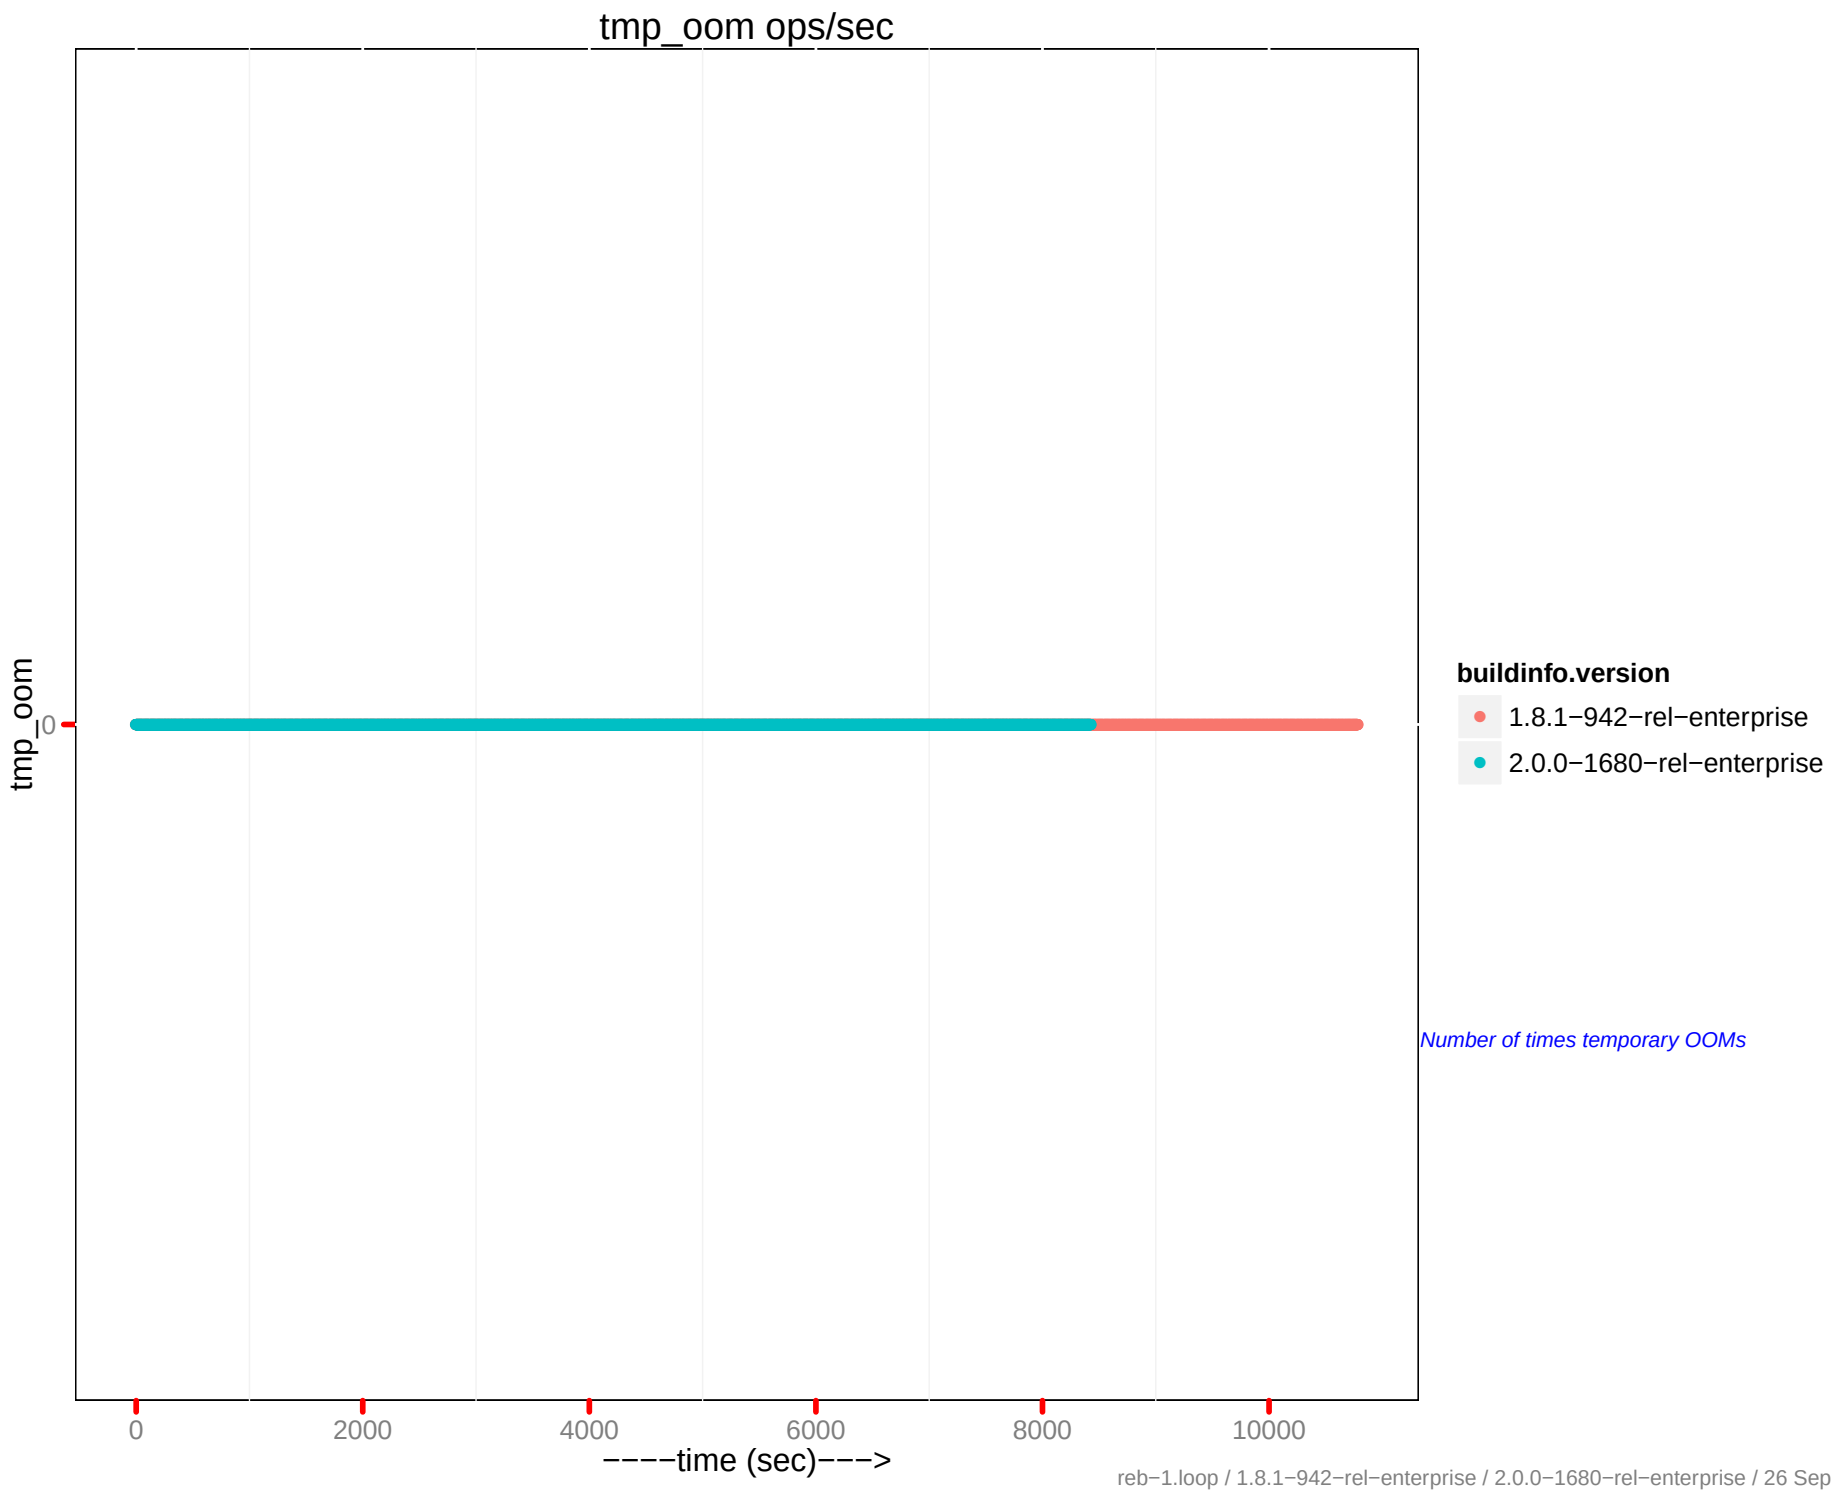
tmp_oom ops/sec
tmp_oom
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of times temporary OOMs
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

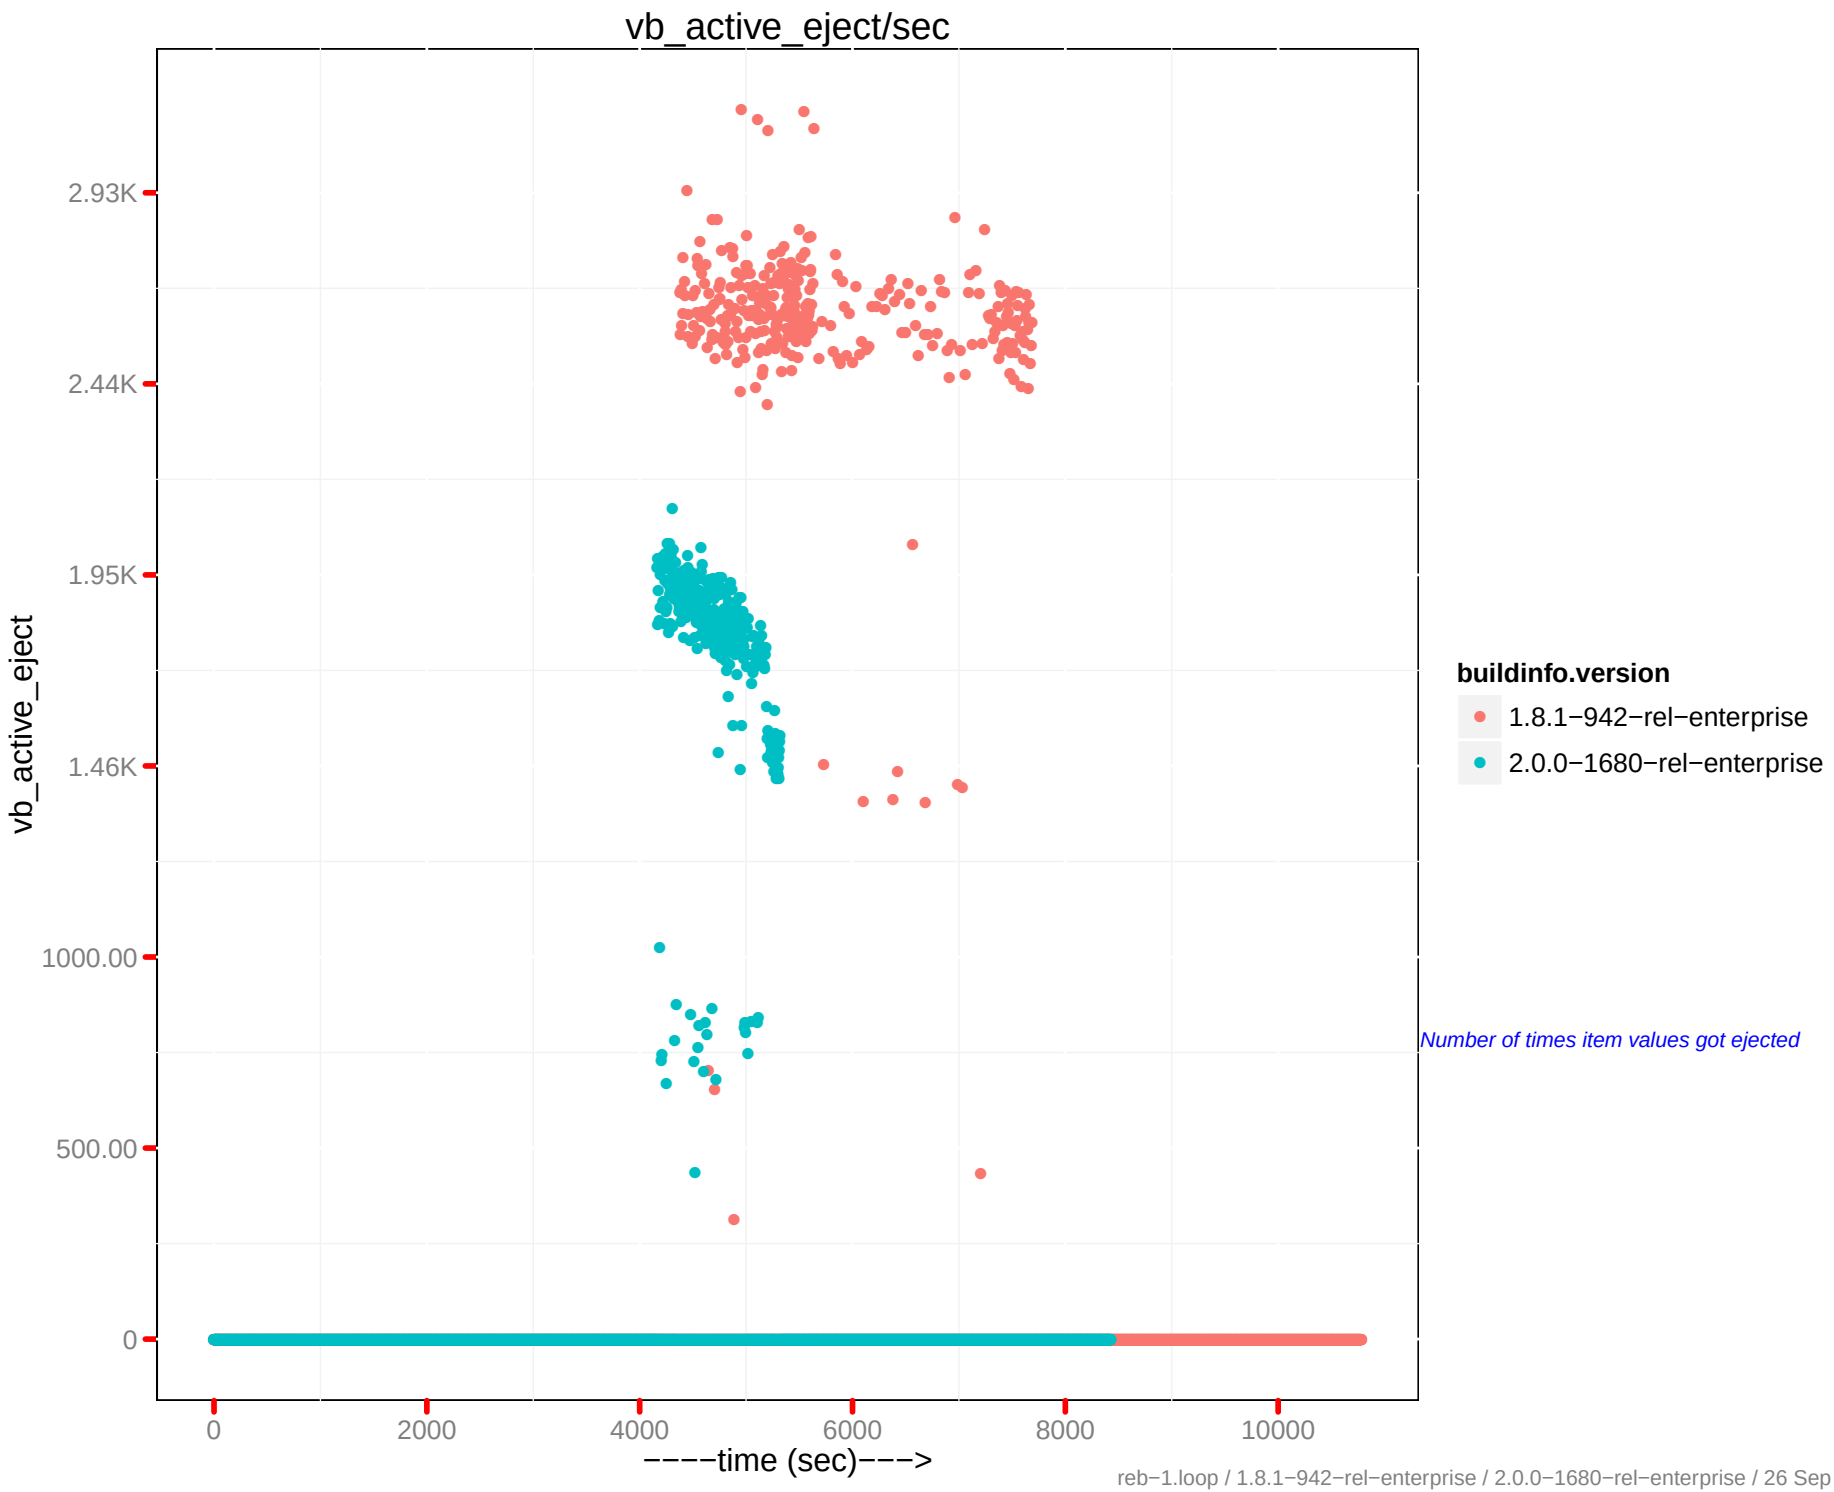
vb_active_eject/sec
vb_active_eject
2.93K
2.44K
1.95K
1.46K
1000.00
500.00
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of times item values got ejected
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

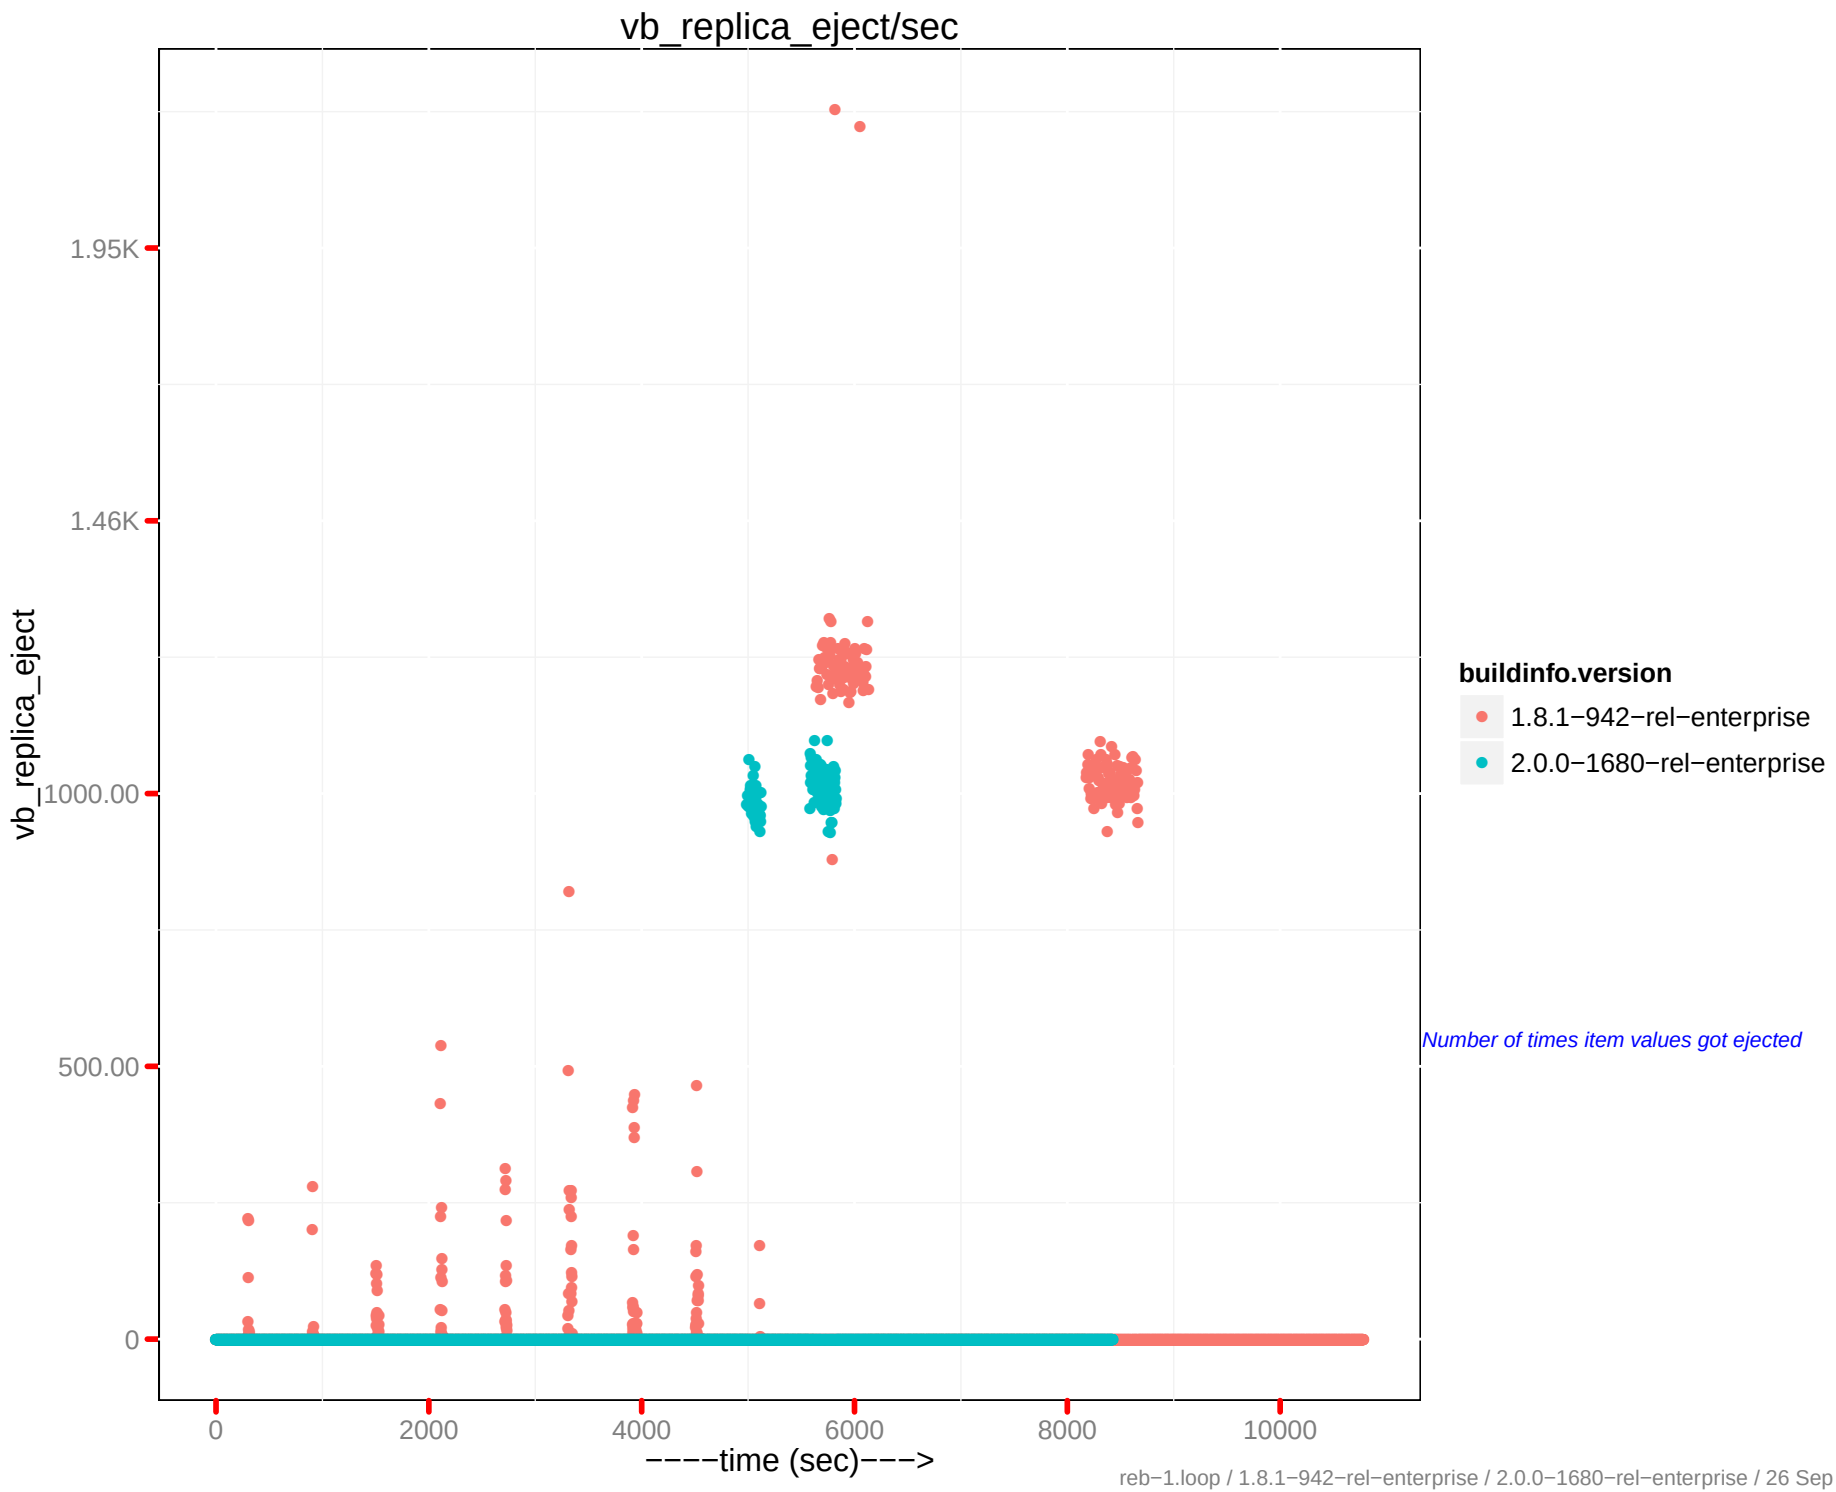
vb_replica_eject/sec
vb_replica_eject
1.95K
1.46K
1000.00
500.00
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of times item values got ejected
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

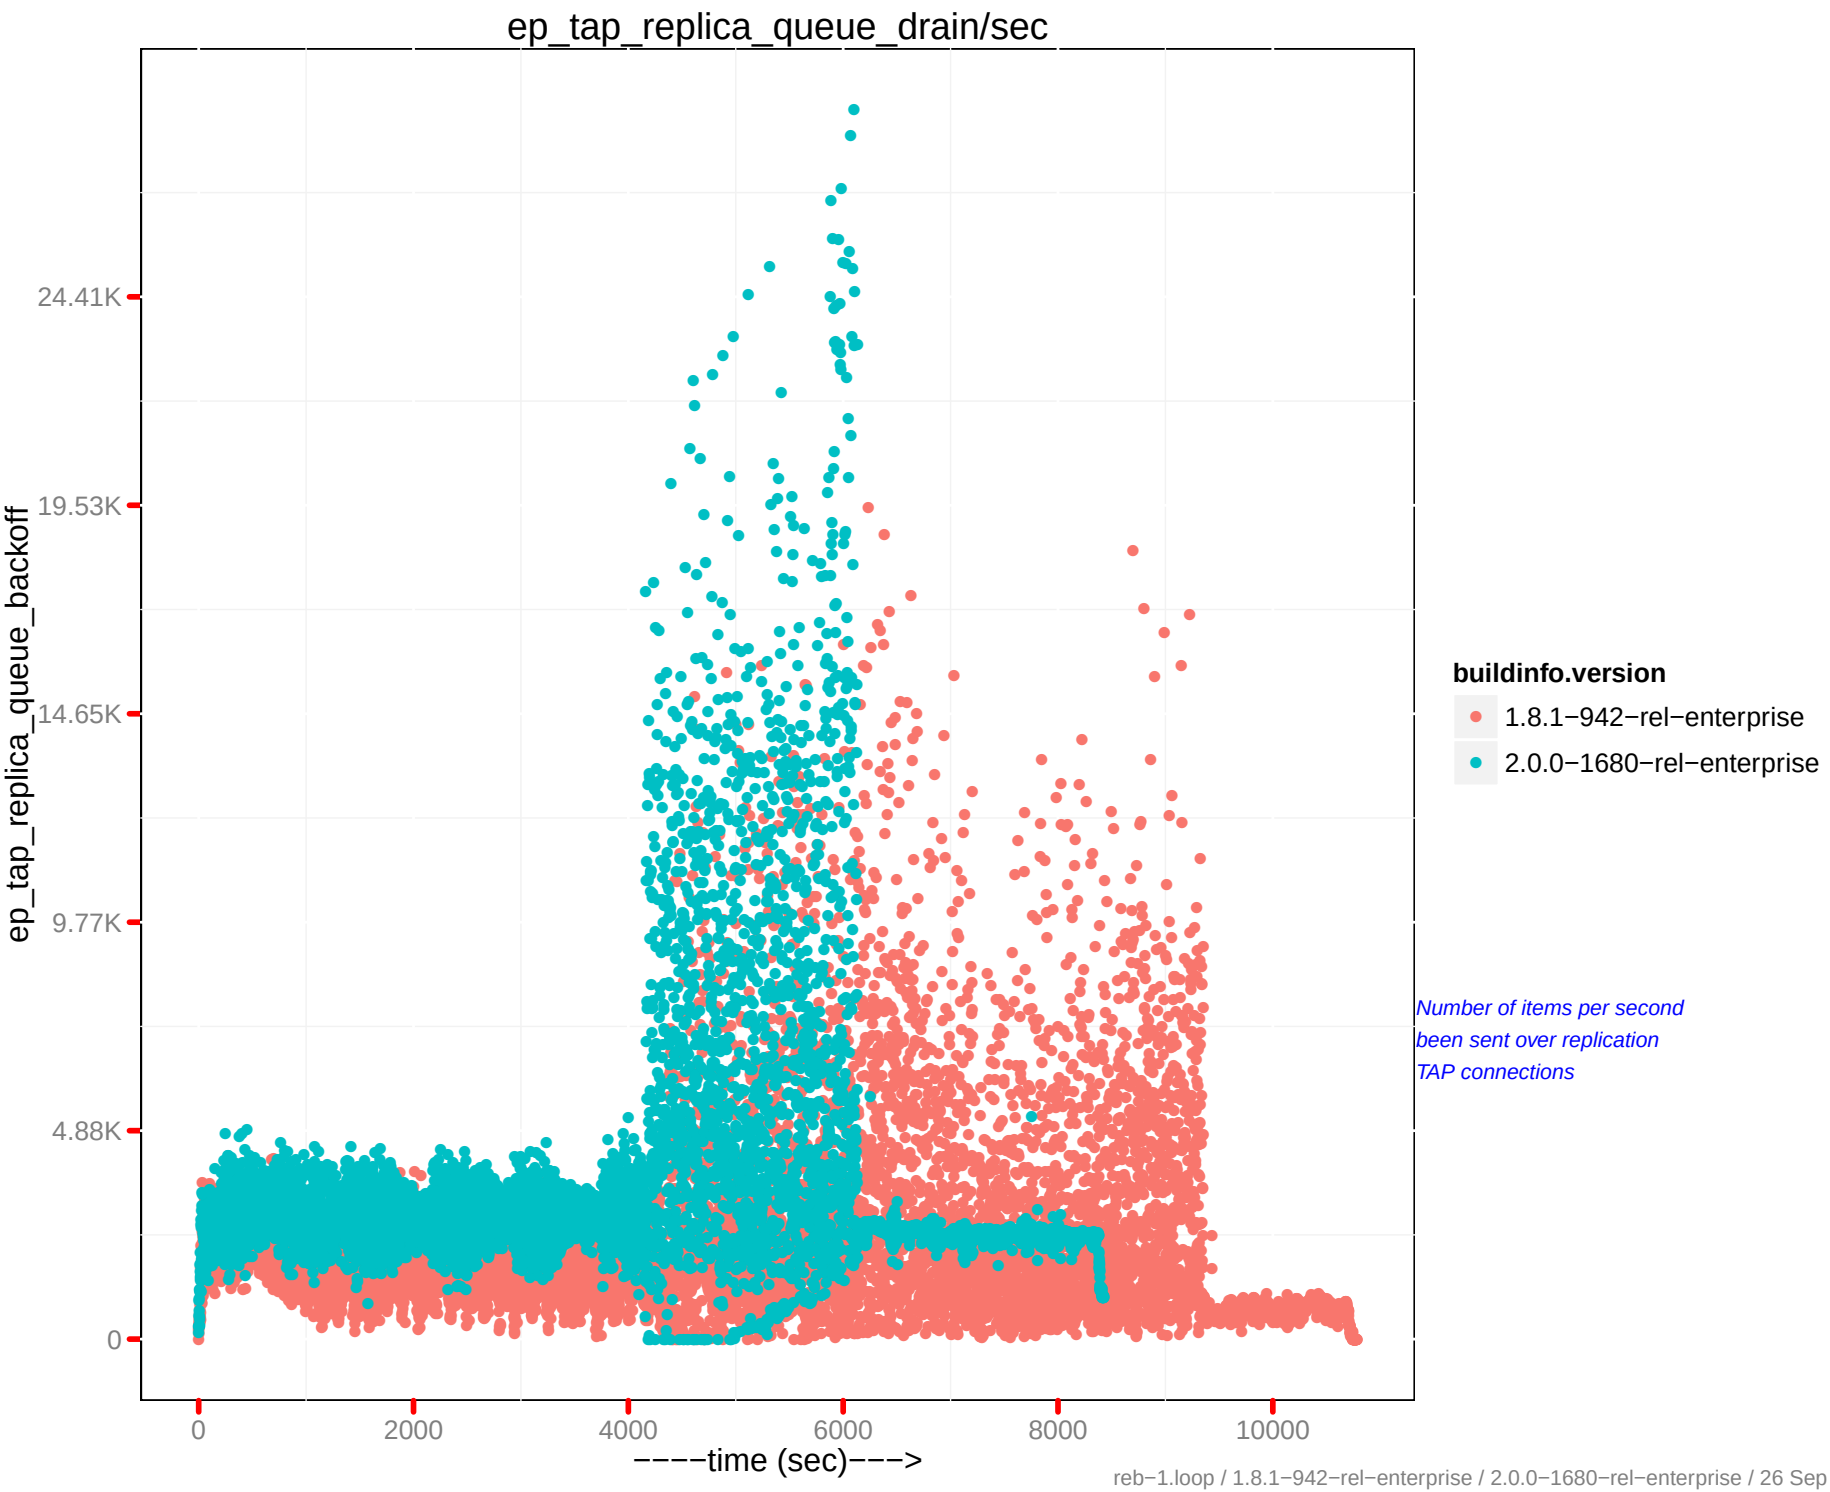
ep_tap_replica_queue_drain/sec
ep_tap_replica_queue_backoff
24.41K
19.53K
14.65K
9.77K
4.88K
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of items per second
been sent over replication
TAP connections
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

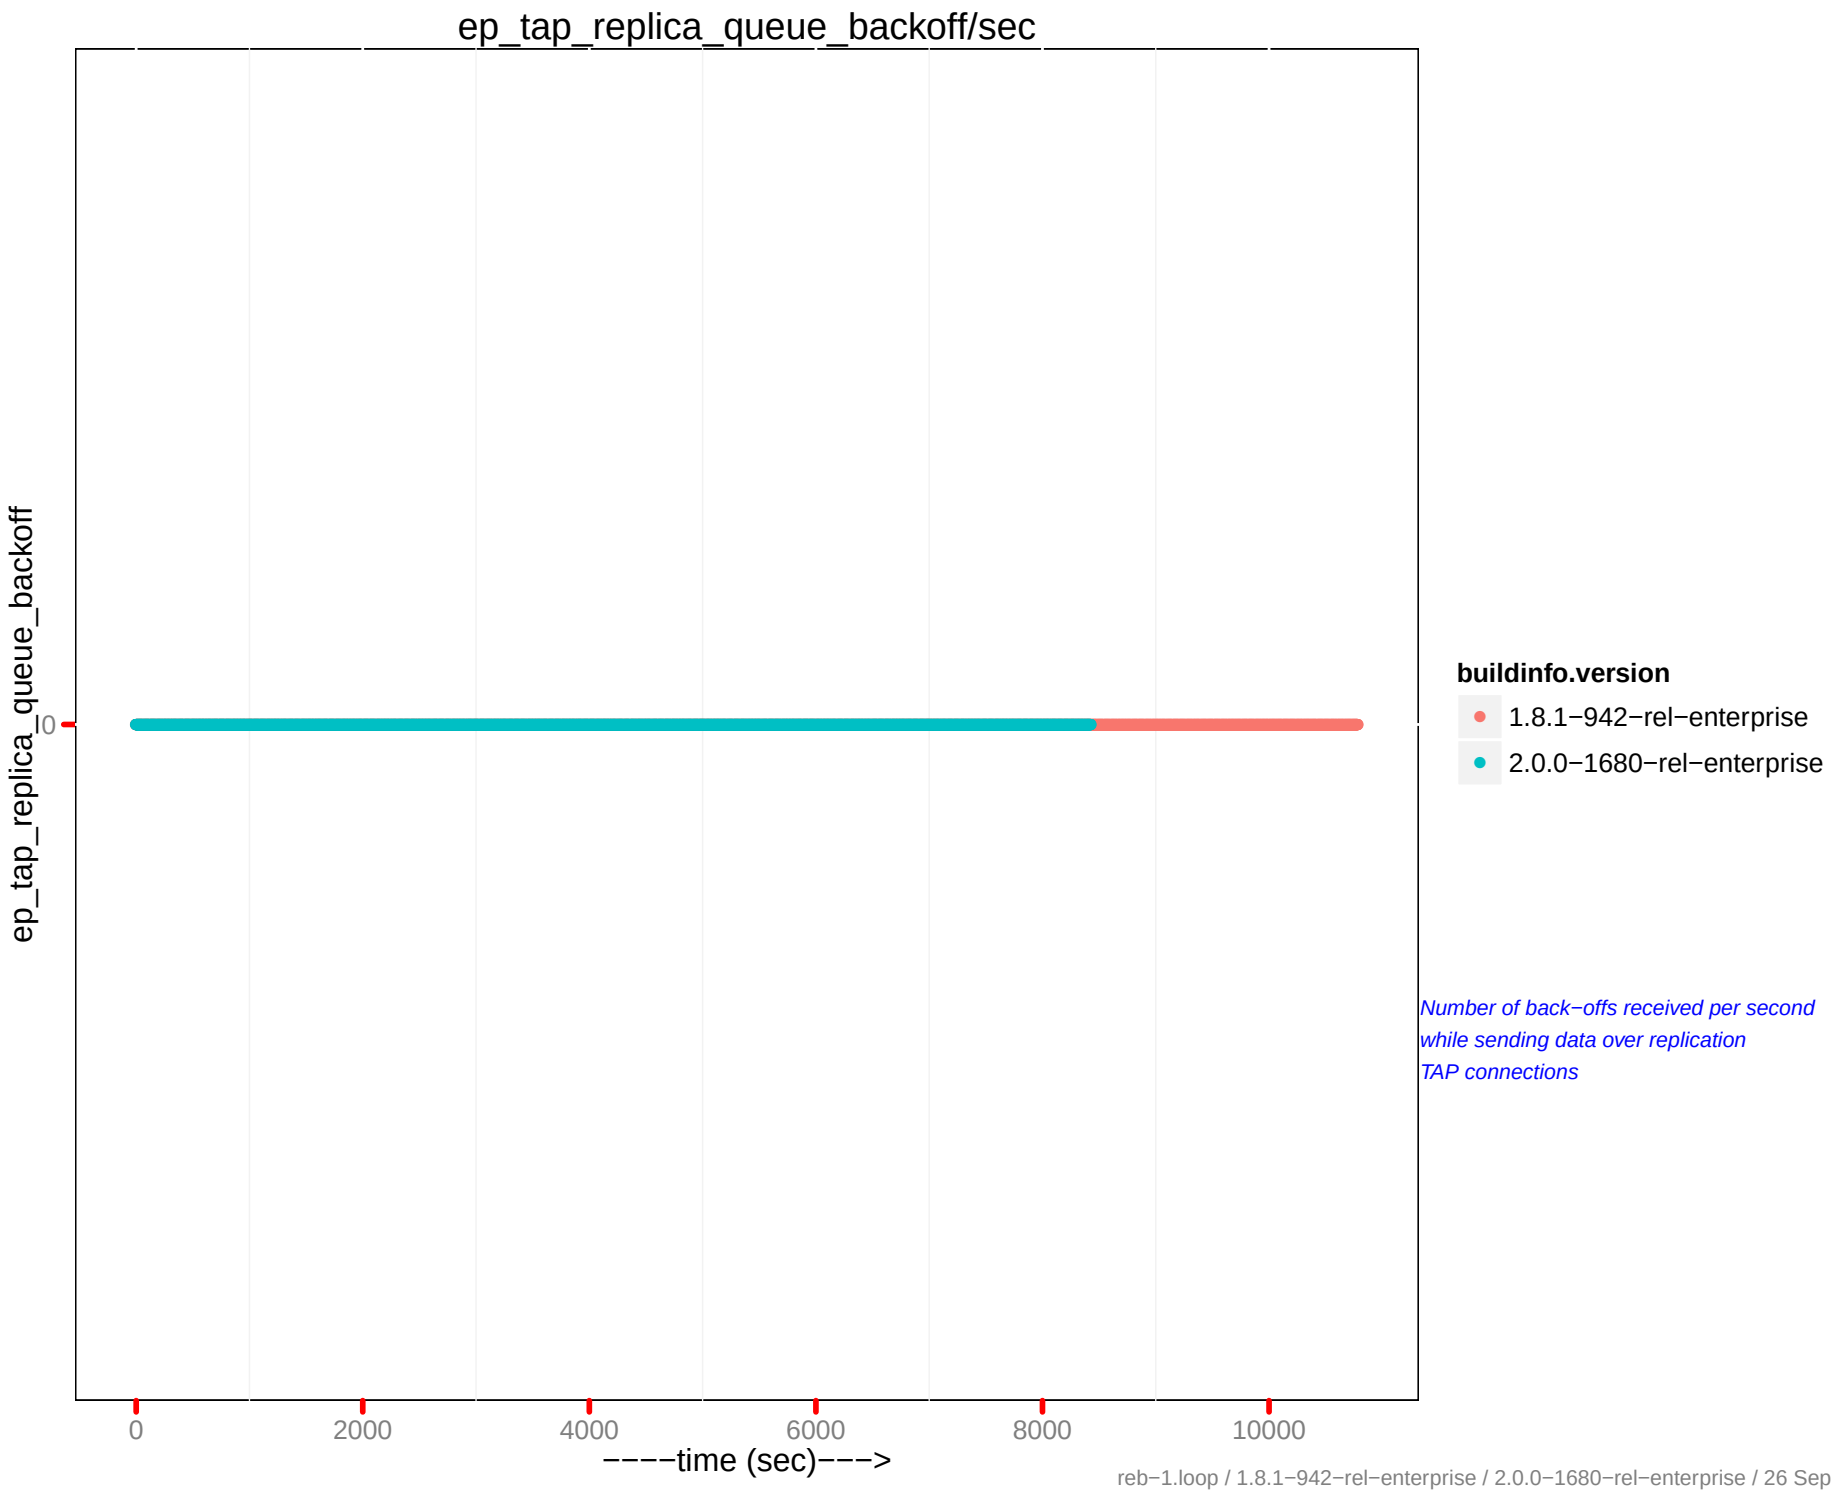
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of back−offs received per second
while sending data over replication
TAP connections
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

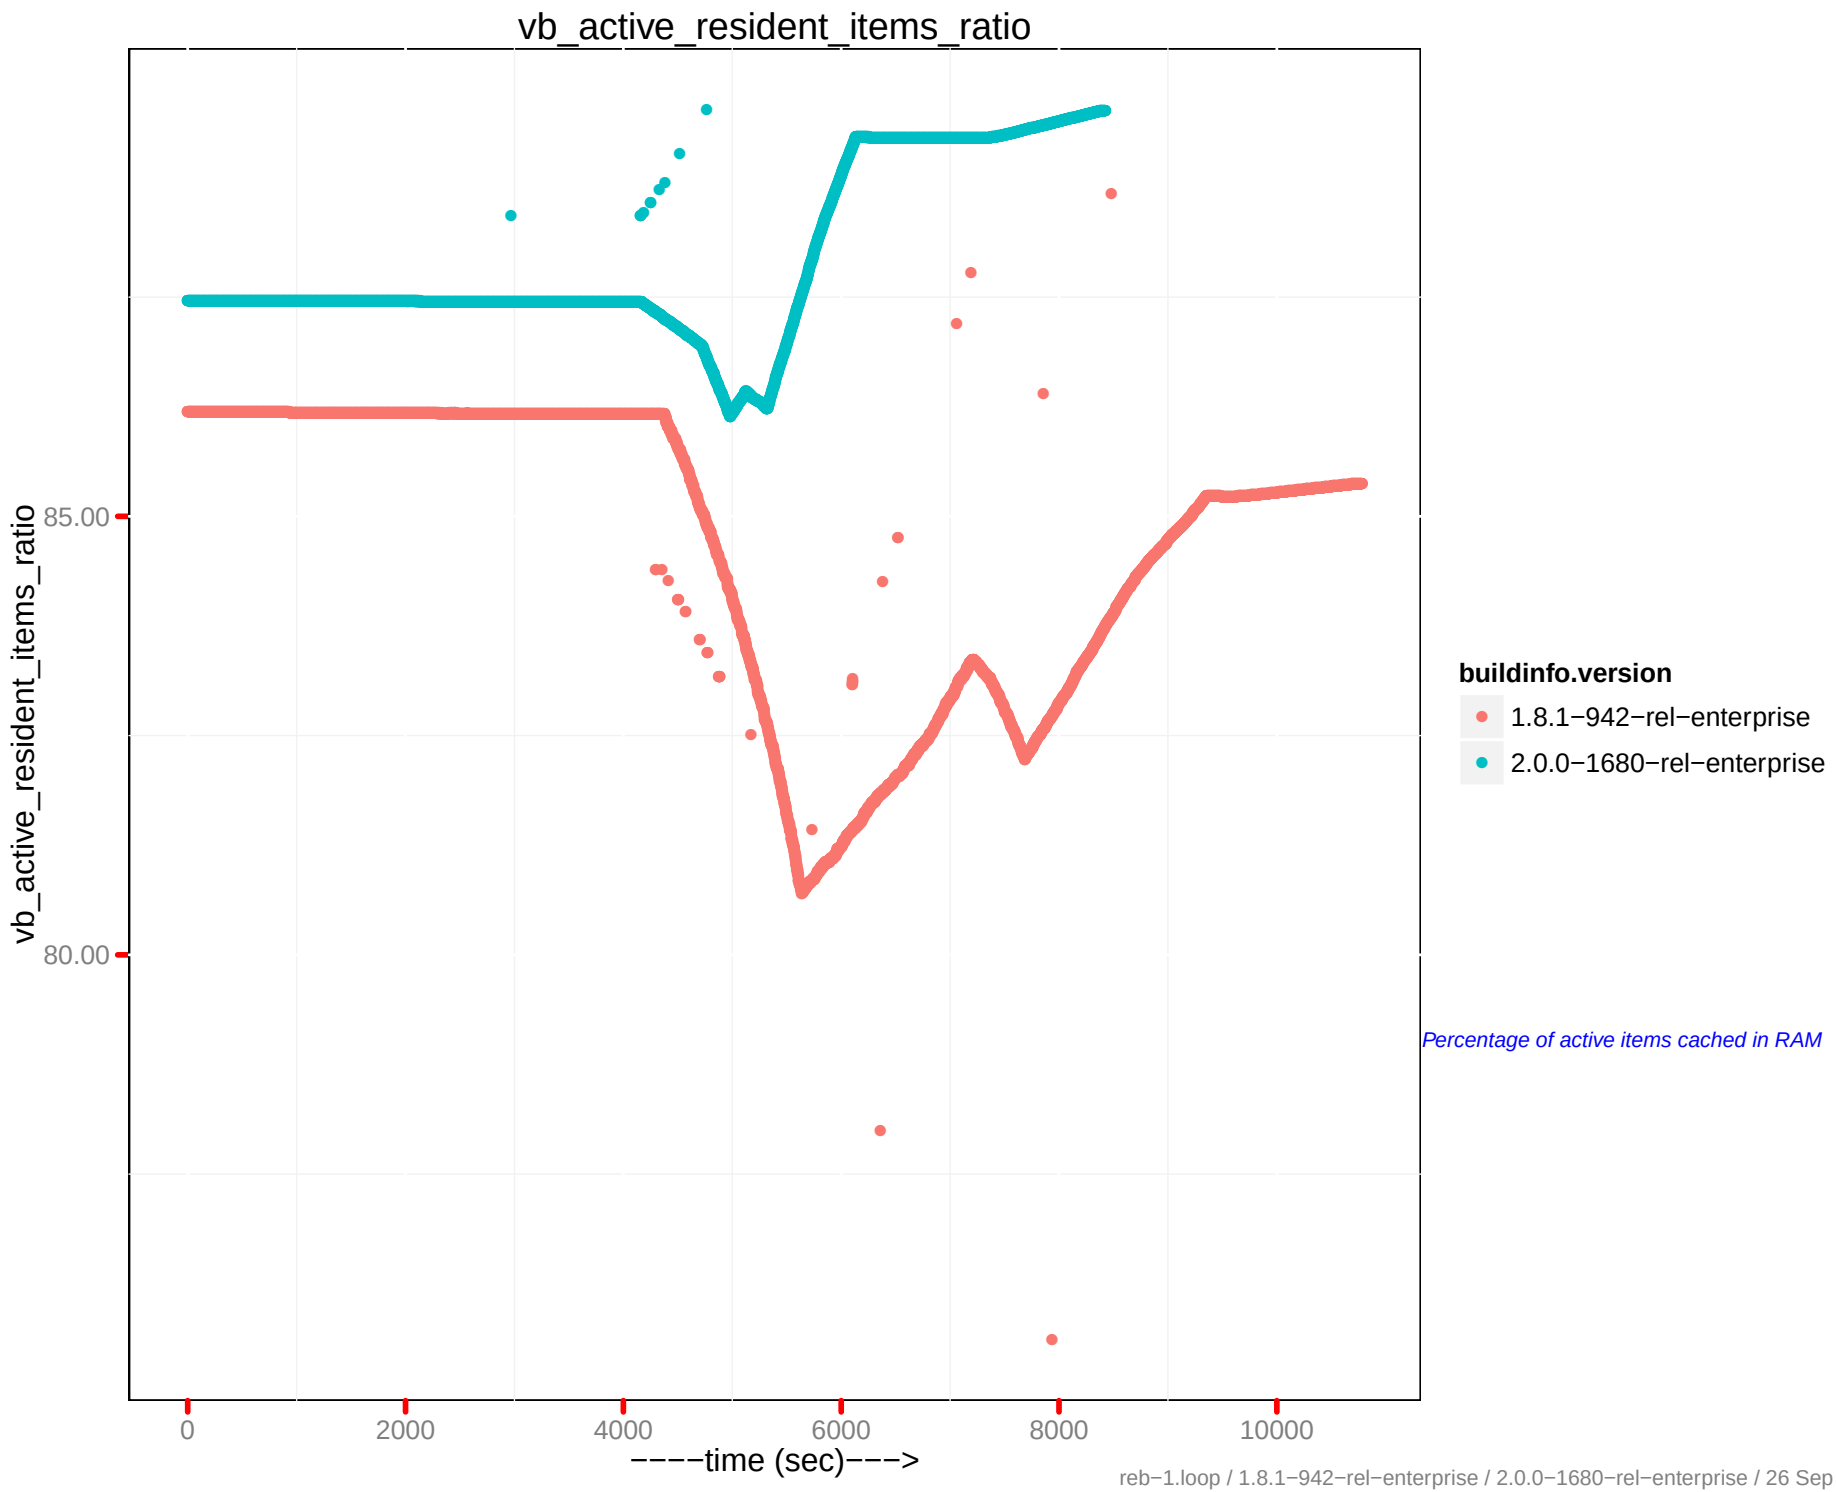
vb_active_resident_items_ratio
vb_active_resident_items_ratio
85.00
80.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Percentage of active items cached in RAM
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

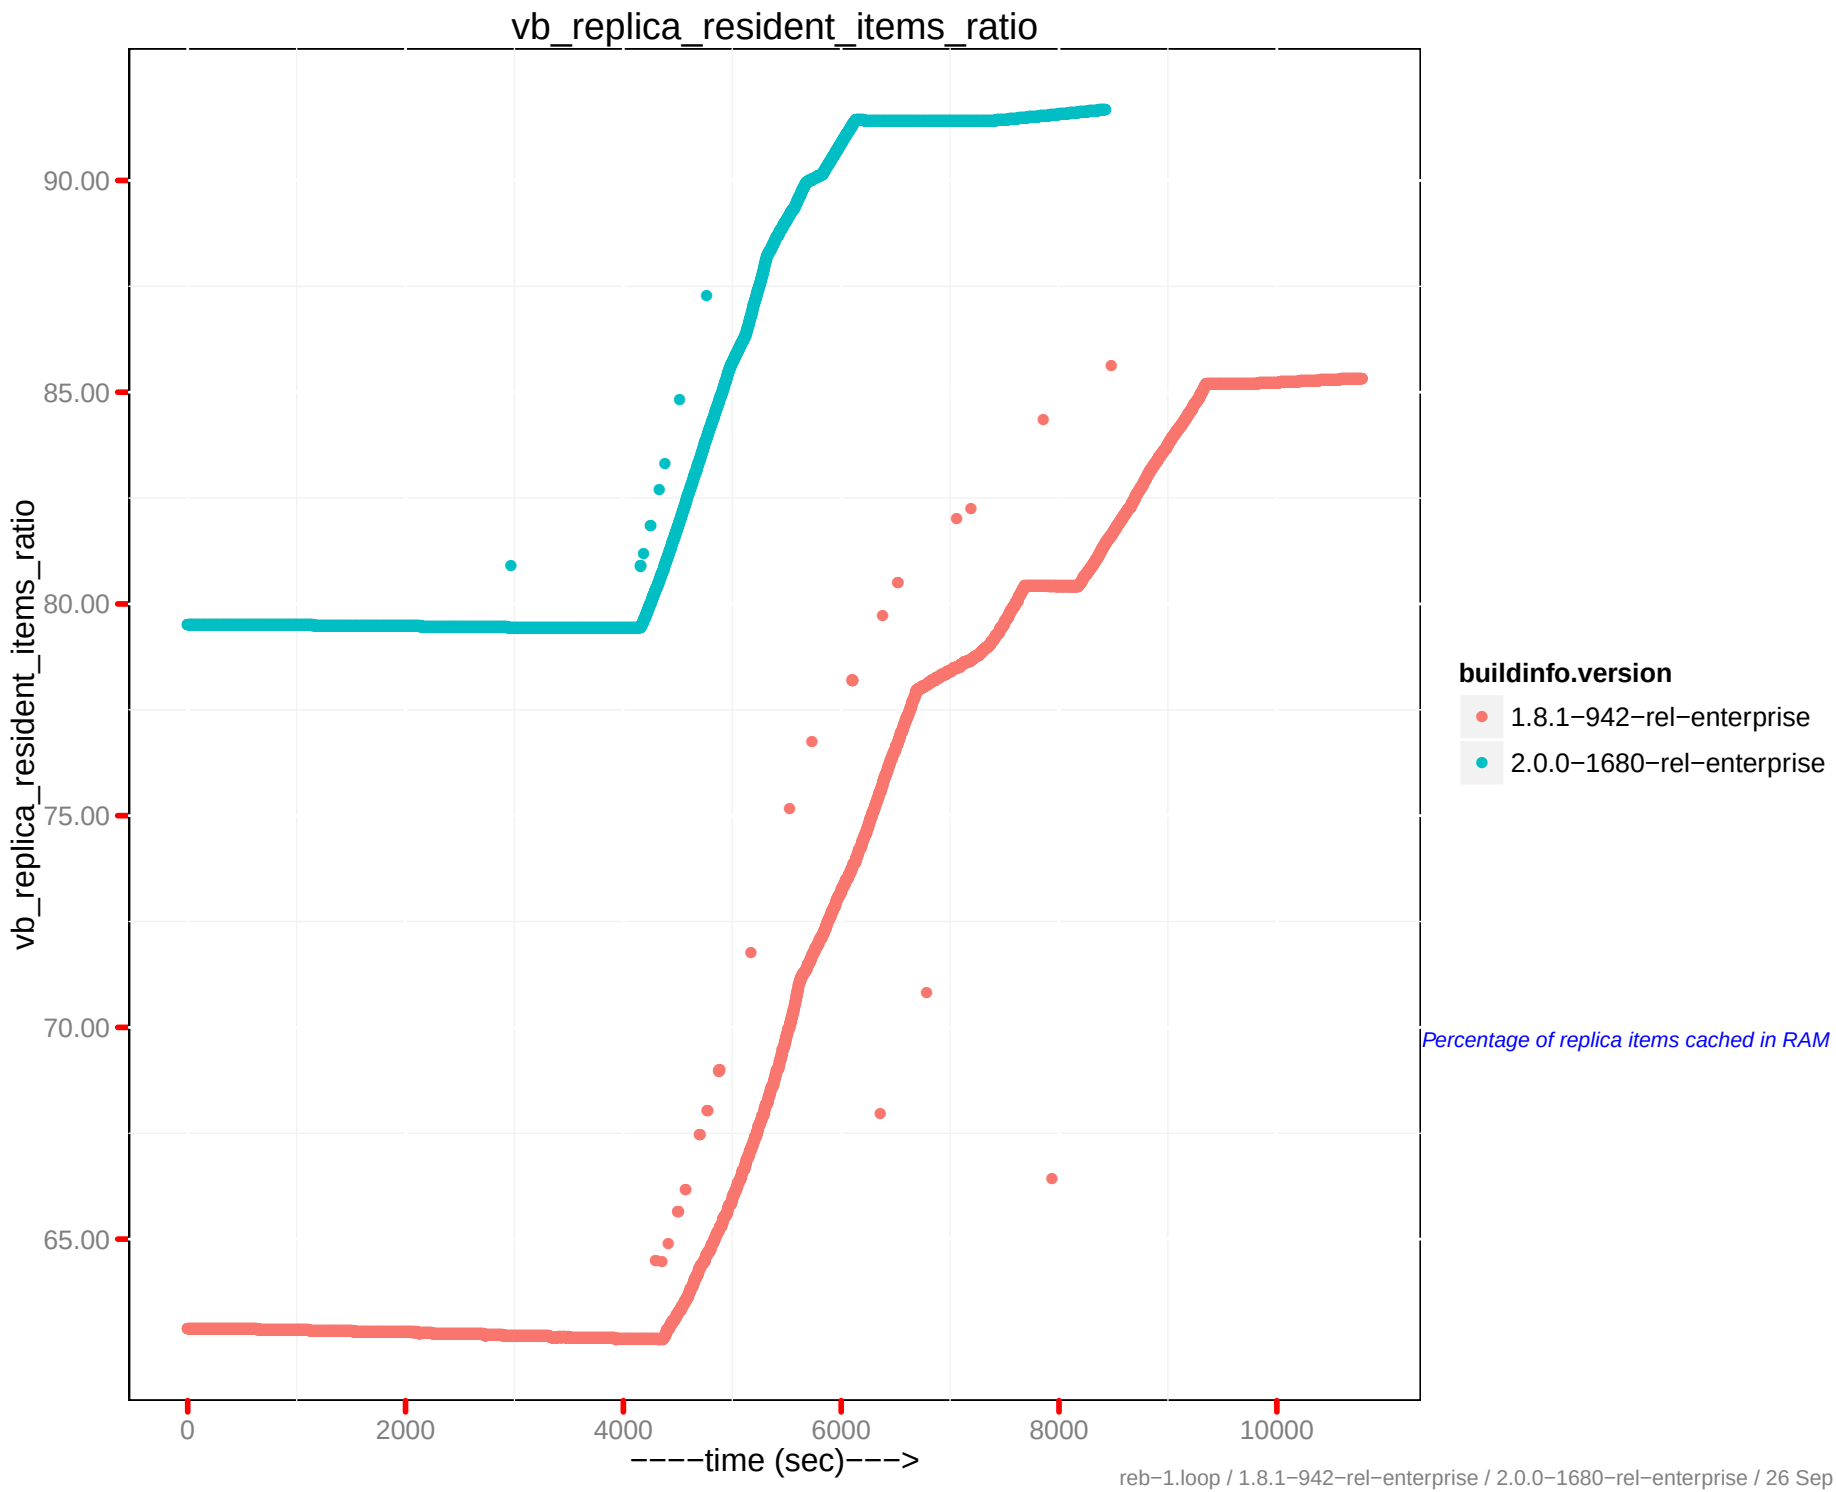
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
90.00
85.00
80.00
75.00
70.00
65.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Percentage of replica items cached in RAM
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

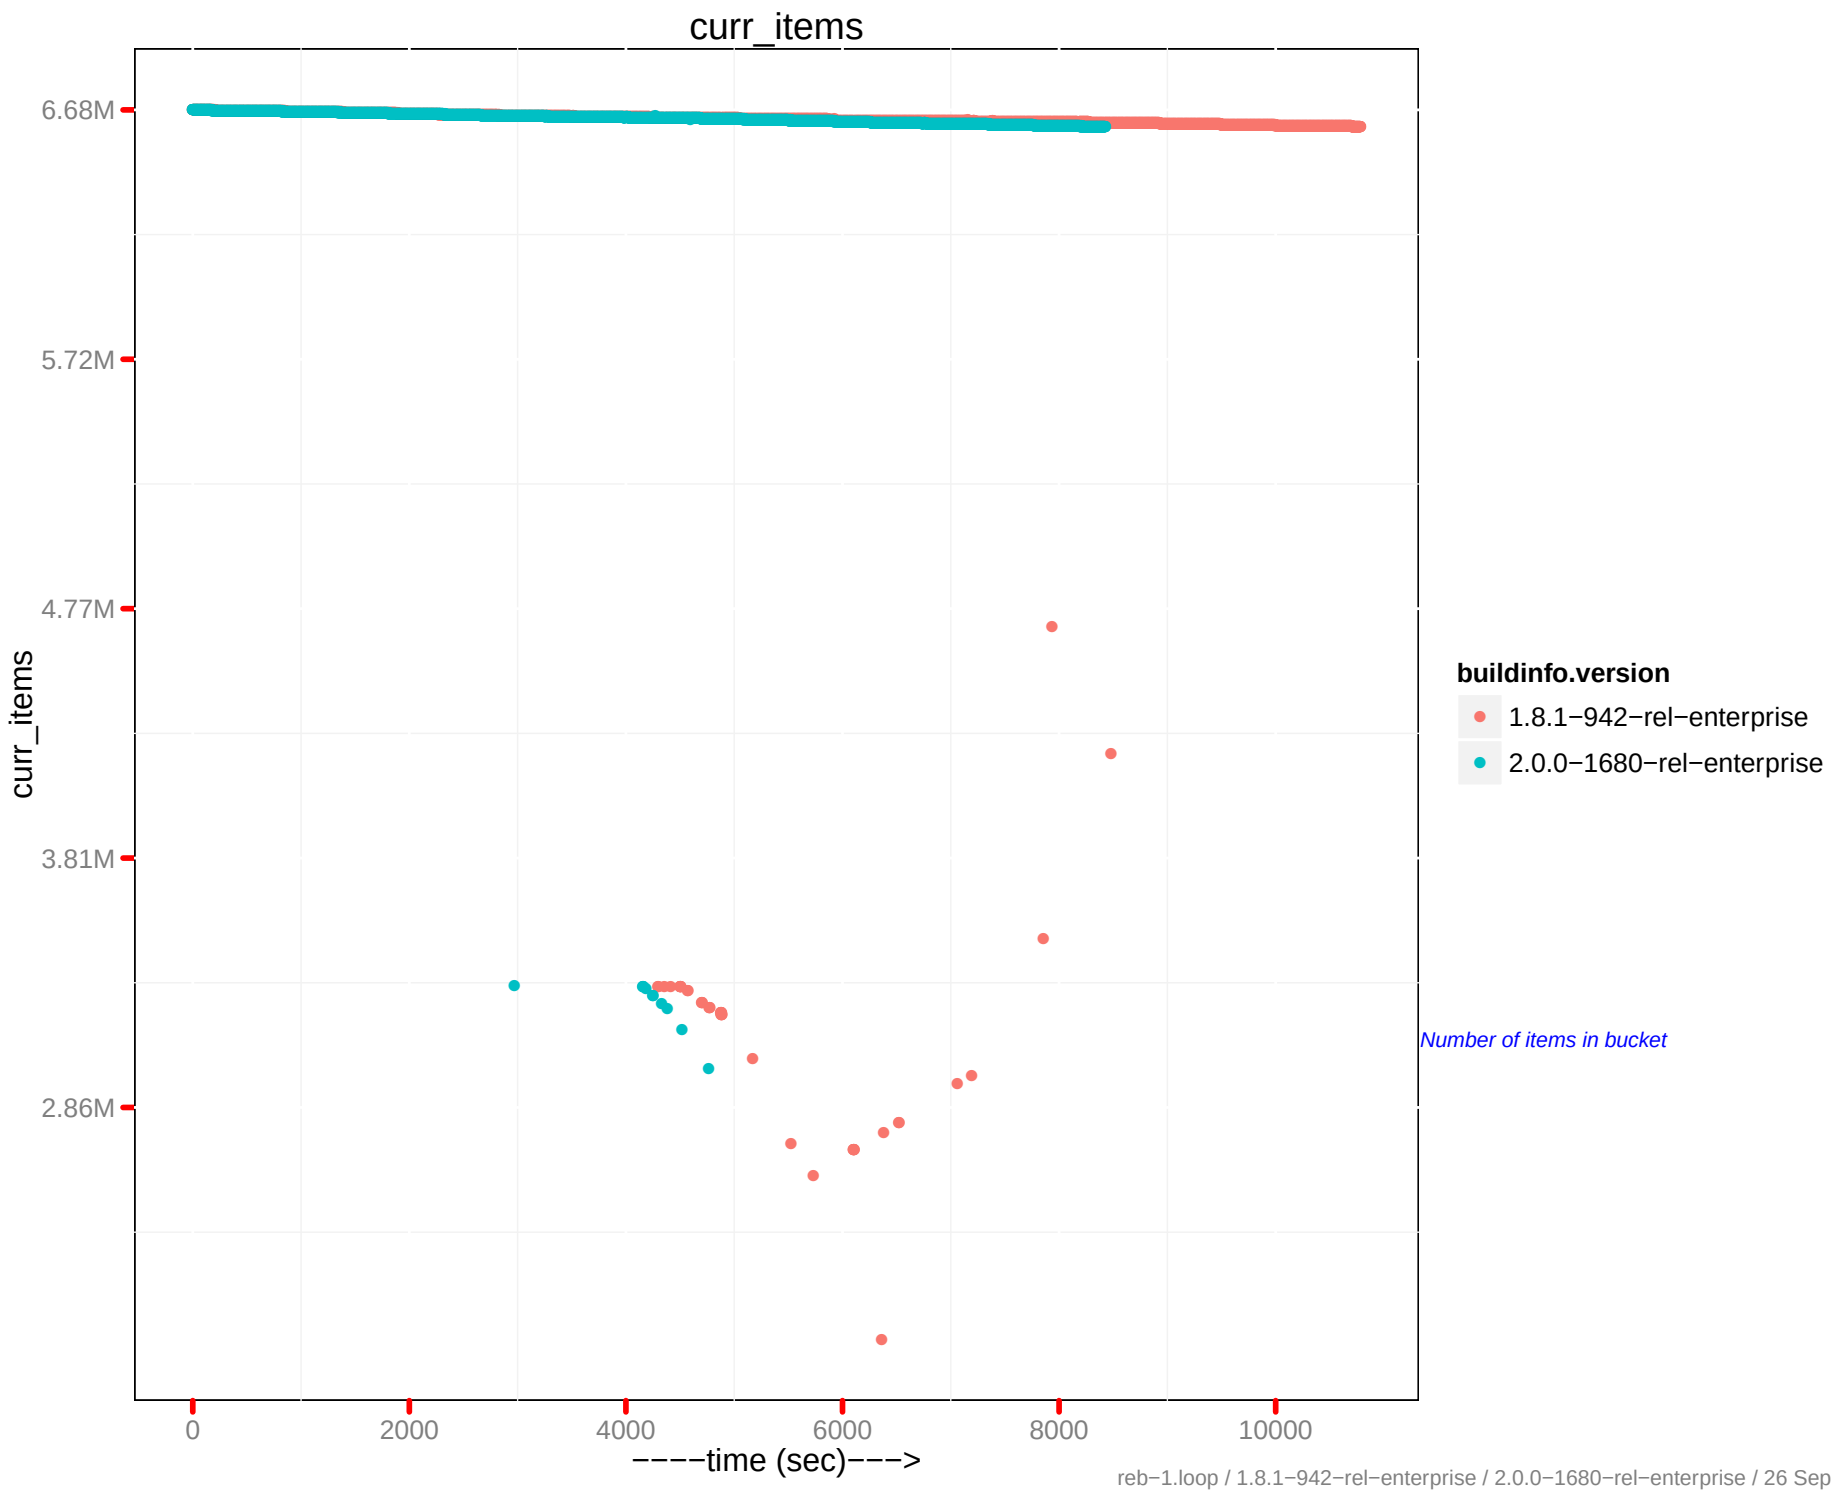
curr_items
6.68M
5.72M
4.77M
3.81M
2.86M
curr_items
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of items in bucket
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

