

SQLServerFast.com

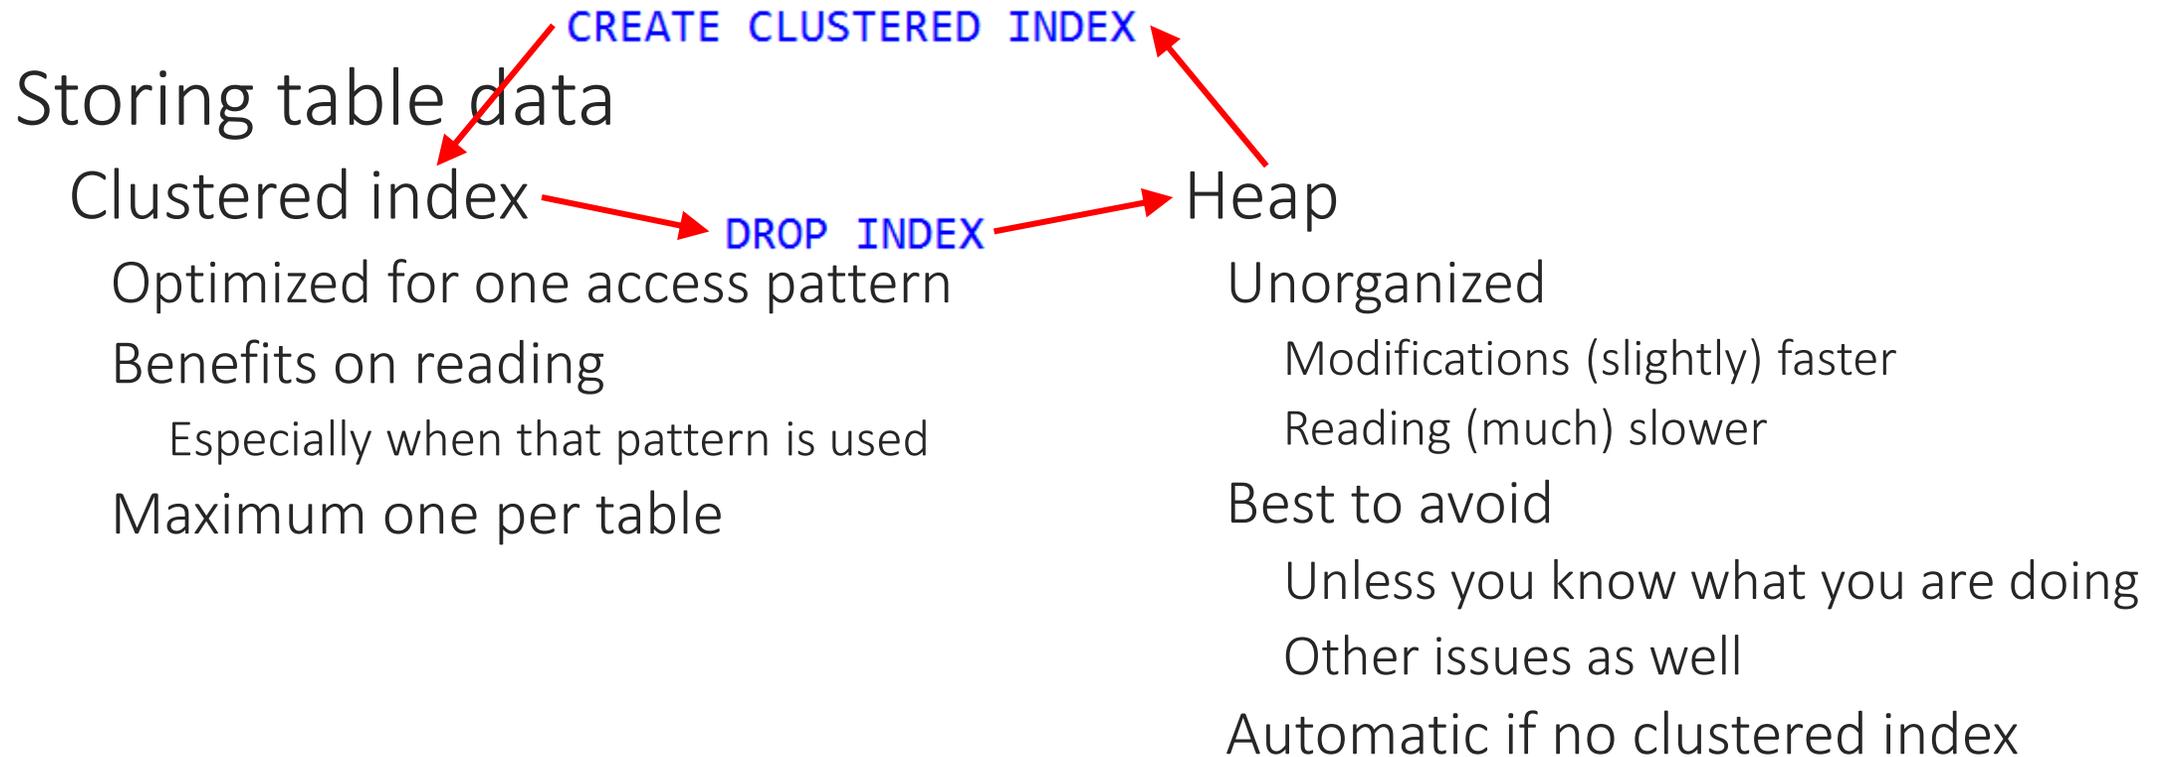
Execution Plan Video Training

Block 2: Reading data

Level: Basic

Chapter 1: Storage structures

Storage structures



Storage structures

Storing table data

- Clustered index

- Heap

- Nonclustered index

