Lenovo Intelligent Computing Orchestration (LiCO)
Product Guide

Lenovo Intelligent Computing Orchestration (LiCO) is a software solution that simplifies the use of clustered computing resources for Artificial Intelligence (AI) model development and training. LiCO interfaces with an open source software orchestration stack, enabling the convergence of AI onto an HPC or Kubernetes-based cluster.

The unified platform simplifies interaction with the underlying compute resources, enabling customers to take advantage of popular open source cluster tools while reducing the effort and complexity of using it for AI.

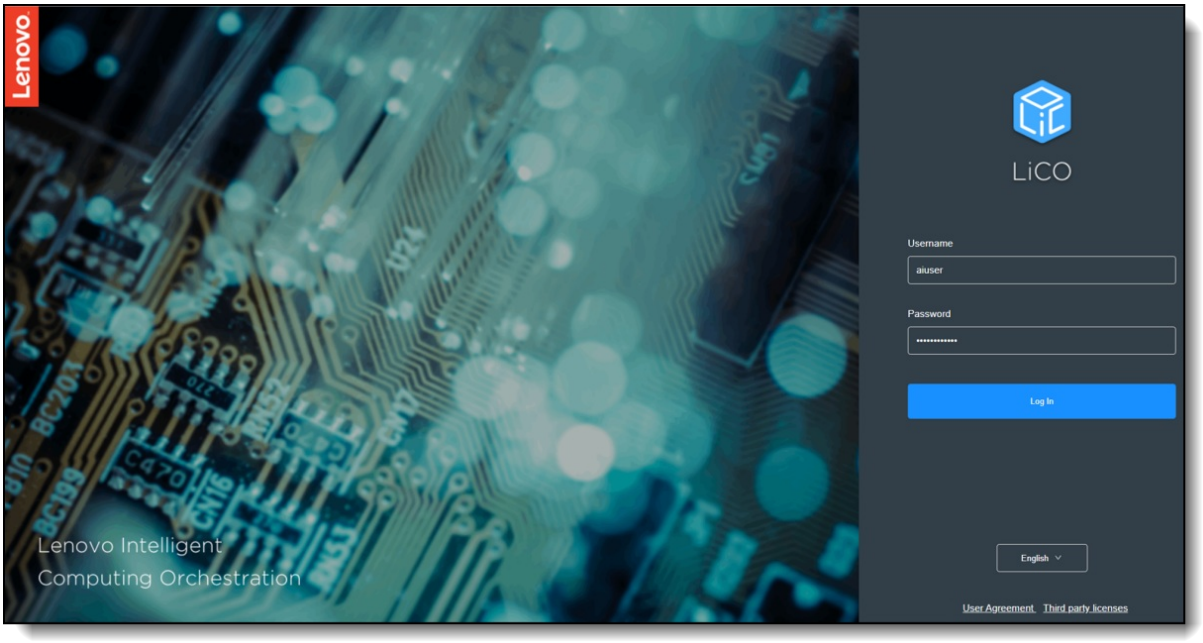

Did You Know?

LiCO enables a single cluster to be used for multiple AI workloads simultaneously, with multiple users accessing the available cluster resources at the same time. Running more workloads can increase utilization of cluster resources, driving more user productivity and value from the environment.

What's new in LiCO 6.2

Lenovo recently announced LiCO Version 6.2, improving the functionality for both AI users, HPC users, and HPC administrators of LiCO, including:

- Support for new ThinkSystem servers (SR670 V2, SR650 V2, SR630 V2, SD650 V2, SD650-N V2)
- Lenovo Accelerated AI for Text Classification
- Trained model packaging into a docker container image
- Intel OneAPI tools and templates (HPC/AI version)
- Cluster View for more detailed resource monitoring (HPC/AI version)

Part numbers

The following table lists the ordering information for LiCO.

Table 1. LiCO HPC/AI version ordering information

Description	LFO	Software CTO	Feature code
Lenovo HPC AI LiCO Software 90 Day Evaluation License	7S090004WW	7S09CTO2WW	B1YC
Lenovo HPC AI LiCO Software w/1 yr S&S	7S090001WW	7S09CTO1WW	B1Y9
Lenovo HPC AI LiCO Software w/3 yr S&S	7S090002WW	7S09CTO1WW	B1YA
Lenovo HPC AI LiCO Software w/5 yr S&S	7S090003WW	7S09CTO1WW	B1YB

Table 2. LiCO K8S/AI ordering information (Kubernetes)

Description	LFO	Software CTO	Feature code
Lenovo K8S AI LiCO Software Evaluation License (90 days)	7S090006WW	7S09CTO3WW	S21M
Lenovo K8S AI LiCO Software 4GPU w/1Yr S&S	7S090007WW	7S09CTO4WW	S21N
Lenovo K8S AI LiCO Software 4GPU w/3Yr S&S	7S090008WW	7S09CTO4WW	S21P
Lenovo K8S AI LiCO Software 4GPU w/5Yr S&S	7S090009WW	7S09CTO4WW	S21Q
Lenovo K8S AI LiCO Software 16GPU upgrade w/1Yr S&S	7S09000AWW	7S09CTO4WW	S21R
Lenovo K8S AI LiCO Software 16GPU upgrade w/3Yr S&S	7S09000BWW	7S09CTO4WW	S21S
Lenovo K8S AI LiCO Software 16GPU upgrade w/5Yr S&S	7S09000CWW	7S09CTO4WW	S21T
Lenovo K8S AI LiCO Software 64GPU upgrade w/1Yr S&S	7S09000DWW	7S09CTO4WW	S21U
Lenovo K8S AI LiCO Software 64GPU upgrade w/3Yr S&S	7S09000EWW	7S09CTO4WW	S21V
Lenovo K8S AI LiCO Software 64GPU upgrade w/5Yr S&S	7S09000FWW	7S09CTO4WW	S21W

Features for LiCO users

Note: There are two distinct versions of LiCO, LiCO HPC/AI (Host) and LiCO K8S/AI, to allow clients a choice for the which underlying orchestration stack is used, particularly when converging AI workloads onto an existing cluster. The user functionality is common across both versions, with minor environmental differences associated with the underlying orchestration being used.

A summary of the differences for user access is as follows:

LiCO K8S/AI version:

- AI framework containers are docker-based and managed outside LiCO in the customer's docker repository
- Custom job submission templates are defined with YAML
- Does not include HPC standard job submission templates

LiCO HPC/AI version:

- AI framework containers are Singularity-based and managed inside the LiCO interface
- Custom job submission templates are defined as SLURM batch scripts
- Includes HPC standard job submission templates

LiCO provides users the following benefits:

- A web-based portal to deploy, monitor and manage AI development and training jobs on a distributed cluster
- Container-based deployment of supported AI frameworks for easy software stack configuration
- Direct browser access to Jupyter notebook instances running on the cluster
- Standard and customized job templates to provide an intuitive starting point for less experienced users
- Lenovo Accelerated AI pre-defined training and inference templates for many common AI use cases
- Lenovo AI Studio end-to-end workflow for Image Classification, Object Detection, Instance Segmentation
- Workflow to define multiple job submissions as an automated workflow to deploy in a single action
- TensorBoard visualization tools integrated into the interface (TensorFlow-based)
- Management of private space on shared storage through the GUI
- Monitoring of job progress and log access

Those designated as LiCO users have access to dashboards related primarily to AI development and training tasks. Users can submit jobs to the cluster, and monitor their results through the dashboards. The following menus are available to users:

- **Home menu for users** – provides an overview of the resources available in the cluster. Jobs and job status are also given, indicating the runtime for the current job, and the order of jobs deployed. Users may click on jobs to access the associated logs and job files. The figure below displays the home menu.

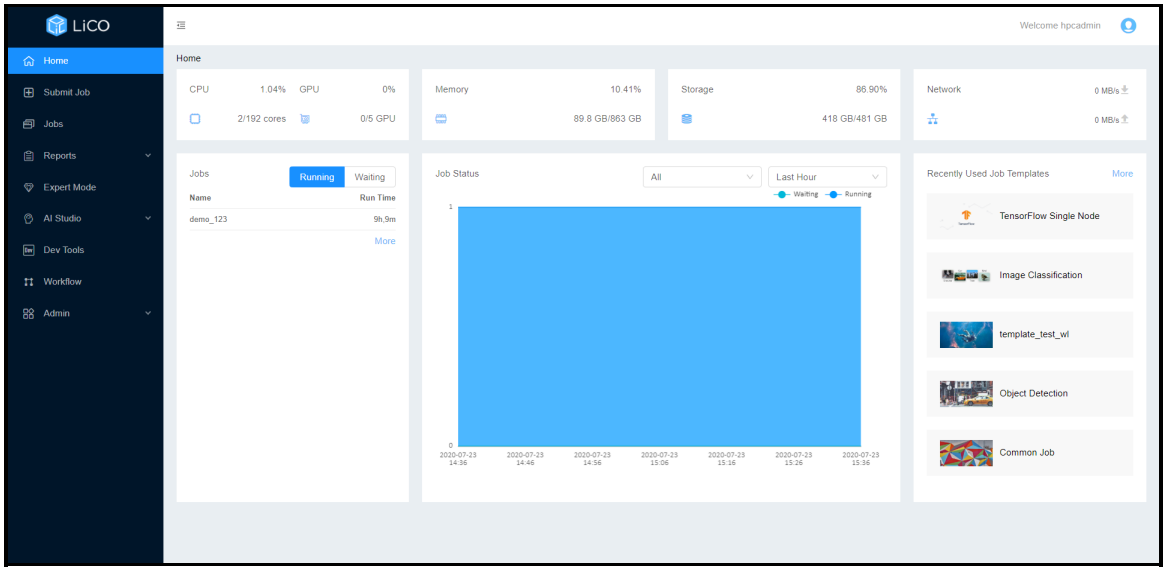

Figure 1. User Home Menu

- **Submit job menu** – allows users to set up a job and submit it to the cluster. The user first picks a job template. After selecting the template, the user gives the job a name and inputs the relevant parameters, chooses the resources to be requested on the cluster and submits it.
Users can take advantage of Lenovo Accelerated AI templates, industry-standard AI templates, submit generic jobs via the Common Job template, as well as create their own templates requesting specified parameters.

  The figure below displays a job template for training with TensorFlow on a single node.

Figure 2. AI Job Template

LiCO also provides TensorBoard monitoring when running certain TensorFlow workloads, as shown in the following figure.

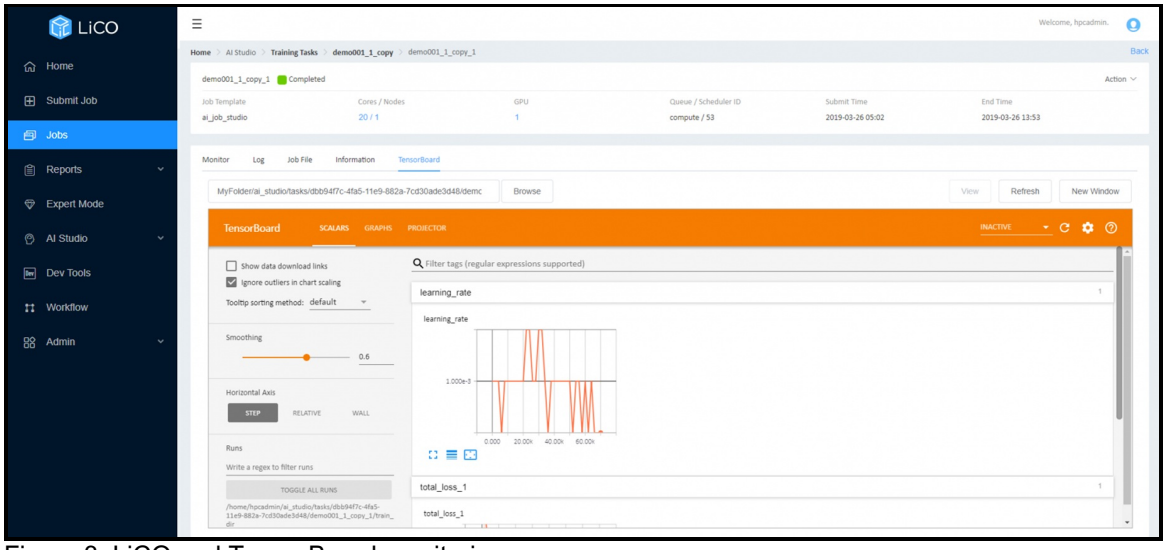

Figure 3. LiCO and TensorBoard monitoring

- **Jobs menu** – displays a dashboard listing jobs and their statuses. In addition, you can select the job and see results and logs pertaining to the job in progress (or after completion). Tags and comments can be added to completed jobs for easier filtering.
- **AI Studio menu** – provides users the ability to label data, optimize hyperparameters, as well as test and publish trained models from within an end-to-end workflow in LiCO. AI Studio supports Image Classification, Object Detection, and Instance Segmentation workflows.
- **Dev Tools menu** – enables users to create, run and view Jupyter notebook instances on the cluster from LiCO for model experimentation and development.
- Workflow menu – allows users to create multi-step jobs that execute as a single action. Workflows can contain serially-executed steps as well as multiple jobs to execute in parallel within a step to take full advantage of cluster resources.
- **Admin menu** – allows users to access a number of capabilities not directly associated with deploying workloads to the cluster, including access to shared storage space on the cluster through a drag-and-drop interface and access to provision API and git interfaces for integration of AI Studio steps into a DevOps environment.

