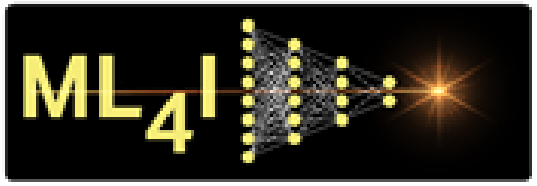

Machine Learning for Industry Forum, August 10 – 12, 2021

Repairing Cardiac Valve Leakage with Patient-specific Data Generation and Machine Learning on an Azure/Google/UberCloud Multi-cloud Environment with SUSE Rancher Kubernetes Management

Customer:

Technology Partners:

Yaghoub Dabiri, Ghassan S. Kassab, and Julius M. Guccione, 3DT Holdings, San Diego, CA
Daniel Gruber and Wolfgang Gentzsch, UberCloud, Los Altos, CA

What is UberCloud ?

- Cloud services provider, on top of AWS, Azure, Google Cloud, etc.
- Providing automated, secure, on-demand, browser-based, self-service **engineering simulation cloud platform** and custom **HPC containers**
- Running on any Linux, in **any cloud**: public, private, hybrid, multi-cloud
- Used by some of the largest companies in the world
- 2015 Azure HPC Partner
- 2018 Largest customer on AWS
- 2020 Google GCP Technology Partner

Goal of this Living Heart Project:

Provide a real-time surgery simulator to heart surgeons

for repairing cardiac valve leakage with a MitraClip device

Machine Learning: 12h simulation => 2 - 3 secs prediction

Focus on practical usability, no supercomputer, just ‘cheap’ HPC cloud

This work was supported by the SBIR Grant Number R43 HL145896.

Heart Failures in the US

- American Heart Association: Heart failure (HF) has increased from **5.7 million (2012) to 6.2 million (2016)** Americans.
- Projected to **increase 46% by 2030**.
- **2030**, estimated total cost: **$70 billion**.
- Need for **innovative treatment** strategies for HF.
- **Computational simulation** provides a virtual platform where the behavior of the heart can be simulated and novel interventions can be assessed and . . .
- . . . To optimize **pharmaceutical** and **device** design and **implantation**.

BUT: FE heart sim too long: 1 heart cycle = 10 – 100 h on workstation

Real-Time Virtual Assessment of Mitral Valve Repair Device Placement

MitraClip (MC) is a recent percutaneous approach whereby the clip is placed in the mitral valve to reduce Mitral Regurgitation (MR).

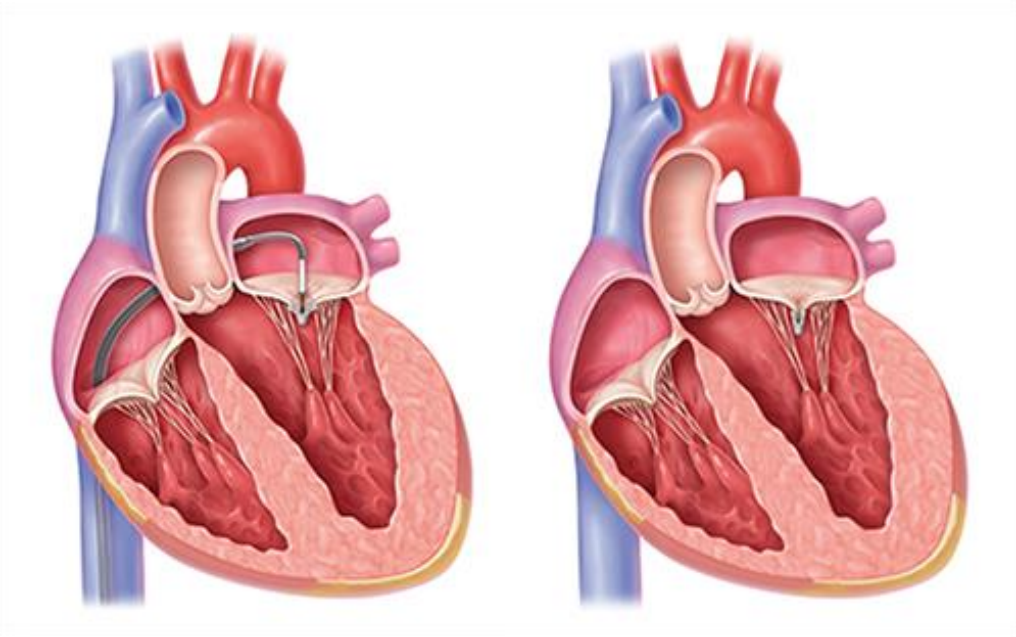

https://www.kherskowitzmd.com/mitraclip-procedure.html

MitraClip device for minimal-invasive catheter-based repair for high-risk mitral regurgitation patients

MitraClip (MC), a small metal clip, allows doctors with catheter-based surgery to repair the mitral valve, brings minimally invasive alternative to open-heart surgery.

Computational Modeling

- Abaqus FE model includes: Left ventricle, Mitral valve and chordae muscles, and blood
- Structural & Smoothed-Particle Hydrodynamics (SPH) modeling
- Models created for different valve geometries and MC scenarios
- UberCloud Engineering Simulation Platform used to run the 2,500 simulations on GCP Google Cloud Platform

