

Lenovo Compute Orchestration in HPC Data Centers with Slurm

Solution Brief

Today almost every business has come to accept that they need to start automating some of their manual business processes. Automating processes reduces the chance of errors and can help businesses become more productive and reliable. Customers are trying to get the best efficiency out of their hardware and sometimes that means they need to handle more work than resources. A job scheduler manages queue(s) of work to support complex scheduling algorithms and supports and optimize resource limits alongside multiple other benefits.

Workload management helps you prioritize tasks more efficiently, improve the balance of work across your team, and create more accurate project schedules. As a result, your team has the capacity and confidence to deliver quality work.

Slurm Workload Manager is a modern, open-source scheduler developed and maintained by SchedMD®. Slurm maximizes workload throughput, scale, reliability, and results in the fastest possible time while optimizing resource utilization and meeting organizational priorities. By automating job scheduling, Slurm is designed to satisfy the demanding needs of high-performance computing (HPC), high throughput computing (HTC) and AI. Slurm's automation capabilities simplify administration, accelerate job execution, and improve end user productivity all while reducing cost and error margins.

Lenovo has an active partnership with SchedMD to enable customers with the right solution to schedule their workloads and optimize their resources.

Slurm is fully integrated into Lenovo Intelligent Computing Orchestration (LiCO) to provide users with an easy way to schedule their HPC and AI workloads and for optimizing them. Customers can benefit from L3 support for Slurm as part of the Lenovo HPC & AI Software Stack offering.

Overview of Slurm Workload Manager

Slurm Workload Manager is the market-leading, free open-source workload manager designed specifically to satisfy the demanding needs of high-performance computing (HPC), high throughput computing (HTC) and AI. Slurm maximizes workload throughput, scale, reliability, and results in the fastest possible time while optimizing resource utilization and meeting organizational priorities.

Slurm is set apart by the way it manages the GPUs. Users can schedule workloads to GPUs similar to CPUs, allowing for greater flexibility and control in how scheduling gets done.

SchedMD is the core company behind Slurm, distributing and maintaining the workload manager software. SchedMD is the sole provider for Slurm support, development, training, installation, and configuration.

Slurm provides three main added values:

- Allocates access to resources
- Provides a framework to run and monitor jobs on allocated nodes
- Manages a job queue for competing resource requests

Slurm capabilities include:

- Allocates and optimizes exclusive/non-exclusive access to ultra-granular resources: Slurm provides this for users and jobs for duration of time for each workload including network topology, fair share scheduling, advanced reservations, preemption, resource limits, and accounting factors.
- Framework for starting, executing, and monitoring workloads on allocated nodes and/or GPUs: Slurm capabilities include accounting for task level in real time, power-consumption, and API usage, as well as automatically re-queuing jobs.
- Simplifies management and arbitrates contention for resources: Slurm manages a queue of pending work according to organizational and work priorities.
- Accelerates processing and throughput of work: Slurm provides extensive policies and algorithms, including those for high throughput computing workloads, to support 1000s of job submissions/second.
- Slurm natively supports elastic and cloud bursting capabilities.

Please reference the [Slurm Workload Manager documentation](#) for more information on these Slurm capabilities.

Customers select and deploy Slurm for:

- Massive scalability to handle performance requirements for large cluster and exascale supercomputer needs spanning HPC and AI workloads with proven reliability
- First class resource management for GPUs and ultra-granular task allocation by specialized resources (cores, GPUs, threads, etc.)
- Unmatched workload throughput, supporting 12 thousand+ jobs per minute, 17 million+ jobs per day, 120 million jobs per week.
- Advanced policies, scheduling algorithms, and queues/partitions management that simplifies and optimizes management of complex workload mix to meet project and organizational priorities
- Optimize utilization to get more value from existing HPC investments
- Agile innovation and integration to meet needs, driven by strong open-source community and partnerships

Power savings

Slurm provides an integrated power saving mechanism for powering down idle nodes. Nodes that remain unused for a configurable period can be placed in a power saving mode, reducing power consumption. A power capping limit to manage power usage is also available.

