

Server administration

General Linux server management notes, not specific to anything in particular.

- [Batch-migrating Gitolite repositories to Gogs](#)
- [What is\(n't\) Docker actually for?](#)
- [Blocking LLM scrapers on Alibaba Cloud from your nginx configuration](#)
- [Dealing with a degraded btrfs array due to disk failure](#)

Batch-migrating Gitolite repositories to Gogs

This article was originally published at

<https://gist.github.com/joepie91/2ff74545f079352c740a>.

NOTE: This will only work if you are an administrator on your Gogs instance, or if an administrator has enabled local repository importing for all users.

First, save the following as `migrate.sh` somewhere, and make it executable (`chmod +x migrate.sh`):

```
HOSTNAME="git.crypto.net"
BASEPATH="/home/git/old-repositories/projects/joepie91"

OWNER_ID="$1"
CSRF=`cat ./cookies.txt | grep _csrf | cut -f 7`

while read REPO; do
  REPO_NAME=`echo "$REPO" | sed "s/\.git\$//"`
  curl "https://$HOSTNAME/repo/migrate" \
    -b "./cookies.txt" \
    -H 'origin: null' \
    -H 'content-type: application/x-www-form-urlencoded' \
    -H "authority: $HOSTNAME" \
    --data "_csrf=$CSRF" \
    --data-urlencode "clone_addr=$BASEPATH/$REPO" \
    --data-urlencode "uid=$OWNER_ID" \
    --data-urlencode "auth_username=" \
    --data-urlencode "auth_password=" \
    --data-urlencode "repo_name=$REPO_NAME" \
    --data-urlencode "description=Automatically migrated from Gitolite"
done
```

Change `HOSTNAME` to point at your Gogs installation, and `BASEPATH` to point at the folder where your Gitolite repositories live on the filesystem. It must be the entire base path - the repository names cannot contain slashes!

Now save the Gogs cookies from your browser as `cookies.txt`, and create a file (eg. `repositories.txt`) containing all your repository names, each on a new line. It could look something like this:

```
project1.git  
project2.git  
project3.git
```

After that, run the following command:

```
cat repositories.txt | ./migrate.sh 1
```

... where you replace `1` with your User ID on your Gogs instance.

Done!

What is(n't) Docker actually for?

This article was originally published at

<https://gist.github.com/joepie91/1427c8fb172e07251a4bbc1974cdb9cd>.

This article was written in 2016. Some details may have changed since.

A brief listing of some misconceptions about the purpose of Docker.

Secure isolation

Some people try to use Docker as a 'containment system' for either:

- Untrusted user-submitted code, or
- Compromised applications

... but Docker explicitly [does not provide that kind of functionality](#). You get essentially the same level of security from just running things under a user account.

If you want secure isolation, either use a full virtualization technology (Xen HVM, QEMU/KVM, VMWare, ...), or a containerization/paravirtualization technology that's explicitly designed to provide secure isolation (OpenVZ, Xen PV, [unprivileged LXC](#), ...)

"Runs everywhere"

Absolutely false. Docker will not run (well) on:

- Old kernels
- OpenVZ
- Non-*nix systems (without additional virtualization that you could do yourself anyway)
- Many other containerized/paravirtualized environments
- Exotic architectures like MIPS

Docker is *just* a containerization system. It doesn't do magic. And due to environmental limitations, chances are that using Docker will actually make your application run in *less* environments.

No dependency conflicts

Sort of true, but misleading. There are *many* solutions to this, and in many cases it isn't even a realistic problem.

- **Compiled languages:** Just compile your binary statically. Same library overhead as when using Docker, less management overhead.
- **Node.js:** Completely unnecessary. Dependencies are *already* local to the project. For different Node.js versions (although you generally shouldn't need this due to LTS schedules and polyfills), [nvm](#).
- **Python:** [virtualenv](#) and [pyenv](#).
- **Ruby:** This one might actually be a valid reason to use *some* kind of containerization system. Supposedly tools like `rvm` exist but frankly I've never seen them work well. Even then, Docker is probably not the ideal option (see below).
- **External dependencies and other stuff:** Usually, isolation isn't necessary, as these applications tend to have extremely lengthy backwards compatibility, so you can just run a recent version.

If you *do* need to isolate something and the above either doesn't suffice or it doesn't integrate with your management flow well enough, you should rather look at something like [Nix/NixOS](#), which solves the dependency isolation problem in a *much* more robust and efficient way, and also [solves the problem of state](#). It does incur management overhead, like Docker would.

Magic scalability

First of all: you probably don't *need* any of this. 99.99% of projects will never have to scale beyond a single system, and all you'll be doing is adding management overhead and moving parts that can break, to solve a problem you never had to begin with.

If you *do* need to scale beyond a single system, even if that needs to be done rapidly, you probably *still* don't get a big benefit from automated orchestration. You set up each server once, and assuming you run the same OS/distro on each system, the updating process will be basically the same for every system. It'll likely take you more time to set up and manage automated orchestration, than it would to just do it manually when needed.

The only usecase where automated orchestration *really* shines, is in cases where you have high *variance* in the amount of infrastructure you need - one day you need a single server, the next day you need ten, and yet another day later it's back down to five. There are extremely few applications that fall into this category, but even if your application does - there have been

automated orchestration systems for a long time (Puppet, Chef, Ansible, ...) that don't introduce the kind of limitations or overhead that Docker does.

No need to rely on a sysadmin

False. Docker is not your system administrator, and you still need to understand what the moving parts are, and how they interact together. Docker is *just a container system*, and putting an application in a container doesn't somehow magically absolve you from having to have somebody manage your systems.

Blocking LLM scrapers on Alibaba Cloud from your nginx configuration

There are currently LLM scrapers running off many Alibaba Cloud IPs, that ignore `robots.txt` and pretend to be desktop browsers. They also generate absurd request rates, to the point of being basically a DDoS attack. One way to deal with them is to simply block *all* of Alibaba Cloud.

This will also block legitimate users of Alibaba Cloud!

Here's how you can block them:

1. Generate a deny entry list at <https://www.enjen.net/asn-blocklist/index.php?asn=45102&type=nginx>
2. Add the entries to your nginx configuration. It goes directly in the `server { ... }` block.

On NixOS

If you're using Nix or NixOS, you can keep the deny list in a separate file, which makes it easier to maintain and won't clutter up your nginx configuration as much. It would look something like this:

```
services.nginx.virtualHosts.<name>.extraConfig = ''
  ${import ./alibaba-blocklist.nix}
  # other config goes here
''
```

... where you replace `<name>` with the name of your hostname.

Dealing with a degraded btrfs array due to disk failure

Forcing a btrfs filesystem to be mounted even though some drives are missing (in a default multi-disk setup, ie. RAID0 for data but RAID1 for metadata):

```
mount -o degraded,ro /path/to/mount
```

This assumes that the mounting configuration is defined in your `/etc/fstab`, and will mount it as read-only in a degraded state. You will be able to browse the filesystem, but any file contents may have unexplained gaps and/or be corrupted. Mostly useful to figure out what data used to be on a degraded filesystem.

Never mount a degraded filesystem as read-write unless you have a very specific reason to need it, and you understand the risks. If applications are allowed to write to it, they can very easily make the data corruption worse, and reduce your chances of data recovery to zero!