

Lenovo

Optimizing Power and Energy in HPC Data Centers with Energy Aware Runtime

Introduces the Energy Aware Runtime (EAR) software

Describes job accounting and system monitoring with EAR

Includes examples of data gathering, analysis and visualization with EAR tools

Presents EAR strategies for cluster wide Energy Optimization and Power Capping

Julita Corbalan
Luigi Brochard
Karsten Kutzer



Abstract

The increasing power consumption of High Performance Computing (HPC) clusters is a concern because of the high cost of electrical power, sustainability considerations and actual limitations in data center power infrastructures. The Energy Aware Runtime (EAR) is an Open Source software developed in a collaboration with Lenovo®, Barcelona Supercomputing Center and Energy Aware Solutions (EAS). EAR provides system and application power monitoring as well as Energy Optimization and Power Capping capabilities and helps HPC data centers to address those challenges.

This paper is intended for HPC system administrators, application developers and data center managers that want to better understand the power usage of their HPC system, optimize the energy consumption of the applications and limit the power consumption of the HPC cluster in operation.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

<http://lenovopress.com>

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction	3
Overview of the Energy Aware Runtime	4
Energy Aware Runtime deep dive	6
Optimizing the HPC system Energy consumption	14
System level power capping – managing power limits with EAR	16
Service and Support for EAR	18
Conclusions and outlook	20
References	21
Change history	21
Authors	21
Notices	23
Trademarks	24

Introduction

High Performance Computing (HPC) clusters, sometimes called supercomputers, consist of hundreds or thousands of compute nodes. The power consumption of these compute nodes has significantly gone up from generation to generation, driven by trends like increasing Thermal Design Power (TDP) of CPUs, more and higher-capacity memory DIMMs per server, and the introduction of GPUs with extreme power density. In addition, saving energy is very important for addressing concerns on constantly increasing cost of electrical power as well as responsible use of resources.

These trends create new challenges for operating and running a HPC system:

- ▶ Previously, the primary efficiency goal when running a HPC system was optimal use of the capital expense (CAPEX). Hence, the target was to achieve a maximum throughput over the lifetime of the system (“time to solution”). This was achieved by optimized scheduling of jobs on the system for highest system usage on one side and running the compute nodes with highest performance on the other side.

Because the rising power consumption of clusters and the increasing cost of energy, the operational cost (OPEX) of running a HPC system today is at a level, where OPEX becomes very significant compared to CAPEX. Hence, a new additional efficiency goal is to maximize the throughput of the system over energy consumed (“energy to solution”). This can be achieved by running the compute nodes not always at peak power, but at energy-optimized performance levels.

Balancing the two goals (“time to solution” and “energy to solution”) is a complex optimization task, which needs customization for each individual application job.

- ▶ Existing data centers have an electrical infrastructure that was usually designed to cope with power requirements at the time it was built. Often it is very expensive or even impossible to enhance the electrical infrastructure to be compliant with requirements for higher power. This sets a hard limit for the maximum electrical power that is available. As a result, customers often get into a situation where they can no longer run their HPC system at peak power.

This may not be an issue in normal operation, when a mix of different applications and jobs with different characteristics run on the system at the same time - not all compute nodes are constantly consuming maximum power, reducing the overall system power consumption below the maximum value. However in situations where very large jobs put maximum load on a huge number of compute nodes in a synchronized way, this may be an issue.

To prevent an overload of the power infrastructure of the data center, a power control system is needed that understands the jobs running on the system as well as their power consumption and the data center power limits.

Energy Aware Runtime (EAR) is Open Source software developed in a collaboration with Lenovo, Barcelona Supercomputing Center and Energy Aware Solutions (EAS). EAR provides a range of functionality, including system power monitoring and collection of application performance metrics.

EAR also provides Energy optimization and Power capping capabilities with integration into workload managers and job schedulers. EAR currently supports integration with Slurm and Altair PBS professional. EAR is a great solution for HPC data centers suffering from the new challenges described above.

The current version of EAR supports Intel and AMD CPUs as well as NVIDIA GPUs. EAR is constantly being enhanced to support other and upcoming technologies as well.

This paper is structured as follows:

- ▶ “Overview of the Energy Aware Runtime” provides an initial overview on the various features that are provided by EAR.
- ▶ “Energy Aware Runtime deep dive”, an in-depth view is given on how the Energy Aware Runtime works internally and on what the EAR core functionalities are.
- ▶ “Optimizing the HPC system Energy consumption”, it is shown how EAR is optimizing energy consumption at the individual node and at the system level.
- ▶ “Controlling HPC system peak system power with EAR” describes how system level power can be controlled and capped by EAR in a highly optimized and dynamic way.

The following resources provide details on Energy Aware Runtime:

- ▶ Energy Aware Solutions
<https://www.eas4dc.com>
- ▶ Energy Aware Runtime Documentation
<https://www.eas4dc.com/documentation>

The functionality and scope described in this paper is based on EAR V4.1.

Overview of the Energy Aware Runtime

The Energy Aware Runtime solution consists of several components, some running locally on the compute nodes like the EAR library and the EAR daemon, and others running centralized on a management server like the EAR database or the EAR global manager. Combined, those components provide services for controlling power and improving the energy efficiency of an existing HPC Data Center.

As shown in Figure 1 on page 5, EAR provides three main added values:

- ▶ Power and environmental system monitoring and job accounting
- ▶ Transparent runtime application performance and power monitoring
- ▶ Dynamic application / cluster energy optimization and system power control




Figure 1 EAR framework for energy management and optimization in HPC data centers

Power and environmental system monitoring and job accounting

System monitoring: EAR provides node and cluster monitoring. The software will report the power consumption for all compute nodes in the cluster together with some additional metrics such as temperature or average CPU frequency. The local records are reported to a database management system (DB). Additionally, data for multiple nodes is aggregated for accelerating the computation of the total power consumption for medium and big clusters.

Job accounting: Apart from system monitoring, which provides data from a hardware perspective, EAR includes job accounting for all the jobs submitted to the cluster. There are two levels of accounting, Basic and Advanced. Basic job accounting is guaranteed for all the jobs and includes the Job IDs (scheduler JOBID, STEP, ID, User ID, Account, job name, etc), the execution time and the energy consumption per node. Advanced accounting is reported when the EAR runtime is used and includes performance metrics such as CPI, memory bandwidth or GPU utilization and more detailed power metrics such as per-GPU power or DRAM and Package power (apart from DC node power).

