High Performance Linux Clusters With Oscar Rocks Openmosix And Mpi Joseph D Sloan

TOWNSEND KYLER

Building a Linux HPC Cluster with XCAT IBM Redbooks
The demand for computational power has significantly increased in the nanoscience era because simulations are usually performed at molecular levels with quantum mechanics. At the same time, the performance and availability of both commodity computers and inexpensive high-speed networking hardware have increased drastically in the past decades. Hence, scientists and engineers working on nanoscience simulations may attempt to build a small- or medium-size high-performance computing cluster at their laboratories. However, building a high-performance computing cluster is not a simple task, especially for nanoscience researchers whose specialties are not on subjects in computer science. Fortunately, with the usage of Diskless Remote Boot for Linux (DRBL), the construction and maintenance of a computing cluster are greatly simplified and, thus, become feasible for researchers without much prior exposure to computer sciences. This monograph presents the procedures to construct a high-performance computing cluster in a way that can be followed by less system-oriented researchers. The interested readers can follow the idea or can use it as a guide to build their own system.
Building Clustered Linux Systems LAP Lambert Academic Publishing
The Linux Enterprise Cluster explains how to take a number of inexpensive computers with limited resources, place them on a normal computer network, and install free software so that the computers act together like one powerful server. This makes it possible to build a very inexpensive and reliable business system for a small business or a large corporation. The book includes information on how to build a high-availability server pair using the Heartbeat package, how to use the Linux Virtual Server load balancing software, how to configure a reliable printing system in a Linux cluster environment, and how to build a job scheduling system in Linux with no single point of failure. The book also includes information on high availability techniques that can be used with or without a cluster, making it helpful for System Administrators even if they are not building a cluster. Anyone interested in deploying Linux in an environment where low cost computer reliability is important will find this book useful. The CD-ROM includes the Linux kernel, ldirectord software, the Mon monitoring package, the Ganglia package, OpenSSH, rsync, SystemImager, Heartbeat, and all the figures and illustrations used in the book.
Building the World's Fastest Linux Cluster Sams Publishing
This IBM® Redbooks® publication provides information about aspects of performing infrastructure health checks, such as checking the configuration and verifying the functionality of the common subsystems (nodes or servers, switch fabric, parallel file system, job management, problem areas, and so on). This IBM Redbooks publication documents how to monitor the overall health check of the cluster infrastructure, to deliver technical computing clients cost-effective, highly scalable, and robust solutions. This IBM Redbooks publication is targeted toward technical professionals (consultants, technical support staff, IT Architects, and IT Specialists) responsible for delivering cost-effective Technical Computing and IBM High Performance Computing (HPC) solutions to optimize business results, product development, and scientific discoveries. This book provides a broad understanding of a new architecture.
High Performance Computing – using the Torque (PBS) resource manager Packt Publishing Ltd
"Linux Clustering" is the premier resource for system administrators wishing to implement clustering solutions on the many types of Linux systems. It guides Linux Administrators through difficult tasks while offering helpful tips and tricks.
High Performance Computing NOITE S.C.
"This Ebook presents state-of-the-art solutions in applications of modern high performance networks. The topics covered in this Ebook include mobile ad-hoc networks, clusters for distance computing, clustering technologies and deployment, emerging wireless"
A General Purpose High Performance Linux Installation Infrastructure Prentice Hall
Create high availability clusters to enhance system performance using CentOS 7 About This Book Master the concepts of high performance and high availability to eliminate performance bottlenecks Maximize the uptime of services running in a CentOS 7 cluster A step-by-step guide that will provide knowledge of methods and approaches to optimize the performance of CentOS clusters Who This Book Is For This book is targeted at system administrators: those who want a detailed, step-by-step guide to learn how to set up a high-availability CentOS 7 cluster, and those who are looking for a reference book to help them learn or refresh the necessary skills to ensure their systems and respective resources are utilized optimally. No previous knowledge of high-availability systems is needed, though the reader is expected to have at least some degree of familiarity with any spin-off of the Fedora family of Linux distributions, preferably CentOS. What You Will Learn Install a CentOS 7 cluster and network infrastructure Configure firewall, networking, and clustering services and settings Set up and test a HAC (high-availability cluster) to host an Apache web server and a MariaDB database server Monitor performance and availability Identify bottlenecks and troubleshoot issues Improve performance and ensure high availability In Detail CentOS is the enterprise level Linux OS, which is 100% binary compatible to Red Hat Enterprise Linux (RHEL). It acts as a free alternative to RedHat's commercial Linux offering, with only a change in the branding. A high performance cluster consists in a group of computers that work together as one set parallel, hence minimizing or eliminating the downtime of critical services and enhancing the performance of the application. Starting with the basic principles of clustering, you will learn the necessary steps to install a cluster with two CentOS 7 servers. We will then set up and configure the basic required network infrastructure and clustering services. Further, you will learn how to take a proactive approach to the split-brain issue by configuring the failover and fencing of the cluster as a whole and the quorum of each node individually. Further, we will be setting up HAC and HPC clusters as a web server and a database server. You will also master the art of monitoring performance and availability, identifying bottlenecks, and exploring troubleshooting techniques. At the end of the book, you'll review performance-tuning techniques for the recently installed cluster, test performance using a payload simulation, and learn the necessary skills to ensure that the systems, and the corresponding resources and services, are being utilized to their best capacity. Style and approach An easy-to-follow and step-by-step guide with hands-on instructions to set up real-world simple cluster scenarios that will start you on the path to building more complex applications on your own.
Deploying Linux on IBM EServer PSeries Clusters "O'Reilly Media, Inc."
For more than 20 years, Network World has been the premier provider of information, intelligence and insight for network and IT executives responsible for the digital nervous systems of large organizations. Readers are responsible for designing, implementing and managing the voice, data and video systems their companies use to support everything from business critical applications to employee collaboration and electronic commerce.
Deploying LINUX on the Desktop John Wiley & Sons
An authoritative guide to today's revolution in "commodity supercomputing, " this book brings together more than 100 of the field's leading practitioners, providing a single source for up-to-the-minute information on virtually every key system issue associated with high-performance cluster computing.
Implementing an IBM High-Performance Computing Solution on IBM Power System S822LC NRC Research Press