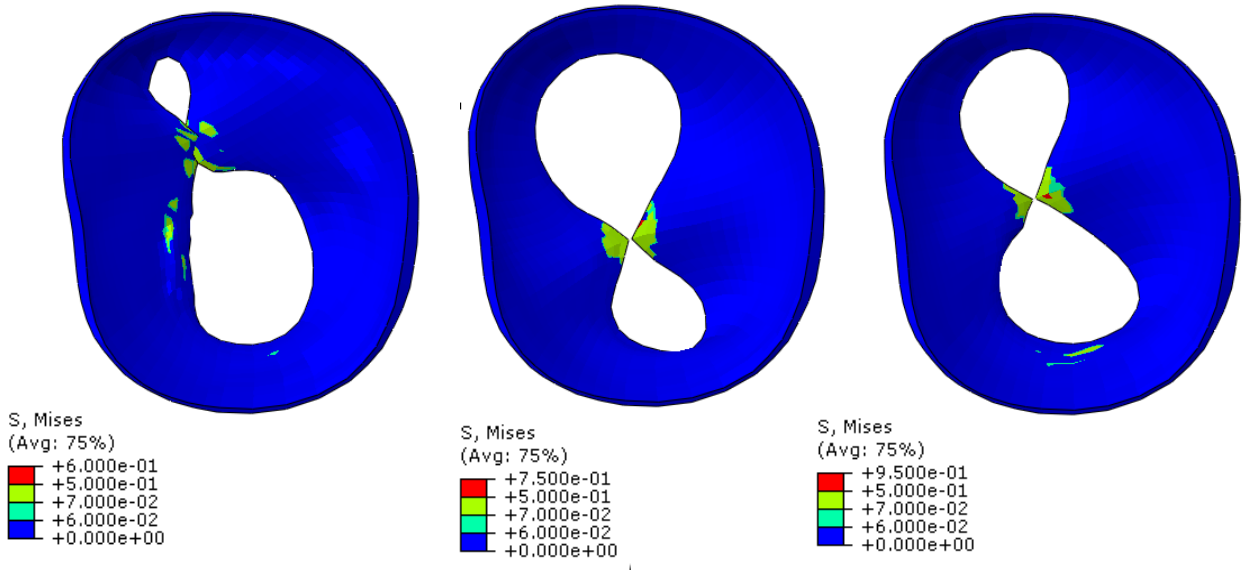

Stresses on the mitral valve for different MC implant placements

Problems with MitraClip (MC)

- Number of Clips
- Location of Clips

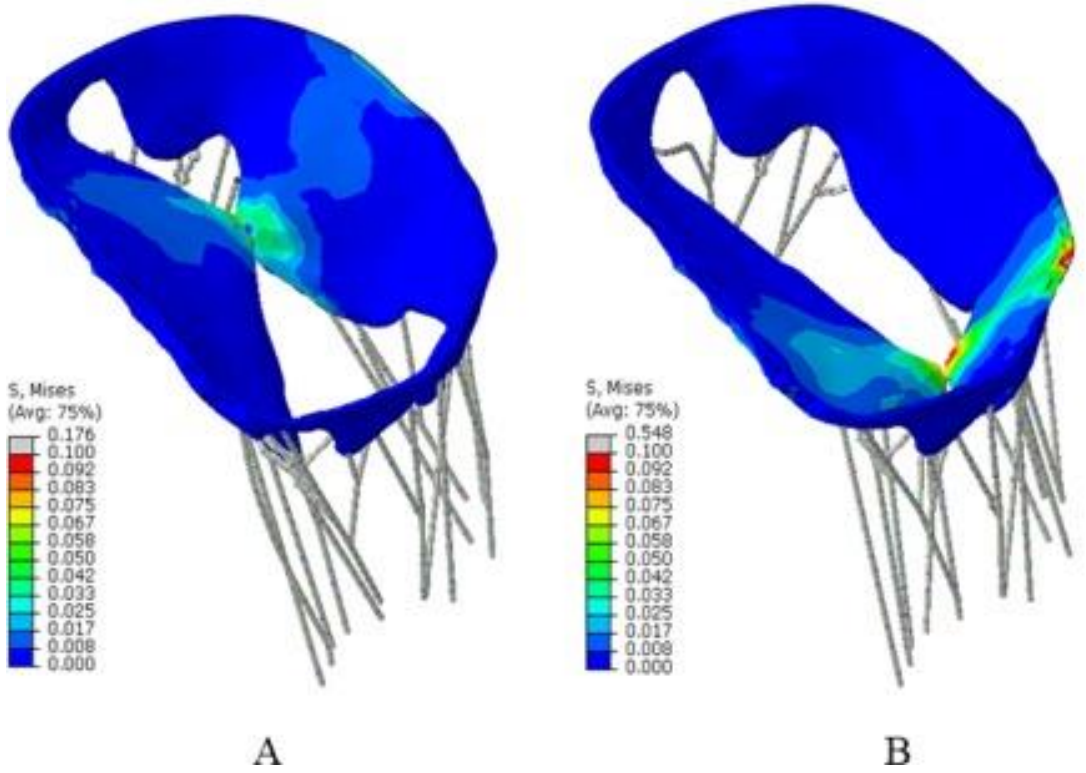

Stresses on the mitral valve for different MC implant placements.

Kamakoti et al. Cardiovasc Sci Rep. 2019; 9: 15823.

Machine Learning Models

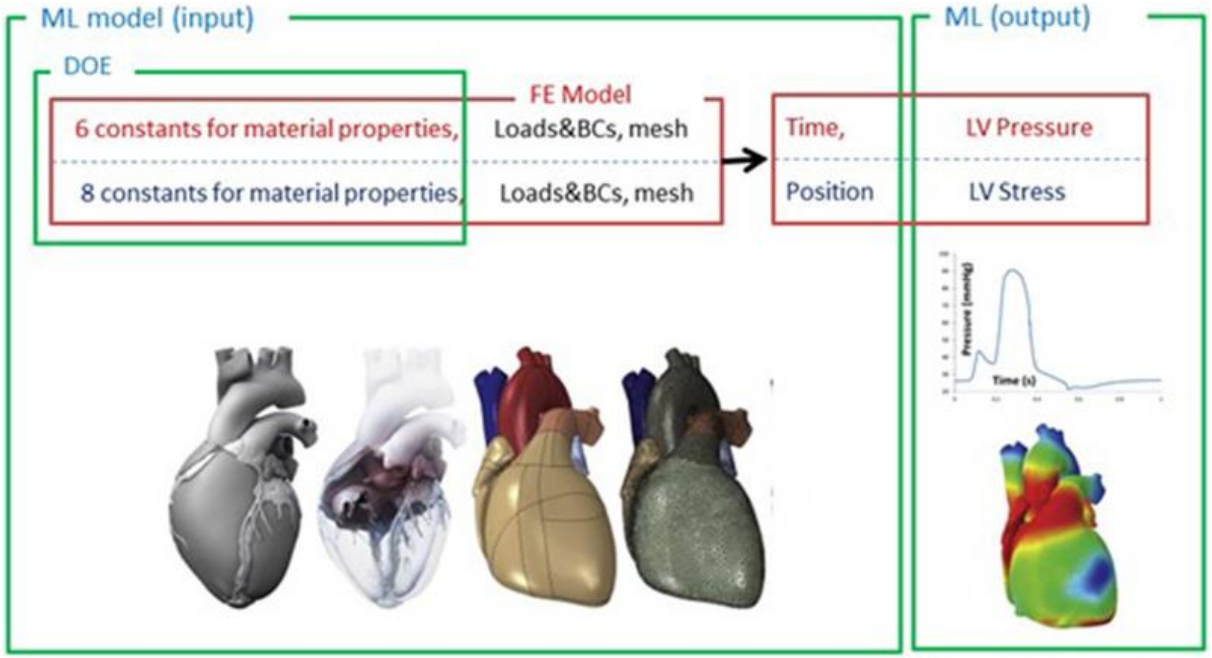

- ML Models evaluated in 2018 by studying Left Ventricle (LV) mechanics*
- Training set obtained by FE simulations for the LV with different material properties
- **XGBoost** eXtreme gradient-boosting decision tree algorithm used for **predictions of pressures, volumes, and stresses**, has been previously recognized in terms of accuracy, flexibility and speed.
- **Cubist** a similar package with smoother predictions, is a **rule-based model** where a tree is grown with terminal leaves containing linear regression models.
- **Next phase:** in addition to XGBoost, we will include Feed-forward DL and recurrent neural network (RNN) with long short-term memory (LSTM) models for LV pressure and volume for optimal MitraClip placement described in the scientific report below**

*) Dabiri et al. Prediction of Left Ventricular Mechanics Using Machine Learning. Front Phys. 2019; 7: 117.
**) Dabiri et al. Application of feed forward and recurrent neural networks in simulation of left ventricular mechanics, Scientific Reports, 2020, 10, 22298.

