

PlaFRIM

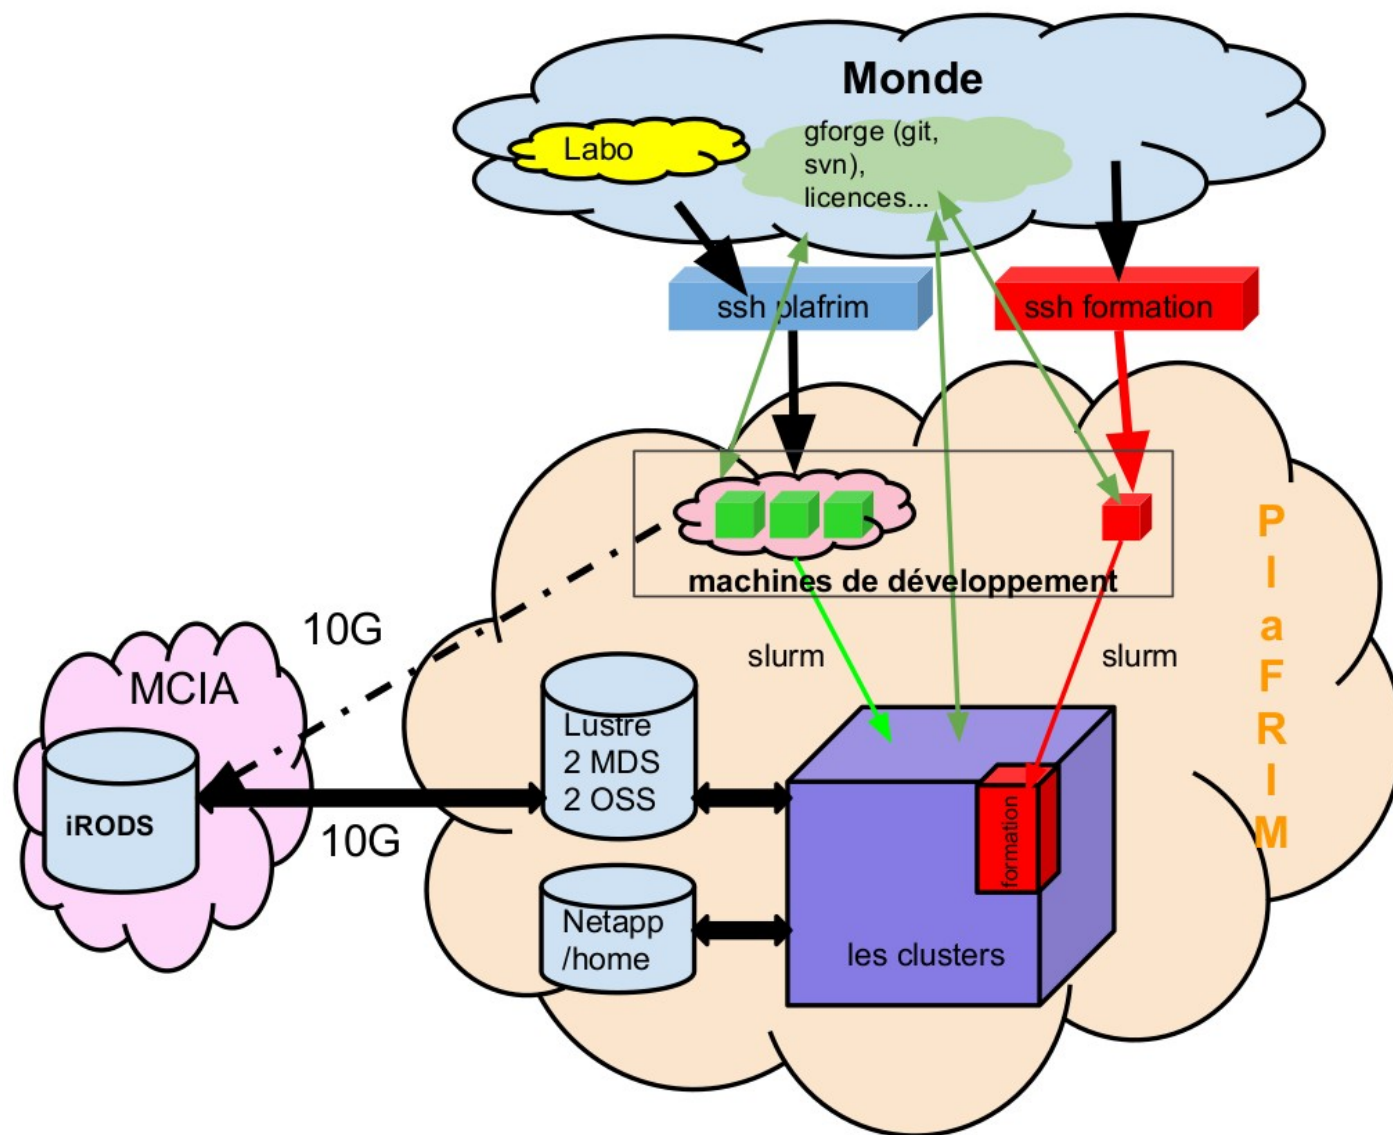
Technical presentation of the platform

Contents

- 01. Overview
- 02. Nodes description
- 03. Networks
- 04. Storage
- 05. Evolutions
- 06. How to acces PlaFRIM ?
- 07. Need Help ?

01

Overview



2 platforms (1 research and 1 formation)

- More than **100 nodes** and **3000 cores**
- More than **35 GPUs** accelerators
- More than **350 users**

Heterogeneous cluster

- Nodes for **modelization** (equivalent nodes of mesocentres or computer center of GENCI)
 - Miriel Cluster (88 nodes)
 - Souris node (1 SGI UV200 96 cores and 3TB memory)
- Nodes for **experimentation** (more innovative nodes)
 - Mistral cluster (18 nodes with 2 x Xeon Phi accelerator)
 - Sirocco cluster (13 nodes with Nvidia GPU accelerator)
 - Kona cluster (4 nodes with KNL)
 - Brise node (1 node with 96 cores and 1TB memory)

The hidden side

- a **security audit** has been carried out from which certain security rules have been applied (access with **SSH key pair**, **isolated network** between formation and research cluster, **internet filtering**,...)
