

Roswell Park IT Research Computing



L. Shawn Matott, PhD
Senior Systems Analyst

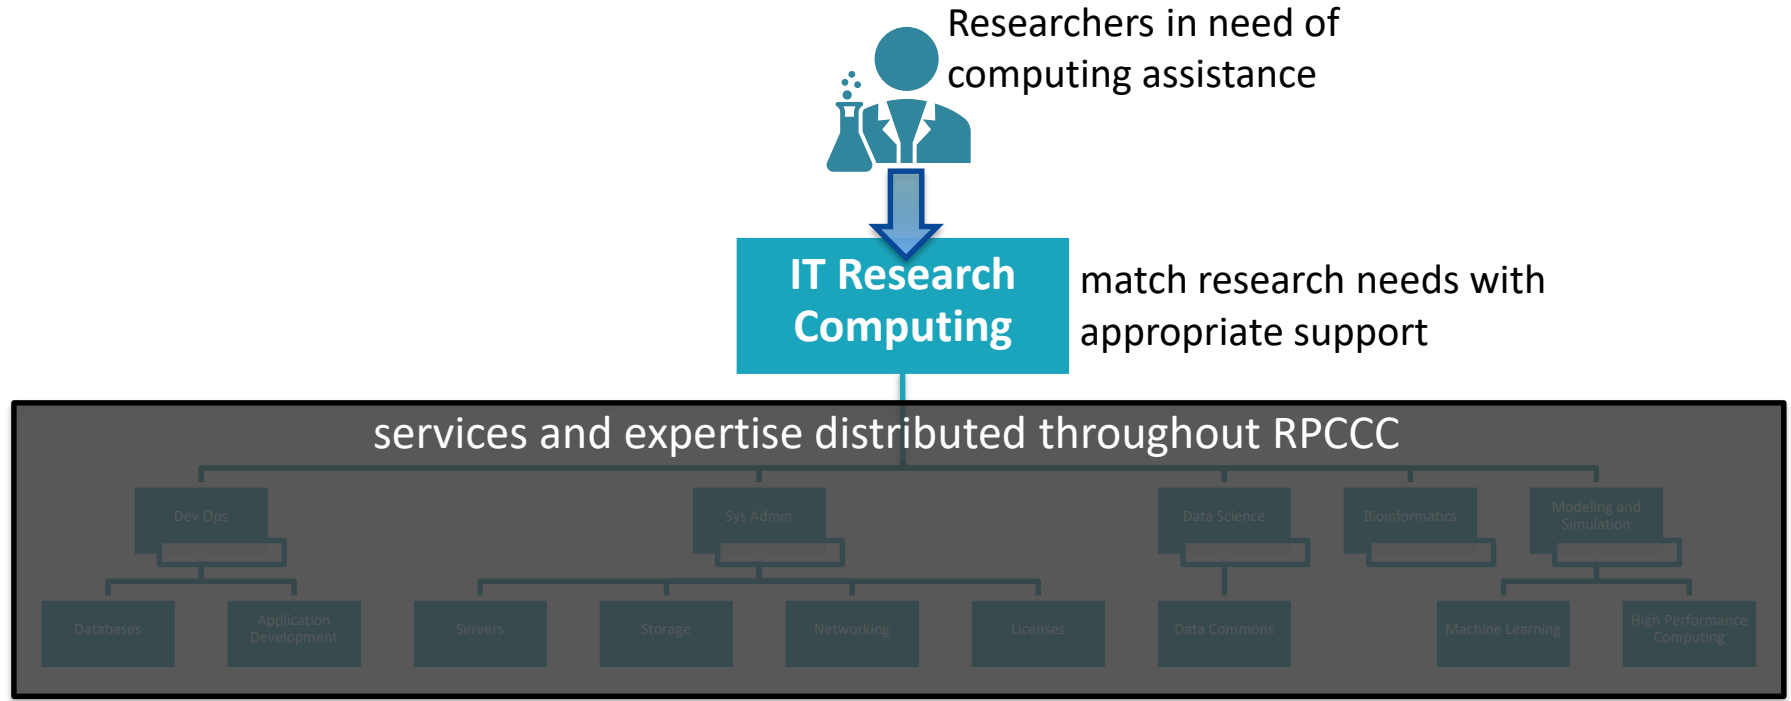
Overview

- IT Research Computing
 - an IT initiative of the new CIO (Tom Furlani)
- Mission
 - *To enable and accelerate research throughout campus by promoting advanced computing and computational science.*