 - Multiple per table allowed

 - Up to 999 per table

 - This is the maximum, not a target

 - Realistically, more than 10 is often already suspect

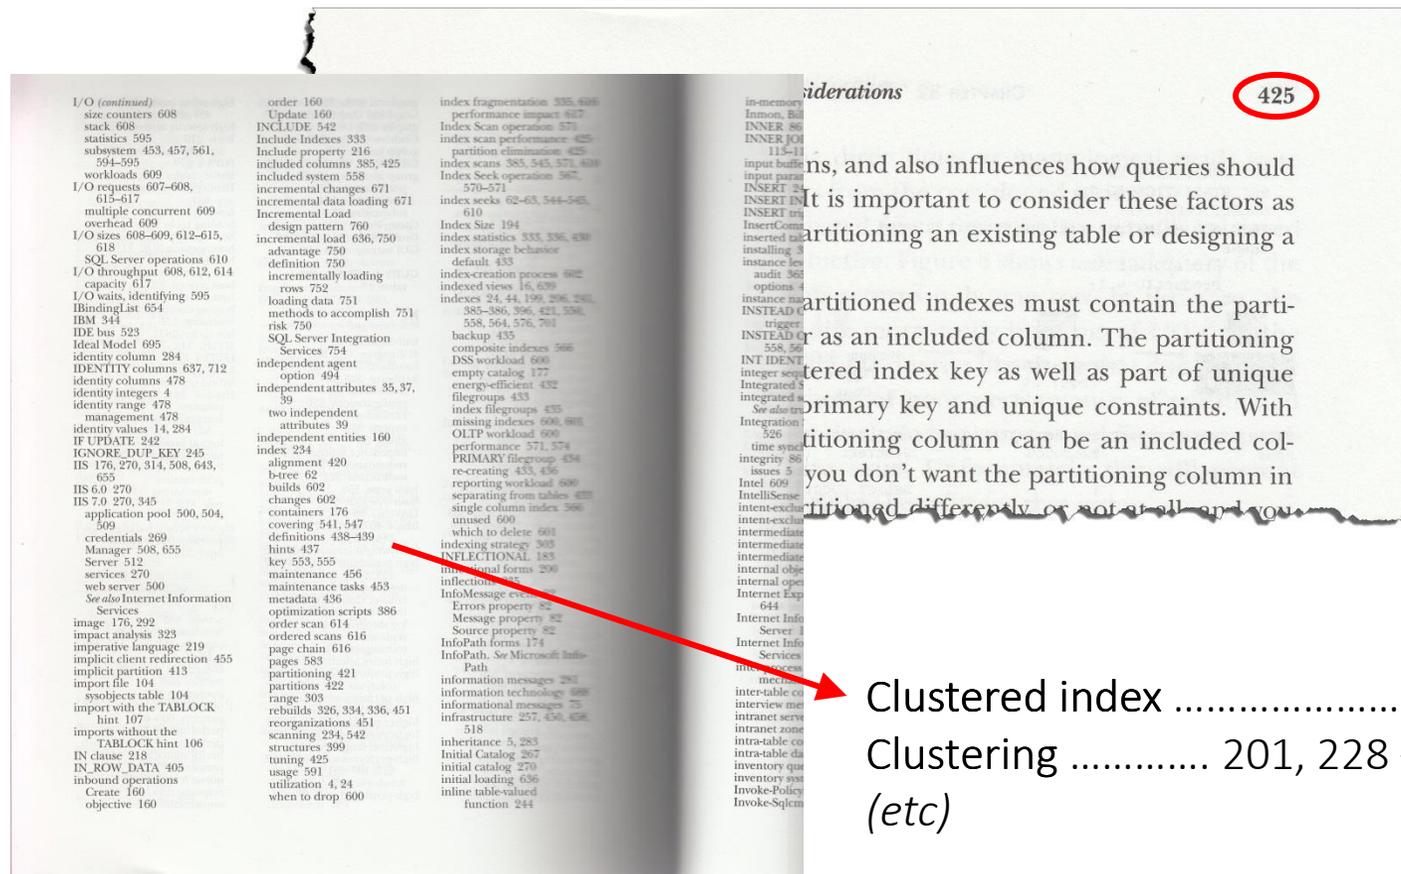
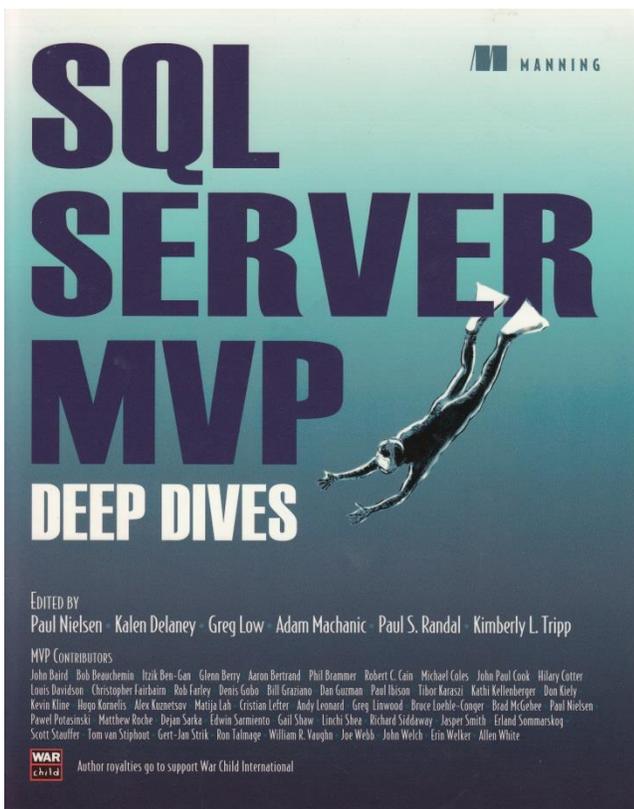
 - Optimized for specific access pattern

