	2.0.0 − 1583	2.0.0 − 1583
Runtime (in hr)	0.76	NA
Avg. Drain Rate	6.79K	NANA
Peak Disk (GB)	56.8	NA
Peak Memory (GB)	16.86	NA
Avg. OPS	2.00K	NANA
Avg. mem memcached (GB)	16.42	NA
Avg. mem beam.smp (MB)	368.84	NA
Latency−get (90th) (ms)	1.56	NA
Latency−get (95th) (ms)	3.03	NA
Latency−get (99th) (ms)	8.45	NA
Latency−set (90th) (ms)	1.62	NA
Latency−set (95th) (ms)	3.05	NA
Latency−set (99th) (ms)	7.02	NA
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	NA
Rebalance Time (sec)	2731.88	NA
Testrunner Version	4e95644	NA

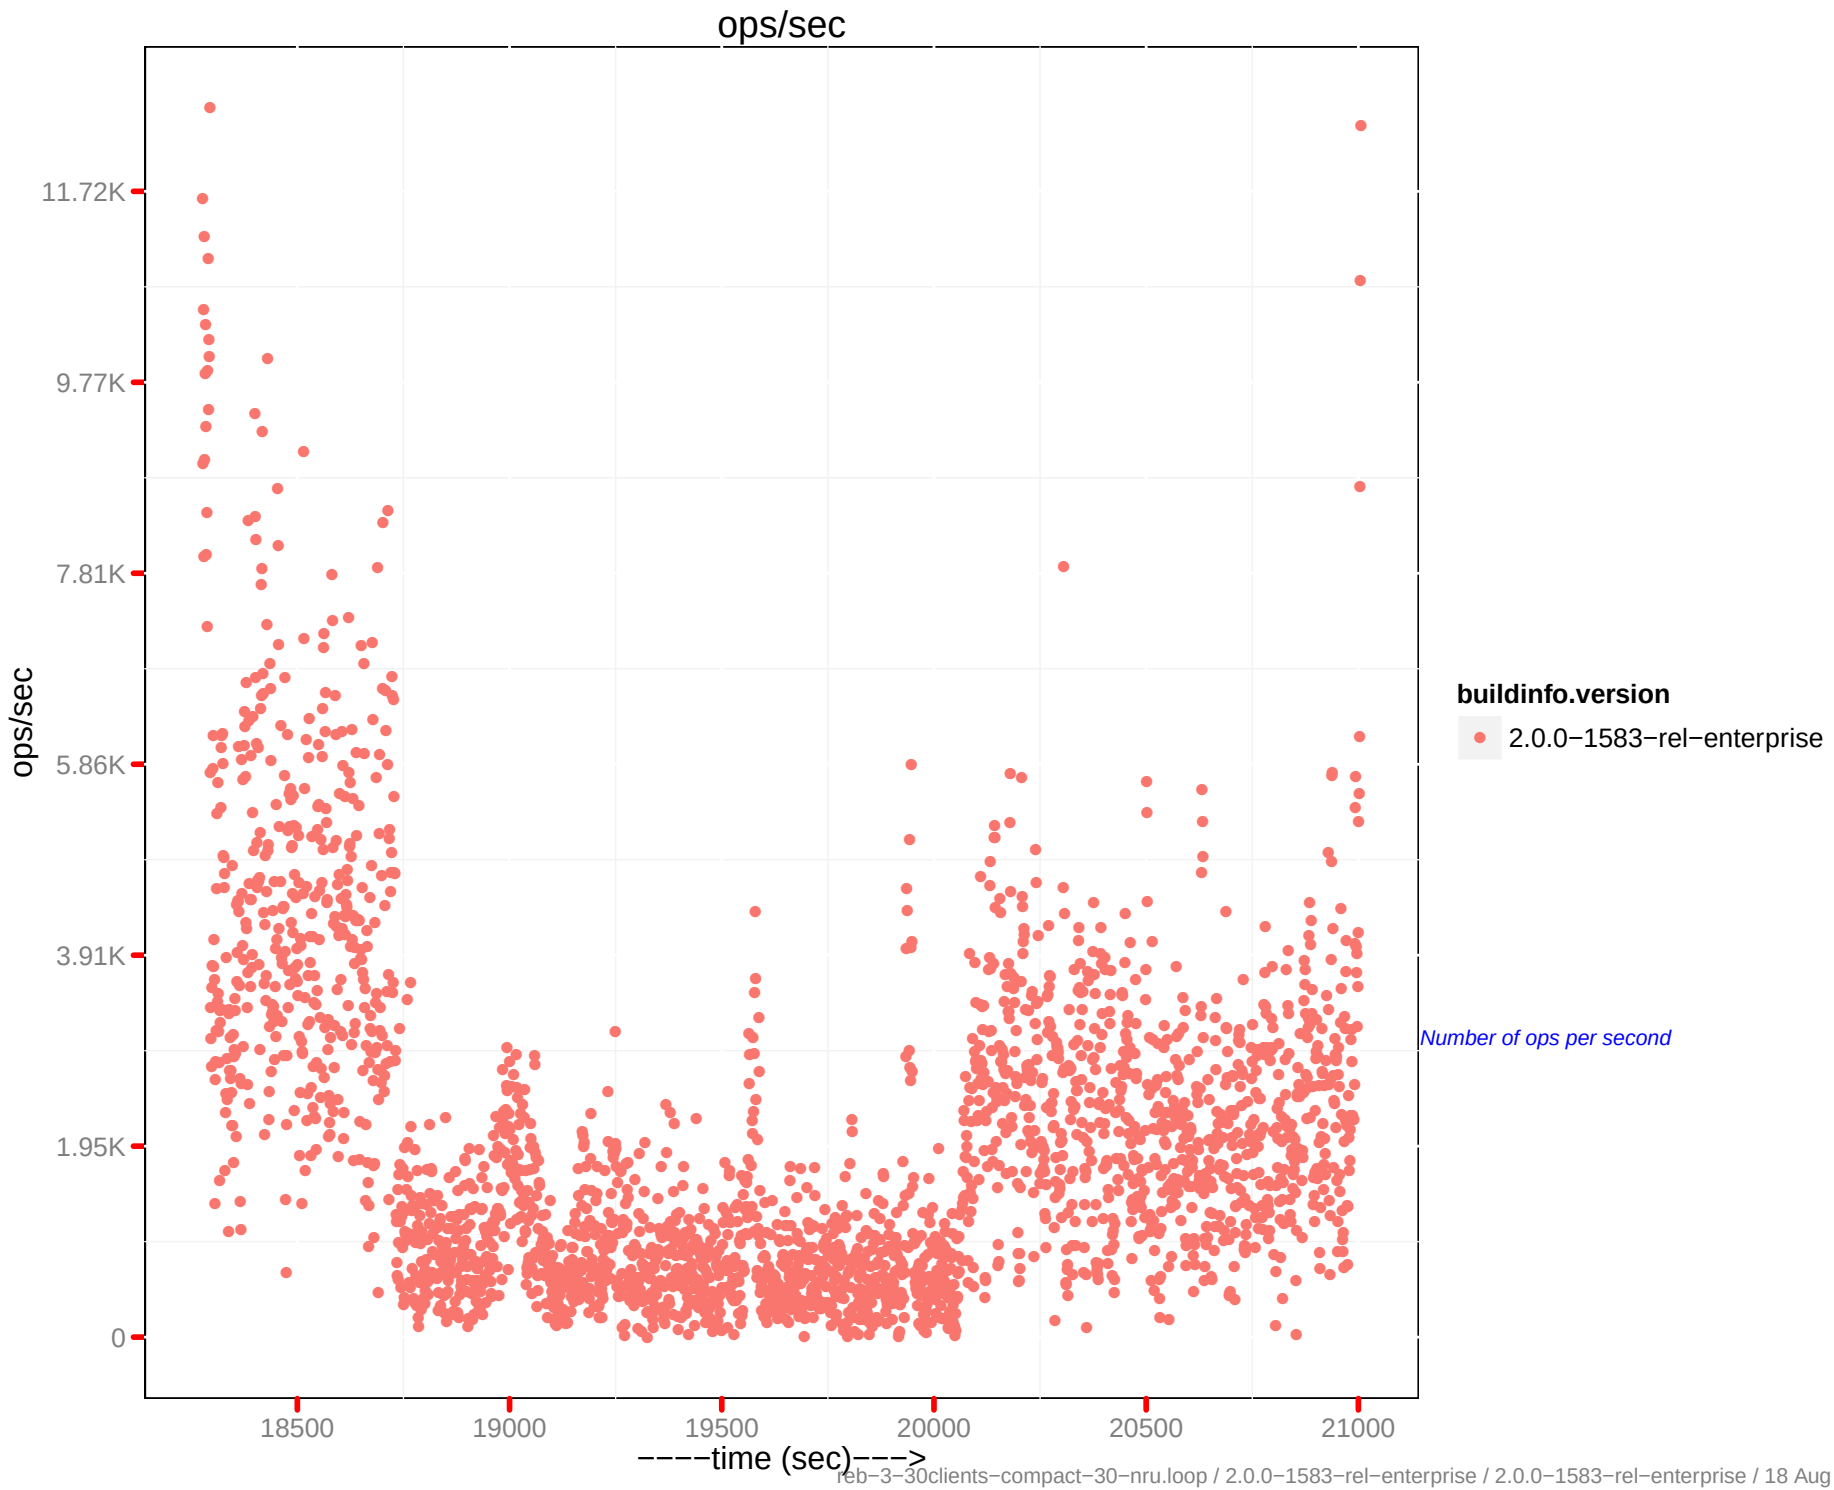
ops/sec
11.72K
9.77K
7.81K
5.86K
3.91K
1.95K
0
ops/sec
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of ops per second
feb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

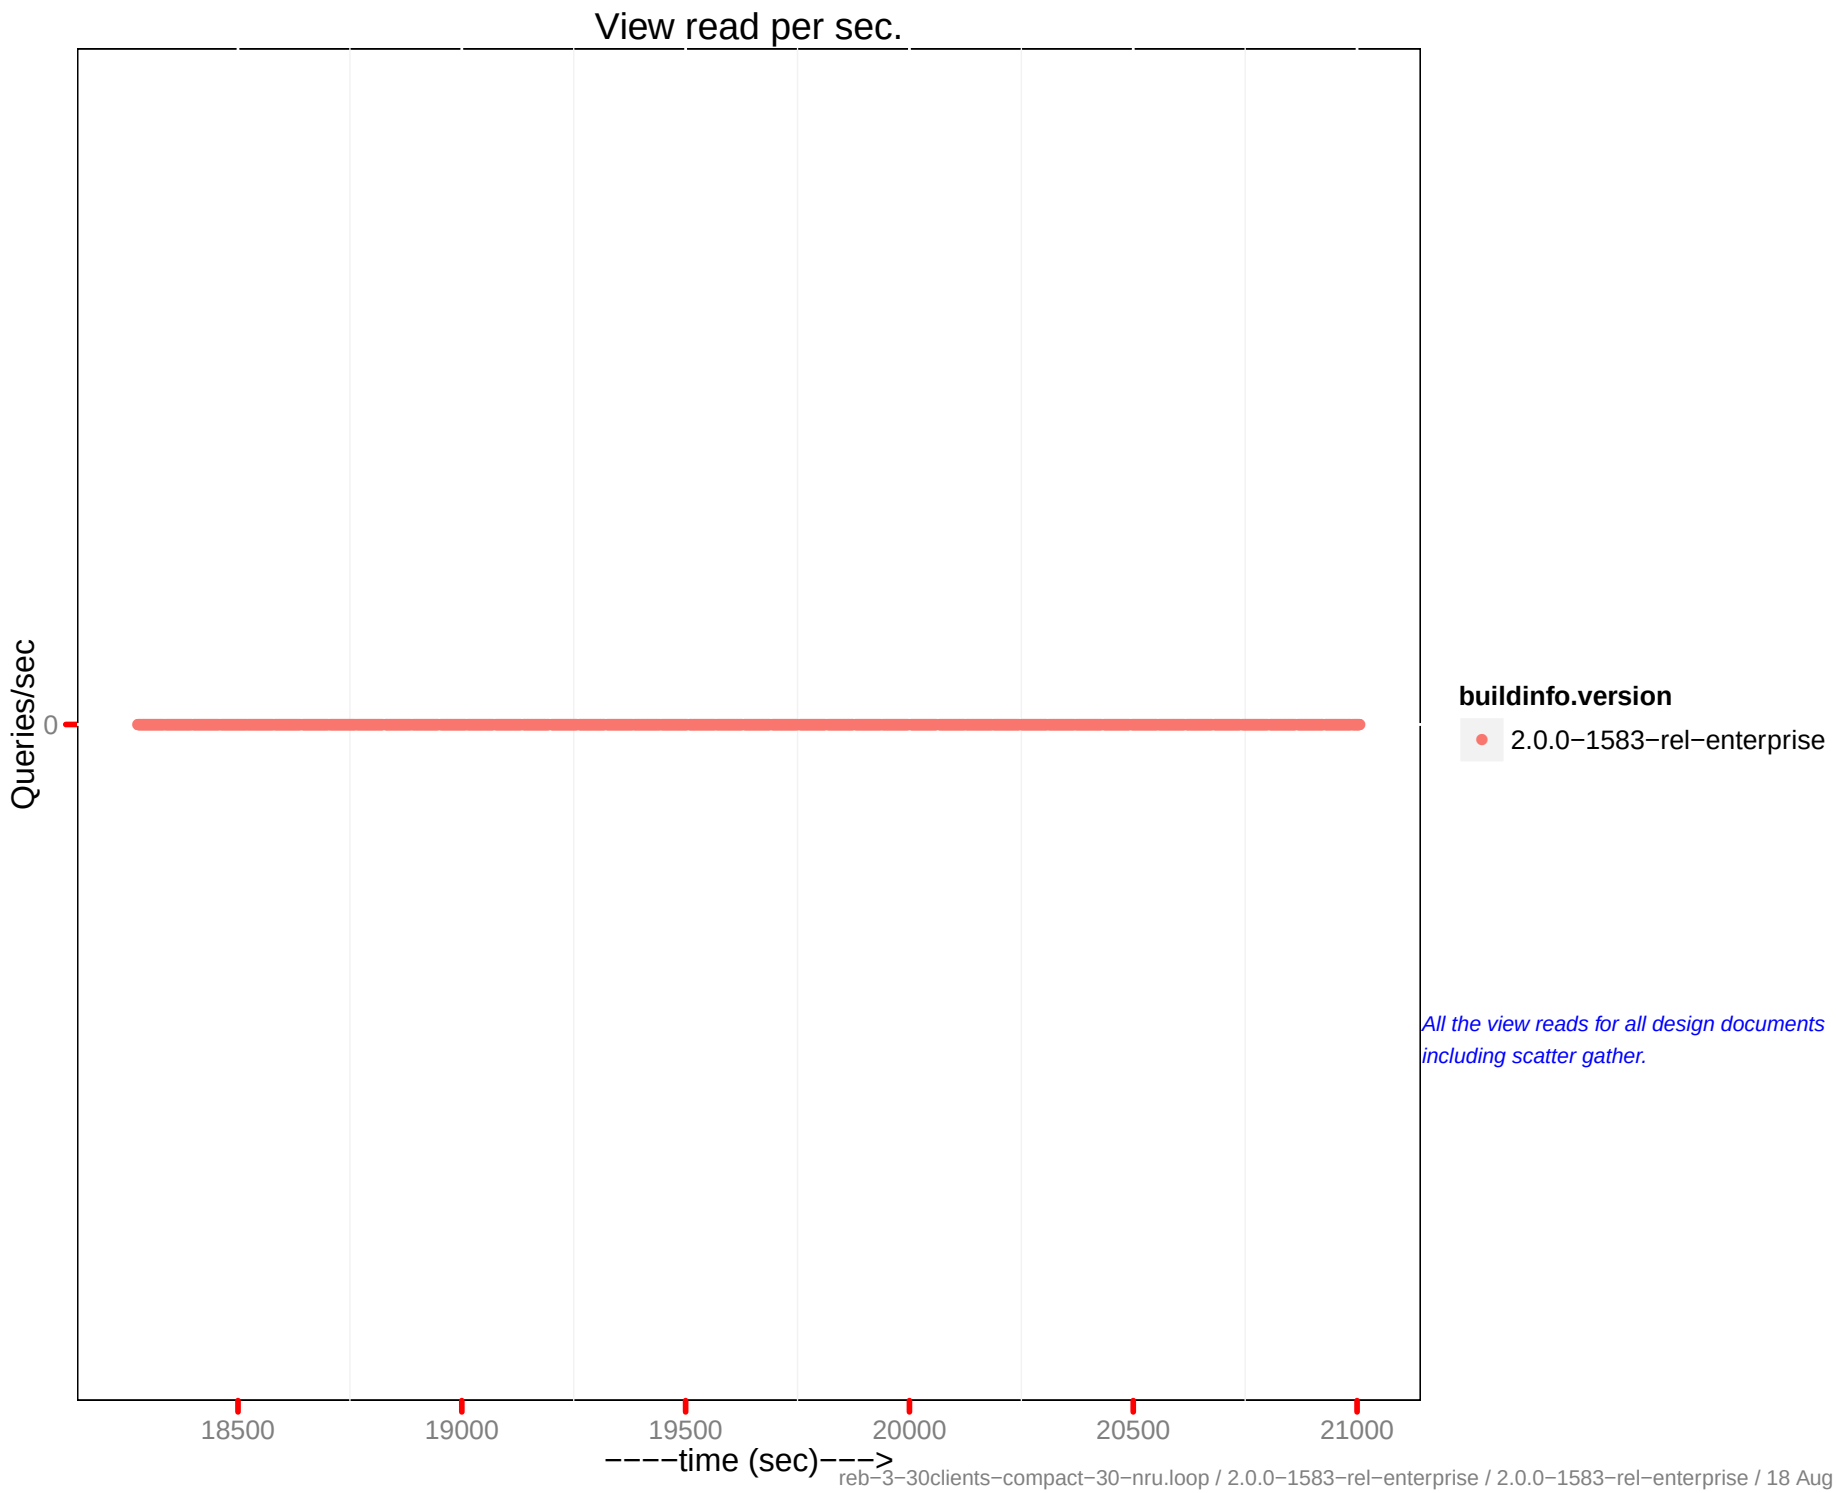
View read per sec.
Queries/sec
0
18500
19000
19500
20000
20500
21000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
All the view reads for all design documents including scatter gather.
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

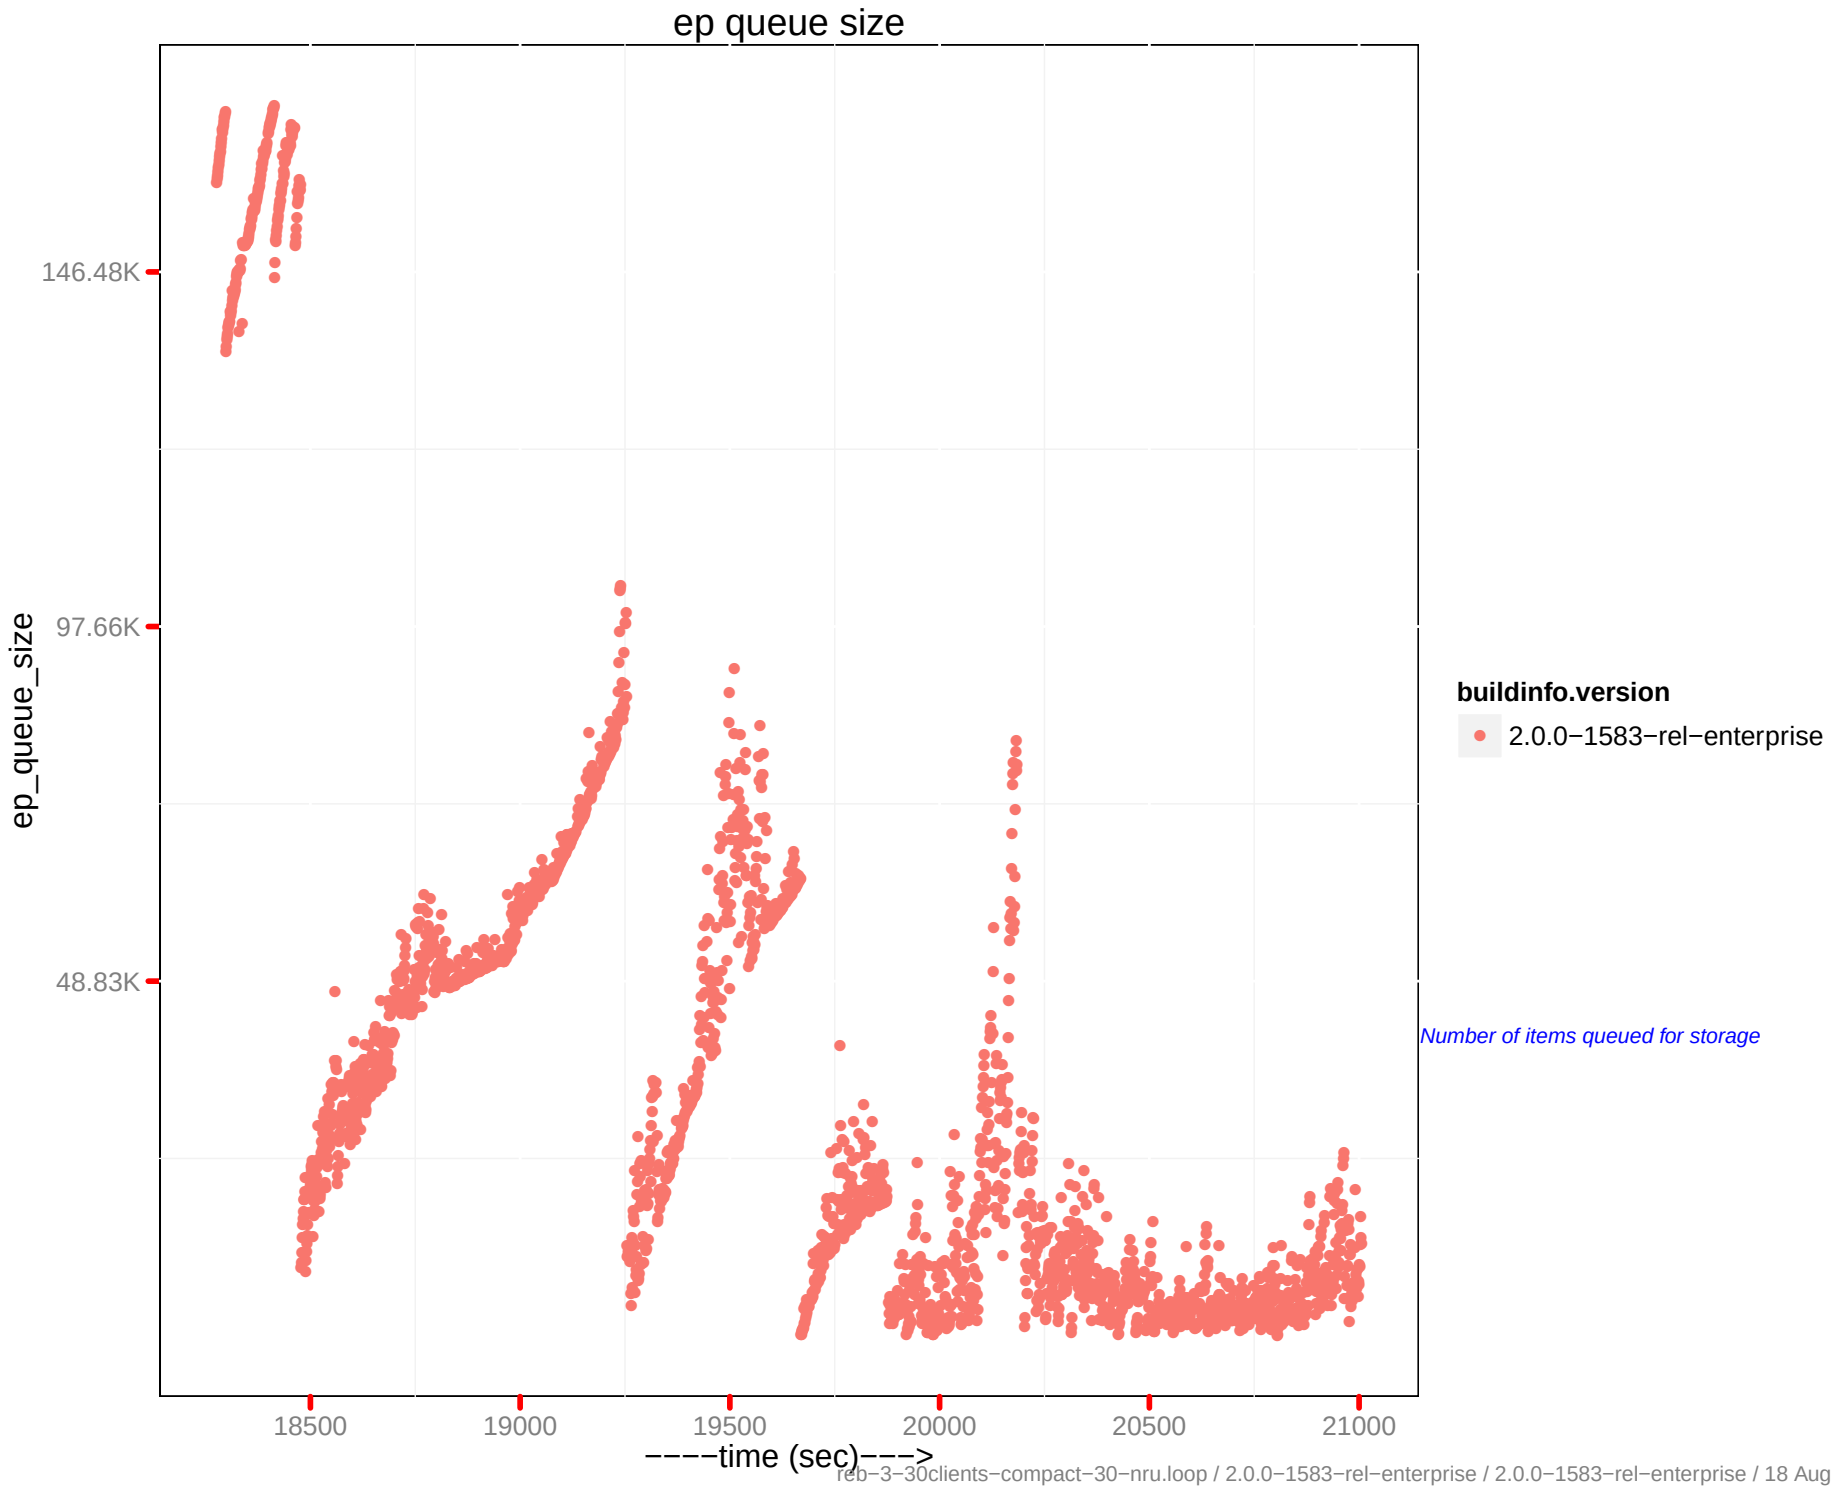
ep queue size
146.48K
97.66K
48.83K
ep_queue_size
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of items queued for storage
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

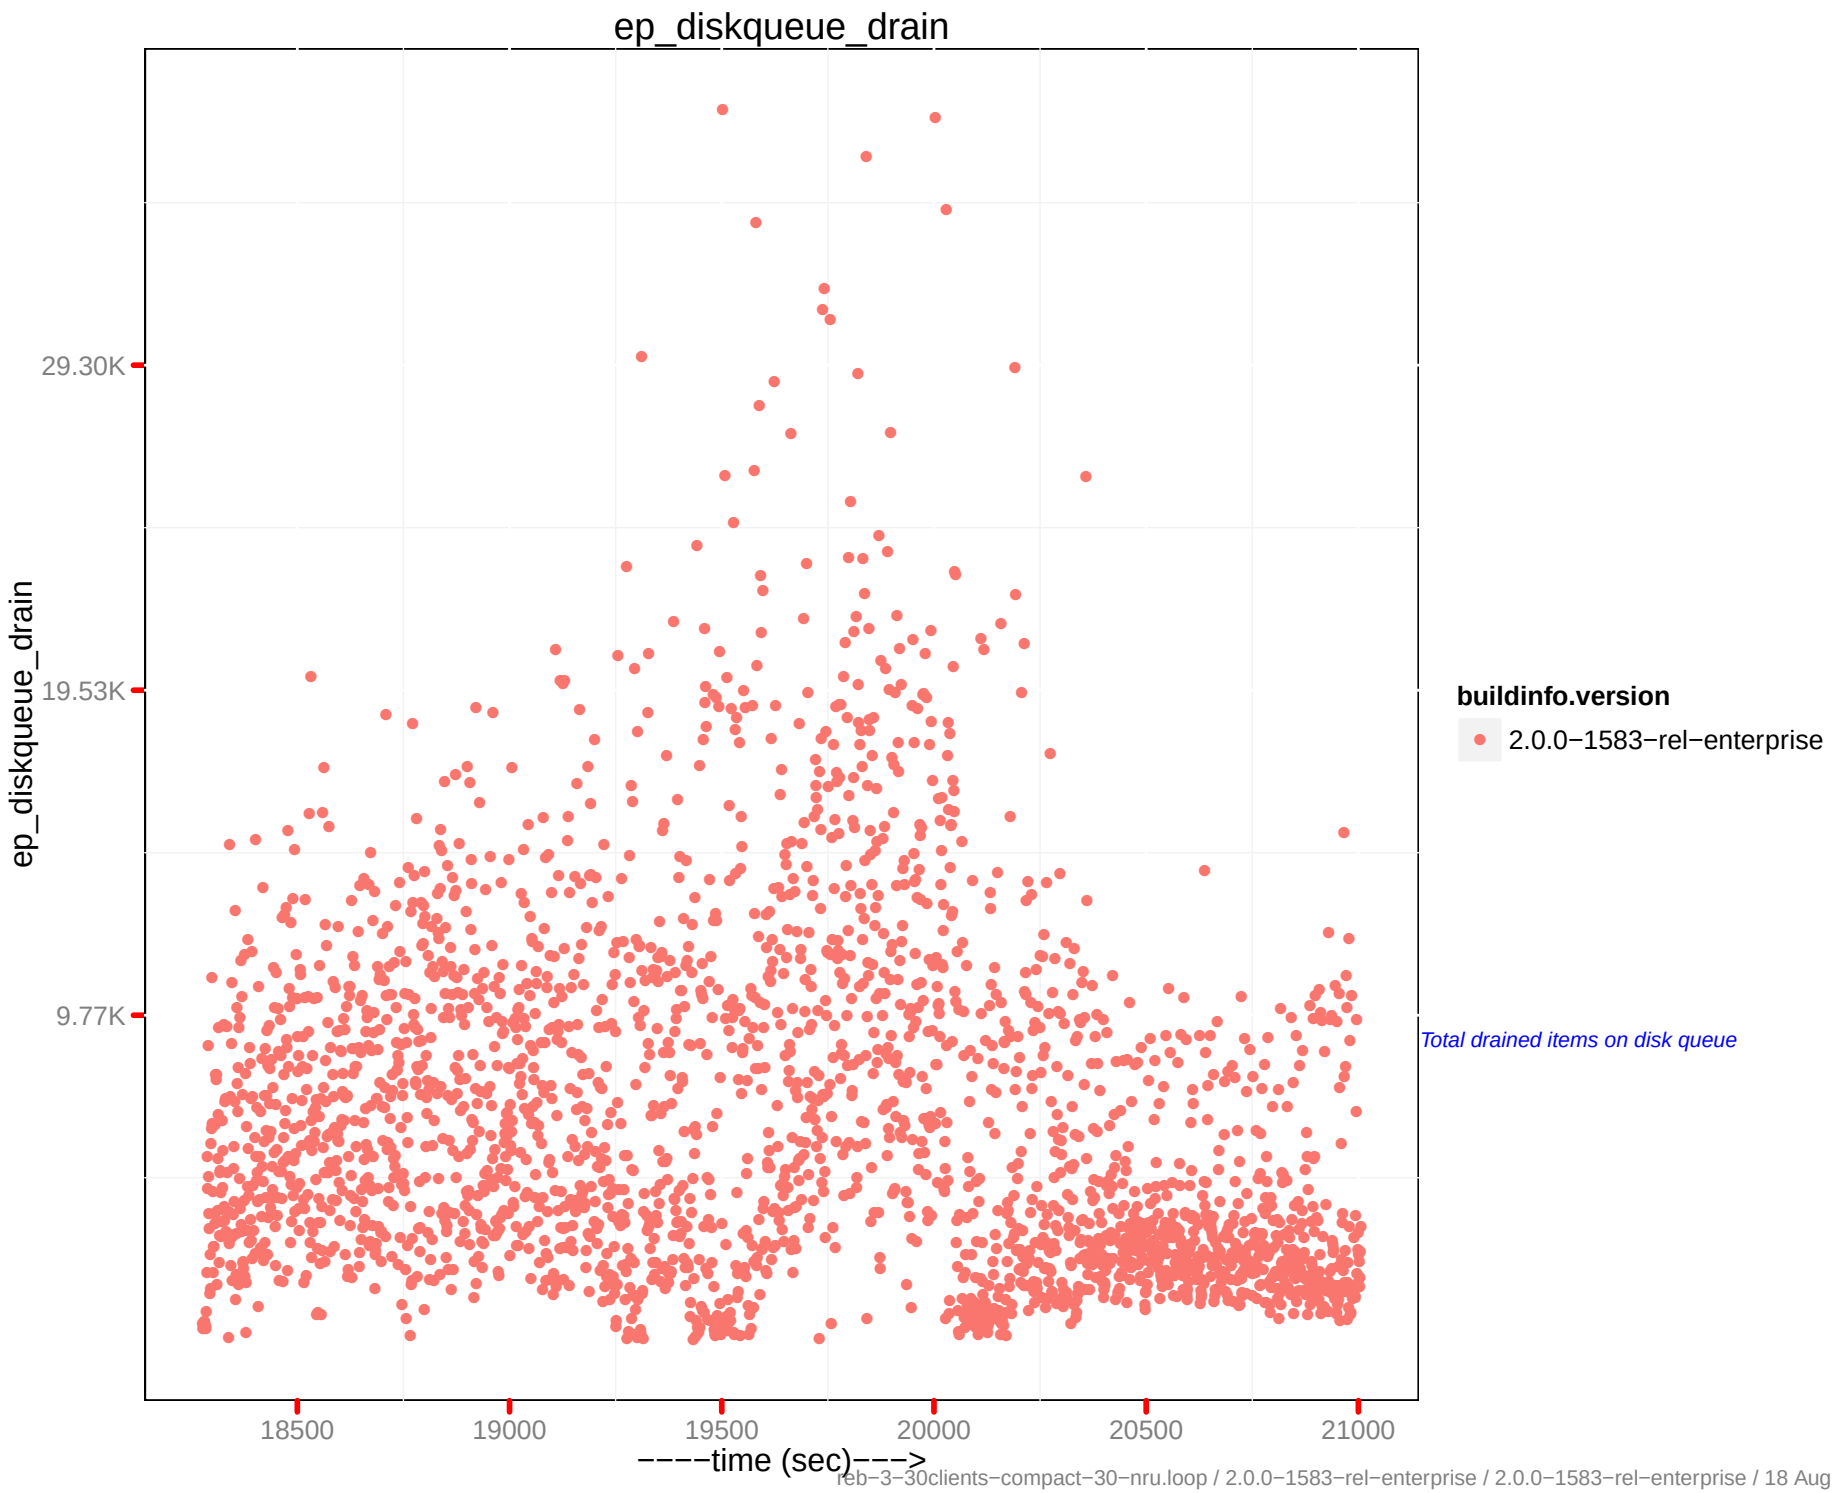
ep_diskqueue_drain
29.30K
19.53K
9.77K
ep_diskqueue_drain
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Total drained items on disk queue
feb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

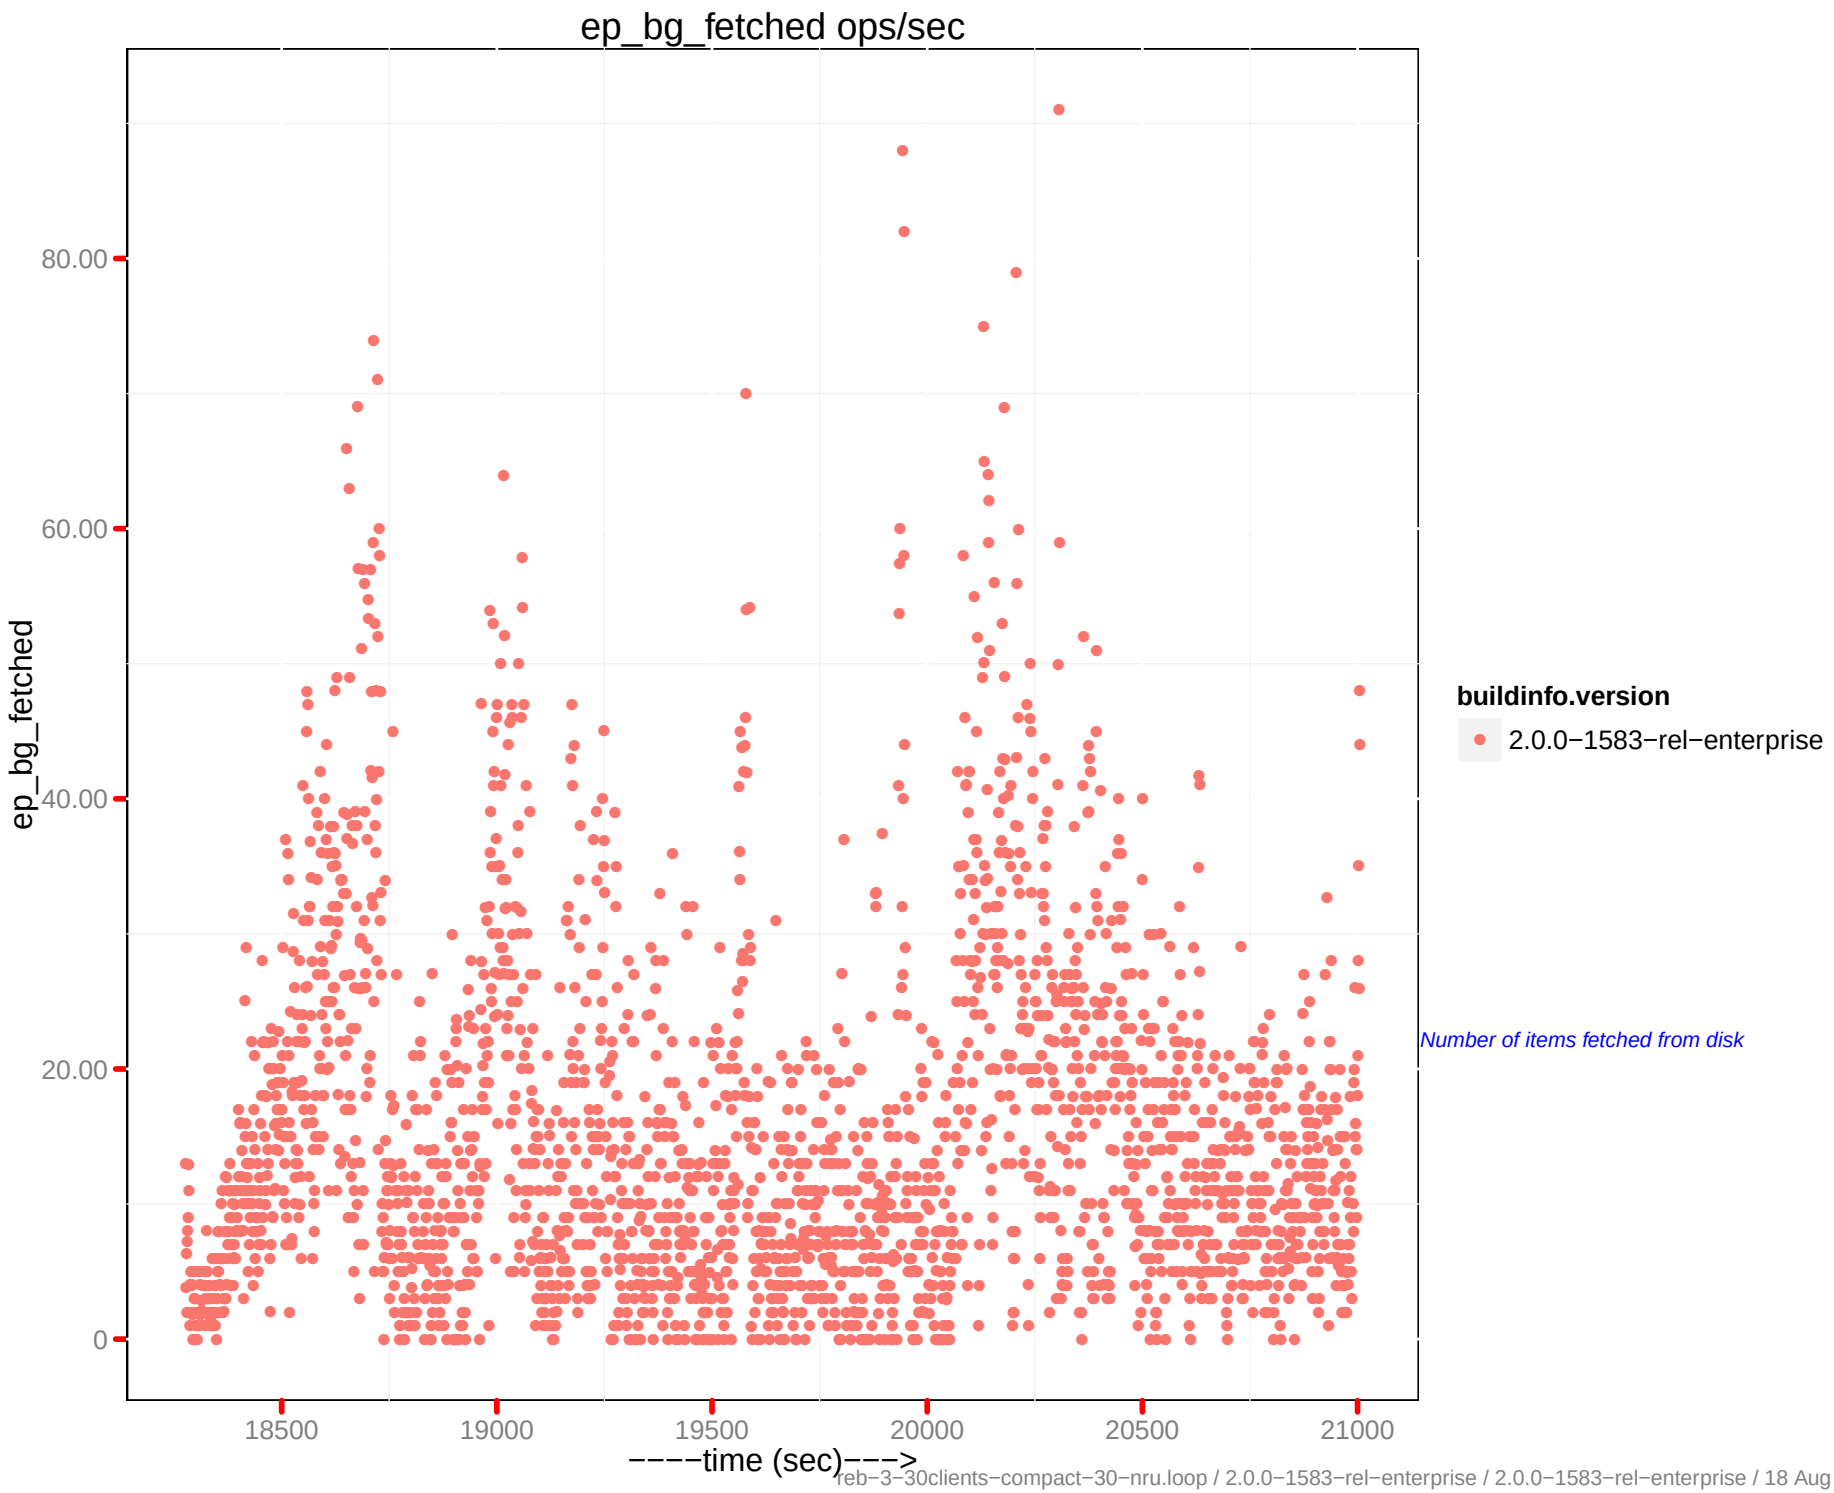
ep_bg_fetched ops/sec
ep_bg_fetched
80.00
60.00
40.00
20.00
0
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of items fetched from disk
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

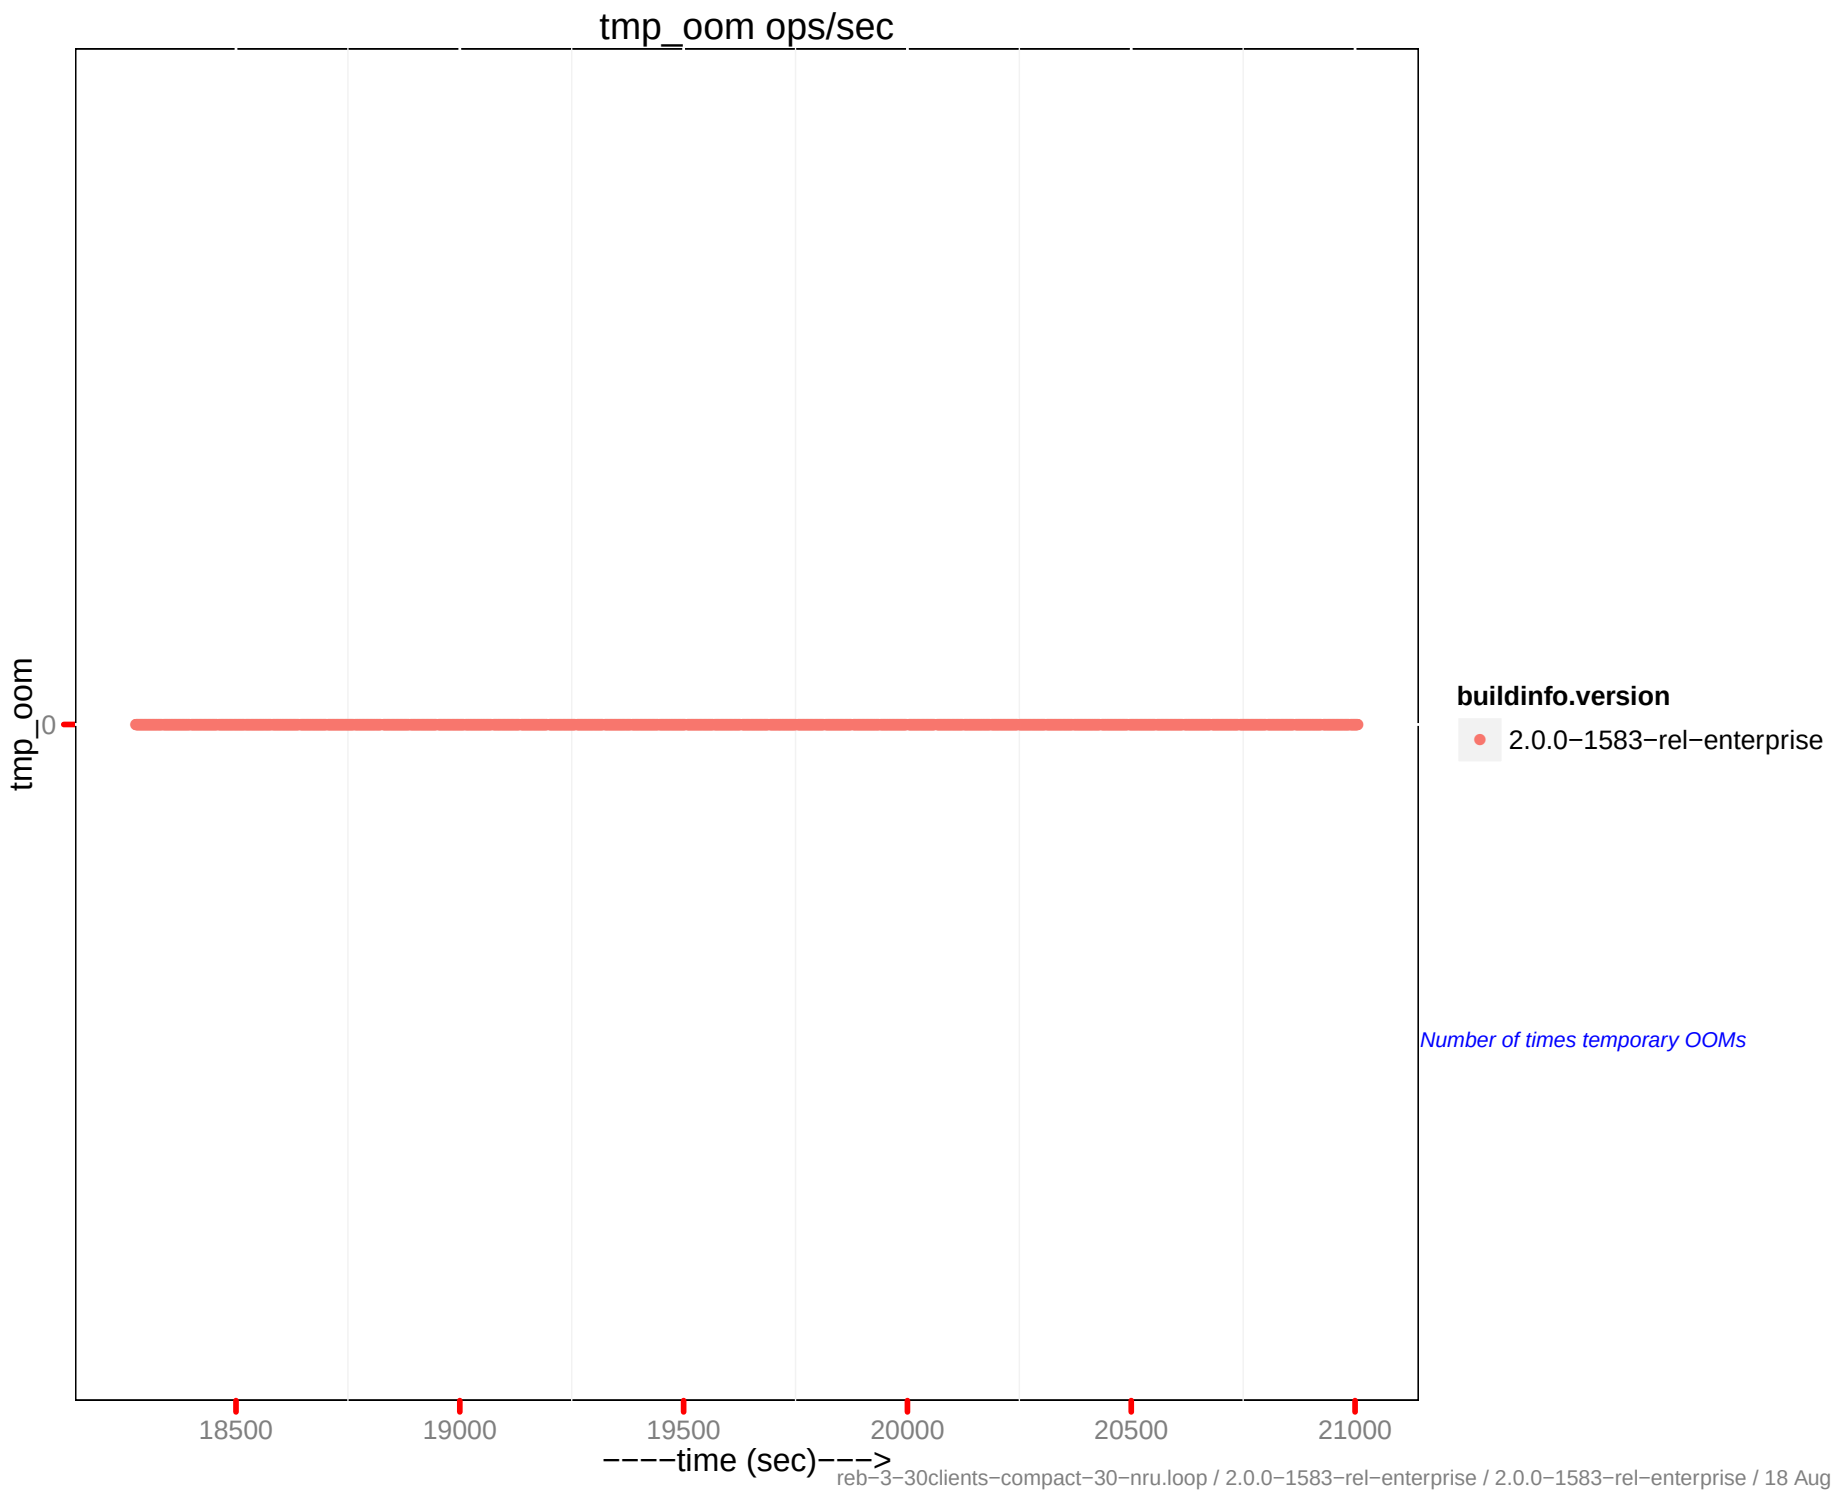
tmp_oom ops/sec
tmp_oom
0
18500
19000
19500
20000
20500
21000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
Number of times temporary OOMs
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

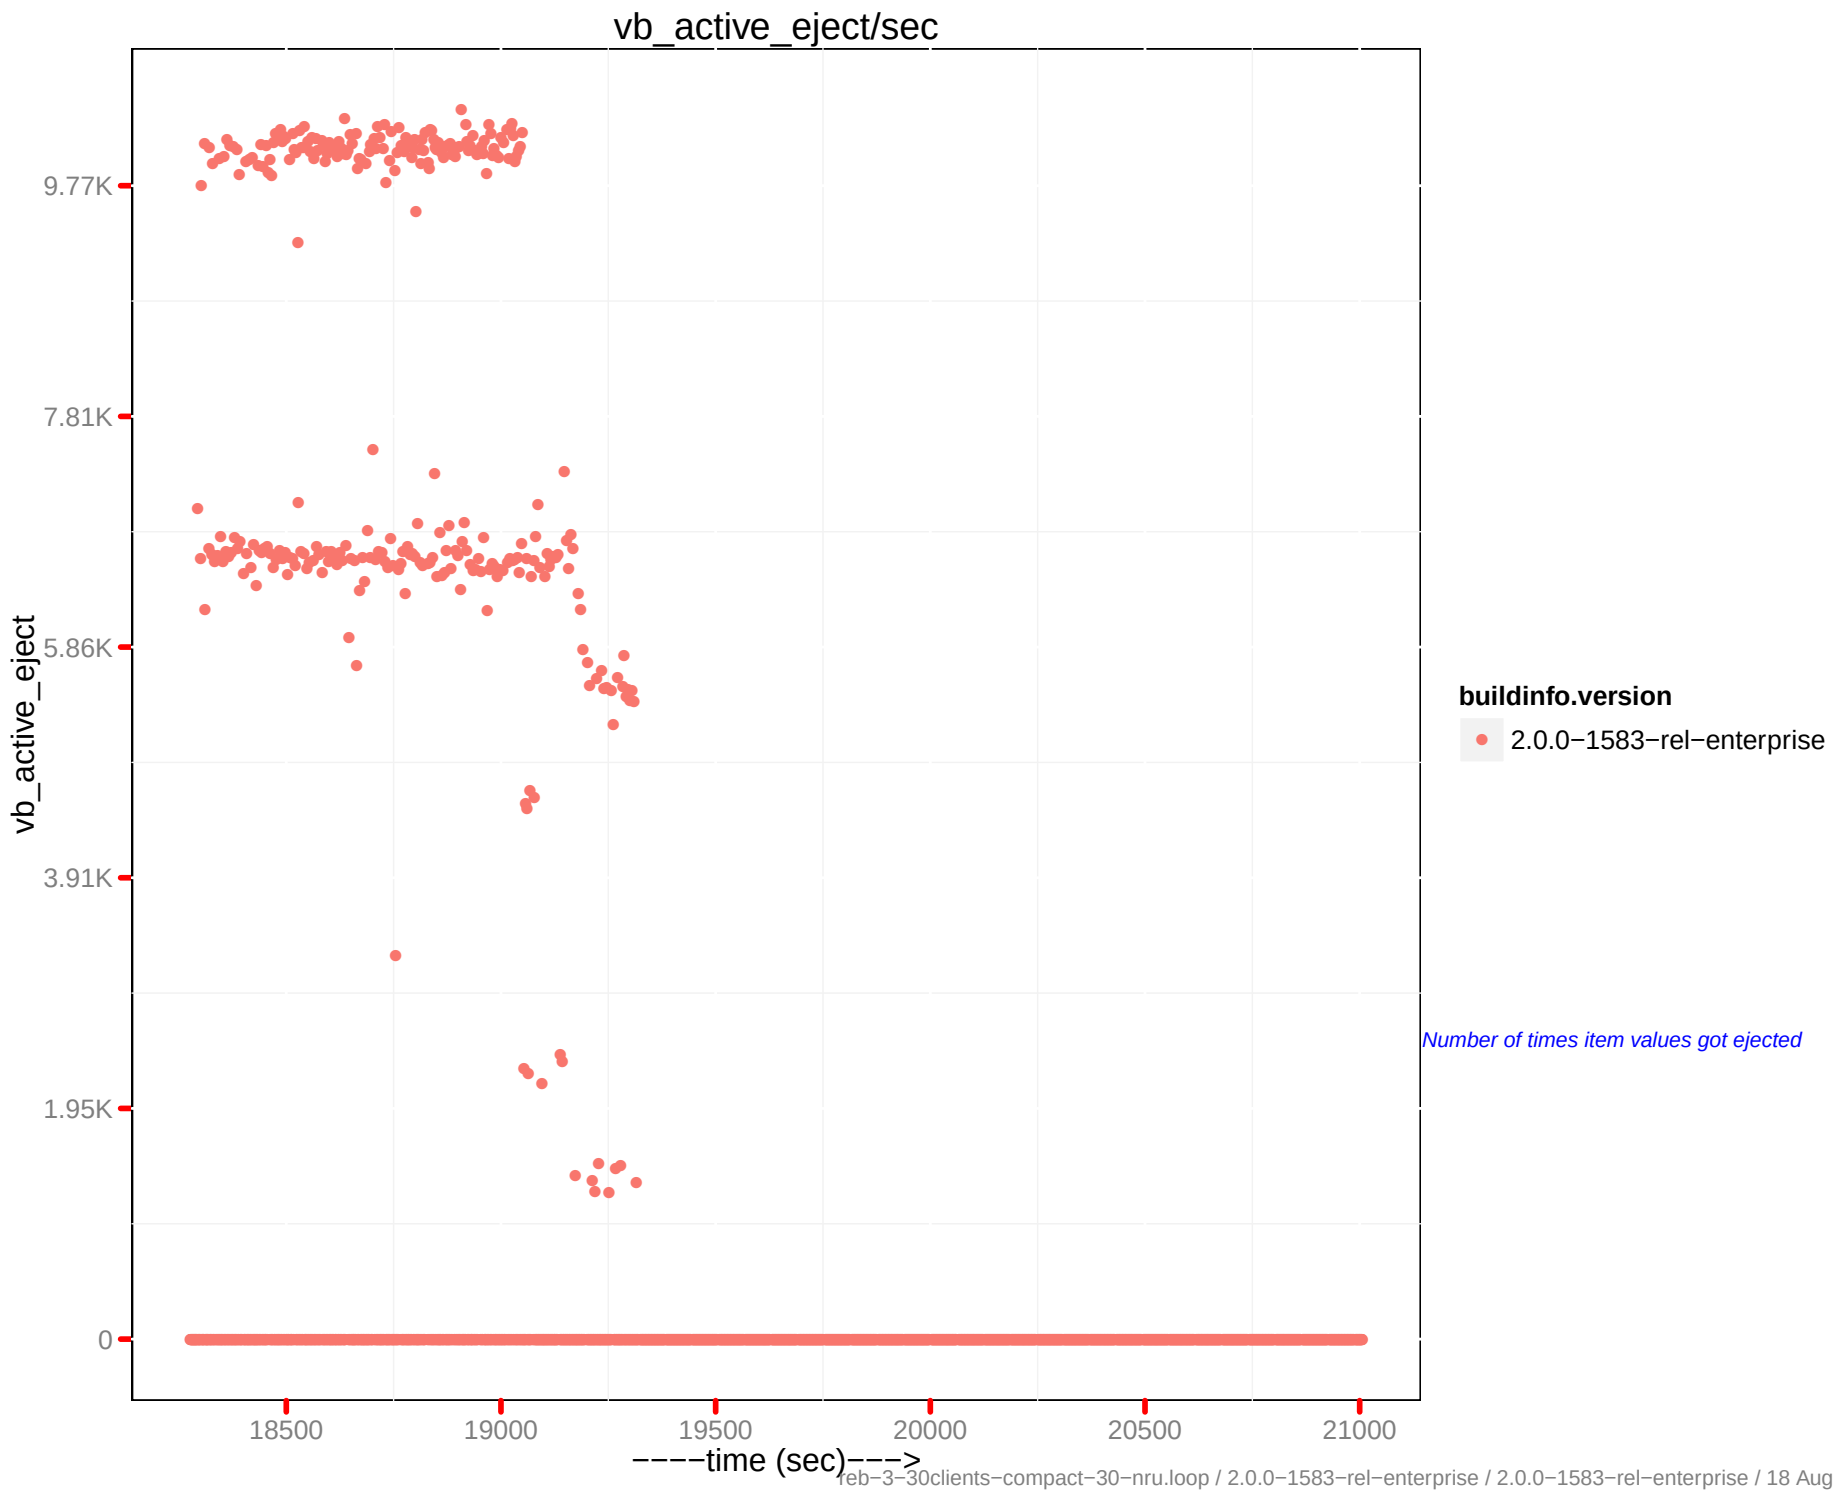
vb_active_eject/sec
9.77K
7.81K
5.86K
3.91K
1.95K
0
vb_active_eject
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of times item values got ejected
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