- 1 **master node** which provides :
 - Name resolution (**DNS**), IP attribution (**DHCP**), **SLURM** server, Server **TFTP**, deployment node tool (**Bright Cluster Manager**), **internet gateway**, and more...
- 2 **hypervisors** (oVirt nodes) to provide infrastructure virtual servers :
 - 2 **LDAP** in master/master mode and 1 LDAP for the formation cluster to
 - 2 server **GUIX** (formation and research clusters)
 - 1 supervision based on **zabbix**
 - 1 **HIDS** (intrusion detection system)
 - 1 **NFS** for the formation cluster
 - Bastion **SSH** (ssh.plafrim.fr, formation.plafrim.fr, ...)
 - And more...

The hidden side

- Storage :
 - **Luster** Parallel File System :
 - 2 MDS servers (in manual failover mode with only 1 MDT)
 - 2 OSS servers (in manual failover mode with 2 OST mounted on each)
 - 2 NetApp which serves **NFS** for /home, /projets and modules. 1 located at Inria and 1 at IMB
- Network :
 - **Compute** network : 4 switchs **10Gbit/s** Ethernet
 - **Management** network : 5 switchs **1Gbit/s** Ethernet
 - **Omnipath** Network : 4 switchs 100Gbit/s
 - **Infiniband** : 7 switchs 40Gbit/s