 - Stores copy of some data

 - Also includes pointer to rest of the data

Storage structures

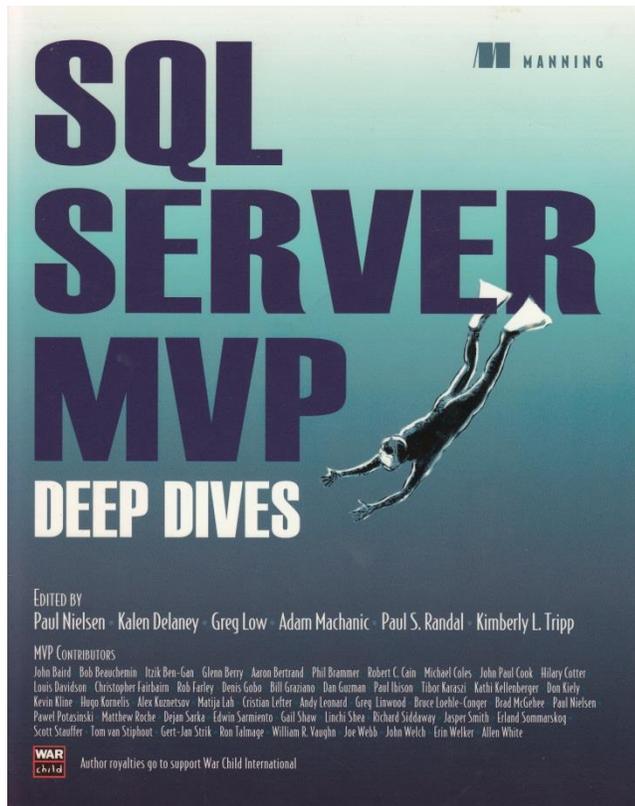
Storing table data



Clustered index 425
Clustering 201, 228 – 230
(etc)

