

Data Model Patterns A Metadata Map The Morgan Kaufmann Series In Data Management Systems

Here you'll find one key to the development of a successful information system: Clearly capture and communicate both the abstract and concrete building blocks of data that describe your organization. In 1995, David Hay published *Data Model Patterns: Conventions of Thought* - the groundbreaking book on how to use standard data models to describe the standard business situations. *Enterprise Model Patterns: Describing the World* builds on the concepts presented there, adds 15 years of practical experience, and presents a more comprehensive view. You will learn how to apply both the abstract and concrete elements of your enterprise's architectural data model through four levels of abstraction: Level 0: An abstract template that underlies the Level 1 model that follows, plus two meta models: • Information Resources. In addition to books, articles, and e-mail notes, it also includes photographs, videos, and sound recordings. • Accounting. Accounting is remarkable because it is itself a modeling language. It takes a very different approach than data modelers in that instead of using entities and entity classes that represent things in the world, it is concerned with accounts that represent bits of value to the organization. Level 1: An enterprise model that is generic enough to apply to any company or government agency, but concrete enough to be readily understood by all. It describes: • People and Organization. Who is involved with the business? The people involved are not only the employees within the organization, but customers, agents, and others with whom the organization comes in contact. Organizations of interest include the enterprise itself and its own internal departments, as well as customers, competitors, government agencies, and the like. • Geographic Locations. Where is business conducted? A geographic location may be either a geographic area (defined as any bounded area on the Earth), a geographic point (used to identify a particular location), or, if you are an oil company for example, a geographic solid (such as an oil reserve). • Assets. What tangible items are used to carry out the business? These are any physical things that are manipulated, sometimes as products, but also as the means to producing products and services. • Activities. How is the business carried out? This model not only covers services offered, but also projects and any other kinds of activities. In addition, the model describes the events that cause activities to happen. • Time. All data is positioned in time, but some more than others. Level 2: A more detailed model describing specific functional areas: • Facilities • Human Resources • Communications and Marketing • Contracts • Manufacturing • The Laboratory Level 3: Examples of the details a model can have to address what is truly unique in a particular industry. Here you see how to address the unique bits in areas as diverse as: • Criminal Justice. The model presented here is based on the "Global Justice XML Data Model" (GJXDM). • Microbiology • Banking. The model presented here is the result of working for four different banks and then adding some thought to come up with something different from what is currently in any of them. • Highways. The model here is derived from a project in a Canadian Provincial Highway Department, and addresses the question "what is a road?"

Thousands of software projects are doomed because they're based on a faulty understanding of the business problem that needs to be solved. *Requirements Analysis: From Business Views to Architecture* is the solution. David C. Hay brings together the world's best requirements analysis practices from two key viewpoints: system development life cycle and architectural framework. Hay teaches you the complete process of defining an architecture - from a full understanding of what business people need to the creation of a complete enterprise architecture.

A quick and reliable way to build proven databases for core business functions Industry experts raved about *The Data Model Resource Book* when it was first published in March 1997 because it provided a simple, cost-effective way to design databases for core business functions. Len Silverston has now revised and updated the hugely successful 1st Edition, while adding a companion volume to take care of more specific requirements of different businesses. This updated volume provides a common set of data models for specific core functions shared by most businesses like human resources management, accounting, and project management. These models are standardized and are easily replicated by developers looking for ways to make corporate database development more efficient and cost effective. This guide is the perfect complement to *The Data Model Resource CD-ROM*, which is sold separately and provides the powerful design templates discussed in the book in a ready-to-use electronic format. A free demonstration CD-ROM is available with each copy of the print book to allow you to try before you buy the full CD-ROM.

Data Modeling Essentials, Third Edition, covers the basics of data modeling while focusing on developing a facility in techniques, rather than a simple familiarization with "the rules". In order to enable students to apply the basics of data modeling to real models, the book addresses the realities of developing systems in real-world situations by assessing the merits of a variety of possible solutions as well as using language and diagramming methods that represent industry practice. This revised edition has been given significantly expanded coverage and reorganized for greater reader comprehension even as it retains its distinctive hallmarks of readability and usefulness. Beginning with the basics, the book provides a thorough grounding in theory before guiding the reader through the various stages of applied data modeling and database design. Later chapters address advanced subjects, including business rules, data warehousing, enterprise-wide modeling and data management. It includes an entirely new section discussing the development of logical and physical modeling, along with new material describing a powerful technique for model verification. It also provides an excellent resource for additional lectures and exercises. This text is the ideal reference for data modelers, data architects, database designers, DBAs, and systems analysts, as well as undergraduate and graduate-level students looking for a real-world perspective. Thorough coverage of the fundamentals and relevant theory. Recognition and support for the creative side of the process. Expanded coverage of applied data modeling includes new chapters on logical and physical database design. New material describing a powerful technique for model verification. Unique coverage of the practical and human aspects of modeling, such as working with business specialists, managing change, and resolving conflict.

