

Bull Super Computer Suite 5



Bull supercomputer suite version 5, or SCS 5, introduces a new approach to extreme computing software solutions. Bull SCS 5 is a scalable, open, and robust software suite that meets the requirements of even the most challenging high performance computing (HPC) environments, which also require enhanced security.

Bull SCS 5 is the result of Atos's long experience in deploying large-scale supercomputers, combined to continued efforts in Research & Development.

Bull supercomputer suite version 5 is designed for every HPC need, from small supercomputers with just a few hundred cores to supercomputers with tens of thousands of nodes. It is cut out to reach performance targets of the order of up to 100 PFlops, based on new-generation CPUs and GPUs. The main goal of Bull SCS 5 is to provide a global high performance supercomputing environment. It includes:

- ▶ a standardized and scalable installation process with an enhanced update solution; mechanisms to ease integration of new hardware;
- ▶ default security with on-time fixes;
- ▶ and support for several user development and execution environments with top performance.

This new generation HPC software suite is a further step towards Exascale computing.

A software suite for your supercomputer and your applications

Bull SCS 5 is a software solution that globally manages supercomputer environments, including data and applications. It is a new product based on a recent Linux distribution (RHEL 7™).

Bull SCS 5 provides a robust, scalable, modular and flexible solution that easily adapts to customer needs. It is closely integrated with the new Bull sequana X1000

series based on the latest available technologies. It supports EDR InfiniBand with the MOFED stack and MPI accelerators MXM/FCA (both from Mellanox®), a complete development suite and Lustre (based on the Intel® solution). All critical services are configured to ensure High Availability.

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

By design, Bull 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 Bull supercomputer suite

As a leading provider of end-to-end Extreme Computing solutions, Atos provides consultancy services to help its customers to take full advantage of the benefits of Bull supercomputer suite throughout their projects: design, deployment and operations. Bull supercomputer suite users can benefit from Atos's worldwide and highly professional support offer, comprising traditional maintenance and upgrade services as well as dedicated, personalized and pro-active services.

Bull SCS 5 key components and features

Operating System

- ▶ Bull supercomputer suite 5 runs on Red Hat Enterprise Linux 7 which has proved its efficiency in HPC environments for years with its previous versions.
- ▶ Bull SCS 5 is the software stack of 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 Bull hardware platforms.
- ▶ Atos - Red Hat customers have access to professional-class worldwide support services provided by high level specialists who have a long experience of deploying large scale supercomputers.

Bull Foundation

This module includes MOFED, PAPI, advanced IPMI tools and specific modules related to Bull advanced products:

- ▶ MOFED is the InfiniBand fabric management (OFED) stack from Mellanox.
- ▶ PAPI (core) is enhanced by Atos to support the latest CPUs 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.

Bull Management Center

Bull Management Center is the administrative component of SCS 5 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 1,000 elements,
- ▶ a master management unit coupled with distributed management units for groups of equipment, for supercomputers with more than 1,000 elements.

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

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 RHEL add-on.

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

Bull OpenMPI

The Bull Open MPI is based on open source MPI stack Open MPI 2.x, which is a standard-compliant library for message passing and hybrid programming.

Bull Open MPI provides key functionalities such as:

- ▶ run time scalability improvement with PMIx support (PMI Exascale),
- ▶ integration and support of Mellanox MXM and FCA accelerators,
- ▶ 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.

Bull Slurm

This batch manager is based on Slurm, the open source resource manager. Major enhancements in version 15.08, to which Bull is a major contributor, include:

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

Heterogeneous resources management and MPMD (Multi Process Multi Data) will be the next step of development for future versions.

Bull Lustre HA

This parallel file system is based on the Intel® Enterprise Edition for Lustre (IEEL) core, providing high performance and large storage solutions for big data workloads.

Extra functionalities were added by Atos for Lustre 2.7:

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

Bull Performance Toolkit

These tools include Bull products such as HPC Toolkit (with Bull extensions), PAPI, and third-party products among which Intel® Parallel Studio XE software:

- ▶ HPCToolkit features Bull 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 available in modern processors, including latest generation thanks to Bull Foundation.
- ▶ A complementary offer for the development environment can be purchased separately: the Intel® Parallel Studio XE development environment software suite.

Bull Accelerator Environment

This module includes specific packaged drivers and additional software for accelerators such as NVIDIA GPUs.

Bull Maintenance Manager

The Maintenance Manager provides the specific Bull tool Argos for the maintenance of your system, to keep it up-to-date and alive with fine tracking of maintenance operations.

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 also 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 mainstream and allows products to be updated during the life cycle of the SCS 5 release.
- ▶ Other components can be provided in Contrib mode (various software for evaluation).

Migration from previous release 4

Systems previously installed with bullx scs 4 can migrate to SCS 5 at the time of its general availability. Full support of bullx scs 4 remains available until the second release of SCS 5 (R2), scheduled for middle of 2017.

Guidelines will be provided by your Atos support team for migration

