

# Account Primer

## Faculty of Mathematics, TU Dortmund

Mathematikrechner-Betriebsgruppe (MRB)  
mrb@mathematik.tu-dortmund.de

August 15, 2018

This document is intended to provide a kick start for anyone who recently acquired an account for the heterogeneous computer system of the Faculty of Mathematics at TU Dortmund University. It briefly describes the main services available at the faculty and how to connect to its network.

If you have never worked with Unix based operating systems before, i.e. you have no experience with either one of Linux, Solaris, Irix, HP UX, Mac OS X and/or you have never used with a command prompt, you will most likely want to follow a short introductory course on Unix in general to get at least a basic grasp of its concepts and possibilities. You will find many books on the topic as well as online resources. We recommend, e.g.:

- The Unix command line - beginner's and advanced level, <http://www.mathematik.uni-dortmund.de/~buijssen/unix-command-line/unix-command-line.html>
- UNIX Tutorial for Beginners, <http://www.ee.surrey.ac.uk/Teaching/Unix/index.html>
- Introduction to UNIX and Linux: Tutorial lectures and exercise sheets, <http://www.doc.ic.ac.uk/~wjk/UnixIntro/index.html>
- Linux-Einführung, presentation by Compute Centre Hannover (in German only), <http://www.rrzn.uni-hannover.de/fileadmin/kurse/material/UnixGK/UnixGK.pdf><sup>1</sup>
- The Linux Command Line, <http://linuxcommand.org/tlcl.php>

---

<sup>1</sup>Additionally available: Linux course exercise plus solutions, <http://www.rrzn.uni-hannover.de/fileadmin/kurse/material/UnixGK/Uebung.pdf> and <http://www.rrzn.uni-hannover.de/fileadmin/kurse/material/UnixGK/Loesungen.pdf>.

**Contents**

<b>1</b>	<b>IGEL thin-clients / Linux terminal service</b>	<b>3</b>
1.1	USB pen drives (new IGEL firmware) . . . . .	4
1.2	USB pen drives (old IGEL firmware) . . . . .	5
<b>2</b>	<b>Mail server</b>	<b>5</b>
2.1	Signature . . . . .	6
2.2	Spam . . . . .	7
2.3	Limitations imposed by mail server . . . . .	7
2.4	Hints . . . . .	7
<b>3</b>	<b>Home and data directory</b>	<b>8</b>
3.1	Backup . . . . .	9
<b>4</b>	<b>Printing</b>	<b>9</b>
<b>5</b>	<b>Software</b>	<b>9</b>
5.1	Editors . . . . .	13
5.2	Browsers . . . . .	14
5.3	E-mail client . . . . .	14
5.4	PDF viewer . . . . .	14
5.5	OpenGL software . . . . .	15
<b>6</b>	<b>Windows server</b>	<b>15</b>
<b>7</b>	<b>Web server</b>	<b>15</b>
<b>8</b>	<b>Compute servers</b>	<b>15</b>
8.1	Available compiler . . . . .	16
8.2	OpenMPI . . . . .	16
8.3	Workload manager Slurm . . . . .	17
8.3.1	Usage of slurm . . . . .	19
<b>9</b>	<b>SSH, SFTP and FTP access</b>	<b>22</b>
9.1	Chain logins . . . . .	23
9.2	Remote desktop . . . . .	23
<b>10</b>	<b>Changing user password and mail forwarding</b>	<b>24</b>

## 1 IGEL thin-clients / Linux terminal service

Typically, professors, research assistants and student workers get an “IGEL UD5 Universal Desktop Thin Client” for accessing the computer systems of the Faculty of Mathematics at TU Dortmund, Germany. An IGEL terminal is a stateless thin-client, i.e. it features a display, keyboard and mouse and USB ports<sup>2</sup> running a small Linux environment. Using the ThinLinc client software<sup>3</sup>, a user connects to one of the Linux terminal servers where the actual desktop session runs. Since summer 2014, these are one of two Linux terminal servers running CentOS 6 named

- `greenhouse1`
- `greenhouse2`

They offer Gnome<sup>4</sup> as window manager.

Relying on thin clients as standard workstations may seem uncomfortable at first, but it does provide several advantages:

1. One can re-connect to a previously started and later suspended ThinLinc session running on the Linux terminal servers from any device that runs the ThinLinc client software. This can be the very same IGEL thin-client terminal used to start the session, but also an IGEL thin-client terminal in another office and even a private notebook/laptop or personal computer at home. This way it is possible to pick up the desktop session at home exactly where you left it in your office and vice versa.
2. If an IGEL thin-client terminal malfunctions, it is very easy to replace it. There is no need to re-install software or transfer data from the old device to the new one.
3. In case of power fluctuations, the IGEL thin-client terminal may reboot, but the desktop session running on the Linux terminal servers typically keeps running due to their having a more robust power supply. This helps to minimise events of inadvertently interrupting long-running calculations due to power malfunctions.
4. Total hardware cost of IGEL thin-client terminals plus the few central Linux terminal and file servers is lower and their administration requires less manpower than acquiring and administrating approximately 100 individual workstations which typically have – given that experience has shown that they are not bought at the same time, but when needed – varying hardware and hardware defects.

A user account is by default valid on all terminal and compute servers (get an overview over all available servers on the server load web page <http://www.mathematik.tu-dortmund.de/~mrb/serverload.html>) and allows log in to all of them except

- the mail server `smtpl` (also known as `hermes`). Users are only able to access their e-mails via the (SSL-encrypted) IMAPS protocol (port 993) and send e-mails via the (SSL-encrypted) SMTPS protocol (port 587). For security reasons, they are not permitted shell access.<sup>5</sup>
- the web server `euler`. Access is granted on request, see section 7.
- the Microsoft Windows terminal servers. Access is granted on request, see section 6..

---

<sup>2</sup>Please note that the USB ports of an IGEL thin-client terminal do not provide enough power to operate external hard drives, CD/DVD/BluRay burners, scanners etc. USB sticks (also known as pendrives), keyboards, mice, smartphones should work, though - if and only if the ThinLinc session was started from an IGEL terminal!

