

Tutor-marked assignment TMA 02

Introduction

This tutor-marked assignment (TM254 TMA 02) must be submitted by 12 noon (UK local time) on the cut-off date **11 March 2021**.

This module requires all assignments to be submitted electronically. To submit an assignment, please follow the link(s) from your StudentHome page to the online TMA service.

If you foresee any difficulty with submitting your assignment on time, you should contact your tutor well in advance of the cut-off date.

For further information about policy, procedure and general submission of assignments please refer to the Assessment Handbook, which can also be accessed via your StudentHome page.

This assignment accounts for 30 per cent of the assessment score for this module. It has seven questions, for which a total of 200 marks are available. TMA02 is marked out of 200 points and is worth double the marks of TMA01 and TMA03. Questions do not all have the same number of marks; the marks allocated to each question or part question are indicated.

Submission document

You are required to produce a single word-processed document containing all your work, which we refer to as your *solution document*.

Make sure that you create and submit your answers in a file with an acceptable file format. Acceptable file formats are those with extensions of .doc, .rtf or .docx.

Preparing your TMA

Installing the Eventbase database

For question 6 in Part 4 of this TMA you will write SQL queries to satisfy data retrieval requests. We've created a sample database against which you can run your queries and generate answers that you can paste into your TMA solution document.

To install the Eventbase database tables, you should:

1. Locate the file eventbase.txt in the Assessment area of the module website; then download this file, placing it in a folder so that it can be evaluated by the Postgresql system using the \i command in psql.
2. Connect to Postgresql using psql: use the default postgres database, default port, and default postgres user with the appropriate password (TM254 if you followed the installation suggestion in the Software Guide).
3. Use the \i eventbase.txt command in psql to load and run the SQL code, this will create the tables and data for the drone sample database. (Note you may need to give the full path to the location where you placed the drone.txt file: \i 'C:\Users\...\eventbase.txt')
4. To test the installation you should run the query `SELECT * FROM client ORDER BY client_id;` in psql. The resulting table should have 6 rows as shown below.

<u>client_id</u>	name	<u>date_of_birth</u>	location
-	-	-	-
C18362	Felicia Hardy	2002-05-03	Scotland
C35242	William Baker	1980-10-23	Scotland
C66243	<u>Aleksei Sytsevich</u>	1997-01-01	Northern Ireland
C72334	Norman Osborne	1993-10-10	Wales
C83463	Calypso <u>Ezili</u>	1988-10-12	England
C94263	Otto Octavius	2001-08-01	Wales
(6 rows)			

You only need to evaluate the eventbase.txt file's content once, the tables will remain in the postgres database if you leave then reconnect using psql.

When working with Question 6 you may want to download the document eventbase_schema (in the assessment resources) as a handy reminder of the tables and foreign keys in the sample database.

Word limits and layout

There are questions in TMA 02 with specified word limits. Exceeding a word limit by more than 10% will lead to a deduction of 2 marks for that question or part question. If a question or part question has a word limit, state how many words you used. You should aim to write clearly and concisely in English using appropriate terminology for the subject material.

Some answers, for example those with SQL table output, require you to copy and paste text from other sources. Please paste this text into the body of your solution document along with the rest of the answer for the relevant question. Do not place any pasted text separate from your main answer, and do not put it in an appendix or additional document(s).

Remember to ensure that you include a header or footer on each page of your solution document that includes your Open University identifier, your name and the page number; this helps if your tutor needs to print your solutions while marking.

Referencing

Where an answer draws on the module parts or the set books you should cite the source, for example (TM254 Block 1 Part 5) or (Agutter, Chapter 1, Stakeholders) or (Hughes, Chapter 1, Section 2.2). You do not need to give a corresponding reference at the end of your solution document.

Any other sources you use should be cited and fully referenced using the OU Harvard style provided on the Library's Referencing and plagiarism page.

Accessibility and software use issues

If you require additional support for software or with accessing the assessment content – for example with diagram interpretation, drawing package use and suchlike – then please inform your tutor at the earliest opportunity.

