

What still needs to be fixed



Apptainer

Make that everyone exports (after adaptation of course) : ~~export~~
~~APPTAINER_CACHEDIR=/scratch/gpfs/\$USER/APPTAINER_CACHE~~ ~~export~~
~~APPTAINER_TMPDIR=/tmp~~ ✓ **done** : Created a /data/apptainer/user.name/.apptainer, migrated the current .apptainer users dirs and added a symlink in their home.

1. To prevent quota explosion : Installed the quota package on disco and chacha :. should we enforce quota on FS level ? All the people who said they would use less than 100GB are using more than what they said (ex martin.barry at 350GB) : Applied quota on / filesystem, not yet on datasets.

Content

1. Explain to PA how to create a proper structure.
2. Where do we put the content of this file, as some information is not intended for the general public
3. How can we make animations on the Wiki using JS + SVG and stuff ? Snow ?
4. ~~Create a proper structure, starting with the infrastructures~~ ✓ **done, after discussion merge into docs: for both groups**
 1. With a section for students
 2. With a section for teachers
5. ~~Make a limited-access location with the critical information~~ ✓ **done**
6. Construct the **tools for teacher** sections with the existing informations

Monitoring

1. alerting and messaging on various metrics (disk space, cpu usage, ...) for the various computational resources (~~chacha, disco,~~ **Done** calypso & others)

Server room

1. Make something nice there. Posters on the walls, screens, stuff.
2. Why is there still a box for a server in the networking lab room ? : **Because we need at least one for sending back in case of support / we needed one for network labs to hide the CTF network setup**
3. Rename networking lab room and change the remplaçant for Darko as well
4. Why only 10 GB for the fiber
5. I don't want a patch panel inside the server rack but outside of it. Space will be premium soon there and we don't know where to put the server rack : ~~Search and buy a patch panel to put in the room~~ Done.

6. Find a proper layout for the server room for accommodating a water-cooled rack and maybe another one in a couple of months
7. If we have Rumba running there, we need some UPS solution.
8. Do a drawing schematic of the future rack, notably for having a proper rumba failover policy
9. Choose new R630 and R730 / R740 for RUMBA main. Budget 3 kFr
10. Do we need a file server from the guys downstairs (baignoire)
11. ~~Remove again the big oven : check with Hervé Girard to store it in 23N322 : this is where a student used it last time (but nicely returned it to N307), why not keep it there ?~~ EDIT: the RoL of N322 is Thomas Sterren, I sent him a message for the oven. (Rémi) / EDIT2 : answer is "we share the room so it will stay in 307. period." : need to find a place to store it ourselves. Moved to 23N321 : They want it back to N307, but for no valid reason : so it will stay in N321 where it is less annoying or if it leaves, it just leaves our rooms for good. EDIT 2025-04-03 : The oven will go in 23N111, the RoL is Cedric (Clivaz?), he should contact us to install the oven in its lab.

Slurm on chacha or disco

1. Make both GPUs available in gres/slurmd confs ✓ **done**
2. ~~Make emails working for start/end of jobs, use an emailer~~ ✓ **done**
3. ~~Find how to do the ressource partitioning with billing credits by user / account~~ ✓ **done** (but still needs tests and real jobs to see how to tweak)
4. Discuss how to allocate credits for users : what about students ?
5. ~~Note everywhere to either remove sshfs for VScode, and give links to properly configure it or no VScode at all :~~ Noted on [runjob](#) and started script to check for .vscode in homedirs : auto-rm in crontab directly ? **Done**
6. For the future jump server, need to test how to restrict ssh access to other servers : via SLURM they might recreate their authorized_keys by running a job writing a .ssh/authorized_keys on the server the job is run. (change .ssh/ permission disabling them to chmod this dir?)

Calypso

1. Reinstall slurm by compiling with all necessary plugins, then package using debuild : https://slurm.schedmd.com/quickstart_admin.html#debuild , then deploy the .deb by Ansible

Rumba

1. Turn on Rumba and install a proper env for us, mainly based on docker as a limited number of members will use it
2. Test backup and replicate ISC / Learn on Rumba
3. Migrate the wiki there
4. Migrate ISC / Learn there ? TBD
5. Have VPS and cloud coder there, please.

