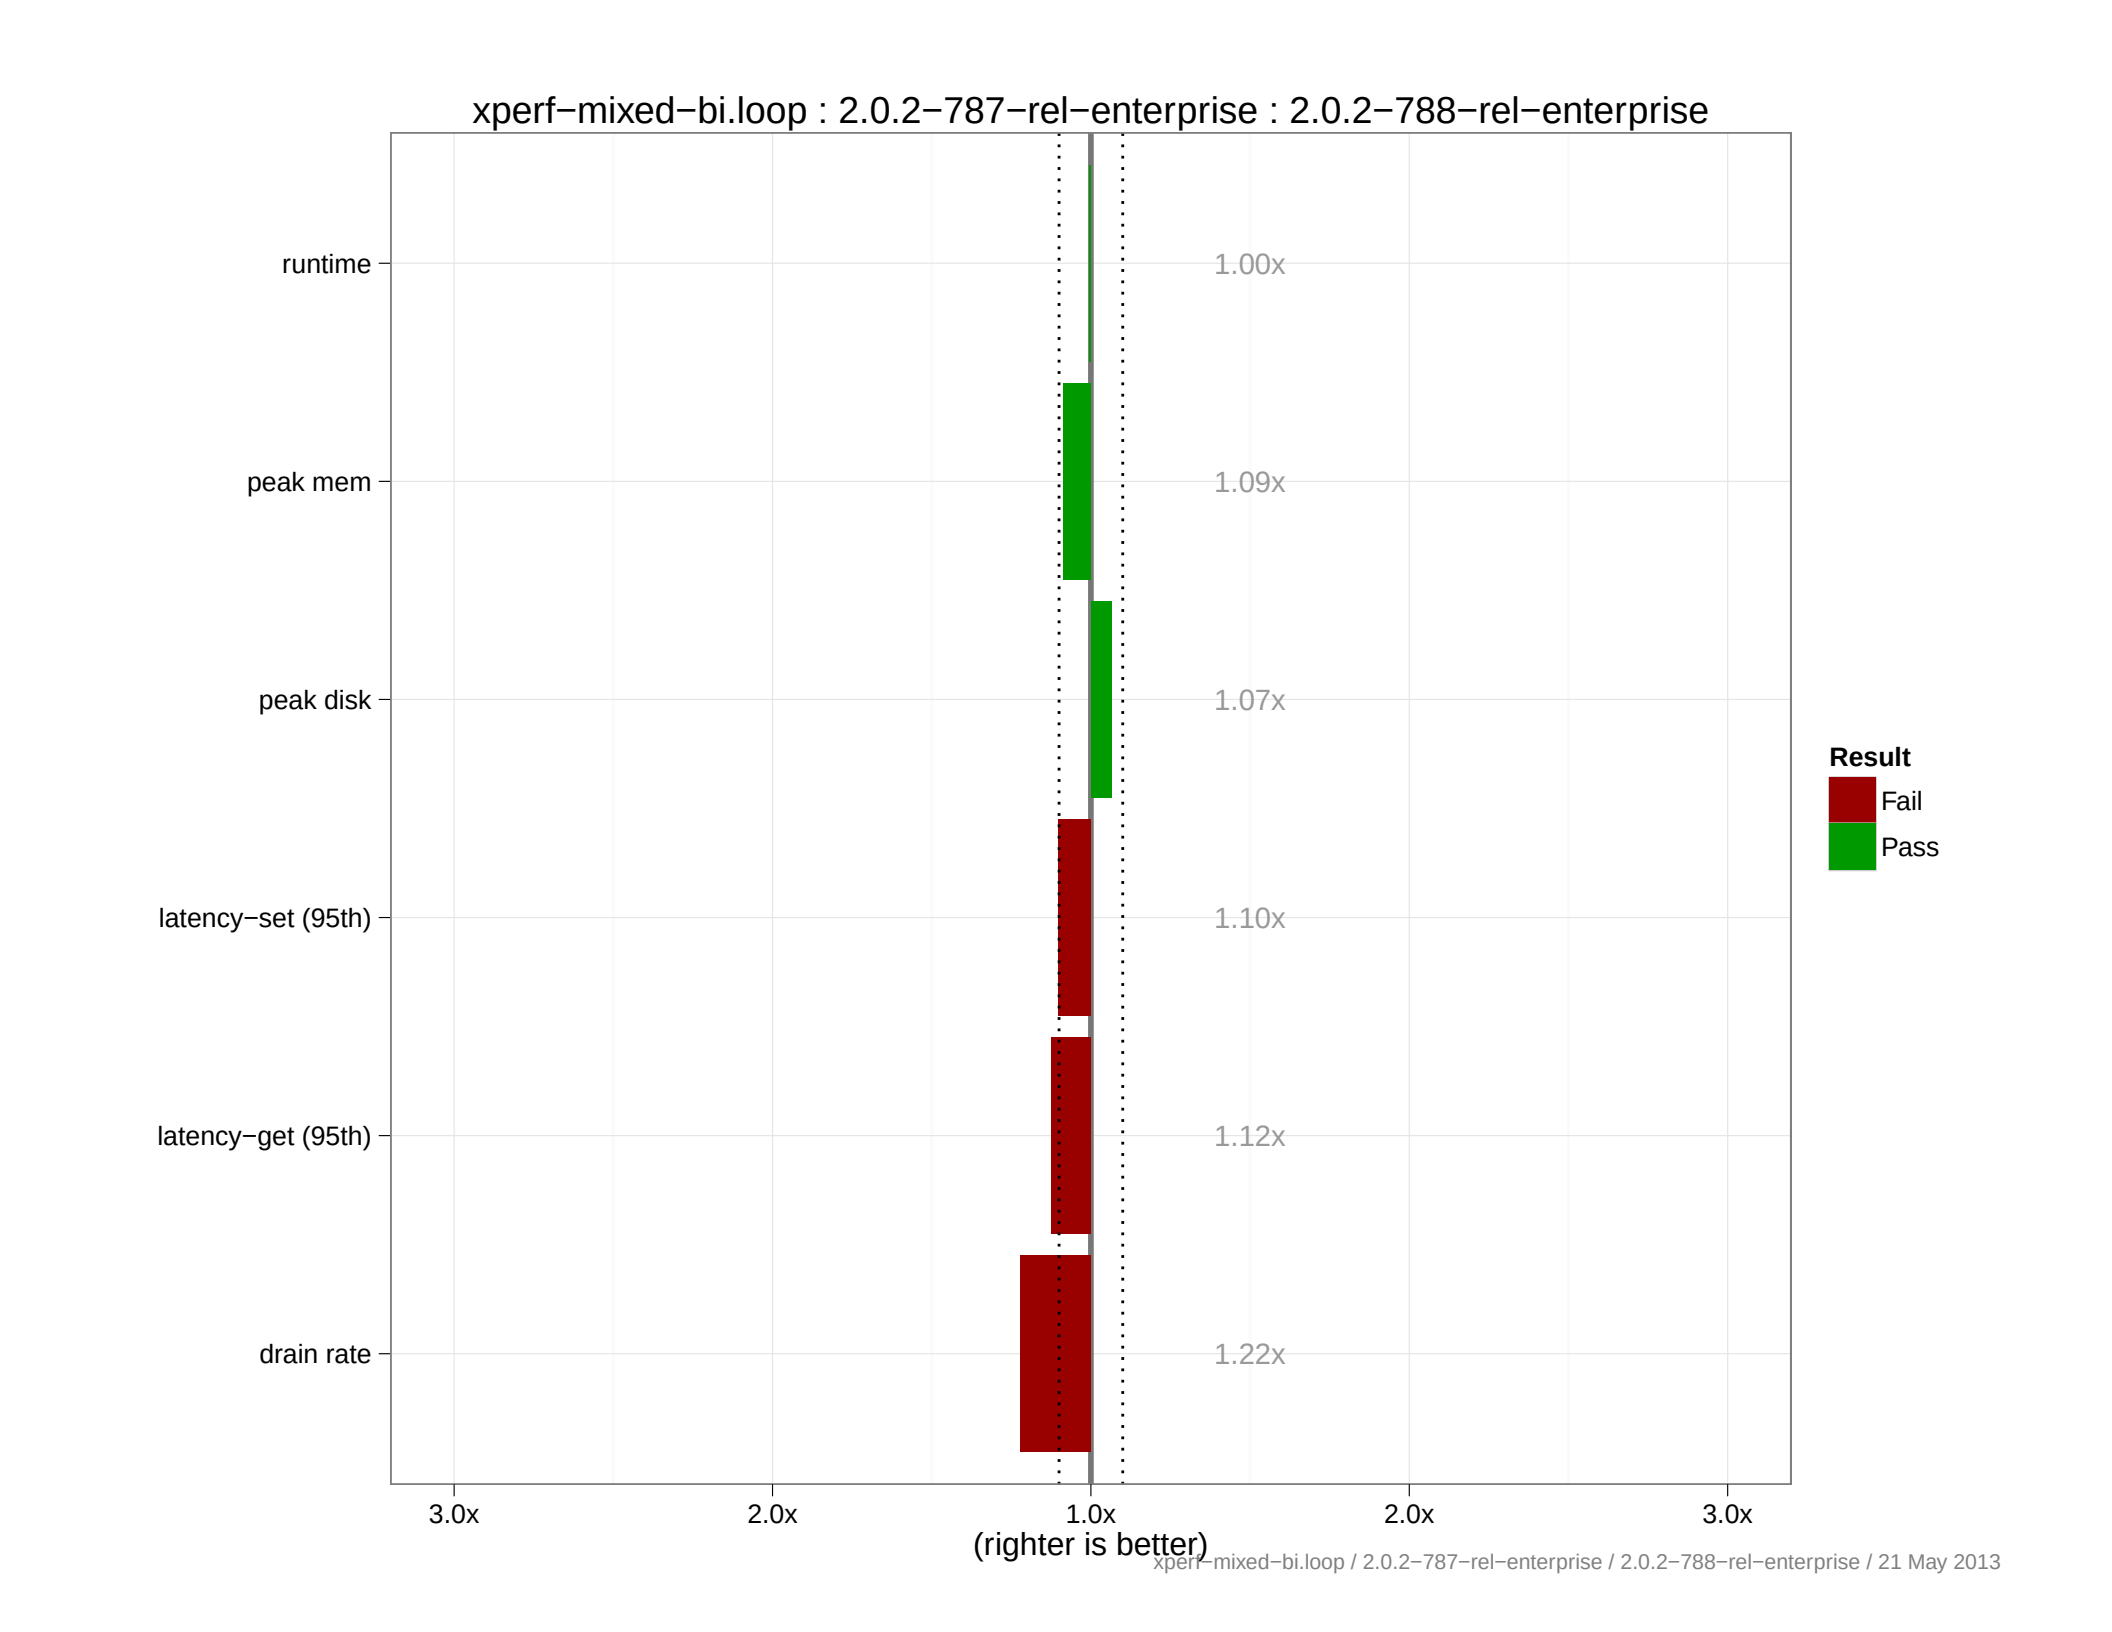
xperf−mixed−bi.loop : 2.0.2−787−rel−enterprise : 2.0.2−788−rel−enterprise
runtime
peak mem
peak disk
latency−set (95th)
latency−get (95th)
drain rate
1.00x
1.09x
1.07x
1.10x
1.12x
1.22x
Result
Fail
Pass
3.0x
2.0x
1.0x
2.0x
3.0x
(righter is better)
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

	2.0.2 − 787	2.0.2 − 788
Runtime (in hr)	6.03	6.05
Avg. Drain Rate	12.55K	10.28K
Peak Disk (GB)	80.58	75.63
Peak Memory (GB)	70124.33	76220.14
Avg. OPS	17.48K	13.33K
Avg. mem memcached (GB)	57154.24	57443.51
Avg. mem beam.smp (MB)	918218.52	931545.69
Avg. CPU rate (%)	82.5	83.11
Latency−get (90th) (ms)	7.13	7.85
Latency−get (95th) (ms)	9.67	10.88
Latency−get (99th) (ms)	22.34	23.72
Latency−set (90th) (ms)	7.35	7.95
Latency−set (95th) (ms)	9.89	10.89
Latency−set (99th) (ms)	22.53	23.48
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
Avg. XDC ops/sec	5342.18	2680.05
Avg. XDC docs to replicate	80163.17	1477713.84
Rebalance Time (sec)	0	0
Testrunner Version	26873dc	f98ef42

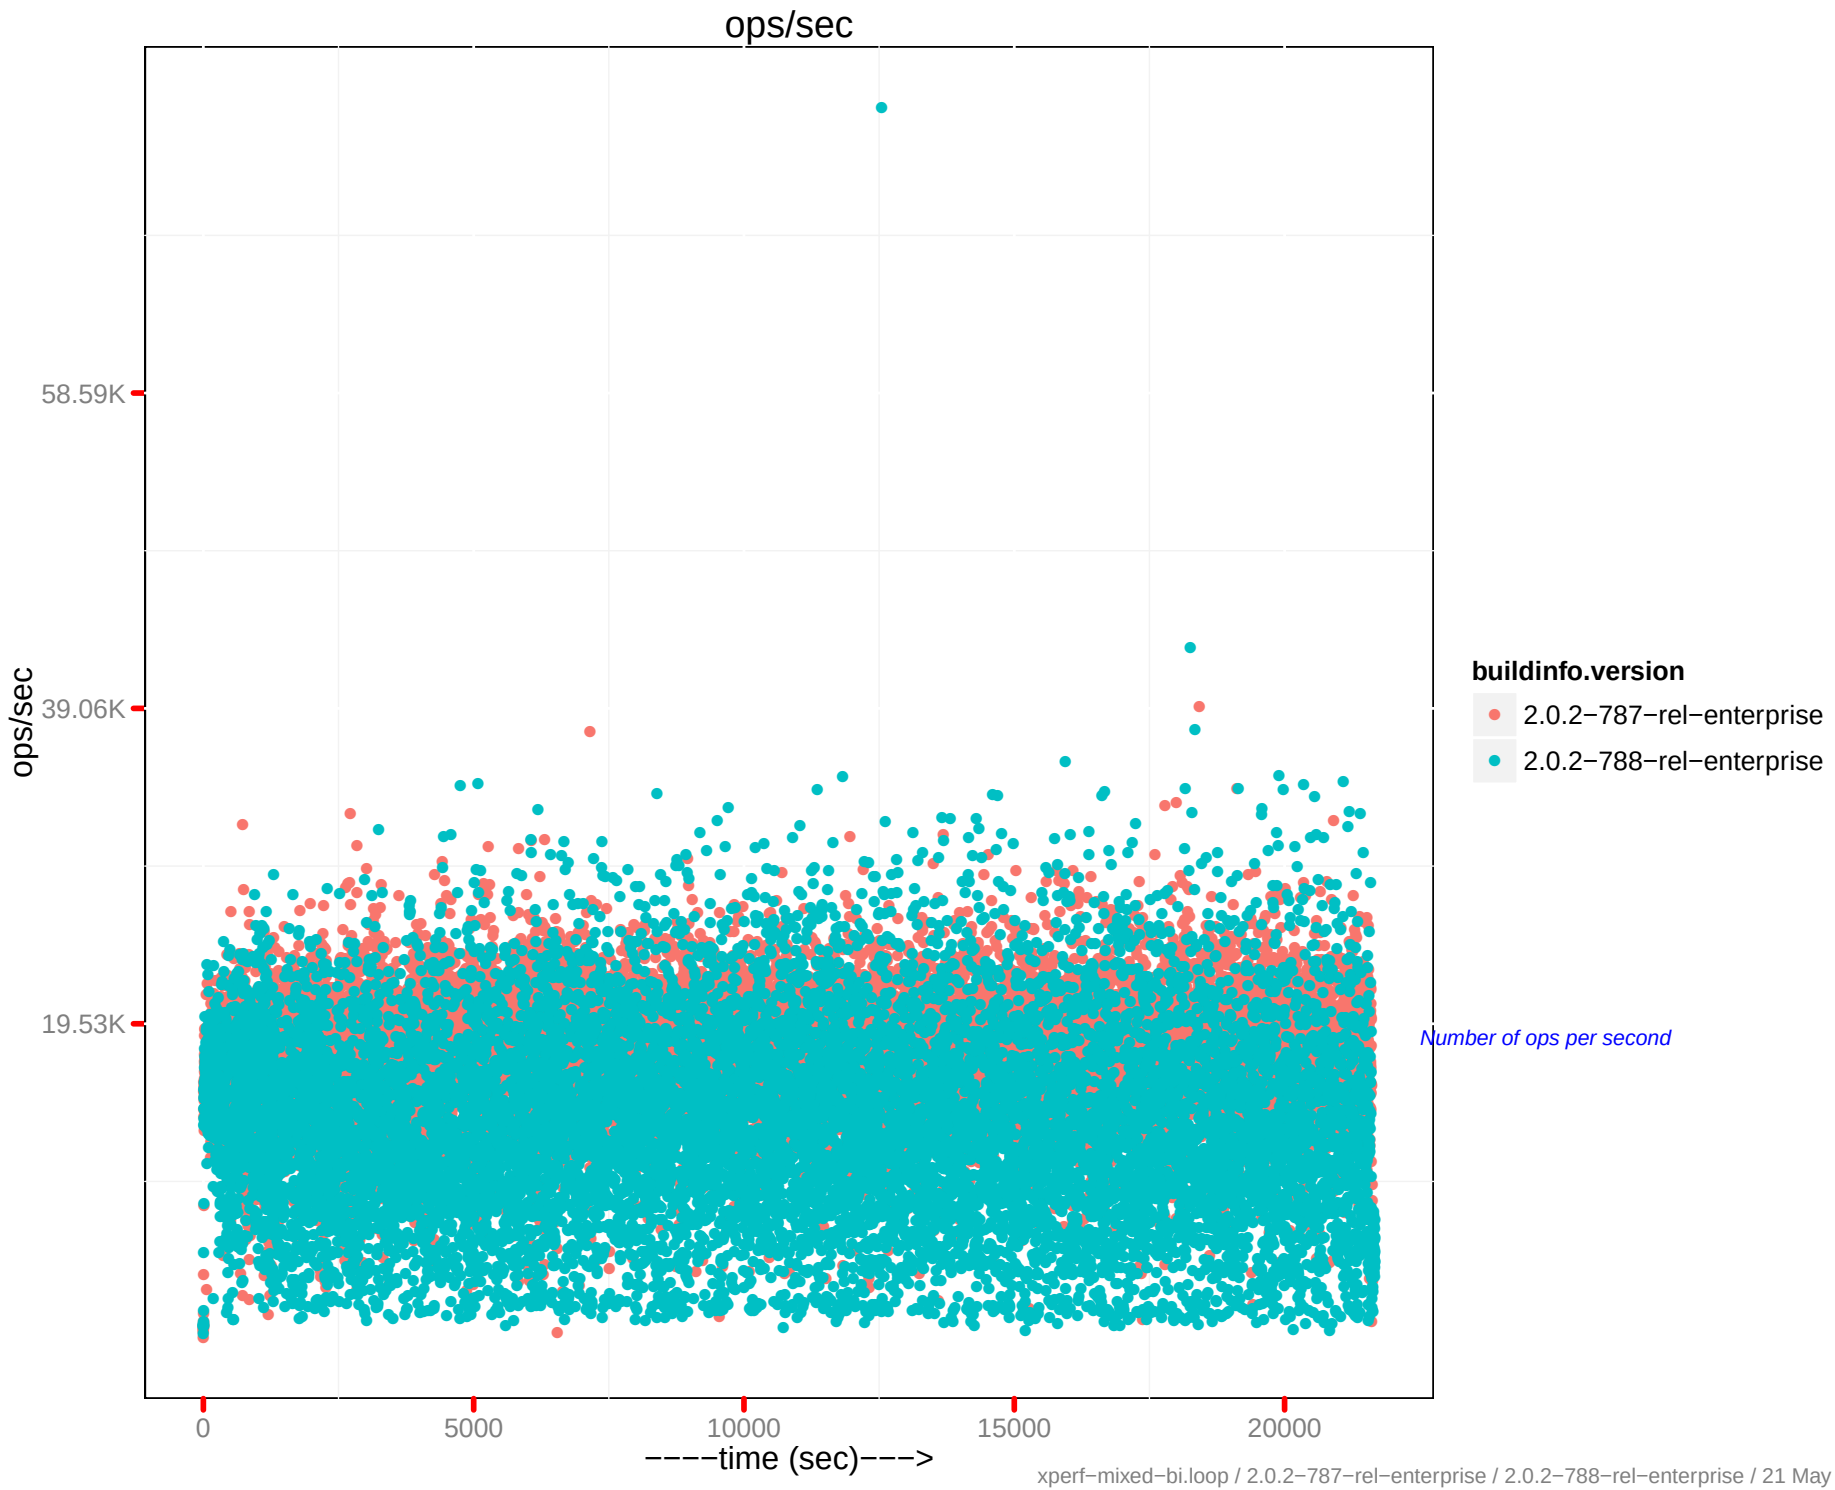
ops/sec
58.59K
39.06K
19.53K
ops/sec
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of ops per second
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

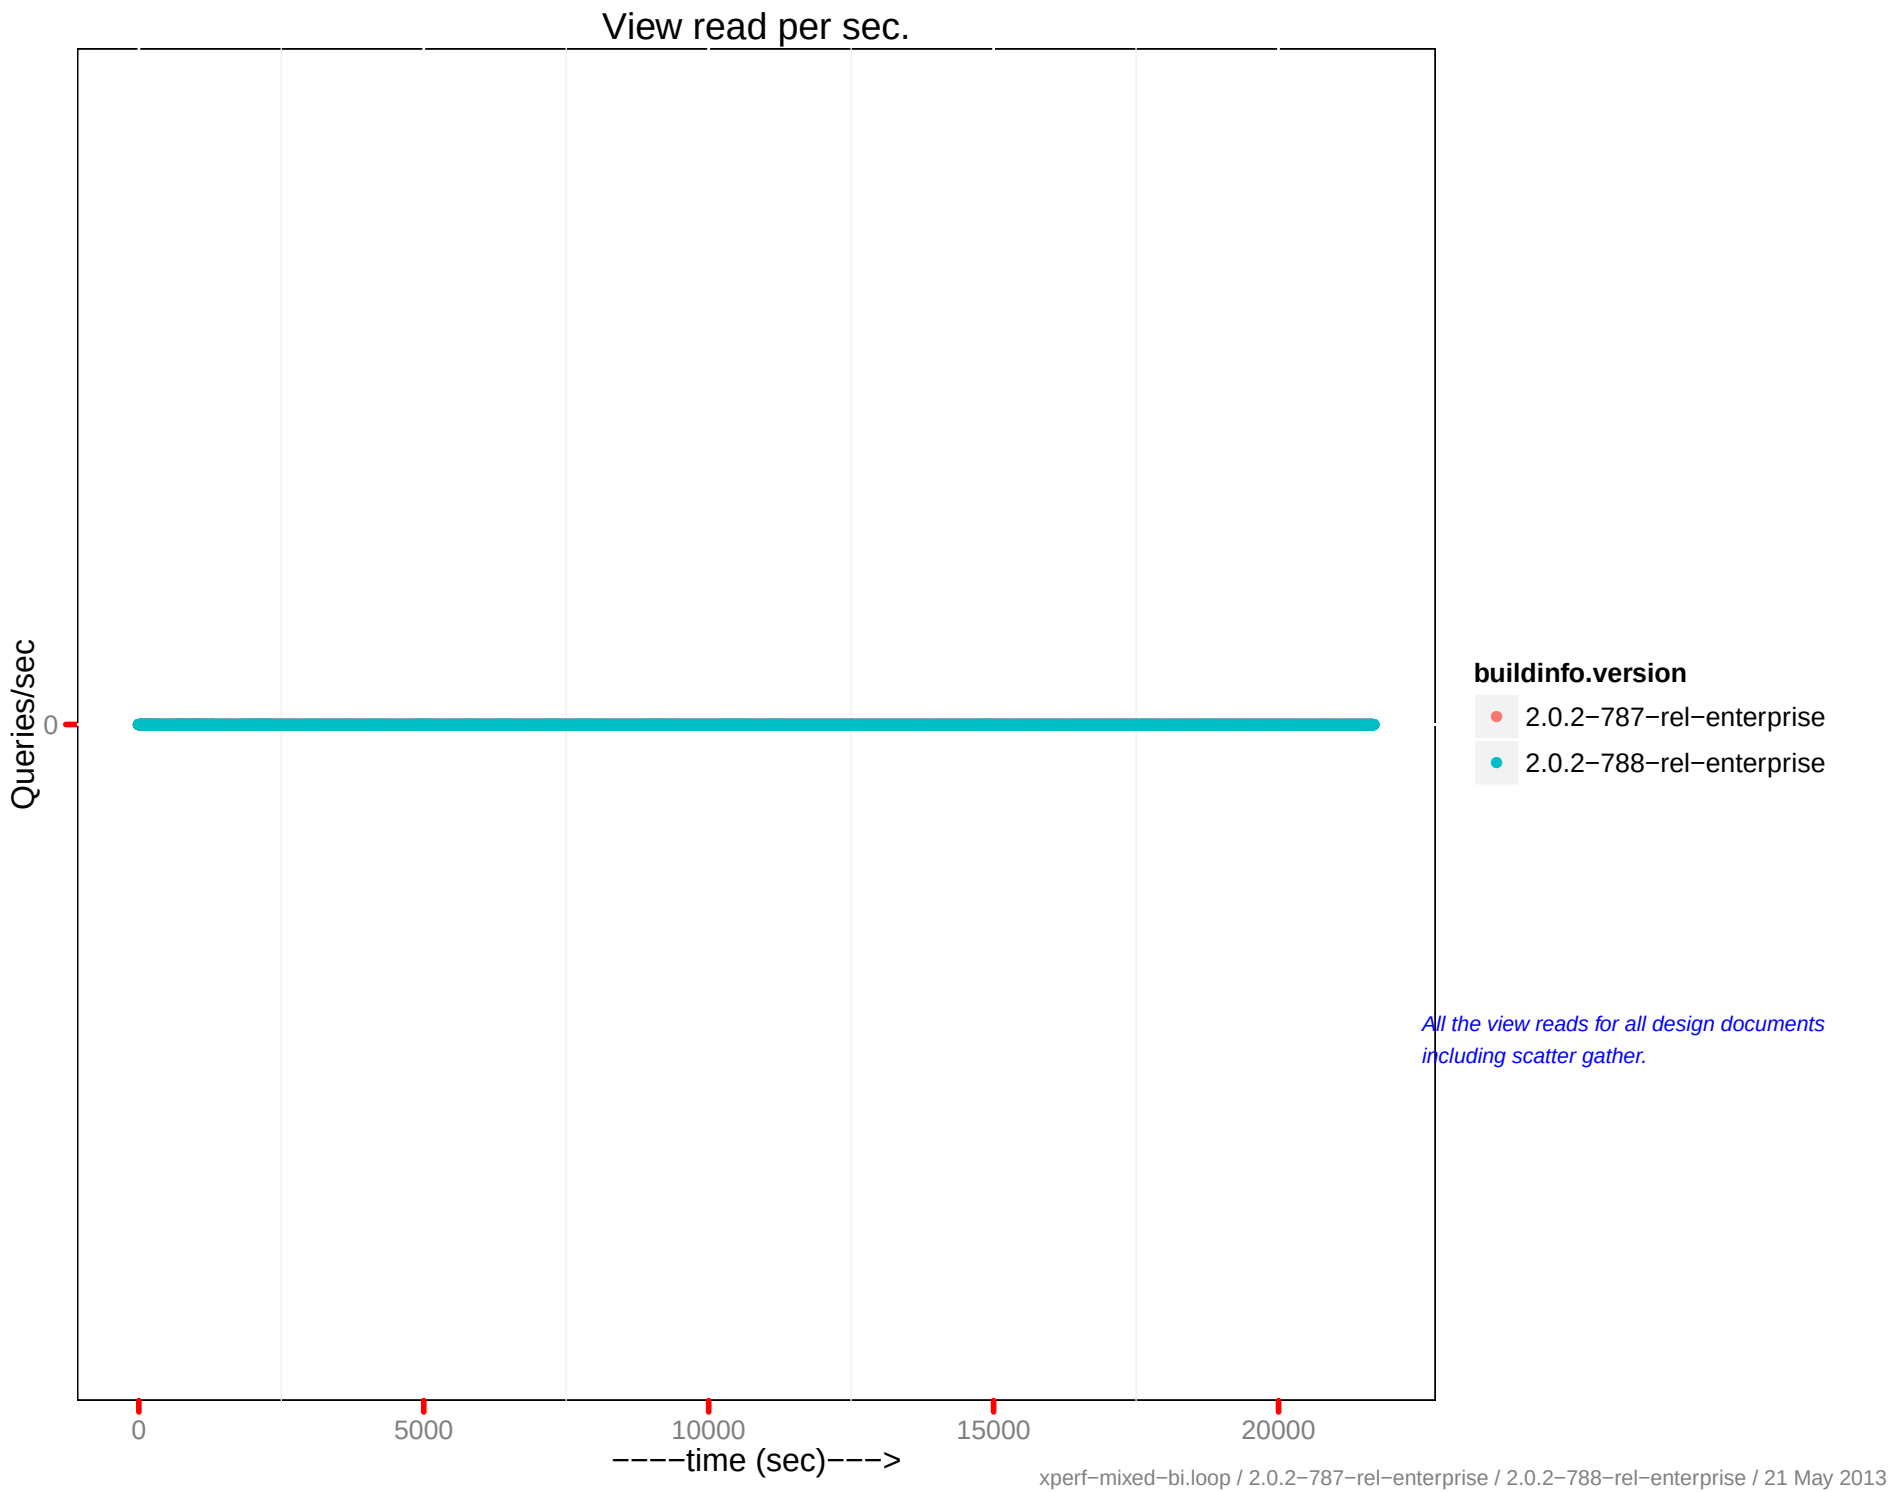
View read per sec.
Queries/sec
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
All the view reads for all design documents including scatter gather.
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

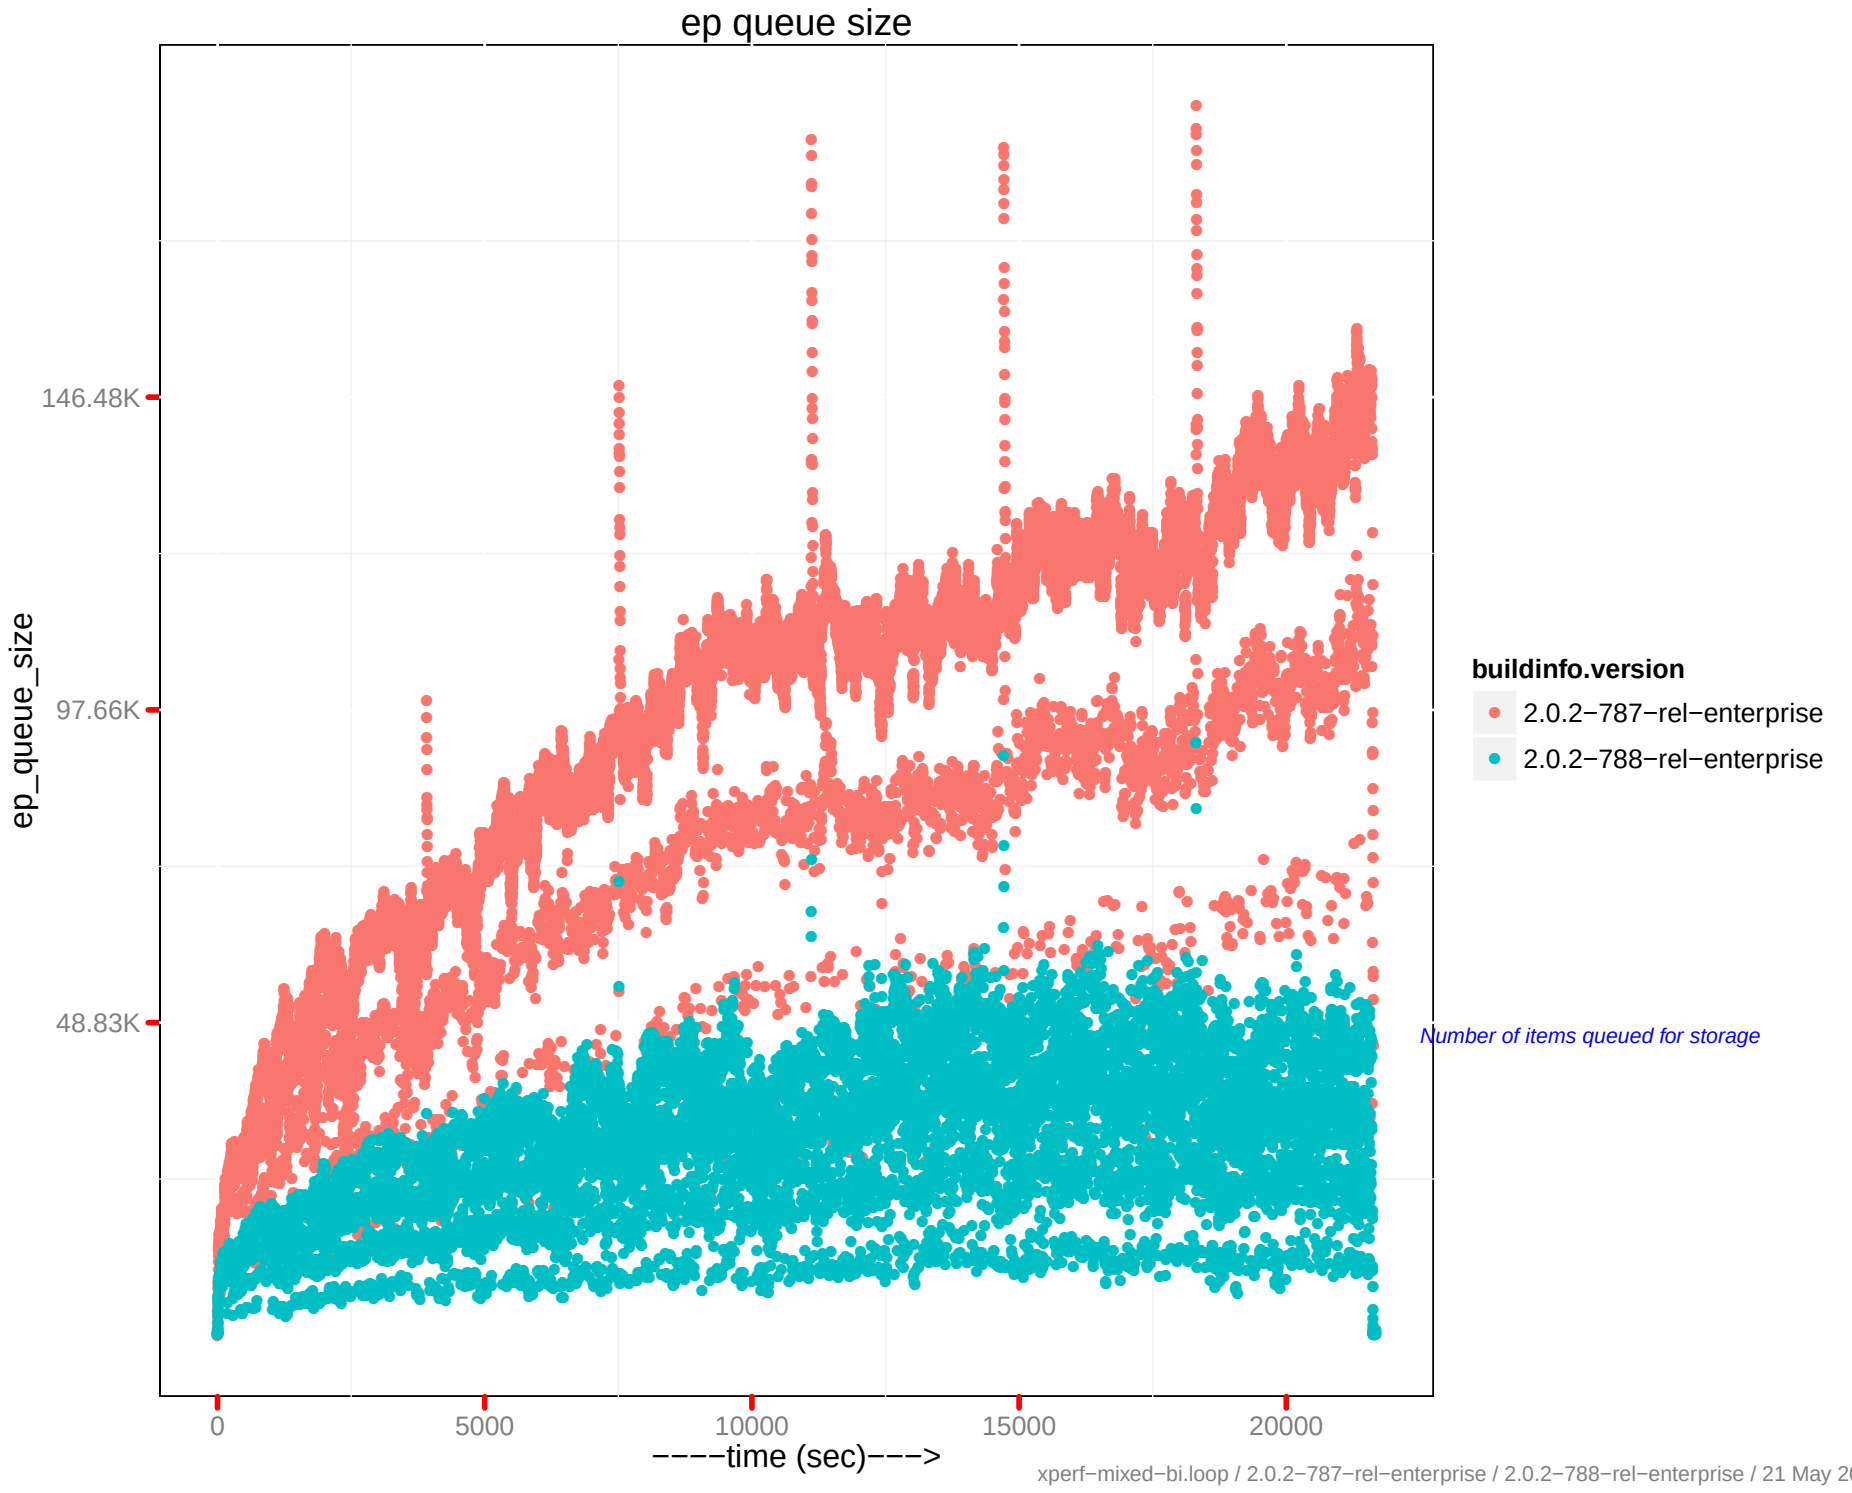
ep queue size
146.48K
97.66K
48.83K
ep_queue_size
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of items queued for storage
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

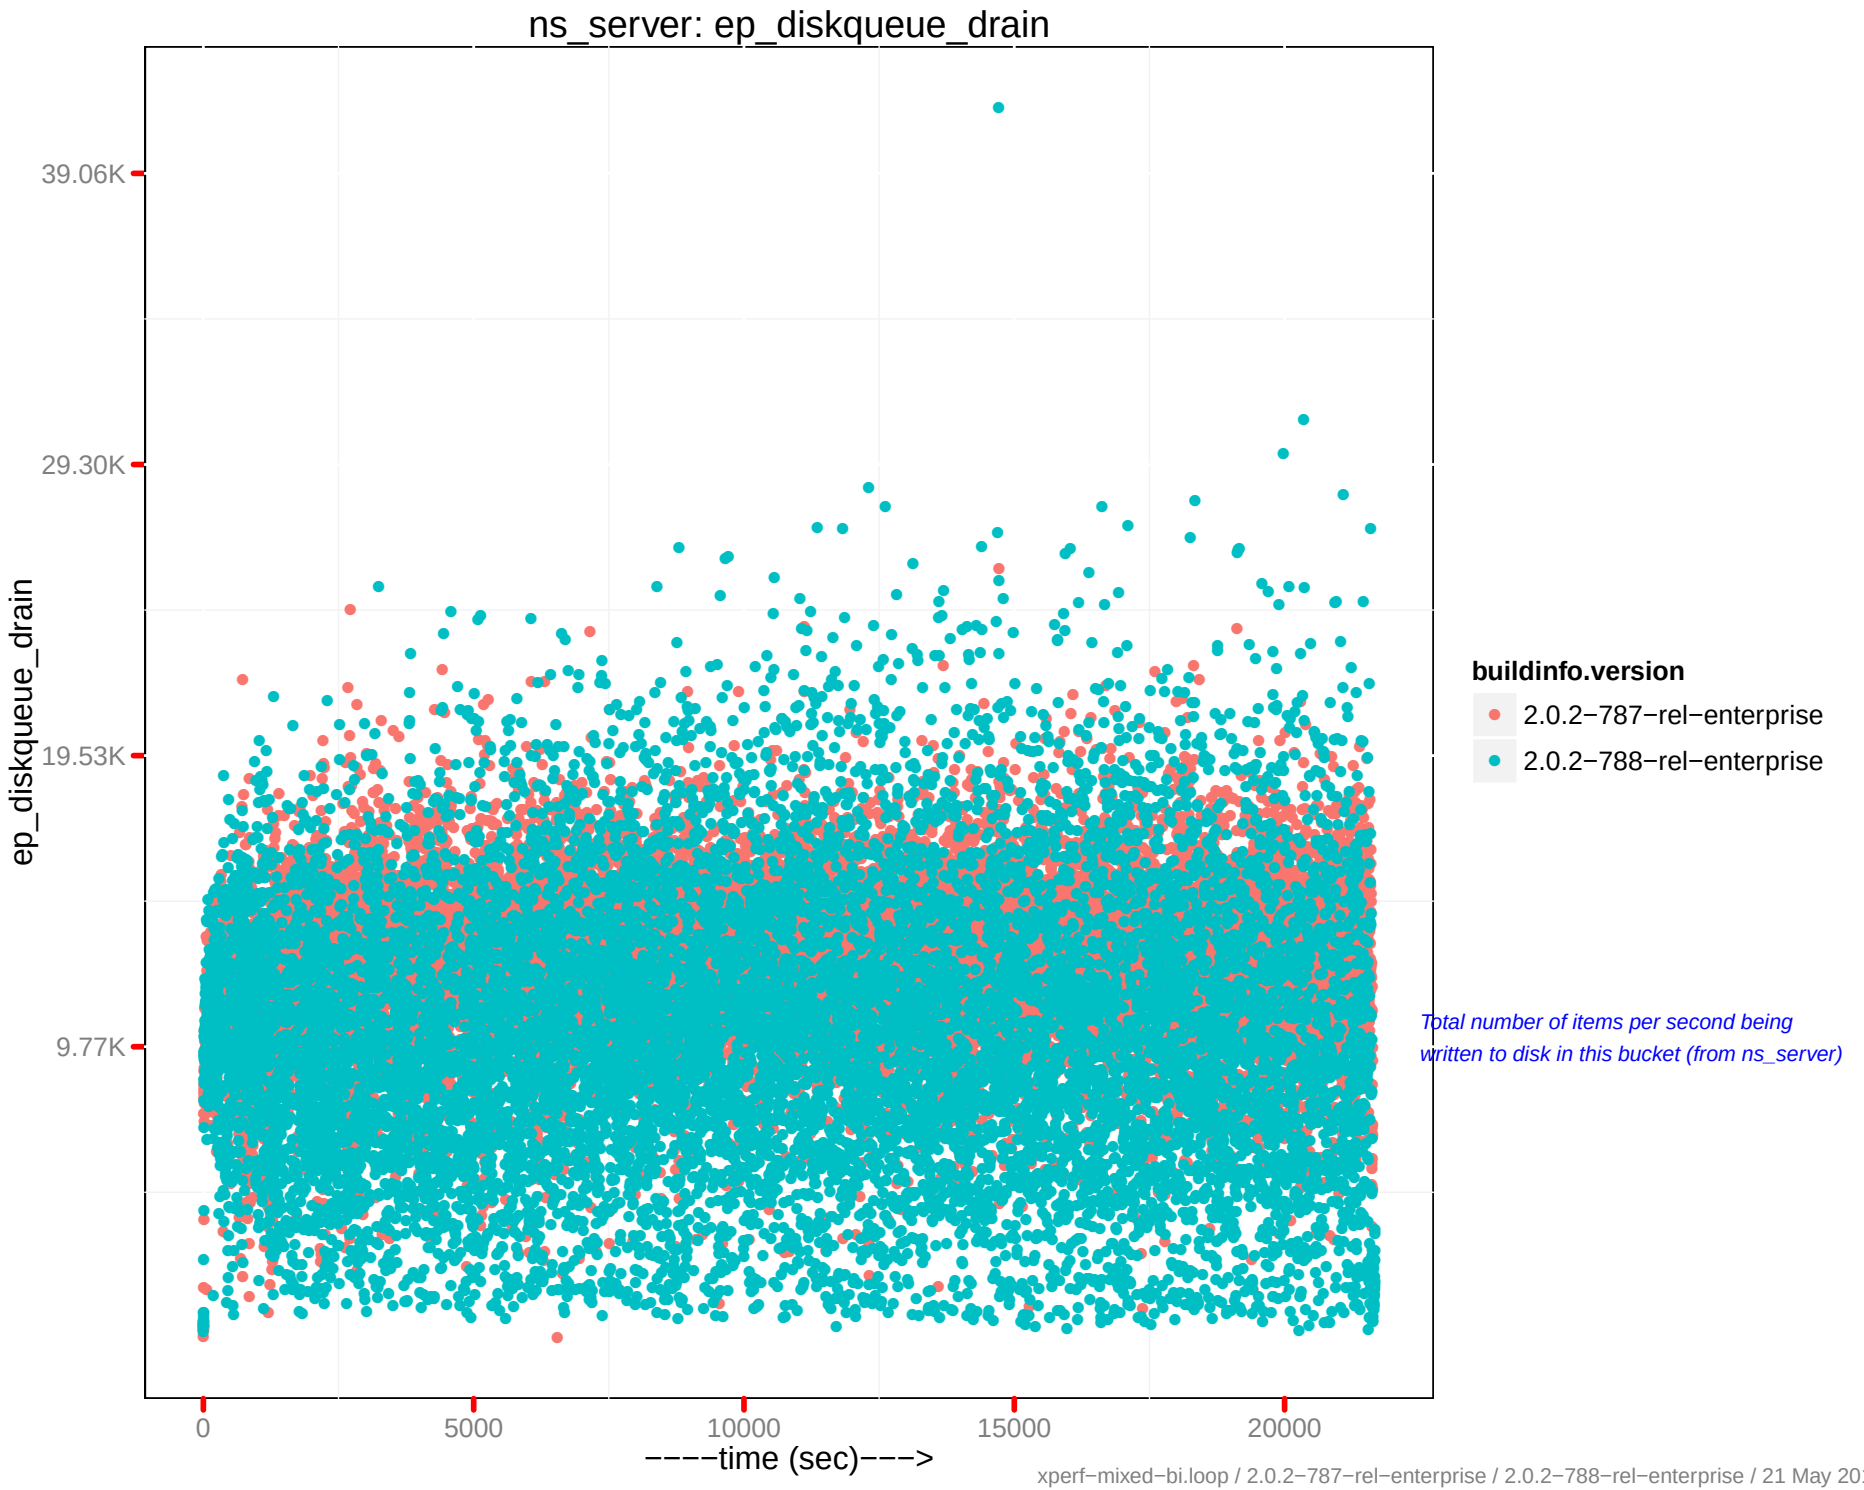
ns_server: ep_diskqueue_drain
39.06K
29.30K
19.53K
9.77K
ep_diskqueue_drain
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Total number of items per second being written to disk in this bucket (from ns_server)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

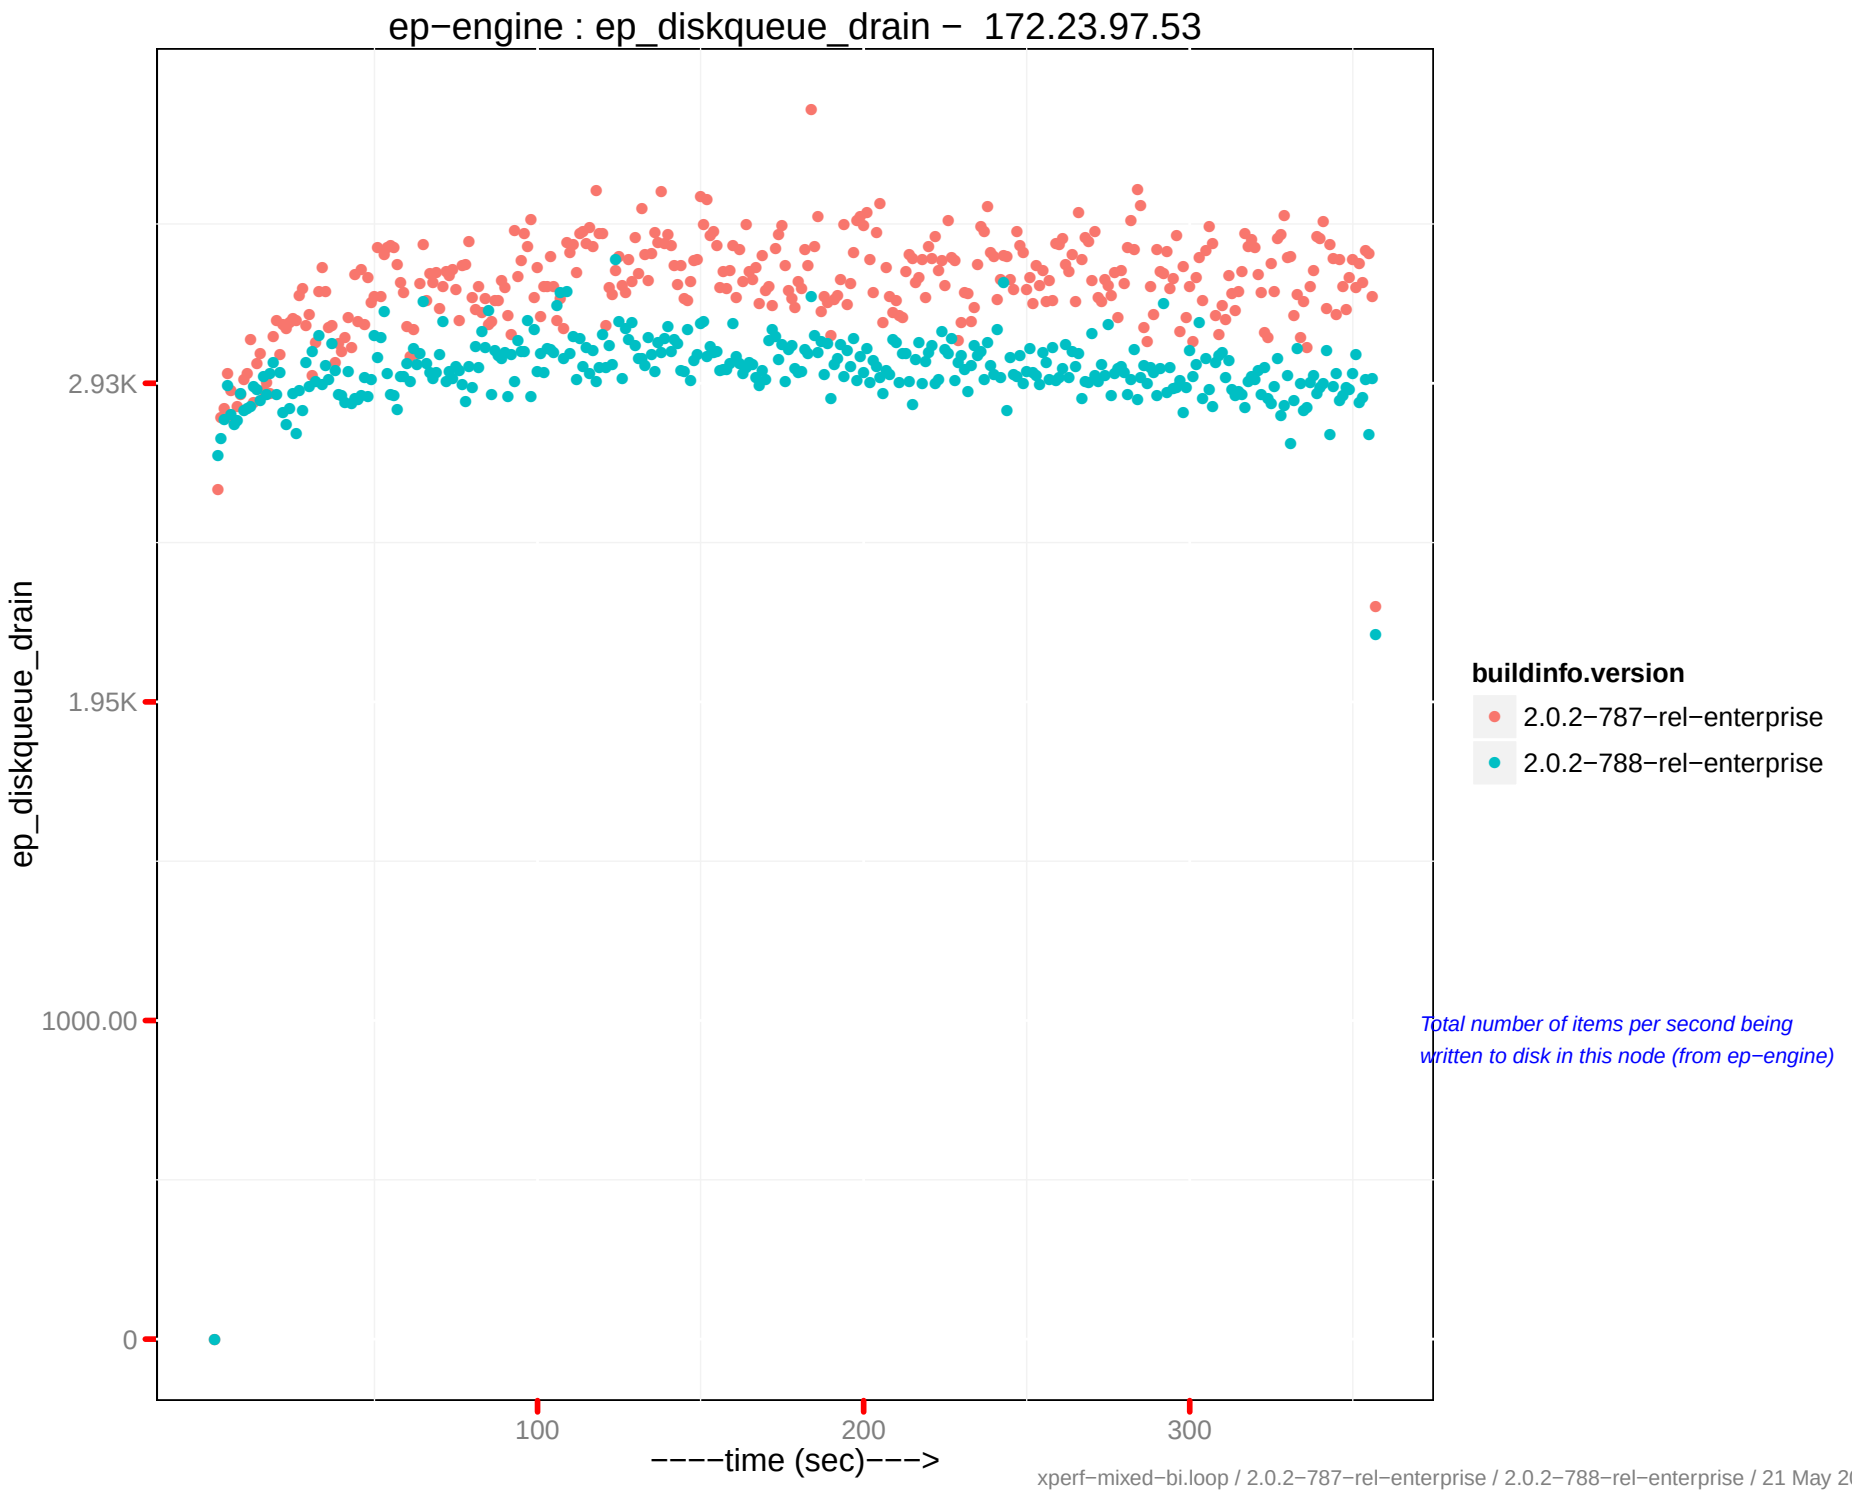
ep−engine : ep_diskqueue_drain − 172.23.97.53
ep_diskqueue_drain
2.93K
1.95K
1000.00
0
100
200
300
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

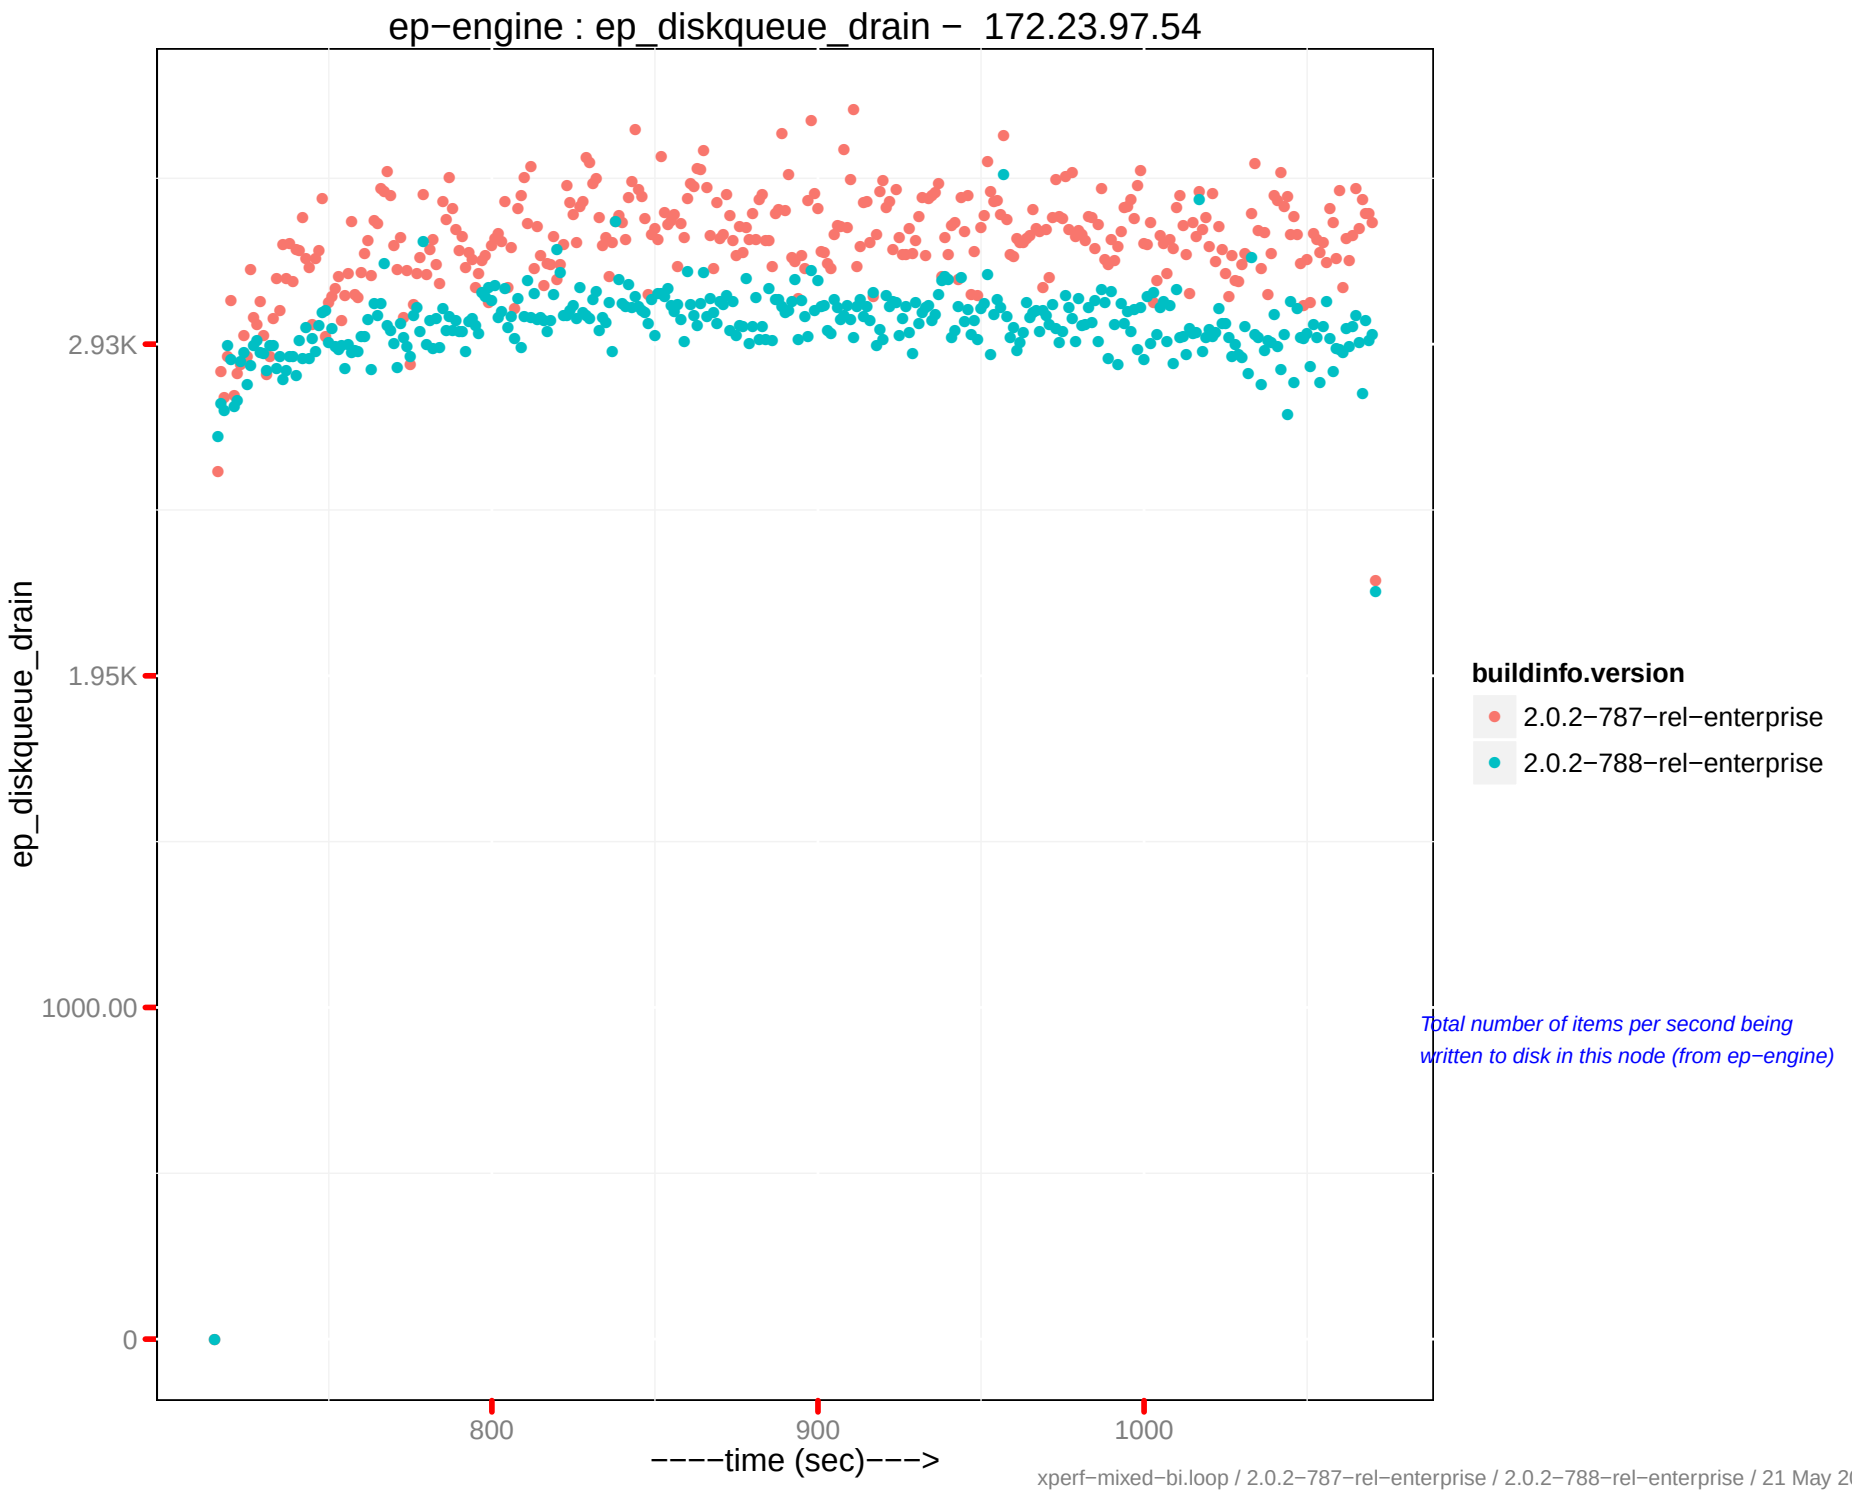
ep−engine : ep_diskqueue_drain − 172.23.97.54
ep_diskqueue_drain
2.93K
1.95K
1000.00
0
800
900
1000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

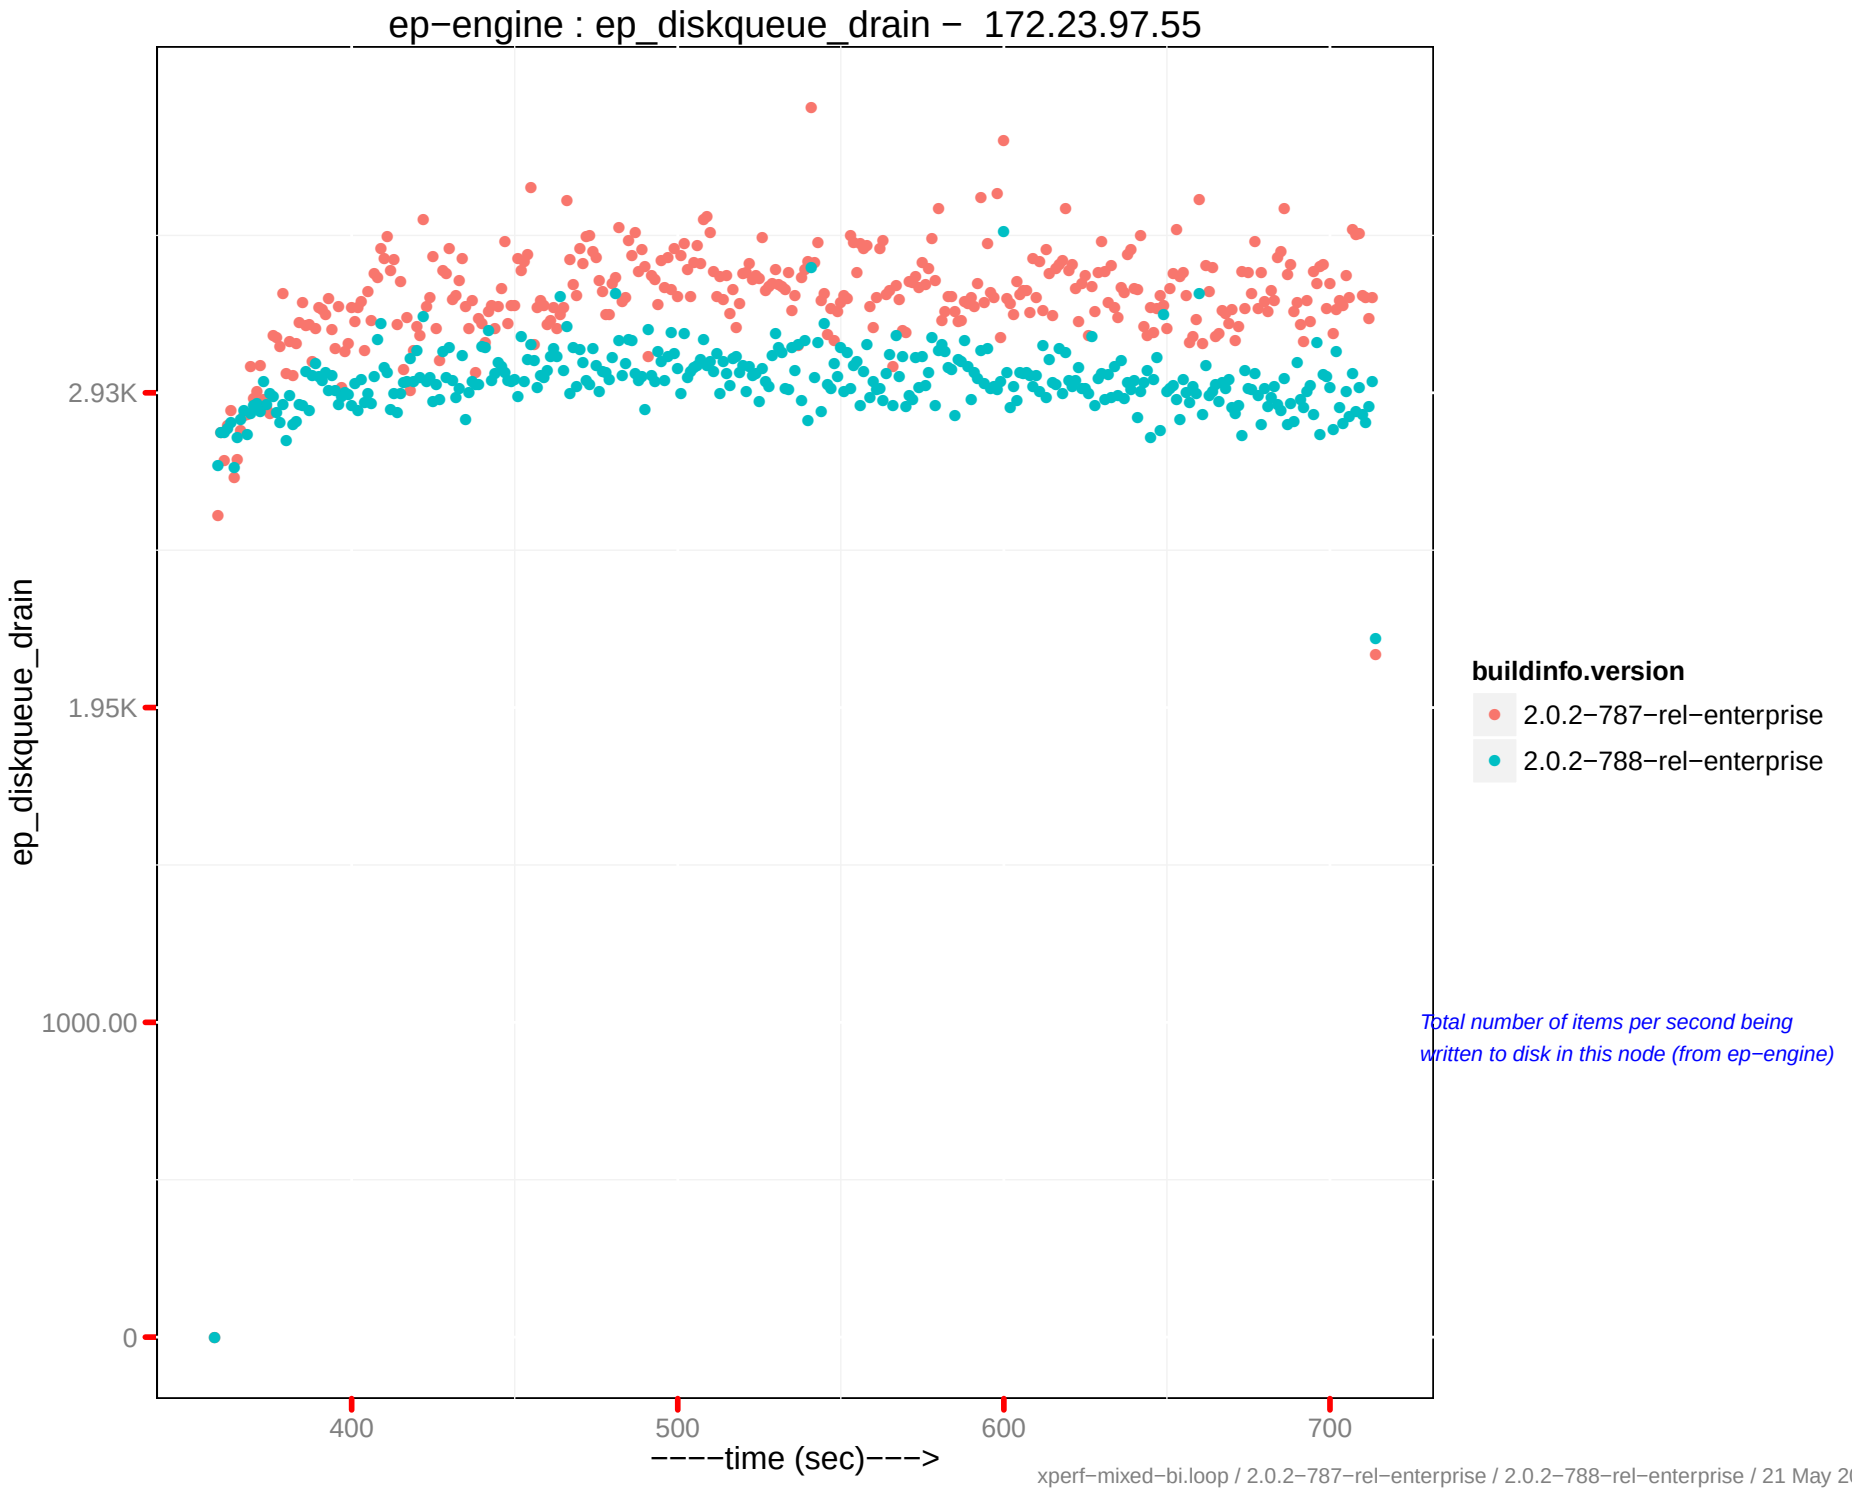
ep−engine : ep_diskqueue_drain − 172.23.97.55
ep_diskqueue_drain
2.93K
1.95K
1000.00
0
400
500
600
700
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

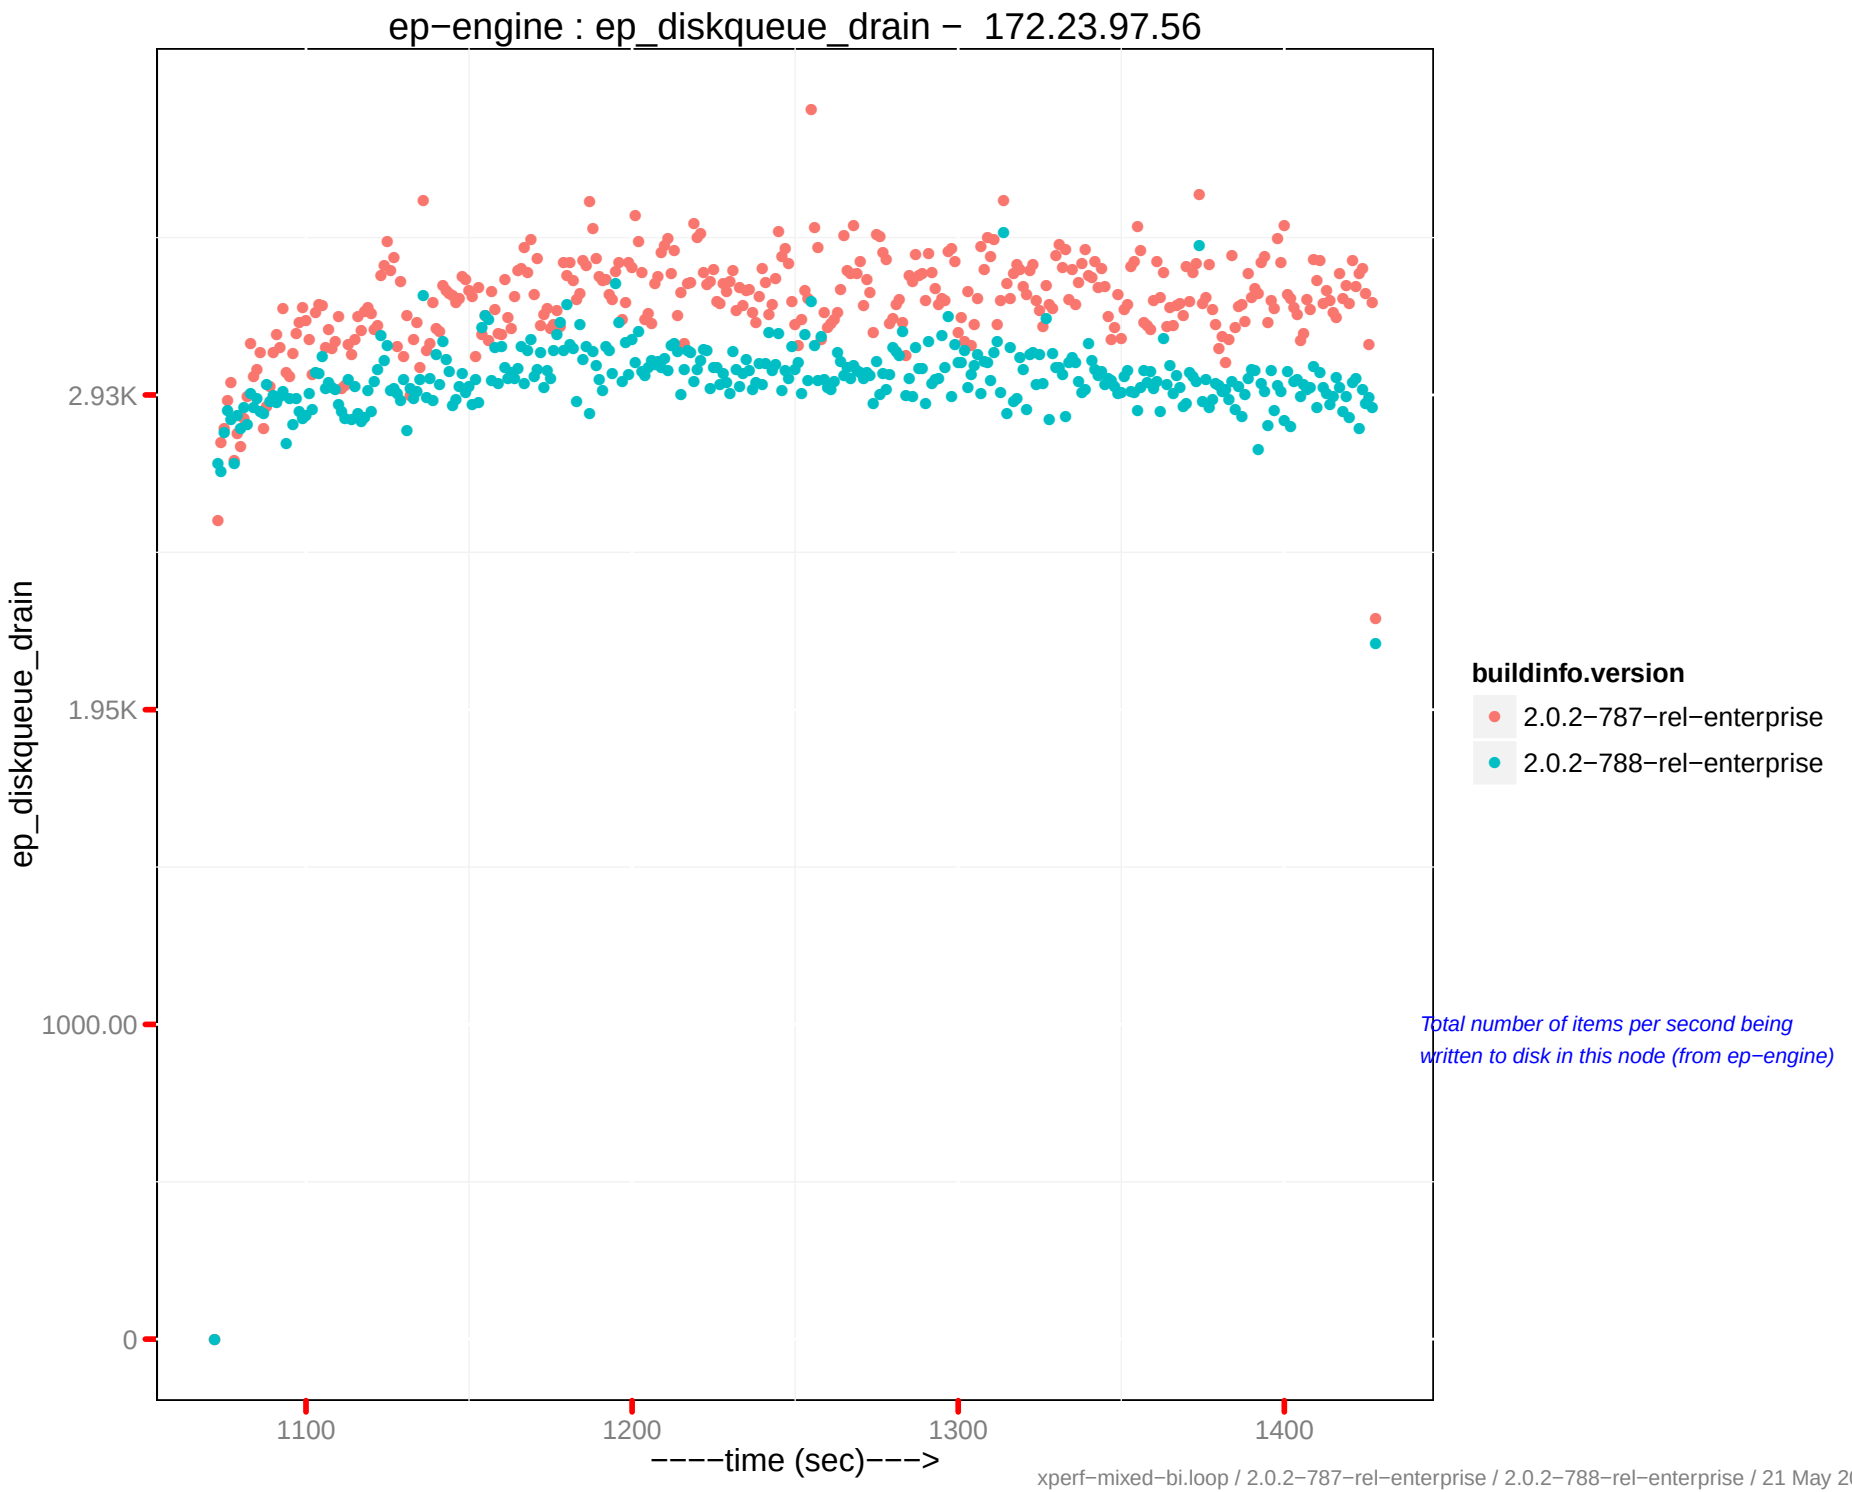
ep−engine : ep_diskqueue_drain − 172.23.97.56
ep_diskqueue_drain
2.93K
1.95K
1000.00
0
1100
1200
1300
1400
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Total number of items per second being written to disk in this node (from ep−engine)
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

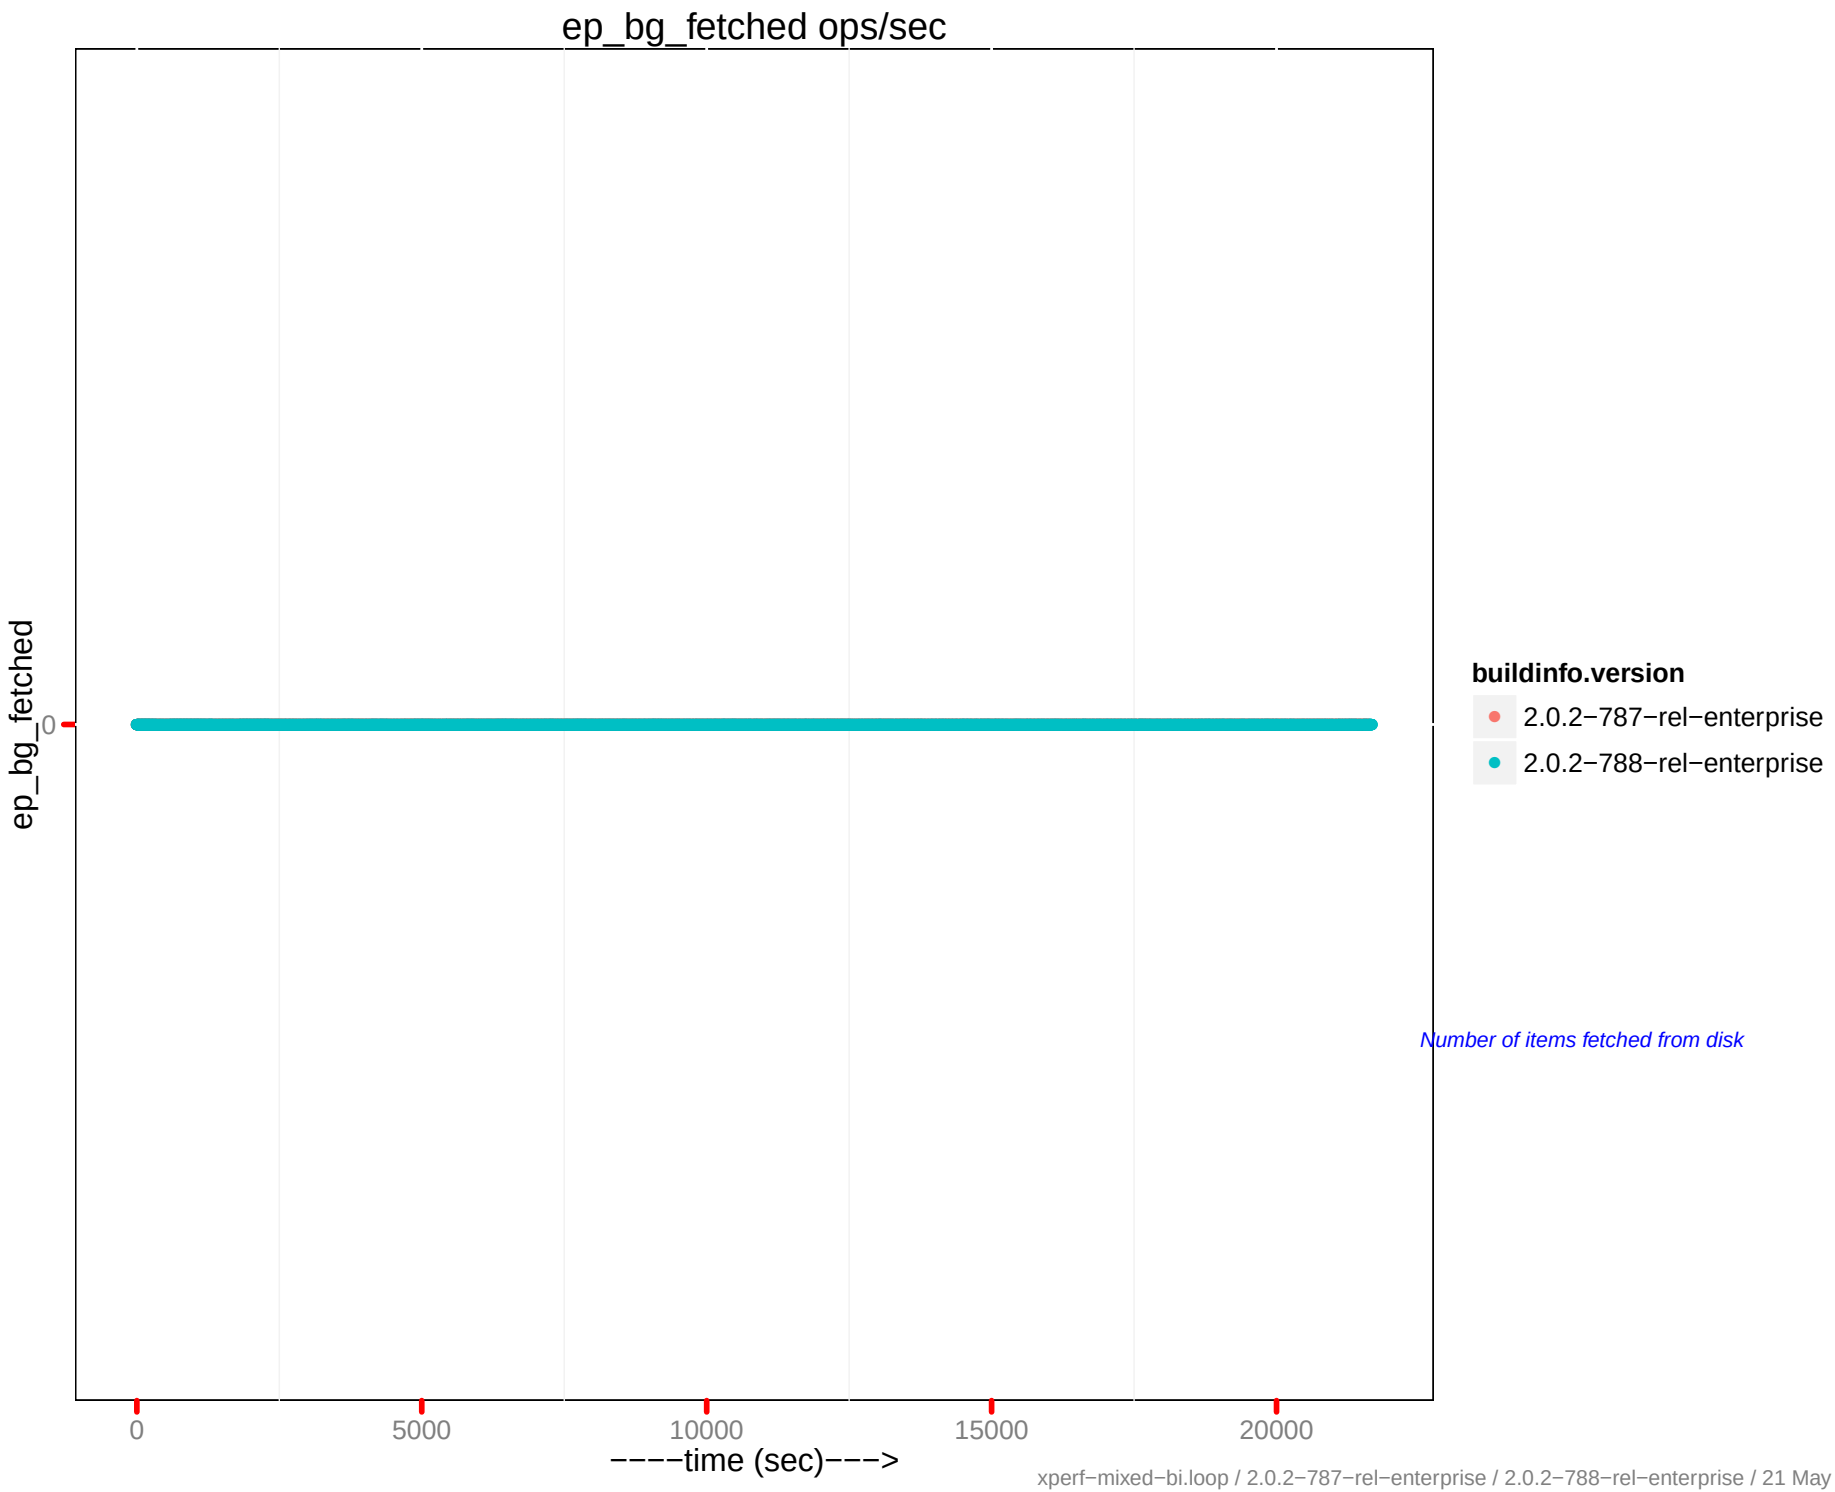
ep_bg_fetched ops/sec
ep_bg_fetched
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of items fetched from disk
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

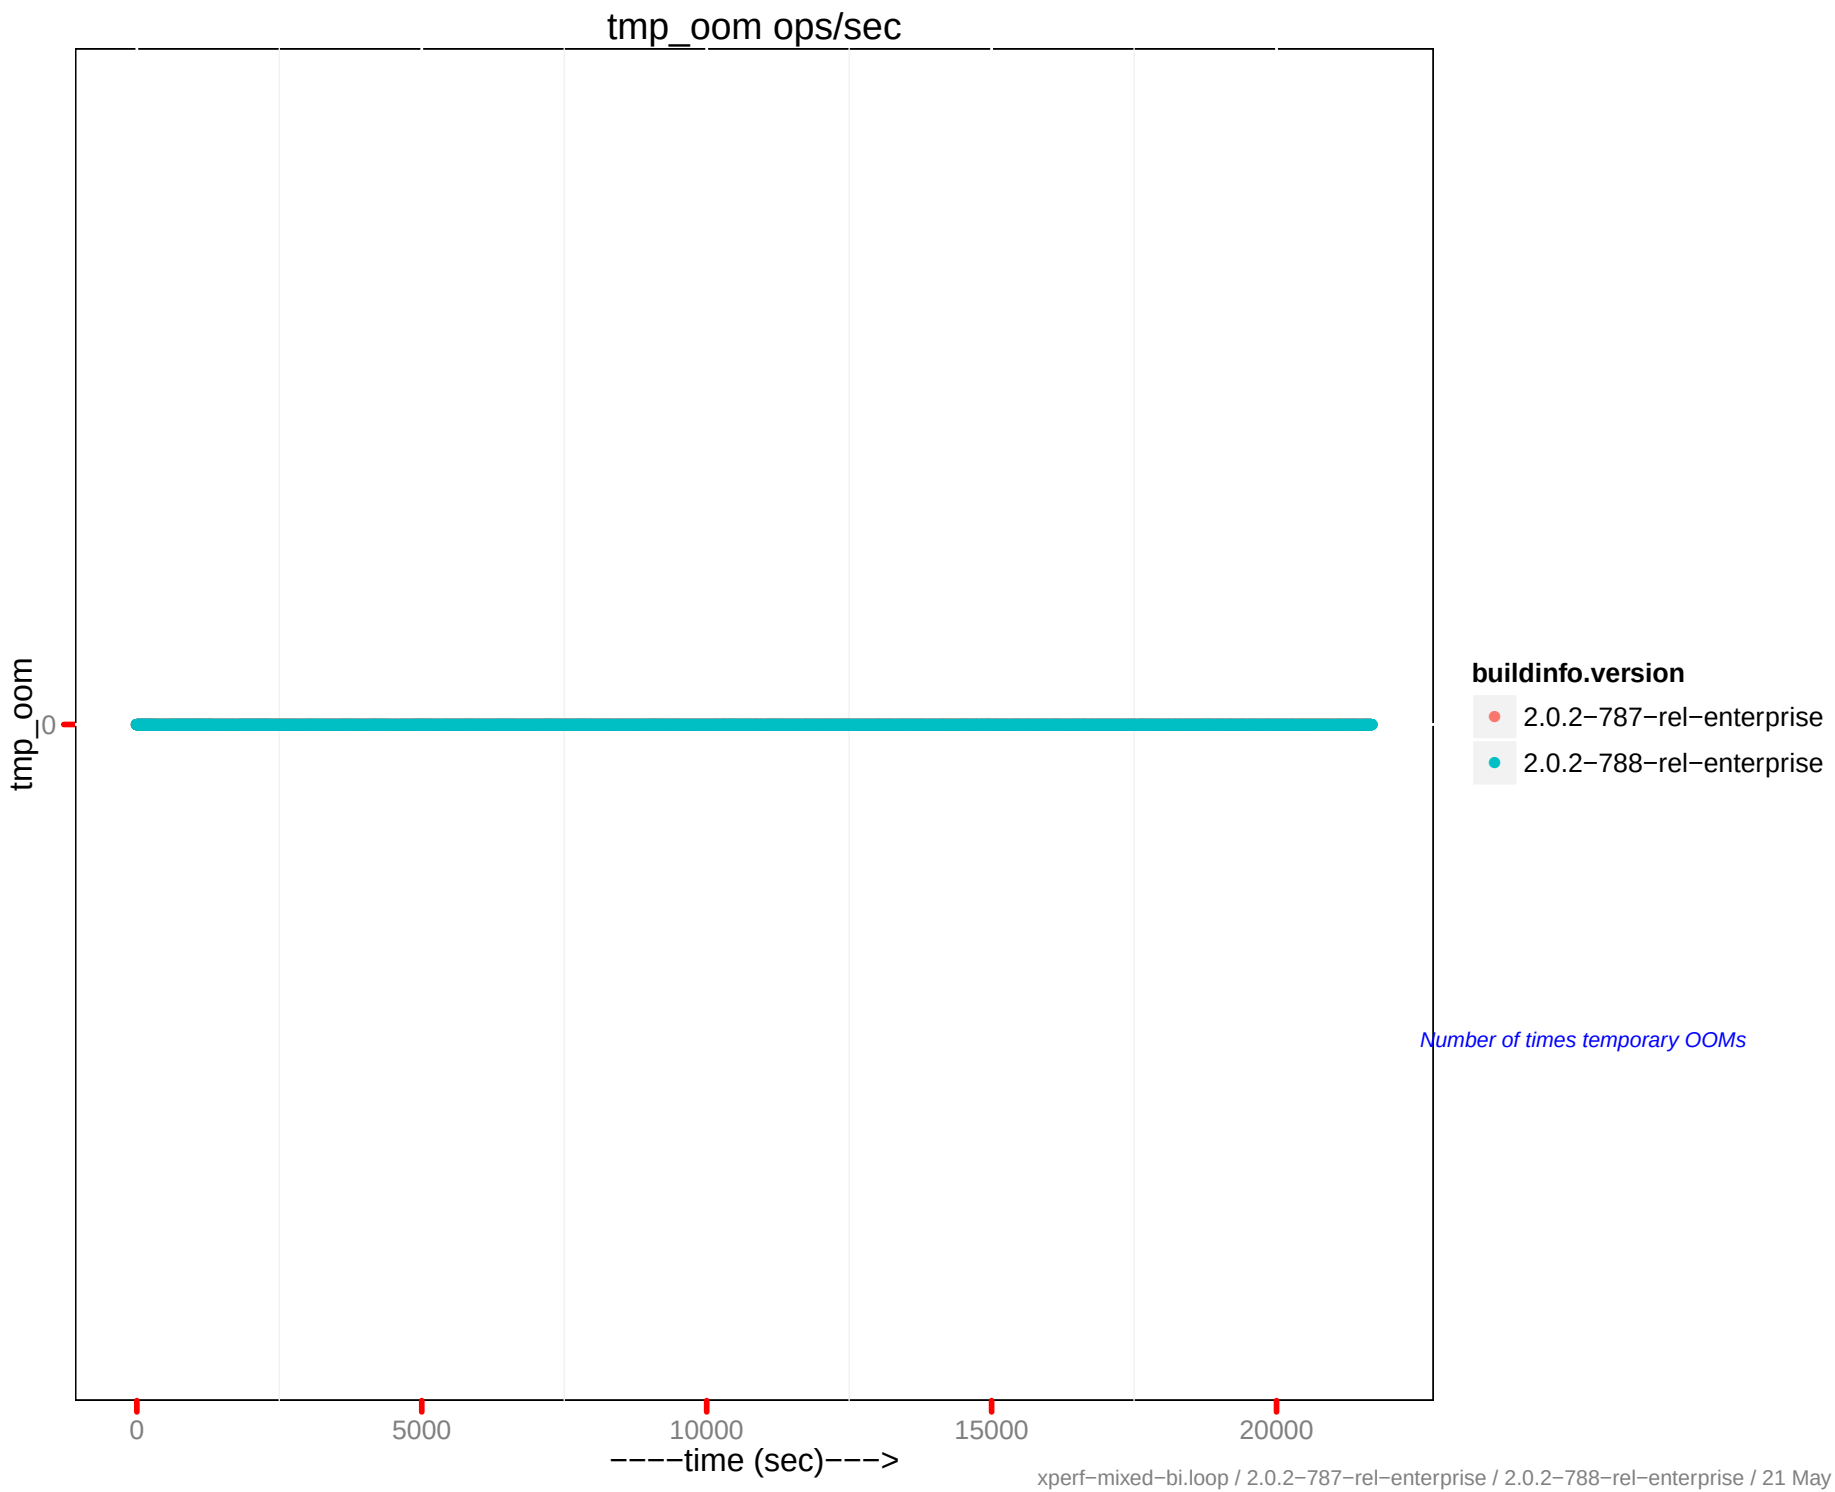
tmp_oom ops/sec
tmp_oom
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of times temporary OOMs
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

