
offline-digital-library Documentation

Release 18.04-amd64

Benjamin Bach, Philip Douglass

Sep 24, 2019

Contents

1	Contents	3
1.1	Setting up a server	3
1.2	Equipment check list when setting up a new center	6
1.3	How-to	8
1.4	Troubleshooting Q&A	12
1.5	Background	13
2	Offline resources	15
3	Indices and tables	19

This project automatically configures a server for a Digital Library. Client computers for students/users are then installed from software images on the server via PXE netboot. The server contains a number of offline educational resources, together with the full Ubuntu software repositories.

This documentation is intended for technicians working in the field to deploy and maintain such Digital Libraries.

<p>Warning: This is a WORK IN PROGRESS Ubuntu 18.04 (Bionic) version. It is not complete yet.</p>

1.1 Setting up a server

1.1.1 Server requirements

The following items are necessary for setting up a server:

- Data drive (at least 1 TB) preloaded with all of the data (if you don't have a data drive already, see these *instructions for creating one*)
- A target computer, which should be a machine with a decent amount of memory (> 4 GB), fast CPU, and stable fans
- A USB flash containing default Ubuntu Desktop (18.04 64-bit)
- A USB flash containing the Offline Digital Library [installscripts](#)
 - If upgrading a previously installed site, copy the old configuration of the centre to a directory on the flash called *installscripts/config/local/name-of-site*

1.1.2 Installation procedure

1. Install Ubuntu 18.04 64-bit on the server, choose to erase everything on the server's disk if there is already something there. Name the root user and choose a password that was agreed with the other technicians (for more discussion of security, see this *description of security policy*).
2. Shutdown the server after installation
3. Install the Data drive in the server case, connect it to a SATA port
4. Boot the system and log in
5. Insert USB flash with installation scripts
6. Make sure that the hard drive is mounted. You can use command "mount" or click the hard drive's icon in the application bar, bottom left

7. Now open a terminal, CTRL+ALT+T
8. Type `sudo bash` to get a command line with administration privileges
9. Go where the USB flash is mounted: `cd /path/to/usb-flash`, for instance “/media/fair/installscripts”.
10. Go to the “installscripts” folder, for instance, `cd offline-digital-library`. This folder should contain “install.sh”, the main program that transforms the newly installed Ubuntu desktop to an ICT centre server.
11. Run `./install.sh` and the installation begins. There will be prompts to enter configuration data at the end of the installation, but the first part will take some minutes, so you can go for tea.
12. At the end, the installation asks for a “Local overlay”, that means that specific configurations for this centre is added. It will ask about each folder that it finds in “installationscripts/config/local”. This should include the centre, that you are working on.

1.1.3 Creating a new Local overlay

1. Go to the installscripts/config/local folder:

```
bash
cd installscripts/config/local
```

2. Look at what’s there and select an existing centre as a template and copy it:

```
cp -R other-centre new-centre
```

3. Re-run the last part of the install.sh job:

```
cd ../../
./install.sh 98-local.sh
```

Changing the local configuration

If you need to change files in the local configurations, you should read the documentation for creating local configurations (*which doesn’t exist yet*). But the overall idea is this:

1. Make changes to the files
2. Re-run the last part of the install.sh job:

```
cd /media/fair/usb-flash/path/to/installscripts/
./install.sh 98-local.sh
```

Re-running the local configuration on all machines

If you have already installed all client machines and have made ammendments to the local configuration (e.g. added a new standard program, user or printer driver), you can re-run the post-installion by running this command on the server:

```
cd /media/fair/usb-flash/path/to/installscripts/
python scripts/deploy/rerun_postinstall.py
```

Warning: Only machines that are switched on and correctly attached to the network will be affected.

