Ubuntu Linux
(and bash)

Unix family tree

- Unix (ca 1970)
  - Berkeley (BSD)
    - Next -> MacOS
- GNU (GNU’s Not Unix, 1983)
- Linux (GNU/Linux, 1990)
  - RedHat
  - Debian
    - Ubuntu
      - Mint
- POSIX

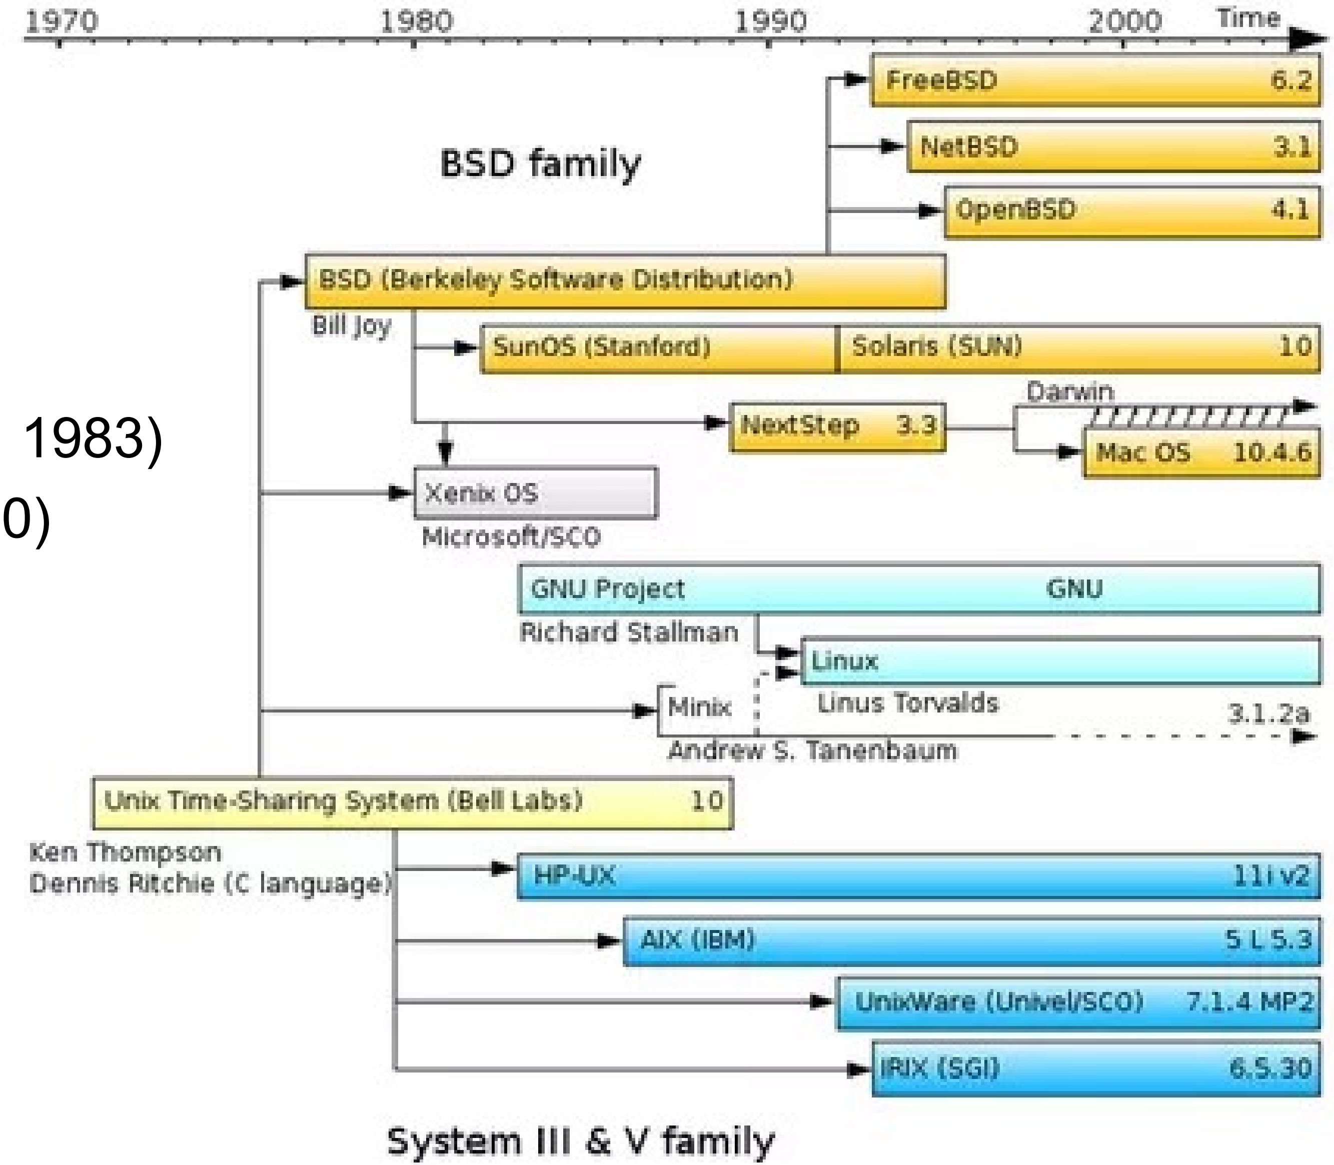

GNU/Linux and Ubuntu

- GNU (GNU’s Not Unix, 1983)
  - Richard Stallmann
  - developing a free Unix
  - full set of utility programs (the «commands»)
  - no kernel finished
- Linux (GNU/Linux, 1990)
  - Linus Torvalds
  - a Unix-like kernel

- Ubuntu (2004)
  - a Linux distribution
  - mostly free and open source
  - a package repository
  - reference OS for OpenStack
  - choice of desktop managers
  - downstream from Debian
  - upstream to Mint
  - versions for server, desktop, IoT
  - many flavours

Accessing an Ubuntu server

- Use a good *terminal window («console»)*
  - most native Linux/Unix terminals are ok
  - on Windows, do *not* use the default (cmd.exe)
    - e.g., install Windows Terminal
- Learn the basic *keyboard shortcuts* (some of these are in your *Linux shell*)
  - cut-and-paste: Ctrl-C [Ctrl-Shift-C], Ctrl-V [Ctrl-Shift-V]
  - command history: Arrow-Up, Arrow-Down,
  - command-line editing
  - scroll up/down
  - etc.
- Avoid remote desktop managers
  - less effective in the long run *& do not teach you scripting*

Ubuntu server first login

- Default set up:
  - all incoming ports closed (managed by OpenStack)
  - default user: ubuntu
  - group: ubuntu
  - home folder: /home/ubuntu (~)
  - sudo to execute commands as root user
  - ~/.bash_profile or ~/.profile
    - script executed when user logs in
    - typically executes ~/.bashrc too

- Default shell:
  - runs in the terminal window
  - bash is default
  - provides
    - built-ins and commands
    - environment variables
    - basic scripting language
    - for- and other loops
    - expressions, tests
  - ~/.bashrc
    - script executed when a non-login
    - shell starts

Command syntax

- Built-in manual:
  - man *command_name*