Lenovo Accelerated AI

Lenovo Accelerated AI provides a set of templates that aim to make AI training and inference simpler, more accessible, and faster to implement. The Accelerated AI templates differ from the other templates in LiCO in that they do not require the user to input a program; rather, they simply require a workspace (with associated directories) and a labelled dataset.

The following use cases are supported with Lenovo Accelerated AI templates:

- Image Classification
- Object Detection
- Instance Segmentation
- Medical Image Segmentation
- Seq2Seq
- Memory Network
- Image GAN
- Text Classification

The following figure displays the Lenovo Accelerated AI templates.

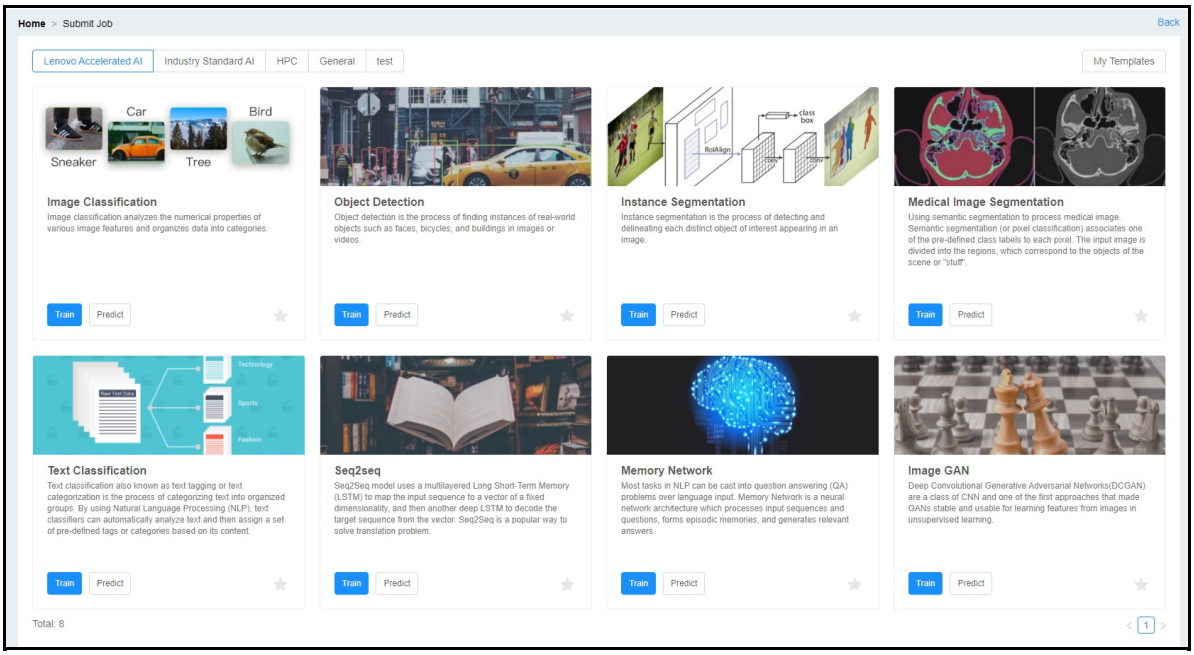

Figure 4. Lenovo Accelerated AI templates

Each Lenovo Accelerated AI use-case is supported by both a training and inference template. The training templates provide parameter inputs such as batch size and learning rate. These parameter fields are pre-populated with default values, but are tunable by those with data science knowledge. The templates also provide visual analytics with TensorBoard; the TensorBoard graphs continually update in-flight as the job runs, and the final statistics are available after the job has completed.

The following figure displays the embedded TensorBoard interface for a job. TensorBoard provides visualizations for TensorFlow jobs running in LiCO, whether through Lenovo Accelerated AI templates or the standard TensorFlow AI templates.

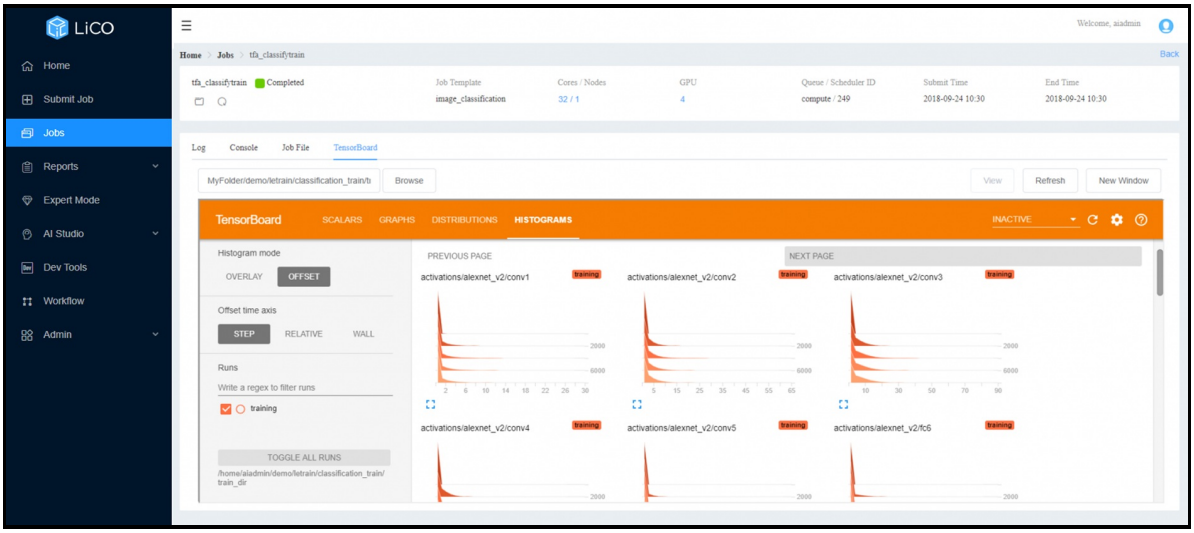

Figure 5: TensorBoard in LiCO

LiCO also provides inference templates which allow users to predict with new data based on models that have been trained with Lenovo Accelerated AI templates. For the inference templates, users only need to provide a workspace, an input directory (the location of the data on which inference will be performed), an output directory, and the location of the trained model. The job will run, and upon completion, the output directory will contain the analyzed data. For visual templates such as Object Detection, images can be previewed directly from within LiCO's Manage Files interface.

The following two figures display an input file to the Object Detection inference template, as well as the corresponding output.

Figure 6: JPG file containing image of cat for input into inference job

Figure 7: LiCO output displaying the section of the JPG containing the cat image

Favorites tab

LiCO allows the user to select frequently used job submission templates as “favorites” to simplify user access. Selecting the star in a template box will add the template to the Favorites tab at the top of the Submit Job screen, which is the default view to the Submit Job tab. If no favorites have been selected, the Favorites tab will not appear. Users may add standard templates, Lenovo Accelerated AI templates, and custom-defined templates to this tab.

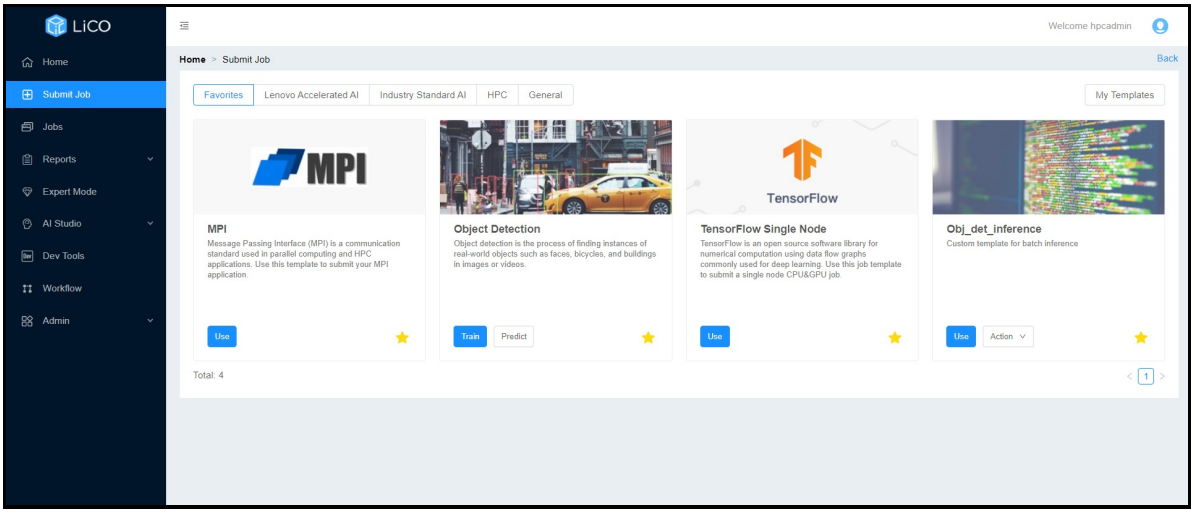

Figure 8. LiCO Favorites tab

AI Studio

LiCO AI Studio provides an end-to-end workflow for Image Classification, Object Detection, and Instance Segmentation, with training based on Lenovo Accelerated AI pre-defined models. A user can import an unprocessed, unlabeled data set of images, label them, train multiple instances with a grid of parameter values, test the output models for validation, and publish to a git repository for use in an application environment. Additionally, users can initiate the steps in AI Studio from a REST API call to take advantage of LiCO as part of a DevOps toolchain.

