

Examcollection

<http://www.ipass4sure.com/examcollection.htm>

70-451

Microsoft

*PRO: Designing Database Solutions and Data Access
Using Microsoft SQL Server 2008*

<http://www.ipass4sure.com/exams.asp?examcode=70-451>

The 70-451 practice exam is written and formatted by Certified Senior IT Professionals working in today's prospering companies and data centers all over the world! The 70-451 Practice Test covers all the exam topics and objectives and will prepare you for success quickly and efficiently. The 70-451 exam is very challenging, but with our 70-451 questions and answers practice exam, you can feel confident in obtaining your success on the 70-451 exam on your FIRST TRY!

Microsoft 70-451 Exam Features

- Detailed questions and answers for 70-451 exam
- Try a demo before buying any Microsoft exam
- 70-451 questions and answers, updated regularly
- Verified 70-451 answers by Experts and bear almost 100% accuracy
- 70-451 tested and verified before publishing
- 70-451 examcollection vce questions with exhibits
- 70-451 same questions as real exam with multiple choice options

Acquiring Microsoft certifications are becoming a huge task in the field of I.T. More over these exams like 70-451 exam are now continuously updating and accepting this challenge is itself a task. This 70-451 test is an important part of Microsoft certifications. We have the resources to prepare you for this. The 70-451 exam is essential and core part of Microsoft certifications and once you clear the exam you will be able to solve the real life problems yourself. Want to take advantage of the Real 70-451 Test and save time and money while developing your skills to pass your Microsoft 70-451 Exam? Let us help you climb that ladder of success and pass your 70-451 now!

DEMO EXAM

For Full Version visit

<http://www.ipass4sure.com/allexams.asp>

Exam Name:	Pro: Designing Database Solutions And Data Access Using Microsoft SQL Server 2008		
Exam Type:	Microsoft		
Exam Code:	70-451	Total Question	146

Question: 1

You are the designer of a sql database on an instance of sql server 2008. your database contains a course table and an instructor table defined as follows:

*Step Up In You Career,
Get Certified Today!*

You want to create a foreign key constraint to associate the two tables and prevent a user from deleting an instructor who is currently assigned to teach a course. Which clause should you include in your foreign key constraint?

- A. On delete no action
- B. On delete cascade
- C. On delete set null
- D. On delete set default

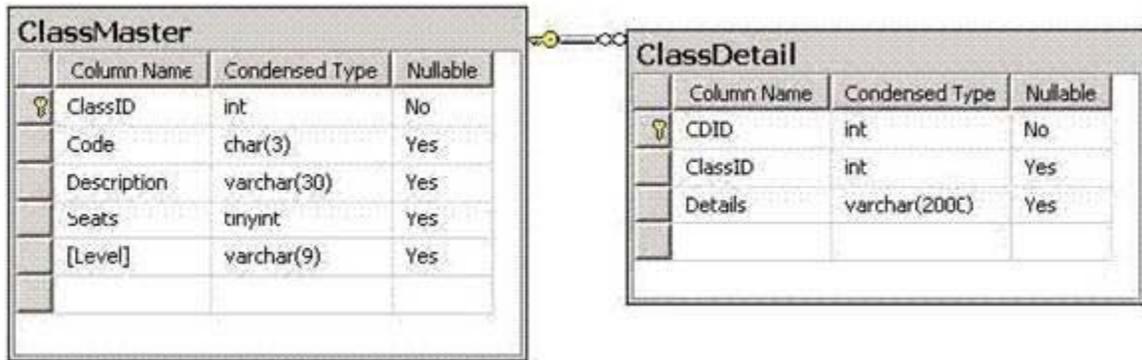
Answer: A

For example, in this scenario, you might use the following statement to create the necessary foreign key constraint: `alter table course add constraint fk_instructor foreign key (unstructured) references instructor (instructorid) on delete no action;` this statement includes the add constraint clause of the alter table statement to add a foreign key constraint to the course table. the references clause specifies the referenced column in the parent table. when you include the on delete no action clause, it ensures that referential integrity is maintained. if a referenced row in the parent table is deleted, then the delete fails if child rows exist. with the given statement, a user could not delete an instructor who was assigned to courses in the course table. you should note that no action is the default for on delete and on update. you should not include the on delete cascade clause in your foreign key constraint. if you did so and a user deleted an instructor that was assigned to teach a course, the delete operation would be performed, and all courses for the instructor would be deleted as well. you should not include the on delete set null clause in your foreign key constraint. if you did so and a user deleted an instructor who was assigned to teach a course, the instructor column would be set to null for all of the instructor's courses in the course table. you should note that if the foreign key consists of multiple columns, all of the columns in the foreign key would be set to null. you should not include the on delete set default clause in your foreign key constraint. if you did so and a user deleted an instructor who was assigned to teach a course, the instructor column would be set to the default value for all of that instructor's courses in the course table. you should note that if the foreign key consists of multiple columns, all of the columns in the foreign key are set to the default value. in addition, you should ensure that each foreign key column has a default definition; otherwise, the column will be assigned a null value.

