

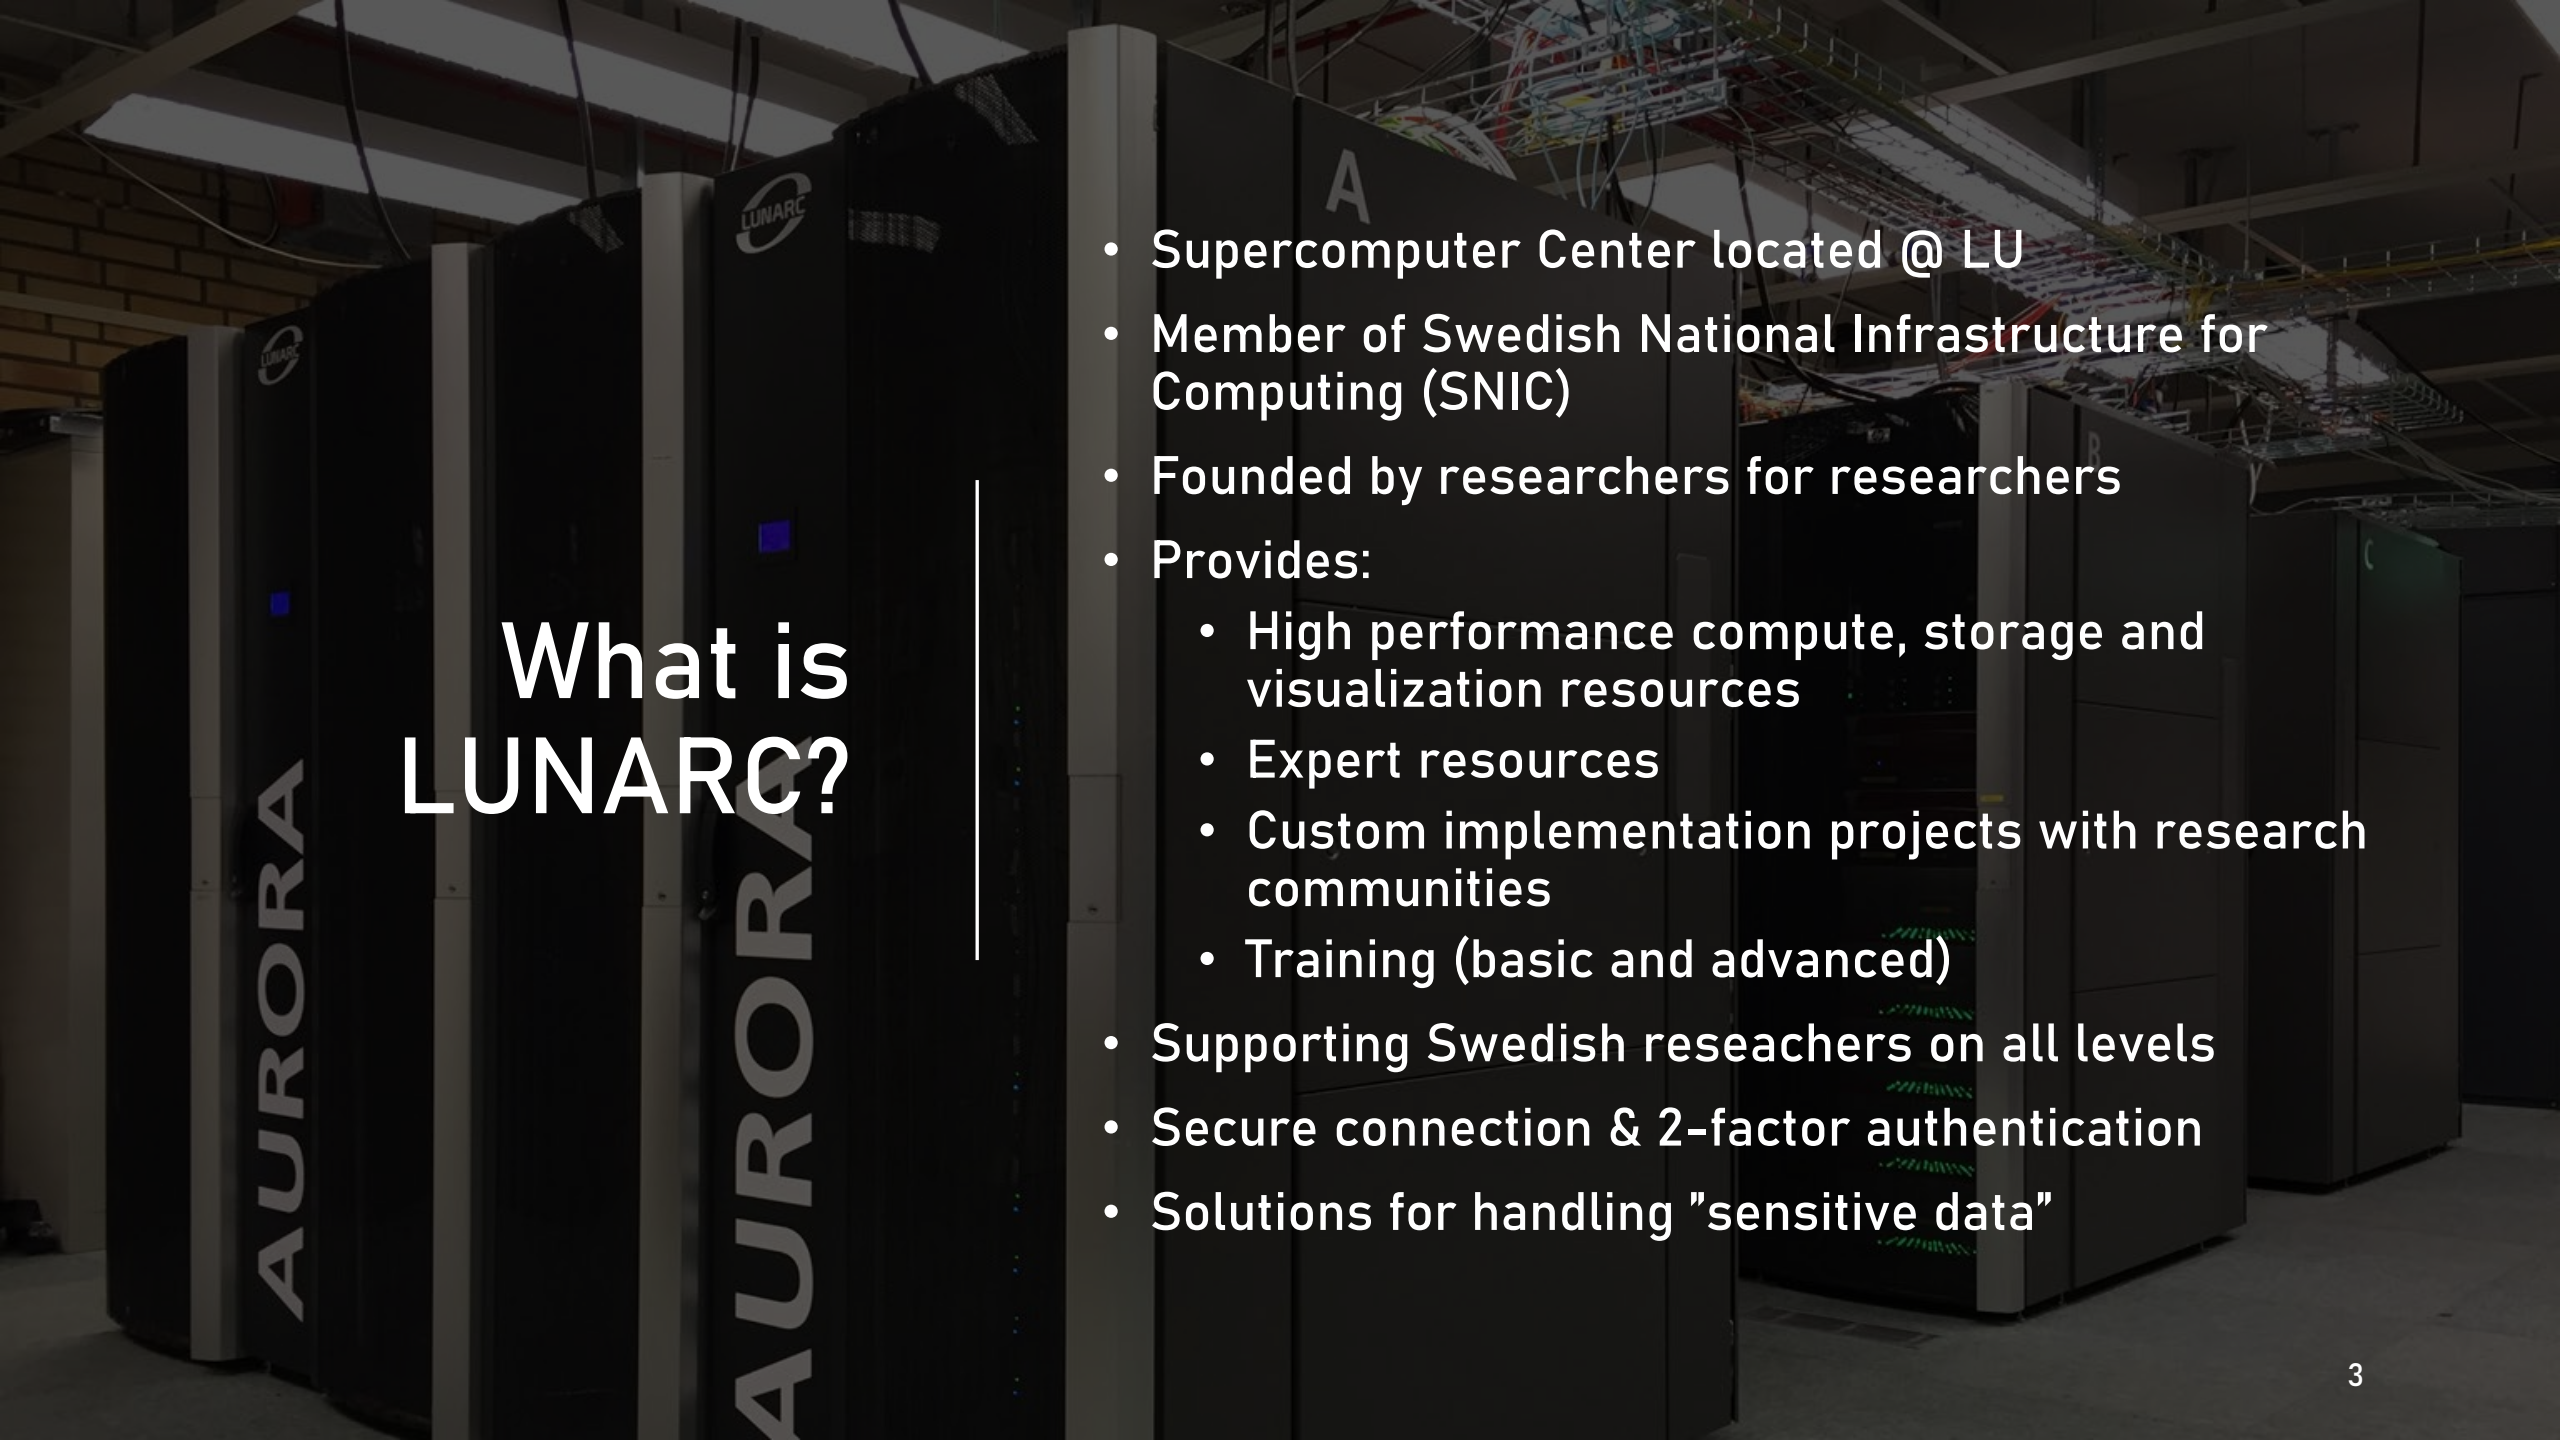


Jonas Lindemann
Anders Follin

Providing a user
friendly and
scalable
desktop
environment for
HPC and
visualisation

Visualisation and interactive HPC at LUNARC

About LUNARC



What is LUNARC?

- Supercomputer Center located @ LU
- Member of Swedish National Infrastructure for Computing (SNIC)
- Founded by researchers for researchers
- Provides:
 - High performance compute, storage and visualization resources
 - Expert resources
 - Custom implementation projects with research communities
 - Training (basic and advanced)
- Supporting Swedish researchers on all levels
- Secure connection & 2-factor authentication
- Solutions for handling "sensitive data"

Remote Visualisation @ LUNARC

LUNARC and Remote Visualisation

Our remote visualisation journey started back in 2010

We early envisioned an increased need for user friendly interactive remote access to HPC systems. The two of us had many years of experience in architecting, developing and implementing graphic sw solutions and rviz environments in both academic and industrial environments.

Two major reasons:

LUNARC and Remote Visualisation

1: Remote (hw accelerated) visualisation.

- Keep large data sets in the datacenter – don't stress network BW & shorten lead times
- Enable parallel compute and remote post processing of huge datasets that won't fit in local WS
- Secure management of sensitive data
- License consolidation
- And many other well known reasons why remote viz is beneficial...

LUNARC and Remote Visualisation

2: Improved usability and workflow

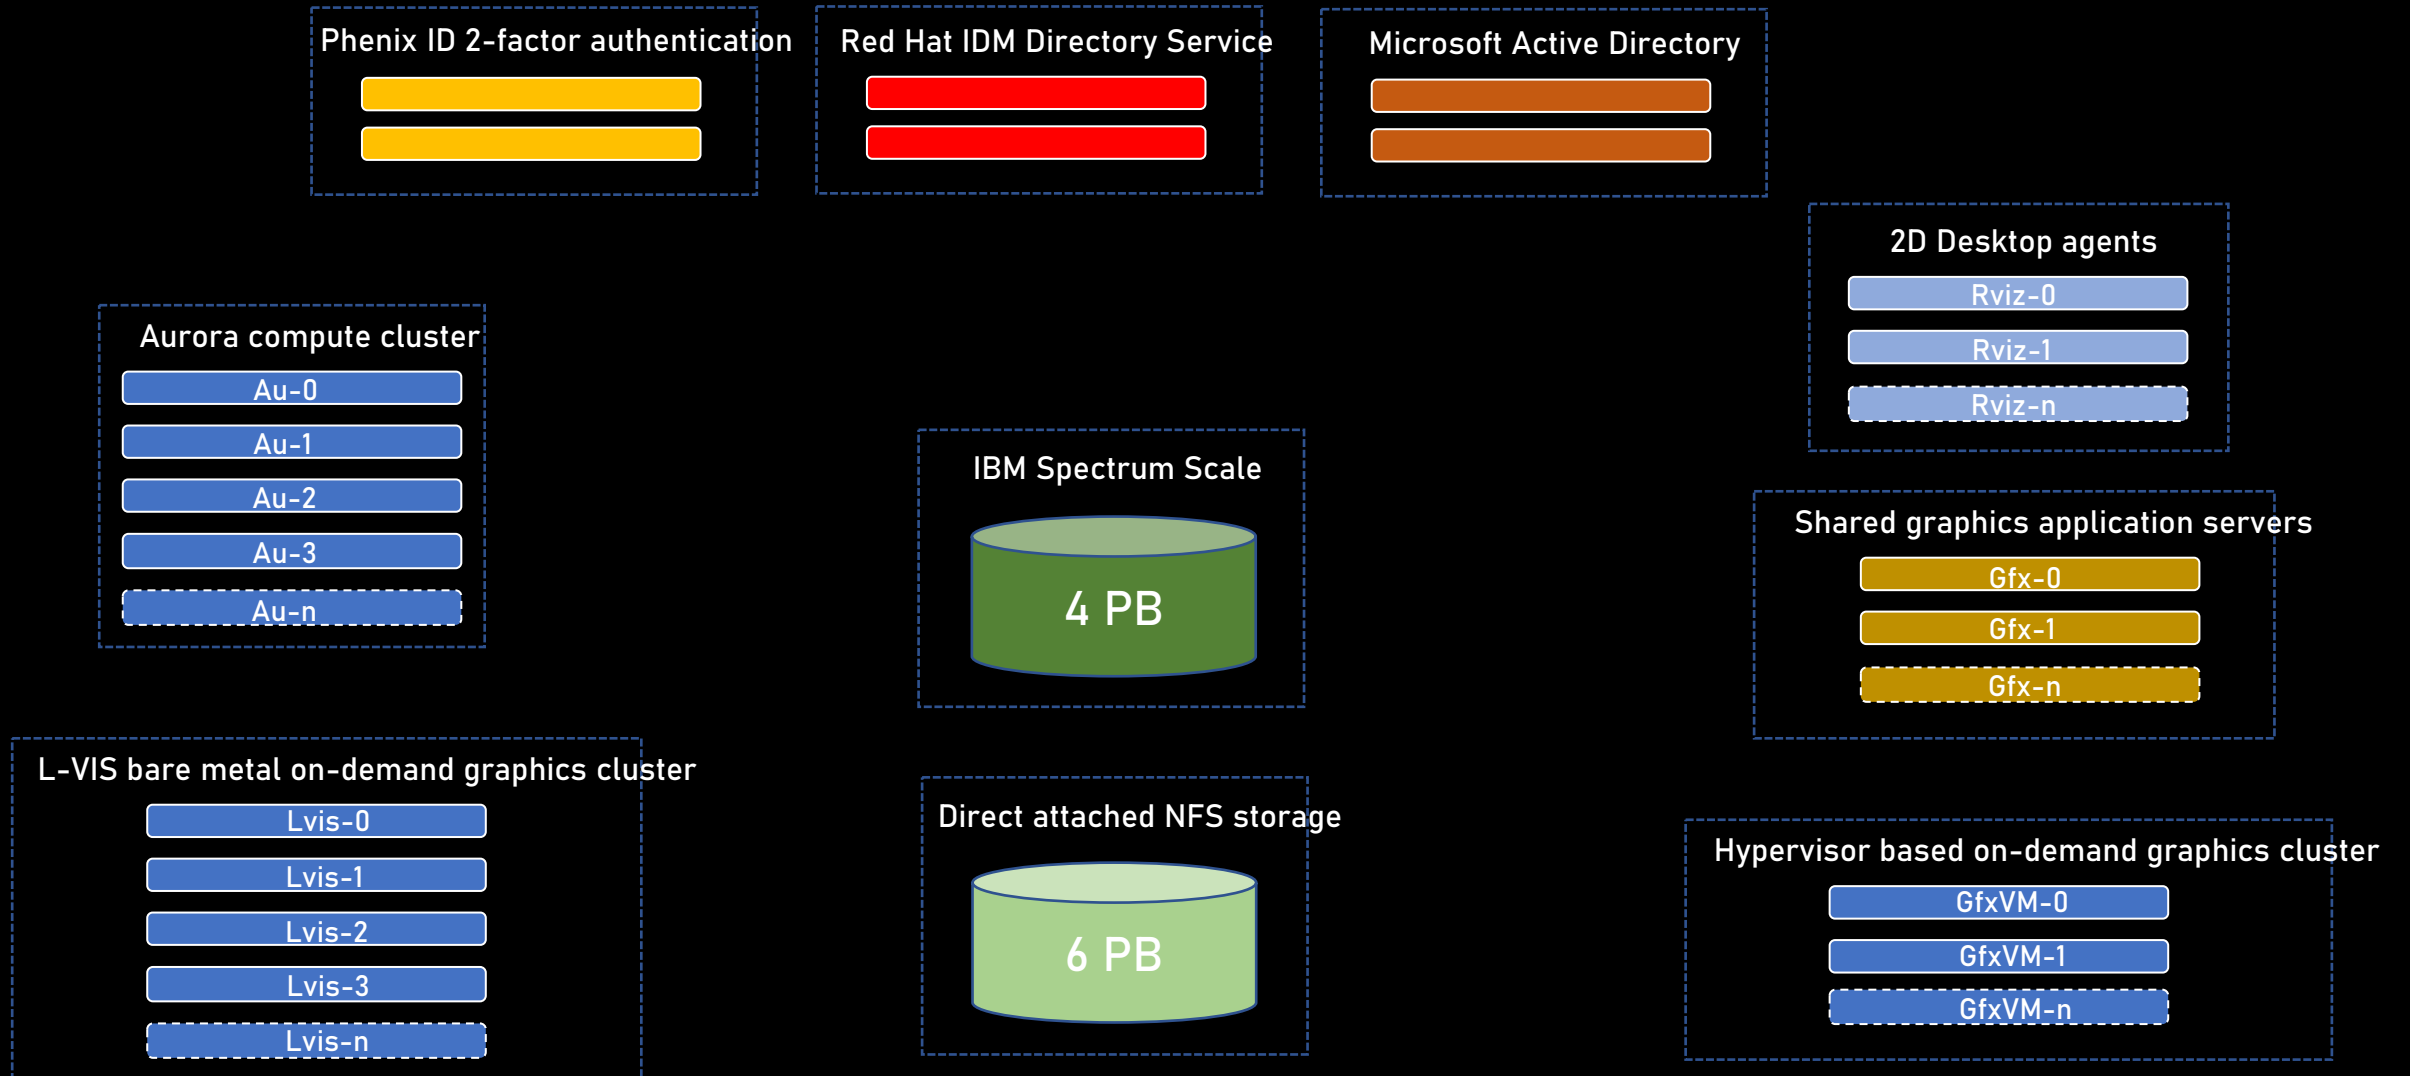
- Adaption of our HPC services to new users familiar with modern desktop environments
- Provide graphical user interfaces to the HPC backend. Both build in DT applications (fm, etc) as well as custom developed tools (module browser, ARC storage explorer, etc)
- Disconnect / reconnect – a massive improvement in usability