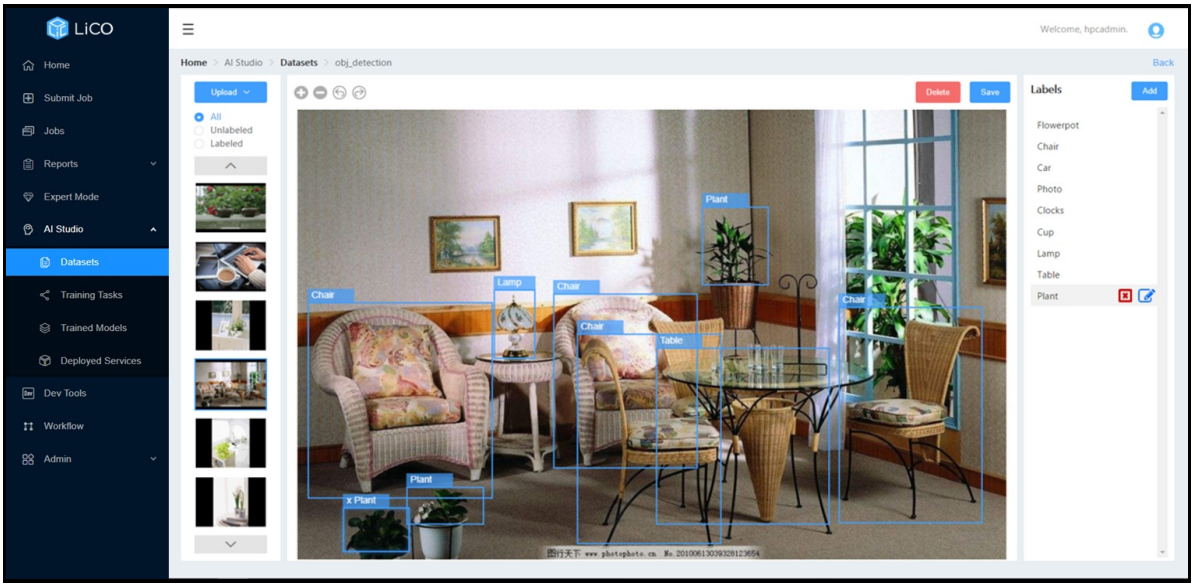

Figure 9. LiCO dataset file with labeled image

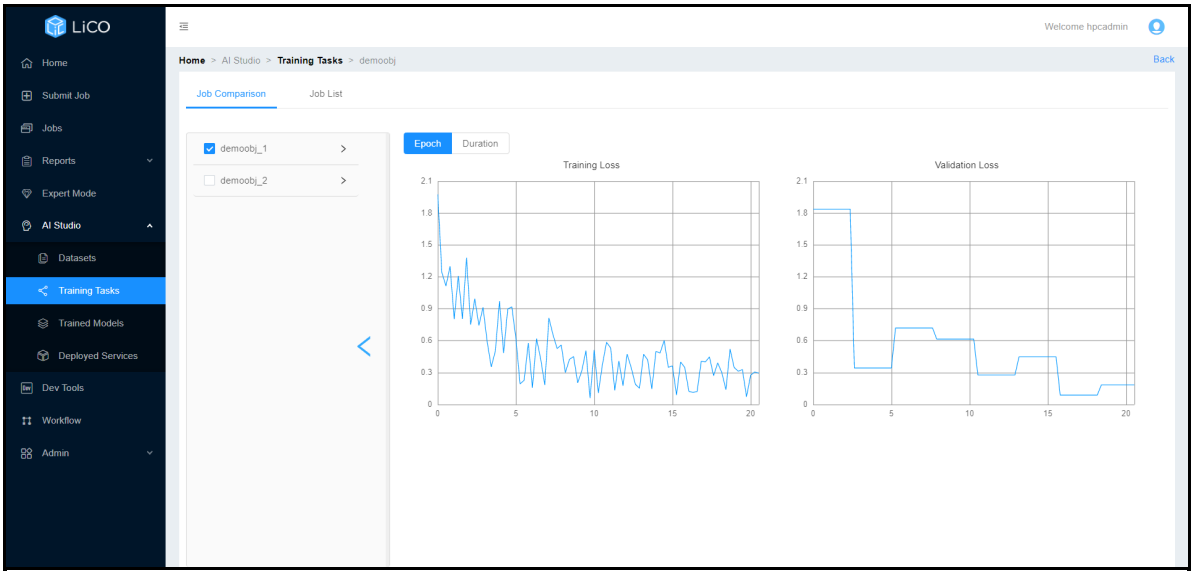

Figure 10. LiCO AI Studio model tuning

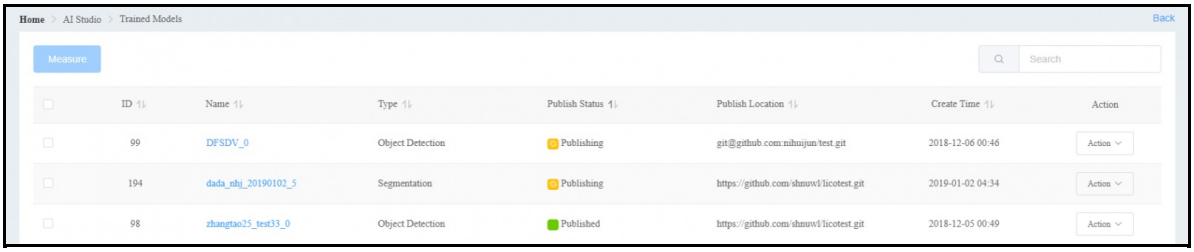

Figure 11. Trained model repository displaying published model location

Dev Tools

LiCO includes the capability to create and deploy instances of Jupyter on the cluster. Users may create multiple instances, to customize for different software environments and projects. At launch of an instance the user can define the amount of compute resource requirements needed (CPU and GPU) to better optimize performance to the task and optimize resource usage on the cluster.

Once a Jupyter instance is created, the user can deploy it to the cluster and use the environment directly from their browser in a new tab. The user can leverage the Jupyter interface directly to upload, download and run code as they normally would, utilizing the shared storage space used for LiCO.

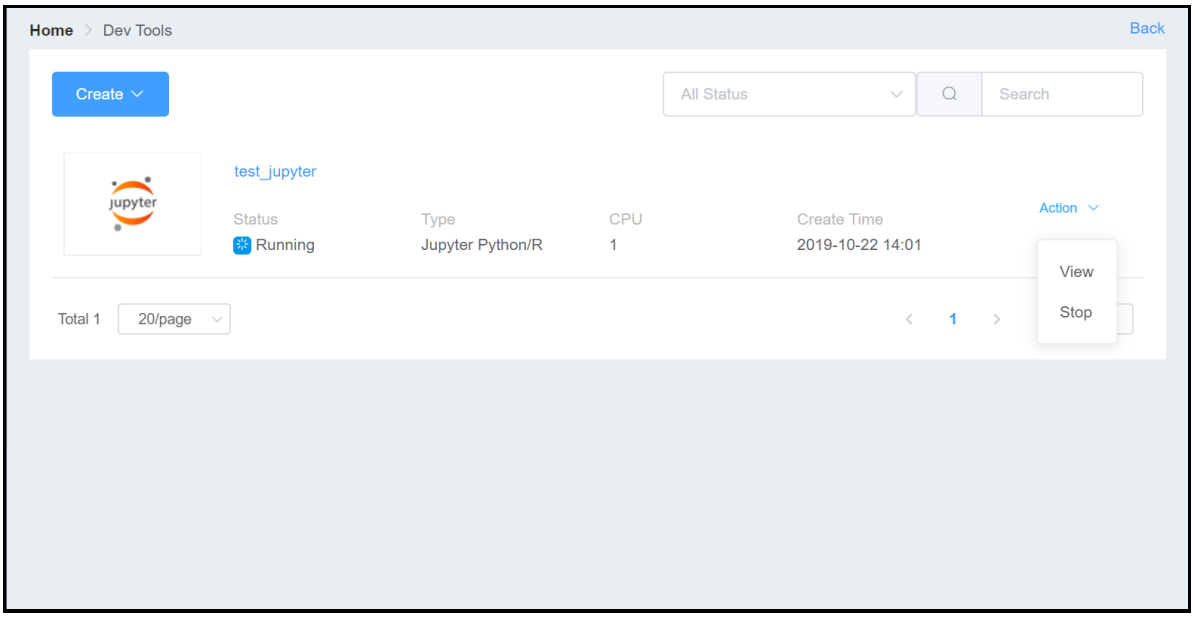

Figure 12. Accessing a running Jupyter instance through the action menu

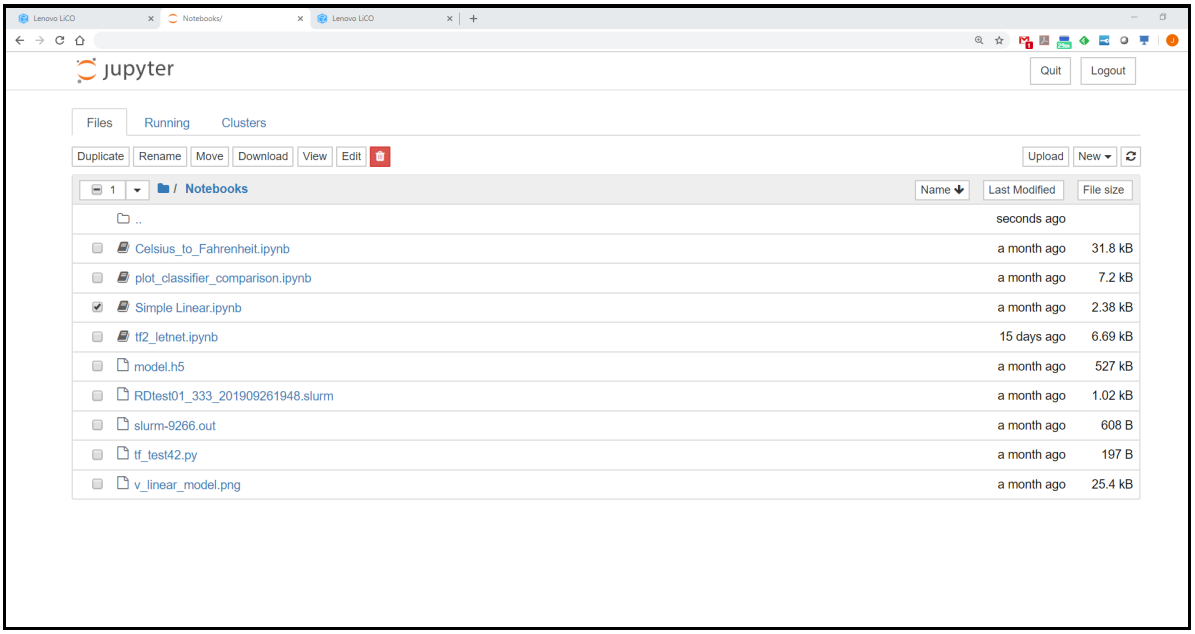

Figure 13. Jupyter instance accessible in new browser tab

Workflow

LiCO provides the ability to define multiple job submissions into a single execution action, called a Workflow. Steps are created to execute job submissions in serial, and within each step multiple job submissions may be executed in parallel. Workflow uses LiCO job submission templates to define the jobs for each step, and any template available including custom templates can be used in a workflow.

