



Database Answers

Learn Data Modelling by Example



Barry Williams

1. Data Modelling at Windsor Castle	3
1.0 Welcome	3
1.1 Introduction	3
1.2 Topics	4
1.3 Let's go to Windsor	4
1.4 Arriving in Windsor	5
1.5 Starting our Data Model	6
1.6 Identifiers and Primary Keys	6
1.7 Relationships and Foreign Keys	7
1.8 Products and Product Types	9
1.9 Products, Types and Product Hierarchies	12
1.10 Types of People	13
1.11 Types of People and Establishments	13
1.12 Visits and Purchases:.....	15
1.13 People and Inheritance	17
1.14 Staff, Establishments and Derived Fields	19
1.15 Reference Data	20
1.16 Bringing it all Together	21
1.17 Top-Level Model with Names Only	22
1.18 Top-Level Model with Key Fields	23
1.19 Ice Cream	25
1.20 What have we learned ?	25
1.21 Please Email Me	25

1. Data Modelling at Windsor Castle

1.0 Welcome

This is a Tutorial on Data Modelling for Young People

It provides a user-friendly introduction to Data Modelling which we hope you find interesting and easy to read.

It covers the basic concepts and has a very 'user-friendly' approach featuring a teddy bear and kitten creating a Data Model on a trip as tourists to Windsor Castle, which is just outside London, England.

You can find this Chapter as a Tutorial on my Database Answers Web Site :-

http://www.databaseanswers.org/tutorial4_data_modelling_dimple_and_toby_visit_windsor_castle/index.htm

You are invited to follow Data Model developments on our Web Site :-

- http://www.databaseanswers.org/data_models/index.htm

You can also join our Database Answers Community

- <http://databaseanswers.ning.com/>

I hope you enjoy this eBook and would be very pleased to have your comments at barryw@databaseanswers.org.

Barry Williams
Principal Consultant
Database Answers Ltd.
London, England

1.1 Introduction

In this Tutorial, we will follow two young Tourists as they visit Windsor Castle just outside London in England and create a Data Model.

Our Tourists are Dimple, a 10-year old girl, who likes sightseeing and ice cream and Toby, Dimple's 12-year-old brother, who likes sightseeing and designing Data Models.

1.1.1 What is this ?

This is a Tutorial on Data Modelling for Young People that represents a typical Data Modelling Project and illustrates the basic principles involved.

1.1.2 Why is it important ?

Data Modelling is important because it is the foundation for so many activities :-

- It provides a vehicle for communication among a wide variety of interested parties, including management, developers, data analysts, DBAs and so on.
- A physical Database can easily be generated from a Data Model using a commercial Data Modelling Tool.

1.1.3 What will I Learn ?

You will learn :-

- how to create a Data Model, starting from scratch
- what a typical Data Model looks like.

1.2 Topics

In this Chapter, we will cover some basic Concepts in Data Modelling :-

- Primary and Foreign Keys
- One-to-Many and Many-to-Many Relationships
- Hierarchies and Inheritance
- Reference Data

1.3 Let's go to Windsor



[Dimple] : Toby, It's great being in London which is so exciting and buzzing.



[Toby] : I'm glad you like it, Dimple. What would you like to do today ?

[Dimple] : Toby, we have seen Buckingham Palace, where the Queen of England lives, and now I'd like to visit Windsor Castle, because it's one of the most popular tourist attractions in England, and it's just a short trip from London.

[Toby] : OK. Let's go ...

We are starting from Buckingham Palace, where the Queen of England lives ...



Toby and Dimple leave London and arrive in Windsor ...

1.4 Arriving in Windsor

[Dimple] Wow, Toby, Windsor has a beautiful **Castle** and a **Royal Park**.



[Toby] Yes, Dimple, and when we look around there are so many Banks, Cafes, Pubs, Restaurants, Shops, Wine Bars and Hospitals !!!

The other thing that we see when we look around is people - lots of people.

So we can start thinking about our Data Model.



1.5 Starting our Data Model

[Dimple] : How do we get started ?

[Toby] : Well, we know that we have People and Places.
The simplest start is to call all these Places Establishments.
Then we simply have different kinds of Establishments.

And we have People - local people, Tourists, students, people passing through, people working here, people here on business and so on.

[Dimple] : Hmmm - so how do we translate what we know to help us get started with our Data Model?

[Toby] : Let's start a diagram with People and Establishments.

This simple diagram is going to grow into a Data Model.



