

Contents: [Dobrica PavlinuÄiÄ 's random unstructured stuff]

- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(MongoDB\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Checkout source\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Install build dependencies\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Build Debian package\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Perl driver\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Binaries\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Queries\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(PXElator audit examples\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Profile\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Indexes\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Comparison with CouchDB\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Disk usage\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Map/Reduce differences\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Blog posts\)](#)
- [Dobrica PavlinuÄiÄ 's random unstructured stuff \(Debian amd64 version\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(build\)](#)
 - ◆ [Dobrica PavlinuÄiÄ 's random unstructured stuff \(run\)](#)

MongoDB

Checkout source

```
dpavlin@t61p:/rest/cvs$ git clone git://github.com/mongodb/mongo.git
Initialized empty Git repository in /rest/cvs/mongo/.git/
remote: Counting objects: 32011, done.
remote: Compressing objects: 100% (9340/9340), done.
remote: Total 32011 (delta 22724), reused 31556 (delta 22412)
Receiving objects: 100% (32011/32011), 20.57 MiB | 1.12 MiB/s, done.
Resolving deltas: 100% (22724/22724), done.
```

Install build dependencies

```
dpavlin@t61p:/rest/cvs/mongo$ sudo apt-get install \
libboost-dev libboost-thread-dev libboost-filesystem-dev libboost-program-options-dev libboost-d
libpcre3-dev xulrunner-dev libreadline-dev
```

Build Debian package

debian/control file needs modification for unstable:

<http://svn.rot13.org/index.cgi/pxelator/view/mongodb/mongo-debian-control-xulrunner.diff>

```
dpavlin@t61p:/rest/cvs$ cd mongo/

# patch source
dpavlin@klin:/rest/cvs/mongo$ patch -p1 < /srv/pxelator/mongodb/mongo-debian-control-xulrunner.diff
patching file debian/control

# clean before new build
dpavlin@t61p:/rest/cvs/mongo$ sudo rm -Rf debian/mongodb

dpavlin@t61p:/rest/cvs/mongo$ time dpkg-buildpackage -rfakeroot -b

...

real    6m16.744s
user    5m41.701s
sys     0m19.393s
```

Perl driver

```
dpavlin@t61p:/rest/cvs$ git clone git://github.com/mongodb/mongo-perl-driver.git
Initialized empty Git repository in /rest/cvs/mongo-perl-driver/.git/
remote: Counting objects: 1782, done.
remote: Compressing objects: 100% (1673/1673), done.
remote: Total 1782 (delta 1122), reused 0 (delta 0)
Receiving objects: 100% (1782/1782), 1.45 MiB | 747 KiB/s, done.
Resolving deltas: 100% (1122/1122), done.

sudo apt-get install libany-moose-perl libdata-types-perl

dpavlin@t61p:/rest/cvs$ cd mongo-perl-driver/

perl Makefile.PL
make test
sudo dh-make-perl
```

Binaries

- <http://debian.rot13.org/>

Queries

PXElator audit examples

```
> use pexlator

> db.audit.group({ key:{ 'package.name':true }, initial:{ count: 0 }, reduce:function(o,p) { p.co

> show profile
```

```

11052ms Sun Jan 31 2010 13:24:47
query pxelator.$cmd ntoreturn:1 reslen:690 nscanned:0
query: { group: { key: { package.name: true }, initial: { count: 0.0 }, ns: "audit", $reduce: fun
  p.count++;
} } } nreturned:1 bytes:674 11052ms

> db.audit.ensureIndex({ 'package.name':true })

> db.audit.group({ key:{ 'package.name':true }, initial:{ count: 0 }, reduce:function(o,p) { p.co

```

no visible speed impact.

We are really interested only in daemons which aren't null:

```

> db.audit.ensureIndex( { daemon: true } )
> db.audit.group({
  key: { daemon:true }
,cond: { daemon: { $exists: true } }
,initial: { count: 0 }
,reduce: function(o,p) { p.count++ }
});

```

dhcp count usage by ip

```

> db.audit.ensureIndex( { "package.name": true } )

> db.audit.group({ key:{ ip:true }, cond: { "package.name": "dhcpd" }, initial: { count: 0 }, red

```

package usage

```

> db.setProfilingLevel(2,1000);

> db.audit.group({ key:{ "package.name":true }, initial:{ count:0 }, reduce:function(o,p){ p.coun

> db.system.profile.find().sort({$natural:-1}).limit(10)

{ "ts" : "Sun Jan 24 2010 15:07:53 GMT+0100 (CET)", "info" : "query pxelator.$cmd ntoreturn:1 res
query: { group: { key: { package.name: true }, initial: { count: 0.0 }, ns: \"audit\", $reduce: f
  p.count++;
} } } nreturned:1 bytes:626 13887ms", "millis" : 13887 }

> db.audit.ensureIndex({ "package.name":true })

> db.audit.group({ key:{ "package.name":true }, initial:{ count:0 }, reduce:function(o,p){ p.coun

```

doesn't help much, because we don't have cond in query.