1.1.4 After installing

Make sure that you have copied the contents of the installscripts to “/root/” so another technician can know what was used for the installation:

```
cd /media/fair/usb-flash/path/to/installscripts/  
cp -R . /root/
```

Make sure to also copy the configuration overlay, for instance:

```
cd /media/fair/usb-flash/path/to/installscripts/config/local  
cp -R . /root/
```

1.1.5 Installing clients (student computers)

Make sure that network booting with PXE is enabled in the BIOS. When you have successfully booted from the network, you will be quickly prompted to choose what operating system to boot. Select the `Automatic Ubuntu install` option which will install everything without asking questions.

After the installation is completed, the client will reboot **but it's not yet finished!** The first time the client boots after installation, it will complete a *post-install* procedure, which is a custom script that will install and configure additional software.

This process takes time, so even if you do not see any progress on the screen for up to 20 minutes, just let the script finish.

Machines that cannot boot from network

Occasionally, you may encounter a computer that is unable to boot with PXE boot. In this case you will need to create a bootable CD / USB flash with gPXE. Booting from this media will substitute for the missing/damaged on-board LAN PXE client. You can obtain the .iso file from:

<https://rom-o-matic.eu/>

- Copy the .iso image to a USB flash, this will overwrite everything of the USB:
- Insert flash, it may be automatically mounted so should be unmounted if it's mounted.
 - To see what is mount: `mount`
 - To unmount: `umount /dev/sdX1`
 - To copy: `sudo dd bs=4M if=Downloads/gpxe.iso of=/dev/sdX`
- Now the flash is bootable

Otherwise, burn the .iso image to a CDR, but remember that you should burn an *image* not a file.

1.1.6 Verifying an installed client

It's important to check that the client has installed everything correctly, because an interruption during the unattended installation will sometimes cause computers to be in a bad state.

Verify for instance:

- When logging in as a student, you receive the Gnome Classic desktop interface, which has the Applications + Places menu in the lower left corner. If you are instead met with the default Unity desktop (with icons vertically along the left side), log back out and select Gnome Fallback (Metacity) instead from the menu next to the username selection.
- That the `student` and `teacher` account are setup correctly, and that there is no password required for the `student` account.
- KA Lite videos are playing correctly.
- That the folder with *Materials For Students* contains files from the server and teachers can write to this folder on the server.

1.2 Equipment check list when setting up a new center

The following is a list of things to remember before embarking on a trip to create a new Digital Library.

Note: An accurate database of the location of equipment facilitates long-term maintenance. Therefore, never let equipment pass through the doors of the warehouse without writing down its ID number and noting the equipments' new location! FAIR denmark is working on new warehouse data management system to make tracking of equipment easy.

1.2.1 Tools

- Network cable crimpers
- Screw drivers, hammer, pliers
- LAN testers

1.2.2 Accessories

Make sure that you are familiar with the layout and facilities of the room in advance: how many electrical outlets are present, the location of the desks, etc. You need to know approximately how many metres of network cable to bring, and how many power extension cables. Sketch a floorplan, with the location of the computers, switches, server, printer.

- **LAN cables (both pre-made, and unfinished rolls of raw cable for DIY)**
 - 1 for each computer
 - 1 for each connection between switches
 - 1 for the server
 - 1 for printer
- **Ethernet plugs (RJ-45) for DIY cables**
 - 2 for each cable
 - + 10-20% extra for mistakes
- Power cables for desktop computers and screens
- Power supplies for laptops

- VGA cables
- Power extension cables, if the electrical outlets are far from the computers
- Projector + long VGA cable and power cable
- Switches with enough ports for computers + server + printer + connection between switches
- Cable clips
- Lots of plastic strips to fixate loose cables
- Labels and pen
- Good mood!

Optional

- Speakers

1.2.3 Spare parts

Hint: It's best to setup and test the server before going. In case you are unable to test the server, make sure to bring an extra server or server data hard drive.

Warning: Computers can be broken in transport. Typically, 1 or 2 machines may be dead on arrival. You should bring one or two extra computers.

Note: Remember to update the inventory! When you are bringing spare parts out, and when you are bringing unused spare parts or broken computers back.

Not everything may go as planned, and its lighter to carry spare parts than spare computers.

