
Hbase Administration Cookbook

Google BigQuery Analytics
Elasticsearch Server
Dive into the Future of Infrastructure
HBase in Action
Distributed Data at Web Scale
Intelligent Systems and Applications
Proceedings of the International Computer
Symposium (ICS) Held at Taichung, Taiwan,
December 12 - 14, 2014
Perform Fast Analytics on Fast Data
Knowledge Graphs and Big Data Processing
2019 Supplement Covering ChessBase 13, 14 &
15
ChessBase Complete
Big Data Analytics with R and Hadoop
Microsoft Azure Essentials - Fundamentals of
Azure
Hadoop Beginner's Guide
Introduction to Hadoop, Spark, and Machine-
Learning
Apache Hadoop 3 Quick Start Guide
OpenStack Operations Guide
Mastering Hadoop 3
Hbase Administration Cookbook

Service Worker Development Cookbook
Second Edition
Big Data For Dummies
Learn about big data processing and analytics
Hadoop: The Definitive Guide
SQL Server 2016 Reporting Services Cookbook
How 45 Successful Companies Used Big Data
Analytics to Deliver Extraordinary Results
Set Up and Manage Your OpenStack Cloud
Distributed Process Coordination
Hadoop MapReduce Cookbook
A Guidebook for Successful Development and
Design
Pentaho Data Integration Cookbook
Hadoop 2.x Administration Cookbook
Covers MongoDB version 3.0
SQL at Any Scale, on Any Storage, in Any
Environment
HBase Essentials
The Definitive Guide
Presto: The Definitive Guide
ZooKeeper
MongoDB in Action

*Hbase
Administration
Cookbook*

*Downloaded
from
tafayor.com
by guest*

KEMP SANTOS

Google BigQuery
Analytics Packt
Publishing Ltd

This book presents the proceedings of the International Computer Symposium 2014 (ICS 2014), held at Tunghai University, Taichung, Taiwan in December. ICS is a biennial

symposium founded in 1973 and offers a platform for researchers, educators and professionals to exchange their discoveries and practices, to share research experiences and to discuss potential new trends in the ICT industry. Topics covered in the ICS 2014 workshops include: algorithms and computation theory; artificial intelligence and fuzzy systems; computer architecture, embedded systems, SoC and VLSI/EDA; cryptography and information security; databases, data mining, big data and information retrieval; mobile computing, wireless communications and vehicular technologies; software engineering and programming

languages; healthcare and bioinformatics, among others. There was also a workshop on information technology innovation, industrial application and the Internet of Things. ICS is one of Taiwan's most prestigious international IT symposiums, and this book will be of interest to all those involved in the world of information technology.

Elasticsearch Server
"O'Reilly Media, Inc."
Imagine what you could do if scalability wasn't a problem. With this hands-on guide, you'll learn how the Cassandra database management system handles hundreds of terabytes of data while remaining highly available across multiple data centers.

This expanded second edition—updated for Cassandra 3.0—provides the technical details and practical examples you need to put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra’s non-relational design, with special attention to data modeling. If you’re a developer, DBA, or application architect looking to solve a database scaling issue or future-proof your application, this guide helps you harness Cassandra’s speed and flexibility. Understand Cassandra’s distributed and decentralized structure Use the Cassandra Query

Language (CQL) and cqlsh—the CQL shell
 Create a working data model and compare it with an equivalent relational model
 Develop sample applications using client drivers for languages including Java, Python, and Node.js
 Explore cluster topology and learn how nodes exchange data
 Maintain a high level of performance in your cluster
 Deploy Cassandra on site, in the Cloud, or with Docker
 Integrate Cassandra with Spark, Hadoop, Elasticsearch, Solr, and Lucene
[Dive into the Future of Infrastructure](#) "O'Reilly Media, Inc."
 Hbase Administration Cookbook
 Packt Publishing Ltd
HBase in Action
 Springer Nature
 Find the right big data