02

Nodes description

→ <https://www.plafrim.fr/en/the-platform/hardware-documentation/>


miriel[001-088] : 88 standards computational nodes

- **2 x 12 cores Haswell** Intel® Xeon® E5-2680 v3 @ 2,5 GHz
- **128GB** RAM (2133 MHz)
- **Infiniband** QDR TrueScale: 40Gbit/s for **miriel[001-088]**
- **OmniPath** 100Gbit/s for **miriel[001-043]**
- Ethernet : 10Gbit/s
- /tmp : ~300GB




brise : 96 cores node

- **4 x 24 cores Broadwell** Intel(R) Xeon(R) CPU E7-8890 v4 @ 2.20GHz
- **1TB** RAM (1600 MHz)
- Ethernet : 10Gbit/s
- /tmp : ~400GB



souris : 96 cores node

- **12 x 8(16HT) cores Ivy Bridge** Intel(R) Xeon(R) CPU E5-4620 v2 @ 2.60GHz
- **3TB** RAM (1333 MHz)
- Ethernet : 10Gbit/s



kona[01-04] : 4 KNL nodes

- **1 x 64 cores Knights Landing** Intel(R) Xeon Phi(TM) CPU 7230 @ 1.30GHz
- **96GB** RAM (2400 MHz) + **16 Go MCDRAM** (7200 MHz)
- OmniPath : 100Gbit/s
- Ethernet : 1Gbit/s
- /tmp : ~600GB

Node	Memory mode	Cluster mode
kona01	flat	quadrant
kona02	cache	quadrant
kona03	flat	snc-4
kona04	cache	snc-4

sirocco[01-13] : Nvidia GPU accelerator nodes

Node	CPU	Memory	GPU	Network
sirocco[01-05]	2 x 12 cores Haswell Intel(R) Xeon(R) CPU E5- 2680 v3 @ 2.50GHz	128GB RAM (2133 MHz)	4 x Nvidia Tesla K40m	Infiniband QDR Mellanox 10Gbit/s Ethernet
sirocco06	2 x 10 cores Ivy Bridge Intel(R) Xeon(R) CPU E5- 2670 v2 @ 2.50GHz	128GB RAM (1866 MHz)	2 x Nvidia Tesla K40m	Infiniband QDR Mellanox 10Gbit/s Ethernet
sirocco[07-13]	2 x 16 cores Broadwell Intel(R) Xeon(R) CPU E5-2683 v4 @ 2.10GHz	256 GB RAM (2133 MHz)	2 x Nvidia Tesla P100	OmniPath 10Gbit/s Ethernet

visu01 : visualization node

- **2 x 10 cores Haswell** Intel(R) Xeon(R) CPU E5-2650 v3 @ 2.30GHz
- **128GB** RAM (2133 MHz)
- 2 x **Nvidia Quadro K4000**
- Ethernet : 10Gbit/s
- /tmp : ~1,8TB