1.6 Identifiers and Primary Keys

[Dimple] : Toby, I am one of these People so how do I create a unique identity for myself to make me different from everybody else ?

[Toby] : We will give everybody a Unique Identifier and every Establishment its own unique Identifier. When we use these we call them **Primary Keys**, and show them in the diagram with a **PK** on the left-hand side.

[Dimple] : That sounds good, Toby, but I don't know what it means.

[Toby] : Well, Dimple, let's look at how we use these Identifiers ...



Lots of People visit Establishments like Starbucks in Windsor ;-0)



1.7 Relationships and Foreign Keys

[Toby] : Dimple, now we can add some interesting details because we know that one person can visit many Establishments.

We also know that one Establishment is visited by many Tourists.

Then we call this a **Many-to-Many Relationship** between People and Establishments.

To make it easier for you to understand I have expanded the **Many-to-Many Relationship** into two different things, which are called **One-to-Many Relationships**.

[Dimple] : So Toby, is that like saying that One Person can make Many Visits to many Establishments ?

[Toby] : Yes, Dimple - that's great - and we can also say that One Establishment can have Visits from many People.

At this point, we can show how all these boxes are related, and that is a very big step, because it takes us to the idea of 'Relationships'.

We can call these boxes 'Tables' - or Entities if we want to speak to professional Data Modellers'.

A Table simply stores data about one particular kind of 'Thing of Interest'.

For example, People or Establishments.

Each record in a Table will be identified by its own unique identifier, which we call the 'Primary Key'.

It is not usually easy to find a specific item of data already in the Table that will always be unique.

For example, in the States, Social Security Numbers are supposed to be unique, but (for various legitimate reasons) that is not always the case.

Also, foreign visitors and tourists will not have SSNs.

Therefore, it is Best Practice to create a new field just for this purpose.

This will be what is called an 'Auto-Increment' data type which will be generated automatically by the DBMS at run-time.

This is called a Surrogate Key and it does not have any other purpose.

It is simply a Key that stands for something else.

It is a meaningless integer that is generated automatically by the Database Management Software, such as Oracle or SQL Server, The values are usually consecutive integers, starting with 1,2,3,4 and so on.

Now we can see how useful our Identifiers can be because we can include the Person and Establishment Identifiers in our Visits table.

Then the Person_ID field becomes a link to a record for a Person in the Person Table.

This link is what is called a **Foreign Key** and we can see it's shown with 'FK' on the left-hand side.



1.8 Products and Product Types

[Dimple] : Toby, when we go into a shop we want to buy something. And there are thousands and thousands of possibilities. How do we deal with all that in our little Data Model ?

[Toby] : Well Dimple, it's really quite easy. It's like all our Modelling where we look for simple patterns that cover many situations.

[Dimple] : Hmm - I don't know what that means. Maybe if you showed me I might understand it.


[Toby] : OK.

Everything that we buy is called a Product, and all we have to do is simply define the type of each Product - such as a Coffee, Muffin or a Newspaper.

Then we draw a little box called Products and say that every Product has a Type. In other words, there is a Relationship between the Products and Product_Types boxes.

The lines are called Relationships and they are very important in Data Modelling. We are now creating an Entity-Relationship Diagram or "ERD".


This diagram shows only a line for the Relationship :-



The symbol at the Products end is called 'crows-feet' and it shows the 'Many' end.

The short straight line at the Product_Types end shows the 'One' end.

In other words, this line shows a 'One-to-Many' Relationship.



Dimple, let me explain about the dotted line. It means that the Relationship results in a 'Foreign Key' in the Products table. This is shown by the 'FK' symbol next to the **product_type_code** field and it means that there is a link back to the Product_Types.

However, the Primary Key is only the `product_id`, and of course, this is shown by the 'PK' symbol next to the **product_id** field.

Later, when we talk about Inheritance, we will use a straight line, in contrast to this dotted line here. This is to show that the Foreign Key field is also a Primary Key.

I have to say something a bit difficult about Primary Keys right now.

In the Products Table, we have to allow for a very large number of Products being stored.

Therefore we use an ID field for the Primary key.

We then create this ID field automatically as a number (called an auto-increment integer).

This number has no meaning and is simply used to identify each record uniquely among possibly millions or hundreds of millions.

However, things are different for 'Type' fields.

These are what we call enumerated data and are typically Reference Data.

They are always relatively small in number and we choose a Code for the Primary Key because we can create them and review them manually.

It also helps us to create a Code that we can use and refer to, in contrast to the ID fields that have no meaning.