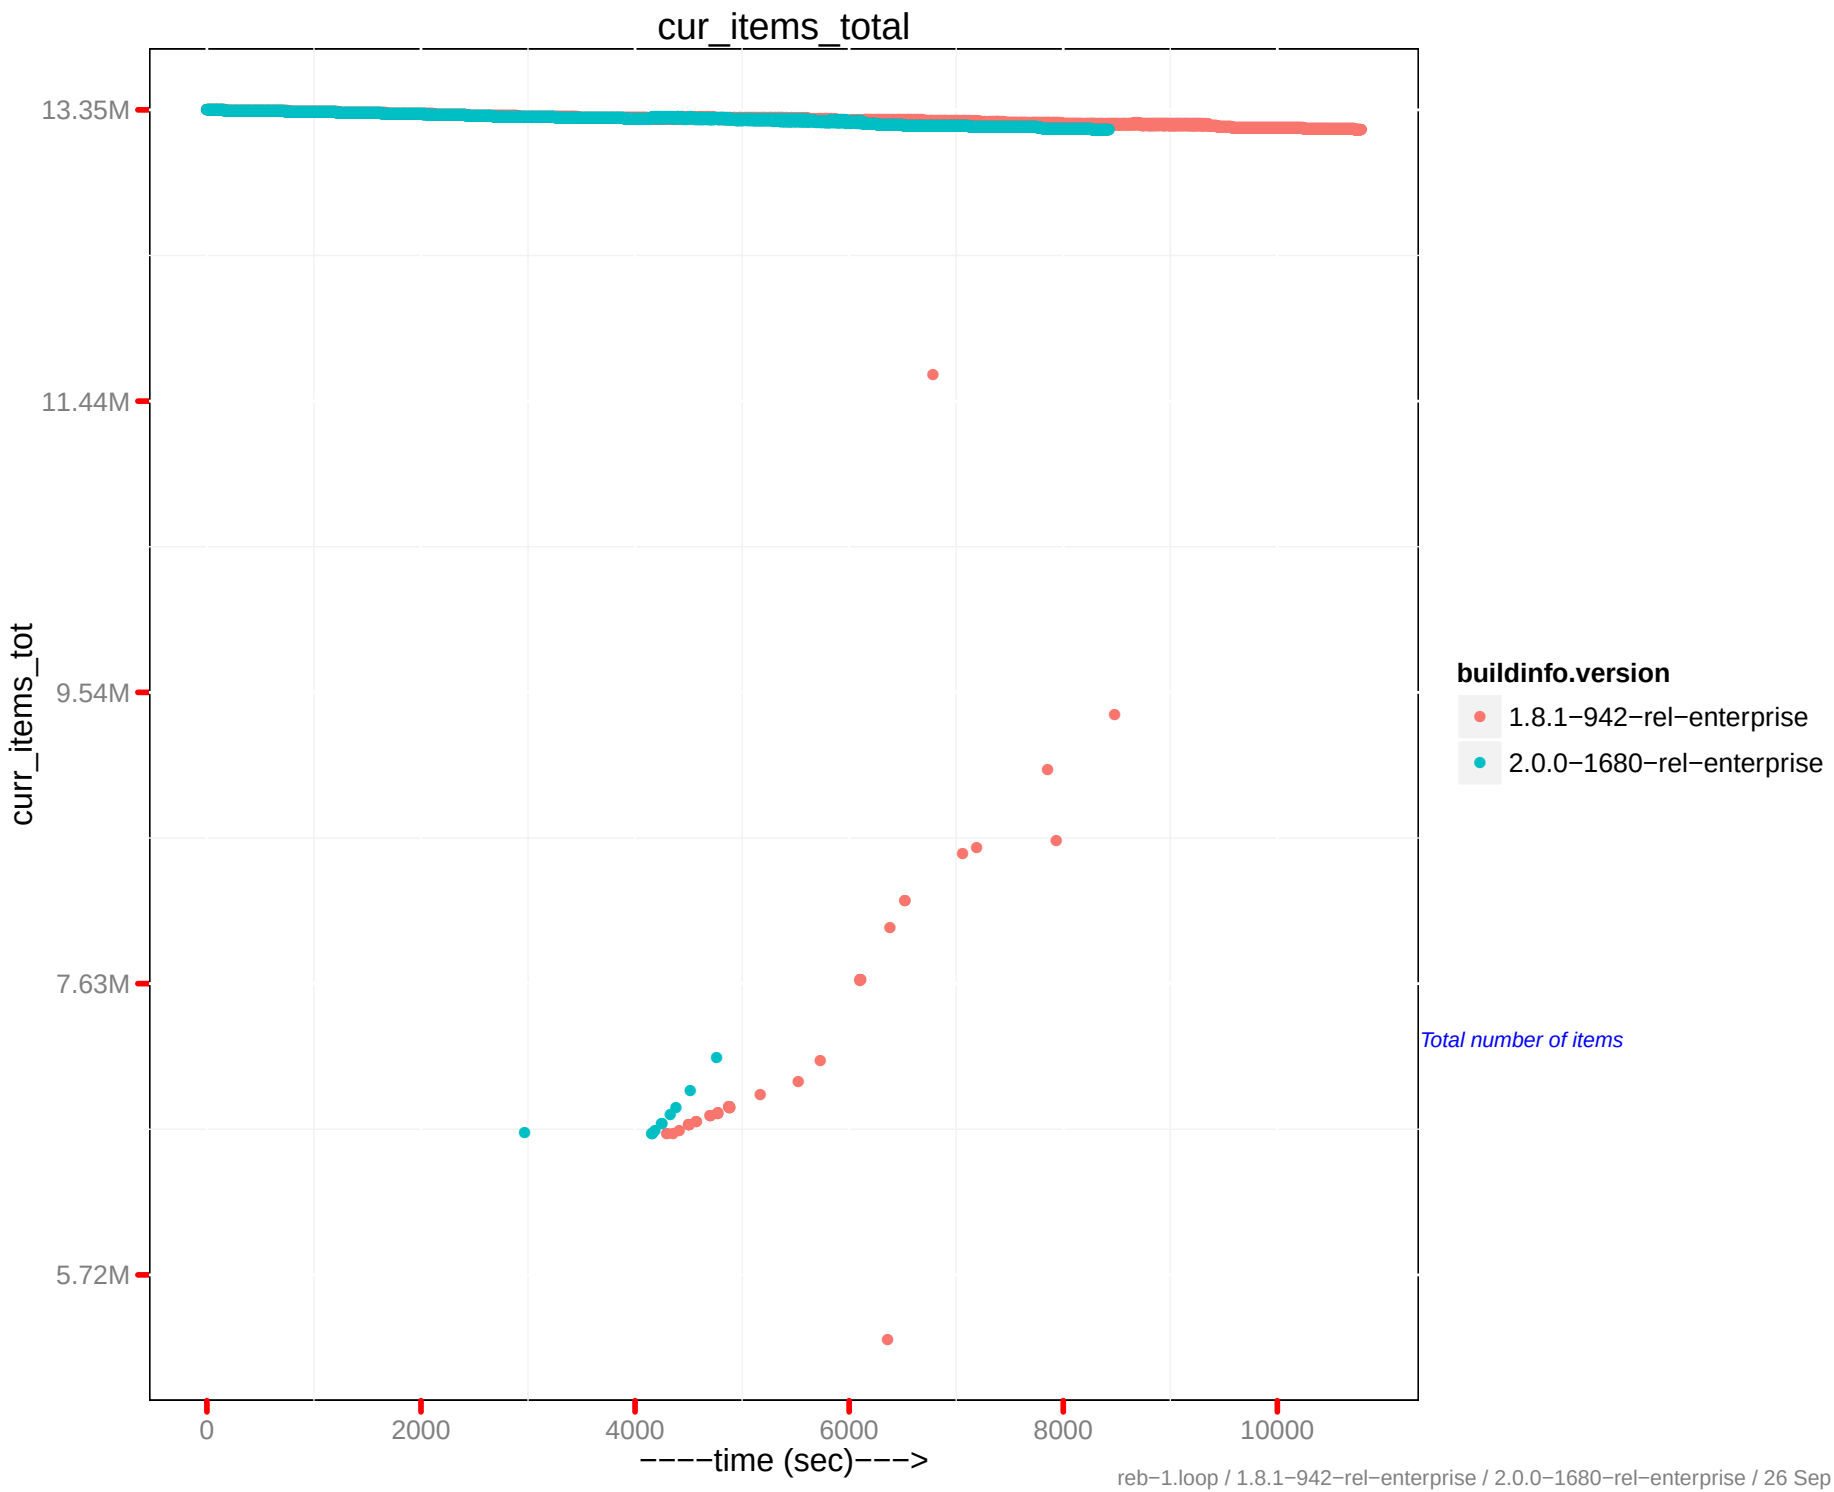
cur_items_total
13.35M
11.44M
9.54M
7.63M
5.72M
curr_items_tot
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Total number of items
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

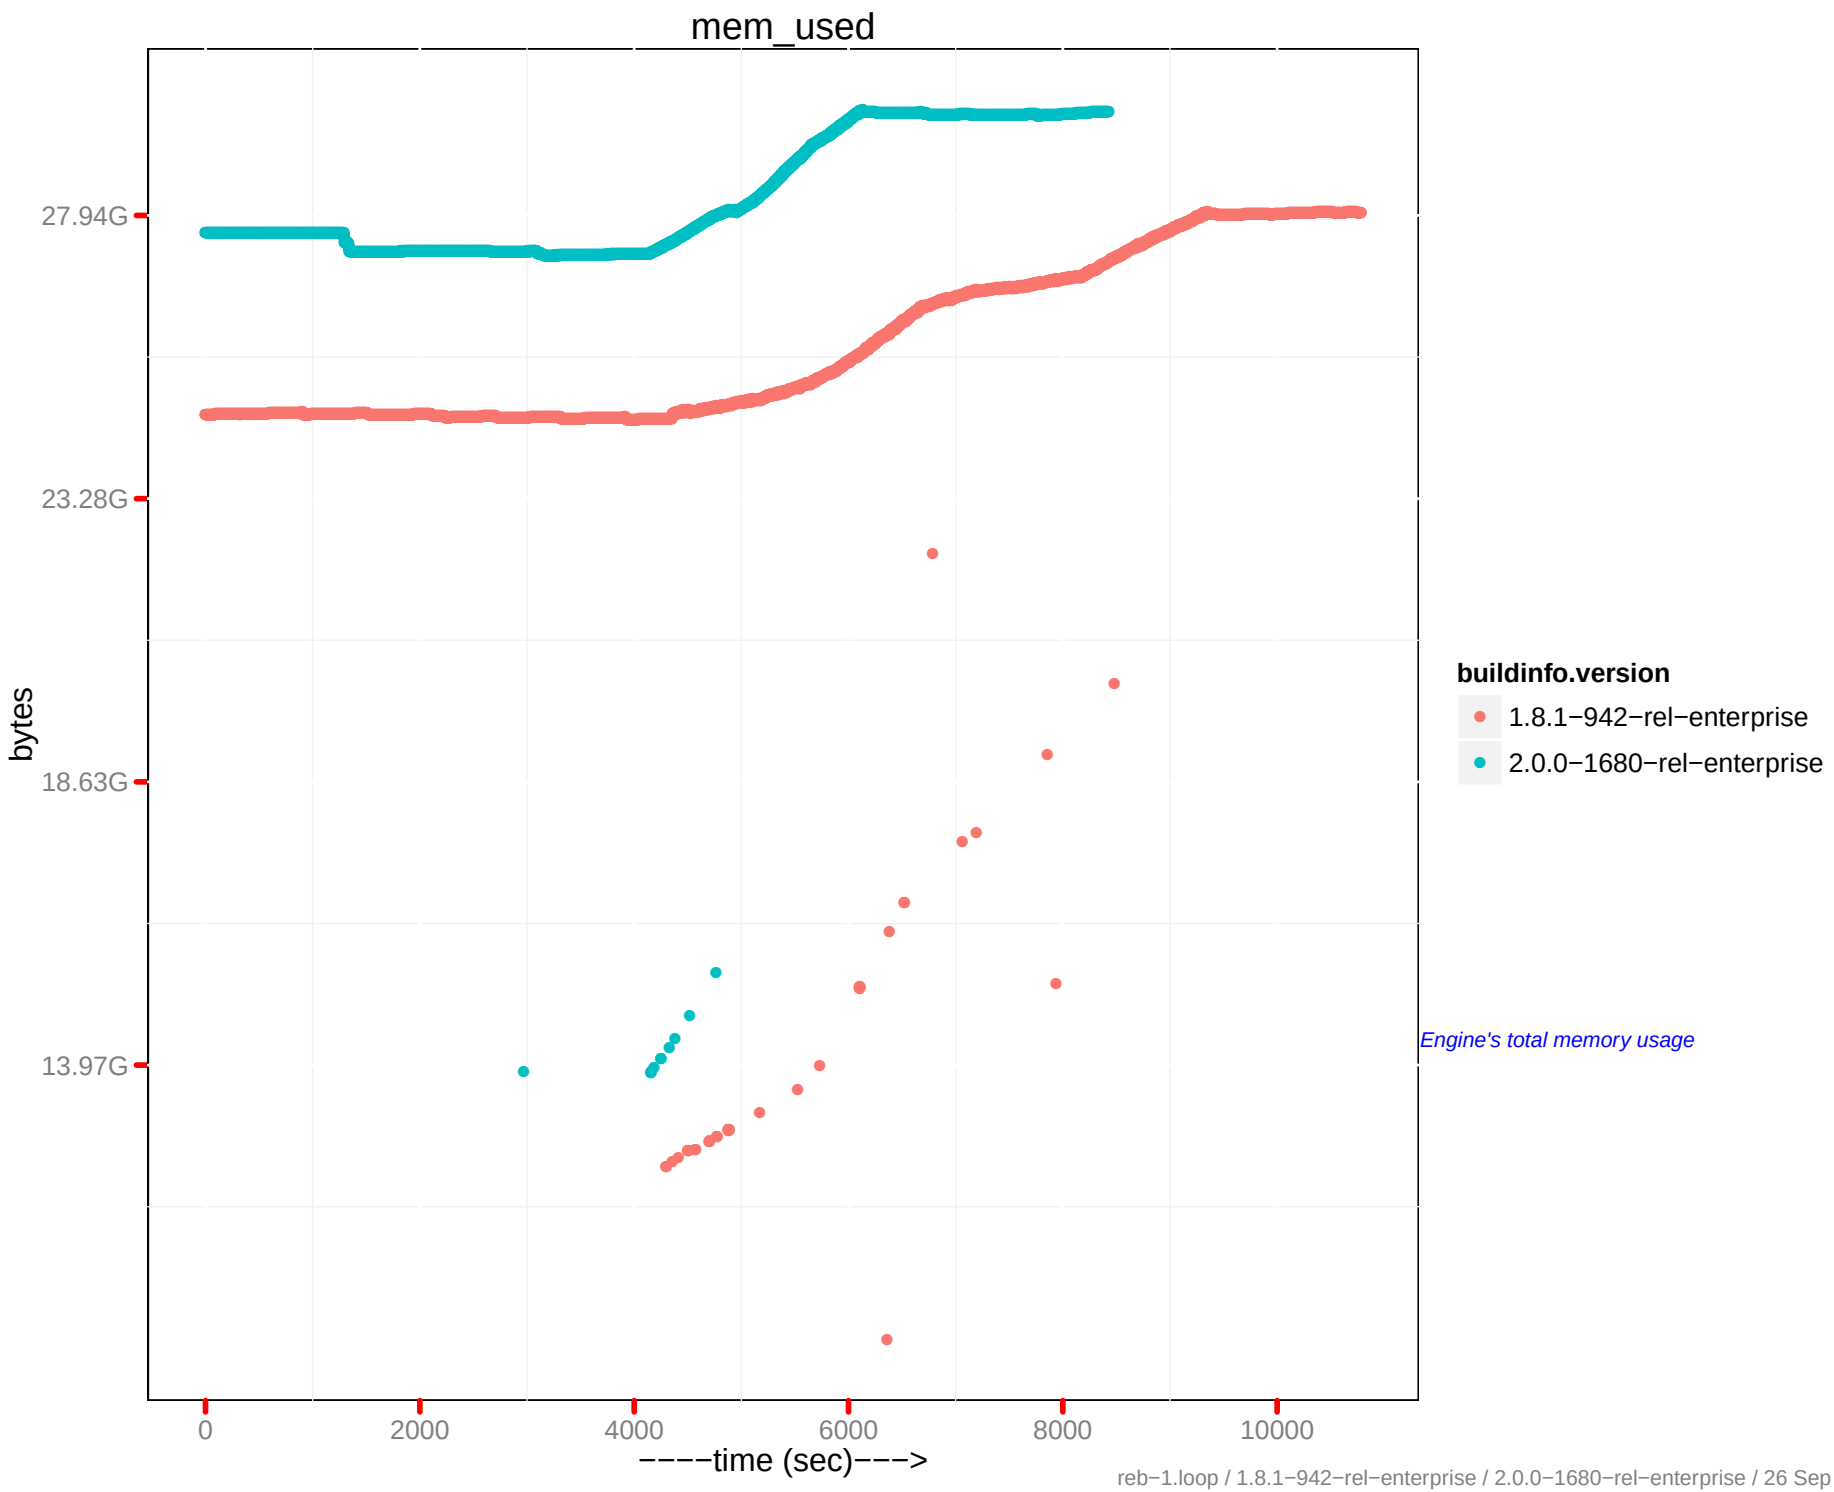
mem_used
27.94G
23.28G
18.63G
13.97G
bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Engine's total memory usage
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

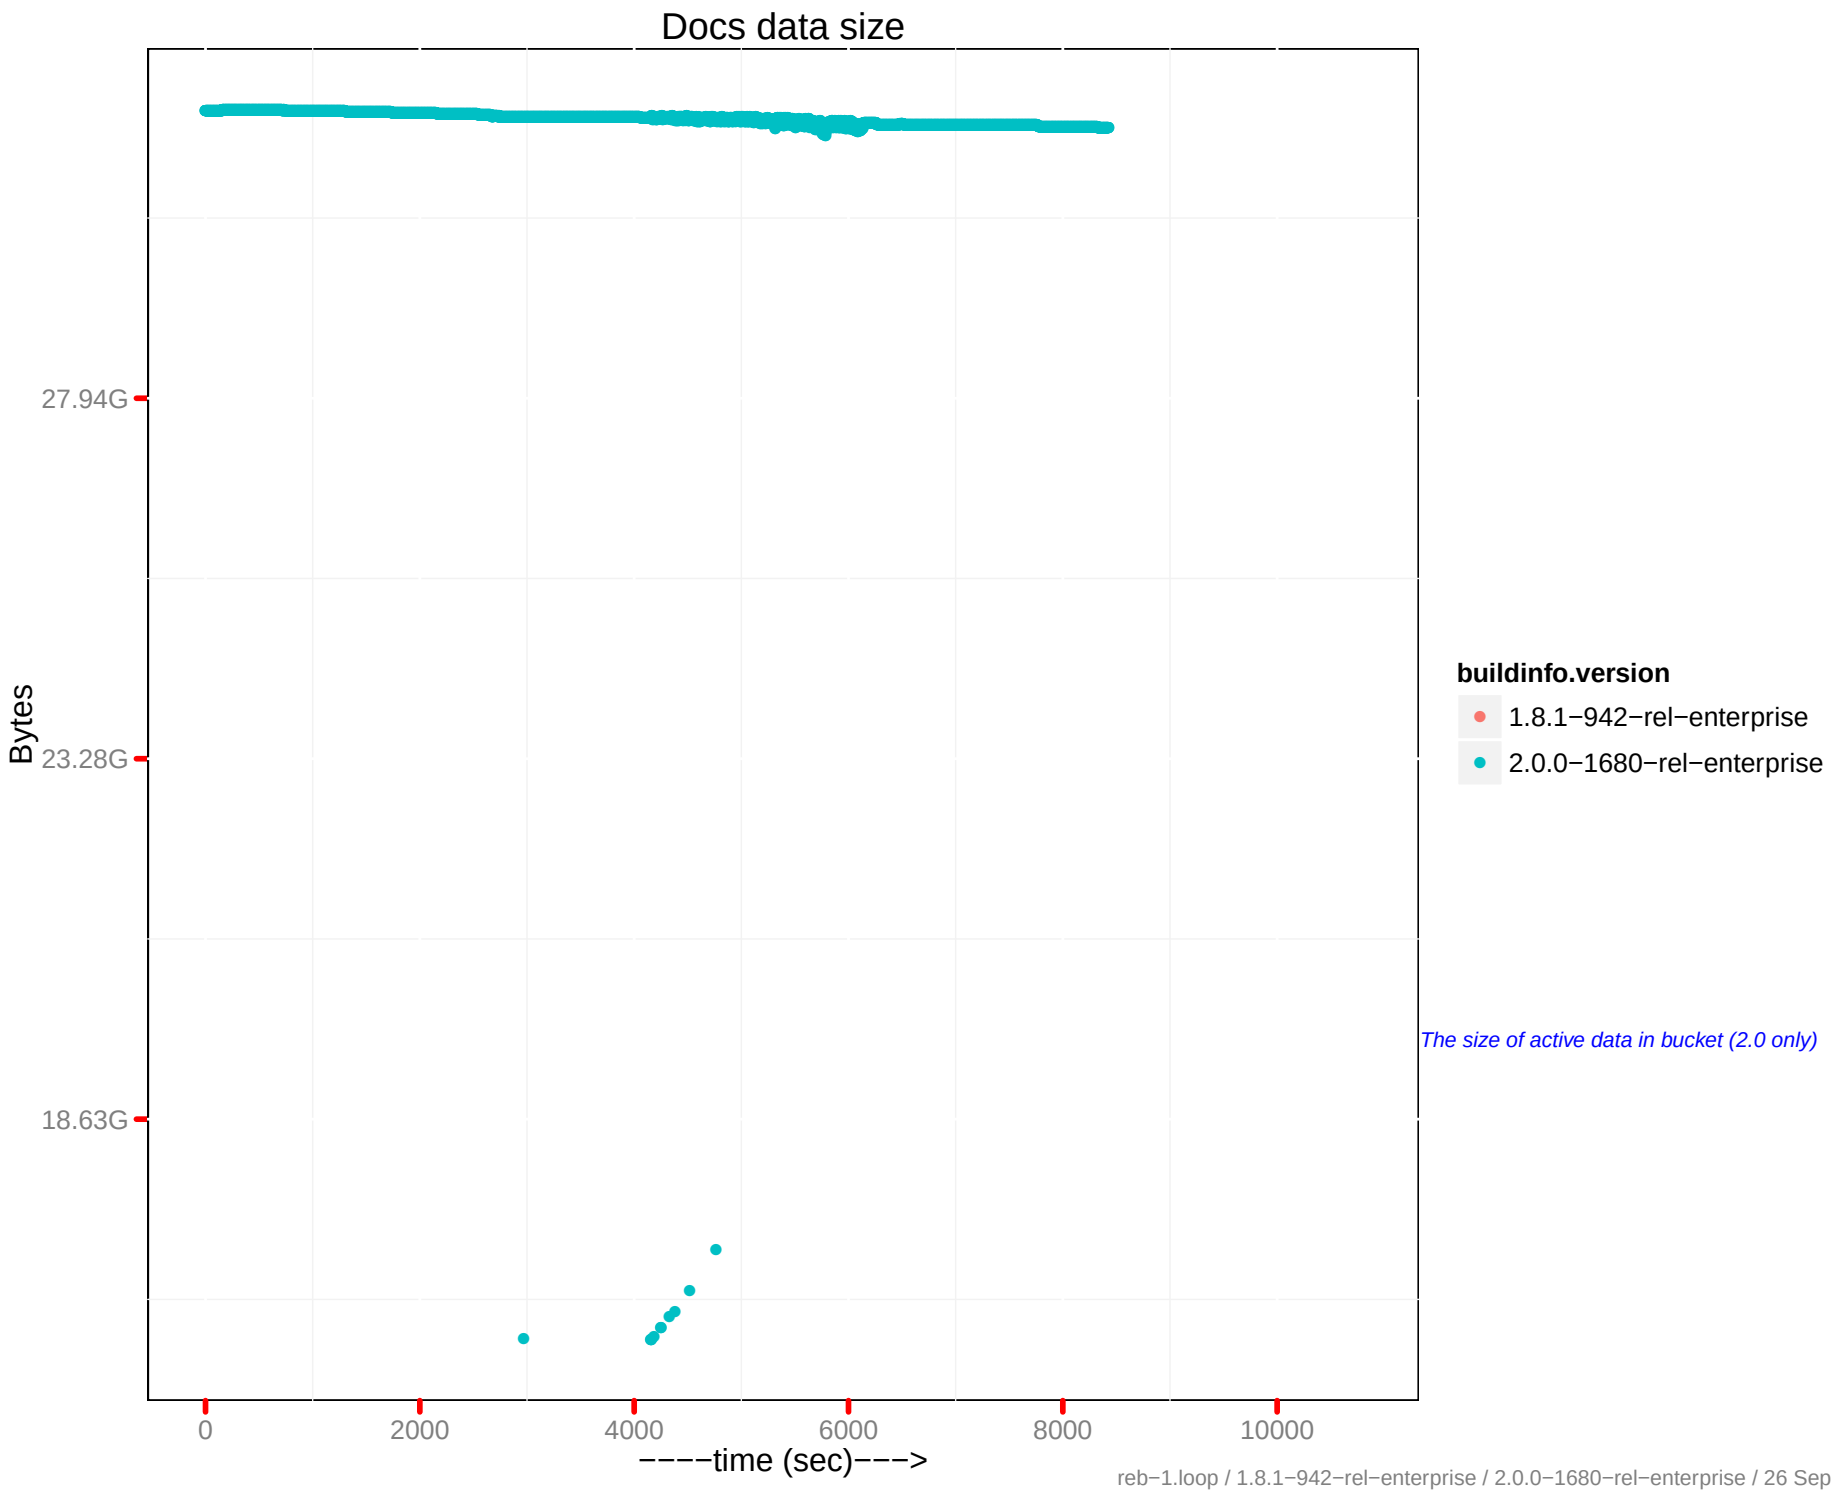
Docs data size
27.94G
23.28G
18.63G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
The size of active data in bucket (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

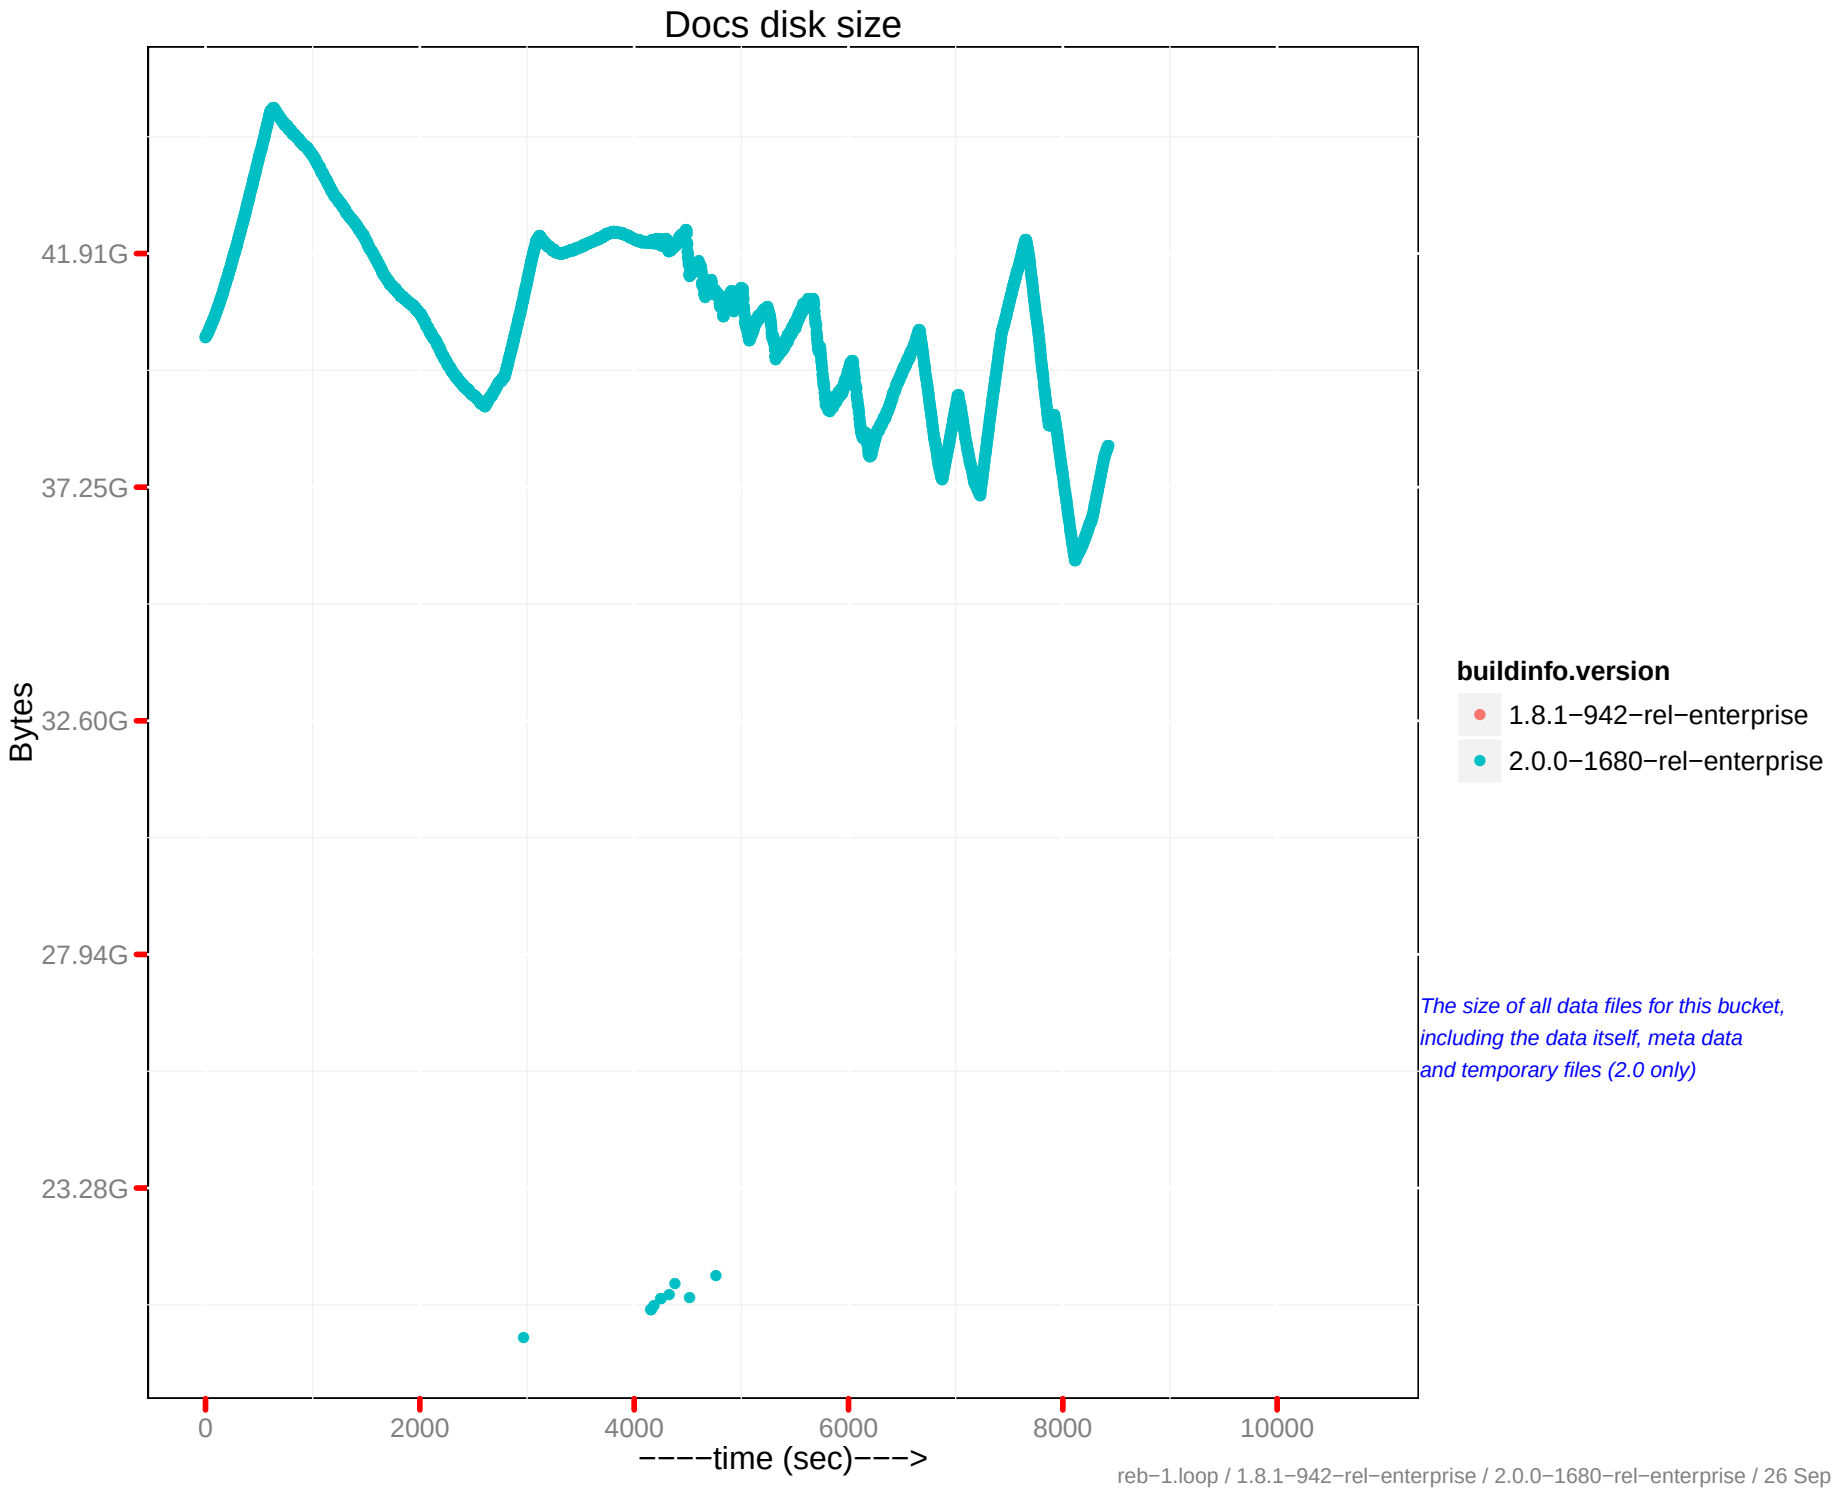
Docs disk size
41.91G
37.25G
32.60G
27.94G
23.28G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

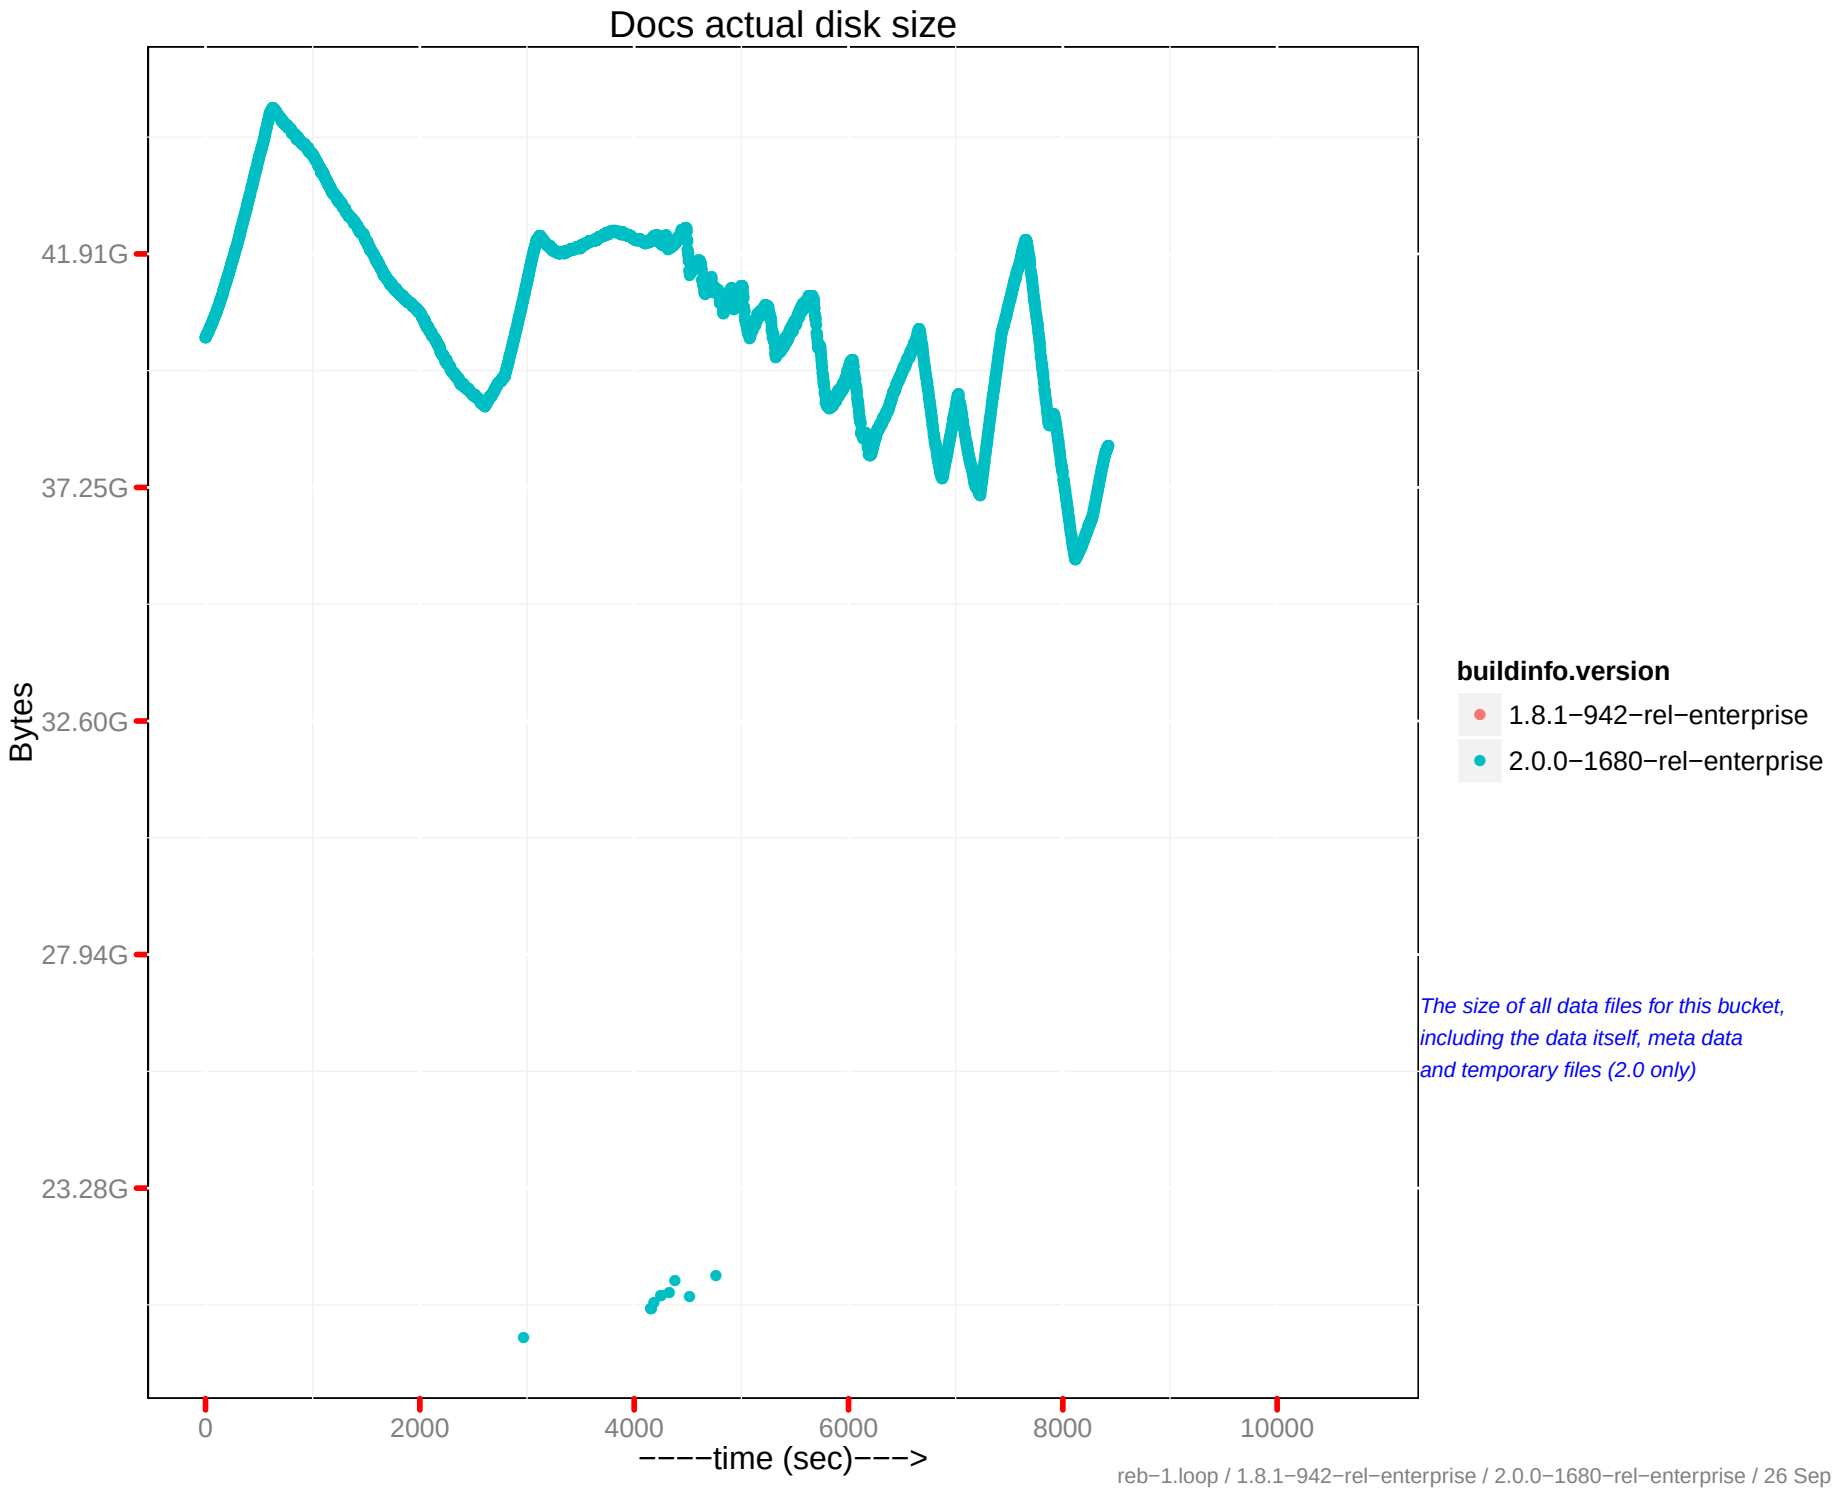
Docs actual disk size
41.91G
37.25G
32.60G
27.94G
23.28G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

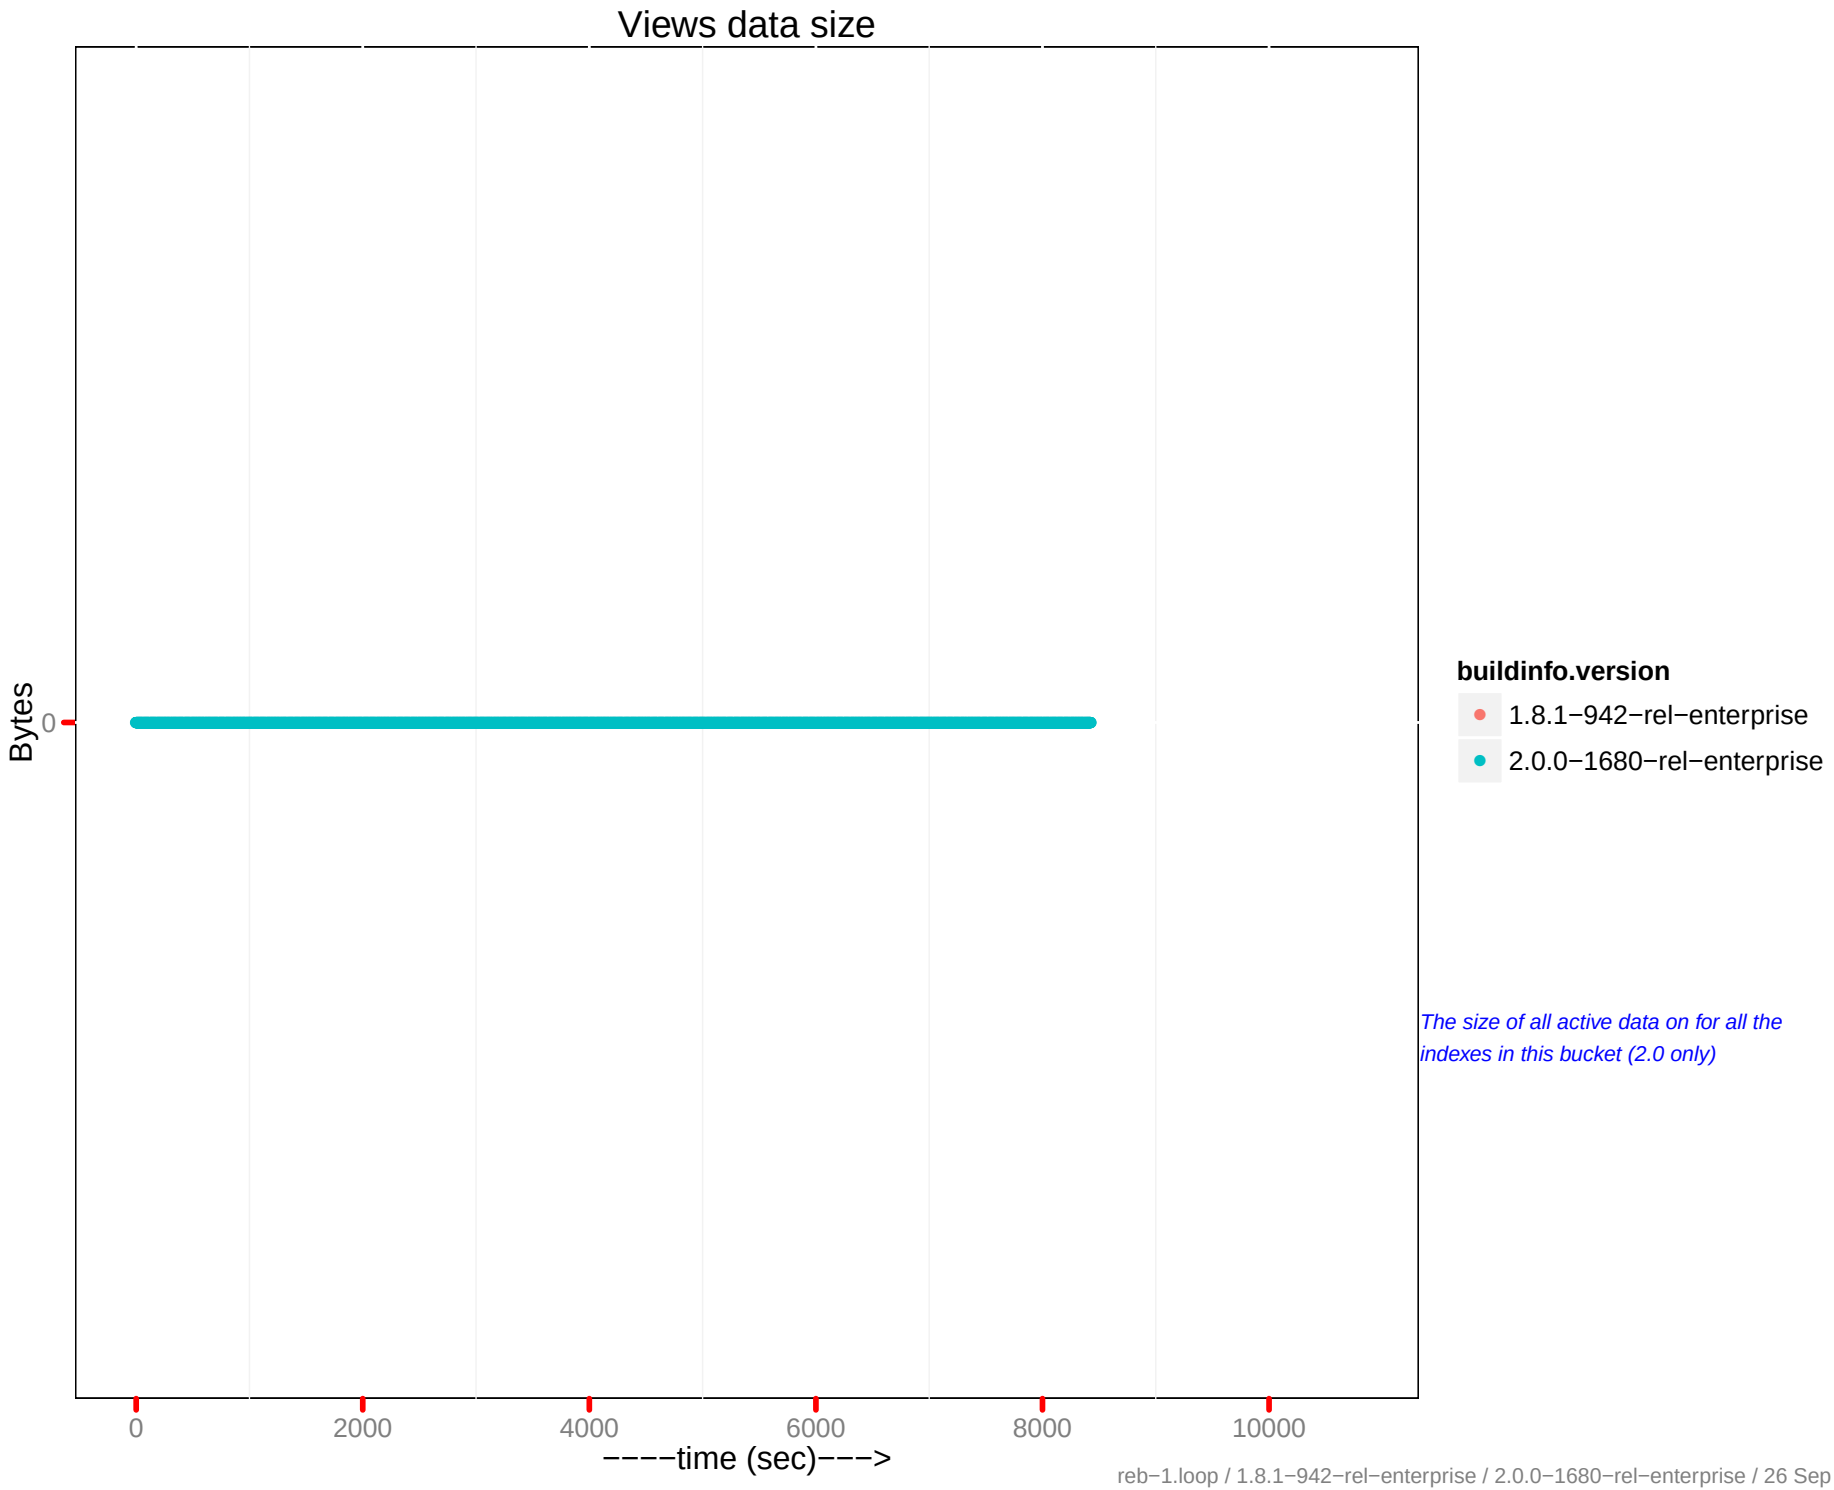
Views data size
Bytes
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

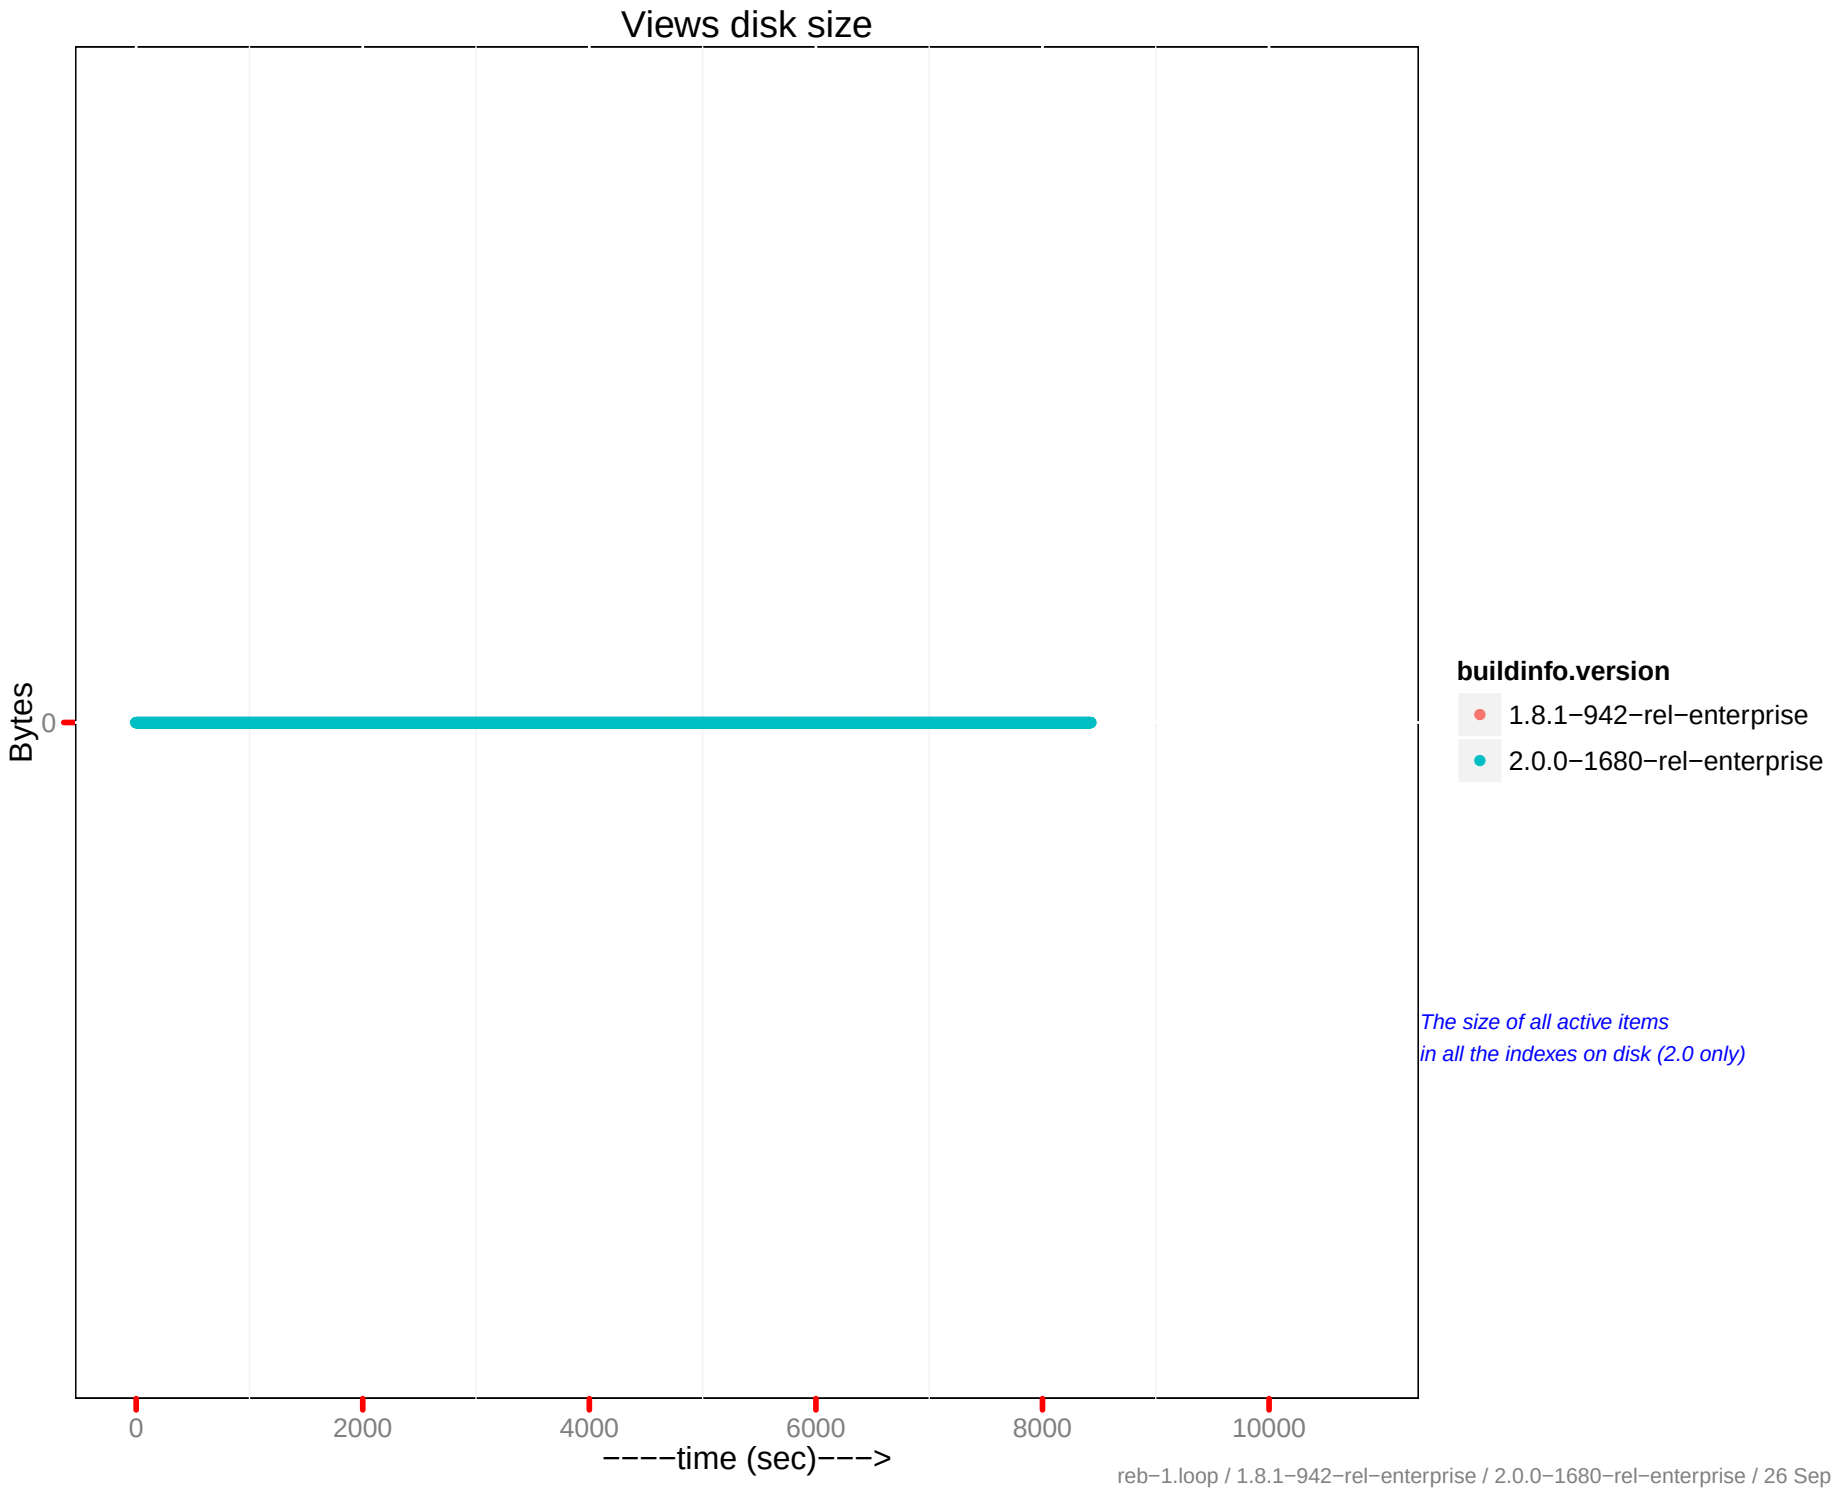
Views disk size
Bytes
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
The size of all active items
in all the indexes on disk (2.0 only)
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

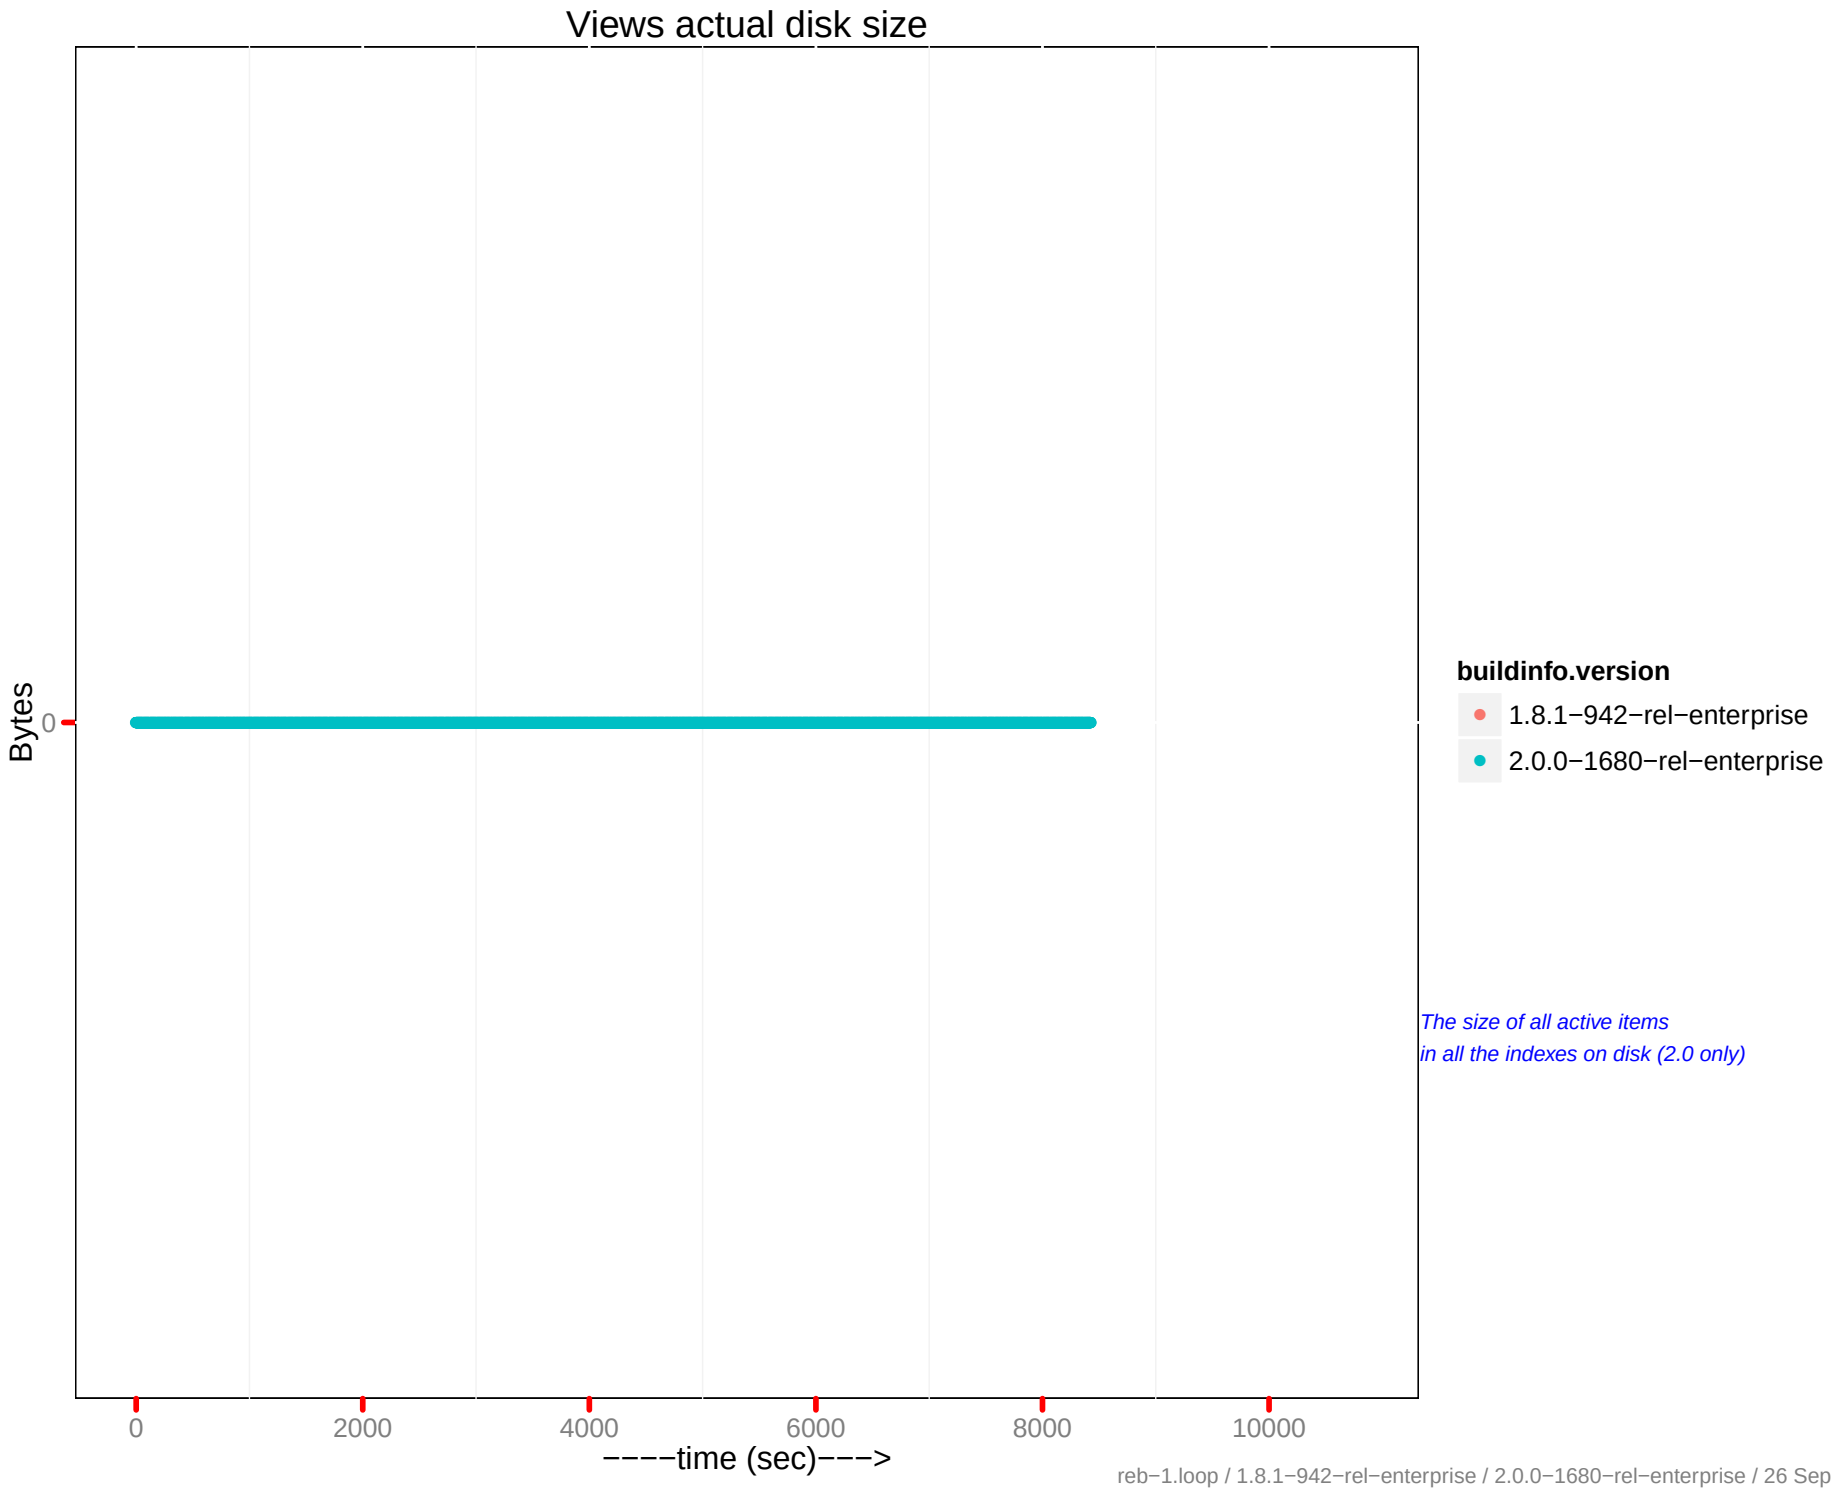
Views actual disk size
Bytes
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