Machine Learning Models

- Evaluation of XGBoost and Cubist for heart mechanics predictions*

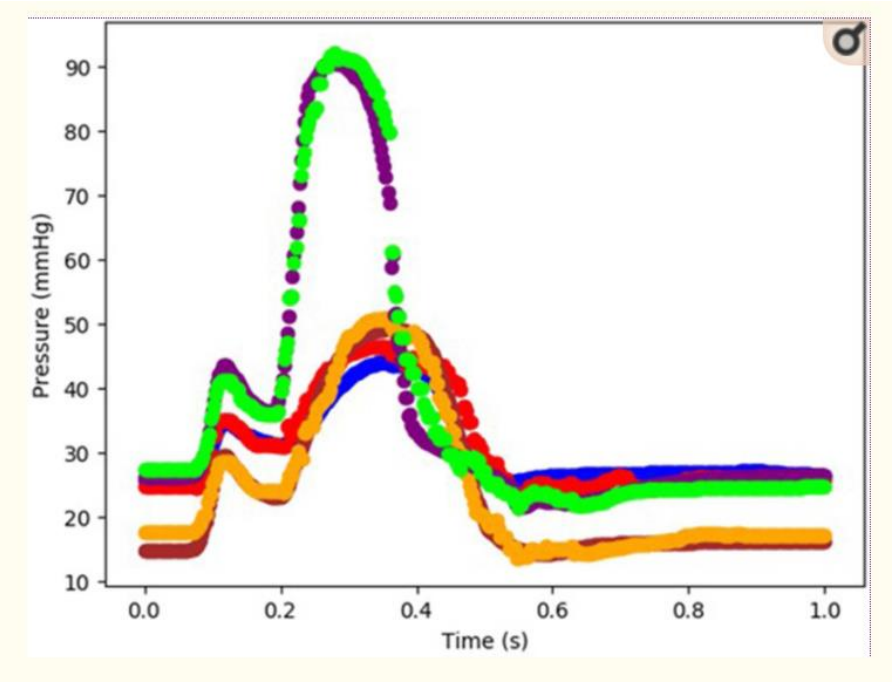

XGBoost: The FE-computed and ML-predicted LV pressure curve based on random selection of mechanical properties
FE: blue, brown, purple, ML: red, orange, green, respectively

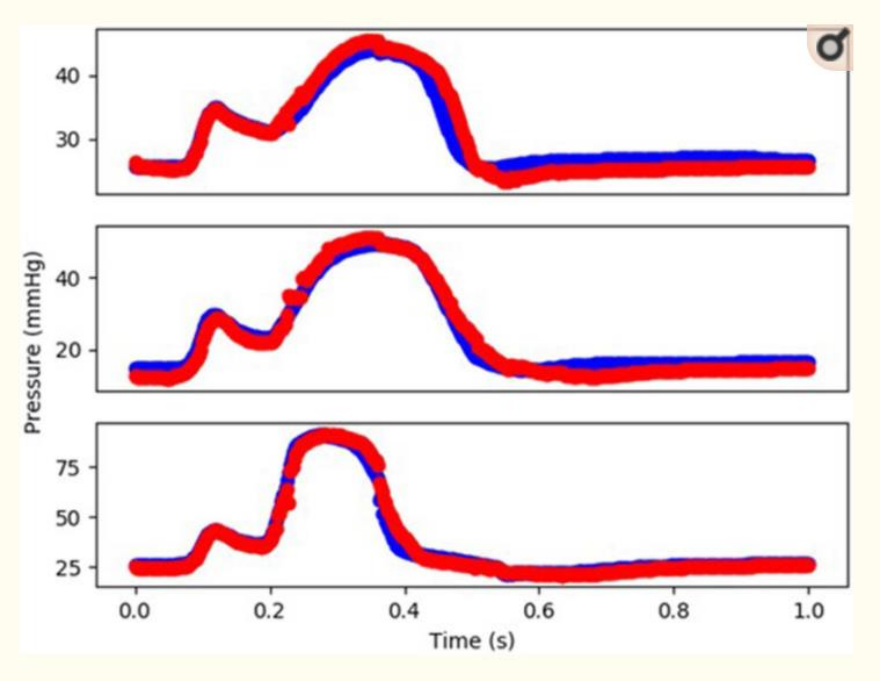

Cubist: The FE-computed and ML-predicted LV pressure curve based on random selection of mechanical properties
FE: blue, ML: red

*) Dabiri et al. Prediction of Left Ventricular Mechanics Using Machine Learning. Front Phys. 2019; 7: 117.
Dabiri et al. Application of feed forward and recurrent neural networks in simulation of left ventricular mechanics, Scientific Reports, 2020, 10, 22298.

Feed-forward DL and Recurrent Neural Network with long short-term memory (LSTM) models

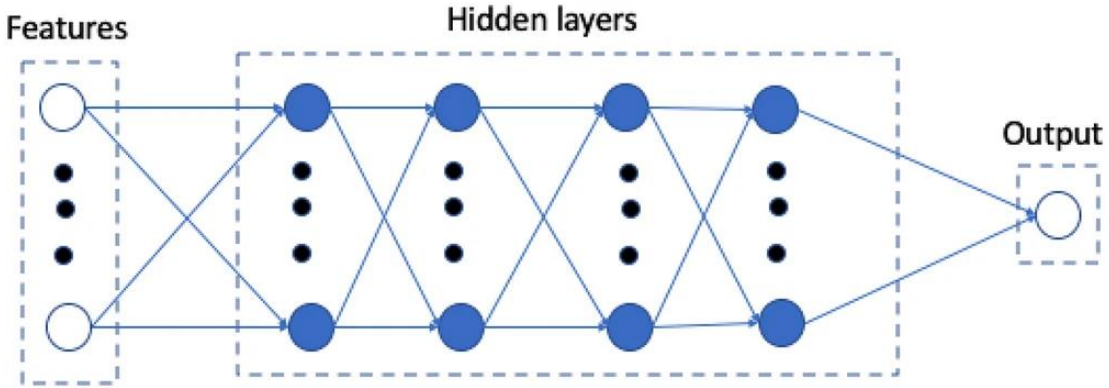

The feed forward DL model used for prediction of LV volume, pressure and stress. Details about the hidden layers in *)