Typical examples would be :-

- Sizes – Small, Medium and Large where we are accustomed to seeing S,M and L.
- Gender – Male and Female, where we use M and F.

- This Menu board at Starbucks shows lots of Products.

We know that they are organised in groups, like Food and Drink, and each of these has more groups and so on, right down to the particular Product, like Caramel Macchiato or a Panini.

This top-down organisations is called a 'Hierarchy' and appears all over the place.

Luckily we can show this very easily and neatly in our Data Model.



1.9 Products, Types and Product Hierarchies

[Dimple] : Toby, when we look closely at the Menu Board to try to decide what to order we can see lots of possibilities

But after a while we can see a pattern that helps us decide.
How do we deal with all that in our little Data Model ?

[Toby] : Well Dimple, it's really quite easy.

We define something called a **Hierarchy**.

Hierarchies are very common and simply mean any situation where there are Parents, Children, Grand-children and so on.

If we look at the Starbucks Menu Board on the right-hand side we can see a simple example of 'Espresso' and under it a number of different drinks.

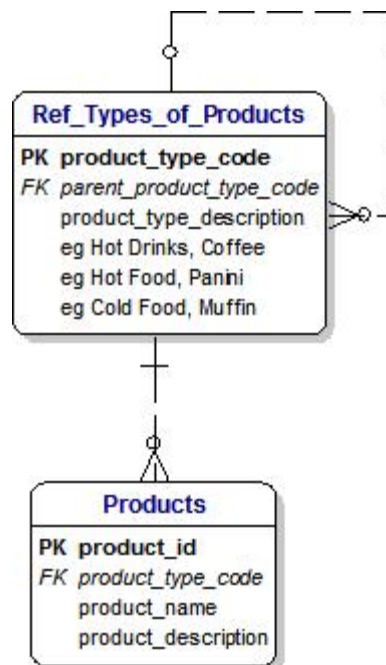
My favourite is Caramel Macchiato.

So in this case, the top-level of our hierarchy is a Product Category called Espresso, and the next level down is a Product called Caramel Macchiato.

[Dimple] : OK. That sounds OK.

[Toby] : Finally, we show this Hierarchy by a dotted line in the top-right hand corner in the Entity called 'Ref_Types_of_Products'.

This is formally called a 'Recursive' or 'Reflexive' relationship and is informally called **'Rabbits-Ears'**



1.10 Types of People

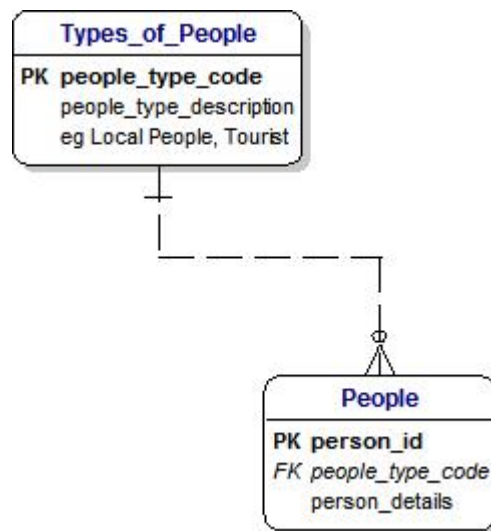
[Dimple] : Toby, that looks OK.

I guess we can deal with Types of People the same way, can we ?

[Toby] : Yes, Dimple, and Types of Establishments as well.

[Dimple] : OK, that sounds sensible. And do they use these Identifiers in a Database ?

[Toby] : Yes, and what is even better is that the Database will automatically generate a new unique Identifier for you and your Visits and Purchases if you want to get a refund later.



1.11 Types of People and Establishments

[Dimple] : Toby, that looks OK.

I guess we can deal with Types of Establishments the same way, can we ?

[Toby] : Yes, Dimple.

[Dimple] : OK, that sounds sensible. And do they use these Identifiers in a Database ?