This IBM® Redbooks® publication demonstrates and documents that IBM Power SystemsTM high-performance computing and technical computing solutions deliver faster time to value with powerful solutions. Configurable into highly scalable Linux clusters, Power Systems offer extreme performance for demanding workloads such as genomics, finance, computational chemistry, oil and gas exploration, and high-performance data analytics. This book delivers a high-performance computing solution implemented on the IBM Power System S822LC. The solution delivers high application performance and throughput based on its built-for-big-data architecture that incorporates IBM POWER8® processors, tightly coupled Field Programmable Gate Arrays (FPGAs) and accelerators, and faster I/O by using Coherent Accelerator Processor Interface (CAPI). This solution is ideal for clients that need more processing power while simultaneously increasing workload density and reducing datacenter floor space requirements. The Power S822LC offers a modular design to scale from a single rack to hundreds, simplicity of ordering, and a strong innovation roadmap for graphics processing units (GPUs). This publication is targeted toward technical professionals (consultants, technical support staff, IT Architects, and IT Specialists) responsible for delivering cost effective high-performance computing (HPC) solutions that help uncover insights from their data so they can optimize business results, product development, and scientific discoveries

Optimizing HPC Applications with Intel Cluster Tools Bentham Science Publishers

Praise for Building Clustered Linux Systems "The author does an outstanding job of presenting a very complicated subject. I very much commend this work. The author sets the pace and provides vital resources and tips along the way. He also has a very good sense of humor that is crafted in the text in such a way that makes the reading enjoyable just when the subject may demand a break. This book should be a requirement for those that are clustering or considering clustering and especially those considering investing a great deal of financial resource toward that goal."--Joe Brazeal, Information Technician III, Southwest Power Pool "This book is for Beginner and Intermediate level system administrators, engineers, and researchers, who want to learn how to build Linux clusters. The book covers everything very well." -Ibrahim Haddad, Senior Researcher, Ericsson Corporate Unit of Research "Nothing that I know of exists yet that covers this subject in as much depth and detail. The practical 'hands-on' approach of this book on how to build a Linux cluster makes this a very valuable reference for a very popular, highly demanded technology." -George Vish, II, Linux Curriculum Program Manager and Senior Education Consultant, HP "In my opinion there is a significant lack of literature on this subject. Most of the currently available books are either dated or do not address the complete picture of the range of decisions that must go into building a Linux cluster. I feel comfortable recommending this to anyone interested in building a Linux cluster to better understand both the technical aspects of building and designing a Linux cluster, but also the business aspects of the same." -Randall Splinter Ph. D., Senior Solution Architect, HP "The author has set a precedent in the cluster design and integration process that is lacking in the industry today." --Stephen Gray, Senior Applications Engineer, Altair Engineering, Incorporated The Practical, Step-by-Step Guide to Building and Running Linux Clusters Low-cost, high-performance Linux clusters are the best solution for an increasingly wide range of technical and business problems. Until now, however, building and managing Linux clusters has required more specialized knowledge than most IT organizations possess. This book dramatically lowers the learning curve, bringing together all the hands-on knowledge and step-by-step techniques you'll need to get the job done. Using practical examples, Robert Lucke simplifies ever ...