Get expert guidance on architecting end-to-end data management solutions with Apache Hadoop. While many sources explain how to use various components in the Hadoop ecosystem, this practical book takes you through architectural considerations necessary to tie those components together into a complete tailored application, based on your particular use case. To reinforce those lessons, the book's second section provides detailed examples of architectures used in some of the most commonly found Hadoop applications. Whether you're designing a new Hadoop application, or planning to integrate Hadoop into your existing data infrastructure, *Hadoop Application Architectures* will skillfully guide you through the process. This book covers: Factors to consider when using Hadoop to store and model data Best practices for moving data in and out of the system Data processing frameworks, including MapReduce, Spark, and Hive Common Hadoop processing patterns, such as removing duplicate records and using windowing analytics Giraph, GraphX, and other tools for large graph processing on Hadoop Using workflow orchestration and scheduling tools such as Apache Oozie Near-real-time stream processing with Apache Storm, Apache Spark Streaming, and Apache Flume Architecture examples for clickstream analysis, fraud detection, and data warehousing

Big data modeling is very challenging to handle using traditional database modeling and management systems. This book will teach you how to model big data using the latest and more efficient tools such as ERWIN, ANACONDA (Python), and WEKA to model data.

Master a graph data modeling technique superior to traditional data modeling for both relational and NoSQL databases (graph, document, key-value, and column), leveraging cognitive psychology to improve big data designs. From Karen Lopez's Foreword: In this book, Thomas Frisendal raises important questions about the continued usefulness of traditional data modeling notations and approaches: Are Entity Relationship Diagrams (ERDs) relevant to analytical data requirements? Are ERDs relevant in the new world of Big Data? Are ERDs still the best way to work with business users to understand their needs?

Are Logical and Physical Data Models too closely coupled? Are we correct in using the same notations for communicating with business users and developers? Should we refine our existing notations and tools to meet these new needs, or should we start again from a blank page? What new notations and approaches will we need? How will we use those to build enterprise database systems? Frisendal takes us through the history of data modeling, enterprise data models and traditional modeling methods. He points out, quite contentiously, where he feels we have gone wrong and in a few places where we got it right. He then maps out the psychology of meaning and context, while identifying important issues about where data modeling may or may not fit in business modeling. The main subject of this work is a proposal for a new exploration-driven modeling approach and new modeling notations for business concept models, business solutions models, and physical data models with examples on how to leverage those for implementing into any target database or datastore. These new notations are based on a property graph approach to modeling data.

Here you will learn how to develop an attractive, easily readable, conceptual, business-oriented entity/relationship model, using a variation on the UML Class Model notation. This book has two audiences: • Data modelers (both analysts and database designers) who are convinced that UML has nothing to do with them; and • UML experts who don't realize that architectural data modeling really is different from object modeling (and that the differences are important). David Hay's objective is to finally bring these two groups together in peace. Here all modelers will receive guidance on how to produce a high quality (that is, readable) entity/relationship model to describe the data architecture of an organization. The notation involved happens to be the one for class models in the Unified Modeling Language, even though UML was originally developed to support object-oriented design. Designers have a different view of the world from those who develop business-oriented conceptual data models, which means that to use UML for architectural modeling requires some adjustments. These adjustments are described in this book. David Hay is the author of Enterprise Model Patterns: Describing the World, a comprehensive model of a generic enterprise. The diagrams were at various levels of abstraction, and they were all rendered in the slightly modified version of UML Class Diagrams presented here. This book is a handbook to describe how to build models such as these. By way of background, an appendix provides a history of the two groups, revealing the sources of their different attitudes towards the system development process. If you are an old-school ER modeler and now find yourself having to come up to speed on UML to get that next job (or keep the current one), this is your guidebook to success. If you are a long time object oriented programmer who has to interact with data modelers, this book is for you too. David has done the hard work of mapping out how to do a logical entity relationship model using standard (and accepted) UML diagram components. This book shows you step-by-step, with ample examples, how to get from here to there with the least pain possible for all concerned. Kent Graziano Certified Data Vault Master and Oracle ACE Past-President of ODTUG & RMOUG Brilliantly organized: three books hidden in one cohesive work. Notwithstanding the tremendous value provided by cross-training data architects/modelers and object modelers/architects, making each better at what they do, Appendix B presents an absolutely awesome concise, yet detailed, history of modeling objects and data that clearly documents the differences in the approaches over the years and helps bring it all into perspective. This book is packed with useful information. Even the footnotes add clarity and offer interesting and often humorous editorial insight making it a fun read. Whatever viewpoint the reader is coming from this book has something to offer as long as the reader maintains an open mind. Roland Berg Senior Architect Diligent Consulting, Inc. San Antonio, Texas

What value does semantic data modeling offer? As an information architect or data science professional, let's say you have an abundance of the right data and the technology to extract business gold—but you still fail. The reason? Bad data semantics. In this practical and comprehensive field guide, author Panos Alexopoulos takes you on an eye-opening journey through semantic data modeling as applied in the real world. You'll learn how to master this craft to increase the usability and value of your data and applications. You'll also explore the pitfalls to avoid and dilemmas to overcome for building high-quality and valuable semantic representations of data. Understand the fundamental concepts, phenomena, and processes related to semantic data modeling Examine the quirks and challenges of semantic data modeling and learn how to effectively leverage the available frameworks and tools Avoid mistakes and bad practices that can undermine your efforts to create good data models Learn about model development dilemmas, including representation, expressiveness and content, development, and governance Organize and execute semantic data initiatives in your organization, tackling technical, strategic, and organizational challenges