<sup>3</sup>The client software is freely available from <http://www.cendio.com/ThinLinc/download> for Windows, Mac OS X and Linux.

<sup>4</sup>See <https://en.wikipedia.org/wiki/GNOME>.

<sup>5</sup>See section 2 for details on accessing e-mail.


Please consider that the Linux terminal servers `greenhouse1/2` are general purpose servers aiming to deliver an office infrastructure! It is severely frowned upon to start highly resource consuming compute jobs on them. If noticed, such processes are typically terminated by an administrator. A number of Linux-based compute server exist that are dedicated to computing purposes, see section 8. For tasks that require a large number of compute cores for a considerable time (e.g. large parameter studies or parallel runs involving more than a dozen compute cores), users are advised to apply for an account on LiDOng, the Linux compute cluster at TU Dortmund.<sup>6</sup>

### 1.1 USB pen drives (new IGEL firmware)

Recent IGEL firmwares (version 5.08.100 and above, release dates in 2016) do not support an automount feature any more, unfortunately.<sup>7</sup> Meaning that one needs to explicitly unmount the USB pen drive before pulling it out of the IGEL terminal, as traditionally the case with operating systems like Linux, Windows and MacOS.

When plugging a USB pen drive into an IGEL thin-client terminal, it gets automatically mounted and its content appears under `$HOME/thindrives/media/[DRIVE NAME]`, if and if only the following conditions are met:

- the USB pen drive has a suitable file system, i.e. FAT16/FAT32/vfat or NTFS
- the **ThinLinc session** was **started from an IGEL** thin-client terminal. Having merely resumed an existing ThinLinc session that did not get started from an IGEL thin-client terminal is not sufficient.

Before removing the USB pen drive from the IGEL terminal, the user needs to explicitly unmount it. The unmount option is only available on the IGEL desktop and involves click in the lower right corner on the eject icon . To access the IGEL desktop when a ThinLinc session is running

- press F8 within the ThinLinc session (or in case you customised the popup menu key in the ThinLinc client's options menu your custom key)
- deactivate the option **Full screen**,
- clickin on the eject icon in the lower right corner, select the option **Safely Remove Hardware**, then choose the corresponding drive
- press F8 within the ThinLinc session,
- activate the option **Full screen**,
- unplug the drive.

Obviously, one can also choose **Disconnect Session** instead of deactivating **Full screen** to access the IGEL desktop, but would need to re-login afterwards to resume the then suspended ThinLinc session.

---

<sup>6</sup>See <https://www.itmc.tu-dortmund.de/cms/de/dienste/hochleistungsrechnen/index.html>.

The cluster is currently administrated by staff that is also part of MRB and as such working on LiDOng is very similar to working on our compute servers. Admission is rather low level, regular progress report as required at other compute facilities is currently not requested nor are there currently financial obligations for LiDOng usage for TU Dortmund employees.

<sup>7</sup>You can check your firmware version by returning to the IGEL terminal desktop, clicking on the icon with yellow background in the lower left corner of the desktop that depicts the outline of a hedgehog's snout and subsequently selecting the "About" menu entry.

## 1.2 USB pen drives (old IGEL firmware)

When plugging a USB pen drive into an IGEL thin-client terminal running a firmware before 5.08.100<sup>8</sup>, it gets automatically mounted, is accessible under `$HOME/thindrives/autofs/storageX` and automatically unmounted after 2 seconds of inactivity, if and if only the following conditions are met:

- the USB pen drive has a suitable file system, i.e. FAT16/FAT32/vfat or NTFS
- the **ThinLinc session** was **started from an IGEL** thin-client terminal

The `X` in `$HOME/thindrives/autofs/storageX` is a automatically incremented number, starting from zero. If merely a single USB pen drive with a single partition is connected, its content is accessible in subdirectory `storage0`. Two USB pen drives with 4 and 2 partitions, respectively, would be mounted as `storage0` up to `storage5`.<sup>9</sup>

The IGEL terminals are configured to automatically trigger unmounting a connected USB pen drive if no process has asked access to `$HOME/thindrives/autofs/storageX` for more than 2 seconds. So, if no file manager accesses `$HOME/thindrives/autofs/storageX` any more and there is no terminal window still running with a working directory `$HOME/thindrives/autofs/storageX` or a subdirectory thereof, a connected and mounted USB pen drive gets disconnected, the subdirectory `storageX` will get removed and it is safe to remove the pen drive from the USB port. Any attempt to access an automatically unmounted USB pen drive later on inside the ThinLinc session (e.g. by invoking

```
ls $HOME/thindrives/autofs/storageX
```

or via a graphical file manager) will cause the IGEL terminal to immediately trigger a re-mount.

## 2 Mail server

Along with a user account<sup>10</sup> a mail account is automatically set up following the scheme<sup>11</sup>

```
<forename>.<surname>@math.tu-dortmund.de
```

By default, every user gets several e-mail aliases as well, all of which redirect to the very same mail account. The aliases are:

```
<forename>.<surname>@mathematik.tu-dortmund.de    <account name>@mathematik.tu-dortmund.de
```

For historic reasons, the `.tu-dortmund.de` suffix can be substituted with `.uni-dortmund.de`, too.

This mail account is *not related to UniMail* service provided by the ITMC<sup>12</sup>. A user is free to set up mail forwarding from UniMail to the Math Faculty mail server or vice versa. Mail forwarding from the Math Faculty mail server to other addresses can be enabled by logging into the web interface <https://admin.mathematik.tu-dortmund.de:11271/> using a user's math account and password.<sup>13</sup>

The mail server provides a webmail interface at

<sup>8</sup>To check firmware version, please refer to footnote 7 on page 4

<sup>9</sup>If the ThinLinc session was, however, started from a non-IGEL device, e.g. a personal computer or laptop, the session subsequently suspended and then resumed again from an IGEL thin-client terminal, an USB pen drive will **not** get mounted automatically! A symptom is that the directory `$HOME/thindrives/autofs` did not get created automatically. In some cases, a file system object named `$HOME/thindrives/cdrom` is shown instead which is inaccessible, i.e. can not be read from nor written to.