LUNARC HPC Desktop 2021

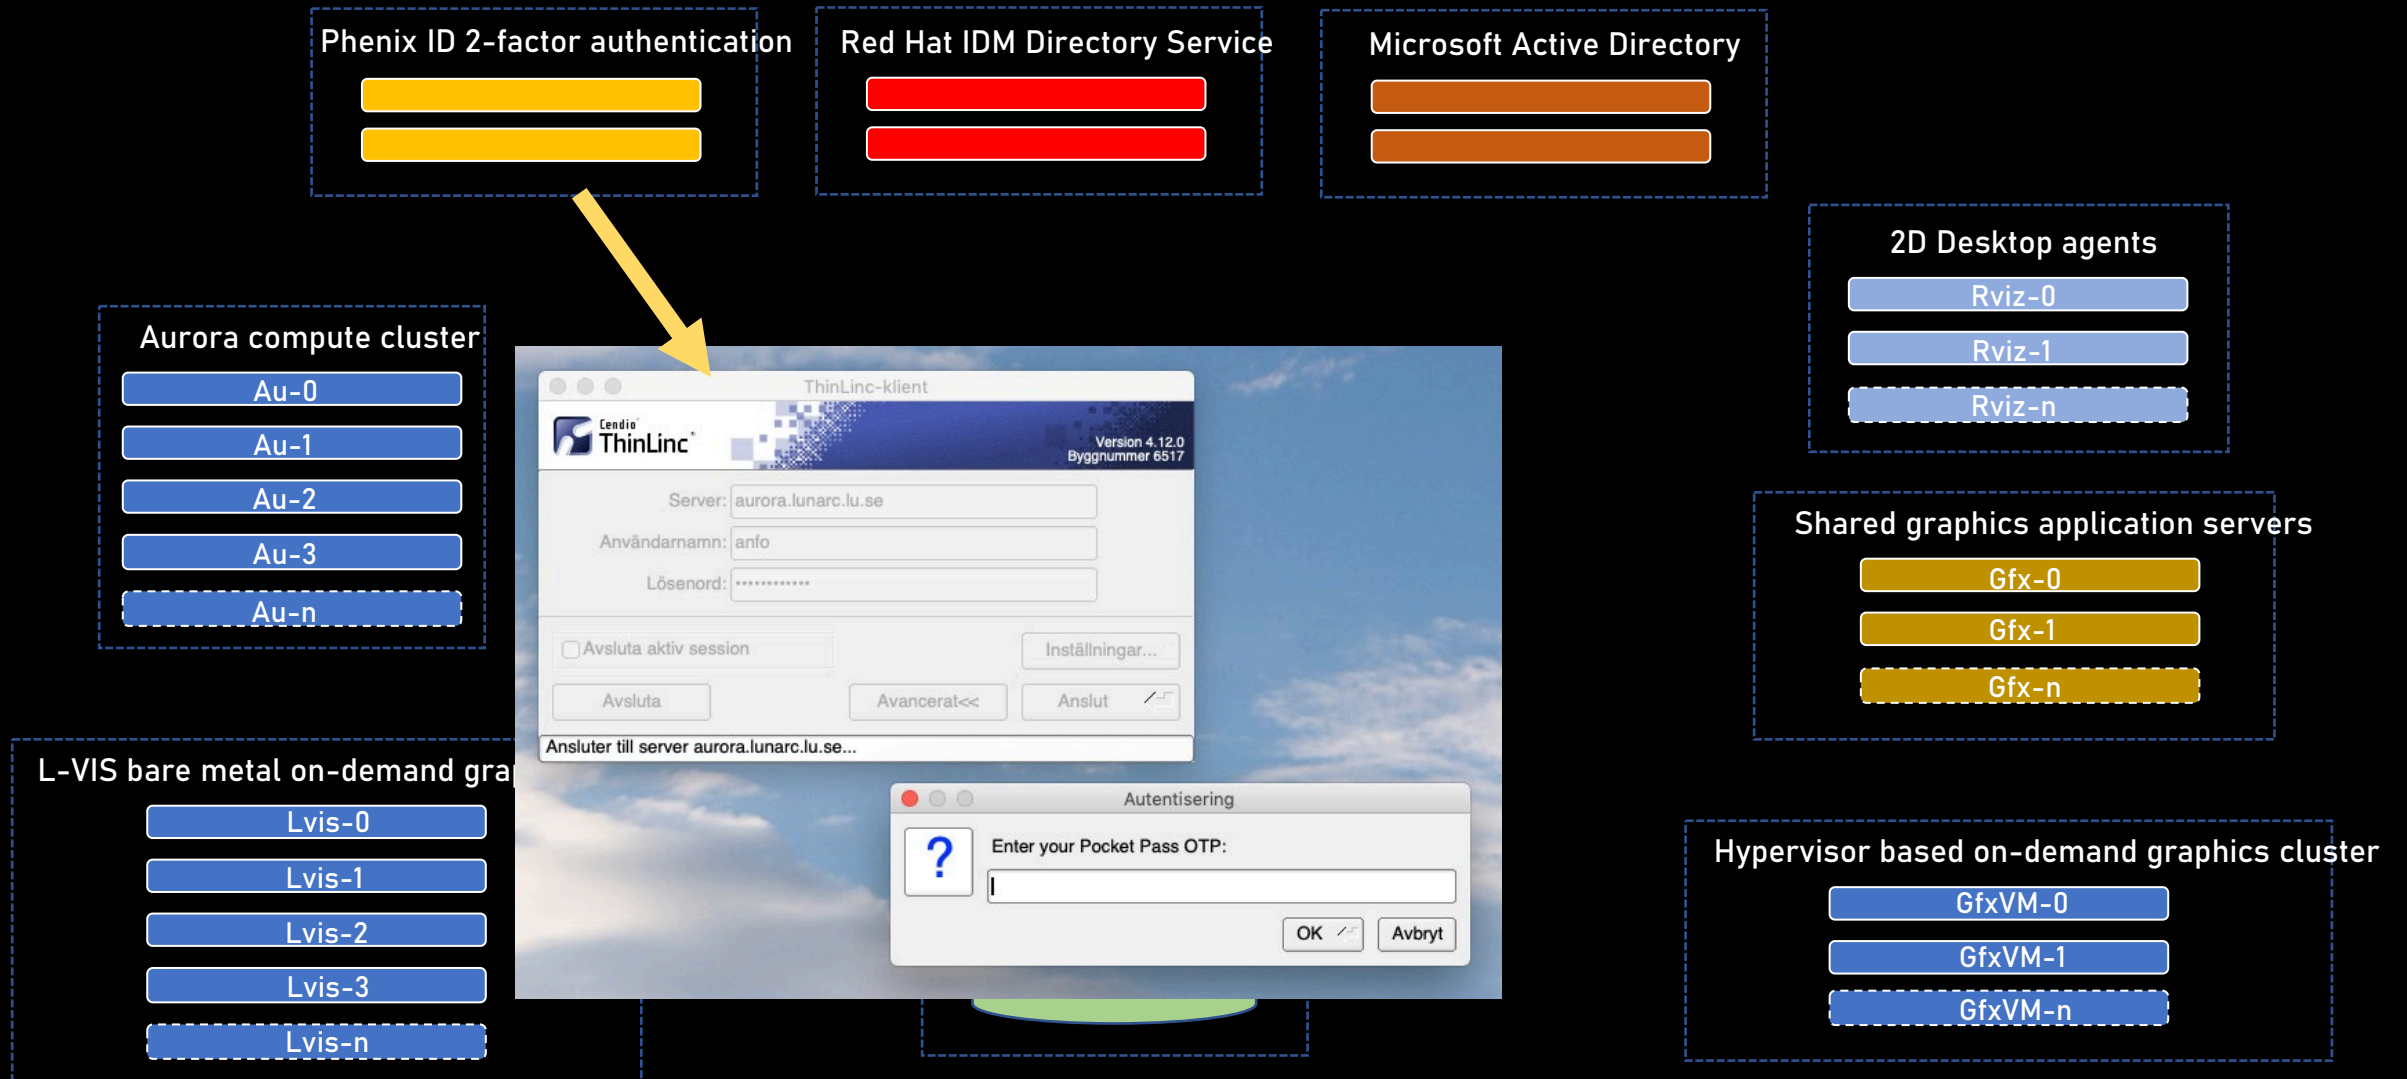
- Remote Desktop solution in production since 2013
- Accessible from all client platforms (Mac, Windows, Linux and iOS)
- Custom architected scalable backend providing both 2D and 3D applications with hardware graphics acceleration
- Custom developed Desktop On-Demand service for requesting dedicated nodes for interactive use managed by queuing system (SLURM)
- Operating system independent solution with hardware accelerated graphics (CentOS, MS Windows, Ubuntu, etc)
- A constantly increasing amount of graphical software available. Both open source and commercial packages installed upon user request

LUNARC High level architecture

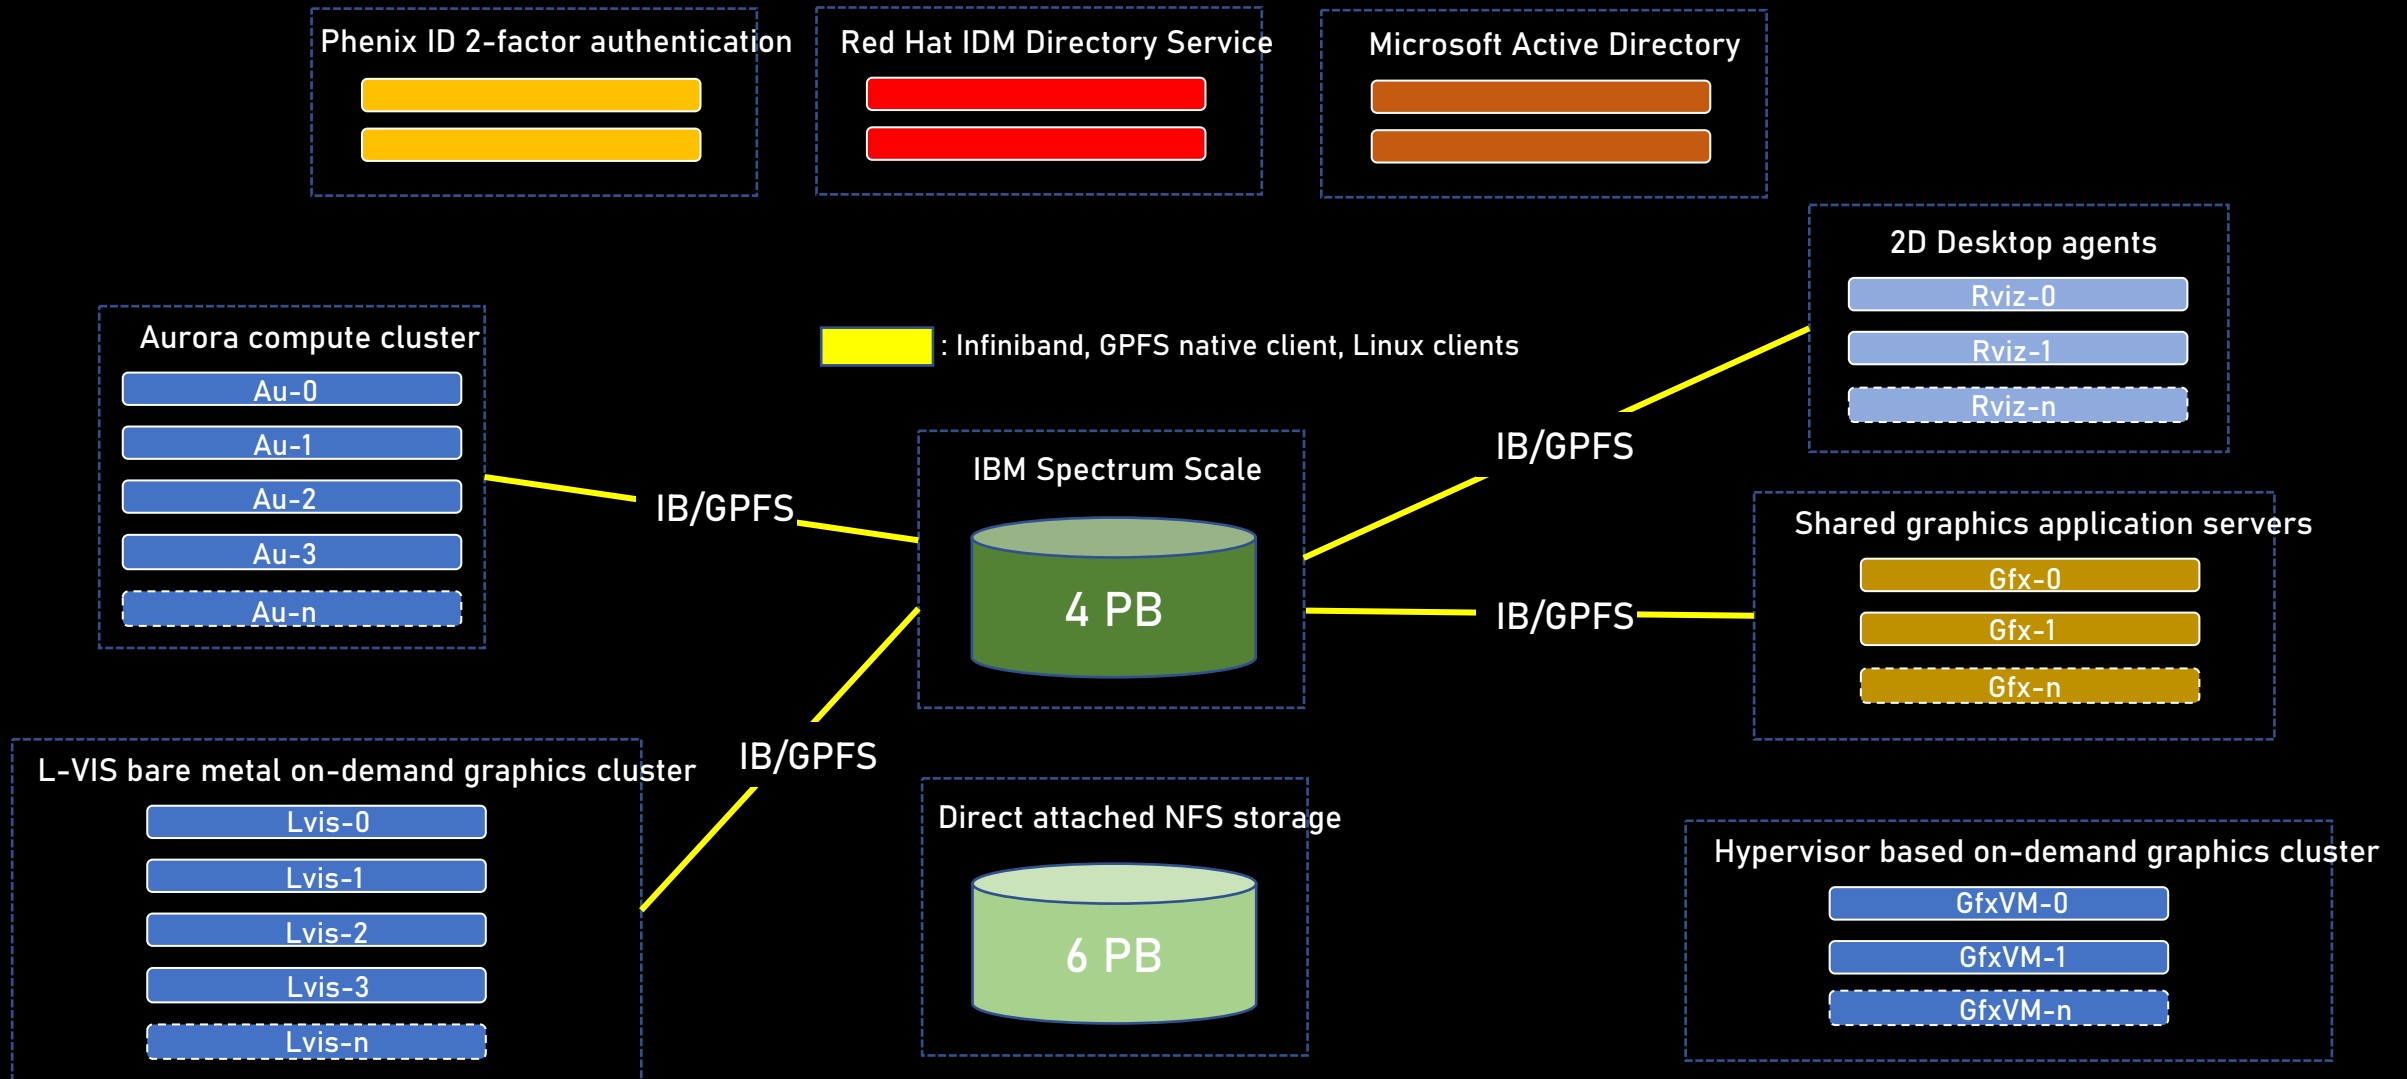
LUNARC Datacenter – High level architecture



