



	2.0.0 – 1792	2.0.0 – 1792
<i>Runtime (in hr)</i>	12.93	NA
<i>Avg. Drain Rate</i>	535.60	NANA
<i>Peak Disk (GB)</i>	135.35	NA
<i>Peak Memory (GB)</i>	20.31	NA
<i>Avg. OPS</i>	7.07K	NANA
<i>Avg. mem memcached (GB)</i>	19.94	NA
<i>Avg. mem beam.smp (MB)</i>	345.63	NA
<i>Avg. CPU rate (%)</i>	74.31	NA
<i>Latency-get (90th) (ms)</i>	1.14	NA
<i>Latency-get (95th) (ms)</i>	7.41	NA
<i>Latency-get (99th) (ms)</i>	75.29	NA
<i>Latency-set (90th) (ms)</i>	0.95	NA
<i>Latency-set (95th) (ms)</i>	1.87	NA
<i>Latency-set (99th) (ms)</i>	5.45	NA
<i>Latency-query (80th) (ms)</i>	24.68	NA
<i>Latency-query (90th) (ms)</i>	31.65	NA
<i>Latency-query (95th) (ms)</i>	40.13	NA
<i>Latency-query (99th) (ms)</i>	146	NA
<i>Latency-query (99.9th) (ms)</i>	1755.21	NA
<i>Avg. QPS</i>	1025.95	NA
<i>Avg. XDC ops/sec</i>	NaN	NA
<i>Avg. XDC queue</i>	NaN	NA
<i>Rebalance Time (sec)</i>	0	NA
<i>Testrunner Version</i>	02c968c	NA


View read per sec.




ep queue size



ns_server: ep_diskqueue_drain




ep-engine : ep_diskqueue_drain - 10.2.1.65




buildinfo.version
● 2.0.0-1792-rel-enterprise

Total number of items per second being written to disk in this node (from ep-engine)

ep-engine : ep_diskqueue_drain - 10.2.1.66




ep-engine : ep_diskqueue_drain - 10.2.1.67



buildinfo.version
● 2.0.0-1792-rel-enterprise

Total number of items per second being written to disk in this node (from ep-engine)


ep-engine : ep_diskqueue_drain - 10.2.1.68




buildinfo.version
● 2.0.0-1792-rel-enterprise

Total number of items per second being written to disk in this node (from ep-engine)

