

Microsoft®

TechEd

2012



Microsoft®



Back to Indexes, the original culprit !

Amit Bansal
CTO, eDominator Systems | Peopleware India
President & Founder, www.SQLServerGeeks.com
www.amitbansal.net

Takeaways...

- ④ How does the optimizer decide which index to use?
- ④ Optimizer's choice of the physical join operator
- ④ Database Tuning Advisor. When? Why? How?
- ④ Getting rid of redundant indexes !



© Roberto Neumiller/SOS SAHEL


Image Source: <http://www.jeffcubos.com/2011/10/18/load-em-up/>

Agenda

- ④ Demo: Indexing Scenario 1
- ④ Optimizer's choice of Physical Join Operator
- ④ Indexing Scenario 2
- ④ Demo: Good Design vs Better Design
- ④ The Optimizer's choice of Indexes
- ④ Demo: Indexing Scenario 3
- ④ How, When & Why did the Indexes become culprits?
- ④ Demo: Clean up act !
- ④ Summary / Call to Action



demo



Indexing Scenario 1

Optimizer's choice of Physical Join Operator

④ NESTED LOOP JOIN

- ④ when the outer input is small and the inner input has an index on the join key
- ④ No restriction on operators; Inputs need not be sorted

④ MERGE JOIN:

- ④ Medium to large inputs
- ④ Requires equality operator & inputs must be sorted on join predicate

④ HASH JOIN:

- ④ Large inputs
- ④ Requires equality operator; Inputs need not be sorted

Agenda

- ④ Demo: Indexing Scenario 1
- ④ Optimizer's choice of Physical Join Operator
- ④ Indexing Scenario 2
- ④ Demo: Good Design vs Better Design
- ④ The Optimizer's choice of Indexes
- ④ Demo: Indexing Scenario 3
- ④ How, When & Why did the Indexes become culprits?
- ④ Demo: Clean up act !
- ④ Summary / Call to Action

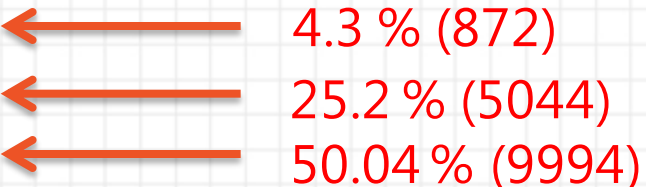
Indexing Scenario 2

```
SELECT C.ContactID, C.FirstName, C.EmailPromotion
FROM Person.Contact2 AS C
WHERE C.FirstName LIKE N'L%'           ← 4.3 % (872)
      AND C.EmailPromotion = 1         ← 25.2 % (5044)
      AND C.ContactID < 10000          ← 50.04 % (9994)
OPTION (MAXDOP 1)
```

--returns 77 out of 19972 records

Indexing Scenario 2

```
SELECT C.ContactID, C.FirstName, C.EmailPromotion
FROM Person.Contact2 AS C
WHERE C.FirstName LIKE 'L%'
      AND C.EmailPromotion = 1
      AND C.ContactID < 10000
OPTION (MAXDOP 1)
```




Condition	Percentage	Count
C.FirstName LIKE 'L%'	4.3 %	872
C.EmailPromotion = 1	25.2 %	5044
C.ContactID < 10000	50.04 %	9994

```
-- Option 1
CREATE INDEX ContactComposite4
ON Person.Contact2(FirstName, EmailPromotion)
-- Option 2
CREATE INDEX ContactComposite5
ON Person.Contact2(EmailPromotion, FirstName)
```



demo



Good Design vs Better Design

The Optimizer's choice of Indexes

- ④ Common advice (as it is):
 - ④ Most selective column should be the first column
 - ④ Does it mean you add the primary key as the leading column of all your indexes?
- ④ SQL maintains HISTOGRAM only for the first column of the index
- ④ Additionally, DENSITY is maintained for all the columns in left-based subset mechanism
- ④ Common advice (as it should be):
 - ④ Most selective column should be the first column when all other column predicates use the equality operator

The Optimizer's choice of Indexes

④ Single column index:

④ True:

- ④ FirstName LIKE 'L%'
- ④ EmailPromotion = 1
- ④ ContactID < 10000

④ False:

- ④ FirstName LIKE '%L'
- ④ ABS(EmailPromotion) = 1
- ④ ContactID + 1 < 10000

The Optimizer's choice of Indexes

⌵ Multi-column index:

⌵ The index can be used to seek on the second column if there is an equality predicate on the first column

⌵ True:

⌵ `FirstName = 'L' AND EmailPromotion = 1`

⌵ Partially True:

⌵ `FirstName LIKE 'L%' AND EmailPromotion = 1`

⌵ False

⌵ `FirstName LIKE '%L' AND EmailPromotion = 1`

The Optimizer's choice of Indexes

⌵ Multi-column index:

⌵ The index can be used to seek on the second column if there is an equality predicate on the first column

⌵ True:

⌵ `FirstName = 'L' AND EmailPromotion = 1`

⌵ Partially True:

⌵ `FirstName LIKE 'L%' AND EmailPromotion = 1`

⌵ False

⌵ `FirstName LIKE '%L' AND EmailPromotion = 1`

The Optimizer's choice of Indexes

⌚ Multi-column index: Our scenario

Predicate

[AdventureWorks].[Person].[contact2].[EmailPromotion] as [c].[EmailPromotion]=(1) AND [AdventureWorks].[Person].[contact2].[ContactID] as [c].[ContactID]<(10000) AND [AdventureWorks].[Person].[contact2].[FirstName] as [c].[FirstName] like N'L%'

Object

[AdventureWorks].[Person].[contact2].[ContactComposite4]
[c]

Output List

[AdventureWorks].[Person].[contact2].ContactID,
[AdventureWorks].[Person].[contact2].FirstName,
[AdventureWorks].[Person].[contact2].EmailPromotion

Seek Predicates

Seek Keys[1]: Start: [AdventureWorks].[Person].[contact2].FirstName, [AdventureWorks].[Person].[contact2].EmailPromotion >= Scalar Operator(N'L'), Scalar Operator((1)), End: [AdventureWorks].[Person].[contact2].FirstName < Scalar Operator(N'M')

Predicate

[AdventureWorks].[Person].[contact2].[ContactID] as [c].[ContactID]<(10000) AND [AdventureWorks].[Person].[contact2].[FirstName] as [c].[FirstName] like N'L%'

Object

[AdventureWorks].[Person].[contact2].[ContactComposite5]
[c]

Output List

[AdventureWorks].[Person].[contact2].ContactID,
[AdventureWorks].[Person].[contact2].FirstName,
[AdventureWorks].[Person].[contact2].EmailPromotion

Seek Predicates

Seek Keys[1]: Prefix: [AdventureWorks].[Person].[contact2].EmailPromotion = Scalar Operator((1)), Start: [AdventureWorks].[Person].[contact2].FirstName >= Scalar Operator(N'L'), End: [AdventureWorks].[Person].[contact2].FirstName < Scalar Operator(N'M')

and the DTA ???




Agenda

- ④ Demo: Indexing Scenario 1
- ④ Optimizer's choice of Physical Join Operator
- ④ Indexing Scenario 2
- ④ Demo: Good Design vs Better Design
- ④ The Optimizer's choice of Indexes
- ④ Demo: Indexing Scenario 3
- ④ How, When & Why did the Indexes become culprits?
- ④ Demo: Clean up act !
- ④ Summary / Call to Action



demo



Indexing Scenario 3

How, When & Why did the indexes become culprits?

- ④ Too many Non-Clustered indexes?
- ④ Did you strike a right balance between SELECT & UPDATE performance?
- ④ Can you get rid of some redundant indexes?



Image source: <http://onlinewebfun.com/1835/truckload/>

How, When & Why did the indexes become culprits?

```
SELECT C.ContactID, C.FirstName, C.EmailPromotion
FROM Person.Contact2 AS C
WHERE C.FirstName LIKE 'L%'
      AND C.EmailPromotion = 1
      AND C.ContactID < 10000
OPTION (MAXDOP 1)
```

```
SELECT C.ContactID, C.FirstName, C.EmailPromotion
FROM Person.Contact2 AS C
WHERE C.FirstName = 'LANE'
      AND C.EmailPromotion =< 1
      AND C.ContactID < 10000
OPTION (MAXDOP 1)
```

```
SELECT c2.FirstName
FROM person.Contact2 AS c2
INNER JOIN
person.contact3 c3
ON C2.FirstName = c3.FirstName
ORDER BY C3.FirstName
```


```
CREATE INDEX ContactComposite5
ON Person.Contact2(EmailPromotion,
FirstName)
```

```
CREATE INDEX ContactComposite4
ON Person.Contact2(firstname,
EmailPromotion)
```

```
CREATE INDEX C2FirstName
ON Person.Contact2(firstname)
```



demo



Clean up act !

Summary / Call to Action

- ④ Browse this recording once again 😊
- ④ Try out the code snippets yourself
- ④ Review your indexing strategies
- ④ Implement the knowledge
- ④ Try out various combinations in your 'test environment'

Takeaways...

- ④ How does the optimizer decide which index to use?
- ④ Optimizer's choice of the physical join operator
- ④ Database Tuning Advisor. When? Why? How?
- ④ Monitor, Monitor, Monitor !

Complete an evaluation on CommNet and enter to win!



Track Resources



SQL Server 2012
Eval Copy



Hands-On Labs

mva

Microsoft Virtual Academy



Get Certified!



@sqlserver
@ms_teched

Resources

mytech·ed
Learn. Connect. Share.

Connect. Share. Discuss.

<http://northamerica.msteched.com>

Microsoft | Learning

Microsoft Certification & Training Resources

www.microsoft.com/learning

Microsoft | TechNet

Resources for IT Professionals


<http://microsoft.com/technet>

 **msdn**

Resources for Developers

<http://microsoft.com/msdn>

Thank you all !



MS Tag

Scan the Tag
to evaluate this
session now on
myTechEd Mobile



Required Slide

*delete this box when your slide is finalized

Your MS Tag will be inserted here during the final scrub.



The Microsoft logo is centered on a dark blue rectangular background. The logo itself is white and features the word "Microsoft" in a bold, italicized, sans-serif typeface. A registered trademark symbol (®) is positioned at the top right of the word.

© 2012 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries. The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.