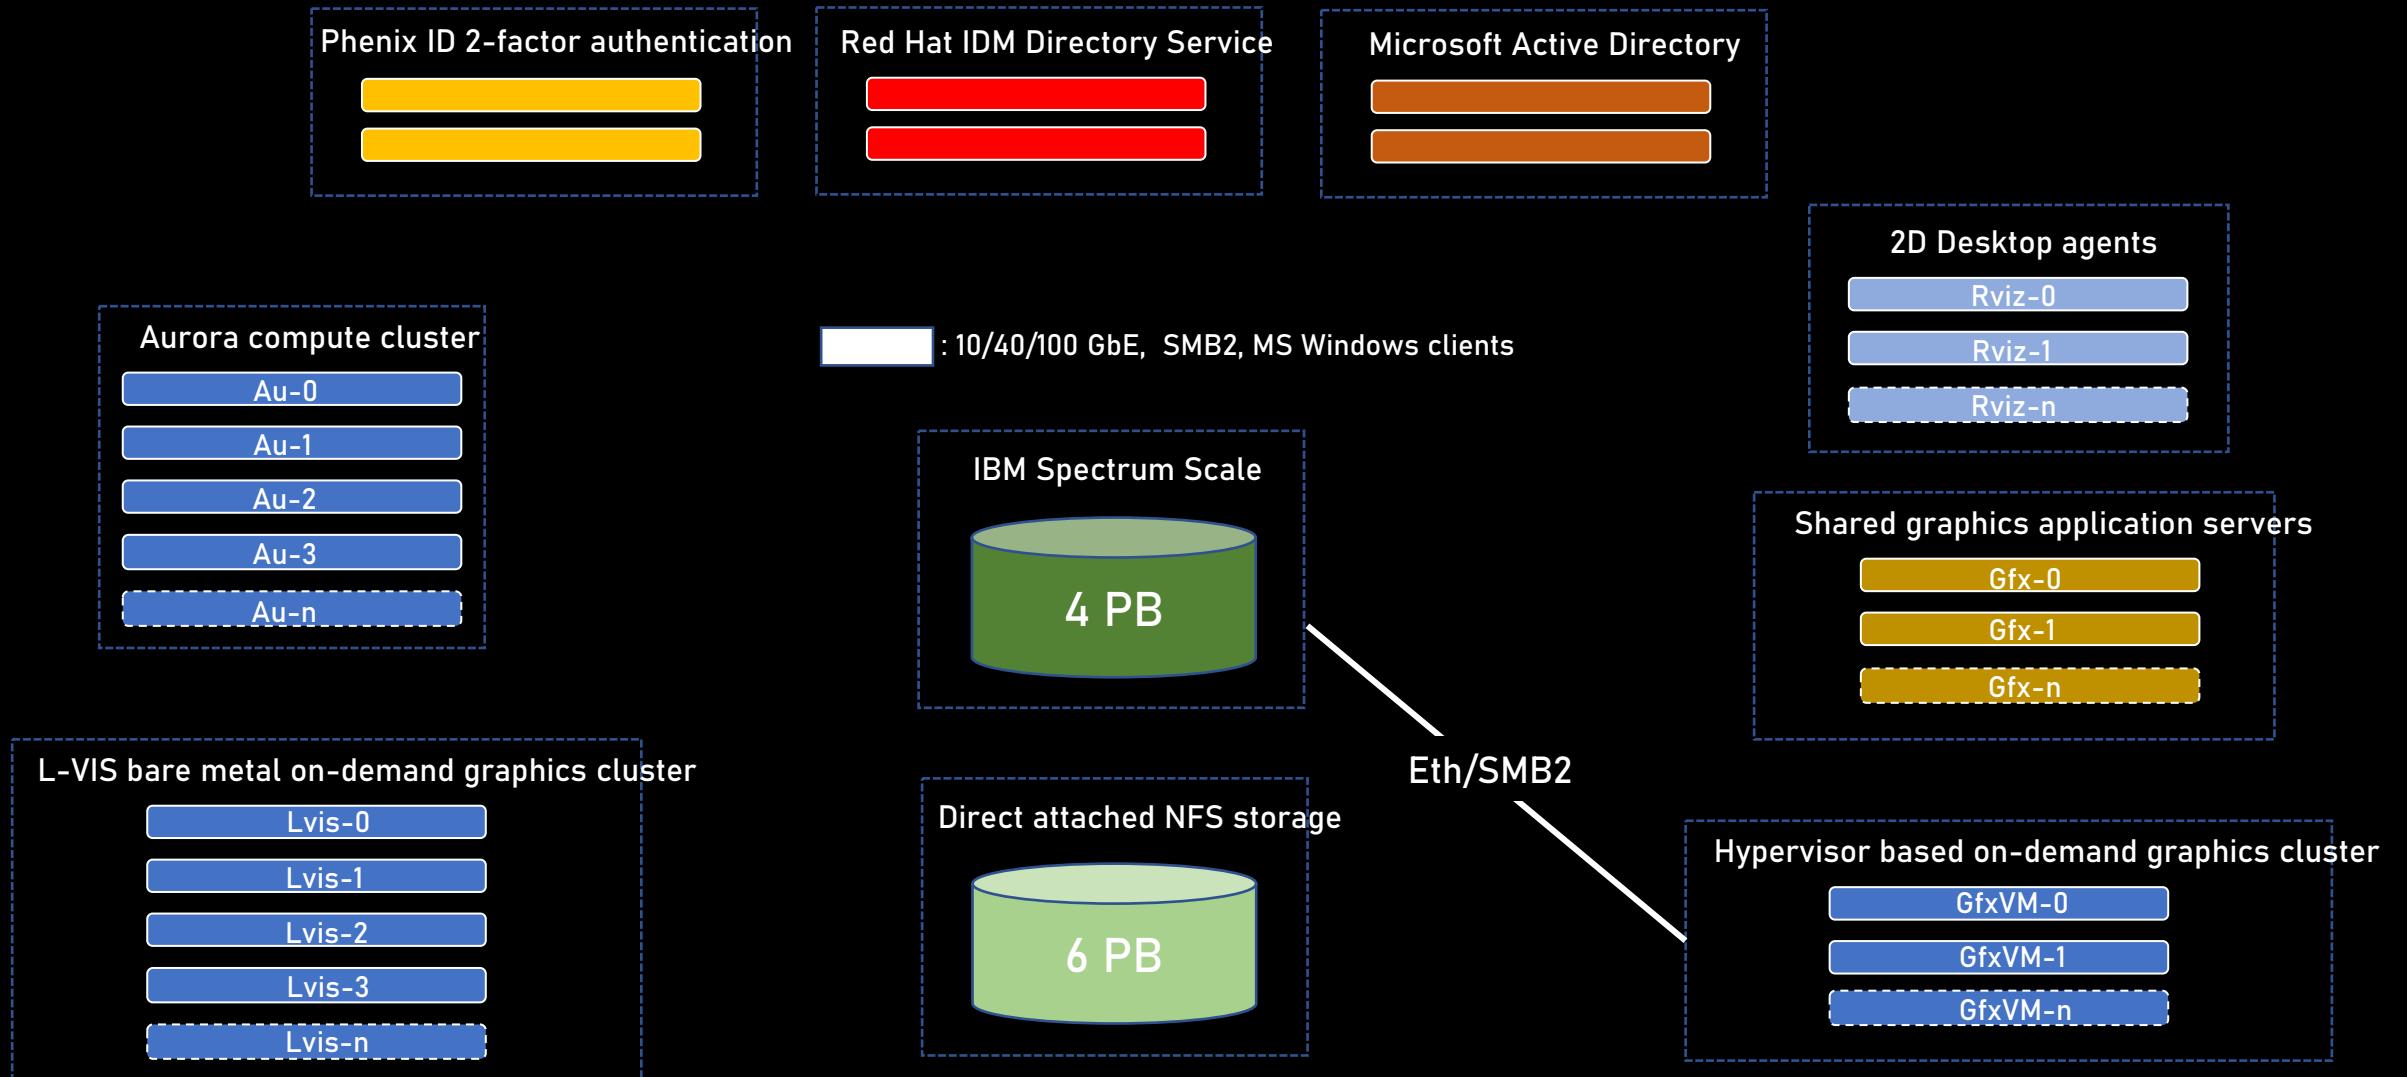
LUNARC Datacenter – High level architecture



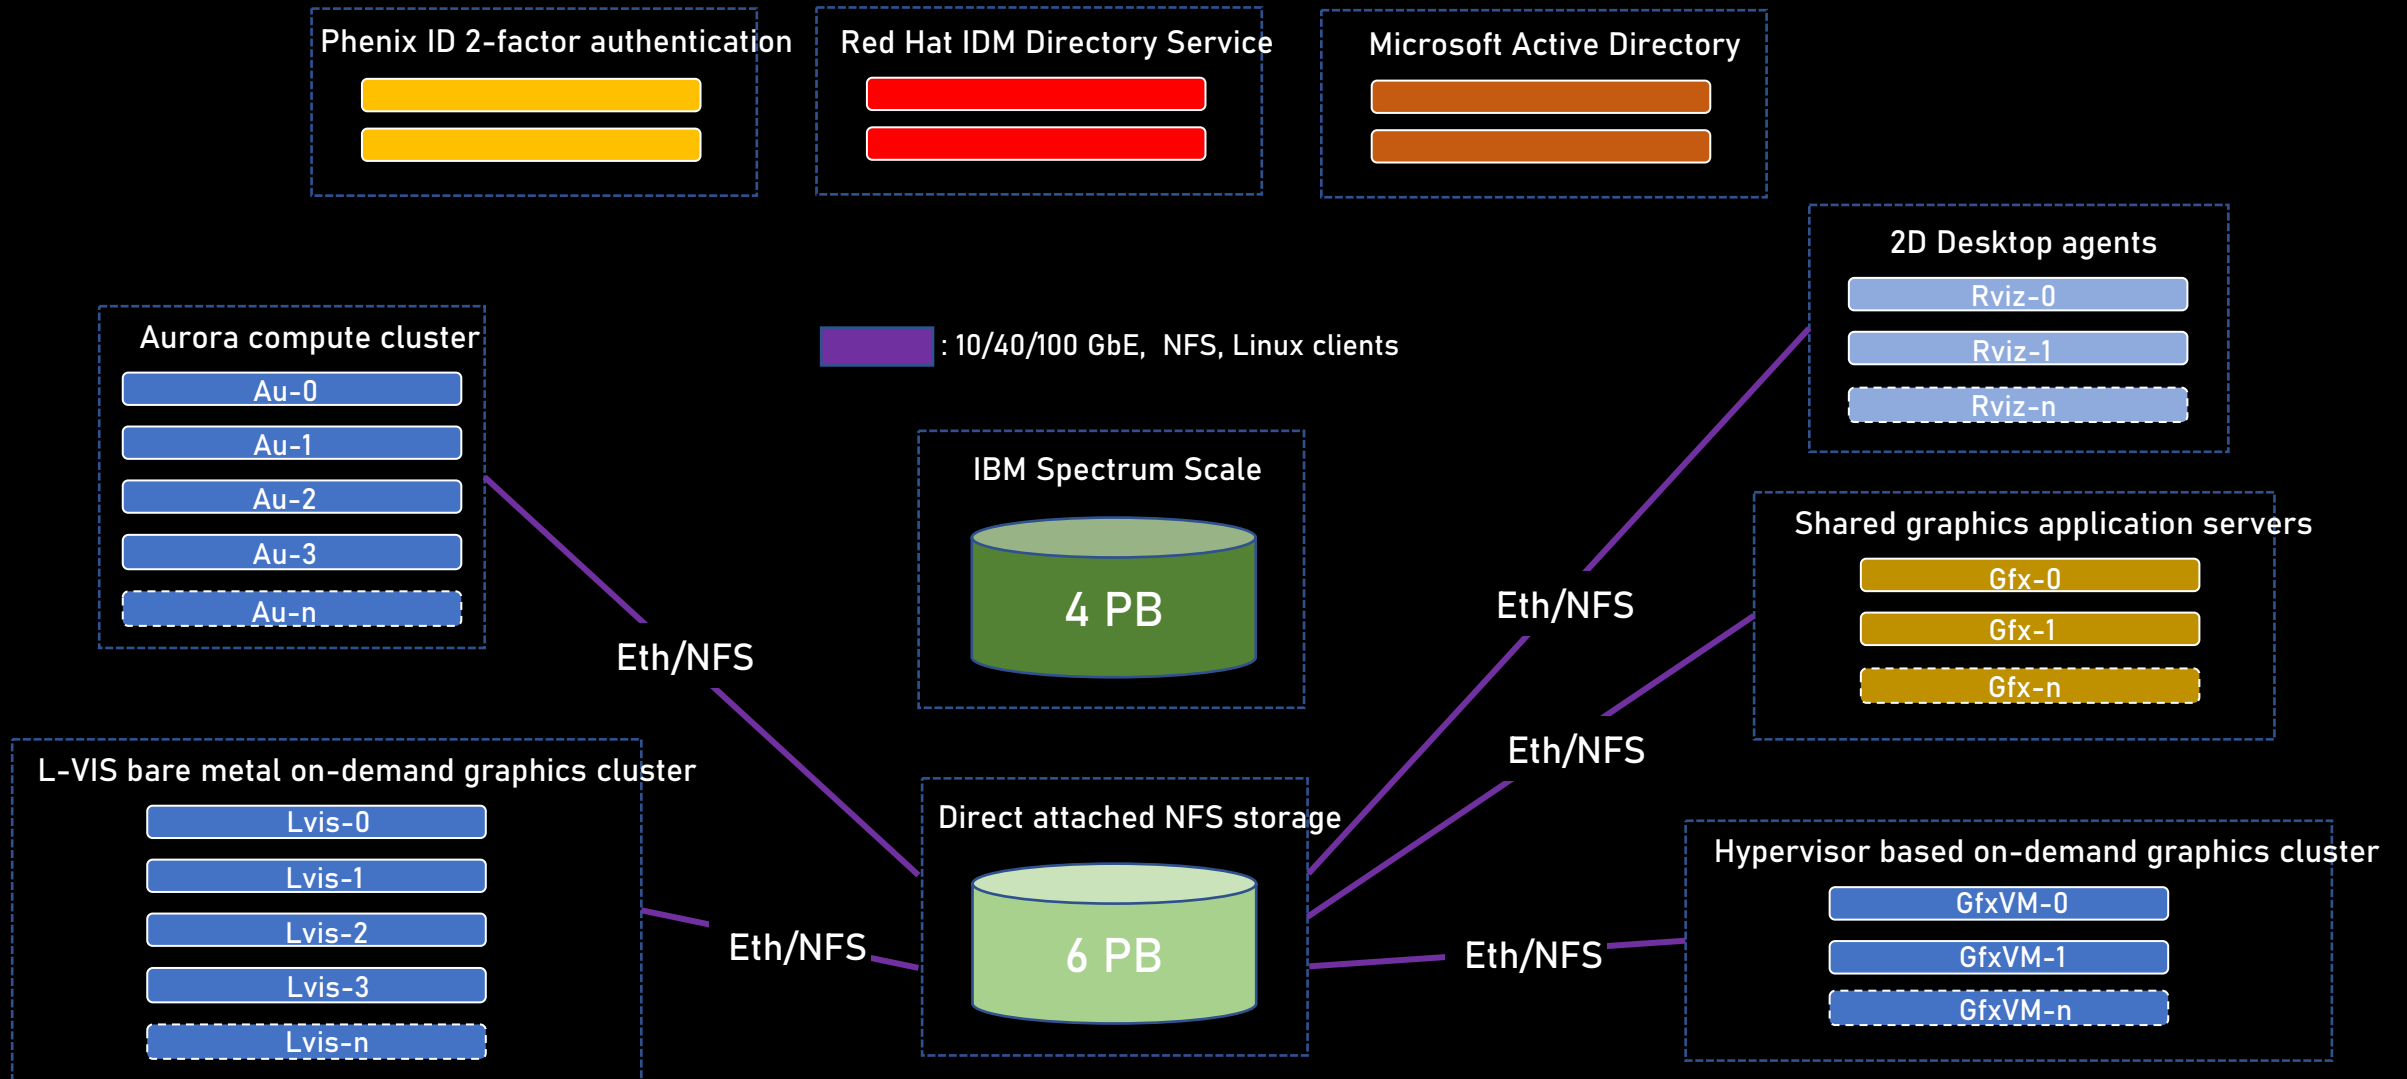
LUNARC Datacenter – High level architecture