The practice of enterprise application development has benefited from the emergence of many new enabling technologies. Multi-tiered object-oriented platforms, such as Java and .NET, have become commonplace. These new tools and technologies are capable of building powerful applications, but they are not easily implemented. Common failures in enterprise applications often occur because their developers do not understand the architectural lessons that experienced object developers have learned. Patterns of Enterprise Application Architecture is written in direct response to the stiff challenges that face enterprise application developers. The author, noted object-oriented designer Martin Fowler, noticed that despite changes in technology--from Smalltalk to CORBA to Java to .NET--the same basic design ideas can be adapted and applied to solve common problems. With the help of an expert group of contributors, Martin distills over forty recurring solutions into patterns. The result is an indispensable handbook of solutions that are applicable to any enterprise application platform. This book is actually two books in one. The first section is a short tutorial on developing enterprise applications, which you can read from start to finish to understand the scope of the book's lessons. The next section, the bulk of the book, is a detailed reference to the patterns themselves. Each pattern provides usage and implementation information, as well as detailed code examples in Java or C#. The entire book is also richly illustrated with UML diagrams to further explain the concepts. Armed with this book, you will have the knowledge necessary to make important architectural decisions about building an enterprise application and the proven patterns for use when building them. The topics covered include · Dividing an enterprise application into layers · The major approaches to organizing business logic · An in-depth treatment of mapping between objects and relational databases · Using Model-View-Controller to organize a Web presentation · Handling concurrency for data that spans multiple transactions · Designing distributed object interfaces

Would you like to use a consistent visual notation for drawing integration solutions? "Look inside the front cover." Do you want to harness the power of asynchronous systems without getting caught in the pitfalls? "See "Thinking Asynchronously" in the Introduction." Do you want to know which style of application integration is best for your purposes? "See Chapter 2, Integration Styles." Do you want to learn techniques for processing messages concurrently? "See Chapter 10, Competing Consumers and Message Dispatcher." Do you want to learn how you can track asynchronous messages as they flow across distributed systems? "See Chapter 11, Message History and Message Store." Do you want to understand how a system designed using integration patterns can be implemented using Java Web services, .NET message queuing, and a TIBCO-based publish-subscribe architecture? "See Chapter 9, Interlude: Composed Messaging." Utilizing years of practical experience, seasoned experts Gregor Hohpe and Bobby Woolf show how asynchronous messaging has proven to be the best strategy for enterprise integration success. However, building and deploying messaging solutions presents a number of problems for developers. " Enterprise Integration Patterns " provides an invaluable catalog of

sixty-five patterns, with real-world solutions that demonstrate the formidable of messaging and help you to design effective messaging solutions for your enterprise. The authors also include examples covering a variety of different integration technologies, such as JMS, MSMQ, TIBCO ActiveEnterprise, Microsoft BizTalk, SOAP, and XSL. A case study describing a bond trading system illustrates the patterns in practice, and the book offers a look at emerging standards, as well as insights into what the future of enterprise integration might hold. This book provides a consistent vocabulary and visual notation framework to describe large-scale integration solutions across many technologies. It also explores in detail the advantages and limitations of asynchronous messaging architectures. The authors present practical advice on designing code that connects an application to a messaging system, and provide extensive information to help you determine when to send a message, how to route it to the proper destination, and how to monitor the health of a messaging system. If you want to know how to manage, monitor, and maintain a messaging system once it is in use, get this book. 0321200683B09122003

Data models are the main medium used to communicate data requirements from business to IT, and within IT from analysts, modelers, and architects, to database designers and developers. Therefore it's essential to get the data model right. But how do you determine right? That's where the Data Model Scorecard® comes in. The Data Model Scorecard is a data model quality scoring tool containing ten categories aimed at improving the quality of your organization's data models. Many of my consulting assignments are dedicated to applying the Data Model Scorecard to my client's data models – I will show you how to apply the Scorecard in this book. This book, written for people who build, use, or review data models, contains the Data Model Scorecard template and an explanation along with many examples of each of the ten Scorecard categories. There are three sections: In Section I, Data Modeling and the Need for Validation, receive a short data modeling primer in Chapter 1, understand why it is important to get the data model right in Chapter 2, and learn about the Data Model Scorecard in Chapter 3. In Section II, Data Model Scorecard Categories, we will explain each of the ten categories of the Data Model Scorecard. There are ten chapters in this section, each chapter dedicated to a specific Scorecard category: · Chapter 4: Correctness · Chapter 5: Completeness · Chapter 6: Scheme · Chapter 7: Structure · Chapter 8: Abstraction · Chapter 9: Standards · Chapter 10: Readability · Chapter 11: Definitions · Chapter 12: Consistency · Chapter 13: Data In Section III, Validating Data Models, we will prepare for the model review (Chapter 14), cover tips to help during the model review (Chapter 15), and then review a data model based upon an actual project (Chapter 16).

What could you do with data if scalability wasn't a problem? With this hands-on guide, you'll learn how Apache Cassandra handles hundreds of terabytes of data while remaining highly available across multiple data centers -- capabilities that have attracted Facebook, Twitter, and other data-intensive companies. Cassandra: The Definitive Guide provides the technical details and practical examples you need to assess this database management system and put it to work in a production environment. Author Eben Hewitt demonstrates the advantages of Cassandra's nonrelational design, and pays special attention to data modeling. If you're a developer, DBA, application architect, or manager looking to solve a database scaling issue or future-proof your application, this guide shows you how to harness Cassandra's speed and flexibility. Understand the tenets of Cassandra's column-oriented structure Learn how to write, update, and read Cassandra data Discover how to add or remove nodes from the cluster as your application requires Examine a working application that translates from a relational model to Cassandra's data model Use examples for writing clients in Java, Python, and C# Use the JMX interface to monitor a cluster's usage, memory patterns, and more Tune memory settings, data storage, and caching for better performance