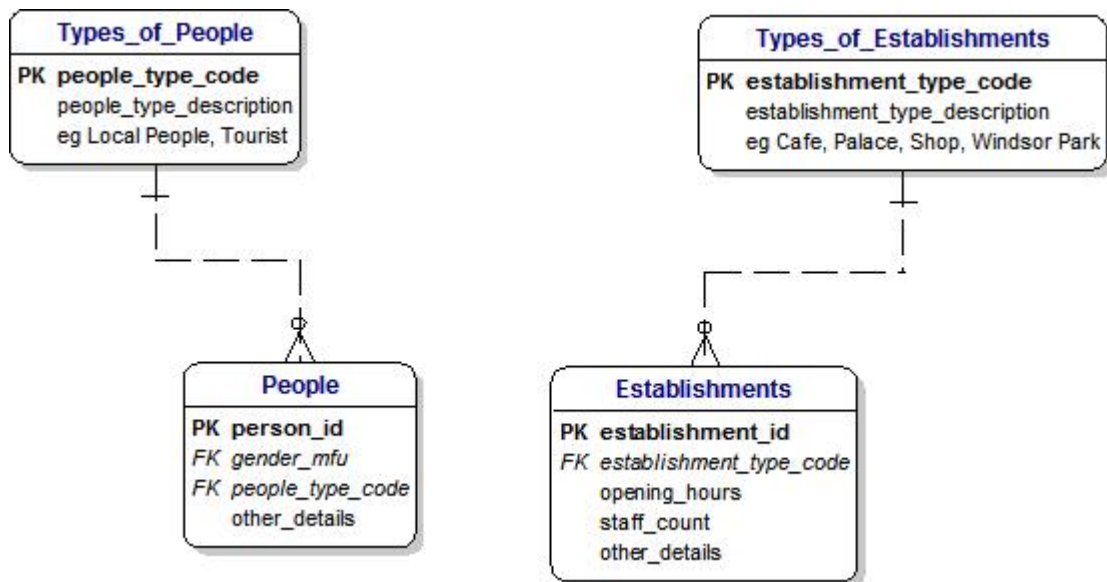
[Toby] : Yes, and we can use our new unique Identifier for you and your Visits and Purchases in case we want to keep track of things.

Like maybe you want to get a refund later so we need to get your details from the Database.

[Toby] : Before we move on, let's talk about Establishments.

One special thing about Windsor is that it has a Castle where the Queen lives and a very large Royal Park, where she keeps Deer.

But when we think about these things, we find that we can simply fit them into our definition of Establishments.



1.12 Visits and Purchases:

Here we can see many visitors to Windsor's Royal Shopping Arcade.



[Dimple] : Toby, with so many People, Establishments and Purchases how do they keep track of everything ?

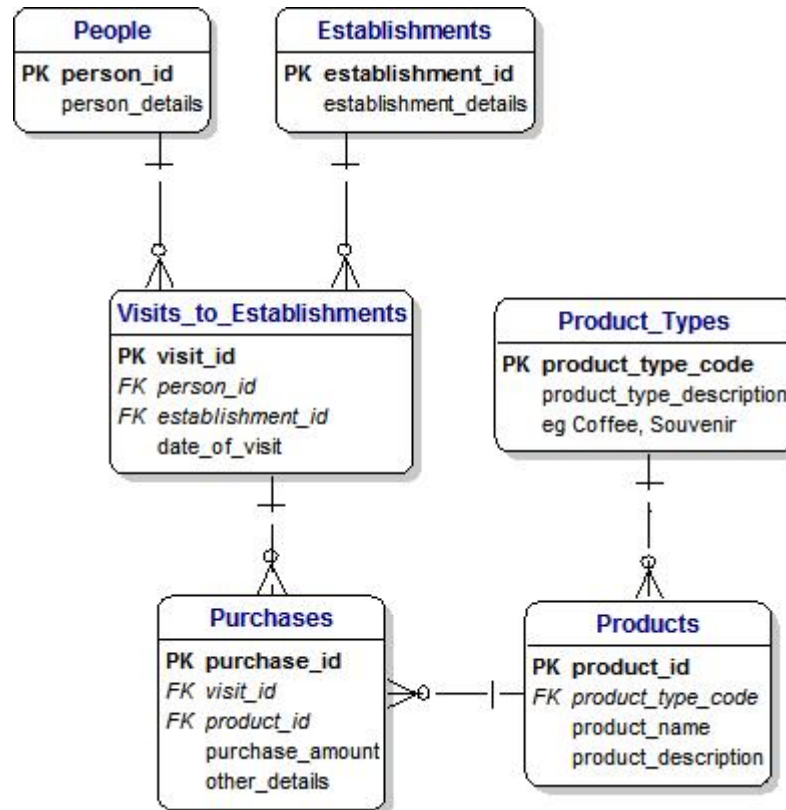
[Toby] : Well, Dimple, by this time, everything has its own Identifier that is used wherever they need to keep track.

