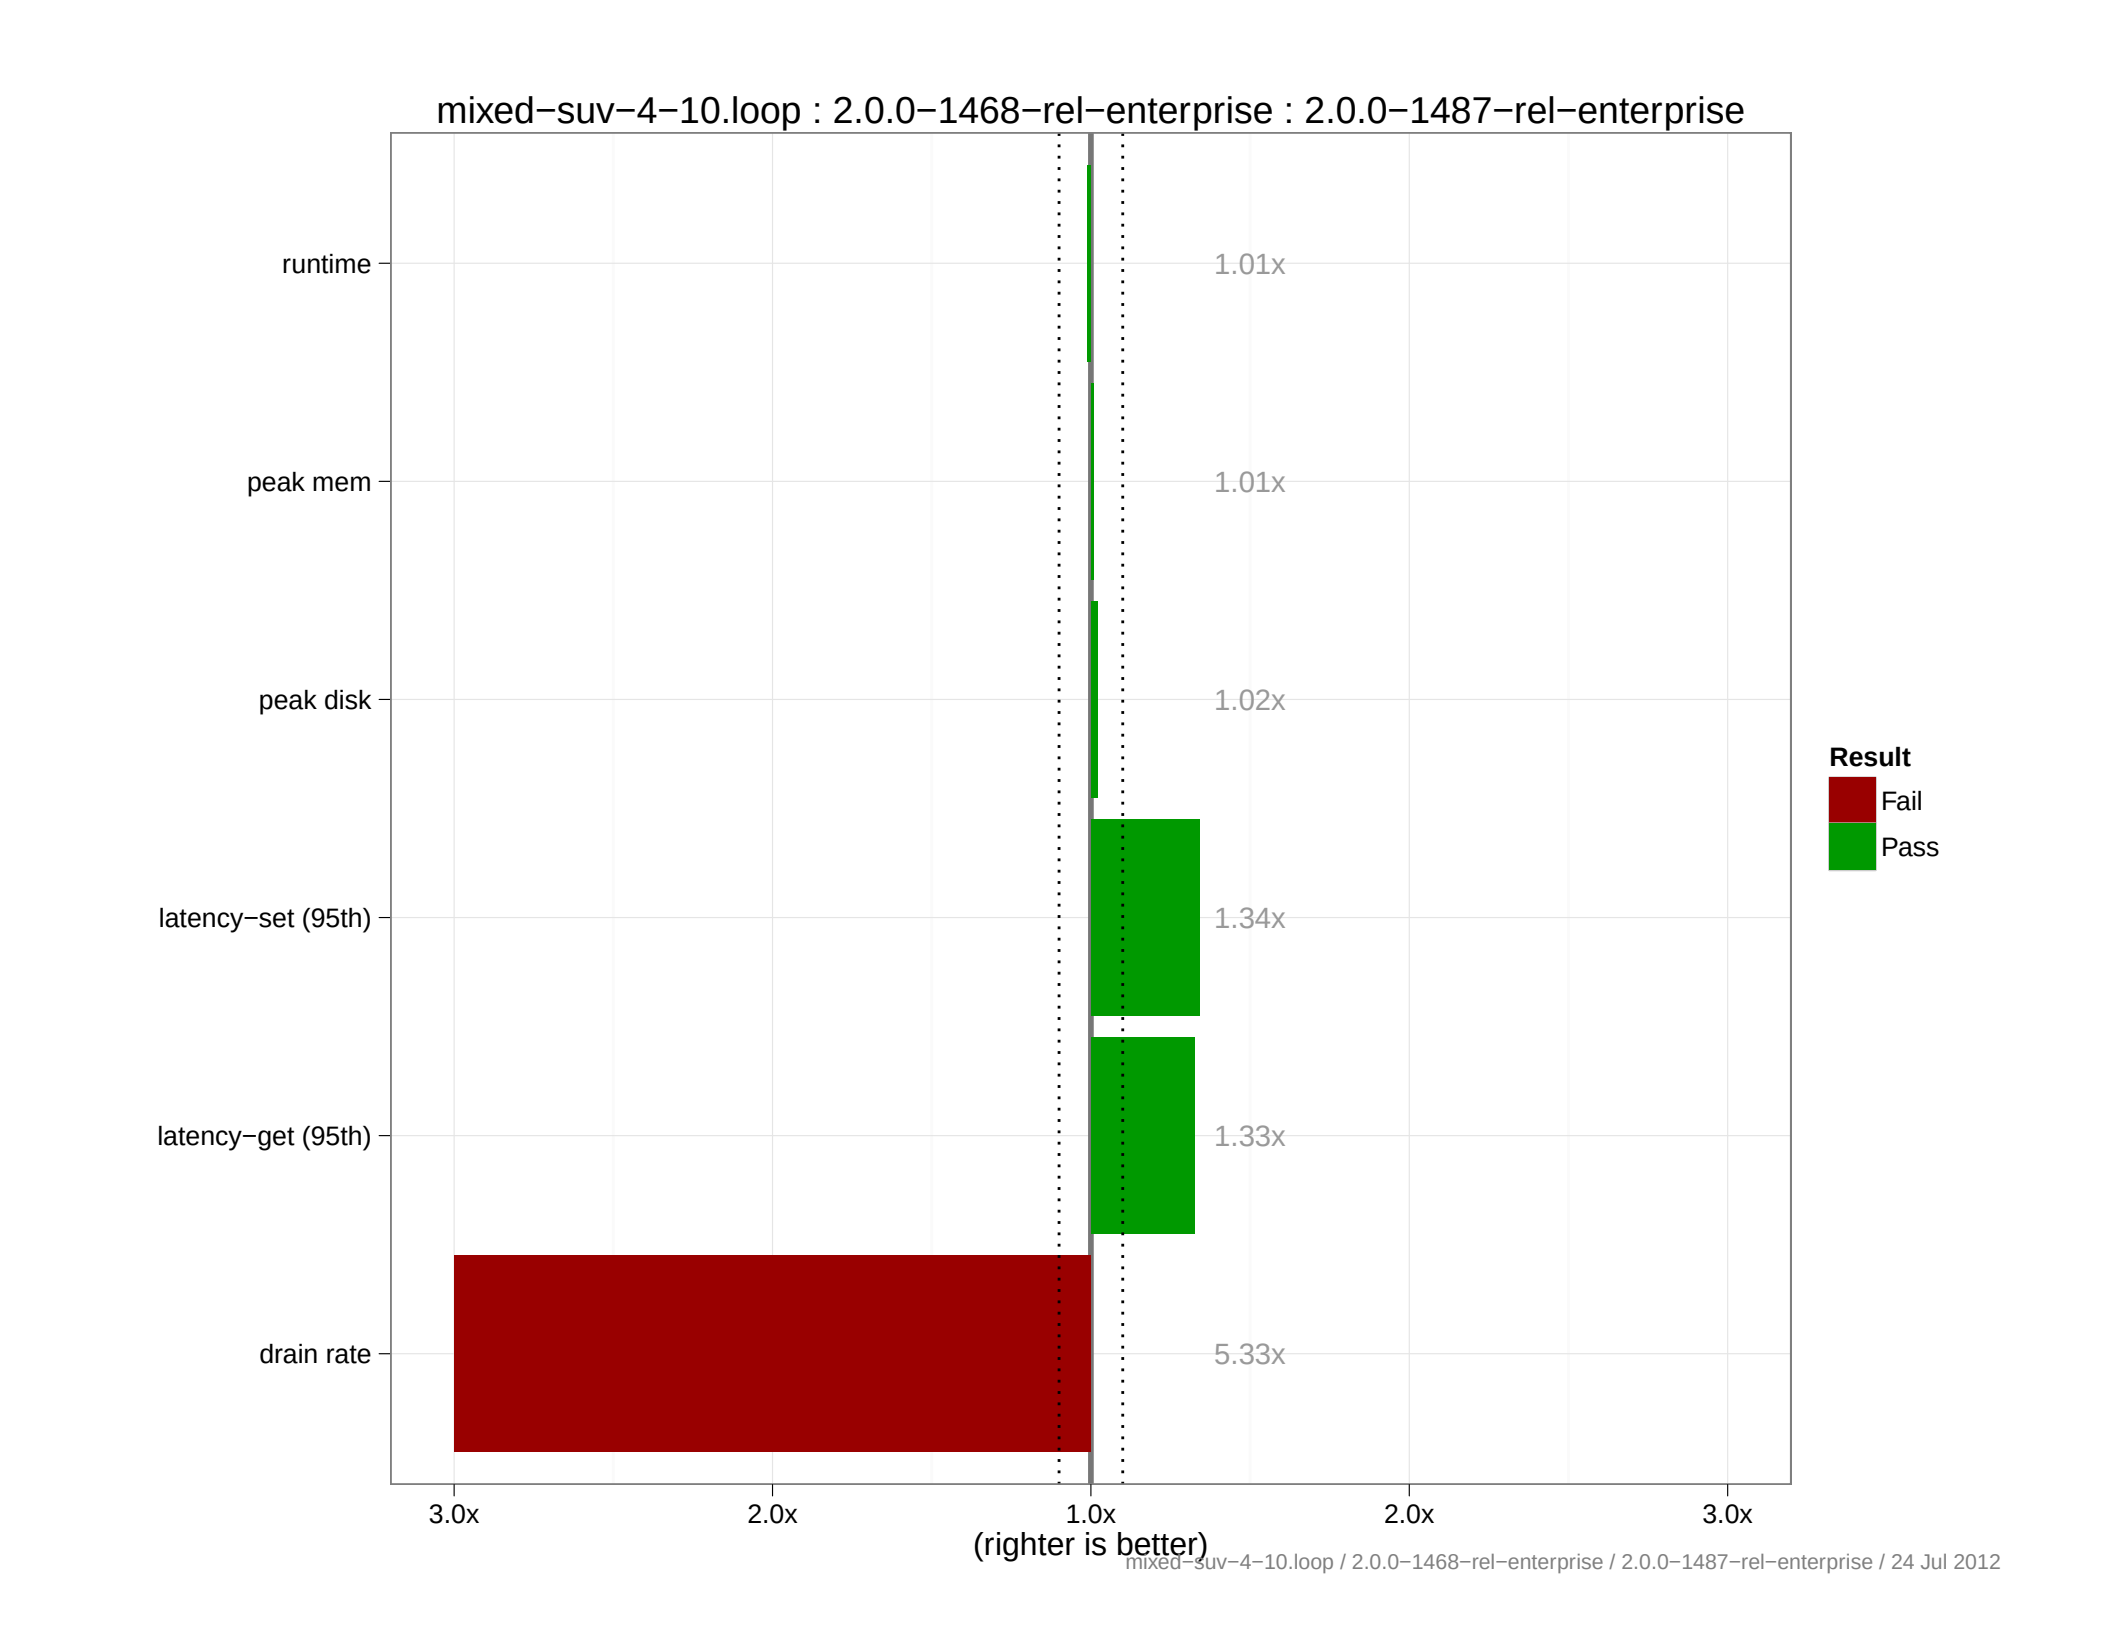
mixed−suv−4−10.loop : 2.0.0−1468−rel−enterprise : 2.0.0−1487−rel−enterprise
runtime
peak mem
peak disk
latency−set (95th)
latency−get (95th)
drain rate
1.01x
1.01x
1.02x
1.34x
1.33x
5.33x
Result
Fail
Pass
3.0x
2.0x
1.0x
2.0x
3.0x
(righter is better)
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

	2.0.0 − 1468	2.0.0 − 1487
Runtime (in hr)	2.89	2.92
Avg. Drain Rate	8.93K	1.68K
Peak Disk (GB)	37.54	36.76
Peak Memory (GB)	16.88	16.72
Avg. OPS	7.47K	7.15K
Avg. mem memcached (GB)	16.47	16.26
Avg. mem beam.smp (MB)	359.61	370.28
Latency−get (90th) (ms)	0.81	0.62
Latency−get (95th) (ms)	1.23	0.93
Latency−get (99th) (ms)	3.81	2.57
Latency−set (90th) (ms)	0.84	0.63
Latency−set (95th) (ms)	1.26	0.94
Latency−set (99th) (ms)	3.69	2.5
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
Rebalance Time (sec)	0	0
Testrunner Version	2e4de23	2718316

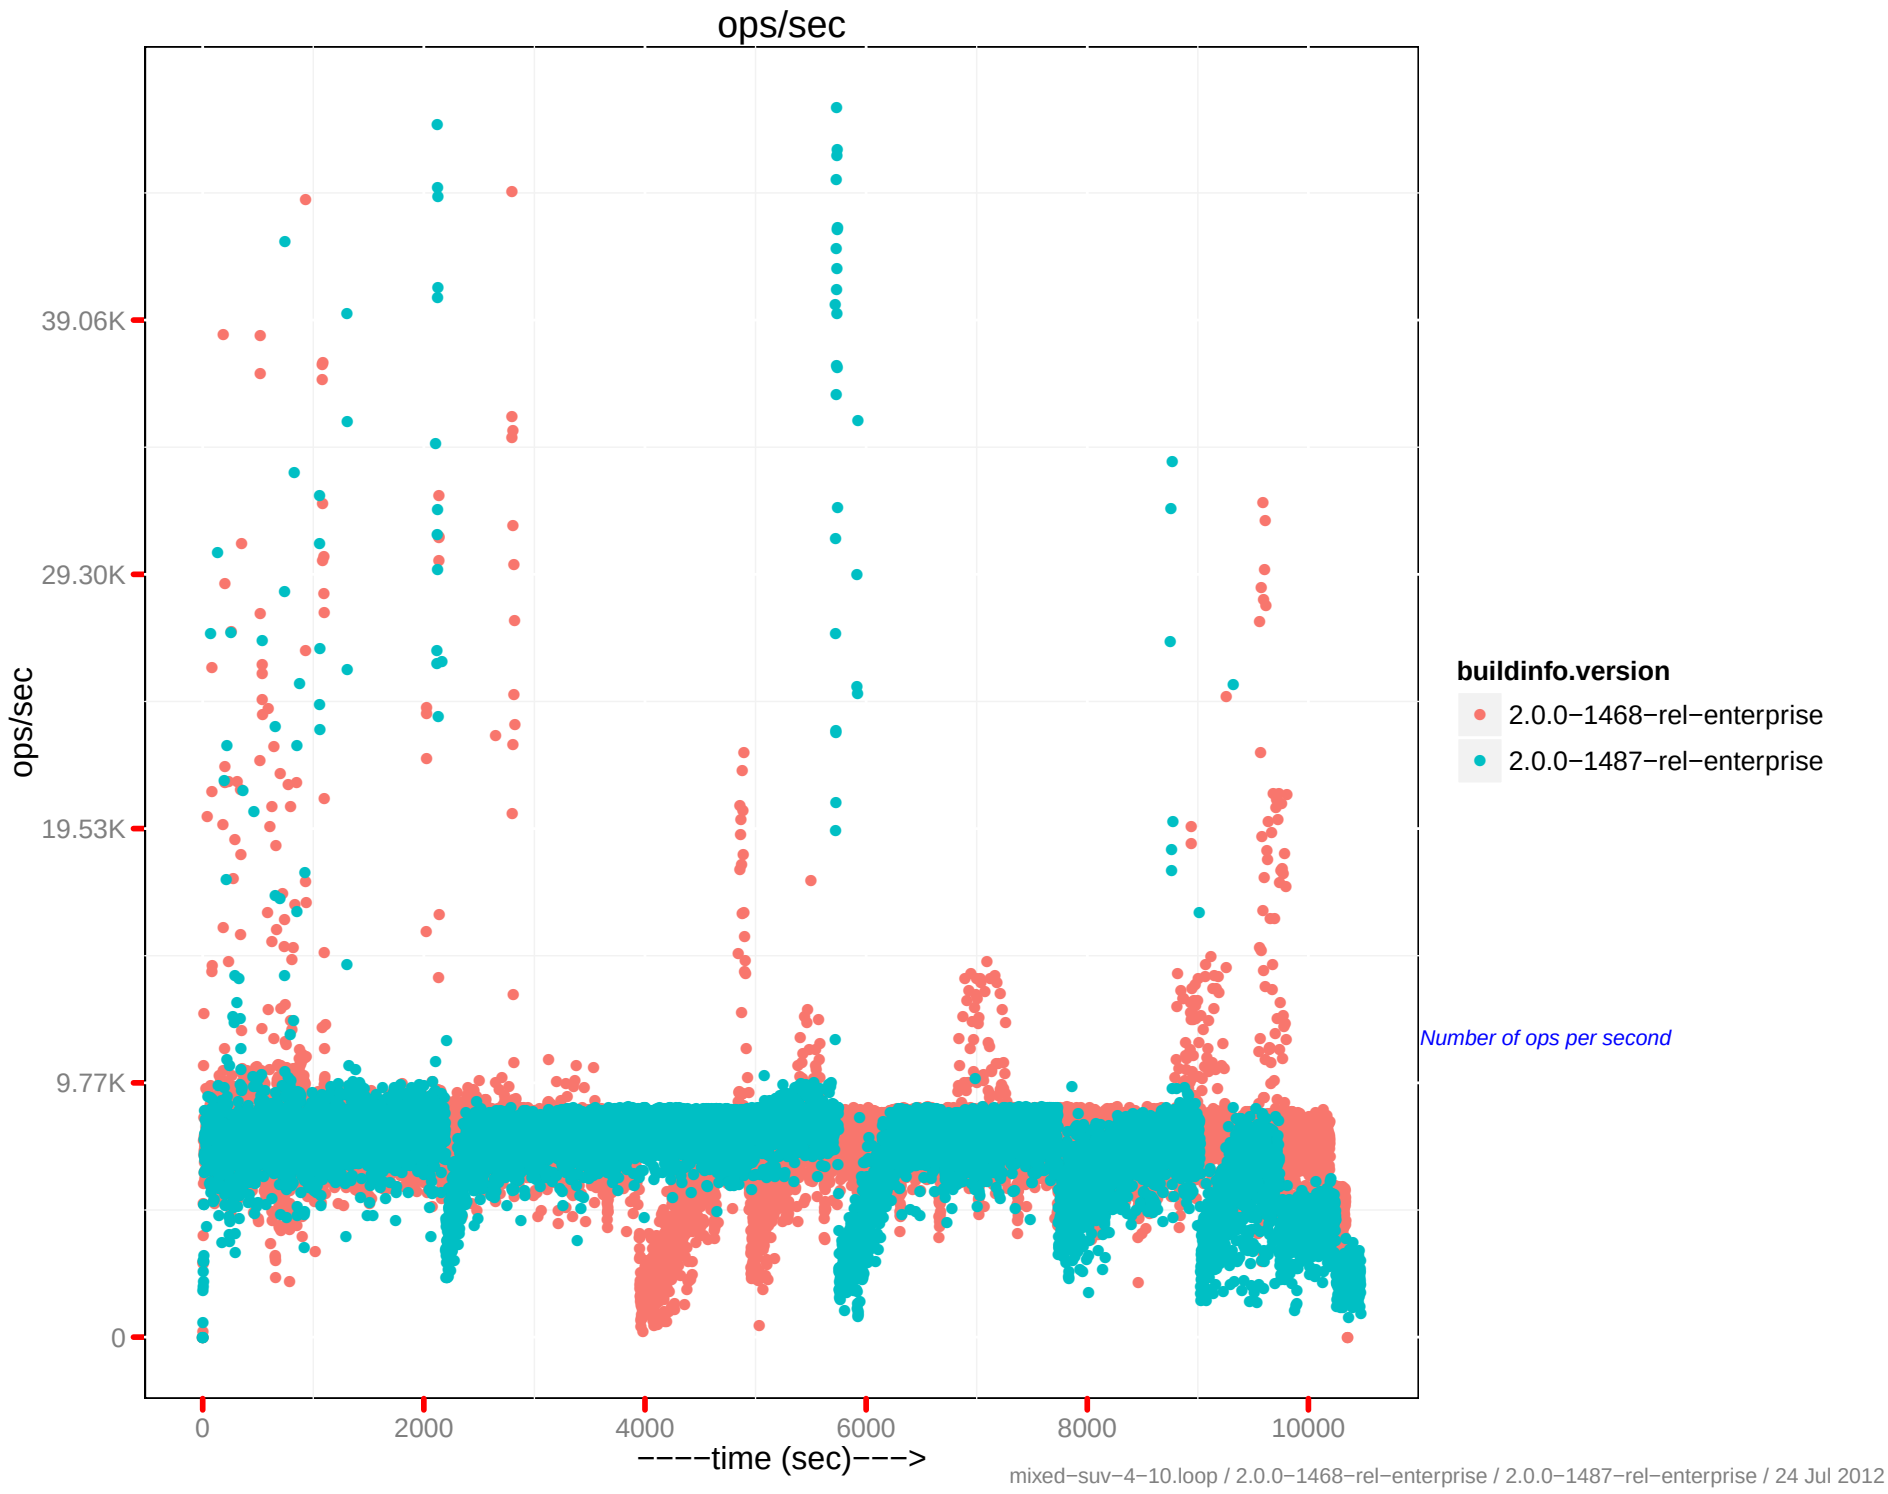
ops/sec
39.06K
29.30K
19.53K
9.77K
0
ops/sec
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of ops per second
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

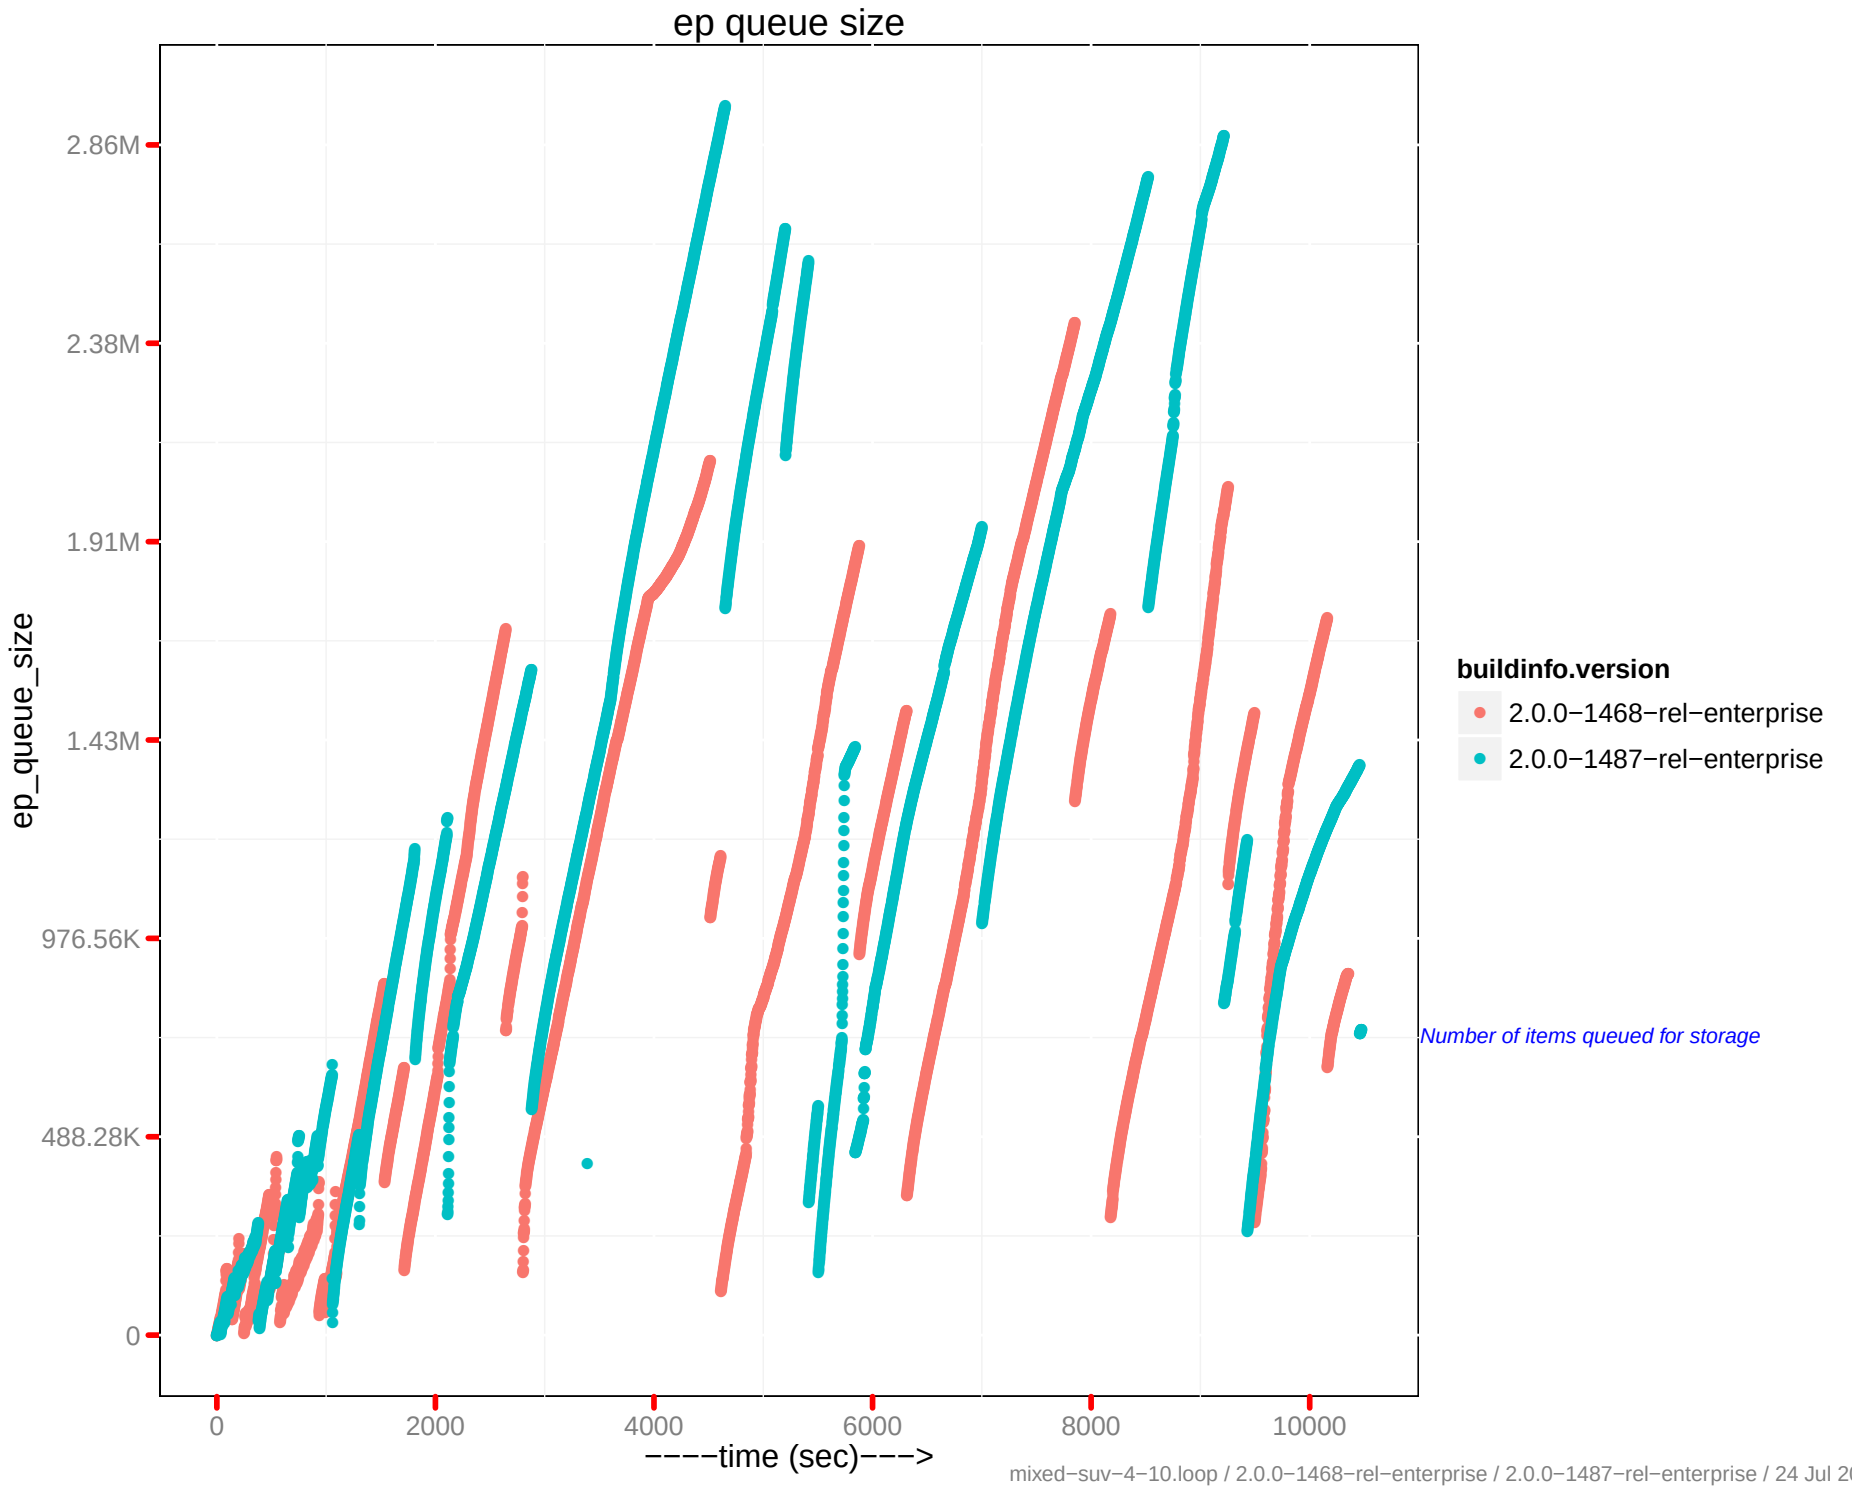
ep queue size
ep_queue_size
2.86M
2.38M
1.91M
1.43M
976.56K
488.28K
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of items queued for storage
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

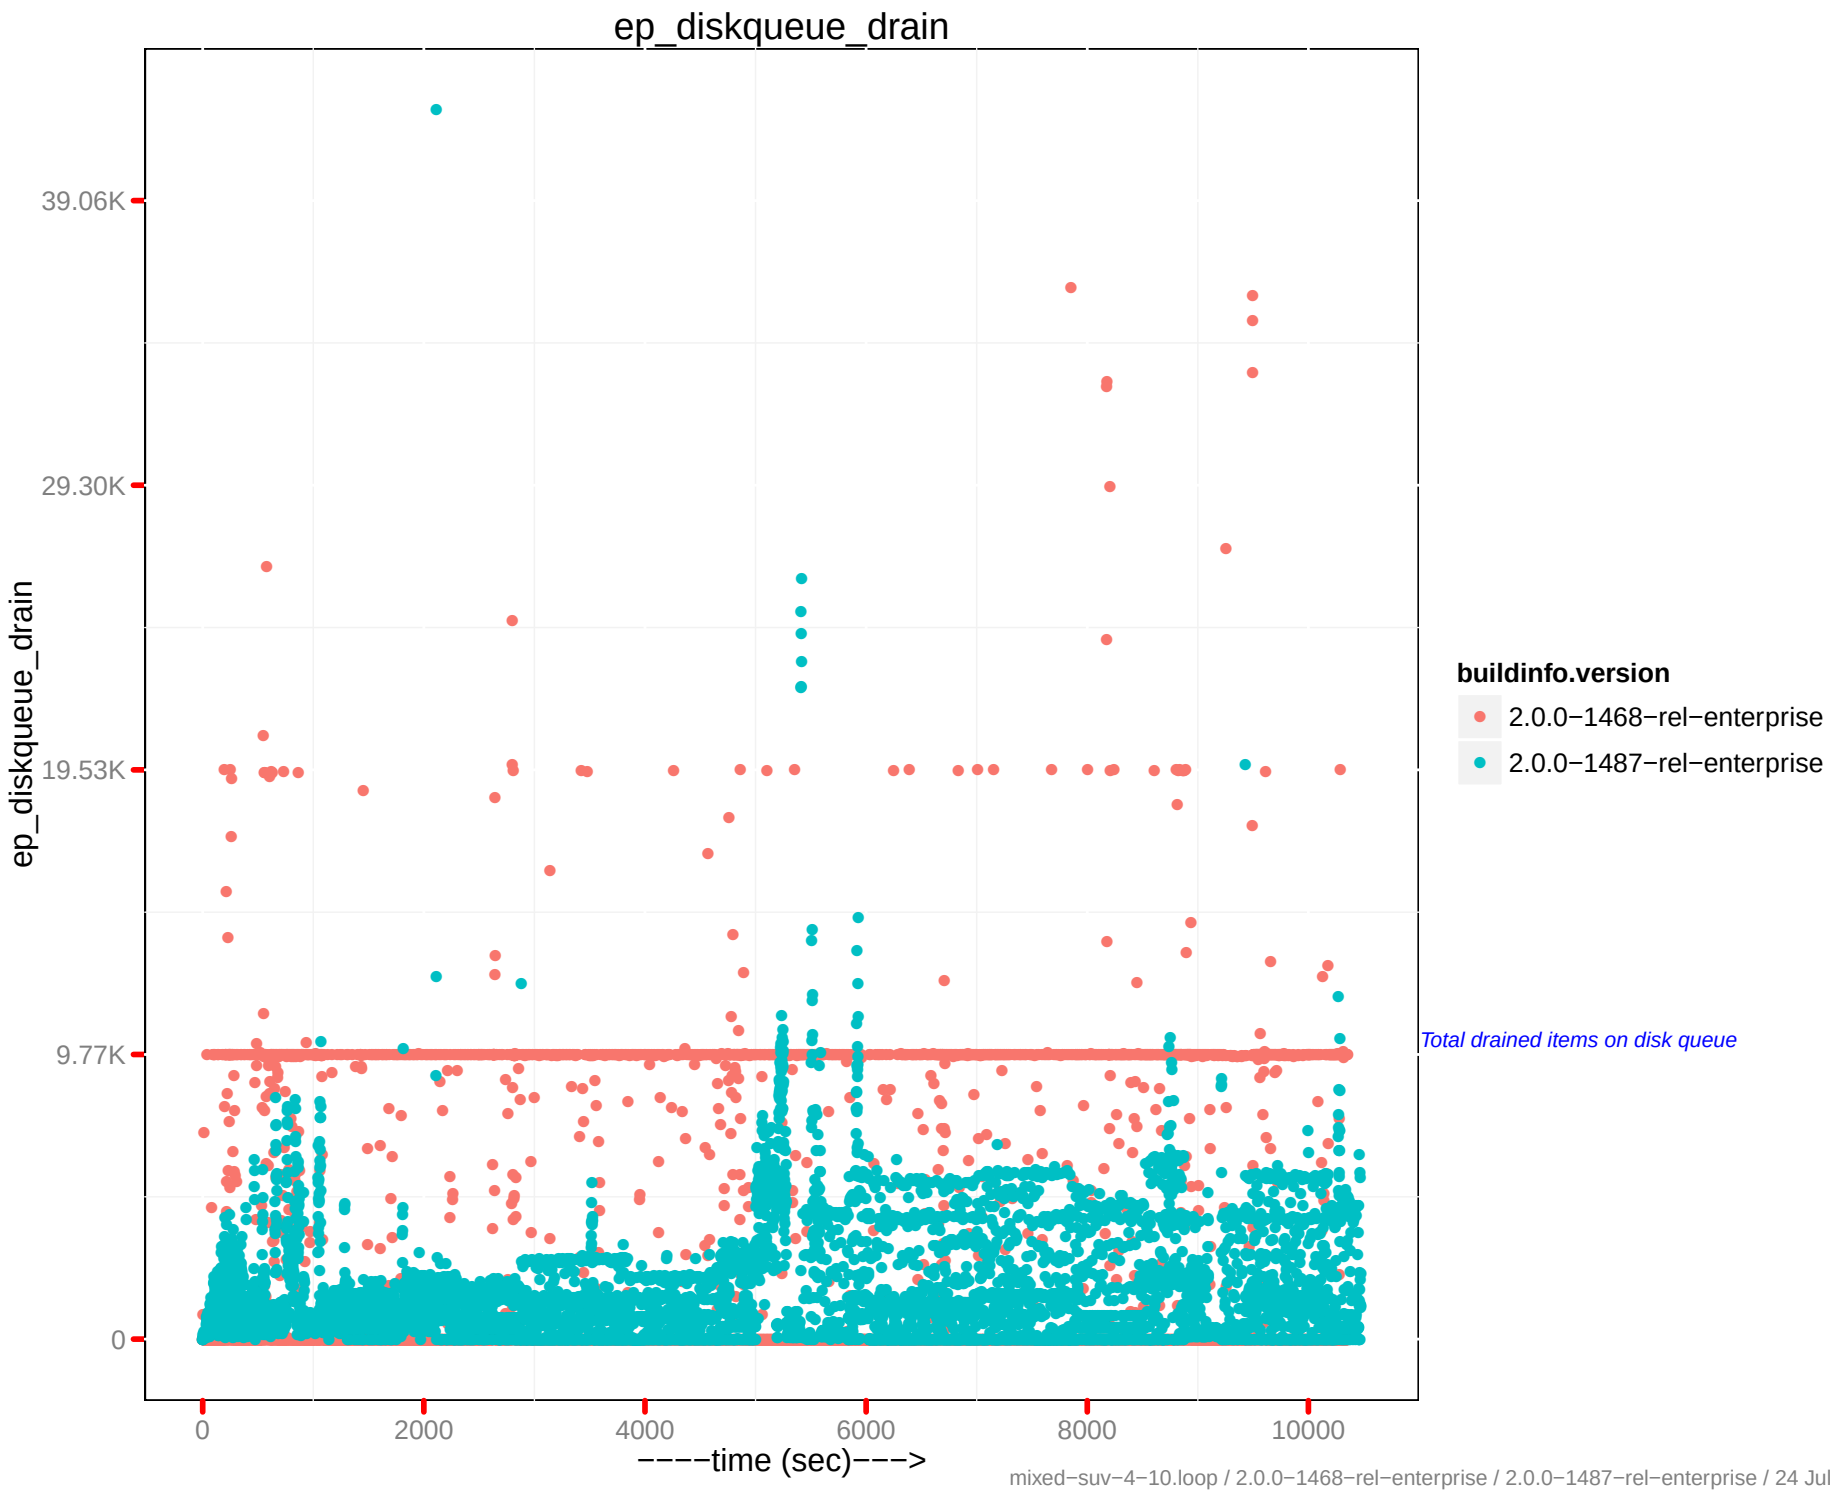
ep_diskqueue_drain
39.06K
29.30K
19.53K
9.77K
0
ep_diskqueue_drain
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Total drained items on disk queue
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

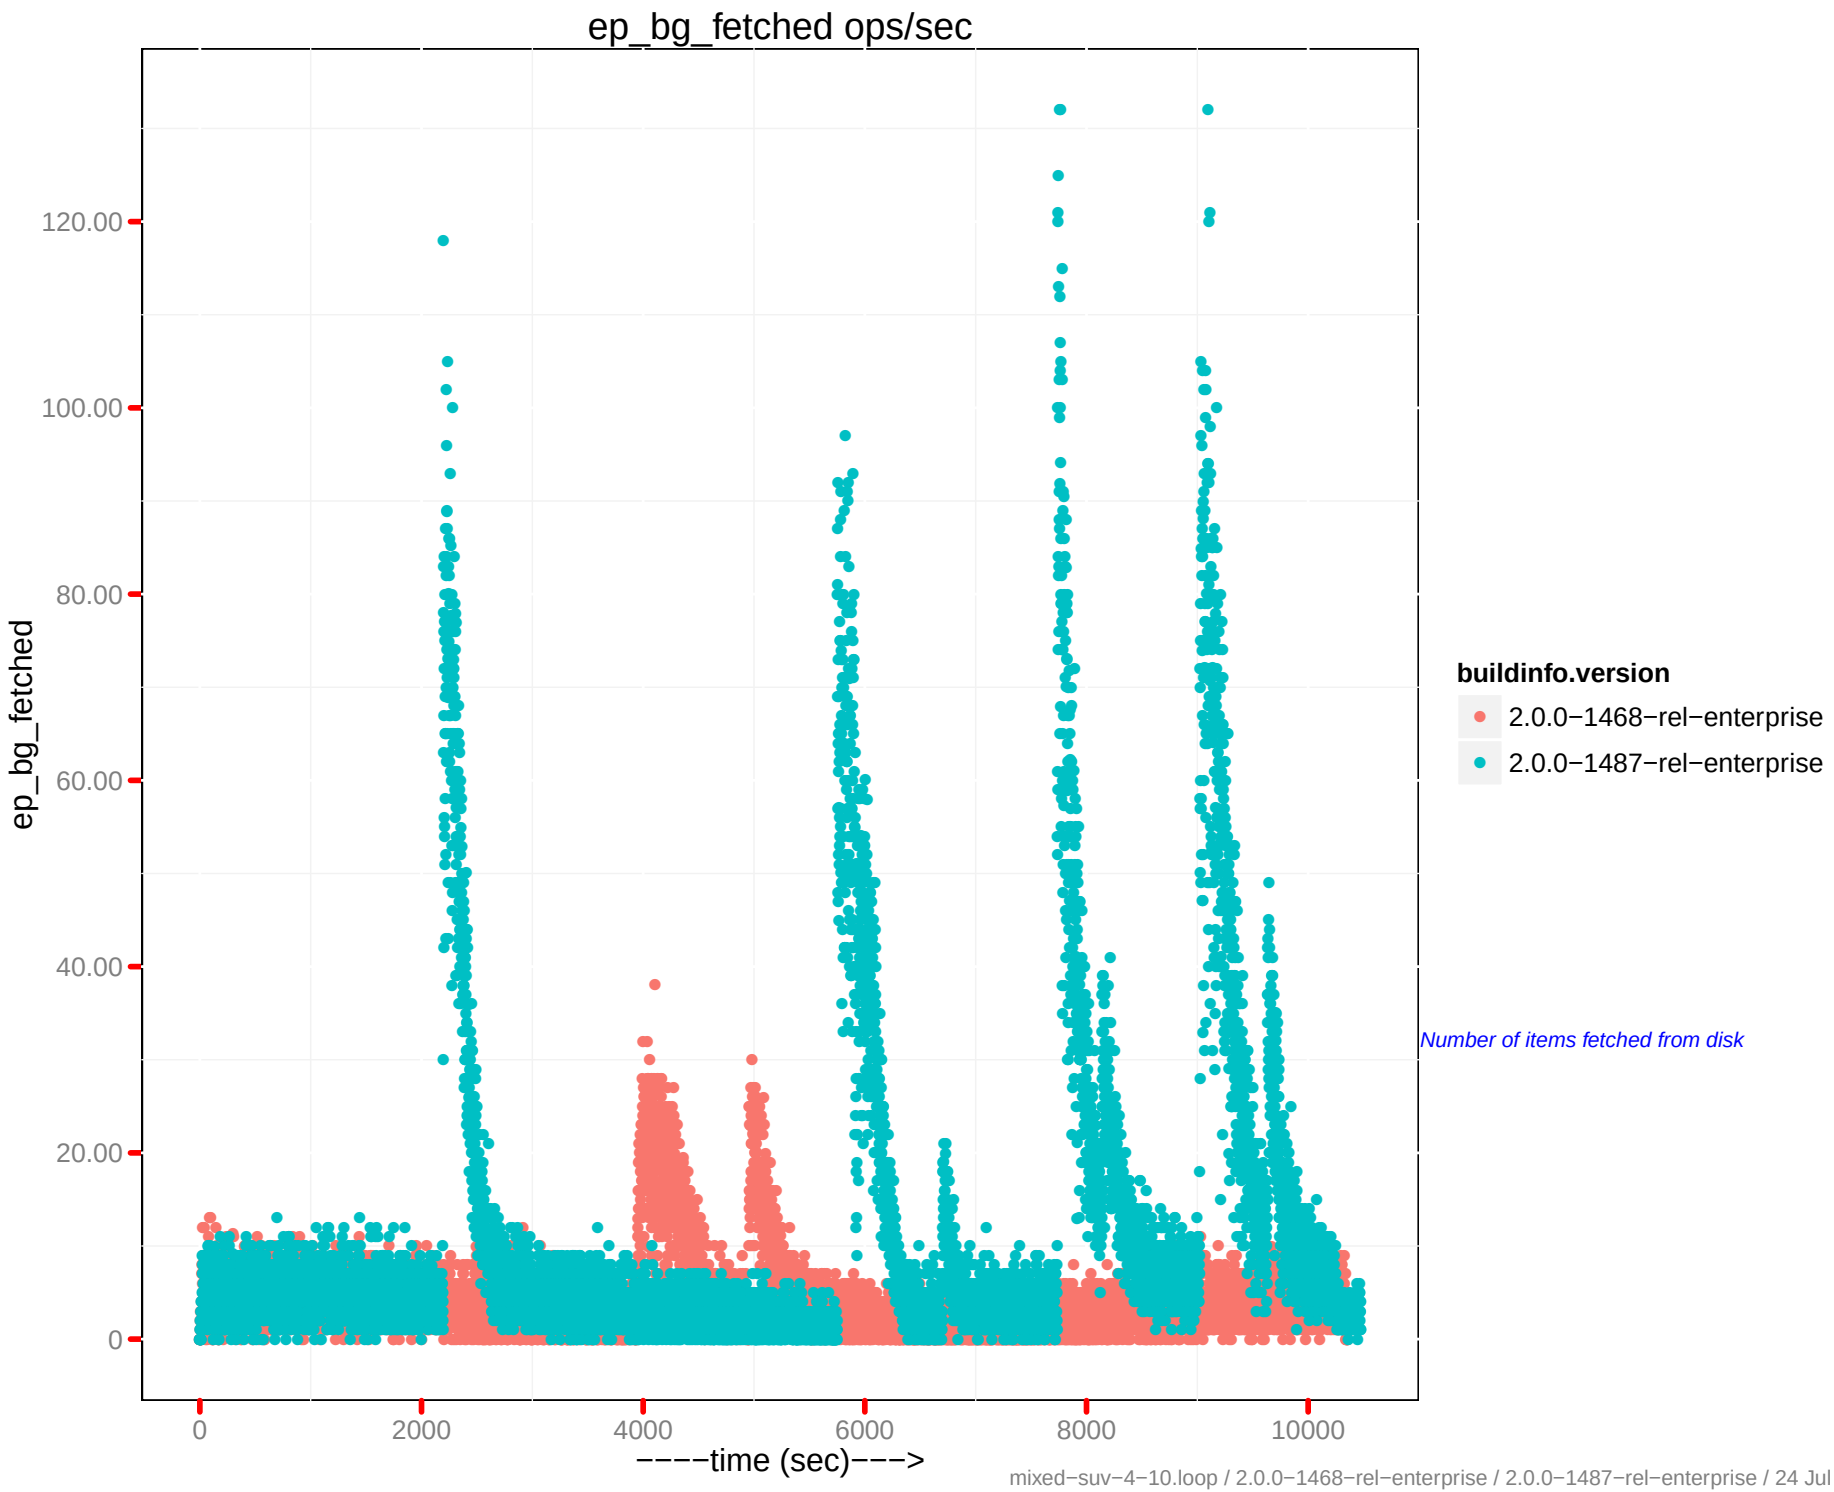
ep_bg_fetched ops/sec
ep_bg_fetched
120.00
100.00
80.00
60.00
40.00
20.00
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of items fetched from disk
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

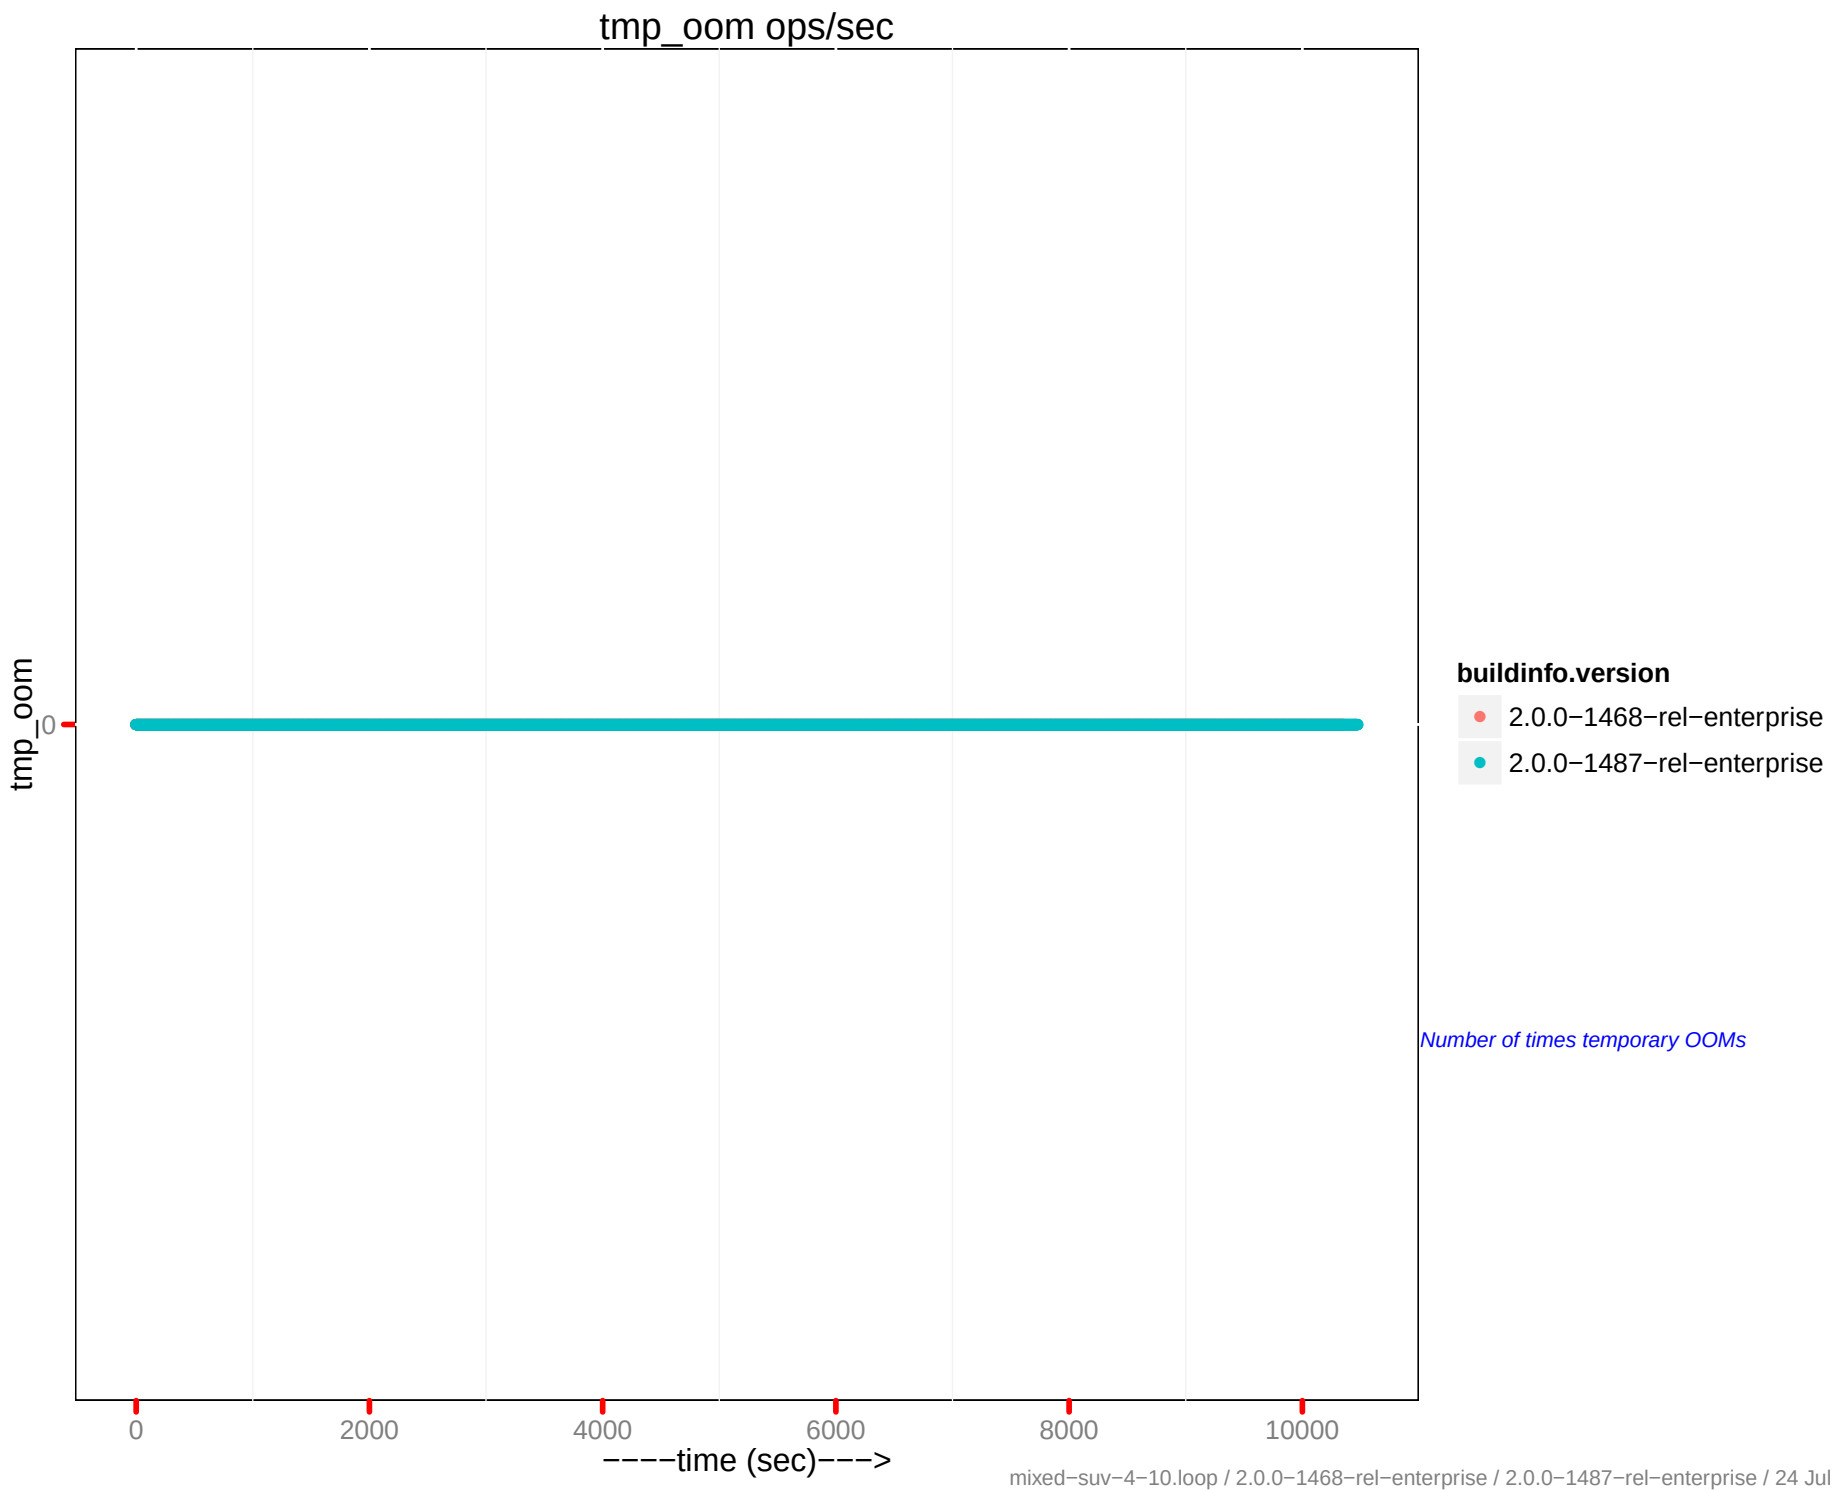
tmp_oom ops/sec
tmp_oom
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
Number of times temporary OOMs
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