Data-driven insights are a key competitive advantage for any industry today, but deriving insights from raw data can still take days or weeks. Most organizations can't scale data science teams fast enough to keep up with the growing amounts of data to transform. What's the answer? Self-service data. With this practical book, data engineers, data scientists, and team managers will learn how to build a self-service data science platform that helps anyone in your organization extract insights from data. Sandeep Uttamchandani provides a scorecard to track and address bottlenecks that slow down time to insight across data discovery, transformation, processing, and production. This book bridges the gap between data scientists bottlenecked by engineering realities and data engineers unclear about ways to make self-service work. Build a self-service portal to support data discovery, quality, lineage, and governance Select the best approach for each self-service capability using open source cloud technologies Tailor self-service for the people, processes, and technology maturity of your data platform Implement capabilities to democratize data and reduce time to insight Scale your self-service portal to support a large number of users within your organization

25 proven patterns for improving data access and application performance Efficient, high-quality data access code is crucial to the performance and usability of virtually any enterprise application--and there's no better way to improve an existing system than to optimize its data access code. Regardless of database engine, platform, language, or application, developers repeatedly encounter the same relational database access challenges. In Data Access Patterns, Clifton Nock identifies 25 proven solutions, presenting each one in the form of a clear, easy-to-use pattern. These patterns solve an exceptionally wide range of problems including creating efficient database-independent applications, hiding obscure database semantics from users, speeding database resource initialization, simplifying development and maintenance, improving support for concurrency and transactions, and eliminating data access bottlenecks. Every pattern is illustrated with fully commented Java/JDBC code examples, as well as UML diagrams representing interfaces, classes, and relationships. The patterns are organized into five categories: Decoupling Patterns: Build cleaner, more reliable systems by decoupling data access code from other application logic Resource Patterns: Manage relational database resources more efficiently Input/Output Patterns: Simplify I/O operations by translating consistently between "physical" relational data and domain object representations of that data Cache Patterns: Use caching strategically, to optimize the tradeoffs between data access optimization and cache overhead Concurrency Patterns: Implement concurrency and transactions more effectively and reliably Data Access Patterns demystifies techniques that have traditionally been used only in the most robust data access solutions--making those techniques practical for every software developer, architect, and designer.

In the era of "Big Data," effective information management has become crucial to the success of virtually all organizations. Unfortunately, few IT practitioners know today's best practices for successfully managing enterprise information resources. "Patterns for Information Management" offers the solution: a multi-disciplinary patterns-based approach that reflects where information comes from, how it is distributed, protected, governed, monitored -- and, ultimately, utilized. Leading IBM information management experts Mandy Chessell and Harald Smith show how to

design, integrate, and optimize "information supply chains" that: * Effectively characterize the information issues associated with today's complex distributed systems* Fully reflect today's service-oriented, message-based environments* Establish an information landscape that is effective, manageable, ordered, and efficient.* Create orderly flows of information which can be reused and synchronized wherever necessary* Support reasoned decision-making about the applicability and consequences of each alternative approach Chessell and Smith complement their plain-English information pattern language with a complete case study demonstrating its real-world applications. Their pioneering guidebook will serve as both an indispensable guide to architects at all levels of experience "and" an outstanding long-term information management reference.

Renowned DAX experts Alberto Ferrari and Marco Russo teach you how to design data models for maximum efficiency and effectiveness. How can you use Excel and Power BI to gain real insights into your information? As you examine your data, how do you write a formula that provides the numbers you need? The answers to both of these questions lie with the data model.

This book introduces the basic techniques for shaping data models in Excel and Power BI. It's meant for readers who are new to data modeling as well as for experienced data modelers looking for tips from the experts. If you want to use Power BI or Excel to analyze data, the many real-world examples in this book will help you look at your reports in a different way—like experienced data modelers do. As you'll soon see, with the right data model, the correct answer is always a simple one! By reading this book, you will:

- Gain an understanding of the basics of data modeling, including tables, relationships, and keys
- Familiarize yourself with star schemas, snowflakes, and common modeling techniques
- Learn the importance of granularity
- Discover how to use multiple fact tables, like sales and purchases, in a complex data model
- Manage calendar-related calculations by using date tables
- Track historical attributes, like previous addresses of customers or manager assignments
- Use snapshots to compute quantity on hand
- Work with multiple currencies in the most efficient way
- Analyze events that have durations, including overlapping durations
- Learn what data model you need to answer your specific business questions

About This Book

- For Excel and Power BI users who want to exploit the full power of their favorite tools
- For BI professionals seeking new ideas for modeling data

This book shows how to apply pattern ideas in business applications. It presents more than 20 structural and behavioral business patterns that use the REA (resources, events, agents) pattern as a common backbone. The developer working on business frameworks can use the patterns to derive the right abstractions and to design and ensure that the meta-rules are followed by the developers of the actual applications. The application developer can use these patterns to design a business application, to ensure that it does not violate the domain rules, and to adapt the application to changing requirements without the need to change the overall architecture.

This work has been revised and updated to provide a comprehensive treatment of database design for commercial database products and their applications. The book covers the basic foundation of design as well as more advanced techniques, and also incorporates coverage of data warehousing and OLAP (On-Line Analytical Processing), data mining, object-relational, multimedia, and temporal/spatial design.

Data Model Patterns: A Metadata Map Elsevier

Introduces concepts for organizing data within a company to make it more accessible and meaningful. The author explains where databases went wrong in the 1990s, describes metadata-based technologies and standards, and illustrates the various implementation options by depicting five distinct metadata solutions for the same problem.