Hannibal

1. Backup DokuWiki : ✓ **Done already, Hannibal has /srv/www completely backedup on the Synolog NAS DS923**
2. Add Ingegamez website on wordpress

Site

1. Proper CSS for title, also for the alignment which is ugly (look at this page!)
2. Editor with no tabs
3. Why is there a search box with the same text ?
4. Rights done properly for every ISC member

TODOs from Remi's papers

Rumba :

1. Create laptops backup space on the Rumba NAS, make the share visible from the wireguard subnet (Volume2)
2. Re-create the Rsync script for Hannibal on the Rumba Synology NAS : test
3. Proxmox to install on the EPYC / Check Windows VM install / Licensing for VDIs
4. Install and test OpenProject
5. Migrate Marks-Crawler (streamlit) from Disco to Rumba

Dance New :

1. Check spares for EPYC servers / order some discs, fans, power supplies
2. Finish the Mellanox switch IP configuration, to put in the new Sinf subnet 10.5.1.148/24 / GW 10.5.1.1 / DNS 10.130.0.11,10.130.1.11
3. Test and configure BeeGFS on the EPYC 48TB storage
4. Configure the storage infiniband network
5. Create SLURM Test Partition : using shard on Disco, or put the current Rumba Dell 7920 with all 3 Nvidia RTX GPUS, and put the test partition on it
6. Change the creation script to make the "Test" partition the default partition when a researcher arrives on the ISC Compute, then when they are ready to run assign the "Dance" partition
7. Apply Data quota on all ISC compute users, not the case for everyone yet
8. Check the BeeGFS quota mechanism to migrate EXT4 quota from current Disco/Chacha to the new Epyc storage
9. Script a wrapper on Apptainer to check execution context and refuse to run directly bare-bone : same for python or other executable to avoid run out of SLURM
10. Migrate current NVMe data disks from Disco/Chacha to BeeGFS when it will be tested and ready also on EPYC
11. Automate file deletion for Standard (Premium too?) researchers to avoid having scratch partition with old tests files / Set in meeting what TTL we want : 2 weeks standard TTL ? More for Premium ?
12. Rename the Filesystem : datasets → workspace? local_workspace? chacha_workspace ? local_scratch ? / shared → network_workspace ? remote_workspace ? remote_scratch ?
13. Migrate Prometheus from Chacha to the new EPYC server / Add some alerting on common checks, disks, jobs outside of slurm etc...
14. Move NVMe disks ? Disco 7TB to Chacha ? / NVMe 3TB from Calypso storage to Disco ? Format as BeeGFS
15. Check for a Modules installation ? Or Apptainer is already fine ? : Install LMOD to allow dynamic lib loading : where to put the terabytes of libraries for Dance ?

Calypso :

1. Finish the local DNS as forwarder either to 172.30.7.1 or another external DNS (port 53 was already opened by YB) : test that when it is given as the default DNS in the Wireguard config it doesn't defeat the Split-Tunneling purpose by sending all DNS request in the tunnel ? Or if only the DNS is sent through the tunnel, maybe live with the fact
2. Test MAAS to auto-provision Calypso[0-14] slaves servers, calypso12-14 are not yet installed, we can test on it
3. Change CUDA12.2 to 12.6 or 13.0 ?
4. Add the FS switch to the stack, to allow wiring all IDRACs and hosts cables to all Calypso servers
5. Migrate the current 10GB router to 25GB router ? NAS can only use ethernet 10GBx2
6. Set somewhere suitable the UID/GID plan (currently only on users.yml in Ansible playbooks) :
 1. users UIDs : 1000 : ubuntu / 1001 pmudry / 1002 remi / 1003-1100 teachers / 10000-10500 ISC students / 10500-11000 other students
 2. services : 7000-8000 services like prometheus / 8000-9000 custom groups for students / 9000-10000 researchers groups
 3. Maybe either one day integrate into HES Active Directory (no direct control of the groups and user management) or recreate a LDAP server in Rumba ?

Backups :

1. Make a coherent backup architecture : which data has to be kept how long
2. Prod backups from Hannibal to put also on the Synology SSD NAS (Rumba)
3. Teachers laptops backup space to create
4. Rumba services backups
5. Clean Moodle courses to make backups smaller : Did the auto-backups reduction option work ? Did the maximum version kept to (5?) works ?
6. Finish to test the script Rsync with ACL permissions / Extended attributes, then replace the hannibal.sh / marcellus.sh scripts on the Desktop NAS / setup on the SSD Rumba NAS, either directly from source at another moment in the night, or as a copy of the the first backup to avoid more I/Os on the prod ?

19.SS09 :

1. Remove blades from the Submer pod ? Isn't it worse to put this expensive opened hardware on cloths ? When does Deepsquare pick up their servers ?
2. Find something to do with the pod when everything else will be out

Other :

1. 23N218 poster to hide the transparent door to all people in the hall / corridor
2. 23N218 add a fridge for tupperwares / coffee machine

Post Mortem Followup :

1. Monitoring ISC/Learn : where to install ? Rumba ? / Dashboard to create / emails alerting to setup
2. Finish the DRP wiki part : <https://wiki.isc-vs.ch/doku.php?id=administratif:processes:drp>
3. Create a PDF + paper DRP to make sure instruction are available even when nothing works : hannibal, wiki, and HES network down

Playbooks :

1. Finish the ISC Compute playbook

From:

<https://wiki.isc-vs.ch/> - **The ISC wiki**

Permanent link:

<https://wiki.isc-vs.ch/doku.php?id=administratif:todo&rev=1758894821>

Last update: **2025/09/26 13:53**