<sup>10</sup>A user account typically consists of up to 8 characters and is in general made up using the first letter of the forename of a person and up to the first seven letters of his/her surname. In case of account name conflicts trailing letters are substituted by numbers, e.g. there `hmueller`, `hmuelle0` etc.

<sup>11</sup>The e-mail address is case insensitive.

<sup>12</sup><http://unimail.tu-dortmund.de/>

<sup>13</sup>This web site is only available from within the university network. When using an external access provider, you will need to establish a VPN or SSH connection first. Otherwise the math firewall will not let you connect to this website.

<https://www.mathematik.tu-dortmund.de/>

and an SSL/TLS-secured IMAP interface:

```
incoming server name: smtp.mathematik.tu-dortmund.de
port:                 993
protocol:             IMAP(S)
user name:            <math account>
connection security:  SSL/TLS
authentication method: normal password
```

The SSL certificate for the mail server should not raise any warnings with current e-mail clients any more! It is signed by the ITMC TU-Dortmund whose certificate's authority is in turn acknowledged by the certificate of Deutsches Forschungsnetz ("German Research Network"), usually abbreviated to DFN whose certificate's authority is in turn acknowledged by the certificate of "Deutsche Telekom Root CA 2". The latter is known by most current e-mail clients and/or operating systems such that the certificate chain will get trusted by your e-mail client. If not, visit <http://www.pki.tu-dortmund.de/> to get a copy of the root certificate "Deutsche Telekom Root CA 2" and follow the instructions given there to install it.

From inside and outside the Faculty of Mathematics network, it is possible to send e-mails using this mail server:

```
outgoing server name: smtp.mathematik.tu-dortmund.de
port:                 587
protocol:             SMTP(S)
user name:            <math account>
connection security:  STARTTLS
authentication method: normal password
```

Alternatively, a user can use the ITMC's UniMail service as a mail relay host:

```
outgoing server:     unimail.uni-dortmund.de
port:                 465
protocol:             SMTP(S)
user name:            <unimail account>
connection security:  SSL/TLS
authentication method: encrypted password
```

## 2.1 Signature

Please be aware that the TU Dortmund – for corporate design and legal reasons – obliges its employees to use a standardised signature, followed by a disclaimer. Please refer to

[https://service.tu-dortmund.de/en\\_GB/group/intra/e-mail-signatur](https://service.tu-dortmund.de/en_GB/group/intra/e-mail-signatur)<sup>14</sup>

for details and a sample signature.

---

<sup>14</sup>At the time of writing this, that web page is only available in German, despite the fact that it is listed as English version.

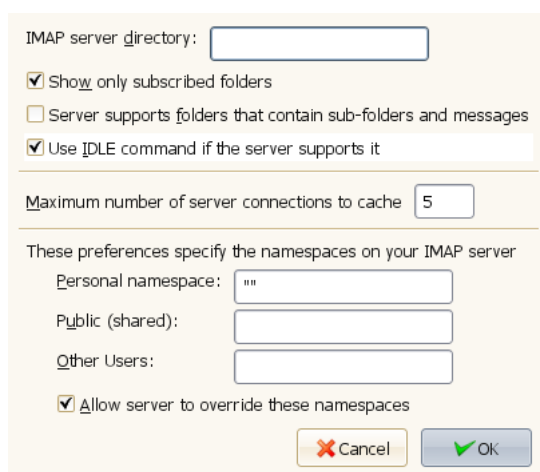


Figure 1: Mozilla Thunderbird dialog to configure mail server settings

## 2.2 Spam

The mail server does not provide a spam detection service on its own. Instead, it relies on the same content security anti-spam software solution the ITMC uses for its separate e-mail service, UniMail<sup>15</sup>. Any e-mail addressed to a mailbox hosted on our mailserver has to first pass through Cisco Ironport servers before being eventually delivered to our server. The ITMC assures that no e-mails get deleted or are withheld there. But mails can get rejected if the sending host is not properly set up as a mail host<sup>16</sup> or is (temporarily) blacklisted due to a bad reputation (for sending massive amounts of spam or malware).

Every e-mail is scanned by Cisco Ironport and a spam score is calculated. For e-mails with a spam score larger than 90 the subject is rewritten: it gets the prefix “[SPAM]”. “SUSPECTED SPAM” is prepended to the subject for e-mails with a spam score larger than 50.

## 2.3 Limitations imposed by mail server

Please note (and keep in mind) that the mail server imposes one restriction on your mail client’s settings: The server does not support folders that contain **both** sub-folders **and** messages. Just either one. Make sure to adjust your mail client accordingly! For Mozilla Thunderbird the setting is done in a dialog that can be reached via the menus **Edit - Account Settings - Server Settings - Advanced**. The option is named **Server supports folders that contain sub-folders and messages**, see figure 1.

## 2.4 Hints

Depending on your mail client you may observe folders to be listed with obscure names. Names like

```
.forward
.history
.profile
```

<sup>15</sup><http://unimail.tu-dortmund.de/>

<sup>16</sup>Technical details: according to ITMC representatives mails are rejected by Cisco Ironport for sending hosts without a fully qualified name or for which a reverse mapping via DNS fails. Additionally, mail receipt is throttled by Cisco Ironport for those servers that send out mails to too many invalid recipients (“Directory Harvest Attack Prevention”).

```
.procmailrc  
.spamassassin  
.subscriptions
```

These files are not supposed to be listed as mail folders as they are system files, but sometimes they are nonetheless displayed in a mail client. There are two possible workarounds: Either explicitly unsubscribe those “mail folders”<sup>17</sup>. Or organise all your mail folders as sub-folders of a single one. For this purpose

1. create a single mail folder, e.g. named `mailfolders`, and move all existing mail folders into this folder. Then,
2. tell your mail client that all e-mails are located in a directory with the name you just chose. In Mozilla Thunderbird, the option is named `IMAP server directory` and can be located via the menus `Edit - Account Settings - Server Settings - Advanced`, see figure 1. Because all system files will remain in their default location, which is not underneath the custom e-mail folder you just created, they will never be displayed again.