With the immense cost savings and scalability the cloud provides, the rationale for building cloud native applications is no longer in question. The real issue is how. With this practical guide, developers will learn about the most commonly used design patterns for building cloud native applications using APIs, data, events, and streams in both greenfield and brownfield development. You'll learn how to incrementally design, develop, and deploy large and effective cloud native applications that you can manage and maintain at scale with minimal cost, time, and effort. Authors Kasun Indrasiri and Sriskandarajah Suhothayan highlight use cases that effectively demonstrate the challenges you might encounter at each step. Learn the fundamentals of cloud native applications Explore key cloud native communication, connectivity, and composition patterns Learn decentralized data management techniques Use event-driven architecture to build distributed and scalable cloud native applications Explore the most commonly used patterns for API management and consumption Examine some of the tools and technologies you'll need for building cloud native systems

Understand the language and vocabulary of Data Architecture. The Data Architecture field is rife with terms that have become "fashionable". Some of the terms began with very specific, specialized, meanings – but as their use spread, they lost the precision of their technical definitions and become, well, "buzzwords". A buzzword is "a word or expression from a particular subject area that has become fashionable because it has been used a lot". Compliance is "the obeying of an accepted principle or instruction that states the way things are or should be done." The assignment is to take buzzwords and follow rules to use them correctly. We cut through the hype to arrive at buzzword compliance – the state where you fully understand the words that in fact have real meaning in the data architecture industry. This book will rationalize the various ways all these terms are defined. Of necessity, the book must address all aspects of describing an enterprise and its data management technologies. This includes a wide range of subjects, from entity/relationship modeling, through the semantic web, to database issues like relational and "beyond relational" ("NoSQL") approaches. In each case, the definitions for the subject are meant to be detailed enough to make it possible to understand basic principles—while recognizing that a full understanding will require consulting the sources where they are more completely described. The book's Glossary contains a catalogue of definitions and its Bibliography contains a comprehensive set of references.

Developing High Quality Data Models provides an introduction to the key principles of data modeling. It explains the purpose of data models in both developing an Enterprise Architecture and in supporting Information Quality; common problems in data model development; and how to develop high quality data models, in particular conceptual, integration, and enterprise data models. The book is organized into four parts. Part 1 provides an overview of data models and data modeling including the basics of data model notation; types and uses of data models; and the place of data models in enterprise architecture. Part 2 introduces some general principles for data models, including principles for developing ontologically based data models; and applications of the principles for attributes, relationship types, and entity types. Part 3 presents an ontological framework for developing consistent data models. Part 4 provides the full data model that has been in development throughout the book. The model was created using Jotne EPM Technologys EDMVisualExpress data modeling tool. This book was designed for all types of modelers: from those who understand data modeling basics but are just starting to learn about data modeling in practice, through to experienced data modelers seeking to expand their knowledge and skills and solve some of the more challenging problems of data modeling. Uses a number of common data model patterns to explain how to develop data models over a wide scope in a way that is consistent and of high quality Offers generic data model templates that are reusable in many applications and are fundamental for developing more specific templates Develops ideas for creating consistent approaches to high quality data models

This IBM Redguide™ publication looks back on the key decisions that made the data lake successful and looks forward to the future. It proposes that the metadata management and governance approaches

developed for the data lake can be adopted more broadly to increase the value that an organization gets from its data. Delivering this broader vision, however, requires a new generation of data catalogs and governance tools built on open standards that are adopted by a multi-vendor ecosystem of data platforms and tools. Work is already underway to define and deliver this capability, and there are multiple ways to engage. This guide covers the reasons why this new capability is critical for modern businesses and how you can get value from it.

Whether you're building a social media site or an internal-use enterprise application, this hands-on guide shows you the connection between MongoDB and the business problems it's designed to solve. You'll learn how to apply MongoDB design patterns to several challenging domains, such as ecommerce, content management, and online gaming. Using Python and JavaScript code examples, you'll discover how MongoDB lets you scale your data model while simplifying the development process. Many businesses launch NoSQL databases without understanding the techniques for using their features most effectively. This book demonstrates the benefits of document embedding, polymorphic schemas, and other MongoDB patterns for tackling specific big data use cases, including: Operational intelligence: Perform real-time analytics of business data Ecommerce: Use MongoDB as a product catalog master or inventory management system Content management: Learn methods for storing content nodes, binary assets, and discussions Online advertising networks: Apply techniques for frequency capping ad impressions, and keyword targeting and bidding Social networking: Learn how to store a complex social graph, modeled after Google+ Online gaming: Provide concurrent access to character and world data for a multiplayer role-playing game

This book analyzes the methods, technologies, standards, and languages to structure and describe data in their entirety. It reveals common features, hidden assumptions, and ubiquitous patterns among these methods and shows how data are actually structured and described independently from particular trends and technologies. Examples of data structuring methods analyzed critically include: Encodings (e.g. Unicode) Identifiers and Identifier systems (e.g. ISBN) File systems Database Systems (record databases, relational databases, NoSQL...) Data structuring languages (JSON, XML, CSV, RDF...) markup languages (SGML, HTML, TEI, Markdown...) Schema languages (BNF, XSD, RDFS, OWL, SQL...) Conceptual modeling languages (ERM, ORM, UML, DSL...) Conceptual diagrams It is shown how particular method of data structuring and description can best be categorized by their primary purpose. The study further exposes five basic paradigms that deeply shape how data is structured and described in practice. The third results is a pattern language of data structuring. Patterns show problems and solutions which occur over and over again in data. Each pattern is described with its benefits, consequences, pitfalls, and relations to other patterns. The results can help to better understand data and its actual forms, both for consumption and creation of data. Possible applications include data analysis, data modeling, data archaeology, and data literacy.

