Compute and Storage

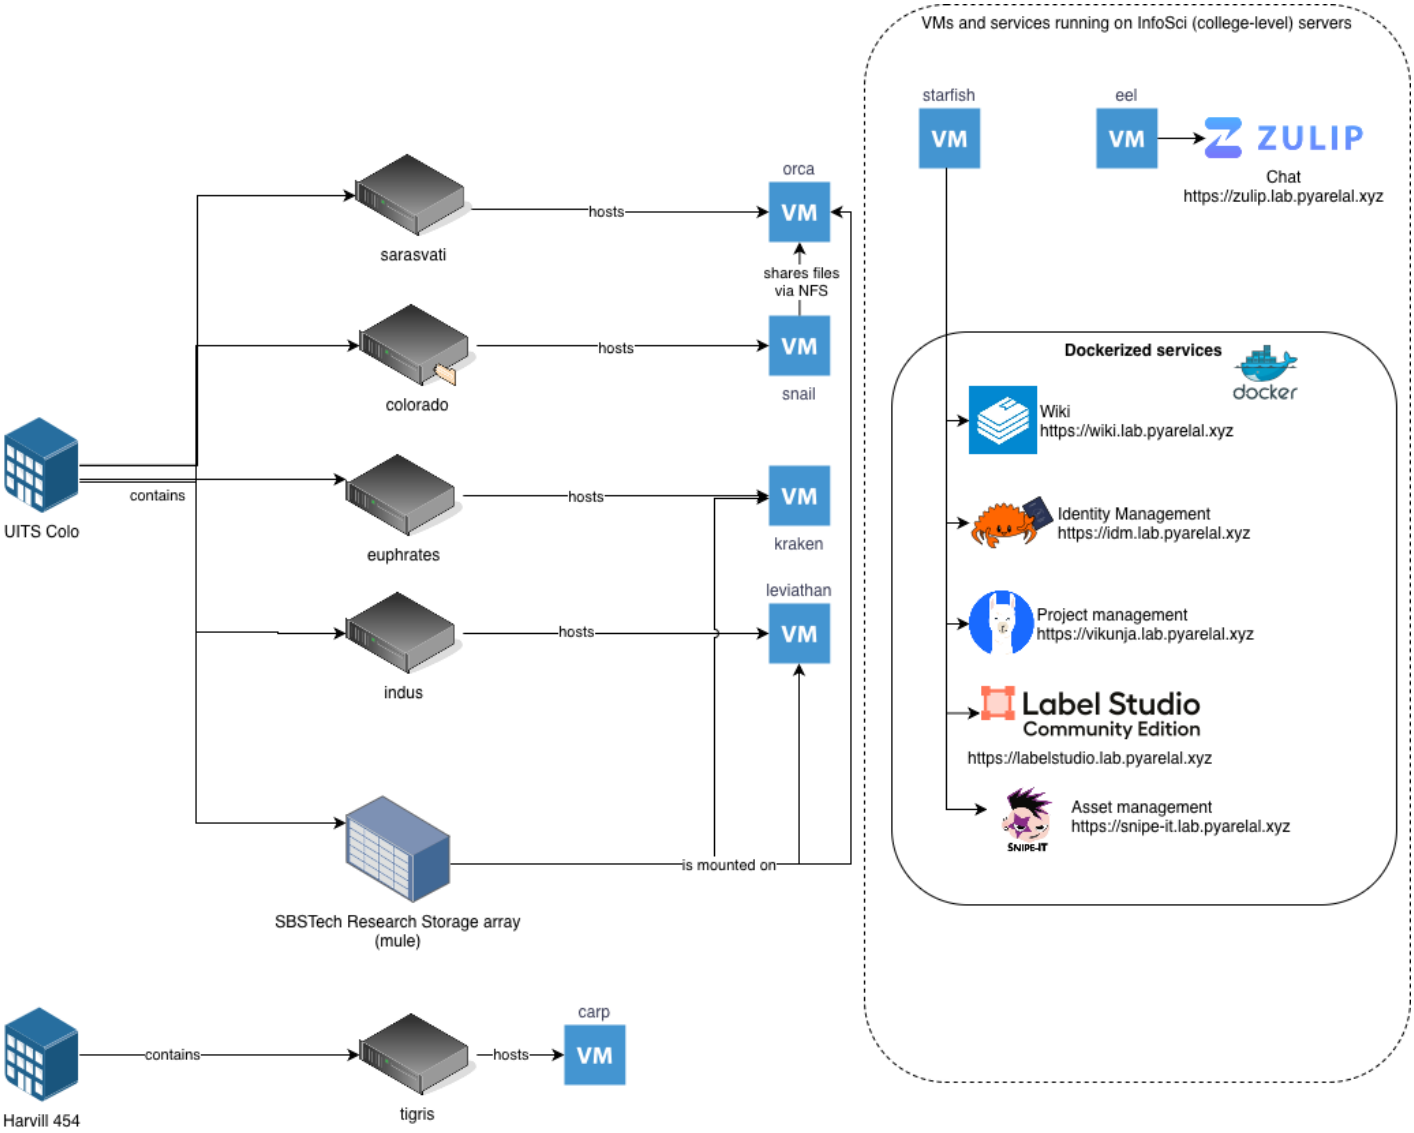

Compute

The ML4AI lab has the following compute VMs:

VM Name	CPU	RAM	GPUs
kraken	AMD EPYC 7662 64-Core Processor (2.0 GHz)	720 GB	2x NVIDIA A100 (40 GB)

VM Name	CPU	RAM	GPUs
leviathan	AMD EPYC 7763 64-Core Processor (2.45 GHz)	720 GB	6x NVIDIA RTX A6000
carp	AMD EPYC-Rome Processor	95 GB	1x NVIDIA GeForce RTX 3090
orca	AMD EPYC 9474f, 48-core, 3.60 GHz, 256MB cache	1.5 TB (tentative)	2x NVIDIA H100 NVL (94GB hbm3, PCIE 5.0 x16)

TODO:

- Add information about venti

Storage

The lab has a 20 TB NFS share mounted at /media/mule on the kraken, leviathan, and orca VMs.

There is a 90 TB NFS share mounted at /media/snail-ssd on the orca VM.

TODO:

- Add information about timelord

TODO:

- Add information about other legacy SISTA systems that are still operational.

Backup

- VMs running on InfoSci servers (e.g., eel, starfish) are backed up every 6 hours
- As of 2025-03-06, VMs running on lab servers (e.g., kraken, leviathan) are not backed up, but the plan is to include them in the backup system in the future.

Revision #26
Created 30 August 2024 03:59:22 by Adarsh Pyarelal
Updated 6 February 2026 18:48:41 by Adarsh Pyarelal