Storage structures

Storing table data



Clustered index **XXX**
Clustering xxx, xxx, xxx

Storage structures

Storing table data

Logical structure

Heap

- All data

- Not organized

Clustered index

- All data

- Organized

Nonclustered index

- Subset of data

- Organized

- Pointer to rest of data

Storage structures

Storing table data

- Logical structure

- Physical structure

 - On-disk rowstore

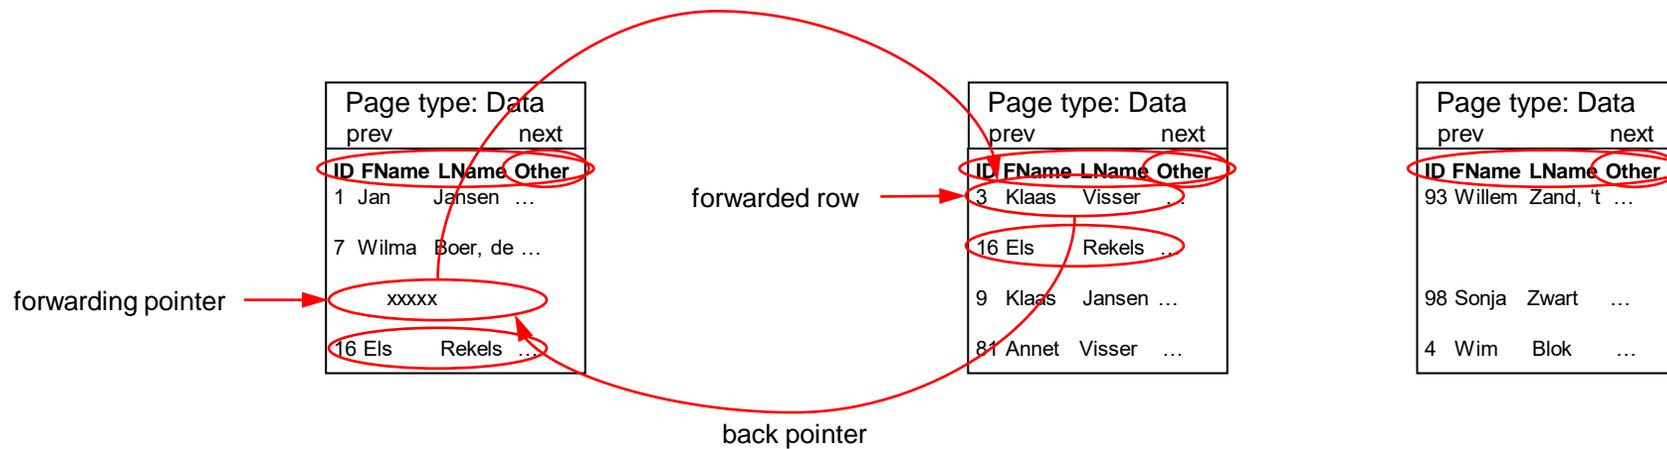
 - Columnstore

 - Memory-optimized

 - Rowstore and columnstore

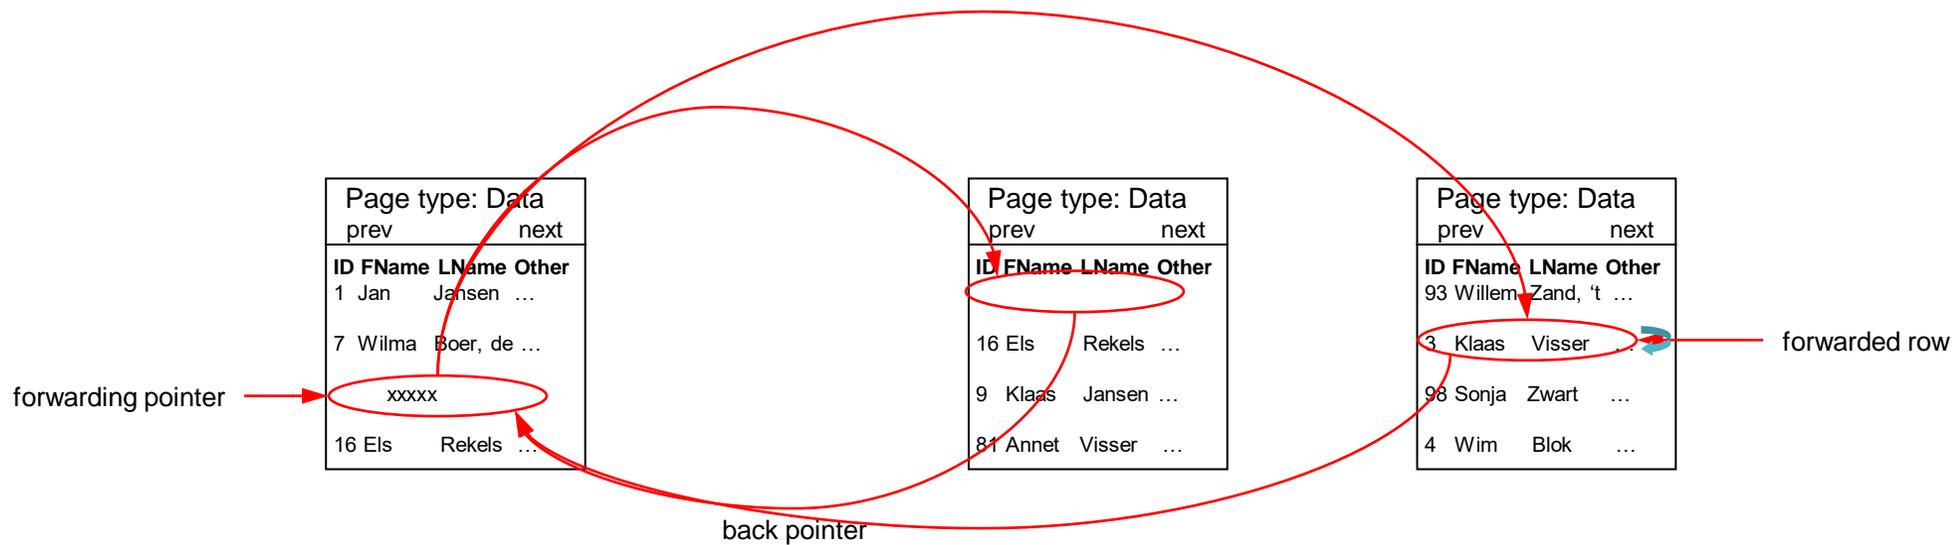
Storage structures