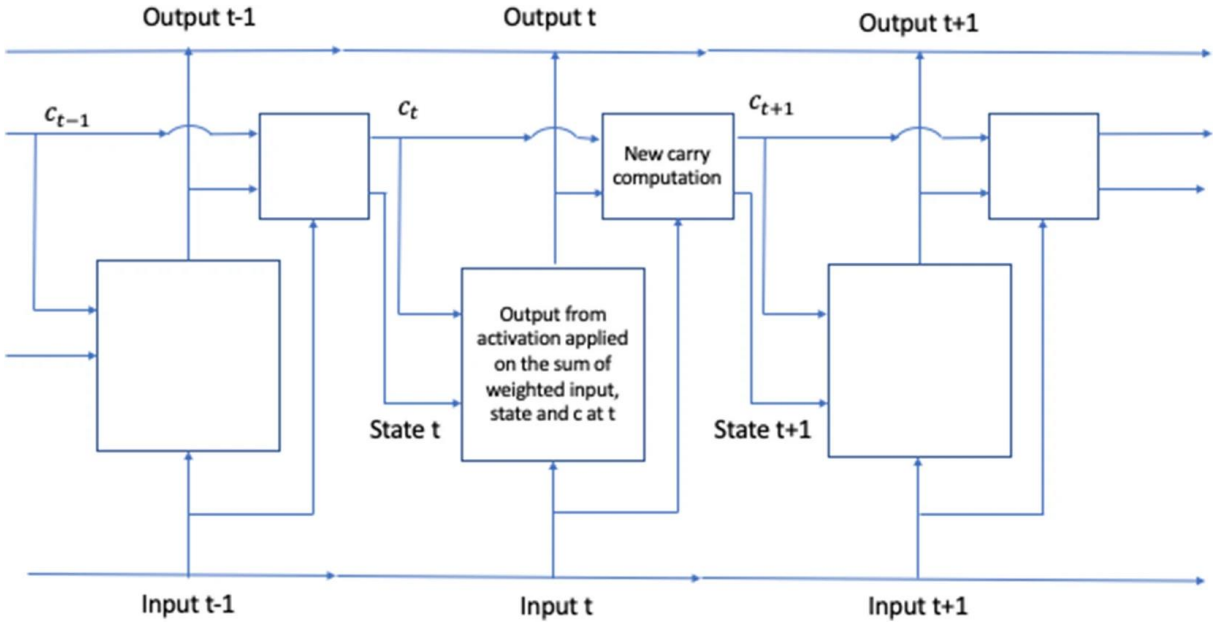

RNN with LSTM. LV pressure and volume from 20 previous time steps and active mechanical properties from current time step were used to predict current LV pressure and volume. Also, the LV endocardium stresses from 20 previous elements as well as passive mechanical properties from current time step were used to predict LV endocardium stresses at current time step. Here c represents carry state19. This model was developed in TensorFlow version 2.2.0. More details in *)

*) Dabiri et al. Application of feed forward and recurrent neural networks in simulation of left ventricular mechanics, Scientific Reports, 2020, 10, 22298.