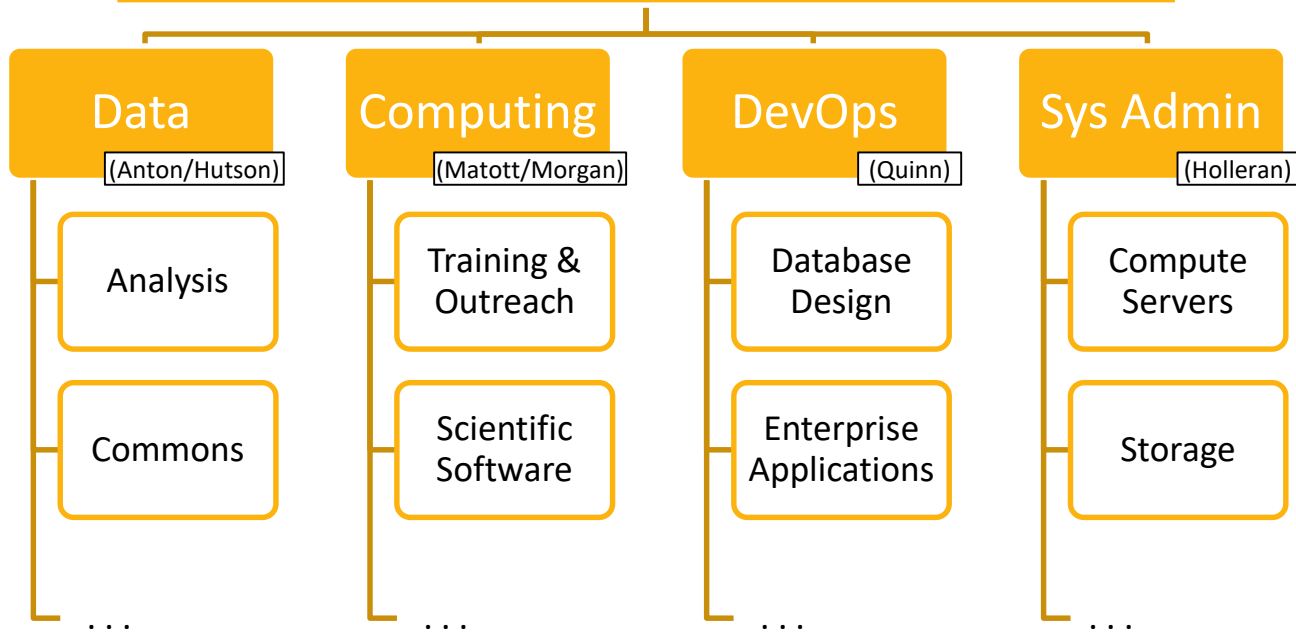
A look inside Roswell's HPC data center.