Question: 2

You manage a database on an instance of sql server 2008 at a large educational institution. You have classmate and class detail tables as shown: (click exhibit to view table design)

Exam Name:	Pro: Designing Database Solutions And Data Access Using Microsoft SQL Server 2008		
Exam Type:	Microsoft		
Exam Code:	70-451	Total Question	146



the user1 user starts a session and executes the following `transact-sql`: `begin transaction update class detail with(tab lock)set classdetail.details = classdetail.details + '. this is a freshman-level class.'` From class master inner join class detail on `classmaster.classid = classdetail.classid where classmaster.level = 'freshman'`; then, user2 starts a session and executes the following `transact-sql`: `begin transaction insert into class master(description, seats, level)values ('calculus i', 25, 'sophomore'),('accounting iii', 80, 'senior'),('world history', 30, 'freshman');``delete from dbo.classdetail where cdid = 2;` `commit transaction` what is the result?

- A. User1's updates are applied first, and then user2's dml operations are performed.
- B. User2's session hangs waiting for user1s update to complete.
- C. User1s updates are not performed, but user2's dml operations complete successfully.
- D. User2's session immediately returns an error.

Answer: B

In this scenario, user1 begins a transaction and executes an update statement, but does not commit the changes. then, user2 begins a transaction, but user1 currently has locks on the table that prevents user1s actions from being performed. user2's session hangs waiting for user1's session to release the locks. Sql server uses locking to control which transactions have access to which resources at any given time. This prevents one transaction from overwriting another transaction's changes. sql server determines the locks that will be acquired based on the type of action being performed. however, you can override default locking behavior if necessary using table hints. the option that states user1's updates are applied first and then user2's dml operations are performed is incorrect. user1 does not commit the changes. therefore, the changes are pending and user2's operations are not performed. the option that states user1's updates are not performed but user2's dml operations complete successfully is incorrect. in this scenario, user1's updates have not been committed and will remain pending, blocking user2's operations from executing. the option that states user2's session immediately returns an error is incorrect. user2's session will wait to obtain the necessary locks to perform the required operations.

Question: 3

You are a database developer on an instance of sql server 2008. you are creating a tracking application. You want to create a column in one of your tables that stores gps latitude and longitude coordinates. which data type should you use for the column?

- A. Xml
- B. Table
- C. Geometry
- D. Geography

Exam Name:	Pro: Designing Database Solutions And Data Access Using Microsoft SQL Server 2008		
Exam Type:	Microsoft		
Exam Code:	70-451	Total Question	146

Answer: D

The geography data type represents data in a round-earth coordinate system. the geography data type is used to represent data as gps latitude and longitude coordinates. you should not use the xml data type because this data type is used to store xml data, such as xml documents or xml document fragments. you should not use the table data type because this data type is used to store rows of data from a result set. the table data type can be processed later as needed, much like a temporary table. you can also pass table-valued parameters into stored procedures and functions for processing. you should not use the geometry data type. the geometry data type represents data in a Euclidean, or flat, coordinate system.

Question: 4

You are a database developer on an instance of sql server 2008. you have a work order table and a workorderdetail table. the work order table contains one row for each work order, and the Workorderdetail table contains line items of work performed for each work order. you want to create a foreign key constraint to relate the two tables. you want to ensure that if a work order identifier in the work order table is updated, that the corresponding rows in the workorderdetail table are also updated to maintain referential integrity. what should you do?

- A. Use a trigger to enforce referential integrity.
- B. Create a check constraint on the work order table.
- C. Include the on update cascade clause in your foreign key constraint.
- D. Include the with check clause in your foreign key constraint.

Answer: C

When you include the on update cascade clause, it ensures that if a referenced row in the parent table is updated, then the corresponding child rows are updated and referential integrity is maintained. this would ensure that if a user updated the identifier of a work order, the corresponding rows in the workorderdetail table would also be updated. you should not use a trigger to enforce referential integrity. triggers should not be used when constraints can accomplish the same task, foreign key constraints are used to enforce referential integrity. you should not create a check constraint on the work order table. a check constraint is used to restrict the data allowed for a column to specific values. a check constraint consists of a Boolean expression that evaluates to either true or false. if the expression evaluates to true, the value is allowed for the column, and if the expression evaluates to false, the value is not allowed for the column, check constraints can be defined at the table level or column level, but only check constraints defined at the table level can use columns other than the constrained column in the constraint expression. you should not include the with check clause in your foreign key constraint because this only enables the constraint. by default, foreign key and check constraints are enabled when they are created. you can use the no check constraint clause of the alter table statement to temporarily disable a foreign key or check constraint if you need to insert rows that violate the constraint.