UberCloud Engineering Simulation Platform on Azure

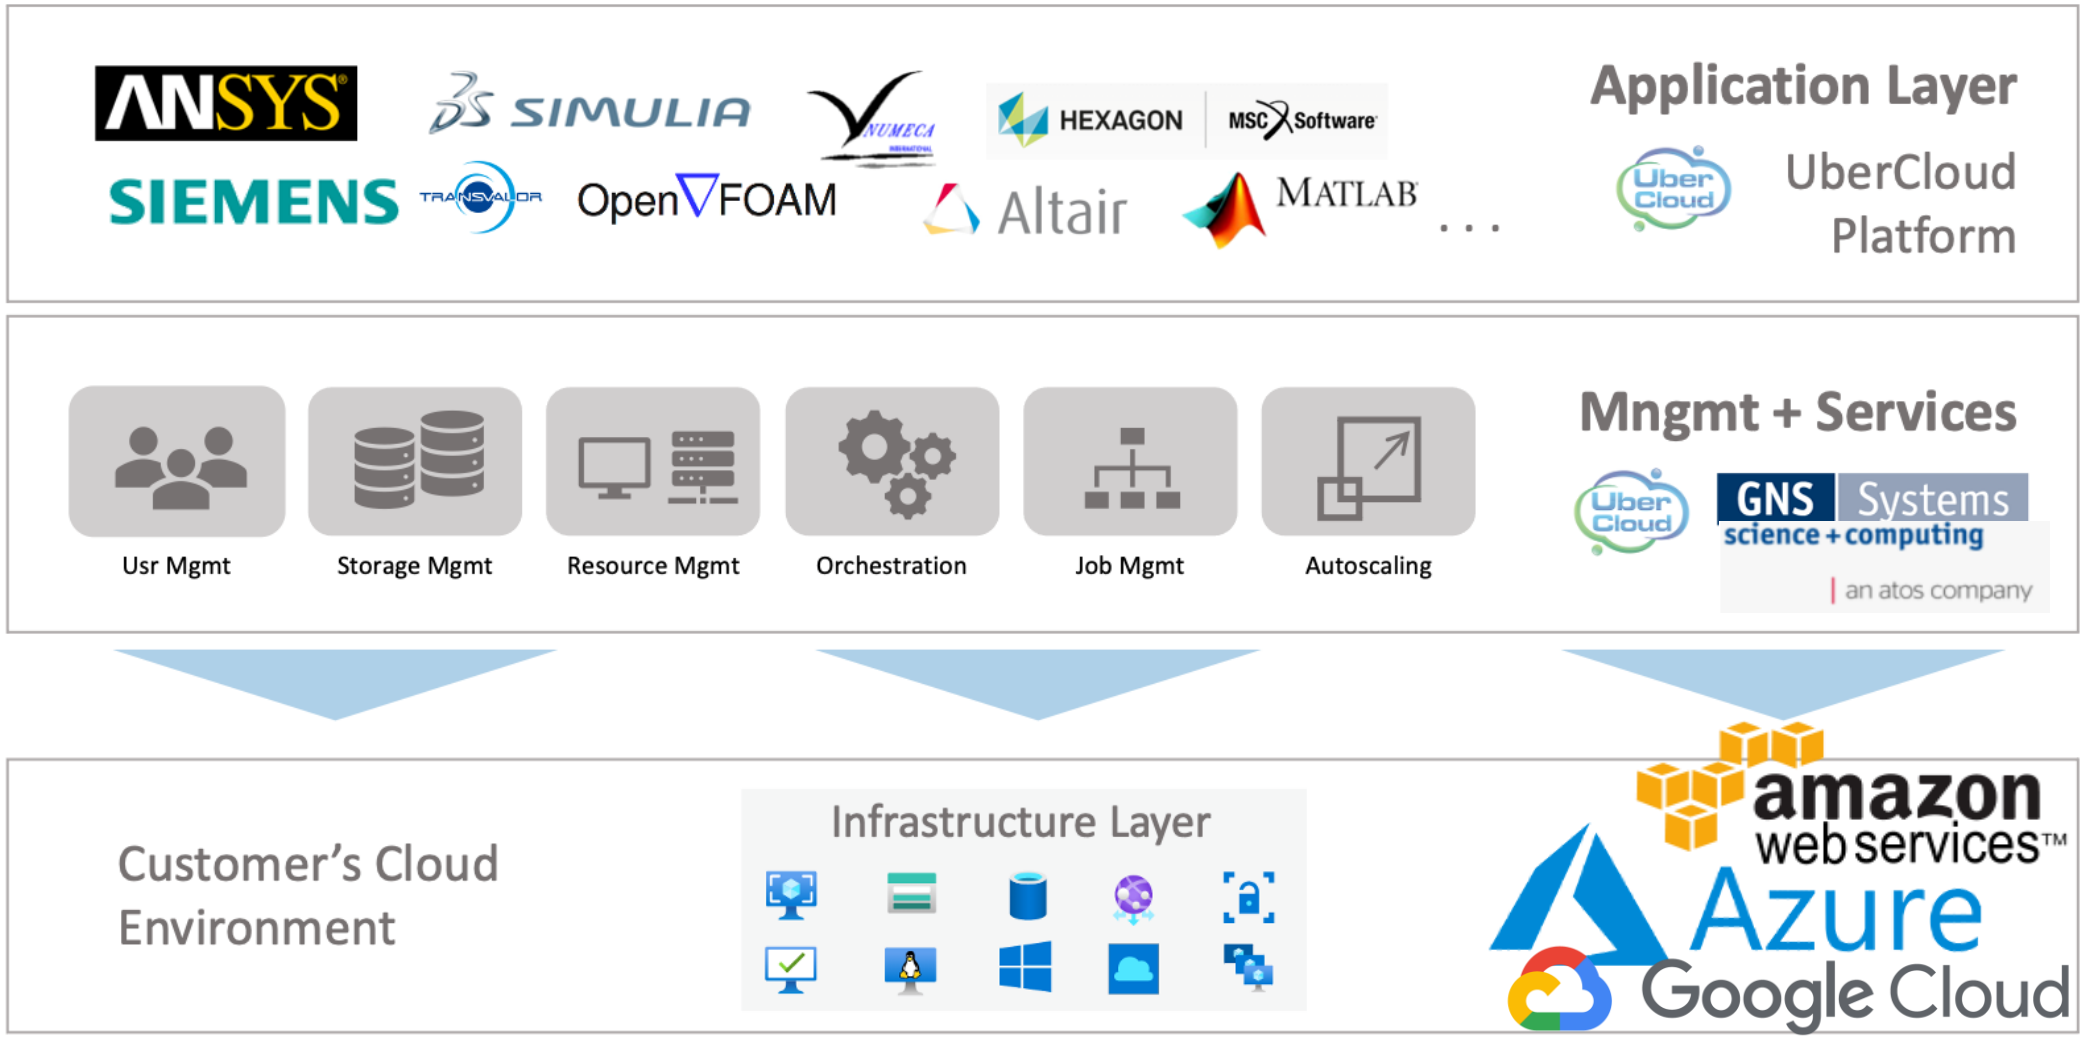

CAE Software Container Technology

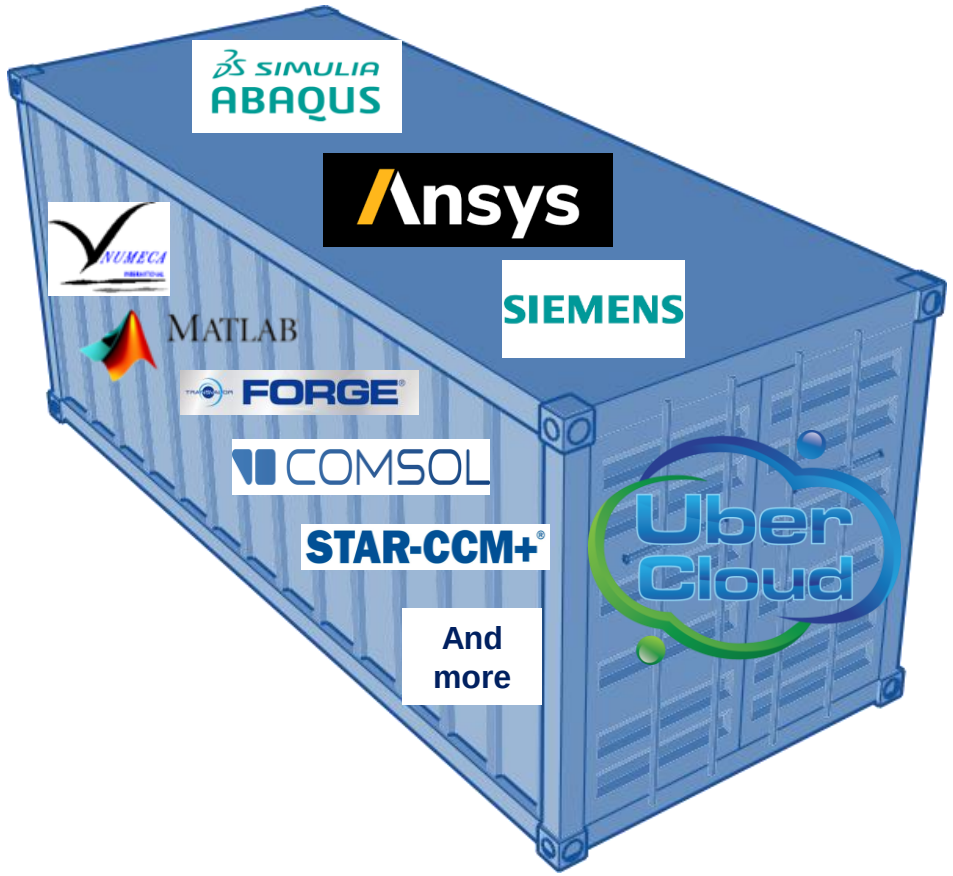

- CAE software containers, **enhanced** for engng
- Applications **pre-installed**, configured, and tested
- Fully managed service on-cloud and on-premise
- On any Linux, VMs, bare metal, OpenStack, vSphere, **Kubernetes**, SUSE Ranger, etc.
- Coupled with Linux and Windows Clients
- Hosting complex application workloads: Multi-Physics, Digital Twins, Data Analytics, AI/ML, Natural Language Processing, Personalized Healthcare, etc.