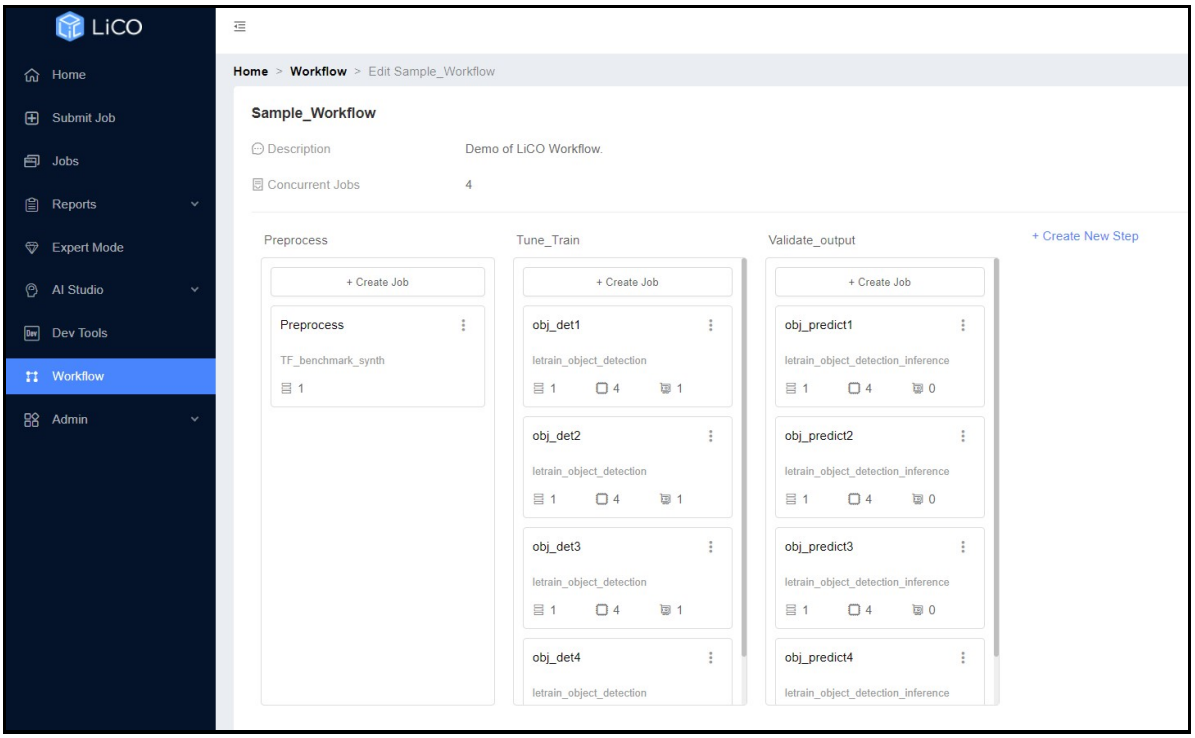

Figure 14. Defining a workflow in LiCO

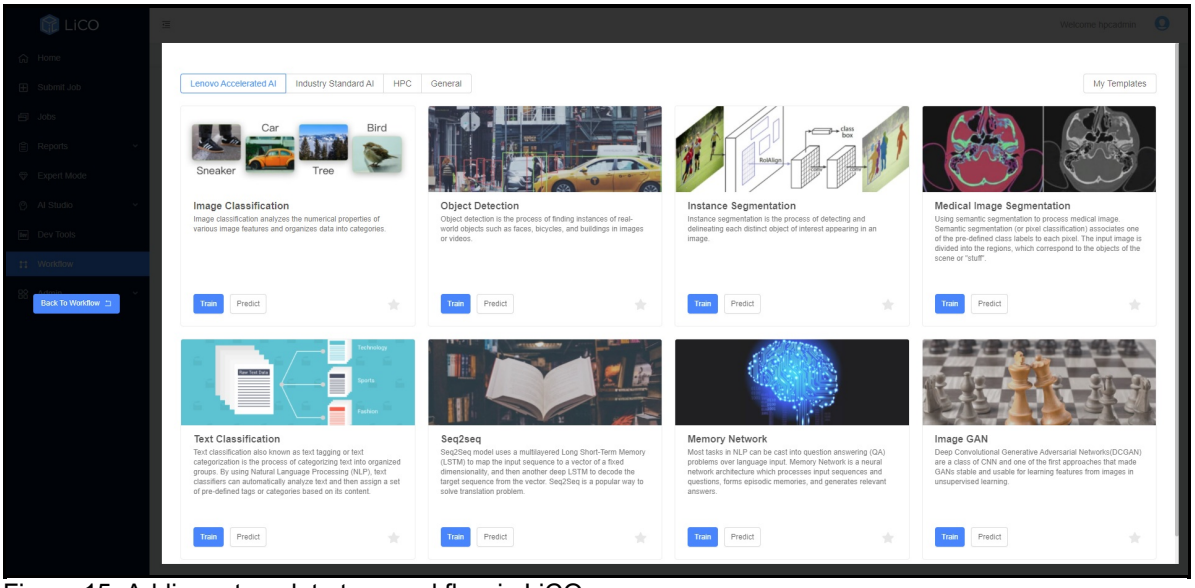

Figure 15. Adding a template to a workflow in LiCO

LiCO workflows allow users to automate the deployment of multiple jobs that may be required for a project, so the user can execute and monitor as a single action. Workflows can be easily copied and edited, allowing users to quickly customize existing workflows for multiple projects.

Admin

The Admin tab for the user provides access to their storage space on the cluster. Through the Manage Files subtab the user can upload, download, cut/copy/paste, preview and edit files on the cluster storage space from within the LiCO portal. The text editor within LiCO allows syntax-aware display and editing based on the file extension. The Admin tab also enables users to publish a trained model to a git repository or as a docker container image.

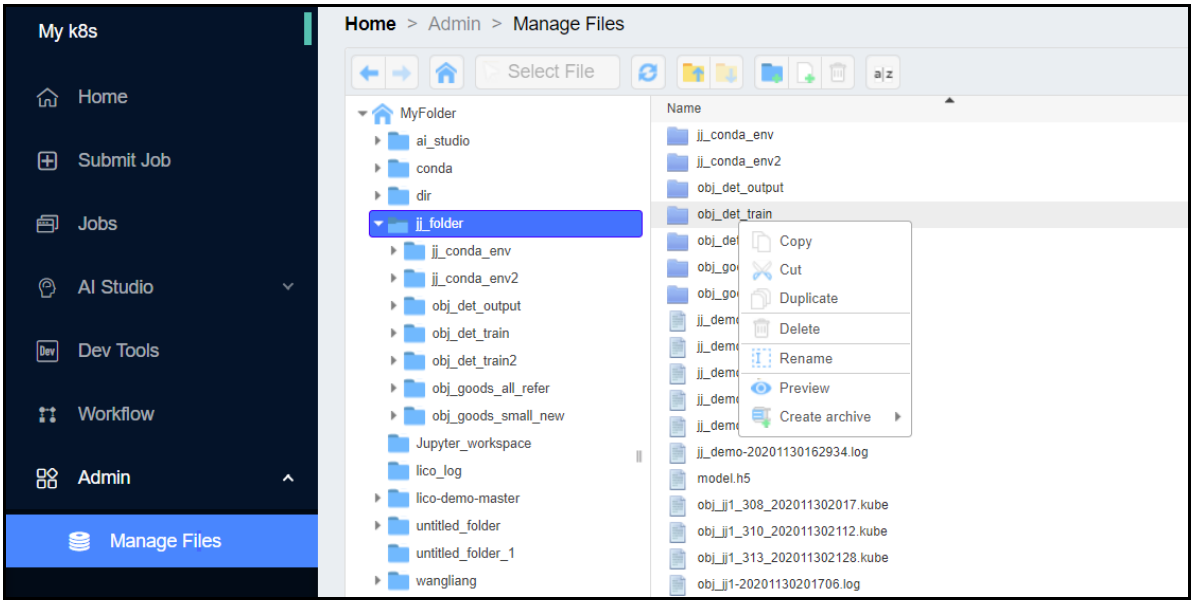

Figure 16. Cluster storage access

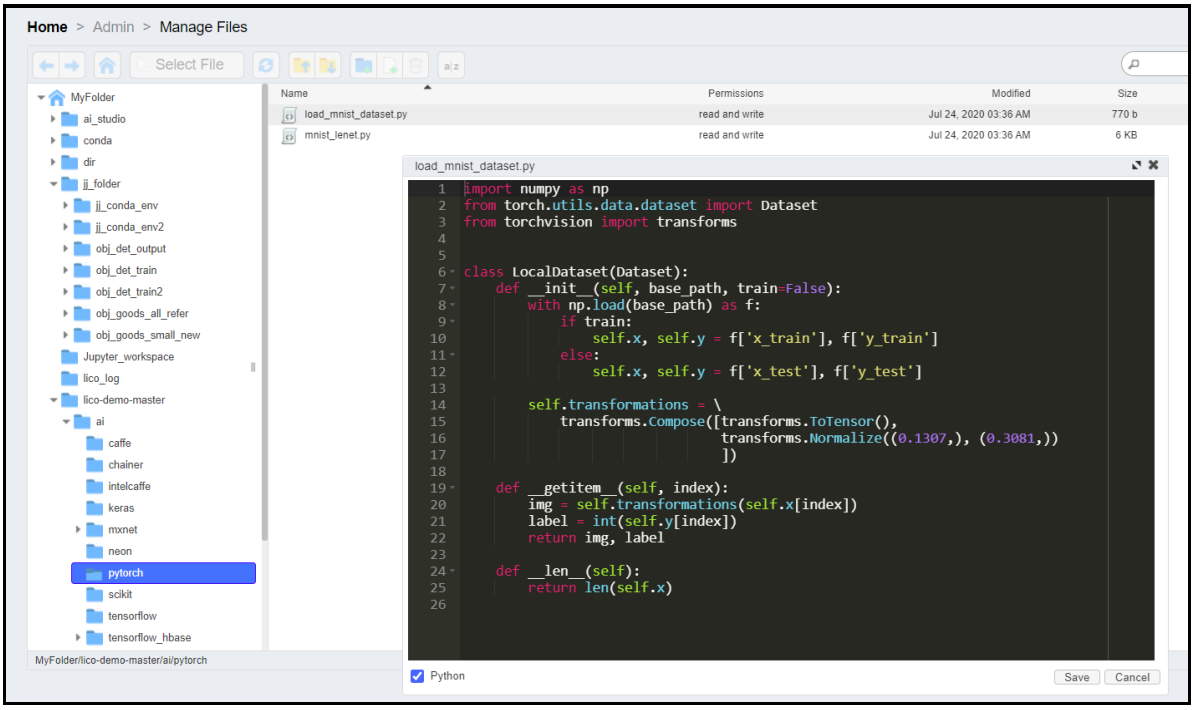

Figure 17. Text file editor

Additional features for LiCO HPC/AI Users

In addition to the user features above, the LiCO HPC/AI version provides additional user capabilities suited to an HPC-based cluster environment.

HPC Runtime Module Management

LiCO HPC/AI version allows the user to pre-define modules and environmental variables to load at the time of job execution through Job submission templates. These user-defined modules eliminate the step of needing to manually load required modules before job submission, further simplifying the process of running HPC workloads on the cluster. Through the Runtime interface, users can choose from the modules available on the system, define their loading order, and specify environmental variables for repeatable, reliable job deployment.

