Evolution of the SWAN service

Diogo Castro
On behalf of the SWAN team

https://cern.ch/swan

Oct 11th, 2019
SWAN Users' Workshop

This workshop is for the SWAN team to learn more about your use cases (…) hearing about how you use SWAN will help us understand how the service should evolve to suit your needs.

Removals

End of life for SLC6

> January 2020
  - Aligned with lxplus

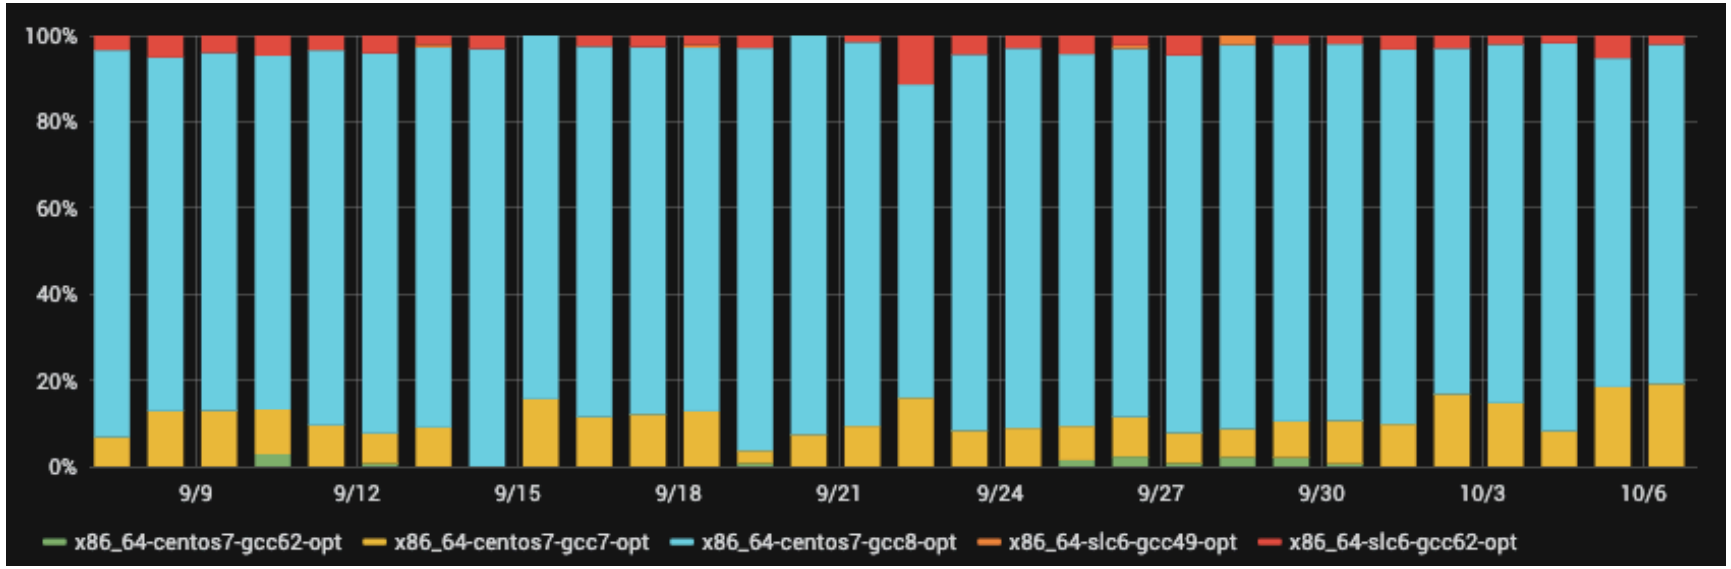

Deprecation of Python 2

> January 2020

> Python 3 as default
  - Old python 2 stacks still available

Deprecation of old LCG stacks

> Ongoing discussion
  - Depending on support from users

> Stacks with low/no usage in the last 30 days:
  - LCG 92 (Python 3), LCG 91 (Python 3), LCG 90 (Python 3), LCG 89 (Python 3), LCG 88 (Python 3), LCG 87, LCG 86, LCG 85

