

NSM Software Stack Details

It is an aggregation of software components that work in tandem to accomplish a given task. The software stack provided with NSM systems has a gamut of software components which meets all the requirements of a user and that of a system administrator.

- **OS: Linux** CentOS (Community Enterprise Operating System) is a Linux distribution that attempts to provide a free, enterprise-class, community-supported computing platform functionally compatible with its upstream source, Red Hat Enterprise Linux (RHEL)
- **Middle-ware Application and components**
 - **File System:**
 - **NFS:** A Network File System (NFS) allows remote hosts to mount file systems over a network and interact with those file systems as though they are mounted locally. This enables system administrators to consolidate resources onto centralized servers on the network.
 - **Local FS (XFS):** XFS is particularly proficient at parallel IO due to its allocation group-based design. This enables extreme scalability of IO threads, file system bandwidth, file and file system size when spanning multiple storage devices.
 - **Lustre:** The Lustre © file system is an open-source, parallel file system that supports many requirements of leadership class HPC simulation environments.
 - **Provisioning:**
 - **OpenHPC(xCAT):** OpenHPC is a collaborative, community effort that initiated from a desire to aggregate a number of common ingredients required to deploy and manage High-Performance Computing (HPC) Linux clusters including provisioning tools, resource management, I/O clients, development tools, and a variety of scientific libraries. Packages provided by OpenHPC have been pre-built with HPC integration in mind with a goal to provide re-usable building blocks for the HPC community.
 - **Resource Management and Job Scheduling:**
 - **SLURM:** Slurm is an open-source, fault-tolerant and highly scalable cluster management and job scheduling system for large and small Linux clusters. Slurm requires no kernel modifications for its operation and is relatively self-contained. As a cluster workload manager, Slurm has three key functions. First, it allocates exclusive and/or non-exclusive access to resources (compute nodes) to users for some duration of time so they can perform work. Second, it provides a framework for starting, executing, and monitoring work (normally a parallel job) on the set of allocated nodes. Finally, it arbitrates contention for resources by managing a queue of pending work.
 - **Cluster Monitoring:**
 - **Ganglia:** Ganglia is a scalable distributed monitoring system for high-performance computing systems such as clusters and Grids. It is based on a hierarchical design targeted at federations of clusters. It leverages widely used technologies such as XML for data representation, XDR for compact, portable data transport, and RRDtool for data storage and visualization.
 - **Nagios:** Nagios offers to monitor and alerting services for servers, switches, applications and services. It alerts users when things go

wrong and alert them a second time when the problem has been resolved.

- **XDMoD:** The Open XDMoD portal provides a rich set of features accessible through an intuitive graphical interface, which is tailored to the role of the user. Metrics provided include a number of jobs, CPU hours consumed, wait time, and wall time, with a minimum, maximum and the average of these metrics, in addition to many others. Metrics are organized by a customizable hierarchy appropriate for your organization.
- **OSTicket:** osTicket is a widely-used open source support ticket system. It seamlessly integrates inquiries created via email, phone and web-based forms into a simple easy-to-use multi-user web interface. Manage, organize and archive all your support requests and responses in one place while providing your customers with accountability and responsiveness they deserve.
- **HPC Programming Tool**
 - **Communication Libraries:**
 - **Intel MPI:** Commercial Compiler support for MPI environment.
 - **MVAPICH2:** The MVAPICH2 software, based on MPI 3.1 standard, delivers the best performance, scalability and fault tolerance for high-end computing systems and servers using Infiniband, Omni-Path, Ethernet/iWARP, and RoCE networking technologies.
 - **OpenMPI:** The Open MPI Project is an open-source Message Passing Interface implementation that is developed and maintained by a consortium of academic, research, and industry partners. Open MPI is, therefore, able to combine the expertise, technologies, and resources from all across the High-Performance Computing community in order to build the best MPI library available. Open MPI offers advantages for system and software vendors, application developers and computer science researchers.
 - **PGAS:** Partitioned Global Address Space is a programming model suited for shared and distributed memory parallel machines, e.g., machines consisting of many (up to hundreds of thousands of) CPUs.
 - **Development Tools:**
 - **Intel Cluster Studio:** Intel Cluster Studio: Intel® Parallel Studio XE Composer Edition provides a software tools environment for developing Fortran, C, and/or C++ code using Intel® Compilers. Intel® Parallel Studio XE Composer Edition also includes Intel® Math Kernel Library, Intel® Performance Primitives, Intel® Threading Building Blocks, and Intel® Data Analytics Acceleration Library (Intel® DAAL). Intel® Parallel Studio XE Professional Edition adds Intel® VTune™ Amplifier XE for performance analysis, Intel® Inspector for correctness analysis, and Intel® Advisor for parallelism discovery. Intel® Parallel Studio XE Cluster Edition adds support for distributed memory computing via Intel® MPI Library, Intel® MPI Benchmarks, and Intel® Trace Analyzer and Collector. Intel® Cluster Checker provides cluster health monitoring tools.
 - **GNU:** The GNU Compiler Collection (GCC) is a compiler system produced by the GNU Project supporting various programming languages. GCC is a key component of the GNU toolchain and the standard compiler for most Unix-like Operating Systems. The Free Software Foundation (FSF) distributes GCC under the GNU General Public License (GNU GPL). GCC has played an important role in the growth of free software, as both a tool and an example.