Similarly, if you are unable to meet any of the submission requirements then please contact your tutor at the earliest opportunity. Your tutor may not be able to access your solution document if you choose an alternative format for submission; with their advice, you may be able to agree an acceptable alternative. The Technical Help forum is also a good place to discuss issues with the process of submission – but please do not discuss the content of TMAs in the module forums.

Learning outcomes

This assignment assesses learning connected with the following outcomes.

Knowledge and understanding

- KU2 Demonstrate knowledge and understanding of capturing and understanding the requirements of stakeholders and be able to satisfy some of those requirements with a database solution.

Cognitive skills

- CS3 Understand the importance of database administration.

Key skills

- KS2 Communicate effectively about service and project management, requirements and database design.

Practical and professional skills

- PPS3 Use a standard database environment to define and implement a database solution.

Part 1: Quiz results (5 marks)

This question asks you to summarise your engagement with the Block 2 quiz questions. You can complete this question after you have attempted the Block 2 quiz.

Note that the marks you will receive for this question are not based on your final quiz score, but on your reflections on engaging with the questions. You are asked to supply a screenshot of the progress you made through the quiz showing that you engaged with all questions, you will not be penalised for any questions you got wrong.

Question 1 (5 marks)

- Take a screenshot of the question navigation panel from the Block 2 quiz, showing the full extent of your engagement with the questions. Your screenshot should look similar to Figure 1. Paste the screenshot into your solution document as your answer to this part of the question.



Figure 1 A screenshot of the question navigation panel (yours may look different from the example shown)

(2 marks)

- You should write no more than **90 words** for this part of the question.

Write a few sentences in which you tell your tutor about your experience of completing the quiz. For instance, you may have found some Block 2 quiz questions particularly interesting or challenging. In that case, tell your tutor why you found these questions challenging or interesting. Think of this as an opportunity to communicate with your tutor about your experience of studying Block 2.

(3 marks)

Part 2: Database management (20 marks)

This question can be answered after you have completed your study of Block 2 Part 3 and any related activities within that part.

It is intended to allow you to show your understanding of the context of data management and the features found in a modern database management system. In addition, you should be able to show that you understand the importance of data independence and how that relates to the three-schema data architecture.

Question 2 (20 marks)

- a. State why data has both cost and value to an organisation.

(2 marks)

- b. *You should write no more than 150 words for this part of the question.*

Briefly explain why a database management system (DBMS) is much more than just a database

(6 marks)

- c. An existing relational database, developed in accordance with the three-schema architecture, is supporting several applications that share data. Some new requirements are identified by the database users that result in new applications being developed, as follows:

1. A new application requires a new table in the database and uses none of the existing tables.
2. A new application is required that uses exactly the same data already being used by an existing application. However, the performance requirements for the new application require the database to use new hardware so that data retrieval is more efficient.
3. A new application uses tables already in the database that are currently in use by different existing applications; no existing application uses all these tables.

For each of these application's external schema requirements:

- i. state what changes would be required (if any) to each of the existing schemas in the three-schema architecture, and state why that change is required
- ii. state whether or not the required schema changes would affect any existing applications and justify your answer.

(4 marks for each application. 12 marks total)

TMA02 Scenario overview

This scenario is supplied to provide context for the TMA questions in Parts 3 and 4, it will help with interpreting the scenario examples used in questions 3-7.

Eventbase is a company that manages event bookings for clients through an online webshop.

Events are one-day events, such as 'Proms in the park', 'Open-air theatre' or 'Films by the lake', taking place around the country, for which a limited number of tickets are made available. The events may be run several times on several different days.

Clients of Eventbase, who must pre-register with the company, can buy a number of tickets for specific events. Booking for an event generates an e-ticket that is sent to the client and used to gain access to the event on the day booked. As well as booking tickets, the clients can also purchase additional enhancements (add-ons) for the event they hold tickets for. These could include picnics, deck-chairs, and gazebos, for example.

Eventbase also keeps details of the event organisers and descriptions of the locations used for the events.

In order to understand how best to improve the design of the webshop, Eventbase maintain usage statistics for their webshop. The webshop has public access areas and areas that are available only to registered clients (who login to access the site).

