

EAR requirements

NOTE: This document includes minimum tested versions for the different requirements. However, this does not imply that EAR will not work with other versions.

HW requirements

- CPUs supported: Intel architectures. Haswell and Skylakes.
- GPUS: Nvidia V100 and T4.
- Platform: Full support with Intel Node Manager compatible systems. Partial support based on RAPL for non INM compatible systems.

SW requirements for EAR compilation

- Lib PAPI. A version compatible with the architecture must be installed. Minimum tested version 5.4.0
- Lib GSL. Minimum tested version 1.6
- SLURM. For automatic job submission. Minimum tested version 17.02.6. Sources are needed for EAR compilation
- MPI installation. Intel and OpenMPI. Minimum tested version for Intel 2017 and for OpenMPI 4.0 (with SLURM support)
- MariaDB or PostgreSQL server installed for database support. Minimum tested version for MariaDB is 5.5. Include files are needed to compile EAR (developer rpm version must be installed)
- Autoconf. Minimum tested version 2.69
- OS: EAR has been tested in CentOS and SUSE

Additional SW requirements

- Grafana server installed for EAR monitoring visualization graphs
- Paraver tool if Paraver traces must be visualized

Kernel drivers required in compute nodes

- CPUFreq and IPMI drivers installed

Additional requirements

- Root permission (sudo) is needed during the pre-test and installation since some of the tests need to be executed through a privileged account, and eard runs as root.
- Specific ear user id is recommended (not mandatory). Only a system user account is needed, not a normal user.

Pots and connectivity

- Ports:
 - 1 TCP port for EARD on the compute nodes.
 - 3 TCP ports for EARDBD on the “service nodes” (depends on the cluster size)
 - 1 TCP port for EARMGD on the “scheduler node”
 - 1 TCP port for DB server

- All these ports are defined in ear.conf. DB port is also specified in ear.conf.
- Default ports are: 50000, 50001, 50002 and 50004
- Grafana server port is not specified in ear.conf since it runs totally independently of EAR (The same server could be used for other purposes, it's not EAR specific)
- Networks
 - EARDs must be able to connect with EARDBDs
 - EARDBDs must be able to connect with DB server
 - Login nodes must be able to connect with DB server
 - Login nodes must be able to connect with compute nodes for econtrol command
 - Grafana server must be able to connect with DB server

EAR Deployment.

- Terminology
 - Login nodes: Nodes where jobs are submitted and ear commands executed
 - Compute nodes: Computational nodes where EARDs are running
 - Service node: Intermediate nodes where specific services are executed, mainly in big clusters
 - Scheduler nodes: Nodes where typically centralized services are executed.

Login nodes deployment

- EAR commands
- EAR slurm plugin
- Special folders and files: EAR_TMP(private) & EAR_ETC paths, and ear.conf
- DB server connection for EAR commands
- Compute nodes connection for econtrol command

Compute nodes deployment

- EARL and EARD must be installed
- EAR slurm plugin
- Special folders and files: EAR_TMP (private) & EAR_ETC paths, and ear.conf and coefficients
- MPI, Libpapi.so and Libmysql (or libpq) must be installed
- CPUFreq and IPMI drivers must be running
- Service nodes connection is needed

Services nodes deployment

- EARDBD must be installed
- Special folders and files: EAR_TMP(private) & EAR_ETC paths, and ear.conf
- Libmysql or libpq must be installed
- Scheduler nodes (were DB server is running) connection is needed

Scheduler nodes deployment

- EARGMD must be installed
- Special folders and files: EAR_TMP(private) & EAR_ETC paths, and ear.conf
- Libmysql or libpq must be installed
- DB server connection is needed

Additional considerations concerning services

- EAR can use the same DB server than SLURM or not, that depends on the customer. Using a single DB server simplifies DB administration but it's not a requirement.
- MariaDB, Grafana, and EARGM don't need to be running on the same node than SLURM, it's not a requirement. It's a customer decision. They are executed on the scheduler node, but it can be different than the slurmctl node.
- The interaction between EAR and SLURM is through the SLURM SPANK plugin, nothing else.

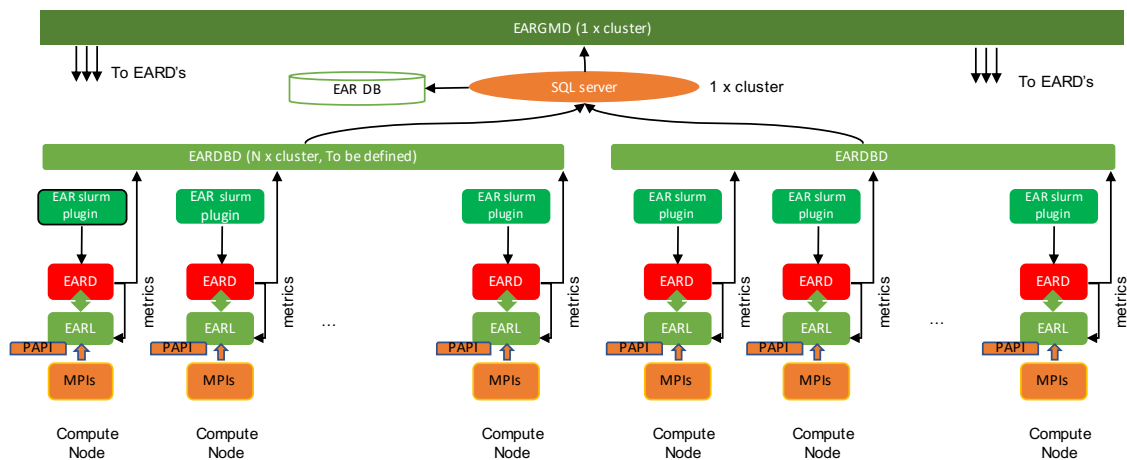


Figure 1: EAR overview