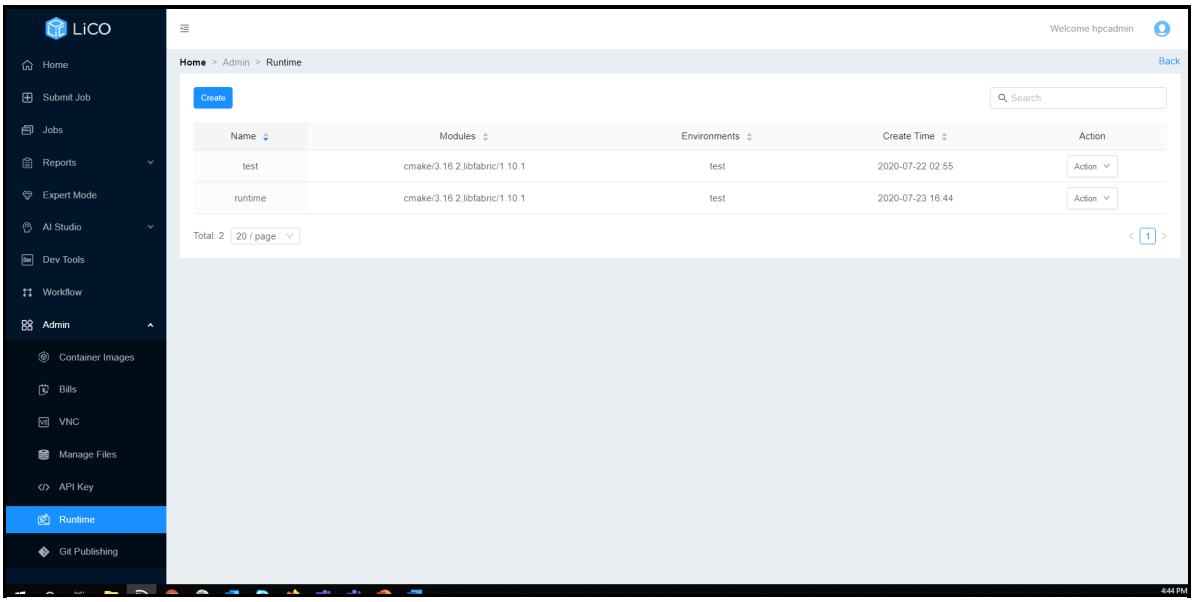

Figure 18. HPC runtime module list

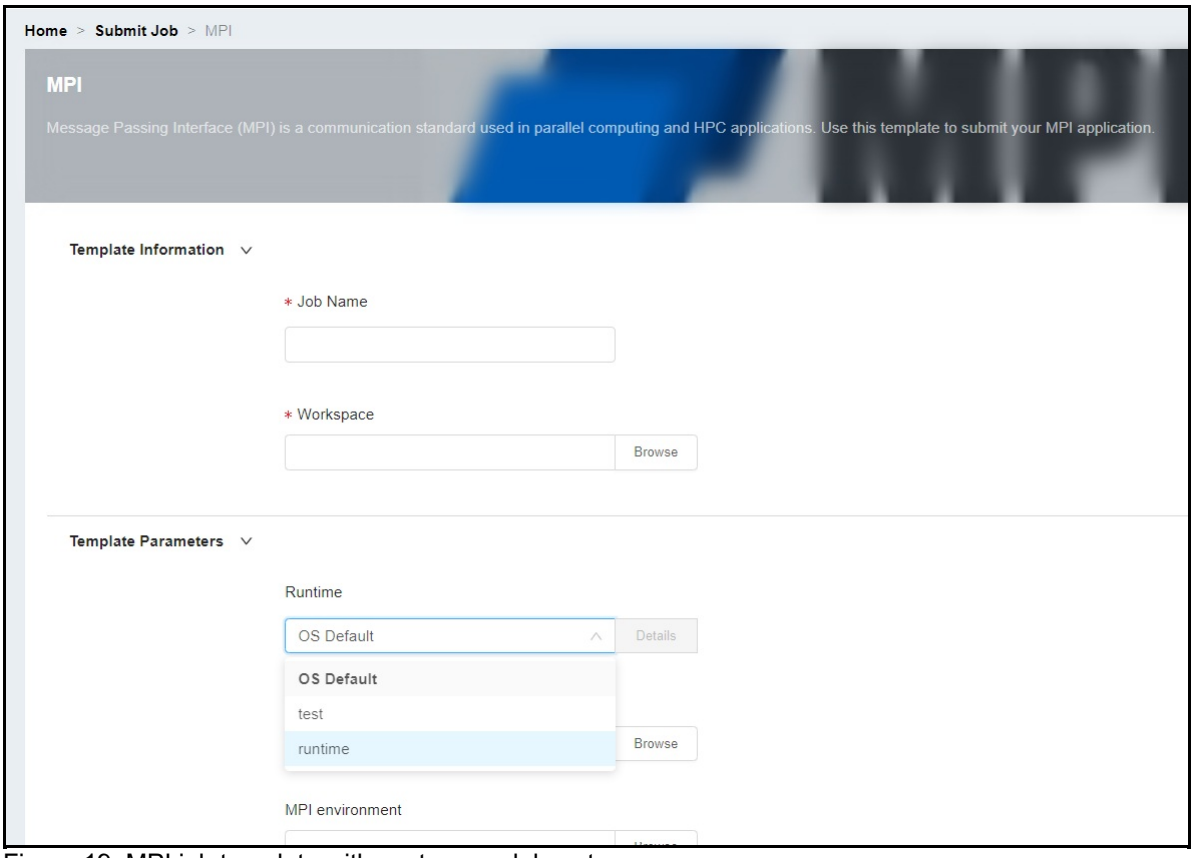

Figure 19. MPI job template with custom module setup

Additional standard templates are provided to support deployment of containerized HPC workloads through Singularity or CharlieCloud. These templates simplify deploying containers for HPC workloads by eliminating the need to create custom runtimes and custom templates for these workloads unless needed for more granularity.

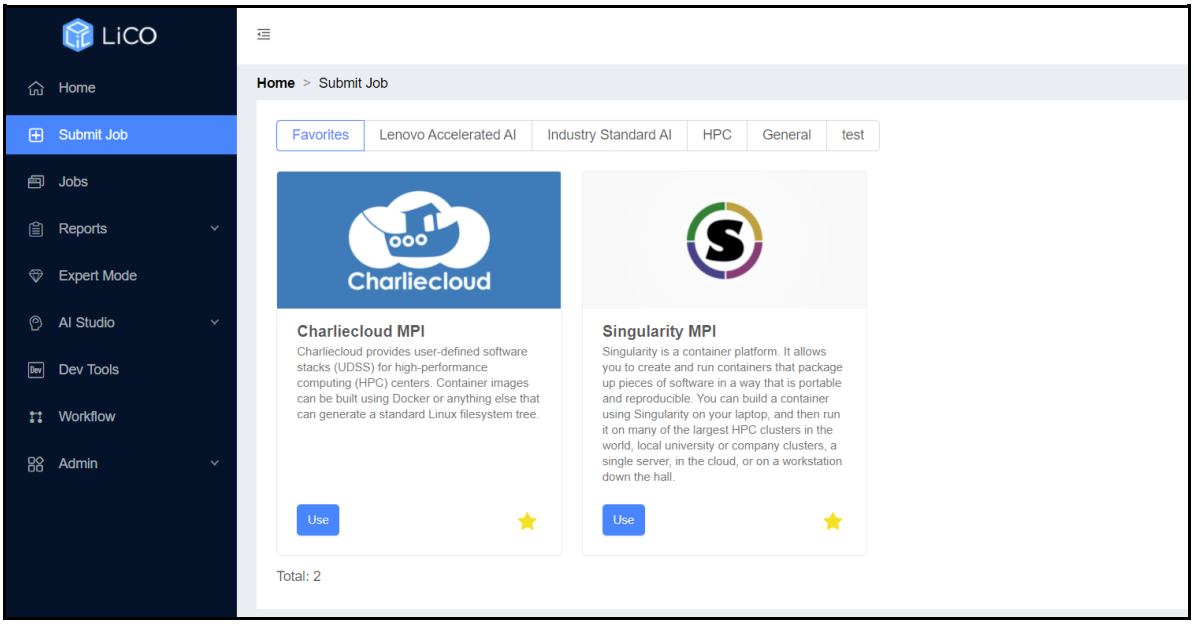

Figure 20. CharlieCloud and Singularity standard job templates

Intel oneAPI

Intel oneAPI is a cross-industry, open, standards-based unified programming model that delivers a common developer experience across accelerator architectures—for faster application performance, more productivity, and greater innovation. LiCO 6.2 provides three new templates based on oneAPI – Intel MPI, Intel MPITune, and Intel OpenMP – that are optimized to run on Intel processors.

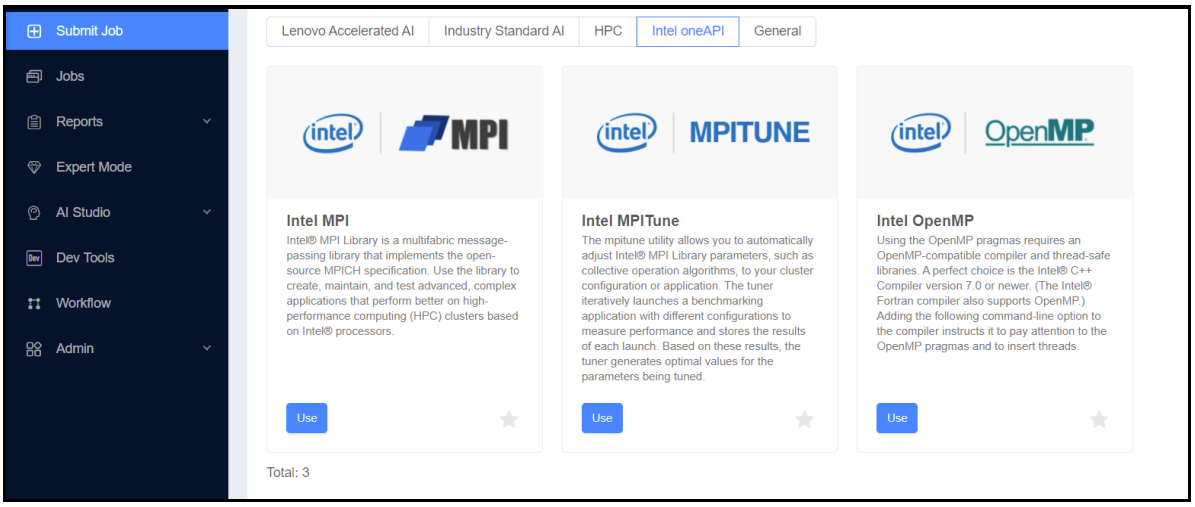

Figure 21. HPC Templates for leveraging Intel oneAPI libraries

Singularity Container Image Management

LiCO HPC/AI version provides both users and administrators with the ability to build, upload and manage application environment images through Singularity containers. These images can support users with AI frameworks and HPC workloads, as well as others. Singularity containers may be built from Docker containers, imported from NVIDIA GPU Cloud (NGC), or other image repositories. Containers created by administrators are available to all users, and users can create container images for their individual use as well. Users looking to deploy a custom image can also create a custom template that will deploy the container and run workloads in that environment.

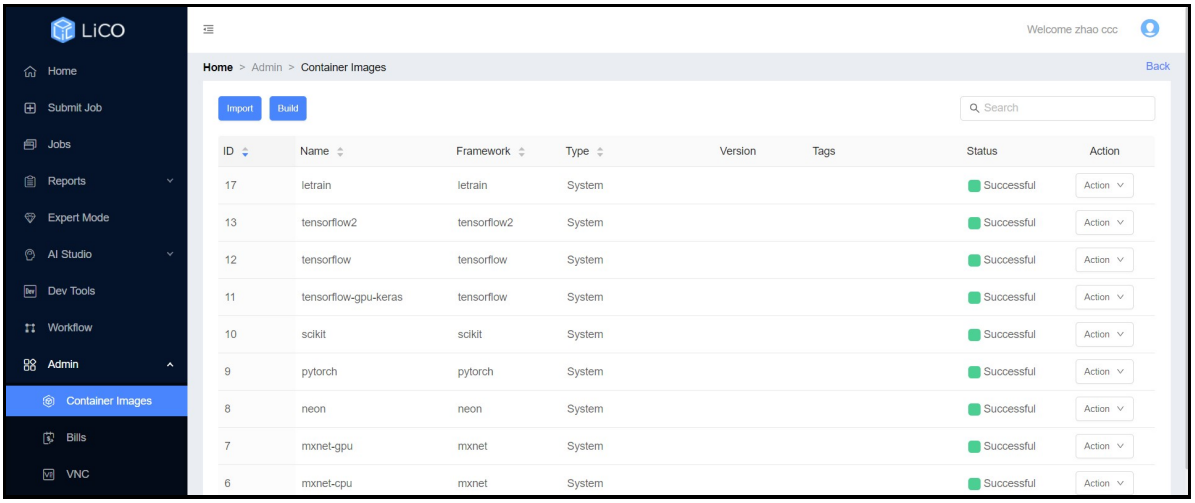

Figure 22. Singularity container management through the Administrator portal

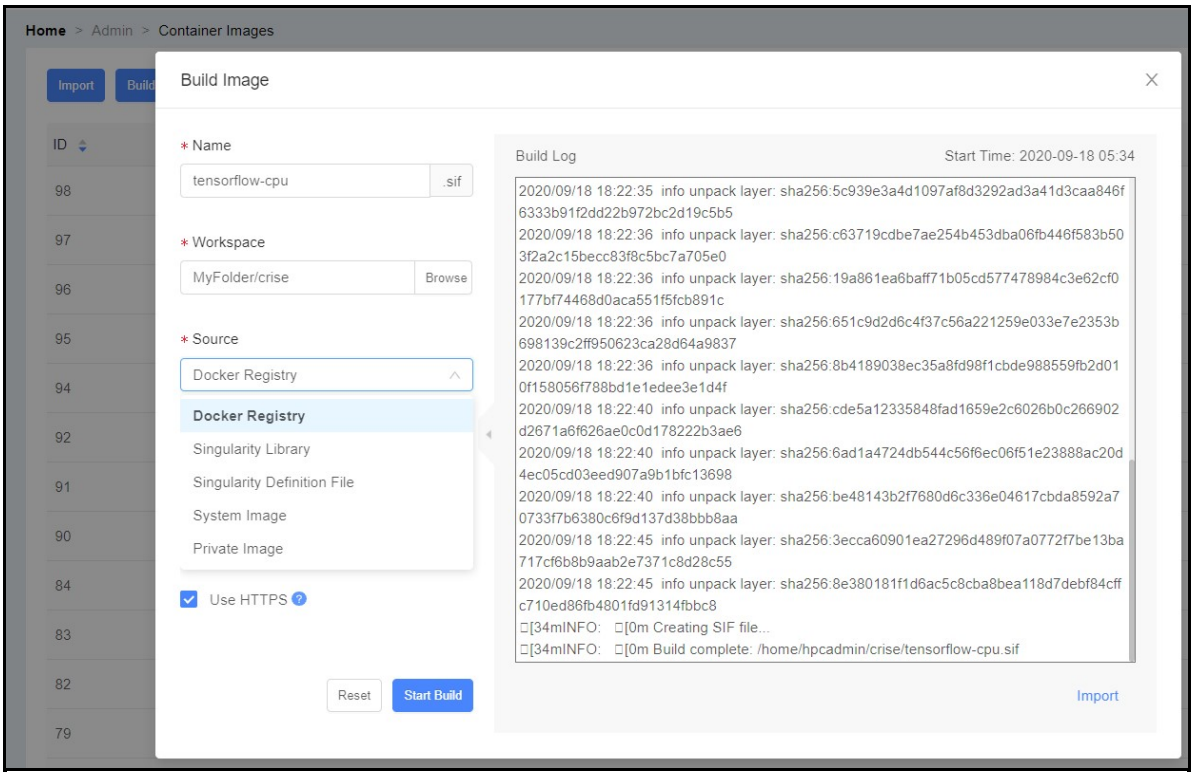

Figure 23. Singularity container building within LiCO

Expert mode

LiCO HPC/AI version provides more experienced cluster users console access to the user space in the LiCO management node, to execute Linux and SLURM commands directly. Expert mode enables users familiar with the underlying cluster orchestration choice in how they work, using either the command line, GUI or both in concert to facilitate their workflow.