Heap table



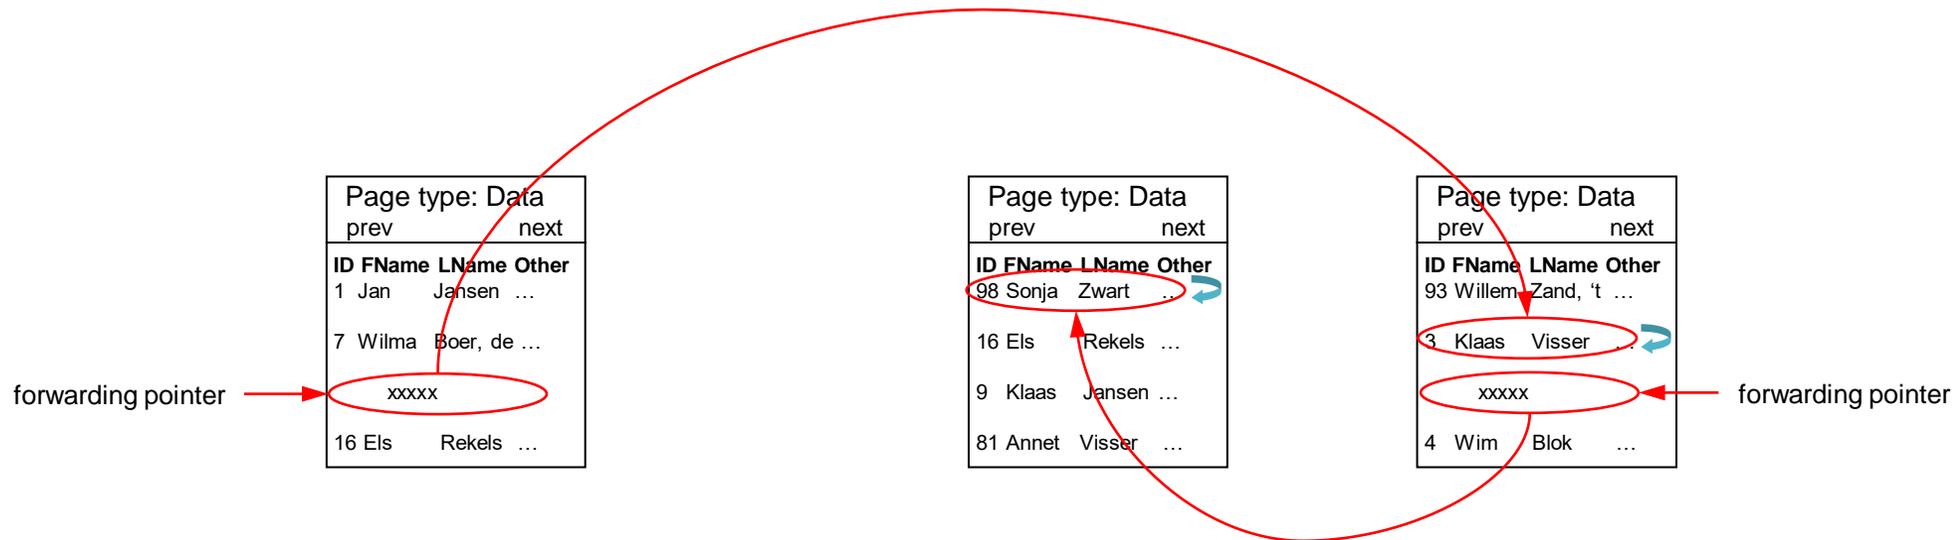
Storage structures

Heap table



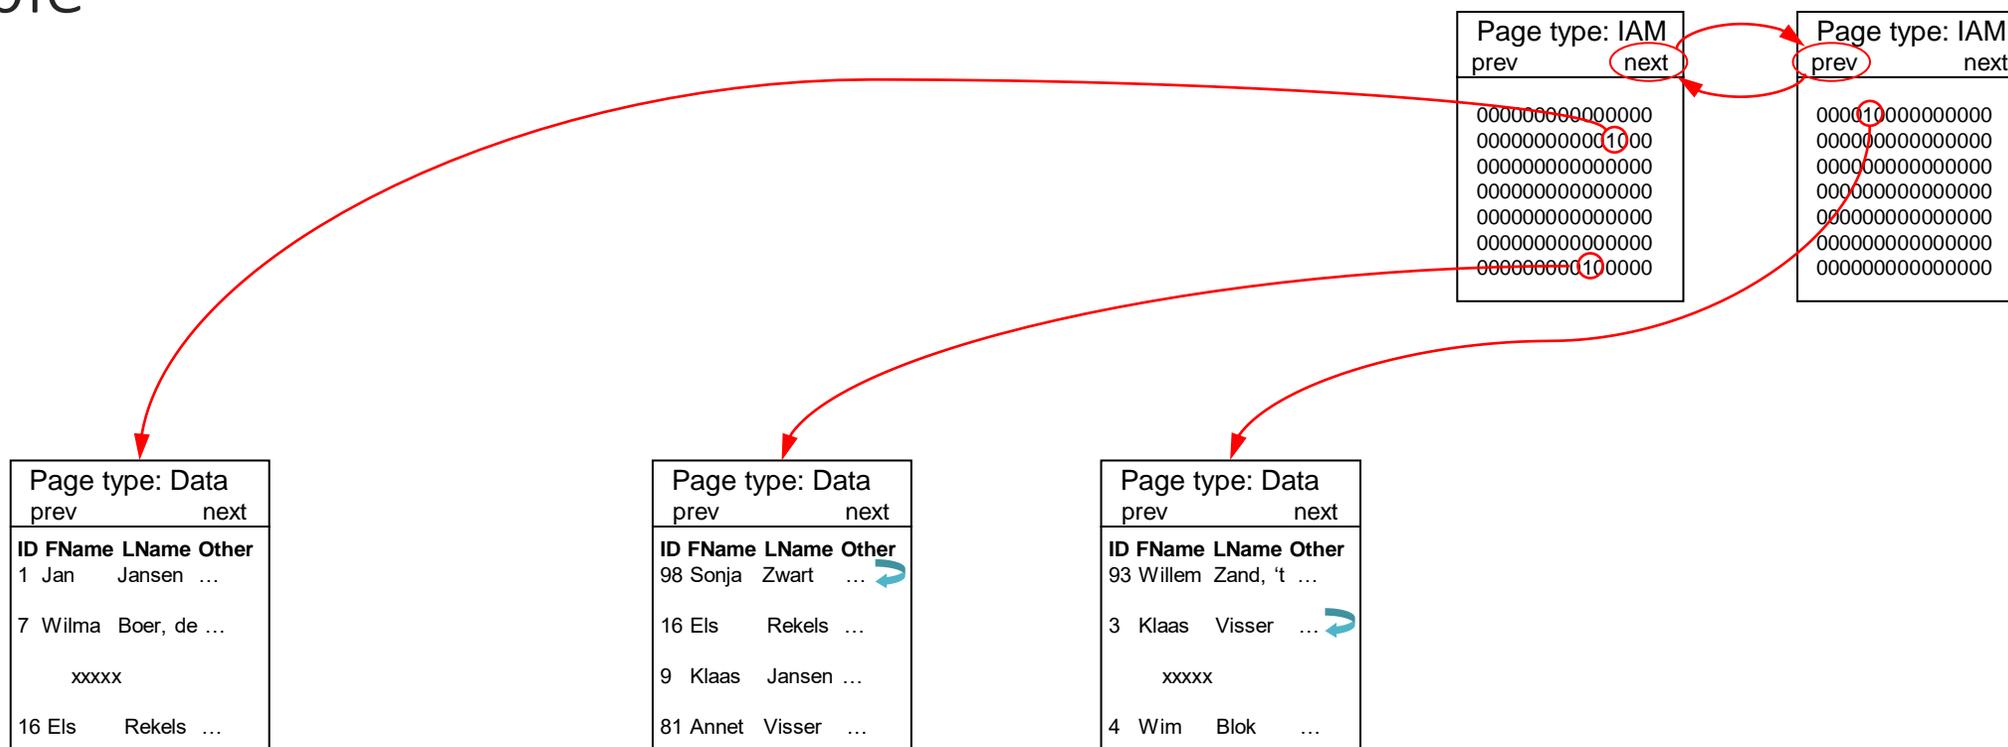
Storage structures

Heap table



Storage structures

Heap table



Storage structures

Clustered index

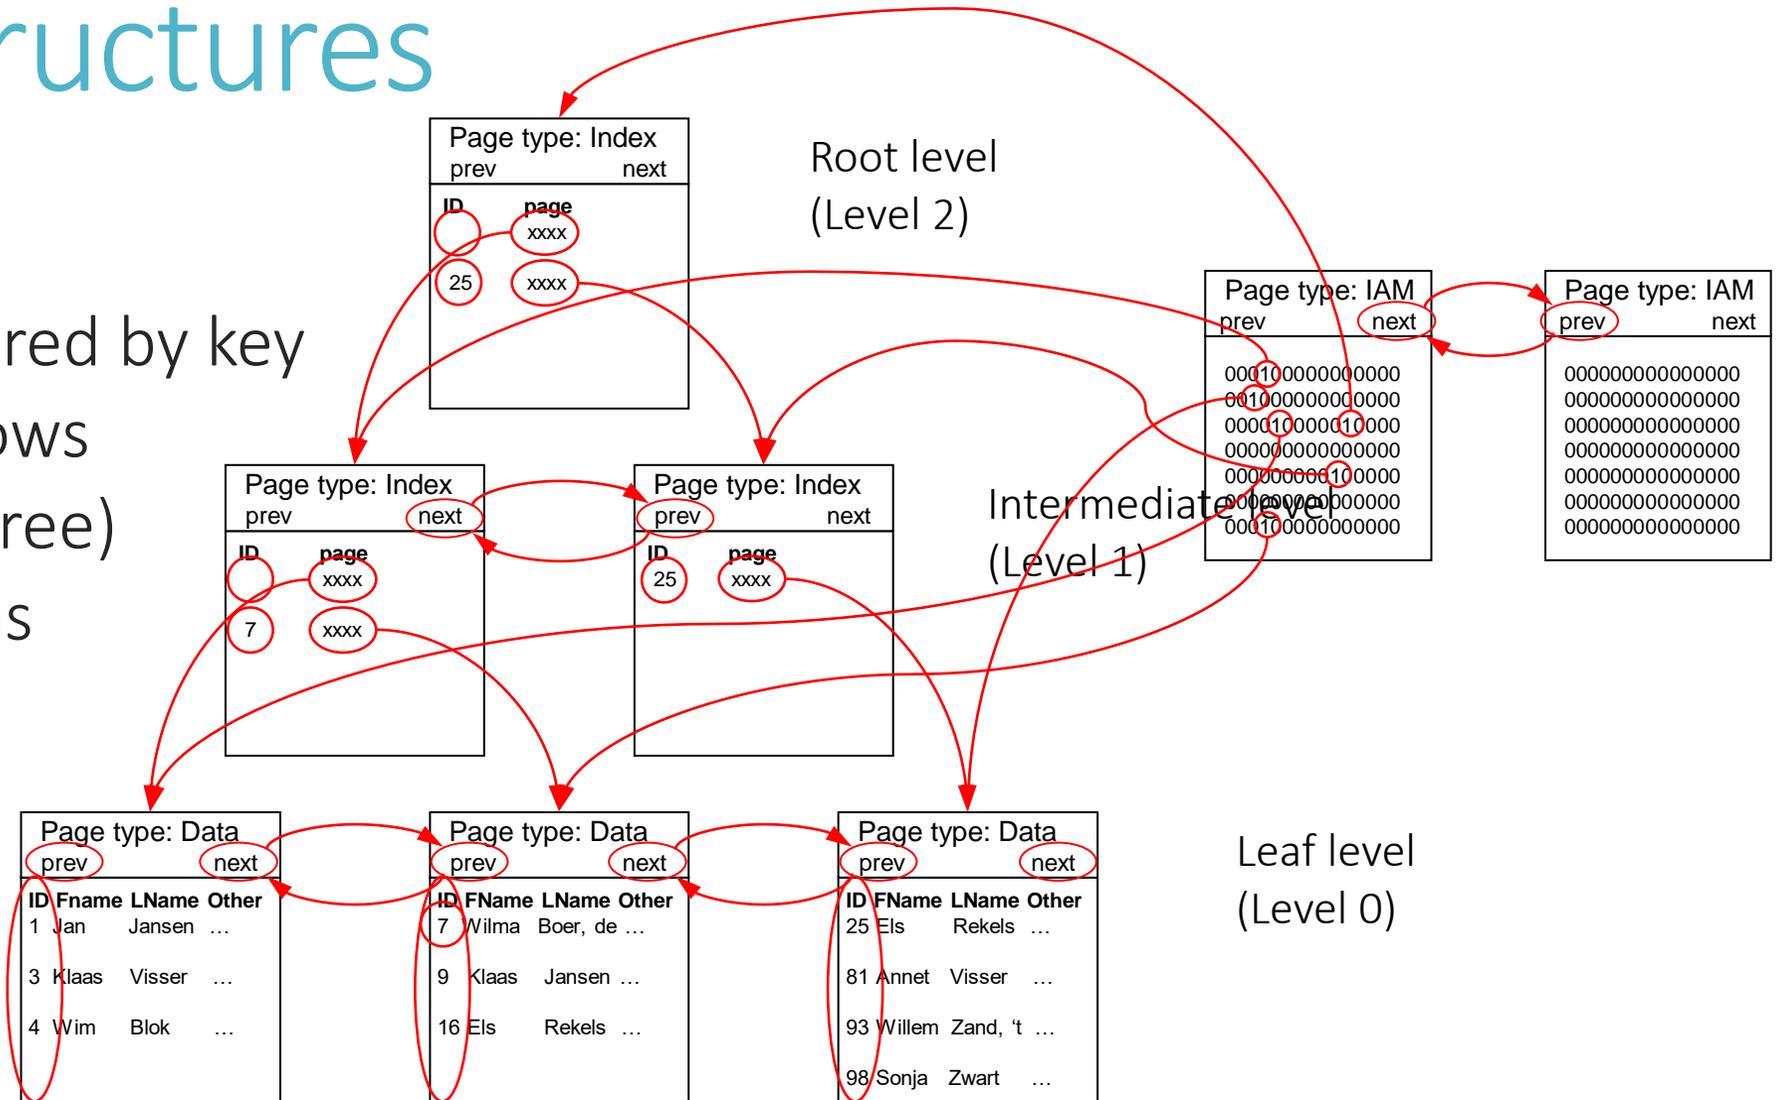
Data pages ordered by key

No forwarded rows

Index pages (B-tree)

As many levels as needed

No maximum



Storage structures

Clustered index

Data pages ordered by key

No forwarded rows

Index pages (B-tree)

As many levels as needed

No maximum

Page type: Index	
prev	next
ID	page
	xxxx
25	xxxx

Page type: IAM	
prev	next
0001000000000000	
0010000000000000	
0000100000100000	
0000000000000000	
0000000001000000	
0000000000000000	
0001000000000000	

Page type: Index	
prev	next
ID	page
	xxxx
7	xxxx

Page type: Index	
prev	next
ID	page
25	xxxx

Page type: Data			
prev	next		
ID	FName	LName	Other
1	Jan	Jansen	...
3	Klaas	Visser	...
4	Wim	Blok	...

Page type: Data			
prev	next		
ID	FName	LName	Other
7	Wilma	Boer, de	...
9	Klaas	Jansen	...
16	Els	Rekels	...

Page type: Data			
prev	next		
ID	FName	LName	Other
25	Els	Rekels	...
81	Annet	Visser	...
93	Willem	Zand, 't	...
98	Sonja	Zwart	...