Profile

```

> db.setProfilingLevel(2,1000);
{ "was" : 2, "ok" : 1 }

```

```
> db.system.profile.find()
```

Indexes

```
> db.system.indexes.find()
{ "name" : "_id_", "ns" : "pxelator.audit", "key" : { "_id" : ObjectId("000000000000000000000000")
{ "ns" : "pxelator.audit", "key" : { "daemon" : true }, "name" : "daemon_" }
{ "ns" : "pxelator.audit", "key" : { "key" : "package.time" }, "name" : "key_" }
{ "ns" : "pxelator.audit", "key" : { "package.name" : true }, "name" : "package.name_" }
```

Comparison with CouchDB

Migrate from CouchDB to MongoDB using

<http://svn.rot13.org/index.cgi/pxelator/view/bin/couchdb2mongodb.pl>

Disk usage

```
root@opr:~# du -hc /var/lib/couchdb/0.9.0/.pxelator* /var/lib/couchdb/0.9.0/pxelator.couch
655M    /var/lib/couchdb/0.9.0/.pxelator_design
23M     /var/lib/couchdb/0.9.0/.pxelator_temp
7.8G    /var/lib/couchdb/0.9.0/pxelator.couch
8.4G    total
```

```
root@opr:~# du -hc /var/lib/mongodb/pxelator.*
65M     /var/lib/mongodb/pxelator.0
129M    /var/lib/mongodb/pxelator.1
257M    /var/lib/mongodb/pxelator.2
513M    /var/lib/mongodb/pxelator.3
513M    /var/lib/mongodb/pxelator.4
513M    /var/lib/mongodb/pxelator.5
17M     /var/lib/mongodb/pxelator.ns
2.0G    total
```

Map/Reduce differences

CouchDB

```
# map
function(doc) {
  if ( doc.package.name == 'dnsd' )
    emit(doc.peerhost,1);
}

# reduce
function (k,v) {
  return sum(v);
}
```

MongoDB

```

> m = function() { emit(this.peerhost,1) }

> r = function(k,vals) { var sum = 0; for (var i in vals) sum += vals[i]; return sum; }

> res = db.audit.mapReduce(m, r, { query:{"package.name":"dnssd"} } )
{
  "result" : "tmp.mr.mapreduce_1264448081_3",
  "timeMillis" : 6040,
  "counts" : {
    "input" : {
      "top" : 0,
      "bottom" : 204293
    },
    "emit" : {
      "top" : 0,
      "bottom" : 204293
    },
    "output" : {
      "top" : 0,
      "bottom" : 22
    }
  },
  "ok" : 1,
}

> db[res.result].find().limit(10)

```

Comparison with ad-hoc query

```

> db.setProfilingLevel(2,1000);

> db.audit.group({ key:{ "peerhost":true }, cond:{ "package.name":"dnssd" },
  initial:{ count:0 }, reduce:function(o,p){ p.count++ } })

> db.system.profile.find().sort({$natural:-1}).limit(10)

{ "ts" : "Mon Jan 25 2010 21:21:11 GMT+0100 (CET)", "info" : "query pxelator.$cmd nreturn:1 res
query: { group: { key: { peerhost: true }, cond: { package.name: \"dnssd\" }, initial: { count: 0.
  p.count++;
} } } nreturned:1 bytes:1132 2161ms", "millis" : 2161 }

```

So, going through server-side JavaScript is **3x performance penalty**

Blog posts

[Search Results for MongoDB](#)

- [It's about system. stupid!](#)
- [MongoDB - so you want fast NoSQL database which you can grok](#)

Debian amd64 version

build

```
root@klin:~/rest/virtual# debootstrap --arch amd64 squeeze ./mongodb-amd64 http://10.60.0.91:3142
root@klin:~/rest/virtual# chroot mongodb-amd64/

root@klin:/# apt-get install \
  git-core locales dpkg-dev debhelper scons \
  libboost-dev libboost-thread-dev libboost-filesystem-dev libboost-program-options-dev libboost-d
  libpcre3-dev xulrunner-dev libreadline-dev

root@klin:/# cd /srv/
root@klin:/srv# git clone git://github.com/mongodb/mongo.git

root@klin:/srv# cd mongo/
root@klin:/srv/mongo# time dpkg-buildpackage -rfakeroot -b
```

run

```
dpavlin@klin:~$ sudo chroot /virtual/mongodb-amd64/ su -c '/usr/bin/mongod --dbpath /var/lib/mon
```