- The shell runs each command in its own child process (shell)
  - with its own inherited environment variables
  - but *usually* not for builtins
- Output-input pipes:
  - cat file | sort | uniq -c
  - cat file | sort | tee file

    # also tee -a

- Input and output redirections
  - command < input > output
  - concatenated output: >>
  - error output: &>

-
  ```
  command \
     over \
     several \
     lines
  ```

Command syntax

- Foreground and background
  - run in foreground: cmd
  - run in background: cmd &
    [1] 12345
  - back to foreground: %1
  - pause foreground: Ctrl-Z
  - [2] 23456
  - unpause to foreground: %1
  - unpause to background: %1&

- Run command inside command:
  - $(...) (or `...`)

Environment variables

- Local to the shell
  - but can be set «globally» in ~/.bash_profile, ~/.profile, or ~/.bashrc
  - recursively inherited by sub-processes spawned (forked) by the shell
- Examples: USER, HOME, PROMPT, PATH, JAVA_HOME...
  - echo $HOME
- To set temporarily (in shell):
  - export JAVA_HOME=/usr/lib/jvm/java-8-openjdk-amd64
  - export PATH=${PATH}:${JAVA_HOME}/bin
- To set permanently (all bash shells):
  - cp ~/.bashrc ~/.bashrc-bkp
  - echo "export JAVA_HOME=/path/to/a/java-version" >> ~/.bashrc
  - echo "export PATH=\${PATH}:\${JAVA_HOME}/bin" >> ~/.bashrc

Environment variables: the PATH

- PATH: where to look for executable command files
  - echo $PATH
  - e.g., /usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin
- Other program folders:
  - /opt
  - /home/ubuntu/bin (/home/USER_NAME/bin)
  - /home/ubuntu/sbin (/home/USER_NAME/sbin)
- Other useful folders:
  - /dev
  - /etc
  - /var/logs

Package management

- apt, e.g.,
  - sudo apt-get upgrade
  - sudo apt-get update
  - sudo apt-get autoremove
  - sudo apt-get install package-name
  - sudo apt-get reinstall package-name
  - sudo apt-get remove package-name
  - sudo apt-get list ...
  - sudo apt-get search ...
  - sudo reboot now
- Also:
  - dpkg, alien, ...

