

## Super Computer Suite 5



BullSequana X Supercomputer Suite version 5 - SCS 5 - introduces a new approach to eXtreme computing software solutions. SCS 5 is a scalable, open, and robust software suite that meets the requirements of the most challenging high-performance computing (HPC) environments, with high security standards.

SCS 5 is the result of Atos's extensive experience in deploying large-scale supercomputers, combined with continued investment in Research & Development.

SCS 5 is designed to address every HPC need, whether it is for a small system with just a few hundred cores or for supercomputers with tens of thousands of nodes. Based on new-generation CPUs and GPUs, it is also designed to reach performance targets of systems positioned in the Top500. The main goal of SCS 5 is to provide a global high-performance supercomputing environment that includes:

- a standardized and scalable installation process with an enhanced update solution;
- mechanisms to ease integration of new hardware;
- a high-level default security including SELinux support and on-time fixes;
- a set of optional software to optimize energy consumption and I/O patterns and help understanding the application behavior;

- a container framework environment for application execution including container management (registry, versioning, resource management integration, etc);
- and of course, support for several user development and execution environments with top performances.

This HPC software suite is another step towards Exascale computing.

### A software suite for your supercomputer and your applications

SCS 5 is a modular software solution that manages supercomputer environments, including data and applications. It is based on an enterprise-class Linux distribution (RHEL 7™).

SCS 5 provides a robust, scalable, and flexible solution that adapts easily to any customer needs. It is closely integrated with the BullSequana X series such as the X1000 or XH2000 and works perfectly on any standard hardware.

SCS 5 supports Bull eXascale Interconnect (BXI), InfiniBand with the Mellanox® Fabric Manager (OFED) and Intel® Omni-Path architecture.

SCS 5 includes a complete profiling environment for MPI, energy and I/O, and standard development suites based on third-party software. The storage solution is based on DDN's Whamcloud Lustre.

All critical services are configured to ensure high availability.

SCS 5 is based on best-of-breed open source and ISV software. The selected open source components are enhanced by Atos's added-value features. ISV components can be added or can fully replace default components, depending on user needs.

By design, SCS 5 is modular, allowing any ISV, open-source, or home-grown software and tools to run seamlessly.

### Atos services to take full advantage of SCS 5

As a leading provider of end-to-end Extreme Computing solutions, Atos provides consultancy services to help its customers take full advantage of the benefits of SCS 5 throughout their projects: design, deployment, and operations. SCS 5 users can benefit from Atos' highly professional worldwide support offer that includes traditional maintenance and upgrade services as well as dedicated, customized, and proactive services.

# SCS 5 key components and features

## Operating System

- SCS 5 runs on Red Hat Enterprise Linux 7 which has proved its efficiency in HPC environments for years.
- SCS 5 is the best software stack choice for HPC applications, thanks to its robustness, scalability, manageability, security, and high availability.
- Atos and Red Hat technical experts have been working closely together for years, to make RHEL the ideal software environment for Atos hardware platforms.
- Atos - Red Hat customers have access to professional-class worldwide support services provided by high-level specialists who have extensive experience in deploying large scale supercomputers.

## SCS 5 Foundation

This module includes Mellanox® OFED, PAPI, advanced IPMI tools, and specific modules related to Atos advanced products:

- Mellanox® OFED is the InfiniBand fabric management stack from Mellanox.
- PAPI (core) is enhanced by Atos to support the latest CPU technologies in the period between CPU introduction and general support by the operating system.
- IPMI tools are delivered with enhanced management functionalities, power management, and inventory.

## SCS 5 Management Center

SCS 5 Management Center is SCS 5's administrative component and integrates all tools needed to install, configure, and manage a supercomputer. Depending on system size, the management will be done by:

- a management unit for supercomputers composed of up to 1,200 elements,
- a master management unit coupled with distributed management units for groups of equipment, for supercomputers with more than 1,200 elements.

The management infrastructure is designed to be scalable with a distributed and hierarchical environment.

A diskless operating system is available to ease deployment and enhance configuration.

The High Availability functionality is introduced for management nodes, thanks to HA support in the RHEL add-on.

Security is greatly improved thanks to SELinux which is activated for supercomputer management and under specific conditions for compute nodes.

The IBMS Interconnect monitoring tool is also provided to debug and tune the topology and routing on InfiniBand, BXI and Intel Omni-Path.

For more information: [hpc@atos.net](mailto:hpc@atos.net)

## SCS5 Open MPI

The SCS 5 Open MPI is based on open source MPI stack Open MPI 3x, which is a standard-compliant library for message passing and hybrid programming.

SCS 5 Open MPI provides key functionalities such as:

- runtime scalability improvement with PMIx support (PMI Exascale),
- integration and support of Mellanox accelerators on IB networks,
- MPI 3.1 standard conformance.
- support of THREAD\_MULTIPLE and Fortran 64-bit integer,
- integration of Portals 4 BTL and MTL for Bull eXascale Interconnect.
- support of MPI+X hybrid communication accelerator.

## SCS 5 Slurm

The batch manager is based on Slurm, the open source resource manager. To fulfil different requirements, SCS provides an LTS and a non-LTS version. Atos is a major Slurm contributor:

- hierarchical implementation based on hardware topology using the interconnect network to improve security and availability;
- support of Kerberos authentication through AUKS module;
- support of SELinux dedicated policy;
- power adaptive scheduling for applications to enhance power capping by managing unused nodes and reducing CPU frequency;
- energy fairshare scheduling based on energy consumption accounting;
- hyperthreading support to extend actual placement (socket and core) to hyperthread level.

## SCS 5 Lustre HA

This parallel file system is based on the Lustre core, providing high performance and large storage solutions for big data workloads.

Extra functionalities were added by Atos:

- integration of Lustre client and router with Mellanox® OFED stack,
- Shine centralized administration tool,
- monitoring with Shinken and Graphite,
- High Availability integration based on pacemaker.

## SCS 5 Performance Toolkit

These tools include Atos products such as Lightweight Profiler, Binding Checker, HPC Toolkit (with extensions), PAPI, and third-party products among which Intel® Parallel Studio software:

- Lightweight Profiler is a first-level profiler for MPI and Open MPI applications.

- Binding Checker is a helpful tool to check correct process binding.
- HPCToolkit features extensions that make it possible to detect processes with various behaviors and to compare successive runs.
- PAPI provides an open source API that gives access to the hardware performance counters.
- A complementary offer for the development environment can be purchased separately such as the Intel® Parallel Studio development environment software suite.

## SCS 5 Data Management

This module includes specific tools to understand and optimize I/O application patterns. Those tools are described in a dedicated data sheet.

## SCS 5 Power Management

Energy management is key for any supercomputer to enhance TCO. SCS 5 Energy Optimizer provides power monitoring interface and power-capping functionalities.

SCS 5 Dynamic Power Optimizer helps to automatically tune application energy footprint by using DVFS. Depending on application phase, the CPU consumption is live-adjusted.

## SCS 5 Containers Framework

Containers are key to the support of hybrid workloads on HPC environments. Singularity runtime is widely used in HPC as it offers a light and effective solution to run containers through Slurm.

In addition, a registry and a container builder are provided to ease adoption by larger communities.

## An optimized Life Cycle Process

The SCS product life cycle is tuned to provide both stability for users who prefer to keep the number of software environment modifications to a minimum, and significant flexibility for those who are interested in having the latest features available.

- A new release of SCS 5 is delivered every 18 months and the main stream (reference version) is stable for 18 months.
- Each release is supported for 24 months.
- The latest stream is based on the main stream and allows products to be updated during the SCS 5 release's life cycle.