- RAM (a variety of types)
- Extra hard drives SATA+IDE
- Graphics cards (in case Nvidia / ATI drivers are buggy)
- Graphics adapters to convert DVI to VGA, and a few DVI cables. Occasionally, a computer won't have a working VGA interface, and DVI can be used instead.

1.2.4 Computers

- Teacher computer (+ screen & keyboard if not using a laptop)
- Server (+ screen & keyboard or KVM (keyborad-video-mouse) switch)
- Clients (+ screens if not laptop). Bring a couple of extras in case you experience problems.
- Keyboards + Mice

Hint: Most mice use USB plugs, and are compatible with all computers. But some mice use round PS-2 plugs, and are only compatible with computers with PS-2 ports. Make sure you know what kind of mice you are bringing, and if they are PS-2 mice, ensure you are bringing compatible computers.

1.3 How-to

1.3.1 Sharing an internet connection

1.3.2 Creating a local overlay

Coming soon!

1.3.3 Cloning hard drives

In order to obtain a data drive for the server, clone a **master** (or **source**) hard drive onto a **target**. There is nothing special about the master, just that it's a hard drive that you keep safe such that other hard drives can be cloned from it.

Tip: Keep your hard drives labeled with a version (or date). Confusion can be quite costly timewise.

Warning: Never deploy your last master copy, it is recommendable to have 2 master copies.

Tip: Before using the a disk copy in a school, consider setting up the server at your own location and testing it there.

As of 2018, the data footprint is roughly 550 GB.

Cloning to same capacity disk

If - say - you have a *source* disk with a 1 TB capacity and you are cloning to a *target* that is also a 1 TB drive, you should use a SATA docking station. This is the much faster option.

In order to optimize the chances for correct copies, make sure to monitor the copying process in case of electrical blackouts. You may also use a UPS to protect against both blackouts and power surges. If you cannot get a UPS, a power surge protector is also helpful, in cases where the SATA docking station is affected by power surges.

Copying to different capacity disk

A step-by-step guide is provided below, but firstly some background...

Cloning to a *larger* target: If - say - you have a *source* disk with a 1 TB capacity and you are cloning to a *target* that is 2 TB, you can clone with a SATA docking station, but you will *not* be able to use the disk at its 2 TB capacity afterwards, it will report as a 1 TB file system (the size of the *source* file system).

Cloning to a *smaller* target: This is **not** possible with a SATA docking station.

When you have different capacity `source` and `target`, the best method is to copy file-by-file using mounted filesystems. You do this by connecting both drives with a USB dock. Use USB 3.0 if available, it is much faster (the USB plug in your computer is usually colored blue for 3.0 ports). You can also connect both drives to the motherboard using a SATA cable.

When both drives are connected, you may have to format the destination drive. It should be labeled `FAIR` and the file system must be `ext4`.

After both drives have been mounted, observe that `source` and `target` are mounted as `FAIR` and `FAIR1`, but which is which? You have to find out manually. The file systems of both drives have identical labels `FAIR`, but the operating system makes sure to mount them in separate locations.

Warning: Before you copy, make sure that you know which is the source and which is the target. Be aware that they can swap if you unplug drives or reboot the computer.

Glossary

- **Drive:** When we just say “drive”, we mean the whole drive and not something specific.
- **Partition table:** A drive will always have one single partition table, which says where the partitions are located.
- **Partition** A section of a hard drive. If you use Windows, you can recall it has **drive letters** such as `C:\`. This actually denotes a partition. On Unix and Linux, partitions can be found anywhere in a different tree structure.
- **Mounting:** A partition is *mounted* on to the root file system, meaning that for instance, you can have a path called `/media/username/drive-label` and it’s our way of saying that when you read or write files anywhere in that hierarchy.
- **Formatting:** The action of formatting something means to create a new file system on a partition.
- **File system:** A file system is a certain way of organizing files and folders on a partition. Different operating systems use different file systems. On Linux, the most normal system is called `Ext4`.
- **Device node:** When a file system isn’t mounted, it is often referred to by its *device node*, such as `/dev/sda1`. This is a special path where you cannot put files or folders. But you can use it to refer to a partition or drive. The disk itself doesn’t have a number and appears like `/dev/sda`, and each partition is referred to by a number, for instance `/dev/sda3`.