### 3 Home and data directory

Every user gets a quota (size) restricted home directory which is subject to automatic backups every workday and a file space that is not quota restricted, but also not backed up.

Regardless to which terminal or compute server, the home and nobackup file space is always the same. A user’s home directory is accessible on all Unix based servers via `/home/user/<account>`. As usual in a Unix environment, an environment variable named `$HOME` points to it.

For user convenience, there typically exists a symbolic link in a user’s home directory to the not quota restricted file space. The symbolic link is strikingly named `nobackup`. In other words, the paths `/home/user/<account>/nobackup` and `/data/<fileserver>/<account>` can be used interchangeably. As indicated by the name of the symbolic link any data stored in `/data/<fileserver>/<account>` will be definitely lost if one accidentally deletes or overwrites it.<sup>18</sup>

Quota limitation for the home directory is typically around 10 GiB. This quota restriction needs, unfortunately, to be deployed for two reasons: Firstly, to avoid single users from being able to overflow the available disk space and by this derogate the work environment of all other users. Secondly, to prevent a too big strain on the ITMC’s backup resources which would lead to the backup procedure not finishing within the 24 hour window.

The data or nobackup file space comprises approximately 10 TB shared among all users. A user is kindly requested to keep disk usage within reasonable limits to not inhibit the work of others by causing this file system to overflow. Being purposely not quota restricted this file space can overflow and has done so in the past which led to untimely ending of simulations that had been running for days and weeks. A top twenty list of users that occupy most disk space in the data file space is available at

<http://www.mathematik.tu-dortmund.de/~mrb/misc/filespaceusagepillory.php>

To reduce file space consumption, users are encouraged to delete or archive and compress old files using compression tools like `gzip`, `bzip2` etc.

---

<sup>17</sup>In Mozilla Thunderbird, from the context menu choose `Subscribe` and remove the tick from the check-box behind any “mail folder” you do not want to be displayed any more. Then click `Ok` or `Apply`.

<sup>18</sup>The files are protected by a reasonably powerful RAID solution, though, to deal with most events of *hardware* failure.



### 3.1 Backup

All data in a user's home directory except for anything stored underneath `$HOME/nobackup` is backed up automatically by the ITMC. The time the backup robot starts the backup can not be configured by the MRB and depends on how long it takes to back up other compute resources at other faculties of TU Dortmund. Having said that experience has shown that backup typically starts between 5.15pm and 5.45pm on weekdays. There is no backup on weekends.

The backup stores the latest version of a file system object (file, directory, symbolic link, etc.) as well as its previous version. Such a previous version, if any, is kept for 30 days. If a file system object is deleted from a user's file space, it can be restored from backup up to 60 days after deletion.

The backup server merely talks to the file servers hosting the home directories – warehouse13 and warehouse15 –, both for backup and restore operations. Given that direct user login to the file servers is restricted for security reasons a user that needs some files or directories to be restored from backup needs to contact the MRB to perform this task. A user is kindly requested to be as precise as possible which files and/or directories need to be restored if it comes to the pinch.

## 4 Printing

Every chair at the Faculty of Mathematics has on average two printers that are connected to the network. They are named according to their features and the room number where they are physically installed. If you are a new user, ask a colleague for the printer names and their location.

The printers are administered through a CUPS print service.

The default printer can be set by means of the environment variable `$PRINTER`. To set the default printer to, e.g., `lsiii_col_dup`, tcsh users invoke<sup>19</sup>

```
setenv PRINTER lsiii_col_dup
```

and bash users

```
export PRINTER=lsiii_col_dup
```

Typically, the system already sets the `$PRINTER` environment variable at login time correctly based on the chair a user belongs to. To overwrite this setting permanently, tcsh users are required to add the abovementioned command to `$HOME/.cshrc` and bash users to `$HOME/.bashrc`, both being files that are parsed automatically at login time after the system settings file have been parsed.

To actually print a postscript or PDF document from the command line, use

```
lpr <file>
```

or specify the printer queue specifically (overriding the `$PRINTER` setting) via

```
lpr -P<printer> <file>
```

Example: `lpr -Plsiii_col_dup <file>`

## 5 Software

Software requirements vary greatly among different chairs and users. In order to cater for those varying needs, all servers rely on a software package called “Environment Modules”<sup>20</sup>. It facilitates the process of

<sup>19</sup>To determine whether your login shell is set to tcsh or bash, please check the environment variable `$SHELL`. It can be queried, e.g., by `echo $SHELL`.

<sup>20</sup>See [https://en.wikipedia.org/wiki/Environment\\_Modules\\_%28software%29](https://en.wikipedia.org/wiki/Environment_Modules_%28software%29).

adapting environment variables (i.e. search paths) in order to be able to locate and start a specific software package in a certain version. Example: The GNU Compiler Collection (gcc) is installed in a variety of versions on the various Linux servers. Older software packages typically require older GCC releases while newer prefer more recent releases. By means of the command

```
module avail gcc
```

a list of all installed gcc versions can be retrieved. To learn more about a specific module, issue its fully qualified name (excluding any string contained in parenthesis, if any)

```
module help cmake/2.8.2
```

This will print a small help screen giving away details about this module. A dedicated software version can be made available by invoking

```
module load cmake/2.8.2
```

To switch from a previously loaded CMake version to another one issue

```
module switch cmake/2.8.2 cmake/2.8.10.2
```

A module can be explicitly made unavailable by means of

```
module unload cmake/2.8.10.2
```

Please continue reading <http://www2.hlrn.de/doc/modules/> for a comprehensive tutorial about the “Environment Modules”.