ExaScale

Slurm workload manager enables unmatched workload throughput across massive numbers of jobs and massive exascale Lenovo HPC & AI infrastructure resources to deliver innovation and insights faster for a competitive edge. Slurm offers massive scalability to handle clusters with more than 100K nodes and GPUs increasing the throughput 5-10x. Slurm's scheduling and resource management capabilities handle both effortlessly, including job arrays to submit millions of tasks in milliseconds and ultra-granular task allocation by specialized resources (cores, GPUs, threads, etc.). The SchedMD team fine-tunes configurations to workload mix and priorities while improving resource consumption efficiency 30-40%.

Distribution media & Software documentation

Slurm Workload Manager is available for download as free, open-source software along with documentation and installation guides at <https://slurm.schedmd.com>.

For information on purchasing SchedMD Slurm Workload Manager Support please contact your Lenovo sales representative. Customers can also access additional details and information on SchedMD Slurm support services at <https://www.schedmd.com/services.php>.

Operating system compatibility

Slurm Workload Manager has been thoroughly tested on most popular Linux distributions. Some features are limited to recent releases and newer Linux kernel versions. Currently supported distributions include:

- Red Hat Enterprise Linux 7, CentOS 7, Scientific Linux 7
- Red Hat Enterprise Linux 8 and RHEL 8 derivatives
- Red Hat Enterprise Linux 9 and RHEL 9 derivatives
- SUSE Linux Enterprise Server 12
- SUSE Linux Enterprise Server 15
- Ubuntu 22.04

Third-party software integration

Slurm Workload Manager's open-source structure, leadership, Plug-in and REST API capabilities allows for flexible integration with hundreds of HPC and AI end user applications, workflow engines and tools across a range of vertical industries. See <https://slurm.schedmd.com/documentation.html> for more information.

Hardware compatibility

Slurm Workload Manager is supported on all Lenovo HPC/AI ThinkSystem Servers.

Slurm Workload Manager Deep Dive

Slurm is a modern scheduler developed with a plug-in-based architecture, allowing it to support both large and small HPC, HTC, and AI environments. Slurm is highly scalable and reliable resulting in the fastest possible compute times of any scheduler on the market. Slurm's plug-in-based architecture enables to load only the components needed to complete specific end user tasks. This lightweight and flexible structure allows for more optimization and control in scheduling operations. Slurm supports the complexities of on-prem, hybrid, or cloud workspaces.

Plugins can add a wide range of features, including resource limit management and accounting, as well as support for advanced scheduling algorithms. Slurm supports the use of GPUs via the concept of Generic Resources (GRES). GRES are computing resources associated with a Slurm node, which can be allocated to jobs and steps.

The most obvious example of GRES use would be GPUs. GRES are identified by a specific name and use an optional plugin to provide device-specific support. The main Slurm cluster configuration file, `slurm.conf`, must explicitly specify which GRES are available in the cluster.

The following is an example of a `slurm.conf` file, which configures four GPUs supporting Multi-Process Service (MPS), with 4 GB of network bandwidth.

```
GresTypes=gpu,mps,bandwidth
NodeName=cn[0-7]
Gres=gpu:tesla:8,gpu:kepler:2,mps:400,bandwidth:lustre:no_consume:4G
```

In addition, Slurm nodes that need to expose GRES to jobs should have a `gres.conf` file. This file describes which types of Generic Resources are available on the node, their count, and which files and processing cores should be used with those resources.

When possible, Slurm automatically determines the GPUs on the system using NVIDIA's NVML library. The NVML library (which powers the `nvidia-smi` tool) numbers GPUs in order by their PCI bus IDs. GPU device files (e.g. `/dev/nvidia1`) are based on the Linux minor number assignment, while NVML's device numbers are assigned via PCI bus ID, from lowest to highest. Mapping between these two is indeterministic and system dependent, and could vary between boots after hardware or OS changes.

[CUDA Multi-Process Service \(MPS\)](#) provides a mechanism where GPUs can be shared by multiple jobs, where each job is allocated some percentage of the GPU's resources. The total count of MPS resources available on a node should be configured in the `slurm.conf` file, for example:

```
NodeName=cn[1-16]
Gres=gpu:2,mps:400
```

Job requests for MPS will be processed the same as any other GRES except that the request must be satisfied using only one GPU per node and only one GPU per node may be configured for use with MPS. For example, a job request for `--gres=mps:50` will not be satisfied by using 20% of one GPU and 30% of a second GPU on a single node. Multiple jobs from different users can use MPS on a node at the same time. Note that GRES types of GPU and MPS cannot be requested within a single job. Also jobs requesting MPS resources can not specify a GPU frequency.