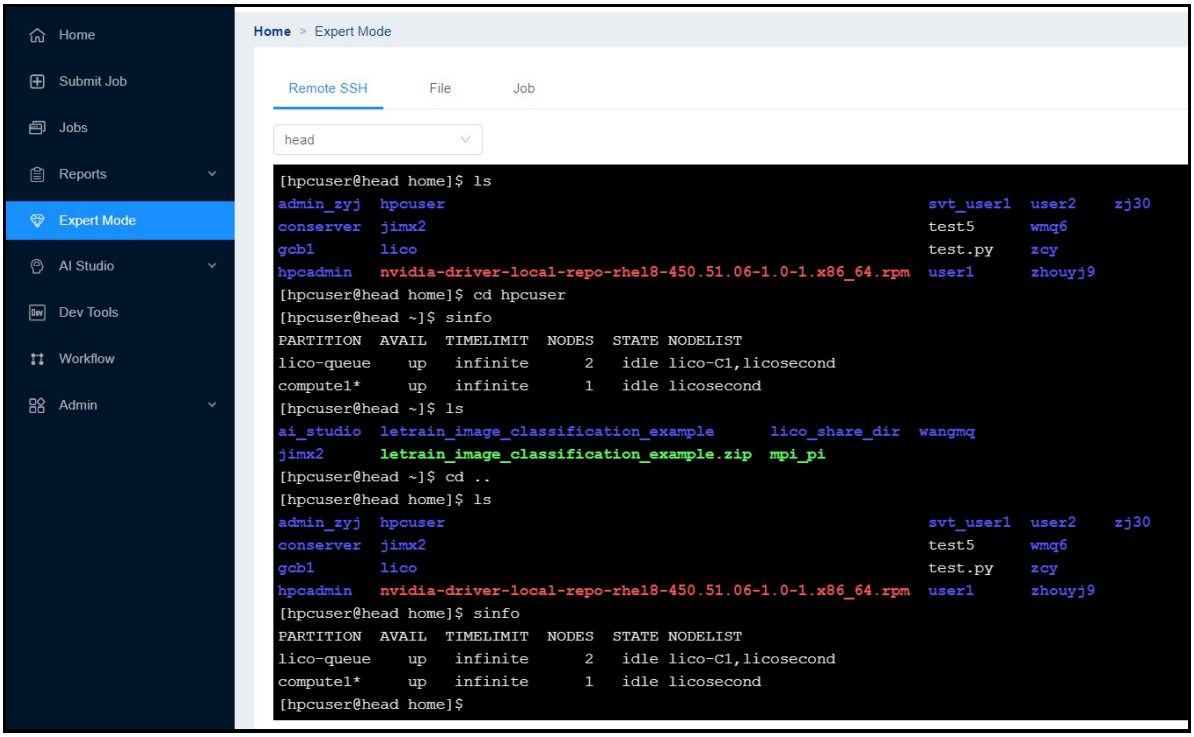

Figure 24. LiCO Expert mode console access

Reports

LiCO HPC/AI version provides expanded billing capabilities and provides the user access to monitor charges incurred for a date range via the Expense Reports subtab. Users can also download daily or monthly billing reports as a .xlsx file from the Admin tab.

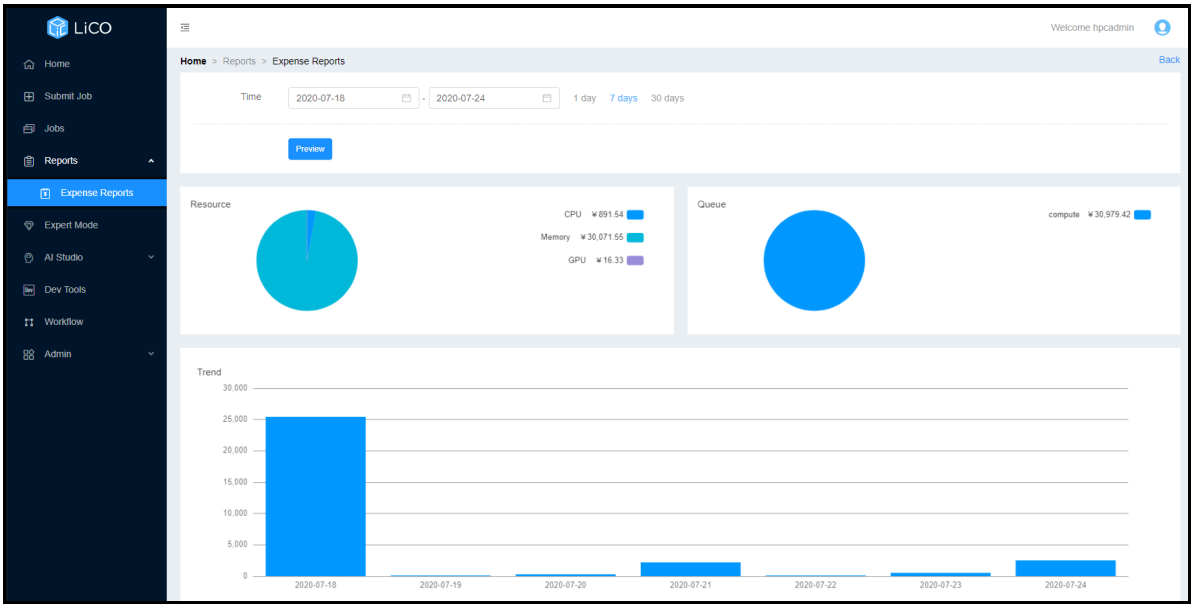

Figure 25. LiCO User view of Expense Reports

Features for LiCO Administrators

Features for LiCO K8S/AI version Administrators

For administrators of a Kubernetes-based LiCO environment, LiCO provides the ability to monitor activity, create and manage users, monitor LiCO-initiated activity, generate job and operational reports, enable container access for LiCO users, and view the software license currently installed in LiCO. LiCO K8S/AI version does not provide resource monitoring for the administrator, resources can be monitored at the Kubernetes level with a tool such as Kubernetes Dashboard.The following menus are available to administrators in LiCO K8S/AI:

- **Home menu for Administrators** – provides an at-a-glance view of LiCO jobs running and operational messages. For monitoring and managing cluster resources, the administrator can use a tool such as Kubernetes dashboard, Grafana, or other Kubernetes monitoring tools.
- **User Management menu** – provides dashboards to create, import and export LiCO users, and includes administrative actions to edit, suspend, or delete
- **Monitor menu** – provides a view of LiCO jobs running, allocating to the Kubernetes cluster, and completed jobs.This menu also allows the administrator to query and filter operational logs.
- **Reports menu** – allows administrators the ability to generate reports on jobs, for a given time interval. Administrators may export these reports as a spreadsheet, in a PDF, or in HTML. The reports menu also allows the administrator to view cluster utilization for a given date range.
- **Admin menu** – Provides the administrator to map container images for use in job submission templates, and download operations and web logs for LiCO.
- **Settings menu** – allows the administrator to view the currently active license for LiCO, including the license key, license tier and expiration date of the license.

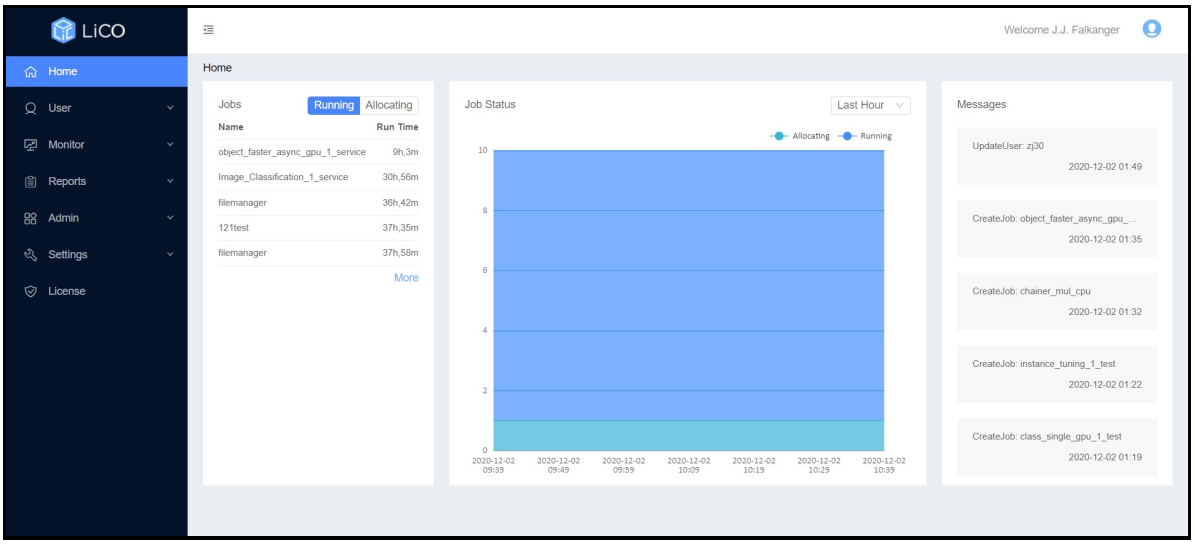

Figure 26. LiCO K8S/AI Administrator Home Menu

Features for LiCO HPC/AI version Administrators

For cluster administrators, LiCO provides a sophisticated monitoring solution, built on OpenHPC tooling. The following menus are available to administrators:

- **Home menu for administrators** – provides dashboards giving a global overview of the health of the cluster. Utilization is given for the CPUs, GPUs, memory, storage, and network. Node status is given, indicating which nodes are being used for I/O, compute, login, and management. Job status is also given, indicating runtime for the current job, and the order of jobs in the queue. The Home menu is shown in the following figure.