Transparent runtime application performance and power monitoring

Application performance and power monitoring is provided by the EAR runtime library and includes additional CPU performance metrics such as cycles per instruction (CPI) and Gflops as well as GPU metrics such as GPU utilization. This runtime solution is dynamic and fully transparent. It does not require user hints, application modifications or recompilation.

The EAR runtime library characterizes the application identifying, measuring, and reporting the main metrics regarding performance and power consumption for a deep understanding of the applications. It allows the administrator to have a better knowledge of the workload, allowing the correlation of performance with power consumption of the system workload.

Moreover, for application developers, it is a powerful tool to analyze real applications with no impact on performance. This information is primarily used by EAR to make decisions for the dynamic application and cluster energy optimization and system power control.

Dynamic application / cluster energy optimization and system power control

The deep application characterization done by the EAR runtime library in terms of power and performance is used by the energy models and energy policies for application energy and power optimization. The dynamic tuning of the energy policies is used by cluster wide policies to implement energy capping at the cluster level.

Apart from energy optimization, EAR implements power capping at the node and cluster level. EAR power capping considers application characteristics and power consumption to guarantee system requirements are met.

The focus of this paper is on those two functions of EAR:

- ▶ Power capping
- ▶ Energy optimization

Energy Aware Runtime deep dive

Energy Aware Runtime (EAR) consists of several components working together to provide monitoring and control capabilities. This section describes concepts of EAR in more detail, and takes a closer look at the individual components and how they work together.

The descriptions in this section are based on the integration of EAR with Slurm as the HPC Cluster Batch Scheduler and Workload Manger. Similar integrations are available for Altair PBS professional.

For details, see the Slurm and Altair web sites:

<https://slurm.schedmd.com/overview.html>

<https://www.altair.com/pbs-professional>

Application performance and power monitoring

The foundation of the EAR capabilities is the ability to monitor the performance of applications which are running on the HPC cluster as jobs controlled by the HPC cluster's Batch Scheduler / Workload Manager. This provides the base functionality for additional features of EAR like energy optimization and power control.

EAR constantly collects metrics like the DC (Direct Current) power consumption of the compute nodes and application performance and stores them in a MYSQL-compatible database (DB). Those metrics are also aggregated to jobs and job steps. This data provides valuable information on how jobs run on the system, how much power and energy they consume, and what their behavior is (e.g., if they are memory or compute bound). Moreover, since EAR already reports data at runtime, this application characterization allows to detect, for example, differences between nodes and/or different application phases.

Figure 2 on page 7 shows the EAR data management architecture and integration with the Slurm job scheduler.

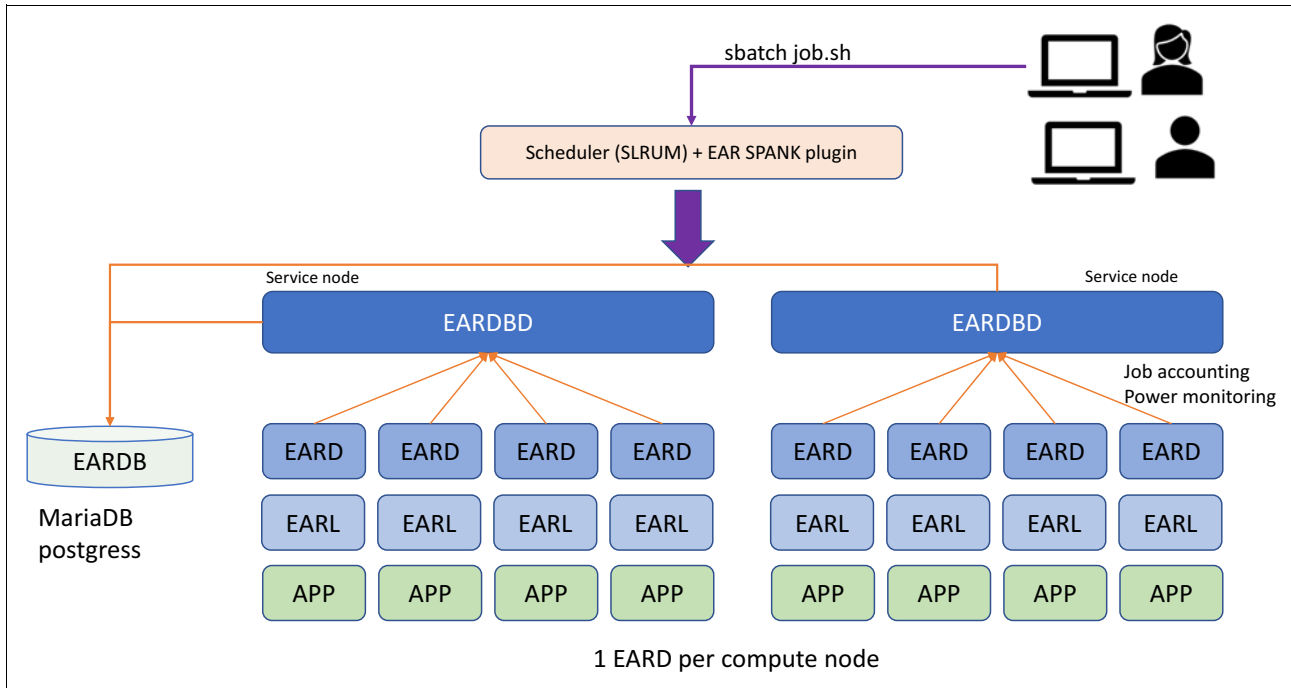


Figure 2 EAR accounting architecture

Components involved in job accounting and node monitoring are the following:

- ▶ EAR runtime library (EARL)
- ▶ EAR DB manager (EARDBD)
- ▶ EAR Daemon (EARD)
- ▶ EAR plugin (EARplug)

EAR offers a highly configurable and extensible infrastructure for energy management since the components are running as Linux services and the energy configuration is set through a global configuration text file (ear.conf).