solution for your business or organization. Big data management is one of the major challenges facing business, industry, and not-for-profit organizations. Data sets such as customer transactions for a mega-retailer, weather patterns monitored by meteorologists, or social network activity can quickly outpace the capacity of traditional data management tools. If you need to develop or manage big data solutions, you'll appreciate how these four experts define, explain, and guide you through this new and often confusing concept. You'll learn what it is, why it matters, and how to choose and implement solutions that work. Effectively managing

big data is an issue of growing importance to businesses, not-for-profit organizations, government, and IT professionals. Authors are experts in information management, big data, and a variety of solutions. Explains big data in detail and discusses how to select and implement a solution, security concerns to consider, data storage and presentation issues, analytics, and much more. Provides essential information in a no-nonsense, easy-to-understand style that is empowering. Big Data For Dummies cuts through the confusion and helps you take charge of big data solutions for your organization. *Distributed Data at Web Scale* John Wiley

& Sons

HBase is a remarkable tool for indexing mass volumes of data, but getting started with this distributed database and its ecosystem can be daunting. With this hands-on guide, you'll learn how to architect, design, and deploy your own HBase applications by examining real-world solutions. Along with HBase principles and cluster deployment guidelines, this book includes in-depth case studies that demonstrate how large companies solved specific use cases with HBase. Authors Jean-Marc Spaggiari and Kevin O'Dell also provide draft solutions and code examples to help you implement your own versions of those use cases, from

master data management (MDM) and document storage to near real-time event processing. You'll also learn troubleshooting techniques to help you avoid common deployment mistakes. Learn exactly what HBase does, what its ecosystem includes, and how to set up your environment Explore how real-world HBase instances were deployed and put into production Examine documented use cases for tracking healthcare claims, digital advertising, data management, and product quality Understand how HBase works with tools and techniques such as Spark, Kafka, MapReduce, and the Java API Learn how to identify the causes and understand the

consequences of the most common HBase issues

Intelligent Systems and Applications

Apress

How to effectively use BigQuery, avoid common mistakes, and execute sophisticated queries against large datasets Google BigQuery Analytics is the perfect guide for business and data analysts who want the latest tips on running complex queries and writing code to communicate with the BigQuery API. The book uses real-world examples to demonstrate current best practices and techniques, and also explains and demonstrates streaming ingestion, transformation via Hadoop in Google Compute engine,

AppEngine datastore integration, and using GViz with Tableau to generate charts of query results. In addition to the mechanics of BigQuery, the book also covers the architecture of the underlying Dremel query engine, providing a thorough understanding that leads to better query results. Features a companion website that includes all code and data sets from the book Uses real-world examples to explain everything analysts need to know to effectively use BigQuery Includes web application examples coded in Python

Proceedings of the International Computer Symposium (ICS) Held at Taichung,

Taiwan, December 12 - 14, 2014

"O'Reilly Media, Inc."

Data is arriving faster than you can process it and the overall volumes keep growing at a rate that keeps you awake at night.

Hadoop can help you tame the data beast.

Effective use of Hadoop however requires a mixture of programming, design, and system administration skills.

"Hadoop Beginner's Guide" removes the mystery from Hadoop, presenting Hadoop and related technologies with a focus on building working systems and getting the job done, using cloud services to do so when it makes sense.

From basic concepts and initial setup through developing applications and

keeping the system running as the data grows, the book gives the understanding needed to effectively use Hadoop to solve real world problems. Starting with the basics of installing and configuring Hadoop, the book explains how to develop applications, maintain the system, and how to use additional products to integrate with other systems. While learning different ways to develop applications to run on Hadoop the book also covers tools such as Hive, Sqoop, and Flume that show how Hadoop can be integrated with relational databases and log collection. In addition to examples on Hadoop clusters on Ubuntu uses of cloud services such as Amazon, EC2 and

Elastic MapReduce are covered.