[Dimple] : OK, that sounds sensible. And do they use these Identifiers in a Database ?

[Toby] : Yes, Dimple, and in this diagram, we can see that we can use the unique Identifiers which are shown as 'PK', for Primary Keys.

We can see that we have a PK for every Entity or Table so we can be pretty sure we can get from any Table to any other Table.

This is called 'navigating' around the Data Model and is a good test for a well-designed Data Model.



1.13 People and Inheritance

[Toby] : Dimple, let's take a closer look at the different types of people we can find in Windsor.

[Dimple] : OK,Toby. I hope I don't have to think too much because I might get a headache?

[Toby] : No, Dimple, I will do the thinking and talking and all you have to do is nod your head when you understand.

[Dimple] : OK, Toby. I promise to do that.

[Toby] : We already said that we have Local People and Tourists.

There are always lots and lots of People visiting Windsor Castle.

When we look at this picture, we can see Ceremonial Guards in ceremonial red uniforms, and a big crowd, with mainly tourists but also staff in shops responsible for controlling the crowd, Tourists, local people and so on.



Some of these Local People are shoppers and some of them will be working in the shops. We will call the workers 'Staff' and we know different things about them than the things we know about the Tourists.

For example, we will probably know the Gender of everybody just by looking at them.

For Staff, we will usually also know their Date of Birth and their Home Address.

In Data Modelling we have a very powerful approach that we call **Inheritance** that we can use here.

If we want to describe this in English, we would say that Staff inherit the People_Type_Code and Gender from the parent Entity of People, and in addition, they have a Date of Birth and Home Address.

For Tourists, we don't know much, except for the Date of their Visit, and maybe, if they buy something in a shop using a Credit Card, then the shop would know the Credit Card details.

For the Ceremonial Guards in red uniforms, we can tell their Rank by looking at their Uniform and maybe it would also tell us which Unit of the Army they belong to.

Does that make sense, Dimple ?

[Dimple] : I think so, Toby.

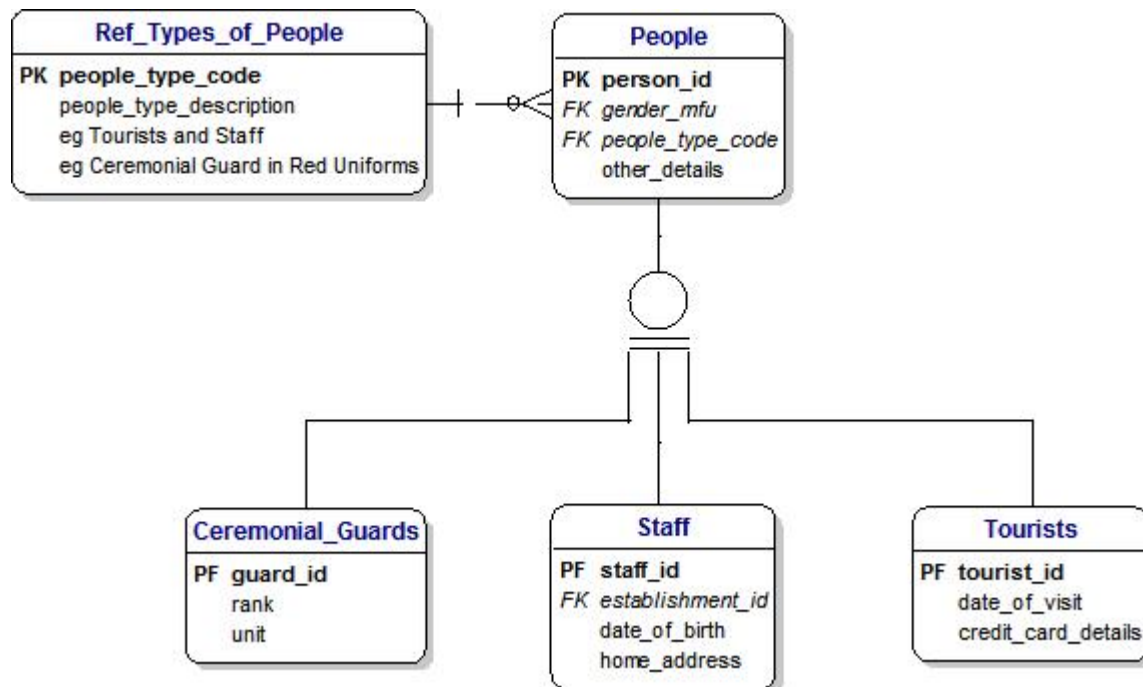
Is it like saying that we inherit having two arms and two legs from our Parents because they have two arms and two legs, but that we have also have things that are just us ?