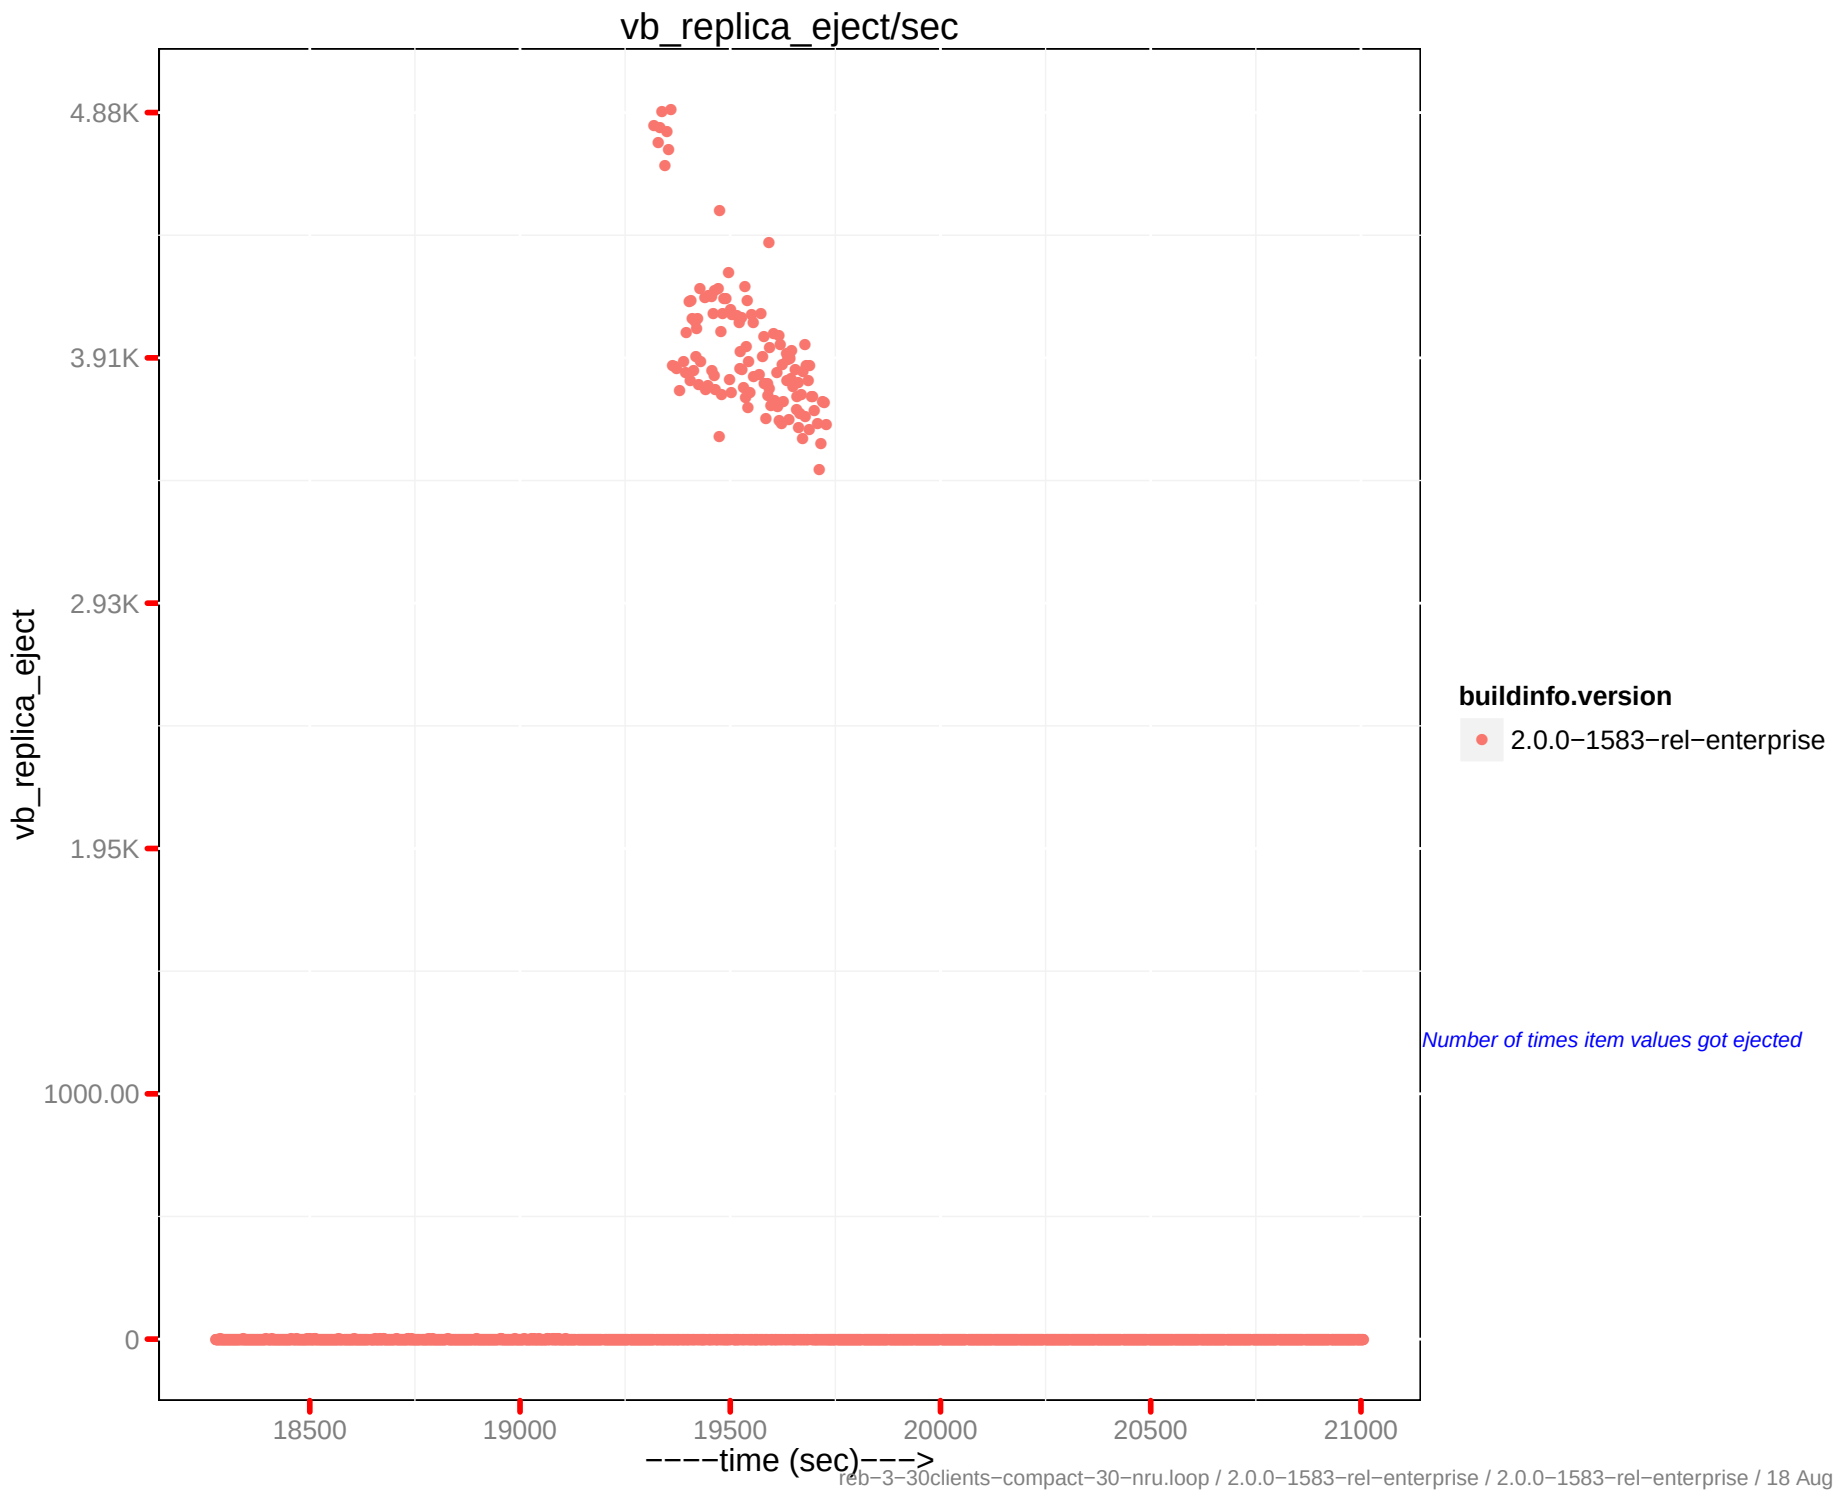
vb_replica_eject/sec
4.88K
3.91K
2.93K
1.95K
1000.00
0
vb_replica_eject
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of times item values got ejected
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

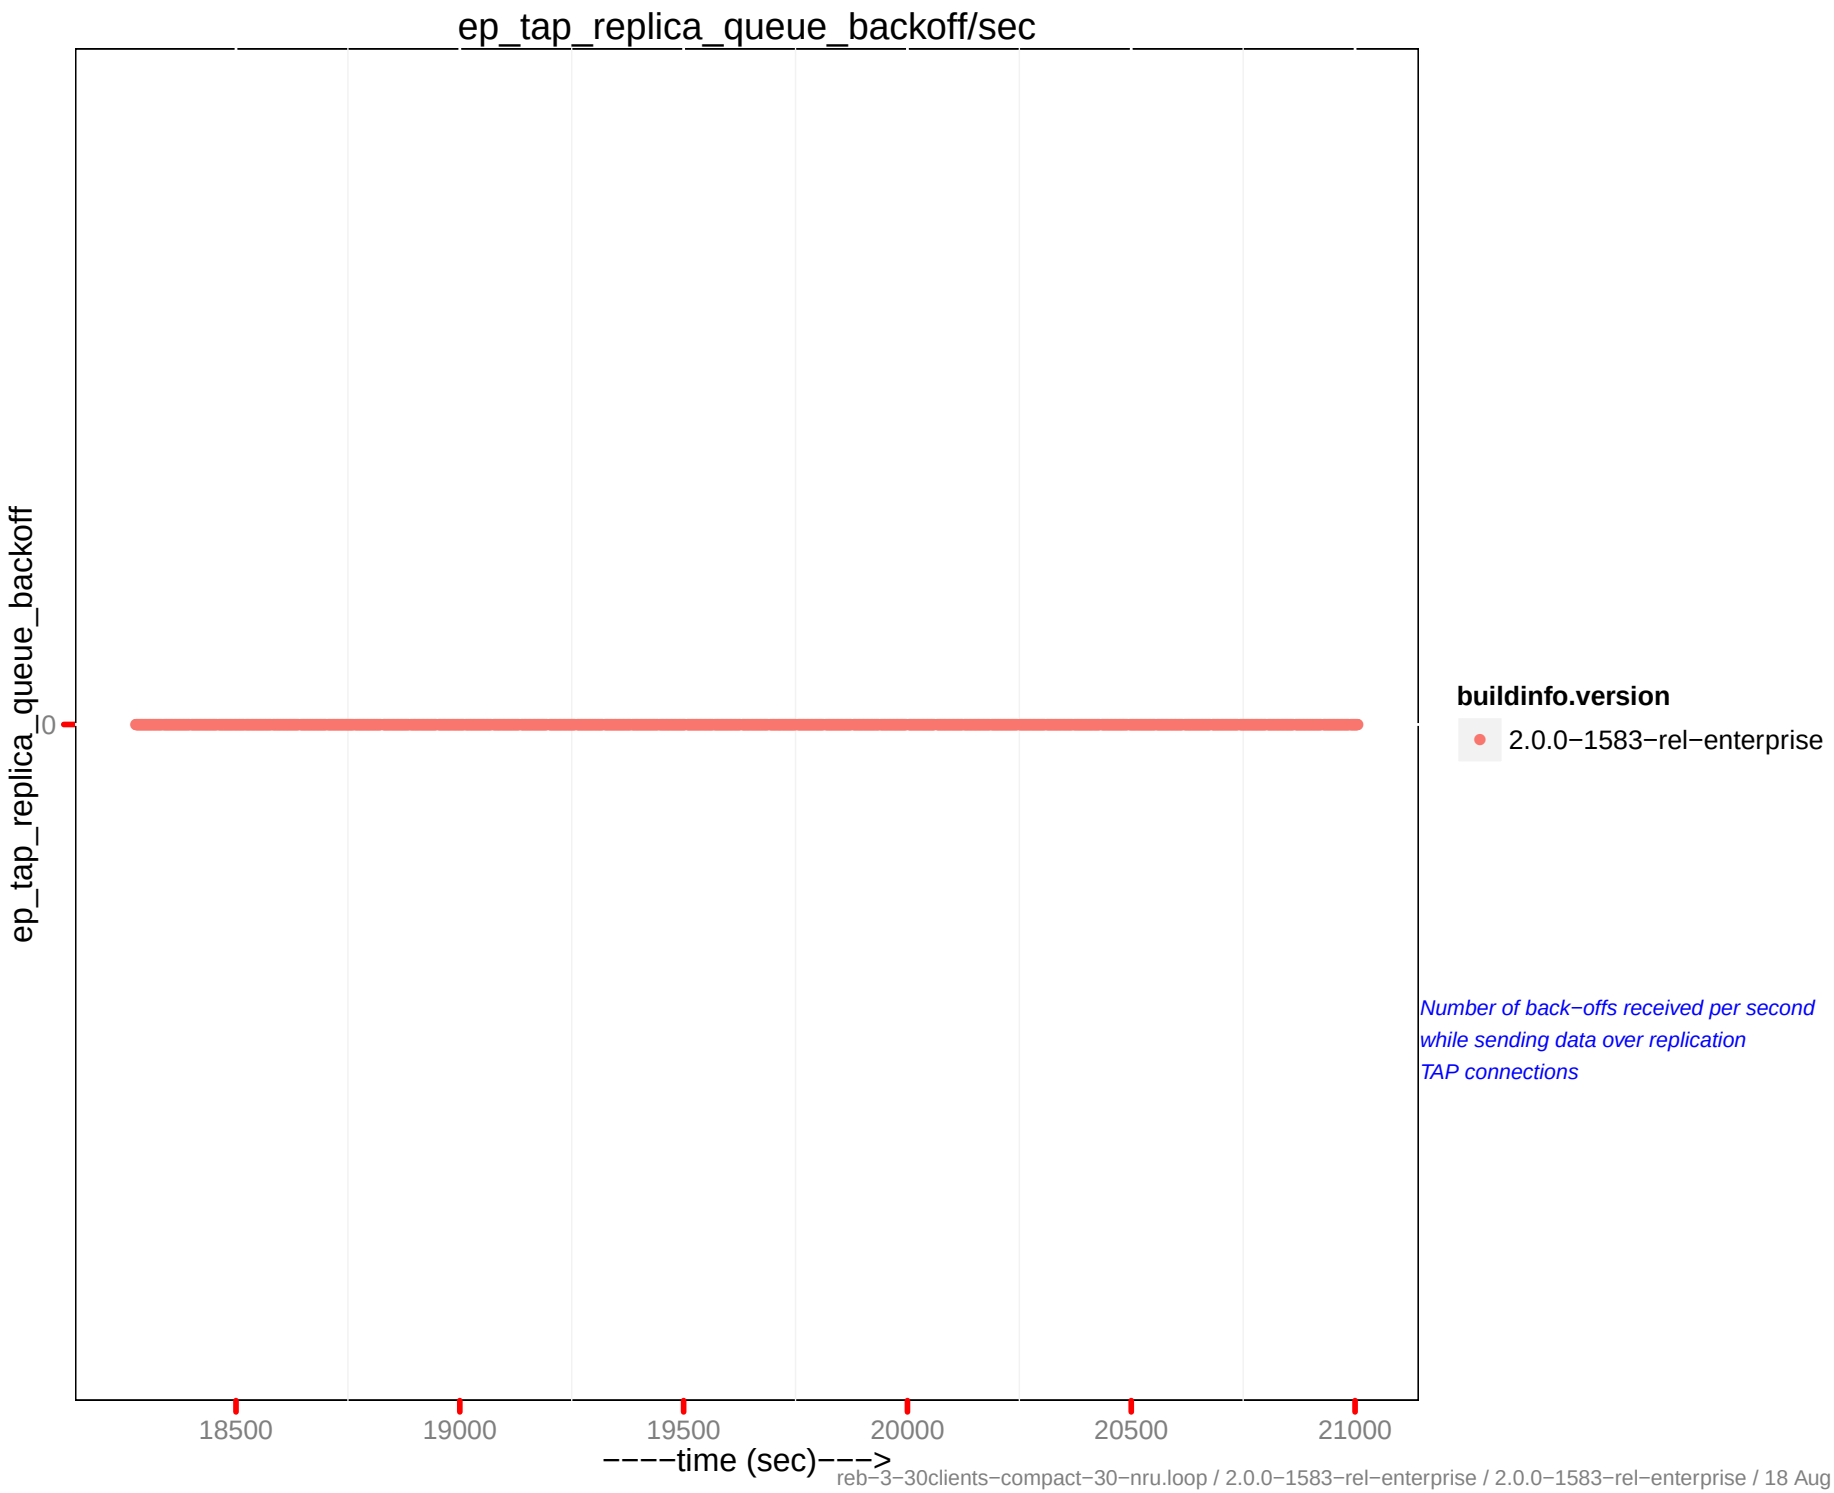
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
0
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of back-offs received per second
while sending data over replication
TAP connections
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

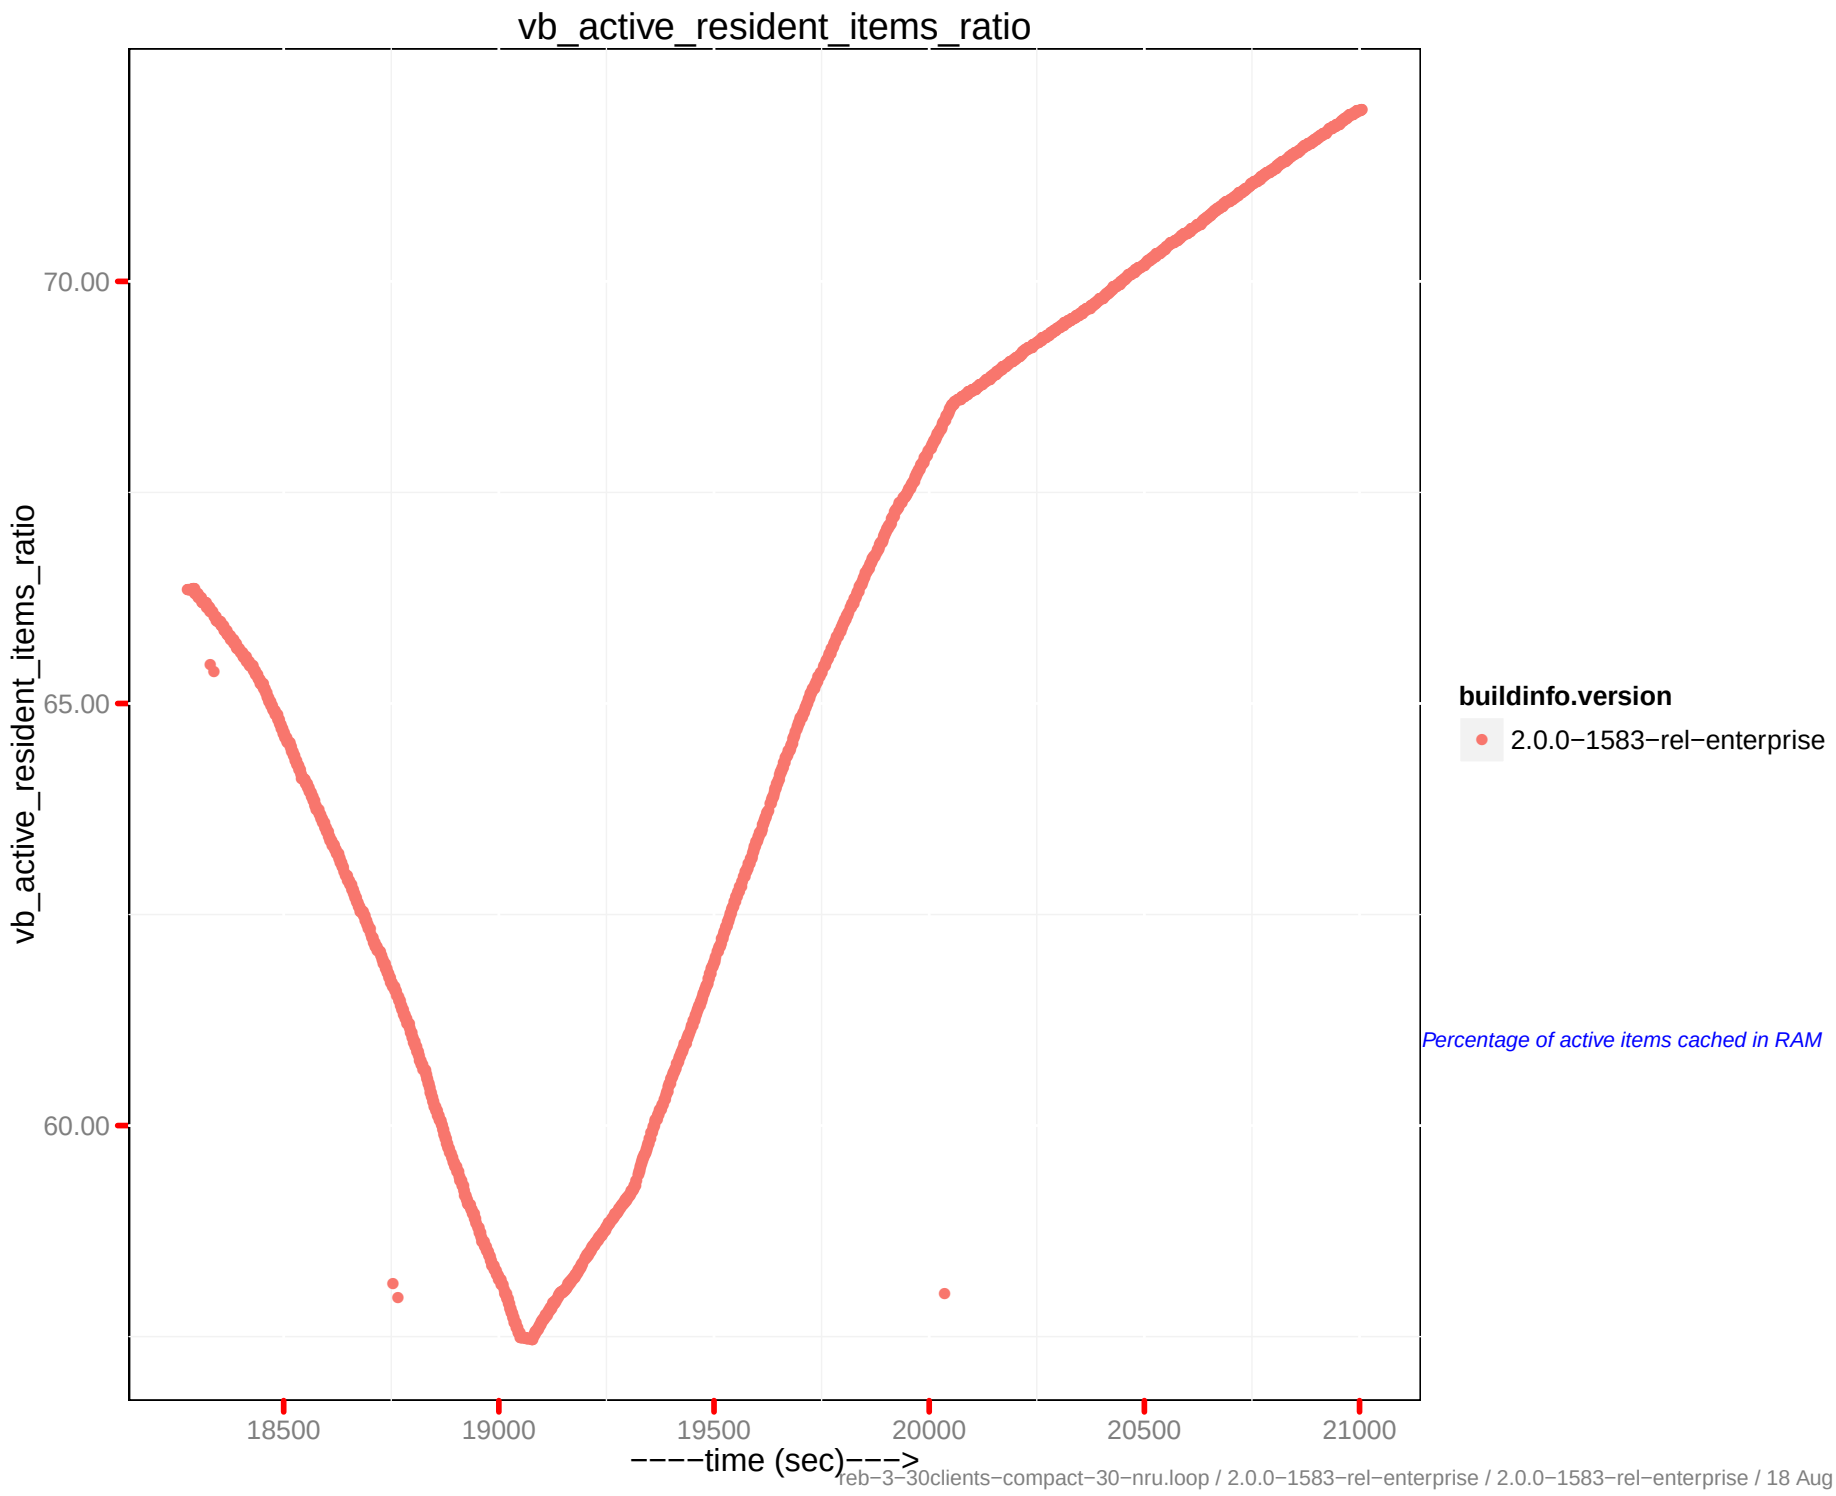
vb_active_resident_items_ratio
70.00
65.00
60.00
vb_active_resident_items_ratio
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Percentage of active items cached in RAM
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

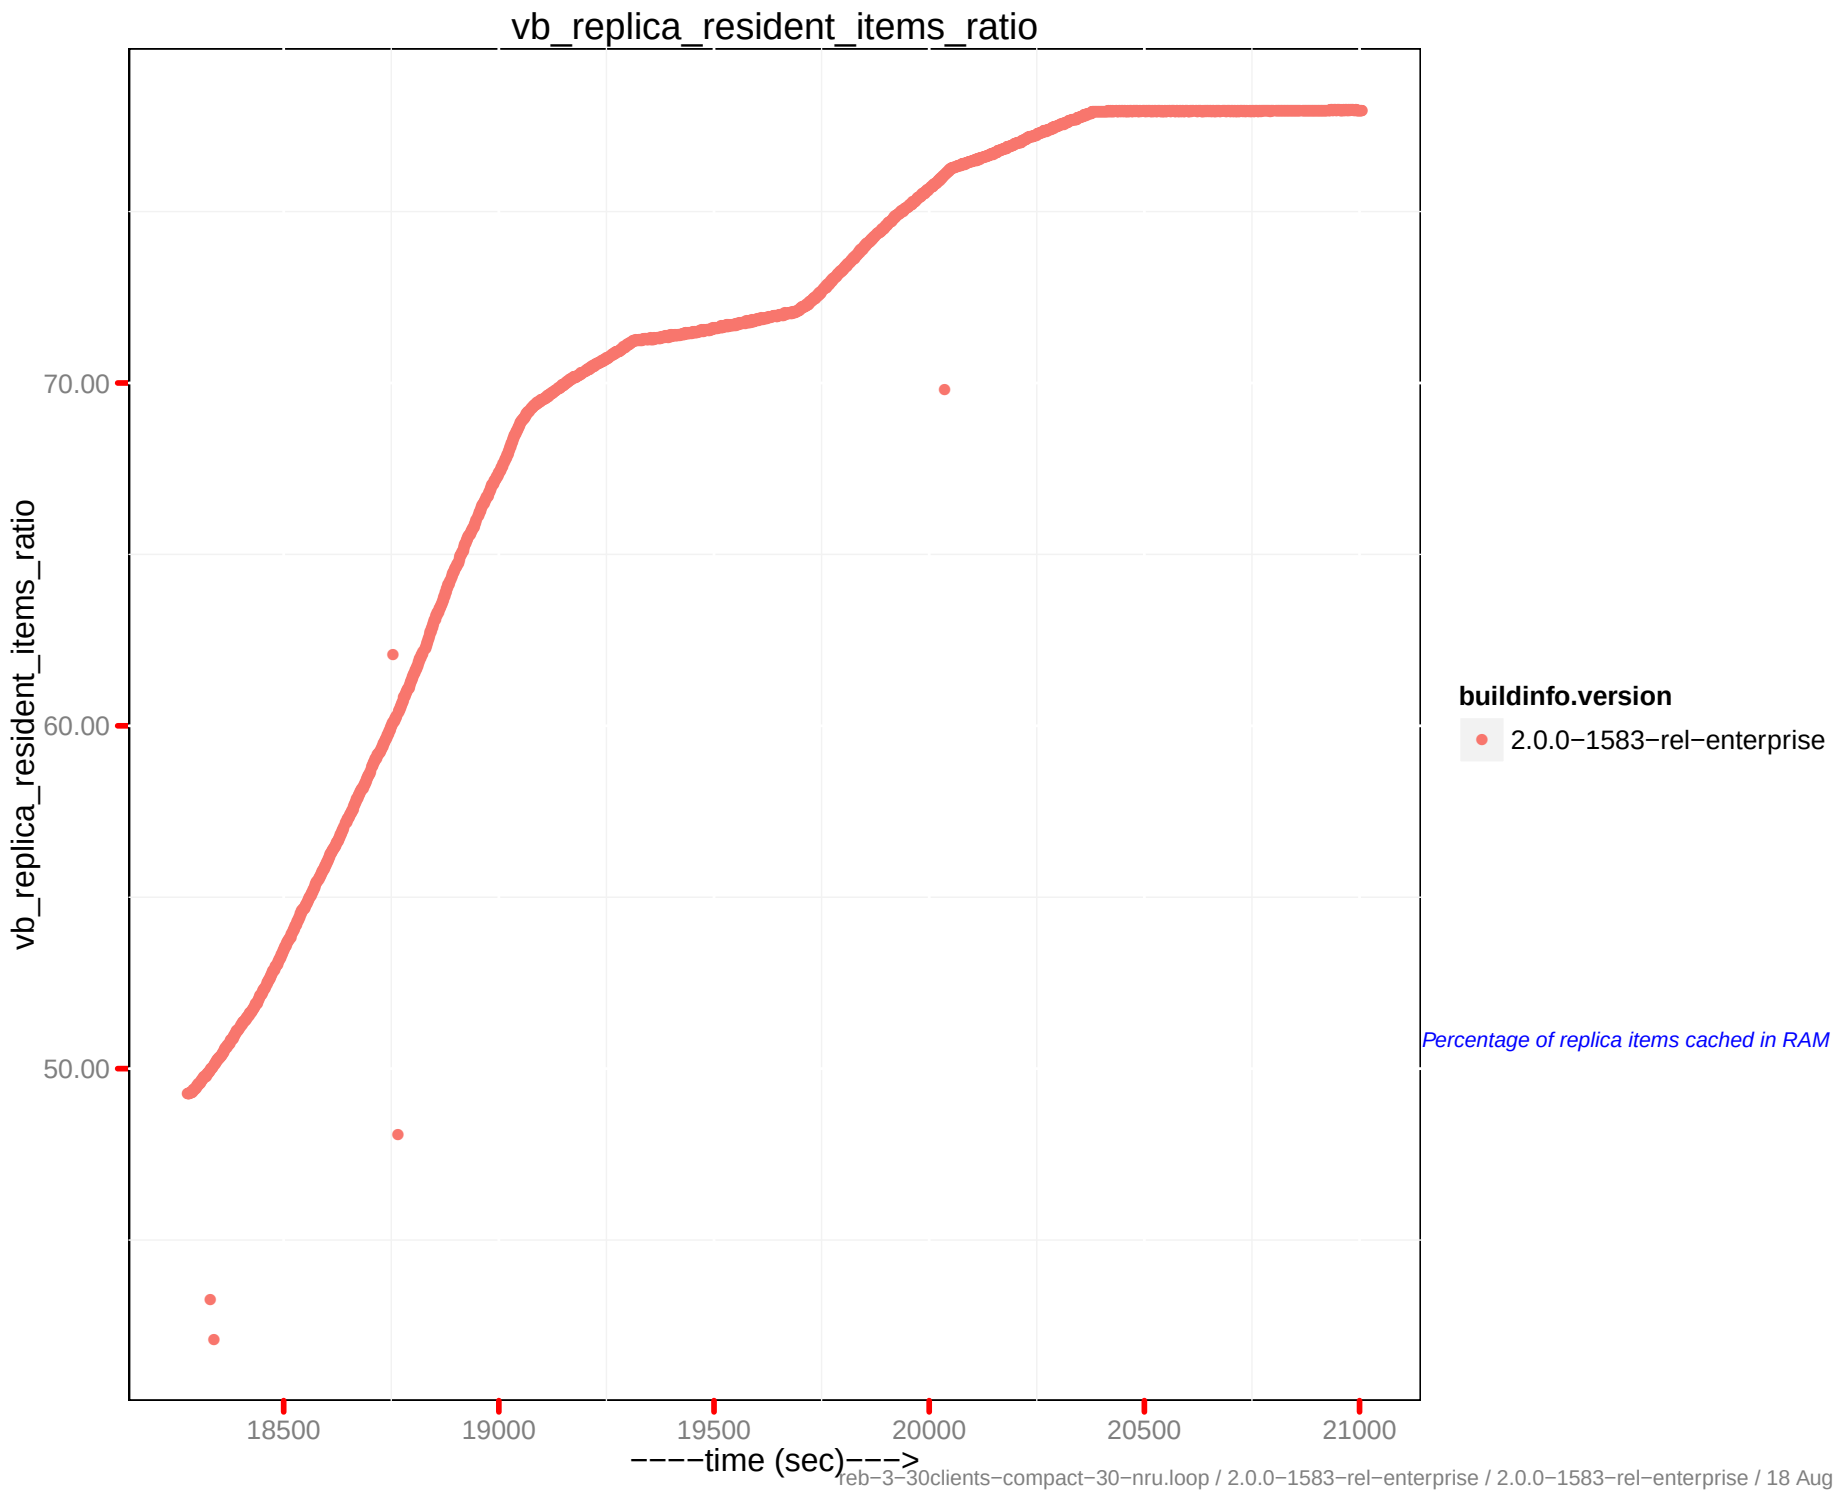
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
70.00
60.00
50.00
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Percentage of replica items cached in RAM
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

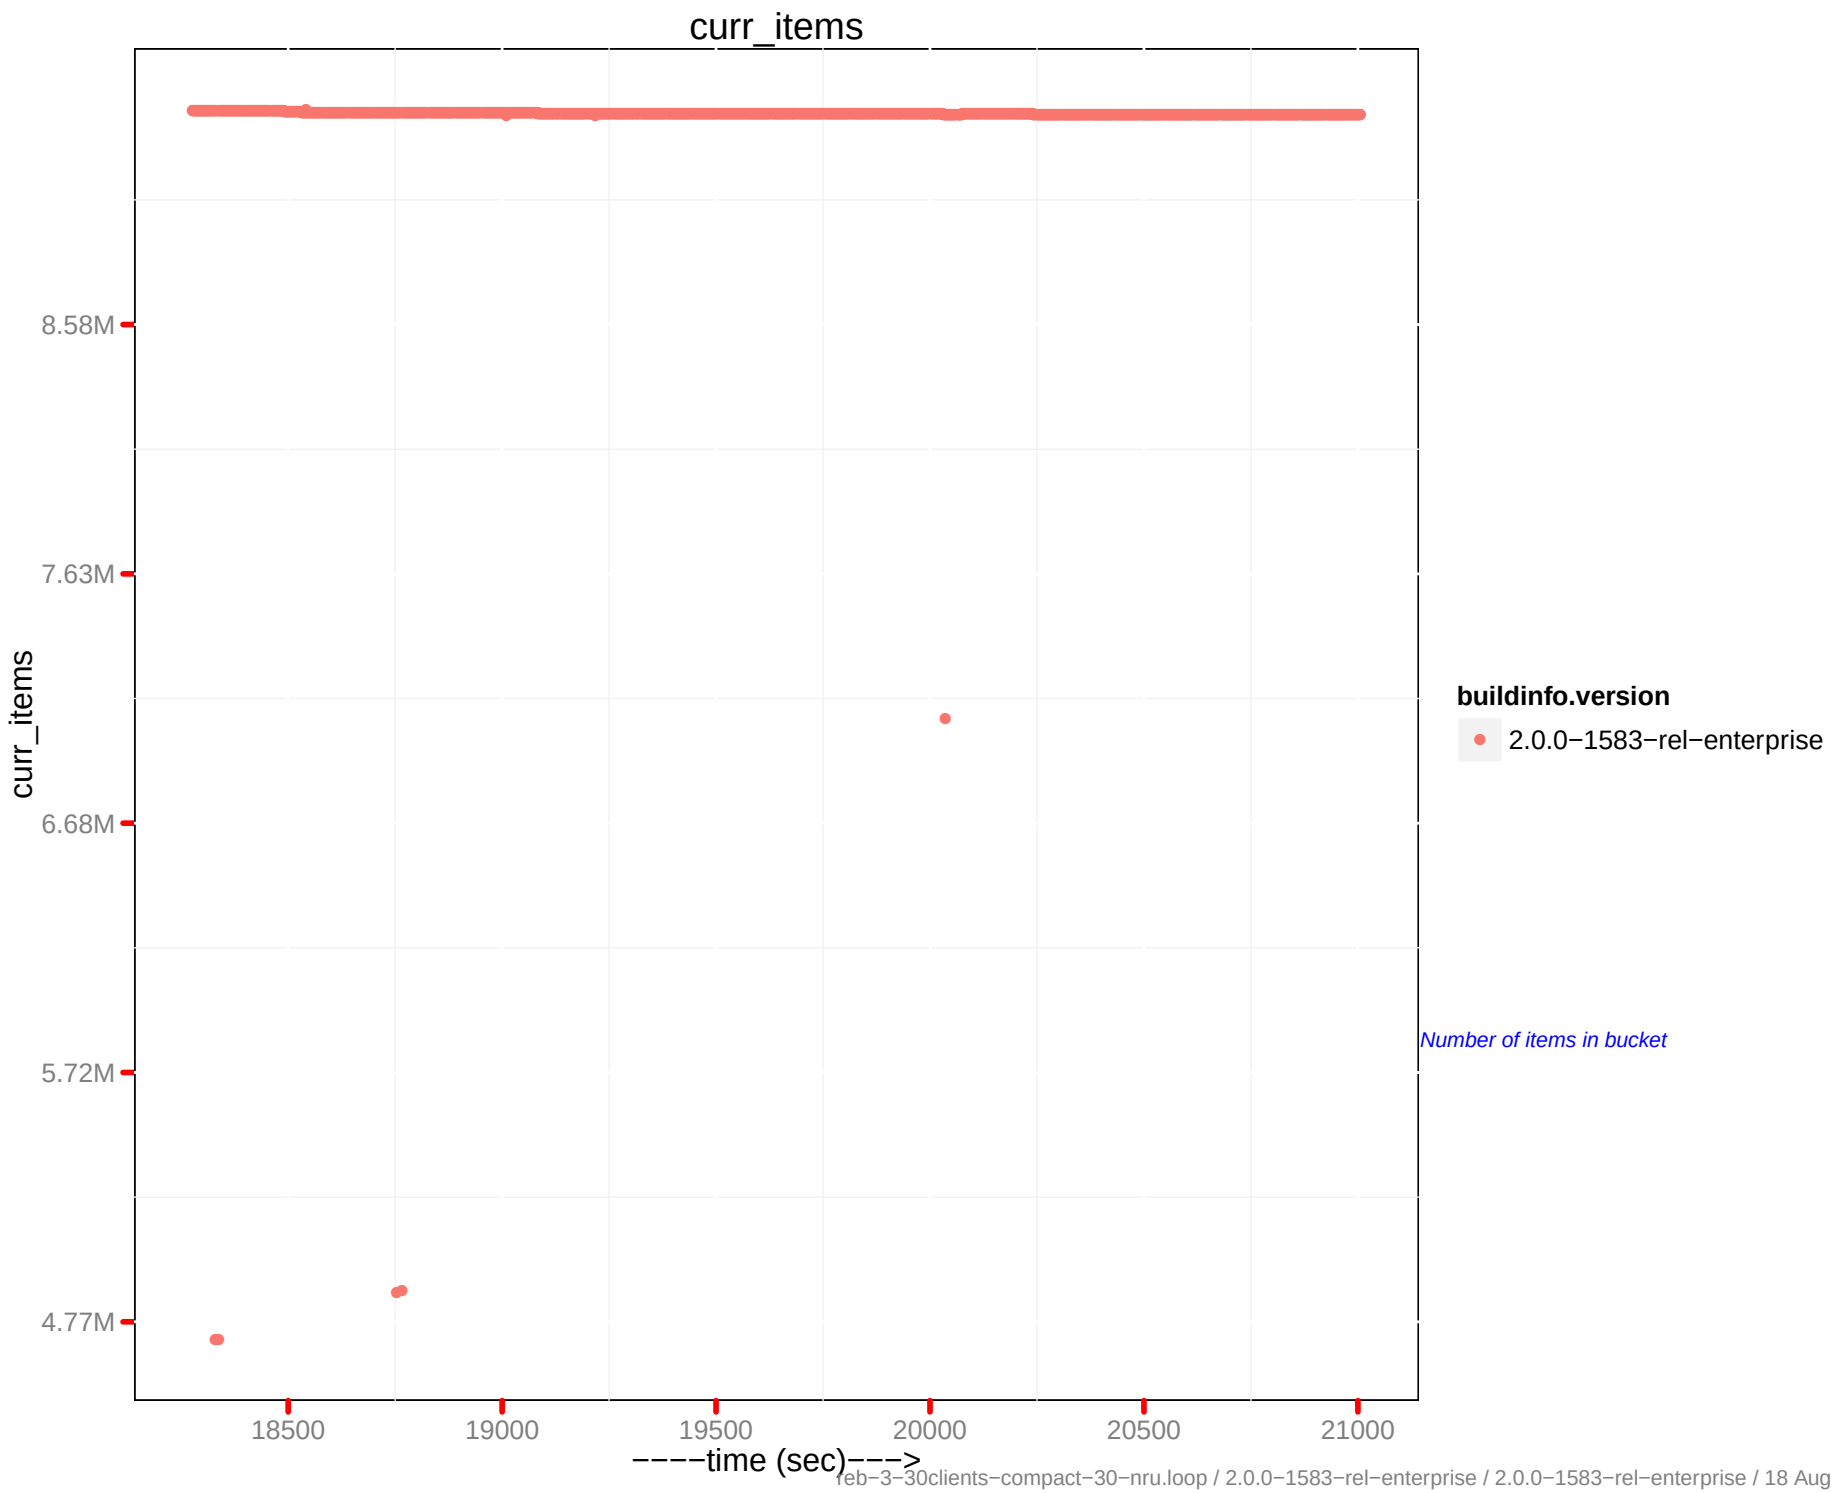
curr_items
8.58M
7.63M
6.68M
5.72M
4.77M
curr_items
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of items in bucket
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

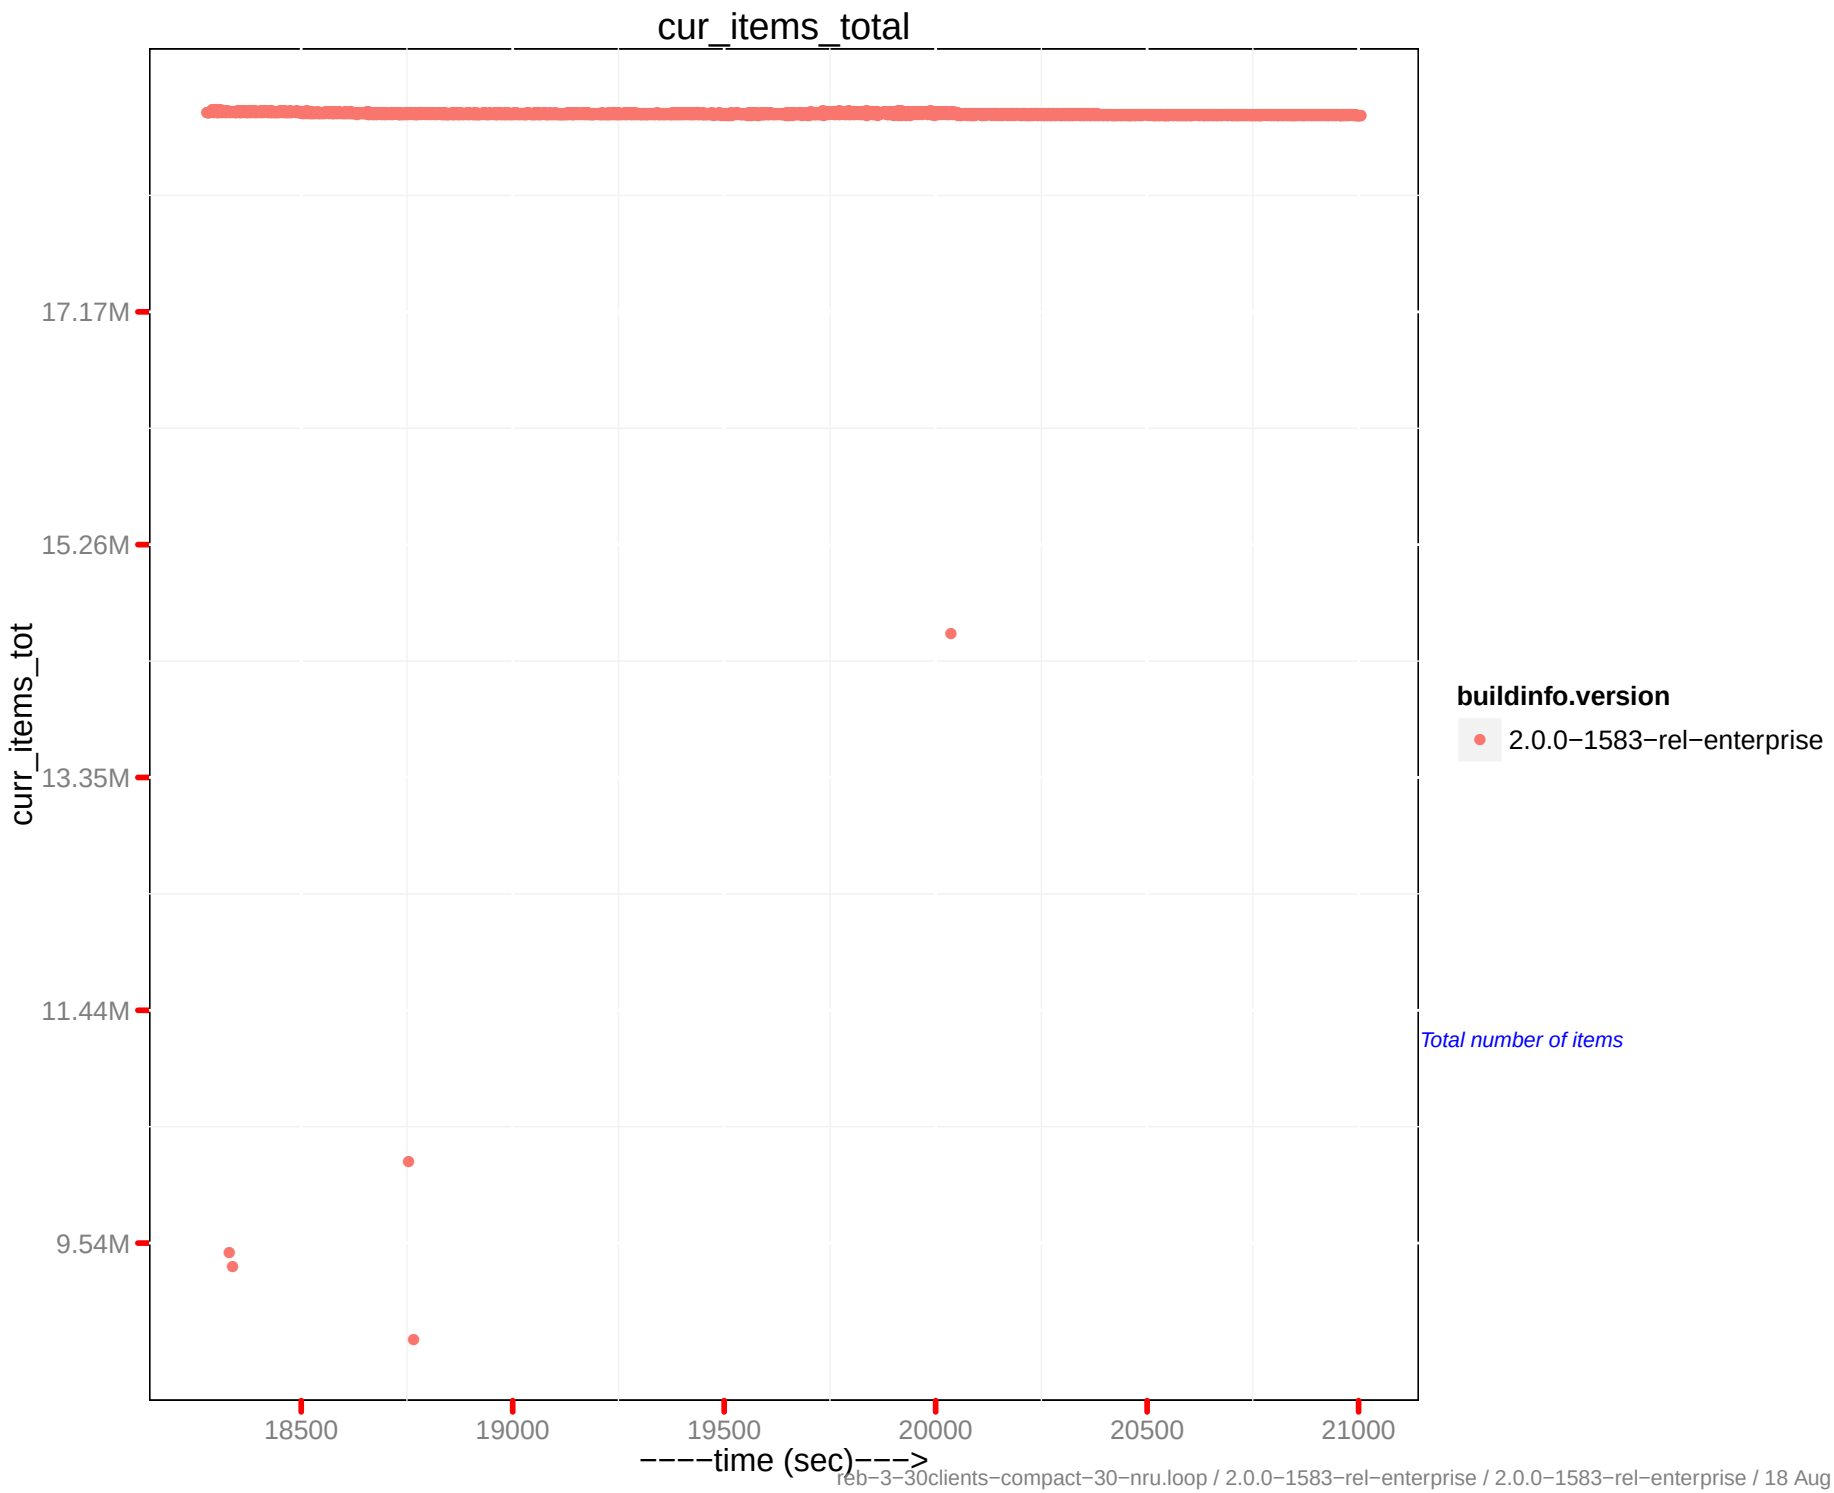
cur_items_total
curr_items_tot
17.17M
15.26M
13.35M
11.44M
9.54M
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Total number of items
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

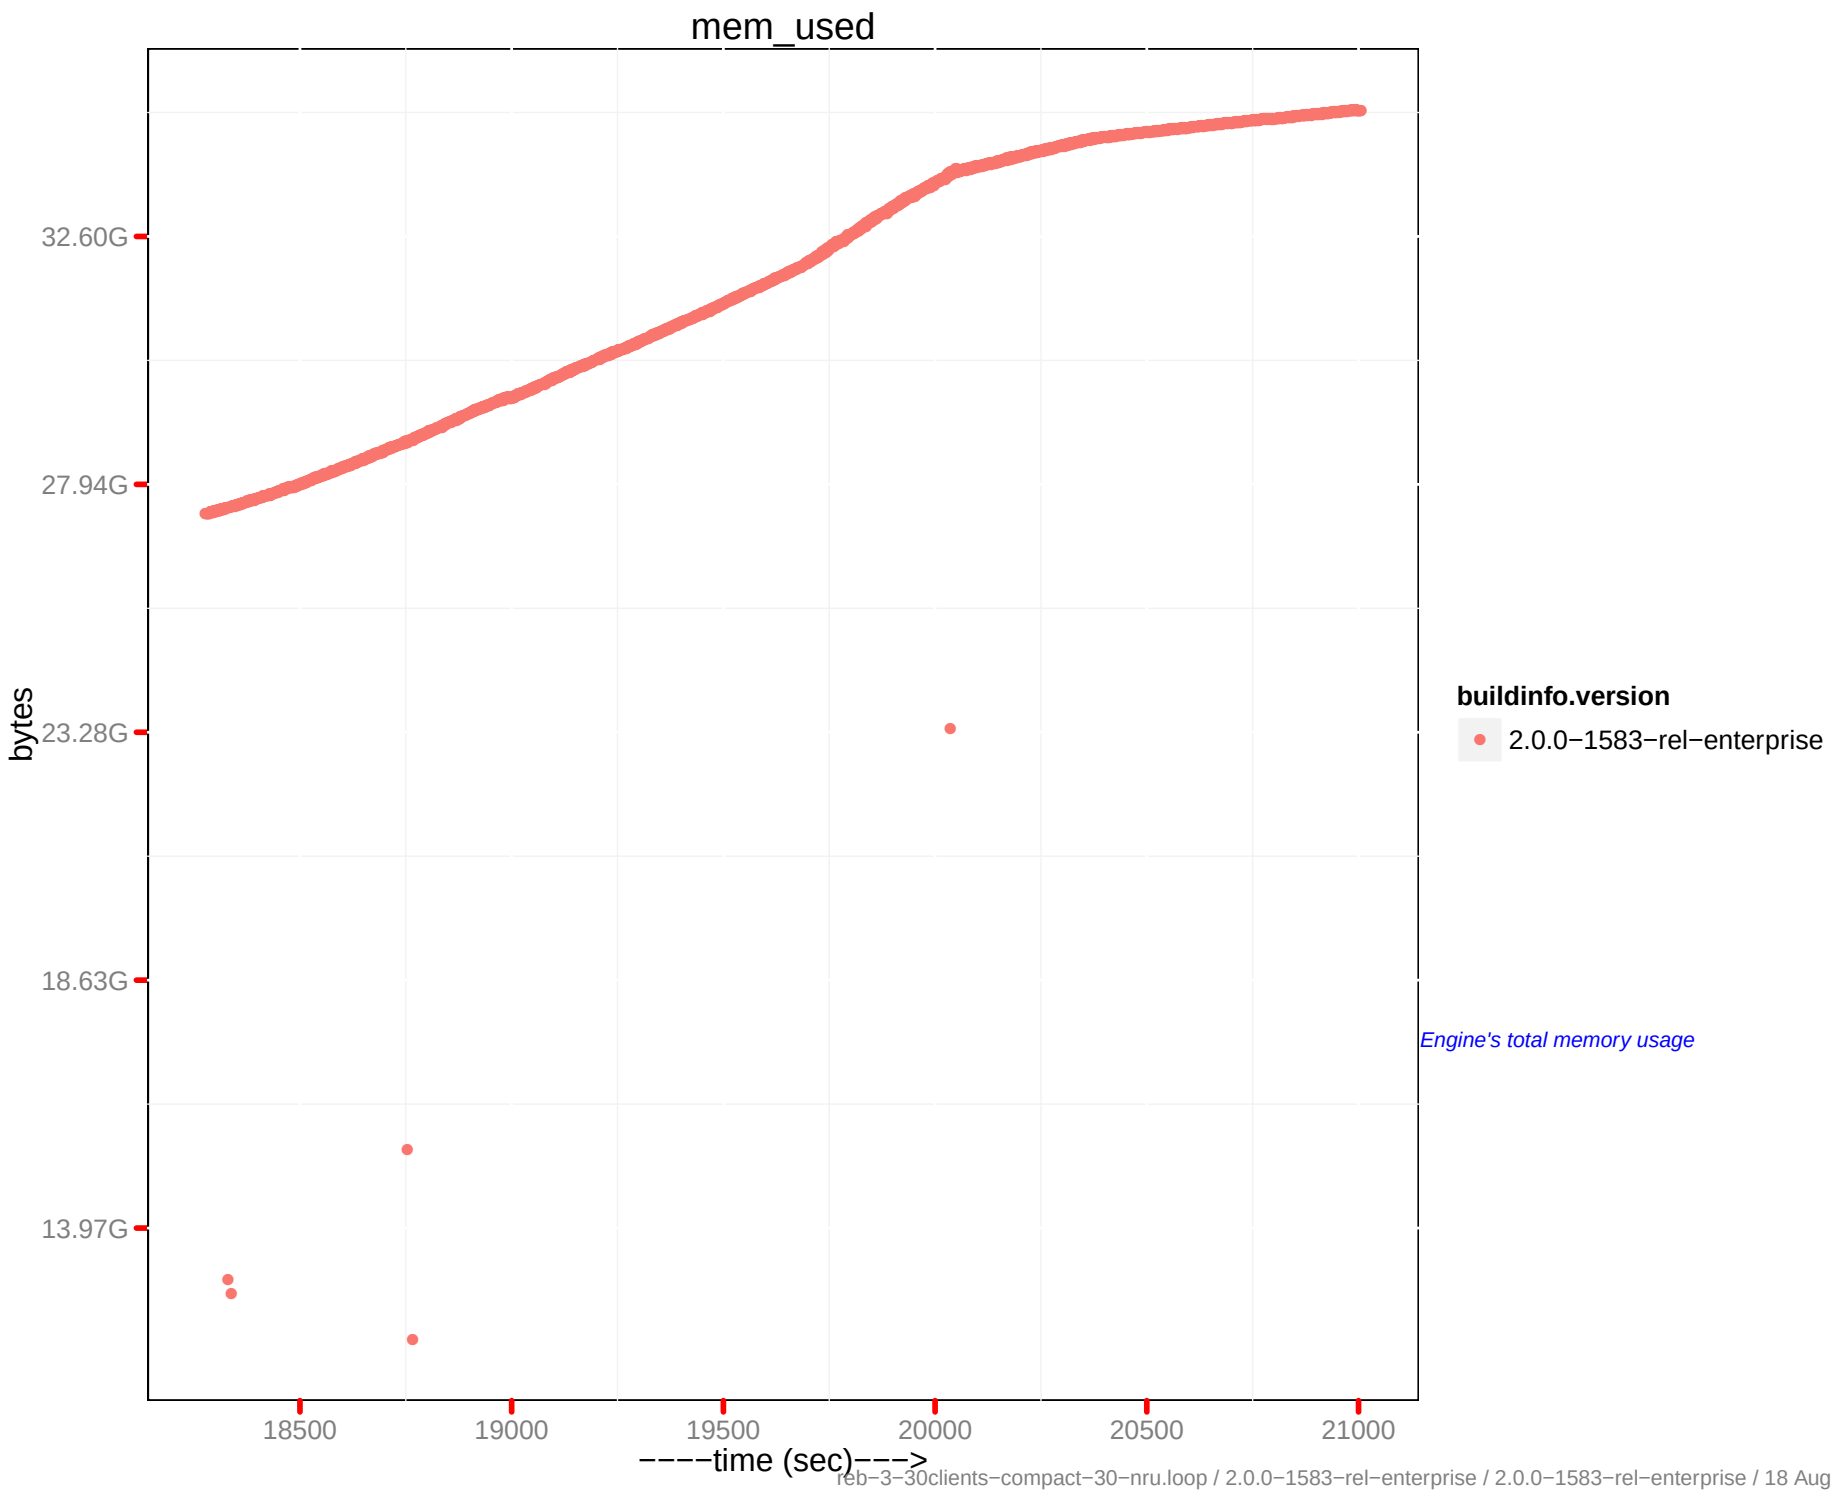
mem_used
32.60G
27.94G
23.28G
18.63G
13.97G
bytes
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Engine's total memory usage
feb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