Building Clustered Linux Systems Elsevier

The author teaches at Wofford College.

High Performance Computing and Communications MIT Press

Pro Linux High Availability Clustering teaches you how to implement this fundamental Linux add-on into your business. Linux High Availability Clustering is needed to ensure the availability of mission critical resources. The technique is applied more and more in corporate datacenters around the world. While lots of documentation about the subject is available on the internet, it isn't always easy to build a real solution based on that scattered information, which is often oriented towards specific tasks only. Pro Linux High Availability Clustering explains essential high-availability clustering components on all Linux platforms, giving you the insight to build solutions for any specific case needed. In this book four common cases will be explained: Configuring Apache for high availability Creating an Open Source SAN based on DRBD, iSCSI and HA clustering Setting up a load-balanced web server cluster with a back-end, highly-available database Setting up a KVM virtualization platform with high-availability protection for a virtual machine. With the knowledge you'll gain from these real-world applications, you'll be able to efficiently apply Linux HA to your work situation with confidence. Author Sander Van Vugt teaches Linux high-availability clustering on training courses, uses it in his everyday work, and now brings this knowledge to you in one place, with clear examples and cases. Make the best start with HA clustering with Pro Linux High Availability Clustering at your side.

Pro Linux High Availability Clustering Apress

Until now, building and managing Linux clusters has required more intimate and specialized knowledge than most IT organizations possess. This book dramatically lowers the learning curve, bringing together all the hands-on knowledge and step-by-step techniques needed to get the job done.

High Performance Linux Clusters with OSCAR, Rocks, OpenMosix, and MPI Apress

Enabling technologies - An overview of cluster computing / Thomas Sterling / - Node Hardware / Thomas Sterling / - Linux / Peter H. Beckman / - Network Hardware / Thomas Sterling / - Network Software / Thomas Sterling / - Setting Up clusters : installation and configuration - How fast is my beowulf? / David Bailey / - Parallel programming / - Parallel programming with MPI / William Gropp / - Advanced topics in MPI programming / William Gropp / - Parallel programming with PVM / Al Geist / - Fault-tolerant and adaptive programs with PVM / Al Geist / - Managing clusters / - Cluster workload management / James Patton Jones / - Condor : a distributed job scheduler / - Maui scheduler : A multifunction cluster scheduler / David B. Jackson / - PBS : portable batch system / James Patton Jones / - PVFS : parallel virtual file system / Walt Ligon / - Chiba city : the Argonne scalable cluster.

The Linux Enterprise Cluster IGI Global

The computer industry today is no longer driven, as it was in the 40s, 50s and 60s, by High-performance computing requirements. Rather, HPC systems, especially Leadership class systems, sit on top of a pyramid investment mode. Figure 1 shows a representative pyramid investment model for systems hardware. At the base of the pyramid is the huge investment (order 10s of Billions of US Dollars per year) in semiconductor fabrication and process technologies. These costs, which are approximately doubling with every generation, are funded from investments multiple markets: enterprise, desktops, games, embedded and specialized devices. Over and above these base technology investments are investments for critical technology elements such as microprocessor, chipsets and memory ASIC components. Investments for these components are spread across the same markets as the base semiconductor processes investments. These second tier investments are approximately half the size of the lower level of the pyramid. The next technology investment layer up, tier 3, is more focused on scalable computing systems such as those needed for HPC and other markets. These tier 3 technology elements include networking (SAN, WAN and LAN), interconnects and large scalable SMP designs. Above these is tier 4 are relatively small investments necessary to build very large, scalable systems high-end or Leadership class systems. Primary among these are the specialized network designs of vertically integrated systems, etc.

High Performance Cluster Computing No Starch Press

Since 2003, with the installation of a 256 processor Linux Networx XEON 1686 system (Powell), Army Research Laboratory (ARL) has been providing large-scale Linux cluster production systems for use within the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP). This initial system was followed in 2004 with the 2,048 processor Linux Networx XEON EM64T (JVN) and 2,304 processor IBM Opteron (Stryker) systems, and in 2007 by the 3,368-core Intel Dempsey (Humvee) and 4,488-core Intel Woodcrest (MJM) systems. These latest systems provide an increased peak performance of over 15 times the original Xeon 1686 system in a four year period. The purpose of this paper is provide a comparative study of the three generations of Linux clusters' capabilities to process some of the most widely used production codes used within the ARL environment. The codes to be benchmarked will include CTH, CFD++, GAMESS, and OVERFLOW. The codes will be run on Powell, JYN and Ping (the Woodcrest testbed machine). The results will focus attention on how architecture enhancements (including CPU speed, memory per node, cache size, and interconnect fabric transfer rates) have affected the overall performance of these systems.