If a user finds himself/herself to frequently load the same number of modules, it is advisable to add the relevant commands to a user’s personal initialisation file. For tcsh users this is `$HOME/.cshrc` and for bash users `$HOME/.bashrc`. A sample `$HOME/.cshrc` that distinguishes between different Unix operating systems is

```
set operatingSystemAndKernel='uname -a'

# Show a welcome message to the user to inform him what is being done here
# automagically on login - BUT only in case of an interactive login. For
# non-interactive logins - think about scp and sftp connections - any output on
# standard output or standard error must be avoided. Otherwise the connecting
# scp/sftp client is very likely to immediately hang up after connecting to the
# server!
set allowMessages=0
if ($?tcsh && $?loginsh) then
    set allowMessages=1
endif

switch ( "$operatingSystemAndKernel" )

case Linux*:
    # Check whether /etc/issue exists. If it does, it contains the name of
    # the Linux distribution and allows this way a finer grained control
    # than simply comparing Linux kernel numbers
    if ( -e /etc/issue ) then
        switch ( "'cat /etc/issue'" )
            case CentOS*release*6*:
                if ($allowMessages) then
                    echo "# (~/.cshrc:) Automatically loading modules for CentOS 6.x system"
                endif
            default:
                # Do nothing
        endswitch
    endif
endcase
endswitch
```

```

# Load Intel C/C++ and Fortran compiler and GotoBLAS1 - to
# be able to compile Featflow2 (relying on the requirements
# of the default build ID)
module load intel/studio-xe/18.0.2.199
module load paraview
breaksw

case CentOS*release*7*:
if ($allowMessages) then
echo "# (~/.cshrc:) Automatically loading modules for CentOS 7.x system"
endif

# Load Intel C/C++ and Fortran compiler and GotoBLAS1 - to
# be able to compile Featflow2 (relying on the requirements
# of the default build ID)
module load intel/studio-xe/18.0.3.222
module load paraview
breaksw

case default:
# unconfigured operating system
if ($allowMessages) then
echo "# (~/.cshrc:) Please note that ~/.cshrc has not been set up yet"
echo "#           for this operating system!"
endif
breaksw
endsw
endif

# Put settings here that hold for all Linux systems, e.g. alias definitions

breaksw

endsw

# Have tab completion be silent, if it can not complete a file name.
set no beep

# Shortcut aliases
alias matlab 'module add matlab; \matlab'

```

Bash shell users have to edit a different startup file, namely `$HOME/.bashrc`. Equivalent instructions for them read like

```

operatingSystemAndKernel='uname -a'

# Show a welcome message to the user to inform him what is being done here
# automagically on login - BUT only in case of an interactive login. For
# non-interactive logins - think about scp and sftp connections - any output on
# standard output or standard error must be avoided. Otherwise the connecting
# scp/sftp client is very likely to immediately hang up after connecting to the
# server!

```

```

# Detect a non-interactive session by trying to remove 'i' from environment
# variable $-.
if [ "${-}" != "${-/i/}" ]; then
  allowMessages=1
else
  allowMessages=0
fi

case "$operatingSystemAndKernel" in

Linux*)
  # Check whether /etc/issue exists. If it does, it contains the name of
  # the Linux distribution and allows this way a finer grained control
  # than simply comparing Linux kernel numbers
  if [ -e /etc/issue ]; then
    case "`cat /etc/issue`" in
      CentOS*release*6*)
        if [ $allowMessages -eq 1 ]; then
          echo "# (~/.bashrc:) Automatically loading modules for CentOS 6.x system"
        fi

        # Load Intel C/C++ and Fortran compiler, Math Kernel library and Maple
        module load intel/studio-xe/18.0.2.199
        module load maple
        ;;

      CentOS*release*7*)
        if [ $allowMessages -eq 1 ]; then
          echo "# (~/.bashrc:) Automatically loading modules for CentOS 7.x system"
        fi

        # Load Intel C/C++ and Fortran compiler, Math Kernel library and Maple
        module load intel/studio-xe/18.0.3.222
        module load maple
        ;;

      default)
        # unconfigured operating system
        if [ $allowMessages -eq 1 ]; then
          echo "# (~/.bashrc:) Please note that ~/.bashrc has not been set up yet"
          echo "#           for this Linux system!"
        fi
        ;;
    esac
  fi

  # Put settings here that hold for all Linux systems, e.g. alias definitions

  ;;

default)

```

```

# unconfigured operating system
if [ $allowMessages -eq 1 ]; then
    echo "# (~/.bashrc:) Please note that ~/.bashrc has not been set up yet"
    echo "#             for this operating system!"
fi
;;
esac

# Have tab completion be silent, if it can not complete a file name.
set bell-style visible

# Shortcut aliases
alias matlab='module add matlab; \matlab'

```

## 5.1 Editors

List of available editors on the Linux terminal servers `greenhouse1/2`:

- `soffice/libreoffice` (= Libre Office in version 4.x)
- `emacs`
- `vi`
- `vim`
- `gedit`
- `kate`
- `kwrite`
- `lyx`
- `nedit`

Via a terminal session on the Windows server<sup>21</sup> or an encrypted console session (`ssh`) on one of the other Linux servers additional editors are available for terminal users:

- Open Office 4.1.1 (from within a windows terminal session, see section 6)
- Microsoft Office Professional Plus 2010 (from within a windows terminal session, see section 6)
- Notepad/Write (from within a windows terminal session, see section 6)
- TeXmaker (from within a windows terminal session, see section 6)
- PSPad editor (from within a windows terminal session, see section 6)
- `kate/kwrite` (from within a `ssh` session to a Linux server)
- LibreOffice 4.x (from within a `ssh` session to a Linux server)
- `vi`
- `vim`
- `gedit`

<sup>21</sup>Which can be started by issuing the command `win.alpha --no-fullscreen`. Remember to contact the MRB to request account setup on the Windows terminal server before trying to connect to it for the first time.