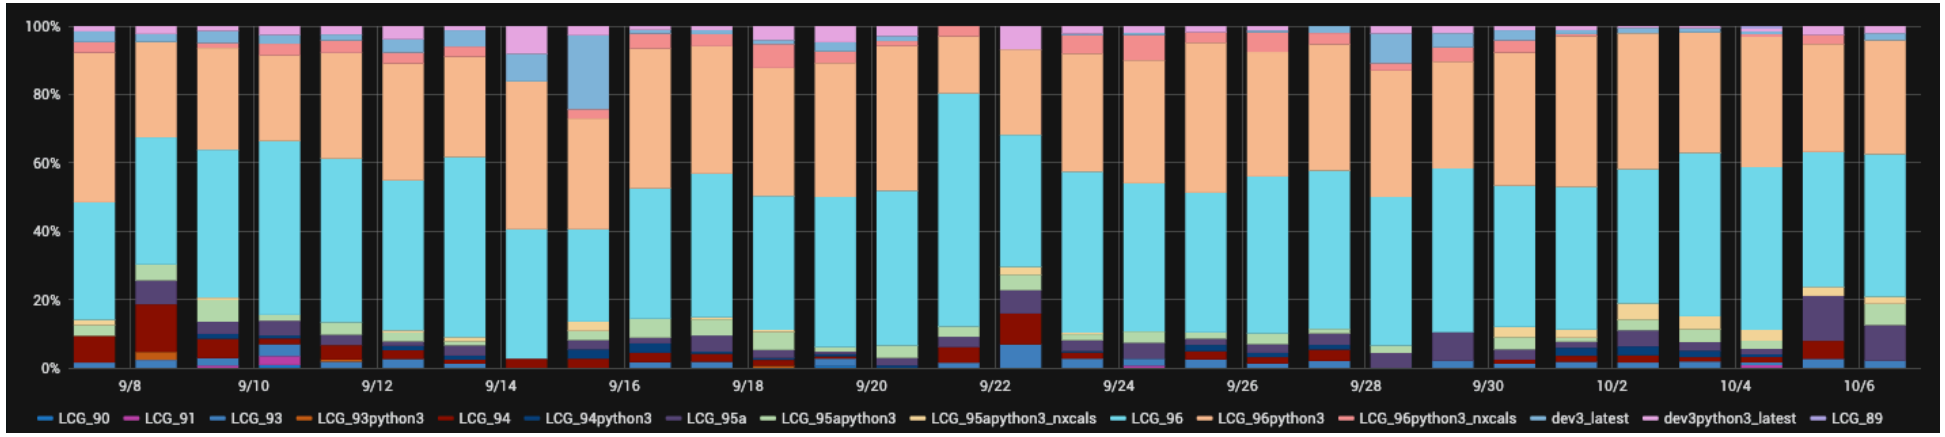

Additions

Jupyterlab

- Next-generation interface for Project Jupyter
  - Concurrent editing
- Missing: porting the current extensions
- Notebooks interface available in parallel during the transition