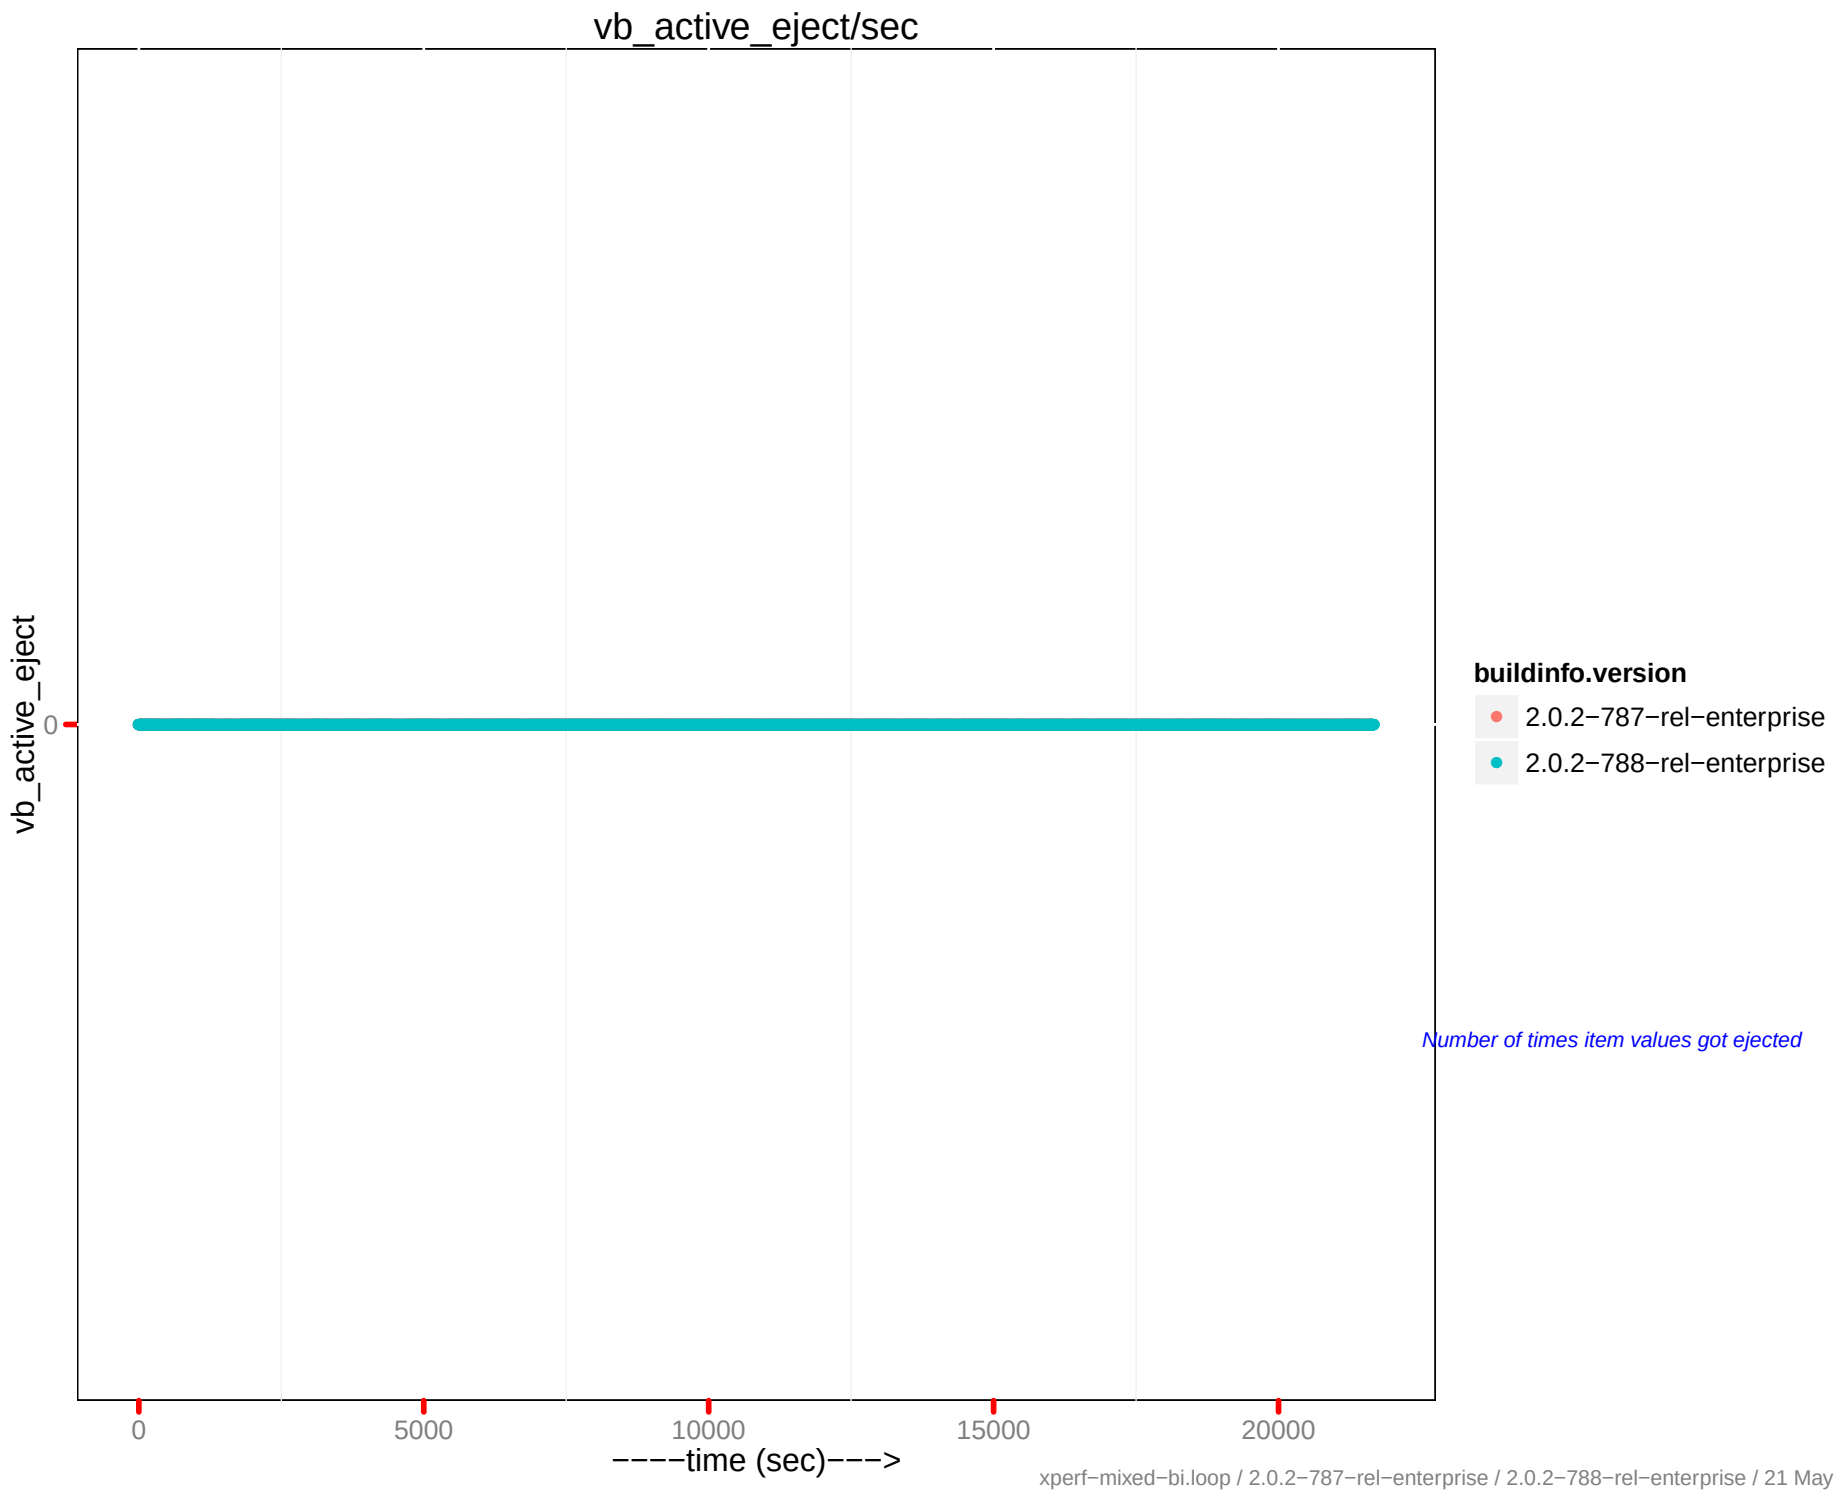
vb_active_eject/sec
vb_active_eject
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of times item values got ejected
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

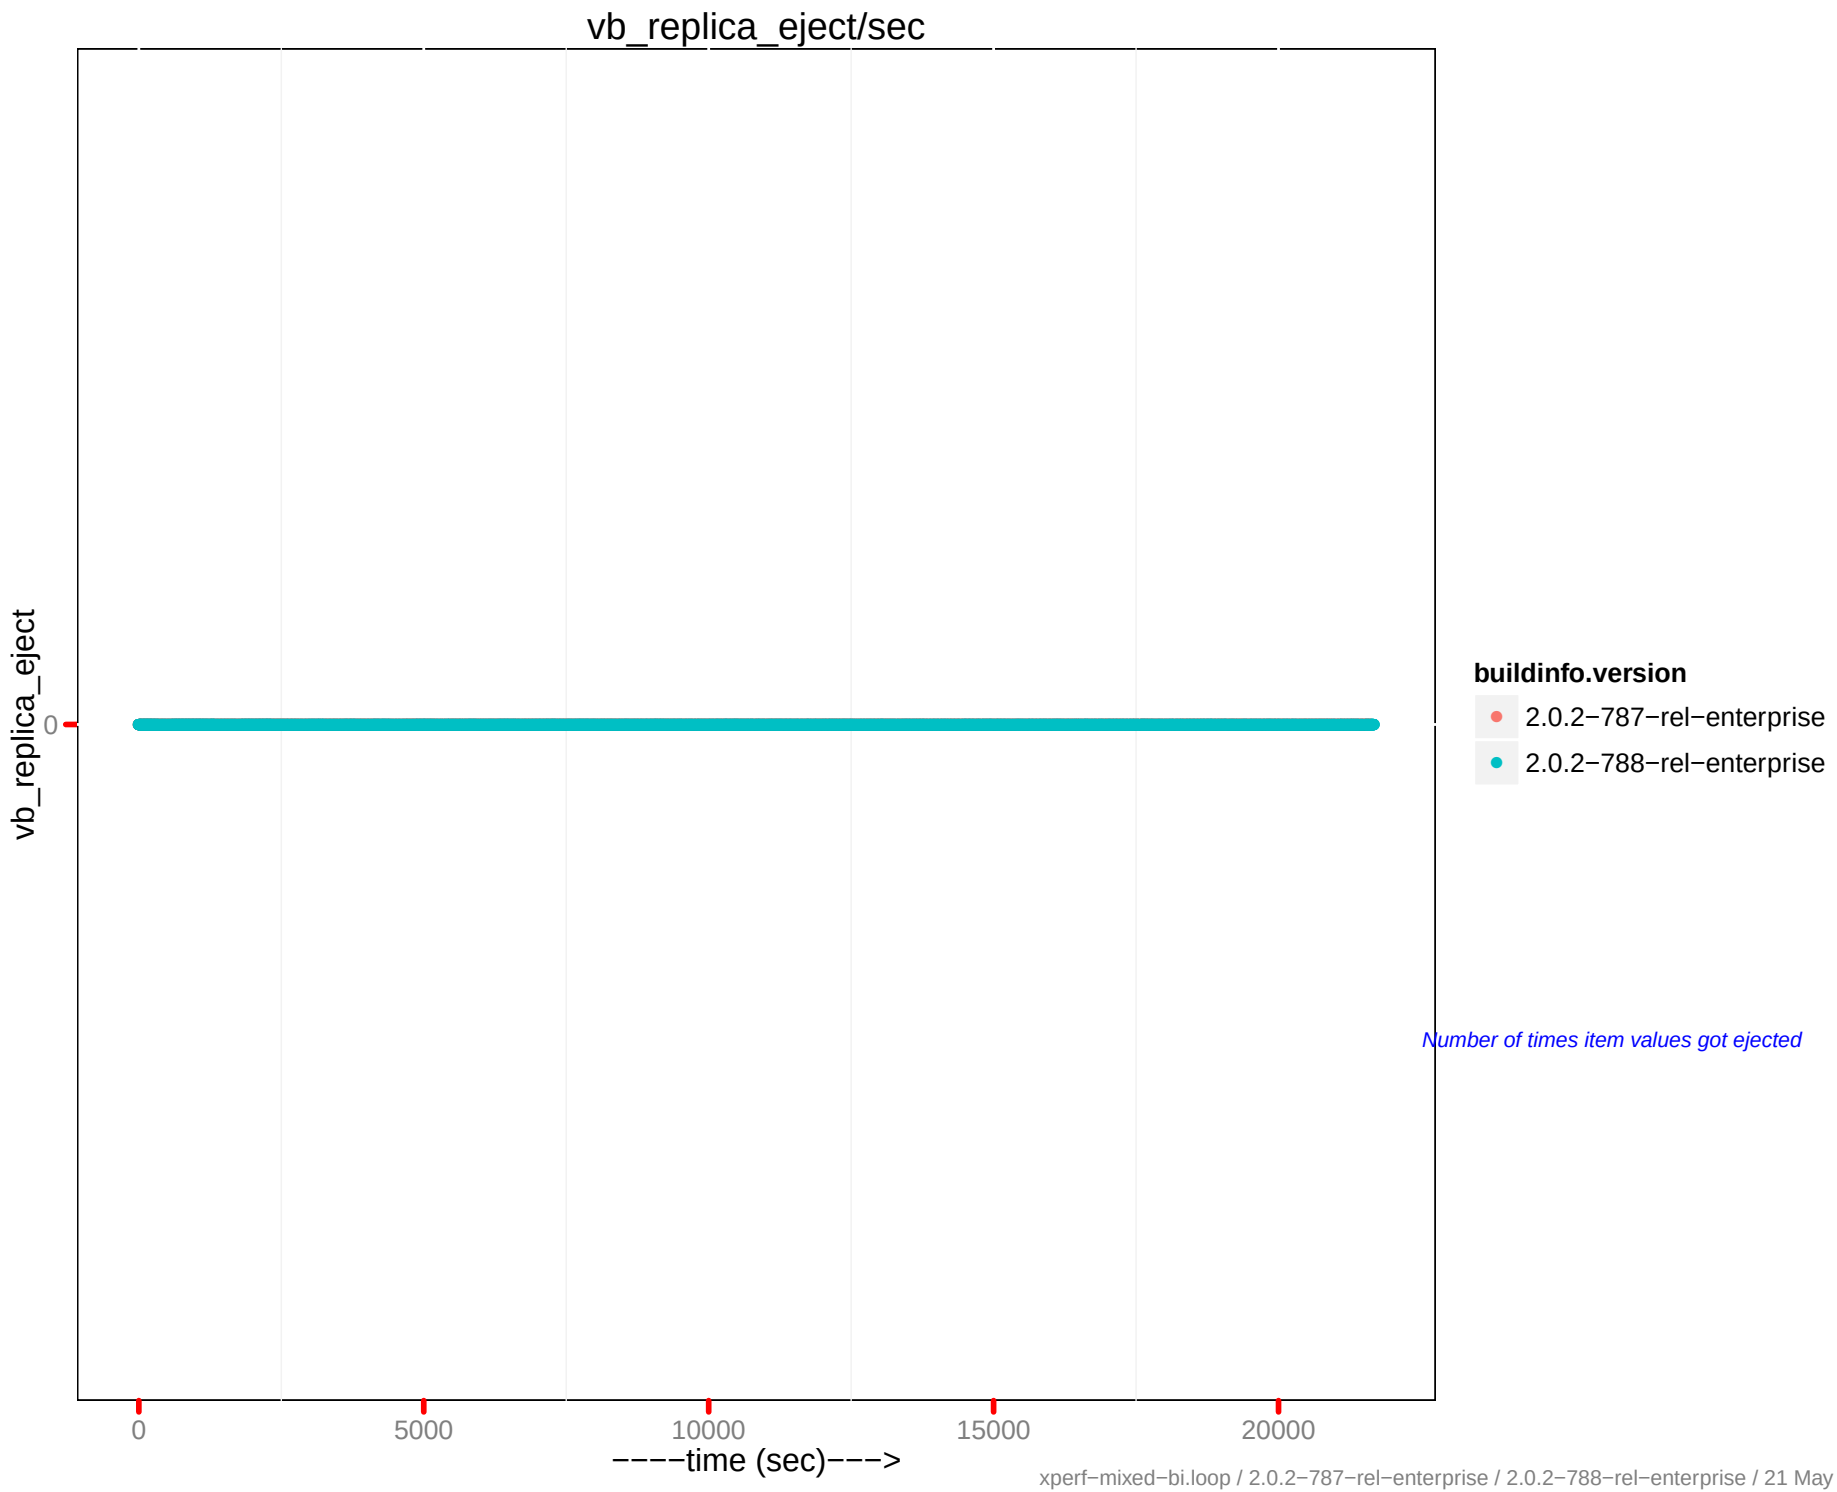
vb_replica_eject/sec
vb_replica_eject
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of times item values got ejected
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013


ep_tap_replica_queue_drain/sec
ep_tap_replica_queue_backoff
29.30K
19.53K
9.77K
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of items per second
been sent over replication
TAP connections
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

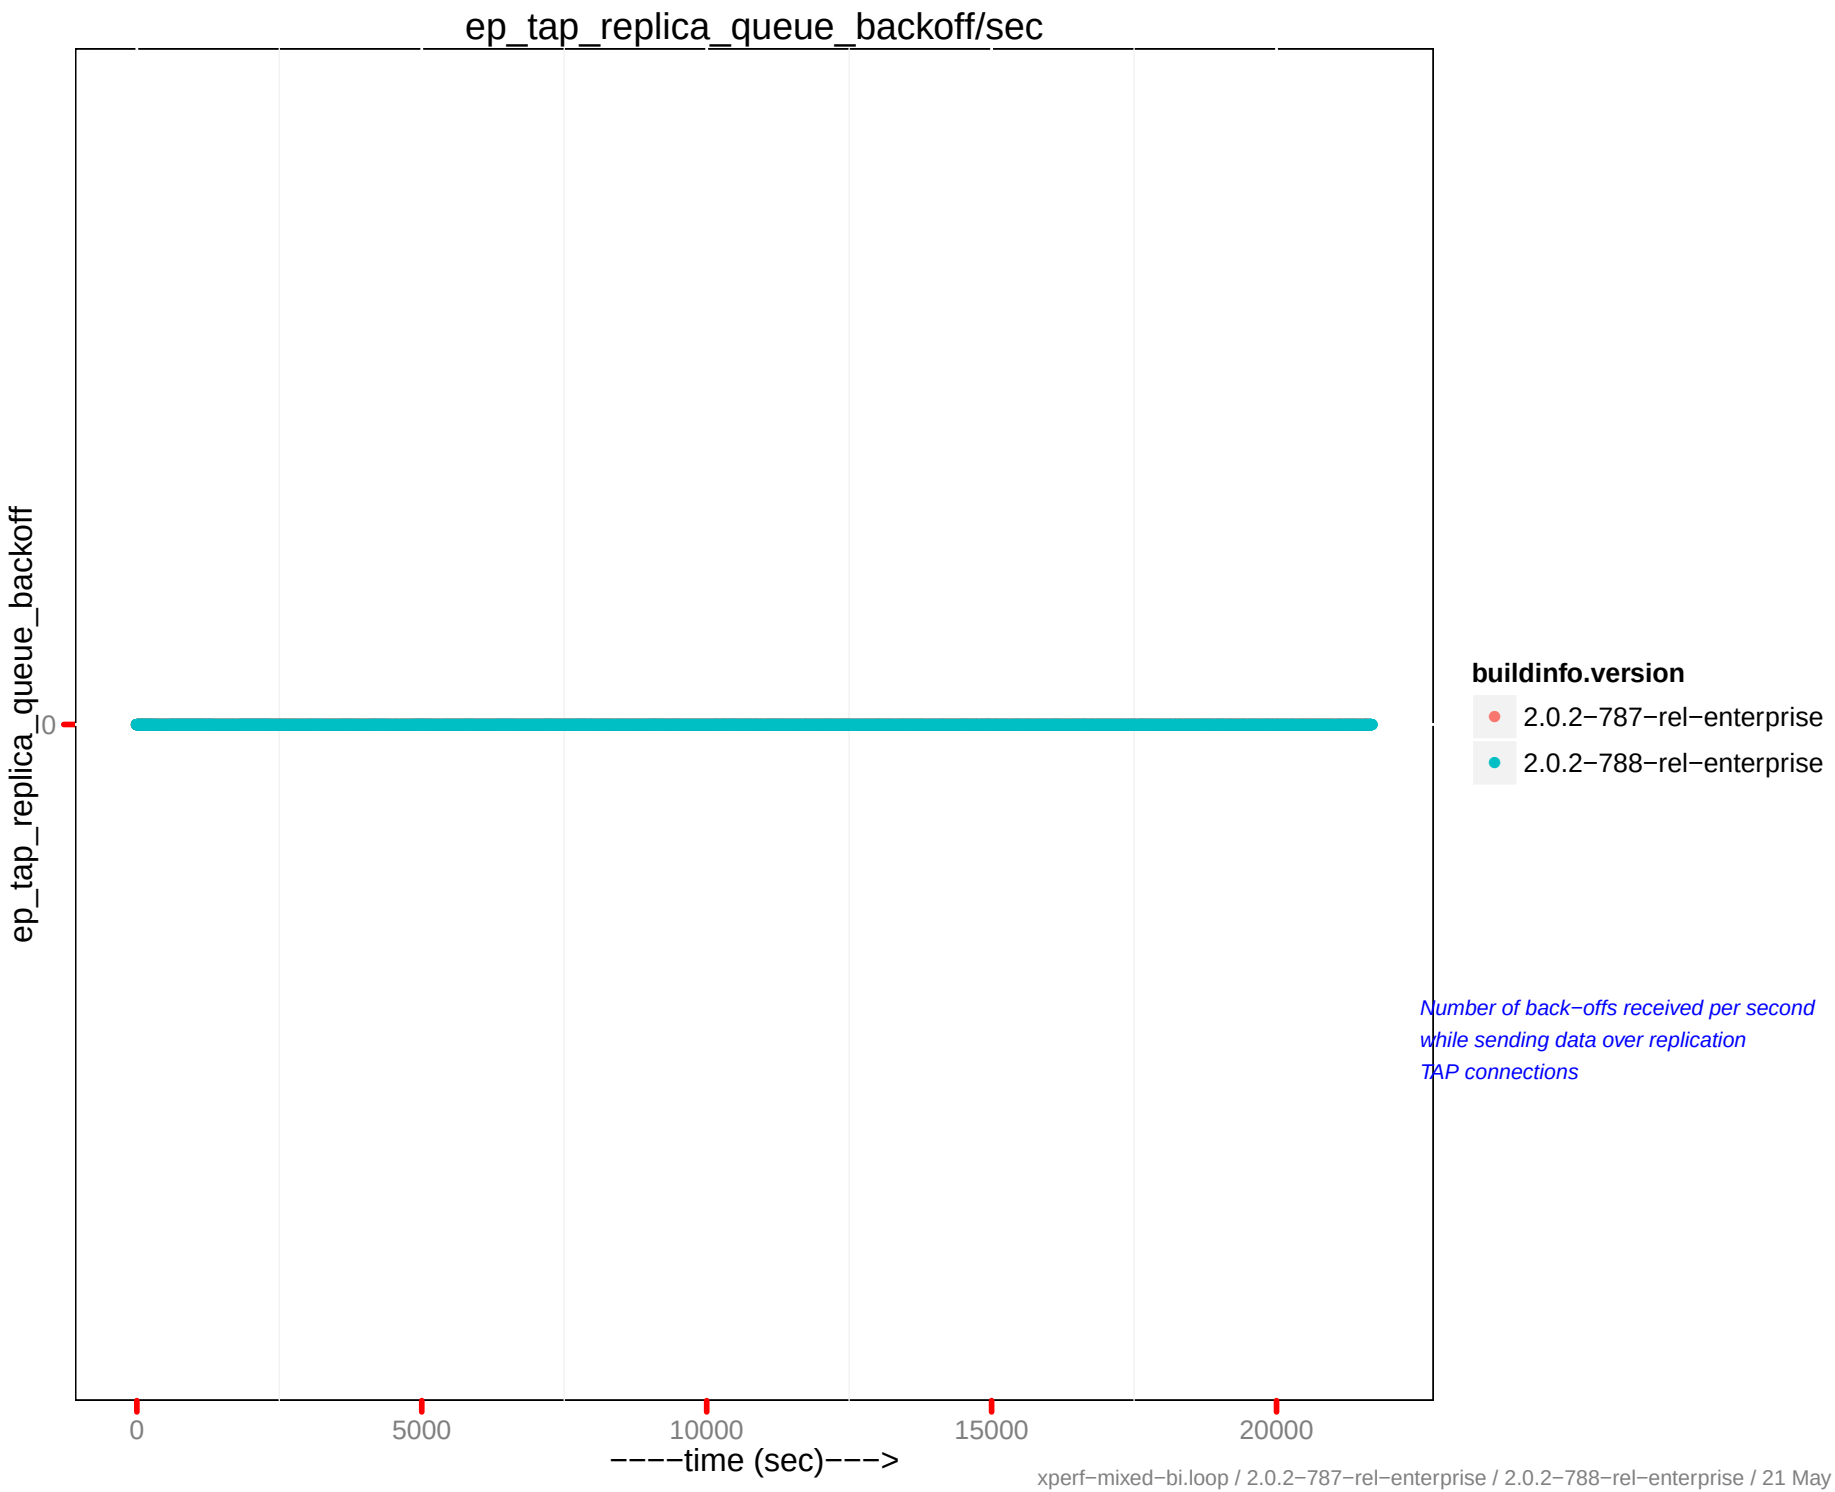
ep_tap_replica_queue_backoff/sec
ep_tap_replica_queue_backoff
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of back-offs received per second
while sending data over replication
TAP connections
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

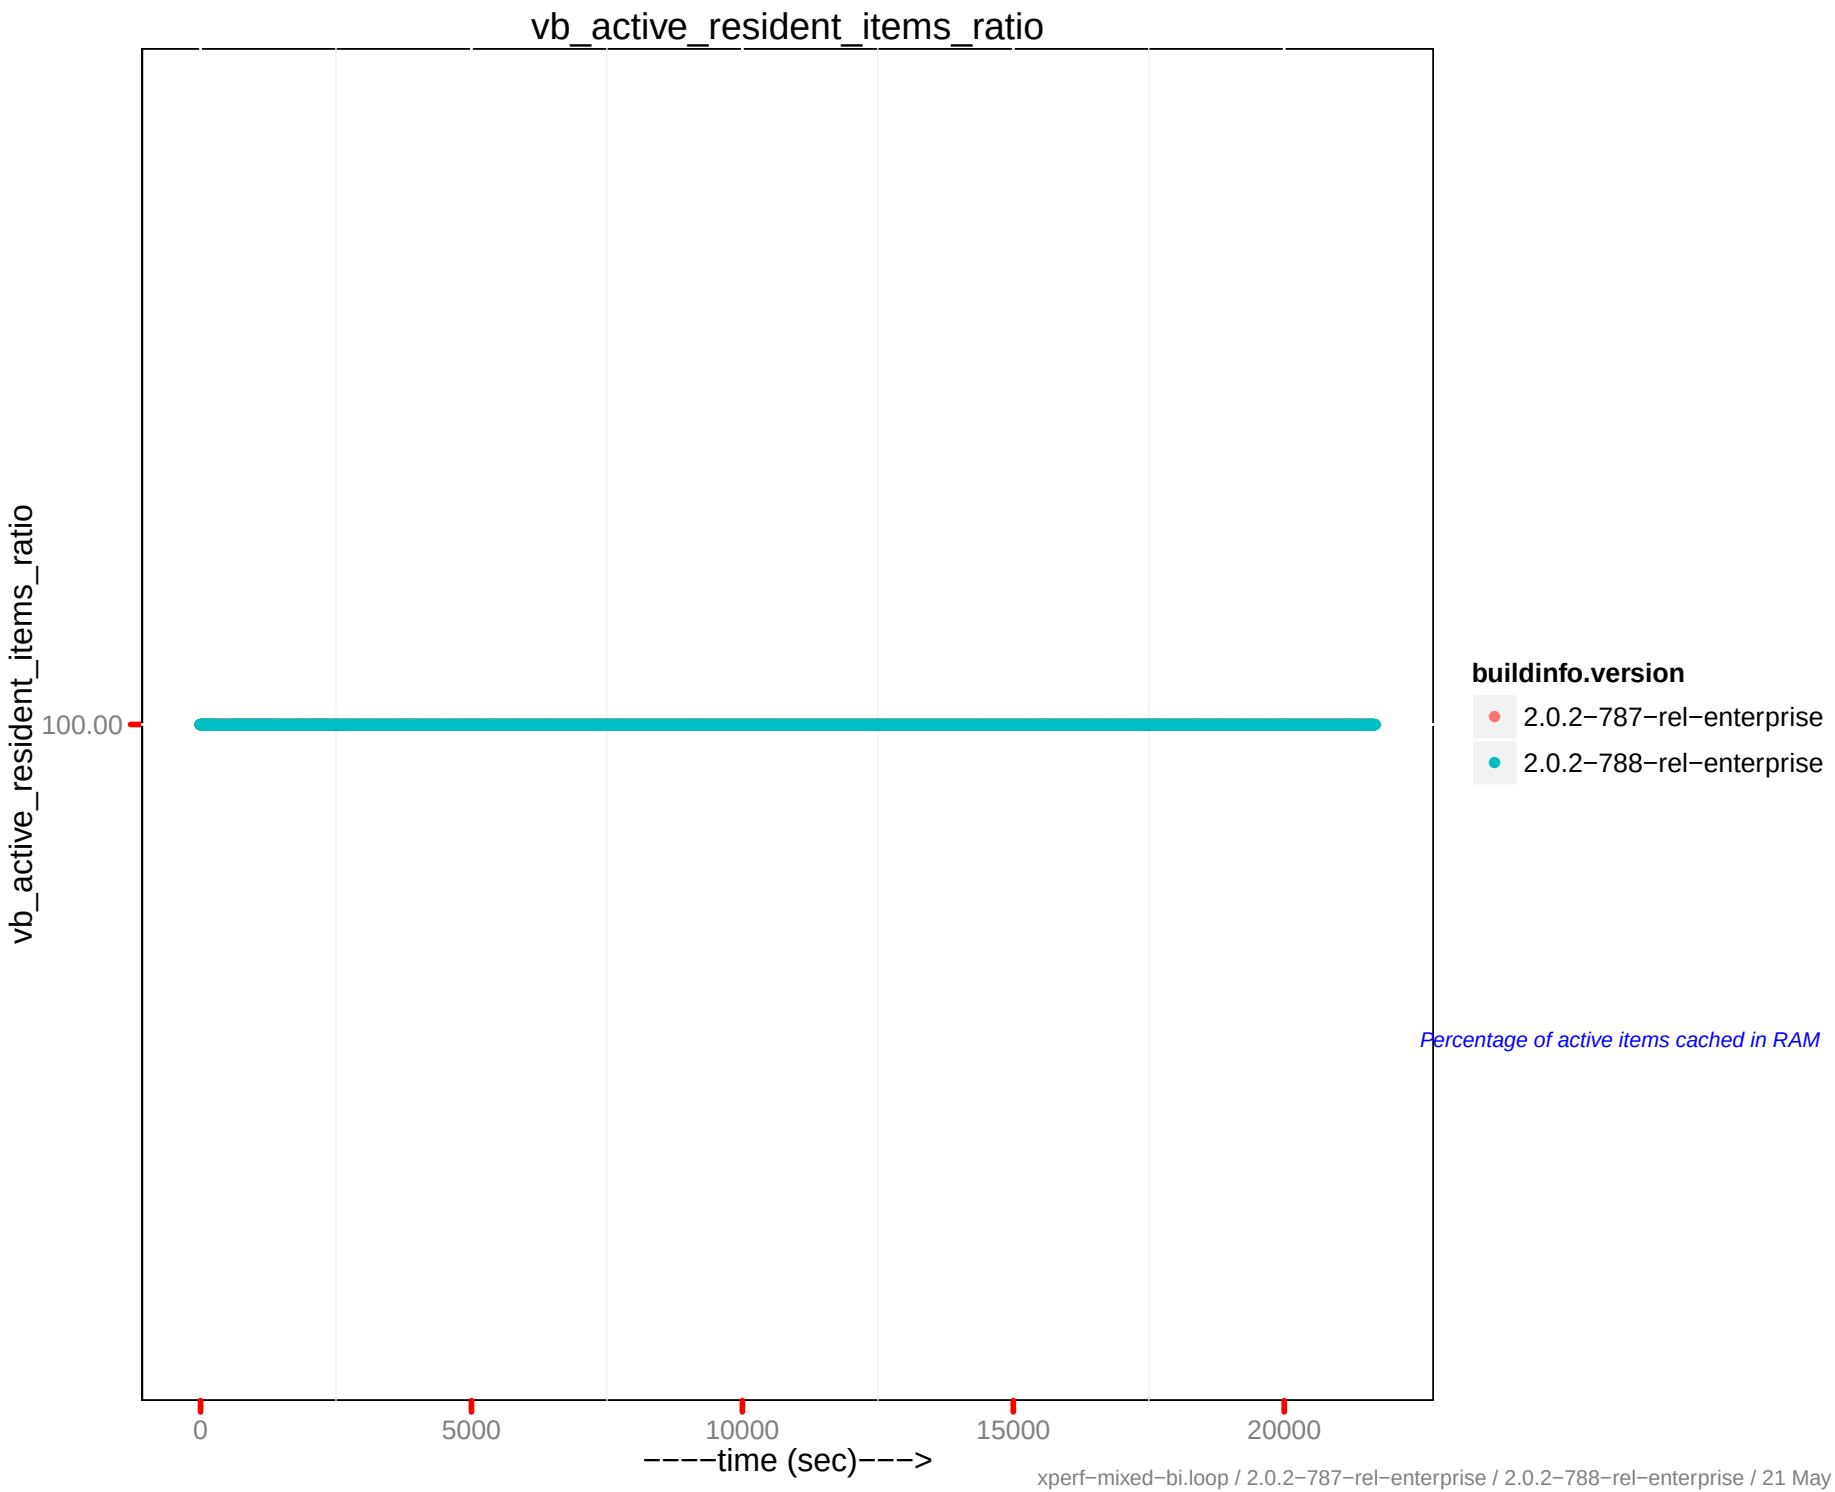
vb_active_resident_items_ratio
vb_active_resident_items_ratio
100.00
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Percentage of active items cached in RAM
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

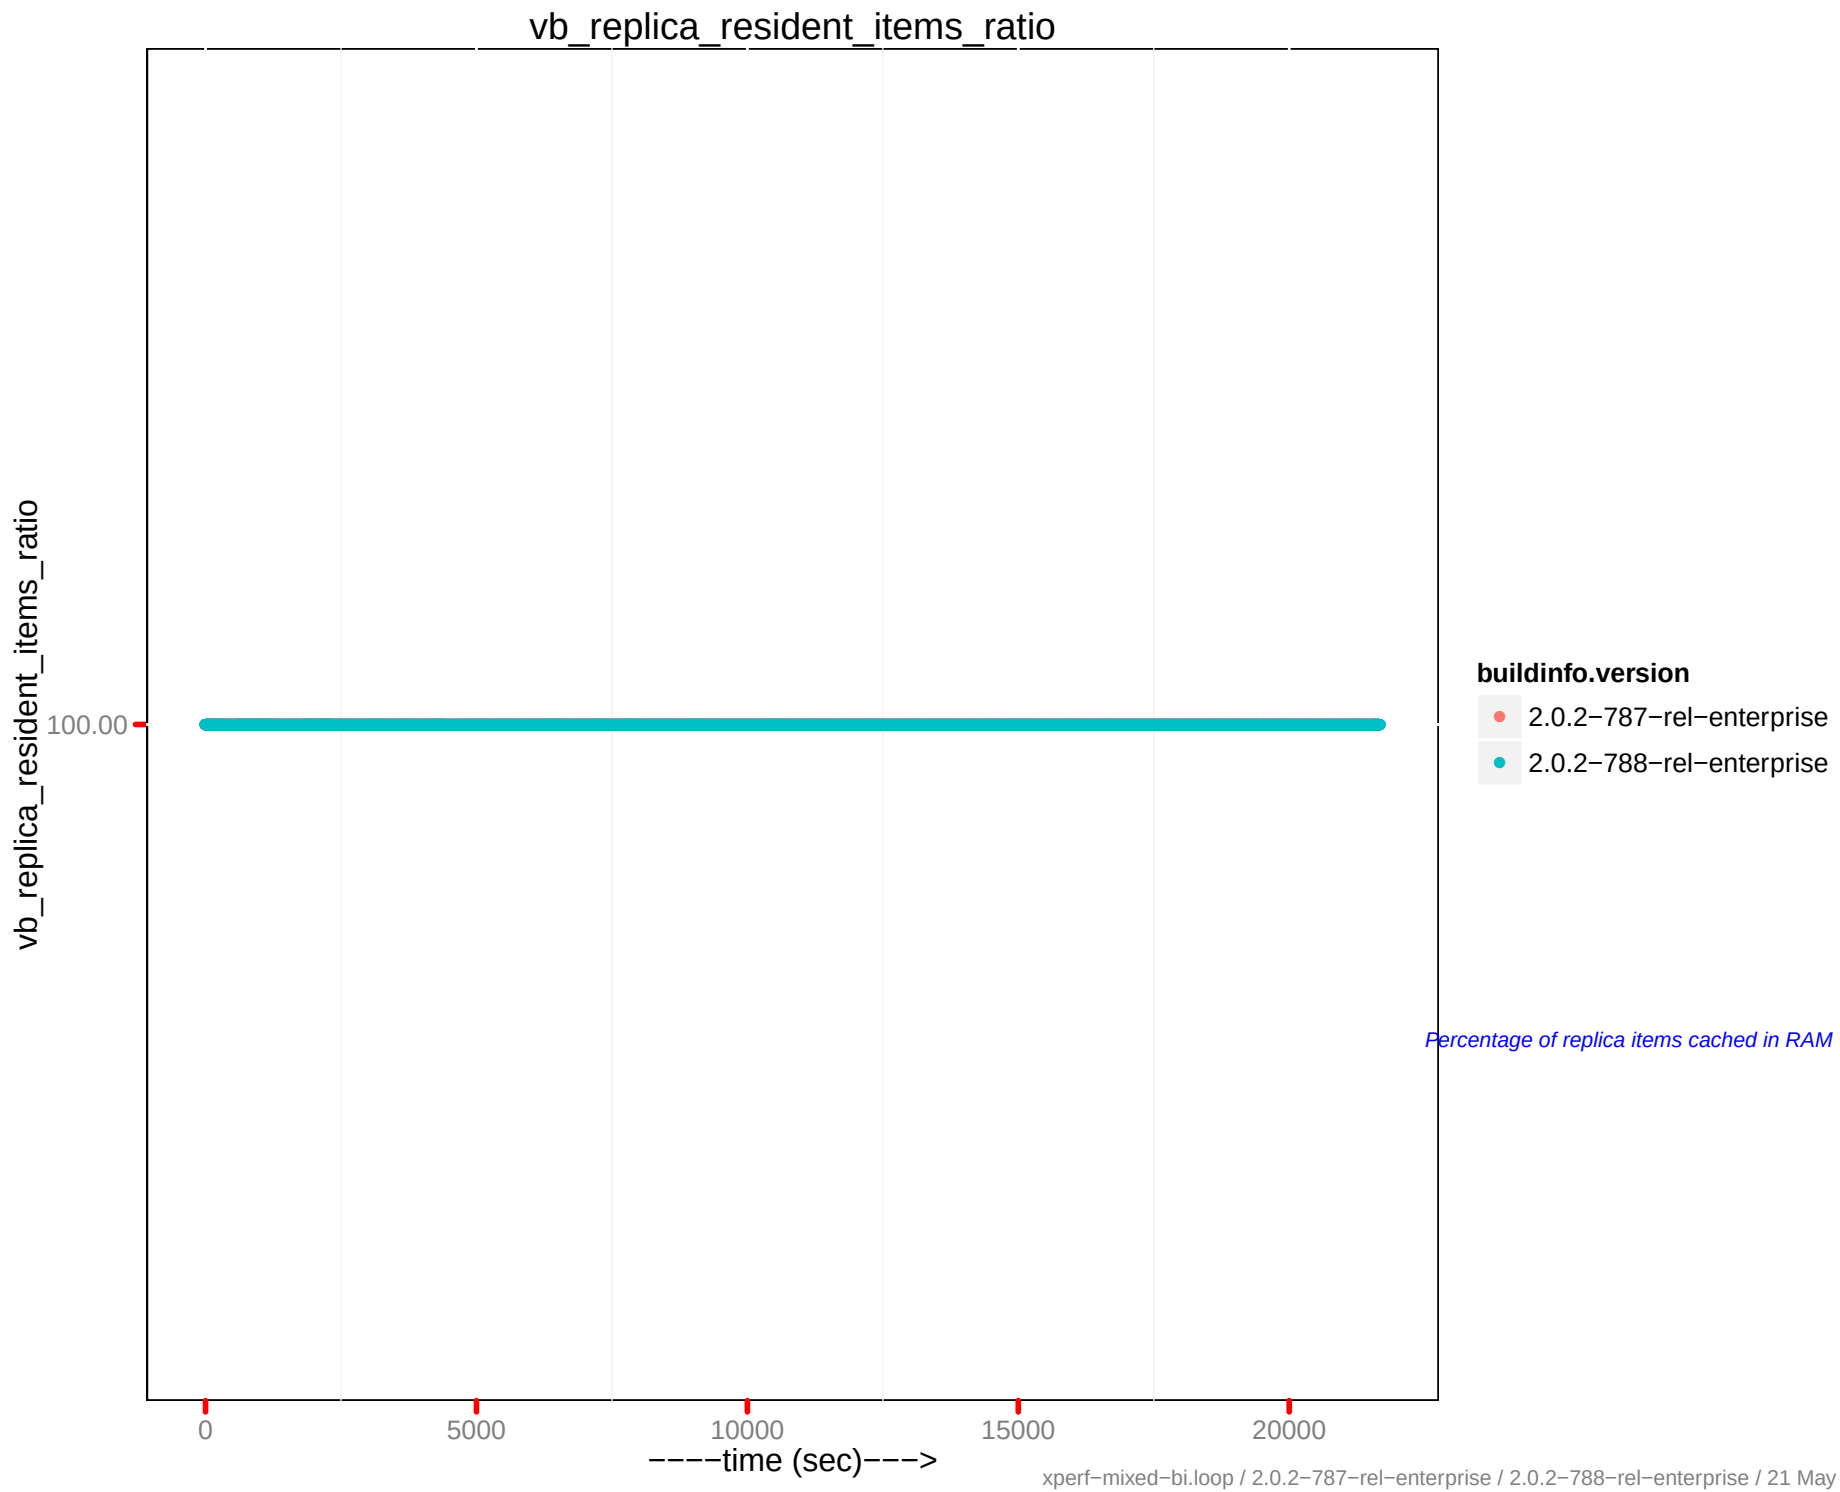
vb_replica_resident_items_ratio
vb_replica_resident_items_ratio
100.00
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Percentage of replica items cached in RAM
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

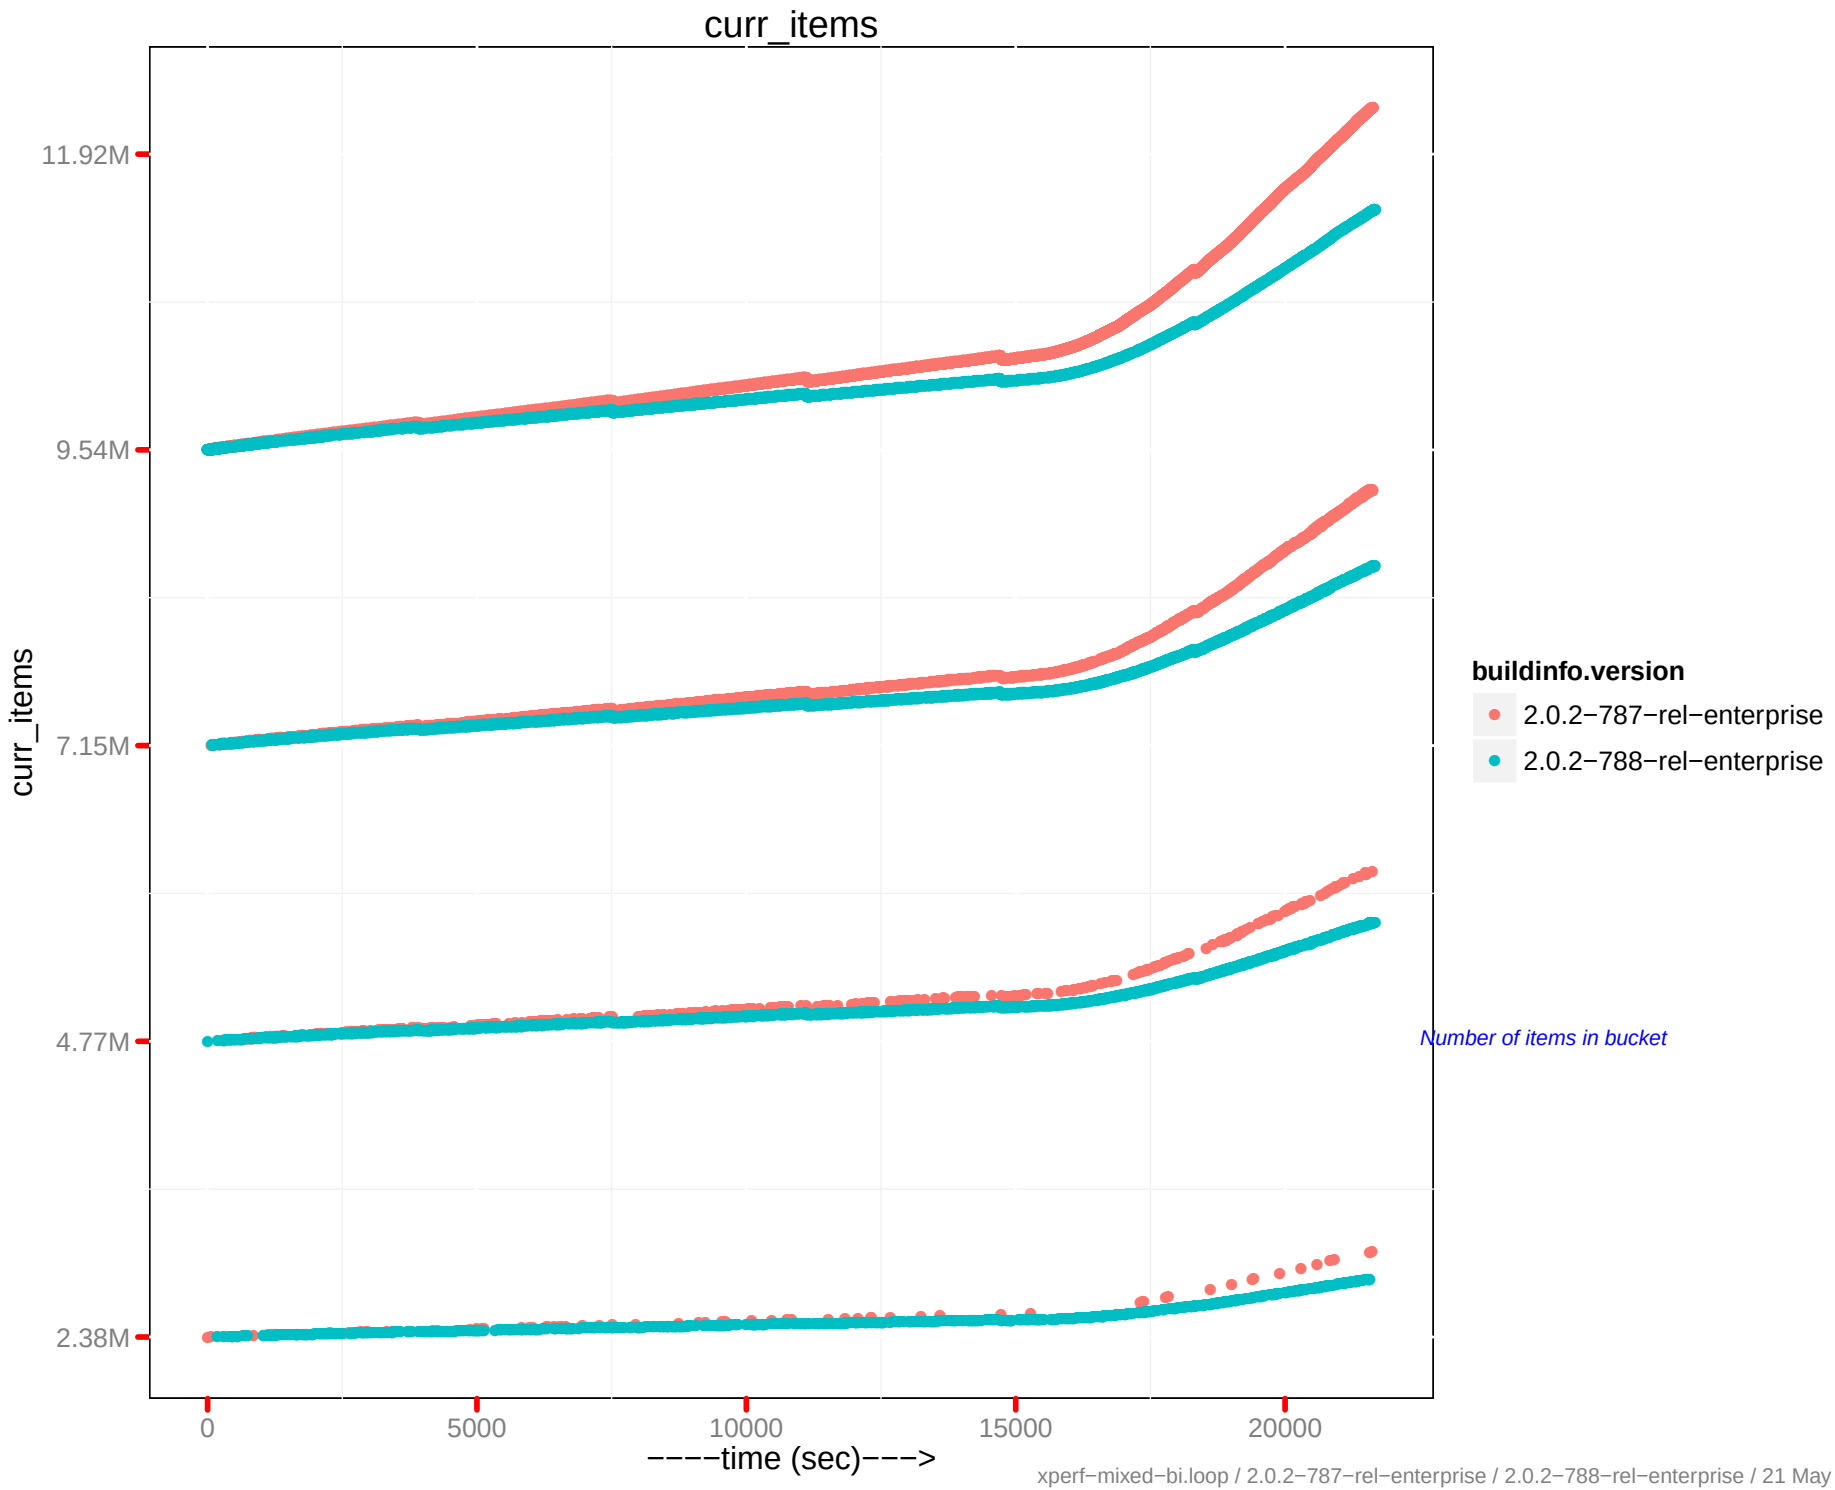
curr_items
11.92M
9.54M
7.15M
4.77M
2.38M
curr_items
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of items in bucket
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

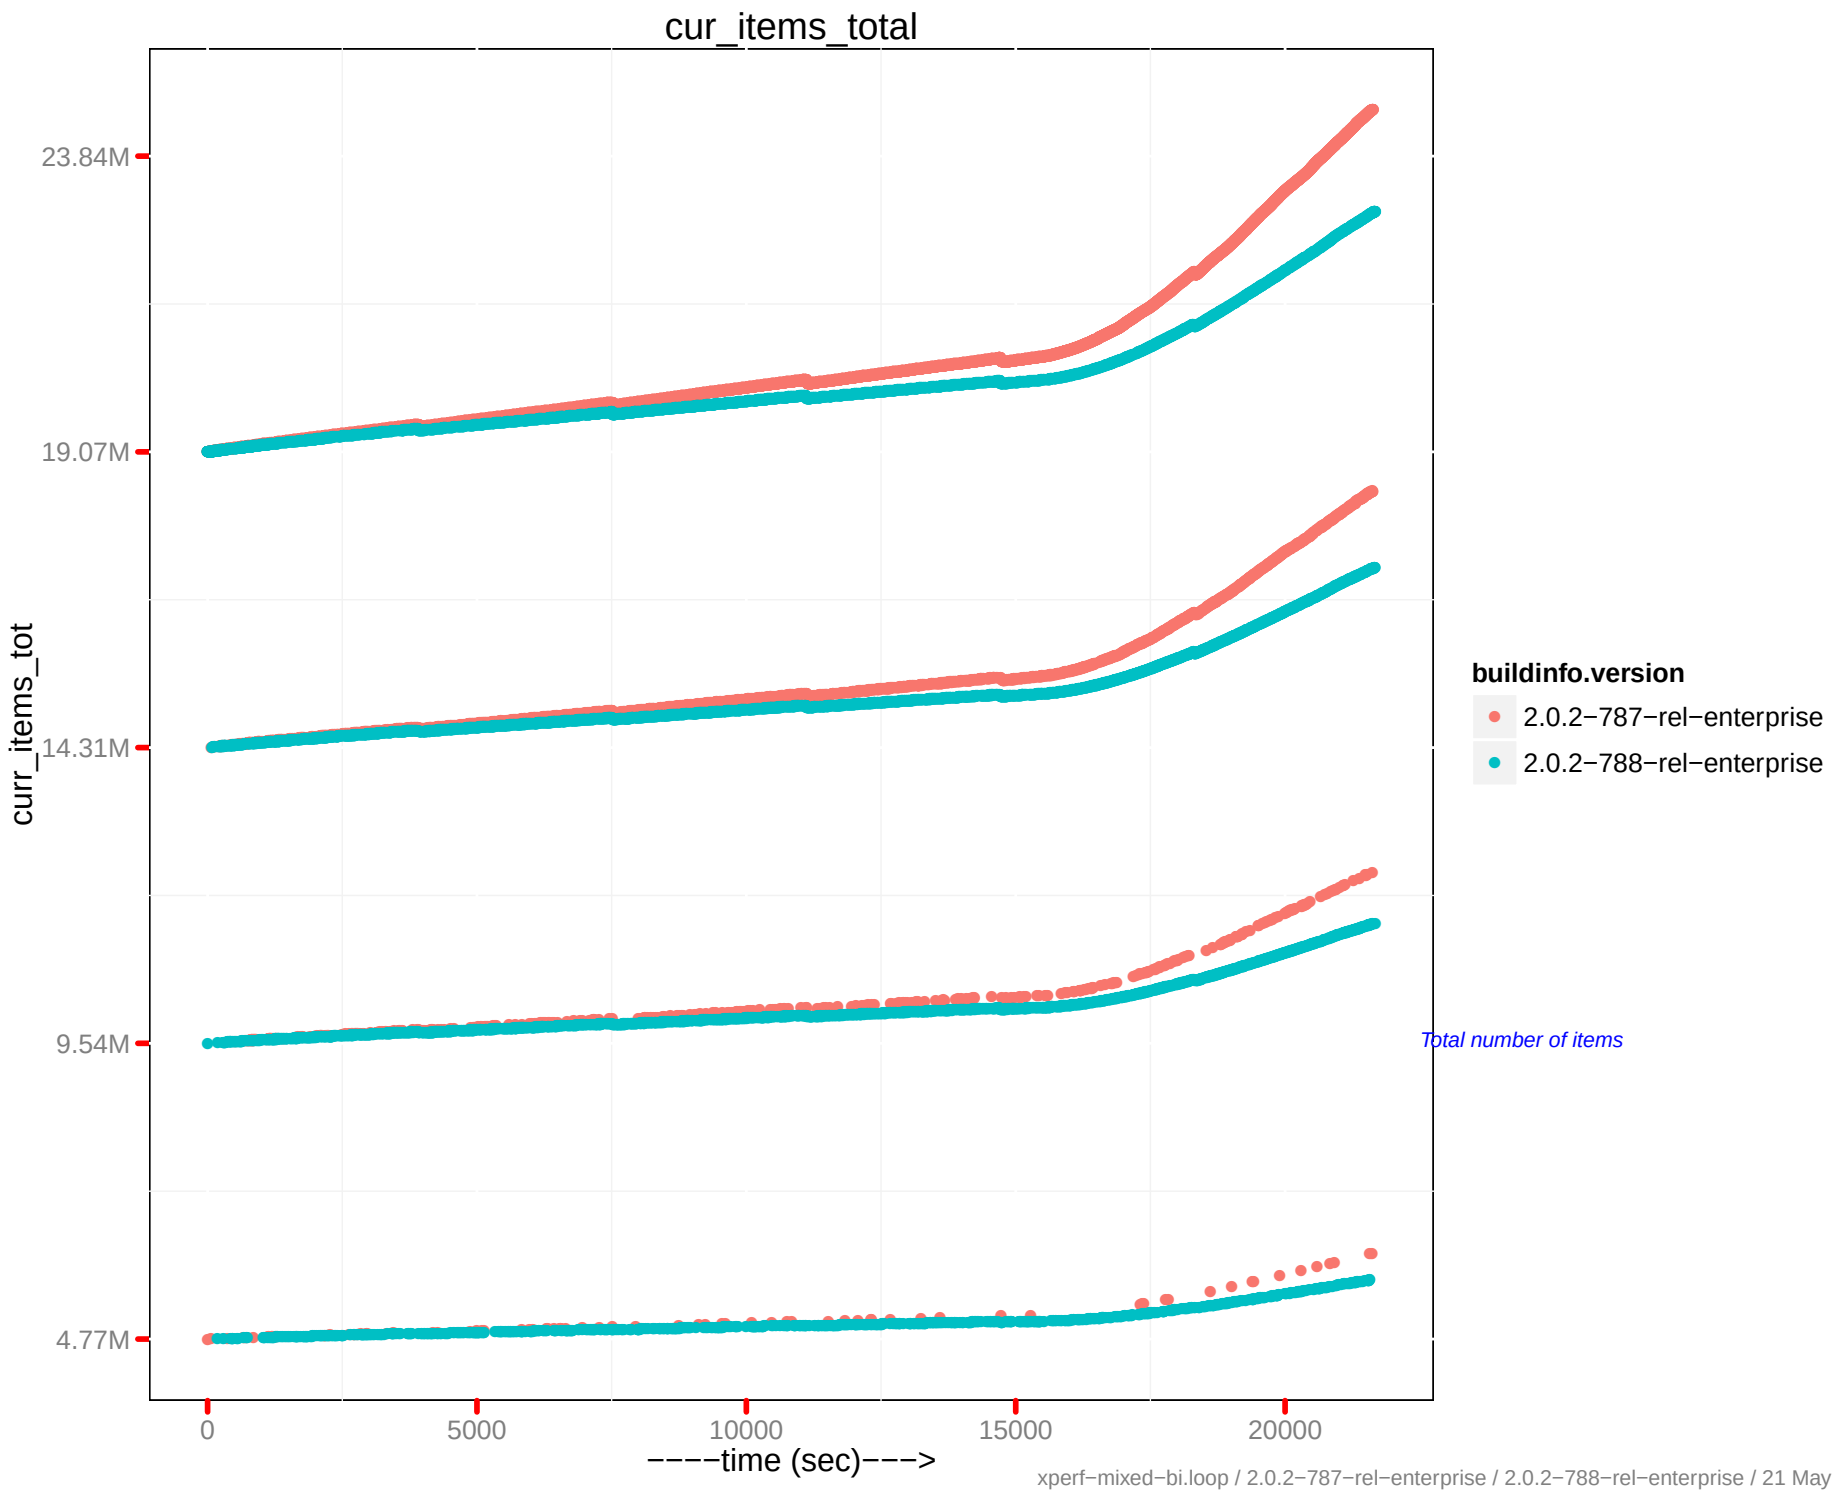
cur_items_total
23.84M
19.07M
14.31M
9.54M
4.77M
curr_items_tot
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Total number of items
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

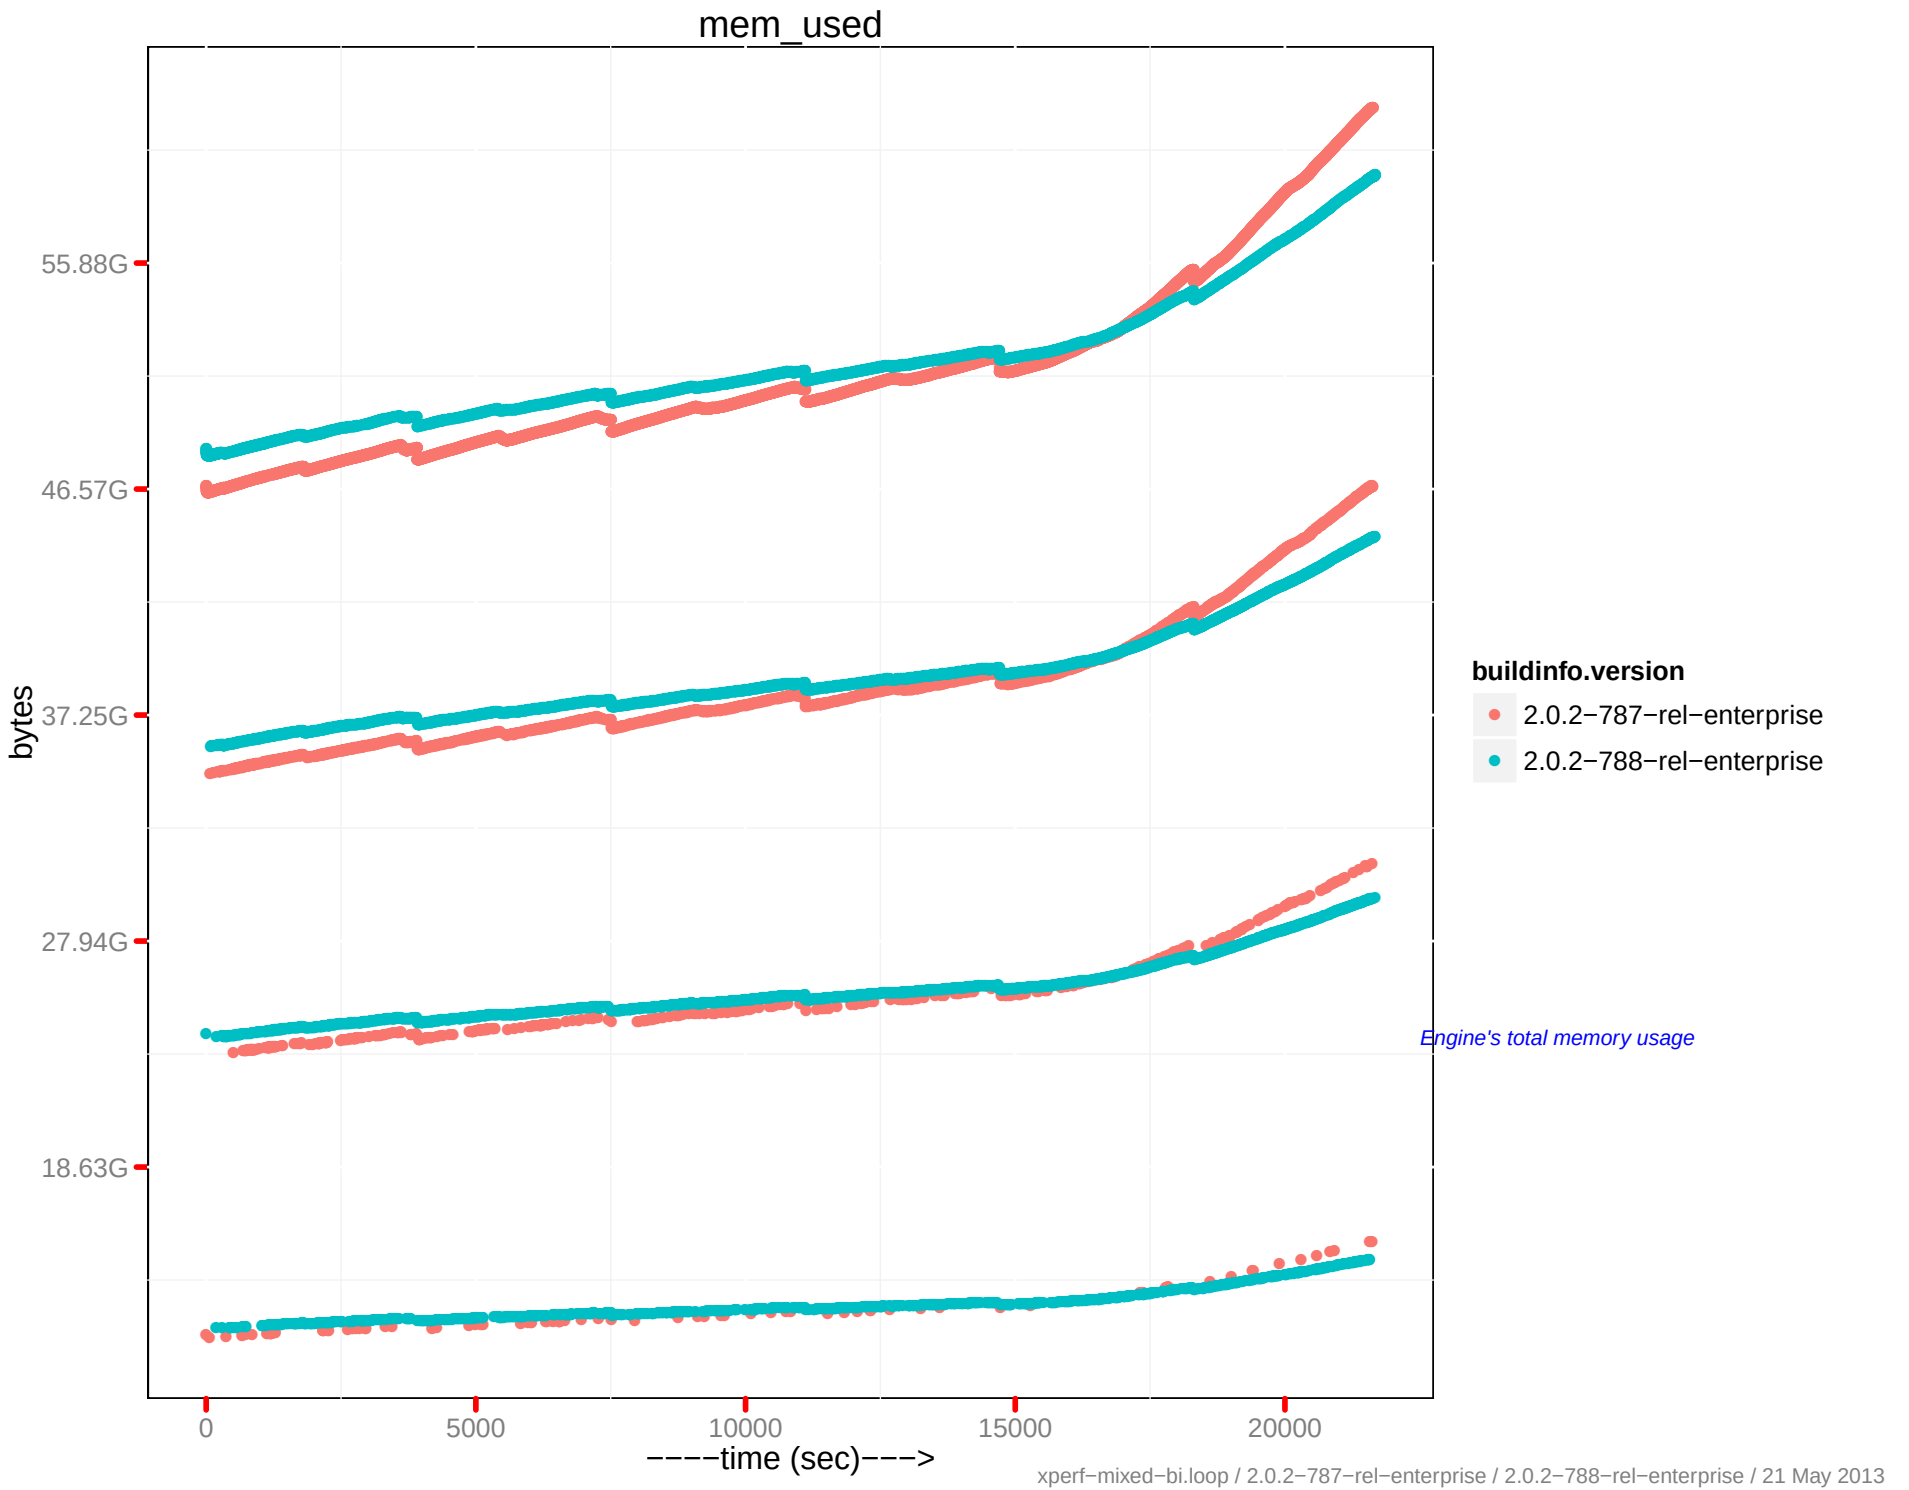
mem_used
55.88G
46.57G
37.25G
27.94G
18.63G
bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Engine's total memory usage
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

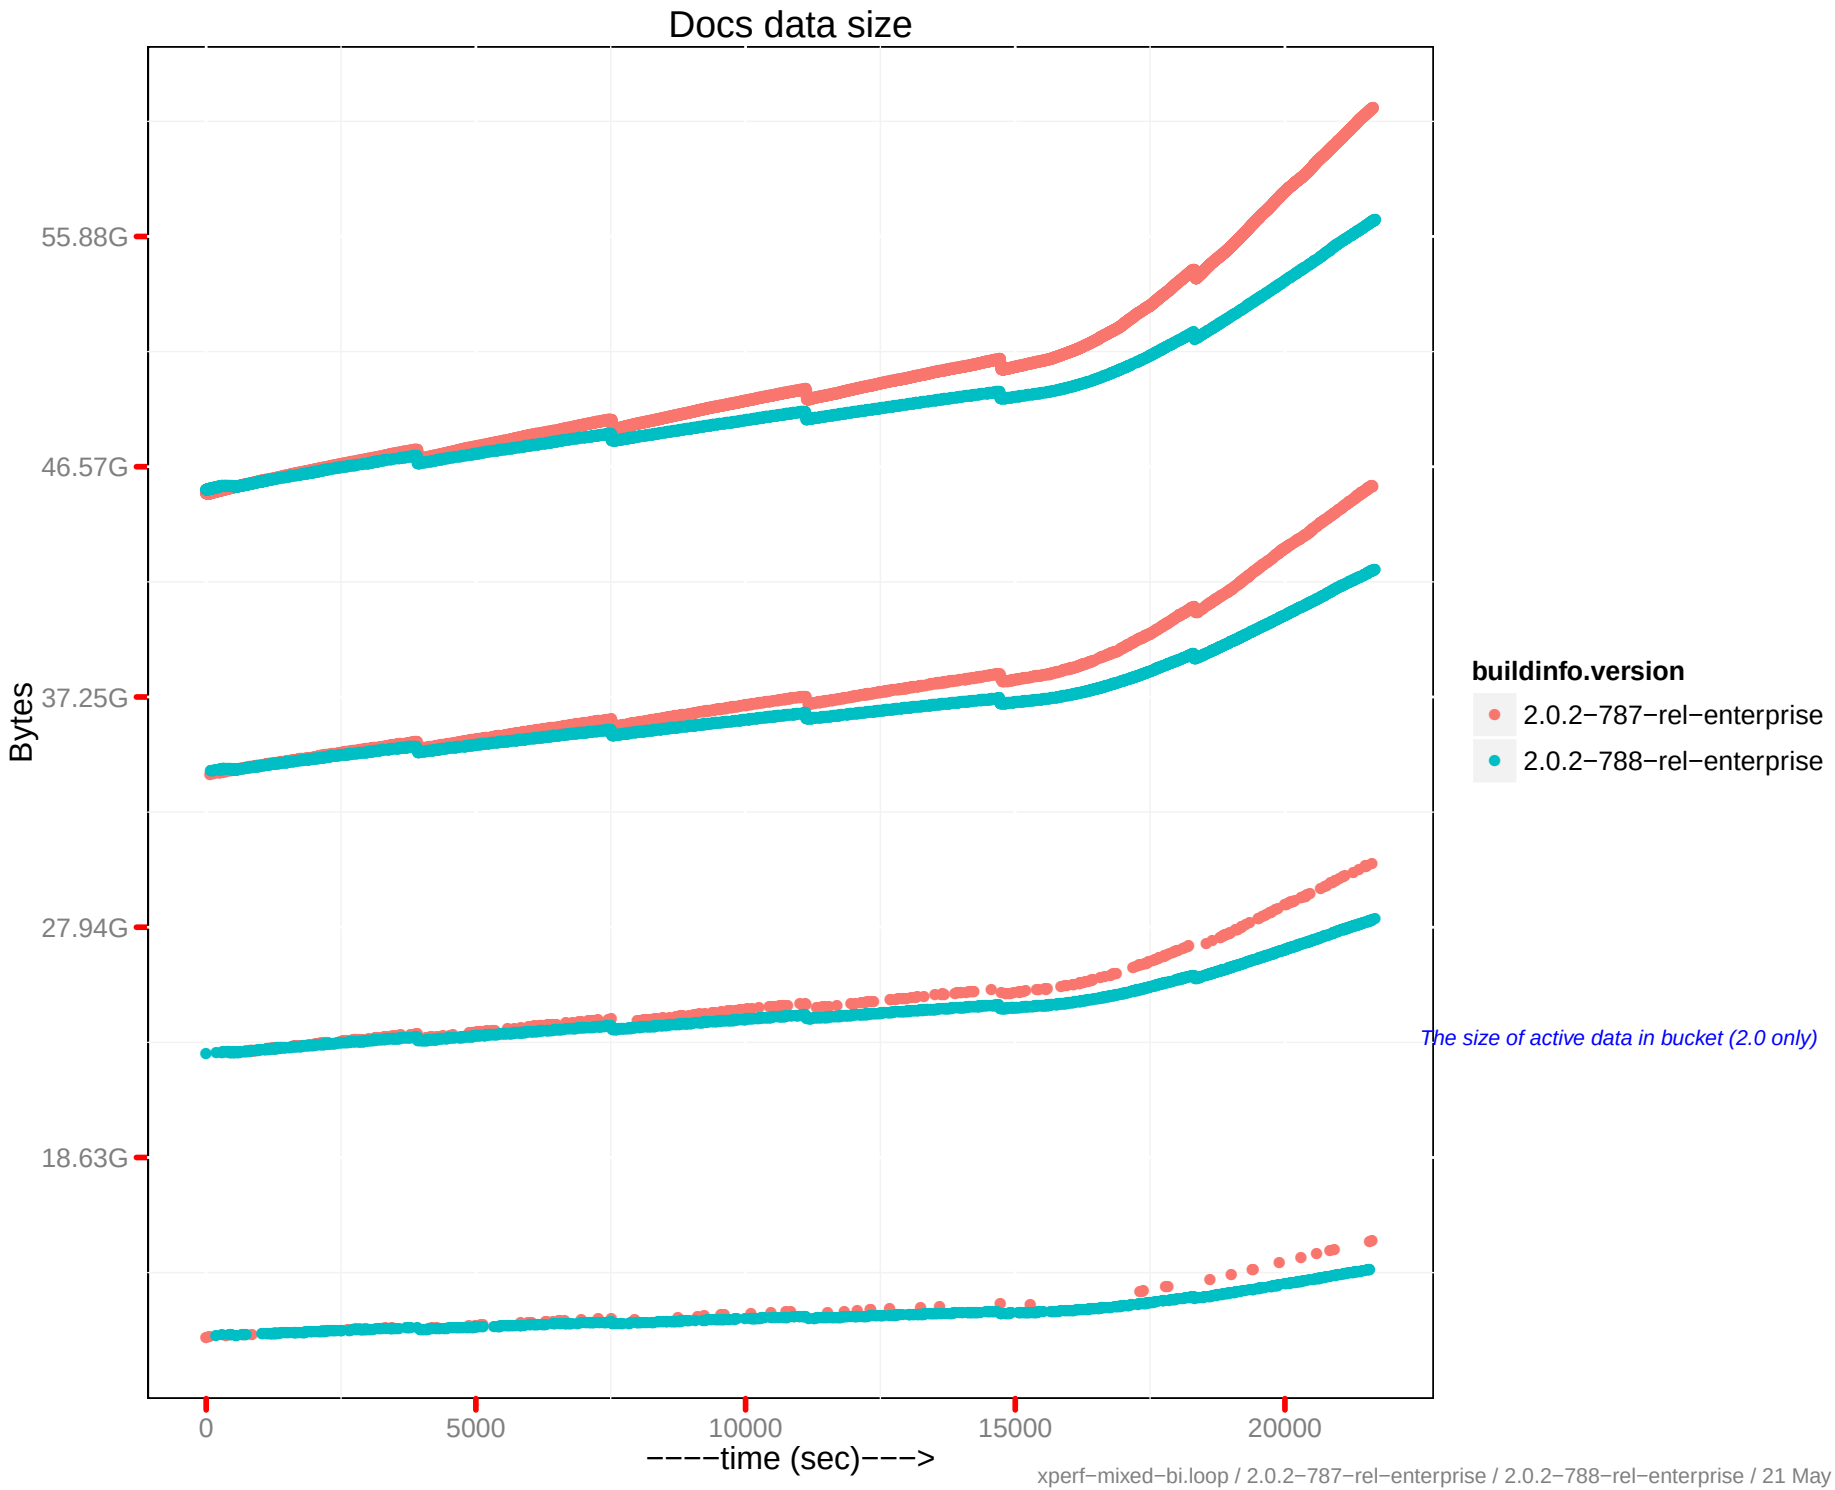
Docs data size
55.88G
46.57G
37.25G
27.94G
18.63G
Bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The size of active data in bucket (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

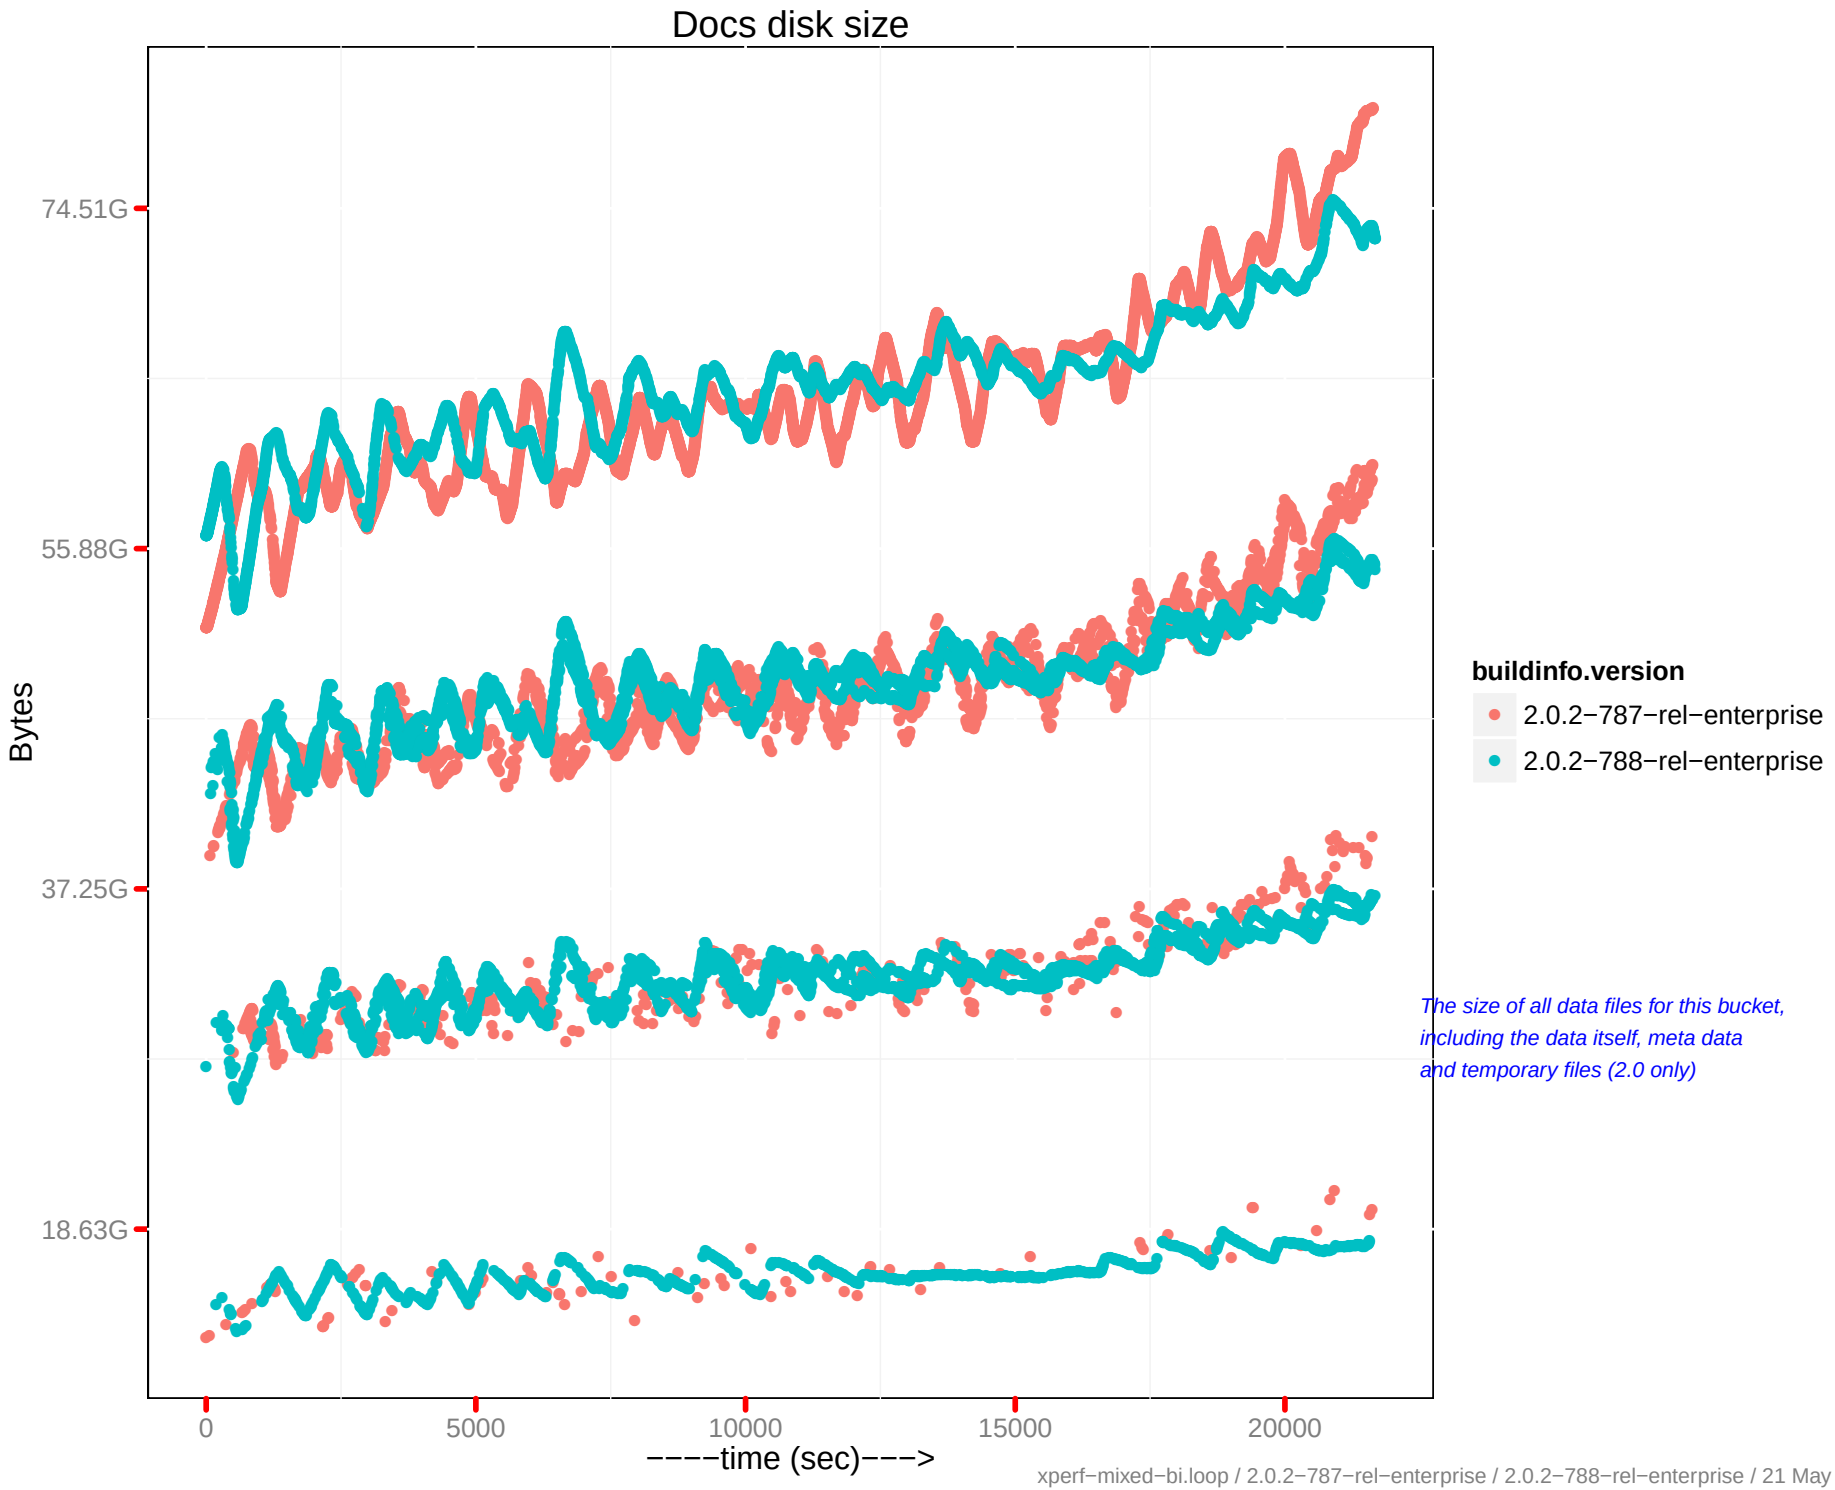
Docs disk size
74.51G
55.88G
37.25G
18.63G
Bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

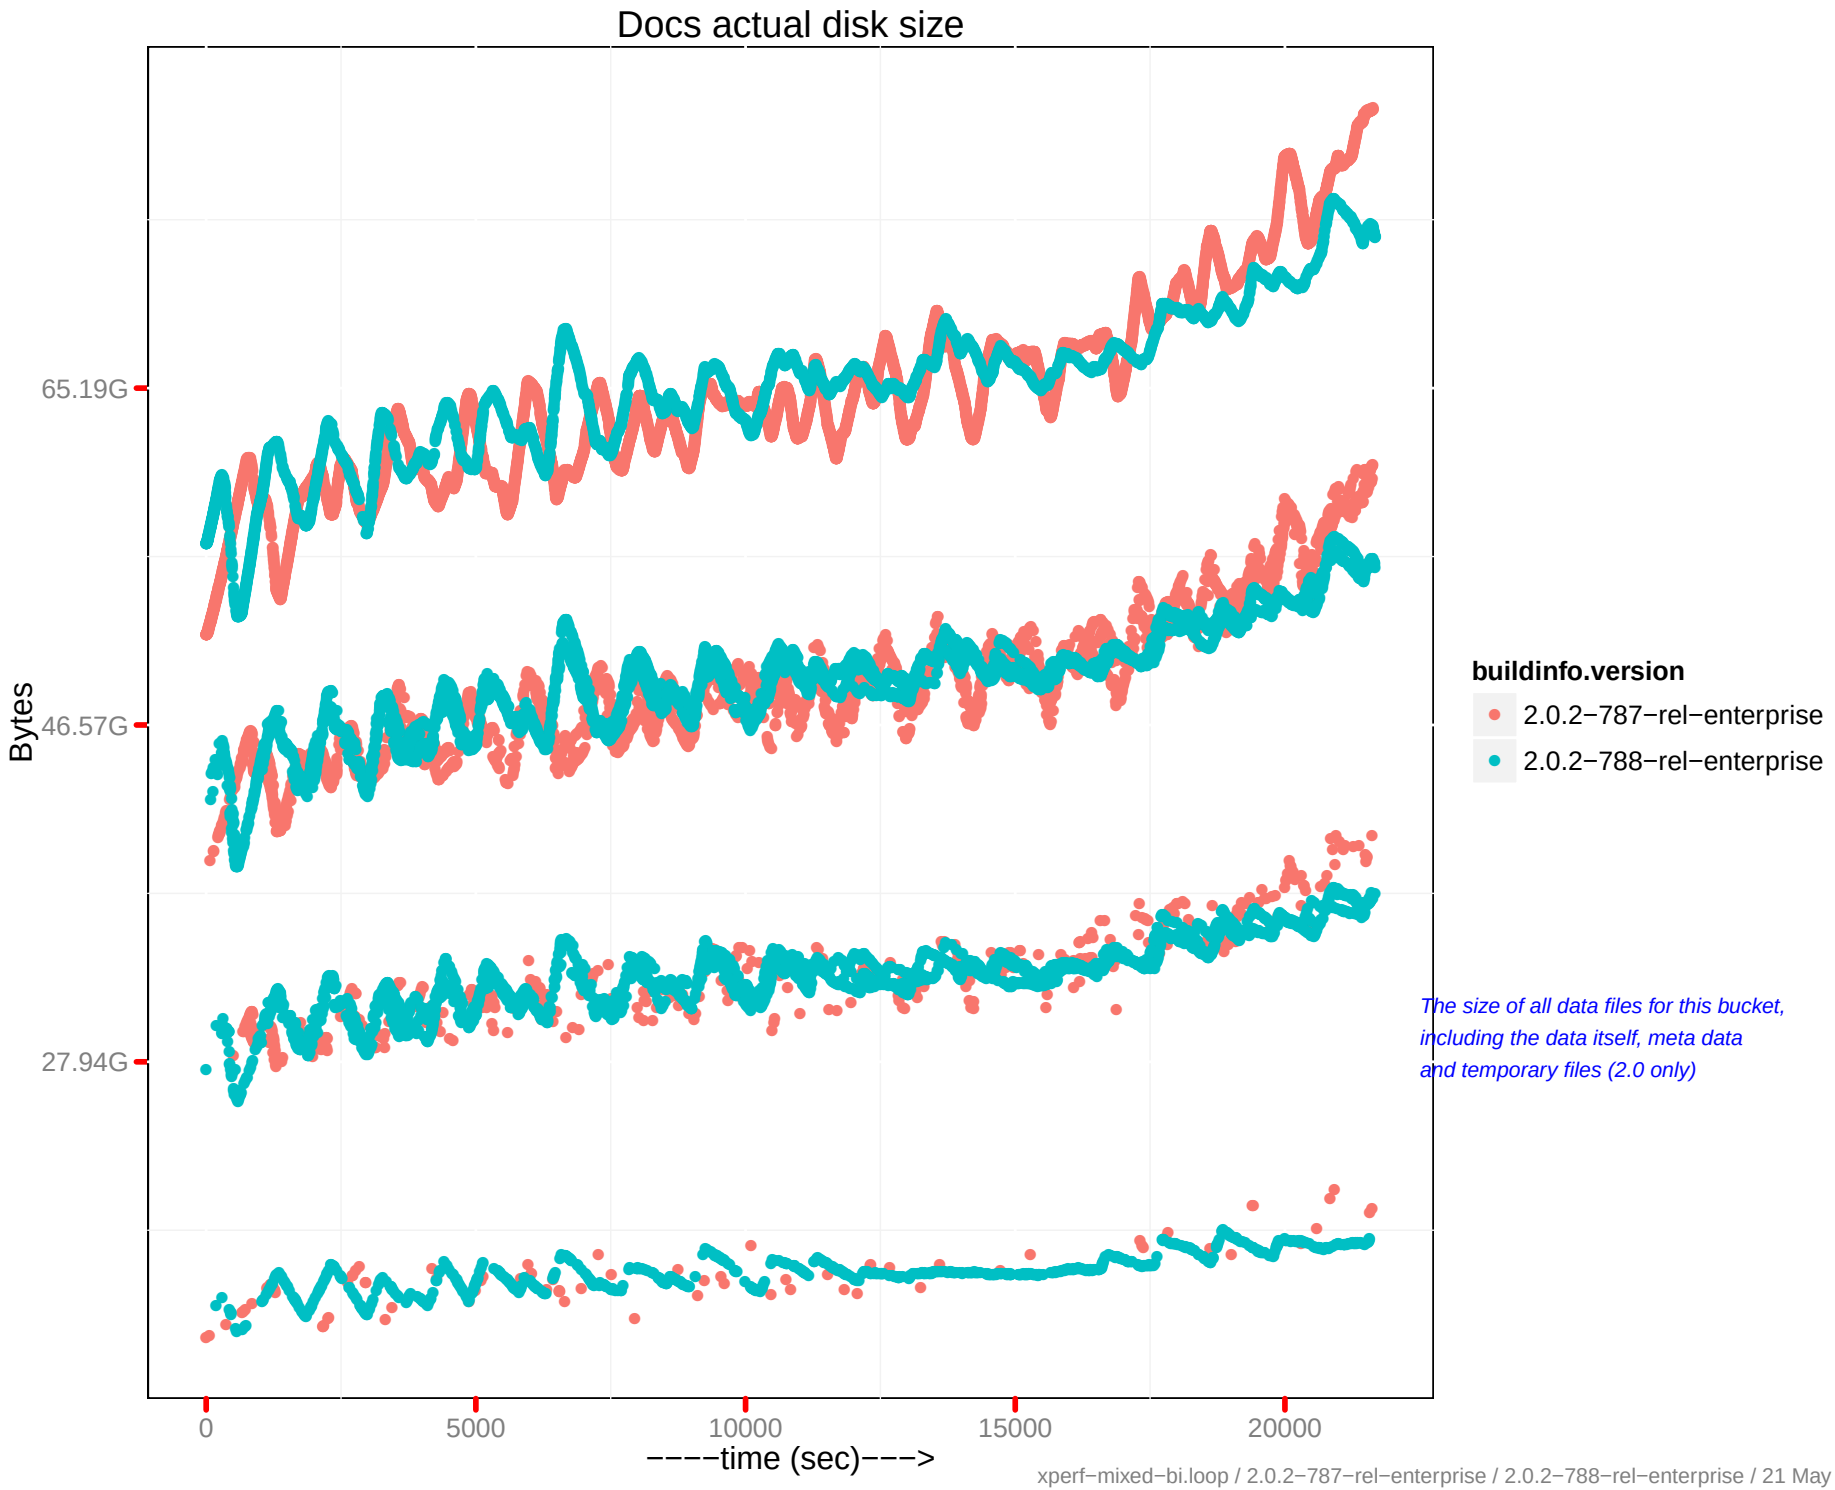
Docs actual disk size
65.19G
46.57G
27.94G
Bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

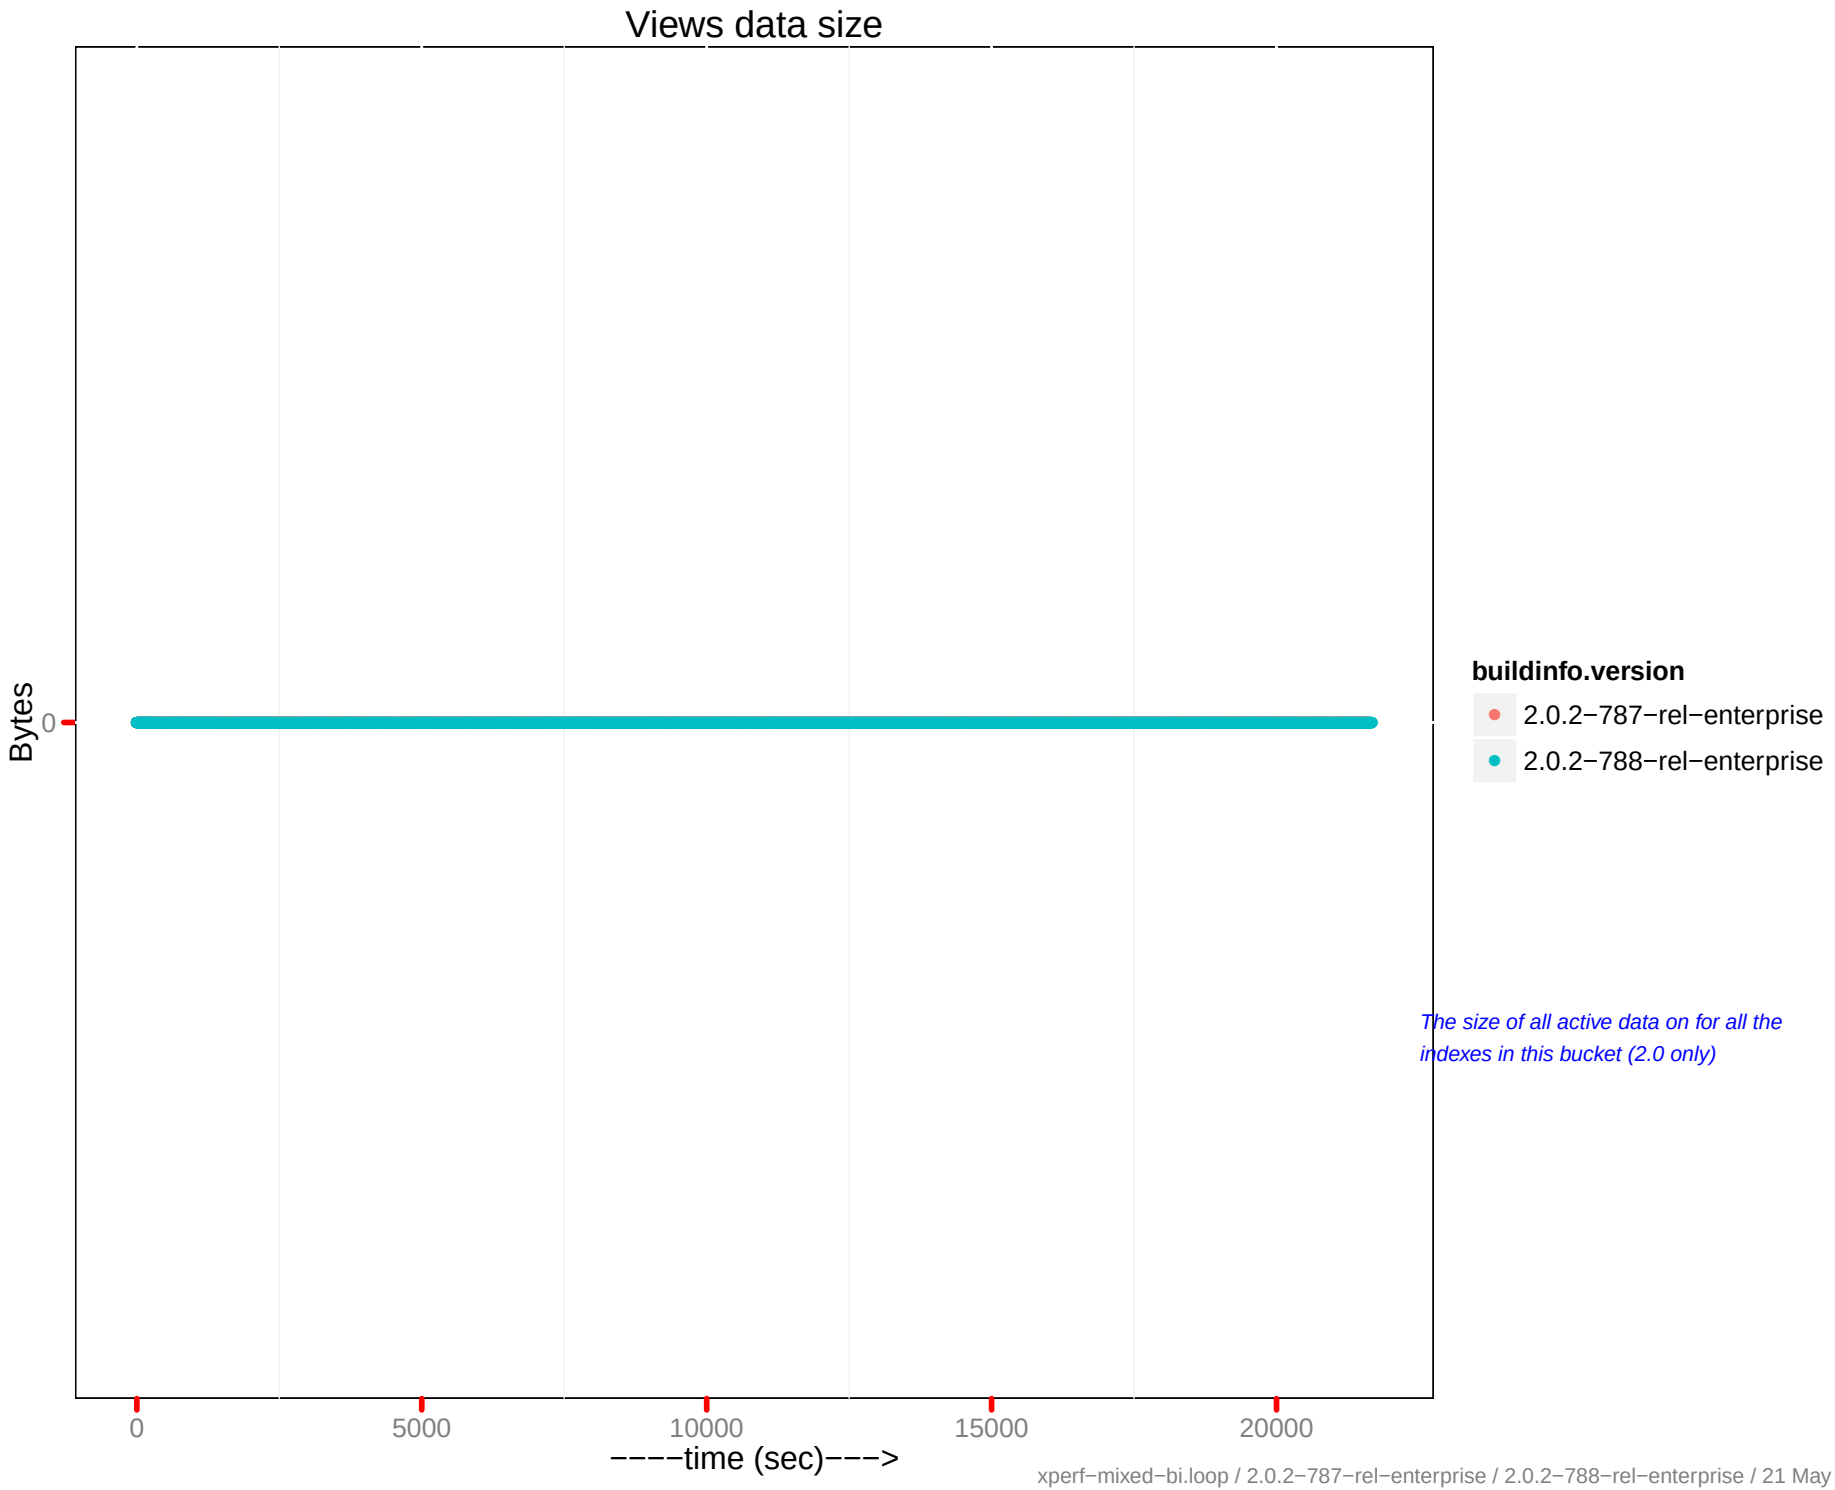
Views data size
Bytes
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The size of all active data on for all the indexes in this bucket (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