A wide range of performance statistics for their website are recorded. These are used, among other things, to evaluate the effectiveness of page designs and to record client's online behaviour (tracking data). The tracking data is based around a viewing action – which is when a user of the webshop (public or client) views a specific webpage. Information about the webpage, who views it, and basic data such as when the viewing action took place, the viewers dwell time on the page, browser type used, etc. is recorded.

Of course, for data held by EventBase, appropriate GDPR permissions have been recorded for each of the users – but that is not in scope for this scenario.

Other details for the scenario are given with the questions that follow.

Note: In order to keep the scenario sufficiently small for the TMA, we've made some simplifications for the scenario (the GDPR permissions being out of scope for the scenario being one example). As you work through the questions in the TMA the specific focus of the question will be made clear and additional information about the scenario will be supplied.

PLEASE ENSURE YOU HAVE INSTALLED THE Eventbase DATABASE BEFORE YOU BEGIN QUESTION 6.

Part 3: Data modelling (55 marks)

The two questions in this part of the TMA are based around Block 2 Parts 4 and 5. They focus on the different ways that data can be modelled, including conceptual data modelling, entity-relationship diagramming and relational representation.

Question 3 is worth 35 marks and Question 4 is worth 20 marks.

Question 3 can be answered after you have completed your study of Block 2 Part 4 and any related activities within that Part. It is intended to allow you to demonstrate your understanding of the purpose of conceptual data modelling and its terminology. You will also demonstrate your ability to produce and interpret entity-relationship notation.

Question 4 can be answered after you have completed your study of Block 2 Part 5 and any related activities within that Part. It is designed to allow you to demonstrate your understanding of the purpose of the logical schema, and the use of the relational model to represent degree and optionality represented by the entity-relationship notation. In addition, the

question allows you to demonstrate your ability to interpret and develop a relational representation from a conceptual data model.

Question 3 (35 marks)

a. Read the following description that extends the Eventbase scenario:

Eventbase needs to keep records of events and the locations at which they are held. Each location has the usual information stored: each has a unique postcode and a (usually) unique name (such as "Parker Manor grounds", but sometimes it's simply something like, 'The Barn' which is why the names are not all unique) and, of course, an address is needed for all locations. Eventbase also need to know the maximum capacity for the location so that it doesn't sell too many tickets for an event, a description of the site accessibility, car park capacity and a general description of onsite facilities, to be used for marketing and event brochures. The events themselves are given a reference number by the Eventbase team, along with a title and the name of the organiser. Each event also has a date and start time, and Eventbase record the total number of tickets available (set by the organiser of the event). Text for event brochures and the website is held for each event – a short version and long version are stored for each event.

List the entity types and relationship types *that can be identified from the information given in the description immediately above*. For the entity types, give their attributes and indicate a possible identifier (primary key) for each entity type.

Use the notation Entity Type(underline>Attribute1, Attribute2, Attribute 3 ...) for entity types, their attributes and the identifier; and the notation Entity Type Relationship Type Entity Type for the relationships.

You are *not* required to draw an entity-relationship diagram, nor are you required to identify any degree or participation conditions for the relationship types.

(10 marks)

b. Figure 2 shows instances of entity types and a relationship type (You have seen similar diagrams in Block 2, Part 4).

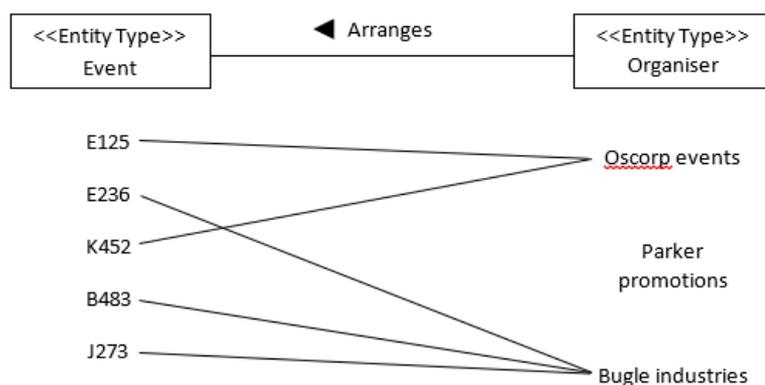


Figure 2 An occurrence diagram

i. How many relationship instances are shown in Figure 2 and how many entity instances, of each type, are involved with these relationship instances?

