Relational Database Design

CS 327E
September 27, 2017

Announcements:

- Athena credits
- Reminder: Lab 2 next week

1) Given the create table statement, which values in the sample data do not satisfy the table definition?

```
create table Pokemon_Characters(
    id int primary key,
    name varchar(50) not null,
    type int,
    height_ft double,
    weight_lbs double,
    health_pts int,
    foreign key (type) references Pokemon_Types(id) on delete cascade
);
```

id	name	type	height_ft	weight_lbs	health_pts
16114	Ponyta	Fire	3.03	66.1	60
5620	Tyranitar	802	6.07	445.3	120
2298	Vaporeon	Water	3.03	63.9	90

A) 'Fire'
B) 802
C) 'Tyranitar'
D) 'Fire' and 'Water'

2) In order to create the table Pokemon_Characters, the referenced table Pokemon_Types must already exist in the database.

```
create table Pokemon_Characters(
   id int primary key,
   name varchar(50) not null,
   type int,
   height_ft double,
   weight_lbs double,
   health_pts int,
   foreign key (type) references Pokemon_Types(id) on delete cascade
);
```

A) True
B) False

3) After creating the table Pokemon_Characters, we decide to add another column to store a pokemon's region. What value will we get if we retrieve the region immediately after modifying the table?

```
create table Pokemon_Characters(
   id int primary key,
   name varchar(50) not null,
   type int,
   height_ft double,
   weight_lbs double,
   health_pts int,
   foreign key (type) references Pokemon_Types(id)
);

alter table Pokemon_Characters add region varchar(30);
```

A) NULL
B) 0
C) ''
D) None

4) What would happen to the data in Pokemon_Characters if we delete the record Pokemon_Types.id = 802?

```
create table Pokemon_Characters(
    id int primary key,
    name varchar(50) not null,
    type int,
    height_ft double,
    weight_lbs double,
    health_pts int,
    foreign key (type) references Pokemon_Types(id) on delete cascade
);
```

id	name	type	height_ft	weight_lbs	health_pts
16114	Ponyta	Fire	3.03	66.1	60
5620	Tyranitar	802	6.07	445.3	120
2298	Vaporeon	Water	3.03	63.9	90

A) No changes to the data in Pokemon_Characters
B) All records with type = 802 would get deleted from Pokemon_Characters
C) The value 802 in Pokemon_Characters.type would get set to NULL
D) The table Pokemon_Characters would get dropped

5) What type of constraint would we use to require unique values in Pokemon_Characters.name?

```
create table Pokemon_Characters(
    id int primary key,
    name varchar(50) not null,
    type int,
    height_ft double,
    weight_lbs double,
    health_pts int,
    foreign key (type) references Pokemon_Types(id) on delete cascade
);
```

id	name	type	height_ft	weight_lbs	health_pts
16114	Ponyta	Fire	3.03	66.1	60
5620	Tyranitar	802	6.07	445.3	120
2298	Vaporeon	Water	3.03	63.9	90

A) DISTINCT
B) PRIMARY KEY
C) UNIQUE
D) CHECK

Demo: Postgres RDS

(see snippets repo for instructions)