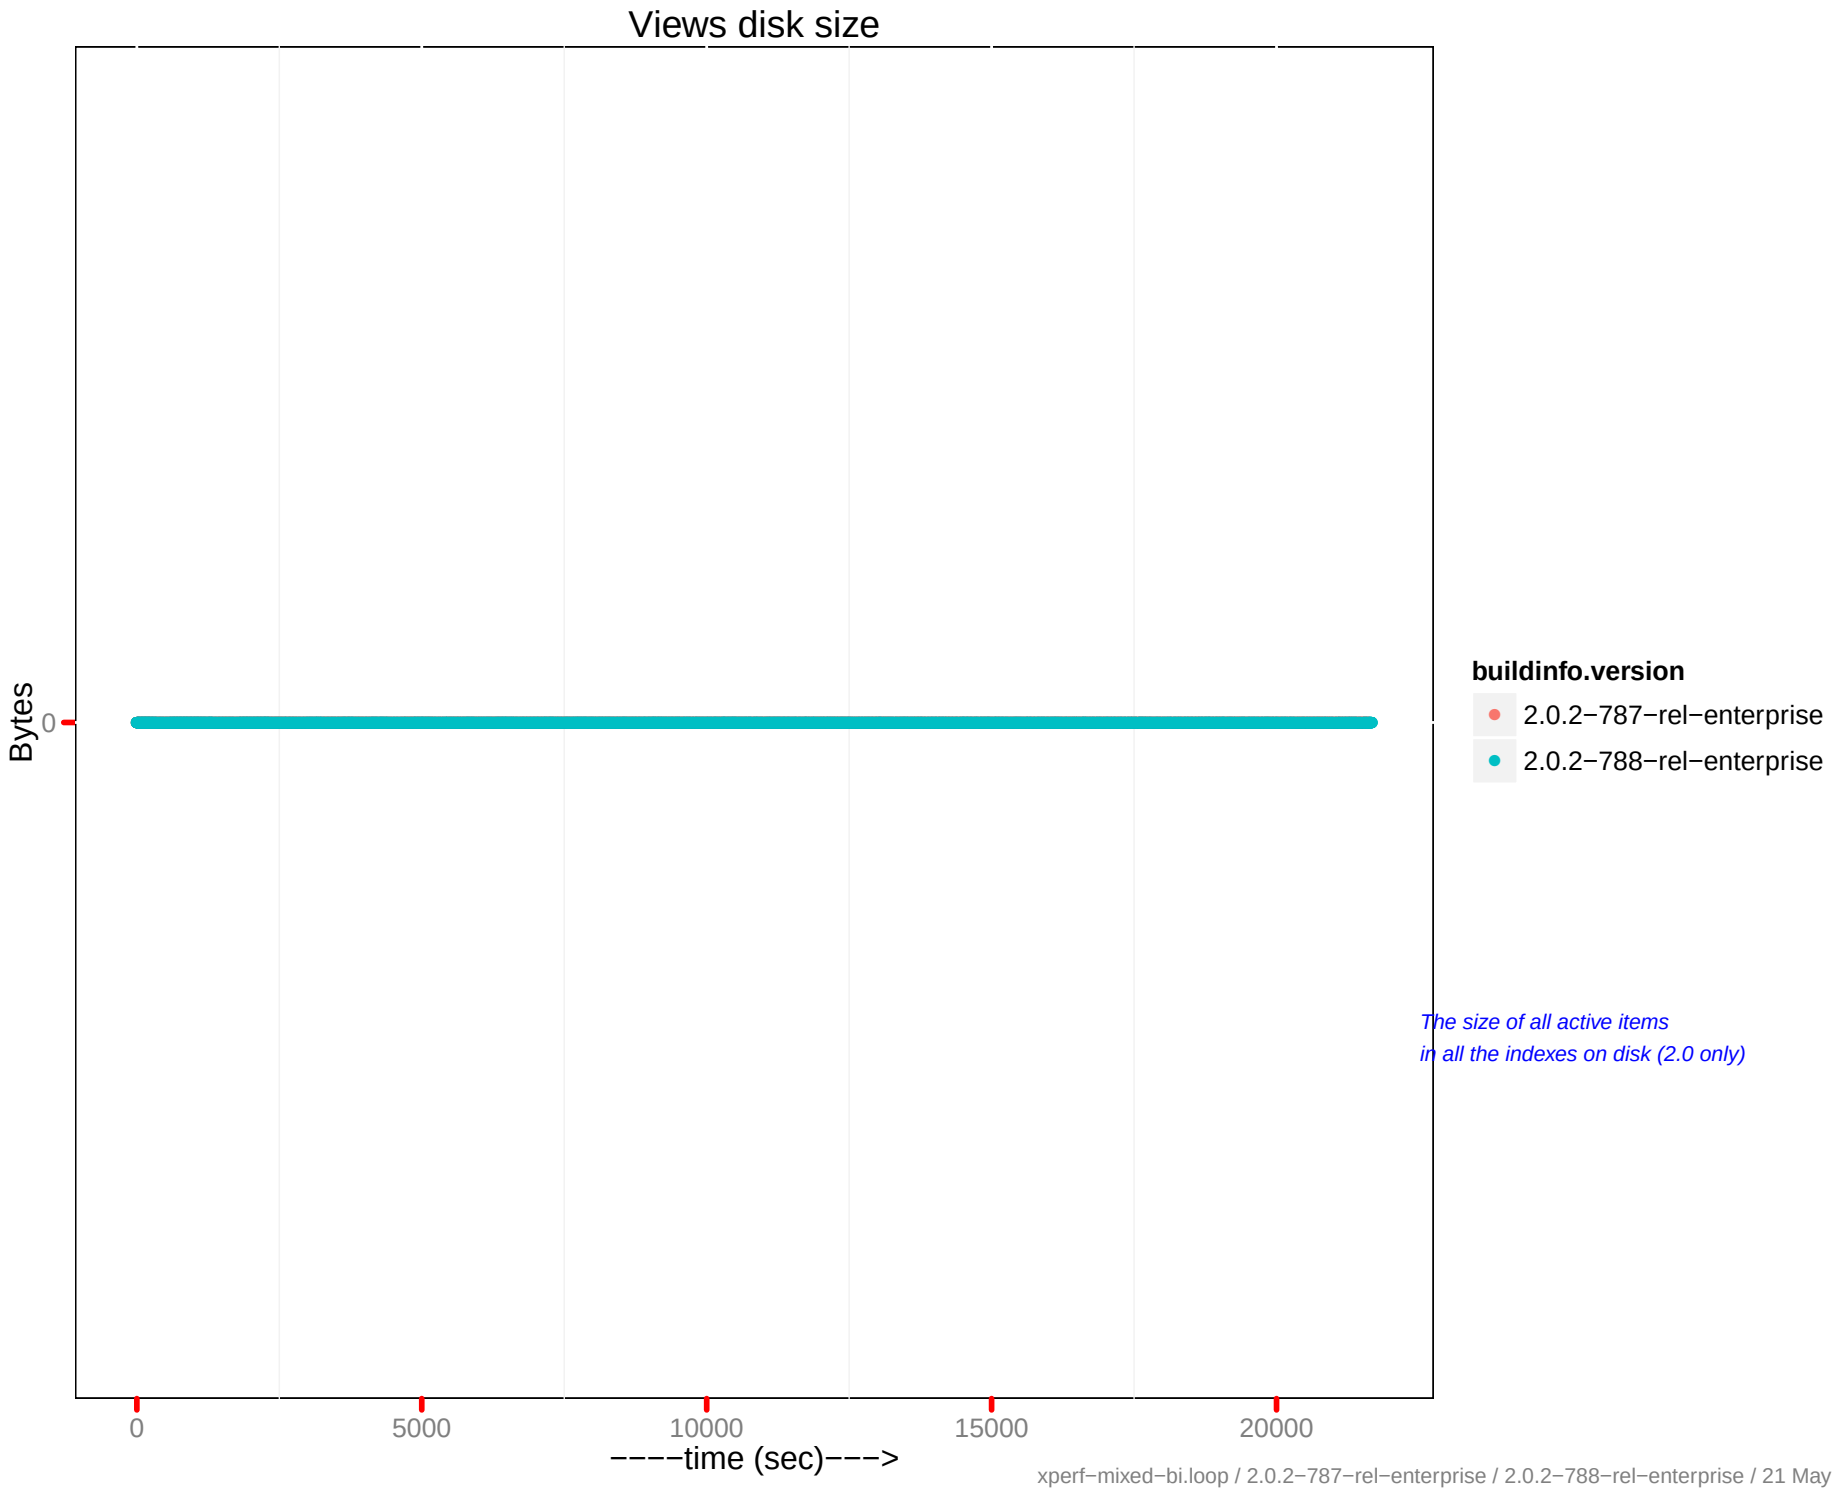
Views disk size
Bytes
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

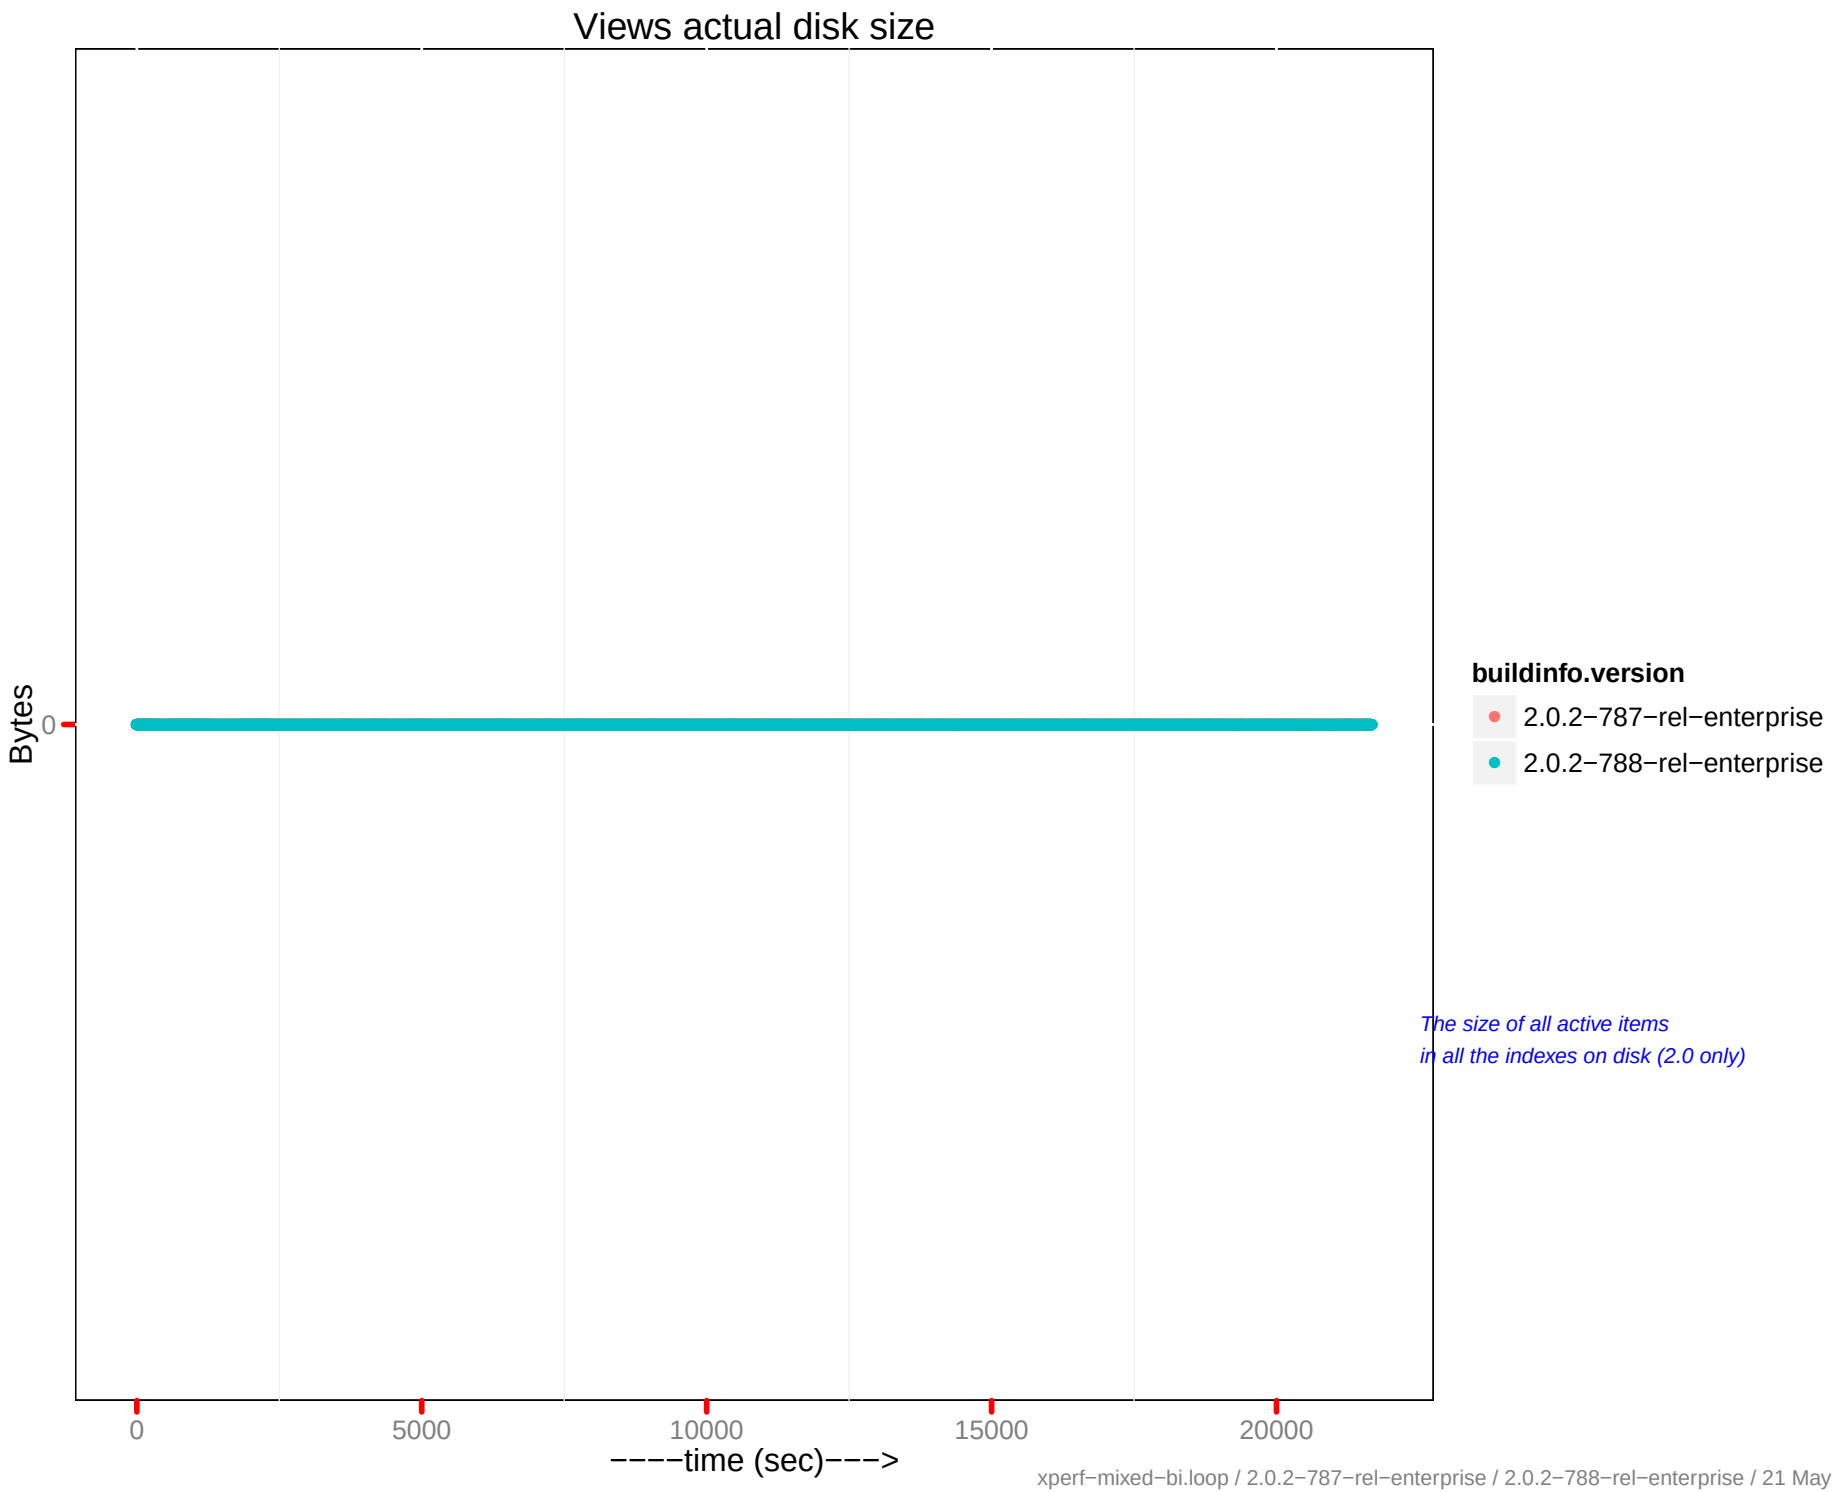
Views actual disk size
Bytes
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The size of all active items
in all the indexes on disk (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

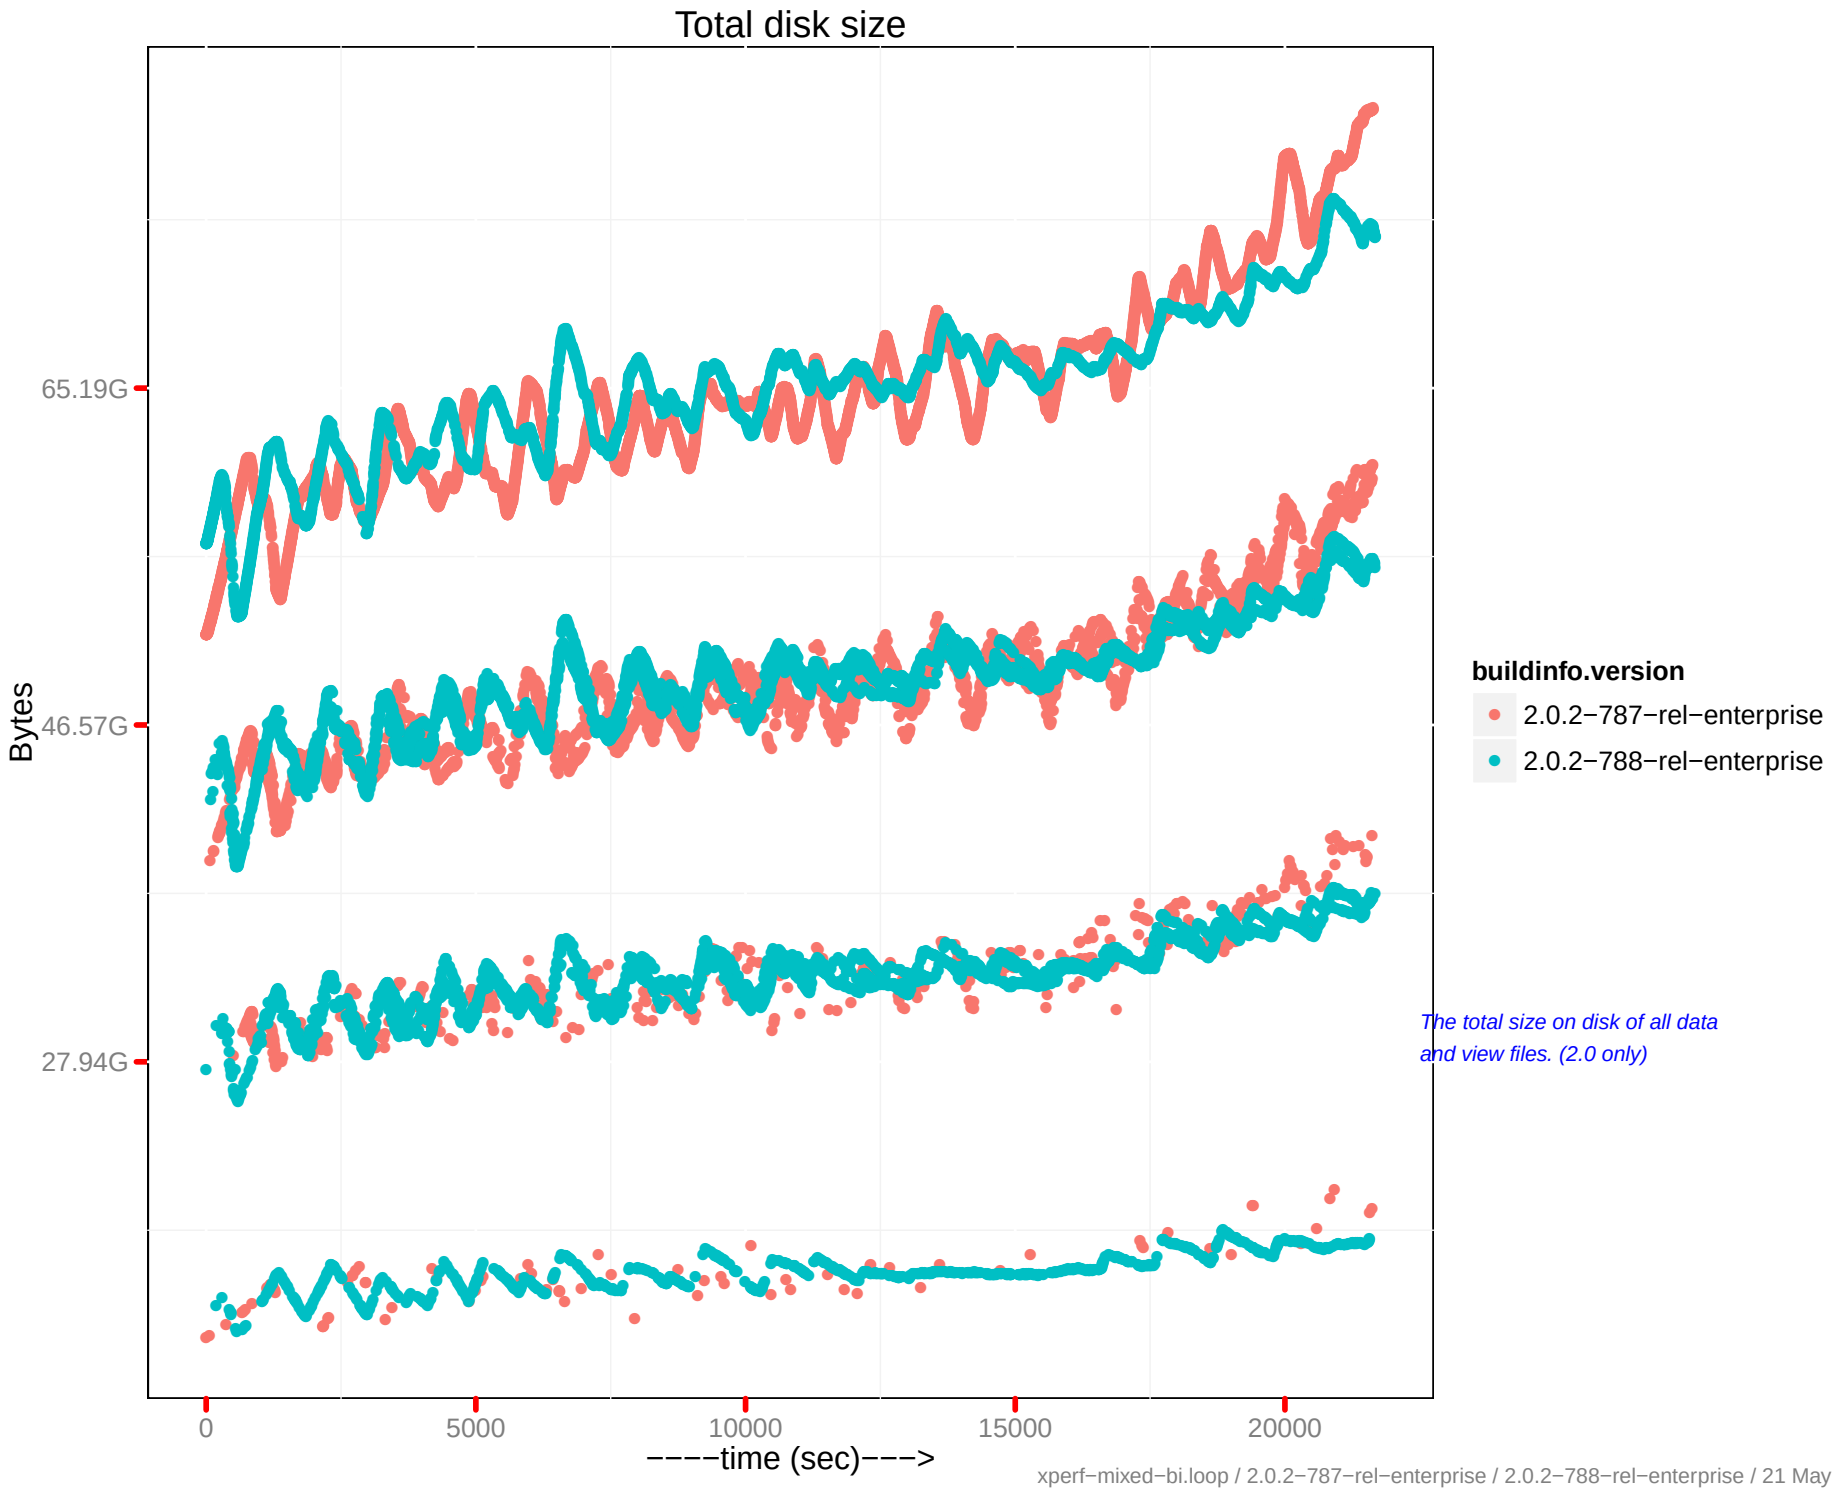
Total disk size
65.19G
46.57G
27.94G
Bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
The total size on disk of all data and view files. (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

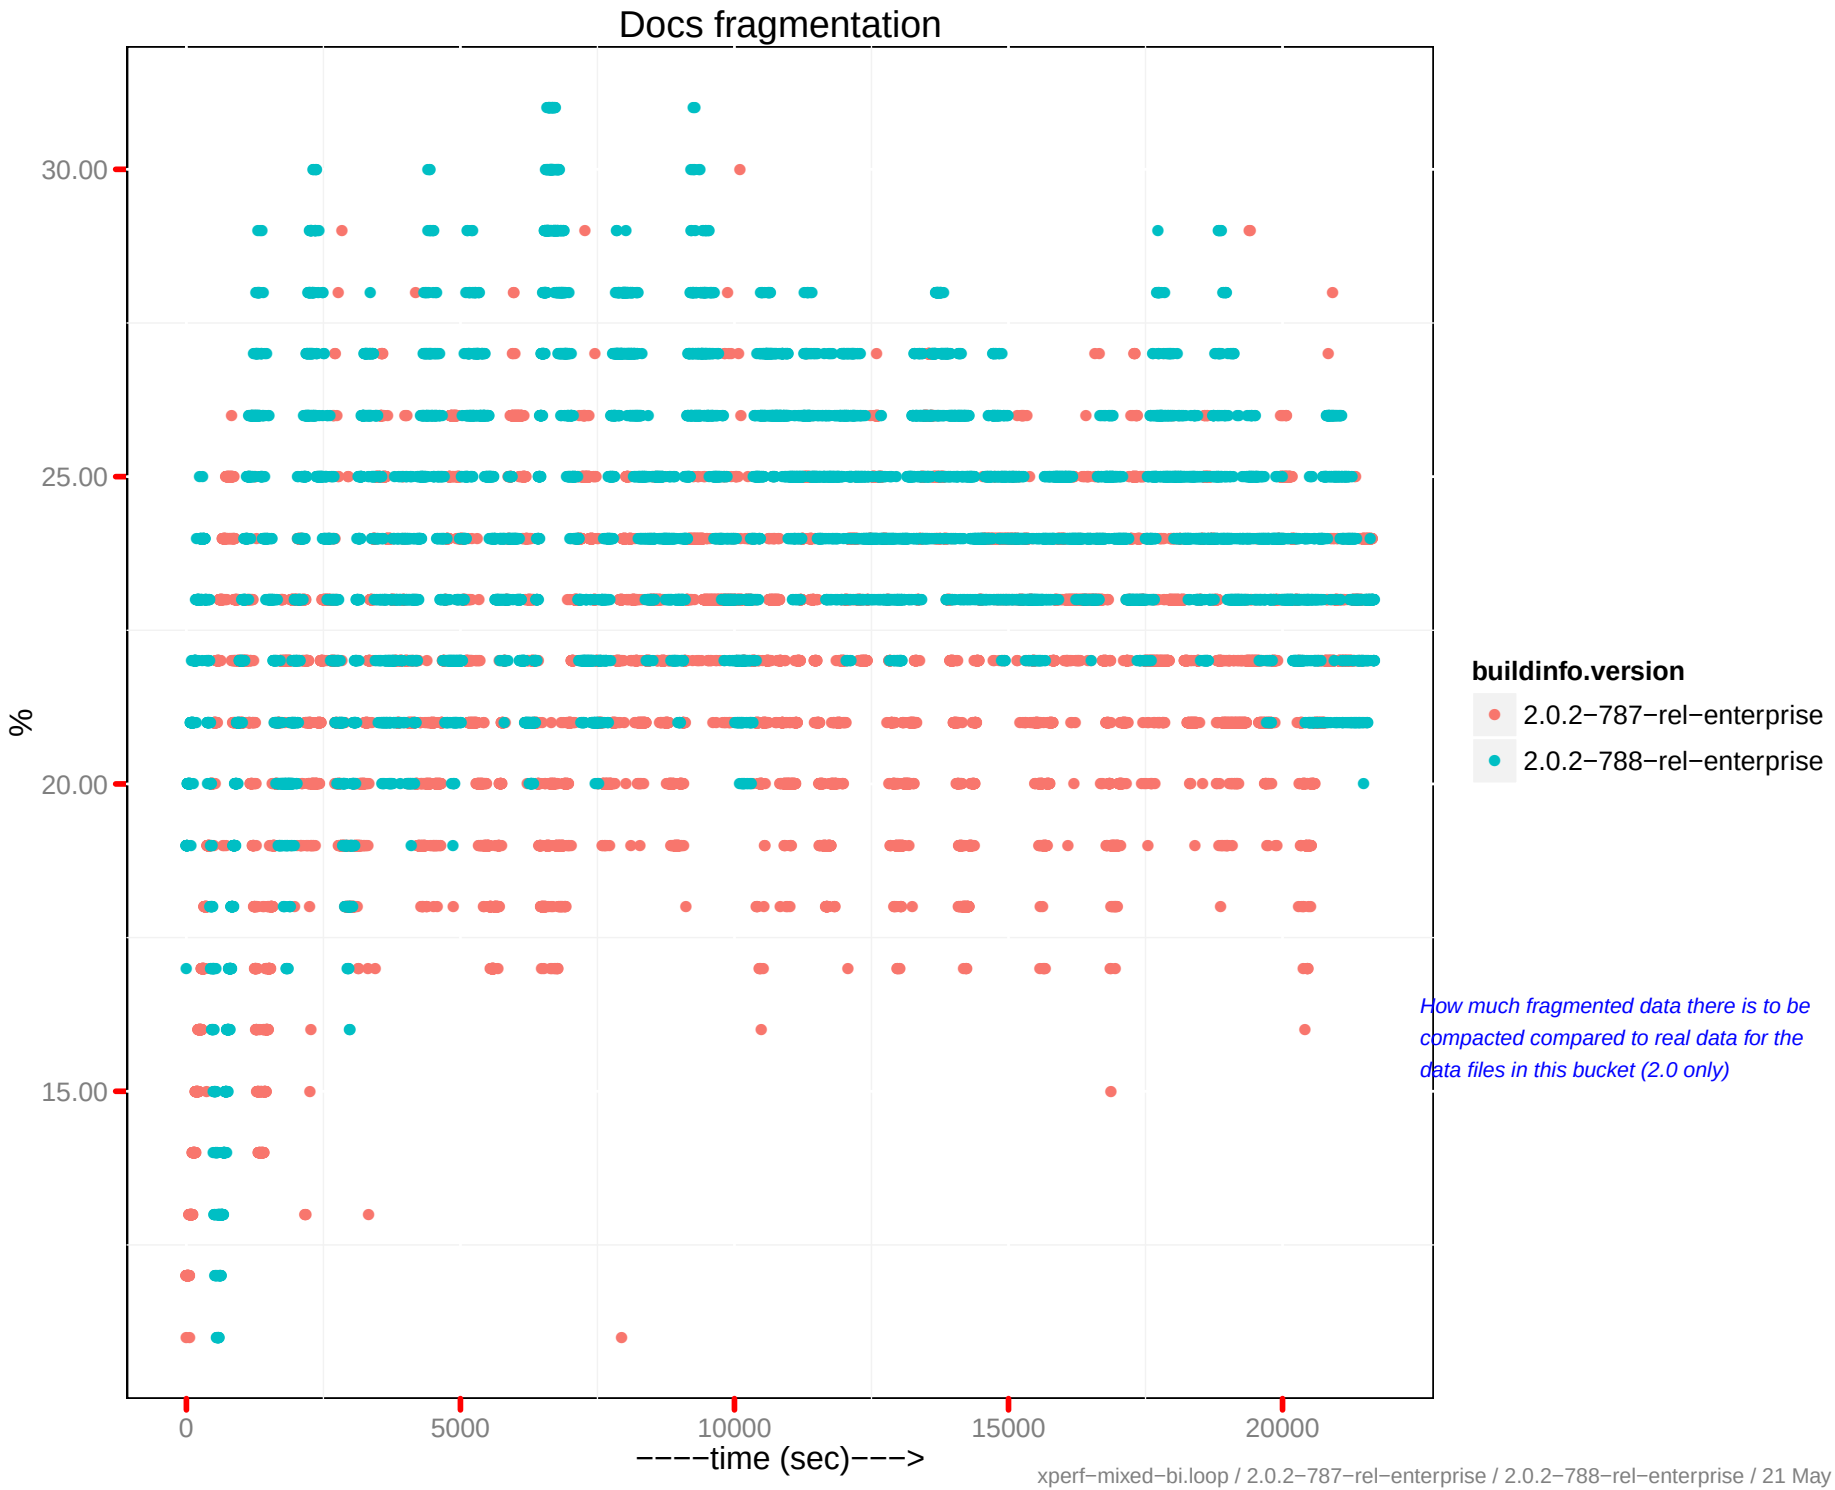
Docs fragmentation
30.00
25.00
20.00
15.00
%
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the data files in this bucket (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

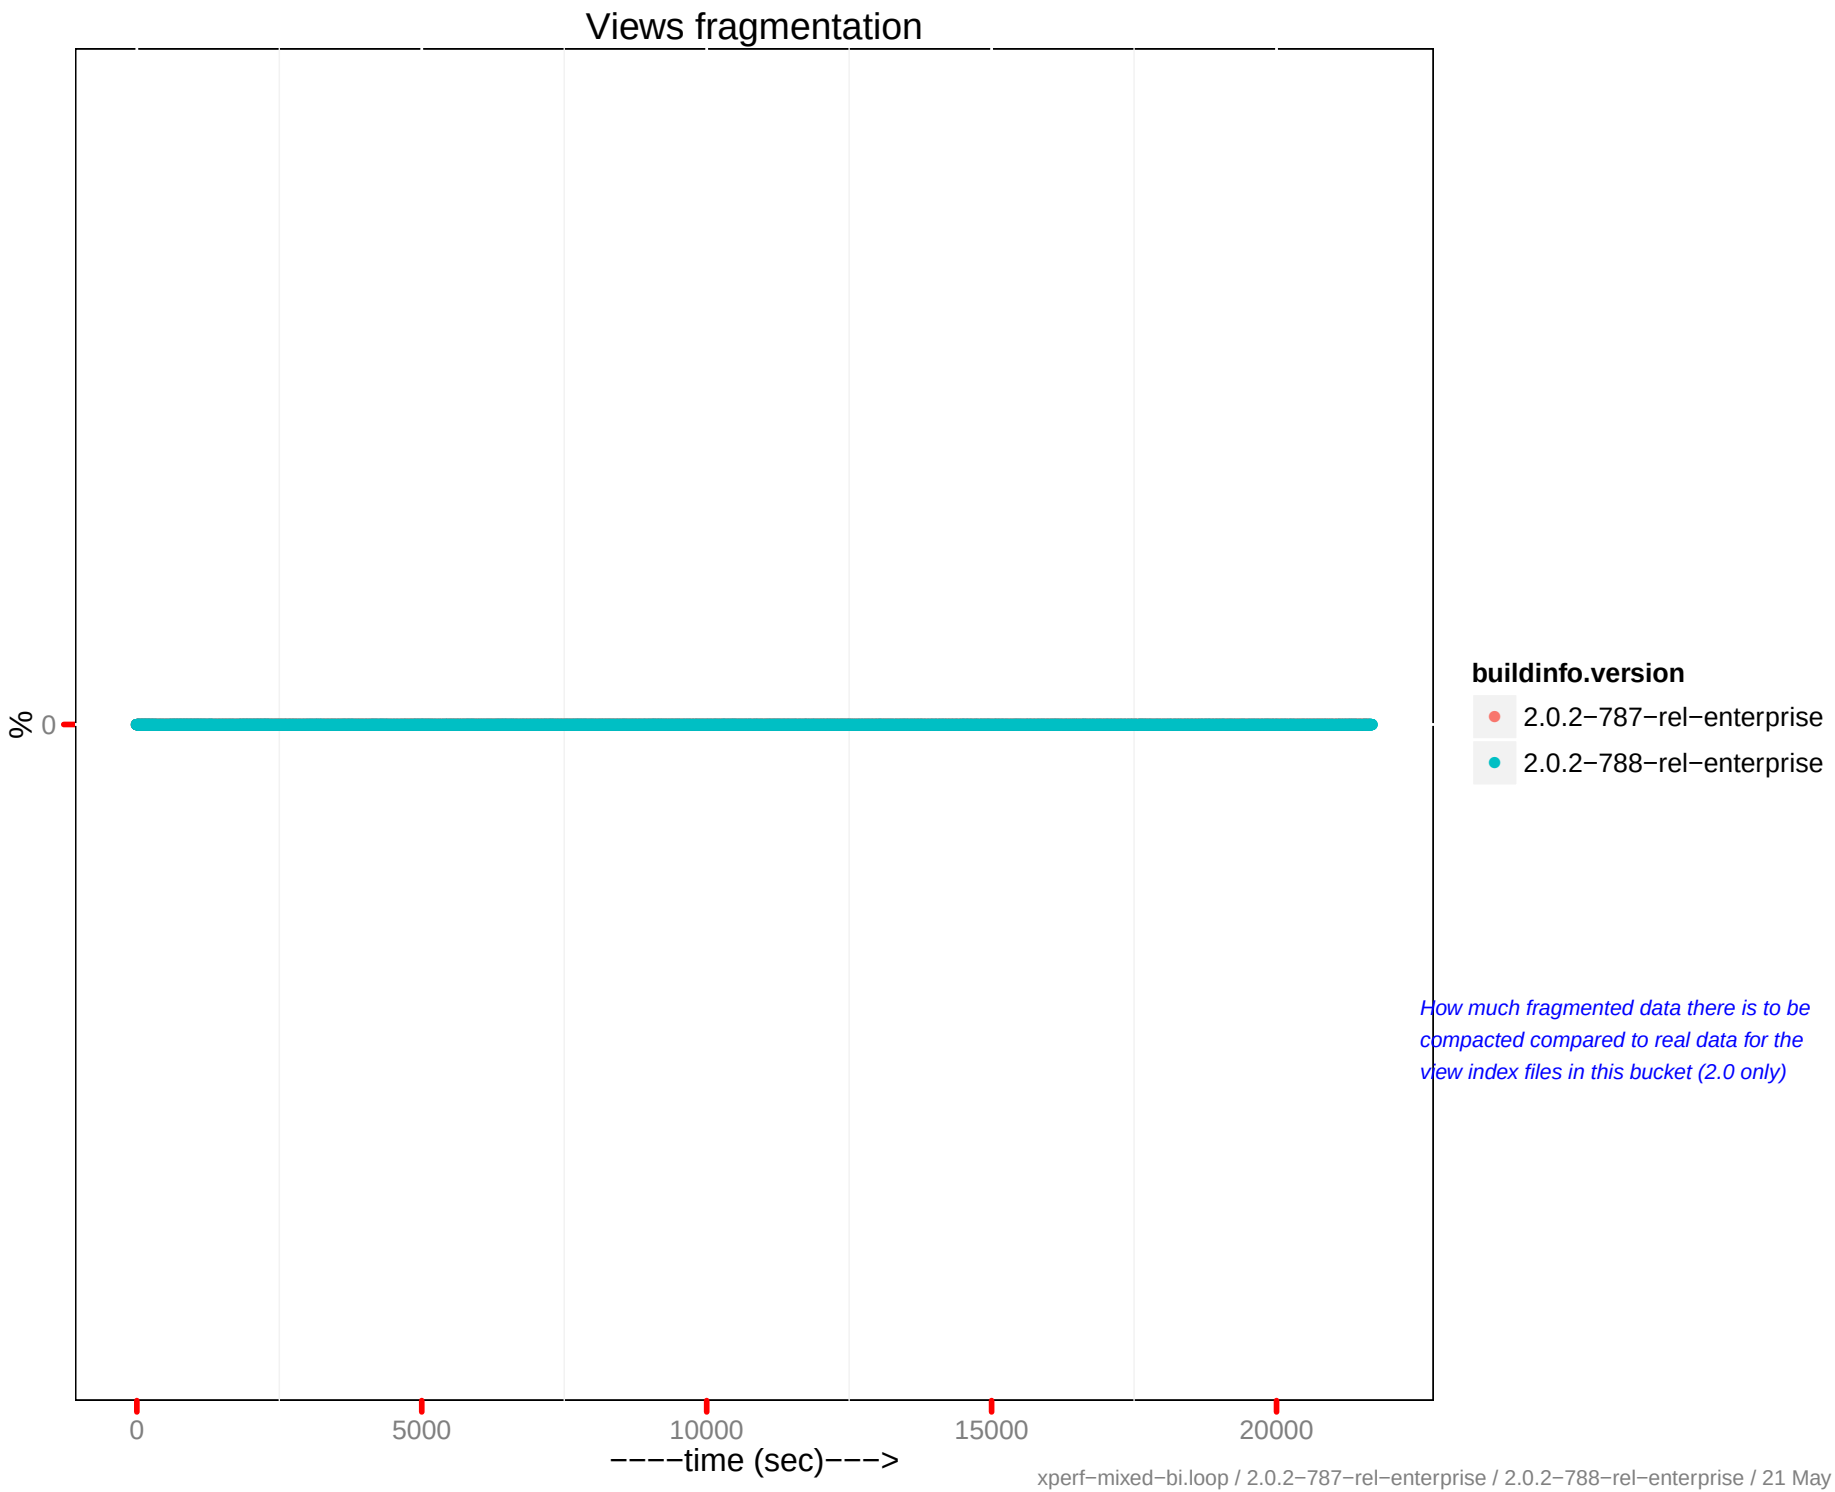
Views fragmentation
% 0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
How much fragmented data there is to be compacted compared to real data for the view index files in this bucket (2.0 only)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013


cmd_get ops/sec
29.30K
19.53K
9.77K
0
ops/sec
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Cumulative number of retrieval reqs
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013


cmd_set ops/sec
19.53K
9.77K
0
ops/sec
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Cumulative number of set reqs
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

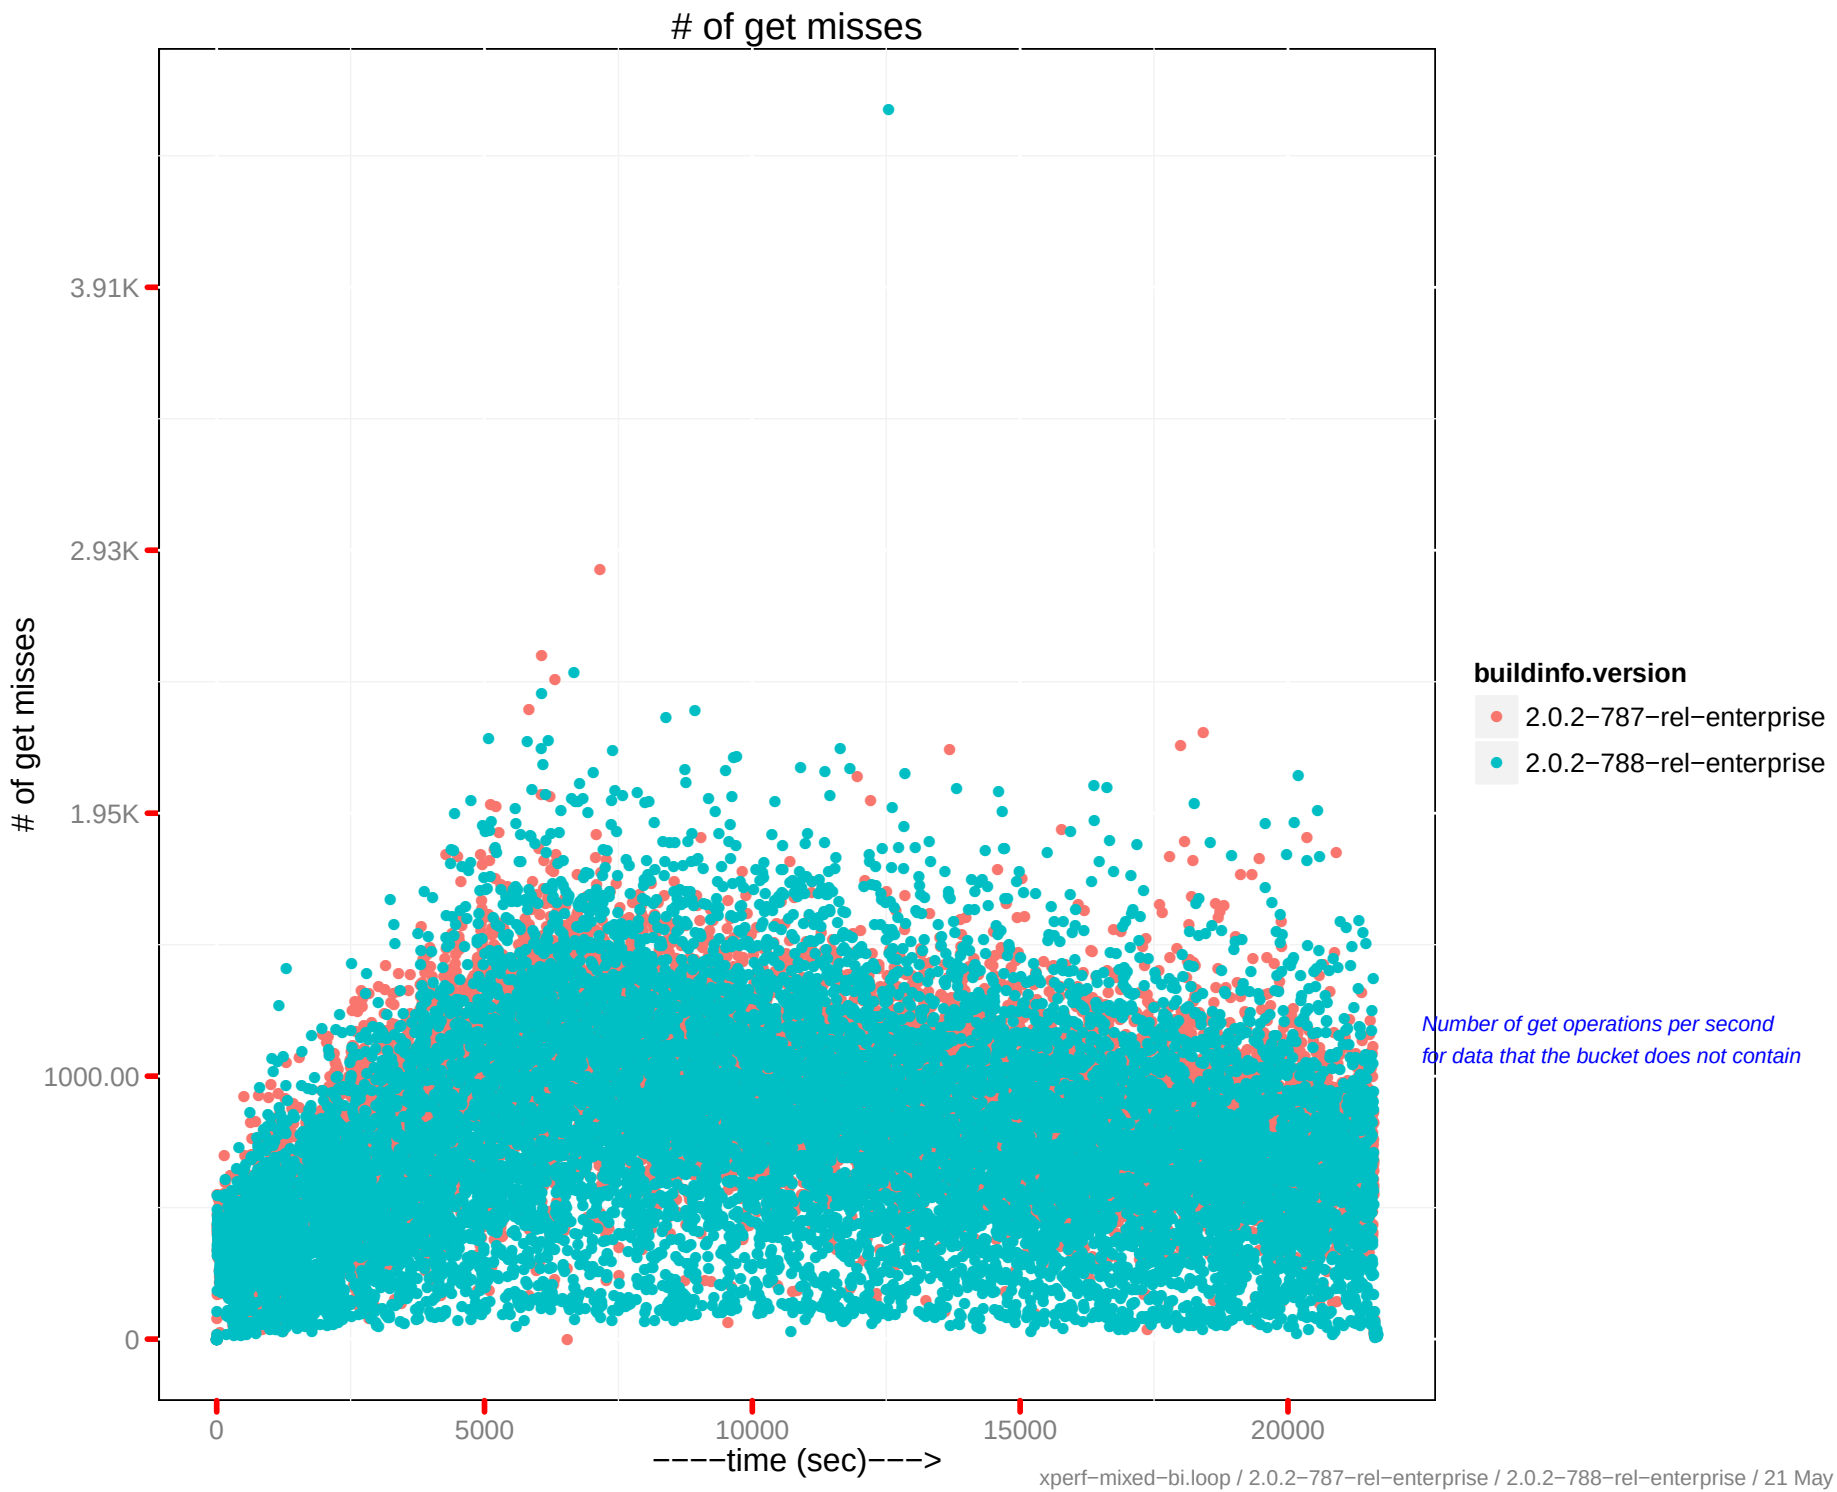
# of get misses
3.91K
2.93K
1.95K
1000.00
0
# of get misses
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of get operations per second
for data that the bucket does not contain
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

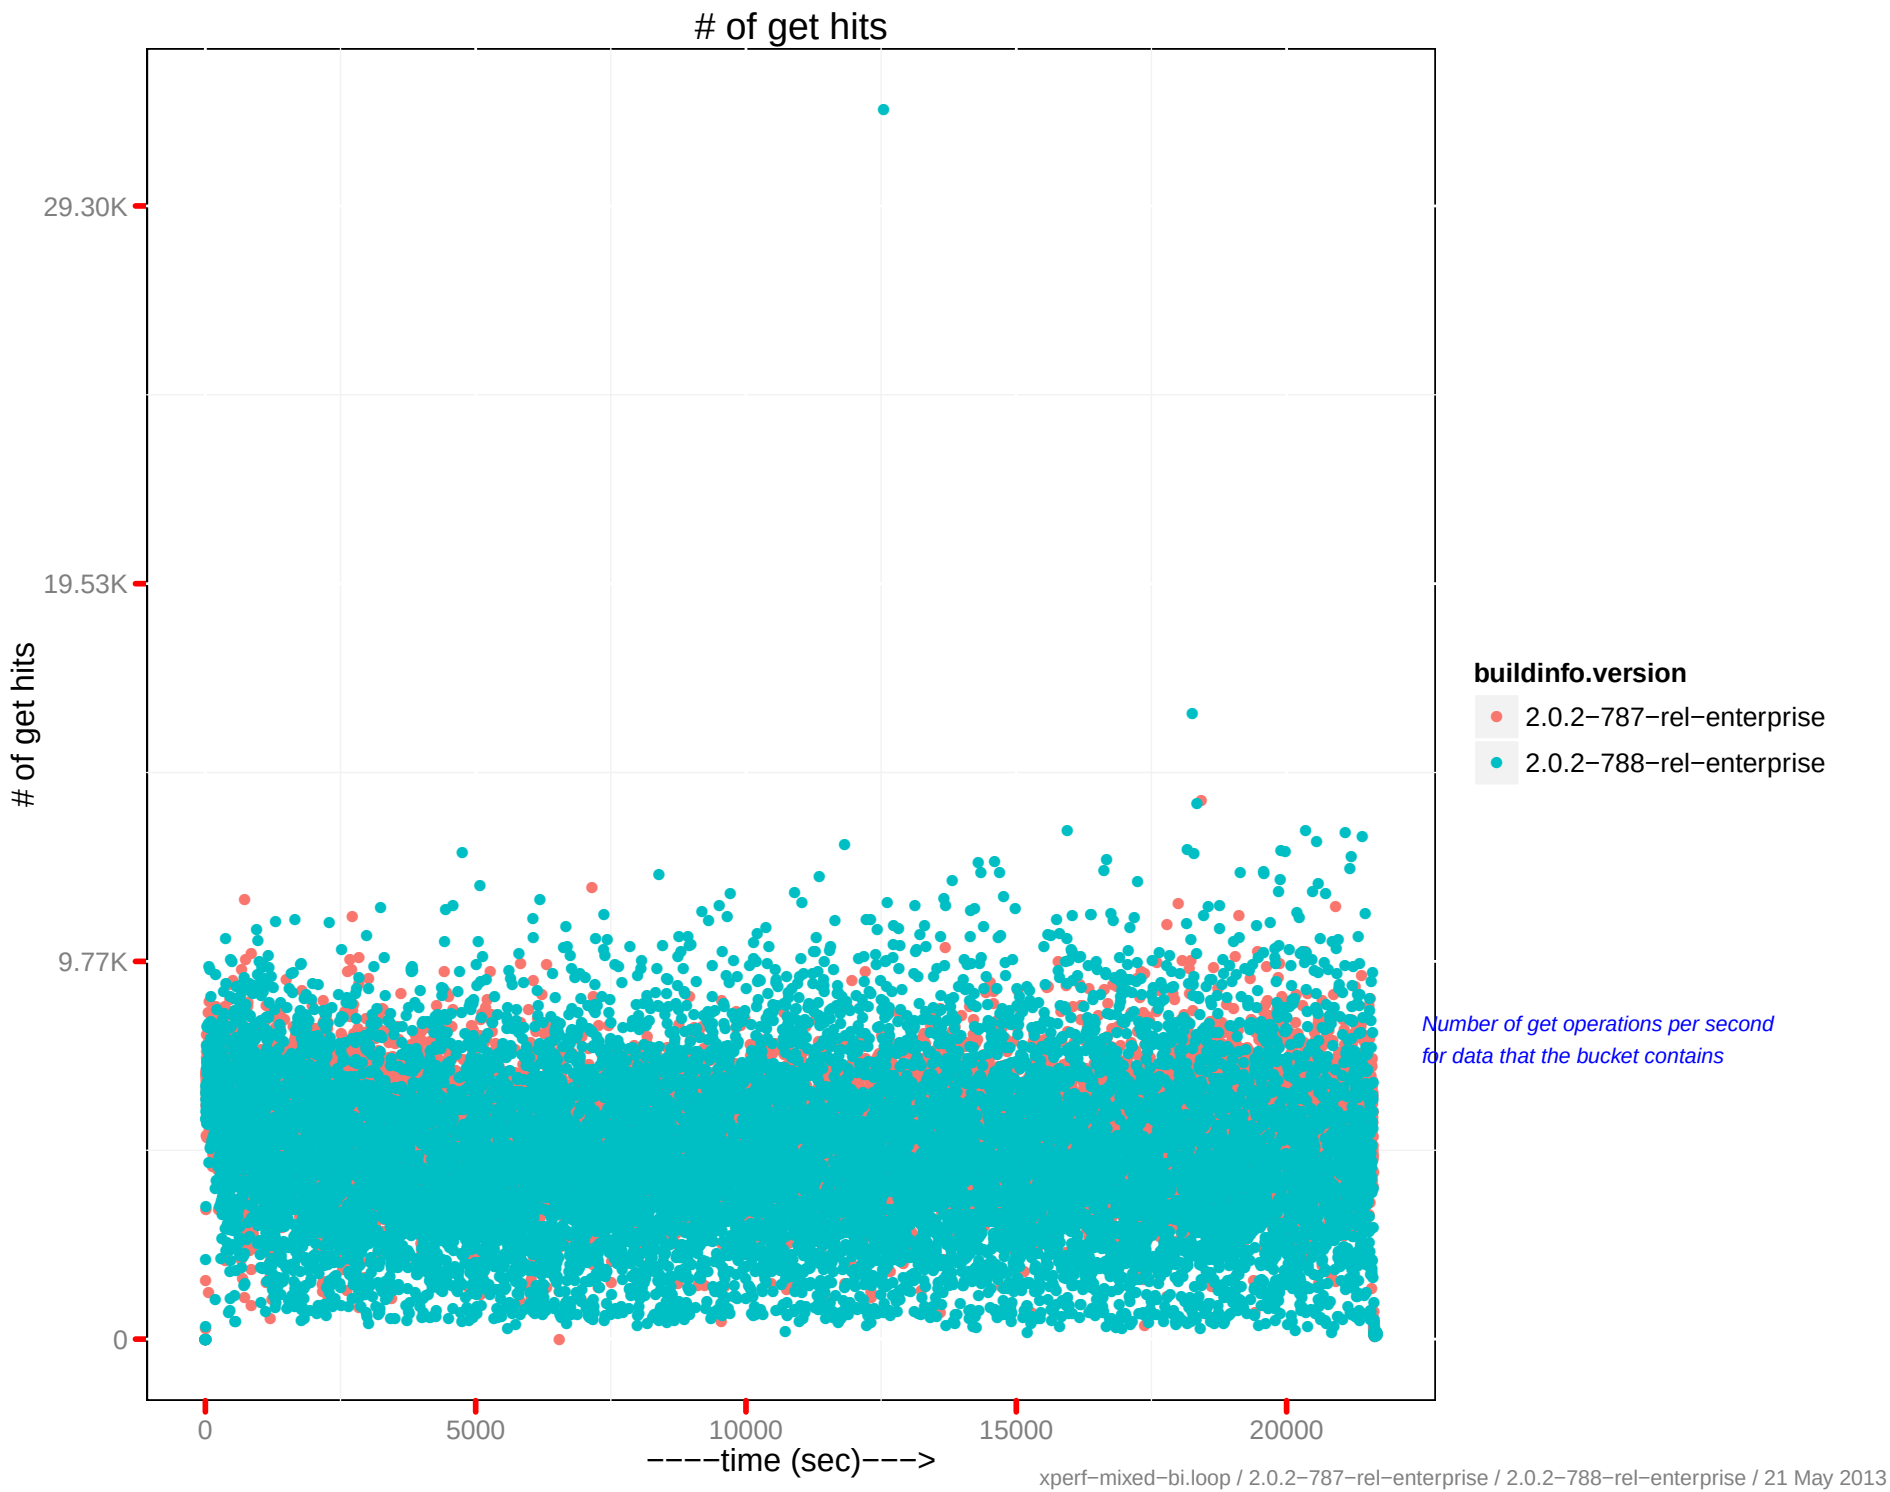
# of get hits
29.30K
19.53K
9.77K
0
# of get hits
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of get operations per second
for data that the bucket contains
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

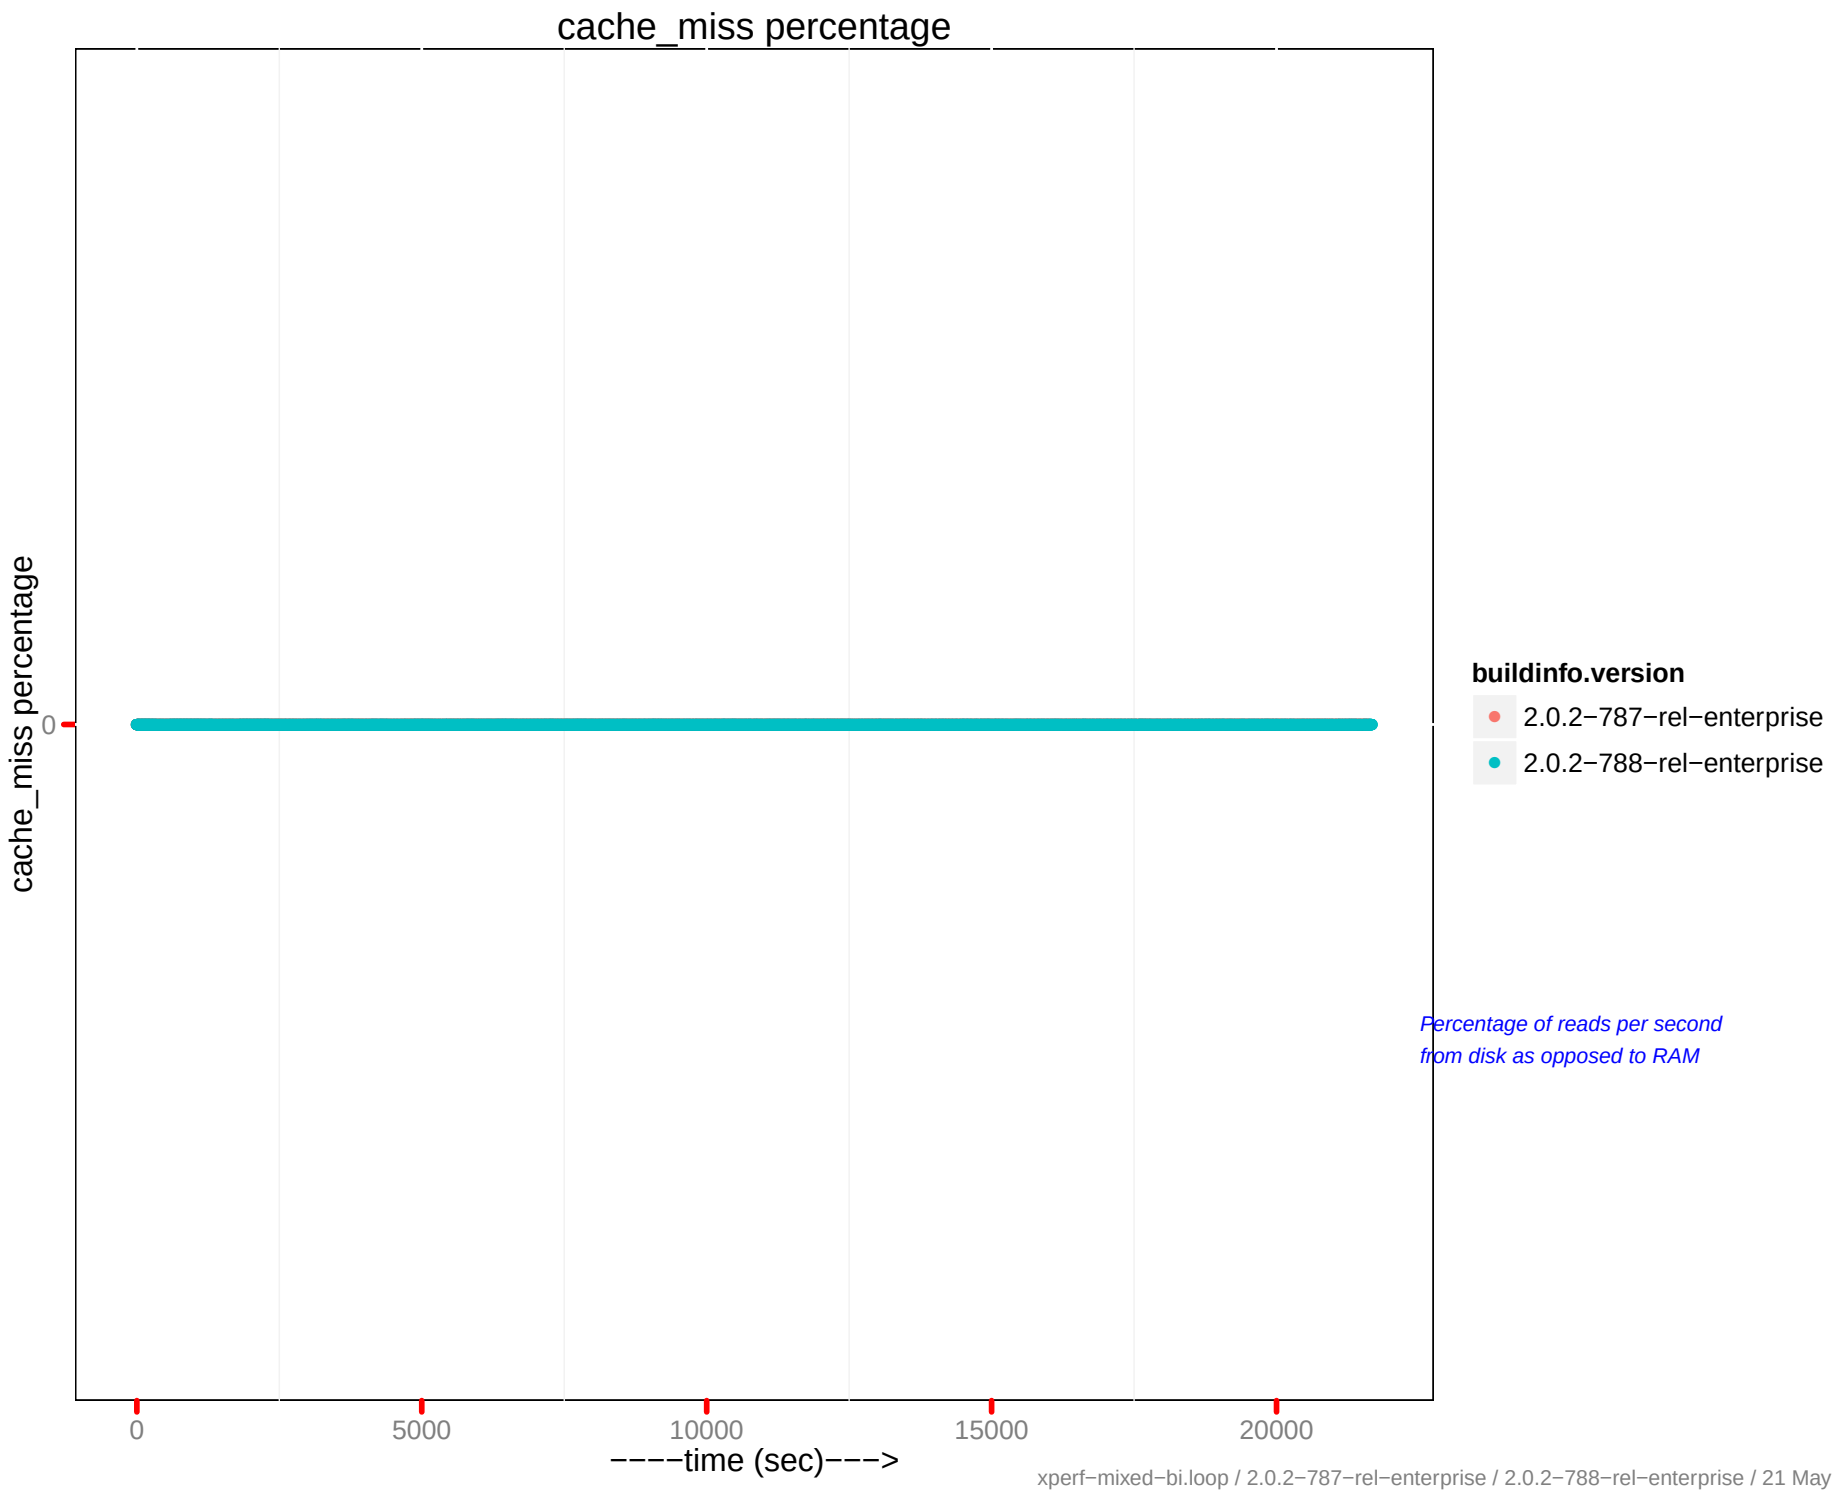
cache_miss percentage
cache_miss percentage
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Percentage of reads per second
from disk as opposed to RAM
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

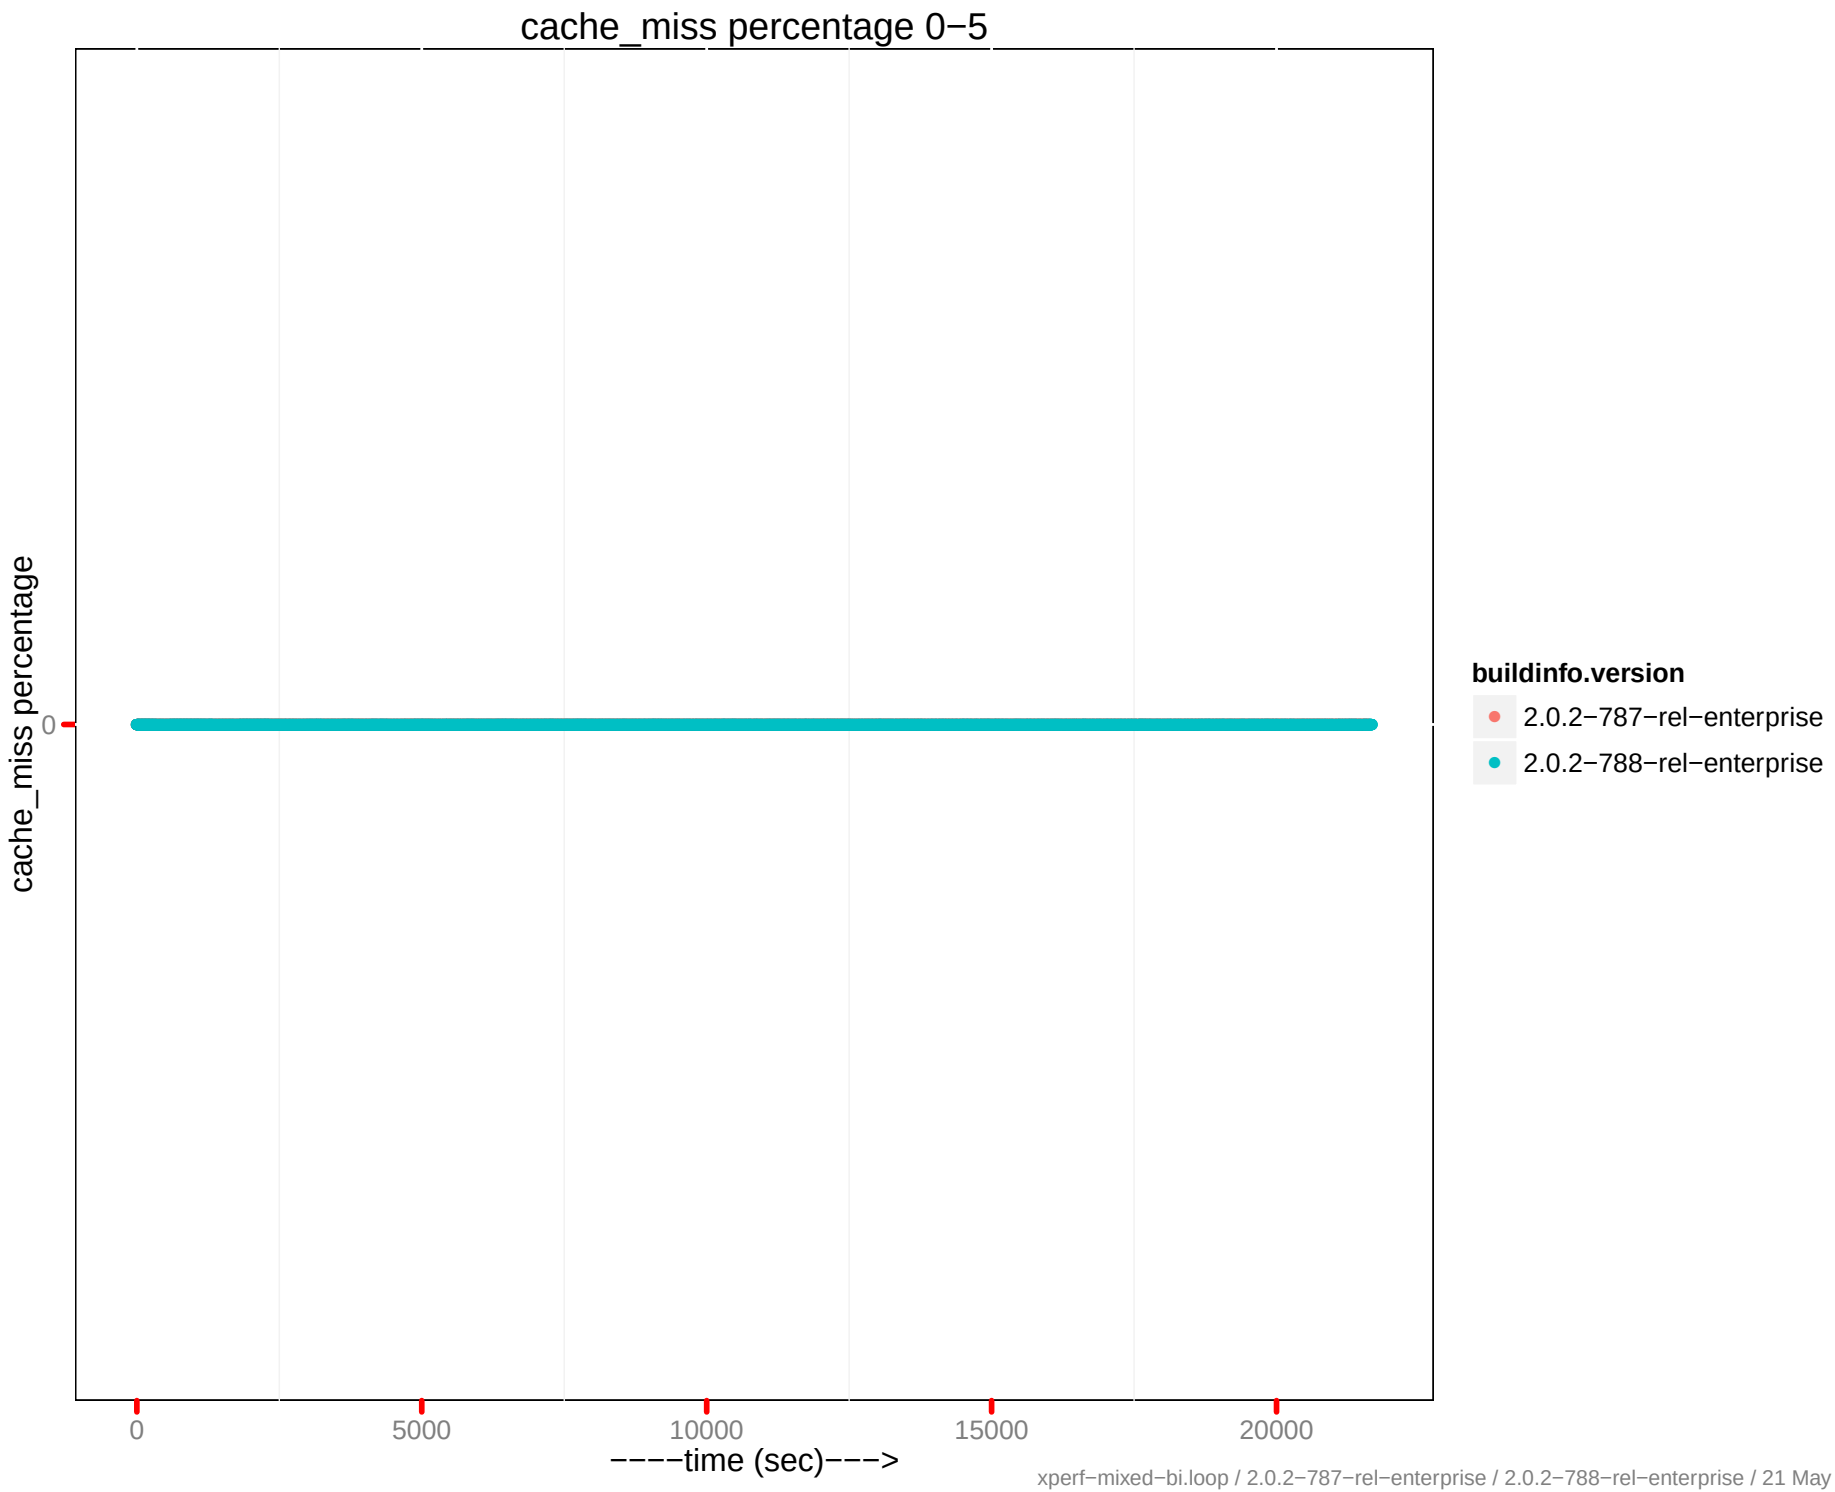
cache_miss percentage 0−5
cache_miss percentage
0
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

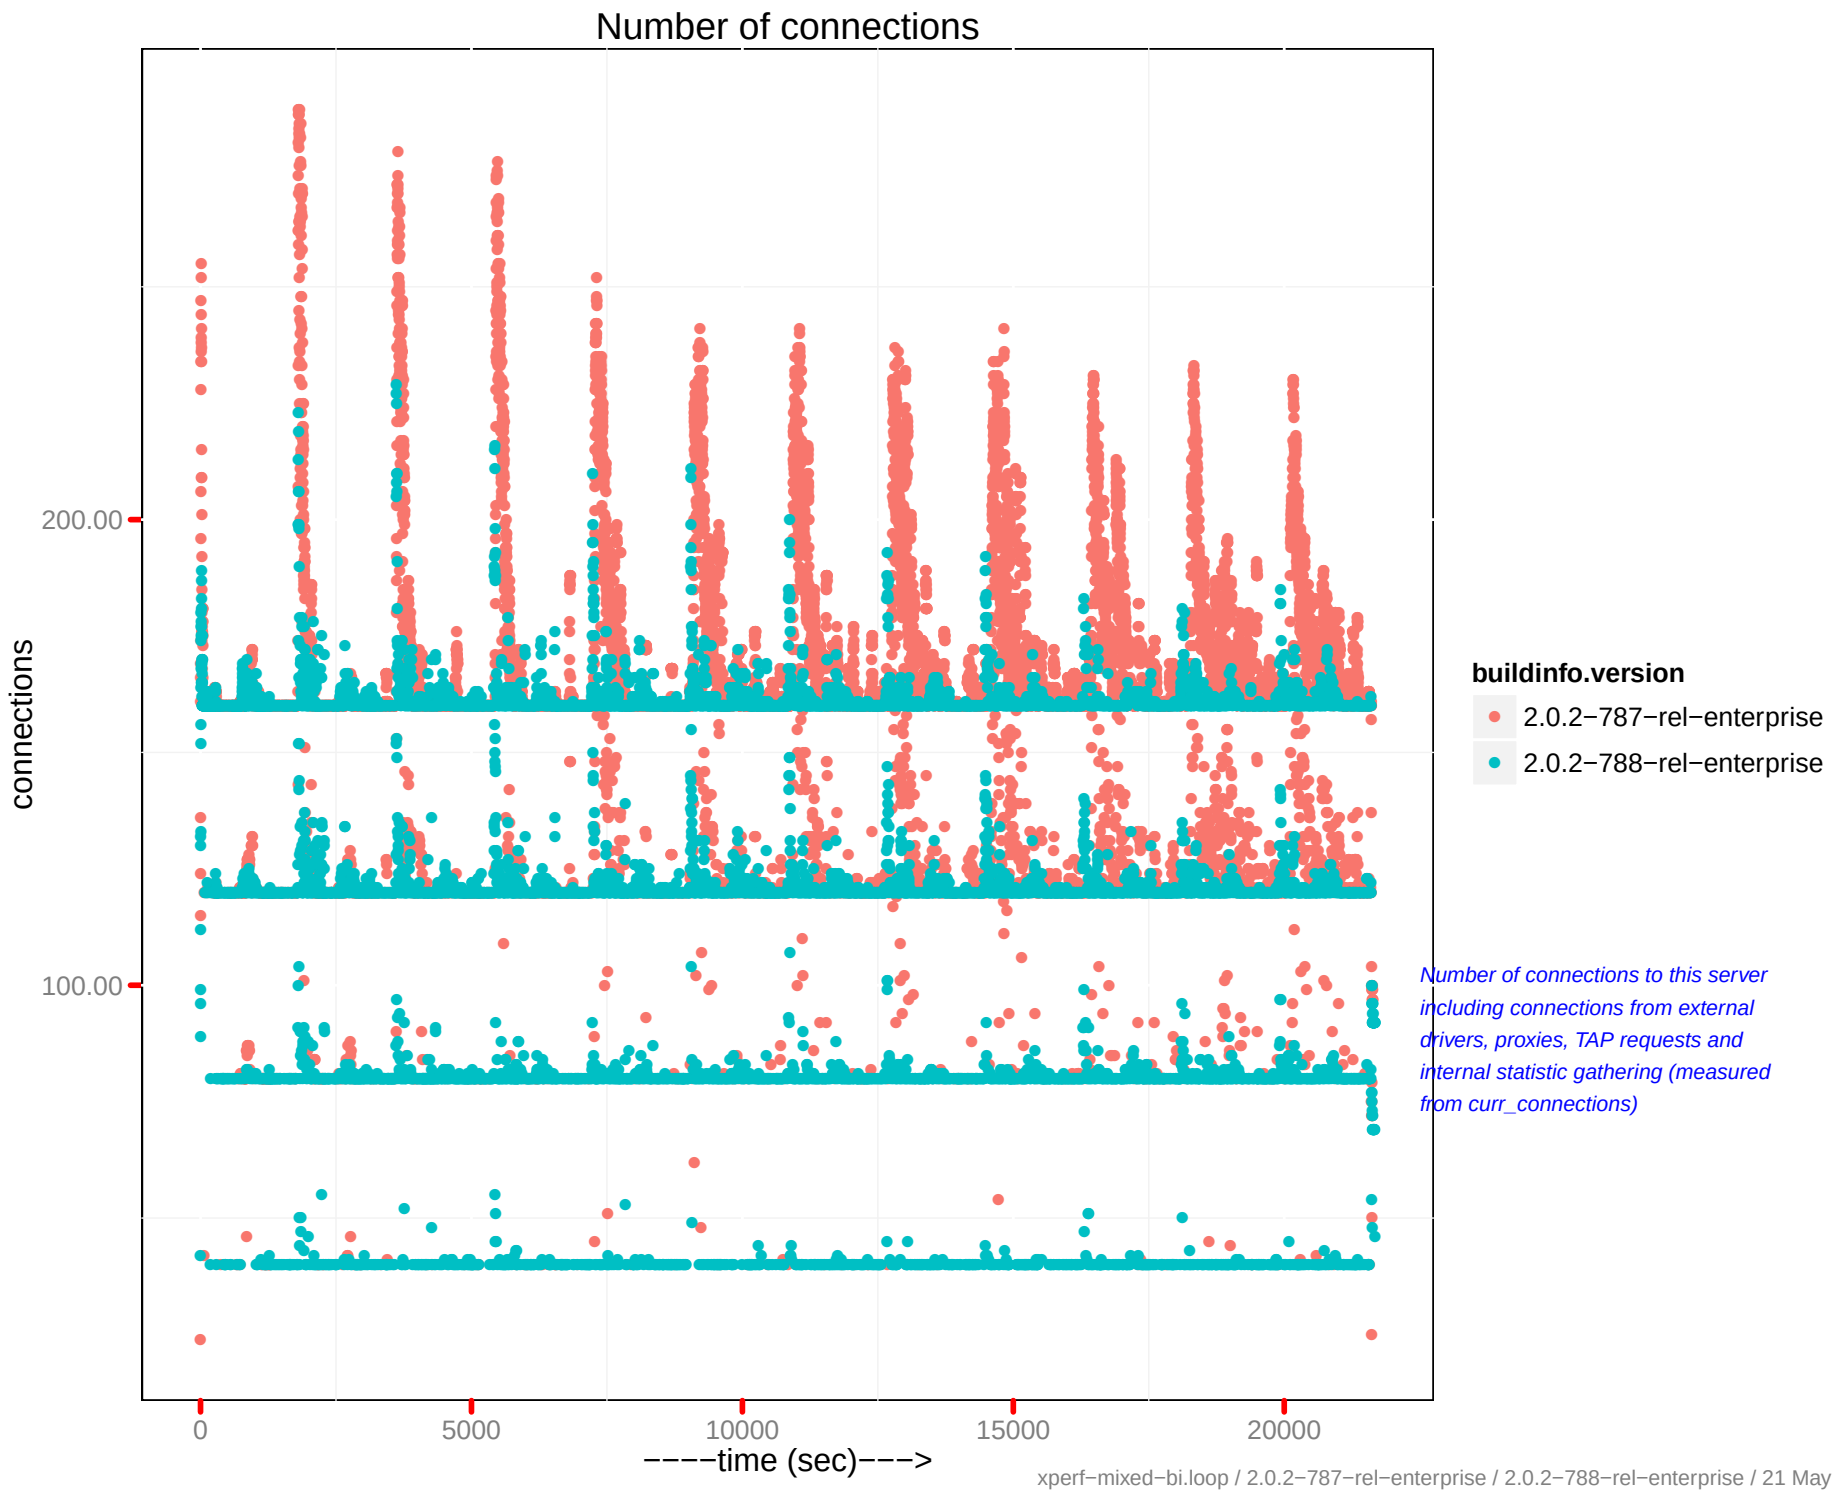
Number of connections
connections
200.00
100.00
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of connections to this server including connections from external drivers, proxies, TAP requests and internal statistic gathering (measured from curr_connections)
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

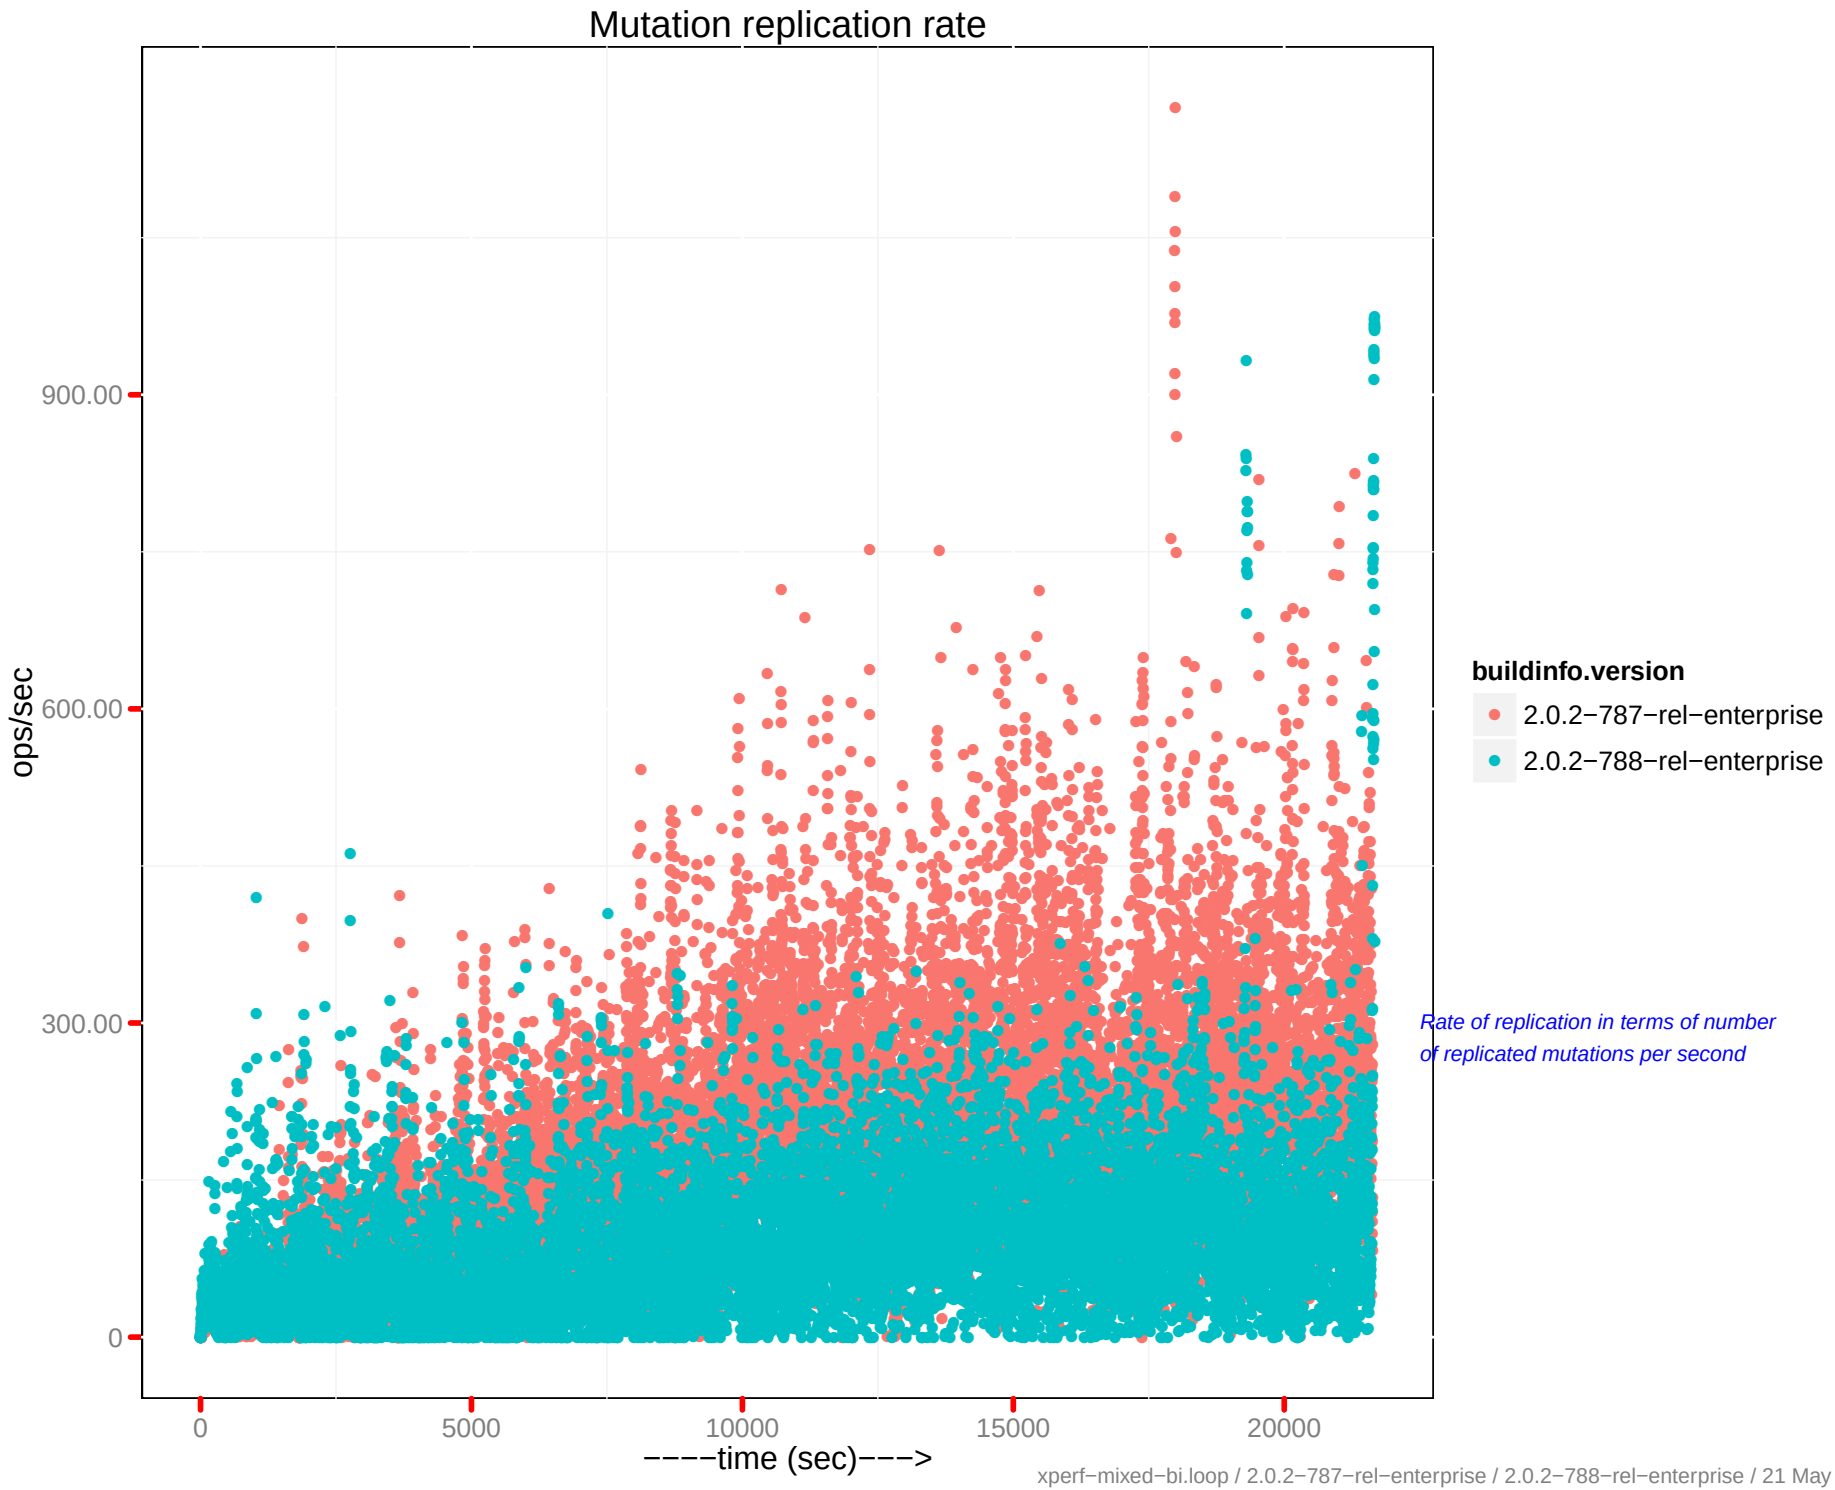
Mutation replication rate
ops/sec
900.00
600.00
300.00
0
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Rate of replication in terms of number
of replicated mutations per second
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

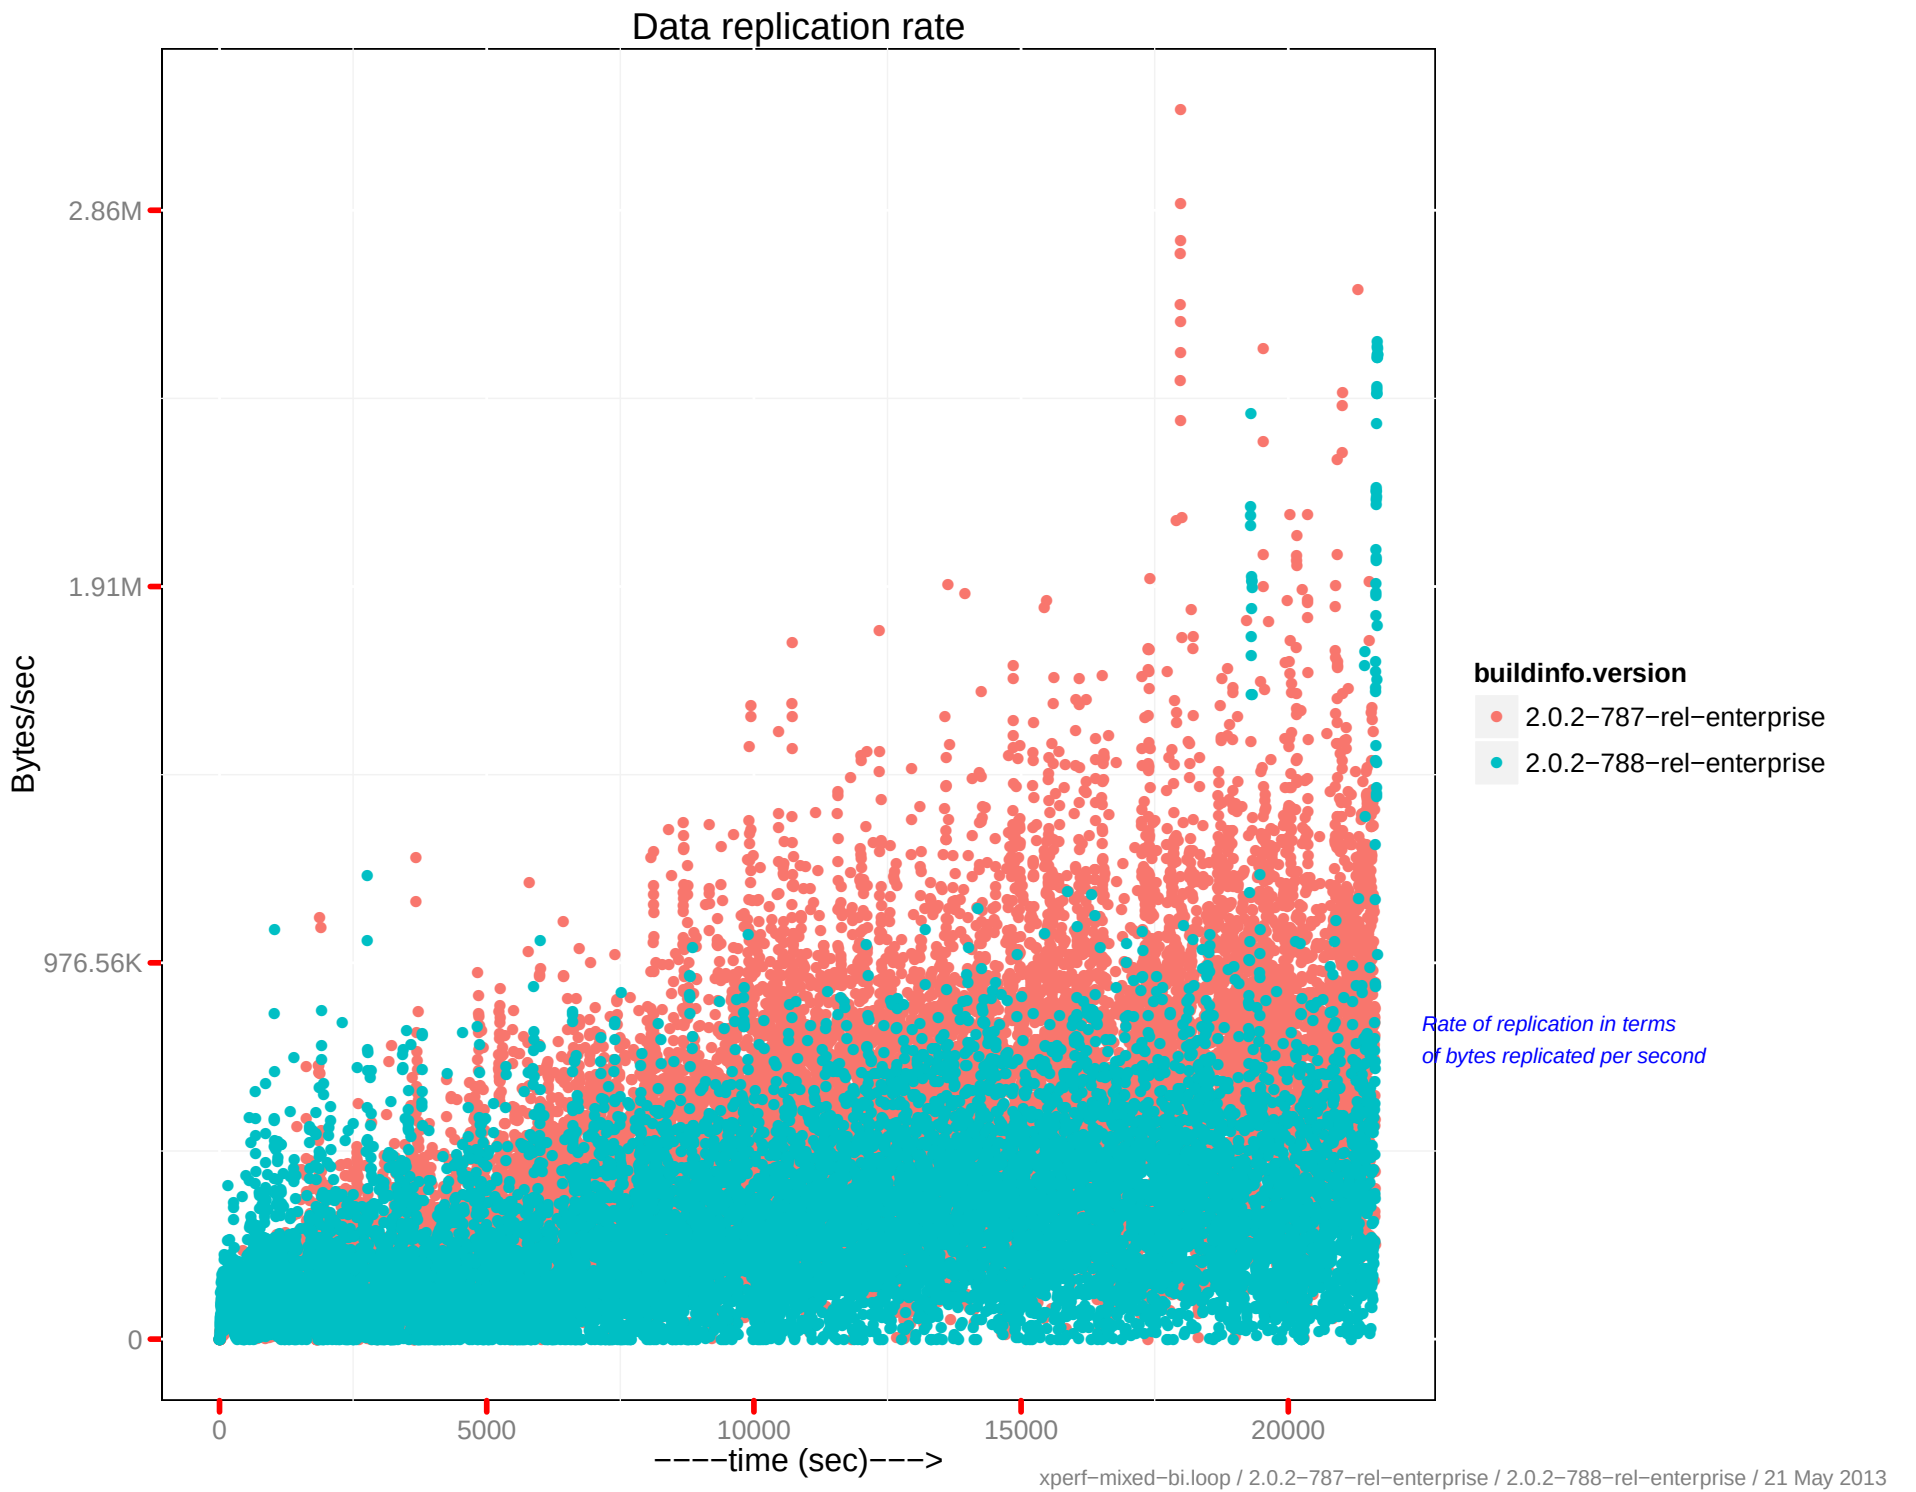
Data replication rate
2.86M
1.91M
976.56K
0
Bytes/sec
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Rate of replication in terms
of bytes replicated per second
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013


ms doc ops latency
7.63M
5.72M
3.81M
1.91M
0
ms
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Weighted average latency in ms of
sending replicated mutations to remote cluster
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013


ms meta ops latency
9.54M
7.15M
4.77M
2.38M
0
ms
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Weighted average latency in ms of sending getMeta and waiting for conflict solution result from remote cluster
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