Main UberCloud's HPC container features

- **Multi-cloud**, public, private, hybrid
- **Multi-container** single-node and multi-node
- **Browser** access with Login and Password
- Any CPUs and GPUs
- **ssh** between containers / Infiniband
- Running within **VMs and on bare metal**
- Up-to-date **Linux** (CentOS, Ubuntu, SUSE,…)
- Screen sharing
- Data storage connectors (to Dropbox, Box.com)
- Emailing features (reporting job status etc.)
- Cloud **licensing** server
- **MPI** libraries, RDMA
- **Fully integrated with UC Engng Sim Platform**
- **Interactive and batch**
- No multi-tenancy
- VCollab data extraction & visualization
- OwnCloud privacy & security for cloud storage
- UberCloud storage containers, NFS, BeeGFS
- Rapid file transfer
- VPN
- NICE DCV virtual desktop and remote viz
- Automated password generation and change
- Performance checking
- Log monitoring
- Automated build/refresh process
- Images: pre-built ISV code & complex scientific and engineering workflows
- SLURM and Grid Engine container integration
- Automated QA testing suite
- Etc. etc.

New Addition to the UberCloud Platform:
SUSE: Powering Innovation With Leadership in Linux & Kubernetes for HPC

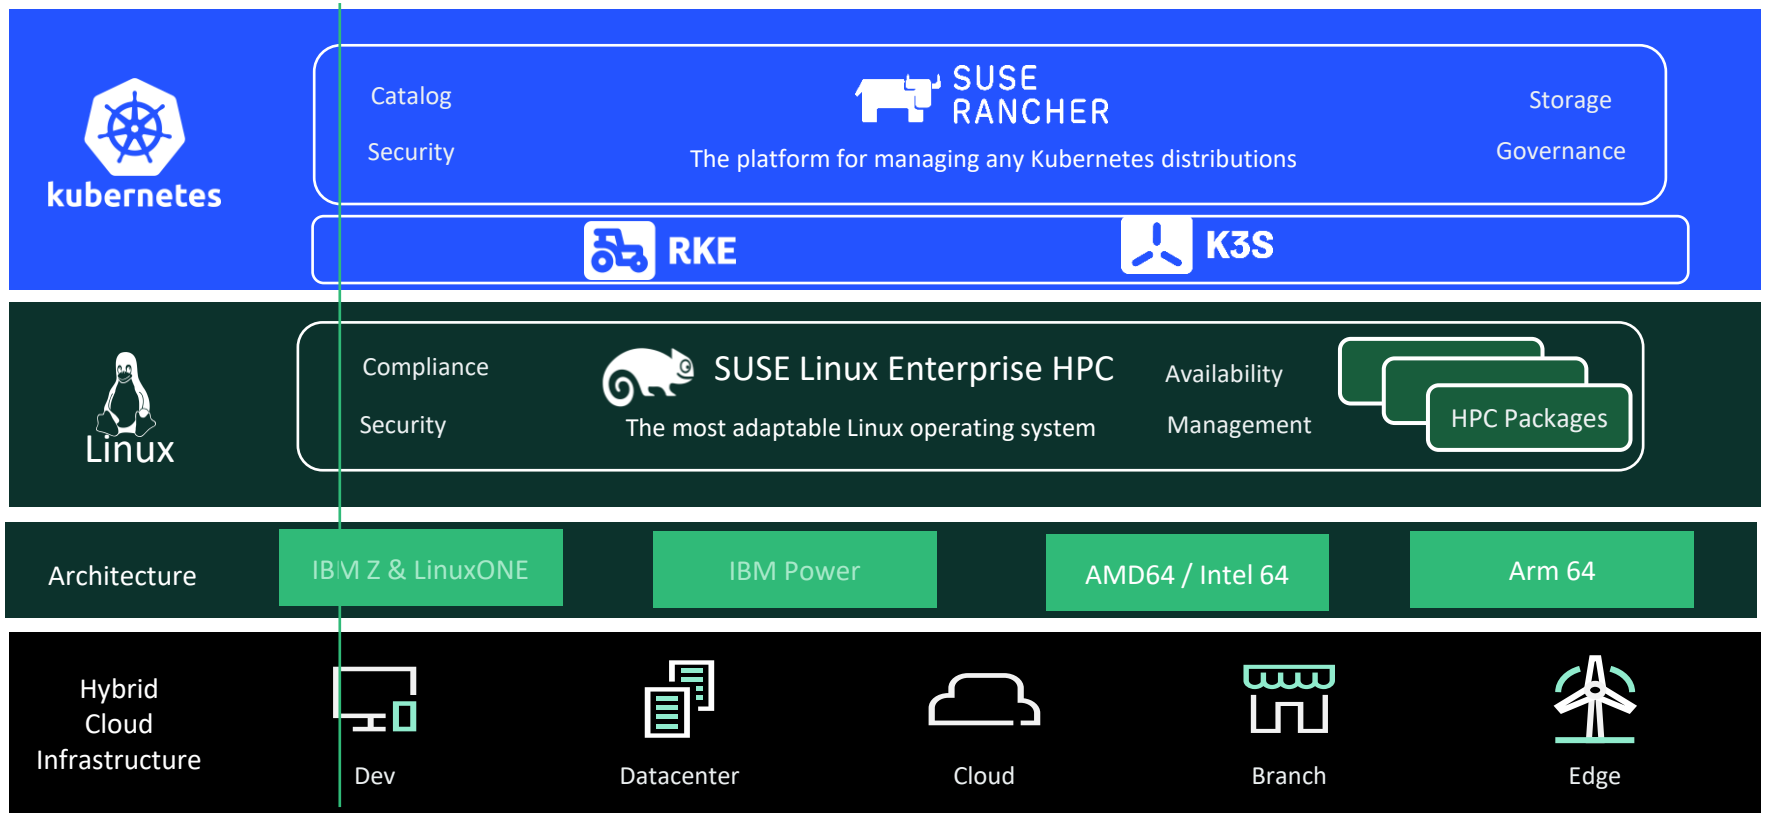

- ✓ **Flexible configuration and deployment**
- ✓ **Easier on-prem, hybrid and cloud management**
- ✓ **Proven HPC track record**
- ✓ **Supported HPC Tools**
- ✓ **Security and governance**
- ✓ **Broad of hardware support**
- ✓ **Strong partnerships and communities**