Storage structures

Clustered index

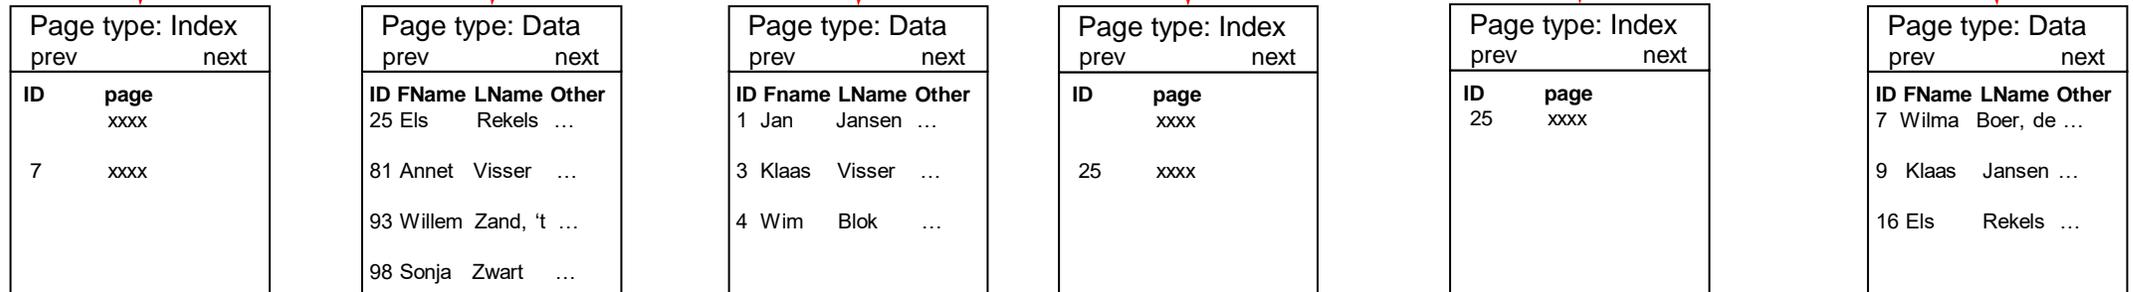
Data pages ordered by key

No forwarded rows

Index pages (B-tree)

As many levels as needed

No maximum



Storage structures

Nonclustered index

Data included

Indexed column(s)

Clustered index key

Overlapping columns stored once

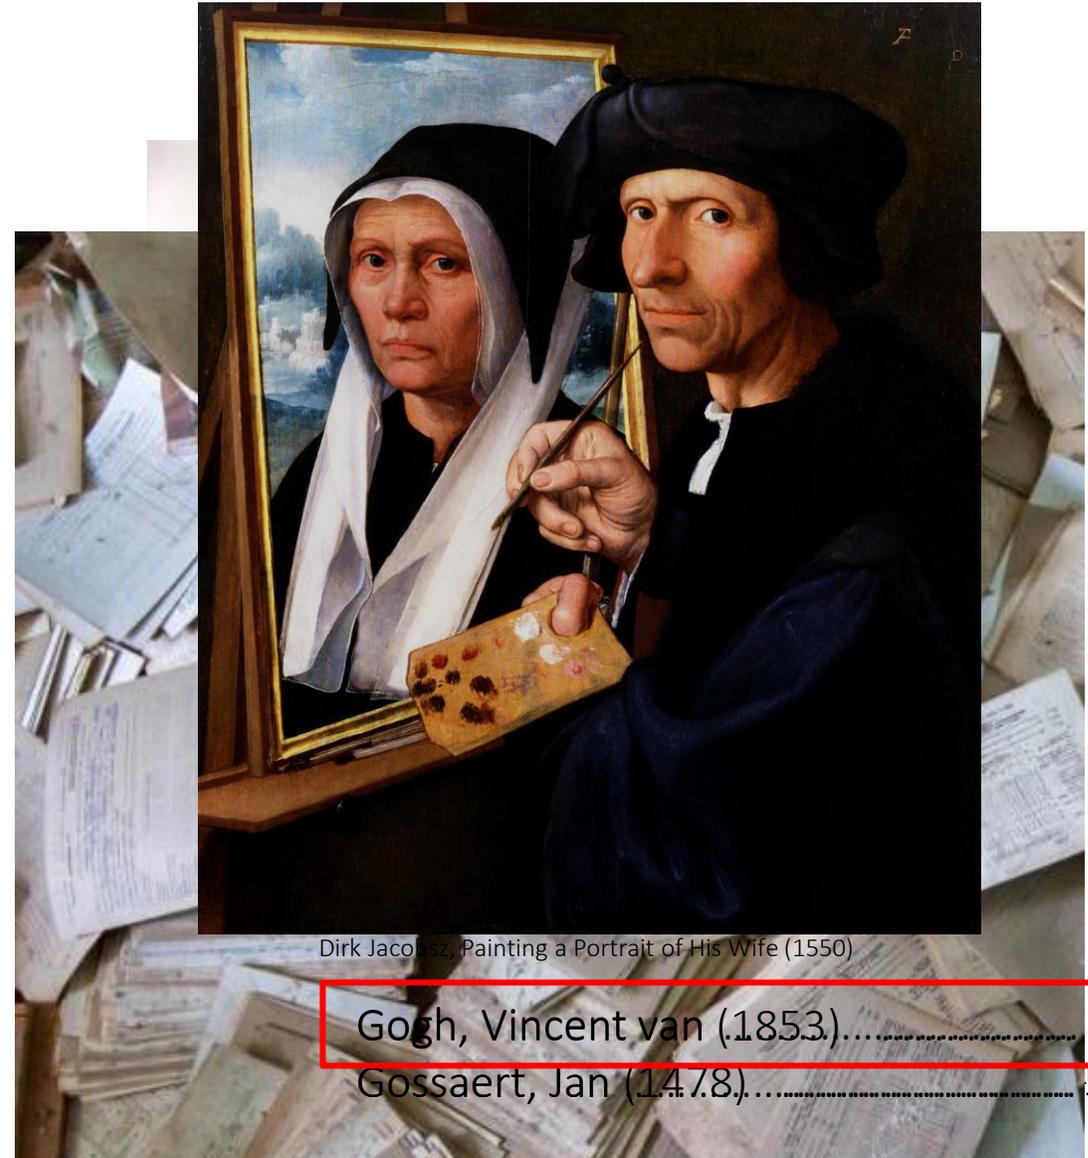
For heaps: RID (row identifier)