- **CUDA Toolkit:** The NVIDIA® **CUDA® Toolkit** provides a development environment for creating high-performance GPU-accelerated applications. The **toolkit** includes GPU-accelerated libraries, debugging and optimization tools, a C/C++ compiler, and a runtime library to deploy your application.
 - **Visualization Tools:**
 - **Ferret:** It is an interactive computer visualization and analysis environment designed to meet the needs of oceanographers and meteorologists analyzing large and complex gridded data sets. Ferret offers a Mathematica-like approach to analysis; new variables may be defined interactively as mathematical expressions involving data set variables. Calculations may be applied over arbitrarily shaped regions. Fully documented graphics are produced with a single command. It runs most UNIX and Linux systems using X Window for display, and on Windows XP/NT/9x.
 - **GrADS:** The Grid Analysis and Display System (GrADS) is an interactive desktop tool that is used for easy access, manipulation, and visualization of earth science data. GrADS has two data models for handling gridded and station data. GrADS supports many data file formats, including binary (stream or sequential), GRIB (version 1 and 2), NetCDF, HDF (version 4 and 5), and BUFR (for station data). GrADS has been implemented worldwide on a variety of commonly used operating systems and is freely distributed over the Internet.
 - **ParaView:** It is an open-source multiple platform application for interactive, scientific visualization. It has a server architecture to facilitate remote visualization of datasets and generates a level of detail (LOD) models to maintain interactive frame rates for large datasets. It is an application built on top of the Visualization Toolkit (VTK) libraries. ParaView is an application designed for data parallelism on shared-memory or distributed-memory multicomputer and clusters. It can also be run as a single computer application.
 - **VMD:** It is an open-source interactive parallel visualization and graphical analysis tool for viewing scientific data. It can be used to visualize scalar and vector fields defined on 2D and 3D structured and unstructured meshes. The visIt was designed to handle very large data set sizes in the terascale range and yet can handle small data sets in the kilobyte range.
 - **Performance Monitoring:**
 - **HPCC:** HPC Challenge Benchmark combines several benchmarks to test a number of independent attributes of the performance of high-performance computer (HPC) systems.
 - **HPCG:** The High-Performance Conjugate Gradients (HPCG) Benchmark project is an effort to create a new metric for ranking HPC systems.
 - **IOR:** IOR is designed to measure parallel file system I/O performance at both the POSIX and MPI-IO level. This parallel program performs writes, reads to/from files under several sets of conditions, and reports the resulting throughput rates.
 - **IMB/OSU:** The Intel® MPI Benchmarks perform a set of MPI performance measurements for point-to-point and global communication operations for a range of message sizes.
- **About C-DAC Component**
 - **HPC Task Automation Scripts:** These scripts are used to do streamline installation across the sites, Avoids any missing parameters or any

- configuration. It also provides few important post-installation functions that allow an administrator to add a new node and new group to the cluster
- **Cluster Checker Scripts:** Cluster Checker verifies the configuration and performance of Linux-based clusters and checks compliance with the Cluster Ready architecture specification. If issues are found, Cluster Checker diagnoses the problems and may provide recommendations on how to repair the cluster.
 - **CHReME:** CHReME is a Web-Based Resource Management Portal that empowers Scientists, Researchers, System administrators, and HPC users with an intuitive GUI for utilizing, managing and monitoring, tracking the HPC workload. By simply logging into the portal, users can access and control their computing resources from anywhere using standards-compliant Web interface via Intranet or the public Internet. CHReME provides a fully customizable easy to use workflow interfaces to the HPC users for state-based pre-processing, execution and post-processing of an application. CHReME's state-based execution notifies the users' with current workflow status which helps in tracing their application execution state and allows the user to rerun the application.
 - **C-CHAKSHU:** C-Chakshu is an HPC multi-cluster monitoring and Management platform which addresses the effective usage of deployed HPC systems by comprehensive monitoring. It provides a unified dashboard over the web for all NSM sites with different geographic location across India. Its web interface makes clusters of different magnitude easy to manage hence facilitating the HPC system administrator which makes it easy for the researchers and scientists of the varied domain to carry out their scientific simulation with minimal efforts.
 - **CAPC:** C -DAC Automatic Parallelizing Compiler, CAPC, is an auto-parallelizing source-to-source compiler which converts a sequential program (e.g. sequential C program) into an equivalent parallel program (e.g. parallel OpenMP/OpenCL program).
 - **SUM:** It is a software which aids computing facility managers to manage users' data and take appropriate decision based on data. It is a user management and analytics portal which helps in streamlining the user approval process. It also helps in decision making by including usage data from the cluster.