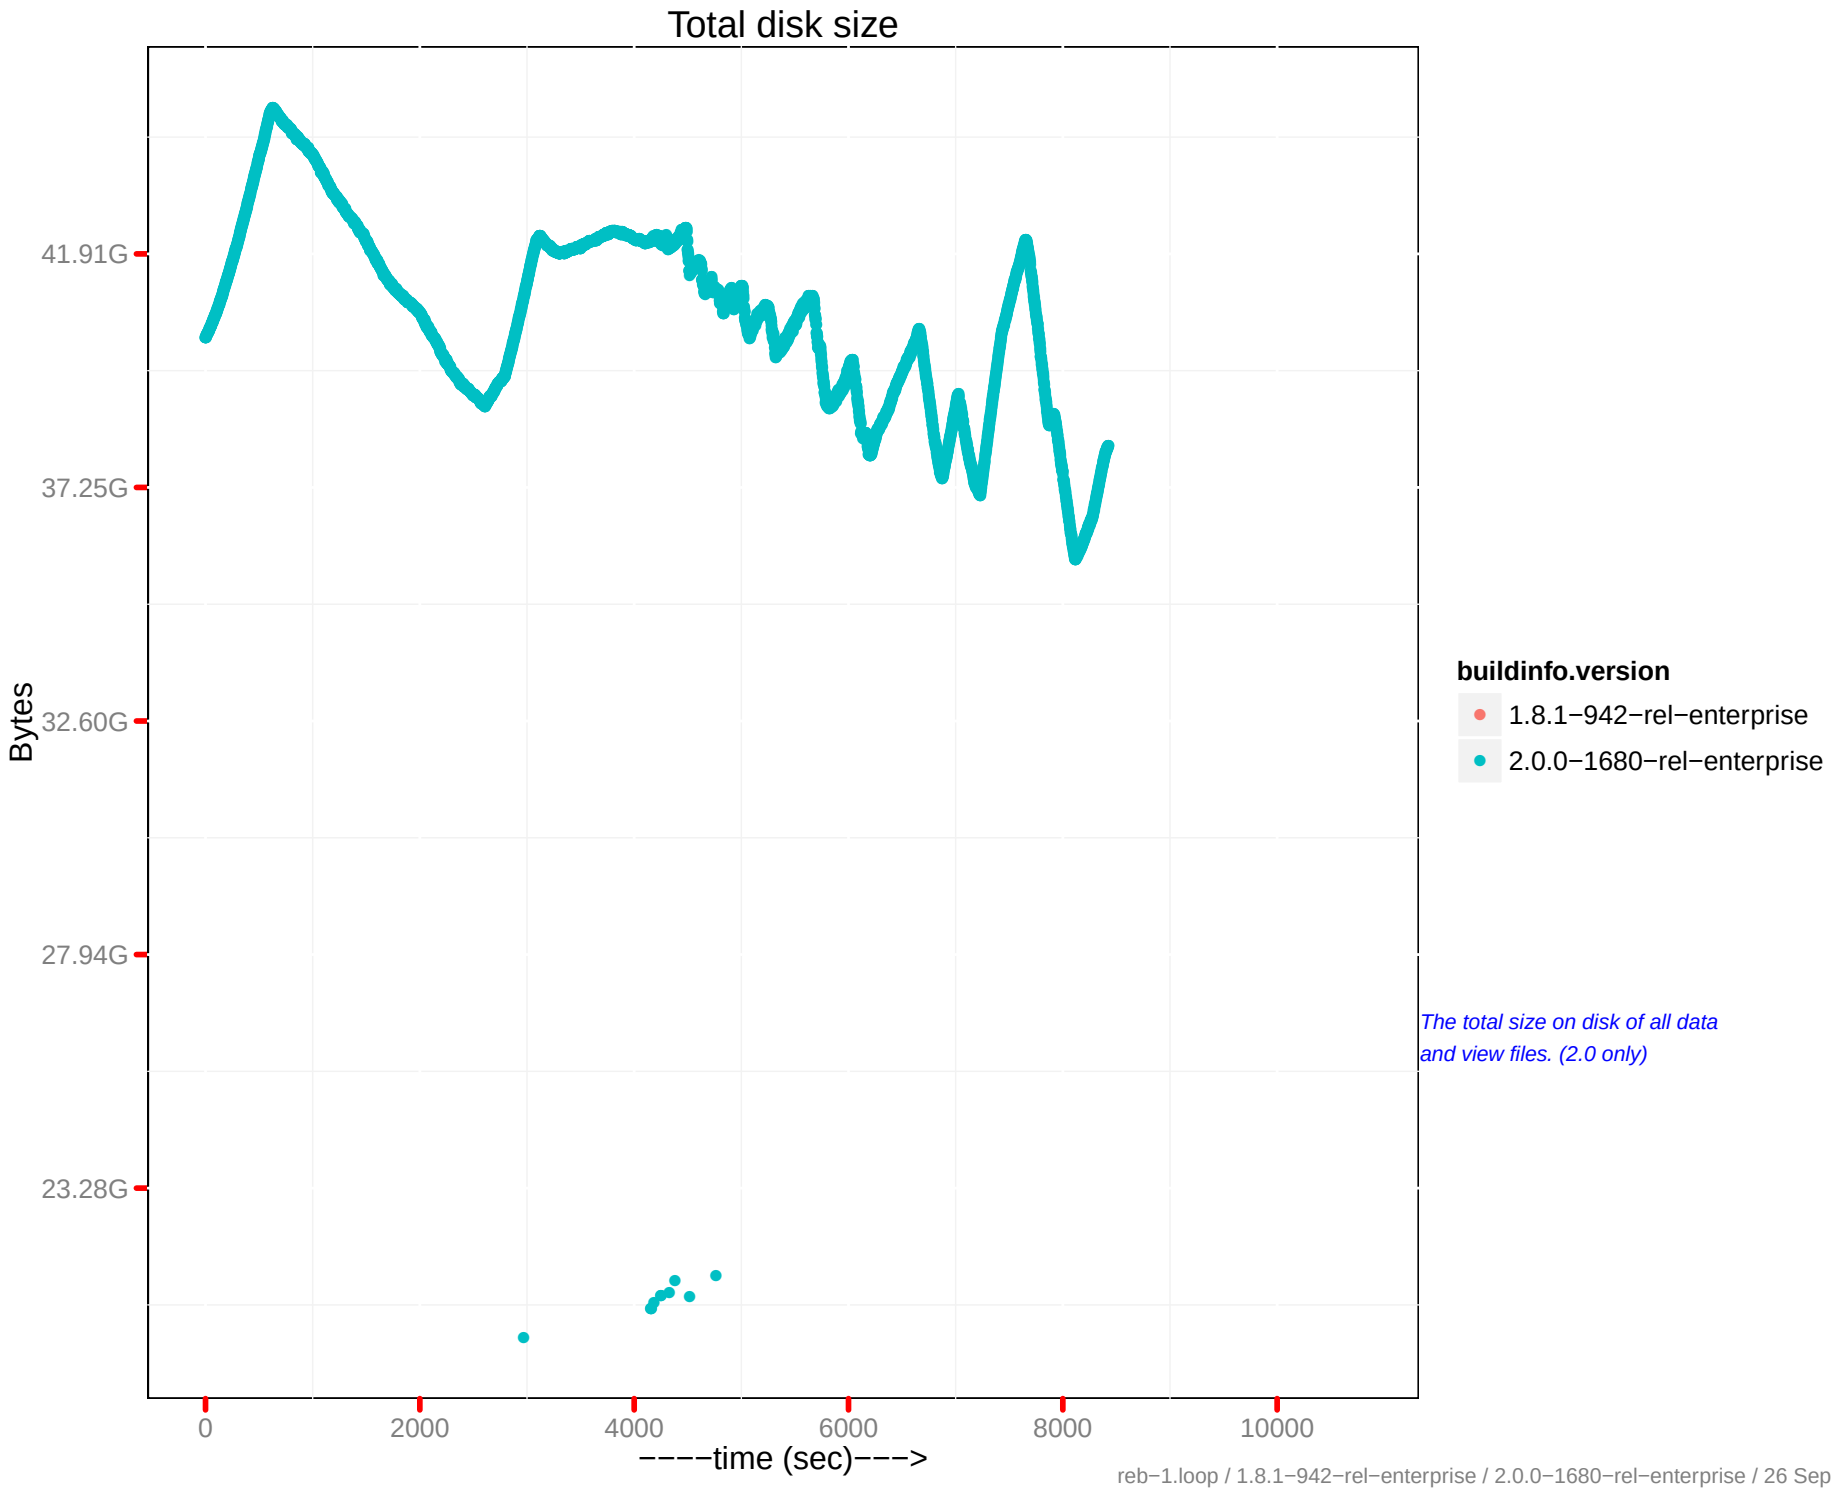
Total disk size
41.91G
37.25G
32.60G
27.94G
23.28G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

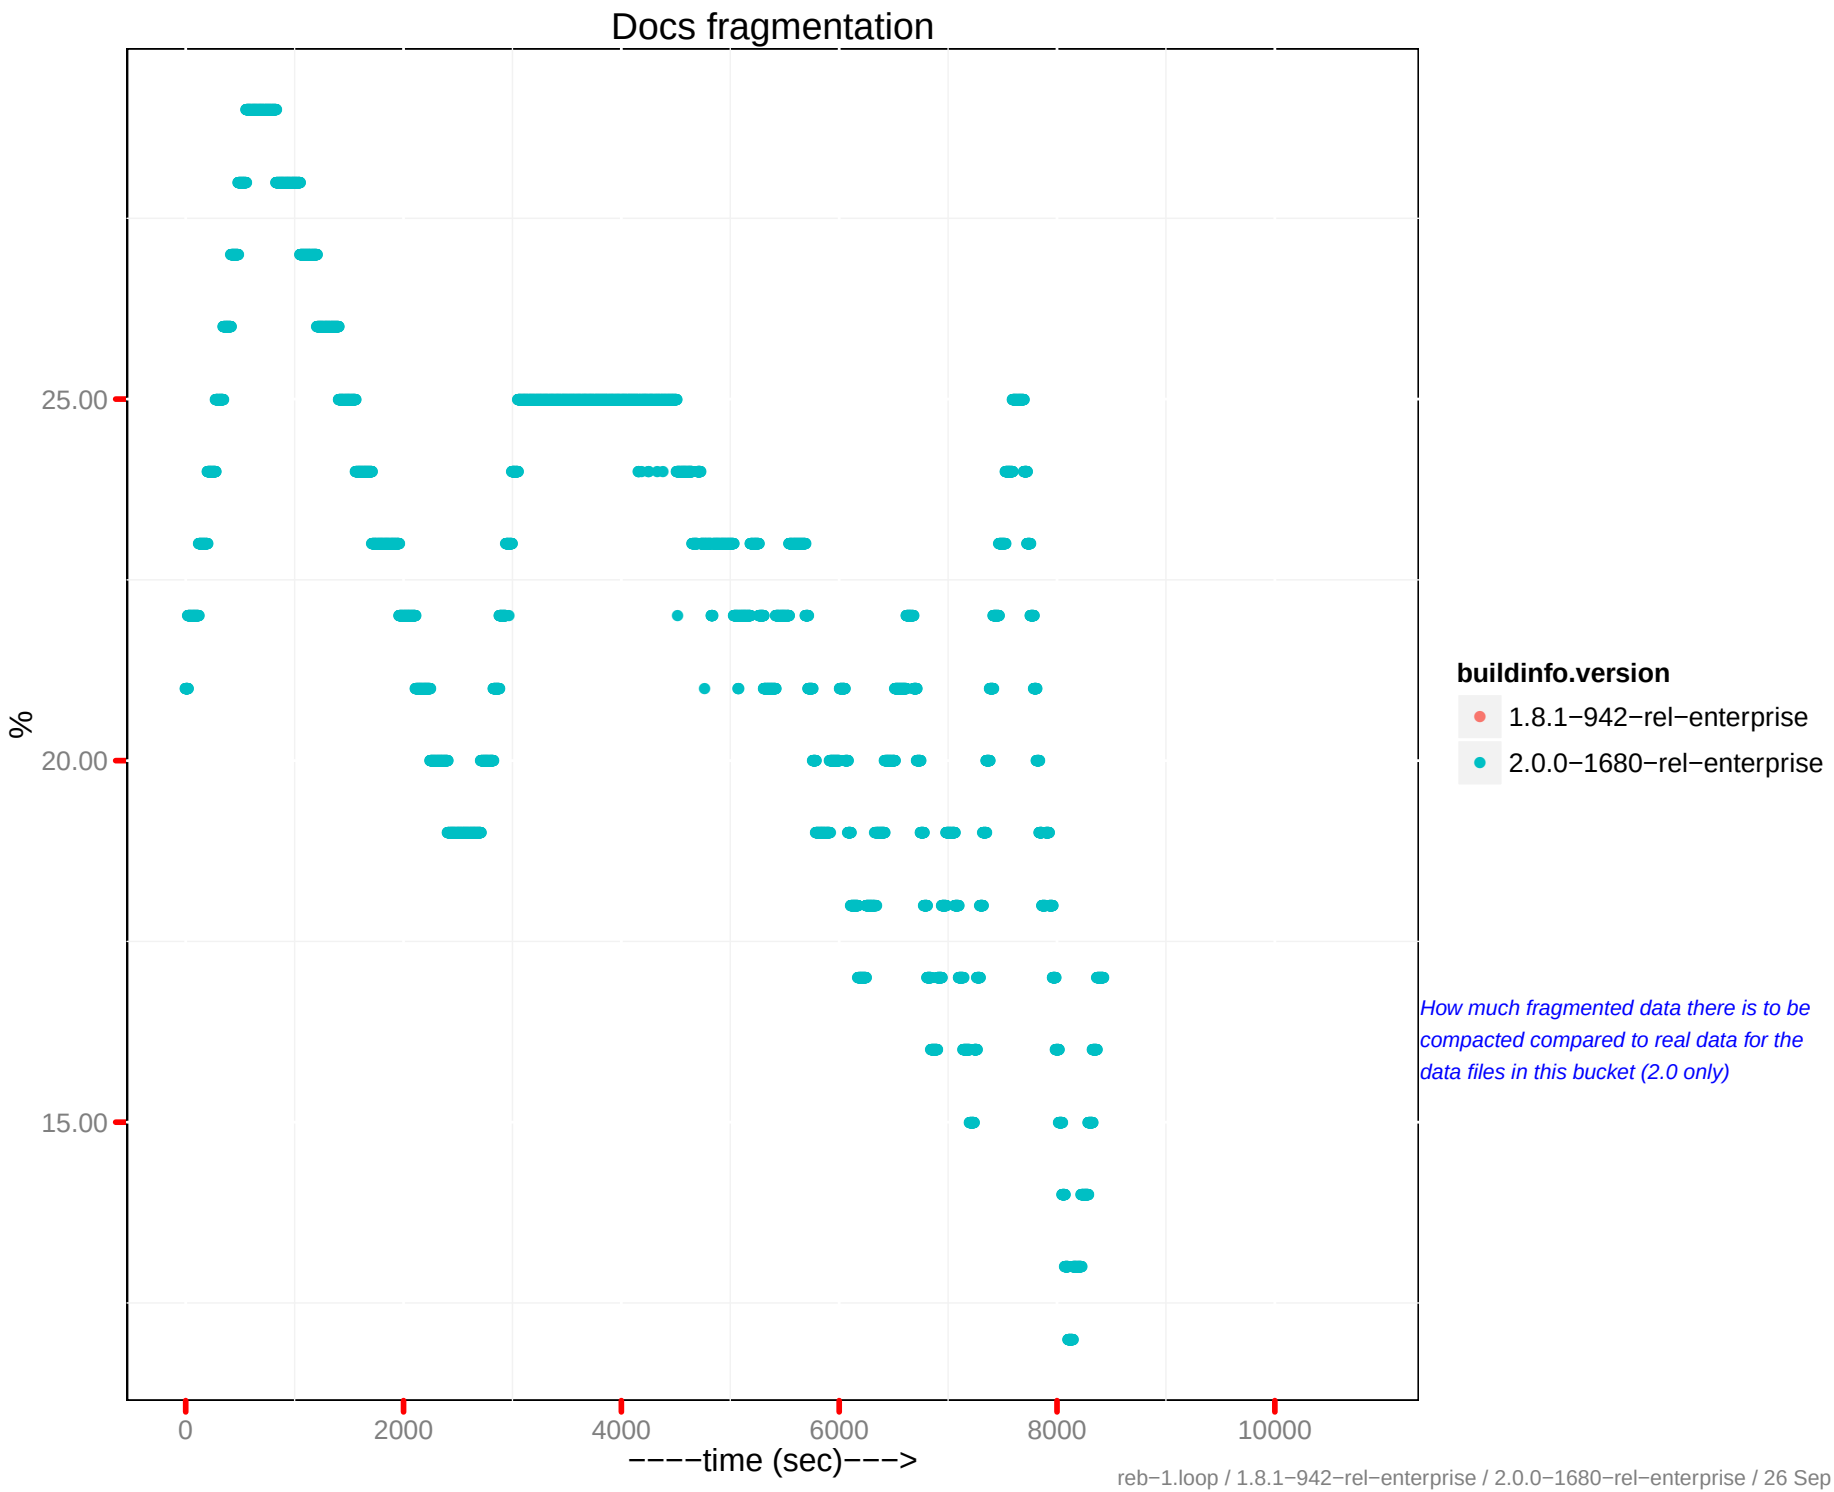
Docs fragmentation
25.00
20.00
15.00
%
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

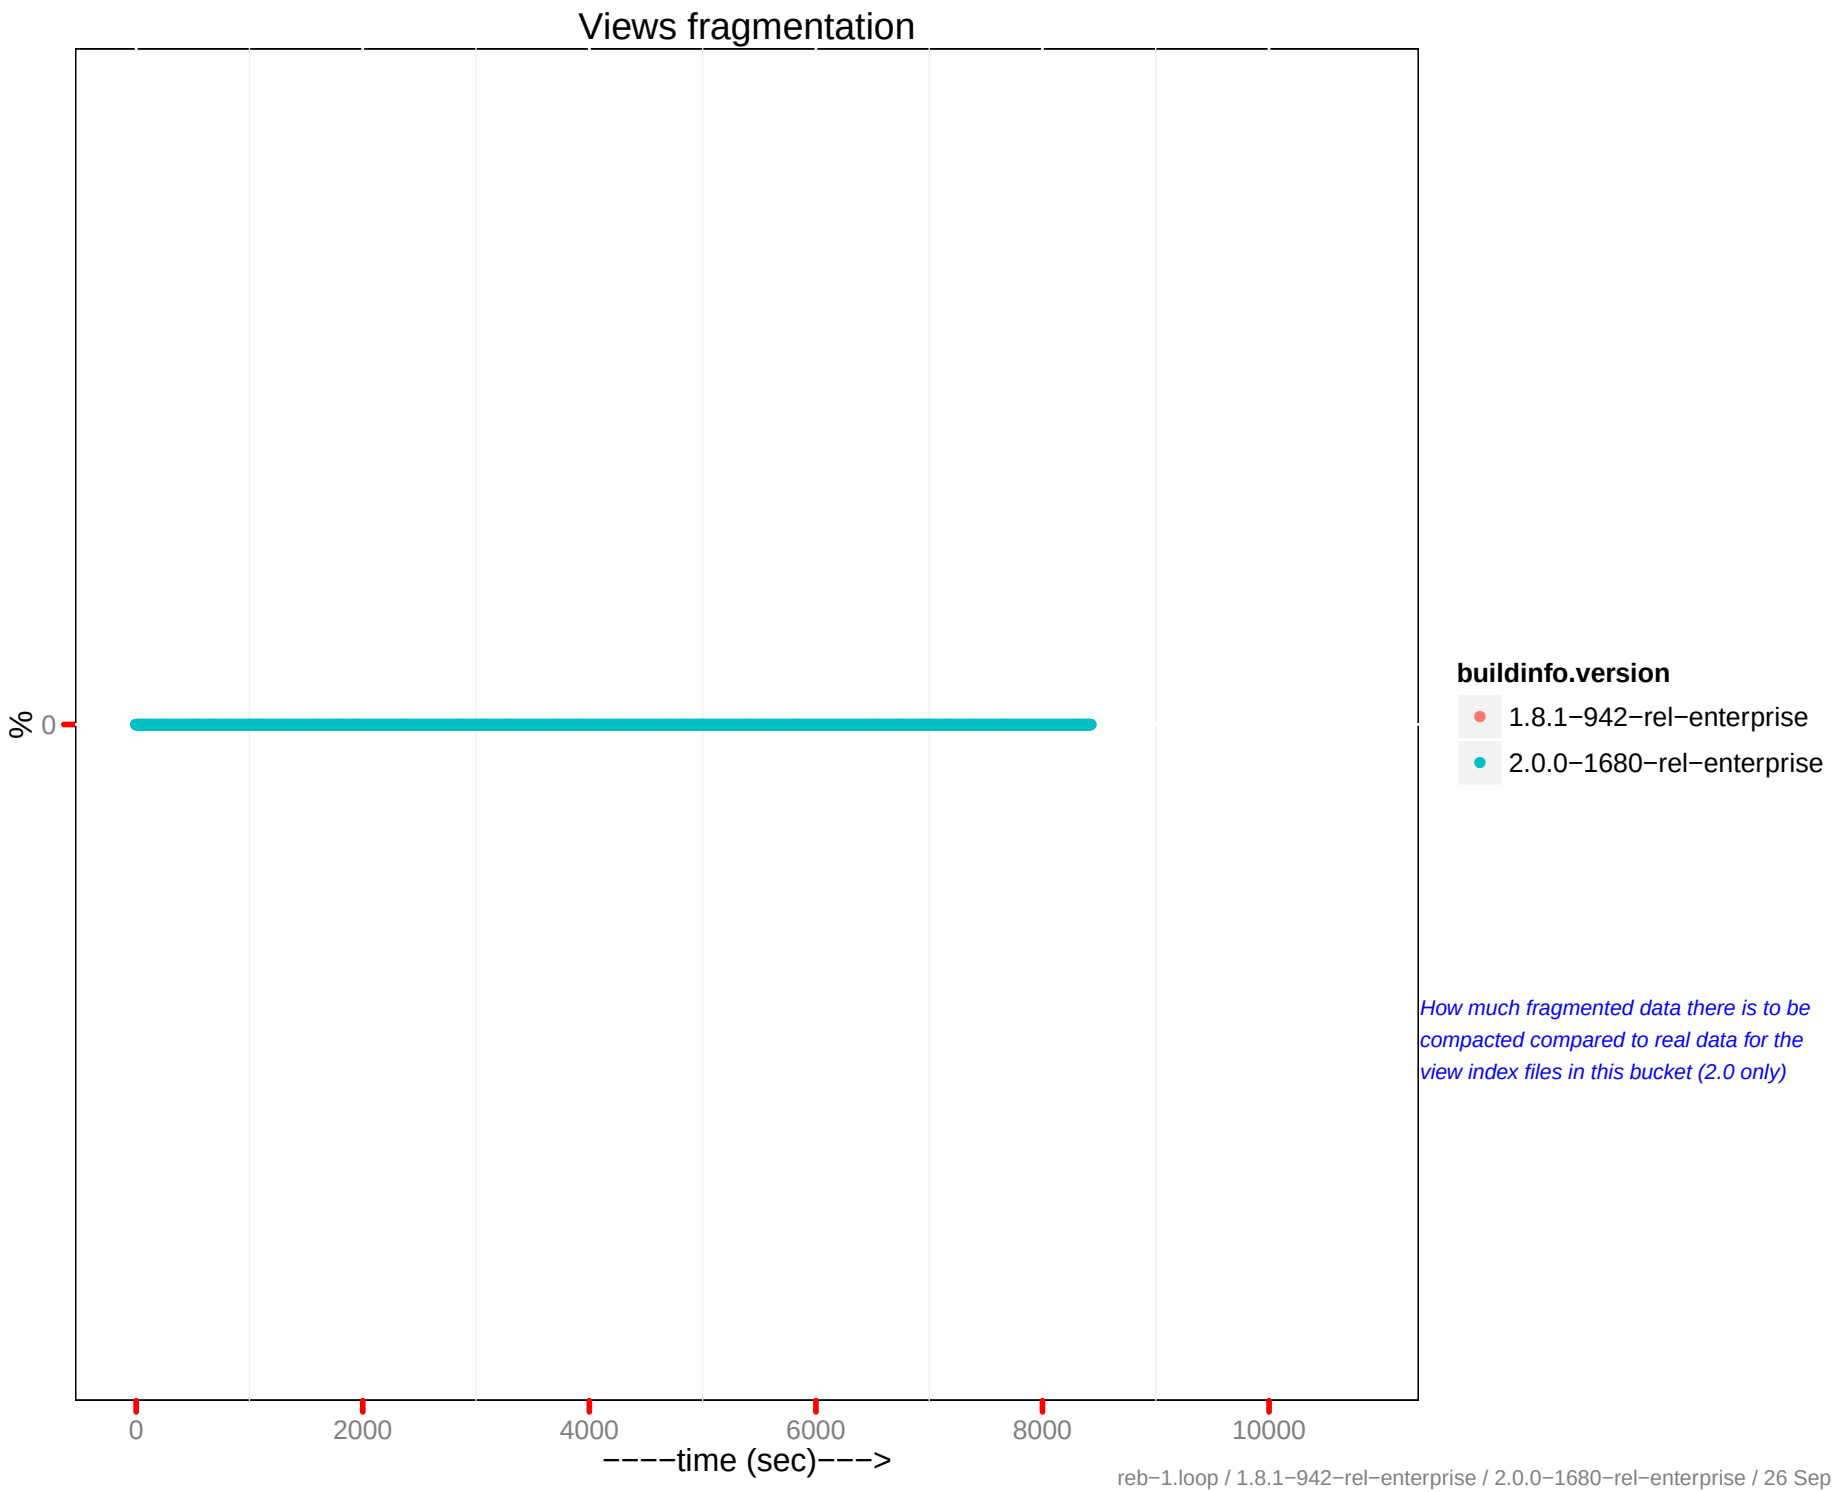
Views fragmentation
%
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

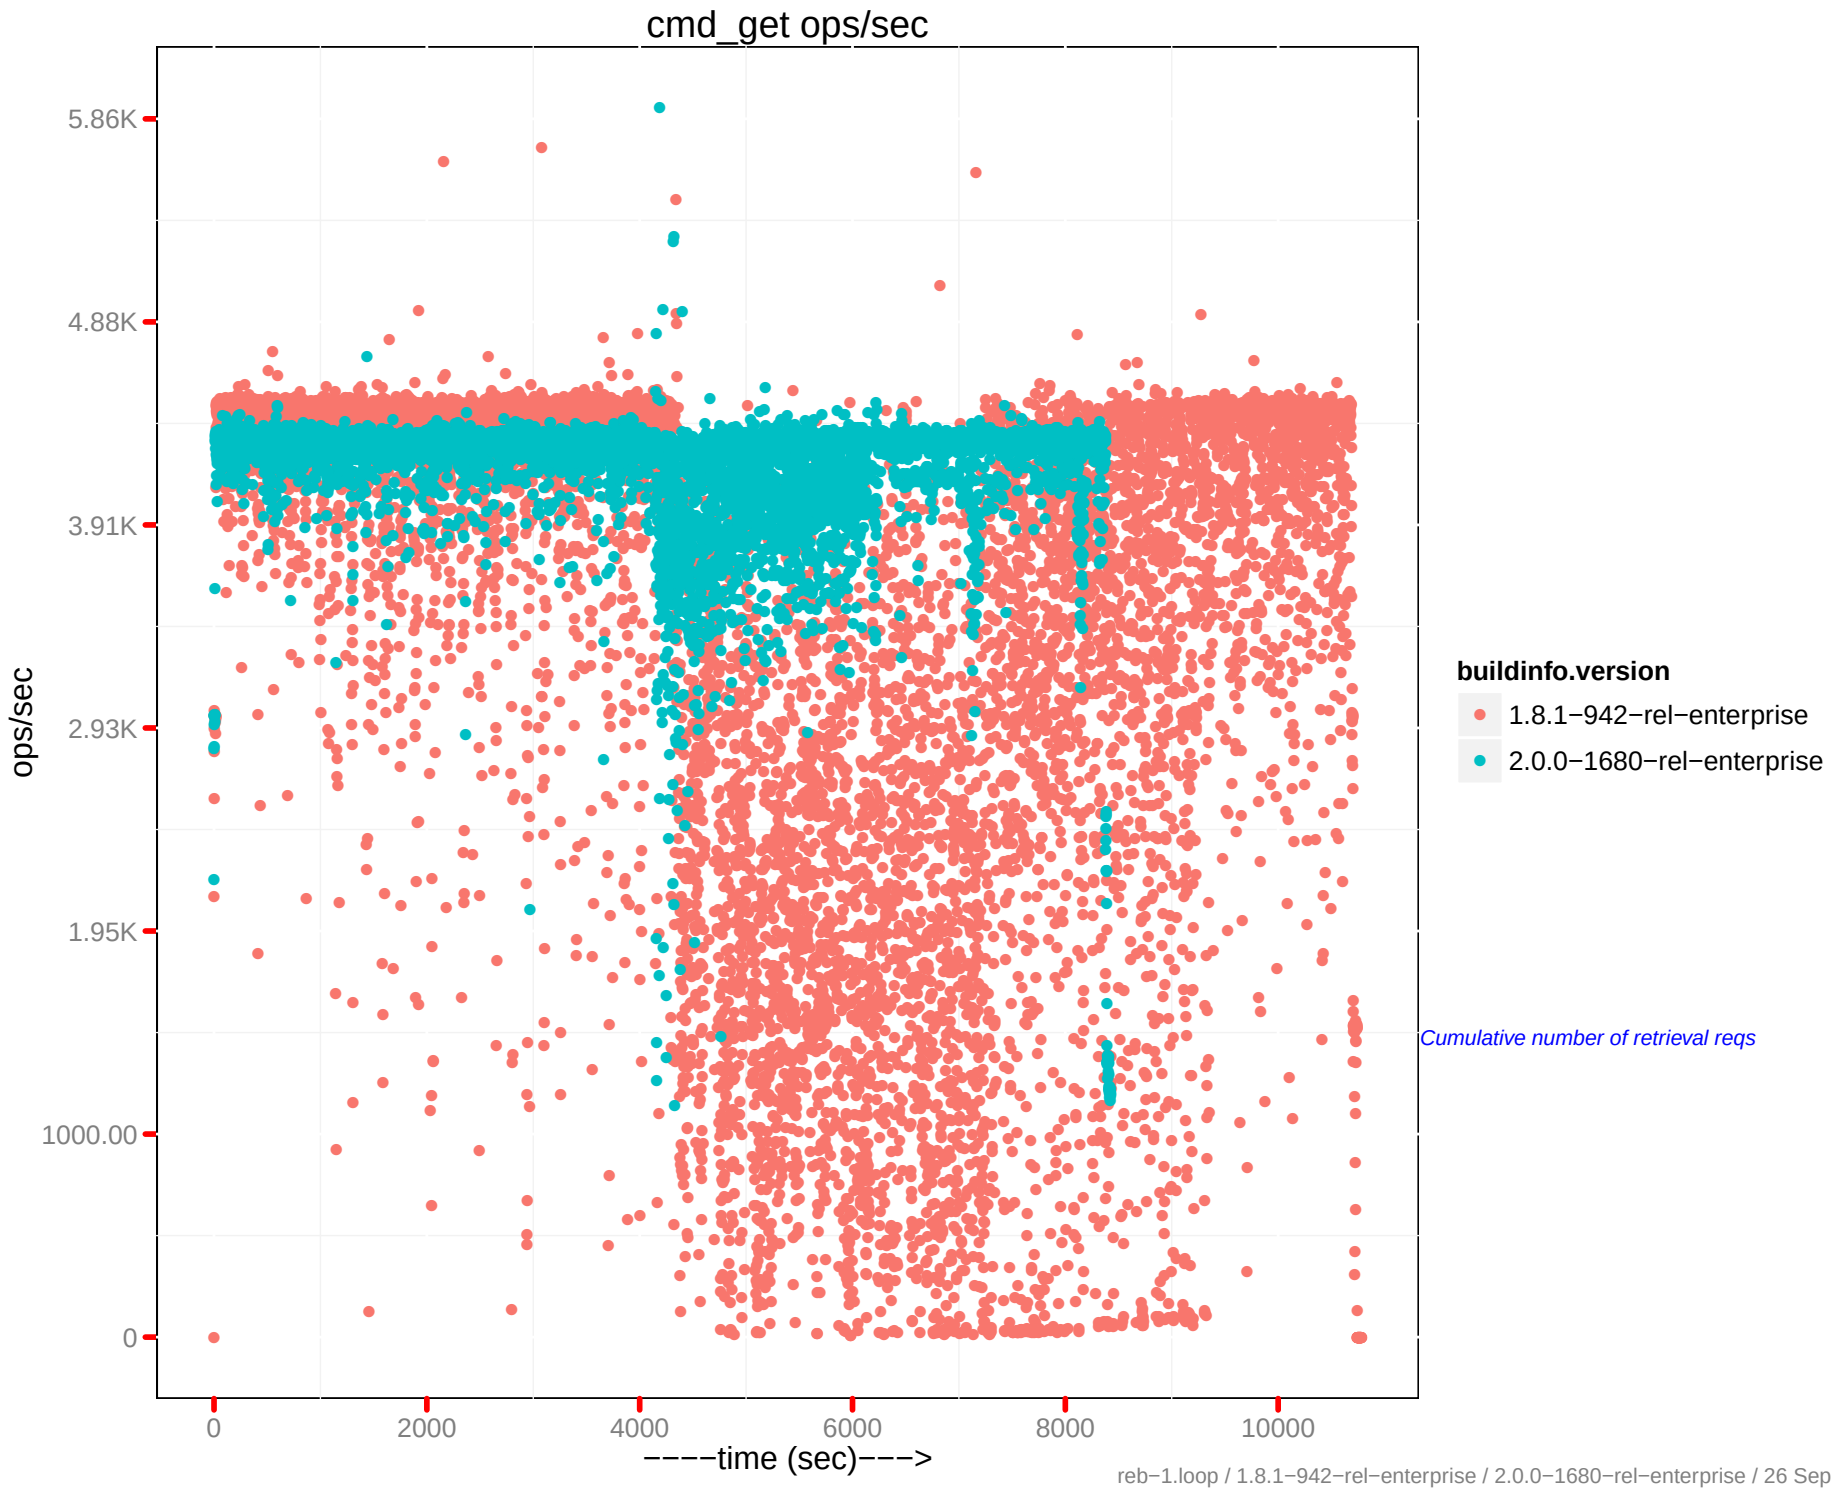
cmd_get ops/sec
5.86K
4.88K
3.91K
2.93K
1.95K
1000.00
0
ops/sec
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Cumulative number of retrieval reqs
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

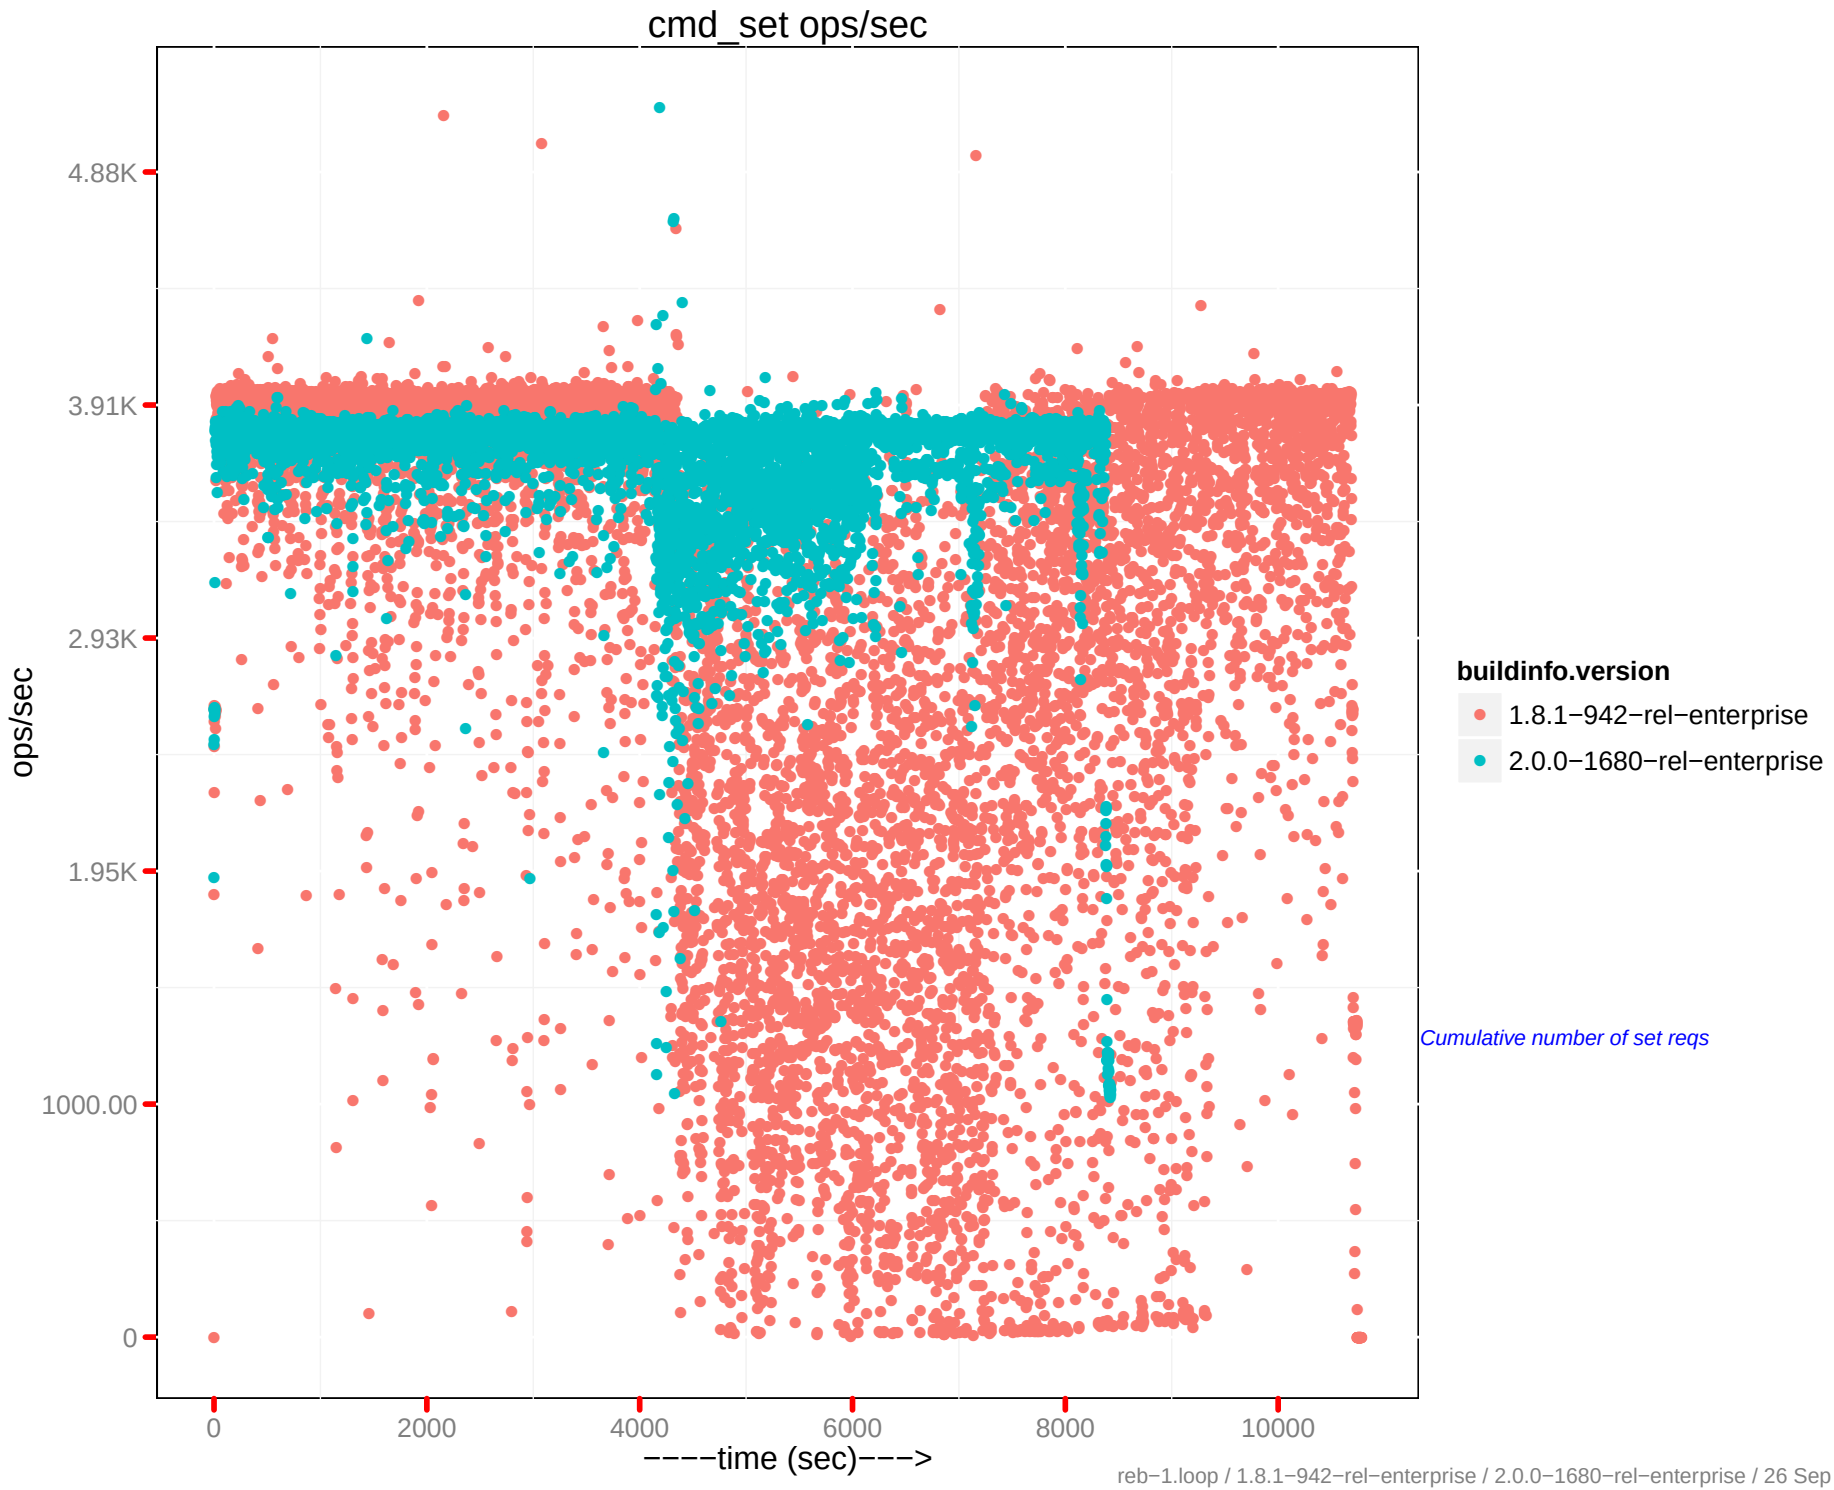
cmd_set ops/sec
4.88K
3.91K
2.93K
1.95K
1000.00
0
ops/sec
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Cumulative number of set reqs
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

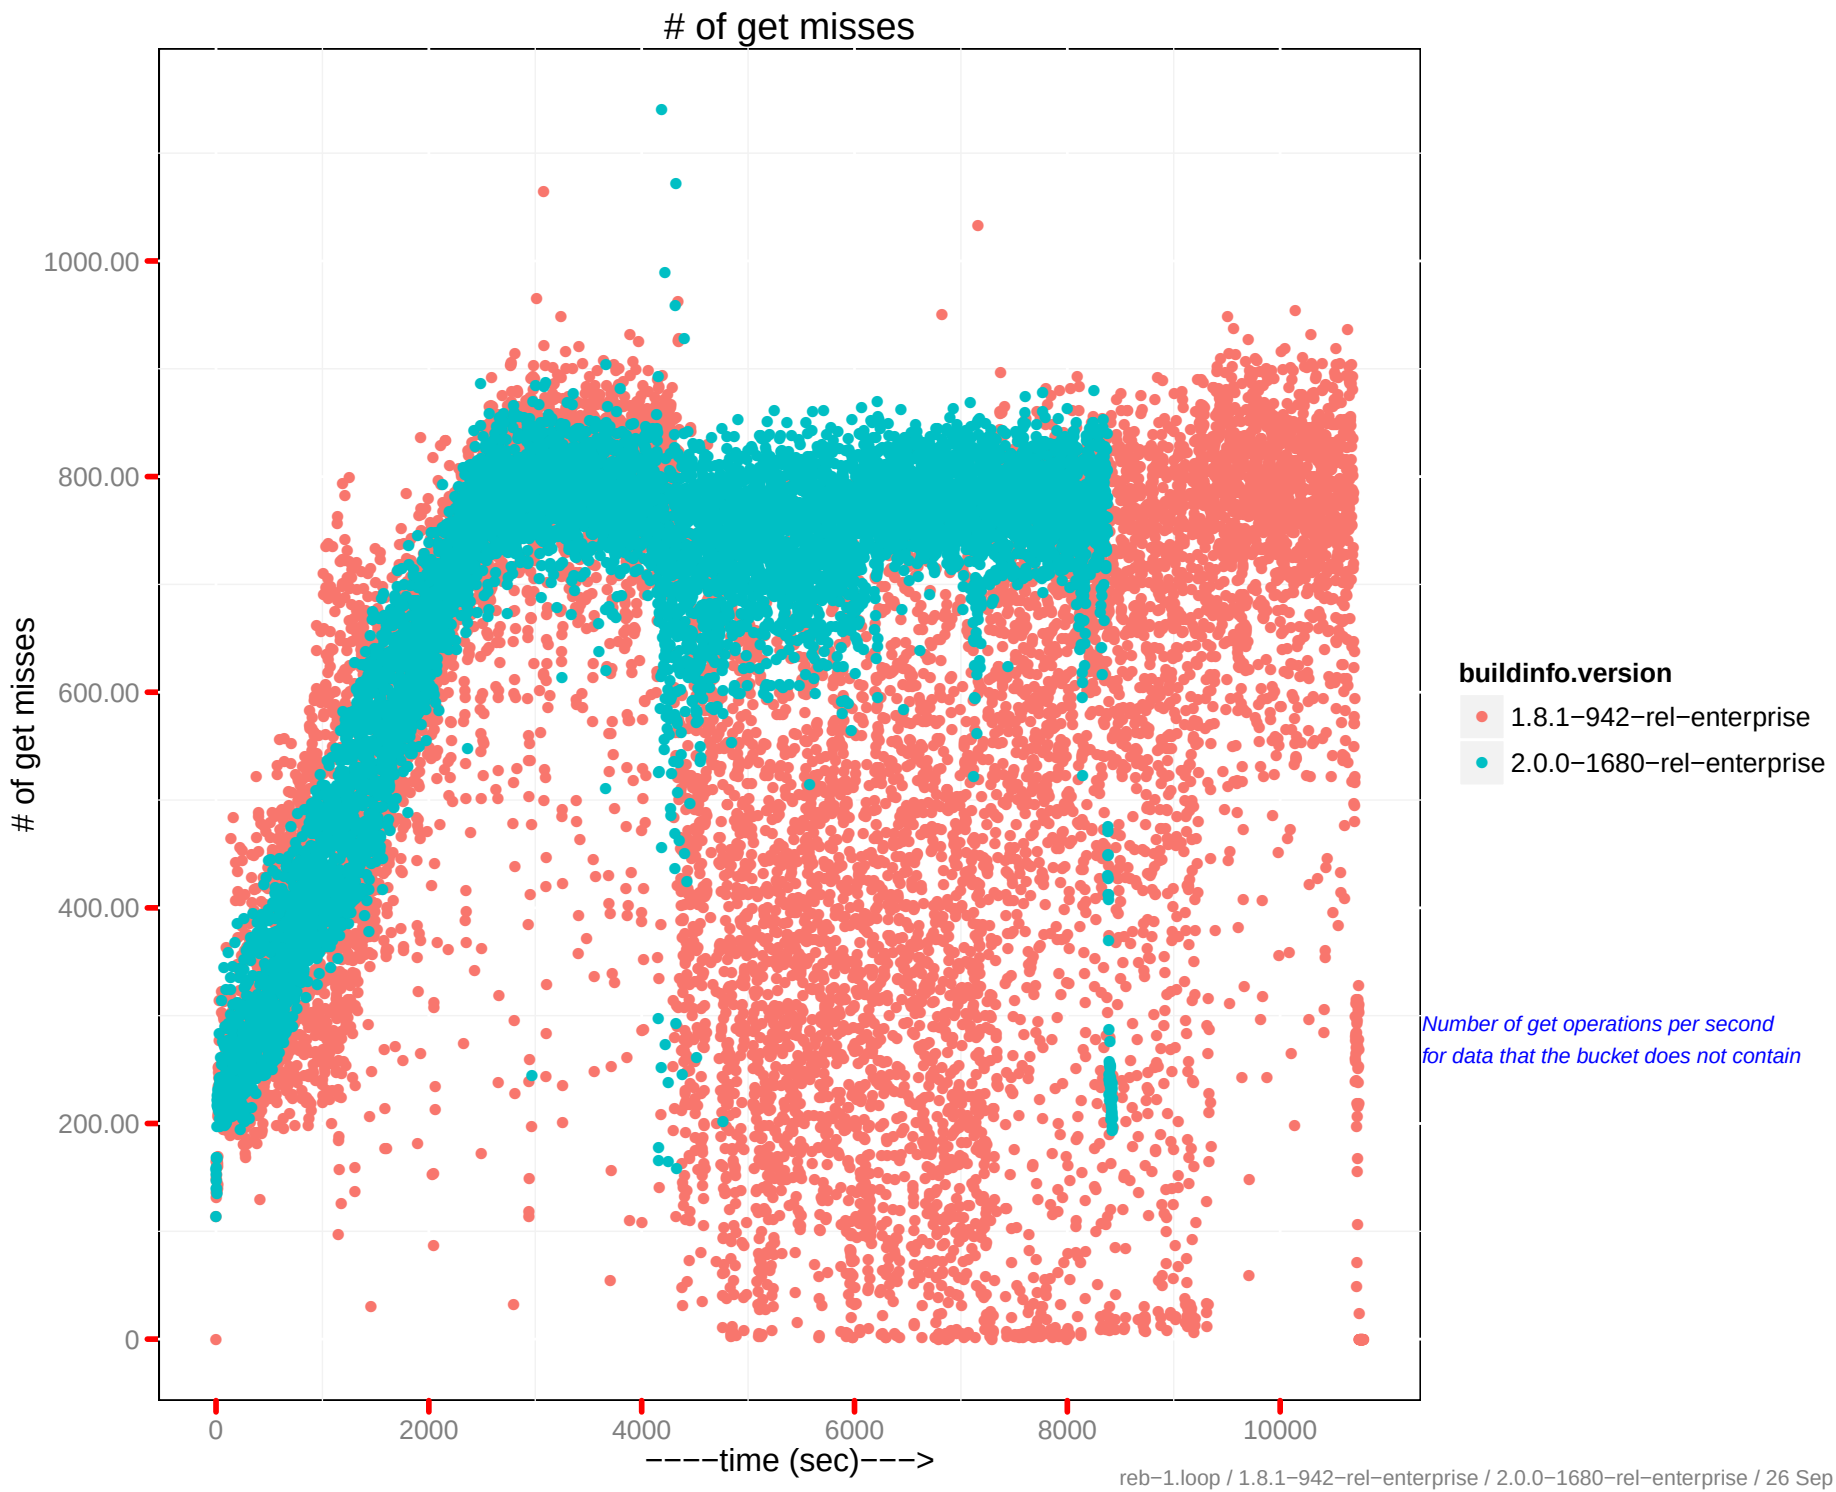
# of get misses
1000.00
800.00
600.00
400.00
200.00
0
# of get misses
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of get operations per second
for data that the bucket does not contain
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

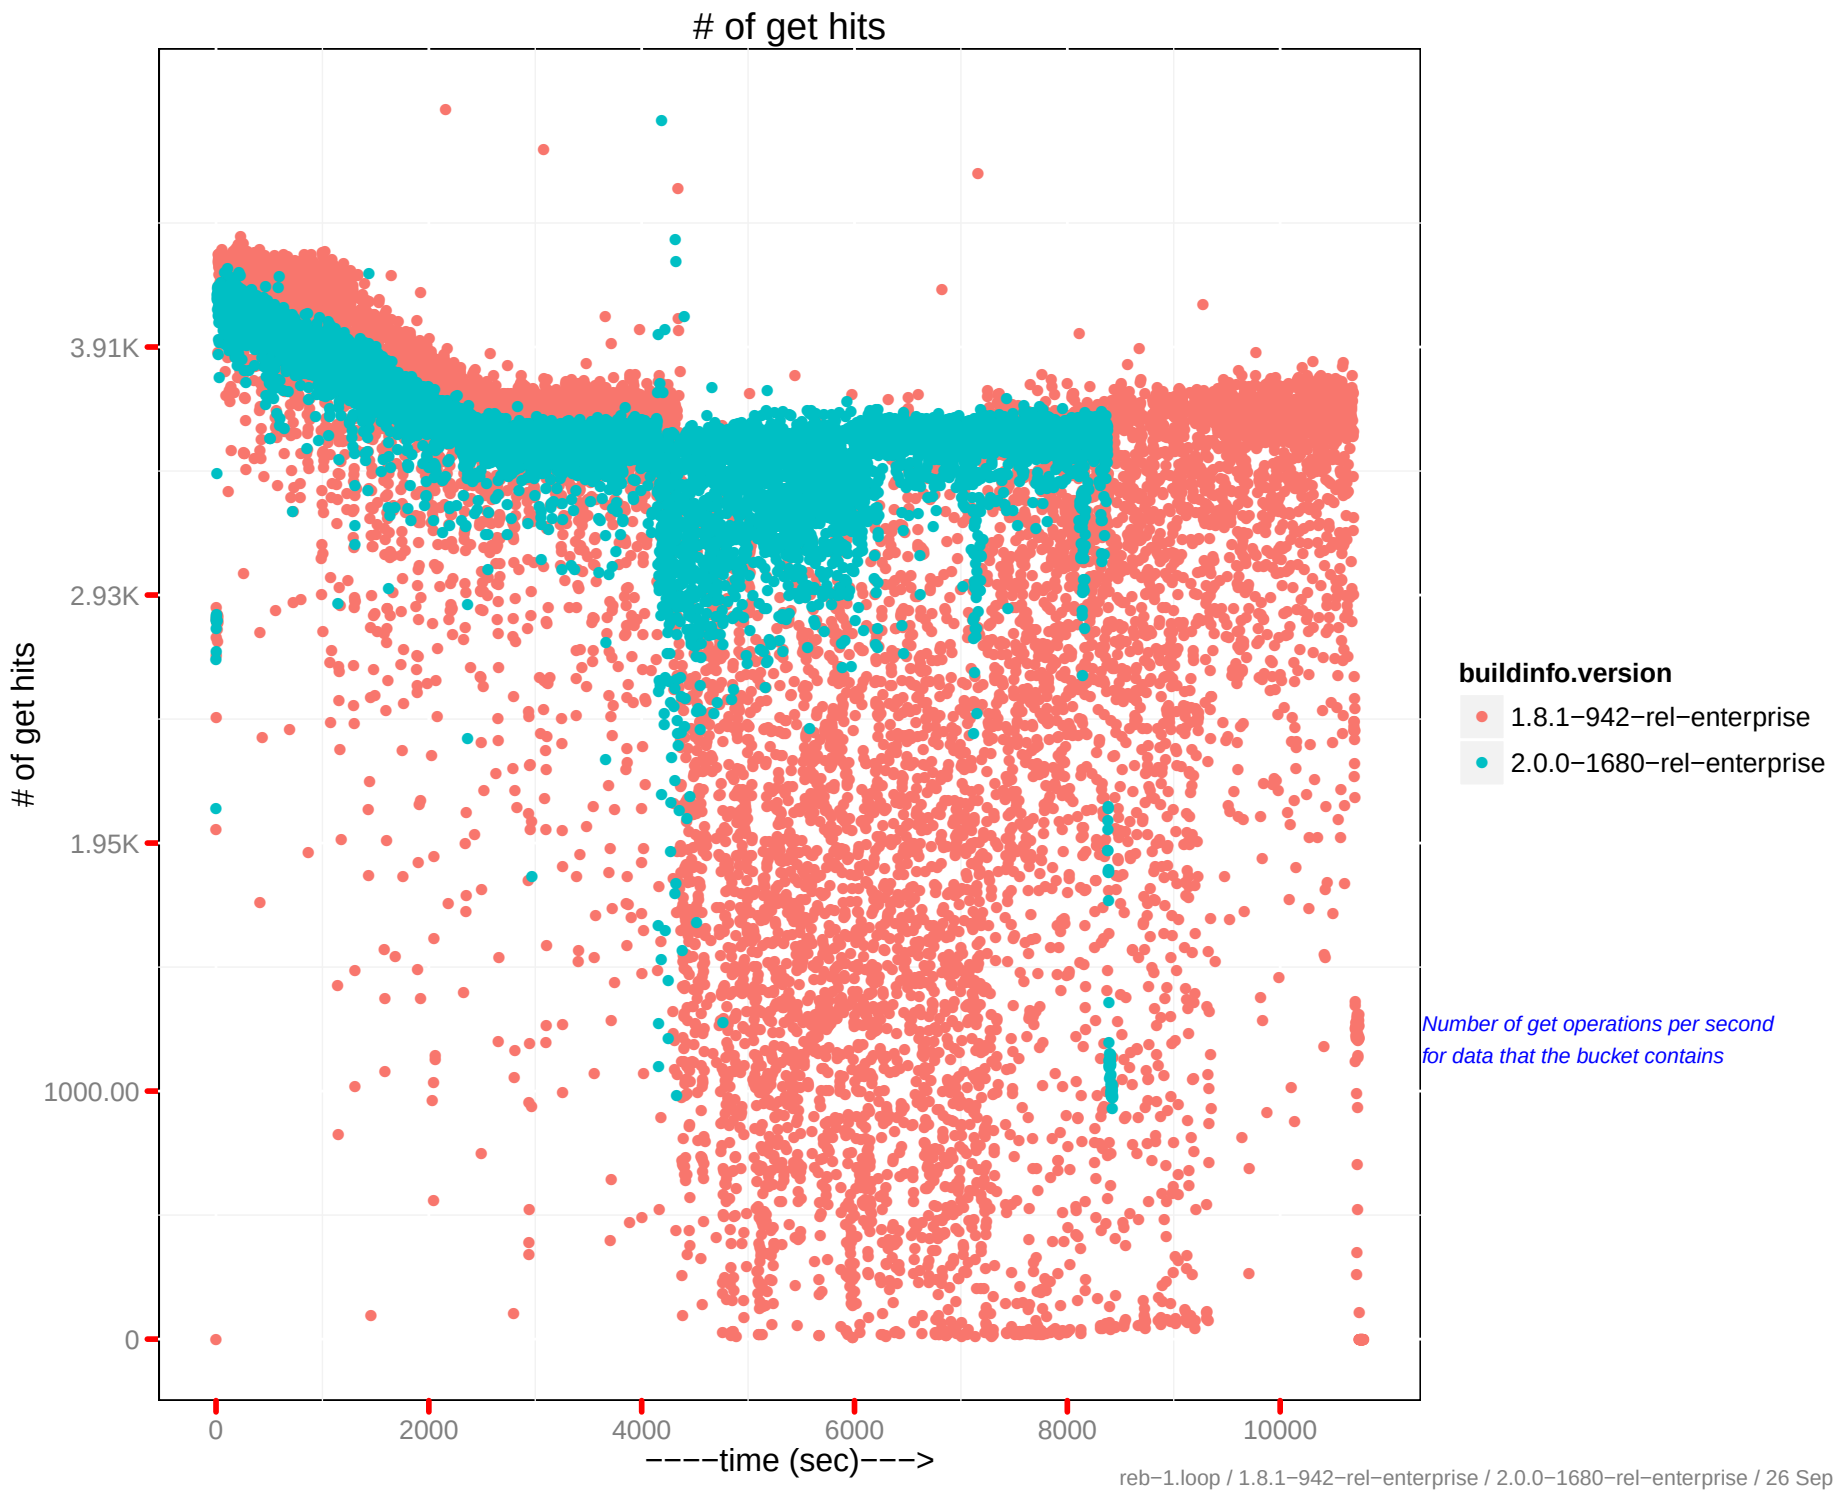
# of get hits
3.91K
2.93K
1.95K
1000.00
0
# of get hits
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of get operations per second
for data that the bucket contains
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

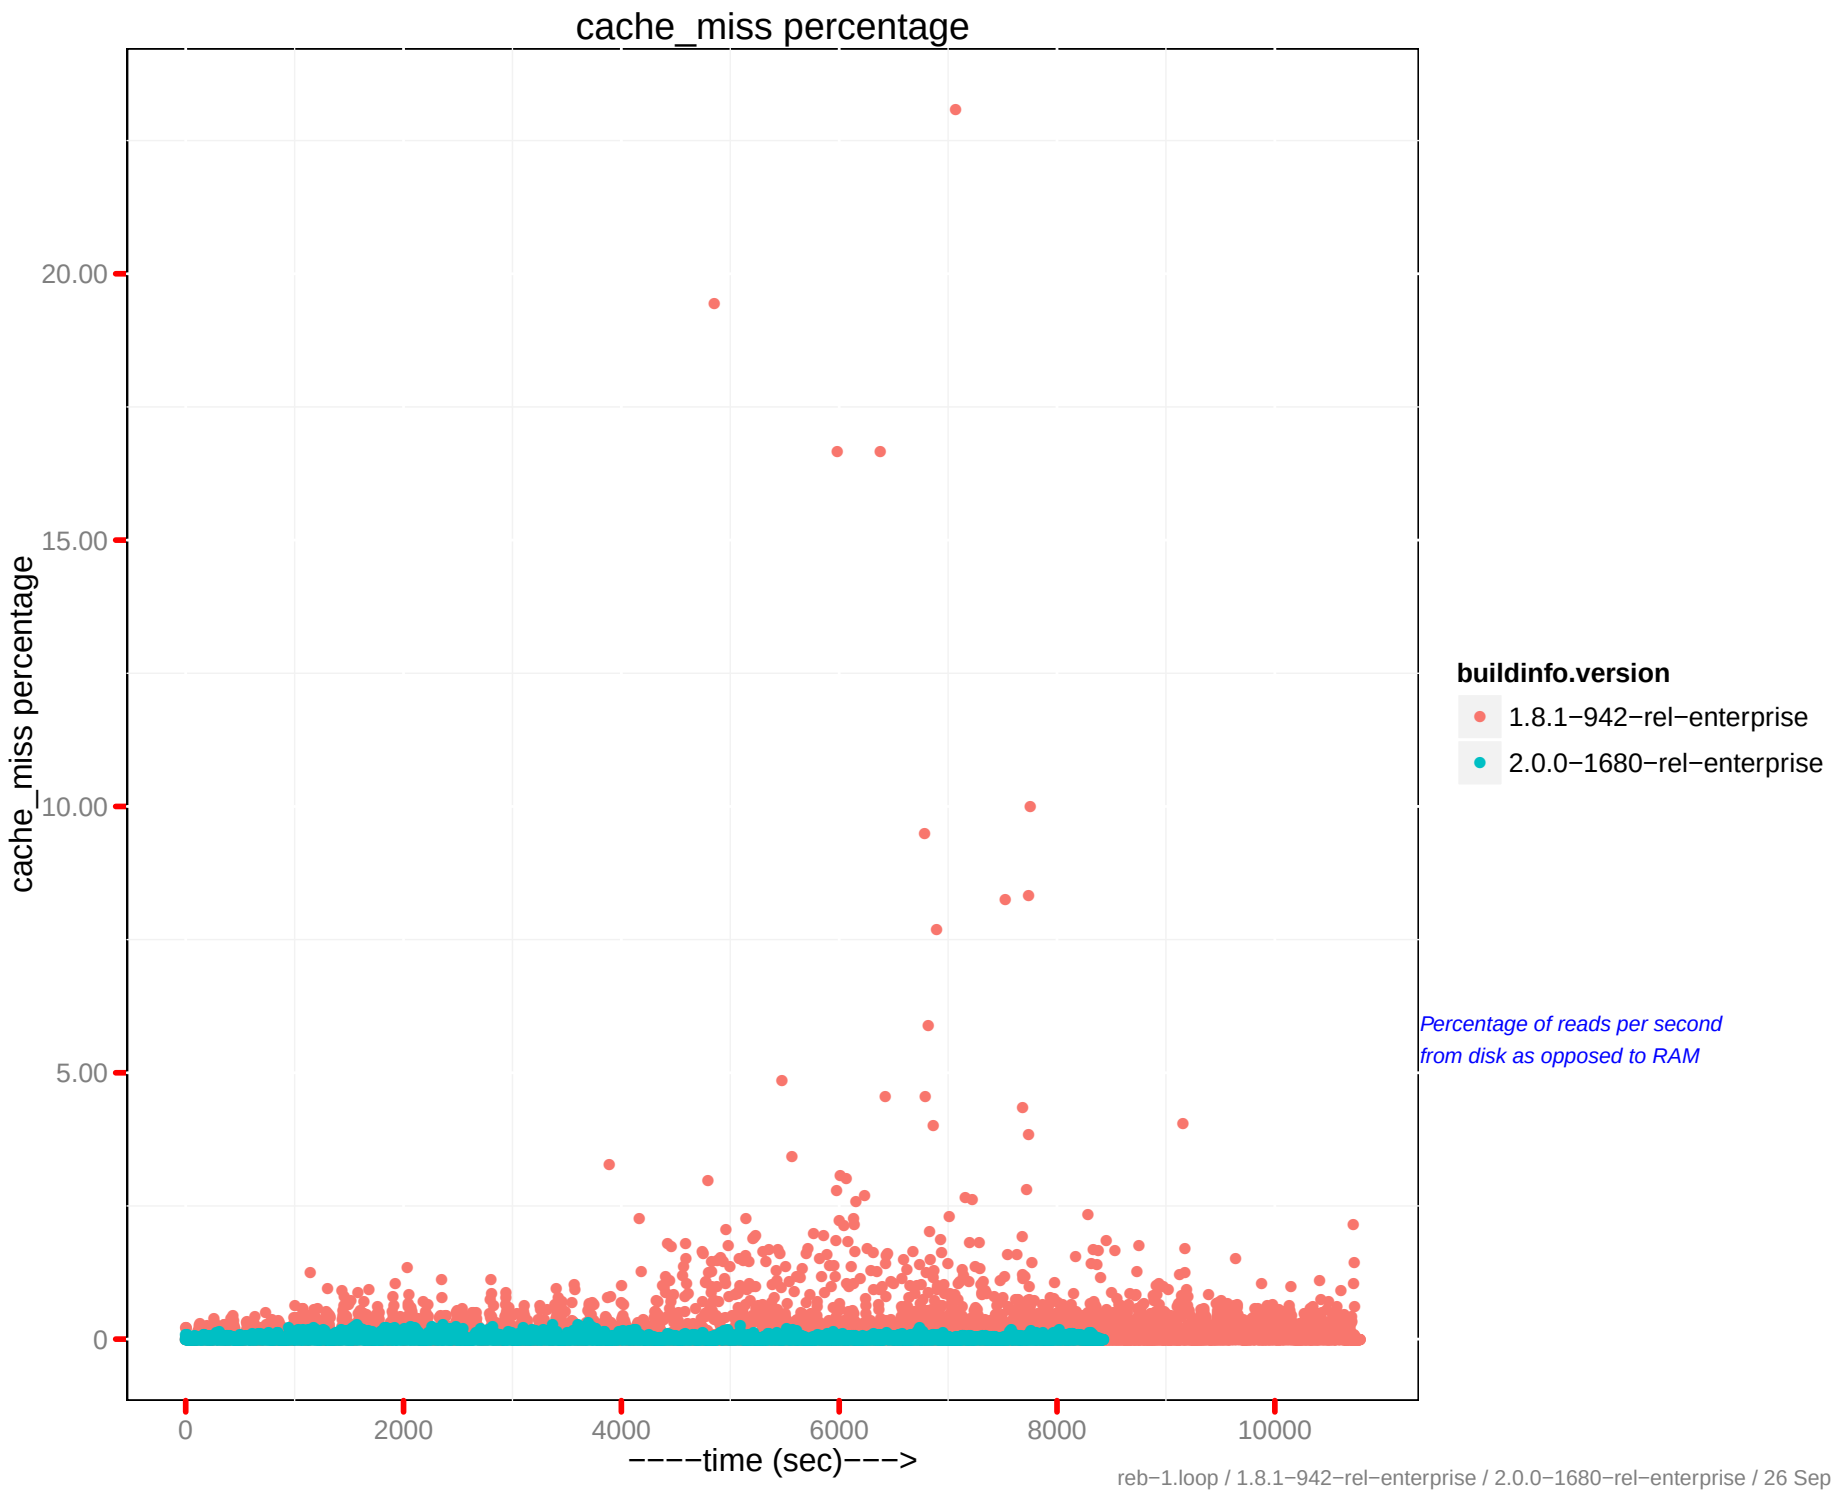
cache_miss percentage
cache_miss percentage
20.00
15.00
10.00
5.00
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Percentage of reads per second
from disk as opposed to RAM
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

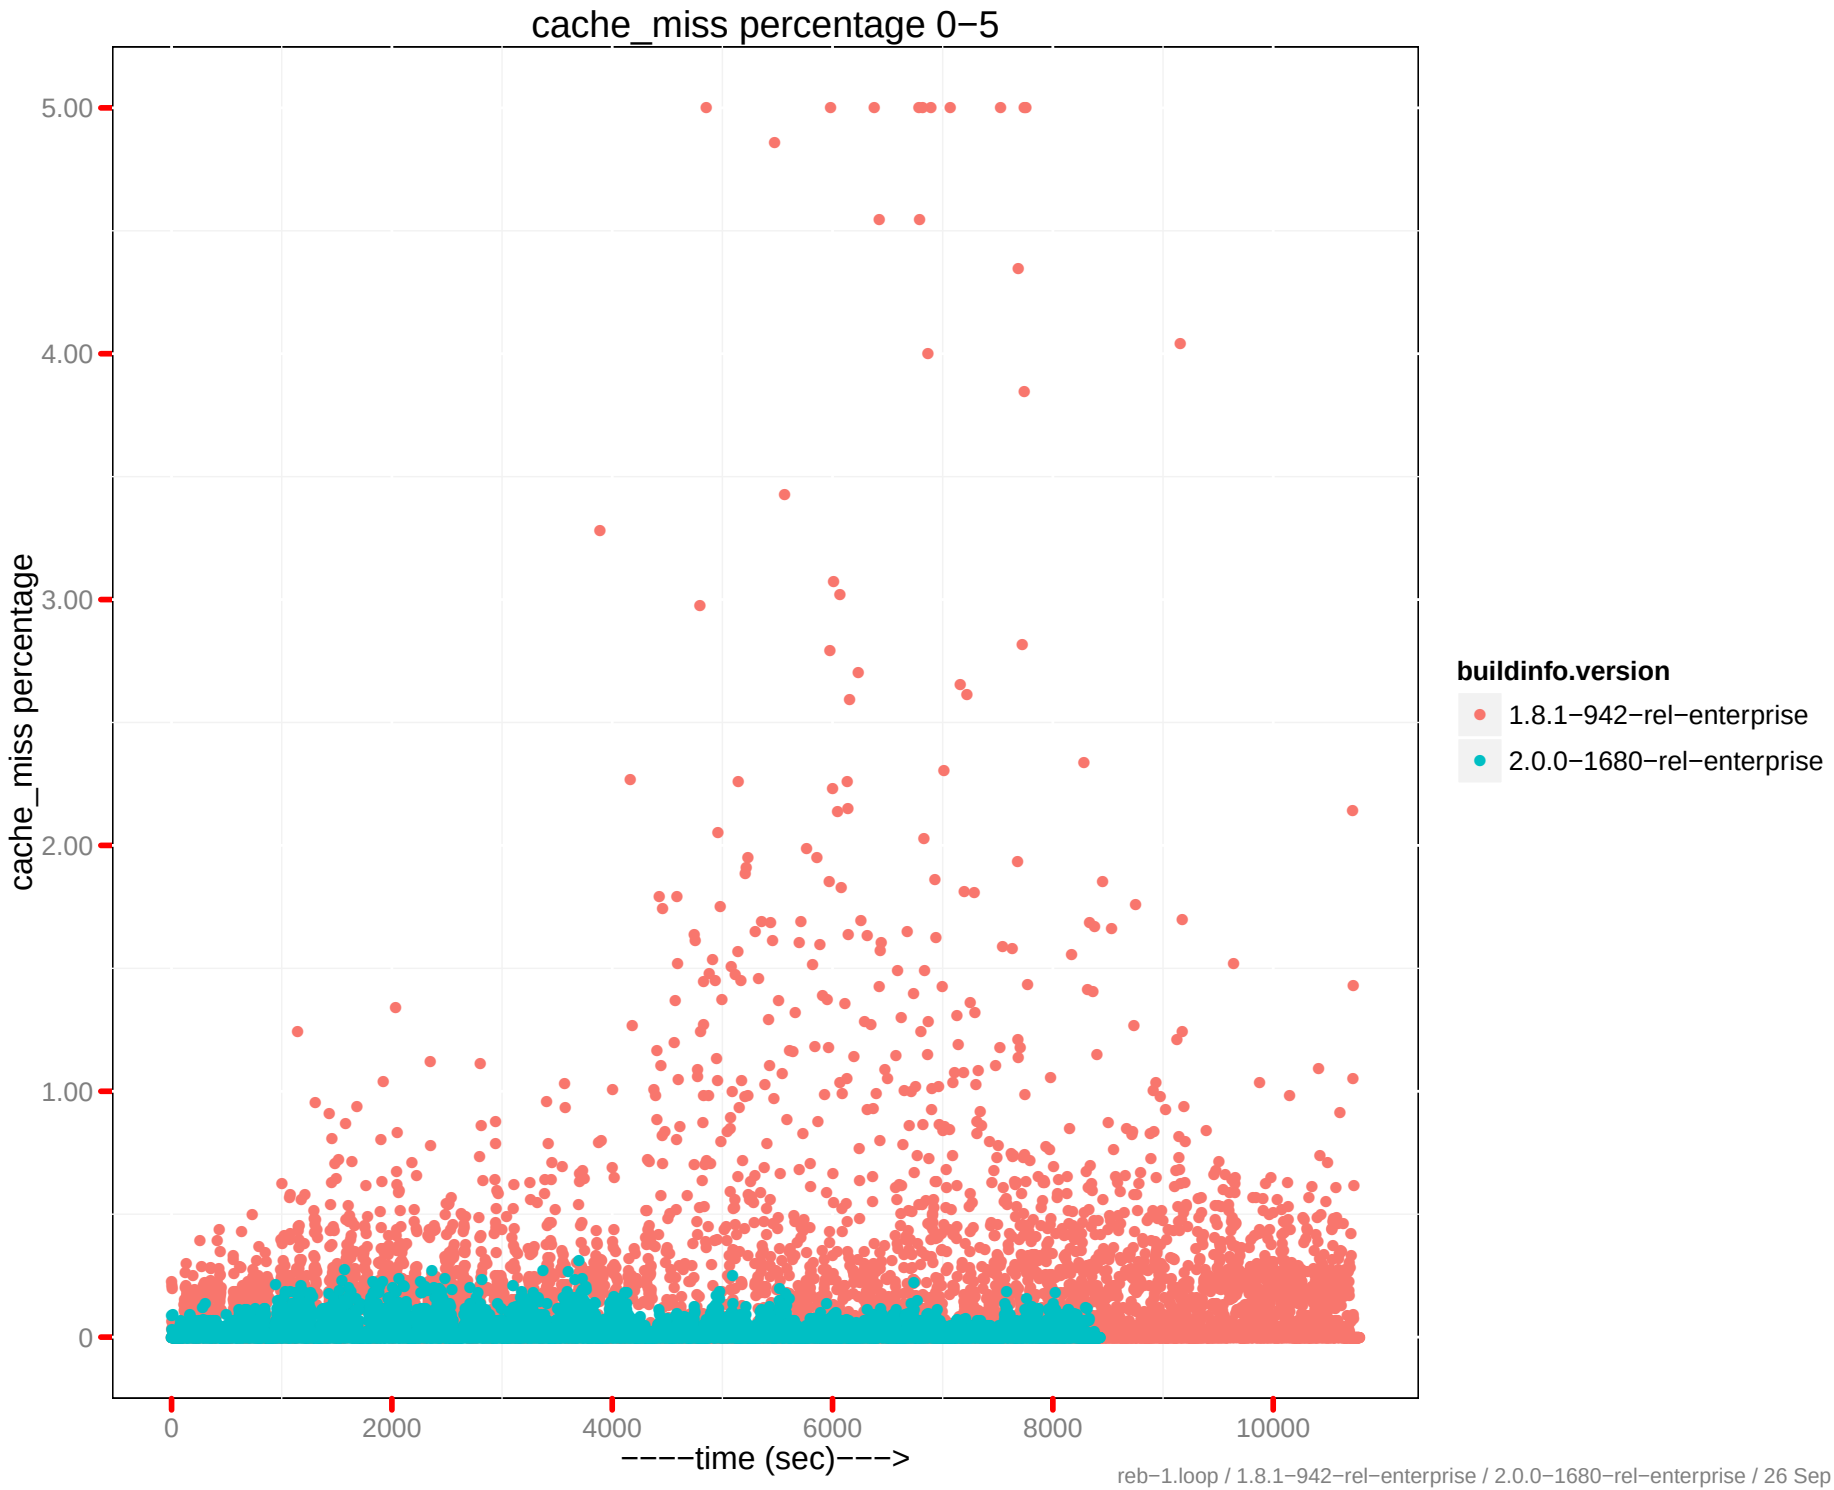
cache_miss percentage 0−5
cache_miss percentage
5.00
4.00
3.00
2.00
1.00
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