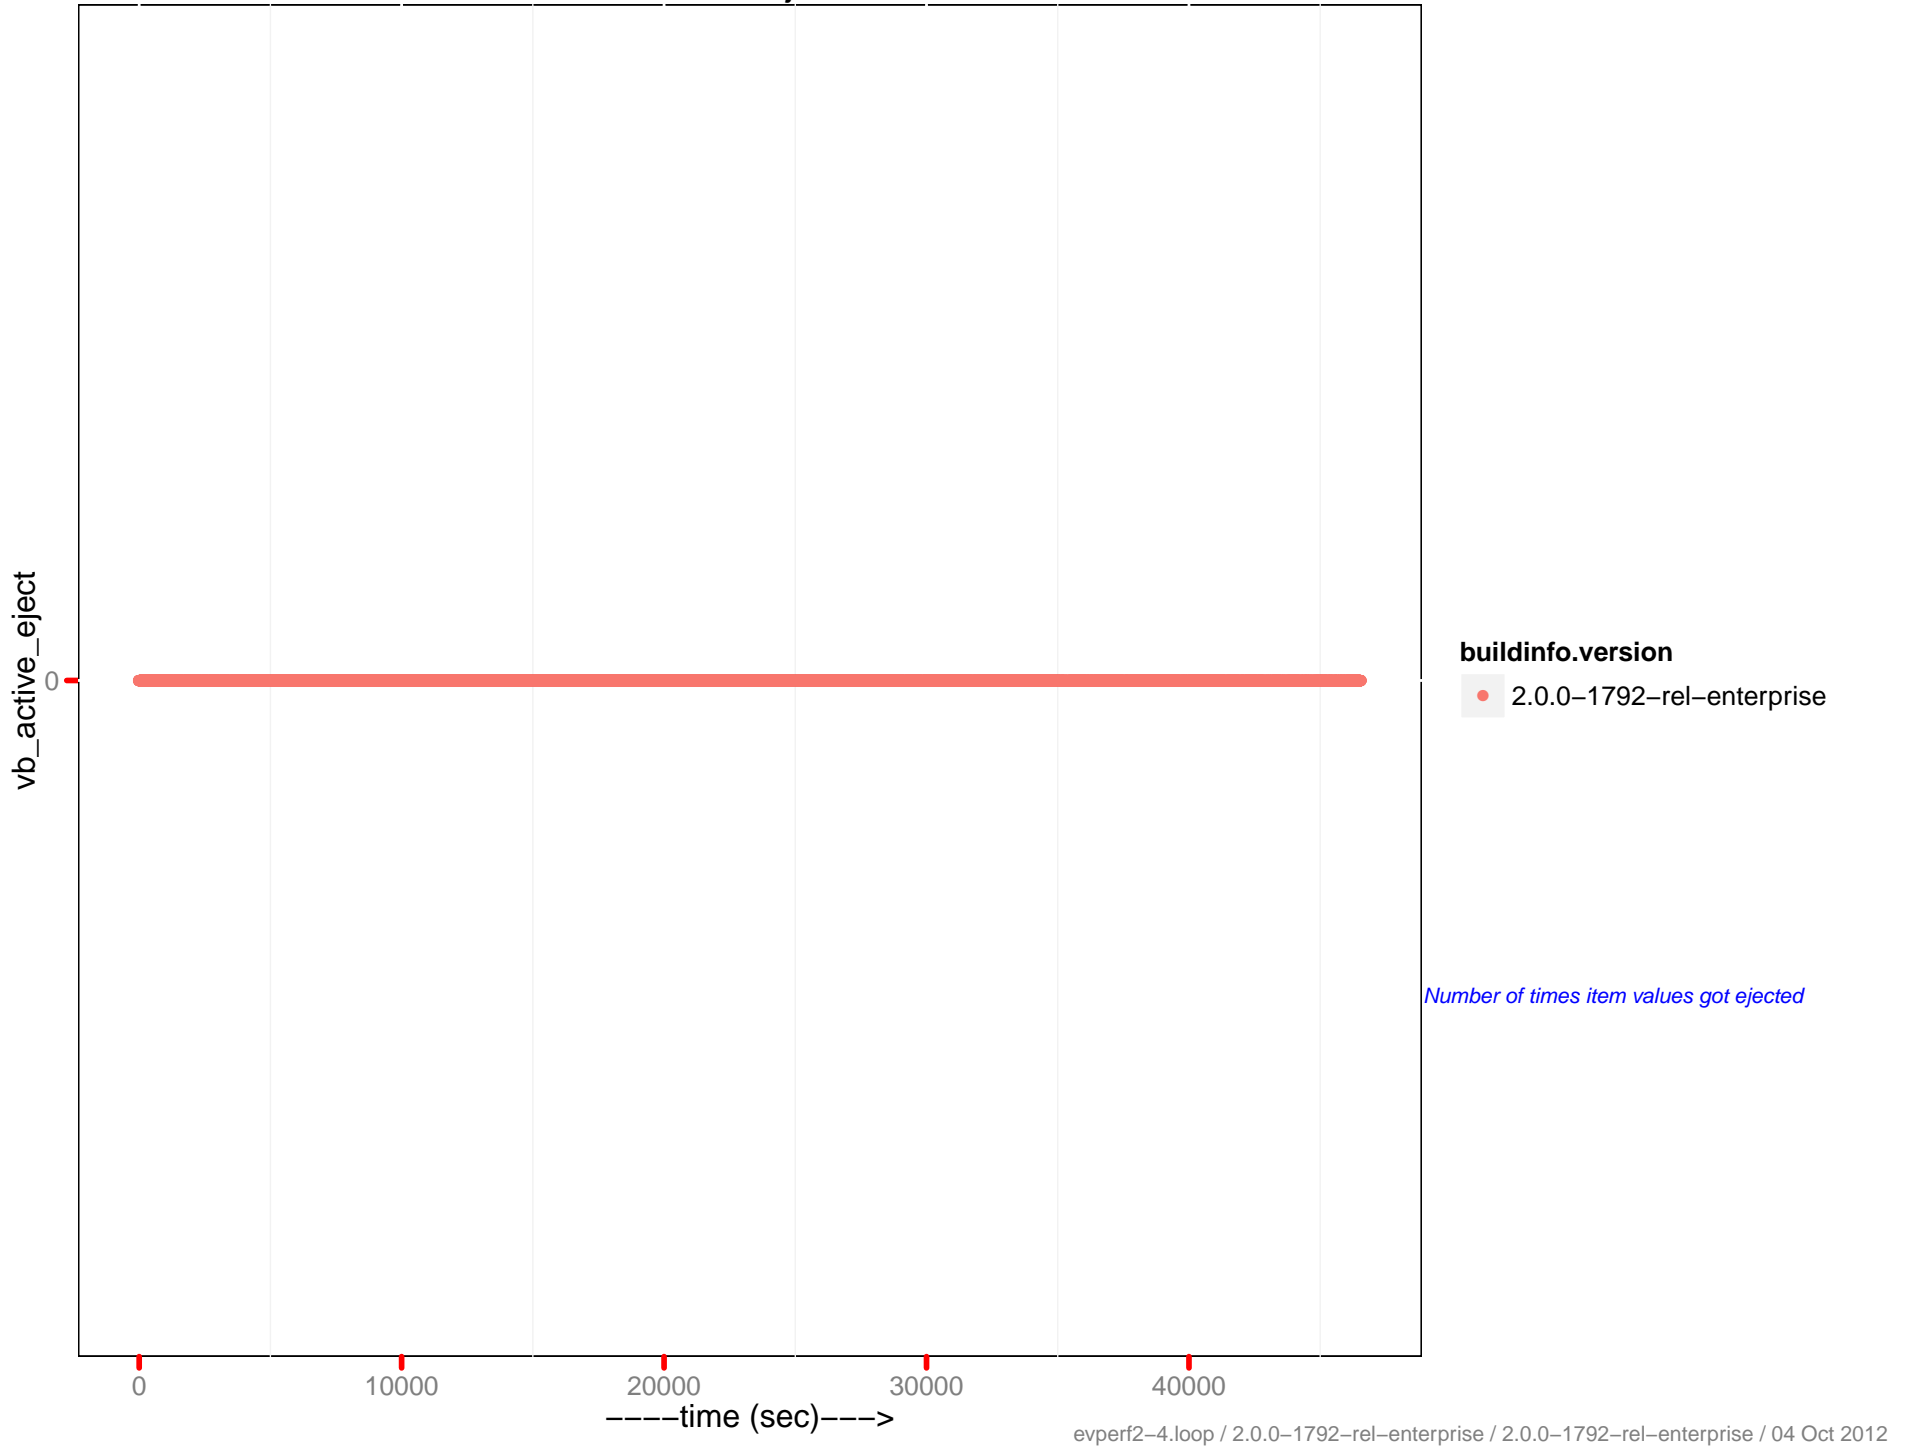
ep_bg_fetched ops/sec




tmp_oom ops/sec




vb_active_eject/sec




vb_replica_eject/sec




ep_tap_replica_queue_drain/sec




ep_tap_replica_queue_backoff/sec




vb_active_resident_items_ratio