Quality of Service (QOS)

End users can specify Quality of Service (QOS) for each job submitted to Slurm. Correct configuration for the quality of service associated with a job ensures the right jobs run at the right time. Slurm quells internal politics and competing workloads by helping businesses schedule competing needs while hitting goals and maintaining processes.

Slurm Architecture

The following figure shows the components of the Slurm architecture.

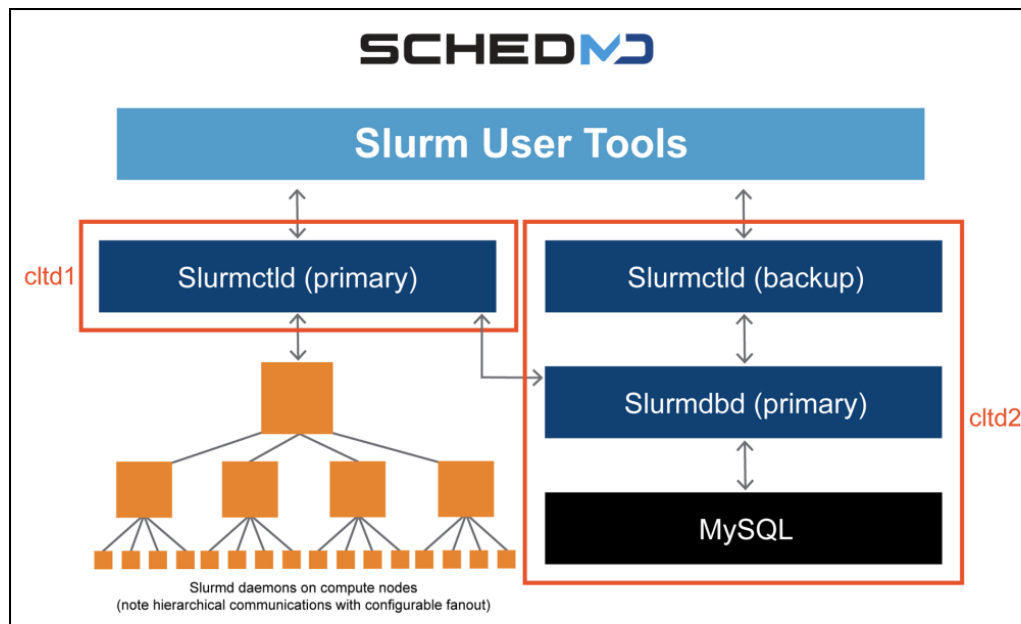


Figure 1. Slurm components

The components are as follows:

- Slurm has a centralized manager, **slurmctld**, to monitor resources and work. There may also be a backup manager to assume those responsibilities in the event of failure.
- Each compute server (node) has a **slurmd** daemon, which can be compared to a remote shell: it waits for work, executes that work, returns status, and waits for more work. The slurmd daemons provide fault-tolerant hierarchical communications.
- There is an optional **slurmdbd** (Slurm DataBase Daemon) which can be used to record accounting information for multiple Slurm-managed clusters in a single database.
- There is an optional **slurmrestd** (Slurm REST API Daemon) which can be used to interact with Slurm through its REST API.
- User tools include:

- **srun** to initiate jobs
- **scancel** to terminate queued or running jobs
- **sinfo** to report system status
- **squeue** to report the status of jobs
- **sacct** to get information about jobs and job steps that are running or have completed.
- **sview** graphically reports system and job status including network topology.
- **scontrol** is an administrative tool to monitor and/or modify configuration and state information on the cluster.
- **sacctmgr** is an administrative tool used to manage the database and is used to identify the clusters, valid users, valid bank accounts, etc. APIs are available for all functions.

Slurm Workload Management for Lenovo HPC & AI Software Stack

Slurm is part of the Lenovo HPC & AI Software Stack, integrated as an open source, agile, modern choice to provide highly scalable, fault-tolerant workload management policies, algorithms and reporting.

The Lenovo HPC & AI Software Stack provides a fully tested and supported, complete but customizable HPC software stack to enable the administrators and users in optimally and environmentally sustainable utilizing their Lenovo Supercomputers.