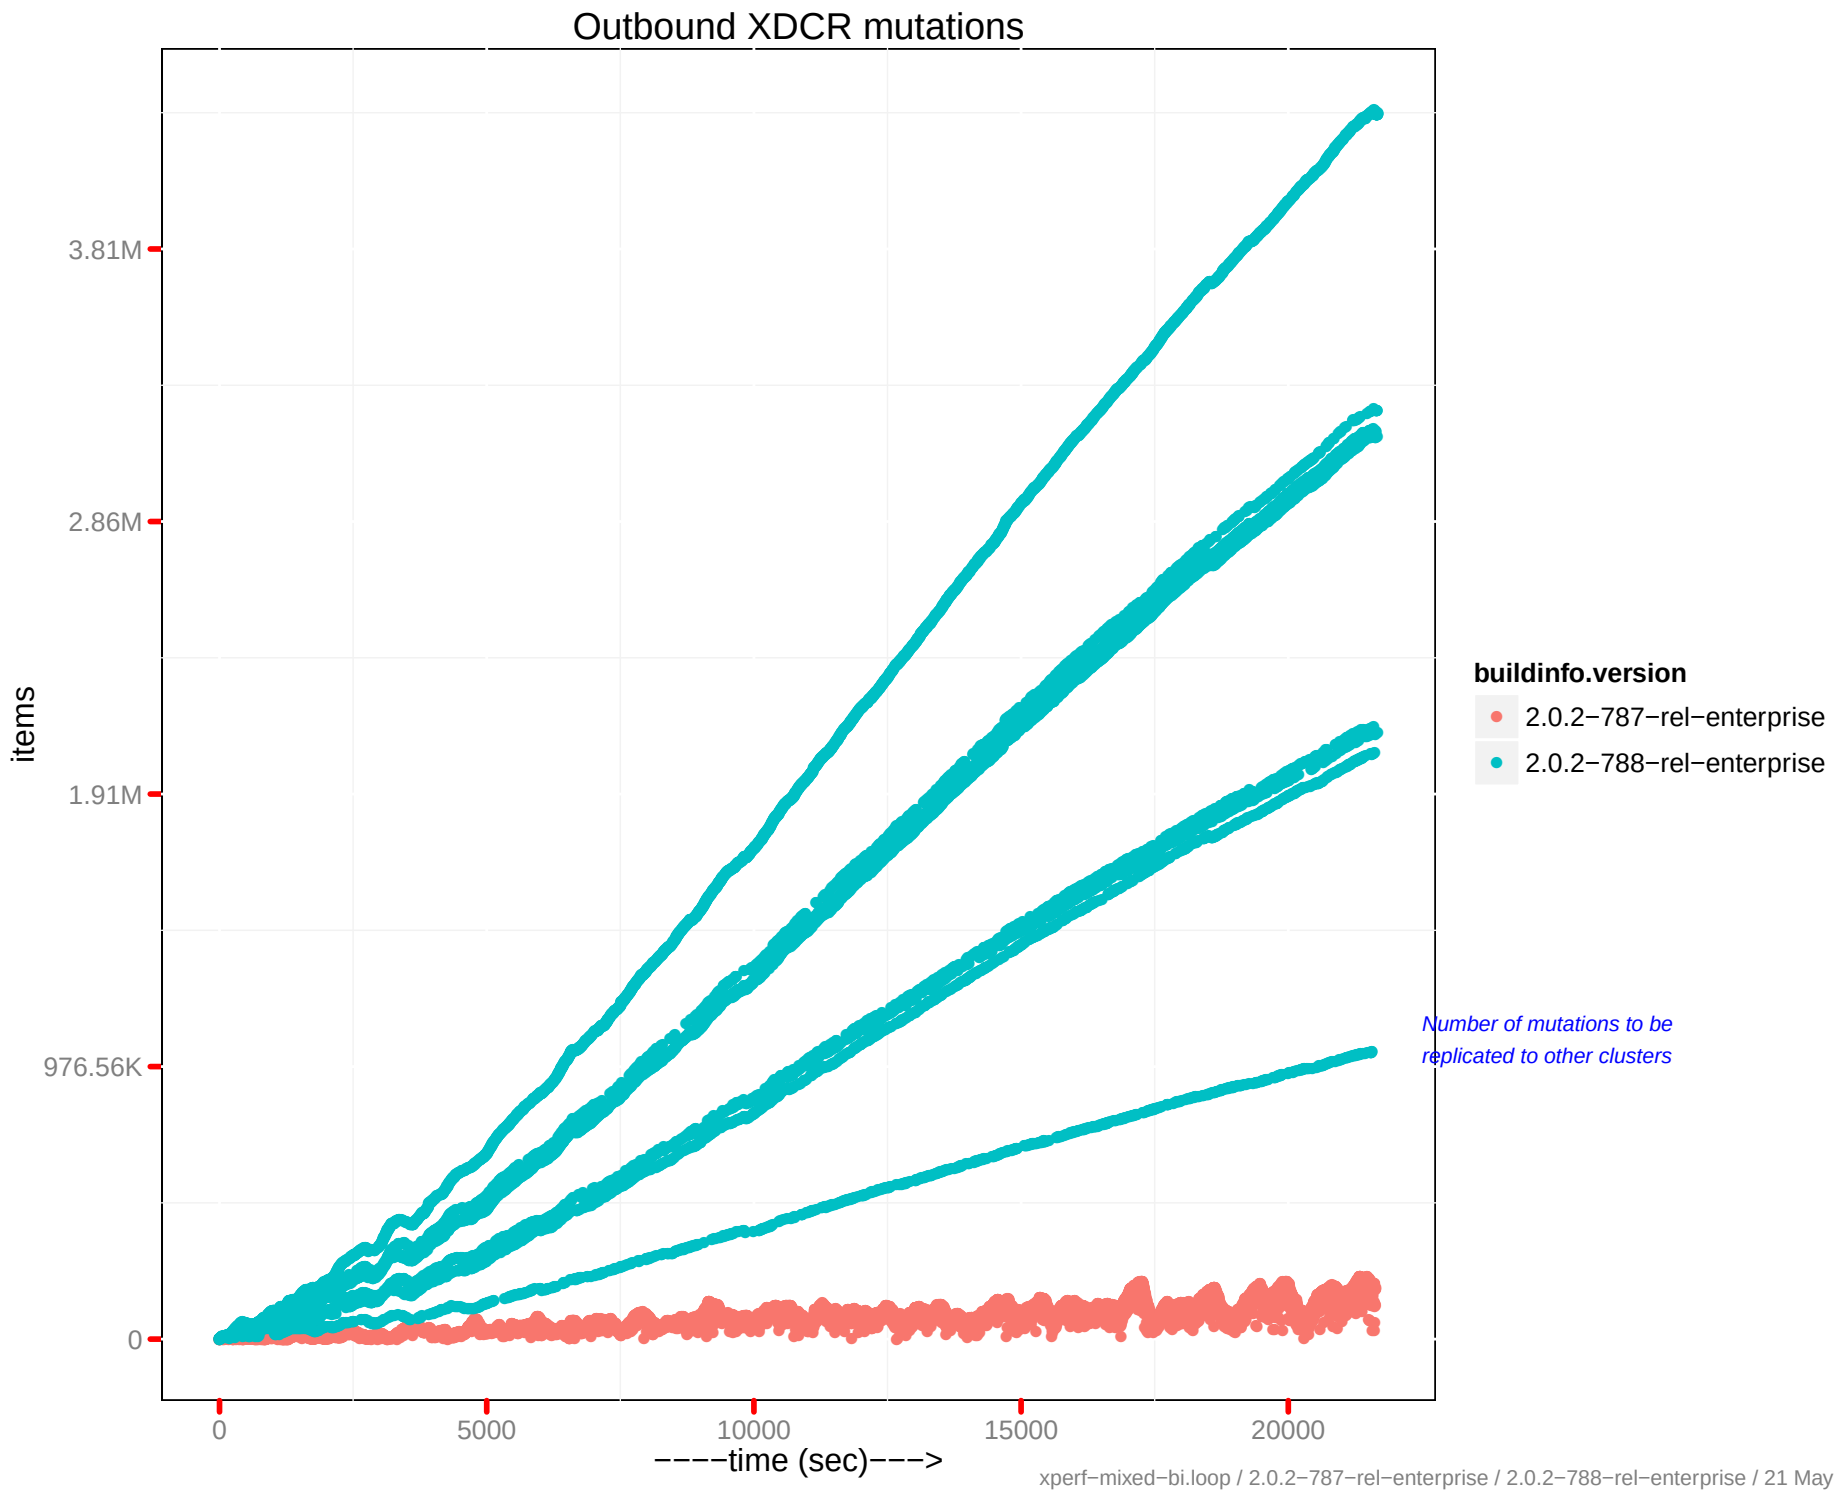
Outbound XDCR mutations
items
3.81M
2.86M
1.91M
976.56K
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of mutations to be replicated to other clusters
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

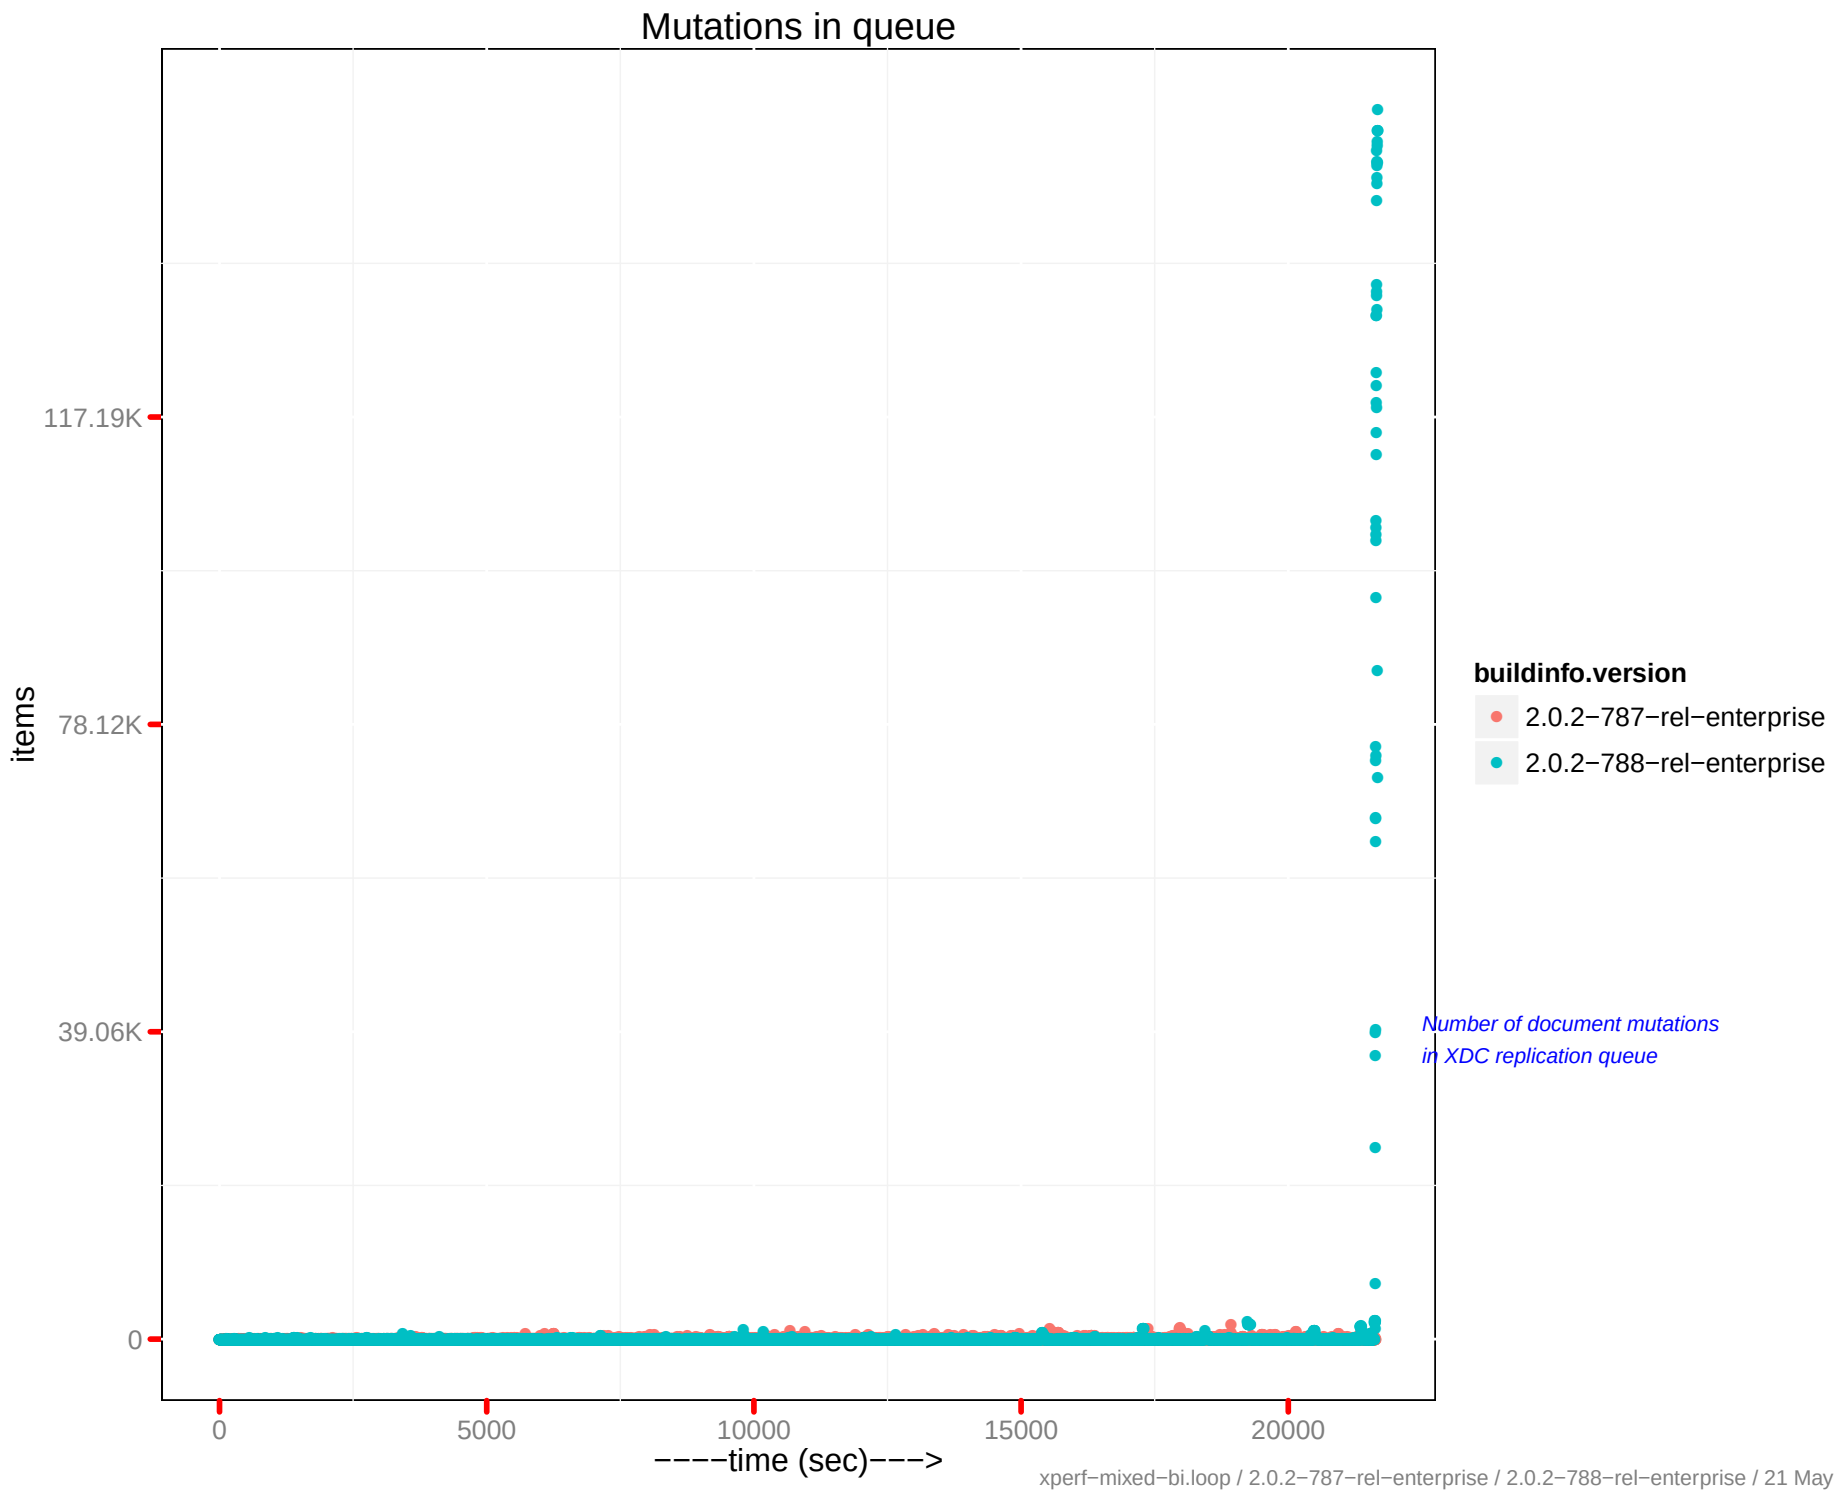
Mutations in queue
items
117.19K
78.12K
39.06K
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of document mutations
in XDC replication queue
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

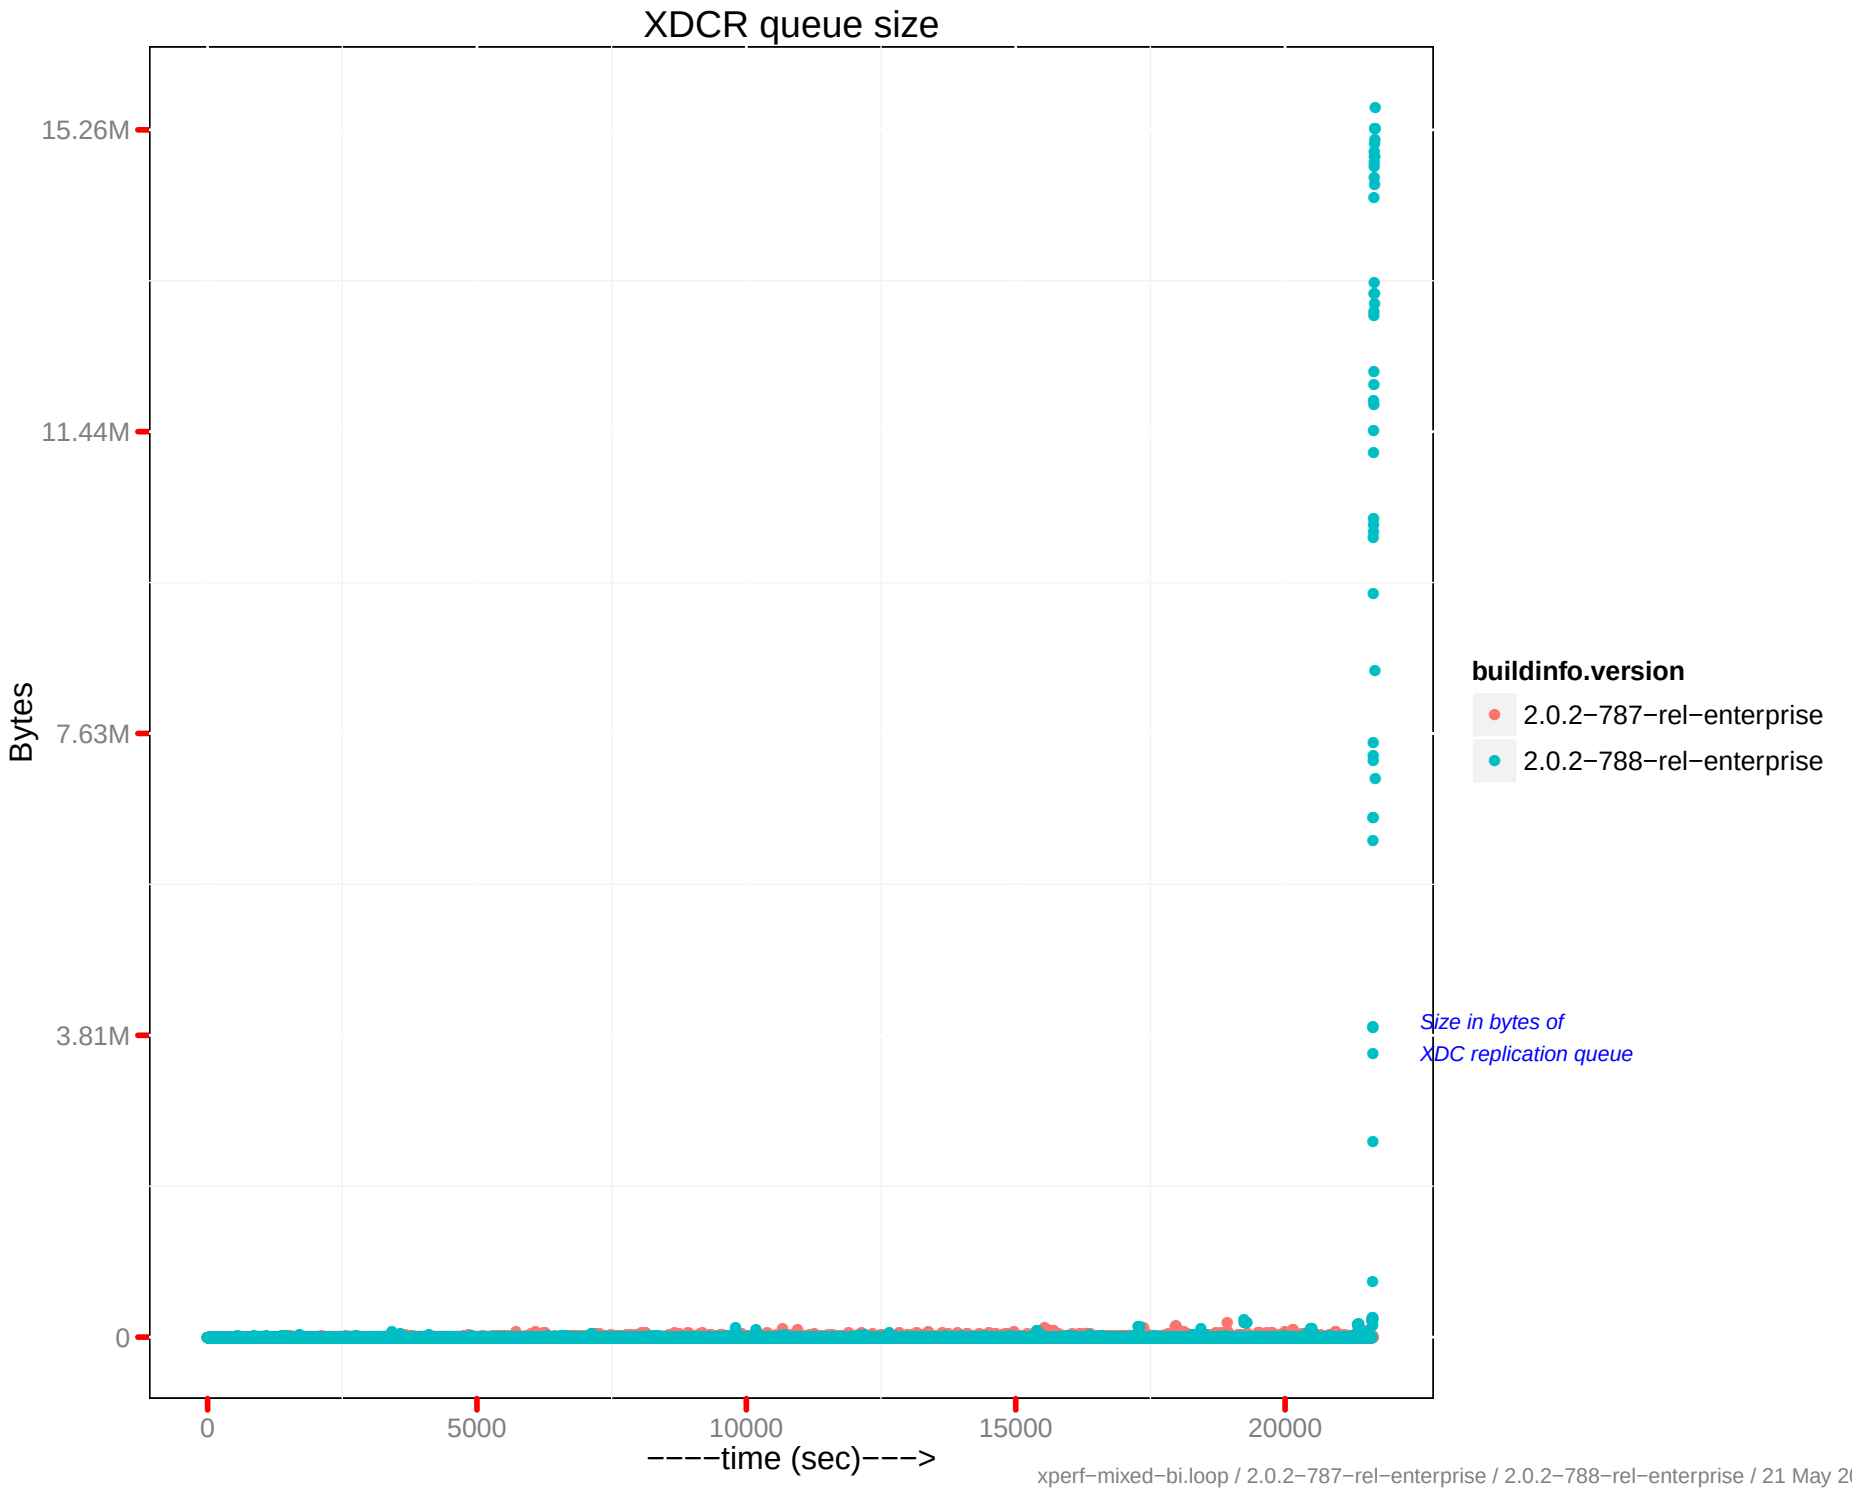
XDCR queue size
15.26M
11.44M
7.63M
3.81M
0
Bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Size in bytes of
XDC replication queue
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

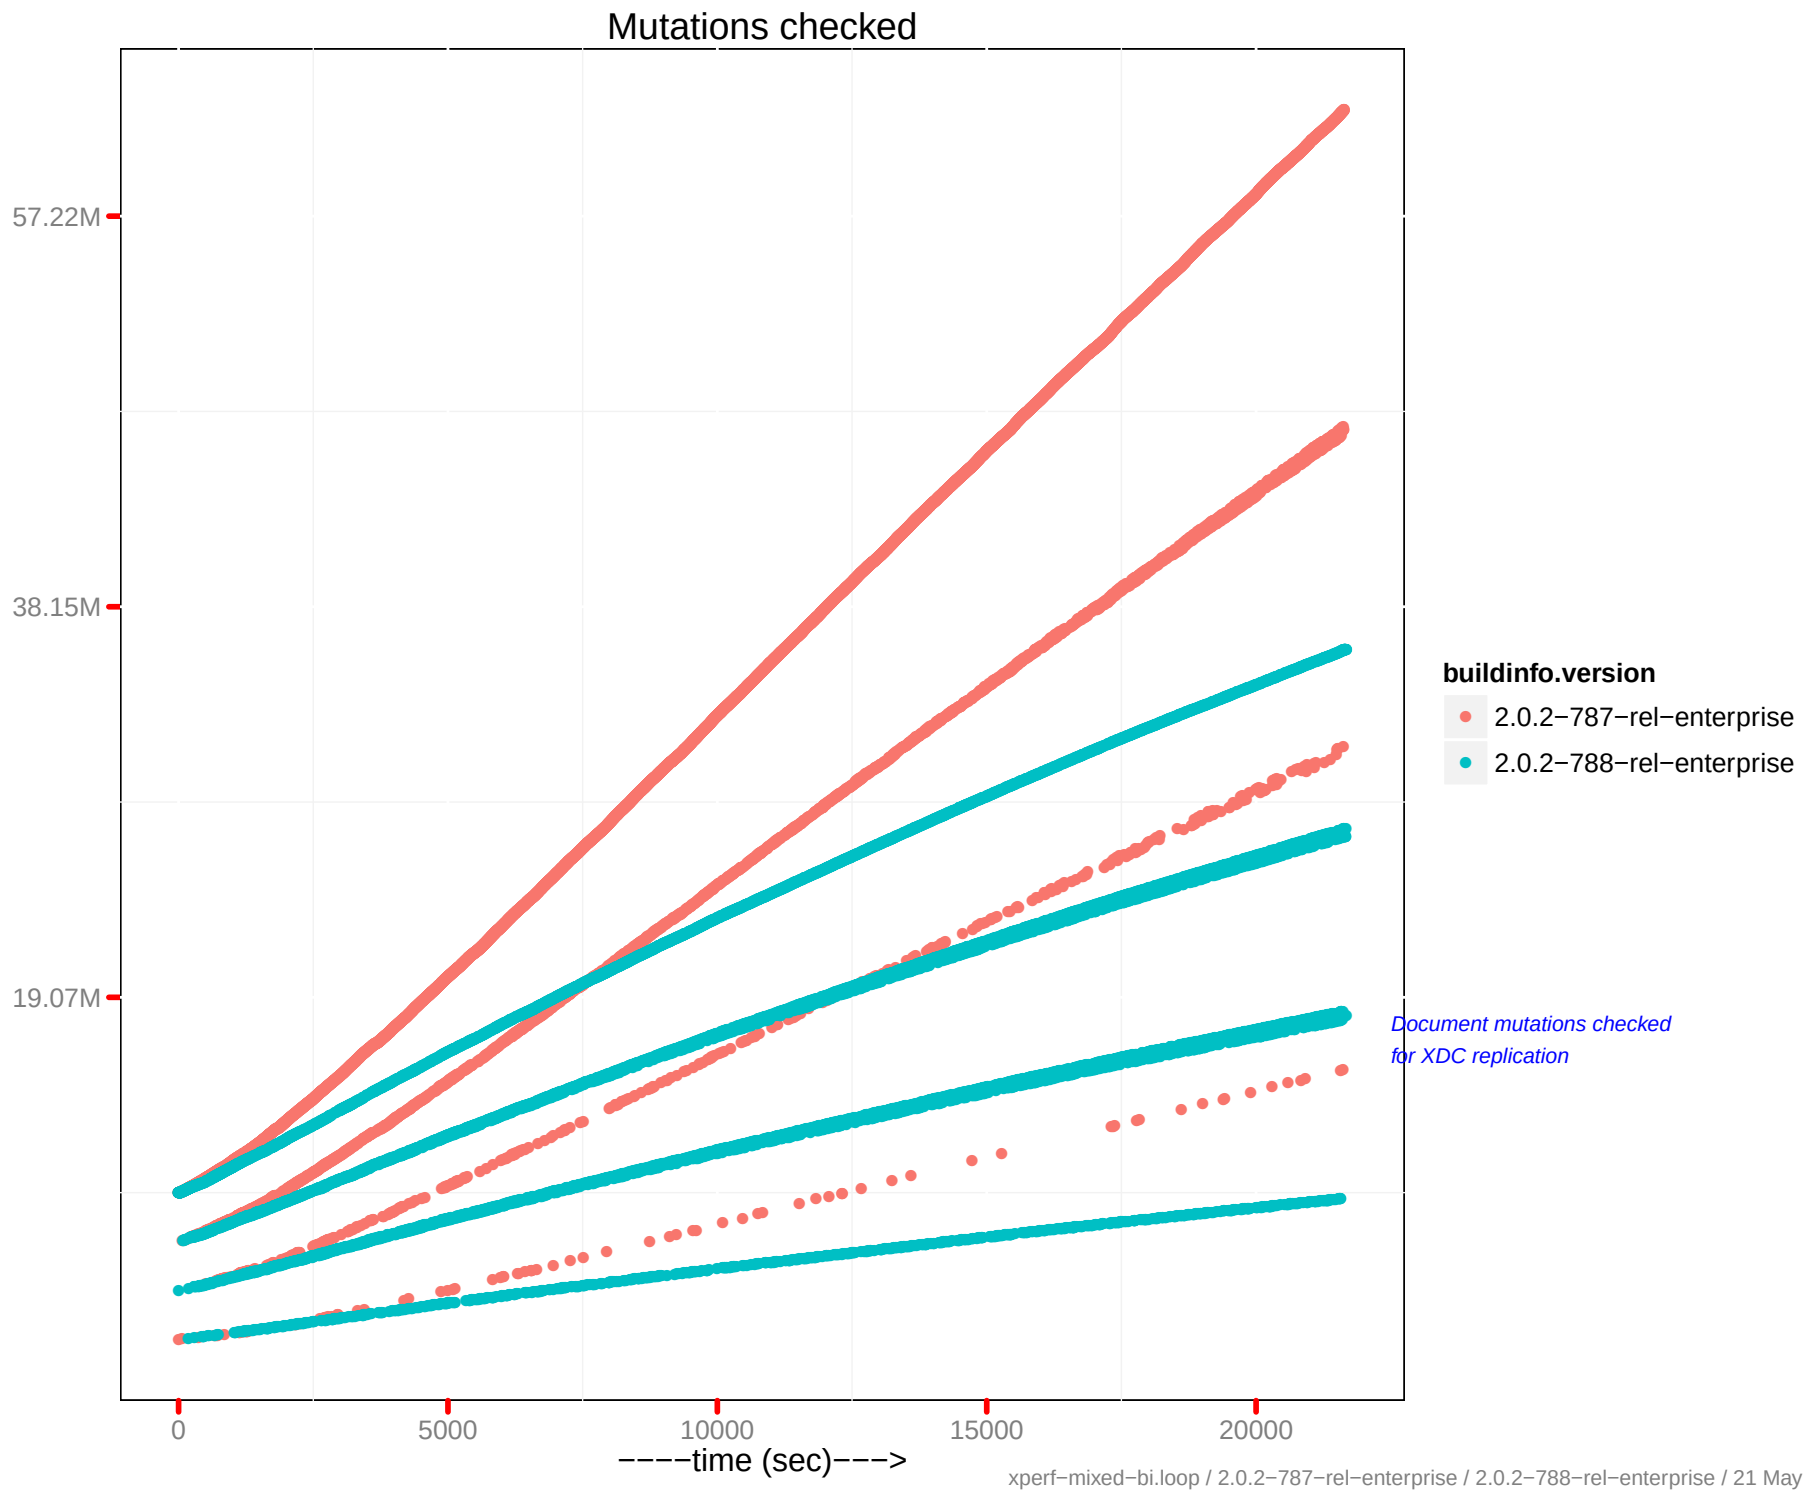
Mutations checked
57.22M
38.15M
19.07M
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Document mutations checked
for XDC replication
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

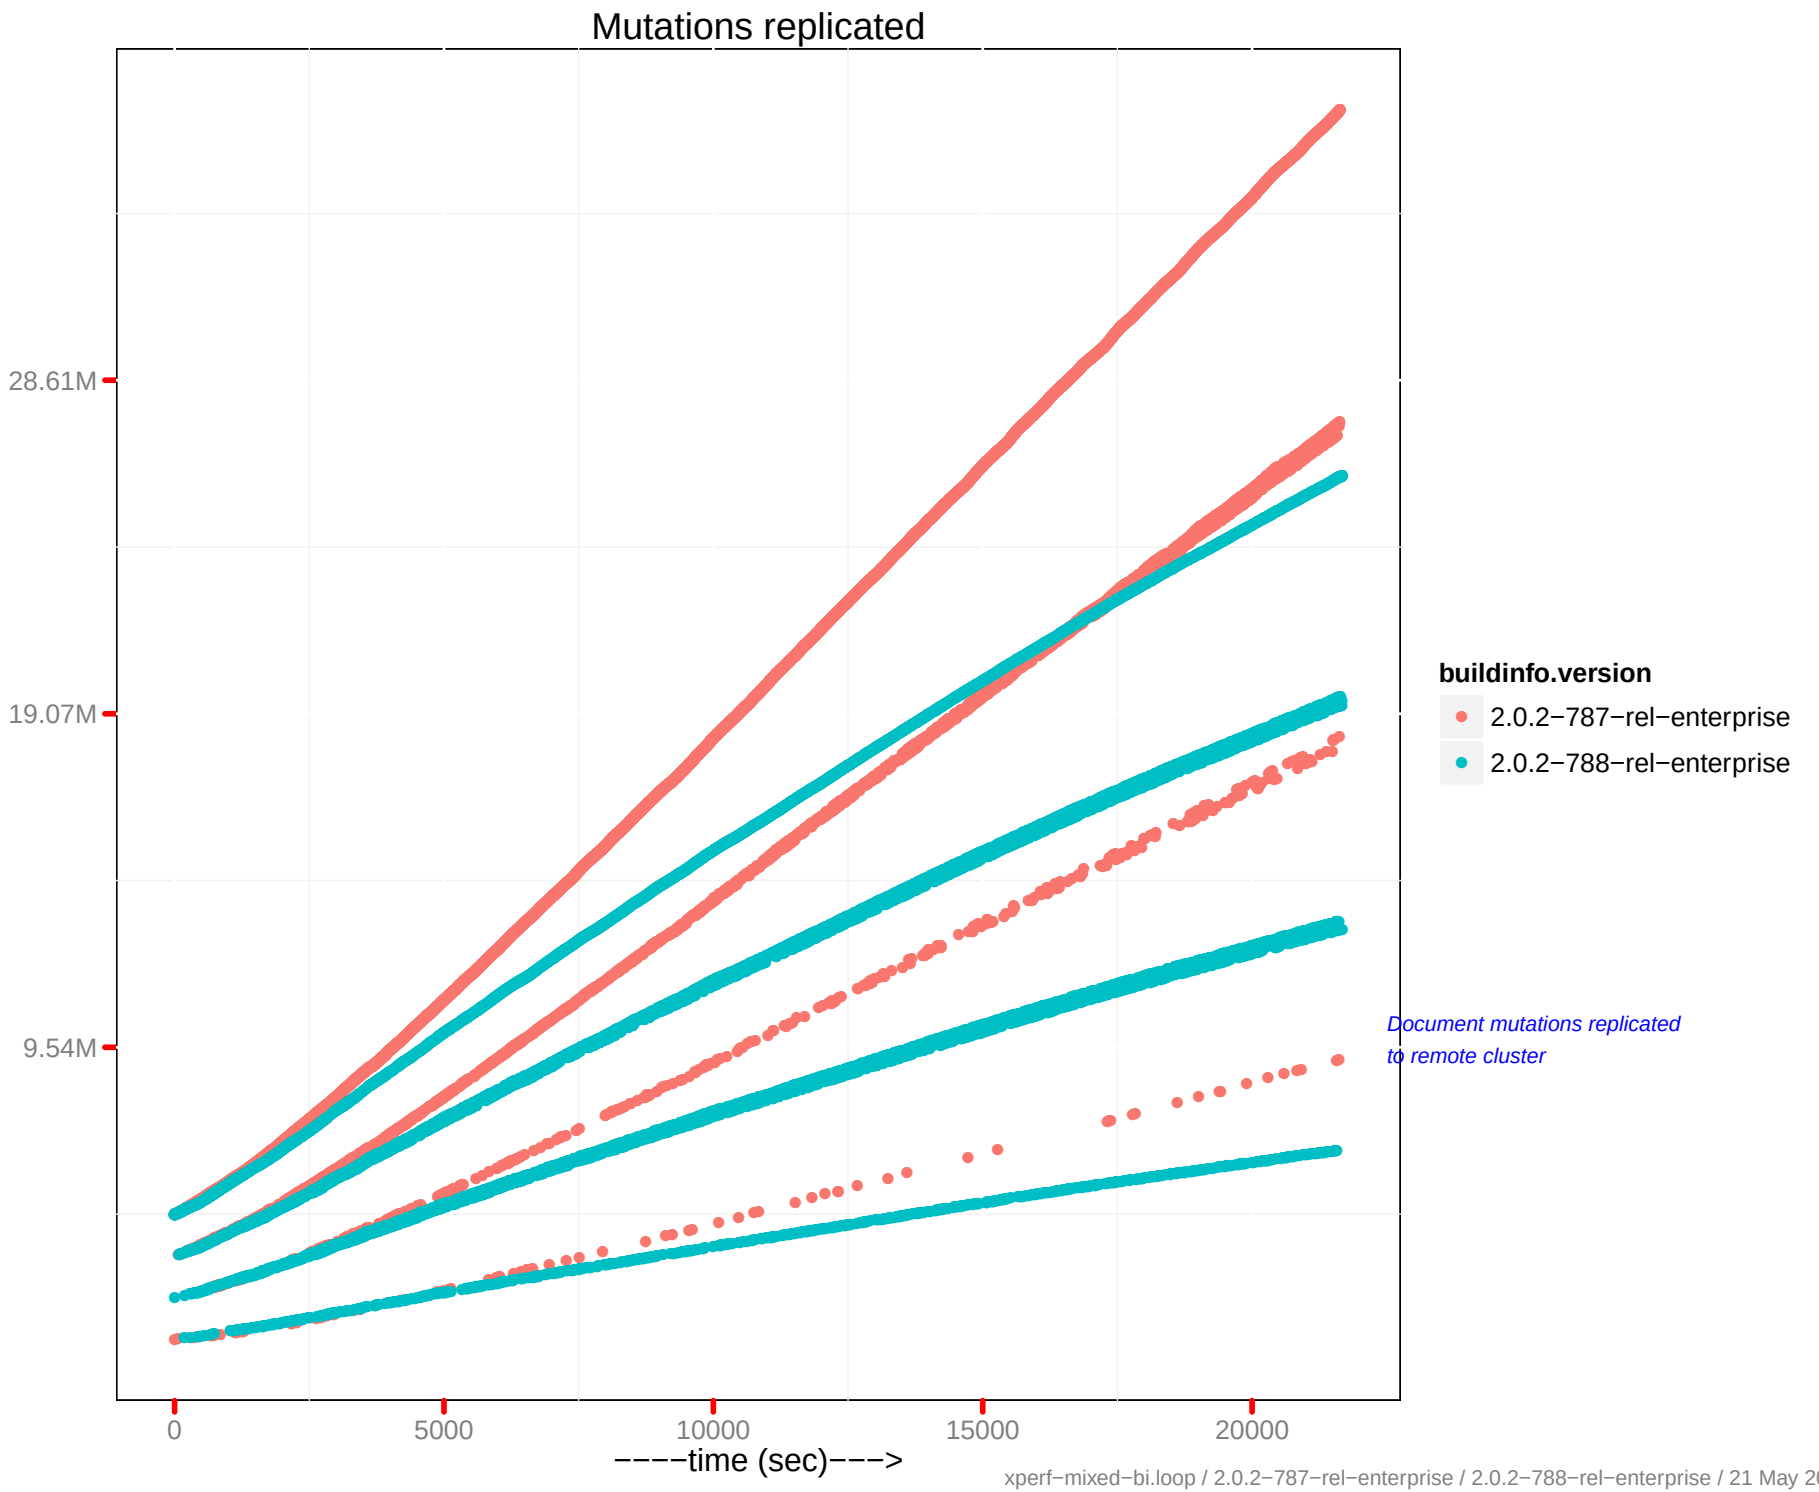
Mutations replicated
28.61M
19.07M
9.54M
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Document mutations replicated
to remote cluster
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

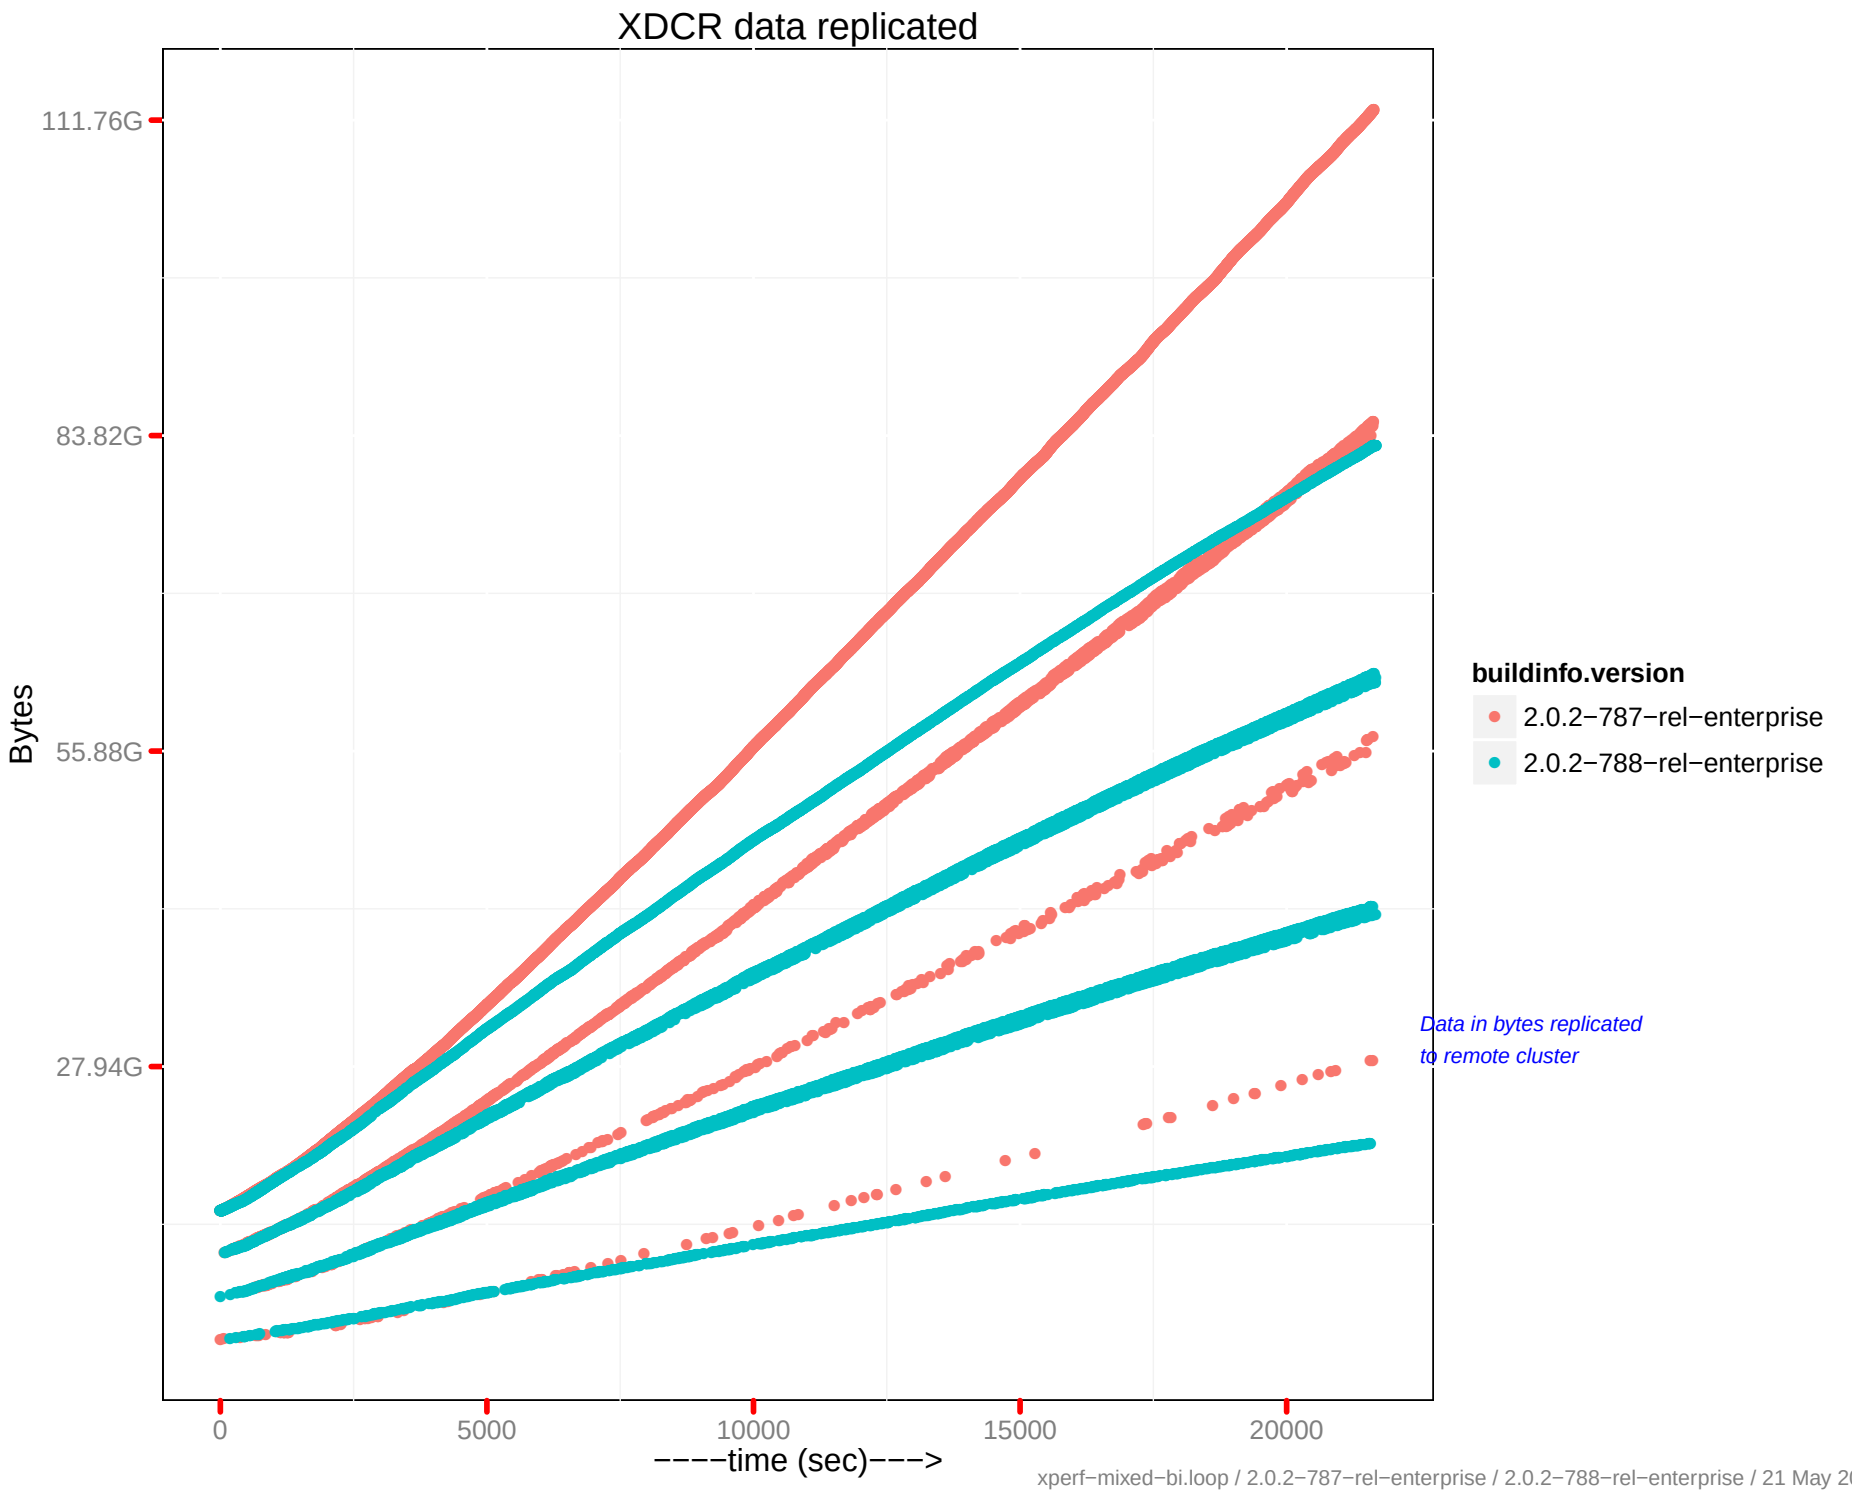
XDCR data replicated
111.76G
83.82G
55.88G
27.94G
Bytes
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Data in bytes replicated
to remote cluster
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

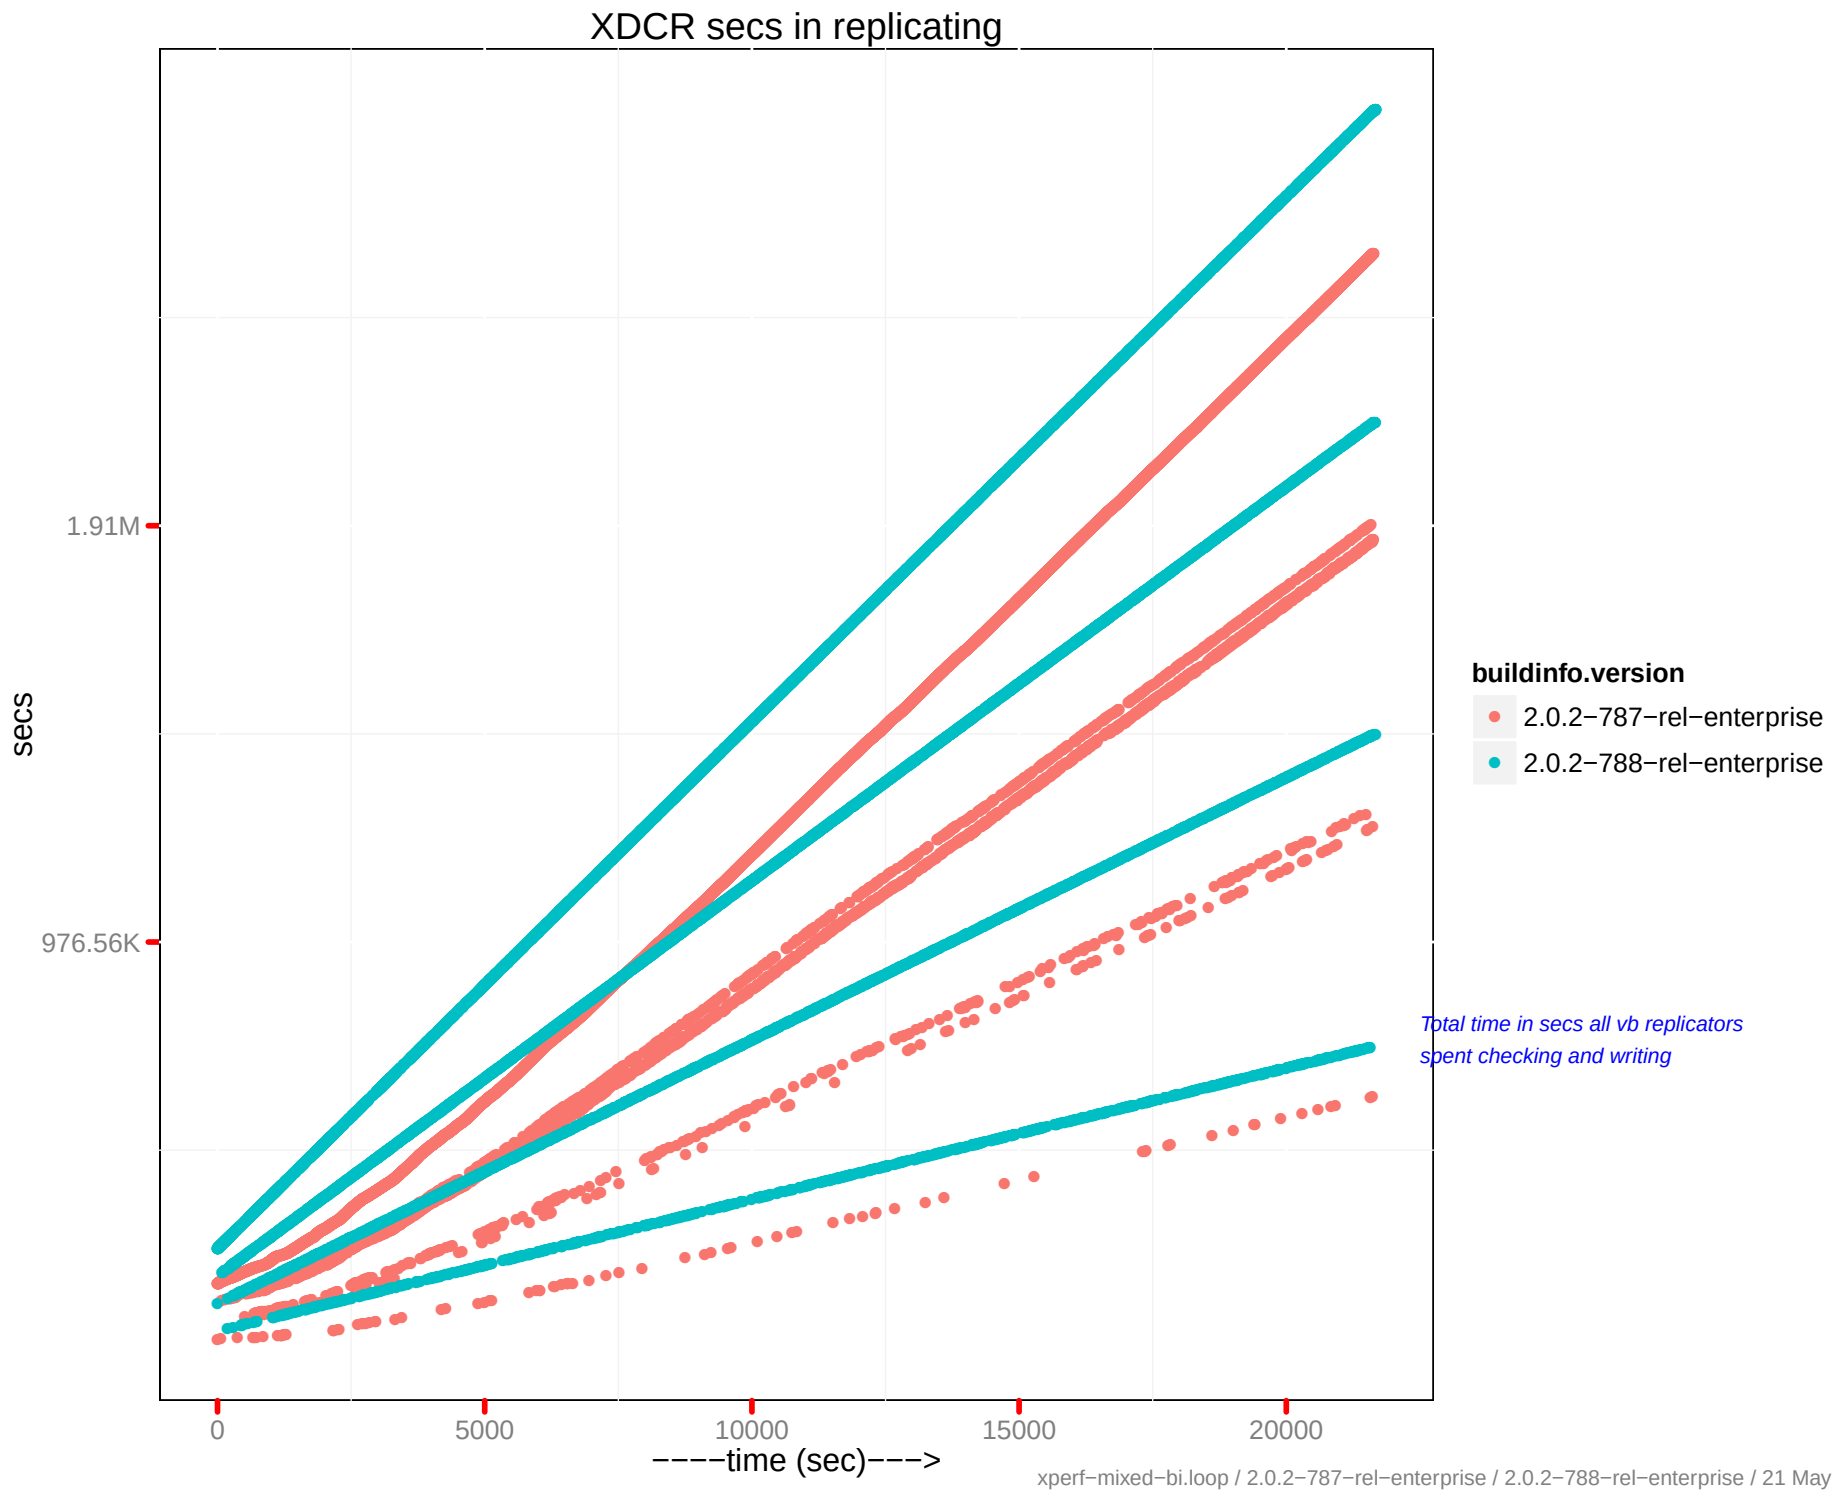
XDCR secs in replicating
secs
1.91M
976.56K
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Total time in secs all vb replicators
spent checking and writing
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

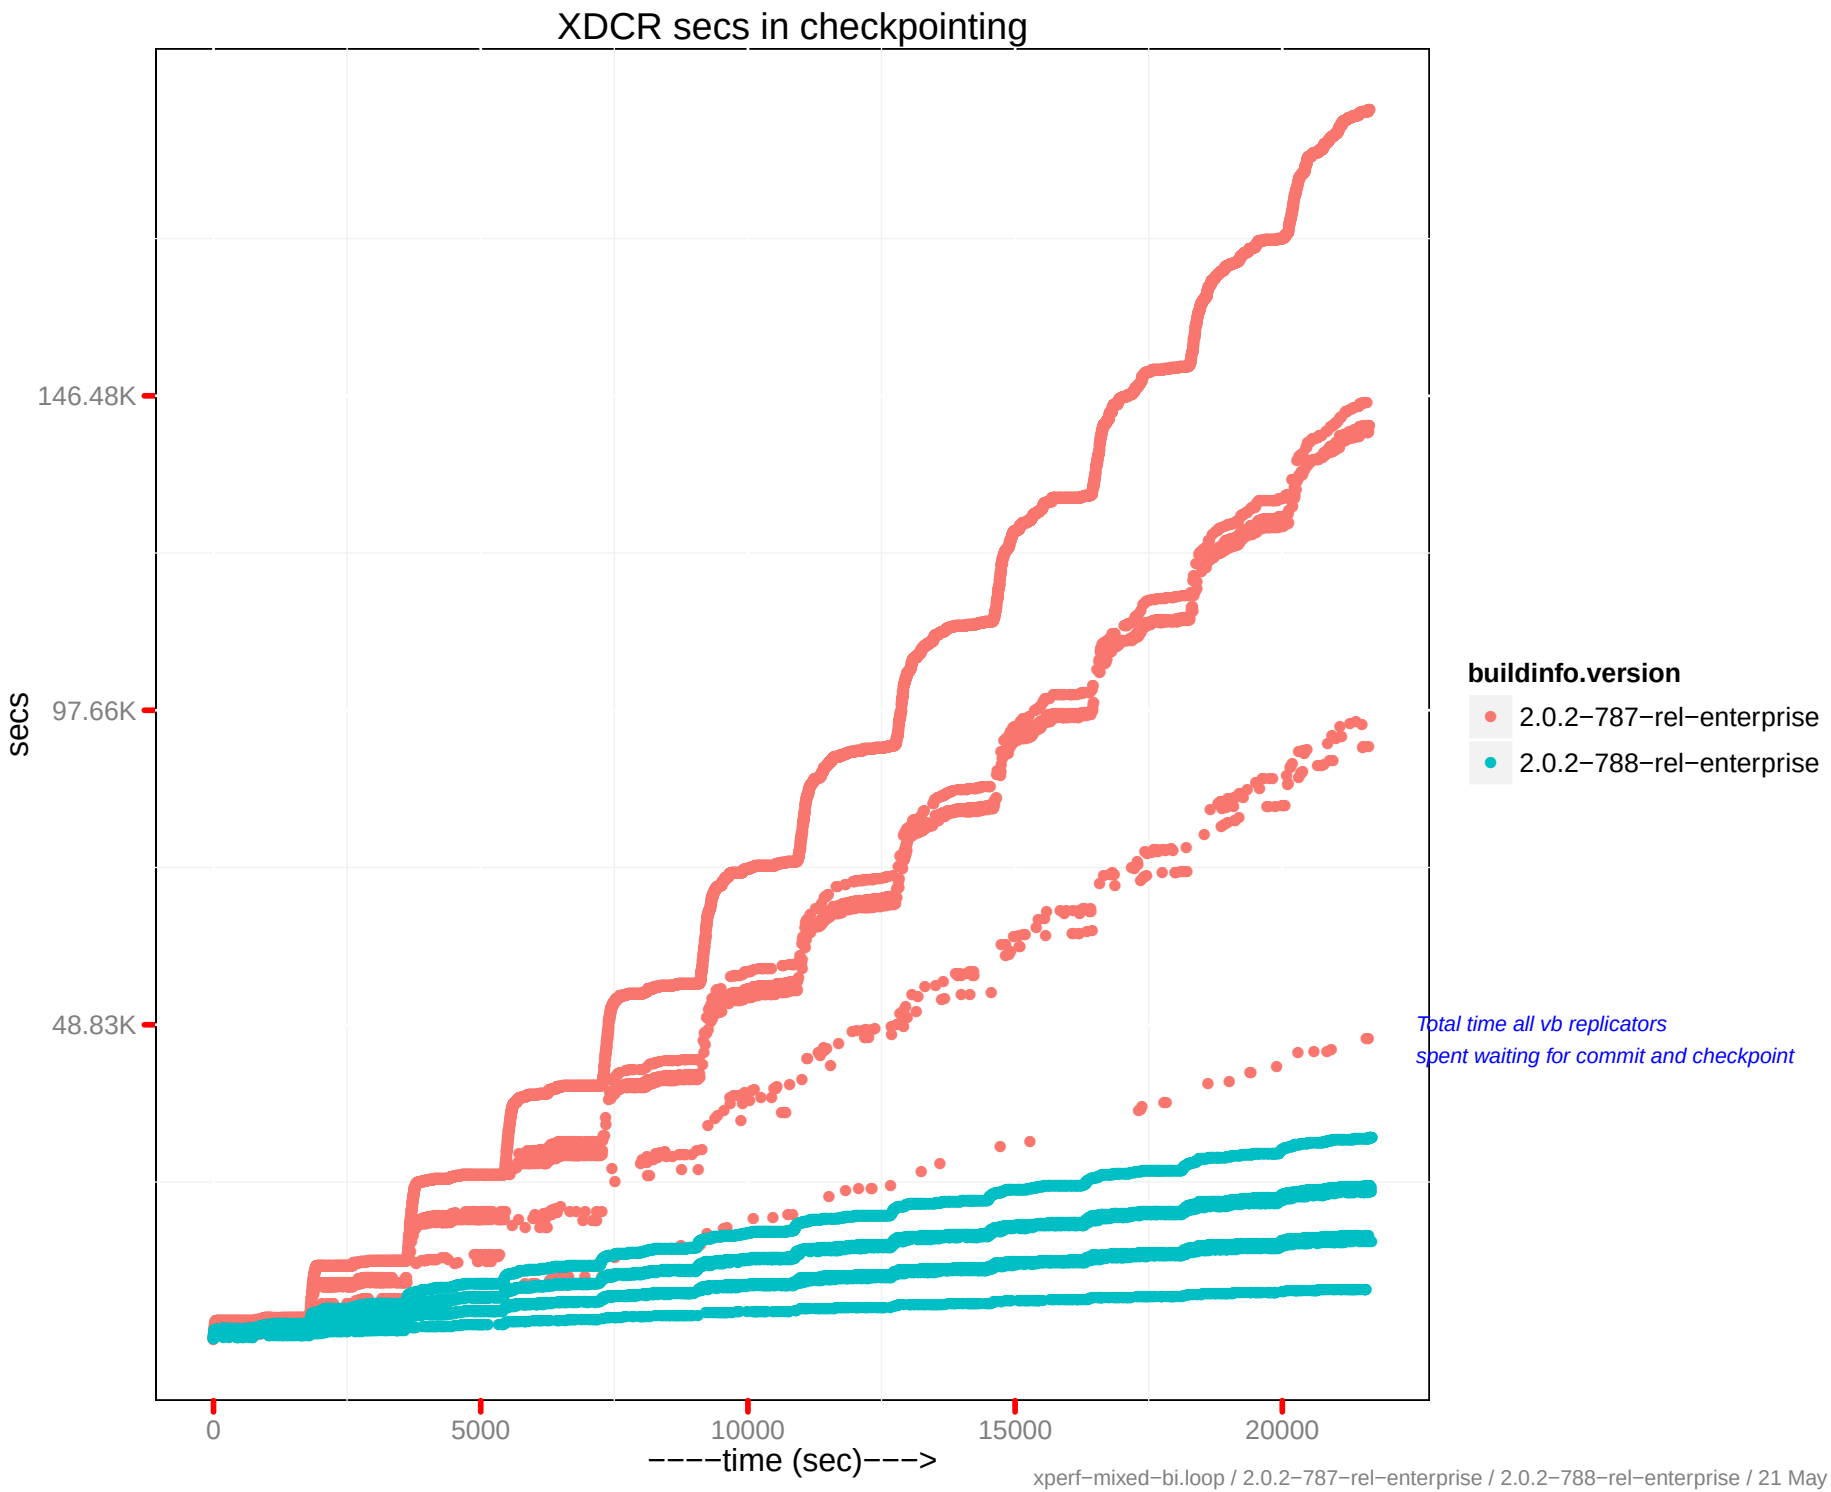
XDCR secs in checkpointing
146.48K
97.66K
48.83K
secs
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Total time all vb replicators
spent waiting for commit and checkpoint
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

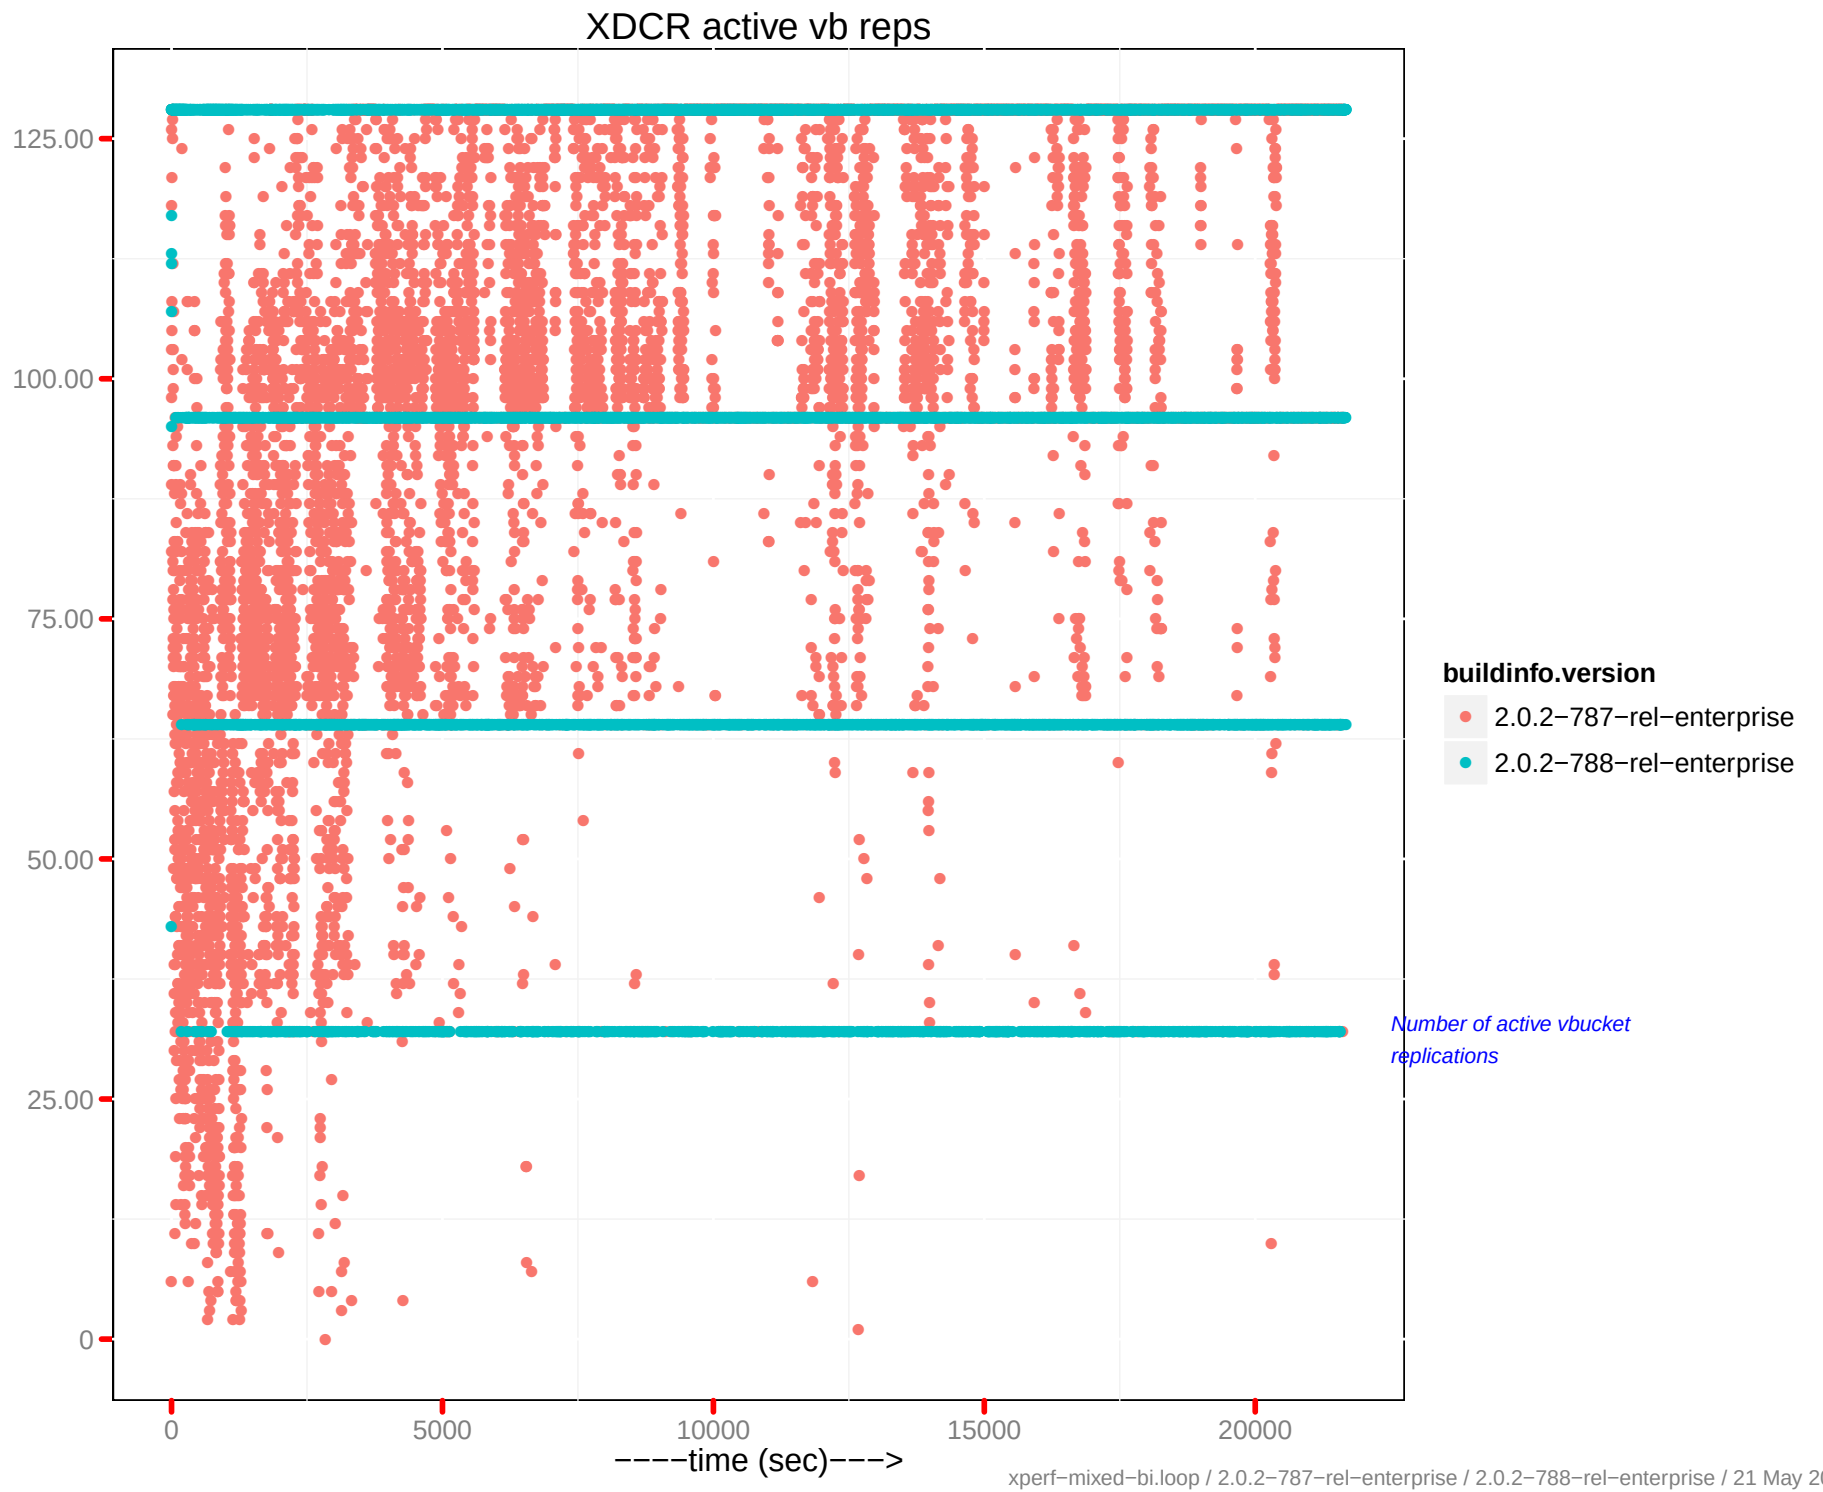
XDCR active vb reps
125.00
100.00
75.00
50.00
25.00
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of active vbucket replications
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

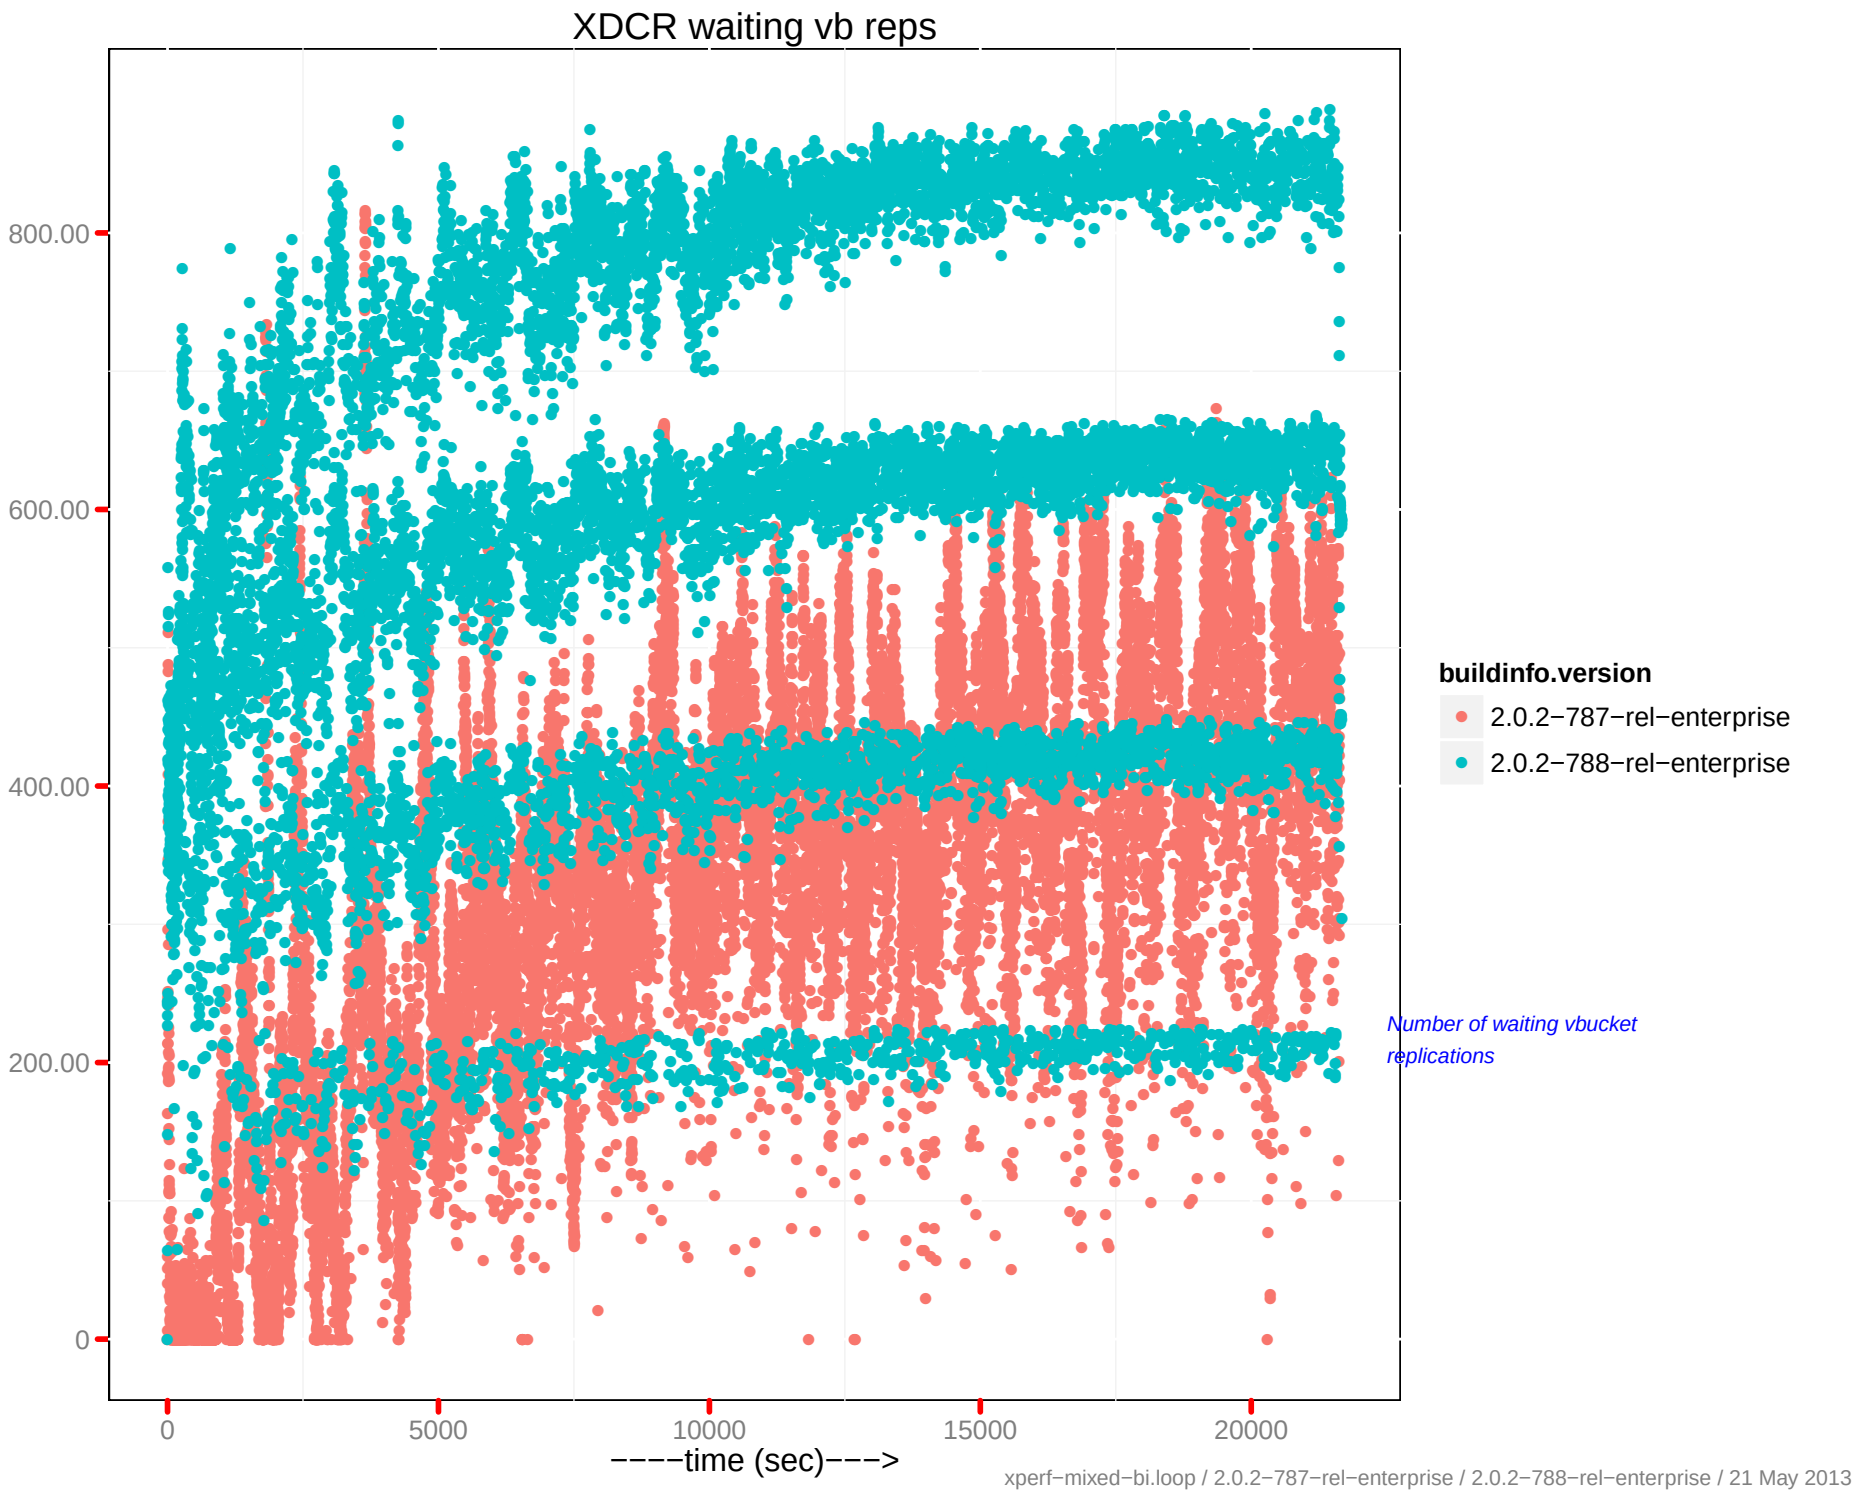
XDCR waiting vb reps
800.00
600.00
400.00
200.00
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of waiting vbucket replications
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

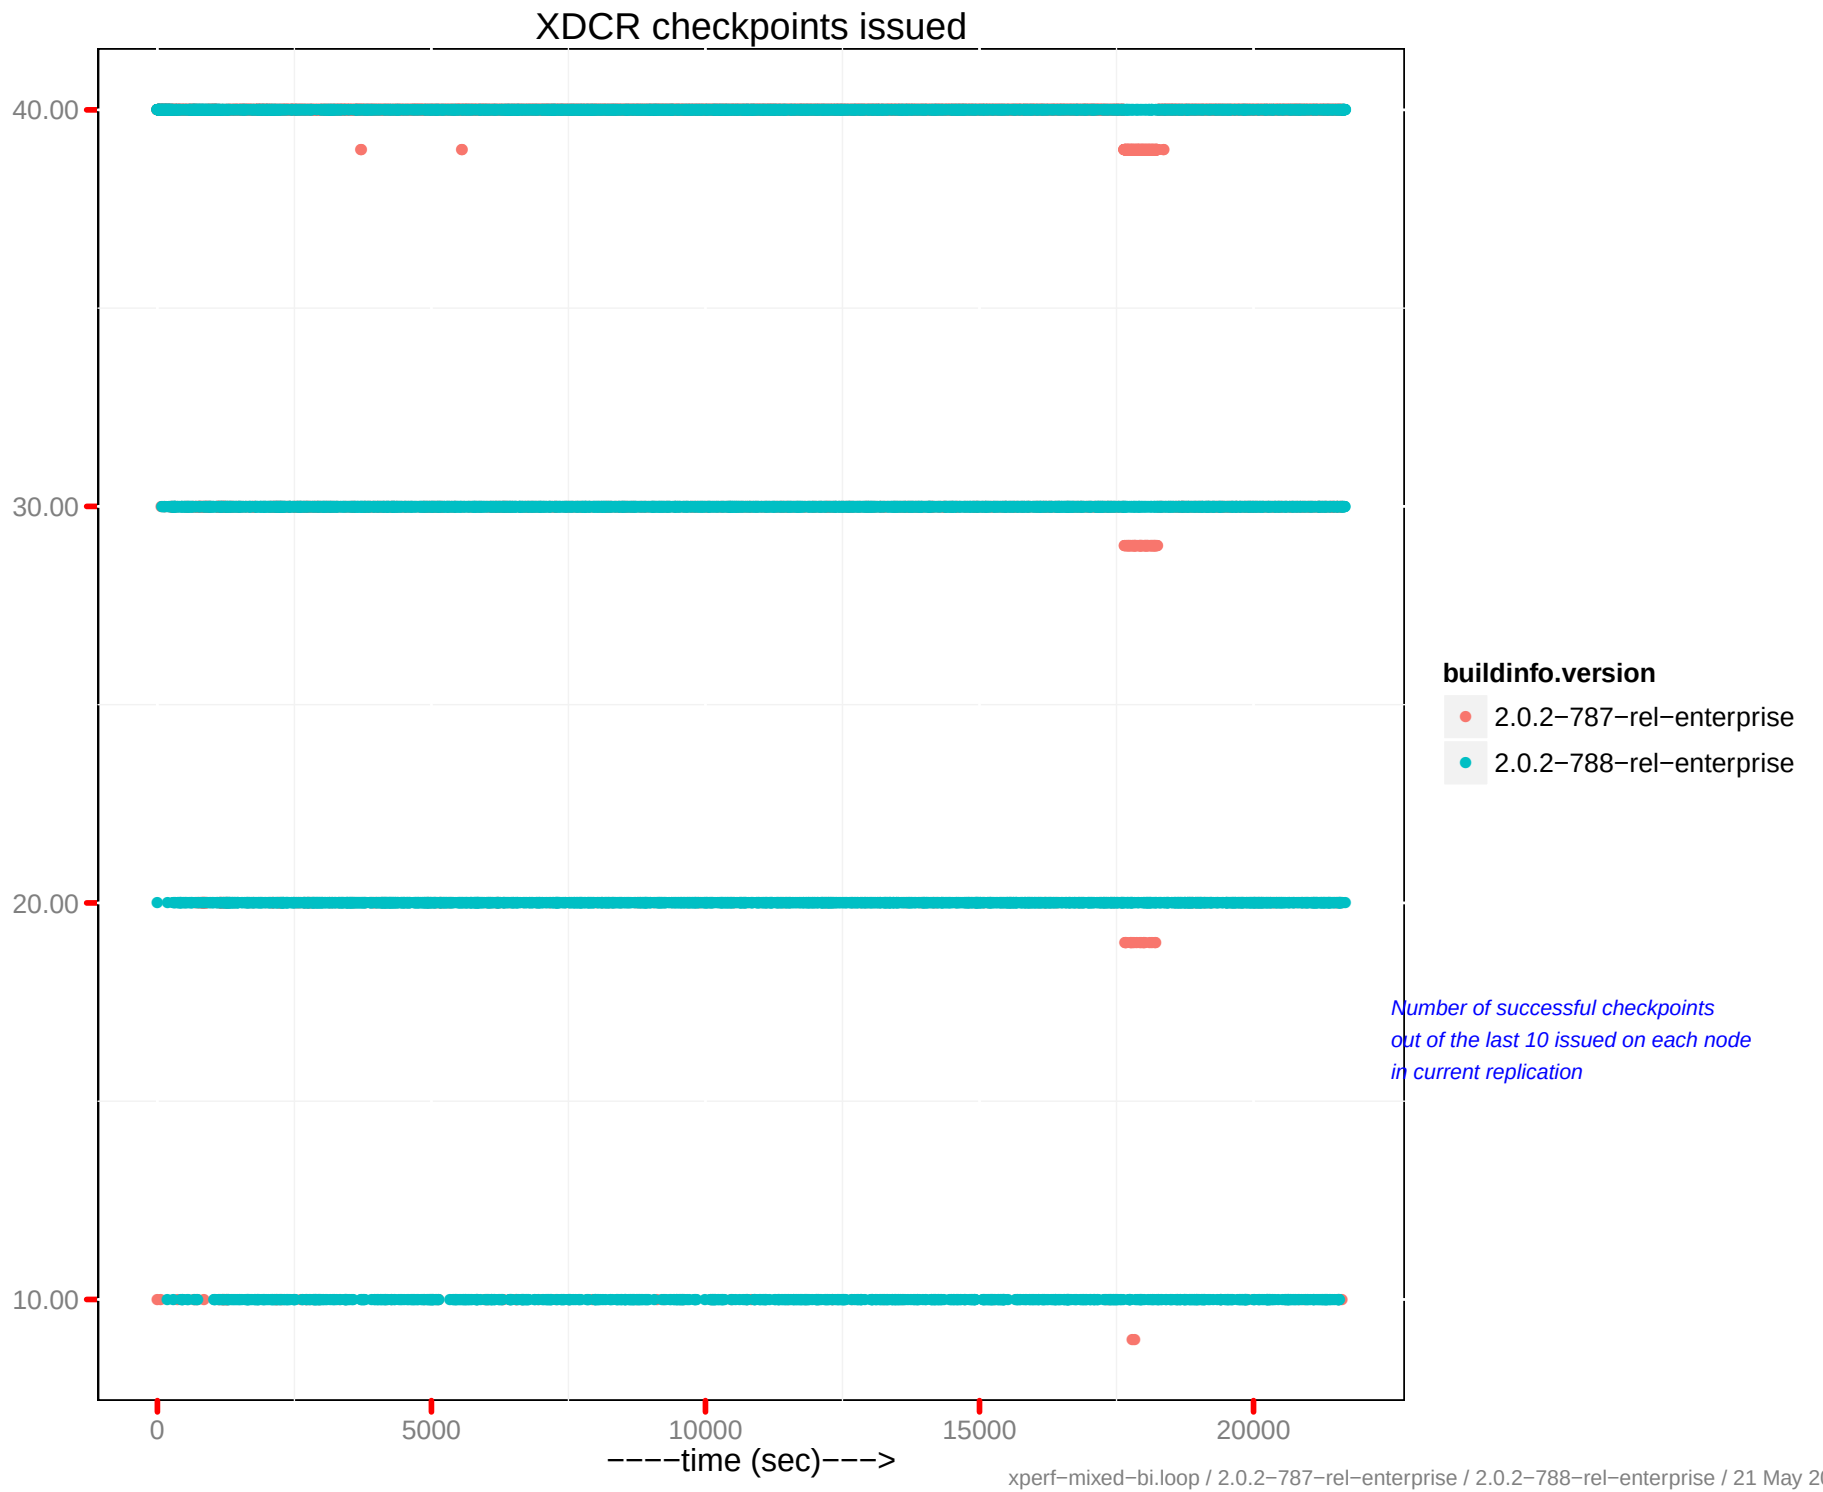
XDCR checkpoints issued
40.00
30.00
20.00
10.00
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of successful checkpoints
out of the last 10 issued on each node
in current replication
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