Perform Fast Analytics on Fast Data Packt Publishing Ltd

A fast paced guide that will help you learn about Apache Hadoop 3 and its ecosystem Key Features Set up, configure and get started with Hadoop to get useful insights from large data sets Work with the different components of Hadoop such as MapReduce, HDFS and YARN Learn about the new features introduced in Hadoop 3 Book Description Apache Hadoop is a widely used distributed data platform. It enables large datasets to be efficiently processed instead of using one large computer to store and process the data. This book will get you

started with the Hadoop ecosystem, and introduce you to the main technical topics, including MapReduce, YARN, and HDFS. The book begins with an overview of big data and Apache Hadoop. Then, you will set up a pseudo Hadoop development environment and a multi-node enterprise Hadoop cluster. You will see how the parallel programming paradigm, such as MapReduce, can solve many complex data processing problems. The book also covers the important aspects of the big data software development lifecycle, including quality assurance and control, performance, administration, and monitoring. You will then learn about the Hadoop ecosystem,

and tools such as Kafka, Sqoop, Flume, Pig, Hive, and HBase. Finally, you will look at advanced topics, including real time streaming using Apache Storm, and data analytics using Apache Spark. By the end of the book, you will be well versed with different configurations of the Hadoop 3 cluster. What you will learn Store and analyze data at scale using HDFS, MapReduce and YARN Install and configure Hadoop 3 in different modes Use Yarn effectively to run different applications on Hadoop based platform Understand and monitor how Hadoop cluster is managed Consume streaming data using Storm, and then analyze it using Spark

Explore Apache Hadoop ecosystem components, such as Flume, Sqoop, HBase, Hive, and Kafka Who this book is for Aspiring Big Data professionals who want to learn the essentials of Hadoop 3 will find this book to be useful. Existing Hadoop users who want to get up to speed with the new features introduced in Hadoop 3 will also benefit from this book. Having knowledge of Java programming will be an added advantage.

Knowledge Graphs and Big Data Processing John Wiley & Sons

Big Data Analytics with R and Hadoop is a tutorial style book that focuses on all the powerful big data tasks that can be achieved by integrating R and Hadoop. This book is

ideal for R developers who are looking for a way to perform big data analytics with Hadoop. This book is also aimed at those who know Hadoop and want to build some intelligent applications over Big data with R packages. It would be helpful if readers have basic knowledge of R.

*2019 Supplement
Covering ChessBase
13, 14 & 15 Packt
Publishing Ltd*

Simplify machine learning model implementations with Spark

About This Book

Solve the day-to-day problems of data science with Spark

This unique cookbook consists of exciting and intuitive numerical recipes

Optimize your work by acquiring, cleaning, analyzing, predicting, and visualizing your data

Who This Book Is For

This book is for Scala developers with a fairly good exposure to and understanding of machine learning techniques, but lack practical implementations with Spark. A solid knowledge of machine learning algorithms is assumed, as well as hands-on experience of implementing ML algorithms with Scala. However, you do not need to be acquainted with the Spark ML libraries and ecosystem.

What You Will Learn

Get to know how Scala and Spark go hand-in-hand for developers when developing ML systems with Spark

Build a recommendation engine that scales with Spark

Find out how to build unsupervised clustering systems to

classify data in Spark
Build machine learning systems with the Decision Tree and Ensemble models in Spark
Deal with the curse of high-dimensionality in big data using Spark
Implement Text analytics for Search Engines in Spark
Streaming Machine Learning System implementation using Spark
In Detail Machine learning aims to extract knowledge from data, relying on fundamental concepts in computer science, statistics, probability, and optimization. Learning about algorithms enables a wide range of applications, from everyday tasks such as product recommendations and spam filtering to cutting edge

applications such as self-driving cars and personalized medicine. You will gain hands-on experience of applying these principles using Apache Spark, a resilient cluster computing system well suited for large-scale machine learning tasks. This book begins with a quick overview of setting up the necessary IDEs to facilitate the execution of code examples that will be covered in various chapters. It also highlights some key issues developers face while working with machine learning algorithms on the Spark platform. We progress by uncovering the various Spark APIs and the implementation of ML algorithms with developing classification systems,

recommendation engines, text analytics, clustering, and learning systems. Toward the final chapters, we'll focus on building high-end applications and explain various unsupervised methodologies and challenges to tackle when implementing with big data ML systems. Style and approach This book is packed with intuitive recipes supported with line-by-line explanations to help you understand how to optimize your work flow and resolve problems when working with complex data modeling tasks and predictive algorithms. This is a valuable resource for data scientists and those working on large scale data projects.