- kate
- kwrite
- nedit

## 5.2 Browsers

List of available browsers on the Linux terminal servers:

- firefox
- konqueror
- chrome

Via a terminal session on the Windows server or an encrypted console session (ssh) on one of the Linux servers additional browsers are available for Linux terminal users:

- Internet Explorer (from within a windows terminal session, see section 6)

## 5.3 E-mail client

List of available e-mail clients on the Linux terminal servers:

- thunderbird
- alpine (text-based)
- emacs (text-based)

## 5.4 PDF viewer

List of available PDF viewers on the Linux terminal servers:

- Adobe Reader `acroread` (via `module` command)
- `evince`
- `ghostview`
- `gv`
- `okular`

## 5.5 OpenGL software

A word of advice for those who want to use OpenGL based visualisation software, e.g. GMV or ParaView: The IGEL thin-client terminals inherently can not provide hardware rendering. OpenGL commands that are typically interpreted by graphic chips (GPUs) need to be translated into software commands, i.e. realised in software rendering. This translation is supposed to happen transparently and works more stable than on the old Sun Ray servers the faculty had in use until summer 2014. But due to software bugs, the translation might not go smoothly and as a result a user's desktop session might crash (resulting in an immediate logout and the loss of unsaved data) when trying to use hardware rendering. Terminal users are hence advised to not use the command `gmvgl` or any of the official ParaView binaries, but their Mesa based counterparts instead: the command line program `gmv` instead of `gmvgl` and ParaView modules that clearly state having been built against Mesa – this holds for any of ThinLinc user, regardless whether he/she connects from an IGEL thin-client terminal or a private computer.

## 6 Windows server

A user account can be, but is not by default, granted access to the terminal servers running Microsoft Windows Server 2008. To access them, the user needs to invoke the command

```
win.alpha
```

on the command line when logged in to one of Linux terminal servers. There is only a limited number of terminal licenses available, so accessing the terminal servers may not be possible at all times.

## 7 Web server

A user's default web page

```
www.mathematik.tu-dortmund.de/de/personen/person/<firstname>+<familyname>.html
```

is automatically generated at night from data stored in the Faculty of Mathematics database. A user can change his/her personal settings like research interests and research tasks, e-mail forward, specify a link to a separate private web site (as well as the account password) via the aforementioned web interface at

```
https://admin.mathematik.tu-dortmund.de:11271/
```

using the math account and password. A dedicated login to the web server is not required for this.

Sometimes it is necessary to add custom web pages and/or upload material to the web server. One way to accomplish this is the mini website system (see next section). Another – providing more freedom, but entailing a greater security risk for the server – is to grant a user account access to the web server (euler) for this matter. Please contact the MRB for that.

## 8 Compute servers

A number of Linux based compute servers is available at the Faculty of Mathematics. Their usage, though, is typically restricted to the members of a particular chair because the servers were project based provided to solve dedicated research problems. For a complete list of compute servers and the mapping compute server to chair see the web site [www.mathematik.tu-dortmund.de/~mrb/serverload.html](http://www.mathematik.tu-dortmund.de/~mrb/serverload.html).<sup>22</sup> The web site

<sup>22</sup>There is a custom command line program available to query all compute servers, `showload.pl`, and another custom command line program to query single compute servers: `askloadserver <host>`. They will display statistics about current load, load average in last 5 and 15 minutes, users logged in and cpu plus main memory details.

also lists the number of resources per compute server. The colours green, yellow and red are intended to indicate how many of the resources of a compute server are already in use. To start a computation on a compute server that exhibits already red load levels is considered imprudent and reckless behaviour and strongly discouraged. If you choose to do so nonetheless, you basically would be penalising yourself as your job is likely to run longer than on a compute server less crowded.

Please, contact the server owning chair (resp. MRB) if you are not a member of that chair and need to allocate its resources (cpu, memory) permanently. Acquirements of new servers must to be discussed with MRB.

## 8.1 Available compiler

Main supported compilers are GCC and Intel, the last release of GCC4 - GCC8 and of Intel 16 - 18 will be installed.

## 8.2 OpenMPI

The current revisions of the main lines 2 and 3 will be installed, version 1 will no longer be installed.

There are implementation and options modules.

### Implementation modules:

```
openmpi/2.1.3
openmpi/3.1.1
openmpi/current-v2
openmpi/current-v3
```

The meta modules `current-v2` and `current-v3` refer to the last revision of that main line.

### Option modules:

```
openmpi/options/cuda/yes
openmpi/options/cuda/no *

openmpi/options/interface/ethernet *
openmpi/options/interface/infiniband

openmpi/options/threading/yes
openmpi/options/threading/no *
```

(Option modules marked with an \* are default)

The choice of the compiler is made by loading the appropriate compiler module, for example:

```
# module add gcc/5.5.0 openmpi/options/threading/yes openmpi/2.1.3
# which mpirun
# /sfw/openmpi/gcc5.5.x/2.1.3/threaded/bin/mpirun
```

If options or compilers are to be changed, the implementation module must first be unloaded. An implication of this scheme is that two different compilers can not be loaded at the same time

```
# module add gcc/4.8.5-centos intel/studio-xe/18.0.3.222
```



so does not work. But if you want to use the Intel for serial compilation or only use the MKL, the following works:

```
# module add gcc/4.8.5-centos serial/intel/studio-xe/18.0.3.222
```

The Intel compiler is now in the search path, but has no access to the MPI bindings. For the use of the MKL with the GCC, the following linker options must be set:

Multithreaded usage of the MKL (the MKL\_NUM\_THREADS environment variable for the number of threads has to be set):

```
-L${MKLRROOT}/lib/intel64 -Wl,--no-as-needed -lmkl_intel_lp64 \
-lmkl_intel_thread -lmkl_core -liomp5 -lpthread -lm -ldl
```

Serial usage of the MKL:

```
-L${MKLRROOT}/lib/intel64 -Wl,--no-as-needed -lmkl_intel_lp64 \
-lmkl_sequential -lmkl_core -lm -ldl
```

### 8.3 Workload manager Slurm

Since the freely available version of torque / maui has not been maintained for years and the ITMC for LiDO3 also relies on slurm, the faculty will follow this step and also use slurm as a scheduling system. This scheduling will be used for all compute servers to provide better utilization and also to prevent users from blocking each other's resources. Unlike torque, a user can never claim more than the resources assigned to him. A direct SSH login will therefore no longer be possible on the servers.

For usage the slurm module must first be loaded:

```
# module add slurm/current
```

The sinfo command gives an overview of the partition structure (analogous to LiDO3):

```
# sinfo
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
short*      up    2:00:00    11  idle arrow,...,mcgee,supergirl
med         up    8:00:00    11  idle arrow,...,mcgee,supergirl
long        up 2-00:00:00    11  idle arrow,...,mcgee,supergirl
ultralong   up 28-00:00:0    11  idle arrow,...,mcgee,supergirl
```

The partitions are used to control the runtime of the jobs, the hardware is selected via so-called constraints.