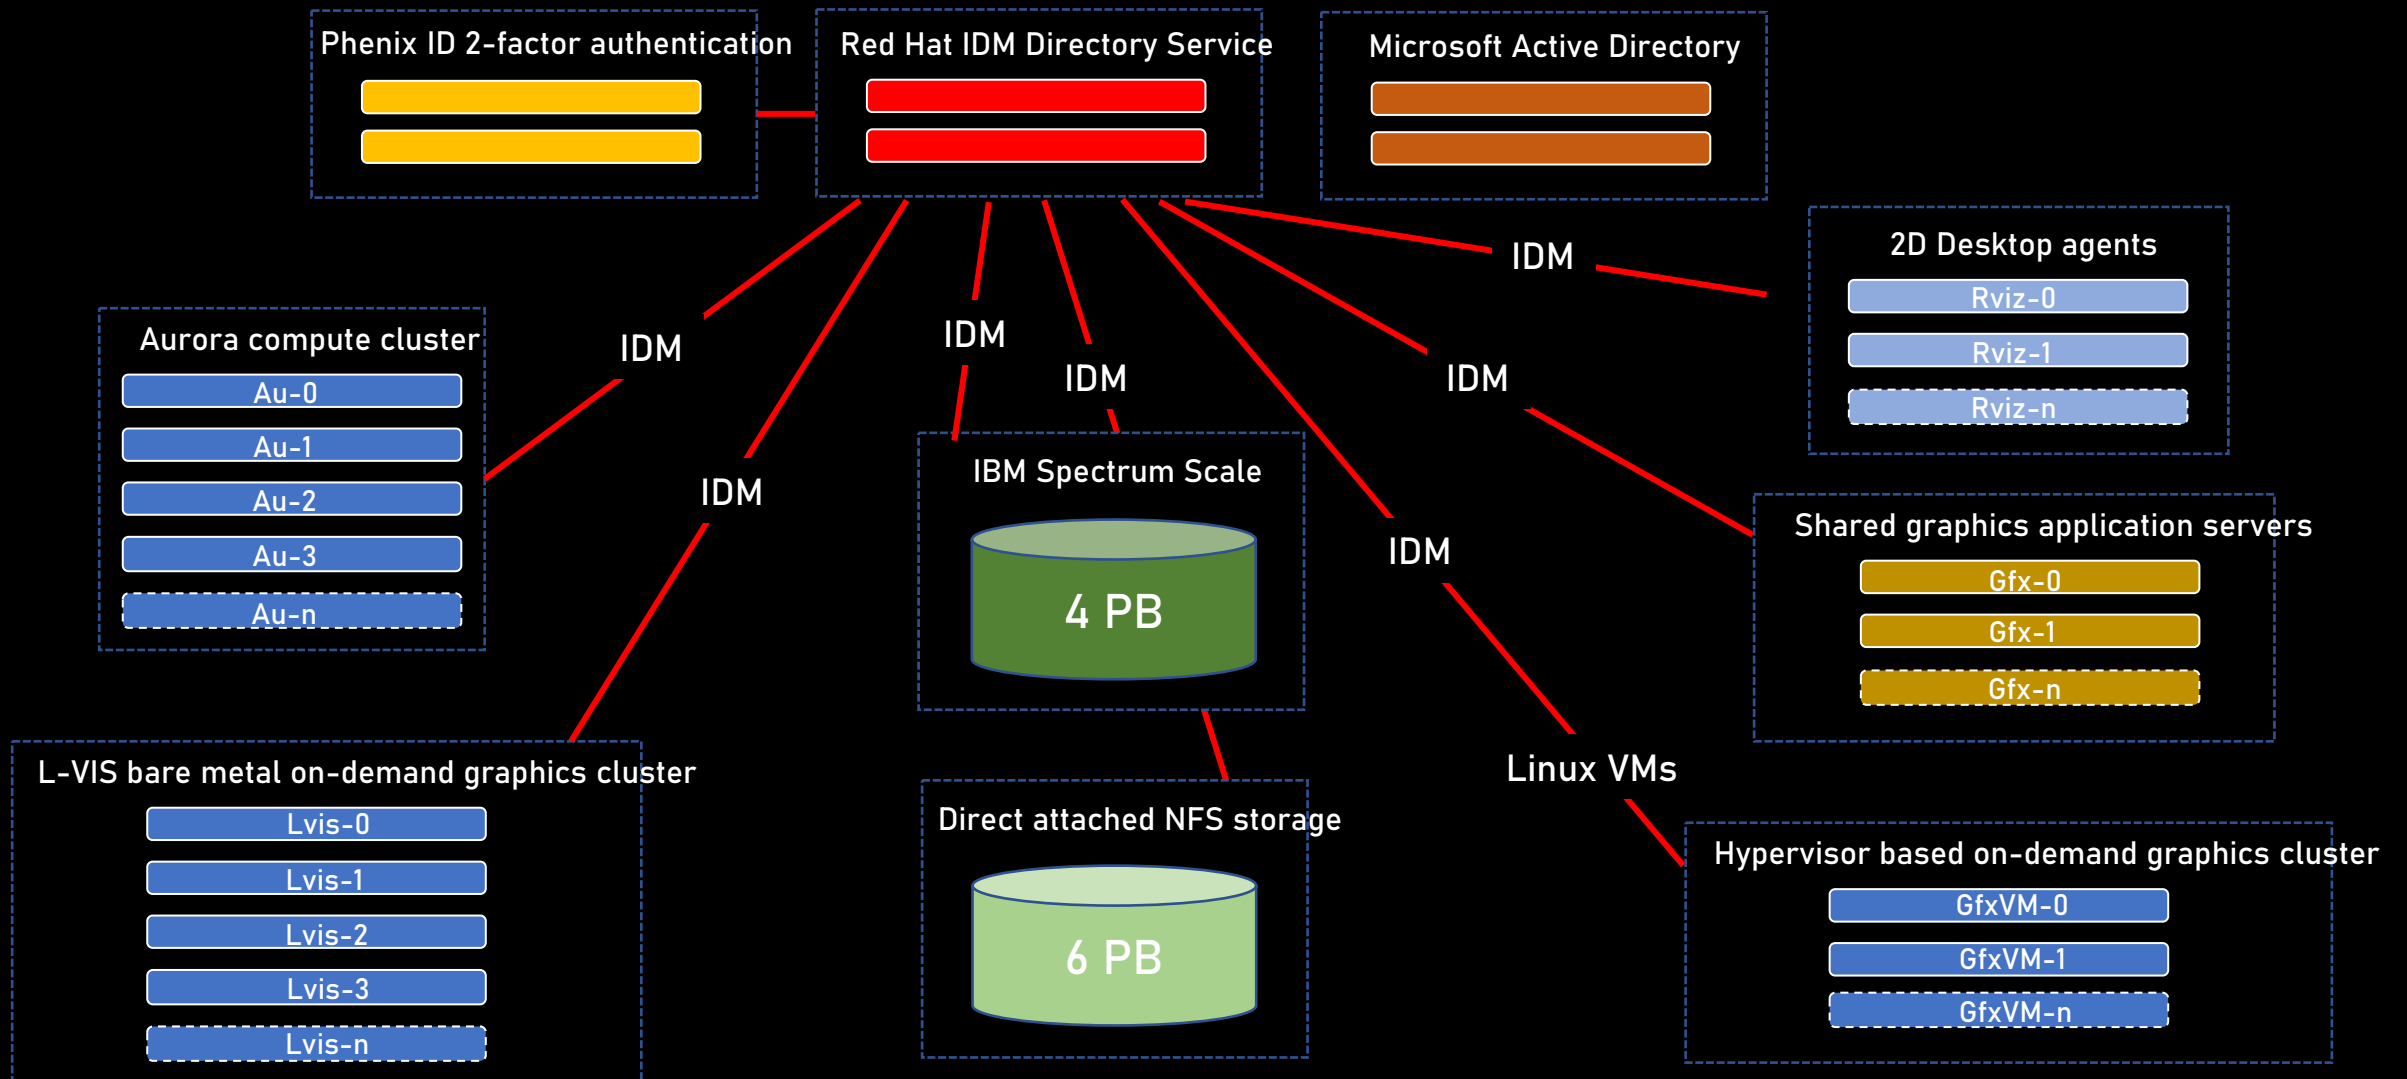
LUNARC Datacenter – High level architecture



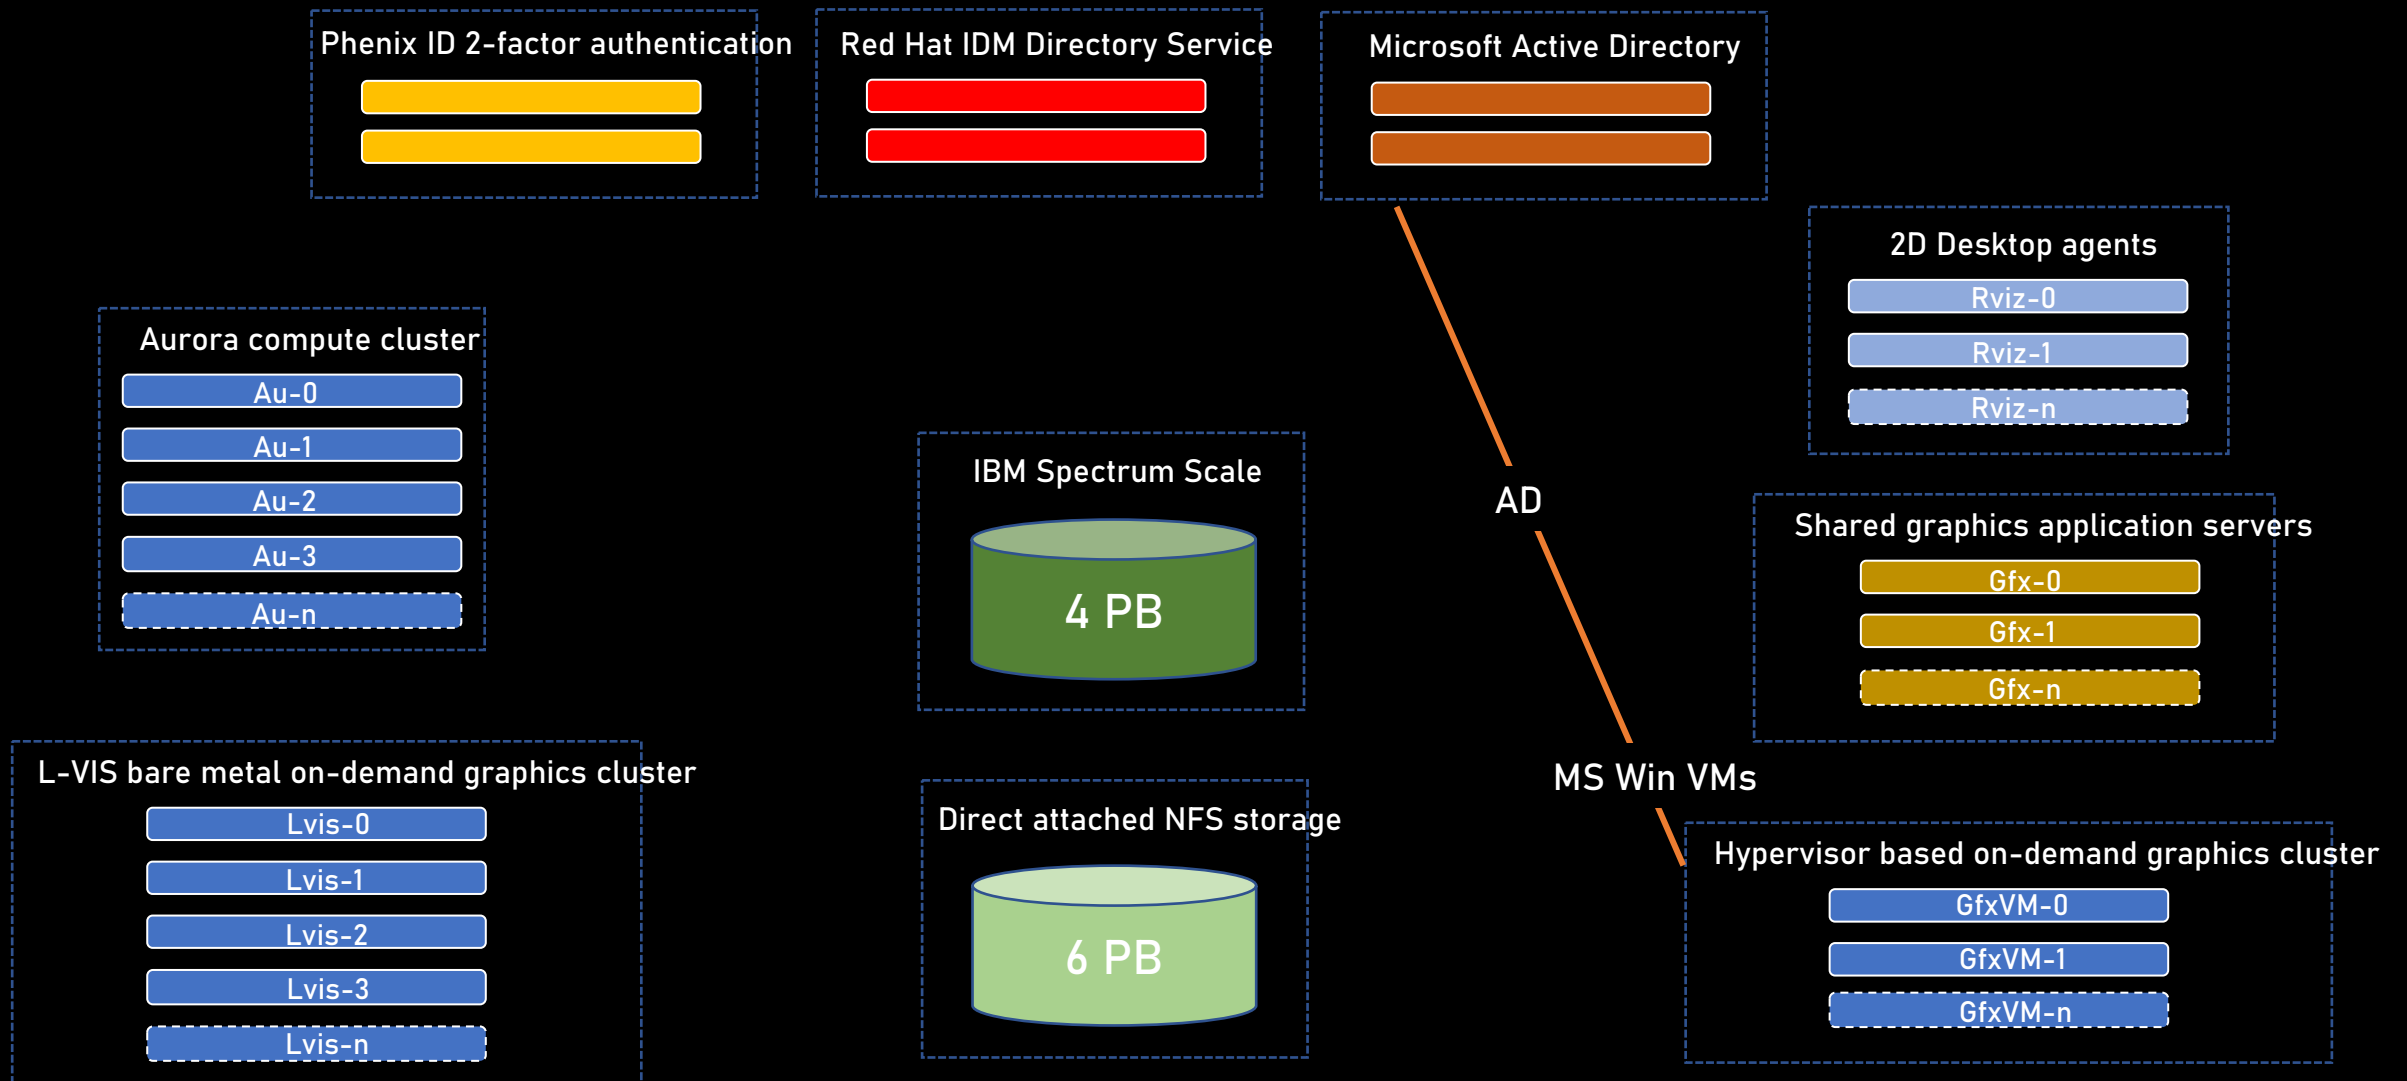
LUNARC Datacenter – High level architecture