EARD is a Linux service which resides on each compute node and provides readings of the average DC node power at high frequency. EARD reports DC power from the available source, hiding the HW specific details about how power is measured.

For example, with the Lenovo ThinkSystem™ SD650 V2 and SD650-N V2, EAR takes advantage of the integrated Fast/Accurate power and energy meter, which is using highly accurate sensors providing 100 power or energy readings per second over the in-band interface. The frequency of EARD DC power monitoring is specified in the ear.conf file. EAR also reports RAPL (DRAM and PACKAGE) and GPU power independently.

EAR reports power and performance from two perspectives: Job accounting and node/system monitoring:

- ▶ Job accounting

For jobs running with the EAR library (EARL), performance, power and energy information is collected at the node level and correlated to the Job ID and Step ID. This information is stored in the EAR database in the form of a "signature". The signature is determined at run-time for each application running as a job on the system. The metrics collected are representative of the performance and power behavior of the application during a specific phase.

All the signatures computed are reported to the DB. Loop signatures are the set of signatures for a given Slurm jobID/StepID computed during the execution of the application. Application signatures are the average Loop signatures over the whole execution time of the application. During the optimization process, EAR uses the Loop signatures to determine the best CPU frequency for running the application at optimal energy efficiency during this phase according to the energy policy selected.

For jobs running without the EAR library, EAR reports a simplified version of the signature which does not include the loop granularity (as provided by the EAR library). In this case, the performance metrics reported are limited to the execution time and energy consumption. EAR does also not perform any energy optimization in this scenario (see section “Optimizing the HPC system Energy consumption”).

► Node/system monitoring

A per-node periodic measurement of DC node power and energy is also provided with the specific CPU and GPU power (when available). This information is used to have a global view of the DC power consumption of the nodes at regular intervals even when the node is idle. The EAR DB manager aggregates the information for all the nodes on each sub-cluster on single records periodically and creates aggregated records in the DB. With these aggregated records the computation of the total cluster power consumption for a given period is very fast because the number of DB records involved is reduced by several orders of magnitude.

Energy reporting and accounting commands

The information collected by EAR and stored in the EAR database can be accessed either directly through the DB with MySQL commands, or easily with two EAR command line tools:

- `eacct` to gather job accounting data
- `ereport` for system power monitoring.

These commands query the SQL DB to gather and process the information applying several filters by jobID, username, application name for the `eacct` command and start time, end time or nodename for the `ereport` command.

Energy accounting: `eacct`

The `eacct` command (Energy accounting) reports accounting data for executed jobs and for running jobs. For example, Job 114277 runs the application GROMACS with 3 job steps. Executing `eacct -j 114277` provides output as shown in Figure 3.

```

user@host EAR]$ eacct -j 114277
JOBID  USER APP  POLICY  NODES  FREQ(GHz)  TIME(s)  POWER(Watts)  GBS  CPI  ENERGY (J)
114277-2 user GROMACS ME  16  2.30  377.65  288.65  8.77  0.67  1744128.94
114277-1 user GROMACS MT  16  2.17  384.38  282.12  8.60  0.67  1735095.66
114277-0 user GROMACS MO  16  2.32  371.12  321.21  8.95  0.66  1907317.41

```

Figure 3 EAR accounting command output

The columns shown in the output are:

- `jobid.stepid`
- `username`
- `application name`
- `EAR policy`
- `number of nodes`
- `average CPU frequency`

- ▶ execution time
- ▶ average DC node power
- ▶ gigabytes/sec to memory (GBS)
- ▶ cycles per instructions (CPI)
- ▶ total energy (Accumulated for all the nodes for this JobID and StepID).

The EAR policies are used to control the energy optimization strategy of the jobs. The currently supported policies are

- ▶ ME="minimize energy to solution"
- ▶ MT="minimize time to solution"
- ▶ "MO=monitoring"

Jobs executed without the EARL are shown as No Policy ("NP"). The number of columns depends on the filter applied to the eacct and the architecture. For example, on systems with GPUs, some additional columns with GPU metrics are shown.

Energy reports: ereport

The **ereport** command (Energy report) generates reports from EAR accounting data of the nodes. It analyzes the energy consumption for a given period of time with some additional criteria such as node name or user name. For example, the **ereport** command in Figure 4 provides information on the DC energy consumed and the average DC node power for all nodes since the 30th of January 2021.

```
[user@host EAR]$ ereport -n all -s 2021-01-30
```

Energy (J)	Node	Avg. Power	Avg. GPU Power
163904695	node1	63	25
115036162	node2	49	19
70435928	node3	27	11
126256145	node4	56	34
115111077	node5	51	32
97747686	node6	44	26
3778456	node7	280	11

Figure 4 EAR Energy reporting command output

EAR database schema

EAR uses a relational, MySQL compatible database (e.g. MariaDB) for storing the information. The database tables are structured into different areas:

- ▶ Job-related tables, which provides information on jobs running with the EAR library
 - Per JobID, StepID and node Signature
 - Per JobID, StepID and node and loop Signature
- ▶ Periodic node and aggregated metrics
 - Per-node monitoring metrics: Power (DC node power, RAPL power, GPU power), temperature, Average CPU frequency, Job IDs (JobID, StepID)
 - Aggregated metrics (DC power aggregated by groups of nodes)
- ▶ Global energy/power measurements and warnings (provided by EARGM)
- ▶ Events (provided by EAR library and EARD)

It is possible to have compute/GPU node DC power data stored at any specific period by setting the EAR polling interval in the ear configuration file, for instance every 60 seconds.

EAR aggregates power per pre-defined groups of nodes automatically at aggregated the period intervals defined at ear.conf. Additional aggregation criteria can be done at the SQL level or through the ereport command (see previous section).

Data visualization

As EAR writes system monitoring and application metrics to a relational database, this also allows visualization of the data with 3rd party tools like Grafana (see <https://grafana.com/>). Figure 5 shows the total cluster power meter generated with data from the EAR DB and visualized with Grafana as an example.