(3 marks)

- ii. Write down sentences, of the form used in Block 2 Part 4, Section 3 (particularly section 3.2.6, Exercise 4.2) that describe the relationship type in Figure 2. Your sentences should capture the degree (multiplicity) and participation conditions (optionality).

(4 marks)

- iii. Draw an entity-relationship diagram that captures the information given in your sentences describing Figure 2. (Your diagram must show multiplicity and optionality.)

(2 marks)

- c. *You should write no more than 60 words for this part of the question.*

Briefly describe how the data descriptions you produced in your answers to parts a) and b) above would be used when working with stakeholders to develop a database.

(4 marks)

- d. A data analyst has been asked to model the data necessary to record site inspections of locations at which events are held.

Consider the following description of location visits which extend the Eventbase scenario. The details in this description have been agreed by the company as an accurate summary of the visits they undertake.

The suitability of each Location used to host events is checked by a site visit. Eventbase want to record details (such as date, time, visiting employee, findings) of each visit. Each visit is made to a single location.

When a new location is identified (a location not previously used for an event) its data is entered into the Eventbase database and it is recorded as 'inactive' and an initial visit is recorded for that location. The initial visit is an in-depth scrutiny of the location to ensure health and safety arrangements are appropriate and that the site accessibility and facilities available meet the Eventbase minimum standards. If the location doesn't meet the standards then a list of required changes is given to the owners and a return visit is arranged; the location status remains inactive. If the required changes aren't made then the location is not accepted by Eventbase and its status is changed to 'do not use'. When the location is approved its status is changed to 'active'.

Each active location in the Eventbase database is also subject to an annual site visit to ensure it is being properly maintained and the facilities remain unchanged. If the location is no longer acceptable the location is recorded as inactive until a later visit approves the location.

A random visit can be made at any active location at any time; this assures the site condition is satisfactory between the regular, annual visits.

The data analyst has considered the information in the above description. He wants to model the relationship between the locations in the database, and the visits used to inspect the facilities, etc. However, he's not sure of the degree and optionality he should show for the relationship.

He's shown three alternatives in Figure 3, (calling them IsMadeTo1, IsMadeTo2, IsMadeTo3).

In Figure 3, details of all visits, initial, annual and random, are in a Visit entity type. Details of all locations are held in a Location entity type.

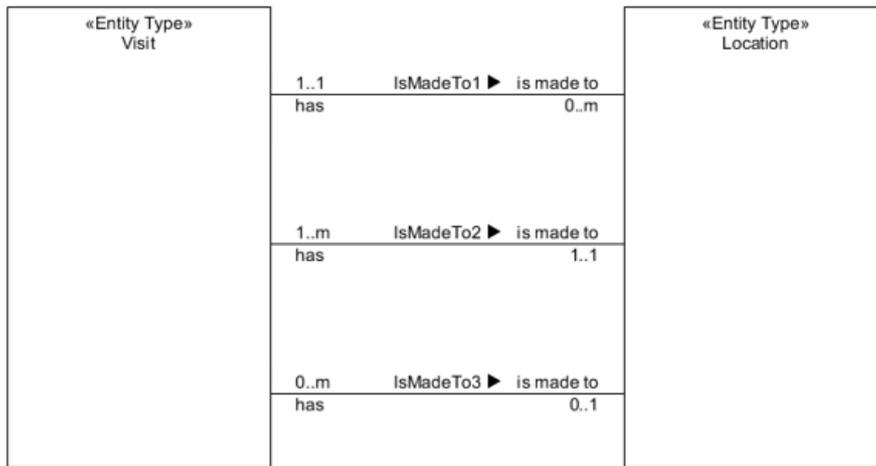


Figure 3 An entity-relationship diagram

Write down sentence pairs, showing the degree and optionality, for each relationship type (IsMadeTo1, IsMadeTo2 and IsMadeTo3) shown on Figure 3 (make sure you label each pair of sentences clearly with the name of the relationship type for those sentences). (You have seen sentences of the form required in Block 2 Part 4 Section 3, particularly section 3.2.6, Exercise 4.2).

