

Linux Performance Tools Brendan Gregg

Right here, we have countless ebook **Linux Performance Tools Brendan Gregg** and collections to check out. We additionally have the funds for variant types and then type of the books to browse. The good enough book, fiction, history, novel, scientific research, as without difficulty as various additional sorts of books are readily easy to use here.

As this Linux Performance Tools Brendan Gregg, it ends taking place innate one of the favored book Linux Performance Tools Brendan Gregg collections that we have. This is why you remain in the best website to see the incredible books to have.

The Art of Capacity Planning - John Allspaw 2008-09-23
Success on the web is measured by usage and growth. Web-based companies live or die by the ability to scale their infrastructure to accommodate

increasing demand. This book is a hands-on and practical guide to planning for such growth, with many techniques and considerations to help you plan, deploy, and manage web application infrastructure. The Art

of Capacity Planning is written by the manager of data operations for the world-famous photo-sharing site Flickr.com, now owned by Yahoo! John Allspaw combines personal anecdotes from many phases of Flickr's growth with insights from his colleagues in many other industries to give you solid guidelines for measuring your growth, predicting trends, and making cost-effective preparations. Topics include: Evaluating tools for measurement and deployment Capacity analysis and prediction for storage, database, and application servers Designing architectures to easily add and measure capacity Handling sudden spikes Predicting exponential and explosive growth How cloud services such as EC2 can fit into a capacity strategy In this book, Allspaw draws on years of valuable experience,

starting from the days when Flickr was relatively small and had to deal with the typical growth pains and cost/performance trade-offs of a typical company with a Web presence. The advice he offers in The Art of Capacity Planning will not only help you prepare for explosive growth, it will save you tons of grief.

Optimizing Oracle Performance - Cary Millsap 2003-09-16

Oracle system performance inefficiencies often go undetected for months or even years--even under intense scrutiny--because traditional Oracle performance analysis methods and tools are fundamentally flawed. They're unreliable and inefficient. Oracle DBAs and developers are all too familiar with the outlay of time and resources, blown budgets, missed deadlines, and

marginally effective performance fiddling that is commonplace with traditional methods of Oracle performance tuning. In this crucial book, Cary Millsap, former VP of Oracle's System Performance Group, clearly and concisely explains how to use Oracle's response time statistics to diagnose and repair performance problems. Cary also shows how "queueing theory" can be applied to response time statistics to predict the impact of upgrades and other system changes. Optimizing Oracle Performance eliminates the time-consuming, trial-and-error guesswork inherent in most conventional approaches to tuning. You can determine exactly where a system's performance problem is, and with equal importance, where it is not, in just a few minutes--even if the

problem is several years old. Optimizing Oracle Performance cuts a path through the complexity of current tuning methods, and streamlines an approach that focuses on optimization techniques that any DBA can use quickly and successfully to make noticeable--even dramatic--improvements. For example, the one thing database users care most about is response time. Naturally, DBAs focus much of their time and effort towards improving response time. But it is entirely too easy to spend hundreds of hours to improve important system metrics such as hit ratios, average latencies, and wait times, only to find users are unable to perceive the difference. And an expensive hardware upgrade may not help either. It doesn't have to be that way. Technological advances have

added impact, efficiency, measurability, predictive capacity, reliability, speed, and practicality to the science of Oracle performance optimization. *Optimizing Oracle Performance* shows you how to slash the frustration and expense associated with unraveling the true root cause of any type of performance problem, and reliably predict future performance. The price of this essential book will be paid back in hours saved the first time its methods are used.

The Art of UNIX Programming - Eric S. Raymond 2003-09-23

The Art of UNIX Programming poses the belief that understanding the unwritten UNIX engineering tradition and mastering its design patterns will help programmers of all stripes to become better programmers. This

book attempts to capture the engineering wisdom and design philosophy of the UNIX, Linux, and Open Source software development community as it has evolved over the past three decades, and as it is applied today by the most experienced programmers. Eric Raymond offers the next generation of "hackers" the unique opportunity to learn the connection between UNIX philosophy and practice through careful case studies of the very best UNIX/Linux programs.

Hands-On System Programming with Linux - Kaiwan N Billimoria 2018-10-31

Get up and running with system programming concepts in Linux Key Features Acquire insight on Linux system architecture and its programming interfaces Get to grips

with core concepts such as process management, signalling and pthreads. Packed with industry best practices and dozens of code examples. **Book Description** The Linux OS and its embedded and server applications are critical components of today's software infrastructure in a decentralized, networked universe. The industry's demand for proficient Linux developers is only rising with time. **Hands-On System Programming with Linux** gives you a solid theoretical base and practical industry-relevant descriptions, and covers the Linux system programming domain. It delves into the art and science of Linux application programming— system architecture, process memory and management, signaling, timers, pthreads, and file IO. This book goes beyond the use API

X to do Y approach; it explains the concepts and theories required to understand programming interfaces and design decisions, the tradeoffs made by experienced developers when using them, and the rationale behind them. Troubleshooting tips and techniques are included in the concluding chapter. By the end of this book, you will have gained essential conceptual design knowledge and hands-on experience working with Linux system programming interfaces. What you will learn: Explore the theoretical underpinnings of Linux system architecture. Understand why modern OSes use virtual memory and dynamic memory APIs. Get to grips with dynamic memory issues and effectively debug them. Learn key concepts and powerful system APIs related to process management. Effectively perform file IO

and use signaling and timers
Deeply understand multithreading concepts, pthreads APIs, synchronization and scheduling
Who this book is for Hands-On System Programming with Linux is for Linux system engineers, programmers, or anyone who wants to go beyond using an API set to understanding the theoretical underpinnings and concepts behind powerful Linux system programming APIs. To get the most out of this book, you should be familiar with Linux at the user-level logging in, using shell via the command line interface, the ability to use tools such as find, grep, and sort. Working knowledge of the C programming language is required. No prior experience with Linux systems programming is assumed.