The command sgetconstraints displays the assignment of the constraints to the respective machines.

```
# sgetconstraints
Constraint arrowverse : arrow flash supergirl

Constraint ether10g : hawaii-five-[0-3] arrow flash supergirl

Constraint geforce : hawaii-five-[0-3]

Constraint gpu-accelerator : hawaii-five-[0-3] graufreud
```

```
Constraint gtx980ti : hawaii-five-[0-3]
Constraint hawaii-5 : hawaii-five-[0-3]
Constraint infiniband : hawaii-five-[0-3]
Constraint ncis : gibbs dinozzo mcgee david
Constraint opt_8220 : gibbs dinozzo mcgee david
Constraint p100 : graufreud
Constraint tesla : graufreud
Constraint xeon_e5-1620-v3 : hawaii-five-[0-3]
Constraint xeon_e5-1620-v4 : graufreud
Constraint xeon_gold-6146 : arrow flash supergirl

Legend:
* constraint arrowverse:
  Systems belonging to the arrowverse cluster

* constraint ether10g:
  Systems equipped with 10G ethernet connection

* constraint geforce:
  Systems equipped with a geforce gpu accelerator

* constraint gpu-accelerator:
  Systems equipped with a gpu accelerator

* constraint gtx980ti:
  Systems equipped with a geforce GTX980 TI accelerator

* constraint hawaii-5:
  Systems belonging to the hawaii-5 cluster

* constraint infiniband:
  Systems equipped with infiniband connection

* constraint ncis:
  Systems belonging to the ncis cluster

* constraint opt_8220:
  Systems equipped with an opteron 8220 cpu

* constraint p100:
  Systems equipped with a tesla P100 accelerator

* constraint tesla:
  Systems equipped with a tesla gpu accelerator
```

```
* constraint xeon_e5-1620-v3:
  Systems equipped with a Xeon E5 1620-v3 cpu

* constraint xeon_e5-1620-v4:
  Systems equipped with a Xeon E5 1620-v4 cpu

* constraint xeon_gold-6146:
  Systems equipped with a Xeon Gold 6146 cpu
```

### 8.3.1 Usage of slurm

#### Interactive session

To start an interactive session you can use the `srun` command

```
# srun -p short --nodes=1 --constraint=arrowverse --mem=128M --pty $SHELL -l
```

This command starts a console on one node (`--nodes`) on the arrowverse cluster (`--constraint`) with 128 megabytes main memory in the short partition.

Alternatively you can use the `sgetlogin*` shortcut.

```
sgetlogin-[short|med|long|ultralong] [CONSTRAINT] [--] [SLURM_OPTIONS]
```

simply type

```
# sgetlogin-short
```

you get a console somewhere on the system with one node.

Important slurm options:

- `--nodes=X` : count of nodes to be reserved
- `--mem=X` : amount of memory per node to be reserved
- `--mem-per-cpu=X` : amount of memory per cpu to be reserved
- `--time=HH:MM:SS` : runtime of the job/session
- `--ntasks-per-node=X` : count of processes per node
- `--cpus-per-task=X` : count of cores per process (multithreading)

#### Use of GPUs

You have to ask for GPU use explicitly, only specify a constraint with available GPUs is not sufficient.

```
# sgetlogin-short hawaii-5
hawaii-five-0# nvidia-smi
No devices were found
```

To reserve a GPU you have to use the `--gres` option.

```
# sgetlogin-short hawaii-5 -- --gres=gpu:1
hawaii-five-0# nvidia-smi
Wed Aug 15 14:27:34 2018
+-----+
| NVIDIA-SMI 396.44                Driver Version: 396.44                |
+-----+-----+-----+-----+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+-----+-----+-----+
|    0  GeForce GTX 980 Ti  On      | 00000000:04:00.0 Off  |          N/A         |
| 22%   32C   P8     16W / 250W |  0MiB /  6083MiB |    0%      Default   |
+-----+-----+-----+-----+-----+-----+
+-----+
| Processes:                         GPU Memory |
|  GPU       PID    Type    Process name                     Usage      |
+-----+-----+-----+-----+-----+-----+
| No running processes found         |
+-----+
```

The syntax for the option is `--gres=gpu:n`, where `n` is the number of GPUs per node to be reserved. For `n=1` this information can be omitted.

### X11-Forwarding

Unconditional login via ssh to a compute server is not possible, you have to start a session first and then login to the assigned node.

```
# sgetlogin-short
hawaii-five-0: #
```

Then in another console you can login to that node.

```
# ssh hawaii-five-0
hawaii-five-0: # xterm
```

Without a running session you will get the message:

```
# ssh hawaii-five-0
Access denied by pam_slurm_adapt: you have no active jobs on this node
Authentication failed.
```

### Batchjobs

Batch jobs can be submitted with

```
sbatch [JOBSCRIPT]
```

can be canceled prematurely with

```
scancel [ID]
```

and their status can be queried using

```
squeue
```

or

```
sfill
```

### Serial jobs

Example `slex_serial.sbs`:

```
#!/usr/bin/env tcsh

#definition of output log file, queue, count of nodes, amount of memory and
#runtime of the job
#SBATCH -o output.log
#SBATCH -p short
#SBATCH --nodes=1 --mem=256G --time=00:30:00

#initialization of module environment
source /sfw/Modules/default/init/tcsh

#load of modules
module purge
module add intel/studio-xe/18.0.3.222 slurm/current

#change to program directory
cd /home/user/cbecker/playground/hpl/bin/LID03

#start of program
./xhpl
```

It is necessary to load the slurm module within the batch job, otherwise the slurm job can not be started.

