



Official Publication of the Northern California Oracle Users Group

# NoCOUG

J O U R N A L

Vol. 23, No. 3 • AUGUST 2009

\$15

## Let Knowledge Ripen at NoCOUG

### Great Expectations

*We interview the newest member of the OakTable network.*

*See page 4.*

### Bind Variable Peeking: Bane or Boon?

*We ask the Oracles.*

*See page 7.*

### First International NoCOUG SQL Challenge

*We have a winner!*

*See page 14.*

***Much more inside . . .***

# The Manga Guide To Databases

A Book Review by Brian Hitchcock

## Details

**Author:** Mana Takahashi and Shoko Azuma

**ISBN:** 978-1-59327-190-9

**Pages:** 224

**Year of Publication:** 2009

**Edition:** 1

**Price:** \$19.95

**Publisher:** No Starch Press



## Summary

**Overall review:** A refreshingly entertaining overview of a topic that is all too often made boring by experts.

**Target audience:** Anyone new to databases.

**Would you recommend to others:** Yes.

**Who will get the most from this book:** Anyone who needs to understand the basic principles of relational databases.

**Is this book platform specific:** No.

**Why did I obtain this book:** While I was attending MySQL Bootcamp, one of the instructors brought this book to class.

## Overall Review

First, what is Manga? I'd say it is a form of graphic novel or comic book from Japan. You may not be familiar with Manga, but you would probably recognize the style of illustration seen in Manga. It has become very popular recently, and now we have a whole series that sets out to teach subjects such as electricity, physics, and databases, and that is illustrated in this style. While it would be easy to dismiss the books in this series as being just for kids, that isn't fair. I got a chance to look at this book while attending a MySQL class and found it interesting. And after reading it, I found it worthwhile.

The material is presented cheerfully, at times bordering on the comical, but the information is accurate and easily understood. For people who aren't already familiar with how relational databases operate, this book is very good. I would recommend it to anyone who needs to understand, at a high level, the most important features of all modern relational database products.

The material is really very serious, but it is presented by means of a simple story, a very Manga-type story of a kingdom that sells fruit and all the problems that come up as their business grows and they try to keep track of everything with paper and pencil. You will not be surprised to find out that just when our heroes appear to have no hope, a database comes to the rescue.

## Preface

Here the author presents the design and structure of the book. We are told that databases are crucial parts of almost all business computer systems and that the true nature of the database is hard to understand. I'm not sure I know the true nature of the database, but it is an interesting concept (something to think about while waiting for support to respond).

## Chapter 1: What Is a Database?

This chapter actually starts with a more important question, namely, why do we need a database? To answer this question, a story is told. It's the story of a kingdom that sells fruit—lots of fruit, and more and more fruit all the time. A princess (Ruruna) is responsible for all aspects of the fruit business and she relies on a humble assistant (Cain). (Do you have any users who act like a prince or princess? I thought so . . .) The princess is struggling to keep up with all the paperwork for the business using her laptop—a laptop with a logo that looks like a piece of fruit, but definitely not an apple.

It is the busy season for the fruit business and the princess is frustrated. There are so many departments, and they sell fruit to many countries. Things are getting out of hand.

Enter the humble assistant with a magic book from the king, a book that describes a secret technology called a "database." Apparently the king found this book in a faraway land. (Someone went to Oracle World on a full pass and attended a lot of marketing presentations.)

Our heroes (the princess and her humble assistant) open the book and release a fairy, a database fairy named Tico. And the database fairy is only visible to those who opened the book. (DBAs appear to be talking to persons who aren't really there all the time . . .)