"This book introduces you to R, RStudio, and the tidyverse, a collection of R packages designed to work together to make data science fast, fluent, and fun. Suitable for readers with no previous programming experience"--

This is the digital version of the printed book (Copyright © 1996). Learning the basics of a modeling technique is not the same as learning how to use and apply it. To develop a data model of an organization is to gain insights into its nature that do not come easily. Indeed, analysts are often expected to understand subtleties of an organization's structure that may have evaded people who have worked there for years. Here's help for those analysts who have learned the basics of data modeling (or "entity/relationship modeling") but who need to obtain the insights required to prepare a good model of a real business. Structures common to many types of business are analyzed in areas such as accounting, material requirements planning, process manufacturing, contracts, laboratories, and documents. In each chapter, high-level data models are drawn from the following business areas: The Enterprise and Its World The Things of the Enterprise Procedures and Activities Contracts Accounting The Laboratory Material Requirements Planning Process Manufacturing Documents Lower-Level Conventions

Data Model Patterns: A Metadata Map not only presents a conceptual model of a metadata repository but also demonstrates a true enterprise data model of the information technology industry itself. It provides a step-by-step description of the model and is organized so that different readers can benefit from different parts. It offers a view of the world being addressed by all the techniques, methods, and tools of the information processing industry (for example, object-oriented design, CASE, business process re-engineering, etc.) and presents several concepts that need to be addressed by such tools. This book is pertinent, with companies and government agencies realizing that the data they use represent a significant corporate resource recognize the need to integrate data that has traditionally only been available from disparate sources. An important component of this integration is management of the "metadata" that describe, catalogue, and provide access to the various forms of underlying business data. The "metadata repository" is essential to keep track of the various physical components of these systems and their semantics. The book is ideal for data management professionals, data modeling and design professionals, and data warehouse and database repository designers. A comprehensive work based on the Zachman Framework for information architecture—encompassing the Business Owner's, Architect's, and Designer's views, for all columns (data, activities, locations, people, timing, and motivation) Provides a step-by-step description of model and is organized so that different readers can benefit from different parts Provides a view of the world being addressed by all the techniques, methods and tools of the information processing industry (for example, object-oriented design, CASE, business process re-engineering, etc.) Presents many concepts that are not currently being addressed by such tools — and should be

Managing Data in Motion describes techniques that have been developed for significantly reducing the complexity of managing system interfaces and enabling scalable architectures. Author April Reeve brings over two decades of experience to present a vendor-neutral approach to moving data between computing environments and systems. Readers will learn the techniques, technologies, and best practices for managing the passage of data between computer systems and integrating disparate data together in an enterprise environment. The average enterprise's computing environment is comprised of hundreds to thousands computer systems that have been built, purchased, and acquired over time. The data from these various systems needs to be integrated for reporting and analysis, shared for business transaction processing, and converted from one format to another when old systems are replaced and new systems are acquired. The management of the "data in motion" in organizations is rapidly becoming one of the biggest concerns for business and IT management. Data warehousing and conversion, real-time data integration, and cloud and "big data" applications are just a few of the challenges facing organizations and businesses today. Managing Data in Motion tackles these and other topics in a style easily understood by business and IT managers as well as programmers and architects. Presents a vendor-neutral overview of the different technologies and techniques for moving data between computer systems including the emerging solutions for unstructured as well as structured data types Explains, in non-technical terms, the architecture and components required to perform data integration Describes how to reduce the complexity of managing system interfaces and enable a scalable data architecture that can handle the dimensions of "Big Data"

Best-selling author and database expert with more than 25 years of experience modeling application and enterprise data, Dr. Michael Blaha provides tried and tested data model patterns, to help readers avoid common modeling mistakes and unnecessary frustration on their way to building effective data models. Unlike the typical methodology book, Patterns of Data Modeling provides advanced techniques for those who have mastered the basics. Recognizing that database representation sets the path for software, determines its flexibility, affects its quality, and influences whether it succeeds or fails, the text focuses on databases rather than programming. It is one of the first books to apply the popular patterns perspective to database systems and data models. It offers practical advice on the core aspects of

applications and provides authoritative coverage of mathematical templates, antipatterns, archetypes, identity, canonical models, and relational database design.

The tremendous growth in the availability of inexpensive computing power and easy availability of computers have generated tremendous interest in the design and implementation of Complex Systems. Computer-based solutions offer great support in the design of Complex Systems. Furthermore, Complex Systems are becoming increasingly complex themselves. This research book comprises a selection of state-of-the-art contributions to topics dealing with Complex Systems in a Knowledge-based Environment. Complex systems are ubiquitous. Examples comprise, but are not limited to System of Systems, Service-oriented Approaches, Agent-based Systems, and Complex Distributed Virtual Systems. These are application domains that require knowledge of engineering and management methods and are beyond the scope of traditional systems. The chapters in this book deal with a selection of topics which range from uncertainty representation, management and the use of ontological means which support and are large-scale business integration. All contributions were invited and are based on the recognition of the expertise of the contributing authors in the field. By collating these sources together in one volume, the intention was to present a variety of tools to the reader to assist in both study and work. The second intention was to show how the different facets presented in the chapters are complementary and contribute towards this emerging discipline designed to aid in the analysis of complex systems.