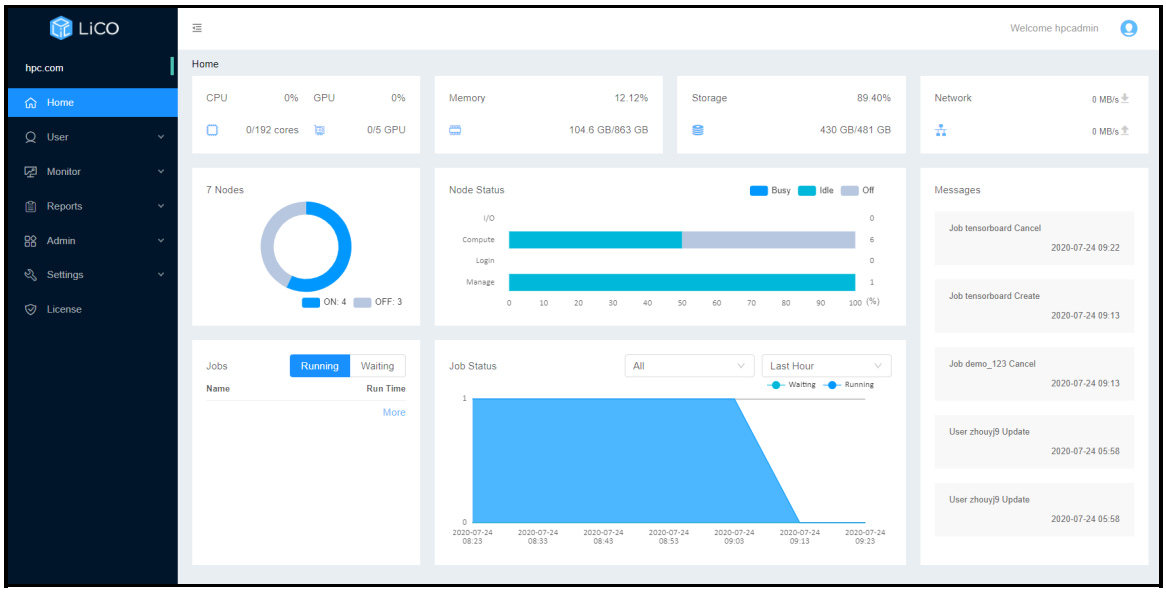

Figure 27. LiCO HPC/AI Administrator Home Menu

- **User Management menu** – provides dashboards to control user groups and users, determining permissions and access levels (based on LDAP) for the organization. Administrators can also control and provision billing groups for accurate accounting.
- **Monitor menu** – provides dashboards for interactive monitoring and reporting on cluster nodes, including a list of the nodes, or a physical look at the node topology. Administrators may also use the Monitor menu to drill down to the component level, examining statistics on cluster CPUs, GPUs, networking, jobs, and operations. Administrators can access alerts that indicate when these statistics reach unwanted values (for instance, GPU temperature reaching critical levels). These alerts are created using the Setting menu. Additionally, a large screen view is available to display a high-level summary of cluster status, and a cluster view is added in LiCO 6.2 for a focused view of compute resource utilization across the cluster. The figures below display the component and alert dashboards.

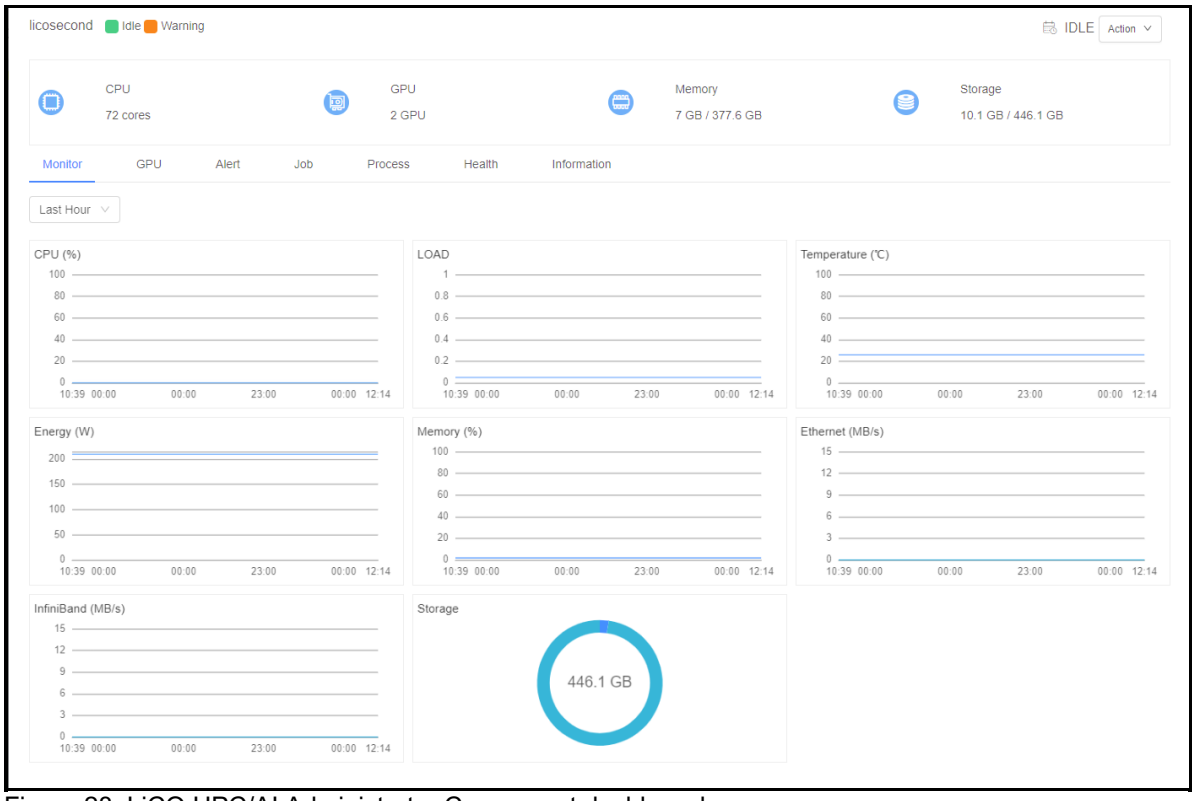

Figure 28. LiCO HPC/AI Administrator Component dashboard

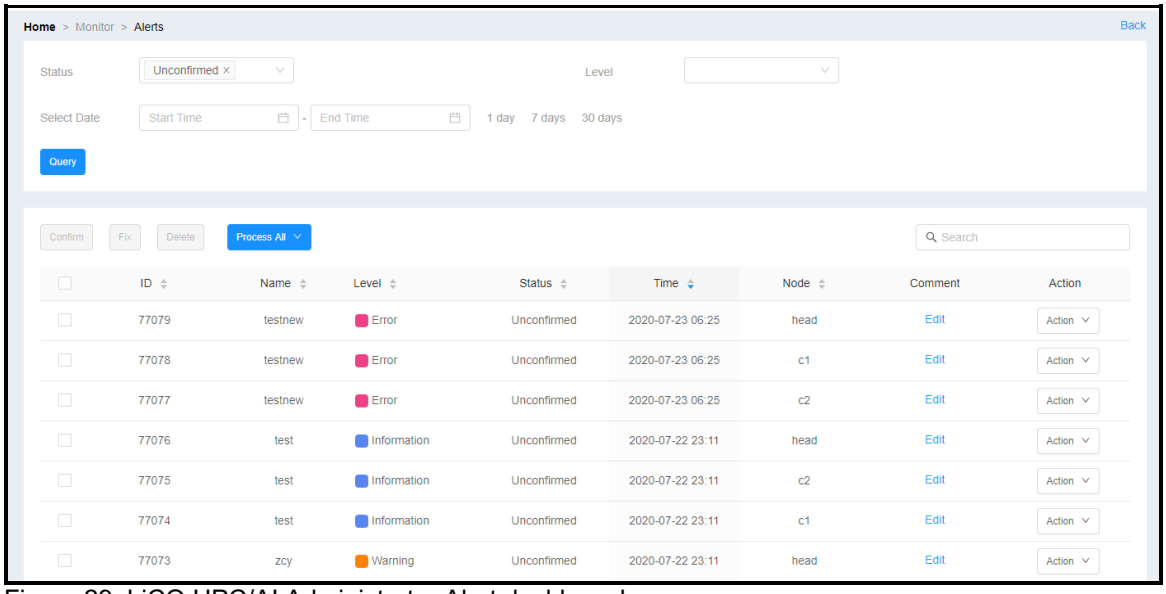

Figure 29. LiCO HPC/AI Administrator Alert dashboard

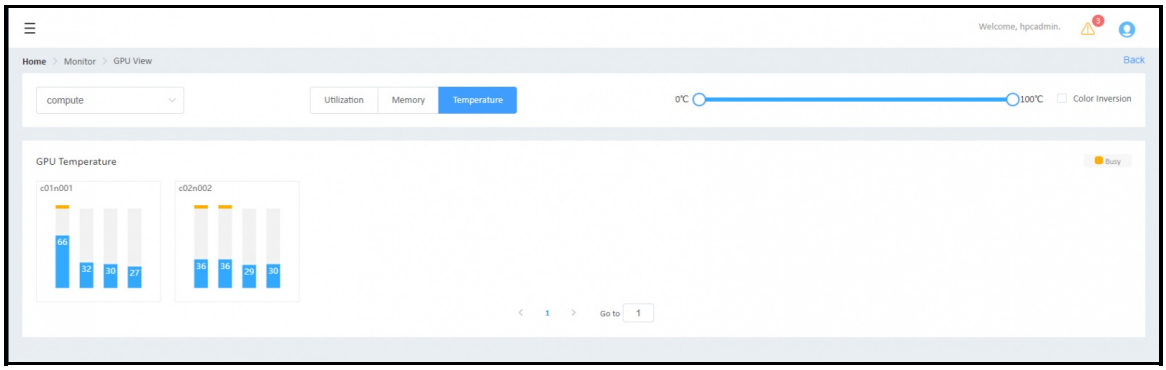

Figure 30. LiCO HPC/AI GPU View dashboard

- **Reports menu** – allows administrators the ability to generate reports on jobs, cluster utilization, alerts, and view current charges and cluster utilization.
- **Admin menu** – Provides the administrator with the capability to create Singularity images for use by all users, generate billing spreadsheets, examine processes and assets, monitor VNC sessions, and download web logs.
- **Settings menu** – allows administrators to set up automated notifications and alerts. Administrators may enable the notifications to reach users and interested parties via email, SMS, and WeChat. Administrators may also enable notifications and alerts via uploaded scripts.
  The Settings menu also allows administrators to create and modify queues. These queues allow administrators to subdivide hardware based on different types or needs. For example, one queue may contain systems that are exclusively machines with GPUs, while another queue may contain systems that only contain CPUs. This allows the user running the job to select the queue that is more applicable to their requirement. Within the Settings menu, administrators can also set the status of queues, bringing them up or down, draining them, or marking them inactive. Administrators can also limit which queues are available to users by user group.
- **License menu** – displays the software licenses active in LiCO including the number of licensed processing entitlements and the expiration date of the license.

Features for LiCO Operators

For the purpose of monitoring clusters but not overseeing user access, LiCO provides the Operator designation. LiCO Operators have access to a subset of the dashboards provided to Administrators; namely, the dashboards contained in the Home, Monitor, and Reports menus:

- **Home menu for operators** – provides dashboards giving a global overview of the health of the cluster. Utilization is given for the CPUs, GPUs, memory, storage, and network. Node status is given, indicating which nodes are being used for I/O, compute, login, and management. Job status is also given, indicating runtime for the current job, and the order of jobs in the queue.
- **Monitor menu** – Dashboard that enables interactive monitoring and reporting on cluster nodes, including a list of the nodes, or a physical look at the node topology. Operators may also use the Monitor menu to drill down to the component level, examining statistics on cluster CPUs, GPUs, jobs, and operations. Operators can access alarms that indicate when these statistics reach unwanted values (for instance, GPU temperature reaching critical levels.) These alarms are created by Administrators using the Setting menu (for more information on the Setting menu, see the Features for LiCO Administrators section.)
- **Reports menu** – allows operators the ability to generate reports on jobs, alerts, or actions for a given time interval. Operators may export these reports as a spreadsheet, in a PDF, or in HTML.

Subscription & support

LiCO HPC/AI is enabled through a per-CPU and per-GPU subscription and support entitlement model, which once entitled for the all the processors contained within the cluster, gives the customer access to LiCO package updates and Lenovo support for the length of the acquired term.

LiCO K8S/AI is enabled through tiered subscription and support entitlement licensing based on the number of GPU accelerators being accessed by running LiCO workloads (tiers are up to 4 GPU in use, up to 16 GPU in use, and up to 64 GPU in use). Additional licensing beyond 64 GPUs can be provided by contacting your Lenovo sales representative.

Lenovo will provide interoperability support for all software tools defined as validated with LiCO, and development support (Level 3) for specific Lenovo-supported tools only. Open source and supported-vendor bugs/issues will be logged and tracked with their respective communities or companies if desired, with no guarantee from Lenovo for bug fixes. Full support details are provided at the support links below for each respective version of LiCO. Additional support options may be available; please contact your Lenovo sales representative for more information.

