

## Perspective

Subqueries are a powerful tool that you can use to solve difficult problems. Before you use a subquery, however, remember that a subquery can often be restated more clearly by using a join. If so, you'll typically want to use a join instead of a subquery.

If you find yourself coding the same subqueries in multiple places, you should consider creating a view for that subquery as described in chapter 12. This will help you develop queries more quickly since you can use the view instead of coding the subquery again. In addition, since views typically execute more quickly than subqueries, this may improve the performance of your queries.

## Terms

subquery	comment
introduce a subquery	pseudocode
nested subquery	common table expression (CTE)
correlated subquery	recursive query
uncorrelated subquery	recursive CTE
inline view	

## Exercises

- Write a SELECT statement that returns the same result set as this SELECT statement, but don't use a join. Instead, use a subquery in a WHERE clause that uses the IN keyword.

```
SELECT DISTINCT vendor_name
FROM vendors JOIN invoices
    ON vendors.vendor_id = invoices.vendor_id
ORDER BY vendor_name
```

- Write a SELECT statement that answers this question: Which invoices have a payment total that's greater than the average payment total for all invoices with a payment total greater than 0?

Return the invoice\_number and invoice\_total columns for each invoice. This should return 20 rows.

Sort the results by the invoice\_total column in descending order.

- Write a SELECT statement that returns two columns from the General\_Ledger\_Accounts table: account\_number and account\_description. Return one row for each account number that has never been assigned to any line item in the Invoice\_Line\_Items table. To do that, use a subquery introduced with the NOT EXISTS operator. This should return 54 rows. Sort the results by the account\_number column.

4. Write a SELECT statement that returns four columns: vendor\_name, invoice\_id, invoice\_sequence, and line\_item\_amount.  
Return a row for each line item of each invoice that has more than one line item in the Invoice\_Line\_Items table. *Hint: Use a subquery that tests for invoice\_sequence > 1.* This should return 6 rows.  
Sort the results by the vendor\_name, invoice\_id, and invoice\_sequence columns.
5. Write a SELECT statement that returns two columns: vendor\_id and the largest unpaid invoice for each vendor. To do this, you can group the result set by the vendor\_id column. This should return 7 rows.  
Write a second SELECT statement that uses the first SELECT statement in its FROM clause. The main query should return a single value that represents the sum of the largest unpaid invoices for each vendor.
6. Write a SELECT statement that returns the name, city, and state of each vendor that's located in a unique city and state. In other words, don't include vendors that have a city and state in common with another vendor. This should return 38 rows.  
Sort the results by the vendor\_state and vendor\_city columns.
7. Use a correlated subquery to return one row per vendor, representing the vendor's oldest invoice (the one with the earliest date). Each row should include these four columns: vendor\_name, invoice\_number, invoice\_date, and invoice\_total. This should return 34 rows.  
Sort the results by the vendor\_name column.
8. Rewrite exercise 7 so it gets the same result but uses an inline view instead of a correlated subquery.
9. Rewrite exercise 5 so it uses a common table expression (CTE) instead of an inline view.