Denormalization

Instacart ERD in 3NF

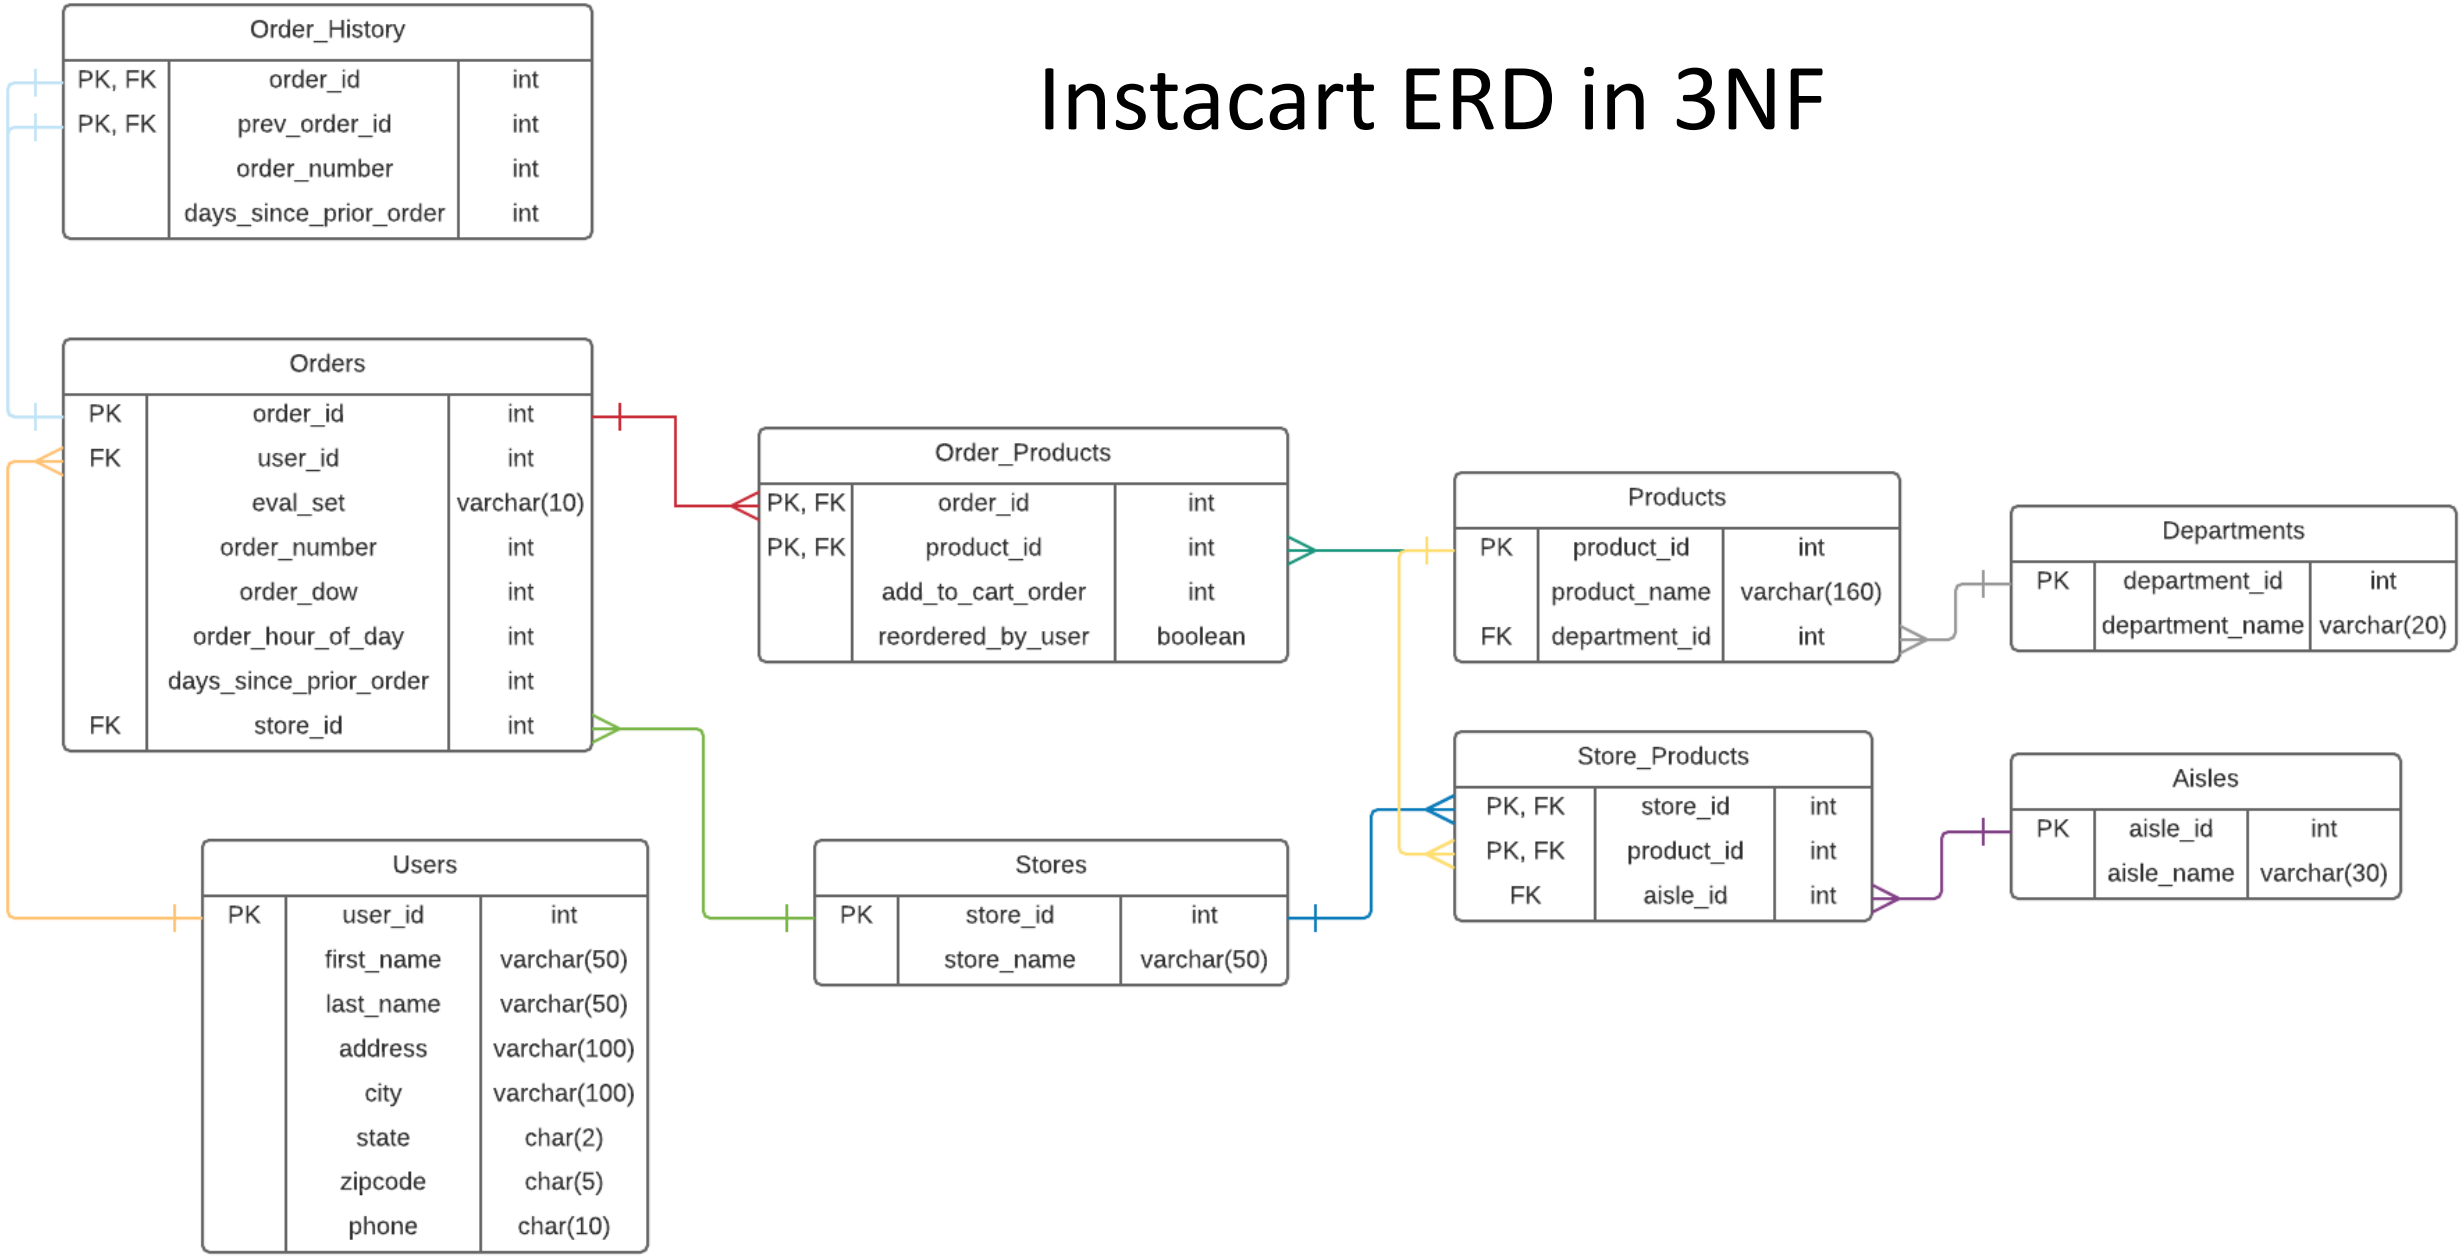

Instacart Dimensional Schema