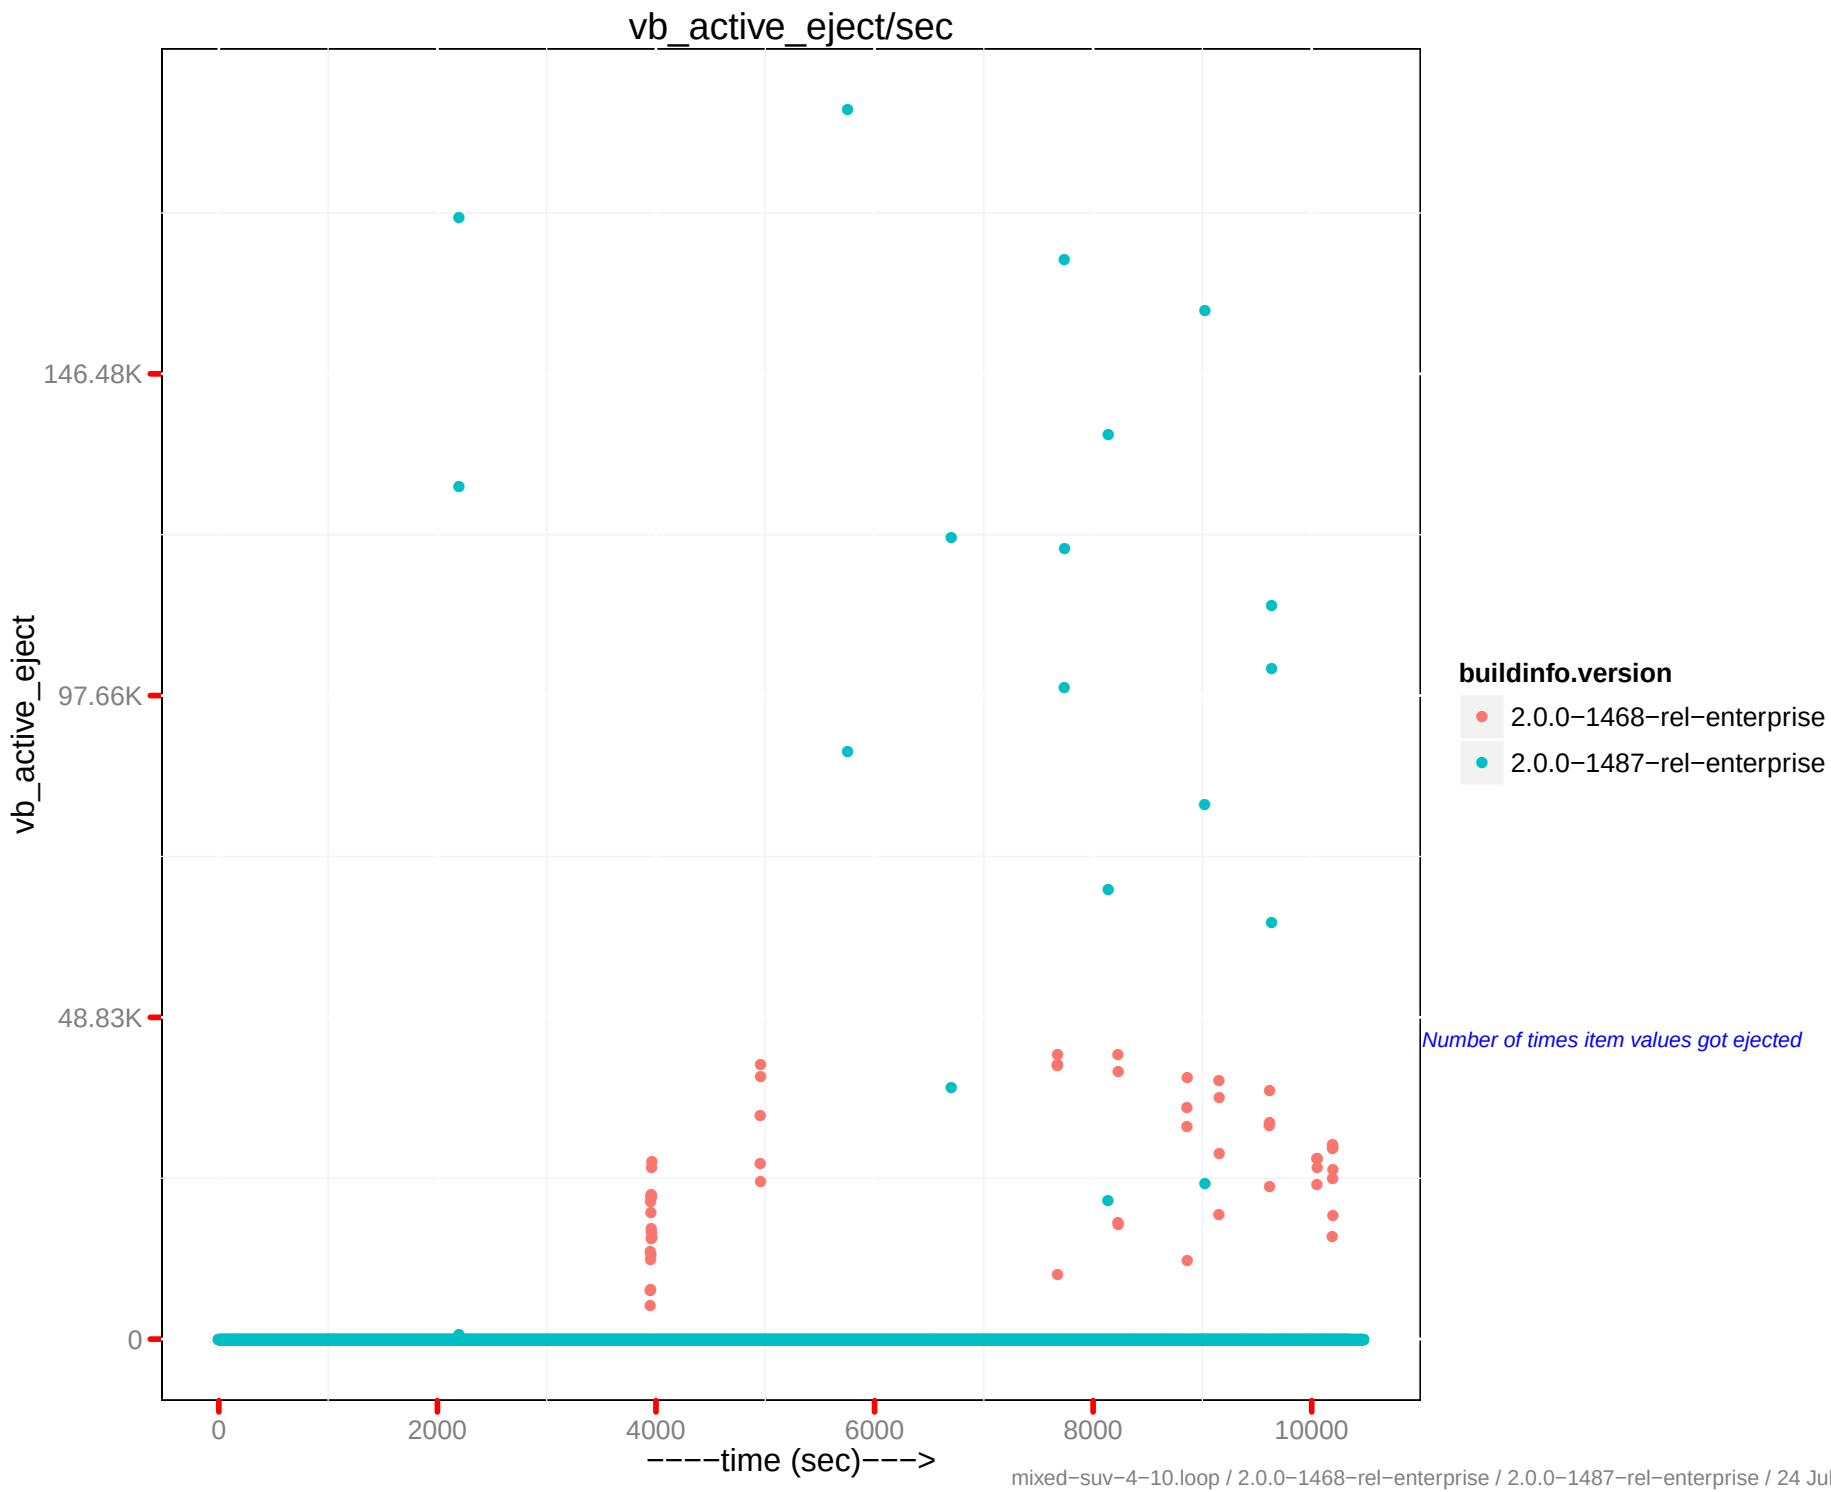
vb_active_eject/sec
vb_active_eject
146.48K
97.66K
48.83K
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
Number of times item values got ejected
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

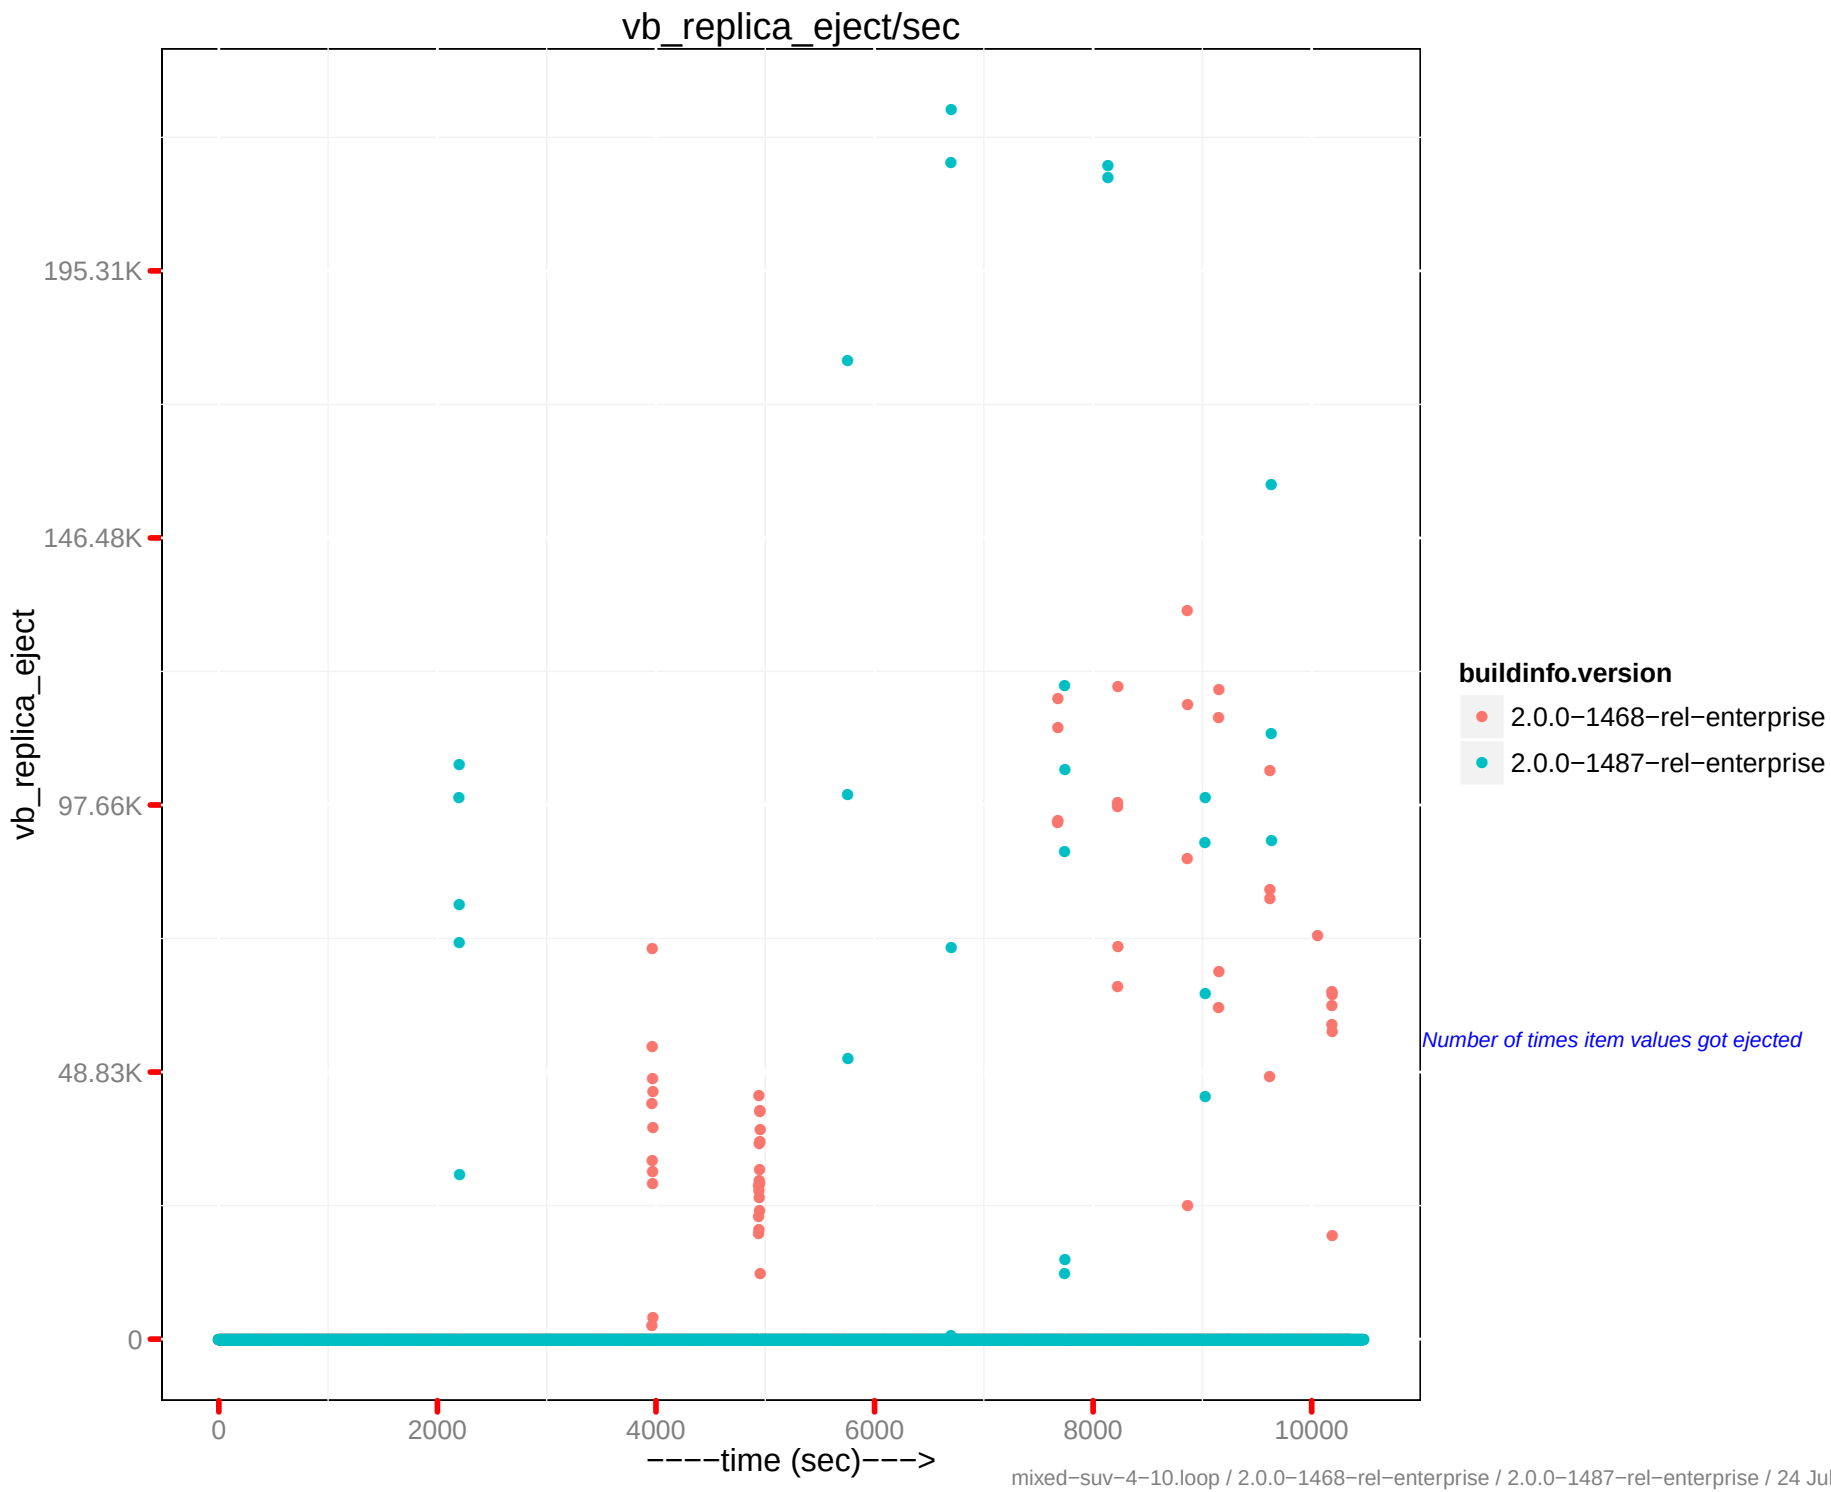
vb_replica_eject/sec
195.31K
146.48K
97.66K
48.83K
0
vb_replica_eject
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of times item values got ejected
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

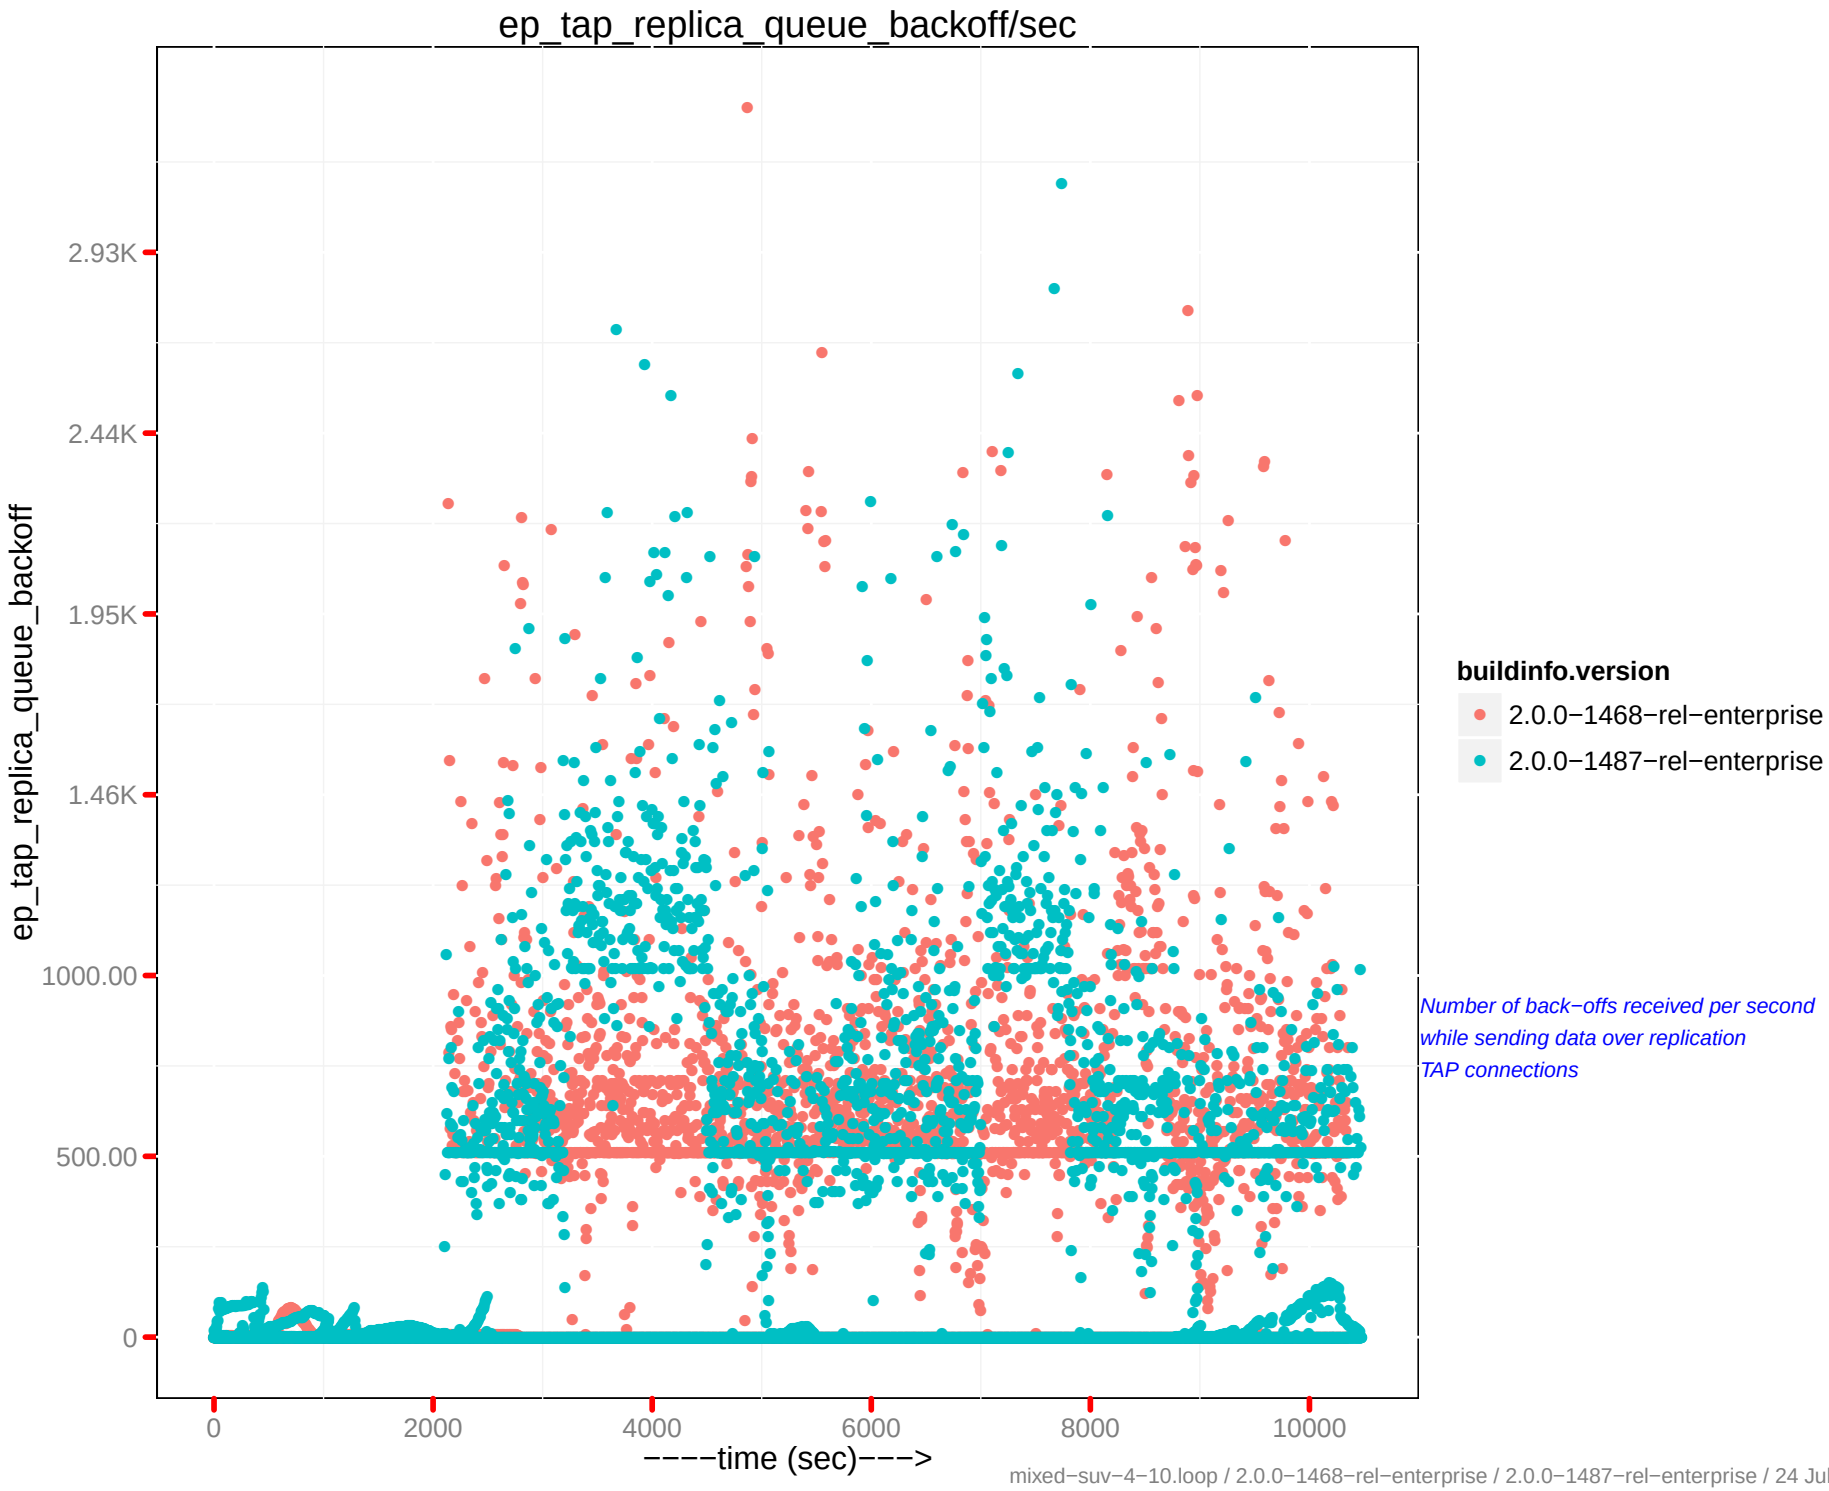
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
2.93K
2.44K
1.95K
1.46K
1000.00
500.00
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
Number of back−offs received per second
while sending data over replication
TAP connections
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

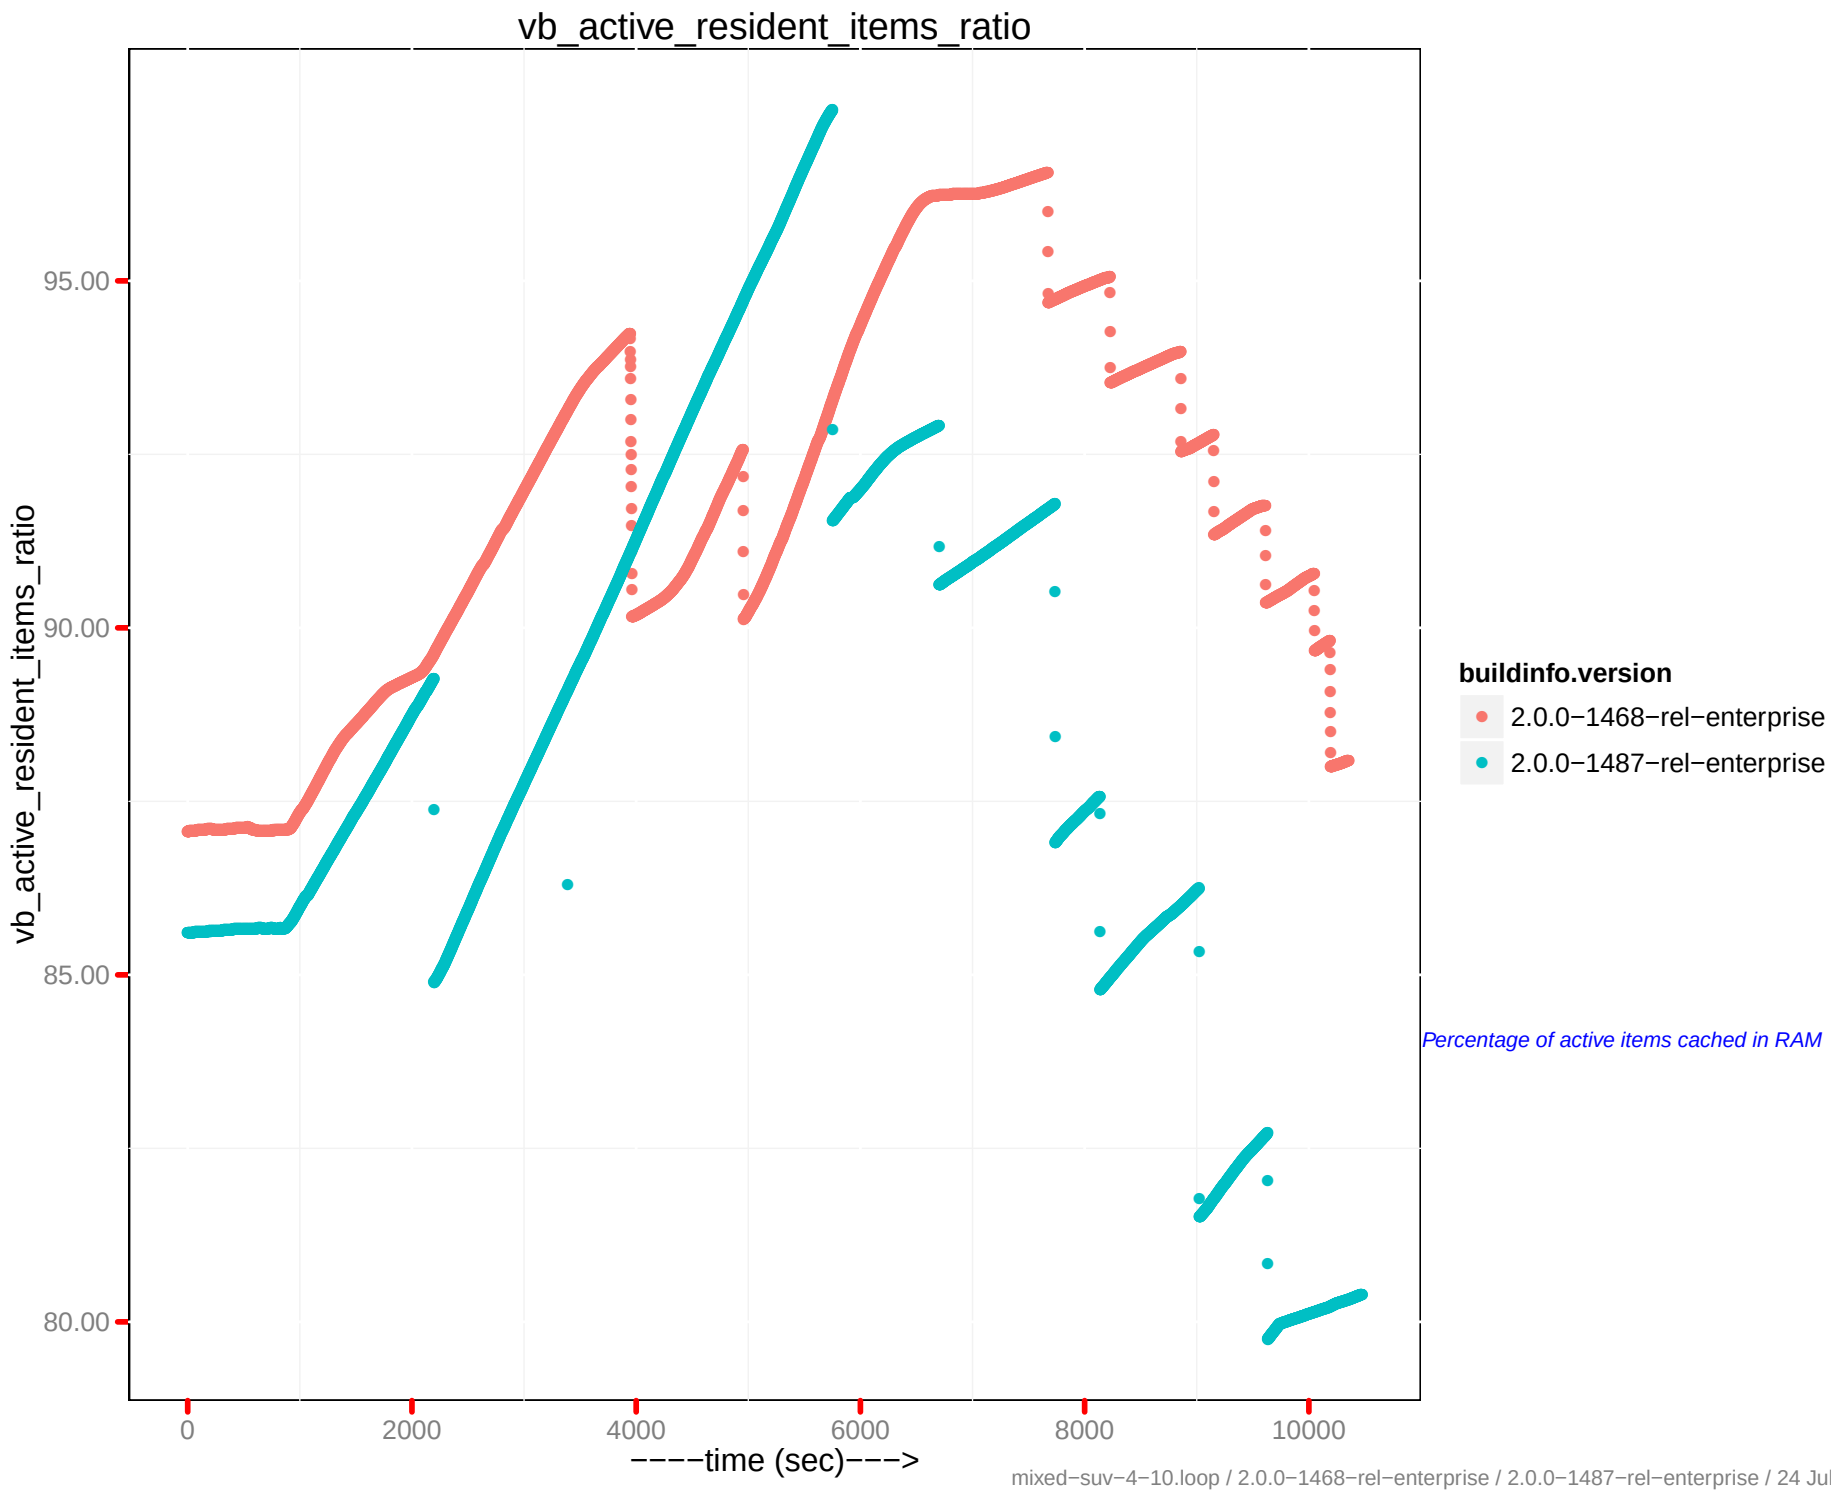
vb_active_resident_items_ratio
95.00
90.00
85.00
80.00
vb_active_resident_items_ratio
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Percentage of active items cached in RAM
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

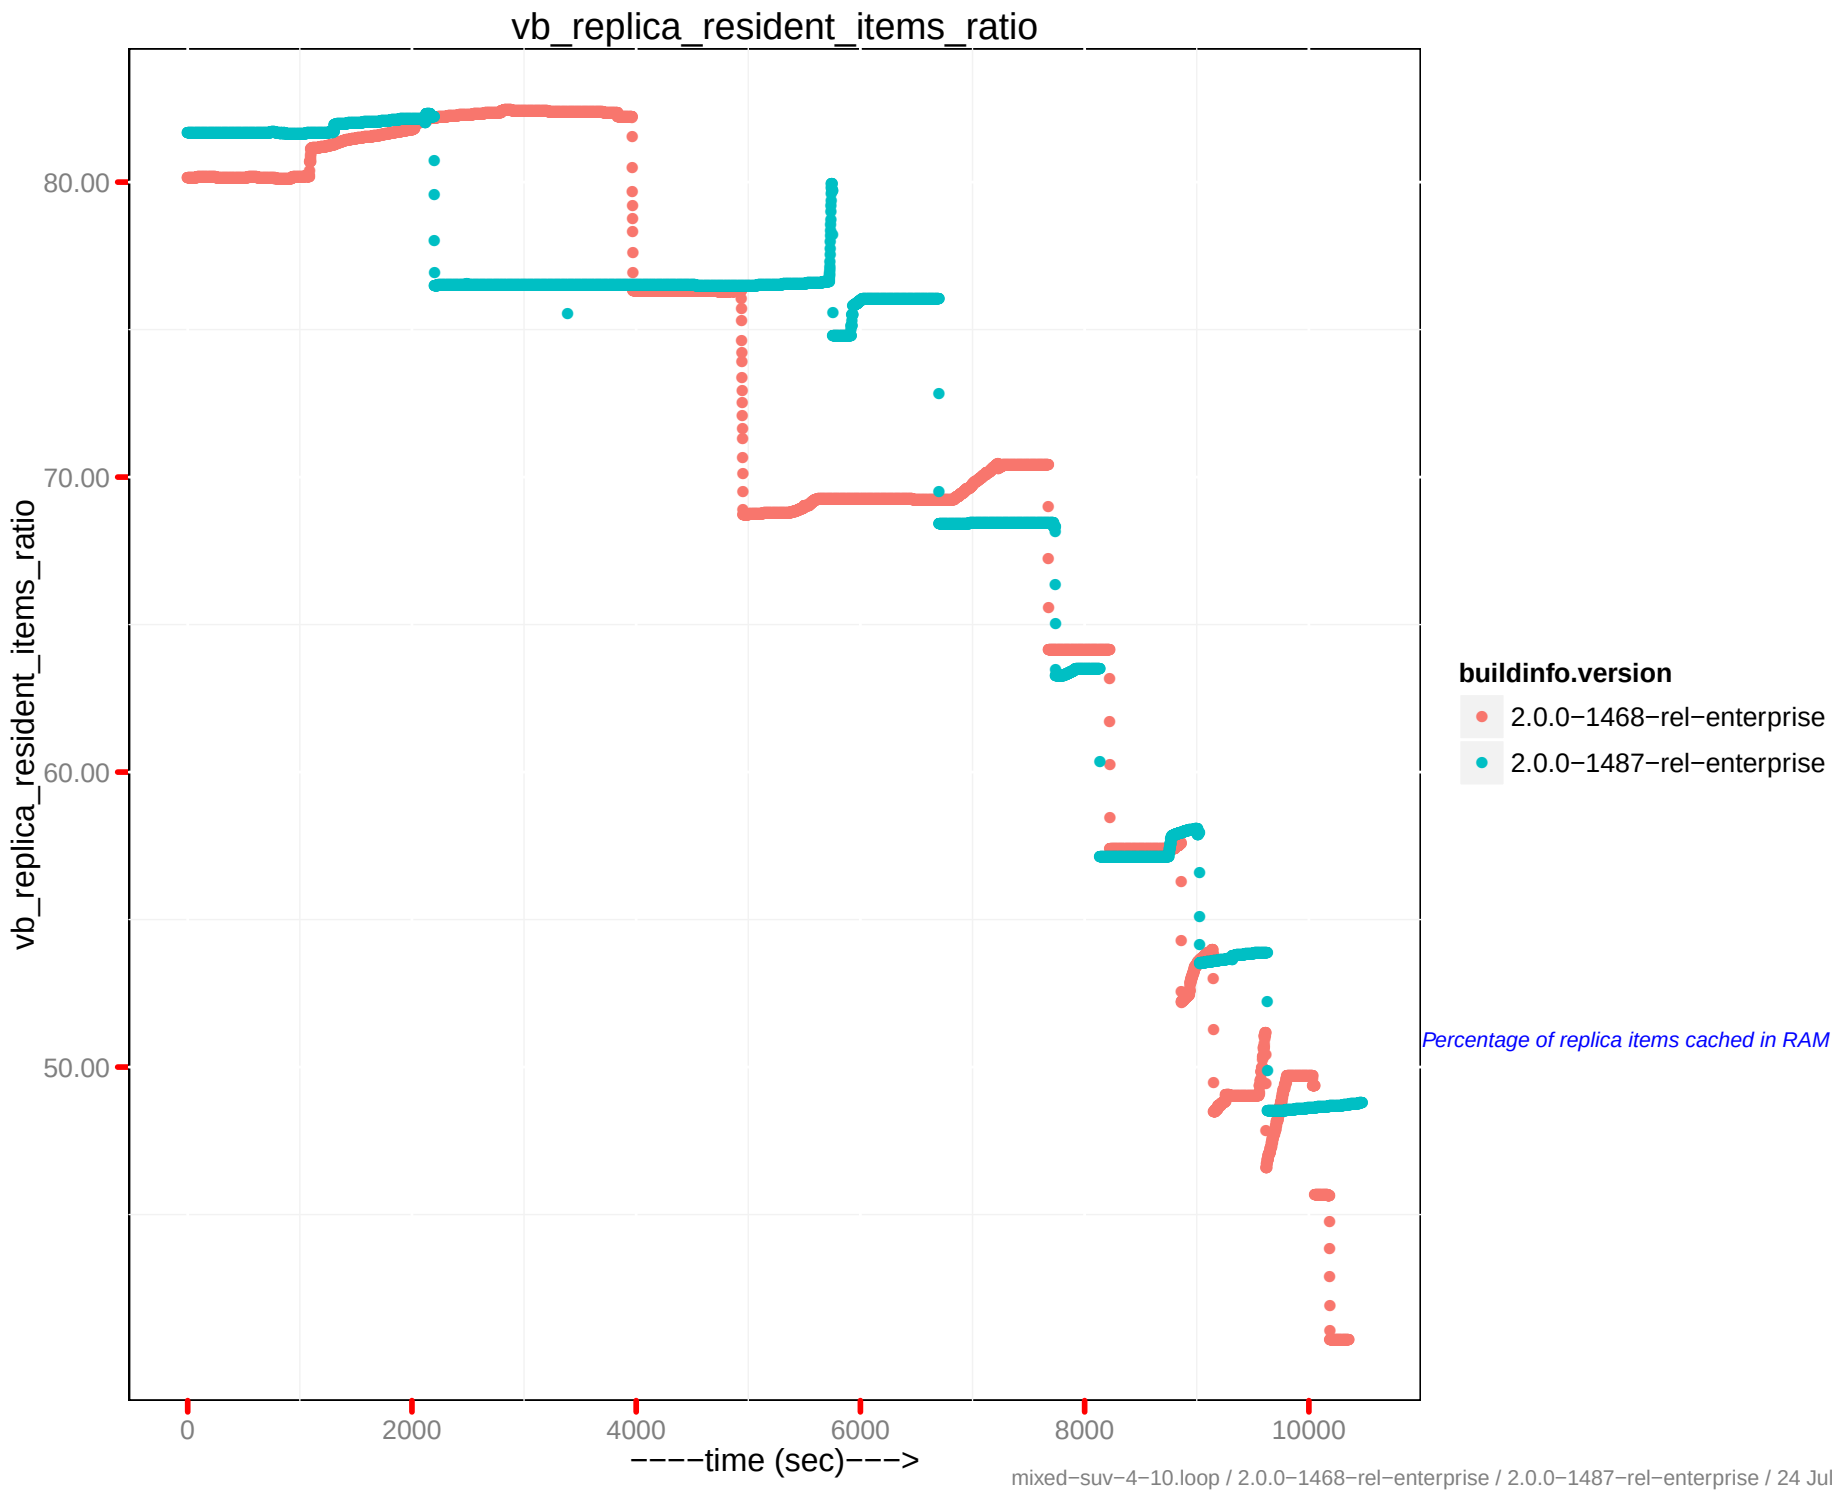
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
80.00
70.00
60.00
50.00
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Percentage of replica items cached in RAM
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

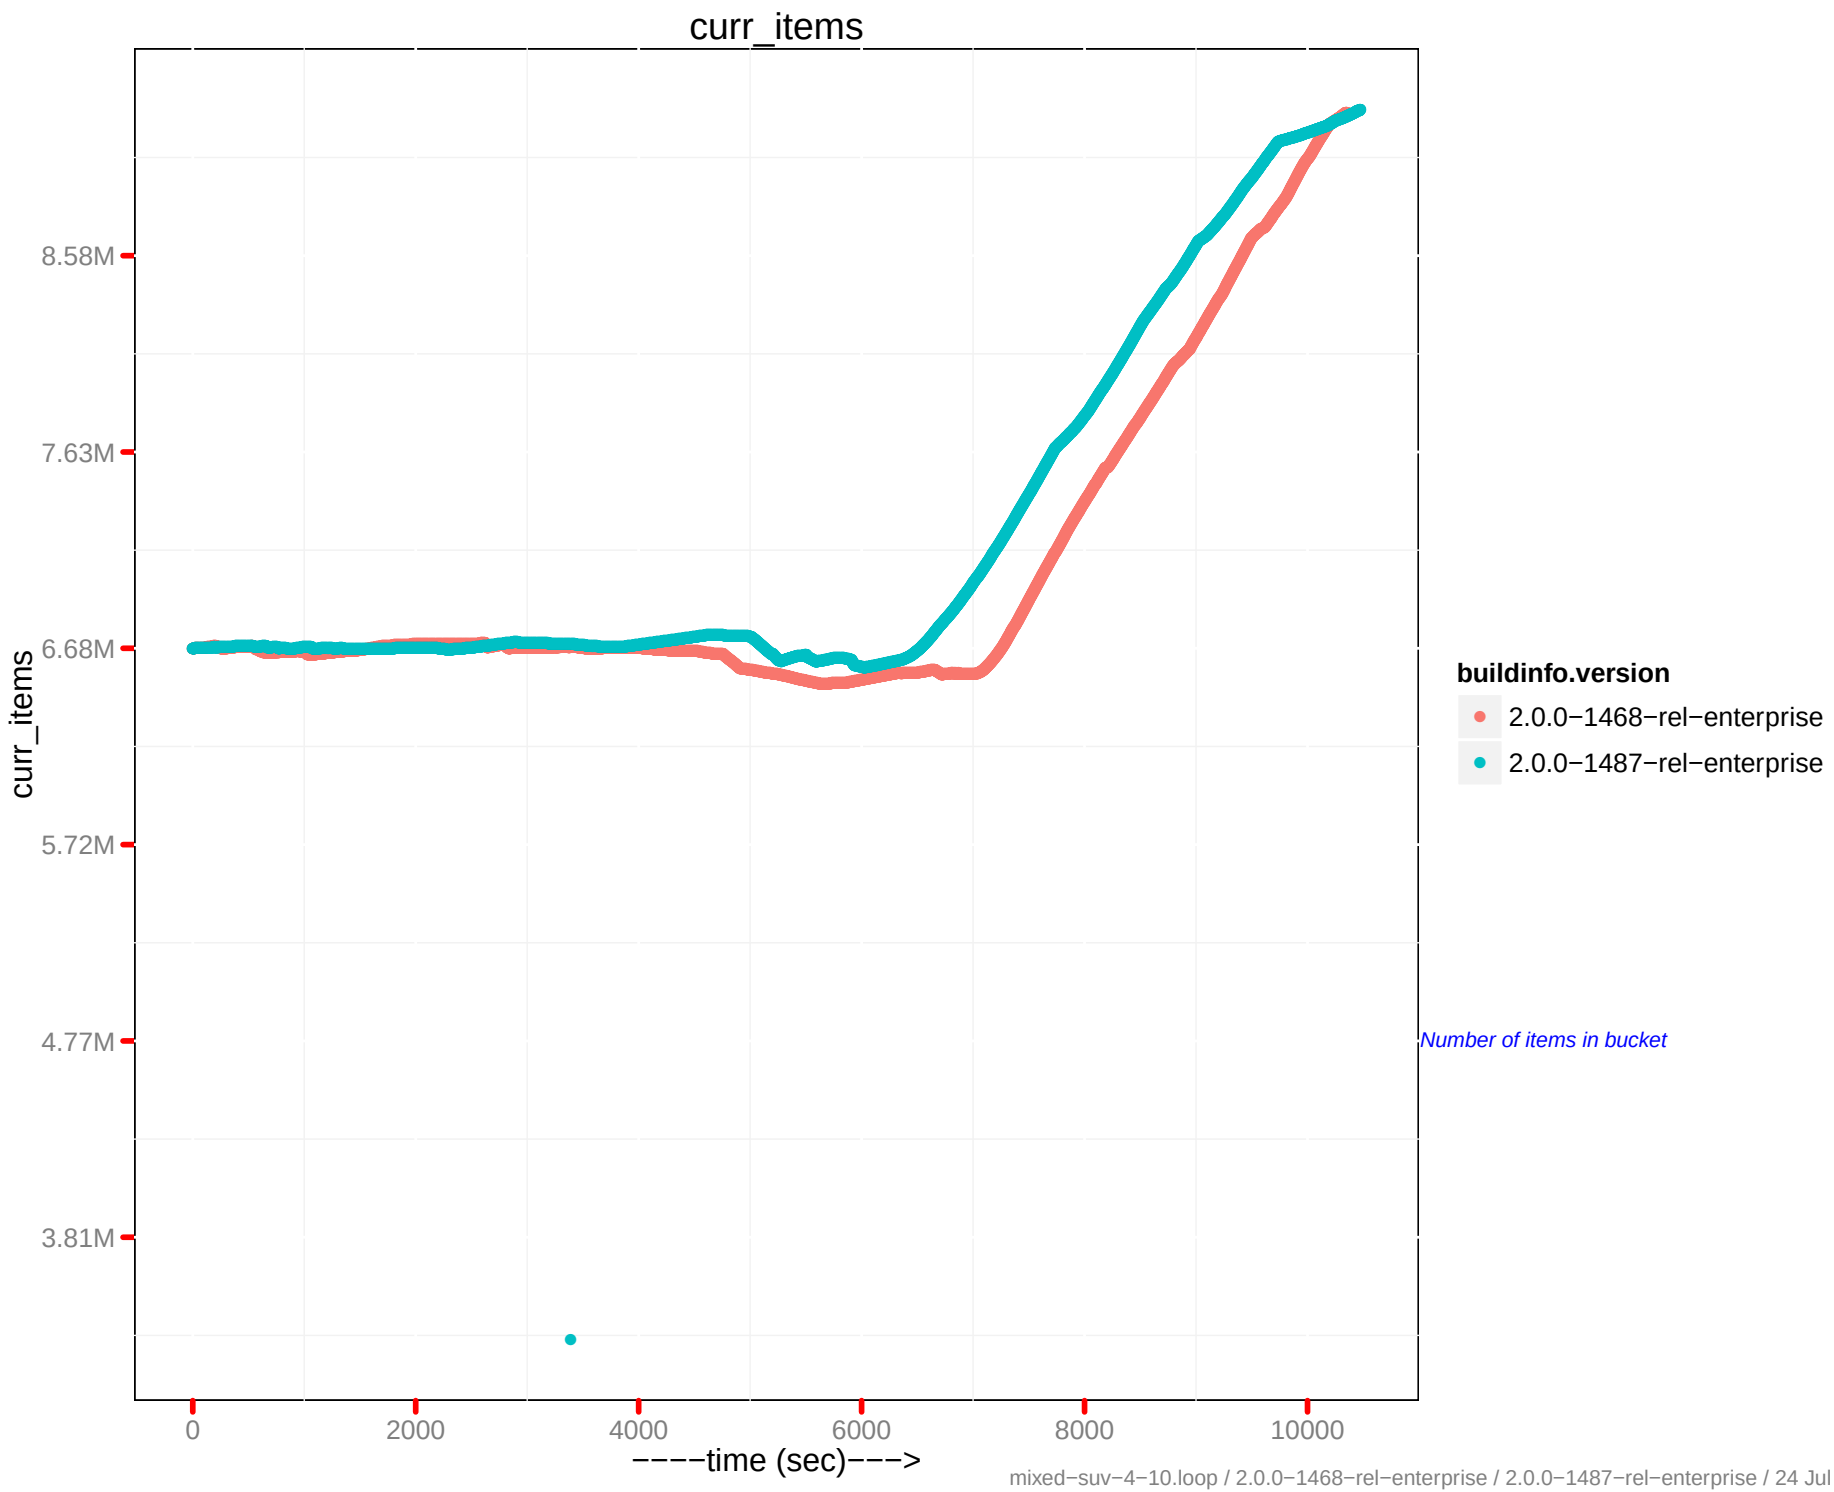
curr_items
8.58M
7.63M
6.68M
5.72M
4.77M
3.81M
curr_items
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of items in bucket
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

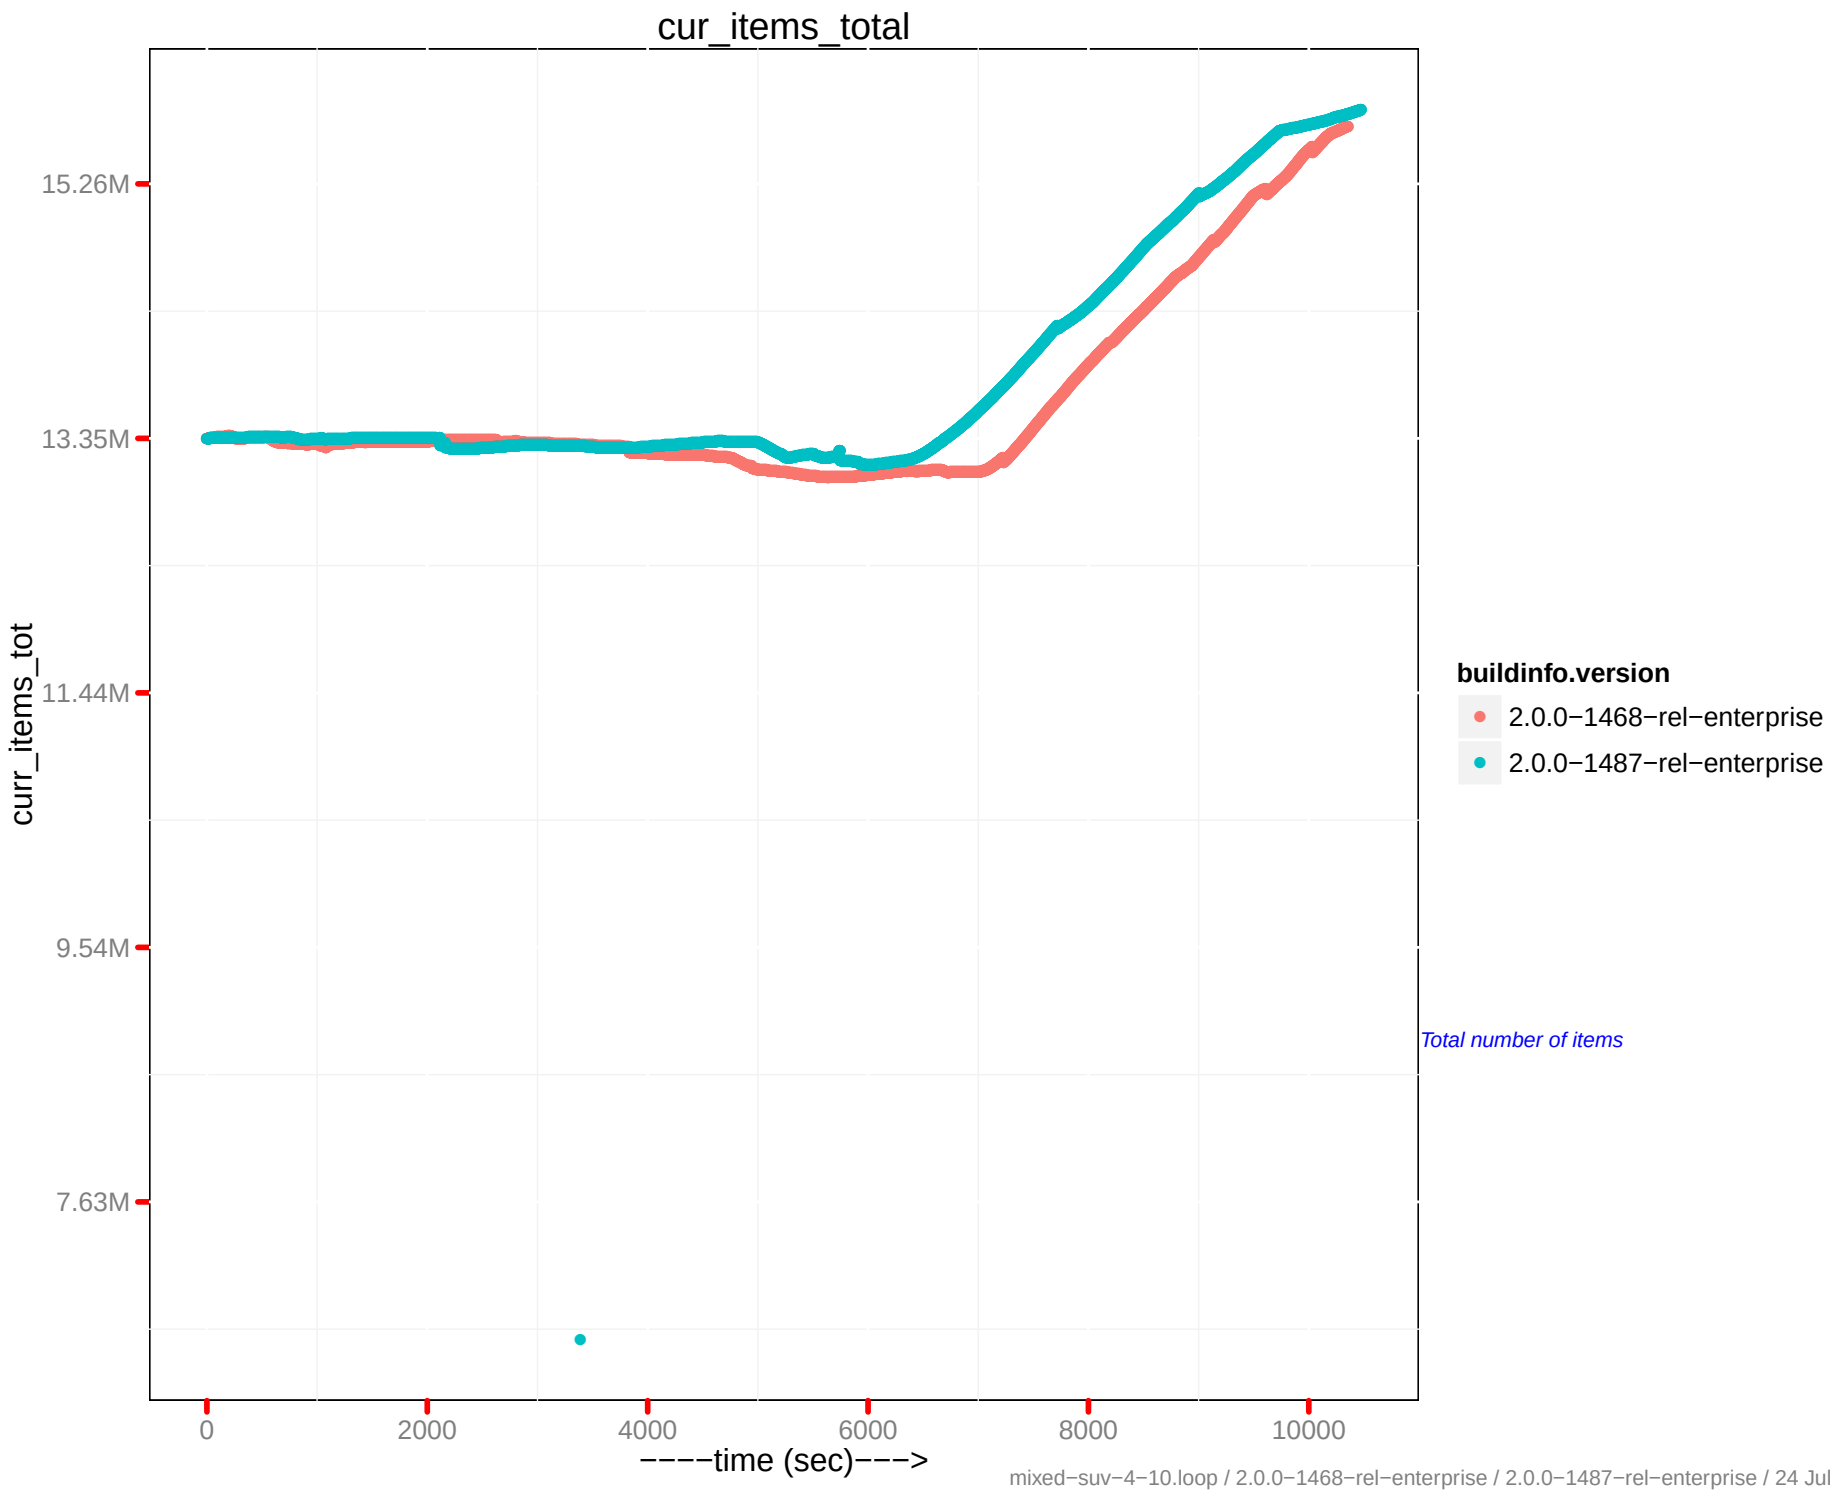
cur_items_total
15.26M
13.35M
11.44M
9.54M
7.63M
curr_items_tot
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Total number of items
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