Beowulf Cluster Computing with Linux Sams Publishing

Modeling and simulation (M&S) needs high-performance computing resources, but conventional supercomputers are both expensive and not necessarily well suited to M&S tasks. Discrete Event Simulation (DES) often involves repeated, independent runs of the same models with different input parameters. A system which is able to run many replications quickly is more useful than one in which a single monolithic application runs quickly. A loosely coupled parallel system is indicated. Inexpensive commodity hardware, high speed local area networking, and open source software have created the potential to create just such loosely coupled parallel systems. These systems are constructed from Linux-based computers and are called Beowulf clusters. This thesis presents an analysis of clusters in high-performance computing and establishes a testbed implementation at the MOVES Institute. It describes the steps necessary to create a cluster, factors to consider in selecting hardware and software, and describes the process of creating applications that can run on the cluster. Monitoring the running cluster and system administration are also addressed.

Linux Cluster Architecture Prentice Hall

This book will meet the needs of those LINUX users who wish to set up a desktop LINUX workstations, and integrate them into their corporate environment. It will provide practical answers to such questions as: a) What tools do I use to fully integrate with the Microsoft Office tool suite? b) How do I set up my email and interact with a Microsoft Exchange Server? c) Where can I obtain, and how do I install, Internet browser plug-ins needed for web access, media playing, and other corporate Internet functionality? Provides a guide to using LINUX on the desktop for the corporate user. It will cover more than basic topics, such as whether to use OpenOffice or use another tool such as Evolution; they will delve into specific configurations necessary to interact efficiently with the Microsoft centric world of the Desktop. This guide will cover those problem areas that arise and discuss how to smooth over the bumps while meeting the goal of using a LINUX desktop.Lastly this book will cover whether or not a complete LINUX solution is available, or if some hybrid desktop will be needed to interact smoothly in the modern corporate computing environment; including a discussion of necessary LINUX growth directions for future expansion and capability.·Reviews real world requirements.·Covers Pure LINUX, and Hybrid Corporate Desktops.·Covers Enabling Tools such as CrossOver Office and the use of Windows Native programs on LINUX.·Reveals Interoperability Concerns.·Implements a solid Corporate Desktop.·Reviews the complete costs of Implementing LINUX as a desktop.

Biomedical Diagnostics and Clinical Technologies: Applying High-Performance Cluster and Grid Computing NOITE S.C.

Since fall 2001, Livermore Computing at Lawrence Livermore National Laboratory has deployed 11 Intel IA-32-based Linux clusters ranging in size up to 1154 nodes. All provide a common programming model and implement a similar cluster architecture. Hardware components are carefully selected for performance, usability, manageability, and reliability and are then integrated and supported using a strategy that evolved from practical experience. Livermore Computing Linux clusters run a common software environment that is developed and maintained in-house while drawing components and additional support from the open source community and industrial partnerships. The environment is based on Red Hat Linux and adds kernel modifications, cluster system management, monitoring and failure detection, resource management, authentication and access control, development environment, and parallel file system. The overall strategy has been successful and demonstrates that world-class high-performance computing resources can be built and maintained using commodity off-the-shelf hardware and open source software.

Automated Installation of Linux High-performance Computing Clusters IBM Redbooks

With more and more and larger and larger Linux clusters, the question arises how to install them. This paper addresses this question by proposing a solution using only standard software components. This installation infrastructure scales well for a large number of nodes. It is also usable for installing desktop machines or diskless Linux clients, thus, is not designed for cluster installations in particular but is, nevertheless, highly performant. The infrastructure proposed uses PXE as the network boot component on the nodes. It uses DHCP and TFTP servers to get IP addresses and a bootloader to all nodes. It then uses kickstart to install Red Hat Linux over NFS. We have implemented this installation infrastructure at SLAC with our given server hardware and installed a 256 node cluster in 30 minutes. This paper presents the measurements from this installation and discusses the bottlenecks in our installation.

Related with High Performance Linux Clusters With Oscar Rocks Openmosix And Mpi Joseph D Sloan:
• Mark Jefferson Science Complex : click here