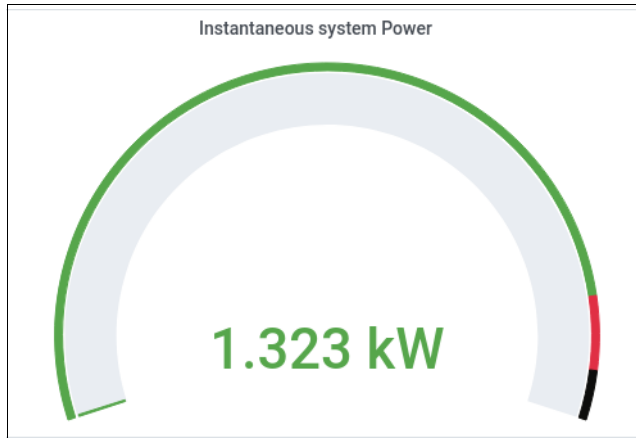


Figure 5 Grafana dashboard showing the EAR reported instantaneous cluster power consumption

Figure 6 shows a Grafana timeline with cluster power consumption over the time.



Figure 6 Grafana dashboard showing EAR reported total cluster power consumption over time

Application metrics can be queried from the DB or reported at runtime by the EARL report plugin and visualized using some visualization tool. Figure 7 on page 11 shows the visualization of some runtime metrics for two executions of GROMACS (GPU version):

- ▶ Per-GPU utilization
- ▶ Per-GPU power
- ▶ Percentage of time consumed in MPI calls




Figure 7 Application metrics runtime visualization

Metrics are displayed with different gradients of colors from yellow (low values) to blue (high values). In this evaluation, we detected some executions were significantly more efficient than others.

The data shown in the graphs in Figure 7 was clearly showing one execution was reporting much more GPU utilization (first vs. second graphs in the figure) and that is also reflected in GPU power (third vs. fourth graphs). We can clearly see how first graph shows more power consumption than second one (green color vs. yellow). These power consumptions reflect the variation measured in GPU utilization.

The third graph clearly shows a high (blue) GPU utilization compared with fourth one.

The last two graphs show the impact of this scenario in the percentage of MPI time. The use case showing high GPU utilization was reporting much less percentage of MPI time (green color) compared with the configuration reporting low GPU utilization. Looking at graphs, what we observed is the improvement in the GPU utilization significantly impacts the CPU computation, making it much more effective and reducing the duration of MPI calls.

EAR data collected can be also converted to other visualization tools such as Paraver (see <https://tools.bsc.es/paraver>). Figure 8 on page 12 shows two application (left and right) metrics (CPI, GBS, power, VPI=%AVX512, CPU frequency) and their variation over time highlighting different phases in the application execution.




Figure 8 Application runtime metrics visualization with Paraver and EAR data

In the case of Paraver, the default gradient of colors goes from yellow/green to dark blue as in Figure 7 on page 11. We can see in the bottom graph, corresponding with the CPU frequency, how EAR has reduced the frequency from the application in the right side (having a high AVX512 utilization) whereas the application in the left side has been executed with high frequency. In this example, the application on the right side is a GROMACS and the one on the left is BT-MZ from the NAS-PB.

Running jobs with EAR

The execution of jobs with EAR is done transparently by scheduler plugins. Support for Slurm is provided through the SPANK plugin (see <https://slurm.schedmd.com/spank.html>), in the case of Altair PBS professional, corresponding hooks are used. The Slurm SPANK plugins are like prolog/epilog and notify the EARD about the new_job/end_jobs events. They are also in charge of defining the environment variables to load the EAR runtime if needed.

Figure 9 shows some of the EAR flags available at submission time when using SLURM.

```
[julitac@ear]$ srun --help
Usage: srun [OPTIONS(0)... [executable(0) [args(0)...]]] [ : [OPTIONS(N)...]] executable(N) [args(N)...]
Options provided by plugins:
  --ear=on|off           Enables/disables Energy Aware Runtime Library
  --ear-policy=type      Selects an energy policy for EAR {type= monitoring,min_energy,min_time}
  --ear-cpufreq=frequency Specifies the start frequency to be used by EAR policy (in KHz)
  --ear-user-db=file     Specifies the file to save the user applications metrics summary
'file.nodename.csv' file will be created per node. If not defined, these files won't be generated.
  --ear-verbose=value   Specifies the level of the verbosity{value=[0..1]}; default is 0
```

Figure 9 EAR flags available when using SLURM

The EAR runtime library is a lightweight solution to analyze applications and to apply optimization strategies. The library is automatically loaded with MPI, OpenMP, MKL, CUDA and python jobs if the scheduler submission tool (srun in the case of Slurm) is used or some tool which interacts with the scheduler like mpirun. If not, the EAR runtime can still be used but it must be loaded explicitly.

For example, for non-parallel applications or parallel applications using different programming models, the name of the application must be specified. For MPI applications using non-compiled approaches, such as python, the MPI version must be specified since it cannot be detected by analyzing the code automatically.

Table 1 Use cases automatically supported by EAR

Use case	Support
MPI (compiled), OpenMP,MKL, CUDA, Python (not MPI)	Yes, 100% automatic using scheduler submission cmds (srun or pbs_attach)
Non parallel	Yes, with explicit request
MPI not compiled (i.e. Python)	Yes, with explicit MPI specification

Even though the EAR software is transparent to users, it also offers some configurable options to advanced users. Those allow additional control, like manually changing options such as the CPU/GPU frequency or executing a job with a specific policy. In the case of Slurm, additional flags are available for the srun/sbatch/salloc commands through a “Slurm Plug-in Architecture for Node and job (K)control” (SPANK) plugin.

One of those flags is --ear-user-db=<file>, which enables the user to specify a file path for a flat-file to store all the job data (including DC node power) at even higher granularity than provided by the EAR database or extra metrics collected by the EARL but not reported because of DB configuration.

Some new options and very specific options for the library are available through environment variables. Environment variables must be used with the other schedulers such as Altair PBS professional which do not support adding new options when submitting jobs. For example, the environment variable SLURM_EAR_GPU_DEF_FREQ is used in Slurm systems to ask for a specific GPU frequency.

The main submission options supported are:

- ▶ CPU frequency
- ▶ GPU frequency
- ▶ Memory frequency limits (max and min)
- ▶ Energy policy

- ▶ Energy policy threshold
- ▶ Disable the library
- ▶ Report flat-text file
- ▶ Load specific report plugin

The last option makes the reporting of data extensible. For example, it allows the user to load a specific report plugin with a specific trace format or to report to a different database.