[Toby] : Yes, Dimple - that's great - let's take a break and do some shopping !!!

[Dimple] : I like the sound of that, Toby. Can I have an ice cream ?

[Toby] : Yes, of course, Dimple – this diagram shows we are doing well.

It show Inheritance between People and the three different types of People :-



We can see a field marked as 'PF' in the three tables for Ceremonial Guards, Staff and Tourists.

This is unusual because it means a field which is a **Primary Key** in the three Tables and also a **Foreign Key** to the People Table.

Therefore, if your first record was a Ceremonial Guard, then we would have a record in the 'People' table with a person_id of 1 and a record in the Ceremonial Guard with a guard_id of 1.

Similarly, if our second record was a member of Staff, we would have a record in the 'People' table with a person_id of 2 and a record in the Staff table with a staff_id of 2.

1.14 Staff, Establishments and Derived Fields

[Dimple] : Toby, how do we specify that Staff must work in some Establishment ?

[Toby] : Dimple, that's a very good question.

Fortunately, the answer is very easy.

We add a One-to-Many Relationship between the Staff and Establishment Entities

In English, we would say that every member of Staff must work in one Establishment and every Establishment can employ many members of Staff.

In the diagram, we show this with a **Foreign Key** by the establishment_id field in the Staff Entity.

So if we look closely at the Staff Entity, we will see '**FK**' by the establishment_id field.

[Dimple] : OK, that sounds good, and I can see how the Identifiers are very important.

[Toby] : I am glad to hear it, Dimple.

There is one more thing I have to say.

We are learning Data Modelling and one important thing about Data Modelling is that it has to follow a set of Rules.

These Rules help us to produce good Data Models and so they are very important.

One of the Rules is that we cannot include any bits of data that can be derived from any other bits of data.

For example, we usually want to know how many people work in a Shop or Cafe.

Therefore we include a **Staff Count** field with the Establishment.

But when it comes to finding the value that goes in here, we will count the records in the Staff Table for each Establishment.


Therefore, it's a **derived Field** and we call it a name that starts with 'derived_' to make things clear.

This is because, according to the rules, we should not include derived fields in our Data Model at this early stage.

I have shown it here simply as an example because it is a situation that occurs quite often so it's good to recognise it when you see it.

Does that sound sensible, Dimple ?

[Dimple] : I suppose so, Toby. But I've got a headache, can we go for an ice cream now ?



1.15 Reference Data

[Toby] : Dimple, you can see that I am using a Gender Table and People Types Table. I have given them both names that begin with 'Ref_' to make it clear that they are Reference Data.

This means that the values don't change much and I can use them to define what the valid values can be.

This is a technique that professional Data Modellers use but we don't need to worry about it today.

[Dimple] : I'm glad to hear it, Toby !!!

Although it isn't difficult to understand and it seems like a good idea.

[Toby] : In our small example, we have only four kinds of Reference Data altogether - Gender, Types of Establishment, People and Products.



1.16 Bringing it all Together

[Toby] : Dimple, if we bring together everything we have talked about, we will see that we have quite a good Data Model that any professional would be proud of.

[Dimple] : OK, Toby. Do you think I will understand it ?

[Toby] : Let me help you by making a list of the **Business Rules** for our Model :-

- People can be either Ceremonial Guards, Staff or Tourists.
- There are a number of Establishments of different Types.
- Tourists can make Visits to Establishments and make Purchases.
- Staff assist the Tourists when they make a Purchase.
- A Purchase involves one Product.

[Toby] : OK, Dimple - we have a very nice Data Model and now we can take the break I promised you.

[Dimple] : That's great, Toby - can I have an ice cream ?

[Toby] : Sure, but before we do I should say something about **PF**, which appears in the Staff Table.

It's unusual and it's called **PF** because it means a field which is a **P**Primary Key in the Staff Table and a **F**oreign Key to the People Table.

[Dimple] : Hmmm, I've got a headache, Toby - can we please go and get an ice cream ?

[Toby] : OK, Dimple. You've been a very good girl and you deserve a break.

You can admire what we have created, which is this very professional-looking Data Model.