Step-by-step

1. Start by plugging in the `source` drive. It will be available in the location `/media/<username>/FAIR`. You can verify by running this from command line:

```
ls /media/<username>/FAIR # Replace <username> with your username.
```

This should display a list of files, such that you can verify that the `source` drive is indeed *mounted* here. If you want to see which drives are currently mounted, and where, run:

```
lsblk # Lists all devices and partitions and where they're mounted
mount # This lists everything mounted but is a bit harder to read
```

2. You need to format a new partition on the `target` drive. Use the **Disks** program that is installed by default in Ubuntu. The partition should be labeled `FAIR` and the file system must be `Ext4`. If there is something already on the drive, the easiest option is to use the *Format Disk* option, which will generate a new *partition table*, on to which you can add a partition using the entire capacity of the drive.

3. After formatting the `target` partition, make sure that it is mounted. You can also use the **Disks** program to mount the partition. Use the `lsblk` command to see which drives and partitions are mounted.
4. When both `source` and `target` drives are mounted and ready, you must use the `rsync` command to copy files. The program makes it possible to resume in case you are interrupted. This is crucial, because depending on the USB or SATA connection, the copying process can take as much as 12 hours, and you may experience power cuts.

The command is as follows, assuming that `source` is mounted in `/media/<username>/FAIR!!`

```
# This command replaces $USER with your username automatically.
rsync -av --stats /media/$USER/FAIR /media/$USER/FAIR1
```

1.3.4 Setting up a Wireless center

Warning: It is not recommended to do Wireless centers, as the installation and booting from networks have been proven very difficult.

Prerequisites

- A wireless access point (AP) with bridging enabled, i.e. not acting as a router
- Server connected to AP.
- Server configured with a local overlay for wireless setups.
- `mini.iso` prepared on a USB for booting clients before installing
 - Obtain `mini.iso`: <http://archive.ubuntu.com/ubuntu/dists/trusty/main/installer-i386/current/images/netboot/>
- Copy the `.iso` image to a USB flash, this will overwrite everything of the USB:
 - Insert flash, it may be automatically mounted so should be unmounted if it's mounted.
 - * To see what is mount: `mount`
 - * To unmount: `umount /dev/sdX1`
 - * To copy: `sudo dd bs=4M if=Downloads/mini.iso of=/dev/sdb`
 - Now the flash is bootable

Installing a client

- Turn on the machine, press for instance F12 and choose to boot from USB
- Press TAB to edit the first “Install” option for Ubuntu
- Delete the `quiet` part.
- Put: `ks=http://192.168.10.1/ks.cfg ksdevice=WLAN0`
- When installing, you should be able to choose the wireless network that you have configured

Configuring DIR-635 access points

1. Reset the device
2. Attach to a machine and obtain DHCP from the AP
3. Connect to 192.168.0.1
4. Setup an un-encrypted wireless, DO NOT REBOOT YET
5. Go to Network
 1. Disable DHCP
 2. Disable DNS relay
 3. Give the router a static IP, i.e. 192.168.10.2 (must be unique to your network!)
6. Go to Advanced and disable features you know are useless.
7. Reboot device
8. Connect to server on one of the Switch ports, not the internet/WAN port

Tip: Use two access points and configure them on separate channels.

Booting from the network

Once the server is installed, and the network is ready, you need to boot from the network.

However, as computers do not support TFTP via WIFI, you need to:

1. Boot from a USB flash
2. Tell the Ubuntu installation program (debinstall) where to find the configuration for automatic installations, and to configure the WIFI in order to retrieve this list.

Creating the bootable USB flash

Retrieve the mini.iso for installing Ubuntu, one which is from the same date as the data drive, but never newer!

- 14.04: <http://archive.ubuntu.com/ubuntu/dists/trusty/main/installer-i386/current/images/netboot/>

Warning: The USB flash you use will be overwritten!