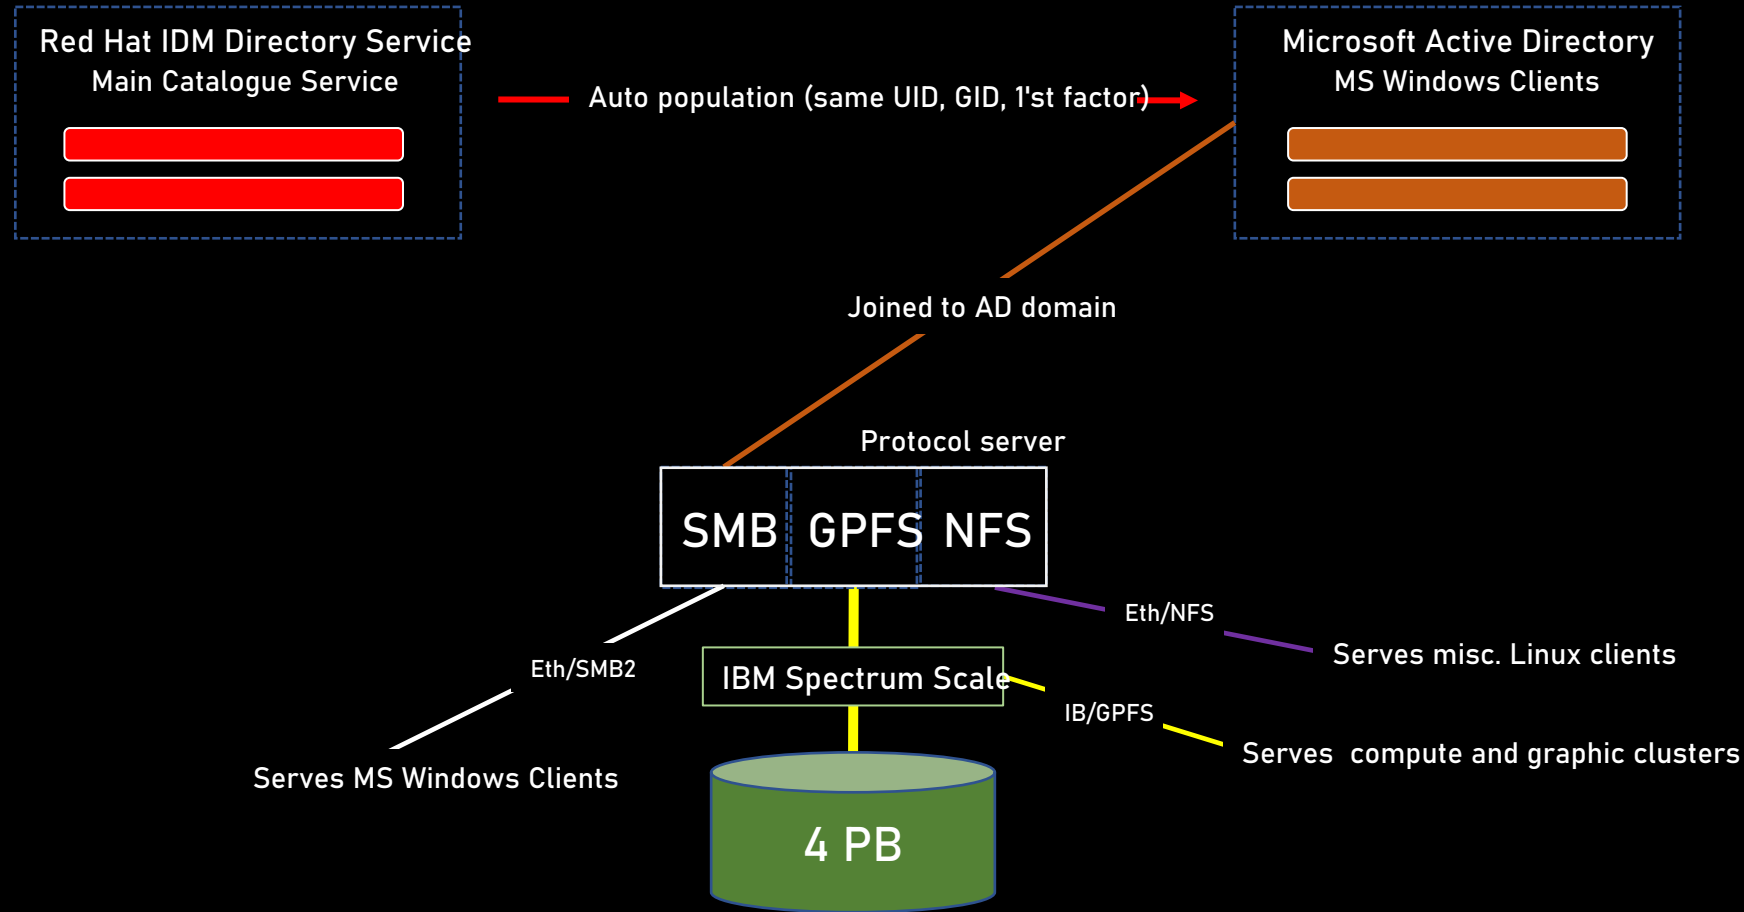
LUNARC Datacenter – High level architecture



LUNARC Datacenter – High level architecture

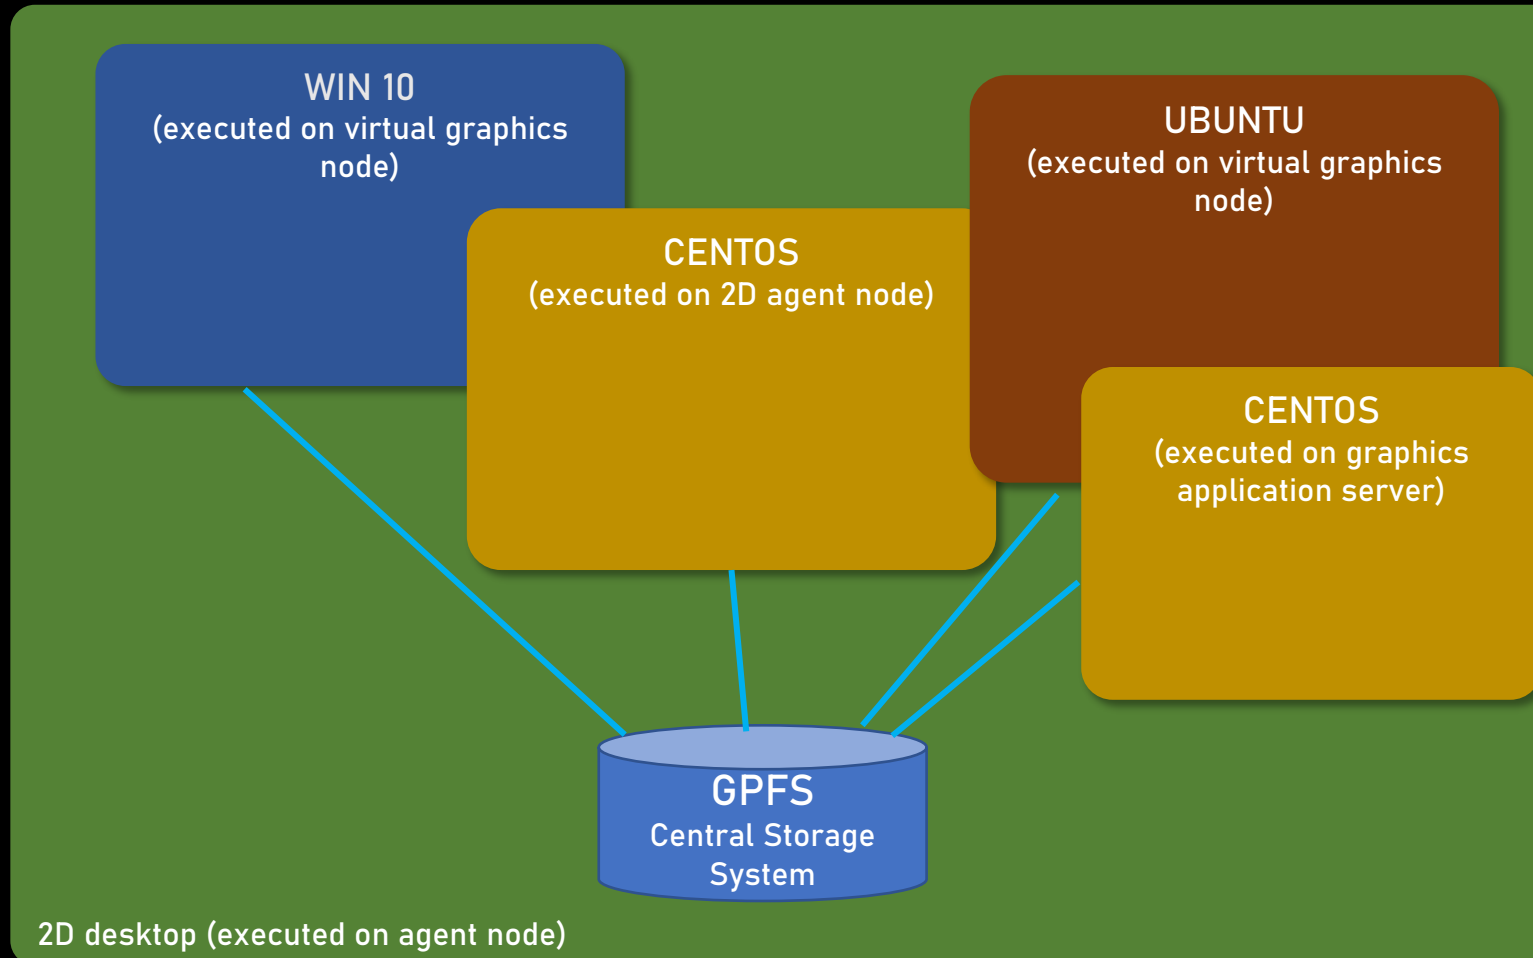


LUNARC Datacenter – High level architecture



LUNARC Datacenter – High level architecture

Final integrated solution. Application windows ingested and assembled into TL desktop



Evolution of interactive HPC and visualisation

Evolution of interactive HPC and visualisation

1. Terminal session (classic ssh, no DT)

```
2. anfo@aurora1:~ (ssh) 🚨
module unload [modulefile] - removes modulefile

As software is installed on the system more modules will become
available.

Other information about using Lunarc systems can be found at
http://www.lunarc.lu.se/Support

NB! The disk space for /home is very limited,
use /lunarc/nobackup/users/<your user name>
for job submission and to store large amounts of data.

Usage and quotas can be checked with the command snicquota.


Basic sample script for running MPI jobs are available in the
directory:

/sw/pkg/submissionsScripts/

For system announcement check http://www.lunarc.lu.se

#####
Compute intensive processes running on Auroras's frontend
are a severe violation of the Lunarc rules.
Such processes will be terminated without warning.
Users violating the Lunarc rules might get their account
suspended.
#####
-bash: warning: setlocale: LC_CTYPE: cannot change locale (UTF-8): No such file or directory
Quota currently unavailable.
Disk quotas for user anfo (uid 1502):
Filesystem      blocks  quota  limit  grace  files  quota  limit  grace
/lunarc/nobackup: 131.1G 500G  550G   -    37184 500000 550000  -
/home:          654M 10240M 20480M 13196 200k  400k

Presto specifics enabled
[anfo@aurora1 ~]$ Hello Lundmarksalen! I'm an old boring terminal window (but still going strong ;) |
```



Evolution of interactive HPC and visualisation

1. Terminal session (classic ssh, no DT)




aurora.lunarc.lu.se

Aurora frontend



Remote terminal user



```

middle@urion [~/bistette] - ssh@urion - ssh
# If software is installed on the system more modules will become
available.

Other information about using Lunarc systems can be found at
http://www.lunarc.lu.se/faq/

!!! The disk space for /home is very limited.
See /home/username/users/your_username
for job submission and to store large amounts of data.

Usage and quotas can be checked with the command urionquota.


Basic template script for running MPI jobs are available in the
directory:
~/templates/submissionScripts/

For system announcements check: http://www.lunarc.lu.se

*****
Complete interactive processing running on Aurora's Frontend
on a secure virtualization of the Lunarc nodes.
Each procedure will be terminated at final network
error (including the user's files might get their account
disabled).


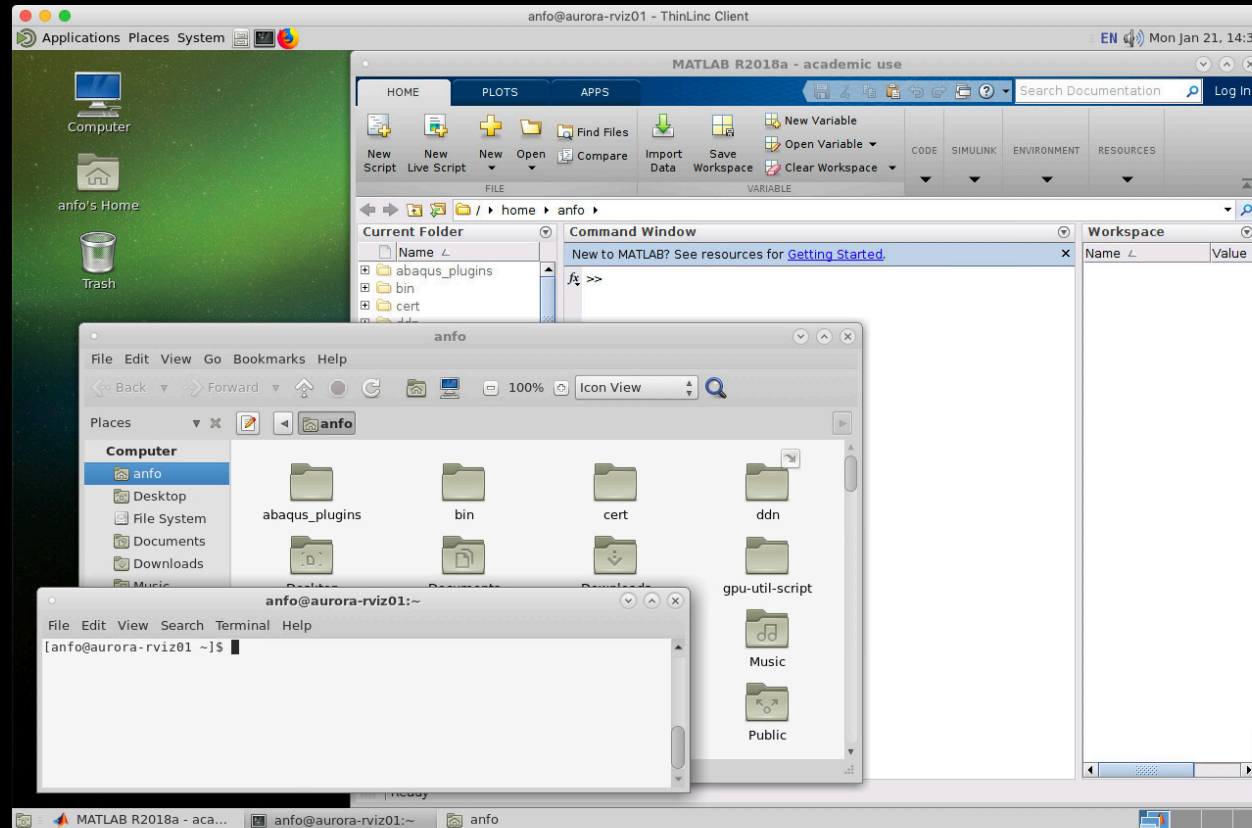
*****
Please settings: setenv LD_LIBRARY_PATH=/usr/lib64:LD_LIBRARY_PATH
Quote currently unavail!
Click inside for user work (111.120.0)
Filesystem      blocks  quota  limit  grace  files  quota  limit  grace
/home/urion     13151  5000  5000  37284  100000  50000
/home/urionproc 6546  1020M 2048M  11506  2000  4000
Please specifics enabled
[middle@urion ~]$ ls -la
total 16
drwxr-xr-x 3 middle middle 4096 Apr  8 11:11 .
-rw-r--r-- 1 middle middle  100 Apr  8 11:11 .bash_history
-rw-r--r-- 1 middle middle  220 Apr  8 11:11 .bashrc
-rw-r--r-- 1 middle middle  832 Apr  8 11:11 .profile

