**TECHNISCHE UNIVERSITEIT DELFT**
**Faculteit Elektrotechniek, Wiskunde en Informatica**

Final

# TI1506 Web and Database Technology

Friday, January 29th 2016
09.00-11.00

**INSTRUCTIONS:**

- This exam consists of 2 parts (DB and Web) and a total of 38 multiple-choice questions. All questions are worth an equal number of points.
- The usage of books, notes, old exams, and other written resources is explicitly **FORBIDDEN** during the exam. The use of electronic aids such as smart-phones, laptops, etcetera, is **ALSO NOT** allowed.
- There is only one right answer for each question. If you think there are more, pick the best one.
- You are not allowed to make corrections on the multiple-choice answer form (MAF). You are therefore advised to first mark the answers on this exam and later copy them to the MAF. If you need to make corrections anyway, ask for a new form and copy all your answers to it.
- You are not allowed to take the exam sheet with you after the exam. We will publish online the text of the exam together with its solutions.
- Note that the order of the answers on your MAF form is not always A-B-C-D.
- Be sure to fill in all header information on the MAF. Enter your *student number* on the form with digits as well as by filling the boxes.
- Sign the MAF. Without your signature, the form is not valid. Since you might forget this at the end, you are advised to do this at the start of the exam.

## Part 1 – Web (17 questions)

**QUESTION 1.** Consider the server-side code below running at `my.site.nl`.

```
1  var express = require("express");
2  var http = require("http");
3  var credentials = require("./credentials");
4  var cookies = require("cookie-parser");
5  var sessions = require("express-session");
6  var app = express();
7  app.use(cookies("my-not-so-scret-secret"));
8  app.use(sessions("my-not-so-scret-secret"));
9  http.createServer(app).listen(3000);
10
11 app.get("/countMe", function (req, res) {
12     var session = req.session;
13     if (session.views) {
14         session.views++;
15         res.send("You have been here " + session.views + " times!");
16     }
17     else {
18         session.views = 1;
19         res.send("This is your first visit!");
20     }
21 });
```

A user starts up his browser (which contains no cookies so far) and accesses `http://my.site.nl:3000/countme`. He then clicks the browser's "*Reload current page*" button **three** times. How many times in the process does the server running at `my.site.nl:3000` send a cookie to the browser?

**A)** 0
**B)** 1
**C)** 3
**D)** 4

**QUESTION 2.** Consider the following node.js code:

```
1  var express = require("express");
2  var http = require("http");
3  var app = express();
4
5  app.use(function (request, response) {
6      response.writeHead(200, { "Content-Type": "text/plain" });
7      response.end("Hello there!");
8  });
9
10 app.use(function (request, response, next) {
11     console.log(request.method + " -> " + request.url);
12     next();
13 });
14
15 app.use(function (request, response, next) {
16     response.writeHead(200, { "Content-Type": "text/plain" });
17     response.end("Hi there!");
18 });
19
20 http.createServer(app).listen(3000);
```

What happens when a user accesses `http://localhost:3000` in the browser (assuming that the node.js script was started on the same machine)?

**A)** On the server-side, the HTTP request is logged to the console; the browser displays "Hi there!".
**B)** On the server-side, the HTTP request is logged to the console; the browser displays nothing.
**C)** On the server-side, the HTTP request is not logged to the console; the browser displays "Hello there!".
**D)** On the server-side, the HTTP request is not logged to the console; the browser displays "Hello there! Hi there!".

**QUESTION 3.** Consider the two files, `foo.js` and `bar.js`:

`foo.js`:

```
1 module.exports.username = "my_username";
2 module.exports.password = "my_password";
3 module.exports.color = "blue";
4 module.exports.color = function () {
5     return "red";
6 };
7 exports = module.exports.color;
```

`bar.js`:

```
1 var foo1 = require("./foo");
2 foo1.password = "my_admin";
3 var foo2 = require("./foo");
4 console.log(foo1.password);
5 console.log(foo2.password);
```

What is the console output when running `node bar.js` ?