Once downloaded to your Ubuntu/Linux computer:

Run the command `df -h` and notice which device name corresponds to your USB. If you get this wrong, you can overwrite your hard drive. The device name is `/dev/sdX`, where X is some letter. `/dev/sda` is your hard drive.

Make sure that no programs are using or blocking the USB flash. Run the command `mount` to see what drives are in use. If you see your USB flash's device name, run the command `umount /dev/sdX1` to *unmount* it.

When you know the device name of the USB, run the following command:

```
dd if=/path/to/iso of=/dev/sdX bs=4096
```

Booting the flash

As you are booting from the flash, you have to edit the default entry by adding the following lines:

```
ks=http://192.168.10.1/ks.cfg ksdevice=wlan0 netcfg/wireless_essid=FAIR
```

Tip: Do not start too many machines, because the wireless network easily gets congested. 5-10 machines is often the limit.

1.4 Troubleshooting Q&A

1.4.1 File system errors on boot

If either the server or student computers are complaining during boot:

- If the BIOS battery is dead, the clock of the computer will reset on reboot. This causes the filesystem to produce errors because timestamps of files are in the future. Resolution: Replace the BIOS battery on the motherboard.
- Electrical blackouts may cause unclean shutdowns of file systems. Select the option to fix the hard drive.

If you want information about a file system error, go to maintenance mode and run this command:

```
# Check all filesystems from /etc/fstab in verbose mode
sudo fsck -A -V
```

1.4.2 Network problems

Symptom: Clients are not able to connect or boot because of network issues. This typically means that while booting, the computer hangs for several minutes with the Ubuntu logo.

If the server is up, and you know that it is accessible directly from its screen and keyboard, then problems are likely caused by network cables or switches. If all computers are unable to connect, try these options:

- Make sure network cables are properly plugged in.
- Check that the switches are on.
- Use a LAN tester to check cables, especially the one from the server.
- Check that cables are not causing loopbacks, for instance if a network cable is connected to a switch at both ends.

1.4.3 Blackouts

While installing the server

The installation can be safely resumed and will pick up from where it left. Just boot the default Ubuntu and re-locate and start `install.sh`, following the normal instructions.

While installing a client

If the installation is disturbed in the first phase, installing the base system (before rebooting the client), then you need to restart “Automatic Ubuntu install”.

If the installation is disturbed after the first reboot, you can do this to save time:

```
# Press CTRL+ALT+F1
# Log in as teacher
# In the command line after logging in, type...
$ sudo bash
# cd /root
# bash rerun-postinstall.sh
```

1.5 Background

The primary goal of the project is to distribute open source software and open access knowledge and media repositories to remote areas without relying on internet connectivity. More universally speaking, the project attempts to do ICT4D, but without any broadband internet connection in the actual deployments of the project.

Secondary goals, are efficient deployment of new centres, and maintainance of existing ones. With the software system, the main deployment effort is physically setting up the ICT centres, the software should run immediately “out of the box”.

The project provides the scripts necessary to configure an ordinary desktop installation of Ubuntu 18.04 to become a network file server, TFTP server, and HTTP intranet server for a variety of offline resources (hereunder the full Ubuntu repositories!).

In order to use the scripts, a (possibly external) hard drive must be created with a number of required and optional resources.

1.5.1 Choice of Linux OS

The project is based on Ubuntu 18.04 LTS. A basic policy is to use Ubuntu’s LTS (Long-Term-Support) after it has been out for at least 6 months because it is the most well-tested, feature-rich, and beginner-friendly distribution available.

1.5.2 Currently

The distribution of Ubuntu 18.04 supports **only 64-bit servers**. Client computers can be either 64-bit or 32-bit.

Warning: On our current plans, 18.04 is the final version that will support 32-bit client computers.

Development effort has been put into modularizing the setup and facilitating customization, so that the project will be usefull beyond FAIR’s deployments.

We are adding functionalities and documentation to support the workflow of technicians working in the field.

