Welcome from the author

This guide for Koha installation was developed by John Sterbenz, Manager of Electronic Collections Administration at the University of Michigan's Kresge Library Services. Kresge's Koha installation is self-hosted, self-installed, self-maintained, and, so far as public access is concerned, allows for users to search our electronic holdings (most of which are provided by aggregators). Staff-side, it's also used to track collection expenditures through use of select Acquisitions elements in a separate instance.

As someone with a self-developed knowledge of *nix systems, there may be better, more efficient ways to get to what these instructions aim to accomplish. This command set, with minimal revision, is what has worked for me after its original construction in January 2016 and its most recent review/revision in July 2019. I would be remiss in not mentioning a (presently out-of-date) guide written by Josh Hertel that I stumbled upon back in late 2015 which proved to be instrumental to my success after a number of failed attempts across many months.

Kresge maintains separate production and development servers that run on University-hosted virtual machines. The machines are identically configured with 4 CPUs, 8 GB of RAM, and 200 GB of storage, with Tivoli Storage Manager used to manage machine-level backups. Both machines store their database on the same server as the Koha software. Production hosts around 300,000 bibliographic records. Currently, both servers only run Koha, but I plan on beginning an investigation of installing CORAL on development later this year.

Production is currently configured to run two instances of Koha over ports 80/8081 and 8082/8083 (public/staff ports for each instance, respectively)

I welcome your thoughts, comments, and questions regarding these instructions--feel free to write me at jsterben@umich.edu.

What will I have when I'm done here?

You'll have a complete (American English) Koha system (version 19.05) that will be ready for logging in for basic, initial configuration, complete through the "Web Installer". This system:

- uses standard port 80 for the OPAC and port 8081 for the staff interface
- uses MariaDB for its SQL server
- has the SQL database on the same server as the Koha software
- does not use SSL, nor has a security certificate installed
- Is not configured to send mail from any application/program
- Has the database performance tweaks memcache and Plack enabled

A final word before you begin...

If you're looking for any sort of authoritative document for establishing a secure production-level server, continue looking. **The author makes no guarantees or assurances regarding the applicability of this document for these purposes and accepts no liability therein. Continue at your own risk.**

If you're looking to see if Koha might be right for your circumstances and are looking for an easy way to get your own Koha server up and running quickly to do some tire-kicking, with (perhaps) real-world bibliographic data and fictional data for patrons, this document is for you!

Assumptions

This document assumes you have root-level access on an Ubuntu 18.04 LTS server (Bionic Beaver) system installation with network access and appropriate ports opened for incoming/outgoing network traffic.

Regarding OS installation

I continue to use the "traditional" Ubuntu 18.04 server install image (ISO file) when needed for OS installs, not the new "Subiquity" installer file now found on the primary (ubuntu.com/download/server) page.

- What's the difference? Versions using the "Subiquity" installer include the word "live" in the ISO file name, e.g. ubuntu-18.04.3-live-server-amd64.iso
- As of this writing, the "traditional" server install image can be found on this page (cdimage.ubuntu.com/releases/18.04.3/release) under "Server install image"

Automatic updates are NOT enabled under "Configuring tasksel".

At Kresge, only "OpenSSH server" is selected on the "Software Selection" screen of the OS installer to allow for SSH terminal access from the start. Note in particular the decision not to install Apache nor a SQL server at this point via any of the other options presented here.

Why Ubuntu over Debian? Kresge runs EZProxy on an Ubuntu installation and we wanted to attempt, to the greatest extent possible, to keep the base systems similar.

A note about Amazon Lightsail

These instructions have also been proven to work with an Amazon Lightsail instance ("OS Only" blueprint) of Ubuntu 18.04 LTS with at least 2 GB RAM, 1 core, and 60 GB of disk space (the standard "$10/month" configuration as of this writing), though I fully expect that performance under anything less than the "$40/month server" (8 GB RAM, 2 cores, 160 GB of disk space) will be less than optimal for processing and working within any collection over 100,000 records.

Note that, under "Networking" for the Lightsail instance and given the ports required based on the configuration specified in this document, the following firewall change is needed:

Application:	All TCP
Protocol:	TCP
Port Range:	0-65535

I have not performed any long-term viability testing on any Lightsail instance I've created. To date, I've not performed any analysis of backup / snapshot instance creation / restoration options and while I'm fairly confident in saying automatic updates are not performed against installed software, I cannot say this for certain based on my work within this option.