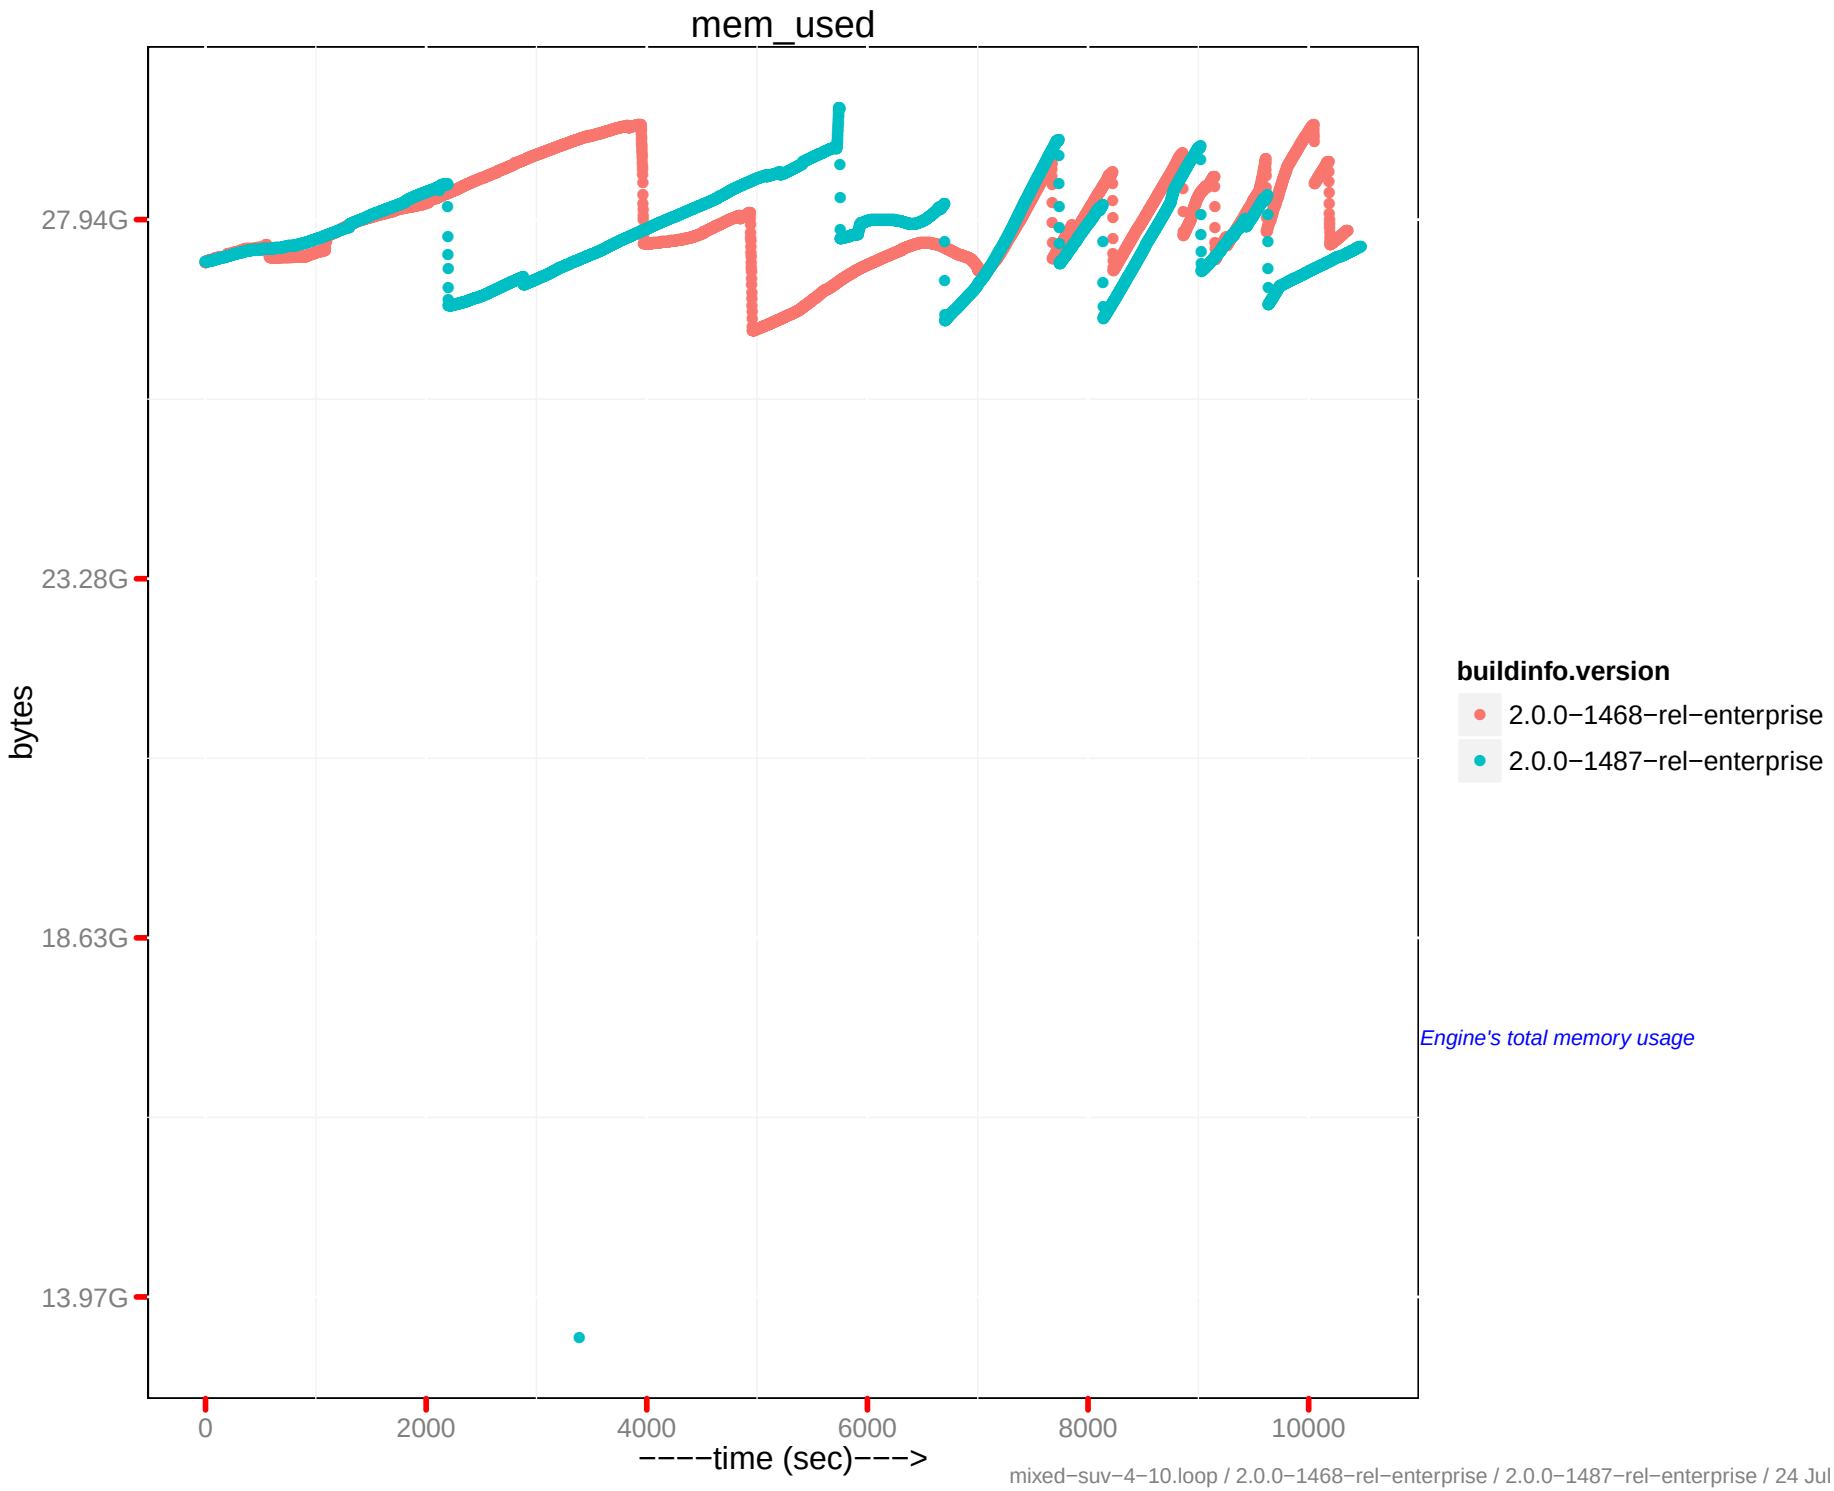
mem_used
27.94G
23.28G
18.63G
13.97G
bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Engine's total memory usage
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

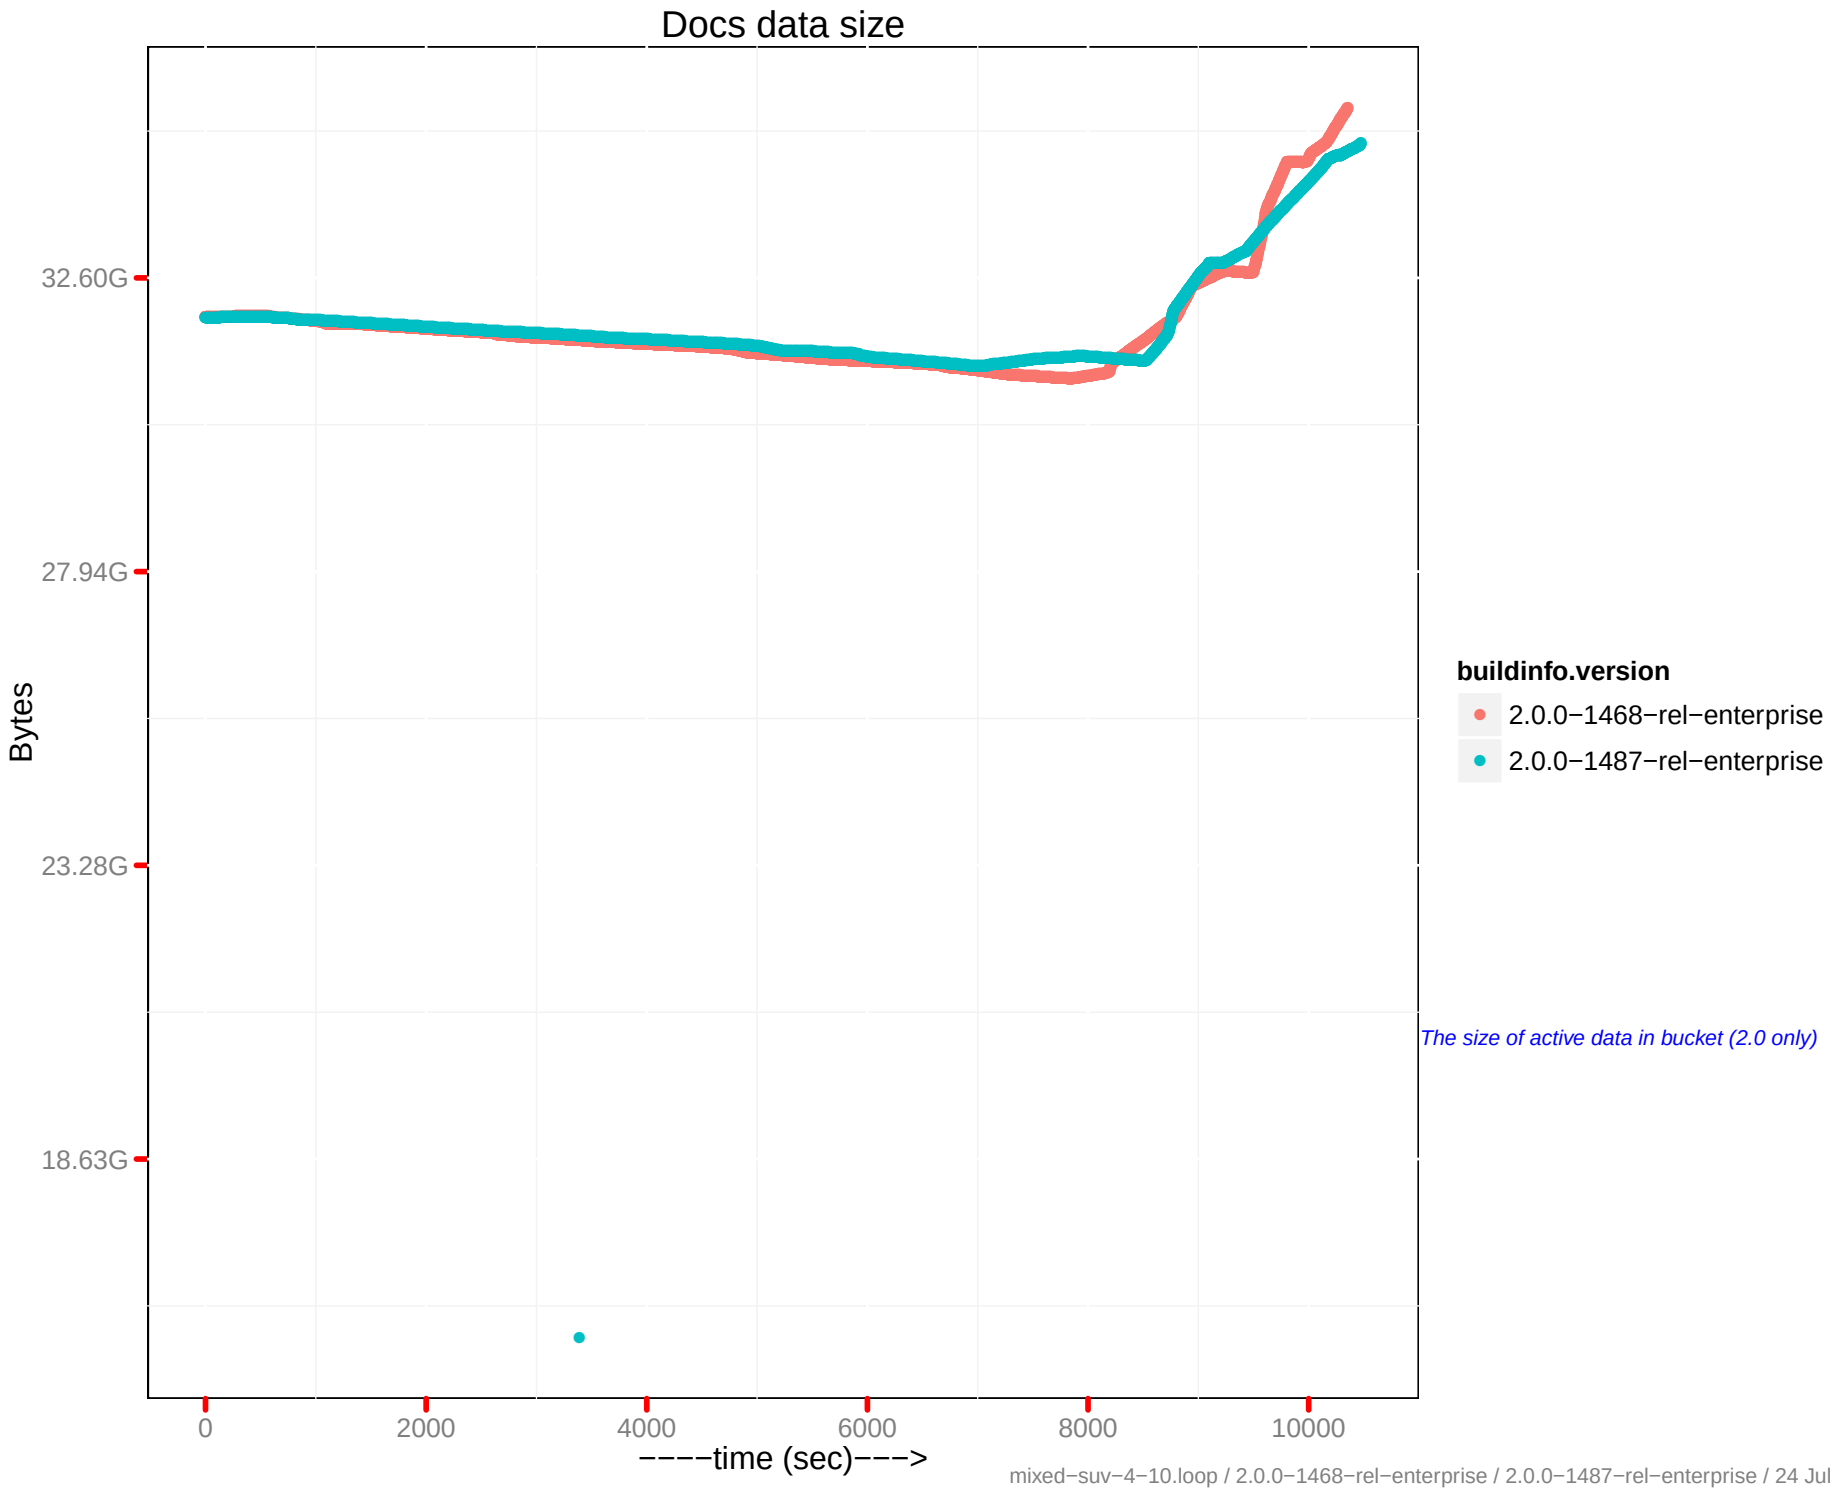
Docs data size
32.60G
27.94G
23.28G
18.63G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
The size of active data in bucket (2.0 only)
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

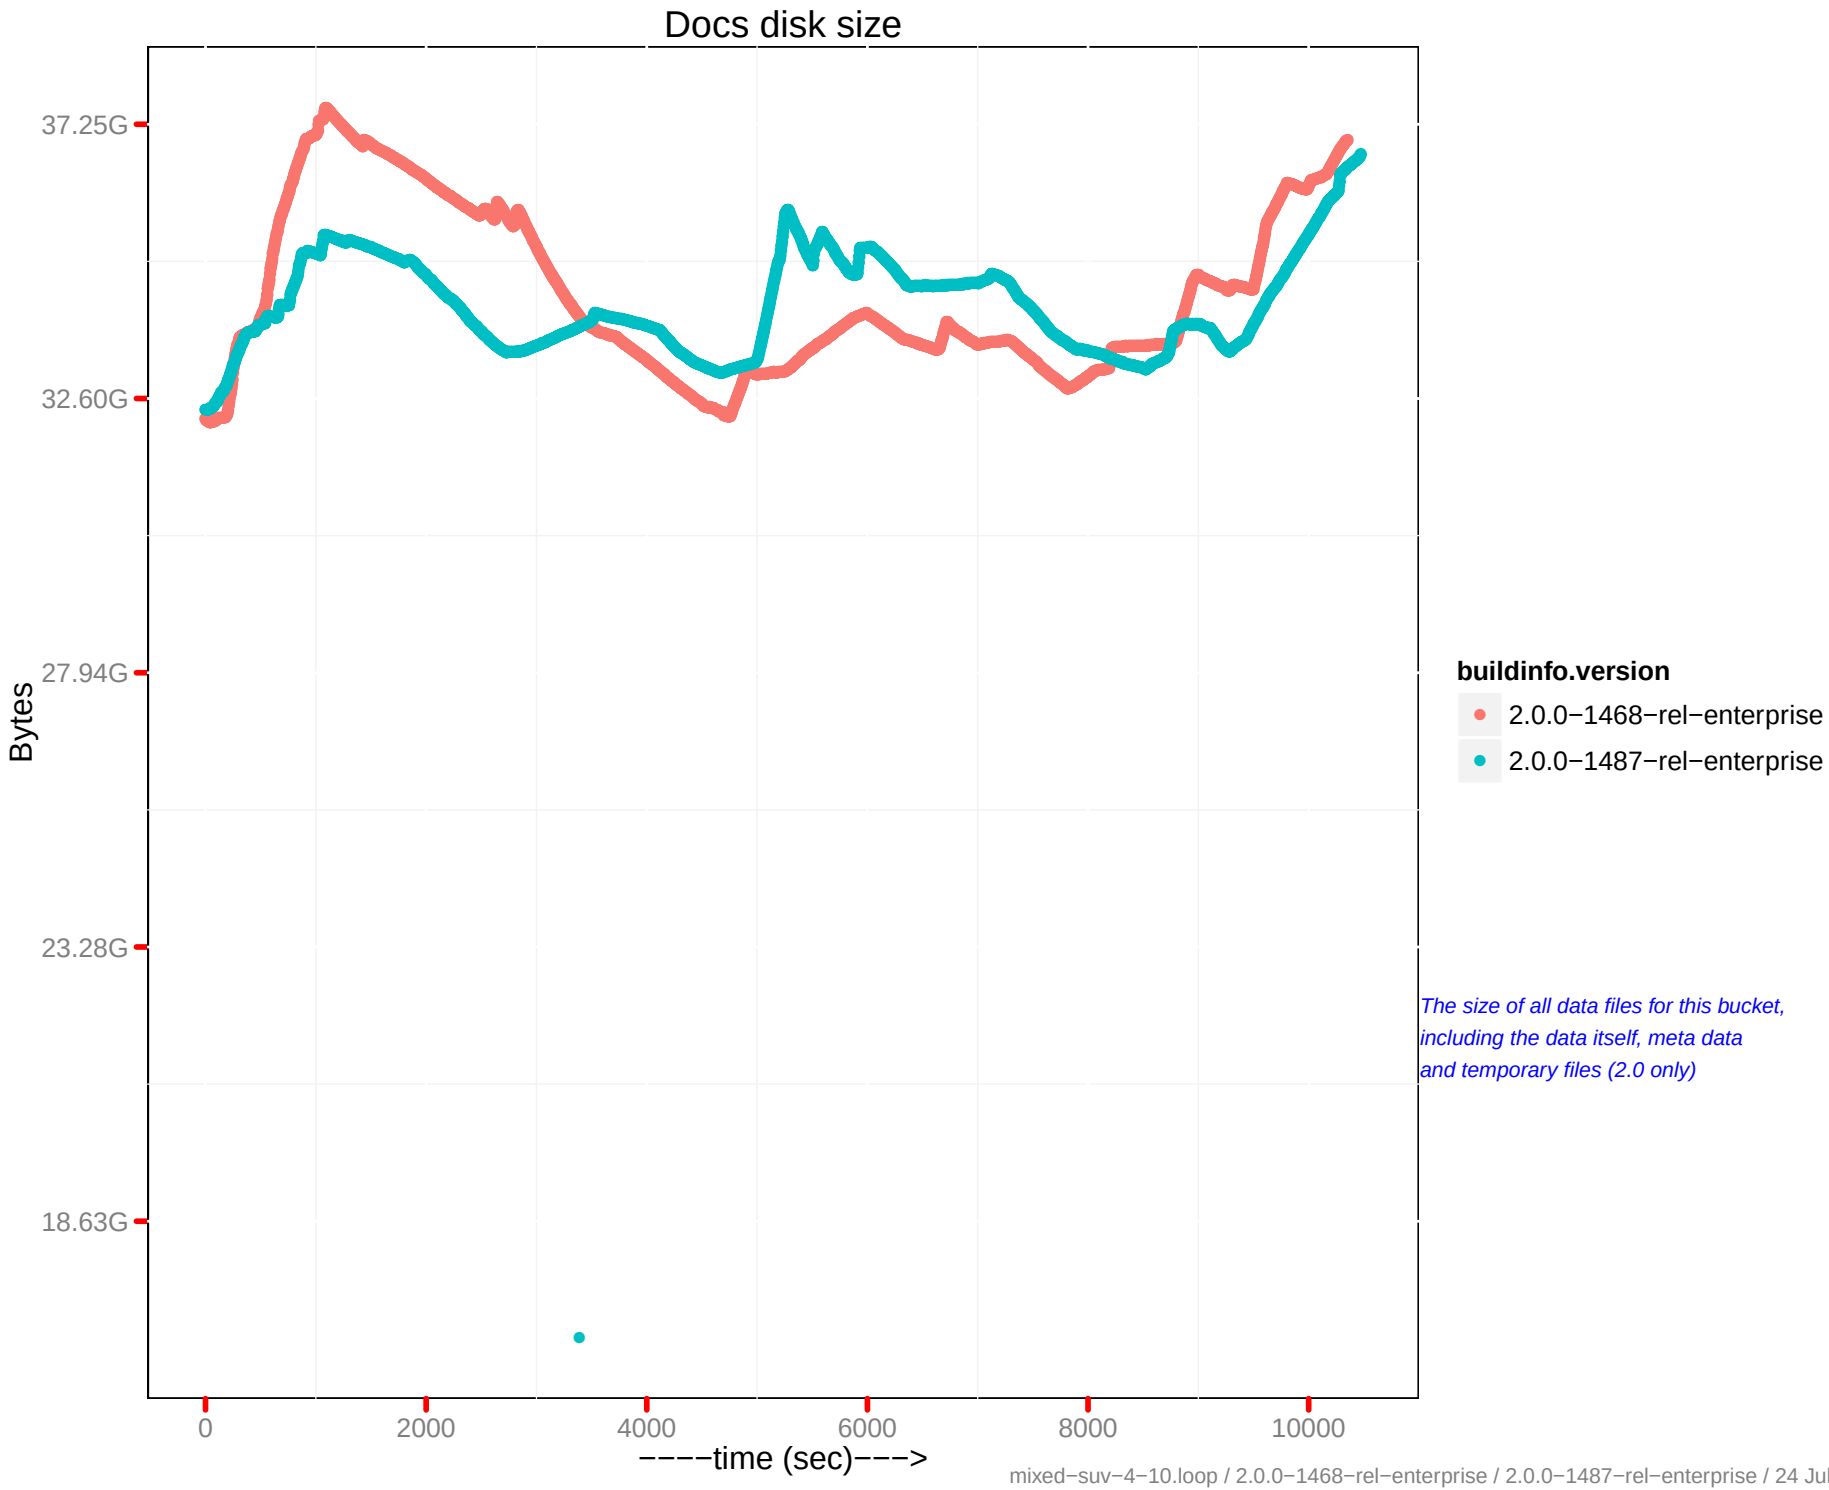
Docs disk size
37.25G
32.60G
27.94G
23.28G
18.63G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

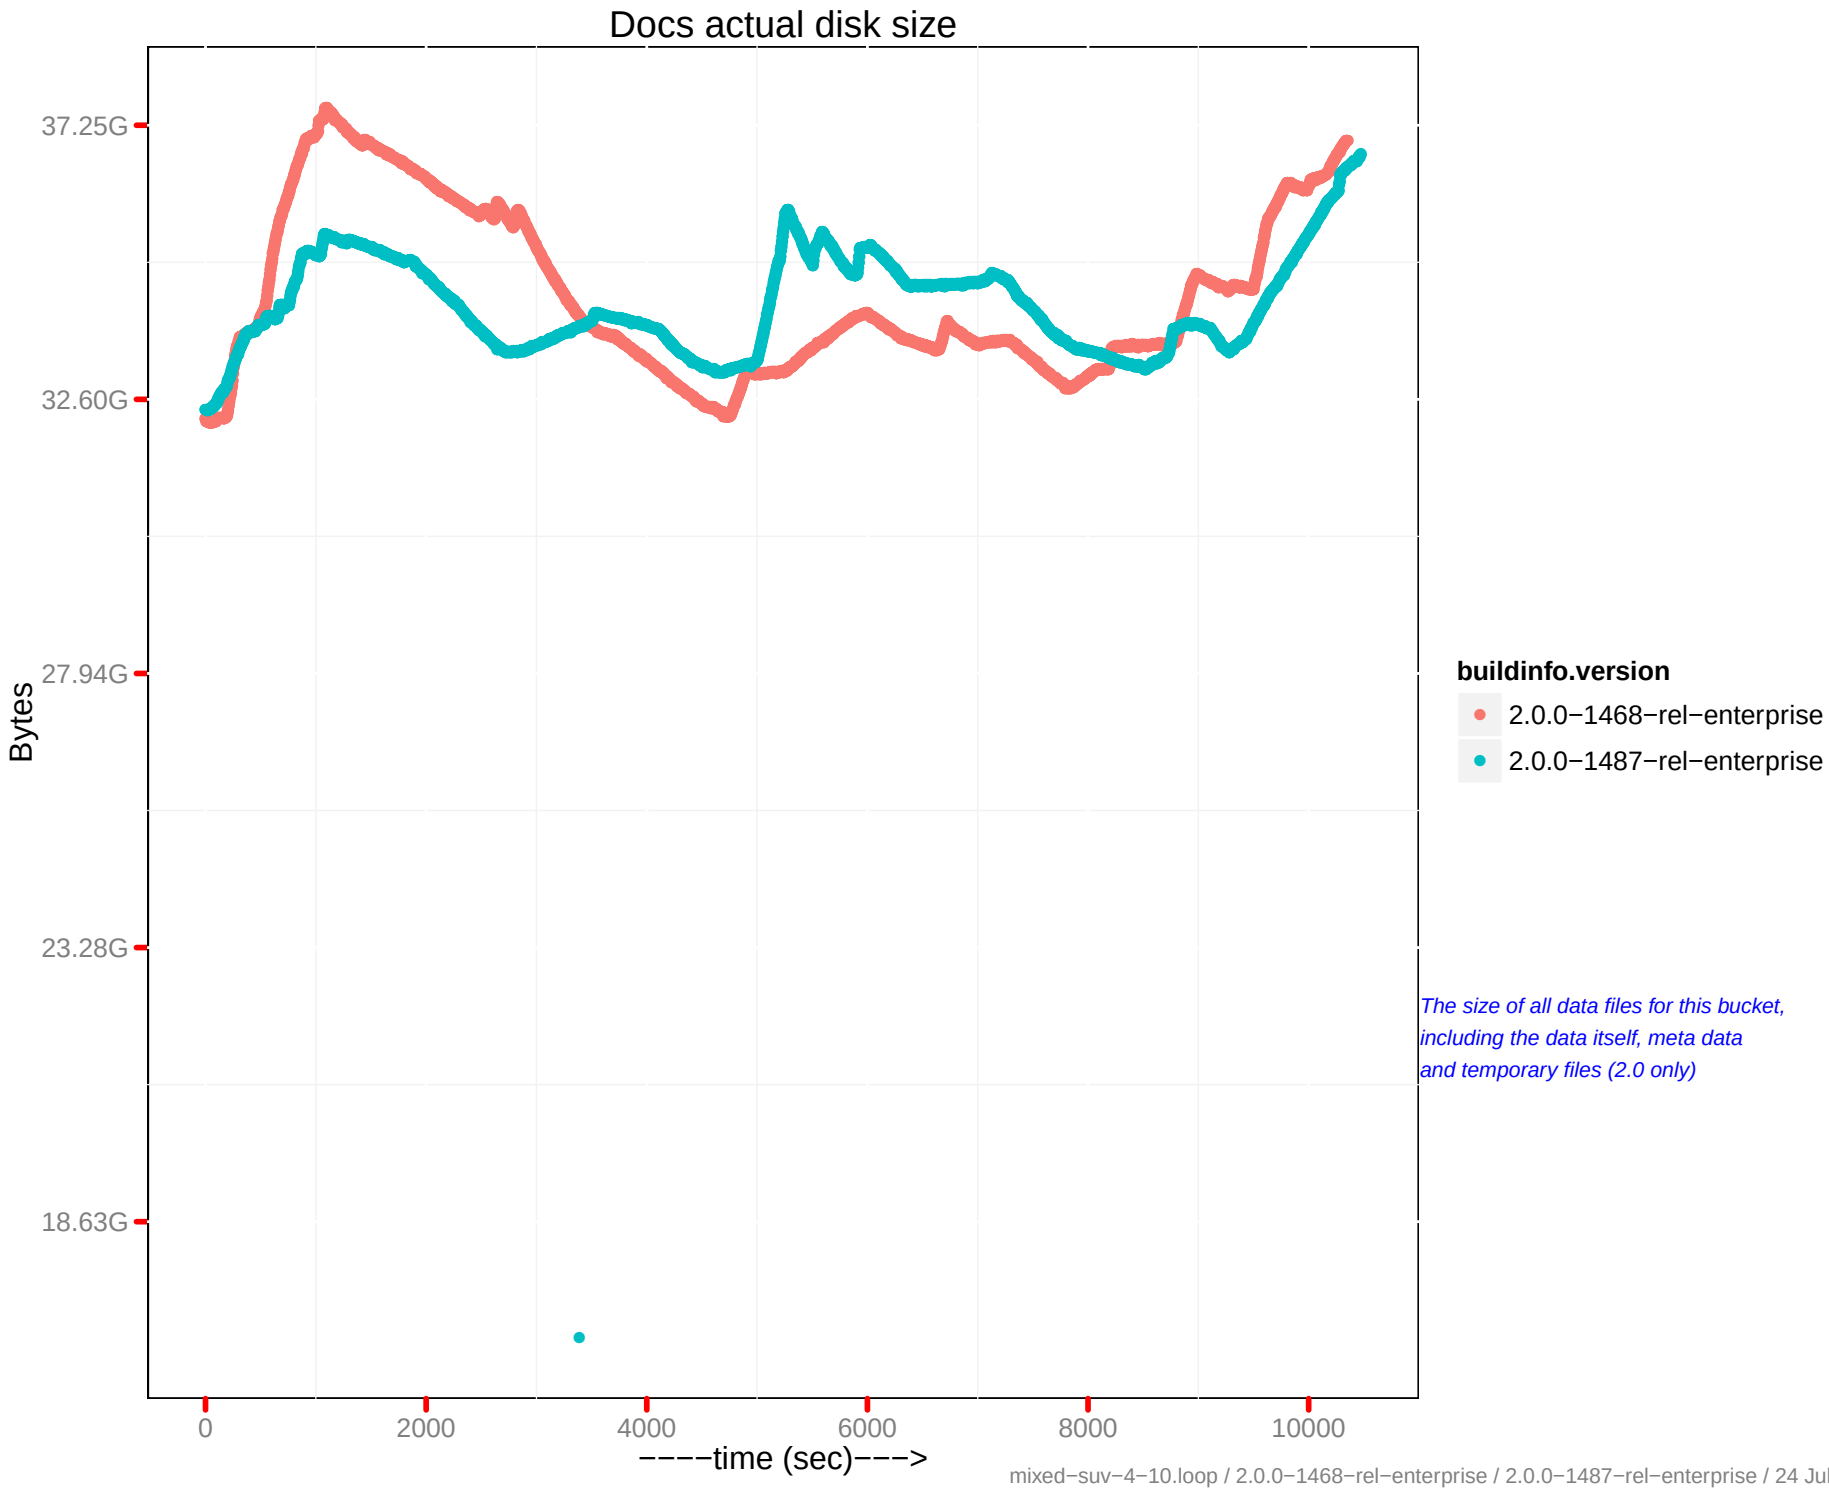
Docs actual disk size
37.25G
32.60G
27.94G
23.28G
18.63G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

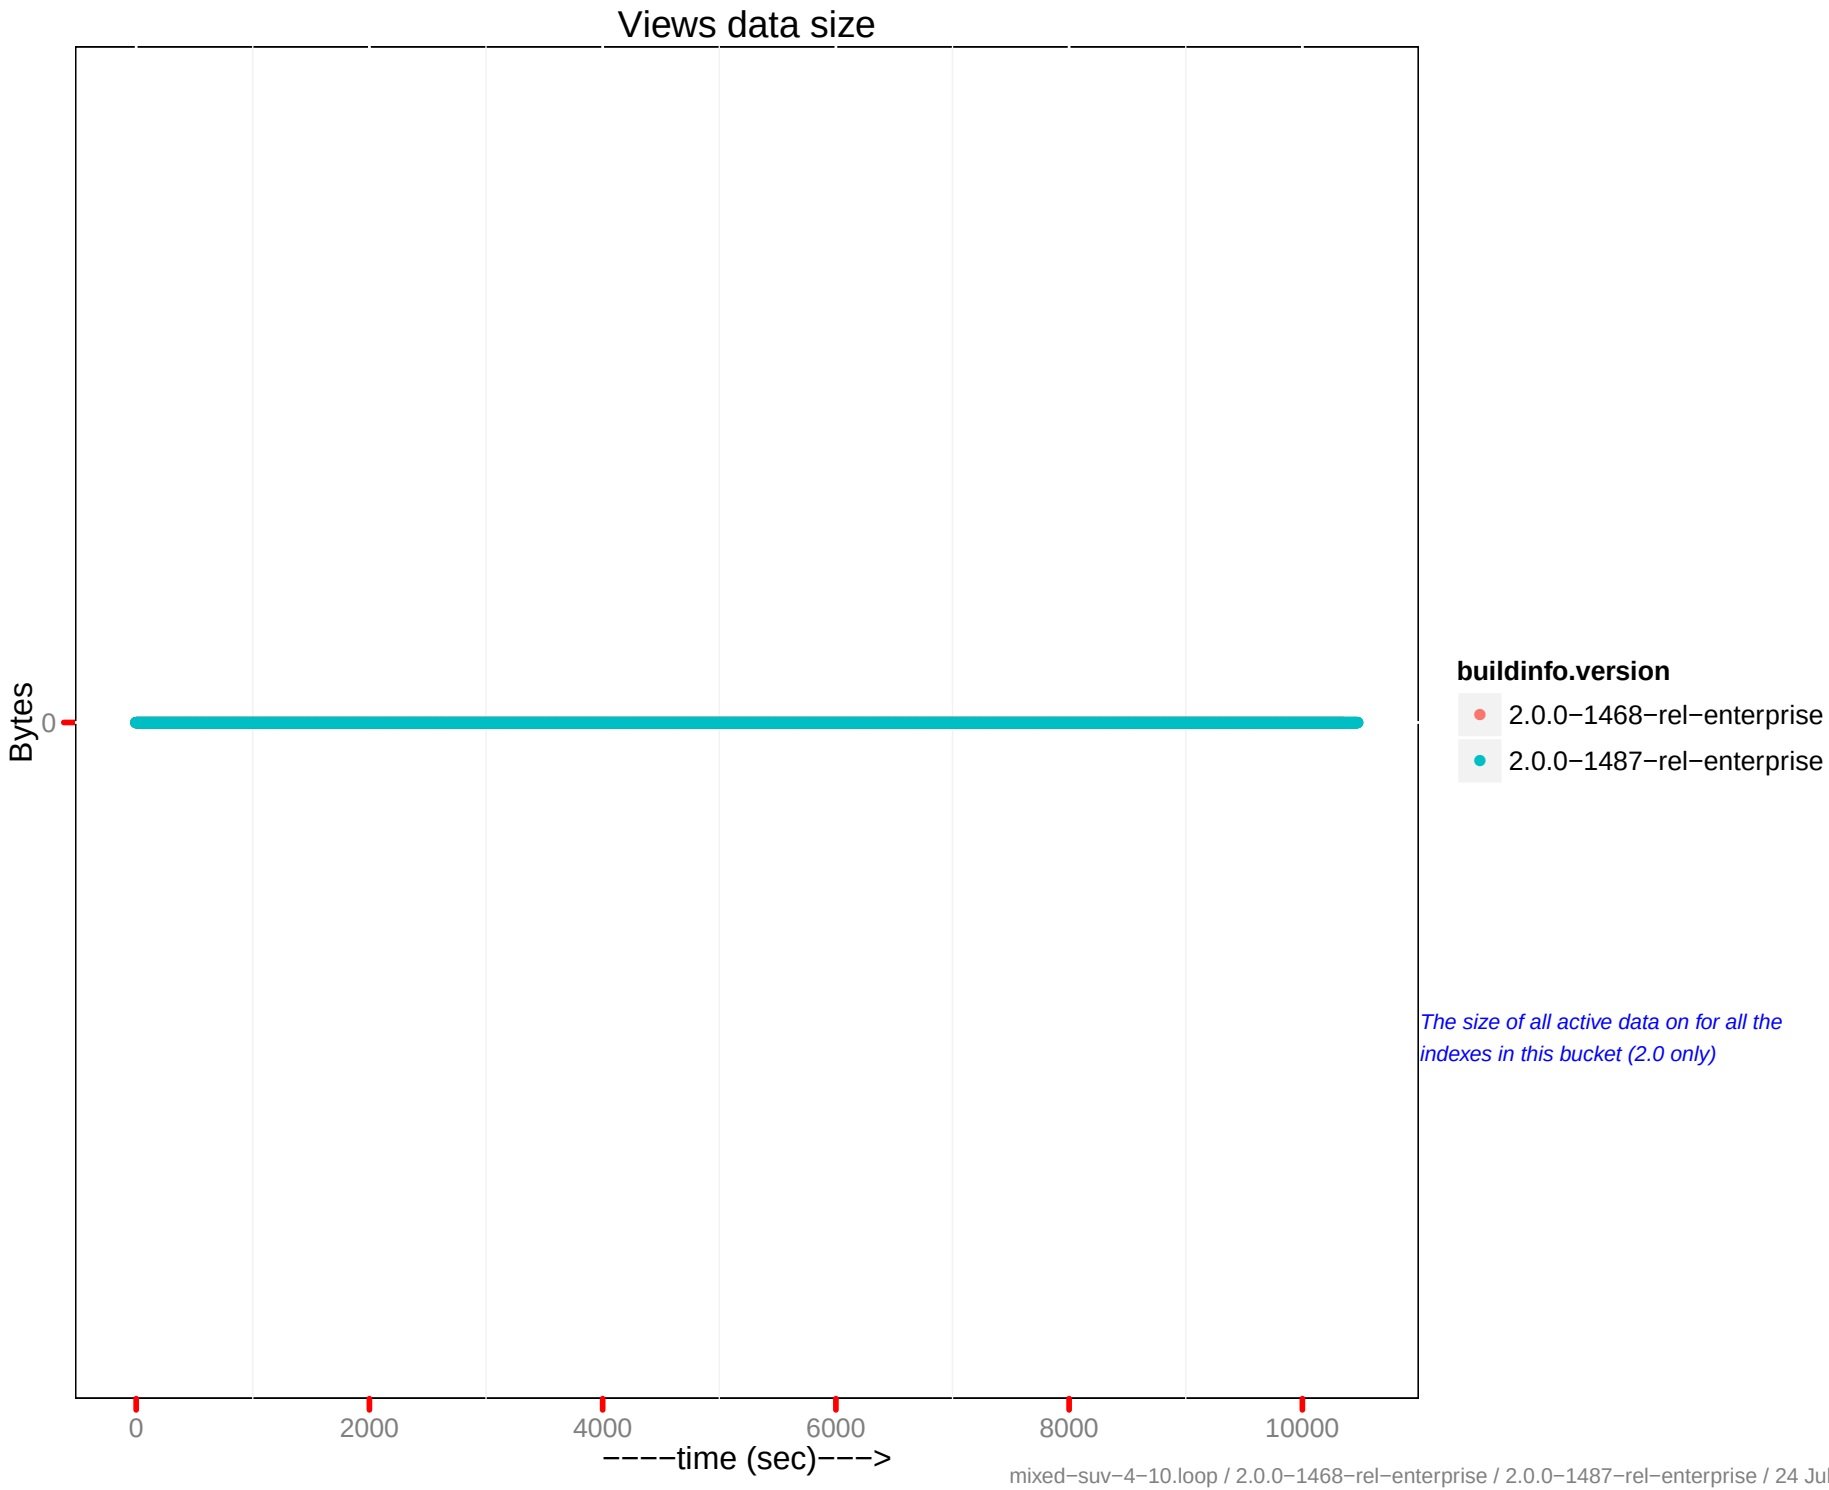
Views data size
Bytes
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

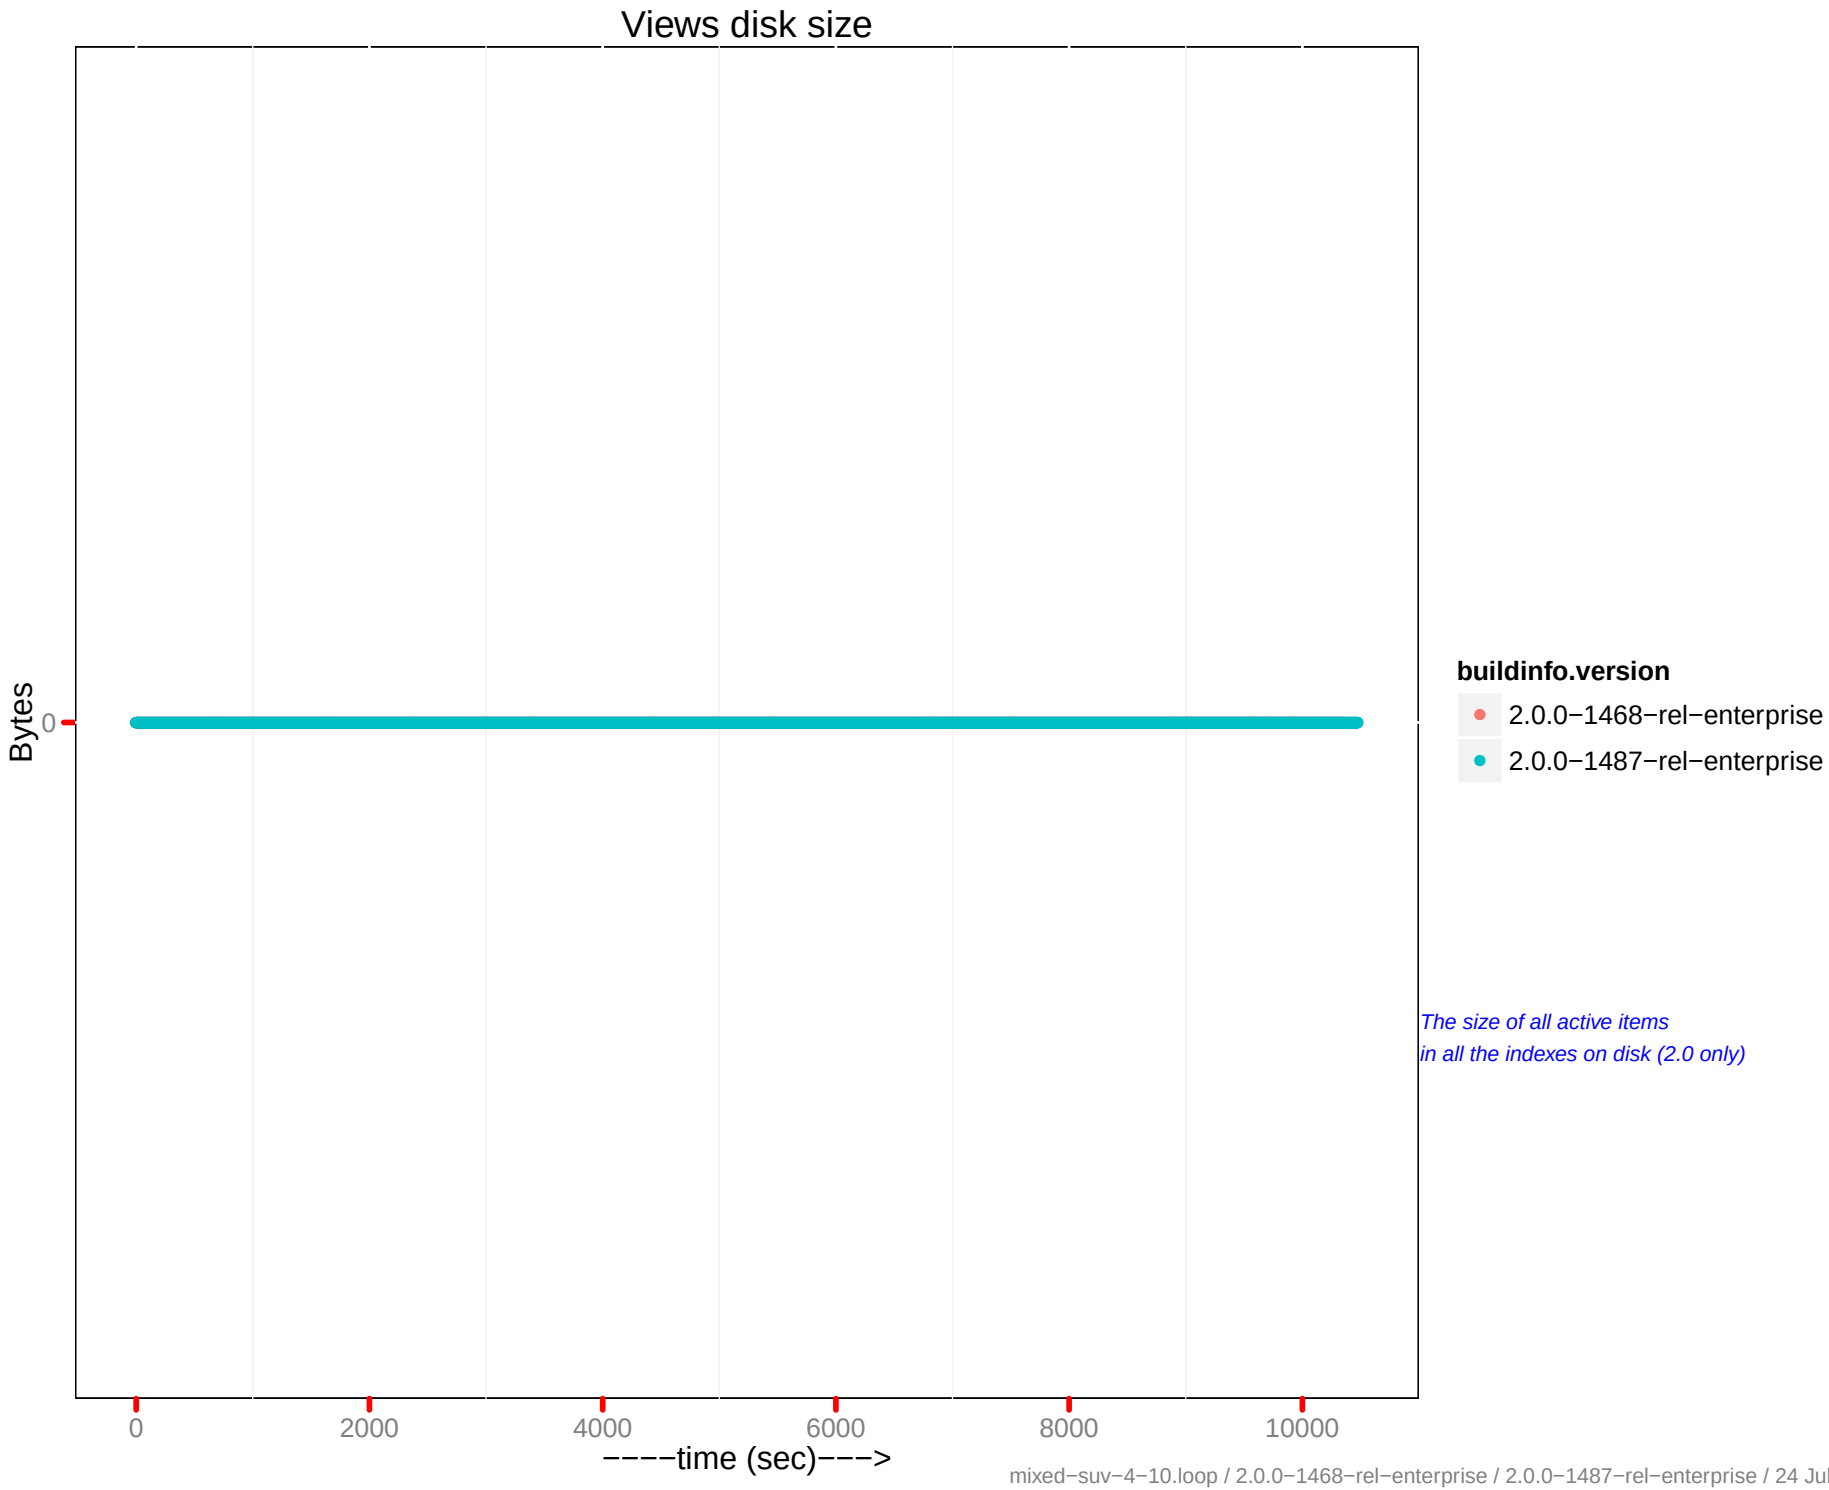
Views disk size
Bytes
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
The size of all active items
in all the indexes on disk (2.0 only)
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

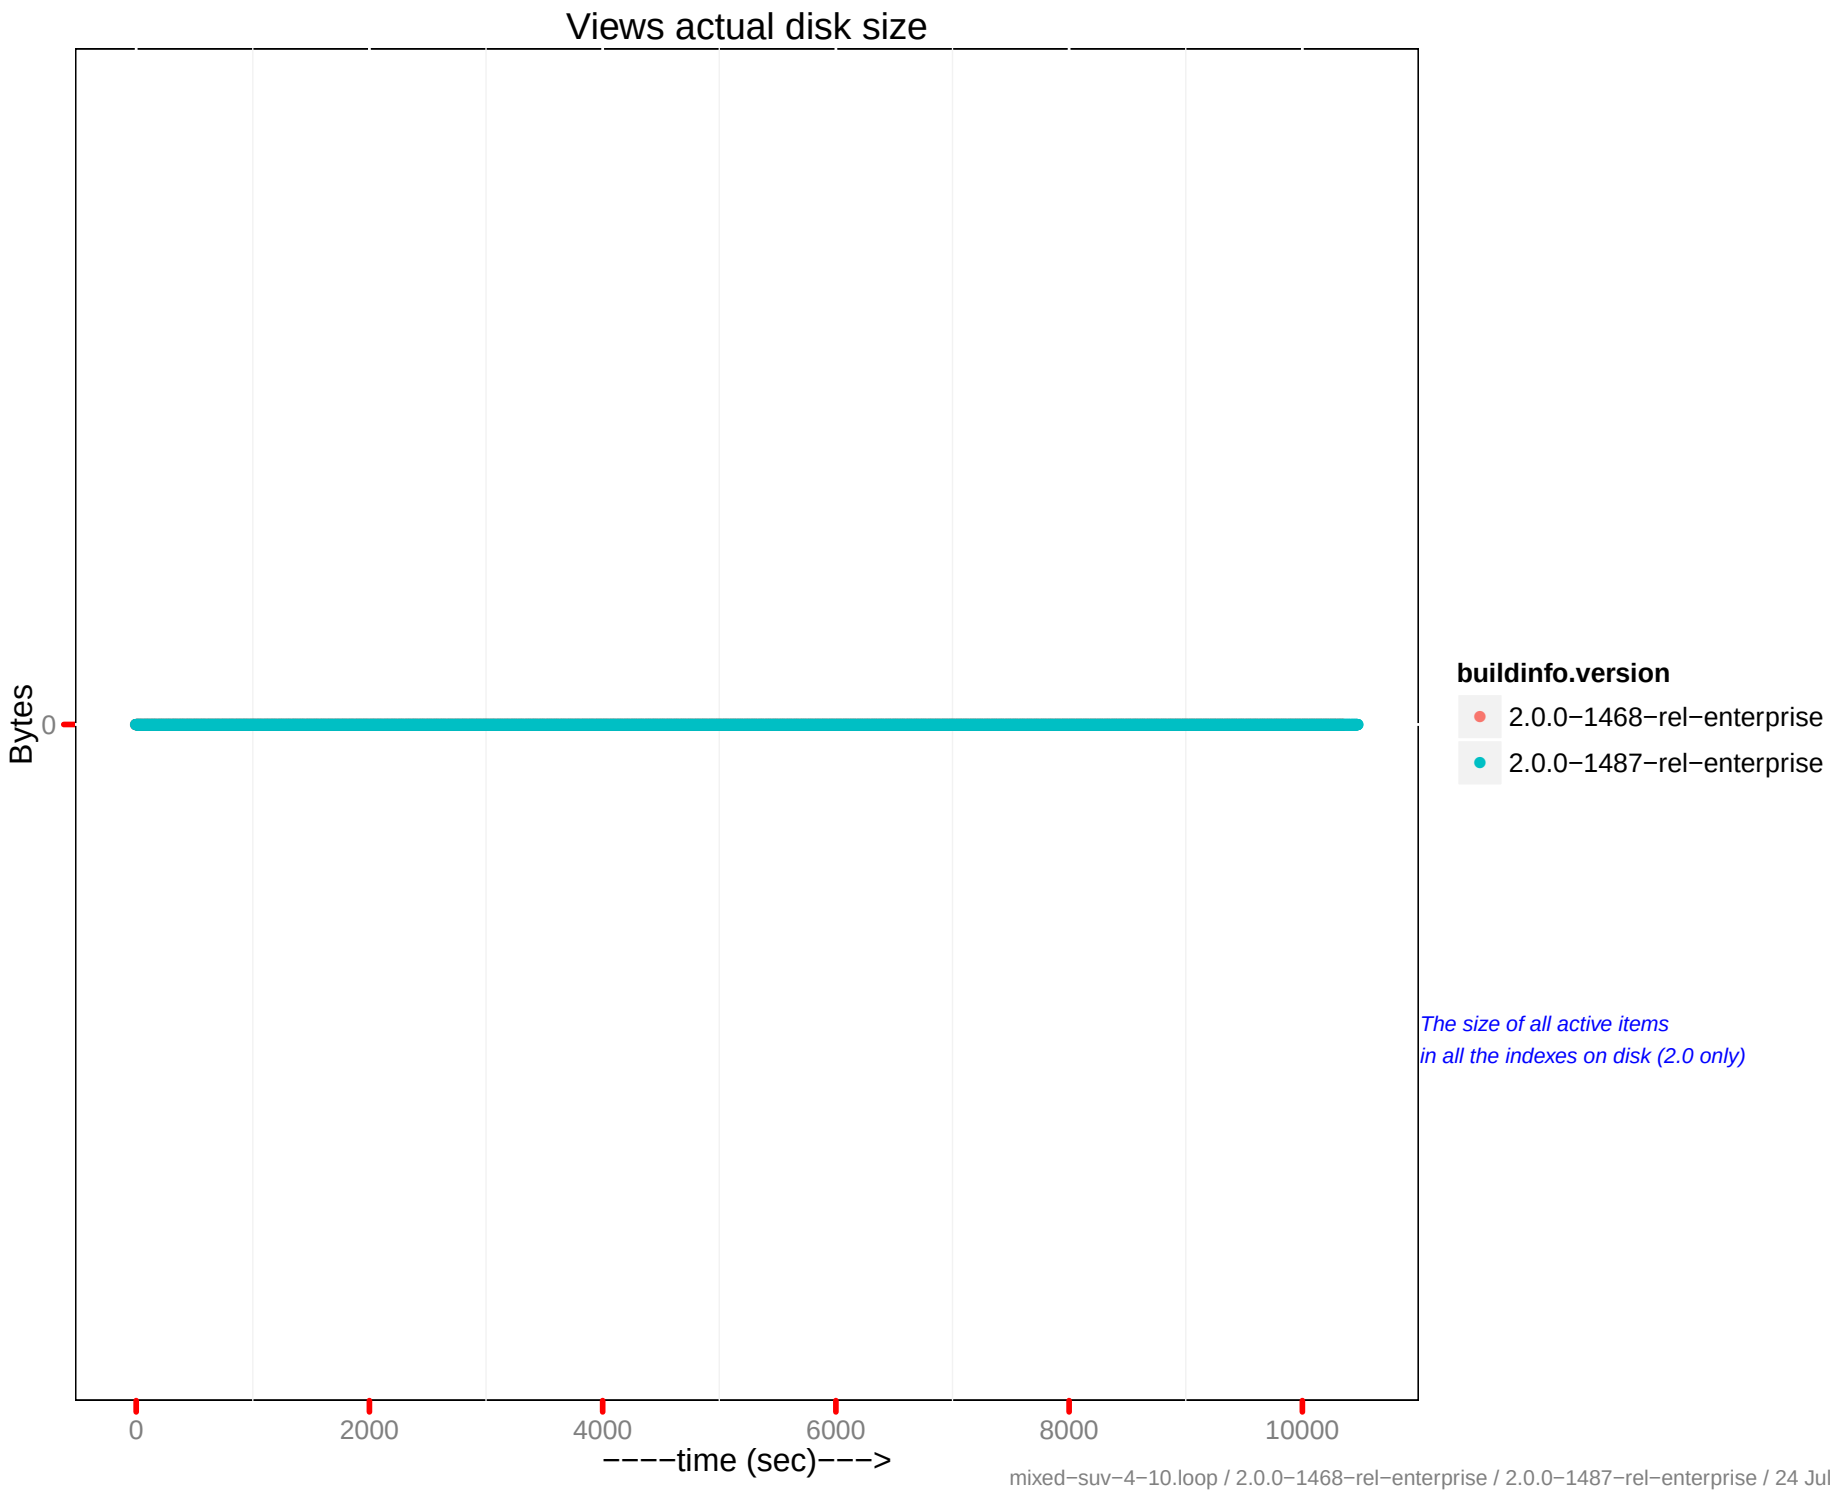
Views actual disk size
Bytes
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
The size of all active items
in all the indexes on disk (2.0 only)
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