**A)** my_admin
my_admin

**B)** my_admin
red

**C)** my_admin
my_password

**D)** my_admin
blue

**QUESTION 4.** Consider the two files, `foo.js` and `bar.js`:

`foo.js`:

```
1 module.exports = function() {
2     return {
3         username : "my_username",
4         password : "my_password",
5         color : "blue"
6     }
7 };
```

`bar.js`:

```
1 var foo1 = require("./foo")();
2 foo1.password = "my_admin";
3 var foo2 = require("./foo")();
4 console.log(foo1.password);
5 console.log(foo2.password);
```

What is the console output when running `node bar.js` ?

**A)** my_admin
my_admin

**B)** my_admin
{}

**C)** my_admin
my_password

**D)** {}
{}

**QUESTION 5.** The browser B currently has no stored cookies. The server sends the following four cookies to B:

```
Set-Cookie: bg=white; Expires=Fri, 01-Apr-2016 21:47:38 GMT; Path=/;
                Domain=tudelft.nl; HttpOnly
Set-Cookie: pref=1; Path=/; Domain=tudelft.nl
Set-Cookie: dom=23; Expires=Thu, 01-Jan-1970 00:00:01 GMT; Path=/;
                Domain=tudelft.nl; HttpOnly
Set-Cookie: view=mobile; Path=/; Domain=tudelft.nl; secure
```

B crashes 10 minutes later and the user restarts B. How many cookies can the user access after the restart with client-side JavaScript (i.e. `document.cookie`)?

**A)** 0
**B)** 1

**C)** 2
**D)** 4

**QUESTION 6.** After accessing `http://www.login.meebo.com` for the first time, the server sent the following four cookies to browser B:

```
Set-Cookie: ID1=32s; Path=/todos; Expires=Fri, 30 Jan 2018 01:01:01 GMT
Set-Cookie: ID2=532; Domain=meebo.com; Path=/;
                  Expires=Fri, 30 Jan 2019 01:01:01 GMT
Set-Cookie: ID3=ssd33dd; Domain=login.meebo.com; Path=/admin;
                  Expires=Fri, 30 Jan 2018 01:01:01 GMT; HttpOnly
Set-Cookie: ID4=bf1; Domain=www.login.meebo.com; Path=/todos;
                  Expires=Fri, 30 Jan 2018 01:01:01 GMT; HttpOnly
```

Next, B tries to access `http://login.meebo.com/admin`. How many of the four cookies are sent back to the server?

**A)** 1
**B)** 2
**C)** 3
**D)** 4

**QUESTION 7.** Consider the following URL route defined in a node.js script:

```
1 app.get("/get(My)?[Tt]*dos+", function (req, res) {
2     /* send HTTP response */
3 });
```

Which of the following routes will **not** match the route defined above?

**A)** `/getMyTodos`
**B)** `/getTdos`
**C)** `/gettodoss`
**D)** `/getMyMyTodos`

**QUESTION 8.** Consider the following node.js script. It should enable the use of routing parameters.

```
1  var express = require("express");
2  var http = require("http");
3  var app = express();
4
5
6  var calendarEntries = {
7     /* missing */
8  };
9
10 app.get("/calendar/:A/:B", function (req, res, next) {
11     var c = calendarEntries[req.params.A][req.params.B];
12     res.send(c);
13 });
14
15 http.createServer(app).listen(3000);
```

How should `calendarEntries` be implemented so that routes of the form `/calendar/today/work` and `/calendar/tomorrow/other` will send the list of correct calendar entries in the HTTP response?