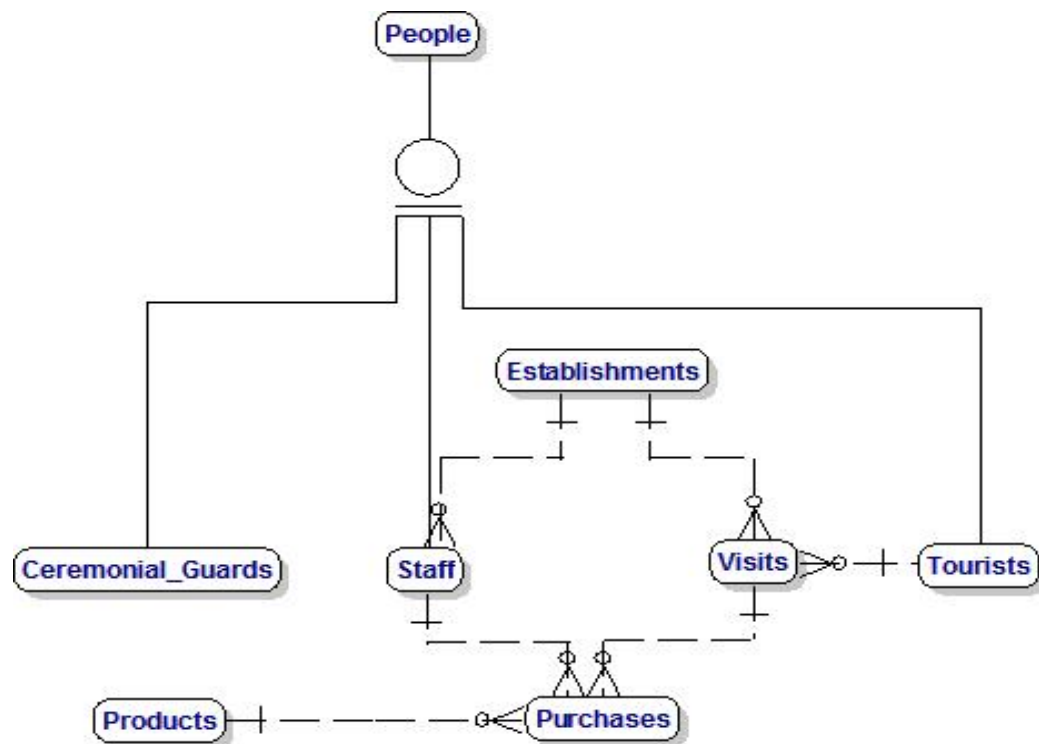
1.17 Top-Level Model with Names Only

We can show our Data Model at the top-level, showing only the names of the 'Things of Interest', which we call Entities or Tables if we are thinking about a Database.

This is suitable for explaining what we saw in Windsor to our family or friends.

If we wanted to describe it, we could simply say :-


- There are lots of People in Windsor, including Ceremonial Guards, Staff and Tourists
- There are also lots of Establishments, like Shops and the Castle !!!
- Tourists made Visits to Establishments where they made Purchases of products.



1.18 Top-Level Model with Key Fields

This is what our Data Model looks like if we show Key fields only and leave out the Reference Data Tables.

This level of display is suitable if we want to confirm to each other how the Tables (or Entities) are related.