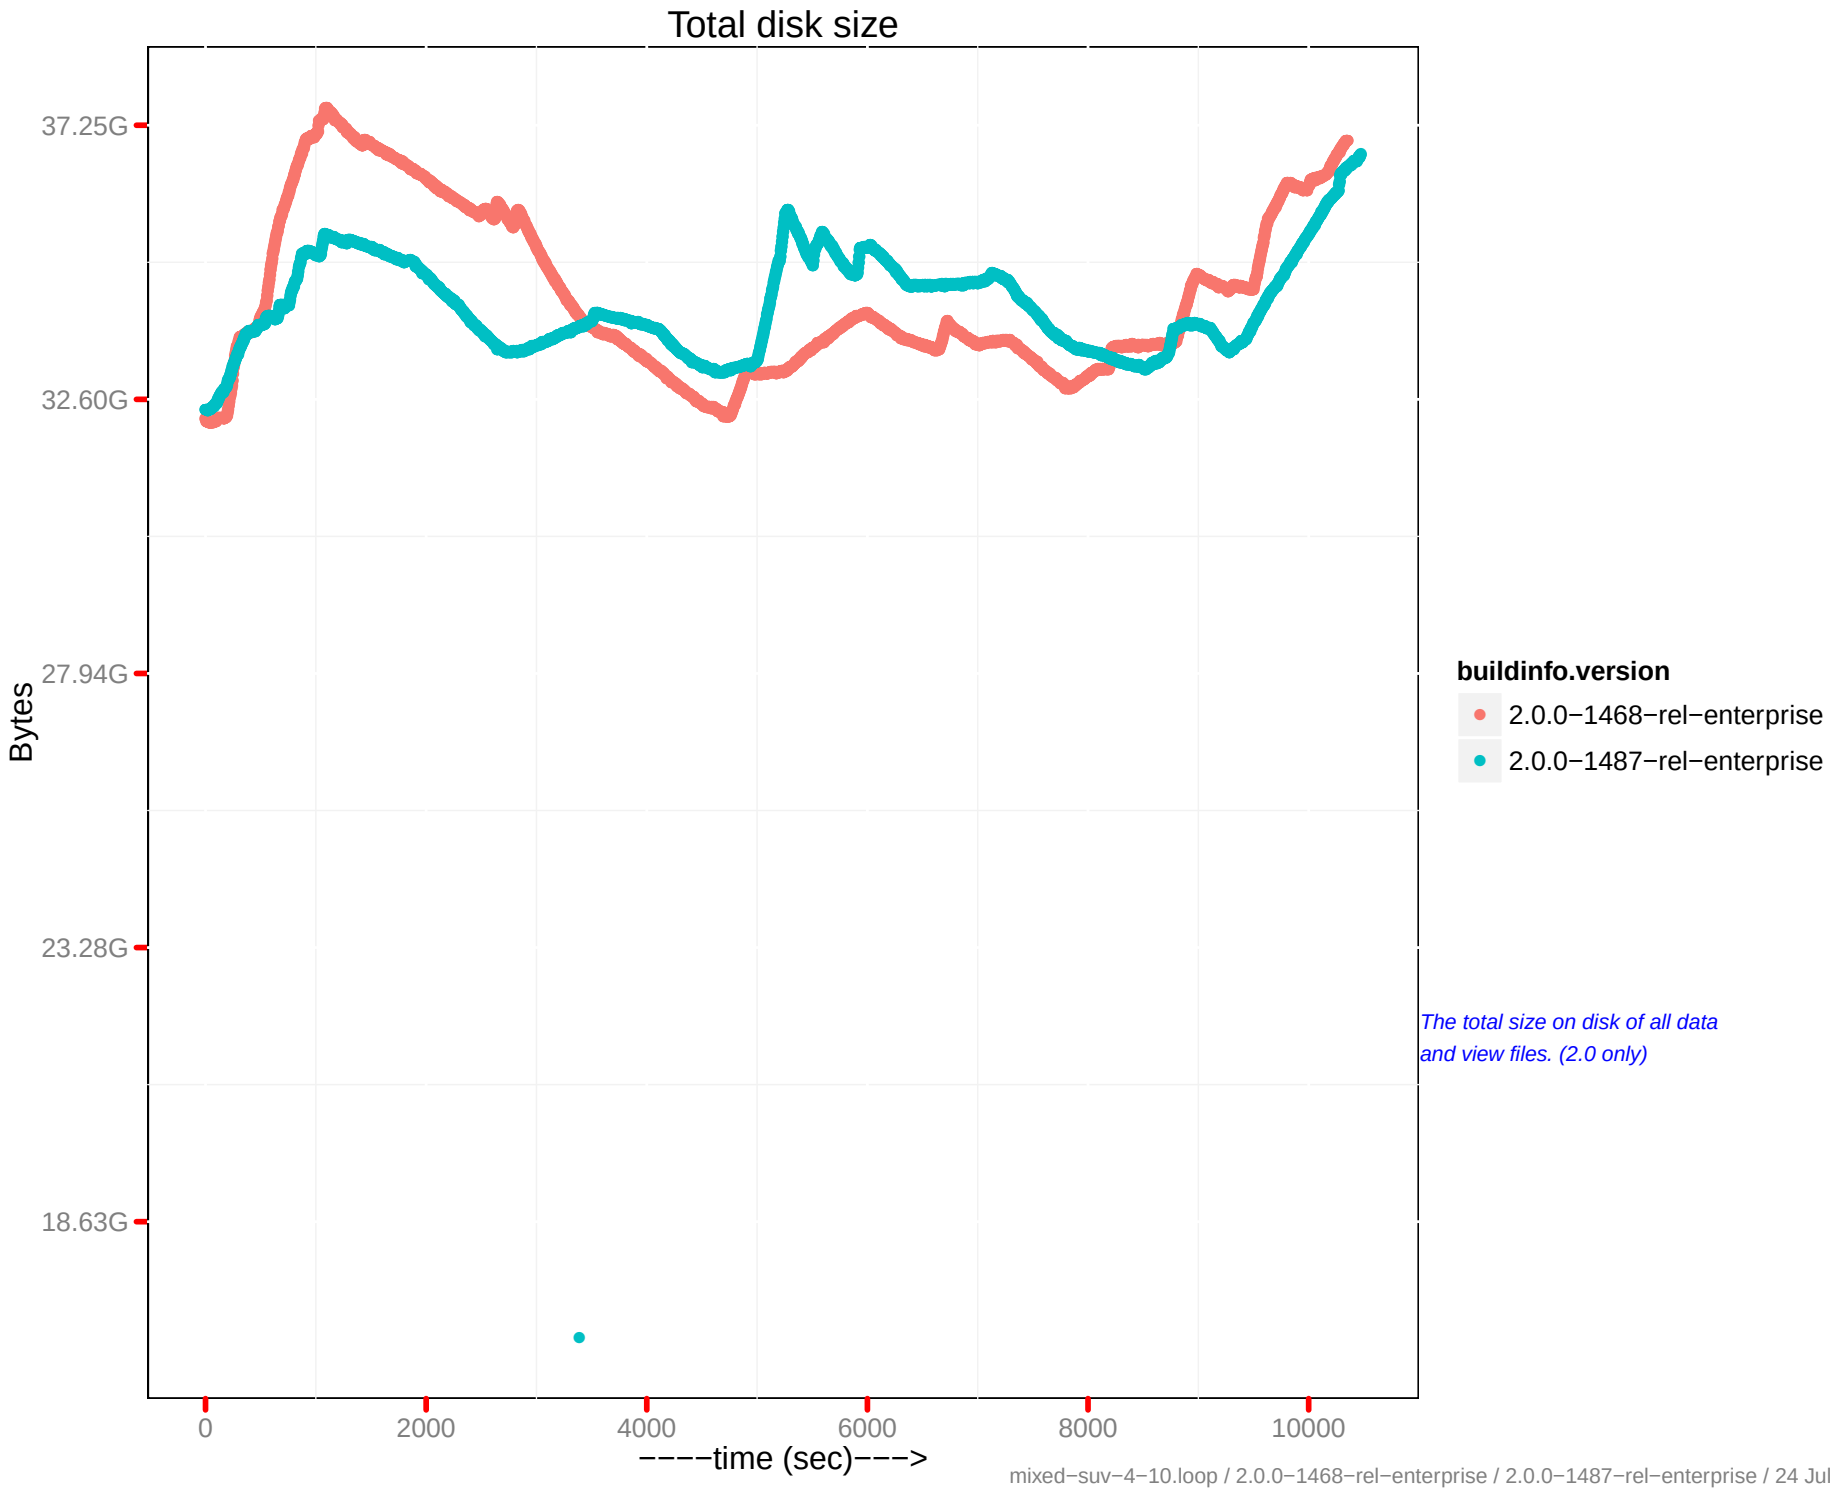
Total disk size
37.25G
32.60G
27.94G
23.28G
18.63G
Bytes
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

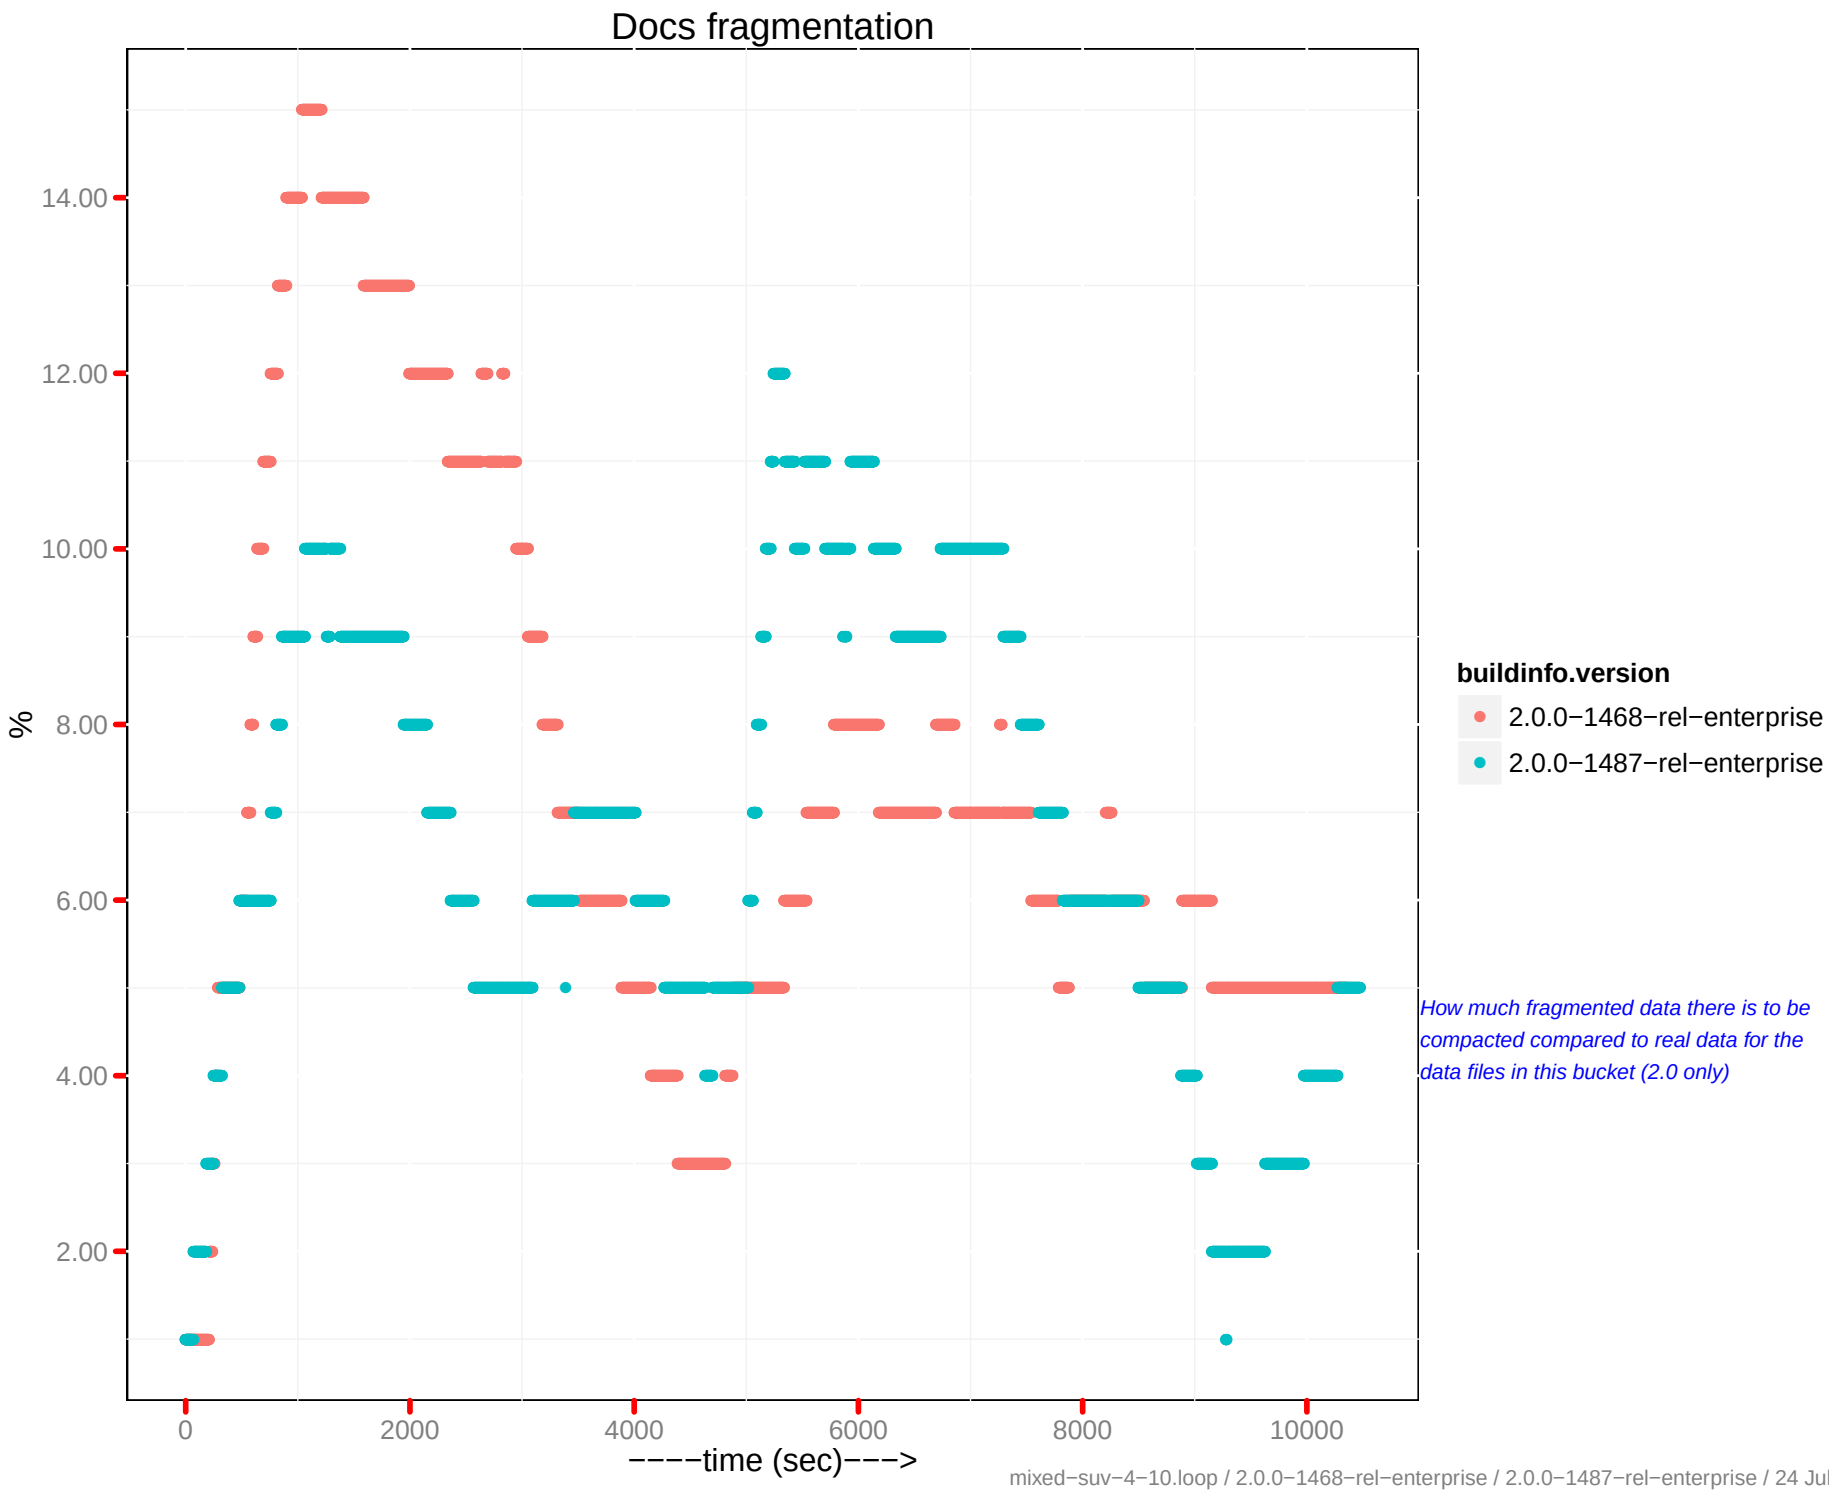
Docs fragmentation
14.00
12.00
10.00
8.00
6.00
4.00
2.00
%
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

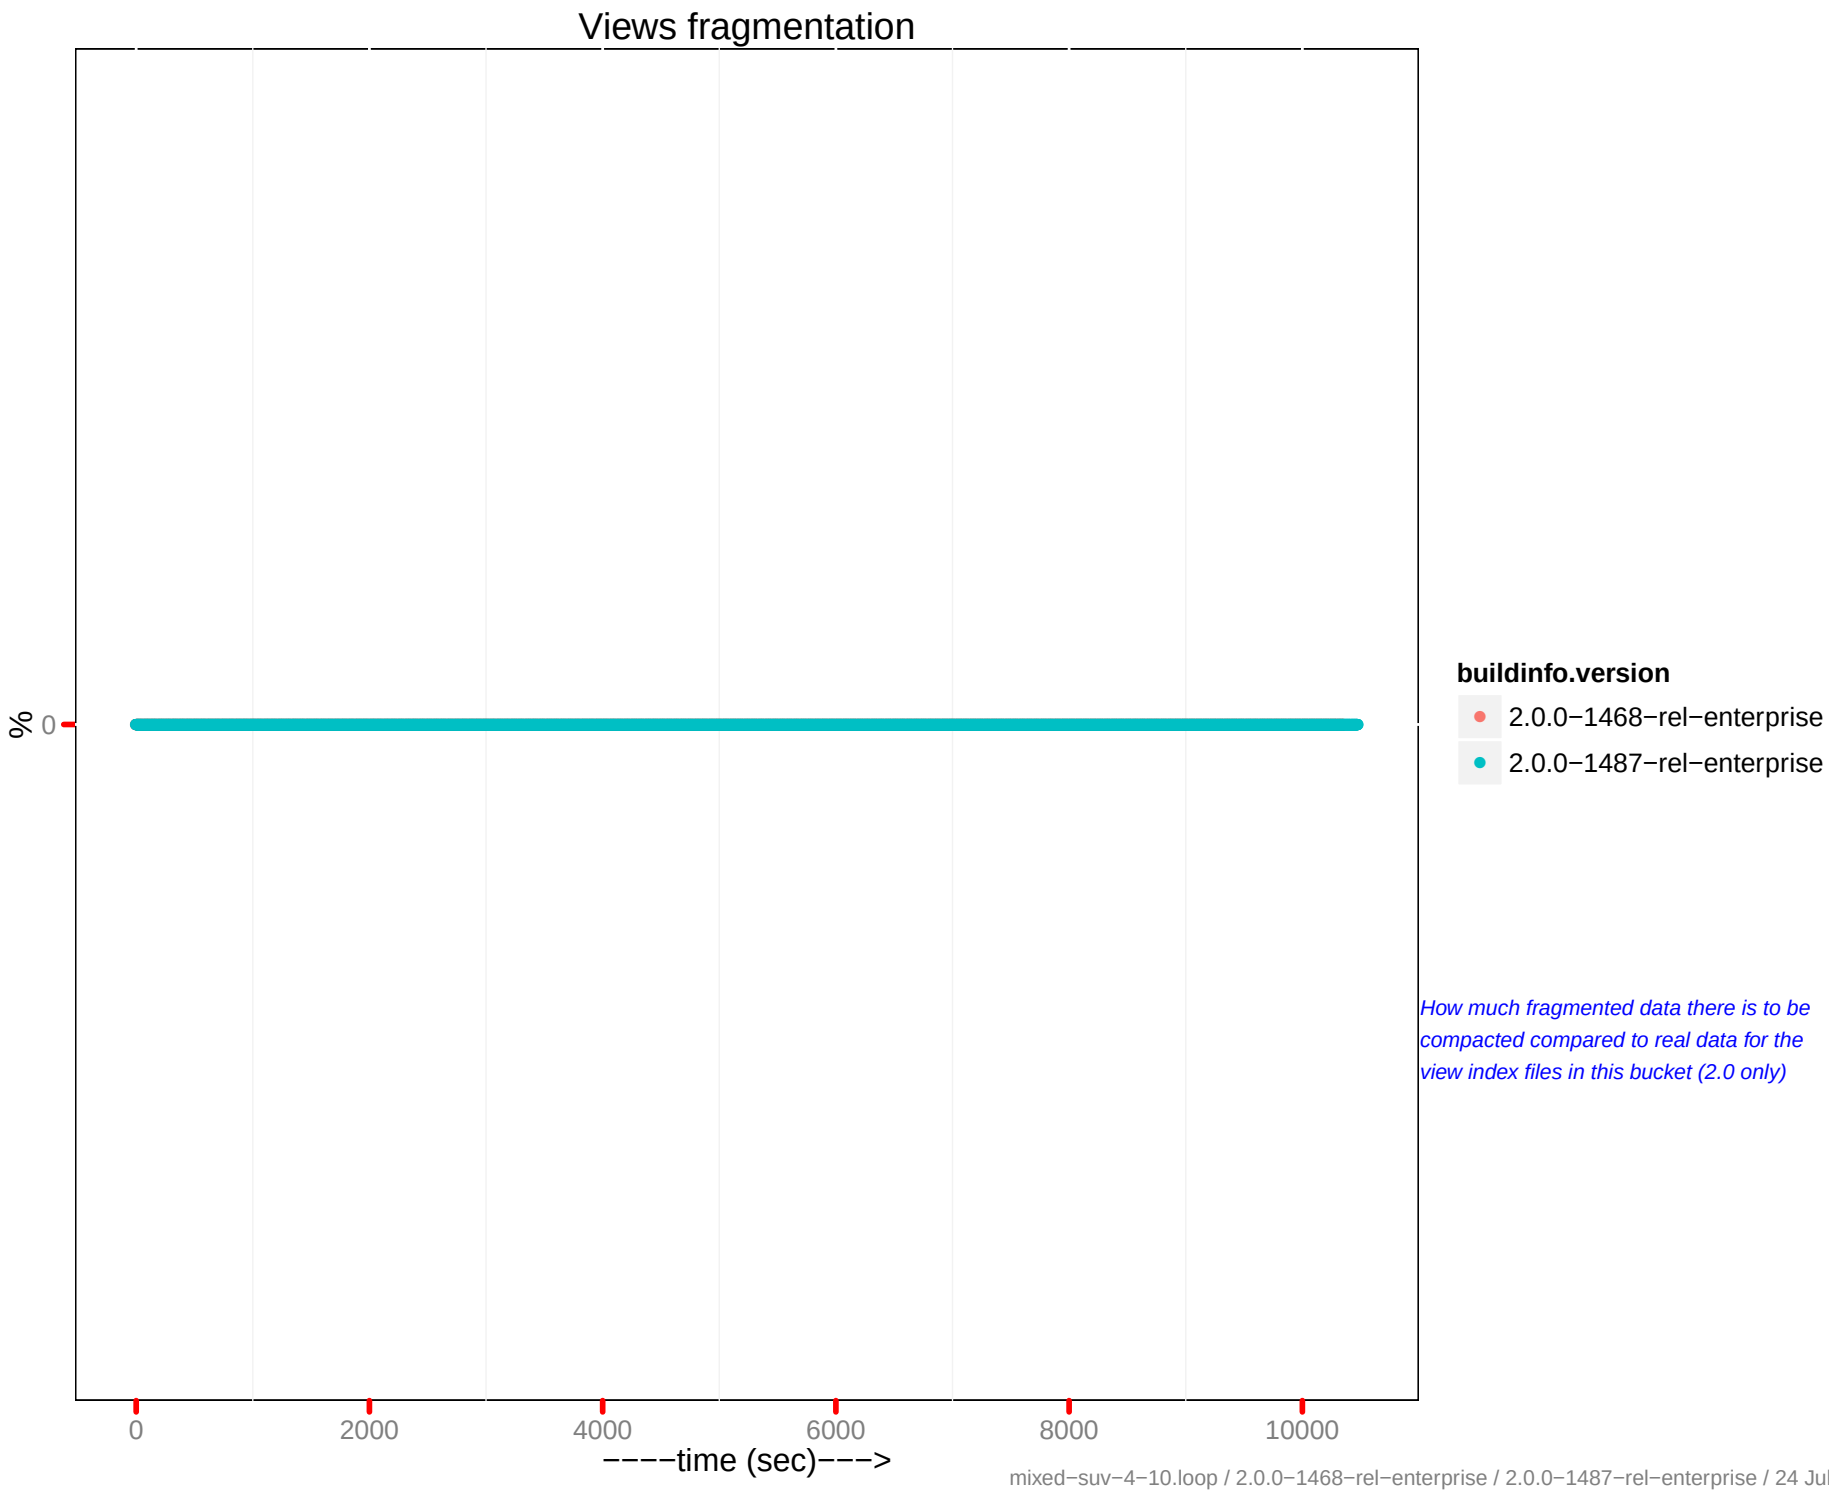
Views fragmentation
%
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

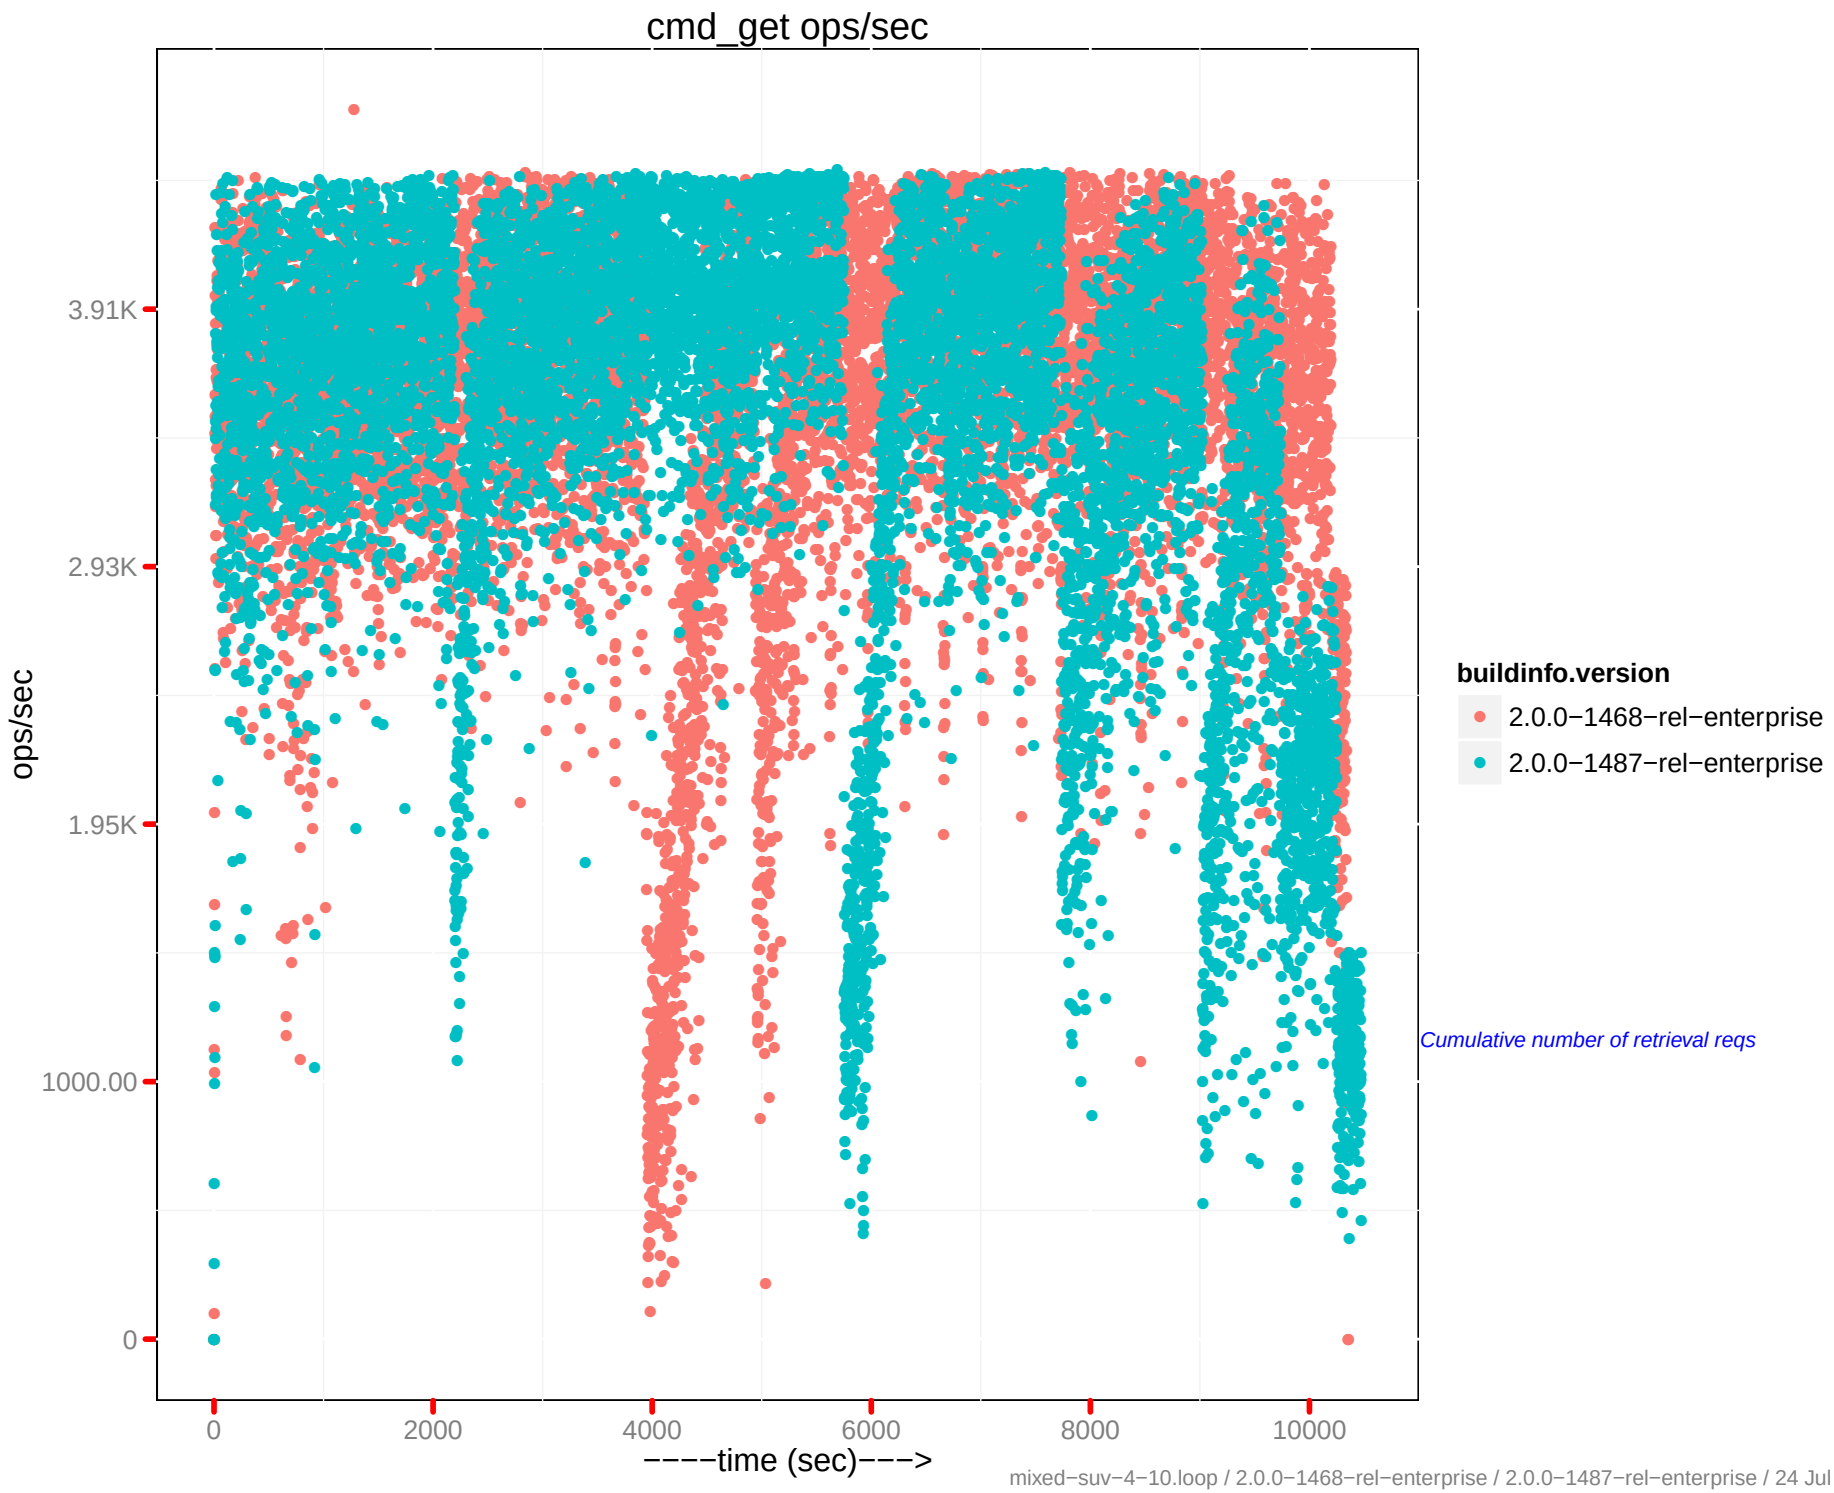
cmd_get ops/sec
ops/sec
3.91K
2.93K
1.95K
1000.00
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
Cumulative number of retrieval reqs
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

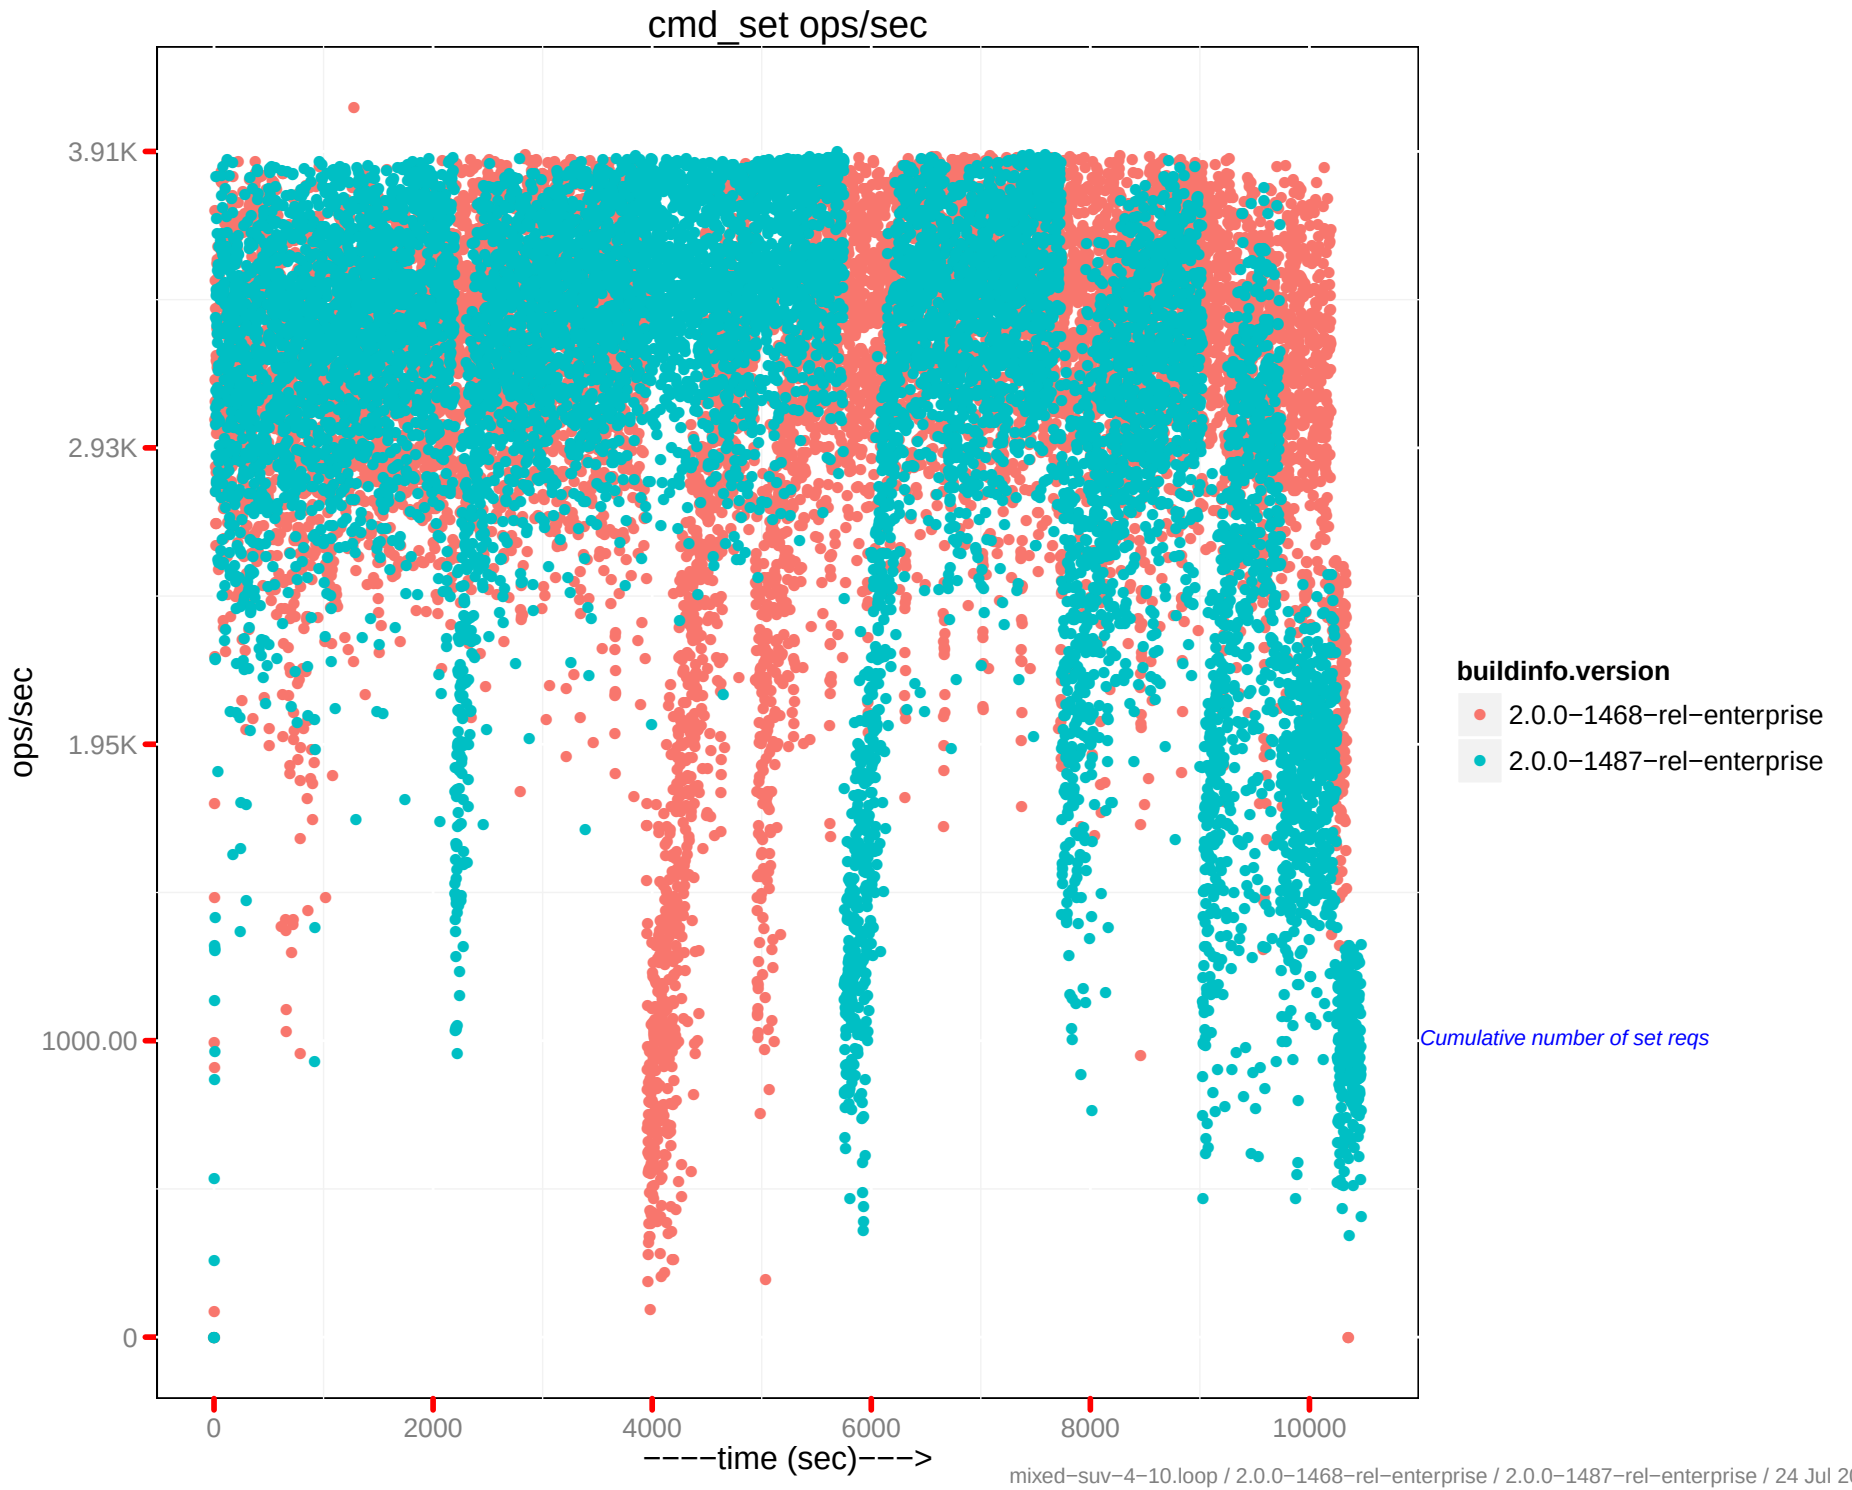
cmd_set ops/sec
3.91K
2.93K
1.95K
1000.00
0
ops/sec
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Cumulative number of set reqs
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

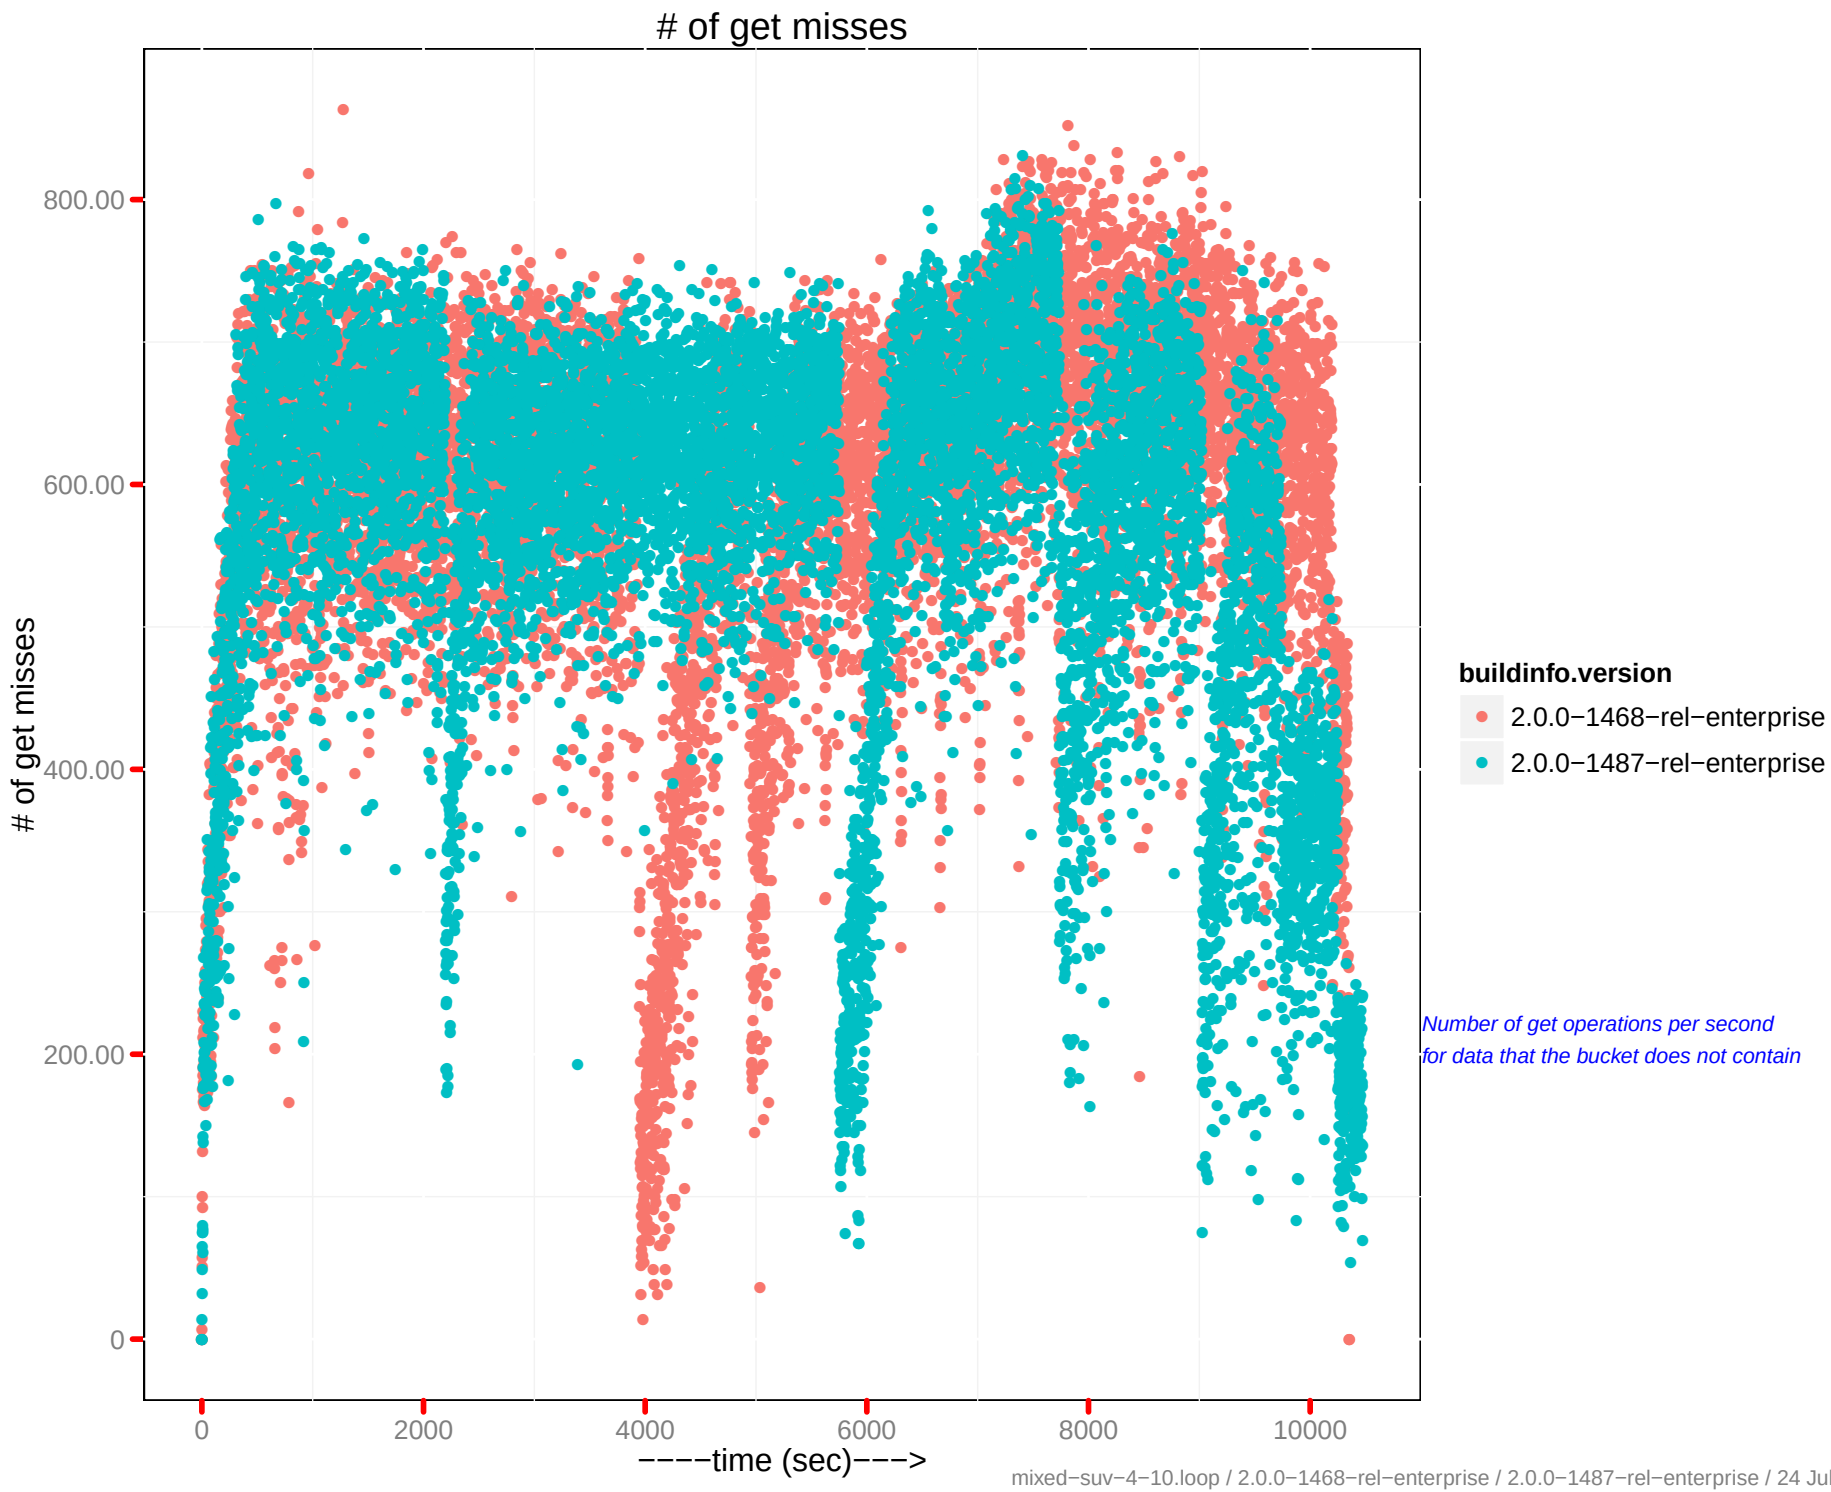
# of get misses
800.00
600.00
400.00
200.00
0
# of get misses
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of get operations per second
for data that the bucket does not contain
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