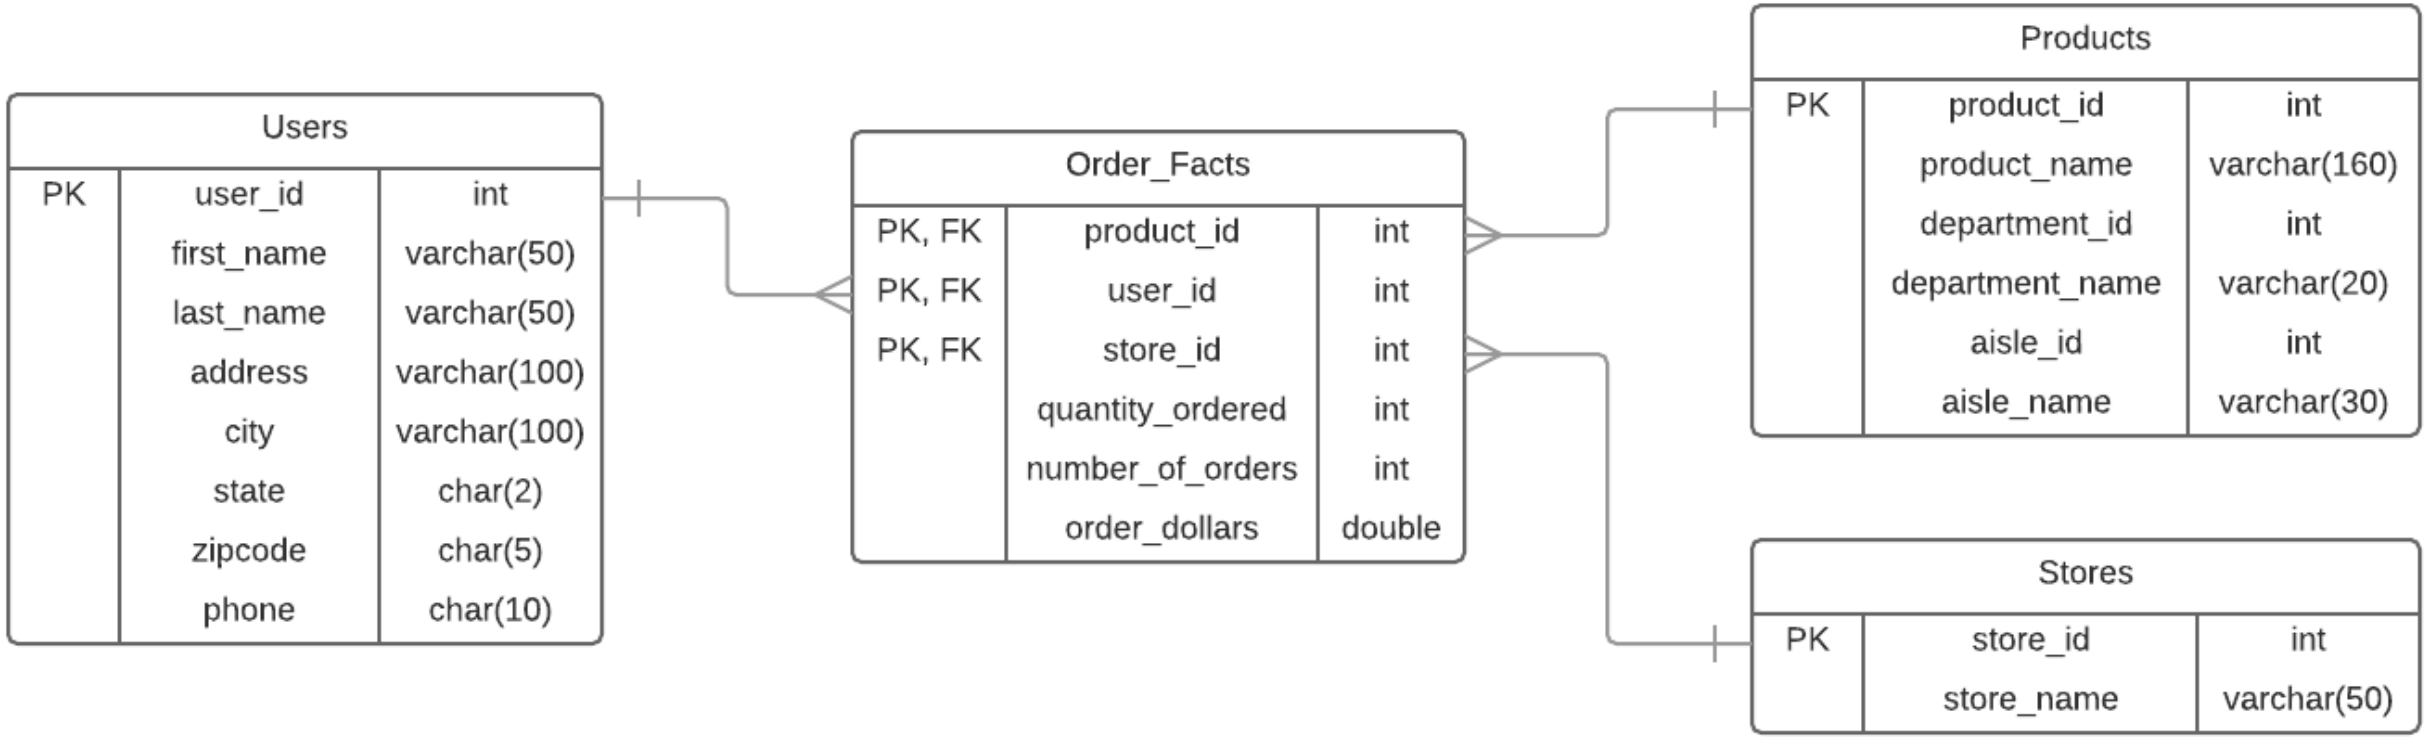

Practice Problem 1: The Instacart business leaders would like to see how these order facts change over time. In particular, they want to analyze **quantity_ordered** by **Day of Week**, **User**, **Product**, and **Store;** they also want to analyze **quantity_ordered** by **Hour of Day**, **User**, **Product**, and **Store**.