1.5.3 Roadmap

The next upgrade will:

- Deprecate 32-bit computers
- <Your ideas here!>

1.5.4 Discussion

We do aim a lot at providing backwards compatible setups. Experiences say that school computers are better off with less machine intensive Linux desktops, and that newer versions do not always provide better performance. For instance, we see that Ubuntu 10.04 is actually providing better responsiveness on older hardware (256 MB, <2 GHz CPU).

As much as Ubuntu, the Linux kernel, X drivers etc. are adding support for new hardware, older hardware is becoming less supported as the test audience shrinks.

Another alternative would be to add installation configurations that gave the choice of Lubuntu or Xubuntu desktops.

1.5.5 Philosophy

Everything in here is implemented in flat-out BASH scripts. It's not as if we didn't feel like using a different language, however to a new system administrator wanting to know what's going on, the choice of a scripting language identical to the command line will make the process of server installation and configuration more transparent.

1.5.6 Security Policy

Linux is a highly secure operating system, which by default restricts permission to change files to a minimum. But software system is only as secure as the people using it. Therefore, your crew of technicians needs to agree on common policies to make sure access to changing the system is only given to trusted parties, who have proper instructions.

Students using client computers are able to save files to their student directory, but otherwise are unable to make persistent changes to the setup. Even this limited permission to create files can be problematic because it can quickly clutter up the desktop. Therefore, to discourage this, tools are provided to clean the student directory, and restore the default setup.

Teachers are given root access to the client computers, enabling them, for example, to install new programs. Teachers are also given user accounts on the server. Teacher are allowed to create, curate, and customize content of the intranet, after being introduced to the content management system. Schools typically do not need root access to the server. This ensures that critical files on the server cannot be corrupted.

Technicians from your team will have root access to the server to perform updates.

This project captures the efforts of FAIR in Malawi. As the number of ICT centres steadily grows, the project adds more efficiency and is improving its architecture to accommodate development of new features and future management of individual centre setups.

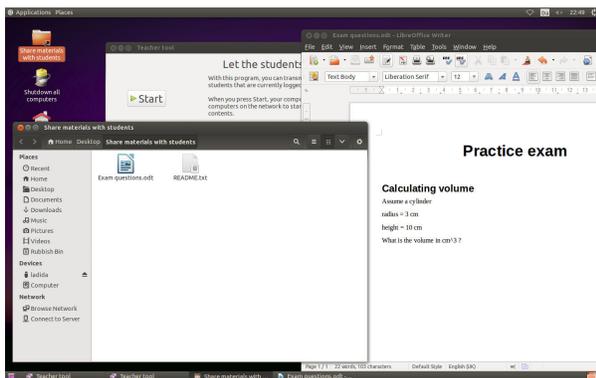
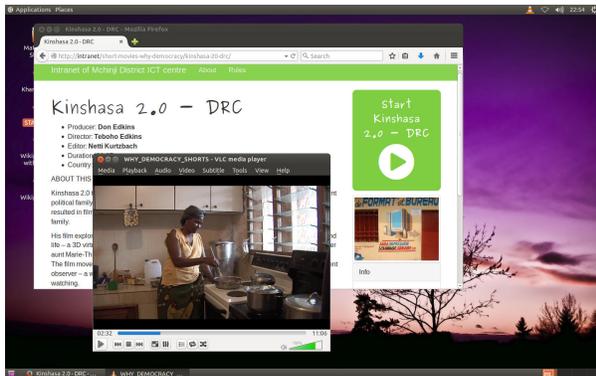
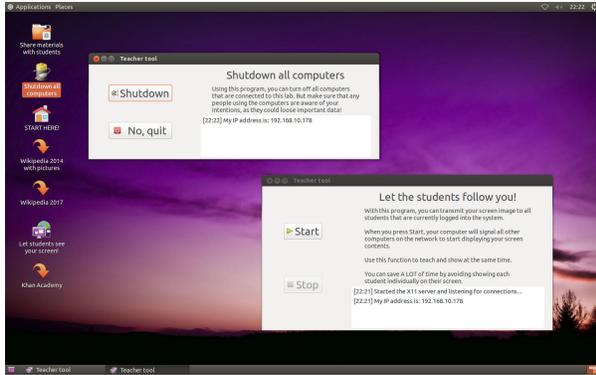
Basically, these scripts deliver a fully-automated configuration of a default Ubuntu desktop to transform it into a server for the ICT centres, provided that it has access to a *data drive* containing a copy of the Ubuntu repositories. The server provides a fully functional Ubuntu repository with thousands of programs, combined with TFTP services for installing client desktops over a LAN.

