



Gaining a Competitive Edge with HPC Cloud Bursting

William J Edsall

The Dow Chemical Company

R&D Information Research – High Performance Computing Team

September 5th 2018

About the presenters

William J. Edsall – The Dow Chemical Company

High Performance Computing Architect

- BA Applied Mathematics
- 11 years HPC experience, 15 years IT
- 3 kids ages 0, 2 and 4
- Self-proclaimed “ex cloud critic”



About Dow High Performance Computing

On premise HPC capabilities


- Two Intel CPU based HPC Clusters
 - >85% AVG utilization
- Nvidia K80 and V100 coprocessors
- High end HPC desktops

HPC workload at Dow

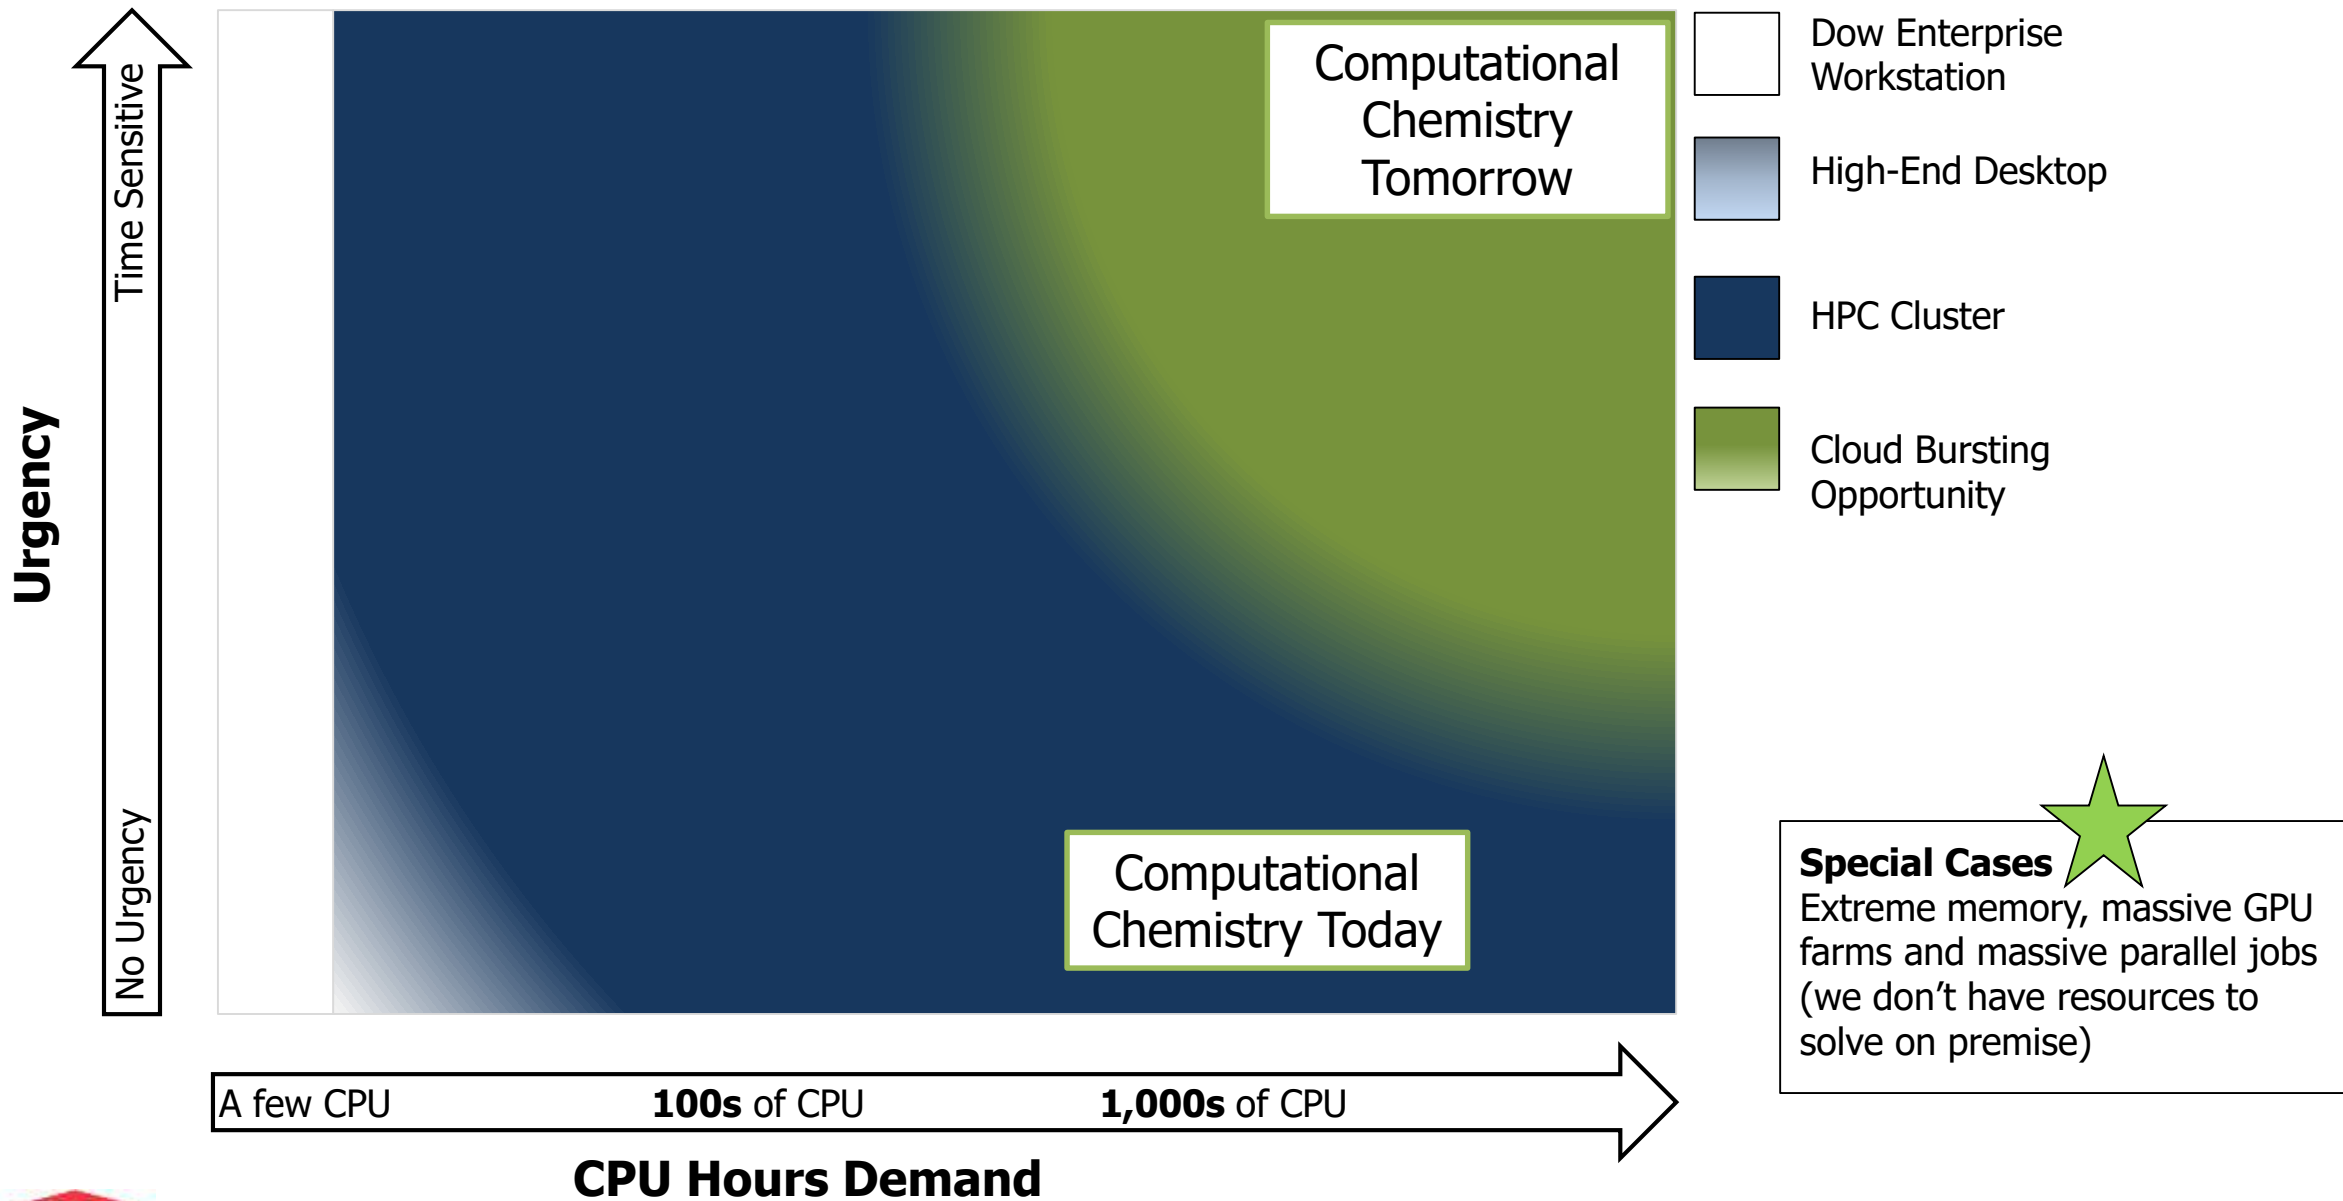
- Computational Chemistry
- Computational Fluid Dynamics
- Molecular Dynamics
- Computer Aided Engineering
- Finite Element Analysis
- Optimization
- Deep Learning / Machine Learning
- X-ray Tomography