**A)**

```
var calendarEntries = {
   today : [
         work: ["TI1506 exam","Sign up for Q3 courses"],
         today: other: ["Shopping","Book holidays"],
   ]
tomorrow : [
         work: ["Buy Q3 books"],
         tomorrow: other: ["Football competition", "Repare bicycle"]
   ]
};
```

**B)**

```
var calendarEntries = {
   today: {
         work: ["TI1506 exam","Sign up for Q3 courses"],
   other: ["Shopping","Book holidays"]
},
      tomorrow: {
         work: ["Buy Q3 books"],
         other: ["Football competition", "Repare bicycle"]
      }
};
```

**C)**

```
var calendarEntries = {
   today: function() {
         work: ["TI1506 exam","Sign up for Q3 courses"],
    other: ["Shopping","Book holidays"]
},
      tomorrow: function() {
         work: ["Buy Q3 books"],
         other: ["Football competition", "Repare bicycle"]
```

```
   }
};
```

**D)**

```
var calendarEntries = {
   return function() {
      today: {
               work: ["TI1506 exam","Sign up for Q3 courses"],
             other: ["Shopping","Book holidays"]
      },
      tomorrow: {
               work: ["Buy Q3 books"],
            other: ["Footbal competition", "Repare bicycle"]
      }
   }
};
```

**QUESTION 9.** Which of the following statements about third-party cookies is **correct**?

- **A)** Third-party cookies and first-party cookies are stored in different cookie storage facilities within the browser.
- **B)** Third-party cookies originate from a different domain than first-party cookies.
- **C)** Third-party cookies are necessary to track users across a single Web application.
- **D)** Third-party cookies have a lower priority than first-party cookies and are returned to the server only after any first-party cookies.

**QUESTION 10.** A server-side application uses sessions instead of cookies to track users. What is the most common approach to determine the end of a session?

- **A)** The cookie the client sends with the final HTTP request to the application contains a special `Sesssion-Ending` cookie field to indicate the end of the session.
- **B)** The server makes an HTTP request to the client every *x* seconds to determine whether the client is still online. If the client is not online, the session ends.
- **C)** The client makes the final HTTP request to the application without returning its session cookie, indicating the end of the session.
- **D)** If x seconds have passed without a request from the client, the server ends the session.

**QUESTION 11.** What is the main purpose of the "consumer secret" (or "client secret") in third-party authentication?

- **A)** It ensures that only authorized applications query the authentication server for an access token.
- **B)** It ensures that the the token issued by the authentication server is encrypted correctly.
- **C)** It enables the consumer of the access token (the application) to compute a checksum of the token – if the checksum is not the same as the consumer secret, the token is considered to have been tampered with by a malicious user.
- **D)** It enables the authentication server to determine which type of access token (*ephemeral* or *continuous*) to return to the consumer application.

**QUESTION 12.** What is the major difference between the commands `npm install --save ABC` and `npm install --save-dev ABC` ?

**A)** The `--save` flag indicates that the package `ABC` will be installed in the global node.js package directory of the machine. The `--save-dev` flag indicates that the package `ABC` will only be installed for particular types of devices (found in the `package.json` file).

**B)** The `--save` flag indicates that the package ABC will be installed and listed in the dependencies of the application. The `--save-dev` flag indicates that the package ABC will be installed and listed in the devDependencies of the application.

**C)** The `--save` flag indicates that the package `ABC` will be installed. The `--save-dev` flag indicates that the package `ABC` will only be installed if the application shows deviant behavior (i.e. throws an error).

**D)** There is no difference between `--save` and `--save-dev`; `--save-dev` is deprecated and in future version of node.js will be replaced by `--save`.

**QUESTION 13.** Consider the following list of abilities a malicious user (the attacker) may have after having managed to intercept all of your server's inbound network traffic:

[1] The attacker can eavesdrop (i.e. read all HTTP requests your server receives)
[2] The attacker can inject additional HTTP requests bound for your server
[3] The attacker can modify HTTP requests bound for your server
[4] The attacker can drop (i.e. delete) HTTP requests bound for your server

You have developed a highly popular Web-based newsticker application that in real-time sends the latest news (received from one hundred different news agencies) to its users when opening the newsticker Website in the browser. Which of the listed abilities is **required** for an attacker to conduct a denial-of-service attack against your application?

