

EVIDEN

Eviden HPC Software Suites

Embracing digital sustainability

High Performance Computing (HPC) workloads have evolved from being “computing-centric” to “data-centric.”

The advent of AI and HPC in the cloud introduces new application paradigms and presents new challenges for both the HPC hardware and software layers. Keeping performance as the cornerstone, the Eviden HPC Software Suites take a holistic approach to maximizing infrastructure efficiency, enhancing scalability and flexibility for a reduced carbon footprint. The convergence of AI and new security mechanisms creates a new level of user experience with reduced TCO. This paper explains the main features and functionalities of the Eviden HPC Software Suites and shares our vision for the significant evolutions coming next in the HPC software arena.

The market and new challenges



From simulation and modeling toward hybrid and intensified workload

HPC continues to be an instrumental component in driving progress in a wide range of scientific and engineering applications around the world. In addition to the race toward Exascale with increasingly powerful processors, the nature of the workload is also evolving — from conventional simulation and modeling to a hybrid workload, including on-premises and cloud applications. There is also a need to integrate, assimilate, and analyze the data from a myriad of IoT sensors. At the same time, the democratization of the AI workload on an HPC infrastructure, the use of AI to enhance workloads, and system availability (maintenance, security, etc.) will soon be mainstream.

A heterogeneous workload brings diverse challenges in the HPC Software Suites

On one hand, next generation HPC must accommodate the complexity of new hybrid workloads with more energy hungry blades, while embracing the advantages that AI offers. On the other hand, more and more organizations worldwide are embarking on the journey toward digital sustainability. In this context, many questions remain:

- How can you craft an efficient and smart software environment to satisfy both needs?
- What new HPC software developments can effectively contribute to reducing power consumption?
- How can HPC software support the evolution of “x-as-a-service” without generating complexity?
- How can HPC software that is designed to optimize job execution avoid wasting resources on unproductive workloads that will not run to completion (failing, wrong result, etc.)?

The Eviden HPC Software Suites: Embracing digital sustainability with a holistic approach



Performance remains key

At the heart of HPC, performance remains the main differentiator. The Eviden HPC Software Suites use containerized or bare-metal jobs to run workloads in order to maximize hardware utilization. The parallel runtime is also critical for maximizing application performance, providing specific tuning algorithms for large-scale applications as well as for collective communications for hybrid applications like MPI/OpenMP. By adopting lightweight profiling tools, clients can get insights on all jobs and information on side effects that may have generated noise or jitter during job execution. With an interconnect management module component, the software also provides an extensive set of hardware, management and monitoring capabilities for maximizing uptime, quick failure handling and optimizing fabric utilization. Maximizing availability and optimizing performance are paramount — not only for extracting value for the workload, but also for reducing down-time and the carbon footprint.



AI converged and AI enhanced

AI workloads are more and more common on HPC infrastructures, taking advantage of the high computation power to train models or for other AI operations. The conjunction of CPU and GPU is also of great benefit, mixing AI and standard HPC workloads. For Eviden, AI is embedded in different layers of the Eviden HPC Software Suites for I/O optimization, predictive maintenance, prescriptive security, runtime settings and workload optimization. Thanks to AI, customers can optimize workload and performance — leveraging historical executions with reduced time-to-result, while optimizing resource usage.



Infrastructure efficiency

To ensure effortless cluster administration, the Eviden Smart Management Center Suite is designed to maximize the efficiency of administrative tasks depending on the scale of the cluster. From a simple, efficient environment for entry level clusters to a fully orchestrated services environment on tier-0 clusters, the Eviden Smart Management Center maximizes availability based on best-in-class technologies. Furthermore, with our Smart Energy Management Suite, customers can monitor, control and dramatically improve their cluster energy efficiency.



Security

The democratization of HPC-as-a-Service (cloud or on-premises model) and the necessity for strong data protection require even higher security features than are found in today's HPC systems. With recent attacks on HPC infrastructure and GDPR regulations introduced in 2020, "security" is not an option, but a legal requirement. As Europe's top cybersecurity firm and number three worldwide for several years, Eviden is working closely with the HPC development team to ensure all measures are taken to build secure, best-in-class implementations. With SELinux support, SSO, trusted certificates and encryption for all critical data exchanges, the management environment is a powerful event and logging solution which allows the security framework to plug-in and perform any advanced/AI security rules with those data. Thanks to these measures, customers can maximize HPC availability and minimize downtime, conserving energy otherwise wasted in managing security incidents and cyberattacks.