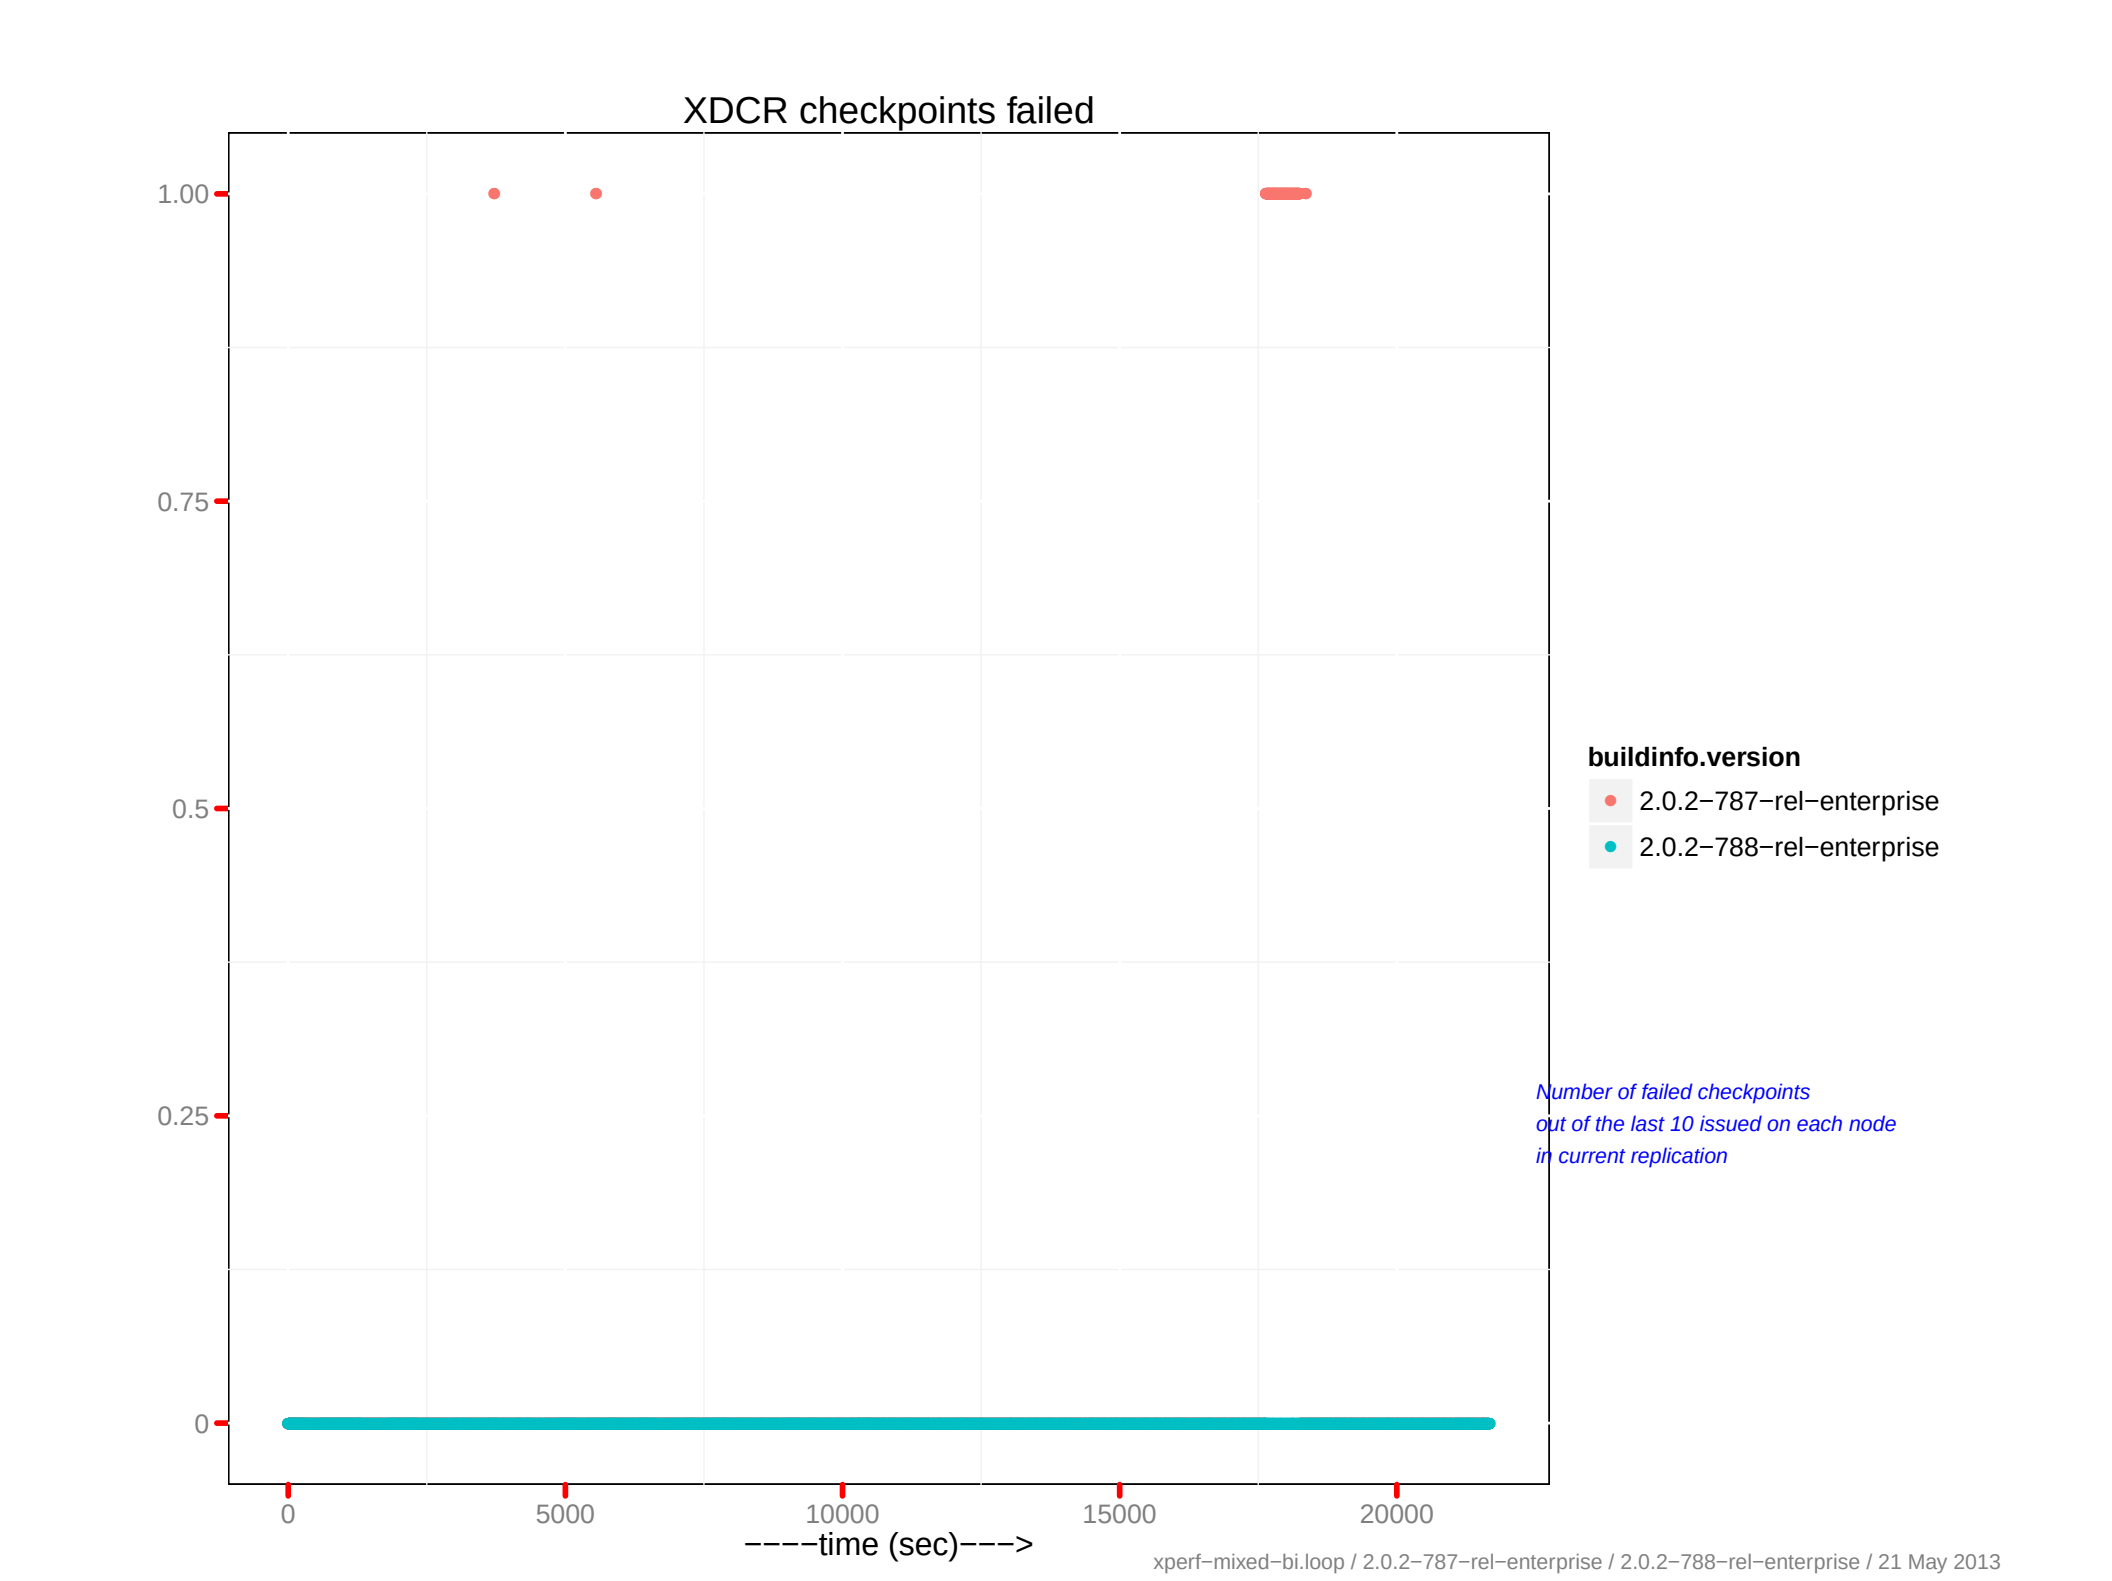
XDCR checkpoints failed
1.00
0.75
0.5
0.25
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of failed checkpoints
out of the last 10 issued on each node
in current replication
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

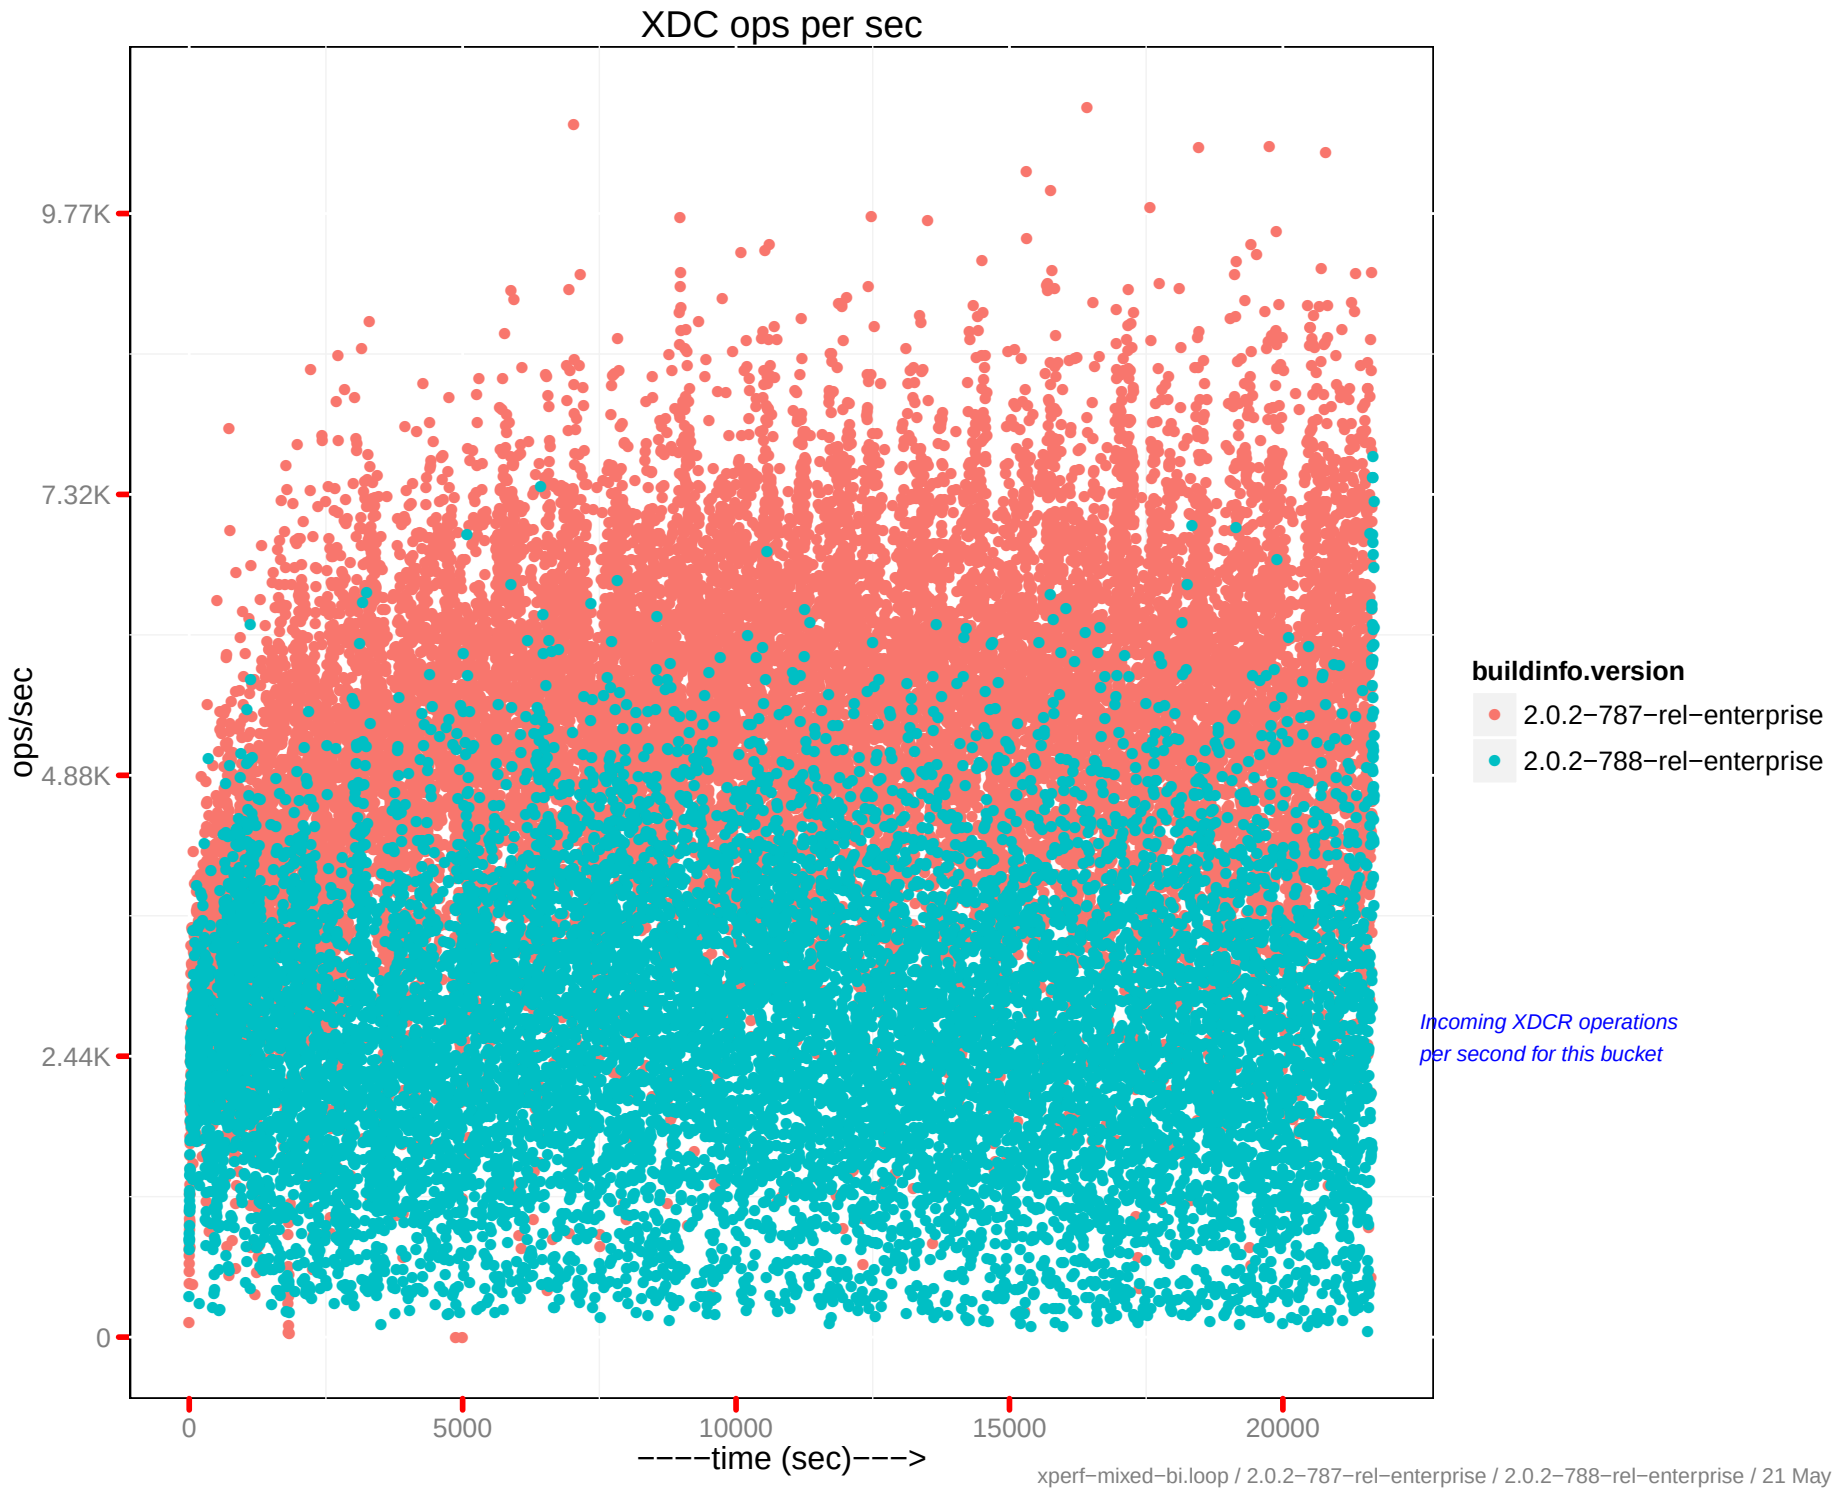
XDC ops per sec
9.77K
7.32K
4.88K
2.44K
0
ops/sec
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Incoming XDCR operations
per second for this bucket
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

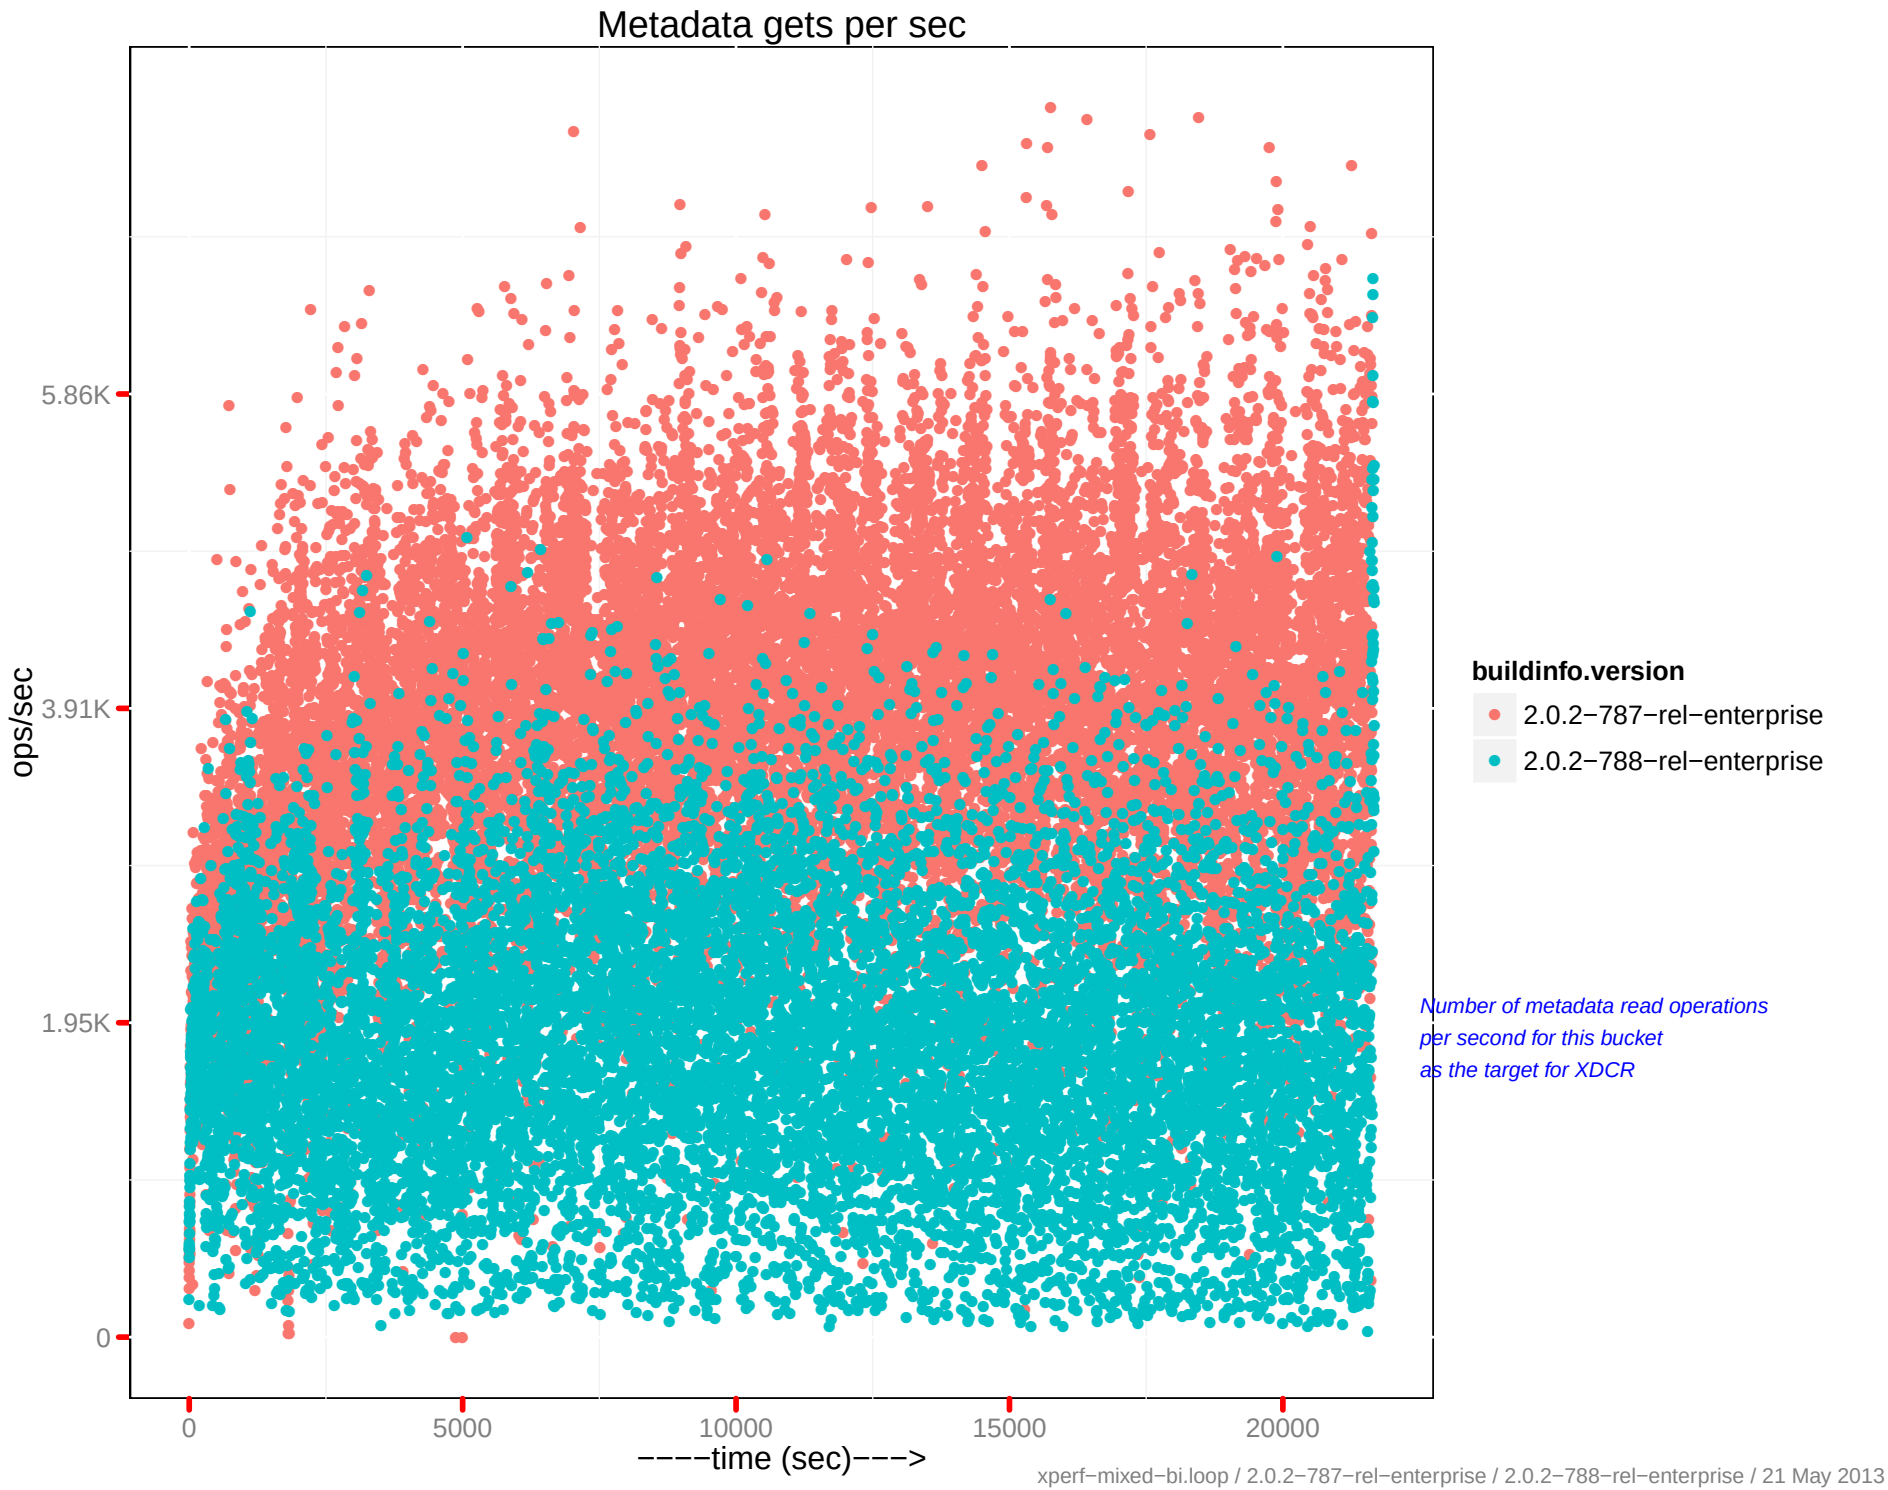
Metadata gets per sec
ops/sec
5.86K
3.91K
1.95K
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of metadata read operations
per second for this bucket
as the target for XDCR
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013


Metadata sets per sec
2.93K
1.95K
1000.00
0
ops/sec
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Number of set operations
per second for this bucket
as the target for XDCR
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

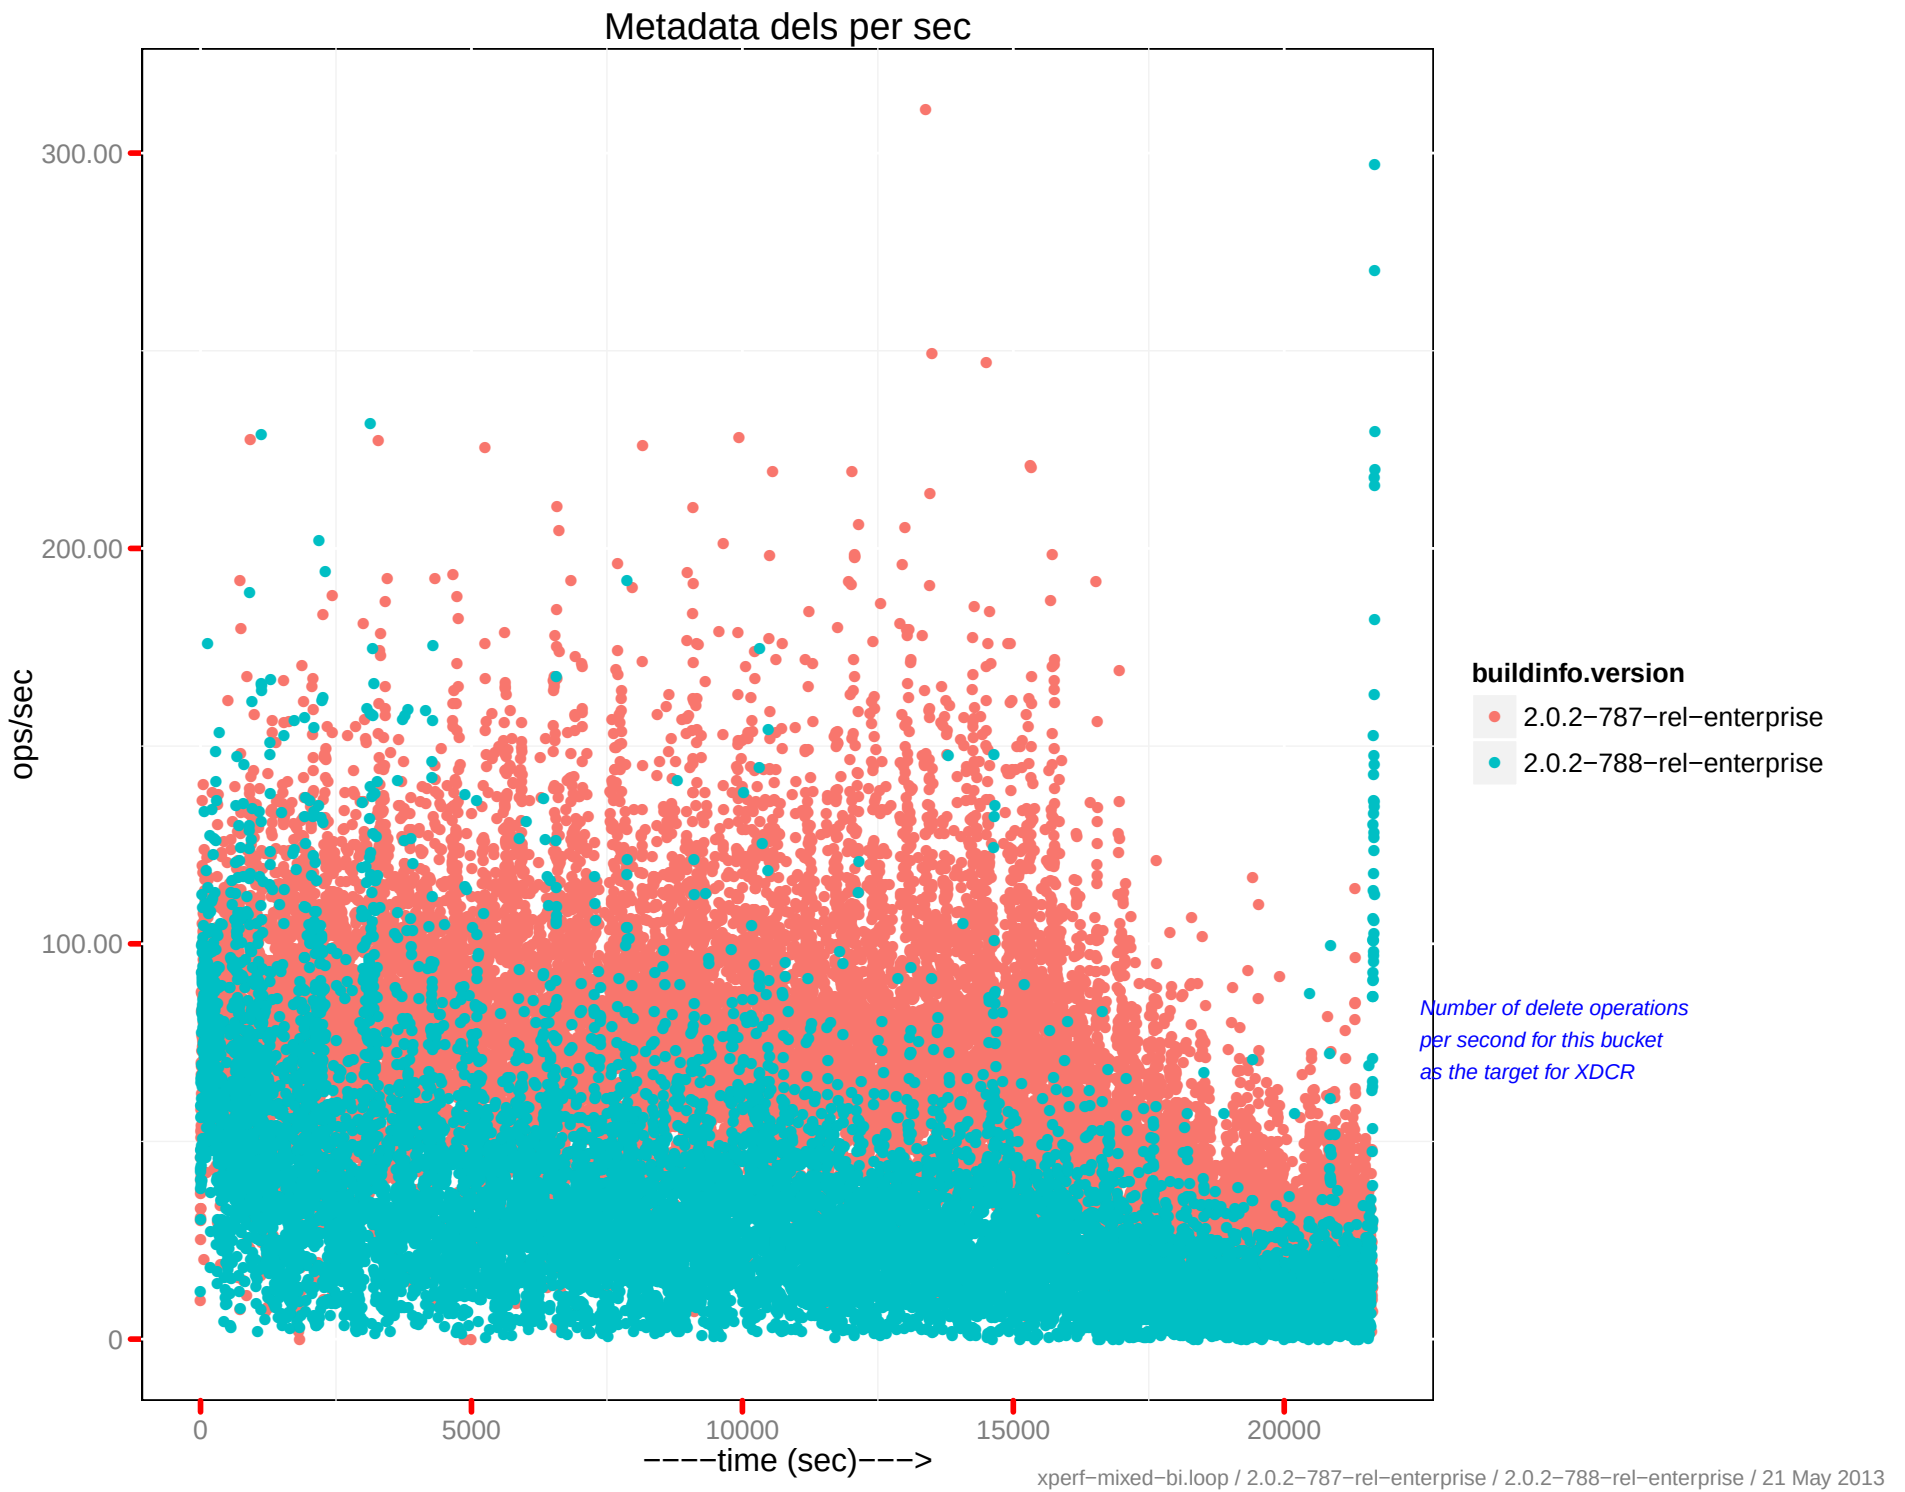
Metadata dels per sec
300.00
200.00
100.00
0
ops/sec
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
Number of delete operations
per second for this bucket
as the target for XDCR
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

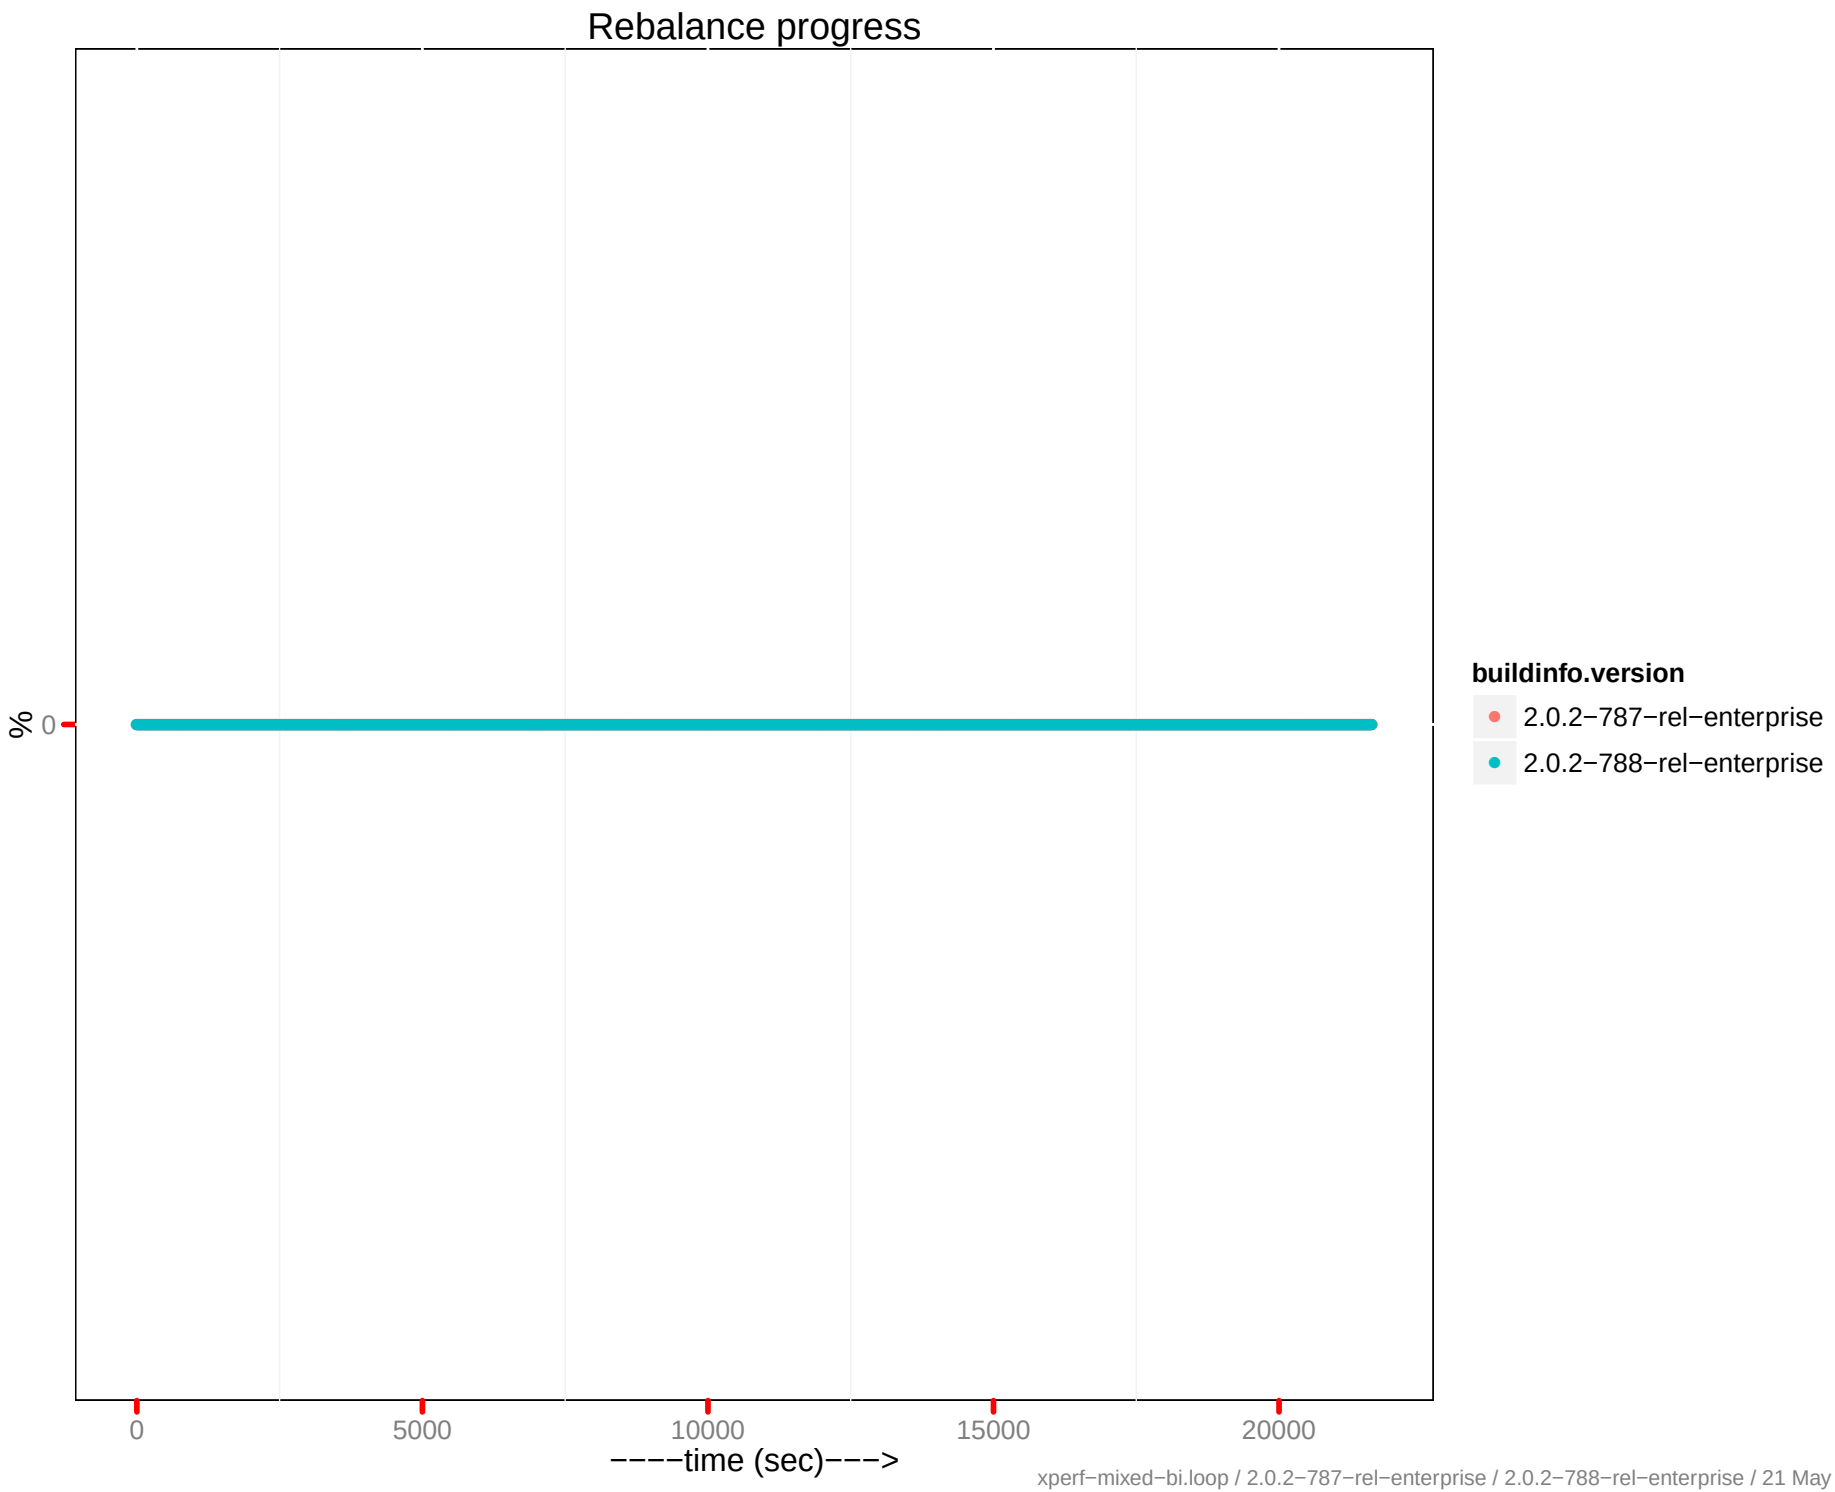
Rebalance progress
%
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

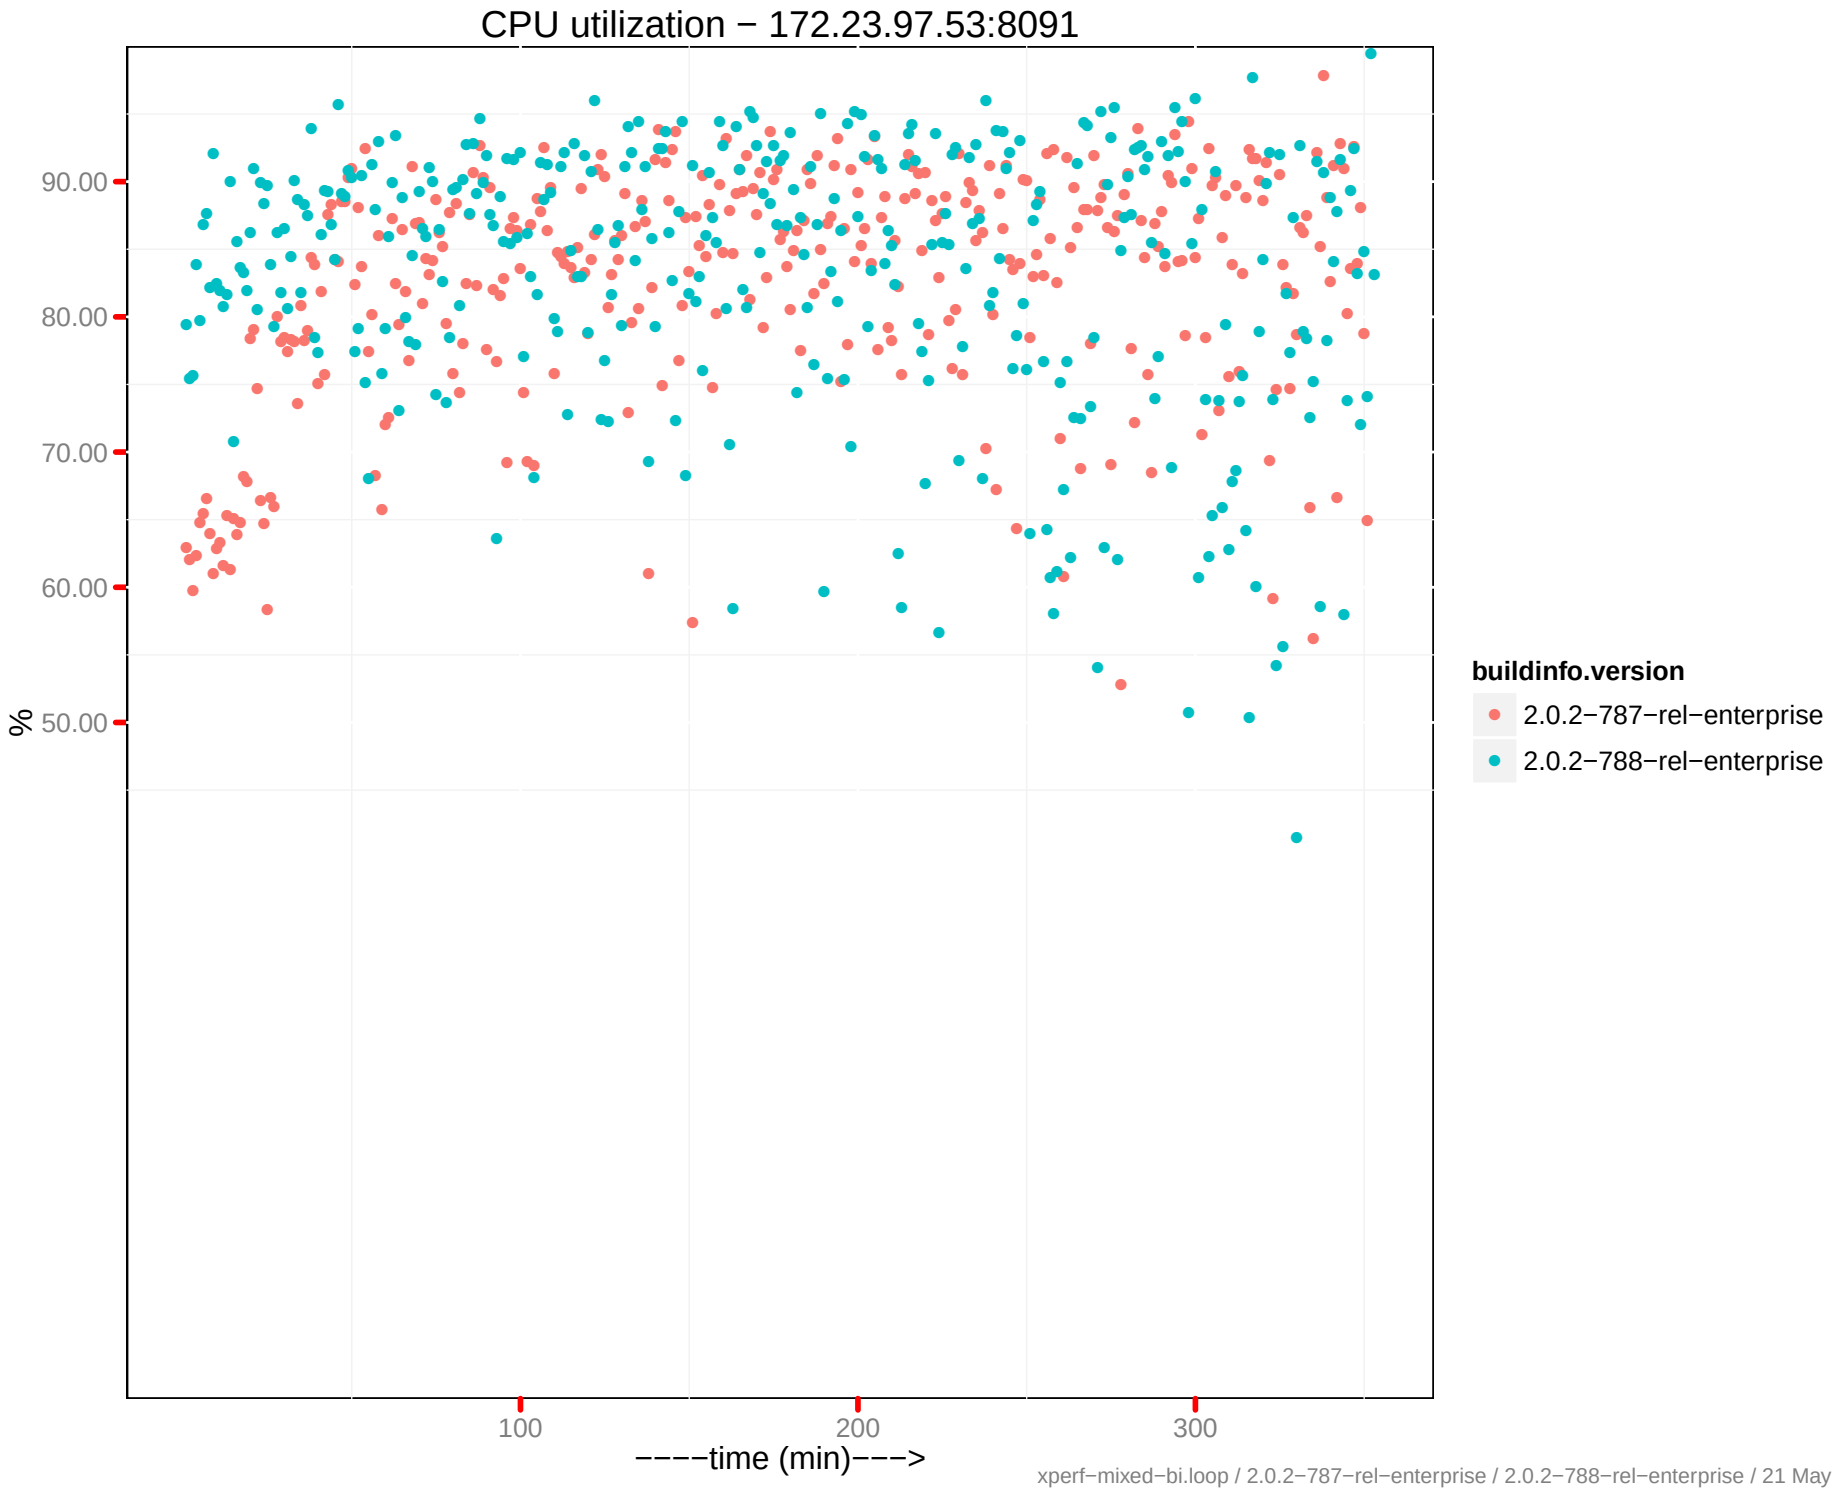
CPU utilization − 172.23.97.53:8091
90.00
80.00
70.00
60.00
50.00
%
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

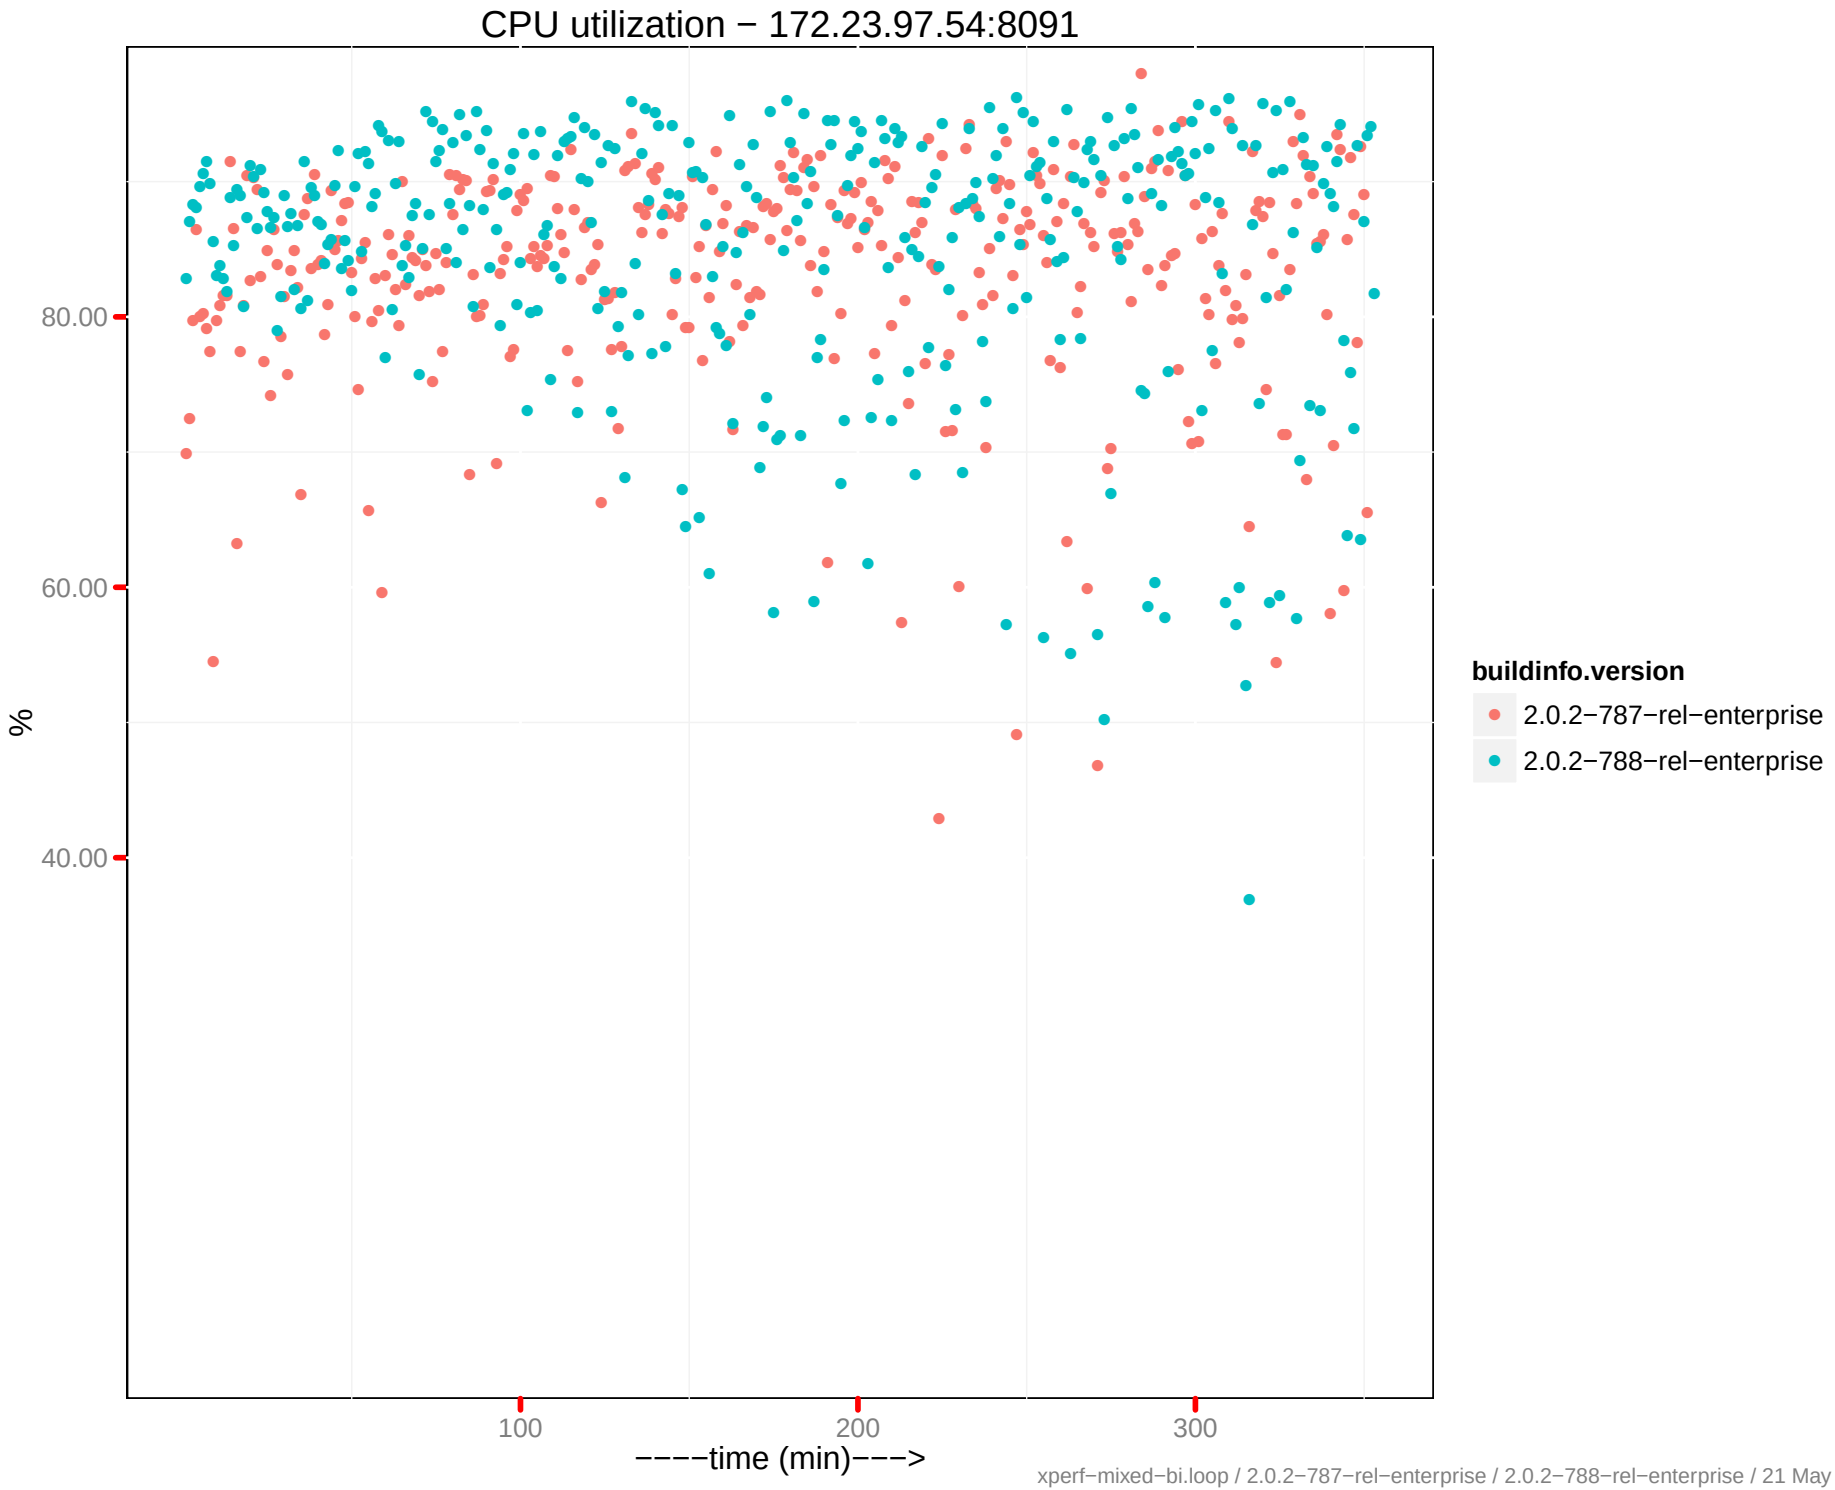
CPU utilization − 172.23.97.54:8091
80.00
60.00
40.00
%
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

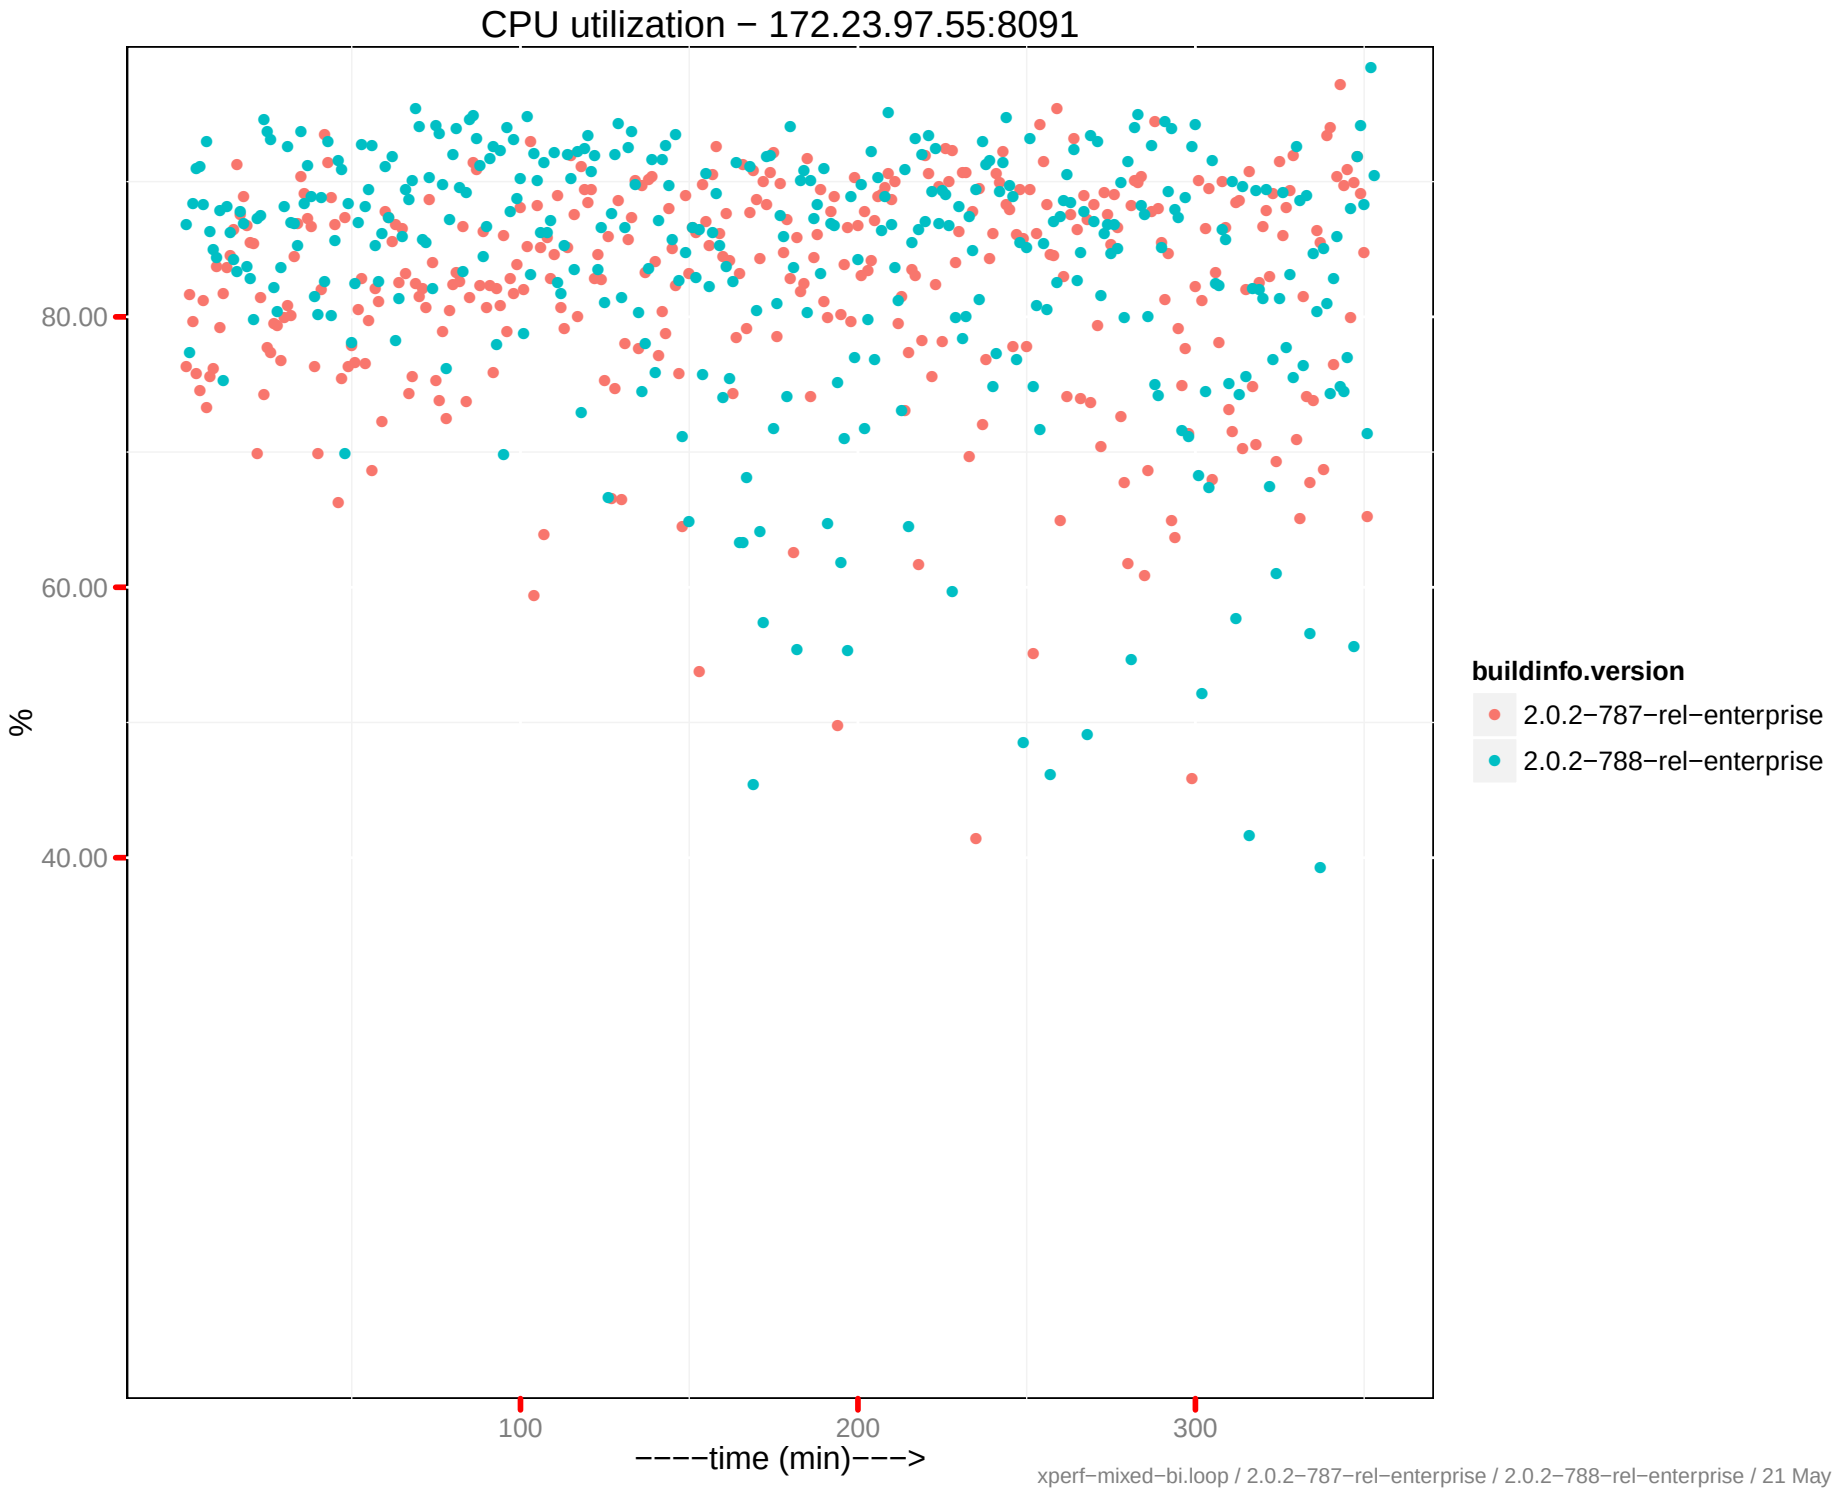
CPU utilization − 172.23.97.55:8091
80.00
60.00
40.00
%
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

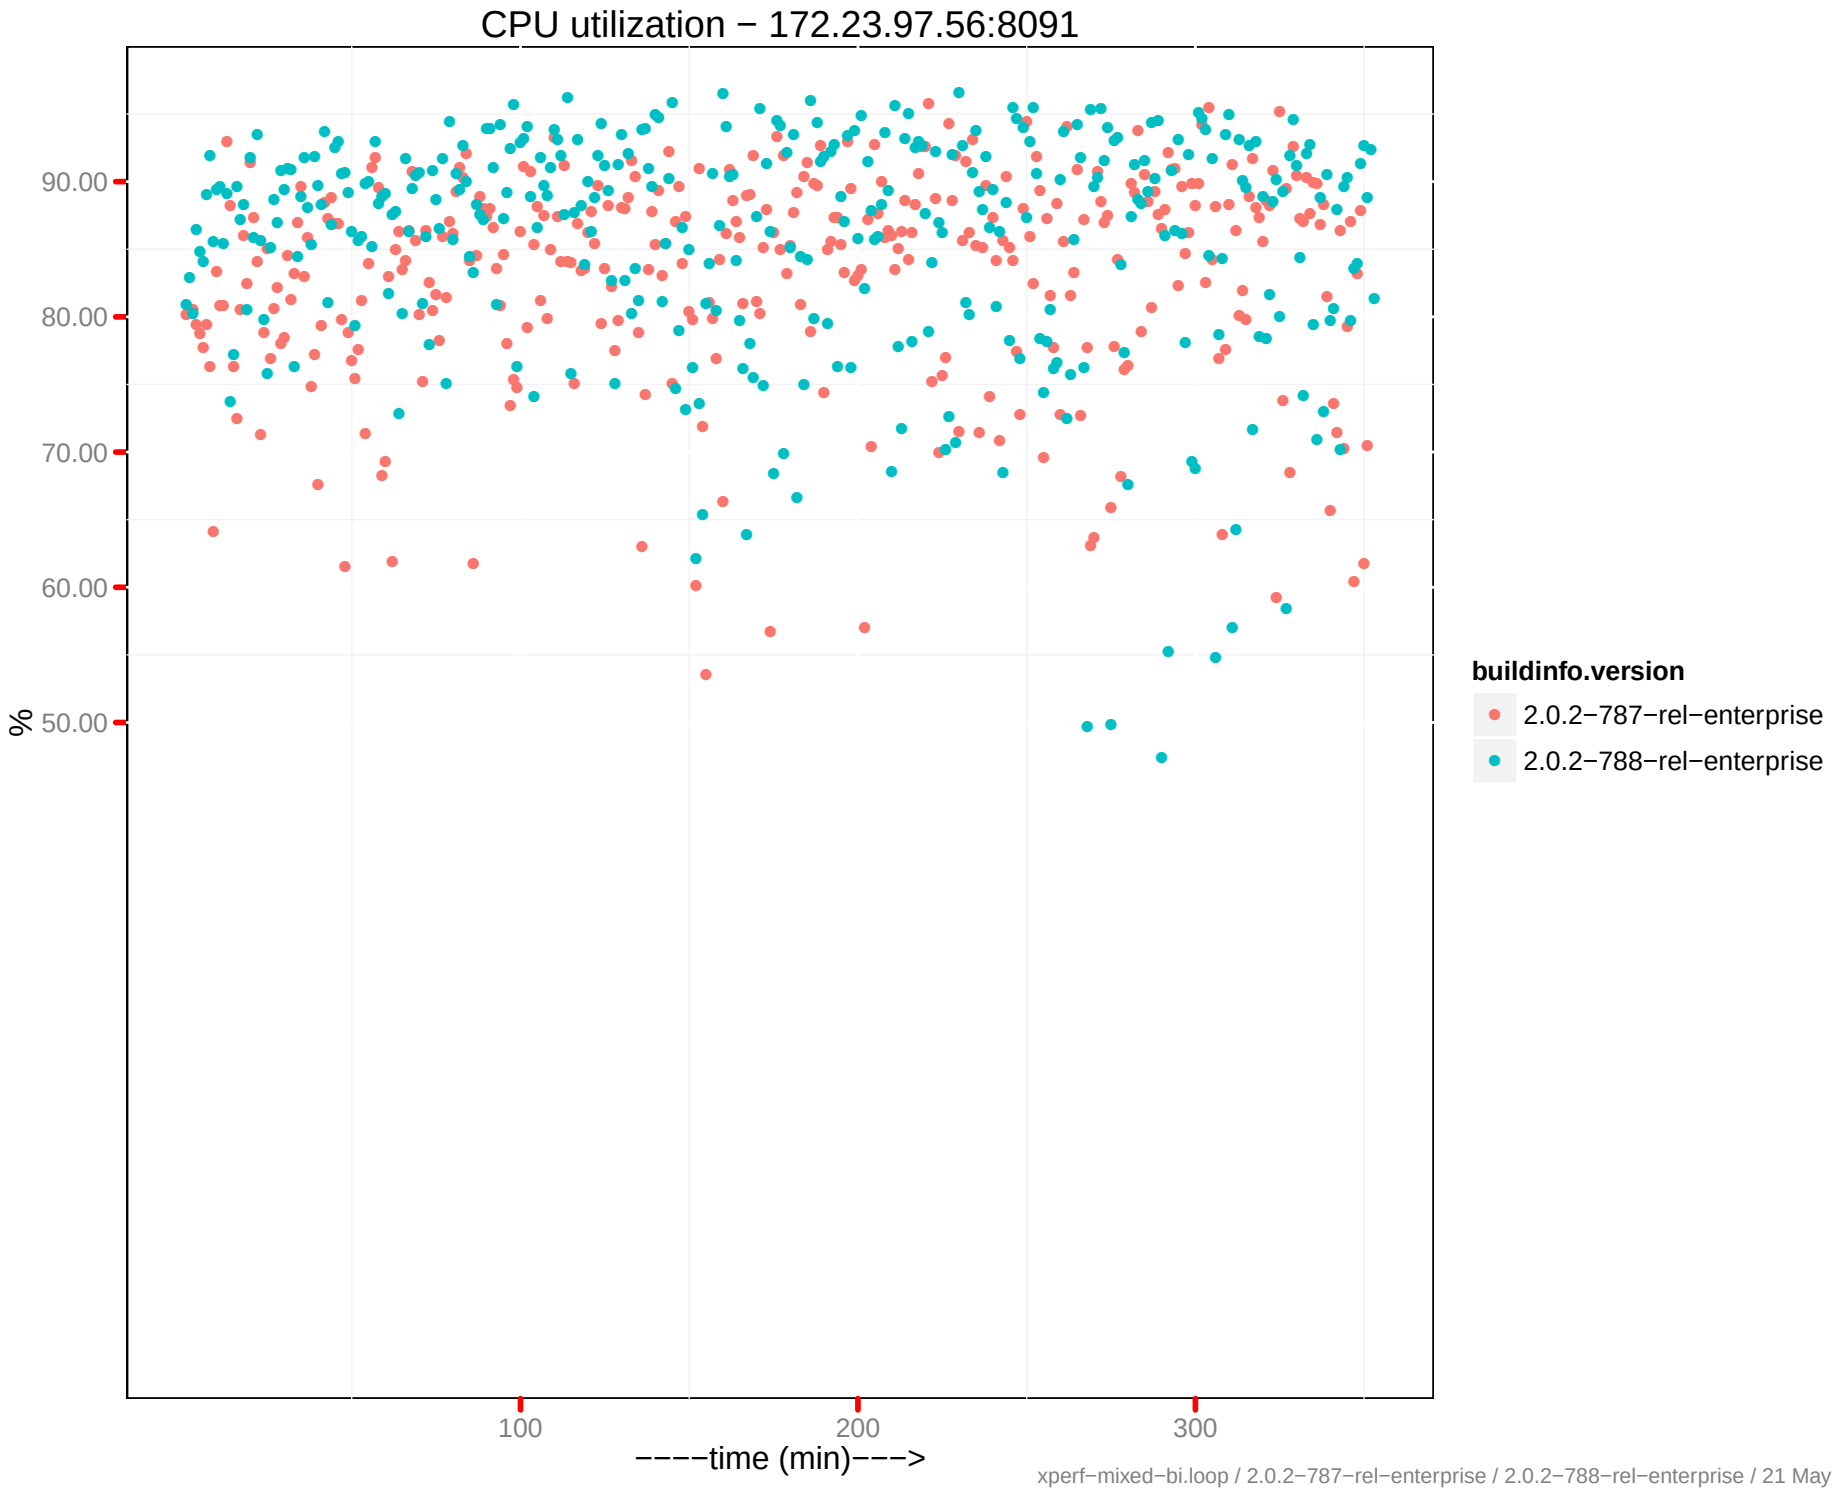
CPU utilization − 172.23.97.56:8091
90.00
80.00
70.00
60.00
50.00
%
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

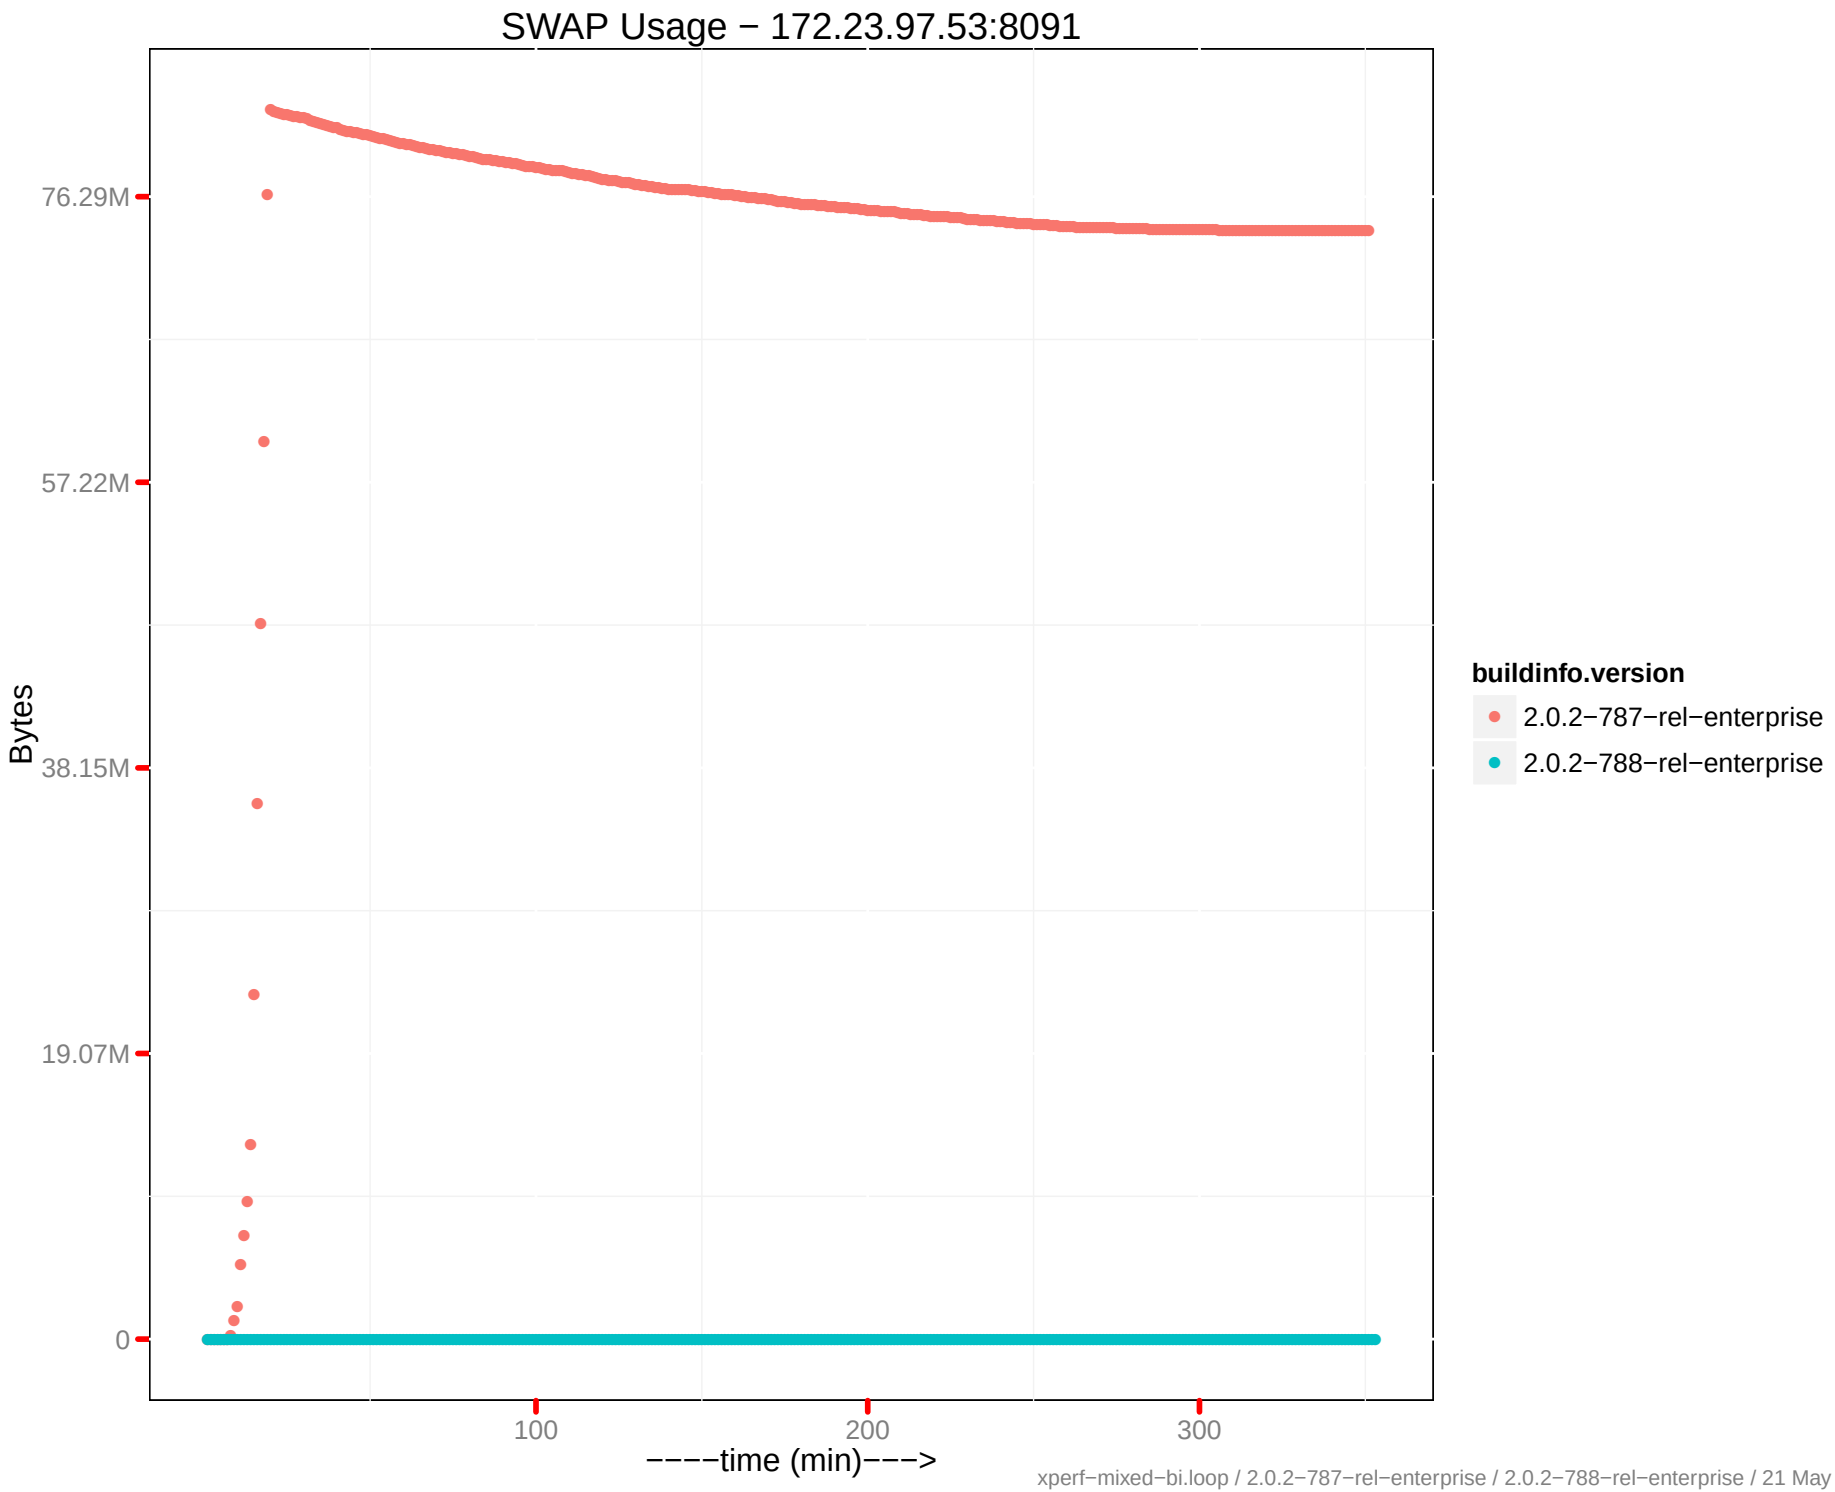
SWAP Usage − 172.23.97.53:8091
76.29M
57.22M
38.15M
19.07M
0
Bytes
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

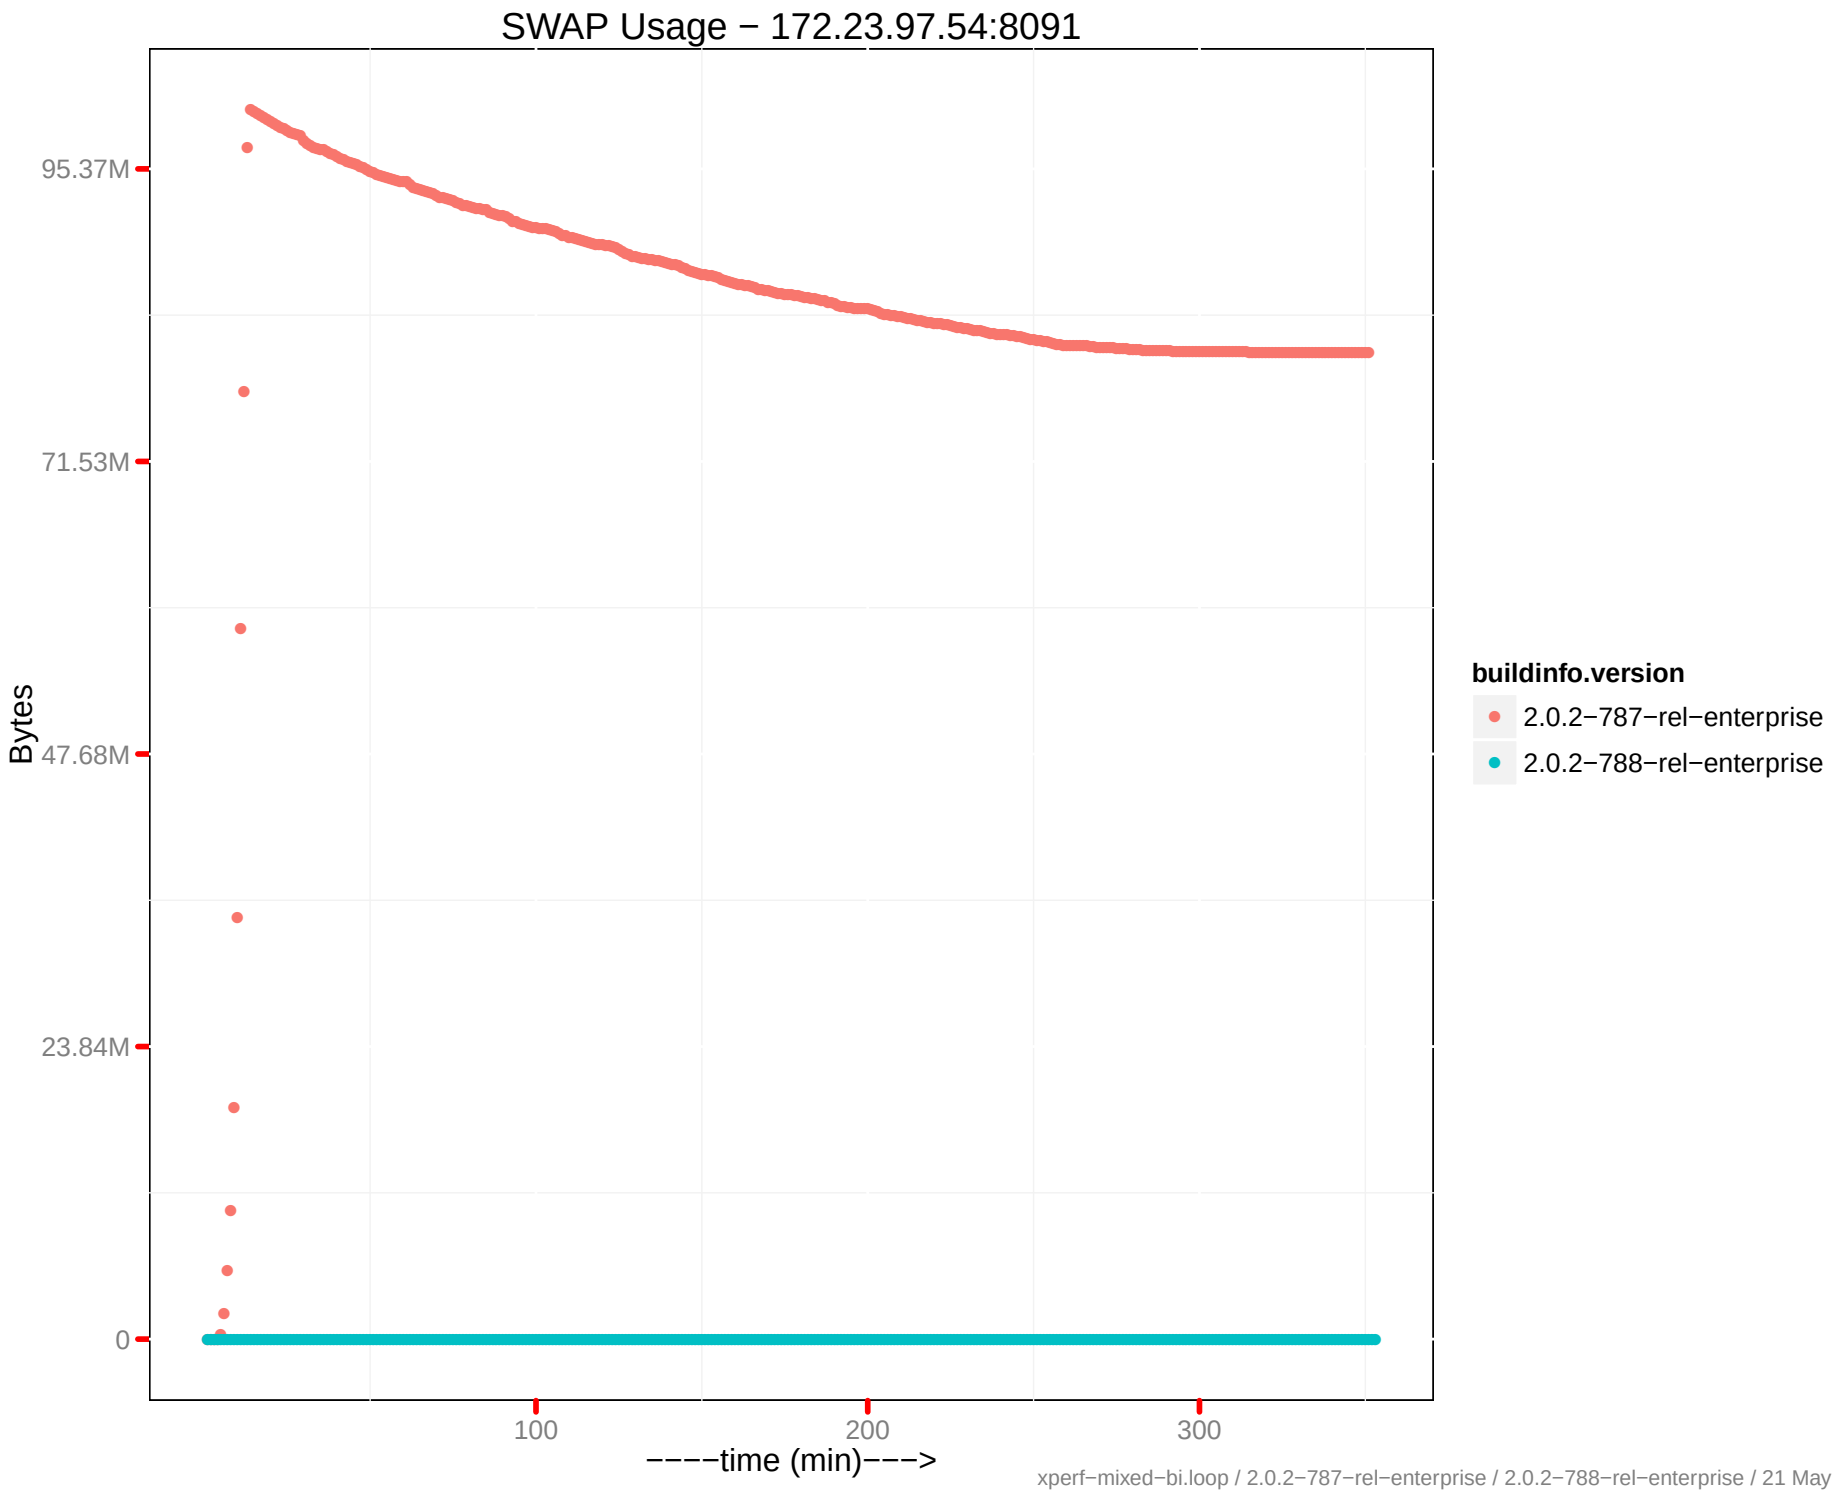
SWAP Usage − 172.23.97.54:8091
95.37M
71.53M
47.68M
23.84M
0
Bytes
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

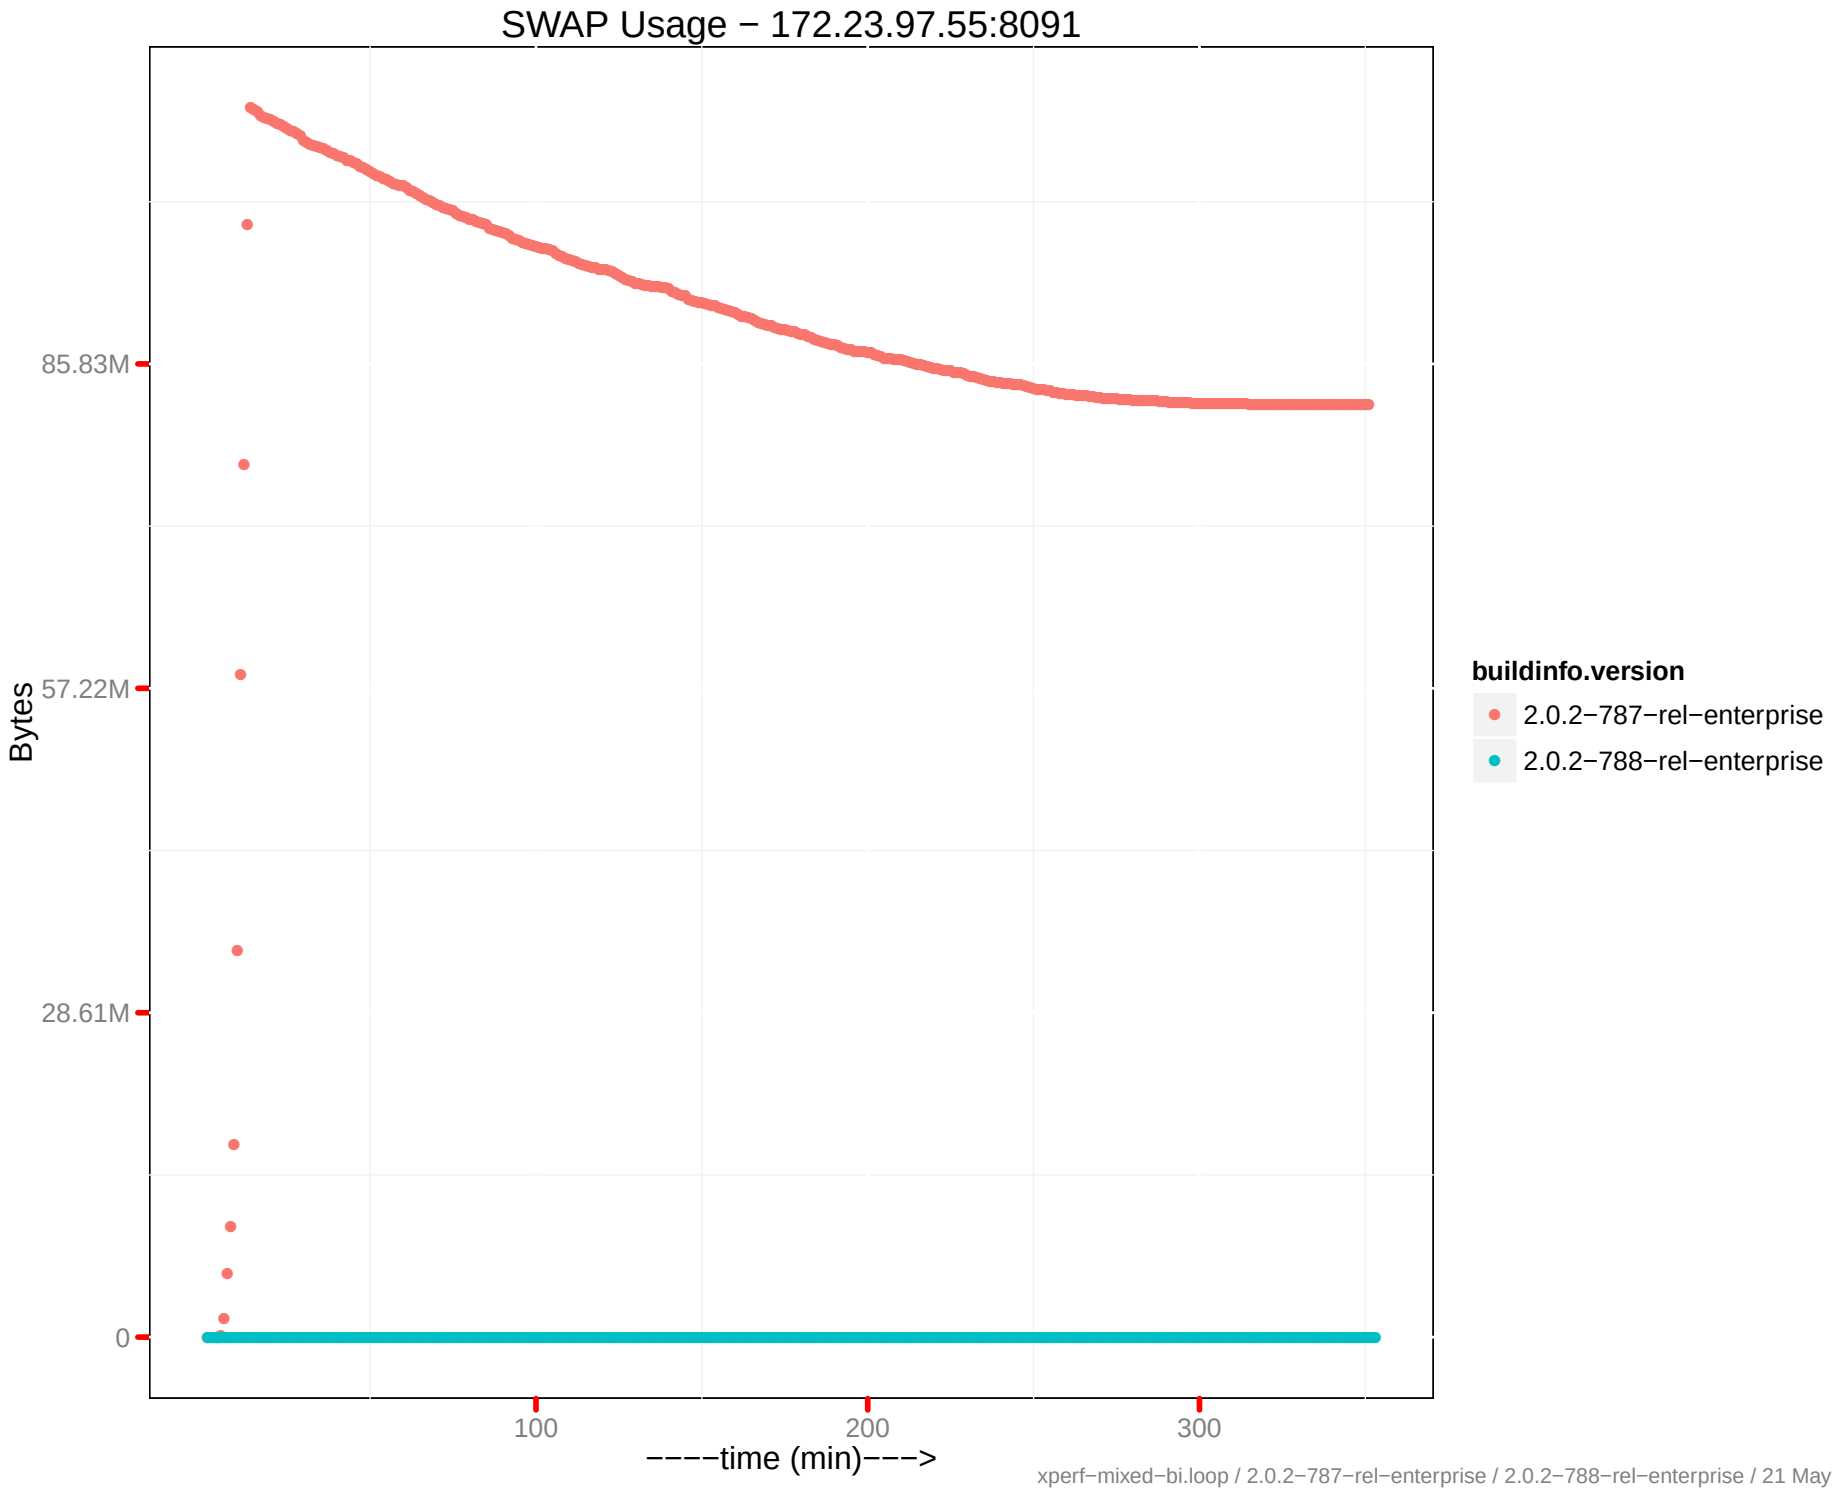
SWAP Usage − 172.23.97.55:8091
85.83M
57.22M
28.61M
0
Bytes
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