**A)** [1] and [2]
**B)** [2] and [3]
**C)** Only [4]
**D)** None of these abilities are required.

**QUESTION 14.** You write a Web application that contains a session management component. Clients who want to use your application have to authenticate and subsequently a session ID is used to guarantee them access to your application. Once they log out of your application, two steps should be implemented: (1) the session ID is to be deleted from the client and (2) the session ID is to be set to expired on the server. Which attack is enabled if the implementation of step (2) is forgotten?

**A)** CSRF
**B)** Reflected XSS
**C)** Insecure Direct Object References
**D)** Missing Function Level Access Control

**QUESTION 15.** You have developed a Web application (e.g. a TODO application) with multiple user accounts. Which threat is your Web application susceptible to if you can manipulate the URL of an account page to access other account pages?

**A)** Reflected XSS
**B)** CSRF
**C)** Insecure Direct Object References
**D)** Injection

**QUESTION 16.** Which of the following approaches is most likely to secure an application against cross-site request forgery?

**A)** Use state-of-the-art encryption algorithms to store the data on the server.
**B)** Use reauthentication and CAPTCHA mechanisms.
**C)** Avoid the use of direct object references; the use of objects should include an authorization subroutine.
**D)** Validate all user input and escape generated output.

**QUESTION 17.** Which of the following approaches is **not** suitable for an attacker to attempt the the injection of malicious data into a server-side application?

**A)** Appcache manipulation
**B)** URL parameter manipulation
**C)** Hidden HTML field manipulation
**D)** Cookie manipulation

## Part 2 – Database (21 questions)

**QUESTION 18.** Which of the following SQL set-operators cannot be expressed using other constructs of the SQL language?

**A)** INTERSECT
**B)** EXCEPT
**C)** UNION
**D)** MINUS

**QUESTION 19.** Which of the following is **not** a property of a Relational Catalogue in RDBMs?

**A)** The Relational Catalogue contains the data dictionary
**B)** The Relational Catalogue contains the conceptual representation of the data in the database
**C)** The Relational Catalogue contains statistics about the data in the database
**D)** The Relational Catalogue contains the description of the data contained in the database

**QUESTION 20.** Which of the following statements best define an `entity` in the EER model?

**A)** A real physical "thing" that is specific to a particular database management system
**B)** Any physical "thing" that exists only in the considered universe of discourse
**C)** A real physical "thing" or a conceptual "thing" that is specific to a particular database management system or technology
**D)** A "thing" which has either a real physical existence (e.g. a car or a student) or a conceptual existence (e.g. a course)

**QUESTION 21.** Consider the EER model. Which type of entity cannot exist in the database unless another type of entity also exists in the database, but does not require that the identifier of that other entity be included as part of its own identifier?

**A)** Weak entity
**B)** Strong entity
**C)** ID-dependent entity
**D)** ID-independent entity

**QUESTION 22.** Which of the following statements best describe the purpose of an EER model?

**A)** The EER model is meant to replace the relational design
**B)** The EER model is meant to enable low level descriptions of data
**C)** The EER model is meant to be close to a user's perception of the data
**D)** The EER model is meant to enable detailed descriptions of data query processing

**QUESTION 23.** What is the **impedance mismatch**?

**A)** It is a term that identifies a set of conceptual and technical difficulties that are often encountered when a RDBMS manages non-relational data.
**B)** It is a property of RDBMS that allows them to effectively deal with programming languages that do not offer SQL data manipulation capabilities.
**C)** It is a property of RDBMS that allows them to efficiently deal with the differences between the relational model and the in-memory data structures of third-party applications.
**D)** It is a term that identifies a set of conceptual and technical difficulties that are often encountered when a RDBMS is used by programming languages that rely on a non-relational data model.