Identify, from the three shown in Figure 3, the relationship type that has the degree and optionality that matches with the information given in the short EventBase scenario above it.

Then justify your answer by identifying the mismatches between each of the other two relationship types shown in Figure 3 and the information in the short EventBase scenario above.

(12 marks)

Question 4 (20 marks)

- a. Primary keys and foreign keys are important in relational representations.

Briefly explain:

- i. what they are, and how they relate to each other.
- ii. why foreign keys and the referential integrity rule are important in representing relationships within a database.
- iii. why declaring a primary key represents a constraint on the values possible in a relation.

(8 marks)

- b. Figure 4 is a fragment of a conceptual data model consisting of an entity-relationship diagram showing entity types Ticket, Client and and the relationship type Buys between them.



Figure 4 A fragment of a conceptual data model

Now,

- i. Convert Figure 4 into a logical schema representation using the relation-for-relationship representation.

And,

- ii. Convert Figure 4 into a logical schema representation using the posted-key representation.

You can use either diagrams or text representations for the logical schema.

(12 marks)

Part 4: SQL (120 marks)

The three questions in this part of the TMA are based around Block 2 Parts 6–9. They focus on the Structured Query Language (SQL) – the standard language for defining, controlling and accessing relational databases.

Question 5 is worth 20 marks, Question 6 is worth 70 marks and Question 7 is worth 30 marks.

Question 5 can be answered after you have completed your study of Block 2 Part 6 and any related activities within that part. It is designed to allow you to demonstrate your understanding of the use of normalisation when checking a data model, and the process of translating between logical and physical schema representations.

Question 6 can be answered after you have completed your study of Block 2 Parts 7 and 8, and any related activities within that part. It is designed to allow you to demonstrate that you can write and interpret SQL queries to manipulate data in a given data model to meet stated requirements, and that you can relate an SQL query to the logical processing model.

Question 7 can be answered after you have completed your study of Block 2 Part 9 and any related activities within that part. It is designed to allow you to demonstrate your understanding of the access and security features of SQL and issues around shared data.

Question 5 (20 marks)

- a. Before developing a database solution, Eventbase had been using a spreadsheet to record basic information about events. A fragment of this spreadsheet is used in this question (see Figure 5). The fragment shows the event id, the date of the event, what it's called, where it will be held with details about the location capacity and accessibility, the event organiser and the contact number for the event organiser.

The company used the following relation as the template for the spreadsheet columns. Each row in the spreadsheet holds data about one event:

Event (EventID, Date, Title, Location, Address, Capacity, AccessRestrictions, OrganiserName, ContactNumber)

The following rules are known to always hold between the attributes of the above relation:

1. A value of EventID uniquely determines a single Date value.
2. A value of EventID uniquely determines a single Title value.
3. A value of EventID uniquely determines a single Location value.
4. A value of EventID uniquely determines a single Capacity value.
5. A value of EventID uniquely determines a single Address value.
6. A value of EventID uniquely determines a single AccessRestrictions value.
7. A value of EventID uniquely determines a single OrganiserName value.
8. A value of EventID uniquely determines a single ContactNumber value.

9. A value of Location uniquely determines a single Capacity value.
10. A value of Location uniquely determines a single Address value.
11. A value of Location uniquely determines a single AccessRestrictions value.
12. A value of OrganiserName uniquely determines a single ContactNumber value.

A table of sample data representing the relation is shown in Figure 5.