Data-governance programs focus on authority and accountability for the management of data as a valued organizational asset. Data Governance should not be about command-and-control, yet at times could become invasive or threatening to the work, people and culture of an organization. Non-Invasive Data Governance™ focuses on formalizing existing accountability for the management of data and improving formal communications, protection, and quality efforts through effective stewarding of data resources. Non-Invasive Data Governance will provide you with a complete set of tools to help you deliver a successful data governance program. Learn how:

- Steward responsibilities can be identified and recognized, formalized, and engaged according to their existing responsibility rather than being assigned or handed to people as more work.
- Governance of information can be applied to existing policies, standard operating procedures, practices, and methodologies, rather than being introduced or emphasized as new processes or methods.
- Governance of information can support all data integration, risk management, business intelligence and master data management activities rather than imposing inconsistent rigor to these initiatives.
- A practical and non-threatening approach can be applied to governing information and promoting stewardship of data as a cross-organization asset.
- Best practices and key concepts of this non-threatening approach can be communicated effectively to leverage strengths and address opportunities to improve.

What people are saying about Search Patterns "Search Patterns is a delight to read -- very thoughtful and thought provoking. It's the most comprehensive survey of designing effective search experiences I've seen." --Irene Au, Director of User Experience, Google "I love this book! Thanks to Peter and Jeffery, I now know that search (yes, boring old yucky who cares search) is one of the coolest ways around of looking at the world." --Dan Roam, author, The Back of the Napkin (Portfolio Hardcover) "Search Patterns is a playful guide to the practical concerns of search interface design. It contains a bonanza of screenshots and illustrations that capture the best of today's design practices and presents a fresh perspective on the broader role of search and discovery." --Marti Hearst, Professor, UC Berkeley and author, Search User Interfaces (Cambridge University Press) "It's not often I come across a book that asks profound questions about a fundamental human activity, and then proceeds to answer those questions with practical observations and suggestions. Search Patterns is an expedition into the heart of the web and human cognition, and for me it was a delightful journey that delivered scores of insights." --Dave Gray, Founder and Chairman, XPLANE "Search is swiftly transforming everything we know, yet people don't understand how mavens design search: by stacking breadcrumbs, scenting widgets, and keeping eyeballs on the engine. I urge you to put your eyeballs on this unique and important book." --Bruce Sterling, Writer, Futurist, and Co-Founder, The Electronic Frontier Foundation "As one who searches a lot (and often ends up frustrated), Search Patterns is a revelation." --Nigel Holmes, Designer, Theorist, and Principal, Explanation Graphics "Search Patterns is a fabulous must-have book! Inside, you'll learn the whys and wheres of practically every modern search design trick and technique." --Jared Spool, CEO and Founder, User Interface Engineering Search is among the most disruptive innovations of our time. It influences what we buy and where we go. It shapes how we learn and what we believe. In this provocative and inspiring book, you'll explore design patterns that apply across the categories of web, ecommerce, enterprise, desktop, mobile, social, and real-time search and discovery. Filled with colorful illustrations and examples, Search Patterns brings modern information retrieval to life, covering such diverse topics as relevance, faceted navigation, multi-touch, personalization, visualization, multi-sensory search, and augmented reality. By drawing on their own experience-as well as best practices and evidence-based research-the authors not only offer a practical guide to help you build effective search applications, they also challenge you to imagine the future of discovery. You'll find Search Patterns intriguing and invaluable, whether you're a web practitioner, mobile designer, search entrepreneur, or just interested in the topic. Discover a pattern language for search that embraces user psychology and behavior, information architecture, interaction design, and emerging technology Boost enterprise efficiency and e-commerce sales Enable mobile users to achieve goals, complete tasks, and find what they need Drive design innovation for search interfaces and applications

This third volume of the best-selling "Data Model Resource Book" series revolutionizes the data modeling discipline by answering the question "How can you save significant time while improving the quality of any type of data modeling effort?" In contrast to the first two volumes, this new volume focuses on the fundamental, underlying patterns that affect over 50 percent of most data modeling efforts. These patterns can be used to considerably reduce modeling time and cost, to jump-start data modeling efforts, as standards and guidelines to increase data model consistency and quality, and as an objective source against which an enterprise can evaluate data models. Praise for The Data Model Resource Book, Volume 3 "Len and Paul look beneath the superficial issues of data modeling and have produced a work that is a must for every serious designer and manager of an IT project." —Bill Inmon, World-renowned expert, speaker, and author on data warehousing and widely recognized as the "father of data warehousing" "The Data Model Resource Book, Volume 3: Universal Patterns for Data Modeling is a great source for reusable patterns you can use to save a tremendous amount of time, effort, and cost on any data modeling effort. Len Silverston and Paul Agnew have provided an indispensable reference of very high-quality patterns for the most foundational types of datamodel structures. This book represents a revolutionary leap in moving the data modeling profession forward." —Ron Powell, Cofounder and Editorial Director of the Business Intelligence Network "After we model a Customer, Product, or Order, there is still more about each of these that remains to be captured, such as roles they play, classifications in which they belong, or states in which they change. The Data Model Resource Book, Volume 3: Universal Patterns for Data Modeling clearly illustrates these common structures. Len Silverston and Paul Agnew have created a valuable addition to our field, allowing us to improve the consistency and quality of our models by leveraging the many common structures within this text." —Steve Hoberman, Best-Selling Author of Data Modeling Made Simple "The large national health insurance company I work at has actively used these data patterns and the (Universal Data Models) UDM, ahead of this book, through Len Silverston's UDM Jump Start engagement. The patterns have found their way into the core of our Enterprise Information Model, our data warehouse designs, and progressively into key business function databases. We are getting to reuse the patterns across projects and are reaping benefits in understanding, flexibility, and time-to-market. Thanks so much." —David Chasteen, Enterprise Information Architect "Reusing proven data modeling design patterns means exactly that. Data models become stable, but remain very flexible to accommodate changes. We have had the fortune of having Len and Paul share the patterns that are described in this book via our engagements with Universal Data Models, LLC. These data modeling design patterns have helped us to focus on the essential business issues because we have leveraged these reusable building blocks for many of the standard design problems. These