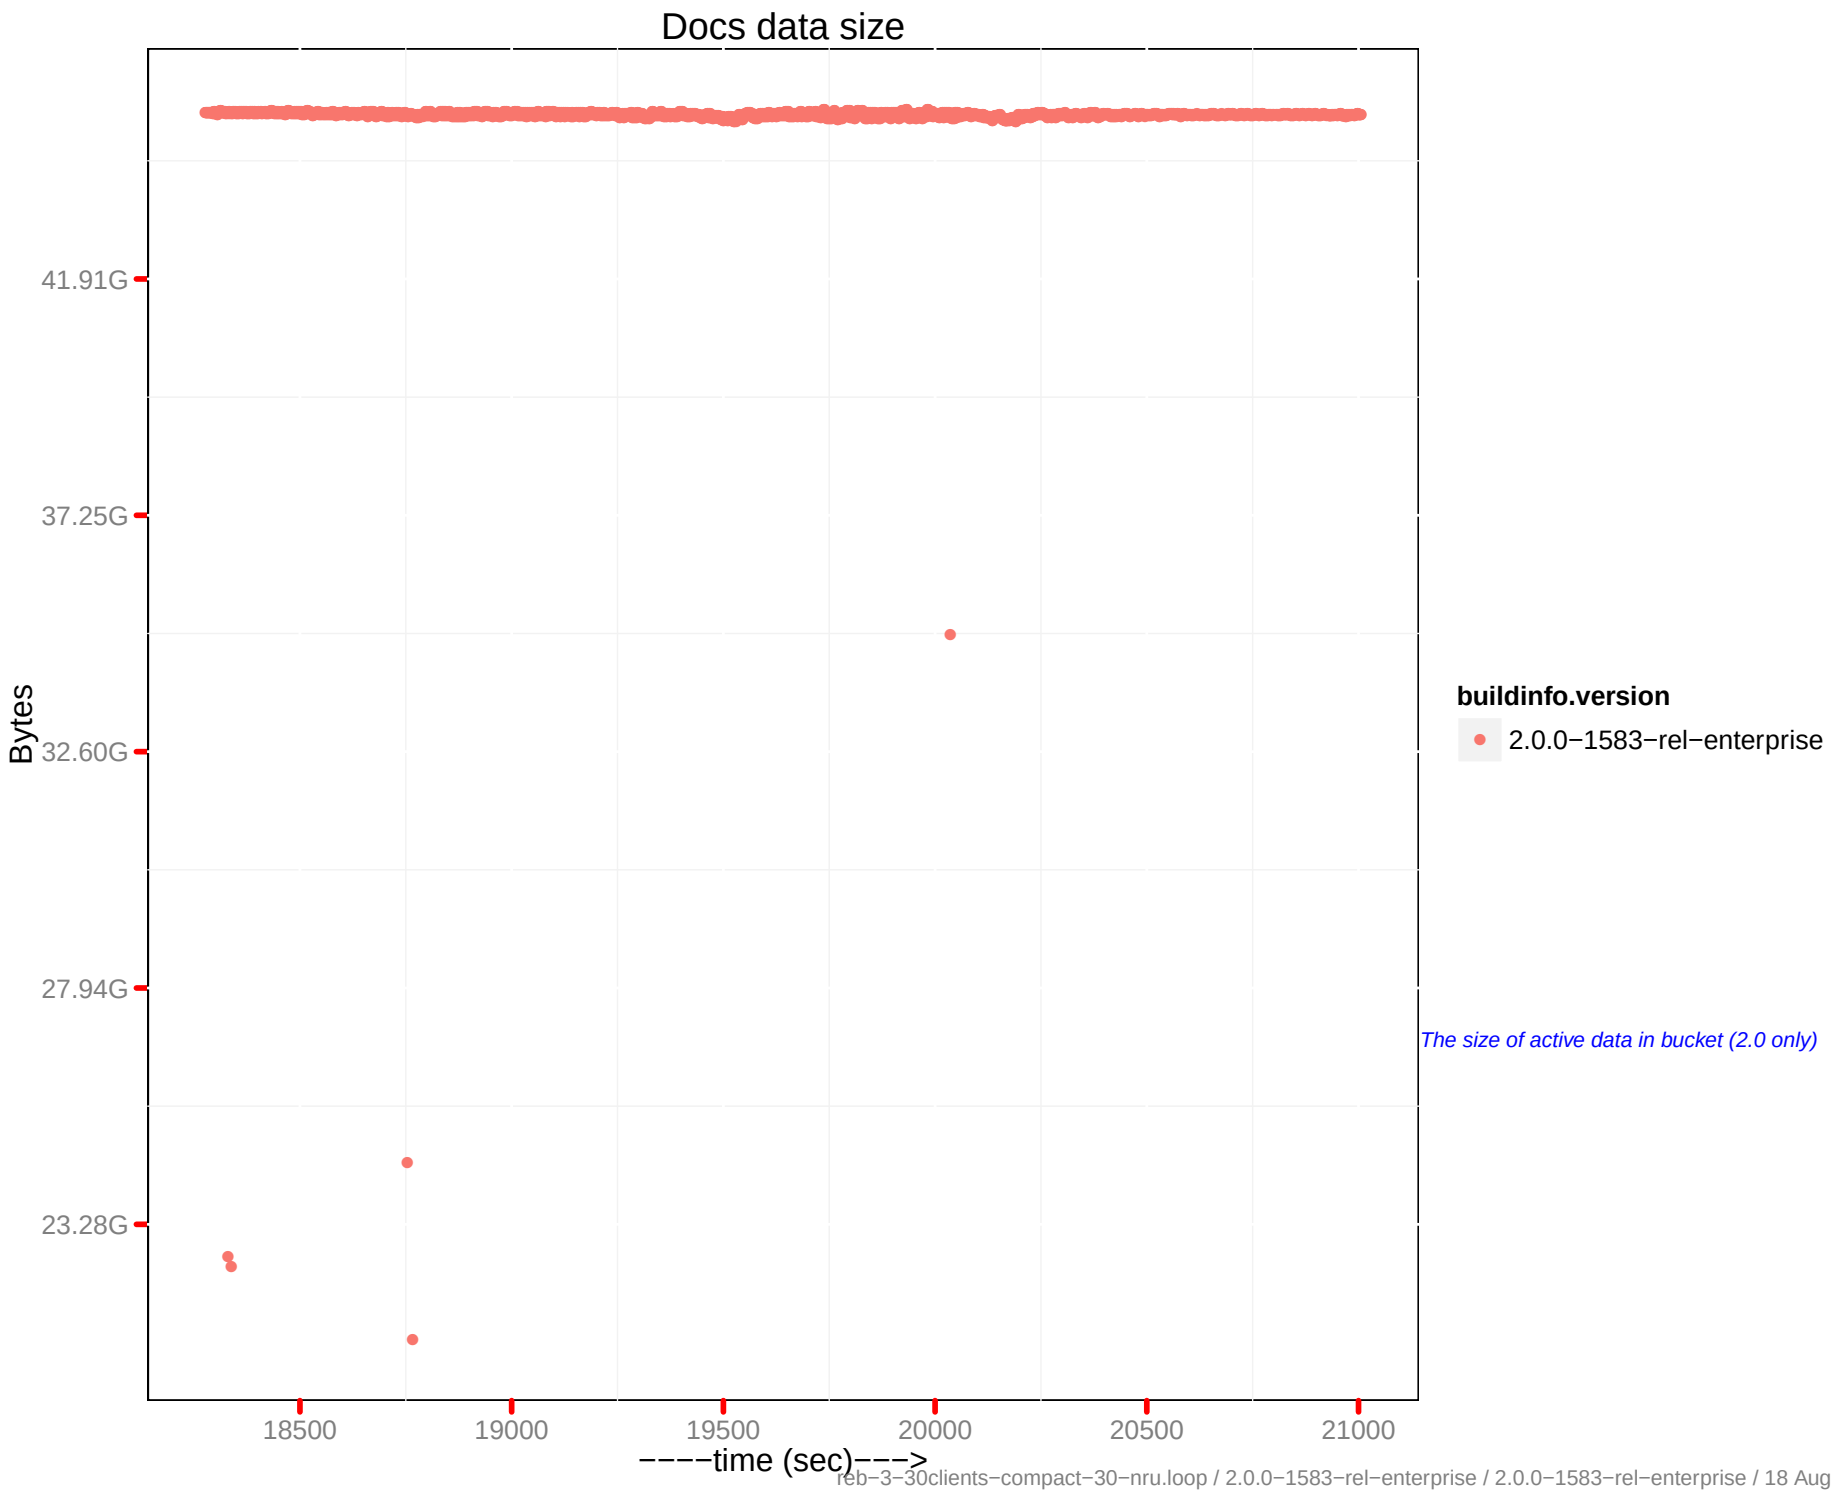
Docs data size
41.91G
37.25G
32.60G
27.94G
23.28G
Bytes
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
The size of active data in bucket (2.0 only)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

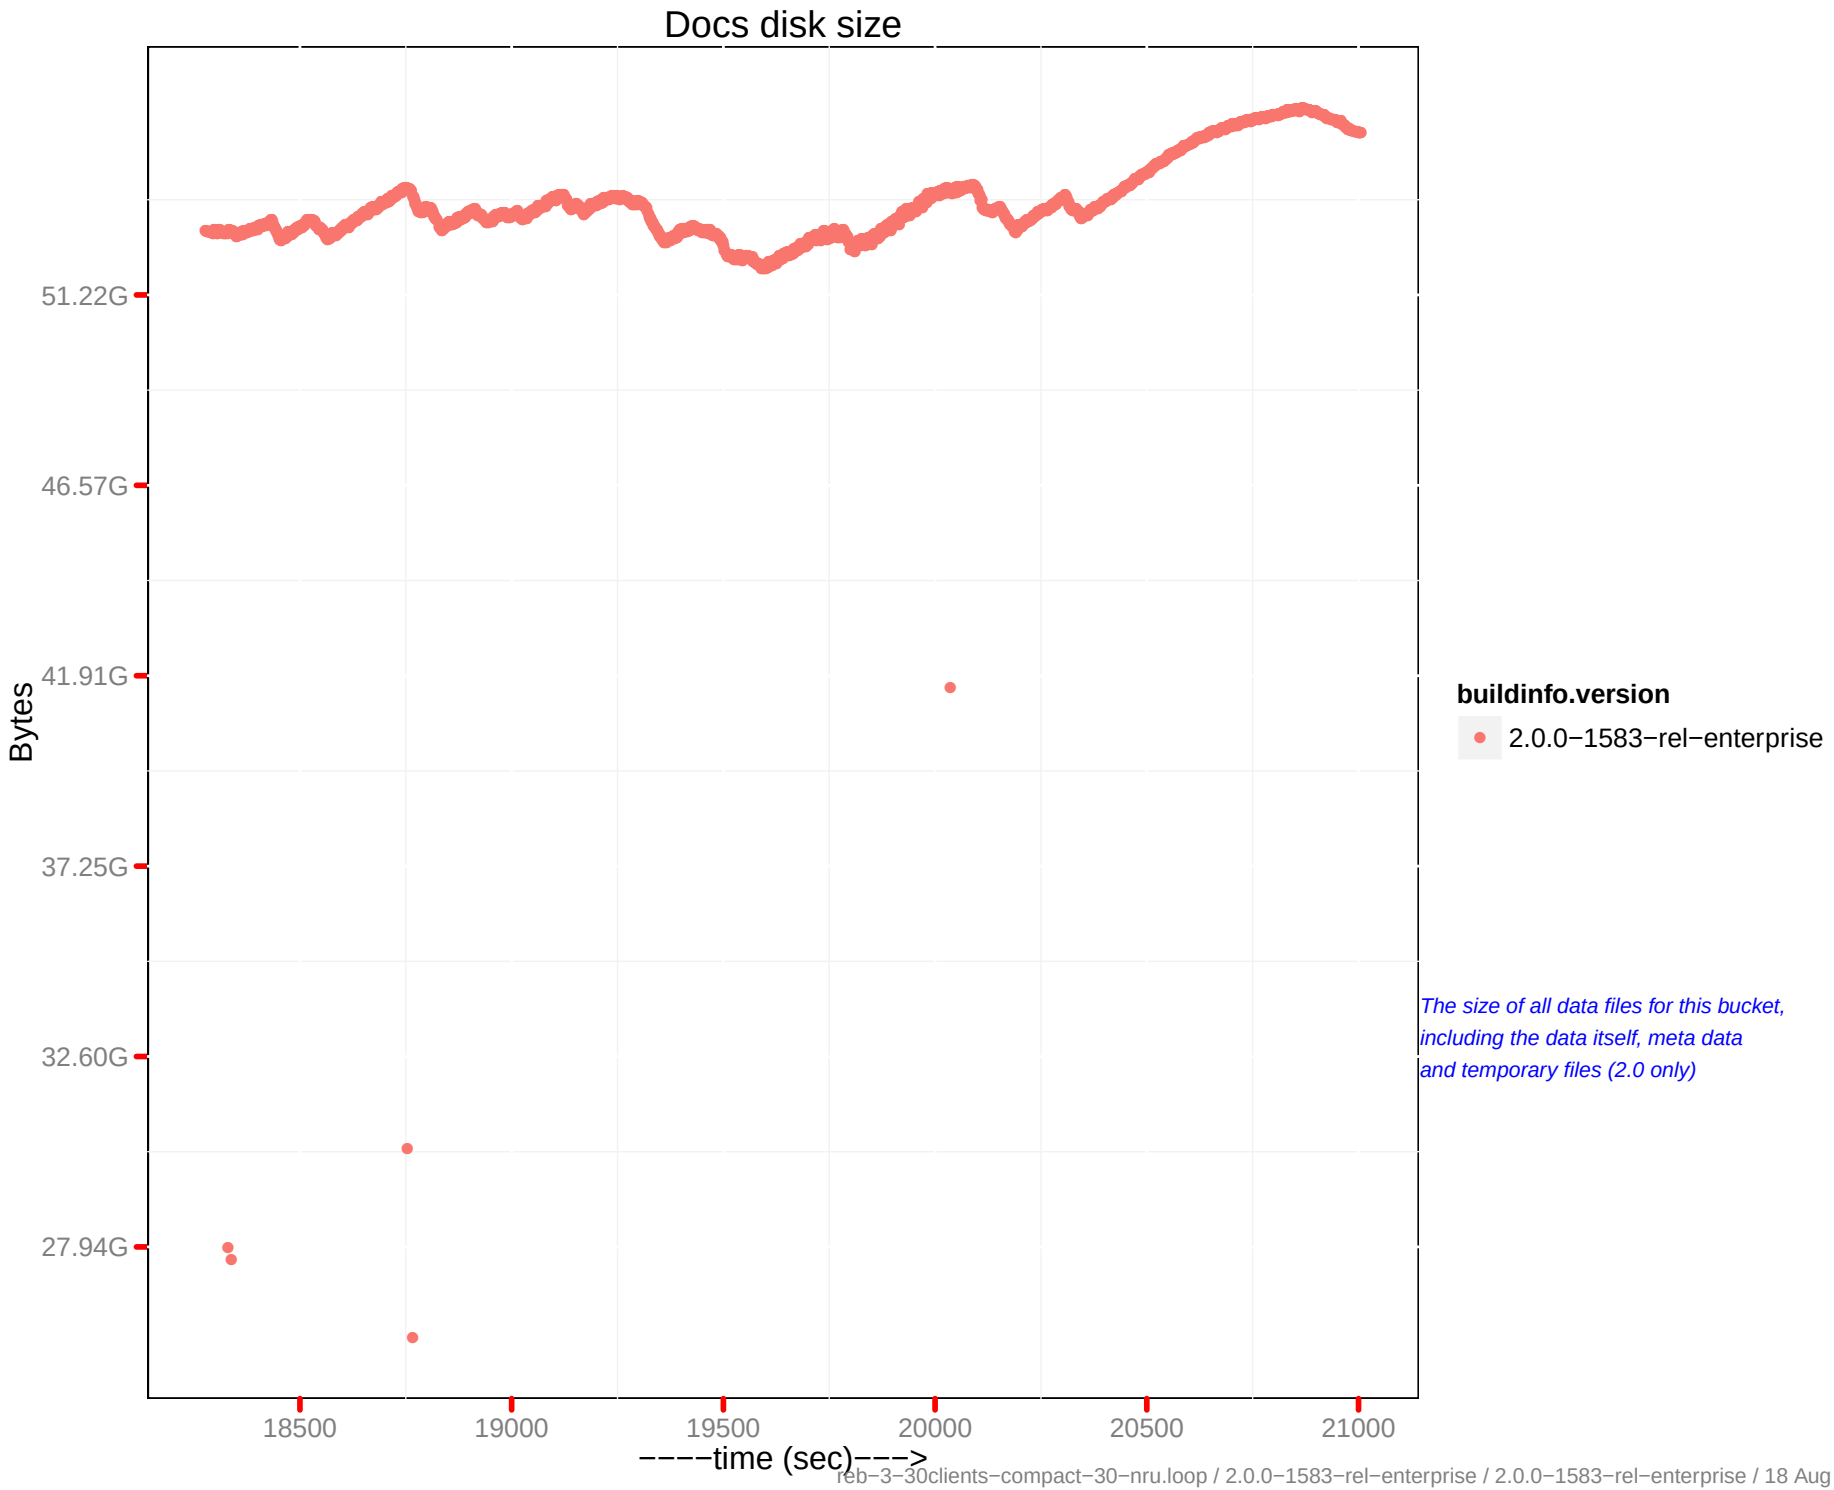
Docs disk size
51.22G
46.57G
41.91G
37.25G
32.60G
27.94G
Bytes
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

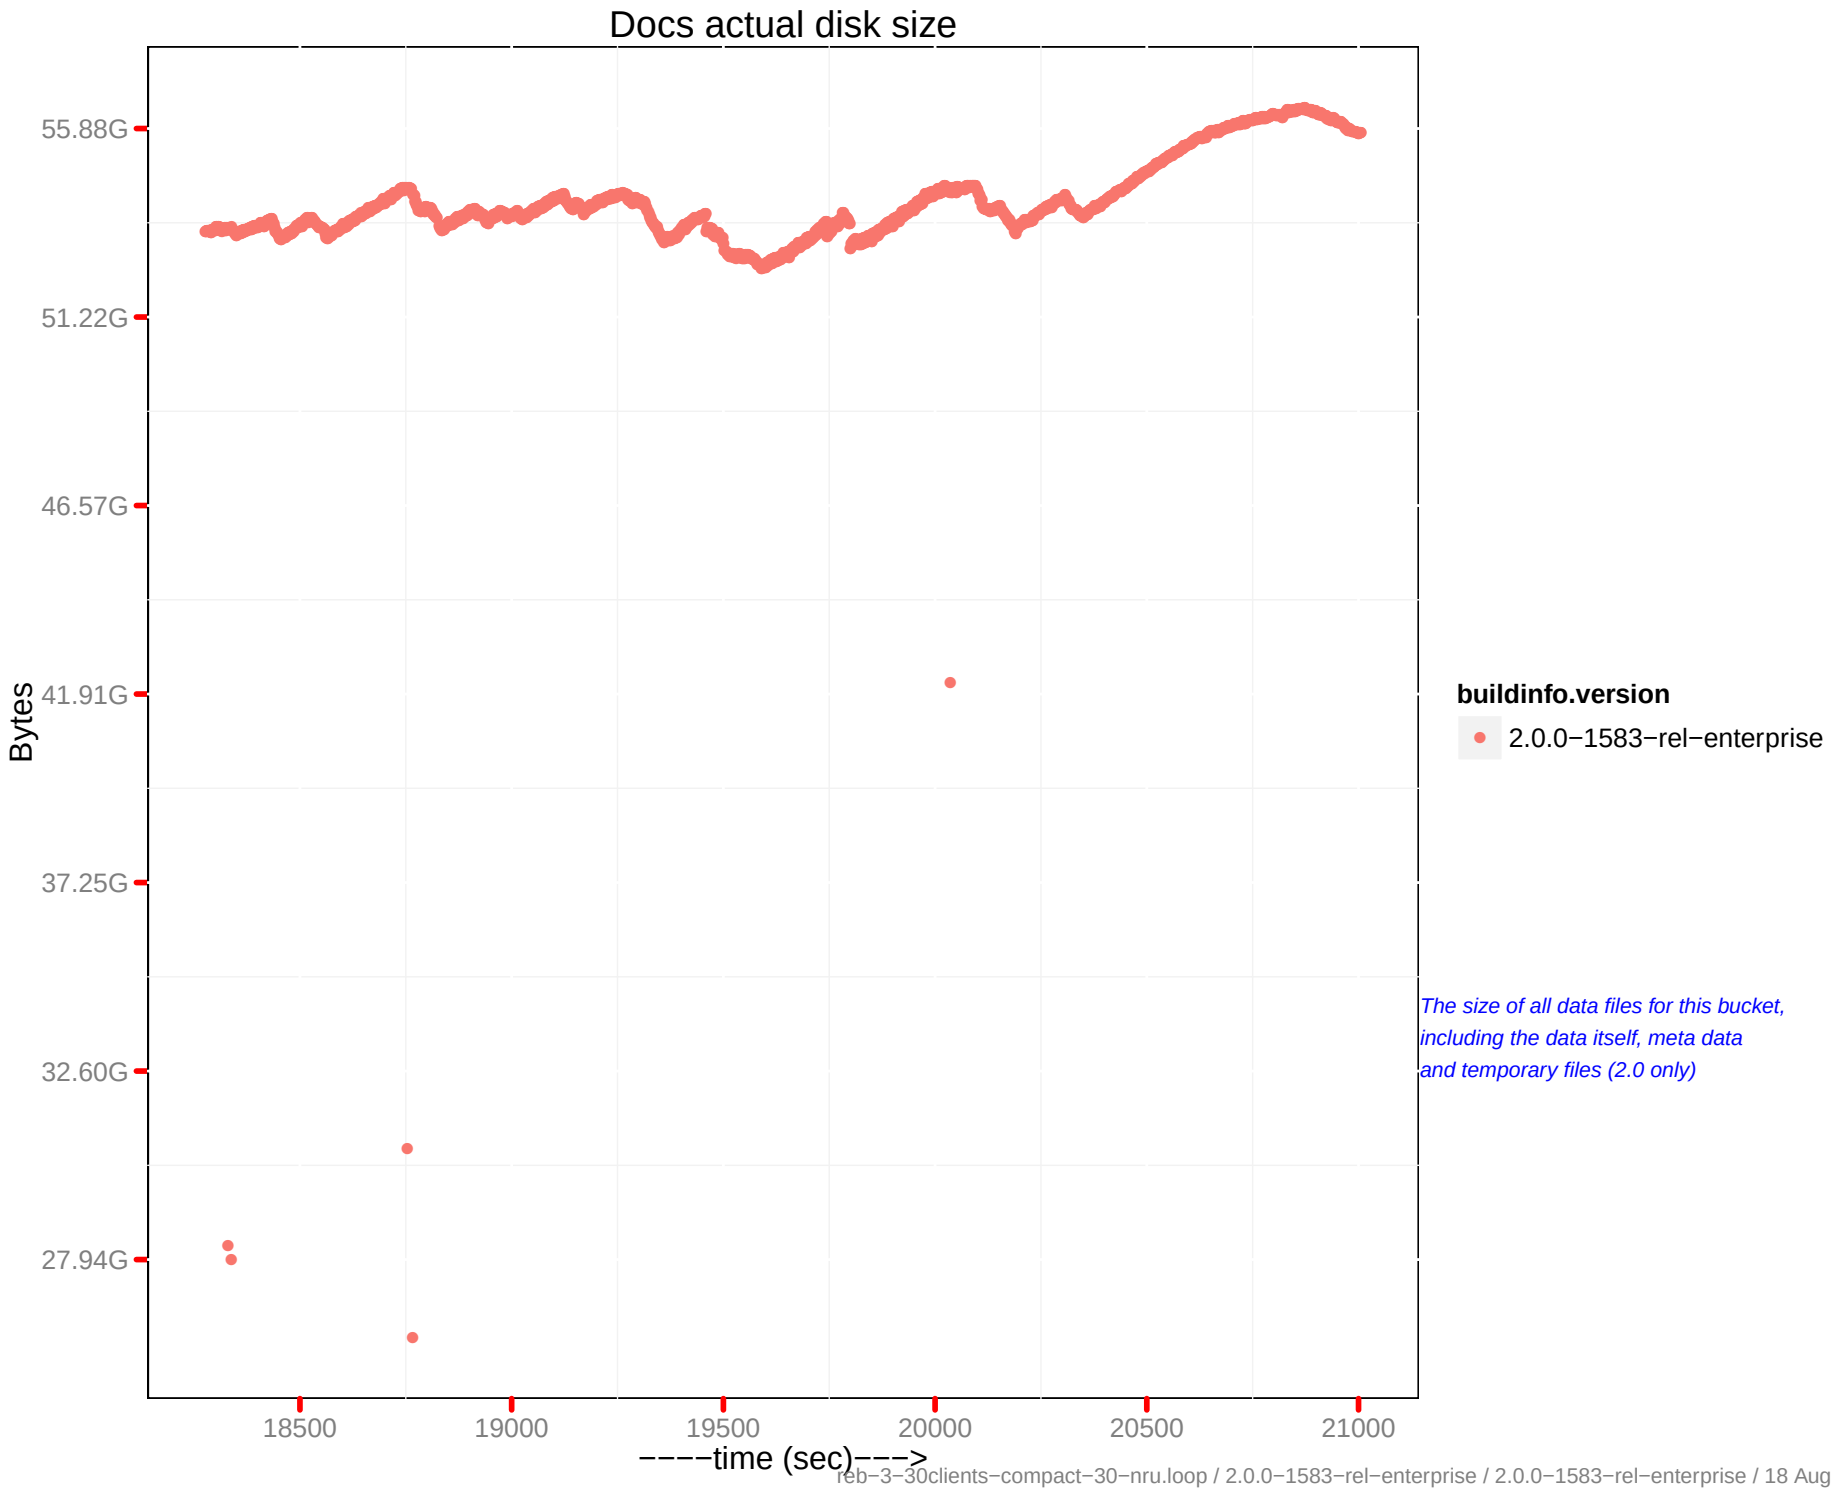
Docs actual disk size
55.88G
51.22G
46.57G
41.91G
37.25G
32.60G
27.94G
Bytes
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

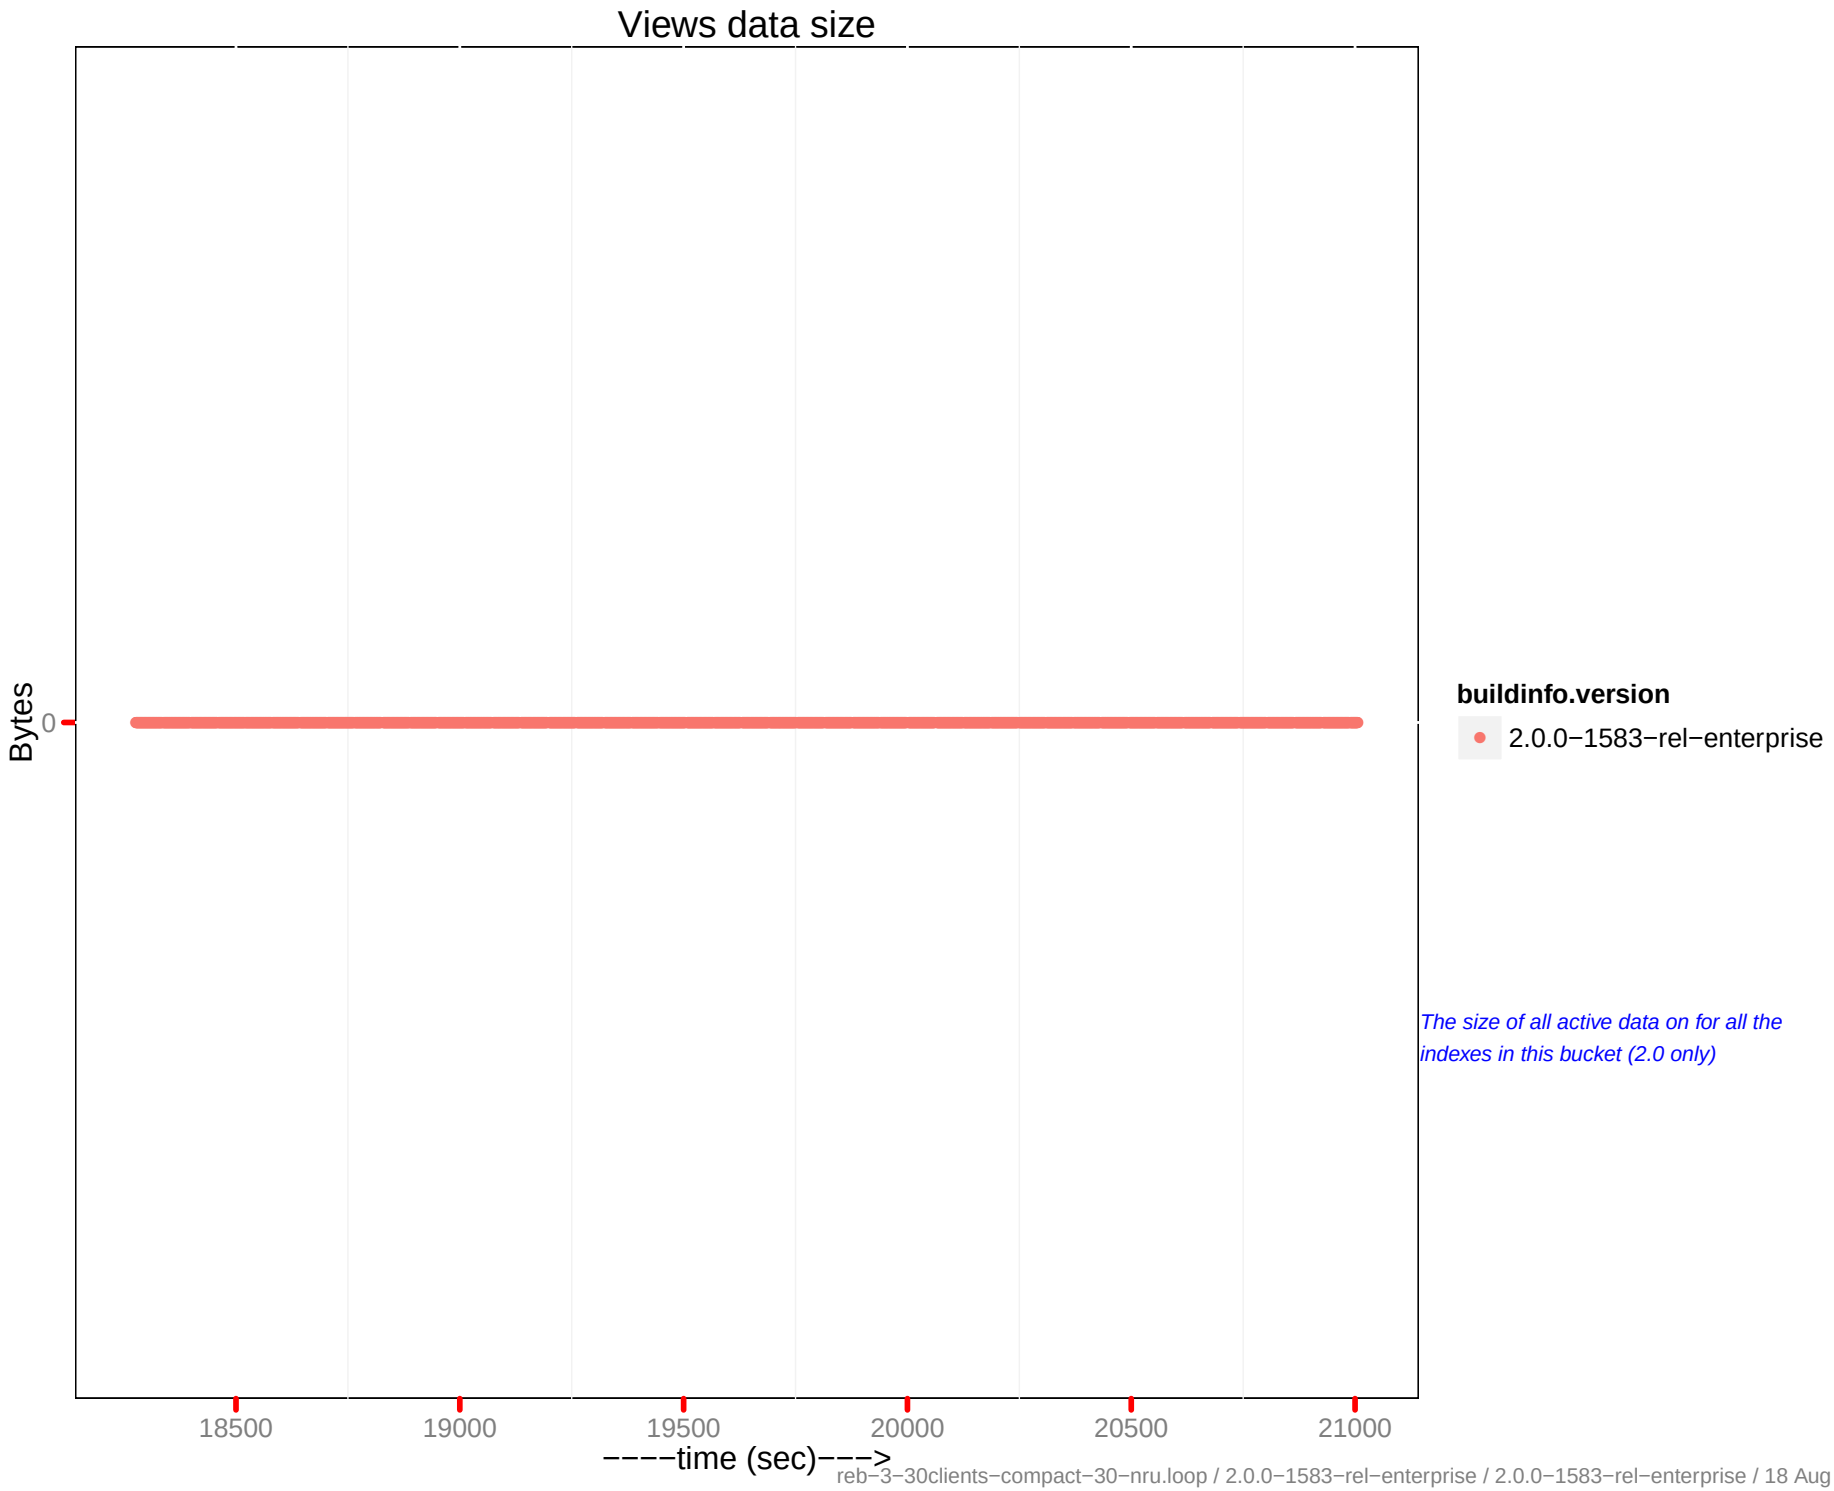
Views data size
Bytes
0
18500
19000
19500
20000
20500
21000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

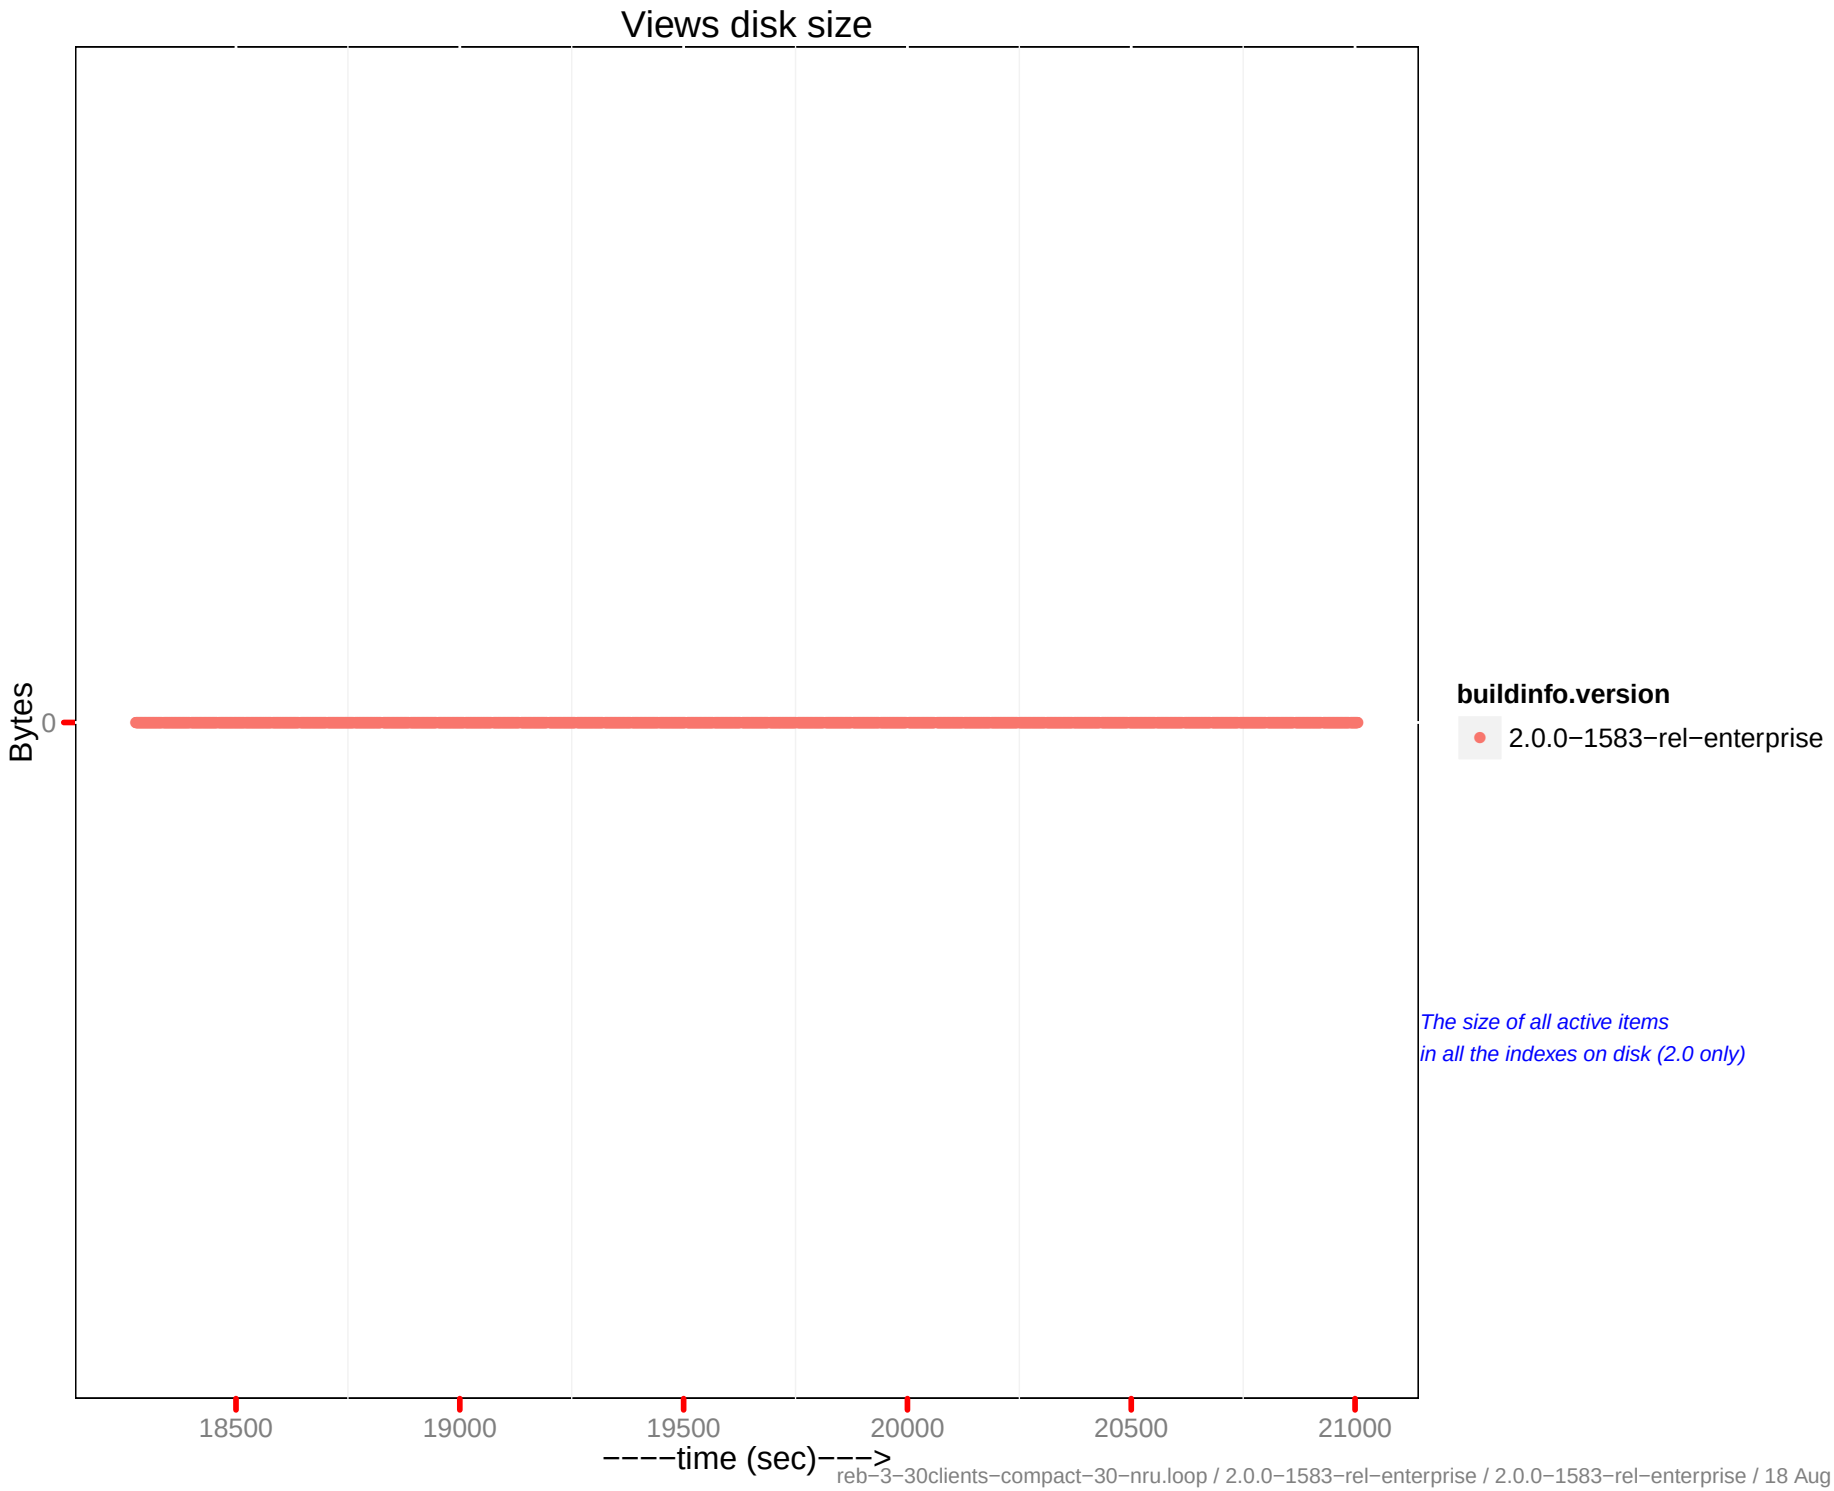
Views disk size
Bytes
0
18500
19000
19500
20000
20500
21000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
The size of all active items
in all the indexes on disk (2.0 only)
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

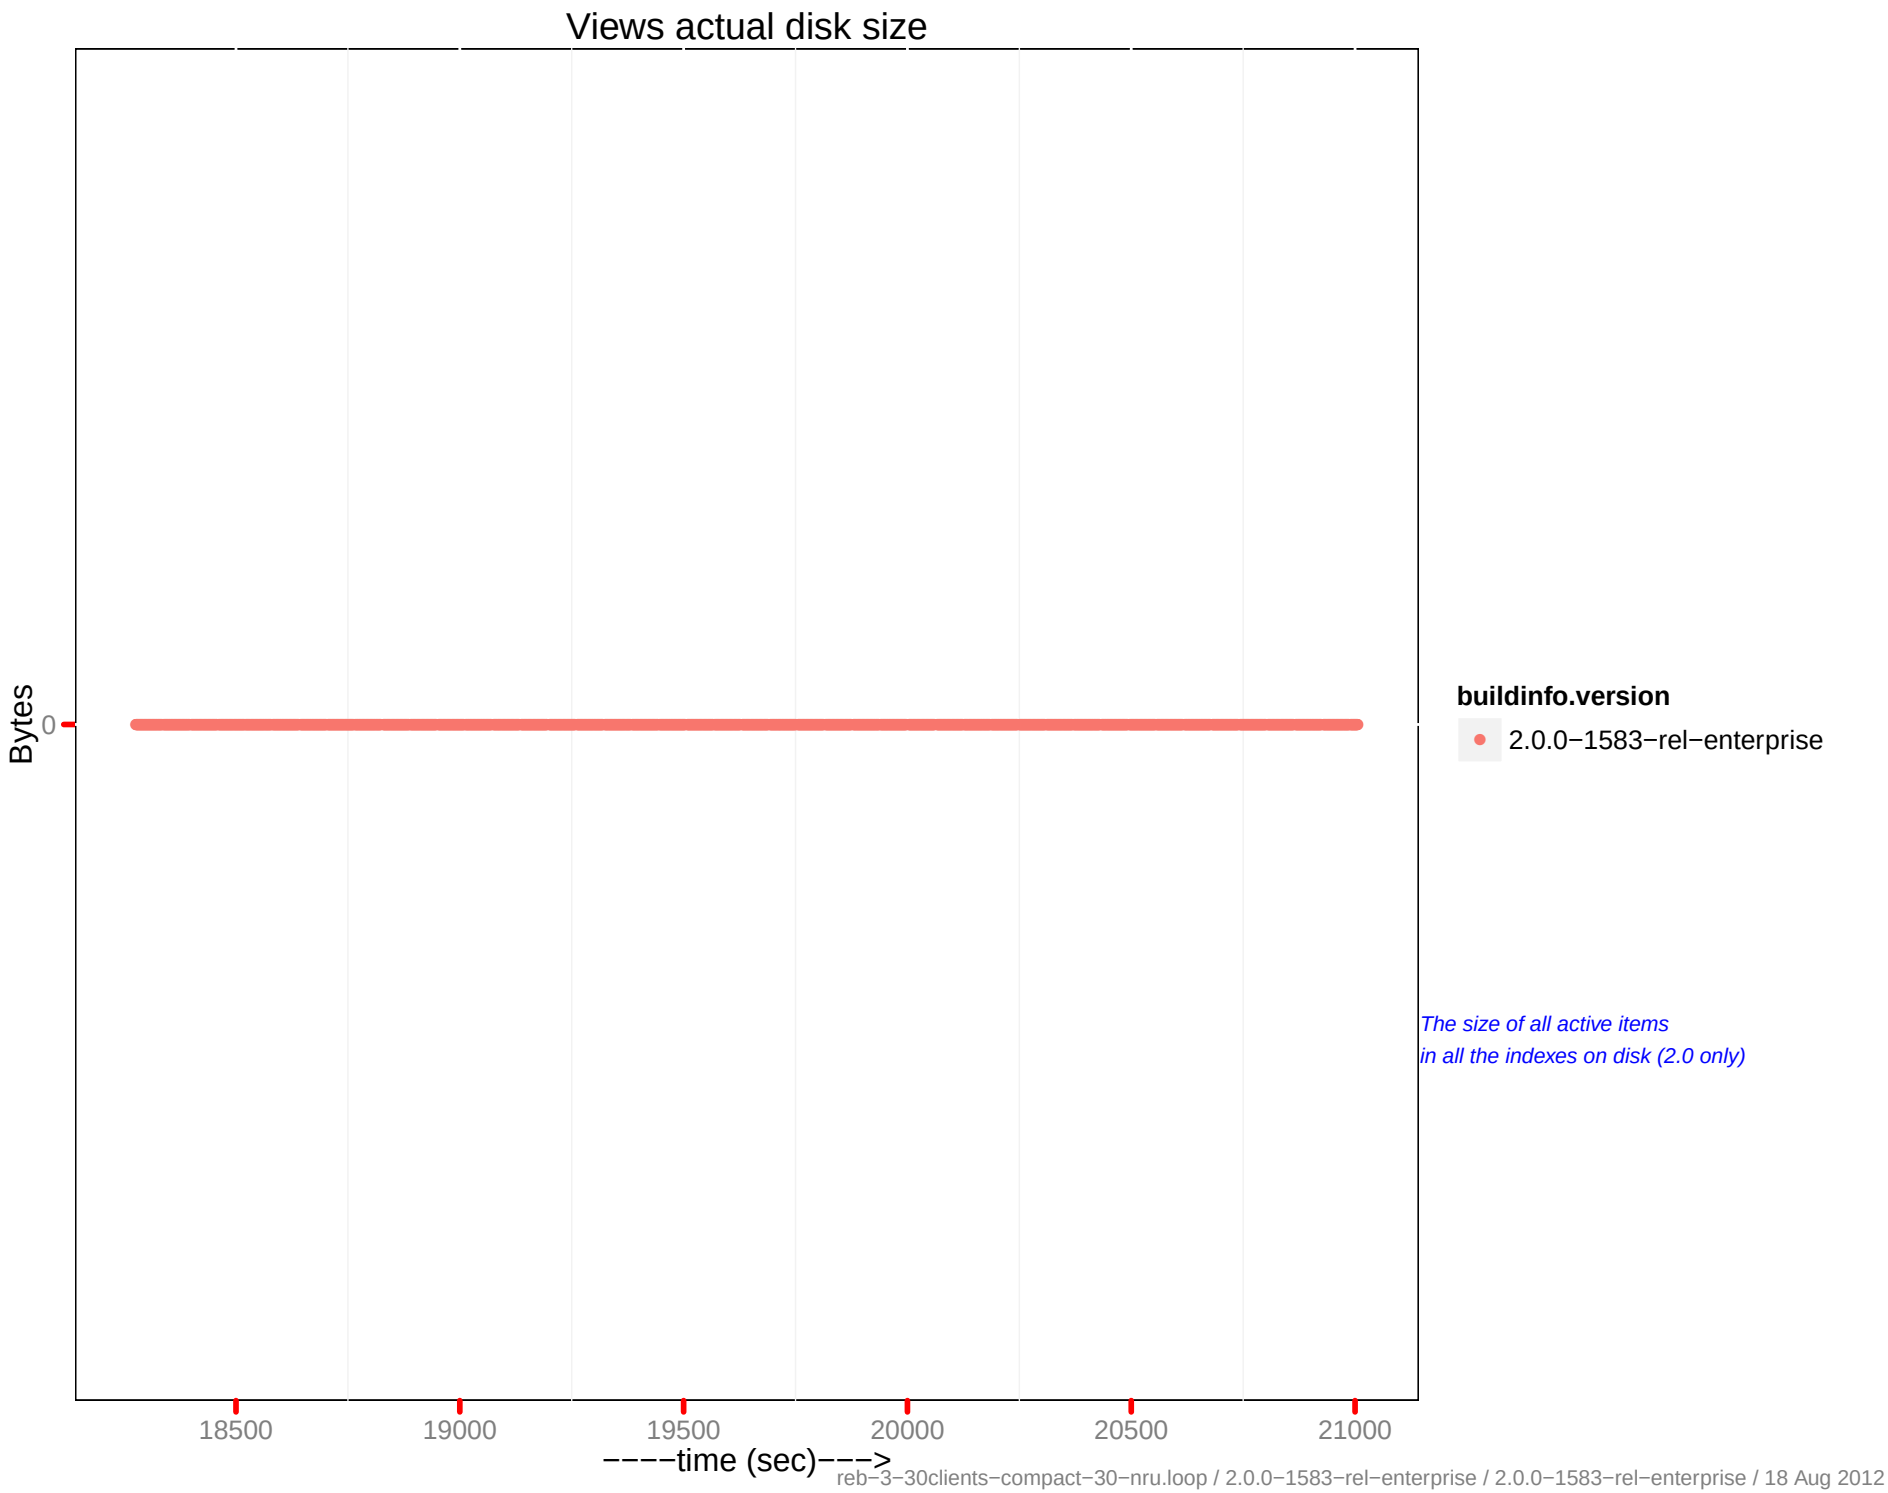
Views actual disk size
Bytes
0
18500
19000
19500
20000
20500
21000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
The size of all active items
in all the indexes on disk (2.0 only)
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

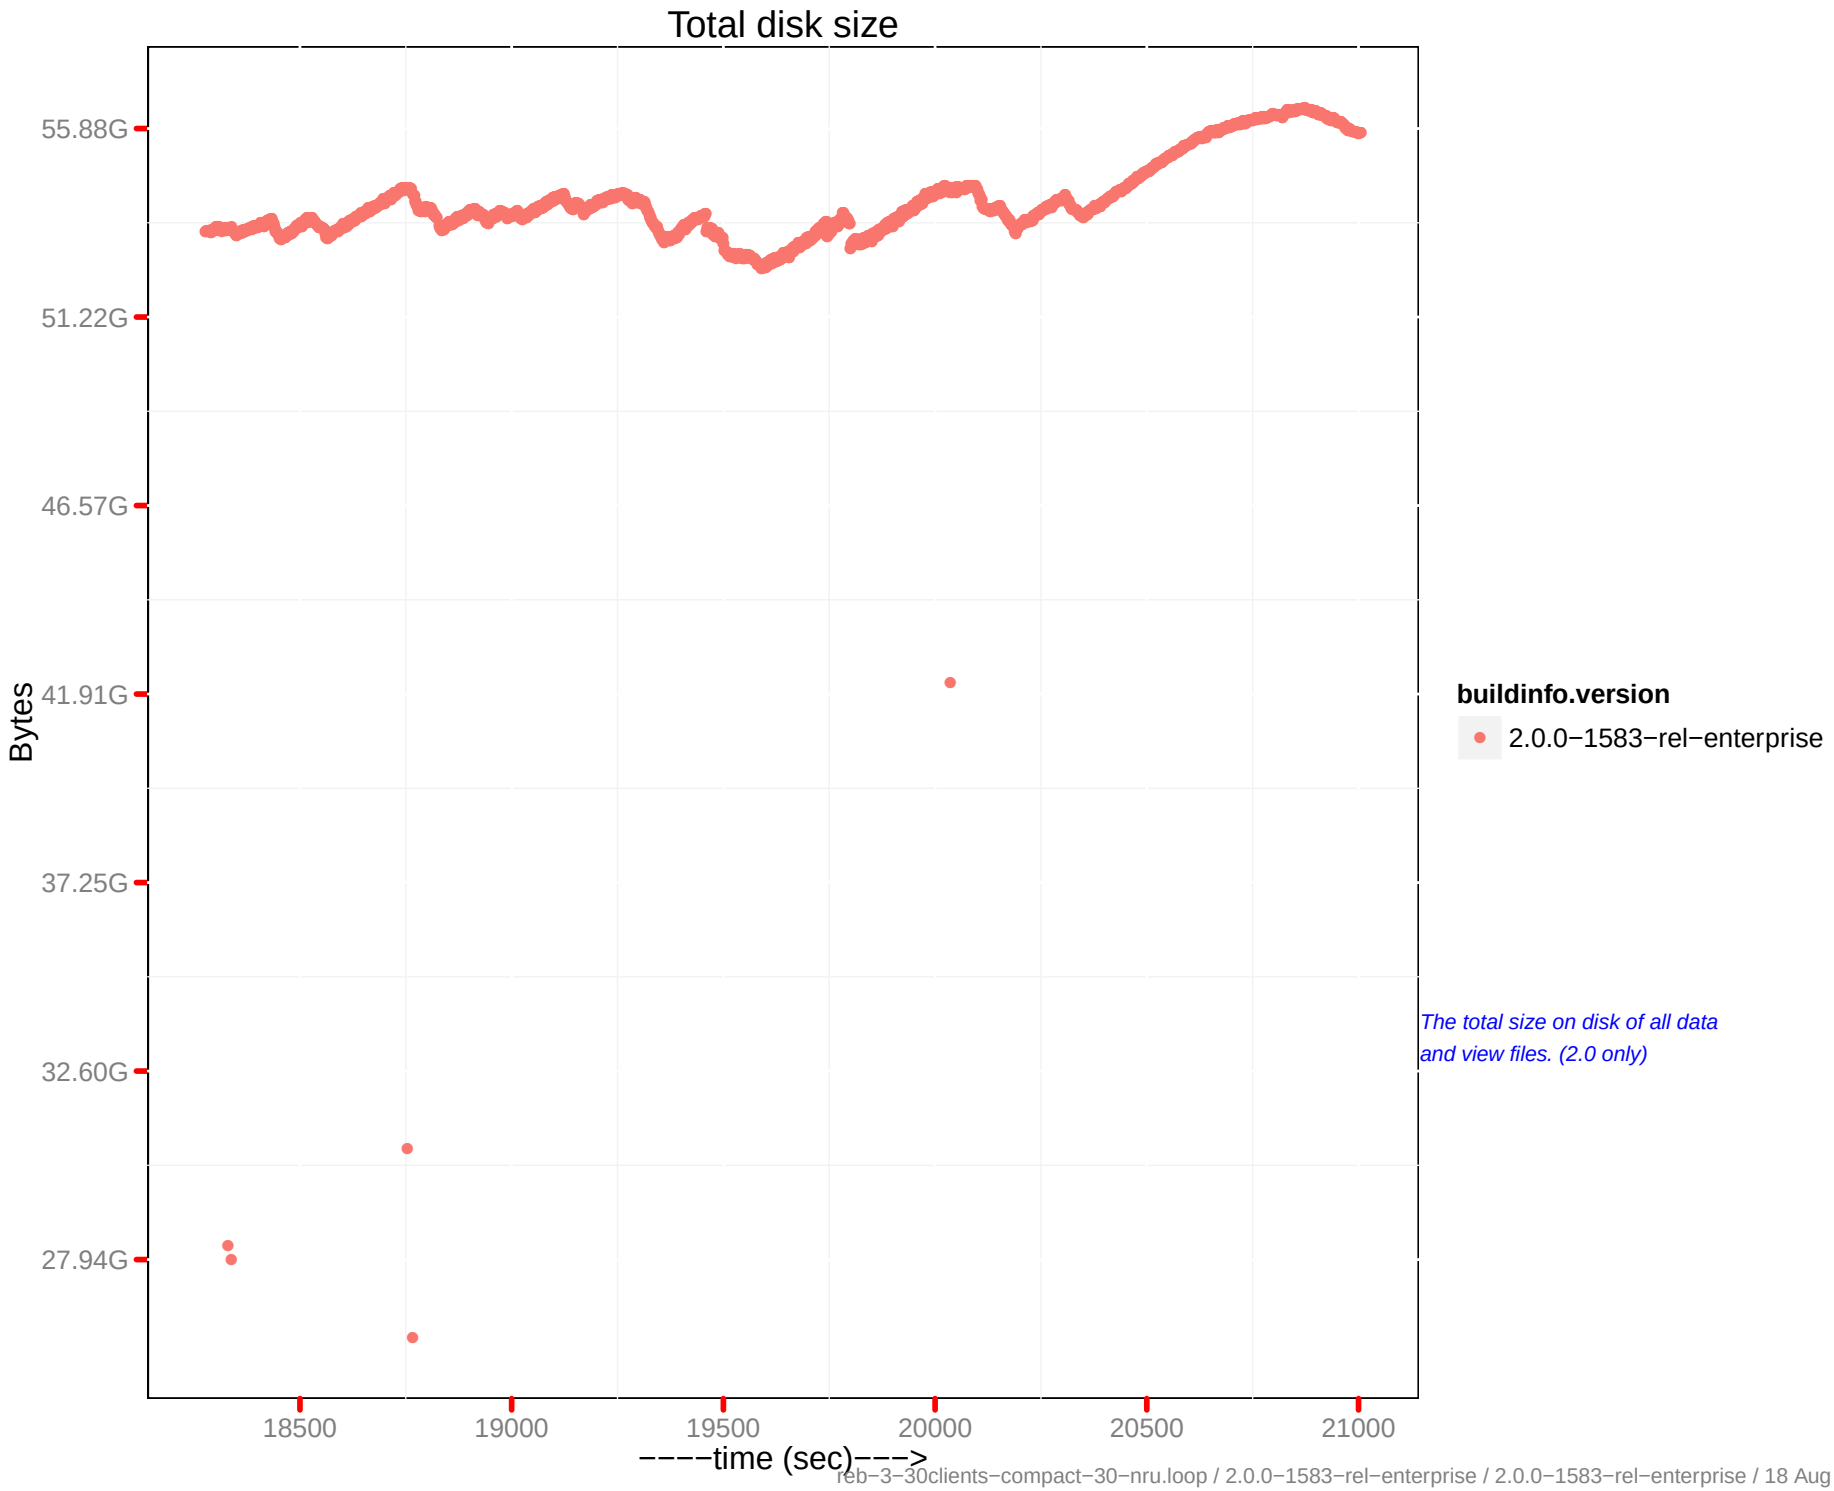
Total disk size
55.88G
51.22G
46.57G
41.91G
37.25G
32.60G
27.94G
Bytes
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