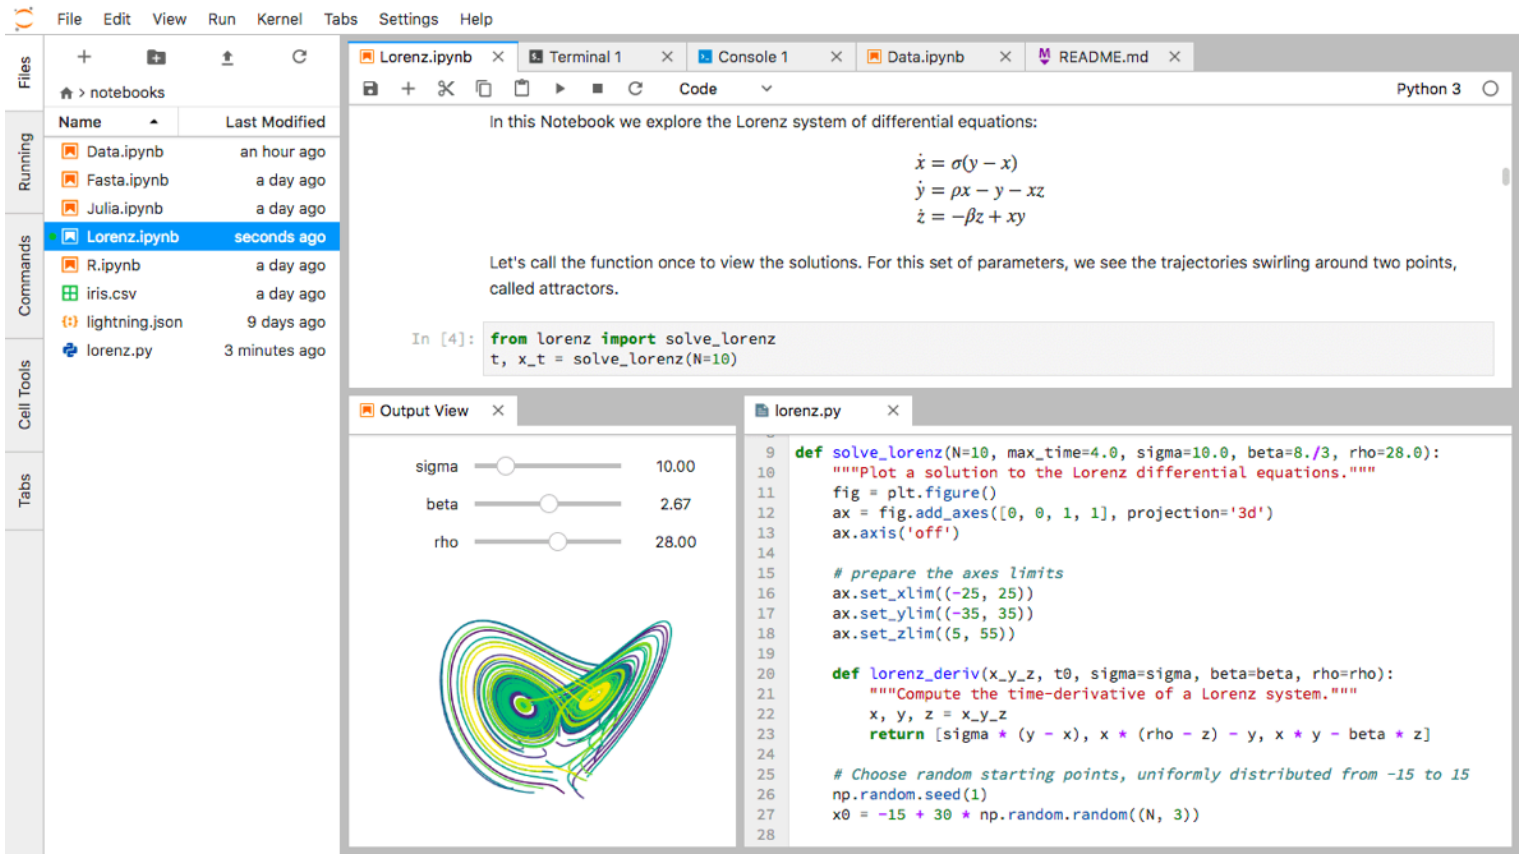

NVidia GPU Support

> Exploitation of container technologies to provide support for NVidia GPUs
  - Already integrated with ScienceBox

> Prototype server for testing purposes
  - NVidia Tesla V100 PCIe 32GB
  - If interested, ask us to join the beta program

> All the packages are provided by CVMFS
  - Including CUDA enabled machine learning software stack
  - TensorBoard for interactive monitoring

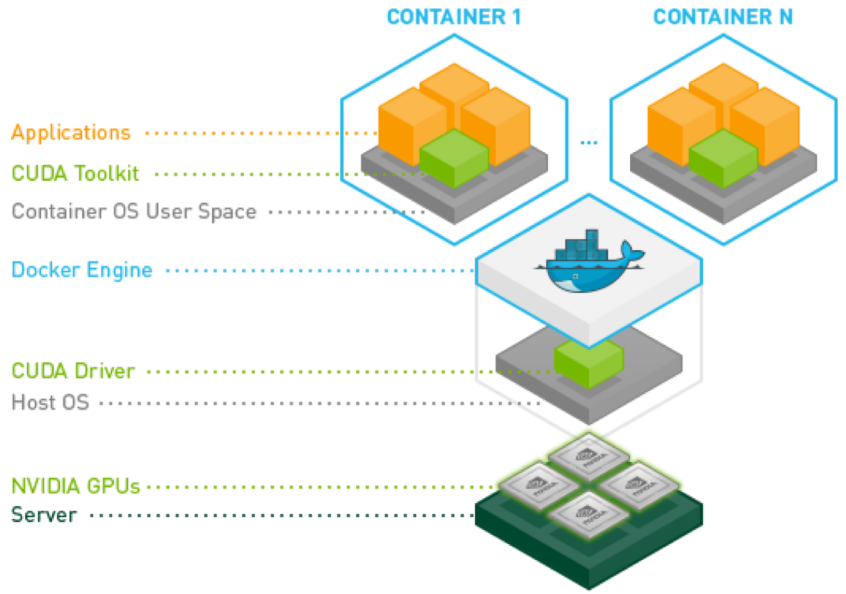

User managed Spark k8s

> Possibility to connect to user managed Kubernetes clusters
  - Offload Spark computations
  - Control and use your own resources
  - Quickly create, use and dispose

