

Book Reviews

ELIZABETH ZWICKY, RIK FARROW, AND MARK LAMOURINE

Numbersense: How to Use Big Data to Your Advantage

Kaiser Fung

McGraw Hill, 2013. 218 pages

ISBN 978-0-07-179966-9

Reviewed by Elizabeth Zwicky

This is a fine book about thinking about numbers, only moderately connected to big data. If you really want to know about big data, you'd be better off with one of the recently reviewed data analysis books; only some of the examples here deal directly with big data issues. On the other hand, if you would like a better idea of how data analysts work and how the news is lying to you, without too many actual numbers, this is a nice start. Seriously, a lot of data analysis is more about skepticism than about numerical manipulation more complex than addition and subtraction, and this kind of introduction will move you towards the right thought habits.

Hyperbole and a Half: Unfortunate Situations, Flawed Coping Mechanisms, Mayhem, and Other Things That Happened

Allie Brosh

Simon and Schuster, 2013. 371 pages

ISBN 978-1-4516-6617-5

Reviewed by Elizabeth Zwicky

This is a purely non-technical book, based on a Web comic. It is one of the funniest descriptions of dog ownership ever, and Allie Brosh is one of the few people who can write about depression in ways that are both funny and true. You don't need to take my word for it, you can just go and search for it. In book format you get more text and no video, but the video is not actually required, while the text is excellent.

Although I read the web comic, I had missed some of these, including some laugh out loud stories. Go get a copy so that you can press it on friends you can't email links to, or just so that when your Internet goes out you can while away the time reading and remembering that every so often the Internet brings us marvelous creators, moments of pure hilarity, and utter poignancy. Or, if you like, so that you know that yours is not the stupidest dog ever.

Systems Performance for Enterprise and Cloud

Brendan Gregg

Prentice Hall, 2013. 735 pages

ISBN 978-0-13-339009-4

Reviewed by Rik Farrow

Although Brendan's book's title refers to performance, the book could just as easily have been called troubleshooting Linux and Solaris systems. And by systems, I do mean everything from caches to distributed services. Brendan writes clearly, leaves nothing out, and is well organized but never boring.

The book begins with four chapters that provide the necessary background for the following eight chapters. Brendan explains, with examples, key concepts, methodology, terminology, kernel internals for the performance analyst, and tools. He uses analogy well: for example, when converting time scales for system latency into human-understandable scales, where if a clock tick is one second, L3 cache access takes 43 seconds, and DRAM access six minutes. The book is full of analogies like this that make the writing easier to comprehend.

The next eight chapters cover particular topic areas, like applications, memory, CPUs, network, and cloud computing, in detail. Brendan designed the book so that it can be used as a reference, and he attempts to future-proof it with the focus on methodology and necessary background. The final chapter contains a lengthy troubleshooting session showing how Brendan uses the methodology that he has described in real life. He ends with a quote from Niels Bohr: "An expert is a person who has made all the mistakes that can be made in a very narrow field." I believe this reflects Brendan's attitude well, in that he speaks from experience and does not talk down to his readers.

There are seven appendices, starting with performance tools for Linux and Solaris and including one of DTrace one-liners. Throughout the book, Brendan shows how tools for both Linux and Solaris-related operating systems are used to display utilization, capacity, saturation, and errors, part of his USE methodology that forms an early step in troubleshooting performance issues.

I found Brendan's writing pleasant to read. The attention to detail is great, and I noticed no typos or mistakes in examples, which was also a real pleasure. I can recommend this book to system designers, system administrators, and programmers because all three groups will benefit by better understanding,

and being able to measure, the many subsystems that are important in systems performance. Although this is not a beginner's book, an intermediate to advanced practitioner will get a lot of benefit from reading it, in whole or in part.

RabbitMQ In Action

Alvaro Videla and Jason J. W. Williams
Manning Publications, 2012. 287 pages
ISBN 978-1935182979

Reviewed by Mark Lamourine

RabbitMQ in Action is subtitled "Distributed Messaging for Everyone," and while that's hyperbolic, it's not without a grain of truth. As the timeline in the first chapter notes, for decades, enterprise quality messaging services have been the purview of a handful of commercial providers. In the past 15 years or so, a handful of attempts have been made to bring an open messaging service to distributed applications with mixed results. RabbitMQ and AMQP (Advanced Message Queuing Protocol) show a lot of promise.

I appreciated the context that the authors provide in the introductory chapters. The verbal and graphical timelines for the inception and development of computer messaging are rooted in the financial industry of the early 1980s and carry through

an attempt at unification in the JMS standard to the emergence of AMQP from the same world of financial applications in the mid-2000s. The authors also explain how the use of messaging technology improves the robustness and flexibility of distributed applications. This is still the first chapter. The last two pages cover installation and initial configuration of the service.

In the remainder of the book, the authors treat service configuration, clustering, and failover and message store persistence (making sure messages en route aren't lost if the node they're on fails). Several chapters relate to application development, focusing on design patterns and writing for failure. RabbitMQ is written in Erlang but, never fear, all of the code for the examples is in common application languages like Python and PHP. In the final chapters, Videla and Williams touch on managing, monitoring, and securing the RabbitMQ service.

The authors write with a clear flowing style that packs a lot of content into an unimposing book without feeling dense or unnecessarily academic. Videla and Williams have created an accessible introduction to messaging technology in general and to AMQP and RabbitMQ in particular. If you're considering writing a modern distributed application, whether for internal business processes or Web applications, *RabbitMQ in Action* is a good place to start.

USENIX Board of Directors

Communicate directly with the USENIX Board of Directors by writing to board@usenix.org.

PRESIDENT

Margo Seltzer, *Harvard University*
margo@usenix.org

VICE PRESIDENT

John Arrasjid, *VMware*
johna@usenix.org

SECRETARY

Carolyn Rowland, *National Institute of Standards and Technology*
carolyn@usenix.org

TREASURER

Brian Noble, *University of Michigan*
noble@usenix.org

DIRECTORS

David Blank-Edelman, *Northeastern University*
dnb@usenix.org

Sasha Fedorova, *Simon Fraser University*
sasha@usenix.org

Niels Provos, *Google*
niels@usenix.org

Dan Wallach, *Rice University*
dwallach@usenix.org

EXECUTIVE DIRECTOR

Casey Henderson
casey@usenix.org