The software stack was designed to abstract the users from the complexity of HPC cluster orchestration and AI workloads management, making open-source HPC software consumable for every customer. With Lenovo, we make the benefits of Slurm even more accessible through a web portal integration that allows your teams to focus on their work while Slurm manages their workloads. Users can use Slurm without the need to learn commands or understand the complexity of the cluster just by using a very intuitive, easy-to-learn interface, which will bring the advantages of Slurm to more customers and users.

These capabilities enable Lenovo customers to achieve faster job processing, optimal utilization of specialized Lenovo system resources, and increased throughput that are all aligned with organizational priorities. Slurm delivers optimized workload performance for all Lenovo supported HPC/AI systems and servers.

LiCO Integration of Slurm

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, and HPC workloads. LiCO interfaces with an open-source software orchestration stack, enabling the convergence of AI onto an HPC cluster.

LiCO supports Slurm Workload Manager as a scheduler for queue management including viewing, creating, and modifying queues. HPC and AI users give their inputs (scripts, containers, resources requested) through LiCO's interface and LiCO creates a Slurm batch script based on the inputs to deploy and manage the workload.

The version of the Slurm supported in LiCO is the version of the latest released package from Lenovo OpenHPC.

Deployment of the solution is simple using LiCO GUI Installer. LiCO GUI Installer is a tool that simplifies HPC cluster deployment and LiCO setup. It runs on the management node, and it can use Lenovo Confluent to deploy the OS on the compute nodes. At the end of the installation, both LiCO and Slurm will be installed in the cluster specified in the installer by the customer.

The Slurm master node and LiCO management node need to be deployed on the same physical node in the cluster.

Workload management with LiCO

Queues allow administrators to subdivide hardware based on different types or needs. Slurm is fully integrated into LiCO, administrators can create and edit queues directly from within LiCO, without going into the Slurm scheduler directly via the console. That simplifies ongoing management for enterprise AI environments without HPC software tool expertise.

For the administrator role in LiCO, the Scheduler page is available, which allows administrators to create, edit, delete queues, set queue state, and set node state.

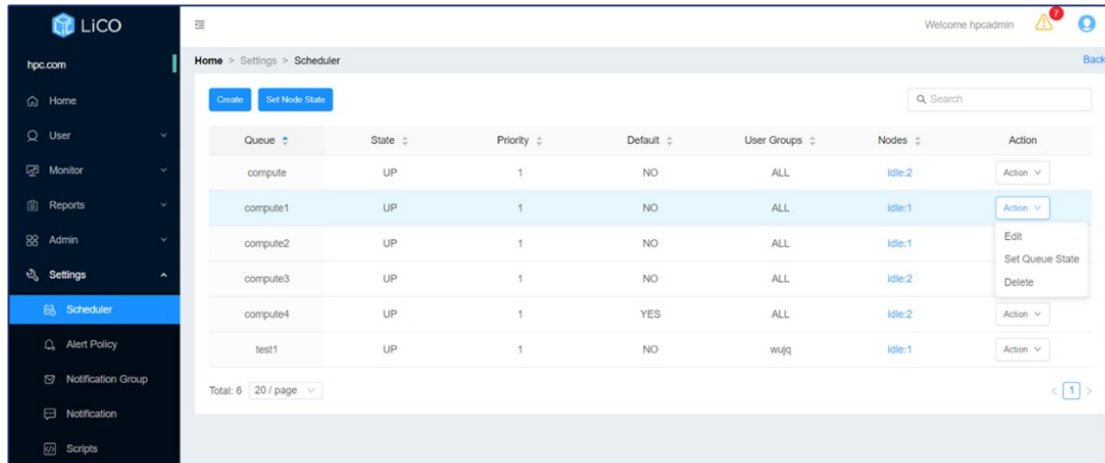


Figure 2. Scheduler management page in LiCO

Queues allow administrators to subdivide hardware based on different types or needs.

Creation of a queue can be done easily from the interface, allowing even users that are not familiar with the command line to perform the task. The interface allows to specify the nodes, the priority, the maximum running time for jobs, to specify whether compute resources (individual CPUs) in that queue can be shared by multiple jobs (an optional job count specifies how many jobs can be allocated to use each resource).

Administrators can set the state for a queue that specifies whether jobs can be allocated to nodes or queued in the same queue.