Question: 5

You are a database developer on an instance of sql server. you have defined an inline table valued function using the following statement: create function udf_getorderdetails (@or did int)returns table as return (select sh.orderdate, sd.*from salesorderheader sh inner join salesorderdetail sd on sh.salesorderid = sd.salesorderid where sd.salesorderid=@or did); you have table named salesorderhistory that contains a salesorderid column. you want to query the salesorderhistory table and use the udf_getorderdetails function to return order details. all sales orders in the salesorderhistory table and the associated order details should be displayed, even if no details exist for a given order. which query should you use?

Exam Name:	Pro: Designing Database Solutions And Data Access Using Microsoft SQL Server 2008		
Exam Type:	Microsoft		
Exam Code:	70-451	Total Question	146

- A. Select * from salesorderhistory h cross apply udf_getorderdetails (h.salesorderid) as d
- B. Select h.*, d.*from salesorderhistory h outer apply udf_getorderdetails (h.salesorderid) as d order by h.salesorderid; select * from udf_getorderdetails(t.salesorderid) as t;
- C. Select h.*, d.* inner join udf_getorderdetails(h.salesorderid) as d on h.salesorderid = d.salesorderid order by h.salesorderid;

Answer: B

You should use the following query:

```
select h.*, d.*from salesorderhistory hotter apply udf_getorderdetails (h.salesorderid) as d order by h.salesorderid; in this scenario, you wanted to join the results of a table-valued function to another table. you can accomplish this using the apply operator in the from clause of your query. the syntax is as follows: select column list from outer table {cross | outer} apply inner_table the inner_table may specify another table or view, but may also specify a function that returns a table type. in this scenario, you specified the udf_getorderdetails function as the inner table. for each row in the query, the udf_getorderdetails function is called using the salesorderid from the outer table. if the function returns rows, these rows are joined with the outer table using union all and are available to the query. you should use outer apply to ensure that all rows in the outer table, salesorderhistory, are returned even if no order details exist. for rows in the sales order history table that are not returned by the function, all order detail columns will be assigned a null value. when using the apply operator, the outer table may also use a table-valued function, but it cannot pass the function columns to the inner_table as arguments. you should not use the query that includes an inner join because this syntax is invalid. this statement will return the following error: msg 4104, level 16, state 1, line 12the multi-part identifier "h.salesorderid" could not be bound. you should not use the query that uses a sub query because it attempts to reference t.salesorderid in the inner query. this statement will generate the following error: msg 4104, level 16, state 1, line 25the multi-part identifier "t.salesorderid" could not be bound. you should not use the query that uses cross apply. a cross apply operator only returns rows that had rows returned by the udf_getorderdetails function. in this scenario, you wanted to return all rows from the sales order history table, even if they contained no order details. a cross apply would not return rows in the salesorderhistory table that did not contain order details.
```

Question: 6

You manage a database on an instance of sql server 2008 for a large sales organization. You have a contact table defined with the following statement: create table contact(contacted int primary key,contacttype nvarchar(10),first name nvarchar(30) not null,lastname nvarchar(50) not null, territory nvarchar (20),region nvarchar(10),repid int,initcontact date time default get date()); you have a stored procedure that accepts a territory and region as input parameters and queries the contact table to return contact information as a table data type. users in the marketing department frequently use the stored procedure to obtain information about contacts for a specified territory and region. each region contains multiple territories. there are generally a small number of contacts within each territory, but each region contains many contacts. the stored procedure is performing poorly, and you want to optimize its performance. which action should you take?

- A. Create a view on the contact table and modify the stored procedure to use the view.
- B. Create a no clustered composite index on the territory and region columns.
- C. Create two no clustered indexes, one on the territory column and one on the region column.
- D. Re-create the stored procedure and include the with schema binding clause.