A note about virtualization products (e.g. VirtualBox)

Early (2015) efforts were concentrated within VirtualBox and success was experienced within this virtual environment. I've not done a Koha install on VirtualBox since moving over to the University's VaaS in 2016 but would expect no problems with an install there.

Convention

Commands to type are shown in `Courier New` font and are preceded by the user command prompt (`$`). Complete commands and supporting text are never displayed on the same line, but results from commands or portions thereof, also in `Courier New`, may share the same line as any supporting comments.

- *`[instance_name]`* refers to a user-selected name for any given instance of Koha. Many simply choose to use `library` for this purpose. Do not include the brackets. Anecdotal evidence suggests that *`[instance_name]`* must begin with a letter.
- Generally speaking, any results from commands are not shown here unless they require examination by you. Many commands will not generate any output at all.
- On occasion given the font and font size used in this document, commands will wrap onto the next line. These do not need to be similarly broken up during entry but there should be a space between the end of the first line of the command and the start of the second.

Koha installation

Open a new session using your favorite SSH client or console option and log in (of course).

(Optional:) Confirm lack of MySQL/MariaDB and Apache on system

There are several ways to do this. Perhaps one of the more consistent in terms of syntax is the Debian package manager:

```
$ dpkg -l | grep -e mysql -e mariadb
```

```
$ dpkg -l | grep -e apache2
```

- If no matches are found against installed packages, no output is generated and another command prompt will appear.
- Additional verification/details can be determined by running these commands:

  `$ mysql --version`

  NOTE two dashes before “version” in this particular command

  - **Absence** will return a list of packages from which to install
  - **Presence** will return something like this:

    ```
    mysql      Ver 14.14 Distrib 5.7.20, for Linux...
    mysql      Ver 15.1 Distrib 10.0.34-MariaDB…
    ```

  `$ apache2 -version`

  NOTE only one dash before “version” in this particular command

  - **Absence** will return a list of packages from which to install
  - **Presence** of Apache will return something like this:

    ```
    Server version: Apache/2.4.18 (Ubuntu)
    Server built:   2019-04-03T13:34:47
    ```

Set up package sources and GPG key

```
$ echo deb http://debian.koha-community.org/koha 19.05 main | sudo tee
/etc/apt/sources.list.d/koha.list
```

As of this writing, Koha version 19.05 is the most recent stable release. `19.05` can be replaced with a different (earlier) version number (not recommended). If you'd like to work within the current development release, replace `19.05` with the word `unstable` (recommended only for the strong and those that likely have more knowledge to the point where they wouldn't require this set of instructions to perform an install.)

```
$ wget -O- http://debian.koha-community.org/koha/gpg.asc | sudo
apt-key add -
```

- Capital letter "O", not number "0".
- Don't forget the dash at the end

Upgrade existing packages and install MariaDB and Koha

```
$ sudo apt-get update
```

```
$ sudo apt-get upgrade
```

```
$ sudo apt-get install mariadb-server
```

- Thorough testing of an installation through the securing of MySQL via `sudo mysql_secure_installation` has not been performed
- While the koha-common packages call for various SQL client packages, it does **not** call for installation of a SQL server

```
$ sudo apt-get install koha-common
```

- This will install a lot of dependencies (Apache among them)

Enable some necessary Apache modules...

```
$ sudo a2enmod rewrite cgi headers proxy proxy_http
```

```
$ sudo service apache2 restart
```

...create a Koha instance...

`$ sudo koha-create --create-db [instance_name]`

- *`[instance_name]`* is a key component in Koha instance creation and is used in many places throughout the system, including the MySQL tables.
  - Example: `$ sudo koha-create --create-db library`
- `koha-create` does some Apache work itself, requiring the previous apache commands to be executed before it can be invoked.
- If installation has gone to form so far, the system will respond:
  `Koha instance is empty, no staff user created.`
  `* Starting Koha indexing daemon for [instance_name]`

...then go back to more necessary Apache configuration

`$ sudo pico /etc/apache2/ports.conf`

- Use your editor of choice as you wish. I still like `pico` and `nano` from my PINE client days of 25 years ago
- `Listen 80` should already be present in the file. Add a new line under it:
  `Listen 8081`
  Save and exit

`$ sudo a2enmod deflate`

- System may report this is already enabled.

`$ sudo a2ensite [instance_name]`

- System may report this is already enabled.