Solaris" Performance And Tools:

Dtrace And Mdb Techniques For Solaris 10 And Opensolaris - McDougall Richard 2007-09

Pro .NET Benchmarking - Andrey Akinshin 2019-06-26

Use this in-depth guide to correctly design benchmarks, measure key performance metrics of .NET applications, and analyze results. This book presents dozens of case studies to help you understand complicated benchmarking topics. You will avoid common pitfalls, control the accuracy of your measurements, and improve performance of your software. Author Andrey Akinshin has maintained BenchmarkDotNet (the most popular .NET library for benchmarking) for five years and covers common mistakes that developers usually make in their

benchmarks. This book includes not only .NET-specific content but also essential knowledge about performance measurements which can be applied to any language or platform (common benchmarking methodology, statistics, and low-level features of modern hardware). What You'll Learn Be aware of the best practices for writing benchmarks and performance tests Avoid the common benchmarking pitfalls Know the hardware and software factors that affect application performance Analyze performance measurements Who This Book Is For .NET developers concerned with the performance of their applications

Linux Observability with BPF - David Calavera 2019-11-14

Build your expertise in the BPF virtual machine in the Linux kernel with this practical guide for systems

engineers. You'll not only dive into the BPF program lifecycle but also learn to write applications that observe and modify the kernel's behavior; inject code to monitor, trace, and securely observe events in the kernel; and more. Authors David Calavera and Lorenzo Fontana help you harness the power of BPF to make any computing system more observable. Familiarize yourself with the essential concepts you'll use on a day-to-day basis and augment your knowledge about performance optimization, networking, and security. Then see how it all comes together with code examples in C, Go, and Python. Write applications that use BPF to observe and modify the Linux kernel's behavior on demand Inject code to monitor, trace, and observe events in the kernel in a

secure way—no need to recompile the kernel or reboot the system Explore code examples in C, Go, and Python Gain a more thorough understanding of the BPF program lifecycle

Performance Analysis and Tuning on Modern CPUs - 2020-11-16

Performance tuning is becoming more important than it has been for the last 40 years. Read this book to understand your application's performance that runs on a modern CPU and learn how you can improve it. The 170+ page guide combines the knowledge of many optimization experts from different industries.

C++ Network Programming, Volume 2 - Douglas Schmidt 2002-10-29

Do you need to develop flexible software that can be customized quickly? Do you need to add the power and efficiency of frameworks to your

software? The ADAPTIVE Communication Environment (ACE) is an open-source toolkit for building high-performance networked applications and next-generation middleware. ACE's power and flexibility arise from object-oriented frameworks, used to achieve the systematic reuse of networked application software. ACE frameworks handle common network programming tasks and can be customized using C++ language features to produce complete distributed applications. C++ Network Programming, Volume 2, focuses on ACE frameworks, providing thorough coverage of the concepts, patterns, and usage rules that form their structure. This book is a practical guide to designing object-oriented frameworks and shows developers how to apply frameworks to concurrent networked applications. C++

Networking, Volume 1, introduced ACE and the wrapper facades, which are basic network computing ingredients. Volume 2 explains how frameworks build on wrapper facades to provide higher-level communication services. Written by two experts in the ACE community, this book contains: An overview of ACE frameworks Design dimensions for networked services Descriptions of the key capabilities of the most important ACE frameworks Numerous C++ code examples that demonstrate how to use ACE frameworks C++ Network Programming, Volume 2, teaches how to use frameworks to write networked applications quickly, reducing development effort and overhead. It will be an invaluable asset to any C++ developer working on networked applications.

Camel Design Patterns - Bilgin Ibryam

Java Performance Tuning - Jack Shirazi 2003-01-21

Helps readers eliminate performance problems, covering topics including bottlenecks, profiling tools, strings, algorithms, distributed systems, and servlets.

Lions' Commentary on UNIX 6th Edition with Source Code - John Lions

1996-01-01

For the past 20 years, UNIX insiders have cherished and zealously guarded pirated photocopies of this manuscript, a "hacker trophy" of sorts. Now legal (and legible) copies are available. An international "who's who" of UNIX wizards, including Dennis Ritchie, have contributed essays extolling the merits and importance of this underground classic.

[Interprocess Communications in Linux](#)

- John Shapley Gray 2003

Gray zeroes right in on the key techniques of processes and interprocess communication from primitive communications to the complexities of sockets. The book covers every aspect of UNIX/Linux interprocess communications in sufficient detail to allow experienced programmers to begin writing useful code immediately.

Understanding the Linux Virtual Memory Manager - Mel Gorman 2004

This is an expert guide to the 2.6 Linux Kernel's most important component: the Virtual Memory Manager.