ChessBase Complete

IOS Press
Summary HBase in Action has all the knowledge you need to design, build, and run applications using HBase. First, it introduces you to the fundamentals of distributed systems and large scale data handling. Then, you'll explore real-world applications and code samples with just enough theory to understand the practical techniques. You'll see how to build applications with HBase and take advantage of the MapReduce processing framework. And along the way you'll learn patterns and best practices. About the Technology HBase is a NoSQL storage system designed for fast, random access to large volumes of data. It

runs on commodity hardware and scales smoothly from modest datasets to billions of rows and millions of columns. About this Book HBase in Action is an experience-driven guide that shows you how to design, build, and run applications using HBase. First, it introduces you to the fundamentals of handling big data. Then, you'll explore HBase with the help of real applications and code samples and with just enough theory to back up the practical techniques. You'll take advantage of the MapReduce processing framework and benefit from seeing HBase best practices in action. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also

available is all code from the book. What's Inside When and how to use HBase Practical examples Design patterns for scalable data systems Deployment, integration, and design Written for developers and architects familiar with data storage and processing. No prior knowledge of HBase, Hadoop, or MapReduce is required. Table of Contents PART 1 HBASE FUNDAMENTALS Introducing HBase Getting started Distributed HBase, HDFS, and MapReduce PART 2 ADVANCED CONCEPTS HBase table design Extending HBase with coprocessors Alternative HBase clients PART 3 EXAMPLE APPLICATIONS HBase

by example:
OpenTSDB Scaling GIS
on HBase PART 4
OPERATIONALIZING
HBASE Deploying
HBase Operations
[Big Data Analytics with
R and Hadoop](#) "O'Reilly
Media, Inc."
Summary MongoDB in
Action, Second Edition
is a completely revised
and updated version. It
introduces MongoDB
3.0 and the document-
oriented database
model. This perfectly
paced book gives you
both the big picture
you'll need as a
developer and enough
low-level detail to
satisfy system
engineers. Purchase of
the print book includes
a free eBook in PDF,
Kindle, and ePub
formats from Manning
Publications. About the
Technology This
document-oriented
database was built for

high availability,
supports rich, dynamic
schemas, and lets you
easily distribute data
across multiple
servers. MongoDB 3.0
is flexible, scalable,
and very fast, even
with big data loads.
About the Book
MongoDB in Action,
Second Edition is a
completely revised and
updated version. It
introduces MongoDB
3.0 and the document-
oriented database
model. This perfectly
paced book gives you
both the big picture
you'll need as a
developer and enough
low-level detail to
satisfy system
engineers. Lots of
examples will help you
develop confidence in
the crucial area of data
modeling. You'll also
love the deep
explanations of each
feature, including

replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation

Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1 GETTING STARTED A database for the modern web MongoDB through the JavaScript shell Writing programs using MongoDB PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY Indexing and query optimization Text search WiredTiger and pluggable storage Replication Scaling your system with sharding Deployment

and administration
Microsoft Azure Essentials - Fundamentals of Azure
Packt Publishing Ltd
Perform fast interactive analytics against different data sources using the Presto high-performance, distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Presto. Initially developed by Facebook, open source Presto is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber,

and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Presto query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Presto's use cases and learn about tools that will help you connect to Presto and query data Go deeper: Learn Presto's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Presto in production: Secure Presto, monitor workloads, tune queries, and connect more applications;

learn how other organizations apply Presto

Hadoop Beginner's Guide "O'Reilly Media, Inc."

As part of Packt's cookbook series, each recipe offers a practical, step-by-step solution to common problems found in HBase administration. This book is for HBase administrators, developers, and will even help Hadoop administrators. You are not required to have HBase experience, but are expected to have a basic understanding of Hadoop and MapReduce.