**QUESTION 24.** Which of the following statement about document databases is **correct**?

**A)** They store data organized as aggregate records accessed by ID values.

**B)** They store data organized as unstructured documents accessed by ID values.
**C)** They store data in a way that is completely opaque (i.e. not visible) to the database, thus allowing optimizations.
**D)** They store data in a way that is visible to the database, but the unstructured nature of the stored document does not allow optimizations.

**QUESTION 25.** Considering the following database schema and SQL query:

```
Movie (Title, Director, Year, Genre)
PlayedIn (Actor, Movie)
```

```
SELECT DISTINCT Actor
FROM PlayedIn AS P1
WHERE NOT EXISTS (SELECT *
                  FROM Movie
                  WHERE Director = 'Woody Allen'
                    AND Title NOT IN (SELECT Movie
                                      FROM PlayedIn AS P2
                                      WHERE P2.Actor = P1.Actor))
```

Which answer best describes the meaning of the query?

**A)** Return the actors that played only in movies directed by Woody Allen.
**B)** Return the actors that played in a movie where Woody Allen also plays a role
**C)** Return the actors that played in all the movies directed by Woody Allen.
**D)** Return the actors that played in at least one movie directed by Woody Allen.

**QUESTION 26.** Considering the following database schema and SQL query:

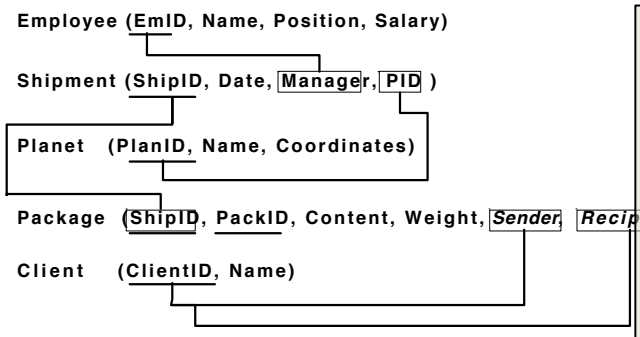


```
SELECT DISTINCT(Employee.Name)
FROM Employee
 JOIN Shipment ON Shipment.Manager = Employee.EmID
 JOIN Package ON Package.ShipID = Shipment.ShipID
WHERE Shipment.ShipID IN (
 SELECT p.ShipID
  FROM Client AS c JOIN Package as p
                    ON c.ClientID = p.Sender
  WHERE c.ClientID = ANY (
    SELECT Client.ClientID
    FROM Client JOIN Package
      ON Client.ClientID = Package.Recipient
    WHERE Package.Weight > 1.5
  )
)
```

Which answer best describes the meaning of the query?

**A)** The query returns the name of the *pilots* that shipped packages of *clients* that previously received **exactly one** package weighting more than 1.5Kg.
**B)** The query returns the name of the *pilots* that shipped packages of *clients* that previously received **at most one** package weighting more than 1.5Kg.
**C)** The query returns the name of the *pilots* that shipped packages of *clients* that previously received **only** packages weighting more than 1.5Kg.
**D)** The query returns the name of the *pilots* that shipped packages of *clients* that previously received **at least** one package weighting more than 1.5Kg.

**QUESTION 27.** Consider the database schema defined in the previous question. Which of the following constraints **cannot** be enforced by a database implementing the schema? Assume no assertions or triggers were defined.

**A)** A `Client` cannot be the recipient of more than one `Package` within the same `Shipment`.

**B)** A `Shipment` must always be directed to a `Planet`.
**C)** A `Shipment` must include at least one `Package`, but it could contain more than one.
**D)** An `Employee` can be the manager of no `Shipments`, but a `Shipments` must be handled by an `Employee`.

---

The following questions are about the transformation into a relational model of the EER diagram depicted in Figure 1, performed according to the standard method described in the book in Chapter 8 and during lectures.