NFS Illustrated - Brent Callaghan 2000

One of the original developers of the NFS and WebNFS offers unique insight into these key technologies, for both

programmers creating and debugging NFS-based applications and network engineers creating new implementations. Readers can gain a deeper understanding of how network file protocols are designed and learn how NFS is implemented on UNIX, Windows NT, Java and web browsers. *Seven NoSQL Databases in a Week* - Xun (Brian) Wu 2018-03-29

A beginner's guide to get you up and running with Cassandra, DynamoDB, HBase, InfluxDB, MongoDB, Neo4j, and Redis Key Features Covers the basics of 7 NoSQL databases and how they are used in the enterprises Quick introduction to MongoDB, DynamoDB, Redis, Cassandra, Neo4j, InfluxDB, and HBase Includes effective techniques for database querying and management Book Description This is the golden age of open source NoSQL

databases. With enterprises having to work with large amounts of unstructured data and moving away from expensive monolithic architecture, the adoption of NoSQL databases is rapidly increasing. Being familiar with the popular NoSQL databases and knowing how to use them is a must for budding DBAs and developers. This book introduces you to the different types of NoSQL databases and gets you started with seven of the most popular NoSQL databases used by enterprises today. We start off with a brief overview of what NoSQL databases are, followed by an explanation of why and when to use them. The book then covers the seven most popular databases in each of these categories: MongoDB, Amazon DynamoDB, Redis, HBase, Cassandra, InfluxDB, and Neo4j. The book doesn't

go into too much detail about each database but teaches you enough to get started with them. By the end of this book, you will have a thorough understanding of the different NoSQL databases and their functionalities, empowering you to select and use the right database according to your needs. What you will learn Understand how MongoDB provides high-performance, high-availability, and automatic scaling Interact with your Neo4j instances via database queries, Python scripts, and Java application code Get familiar with common querying and programming methods to interact with Redis Study the different types of problems Cassandra can solve Work with HBase components to support common operations such as creating tables and reading/writing data Discover data models and work

with CRUD operations using DynamoDB Discover what makes InfluxDB a great choice for working with time-series data Who this book is for If you are a budding DBA or a developer who wants to get started with the fundamentals of NoSQL databases, this book is for you. Relational DBAs who want to get insights into the various offerings of popular NoSQL databases will also find this book to be very useful.

Systems Performance - Brendan Gregg
2014

The Complete Guide to Optimizing Systems Performance Written by the winner of the 2013 LISA Award for Outstanding Achievement in System Administration Large-scale enterprise, cloud, and virtualized computing systems have introduced serious performance challenges. Now,

internationally renowned performance expert Brendan Gregg has brought together proven methodologies, tools, and metrics for analyzing and tuning even the most complex environments. Systems Performance: Enterprise and the Cloud focuses on Linux® and Unix® performance, while illuminating performance issues that are relevant to all operating systems. You'll gain deep insight into how systems work and perform, and learn methodologies for analyzing and improving system and application performance. Gregg presents examples from bare-metal systems and virtualized cloud tenants running Linux-based Ubuntu®, Fedora®, CentOS, and the illumos-based Joyent® SmartOS™ and OmniTI OmniOS®. He systematically covers modern systems performance, including the “traditional” analysis of CPUs,

memory, disks, and networks, and new areas including cloud computing and dynamic tracing. This book also helps you identify and fix the “unknown unknowns” of complex performance: bottlenecks that emerge from elements and interactions you were not aware of. The text concludes with a detailed case study, showing how a real cloud customer issue was analyzed from start to finish. Coverage includes • Modern performance analysis and tuning: terminology, concepts, models, methods, and techniques • Dynamic tracing techniques and tools, including examples of DTrace, SystemTap, and perf • Kernel internals: uncovering what the OS is doing • Using system observability tools, interfaces, and frameworks • Understanding and monitoring

application performance • Optimizing CPUs: processors, cores, hardware threads, caches, interconnects, and kernel scheduling • Memory optimization: virtual memory, paging, swapping, memory architectures, busses, address spaces, and allocators • File system I/O, including caching • Storage devices/controllers, disk I/O workloads, RAID, and kernel I/O • Network-related performance issues: protocols, sockets, interfaces, and physical connections • Performance implications of OS and hardware-based virtualization, and new issues encountered with cloud computing • Benchmarking: getting accurate results and avoiding common mistakes This guide is indispensable for anyone who operates enterprise or cloud environments: system, network,

database, and web admins; developers; and other professionals. For students and others new to optimization, it also provides exercises reflecting Gregg's extensive instructional experience.

Linkers and Loaders - John R. Levine
2000

"I enjoyed reading this useful overview of the techniques and challenges of implementing linkers and loaders. While most of the examples are focused on three computer architectures that are widely used today, there are also many side comments about interesting and quirky computer architectures of the past. I can tell from these war stories that the author really has been there himself and survived to tell the tale." -Guy Steele
Whatever your programming language, whatever

your platform, you probably tap into linker and loader functions all the time. But do you know how to use them to their greatest possible advantage? Only now, with the publication of *Linkers & Loaders*, is there an authoritative book devoted entirely to these deep-seated compile-time and run-time processes. The book begins with a detailed and comparative account of linking and loading that illustrates the differences among various compilers and operating systems. On top of this foundation, the author presents clear practical advice to help you create faster, cleaner code. You'll learn to avoid the pitfalls associated with Windows DLLs, take advantage of the space-saving, performance-improving techniques supported by many modern linkers, make the best use of the

UNIX ELF library scheme, and much more. If you're serious about programming, you'll devour this unique guide to one of the field's least understood topics. Linkers & Loaders is also an ideal supplementary text for compiler and operating systems courses. Features:

- * Includes a linker construction project written in Perl, with project files available for download.
- * Covers dynamic linking in Windows, UNIX, Linux, BeOS, and other operating systems.
- * Explains the Java linking model and how it figures in network applets and extensible Java code.
- * Helps you write more elegant and effective code, and build applications that compile, load, and run more efficiently.