- sources.list
  - files in /etc/apt manage
    - trusted package sources (URLs)
    - their public GPG keys
- GnuPG (GPG)
  - implementation of the OpenPGP standard
- PGP (Pretty Good Privacy)
  - software and model for authentication and privacy

Common commands/builtins

- Files and folders:
  - ls [-ldtR]
  - ln [-fns]
  - mkdir [-pm]
  - cd
  - rm [-rf]
  - rmdir [-f]
  - touch
  - chmod [-R]
  - chown [-R]
- File contents:
  - cat
  - grep
  - head
  - tail
  - cut
  - emacs or nano
  - sed [-i]
- Manual:
  - man

Common commands/builtins

- Network:
  - wget
  - ip address
  - ping
- Environment;
  - echo $ENV_VAR
  - envsubst < input > output
- Other:
  - date
  - tee
  - tr
  - paste [-sd]

- Archives:
  - tar [xzvf] (to extract from file)
  - tar [czvf] (to compress to file)
- Examples:
-
  ```
  $ wget https://dlcdn.apache.org/hadoop/common/hadoop-3.3.4/hadoop-3.3.4.tar.gz
  $ tar xzvf hadoop-3.3.4.tar.gz
  $ ln -fns hadoop-3.3.4 hadoop
  ```
-
  ```
  $ date --iso-8601=minutes --utc
  ```
-
  ```
  $ cp ~/.bashrc ~/.bashrc-bkp-$(date --iso-8601=minutes --utc)
  ```

Bash loops

- Syntax:

```
for var in el1 el2 ... eln ; do
    cmd1 $var ;
    cmd2 $var ;
done
```

- Nested list elements:

```
for name in hund ka{tt,mel{,eon}} ;
do
    echo $name ;
done
```

- Bash expressions:

```
i=1;
for word in abra ka dabdra ; do
    echo "$i $word" ;
    ((i++)) ;
done
```

SSH and SCP

- Private and public keys
  - ssh-keygen -b 4096 -f info319-spark-cluster # leave password empty
  - move keys to ~/.ssh
  - check modes: 700 for folder, 600 for files
- ~/.ssh/config:

```
Host spark-driver
    Hostname 2001:700:2:8301::1111  # the new IPv6 address here
    User ubuntu
    IdentityFile ~/.ssh/info319-spark-cluster
    ProxyJump YOUR_USERNAME@login.uib.no
    StrictHostKeyChecking no
    UserKnownHostsFile=/dev/null
```

- scp local_file remote_host: # many variants

More SSH configuration

- "~/.ssh/config", wildcard notation:

```
Host spark-*
    User ubuntu
    Port 22
    IdentityFile ~/.ssh/info319-spark-cluster
    ProxyJump USER@login.uib.no
    StrictHostKeyChecking no
    UserKnownHostsFile=/dev/null

Host spark-driver
    Hostname 2001:700:2:8301::1111

Host spark-worker-1
    Hostname 2001:700:2:8301::2222
```

- Multiplexing sockets:

```
$ mkdir -m 0700 ~/.ssh/controlmasters
```

- And add lines like these to "~/.ssh/config":

```
Host login.uib.no
    User YOUR_USERNAME
    ControlPath ~/.ssh//controlmasters/%r@%h:%p
    ControlMaster auto
    ControlPersist 10m
```

- Alternatives:
  - /{dev,var}/shm/controlmasters
  - temporary, so you need to recreate them in the login file, e.g.:

    ```
    $ echo "mkdir -p -m /dev/shm/controlmasters" >> ~/.profile
    ```

Devices and disks

- New empty volumes must be partitioned, formatted and mounted:
  - ls /dev/*   # list attached devices
  - mounted   # list mounted devices
  - sudo gdisk /dev/sdb   # partition (sdb is an example)
  - sudo mkfs.ext4 /dev/sdb1   # format (sdb1 is an example)
  - mount temporarily
    - sudo mount /dev/sdb1 volume
    - sudo chown -R ubuntu:ubuntu volume
  - mount permanently
    - blkid   # to find device id
    - cat >> /etc/fstab

      ```
      UUID="d0fb-...-1552" /home/ubuntu/volume ext4 defaults 0 0
      Ctrl-D
      ```