Executive Summary

GOAL: Develop processes and leverage a hybrid On Premise / Cloud Bursting environment by 2019.



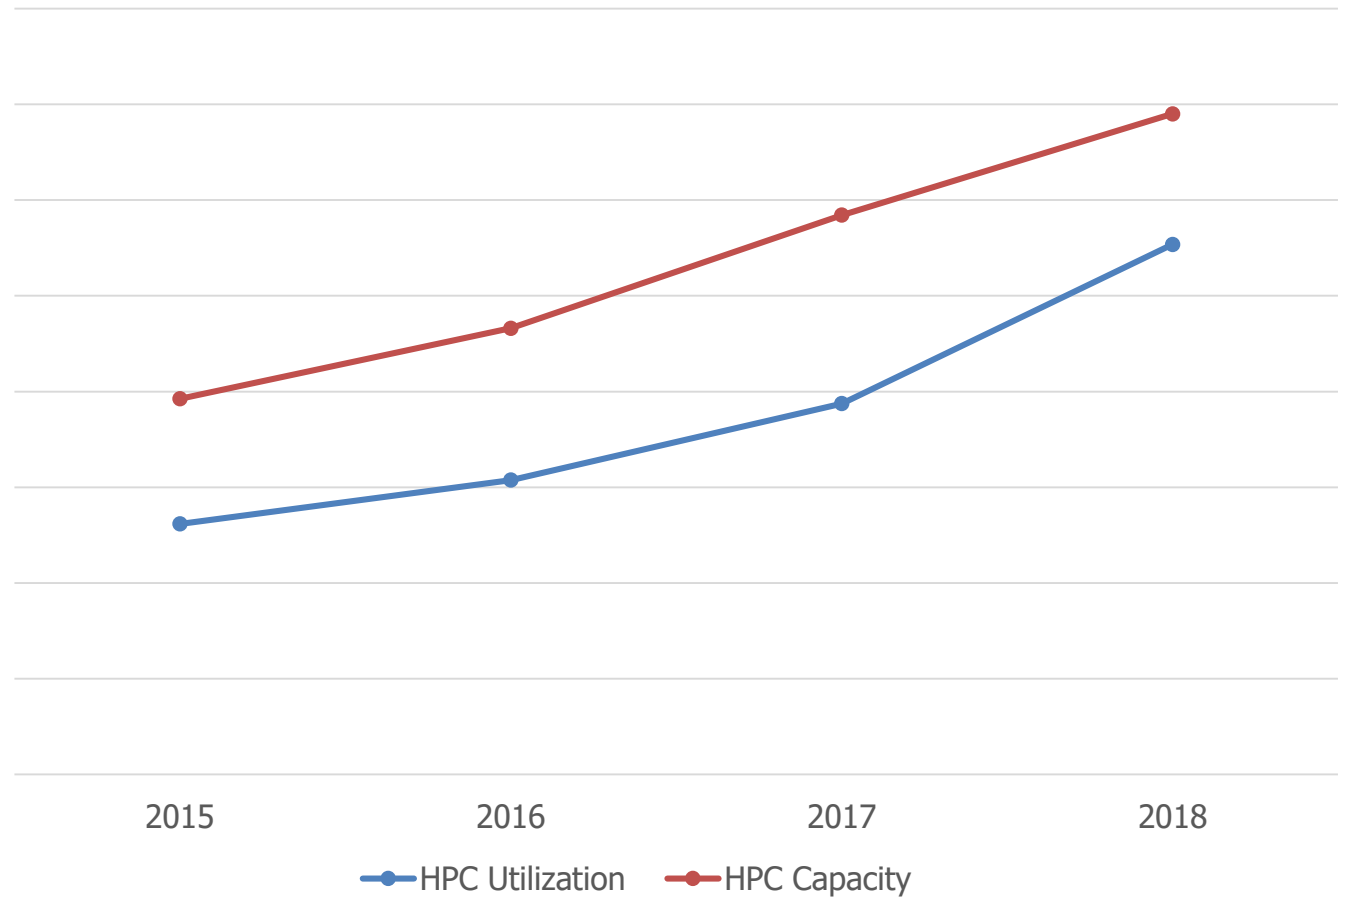
Cloud Bursting Opportunity



Challenges with on premise HPC Model

- True demand difficult/impossible to capture
- Utilization always follows capacity trend

Dow HPC Capacity / Utilization Trend 2015-2018



Challenges With HPC Cloud Bursting Model

HPC Cloud Data Security Model

Cost Recovery Model

Azure Active Directory Integration

Independent Software Vendor (ISV) Licensing

Data Transfer (Dow to/from Azure)

Rescale Cloud Bursting Proof of Concept at Dow



~2x the total CPUs on premise at Dow

Computational Chemistry

- 1,400 molecules from ChEMBL database
- Each molecule calculated with 7 CPU using Azure H16r instance
- Entire ~9,800 CPU cluster running in <30 minutes
 - 0 queue time (no queue exists)
- Cluster was shut down as jobs finished
 - Charges stopped automatically
- **Goal met in 5 hours**

Work performed by Dr. Peter Margl, The Dow Chemical Company

~2x larger than the biggest parallel job at Dow

Computational Fluid Dynamics

- Ansys FLUENT simulation
- Scaled up 20 -> 200 -> 2,000 CPUs
- Entire 2,000 CPU run on Azure H16r instances
 - One Infiniband fabric
- Cluster was stood up – job was verified and then stopped
- **Goal met in 15 minutes**