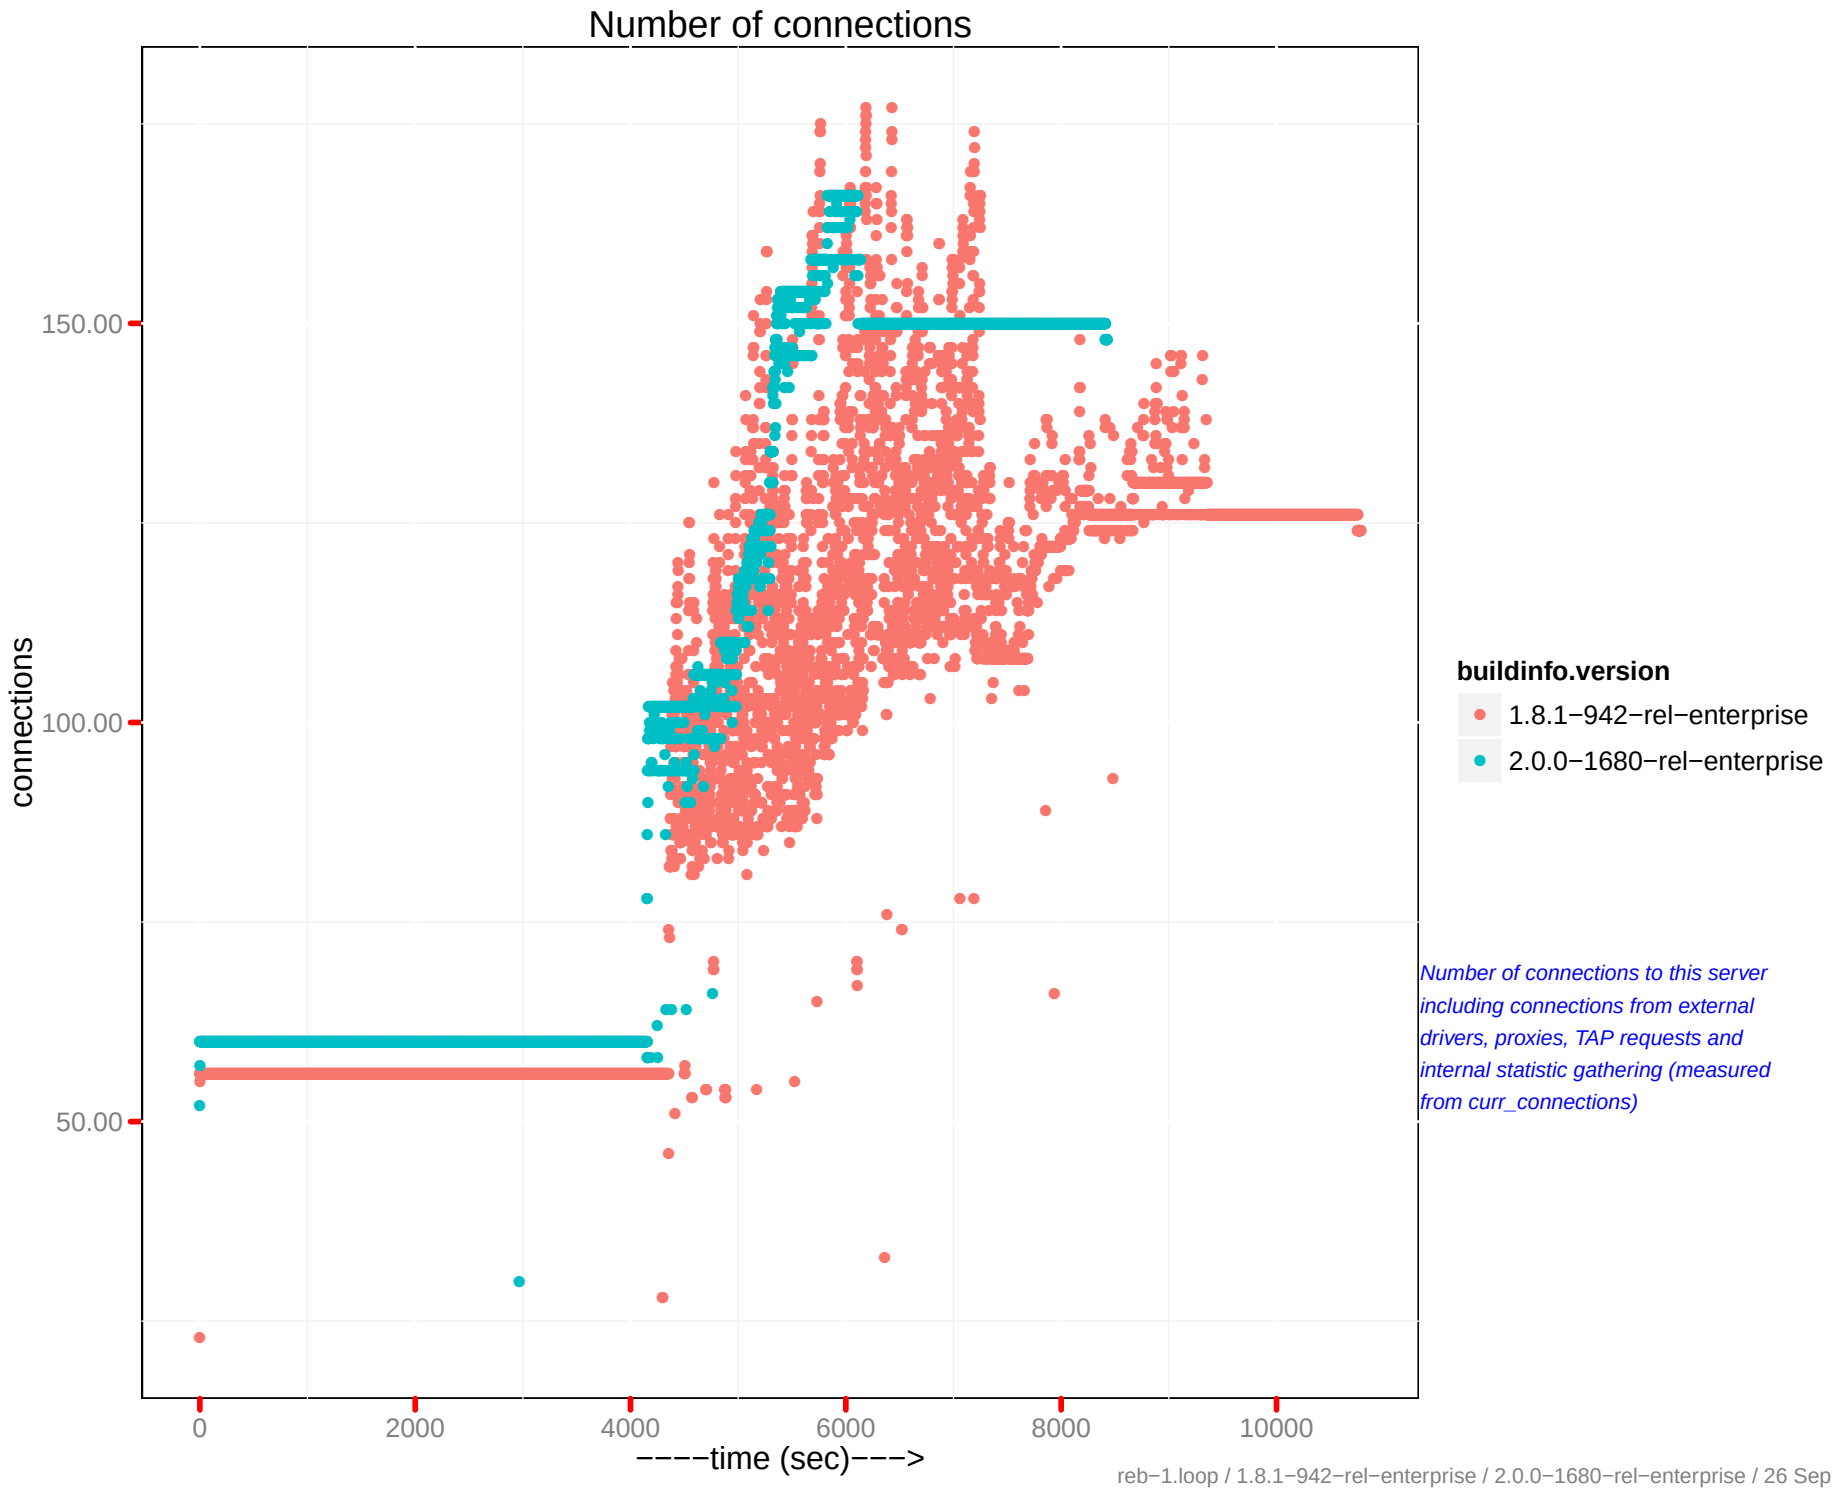
Number of connections
connections
150.00
100.00
50.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of connections to this server including connections from external drivers, proxies, TAP requests and internal statistic gathering (measured from curr_connections)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

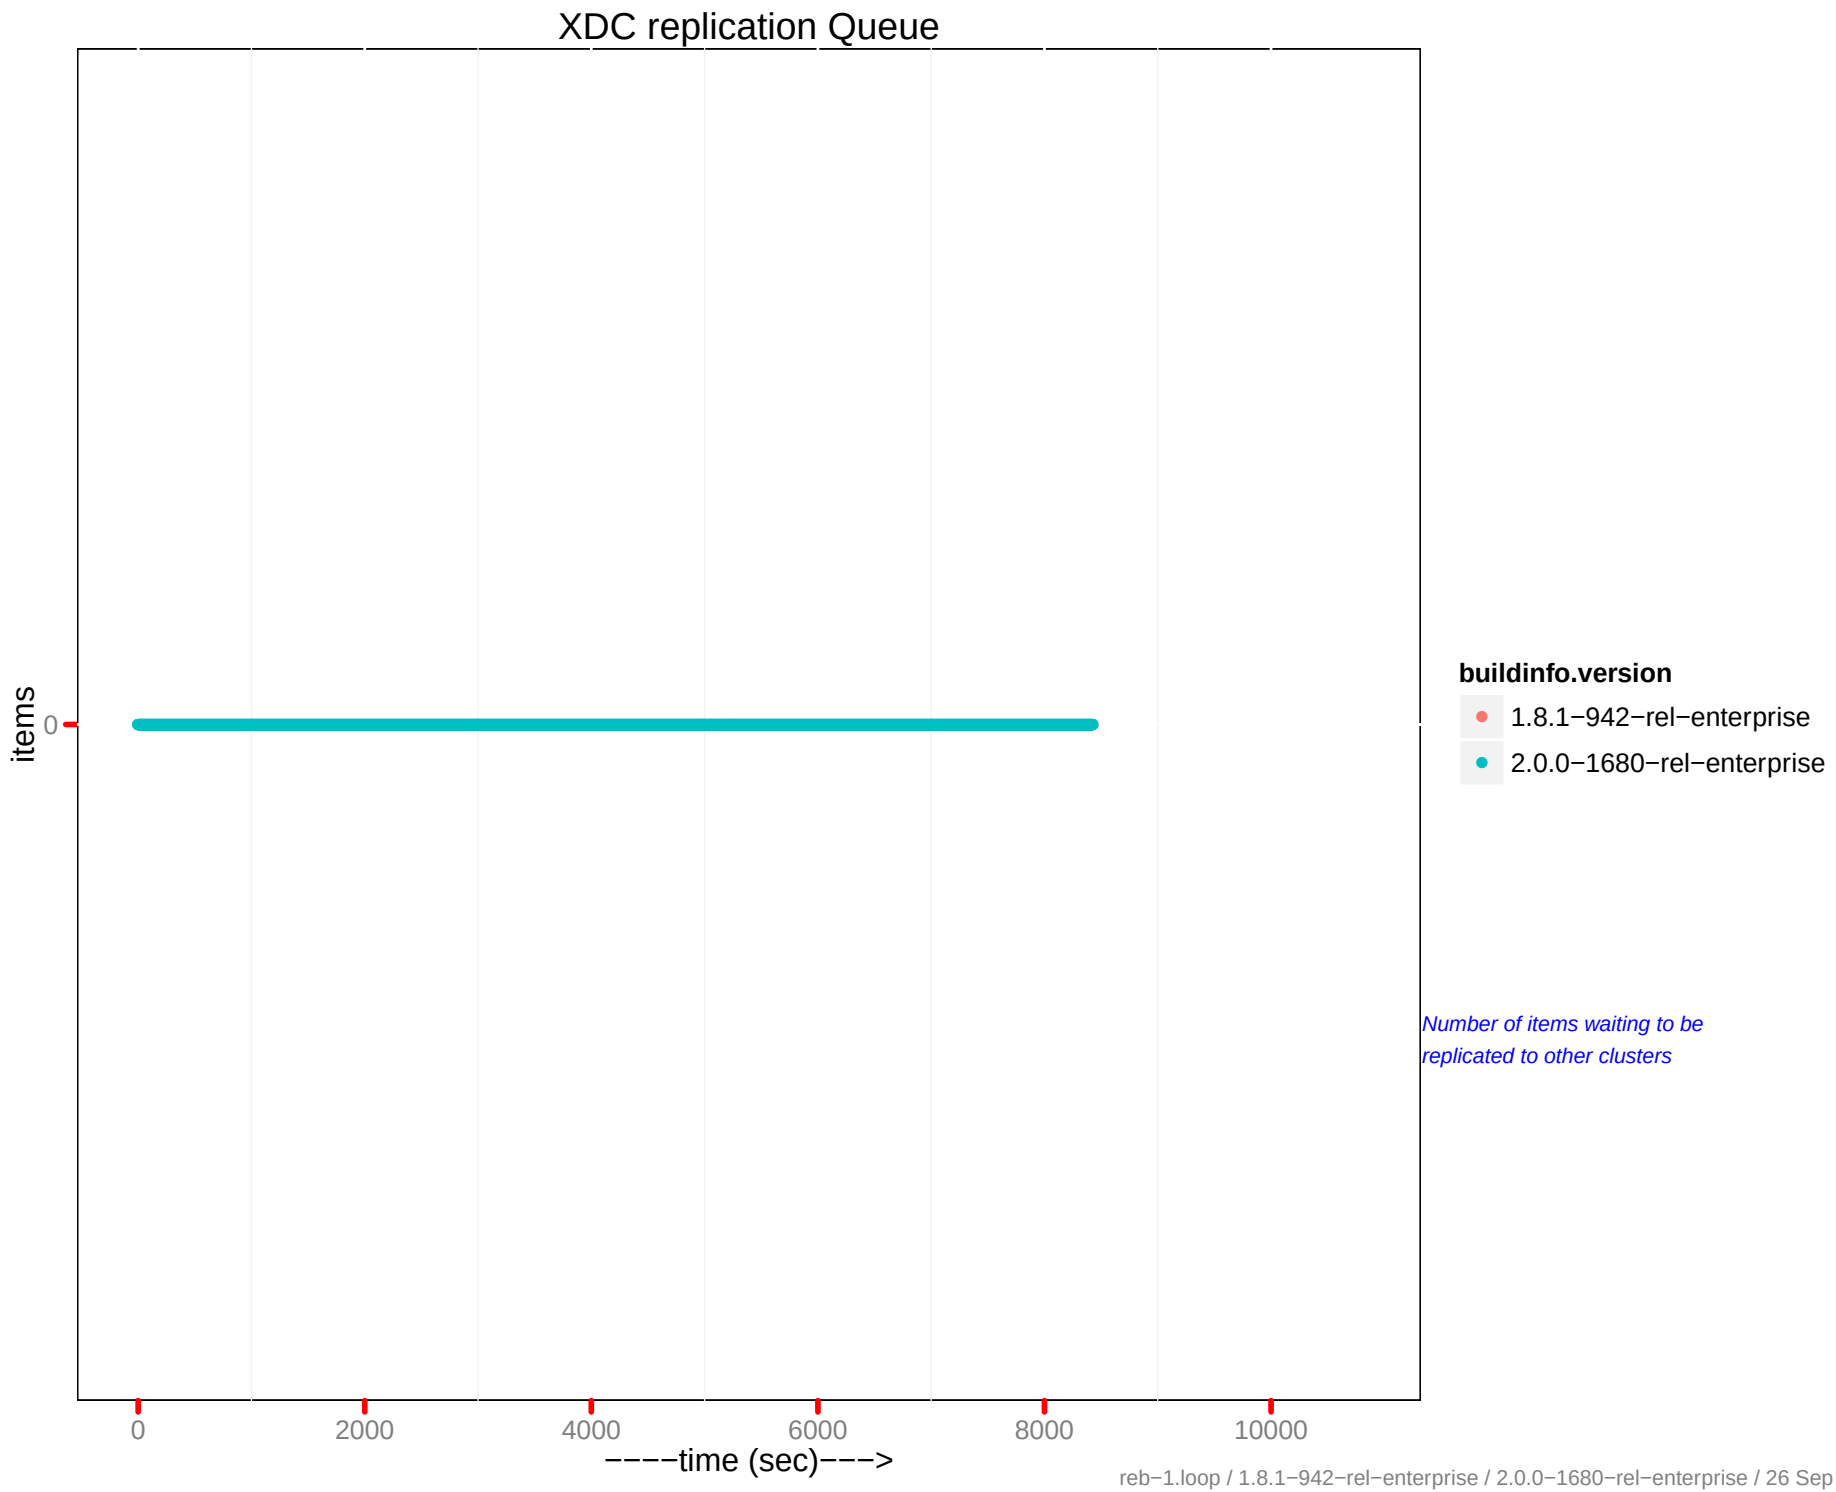
XDC replication Queue
items
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of items waiting to be replicated to other clusters
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

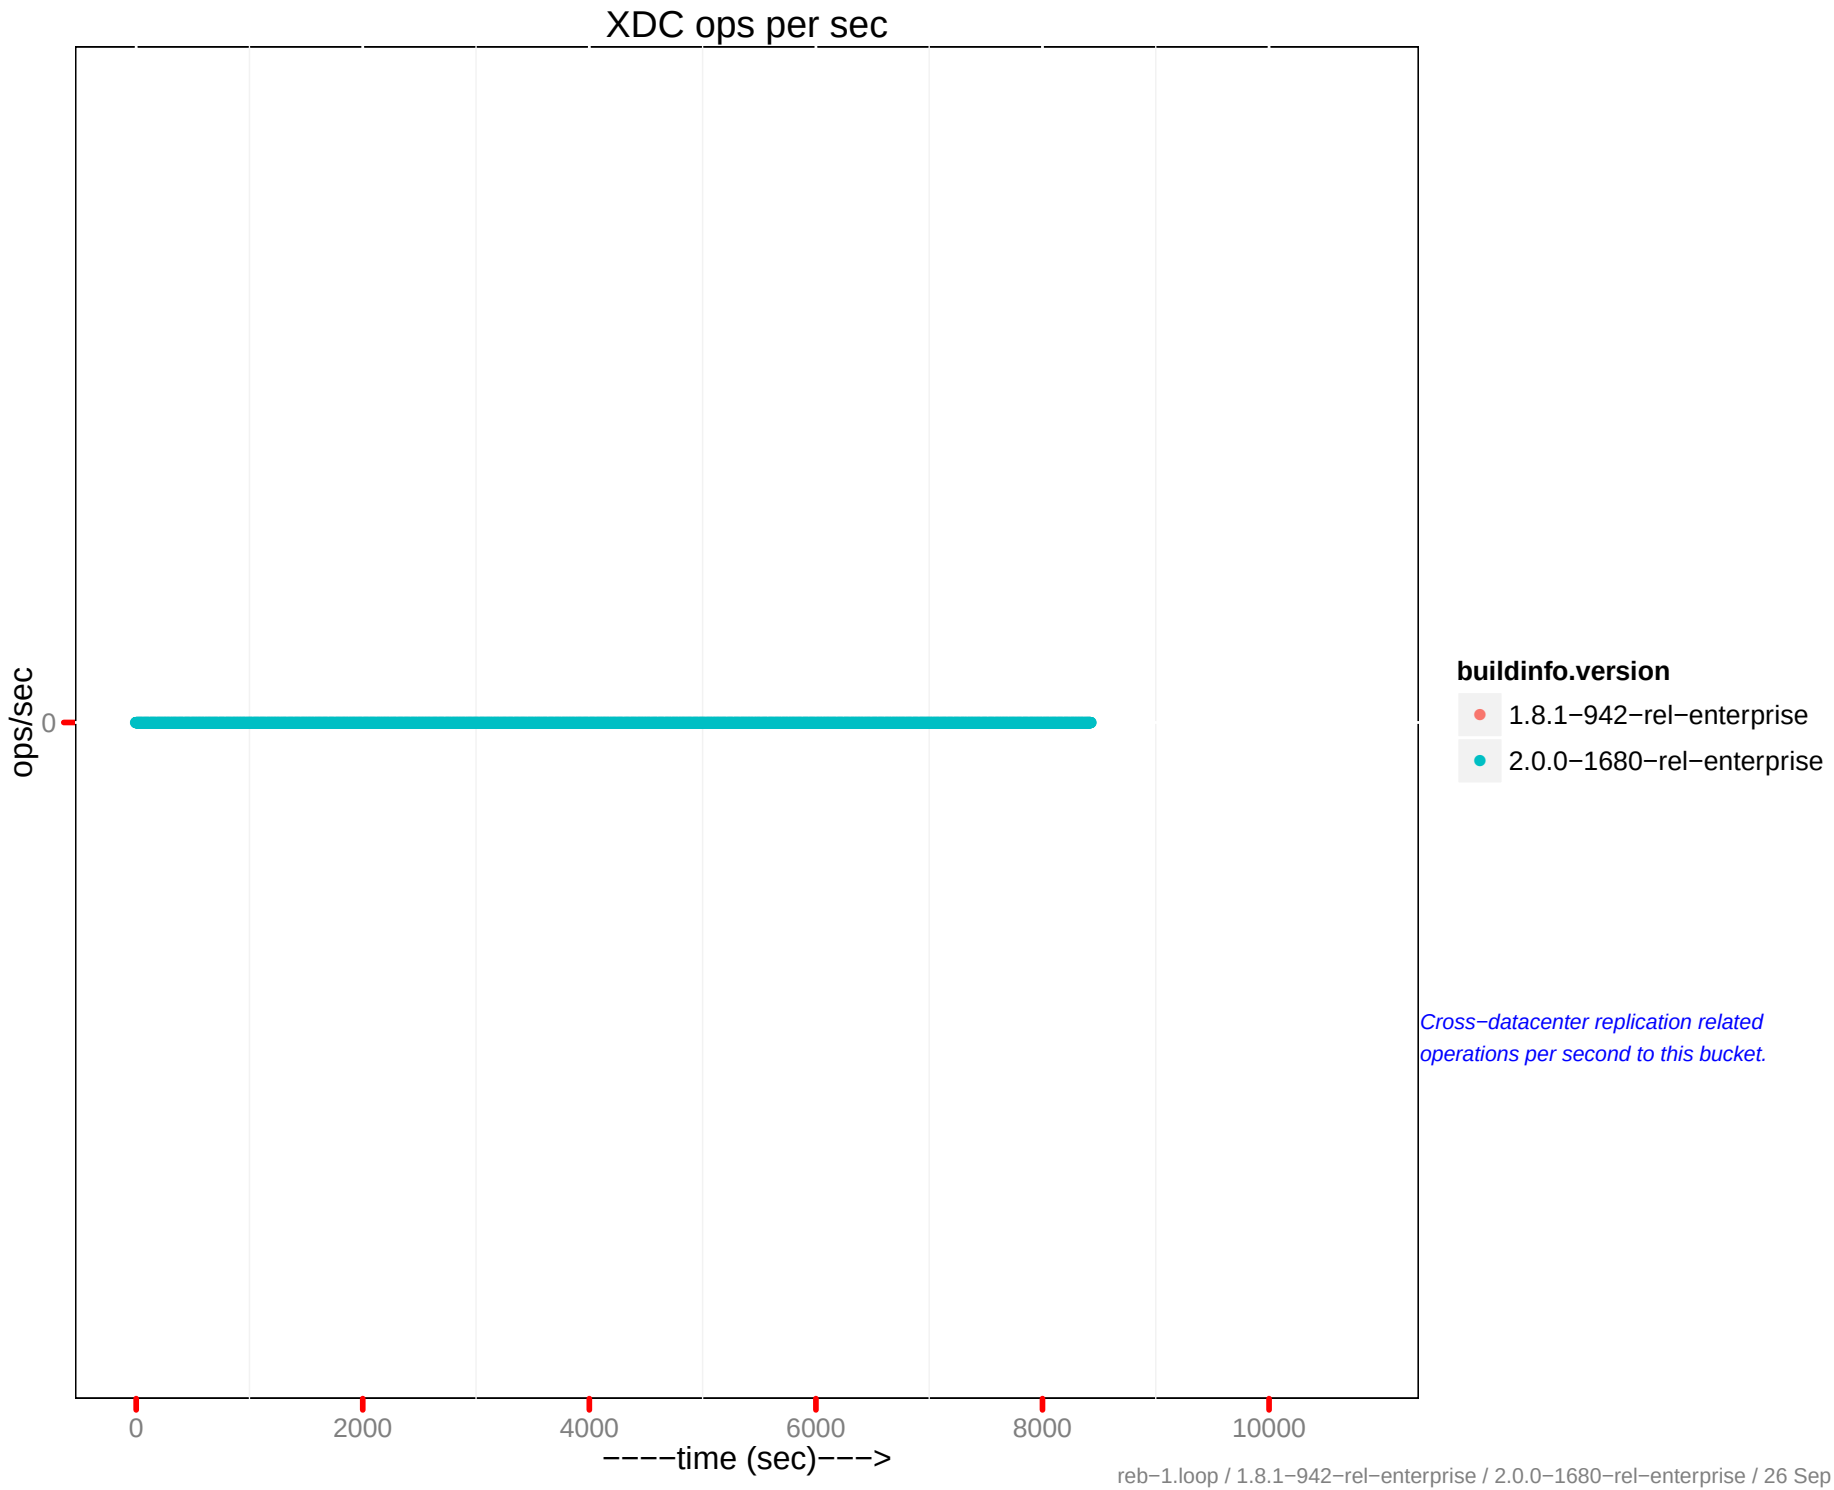
XDC ops per sec
ops/sec
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Cross−datacenter replication related
operations per second to this bucket.
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

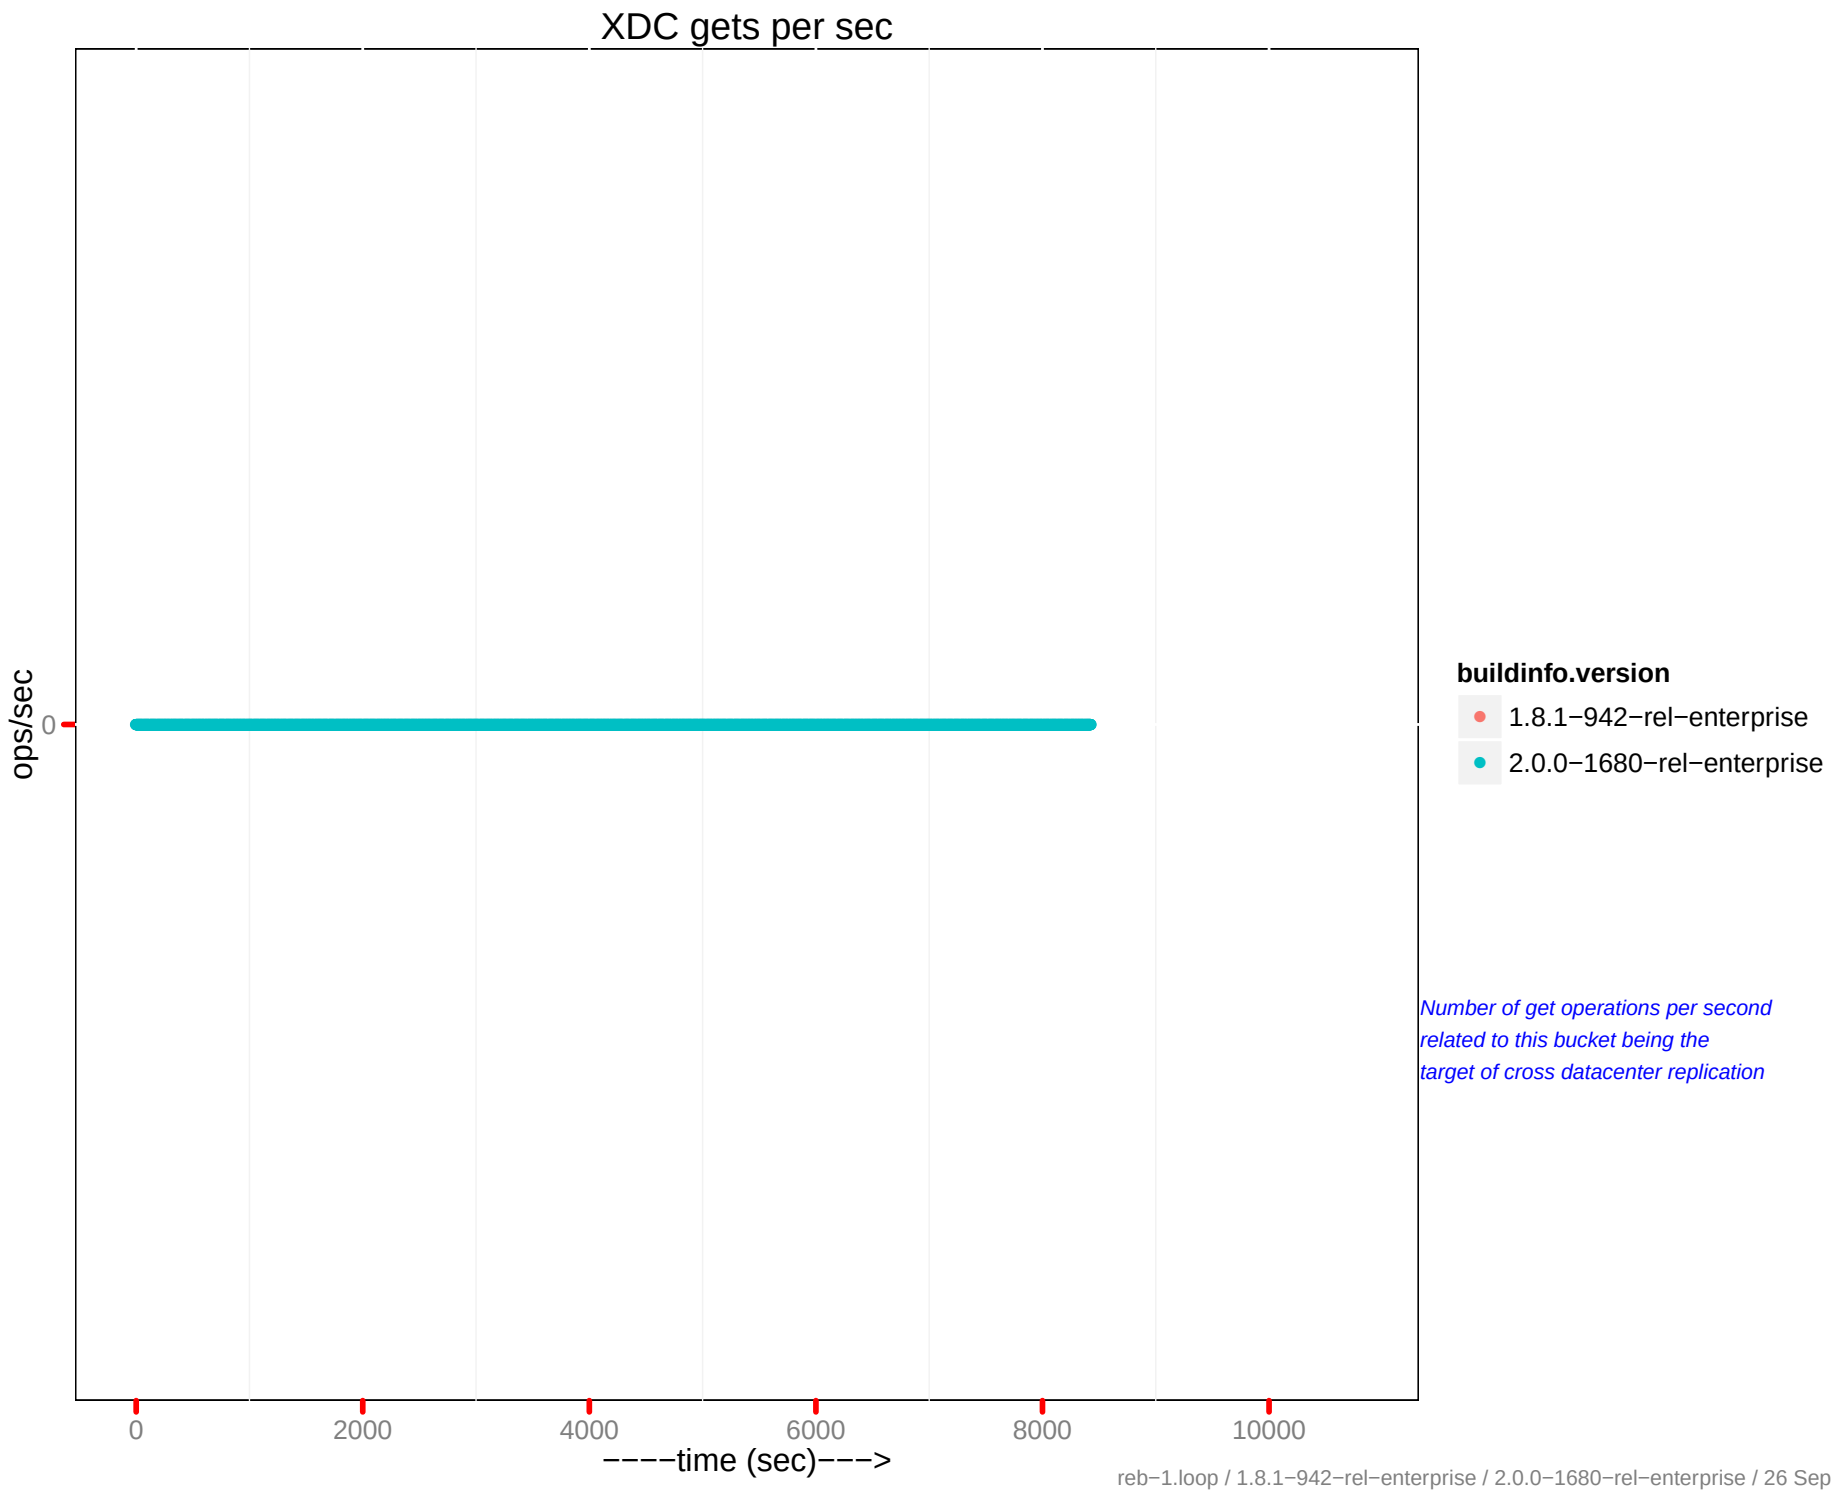
XDC gets per sec
ops/sec
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of get operations per second related to this bucket being the target of cross datacenter replication
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

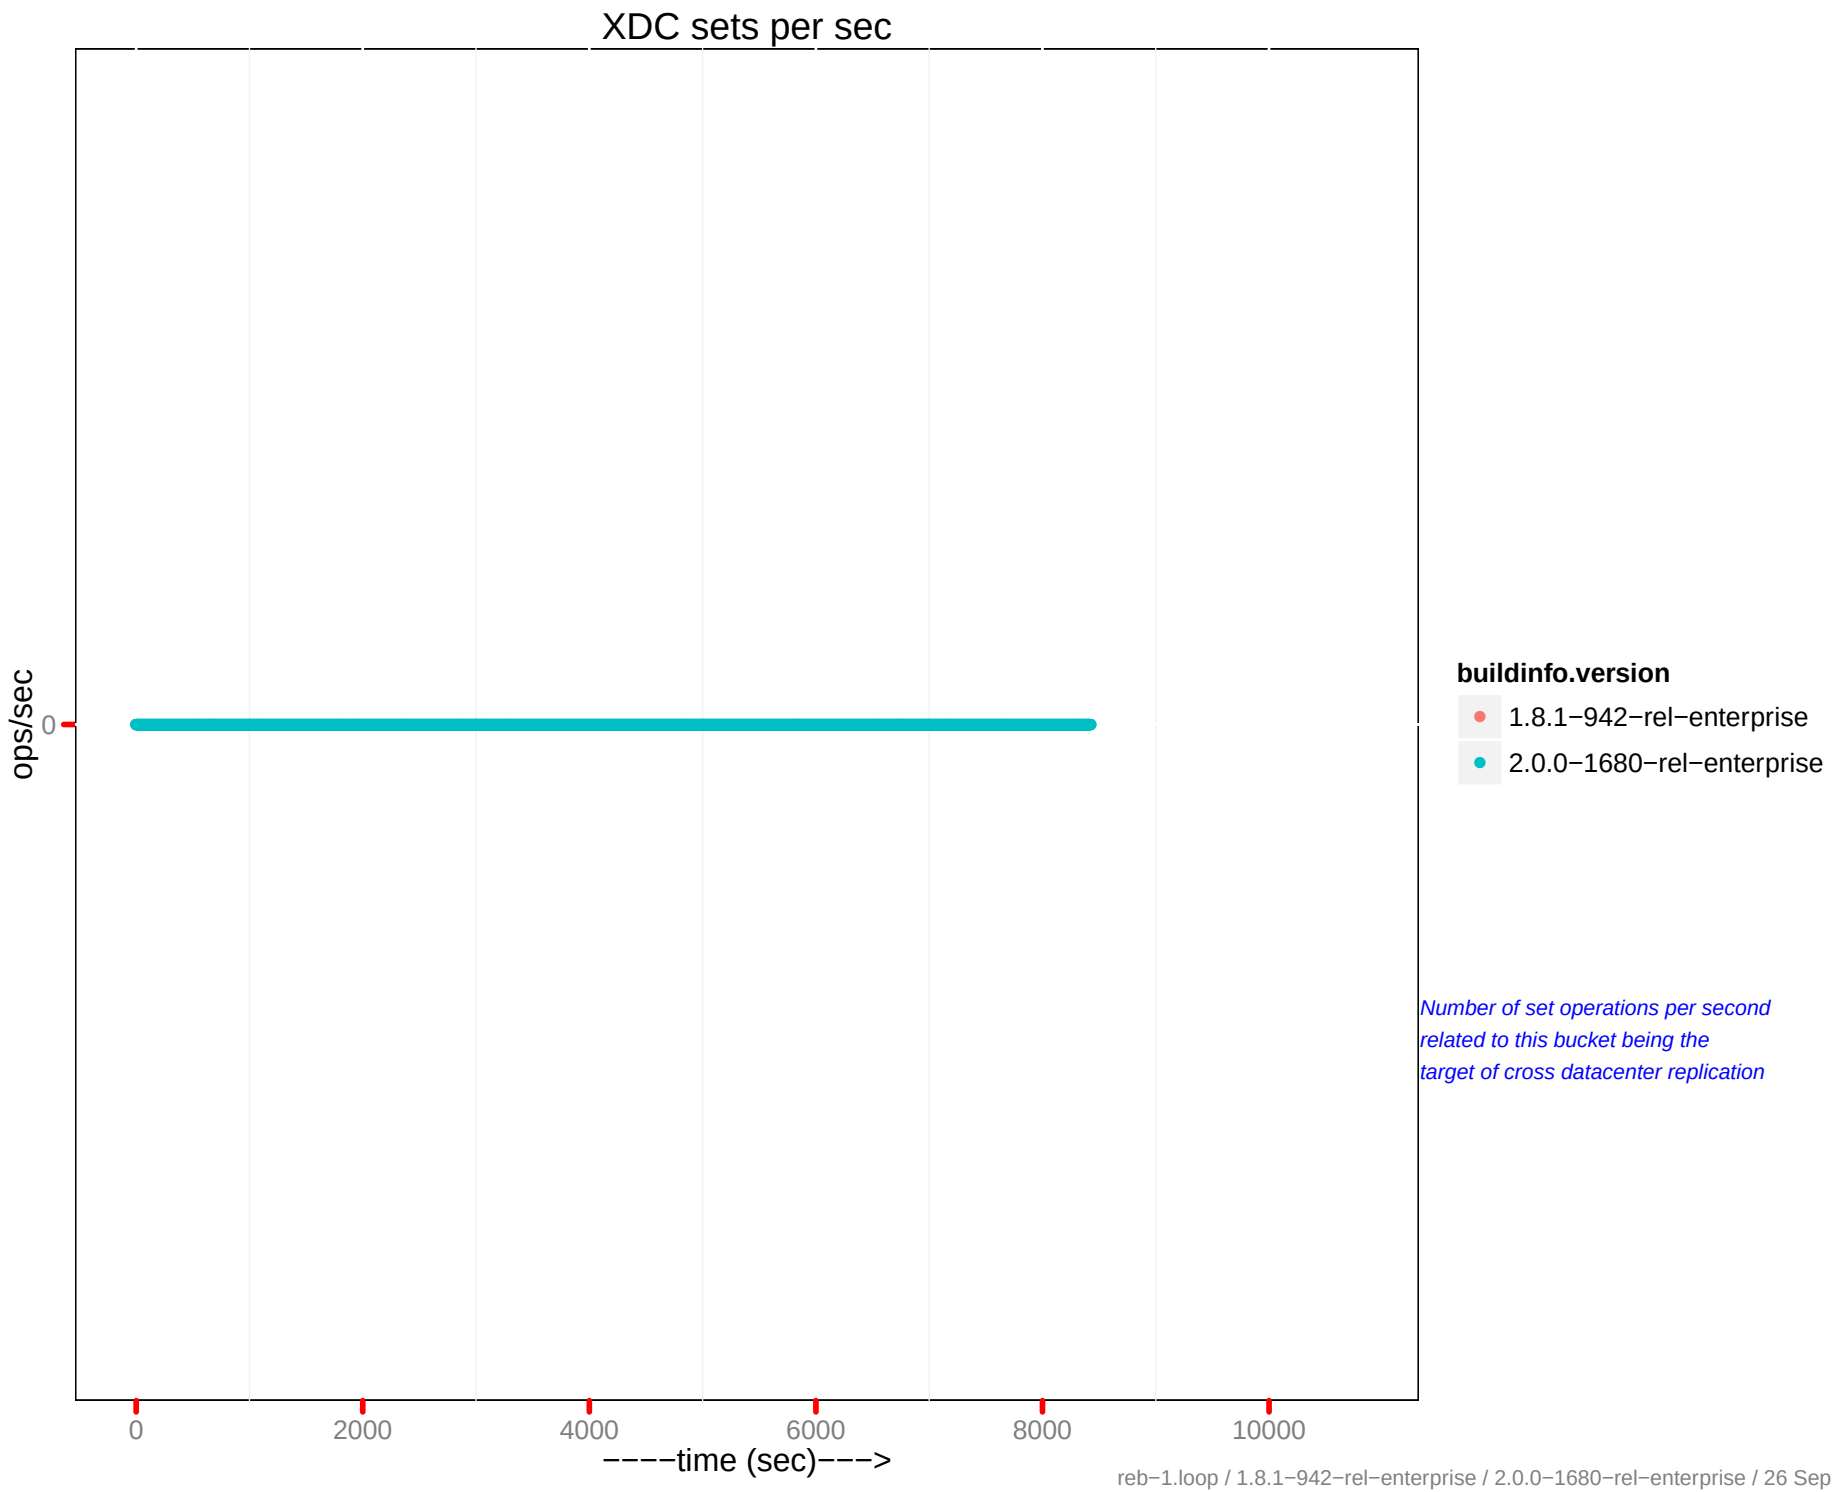
XDC sets per sec
ops/sec
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Number of set operations per second
related to this bucket being the
target of cross datacenter replication
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

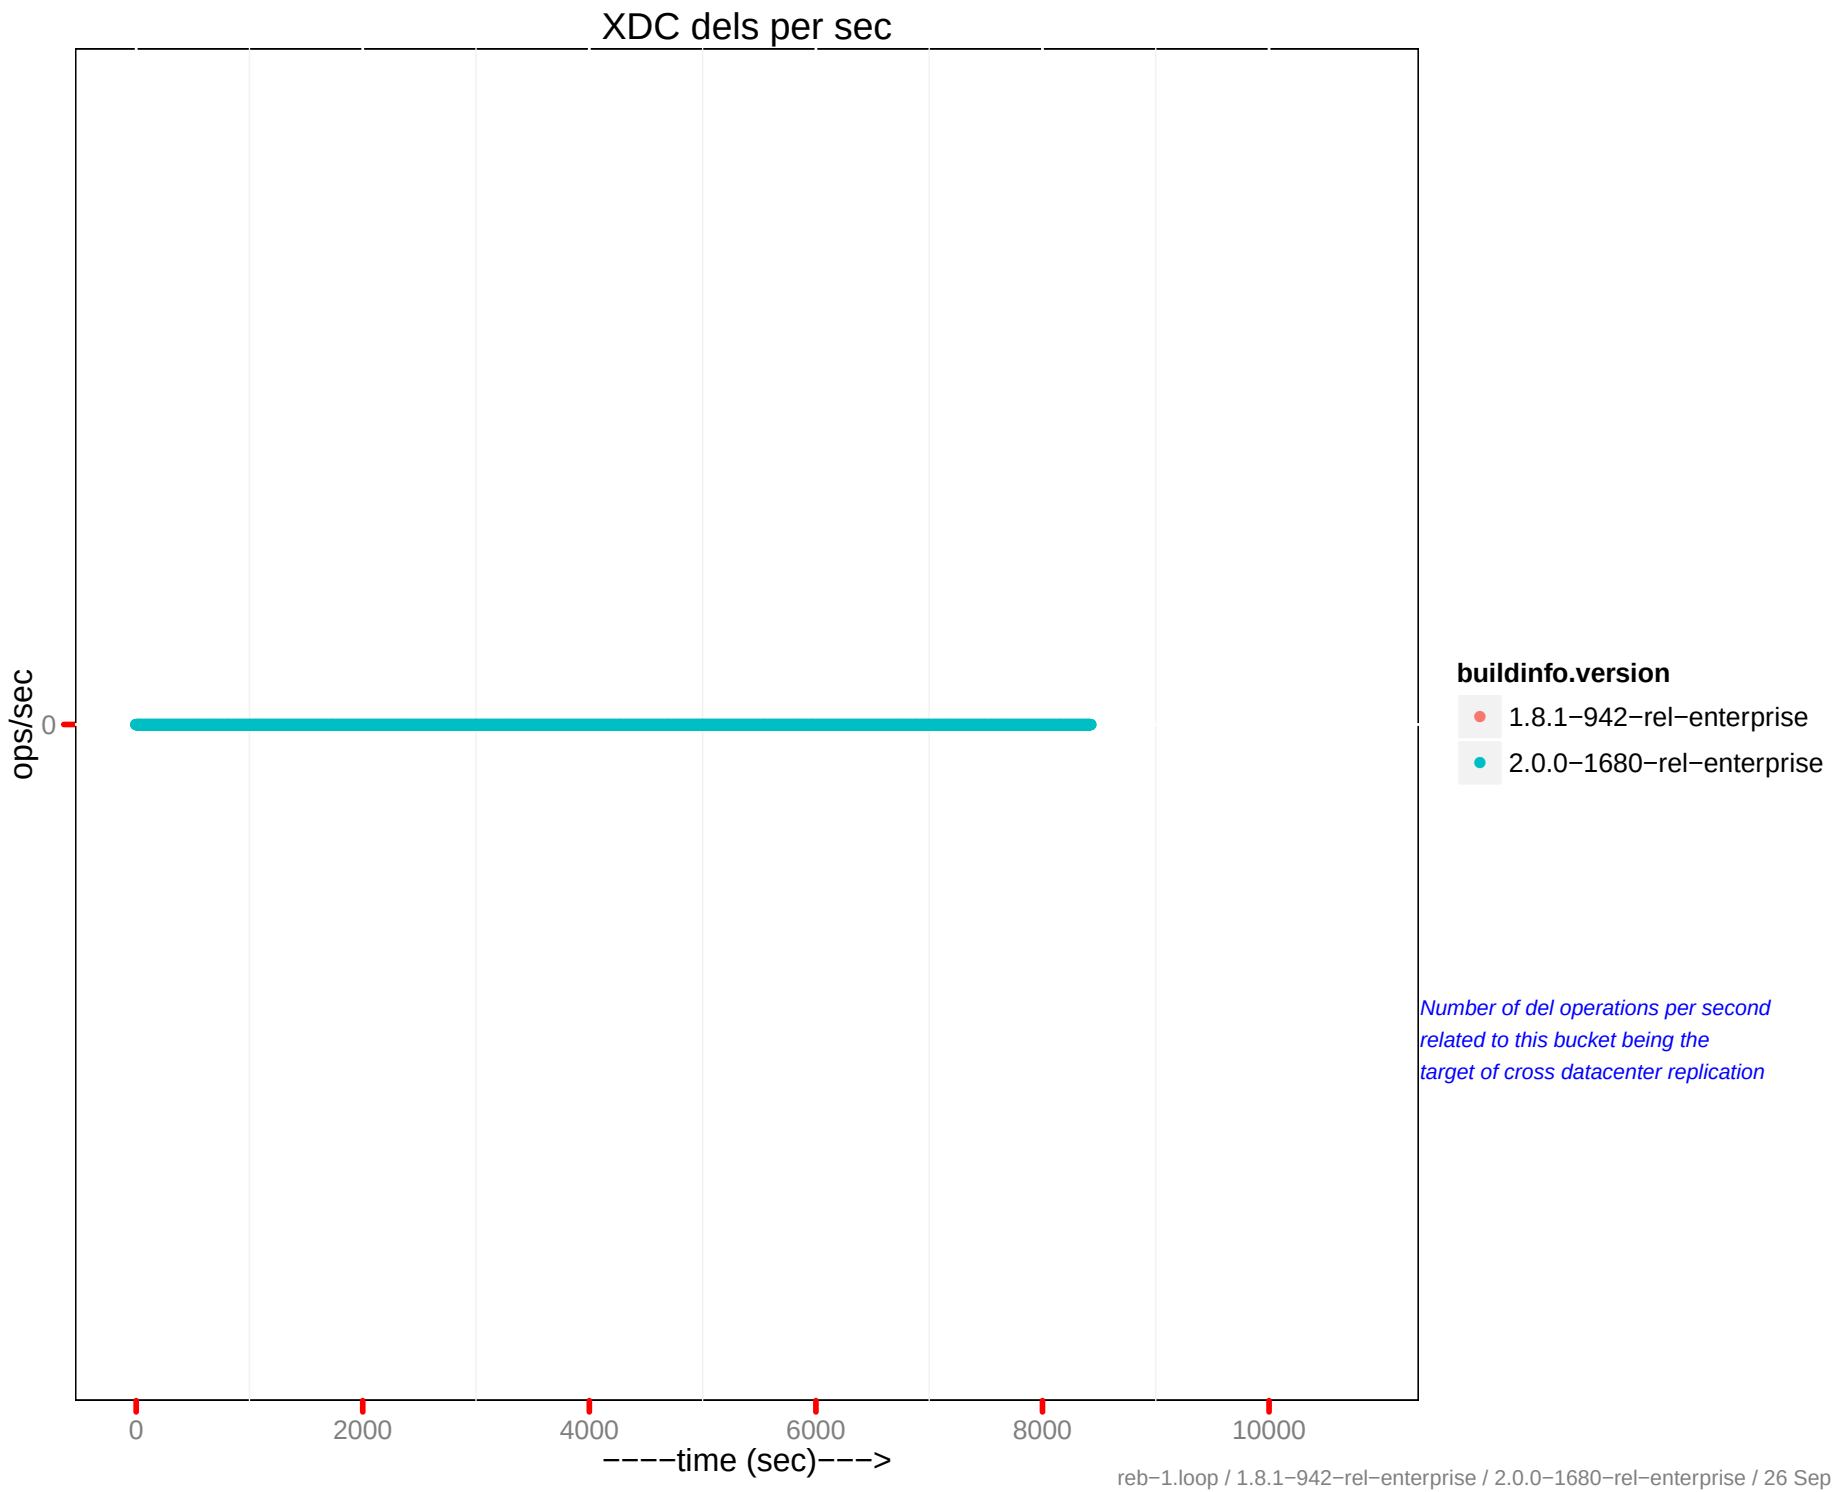
XDC dels per sec
ops/sec
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Number of del operations per second
related to this bucket being the
target of cross datacenter replication
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

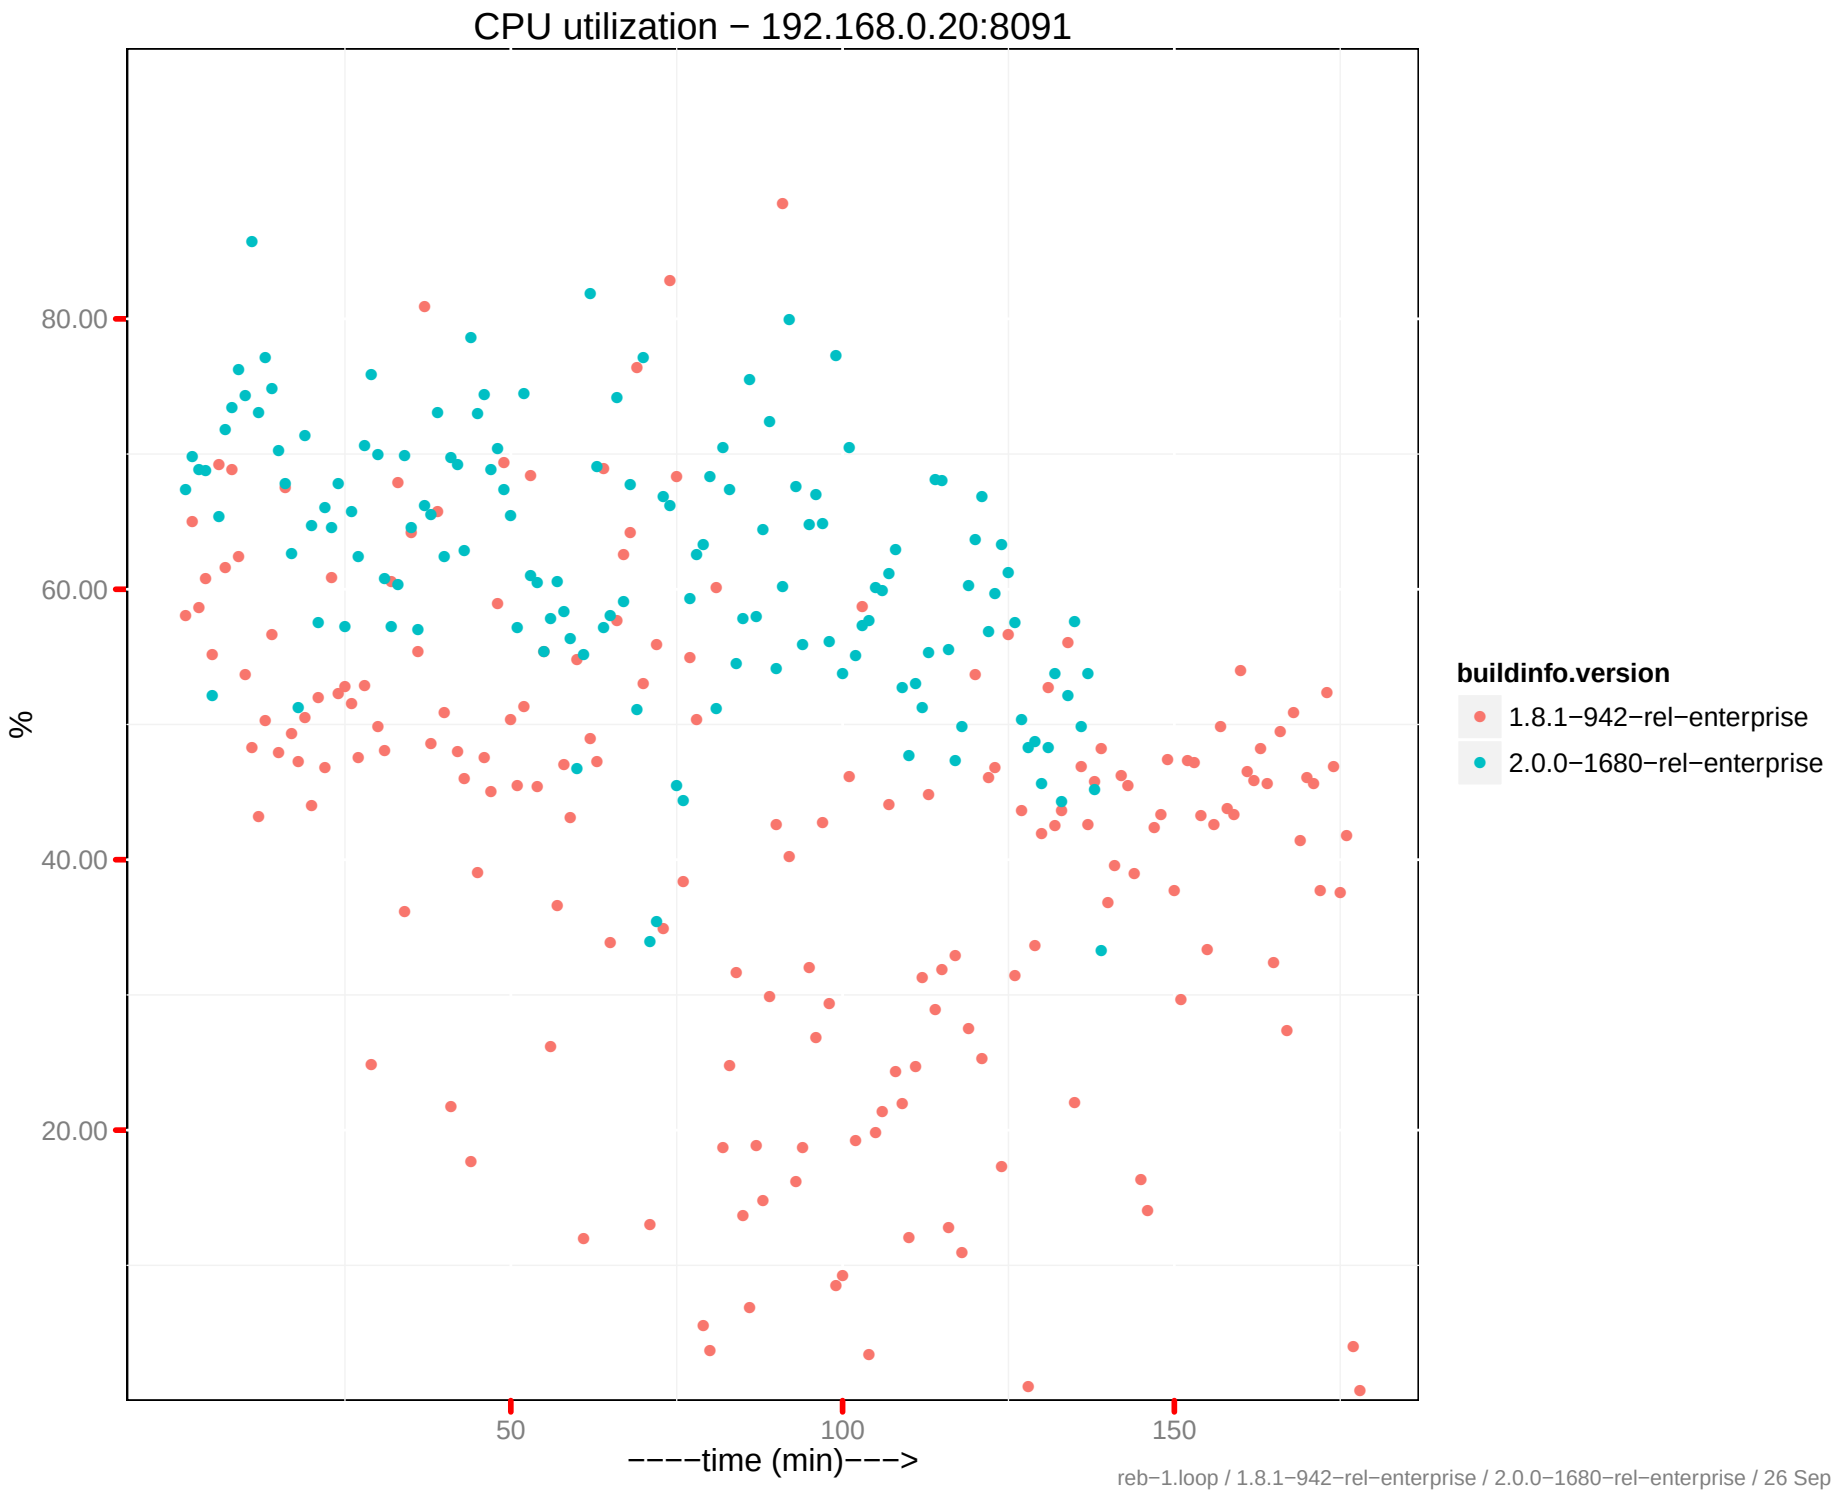
CPU utilization − 192.168.0.20:8091
80.00
60.00
40.00
20.00
%
50
100
150
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

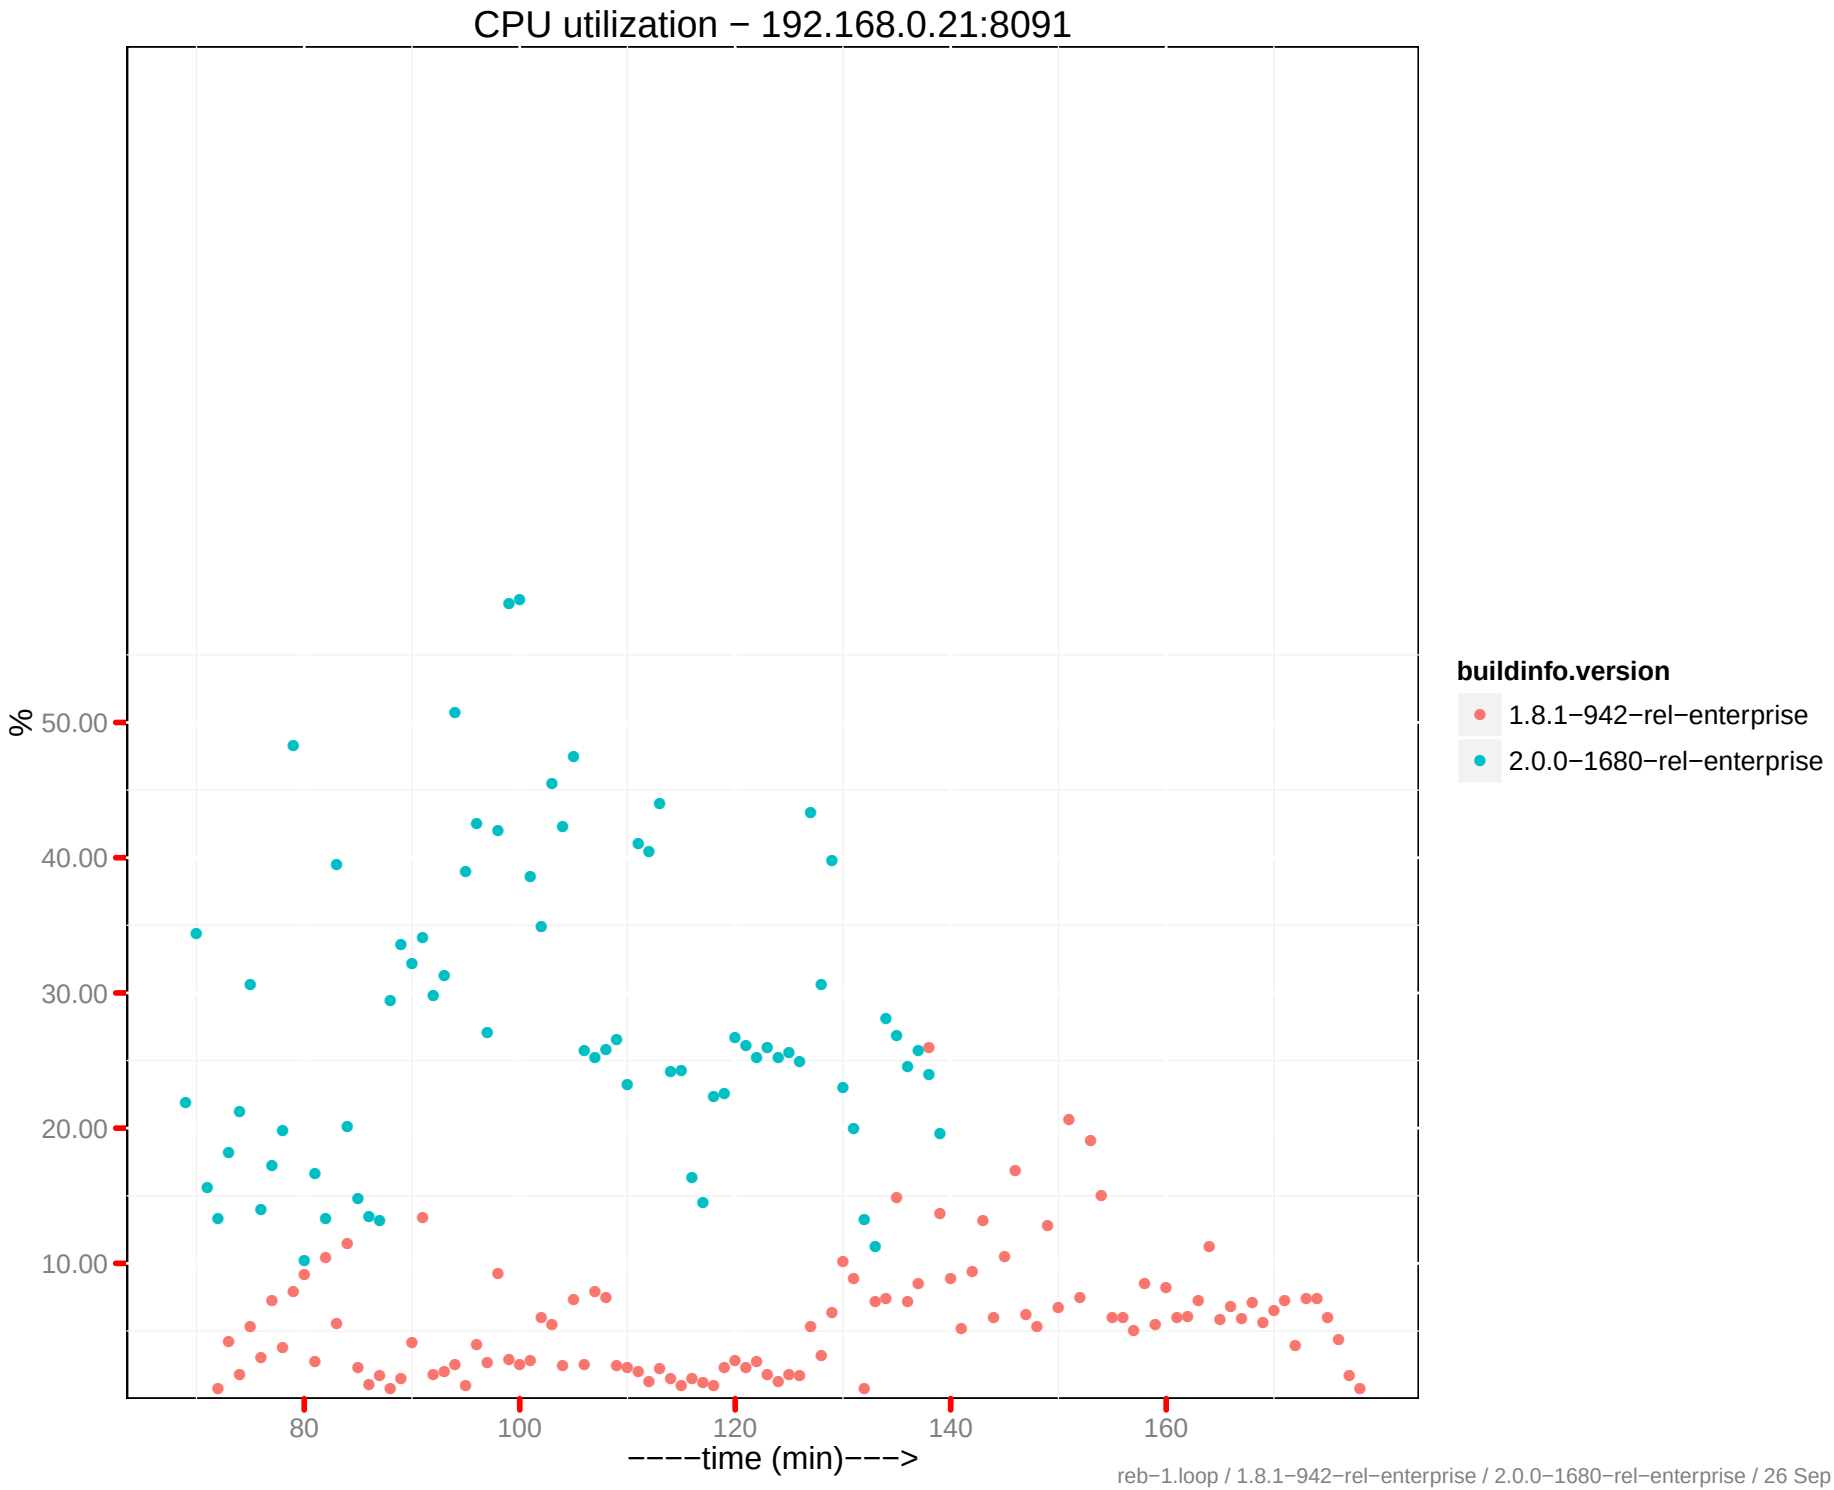
CPU utilization − 192.168.0.21:8091
%
50.00
40.00
30.00
20.00
10.00
80
100
120
140
160
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