For more experienced users, command line actions can be performed. In queue management, the current user can log in to the head node and utilize Slurm scheduler command lines. After completing the creation of the queue and restarting the Slurm-related services, the newly created queue can be viewed on the webportal interface. The queue may be modified lately by changing the configuration file of Slurm.

From the same page in the webportal, users can edit, delete, set the state of a queue, and set the node state.

The screenshot shows a 'Create Queue' dialog box with the following fields and values:

- Queue: (empty text input)
- Nodes: (empty text input)
- Default:
- Priority: 1
- Max Time: UNLIMITED
- Over Subscribe: YES (dropdown menu)
- Over Subscribe Value: 4
- User Groups: All (dropdown menu)
- State: UP (dropdown menu)

Buttons: Cancel, Submit

Figure 3. Creating a queue in LiCO

Slurm Quality of Service

Some customers, especially AI cluster customers, have not used an HPC scheduler before, but they have the need to limit computing resource usage for different users or groups. LiCO supports complex scheduling strategy with Slurm Quality of Service to help fulfil that need.

Quality of Service (QoS) is a method provided by Slurm to define cluster computing resource limit rules. Users can specify a QoS for each job submitted to Slurm. The administrator can dynamically adjust the QoS through the Slurm command to realize the operation and maintenance of the cluster. In order to make the QoS rules take effect, it is necessary to establish relationships with queues (Partition), user groups (Account), and users (User).

Common QoS configurations are:

- Used to limit resources and runtime. This type of QoS is often associated with queues.
- Used to limit the number of jobs. This type of QoS is often associated with user groups. Using the above two types of QoS configurations can basically meet the requirements of most resource constraints, and you can use the command line to adjust QoS to meet O&M requirements.

Note: The current version of LiCO does not support setting the QoS of Slurm through the web interface, which needs to be done through the command line.

Job management

Job management can be performed on the LiCO interface. An administrator can view and act on a job by giving commands to the scheduler.

Note: If a job is submitted through Slurm command lines, it will not start billing on the LiCO system.

Slurm Workload Manager support

While Slurm is an open-source workload manager, establishing best practices for scheduling optimization requires expert training and support.

Lenovo provides support and services in partnership with SchedMD.

SchedMD is the core company behind Slurm, distributing and maintaining the workload manager software. SchedMD is also the sole provider for Slurm support, development, training, and configuration. Secure and enhance your Slurm environment, with SchedMD, so that your software is risk averse and performs the way you need it to.

Security

The best way to secure your Slurm environment is through SchedMD support. When a security breach occurs, end users can receive information about hot fixes from SchedMD engineers immediately.

Lenovo will provide interoperability support for all software tools defined as validated with LiCO, Slurm included. 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.

Support of LiCO software requires Subscription and Support purchased together with your solution. Please contact your Lenovo sales representative for further details.

- Lenovo provides support in English globally and in Chinese for China (24x7)
- Support response times are as follows:
 - Severity 1 issues response is 1 business day
 - Other issues: 3 business days

Technical Support for software means the provision of telephone or web-based technical assistance by Lenovo to Customer's technical contact(s) with respect to any software defects, errors and product problems exhibited on supported Lenovo configurations.

Technical support does not cover help with the initial installation of the product, software how-to, training and or configuring the production environment. If you need assistance in these areas, please contact your Lenovo Sales Representative or Lenovo Business Partner for the best service offering.

Lenovo contacting support

To contact Lenovo support:

1. Go to <https://datacentersupport.lenovo.com>
2. Search for **LiCO - Lenovo Intelligent Computing Orchestration** in the search field or select **Solutions and Software -> LiCO - Lenovo Intelligent Computing Orchestration**.
3. You will see both Lenovo phone numbers for support for your country and also the option to submit an electronic support ticket

For more details and additional information on the supported components please refer to <https://support.lenovo.com/us/en/solutions/HT507011>.

SchedMD is the core developer and services provider for Slurm providing support, consulting, configuration, development and training services to accelerate workload results with proven best practices and innovation. SchedMD offers 5-10x more complex HPC and AI scheduling experience, including half of the biggest systems in the TOP500, to optimize the speed, throughput and resource utilization for each unique workload mix so organizations can get more results faster and easier. SchedMD expert services enable organizations to quickly implement, maximize throughput, manage complexity, and easily grow their high-performance workloads on Lenovo HPC and AI solutions.