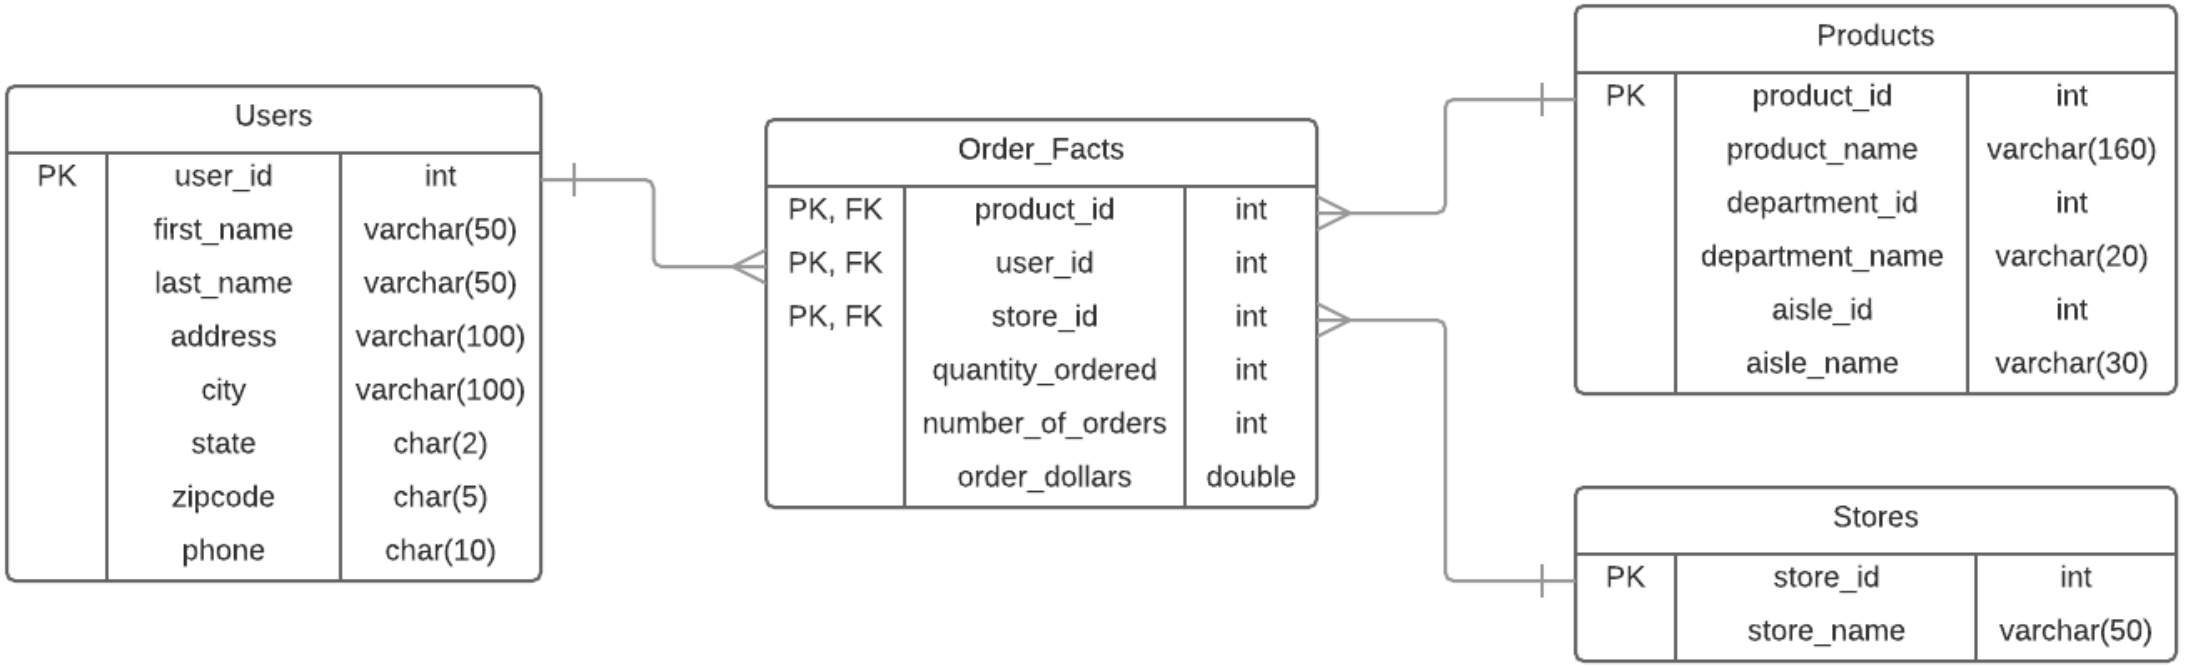

Practice Problem 1: The Instacart business leaders would like to see how these order facts change over time. In particular, they want to analyze **quantity_ordered** by **Day of Week**, **User**, **Product**, and **Store;** they also want to analyze **quantity_ordered** by **Hour of Day**, **User**, **Product**, and **Store**.