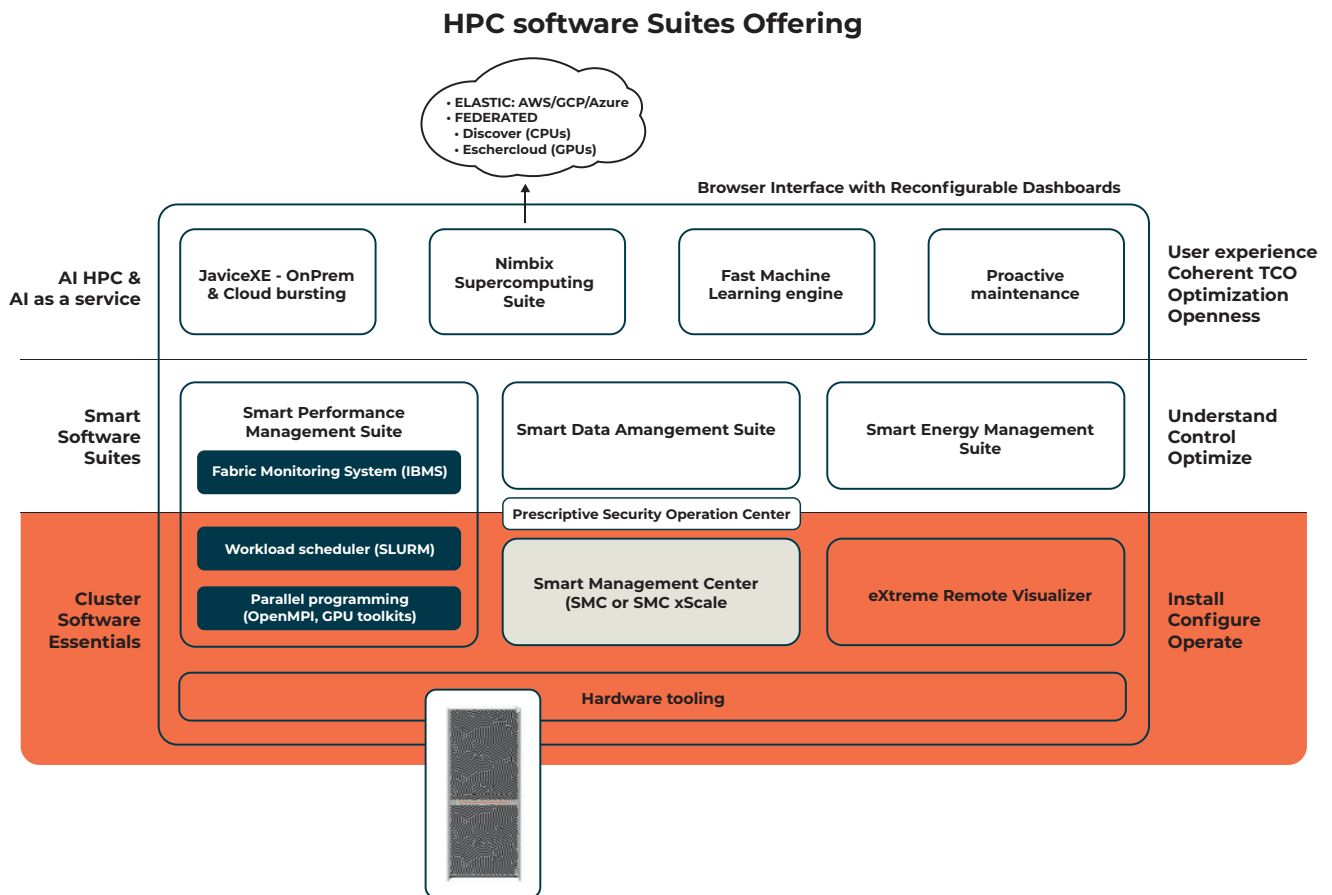


Scalability and flexibility

"Build to last" has been one of the primary approaches adopted to ensure digital sustainability. The Eviden HPC Software Suites have taken a similar approach. The Eviden HPC Software Suites work with all types of HPC, including entry-level, medium range and the Top 10 supercomputers with a lifespan of ten years. It is Exascale-ready, taking into consideration the application evolution towards containerization. The solution is designed to use a reduced set of dedicated management servers while delivering the maximum availability — regardless of the size of the supercomputer. Maintenance operations are mainly performed live, without any impact on the running applications.

The Eviden HPC Software Suites: Key benefits and features

The Eviden HPC Software Suites are composed of 3 stacks as shown below



Layer A: HPC Smart Management Center used to install, configure and operate

Openness: Standard software components embrace a broad ecosystem with large communities. The software is production-ready and supported by Eviden professional services.

High availability: Maximize uptime and minimize downtime even during maintenance operations, by leveraging microservice orchestration and shared storage.

Scalability: Easy to extend regardless of the cluster size. Thanks to centralized management, it is very easy to manage a unique system whatever its size (a 10k-node or 100-node system), eliminating complexity related to cluster extension.

Efficiency: Optimize the overall energy consumption, avoiding any waste of resources on the infrastructure.

