

| | |
|-------------------------|--|
| 1. Record Nr. | UNINA9910734841203321 |
| Autore | Borin Edson |
| Titolo | High Performance Computing in Clouds : Moving HPC Applications to a Scalable and Cost-Effective Environment / / edited by Edson Borin, Lúcia Maria A. Drummond, Jean-Luc Gaudiot, Alba Melo, Maicon Melo Alves, Philippe Olivier Alexandre Navaux |
| Pubbl/distr/stampa | Cham : , : Springer International Publishing : , : Imprint : Springer, , 2023 |
| ISBN | 3-031-29769-5 |
| Edizione | [1st ed. 2023.] |
| Descrizione fisica | 1 online resource (337 pages) |
| Altri autori (Persone) | DrummondLúcia Maria A GaudiotJean-Luc MeloAlba Melo AlvesMaicon NavauxPhilippe Olivier Alexandre |
| Disciplina | 004.6782 |
| Soggetti | Cloud computing Big data Machine learning Cloud Computing Big Data Machine Learning |
| Lingua di pubblicazione | Inglese |
| Formato | Materiale a stampa |
| Livello bibliografico | Monografia |
| Nota di contenuto | Chapter. 1. Why move HPC applications to the Cloud? -- Part. I. Foundations -- Chapter. 2. What is Cloud Computing? -- Chapter. 3. What do HPC applications look like? -- Part. II. Running HPC Applications in Cloud -- Chapter. 4. Deploying and Configuring Infrastructure -- Chapter. 5. Executing Traditional HPC Application Code in Cloud with Containerized Job Schedulers -- Chapter. 6. Designing Cloud-friendly HPC Applications -- Chapter. 7. Exploiting Hardware Accelerators in Clouds -- Part III. Cost and Performance Optimizations -- Chapter. 8. Optimizing Infrastructure for MPI Applications -- Chapter. 9. Harnessing Low-Cost Virtual Machines on the Spot -- Chapter. 10. Ensuring Application Continuity with Fault |

Tolerance Techniques -- Chapter. 11. Avoiding Resource Wastage -- Part. IV. Application Study Cases -- Chapter. 12. Biological Sequence Comparison on Cloud-based GPU Environment -- Chapter. 13. Oil & Gas Reservoir Simulation in the Cloud -- Chapter. 14. Cost effective deep learning on the cloud -- Appendix A. Deploying an HPC cluster on AWS -- Appendix B. Configuring a cloud-deployed HPC cluster.

Sommario/riassunto

This book brings a thorough explanation on the path needed to use cloud computing technologies to run High-Performance Computing (HPC) applications. Besides presenting the motivation behind moving HPC applications to the cloud, it covers both essential and advanced issues on this topic such as deploying HPC applications and infrastructures, designing cloud-friendly HPC applications, and optimizing a provisioned cloud infrastructure to run this family of applications. Additionally, this book also describes the best practices to maintain and keep running HPC applications in the cloud by employing fault tolerance techniques and avoiding resource wastage. To give practical meaning to topics covered in this book, it brings some case studies where HPC applications, used in relevant scientific areas like Bioinformatics and Oil and Gas industry were moved to the cloud. Moreover, it also discusses how to train deep learning models in the cloud elucidating the key components and aspects necessary to train these models via different types of services offered by cloud providers. Despite the vast bibliography about cloud computing and HPC, to the best of our knowledge, no existing manuscript has comprehensively covered these topics and discussed the steps, methods and strategies to execute HPC applications in clouds. Therefore, we believe this title is useful for IT professionals and students and researchers interested in cutting-edge technologies, concepts, and insights focusing on the use of cloud technologies to run HPC applications.