Tico offers to help our heroes with the supernatural powers of the book in a proper manner. Tico says the first step is to create a database, at which point the princess asks, what is a database? The database fairy is eager to explain, and then observes that our heroes are trying to manage values and numbers, customers and sales, files and departments, and all in an uncoordinated fashion. Data is duplicated in each department, for example. This can create problems. It seems the princess had a crisis just the other day. (This is sounding more and more familiar.) Specific examples of the daily crises are given: the price of apples went up and this required updating data in multiple places. One department didn't update their data, one department updated the price incorrectly, and there was lots of time spent finding and correcting all these errors. Tico offers insight: "You will be tormented by data management even if you do your best, won't you." This hits close to home.

Tormented. The database fairy speaks the truth.

But Tico is here to help, explaining that a database is a system where data can be shared by everyone (cue the SOX fairy . . .) and, listen carefully, if you use a database you “would not have to keep useless data.” Unfortunately, Tico doesn’t explain that last part. Upon hearing this, the princess wants a database very much and now wants to create one and asks Tico, “Can’t you just do it?” Sounds like a consulting gig to me, but Tico replies that the princess and her assistant have to create the database themselves. Tico can only offer guidance. Sounds like a consultant to me.

At this point the Manga story pauses and we have several pages where the ideas introduced so far are described in a more classical style without comic book artwork. This section covers how data is duplicated among departments, how data conflicts arise, update difficulties, and how a database handles all of these issues.

### **Chapter 2: What Is a Relational Database?**

---

This chapter starts, after a brief bit of drama with a prince from a neighboring country (who has no interest in databases, only the princess), by defining some database terms. Tico, the database fairy, explains that this knowledge is necessary before designing a database. Terms such as “record,” “field,” and “product code” are discussed. You can see how this is leading to normalization. We also see that duplicated records are not good. The need for NULLs is brought up and covered very briefly—if only it were this simple.

Next we are told about different data models used in databases. Tico, with a BOINK and a SHAZAM (no, I’m not making this up!) causes a hierarchical data model to appear, followed, with similar sound effects, by a network data model and finally a relational model. The first two models are left simply as diagrams with little explanation, while the relational model is explained in detail. Tables with keys and relational operations are all covered. And then it is time to leave the story to give more detailed explanations and examples of all that was covered so far. The types of data models are further explained. The way the various operations work is illustrated with tables and data for the kingdom’s fruit business. At the end of this section we are told that “The Relational Database Prevails!”—and indeed it does!

### **Chapter 3: Let’s Design a Database!**

---

Now that the princess Ruruna and her assistant Cain have the necessary knowledge, it is time to design the database for the kingdom. Tico appears to present an ER diagram. Entities and relationships are discussed. A many-to-many (way too many?) relationship is illustrated between fruit and export destinations. I like the map showing a horse-drawn cart overloaded with fruit exiting the kingdom. My databases are never this interesting. Tico explains all this by using a chalkboard, a nice retro touch.

Next up is normalization, which requires making more tables. The princess is not happy; she had just one table for all the product sales and was satisfied. But Tico shows us why more tables are better with the example of updating the price of a product. Better to have the price of each product in a

separate table and link that to a table of sales.

Tico is shown leading our heroes up “Mt. Relational Database” as the levels of normalization are explained. Everyone is very tired when this is completed. Having arrived at the top of the mountain, with third normal form, it is again time for a more detailed explanation of the various concepts covered. More traditional examples are given to further illustrate relationships and normalization.

### **Chapter 4: Let’s Learn About SQL!**

---

This chapter opens with more Manga-style drama. The princess and her assistant are strolling through the town. The relationship between them isn’t clear, but the princess obviously perceives that they have some sort of connection. The prince appears and suddenly the princess wants to go into a café with Cain (and not the prince) to “study databases for a while.” I guess this is one way to deal with your relationship issues. And Cain is clueless. But so it goes. Anyway, with this drama out of the way we can now focus on SQL. Why the prince isn’t bright enough to pursue the princess into the café is not explained.