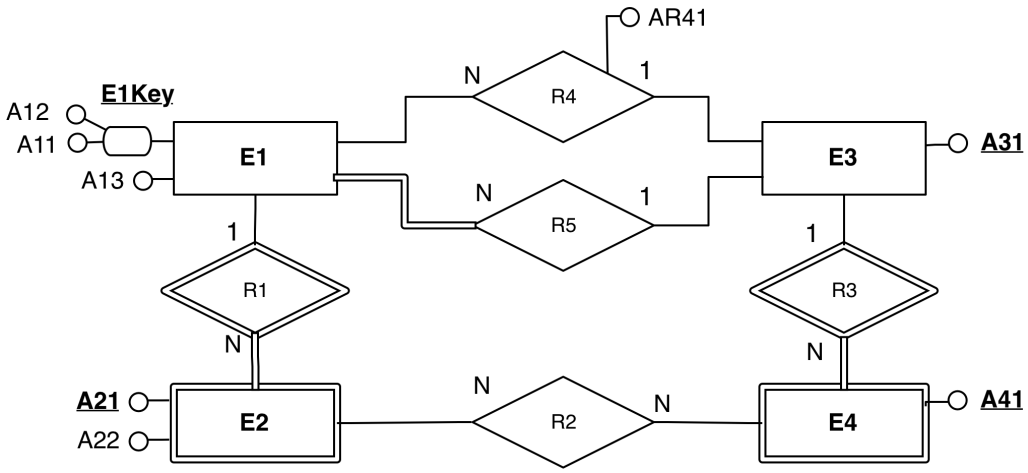


Figure 1 – EER Model #1

**QUESTION 28.** Which of the following statements about the EER Schema in Figure 1 are **correct**?

1. The entity `E4` is transformed in a table with only one attribute.
2. Every instance in the `E1` table will be associated with at least one instance in table `E3`.

**A)** None
**B)** Only [1]
**C)** Only [2]
**D)** Both

**QUESTION 29.** Which of the following statements about the EER Schema in Figure 1 are **correct**?

1. Every instance in table `E3` associated with an instance in table `E1` will also be associated with an instance in table `E4`
2. When creating the relational schema, the relationship `R5` is redundant and therefore will be discarded.

**A)** None
**B)** Only [1]
**C)** Only [2]
**D)** Both

**QUESTION 30.** Which is the minimum number of tables resulting from the transformation of the EER schema in Figure 1 into a logical schema?

**A)** 5
**B)** 6
**C)** 7
**D)** 8

**QUESTION 31.** How many attributes will compose the primary key of table `E2` in the resulting logical schema?

**A)** 2
**B)** 3
**C)** 4
**D)** 5

**QUESTION 32.** How many attributes in table `E1` will be allowed to assume `NULL` values in the resulting logical schema?

**A)** 1
**B)** 2
**C)** 3
**D)** 4

---

The next 3 questions are related to the EER diagram depicted in Figure 2.

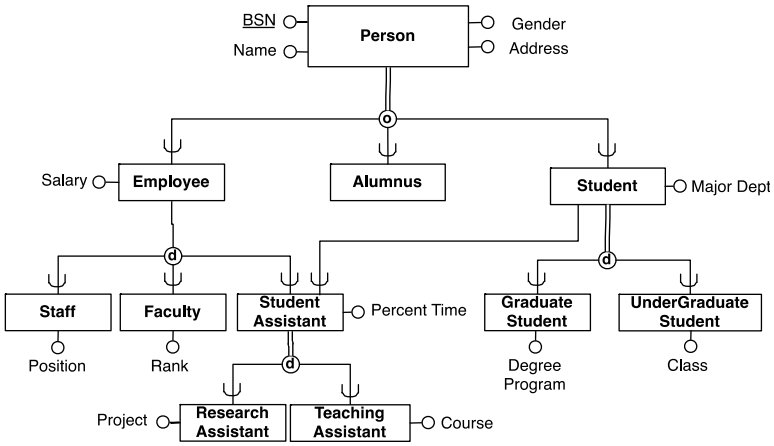


**Figure 2 – EER Model #2**

**QUESTION 33.** Consider the diagram in Figure 2**.** Which of the following statements are **correct?**

1. The diagram in Figure 2 is a specialization lattice because the `Student Assistant` entity is a subset of two entities participating in a total specialization relationship.
2. The diagram in Figure 2 is a specialization lattice because every entity participates as a sub-entity in only one specialization relationship.