> Share access with other users

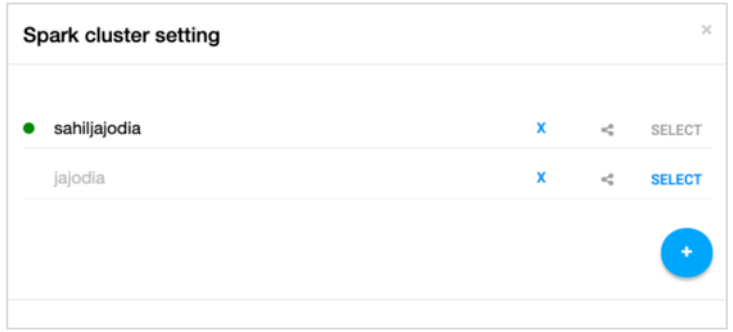

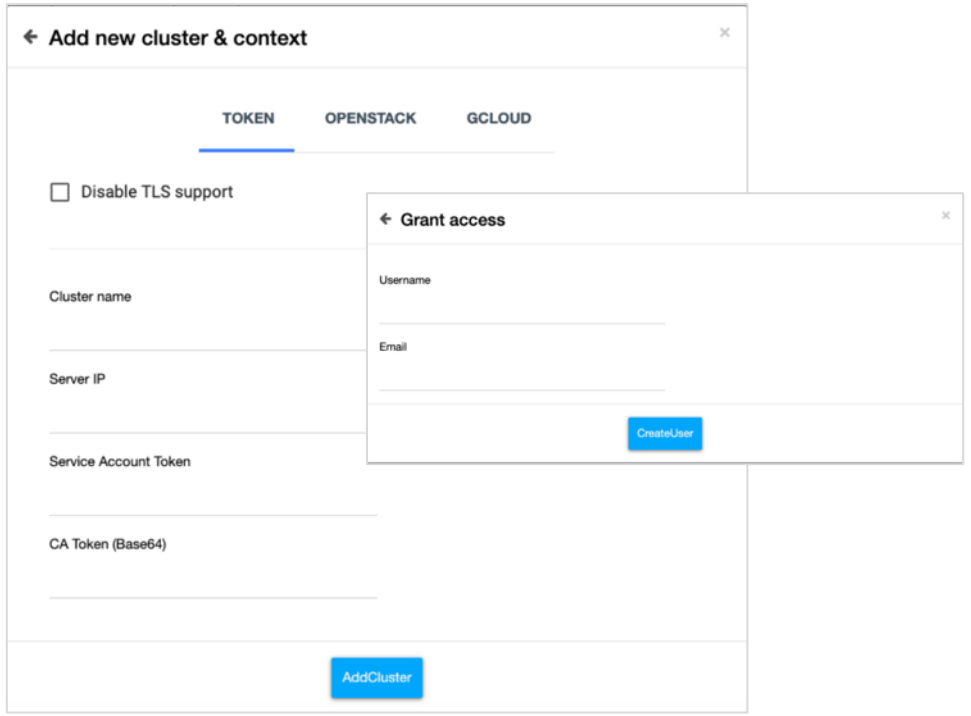

In exploration

Batch submission

- Ongoing effort: submit **batch jobs** from the notebook
  - Using Ganga
  - Monitoring display
  - Jobs tab

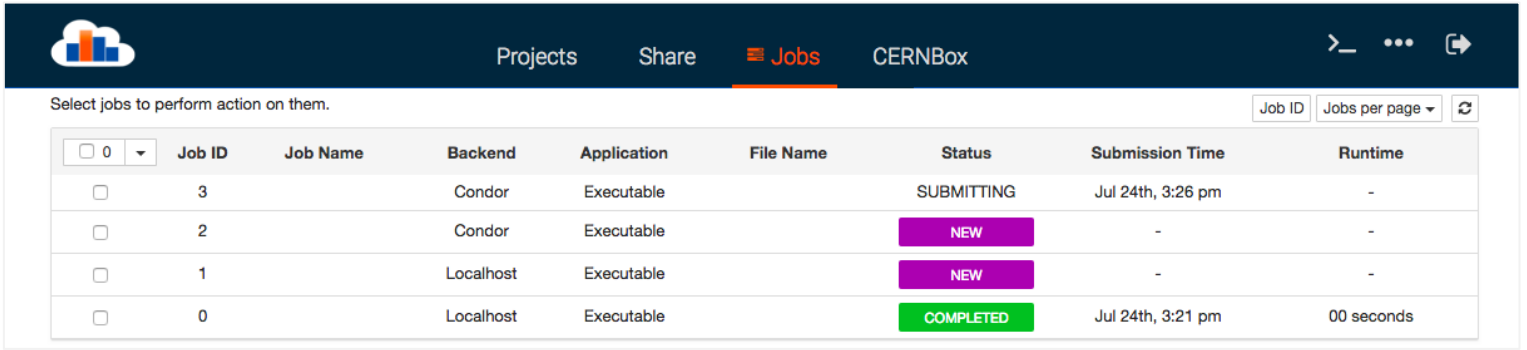


```
In [6]: %%ganga
        j=ganga.Job()
        j.submit()
```