Customers add SchedMD Slurm support for:

- Consulting and implementation expertise that speeds custom configuration tuning to increase throughput and utilization efficiency on complex and large-scale systems
- Support services that ensure cluster workload management configuration continually processes workload at peak levels as mix and scale evolve for improved productivity
- Unique HPC workload expertise that spans Slurm development, software & computer engineering,

and systems administration for resolution speed and quality, without escalation delays

- Best practices for optimal workload performance from an expert team with 5-10x more complex HPC scheduling experience, including half of the biggest systems in the TOP500
- Tailored Slurm expert training that empowers users on harnessing Slurm capabilities with Lenovo exascale technologies to speed projects and increase adoption
- Cloud workload management expertise and proven best practices to ensure optimal performance and streamlined management for workloads across on-prem systems, public and private clouds

SchedMD Slurm Support for Lenovo HPC Systems

Slurm Support and services are part of the Lenovo HPC & AI software stack, integrated as open source, flexible, and modern choices to manage complex workloads for faster processing and optimal utilization of the large-scale and specialized high-performance and AI resource capabilities needed per workload provided by Lenovo systems.

SchedMD Slurm Support service capabilities for Lenovo HPC systems include:

- Level 3 Support: High-performance systems must perform at high utilization and performance to meet end users and management return on the investment expectations. Customers covered by a support contract can reach out to SchedMD Slurm experts to help resolve complex workload management issues and questions quicker, including detailed answers directly from the Slurm Development team, instead of taking weeks or even months to try to resolve them in-house.
- Configuration assistance: Valuable assistance when the customer system is initially being configured to use Slurm or being modified as requirements change. Customers can review cluster requirements, operating environment, and organizational goals with a Slurm engineer who will work with them to optimize the configuration to achieve their needs.
- Tailored Slurm Training: SchedMD can provide customized Slurm training for Lenovo customers who need or desire it. Lenovo representatives or customers can contact Jess Arrington at SchedMD, jess@schedmd.com, to request a training quote. A customer scoping call before the remote or onsite training ensures coverage of their specific use cases and needs for the in-depth and comprehensive Slurm technical training is delivered in a hands-on lab workshop format to ensure users are empowered on Slurm best practices as well as site-specific use cases and configuration optimization.

Slurm support and services part numbers

The following Lenovo product numbers can be used to purchase SchedMD L3 Slurm support, consulting, and training services.

Table 1. Slurm support and services part numbers

Part number	Description
SchedMD Slurm Services for Lenovo HPC Systems	
7S09001MWW	SchedMD Slurm Onsite or Remote 3-day Training
7S09001NWW	SchedMD Slurm Consulting w/Sr.Engineer 2REMOTE Sessions
SchedMD Slurm Support for Lenovo HPC Systems	
7S09001PWW	SchedMD L3 Slurm support up to 100 Sockets/GPUs 3Y
7S09001QWW	SchedMD L3 Slurm support up to 100 Sockets/GPUs 5Y
7S09001RWW	SchedMD L3 Slurm support up to 100 Sockets/GPUs additional 1Y
7S09001SWW	SchedMD L3 Slurm support 101-1000 Sockets/GPUs 3Y
7S09001TWW	SchedMD L3 Slurm support 101-1000 Sockets/GPUs 5Y
7S09001UWW	SchedMD L3 Slurm support 101-1000 Sockets/GPUs additional 1Y
7S09001VWW	SchedMD L3 Slurm support 1001-5000+ Sockets/GPUs 3Y
7S09001WWW	SchedMD L3 Slurm support 1001-5000+ Sockets/GPUs 5Y
7S09001XWW	SchedMD L3 Slurm support 1001-5000+ Sockets/GPUs additional 1Y
7S09001YWW	SchedMD L3 Slurm support up to 100 Sockets/GPUs 3Y EDU&GOV
7S09001ZWW	SchedMD L3 Slurm support up to 100 Sockets/GPUs 5Y EDU&GOV
7S090022WW	SchedMD L3 Slurm support up to 100 Sockets/GPUs additional 1Y EDU&GOV
7S090023WW	SchedMD L3 Slurm support 101-1000 Sockets/GPUs 3Y EDU&GOV
7S090024WW	SchedMD L3 Slurm support 101-1000 Sockets/GPUs 5Y EDU&GOV
7S090026WW	SchedMD L3 Slurm support 101-1000 Sockets/GPUs additional 1Y EDU&GOV
7S090027WW	SchedMD L3 Slurm support 1001-5000+ Sockets/GPUs 3Y EDU&GOV
7S090028WW	SchedMD L3 Slurm support 1001-5000+ Sockets/GPUs 5Y EDU&GOV
7S09002AWW	SchedMD L3 Slurm support 1001-5000+ Sockets/GPUs additional 1Y EDU&GOV