DTrace - Brendan Gregg 2011-03-18
The Oracle Solaris DTrace feature

revolutionizes the way you debug operating systems and applications. Using DTrace, you can dynamically instrument software and quickly answer virtually any question about its behavior. Now, for the first time, there's a comprehensive, authoritative guide to making the most of DTrace in any supported UNIX environment--from Oracle Solaris to OpenSolaris, Mac OS X, and FreeBSD. Written by key contributors to the DTrace community, DTrace teaches by example, presenting scores of commands and easy-to-adapt, downloadable D scripts. These concise examples generate answers to real and useful questions, and serve as a starting point for building more complex scripts. Using them, you can start making practical use of DTrace immediately, whether you're an

administrator, developer, analyst, architect, or support professional. The authors fully explain the goals, techniques, and output associated with each script or command. Drawing on their extensive experience, they provide strategy suggestions, checklists, and functional diagrams, as well as a chapter of advanced tips and tricks. You'll learn how to Write effective scripts using DTrace's D language Use DTrace to thoroughly understand system performance Expose functional areas of the operating system, including I/O, filesystems, and protocols Use DTrace in the application and database development process Identify and fix security problems with DTrace Analyze the operating system kernel Integrate DTrace into source code Extend DTrace with other tools This book will help

you make the most of DTrace to solve problems more quickly and efficiently, and build systems that work faster and more reliably.

Linux Kernel Programming - Kaiwan N Billimoria 2021-03-19

Learn how to write high-quality kernel module code, solve common Linux kernel programming issues, and understand the fundamentals of Linux kernel internals Key Features Discover how to write kernel code using the Loadable Kernel Module framework Explore industry-grade techniques to perform efficient memory allocation and data synchronization within the kernel Understand the essentials of key internals topics such as kernel architecture, memory management, CPU scheduling, and kernel synchronization Book Description Linux

Kernel Programming is a comprehensive introduction for those new to Linux kernel and module development. This easy-to-follow guide will have you up and running with writing kernel code in next-to-no time. This book uses the latest 5.4 Long-Term Support (LTS) Linux kernel, which will be maintained from November 2019 through to December 2025. By working with the 5.4 LTS kernel throughout the book, you can be confident that your knowledge will continue to be valid for years to come. You'll start the journey by learning how to build the kernel from the source. Next, you'll write your first kernel module using the powerful Loadable Kernel Module (LKM) framework. The following chapters will cover key kernel internals topics including Linux kernel architecture, memory

management, and CPU scheduling. During the course of this book, you'll delve into the fairly complex topic of concurrency within the kernel, understand the issues it can cause, and learn how they can be addressed with various locking technologies (mutexes, spinlocks, atomic, and refcount operators). You'll also benefit from more advanced material on cache effects, a primer on lock-free techniques within the kernel, deadlock avoidance (with lockdep), and kernel lock debugging techniques. By the end of this kernel book, you'll have a detailed understanding of the fundamentals of writing Linux kernel module code for real-world projects and products. What you will learnWrite high-quality modular kernel code (LKM framework) for 5.x kernelsConfigure and build a

kernel from sourceExplore the Linux kernel architectureGet to grips with key internals regarding memory management within the kernelUnderstand and work with various dynamic kernel memory alloc/dealloc APIsDiscover key internals aspects regarding CPU scheduling within the kernelGain an understanding of kernel concurrency issuesFind out how to work with key kernel synchronization primitivesWho this book is for This book is for Linux programmers beginning to find their way with Linux kernel development. If you're a Linux kernel and driver developer looking to overcome frequent and common kernel development issues, or understand kernel internals, you'll find plenty of useful information. You'll need a solid foundation of Linux CLI and C

programming before you can jump in.

Linux in Action - David Clinton

2018-08-19

Summary Linux in Action is a task-based tutorial that will give you the skills and deep understanding you need to administer a Linux-based system. This hands-on book guides you through 12 real-world projects so you can practice as you learn. Each chapter ends with a review of best practices, new terms, and exercises. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology You can't learn anything without getting your hands dirty— including Linux. Skills like securing files, folders, and servers, safely installing patches and applications, and managing a network are required for any serious

user, including developers, administrators, and DevOps professionals. With this hands-on tutorial, you'll roll up your sleeves and learn Linux project by project. About the Book Linux in Action guides you through 12 real-world projects, including automating a backup-and-restore system, setting up a private Dropbox-style file cloud, and building your own MediaWiki server. You'll try out interesting examples as you lock in core practices like virtualization, disaster recovery, security, backup, DevOps, and system troubleshooting. Each chapter ends with a review of best practices, new terms, and exercises. What's inside
Setting up a safe Linux environment
Managing secure remote connectivity
Building a system recovery device
Patching and upgrading your system

About the Reader No prior Linux admin experience is required. About the Author David Clinton is a certified Linux Server Professional, seasoned instructor, and author of Manning's bestselling Learn Amazon Web Services in a Month of Lunches. Table of Contents
Welcome to Linux
Linux virtualization: Building a Linux working environment
Remote connectivity: Safely accessing networked machines
Archive management: Backing up or copying entire file systems
Automated administration: Configuring automated offsite backups
Emergency tools: Building a system recovery device
Web servers: Building a MediaWiki server
Networked file sharing: Building a Nextcloud file-sharing server
Securing your web server
Securing network connections: Creating a VPN

or DMZ System monitoring: Working with log files
Sharing data over a private network
Troubleshooting system performance issues
Troubleshooting network issues
Troubleshooting peripheral devices
DevOps tools: Deploying a scripted server environment using Ansible
BPF Performance Tools - Brendan Gregg 2020