<u>EventID</u>	<u>Date</u>	<u>Title</u>	<u>Location</u>	<u>Address</u>	<u>Capacity</u>	<u>AccessRestrictions</u>	<u>OrganiserName</u>	<u>ContactNumber</u>
E00282	10/09/22	Proms in the Park	Park Lane	11 Rustic Town	250	Fully accessible.	Parker Promotions	00123 382112
E00275	1/09/22	Stars by Starlight	Haven Field	10, Half-acre Wood	300	Fully accessible.	Bugle industries	555 12381
E00271	1/08/22	Twilight Tales	Lakeside Manor	18, The Lake	145	Long steep path to lakeside. Site vehicle available to assist.	Bugle industries	555 12381
E00283	10/09/22	Tea at Tall Trees	Tall Trees	Windy Lane	60	Disabled access to carpark at side of site.	Parker Promotions	00123 382112
E00284	11/09/22	Proms by the Lake	Lakeside Manor	18, The Lake	145	Long steep path to lakeside. Site vehicle available to assist.	Oscorp events	00551 99231
E00285	14/06/22	Tea Dance in the Park	Park Lane	11 Rustic Town	250	Fully accessible.	Parker Promotions	00123 382112
E00290	15/09/22	Stars by Starlight	Haven Field	10, Half-acre Wood	300	Fully accessible.	Bugle industries	555 12381
E00291	17/09/22	Proms by the Lake	Lakeside Manor	18, The Lake	145	Long steep path to lakeside. Site vehicle available to assist.	Bugle industries	555 12381

Figure 5 The relation Event represented by a table of sample data

a.

- i. Using the data in Figure 5, and the 'uniquely determines' rules 1-12, describe, with an example, how an update anomaly might arise.

(4 marks)

- ii. Briefly explain why Event is in first normal form.

(3 marks)

- iii. Briefly explain why Event is in second normal form.

(2 marks)

- iv. Event is not in third normal form. Identify what prevents Event being in third normal form and then split the relevant attributes into new tables so that the resulting tables are in third normal form.

(5 marks)

(14 marks)

- b. Convert the relational representation shown in Figure 6 to appropriate SQL Data Definition Language statements to create tables representing the same data requirements.

Recall, from the TMA02 scenario overview, that an add-on is an additional item such as a picnic, deck-chairs or gazebo booked by a client for an event. In the table below the AddOnID is the identifier for each additional item. You can assume that the referenced domains have already been defined in SQL as: AOID, EID, CID and NameString, and that the SQL INTEGER datatype is available. (You are not required to define any other tables from the Eventbase scenario).

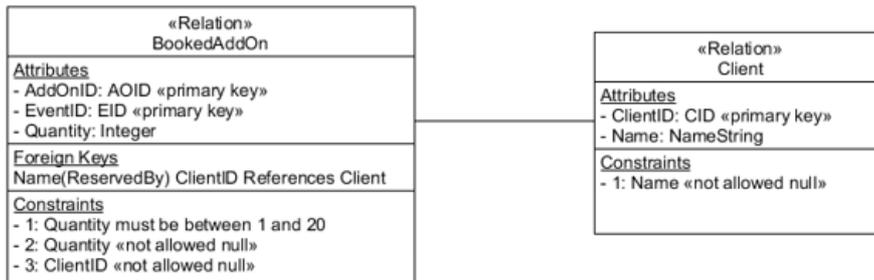


Figure 6 A relational model of the relationship ReservedBy between the relations BookedAddOn and Client.

(6 marks)

Question 6 (70 marks)

This question requires you to use the tables of the Eventbase described in the introduction to the TMA. You will need to install the tables and data by running the eventbase.txt file in psql which will create tables according to the schema in Figure 7 and insert data into these tables.