Backend: LOCALHOST Application: EXECUTABLE Splitter: None 0 SUBJOBS

Job ID	Job Name	Status	Subjobs	Submission Time	Runtime
0		COMPLETED	No Subjobs	Jul 24th, 3:21 pm	00 seconds

Configurable software environment

- Adding support for Conda environments
  - Linked to Projects
  - Sharable
- Easy installation of extra packages
  - Clone/import Projects and install the software automatically
- Still a proof of concept
  - Integration with EOS is starting

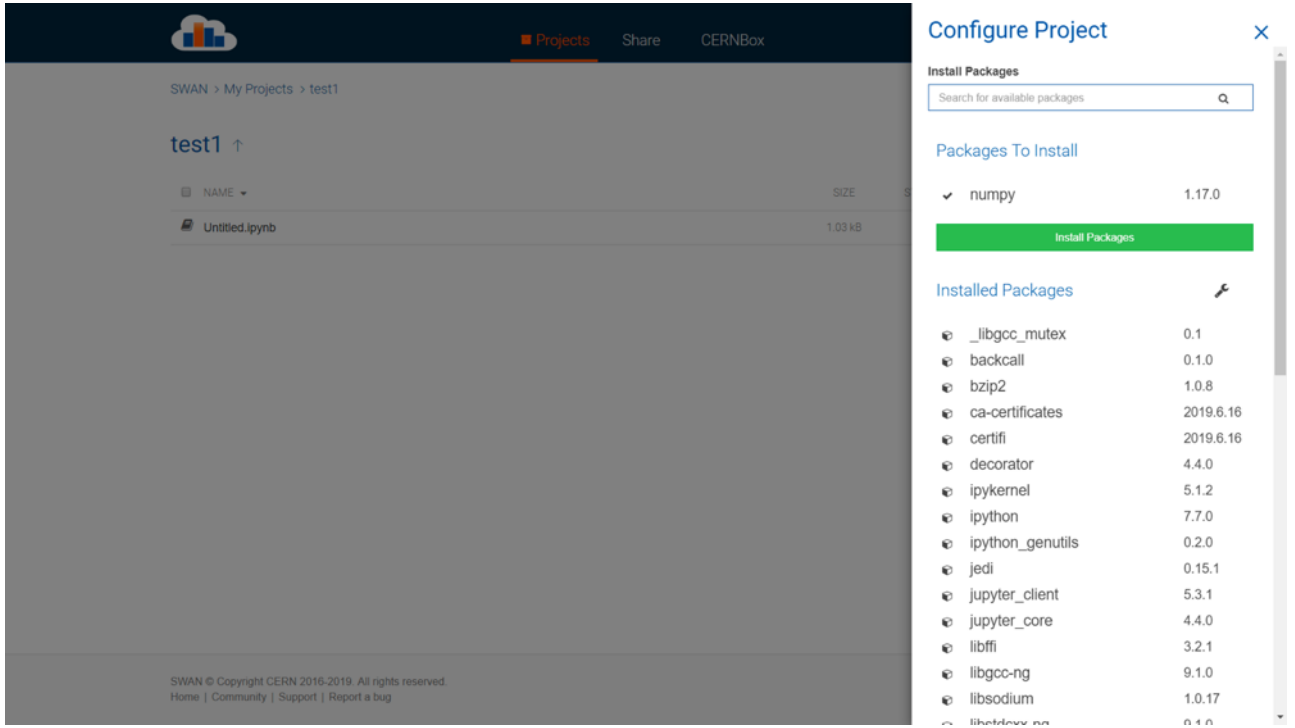

Changes

Move to Kubernetes

> Ongoing effort
- Based on ScienceBox and upstream
- Pilot infrastructure already in testing

> Improve and modernize SWAN infrastructure
- Replicated, highly-available containers
- Add capacity in minutes to support spikes in service utilization
- Ability to roll out updates with no impact on service capacity
- Leverage on Cloud Containers service by IT-CM

Future improvements

Thank you

Diogo Castro
diogo.castro@cern.ch