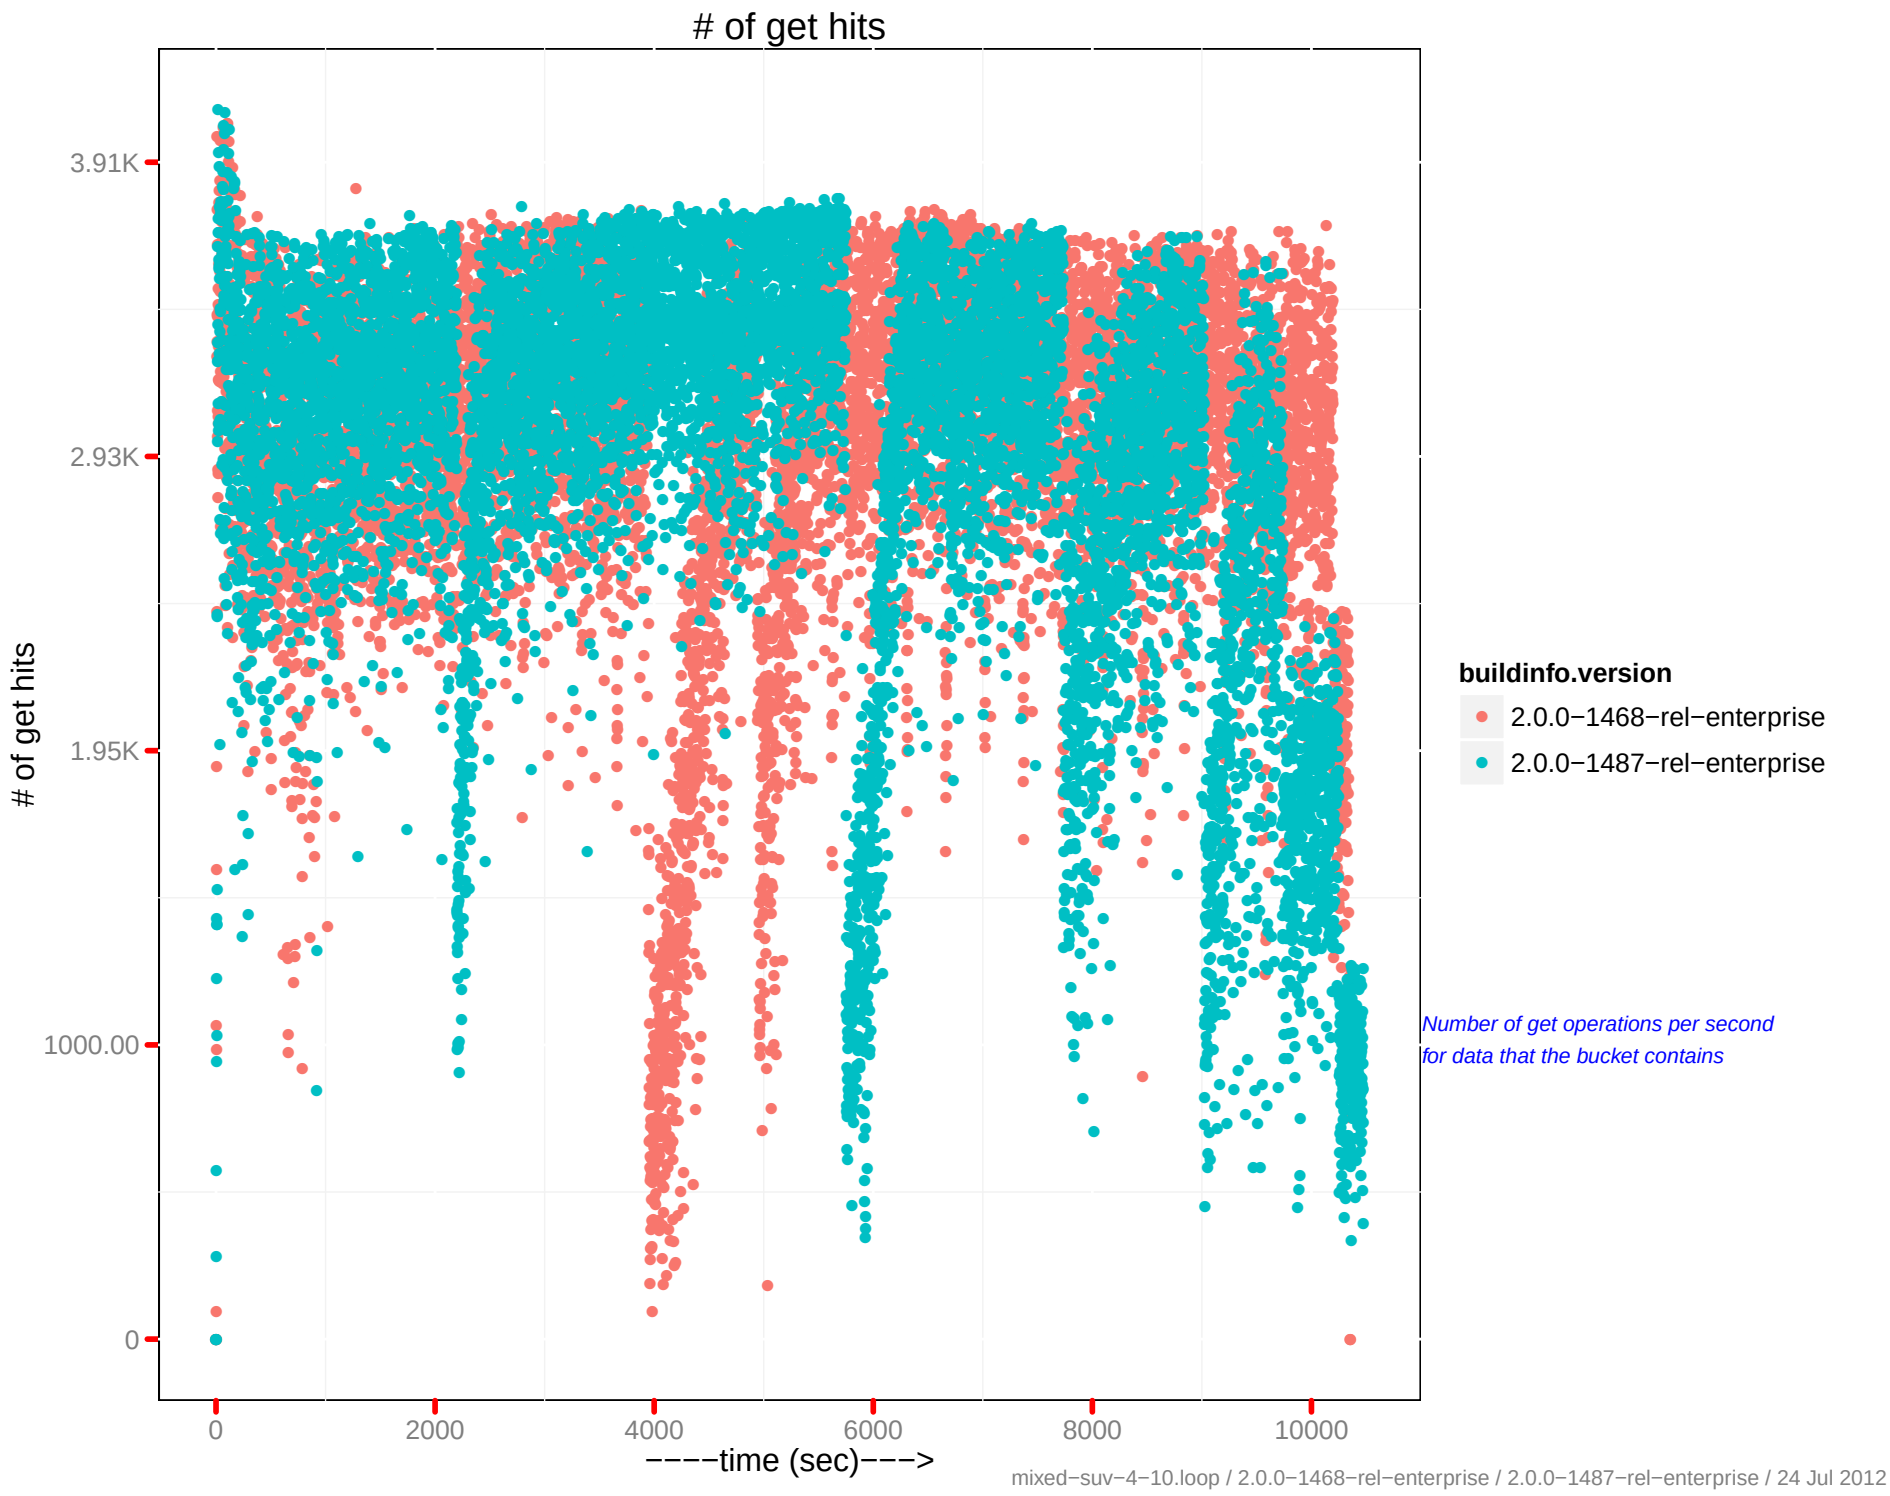
# of get hits
3.91K
2.93K
1.95K
1000.00
0
# of get hits
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Number of get operations per second
for data that the bucket contains
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

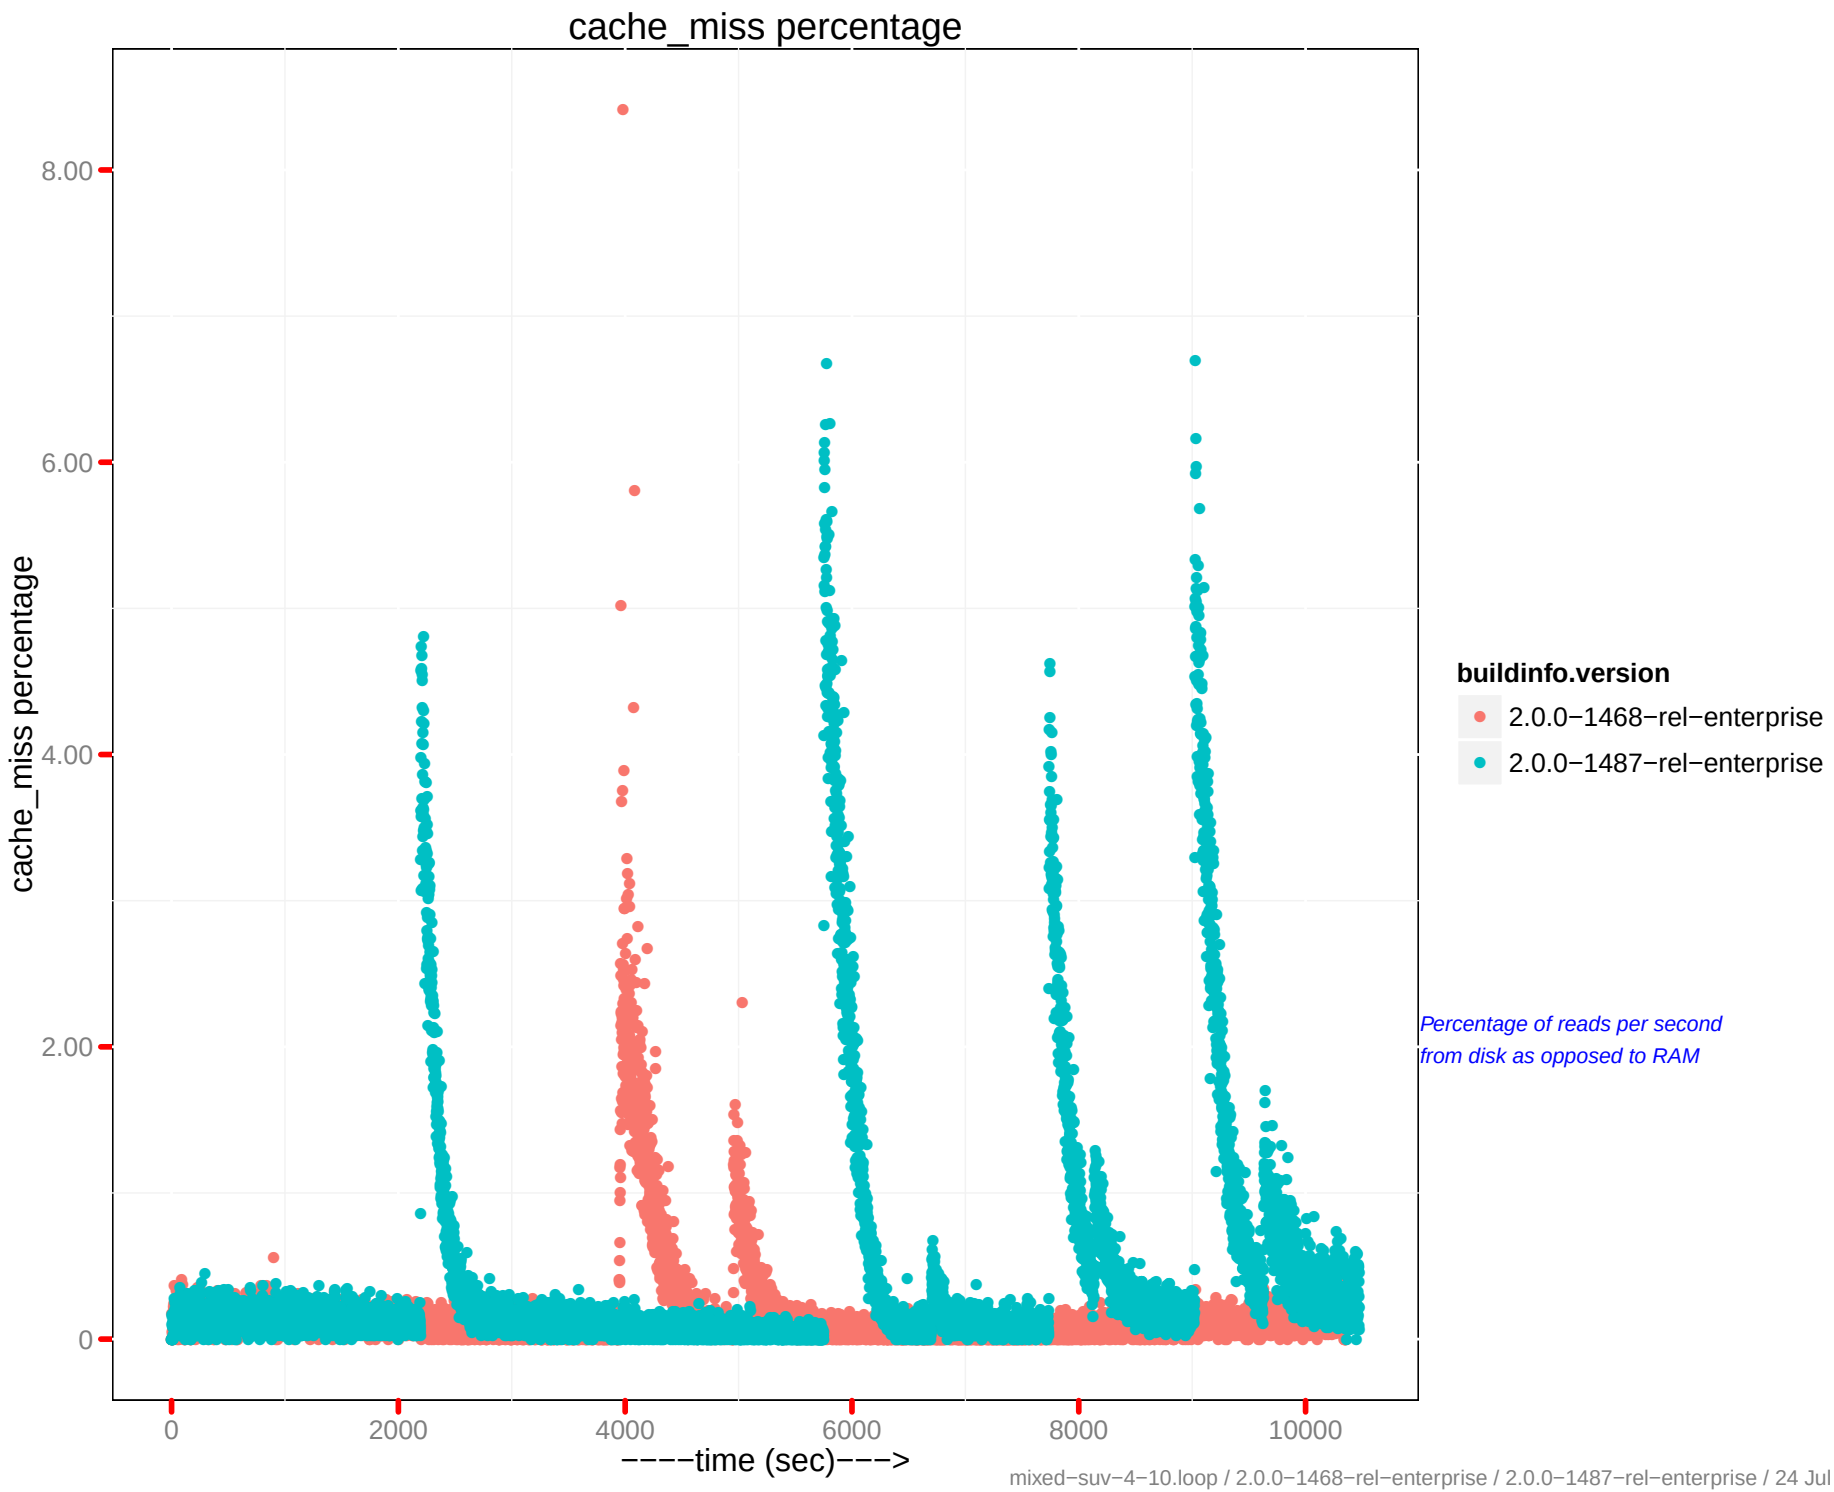
cache_miss percentage
8.00
6.00
4.00
2.00
0
cache_miss percentage
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Percentage of reads per second
from disk as opposed to RAM
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

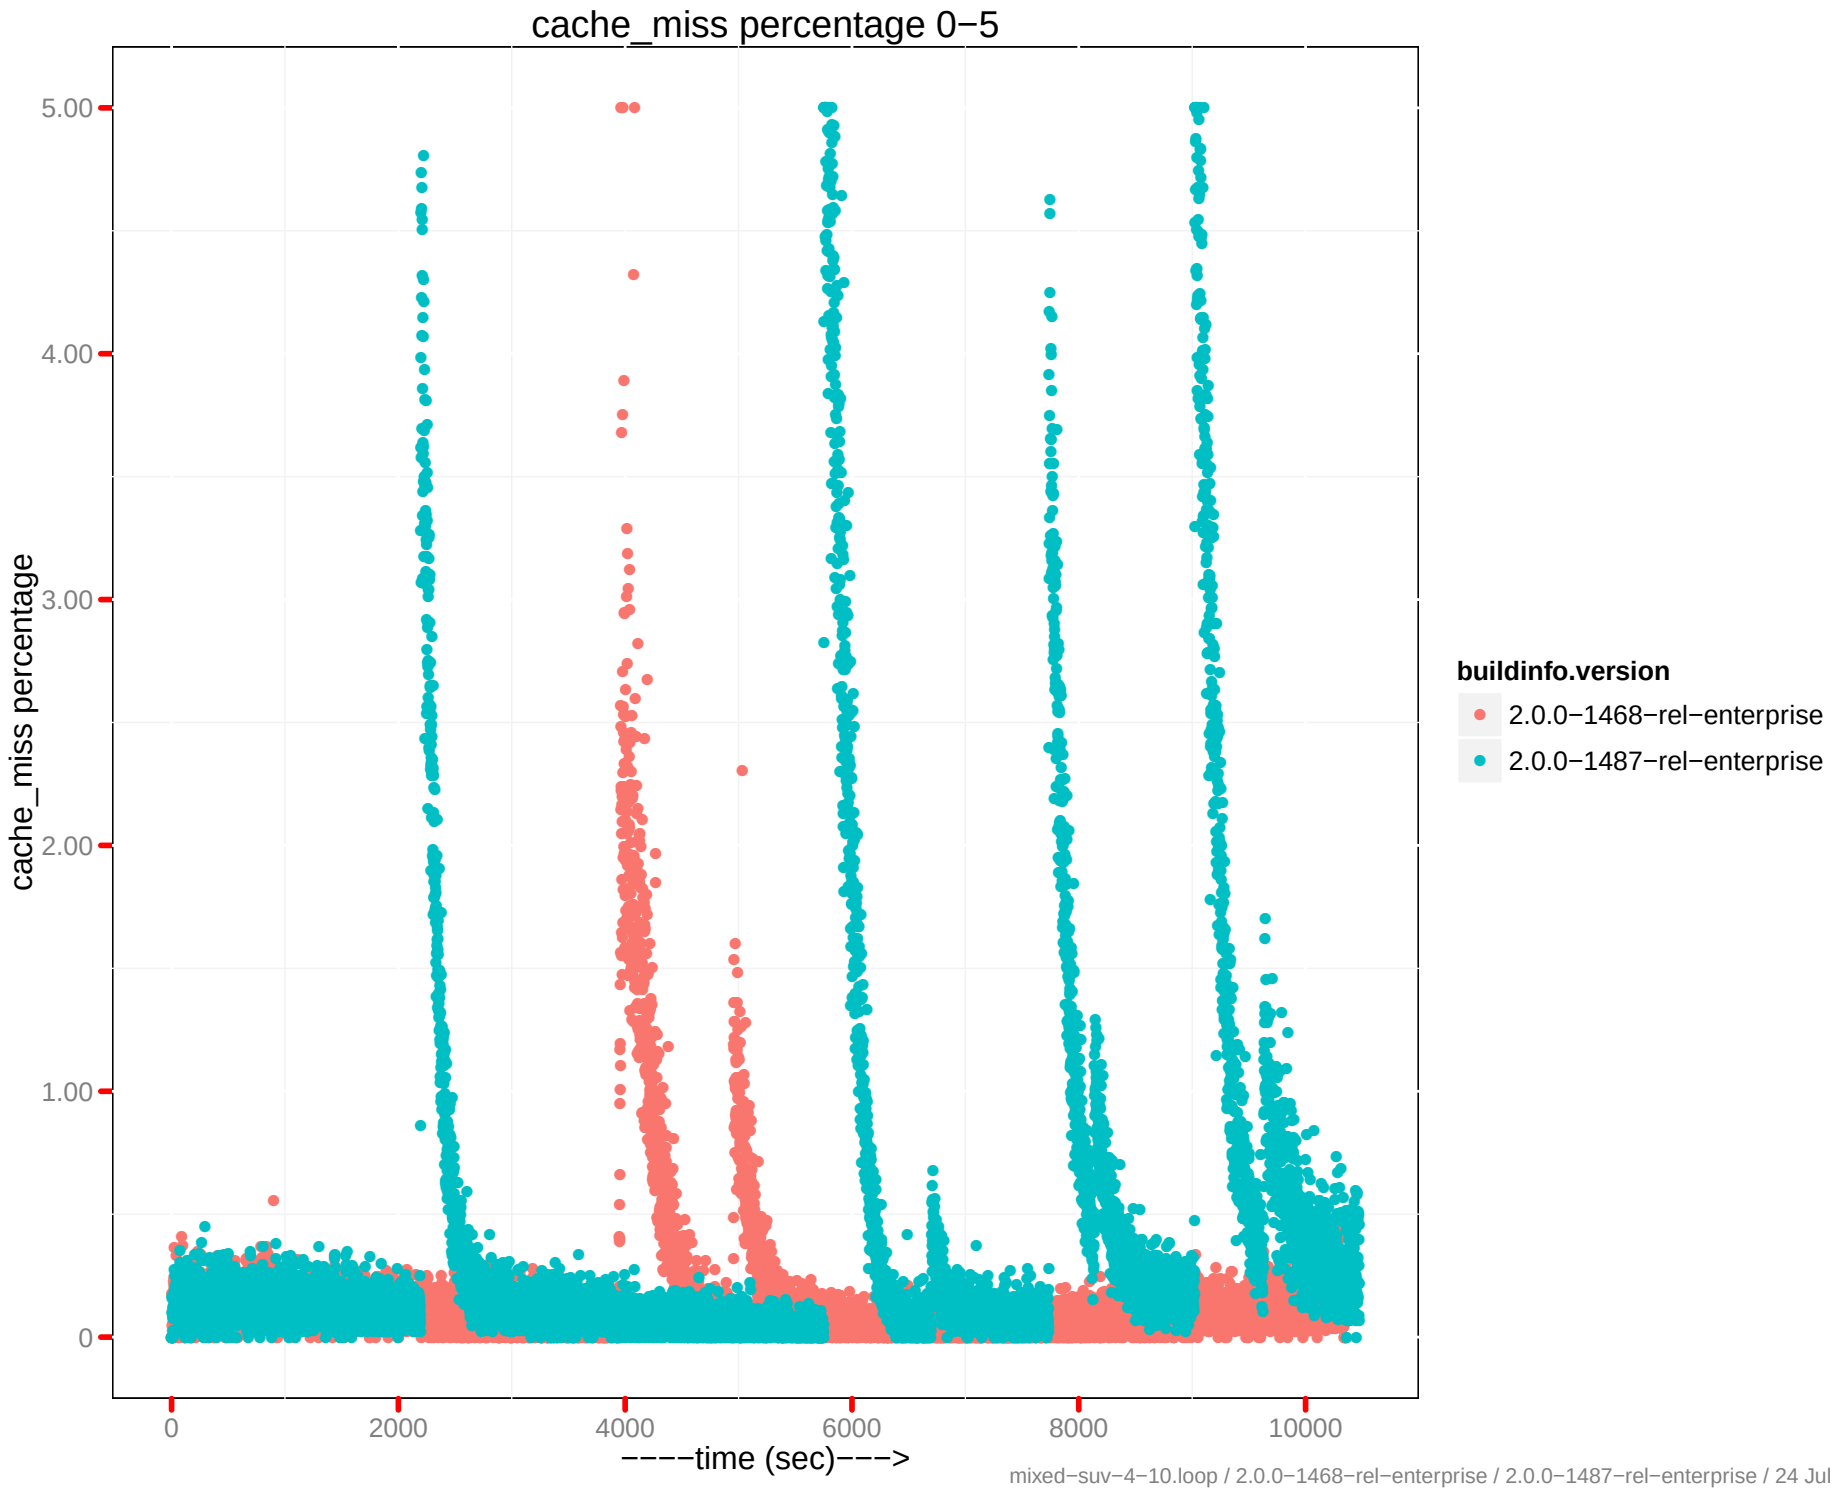
cache_miss percentage 0−5
cache_miss percentage
5.00
4.00
3.00
2.00
1.00
0
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

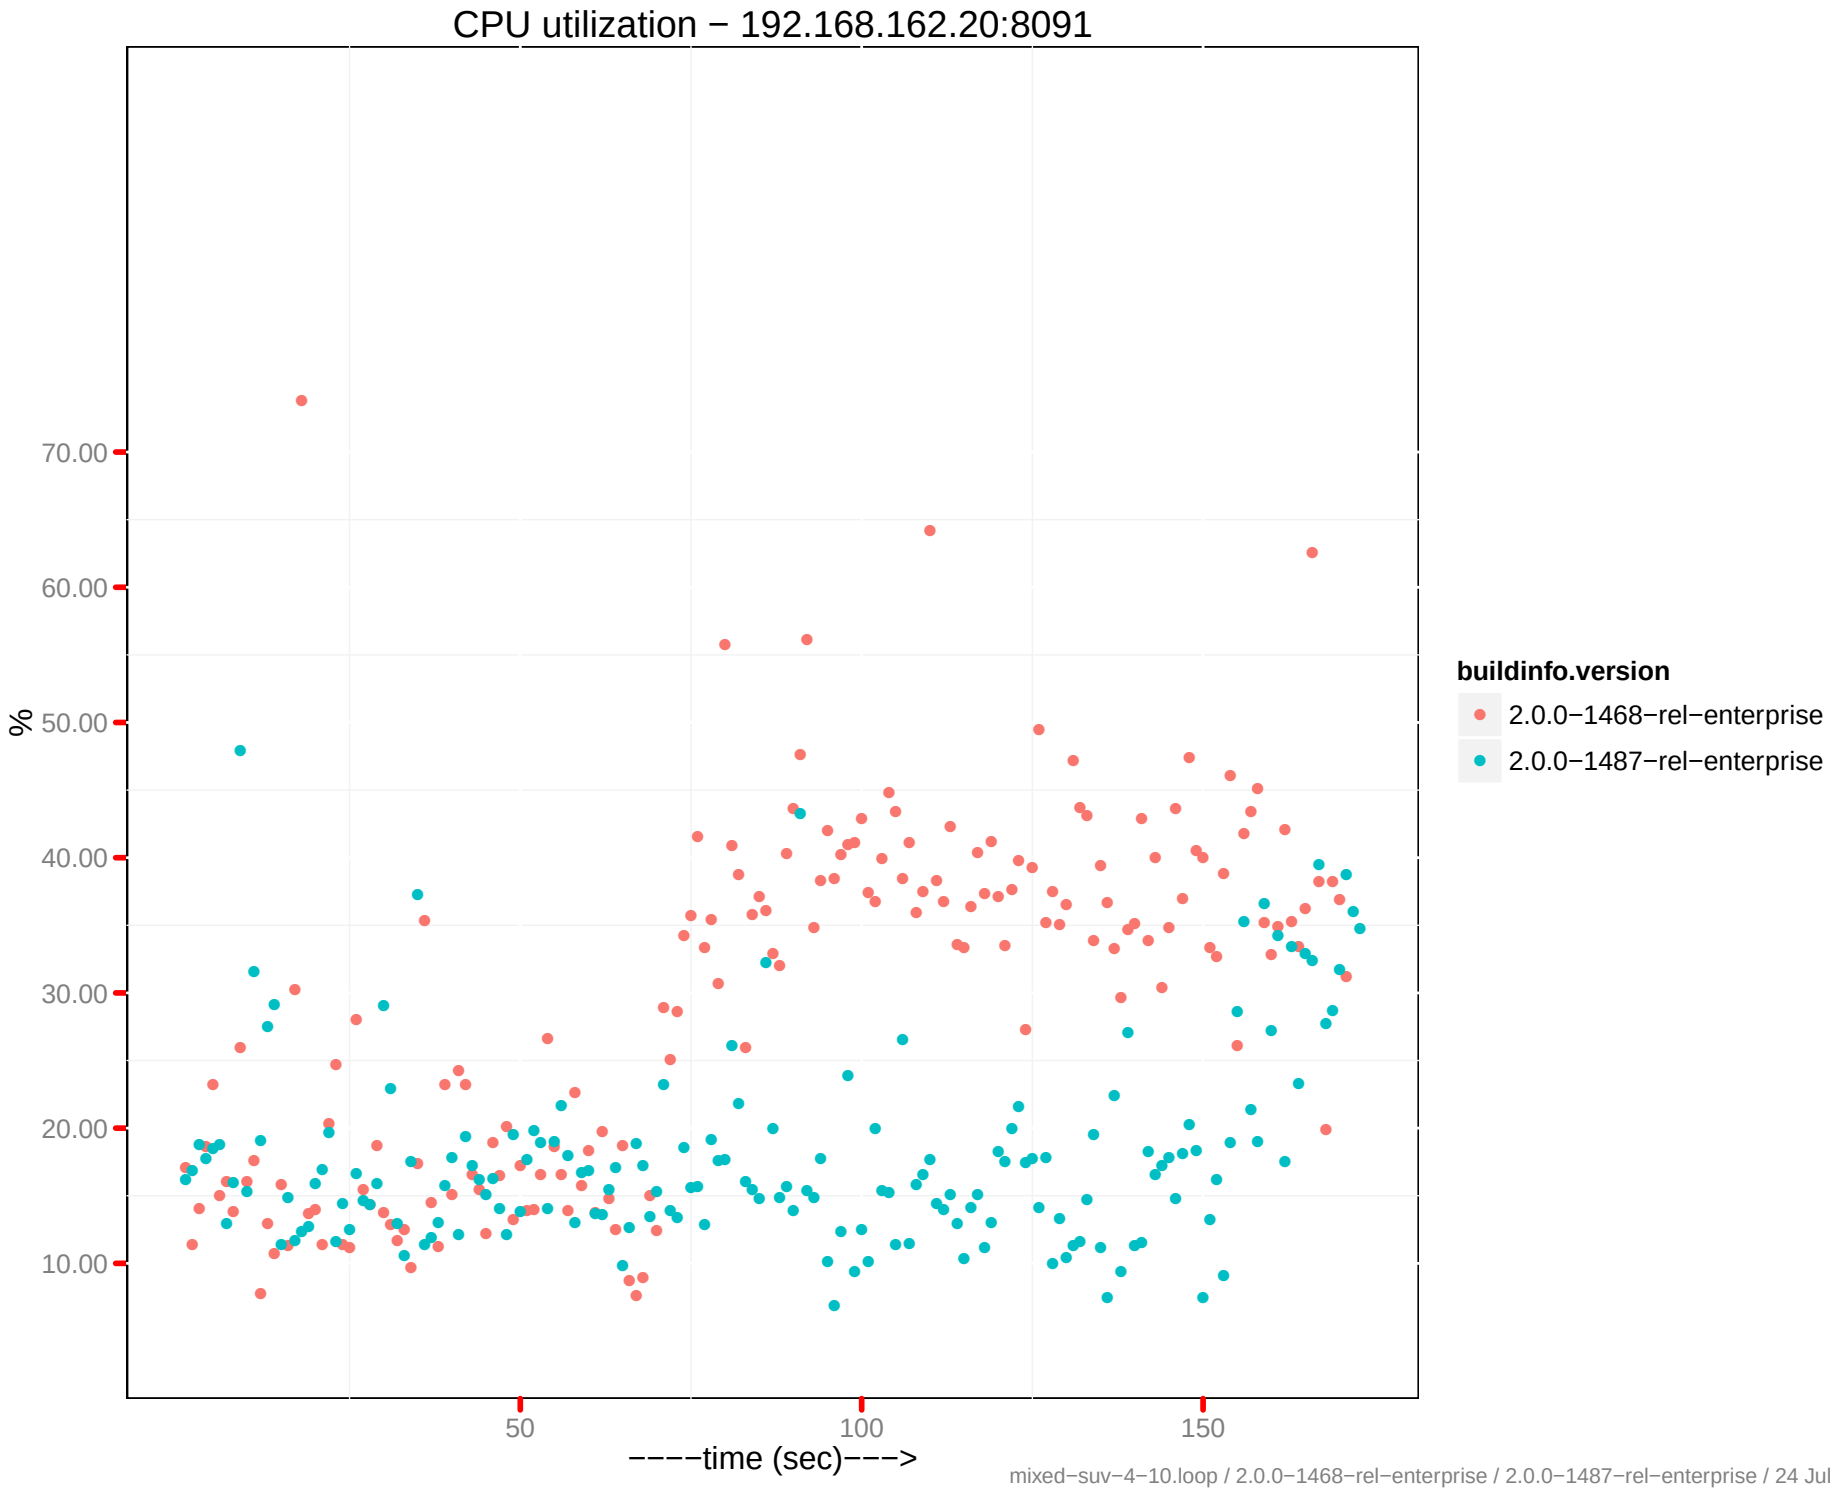
CPU utilization − 192.168.162.20:8091
70.00
60.00
50.00
40.00
30.00
20.00
10.00
%
50
100
150
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

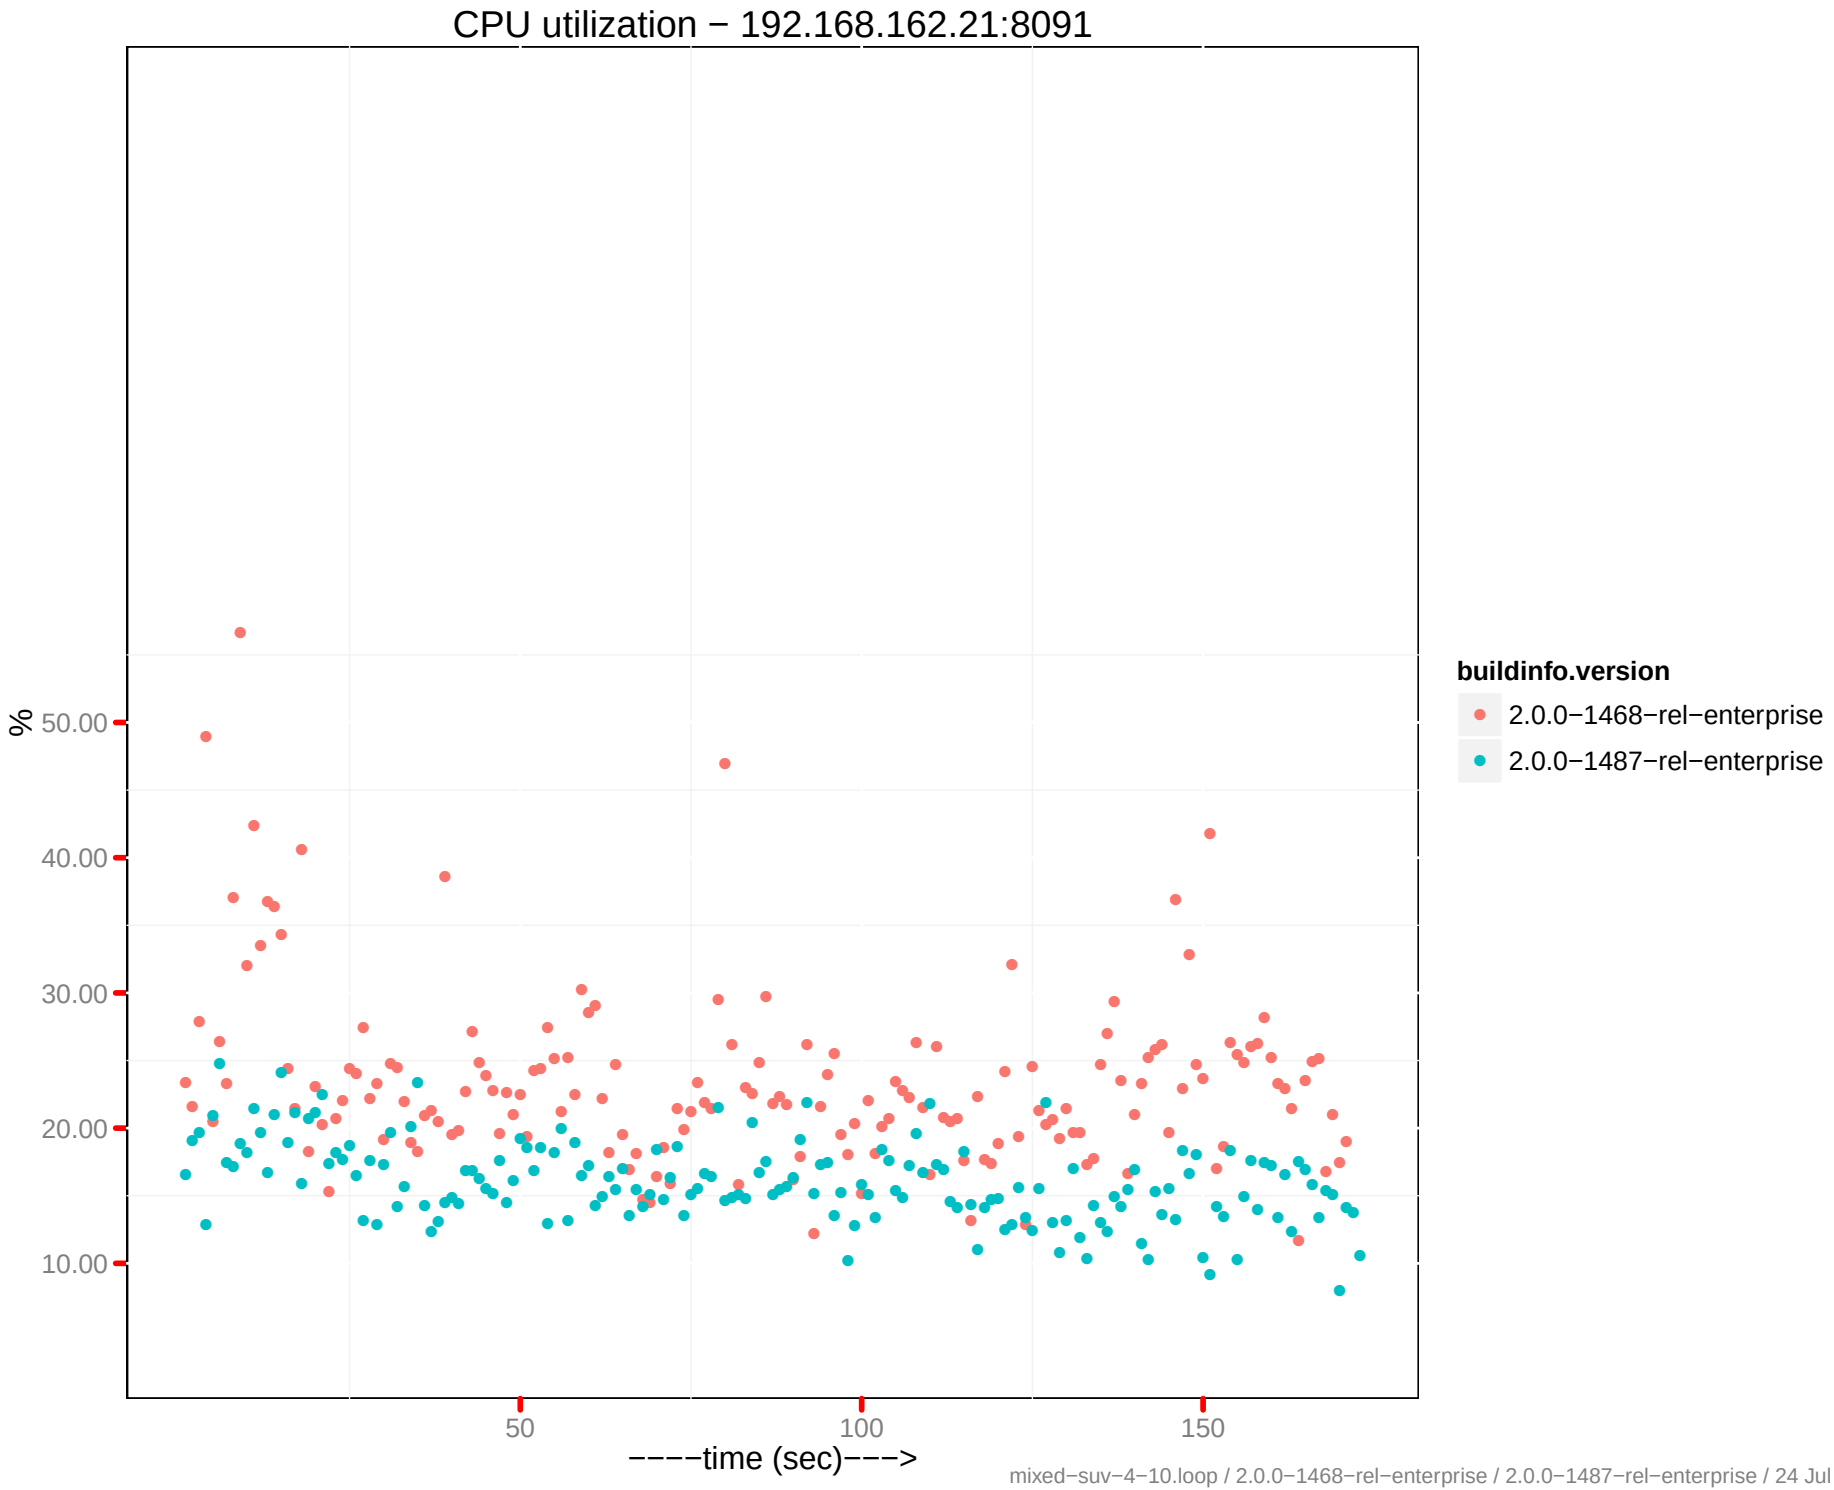
CPU utilization − 192.168.162.21:8091
%
50.00
40.00
30.00
20.00
10.00
50
100
150
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

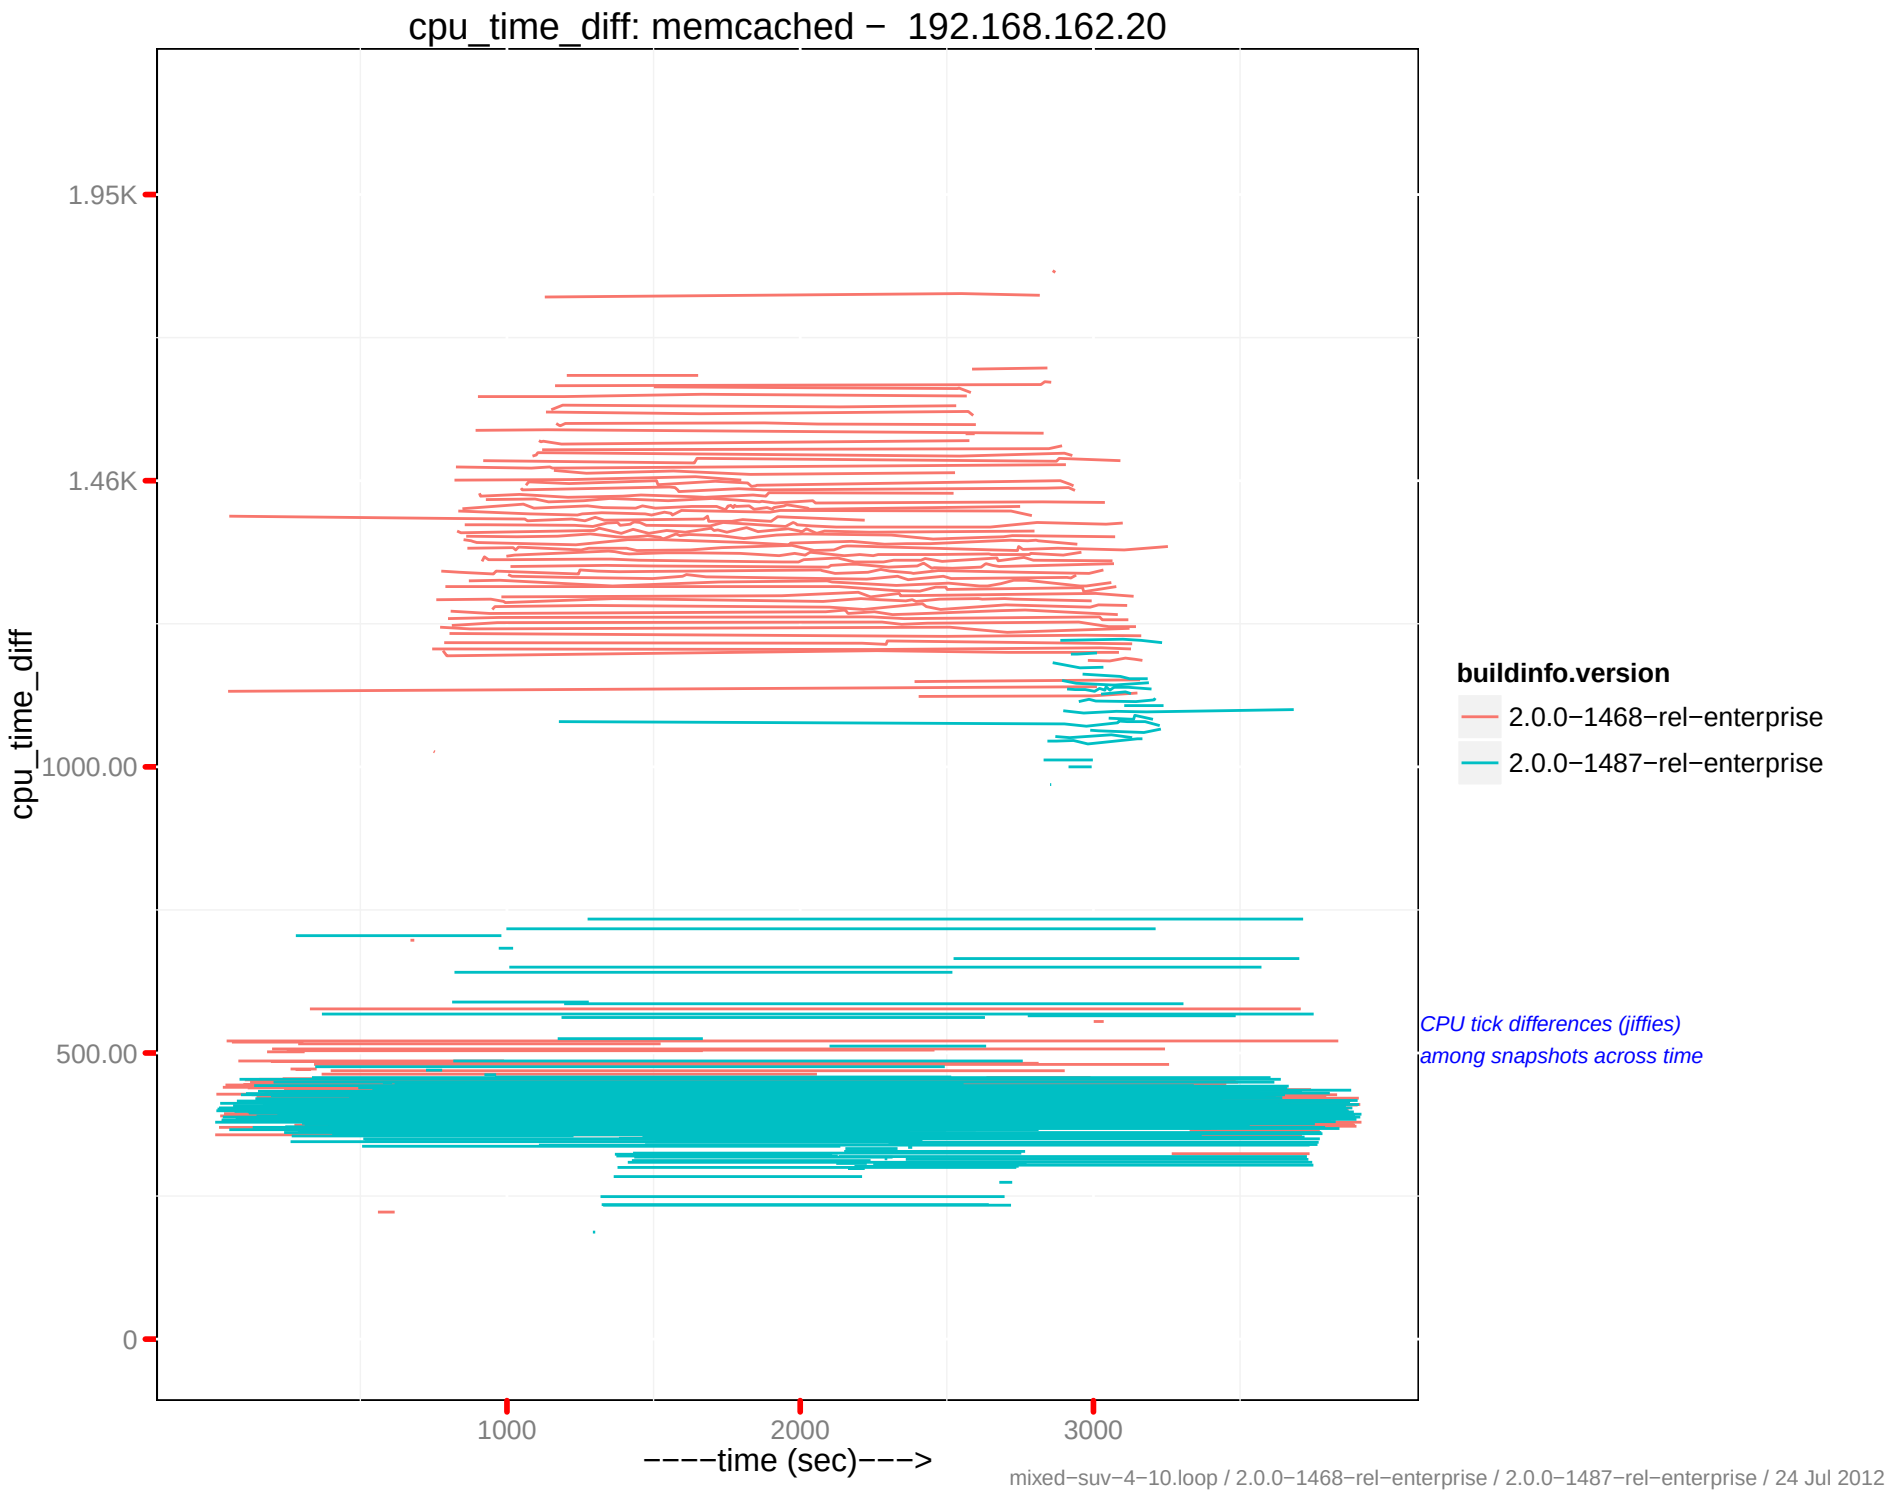
cpu_time_diff: memcached − 192.168.162.20
1.95K
1.46K
1000.00
500.00
0
cpu_time_diff
1000
2000
3000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

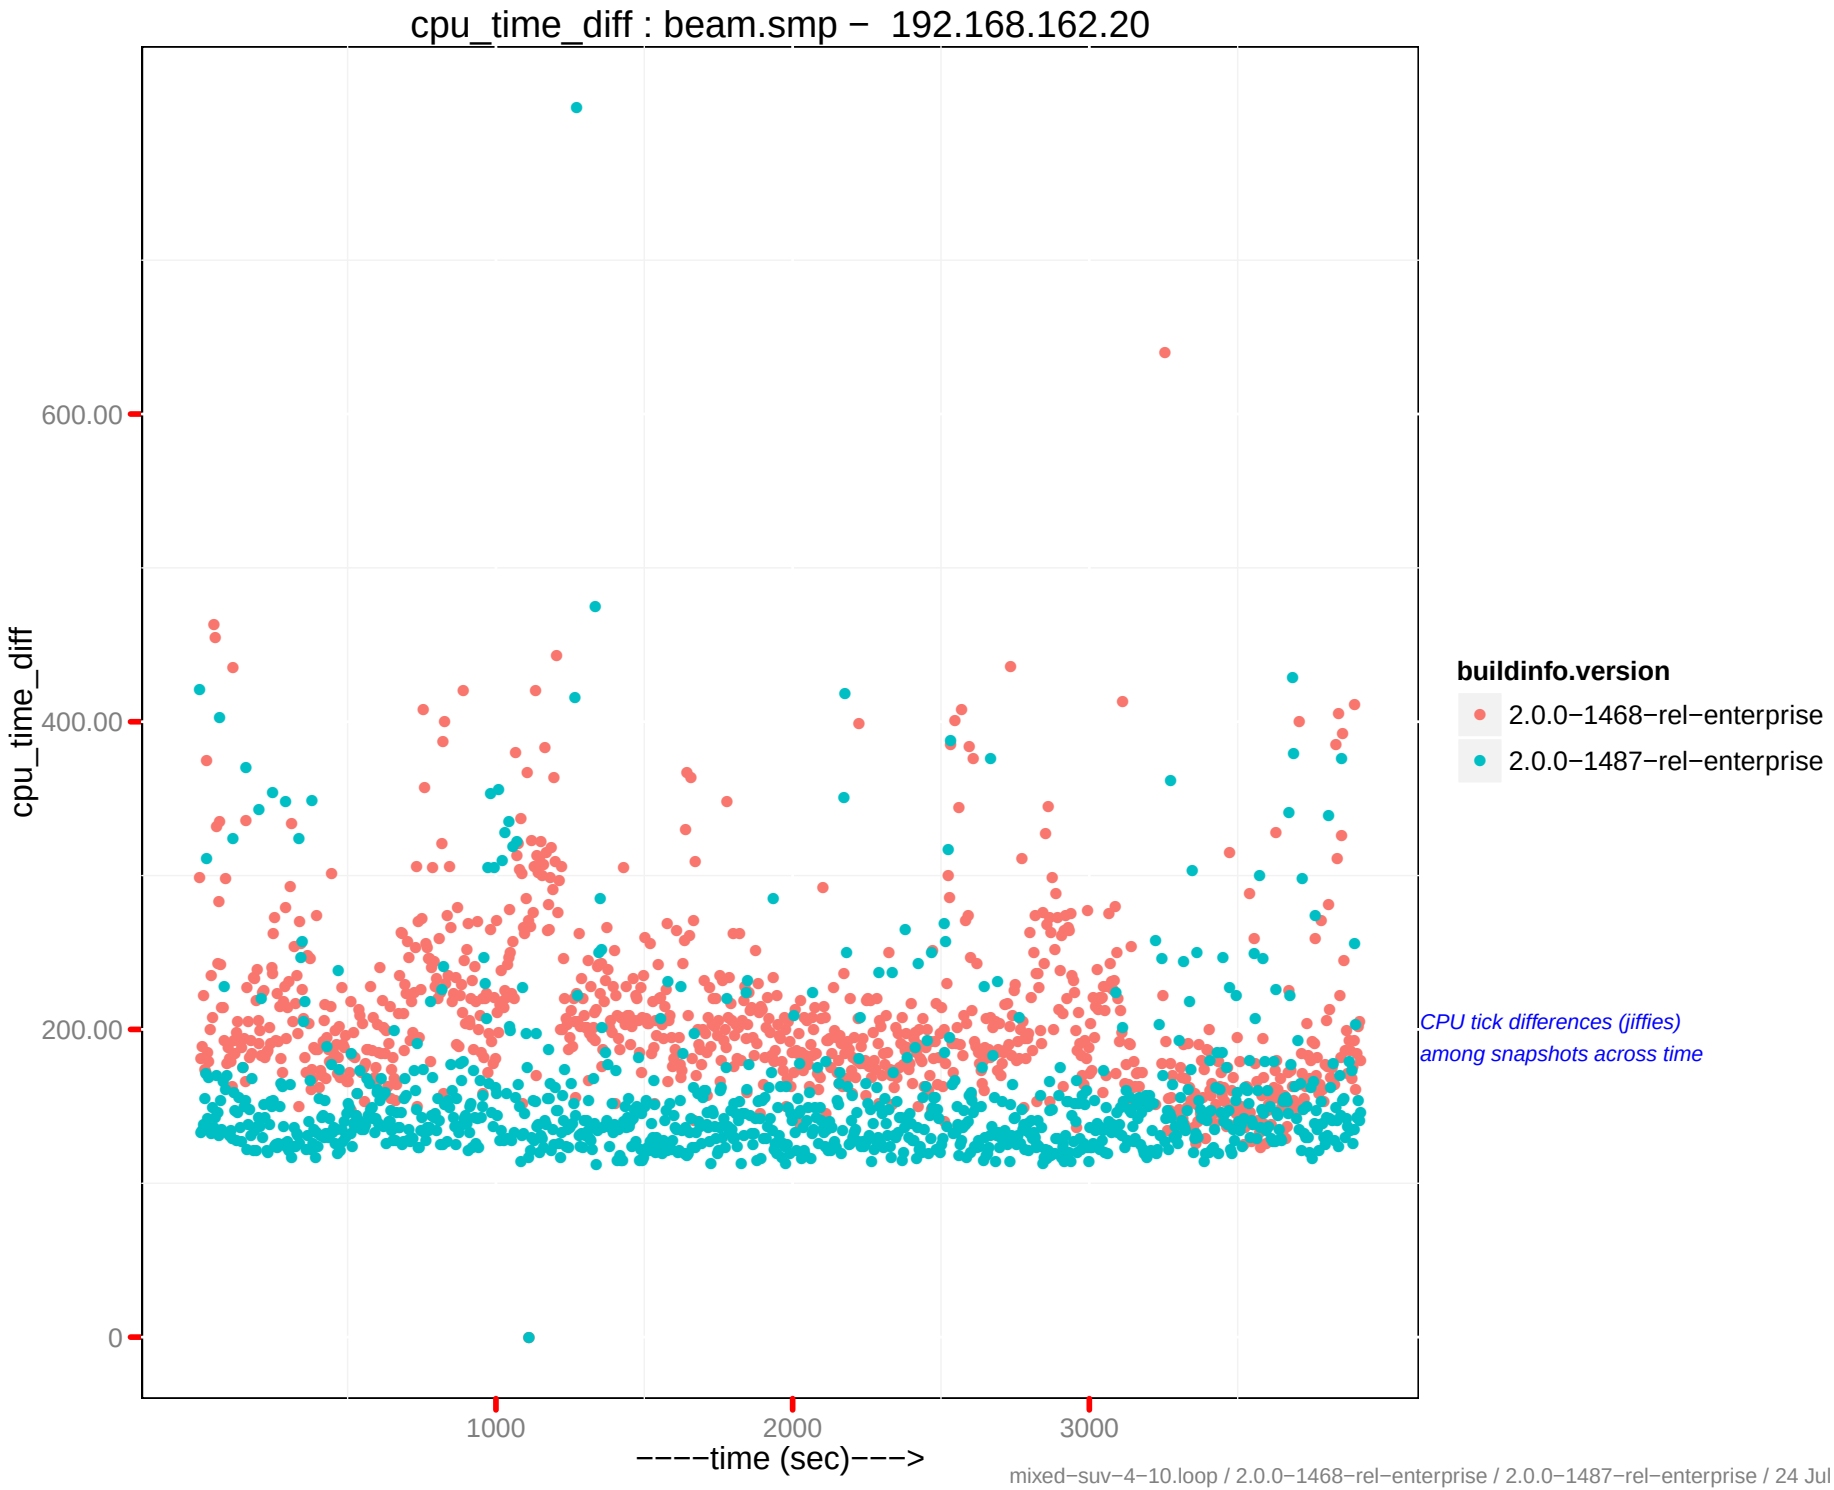
cpu_time_diff : beam.smp − 192.168.162.20
cpu_time_diff
600.00
400.00
200.00
0
1000
2000
3000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