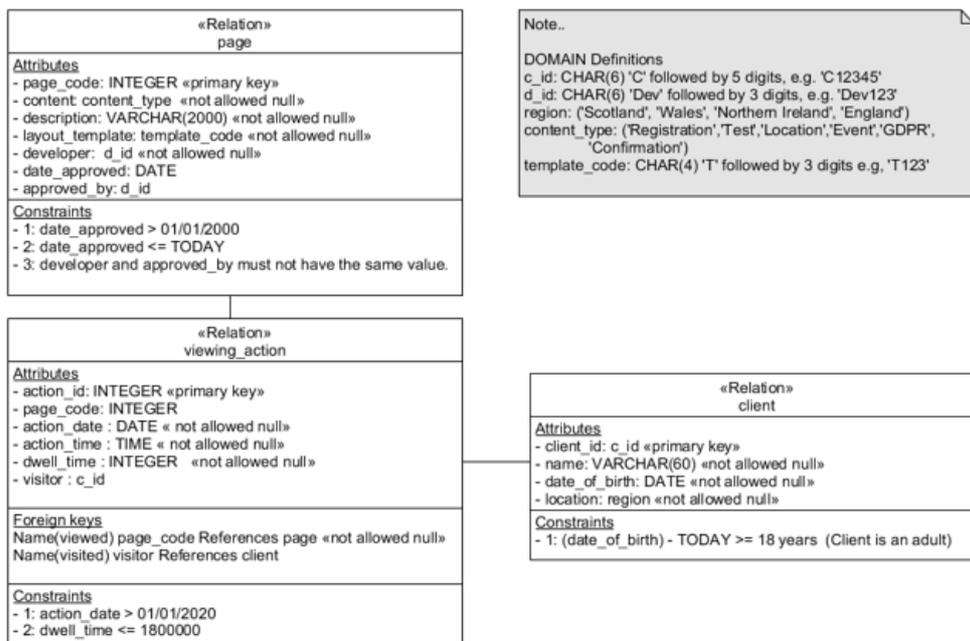


Figure 7 A schema and description of the tables defined in the eventbase.txt file (note this is also given in the Eventbase_schema document).

Eventbase schema description

The Eventbase website consists of HTML pages - each page has a unique identifying `page_code`, and is characterised as being a particular type of page (e.g. Registration). The page has a more detailed description of its content (e.g. allows date/location selection). Each page is built using a particular webpage template by a developer and the page is approved, for use on the website, on a given date by another identified developer (a developer cannot approve the same webpages they have developed).

Clients are given a unique client identifier when they register with Eventbase. Eventbase records their name, date of birth and home location. (Other data would also be recorded; but for this exercise this is sufficient).

A viewing action is recorded when a visitor is served a page on the website, the action date and start time are recorded, along with the dwell time on that page. The dwell time is the amount of time the visitor remains viewing that page, recorded in milliseconds. There is a maximum, 30 minutes (1,800,000 milliseconds), for the dwell time, beyond which the time is recorded as 1,800,000 milliseconds. Each `viewing_action` is given a unique identifier automatically by the system.

A visitor to the website is a client (logged onto the site), in which case their `client_id` is recorded, or is a guest (not logged onto the site) in which case the `client_id` is recorded as NULL. (Note in a real application most companies would give some form of tracking identifier, even to guest visitors, so they could see extract data about every individual guest's behaviour.)

In your solution document please include your SQL statements and a copy of the output produced by the evaluation of each statement. It is important that you include the output otherwise your tutor will not know that your statement produced any output.

a. Write SQL statements that answer each of the requests for data in the Eventbase.txt tables:

- i. Produce a list of the developer identifiers that are recorded as having developed or approved at least one page. The list should not include duplicate values.

(3 marks)

- ii. For each page that has not yet been approved, give the page code and developer. (Note unapproved pages have no `date_approved` and no `approved_by` values). Rename the `page_code` column as 'unapproved' in the resulting table.

(4 marks)

- iii. How many pages are recorded where the content type is 'Registration'?

(3 marks)

- iv. How many different developers have approved at least one page?

(3 marks)

- v. List the page code, content type, and layout template for all pages where the description includes mention of 'index'.

(3 marks)

- vi. For each viewing action, give the action id, page code and dwell time for that viewing action, sort the resulting table by order of dwell_time, lowest to highest.

(4 marks)

- vii. List the client id, name and date of birth for clients who have a location of Scotland or Wales, the list should be in alphabetic order of name.

(6 marks)

- viii. For each page, list the page number and the number of viewing actions for that page.

(4 marks)

- ix. List the page code for every page that has never been viewed.

(4 marks)

- x. Which pages, with the template code 'T233', have never been viewed for longer than 1 second, in any viewing action?