Integration of EAR with Lenovo intelligent Computing Orchestration (LiCO)

Lenovo Intelligent Computing Orchestration (LiCO) is a software solution that simplifies the use of clustered computing resources for HPC workloads and 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.

Details about the offering can be found in the LiCO Product Guide:

<https://lenovopress.lenovo.com/1p0858-lenovo-intelligent-computing-orchestration>

LiCO exposes EAR deployment options within the standard MPI template, allowing users to take advantage of the capability for MPI workloads.

Once the workload has been profiled through a learning phase, EAR will minimize CPU frequency to reduce energy consumption while maintaining a set threshold of performance. This is particularly helpful where MPI applications may not take significant advantage of higher clock frequencies, so the frequency can be reduced to save energy while maintaining expected performance.

Users can select EAR options at job submission in the standard MPI template, either to run the default set by the administrator, minimum time to solution, or minimum energy. Administrators can set the policies and thresholds for EAR usage within the LiCO Administrator portal, as well as which users are authorized to use EAR.

Optimizing the HPC system Energy consumption

The major benefit of EAR is the optimization of the energy consumption of the HPC cluster. This is done at two levels: first, there is an energy optimization at the individual compute node level, which is provided by the EAR library. Second, there is the system level energy optimization, which makes sure the system stays within a certain pre-defined energy limit over time – this is called “system level energy capping”.

Node level energy optimization

Figure 10 on page 15 shows the iterative process for energy optimization in the EAR library. There are four main stages applied at runtime to continuously monitor and optimize the applications: The Dynamic application monitoring, the computation of the Loops signature (performance and power metrics), the Signature classification and the core of the optimization when energy models and energy policies are applied.

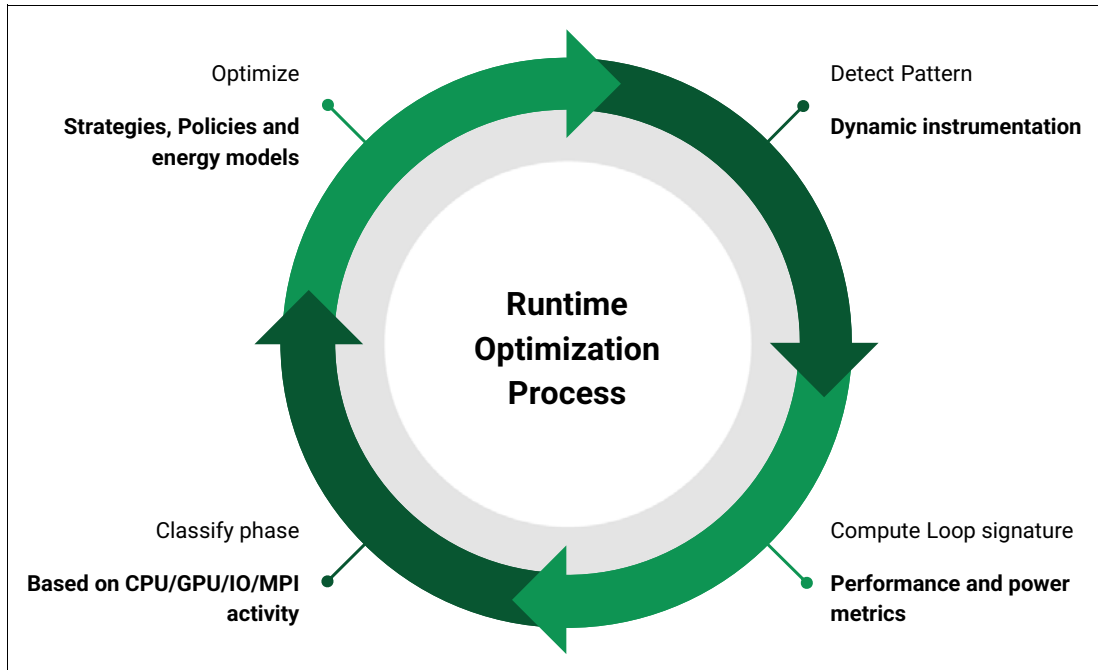


Figure 10 EAR Runtime optimization strategy

EAR provides energy optimization through the selection of an energy policy: minimize energy to solution (ME) or minimize time to solution (MT). This feature is provided without the requirement for any source code modification or user input. To optimize job and system energy usage, EAR dynamically and periodically adjusts the CPU, memory, and GPU frequencies to optimize the energy consumed by the job. For example, a memory-bound job may not suffer a lot from reduced CPU frequency, hence saving energy with a low impact on performance; Energy optimization has been tested on a large selection of applications and has shown average energy savings on Intel Xeon CPU architecture of 11%.

The latest version of EAR supports the simultaneous execution of multiple jobs in the same node. EAR uses a CPU power model to distribute the power consumption between the different jobs sharing the node.

System level energy capping

System (or cluster) level energy capping is an option to maintain the energy consumption of the cluster within a predefined upper bound over a given period. Energy capping is therefore a way to make sure that the average power consumption of the system stays within a predefined power envelope and is designed to be applied at a long-term scale. This allows to compute the cluster energy consumption using the aggregated power data in the database, instead of using live data from the distributed EAR infrastructure. This design significantly reduces the network traffic and cost of the computation of the total energy since only a few database records are involved.

The system level energy capping feature is provided by the EAR Global Manager (EARGM). The EARGM periodically controls the cluster energy consumption for a defined period of time and compares it with the defined limits (both are set in the EAR global configuration file `ear.conf`). EARGM automatically adapts system settings in coordination with the EAR library (EARL) and the EAR daemon (EARD). Since EARL is aware of application characteristics, it can react to the different EARGM warning levels based on application characteristics and the energy efficiency measured. EARGM is a lightweight solution for global control since it

centralizes the problem detection, while the specific actions to control the energy are distributed.

Figure 11 shows the main flow of data and actions for energy capping implementation.