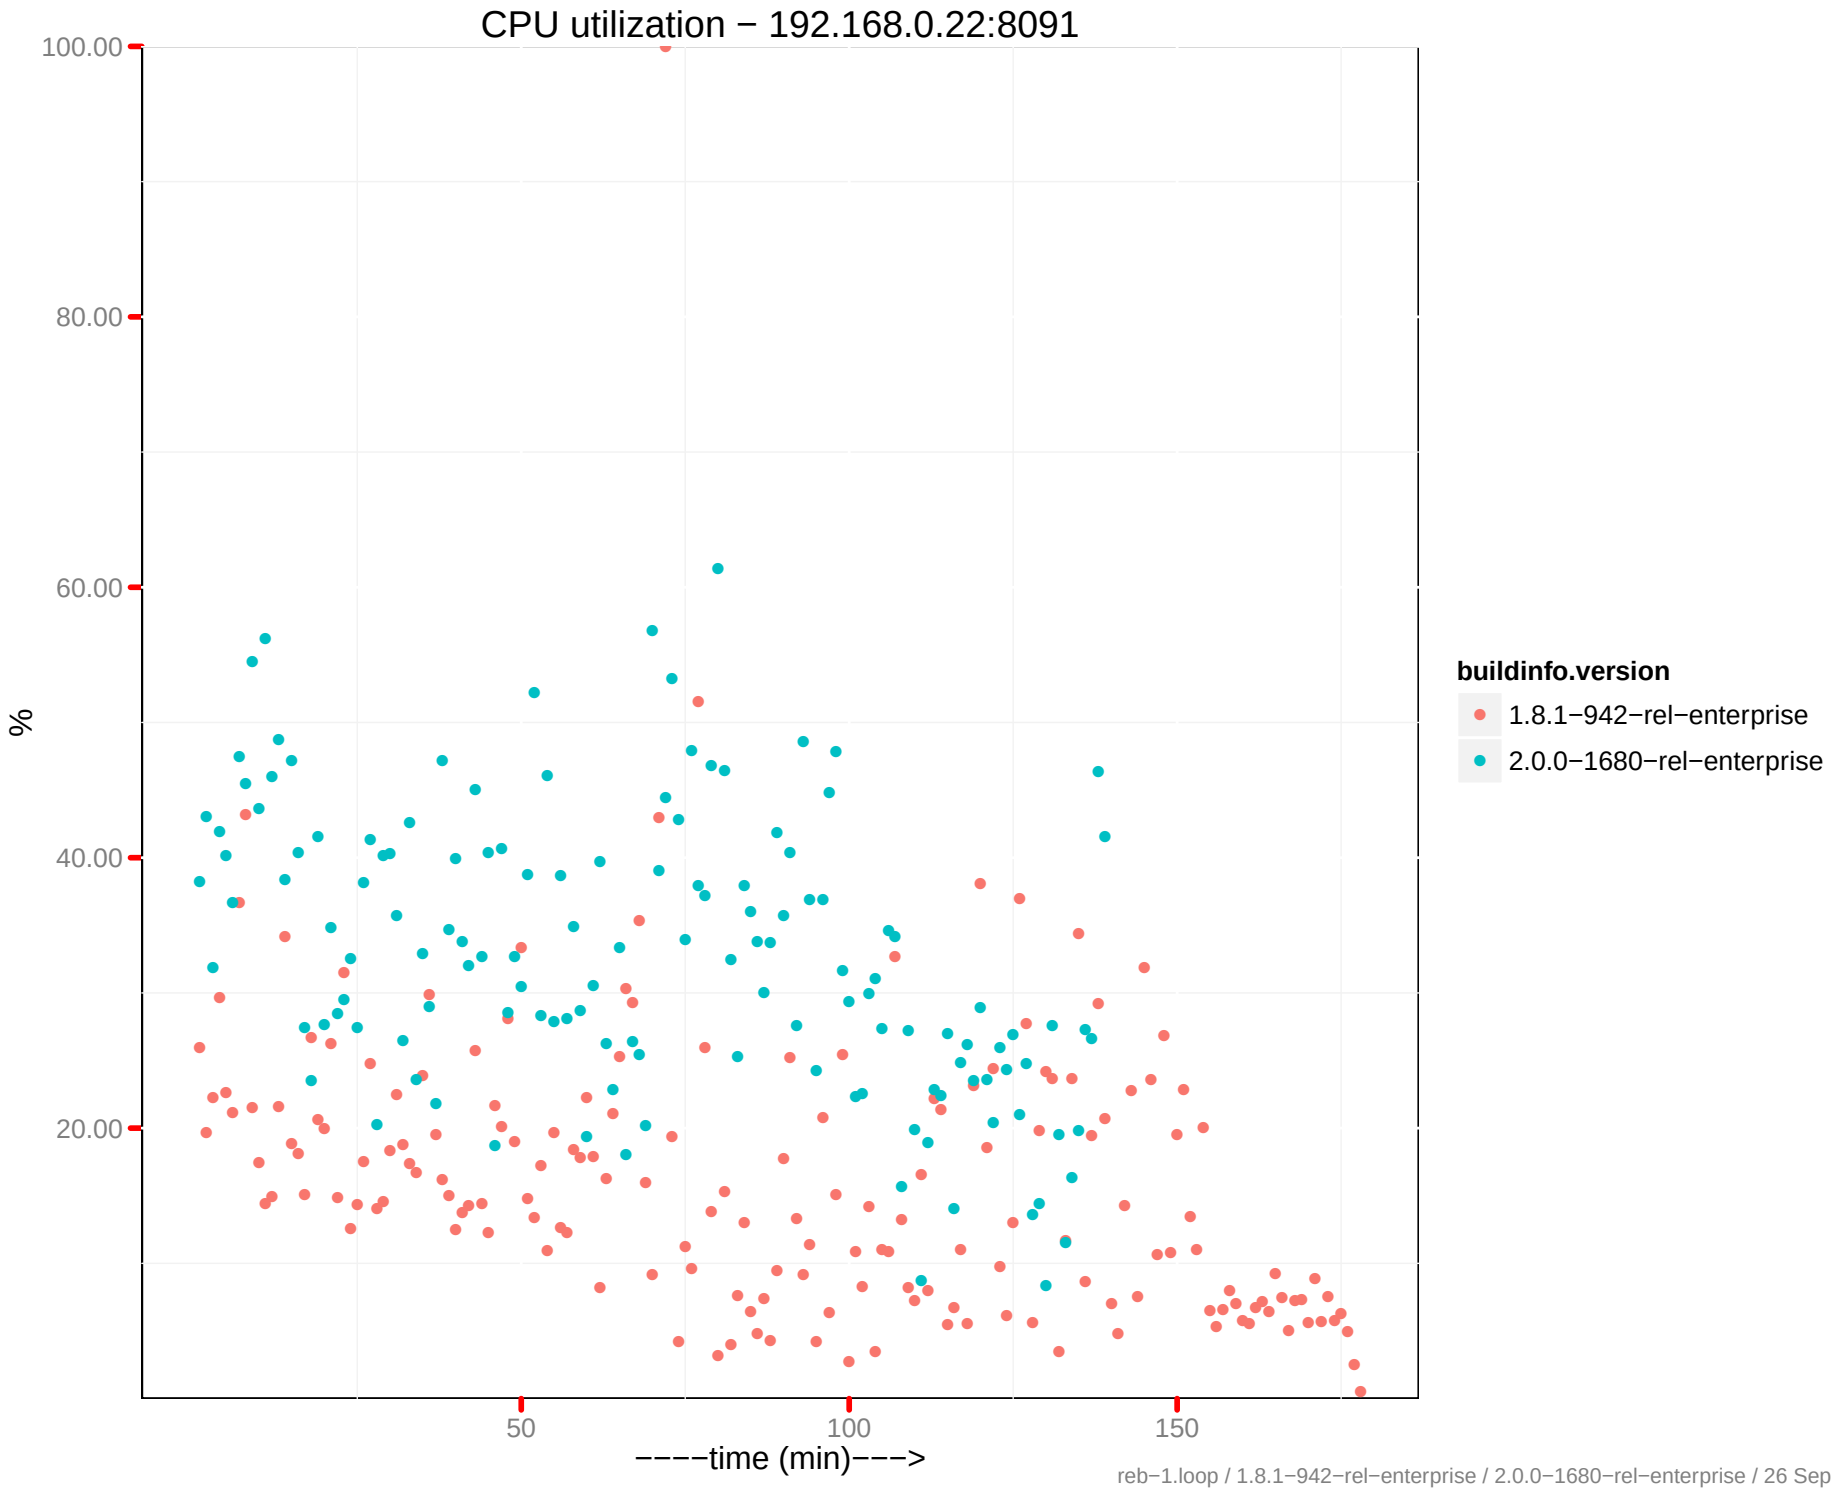
CPU utilization − 192.168.0.22:8091
100.00
80.00
60.00
40.00
20.00
%
50
100
150
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

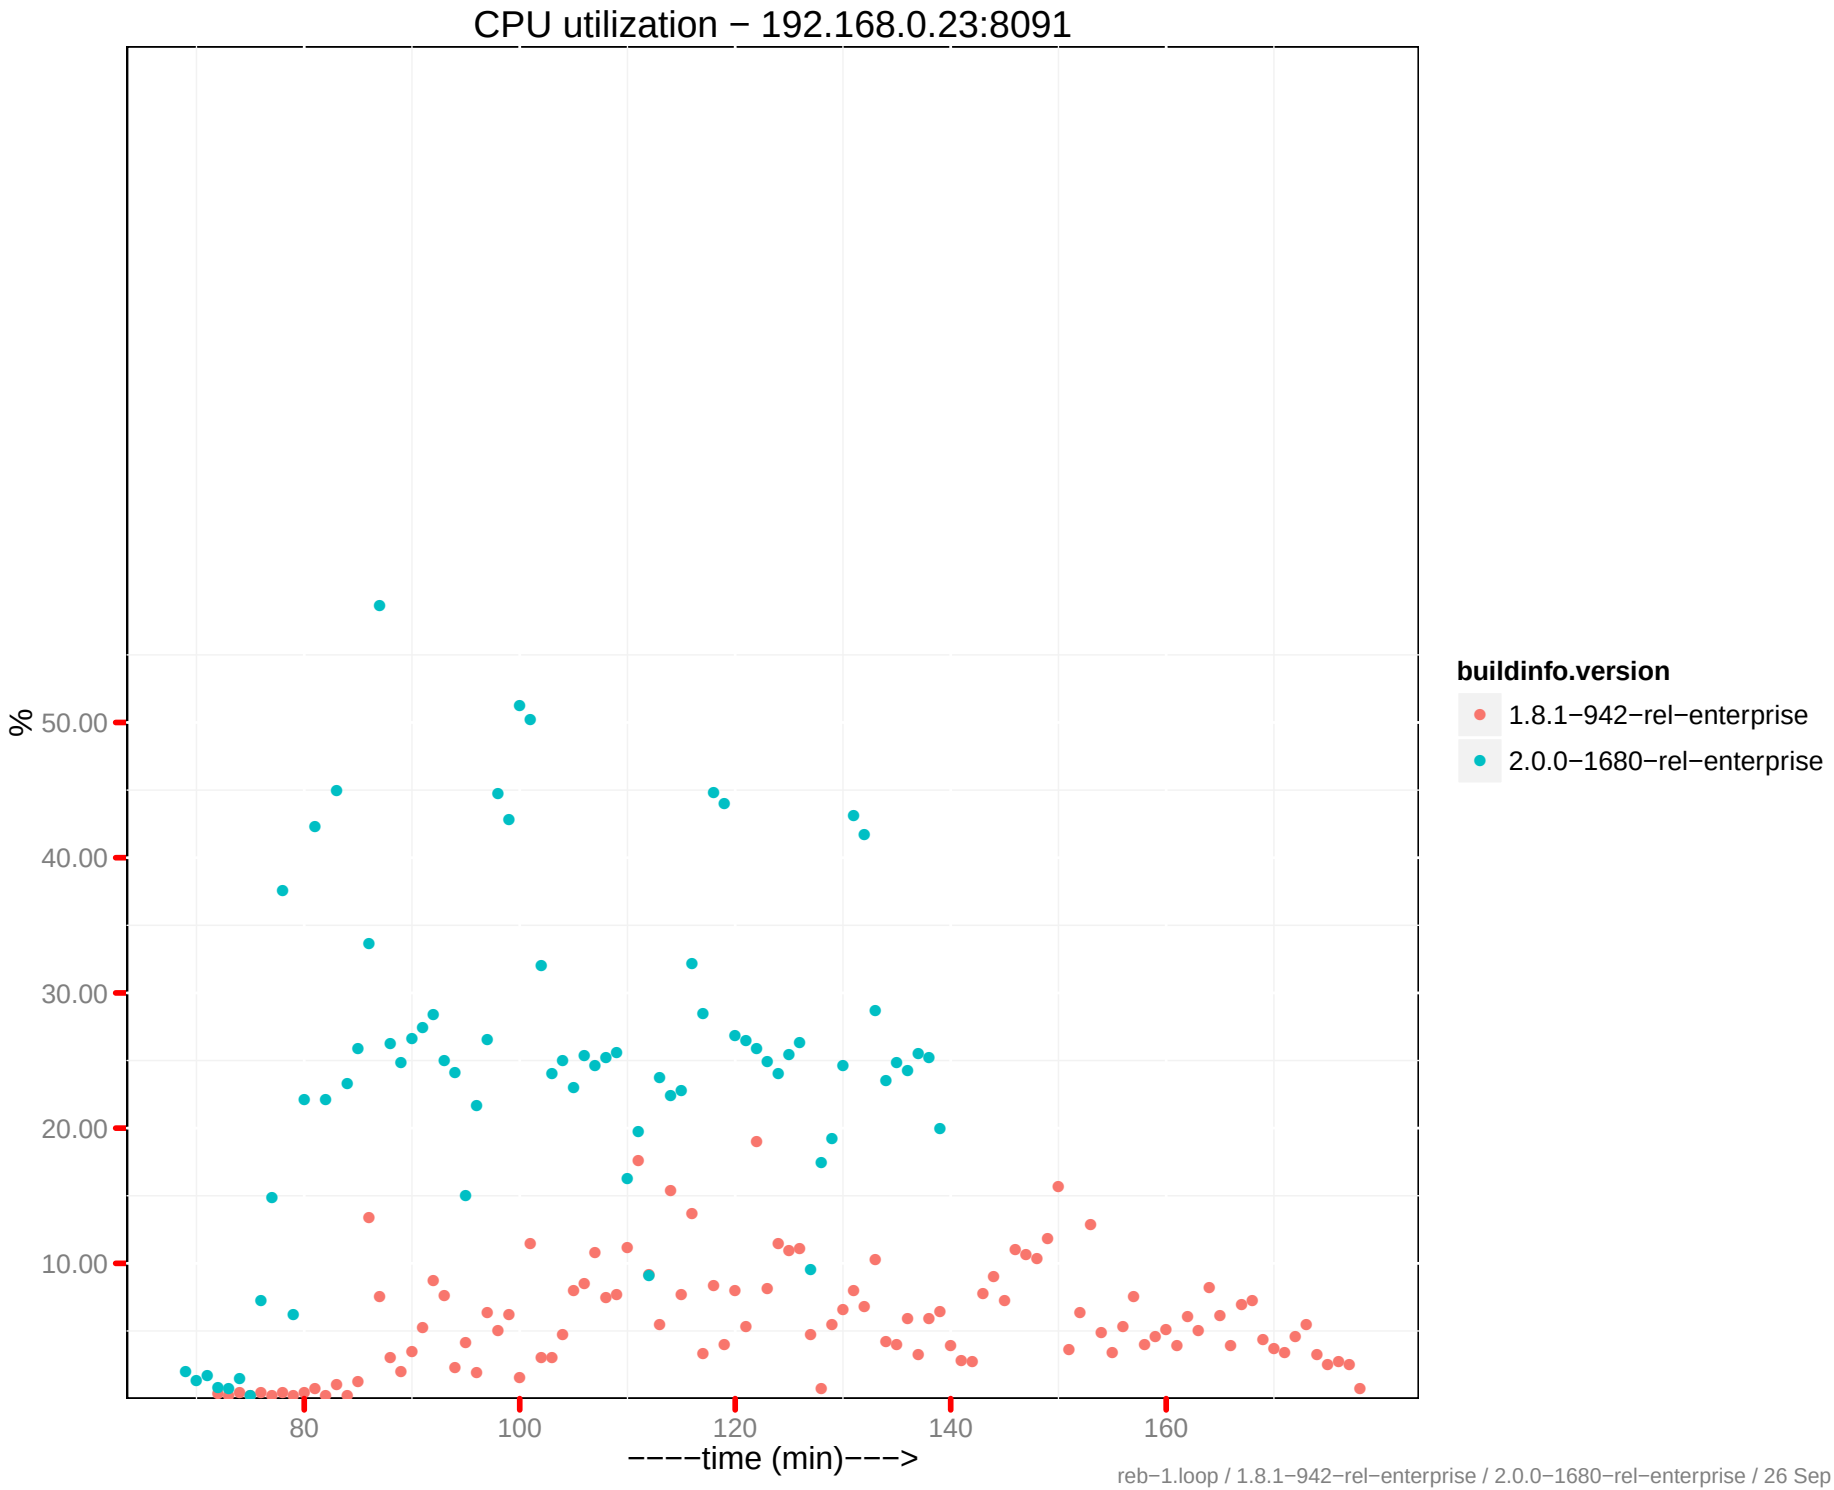
CPU utilization − 192.168.0.23:8091
%
50.00
40.00
30.00
20.00
10.00
80
100
120
140
160
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

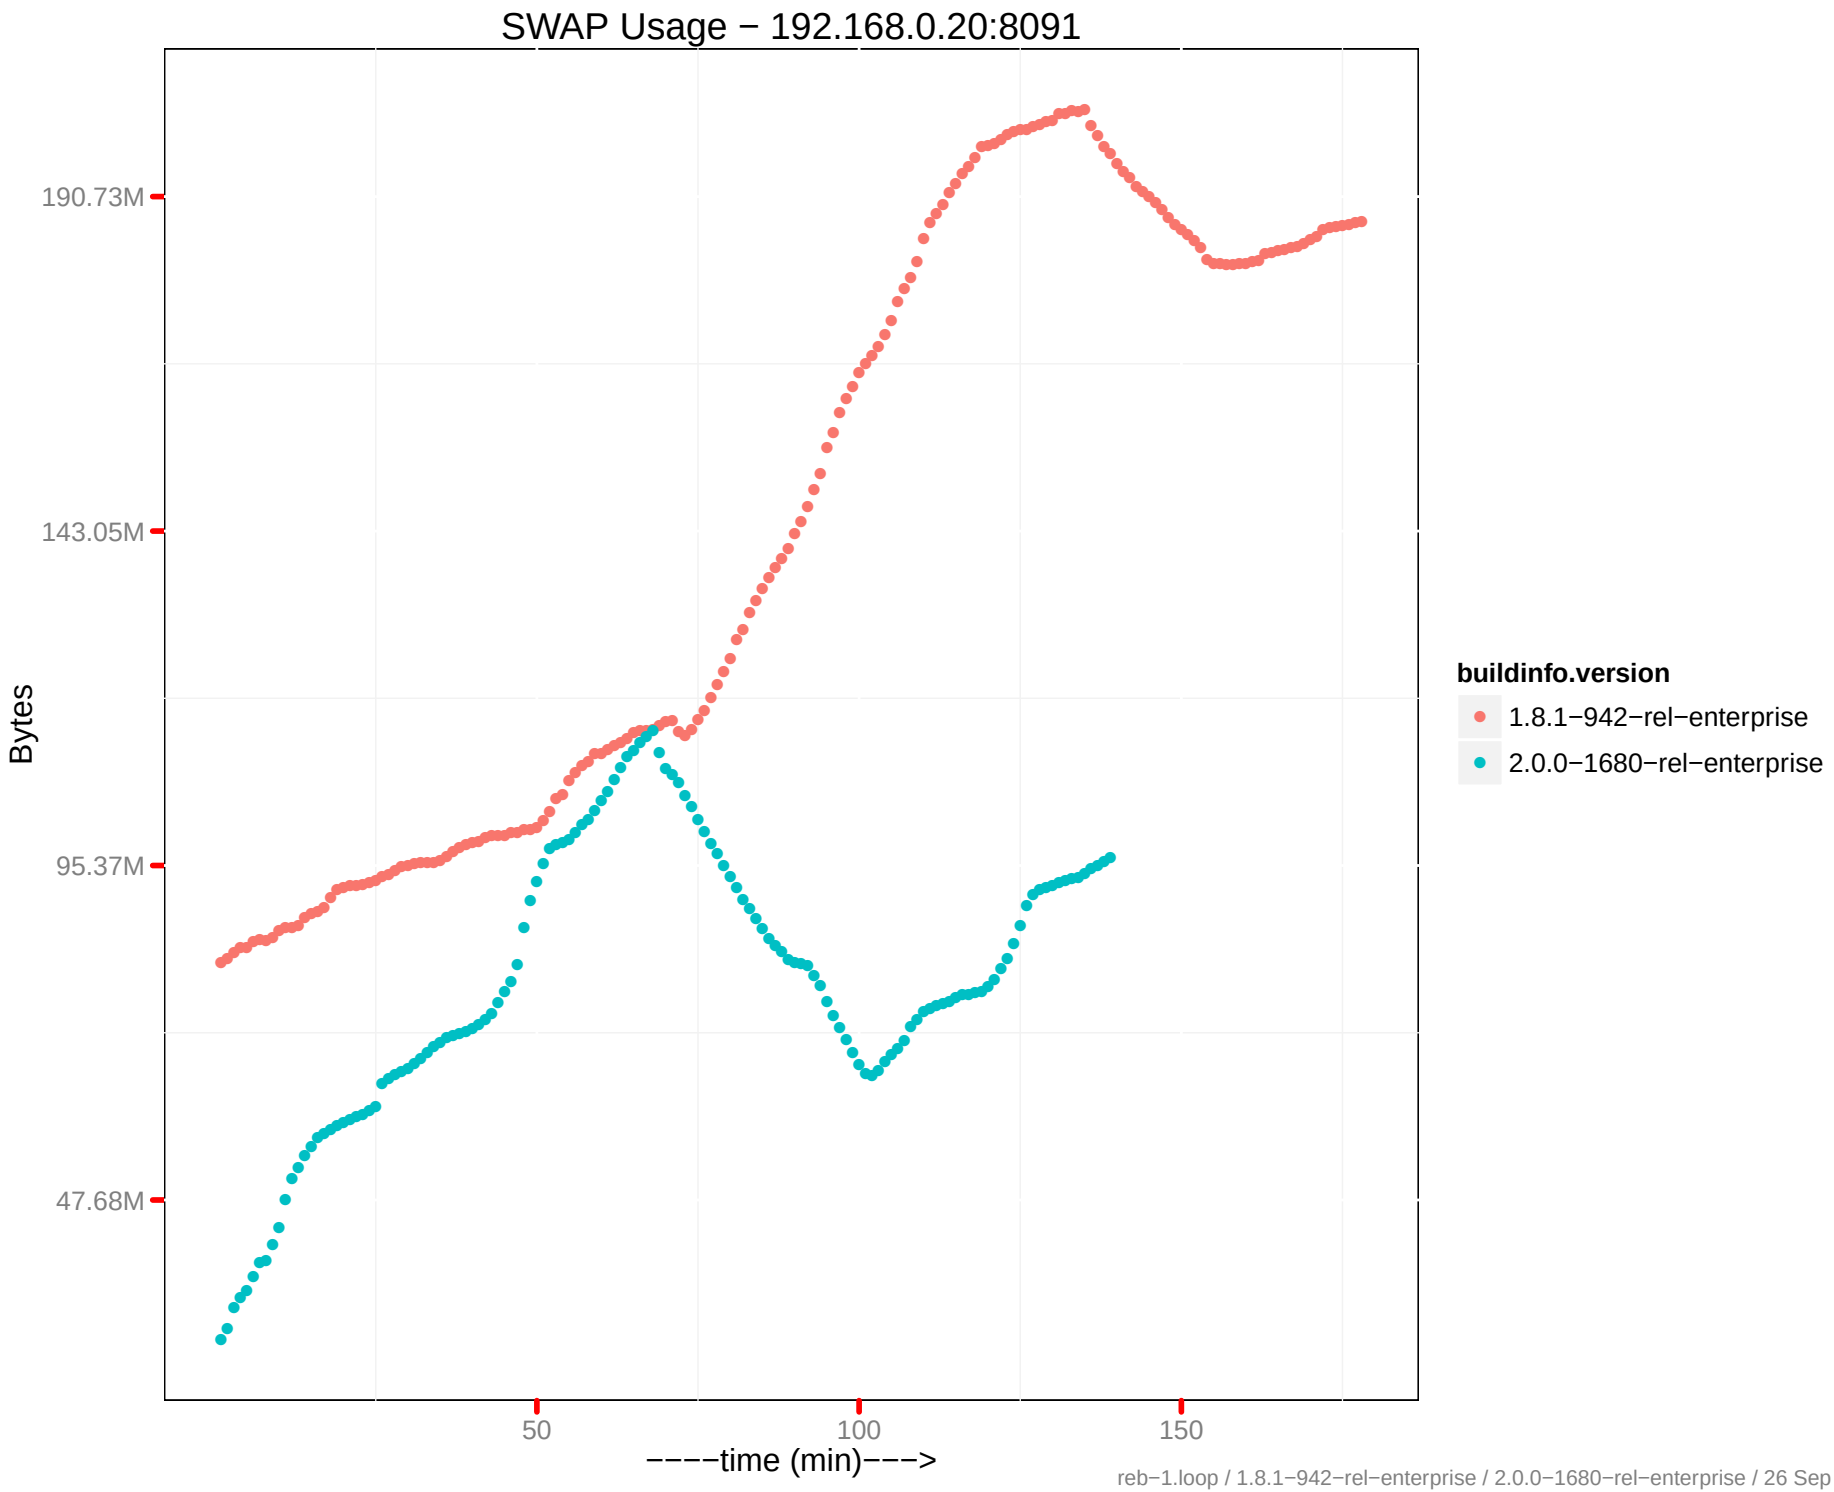
SWAP Usage − 192.168.0.20:8091
190.73M
143.05M
95.37M
47.68M
Bytes
50
100
150
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

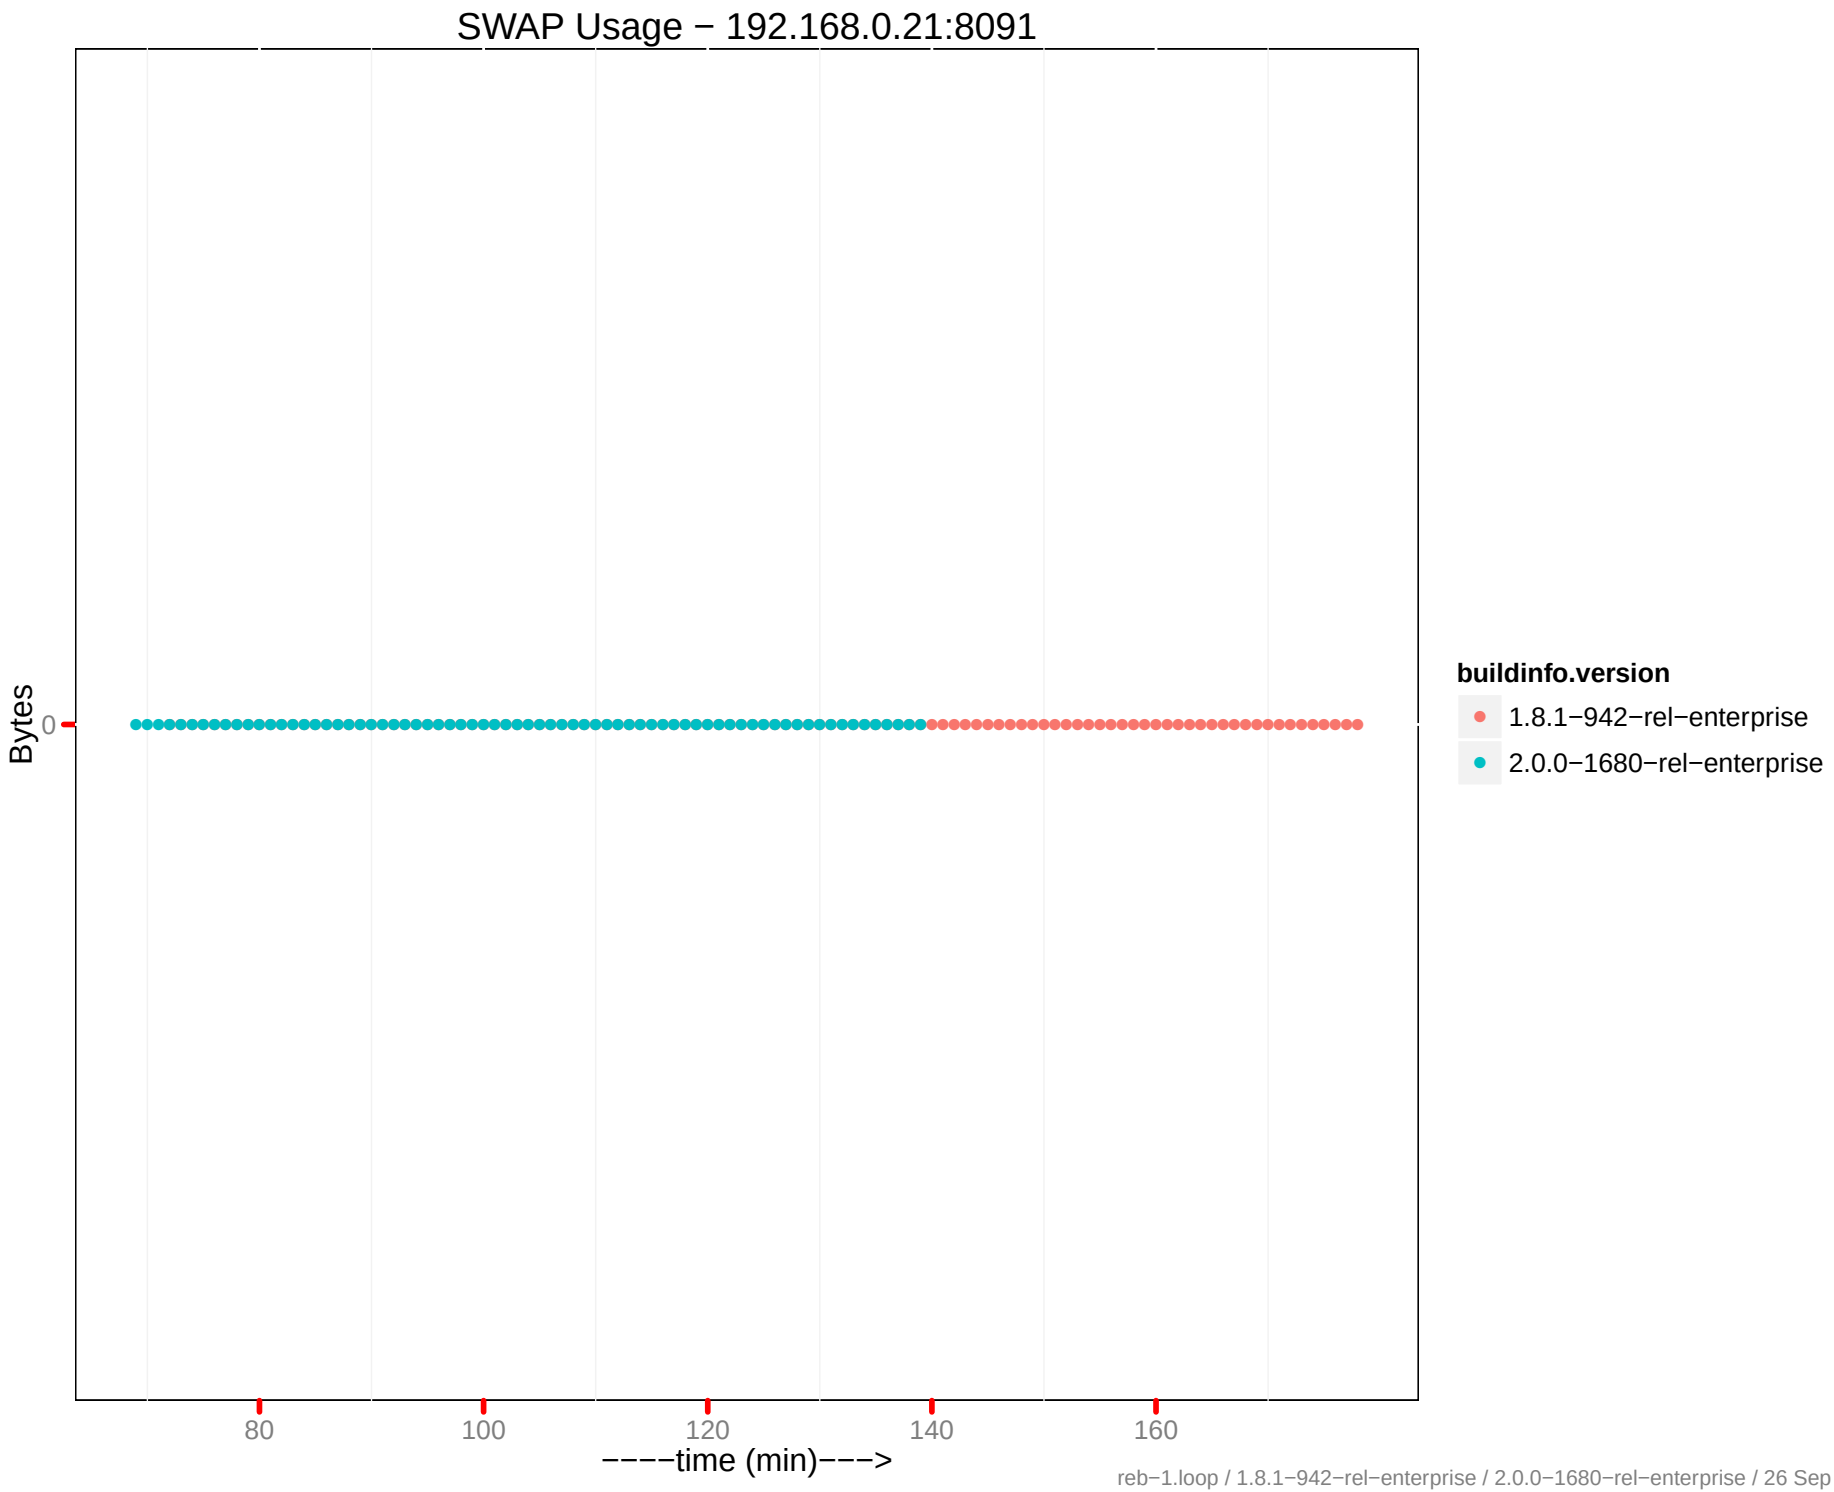
SWAP Usage − 192.168.0.21:8091
Bytes
0
80
100
120
140
160
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

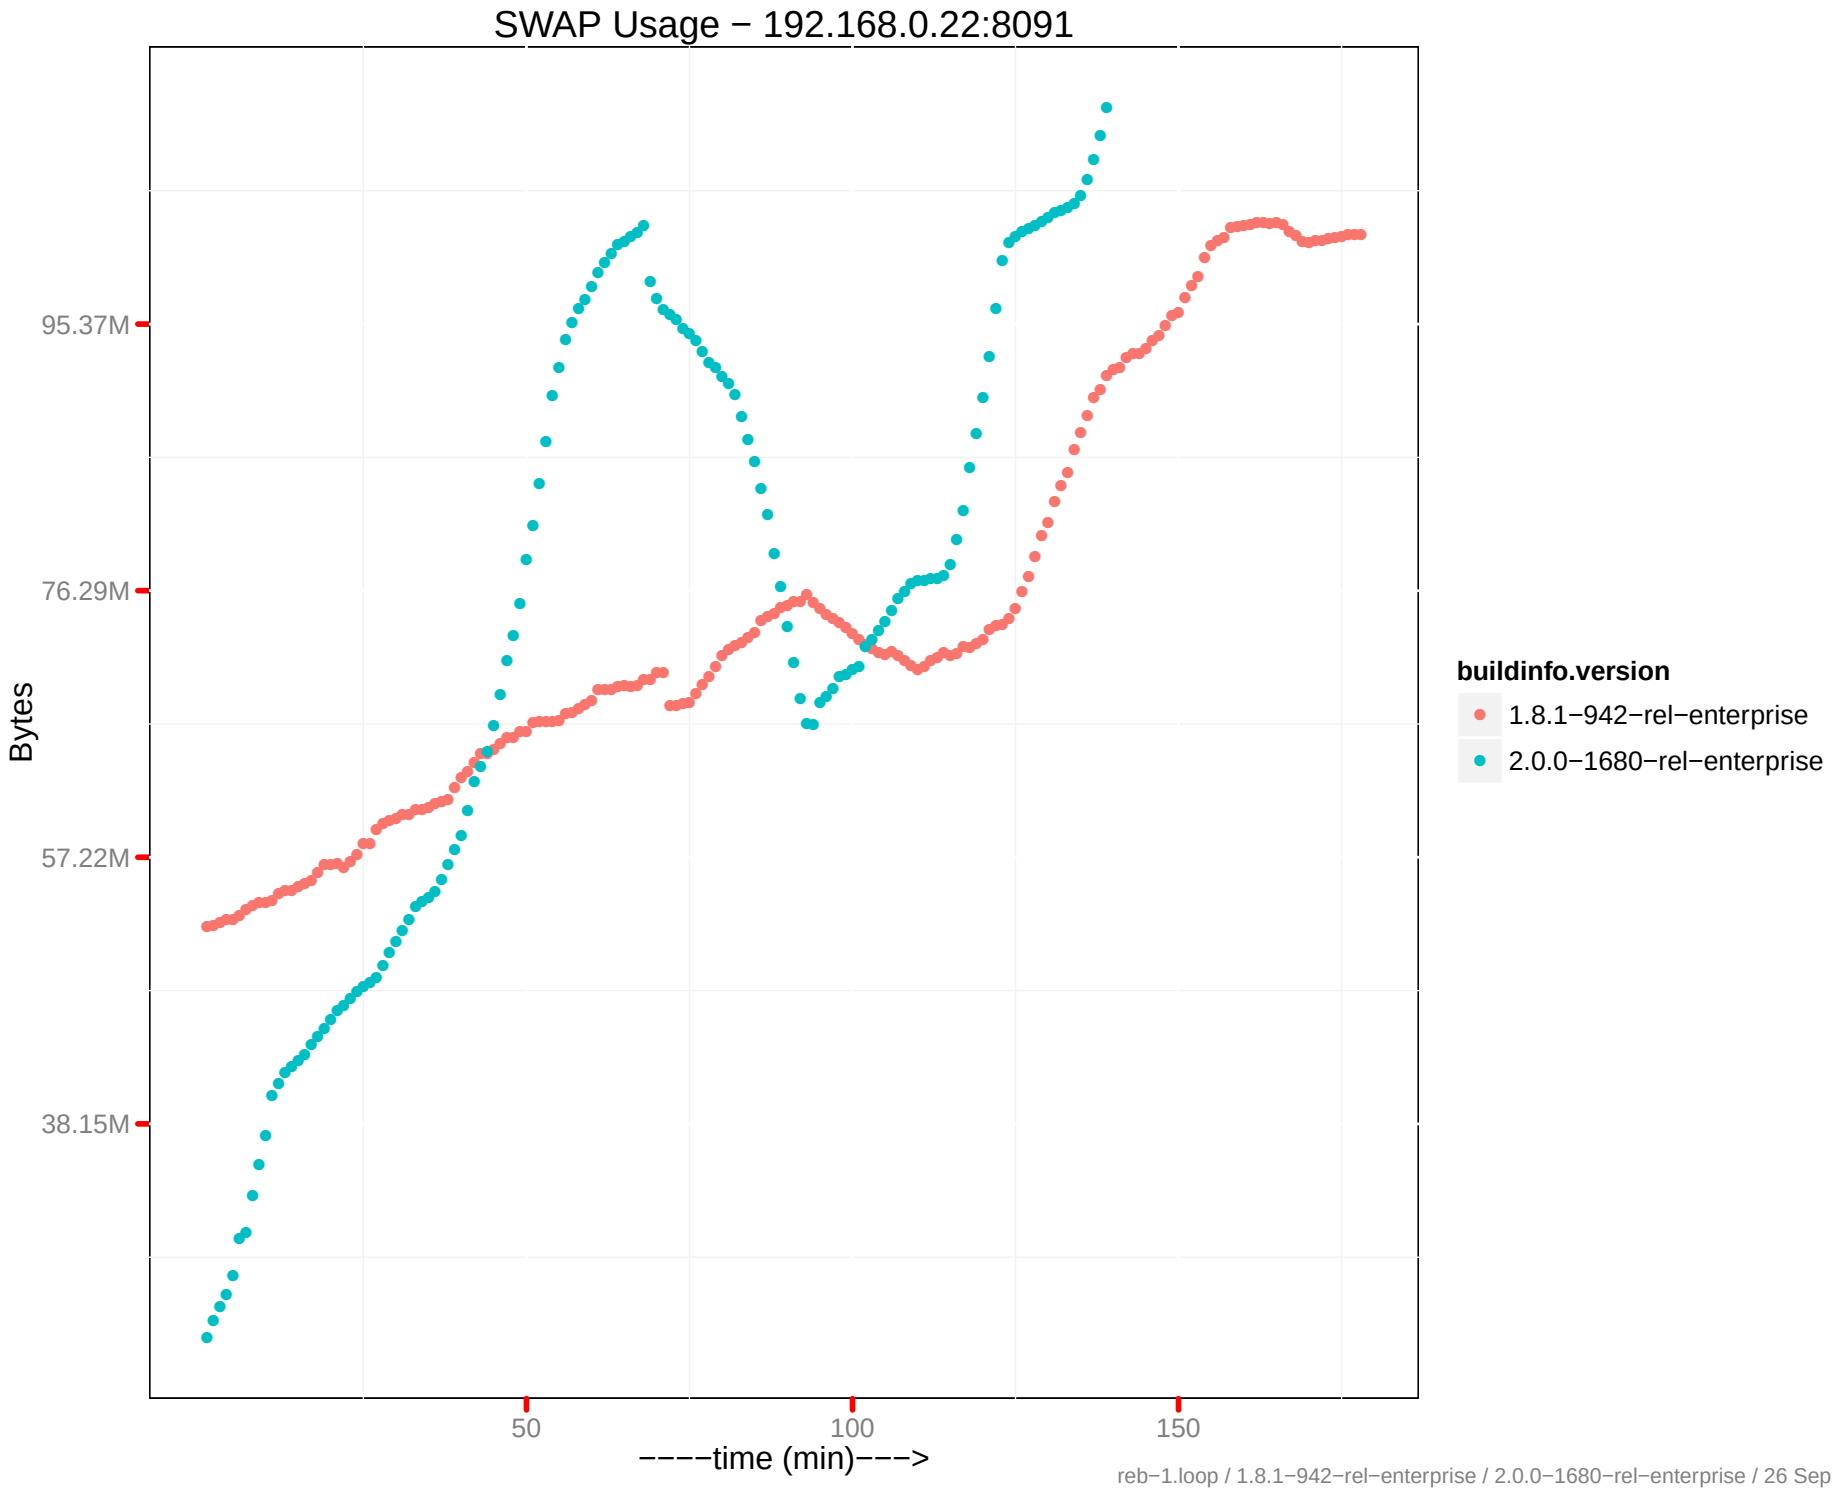
SWAP Usage − 192.168.0.22:8091
95.37M
76.29M
57.22M
38.15M
Bytes
50
100
150
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

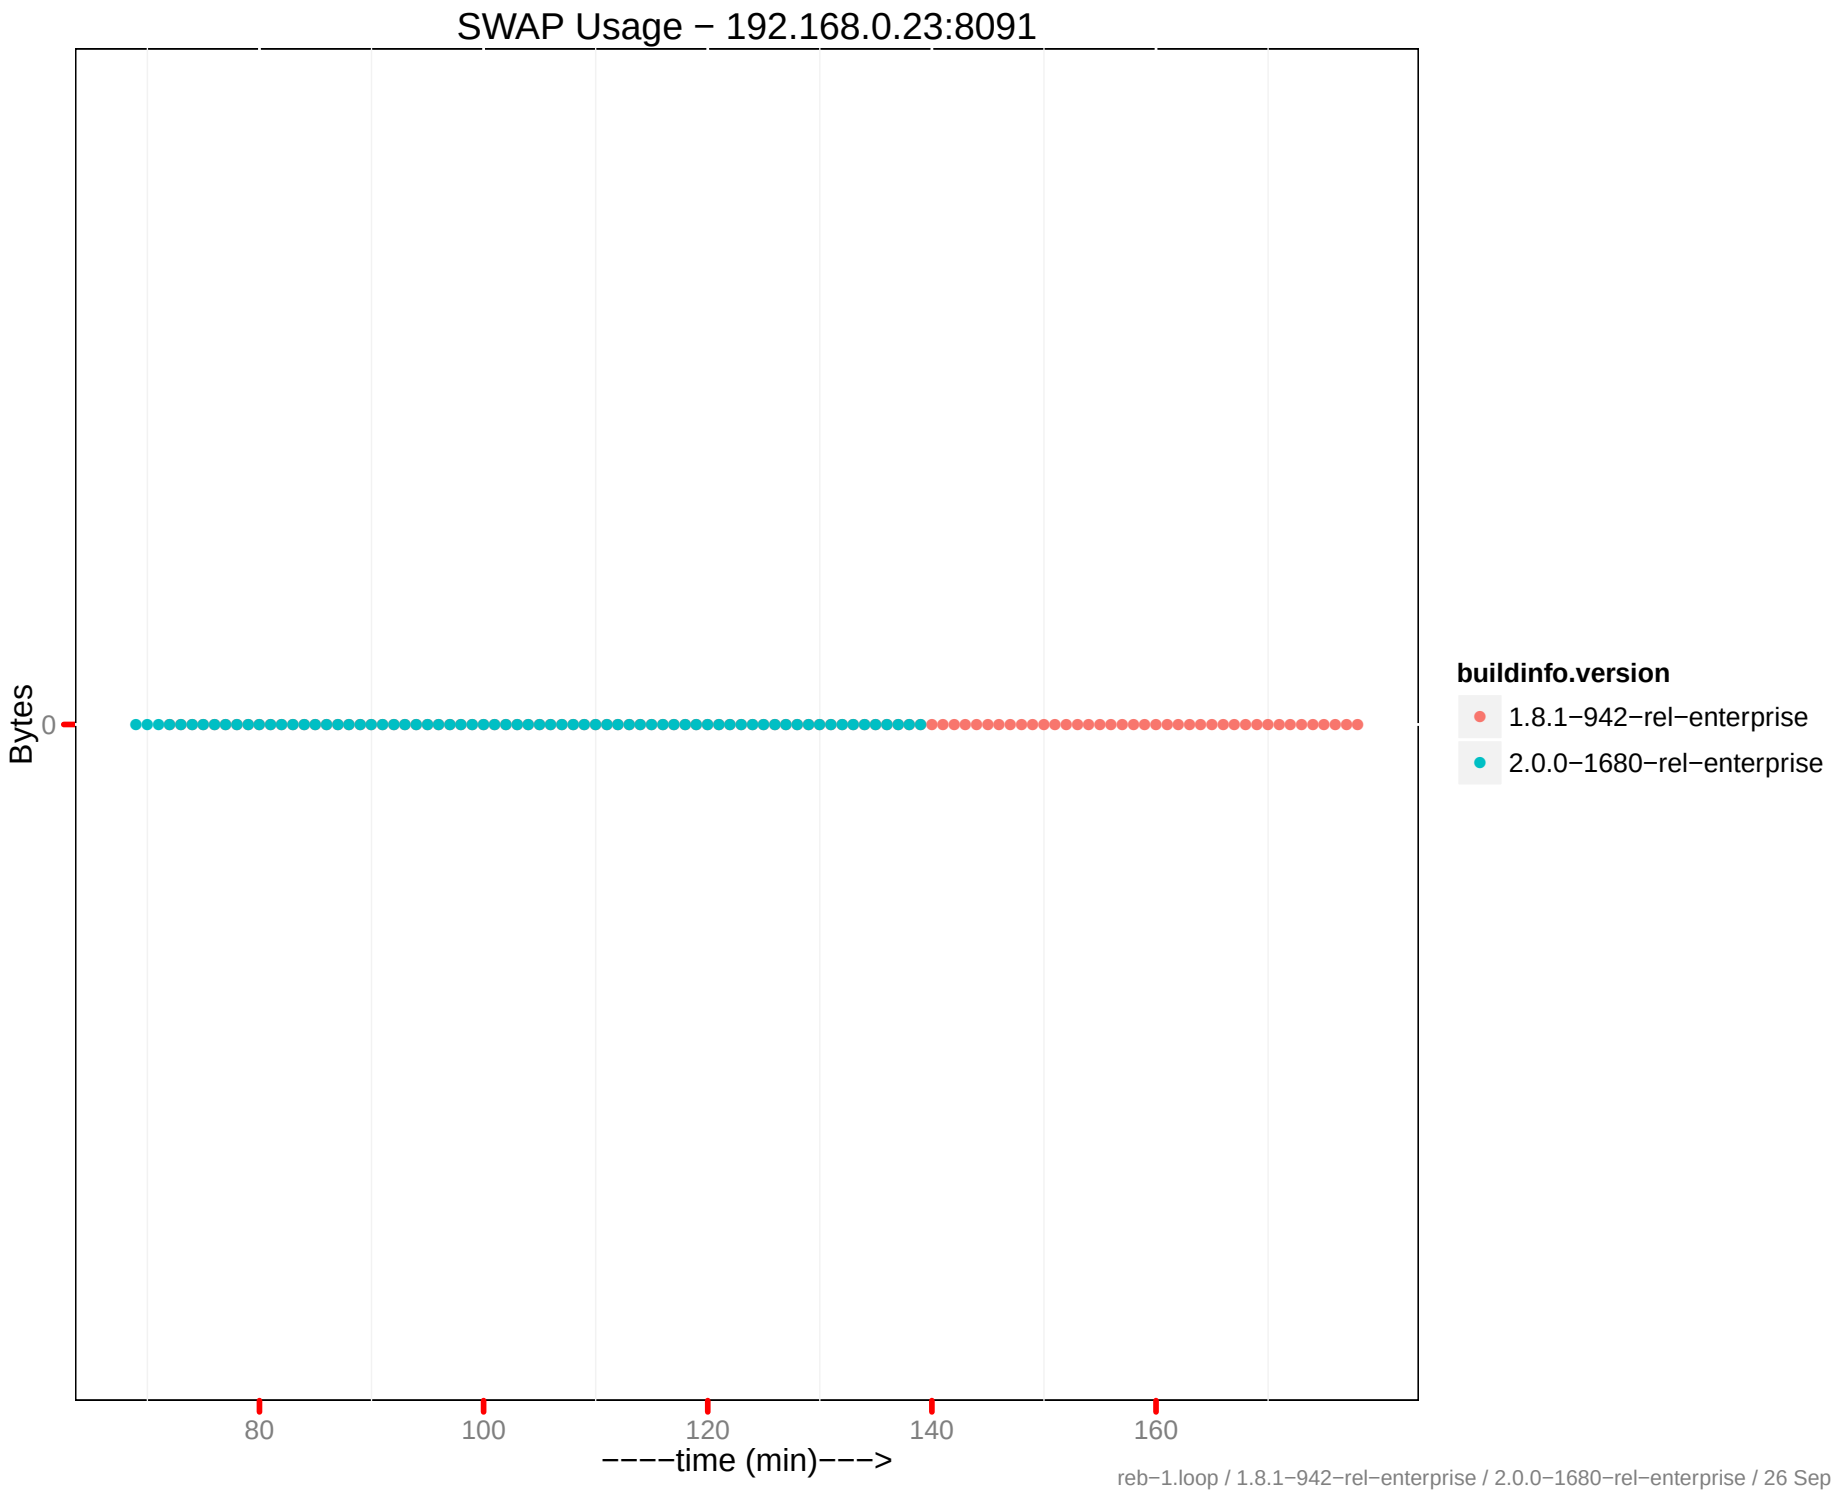
SWAP Usage − 192.168.0.23:8091
Bytes
0
80
100
120
140
160
−−−−time (min)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

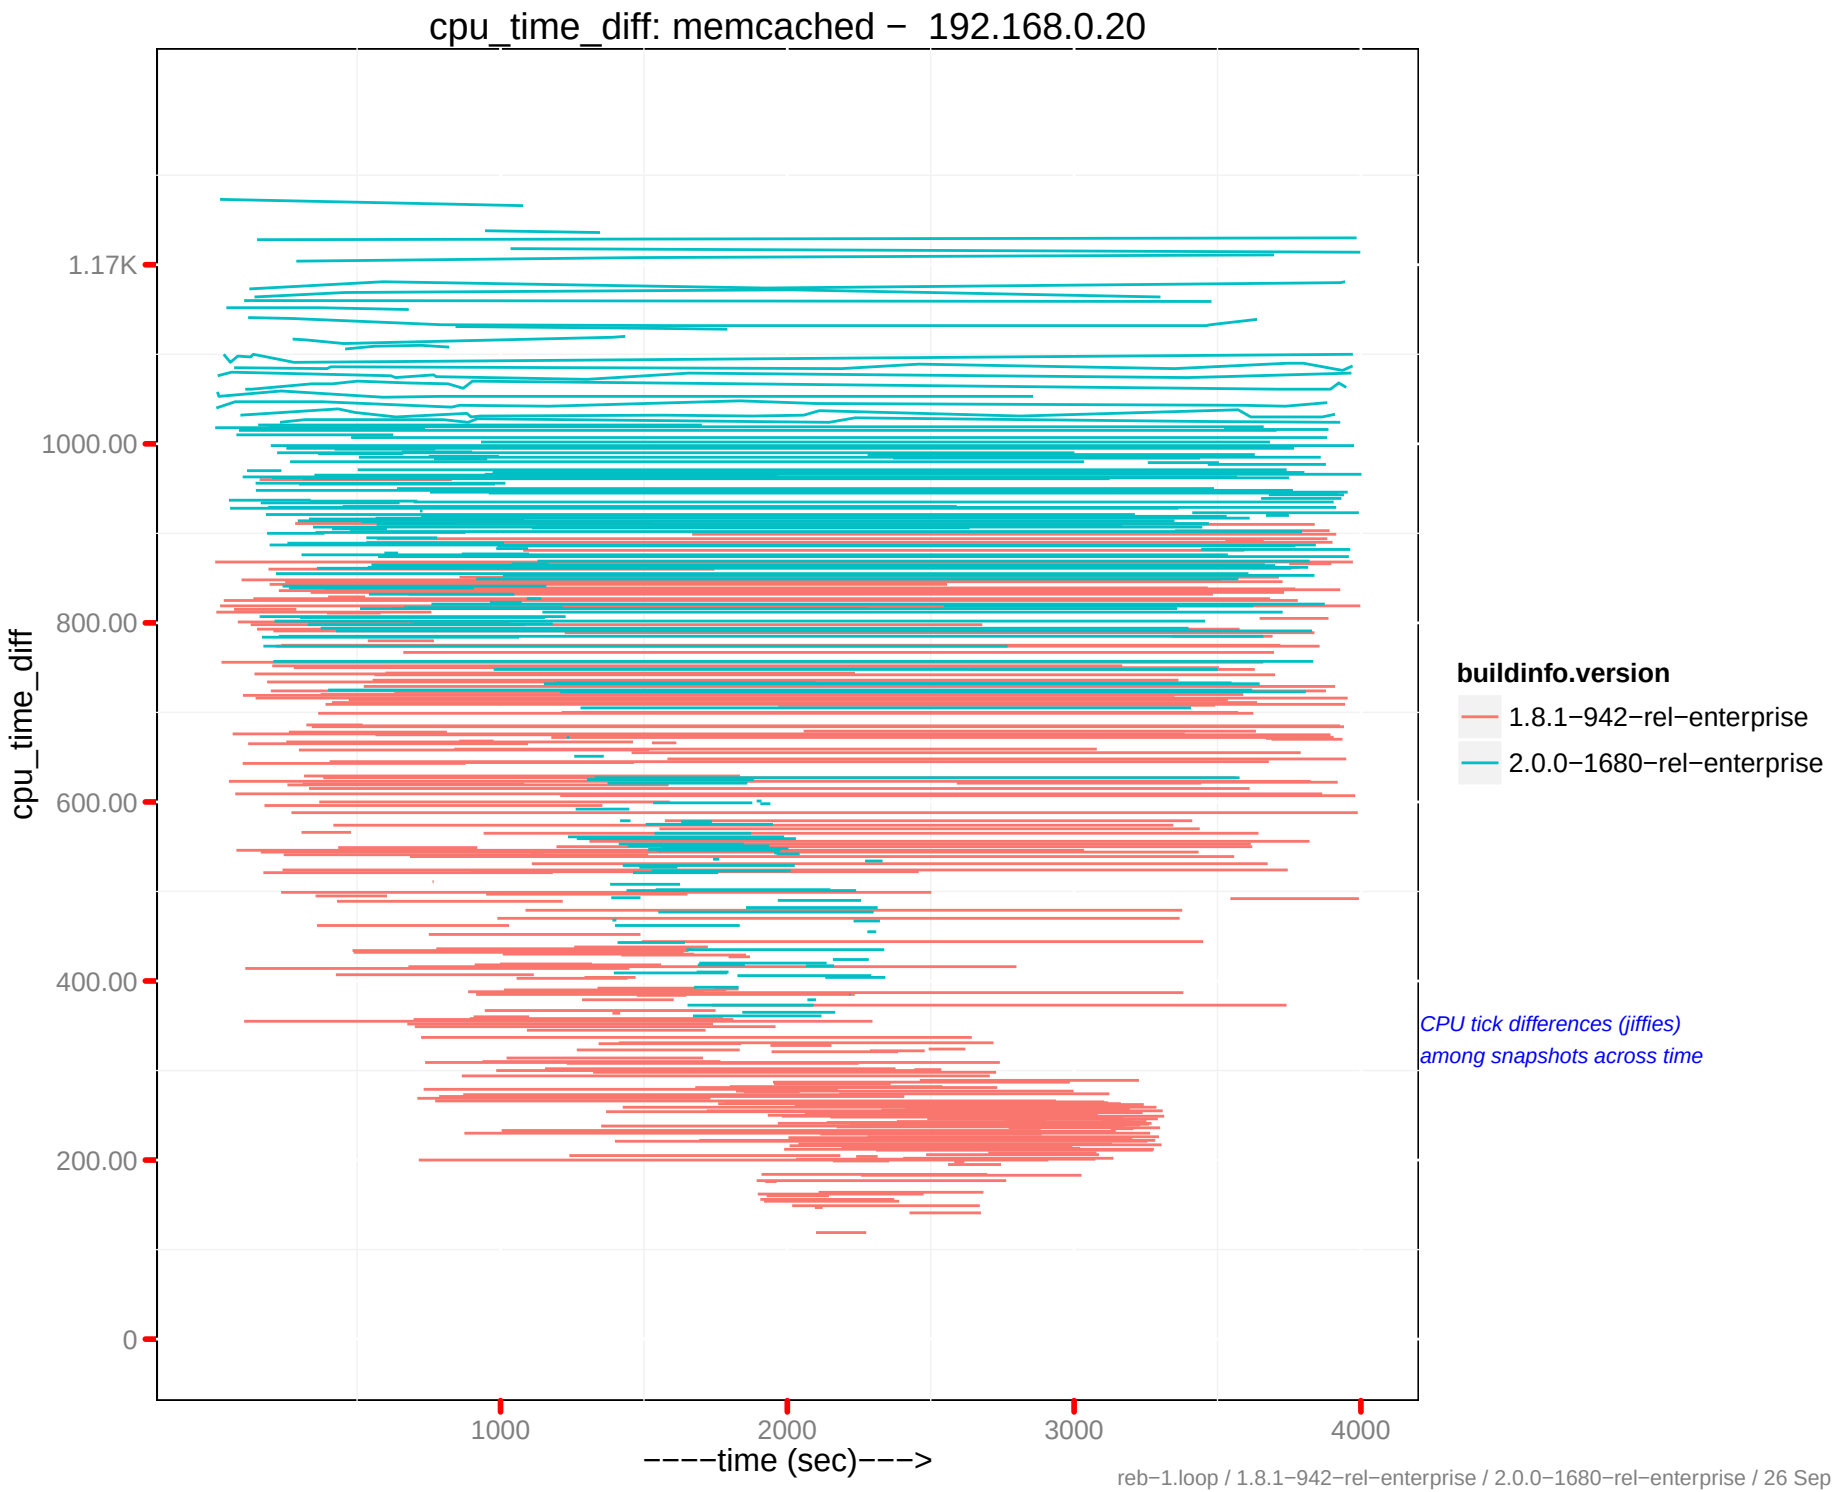
cpu_time_diff: memcached − 192.168.0.20
1.17K
1000.00
800.00
600.00
400.00
200.00
0
cpu_time_diff
1000
2000
3000
4000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

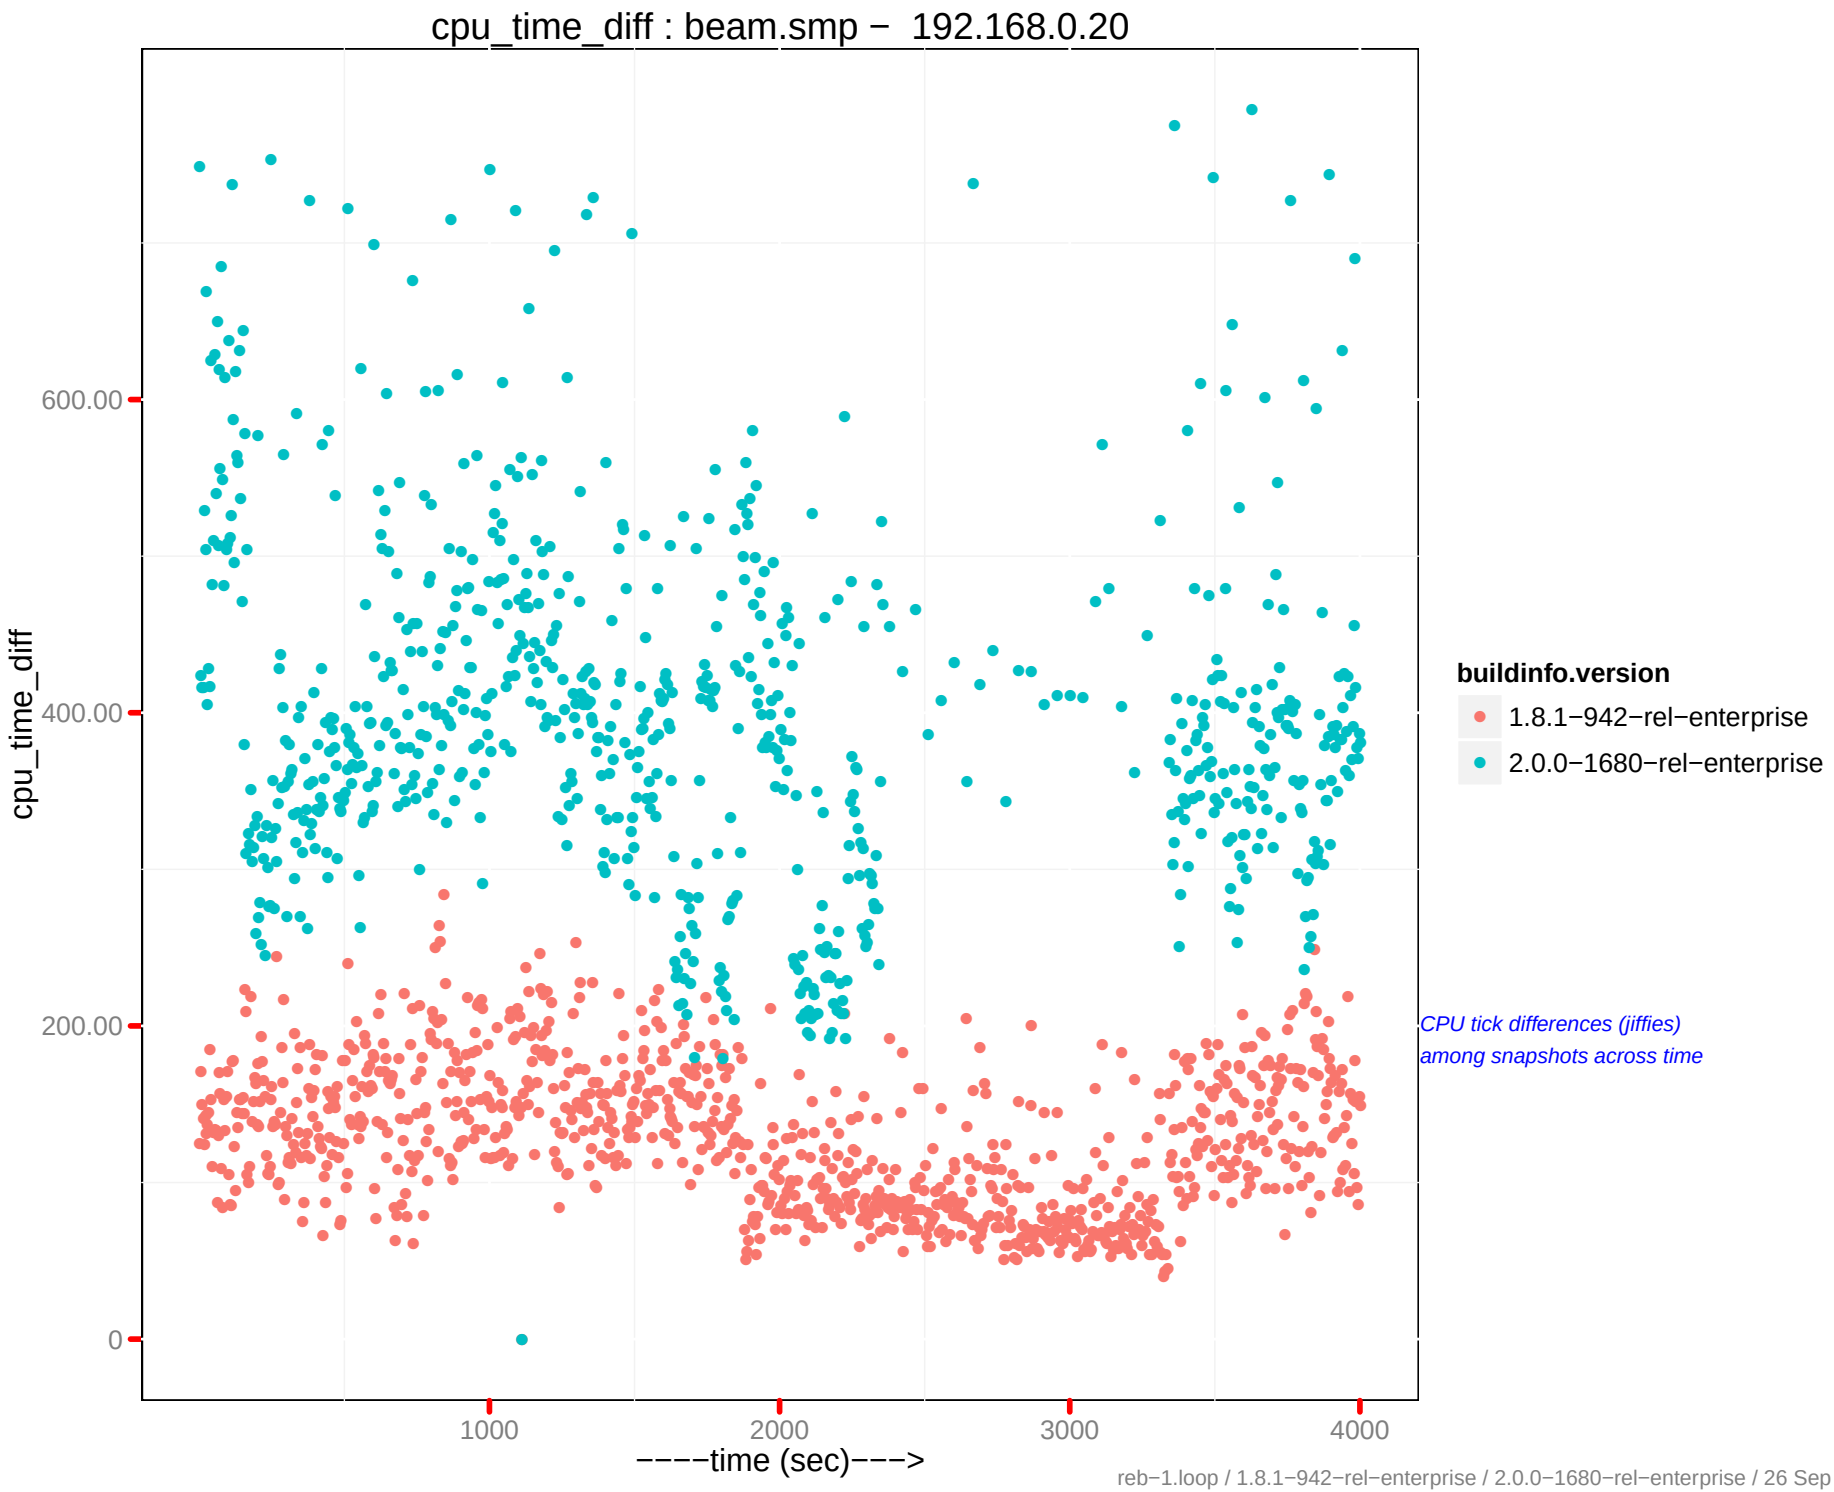
cpu_time_diff : beam.smp − 192.168.0.20
cpu_time_diff
600.00
400.00
200.00
0
1000
2000
3000
4000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