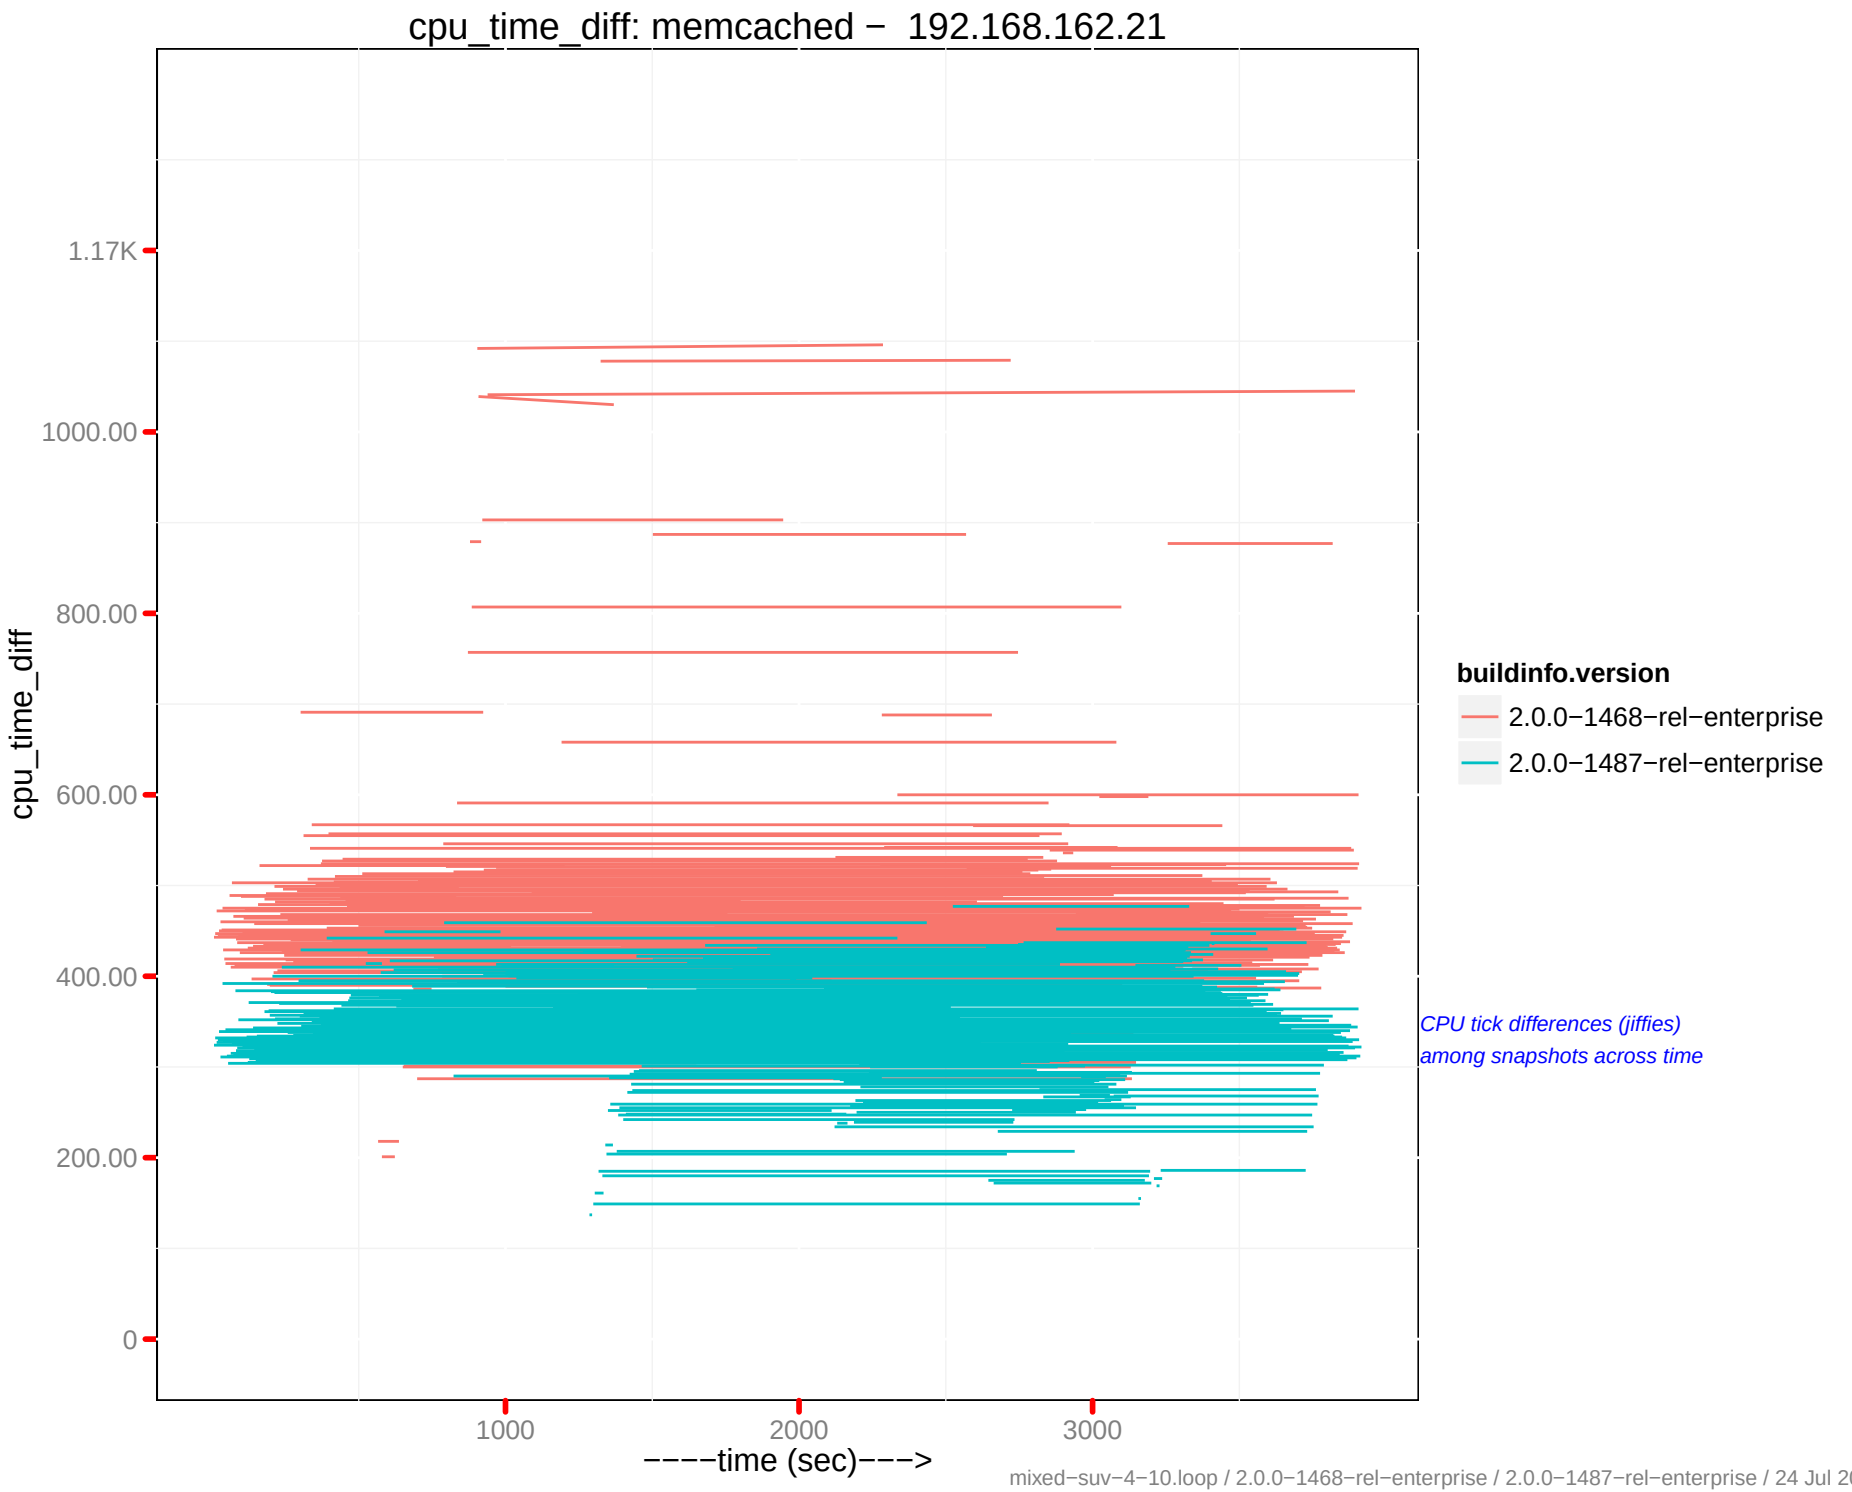
cpu_time_diff: memcached − 192.168.162.21
1.17K
1000.00
800.00
600.00
400.00
200.00
0
cpu_time_diff
1000
2000
3000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

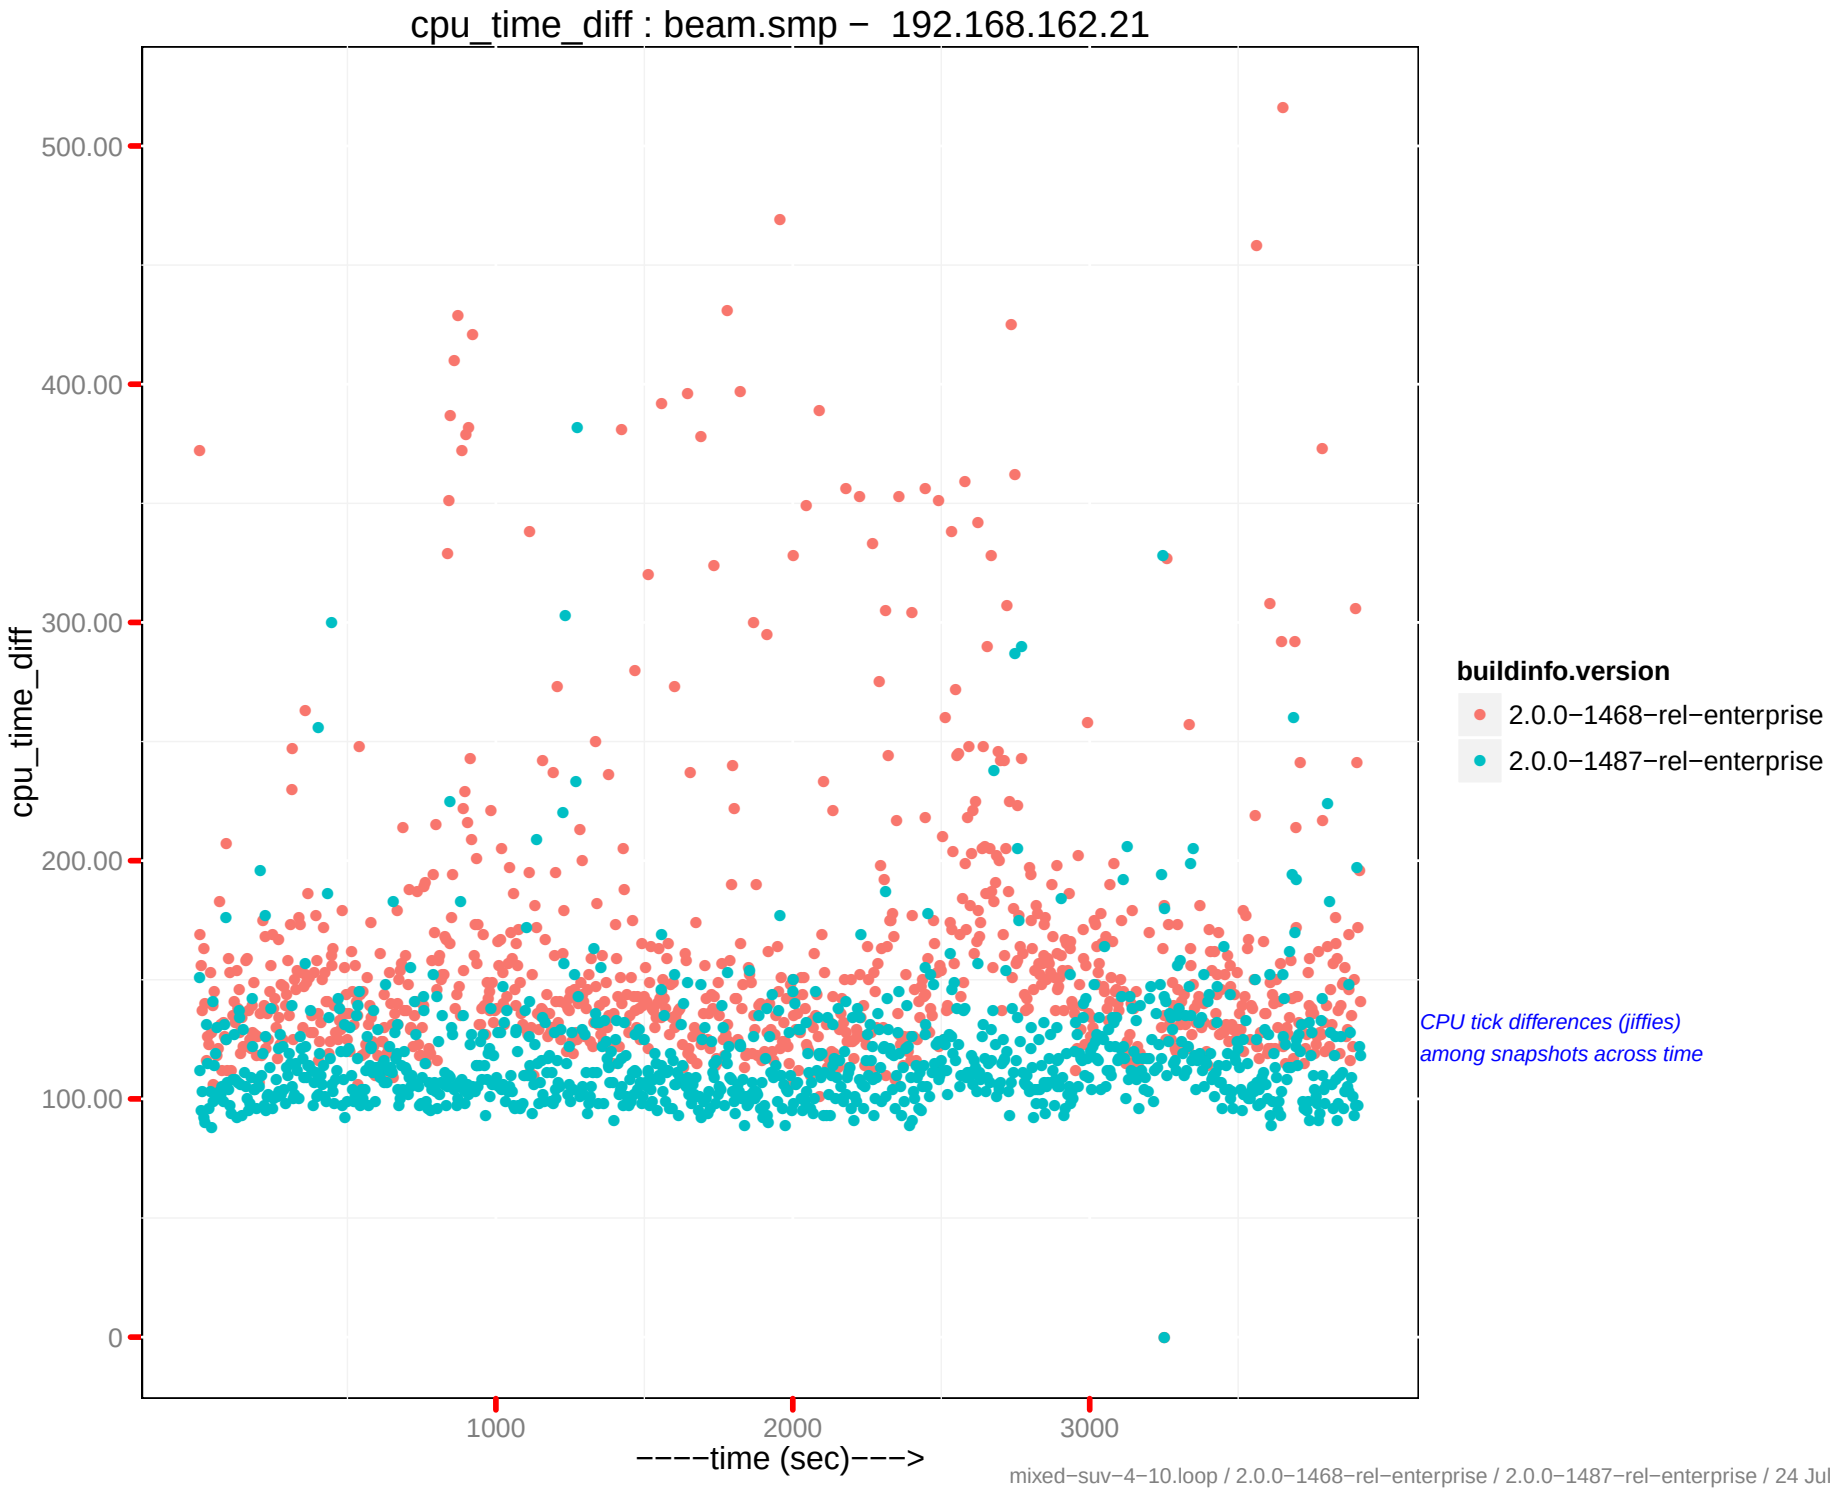
cpu_time_diff : beam.smp − 192.168.162.21
cpu_time_diff
500.00
400.00
300.00
200.00
100.00
0
1000
2000
3000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

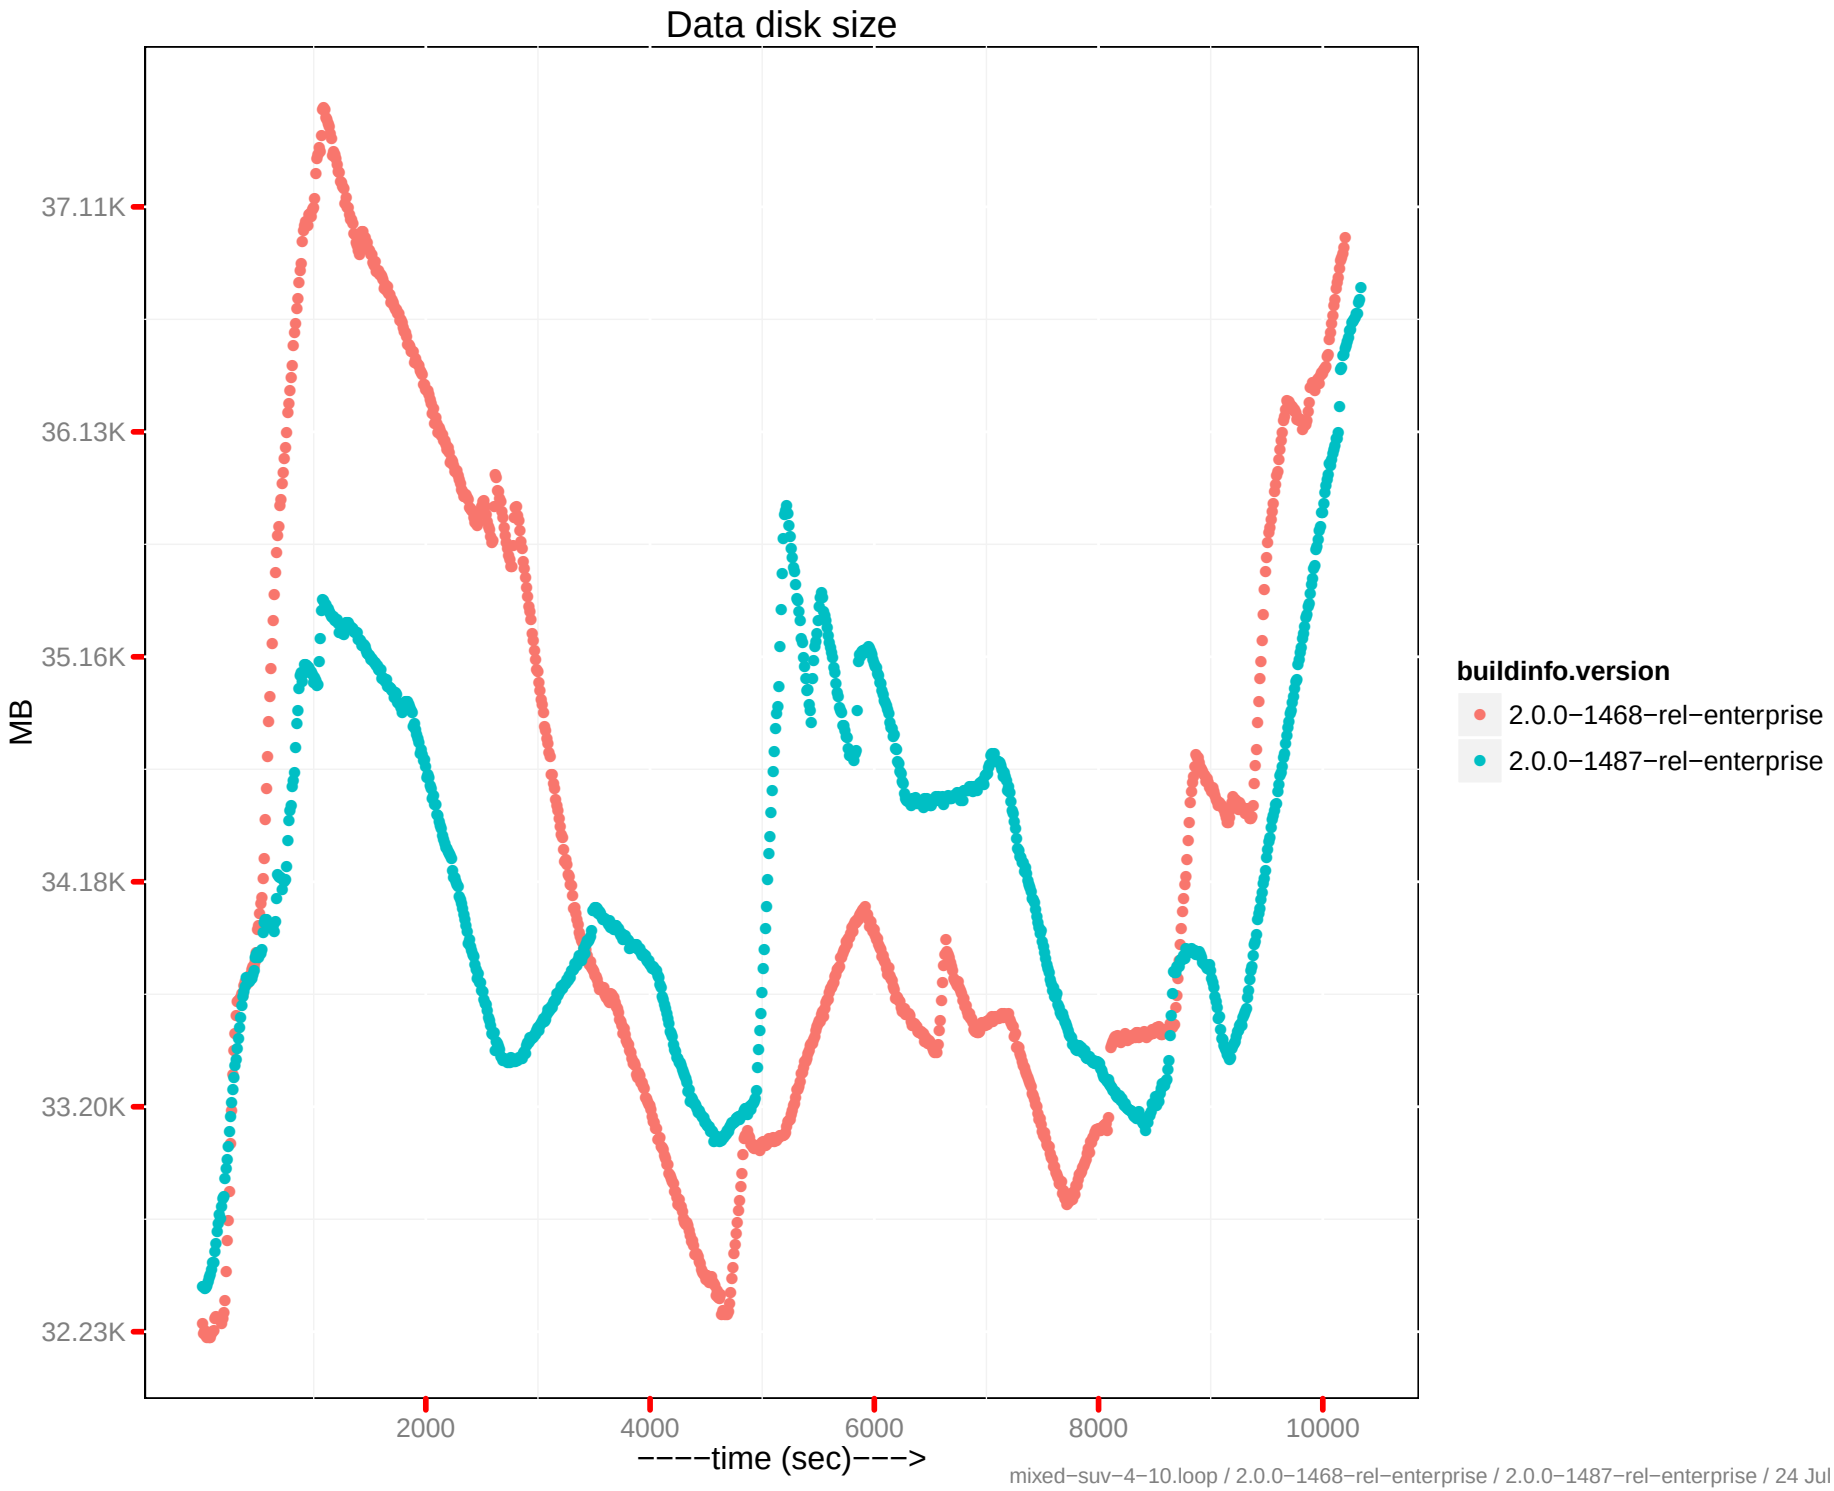
Data disk size
37.11K
36.13K
35.16K
34.18K
33.20K
32.23K
MB
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

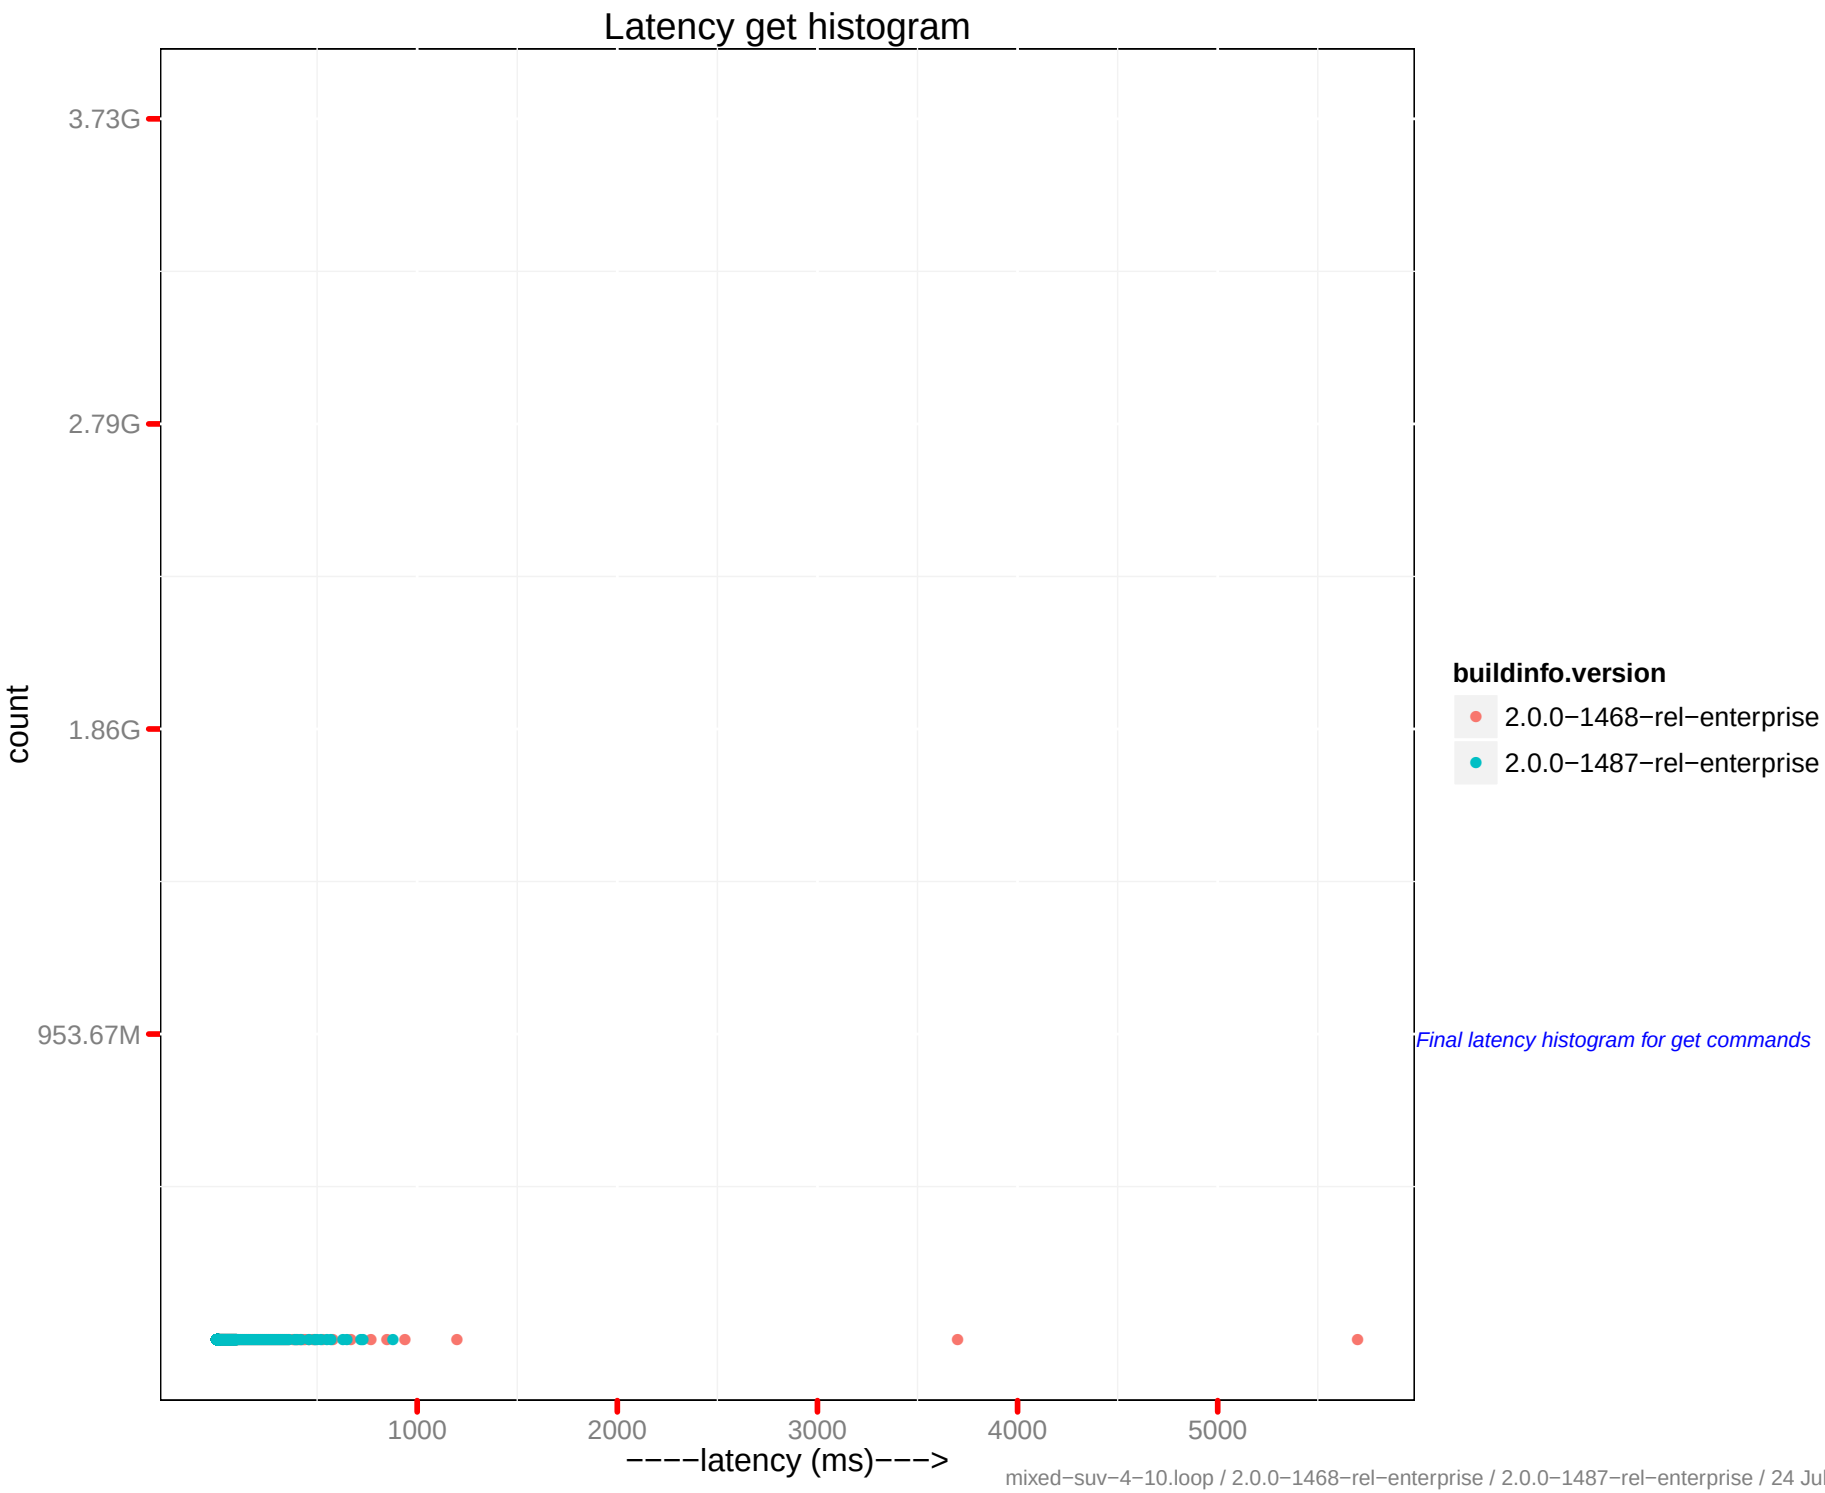
Latency get histogram
3.73G
2.79G
1.86G
953.67M
count
1000
2000
3000
4000
5000
----latency (ms)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Final latency histogram for get commands
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

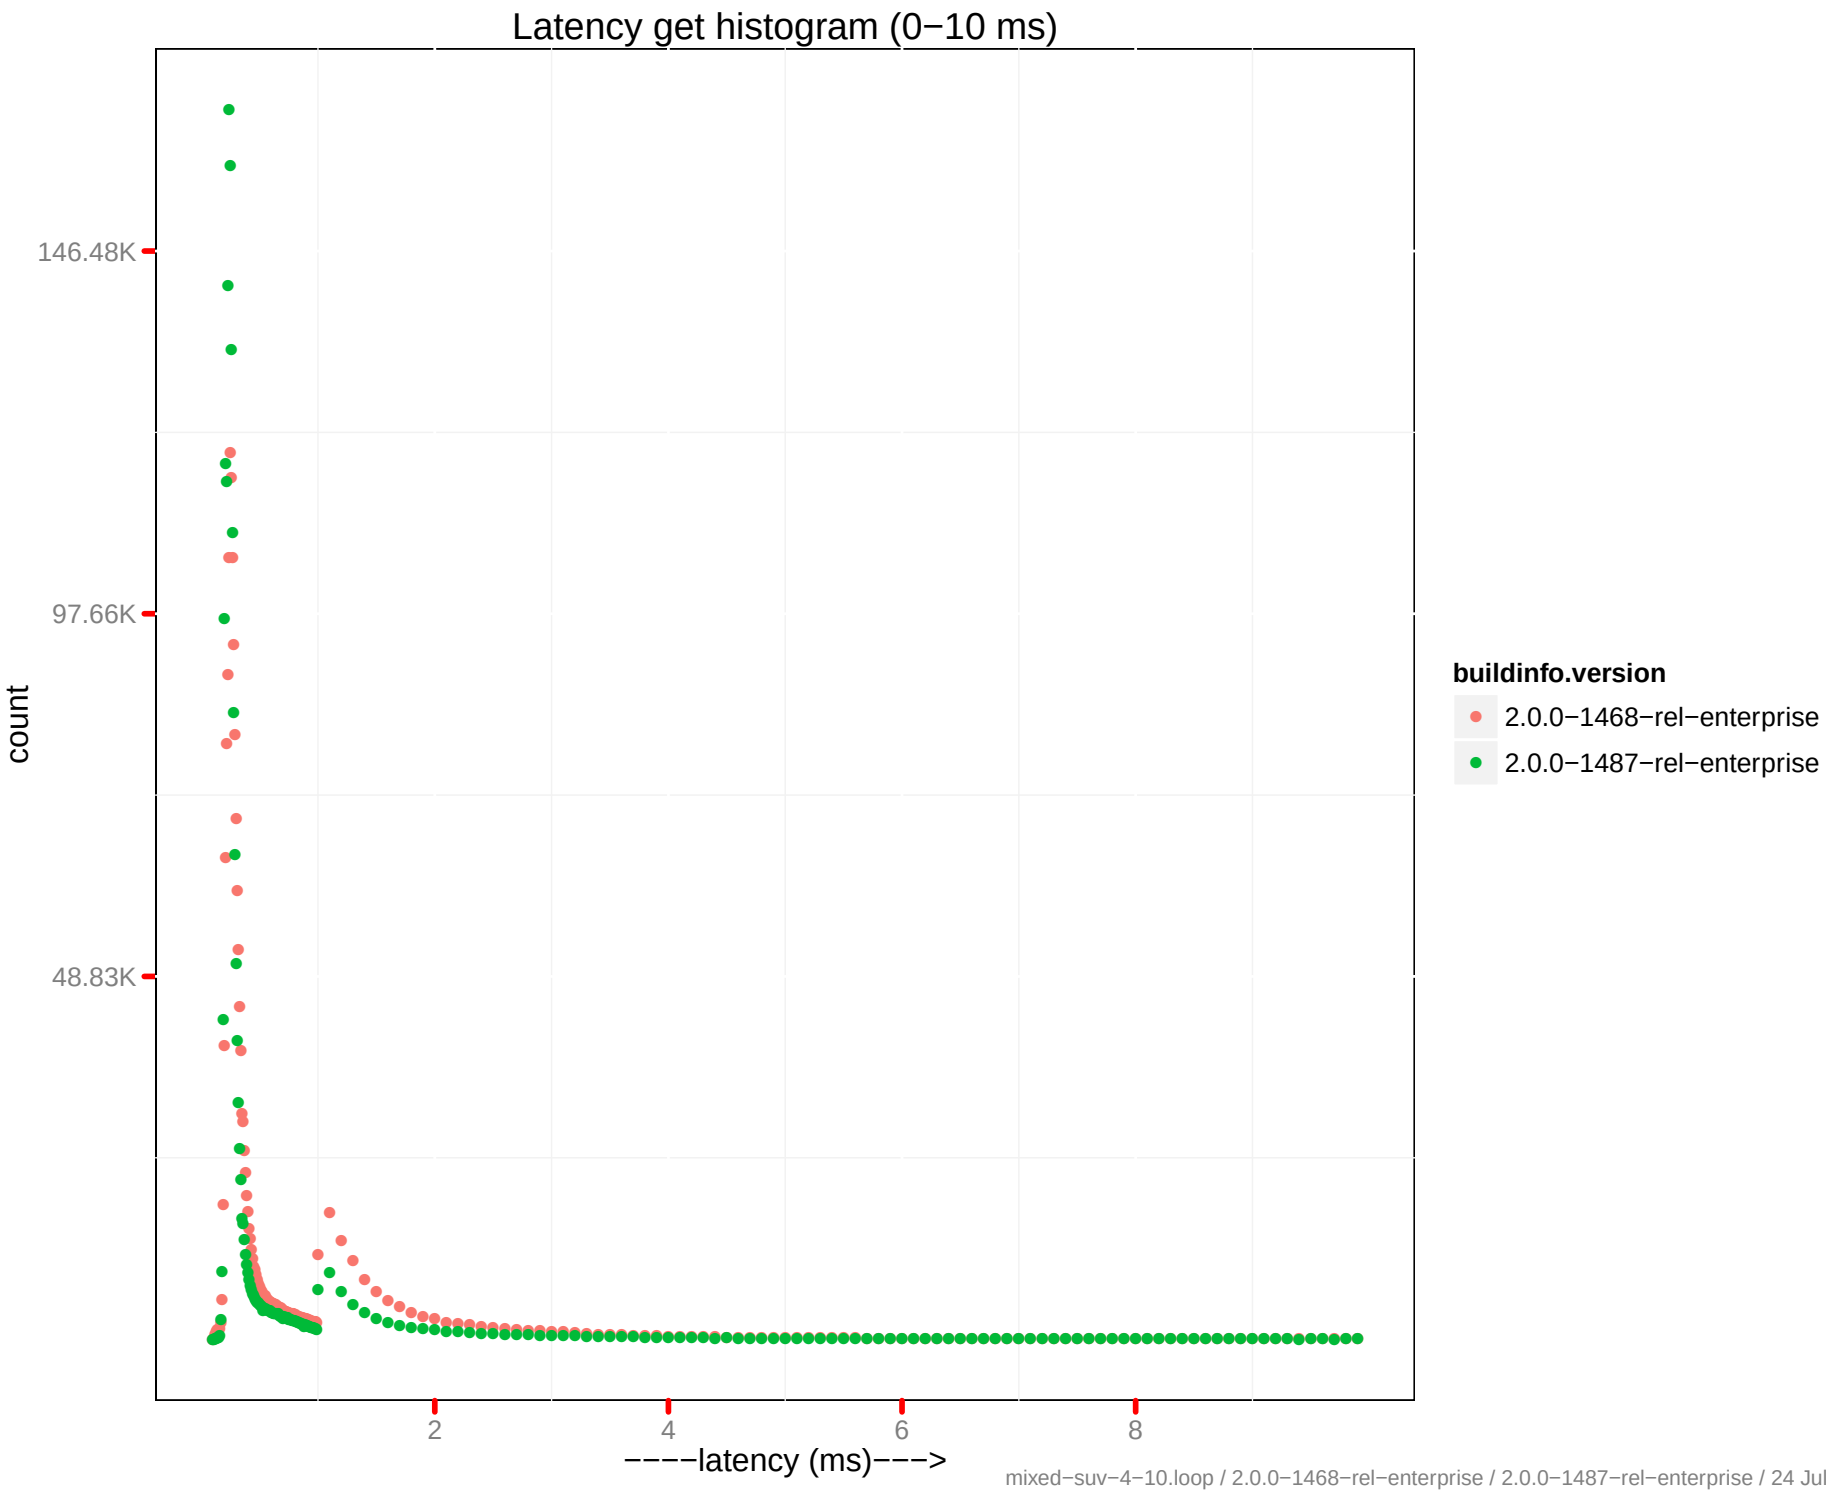
Latency get histogram (0−10 ms)
146.48K
97.66K
48.83K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

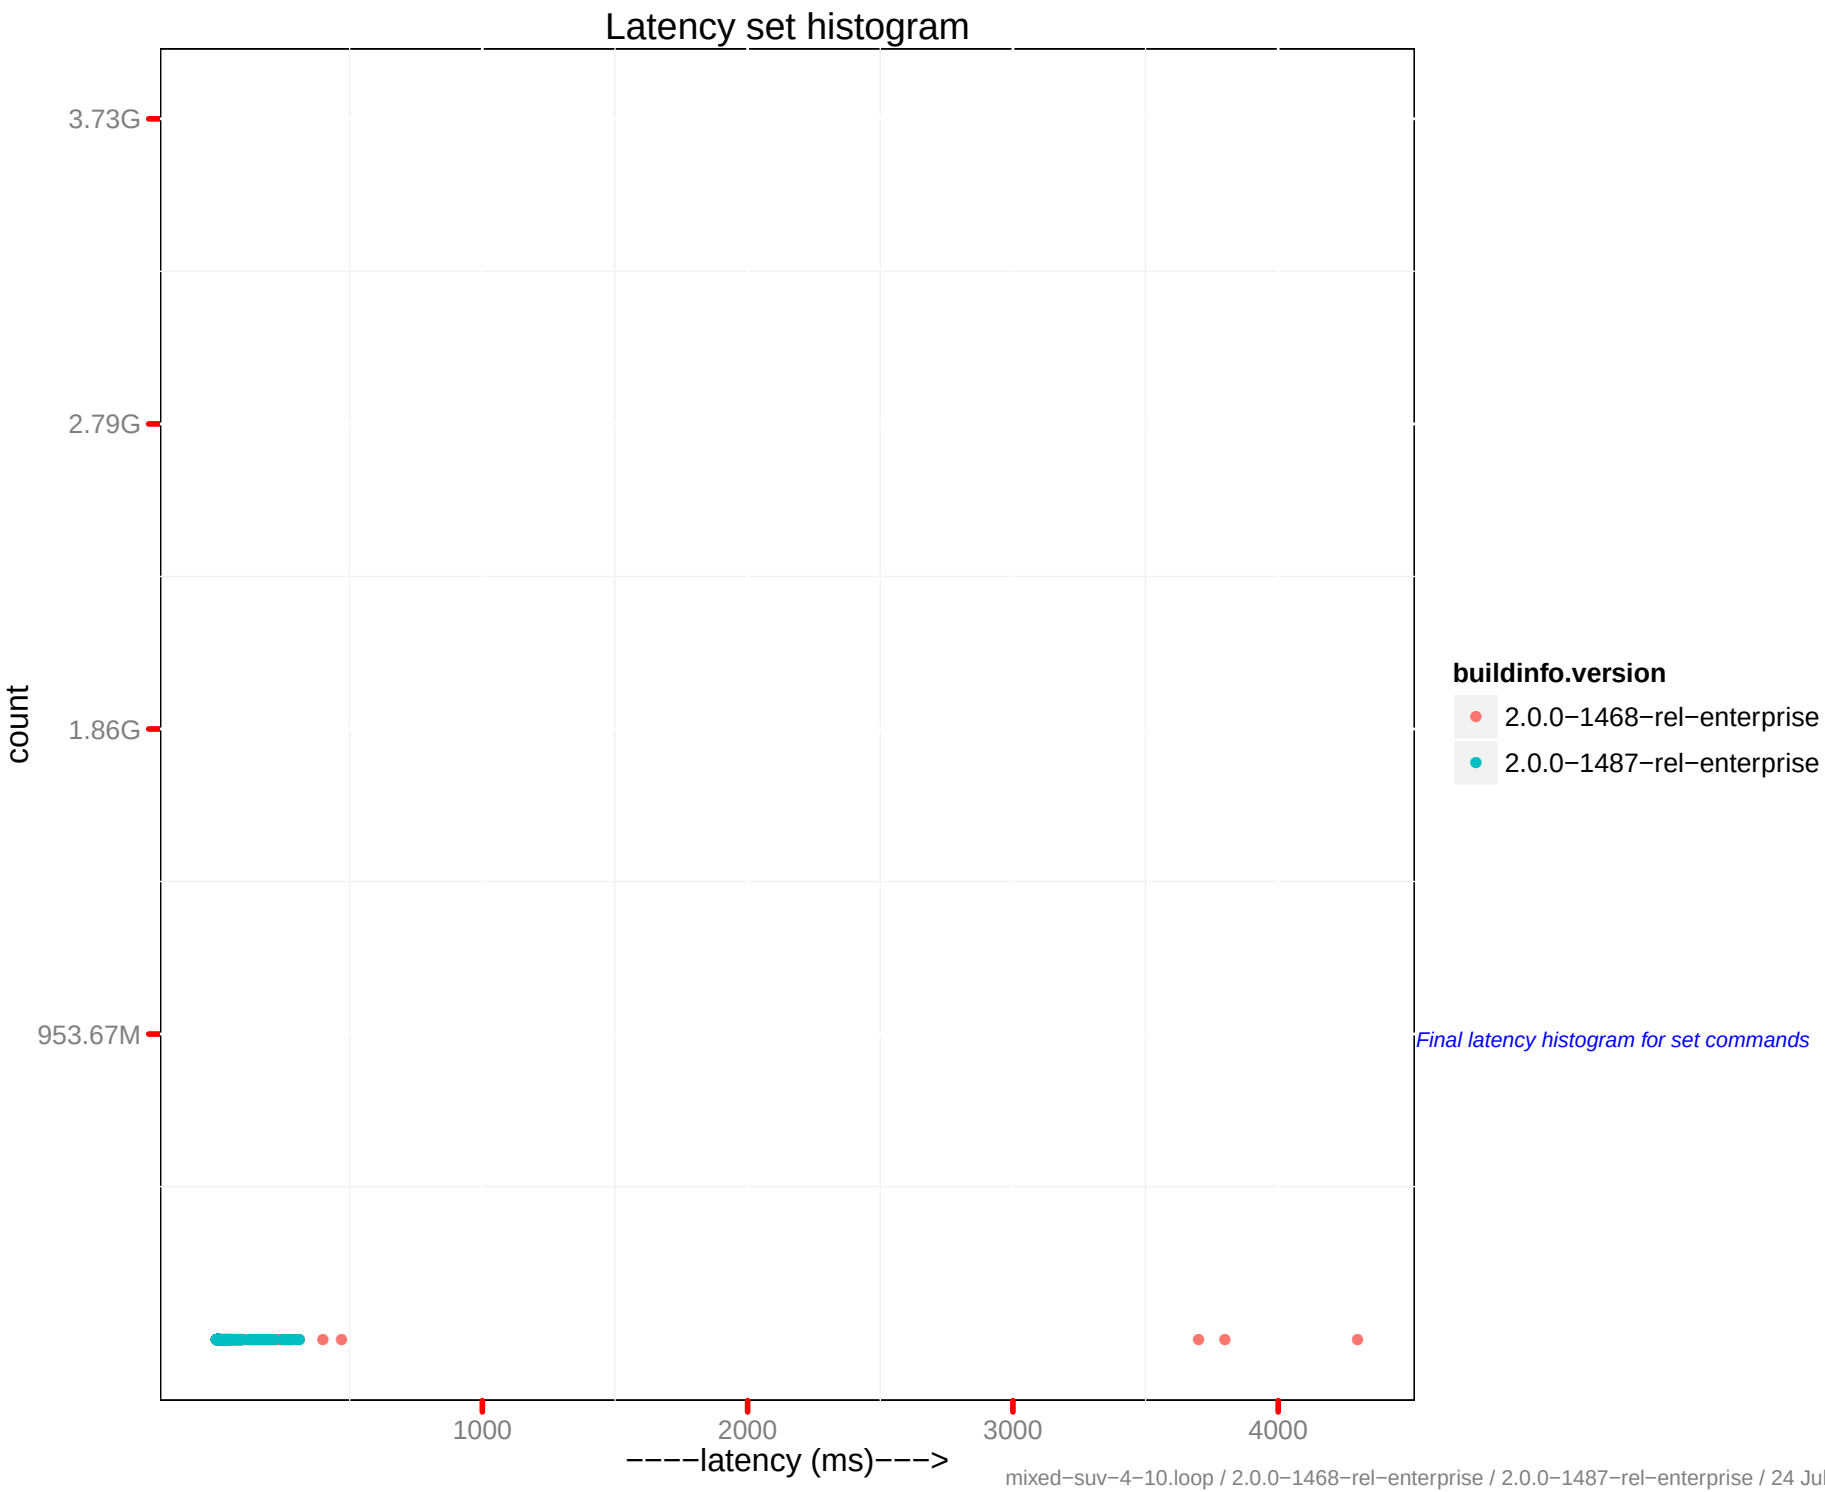
Latency set histogram
3.73G
2.79G
1.86G
953.67M
count
1000
2000
3000
4000
----latency (ms)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
Final latency histogram for set commands
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

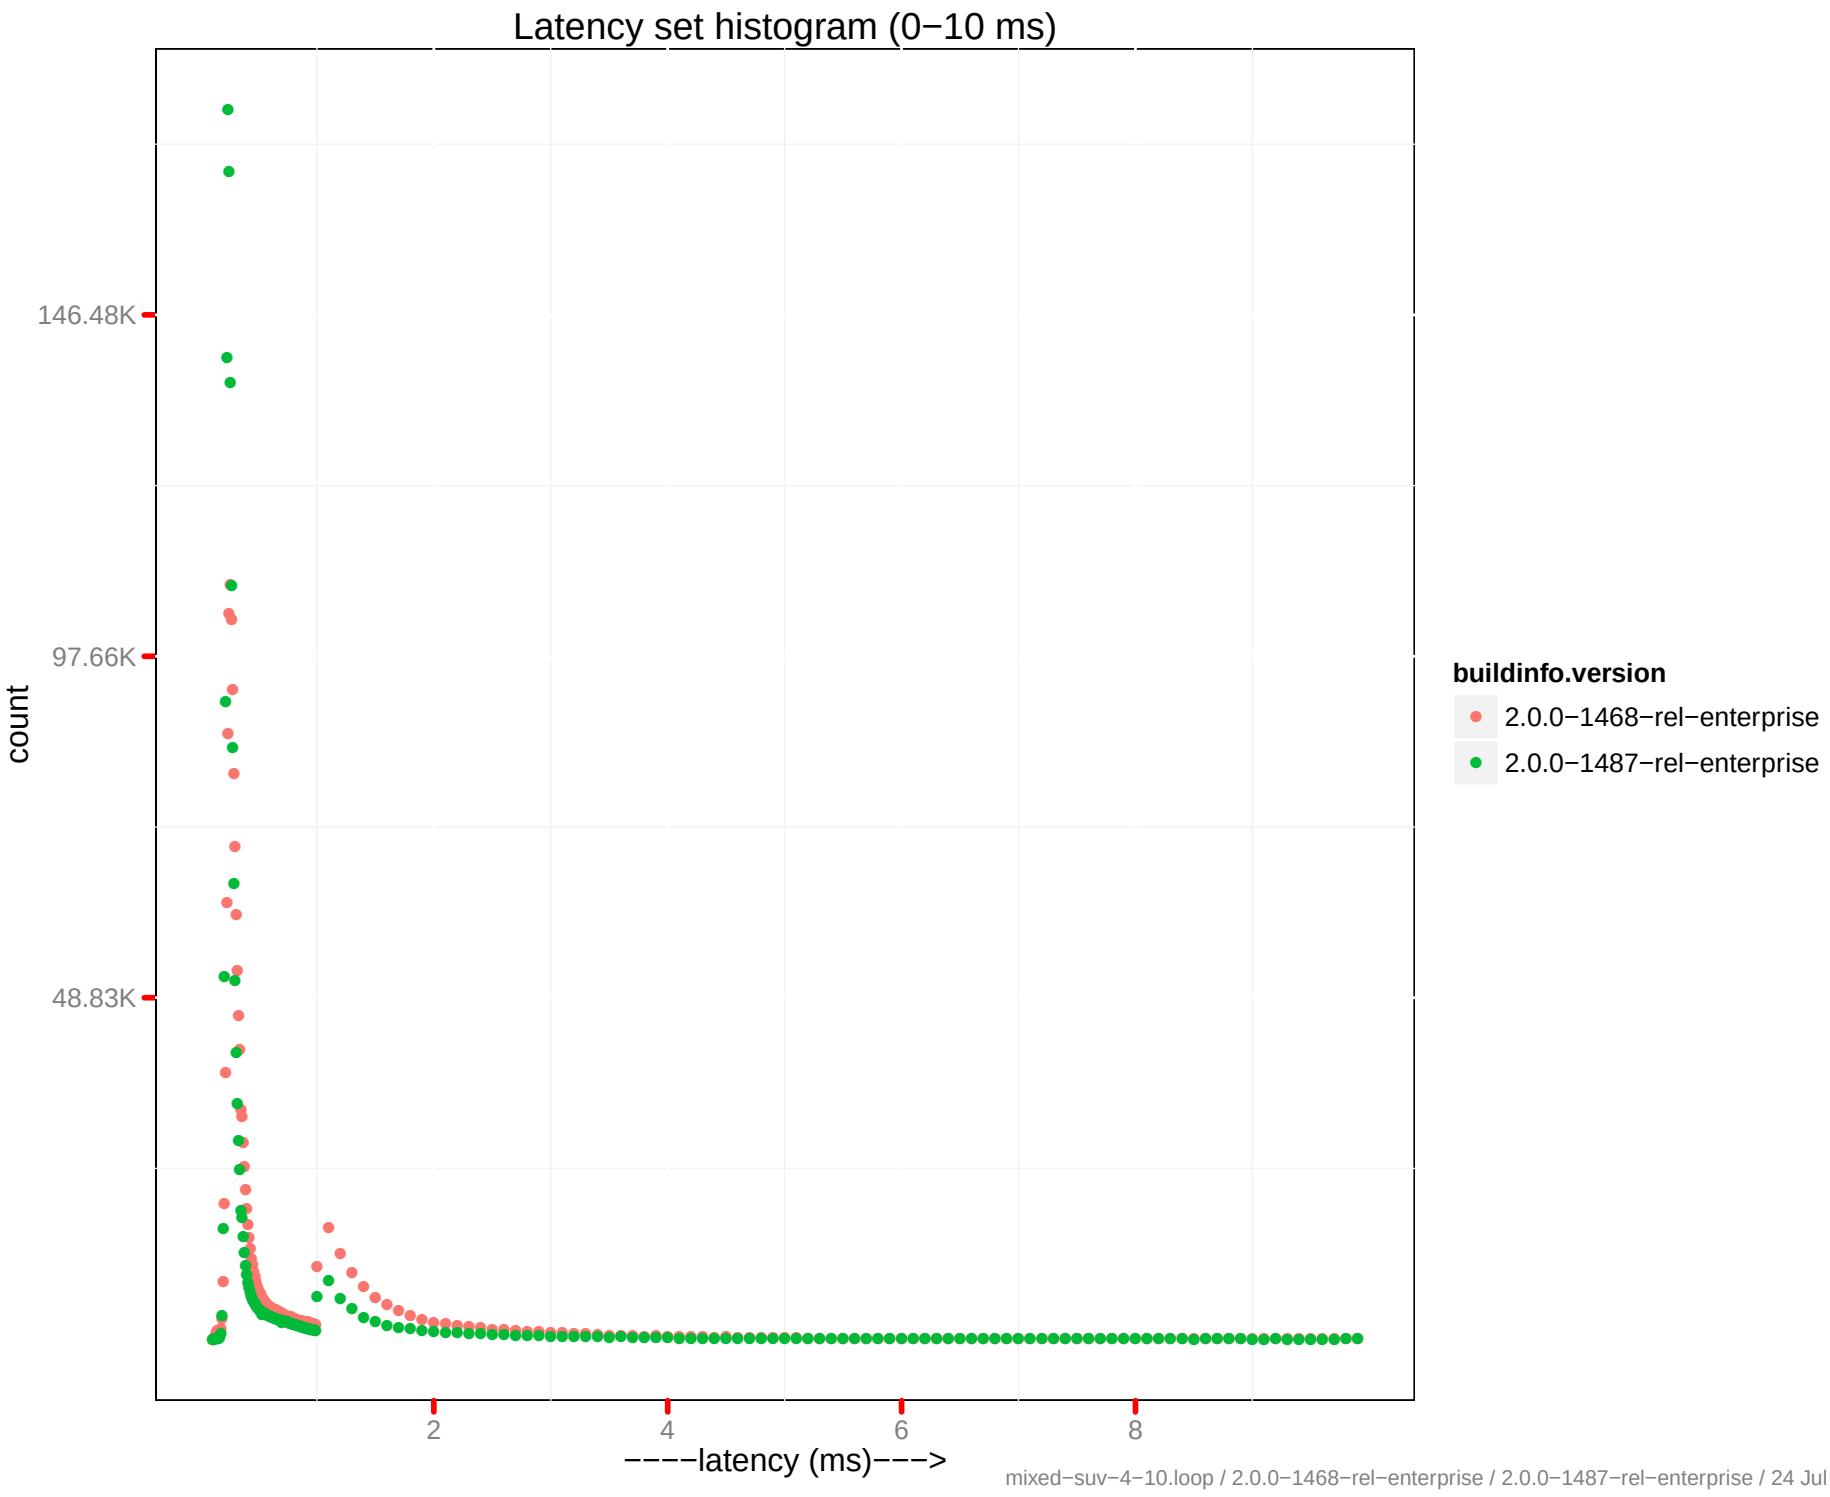
Latency set histogram (0−10 ms)
146.48K
97.66K
48.83K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