[Introduction to Hadoop, Spark, and Machine-Learning](#)

Simon and Schuster
Moving beyond MapReduce - learn resource management and big data

processing using YARN
About This Book Deep dive into YARN

components, schedulers, life cycle management and security architecture
Create your own Hadoop-YARN applications and integrate big data technologies with YARN
Step-by-step guide to provision, manage, and monitor Hadoop-YARN clusters with ease
Who This Book Is For This book is intended for those who want to understand what YARN is and how to efficiently use it for the resource management of large clusters. For cluster administrators, this book gives a detailed explanation of provisioning and managing YARN clusters. If you are a Java developer or an open source

contributor, this book will help you to drill down the YARN architecture, write your own YARN applications and understand the application execution phases. This book will also help big data engineers explore YARN integration with real-time analytics technologies such as Spark and Storm. What You Will Learn Explore YARN features and offerings Manage big data clusters efficiently using the YARN framework Create single as well as multi-node Hadoop-YARN clusters on Linux machines Understand YARN components and their administration Gain insights into application execution flow over a YARN cluster Write your own distributed application and execute it over

YARN cluster Work with schedulers and queues for efficient scheduling of applications Integrate big data projects like Spark and Storm with YARN In Detail Today enterprises generate huge volumes of data. In order to provide effective services and to make smarter and more intelligent decisions from these huge volumes of data, enterprises use big-data analytics. In recent years, Hadoop has been used for massive data storage and efficient distributed processing of data. The Yet Another Resource Negotiator (YARN) framework solves the design problems related to resource management faced by the Hadoop 1.x framework by

providing a more scalable, efficient, flexible, and highly available resource management framework for distributed data processing. This book starts with an overview of the YARN features and explains how YARN provides a business solution for growing big data needs. You will learn to provision and manage single, as well as multi-node, Hadoop-YARN clusters in the easiest way. You will walk through the YARN administration, life cycle management, application execution, REST APIs, schedulers, security framework and so on. You will gain insights about the YARN components and features such as ResourceManager, NodeManager, ApplicationMaster,

Container, Timeline Server, High Availability, Resource Localisation and so on. The book explains Hadoop-YARN commands and the configurations of components and explores topics such as High Availability, Resource Localization and Log aggregation. You will then be ready to develop your own ApplicationMaster and execute it over a Hadoop-YARN cluster. Towards the end of the book, you will learn about the security architecture and integration of YARN with big data technologies like Spark and Storm. This book promises conceptual as well as practical knowledge of resource management using YARN. Style and approach Starting with

the basics and covering the core concepts with the practical usage, this tutorial is a complete guide to learn and explore YARN offerings. [Apache Hadoop 3 Quick Start Guide](#) Packt Pub Limited Design, deploy, and maintain your own private or public Infrastructure as a Service (IaaS), using the open source OpenStack platform. In this practical guide, experienced developers and OpenStack contributors show you how to build clouds based on reference architectures, as well as how to perform daily administration tasks. Designed for horizontal scalability, OpenStack lets you build a cloud by integrating several technologies. This

approach provides flexibility, but knowing which options to use can be bewildering. Once you complete this book, you'll know the right questions to ask while you organize compute, storage, and networking resources. If you already know how to manage multiple Ubuntu machines and maintain MySQL, you're ready to: Set up automated deployment and configuration Design a single-node cloud controller Use metrics to improve scalability Explore compute nodes, network design, and storage Install OpenStack packages Use an example architecture to help simplify decision-making Build a working environment to explore an IaaS cloud Manage users, projects, and

quotas Tackle
maintenance,
debugging, and
network
troubleshooting
Monitor, log, backup,
and restore

**OpenStack
Operations Guide**

"O'Reilly Media, Inc."

If you are an
administrator or
developer who wants
to enter the world of
Big Data and BigTables
and would like to learn
about HBase, this is
the book for you.