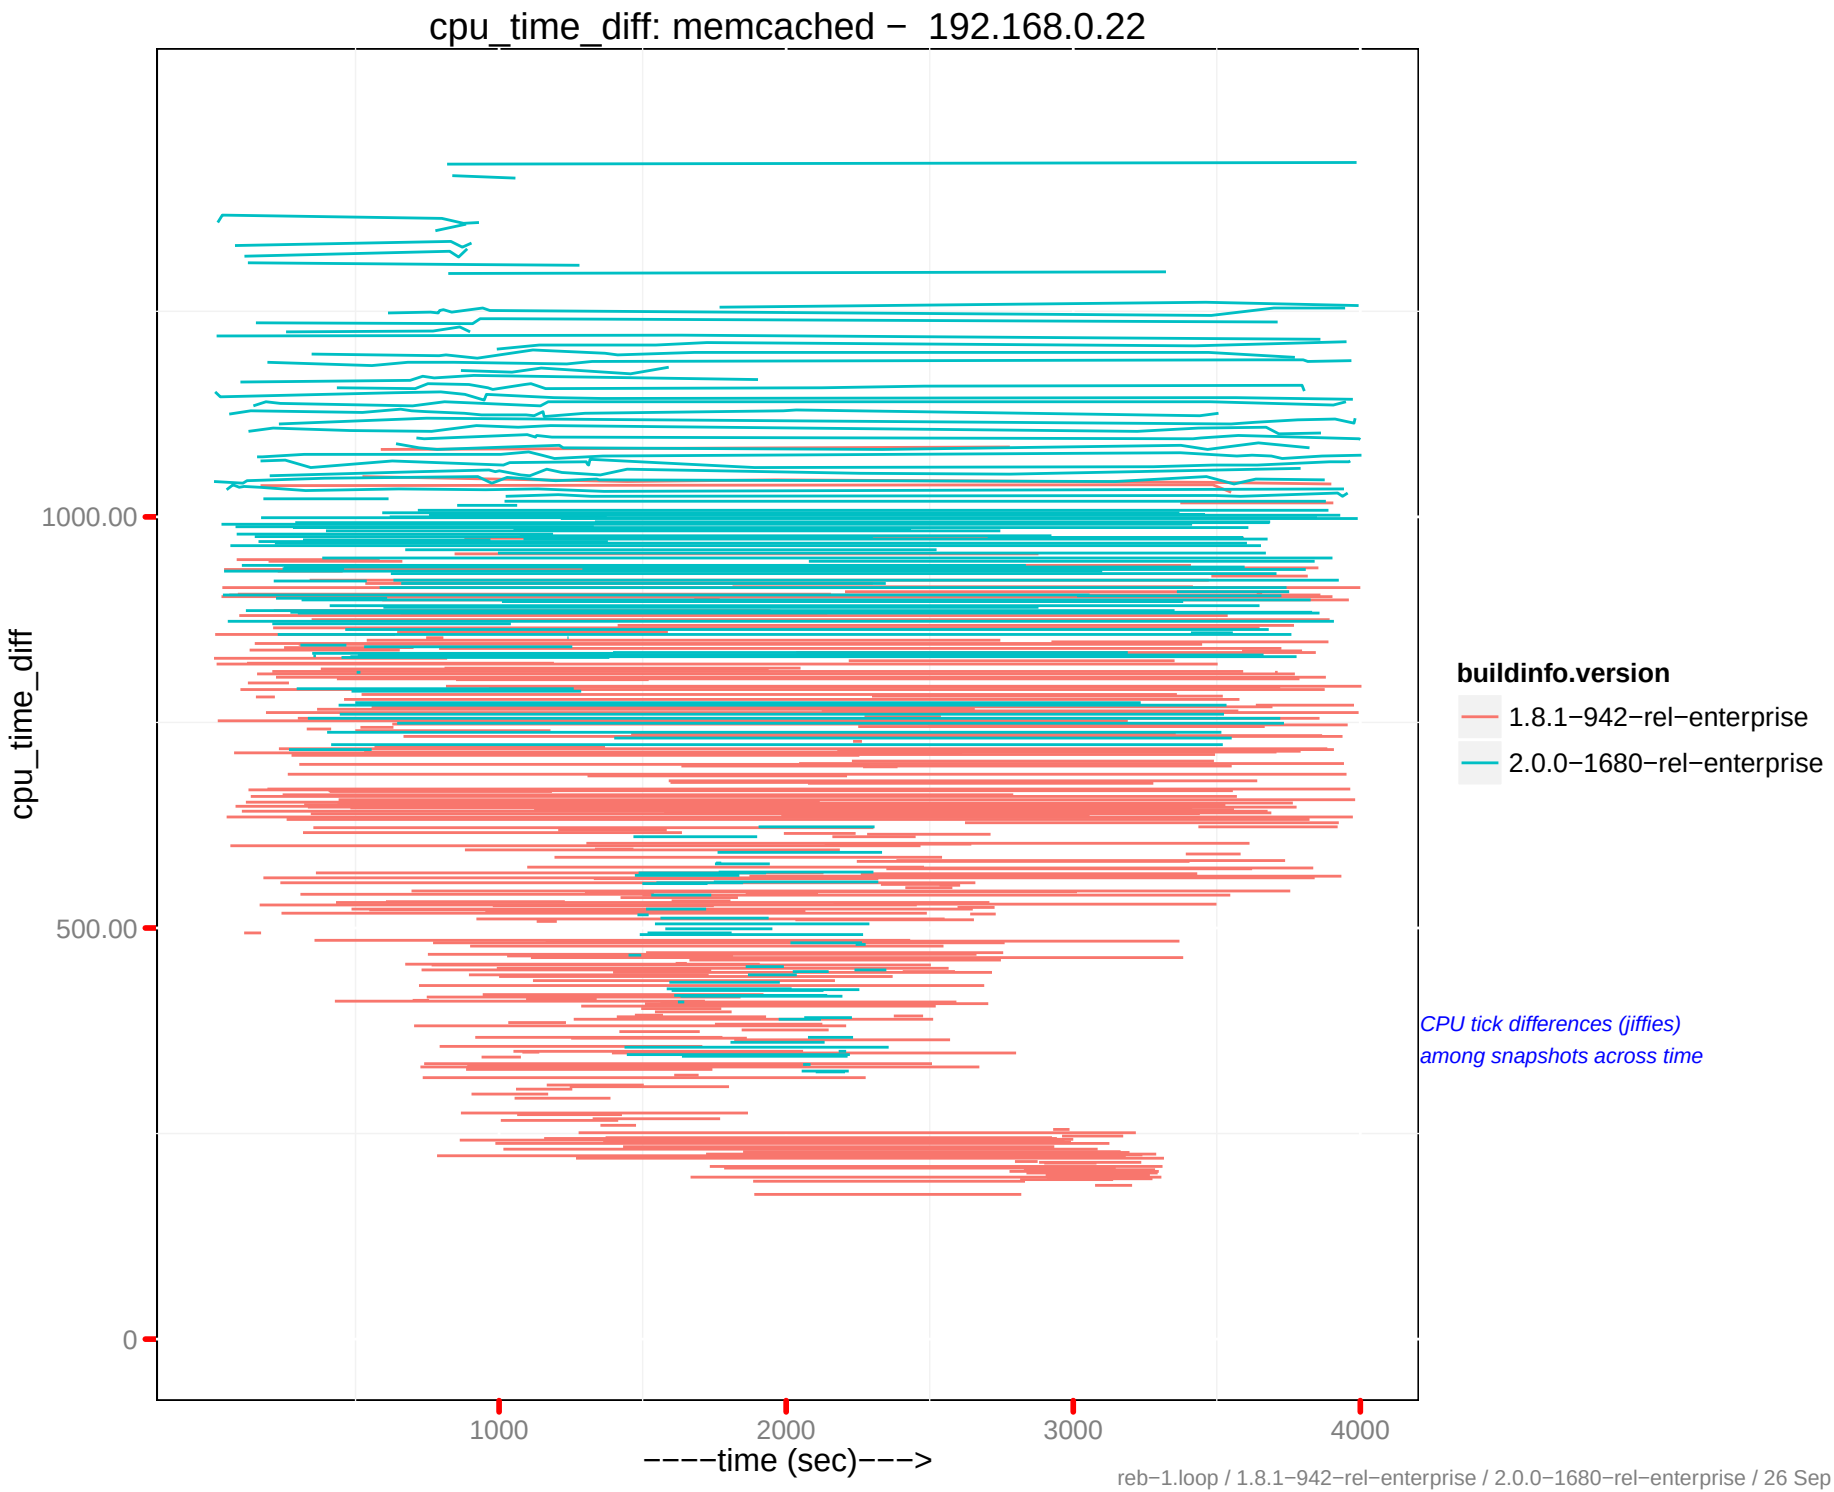
cpu_time_diff: memcached − 192.168.0.22
cpu_time_diff
1000.00
500.00
0
1000
2000
3000
4000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

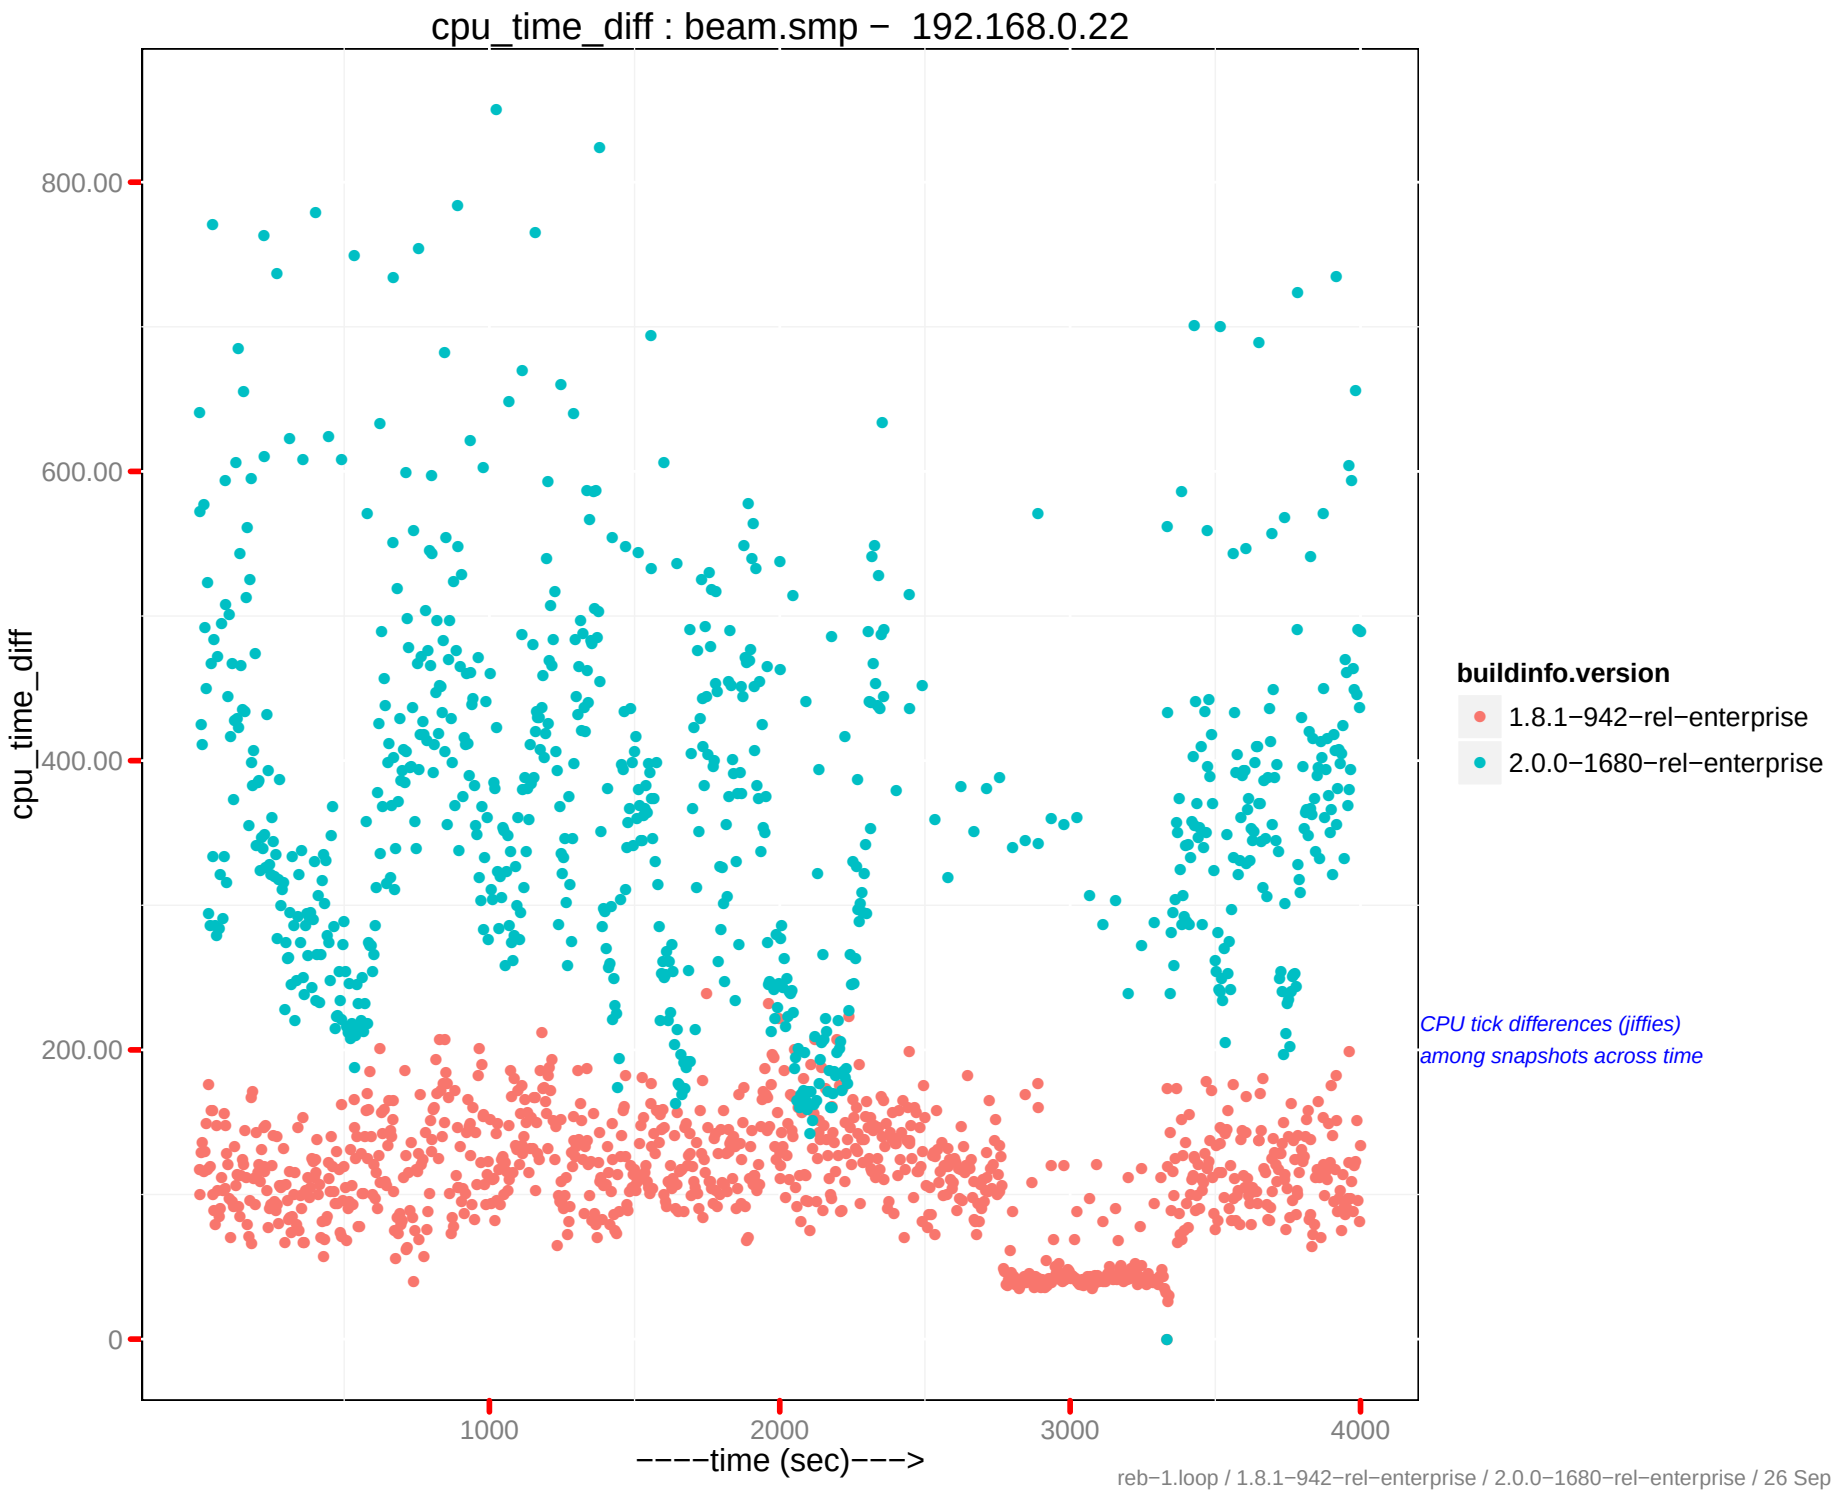
cpu_time_diff : beam.smp − 192.168.0.22
cpu_time_diff
800.00
600.00
400.00
200.00
0
1000
2000
3000
4000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

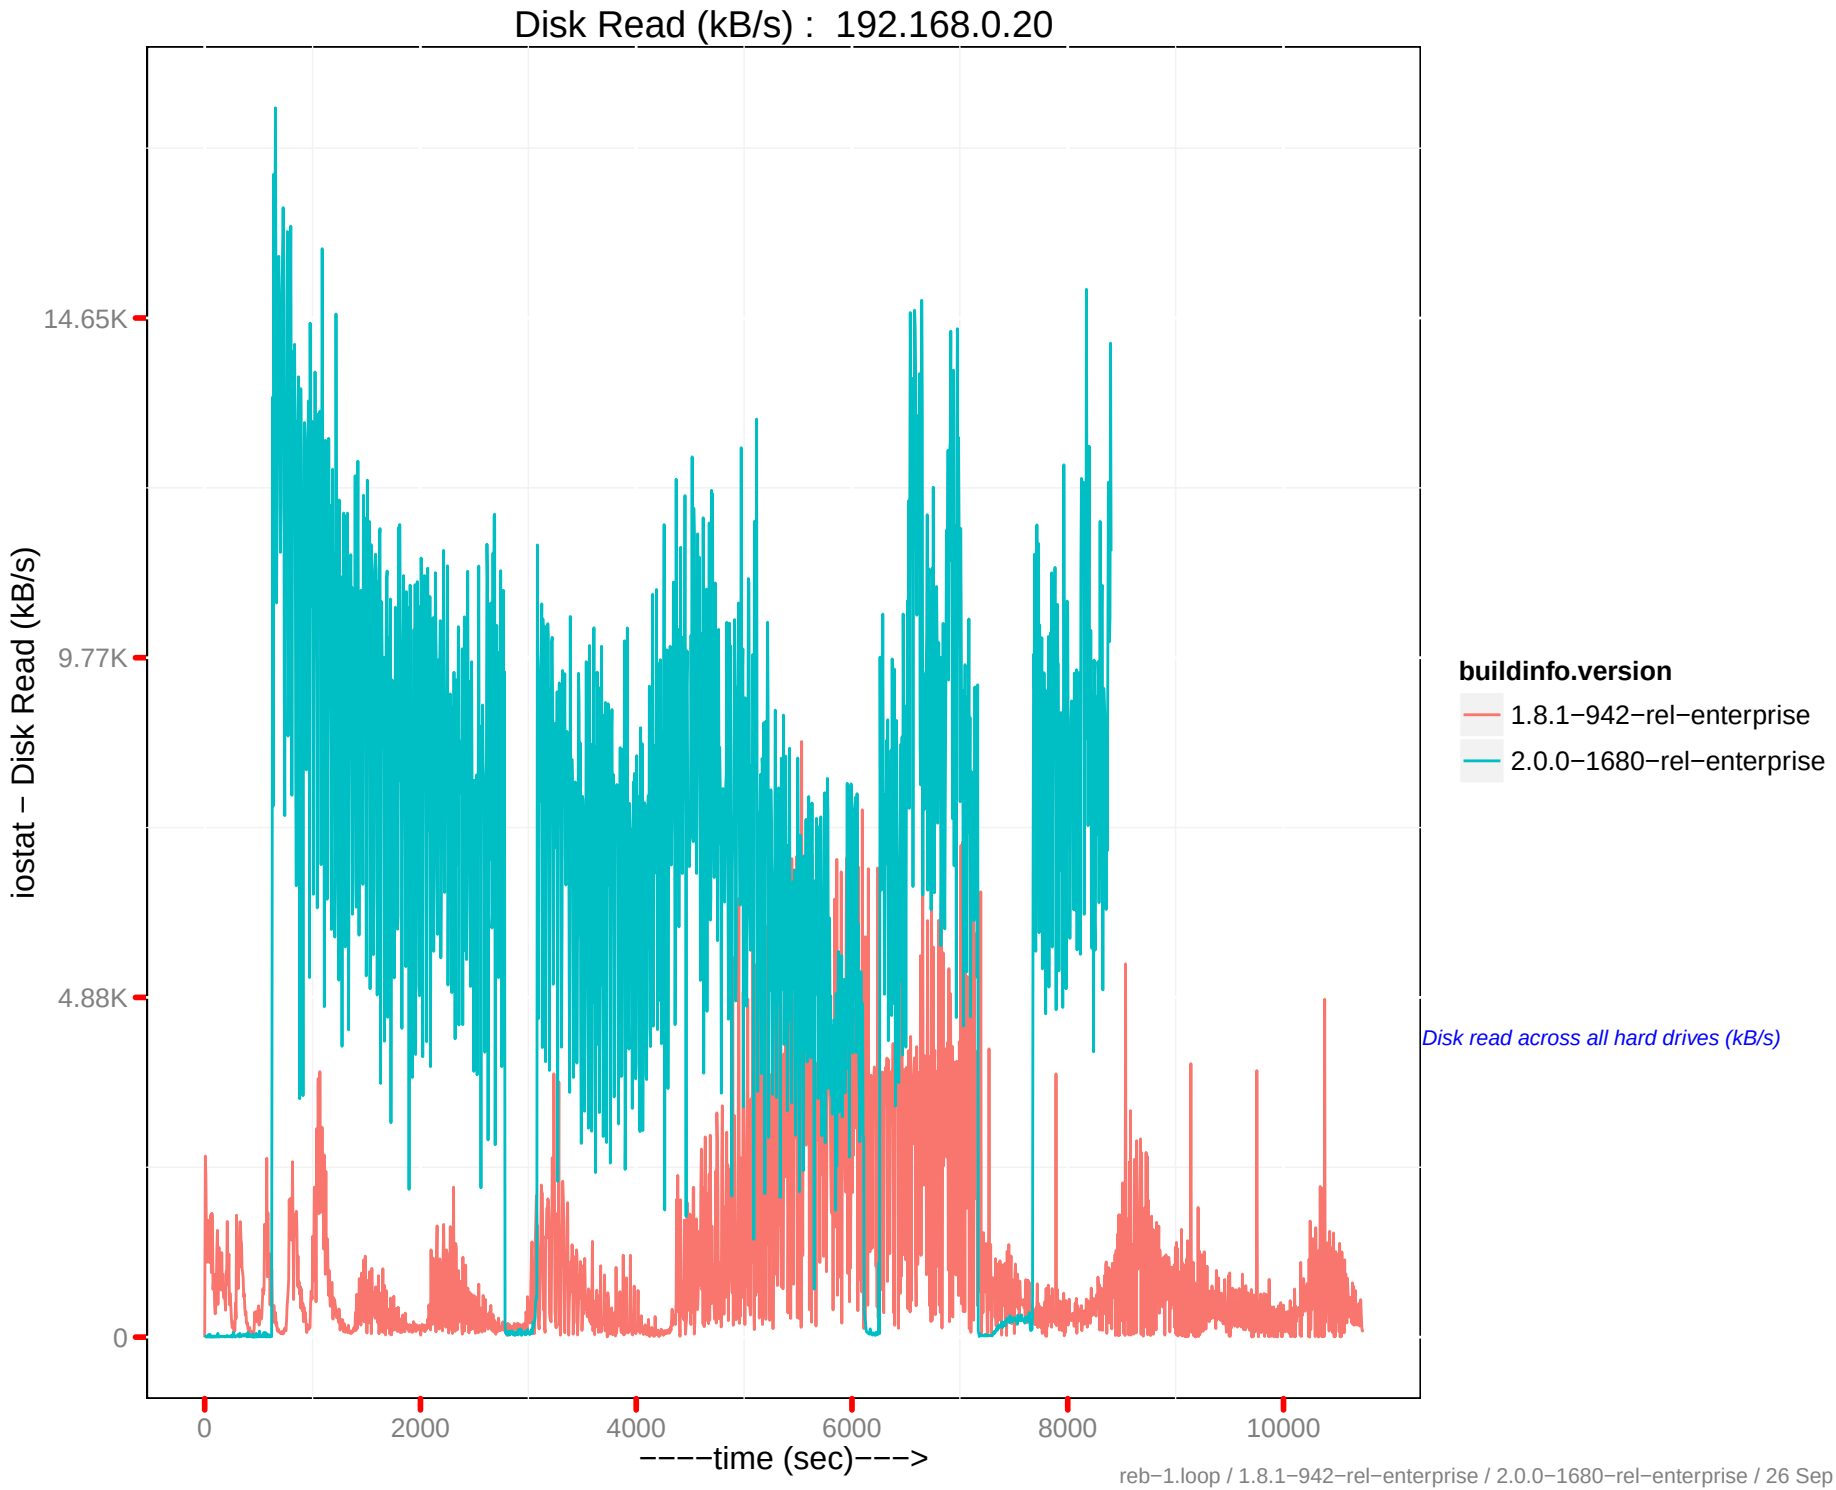
Disk Read (kB/s) : 192.168.0.20
iostat – Disk Read (kB/s)
14.65K
9.77K
4.88K
0
0
2000
4000
6000
8000
10000
––––time (sec)––––>
buildinfo.version
1.8.1–942–rel–enterprise
2.0.0–1680–rel–enterprise
Disk read across all hard drives (kB/s)
reb–1.loop / 1.8.1–942–rel–enterprise / 2.0.0–1680–rel–enterprise / 26 Sep 2012

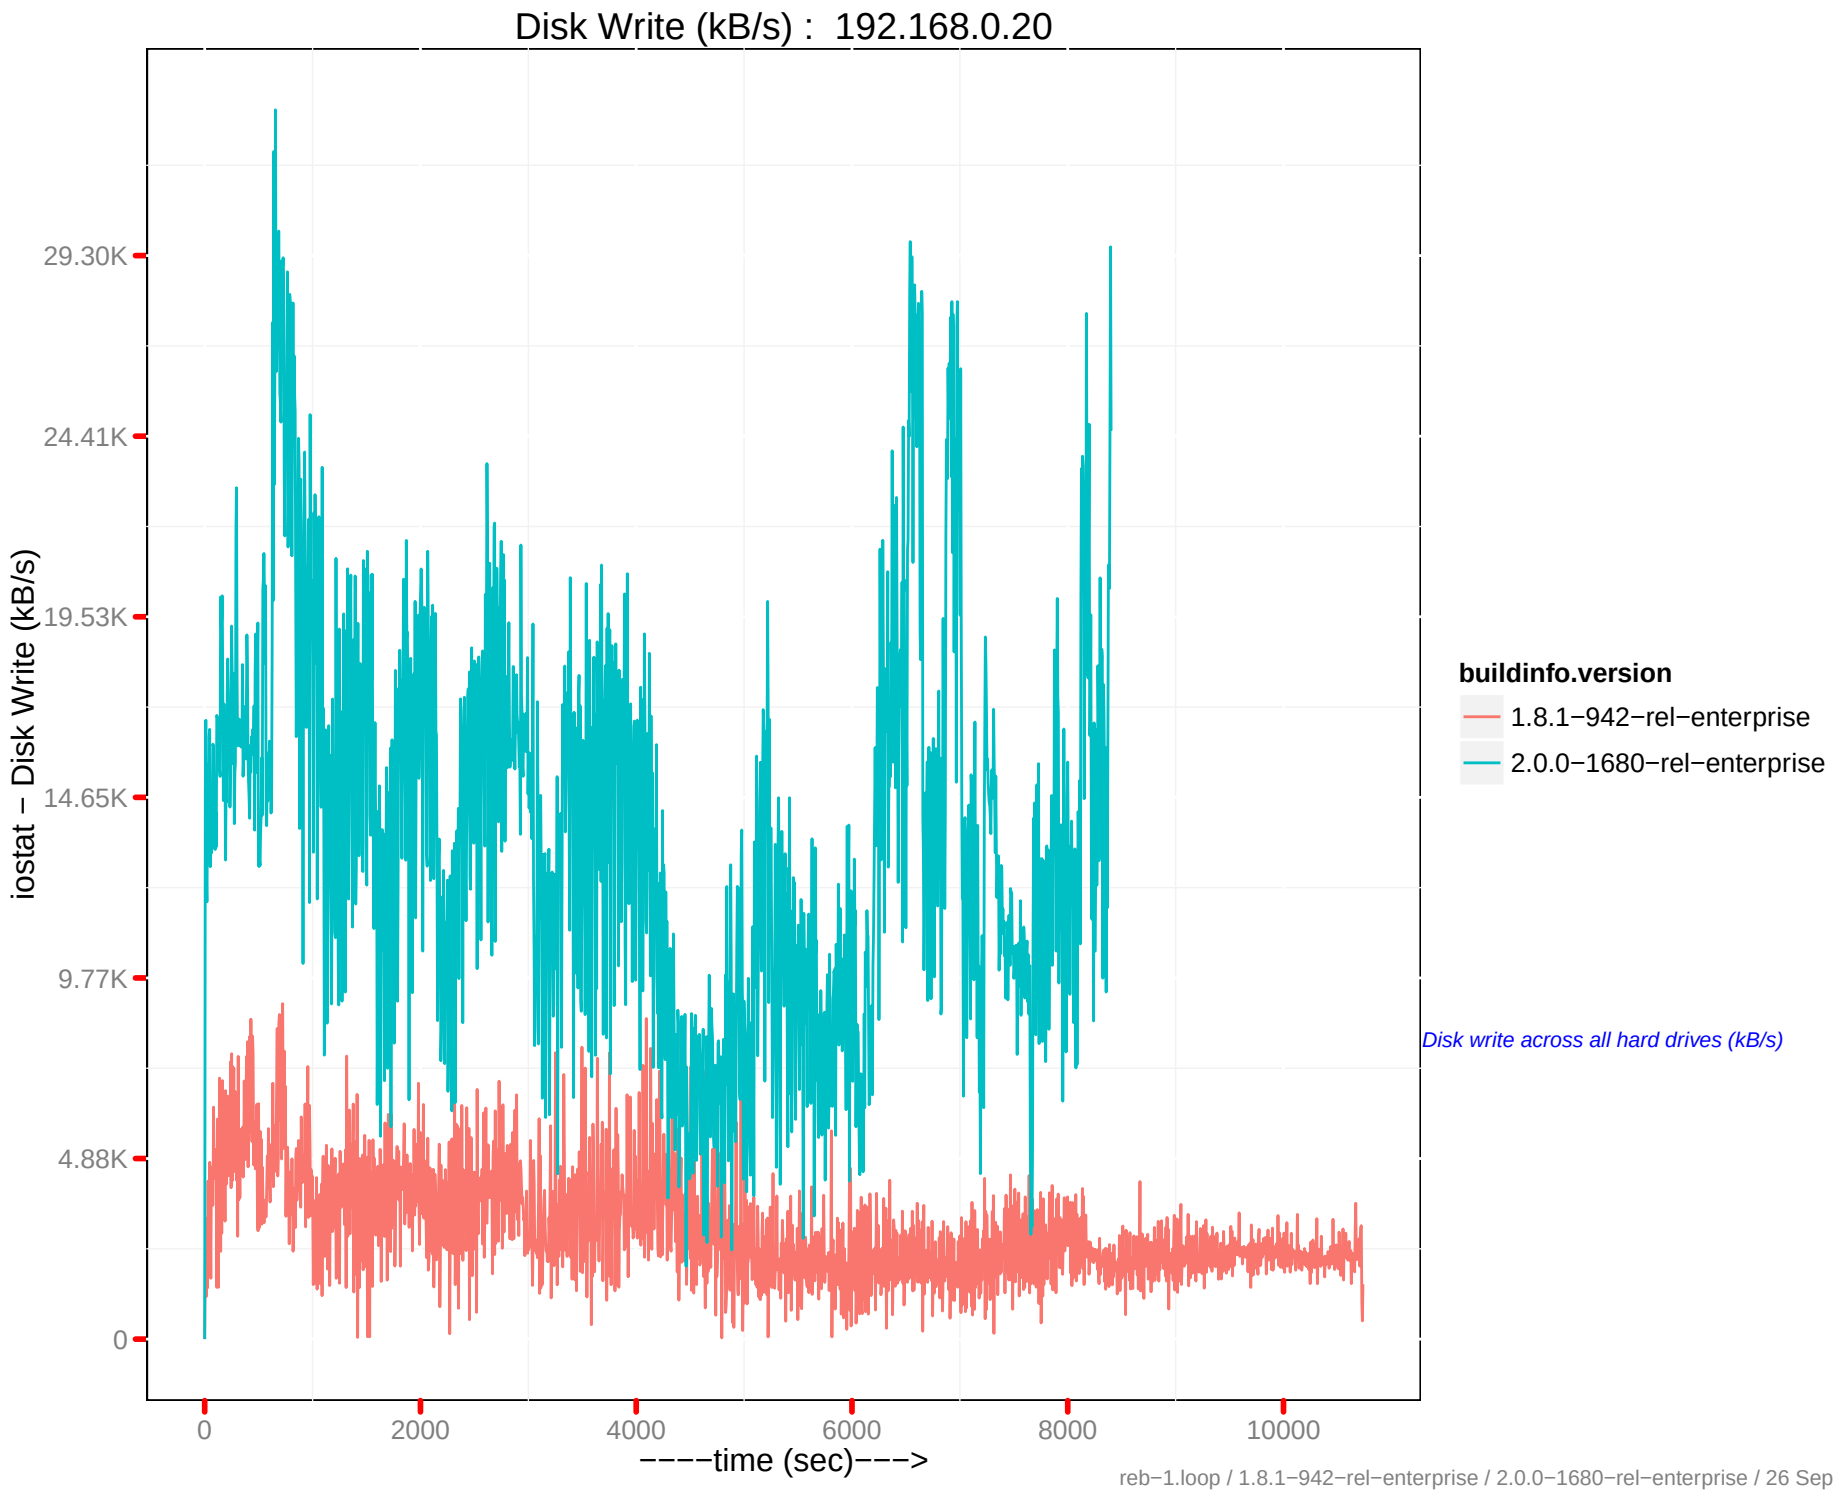
Disk Write (kB/s) : 192.168.0.20
iostat – Disk Write (kB/s)
29.30K
24.41K
19.53K
14.65K
9.77K
4.88K
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Disk write across all hard drives (kB/s)
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

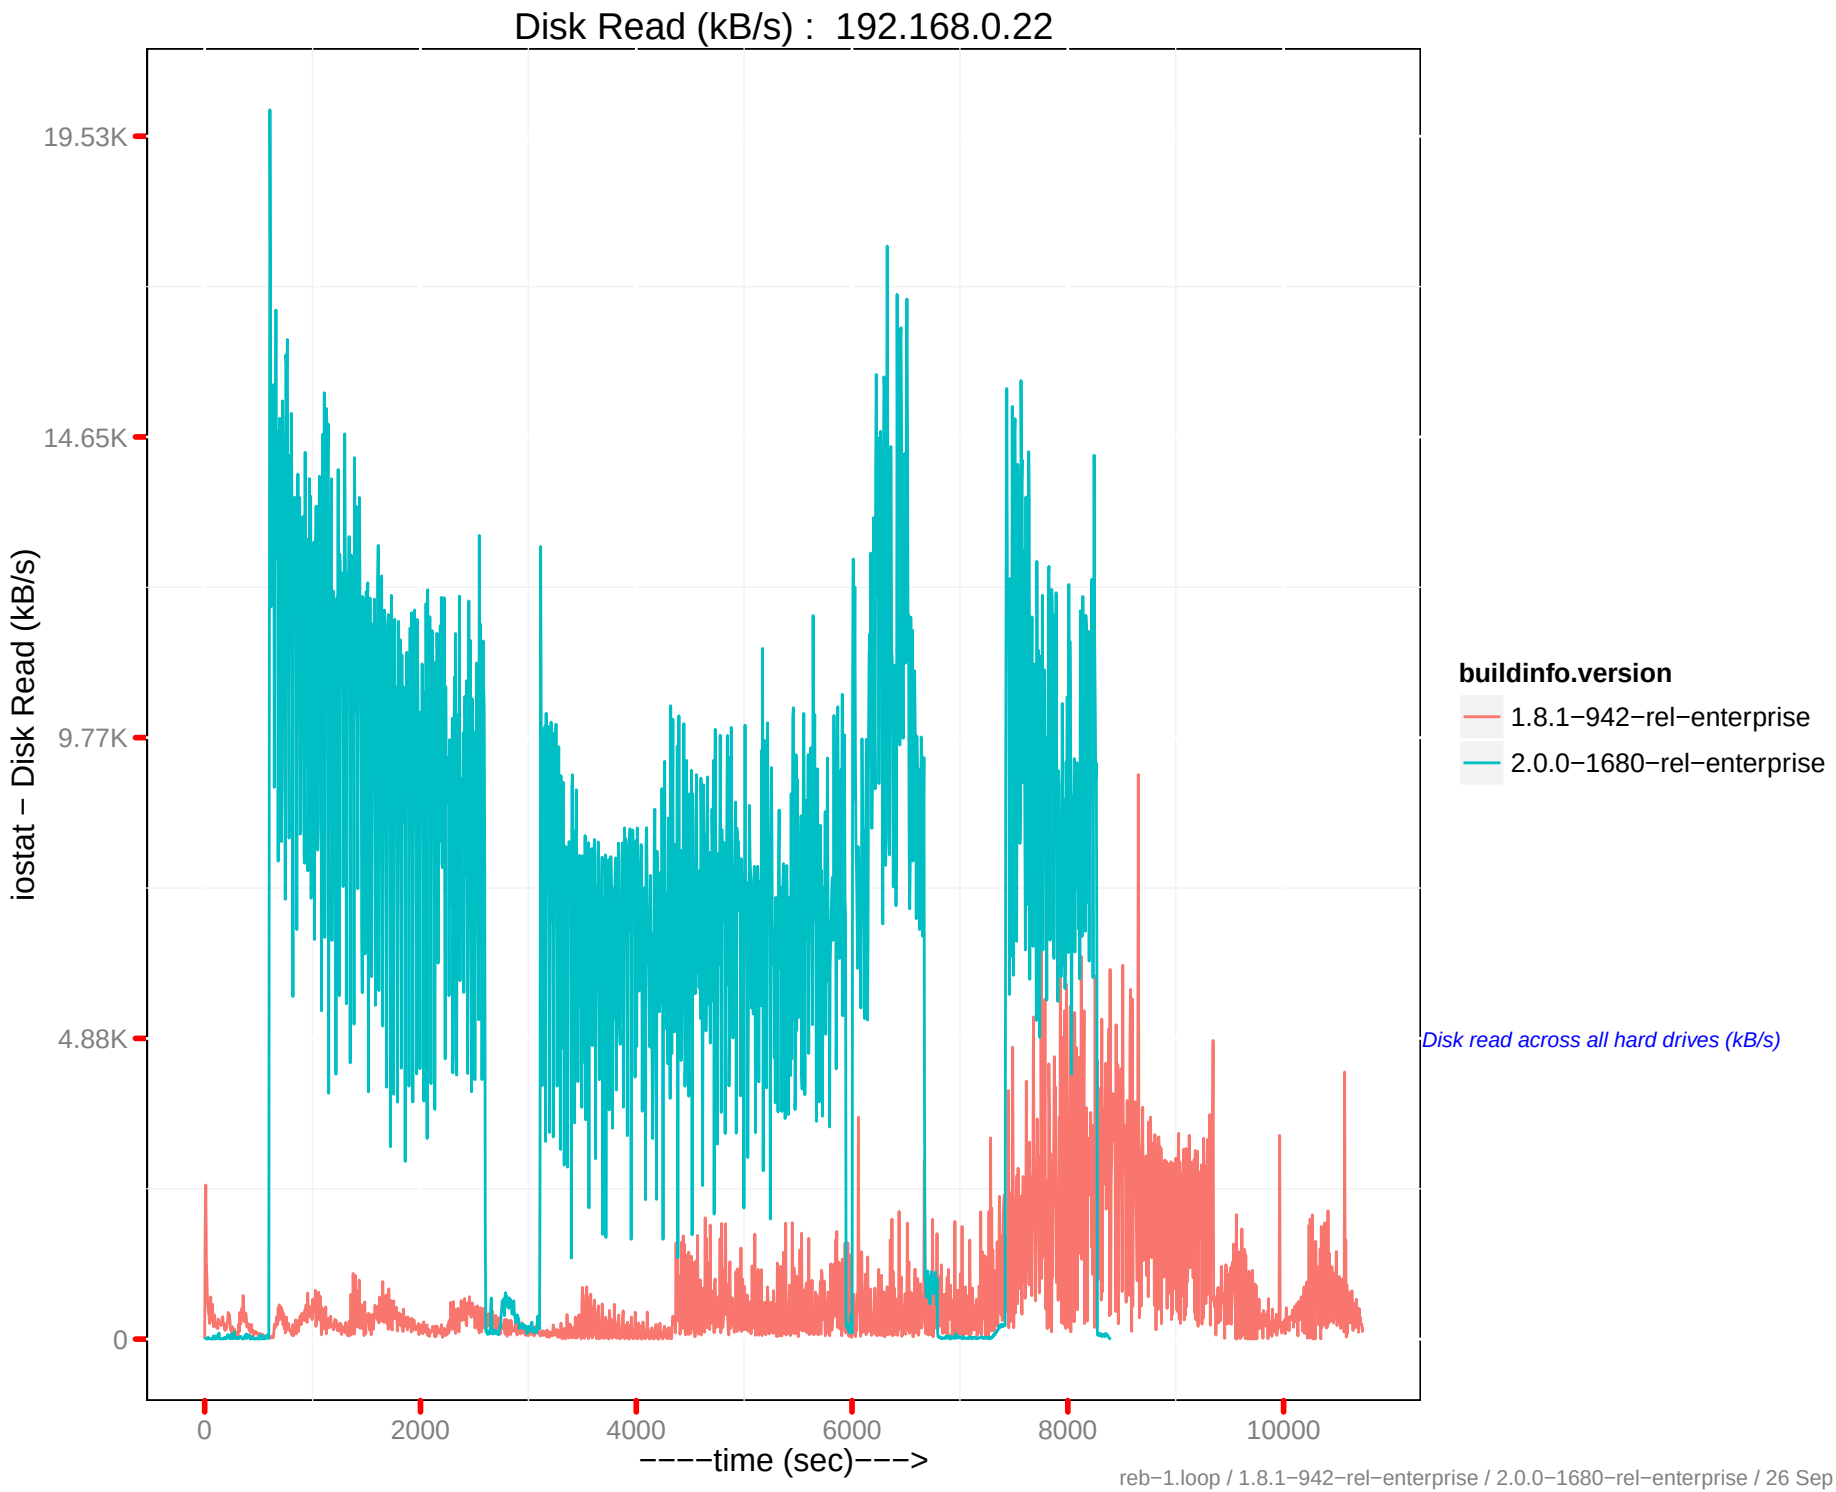
Disk Read (kB/s) : 192.168.0.22
19.53K
14.65K
9.77K
4.88K
0
iostat – Disk Read (kB/s)
0
2000
4000
6000
8000
10000
––––time (sec)––––>
buildinfo.version
1.8.1–942–rel–enterprise
2.0.0–1680–rel–enterprise
Disk read across all hard drives (kB/s)
reb–1.loop / 1.8.1–942–rel–enterprise / 2.0.0–1680–rel–enterprise / 26 Sep 2012

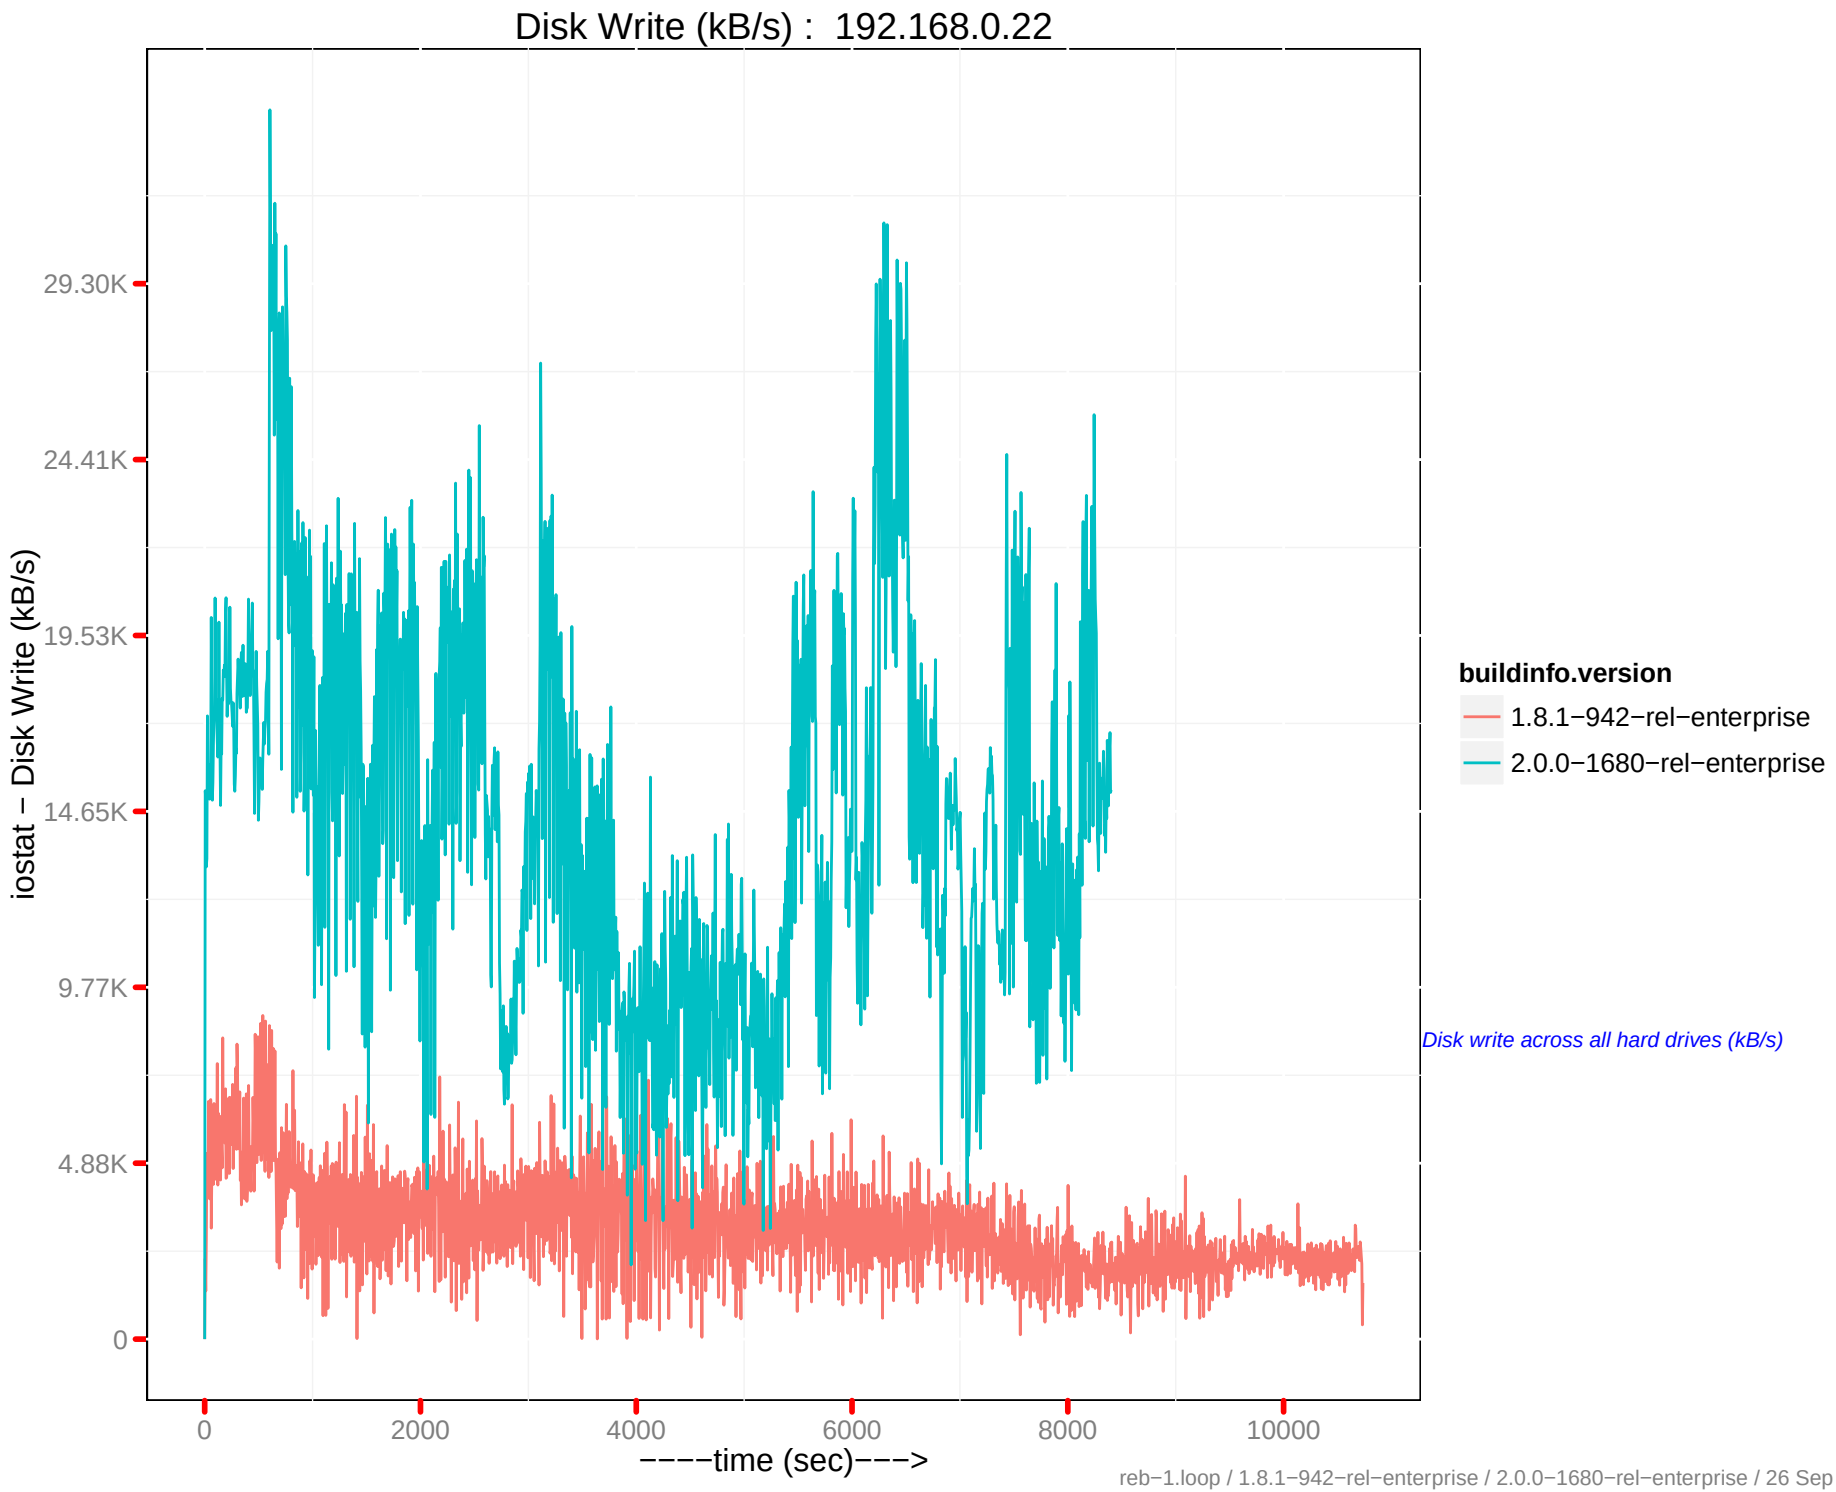
Disk Write (kB/s) : 192.168.0.22
iostat − Disk Write (kB/s)
29.30K
24.41K
19.53K
14.65K
9.77K
4.88K
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
Disk write across all hard drives (kB/s)
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

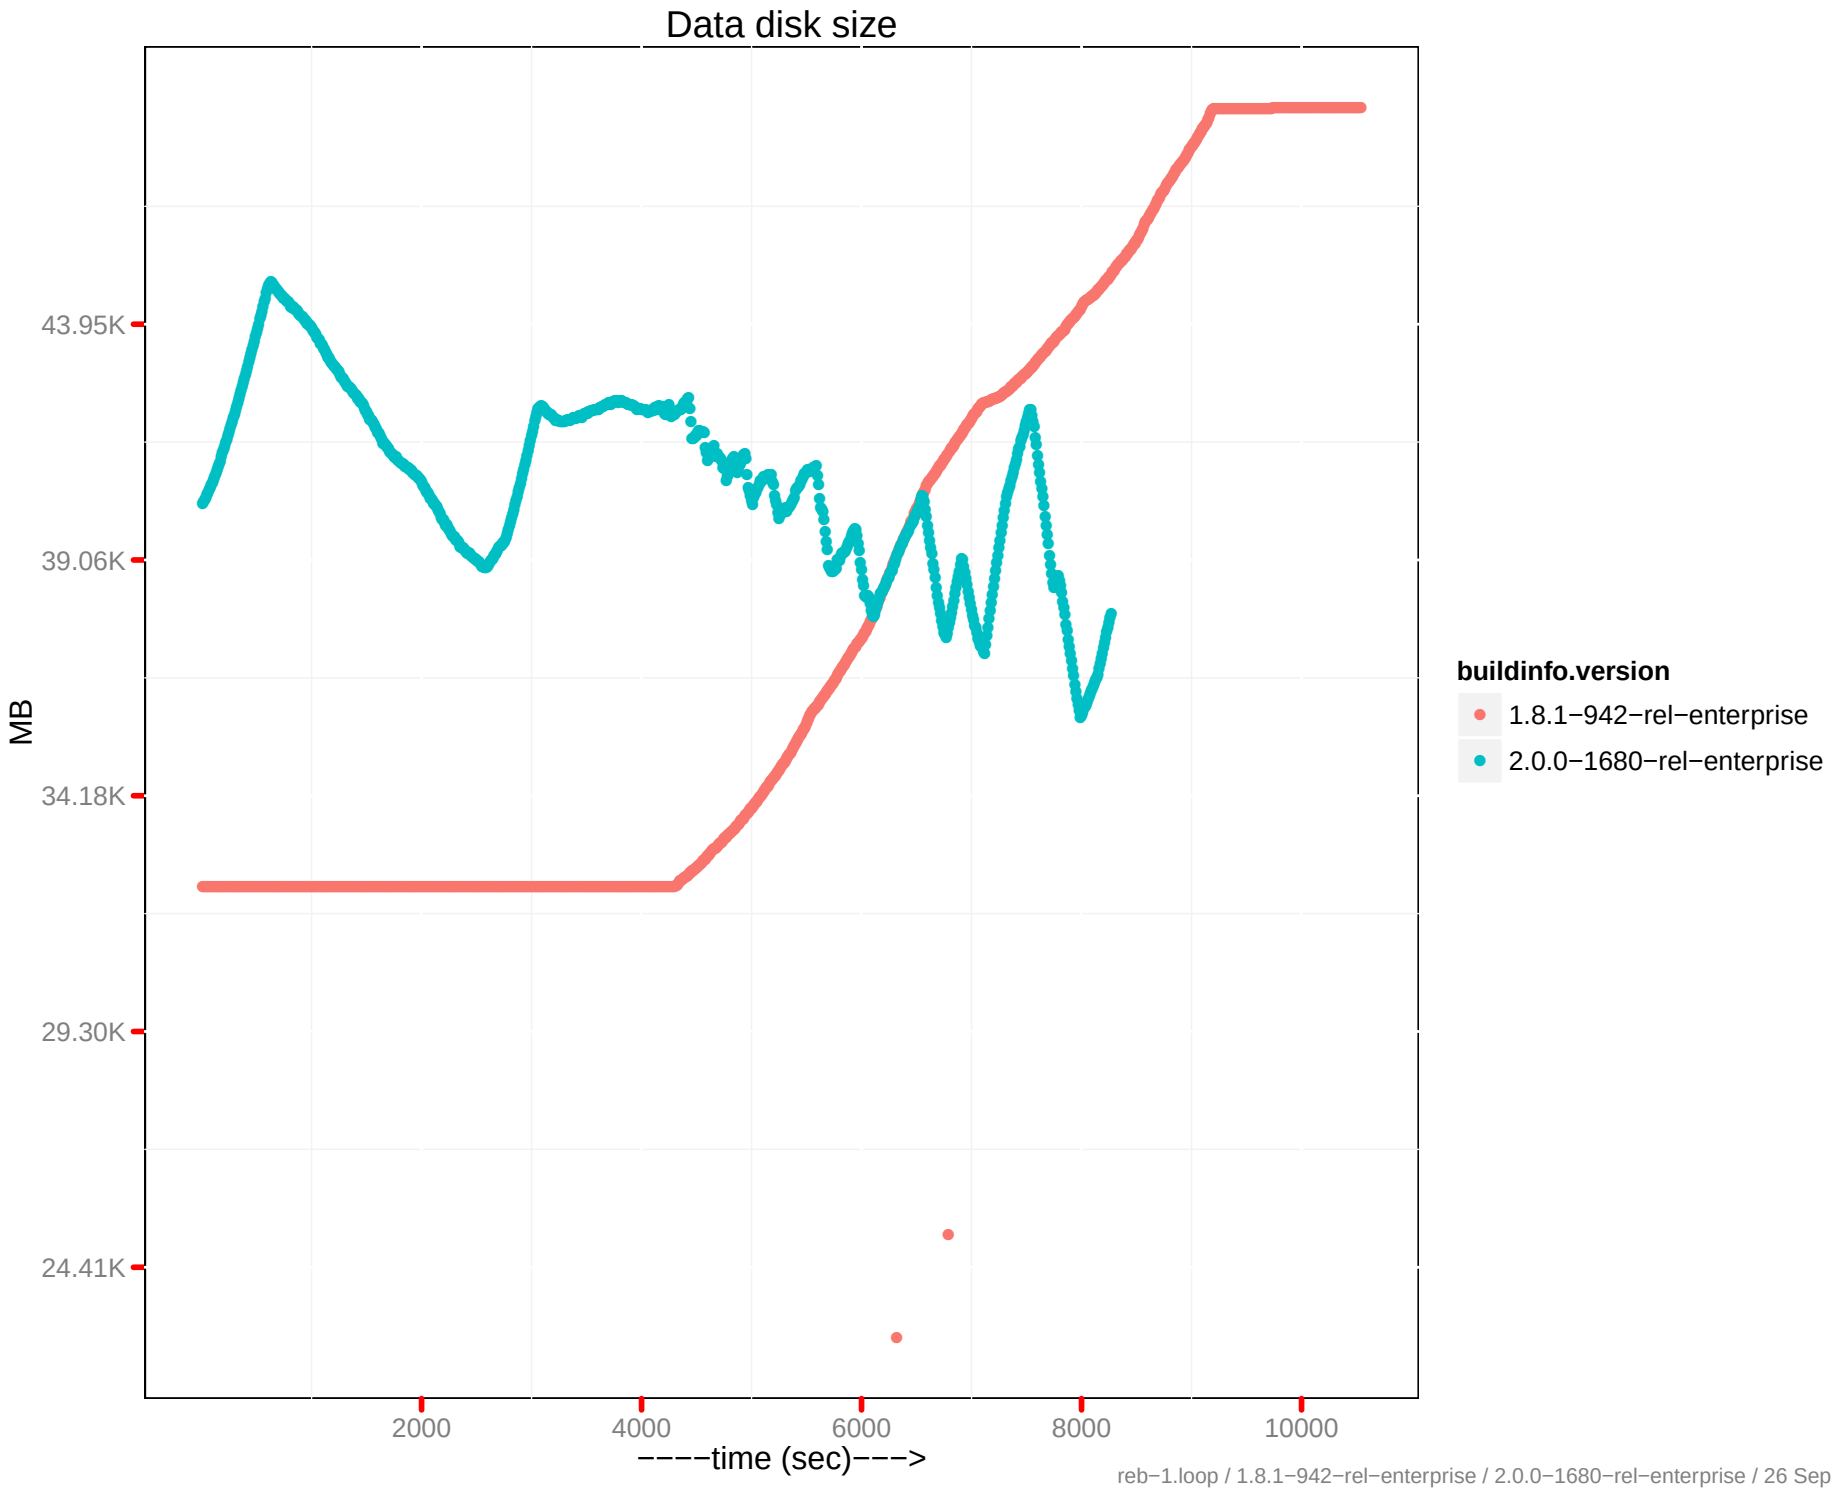
Data disk size
43.95K
39.06K
34.18K
29.30K
24.41K
MB
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

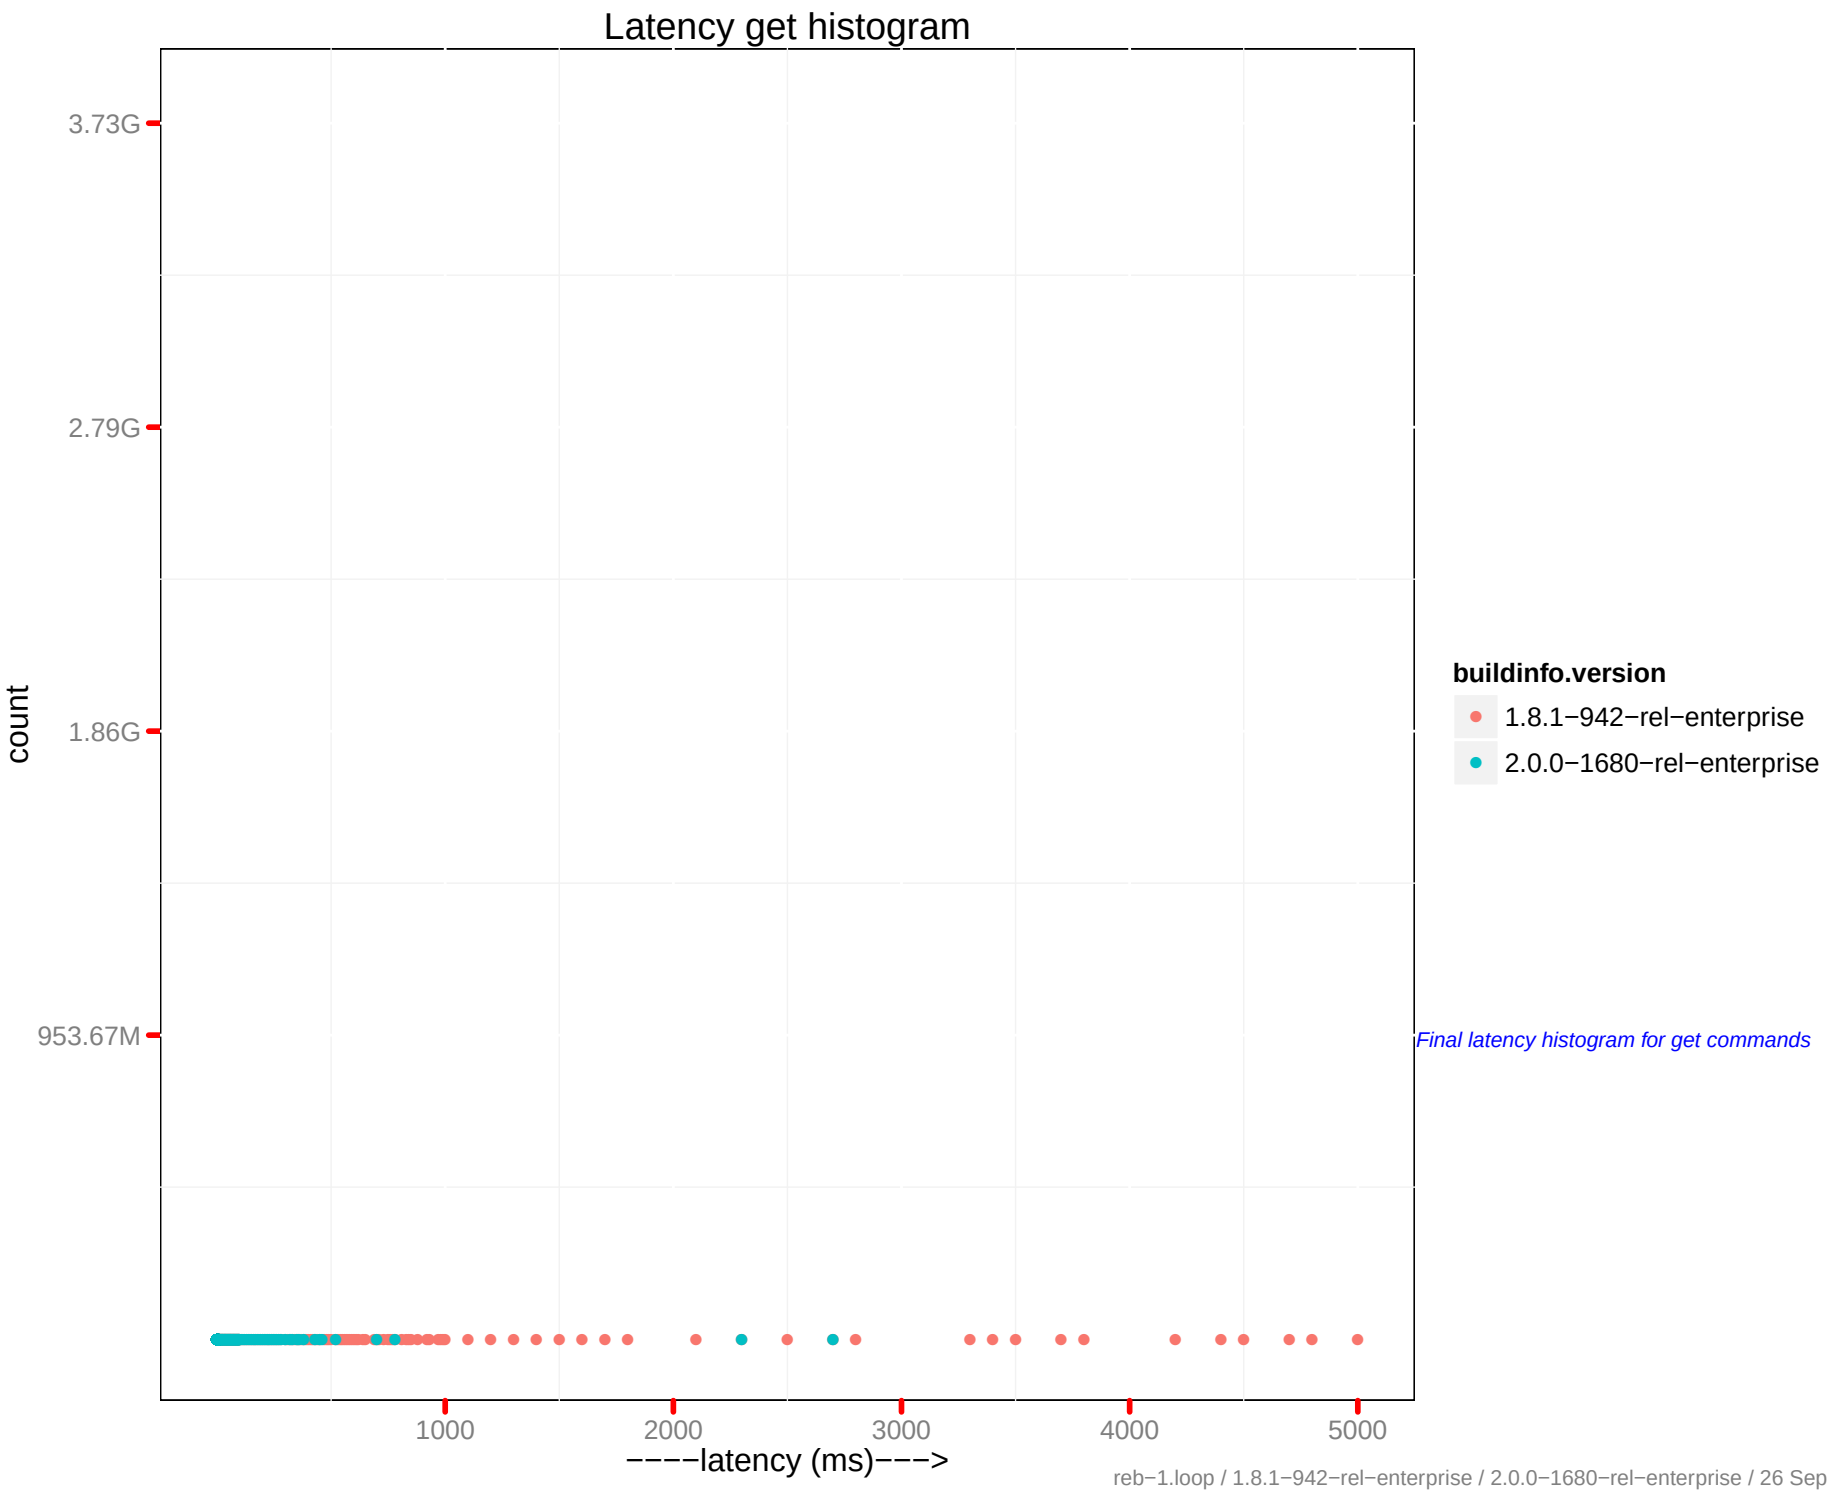
Latency get histogram
3.73G
2.79G
1.86G
953.67M
count
1000
2000
3000
4000
5000
----latency (ms)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Final latency histogram for get commands
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