UberCloud's Next Generation HPC Application Platform

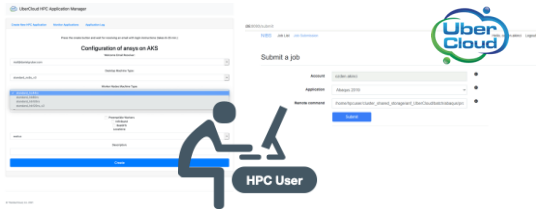

Self-service User Interface for the HPC engineer

uc *container* and *command line tool* for integration in **any HPC cluster / CI/CD / dev-ops pipeline /** container based workflow engine / job submission portal etc.

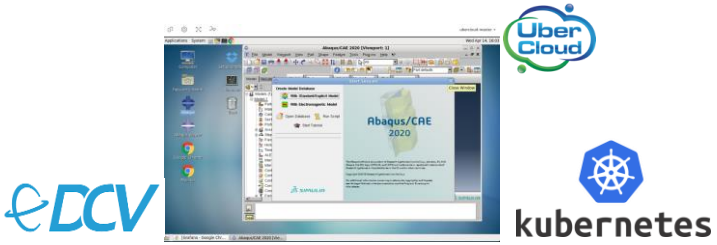

Cloud Desktop including dedicated, scalable, self-service, zero-maintenance HPC cluster.

Batch Job Processing

kubernetes kubernetes kubernetes kubernetes

Multi-Cloud Operations Platform

Cloud Provider Portal

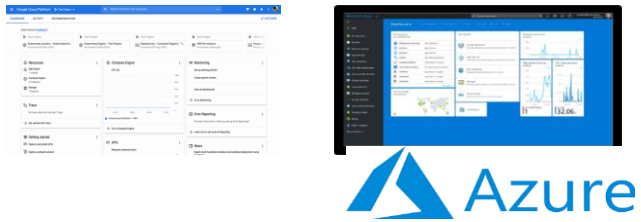

Enterprise / HPC IT Integration

On-demand Multi-Cloud HPC Resources

Day 2 Operations (Monitoring, Logging, ...)