vb_replica_resident_items_ratio




curr_items




cur_items_total




mem_used




Docs data size




Docs disk size




Docs actual disk size




Views data size




Views disk size




Views actual disk size



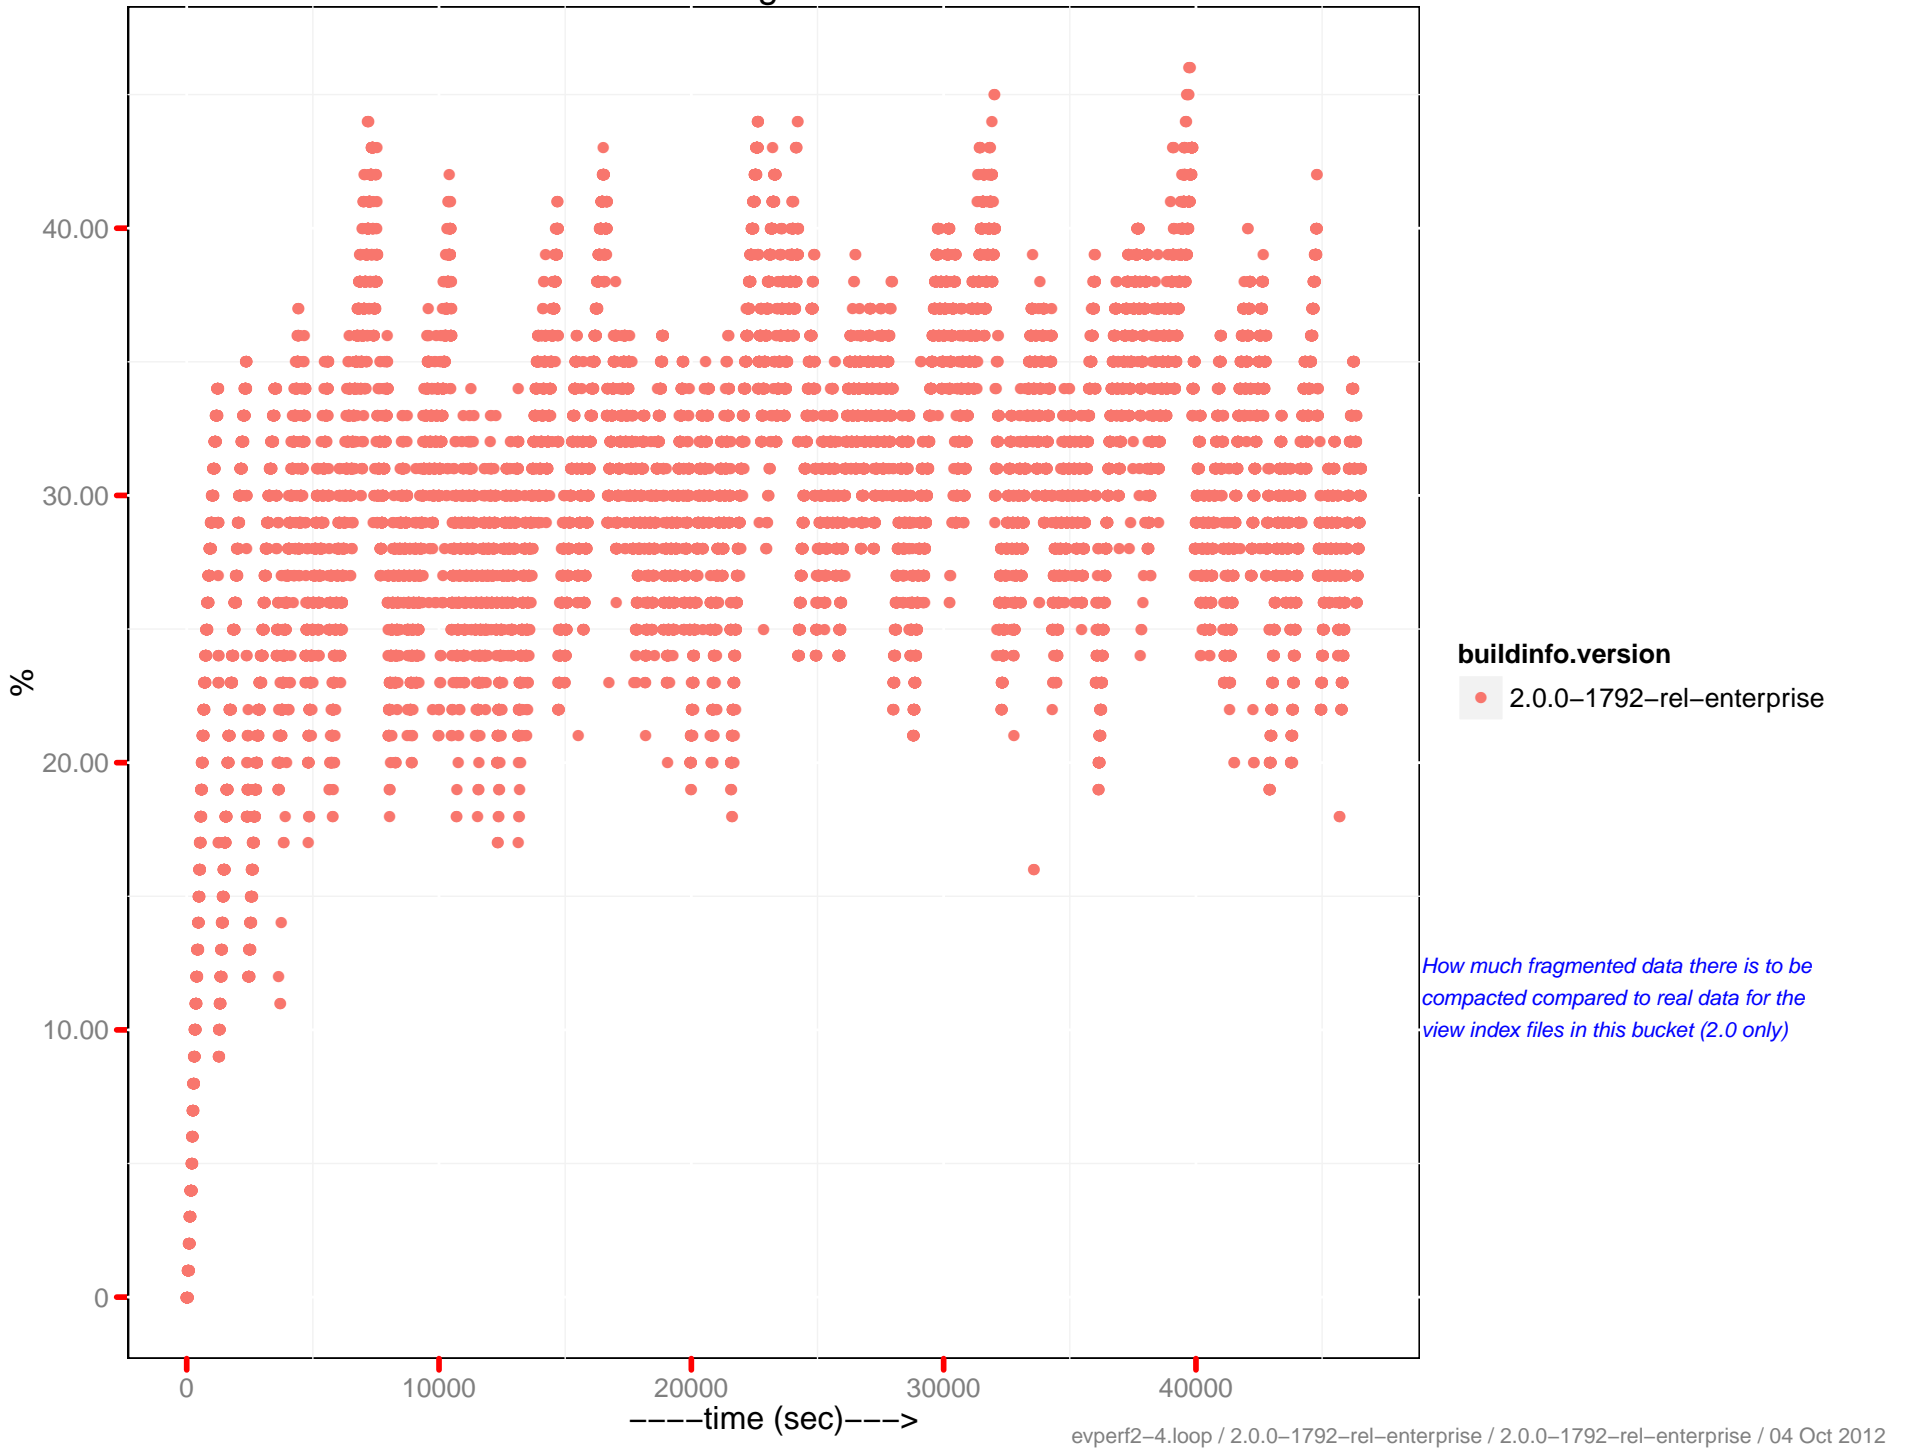
Total disk size




Docs fragmentation




Views fragmentation




cmd_get ops/sec




cmd_set ops/sec




of get misses




of get hits




cache_miss percentage




cache_miss percentage 0-5




Number of connections




XDC replication Queue




XDC ops per sec



XDC gets per sec




XDC sets per sec



Number of set operations per second related to this bucket being the target of cross datacenter replication