IT Research Computing – User Perspective




Components of IT Research Computing

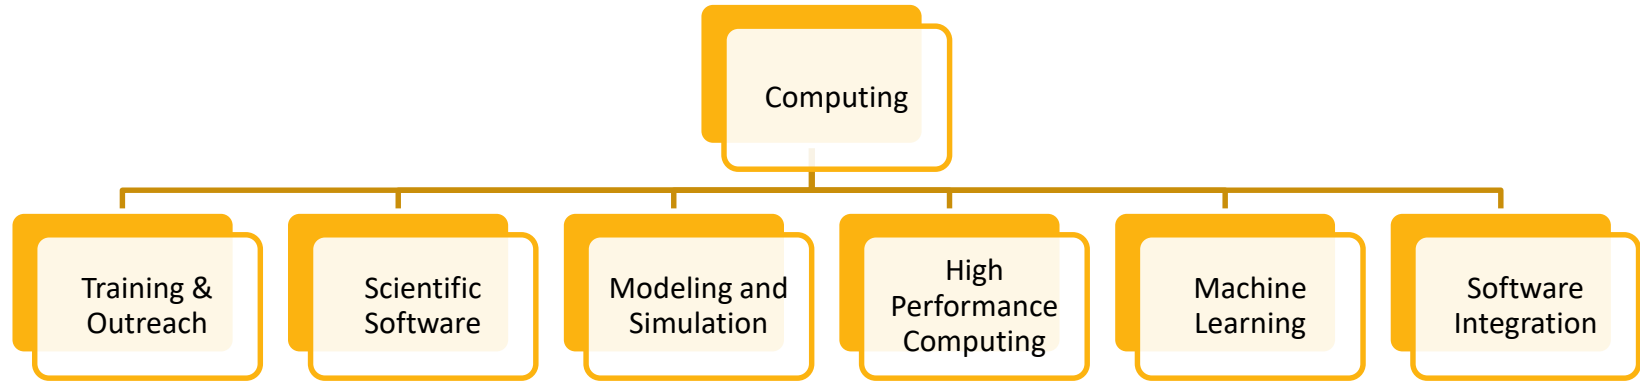
IT Research Computing



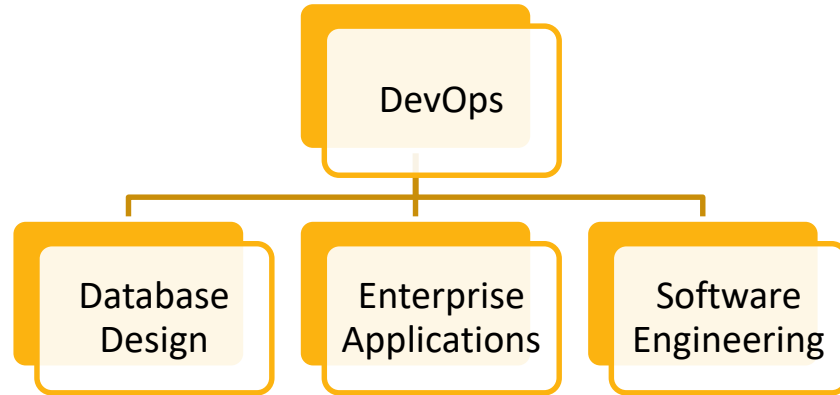
IT Research Computing – Data Services




IT Research Computing – Computing Services



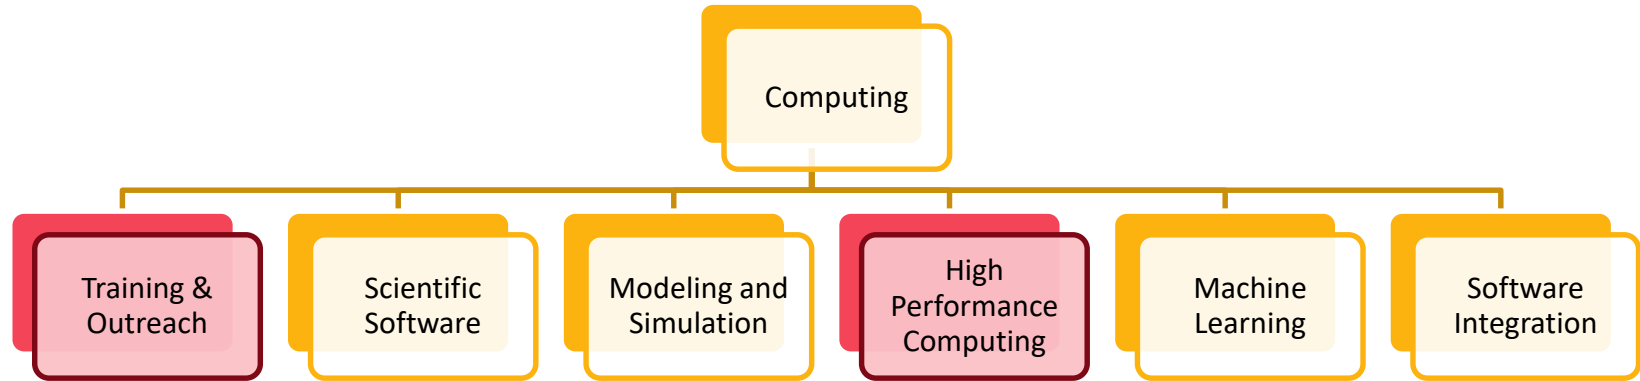
IT Research Computing – DevOps Services



IT Research Computing – Sys Admin Services




IT Research Computing – HPC Training



Computing Assets

- Roswell HPC

- 1,600 processors, 450 TB Lustre storage
- Software modules include: Grace, EGSnrc, Keras, and TensorFlow
- OnDemand web interface
<http://u1.roswellpark.org/pun/sys/dashboard/>
- e-mail ServiceDesk@RoswellPark.org to request access

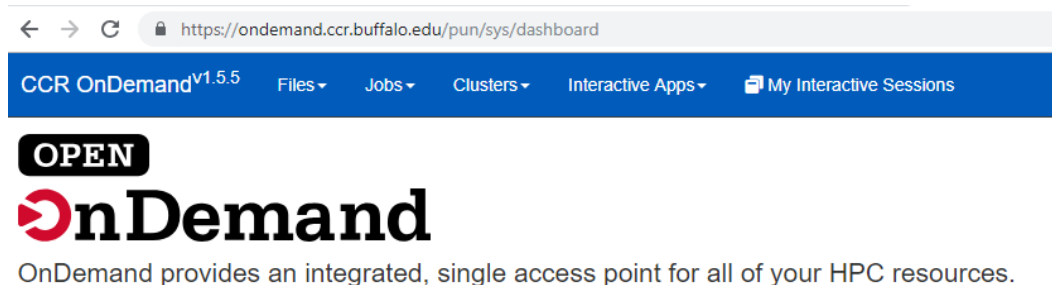


OnDemand provides an integrated, single access point for all of your HPC resources.

Computing Assets

- UB CCR

- 20,000+ processors
4.2 PB storage
- GPU hardware for machine learning acceleration
- 400+ S/W packages
- OnDemand web interface
<https://ondemand.ccr.buffalo.edu/pun/sys/dashboard/>
- Visit www.buffalo.edu/ccr/support/ccr-help/accounts.html to request access
- Currently: freely available to UB-affiliated researchers
- Coming soon: a portion of CCR will be available to *all* Roswell researchers!



Training

- On-boarding of new HPC users
 - Offered quarterly
May, August, November, & February
 - Initial login
 - Remote access
 - OnDemand interface
 - Job submission
 - Interactive applications
 - File management
 - Backup, archiving, security, etc.
 - How to request additional help and/or training



Training

- Specialized training