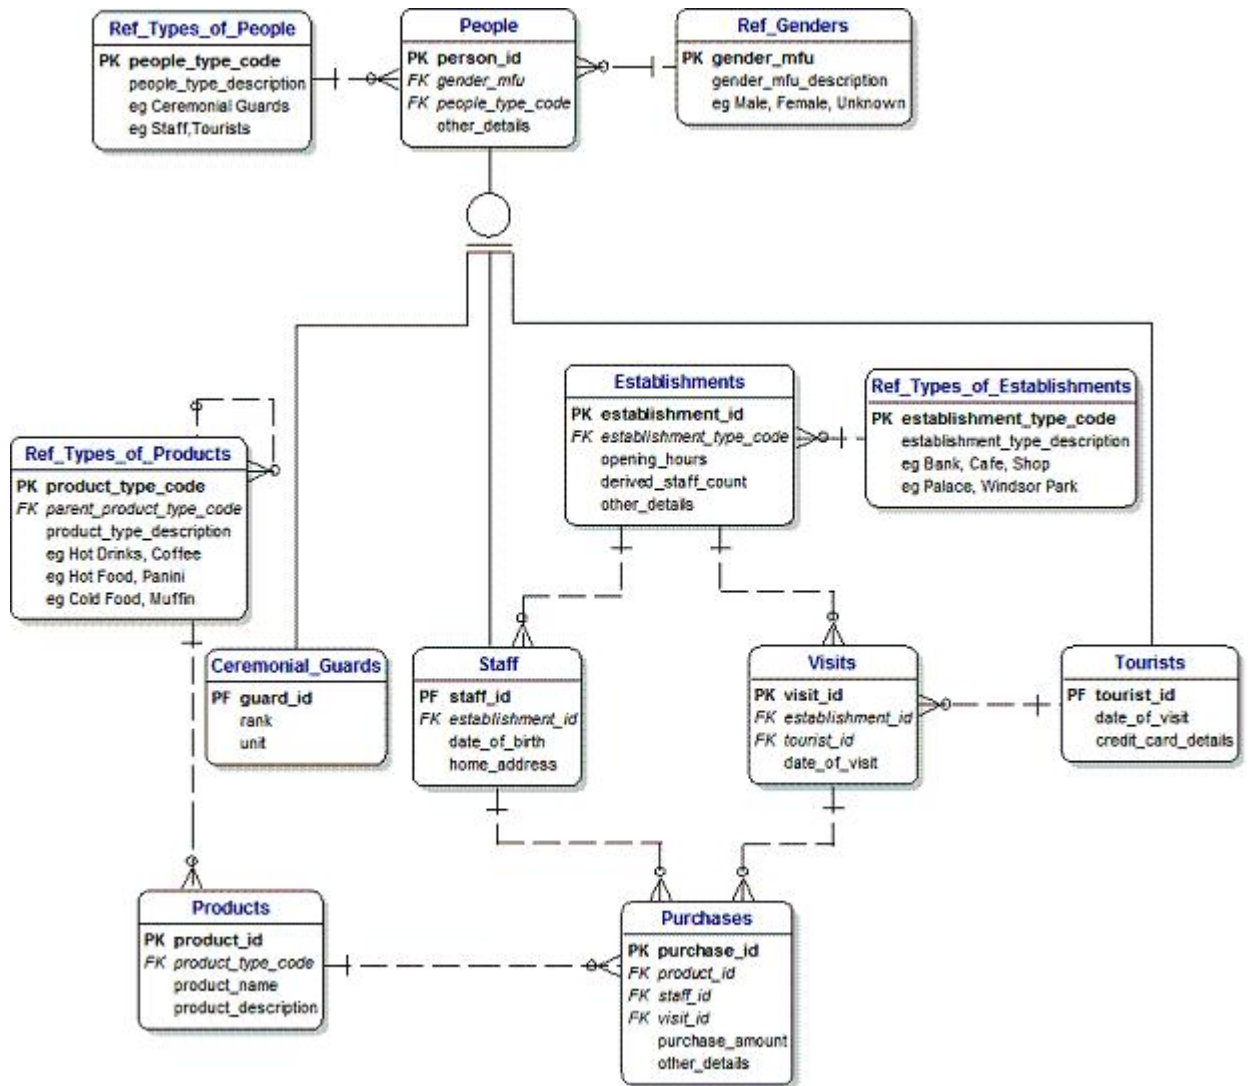


1.19 Top-Level Model with all Details

Finally, this is what our Data Model looks like if we show the Key fields, all the data items only and the Reference Data Tables.

You can see that the amount of detail involved makes it more difficult to understand what's going on and to identify what is important.

This level of display is suitable if we want to talk about details and develop a Database from our Data Model.



1.19 Ice Cream

[Toby] : Dimple, I 've got some wonderful news for you.

[Dimple] : I'm glad to hear it, Toby - what is it ?

[Toby] : I have found your favourite **Baskin-Robbins** ice cream here in Windsor ;)

[Dimple] : Toby, are you teasing me ?

[Toby] : No, Dimple - look, there it is across the road from Windsor Castle !!

[Dimple] : Wow - that's great, so I can have my favourite **Butter Pecan ice-cream**.



1.20 What have we learned ?

In this Chapter, we have learned how to think like a Data Modeller and how to gradually put together a Data Model in our heads.

We know that if we get in the habit of doing this regularly it gets easier and more natural and soon we will be seeing the world around us as pieces of a Data Model that we can fit together like a jigsaw puzzle.

1.21 Please Email Me

I hope you have found this Tutorial fun and useful.

We would be very pleased to have your comments – do you like this Chapter or are there any changes you would recommend ?

You can email us at barryw@databaseanswers.org



Barry Williams

Principal Consultant
Database Answers Ltd.
London, England