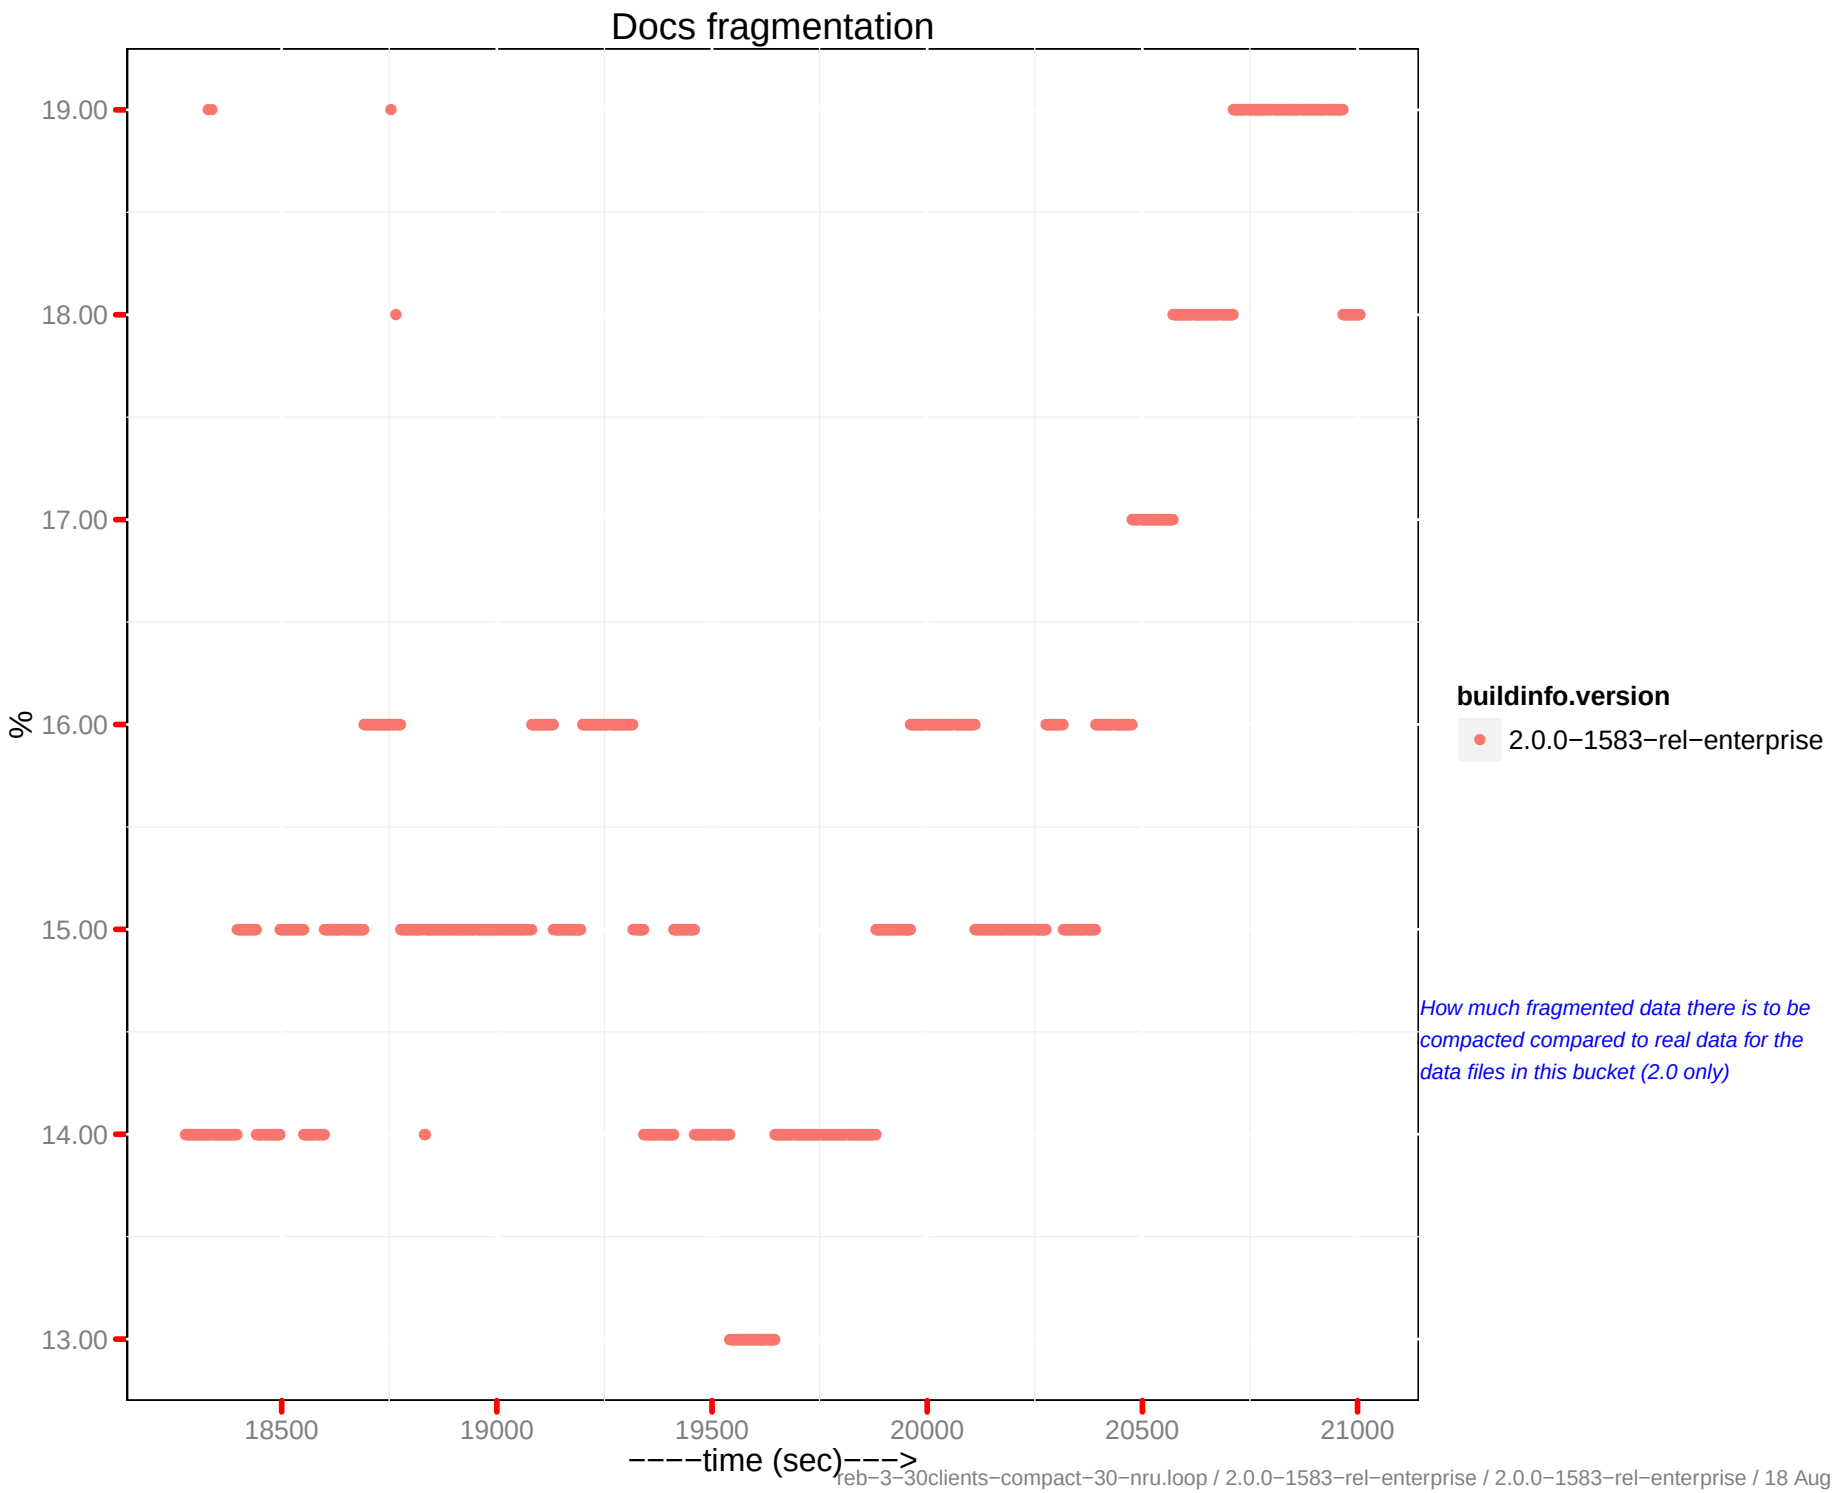
Docs fragmentation
19.00
18.00
17.00
16.00
15.00
14.00
13.00
%
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

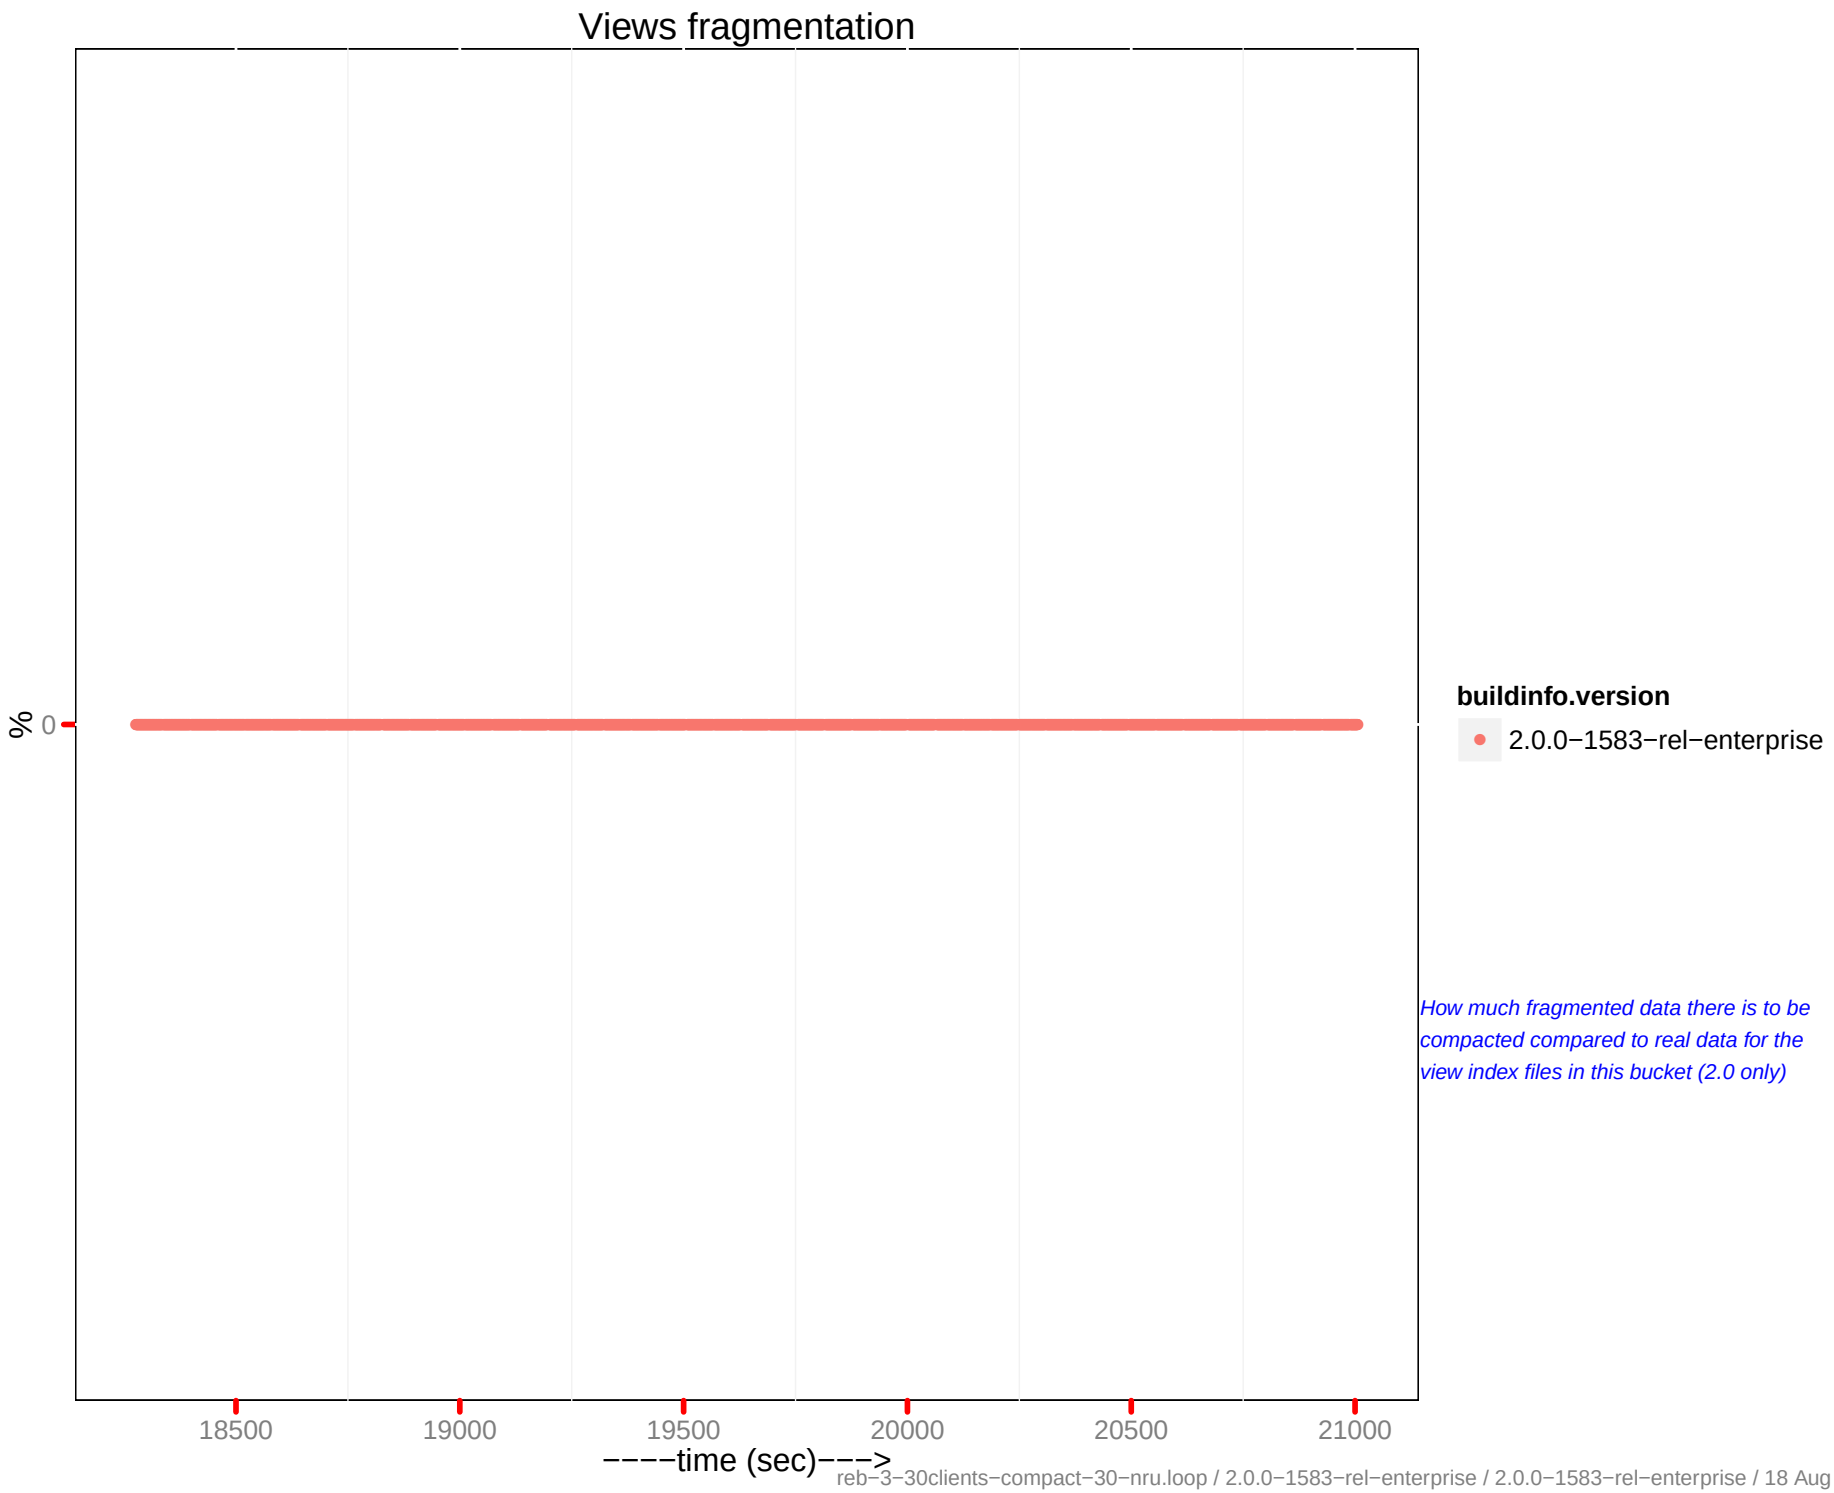
Views fragmentation
% 0
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

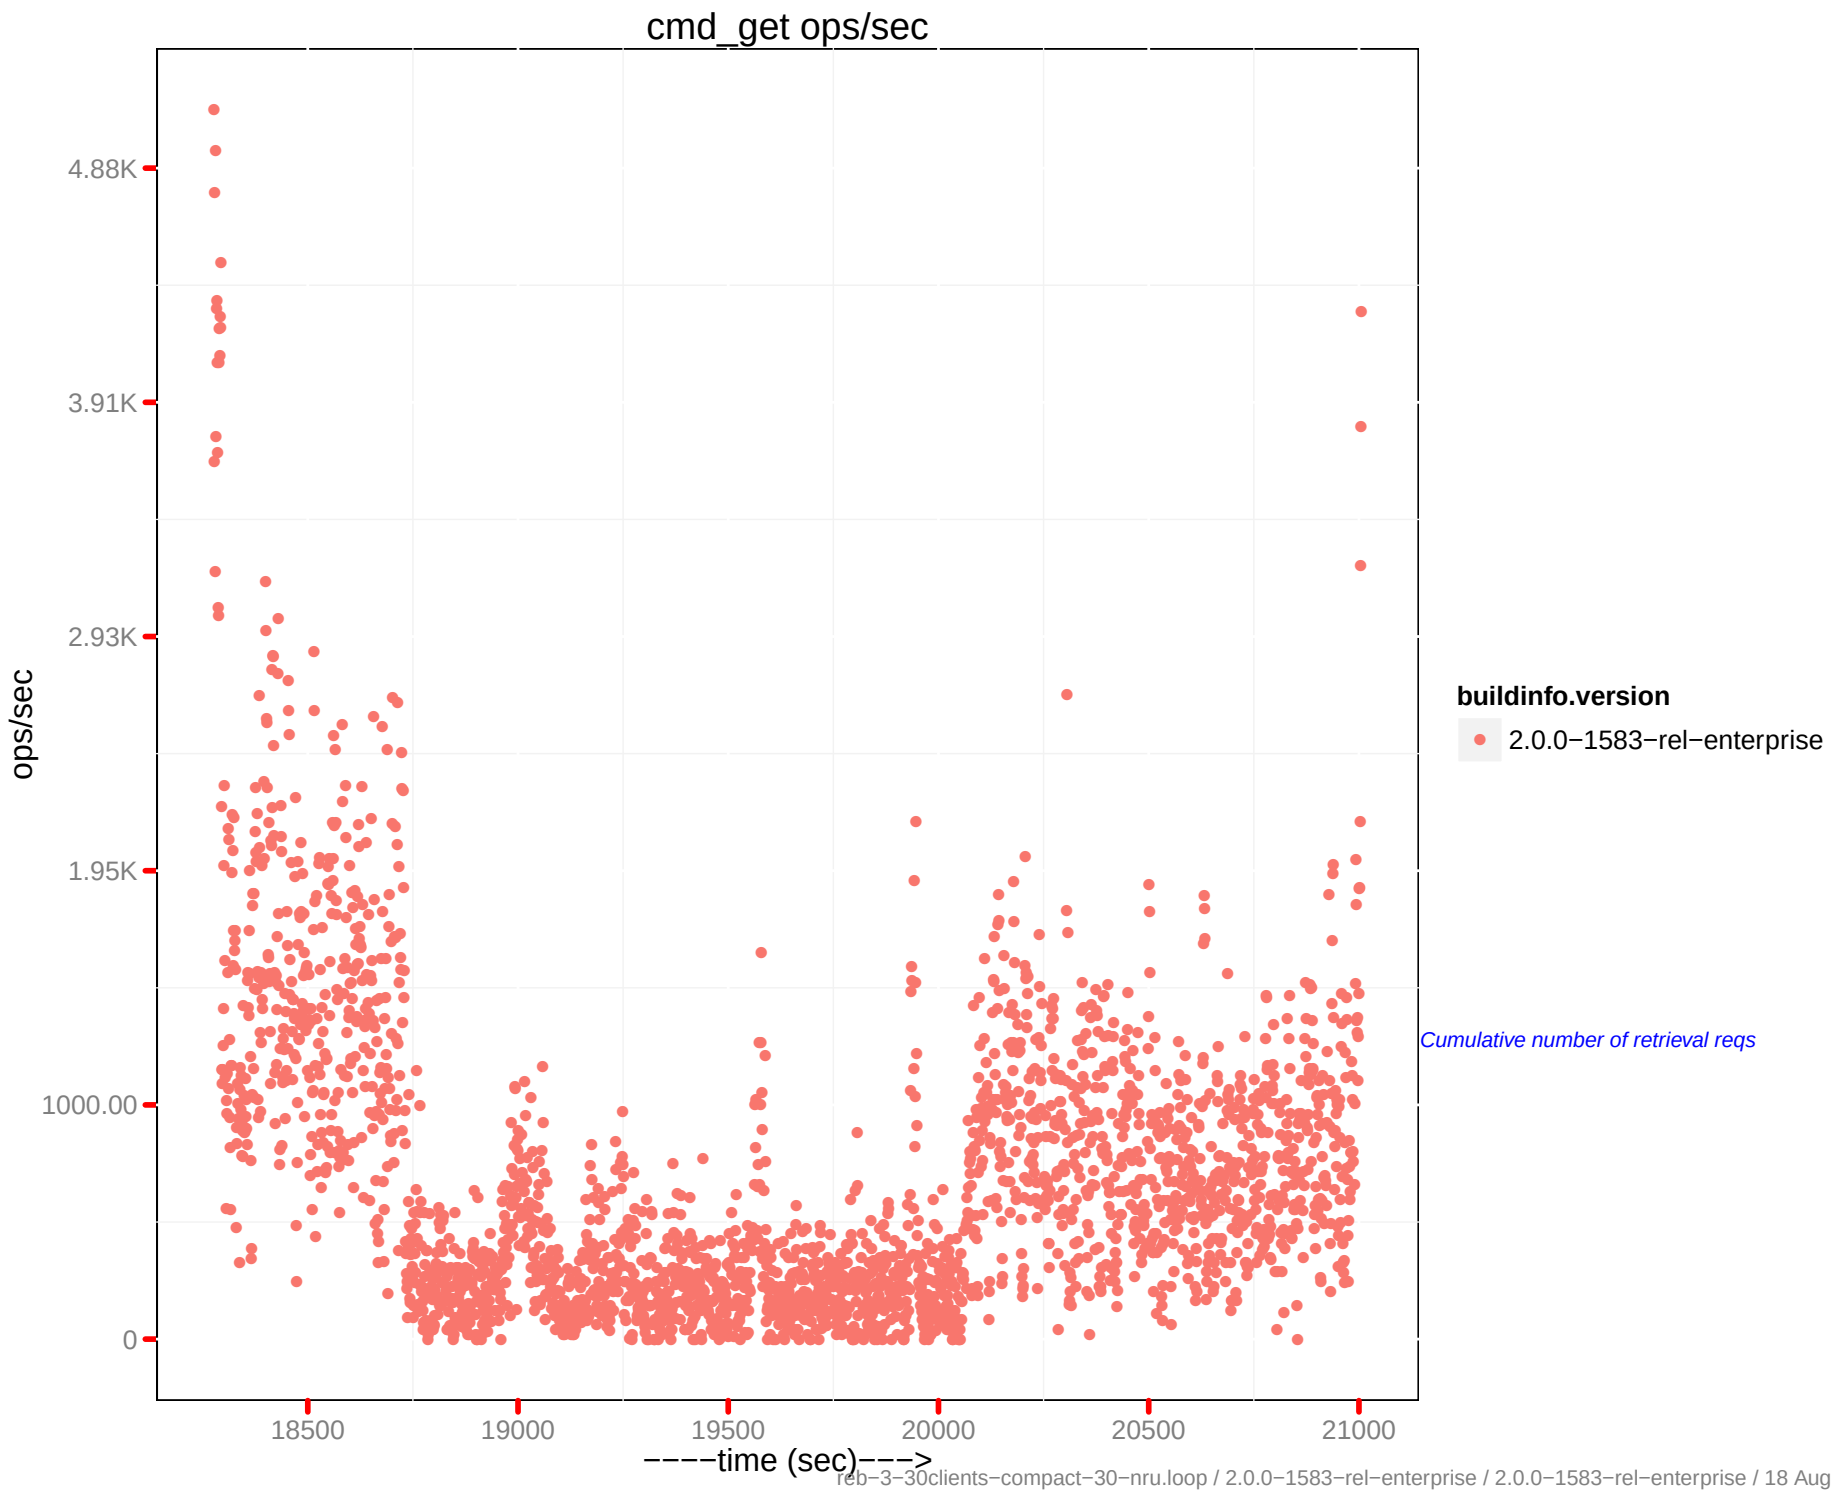
cmd_get ops/sec
4.88K
3.91K
2.93K
1.95K
1000.00
0
ops/sec
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Cumulative number of retrieval reqs
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

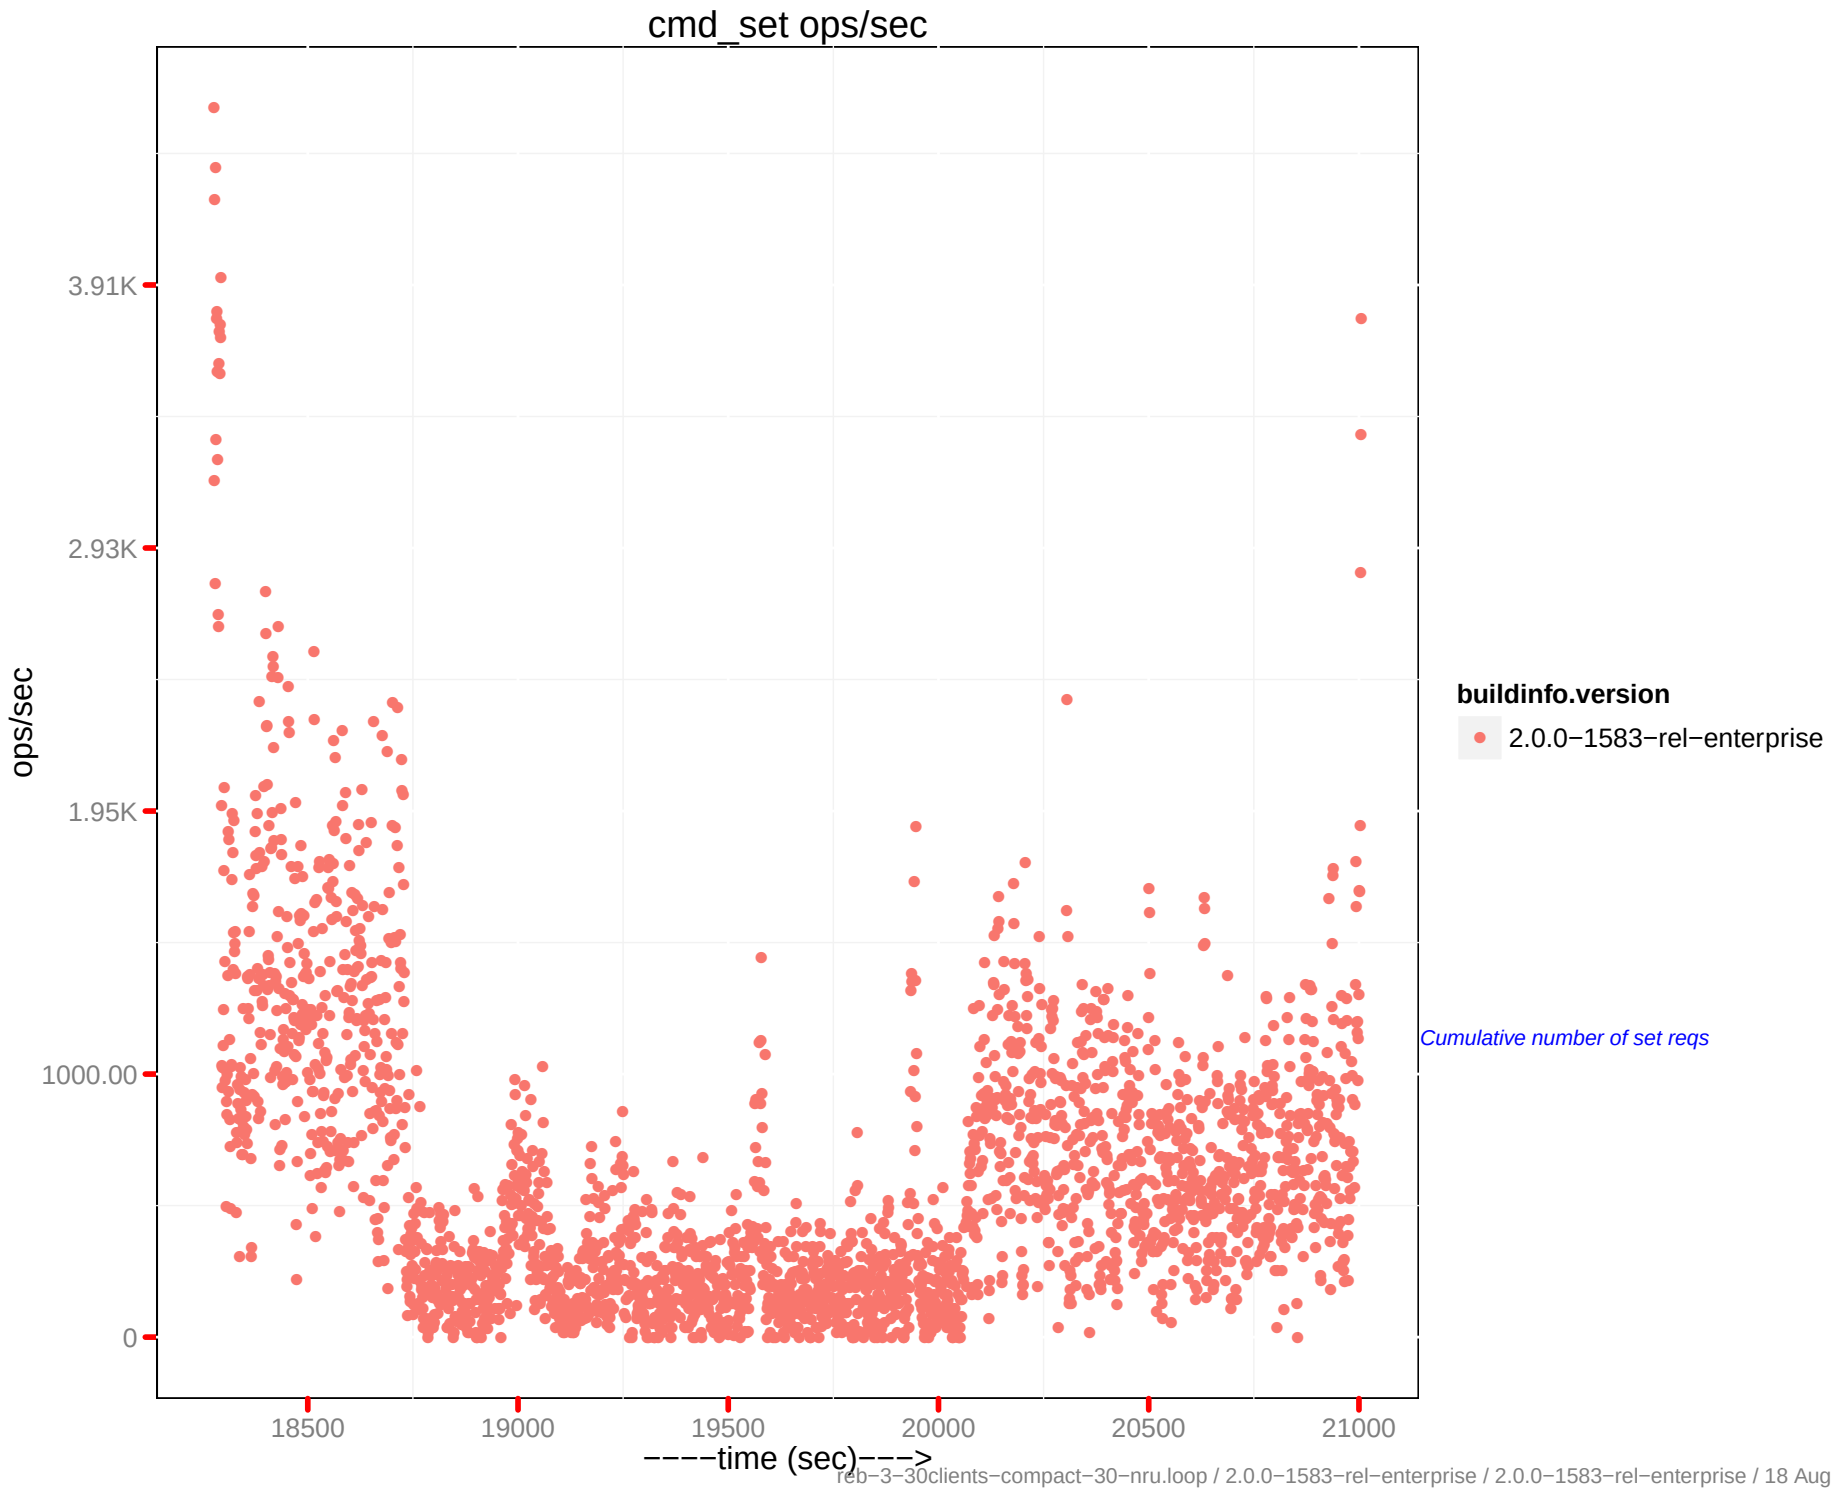
cmd_set ops/sec
ops/sec
3.91K
2.93K
1.95K
1000.00
0
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Cumulative number of set reqs
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

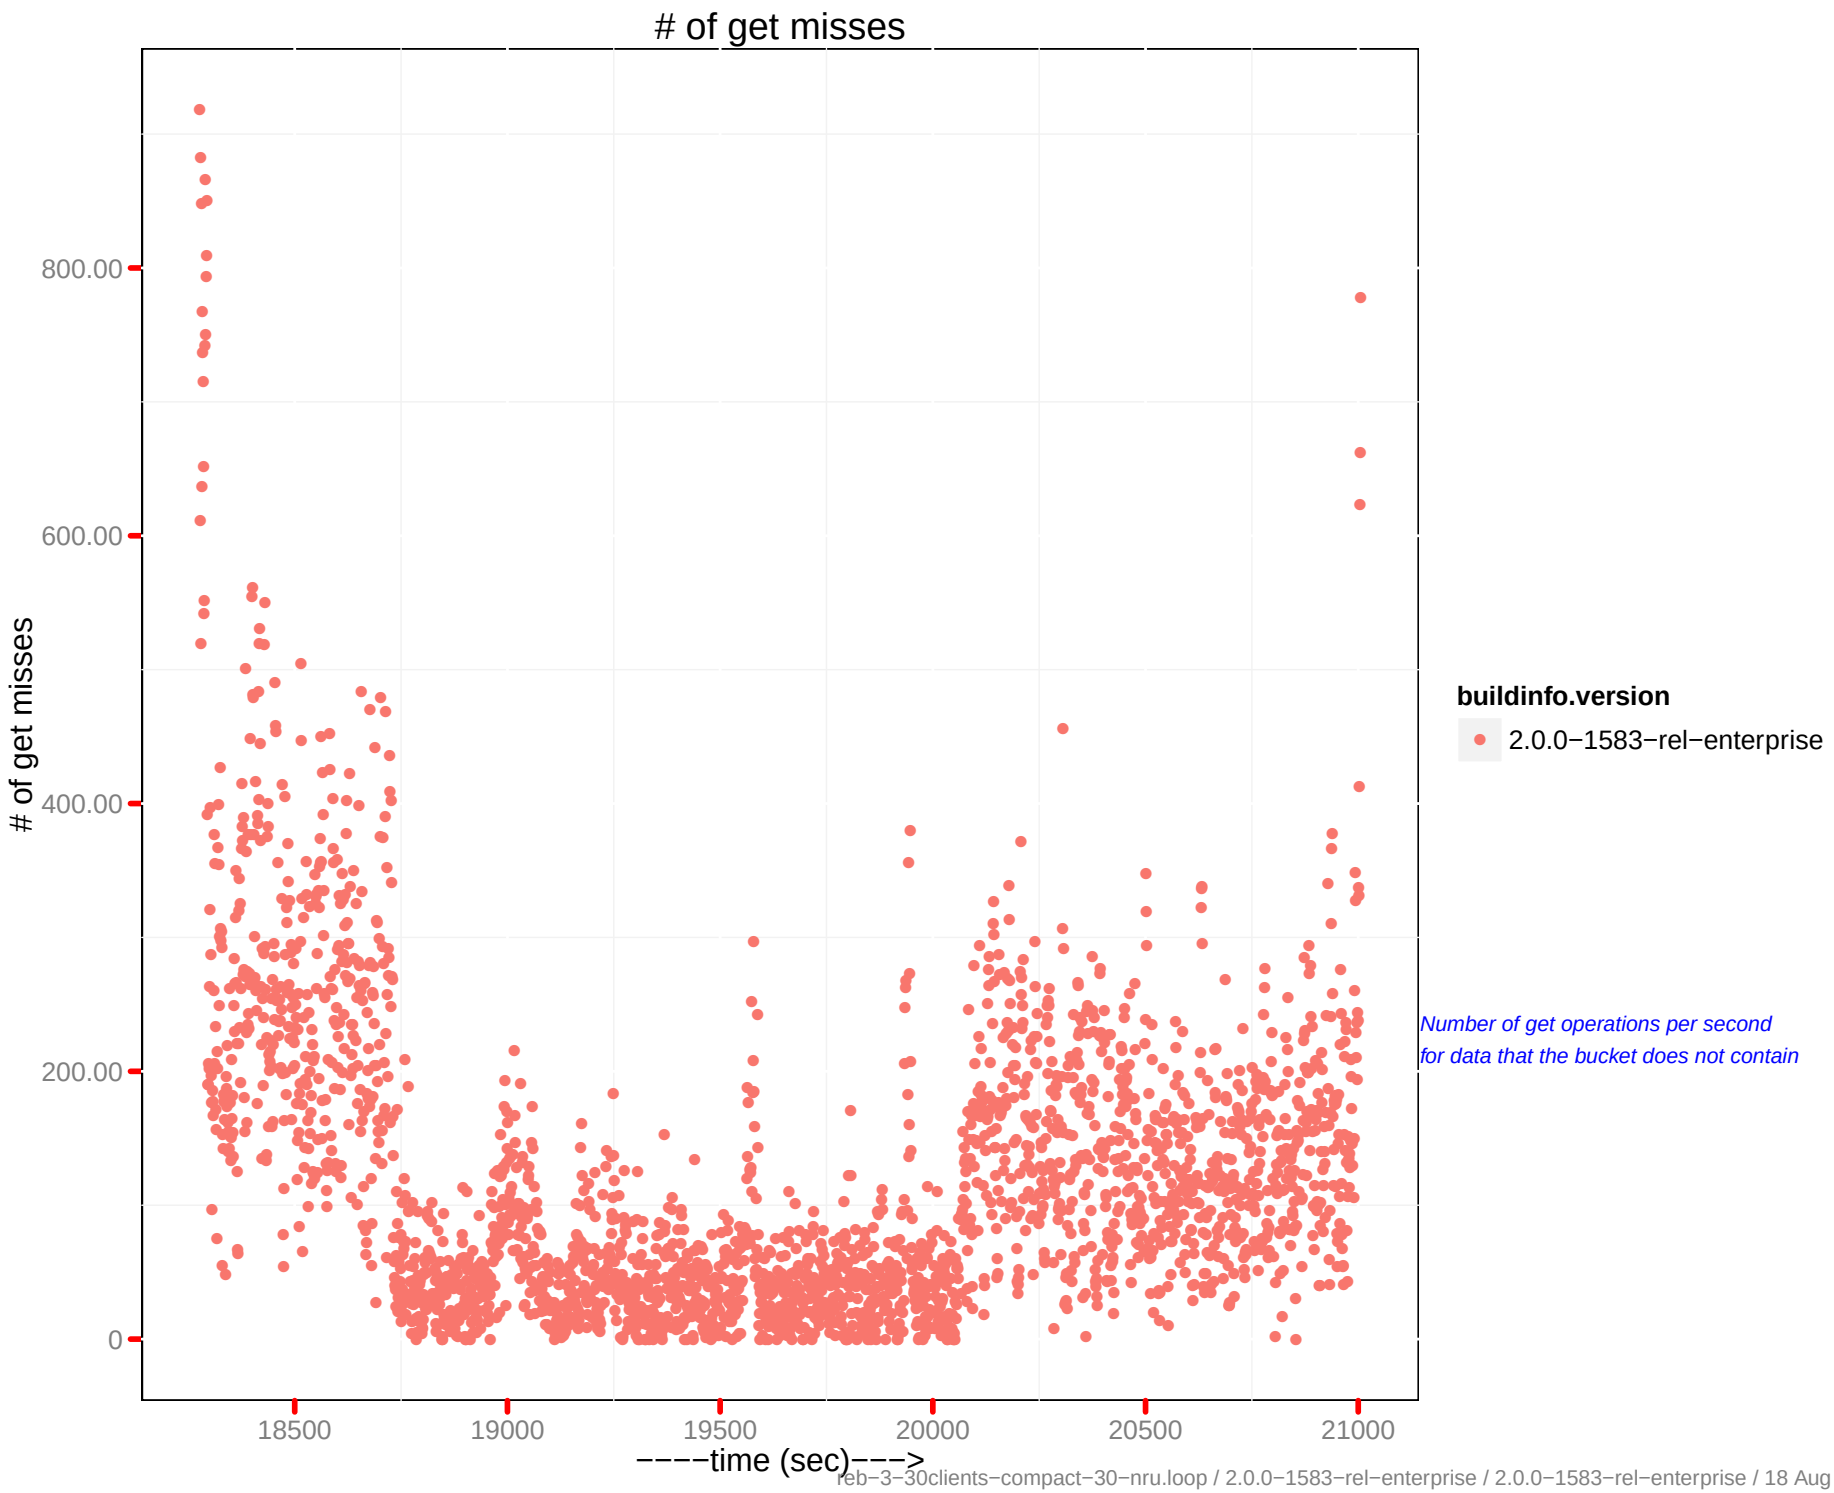
# of get misses
800.00
600.00
400.00
200.00
0
# of get misses
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of get operations per second
for data that the bucket does not contain
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

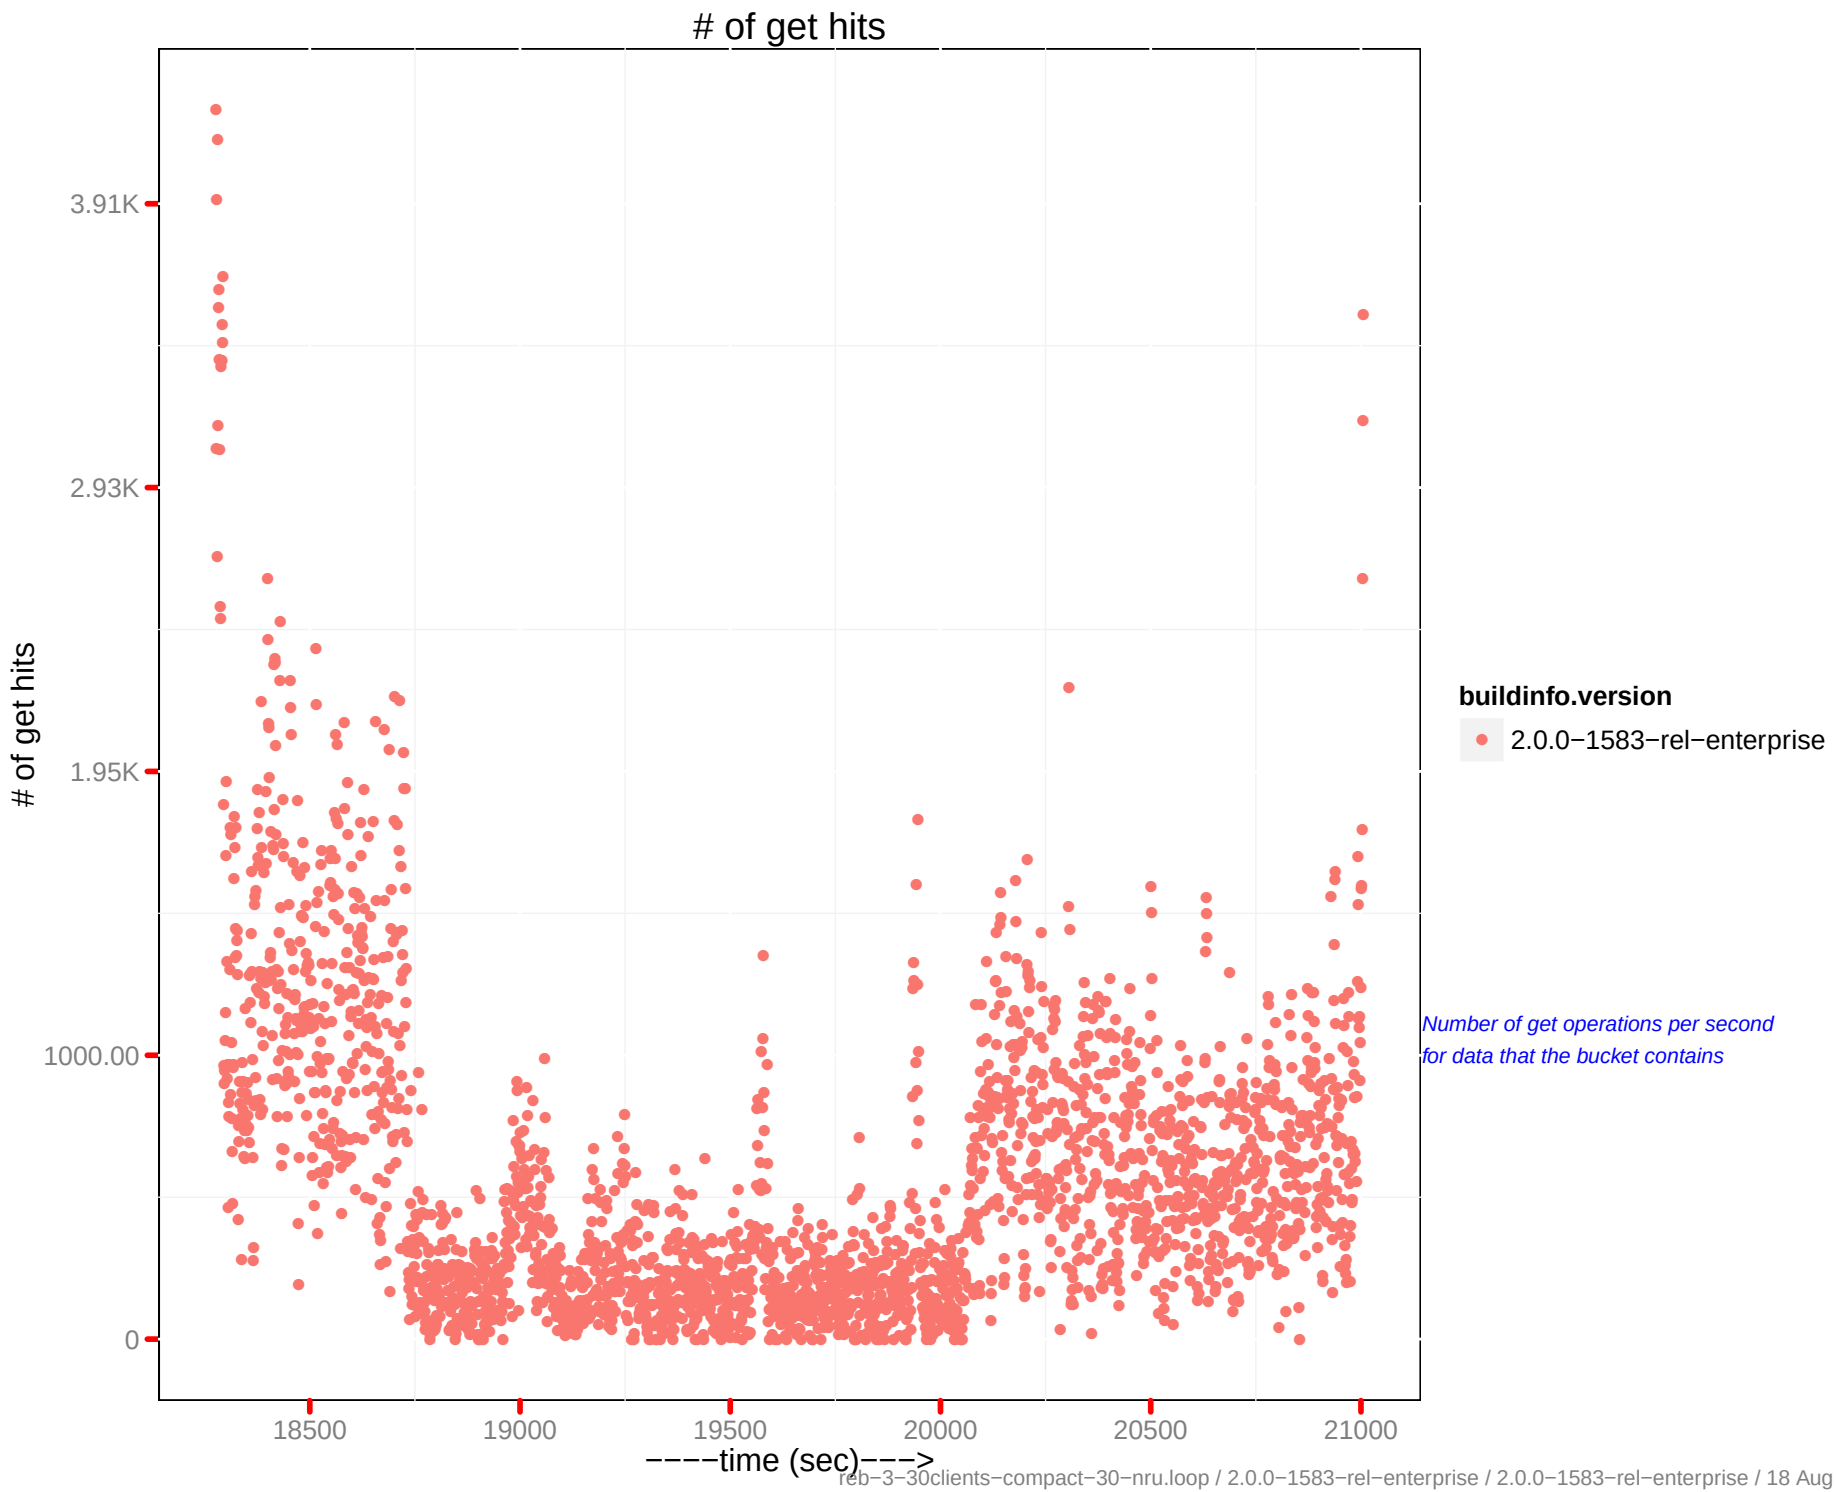
# of get hits
3.91K
2.93K
1.95K
1000.00
0
# of get hits
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of get operations per second
for data that the bucket contains
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

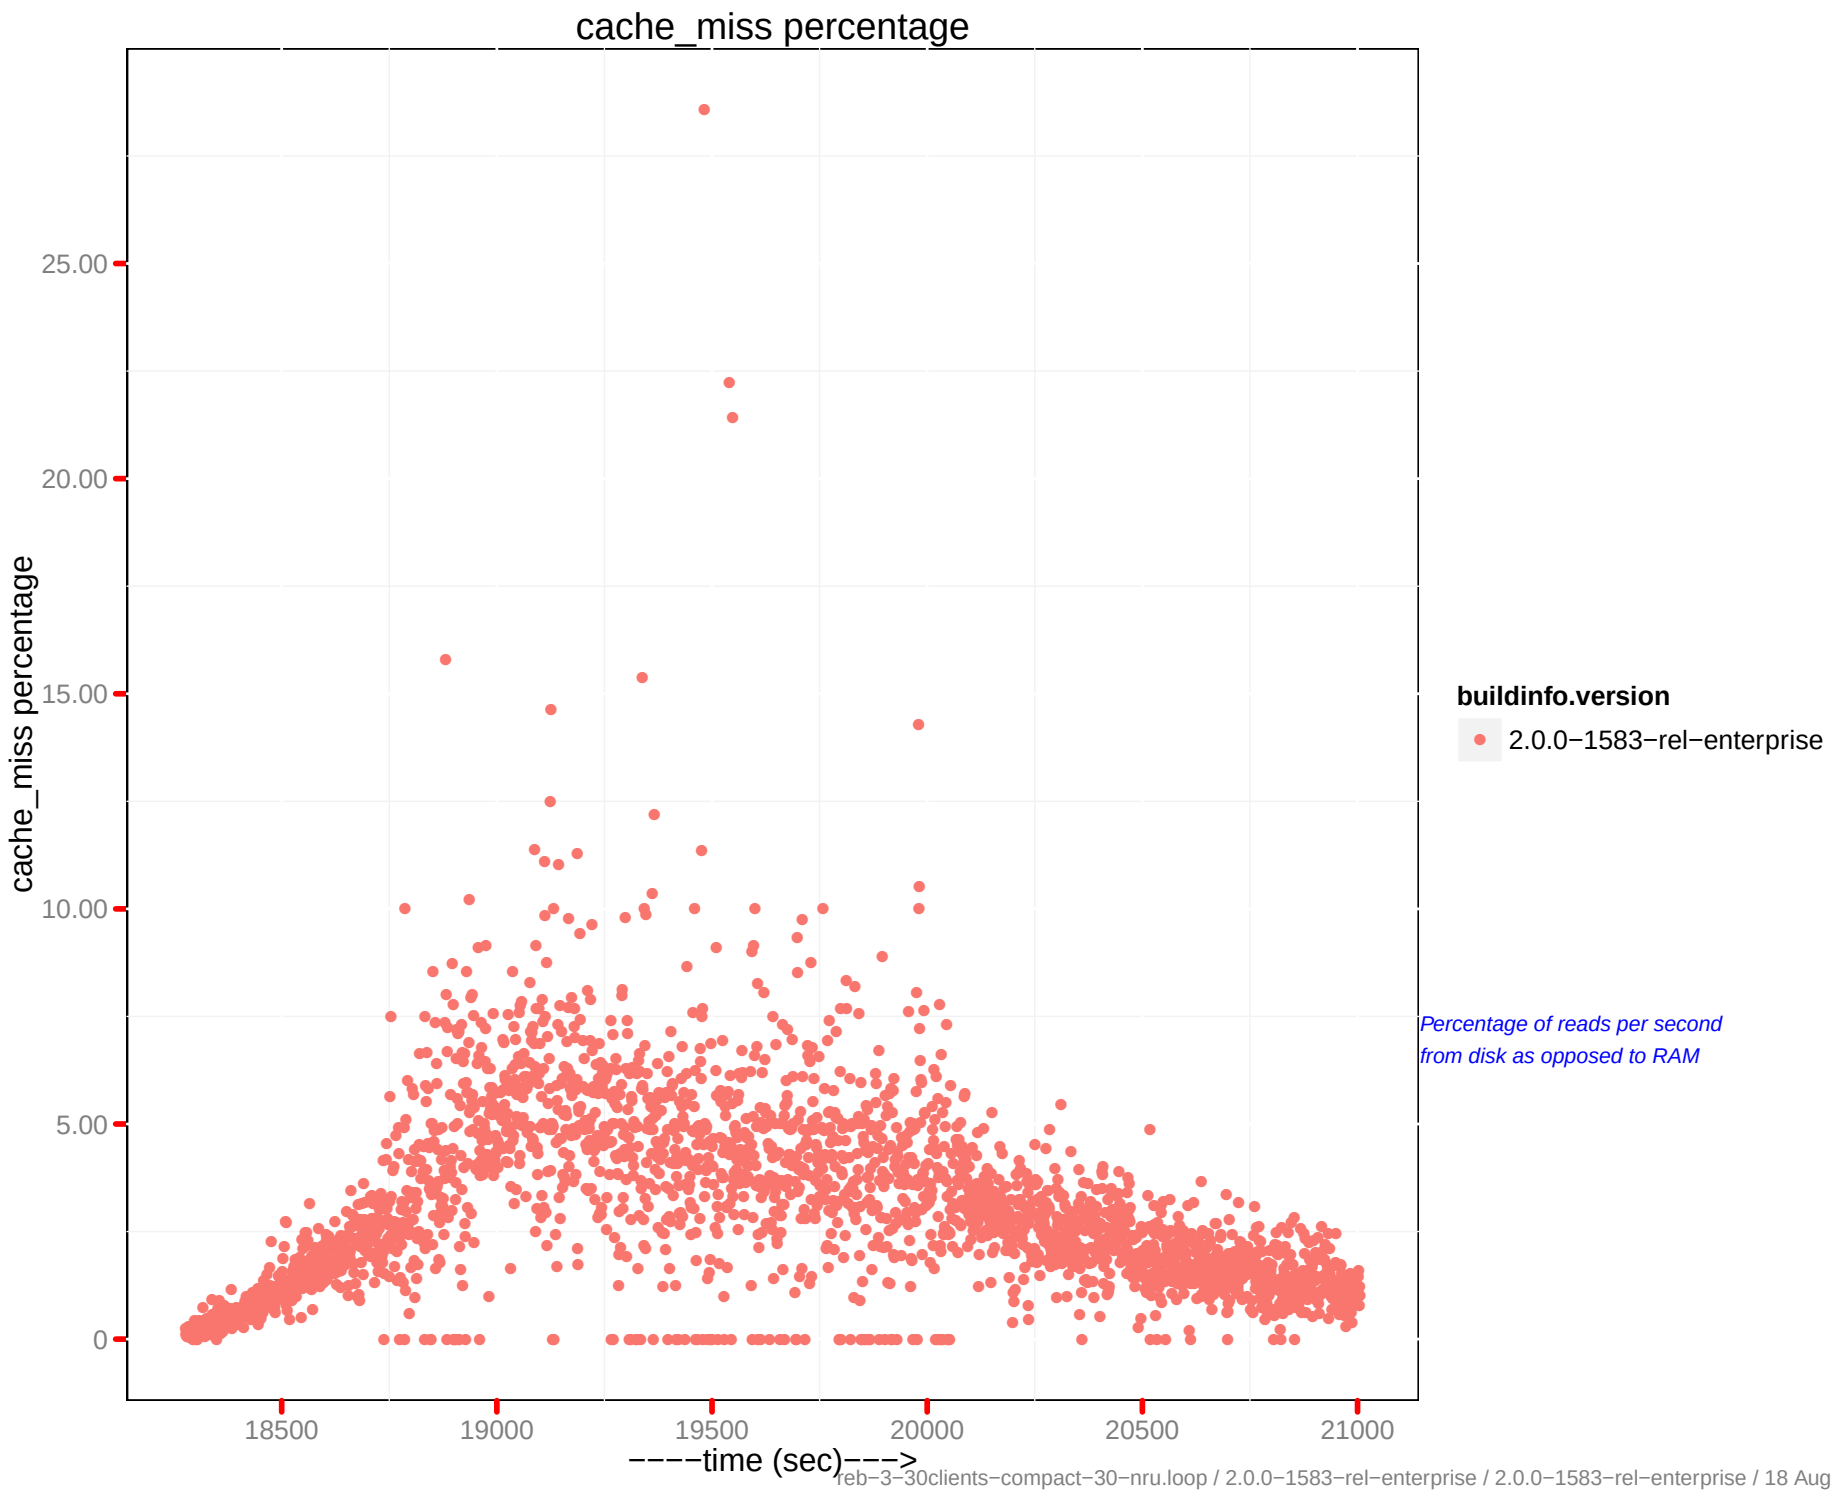
cache_miss percentage
cache_miss percentage
25.00
20.00
15.00
10.00
5.00
0
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Percentage of reads per second
from disk as opposed to RAM
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