```



Evolution of interactive HPC and visualisation

2. Distributed 2D Desktop (TL 2D DT with custom menus)



Evolution of interactive HPC and visualisation

2. Distributed 2D Desktop (TL 2D DT with custom menus)



aurora.lunarc.lu.se

2D Desktop Agent 1


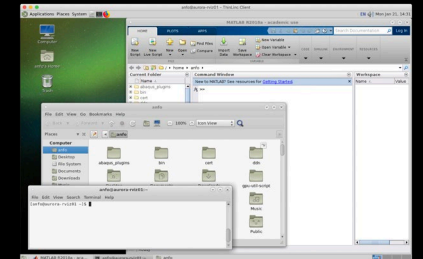

2D Desktop Agent 2

2D Desktop Agent n

Scalable load balancing
2D backend


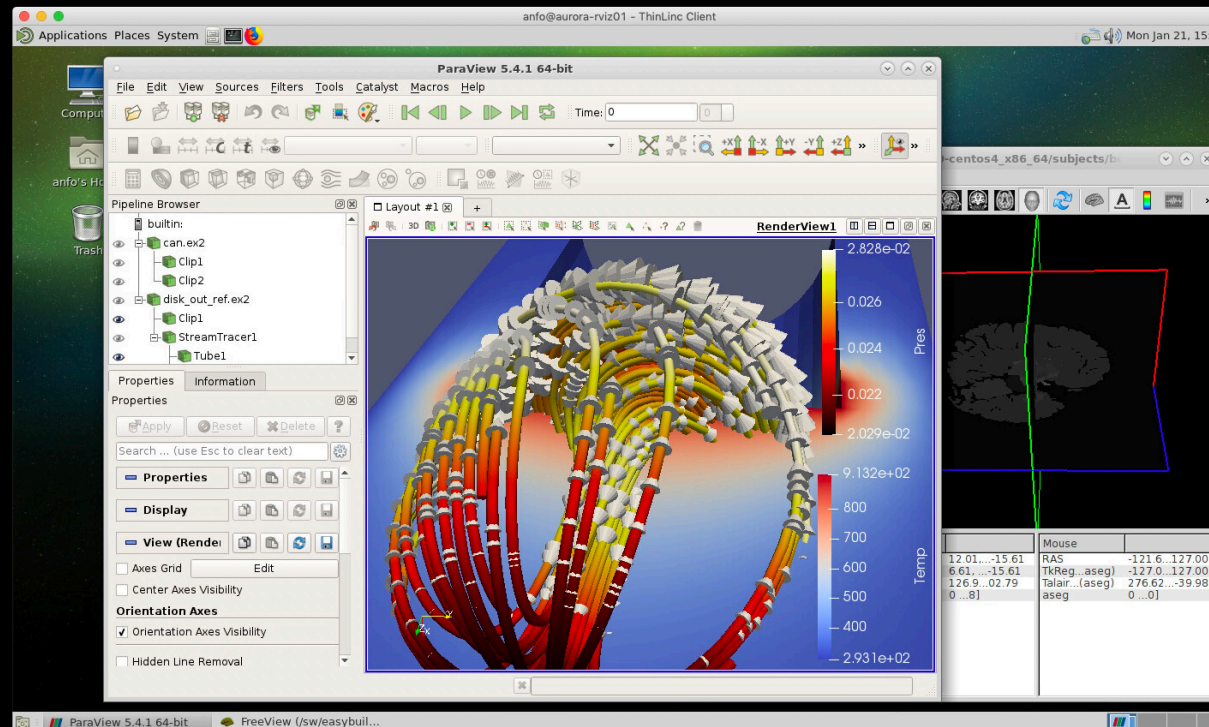
Desktop session

Remote LUNARC
Desktop user



Evolution of interactive HPC and visualisation

3. Accelerated 3D Desktop (2D DT + gfx app servers)




Evolution of interactive HPC and visualisation

3. Accelerated 3D Desktop (2D DT + gfx app servers)

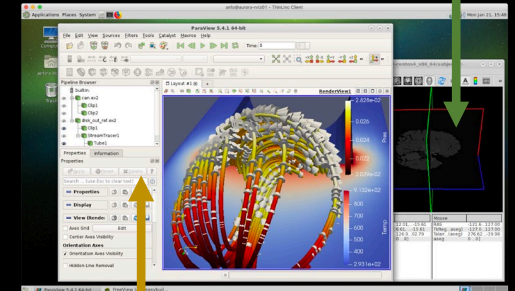


Remote LUNARC Desktop user


Desktop session



Executed on 2D DT Agent



Executed on 3D Appl Server. Accelerated pixels assembled into 2D DT




Evolution of interactive HPC and visualisation

3. Accelerated 3D Desktop (2D DT + gfx app servers)

How to get accelerated pixels from application servers to the remote user across the network?

NV QUADRO (WORKSTATION CARD)


- On board fans (not intended for data center use)
- Physical video output



This is not an option
(we can't pull an DVI
extension cable
across the campus)

NV V100 (SERVER CARD)


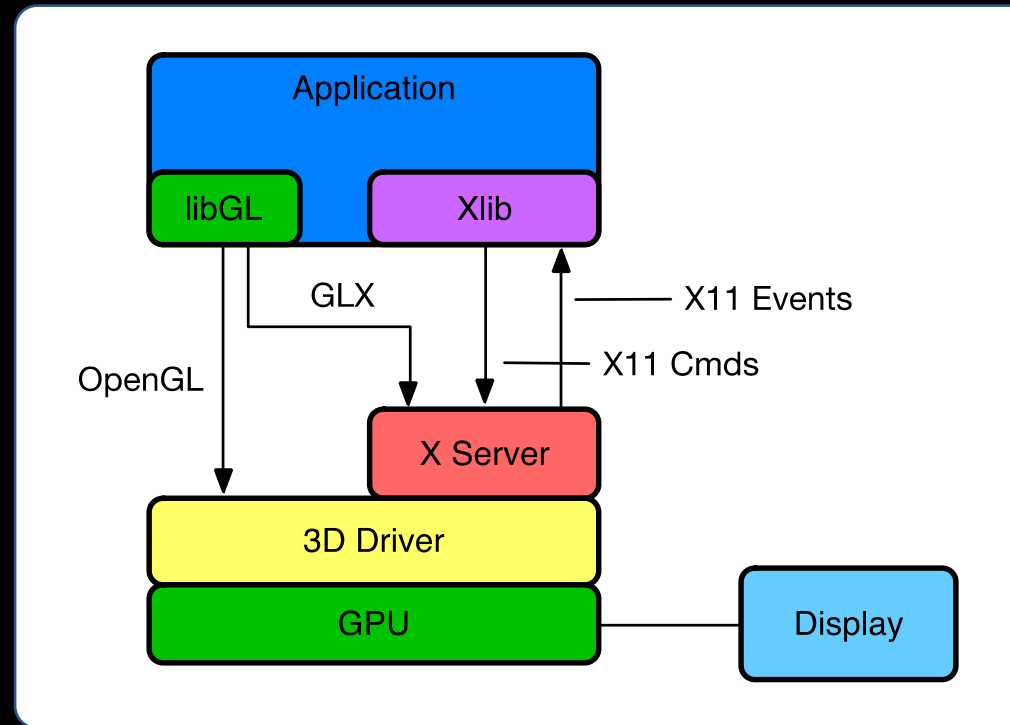
- Passive cooled
- No physical video output



Evolution of interactive HPC and visualisation

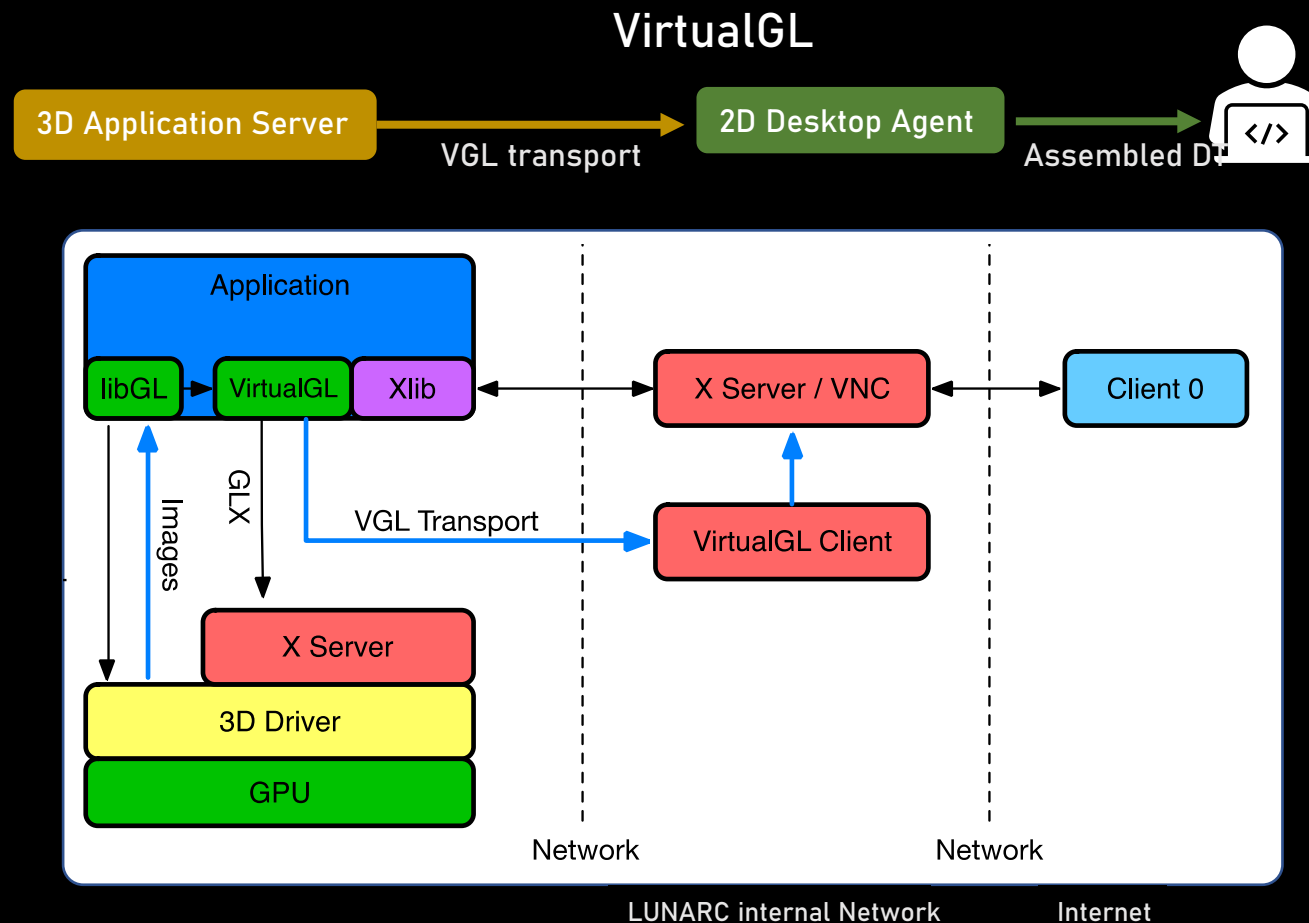
3. Accelerated 3D Desktop (2D DT + gfx app servers)

3D OpenGL on a workstation (with local display)




Evolution of interactive HPC and visualisation

3. Accelerated 3D Desktop (2D DT + gfx app servers)

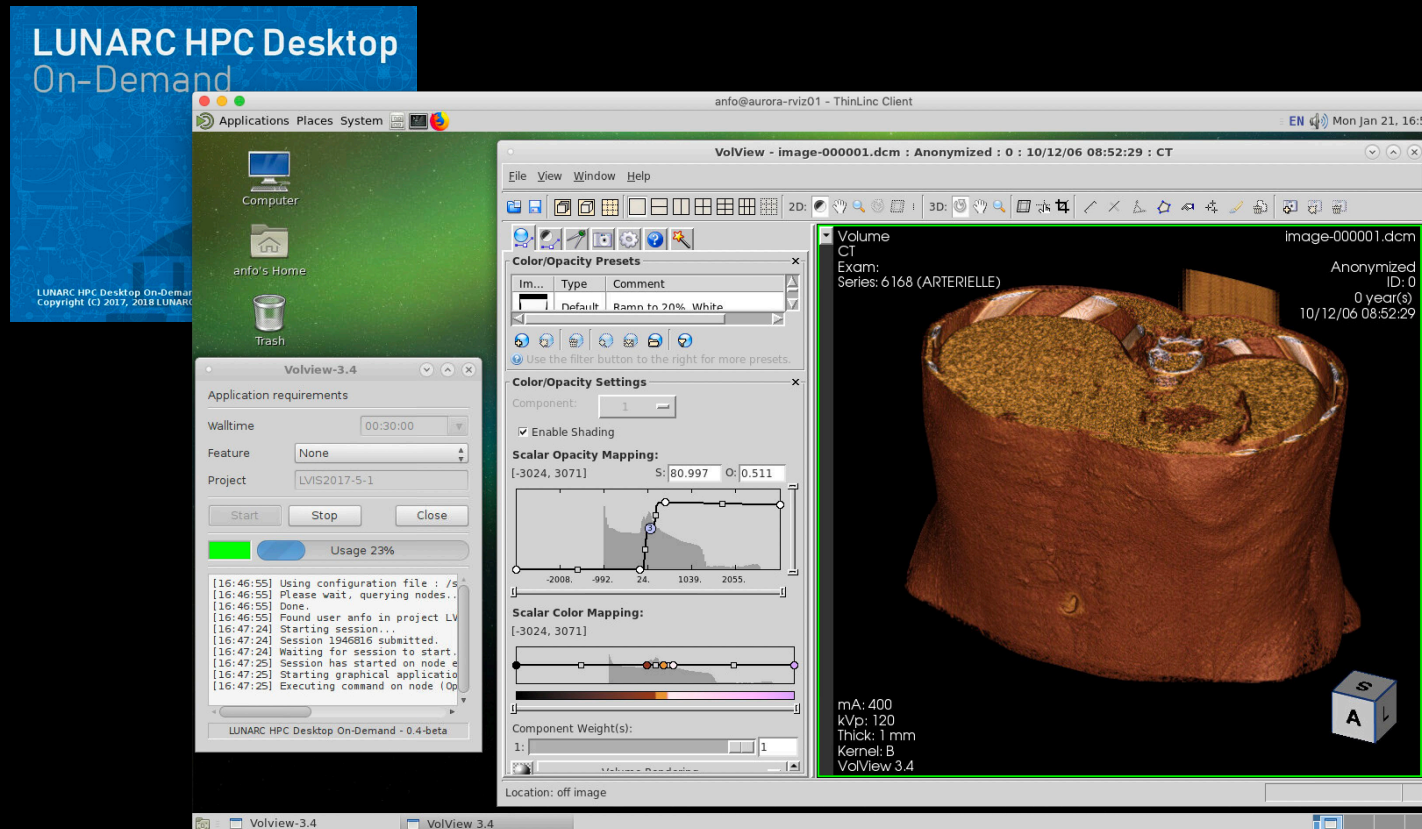


- OpenGL commands are redirected to the application server GPU and rendered locally into a pbuffer instead of being sent to remote (or local) Xserver
- Only rendered frames sent across network to host
- Linux apps only




Evolution of interactive HPC and visualisation

4. Accelerated 3D Desktop (LUNARC HPC Desktop On-Demand gen1)


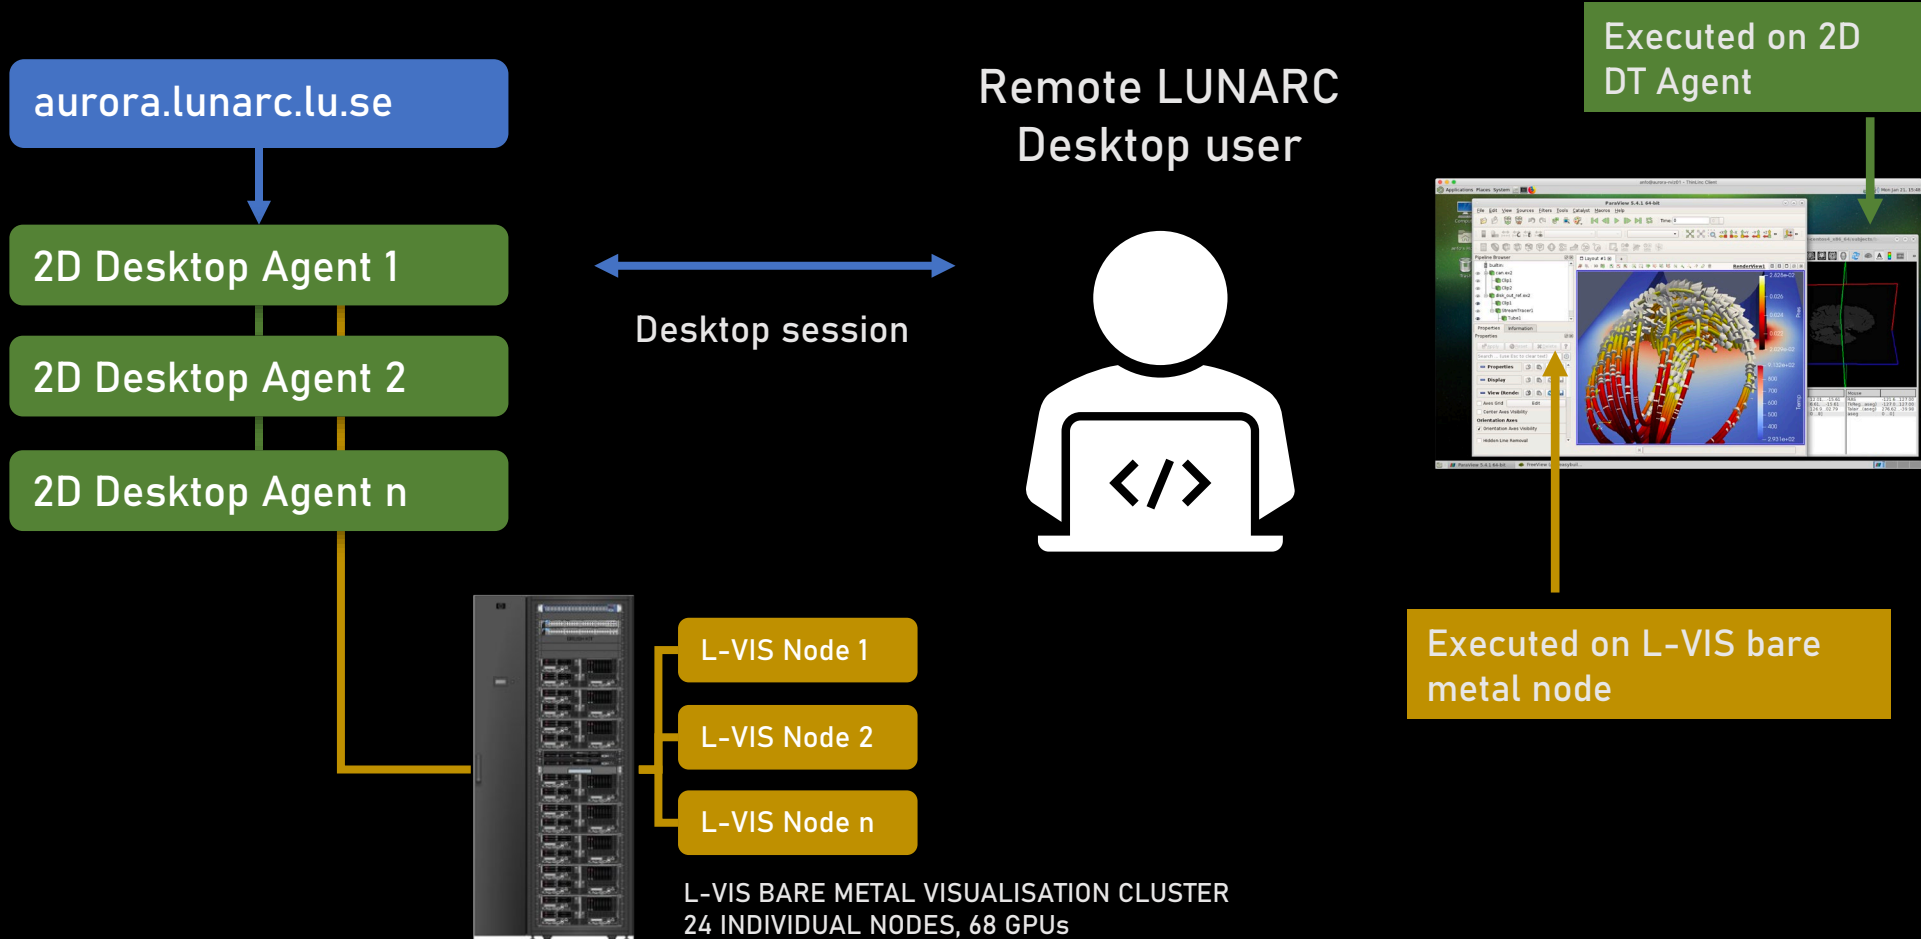


- User interface for specifying resource allocation through SLURM
- User "owns" 100% graphics server resources during walltime i.e. no resource sharing (RAM, CPU, GPU)
- Bare metal servers with hardware accelerated graphics