Problem-solving in High Performance Computing - Igor Ljubuncic 2015-09-01
Problem-Solving in High Performance Computing: A Situational Awareness Approach with Linux focuses on understanding giant computing grids as cohesive systems. Unlike other titles on general problem-solving or system administration, this book offers a cohesive approach to complex, layered environments,

highlighting the difference between standalone system troubleshooting and complex problem-solving in large, mission critical environments, and addressing the pitfalls of information overload, micro, and macro symptoms, also including methods for managing problems in large computing ecosystems. The authors offer perspective gained from years of developing Intel-based systems that lead the industry in the number of hosts, software tools, and licenses used in chip design. The book offers unique, real-life examples that emphasize the magnitude and operational complexity of high performance computer systems. Provides insider perspectives on challenges in high performance environments with thousands of servers, millions of cores,

distributed data centers, and petabytes of shared data Covers analysis, troubleshooting, and system optimization, from initial diagnostics to deep dives into kernel crash dumps Presents macro principles that appeal to a wide range of users and various real-life, complex problems Includes examples from 24/7 mission-critical environments with specific HPC operational constraints
Linux Kernel Networking - Rami Rosen
2014-02-28

Linux Kernel Networking takes you on a guided in-depth tour of the current Linux networking implementation and the theory behind it. Linux kernel networking is a complex topic, so the book won't burden you with topics not directly related to networking. This book will also not overload you with cumbersome line-by-line code

walkthroughs not directly related to what you're searching for; you'll find just what you need, with in-depth explanations in each chapter and a quick reference at the end of each chapter. Linux Kernel Networking is the only up-to-date reference guide to understanding how networking is implemented, and it will be indispensable in years to come since so many devices now use Linux or operating systems based on Linux, like Android, and since Linux is so prevalent in the data center arena, including Linux-based virtualization technologies like Xen and KVM.
Embedded Linux Development Using Yocto Project Cookbook - Alex González
2018-01-25
Over 79 hands-on recipes for professional embedded Linux developers to optimize and boost

their Yocto Project know-how Key Features Optimize your Yocto setup to speed up development and debug build issues Use what is quickly becoming the standard embedded Linux product builder framework—the Yocto Project Recipe-based implementation of best practices to optimize your Linux system Book Description The Yocto Project has become the de facto distribution build framework for reliable and robust embedded systems with a reduced time to market. You'll get started by working on a build system where you set up Yocto, create a build directory, and learn how to debug it. Then, you'll explore everything about the BSP layer, from creating a custom layer to debugging device tree issues. In addition to this, you'll learn how to add a new software layer, packages, data,

scripts, and configuration files to your system. You will then cover topics based on application development, such as using the Software Development Kit and how to use the Yocto project in various development environments. Toward the end, you will learn how to debug, trace, and profile a running system. This second edition has been updated to include new content based on the latest Yocto release. What you will learn Optimize your Yocto Project setup to speed up development and debug build issues Use Docker containers to build Yocto Project-based systems Take advantage of the user-friendly Toaster web interface to the Yocto Project build system Build and debug the Linux kernel and its device trees Customize your root filesystem with already-supported and

new Yocto packages Optimize your production systems by reducing the size of both the Linux kernel and root filesystems Explore the mechanisms to increase the root filesystem security Understand the open source licensing requirements and how to comply with them when cohabiting with proprietary programs Create recipes, and build and run applications in C, C++, Python, Node.js, and Java Who this book is for If you are an embedded Linux developer with the basic knowledge of Yocto Project, this book is an ideal way to broaden your knowledge with recipes for embedded development.
Solaris Internals - Richard McDougall
2006-07-10

"The Solaris™Internals volumes are simply the best and most comprehensive treatment of the

Solaris (and OpenSolaris) Operating Environment. Any person using Solaris--in any capacity--would be remiss not to include these two new volumes in their personal library. With advanced observability tools in Solaris (likeDTrace), you will more often find yourself in what was previously unchartable territory. Solaris™ Internals, Second Edition, provides us a fantastic means to be able to quickly understand these systems and further explore the Solaris architecture--especially when coupled with OpenSolaris source availability." --Jarod Jenson, chief systems architect, Aeysis "The Solaris™ Internals volumes by Jim Mauro and Richard McDougall must be on your bookshelf if you are interested in in-depth knowledge of Solaris operating system internals

and architecture. As a senior Unix engineer for many years, I found the first edition of Solaris™ Internals the only fully comprehensive source for kernel developers, systems programmers, and systems administrators. The new second edition, with the companion performance and debugging book, is an indispensable reference set, containing many useful and practical explanations of Solaris and its underlying subsystems, including tools and methods for observing and analyzing any system running Solaris 10 or OpenSolaris." --Marc Strahl, senior UNIX engineer Solaris™ Internals, Second Edition, describes the algorithms and data structures of all the major subsystems in the Solaris 10 and OpenSolaris kernels. The text has been extensively revised

since the first edition, with more than 600 pages of new material. Integrated Solaris tools and utilities, including DTrace, MDB, kstat, and the process tools, are used throughout to illustrate how the reader can observe the Solaris kernel in action. The companion volume, Solaris™ Performance and Tools, extends the examples contained here, and expands the scope to performance and behavior analysis. Coverage includes: Virtual and physical memory Processes, threads, and scheduling File system framework and UFS implementation Networking: TCP/IP implementation Resource management facilities and zones The Solaris™ Internals volumes make a superb reference for anyone using Solaris 10 and OpenSolaris. *Linux Internals Simplified* - Haris