→ <https://www.plafrim.fr/en/post-processing3d-visualization/>


03

Networks

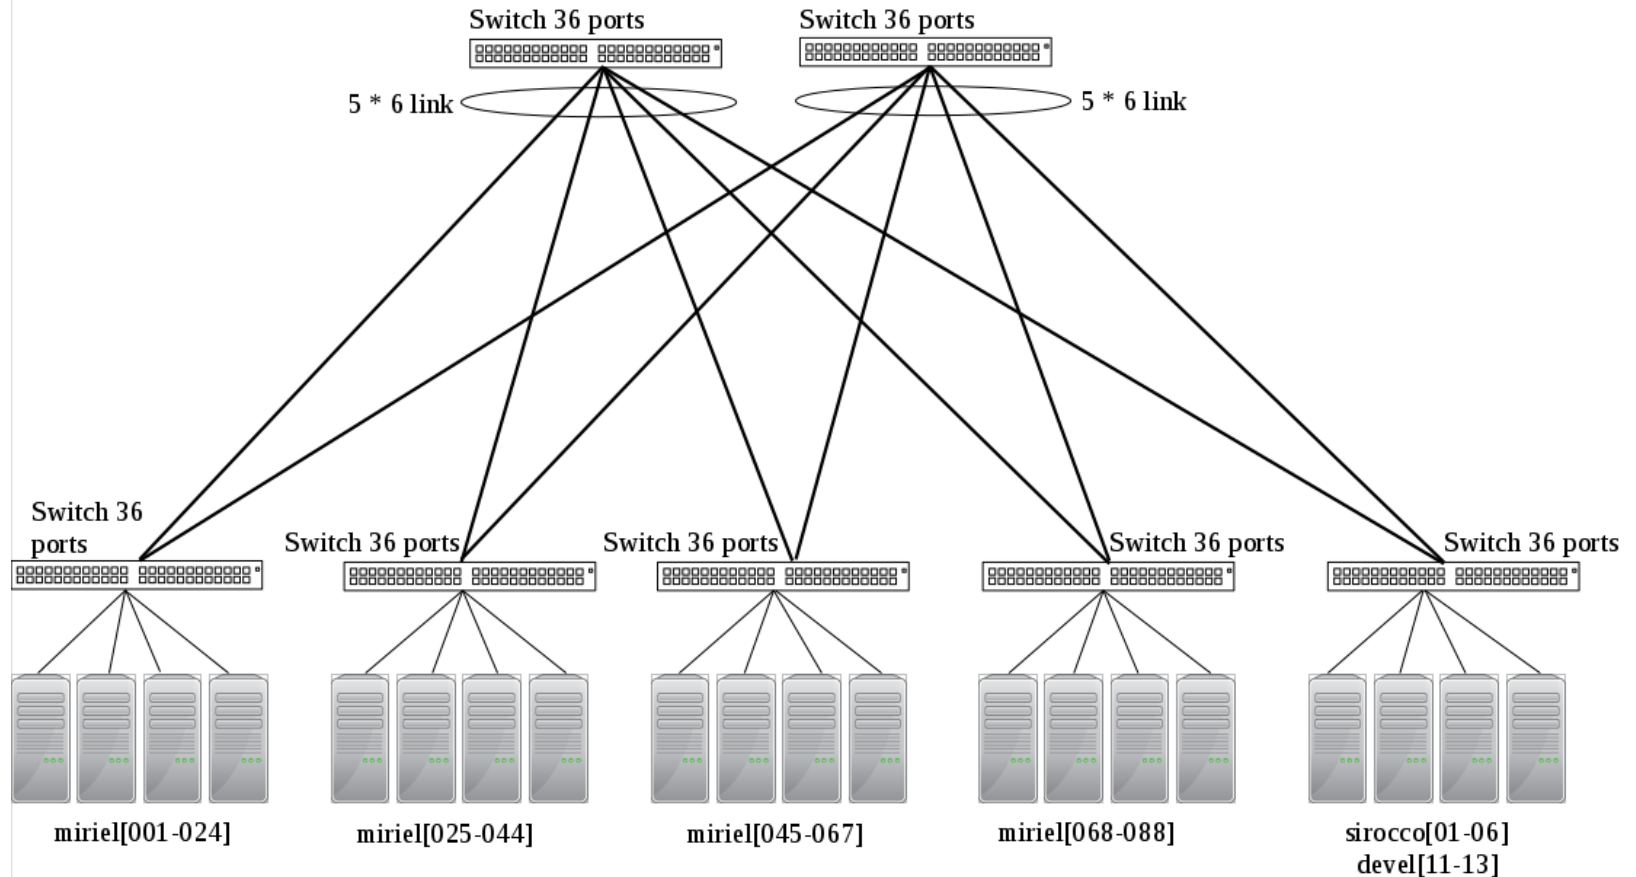
Multiples networks

- **Omnipath** 100Gbit/s network
- **Infiniband** QDR 40Gbit/s network : **Qlogic** and **Mellanox**
- **Ethernet** 10Gbit/s network

Omnipath Network



Infiniband Network



04

Storage

→ <https://www.plafrim.fr/en/the-platform/faq-en/>

Storage available

Name	Max size	Deletion	Hardware Protection	Backup	Use	How to obtain
/home	20GB / user	never	ON	ON	individual	automatic
/projets	200GB / user	never	ON	ON	group	On demand
/lustre	1TB / user 400 000 files / user	If needed	ON	OFF	individual	automatic
/tmp	Variable	At reboot and if needed	OFF	OFF	individual	automatic

05

Evolutions


- Before the end of 2018 :
 - 1 **ARM** node with **2 x 28 cores Cavium ThunderX2(R)** CPU CN9975 v2.1 @ 2.0GHz and **256GB** memory.
 - 4 nodes with **2 x 16 cores Skylake** Intel Xeon Gold 6142 2,6 GHz, **2x Nvidia V100** GPU and **NVMe** and **384GB** memory
 - 1 nodes with **2 x 20 cores Skylake** Intel Xeon Gold 6148 2,4 GHz, **2x Nvidia V100** GPU and **1TB** memory

- Second half of 2019 :
 - Renewal of computational nodes cluster

06

How to access PlaFRIM ?

