

# 2018 Annual Maintenance

During the summer a significant amount of work will be going on on the clusters. There are two periods of downtime (July and August) that you should be aware of. More details are given below.

- [Maintenance Windows](#)
  - [9th to 17th of July](#)
  - [August 14th \(Deneb\) and August 21st \(Fidis\)](#)
- [Details](#)
  - [Operating System and SLURM scheduler](#)
  - [New Software Stack](#)
  - [Build nodes and serial nodes](#)

## Maintenance Windows

### 9th to 17th of July

During this period all the clusters will be unavailable. We will be carrying out work including upgrading the OS and the software stack.



As the operating system version, the scheduler and the software stack will all change the queues will be emptied. Any jobs still in the queue on the 9th of July will be cancelled.

### August 14th (Deneb) and August 21st (Fidis)

The standard (every three months) maintenance days remain in place to allow for any corrections required following the major upgrades.

## Details

### Operating System and SLURM scheduler

In order to provide a modern environment and apply performance, security and bug fixes we will be changing the operating system and batch system.

- Deneb and Fidis will be upgraded to the latest RedHat 7 release
- The SLURM scheduler will be upgraded to version 17.11
- CUDA on the GPU nodes will be upgraded to version 9
- The Infiniband stacks will be upgraded to the latest versions
- Tasks requiring less than a whole node will be automatically sent to the serial partition via a job submit plugin

### SSH Keys changed

As a consequence of the reinstallation the ssh host keys have changed for all the nodes (login and compute). That's why you see the "WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!" message. First you have to remove the offending key, then you must check the new key against the list below before accepting the new one, the type picked up (rsa, dsa) depends on your ssh client version.

Check the current saved key

```
ssh-keygen -F fidis.epfl.ch -l
```

Remove the old key

```
ssh-keygen -R fidis.epfl.ch
```

Here are the keys for the login nodes:

Fidis		
DSA	SHA256:HvbcrjP3X29Bn0g1OgXB+vHLEjWuSWvoaf/vLdZoe/Q	MD5:7a:a0:c8:f2:e9:2d:7f:e7:66:69:4a:d3:b8:dd:22:bd
ECDSA	SHA256:mLntSNT2Rz69Nba7jOyWt6r+kwF4zpmZZlyrRGfRgXw	MD5:5a:c4:54:4d:c5:44:9c:f5:10:1d:e8:17:09:ab:27:4e
ED25519	SHA256:Syy7jQvfKZKMUcWXBm+MIzLxPyVxsdGEiSCylr0kmQI	MD5:aa:3c:6e:18:9d:95:3f:cc:73:77:a8:f9:1a:fb:b7:9a
RSA	SHA256:mqLUyQMqKNzC8b+HgBPvf3QRnojeZWFJNWHCT7dE1FM	MD5:2f:d3:37:32:88:4f:e6:70:29:15:d8:48:10:b4:d5:1f

Deneb1		
DSA 5b	SHA256:jjg2UsZi+5vn4JVSHhHO+8nCbJe+PYeus5oXDy5HrIg	MD5:59:d4:c1:a0:97:bc:90:ae:69:46:af:03:41:3c:3b:
ECDSA 3f	SHA256:/4HY5EVScDLu/Btzr0Rr7b+G9vcZtqp2a3P6NoH5Ey4	MD5:95:76:2b:0d:10:ad:2d:be:ca:cd:8e:37:3b:f5:56:
ED25519 ef	SHA256:fUxBEzD/f9VKiEzApNQq6bs2p/9d1vUcYeplHrzNVfc	MD5:f0:03:14:d8:6c:a0:66:61:b9:65:18:04:b2:5a:27:
RSA 15	SHA256:NbpSOygh72m2b4xb4bptwv0of8RXhHiYMJeuUsG4hcA	MD5:b8:02:e0:27:8b:75:ce:9a:3f:5a:78:18:de:7e:3a:

Deneb2		
DSA	SHA256:jjg2UsZi+5vn4JVSHhHO+8nCbJe+PYeus5oXDy5HrIg	MD5:59:d4:c1:a0:97:bc:90:ae:69:46:af:03:41:3c:3b:5b
ECDSA	SHA256:kYv/aUHYRlv7NUGDqieMEET9EcB3Uk4CDVSAtuoJ5c	MD5:b6:53:ec:b0:a8:aa:a9:8c:76:e5:f0:ac:d5:8d:88:11
ED25519	SHA256:R/BxGi0pI1W9s7o2xNOcxcHuo4fneu+hWAgHR1nimDo	MD5:04:f6:b9:f1:77:b2:a6:c3:28:7c:ac:b0:30:93:59:4c
RSA	SHA256:NbpSOygh72m2b4xb4bptwv0of8RXhHiYMJeuUsG4hcA	MD5:b8:02:e0:27:8b:75:ce:9a:3f:5a:78:18:de:7e:3a:15

## Storage

We will carry on upgrades and performance improvements to the filesystems:

- Upgrade of GPFS to version 5.0.1-1
- Reconfigure remote clustering in GPFS, in particular:
  - use the modern GPFS key exchange protocol (remove `nistCompliance=off`)
- Improve performance between Fidis and the shared storage, in particular:
  - increase usable bandwidth (MTU >1500)
  - balance usage of the redundant links (proxy-arp and lacp hash function)

We will also complete the migration to the new global storage system by physically removing the old GSS storage.

## New Software Stack

The *software stack* is everything you see and load with the `modules` command. Every summer we release a new bundle of packages with the main changes being the compiler and MPI versions as well as newer versions of many packages.

### Summary

- try the new software stack release now available as *future*: `slmodules -r future`
- let us know if you use any packages/libraries not yet available in *future* or the current default modules (release *stable*)
- contact us for any questions or issues by sending an email to [1234@epfl.ch](mailto:1234@epfl.ch) starting the subject with 'HPC'

## Future

During the maintenance period in July (9-16) the next release of our software stack will be promoted to stable.

As of today the new release is available under the name 'future' and can already be used by issuing the following command before loading any modules:

```
$ slmodules -r future
$ module load <module name>
```

We encourage you to try the new release as soon as possible and let us know (by opening a ticket) if you find any issues or missing software packages.

With this release newer versions of almost all packages and libraries are available. Please be aware that for packages with multiple package versions only one of the versions will be kept once the 'future' release becomes the default one.

In particular the supported compilers are now

- **Intel 2018 update 2**
- **GCC 6.4.0**
- **GCC 7.3.0**

## Note on Future and operating systems

If you compile code now using "Future" it is not impossible that it will need to be recompiled after July 16th, although the software release is the same, the underlying operating system and scheduler versions will be different.

## Current release (Stable)

The current stable release will still be available under the name 'deprecated'.

To use what is now the stable release after the 16th of August you will need to change to the 'deprecated' release before loading any modules:

```
$ sload -r deprecated
$ module load <module name>
```

## Deprecated

Please also be aware that the current *deprecated* set of software modules will no longer be available.

If you are still using any software packages or libraries from the current *deprecated* release please check the *future* release today to see if it is available there, and let us know immediately if it is not.

## Build nodes and serial nodes

### Build Nodes

There is now an additional "build" partition available on the clusters to allow you to compile codes for the different architectures. Please note that, unlike the "debug" partition, these nodes are shared. To make use of them please use the Sinteract tool using the "-s" flag to specify the architecture and select the appropriate number of CPUs and amount of memory.

These nodes also have a slightly different operating system configuration to allow debuggers such as GDB to work correctly.

### Serial Nodes

In order to increase system stability the nodes in the serial partition on Deneb now have two CPUs reserved for system use. This means that a maximum of 14 cores are available per serial node.

### Large Memory Nodes on Deneb

The nodes with 256 and 512 GB of memory on Deneb are now part of the serial partition as they are intended for non MPI tasks.

If you require access to these nodes then you should explicitly ask for the serial partition with "-p serial".

For large memory MPI calculations you should use Fidis which has 72 nodes with 256GB of memory connected via fully non-blocking FDR Infiniband.