 - Offered on request
 - Training on running specific software tools and workflows in a cluster environment
- Examples:
 - R, Bioconductor
 - Bioinformatics pipelines
 - Engineering codes (COMSOL, ANSYS, etc.)
 - EGSnrc and dosimetry
 - Machine learning: TensorFlow, Keras, etc.
 - Python, MatLab, C/C++




Source: HP



Source: LLNL

Training Examples - Python

- Intro. To Python programming
 - Basic syntax and unique features
 - Creating functions and loading modules
- Python for data science
 - Pandas, StatsModels, PyPlot, NumPy/SciPy
- Python for machine learning
 - Machine learning basics
 - TensorFlow, Keras, PyTorch
 - Image classification, natural language processing
- Python for HPC
 - Multiprocessing, MPI4Py



Pandas



Technical Support - Helpdesk

- Available through ServiceNow
 - Request IT, Fix IT, and Know IT
- Use Request IT to:
 - Request software installation
 - Request new accounts, storage, training, etc.
- Use Fix IT for:
 - Troubleshooting failed jobs, storage quotas, hardware failures, etc.
 - Password resets, account problems, etc.
- Use Know IT for:
 - Online documentation and FAQs
 - Archives of workshops and seminars



Source: intel/facebook

A typical IT Research Computing Engagement

Discovery
Initial meeting between
research group and IT RC.

Preparation
IT RC performs S/W installation
and template development to
meet needs of research group.

Setup/Training
Account creation and
training of new users.

Data Migration
Research group migrates
data into HPC storage.

Ongoing Research
Ongoing usage of
compute and storage
assets with periodic IT
support requests.