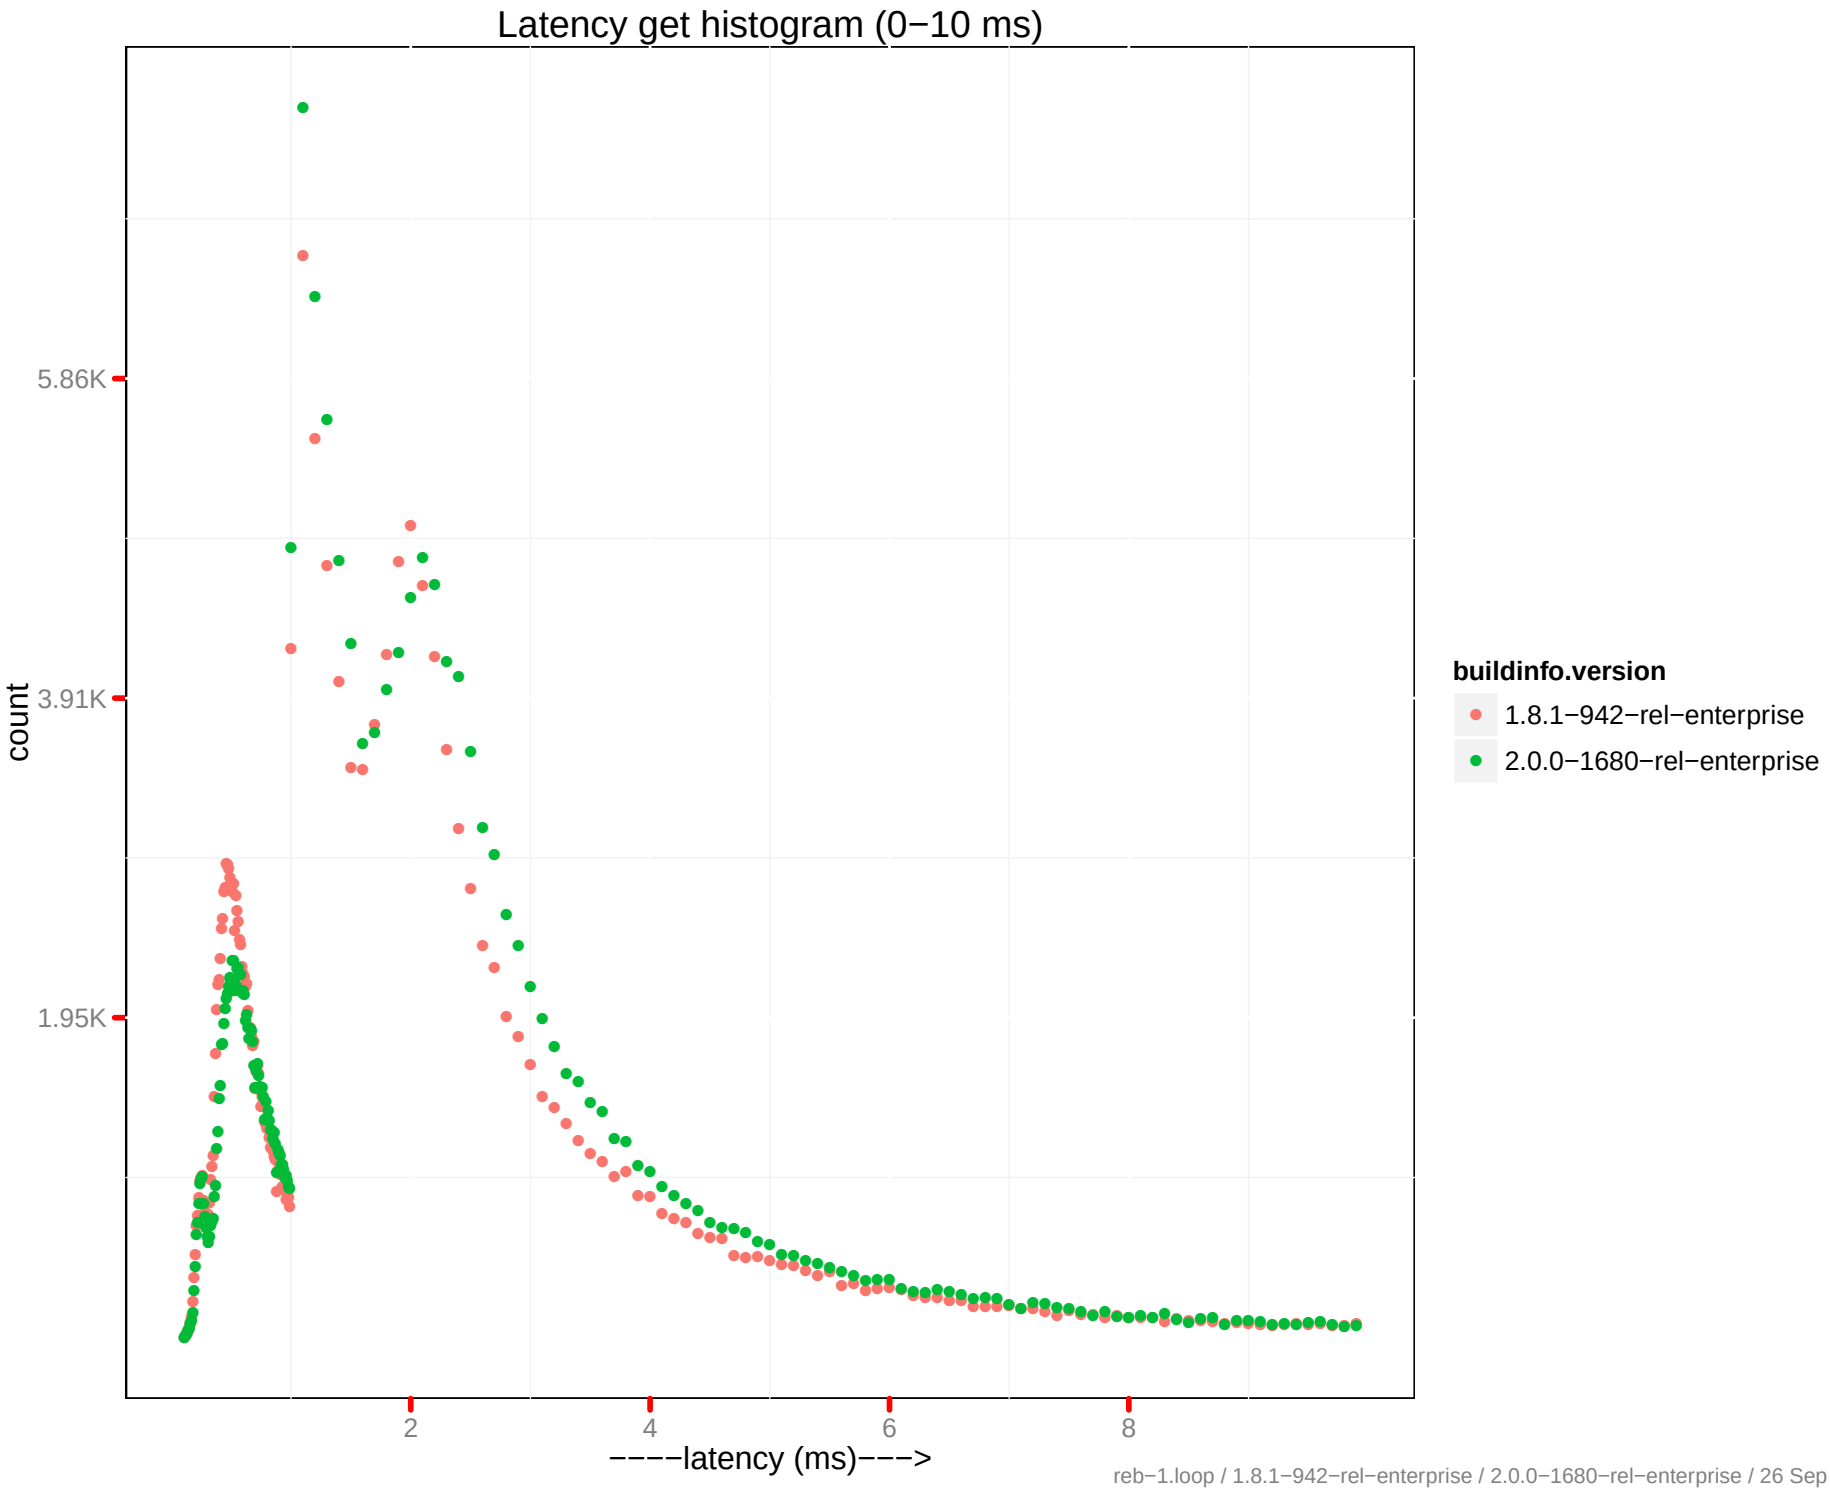
Latency get histogram (0−10 ms)
count
5.86K
3.91K
1.95K
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

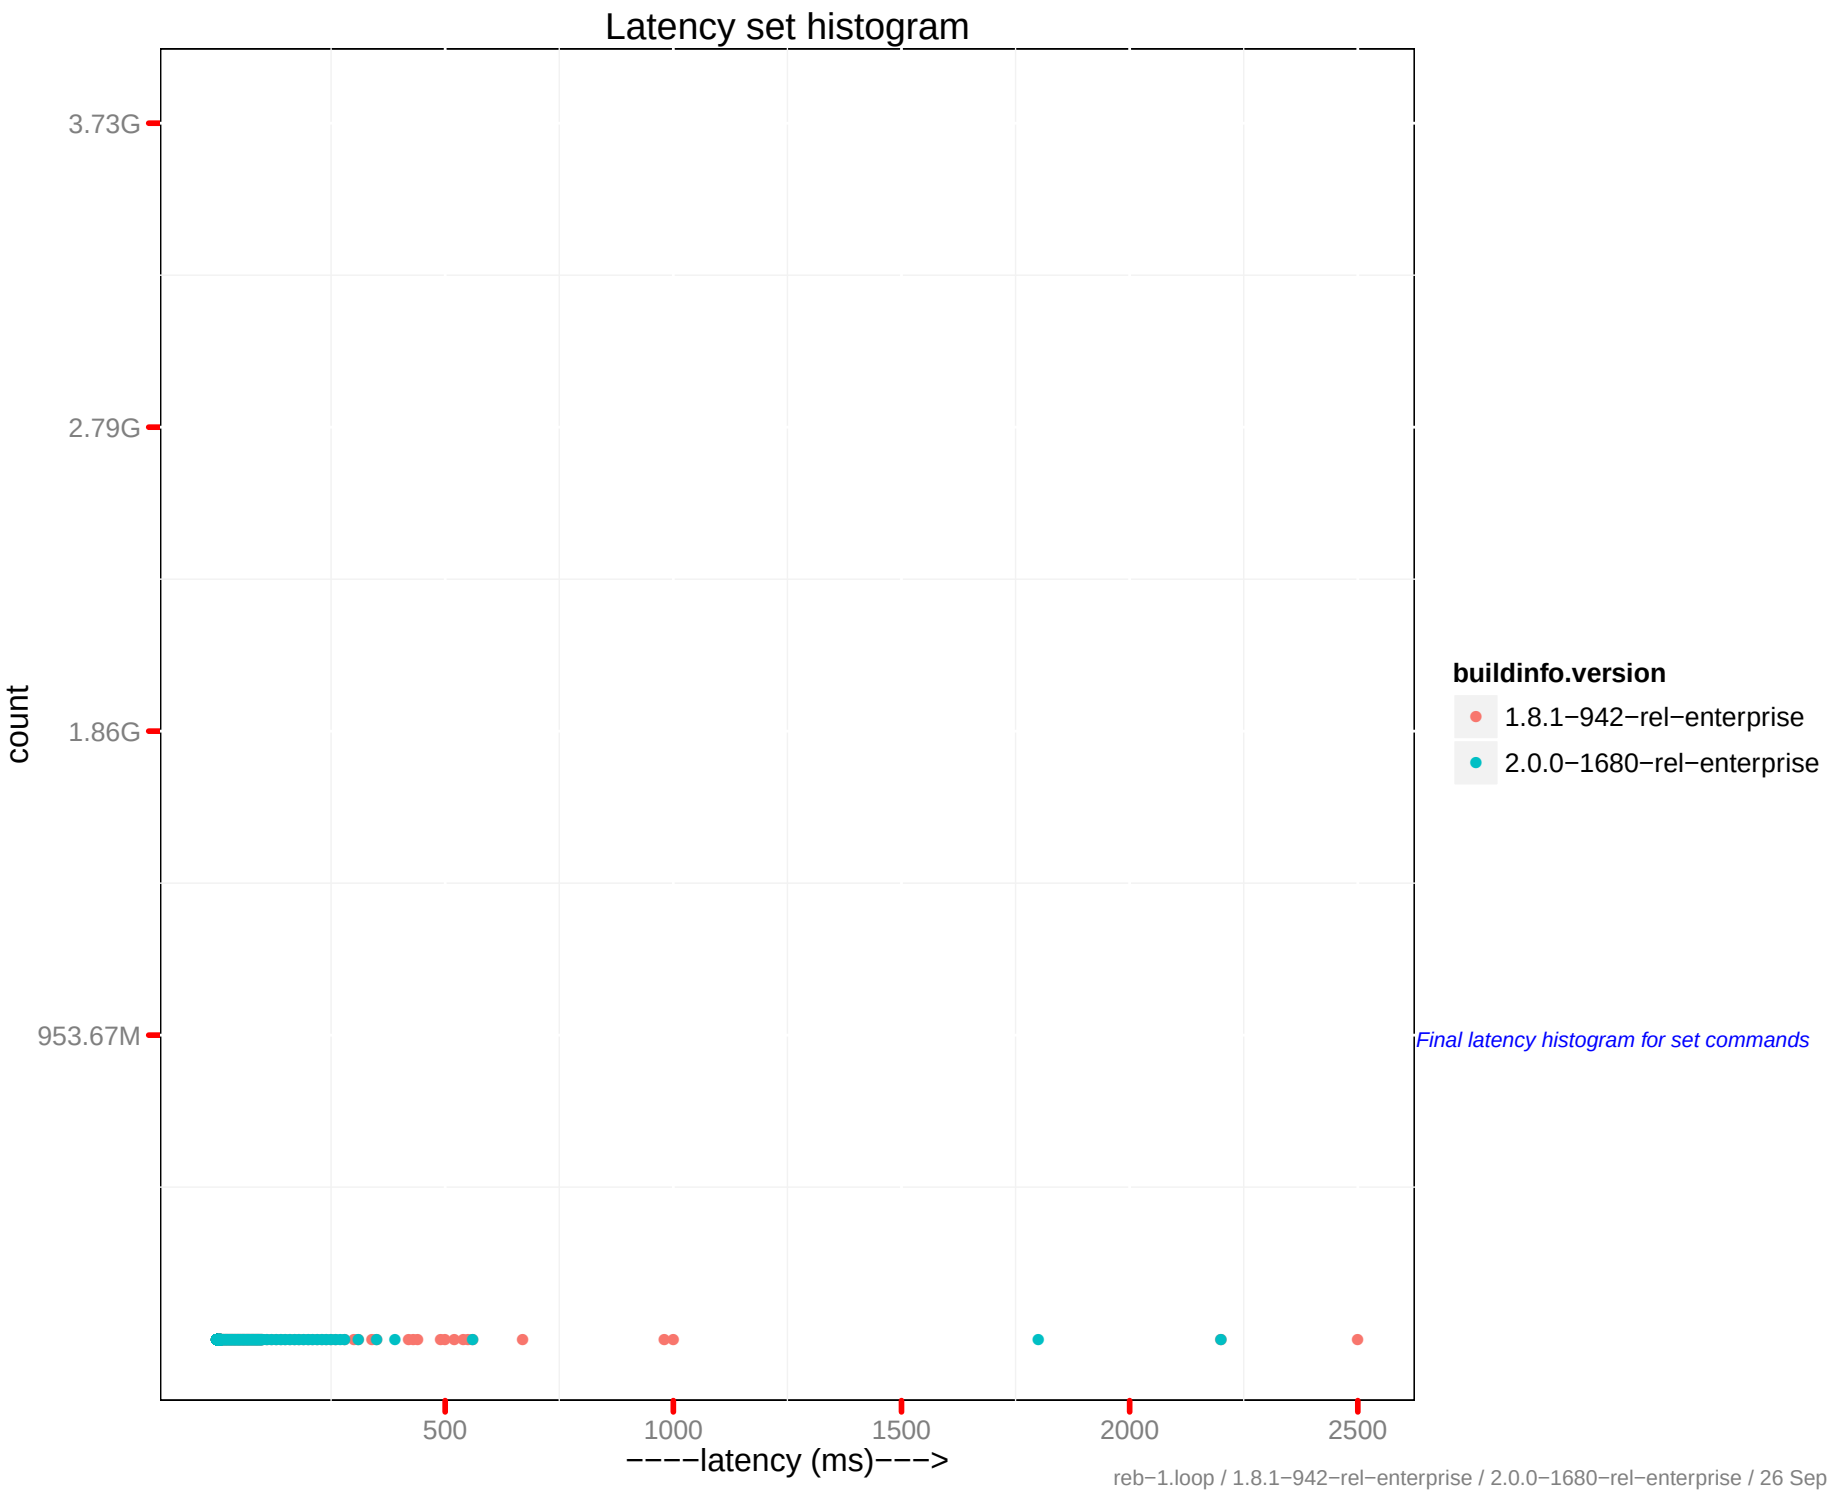
Latency set histogram
3.73G
2.79G
1.86G
953.67M
count
500
1000
1500
2000
2500
----latency (ms)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
Final latency histogram for set commands
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

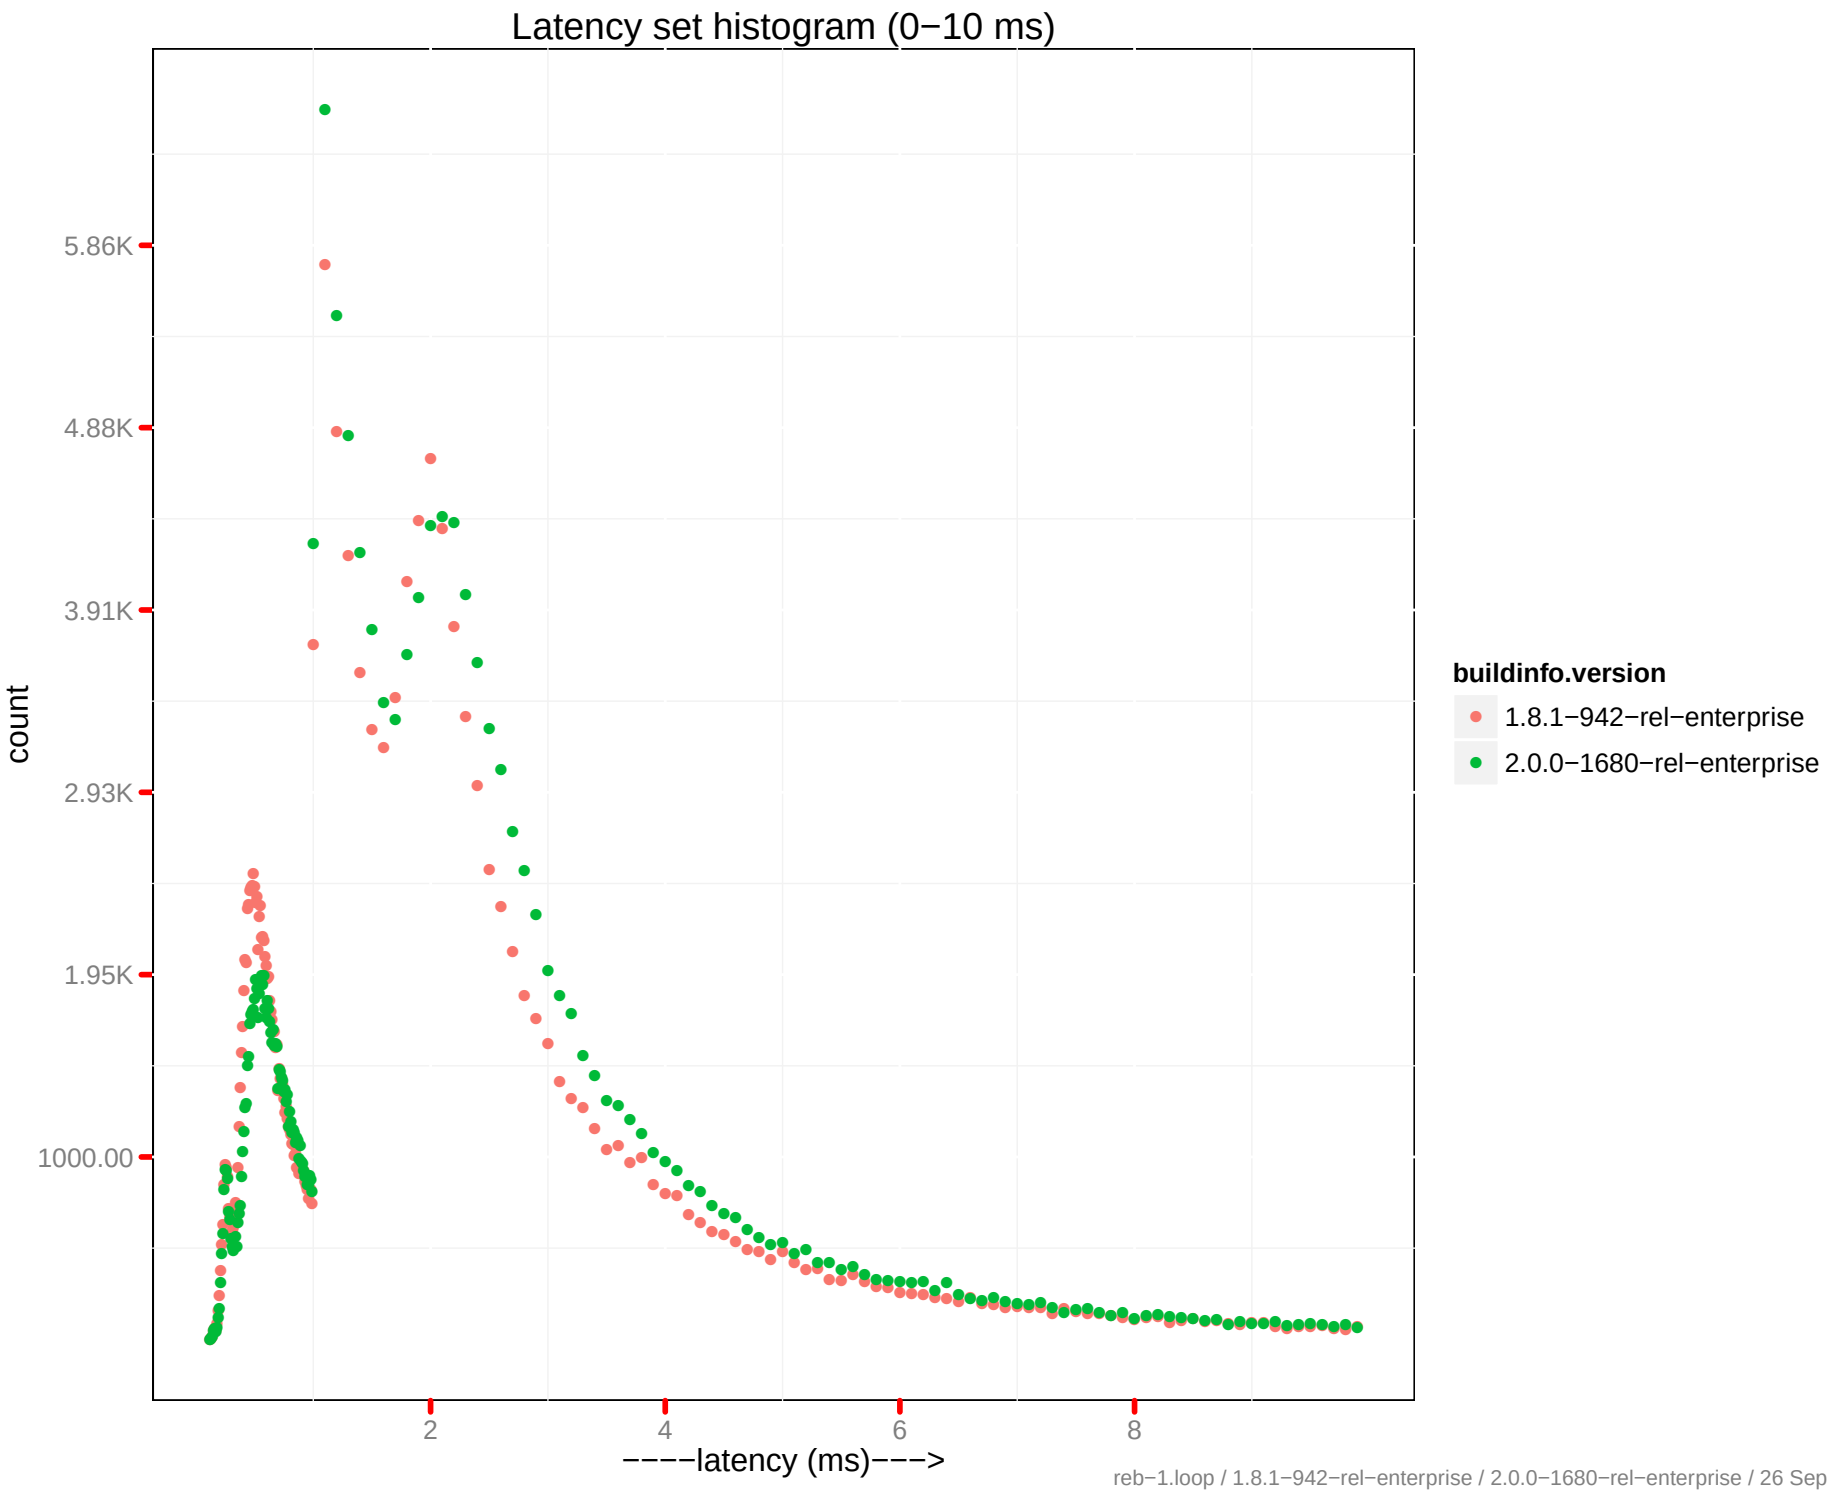
Latency set histogram (0−10 ms)
count
5.86K
4.88K
3.91K
2.93K
1.95K
1000.00
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

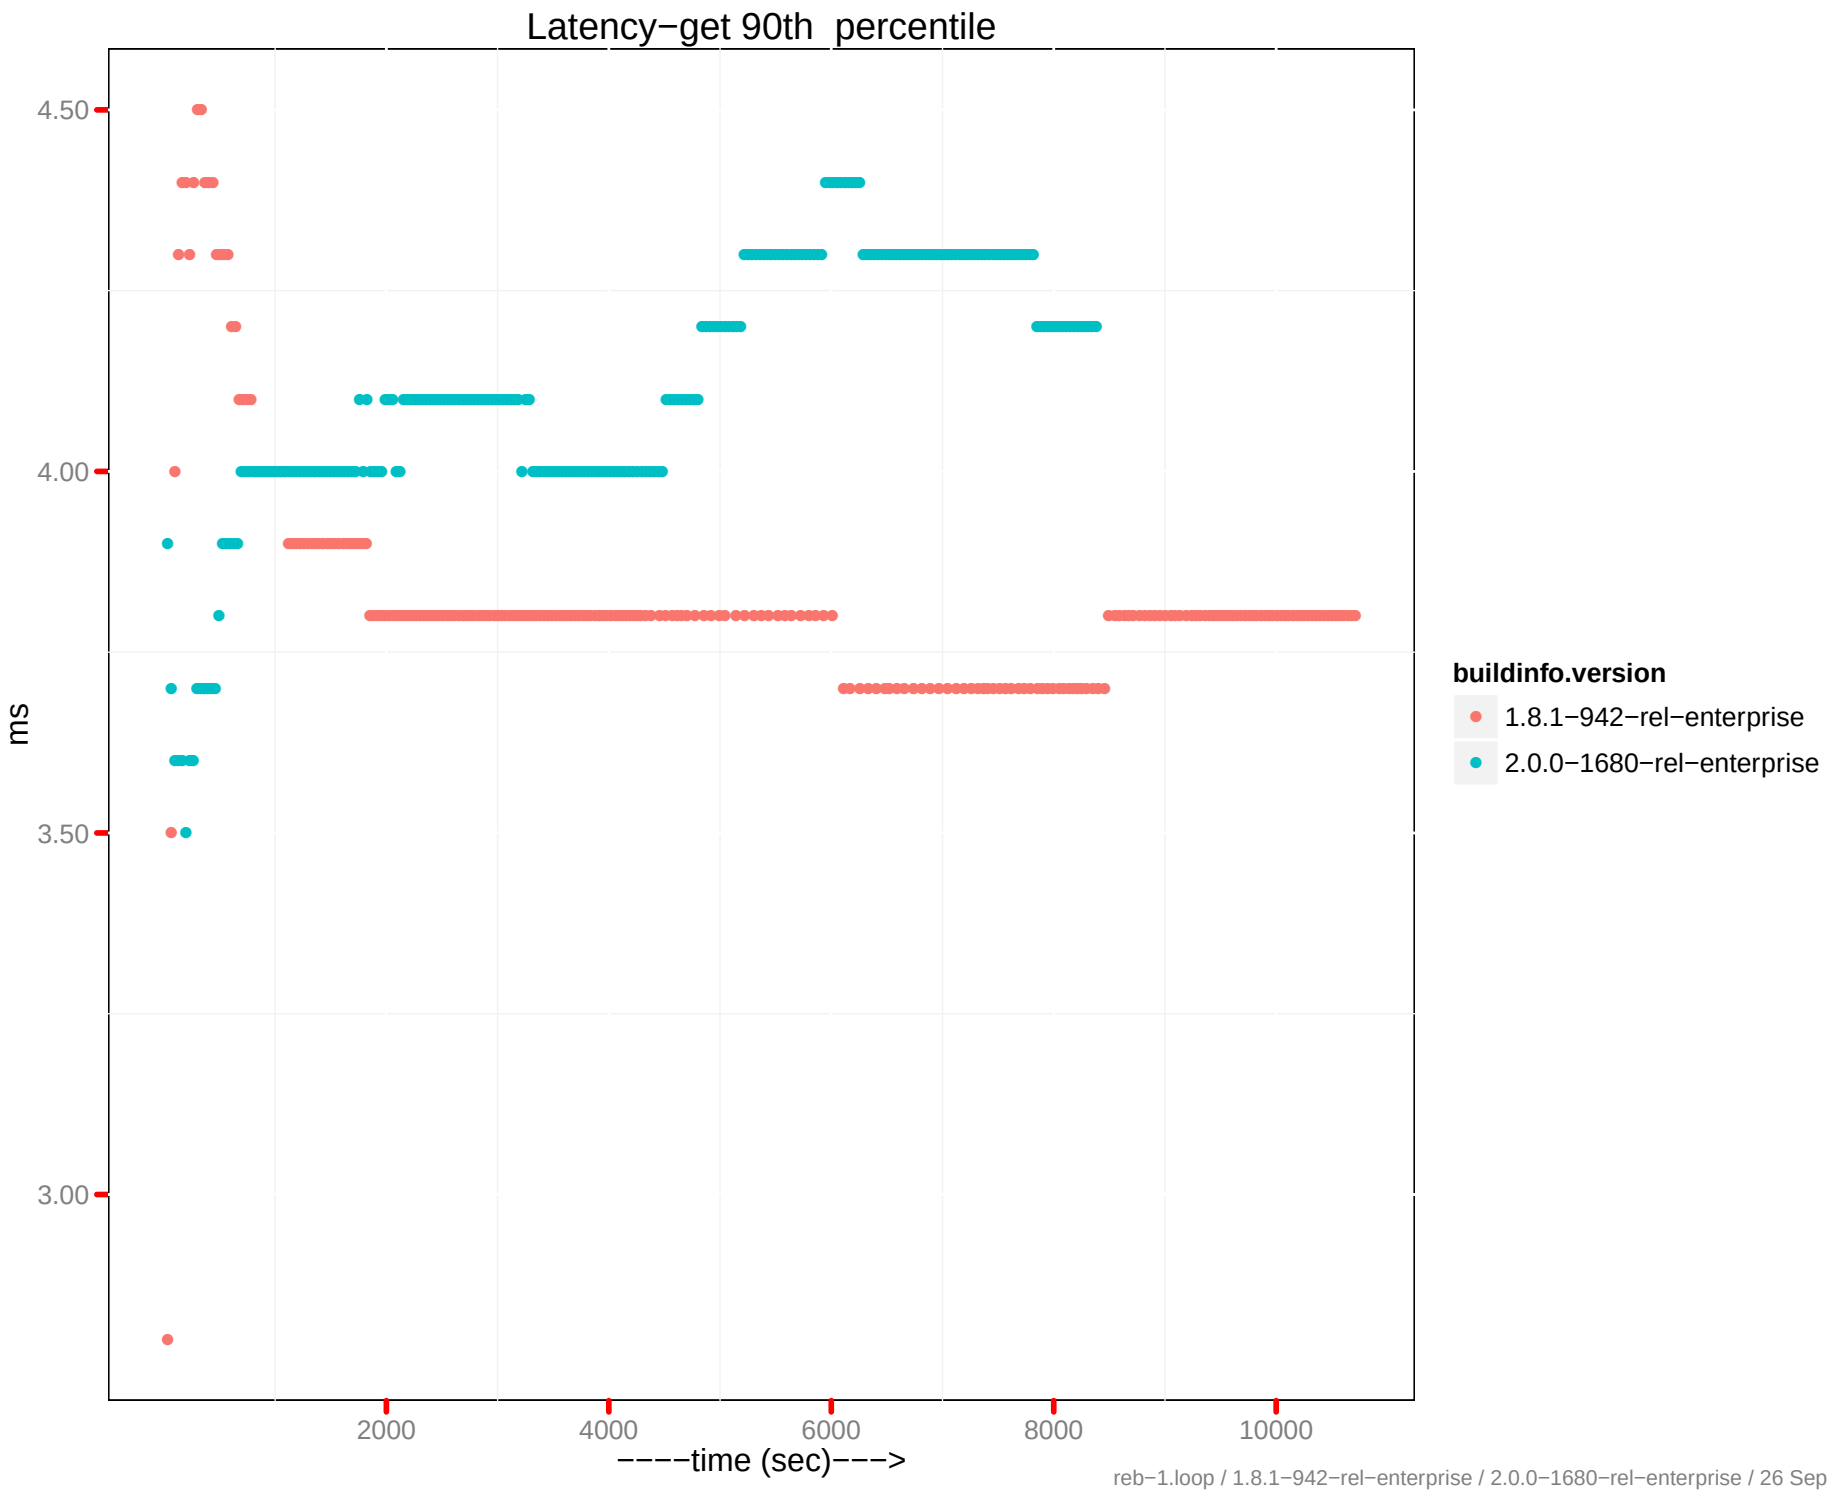
Latency−get 90th_percentile
4.50
4.00
3.50
3.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

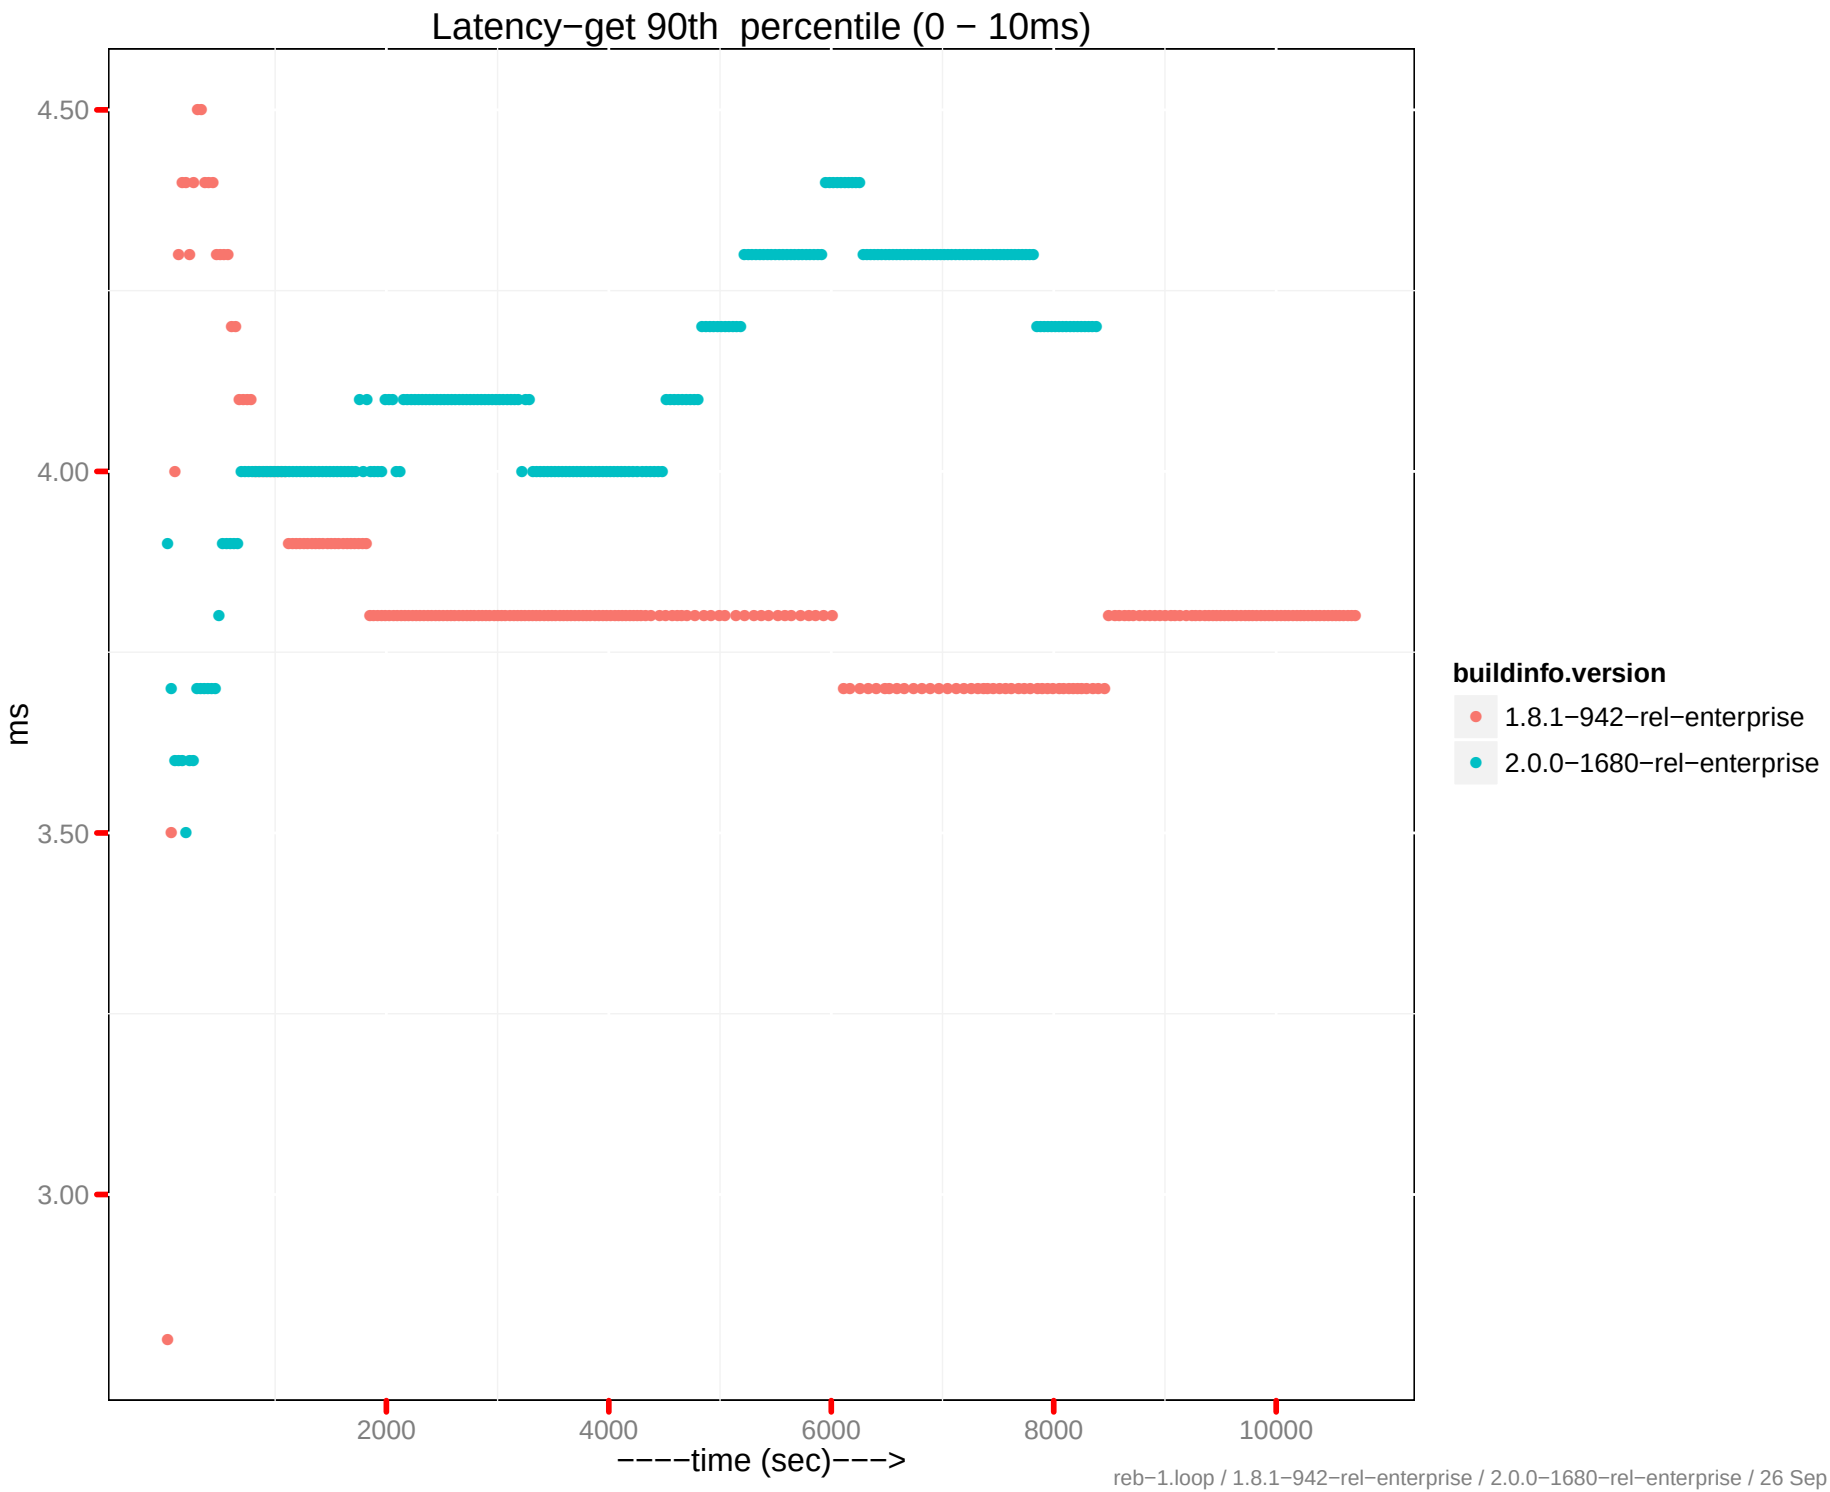
Latency−get 90th percentile (0 − 10ms)
4.50
4.00
3.50
3.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

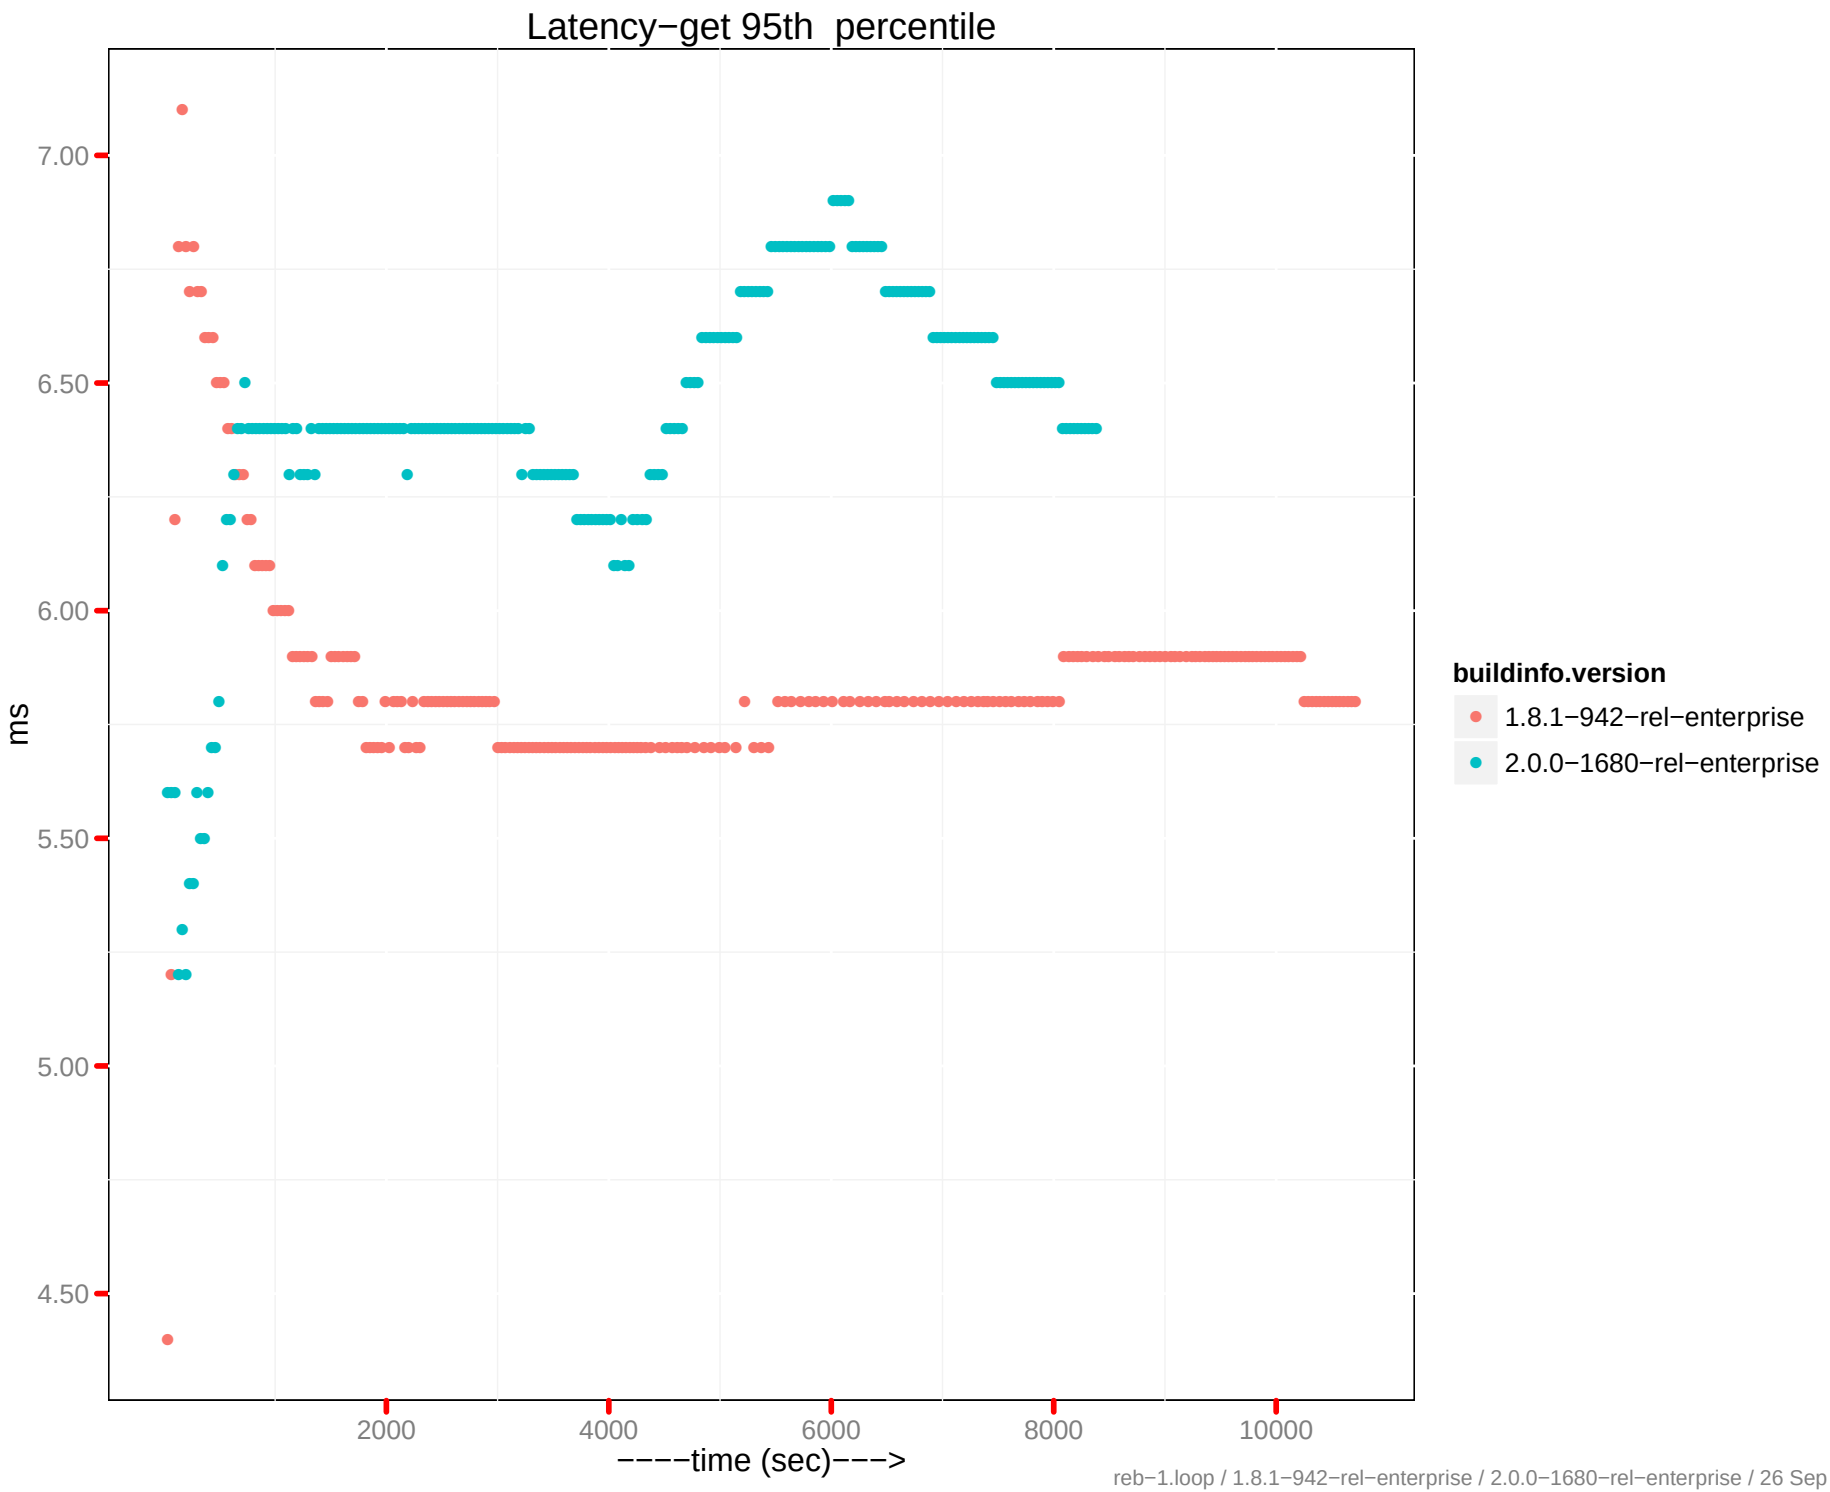
Latency−get 95th percentile
7.00
6.50
6.00
5.50
5.00
4.50
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

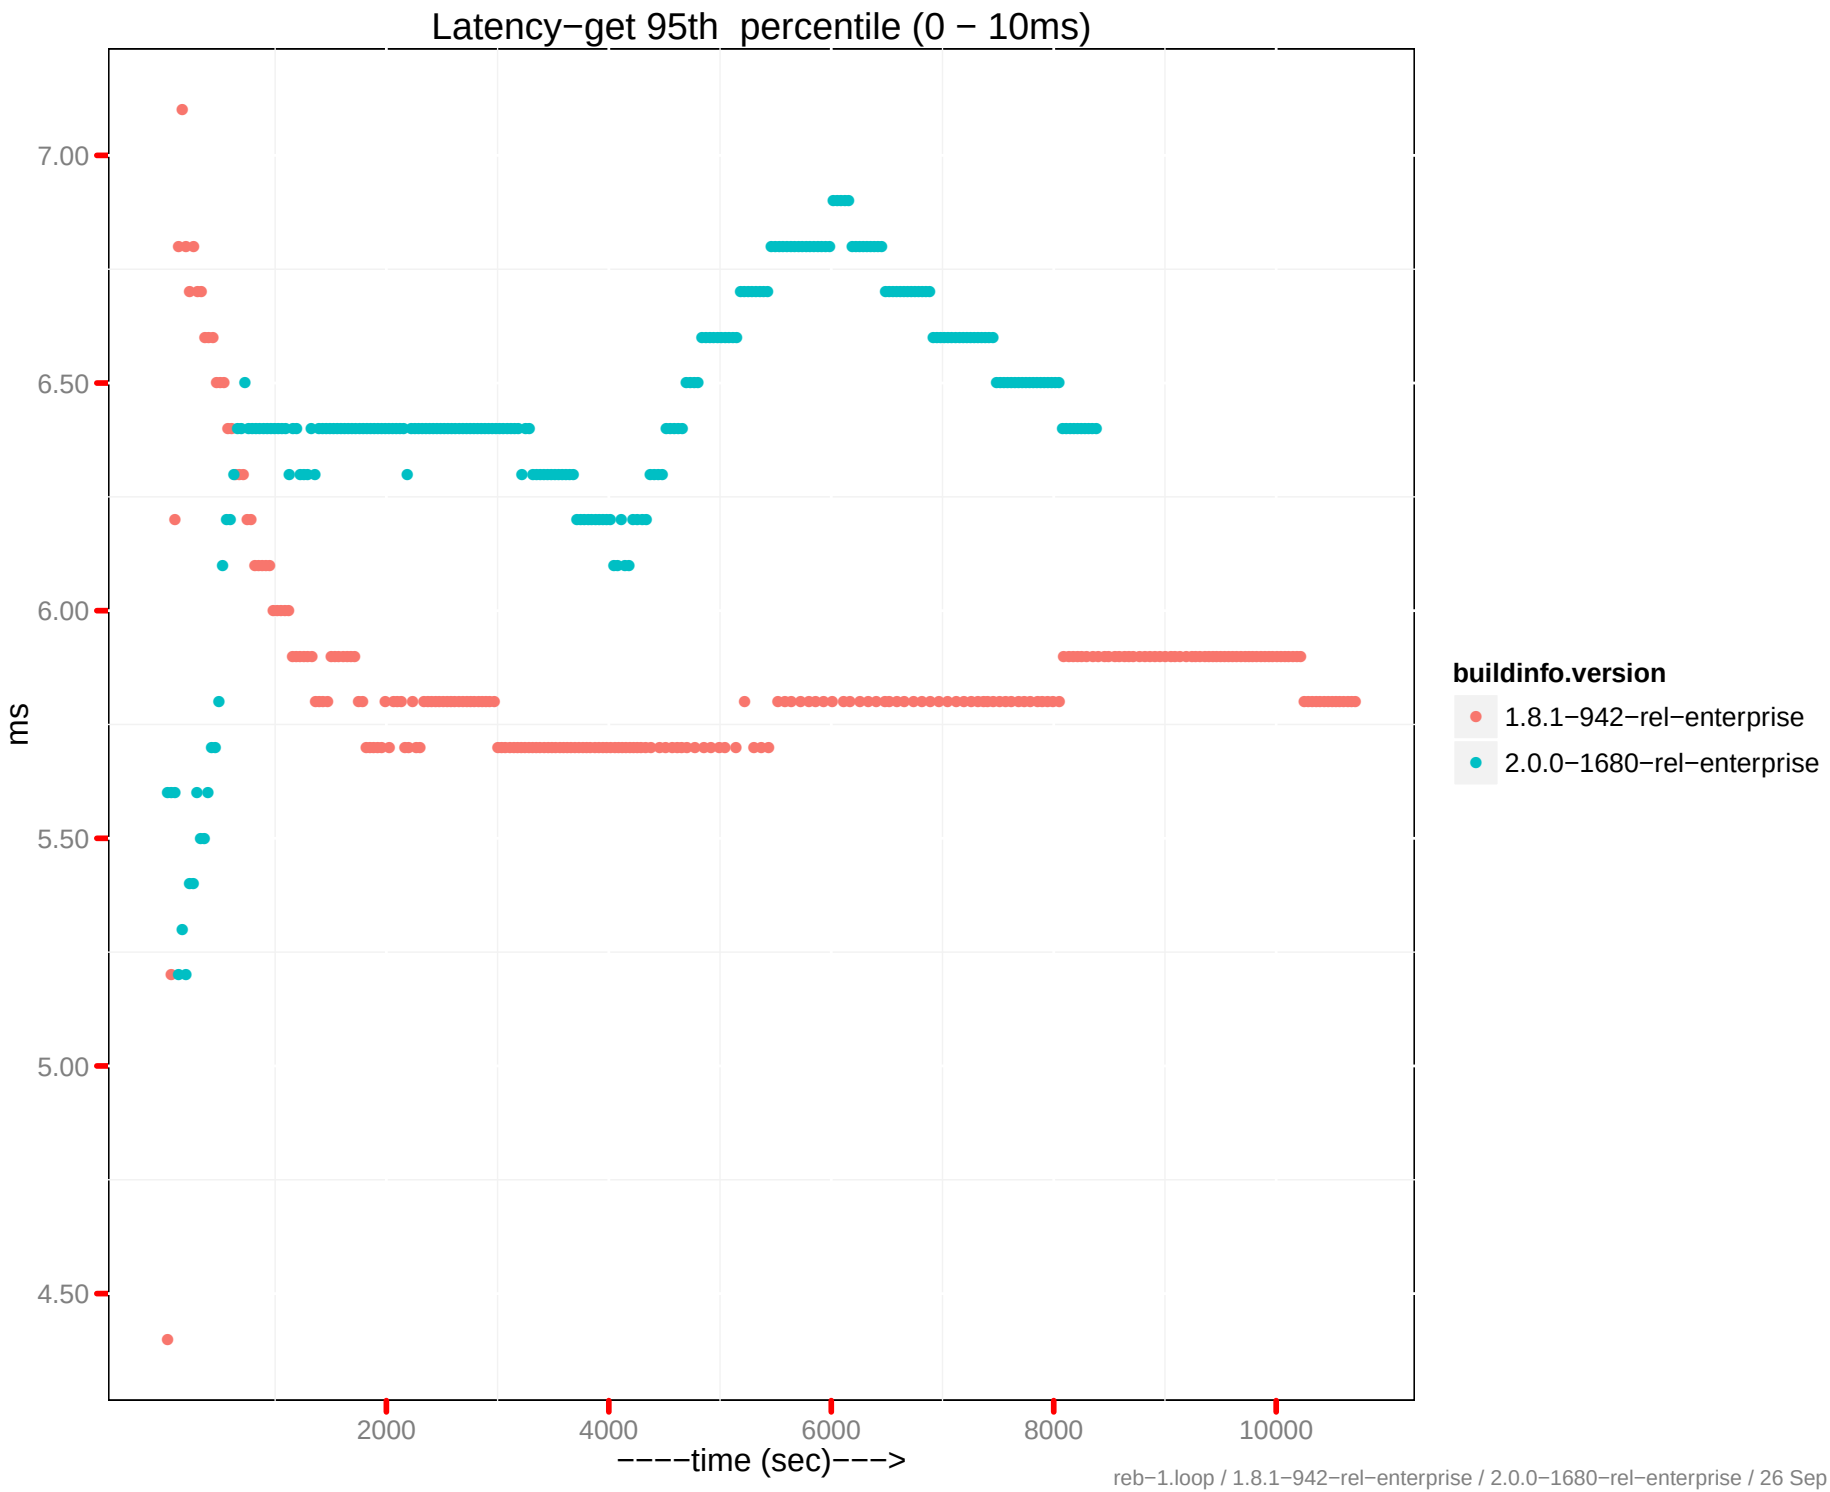
Latency−get 95th percentile (0 − 10ms)
7.00
6.50
6.00
5.50
5.00
4.50
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

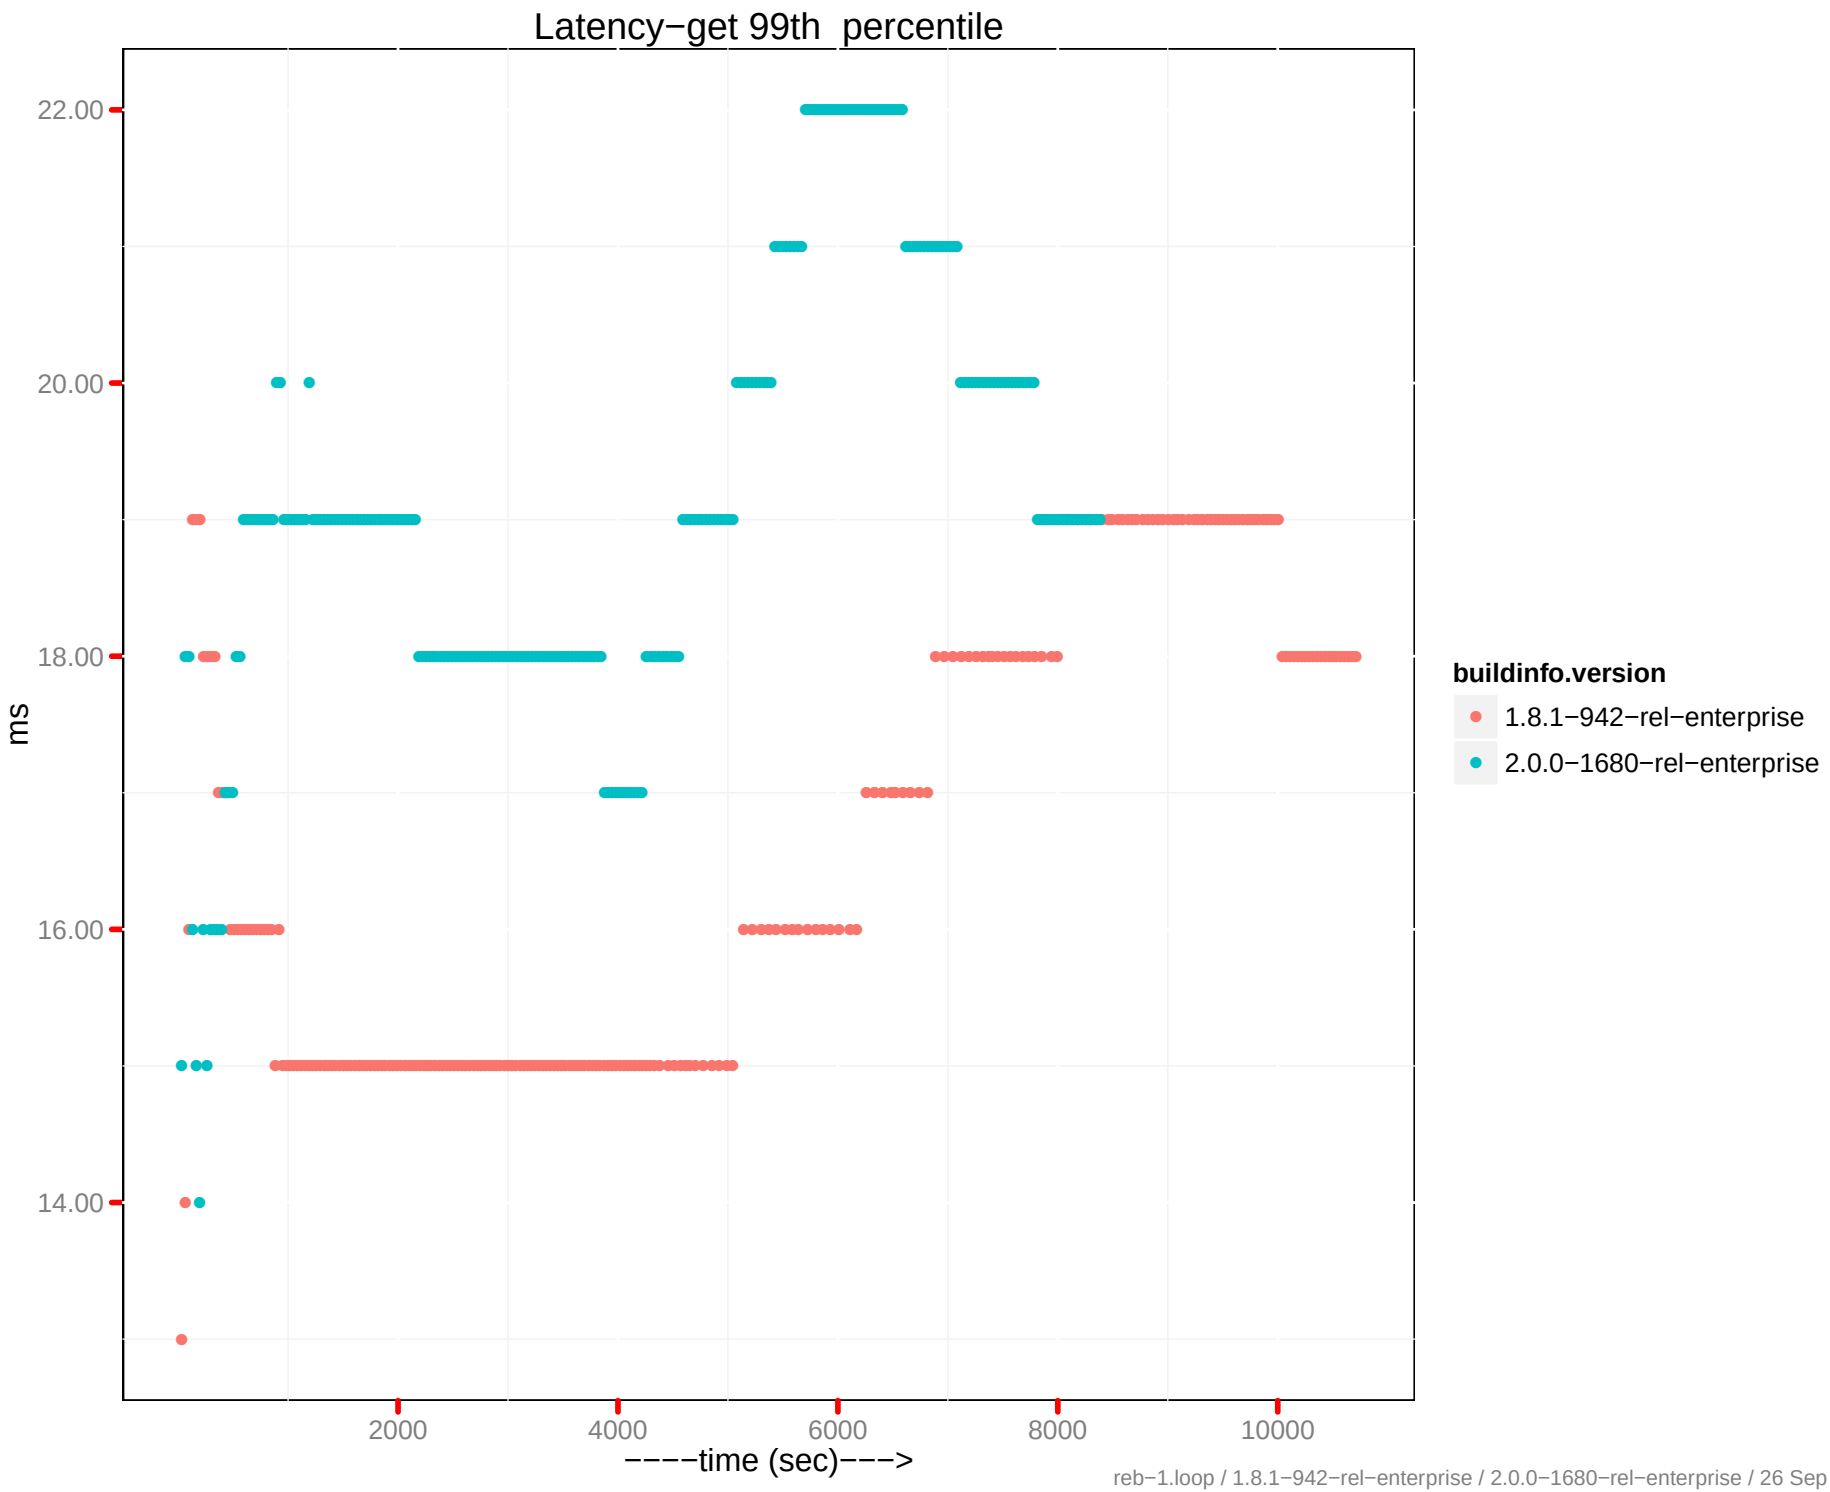
Latency−get 99th percentile
22.00
20.00
18.00
16.00
14.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