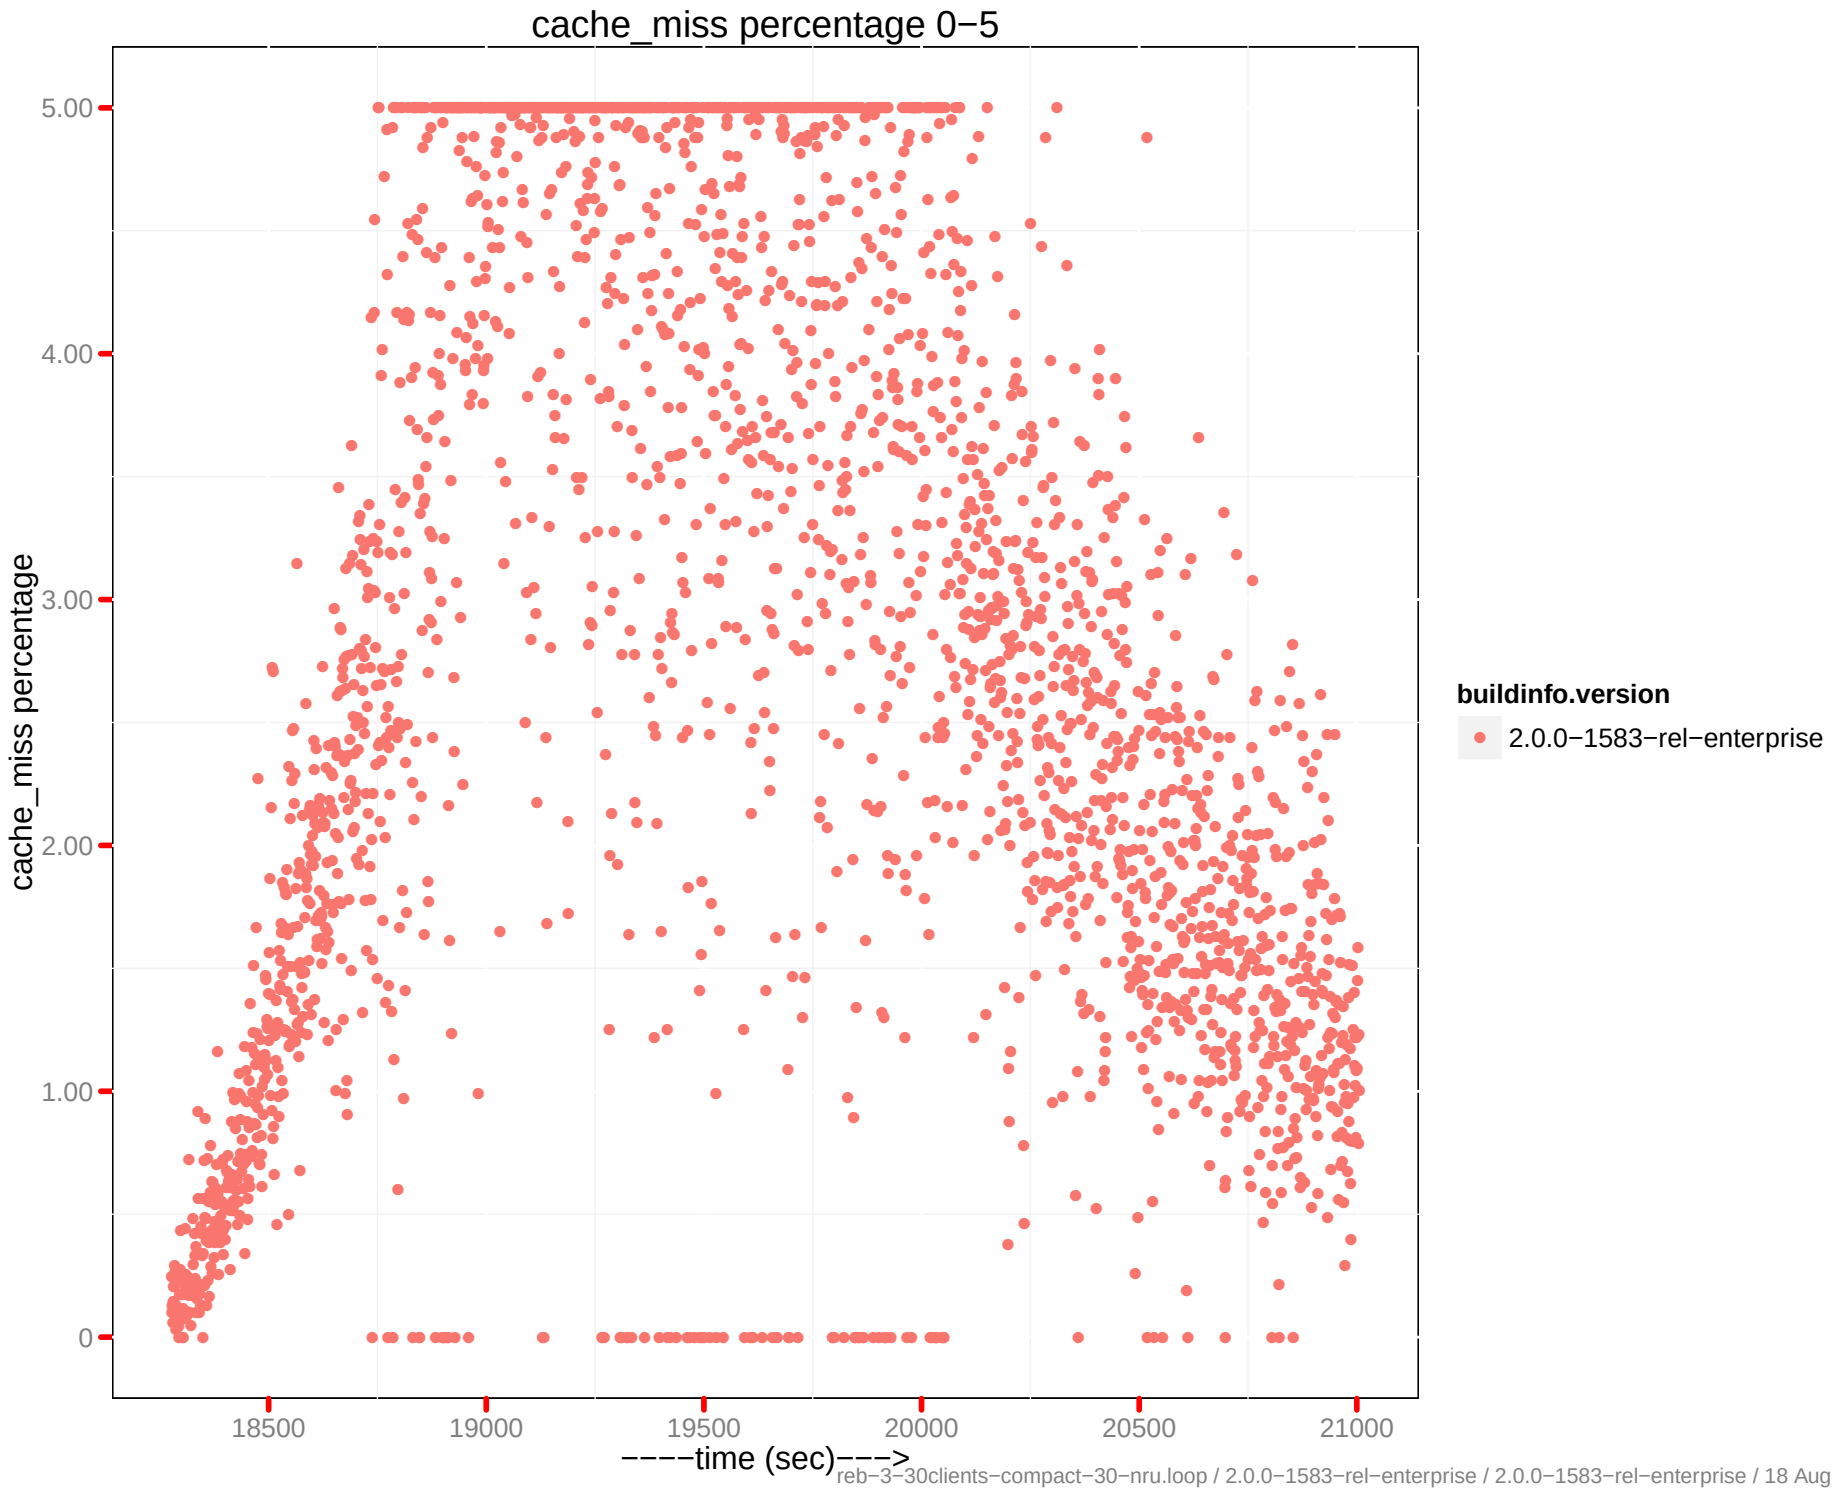
cache_miss percentage 0−5
cache_miss percentage
5.00
4.00
3.00
2.00
1.00
0
18500
19000
19500
20000
20500
21000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

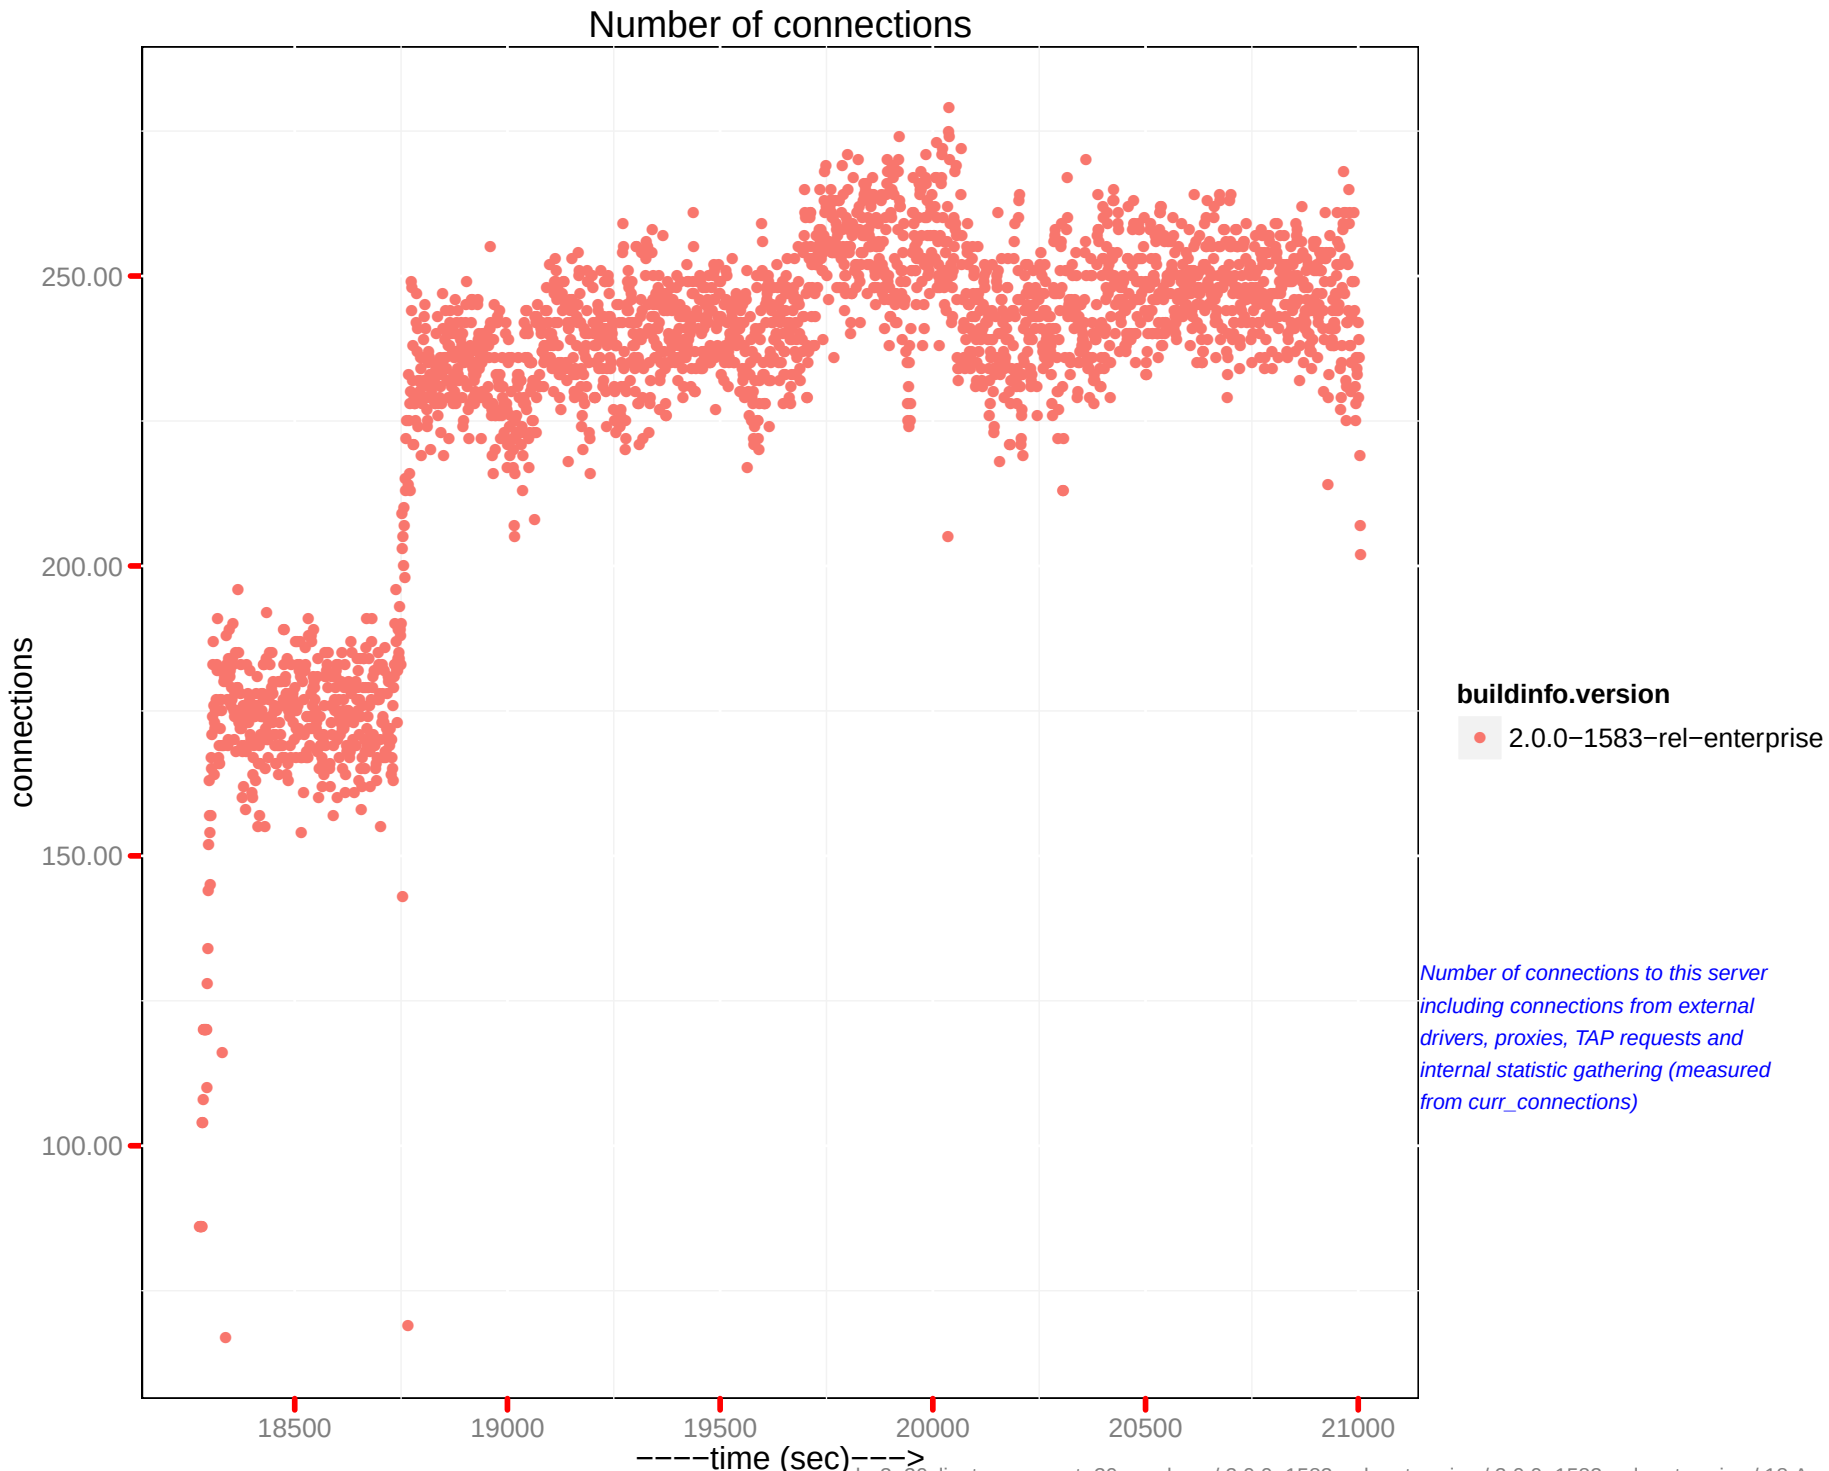
Number of connections
connections
250.00
200.00
150.00
100.00
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Number of connections to this server including connections from external drivers, proxies, TAP requests and internal statistic gathering (measured from curr_connections)
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

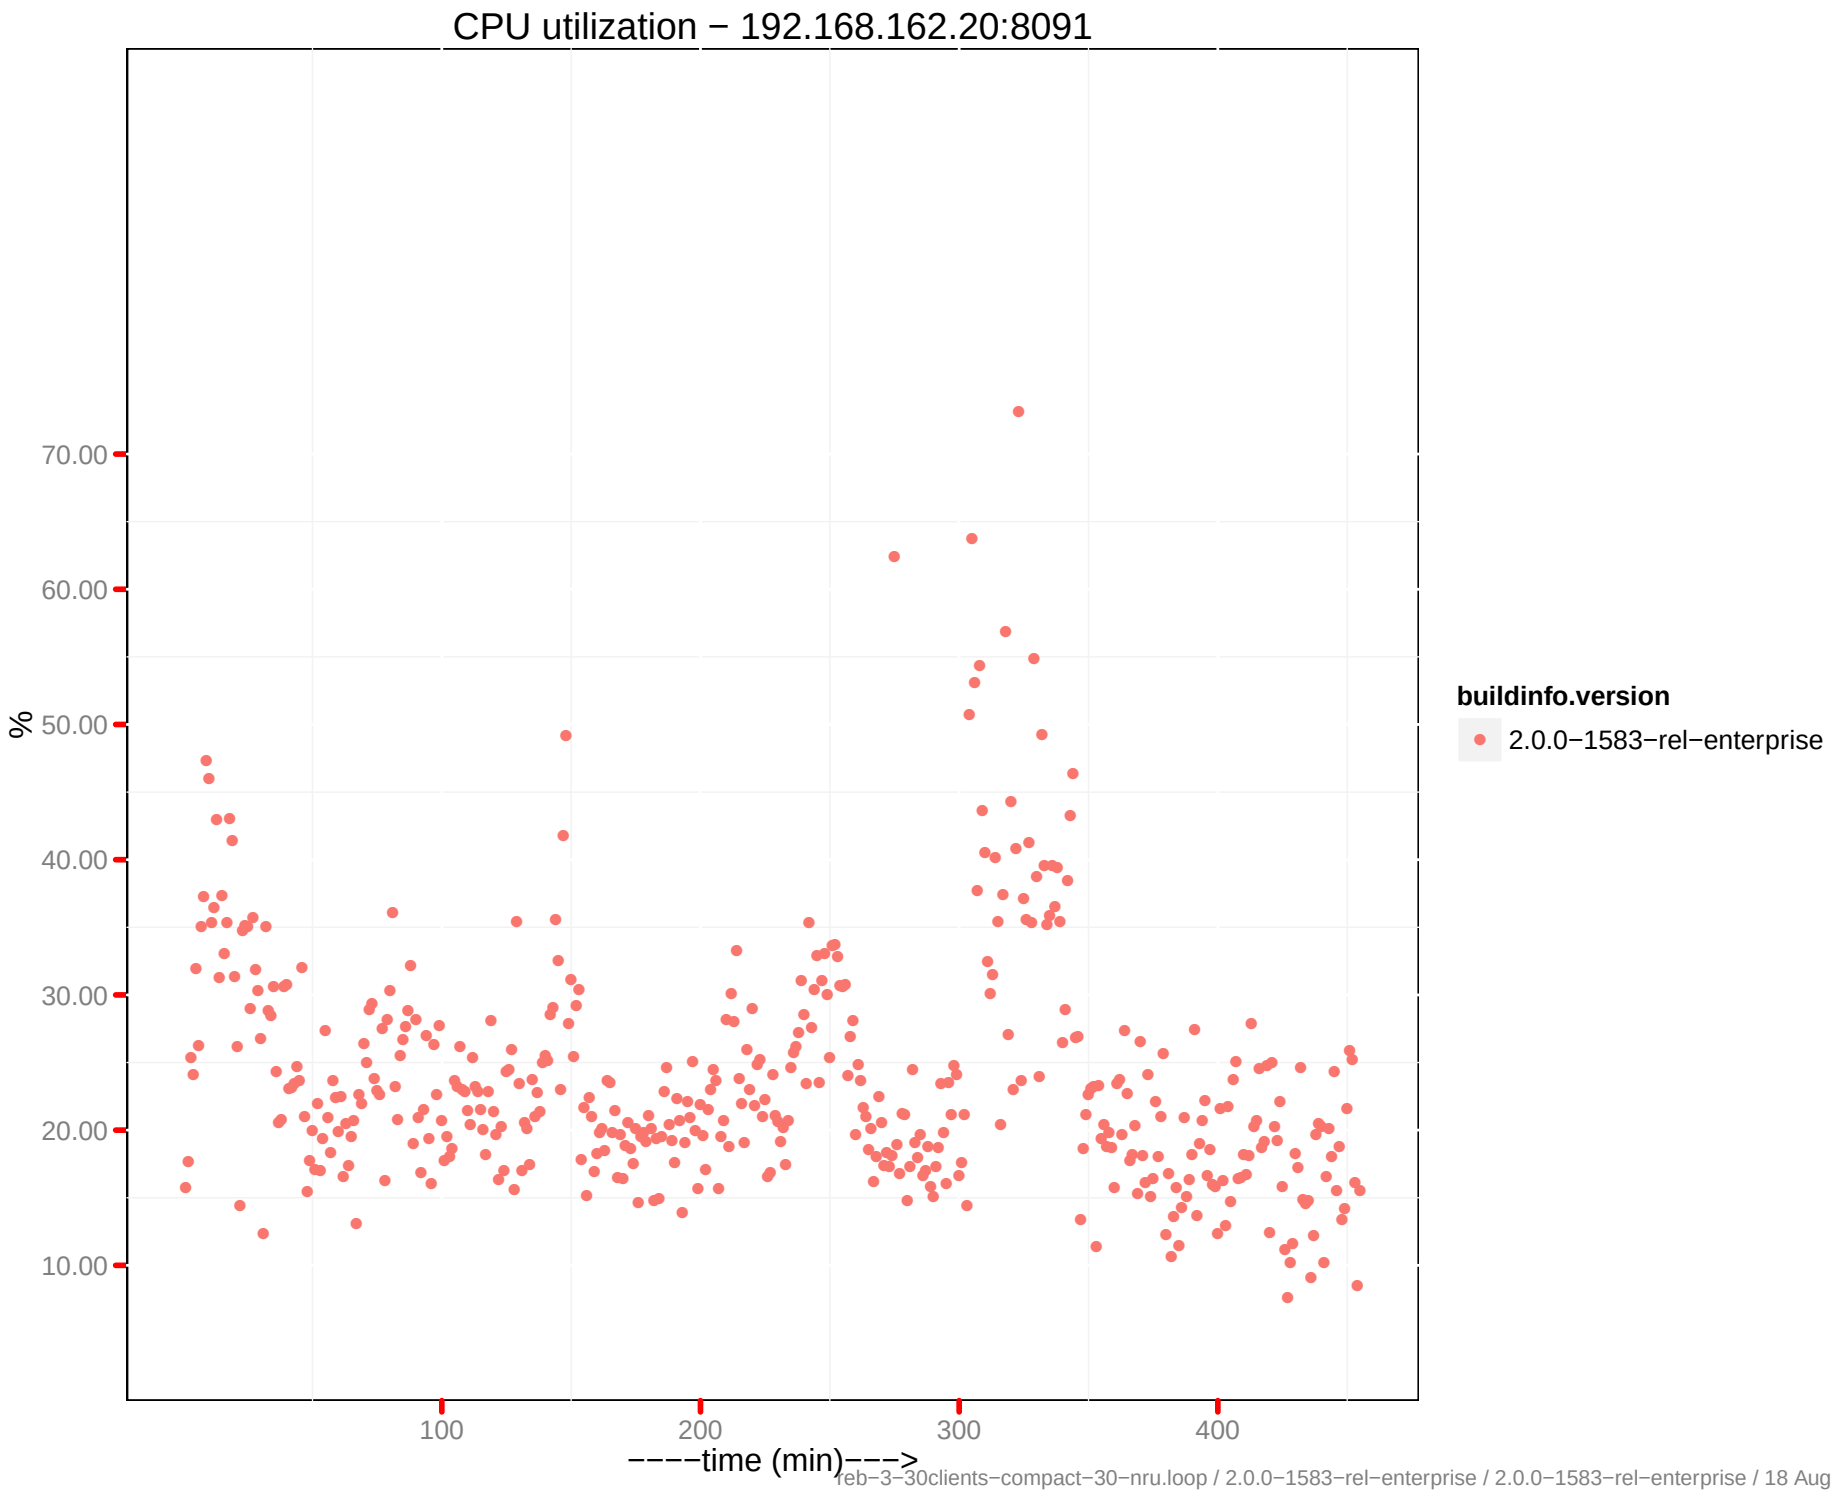
CPU utilization − 192.168.162.20:8091
70.00
60.00
50.00
40.00
30.00
20.00
10.00
%
100
200
300
400
−−−−time (min)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

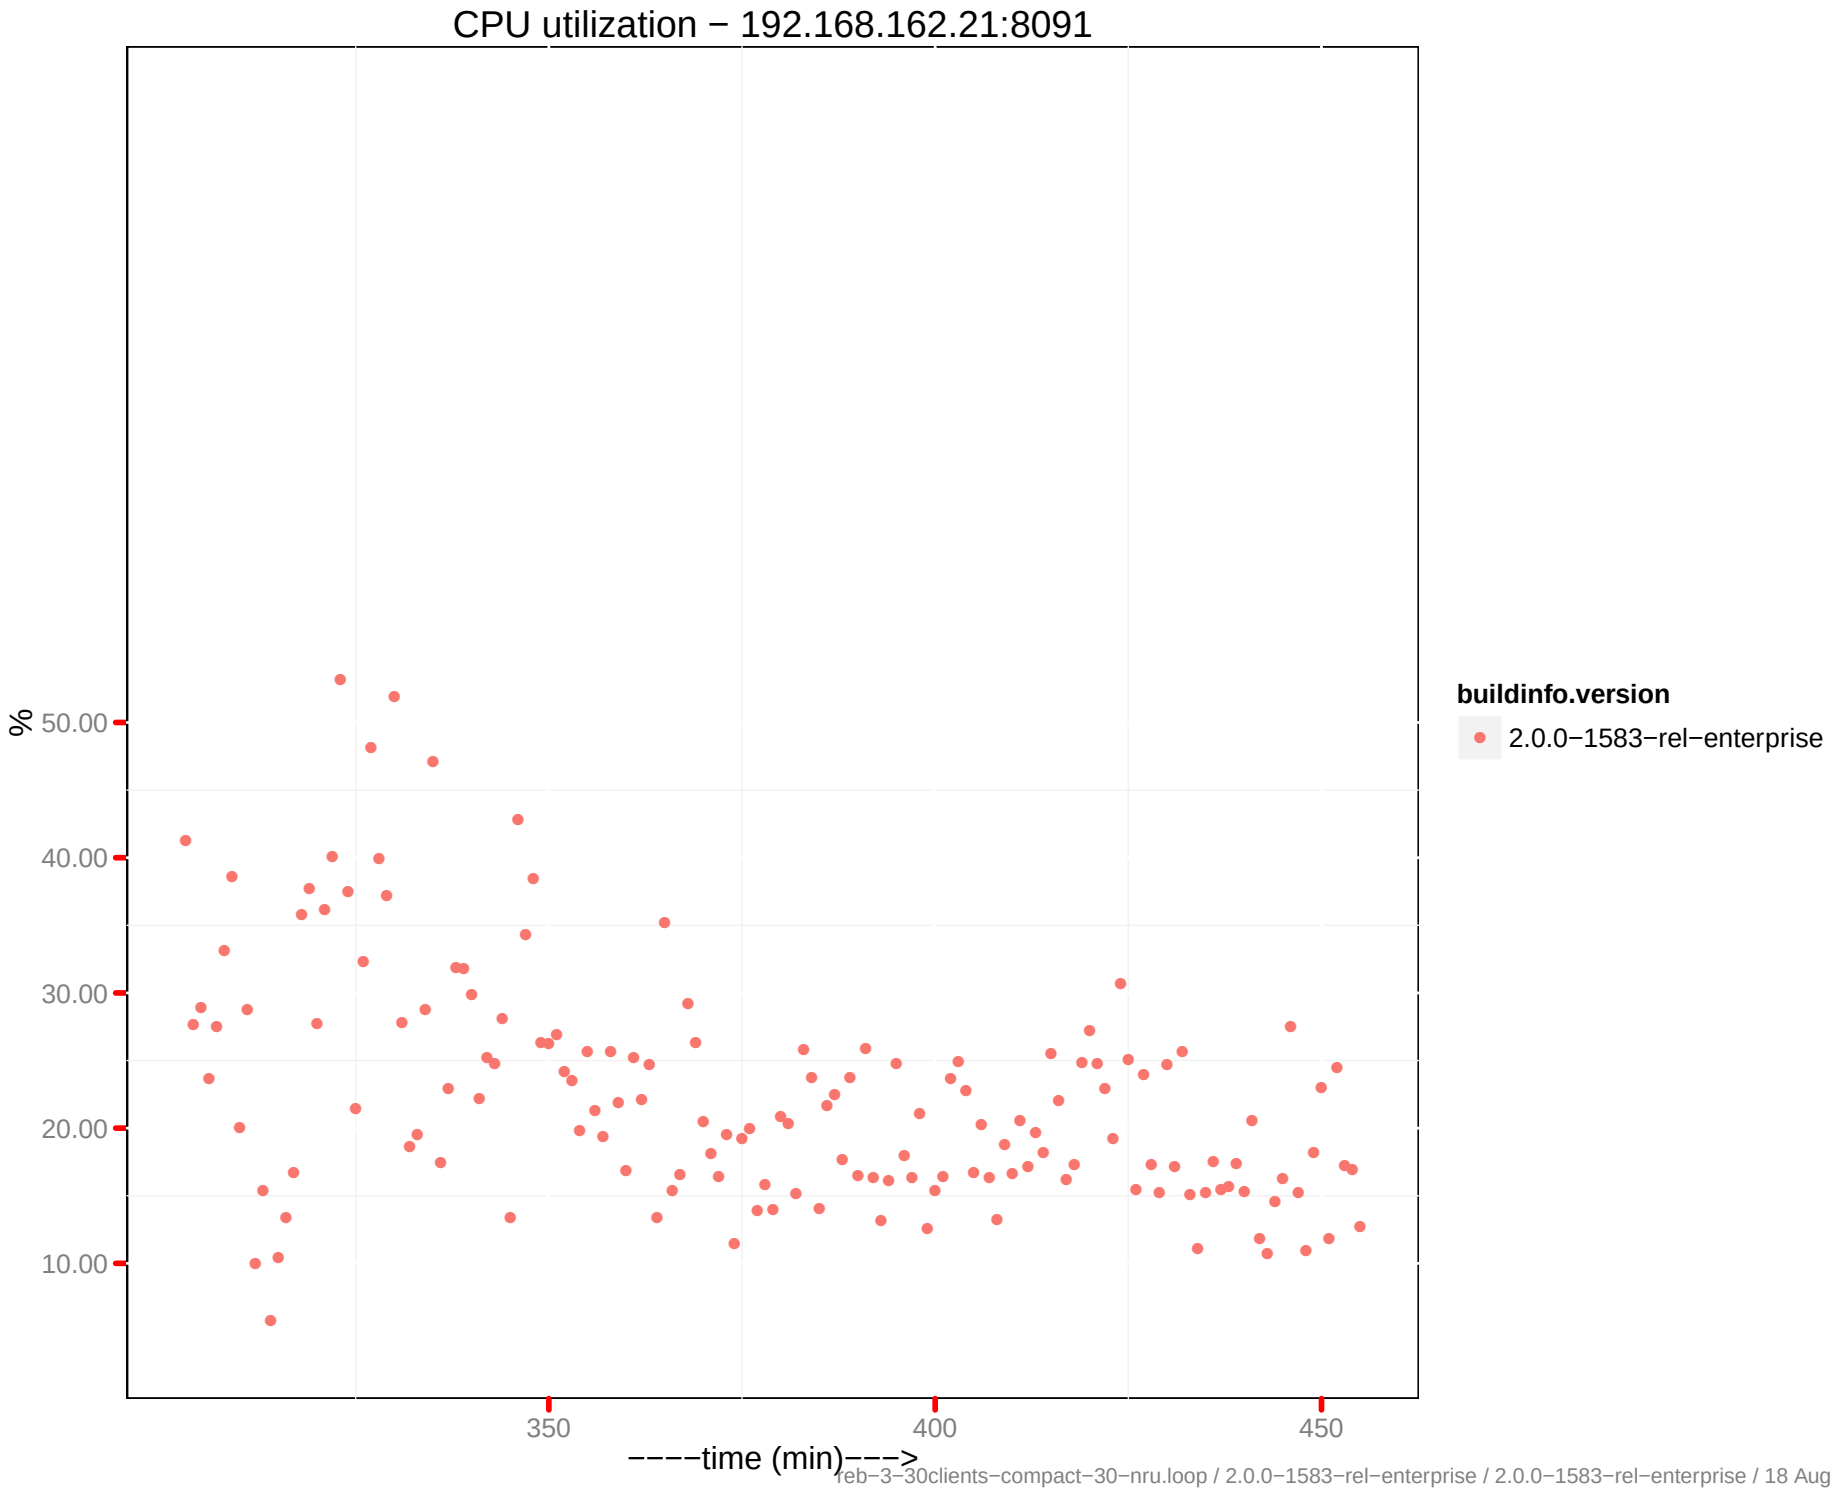
CPU utilization − 192.168.162.21:8091
%
50.00
40.00
30.00
20.00
10.00
350
400
450
−−−−time (min)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

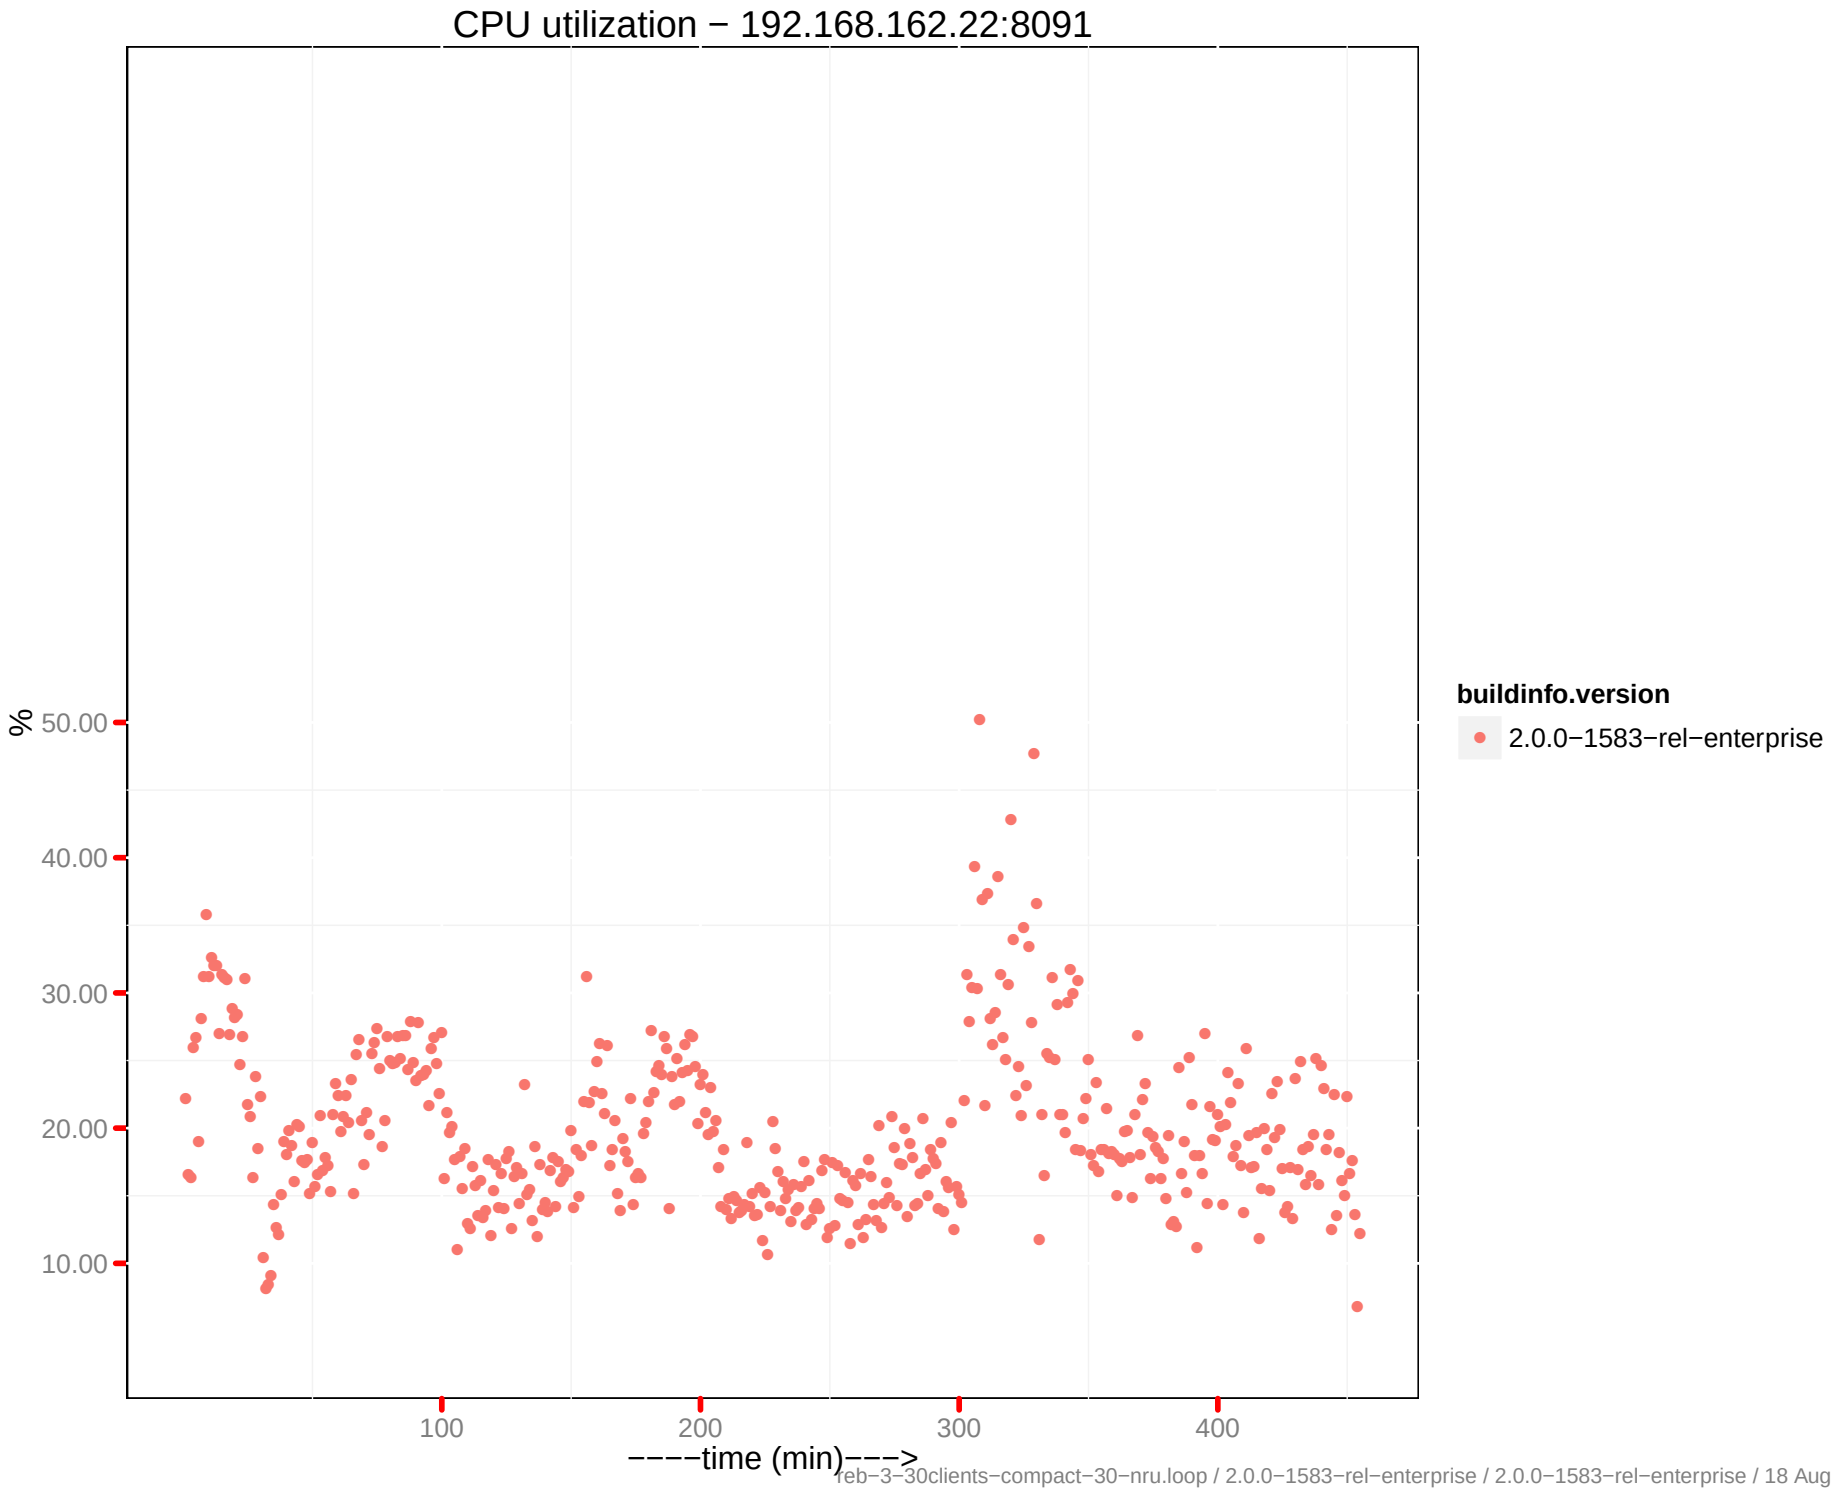
CPU utilization − 192.168.162.22:8091
%
50.00
40.00
30.00
20.00
10.00
100
200
300
400
−−−−time (min)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

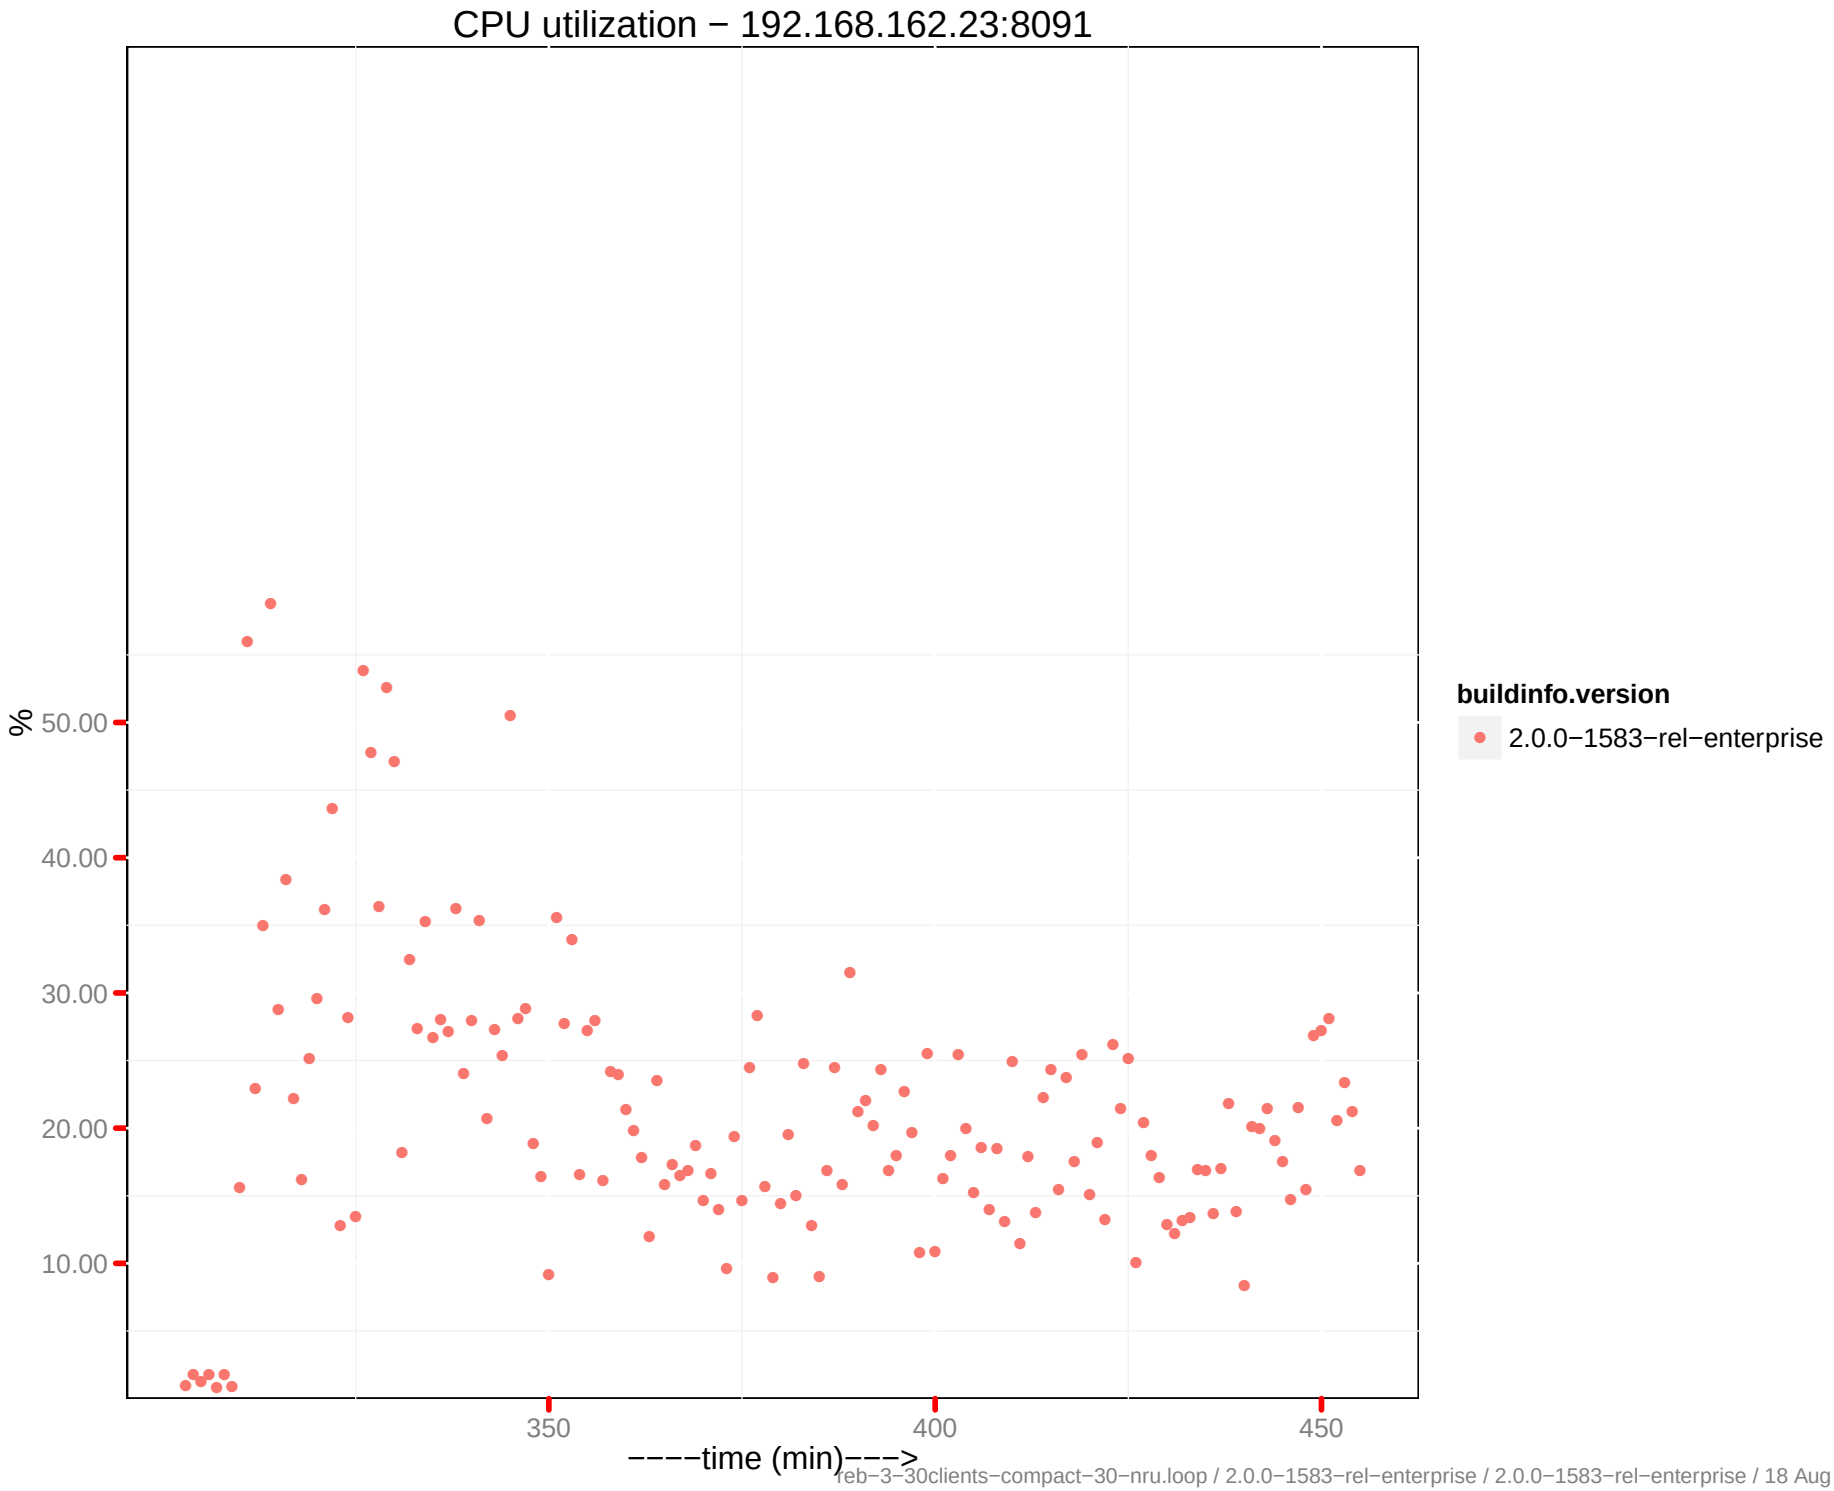
CPU utilization − 192.168.162.23:8091
%
50.00
40.00
30.00
20.00
10.00
350
400
450
−−−−time (min)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

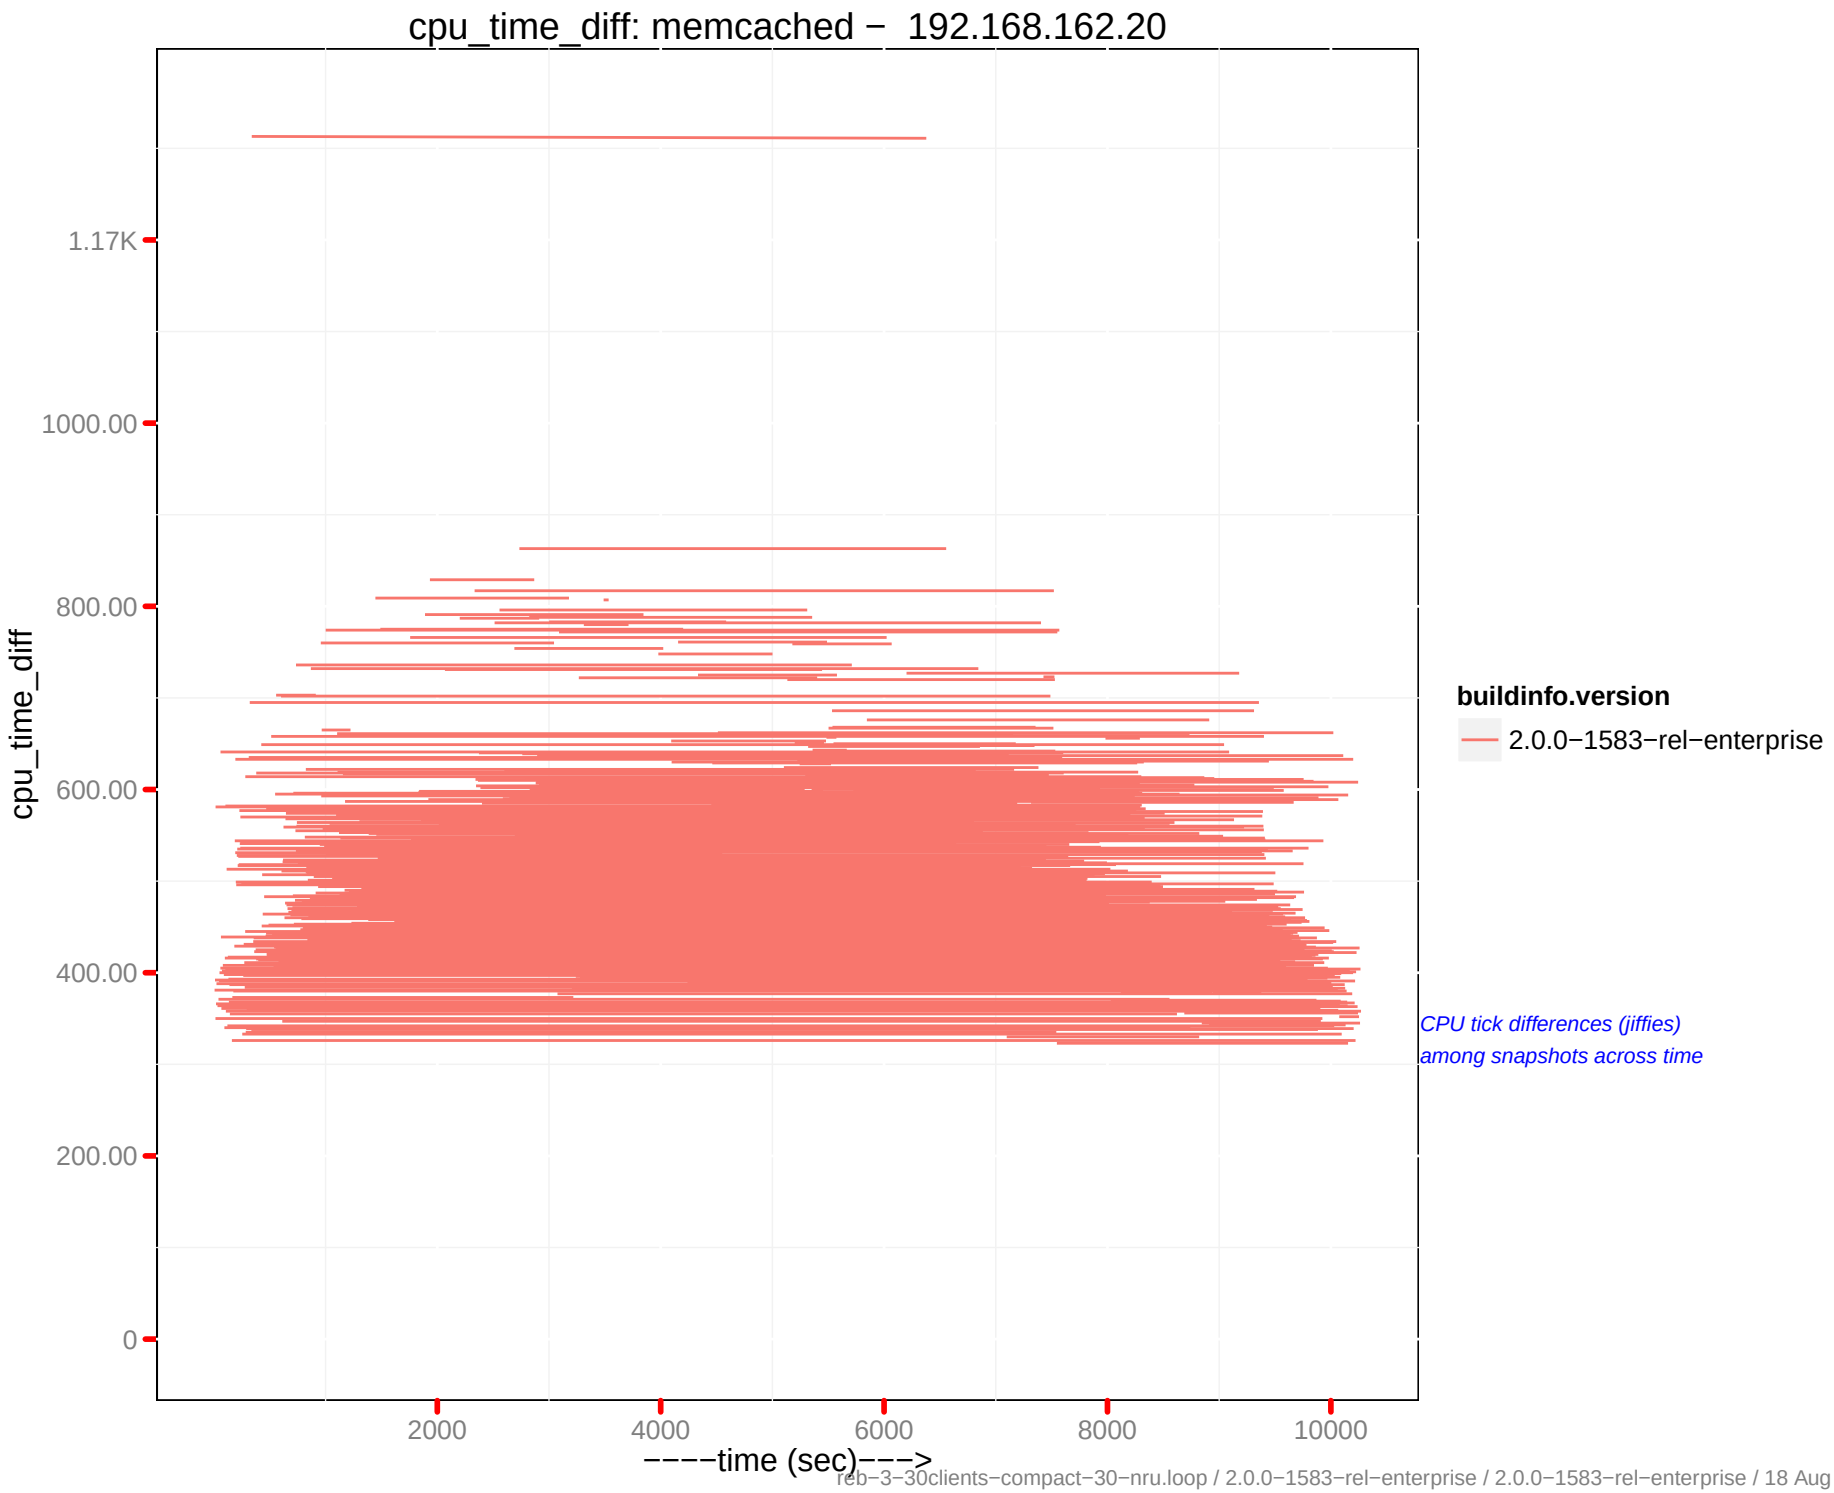
cpu_time_diff: memcached − 192.168.162.20
1.17K
1000.00
800.00
600.00
400.00
200.00
0
cpu_time_diff
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