Tico re-emerges from the database book and explains that we need SQL to use the database. Both the princess and Cain look at Tico and in unison ask “Squeal?” Indeed, there is lots of squealing when users start using SQL. Tico explains that when you go to the Swimmy region “across the sea” you need to speak the Swimmy language, and therefore, when you want to have a conversation with a database you need SQL. Ah, sure, the Swimmy Region and Squeal. It’s all coming into focus, isn’t it?

But seriously, we learn how to use SQL to enter data into and retrieve data from the tables we created earlier to model the entities and relationships for the kingdom’s fruit business. When Tico directs our heroes to retrieve data from the product table, both Princess Ruruna and Cain begin to pray for the data to be retrieved. Squealing and praying—clearly Tico has worked with users for a long time! After explaining how to get data from a single table using wildcards and how to sort the data that is returned, joins are discussed. Tico also tells us that a table must have a primary key and that this will prevent the entry of incorrect data. I would offer that incorrect data can be entered at any time, the definition of what would make data incorrect is complicated, and that a primary key really only prevents entering duplicate key values, but Tico seems pretty definite about this so I’ll keep my opinions to myself. Again it is time for a more detailed explanation of SQL, comparison operators, searches, functions, subqueries, and how to create tables. All the basics are covered. With a description of views and several examples for us to study, we are done.

### **Chapter 5: Let’s Operate a Database!**

---

Suddenly, the quiet kingdom has an IT department! We see nameless drones pounding at keyboards (click, click, clack, click) surrounded by lots of screens. While our heroes have learned a lot about databases, they have “some worries.” Cain has done some research and wonders how a database can let a large number of users access data at the same time. But not to worry, he has a set of charts that he will present to us, with a wonderful title of “Database Theater.” My Oracle classes were

never this entertaining. There is much excitement (clapping and a “Whoopie!”) as the show begins. The first chart is presented: “One day, Andy and Becky accessed the database at the same time”—the chart shows two users (stick figures), each with a pointed hat, one labeled “A” and the other “B.” Do your users wear silly hats? Do you see them as cartoonish stick figures?

It turns out that UserA and UserB both read from the database that a total of 30 apples were in stock, but UserA added 10 to that amount while UserB did the same thing. We can see that 50 apples total are now in stock, but the database only reports 40. The princess is quick to accuse Cain of eating the 10 missing apples (Cain is a DBA, he just doesn’t realize it yet).

Cain has more charts of course, and he is quick to offer that a database is designed to process data operations in “lumps.” Indeed, this is the way I’ve always thought of it. Lumps, swirling silently in space. And these “lumpy” operations are called “transactions.” I will assume that a follow-up book will be needed to explain what happens when your database has too many lumps at one time. A lumpy Method-R perhaps? The Automated Lump Repository can’t be far behind, but it will probably require an additional license fee. Cain’s charts provide a very clear and detailed explanation of how transactions prevent UserA and UserB from having a conflict, and locks come into the picture. The princess is impressed with Cain. There are different kinds of locks. Tico whispers in the princess’s ear that Cain is “dependable.” Deadlocks are covered in more charts, and a cancelled transaction is rolled back while completed transactions are committed. But wait, Database Theater is about to provide more relationship drama as the prince appears with a relevant comment about concurrency. Who knew? It seems the prince has appeared because he has invoices from the kingdom that are a real mess, and also because “a database is a nasty thing.” Cain (ever more the hapless DBA) opines that “someone with malicious intent might have performed an unauthorized data overwrite.” What? There’s no chance that a user with a silly hat put in some incorrect data? My money is on UserZ—I never trusted that user. But back to the relationship drama. The prince thinks that as compensation for this inconvenience, the princess should marry him and leave the kingdom to live in his country. Cain isn’t sure why he is upset by this (clueless) but he keeps quiet. Cain proposes that they set up usernames and passwords to control user access. And they will grant specific rights to access specific data to specific users. The princess thinks this is clever. The prince tries to steer the princess back to the subject of marriage, but Cain isn’t done. He also wonders what will happen as more and more users access the growing database. Searches could become slow. The prince asks if it is “safe to trust a database,” foreshadowing future conflicts between his country and the kingdom over things like medical records privacy, et al, but that is getting ahead of the current drama. But Cain (the once and future DBA) is on a roll. He offers that indexing seems to be promising. The prince wants to discuss marriage, but Cain rolls on, explaining all about indexes. The prince and Cain then discuss disaster recovery and logs that keep track of all data operations performed in the database. It isn’t clear why the prince is suddenly able to discuss database issues. Perhaps his country needs a