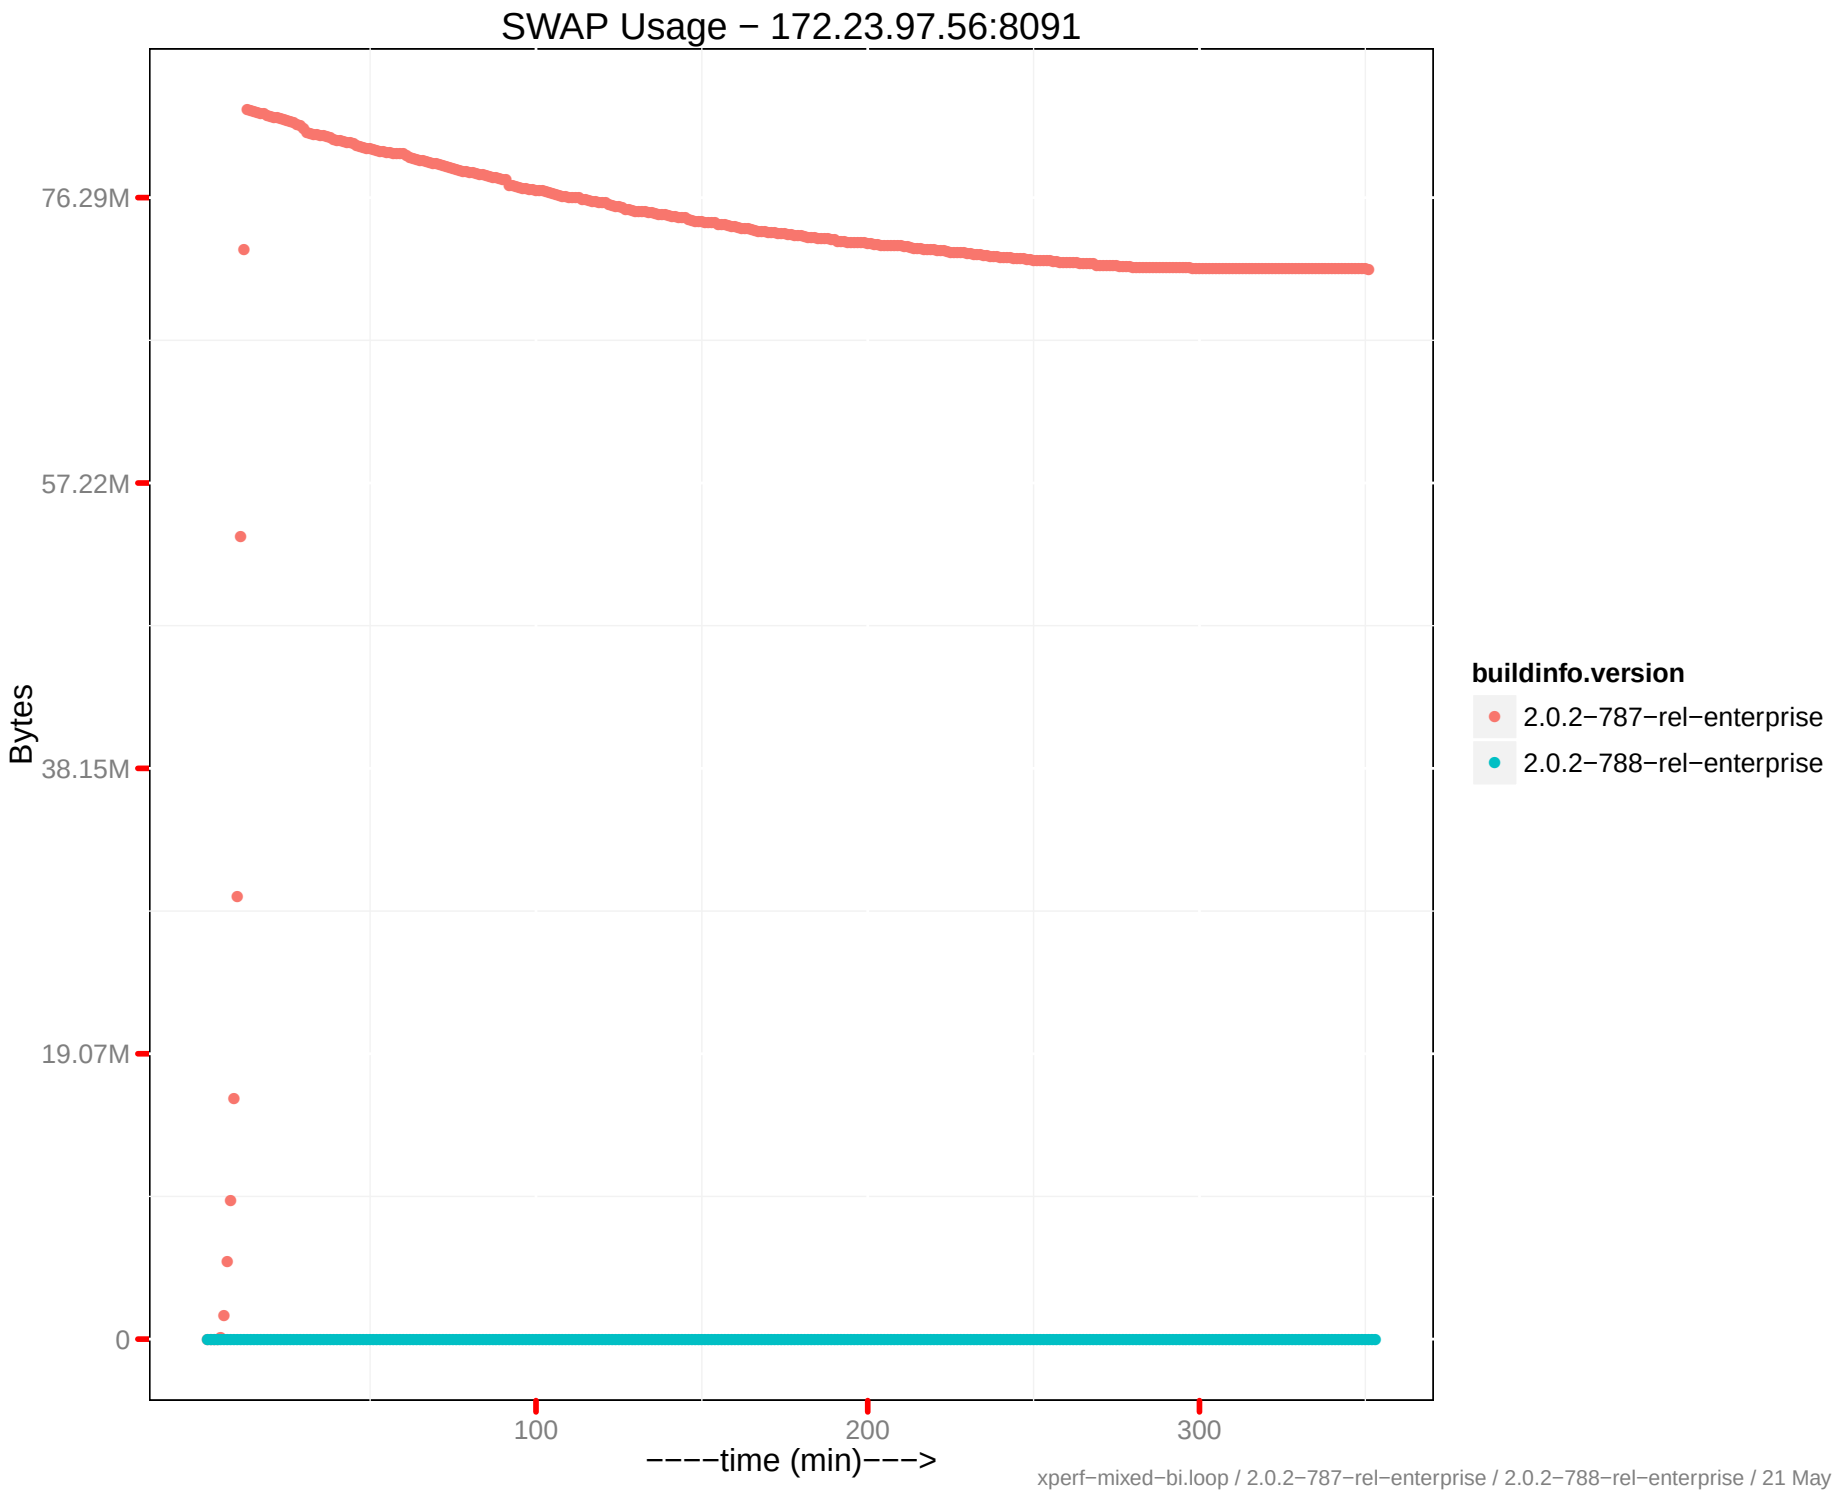
SWAP Usage − 172.23.97.56:8091
76.29M
57.22M
38.15M
19.07M
0
Bytes
100
200
300
−−−−time (min)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

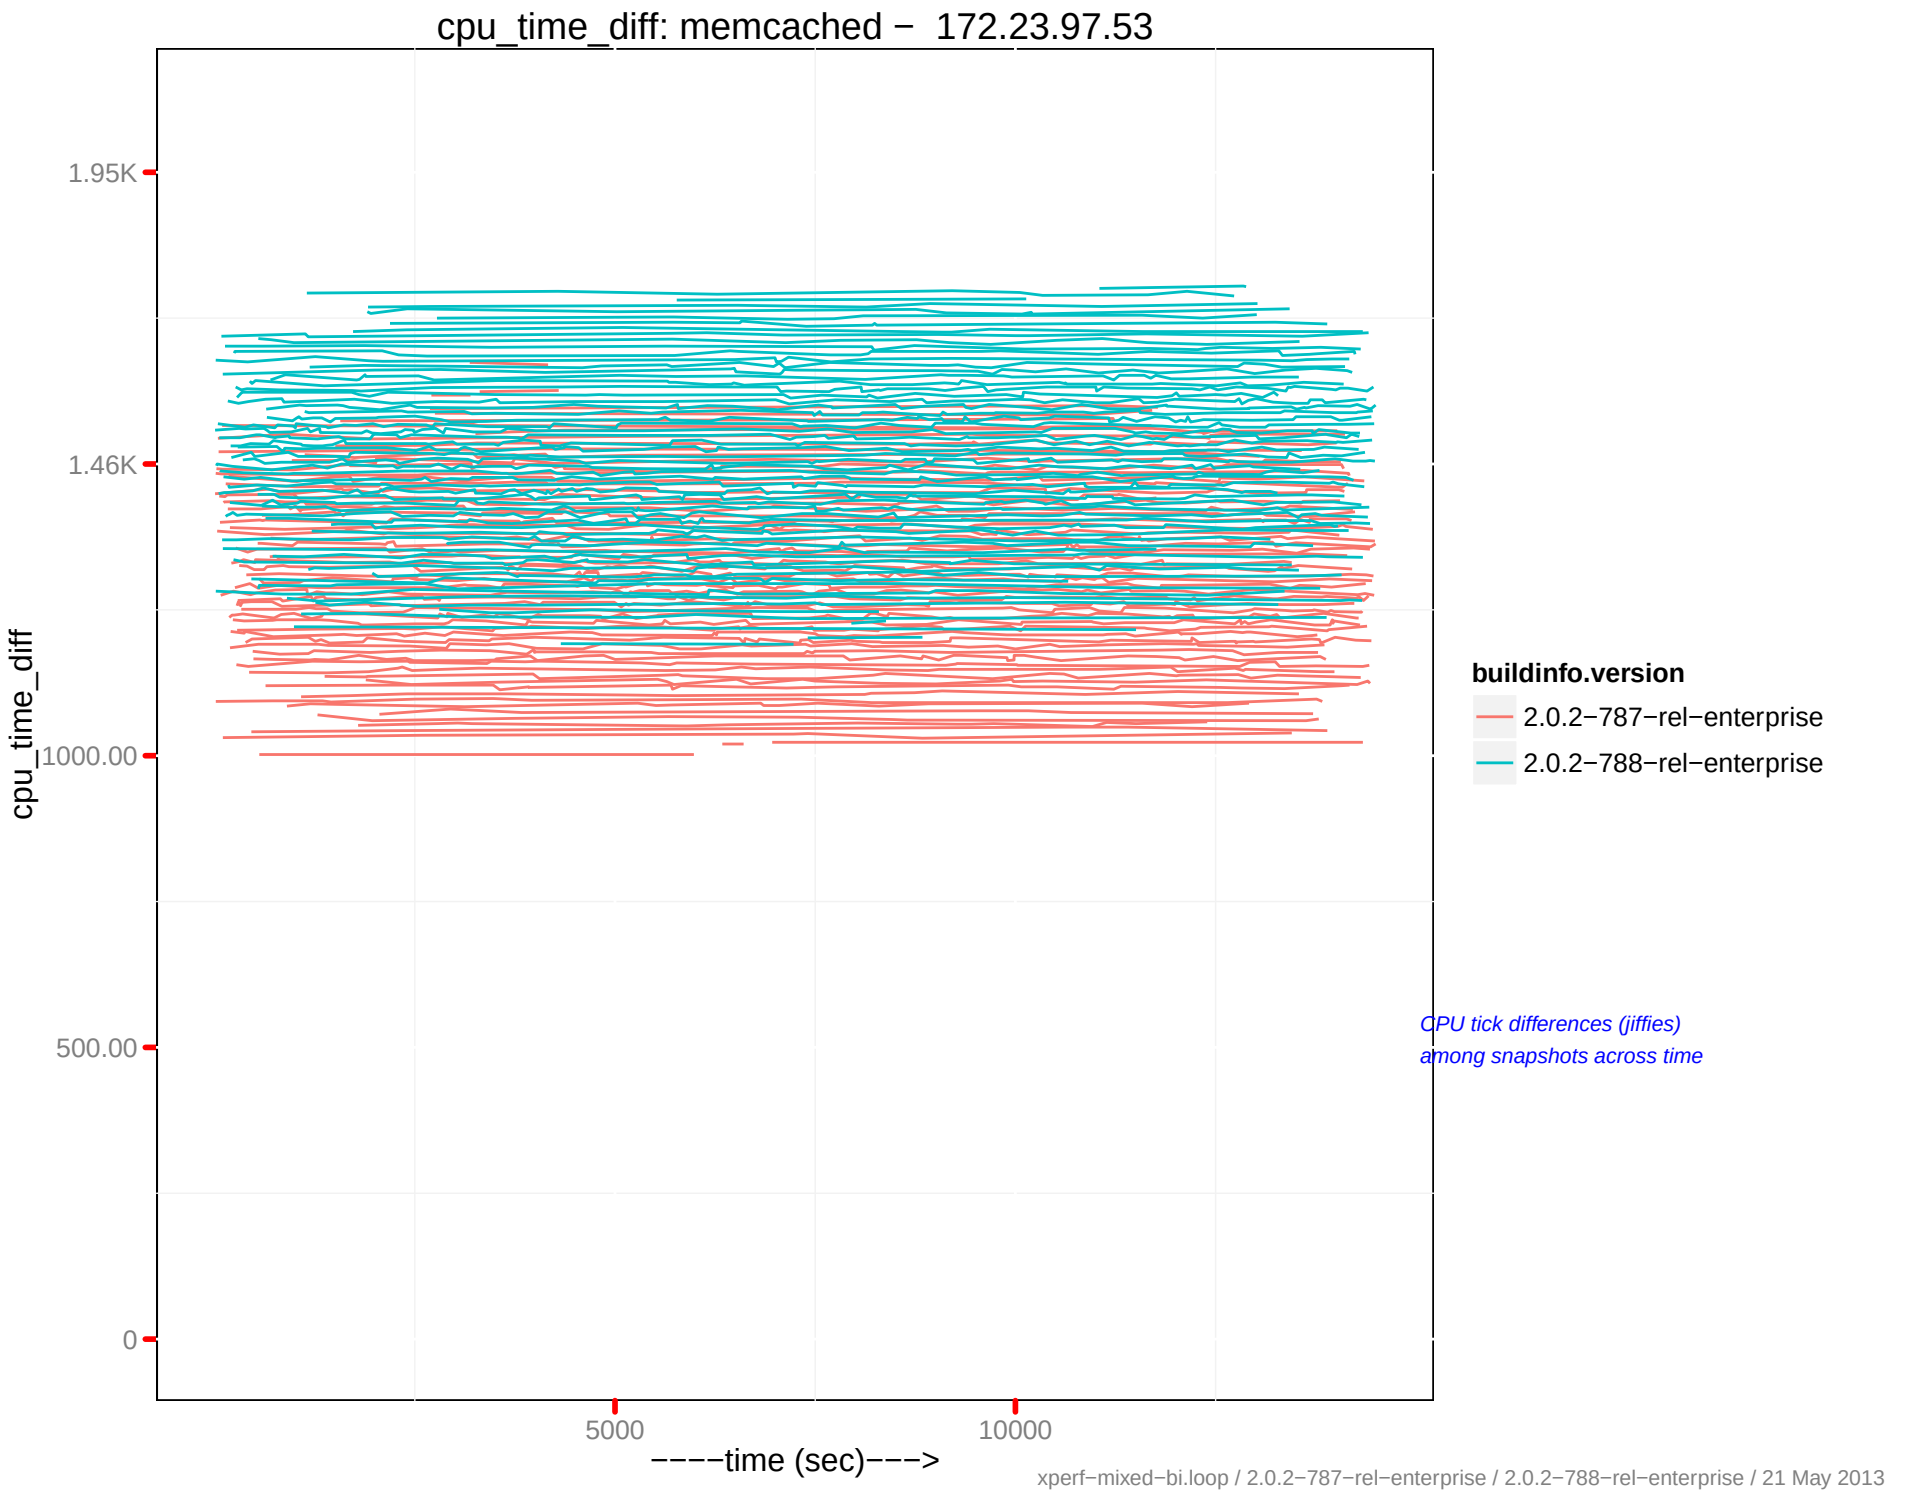
cpu_time_diff: memcached − 172.23.97.53
1.95K
1.46K
1000.00
500.00
0
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

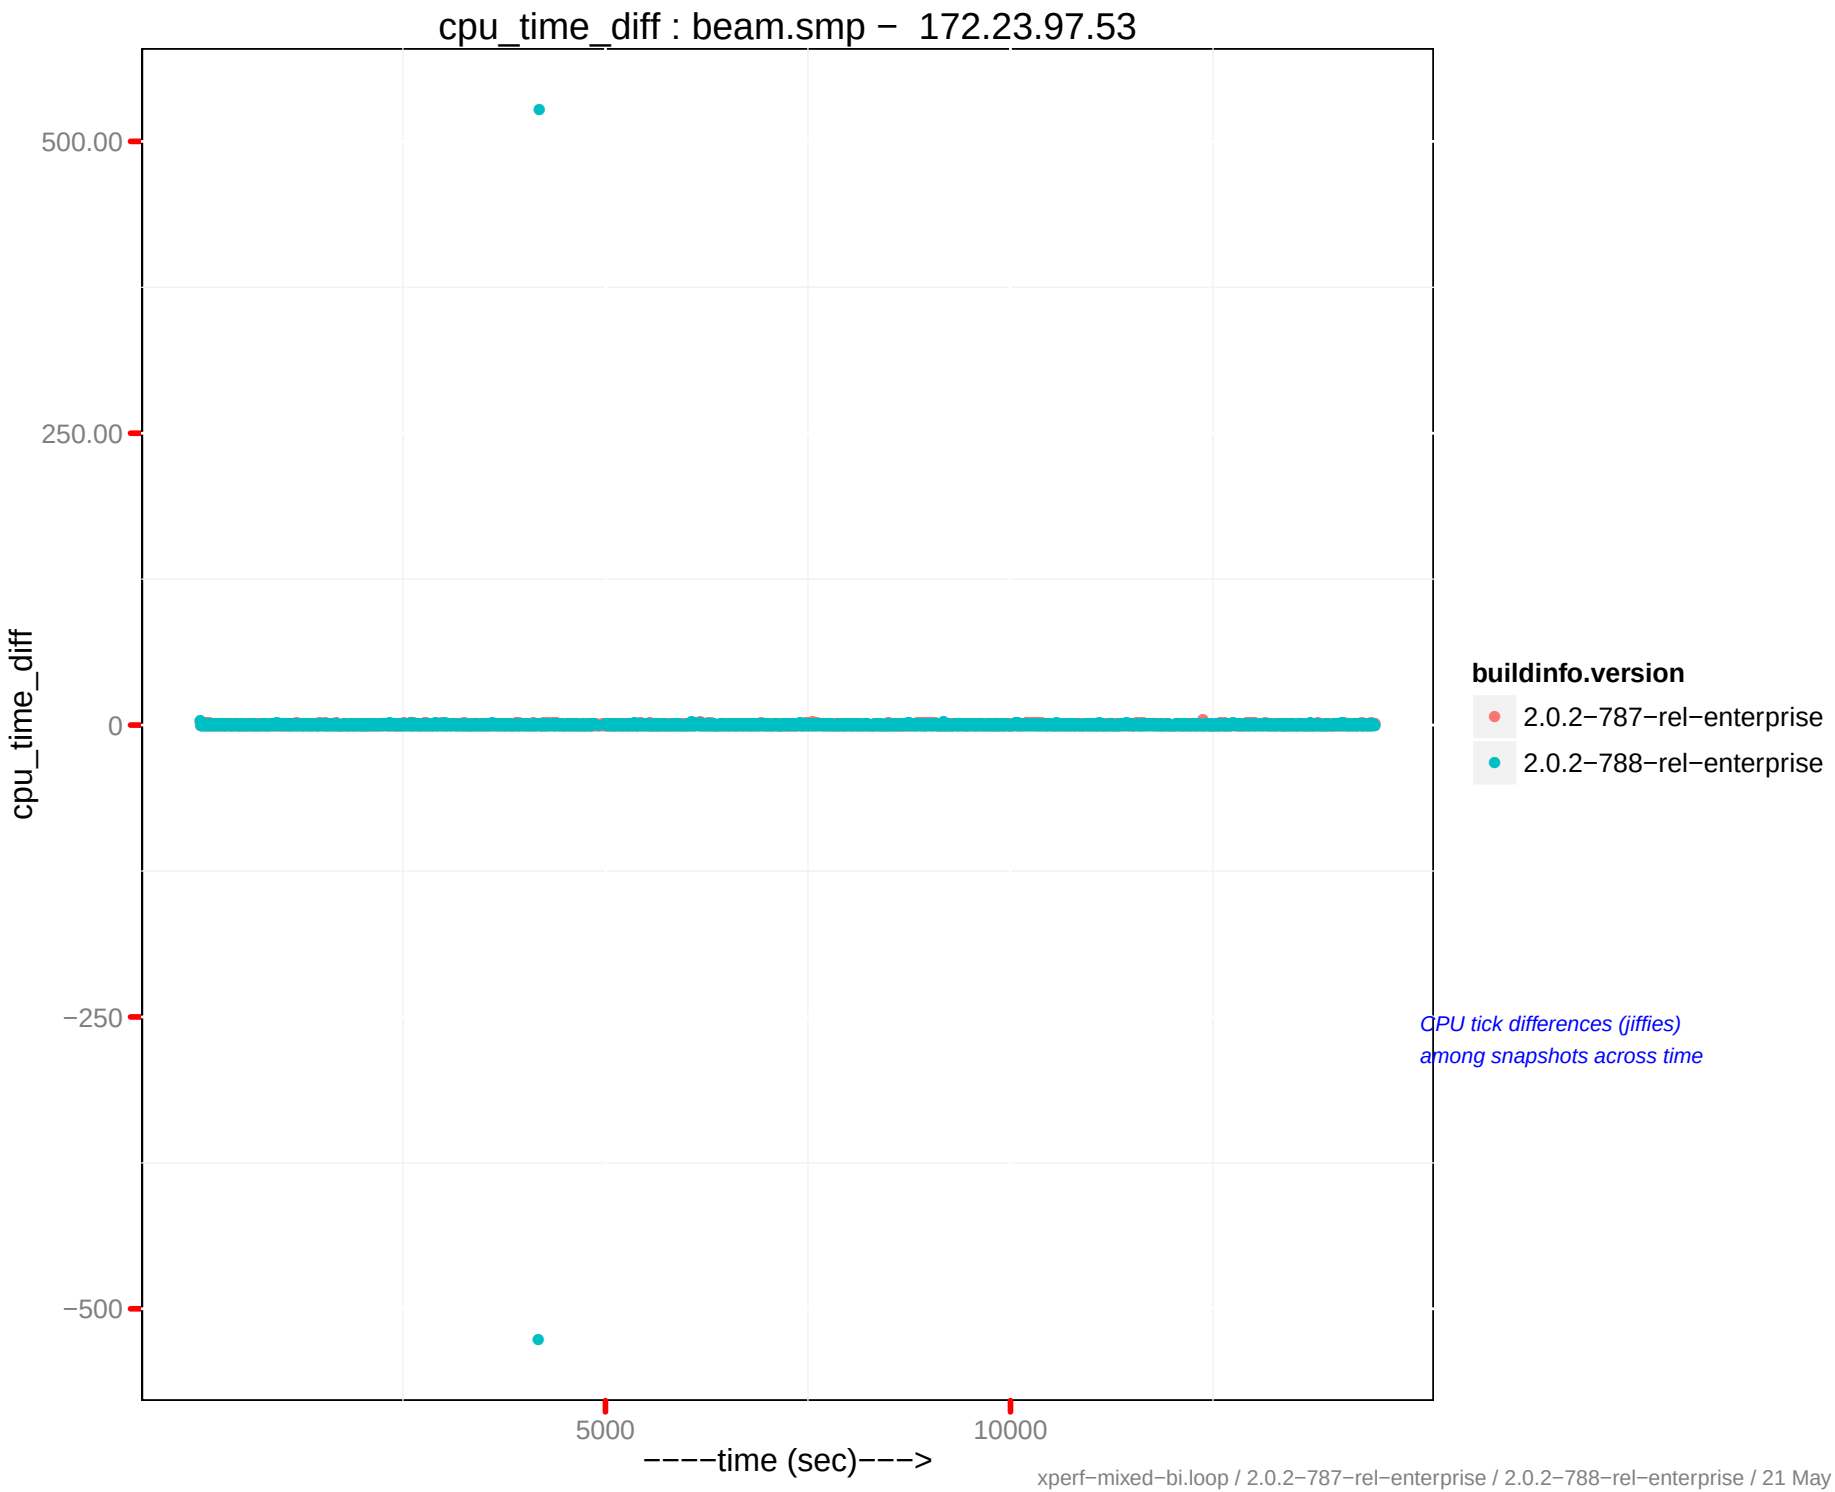
cpu_time_diff : beam.smp − 172.23.97.53
500.00
250.00
0
−250
−500
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

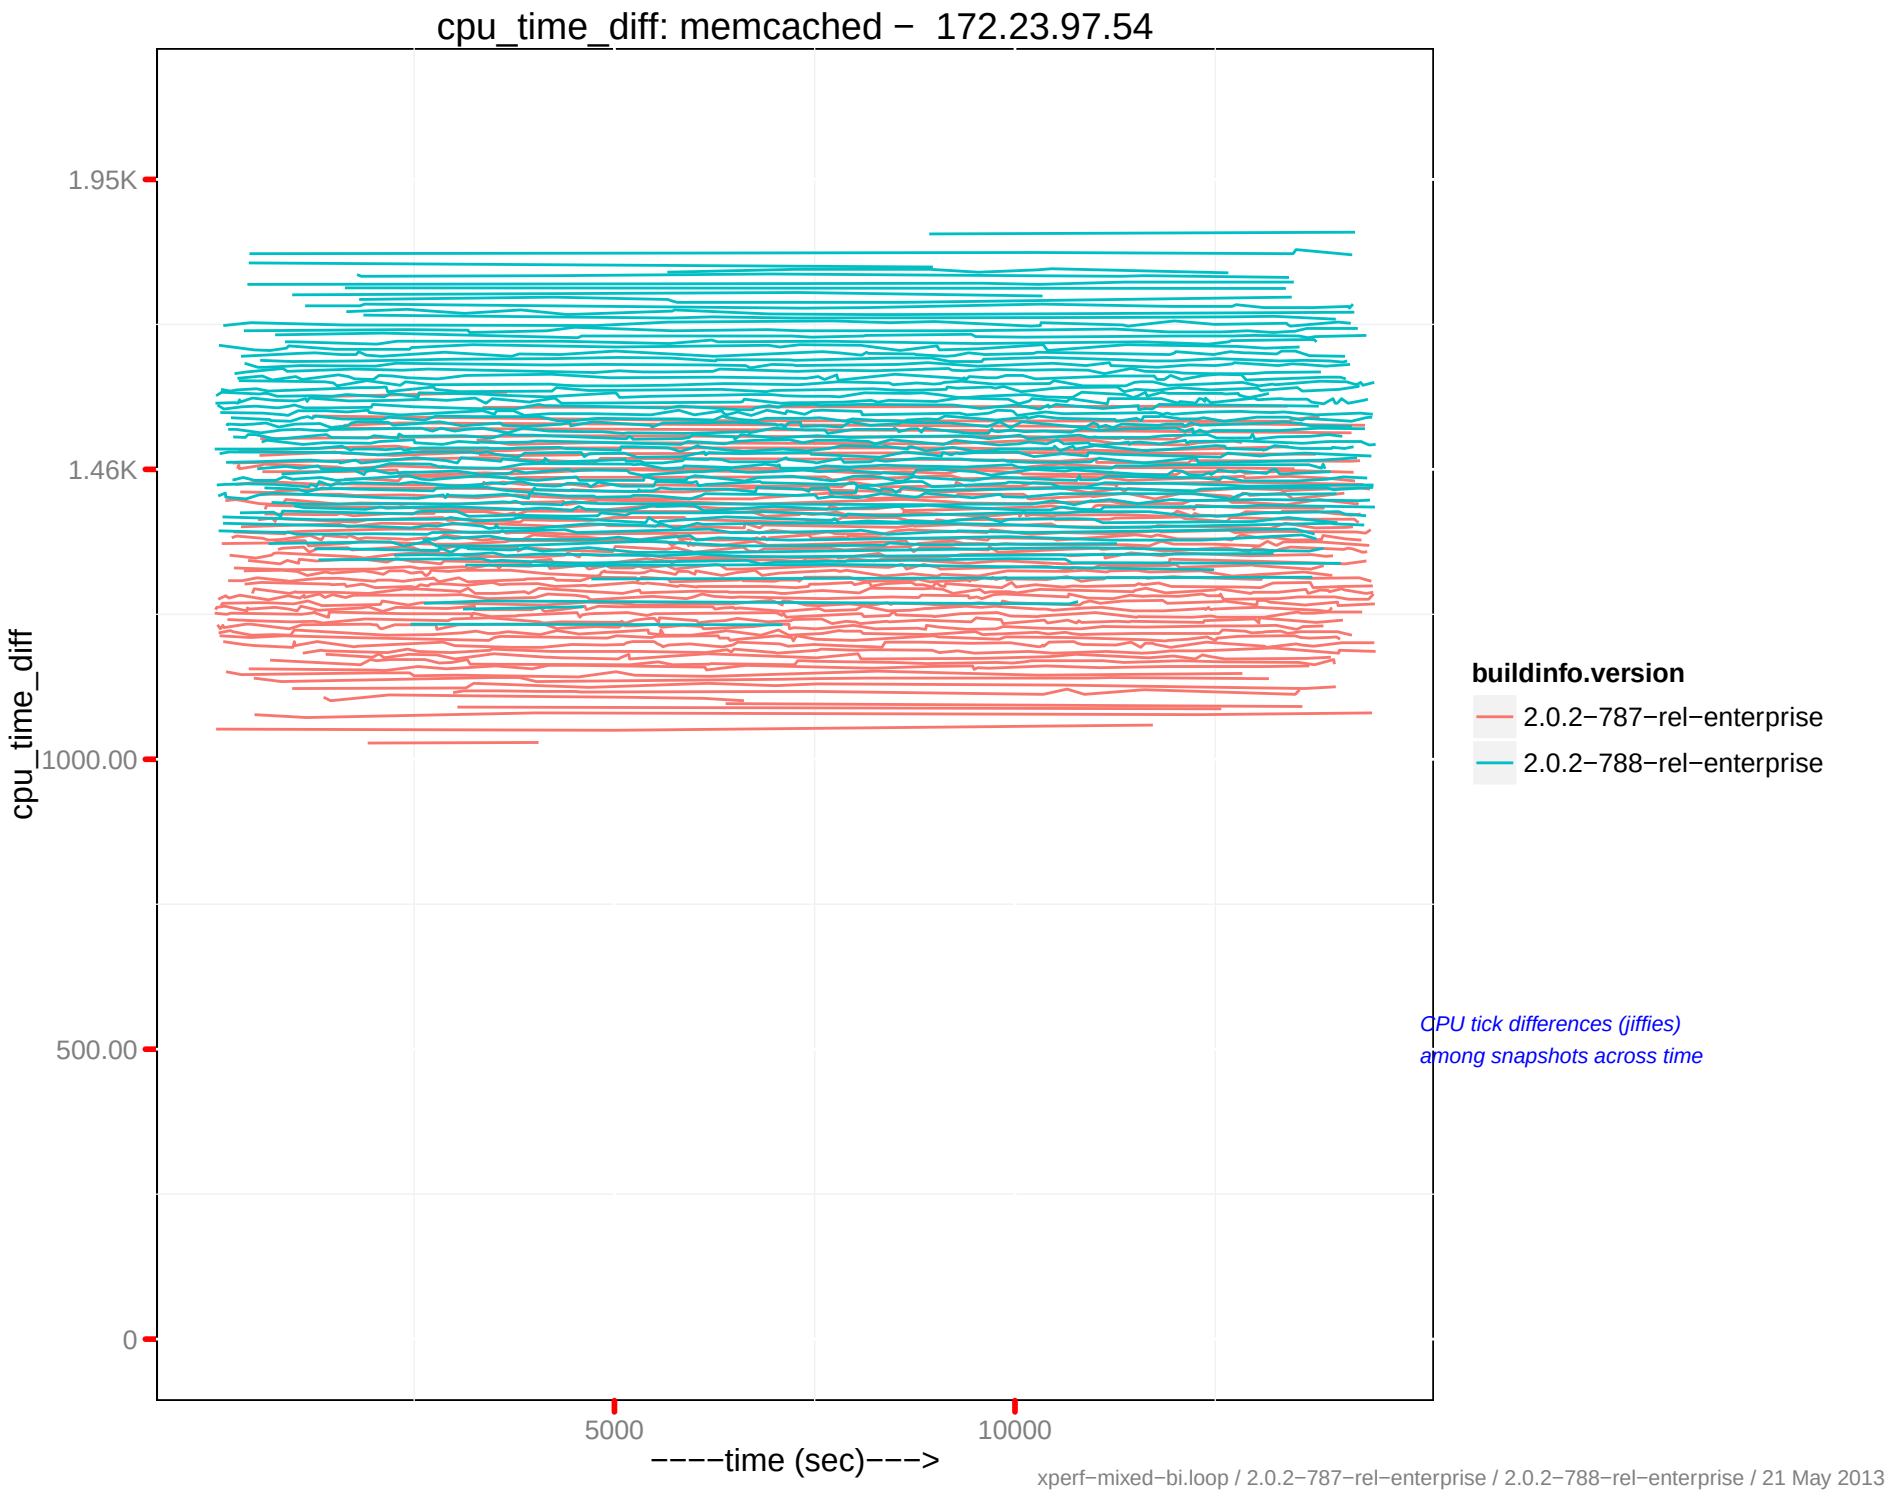
cpu_time_diff: memcached − 172.23.97.54
1.95K
1.46K
1000.00
500.00
0
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

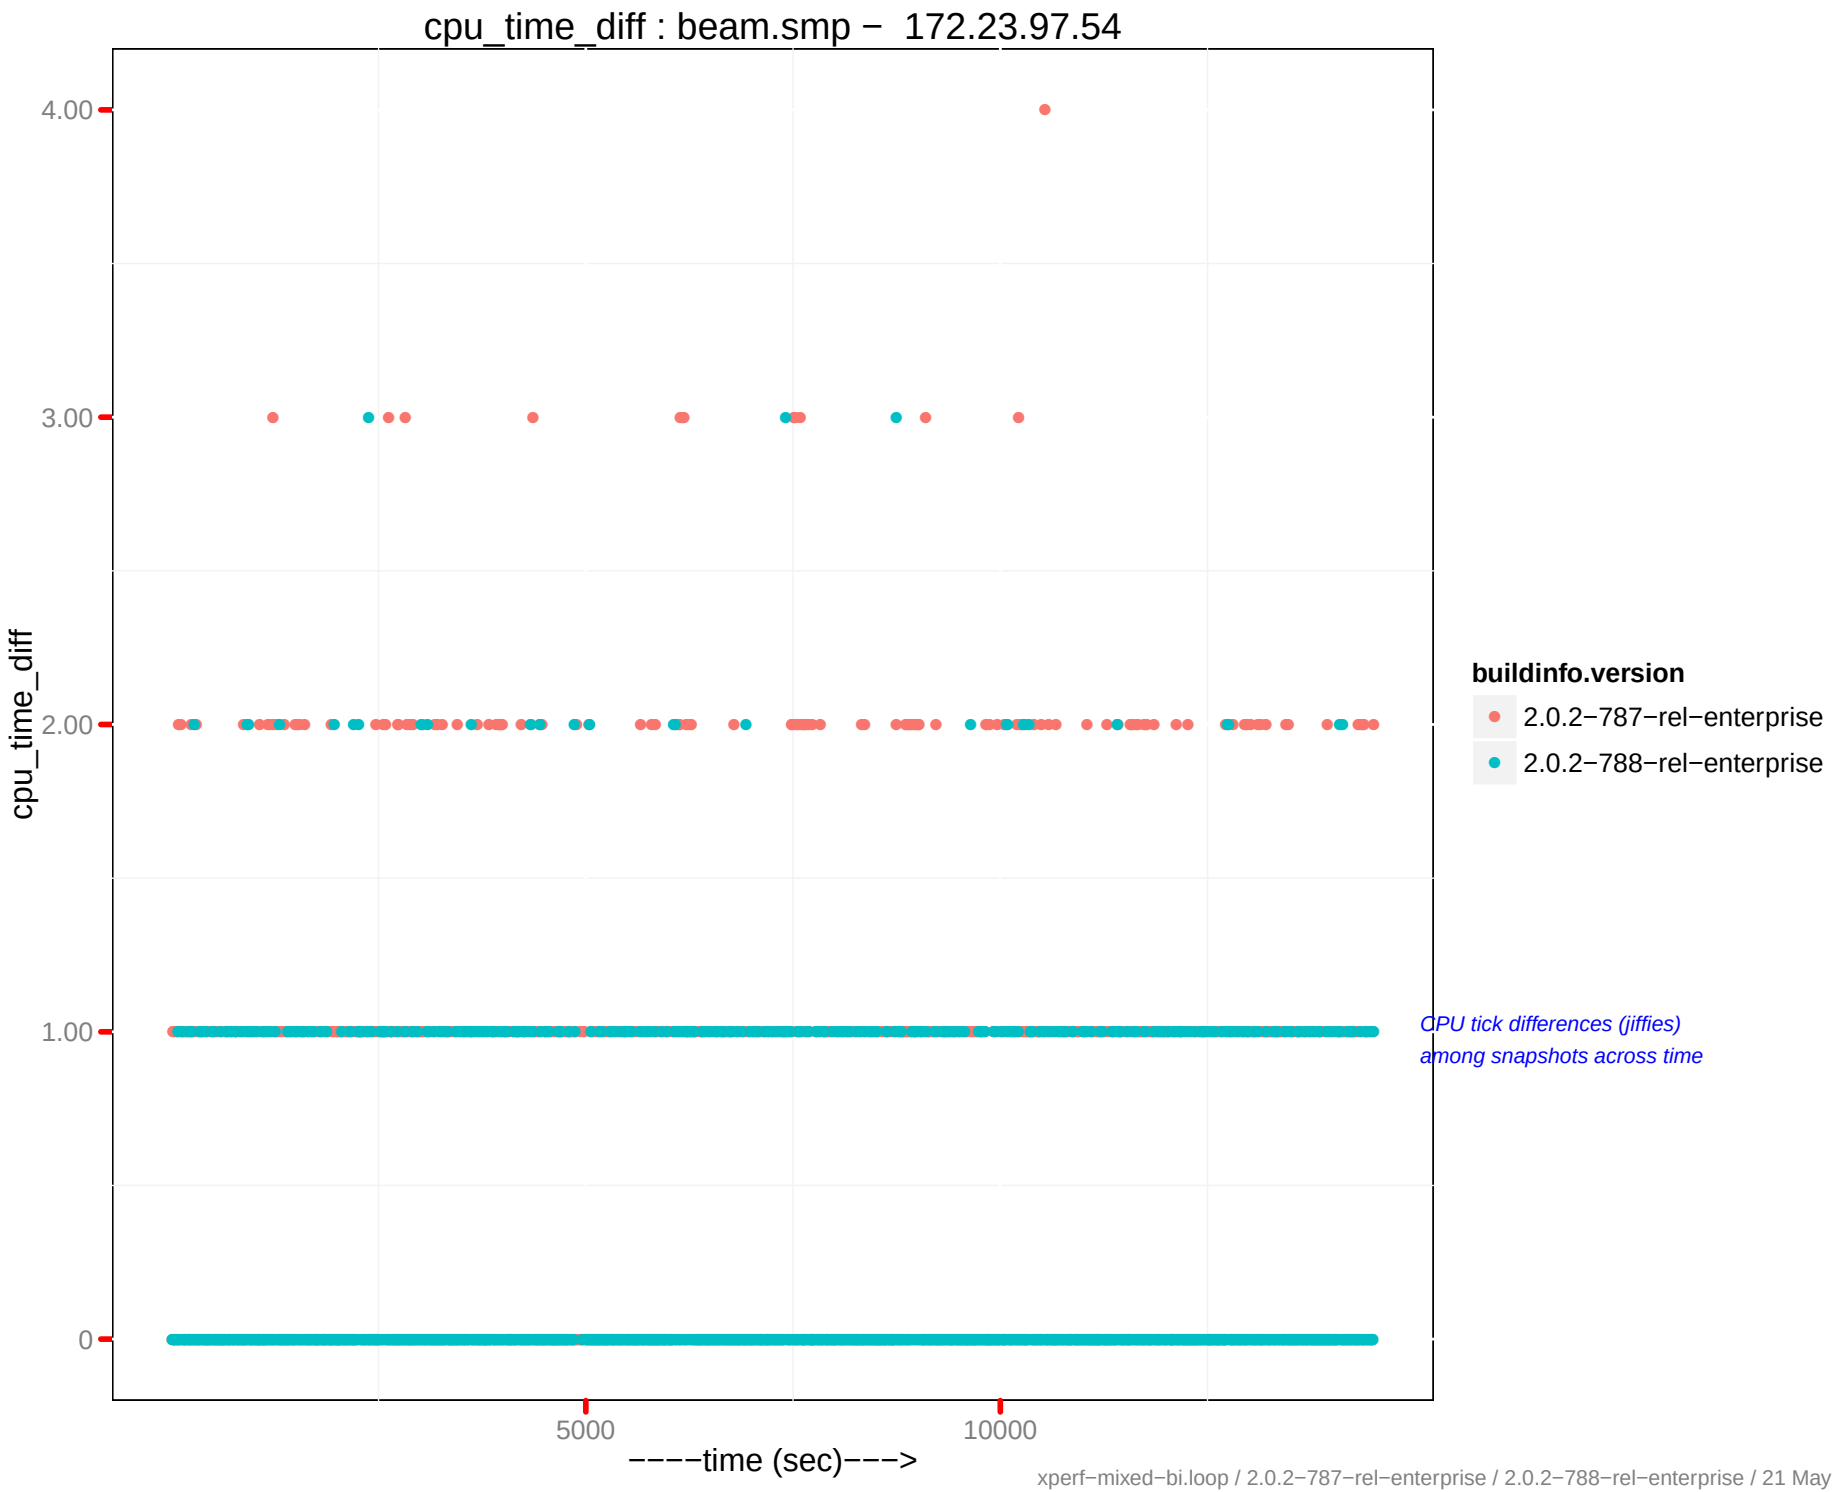
cpu_time_diff : beam.smp − 172.23.97.54
4.00
3.00
2.00
1.00
0
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

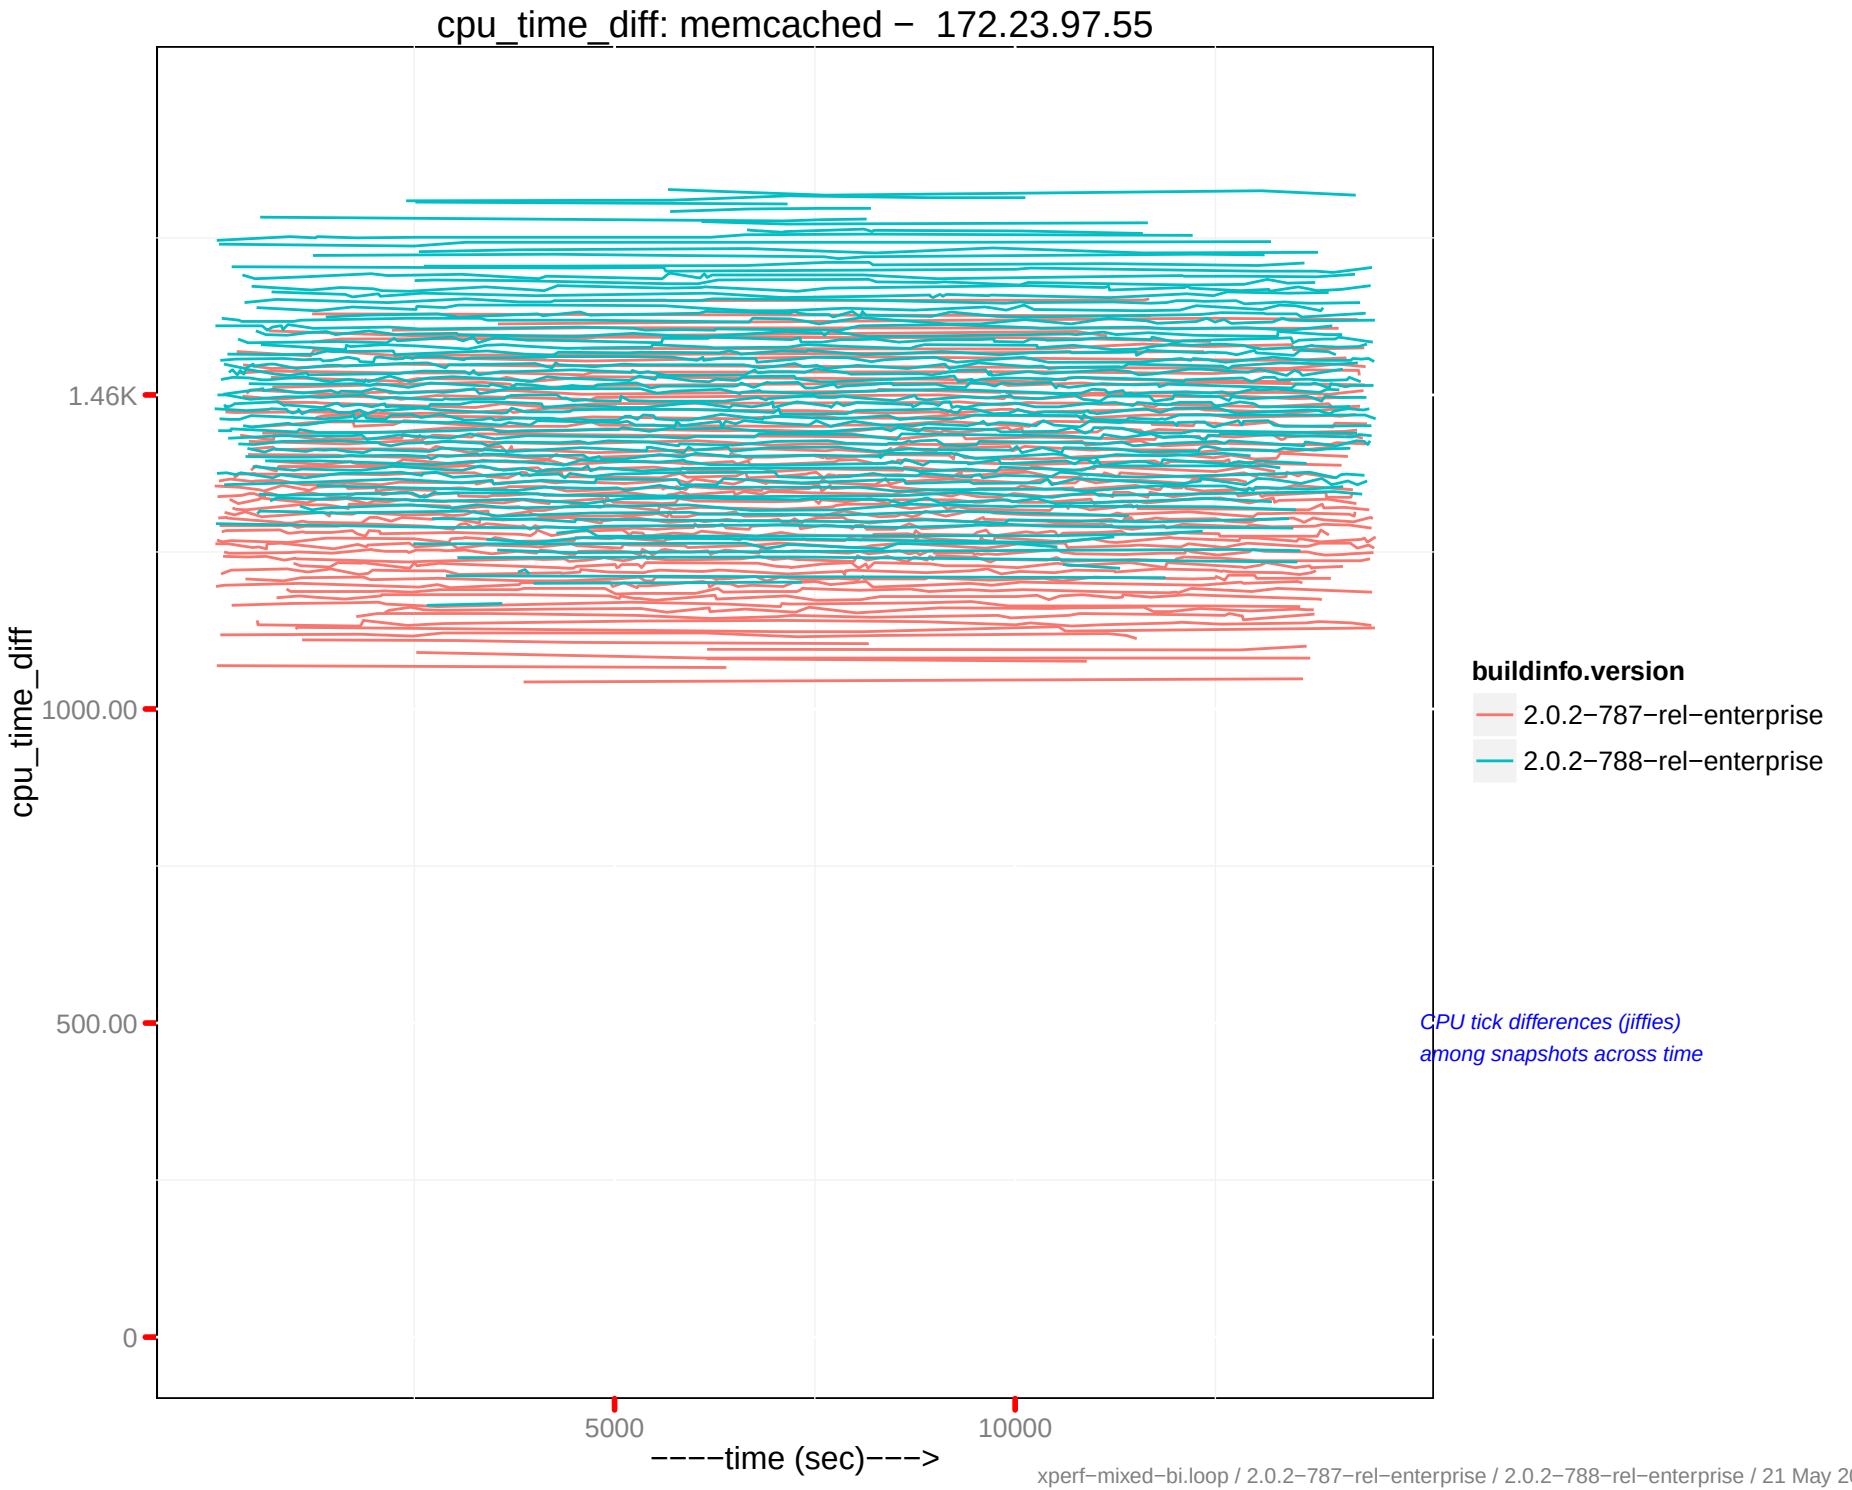
cpu_time_diff: memcached − 172.23.97.55
1.46K
1000.00
500.00
0
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

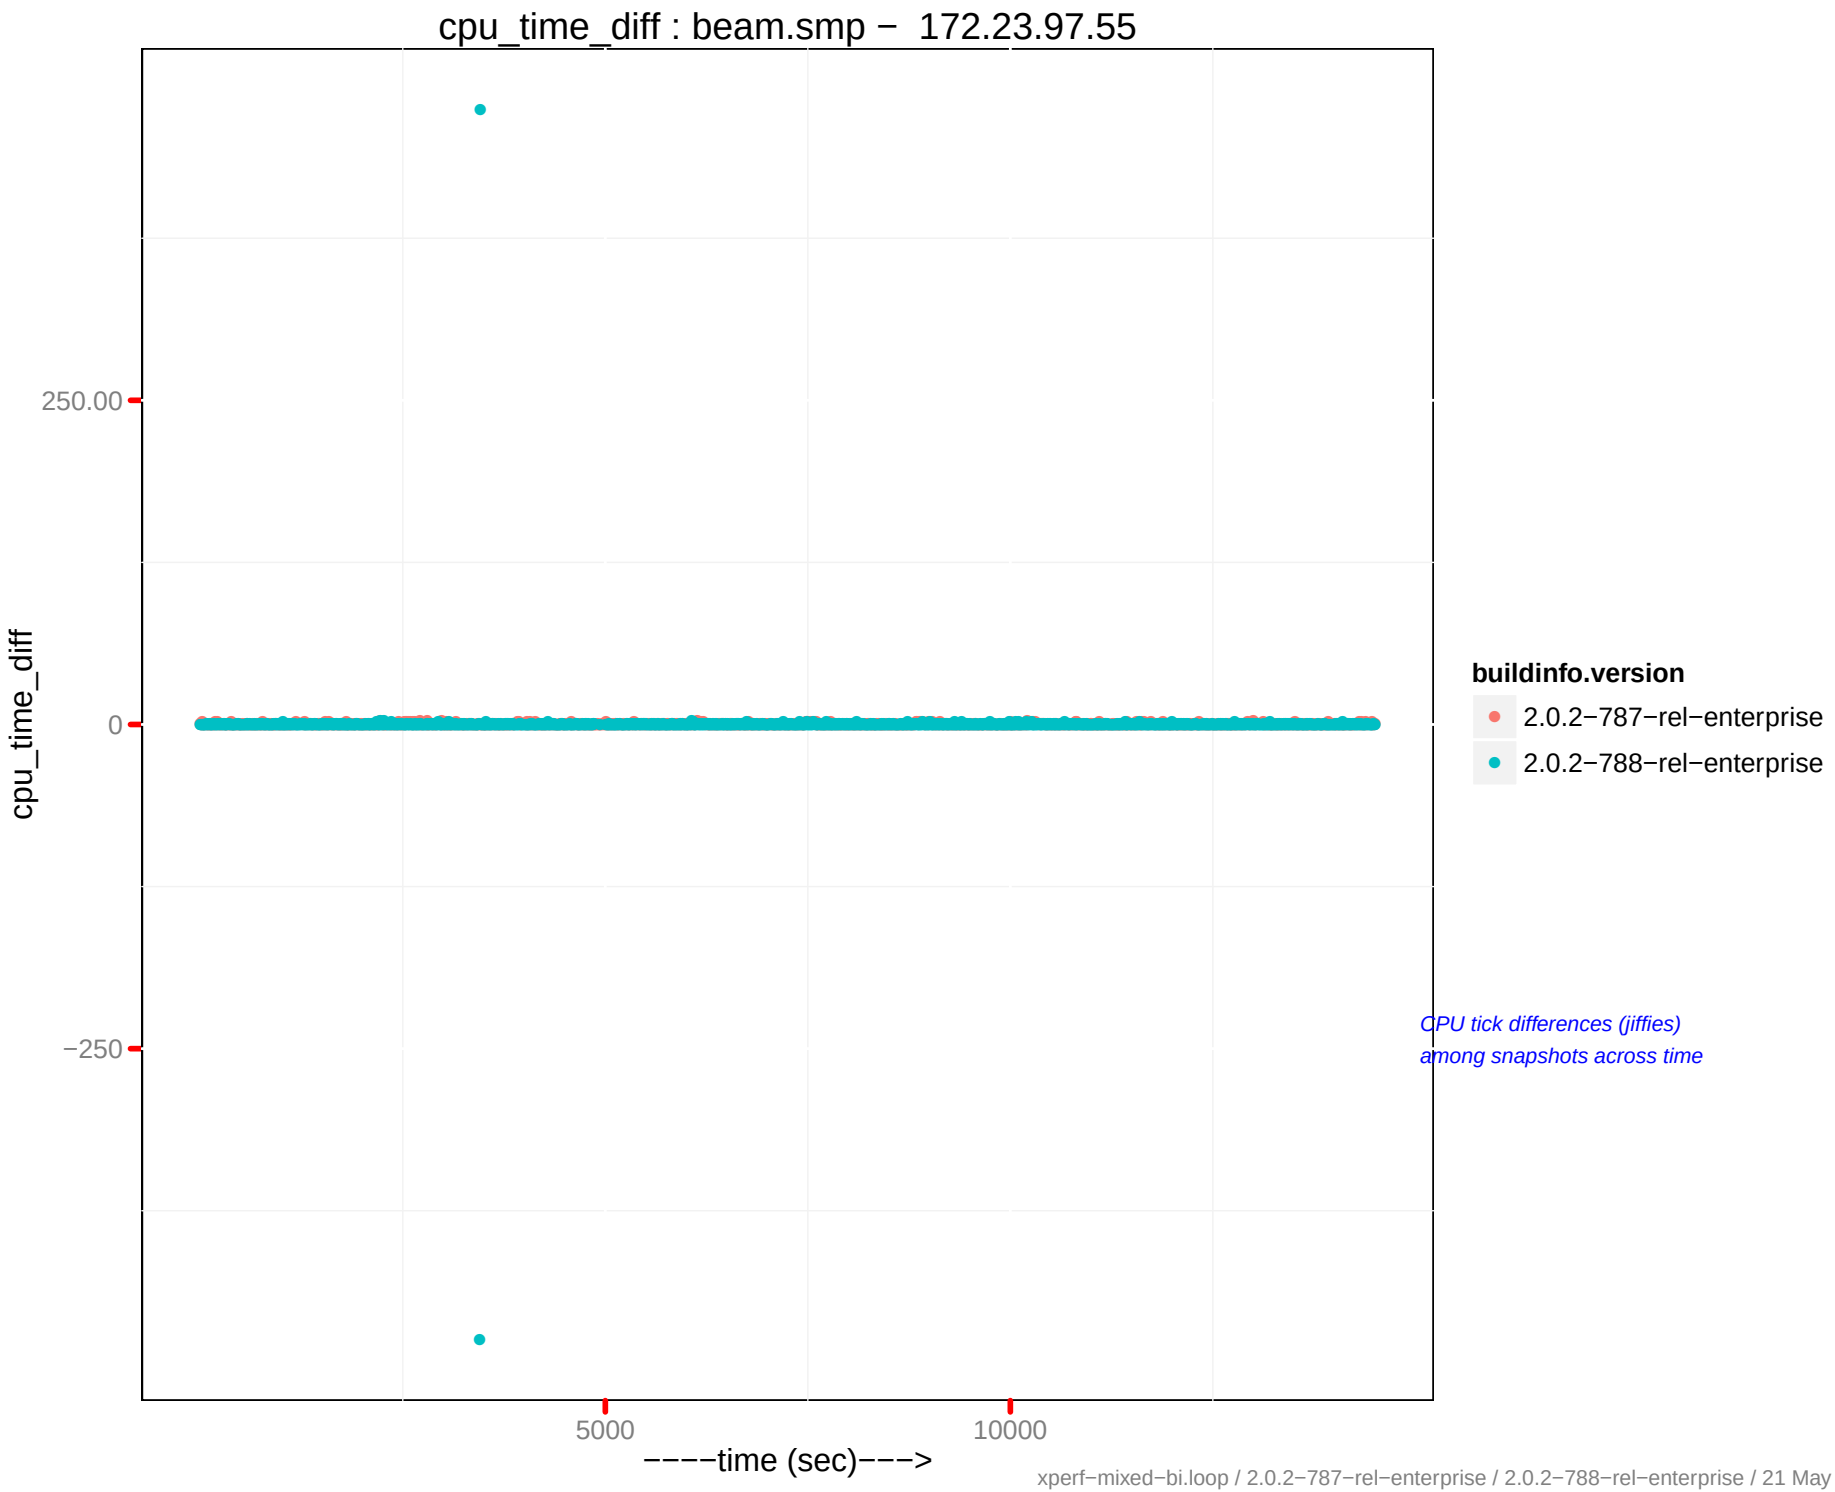
cpu_time_diff : beam.smp − 172.23.97.55
250.00
0
−250
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

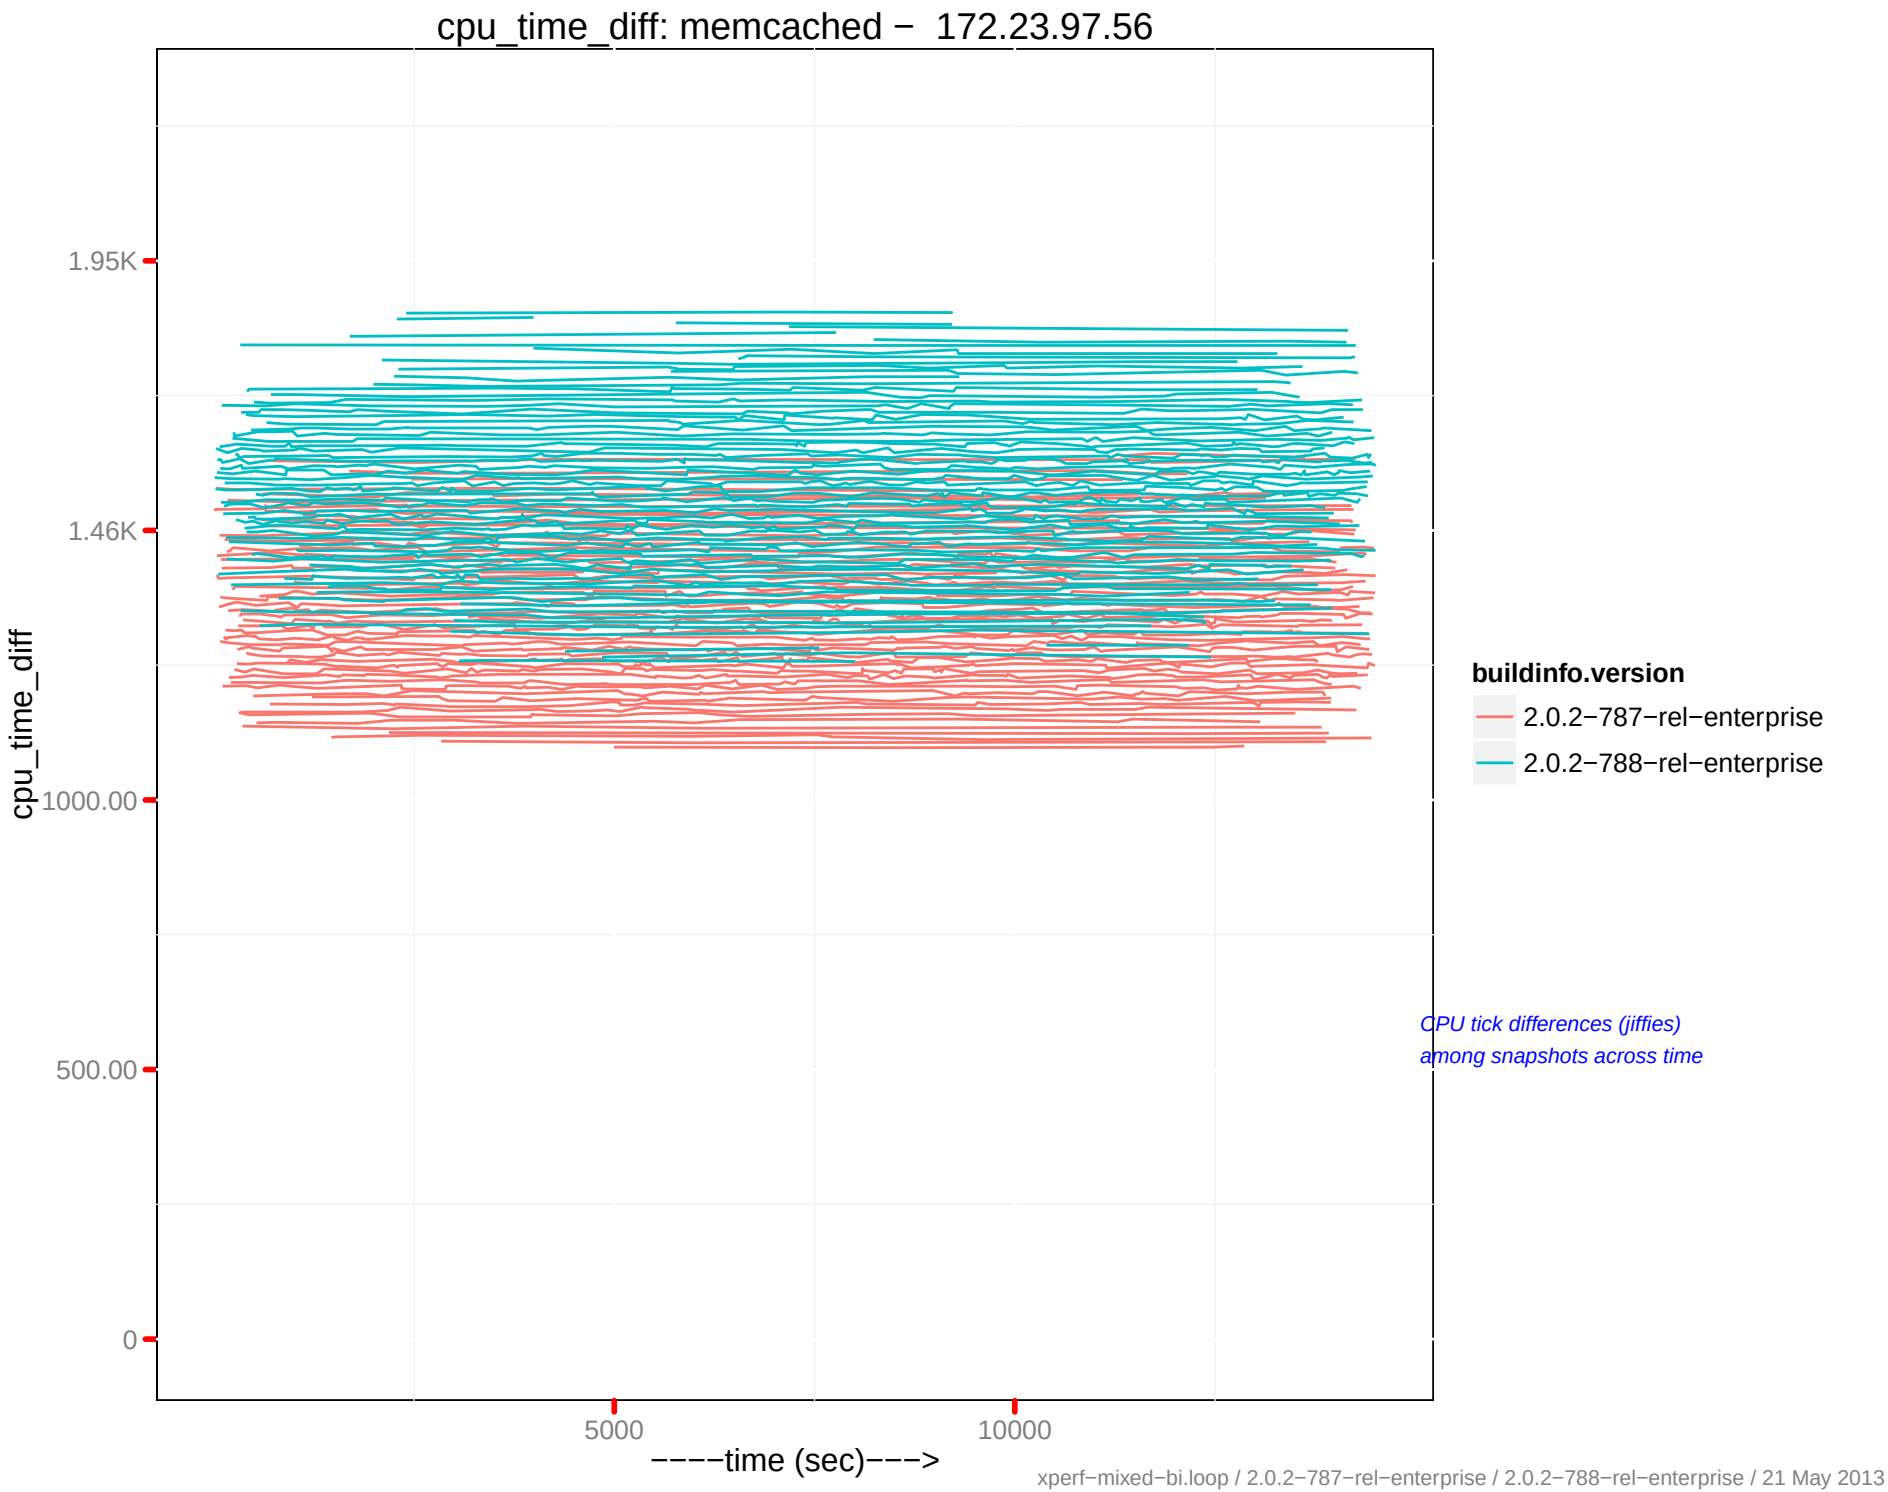
cpu_time_diff: memcached − 172.23.97.56
1.95K
1.46K
1000.00
500.00
0
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

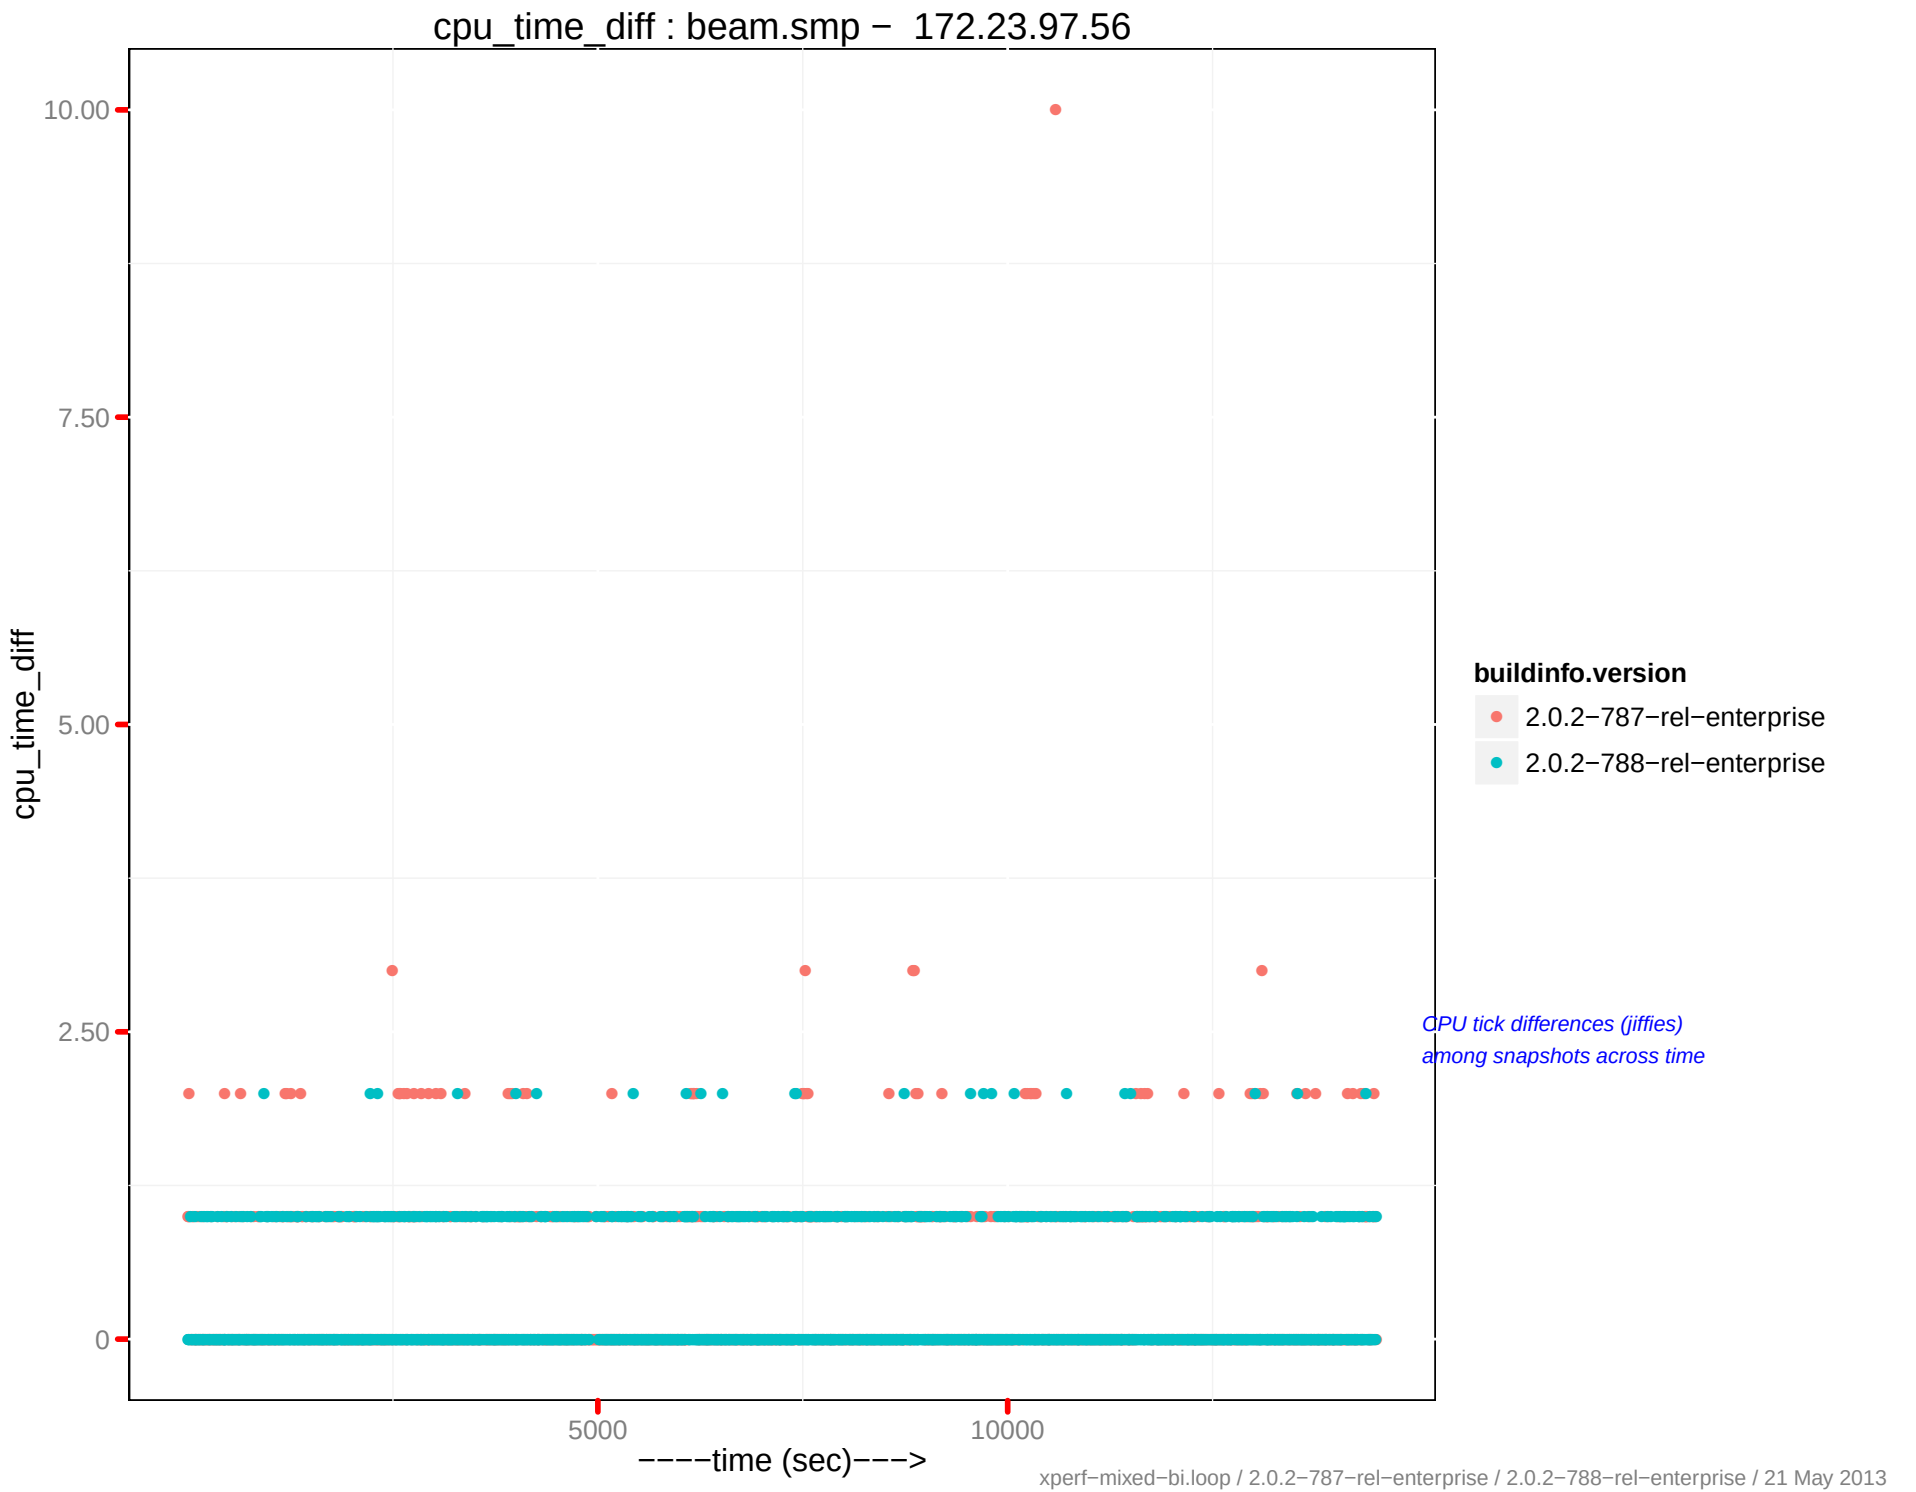
cpu_time_diff : beam.smp − 172.23.97.56
10.00
7.50
5.00
2.50
0
cpu_time_diff
5000
10000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
CPU tick differences (jiffies)
among snapshots across time
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

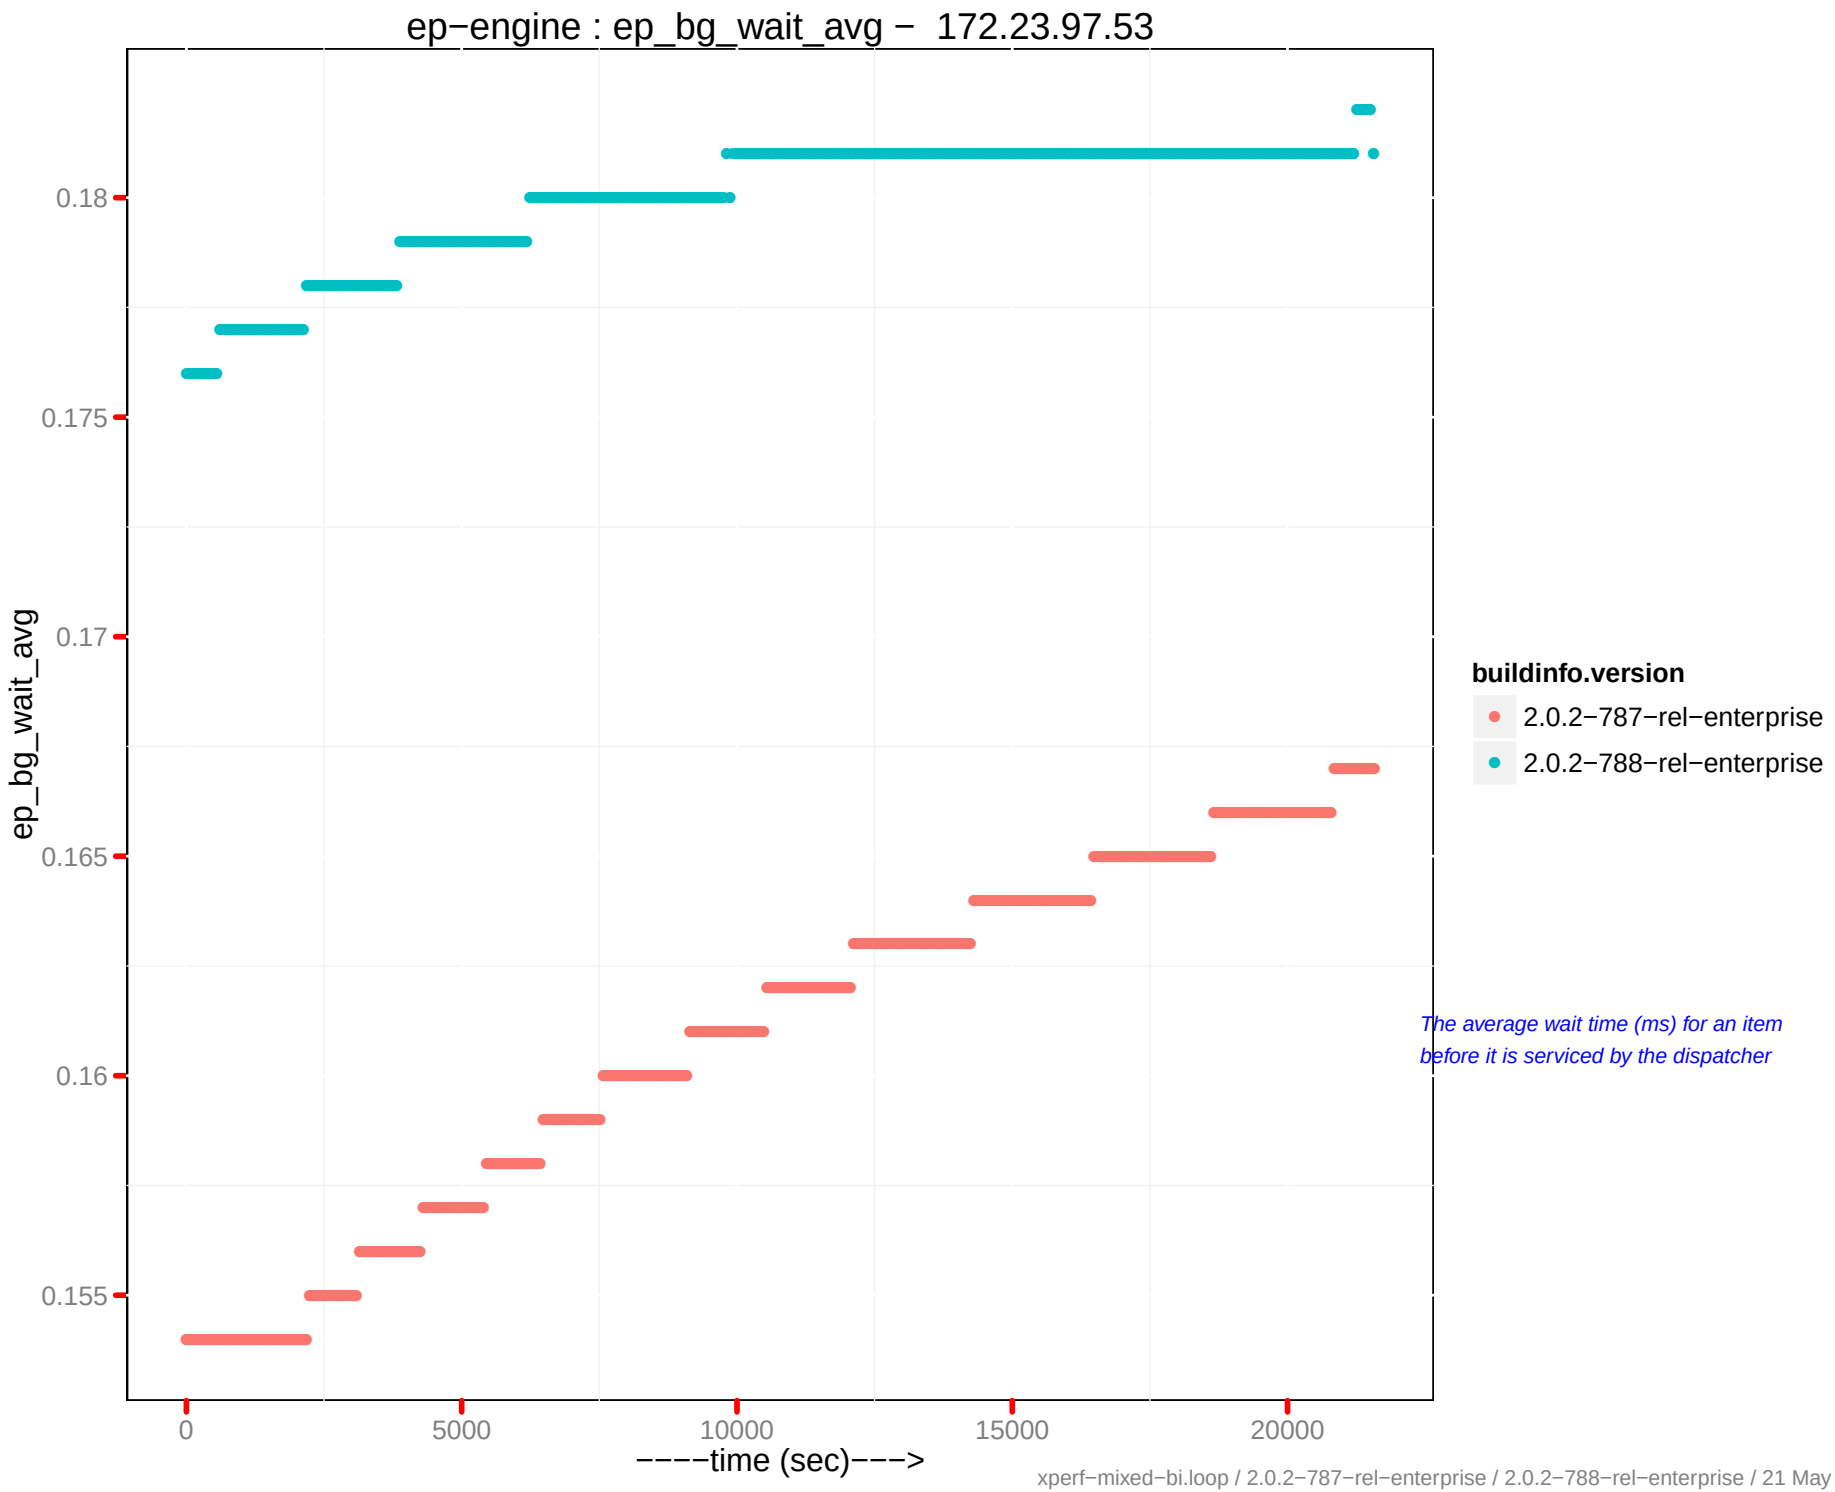
ep−engine : ep_bg_wait_avg − 172.23.97.53
ep_bg_wait_avg
0.18
0.175
0.17
0.165
0.16
0.155
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

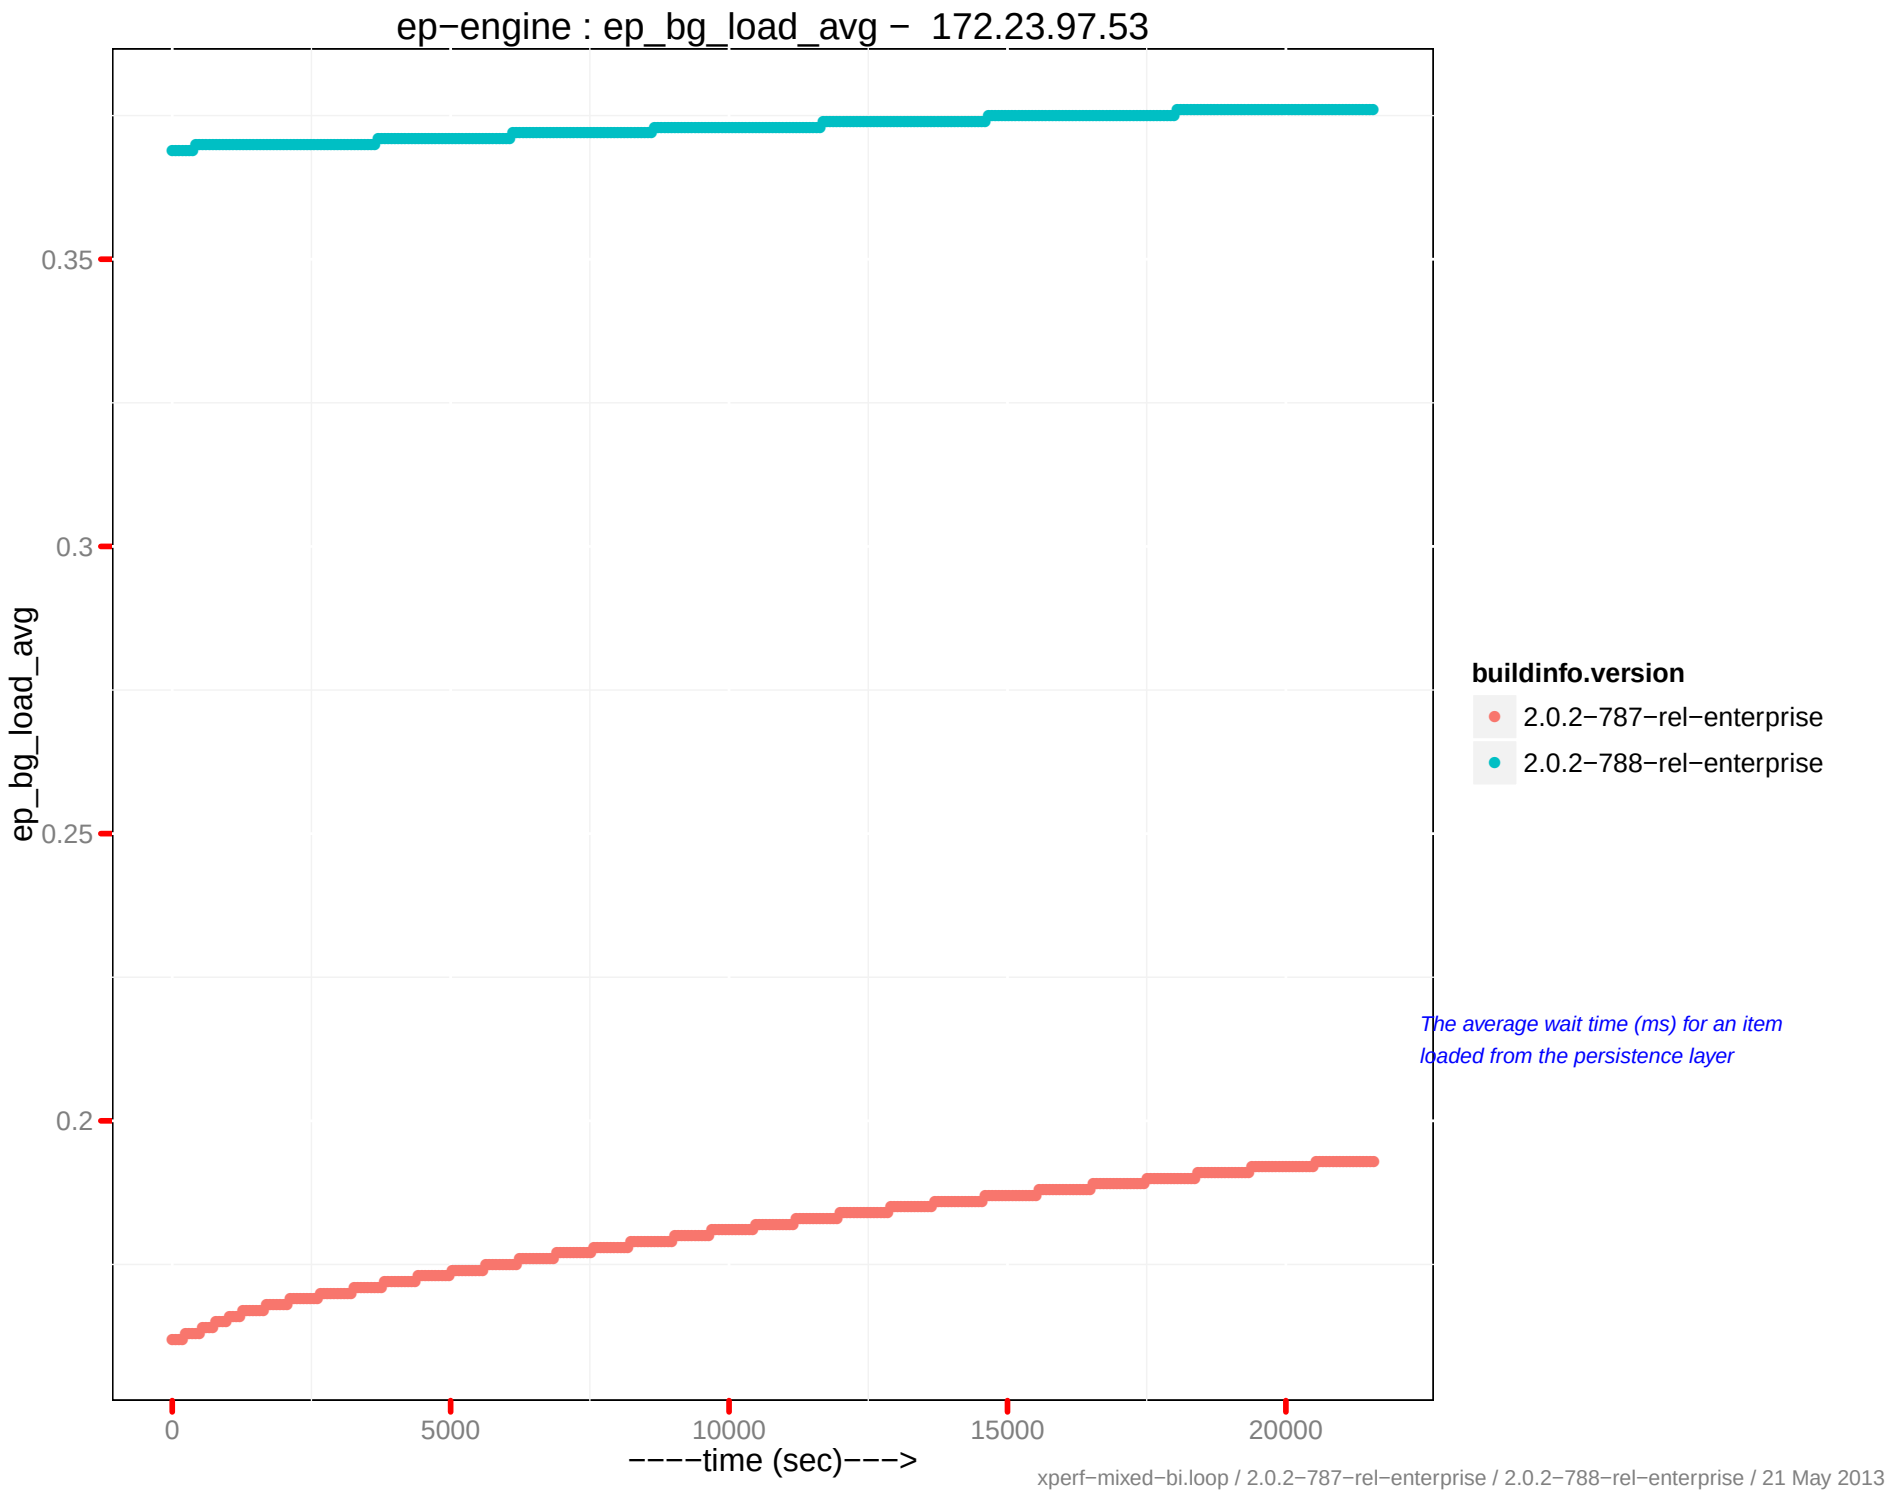
ep−engine : ep_bg_load_avg − 172.23.97.53
ep_bg_load_avg
0.35
0.3
0.25
0.2
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

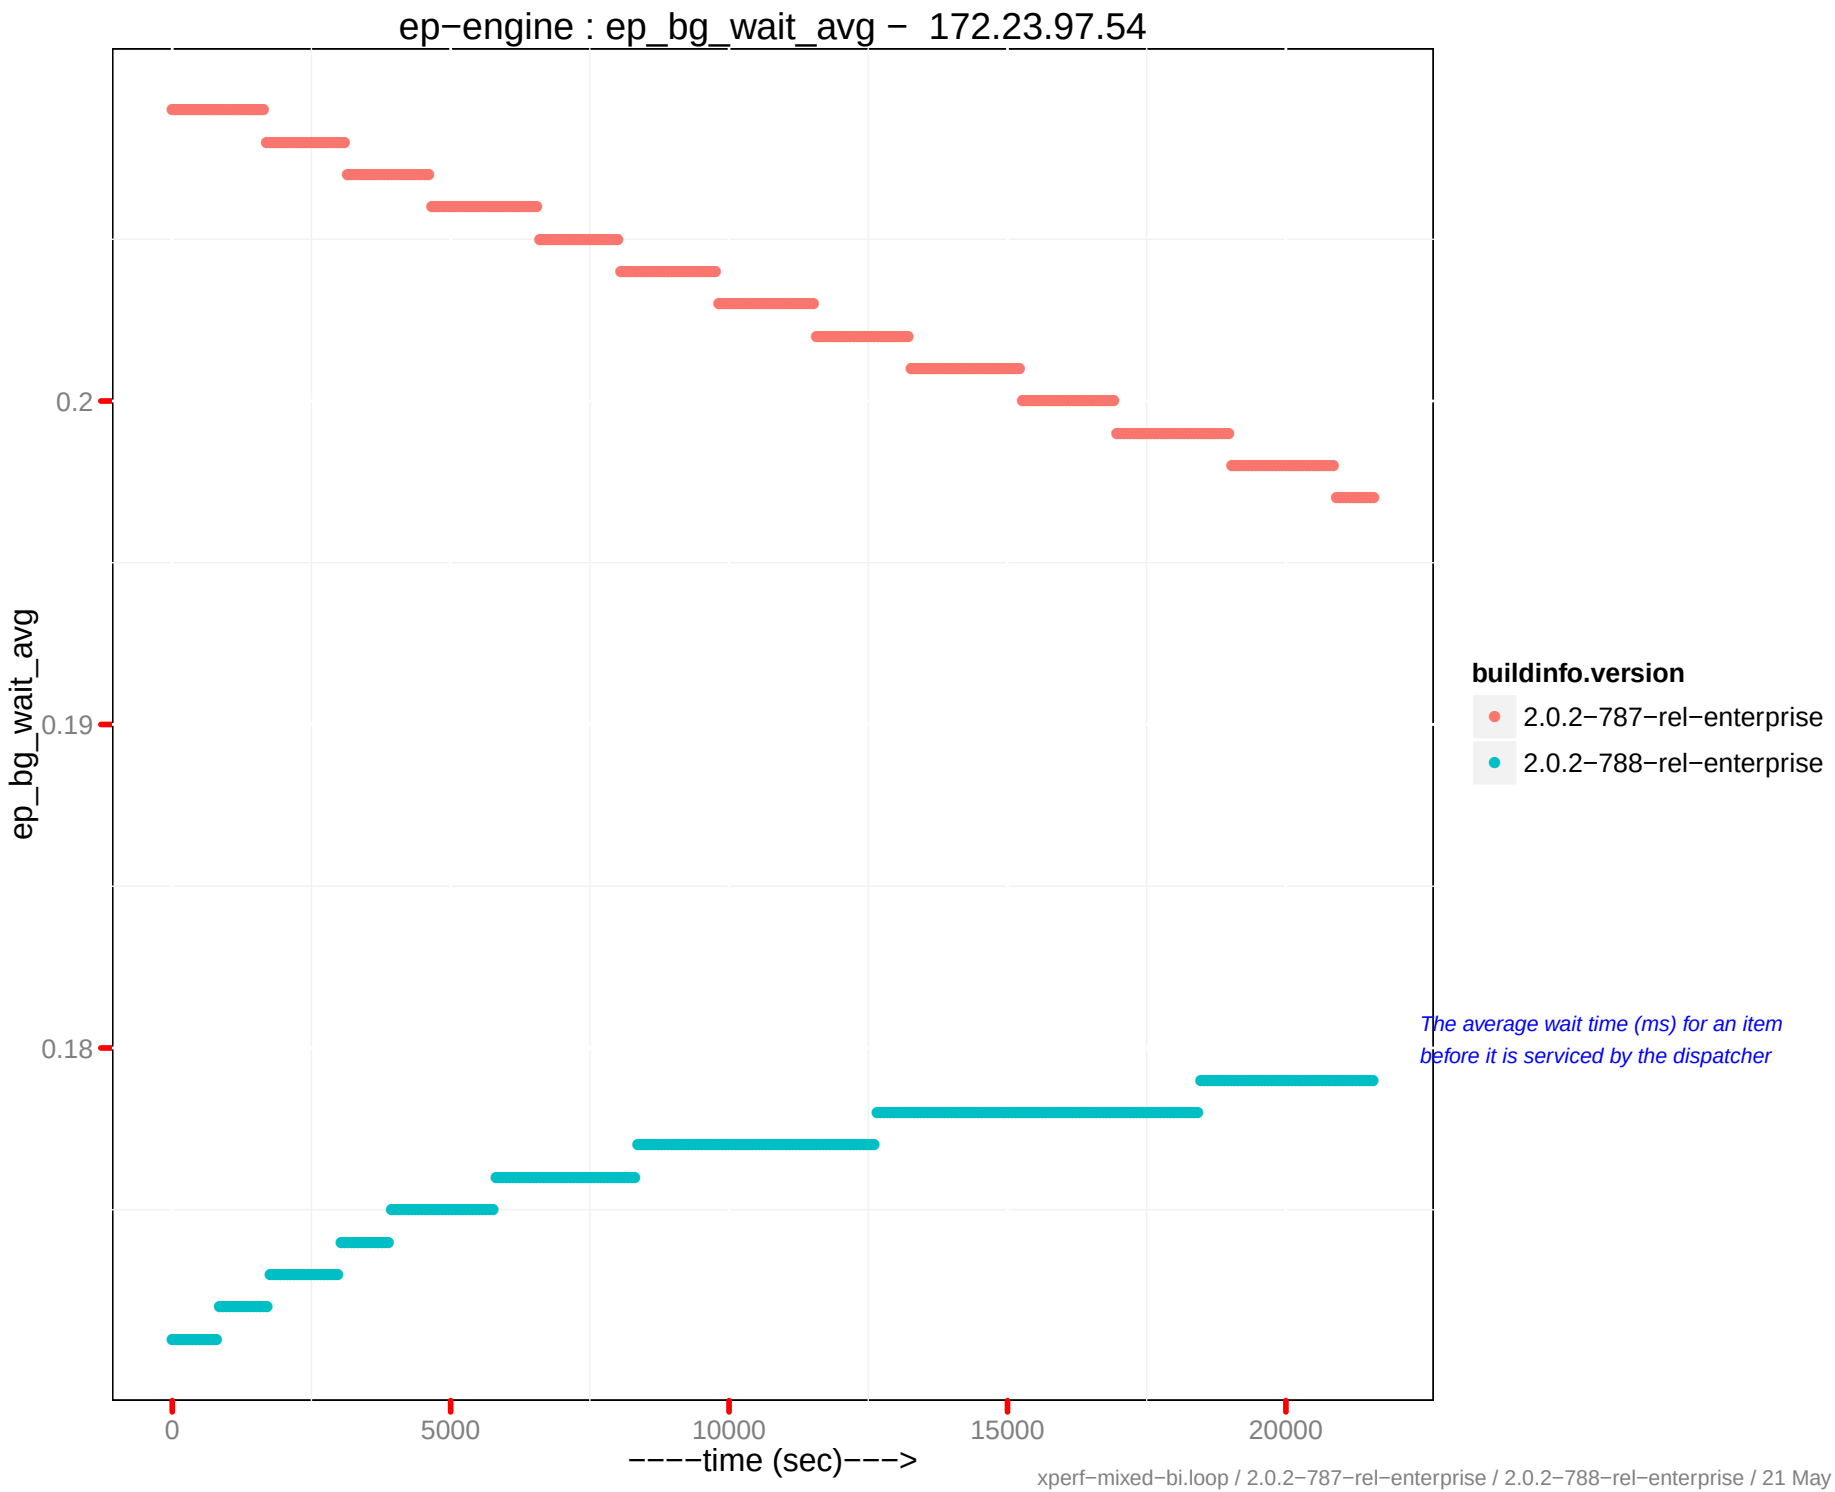
ep−engine : ep_bg_wait_avg − 172.23.97.54
ep_bg_wait_avg
0.2
0.19
0.18
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