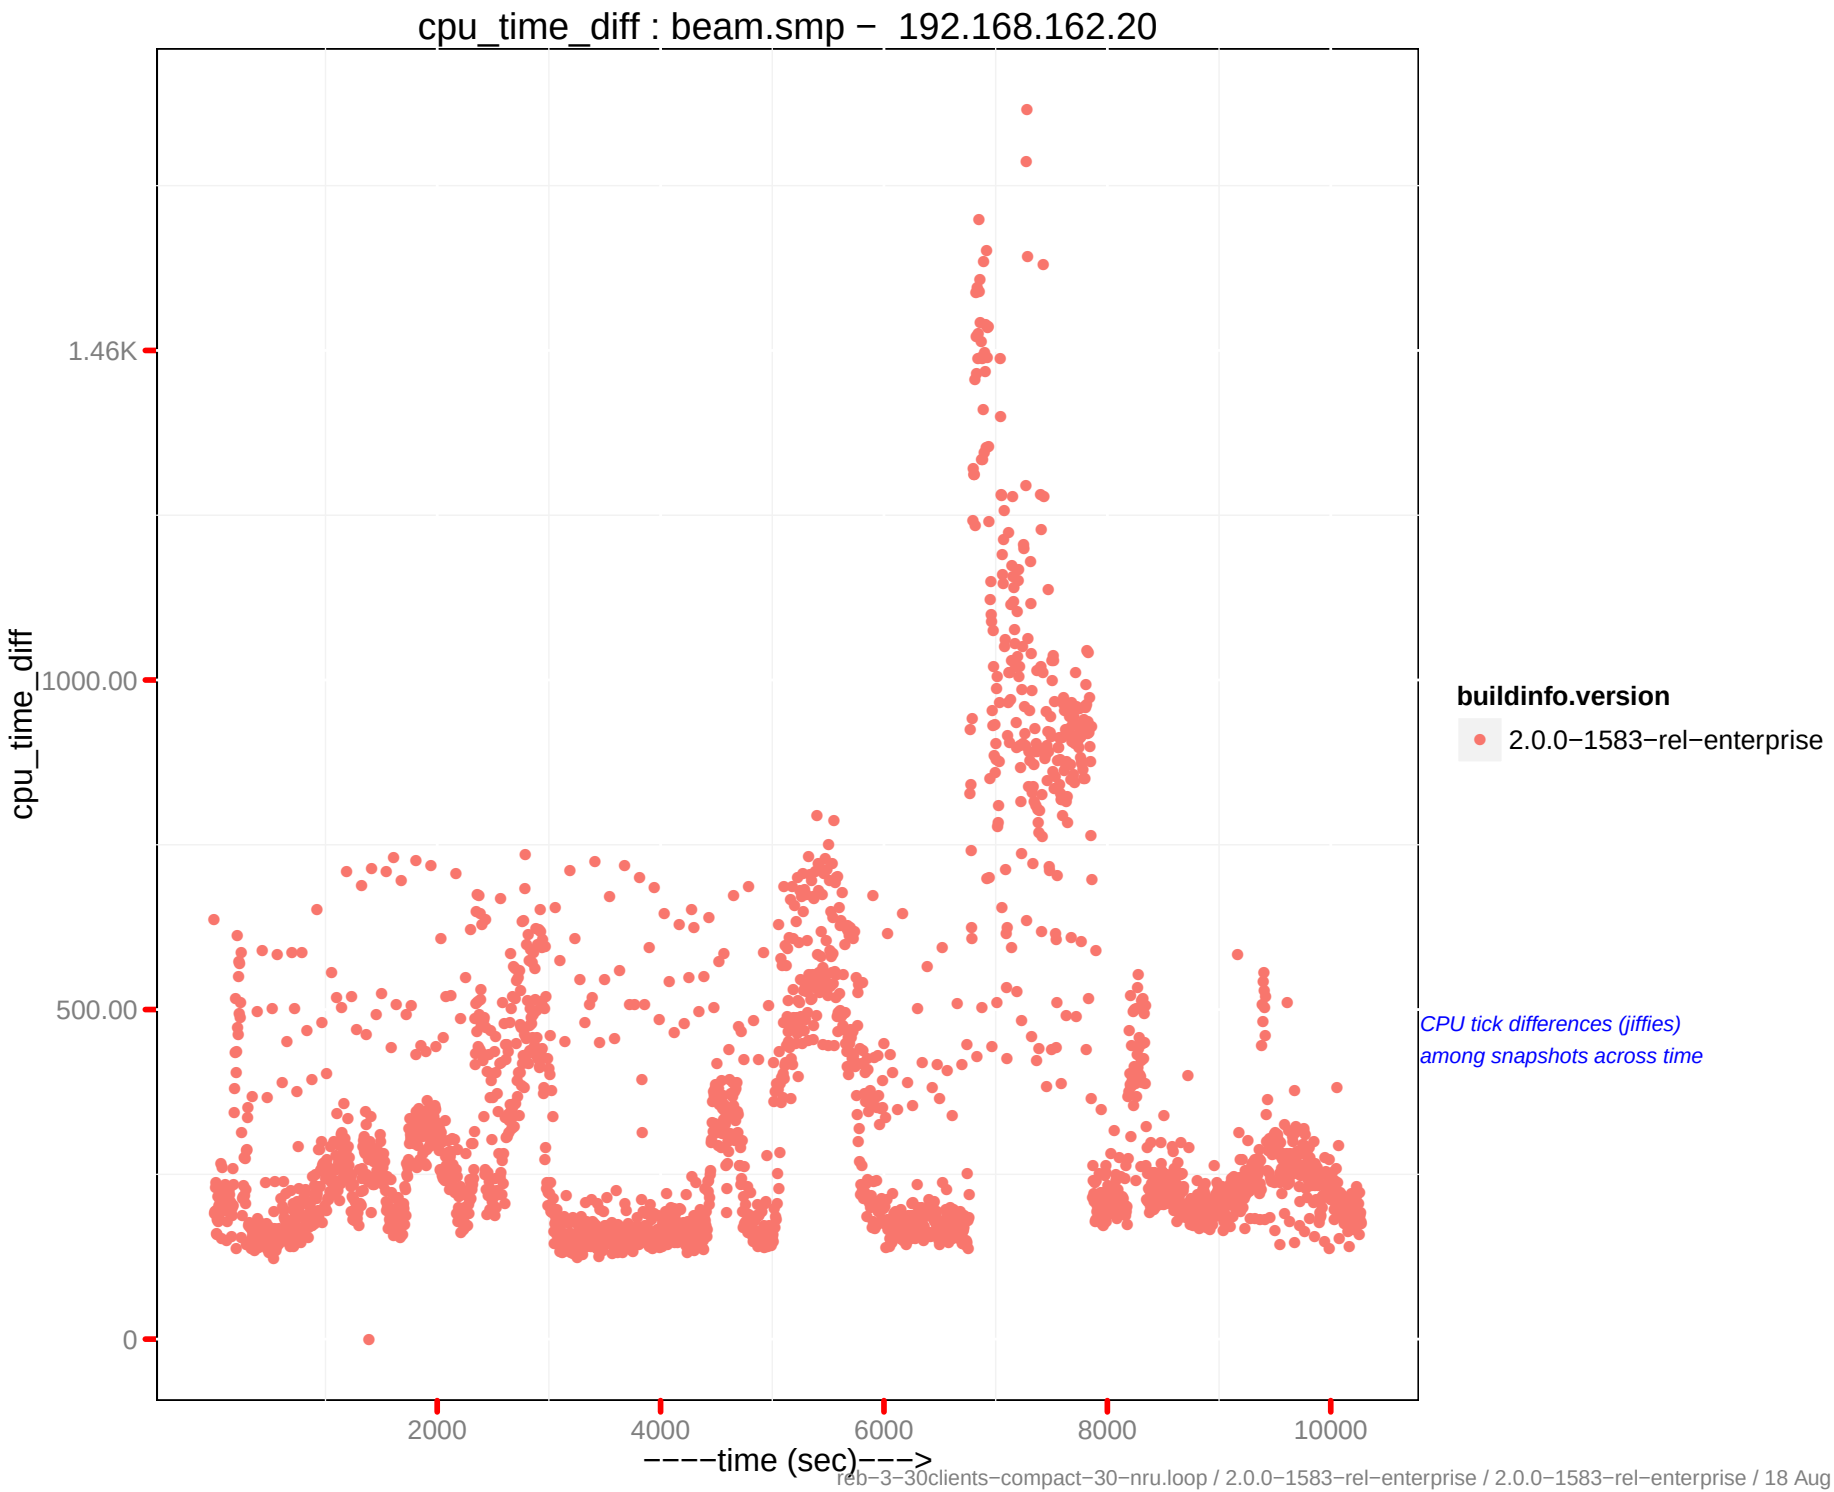
cpu_time_diff : beam.smp − 192.168.162.20
cpu_time_diff
1.46K
1000.00
500.00
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

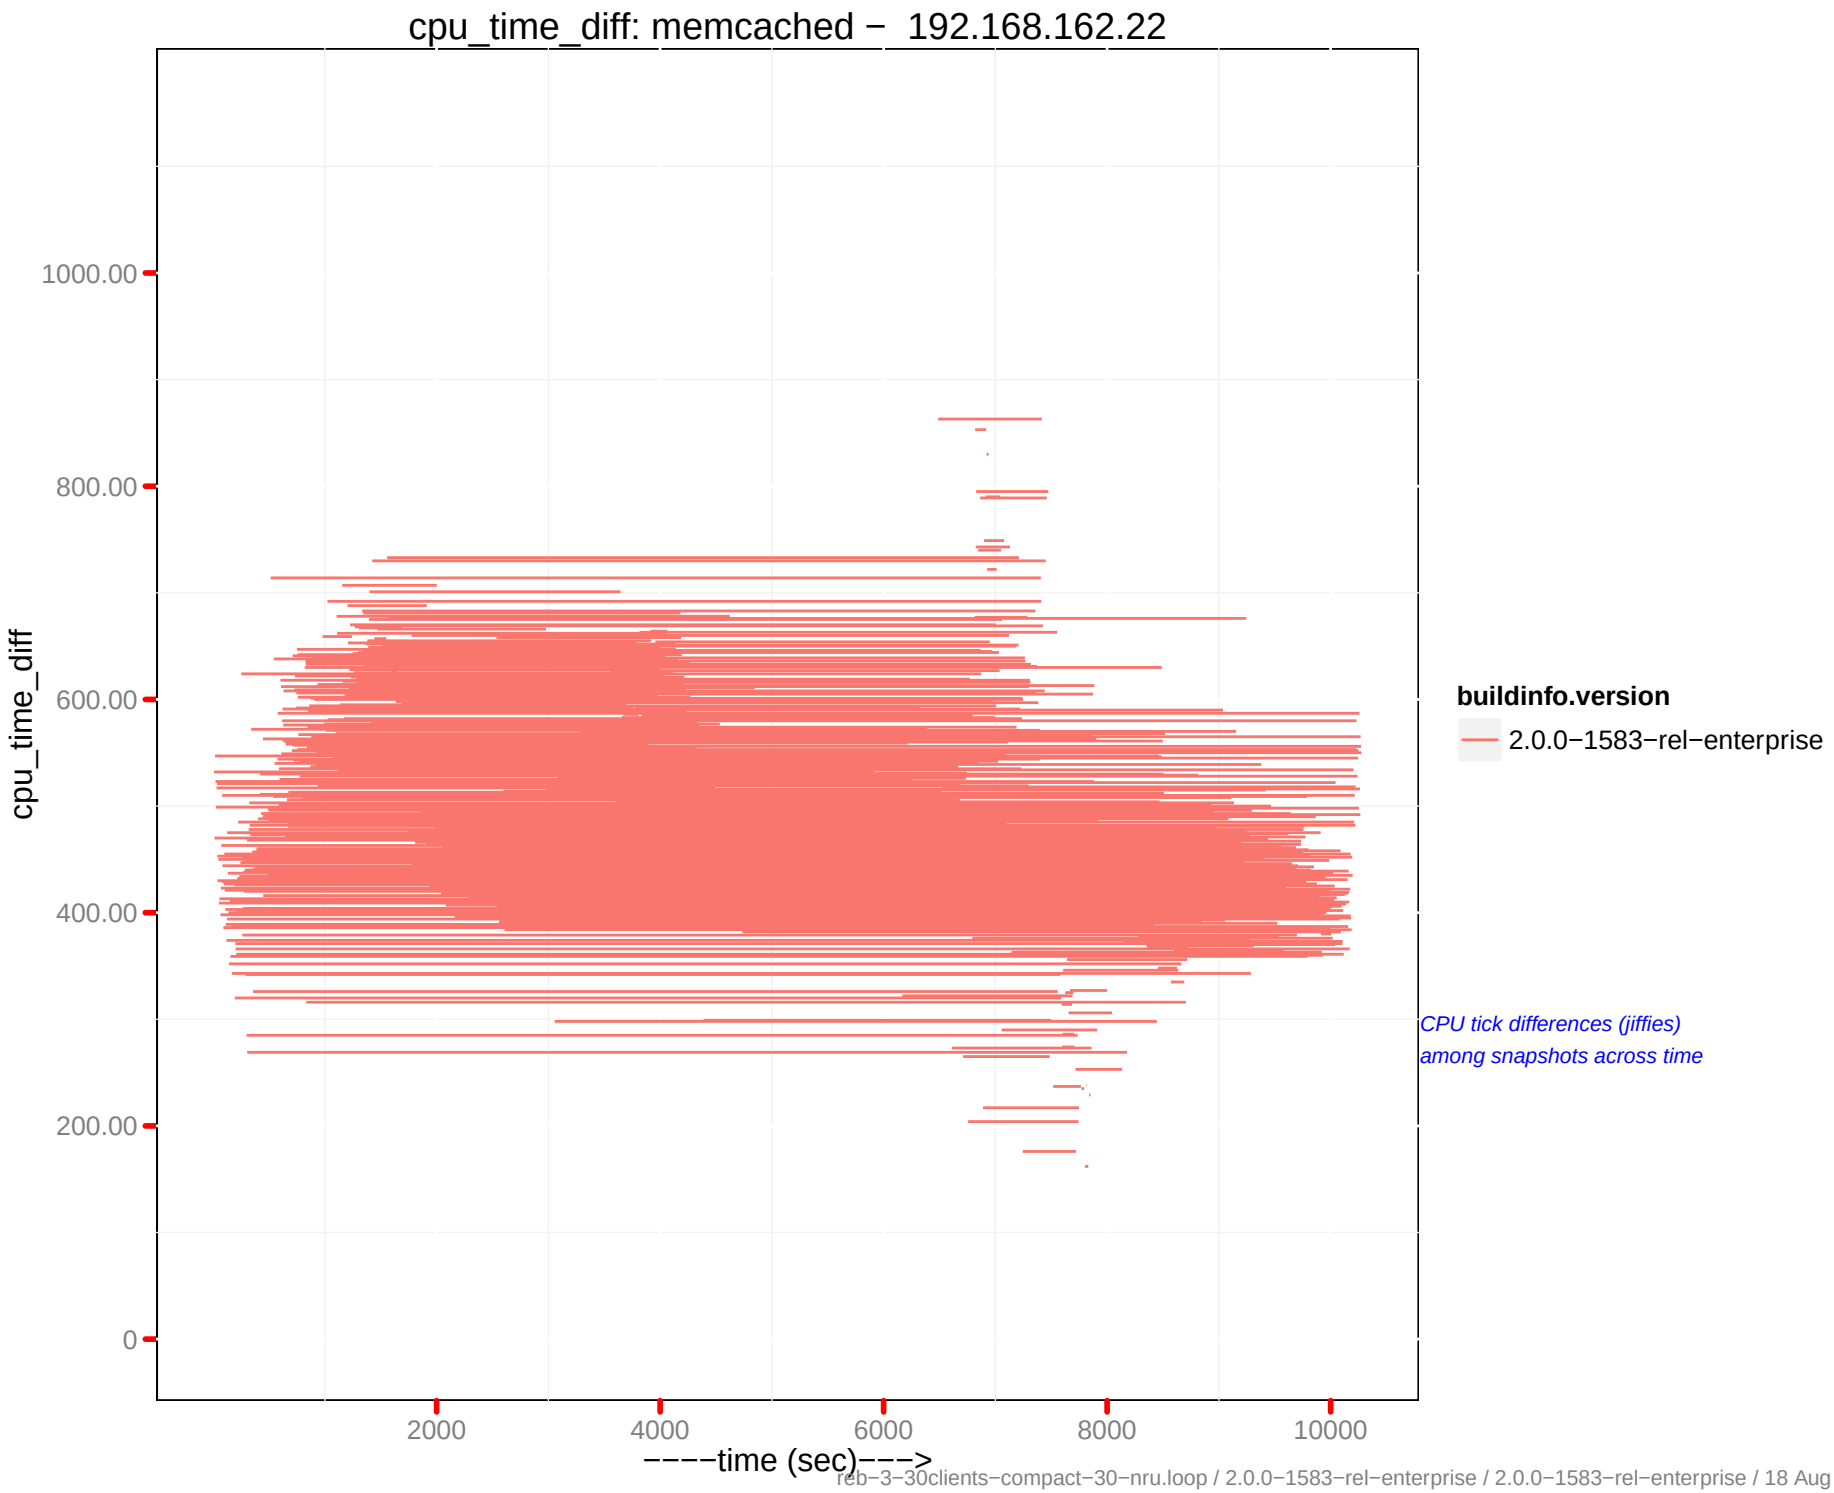
cpu_time_diff: memcached − 192.168.162.22
cpu_time_diff
1000.00
800.00
600.00
400.00
200.00
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

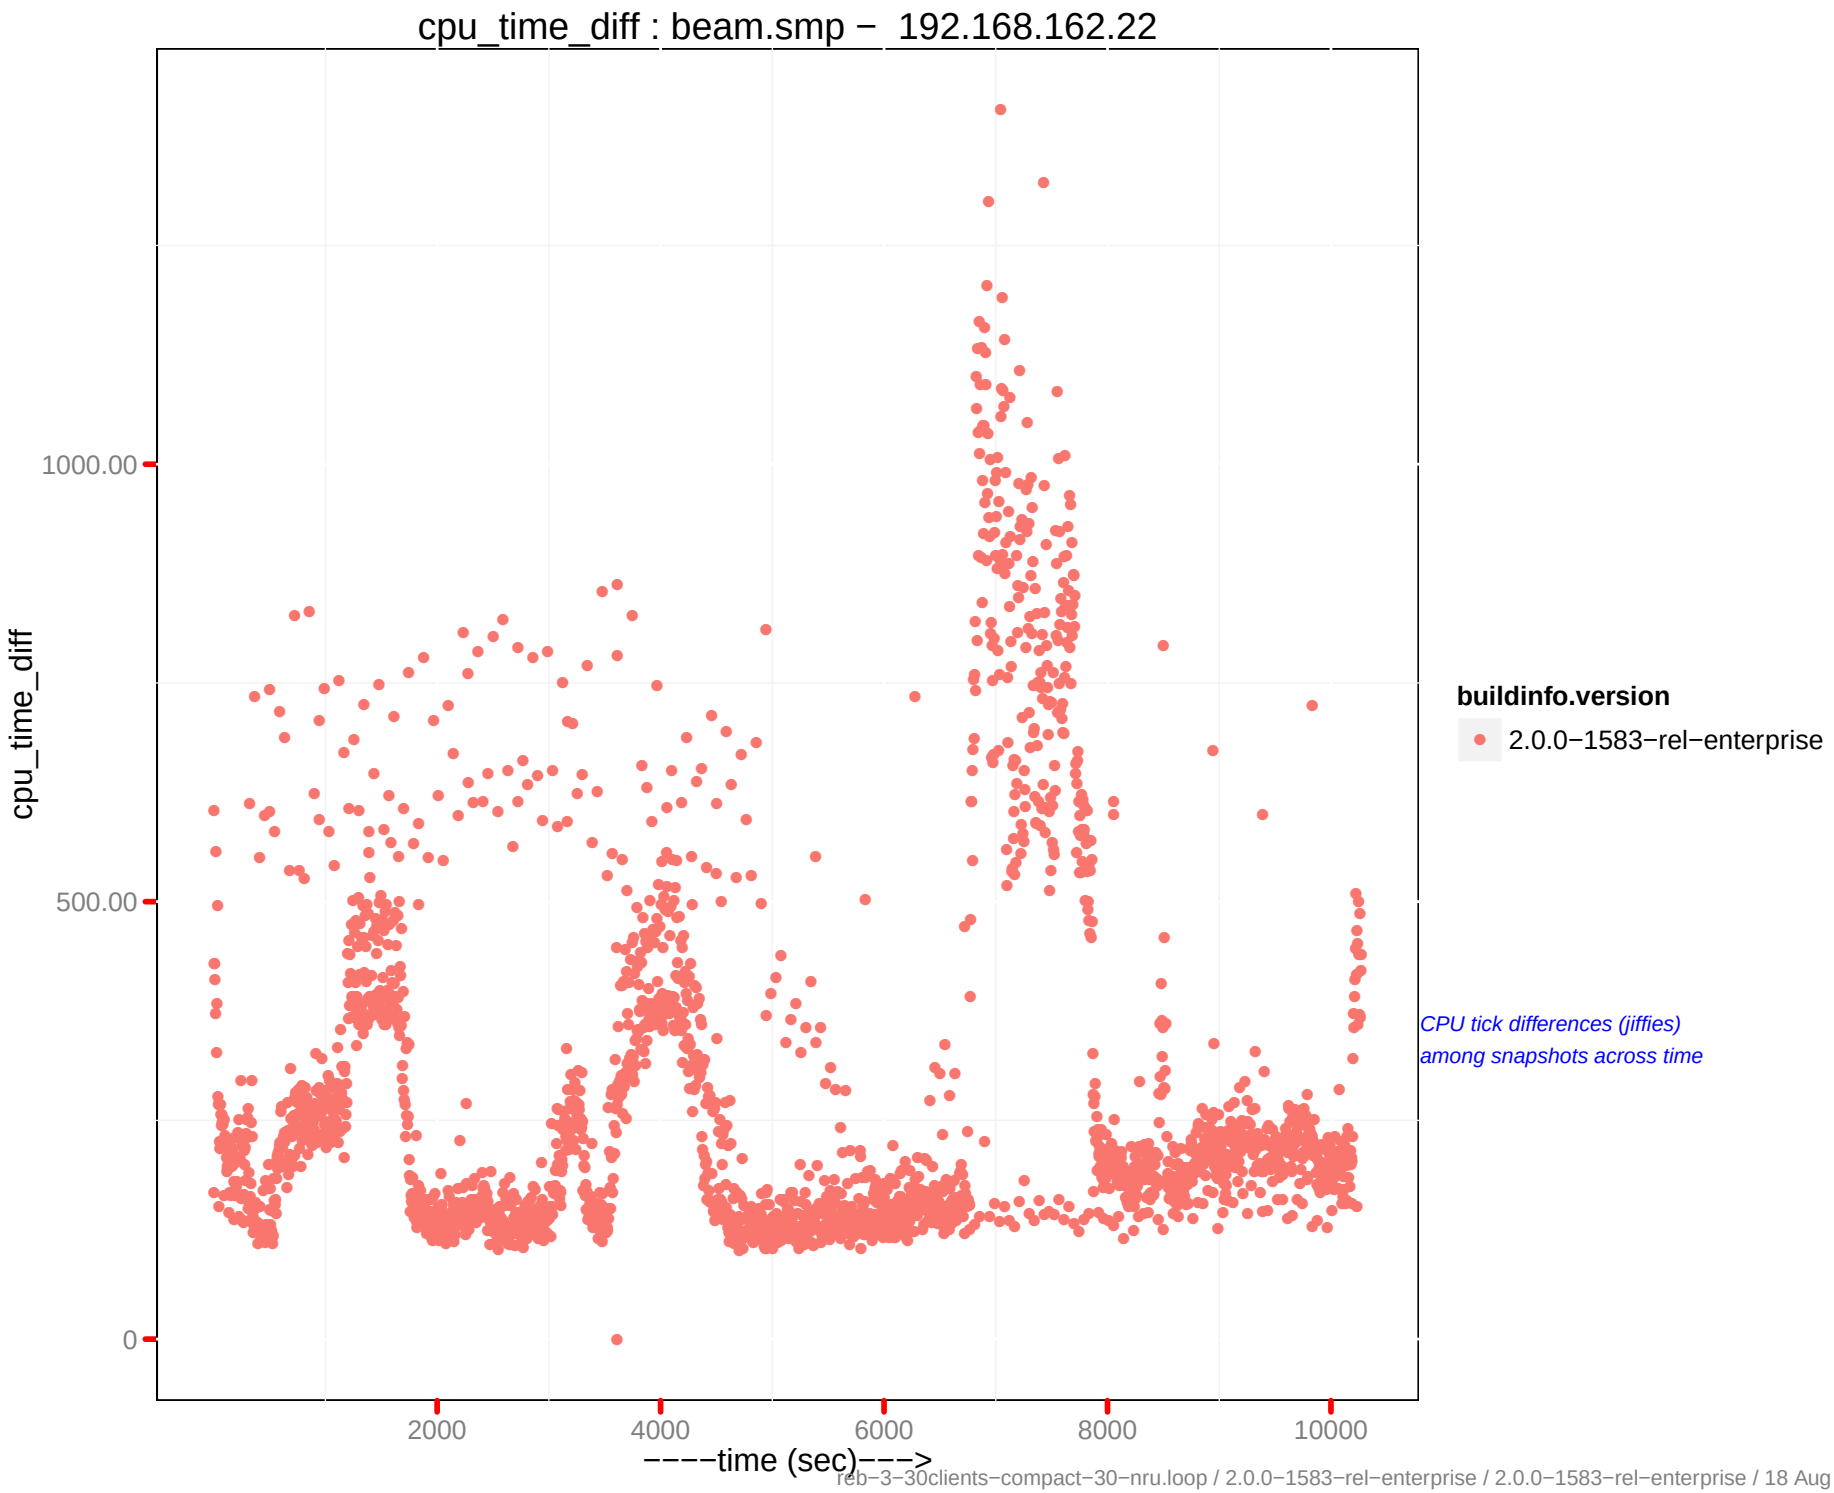
cpu_time_diff : beam.smp − 192.168.162.22
cpu_time_diff
1000.00
500.00
0
2000
4000
6000
8000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

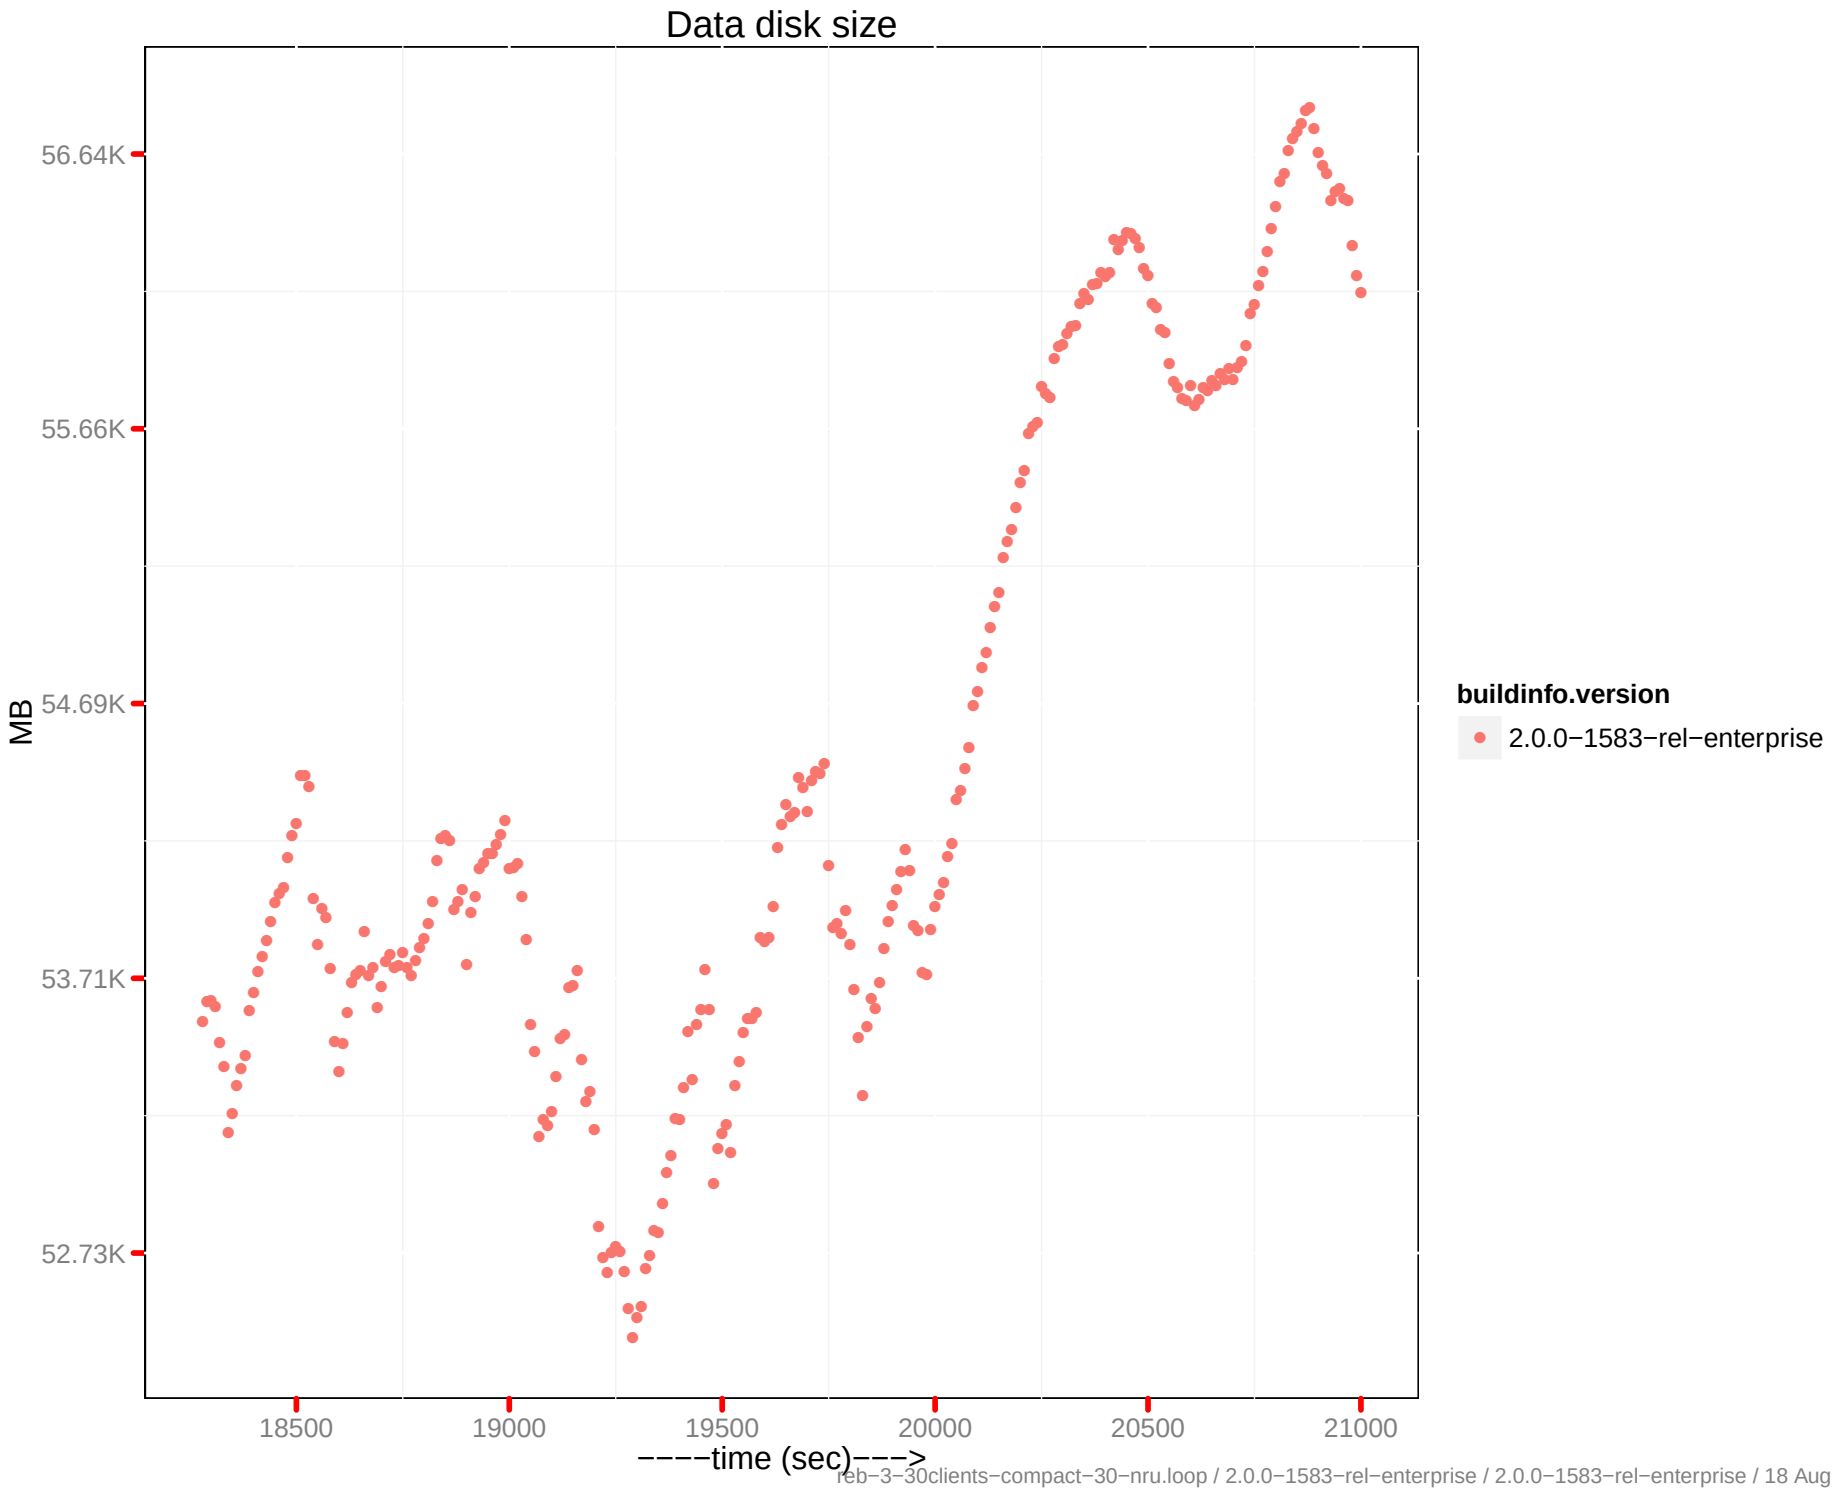
Data disk size
56.64K
55.66K
54.69K
53.71K
52.73K
MB
18500
19000
19500
20000
20500
21000
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

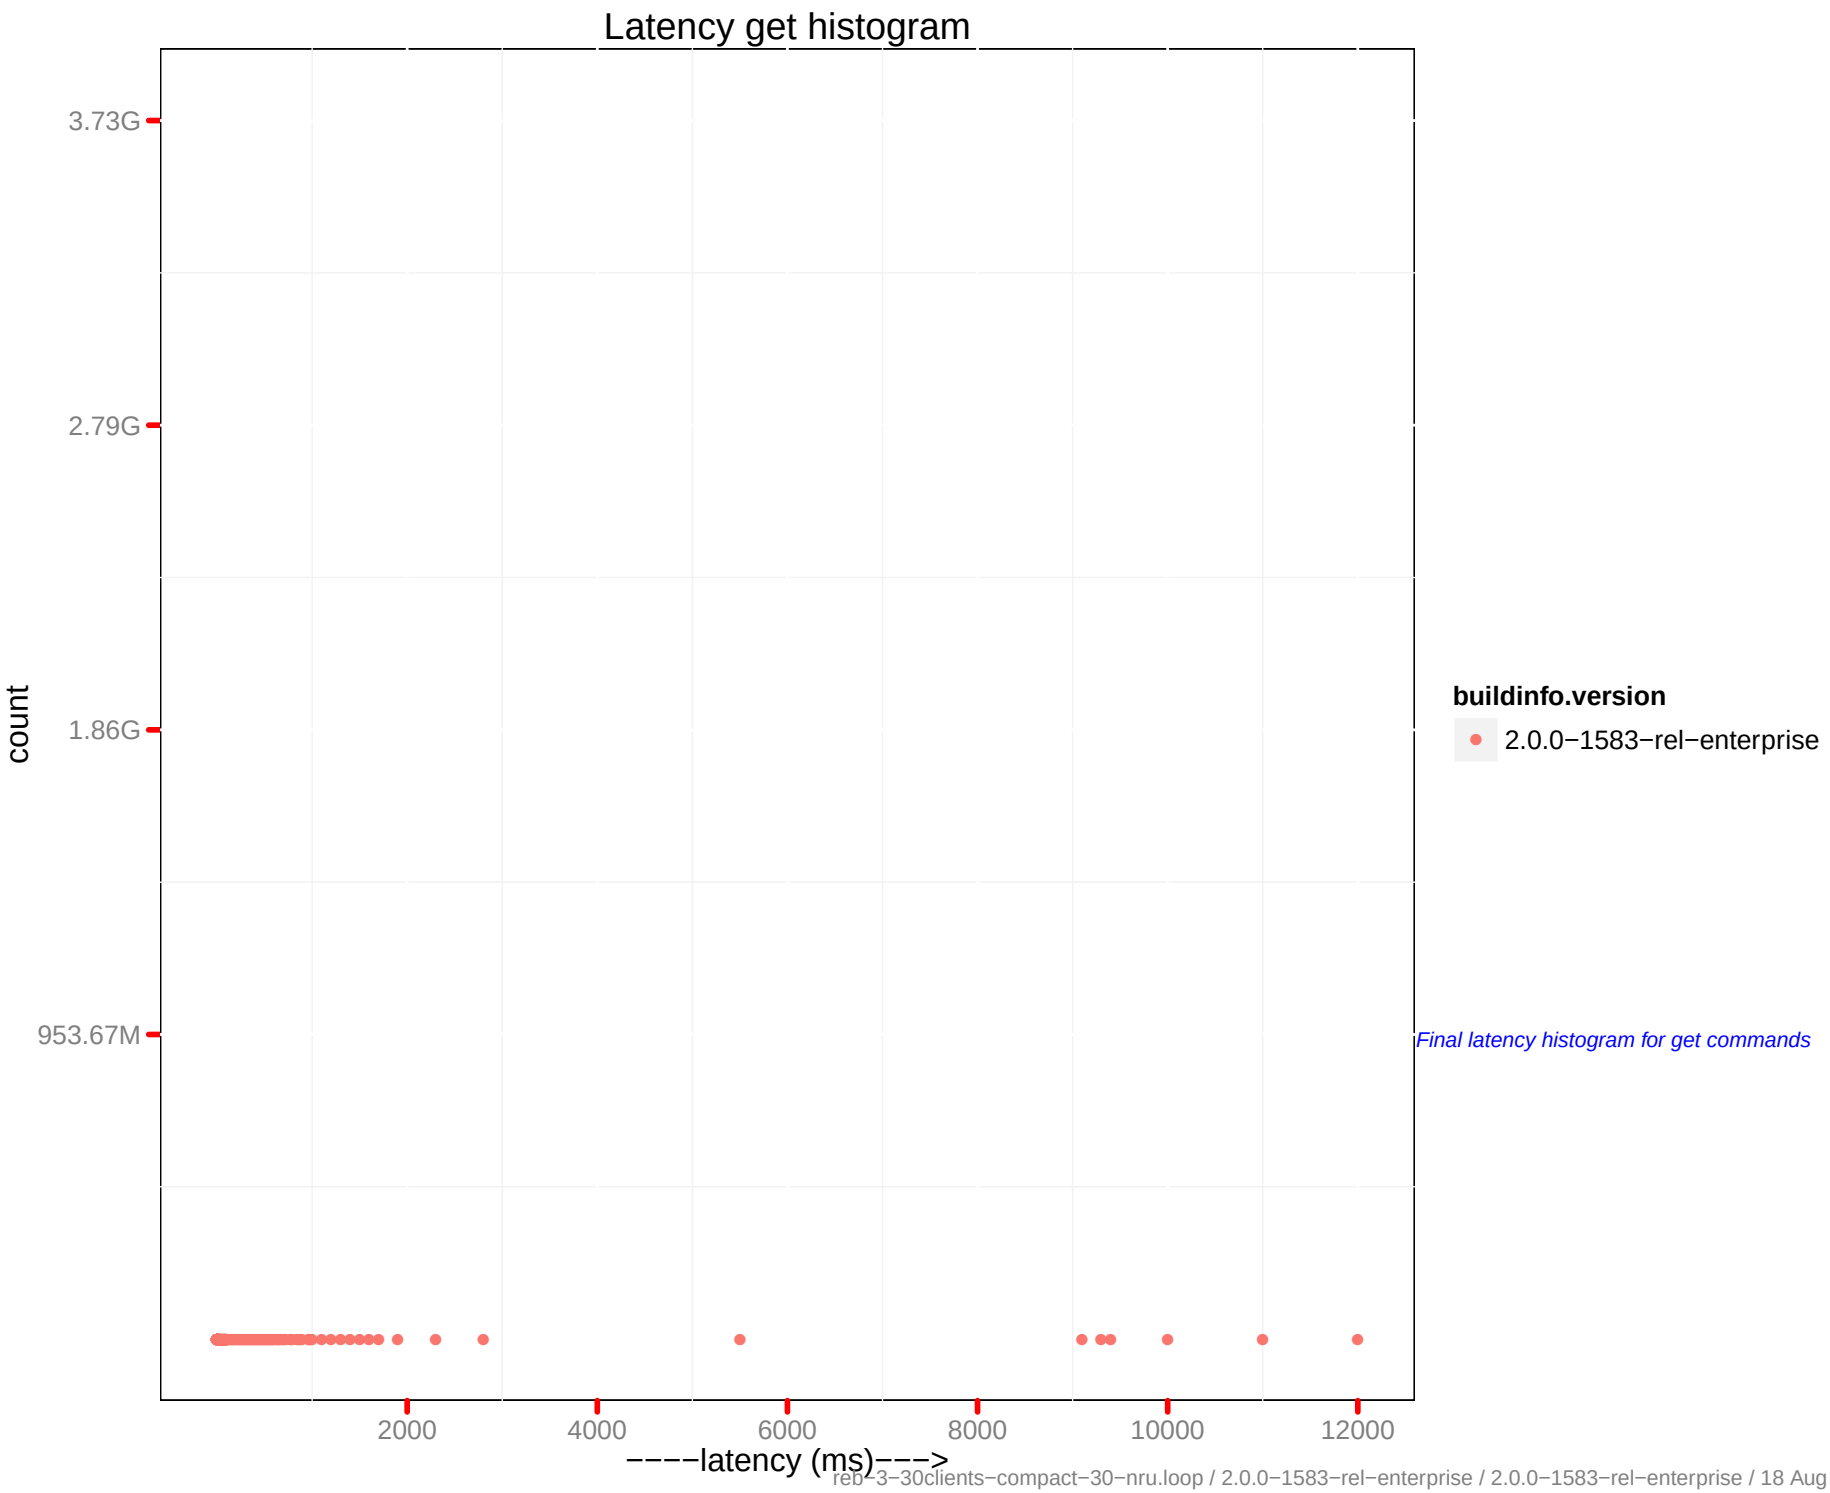
Latency get histogram
3.73G
2.79G
1.86G
953.67M
count
2000
4000
6000
8000
10000
12000
----latency (ms)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Final latency histogram for get commands
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

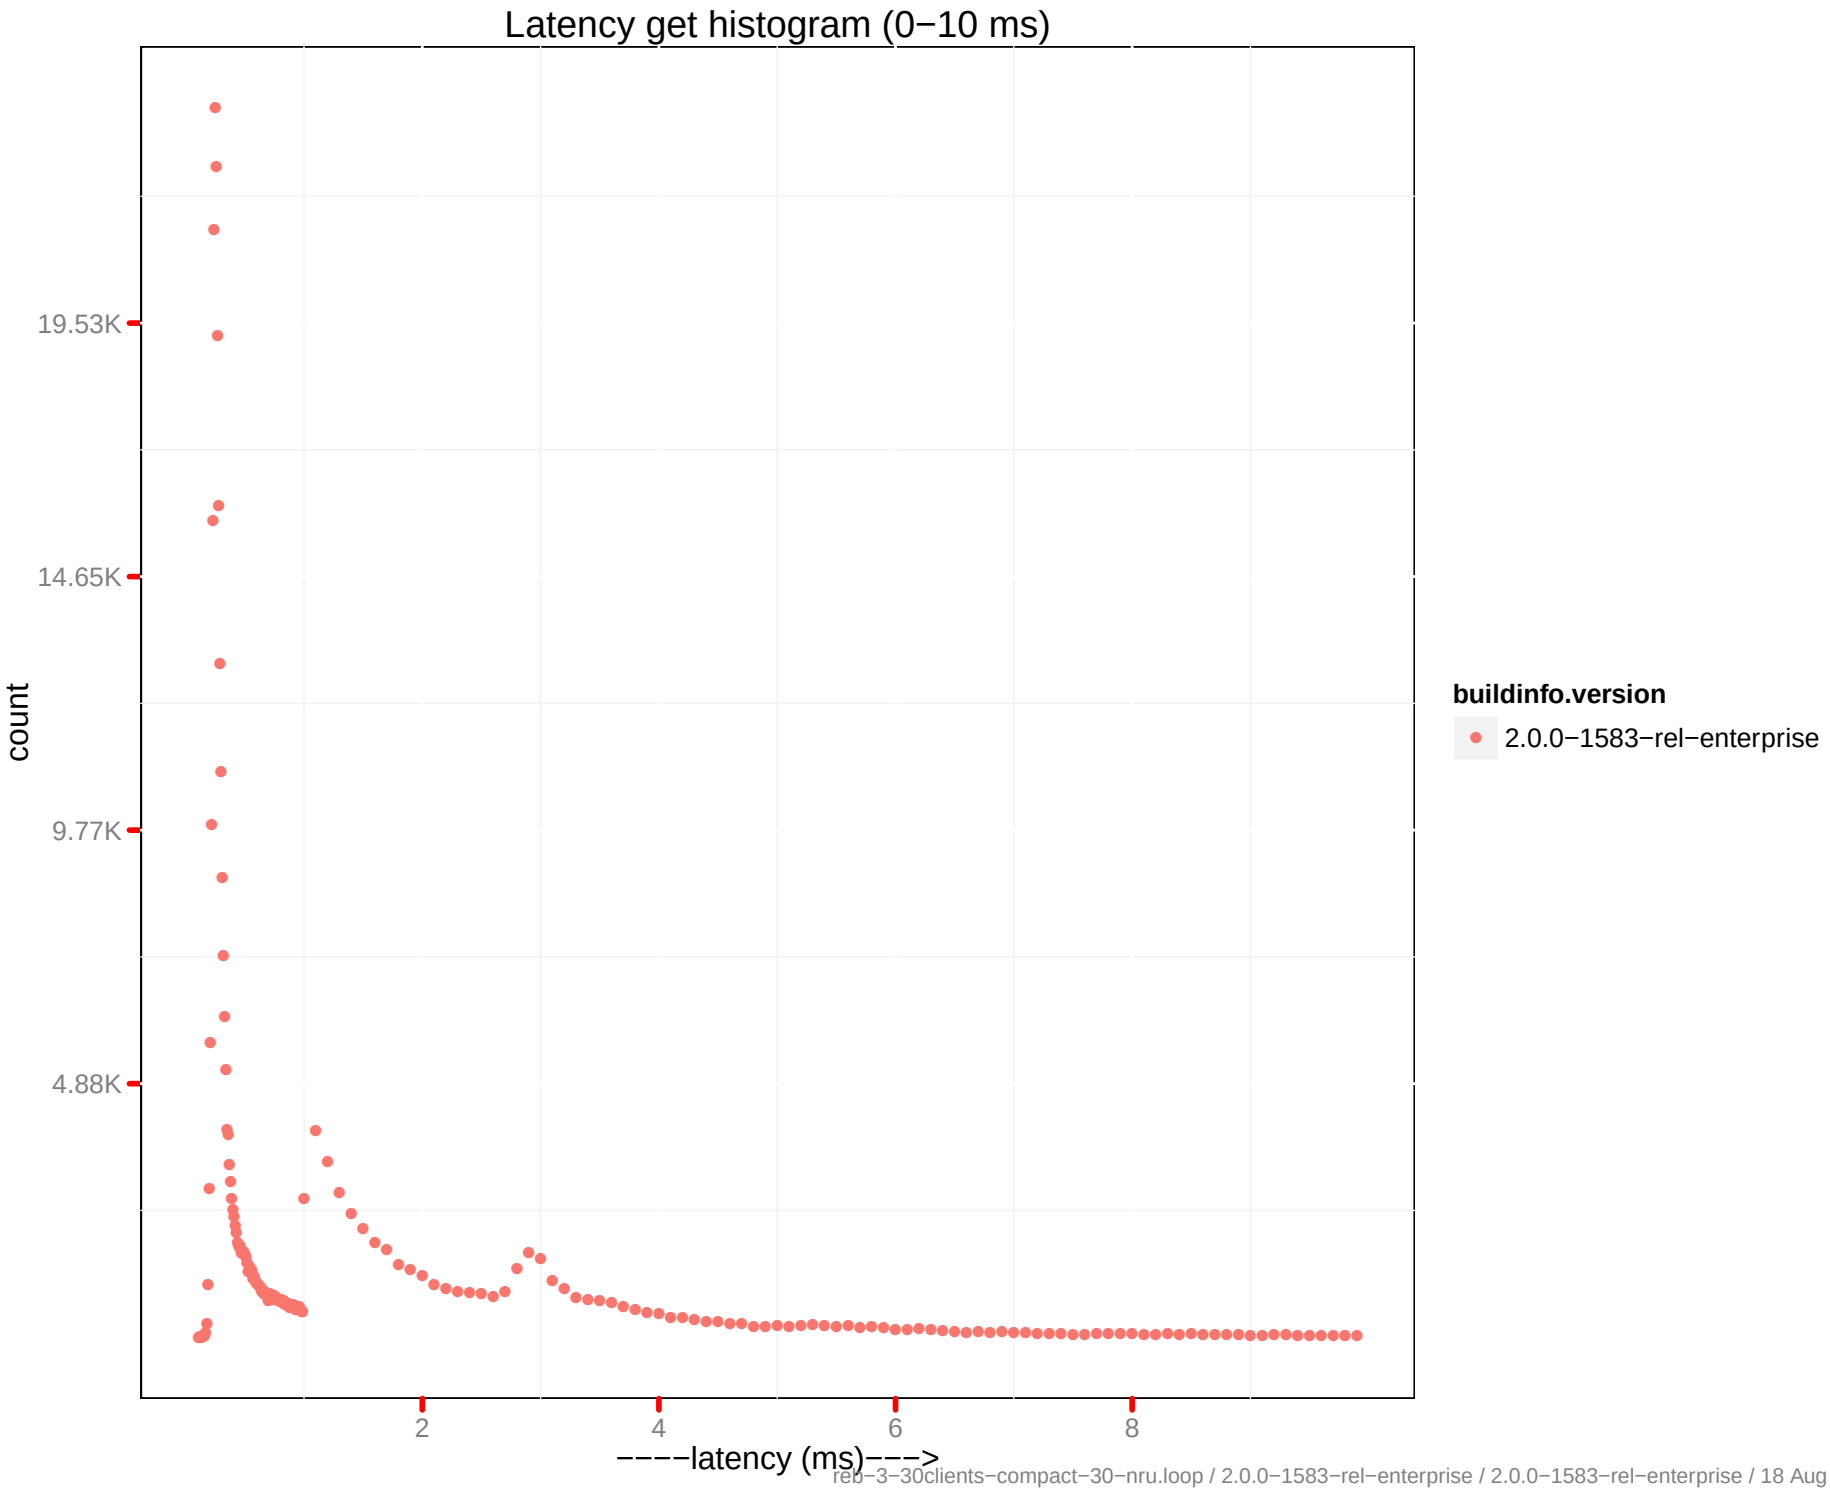
Latency get histogram (0−10 ms)
19.53K
14.65K
9.77K
4.88K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

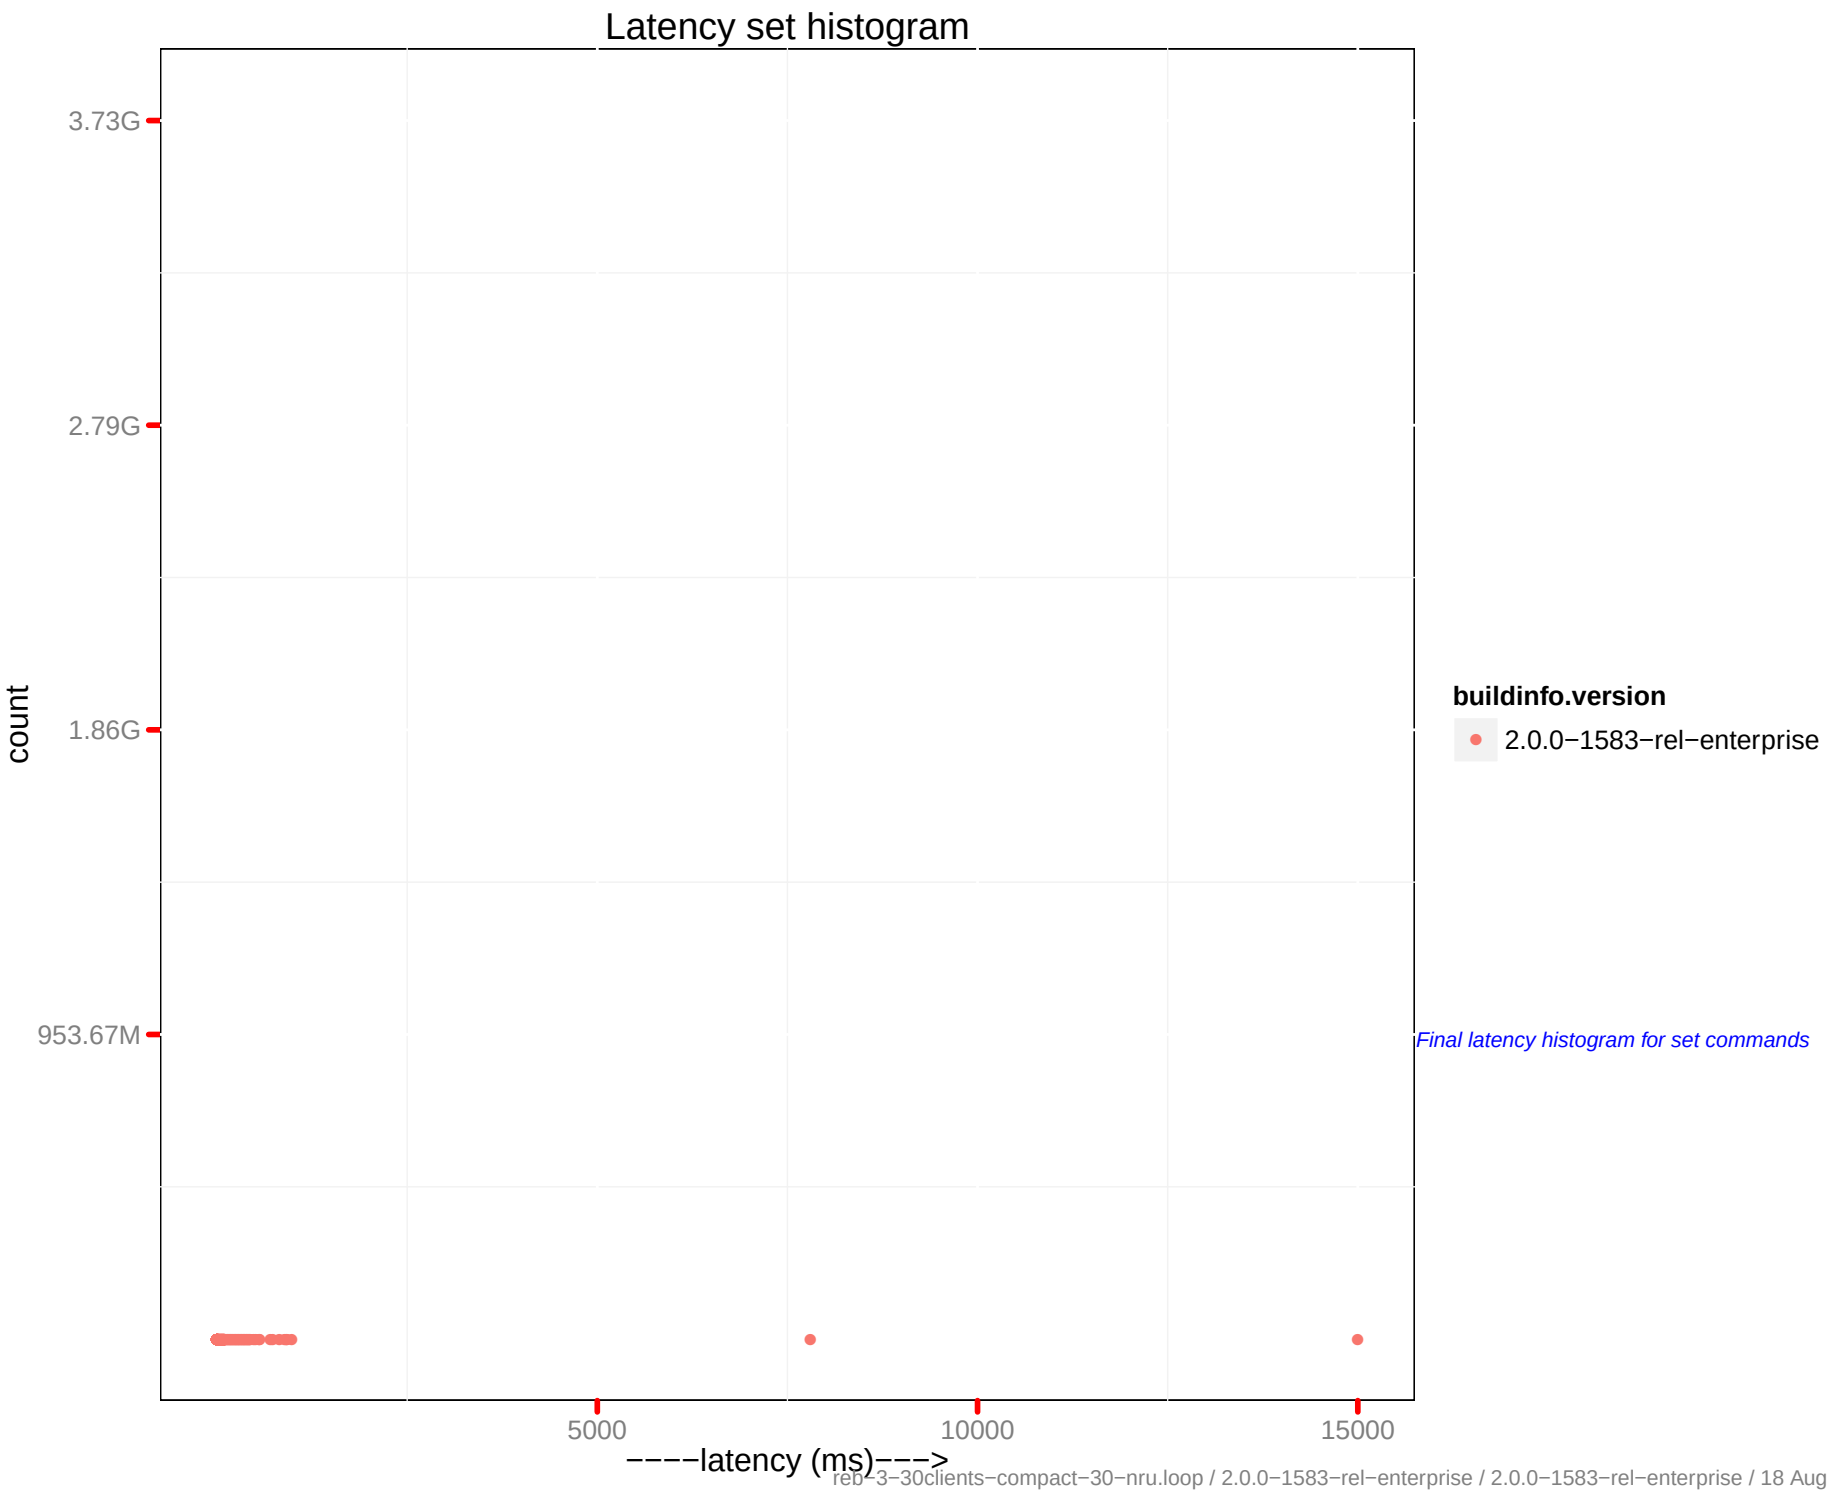
Latency set histogram
3.73G
2.79G
1.86G
953.67M
count
5000
10000
15000
----latency (ms)--->
buildinfo.version
2.0.0-1583-rel-enterprise
Final latency histogram for set commands
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