Work performed by Dr. Chang Kai (Lance) Wu, The Dow Chemical Company

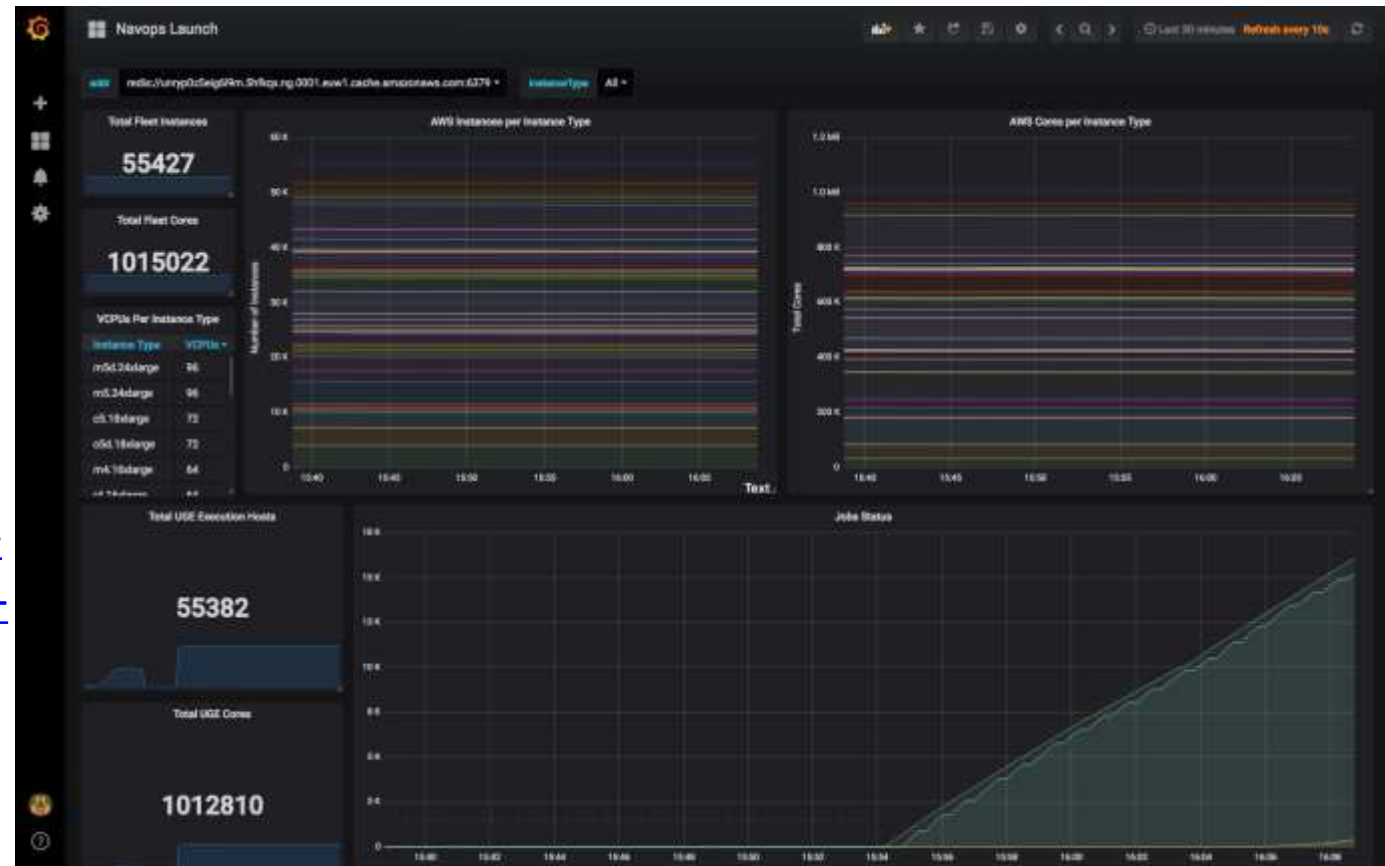
Scale Beyond our POC

Massive Scale Example: Univa (June 27th 2018)

- “Univa leveraged AWS to deploy **1,015,022 cores** in a single Univa Grid Engine cluster”
- “The cluster was built in approximately 2.5 hours”

Reference Link:

<https://www.hpcwire.com/off-the-wire/univa-deploys-over-1m-cores-in-a-single-grid-engine-cluster-using-aws/>



Q&A