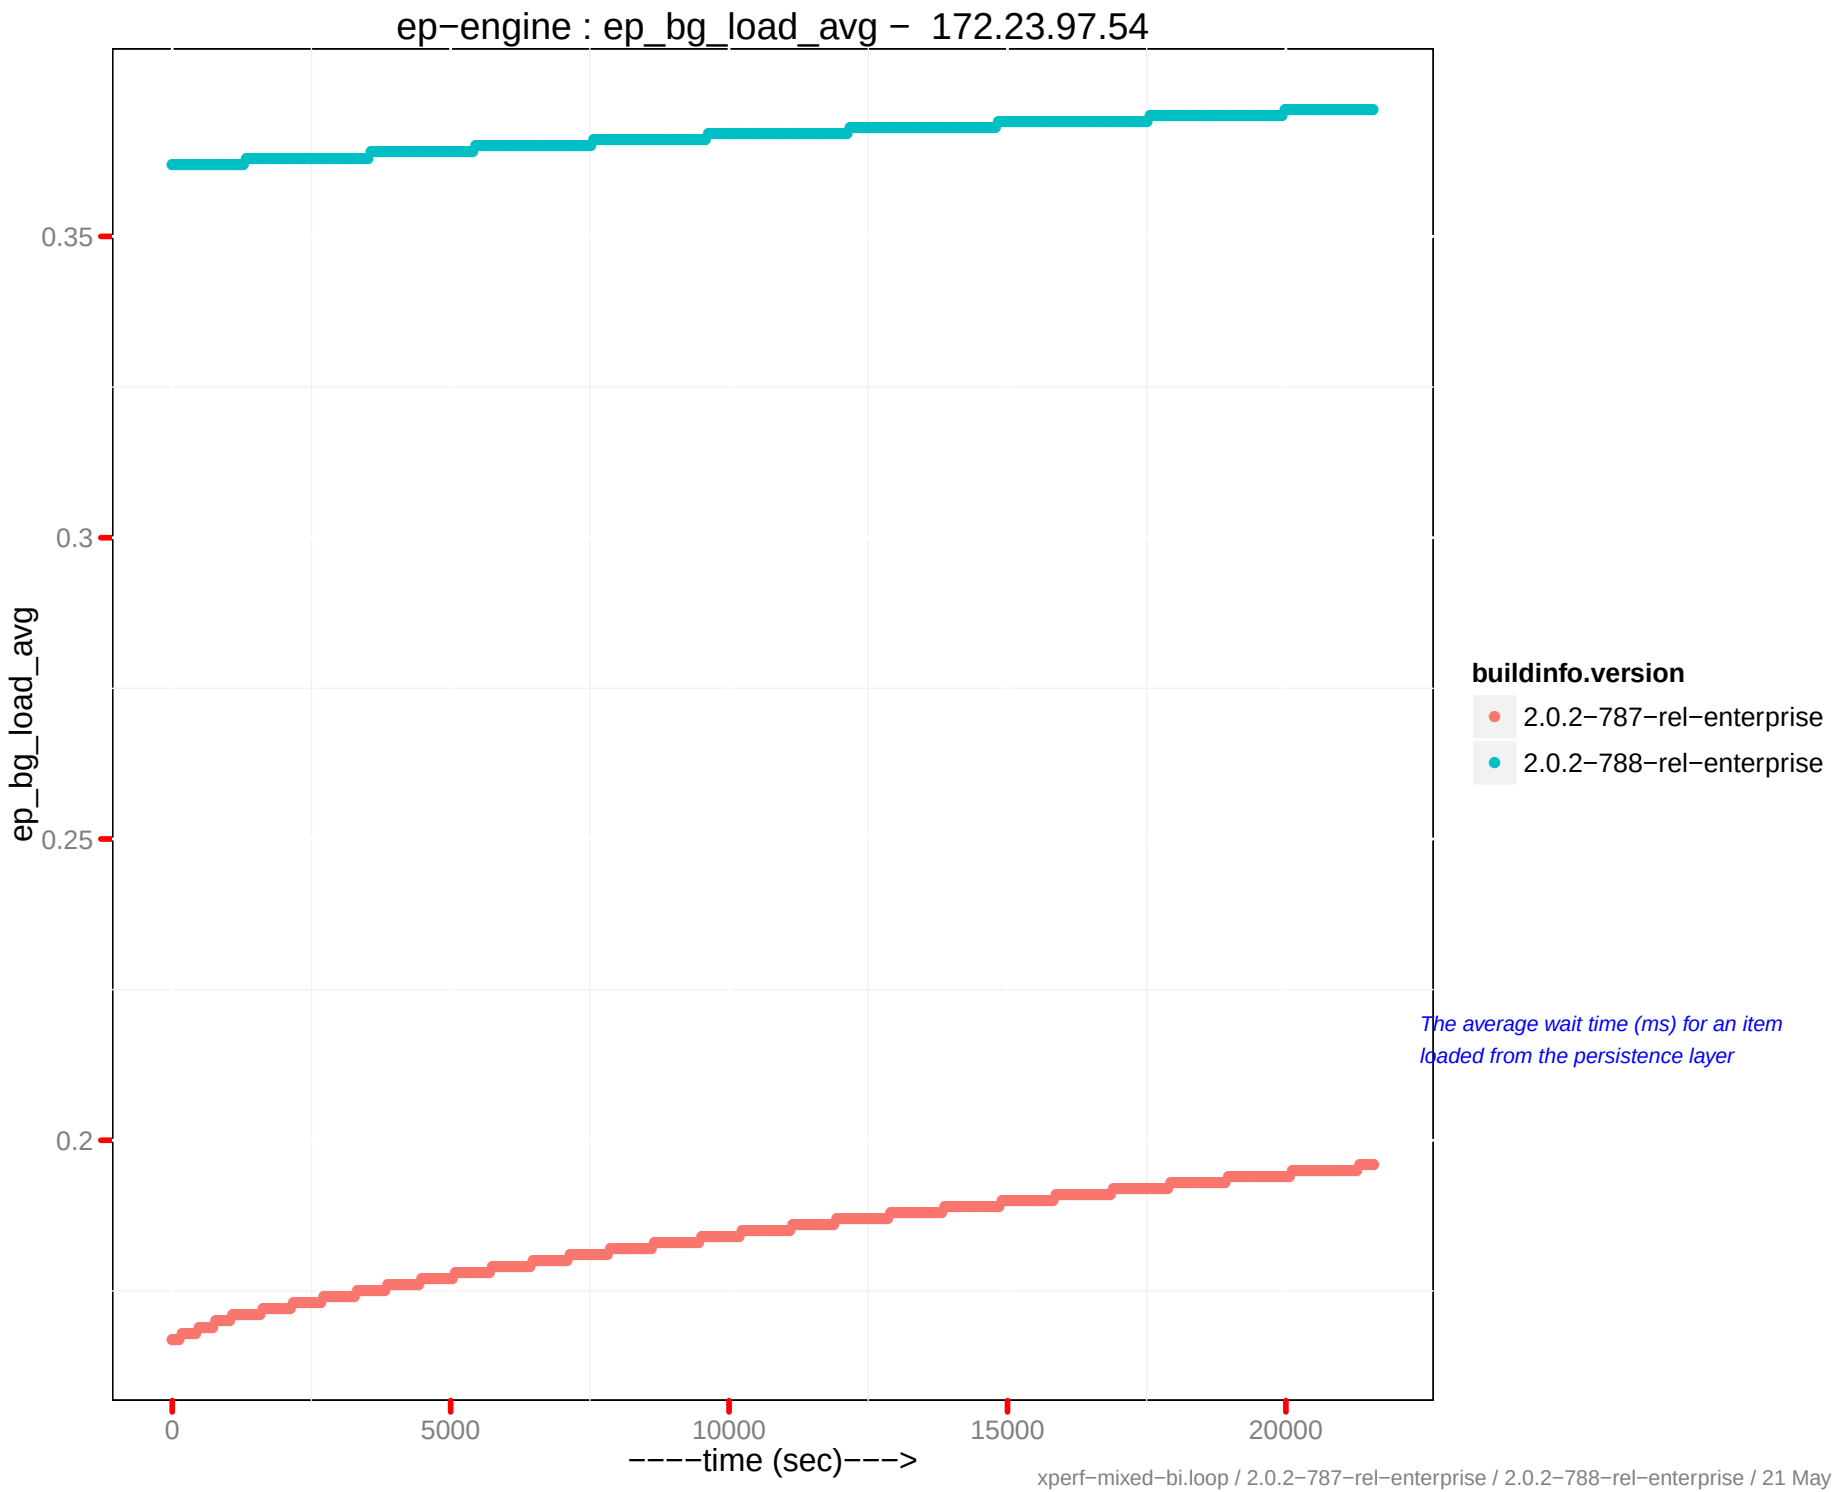
ep−engine : ep_bg_load_avg − 172.23.97.54
ep_bg_load_avg
0.35
0.3
0.25
0.2
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

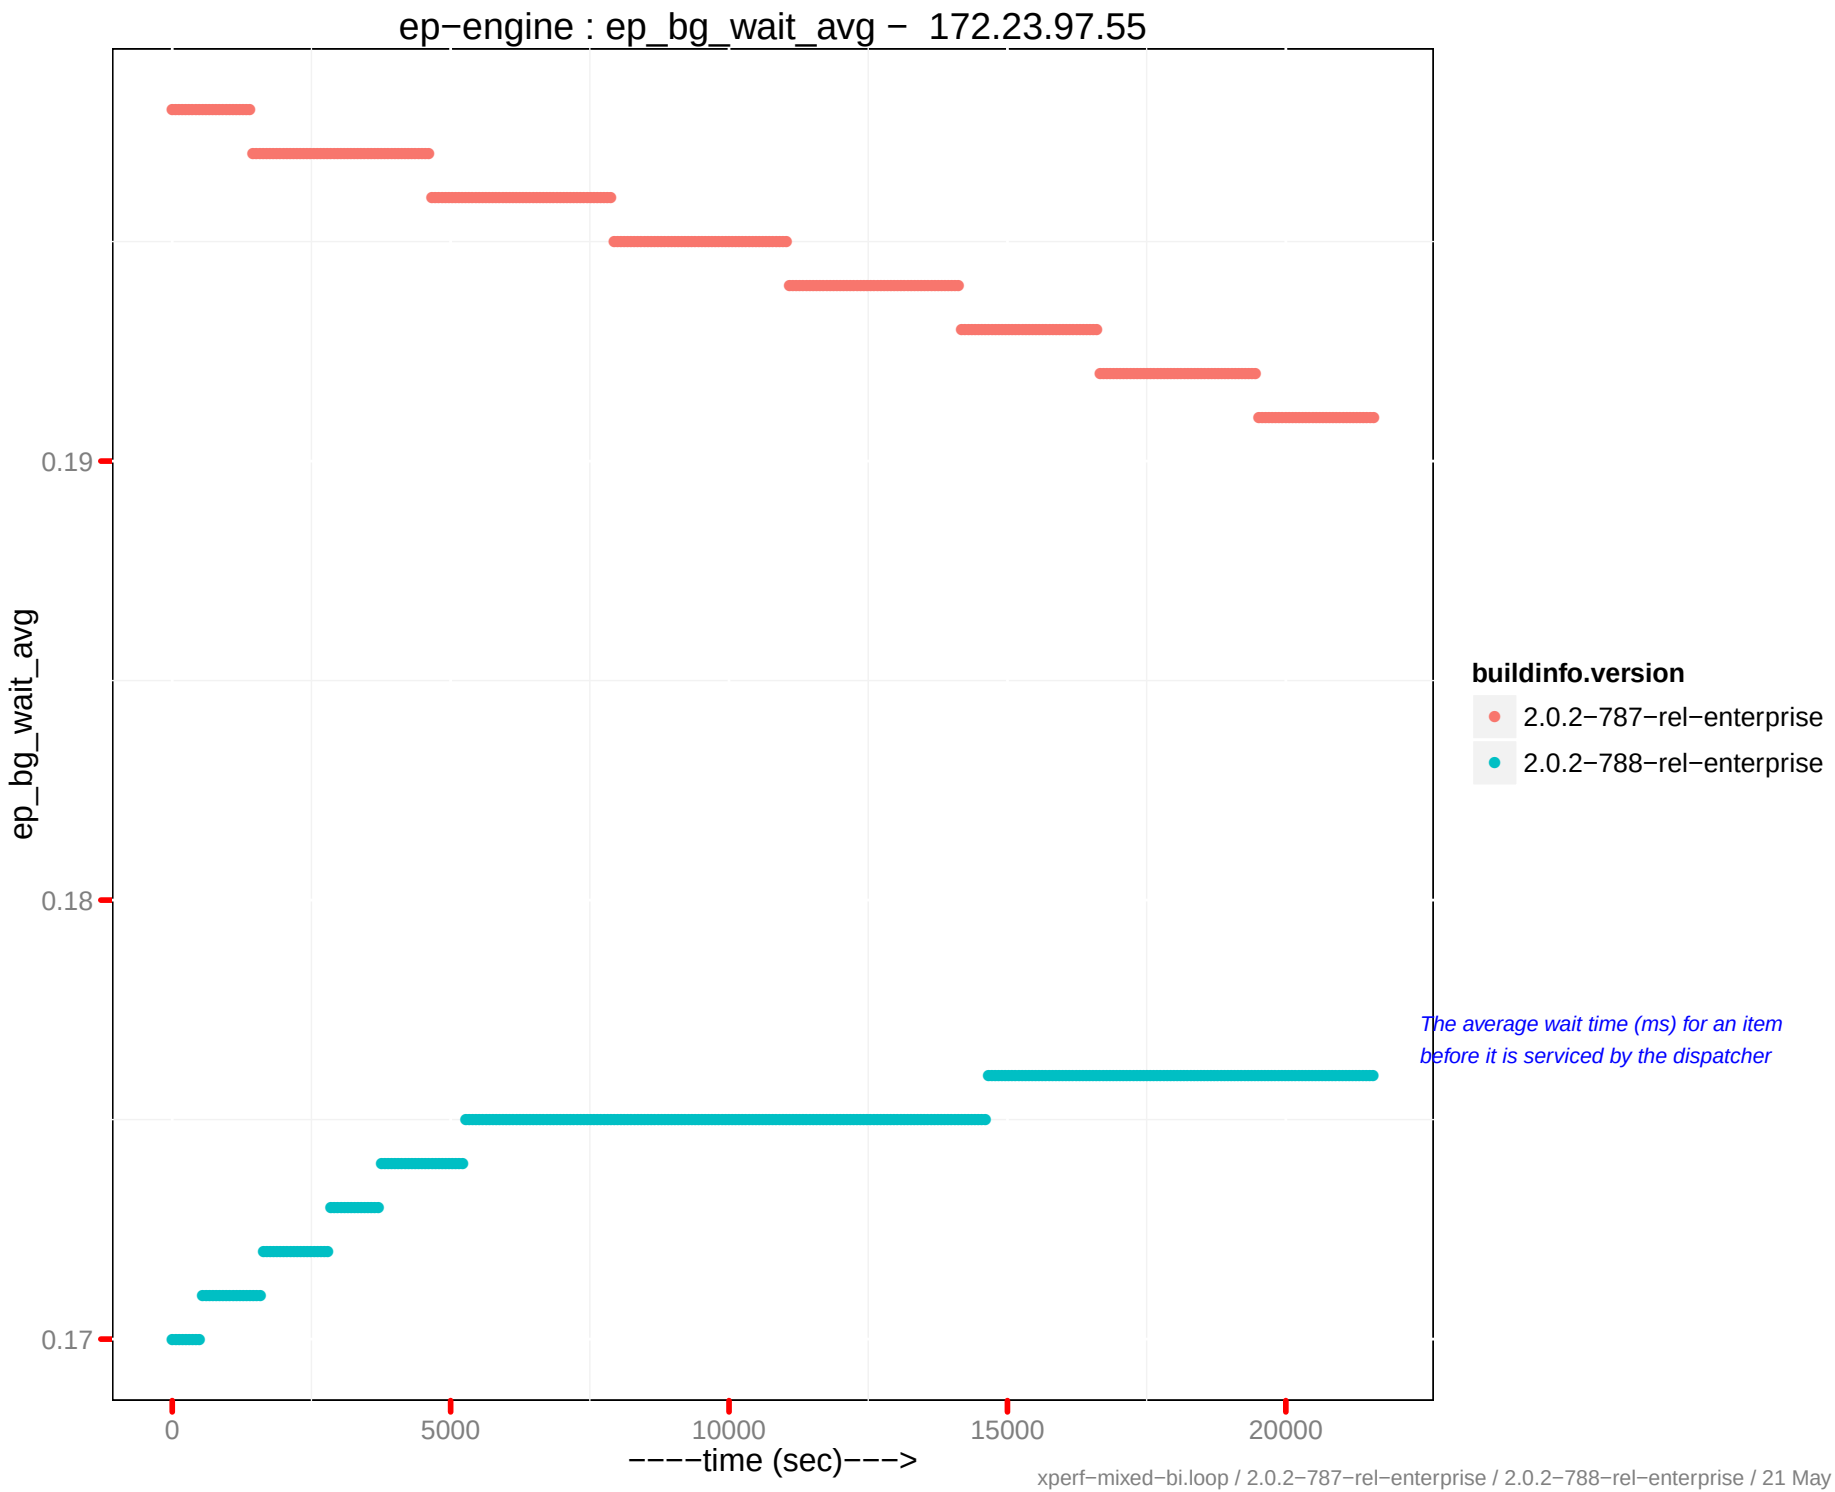
ep−engine : ep_bg_wait_avg − 172.23.97.55
ep_bg_wait_avg
0.19
0.18
0.17
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

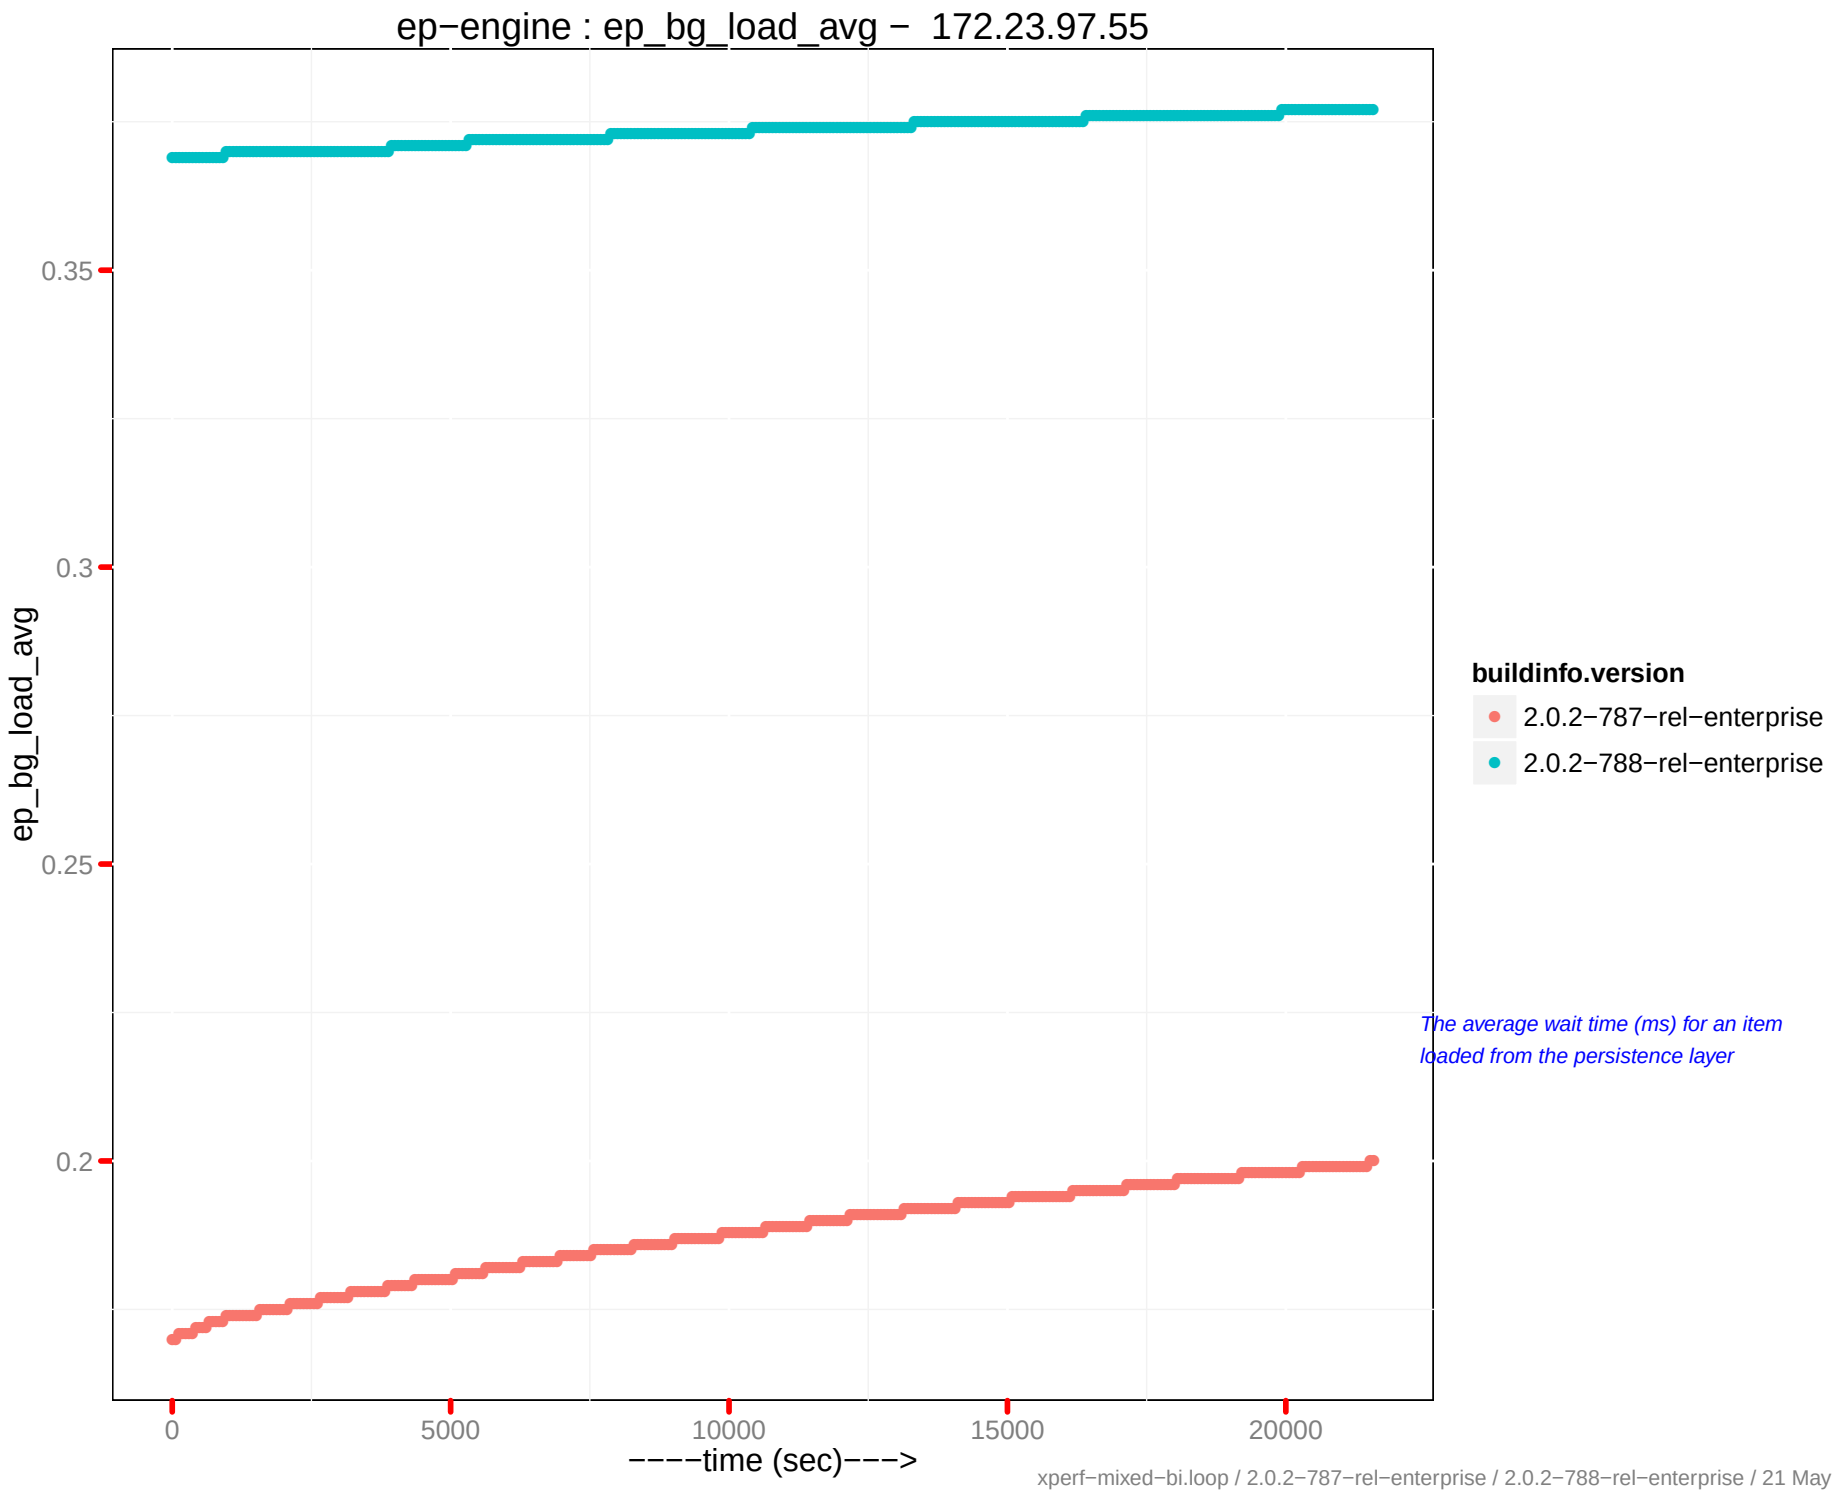
ep−engine : ep_bg_load_avg − 172.23.97.55
ep_bg_load_avg
0.35
0.3
0.25
0.2
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

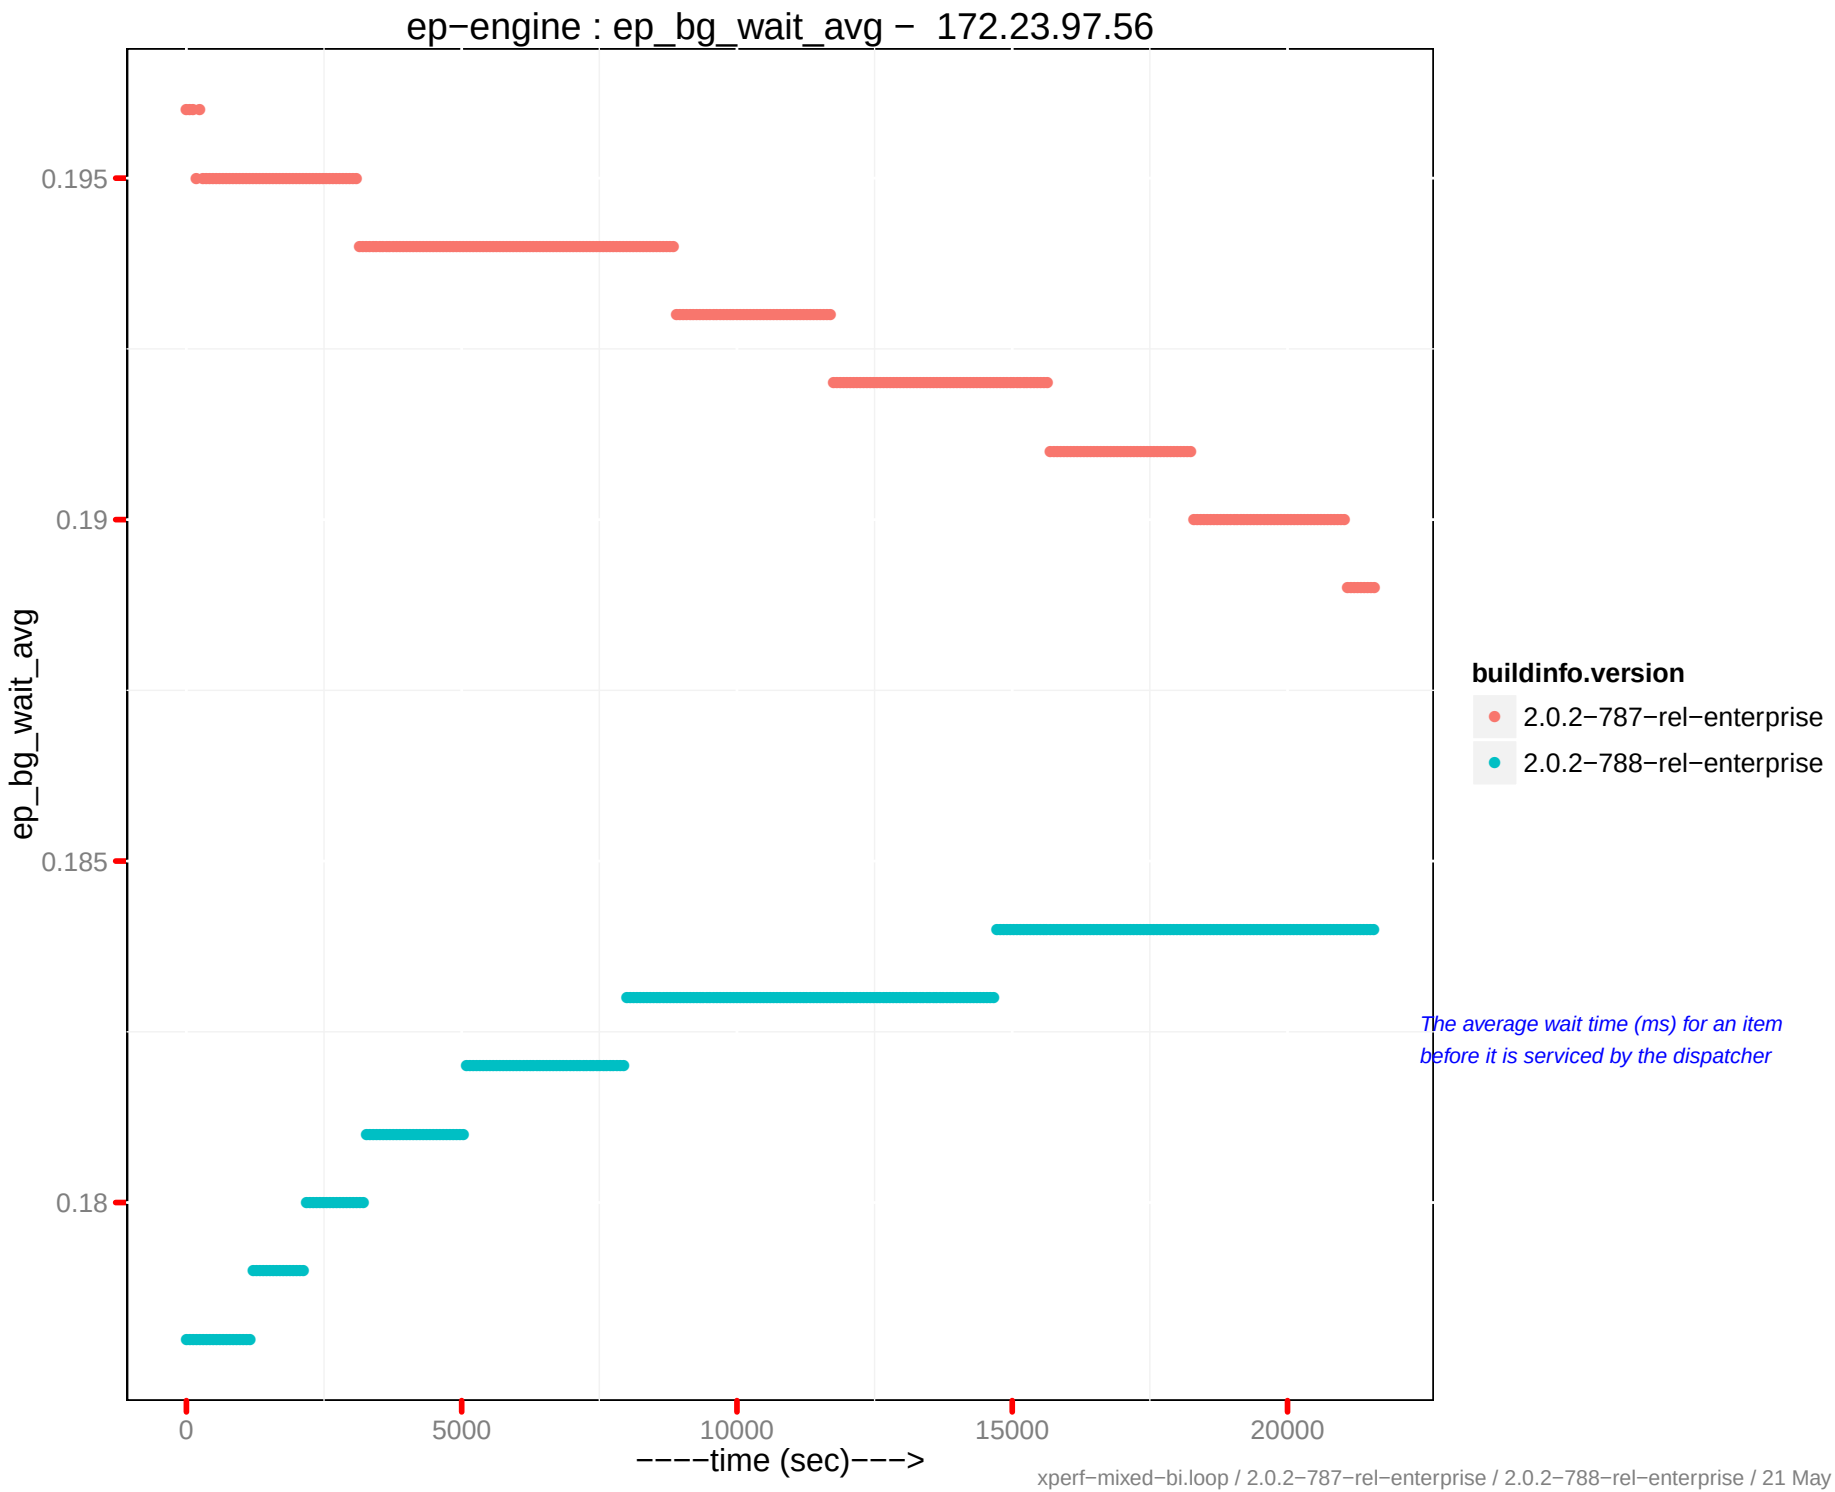
ep−engine : ep_bg_wait_avg − 172.23.97.56
ep_bg_wait_avg
0.195
0.19
0.185
0.18
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item
before it is serviced by the dispatcher
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

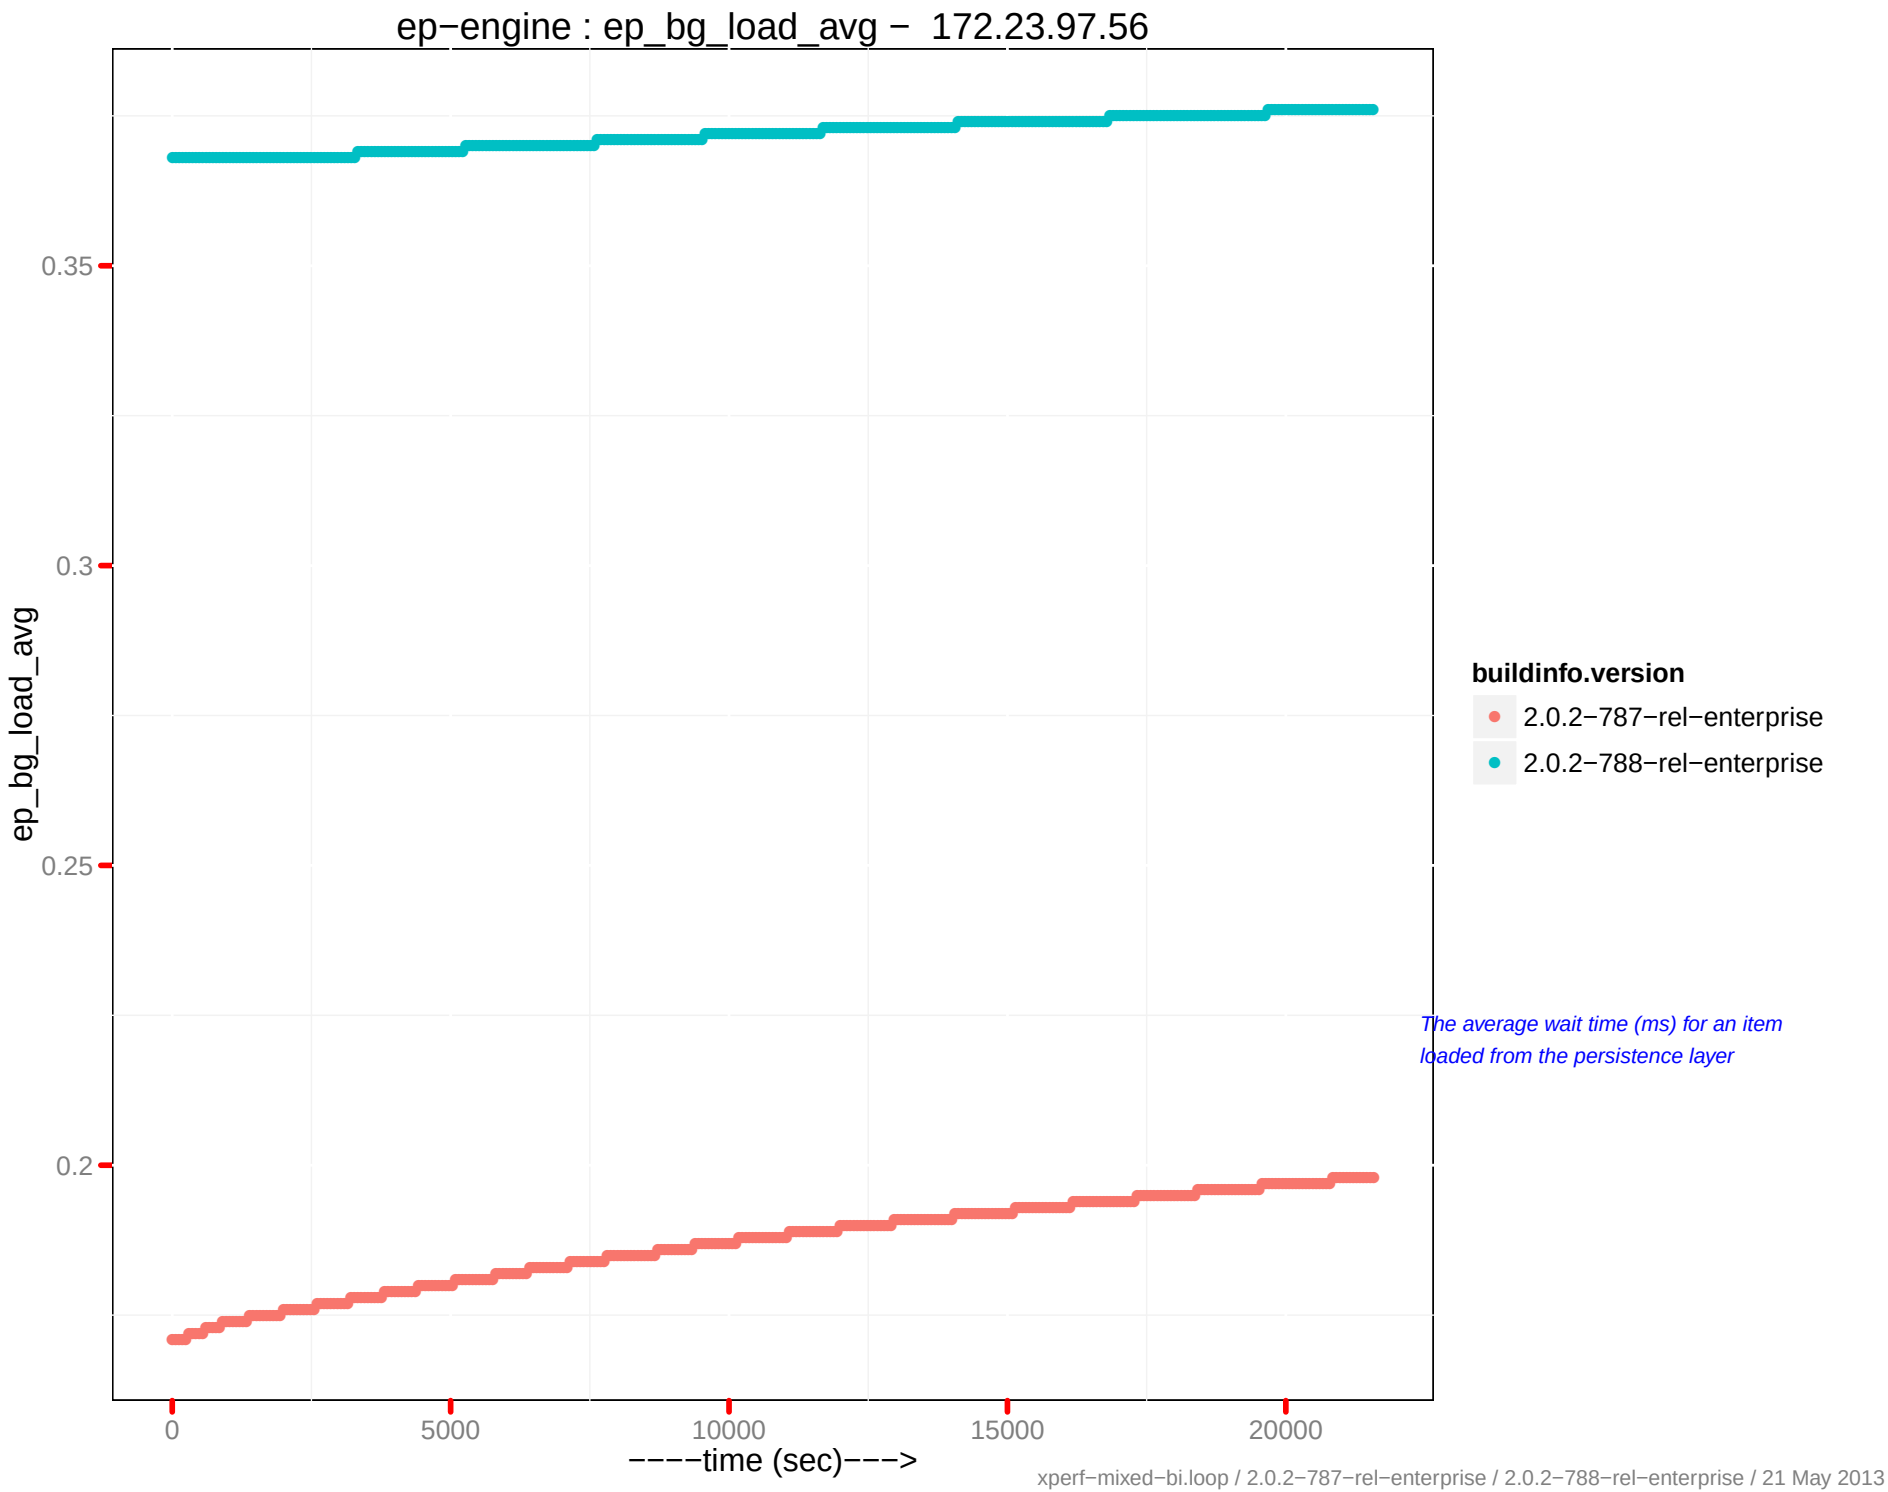
ep−engine : ep_bg_load_avg − 172.23.97.56
ep_bg_load_avg
0.35
0.3
0.25
0.2
0
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
The average wait time (ms) for an item loaded from the persistence layer
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

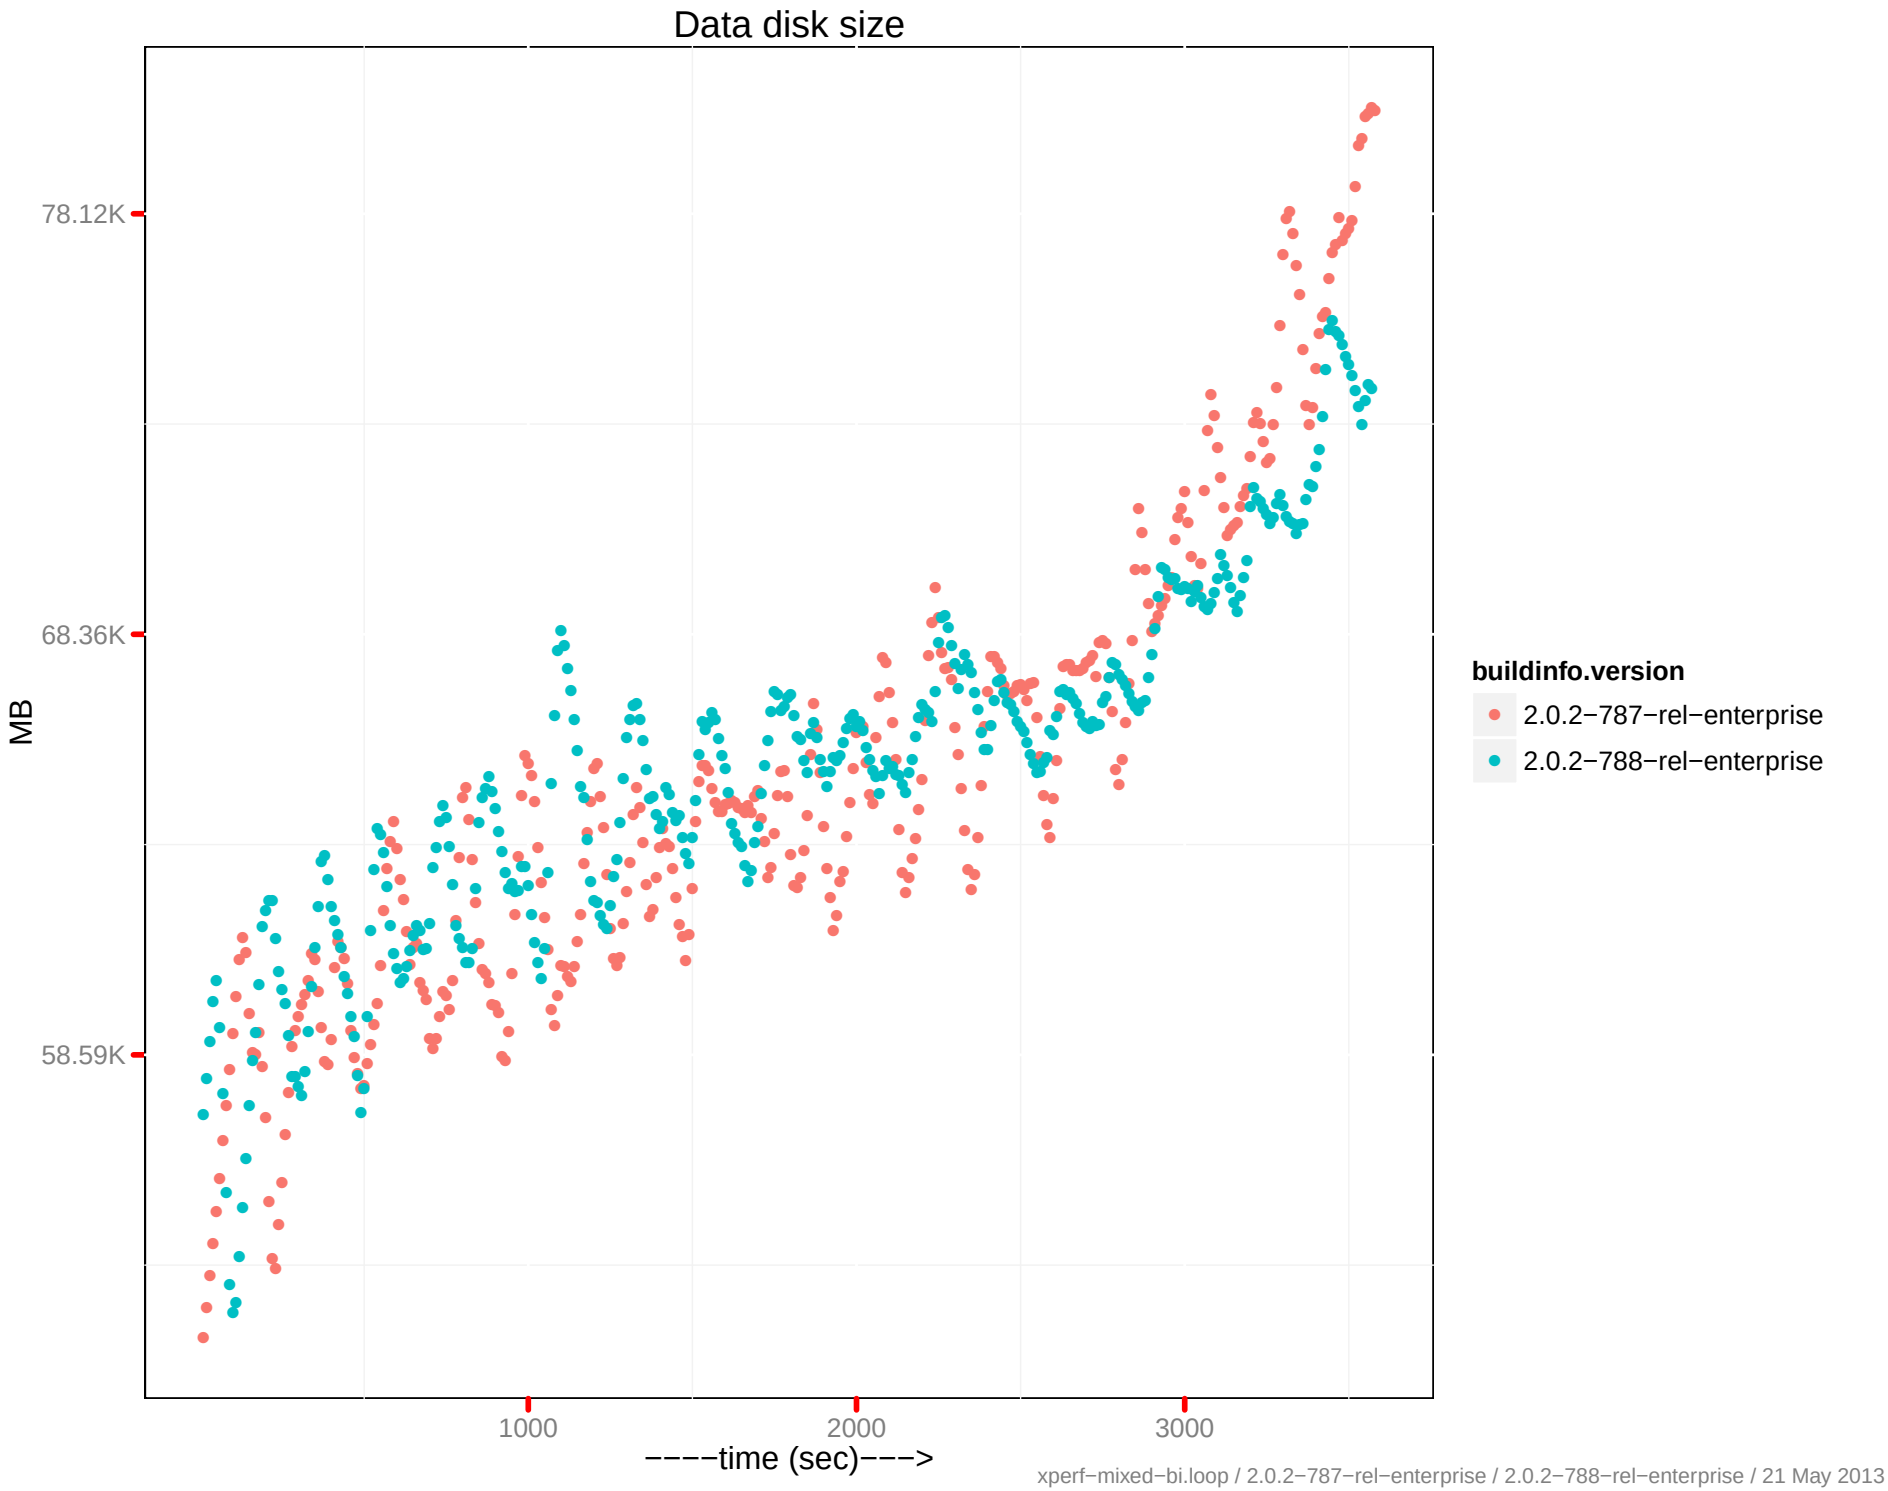
Data disk size
78.12K
68.36K
58.59K
MB
1000
2000
3000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

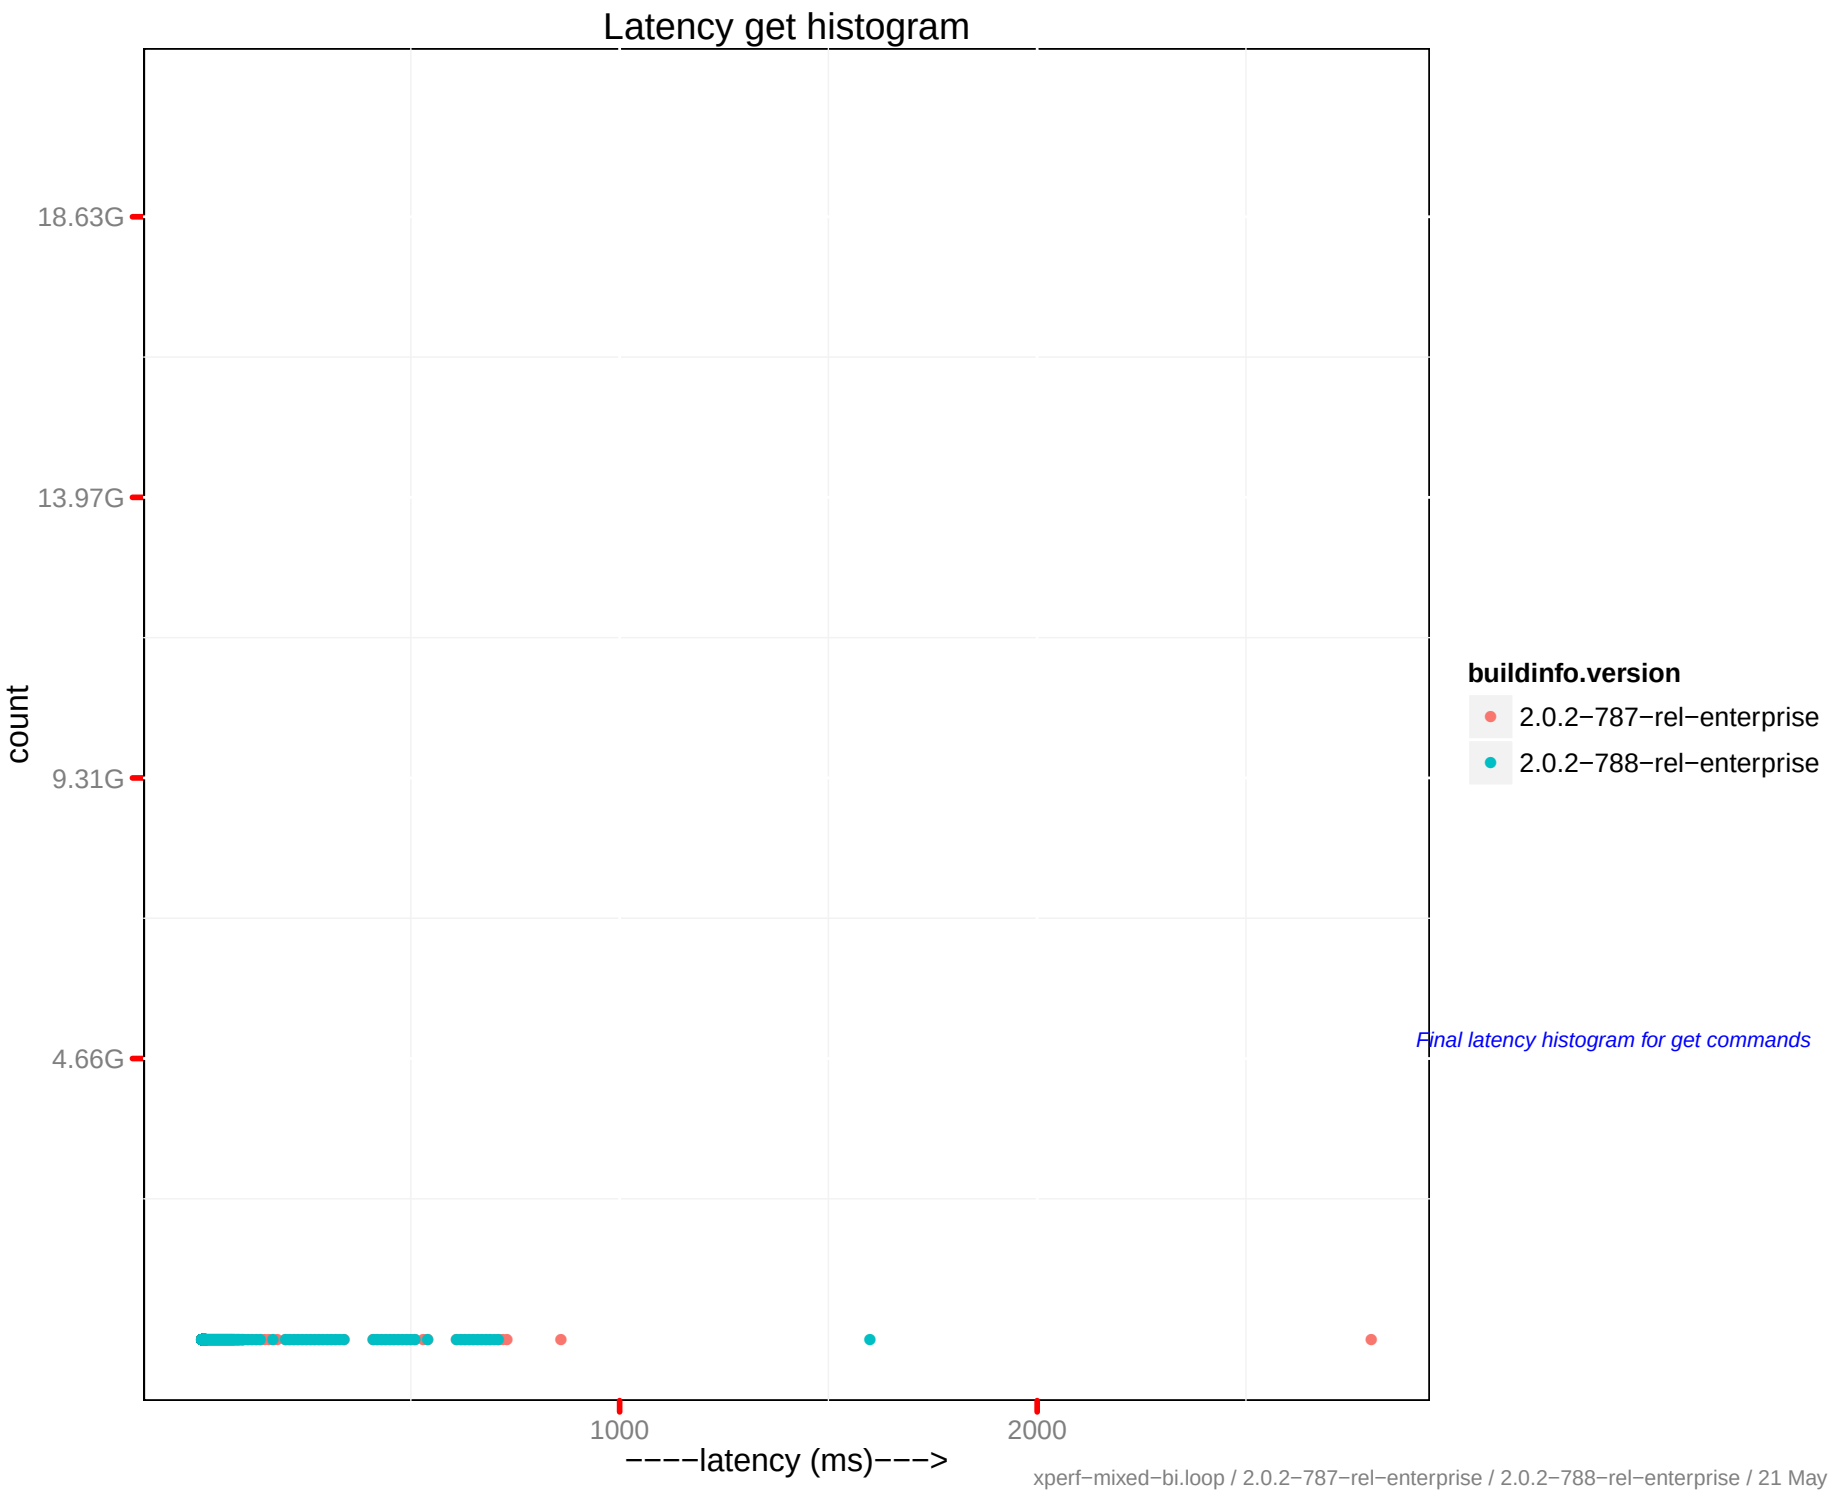
Latency get histogram
18.63G
13.97G
9.31G
4.66G
count
1000
2000
----latency (ms)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Final latency histogram for get commands
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

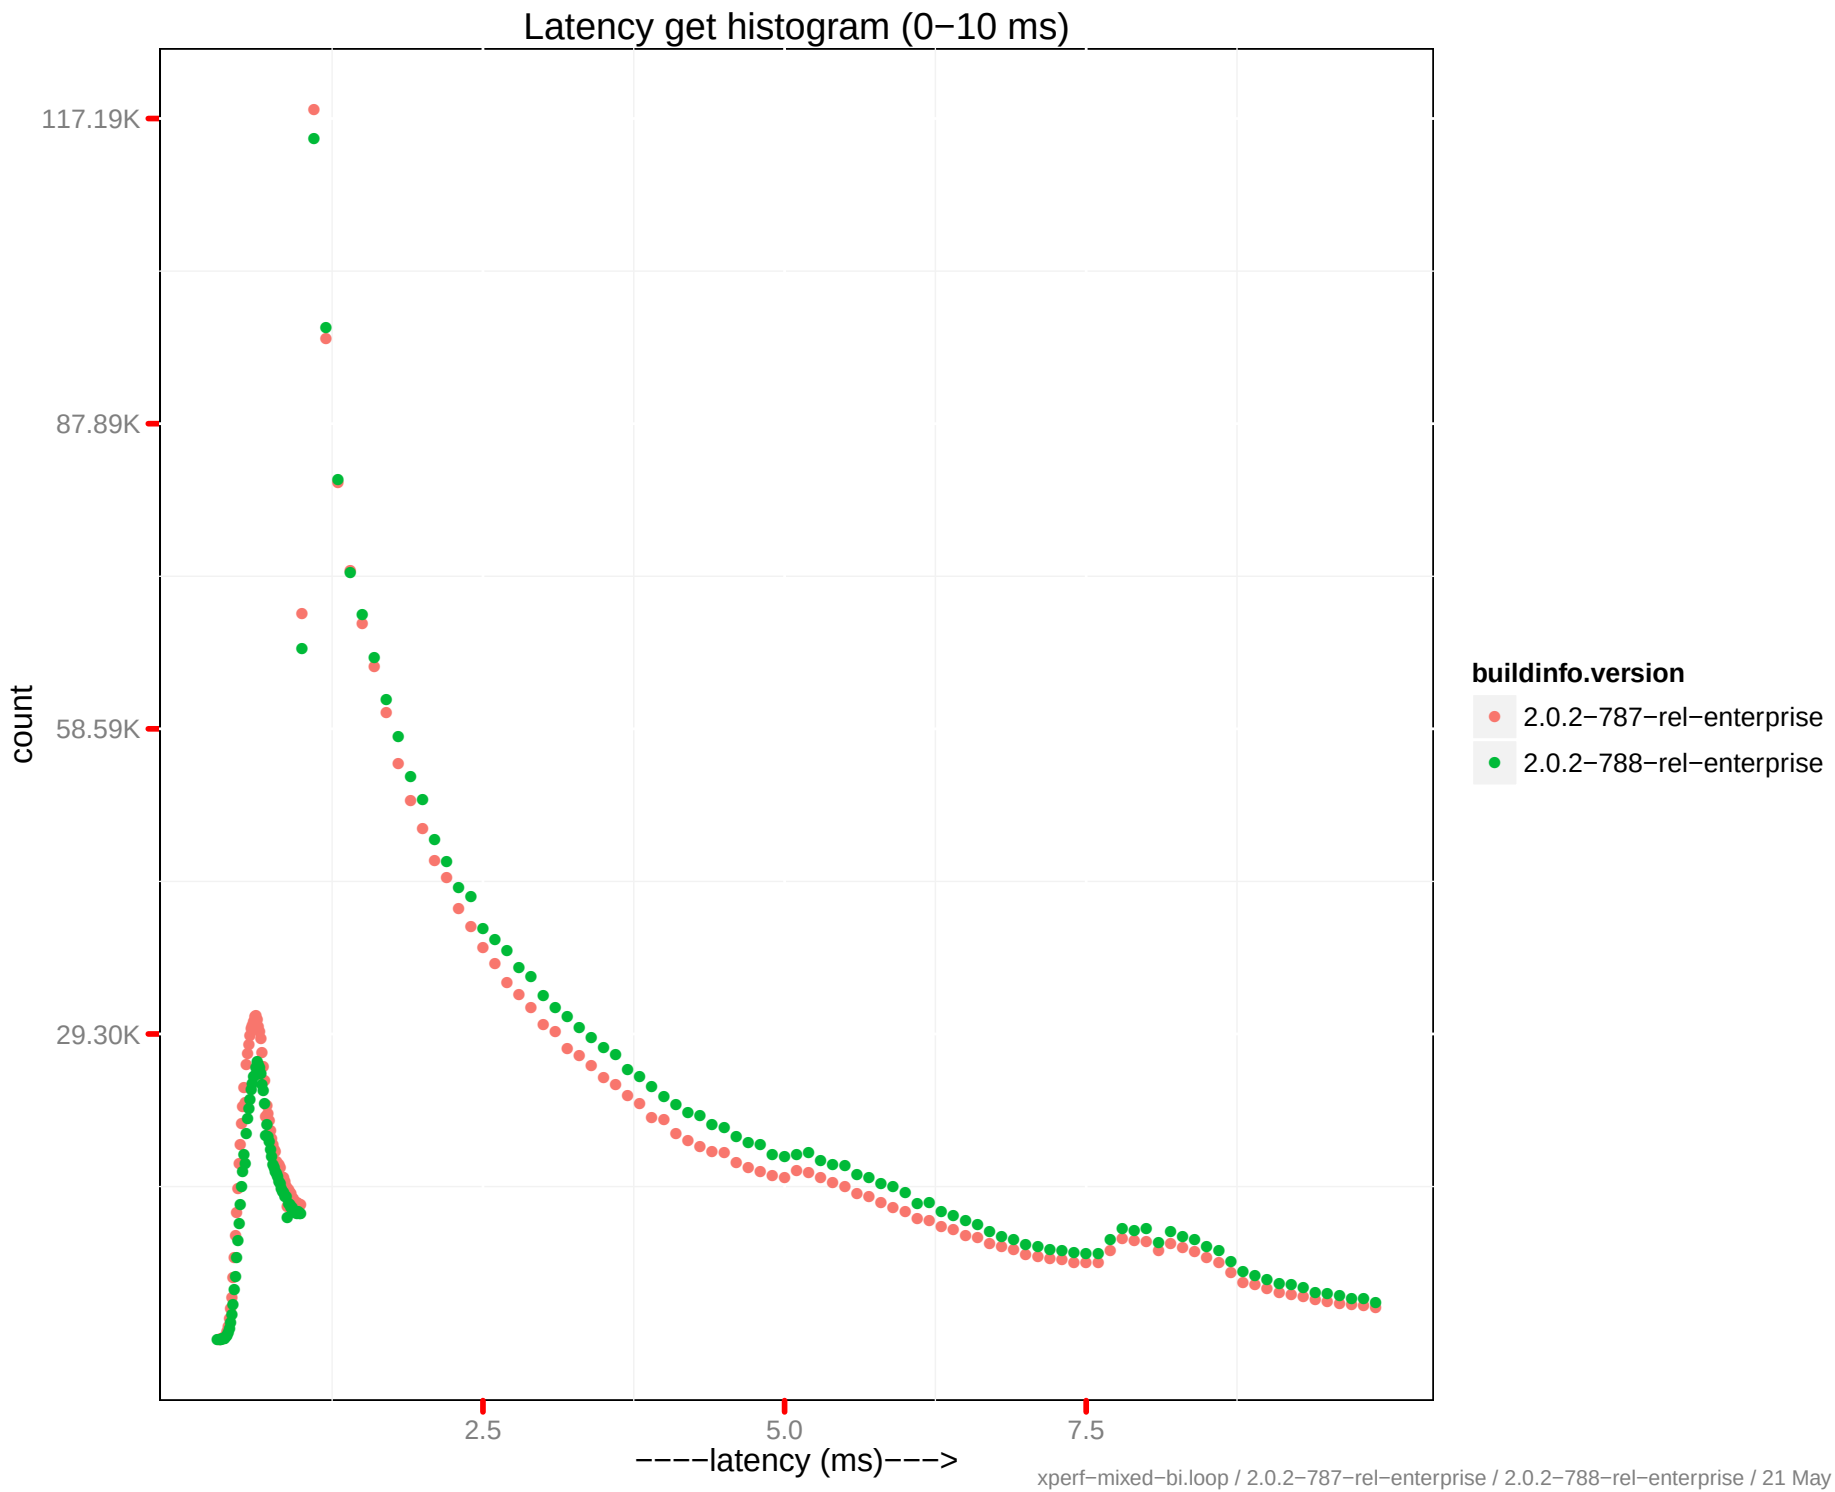
Latency get histogram (0−10 ms)
117.19K
87.89K
58.59K
29.30K
count
2.5
5.0
7.5
−−−−latency (ms)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

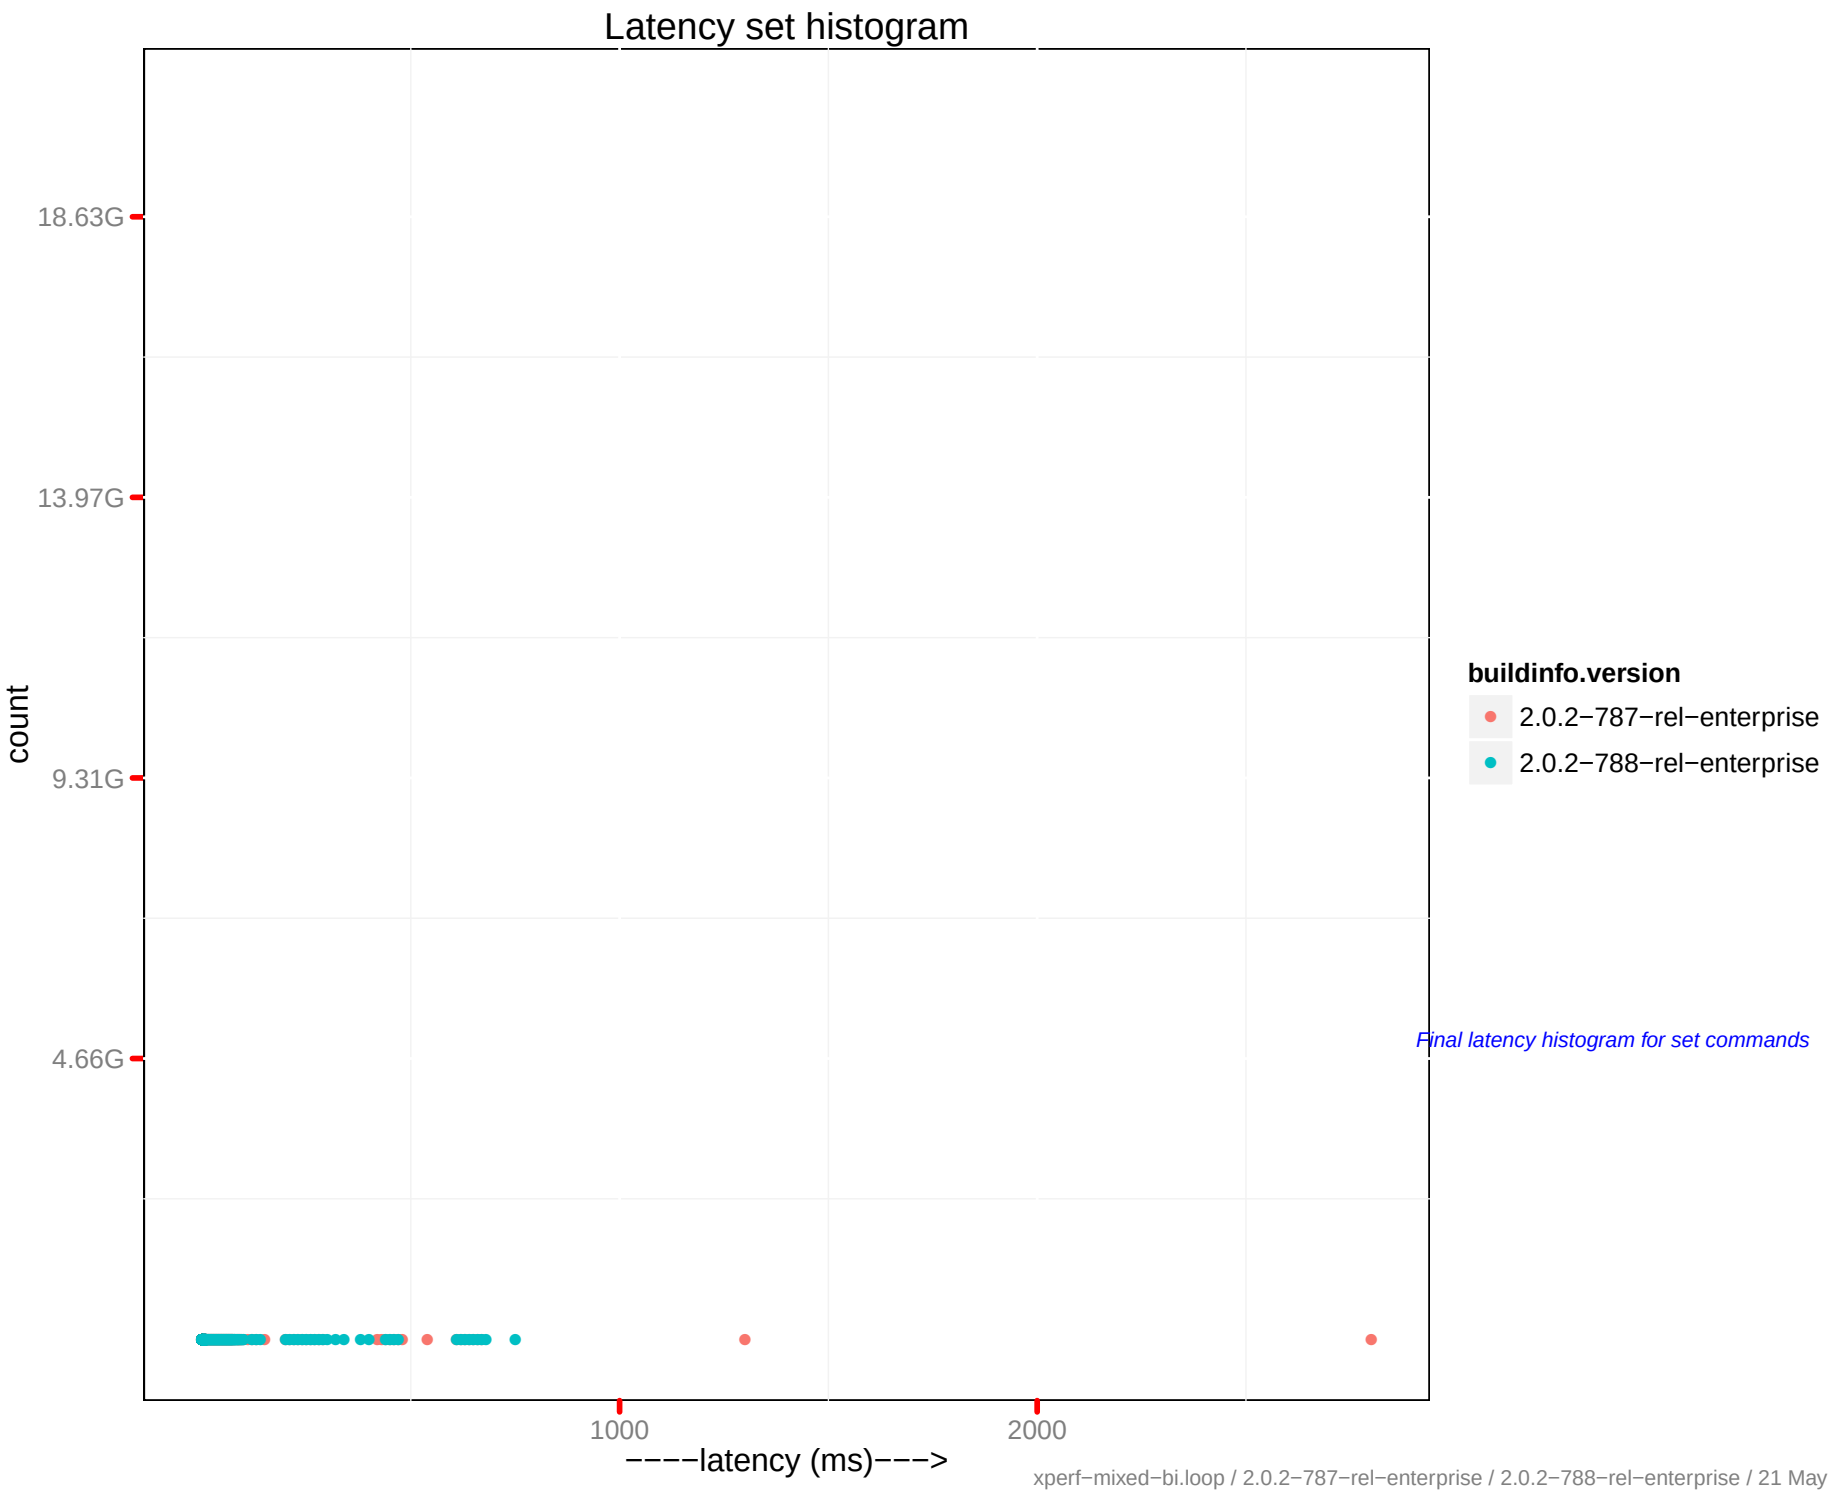
Latency set histogram
18.63G
13.97G
9.31G
4.66G
count
1000
2000
----latency (ms)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
Final latency histogram for set commands
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013

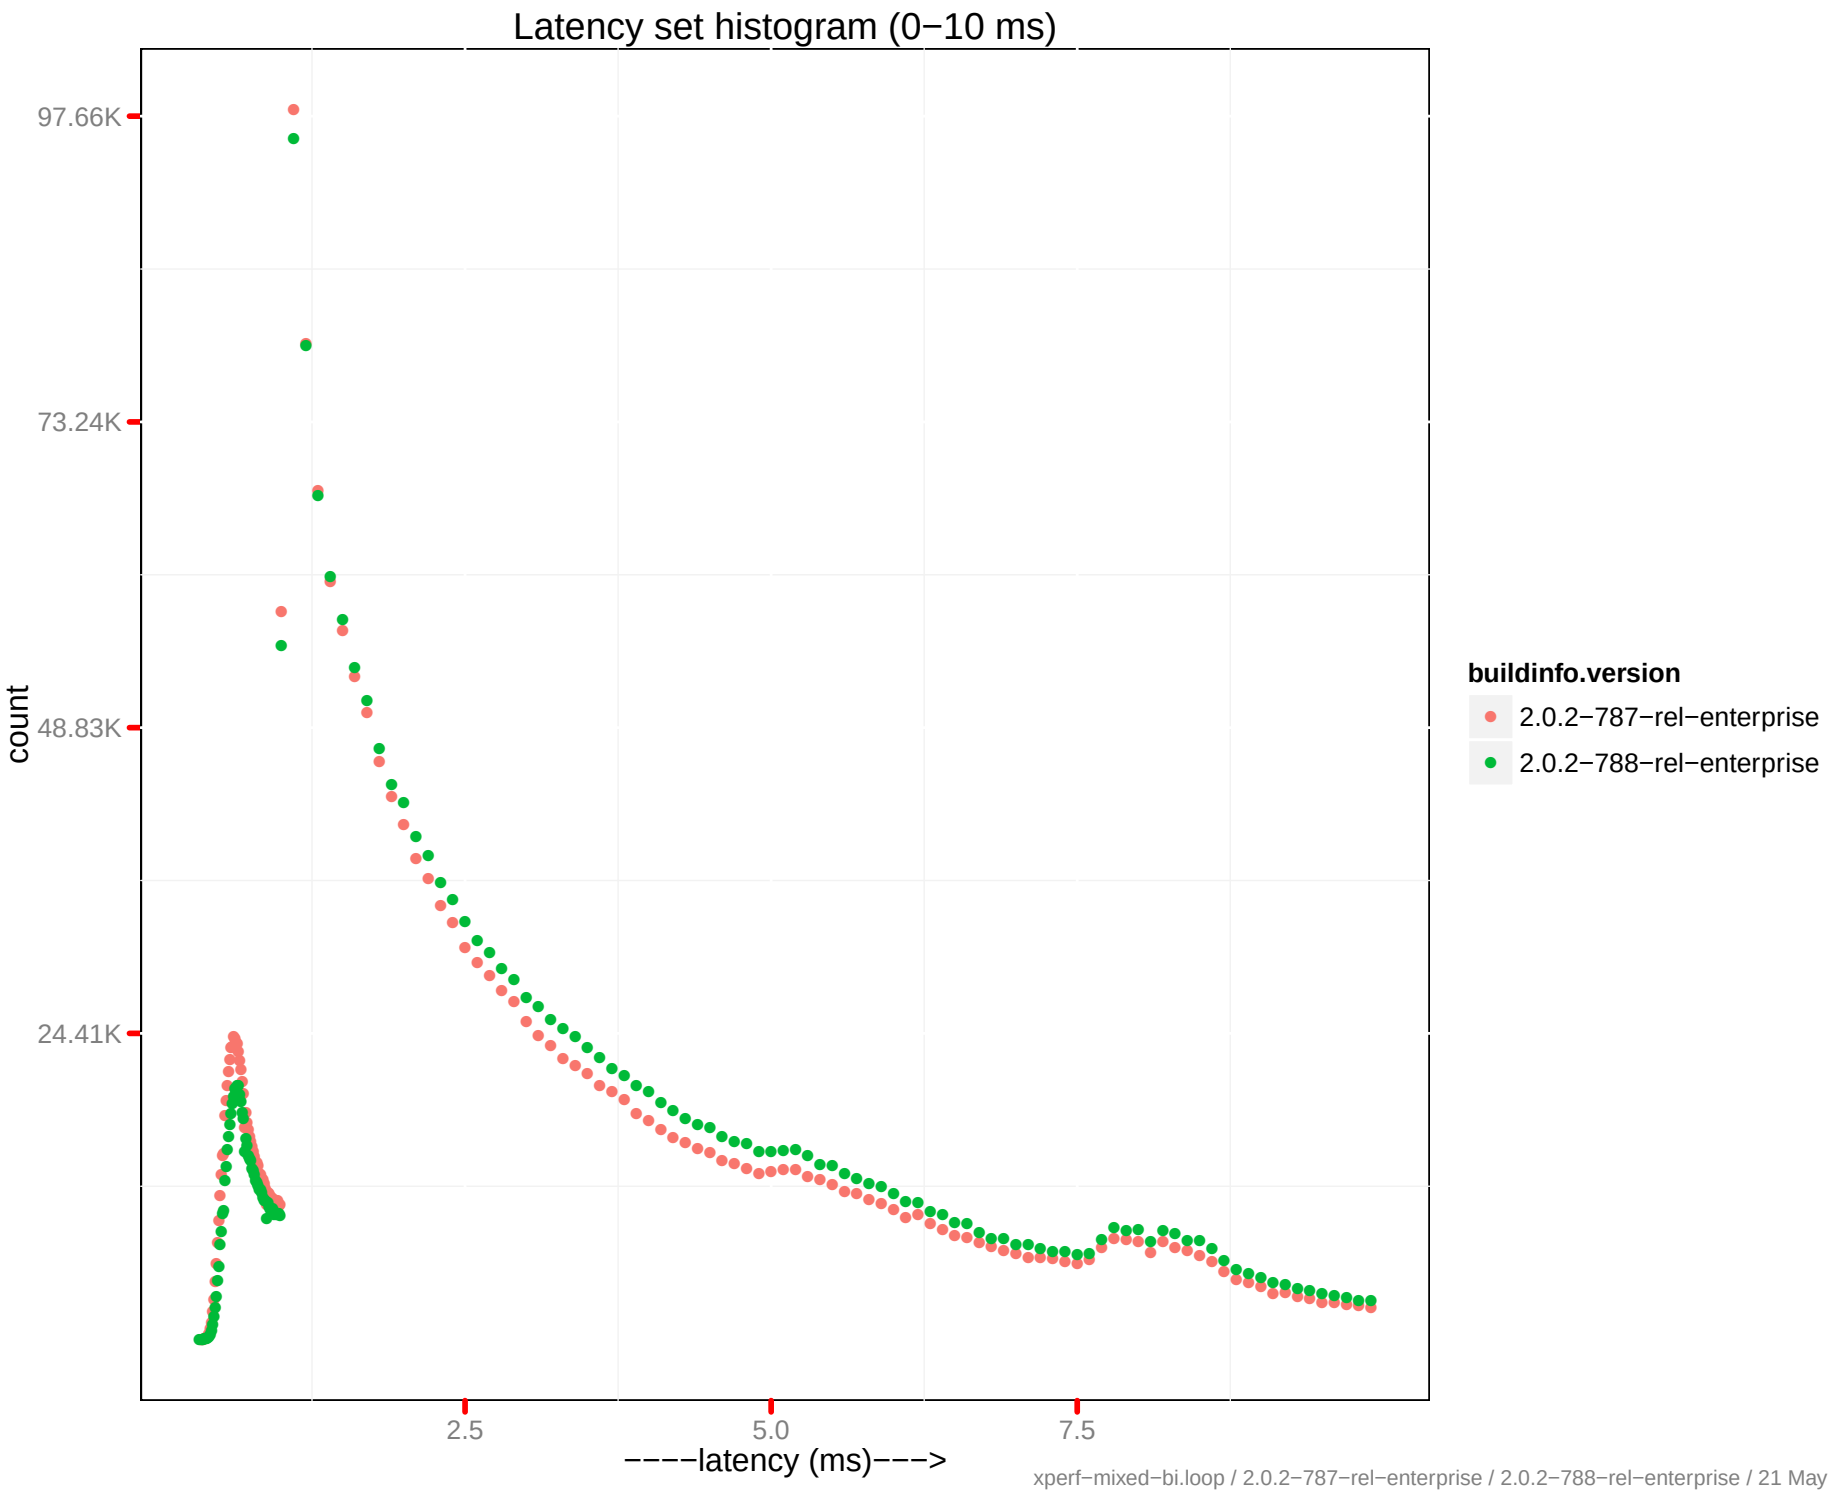
Latency set histogram (0−10 ms)
97.66K
73.24K
48.83K
24.41K
count
2.5
5.0
7.5
−−−−latency (ms)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

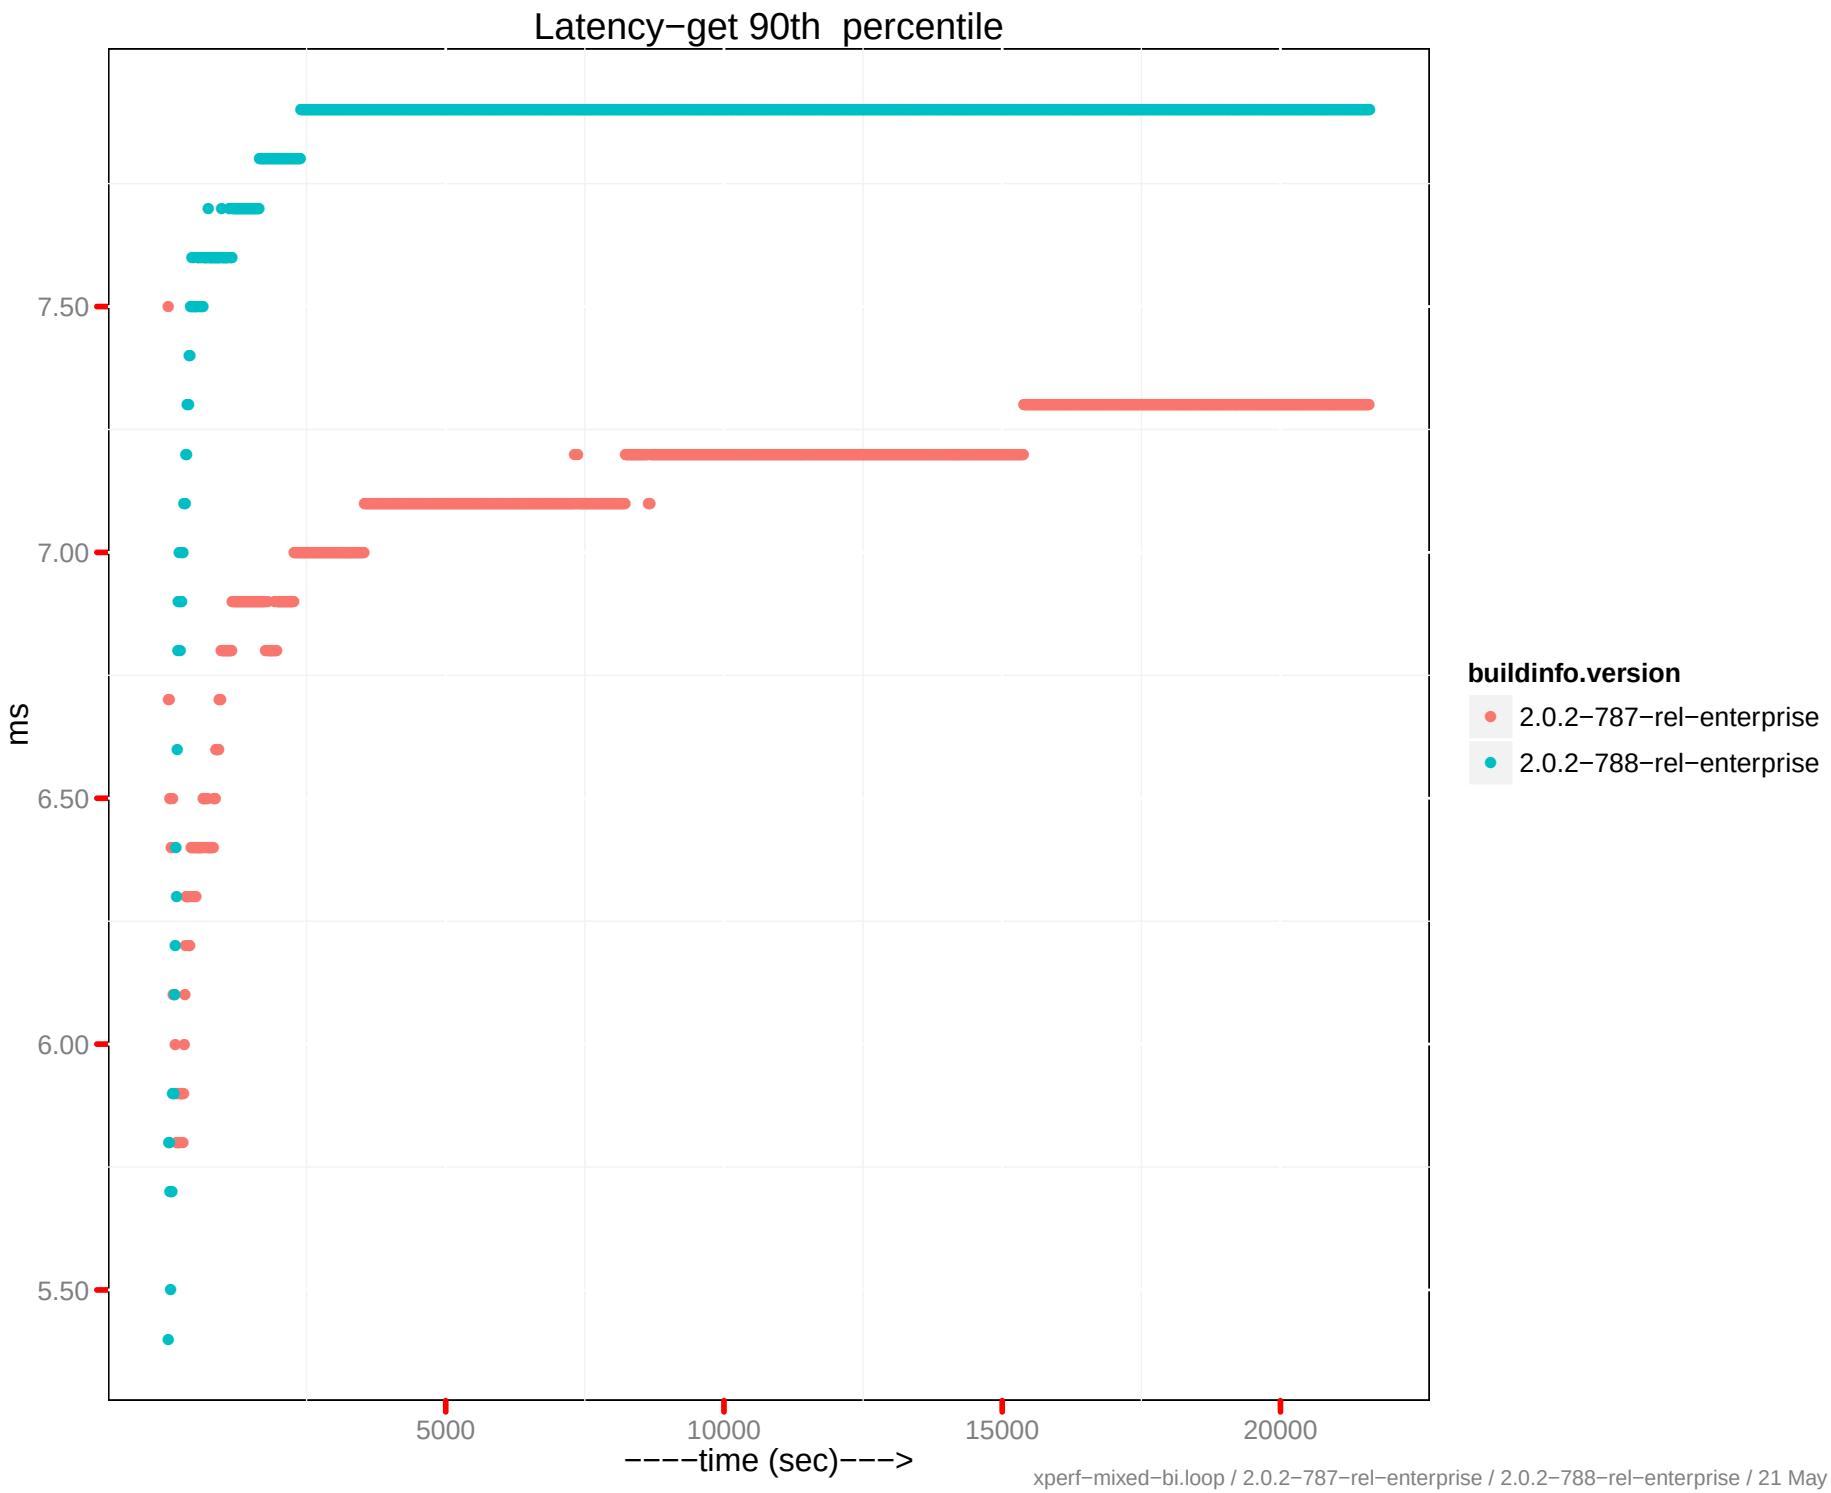
Latency−get 90th percentile
7.50
7.00
6.50
6.00
5.50
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