XDC dels per sec




buildinfo.version


- 2.0.0-1792-rel-enterprise

*Number of del operations per second
related to this bucket being the
target of cross datacenter replication*


XDCR secs working




XDCR secs committing




CPU utilization – 10.2.1.65:8091




CPU utilization – 10.2.1.66:8091




CPU utilization – 10.2.1.67:8091




CPU utilization – 10.2.1.68:8091




SWAP Usage - 10.2.1.65:8091




SWAP Usage - 10.2.1.66:8091




SWAP Usage - 10.2.1.67:8091




SWAP Usage - 10.2.1.68:8091




cpu_time_diff: memcached - 10.2.1.65




cpu_time_diff : beam.smp - 10.2.1.65




cpu_time_diff: memcached - 10.2.1.66




cpu_time_diff : beam.smp - 10.2.1.66




cpu_time_diff: memcached - 10.2.1.67




cpu_time_diff : beam.smp - 10.2.1.67




cpu_time_diff: memcached - 10.2.1.68




cpu_time_diff : beam.smp - 10.2.1.68




Disk Read (kB/s) : 10.2.1.65




Disk Write (kB/s) : 10.2.1.65




Data disk size




Latency get histogram




Latency get histogram (0-10 ms)




Latency set histogram




Latency set histogram (0–10 ms)




Latency query histogram




Latency-get 90th percentile




Latency-get 90th percentile (0 - 10ms)




Latency-get 95th percentile




Latency-get 95th percentile (0 - 10ms)




Latency-get 99th percentile




Latency-get 99th percentile (0 - 10ms)




Latency-set 90th percentile




Latency-set 95th percentile (0 - 10ms)




Latency-set 99th percentile




Latency-query 80th percentile




Latency-query 80th percentile (0 - 10ms)




Latency-query 90th percentile




Latency-query 90th percentile (0 - 10ms)




Latency-query 95th percentile




Latency-query 95th percentile (0 - 10ms)




Latency-query 99th percentile (0 - 10ms)




Latency-query 99.9th percentile




Latency-query 99.9th percentile (0 - 10ms)




Query throughput




Indexing time – 10.2.1.65




Indexing time – 10.2.1.66




Indexing time – 10.2.1.67




Indexing time – 10.2.1.68




Indexing time (0-5 sec) - 10.2.1.65




Indexing time (0-5 sec) - 10.2.1.66



Indexing time (0-5 sec) - 10.2.1.67



Indexing time (0-5 sec) - 10.2.1.68



```
evperf2-4.conf
# "EVPERF'2" view performance test:
# 3 ddocs with 8 views per ddoc
# 20M initial items
# 25GB RAM quota (32GB total RAM)
# DGM
# 30 clients
# 8K ops/sec total background load (memcached commands)
# 80% reads, 20% write (12% updates/deletes, 8% inserts)
# Cache miss ratio < 1%
# Stop on 36M total queries (tuned to be > 14 hours)

performance.ipperf.MultiClientTests.test_evperf2

params:

# general
batch=50
kind=json
mem_quota=25000
loglevel=error

# load phase
items=20000000
hot_init_items=16000000

# access phase
ratio_sets=0.2
ratio_misses=0.04
ratio_creates=0.40
ratio_deletes=0.50
ratio_hot=0.2
ratio_hot_gets=0.975
ratio_hot_sets=0.975
ratio_expirations=0.0
bg_max_ops_per_sec=265
fg_max_ops=56000000
total_clients=30
start_delay=5

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

```
vesta.ini
[global]
username:root
password:couchbase
port:8091
data_path:/data
index_path:/data2

[servers]
1:10.2.1.65
2:10.2.1.66
3:10.2.1.67
4:10.2.1.68

[clients]
1:10.2.1.60

[membase]
rest_username:Administrator
rest_password:password

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