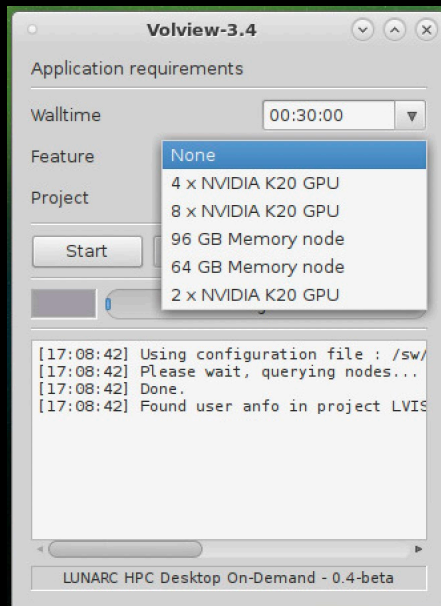
Evolution of interactive HPC and visualisation

4. Accelerated 3D Desktop (LUNARC HPC Desktop On-Demand gen1)


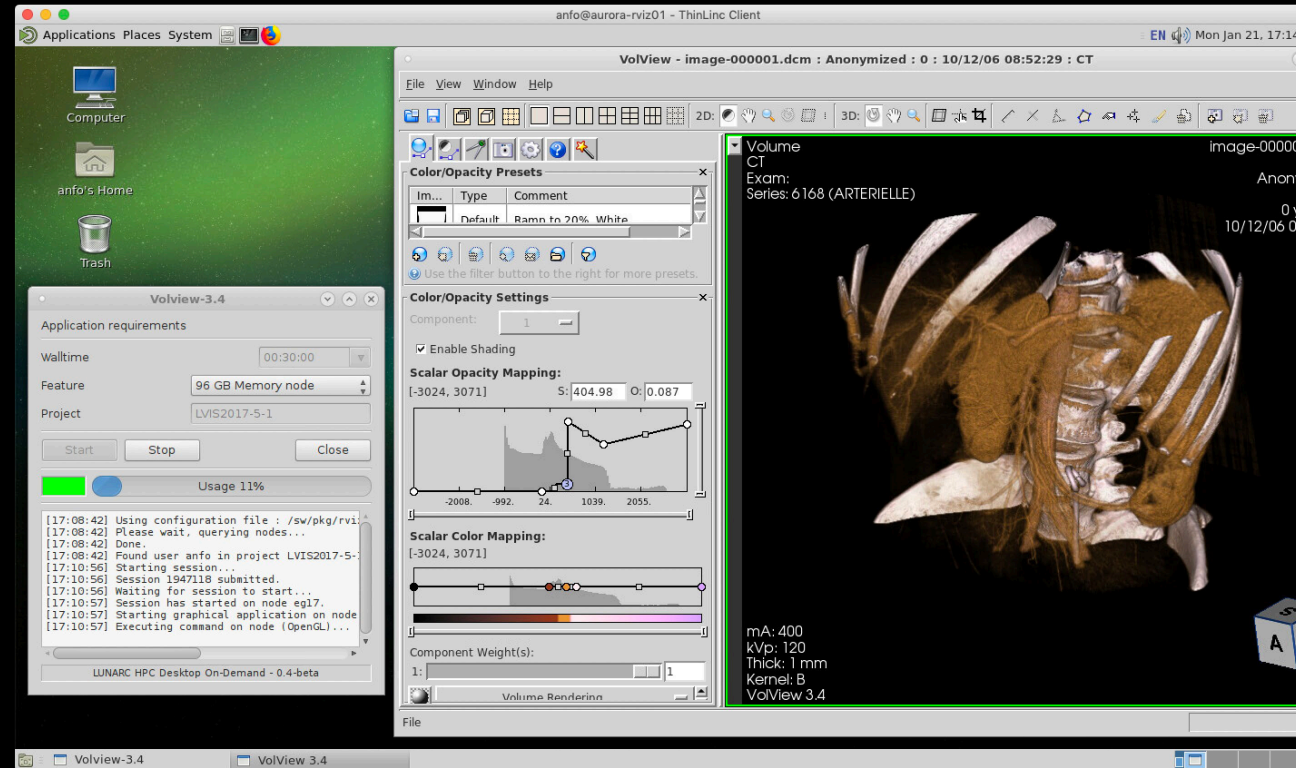


Evolution of interactive HPC and visualisation

4. Accelerated 3D Desktop (LUNARC HPC Desktop On-Demand gen1)

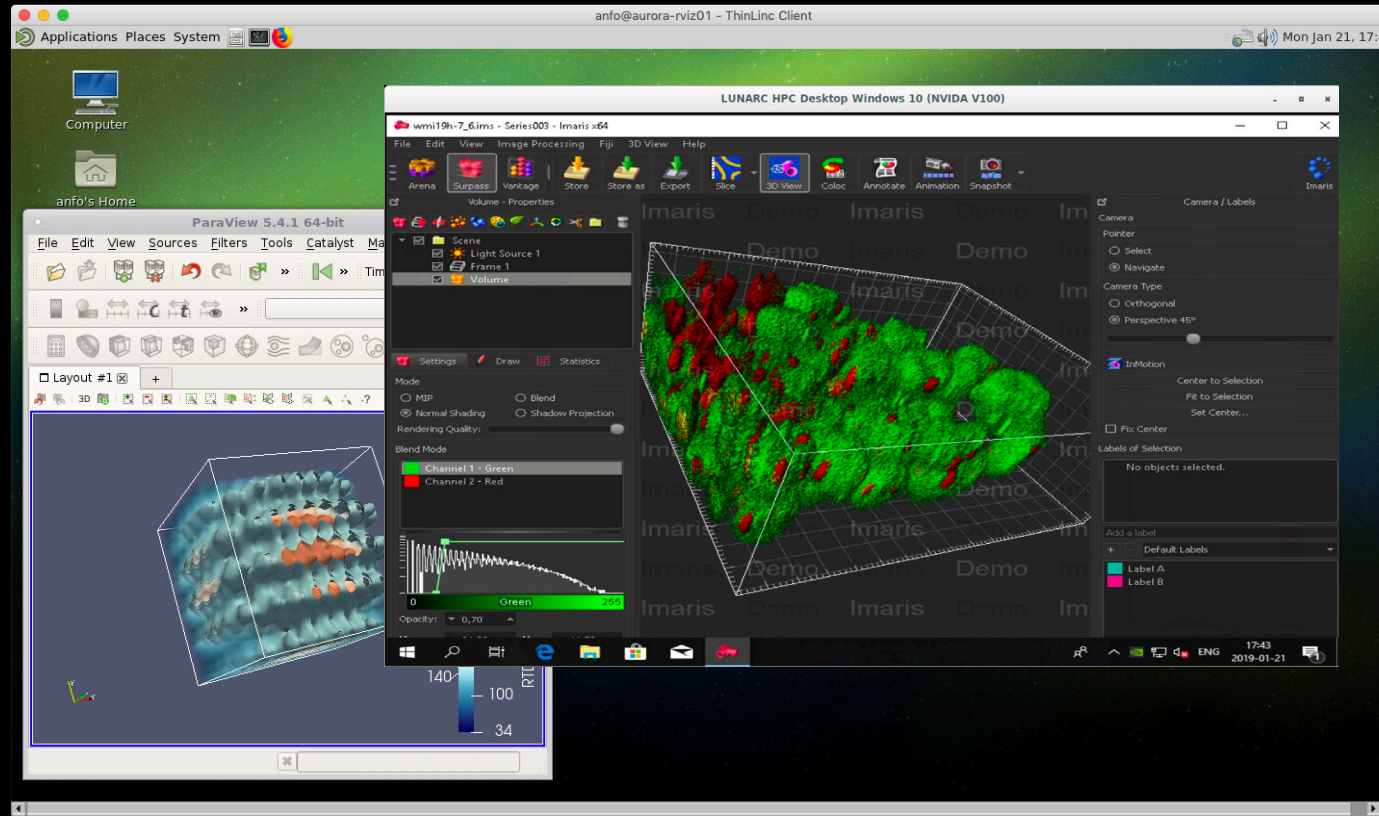


L-VIS node with selectable hardware features allocated for a certain time slot.




Evolution of interactive HPC and visualisation

5. Accelerated 3D Desktop (LUNARC HPC Desktop On-Demand gen2)


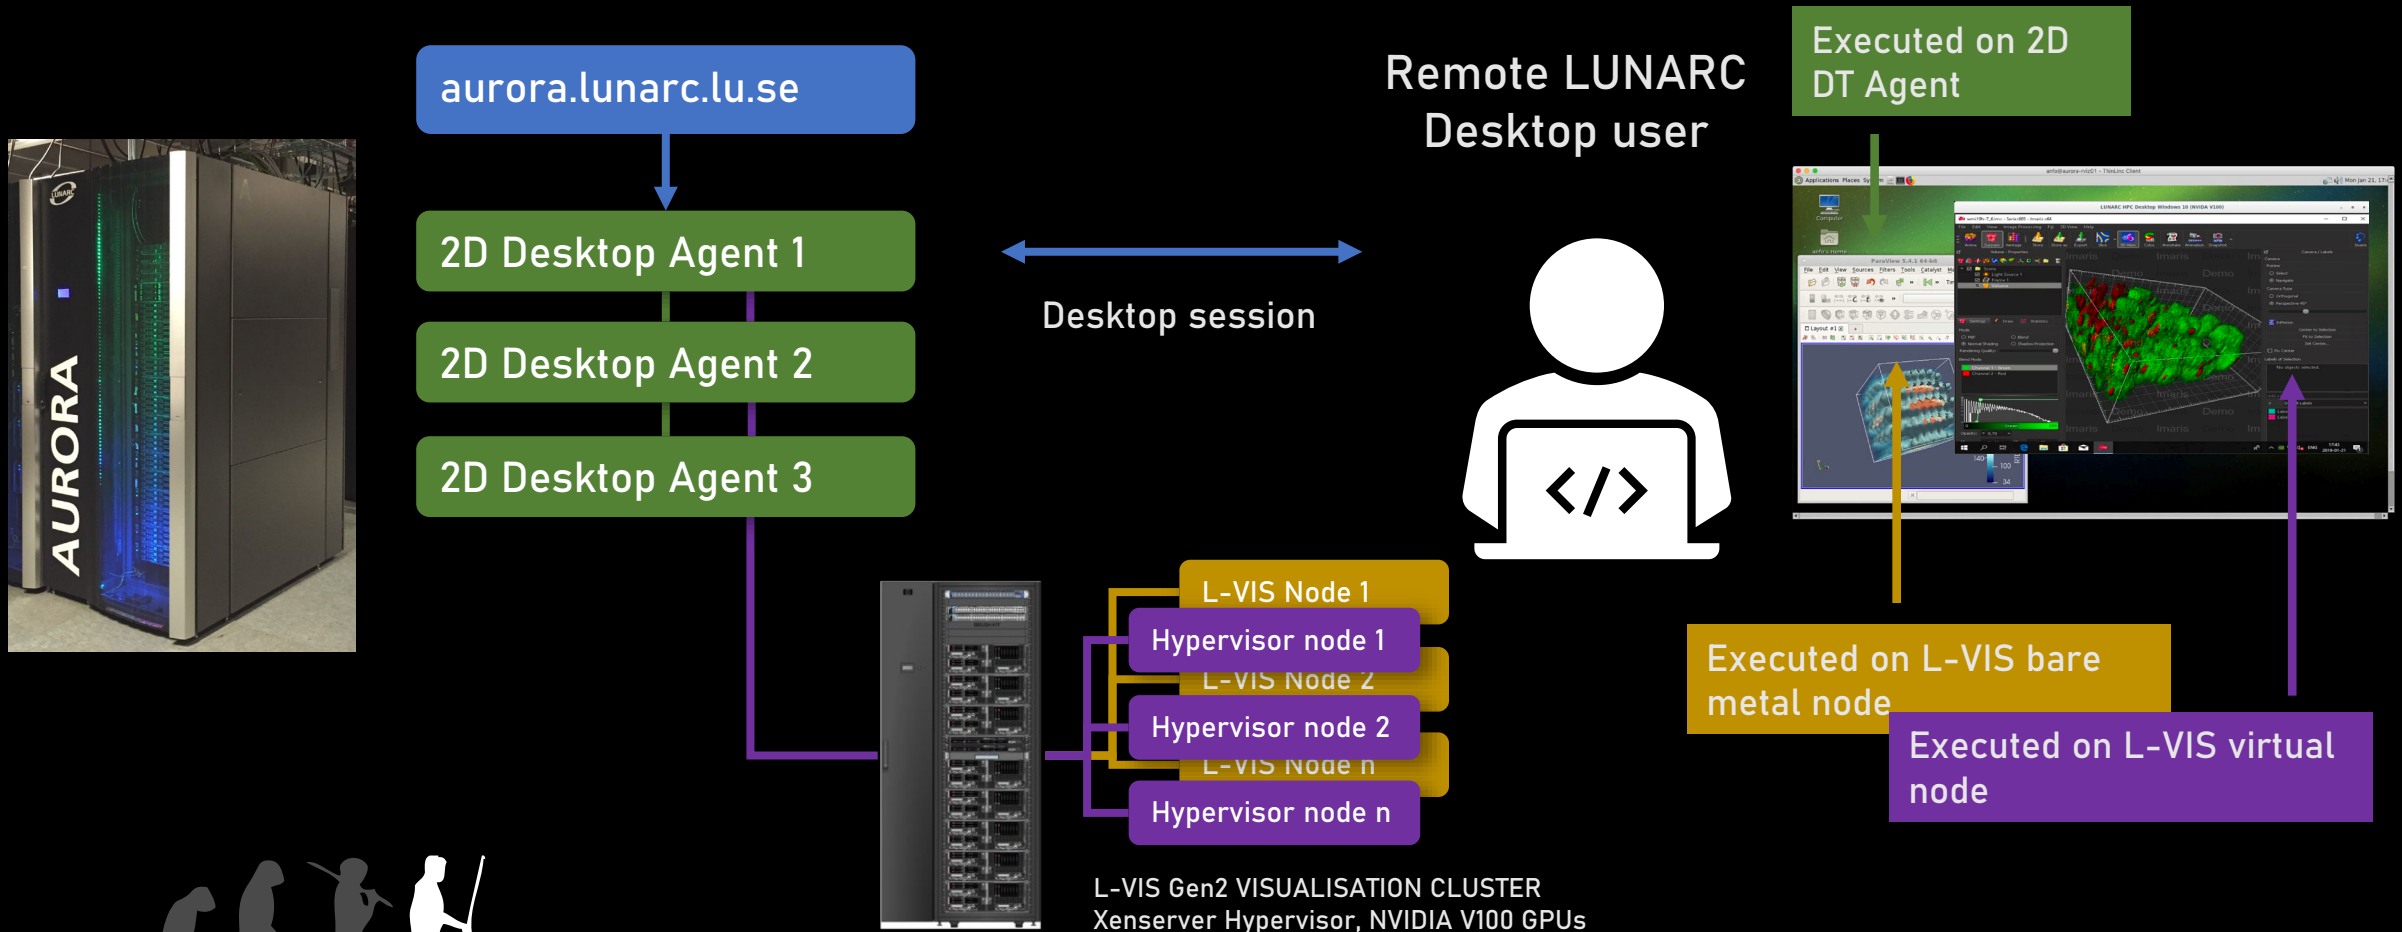


- Operative System Independent (MS Windows 10, Ubuntu, etc)
- Virtual machine resource management using Slurm scheduler
- 2D, 3D app server & VM hardware accelerated graphics assembled into one desktop



