Título:

Autor: **Precio:** $580.00

Editorial: **Año:** 2004

Tema: **Edición:** 1ª

Sinopsis   **ISBN:** 9780596005702

To the outside world, a "supercomputer" appears to be a single system. In fact, it's a cluster of computers that share a local area network and have the ability to work together on a single problem as a team. Many businesses used to consider supercomputing beyond the reach of their budgets, but new Linux applications have made high-performance clusters more affordable than ever. These days, the promise of low-cost supercomputing is one of the main reasons many businesses choose Linux over other operating systems.

This new guide covers everything a newcomer to clustering will need to plan, build, and deploy a high-performance Linux cluster. The book focuses on clustering for high-performance computation, although much of its information also applies to clustering for high-availability (failover and disaster recovery). The book discusses the key tools you'll need to get started, including good practices to use while exploring the tools and growing a system. You'll learn about planning, hardware choices, bulk installation of Linux on multiple systems, and other basic considerations. Then, you'll learn about software options that can save you hours--or even weeks--of deployment time.

Since a wide variety of options exist in each area of clustering software, the author discusses the pros and cons of the major free software projects and chooses those that are most likely to be helpful to new cluster administrators and programmers. A few of the projects introduced in the book include:

MPI, the most popular programming library for clusters. This book offers simple but realistic introductory examples along with some pointers for advanced use.
OSCAR and Rocks, two comprehensive installation and administrative systems
openMosix (a convenient tool for distributing jobs), Linux kernel extensions that migrate processes transparently for load balancing
PVFS, one of the parallel filesystems that make clustering I/O easier
C3, a set of commands for administering multiple systems
Ganglia, OpenPBS, and cloning tools (Kickstart, SIS and G4U) are also covered. The book looks at cluster installation packages (OSCAR & Rocks) and then considers the core packages

individually for greater depth or for folks wishing to do a custom installation. Guidelines for debugging, profiling, performance tuning, and managing jobs from multiple users round out this immensely useful book.

About the Author
Joseph D. Sloan has been working with computers since the mid-1970s. He began using Unix as a graduate student in 1981, first as an applications programmer and later as a system programmer and system administrator. Since 1988 he has taught computer science, first at Lander University and more recently at Wofford College where he can be found using the software described in this book.