LiCO can be acquired as part of a Lenovo Scalable Infrastructure (LeSI) solution or for “roll your own” (RYO) solutions outside of the LeSI framework, and LiCO software package updates are provided directly through the Lenovo Electronic Delivery system. More information on LeSI is available in the LeSI product guide, available from https://lenovopress.com/lp0900.

Validated software components

LiCO’s software packages are dependent on a number of software components that need to be installed prior to LiCO in order to function properly. Each LiCO software release is validated against a defined configuration of software tools and Lenovo systems, to make deployment more straightforward and enable support. Other management tools, hardware systems and configurations outside the defined stack may be compatible with LiCO, though not formally supported; to determine compatibility with other solutions, please check with your Lenovo sales representative.

The following software components are validated by Lenovo as part of the overall LiCO software solution entitlement:

LiCO HPC/AI version support

- Lenovo Development Support (L1-L3)
  - Graphical User Interface: LiCO
  - System Management & Provisioning: xCAT/Confluent
- Lenovo LiCO HPC/AI Configuration Support (L1 only)
  - Job Scheduling & Orchestration: SLURM; Torque/Maui (HPC only)
  - System Monitoring: Icinga v2
  - Container Support (AI): Singularity, CharlieCloud
  - AI Frameworks (AI): Caffe, Intel-Caffe, TensorFlow, MxNet, Neon, Chainer, Pytorch, Scikit-learn

The following software components are validated for compatibility with LiCO HPC/AI:

- Supported by their respective software provider
  - Operating System: CentOS/RHEL 8.3, SUSE SLES 15.2
  - File Systems: IBM Spectrum Scale, Lustre
  - Job Scheduling & Orchestration: IBM Spectrum LSF v10
  - Development Tools: GNU compilers, Intel Cluster Toolkit

LiCO K8S/AI version support

- Lenovo Development Support (L1-L3)
  - Graphical User Interface: LiCO
- Lenovo LiCO K8S/AI Configuration Support (L1 only)
  - AI Frameworks (AI): Caffe, Intel-Caffe, TensorFlow, MxNet, Neon, Chainer, Pytorch, Scikit-learn

Supported servers (LiCO HPC/AI version)

The following Lenovo servers are supported to run with LiCO HPC/AI. This server must run one of the supported operating systems as well as the validated software stack, as described in the Validated Software Components section.

- **ThinkSystem SR670 V2 –** The Lenovo ThinkSystem SR670 V2 is a versatile GPU-rich 3U rack server that supports eight double-wide GPUs including the new NVIDIA A100 and A40 Tensor Core GPUs, or the NVIDIA HGX A100 4-GPU offering with NVLink and Lenovo Neptune hybrid liquid-to-air cooling. The server is based on the new third-generation Intel Xeon Scalable processor family (formerly codenamed "Ice Lake"). The server delivers optimal performance for Artificial Intelligence (AI), High Performance Computing (HPC) and graphical workloads across an array of industries. For more information, see the SR670 V2 product guide.
- **ThinkSystem SD650 V2 –** The ThinkSystem SD650 V2 server is the next-generation high-performance server based on Lenovo's fourth generation Lenovo Neptune™ direct water cooling platform. With two third-generation Intel Xeon Scalable processors, the ThinkSystem SD650 V2 server combines the latest Intel processors and Lenovo's market-leading water cooling solution, which results in extreme performance in an extreme dense packaging, supporting your application From Exascale to Everyscale™. For more information, see the SD650 V2 product guide.
- **ThinkSystem SD650-N V2 –** The ThinkSystem SD650-N V2 server is the next-generation high-performance GPU-rich server based on Lenovo's fourth generation Lenovo Neptune™ direct water cooling platform. With four NVIDIA A100 SXM4 GPUs and two third-generation Intel Xeon Scalable processors, the ThinkSystem SD650-N V2 server combines advanced NVIDIA acceleration technology with the latest Intel processors and Lenovo's market-leading water cooling solution, which results in extreme performance in an extreme dense packaging supporting your accelerated application From Exascale to Everyscale™. For more information, see the SD650-N V2 product guide.
- **ThinkSystem SR650 V2 –** The Lenovo ThinkSystem SR650 V2 is an ideal 2-socket 2U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR650 V2 is a very configuration-rich offering, supporting 28 different drive bay configurations in the front, middle and rear of the server and 5 different slot configurations at the rear of the server. This level of flexibility ensures that you can configure the server to meet the needs of your workload. For more information, see the SR650 V2 product guide.
- **ThinkSystem SR630 V2 –** The Lenovo ThinkSystem SR630 V2 is an ideal 2-socket 1U rack server designed to take full advantage of the features of the 3rd generation Intel Xeon Scalable processors, such as the full performance of 270W 40-core processors, support for 3200 MHz memory and PCIe Gen 4.0 support. The server also offers onboard NVMe PCIe ports that allow direct connections to 12x NVMe SSDs, which results in faster access to store and access data to handle a wide range of workloads. For more information, see the SR630 V2 product guide.
- **ThinkSystem SD530** – The Lenovo ThinkSystem SD530 is an ultra-dense and economical two-socket server in a 0.5U rack form factor. With up to four SD530 server nodes installed in the ThinkSystem D2 enclosure, and the ability to cable and manage up to four D2 enclosures as one asset, you have an ideal high-density 2U four-node (2U4N) platform for enterprise and cloud workloads. The SD530 also supports a number of high-end GPU options with the optional GPU tray installed, making it an ideal solution for AI Training workloads. For more information, see the SD530 product guide.
- **ThinkSystem SD650** – The Lenovo ThinkSystem SD650 direct water cooled server is an open, flexible and simple data center solution for users of technical computing, grid deployments, analytics workloads, and large-scale cloud and virtualization infrastructures. The direct water cooled solution is designed to operate by using warm water, up to 50°C (122°F). Chillers are not needed for most customers, meaning even greater savings and a lower total cost of ownership. The ThinkSystem SD650 is designed to optimize density and performance within typical data center infrastructure limits, being available in a 6U rack mount unit that fits in a standard 19-inch rack and houses up to 12 water-cooled servers in 6 trays. For more information, see the SD650 product guide.
- **ThinkSystem SR630** – Lenovo ThinkSystem SR630 is an ideal 2-socket 1U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR630 server is designed to handle

a wide range of workloads, such as databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), infrastructure security, systems management, enterprise applications, collaboration/email, streaming media, web, and HPC. For more information, see the SR630 product guide.

- **ThinkSystem SR650** – The Lenovo ThinkSystem SR650 is an ideal 2-socket 2U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR650 server is designed to handle a wide range of workloads, such as databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), enterprise applications, collaboration/email, and& business analytics and big data. For more information, see the SR650 product guide.
- **ThinkSystem SR670** – The Lenovo ThinkSystem SR670 is a purpose-built 2 socket 2U accelerated server, supporting up to 8 single-wide or 4 double-wide GPUs and designed for optimal performance required by both Artificial Intelligence and High Performance Computing workloads. Supporting the latest NVIDIA GPUs and Intel Xeon Scalable processors, the SR670 supports hybrid clusters for organizations that may want to consolidate infrastructure, improving performance and compute power, while maintaining optimal TCO. For more information, see the SR670 product guide.
- **ThinkSystem SR950** – The Lenovo ThinkSystem SR950 is Lenovo's flagship server, suitable for mission-critical applications that need the most processing power possible in a single server. The powerful 4U ThinkSystem SR950 can expand from two to as many as eight Intel Xeon Scalable Family processors. The modular design of SR950 speeds upgrades and servicing with easy front or rear access to all major subsystems that ensures maximum performance and maximum server uptime. For more information, see the SR950 product guide.
- **ThinkSystem SR655** – The Lenovo ThinkSystem SR655 is a 1-socket 2U server that features the AMD EPYC 7002 "Rome" family of processors. With up to 64 cores per processor and support for the new PCIe 4.0 standard for I/O, the SR655 offers the ultimate in single-socket server performance. ThinkSystem SR655 is a multi-GPU optimized rack server, providing support for up to 6 low-profile GPUs or 3 double-wide GPUs. For more information, see the SR655 product guide.
- **ThinkSystem SR635** – The Lenovo ThinkSystem SR635 is a 1-socket 1U server that features the AMD EPYC 7002 "Rome" family of processors. With up to 64 cores per processor and support for the new PCIe 4.0 standard for I/O, the SR635 offers the ultimate in single-socket server performance. For more information, see the SR635 product guide.
- **ThinkSystem SR645** – The Lenovo ThinkSystem SR645 is a 2-socket 1U server that features the AMD EPYC 7002 "Rome" family of processors. With up to 64 cores per processor and support for the new PCIe 4.0 standard for I/O, the SR645 offers the ultimate in two-socket server performance in a space-saving 1U form factor. For more information, see the SR645 product guide.
- **ThinkSystem SR665** – The Lenovo ThinkSystem SR665 is a 2-socket 2U server that features the AMD EPYC 7002 "Rome" family of processors. With support for up to 8 single-wide or 3 double-wide GPUs, up to 64 cores per processor and support for the new PCIe 4.0 standard for I/O, the SR665 offers the ultimate in two-socket server performance in a 2U form factor. ThinkSystem SR665 is a multi-GPU optimized rack server, providing support for up to 8 low-profile GPUs or 3 double-wide GPUs. For more information, see the SR665 product guide.
- **ThinkSystem SR850** – The Lenovo ThinkSystem SR850 is a 4-socket server that features a streamlined 2U rack design that is optimized for price and performance, with best-in-class flexibility and expandability. The SR850 now supports second-generation Intel Xeon Scalable Family processors, up to a total of four, each with up to 28 cores. The ThinkSystem SR850's agile design provides rapid upgrades for processors and memory, and its large, flexible storage capacity helps to keep pace with data growth. For more information, see the SR850 product guide.

Additional Lenovo ThinkSystem and System x servers may be compatible with LiCO. Contact your Lenovo sales representative for more information.

LiCO Implementation services

Customers who do not have the cluster management software stack required to run with LiCO may engage Lenovo Professional Services to install LiCO and the necessary open-source software. Lenovo Professional Services can provide comprehensive installation and configuration of the software stack, including operation verification, as well as post-installation documentation for reference. Contact your Lenovo sales representative for more information.

Client PC requirements

A web browser is used to access LiCO's monitoring dashboards. To fully utilize LiCO's monitoring and visualization capabilities, the client PC should meet the following specifications:

- Hardware: CPU of 2.0 GHz or above and 8 GB or more of RAM
- Display resolution: 1280 x 800 or higher
- Browser: Chrome (v62.0 or higher) or Firefox (v56.0 or higher) is recommended

Related links

For more information, see the following resources:

- LiCO website:
  https://www.lenovo.com/us/en/data-center/software/Lenovo-Intelligent-Computing-Orchestration/p/WMD00000356
- LiCO HPC/AI (Host) Support website:
  https://support.lenovo.com/us/en/solutions/HT507011
- LiCO K8S/AI (Kubernetes) Support website:
  https://support.lenovo.com/us/en/solutions/HT509422
- Lenovo DSCS configurator:
  https://dcsc.lenovo.com
- Lenovo AI website:
  https://www.lenovo.com/us/en/data-center/solutions/analytics-ai/
- Lenovo HPC website:
  https://www.lenovo.com/us/en/data-center/solutions/hpc/
- LeSI website:
  https://www.lenovo.com/us/en/data-center/servers/high-density/Lenovo-Scalable-Infrastructure/p/WMD00000276
- OpenHPC User Resources:
  https://github.com/openhpc/ohpc/wiki/User-Resources
- Intel oneAPI
  https://software.intel.com/content/www/us/en/develop/tools.html

Related product families

Product families related to this document are the following:

- Artificial Intelligence
- High Performance Computing

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document, LP0858, was created or updated on June 1, 2021.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  http://lenovopress.com/LP0858
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at http://lenovopress.com/LP0858.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
From Exascale to Everyscale
Lenovo Neptune
System x®
ThinkSystem

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Other company, product, or service names may be trademarks or service marks of others.