

Launch Your HPC Workflows With HPC E4S

Adaptive Computing and ParaTools partner to supply uniform multi-cloud provider access to launch your HPC workflows with the Extreme-scale Scientific Software Stack (E4S) via Adaptive Computing's HPC Cloud On-Demand Data Center (ODDC) web interface using custom images. Commercial images are available on all cloud platforms via the ODDC web interface.

E4S

The Extreme-scale Scientific Software Stack (E4S) is a community effort to provide open source software packages for developing, deploying and running scientific applications on high-performance computing (HPC) platforms. E4S provides for both source builds and container builds of a broad collection of HPC and AI/ML software packages targeting GPUs from NVIDIA, AMD, and Intel.



- Extreme-scale Scientific Software Stack (E4S)
- E4s with ODDC provides for the 3 bullets below:
- Commercial and Open Source HPC Applications
- Custom HPC Applications
- GPU Support for AI / ML Workloads

Department of Energy's Extreme Scale

E4S is a curated, [Spack](#) based collection of HPC packages including TAU, PETSc, Trilinos, MPICH, OpenMPI, HDF5, NetCDF, WarpX as well as AI/ML packages including TensorFlow and PyTorch.

- E4S has 100+ HPC and AI/ML tools.
- E4S supports cloud hardware resources (GPUs, network cards) with integration with workflow management systems (MOAB/SLURM/Torque).
- Multiple MPI implementations are available through E4S.
- Cost varies between cloud providers and a single interface to multiple cloud providers with specialized hardware (GPUs, network adapters) supported by E4S reduces cost.
- Access SLURM and Moab via E4S.
- Use your own resources and launch workflows and E4S through the adaptive ODDC interface.

The HPC Cloud On-Demand Data Center (ODDC) comes from the creators of Moab/TORQUE to provide seamless integration with multiple Cloud Service Provider (CSP) infrastructure resources.

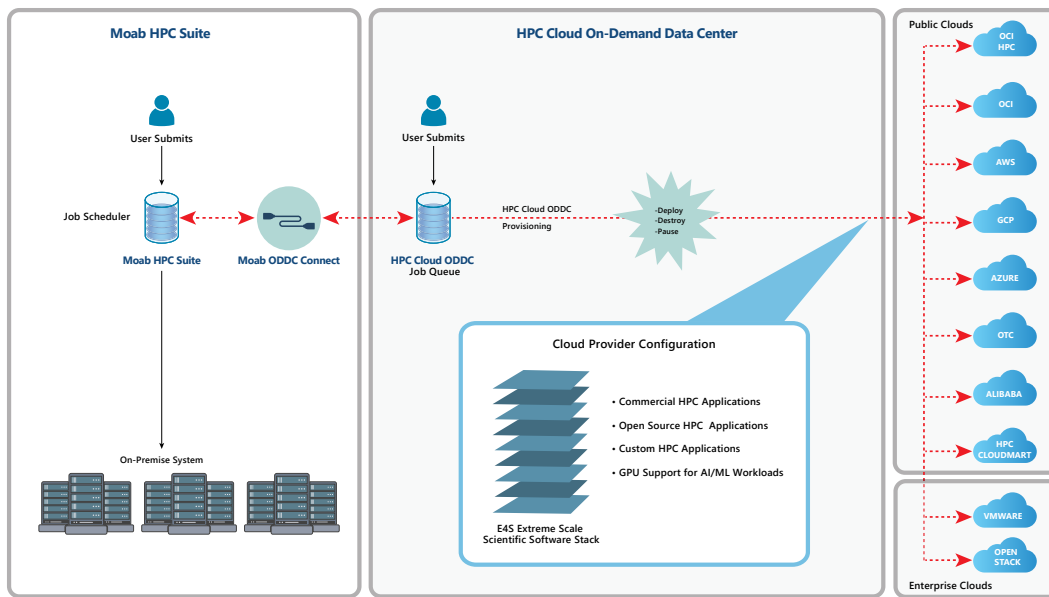
- With ODDC and E4S you have the power to make custom images that are optimized to the application and derived from commercial images on major CSPs (e.g. AWS).
- Multi-cloud provider access through the ODDC web interface.
- Secure access to integrate on-premise and remote cloud resources.
- MOAB/SLURM/TORQUE scheduler integration with MPI environments.
- Transparently access cloud resources from on-premise resources. Seamless integration between on-premise infrastructure and the commercial cloud infrastructure.
- Support for GPUs and AI/ML applications.

Launch your HPC Workflows with E4S via Adaptive Computing's HPC Cloud On-Demand Data Center (ODDC) Web Interface

Adaptive's HPC Cloud On-Demand Data Center is a powerful cloud management solution that does not require cloud expertise to use.

- The HPC Cloud ODDC is a cloud enablement solution that provides an easy way to spin up temporary or persistent HPC cloud infrastructure resources quickly, inexpensively, and on demand. ODDC automates HPC cloud infrastructure deployment, monitoring, scaling, and shutdown of cloud resources, assuring that you only pay for cloud resources while in use.
- Hardware refreshes become unnecessary due to the availability of on-demand cloud nodes when required.
- Significant savings of cloud usage costs.
- ODDC reduces operation costs by automatic termination of expensive cloud resources when the job completes. (GPU nodes) Don't even think of using GPU nodes unless you are using ODDC, because ODDC takes the GPU nodes down when the job completes.
- Seamless integration with multi-node cloud executions.
- Secure access to cloud resources - ODDC requires you to utilize keys that are 4096-bit class security.
- Security is available through ODDC via CSP security levels (GovCloud level secure). ODDC falls within the CSP security compliance.

Adaptive Computing's HPC Cloud On-Demand Data Center is the most efficient way to use E4S



HPC Cloud ODDC technical approach to integrating the customer's on-premise assets with those within the Cloud Service Providers

About ParaTools

ParaTools improves the performance and productivity of the world's fastest supercomputers, and provides consulting expertise in high performance computing to clients around the world. ParaTools, Inc. was founded by Dr. Allen Malony and Dr. Sameer Shende in 2004. From a single office in the Riverfront Research Park at University of Oregon, ParaTools has grown to a multi-national business venture with offices in France and across the United States. ParaTools is devoted to the development and application of tools for parallel computing systems, with specific expertise in high-performance computing and performance analysis.

About Adaptive Computing

Adaptive Computing is a global software company headquartered in Naples, Florida, USA that has provided advanced applications and tools to the world's largest High-Performance Computing installations for over 2 decades. We have a rock-solid industry reputation in HPC Workload Scheduling and Cloud Management Solutions and continue to solve IT management challenges in scale-intensive and complex environments worldwide. Adaptive Computing works with large commercial enterprises, government labs, and academic institutions. Some of the world's largest clusters, grids, and data centers use Adaptive's Moab HPC Suite and HPC Cloud On-Demand data Center to maximize performance and value, simplify management, and create a competitive advantage. Our mission is to help organizations to enhance performance, improve efficiency and reduce costs.