Engineer‘s simulation workflow on Azure

starts desktop

gets login instructions per mail

works with the HPC applications

resizes cluster if required / runs simulations

shuts down cloud based infrastructure

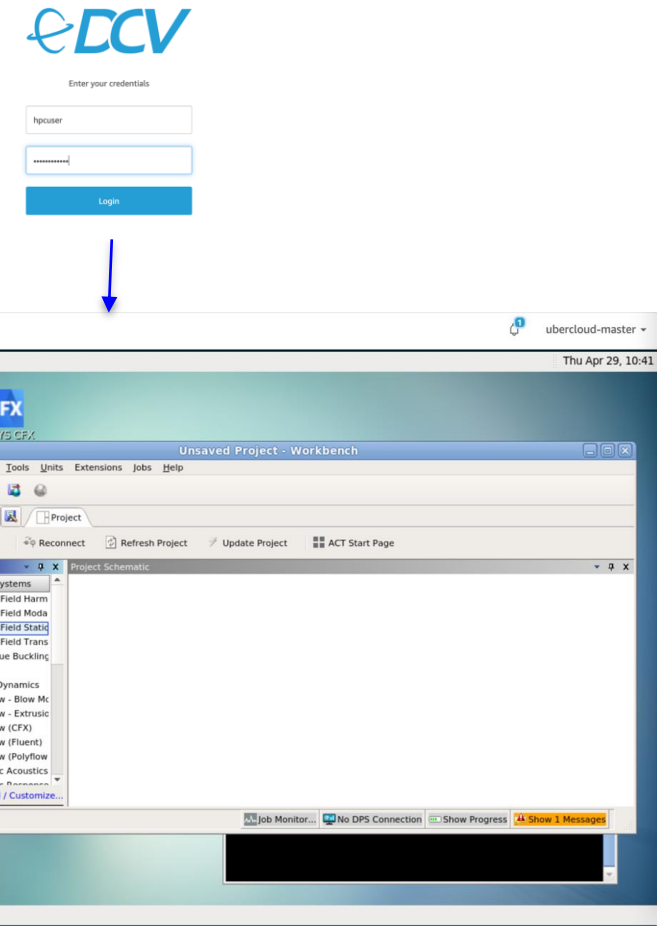

SUSE Rancher Monitoring

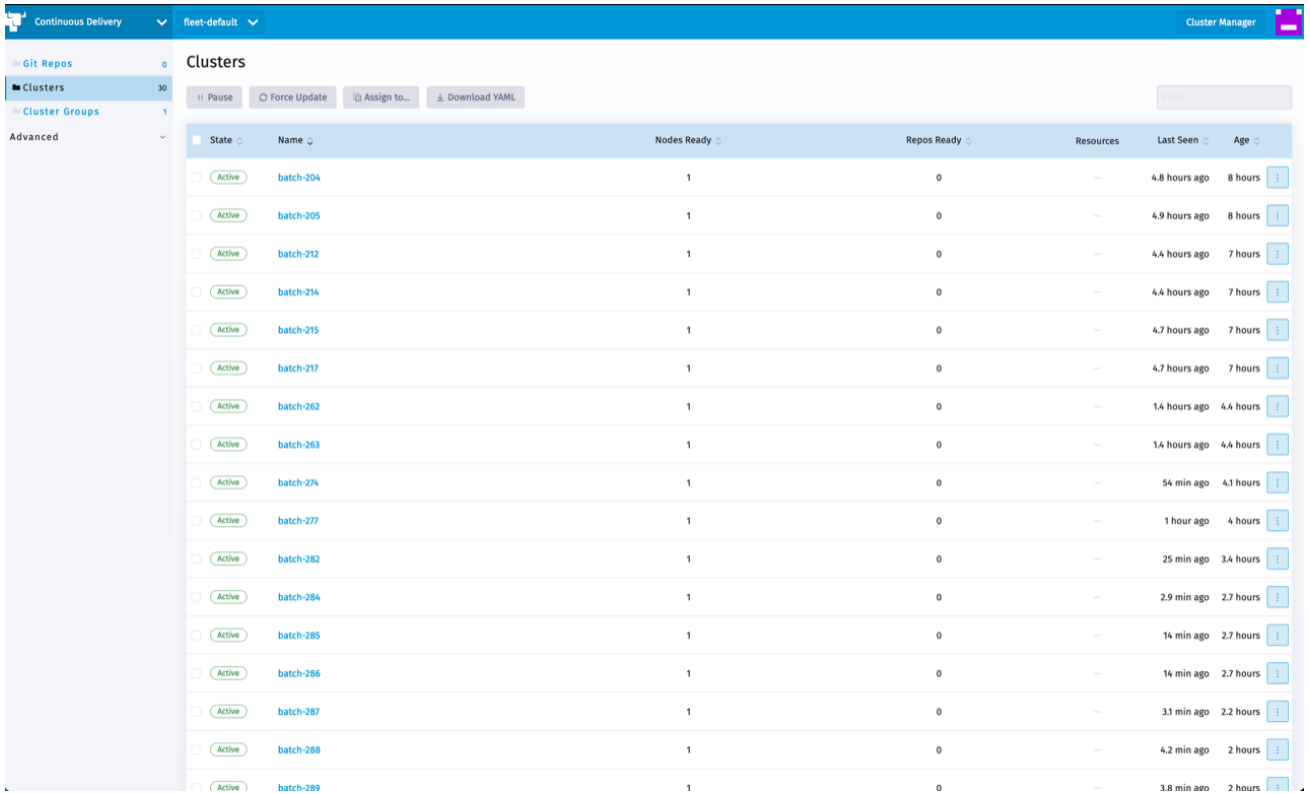

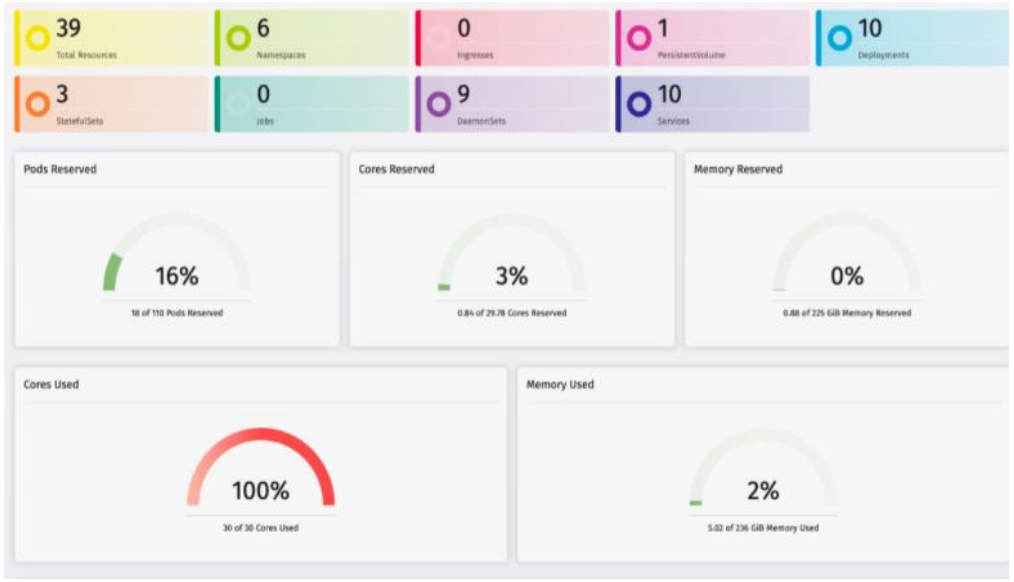

Living Heart Project Setup for 2500 Jobs on 2500 Clusters

- UberCloud Abaqus **desktop** (with GPU) for file staging and exploration of the simulation progress by the engineer

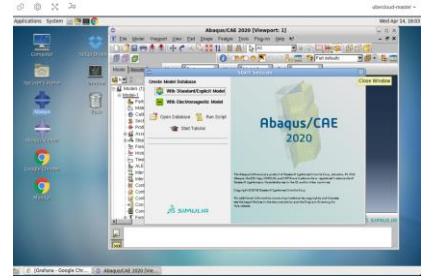

- **Batch processing** integrated in an HPC workload manager for job-throttling / queueing, and license management. Each job creates a Kubernetes cluster, runs simulation, and deletes the infrastructure.

```
uc create batch_N
uc run --clusterID batch_N -- abaqus inp=$INPUT job=$JOBNAME cpus=$CPUS
double=both interactive
uc delete batch_N
```

- In total 2500 Kubernetes clusters have been created. All of the clusters have been registered at SUSE's Rancher, allowing us to constantly monitor all infrastructure activities in any detail.

RANCHER

In Summary: Azure / SUSE / UberCloud Platform

- **HPC Containers:** unique, universal, standard, integrated with OpenStack, VMware, Kubernetes, SUSE Rancher, and many more
- **Public, private, hybrid, and multi-**cloud environments with **ONE set of tools**, fully integrated, providing the same look-and-feel experience everywhere
- Running in **customer's** own discounted cloud **account**
- **On dedicated,** secure HPC servers, no multi-tenancy
- Engineer has **sole control (rental car model),** 'owns' hw/sw stack, no 3rd-party interference/control
- **Customization** of complex engineering workflows, such as:
- Custom HPC, multi-physics, digital twins, personalized healthcare, big data analytics, AI/ML, NLP, . . . Anything data / compute intensive
- **Fully integrated and automated solution** allows for fast POCs and cloud on-boarding

Project Summary

- **Living Heart Project (LHP):** Repairing cardiac valve leakage (heart's mitral valve leaflets don't close tightly), most common cardiac valvular disease, with 5+ million people in the U.S.
- **Solution:** MitraClip (MC), a small metal clip allows doctors with catheter-based surgery to repair the mitral valve, brings minimally invasive alternative to open-heart surgery.
- **Multi-cloud Platform:** Compute: GCP+UC Platform, K8s Management: SUSE Rancher in UC subscription, on **any** cloud, Abaqus container + DCV remote viz, Dassault's LHP model
- **Approach:** One Abaqus job on engineer's on-prem **workstation*** = 12 hours

2,500 jobs on workstation = 1,250 days

For less than $20K

One Abaqus job on **GCP** on 1 compute node = 4 hours

On **GCP** 50 x 50 jobs in parallel = 9 days

=> ML prediction with 95+ % accuracy = 2 secs

*) Compute node: two Intel Xeon E5-2680 v4 processors, each with 14 cores

Thank You

Contact us at help@TheUberCloud.com

Gartner 2015 Cool Vendor