DBA as well? We are told about rolling forward (and back) transactions as part of disaster recovery. And now, the prince returns to the marriage issue, but the princess suddenly announces her “passion” for Cain and their database.

This is followed by detailed descriptions of transactions, atomicity, concurrency, and query optimization.

## Chapter 6: Databases Are Everywhere!

---

The Manga drama is pretty much over. Now that the princess has a DBA under her control (marriage is a low-cost way to ensure in-house DBA services) she can focus on running her empire. Tico tells our heroes that databases are used in other countries for things like banks, railways and even the Web! It seems that, in other countries, you can by any book you want from the Web! Imagine that! It seems that you can enter search criteria into a web page and that information is used by a database on the Web to find and return data using SQL. Princess Ruruna is shocked to find that her father has a book for sale online titled *Fruit Love*. But we move on. This leads to a discussion of multiple databases to share the load (distributed database, failover) and hints at the issues of keeping multiple databases in sync so that any one of them can take over if the others fail. Stored procedures are described as a way to reduce the burden on the network by doing lots of the processing in the server instead of sending SQL and results back and forth. Triggers are next. Then princess Ruruna concludes that she will, by using databases, build a wonderful country. Wow! These databases can work miracles!

The final section provides more details about databases on the Web, stored procedures, partitioning data, two-phase commit, replication, and more.

## Conclusion

---

Despite what you might think, the material is very accurate and does bring out the important issues. The style is very different from all other database books I’ve seen, and that is very refreshing.

Perhaps I will be replaced by the database fairy—oh wait, that’s right, the fairy can’t do the actual work. Perhaps this fantasy kingdom looks a lot more like my workplace than at first glance?

I think this book has another message, one that is very important for all of us in the database management world. This series of books (search Amazon.com for *Manga Guide*) covers very basic topics such as electricity, physics, statistics, and molecular biology (and Sudoku!). While Manga is enjoyed by many people of all ages, it is primarily aimed at young people. Note the range of subjects that young people are interested in these days, and then realize that databases are right in there with molecular biology (and Sudoku!). What we do all day has become relevant to a much wider audience than we realize. Many high school students have had exposure to MySQL databases as part of websites they frequent. These days, databases really are a part of everyday life and are becoming more important all the time. The fact that a publisher wants a series of Manga titles to cover databases is something worth knowing. The impact of the systems we create and support is huge and growing.

(continued on page 17)

*(continued from page 13)*

Finally, if you should want to read some Manga, I recommend a series titled DeathNote. Full Metal Alchemist is also very good. Another Manga series called Ghost in the Shell is great, but is only available in Japanese. This series is available in the Anime (animation) form in English. You might want to read these, if only so that you can be slightly less un-hip than you currently are. ▲

---

*Brian Hitchcock has worked at Sun Microsystems in Newark, California, for the past 11 years. He is a member of a DBA team that supports 2400+ databases for many different applications at Sun. He frequently handles issues involving tuning, character sets, and Oracle applications. Other interests include Formula One racing, finishing his second Tiffany Wisteria lamp, Springbok puzzles, Märklin model trains, Corel Painter 8, and watching TV (TiVo rules!). Previous book reviews by Brian and his contact information are available at [www.brianhitchcock.net](http://www.brianhitchcock.net).*

*Copyright © 2009, Brian Hitchcock*

*(continued from page 26)*