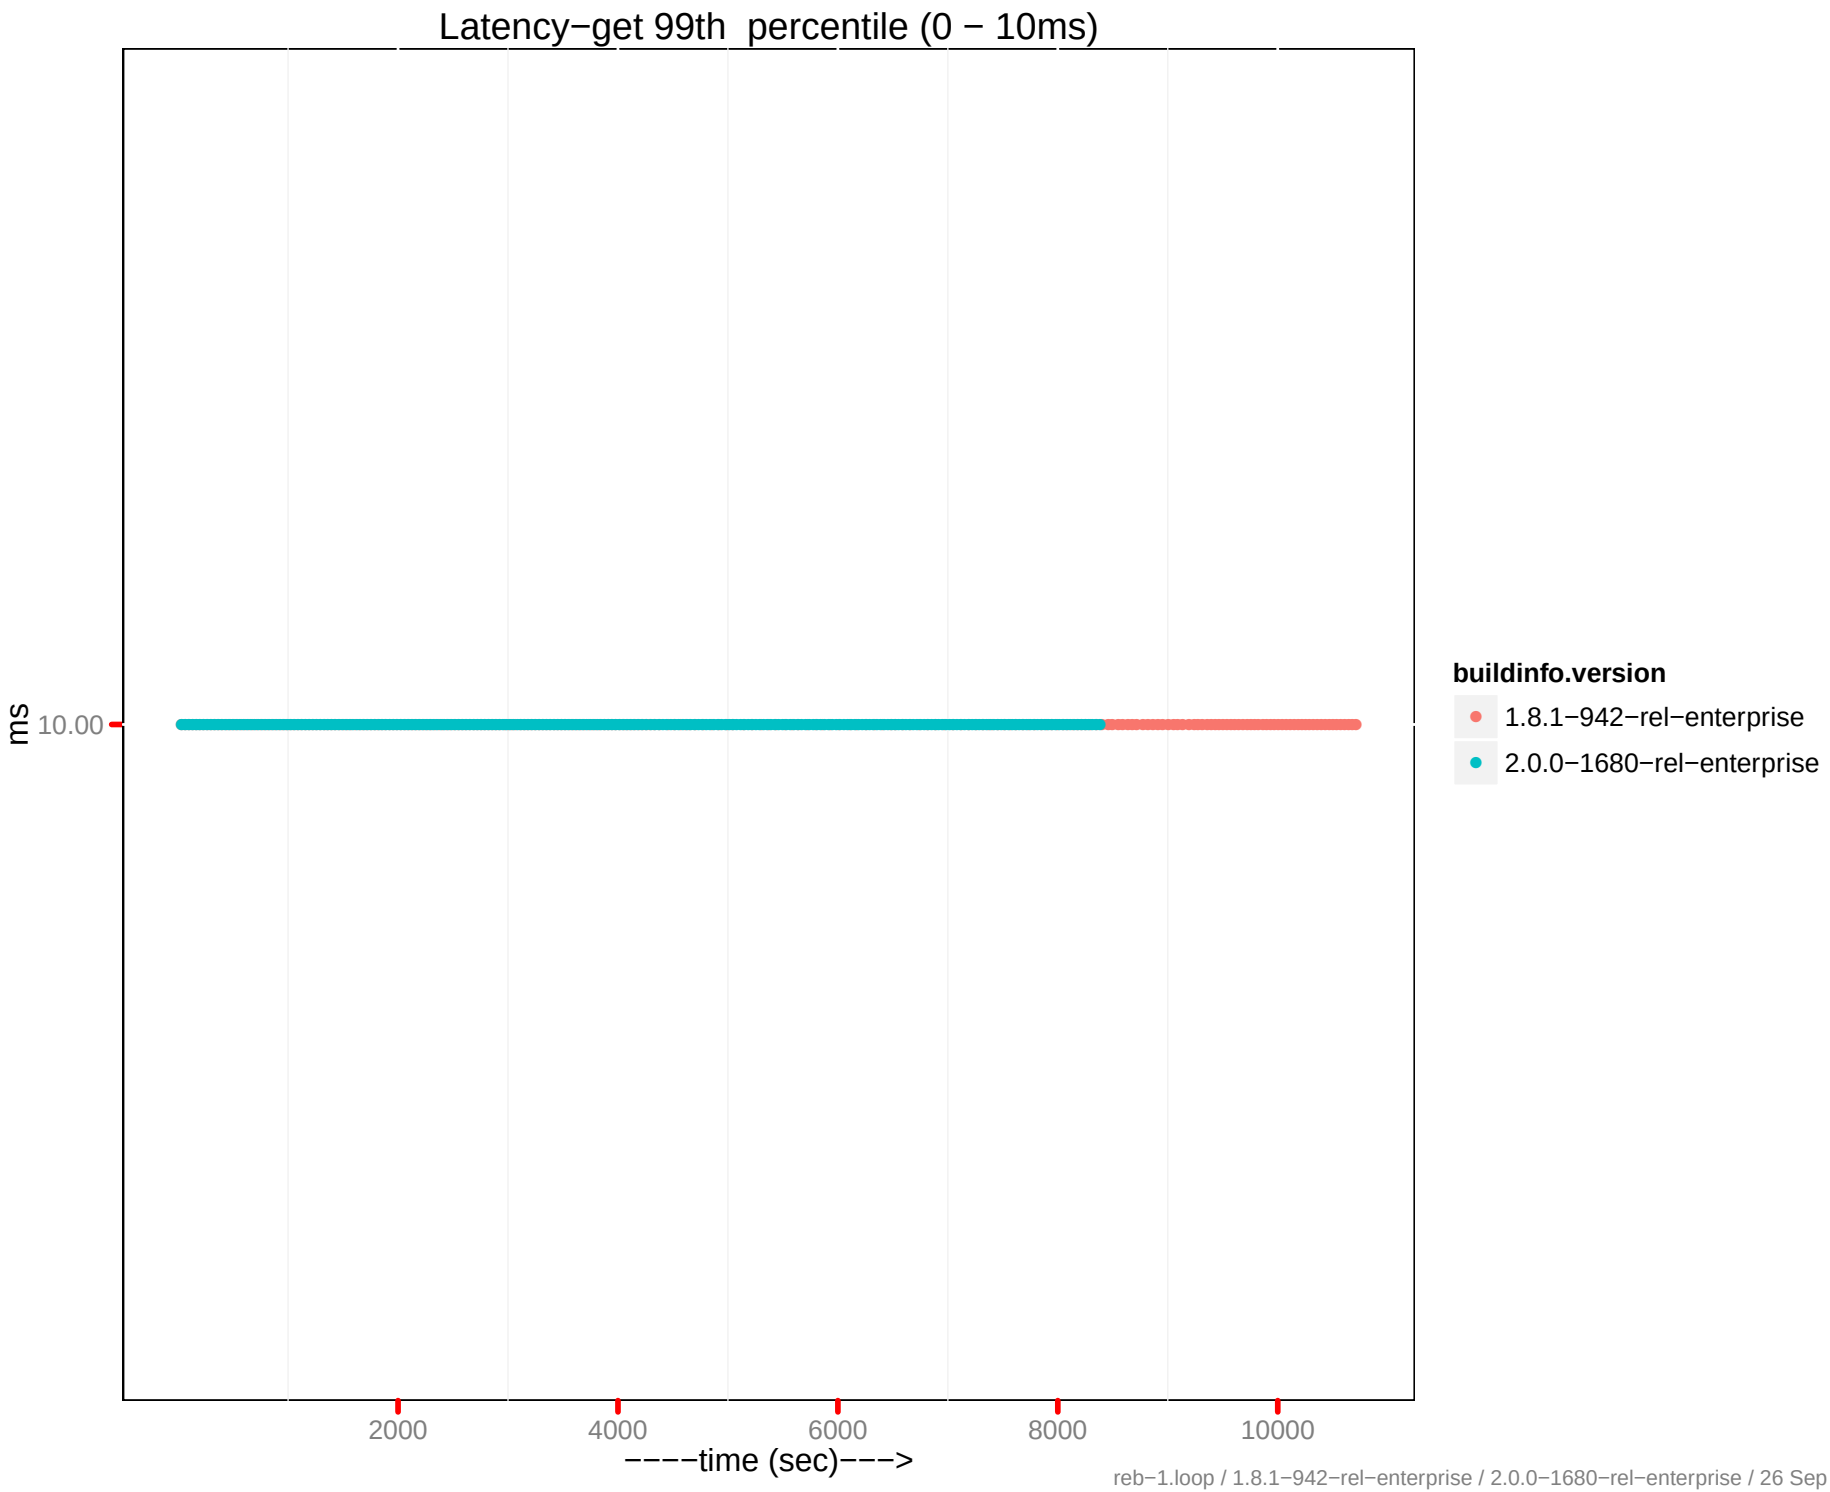
Latency−get 99th percentile (0 − 10ms)
ms
10.00
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

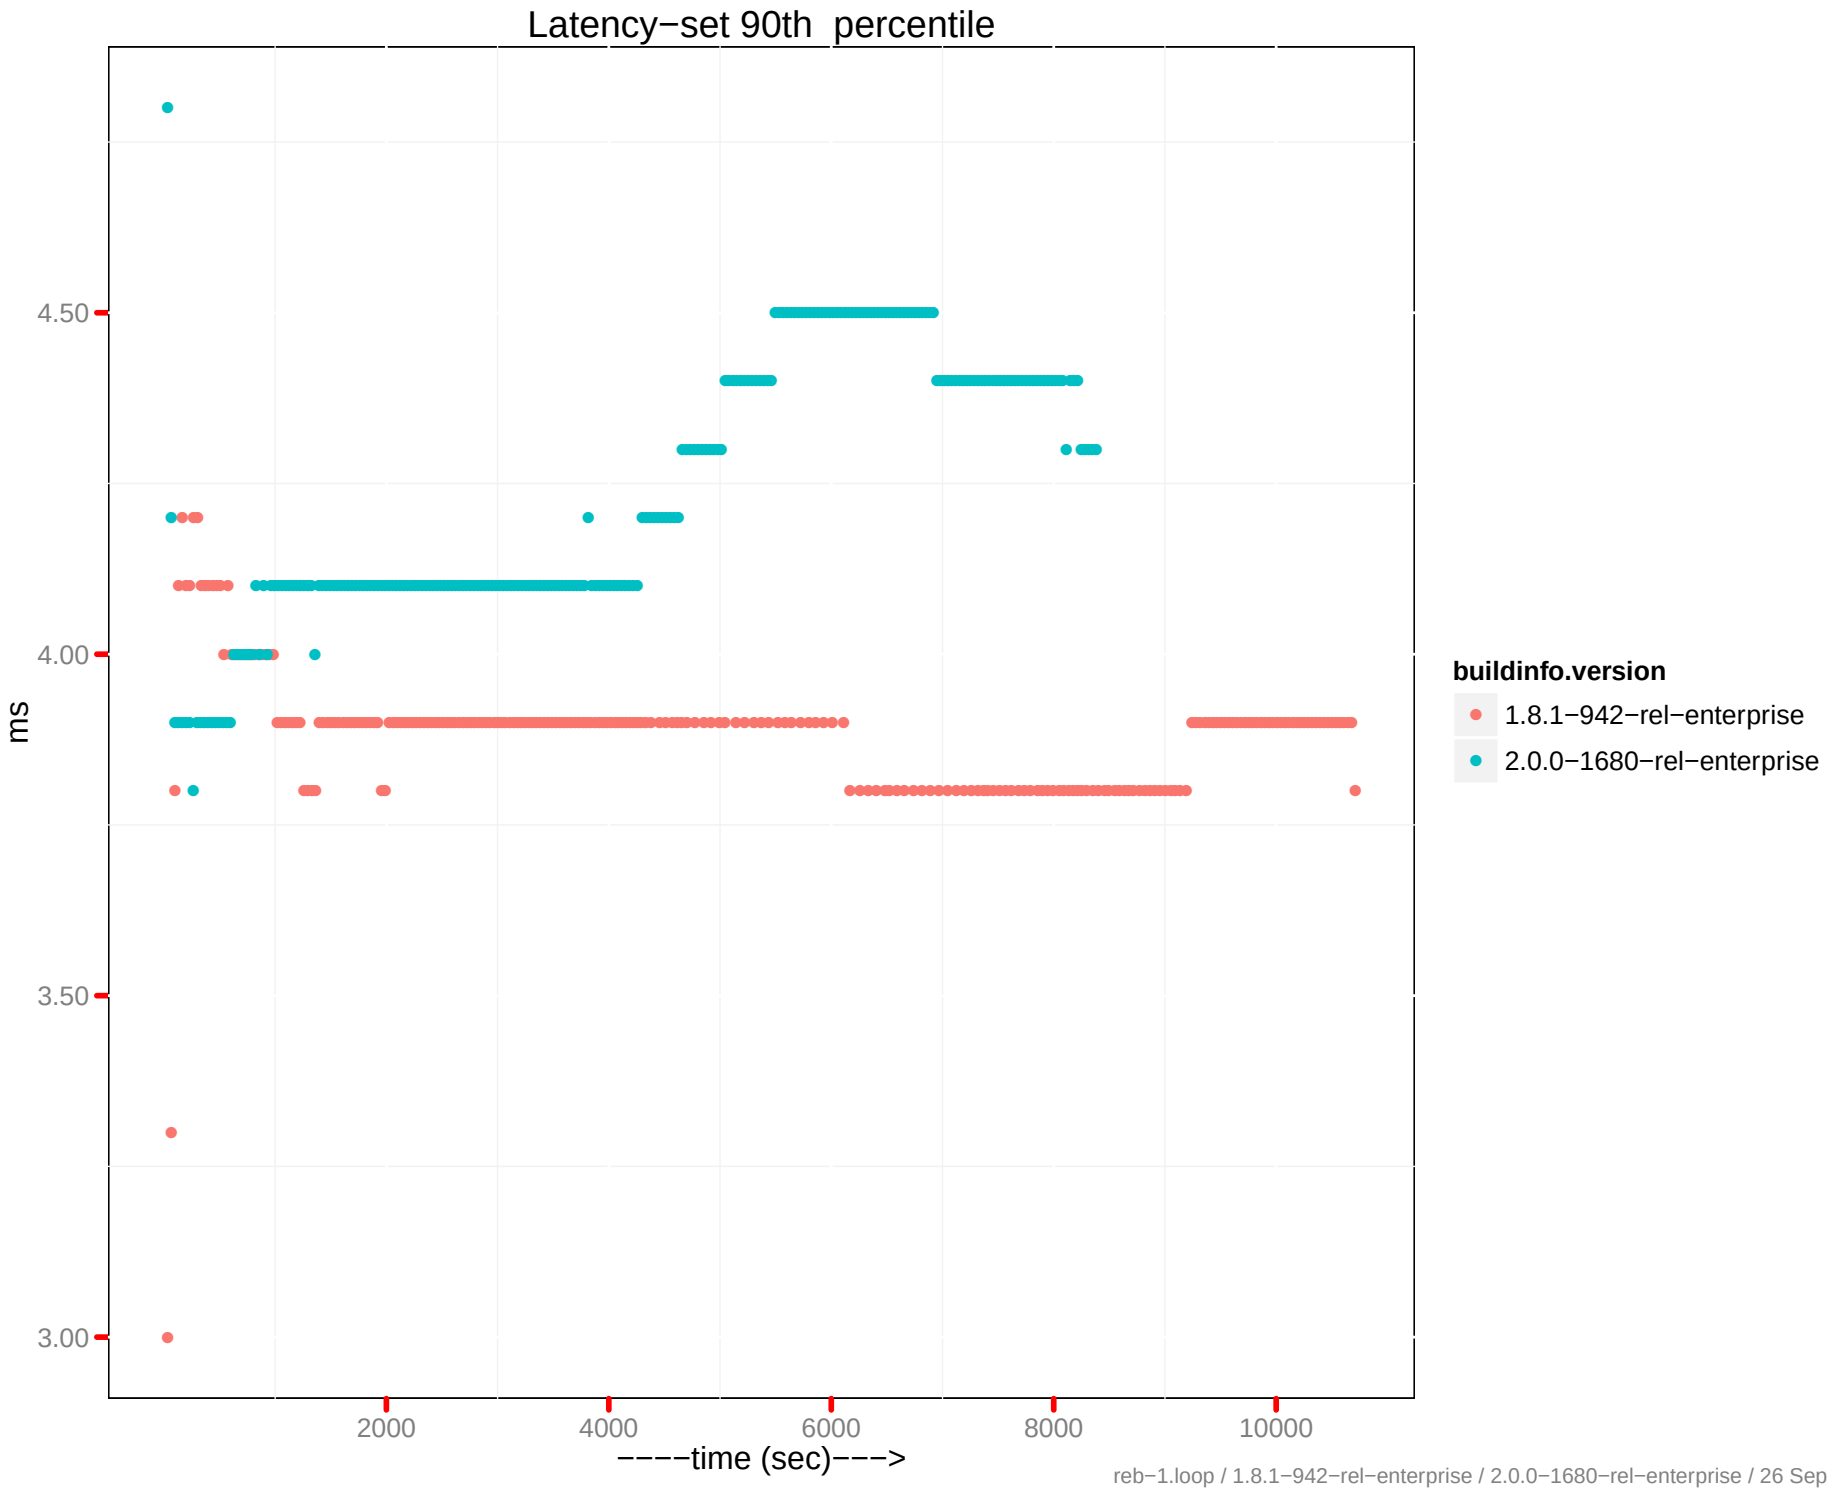
Latency−set 90th_percentile
4.50
4.00
3.50
3.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

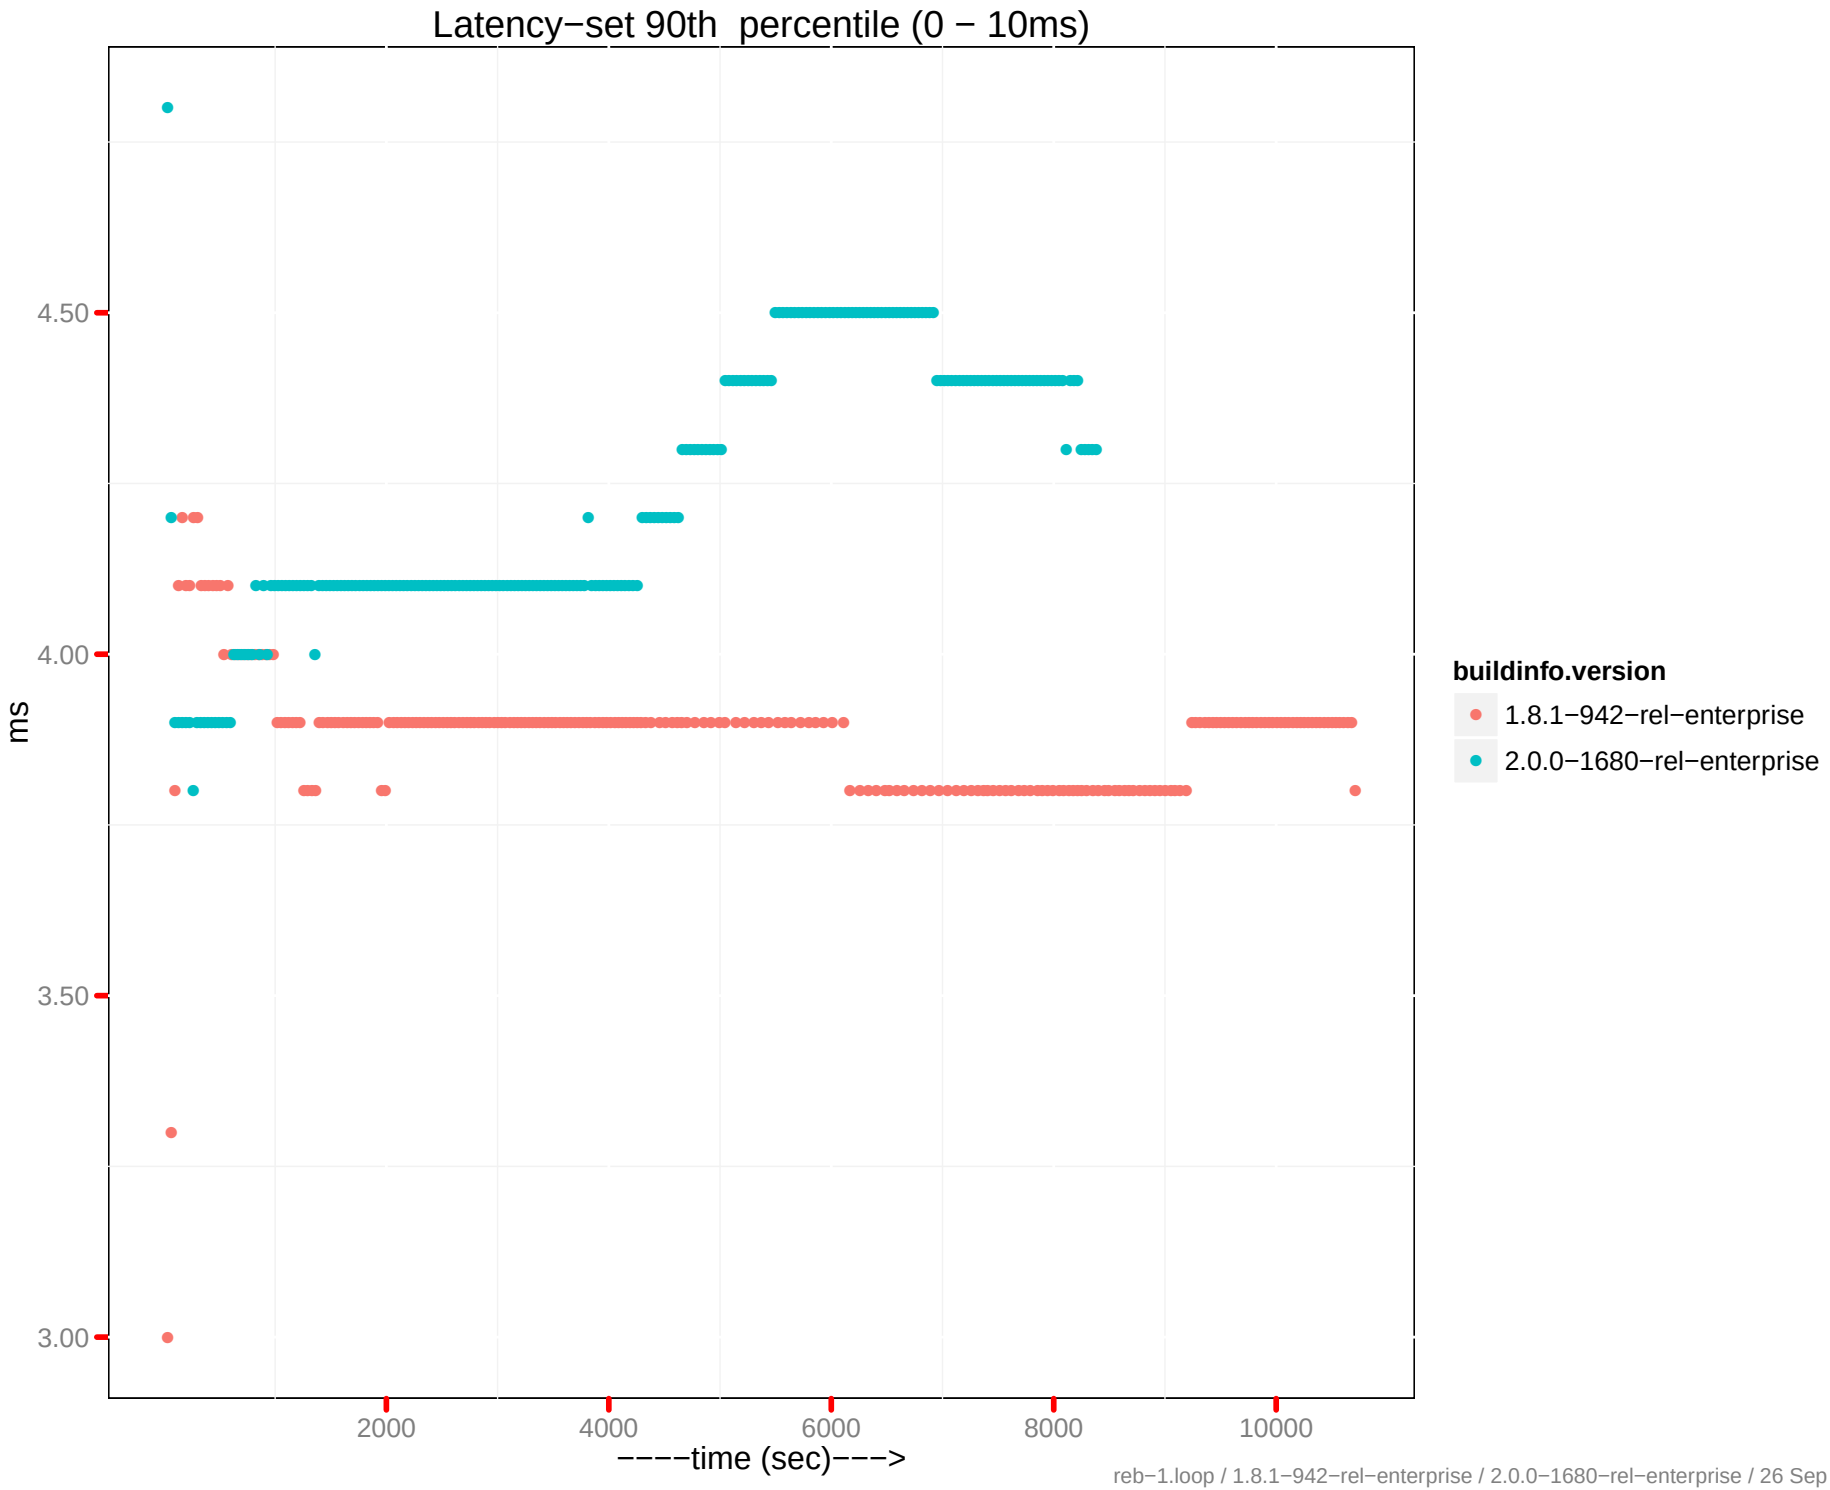
Latency−set 90th percentile (0 − 10ms)
4.50
4.00
3.50
3.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

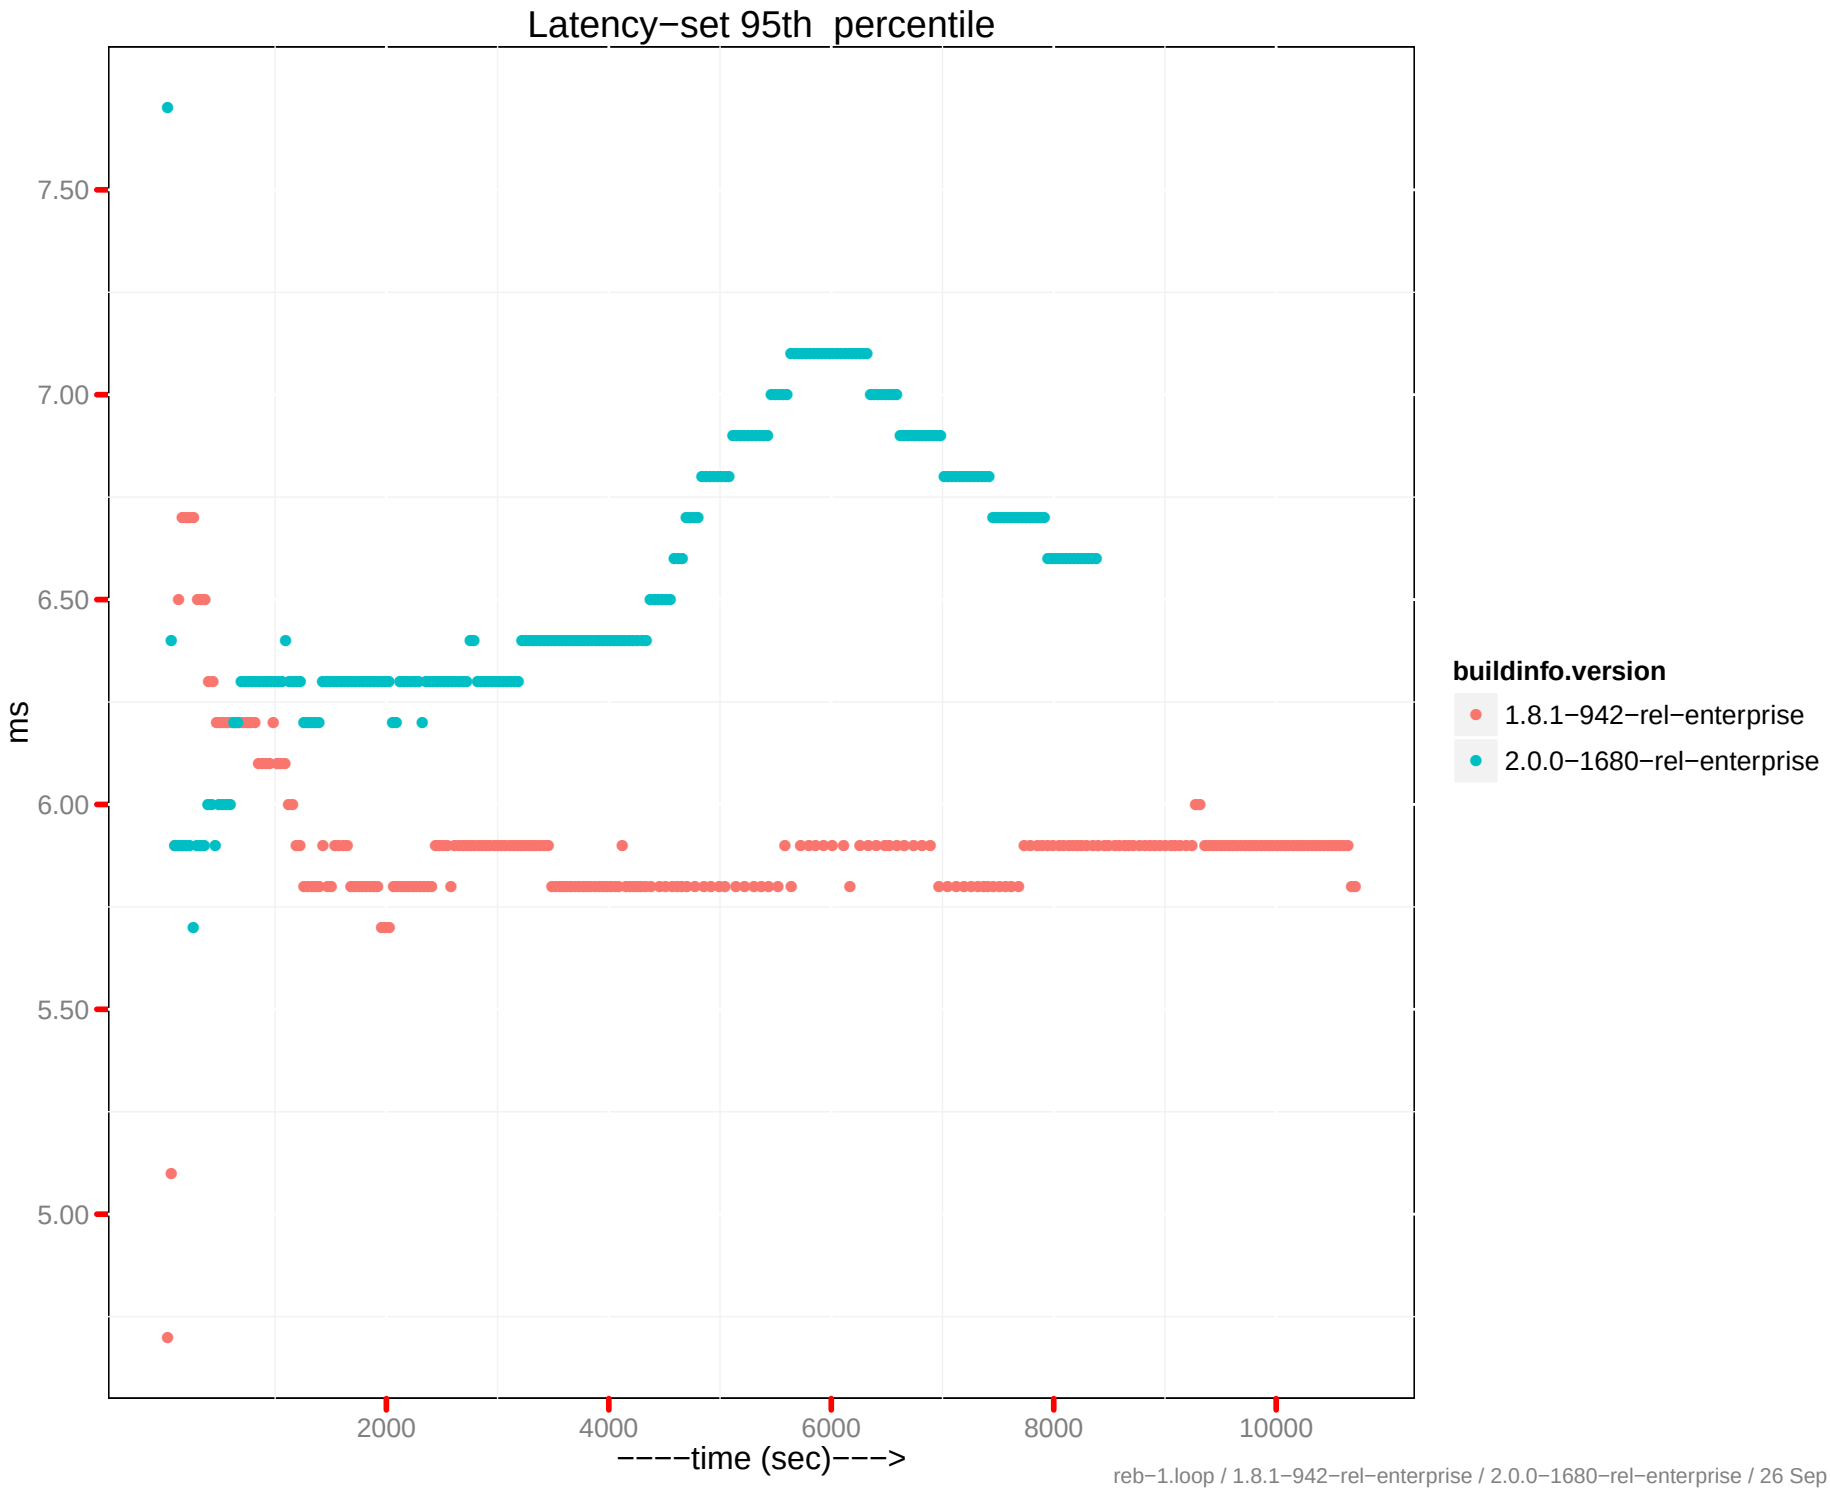
Latency−set 95th_percentile
7.50
7.00
6.50
6.00
5.50
5.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

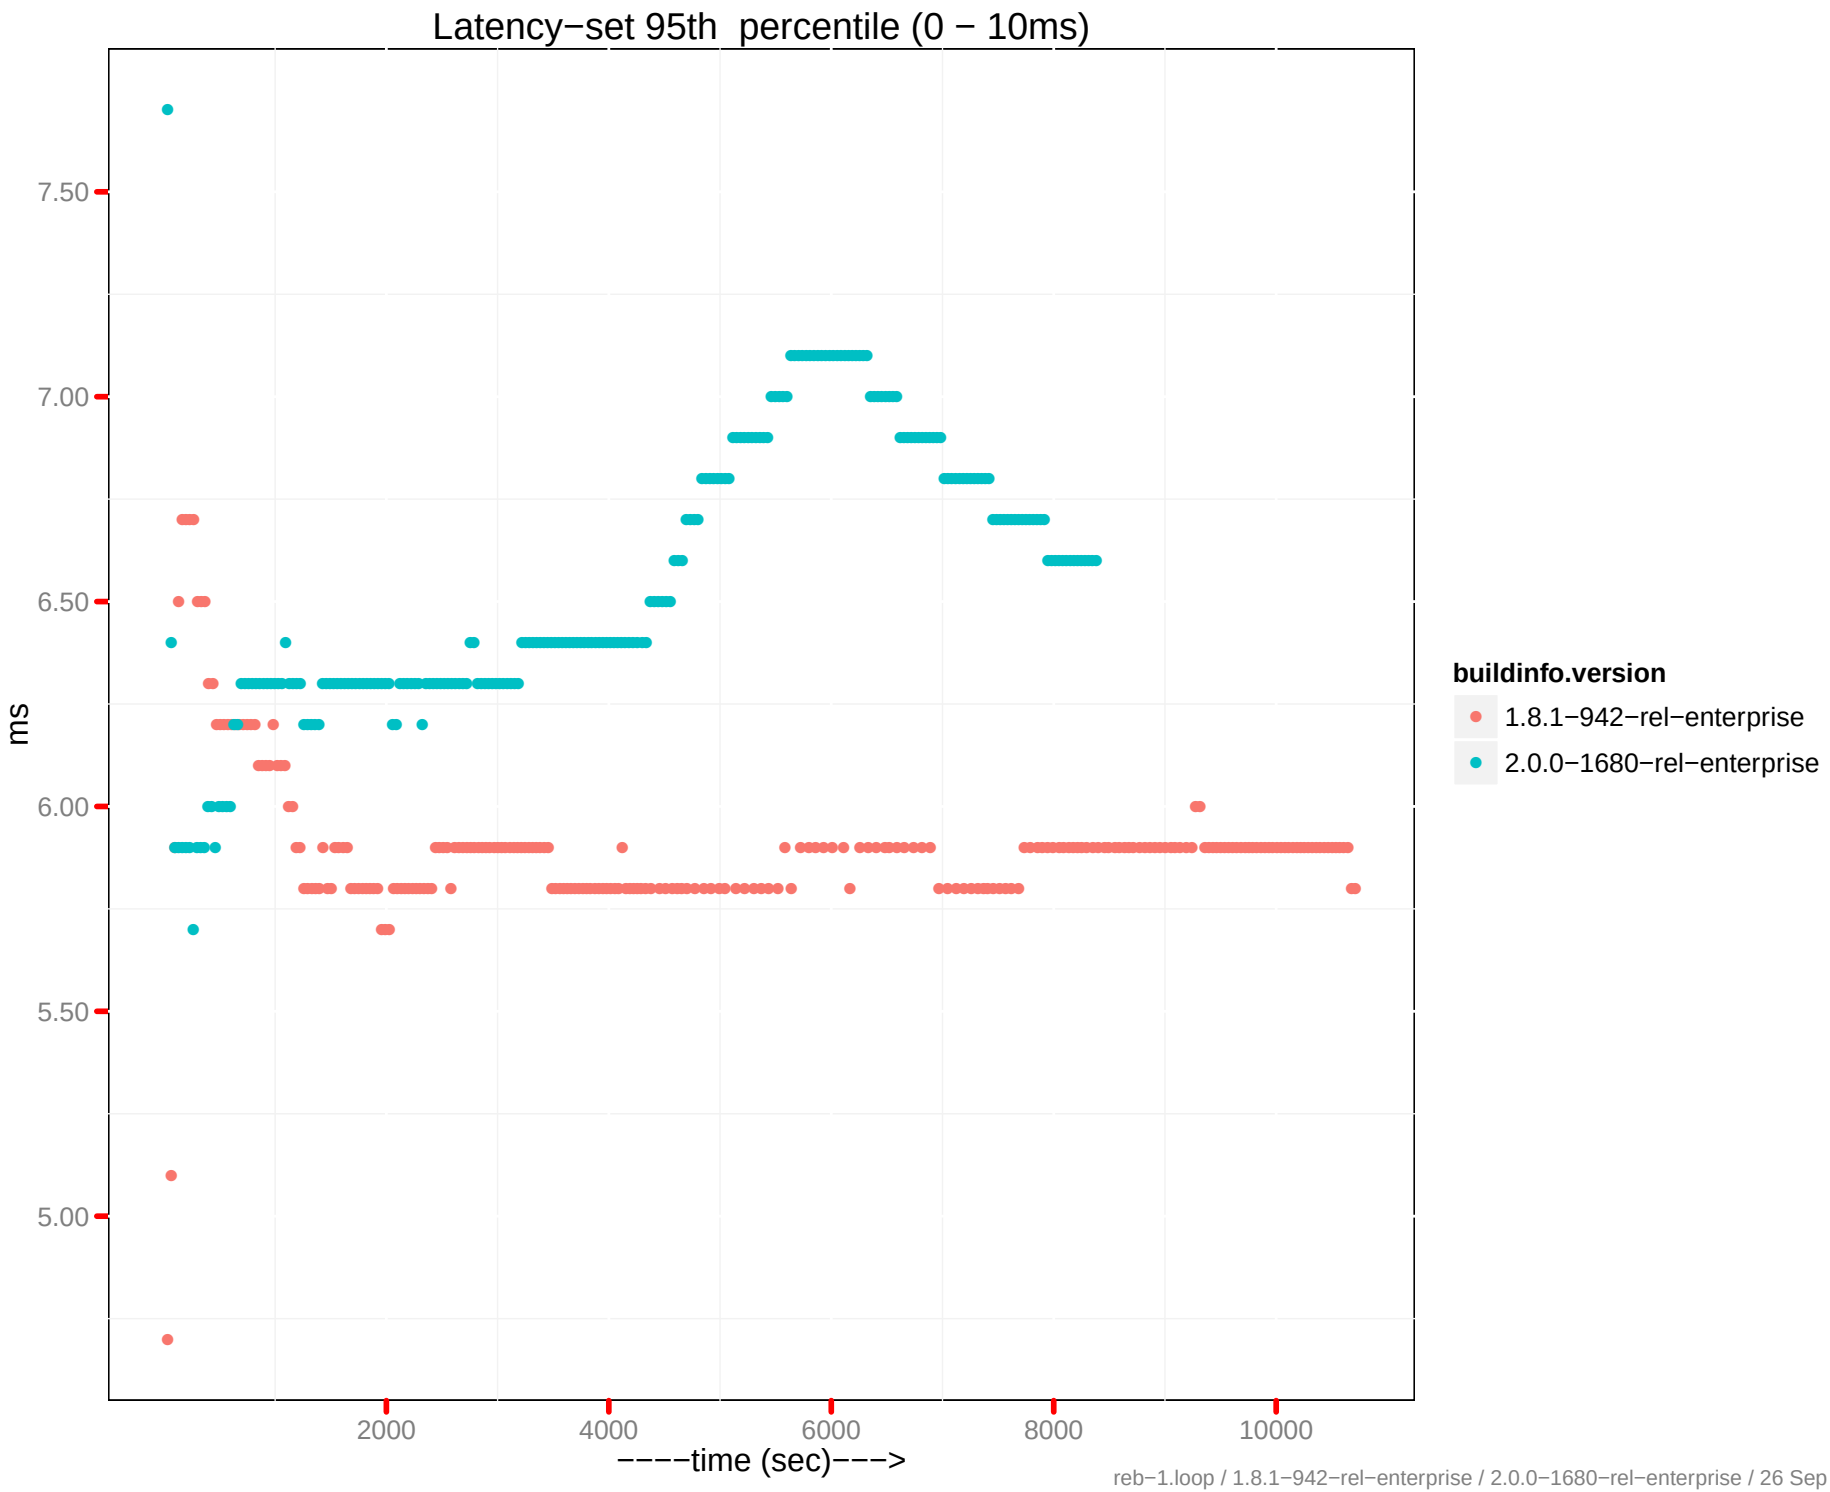
Latency−set 95th percentile (0 − 10ms)
7.50
7.00
6.50
6.00
5.50
5.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

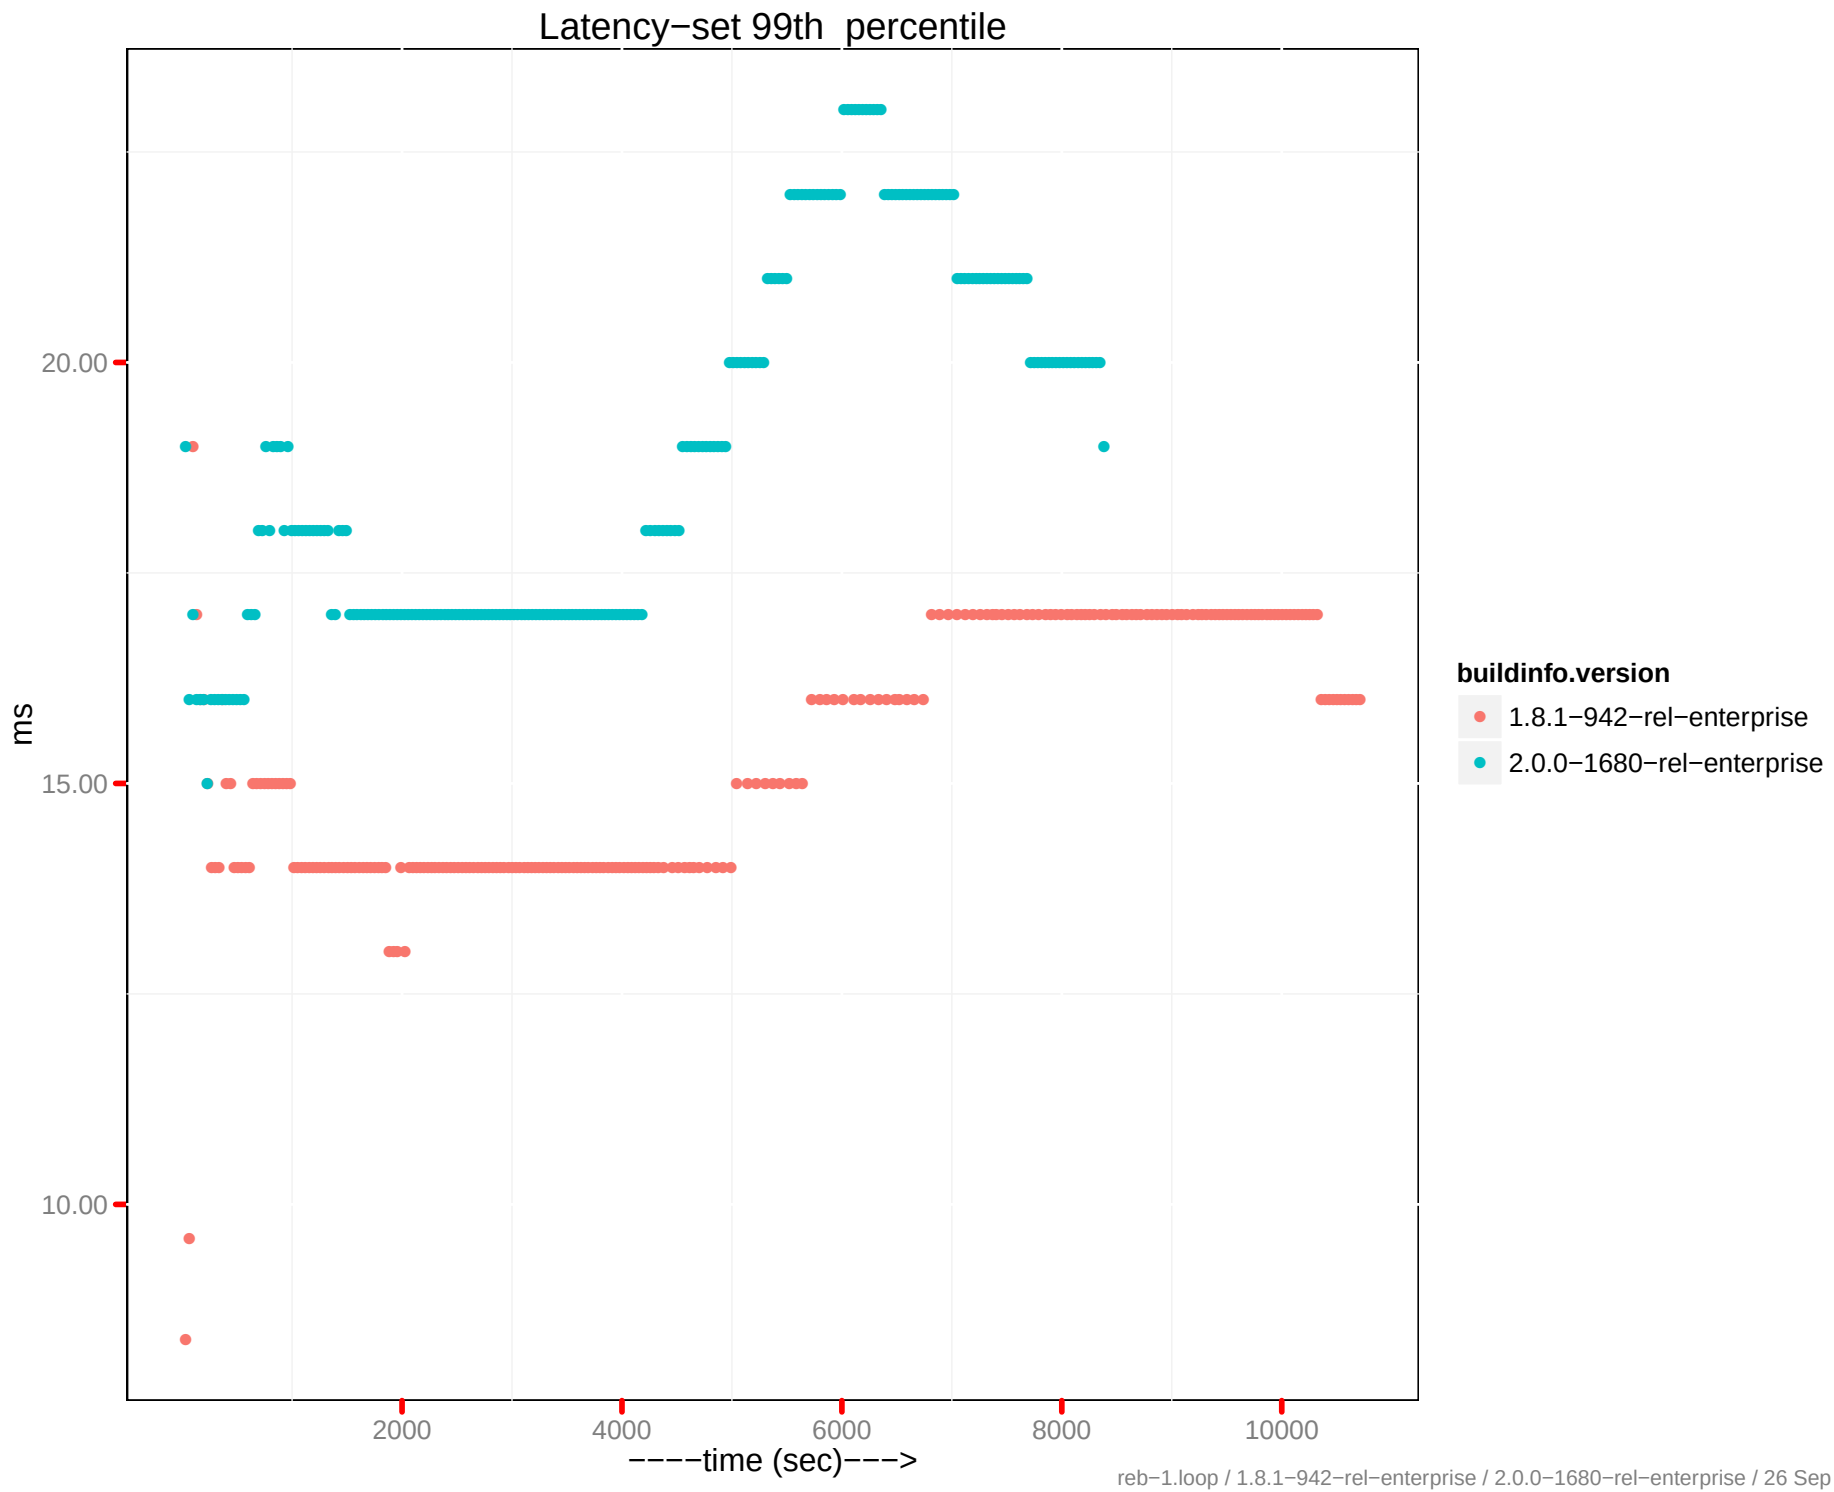
Latency−set 99th percentile
20.00
15.00
10.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

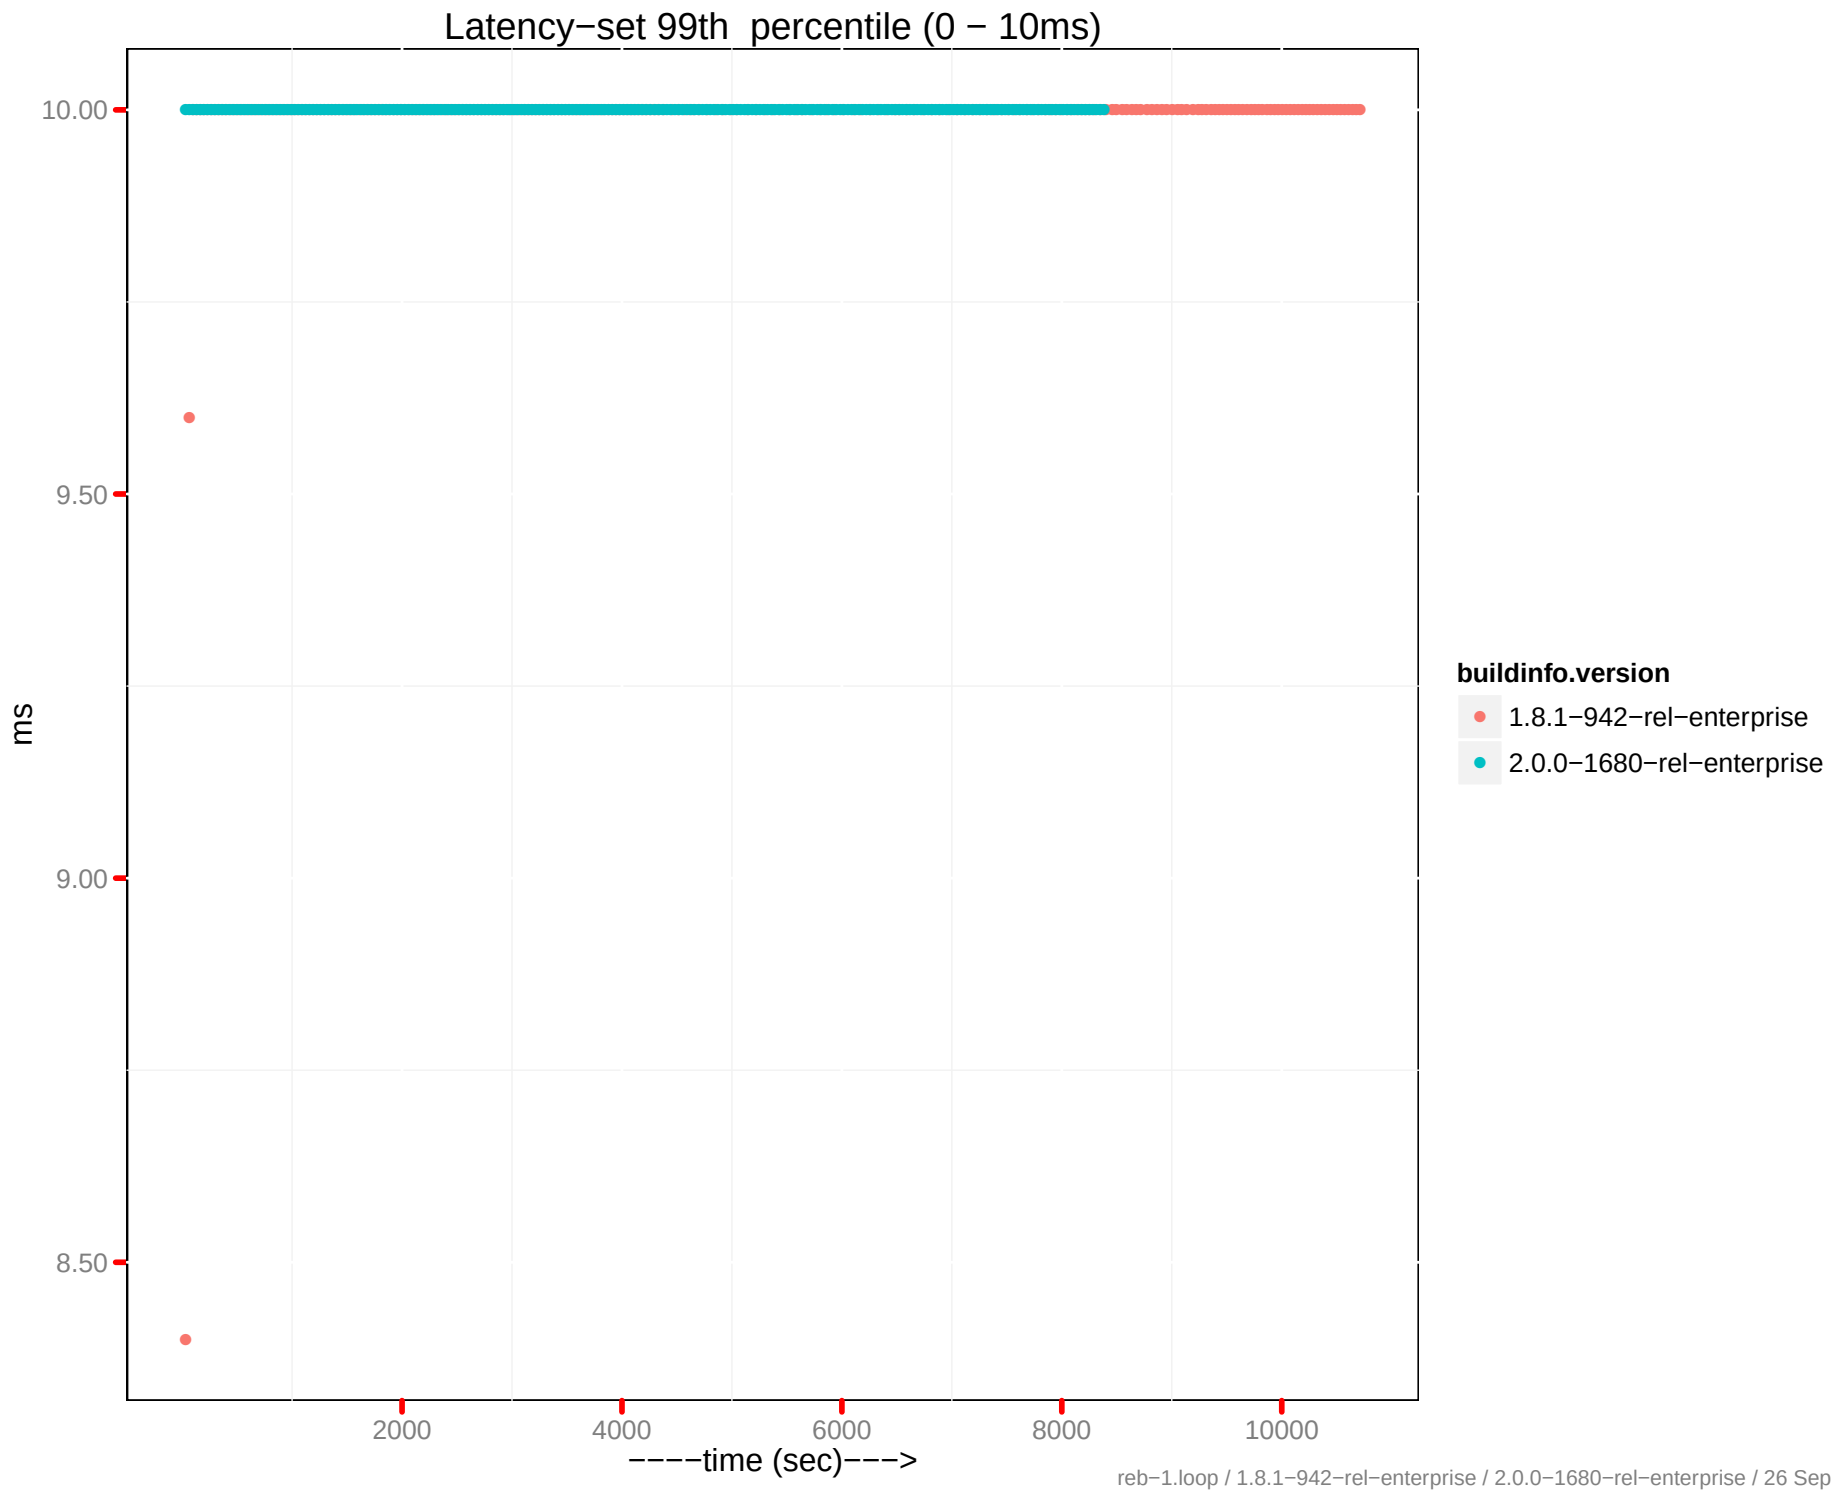
Latency−set 99th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
1.8.1−942−rel−enterprise
2.0.0−1680−rel−enterprise
reb−1.loop / 1.8.1−942−rel−enterprise / 2.0.0−1680−rel−enterprise / 26 Sep 2012

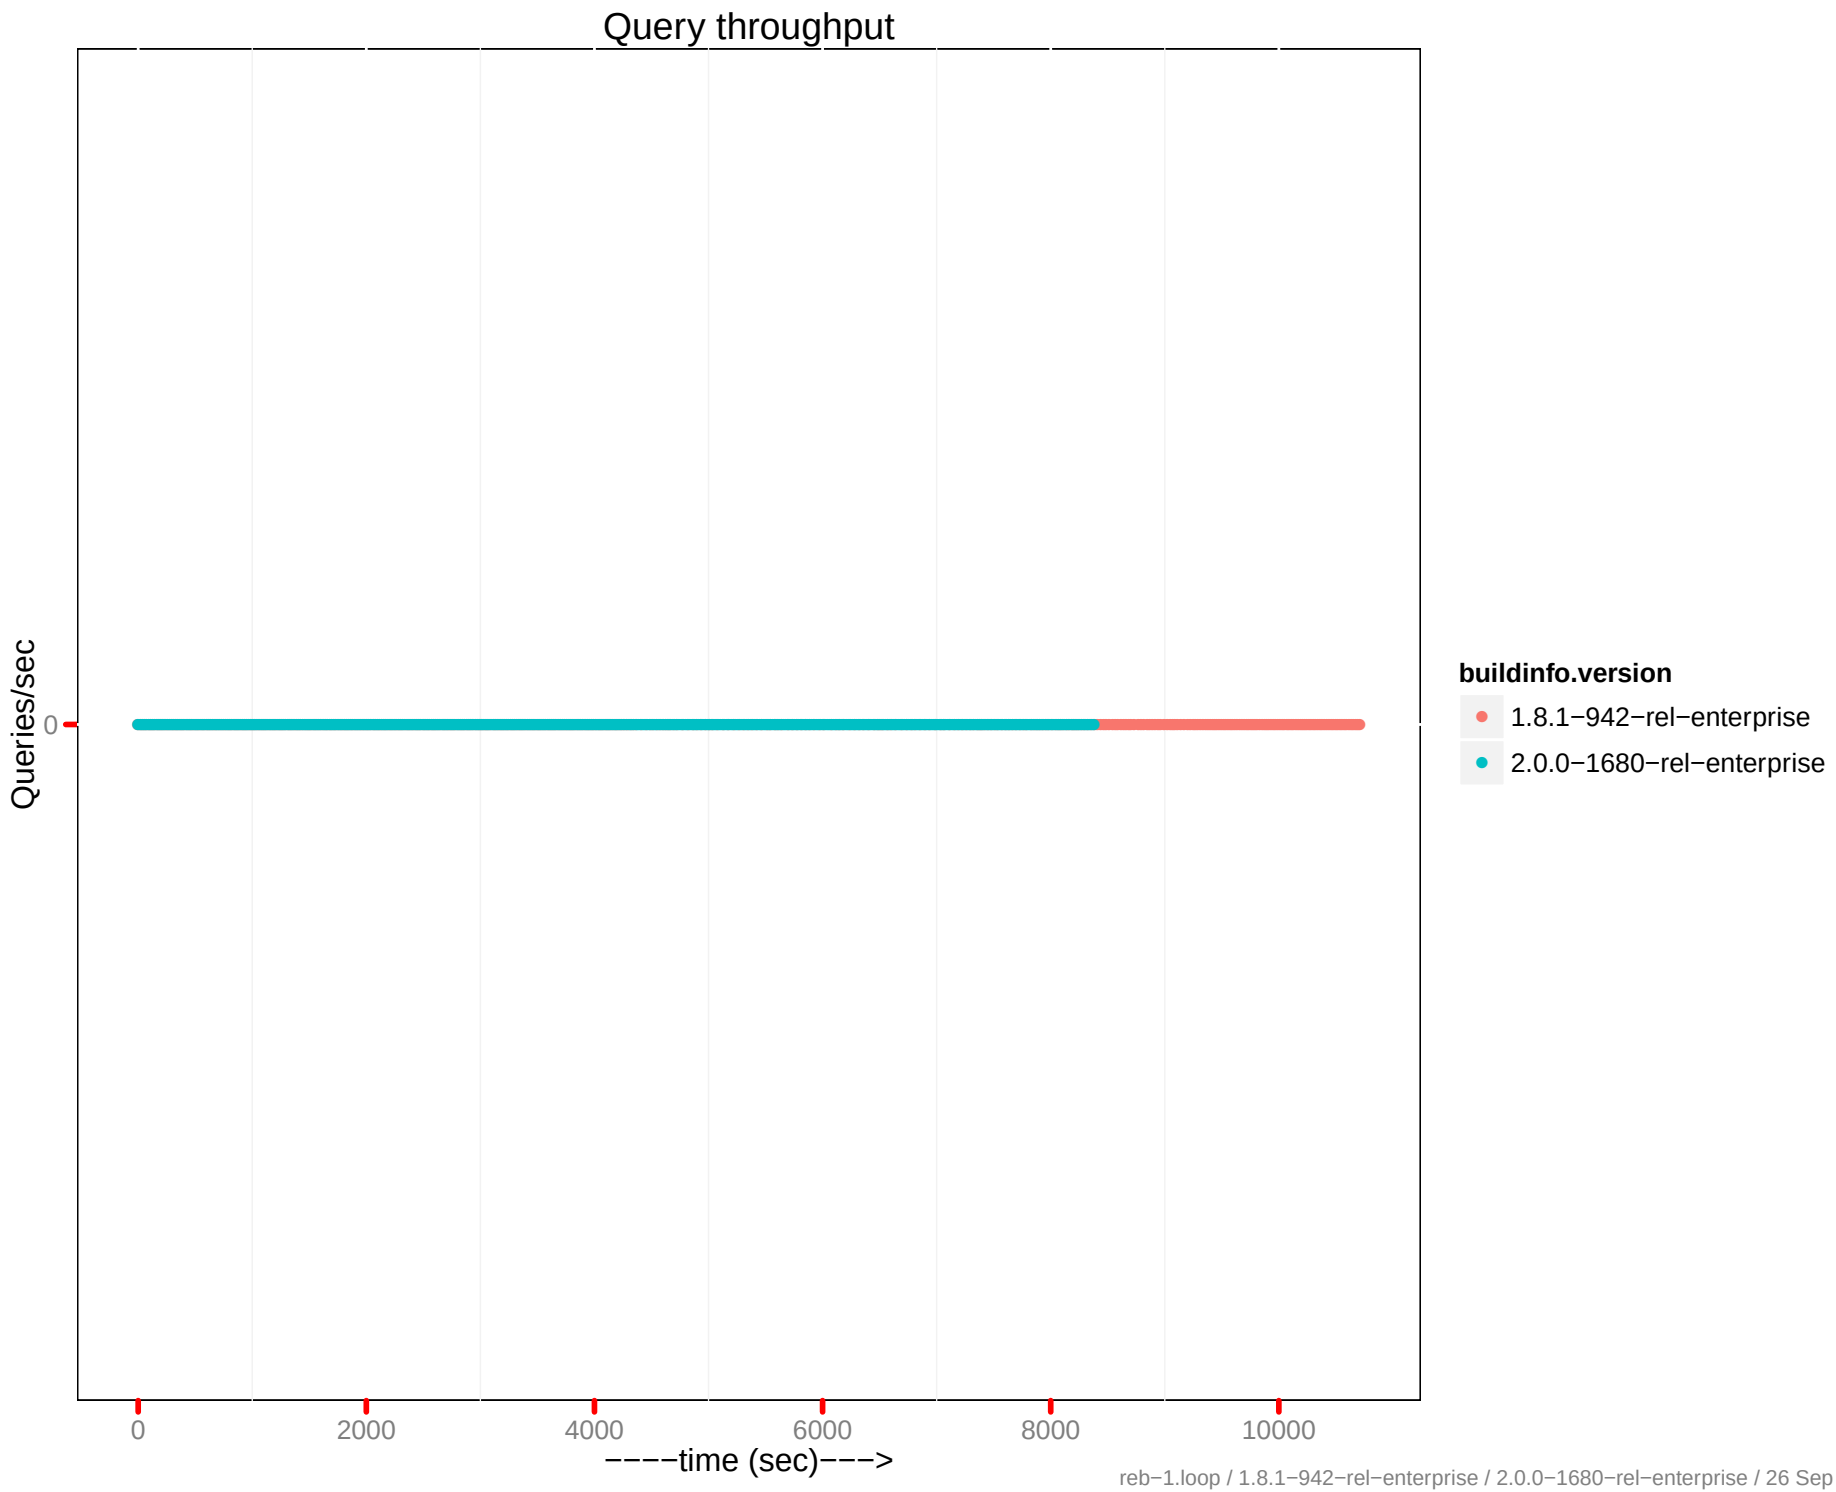
Query throughput
Queries/sec
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
1.8.1-942-rel-enterprise
2.0.0-1680-rel-enterprise
reb-1.loop / 1.8.1-942-rel-enterprise / 2.0.0-1680-rel-enterprise / 26 Sep 2012

```
reb-1.conf
# rebalance mixed 7M load, 1M hot reload, 3M access creates
# speed limit = 3k
#
performance.eperf.EPerfClient.test_eperf_rebalance

params:

# general
batch=50
kind=nonjson
mem_quota=20000
db_compaction=30

# load phase
hot_init_items=1000000
items=7000000

# access phase
# Read:Insert:Update:Delete Ratio = 50:4:40:6.
ratio_sets=0.5
ratio_misses=0.05
ratio_creates=0.08
ratio_deletes=0.13
ratio_hot=0.05
ratio_hot_gets=0.99
ratio_hot_sets=0.99
ratio_expirations=0.03
max_creates=3000000

# rebalance
rebalance_after=1500000
num_nodes_after=3
reb_max_retries=5

# control (defaults: pytests/performance/perf_defaults.py)
load_wait_until_drained=1
loop_wait_until_drained=0
mcsoda_heartbeat=3
mcsoda_max_ops_sec=300
tear_down=1
tear_down_proxy=1
tear_down_bucket=0
tear_down_cluster=1
tear_down_on_setup=0
```

```
tahoe-dedicated.ini
[global]
username:root
password:couchbase
port:8091
data_path:/data

[servers]
1:192.168.0.20
2:192.168.0.21
3:192.168.0.22
4:192.168.0.23

[clients]
1:192.168.0.24
2:192.168.0.25
3:192.168.0.26
4:192.168.0.27
5:192.168.0.28
6:192.168.0.29

[membase]
rest_username:Administrator
rest_password:password

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