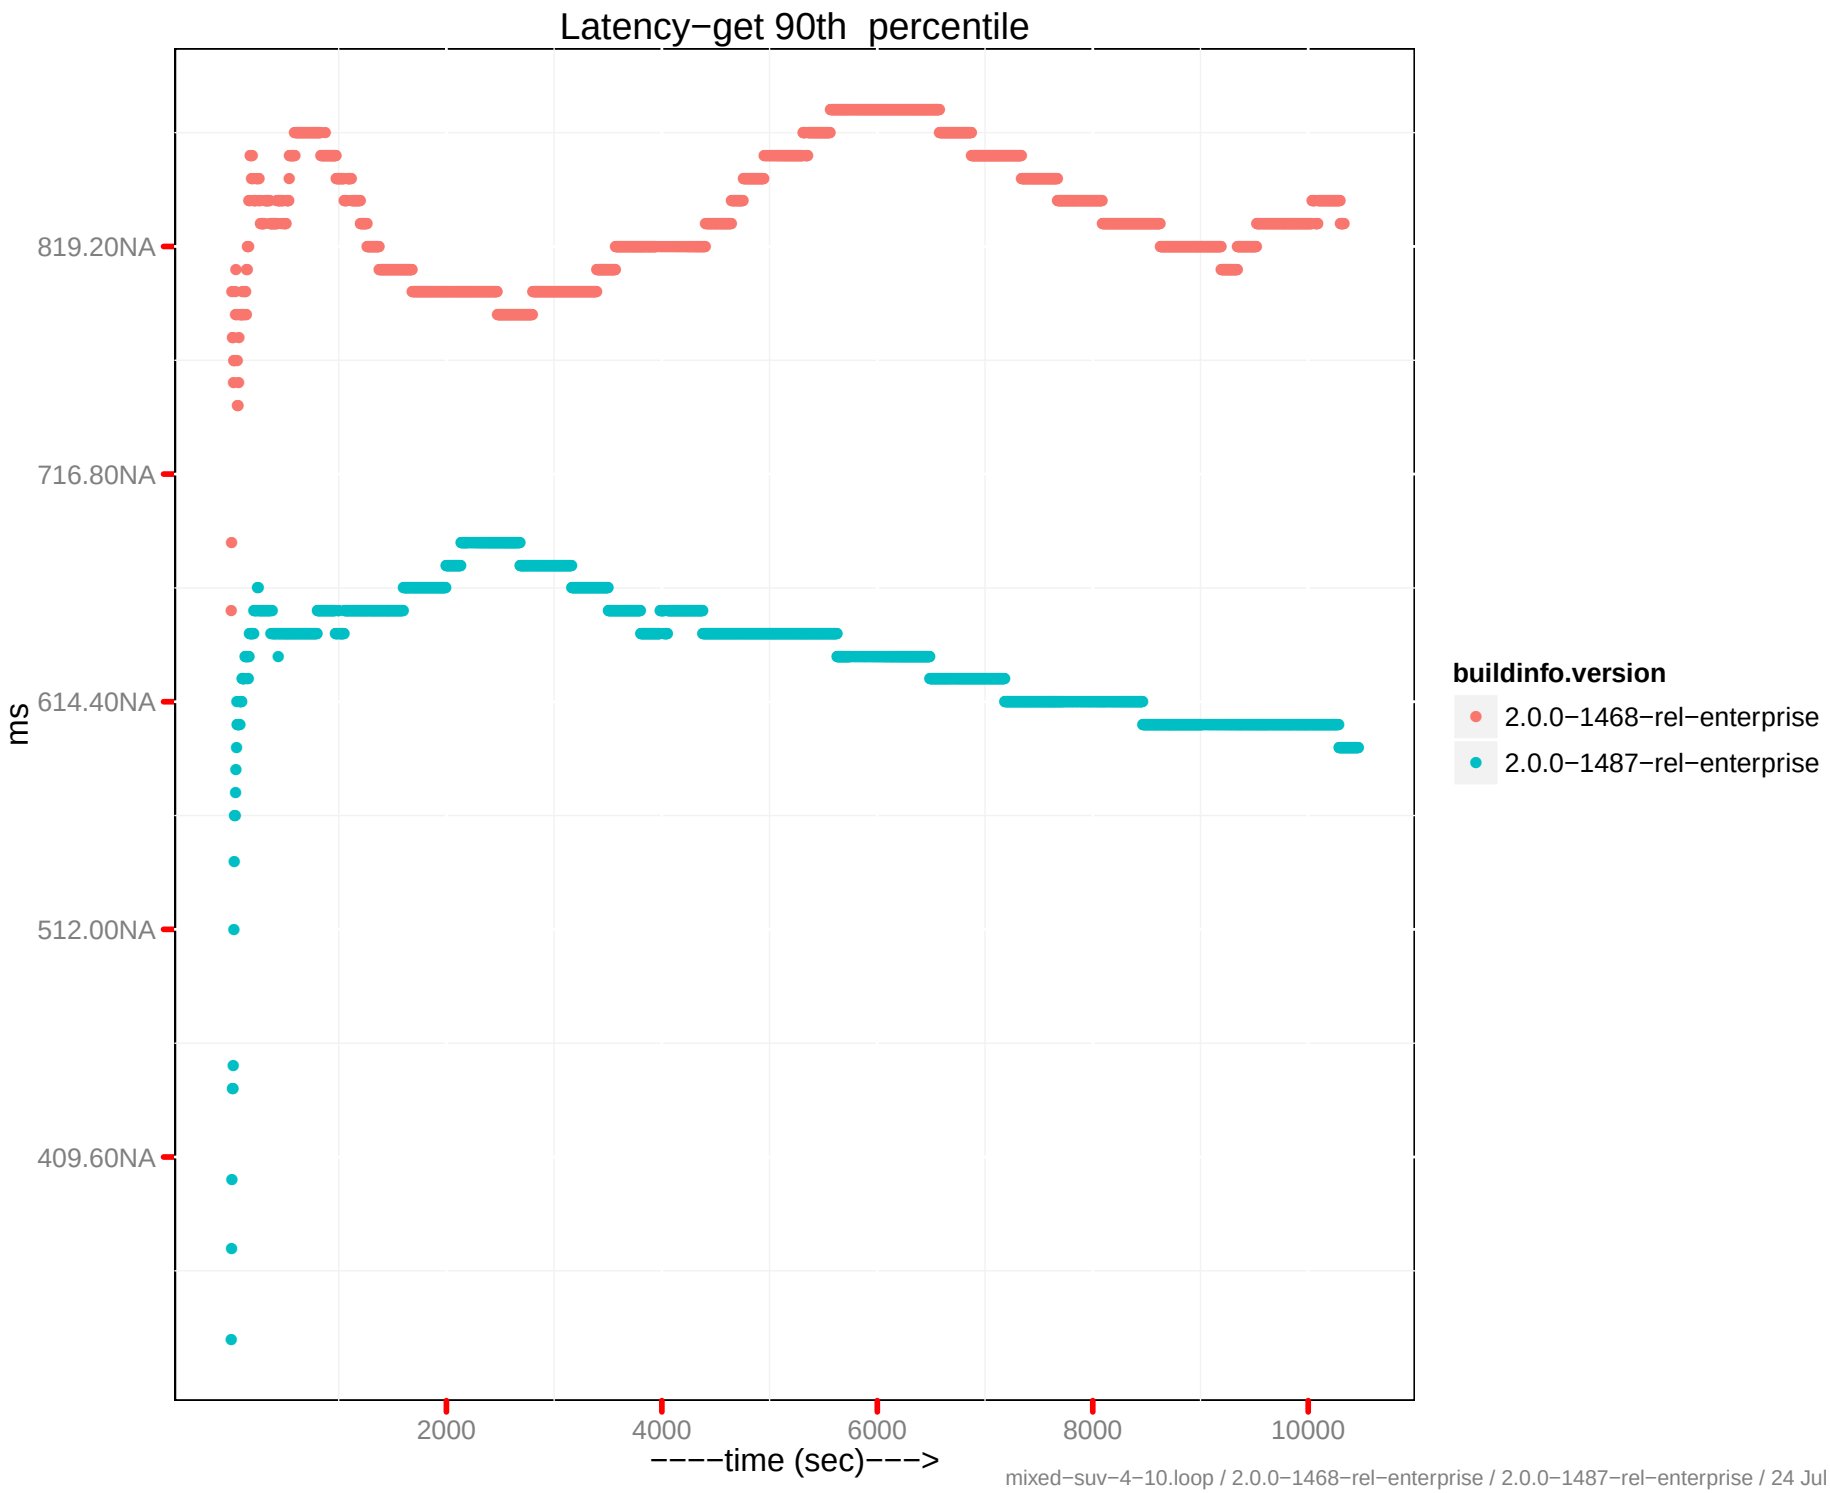
Latency−get 90th percentile
819.20NA
716.80NA
614.40NA
512.00NA
409.60NA
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

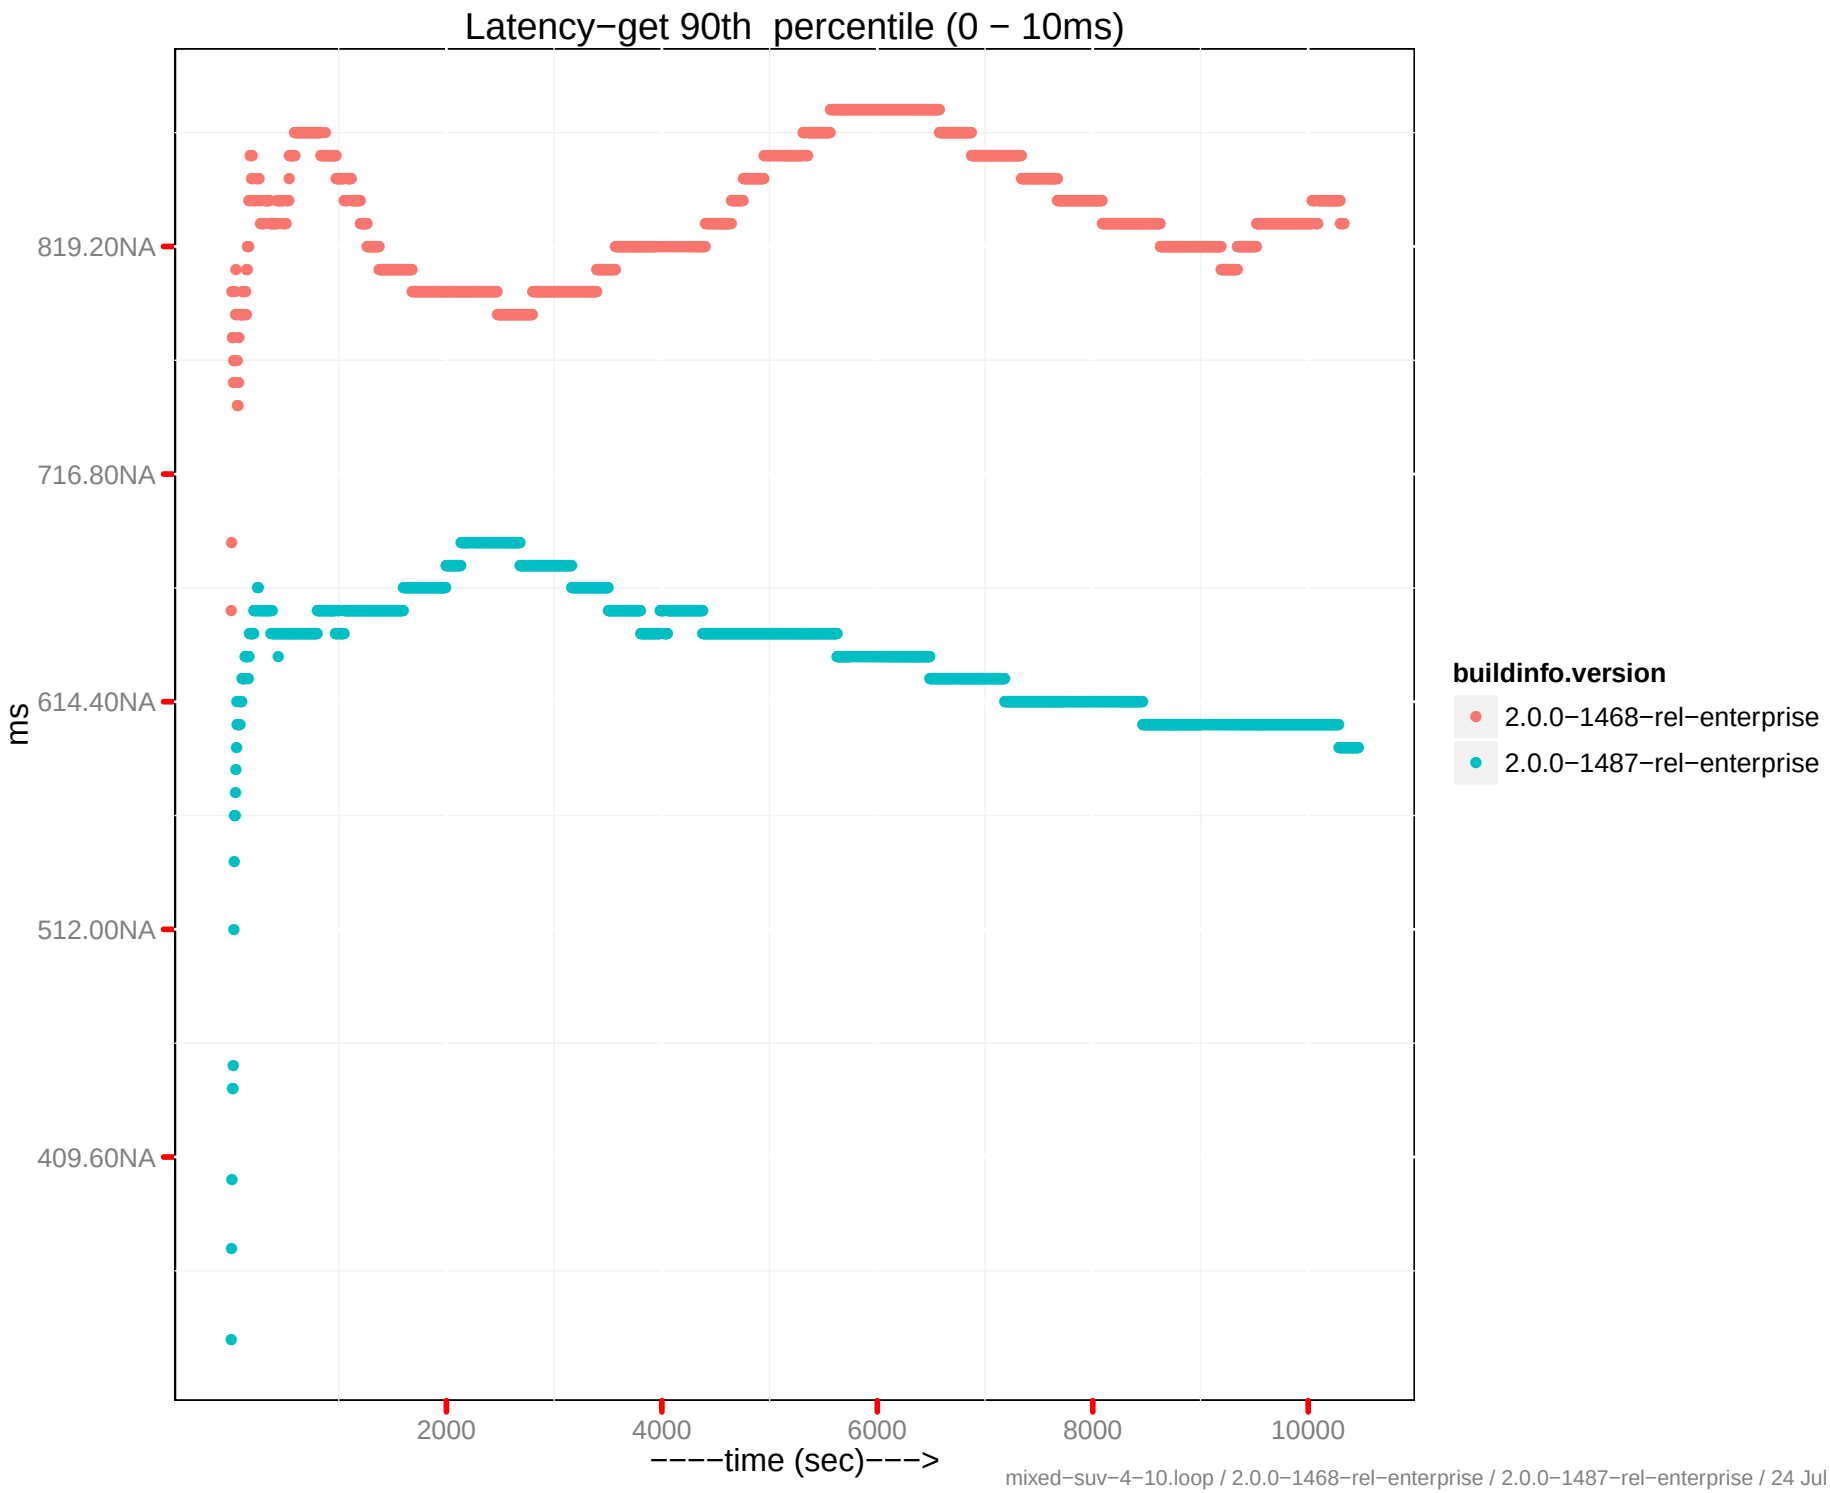
Latency−get 90th percentile (0 − 10ms)
819.20NA
716.80NA
614.40NA
512.00NA
409.60NA
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

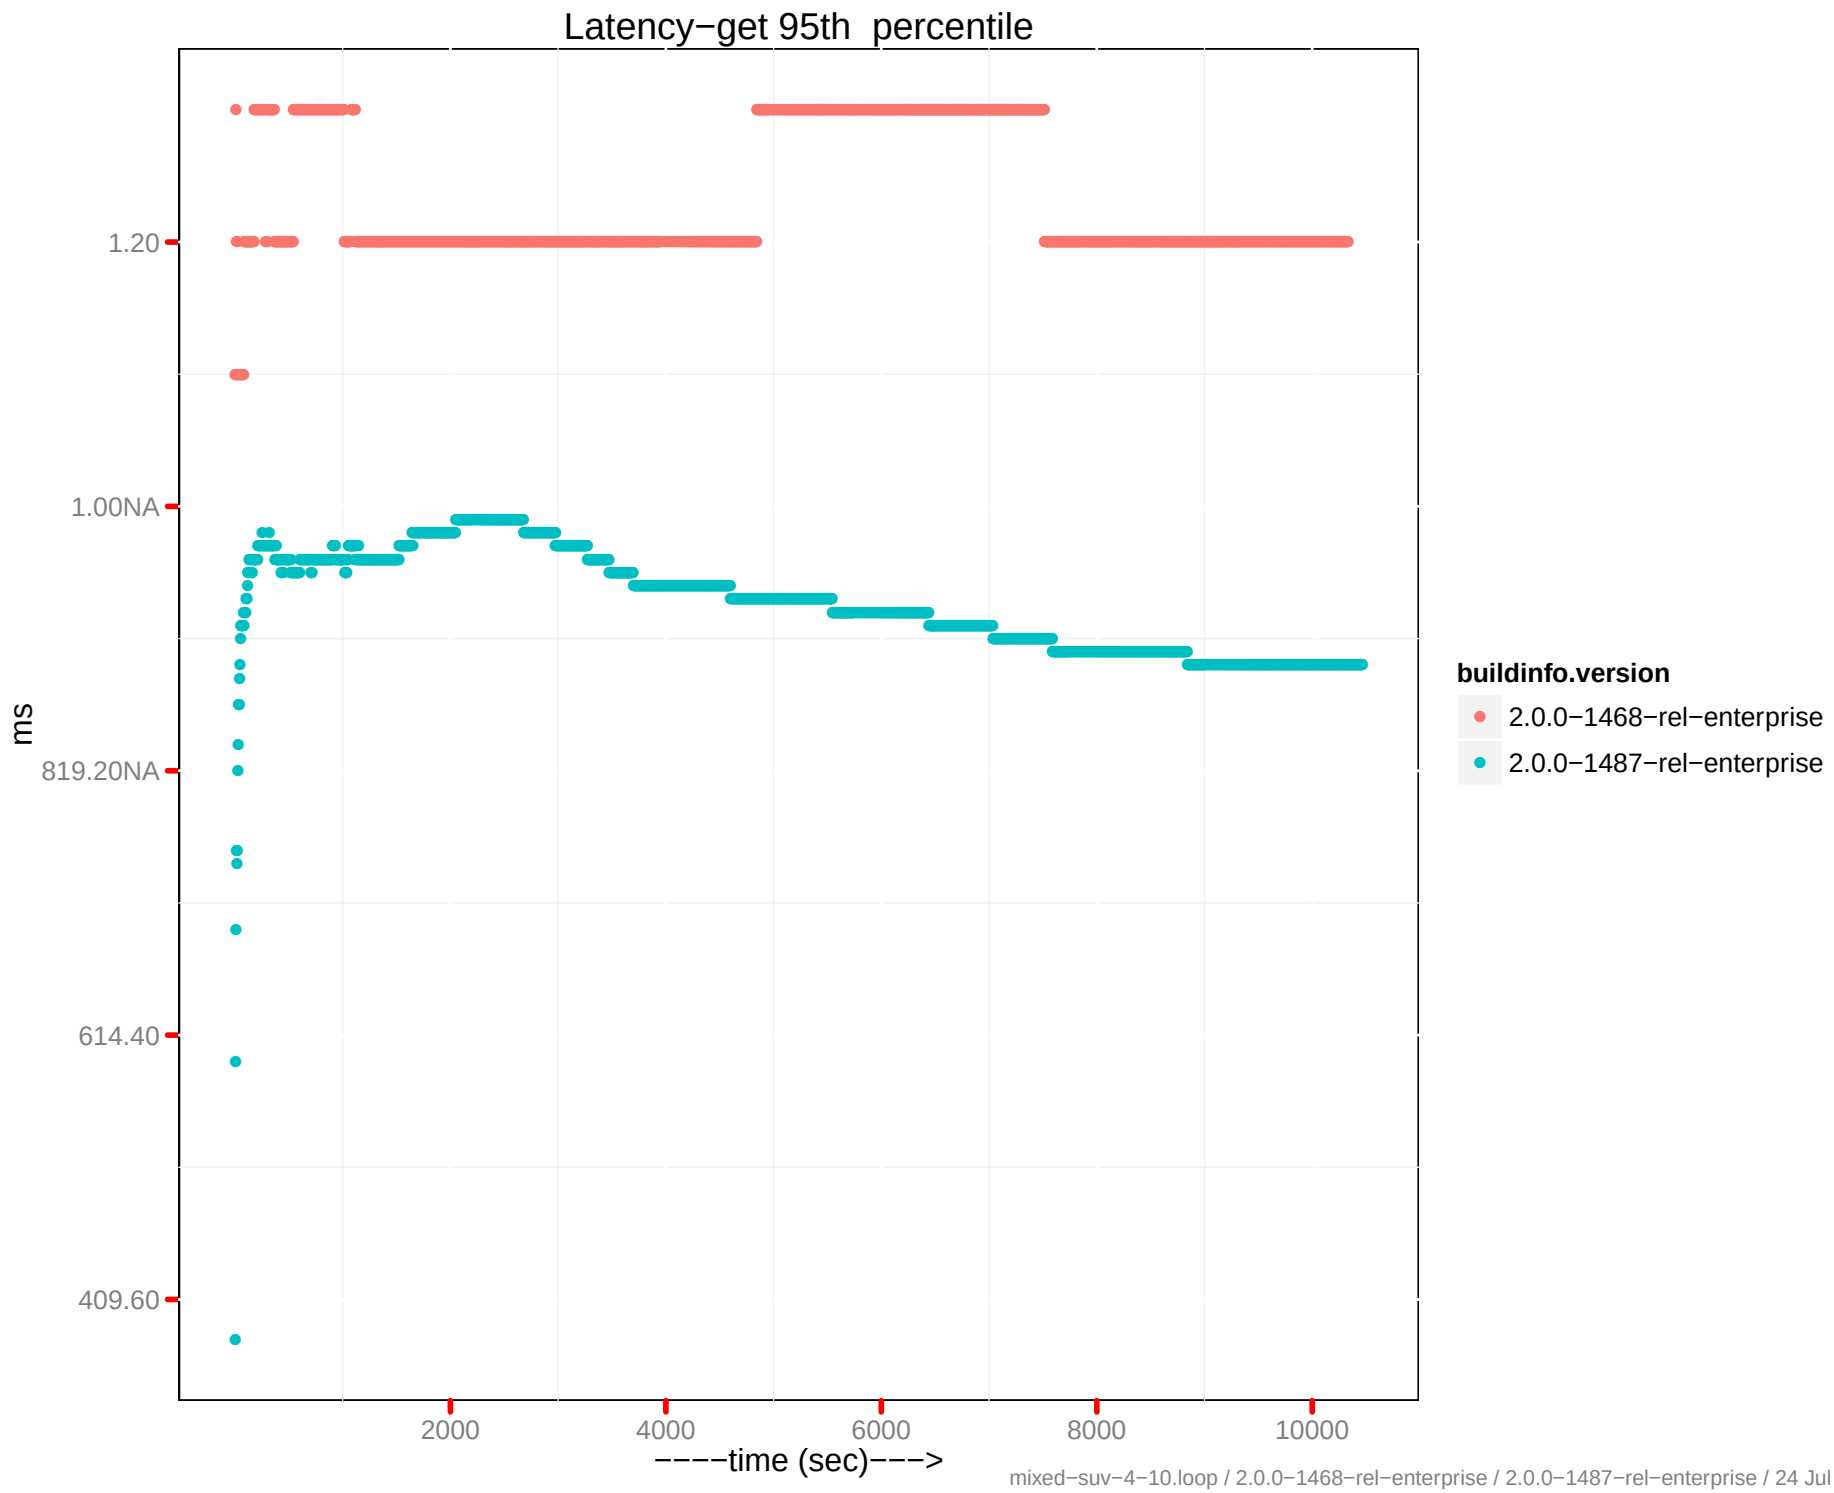
Latency−get 95th percentile
1.20
1.00NA
819.20NA
614.40
409.60
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

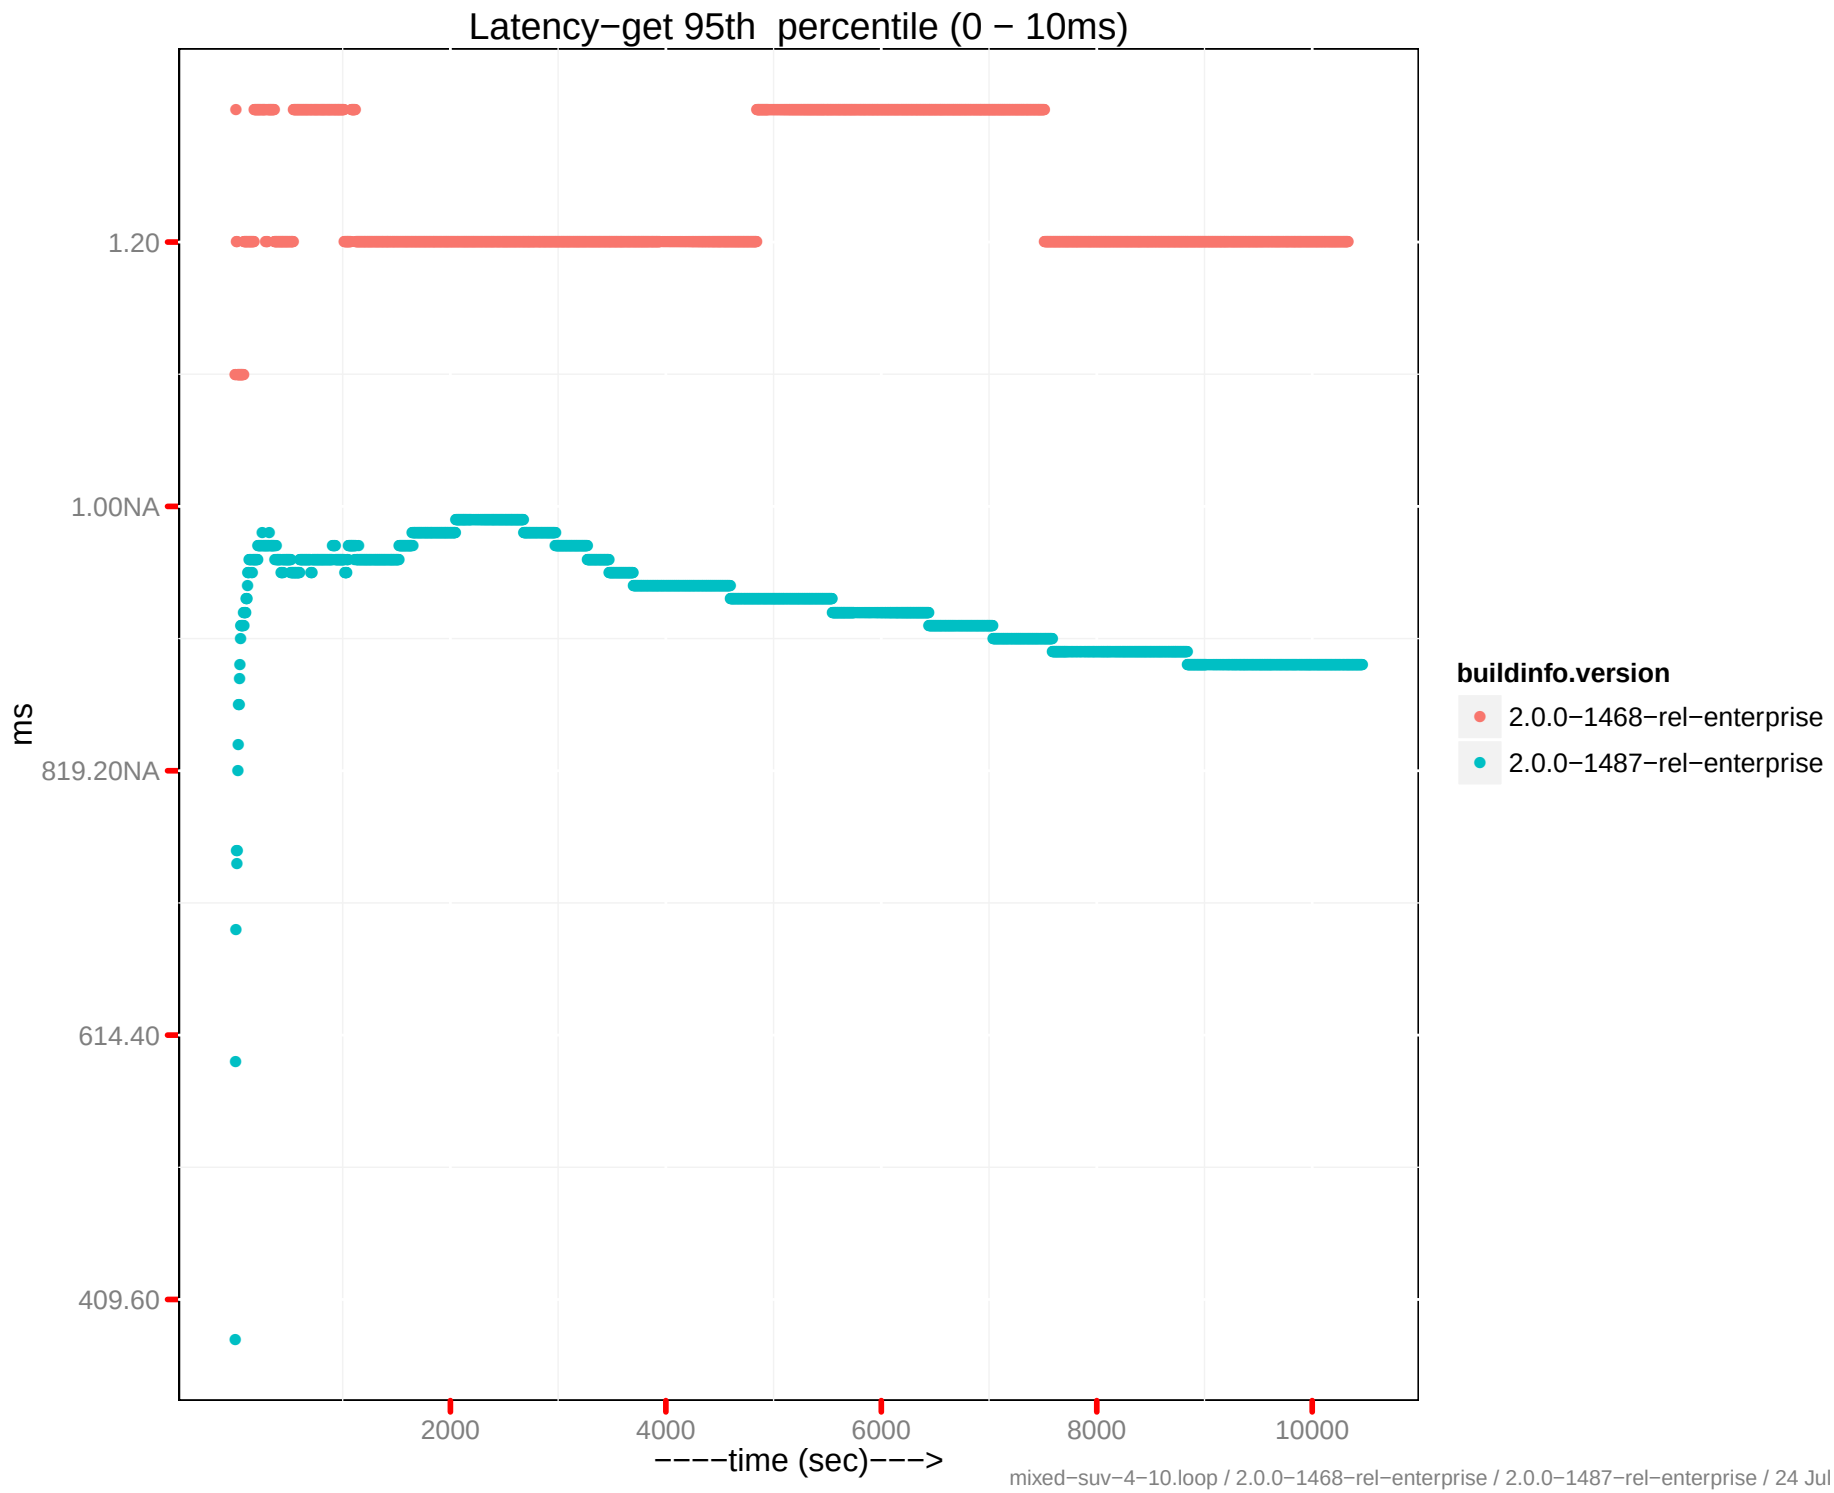
Latency−get 95th percentile (0 − 10ms)
1.20
1.00NA
819.20NA
614.40
409.60
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

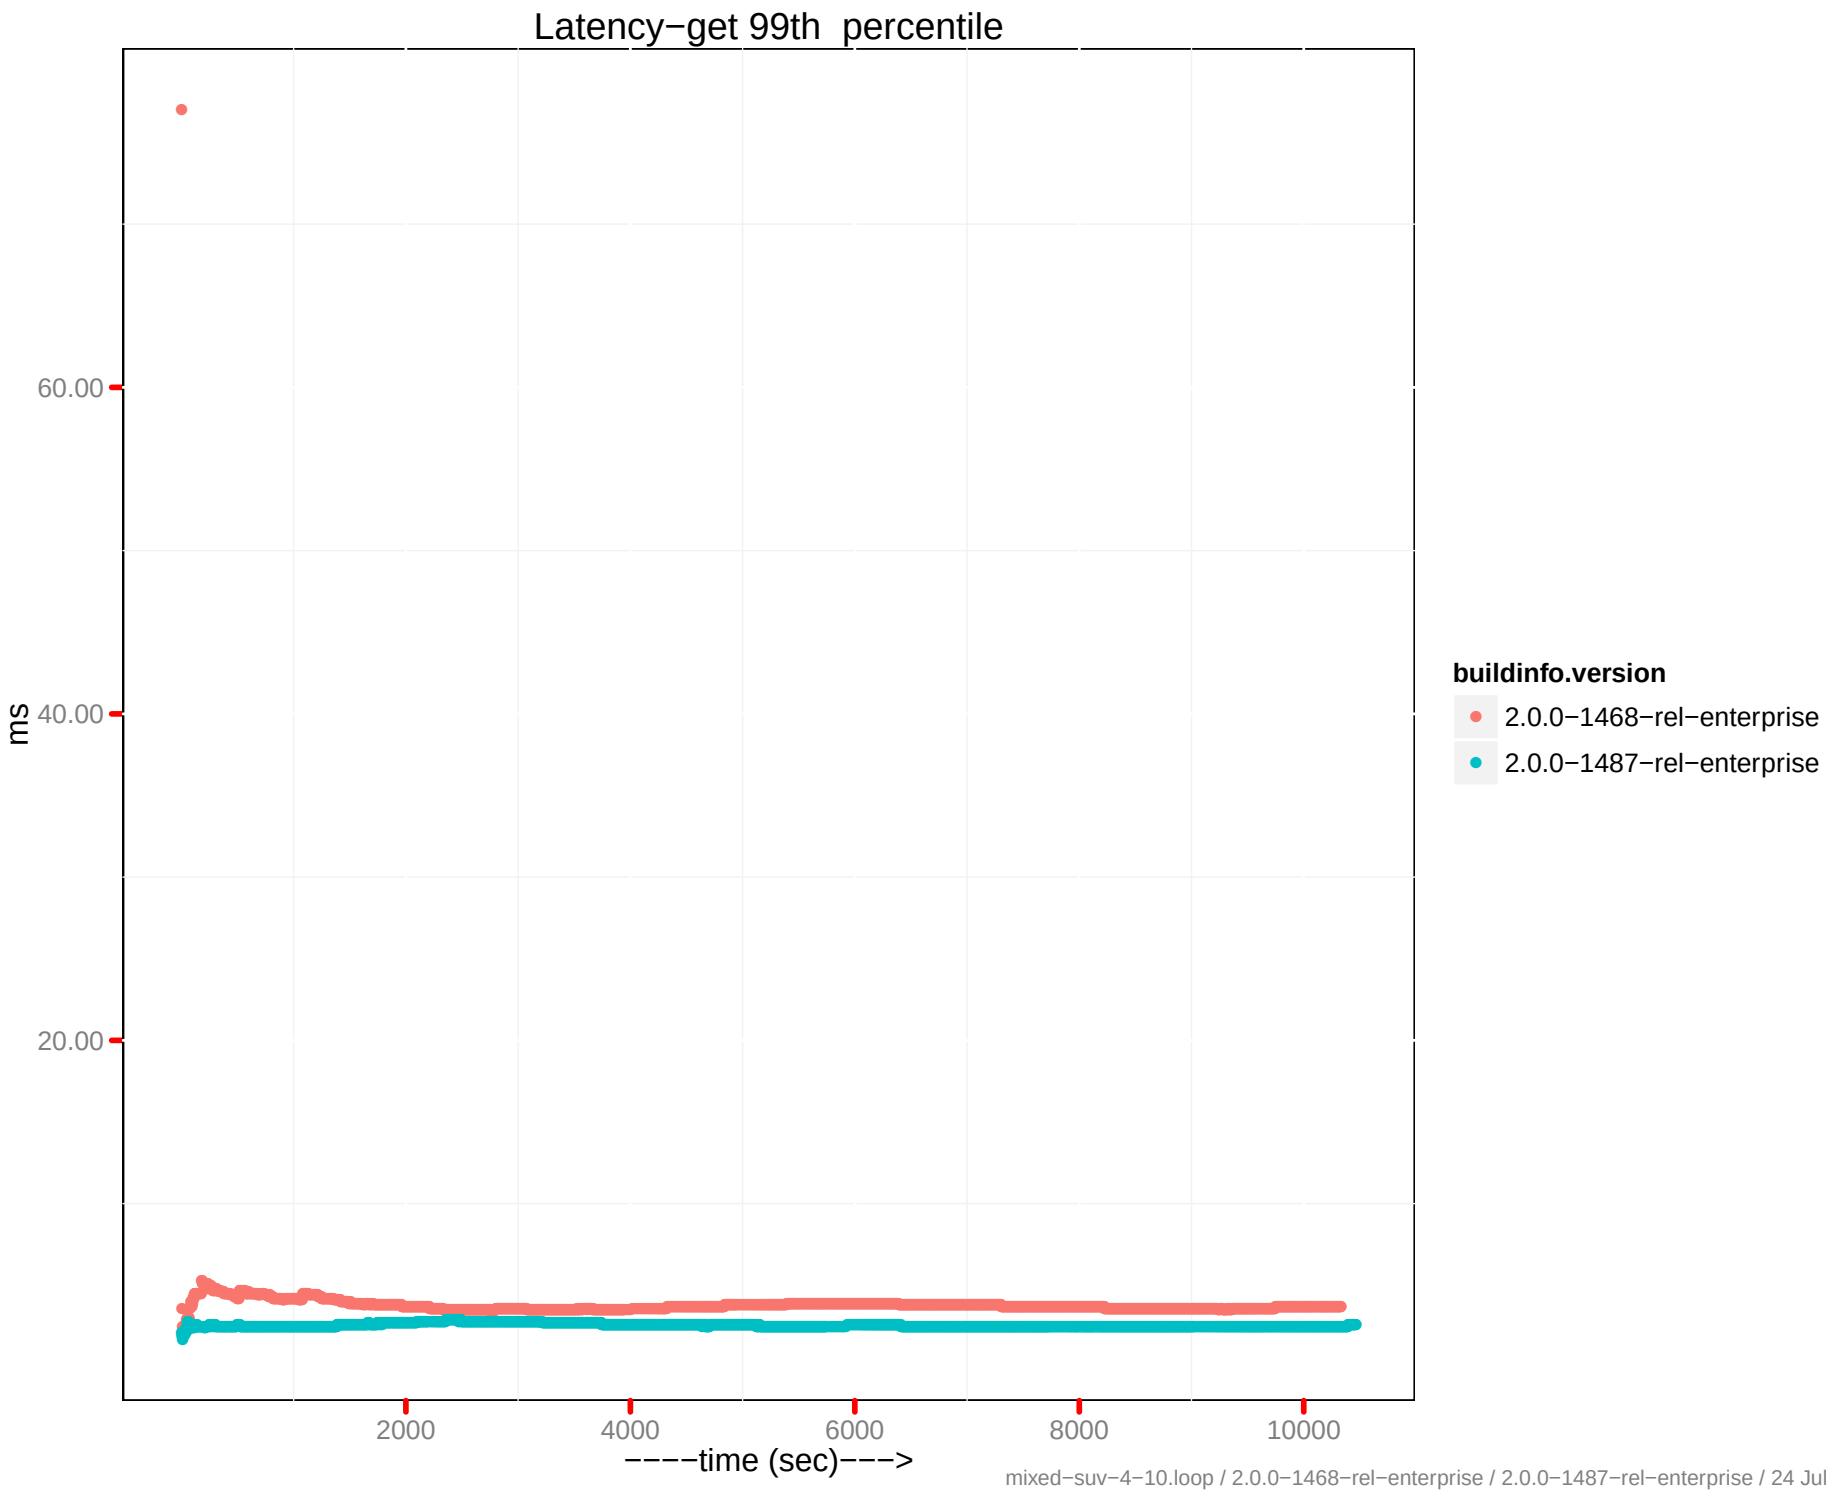
Latency−get 99th percentile
60.00
40.00
20.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

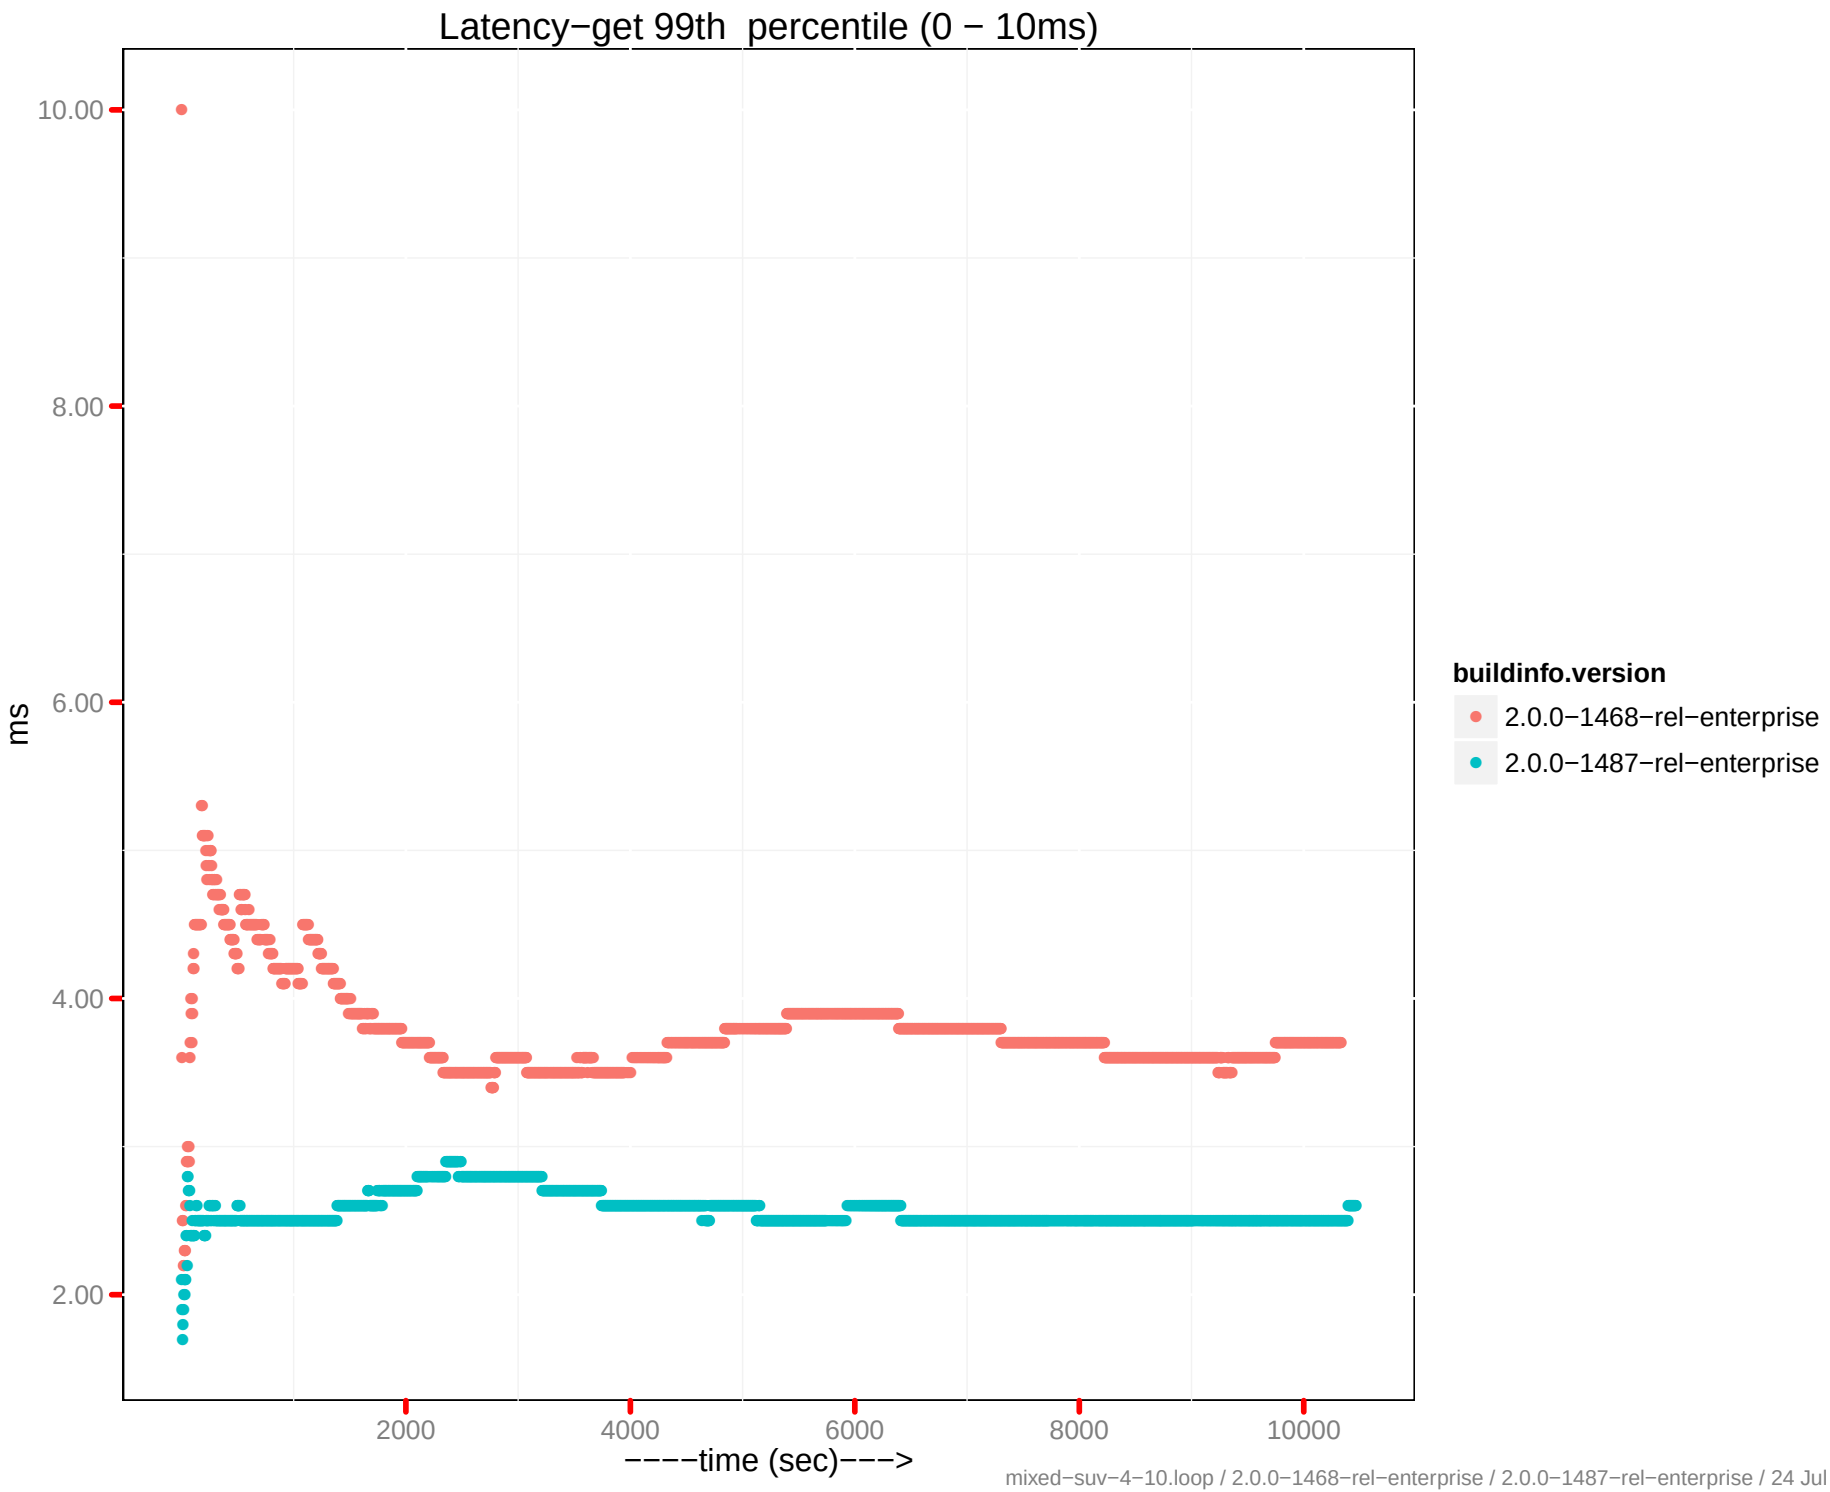
Latency−get 99th percentile (0 − 10ms)
10.00
8.00
6.00
4.00
2.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