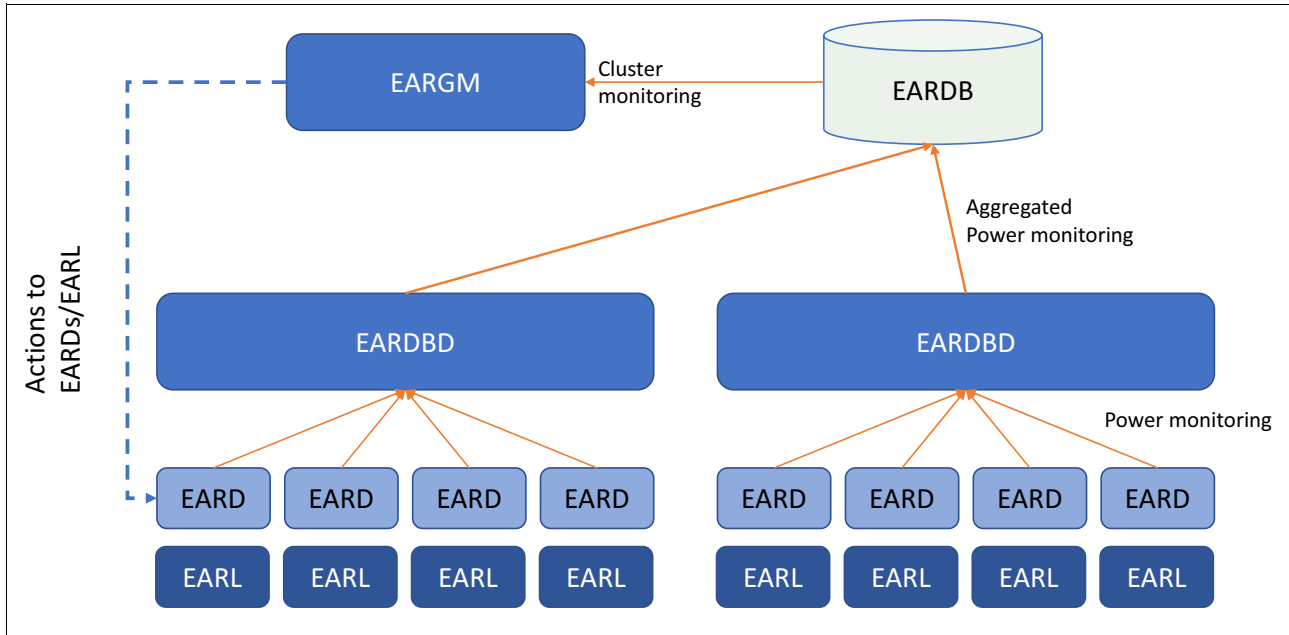


Figure 11 EAR energycap architecture for HPC system level energy capping

System level power capping – managing power limits with EAR

Based on the capability to manage the power of the individual systems, EAR also provides the ability for system level power capping. This is done in a very intelligent, optimized, and dynamic way, considering the user jobs that are currently running on the system and their individual energy efficiency at different levels of power consumption.

The difference between system level energy capping (as discussed in the previous section) and system level power capping is the reaction time:

- ▶ For energy capping, the reaction time can be slow, in terms of hours. Hence, EAR must not detect the event immediately. For reaction, EAR can dynamically modify general energy optimization parameters and for slowly trending the energy consumption towards the target again.
- ▶ For power capping, the reaction time must be very fast, in terms of sub-second. Hence, hard power limits must be set on the individual compute nodes to guarantee, that the power consumption of the system does not exceed the set power limit of the system. EAR is managing and optimizing those power limits at a higher level.

System (or cluster) level power capping is an option to maintain the instantaneous power consumption of the cluster within a predefined upper bound at any given time. Power capping is provided by the EAR Global Manager (EARGM) which monitors and controls the power consumed in the system. A different architecture than what is used for system level energy capping is needed to guarantee the system responsiveness and the scalability. The system level power capping is provided by the collaboration of EARGM and the EAR daemons (EARDs) running on the compute nodes.

EARGM is responsible for power reallocation in the cluster, bringing the intelligence to the power capping solution. EARGM periodically collects the cluster power status. The cluster power status includes the list of compute nodes which require more power and compute nodes that can don't need their budget and can share some power. This information is provided by EARDs. In case of large clusters, a hierarchy of EARGMs can be used, with different power limits for each EARGM if needed.

EARDs bring the low level and fast capabilities to guarantee the node powercap control, which are the pre-requisite for system level power capping. Different powercap limits per CPU type can be specified at node level. The EARDs apply power limits by using the available underlying technology, for example DVFS for the CPU domain and NVML capabilities for power management of NVIDIA GPUs.

The EARD powercap algorithm also balances the power allocated to the CPU and GPU within the system, taking into account the requested CPU frequency and the GPU utilization. The requested CPU frequency is a hint provided by the EAR library. In those cases where the hardware does not provide a power limit capability, EARDs uses a pro-active mechanism by actively checking the power consumption and forcing the power limit by reducing the CPU or the GPU frequency, implementing the power limits through DVFS.

Three possible powercap configurations are supported:

- ▶ Node power cap only: All the compute nodes receive a static power allocation. EARDs are responsible for guaranteeing the power cap. In that case, the EARGM is not involved.
- ▶ Cluster [SOFT] power cap: This option guarantees the cluster power limit is not exceeded for more than the time specified by the cluster power cap monitoring period. Compute nodes run with the powercap disabled and the EARGM monitors the cluster power consumption. If the cluster power consumption exceeds the limit, the node power cap is enabled. EARGM enable/disable node power cap depending on the cluster status. This option reduces the potential node overhead to guarantee the powercap when it's not really needed.
- ▶ Cluster [HARD] power cap: This option guarantees the cluster power limit is never exceeded. Compute nodes run with the power cap always enabled and the EARGM re-allocates power between compute nodes from nodes with fewer requirements to highly power demanding nodes.

Figure 12 on page 18 shows EAR design for energy and power capping. In this example, there are 3 EARGMs running in the system, each one taking care of N nodes. The EARGM taking care of the top level in the EARGM hierarchy is called the Meta-EARGM.