(5 marks)

- xi. For each developer who has develop at least one page, list the developer and the total dwell time spent viewing all pages designed by that developer. Show the list in order of total dwell time, highest to lowest.

(5 marks)

(44 marks)

- b. What request does the following SQL answer? (Note that a request should be an English language question or description like the ones in part a), *not* an account of how the data is processed to produce the result.)

```
SELECT page.page_code, developer
FROM page JOIN viewing_action
ON (page.page_code = viewing_action.page_code)
WHERE dwell_time < 1800000
GROUP BY page.page_code, developer
HAVING (AVG(dwell_time) < 3000);
```

(6 marks)

- c. Describe, using the logical processing model, the evaluation of the following query:

```
SELECT page_code
FROM viewing_action
WHERE (dwell_time >= 15000)
GROUP BY page_code
HAVING count(visitor) > 3;
```

(8 marks)

d. Consider the information:

A new page, given the page code 10001231, is created. It has content type 'Test' and a description of 'exploring a new page construction technique'. It uses a 'T394' template type. It was developed by developer 'Dev311' but has not yet been approved.

Add this information to the maintenance table using an SQL INSERT statement, then write a query to show that your insertion was successful.

(3 marks)

e. Describe the effect of the NATURAL JOIN operator in SQL?

What are the risks in coding an SQL query in an application that includes the NATURAL JOIN operator?

(4 marks)

f. The company wants to maintain a table showing the minimum and average dwell times for each page that has been viewed and the number of times the page has been viewed.

Create an SQL VIEW with the name `dwell_time_stats` that selects rows from the `viewing_action` table. The view should show the `page_code`, minimum dwell time with the heading `min_dwell`, the maximum dwell time with the heading `max_dwell`, and the number of times the page has been viewed with the heading `view_count`.

Then, write a query showing the data the view defines.

Your solution document should contain the CREATE VIEW statement, the query and the output of the query.

(5 marks)

Question 7 (30 marks)

a. *You should write no more than 300 words for this part of the question.*

In the full Eventbase client record keeping system, as well as basic name and date of birth data there is more confidential and sensitive information held about the clients – including addresses, credit card data, health/accessibility requirements (to be able to advise on location accessibility and/or event first aid facilities), car registration details for car-park booking, and so on.

Eventbase exchange data with outside companies for some of the services they arrange for their clients. For example:

- when they send brochures to clients they send the name and address data to a bulk mailing company,
- when a client books car-parking at a location, the local car-park management company are sent the client name and car registration data for that event.

When obtaining GDPR permissions, the clients are informed that Eventbase will hold all client data securely on their data servers and itemise the specific data items that will be shared with outside companies.

Considering the security and data management facilities for relational databases **presented in this module**, describe how the company could ensure the client information can be adequately protected.

Your description should consider the full client table (which would contain all the details of their clients) and how only selected data can be made available to the applications and users that share data with the outside companies. You should assume that different users (marketing, sales, event handling, etc.) should only be given access to the data relevant to their roles within Eventbase.

You should draw on the principle of least privilege and the rules for controlled sharing when answering this question.

You do not need to demonstrate any SQL code for your answer to this question, but you might refer to SQL features in your answer.

(15 marks)

b. *You should write no more than 300 words for this part of the question.*

During the initial prototype evaluation of the webshop, Eventbase discovered that some clients would add tickets to their online basket and then browse through the available add-ons for the event adding some add-ons to the basket. However, when they went to pay for the content of their basket they would be told that the event tickets could not be purchased as they had already been sold and the entire basket sale is rejected. This occurred most often for popular events when only a few tickets remained available and when the user had spent some considerable time browsing through the add-ons. When the data was examined, the tickets that had been in the user's basket were showing as sold to another user.

Briefly explain why the Eventbase webshop application described above could use a transaction management system capable of realising the ACID properties for transactions to prevent the above problem arising. You should include a description of how the above problem might have occurred, and which of the four ACID properties this particular example demonstrates.

Why are the other three ACID properties important to Eventbase when developing the webshop? Give examples of the potential problems they could prevent.

You do not need to write any SQL code for your answer to this question, but you may refer to SQL features in your answer.

(15 marks)