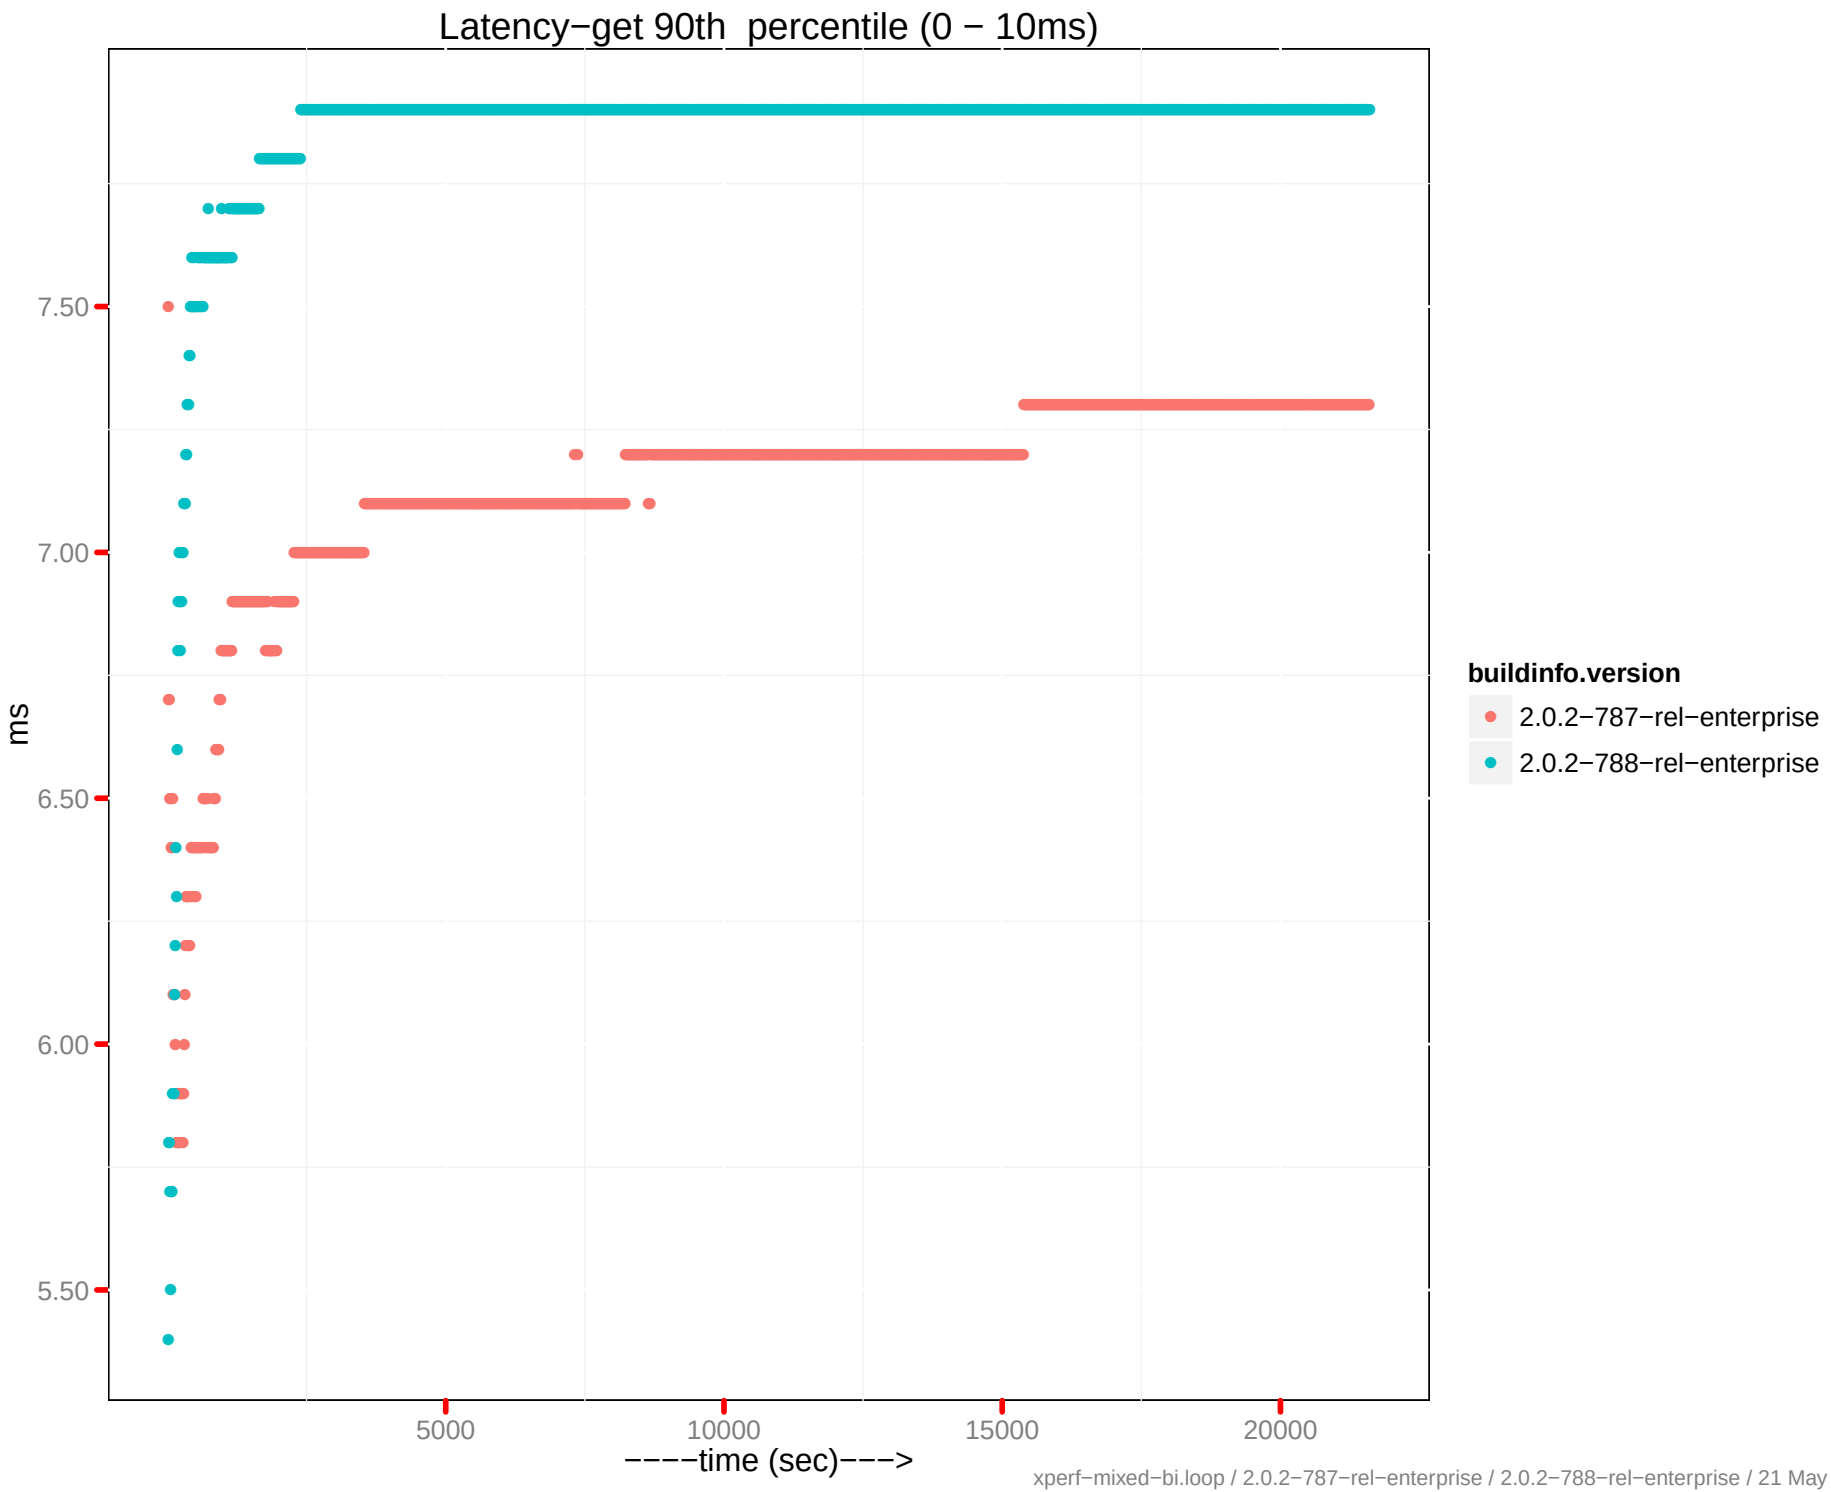
Latency−get 90th percentile (0 − 10ms)
7.50
7.00
6.50
6.00
5.50
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

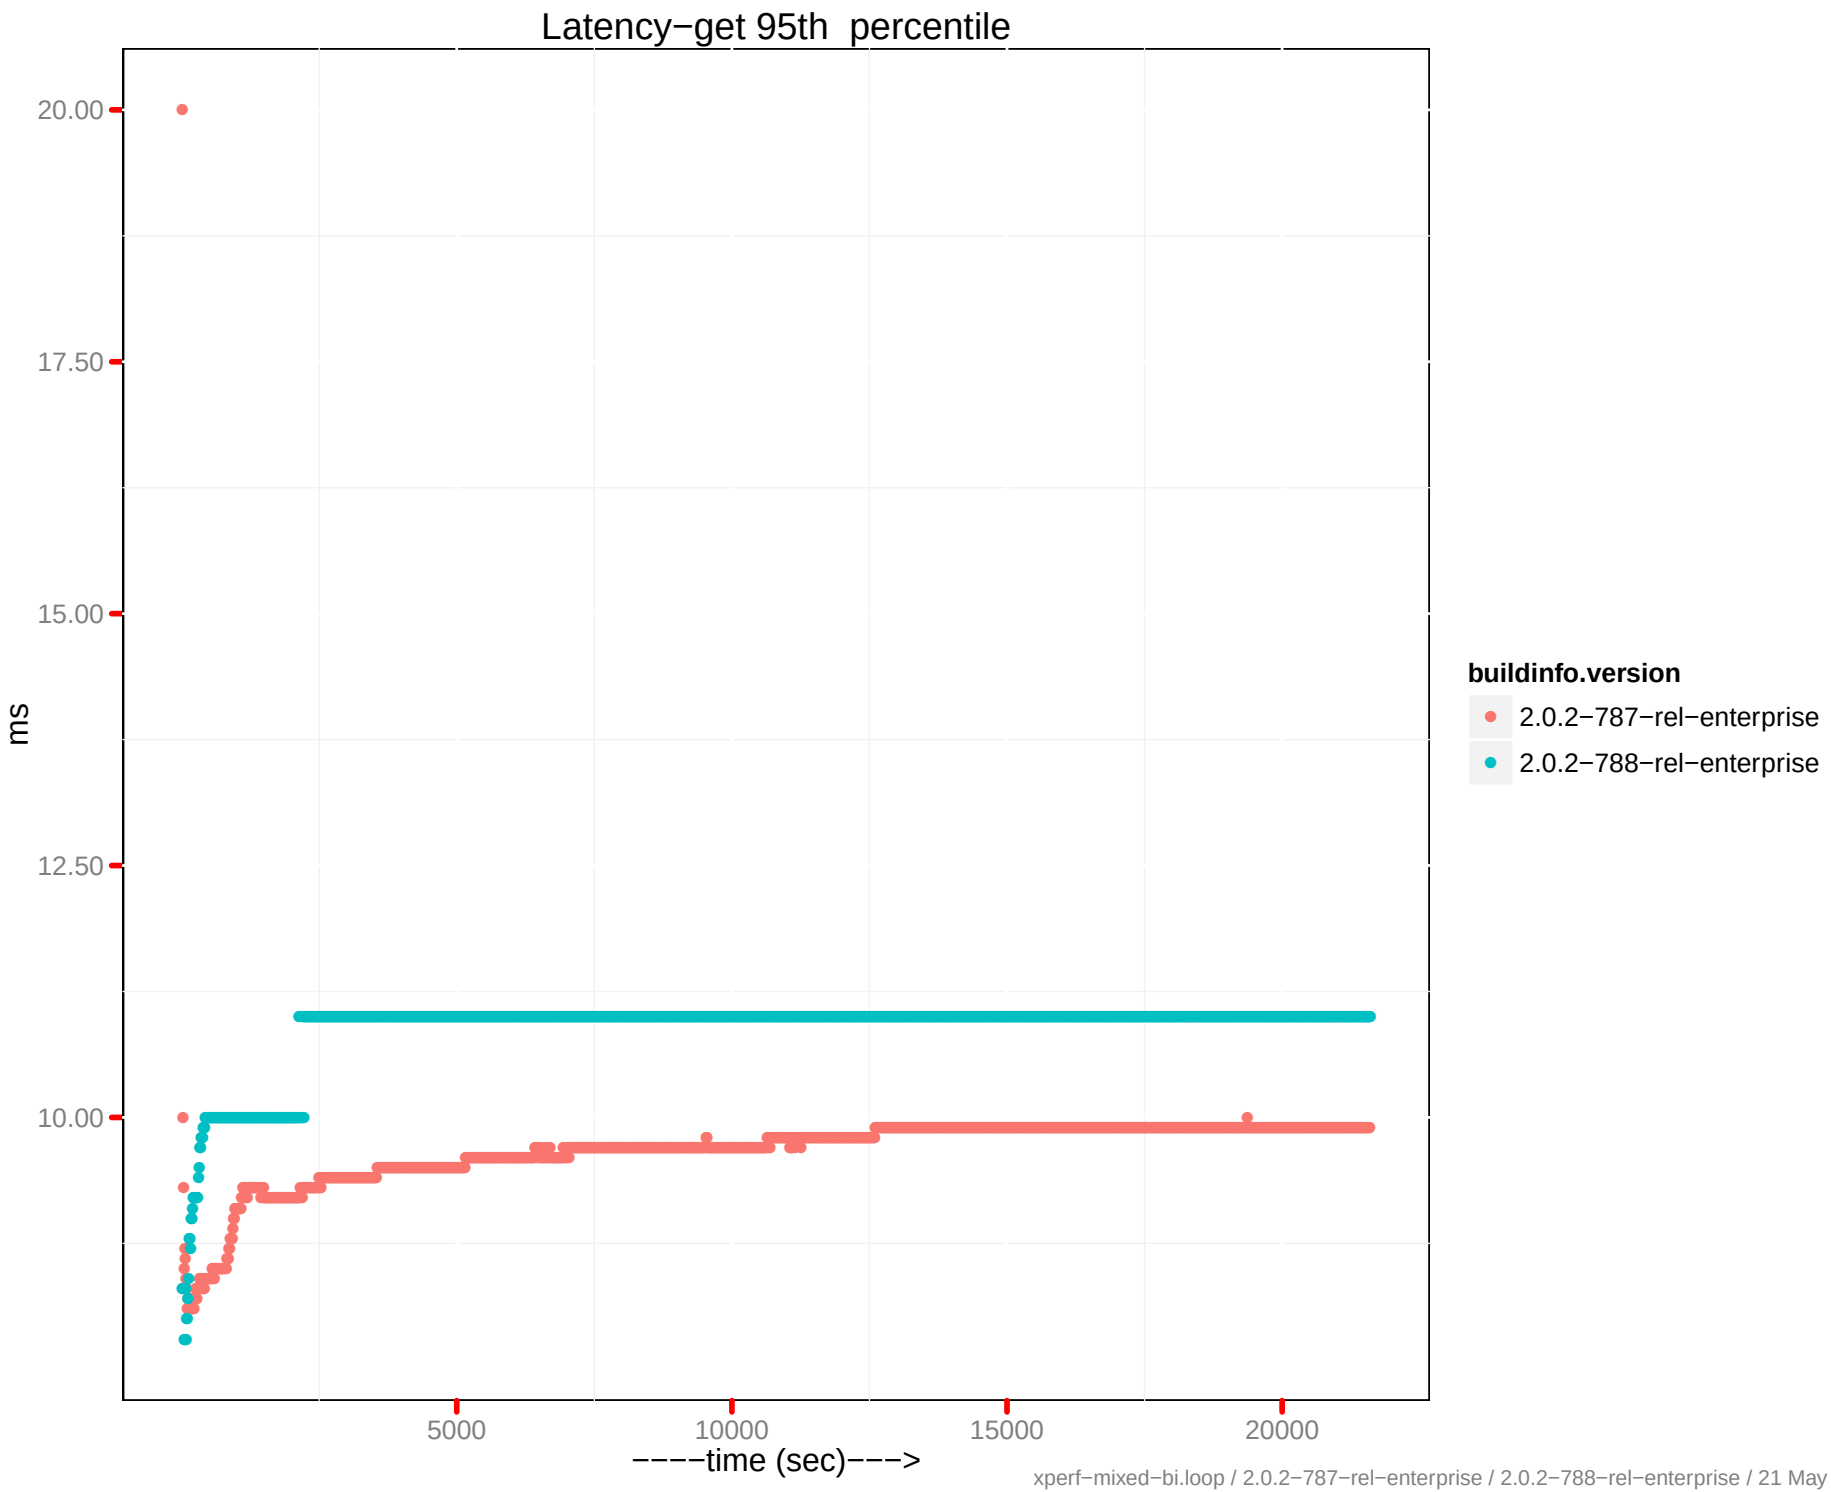
Latency−get 95th percentile
20.00
17.50
15.00
12.50
10.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

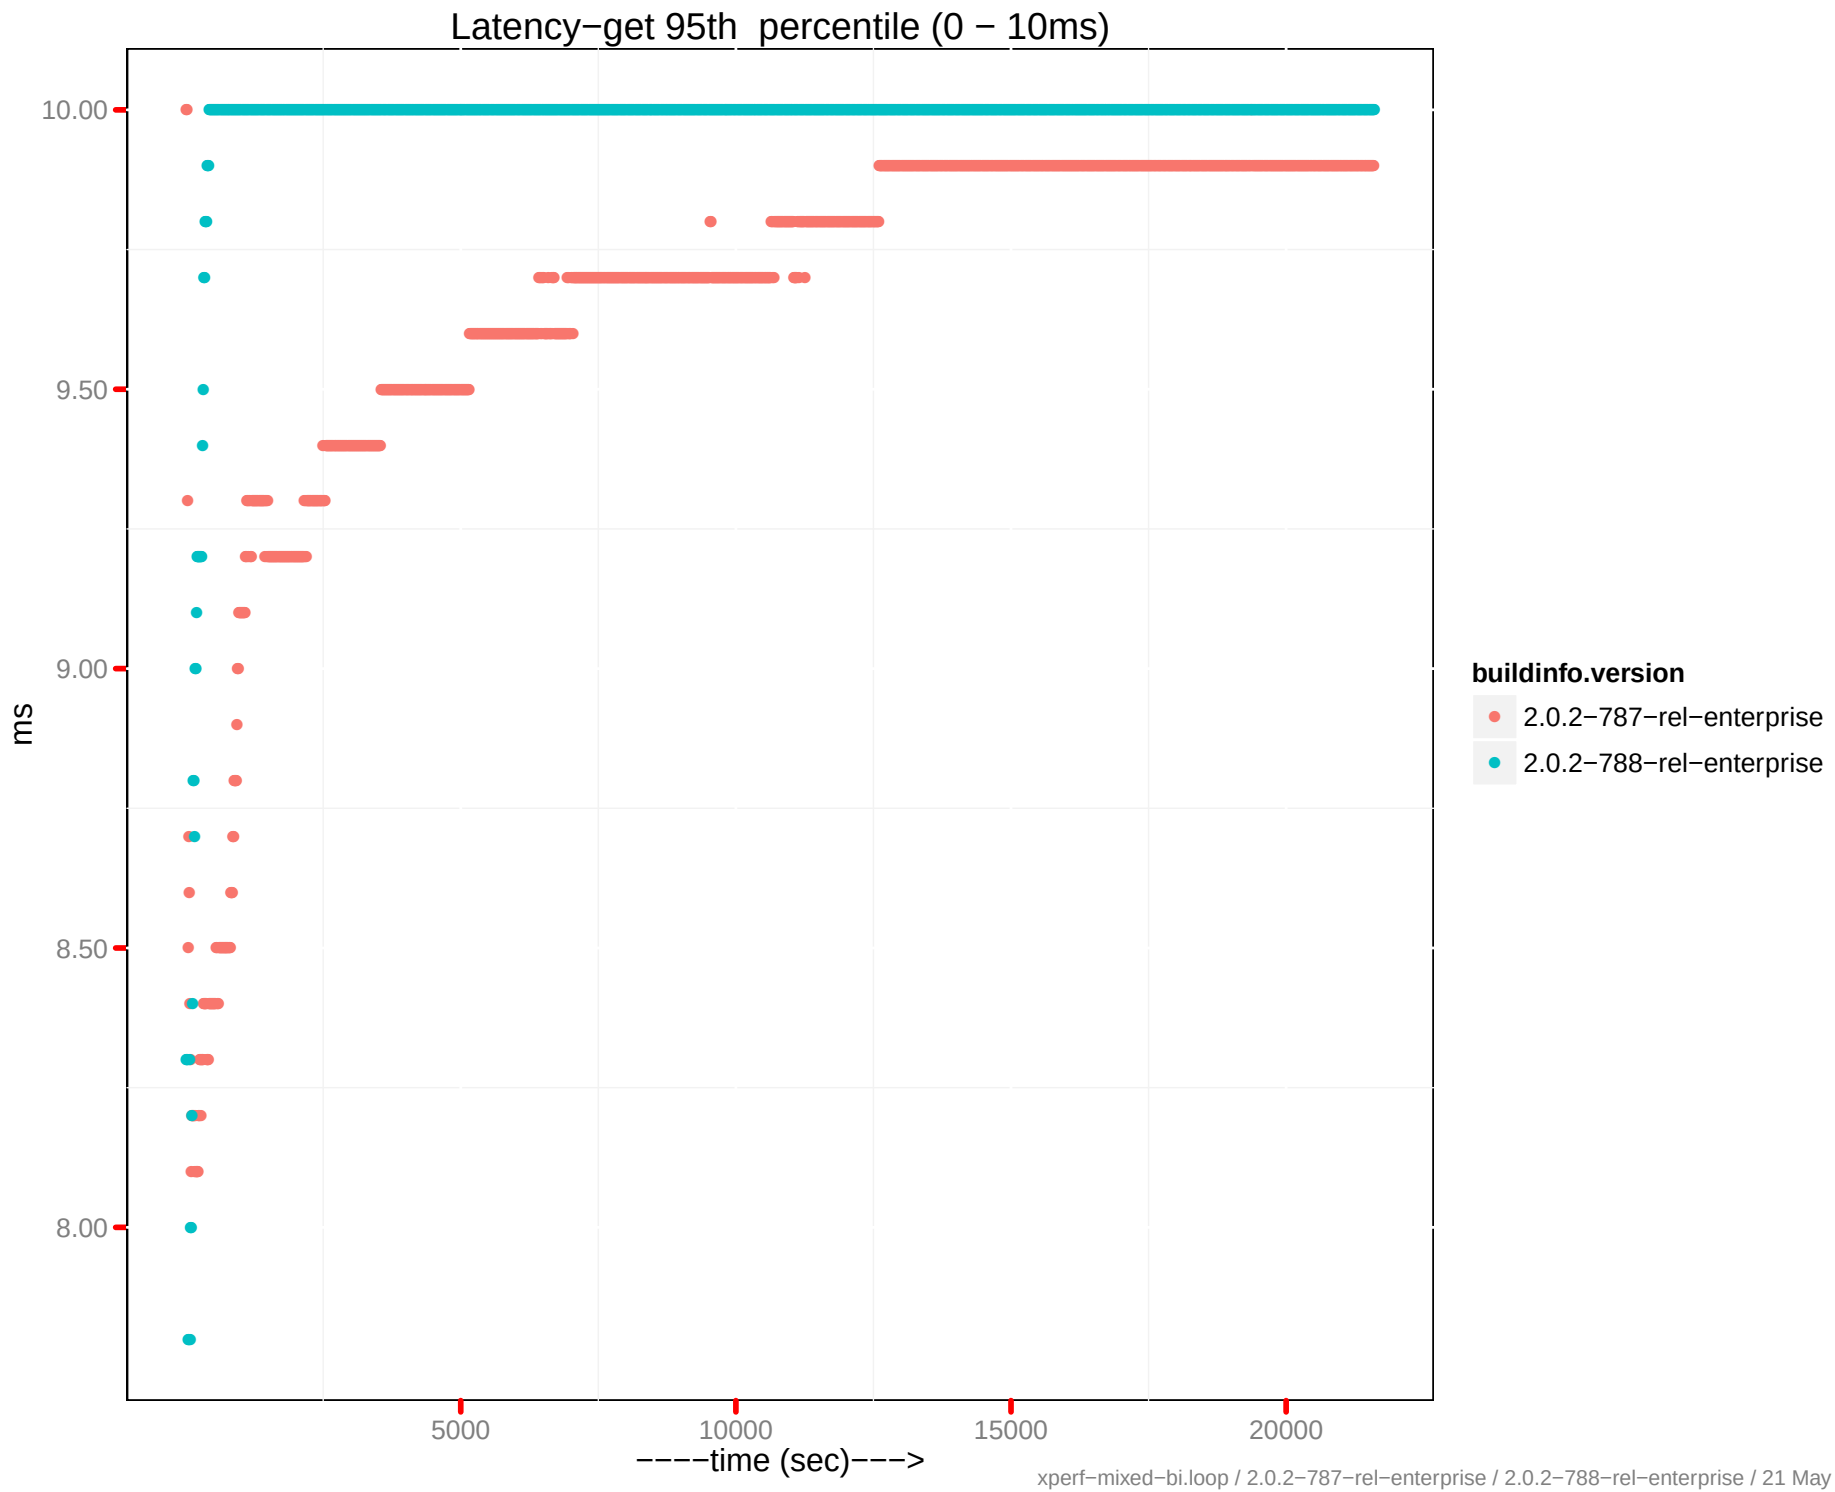
Latency−get 95th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
8.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

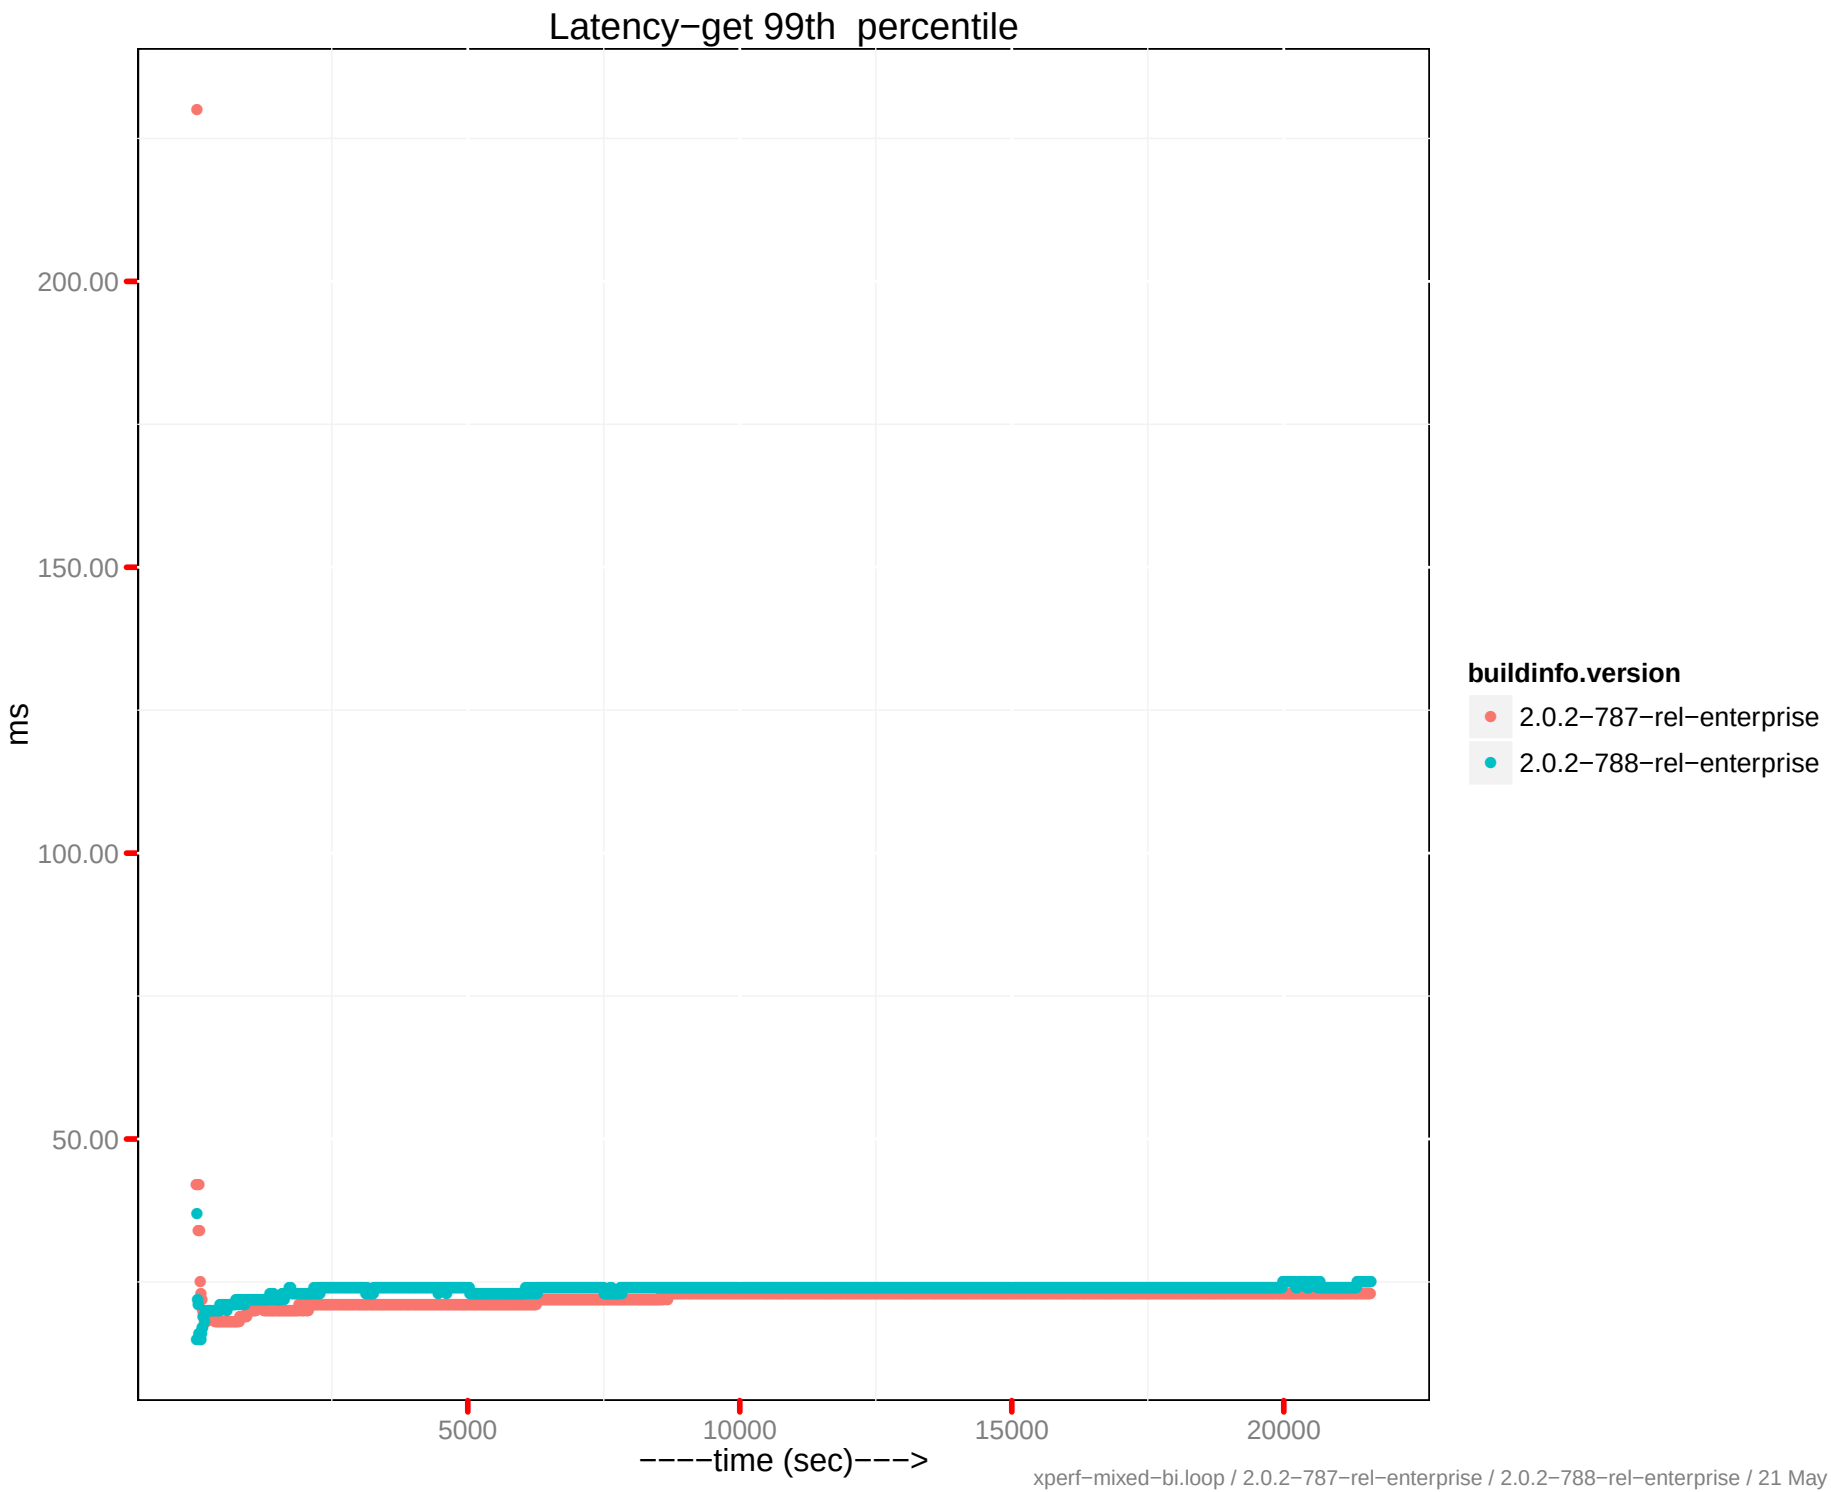
Latency−get 99th percentile
200.00
150.00
100.00
50.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

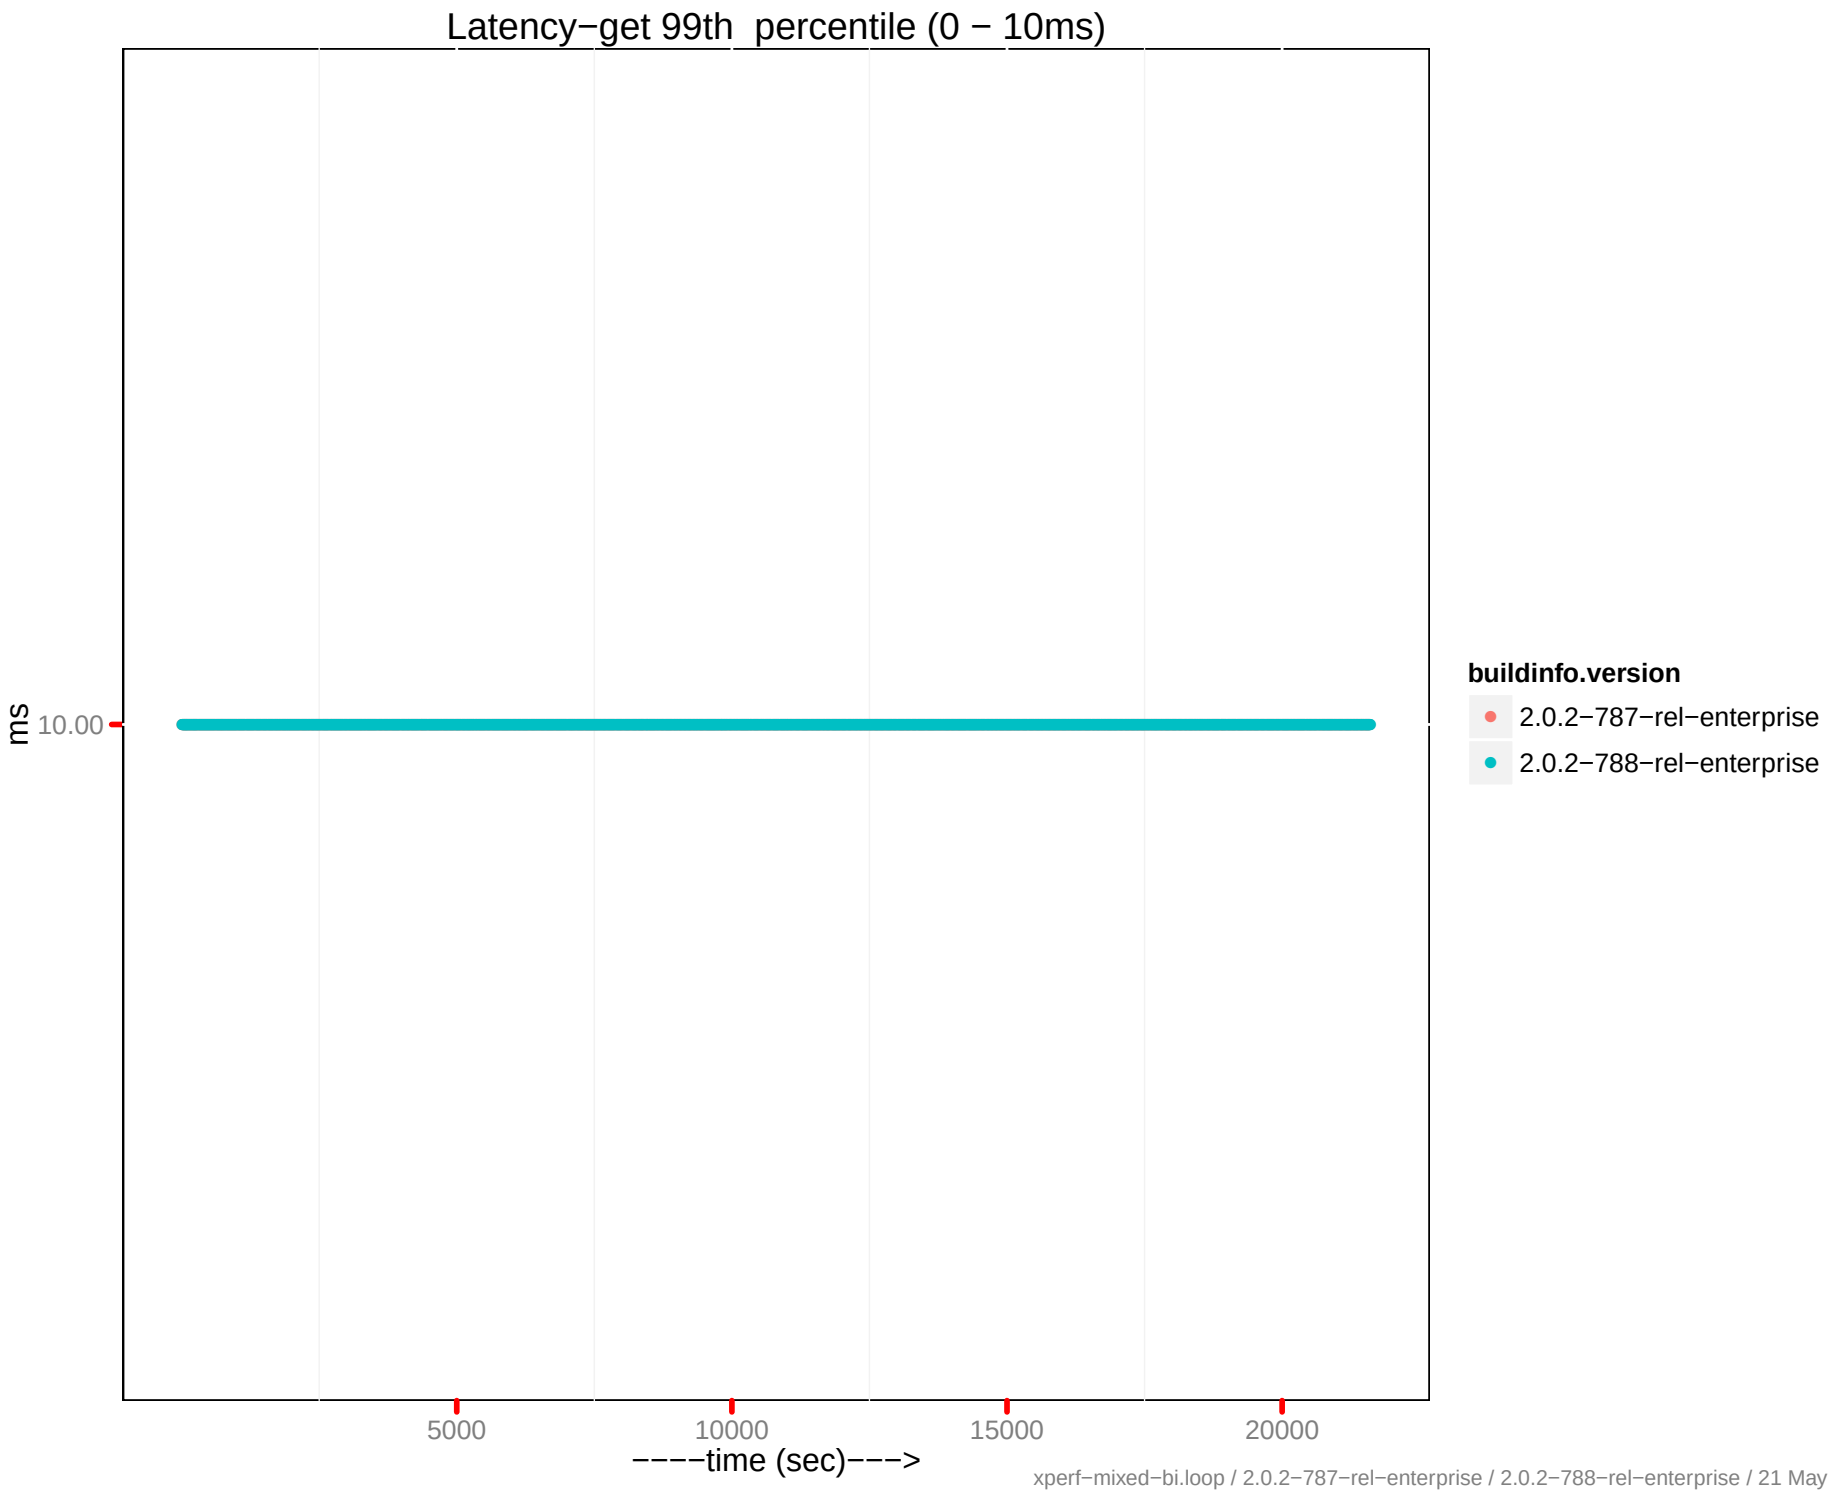
Latency−get 99th percentile (0 − 10ms)
ms
10.00
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

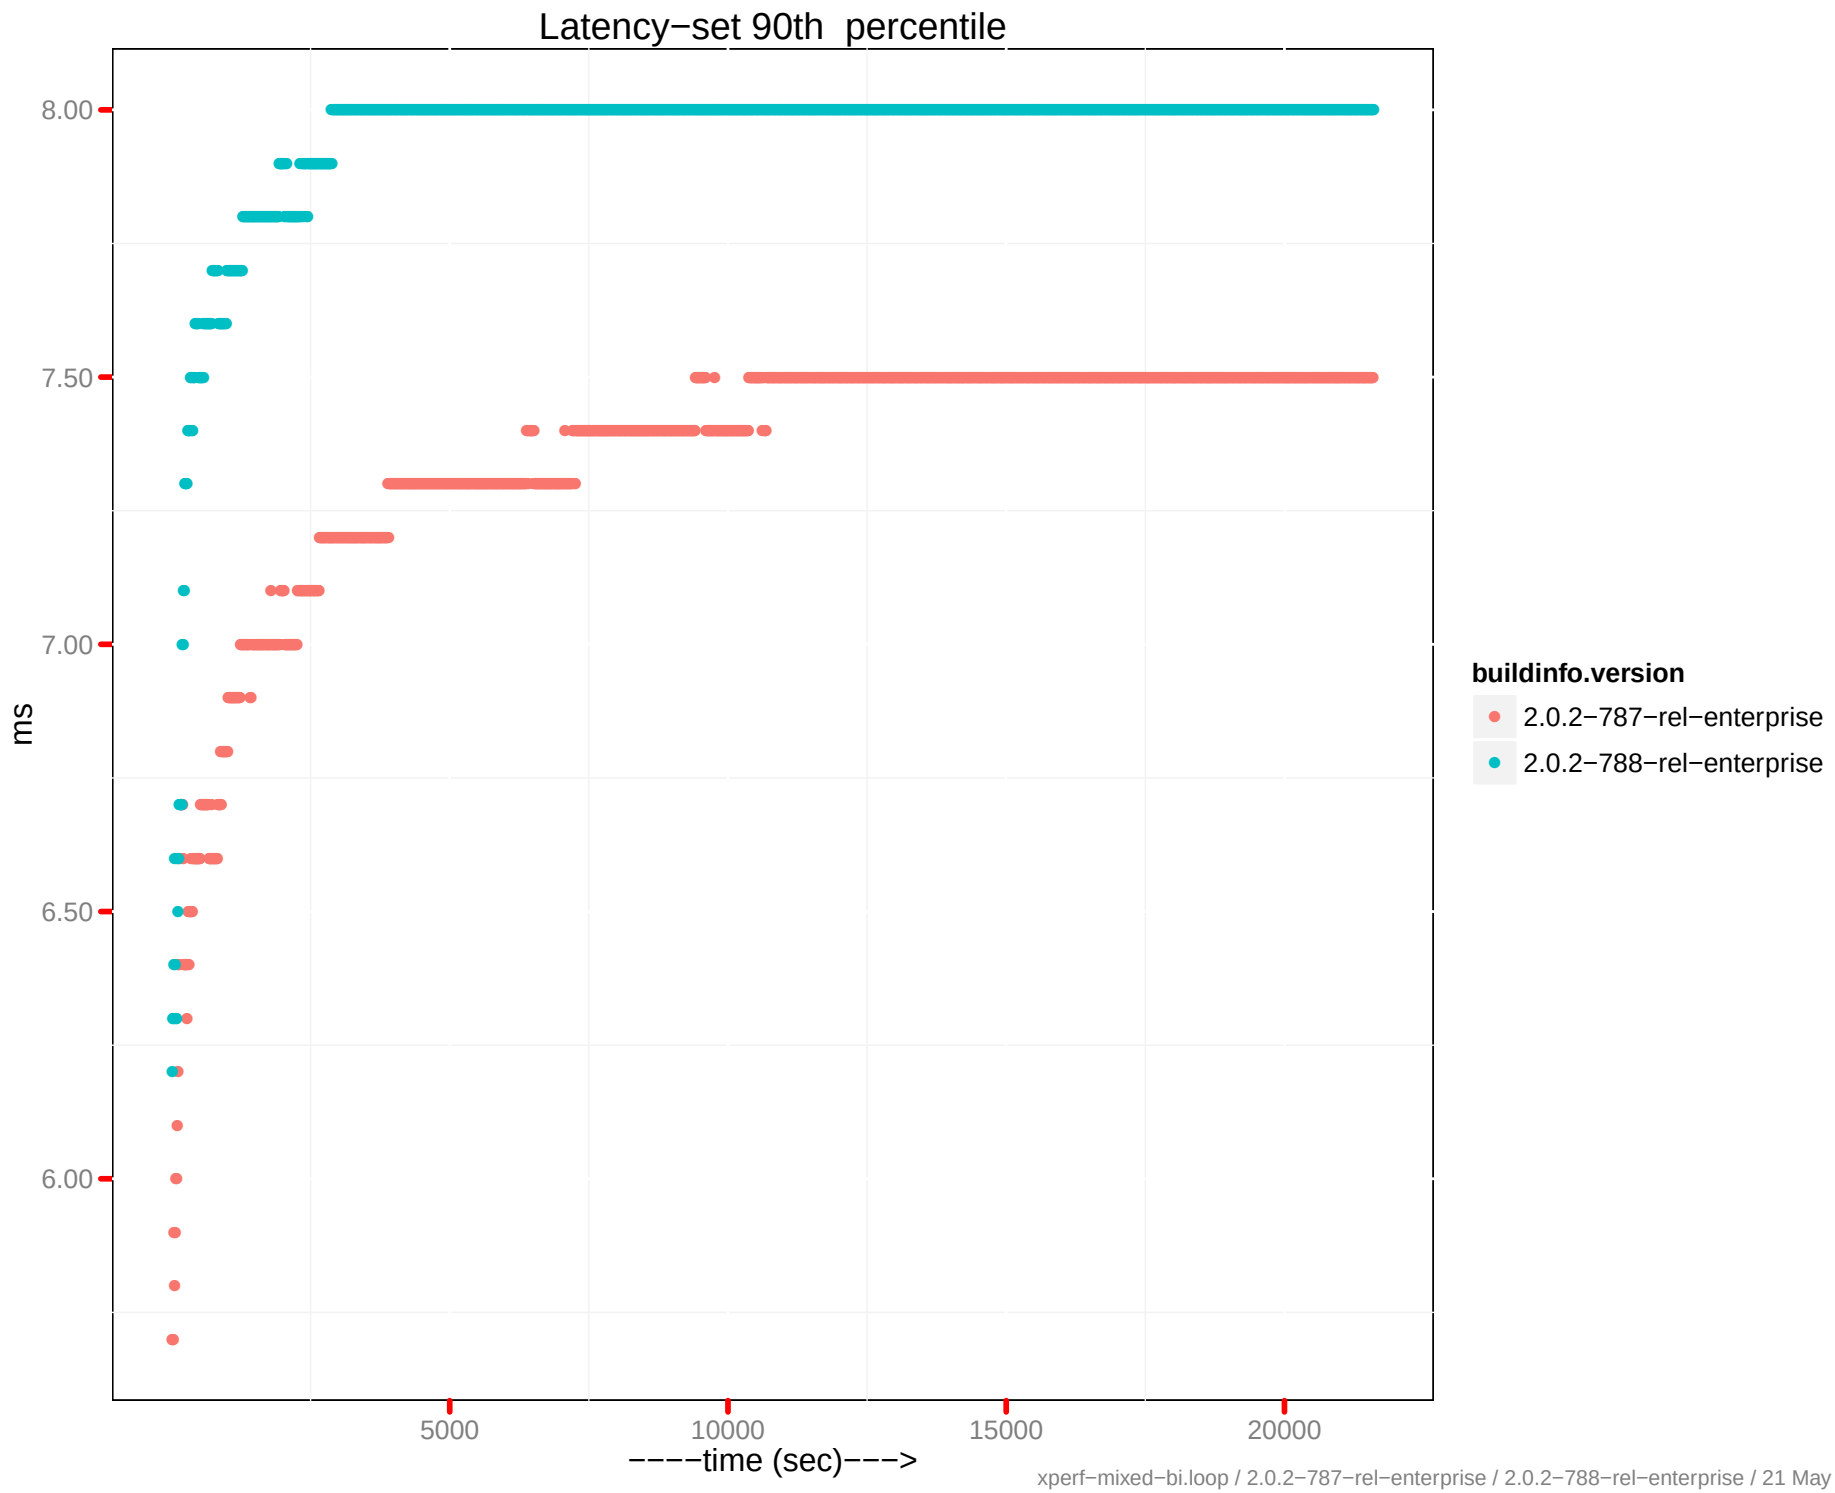
Latency−set 90th percentile
8.00
7.50
7.00
6.50
6.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

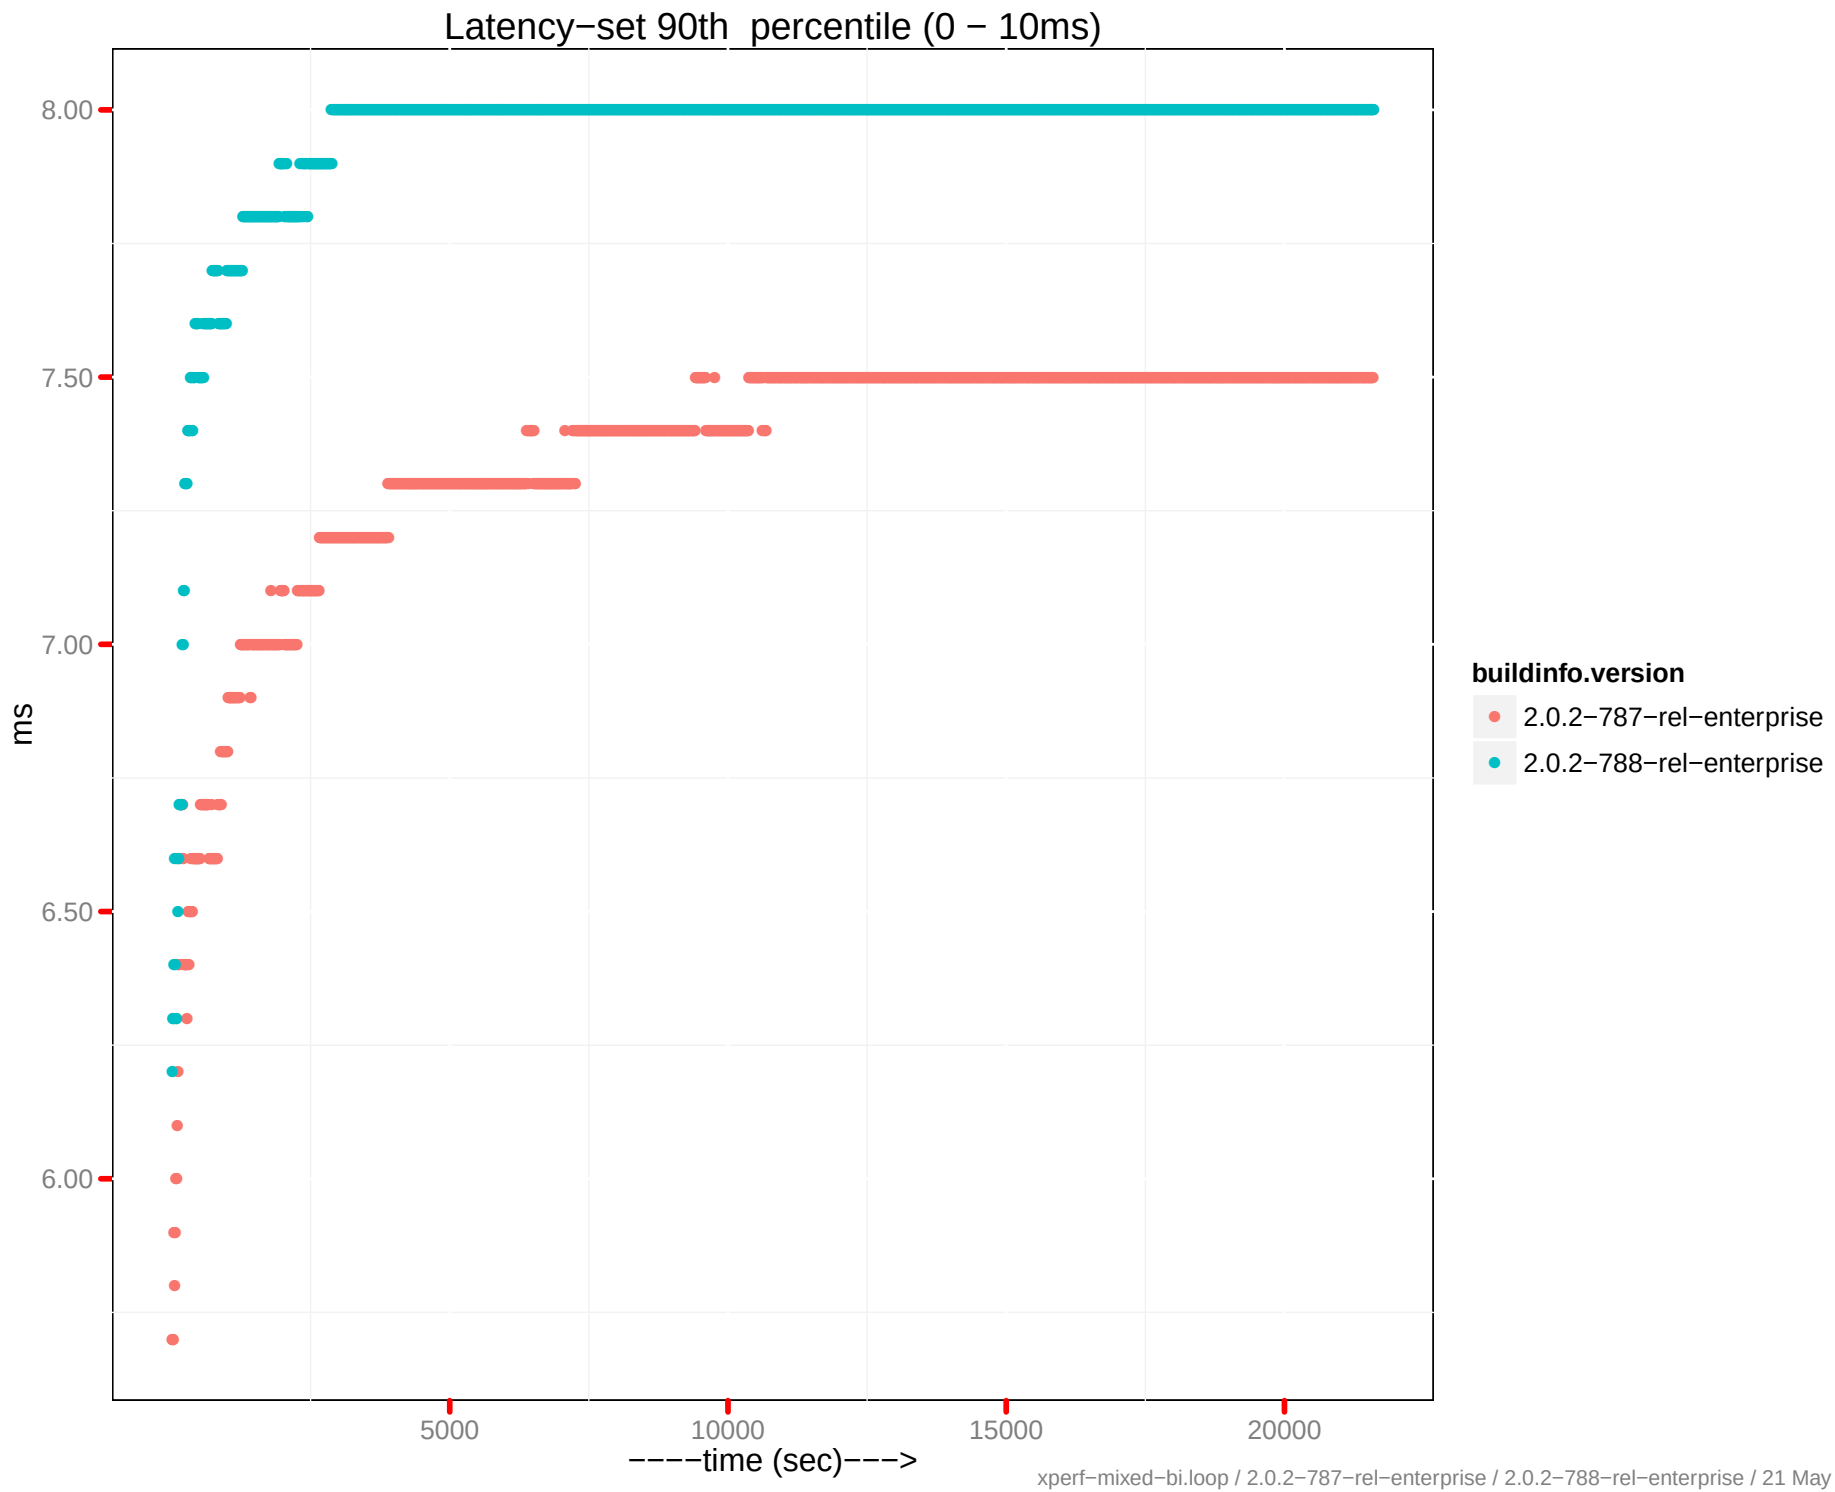
Latency−set 90th percentile (0 − 10ms)
8.00
7.50
7.00
6.50
6.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

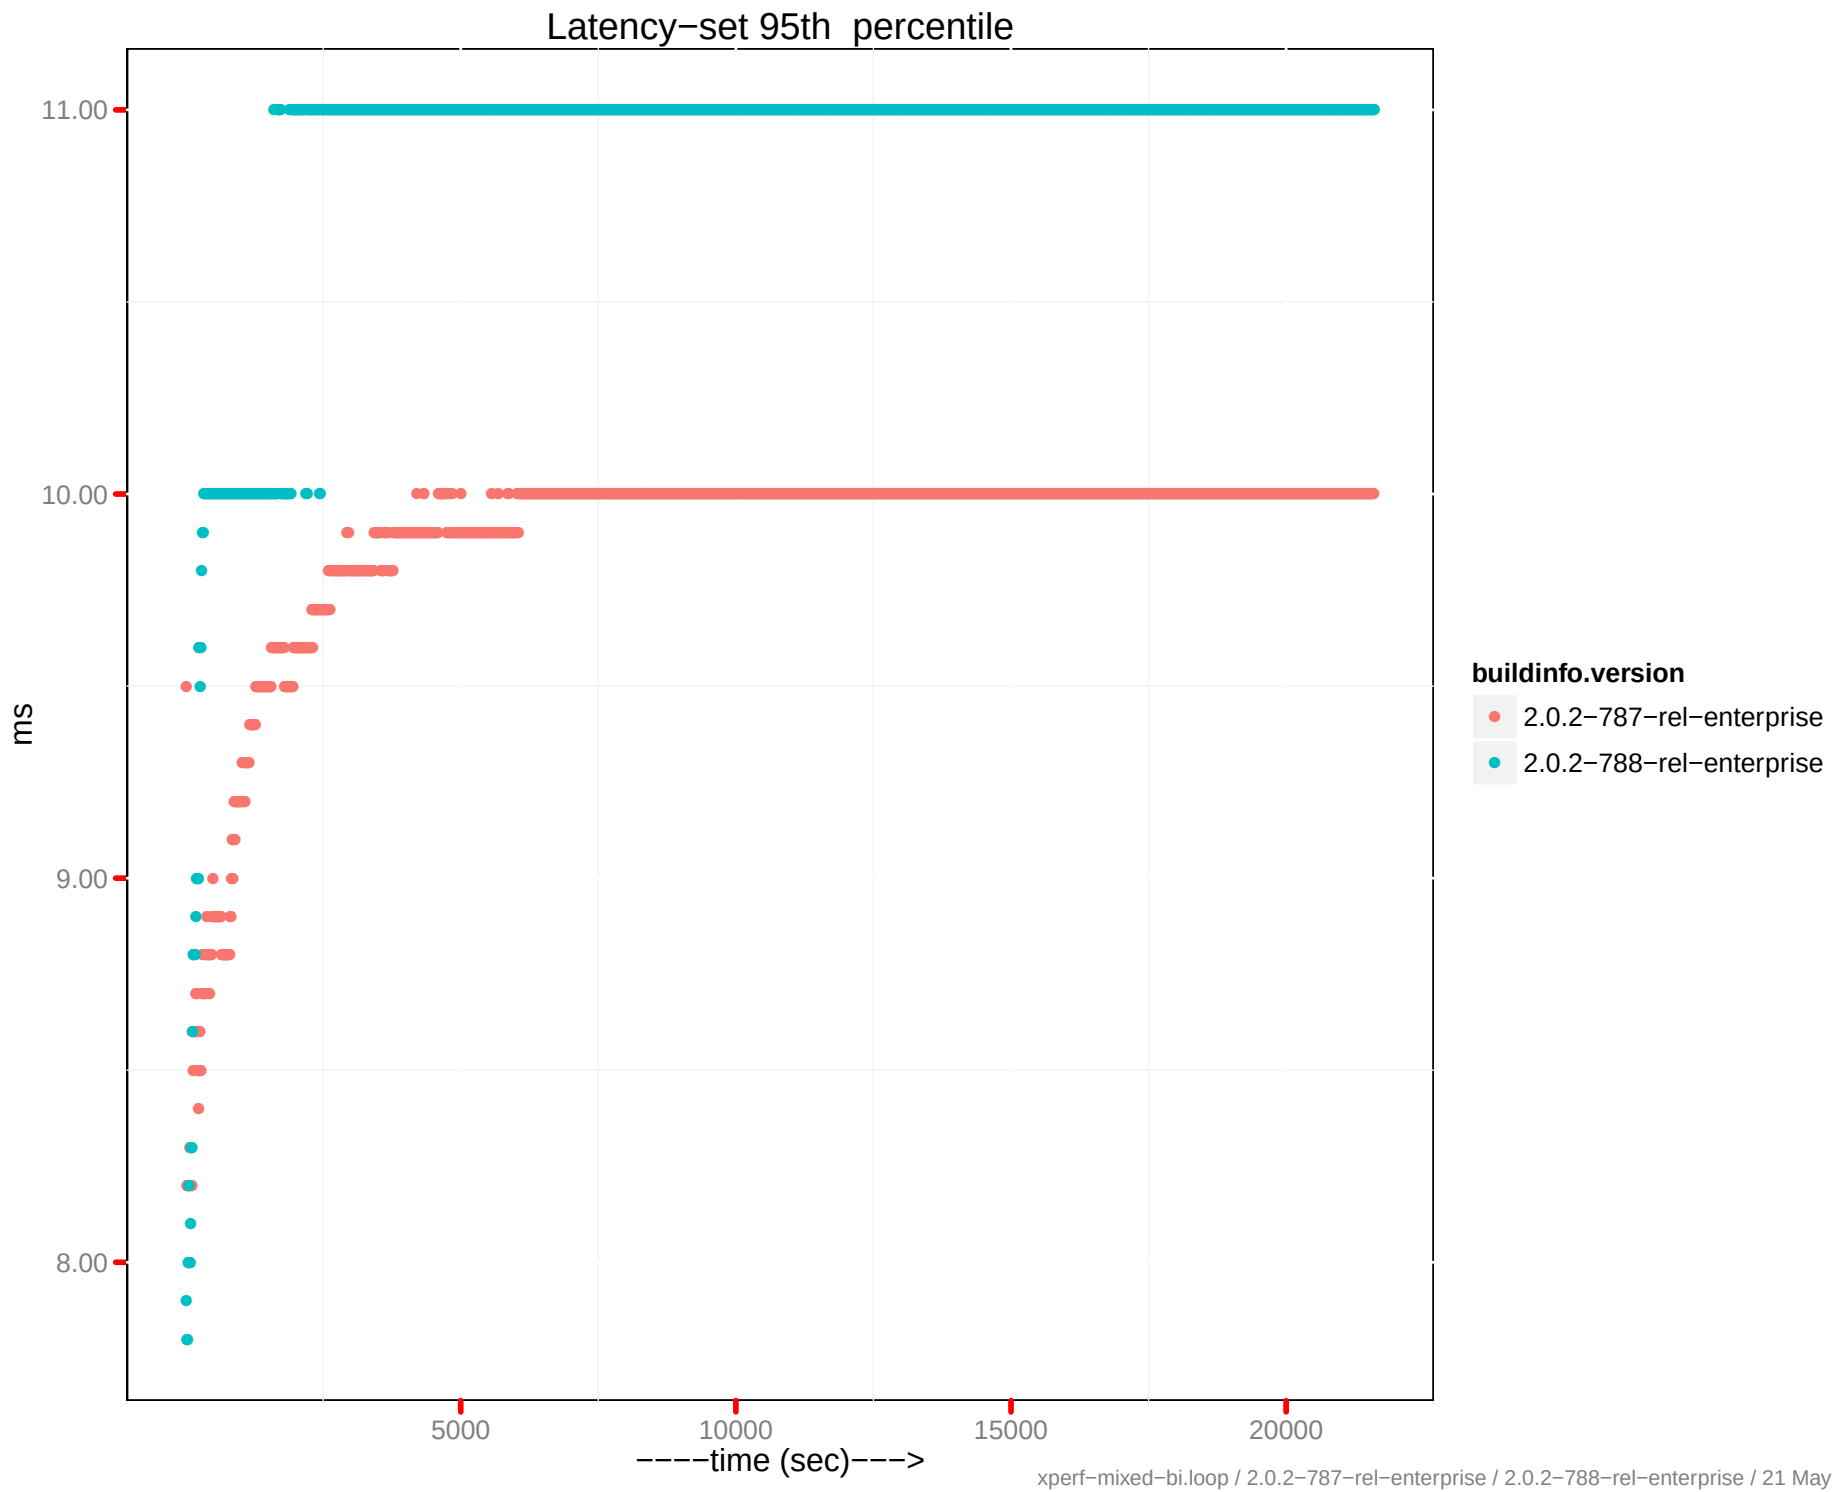
Latency−set 95th percentile
11.00
10.00
9.00
8.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

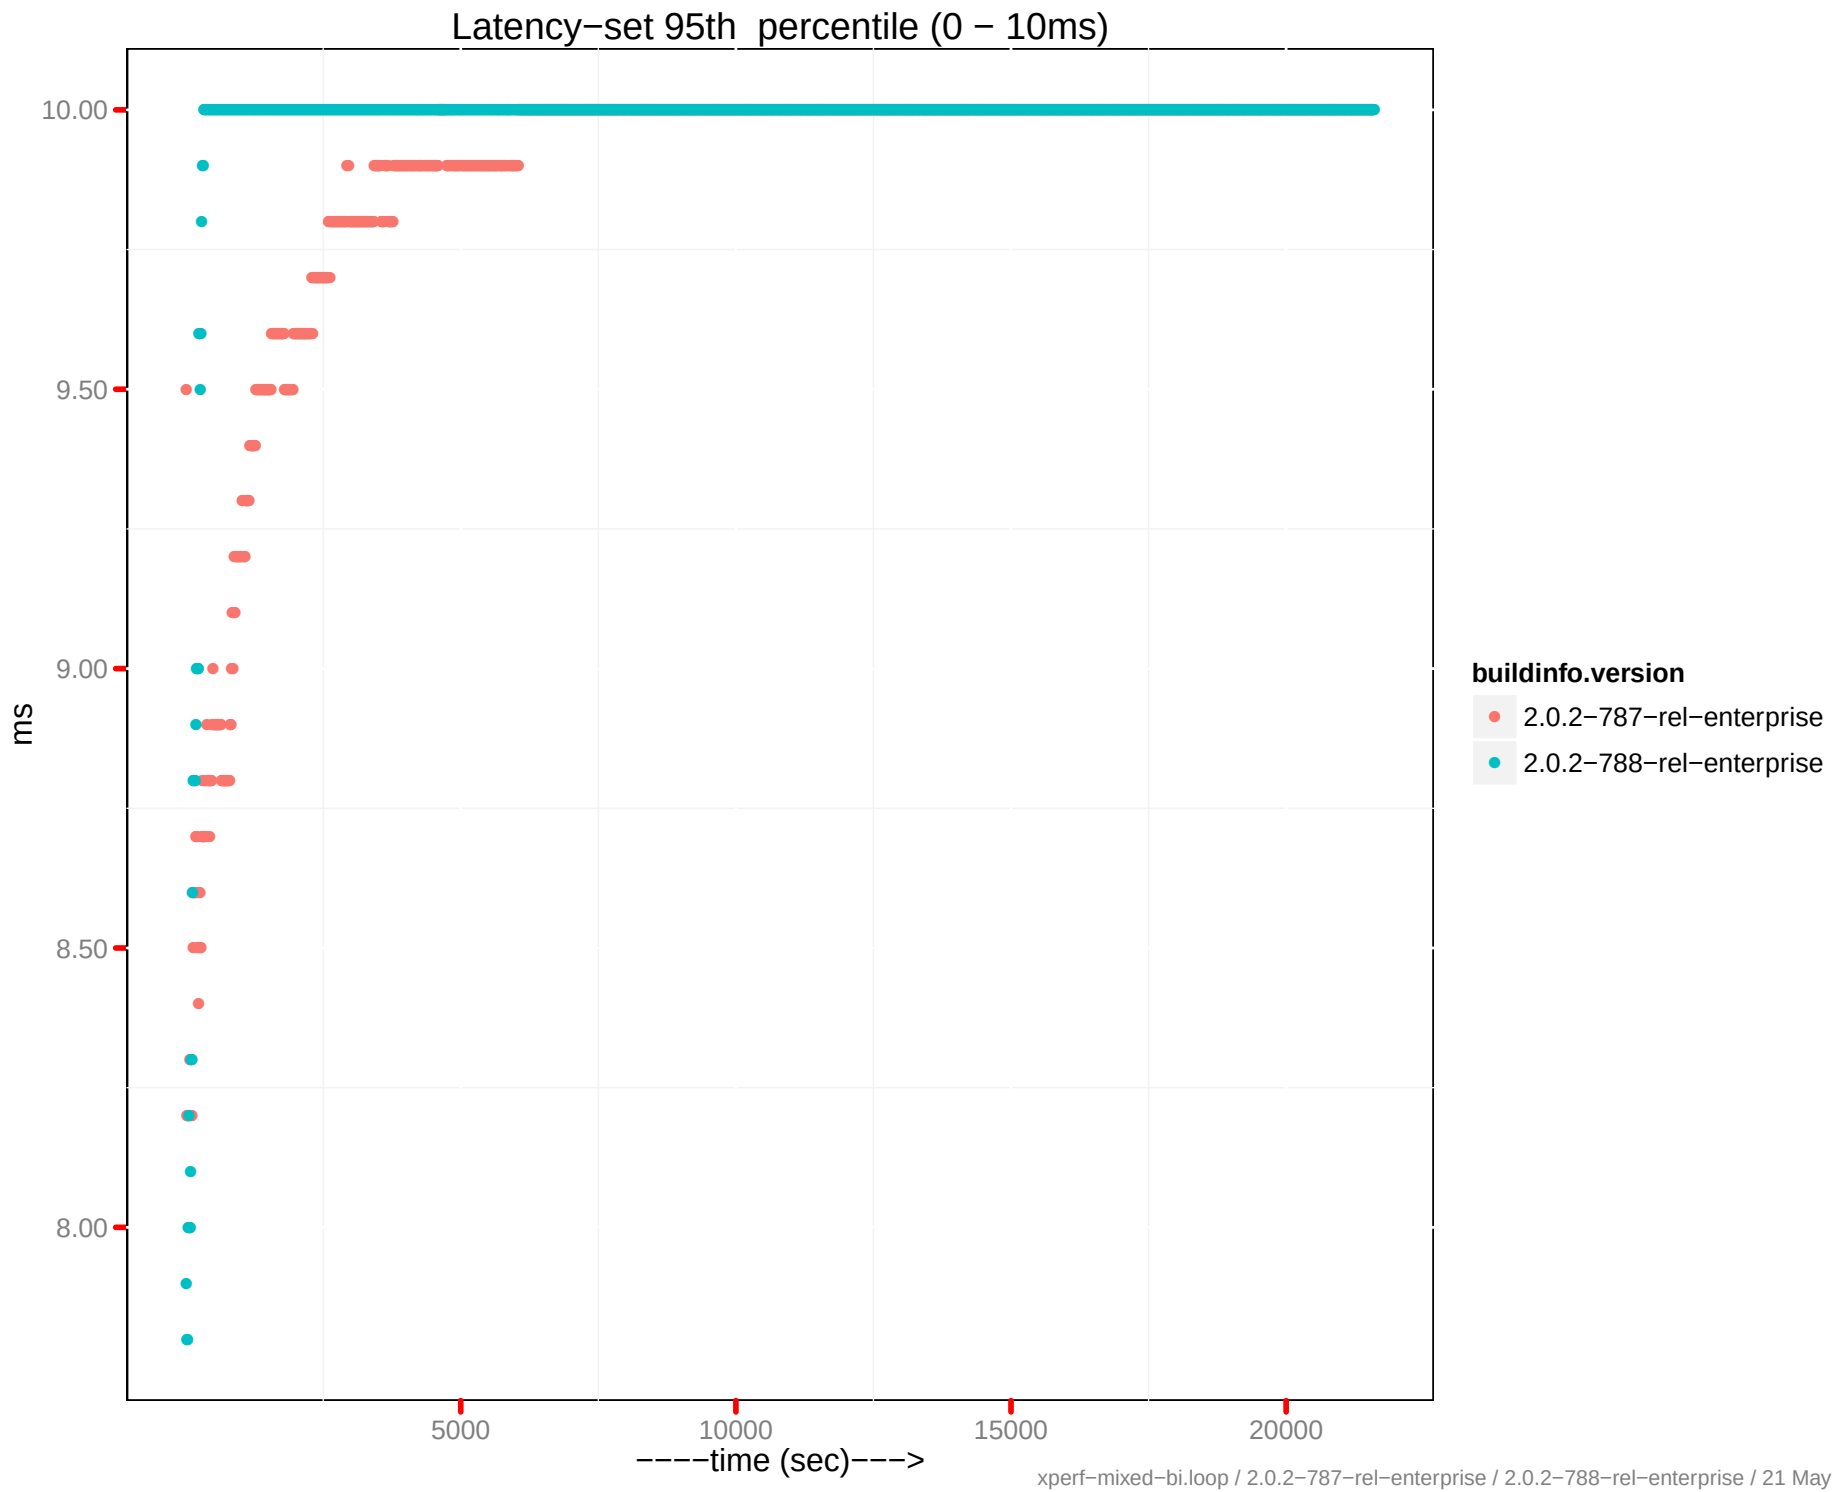
Latency−set 95th percentile (0 − 10ms)
10.00
9.50
9.00
8.50
8.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

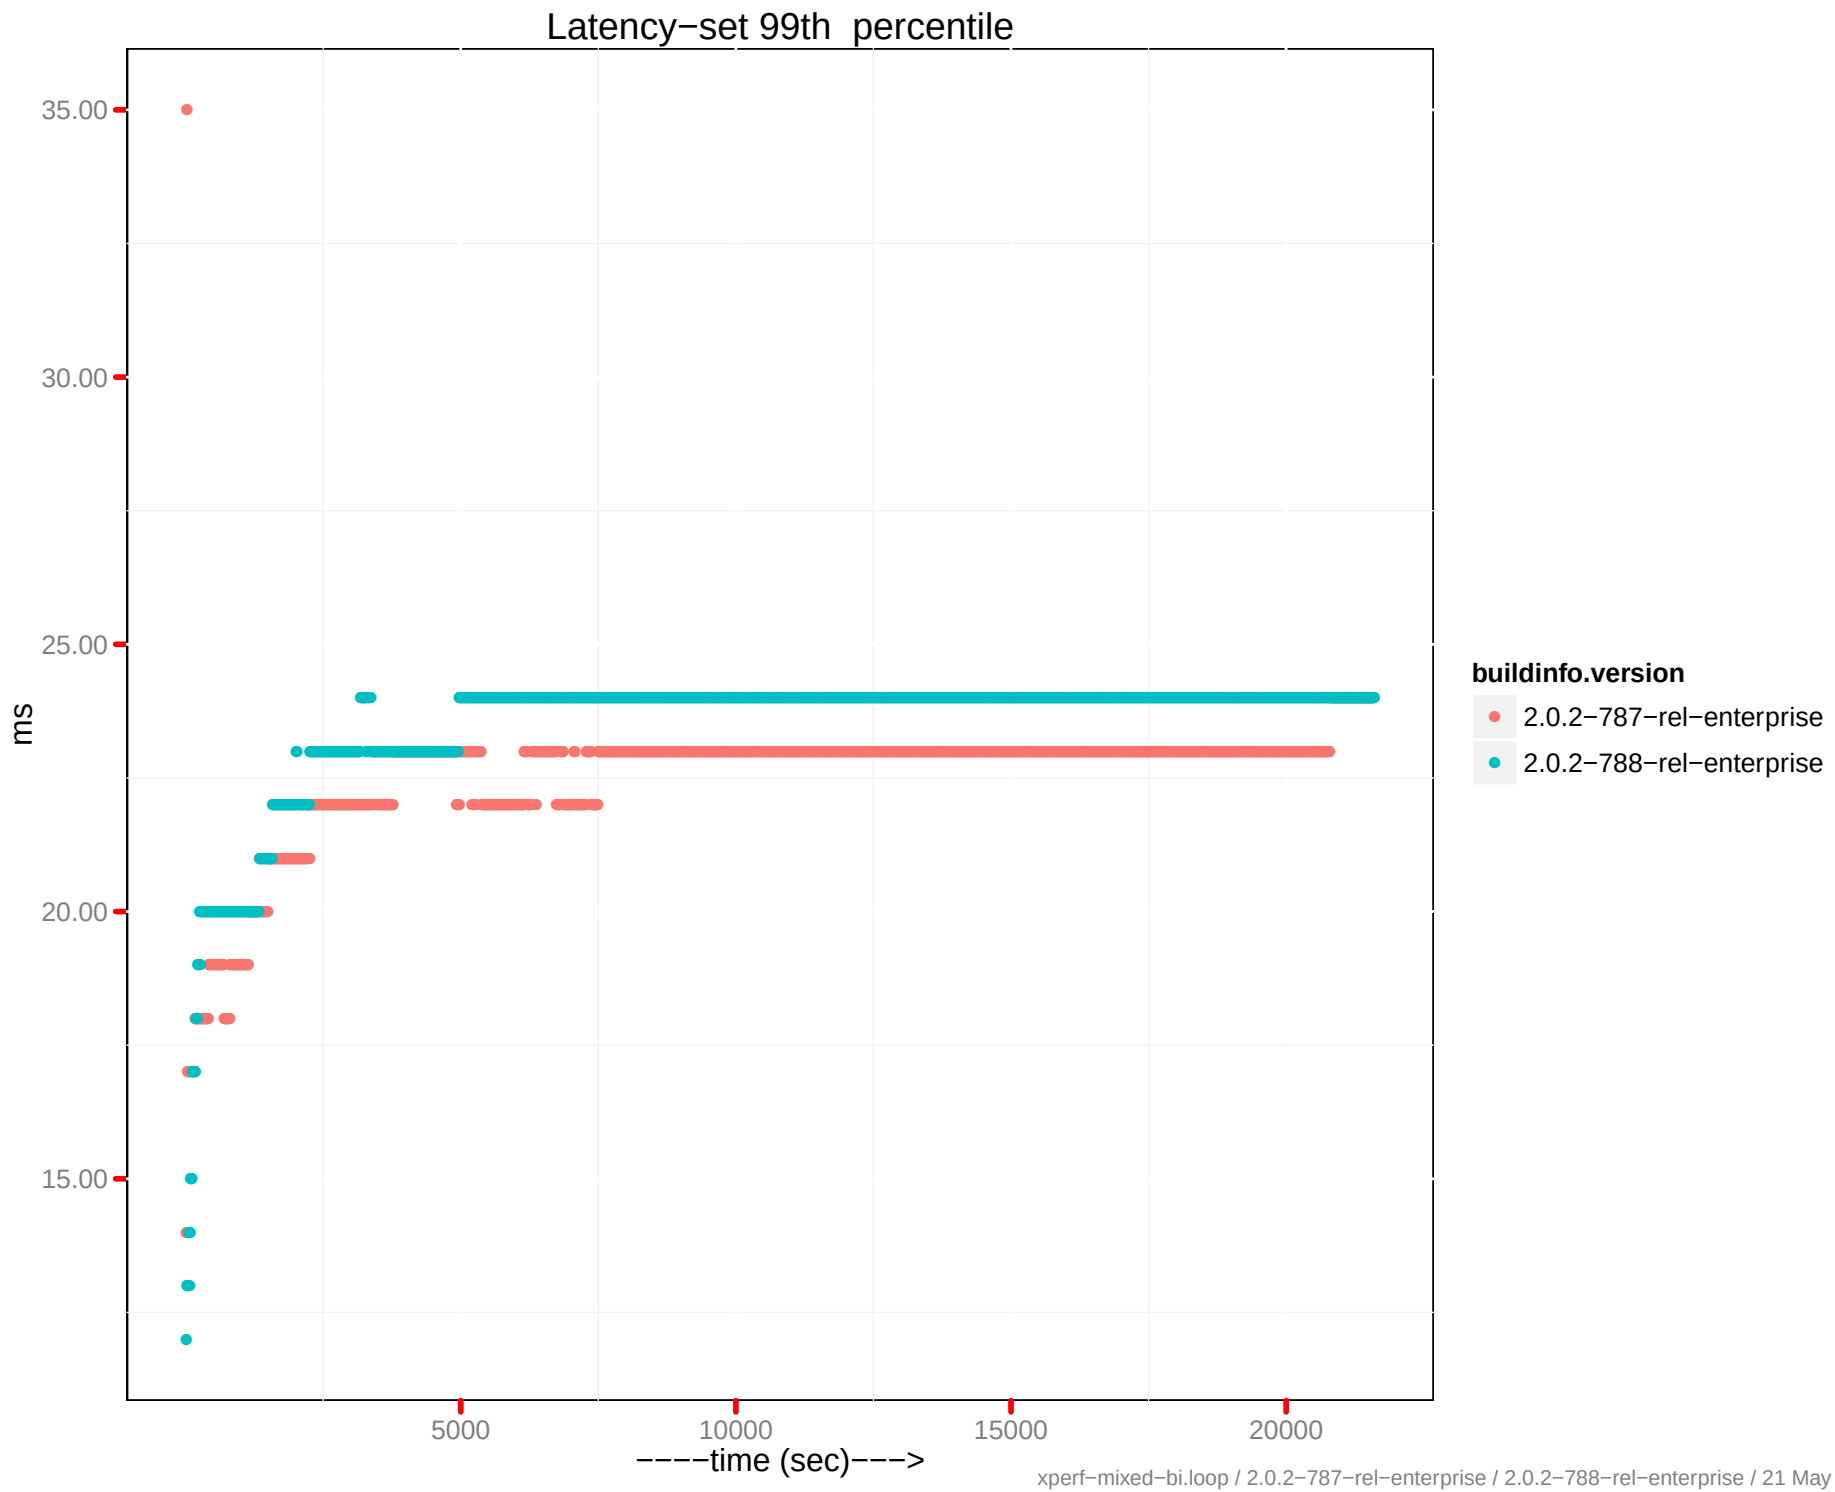
Latency−set 99th percentile
35.00
30.00
25.00
20.00
15.00
ms
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

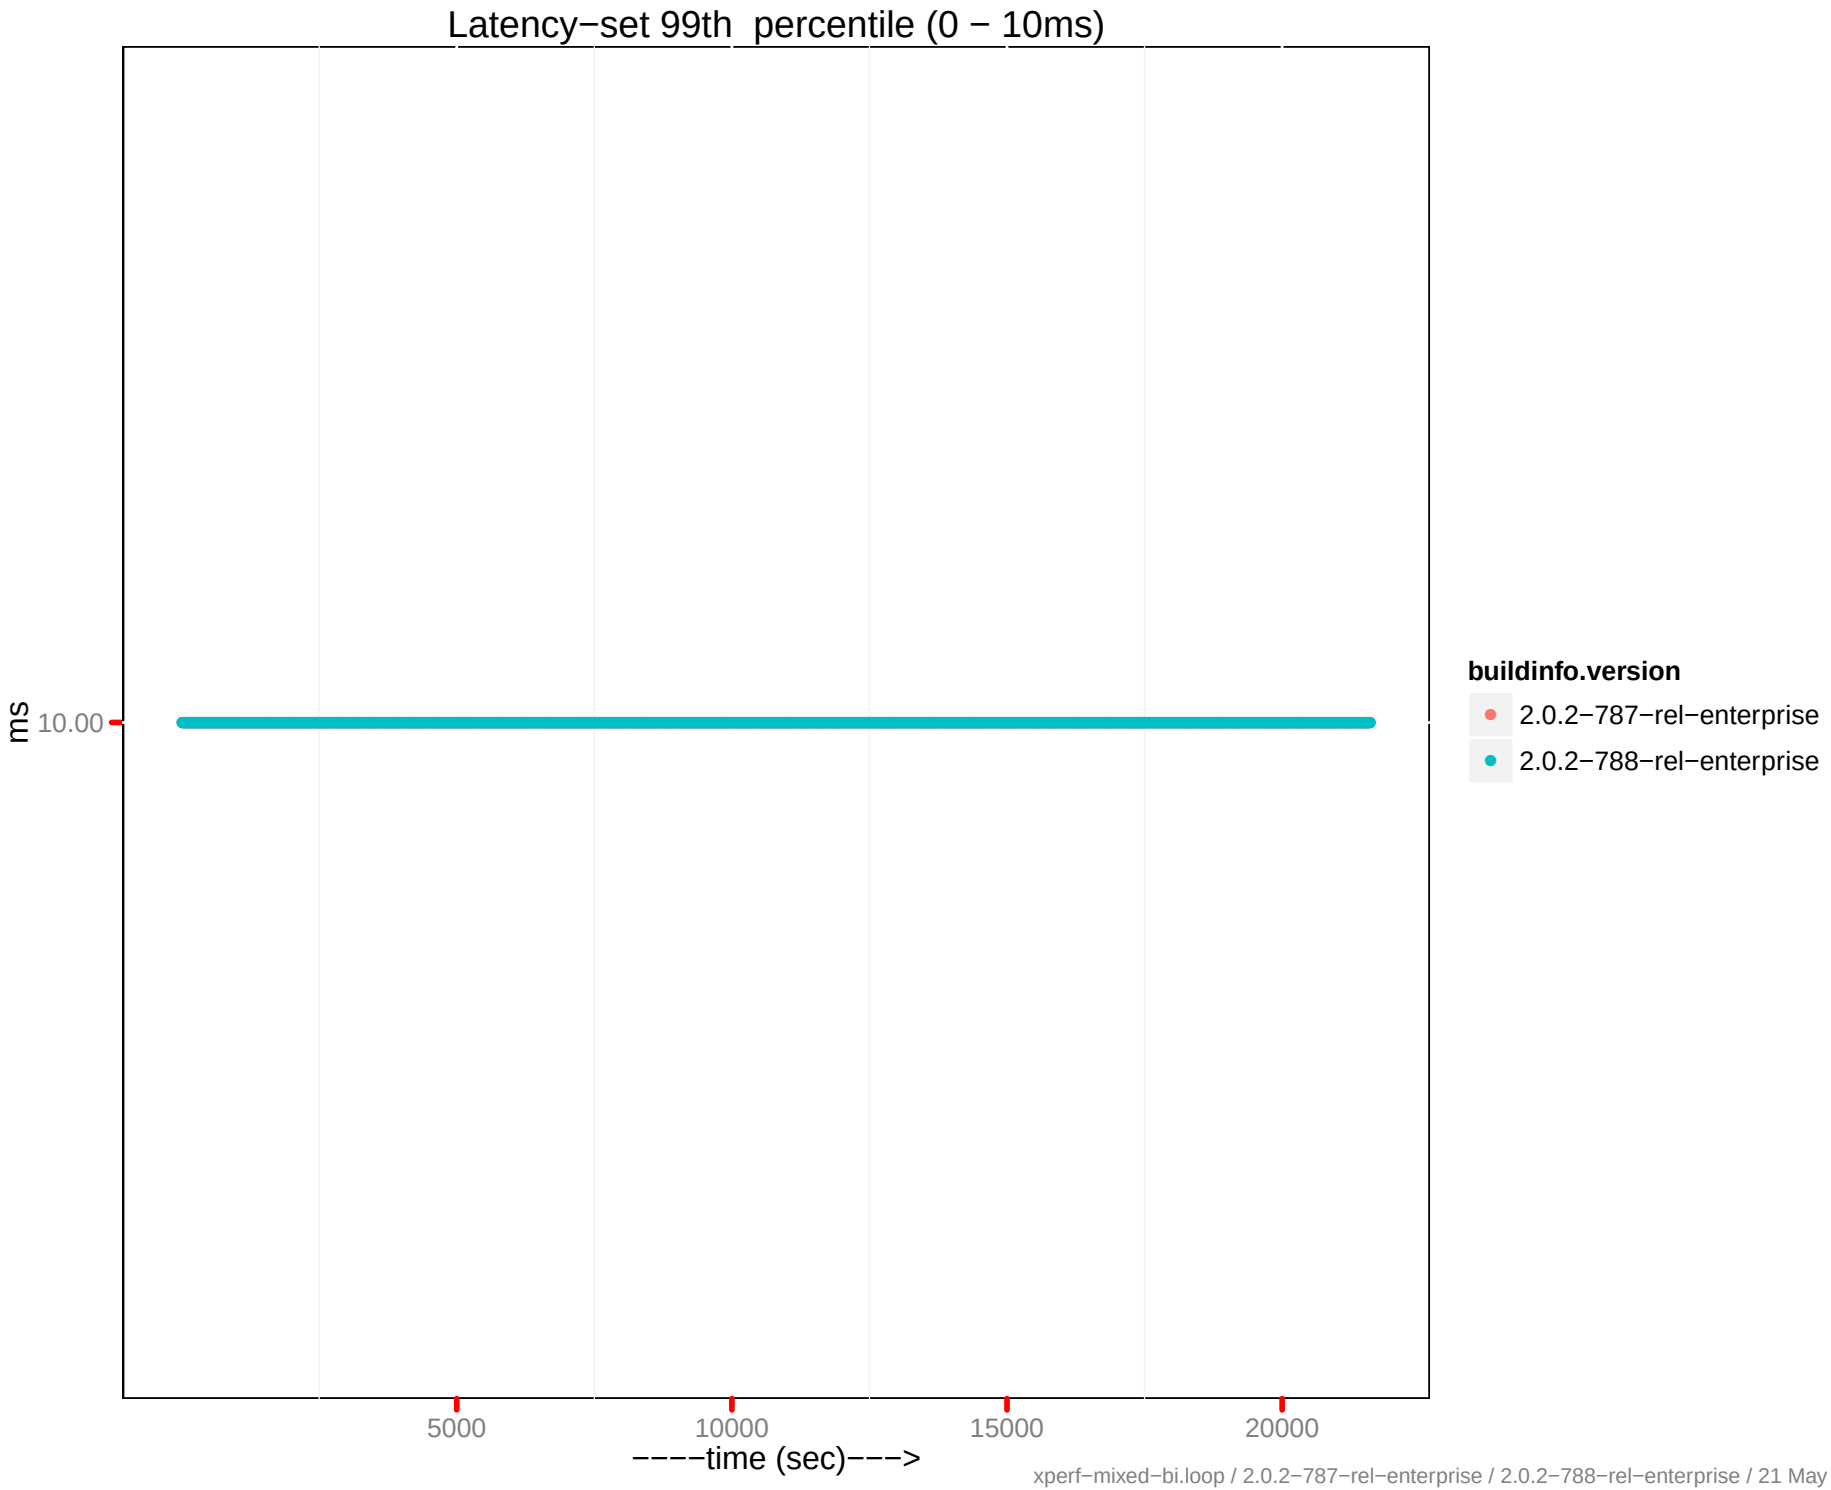
Latency−set 99th percentile (0 − 10ms)
ms
10.00
5000
10000
15000
20000
−−−−time (sec)−−−>
buildinfo.version
2.0.2−787−rel−enterprise
2.0.2−788−rel−enterprise
xperf−mixed−bi.loop / 2.0.2−787−rel−enterprise / 2.0.2−788−rel−enterprise / 21 May 2013

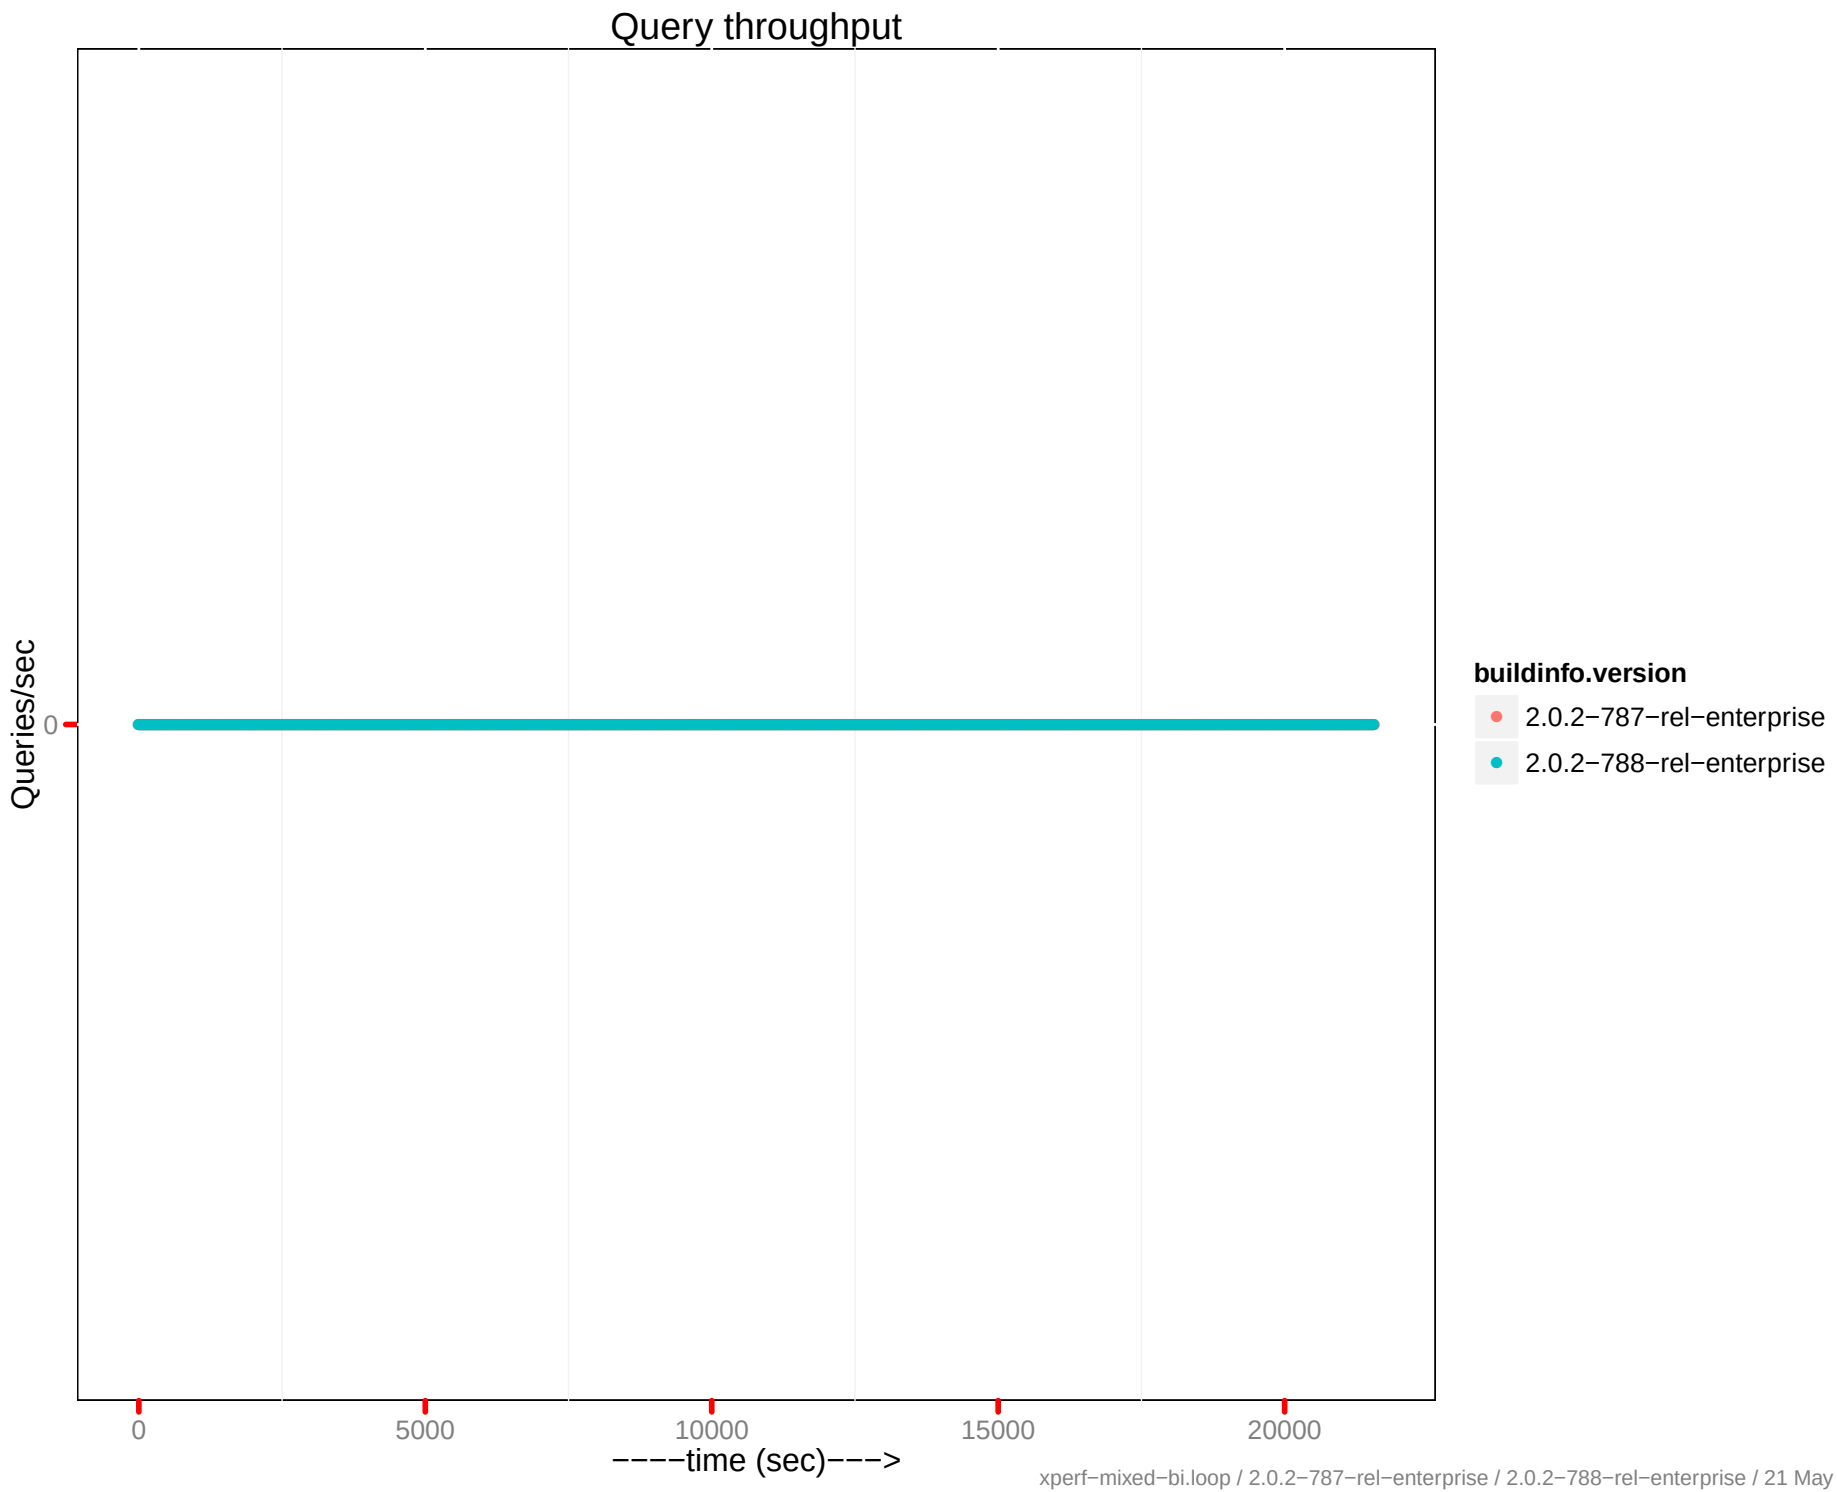
Query throughput
Queries/sec
0
0
5000
10000
15000
20000
----time (sec)--->
buildinfo.version
2.0.2-787-rel-enterprise
2.0.2-788-rel-enterprise
xperf-mixed-bi.loop / 2.0.2-787-rel-enterprise / 2.0.2-788-rel-enterprise / 21 May 2013