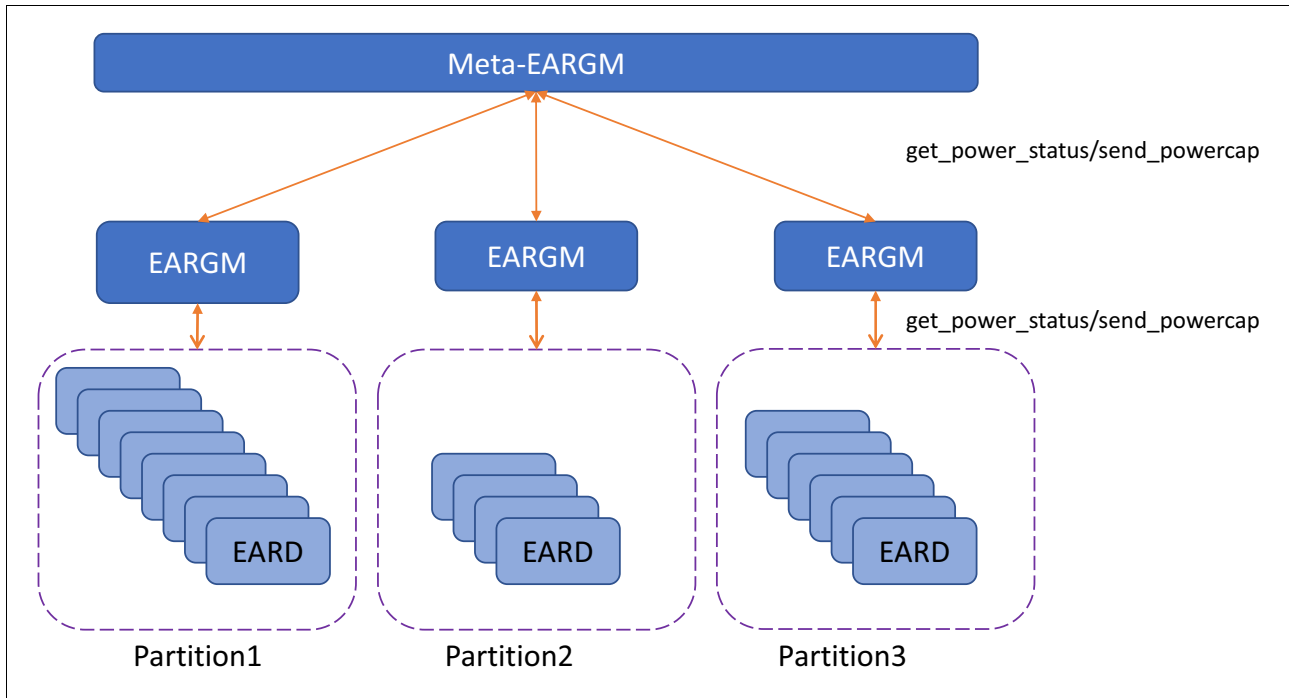


Figure 12 EAR powercap hierarchical architecture for HPC data centers

For power capping, EARGM provides an initial upper power bound per node. This upper bound is derived from the configuration file, either from the per-node power limit, or from the cluster power cap, which is homogeneously distributed across the compute nodes. Then the power limit is periodically re-distributed to optimize the system power budget while still minimizing the application performance degradation in coordination with EARL and EARD. Each component is responsible for implementing the power cap API, to implement power balance (if needed), and to guarantee the allocated node power budget is not exceeded.

The power cap API has two main functions:

- ▶ To get the power cap status
- ▶ To receive the power limits

The power cap status includes hints about power requirements. On heterogeneous clusters, EAR implements the power balance between the CPU and the GPU devices.

On top of optimizing the power and energy usage, EAR can also help improve the overall system performance and reliability. Given EARD is continuously monitoring node metrics, hardware reliability tests are periodically done. When a HW problem is detected, a notification is sent to syslogs and DB. These events could be automatically triggered by the system for an immediate reaction and/or analyzed.

More details in the Wiki on GitHub:

<https://github.com/BarcelonaSupercomputingCenter/ear>

Service and Support for EAR

The Energy Aware Runtime is Open Source under BSD-3 license and EPL-1.0. For professional use cases in production environments, installation and support service are available.

Commercial support as well as implementation services for EAR can be purchased from Lenovo and is delivered through Energy Aware Solutions (EAS). There are three different distributions of EAR: Detective Pro, Optimizer and Optimizer Pro. Detective Pro provides the basic monitoring and accounting capabilities, Optimizer adds the energy optimization and Optimizer Pro the power capping features as shown in Figure 13.

	Detective Pro	Optimizer	Optimizer Pro
System monitoring & reporting	✓	✓	✓
Basic Job Accounting	✓	✓	✓
Advanced Job Accounting	✓	✓	✓
Energy Optimization		✓	✓
Power Capping			✓

Figure 13 Comparison of EAR Detective Pro, Optimizer and Optimizer Pro features

Table 2 lists the Lenovo product numbers can be used to purchase EAS support and implementations services.

Table 2 Lenovo product numbers for Energy Aware Runtime

Part number	Description
7S09001KWW	EAR Energy Detective Pro Worldwide Remote Installation and Training for AMD or Intel CPUs
7S09002XWW	EAR Energy Detective Pro 3-years Worldwide Remote support for AMD or Intel CPUs (flat fee)
7S09002YWW	EAR Energy Detective Pro 5-years Worldwide Remote support for AMD or Intel CPUs (flat fee)
7S09001LWW	EAR Energy Detective Pro 1-year Worldwide Remote support for AMD or Intel CPUs (flat fee)
7S09002WWW	EAR Energy Detective Pro Worldwide Remote Installation and Training for AMD or Intel CPUs + NVIDIA GPUs
7S09002ZWW	EAR Energy Detective Pro 1-year Worldwide Remote support for AMD or Intel CPUs + NVIDIA GPUs (flat fee)
7S090030WW	EAR Energy Detective Pro 3-year Worldwide Remote support for AMD or Intel CPUs + NVIDIA GPUs (flat fee)
7S090031WW	EAR Energy Detective Pro 5-year Worldwide Remote support for AMD or Intel CPUs + NVIDIA GPUs (flat fee)
7S090032WW	EAR Energy Optimizer 1-year Support Entitlement for Energy Monitoring and Optimization per system power rating
7S090033WW	EAR Energy Optimizer 3-year Support Entitlement for Energy Monitoring and Optimization per system power rating
7S090034WW	EAR Energy Optimizer 5-year Support Entitlement for Energy Monitoring and Optimization per system power rating
7S09001JWW	EAR Energy Optimizer Pro 1-year Support Entitlement for Energy Monitoring, Optimization and Power Capping per system power rating