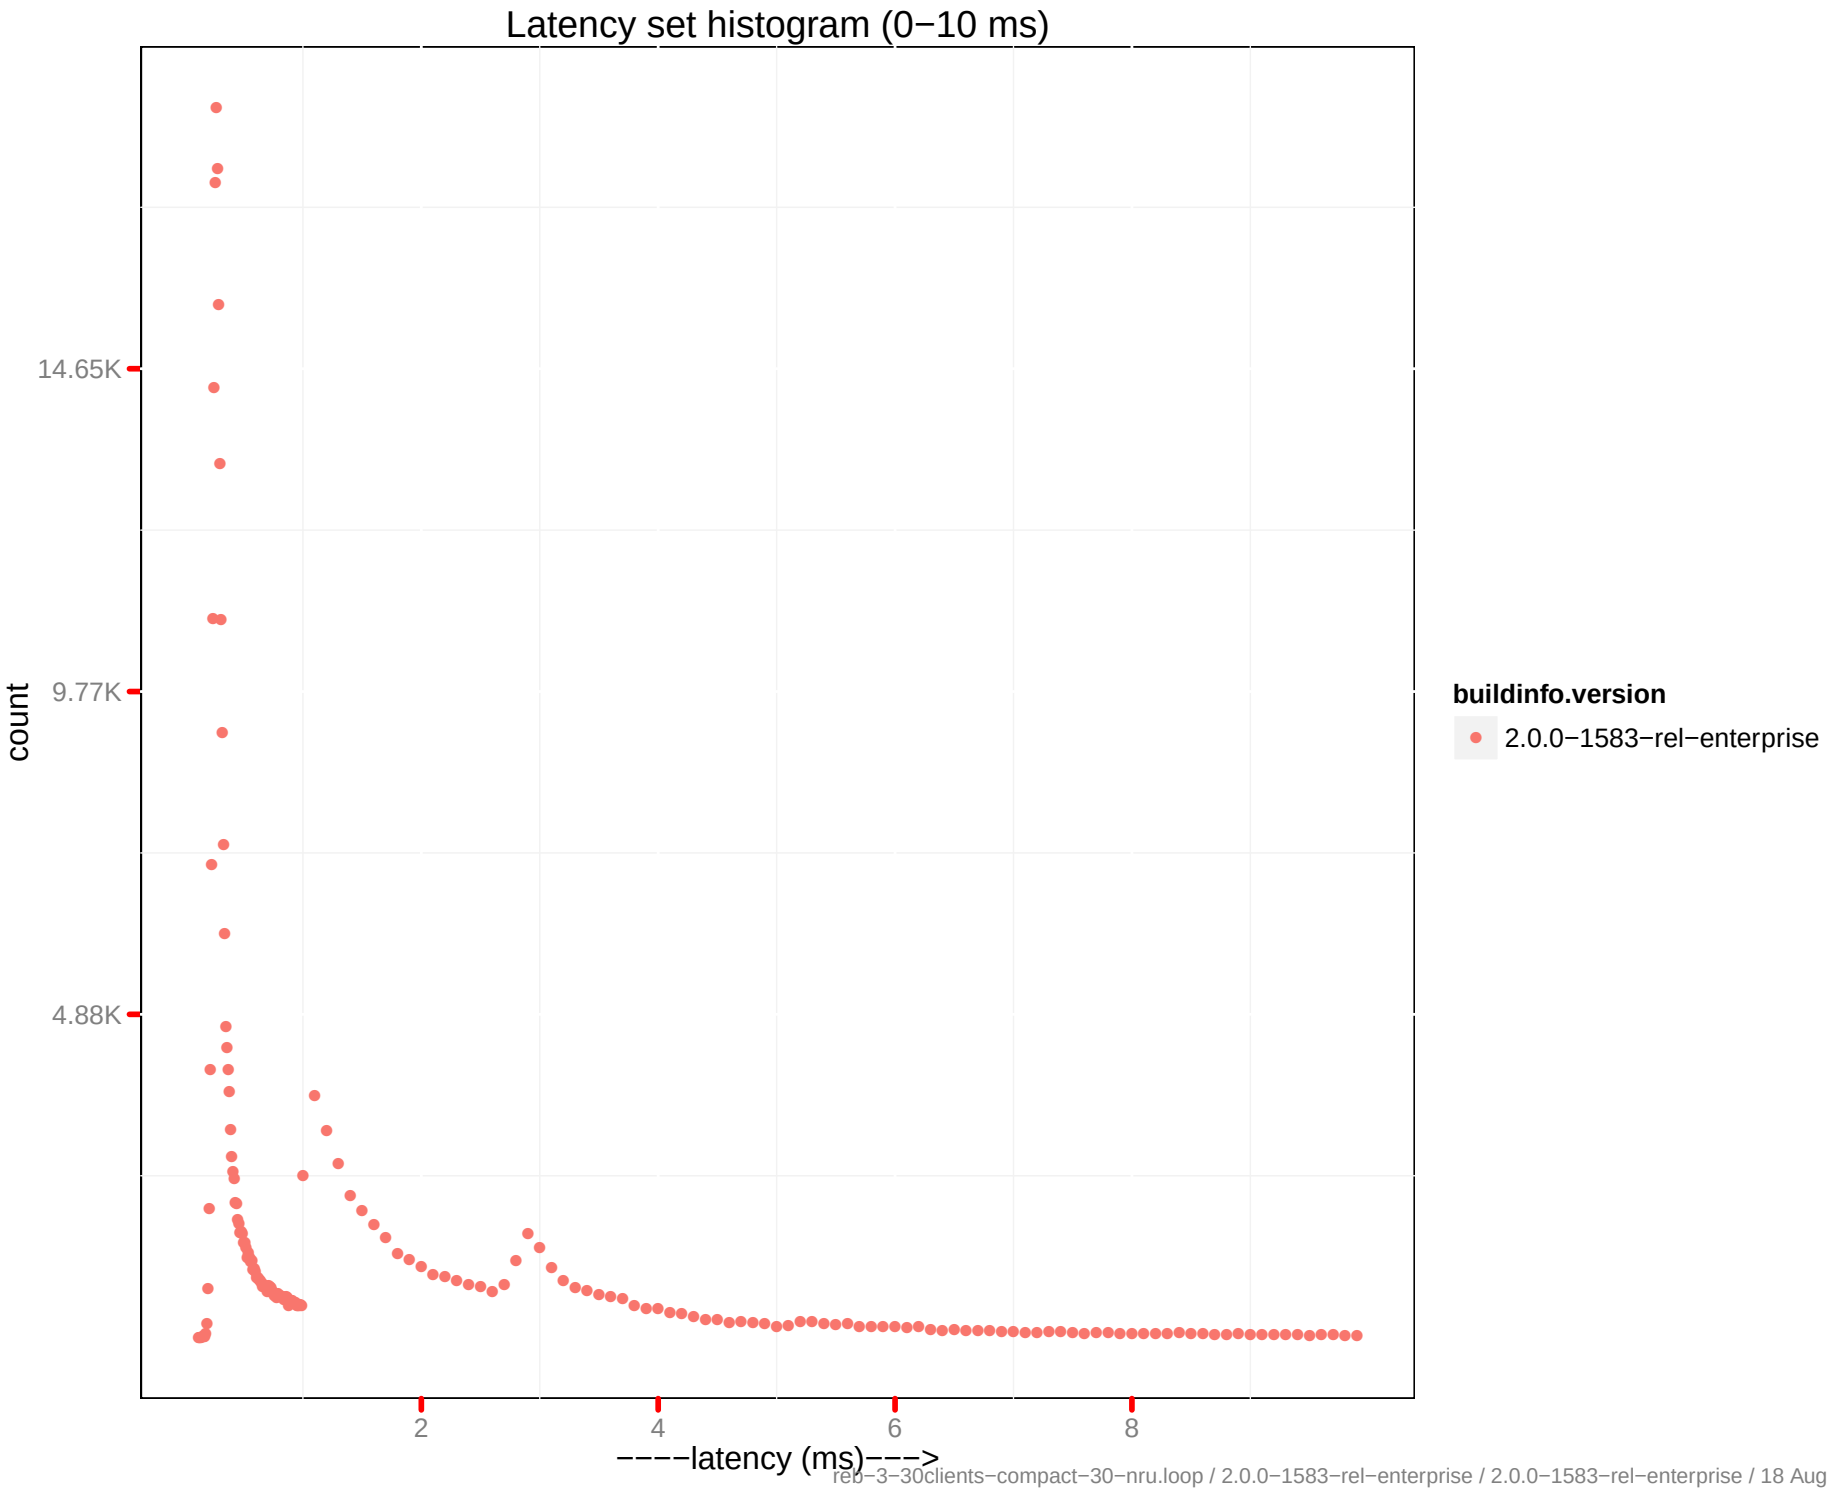
Latency set histogram (0−10 ms)
14.65K
9.77K
4.88K
count
2
4
6
8
−−−−latency (ms)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

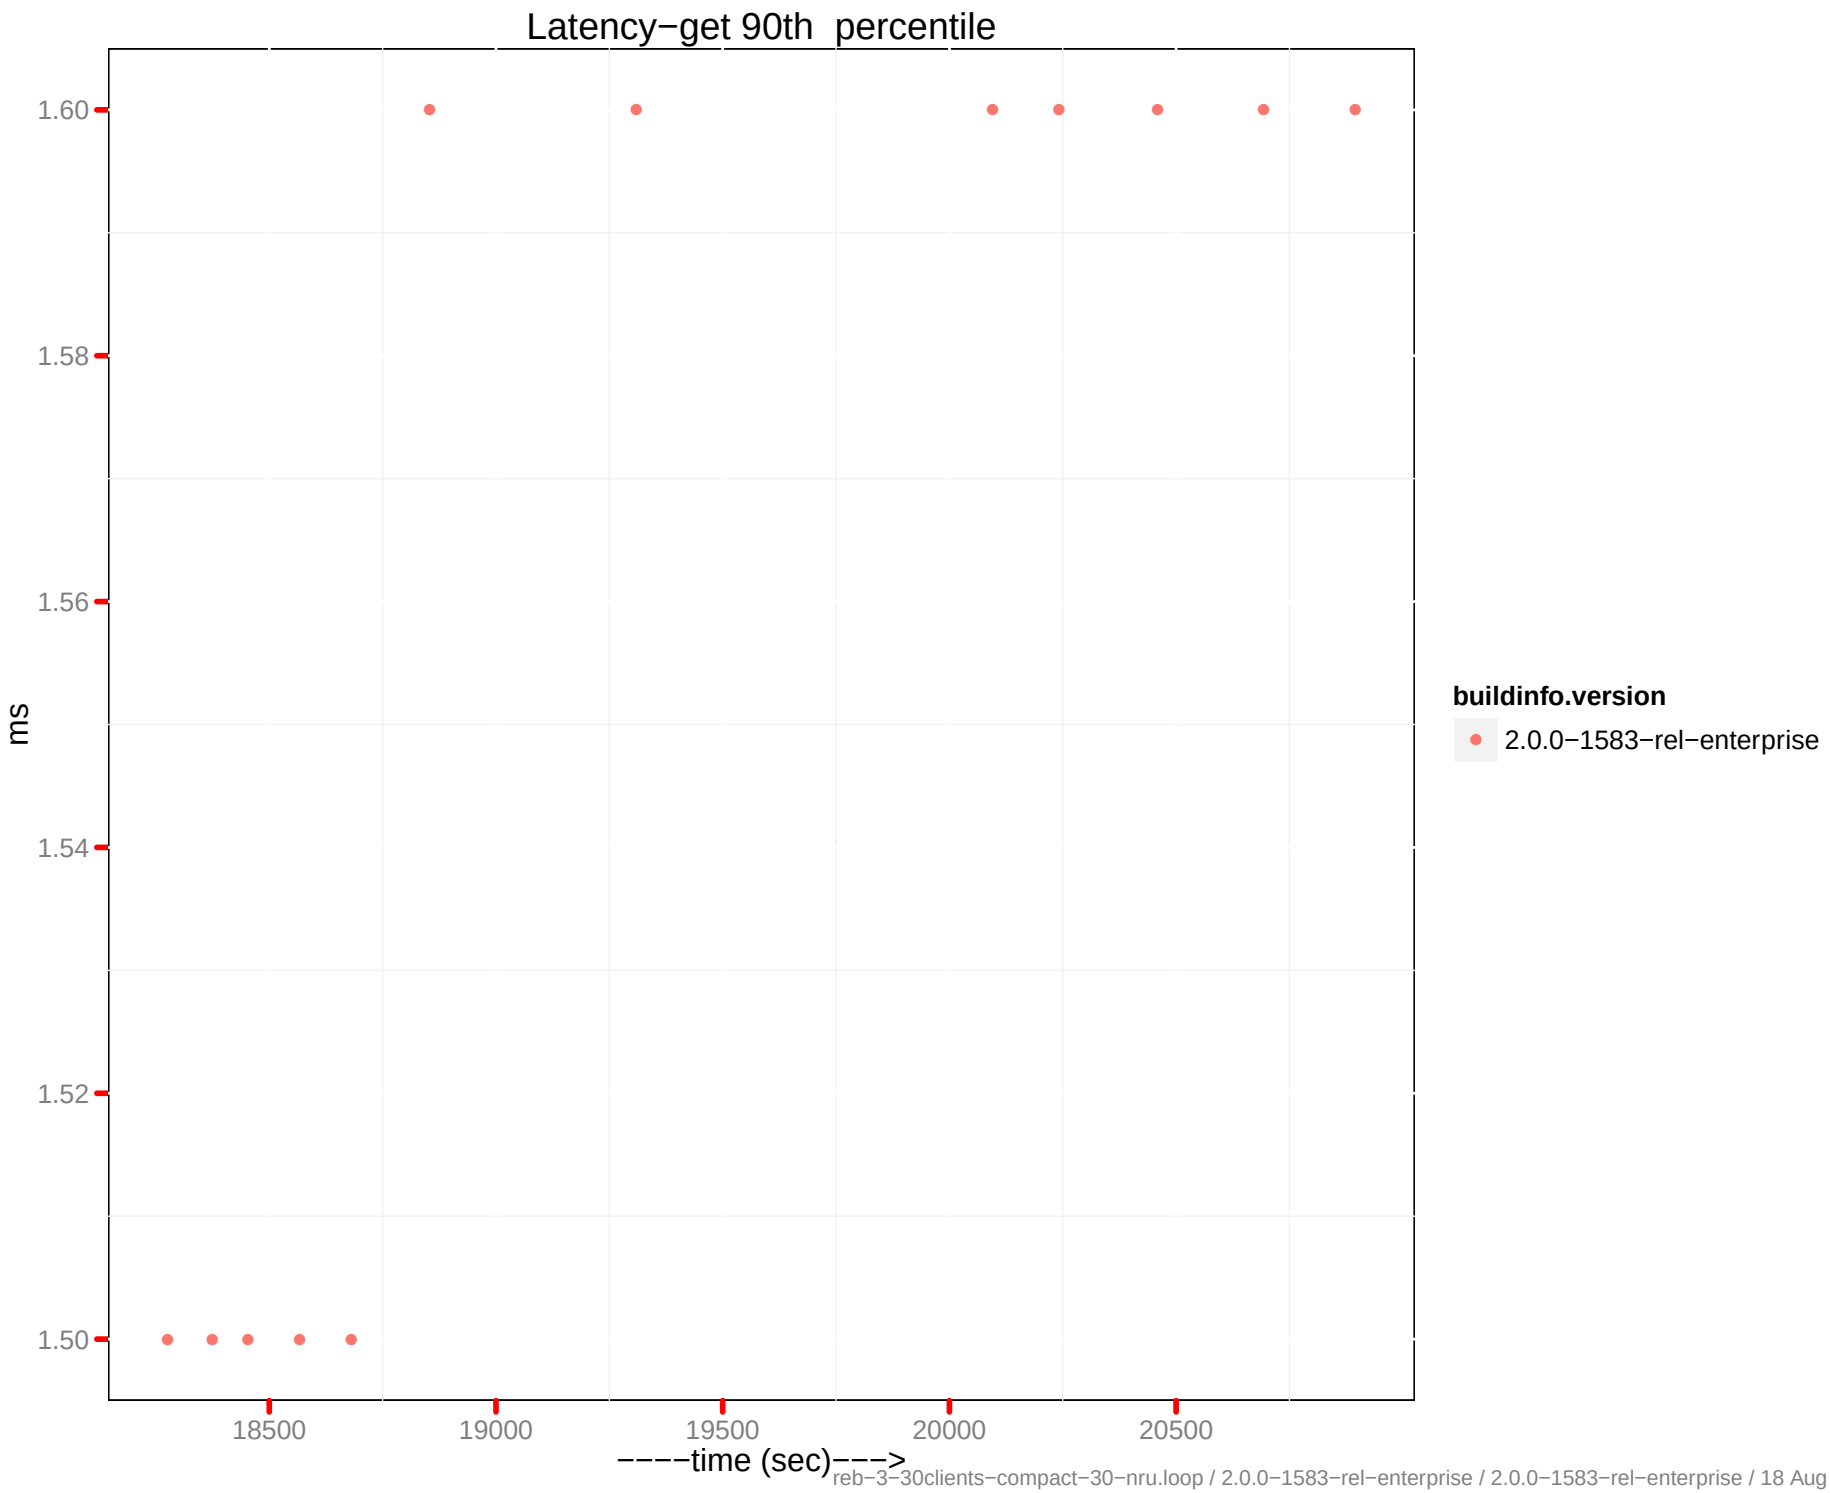
Latency−get 90th percentile
1.60
1.58
1.56
1.54
1.52
1.50
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

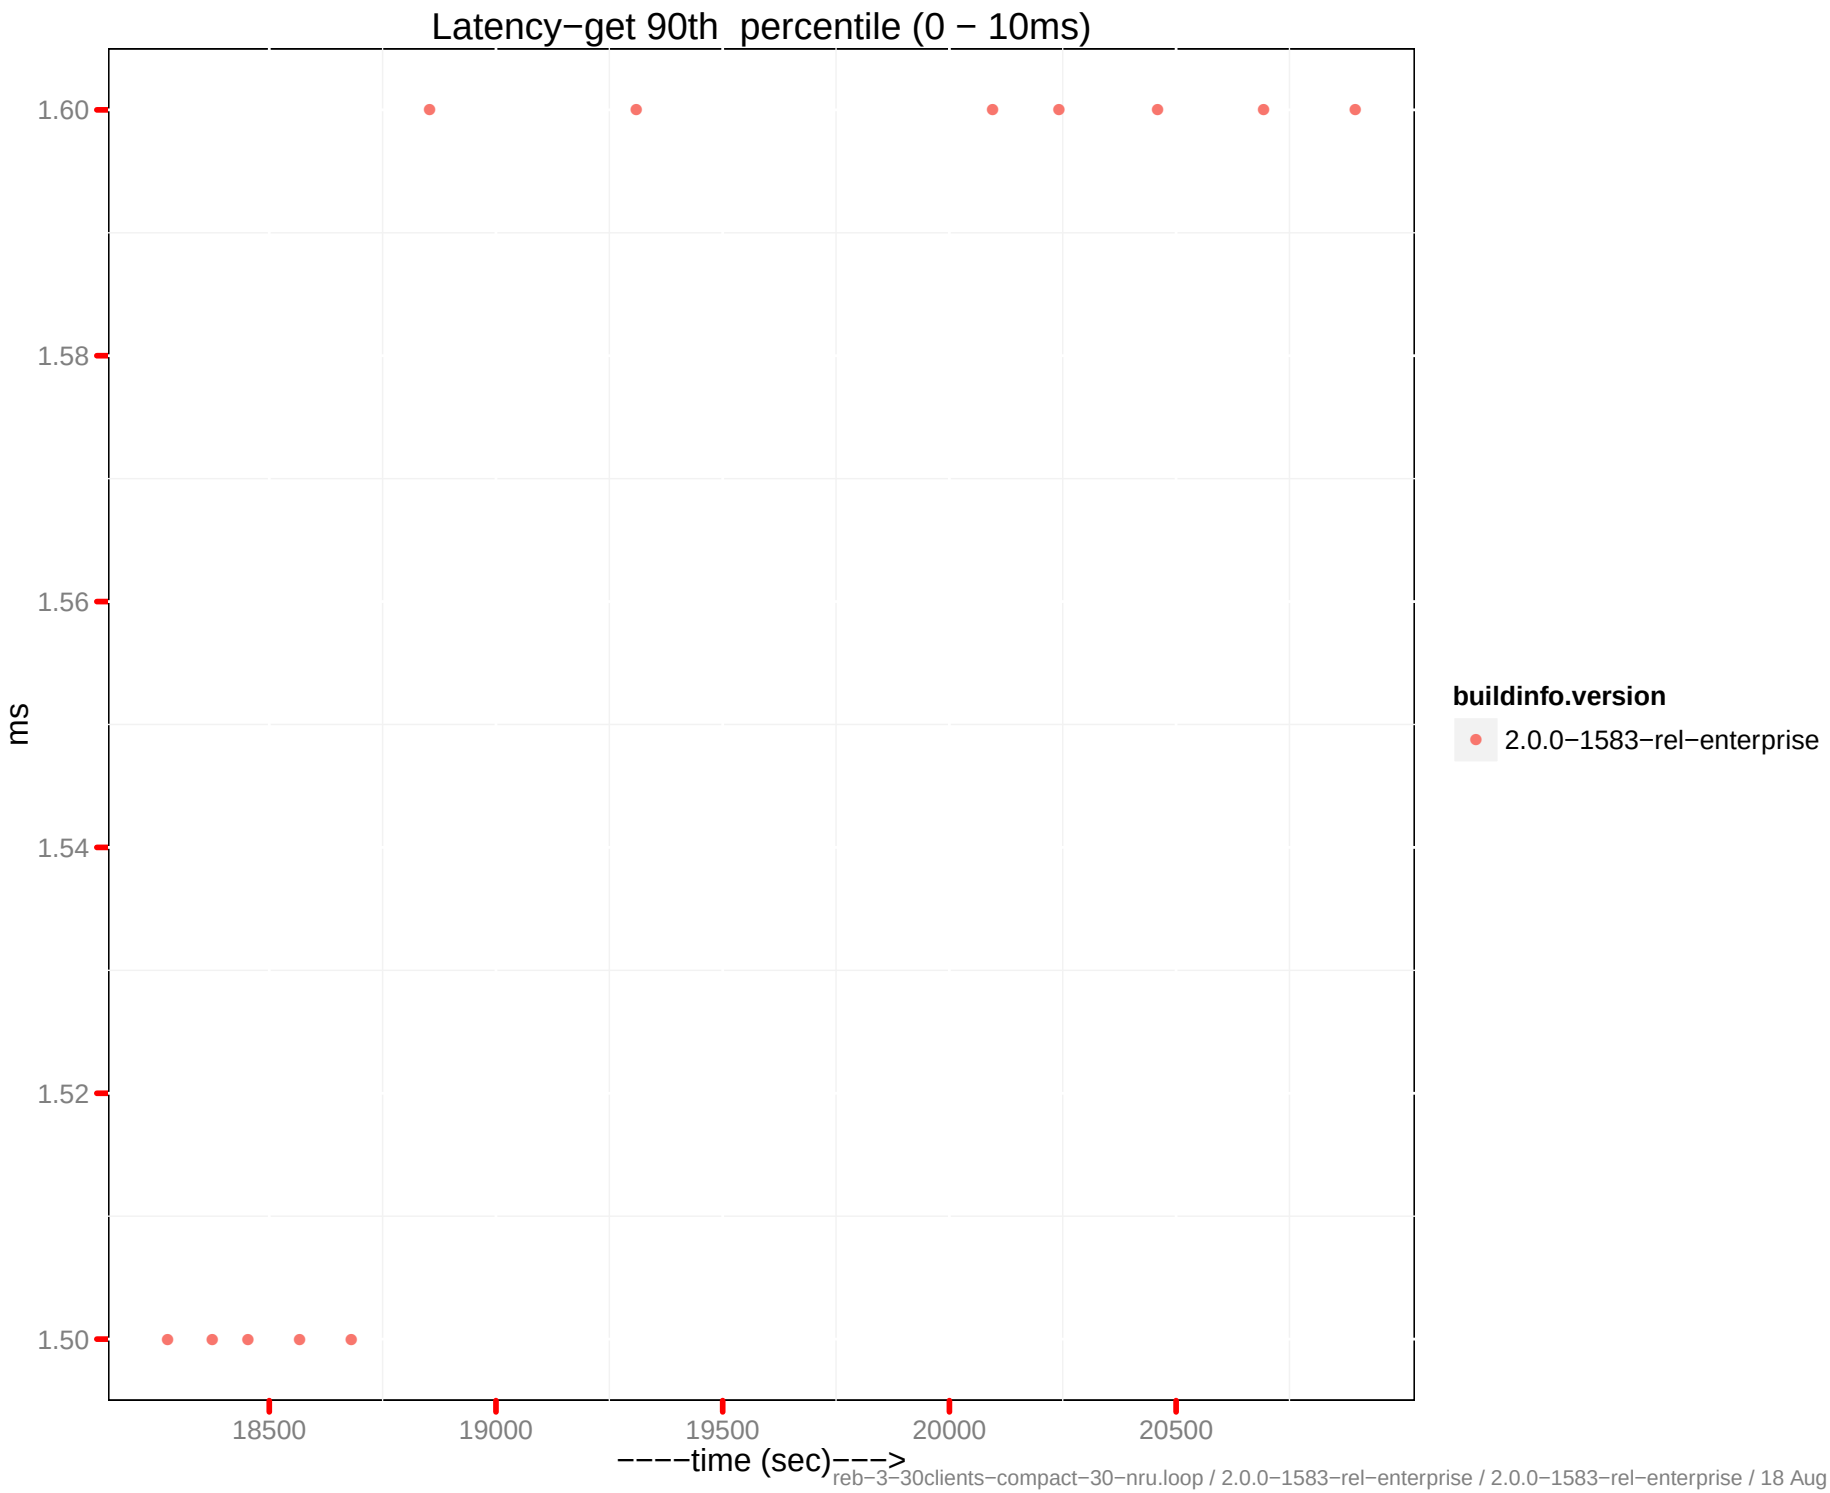
Latency−get 90th percentile (0 − 10ms)
1.60
1.58
1.56
1.54
1.52
1.50
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

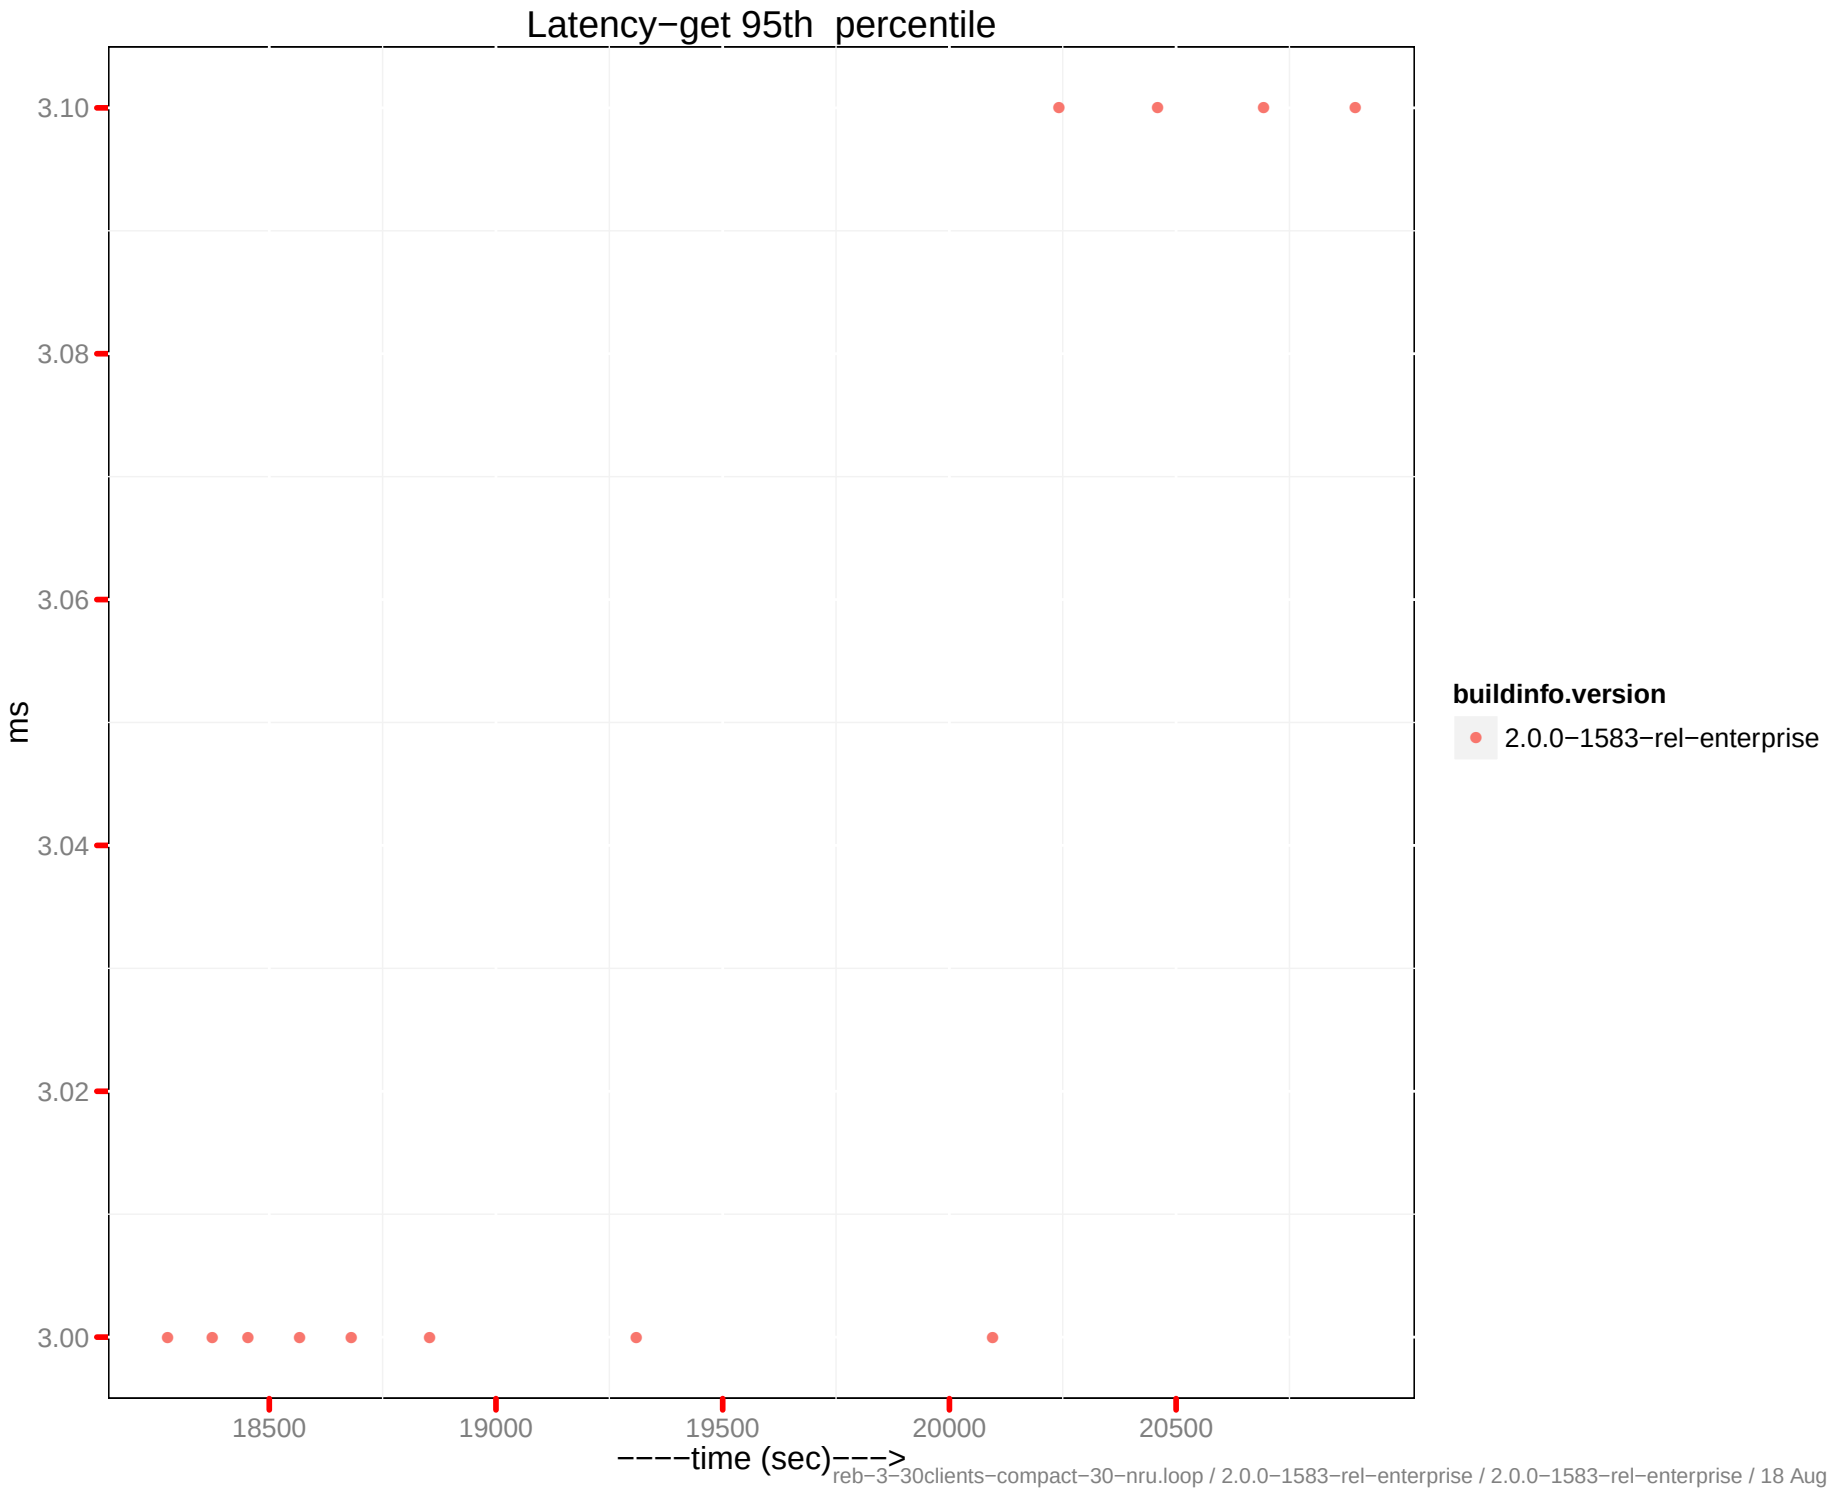
Latency−get 95th percentile
3.10
3.08
3.06
3.04
3.02
3.00
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

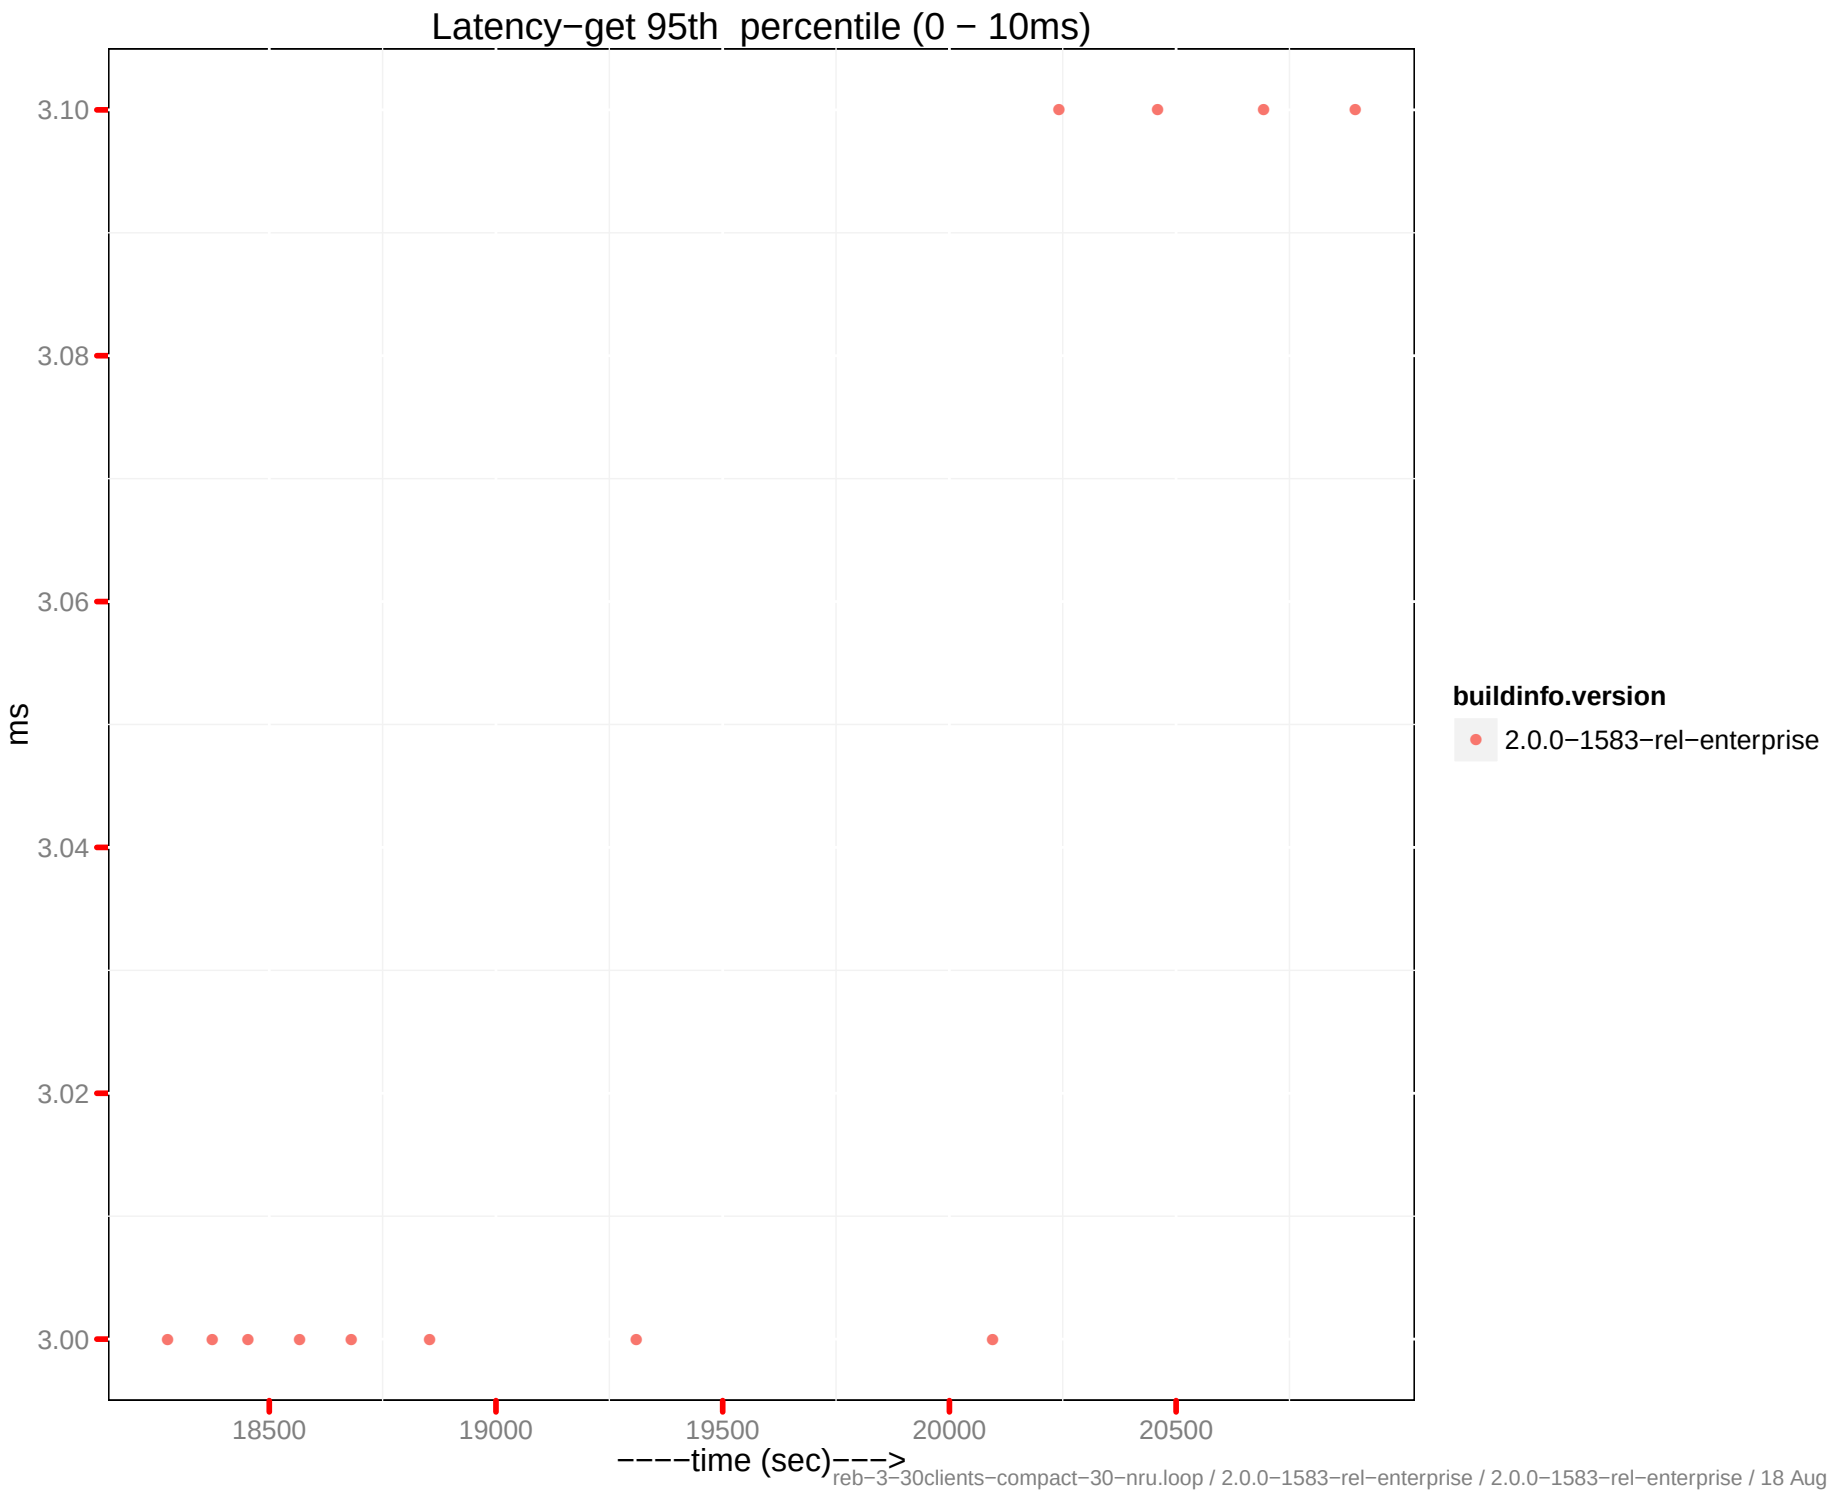
Latency−get 95th percentile (0 − 10ms)
3.10
3.08
3.06
3.04
3.02
3.00
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

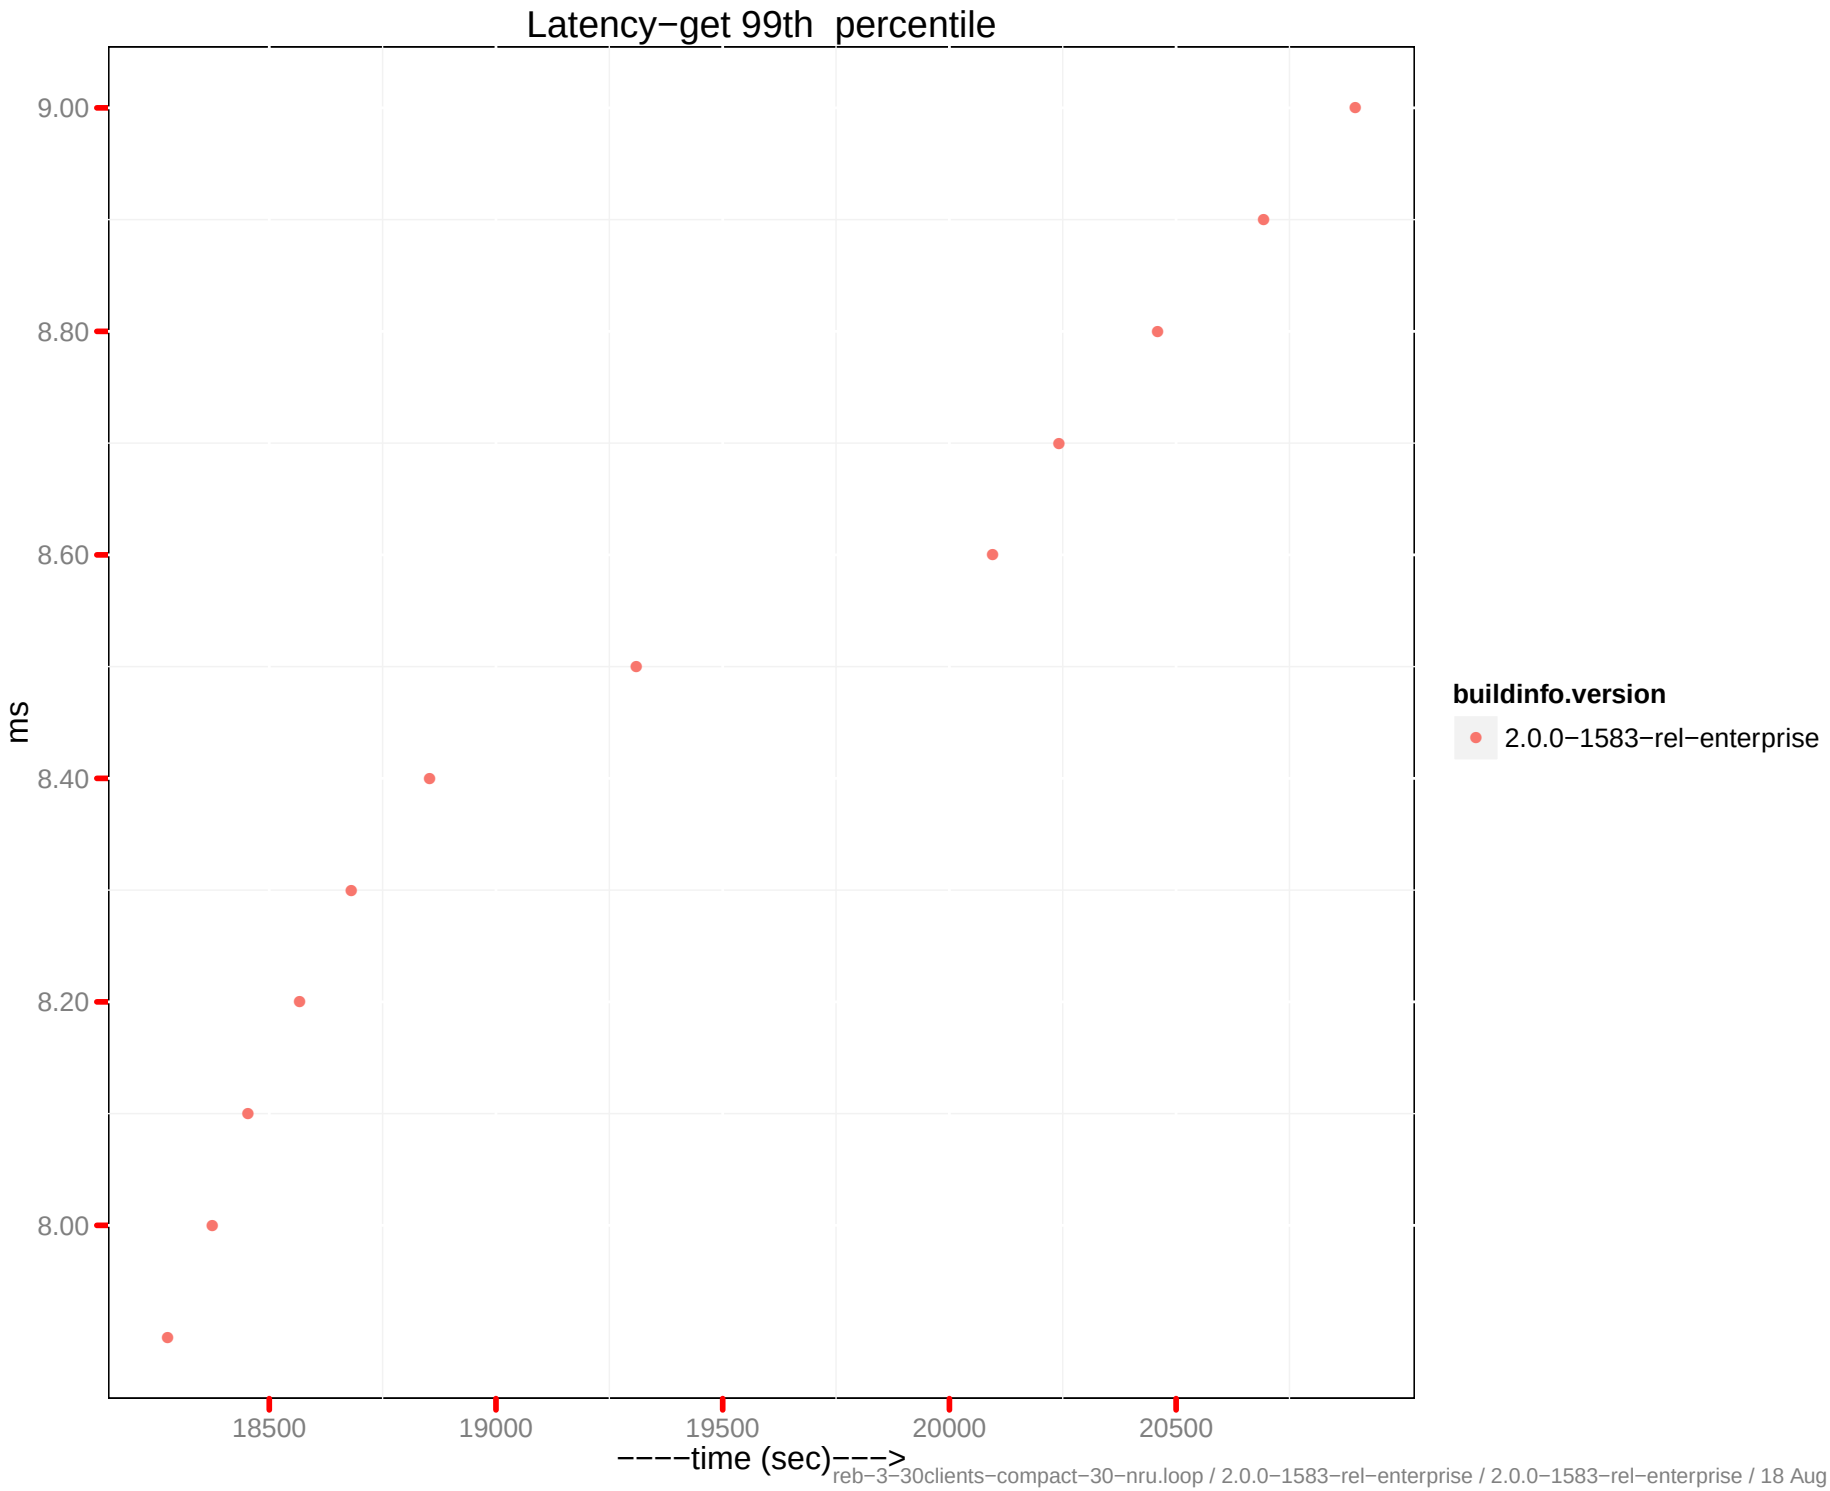
Latency−get 99th percentile
9.00
8.80
8.60
8.40
8.20
8.00
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

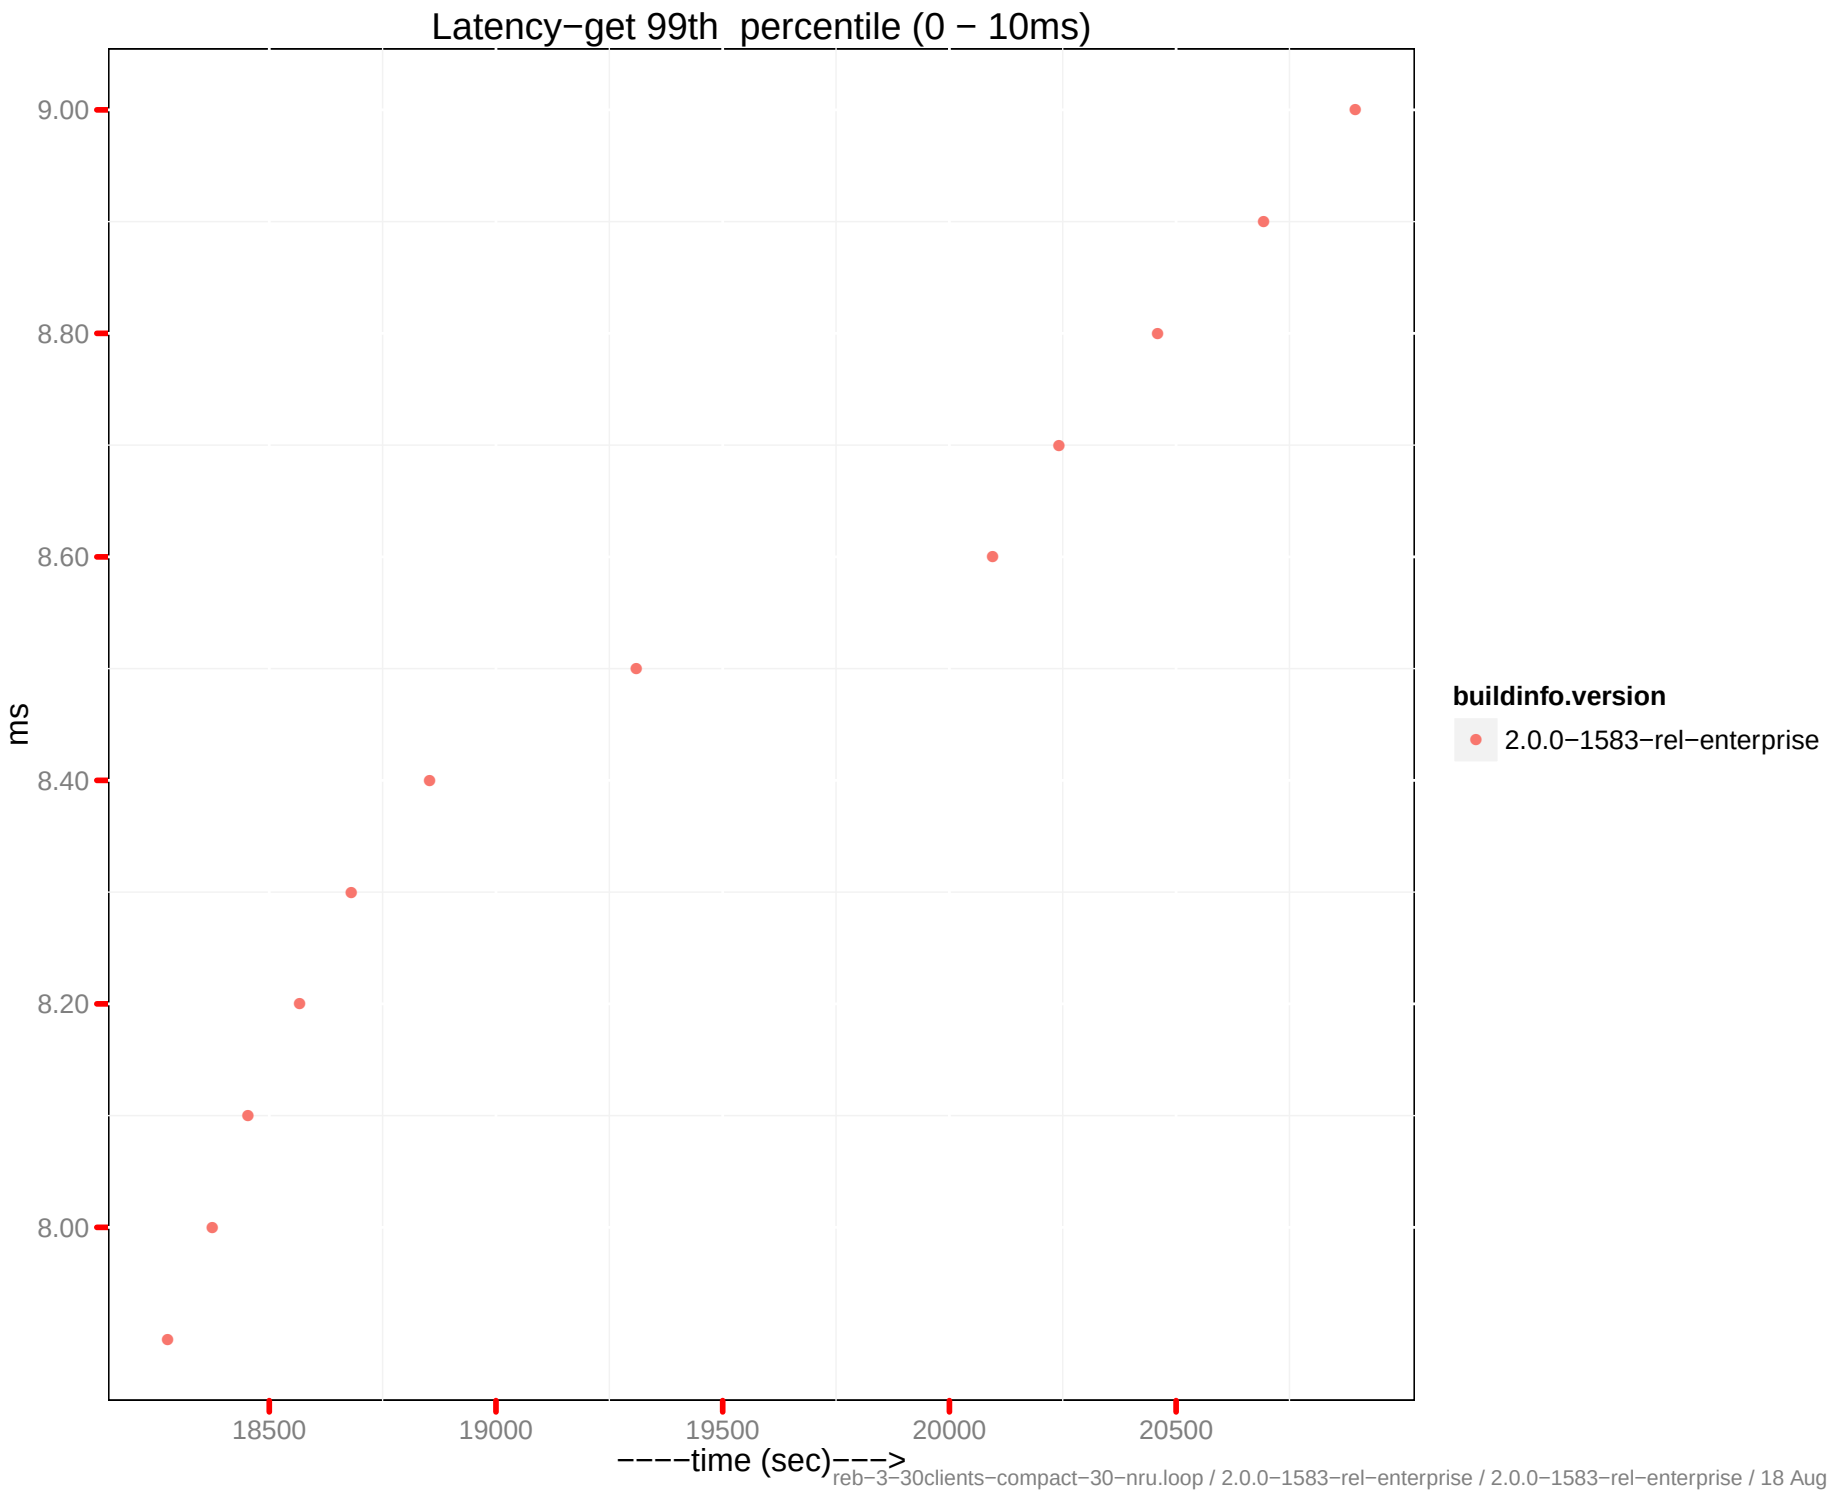
Latency−get 99th percentile (0 − 10ms)
9.00
8.80
8.60
8.40
8.20
8.00
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