Notes:

- Each 1 GPU counts as 1 socket in quoting support for Slurm
- Request special quotes from Jess Arrington at SchedMD, jess@schedmd.com, to do a manual exception on the appropriate SKU above
- SchedMD Slurm Onsite or Remote 3-day Training: in-depth and comprehensive site-specific technical training. **Can only be added to a support purchase.**
- SchedMD Slurm Consulting w/ Sr. Engineer REMOTE Sessions (Up to 8 hrs): review initial Slurm setup, in-depth technical chats around specific Slurm topics, and review site config for optimization & best practices (**Required with support purchase, cannot be purchased separately**)
- *SchedMD Slurm Consulting w/ Sr. Engineer REMOTE Sessions (Up to 8 hrs)* option **must be selected** and locked in for every SchedMD support selection.
- *SchedMD Slurm Onsite or Remote 3-day Training* option **must be selected** and locked in for every SchedMD Commercial support selection. Optional for Edu & Government support selections.

For Commercial options: Remote consulting session and training required to be quoted with support.

For EDU & Government: Remote consulting session required to be quoted with support, training optional.

SchedMD Slurm L3 support is part of the Lenovo HPC & AI Software Stack. The Lenovo HPC & AI Software Stack is built on the most widely adopted and maintained HPC community software for orchestration and management. It integrates third-party components especially around programming environments and performance optimization to complement and enhance the capabilities, creating the organic umbrella in software and service to add value for our customers. Slurm L3 Support is a key component of the HPC & AI Software Stack and can be added to the BOM through DCSC (Lenovo Data Center Solution Configurator) or through x-config configurator.