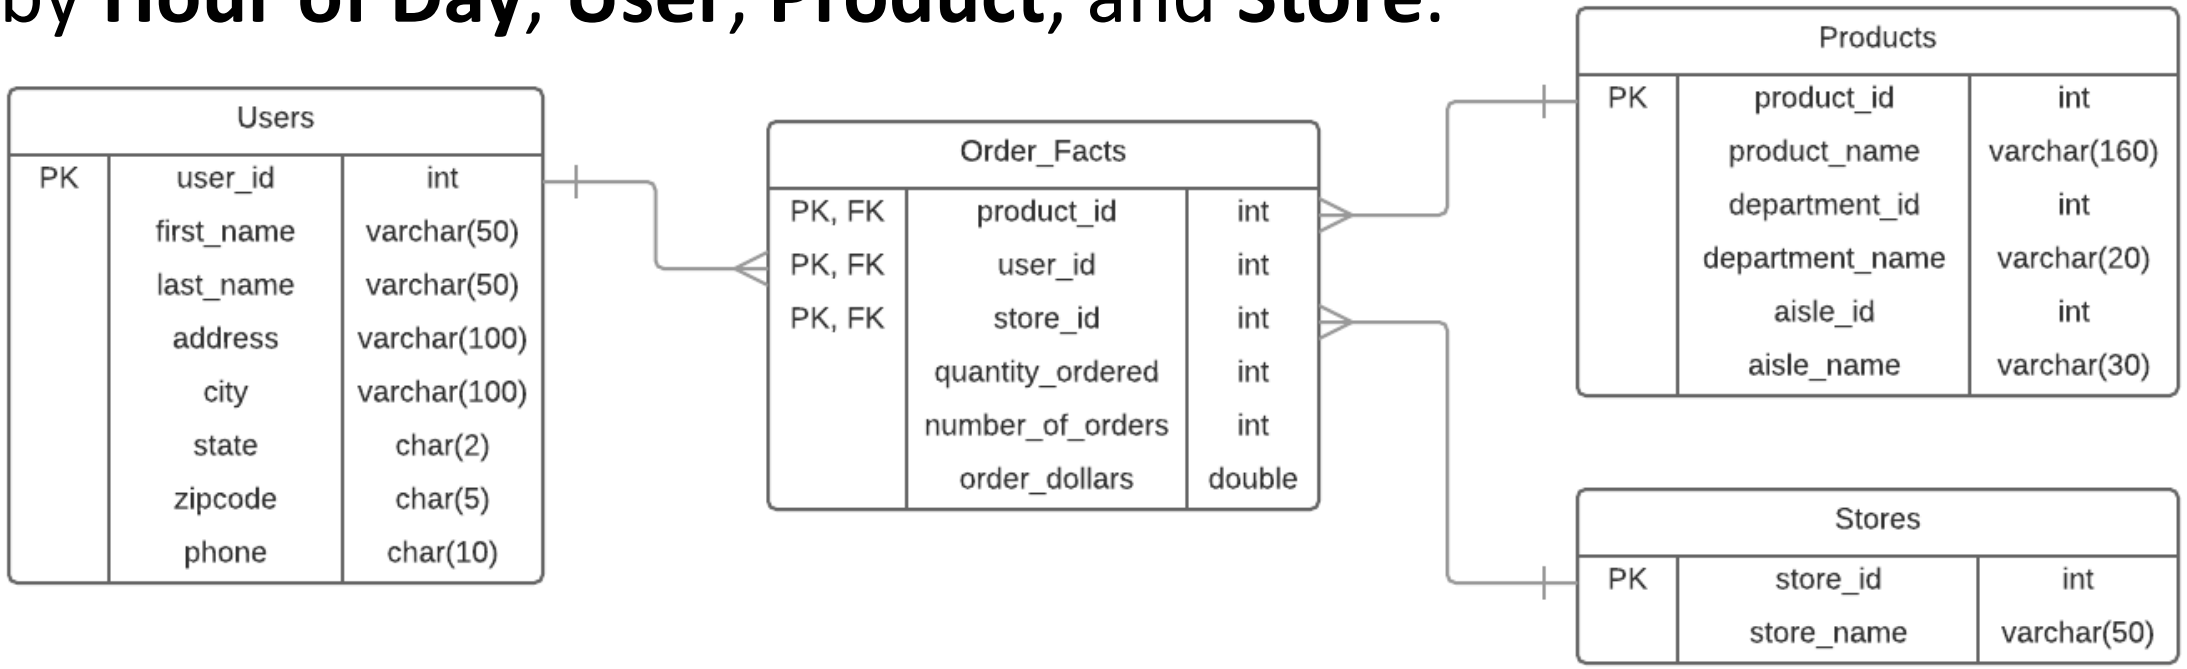

How many new Order_Facts tables are needed to support the analyses?

A) 0 B) 1 C) 2 D) > 2

Demo: DDL and copy command

(see snippets repo for code)