Part number	Description
7S090037WW	EAR Energy Optimizer Pro 3-years Support Entitlement for Energy Monitoring, Optimization and Power Capping per system power rating
7S090038WW	EAR Energy Optimizer Pro 5-years Support Entitlement for Energy Monitoring, Optimization and Power Capping per system power rating
7S090035WW	EAR Energy Optimizer Worldwide Remote Installation and Training for AMD or Intel CPUs
7S090036WW	EAR Energy Optimizer Worldwide Remote Installation and Training for AMD or Intel CPUs + NVIDIA GPUs
7S09001GWW	EAR Energy Optimizer Pro Worldwide Remote Installation and Training for AMD or Intel CPUs
7S09001HWW	EAR Energy Optimizer Pro Worldwide Remote Installation and Training for AMD or Intel CPUs + NVIDIA GPUs

The number of required licenses depends on an assumed power rating of the cluster. For details, get back to your Lenovo sales representative.

Conclusions and outlook

The Energy Aware Runtime is a great tool to address the challenges that High Performance Computing data centers are facing today: optimizing the energy consumption as well as controlling the data center power budget.

Since EAR development started in 2016, EAR was improved along four dimensions.

- ▶ First dimension is the type of processor/platform EAR supports, like adding support for AMD CPU and NVIDIA GPU.
- ▶ Second dimension is the type of applications, like extending the support from MPI only applications to nearly all type of workloads including AI applications with Python and Tensorflow/PyTorch.
- ▶ Third dimension is EAR robustness and ease of use running on production systems.
- ▶ Fourth dimension is working in collaboration with our customers to understand how they use EAR and what they'd like to do which is not possible today.

EAR will continue to expand along these four dimensions.

Regarding the type of processor, support for ARM processors is under evaluation. Extending EAR support to other types of GPUs is also important: Intel GPU support is on the EAR roadmap; AMD GPU support is being evaluated and NVIDIA GPU support continues and is planned to be extended to latest generations.

Looking at the type of applications supported, we are evaluating EAR support for containers and workflows.

With respect to robustness and ease of use, the teams at Lenovo and EAS are constantly working with customers and feed the EAR roadmap with their feedback.

Stay tuned!

References

- ▶ Energy Aware Solutions
<https://www.eas4dc.com>
- ▶ Energy Aware Runtime Documentation
<https://www.eas4dc.com/documentation>
- ▶ J. Corbalan, O. Vidal, L. Alonso and J. Aneas, "Explicit uncore frequency scaling for energy optimisation policies with EAR in Intel architectures," 2021 IEEE International Conference on Cluster Computing (CLUSTER), 2021, pp. 572-581, doi: 10.1109/Cluster48925.2021.00089.
- ▶ J. Corbalan, L. Alonso, J. Aneas and L. Brochard, "Energy Optimization and Analysis with EAR," 2020 IEEE International Conference on Cluster Computing (CLUSTER), 2020, pp. 464-472, doi: 10.1109/CLUSTER49012.2020.00067.
- ▶ Slurm web page
<https://slurm.schedmd.com/overview.html>
- ▶ Slurm SPANK plugin documentation
<https://slurm.schedmd.com/spank.html>
- ▶ Altair PBS professional
<https://www.altair.com/pbs-professional>
- ▶ Lenovo ThinkSystem SD650 V2
<https://lenovopress.lenovo.com/lp1395-thinksystem-sd650-v2-server>
- ▶ Lenovo ThinkSystem SD650-N V2
<https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>
- ▶ Grafana Labs
<https://grafana.com/>
- ▶ Paraver visualization tool
<https://tools.bsc.es/paraver>
- ▶ Lenovo Intelligent Computing Orchestration (LiCO)
<https://lenovopress.lenovo.com/lp0858-lenovo-intelligent-computing-orchestration>

Change history

December 13, 2022

- ▶ Added additional Energy Aware Runtime part numbers to Table 2 on page 19

First published: November 2, 2022

Authors

Julita Corbalán is the team leader of the System software for energy management in HPC group at Barcelona Supercomputing Center (BSC) and the co-founder of Energy Aware Solutions (EAS) – the company that drives the development of the Energy Aware Runtime

(EAR) with Lenovo and provides installation services and technical support. Her research interests include processor management of parallel applications, parallel runtimes, scheduling policies and energy efficiency solutions for data centers. She received the engineering degree in computer science in 1996 and the PhD degree in computer science in 2002, both from the Technical University of Catalunya (UPC), Spain.

Luigi Brochard started his career at IBM as HPC architect and became IBM Distinguished Engineer in 2004. In 2011, he created the Energy Aware Scheduling technology and after moving to Lenovo in 2015 he started the collaboration with Barcelona Supercomputing Center (BSC) to develop Energy Aware Run time (EAR). He retired from Lenovo in 2018, co-authored the book “Energy-Efficient Computing and Data Centers” (published by Wiley in 2019) and co-founded Energy Aware Solutions (EAS) in 2020. Luigi holds a Ph.D. in Applied Mathematics and an HDR in Computer Science from Pierre et Marie Curie University in Paris, France.

Karsten Kutzer is a Principal Technical Consultant at Lenovo, working as an architect in High Performance Computing (HPC). He has been working in HPC projects for more than 20 years, starting in IBM as a service specialist implementing supercomputers and then moving into a pre-sales architect position. Since he moved to Lenovo in 2015, he is leading the design and implementation of huge HPC clusters. He has been working with EAS on the first customer implementation of the Energy Aware Runtime. Karsten holds a Bachelor in Technical Computer Sciences from Berufsakademie Mannheim, Germany.

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.
1009 Think Place - Building One
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 was created or updated on December 9, 2022.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/lp1646>

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available from <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

Lenovo(logo)®

ThinkSystem™

The following terms are trademarks of other companies:

Intel, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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