Kundan Prasad 2020-05-02

"Linux internals simplified" is a book which discusses the basics of Linux kernel internals in a code driven approach. It picks the major subsystems of the kernel which are important, and tries to simplify its internal working and data structures. As such, this book is aimed at engineers who wish to start learning about the Linux kernel. This book starts with the basic steps to acquire the Linux kernel code. It then shows ways of customizing the build options and lastly kernel compilation. Next it looks at a number of hacking tools which will help one to debug and trace in a live Linux system. Practical examples of ftrace, kprobes and crash tool are discussed. These tools are useful in trying to understand the way the

Linux system works. Chapter 3 discusses the details of a running process in a Linux system. It touches topics such as address spaces of a running process, user and kernel spaces, system calls, Linux process descriptor, Linux process creation, and so on. This chapter builds a foundation of a program in execution in the Linux system. Once the reader knows about the running processes, chapter 4 discusses about the Linux process scheduling subsystem. This chapter discusses different data structures and code paths of the Linux scheduler, which controls the scheduling of processes in the Linux system. Chapter 5 discusses Interrupts, which play a significant role in the Linux operating system. The chapter discusses edge and level triggered interrupts, interrupt

handlers and their registration, shared interrupt handlers, and so on. It also shows the ftrace of the `do_irq` function. Chapter 6 discusses the signal subsystem. It starts with a little introduction of the design of the signal subsystem. It then traces the code execution of delivering and handling of signals in the Linux kernel. The chapter then discusses signal overloading and how it is performed, while exploring the kernel code which handles this. Chapter 7 covers Linux synchronization primitives, and why they are needed. It shows the detailed implementation of primitives like atomic variables, spinlocks, semaphores and mutexes in the Linux kernel. Chapter 8 discusses various ways of Linux kernel memory allocation. It discusses Buddy

allocator, Resource map allocator and Slab allocator. It discusses various APIs used for these allocators (`alloc_page/s`, `kmem_cache_alloc`, `kmalloc` etc.). It also discusses how user space `malloc` results in memory allocation in the Linux kernel. Chapter 9 discusses the Linux dynamic modules, Linux character driver framework, internal functions which are used while creating a character driver, UDEV events and IOCTL interface. It also discusses Linux device model. It discusses example of bus, device and `device_driver` components. It illustrates device model when used in PCI BUS. Chapter 10 covers the subsystem related to block IOs. It starts with an introduction of filesystem and its purpose. It then traces the path an IO takes, right

from the "write()" system call, to the moment it gets written to the disk. The chapter covers basic data structures and design elements while going down the IO stack.

Understanding Software Dynamics -

Richard Sites 2021-10-15

Troubleshoot and Optimize Complex, Time-Constrained Software From mobile and cloud apps to video games to driverless vehicle control, more and more software is time-constrained: It must deliver reliable results seamlessly, consistently, and virtually instantaneously. If it doesn't, customers are unhappy--and sometimes lives are put at risk. When time-constrained software underperforms or fails, software professionals must quickly identify and address the root causes. This is difficult and, historically, few

tools have been available to help. In *Understanding Software Dynamics*, performance expert Richard L. Sites tackles the problem head on, offering expert methods and advanced tools for understanding complex time-constrained software dynamics, improving reliability, and troubleshooting challenging performance problems. Sites draws on several decades of experience pioneering software performance optimization, as well as extensive experience teaching graduate-level developers. He introduces principles and techniques for use in any environment, from embedded devices to datacenters, illuminating them with examples based on x86 or ARM processors running Linux and linked by Ethernet. He also guides readers through building and applying a

powerful, new, extremely low-overhead open-source software tool, KUtrace, to precisely trace executions on every CPU core. Using insights gleaned from this tool, readers can apply nuanced solutions--not merely brute-force techniques such as turning off caches or cores. Measure and address issues associated with CPUs, memory, disk/SSD, networks, and their interactions Fix programs that are always too slow, and those that sometimes lag for no apparent reason Design useful observability, logging, and time-stamping capabilities into your code Reason more effectively about performance data to see why reality differs from expectations Identify problems such as excess execution, slow instruction execution, waiting for resources, and software locks Understanding Software

Dynamics will be valuable to experienced software professionals, including application and OS developers, hardware and system architects, real-time system designers, and game developers, as well as advanced students. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Linux Network Administrator's Guide - Olaf Kirch 2000

This introduction to networking on Linux now covers firewalls, including the use of ipchains and Netfilter, masquerading, and accounting. Other new topics in this second edition include Novell (NCP/IPX) support and INN (news administration).

Implementing Service Level Objectives - Alex Hidalgo 2020-08-05

Although service-level objectives (SLOs) continue to grow in importance, there's a distinct lack of information about how to implement them. Practical advice that does exist usually assumes that your team already has the infrastructure, tooling, and culture in place. In this book, recognized SLO expert Alex Hidalgo explains how to build an SLO culture from the ground up. Ideal as a primer and daily reference for anyone creating both the culture and tooling necessary for SLO-based approaches to reliability, this guide provides detailed analysis of advanced SLO and service-level indicator (SLI) techniques. Armed with mathematical models and statistical knowledge to help you get the most out of an SLO-based approach, you'll learn how to build

systems capable of measuring meaningful SLIs with buy-in across all departments of your organization. Define SLIs that meaningfully measure the reliability of a service from a user's perspective Choose appropriate SLO targets, including how to perform statistical and probabilistic analysis Use error budgets to help your team have better discussions and make better data-driven decisions Build supportive tooling and resources required for an SLO-based approach Use SLO data to present meaningful reports to leadership and your users