`$ sudo pico /etc/apache2/sites-enabled/[instance_name].conf`

- Change port number referenced in `<VirtualHost>` under **`# Intranet`** to `8081`
- `<VirtualHost>` under **`# OPAC`** should remain `80`, as you would expect
  Save and exit

Run a few more Koha-specific commands and restart Apache...

```
$ sudo koha-plack --enable [instance_name]

$ sudo koha-plack --start [instance_name]

$ sudo service apache2 restart
```

...then go to a browser

Call up the websites being served over ports 80 and 8081. Port 80 is likely still the Apache “It works!” site. Port 8081 will likely show the Koha Web Installer login screen.

Back in the console/SSH terminal...

Two more commands will do the trick

```
$ sudo a2dissite 000-default

$ sudo service apache2 reload
```

...then back to the browser

Calling up websites on the server over port 80 (OPAC) now should show a site indicating the Koha catalog is offline for maintenance.

Calling up the website over port 8081 should (still) show the Web Installer login screen.

Success! But what do I do from here?

This was, roughly, as far as I had gotten on January 6, 2015 (well, on the port 80 side...I had no idea about the Web Installer at the time). I had no idea how to proceed (How do I log in to this system? For that matter, how do I create a login for it?), and it wouldn’t be until ten months later when I’d get the chance to revisit all of this again that I’d find Josh Hertel’s document and get the final piece of the puzzle I needed.

Back in the console/SSH, retrieve system-established password for created Koha instance

```
$ sudo xmlstarlet sel -t -v 'yazgfs/config/pass'
/etc/koha/sites/[instance_name]/koha-conf.xml
```

- Characters surrounding `yazgfs/config/pass` are identical (single quotation marks) despite how they are rendered here
- In more recent versions, it appears that this system-established password always uses a `@` (at sign) as its final (16th) character
- This command parses the `koha-conf.xml` file and retrieves the (16-character) system-generated password created for user `koha_[instance_name]` when `koha-create` is run. This is needed to start the Koha configuration process via the web installer. Mind your l, 1, O and 0s! Continue on!

Run Koha Web Installer

Point your favorite web browser to

`<machine name or IP>:8081`

if you're not already there from the earlier steps. Login using `koha_[instance_name]` and the 16-character password retrieved from the last command.

English (en) is the only installation language available.

Follow the prompts and click through. A number of system checks are performed. Selecting many of these options will pre-populate tables and system options in your Koha instance with "typical" values, which can be customized later through the staff interface.

There are two sets of "default settings" to review during installation (MARC 21 framework, Other data). Both are on the same page of the installer. Each set has several mandatory sets pre-selected. These should not be unchecked, but you're free to decide which of the optional settings you'd like to have configured. By default, none of the optional settings are selected.

Generally, Kresge selects the following optional settings:
MARC 21
Selected matching rules… (*marc21_default_matching_rules*)
Simple MARC 21 bibliographic frameworks… (*marc21_simple_bib_frameworks*)

Other data
Some basic default authorised values… (*auth_val*)
CSV profiles (*csv_profiles*)
Coded values conforming to the … holdings statements (*marc21_holdings_coded_values*)
MARC code list for relators… (*marc21_relatorterms*)
A set of default item types (*sample_itemtypes*)
Allow access to the following serviers to search… (*sample_z3950_servers*)

- I've found it helpful to **always** select "*auth_val*" and **never** select "*sample_libraries*" from the "Other data" options
- In more recent releases, the parenthetical qualifiers do not appear on the selection screen for "Other data". Rather they appear on the confirmation screen.

You will be prompted for anything required by the installer for core functionality but not installed with any of the mandatory or optional settings. Based on Kresge Library selections earlier, this includes:

- the configuration of a library name/code
- A patron category/code for the initial (staff super)user

before providing details on the first patron to be created in the new Koha instance, which will automatically be granted “superlibrarian” privileges.

Note your Koha username here:_____

Note your Koha password here:_____

You’ll need these in a few minutes, obviously.

A circulation rule must also be defined. All defaults here can be accepted and are easily modified later.

You are now ready to log into your new Koha instance using the username and password generated during the Web Installer process

“Start using Koha”!

In a browser window, go back to

`<machine name or IP>:8081`

(You should have ended up back at a standard Koha login screen after running the Web Installer.)

Use the login/password combination created during the Web Installer process.

The `koha_[instance_name]` login/password should not work on the standard Koha login screen. **DO KEEP IT**, as it will be needed any time an update to the database schema is performed. Selected elements of the Web Installer will be triggered after a Koha software update to update existing tables/records with the new database schema.