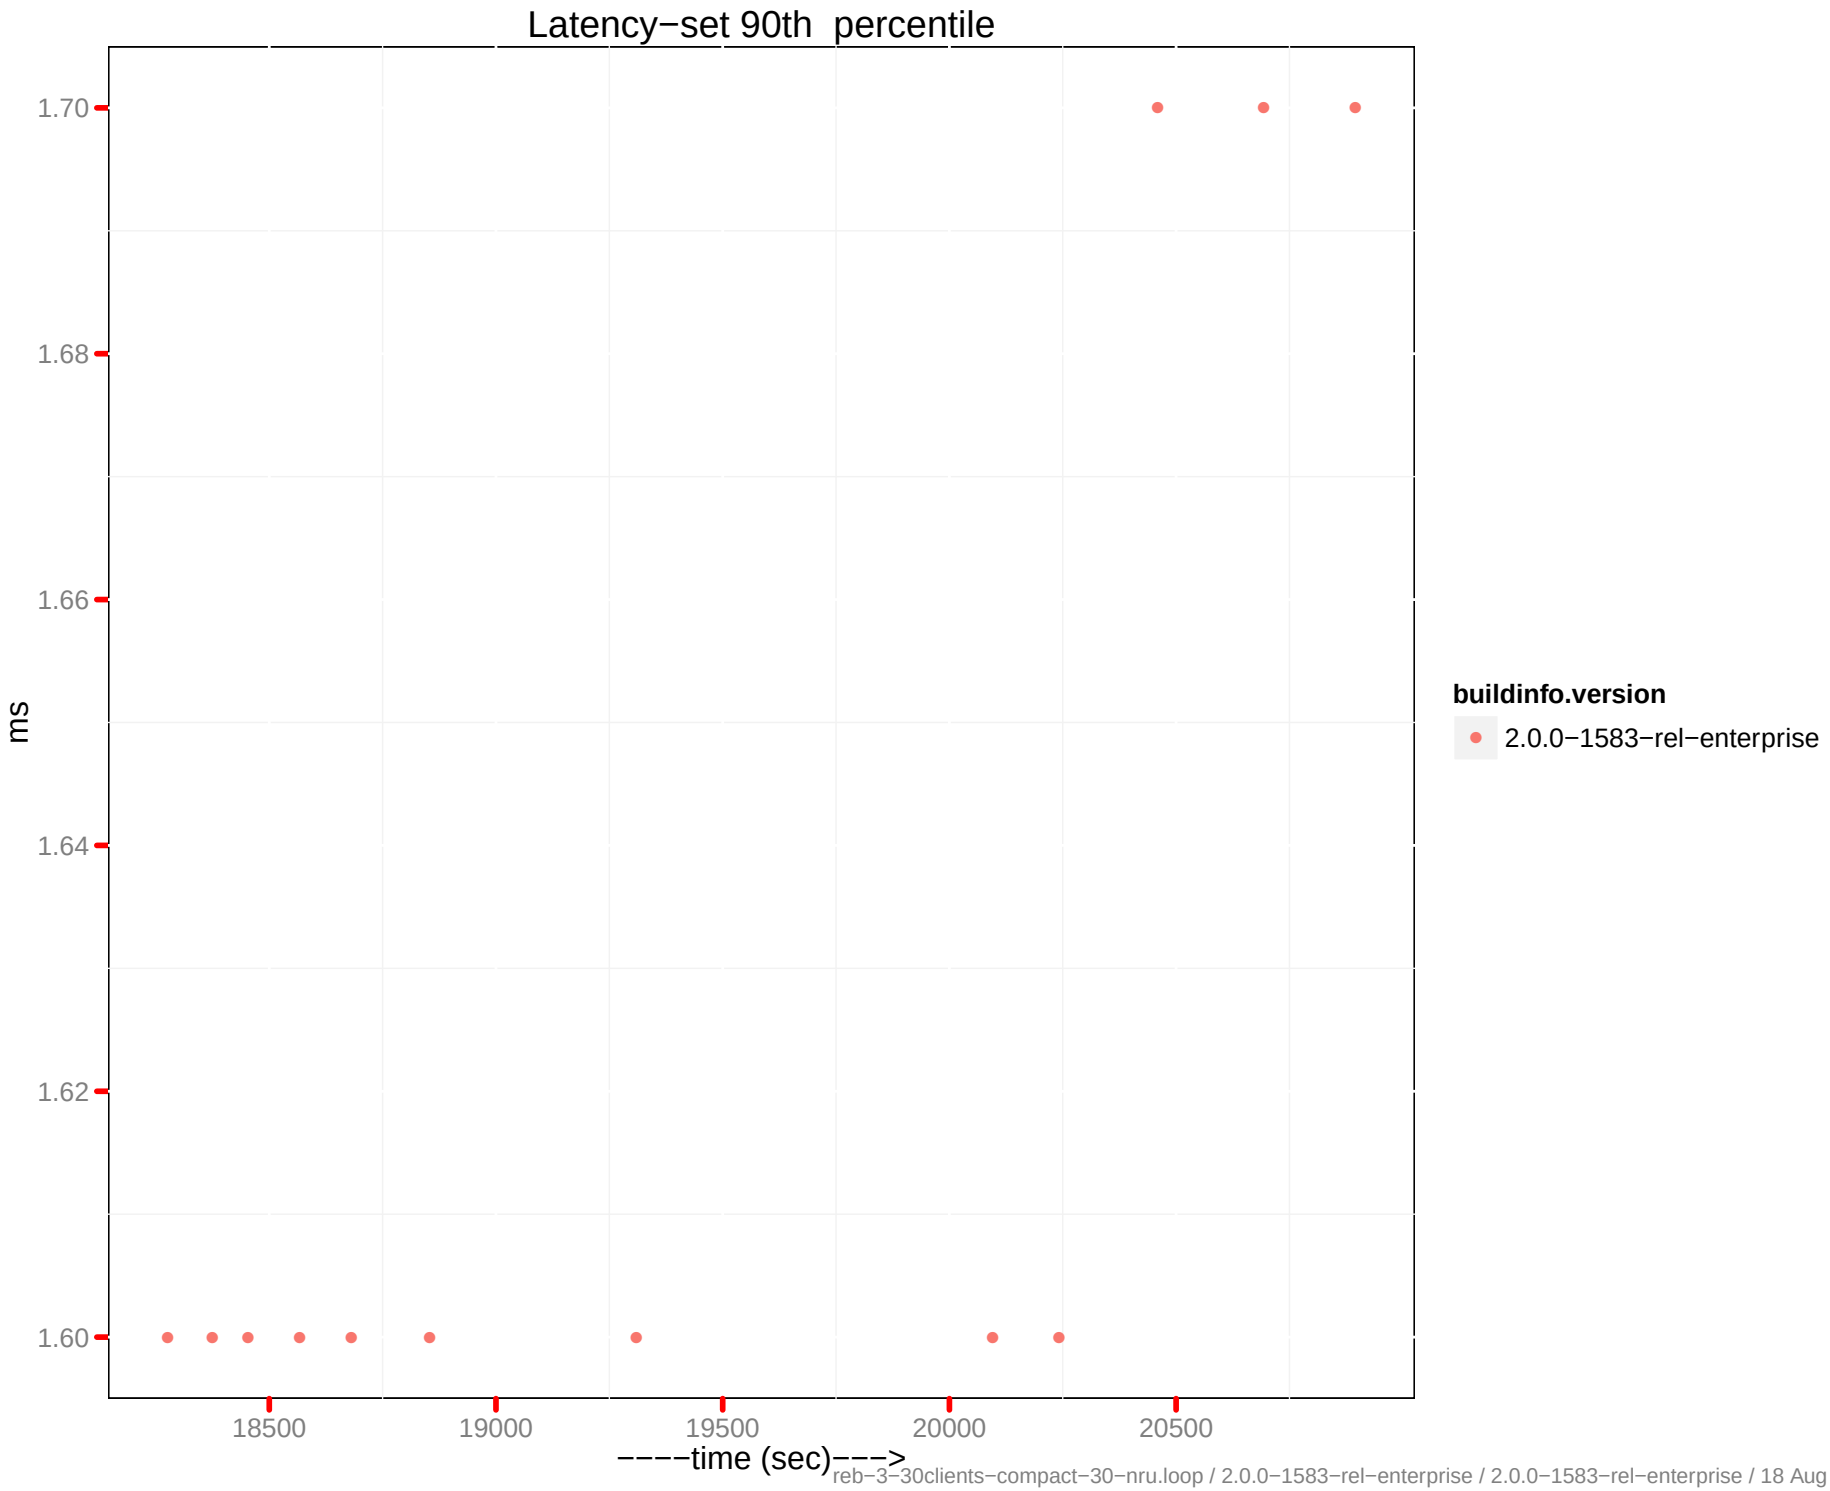
Latency−set 90th percentile
1.70
1.68
1.66
1.64
1.62
1.60
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

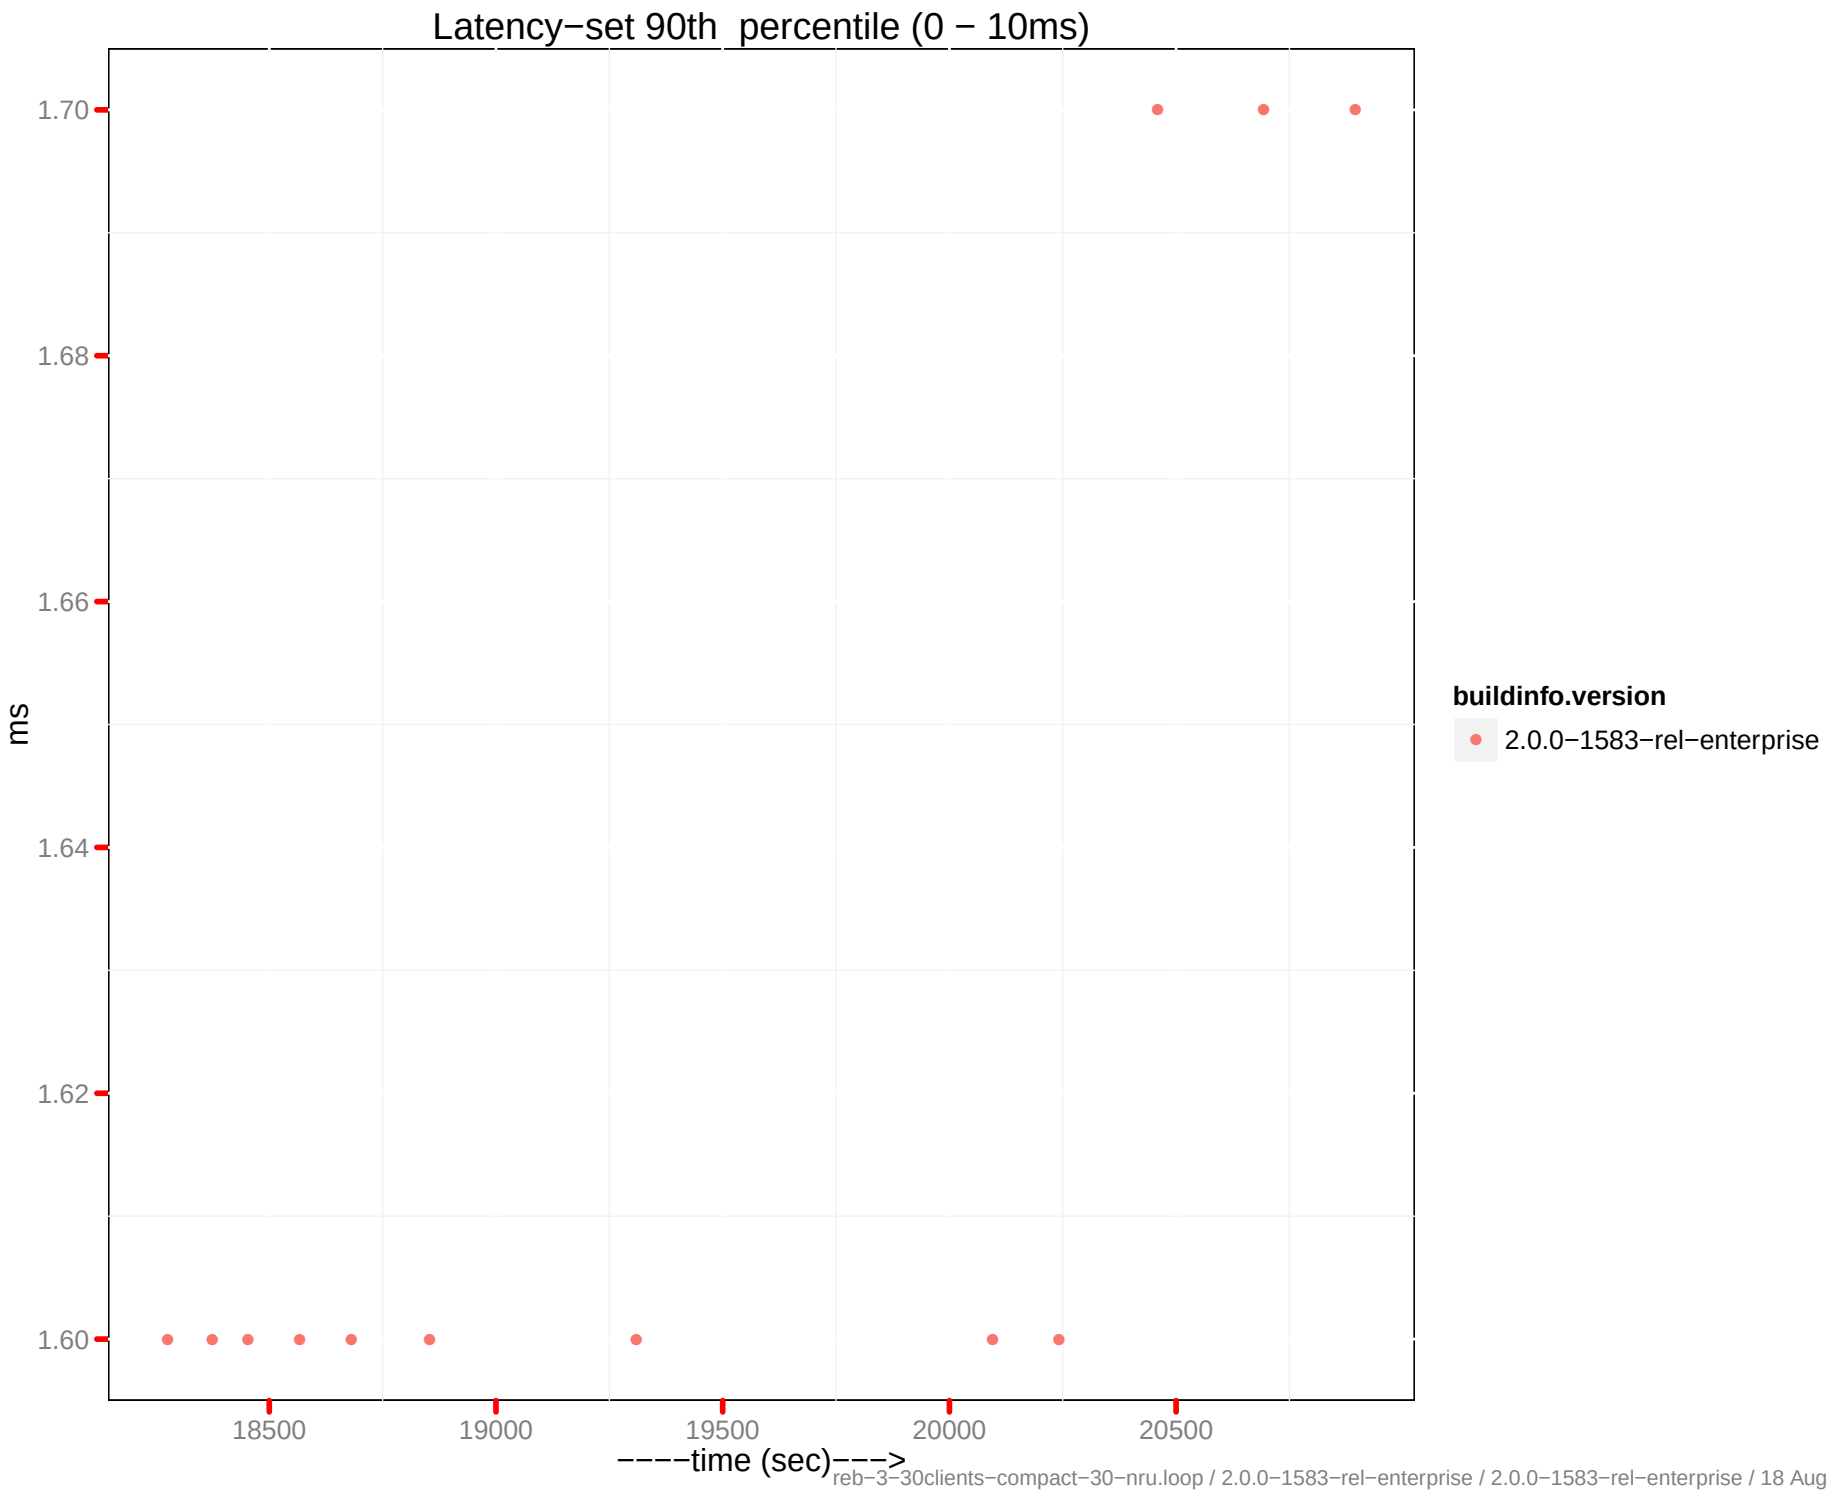
Latency−set 90th percentile (0 − 10ms)
1.70
1.68
1.66
1.64
1.62
1.60
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

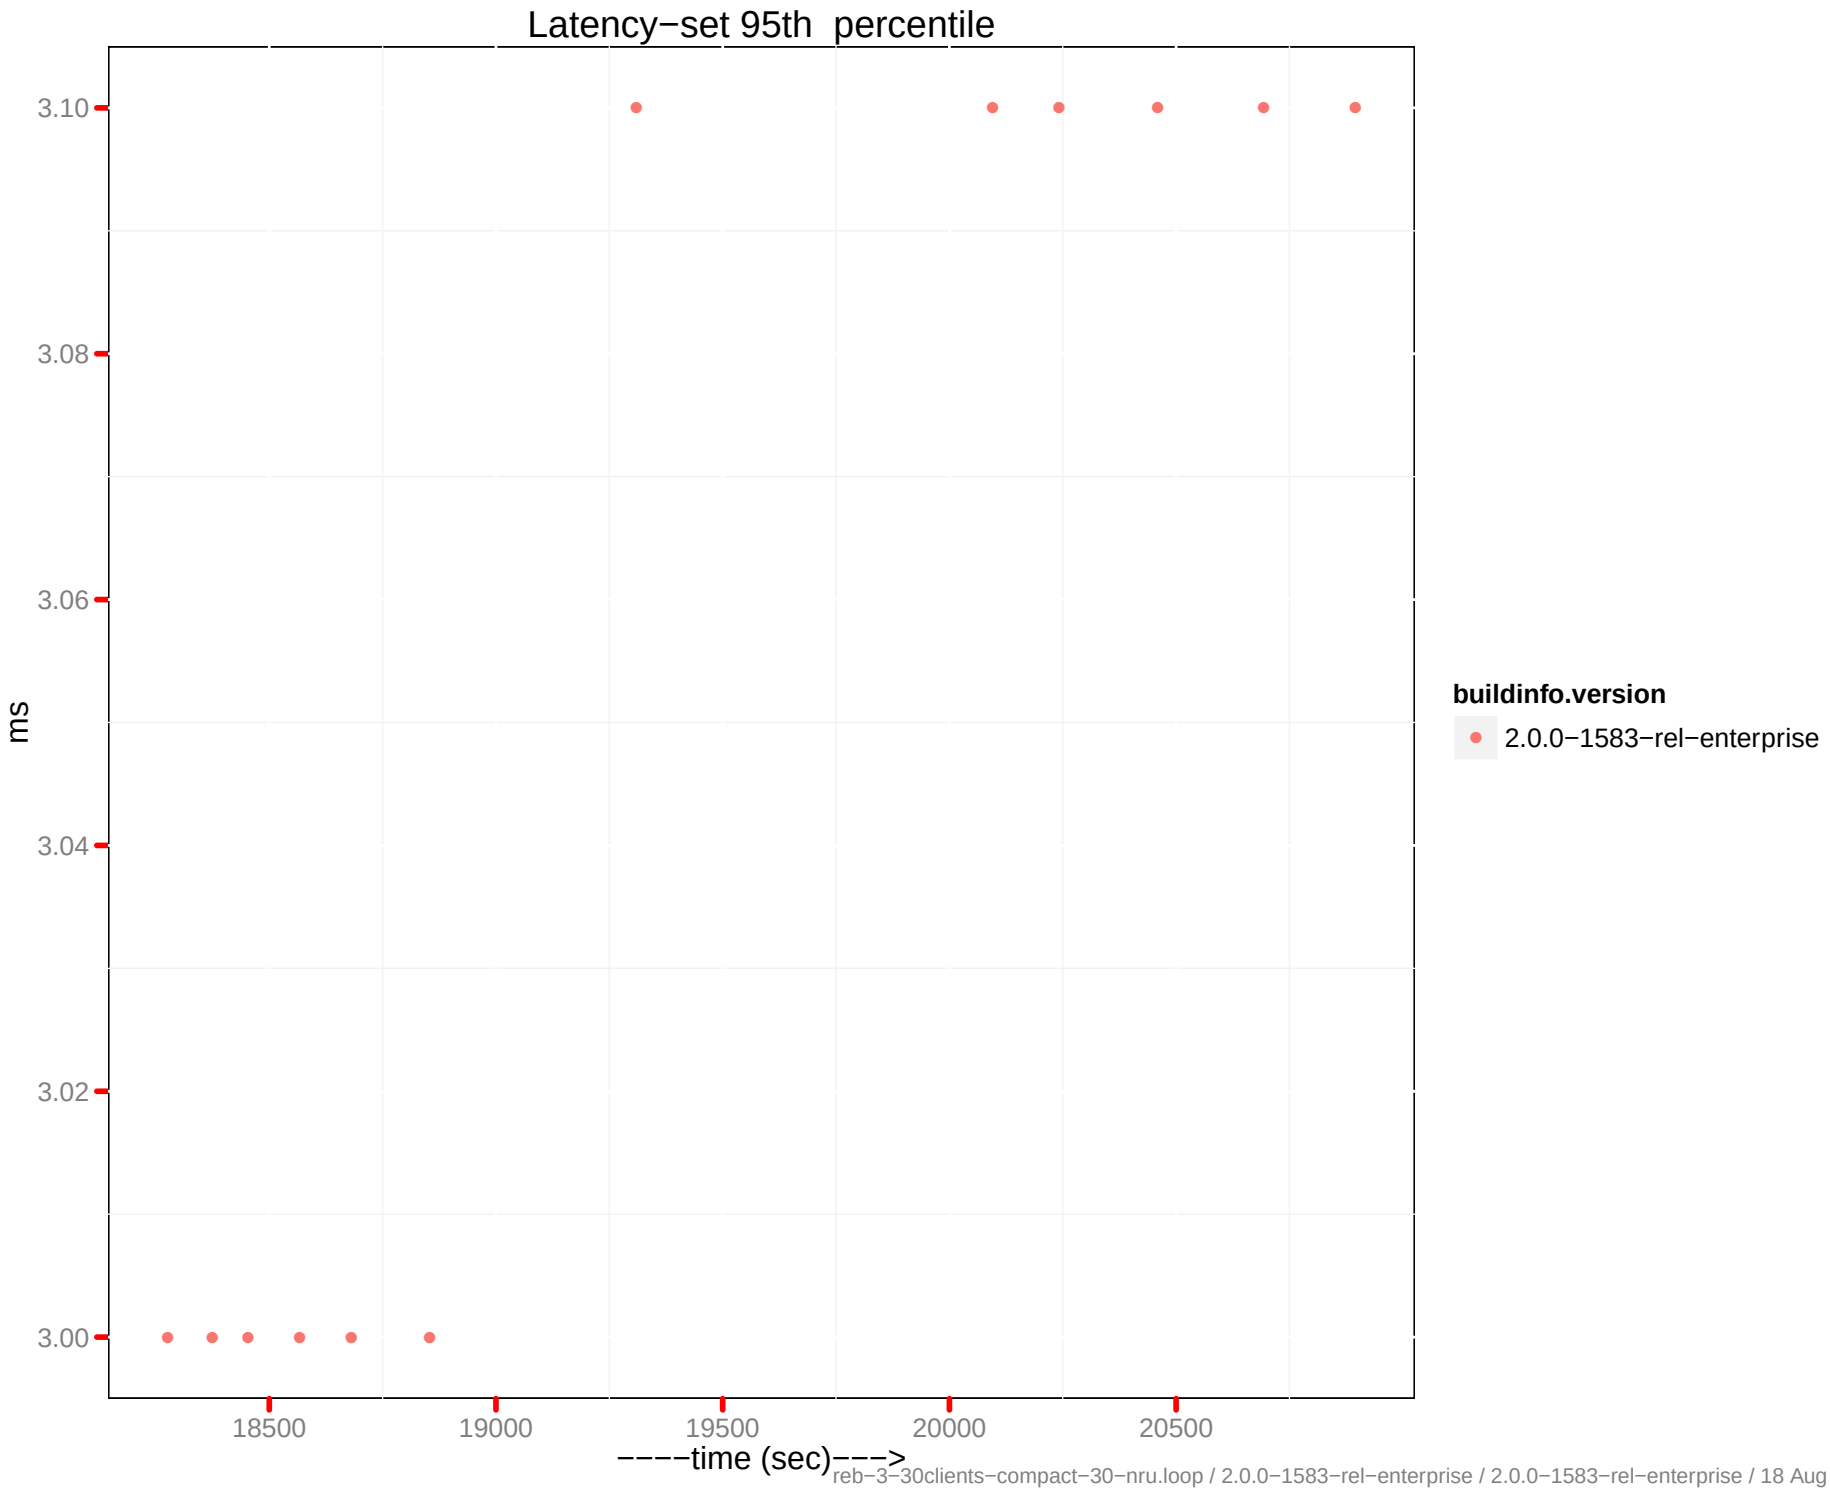
Latency−set 95th percentile
3.10
3.08
3.06
3.04
3.02
3.00
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

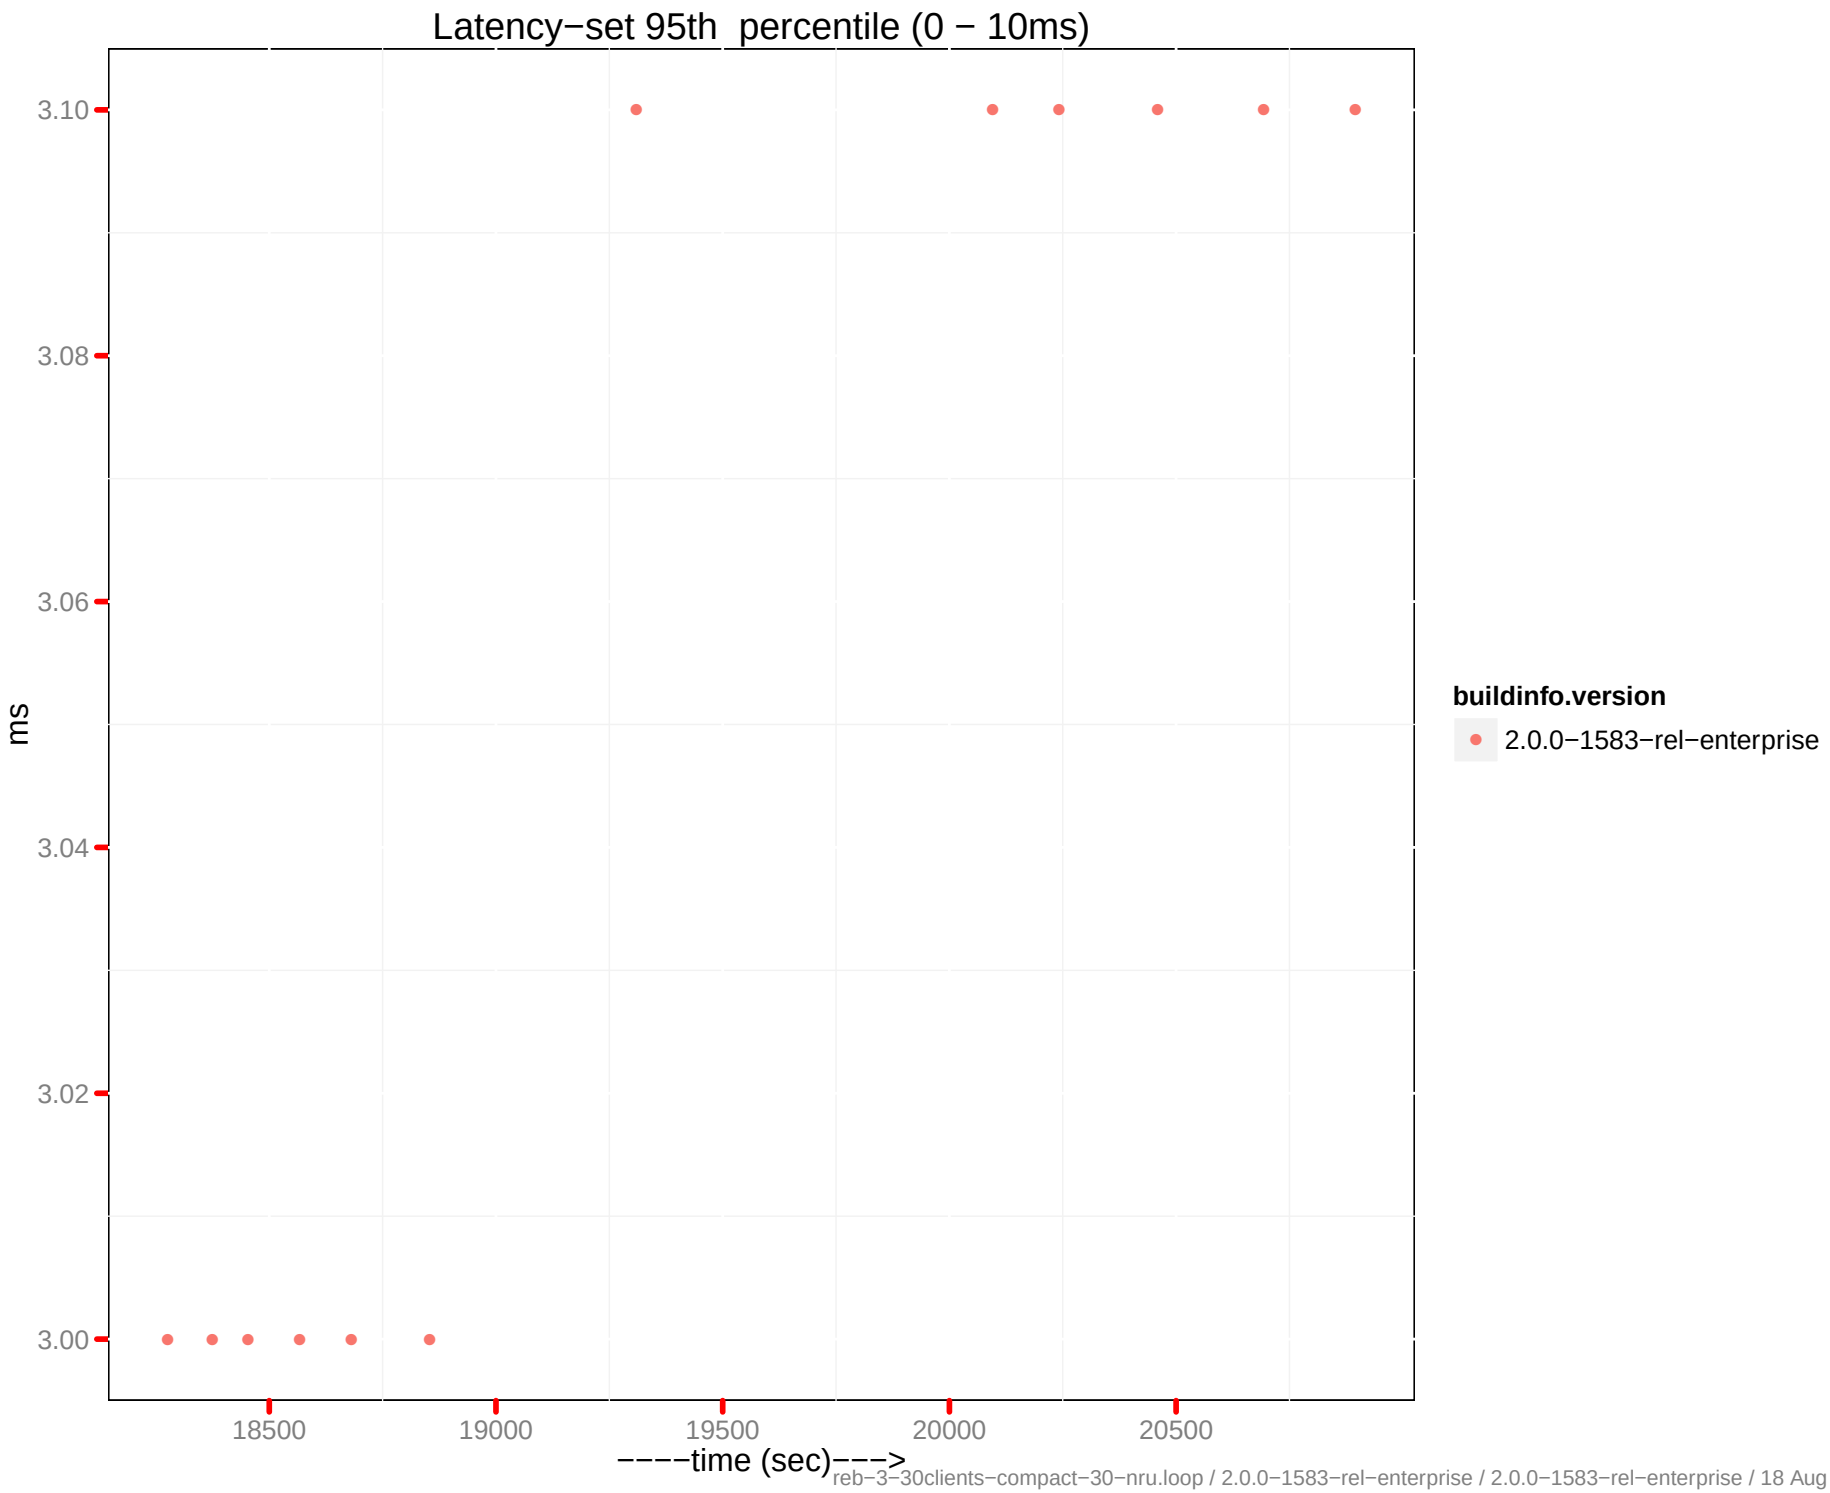
Latency−set 95th percentile (0 − 10ms)
3.10
3.08
3.06
3.04
3.02
3.00
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

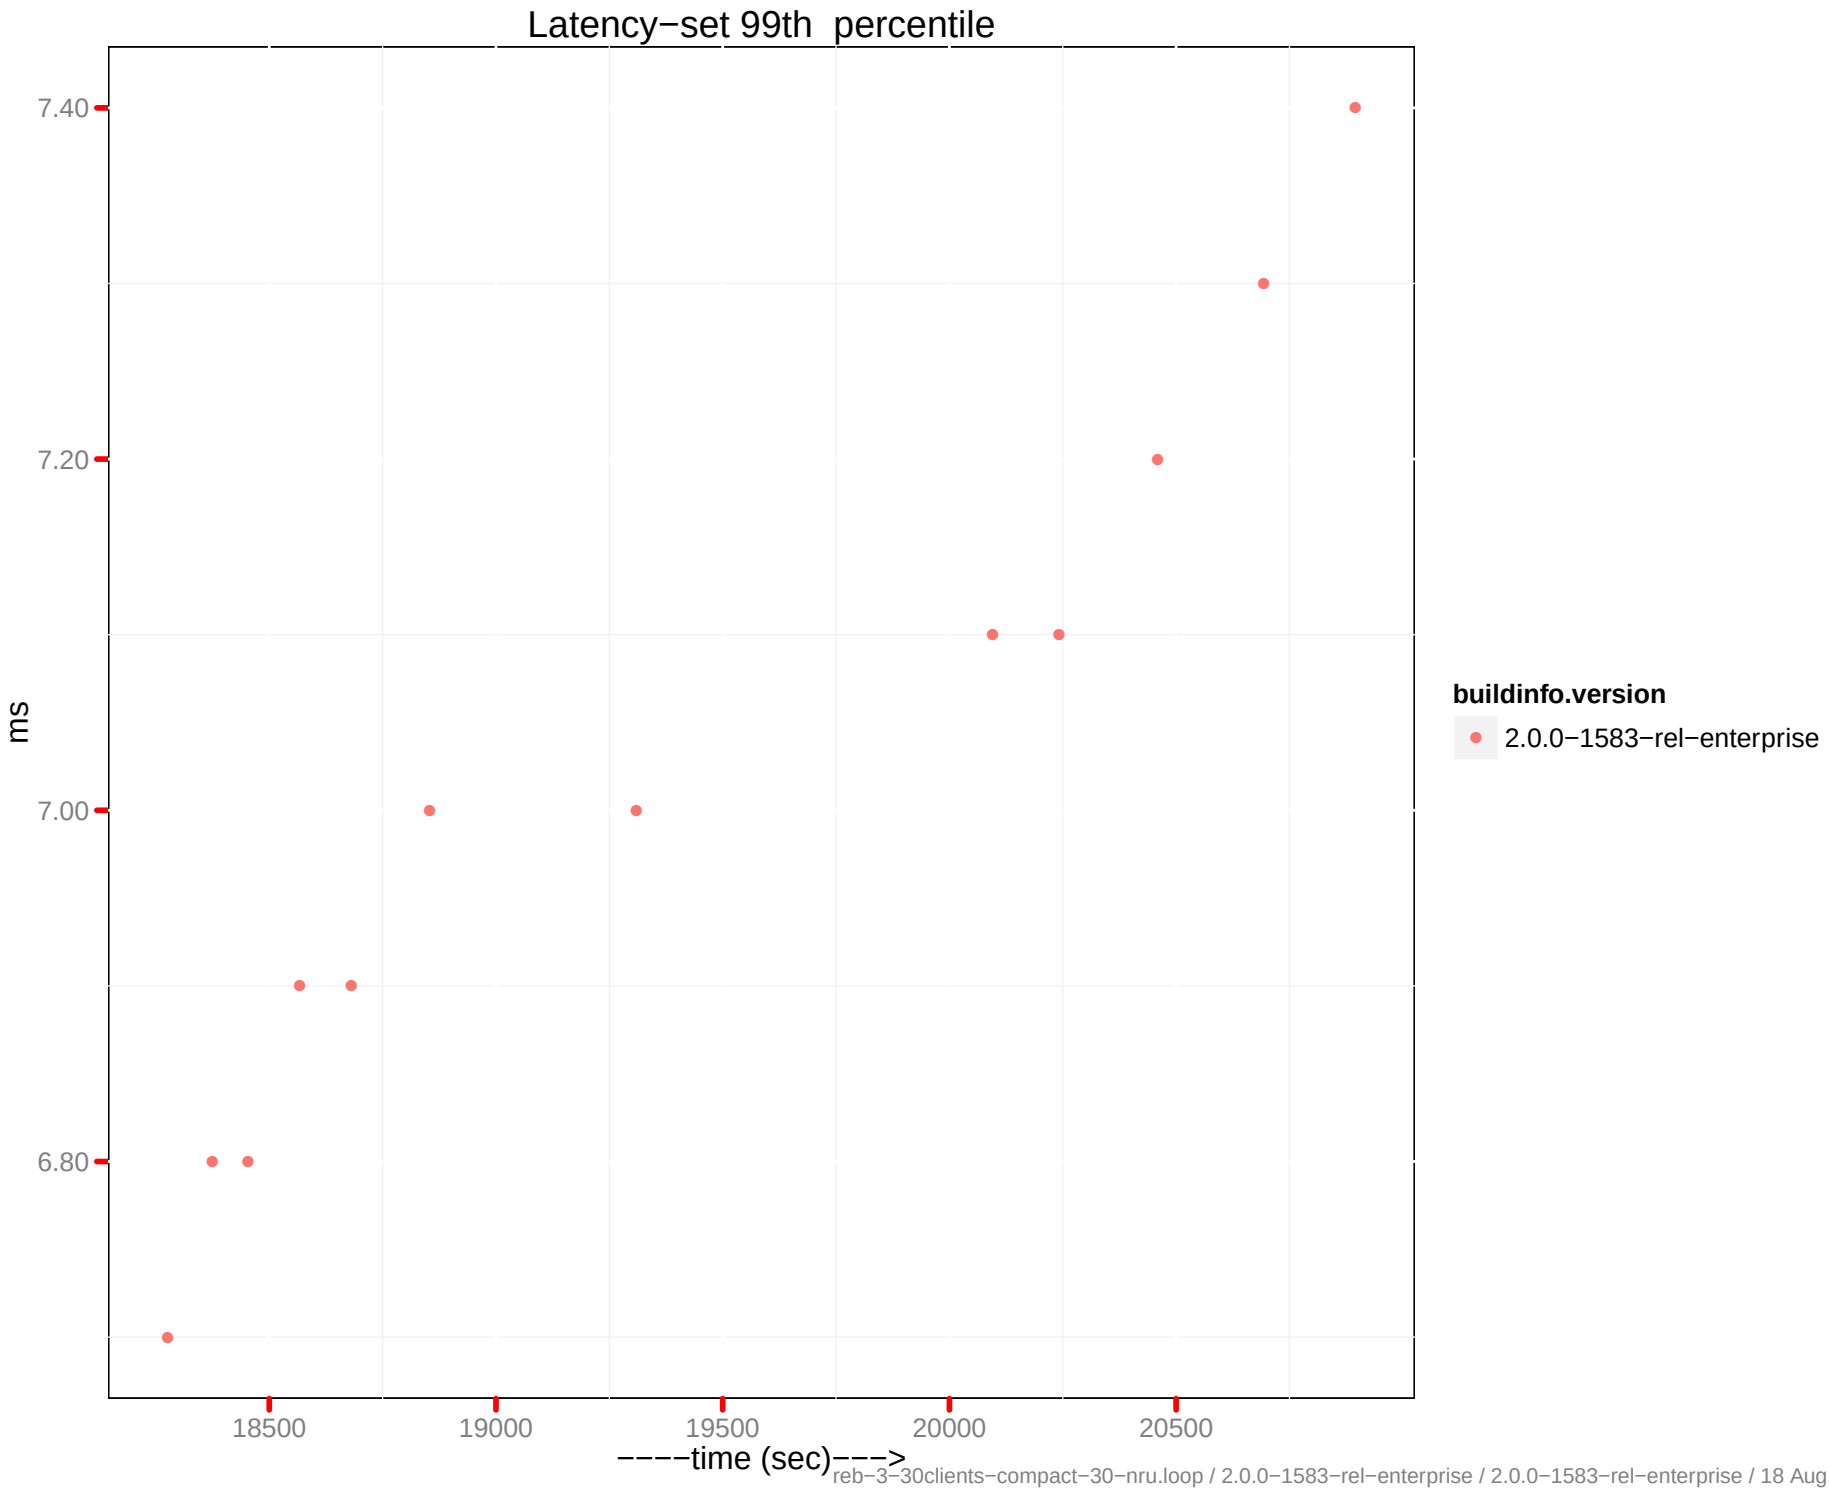
Latency−set 99th percentile
7.40
7.20
7.00
6.80
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

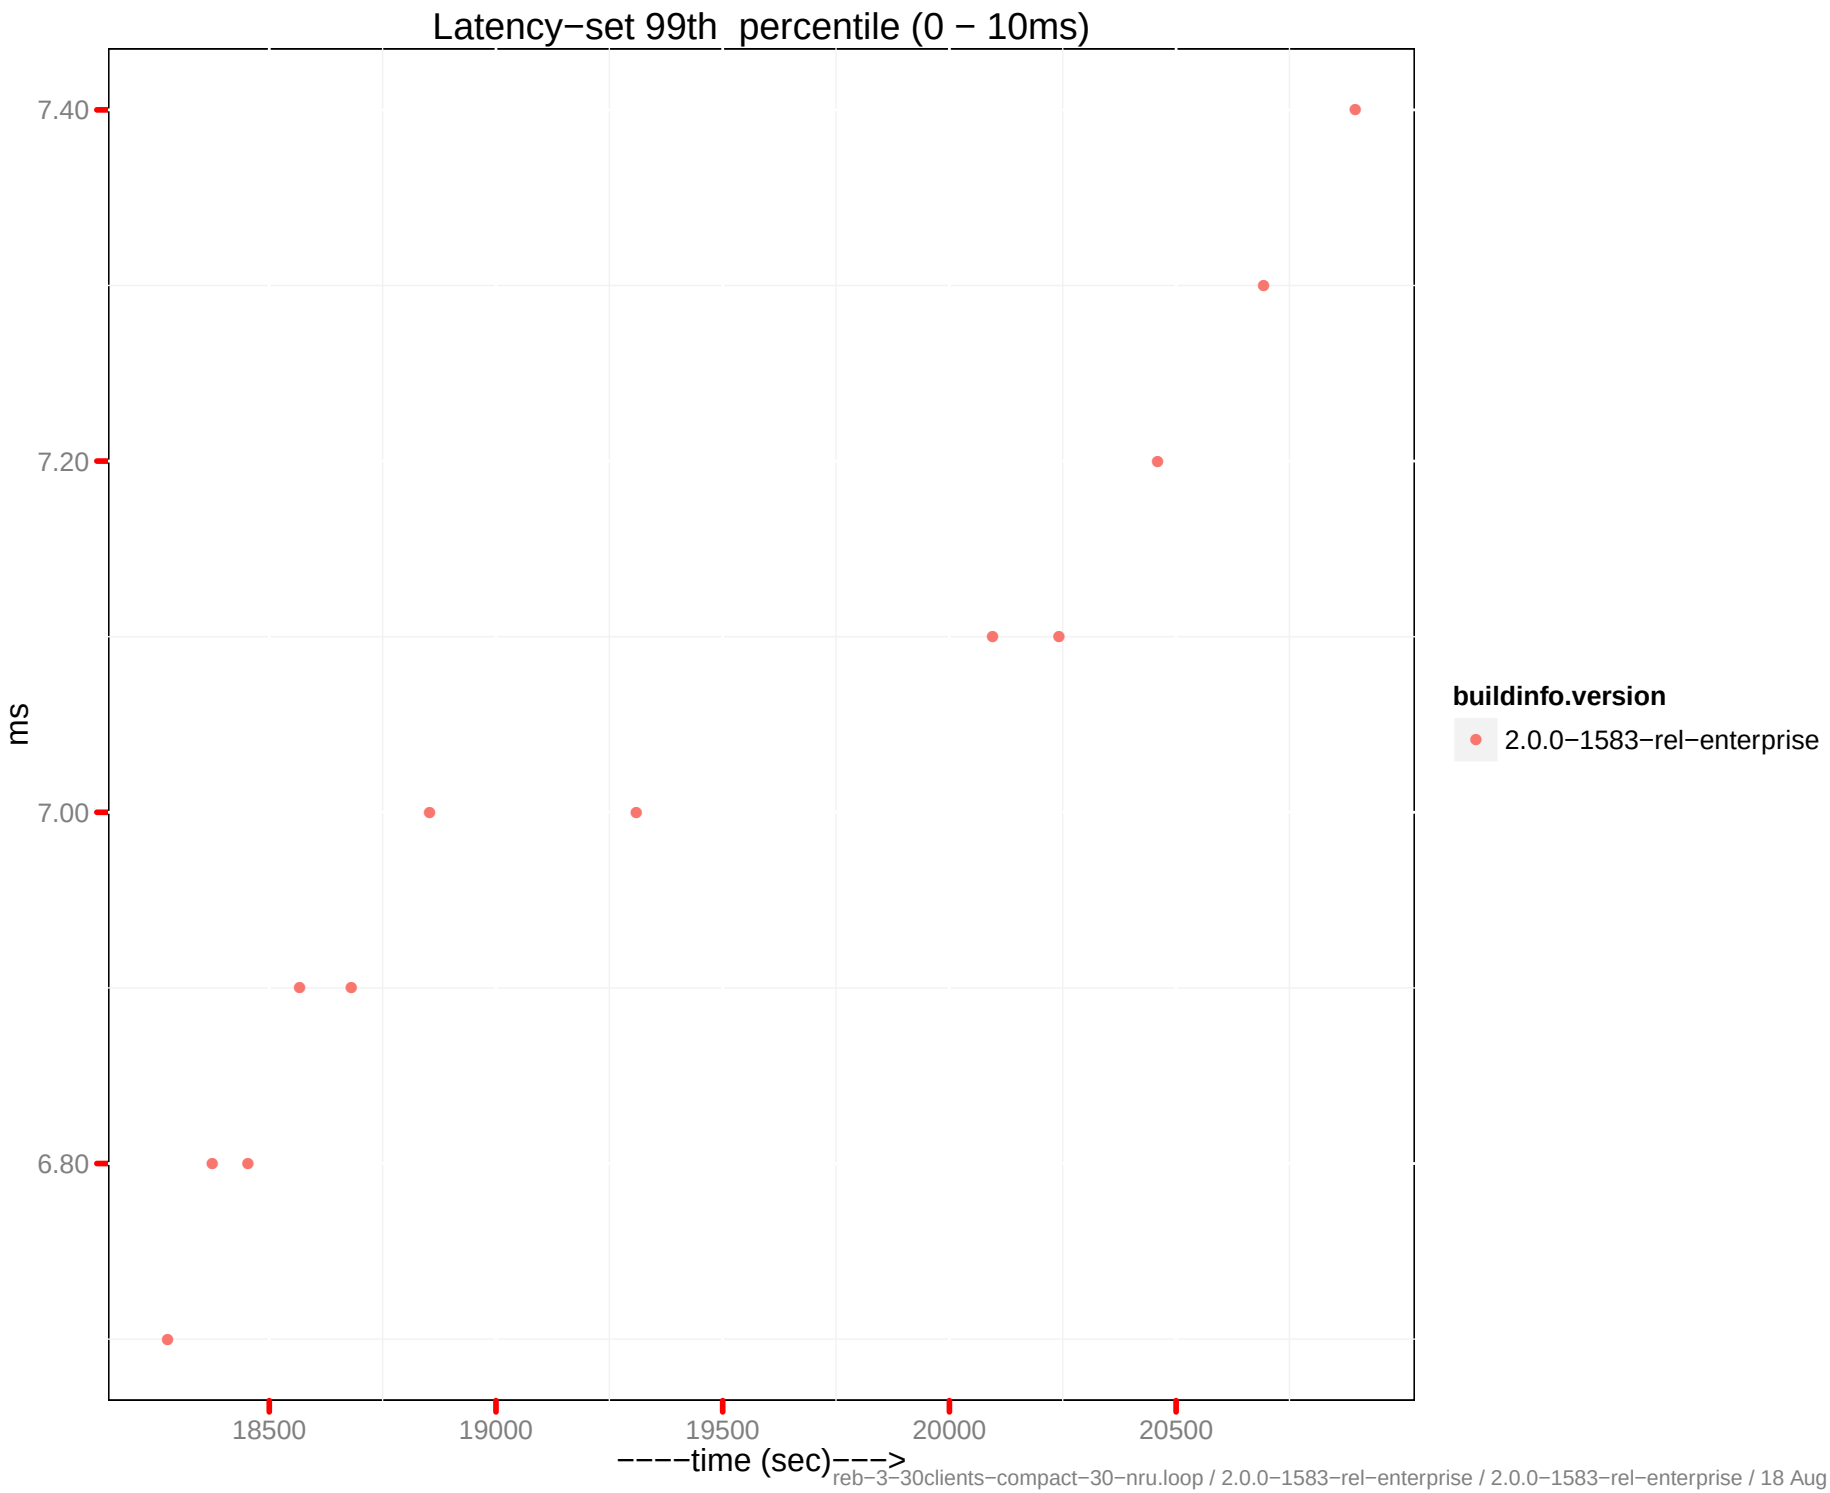
Latency−set 99th percentile (0 − 10ms)
7.40
7.20
7.00
6.80
ms
18500
19000
19500
20000
20500
−−−−time (sec)−−−>
buildinfo.version
2.0.0−1583−rel−enterprise
reb−3−30clients−compact−30−nru.loop / 2.0.0−1583−rel−enterprise / 2.0.0−1583−rel−enterprise / 18 Aug 2012

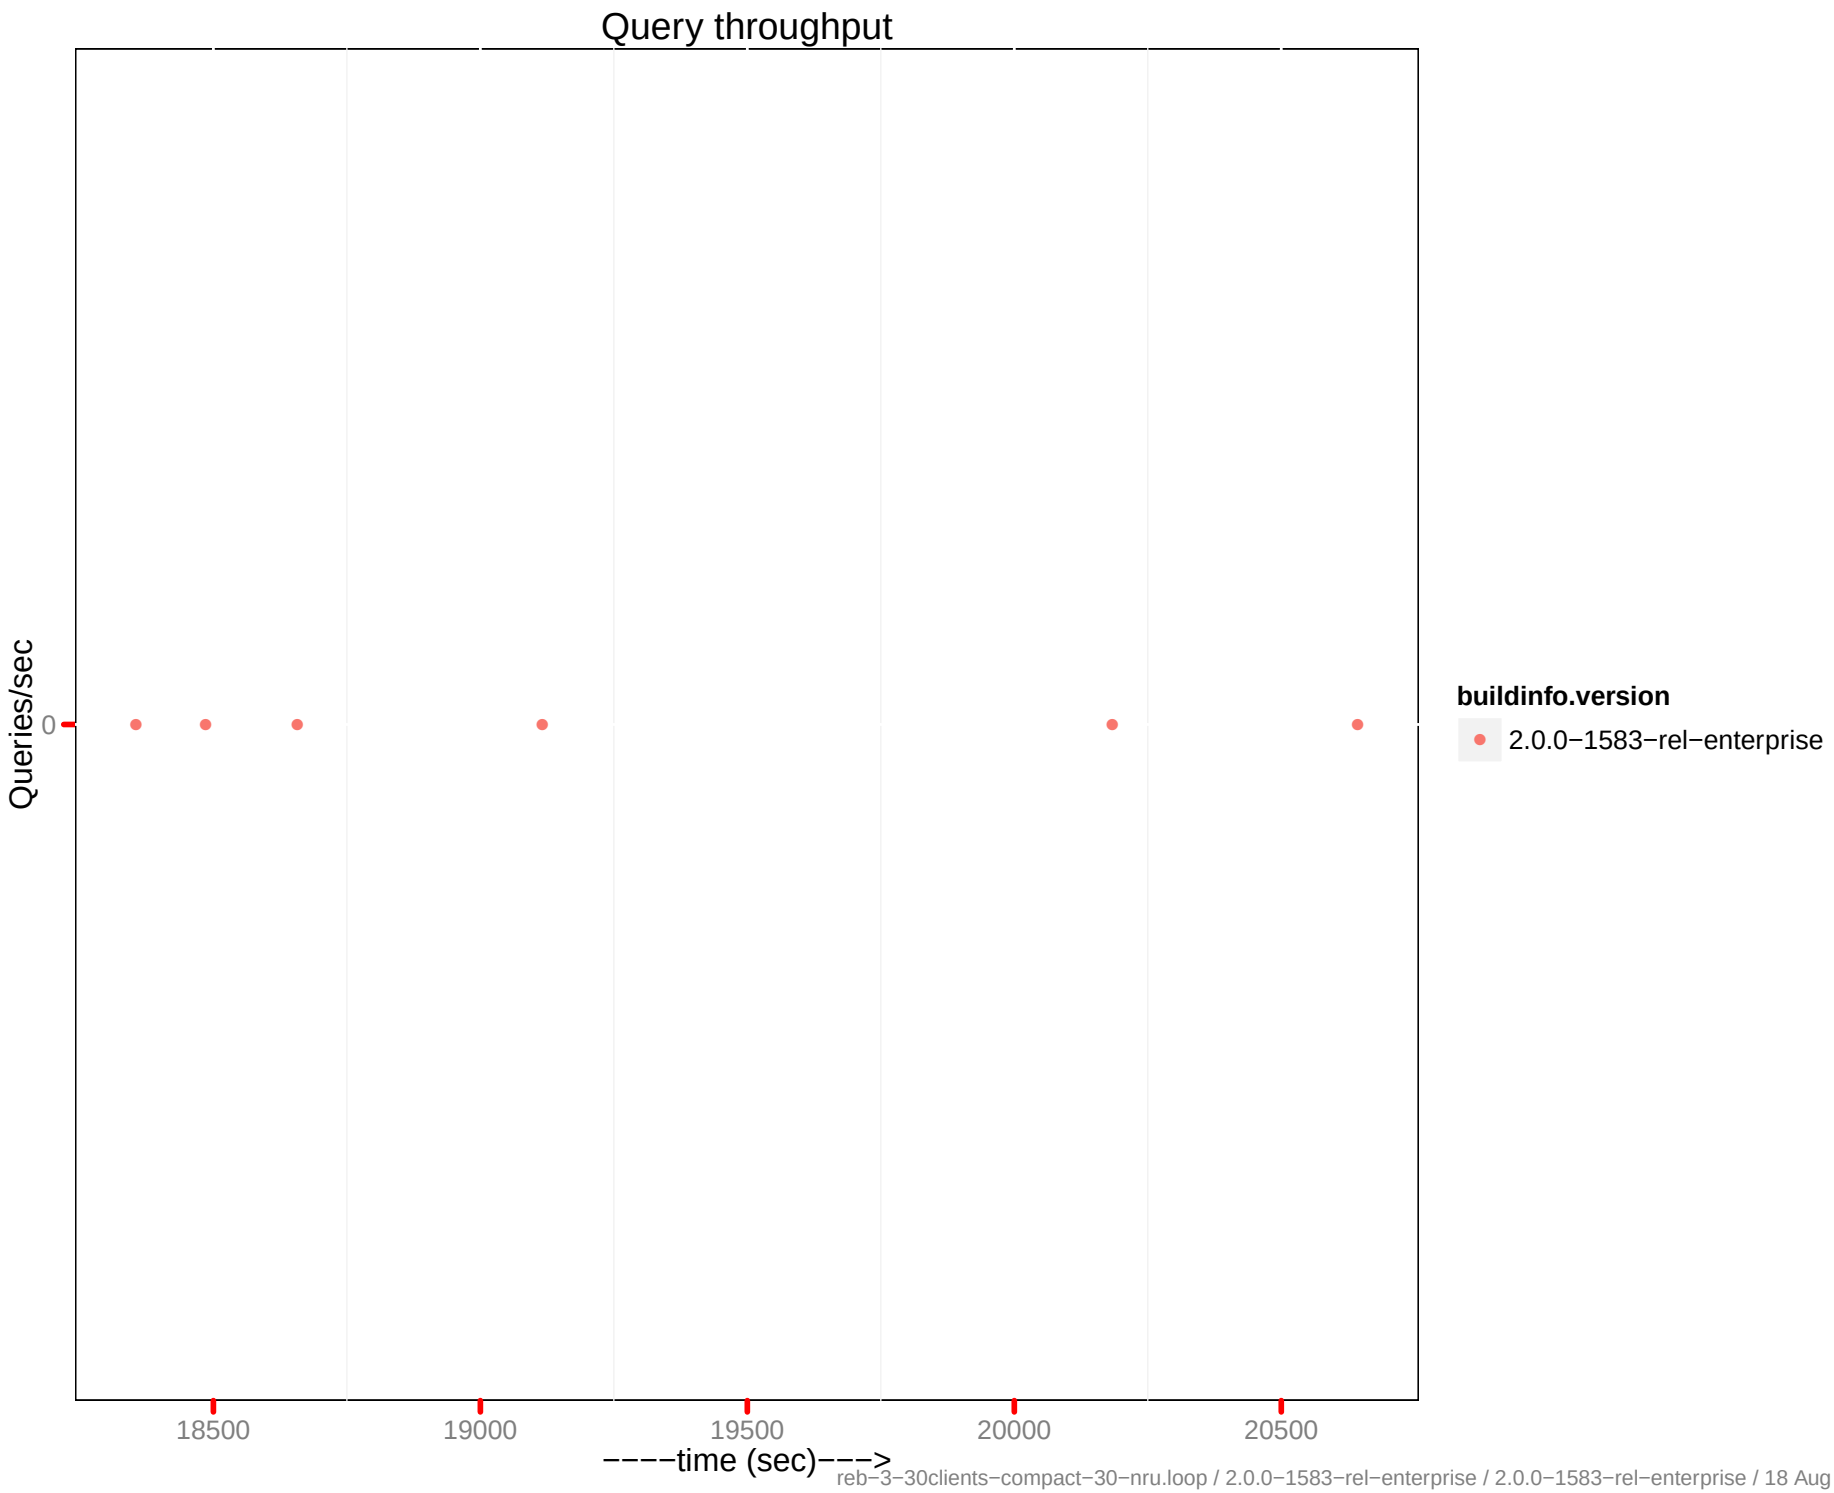
Query throughput
Queries/sec
0
18500
19000
19500
20000
20500
----time (sec)--->
buildinfo.version
2.0.0-1583-rel-enterprise
reb-3-30clients-compact-30-nru.loop / 2.0.0-1583-rel-enterprise / 2.0.0-1583-rel-enterprise / 18 Aug 2012

```
reb-3-30clients-compact-30-nru.conf
# rebalance mixed 10M load, 1M hot reload, 6M access creates
#
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=10000000

# 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=9000000

# rebalance
nru_task=1
nru_reb_delay=3600
reb_max_retries=5
num_nodes_after=3

# 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
```

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

[servers]
1:192.168.162.20
2:192.168.162.21
3:192.168.162.22
4:192.168.162.23

[clients]
1:192.168.162.24
2:192.168.162.25
3:192.168.162.26
4:192.168.162.27
5:192.168.162.28
6:192.168.162.29

[membase]
rest_username:Administrator
rest_password:password

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