TACC Quick Start Guide

What is TACC

The Texas Advanced Computing Center (TACC) is a Research Center at the University of Texas at Austin. TACC designs and operates some of the world's most powerful computing resources, with the mission to “enable discoveries that advance science and society through the application of advanced computing technologies.” The center’s comprehensive cyberinfrastructure ecosystem includes high performance computing (HPC), visualization, data analysis, storage, archive, cloud, data-driven computing, connectivity, tools, APIs, algorithms, consulting, and software. For more information, please visit: https://www.tacc.utexas.edu/

Eligibility

All UTRGV faculty, students, and staff are eligible to create a TACC account and use a variety of TACC resources. As a general rule, faculty, postdocs, and senior staff scientists are considered to be Principle Investigators (PIs) in the view of TACC. PIs are eligible to request CPU hours and storage space on TACC computing resources through the Projects and Allocations system (see below). Graduate students are not considered to be PIs in the view of TACC, and therefore cannot request CPU hours or storage space. A registered PI (e.g. the graduate student’s advisor) must request the appropriate resources and add the student as an approved user. See: https://portal.tacc.utexas.edu/tutorials/managing-allocations

Creating an Account

Creating a TACC account is fast and simple using your UTRGV credentials. To begin, please visit the following link: https://portal.tacc.utexas.edu/utdr

Fill out all fields in the account request form, and if appropriate, check the box labeled “I am PI Eligible.” Watch your e-mail for a message from TACC asking you to confirm the account creation by clicking a link. Please read all of the instructions in the e-mail carefully. The final step of the account creation process is to use your new credentials to log in to the TACC User Portal: https://portal.tacc.utexas.edu/

From the TACC User Portal, sign out if you are already signed in, then click on the large green ‘Log in with TACC Account’ button near the top right. The username and password used to log in to the TACC User Portal will be the same username and password you use to log in to the computing systems. You will not be able to request allocations or be added to allocations until you log in to the portal for the first time.

Projects and Allocations

TACC grants CPU hours – also called Service Units (SUs) – to users based on a Project and Allocation system. A project is a group of faculty, students, and / or staff working on a common, specific research goal. Each project has one designated PI, and each user may be associated with multiple different projects. An allocation is a specific allotment of resources (including SUs and / or storage space) to be shared by all users associated with a project. Each project may have multiple allocations, but each allocation may only be associated with one project.

Detailed instructions exist on the TACC website to create and manage Projects and Allocations. Please see the following:

https://portal.tacc.utexas.edu/allocations-overview

https://portal.tacc.utexas.edu/tutorials/managing-allocations

Training and Further Help

New TACC users are encouraged to attend TACC training: https://portal.tacc.utexas.edu/training/ when possible. A variety of courses are offered, covering topics in high-performance computing, scientific visualization, and distributed and collaborative computing. Course attendees are provided with all presentation materials for future reference The majority of training is hosted in Austin, TX, and users typically can register to attend remotely via webcast.

TACC Quick Start Guide

TACC staff members also make available and maintain a set of user guides: https://portal.tacc.utexas.edu/user-guides to assist users to continue to improve their knowledge of the resources. During normal business hours, personal assistance from the TACC consulting staff can be requested via the TACC Consulting system https://portal.tacc.utexas.edu/tacc-consulting available through the TACC User Portal.

System Guides

Lonestar 5 (Cray XC40 Cluster with Intel Xeon)
HPC, Remote Visualization

The Lonestar 5 (LS5) system is designed for academic researchers in Austin and across Texas. The machine is a Cray XC40 customized by the TACC staff to provide a unique software environment designed to meet the needs of our diverse user community. Lonestar 5 provides 1252 twenty-four core general compute nodes (for a total of 30,000 processor cores), 16 GPU nodes, and 10 large memory nodes. The system is configured with over 80 TB of memory and 5PB of disk storage, and has a peak performance of 1.2PF.

Stampede (Dell Power Edge Cluster with Intel Xeon Phi coprocessors)
HPC, Remote Visualization, Data Analysis, Data Intensive Computing

Stampede entered production in January 2012 as a 6,400+ node cluster of Dell PowerEdge server nodes featuring Intel Xeon E5 Sandy Bridge host processors and the Intel Knights Corner (KNC) coprocessor, the first generation of processors based on Intel's Many Integrated Core (MIC) architecture. Stampede's 2016 Intel Knights Landing (KNL) Upgrade prepares the way for Stampede 2 by adding 508 Intel Xeon Phi 7250 second-generation KNL MIC compute nodes. The KNL represents a radical break with the first generation KNC MIC coprocessor. Unlike the legacy KNC, a Stampede KNL is not a coprocessor: each KNL is a stand-alone, self-booting processor that is the sole processor in its node.