Linux Administration Handbook - Evi Nemeth 2006-10-30

"As this book shows, Linux systems are just as functional, secure, and reliable as their proprietary counterparts. Thanks to the ongoing

efforts of thousands of Linux developers, Linux is more ready than ever for deployment at the frontlines of the real world. The authors of this book know that terrain well, and I am happy to leave you in their most capable hands.” –Linus Torvalds “The most successful sysadmin book of all time—because it works!” –Rik Farrow, editor of ;login: “This book clearly explains current technology with the perspective of decades of experience in large-scale system administration. Unique and highly recommended.” –Jonathan Corbet, cofounder, LWN.net “Nemeth et al. is the overall winner for Linux administration: it’s intelligent, full of insights, and looks at the implementation of concepts.” –Peter Salus, editorial director, Matrix.net Since 2001, Linux Administration Handbook has

been the definitive resource for every Linux® system administrator who must efficiently solve technical problems and maximize the reliability and performance of a production environment. Now, the authors have systematically updated this classic guide to address today’s most important Linux distributions and most powerful new administrative tools. The authors spell out detailed best practices for every facet of system administration, including storage management, network design and administration, web hosting, software configuration management, performance analysis, Windows interoperability, and much more. Sysadmins will especially appreciate the thorough and up-to-date discussions of such difficult topics such as DNS, LDAP, security, and the

management of IT service organizations. Linux® Administration Handbook, Second Edition, reflects the current versions of these leading distributions: Red Hat® Enterprise Linux® Fedora™ Core SUSE® Linux Enterprise Debian® GNU/Linux Ubuntu® Linux Sharing their war stories and hard-won insights, the authors capture the behavior of Linux systems in the real world, not just in ideal environments. They explain complex tasks in detail and illustrate these tasks with examples drawn from their extensive hands-on experience.

Hello, Startup - Yevgeniy Brikman
2015-10-21

This book is the "Hello, World" tutorial for building products, technologies, and teams in a startup environment. It's based on the experiences of the author, Yevgeniy

(Jim) Brikman, as well as interviews with programmers from some of the most successful startups of the last decade, including Google, Facebook, LinkedIn, Twitter, GitHub, Stripe, Instagram, AdMob, Pinterest, and many others. Hello, Startup is a practical, how-to guide that consists of three parts: Products, Technologies, and Teams. Although at its core, this is a book for programmers, by programmers, only Part II (Technologies) is significantly technical, while the rest should be accessible to technical and non-technical audiences alike. If you're at all interested in startups—whether you're a programmer at the beginning of your career, a seasoned developer bored with large company politics, or a manager looking to motivate your

engineers—this book is for you.
Data Plane Development Kit (DPDK) -
Heqing Zhu 2020-11-19

This book brings together the insights and practical experience of some of the most experienced Data Plane Development Kit (DPDK) technical experts, detailing the trend of DPDK, data packet processing, hardware acceleration, packet processing and virtualization, as well as the practical application of DPDK in the fields of SDN, NFV, and network storage. The book also devotes many chunks to exploring various core software algorithms, the advanced optimization methods adopted in DPDK, detailed practical experience, and the guides on how to use DPDK.

Container Security - Liz Rice
2020-04-06

To facilitate scalability and resilience, many organizations now run applications in cloud native environments using containers and orchestration. But how do you know if the deployment is secure? This practical book examines key underlying technologies to help developers, operators, and security professionals assess security risks and determine appropriate solutions. Author Liz Rice, Chief Open Source Officer at Isovalent, looks at how the building blocks commonly used in container-based systems are constructed in Linux. You'll understand what's happening when you deploy containers and learn how to assess potential security risks that could affect your deployments. If you run container applications with kubectl or docker and use Linux

command-line tools such as ps and grep, you're ready to get started. Explore attack vectors that affect container deployments Dive into the Linux constructs that underpin containers Examine measures for hardening containers Understand how misconfigurations can compromise container isolation Learn best practices for building container images Identify container images that have known software vulnerabilities Leverage secure connections between containers Use security tooling to prevent attacks on your deployment

Systems Performance - Brendan Gregg
2020-10-30

Systems performance analysis and tuning lead to a better end-user experience and lower costs, especially for cloud computing environments that charge by the OS

instance. Systems Performance, 2nd Edition covers concepts, strategy, tools, and tuning for operating systems and applications, using Linux-based operating systems as the primary example. World-renowned systems performance expert Brendan Gregg summarizes relevant operating system, hardware, and application theory to quickly get professionals up to speed even if they've never analyzed performance before, and to refresh and update advanced readers' knowledge. Gregg illuminates the latest tools and techniques, including extended BPF, showing how to get the most out of your systems in cloud, web, and large-scale enterprise environments. He covers these and other key topics: Hardware, kernel, and application internals, and how they perform Methodologies

for rapid performance analysis of complex systems Optimizing CPU, memory, file system, disk, and networking usage Sophisticated profiling and tracing with perf, Ftrace, and BPF (BCC and bpftrace) Performance challenges associated with cloud computing hypervisors Benchmarking more effectively Fully updated for current Linux operating systems and environments, Systems Performance, 2nd Edition addresses issues that apply to any computer system. The book will be a go-to reference for many years to come and recommended reading at many tech companies, like its predecessor first edition.