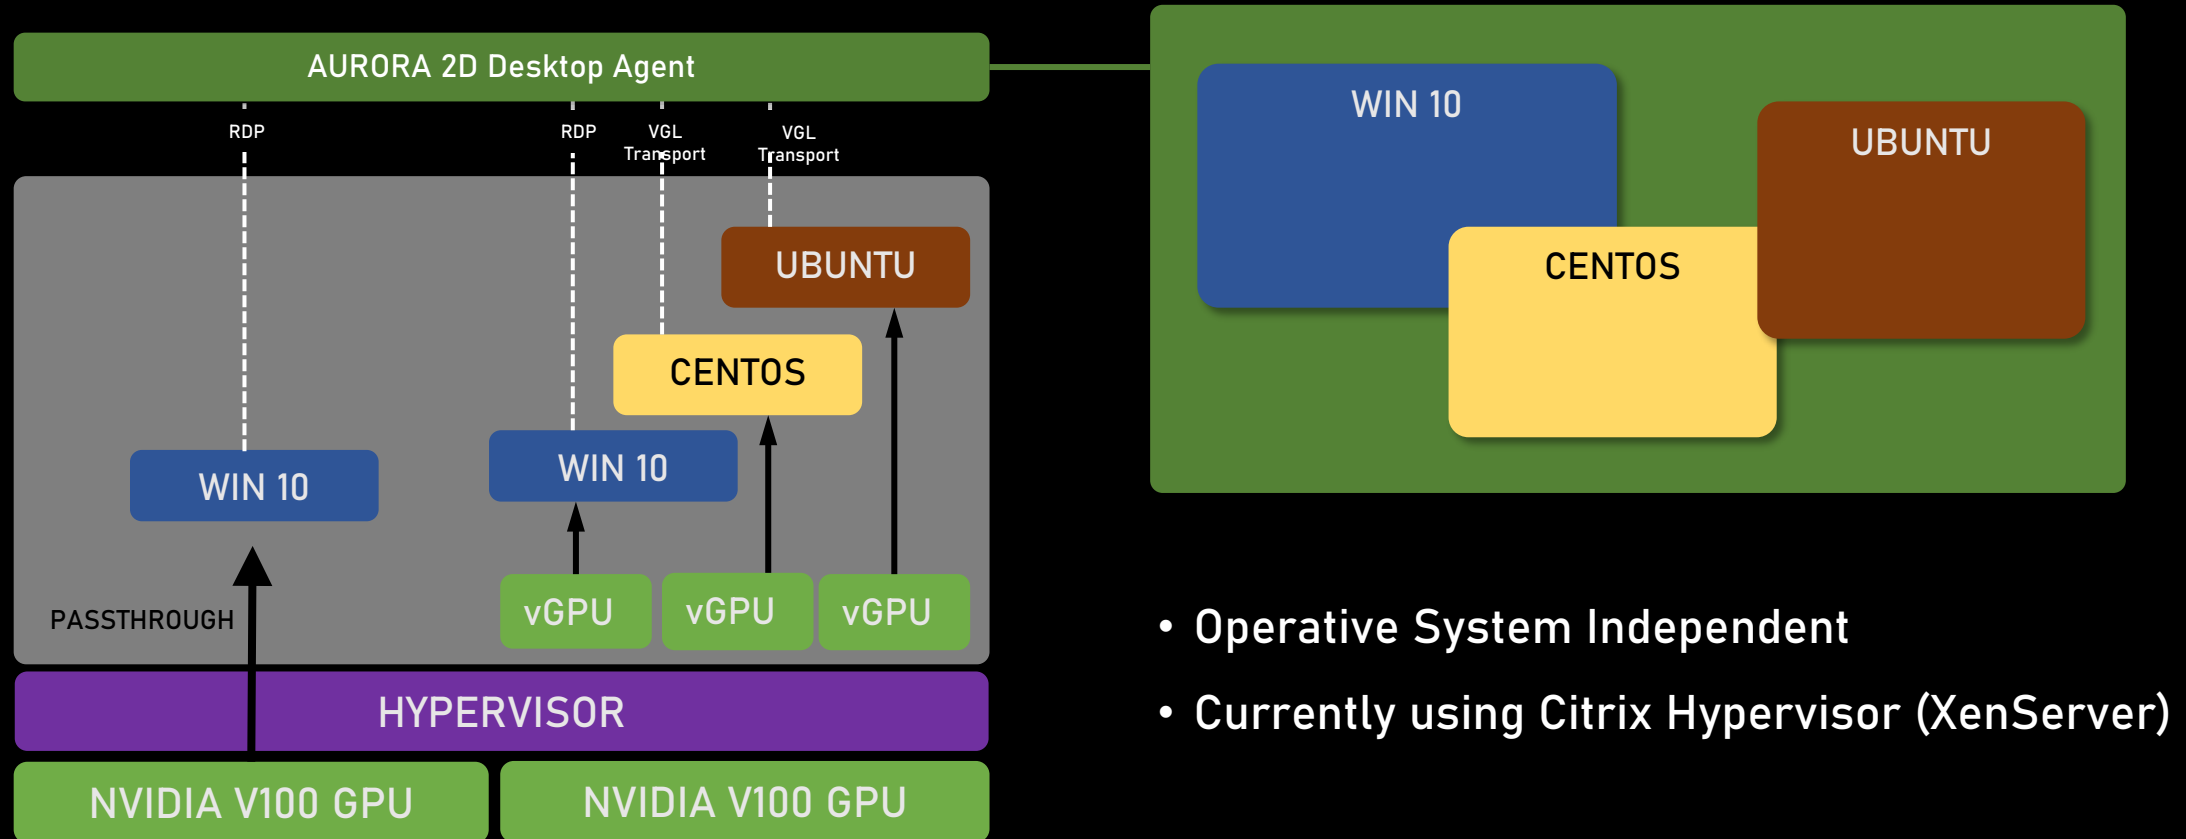
Evolution of interactive HPC and visualisation

5. Accelerated 3D Desktop (LUNARC HPC Desktop On-Demand gen2)



Evolution of interactive HPC and visualisation

Hypervisor node architecture



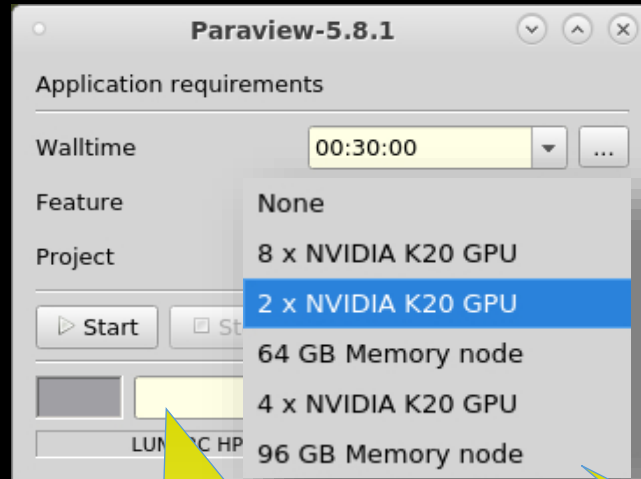
LUNARC HPC Desktop On-Demand

Architecture

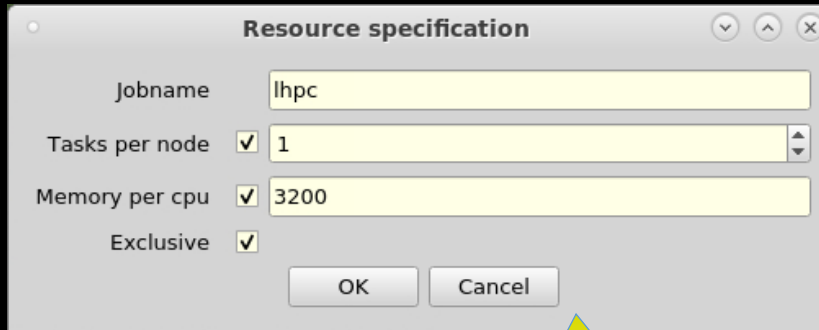
LUNARC HPC Desktop On-Demand

- Graphical client for launching, monitoring and user resource selection
- All interactive applications launched through SLURM
- Supports launching of
 - Interactive graphical applications
 - Interactive graphical applications with hardware acceleration
 - Jupyter Notebooks with automatic connection with automatic start of browser against notebook
 - Windows based VMs through SLURM

GfxLauncher – Application launcher



Usage bar displaying how much of the session time that has been used.

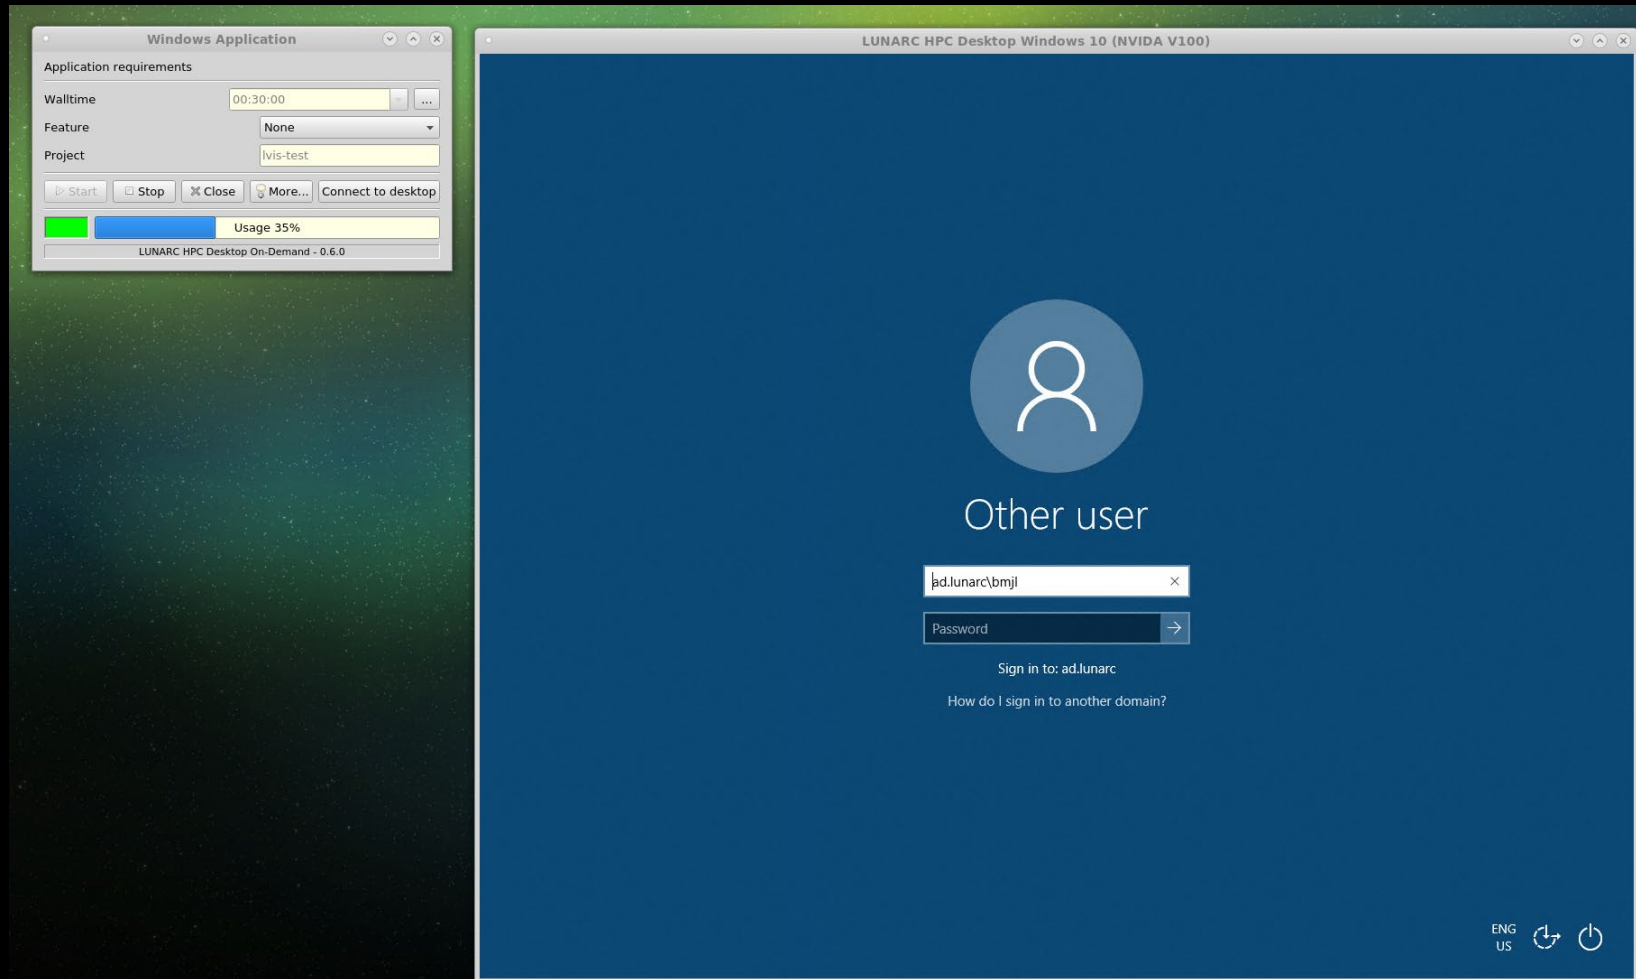


Detailed resource specification dialog.

Lists available resources in SLURM in clear text

- Launches graphical sessions through SLURM
- Configures and sets up sessions using hardware accelerated graphics
- Users gets dedicated server / part of server
- SLURM controls resource usage and session time

GfxLauncher – Window sessions



Case study

3D Scanning and reconstruction in Archeology


New users with large needs

Background


- Archeology in Lund was an early adopter of 3D photogrammetry and 3D reconstruction
- In the beginning only a limited set of photographs were used.
- Increased resolution of cameras produced larger images
- Drone photography could produce aerial images in an automated way.
 - Create large 3D models of large areas
 - Generate a lot of images
- Slowly the demands for compute and storage grew
- Mainly uses interactive software and an interactive workflow for 3D reconstruction.

Bild på monitor – Skriptat arbetsflöde











A typical field study



- Computer
- bmjl's Home
- Old Firefox L
- ABAUQ5 (SLURM)
- ParaView (SLURM)
- VMD (SLURM)
- Trash
- SNIC Storage Explorer
- Spyder - Anaconda3
- Matlab 2018a
- Matlab 2018a - Hardware OpenGL acceleration

Photoscan-1.4.2

Application requirements

Walltime: 06:00:00

Feature: 8 x NVIDIA K20 GPU

Project: LVIS2017-5-1

Start Stop Close

Usage 13%

Logs:

- [15:49:07] Please wait, querying nodes...
- [15:49:07] Done.
- [15:49:07] Found user bmjl in project LVIS
- [15:49:17] Starting session...
- [15:49:17] Session 2366584 submitted.
- [15:49:17] Waiting for session to start...
- [15:49:18] Session has started on node ep2
- [15:49:18] Starting graphical application
- [15:49:18] Executing command on node (Open

LUNARC HPC Desktop On-Demand - 0.4-beta

Rya_02.psx* — Agisoft PhotoScan Professional

File Edit View Workflow Model Photo Ortho Tools Help

Workspace

- Workspace (1 chunks, 987 cameras)
- Copy of Chunk 1 (987 cameras, 605,5
- Cameras (987/987 aligned)
- Tie Points (605,953 points)
- Dense Cloud (52,644,162 points, Medi
- 3D Model (3,509,609 faces)

Perspective 30°

faces: 3,509,609 vertices: 1,759,350

Property Value

Model

- Faces 3,509,609
- Vertices 1,759,350
- Vertex colors 3 bands, uint8
- Texture 4,096 x 4,096

Reconstruction parameters

- Surface type Arbitrary
- Source data Dense
- Interpolation Enabled
- Quality Medium
- Depth filtering Aggressive
- Face count 3,509,610
- Processing time 33 minutes ...

Texturing parameters

- Mapping mode Generic
- Blending mode Mosaic
- Texture size 4,096 x 4,096

Console

```

2019-05-13 15:49:30 Finished processing in 0.733415 sec (exit code 1)
2019-05-13 15:50:32 Error: Can't open file: No such file or directory (2): /lunarc/nobackup/users/bmjl/church/Rya_02.files/1/0/dense_cloud/dense_cloud.oc3
2019-05-13 15:51:17 Error: Can't open file: No such file or directory (2): /lunarc/nobackup/users/bmjl/church/Rya_02.files/1/0/dense_cloud/dense_cloud.oc3
  >>>
  
```

Workspace Reference Photos Console Jobs


Preferences Help About Remove From Panel Move Lock To Panel



Redigera Egenskaper

| | Storlek | Ändrad |
|------------------|----------|------------------|
| 349.JPG | 5 134 KB | 2019-03-29 16:43 |
| 350.JPG.filepart | 3 006 KB | 2019-05-02 21:46 |

- sese_cpp_course
- sese_exercises
- storage_dir
- sw
- Templates
- test42
- test3412
- testcpp
- testing
- testing3423
- Untitled Folder



Data and compute requirements from a single site

- ~1000 images - 9 GB of data
- Processing time on laptop 7 days (if it fits in memory)
- Verifying data in the field not possible

LUNARC as a field tool

- Processing using LUNARC HPC On-demand
 - Minutes for a course mesh for coverage
 - 4 hours for complete model using GPU:s (older ones)
- Verification of data coverage is now available in the field
- Data processing and analysis can be started in the field
- Upload of data can be a problem depending on location
 - 5G will enable faster uploads
 - For slower connection uploads can be done between sessions
 - Even with slower speeds: "Hours instead of days" for computation
- Bring your favorite data center into the field

DEMO

Thank you!

- For more information
 - jonas.lindemann@lunarc.lu.se
 - anders.follin@lunarc.lu.se
 - www.lunarc.lu.se