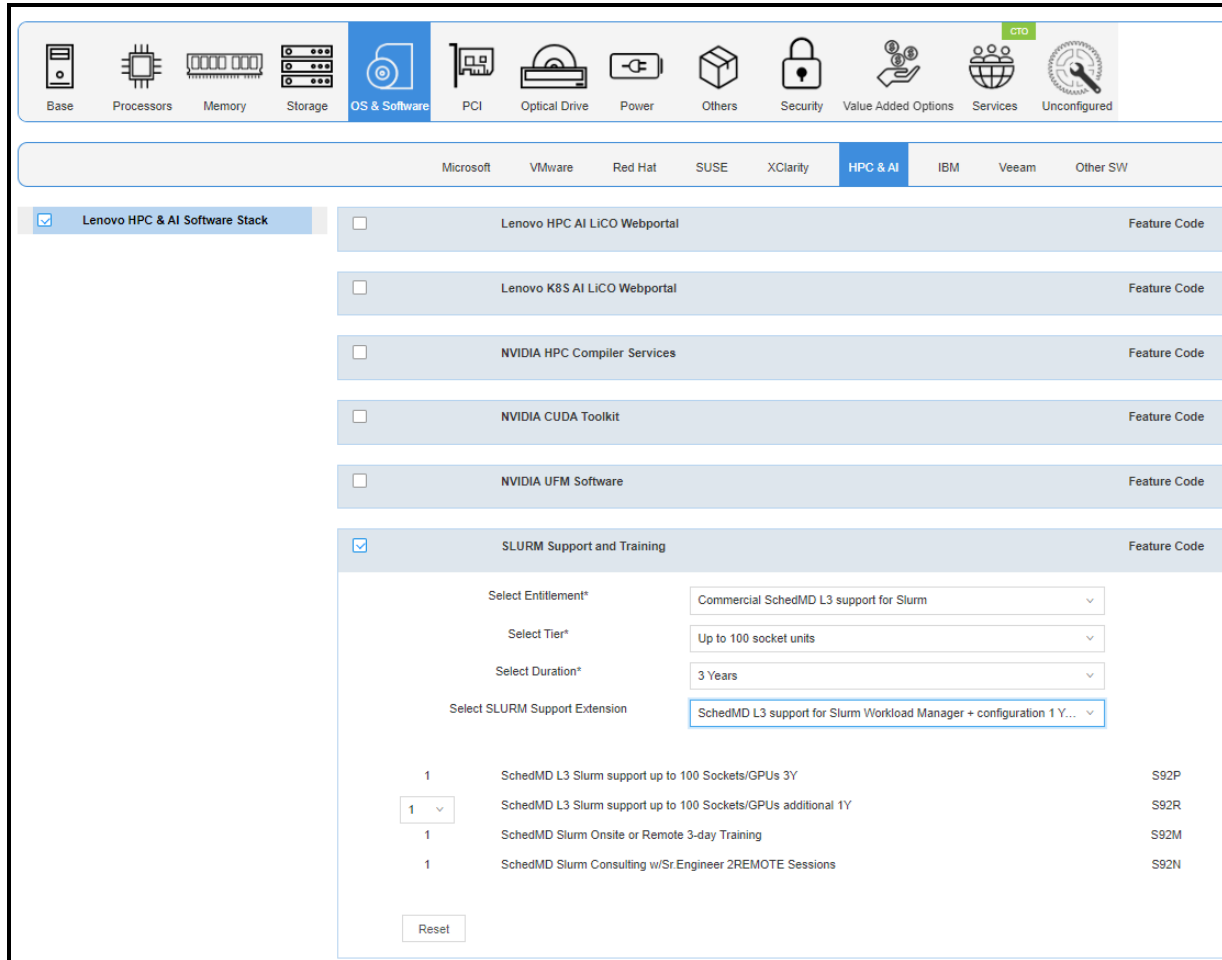


Figure 4. Configuring Slurm Support and Services through Lenovo DCSC

Summary

Slurm Workload Manager is a great tool to address challenges that High-Performance Computing data centers are facing today: automating processes as well as optimizing resource utilization, helping businesses become more productive and reliable.

Together, Lenovo and SchedMD, Lenovo offers a fully integrated, easy-to-use solution, thoroughly tested and supported, so that the customers can focus on their work without the burden of managing the complexity of an HPC/AI cluster.

With respect to robustness and ease of use, the teams at Lenovo and SchedMD are constantly working with customers and feeding the roadmap with their feedback to achieve the fastest time to results for our customers. Stay tuned!

For more information

To learn more about the Lenovo computing orchestration and Slurm, use the following links or contact your Lenovo Sales Representative or Business Partner:

- SchedMD Slurm Commercial Support and Development
<https://www.schedmd.com/>
- Slurm Documentation
<https://slurm.schedmd.com/>
- Lenovo Intelligent Computing Orchestration (LiCO) product guide
<https://lenovopress.lenovo.com/lp0858-lenovo-intelligent-computing-orchestration-lico>
- Lenovo HPC & AI Software Stack product guide
<https://lenovopress.lenovo.com/lp1651-lenovo-hpc-ai-software-stack>
- Lenovo DCSC configurator
<https://dcsc.lenovo.com>

Authors

Ana Irimiea is a Lenovo AI Systems & Solutions Product Manager. She defines marketing requirements and participates in execution through the operations process for each solution release, to lead the next generation of AI and HPC platforms and solutions. She is the product manager for the Lenovo GPU-rich servers and HPC&AI Software Stack. Ana is currently based in Bucharest, Romania and she holds a Bachelor of Systems Engineering degree and a Master in Management in IT from the Polytechnic University of Bucharest.

Victoria Hobson is the Vice President of Marketing at SchedMD. She establishes brand vision and development while executing supporting marketing activities across the company. With a background in software marketing, specializing in partner relationships, Victoria heads the marketing relationship between SchedMD and Lenovo. Victoria is currently based in Greenville, South Carolina and holds a BA in Media Communications with a Minor in International Business from Brigham Young University.

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.

© Copyright Lenovo 2025. All rights reserved.

This document, LP1701, was created or updated on March 3, 2023.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1701>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1701>.

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®

ThinkSystem®

The following terms are trademarks of other companies:

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.