Factory Man - Beth Macy 2014-07-15
The instant New York Times bestseller about one man's battle to save hundreds of jobs by demonstrating the

greatness of American business. The Bassett Furniture Company was once the world's biggest wood furniture manufacturer. Run by the same powerful Virginia family for generations, it was also the center of life in Bassett, Virginia. But beginning in the 1980s, the first waves of Asian competition hit, and ultimately Bassett was forced to send its production overseas. One man fought back: John Bassett III, a shrewd and determined third-generation factory man, now chairman of Vaughan-Bassett Furniture Co, which employs more than 700 Virginians and has sales of more than \$90 million. In FACTORY MAN, Beth Macy brings to life Bassett's deeply personal furniture and family story, along with a host of characters from an industry that was as cutthroat as

it was colorful. As she shows how he uses legal maneuvers, factory efficiencies, and sheer grit and cunning to save hundreds of jobs, she also reveals the truth about modern industry in America.

BPF Performance Tools - Brendan Gregg
2019-11-06

BPF and related observability tools give software professionals and students alike unprecedented visibility into software, helping them analyze operating system and application performance, troubleshoot code, and strengthen security. BPF Performance Tools: Linux System and Application Observability is the industry's most comprehensive guide to using these tools for observability. Brendan Gregg, author of the industry's definitive guide to system performance, introduces

powerful new methods and tools for doing analysis that leads to more robust, reliable, and safer code. This authoritative guide: Explores a wide spectrum of software and hardware targets Thoroughly covers open source BPF tools from the Linux Foundation iovisor project's bcc and bpfftrace repositories Summarizes performance engineering and kernel internals you need to understand Provides and discusses 150+ bpfftrace tools, including 80 written specifically for this book: tools you can run as-is, without programming - or customize and develop further, using diverse interfaces and the bpfftrace front-end Students will learn how to use BPF (eBPF) tracing tools to analyze CPUs, memory, disks, file systems, networking, languages, applications, containers,

hypervisors, security, and the Linux kernel. Students will move from basic to advanced tools and techniques, producing new metrics, stack traces, custom latency histograms, and more. It's like having a superpower: with Gregg's guidance and tools, students can analyze virtually everything that impacts system performance, so they can improve virtually any Linux operating system or application. Deeper, more in-depth coverage than any other eBPF resource Quickly analyze everything that impacts Linux system performance: ask questions and get fast answers in production environments Learn by example, with tools you can use to find performance wins and then customize for even more power Covers invaluable, in-demand technology: eBPF was the subject of over two dozen talks at the recent

Linux Plumbers developer's conference Downloadable source code includes 80+ new BPF analysis tools created for this book Use BPF/eBPF tracing and observability tools to improve system performance, reduce costs, resolve software issues, and gain unprecedented visibility into running systems Deeper, more in-depth coverage than any other eBPF resource Quickly analyze everything that impacts Linux system performance: ask questions and get fast answers in production environments Learn by example, with tools you can use to find performance wins and then customize for even more power Covers invaluable, in-demand technology: eBPF was the subject of over two dozen talks at the recent Linux Plumbers developer's conference Downloadable source code includes 80+

new BPF analysis tools created for this book

Is Parallel Programming Hard - Paul E. McKenney 2015-06-13

BPF Performance Tools - Brendan Gregg 2019-11-27

BPF and related observability tools give software professionals unprecedented visibility into software, helping them analyze operating system and application performance, troubleshoot code, and strengthen security. **BPF Performance Tools: Linux System and Application Observability** is the industry's most comprehensive guide to using these tools for observability. Brendan Gregg, author of the industry's definitive guide to system performance, introduces powerful new methods and tools for doing analysis

that leads to more robust, reliable, and safer code. This authoritative guide: Explores a wide spectrum of software and hardware targets Thoroughly covers open source BPF tools from the Linux Foundation iovisor project's bcc and bpftrace repositories Summarizes performance engineering and kernel internals you need to understand Provides and discusses 150+ bpftrace tools, including 80 written specifically for this book: tools you can run as-is, without programming – or customize and develop further, using diverse interfaces and the bpftrace front-end You'll learn how to use BPF (eBPF) tracing tools to analyze CPUs, memory, disks, file systems, networking, languages, applications, containers, hypervisors, security, and the Linux kernel. You'll move

from basic to advanced tools and techniques, producing new metrics, stack traces, custom latency histograms, and more. It's like having a superpower: with Gregg's guidance and tools, you can analyze virtually everything that impacts system performance, so you can improve virtually any Linux operating system or application.

Text Analytics with Python - Dipanjan Sarkar 2019-05-21

Leverage Natural Language Processing (NLP) in Python and learn how to set up your own robust environment for performing text analytics. This second edition has gone through a major revamp and introduces several significant changes and new topics based on the recent trends in NLP. You'll see how to use the latest state-of-the-art frameworks in NLP,

coupled with machine learning and deep learning models for supervised sentiment analysis powered by Python to solve actual case studies. Start by reviewing Python for NLP fundamentals on strings and text data and move on to engineering representation methods for text data, including both traditional statistical models and newer deep learning-based embedding models. Improved techniques and new methods around parsing and processing text are discussed as well. Text summarization and topic models have been overhauled so the book showcases how to build, tune, and interpret topic models in the context of an interest dataset on NIPS conference papers. Additionally, the book covers text similarity techniques with a real-world example of movie

recommenders, along with sentiment analysis using supervised and unsupervised techniques. There is also a chapter dedicated to semantic analysis where you'll see how to build your own named entity recognition (NER) system from scratch. While the overall structure of the book remains the same, the entire code base, modules, and chapters has been updated to the latest Python 3.x release. What You'll Learn • Understand NLP and text syntax, semantics and structure•

Discover text cleaning and feature engineering• Review text classification and text clustering • Assess text summarization and topic models• Study deep learning for NLP Who This Book Is For IT professionals, data analysts, developers, linguistic experts, data scientists and engineers and basically anyone with a keen interest in linguistics, analytics and generating insights from textual data.