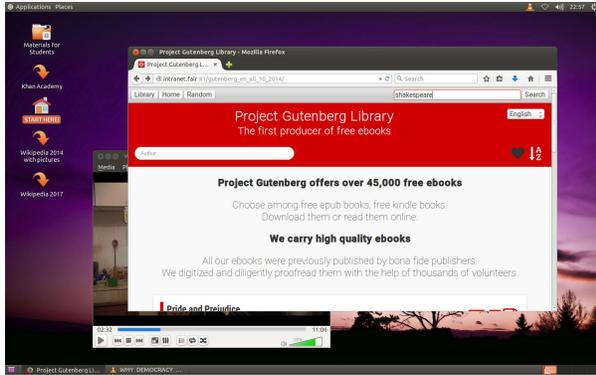
Offline resources

In addition to LAN-managed installations of Ubuntu, the server can also provide a range of contemporary offline projects, such as Kolibri and Kiwix (including Wikipedia), together with a *Intranet* service that collects various educational resources for access via HTTP over the local LAN.

The screenshot shows the homepage of the 'Intranet of Mohinji District ICT centre'. The page has a green header with navigation links for 'Home', 'About', and 'Index'. Below the header, a welcome message reads: 'Welcome to your school's local Intranet. Below is a list of various collections of resources, some of them quite vast. Click each of them to read more about their contents before clicking on to explore.' The main content area is a grid of 24 resource tiles, each with a title, a brief description, and a small icon. The tiles are arranged in 8 rows and 3 columns:

- Wikipedia**: World's largest encyclopaedia right here on your pc. (Icon: globe)
- Khan Academy**: Learn through videos and exercises on the computer. (Icon: leaf)
- TED Talks**: World's leading thinkers sharing their latest ideas. (Icon: TED logo)
- Camara training**: ICT training basic to advanced. (Icon: computer monitor)
- Getting started With Ubuntu**: Learn everything about the new Ubuntu 14.04. (Icon: Ubuntu logo)
- NICE**: Civic education reports, articles, and lists-just at your hand. (Icon: NICE logo)
- OER4Schools**: Teacher training materials. (Icon: graduation cap)
- Project Gutenberg**: 43,000 high quality eBooks. (Icon: book)
- Short movies - Why Democracy?**: Watch, learn, discuss these global short movies about democracy. (Icon: 'WHY DEMOCRACY?' logo)
- Short movies - Why Poverty?**: Watch, learn, discuss these global short movies about poverty. (Icon: 'WHY POVERTY?' logo)
- Why Democracy?**: Watch documentary movies about democracy from around the world. (Icon: 'WHY DEMOCRACY?' logo)
- Why Poverty?**: Watch and discuss documentary movies about the causes of poverty. (Icon: 'WHY POVERTY?' logo)
- Wikibooks**: Textbook collection of over 2,400 books with 33,000 pages. (Icon: book)
- Wikiquote**: Over 24,303 famous quotes from around the world. (Icon: speech bubble)
- Wikisource**: A free collection of 300,000 original textbooks. (Icon: book)
- Wikispecies**: Animals, Plants, Fungi, Bacteria, Protists, and all other life. (Icon: globe)
- Wikivoyage**: The FREE worldwide travel guide -where would you like to Travel? (Icon: globe)
- Wiktionary**: 1,000,000 entries with English definitions from 100 languages. (Icon: book)





CHAPTER 3

Indices and tables

- `genindex`
- `modindex`
- `search`