design patterns have also helped us to evaluate the quality of data models for their intended purpose. Many times there are a lot of enhancements required. Too often the very specialized business-oriented data model is also implemented physically. This may have significant drawbacks to flexibility. I'm looking forward to increasing the data modeling design pattern competence within Nokia with the help of this book." —Teemu Mattelmaki, Chief Information Architect, Nokia "Once again, Len Silverston, this time together with Paul Agnew, has made a valuable contribution to the body of knowledge about datamodels, and the act of building sound data models. As a professional d

Proven data models that save companies' time and money developing data architectures, databases and data warehouses. The Data Model Resource Book provides a common set of data models for specific functions common to most businesses, such as sales, marketing, order processing, contracts, shipments, invoicing, work efforts, budgeting, and accounting. Readers can apply more than one data model to their own company to meet specific data needs. A CD-ROM ,sold separately, contains the SQL for all of the data models contained in the book in Oracle, Sybase, Informix, Microsoft SQL Server and ANSI compatible formats.

Written in plain English and based on successful client engagements, Data Modeling of Financial Derivatives: A Conceptual Approach introduces new and veteran data modelers, financial analysts, and IT professionals to the fascinating world of financial derivatives. Covering futures, forwards, options, swaps, and forward rate agreements, finance and modeling expert Robert Mamayev shows you step-by-step how to structure and describe financial data using advanced data modeling techniques. The book introduces IT professionals, in particular, to various financial and data modeling concepts that they may not have seen before, giving them greater proficiency in the financial language of derivatives—and greater ability to communicate with financial analysts without fear or hesitation. Such knowledge will be especially useful to those looking to pick up the necessary skills to become productive right away working in the financial sector. Financial analysts reading this book will come to grips with various data modeling concepts and therefore be in better position to explain the underlying business to their IT audience. Data Modeling of Financial Derivatives—which presumes no advanced knowledge of derivatives or data modeling—will help you: Learn the best entity–relationship modeling method out there—Barker's CASE methodology—and its application in the financial industry Understand how to identify and creatively reuse data modeling patterns Gain an understanding of financial derivatives and their various applications Learn how to model derivatives contracts and understand the reasoning behind certain design decisions Resolve derivatives data modeling complexities parsimoniously so that your clients can understand them intuitively Packed with numerous examples, diagrams, and techniques, this book will enable you to recognize the various design patterns that you are most likely to encounter in your professional career and apply them successfully in practice. Anyone working with financial models will find it an invaluable tool and career booster.

With Learning JavaScript Design Patterns, you'll learn how to write beautiful, structured, and maintainable JavaScript by applying classical and modern design patterns to the language. If you want to keep your code efficient, more manageable, and up-to-date with the latest best practices, this book is for you. Explore many popular design patterns, including Modules, Observers, Facades, and Mediators. Learn how modern architectural patterns—such as MVC, MVP, and MVVM—are useful from the perspective of a modern web application developer. This book also walks experienced JavaScript developers through modern module formats, how to namespace code effectively, and other essential topics. Learn the structure of design patterns and how they are written Understand different pattern categories, including creational, structural, and behavioral Walk through more than 20 classical and modern design patterns in JavaScript Use several options for writing modular code—including the Module pattern, Asynchronous Module Definition (AMD), and CommonJS Discover design patterns implemented in the jQuery library Learn popular design patterns for writing maintainable jQuery plug-ins "This book should be in every JavaScript developer's hands. It's the go-to book on JavaScript patterns that will be read and referenced many times in the future."—Andrée Hansson, Lead Front-End Developer, presis!

In 1995, David Hay published "Data Model Patterns: Conventions of Thought" -- the groundbreaking book on how to use standard data models to describe the standard business situations. This book builds on the concepts presented there, adds 15 years of practical experience, and presents a more comprehensive view. You will learn how to apply both the abstract and concrete elements of your enterprise's architectural data model through four levels of abstraction: Level 0: An abstract template that underlies the Level 1 model that follows, plus two meta models; Level 1: An enterprise model that is generic enough to apply to any company or government agency, but concrete enough to be readily understood by all; Level 2: A more detailed model describing specific functional areas; Level 3: Examples of the details a model can have to address what is truly unique in a particular industry.

As data management and integration continue to evolve rapidly, storing all your data in one place, such as a data warehouse, is no longer scalable. In the very near future, data will need to be distributed and available for several technological solutions. With this practical book, you'll learn how to migrate your enterprise from a complex and tightly coupled data landscape to a more flexible architecture ready for the modern world of data consumption. Executives, data architects, analytics teams, and compliance and governance staff will learn how to build a modern scalable data landscape using the Scaled Architecture, which you can introduce incrementally without a large upfront investment. Author Piethein Strengholt provides blueprints, principles, observations, best practices, and patterns to get you up to speed. Examine data management trends, including technological developments, regulatory requirements, and privacy concerns Go deep into the Scaled Architecture and learn how the pieces fit together Explore data governance and data security, master data management, self-service data marketplaces, and the importance of metadata

[Copyright: 4ad71e542a387bedc30a6998d9cf985b](#)