Need Help?
We've got you Covered

You don't have to be an experienced programmer or developer to use the HPC. It runs on Linux, but in most cases, users can learn everything they need to know in a short tutorial with our support team. We offer customized training and orientations to meet the unique needs of each user.

How to get Started

- Request an account (freely available to FDA staff or contractors in any Center)
- Attend the new user orientation to learn best practices and parallelization techniques
- Set up your program in a 1 on 1 consultation, if needed

HPC Environments at FDA

- The HPC has the Authority to Operate (ATO) at a moderate-level security impact.
- It is available for both regulatory and research projects.
  - ⇒ CFSAN RAVEN
  - ⇒ CDRH Betsy/Bluefin
  - ⇒ CBER HIVE
  - ⇒ AWS Cloud HPC

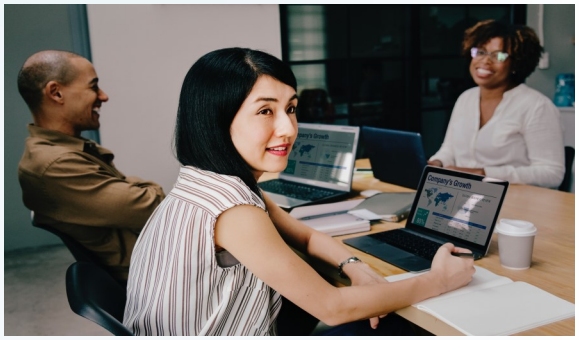

Visit our Webpage

For more information or to request services on the available HPC at the FDA, please visit the FDA HPC website:

http://inside.fda.gov:9003/ProgramsInitiatives/ScientificComputing/HighPerformanceComputing/

High Performance Computing at FDA

A service of the Office of Science and Engineering Laboratories

Powering Innovation in FDA Regulatory Science

About us

The High Performance Computing services at FDA support research and regulatory activities undertaken by scientists in all FDA centers. It has been in operation since 2011.

Features

- Parallel computing clusters containing thousands of CPUs
- A high-speed internal network for lightning fast processing on the most complex jobs
- Storage capacity for massive datasets
- Access via FDA's administrative or scientific network
- Reductions in runtime on the order of 55 months down to a single day are not unusual
- Training and specialized consultation in best practices for HPC use
- Expert support in migration of scientific applications to an HPC environment
- Hosts numerous high-quality open-source scientific tools and applications
- Cloud HPC ecosystem allows for cross-center data sharing for regulatory science

CBER HIVE

The CBER High-performance Integrated Virtual Environment (HIVE) is a specialized platform developed in CBER, containing a sequence read archive linked directly to the CBER HPC. This environment provides web access for authorized users to deposit, retrieve, annotate and compute on next-generation sequencing (NGS) data, and to analyze the outcomes using web-interface visual environments built in collaboration with research scientists and regulatory personnel. HIVE has additional applications for analysis of post market, adverse events, and metagenomic data.

CDRH Betsy/Bluefin

The CDRH HPC environment consists of two semi-autonomous computational clusters. The Betsy cluster has over 3000 processing cores in 350+ compute nodes. The Bluefin cluster, which is housed in a mobile container, provides 920 additional cores.

Applications include large-scale modeling and simulations, genomic analysis, computational physics, molecular and fluid dynamics, Bayesian analysis, semantic data mining, and many others.

CFSAN RAVEN

The CFSAN HPC environment consists of the RAVEN Cluster containing 128 nodes, 2,560 processing cores, and 2.0 PB data storage. Major software applications include GIMS, CLC Bio, and other software for processing of genomic data from sequencers of Pacific Biosciences, Illumina and Oxford Nanopore.

Moving FDA Regulatory Science Forward

Examples of supporting FDA research:

Researchers in CDRH/OSEL used the HPC to create a whole-heart simulation of ventricular fibrillation. They developed a high-resolution anatomically detailed model of 21 million elements.

Using DOCK6, researchers with NCTR's Interactome Project analyzed 31 million drug-protein interactions to create DRAR-CPI, a tool that can help predict adverse drug reactions and identify drug repositioning opportunities.

Supporting the Genome in a Bottle Consortium Collaboration, the HPC was instrumental in developing whole human genome reference material, now available through NIST.

CFSAN researchers are utilizing both Raven and cloud HPC model with GenomeTrakr to complete food-borne pathogen genome analysis in 60 minutes.