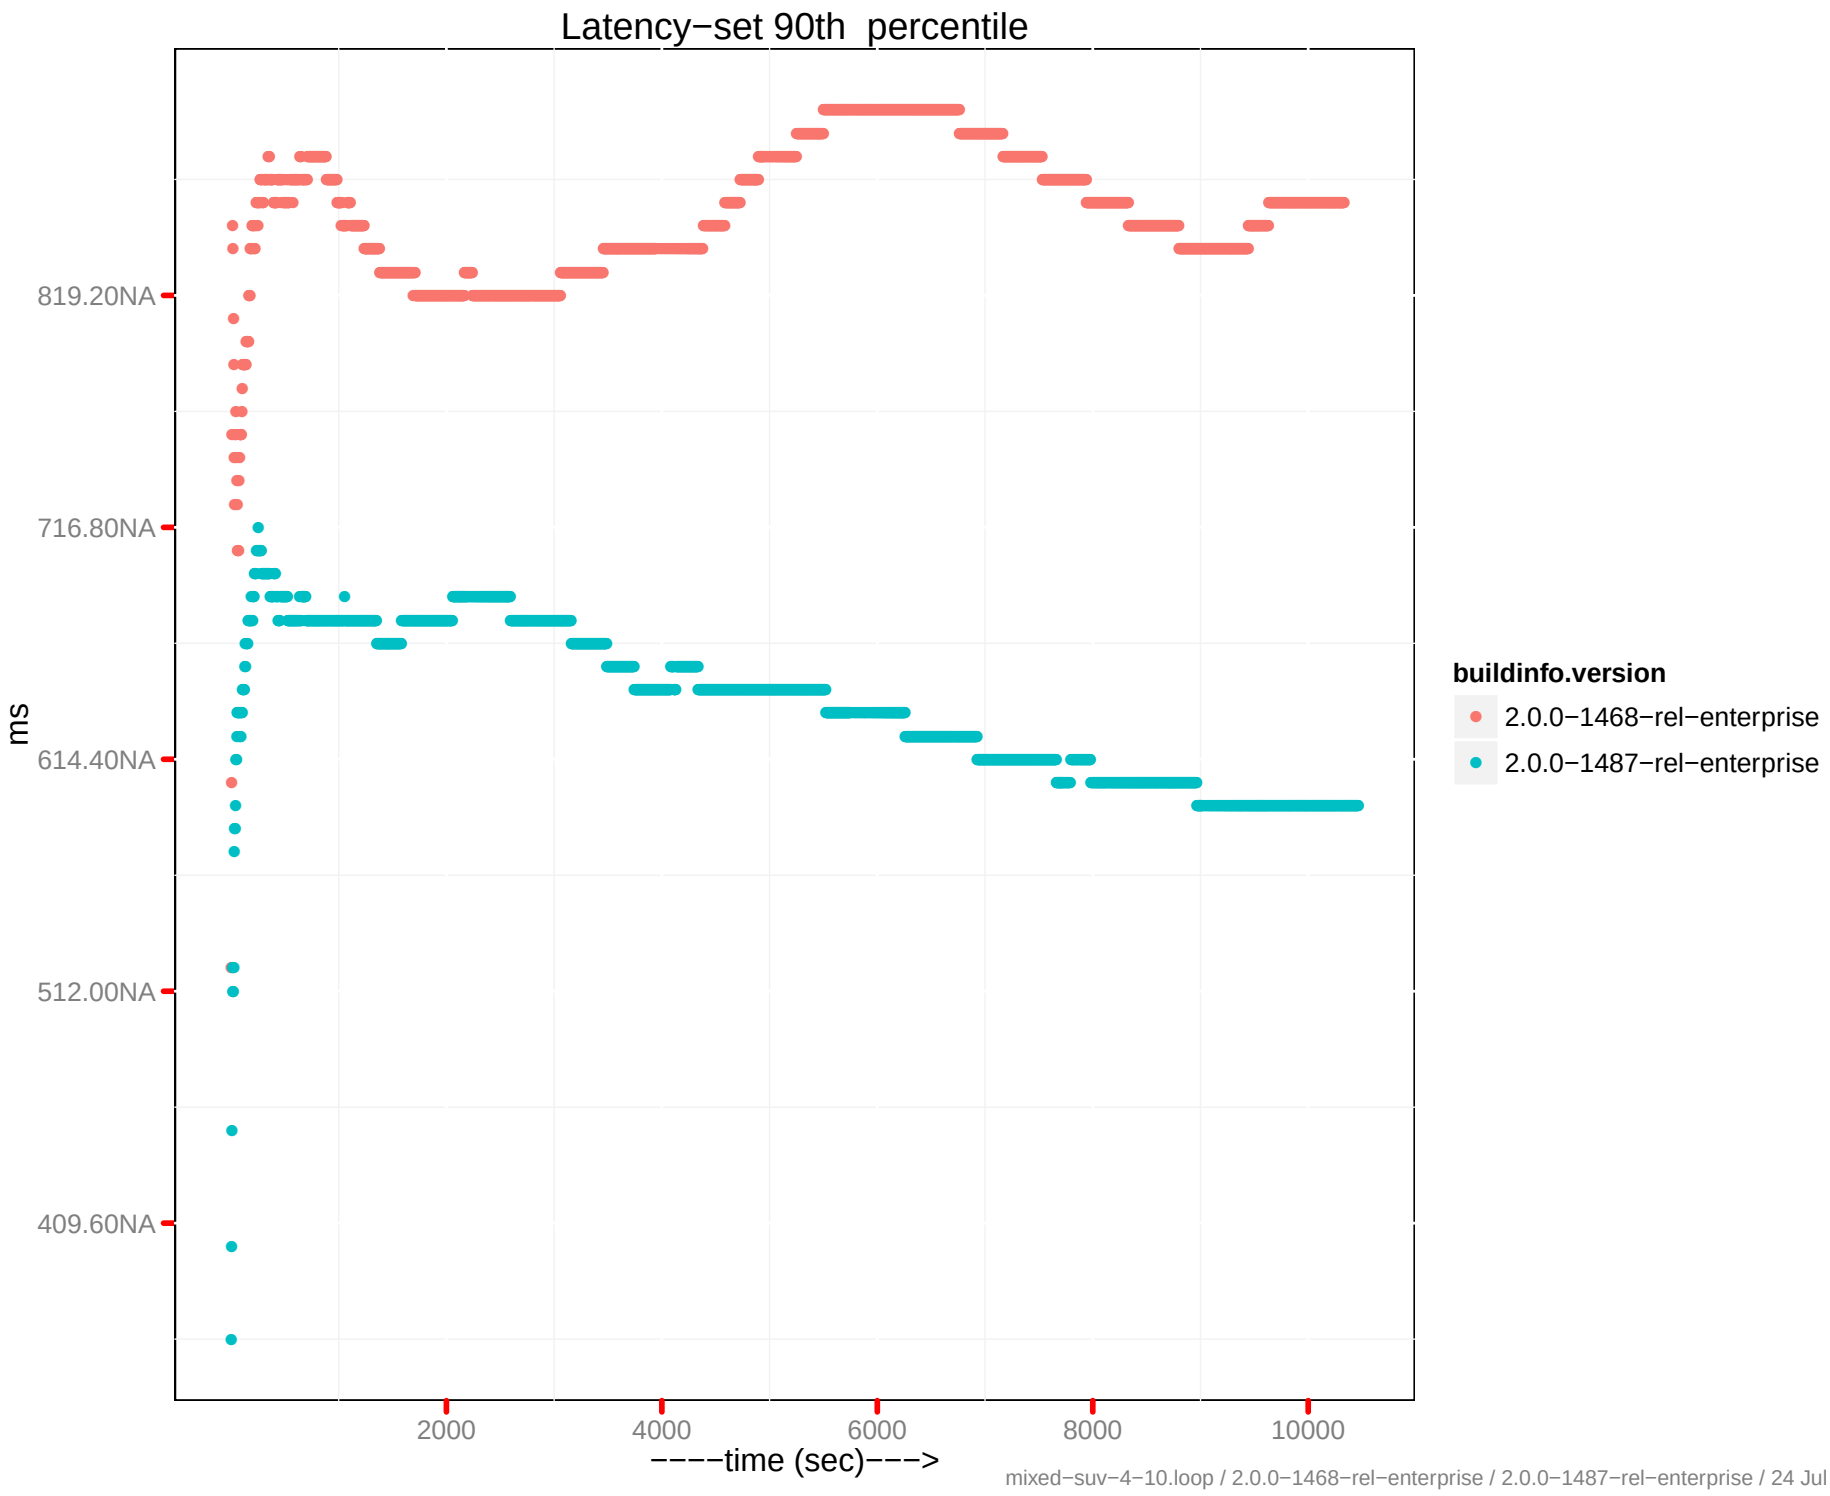
Latency−set 90th percentile
819.20NA
716.80NA
614.40NA
512.00NA
409.60NA
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

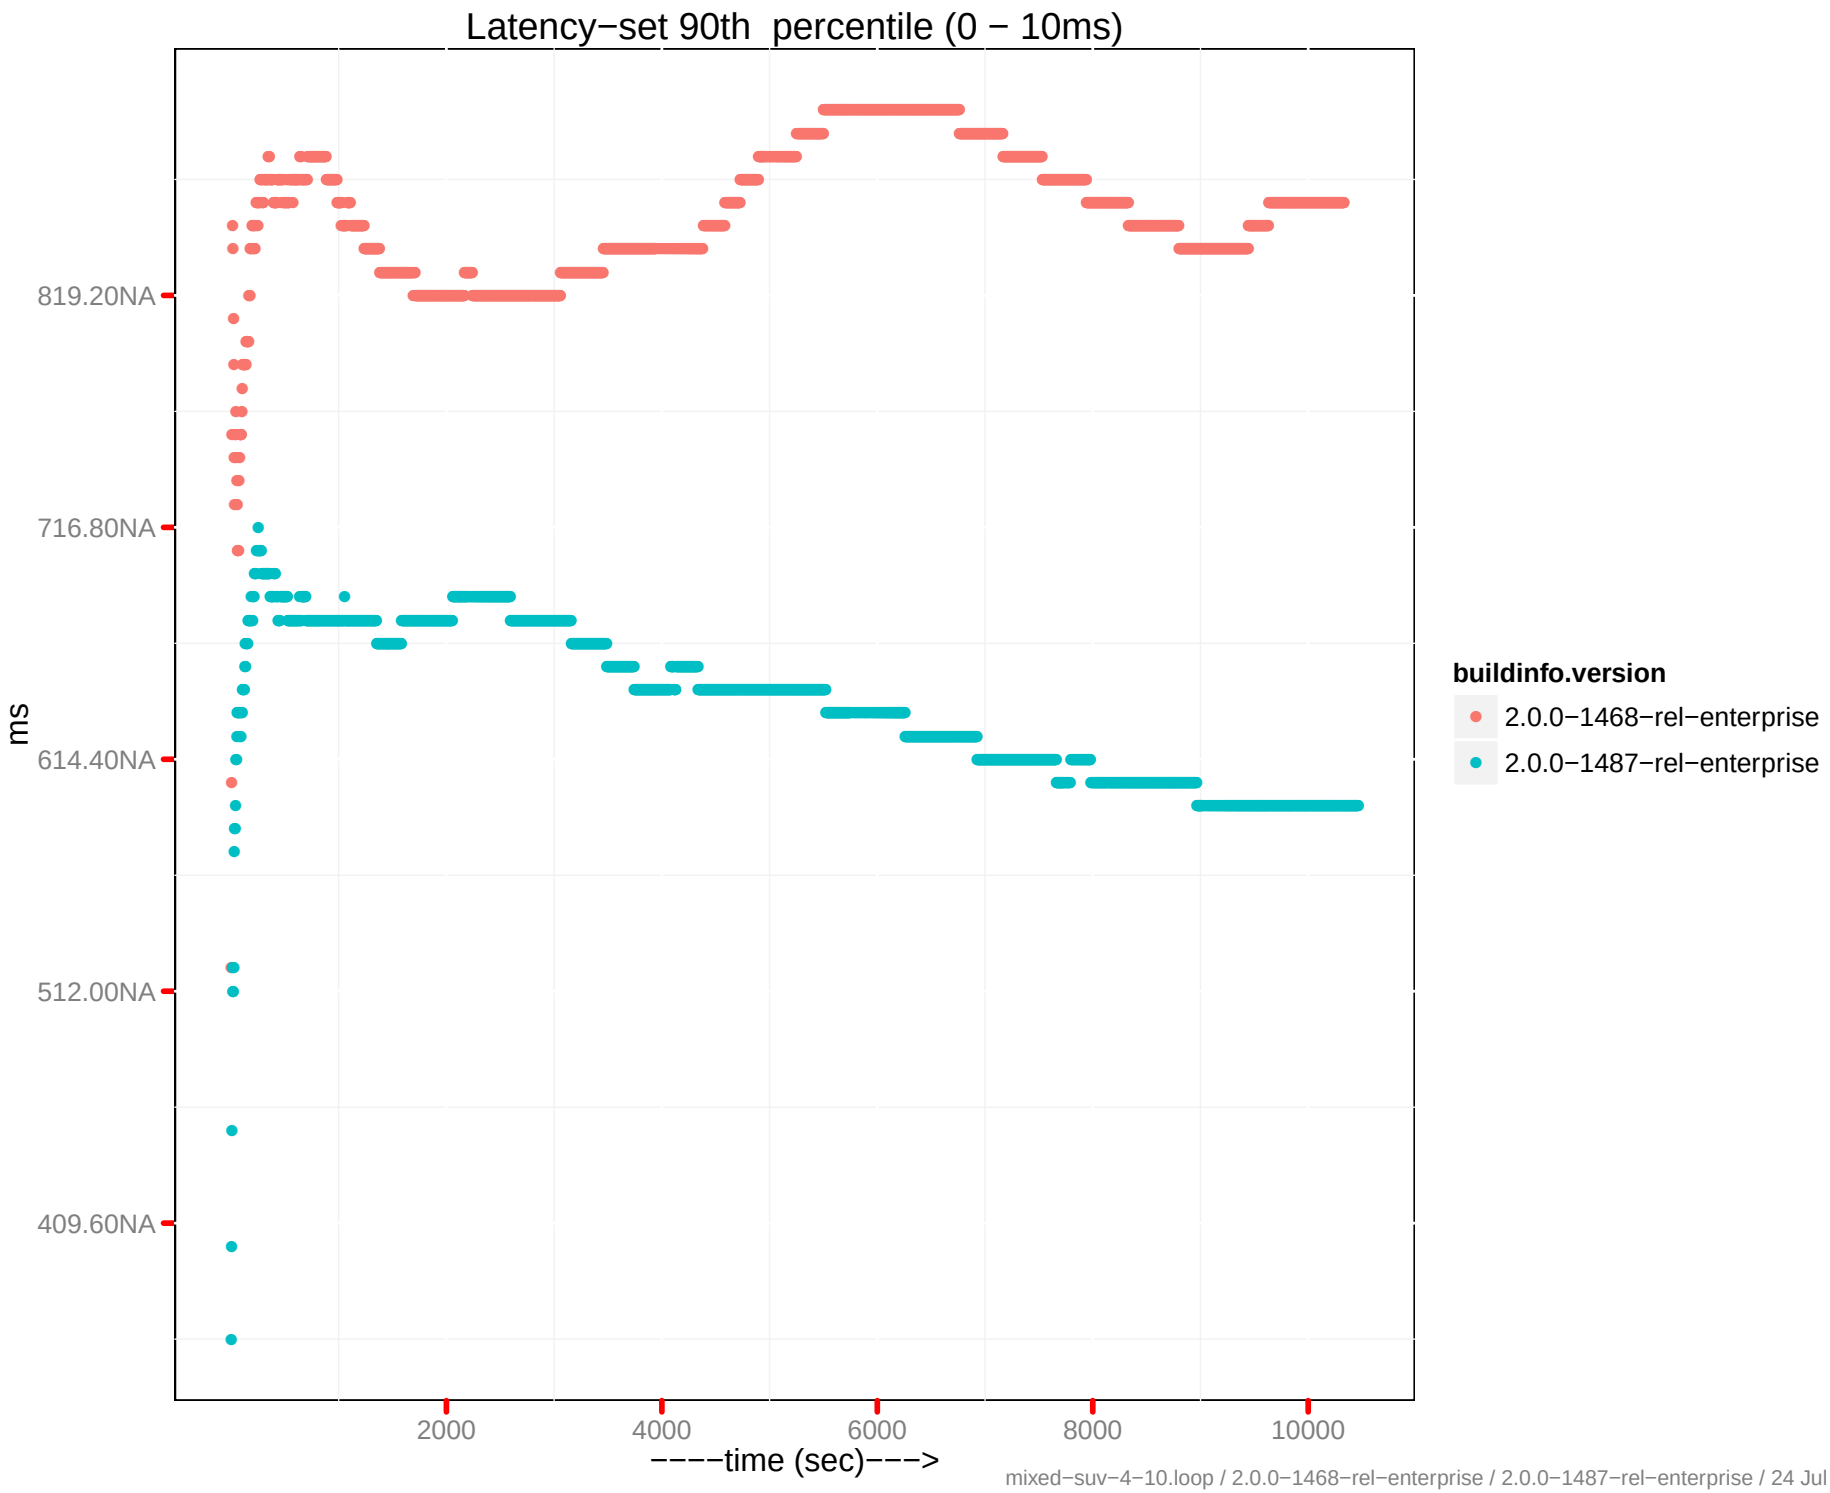
Latency−set 90th percentile (0 − 10ms)
819.20NA
716.80NA
614.40NA
512.00NA
409.60NA
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

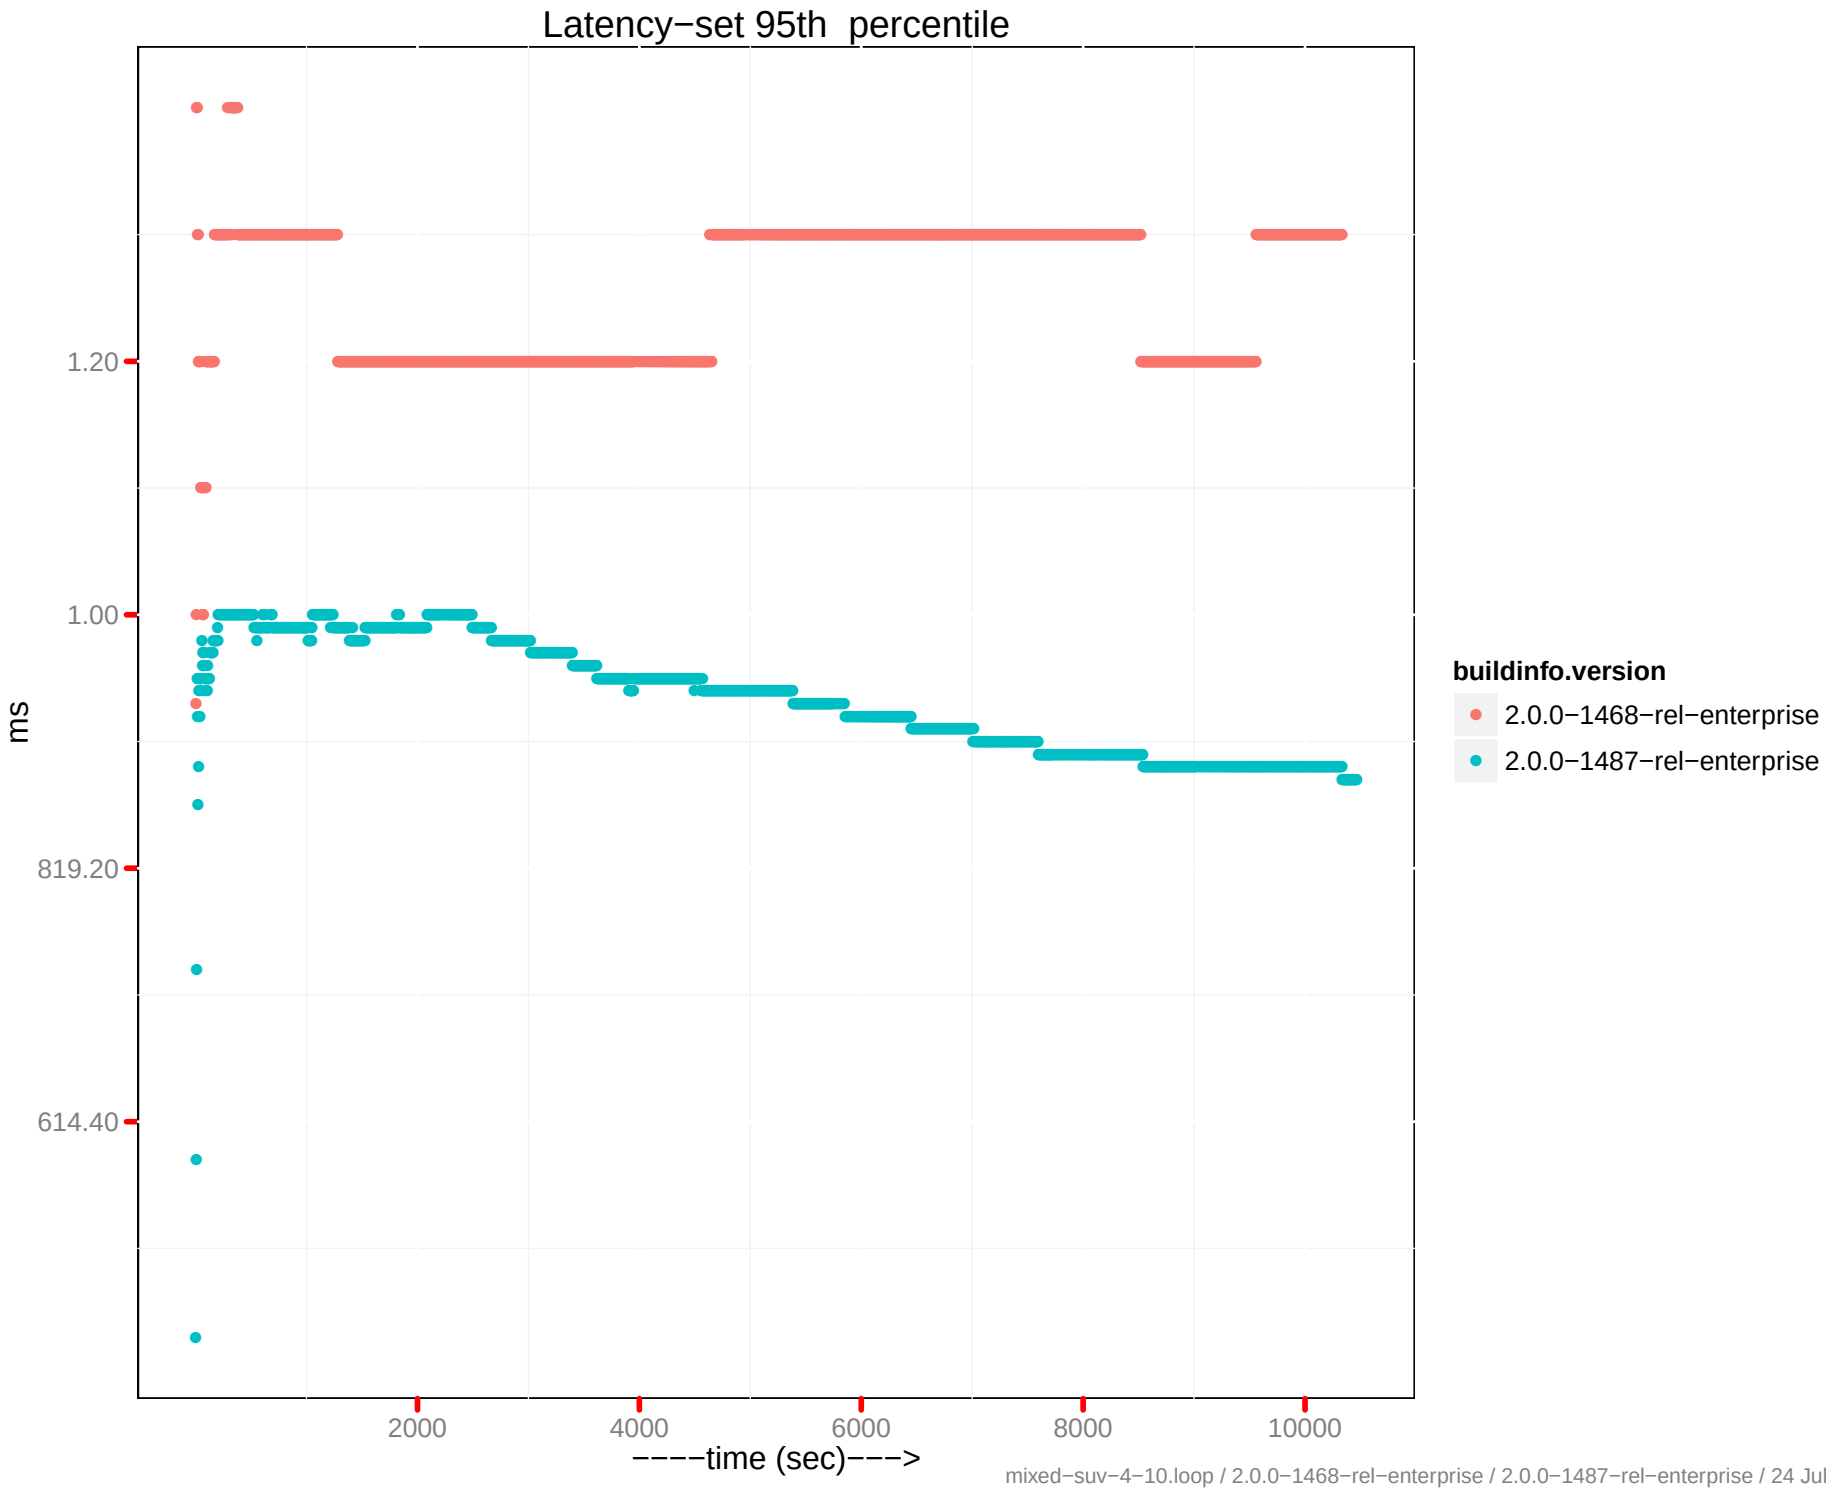
Latency−set 95th percentile
1.20
1.00
819.20
614.40
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

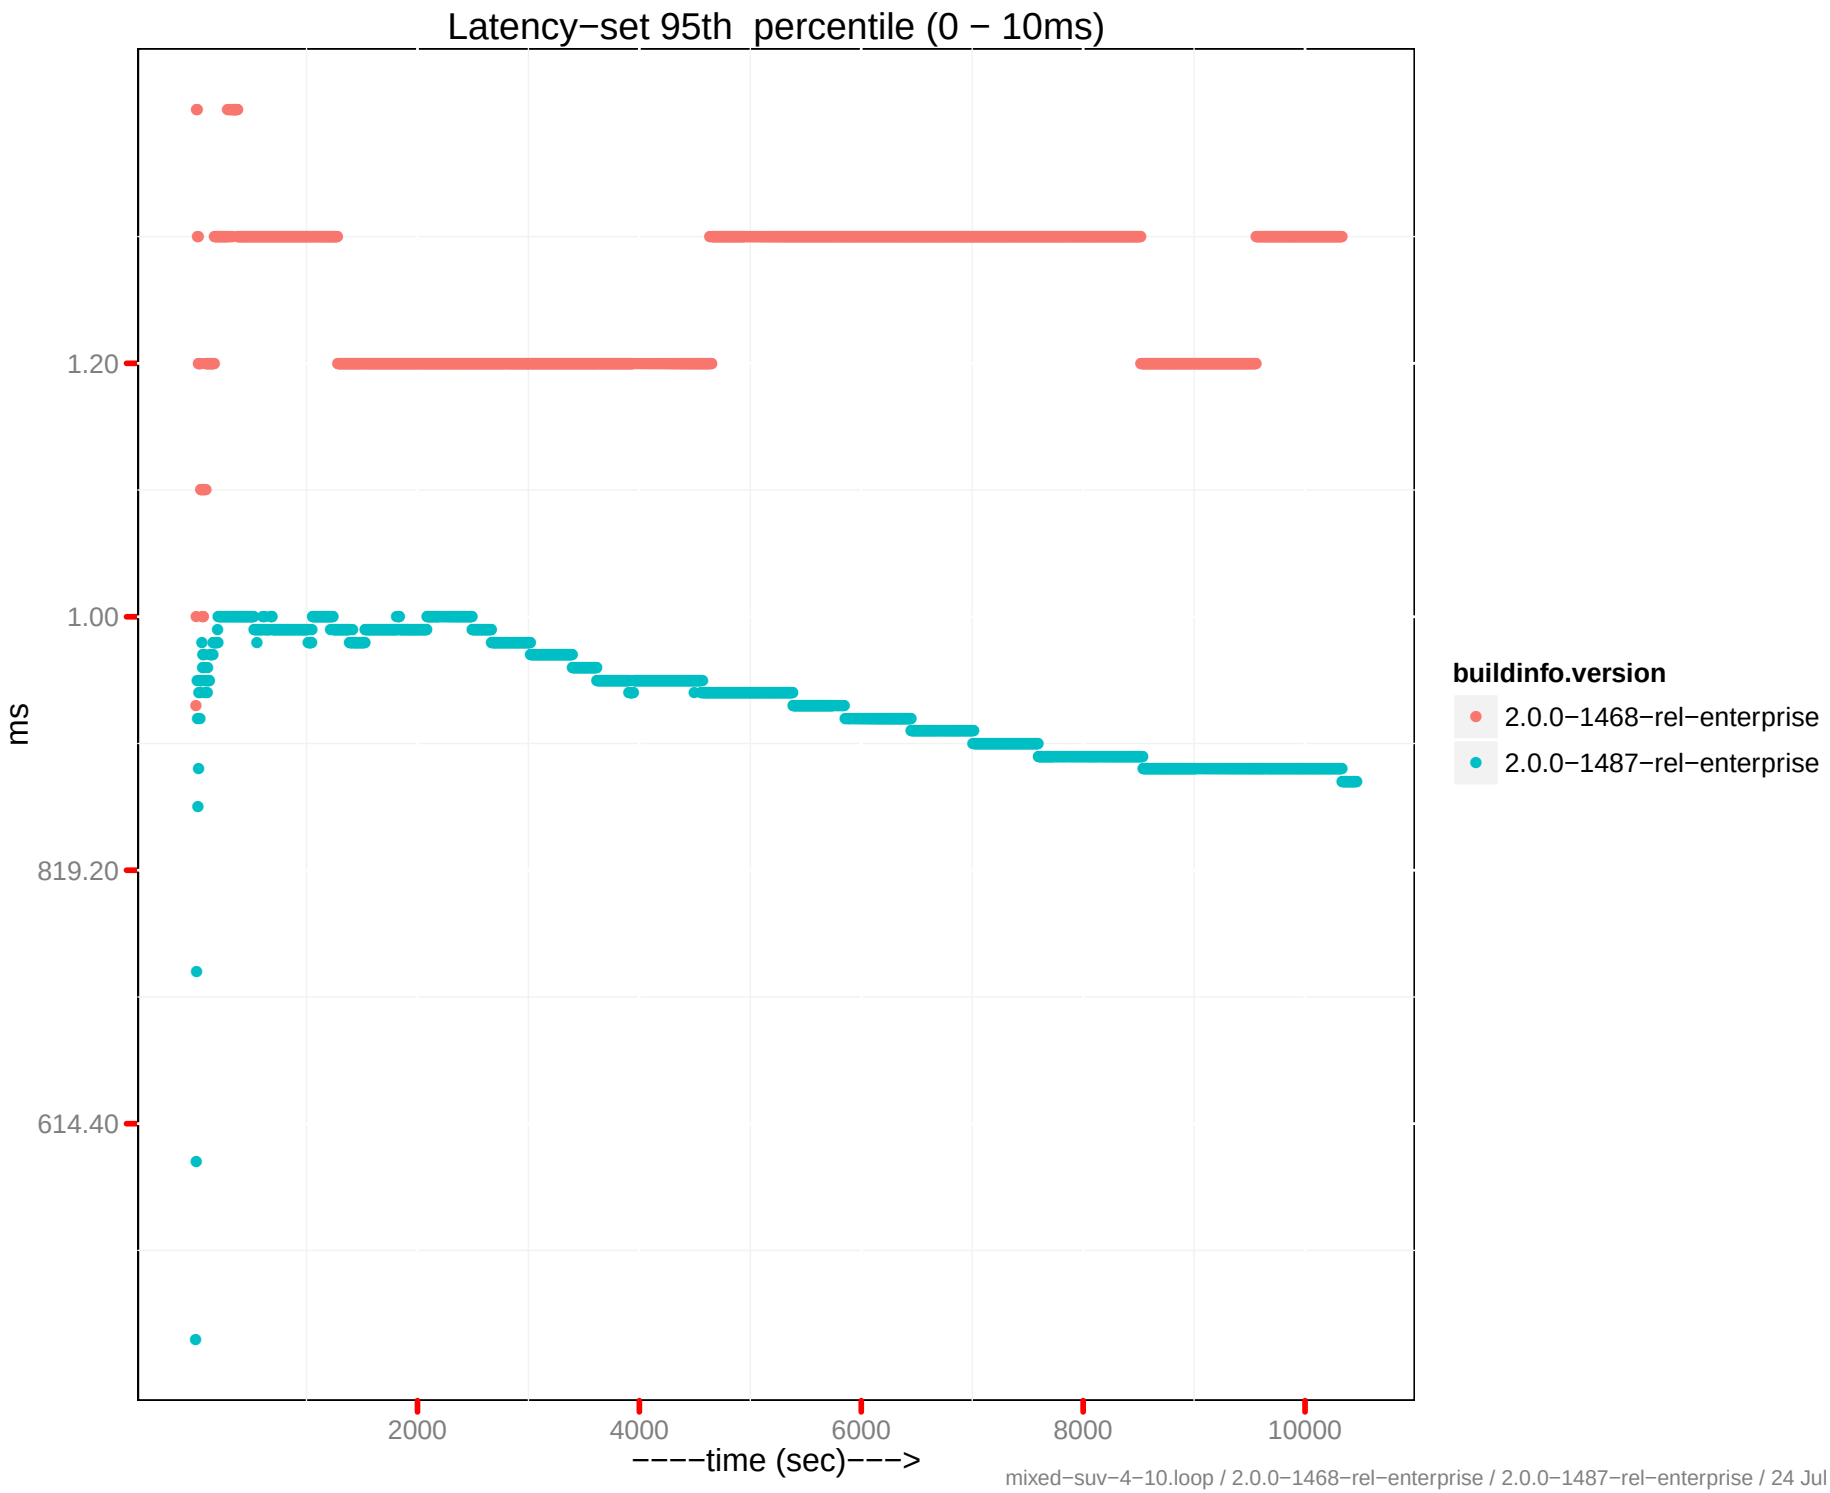
Latency−set 95th percentile (0 − 10ms)
1.20
1.00
819.20
614.40
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

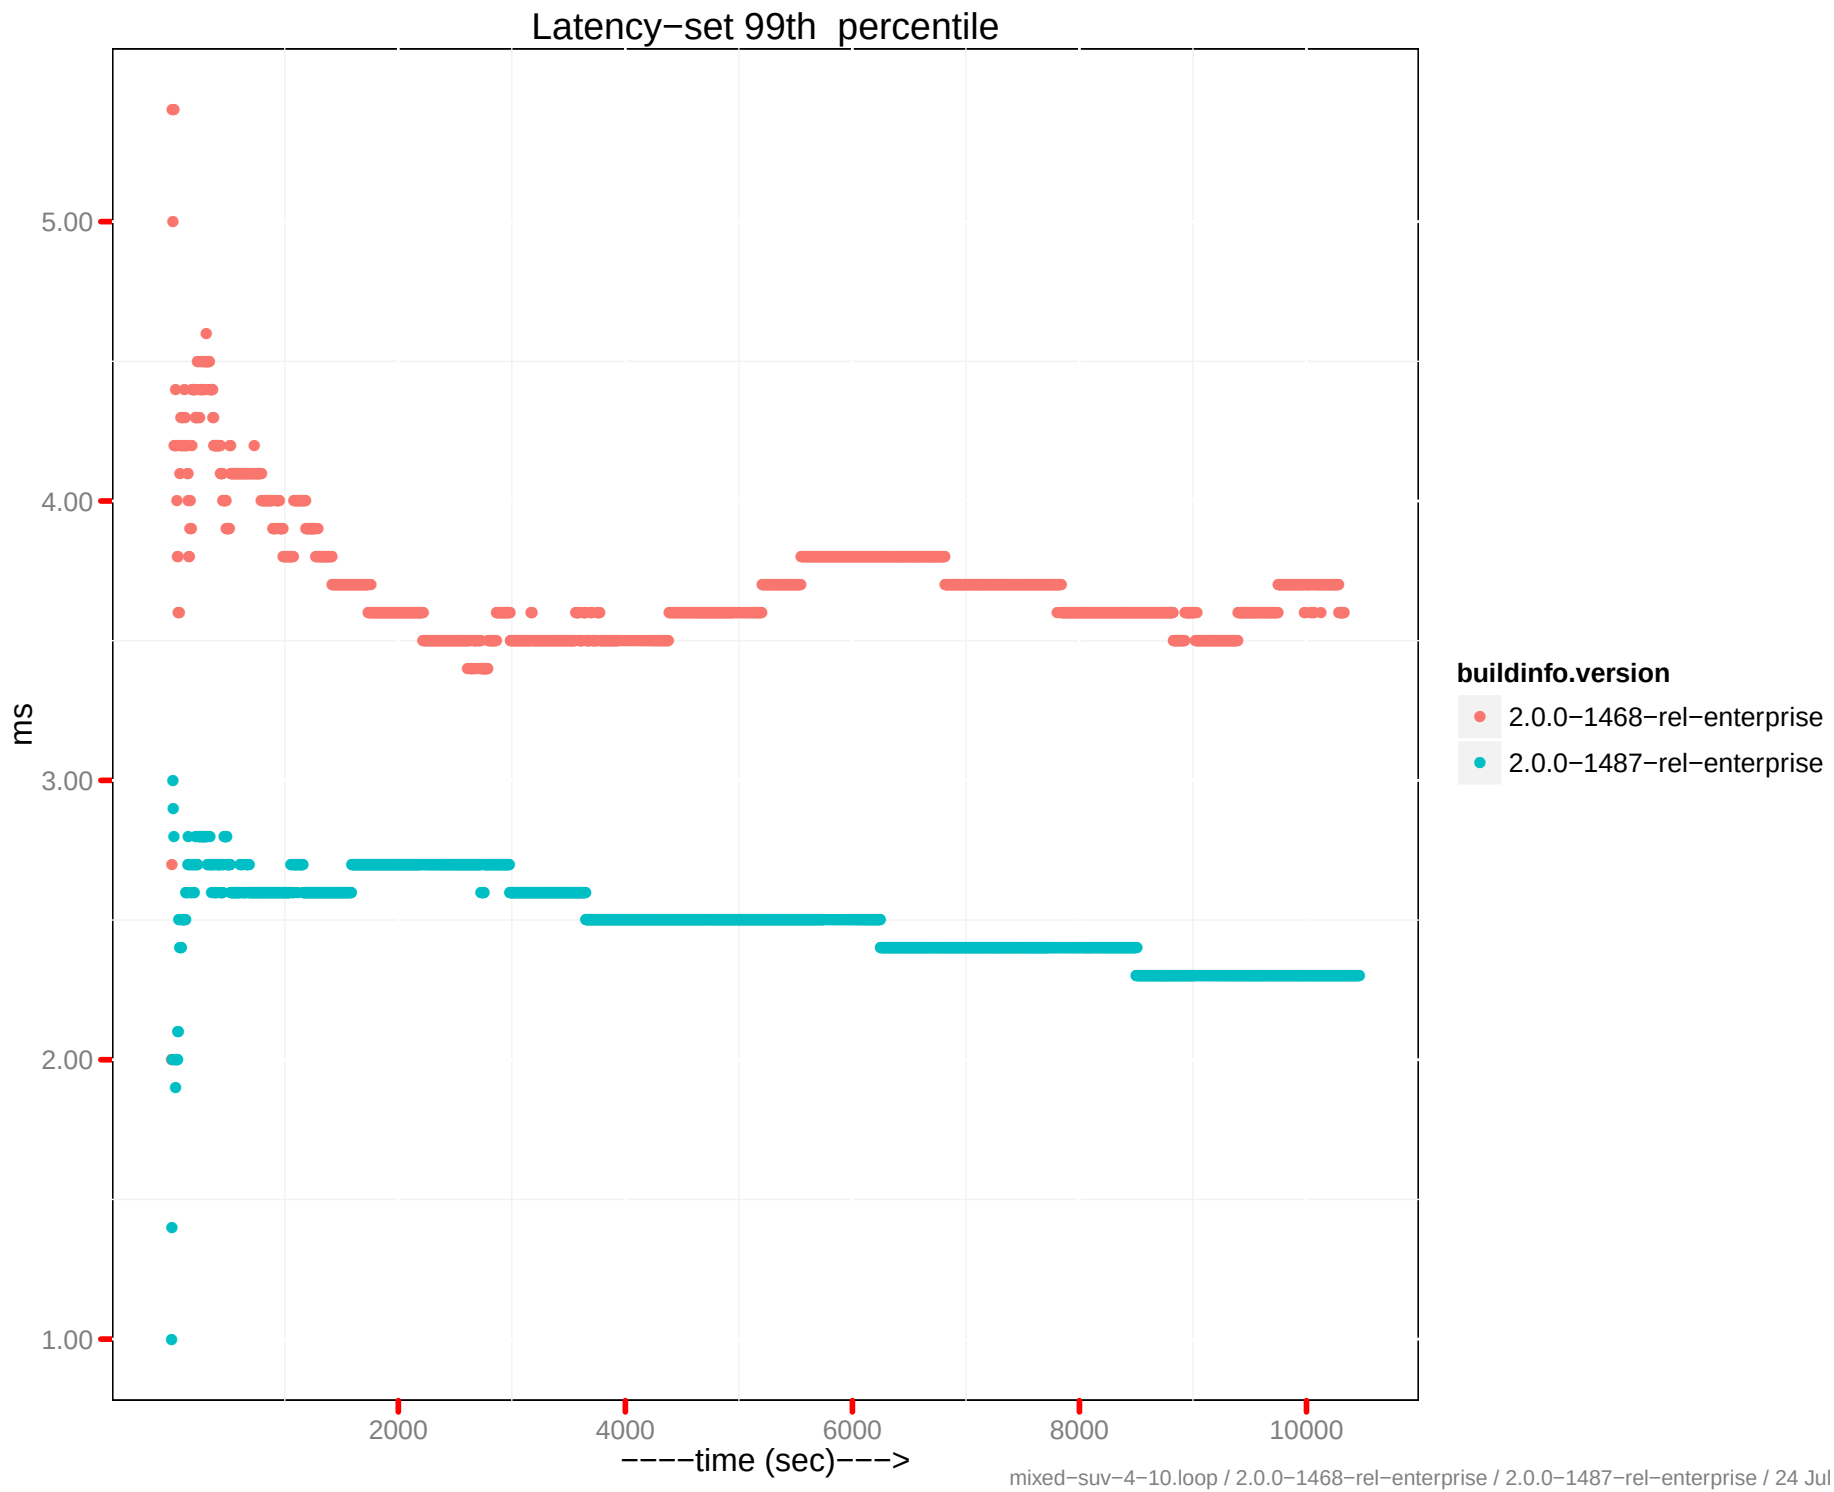
Latency−set 99th percentile
5.00
4.00
3.00
2.00
1.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

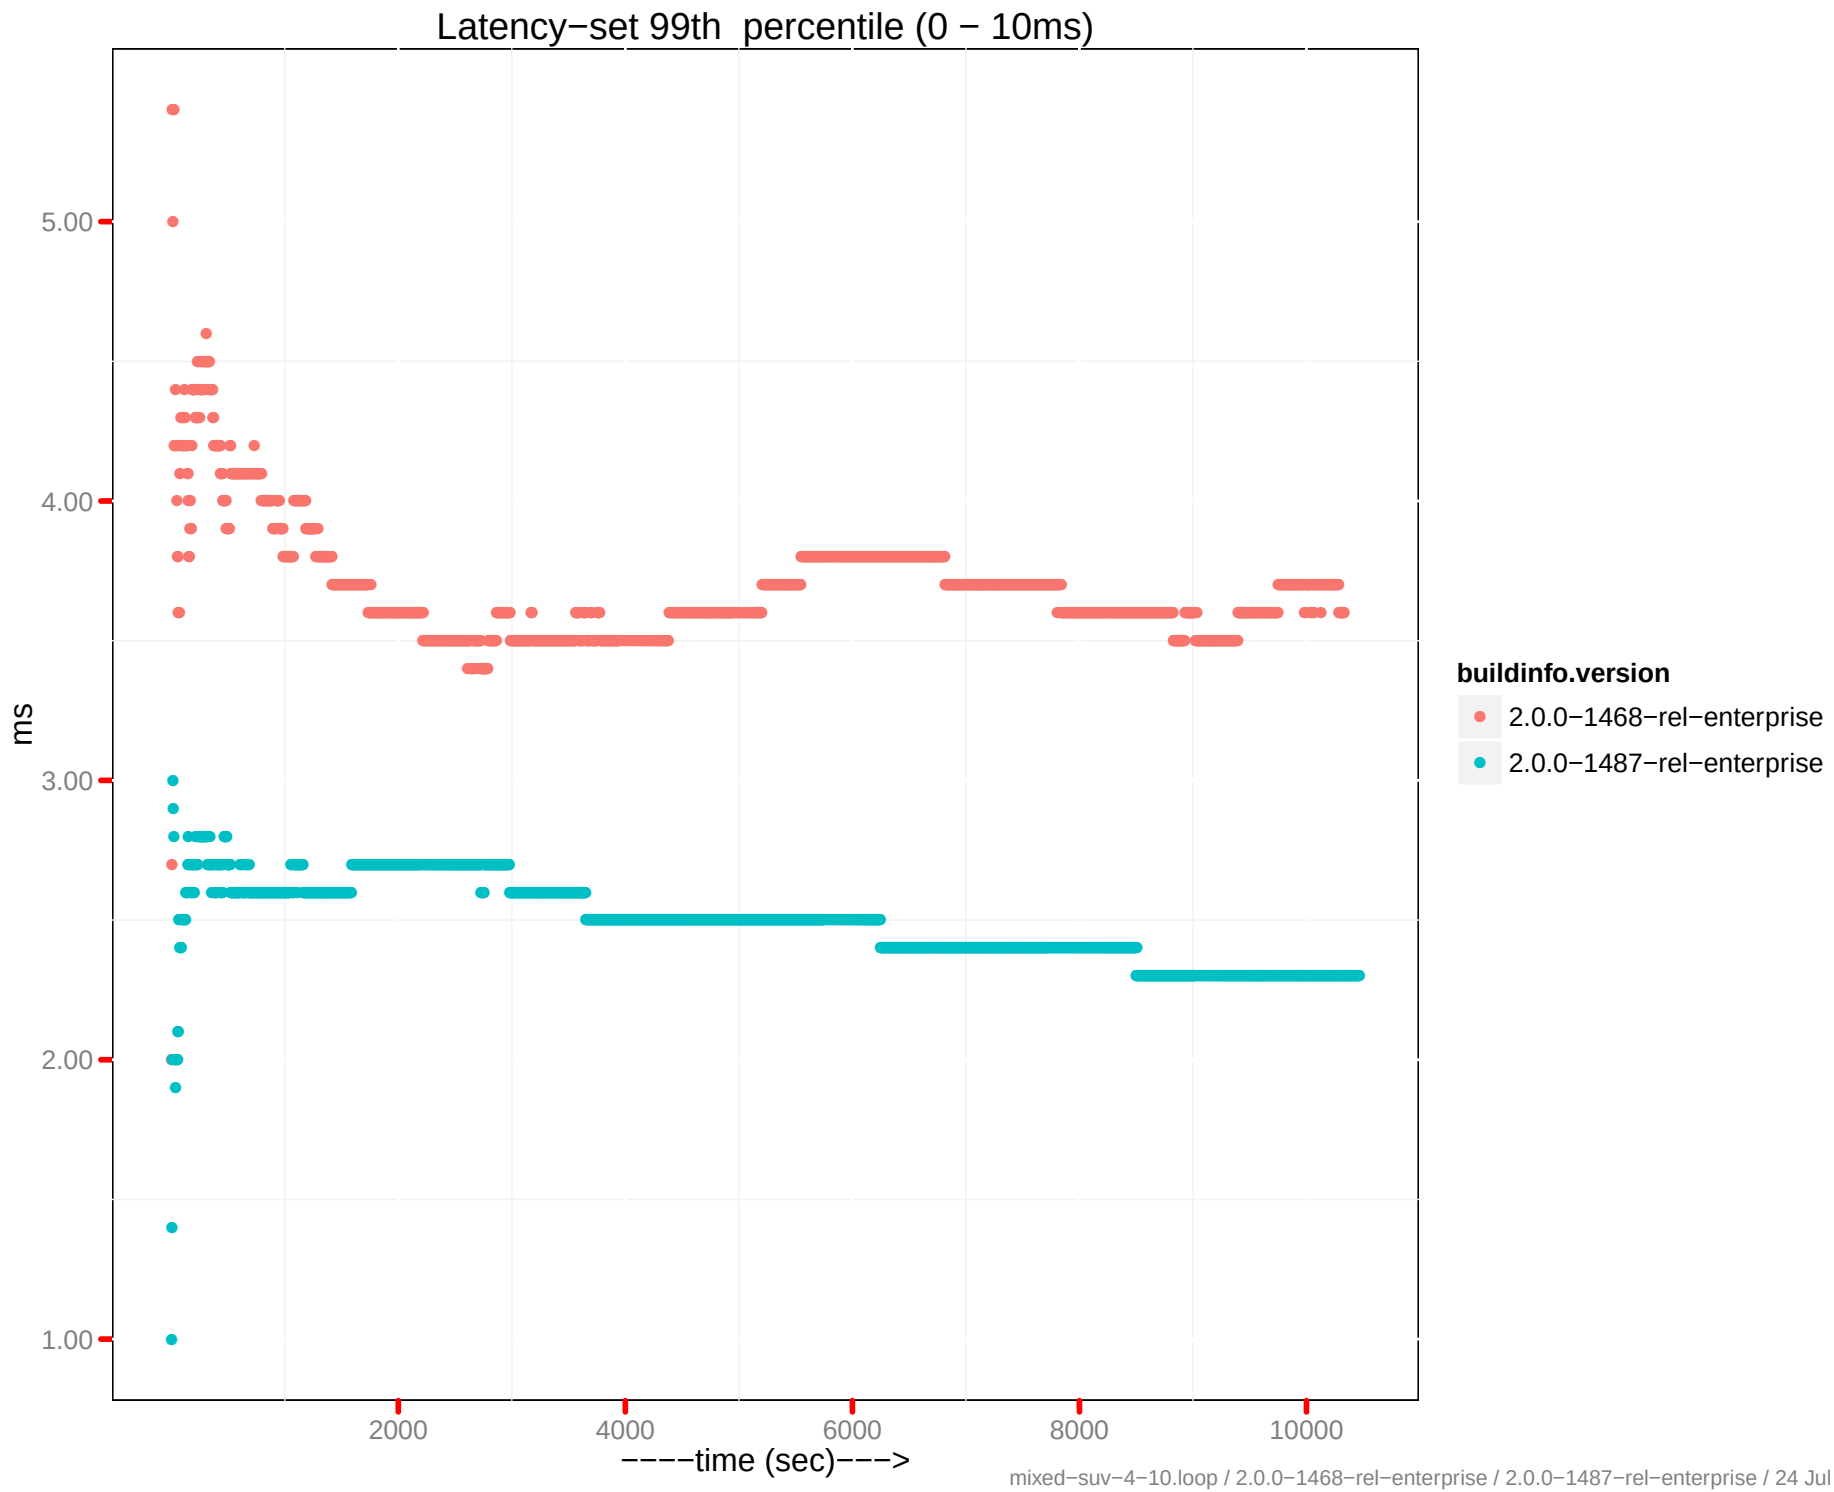
Latency−set 99th percentile (0 − 10ms)
5.00
4.00
3.00
2.00
1.00
ms
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1468−rel−enterprise
2.0.0−1487−rel−enterprise
mixed−suv−4−10.loop / 2.0.0−1468−rel−enterprise / 2.0.0−1487−rel−enterprise / 24 Jul 2012

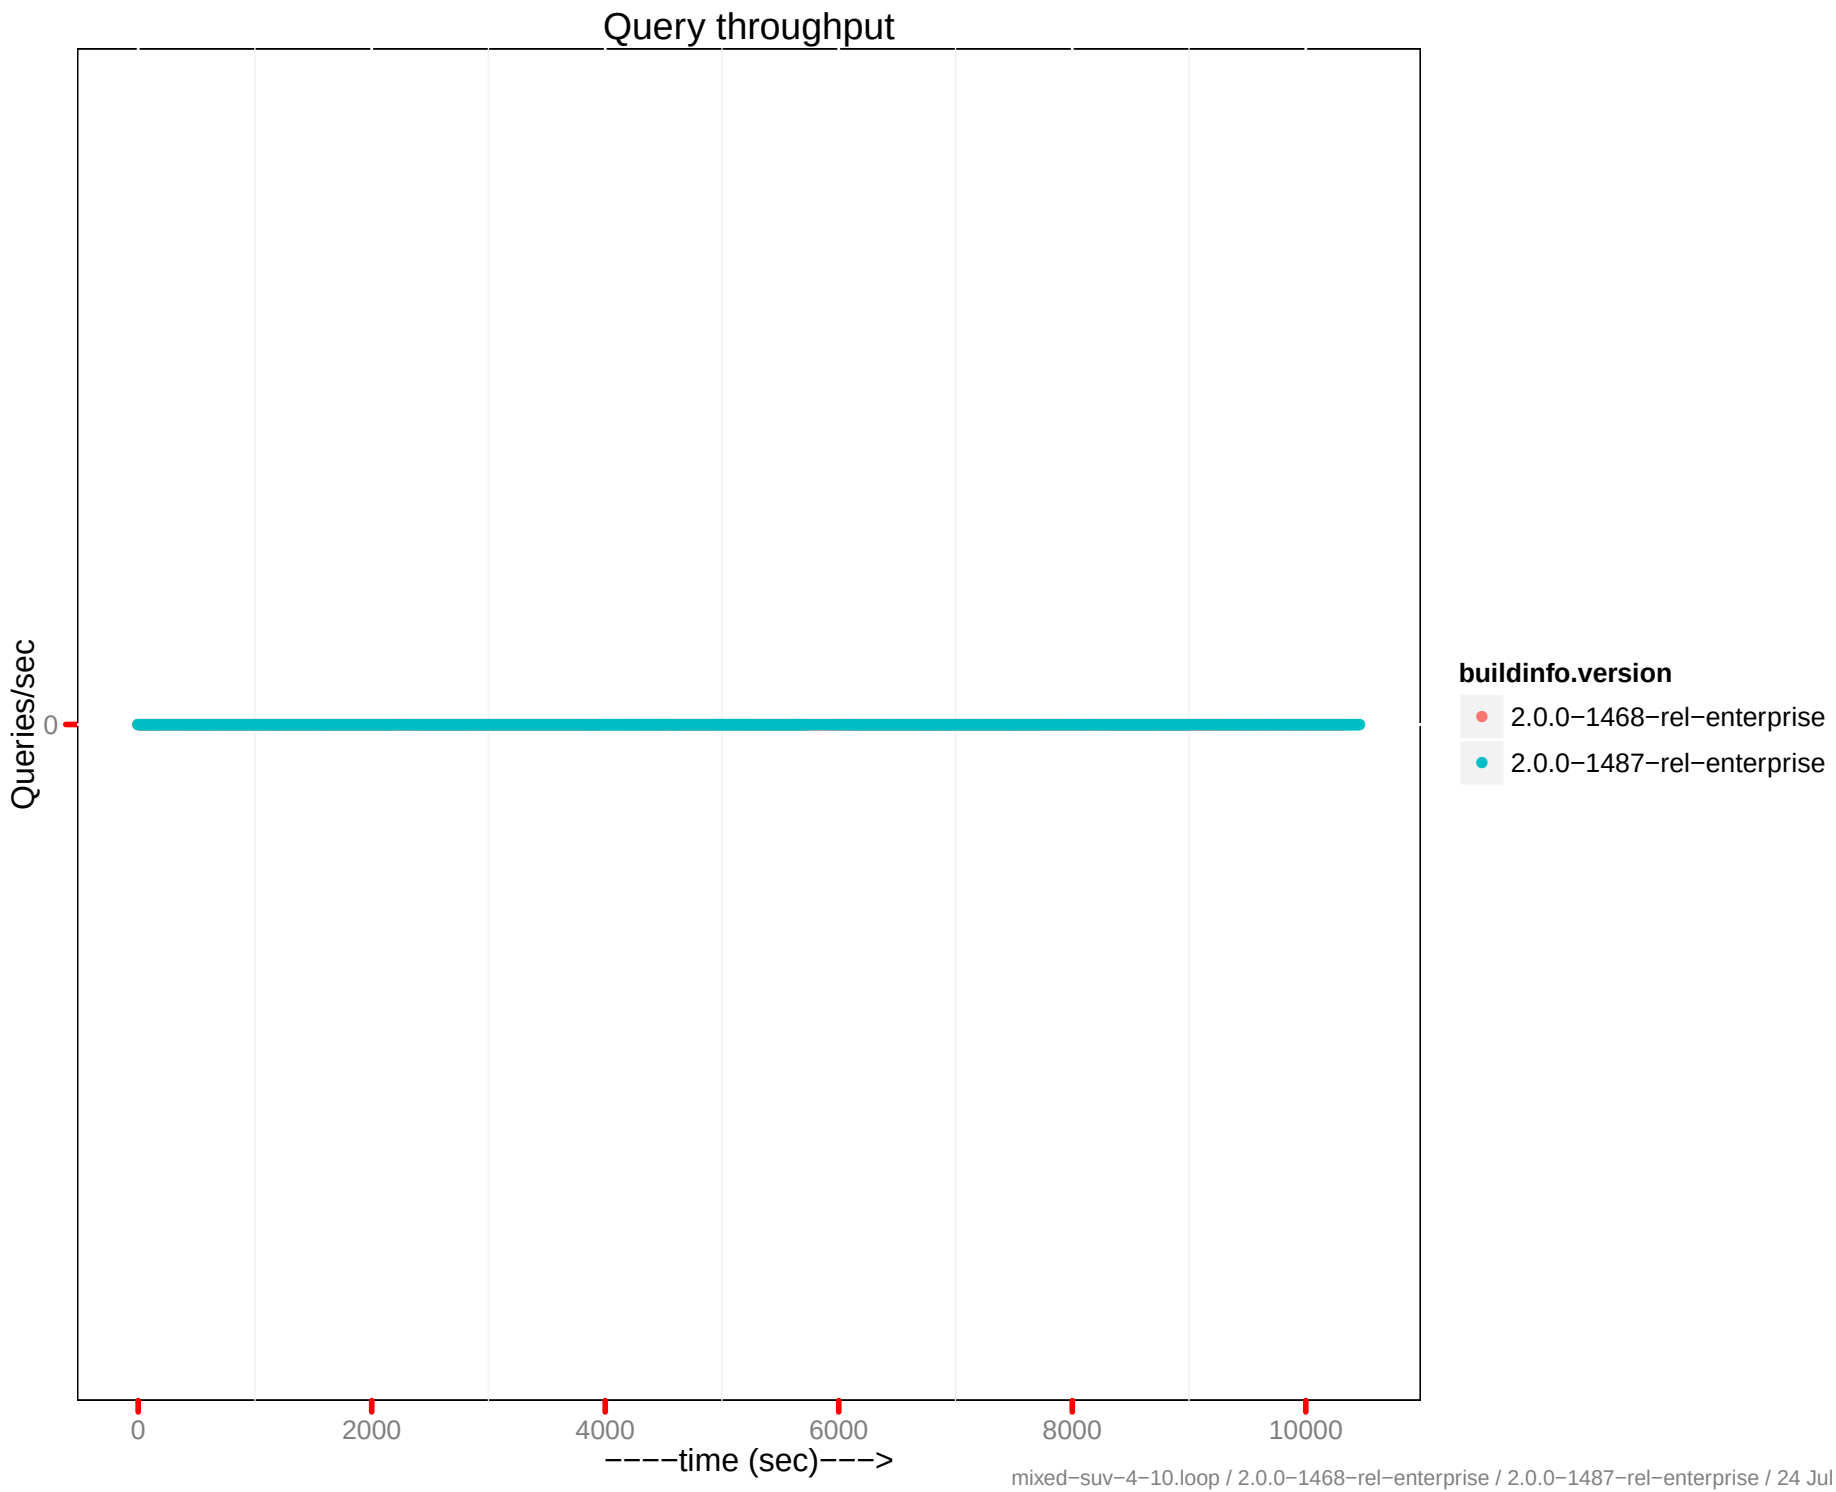
Query throughput
Queries/sec
0
0
2000
4000
6000
8000
10000
----time (sec)--->
buildinfo.version
2.0.0-1468-rel-enterprise
2.0.0-1487-rel-enterprise
mixed-suv-4-10.loop / 2.0.0-1468-rel-enterprise / 2.0.0-1487-rel-enterprise / 24 Jul 2012

```
mixed-suv-4-10.conf
# mixed 7M load, 1M hot reload, 3M access creates
# speed limit = 3k
#
performance.eperf.EPerfClient.test_eperf_mixed

params:

# general
batch=50
kind=nonjson
mem_quota=20000
spec=mixed-suv-4-10

# 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

# control (defaults: pytests/performance/perf_defaults.py)
load_wait_until_drained=1
loop_wait_until_drained=0
mcsoda_heartbeat=3
mcsoda_max_ops_sec=3000
tear_down=1
tear_down_proxy=1
tear_down_bucket=0
tear_down_cluster=1
tear_down_on_setup=0
```

```
hummer-dedicated-1.ini
[global]
username:root
password:couchbase
port:8091
data_path:/data

[servers]
1:192.168.162.20
2:192.168.162.21

[clients]
1:192.168.162.24
2:192.168.162.25
3:192.168.162.26

[membase]
rest_username:Administrator
rest_password:password

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