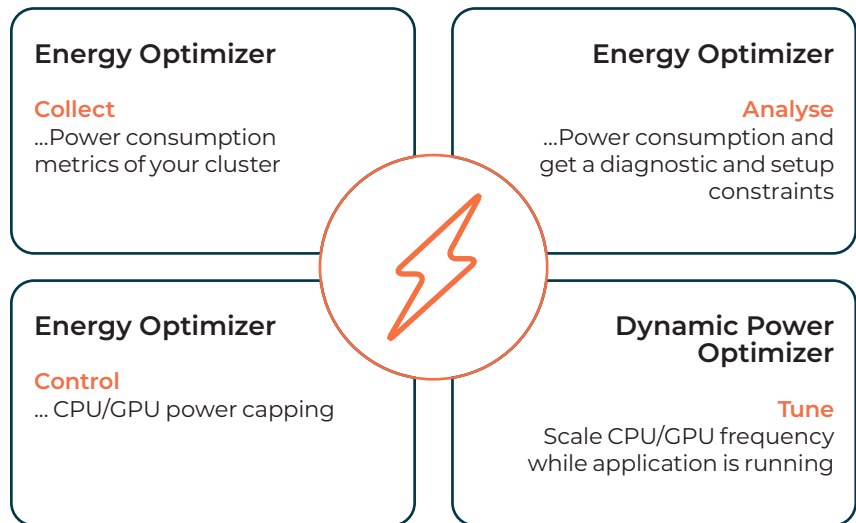
Security: Capable of gathering all events (logs, metrics, etc.) from the supercomputer with near real-time analysis. Ensures detection of any unwanted behavior, enabling clients to take appropriate countermeasures to keep the system running safely.

Layer B: HPC-focused applications to manage energy, data and HPC cluster performance

Smart Energy Management Suite:

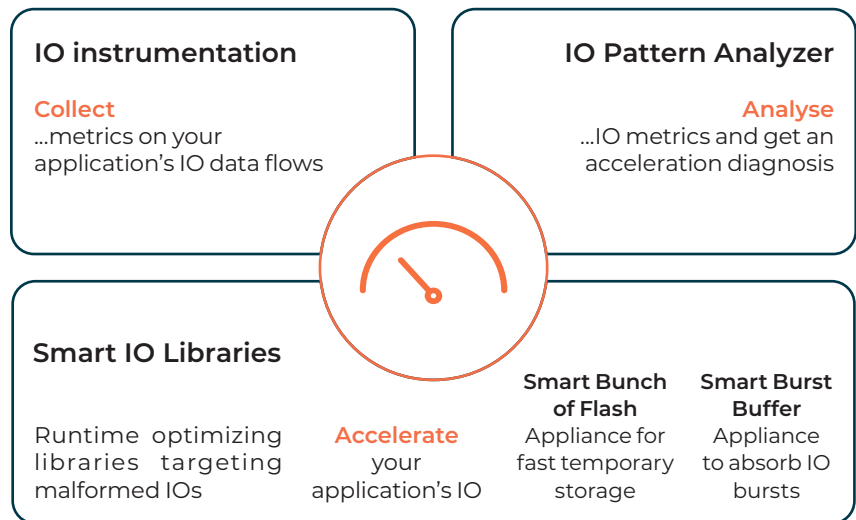
This module is dedicated to understanding, controlling and optimizing energy consumption. It is composed of two components:

1. An Energy Optimizer that gathers energy and temperature metrics from all components of the supercomputer and links them with the job execution data to provide per-job information.
2. A Dynamic Power Optimizer that controls energy consumption through a power capping policy, keeping the system under pre-defined constraints. It also enables application behavior profiling in order to minimize its energy footprint, without any application modification.



Smart Data Management Suite:

Understanding and optimizing I/O patterns is a very complex task. This module brings a complete framework to gather data from jobs about I/O calls (IO instrumentation), represent I/O exchanges and detect abnormal or sub-optimal patterns (IO pattern analyzer). Based on this information, optimization can be performed directly on application code or managed through a smart IO optimizer that dynamically changes malformed I/O thanks to machine learning algorithms. This approach can boost applications up to 30%, a significant gain in both performance and energy efficiency. This solution can also be enhanced by a hardware and software solution named Smart Burst Buffer/ Bunch of Flashes to speed up I/O calls, thanks to an additional level of storage caching.



Smart Performance Management Suite:

This module brings together all the required components to ensure maximum performance of the supercomputer job executions.