→ <https://www.plafrim.fr/en/the-platform/faq-en/>



1. if you don't have one yet, create your ssh key pair with ssh-keygen
2. ask your account on : <https://www.plafrim.fr/en/connection/registration/>
3. Your ssh client must use a "ProxyCommand" to reach the target server

Sample configuration of **.ssh/config** to reach plafrim on port 22 :
(replace LOGIN_PLAFRIM with your actual login)

Host plafrim

User LOGIN_PLAFRIM

ForwardAgent yes

ForwardX11 yes

ProxyCommand ssh -A -l LOGIN_PLAFRIM ssh.plafrim.fr -W plafrim:22

Check that your private key is loaded with ssh-add -l. If not, load it with ssh-add ~/.ssh/private_key

Then use **ssh LOGIN_PLAFRIM@plafrim**

07

Need Help ?

If you need Help

1. The FAQ : <https://www.plafrim.fr/en/the-platform/faq-en/>
2. PlaFRIM support : plafrim-support@inria.fr


For technical problems (access, account, administration, modules ...)

3. Users community : plafrim-users@inria.fr

If you need help about more specific use or want to share with others users platform :
contact plafrim-users@inria.fr or view archives here
<https://sympa.inria.fr/sympa/arc/plafrim-users>

Some tools

- Performance monitoring : <https://plafrim.fr/ganglia>
- Jobs monitoring : <https://www.plafrim.fr/en/the-platform/jobs-monitoring/>



► Refresh

REFERENCES: FREE USED DRAIN UNKNOWN

brise					
kona01	kona02	kona03	kona04		
miriel001	miriel002	miriel003	miriel004	miriel005	miriel006
miriel007	miriel008	miriel009	miriel010	miriel011	miriel012
miriel013	miriel014	miriel015	miriel016	miriel017	miriel018
miriel019	miriel020	miriel021	miriel022	miriel023	miriel024
miriel025	miriel026	miriel027	miriel028	miriel029	miriel030
miriel031	miriel032	miriel033	miriel034	miriel035	miriel036
miriel037	miriel038	miriel039	miriel040	miriel041	miriel042
miriel043	miriel044	miriel045	miriel046	miriel047	miriel048
miriel049	miriel050	miriel051	miriel052	miriel053	miriel054
miriel055	miriel056	miriel057	miriel058	miriel059	miriel060
miriel061	miriel062	miriel063	miriel064	miriel065	miriel066
miriel067	miriel068	miriel069	miriel070	miriel071	miriel072
miriel073	miriel074	miriel075	miriel076	miriel077	miriel078
miriel079	miriel080	miriel081	miriel082	miriel083	miriel084
miriel085	miriel086	miriel087	miriel088		

Job ID	Account	Nodes
1011173	mdp-imb	Unknown
1011172	mdp-imb	miriel039
1011165	tadaam	miriel[045-055,075-079]
1011144	logis-br	miriel[001-003,082-086]
1011128	camus	Unknown
1011127	camus	Unknown
1011095	camus	sirocco02
1011094	camus	sirocco01
1011076	flowers	miriel073
1011030	realopt	miriel[031-038]
1011005	mnmemosyne	sirocco09
1010883	memphis	Unknown
1010882	memphis	Unknown
1010670	flowers	miriel008
1010623	memphis	Unknown
1010531	memphis	Unknown

Some tools

- **GUIX** to easy install your packages (→ the last presentation of the day by L. Courtès)
- **Modules** to modify your environment and load/unload programs (→ next presentation by F. Rué)
- **SLURM** to submit jobs on the cluster (→ next presentation by F. Rué)

Thank you!

Follow us on www.plafrim.fr