Answer: B

To optimize the query in the stored procedure, you should create a composite index on the two columns. In this scenario, each region contains many contacts, and each territory contains only a

Exam Name:	Pro: Designing Database Solutions And Data Access Using Microsoft SQL Server 2008		
Exam Type:	Microsoft		
Exam Code:	70-451	Total Question	146

few. because there are fewer region values than territory values, you should specify the territory column first when creating the index. with the composite index, sql server will find the rows that match the territory and region criteria in the where clause with a single index search. the search will be performed only by reading the index pages, and the data pages will be accessed only once to retrieve the result set. you might create the index using the following statement:
create index ix_contact_territoryregion on contact (territory, region);using this index would generally offer better performance than if you had created the index with the region column first, because the territory column contains more unique values. you should not create a view on the contact table and modify the stored procedure to use the view. using a view would not improve performance. the stored procedure would still have to query the view, and the view would access the underlying table. you should not create two no clustered indexes, one on the territory column and one on the region column. in this scenario, a single composite no clustered index would provide better performance. although sql server can use more than one index on a table, it is unlikely that the index on the region column would be used in this scenario because of the low uniqueness of the column's values. you should not re-create the stored procedure including the with schema binding clause because the with schema binding clause is not allowed for a stored procedure. in addition, it has no affect on performance. attempting to create a stored procedure including with schema binding will generate an error similar to the following: msg 487, level 16, state 1, procedure p, line 3an invalid option was specified for the statement "create/alter procedure". schema binding is used when creating a view or function. you can include the with schema binding clause in a create view statement to ensure that no base tables on which the view is based are dropped or modified in a way that might affect the view's definition. to drop base tables or make such modifications, a user would first need to drop the view, alter the view omitting schema binding, or alter the view to remove any unwanted dependencies. you can also include the uith schehabinding clause when creating a function to ensure base tables are not changed in such a way that would render the function unusable.

Question: 7

You are a database developer on an instance of sql server 2008. you need to create a table to store information for construction projects undertaken by your company. you want to create a column to store a text narrative for each project. you want to minimize storage space, but allow the narrative column to store up to 10000 bytes of data. which column definition should you use for the narrative column?

- A. Narrative text
- B. Narrative varbinary(max)
- C. Narrative varchar(max)
- D. Narrative varchar(10000)

Answer: C

This will create the narrative column as a large-value type that can store up to 2 gb of data. you should not define the narrative column as a text data type because using the text data type should be avoided. a text data type can support 10000 bytes of data, but cannot be indexed. you should use the large-value type varchar(max) instead. you should not define the narrative column as a varbinary (max) data type. a varbinary data type is used to store binary data, not text data. you should not define the narrative column using narrative varchar (10000) because this would generate an error. even though using large-value types allows you to store more than 8060 bytes within a single row, each individual column can be no larger than 8000 bytes.

Question: 8

You are a developer on an instance of sql server 2008. you are designing an event table. the event date column must contain a value. in addition, for each event, the event date must be a



Pass4sure Certification Exam Features;

- Pass4sure offers over **2500** Certification exams for professionals.
- More than **98,800** Satisfied Customers Worldwide.
- Average **99.8%** Success Rate.
- Over **120** Global Certification Vendors Covered.
- Services of **Professional & Certified Experts** available via support.
- Free **90 days** updates to match real exam scenarios.
- **Instant Download Access!** No Setup required.
- Price as low as **\$19**, which is 80% more **cost effective** than others.
- **Verified answers** researched by industry experts.
- Study Material **updated** on regular basis.
- Questions / Answers are downloadable in **PDF** format.
- Mobile Device Supported (**Android, iPhone, iPod, iPad**)
- **No authorization** code required to open exam.
- **Portable** anywhere.
- **Guaranteed Success.**
- **Fast**, helpful support **24x7**.



View list of All certification exams offered;
<http://www.ipass4sure.com/all exams.asp>

View list of All Study Guides (SG);
<http://www.ipass4sure.com/study-guides.asp>

View list of All Audio Exams (AE);
<http://www.ipass4sure.com/audio-exams.asp>

Download Any Certification Exam DEMO.
<http://www.ipass4sure.com/samples.asp>

To purchase Full version of exam click below;
<http://www.ipass4sure.com/all exams.asp>

3COM	CompTIA	Filemaker	IBM	LPI	OMG	Sun
ADOBE	ComputerAssociates	Fortinet	IISFA	McAfee	Oracle	Sybase
APC	CWNP	Foundry	Intel	McData	PMI	Symantec
Apple	DELL	Fujitsu	ISACA	Microsoft	Polycom	TeraData
BEA	ECCouncil	GuidanceSoftware	ISC2	Mile2	RedHat	TIA
BICSI	EMC	HDI	ISEB	NetworkAppliance	Sair	Tibco
CheckPoint	Enterasys	Hitachi	ISM	Network-General	SASInstitute	TruSecure
Cisco	ExamExpress	HP	Juniper	Nokia	SCP	Veritas
Citrix	Exin	Huawei	Legato	Nortel	See-Beyond	Vmware
CIW	ExtremeNetworks	Hyperion	Lotus	Novell	Google	

and many others.. See complete list [Here](#)