Corral (Petabyte-scale Dell / DDN storage cluster)
Storage, Data Management

Corral leads the way in the preservation and sharing of data for researchers. This storage and data management resource is designed and optimized to support large-scale collections and a collaborative research environment. Corral consists of two 6-petabyte (>4 petabytes usable capacity) disk arrays, installed at UT Austin and UT Arlington, along with 24 Dell servers providing high-performance storage and services for all types of digital research data.

Maverick (HP/NVIDIA Visualization and Data Analysis Cluster)
Interactive Visualization, Data Analytics, GPGPU Calculations

Maverick, an HP/NVIDIA Interactive Visualization and Data Analytics System, combines capacities for interactive advanced visualization and large-scale data analytics as well as traditional high performance computing. Recent exponential increases in the size and quantity of digital datasets necessitate new systems such as Maverick, capable of fast data movement and advanced statistical analysis. Maverick consists of NVIDIA K40 GPU for remote visualization and GPU computing.

Ranch (Oracle Sun Starage Archive Manager Filesystem)
Mass Archival Storage

Ranch, a long-term mass storage solution, has 2 petabytes of online storage for data transfer and capacity for 160 petabytes of offline tape storage. By providing a massive file system for archival purposes, this system serves the HPC and Visualization communities.

Jetstream (User-Customizable Dell Cloud Resource)
Cloud Computing

TACC Quick Start Guide

Jetstream is the first user-friendly, scalable cloud environment for XSEDE. The system enables researchers working at the "long tail of science" and the creation of truly customized virtual machines and computing architectures. It has a web-based user interface integrated with XSEDE via Globus Auth. The architecture is derived from the team's collective experience with CyVerse Atmosphere, Chameleon and Quarry. The system also fosters reproducible, sharable computing with geographically isolated clouds located at Indiana University and TACC.

Hikari (HPE Apollo 8000 system)
HPC, Secure / Compliant Computing

Hikari, an HPE Apollo 8000 system, is the result of a collaboration between TACC, the New Energy and Industrial Technology Development Organization (NEDO), a Japanese government agency, and NTT FACILITIES INC. The system consists of 432 nodes, each with dual Intel 12-core "Haswell" Xeon processors, and delivers a peak performance of over 400TF. Hikari will be used to support TACC's secure/compliant computing needs (HIPAA/FISMA compliant data), as well as the increasing demand for fast turnaround on the many jobs received through web APIs from Science Gateways.

Wrangler (Dell / EMC Data Analytics System)
Data Analysis, Data Management

The Wrangler Data Analysis and Storage* system is designed for the needs for modern data researchers. Wrangler's unique architecture handles the many aspects of the volume, velocity, and variety that can make digital data research difficult to handle on standard high performance systems. The system is designed around a 0.5 PB high speed flash storage system that can be used to handle data analysis and processing workflows not practical on other systems with slower spinning disks or significantly smaller internal SSD storage devices.

See more TACC systems here: https://www.tacc.utexas.edu/systems/overview

TACC Quick Start Guide

Examples of project use on the various systems

FIELD(S) OF SCIENCE	PROJECT TITLE	TACC RESOURCES
Fluid, Particulate, and Hyrdaulic Systems	Direct Numerical Simulation of Turbulent Boundary Layer Control and Rarefied Atmospheric Flows	Stampede, Lonestar5
Earth Sciences (EAR)	Static and Time Variable Gravity Determination from GRACE Satellite Data	Lonestar5, Stampede
Ocean Sciences (OCE)	SATELLITE DATA ASSIMILATION IN THE GFDL OCEAN GENERAL CIRCULATION MODEL	Lonestar5
Advanced Scientific Computing (ASC)	Scientific Computing Support	Hikari, Stampede, Lonestar5, Wrangler, Maverick
Condensed Matter Physics	High-Throughput Ab Initio-Based Methods in Nanoscale Binary Alloys	Stampede, Lonestar5
Geophysics	pwmodel	Lonestar5, Stampede
COMPUTER AND INFORMATION SCIENCE AND ENGINEERING (CISE)	Generation of BioPhysical Properties of Molecular and Macromolecular Complexes	Maverick, Lonestar5, Stampede
Physics (PHY)	Plasma Dynamics and Confinement	Lonestar5, Maverick, Ranch, Stampede
Computational Mathematics	Deterministic Boltzmann Poisson Solvers for the modeling on nano-devices	Stampede
Atmospheric Chemistry	meteorological/photochemical simulation of central/eastern texas	Lonestar5, Corral, Stampede
Theoretical Physics	Electronic Properties of Two-Dimensional Crystals	Lonestar5, Stampede
Extragalactic Astronomy and Cosmology	Measuring Dark Halo Profiles in Galaxies	Stampede, Corral
Physical Chemistry	Molecular dynamics simulations of protein unfolding and other mechanical processes involving biopolymers	Lonestar5, Stampede
Visualization, Graphics and Image Processing	Vis Management	Lonestar5, Stampede, Corral, Ranch, Maverick, VisLab, Wrangler
BIOLOGICAL SCIENCES (BIO)	Protein identification via tandem mass spectrometry	Ranch, Lonestar5, Stampede
Atmospheric Sciences (ATM)	Long-Term Ensemble Climate Modeling	Lonestar5, Ranch, Stampede
Computer and Computation Research (CCR)	Identifying simulation points for SPEC CPU 2006 benchmarks	Rustler, Maverick, Stampede, Lonestar5
Materials Research (DMR)	A multiscale simulation approach for the properties of polymer nanocomposites	Lonestar5, Ranch, Stampede
Chemical Reaction Processes	Design of Materials for Energy Conversion and Storage	Stampede, Lonestar5