### Parallel job singlethreaded

Example `slex_parallel_single.sbs`:

```
#!/usr/bin/env tcsh

#--nodes specifies the number of physical machines, while
#--ntasks-per-node specifies the number of processes per node
#SBATCH -o output.log
#SBATCH -p short
#SBATCH --nodes=3 --ntasks-per-node=4 --constraint=arrowverse --mem=32G --time=00:30:00

source /sfw/Modules/default/init/tcsh

module purge
module add intel/studio-xe/18.0.3.222 openmpi/3.1.1 slurm/current

cd /home/user/cbecker/playground/hpl/bin/LID03
```

```
#the mpirun has to be started with the -np parameter and overall
#count of processes (3 nodes * 4 ntasks-per-node = 12)
mpirun -np 12 xhpl
```

### Parallel jobs multithreaded

Example `slex_parallel_multi.sbs`

```
#!/usr/bin/env tcsh

#--cpus-per-task specifies the number of threads per core / task
#SBATCH -o output.log
#SBATCH -p short
#SBATCH --nodes=3 --ntasks-per-node=4 --cpus-per-task=4
#SBATCH --constraint=arrowverse --mem=256G

source /sfw/Modules/default/init/tcsh

module purge
module add intel/studio-xe/18.0.3.222 openmpi/3.1.1 slurm/current

#specify the number of threads each MKL process should use
setenv MKL_NUM_THREADS 4

cd /home/user/cbecker/playground/hpl/bin/LID03

#The MKL_NUM_THREADS variable must be passed to the mpirun command, otherwise it will
#not arrive at the node
mpirun -x MKL_NUM_THREADS -np 12 xhpl
```

## 9 SSH, SFTP and FTP access

The only means to connect to the network of the Faculty of Mathematics is the SSH protocol. FTP access is purposely not available. Because of abuse in the past, the default SSH port of 22 is blocked by the firewall and hence not accessible from outside the math network. Use the port 22122 instead.

Free and commercial SSH/SCP/SFTP clients are available for all operating systems. Prominent examples for Microsoft Windows platforms are

- SSH access
  - KiTTY, <http://www.9bis.net/kitty/>
  - PuTTY, <http://www.chiark.greenend.org.uk/~sgtatham/putty/>
  - SSH client shipped with Cygwin
- for SCP/SFTP access
  - WinSCP, <http://winscp.net>
  - FileZilla, <http://filezilla-project.org>

Please ask the MRB if you need help choosing a program.

Please note that another popular SSH client for Windows, namely the last non-commercial version of **SSH Secure Shell**<sup>23</sup> got released in 2003 and is now too outdated. It can not be used to access our servers as it does not support the required stronger encryption algorithms.

When trying to log in from a IP address outside of the Math network, you have 3 attempts to log in successfully. After 3 failed attempts, the IP address is banned from further logins for 30 minutes in order to obstruct brute force SSH password attacks.

## 9.1 Chain logins

SSH and SCP/SFTP are installed on all our servers and workstations permitting a user to log in from a server A to another server B. It is obviously possible on server B to use ssh again to log in to another machine C or even back to A again. This situation is referred to as a “chain login”. Resorting to chain logins is strongly discouraged!

Why is this a bad habit? Firstly, because all data (your keyboard input as well as any screen output) needs to be repeatedly encrypted. Once for every machine you ssh-d through. This obviously slows down response time and is simply a waste of resources and network bandwidth.

Secondly, a chain login heightens the risk of loosing a unsaved changes in an editor running on a remote host or forfeiting a long-running simulation. Obviously, every server in your login chain needs to keep running as long as your editor session or simulation does. Admittedly, servers tend to be online 24/7, but their operating system needs to be patched regularly for security and bug-fixing reasons and some of these patches require a reboot. Beyond that, you might want to keep in mind that even server hardware is not perfect and may fail. For instance due to power fluctuations that last even less than a second which do, on average, occur once or twice a year and might cause a server to be abruptly turned off. Servers with powerful GPUs (ashenvale, azeroth, kalimdor, stonetalon) are reboot prone, too, because the way these GPUs are used at our department regularly reveals bugs in the GPU driver, bugs which are merely remediable by a reboot.

## 9.2 Remote desktop

Beside plain terminal sessions through the SSH protocol, graphical sessions are possible, too. The terminal servers **greenhouse1/2**, all Linux workstations as well as a few Linux compute servers run a 50-user license of Cendio ThinLinc each. The software enables a user to work from his personal computer or laptop at home on the Faculty of Mathematics servers as if he/she were logged in regularly in his office. The client software is available for most operating systems and downloadable free of charge via

<http://www.cendio.com/downloads/clients/>

A ThinLinc client is configured as follows: As server to connect to give the fully qualified server name

```
<server>.mathematik.tu-dortmund.de
```

as well as your user account name and password in the appropriate fields. There are two non-default settings one needs to specify in the advanced dialog that pops up when hitting the “Options” button. In the “Security” tab set the SSH port to 22122 and in the “Local Devices” tab disable export of the local sound device<sup>24</sup>. Please note that with ThinLinc one does not need to end a user session by logging out. Such a terminal session can also be suspended and resumed later, even if the re-connect happens from a different hardware, operating system or at a different screen resolution.

<sup>23</sup>Search FTP sites for a binary named SSHSecureShellClient-3.2.9.exe

<sup>24</sup>Note that this option is unavailable in the ThinLinc client for Mac OS X at the time of writing.

## 10 Changing user password and mail forwarding

A single password for all the servers and IT services at the Faculty of Mathematics (“single sign-on”) is used. A user can change his/her account’s password by logging in to

`https://admin.mathematik.tu-dortmund.de:11271/`

selecting “Benutzer” – “Passwort ändern” and entering and confirming a new one.<sup>25</sup> Unlike in previous years, this change takes effect immediately. In other words, from that moment on the new password is required to access the mail server, the web server and all other 70+ compute servers at the Faculty of Mathematics.

Mail forwarding settings can be set and altered using the same URL via the menu “Benutzer” – “Mailumleitung ändern”.

---

<sup>25</sup>Note that a few Windows installations (with name prefix “unimatrix”) still use their own, separate authentication server. If you have an account there, changing your Math network user password via the mentioned database interface URL will not automatically update the password information on those Windows servers, too, unfortunately.