Mastering Hadoop 3

McGraw-Hill Education
Big Data Analytics with
Spark is a step-by-step
guide for learning
Spark, which is an
open-source fast and
general-purpose
cluster computing
framework for large-
scale data analysis.
You will learn how to
use Spark for different
types of big data

analytics projects,
including batch,
interactive, graph, and
stream data analysis
as well as machine
learning. In addition,
this book will help you
become a much
sought-after Spark
expert. Spark is one of
the hottest Big Data
technologies. The
amount of data
generated today by
devices, applications
and users is exploding.
Therefore, there is a
critical need for tools
that can analyze large-
scale data and unlock
value from it. Spark is
a powerful technology
that meets that need.
You can, for example,
use Spark to perform
low latency
computations through
the use of efficient
caching and iterative
algorithms; leverage
the features of its shell
for easy and

interactive Data analysis; employ its fast batch processing and low latency features to process your real time data streams and so on. As a result, adoption of Spark is rapidly growing and is replacing Hadoop MapReduce as the technology of choice for big data analytics. This book provides an introduction to Spark and related big-data technologies. It covers Spark core and its add-on libraries, including Spark SQL, Spark Streaming, GraphX, and MLlib. Big Data Analytics with Spark is therefore written for busy professionals who prefer learning a new technology from a consolidated source instead of spending countless hours on the Internet trying to pick

bits and pieces from different sources. The book also provides a chapter on Scala, the hottest functional programming language, and the program that underlies Spark. You'll learn the basics of functional programming in Scala, so that you can write Spark applications in it. What's more, Big Data Analytics with Spark provides an introduction to other big data technologies that are commonly used along with Spark, like Hive, Avro, Kafka and so on. So the book is self-sufficient; all the technologies that you need to know to use Spark are covered. The only thing that you are expected to know is programming in any language. There is a critical shortage of people with big data

expertise, so companies are willing to pay top dollar for people with skills in areas like Spark and Scala. So reading this book and absorbing its principles will provide a boost—possibly a big boost—to your career.

Hbase Administration Cookbook "O'Reilly Media, Inc."

This book is intended for developers and Big Data engineers who want to know all about HBase at a hands-on level. For in-depth understanding, it would be helpful to have a bit of familiarity with HDFS and MapReduce programming concepts with no prior experience with HBase or similar technologies. This book is also for Big Data enthusiasts and database developers who have worked with other NoSQL databases

and now want to explore HBase as another futuristic, scalable database solution in the Big Data space.

[Service Worker](#)

[Development](#)

[Cookbook](#) SCB

Distributors

Updating a Classic!

Five years ago, when ChessBase Complete was released, it was an instant sensation. For the first time ever, a comprehensive manual for one of the most popular chess programs was available. It covered ChessBase through version 12. Since then, ChessBase has introduced three new versions and a remarkable (and free) suite of online tools for the world wide web. Many new and powerful functions have been added, all

with an eye towards ease-of-use. This Supplement updates the original ChessBase Complete, and once again gives the chess community the opportunity to make the most of this remarkable software. Searches are much more powerful, analysis has been automated, and we can now take full advantage of storing and sharing our data in a “chess cloud.” While this Supplement should be helpful to all ChessBase users, it builds upon the original edition. If you do not already have the first edition, you may wish to consider getting it. ChessBase Complete and this Supplement should make your chess time much more productive and enjoyable and, with the

software at hand, vastly accelerate your chess improvement.

About the Author: Jon Edwards recently qualified for the World Correspondence Chess Championship final round. He won the 10th United States Correspondence Championship in 1997 and the 8th North American Invitational Correspondence Chess Championship in 1999. He is a four-time winner of the APCT (American Postal Chess Tournaments) Championship and has been awarded the APCT Game of the Year Award twice. He received his correspondence International Master (IM) in 1997, his Senior International Master (SIM) in 1999. He is currently fighting for his final grandmaster

norm in the prestigious ICCF Spanish Masters. He has competed on the United States Correspondence Chess Olympiad team competing, reaching the final round. His correspondence ICCF rating places him the top 100 correspondence chess players worldwide. In addition to the extremely popular ChessBase Complete, Jon has written more than a dozen chess books, including The

Chess Analyst (Thinkers Press 1999) which chronicles the success in the US championship; Teach Yourself Visually: Chess (Wiley 2006); a photographically based chess primer; and Sacking the Citadel: The History, Theory, and Practice of the Classic Bishop Sacrifice (Russell Enterprises 2011). He also writes a regular column on Chess Technology for the American Chess Magazine.