**A)** None
**B)** Only [1]
**C)** Only [2]
**D)** Both

**QUESTION 34.** Consider the diagram in Figure 2. Which of the following statements are **correct**?

1. A member of the `Research Assistant` entity can also be a member of the `Alumnus` entity
2. A member of the `Faculty` entity can also be a member of the `Graduate Student` entity

**A)** None

**B)** Only [1]
**C)** Only [2]
**D)** Both

**QUESTION 35.** Consider the diagram in Figure 2. Which of the following statements are **correct**?

1. Due to multiple inheritance, an instance of the `Teaching Assistant` entity includes the `Salary` attribute and the `Major Dept` attribute
2. Due to multiple inheritance, a member of the `Teaching Assistant` entity includes two `Name` attributes.

**A)** None
**B)** Only [1]
**C)** Only [2]
**D)** Both

---

The next 3 questions are related to the EER diagram in drawn according to the notation used in the course's slides. The diagram represents a database about programming languages.

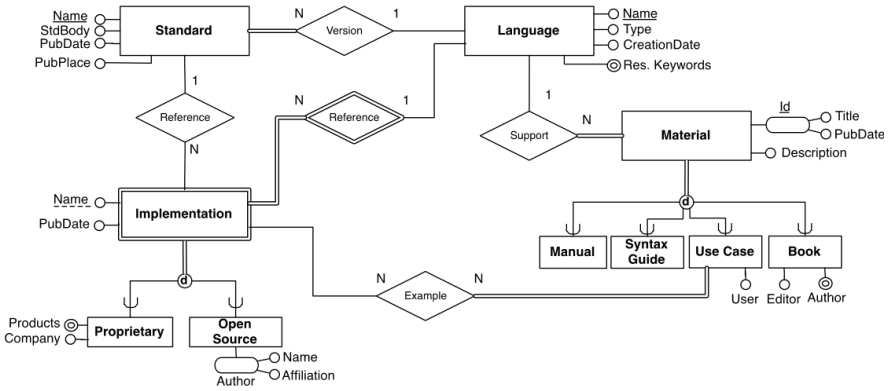


**Figure 3 – EER Model #2**

**QUESTION 36**. Consider the EER diagram in Figure 3. Which of the following constraints is not expressed?

**A)** An `Implementation` of a given programming language must be either `Open Source` or `Proprietary`.
**B)** A programing `Language`'s syntax might be described in a `Syntax Guide`.
**C)** A support `Material` must be related to at most one programming `Language`.
**D)** A `Standard` must be described in at least one `Book`.

**QUESTION 37**. Consider the EER diagram in Figure 3. How many *internal* identifiers are defined?

**A)** 2
**B)** 3
**C)** 4
**D)** 5

**QUESTION 38**. Consider the restructuring of the EER diagram in Figure 3, performed as a preliminary operation functional to the translation of the EER model into a relational model. The

restructuring is performed according to the standard method presented during lectures, but you are asked to use a *parent-collapsing* specialization removal method. Which of the following statements **correctly describe** the resulting EER diagram?

1. The `Material` entity participates in the `Example` relationship with a partial participation constraint.
2. The `Implementation` entity includes 5 attributes, but only the `Author` attribute becomes optional because it is multi-valued
3. The `Implementation` entity features 6 attributes, but only 3 become optional
4. The `Example` relationship changes its cardinality from N:N to 1:N because not all the instances from the Material entities will be related to an instance of the `Implementation` entity

**A)** Statements [1] and [2]
**B)** Statements [1] and [3]
**C)** Statements [2] and [4]
**D)** Statements [2] and [3]