TACC Quick Start Guide

Interfacial, Transport, and Separation Processes	Understanding the Rheology of Gels from their Potential Energy Landscapes	Lonestar5, Stampede
Performance Evaluation and Benchmarking	TACC Benchmark Account	Stampede, Wrangler
Information, Robotics and Intelligent Systems (IRI)	Learning Agents in Dynamic, Collaborative, and Adversarial Multiagent Environments	Maverick
Computer Systems Architecture	Research in high performance microarchitecture	Lonestar5, Stampede
Nuclear Physics	Simulation for Heavy Ion Collision Experiment	Stampede, Corral
Cross-Disciplinary Activities (CDA)	Temporal and Spatial Distribution of Injectable Particle Introduced in Blood Flow	Lonestar5, Maverick, Stampede, Ranch
Software Engineering	PARSE: Parallel Symbolic Execution	Stampede
Biochemistry and Molecular Structure and Function	Investigation of molecular interactions between Skp2 and novel inhibitors	Lonestar5
Astronomical Sciences (AST)	Supernova Asymmetries	Maverick, Stampede
Biophysics	Elber stampede project	Lonestar5, Stampede
Physical Oceanography	Dynamics of Internal Ocean-type Waves	Stampede
Molecular and Cellular Biosciences (MCB)	Energetics of the Mechano-Chemical Coupling in DNA-Protein and DNA-Nanoparticle Complexes	Ranch, Stampede
Chemical and Thermal Systems (CTS)	High-Resolution Computations for Studies of Scaling in Turbulence and Turbulent Mixing	Maverick, Stampede, Ranch
Mathematical Sciences (DMS)	Hurricane Storm Surge Simulations on Petascale Computers	Stampede, Ranch
Electrical and Communications Systems (ECS)	Accelerating Nano-Scale Transistor Innovation Through Petascale Simulation	Stampede, Ranch
Chemistry (CHE)	Salt Effects in Peptides and Nucleic Acids	Stampede, Ranch
Solid State Physics	Quantum Transport in Transition Metal Dichalcogenides	Stampede
Software Development	Ranger Technology Insertion	Wrangler
Analytical and Surface Chemistry	Computer Simulation of Atomistic Chemical Dynamics	Lonestar5, Stampede
MATHEMATICAL AND PHYSICAL SCIENCES (MPS)	Electronic Structure Calculations on Metal Clusters Using Spin-Orbit Configuration Interaction	Lonestar5, Stampede
Numeric and Symbolic Computation	An On-demand Test Problem Server	Lovett
Metals, Ceramics, and Electronic Materials	ab initio molecular dynamics simulation of structure and dynamics in aluminosilicate glasses	Stampede
Seismology	3-D seismic simulation, imaging and inversion	Lonestar5, Stampede, Ranch

TACC Quick Start Guide

GEOSCIENCES (GEO)	Systematic Development of a Subgrid Scaling Framework to Improve Land Simulation	Stampede
Stellar Astronomy and Astrophysics	Progress Towards a Comprehensive Theory of Star Formation- From Brown Dwarfs to High Mass Stars and On To Giant Molecular Clouds	Maverick, Ranch, Stampede
Gravitational Physics	Studies In Theoretical Astrophysics and General Relativity	Stampede, Ranch
Solar Terrestrial Research	Center for Integrated Space Weather Modeling	Stampede, Wrangler, Ranch
Genetics and Nucleic Acids	Analysis of next-generation sequencing data	Corral, Lonestar5, Stampede