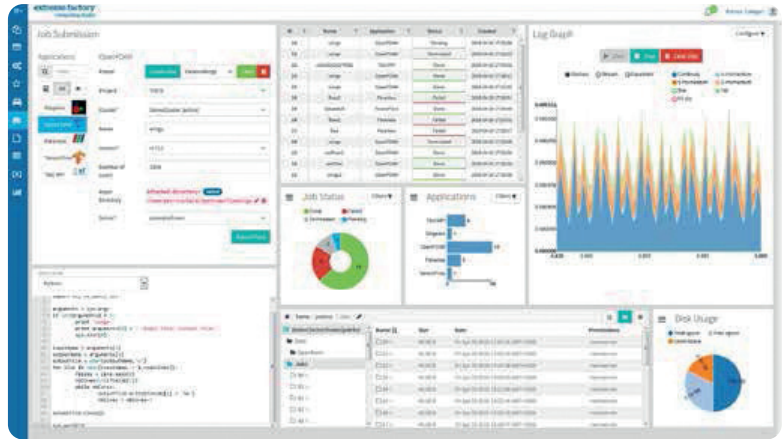
- The Slurm resource and batch manager, with additional Eviden-developed tooling.
- The latest version of Open MPI, with pre-defined profiles and embedded accelerators for collectives, hybrid communication, etc.
- A first-level MPI and OpenMP profiler that can be enabled for all applications without performance impact.

- An optional application container solution based on Sylabs' Singularity.

All these components are designed to maximize the performance and availability of the overall solution, minimizing the downtime and reducing energy consumption by optimizing resource usage.

eXtreme Remote Visualizer:

The HPC, AI & quantum computing portal module is the cornerstone for “x-as-a service” applications, enabling scientists to get an easy-to-use abstraction for job submission, tracking and result gathering. It supports most job schedulers, ready to expose most HPC and viewing applications. The Remote Visualizer uses HPC cluster resources to run graphical applications on the supercomputer resources, with a view similar to that shown on a local desktop/laptop. It avoids data transfer, consolidates graphical resources and ensures security by keeping the data in the cluster.



Layer C: Transverse applications and broader business applications

On top of layers A & B, clients can select relevant software “adds-on” that leverages AI to manage data assimilation and security challenges — further enhancing the user experience, optimizing TCO and promoting openness. These add-ons provide different features, including:

- The Eviden ThinkAI Suite, which provides a toolset and environment necessary to scope, develop, roll-out and manage AI applications — especially ML/DL.
- The Eviden Proactive Maintenance, a framework that enables detection, prescription and investigation of hardware and software infrastructure incidents. Based on historical executions and AI-based algorithms, Proactive Maintenance can understand system behavior and detect anomalies and outliers before they occur. By

anticipating incidents and providing faster root cause analysis, the overall availability is dramatically increased. Hence, Proactive Maintenance can generate a significant carbon footprint reduction by avoiding unproductive jobs which fail in mid-execution.

- A prescriptive Security Operational Center (SOC) based on AI to detect unknown attack patterns and detect modified application code (hidden mining, etc.) that run on the supercomputer.
- Hybrid cloud support, which allows for multisite clusters and running applications on cloud providers (direct execution or bursting). In this context, the Eviden software suites can run on external resources and on the local cluster environment

Next steps

The software environment is critical to the optimization of HPC resource usage in terms of performance, power consumption and availability. The openness of HPC infrastructure is generating new use cases such as HPC, AI and Big Data using different frameworks. At the same time, this openness also comes with its own set of constraints, such as interactive jobs, multi-steps jobs with different hardware resources, etc. These operations must be as easy to run as on a standard personal computer.

From an administrative perspective, the flexibility and robustness of the solution is mandatory. From an end-user perspective, the ease of use and the ability to run an application smoothly are critical.

In this context, the security of execution (data, application, results) is a real challenge — especially in an open

environment involving multiple sites, multiple clouds and multiple external user environments. The ability to ensure security with data segregation, exchange encryption and a software-defined environment will become the de facto standard.

Furthermore, the imperative to achieve “digital sobriety” will be a key differentiator based on multiple criteria such as overall availability (downtime), the ratio of successful jobs, the most appropriate selection of hardware and software resources for a given job depending on its profile, and the dynamic adaptation of power usage according to job phases (CPU, memory, I/O).


The HPC software environment is not only paramount in the performance race, but is also instrumental in reducing the carbon footprint for the overall infrastructure.

Learn more at eviden.com/solutions/advanced-computing or contact us

Jean-Olivier Gerphagnon
Software Principal Architect
HPC, AI & QC
jean-olivier.gerphagnon@eviden.com

Matthieu Isoard
Software Product Manager
HPC, AI & QC
matthieu.isoard@eviden.com

Olivia-Yongmei Ding
HPC & QC Marketing
olivia-yongmei.ding@eviden.com



Connect with us

in /in/eviden

X @EvidenLive

@ @evidenlive

▶ /EvidenLive

eviden.com

Eviden is a registered trademark © Copyright 2023, Eviden SAS – All rights reserved



ECT-231103-SB-HPC-Software-Suites-position-paper