File number

Page number

Row slot number

INCLUDEd columns



Storage structures

Nonclustered index

Data included

Indexed column(s)

Clustered index key

Overlapping columns stored once

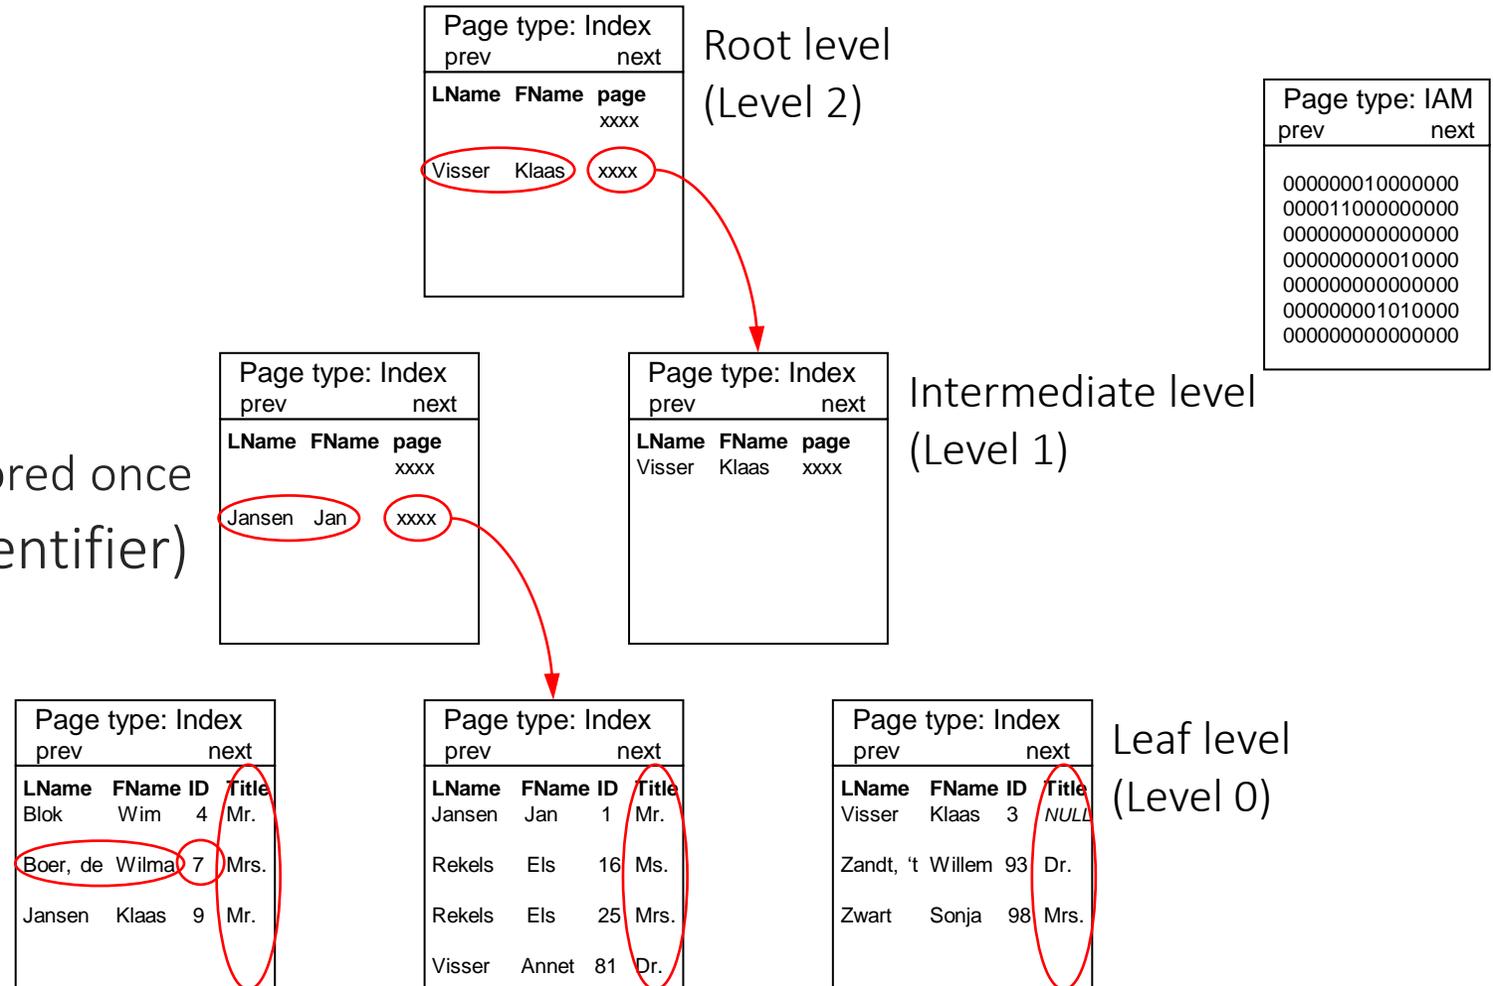
For heaps: RID (row identifier)

File number

Page number

Row slot number

INCLUDEd columns



Storage structures

Nonclustered index

Data included

Indexed column(s)

Clustered index key

Overlapping columns stored once

For heaps: RID (row identifier)

INCLUDEd columns

Page type: Index		
prev	next	
LName	FName	page
		xxxx
Visser	Klaas	xxxx

Root level
(Level 2)

Page type: IAM	
prev	next
000000010000000	
000011000000000	
000000000000000	
000000000010000	
000000000000000	
000000001010000	
000000000000000	

Page type: Index		
prev	next	
LName	FName	page
		xxxx
Jansen	Jan	xxxx

Page type: Index		
prev	next	
LName	FName	page
Visser	Klaas	xxxx

Intermediate level
(Level 1)

Page type: Index		
prev	next	
LName	FName	RID
Blok	Wim	1,95,4
Boer, de	Wilma	1,18,2
Jansen	Klaas	1,31,3

Page type: Index		
prev	next	
LName	FName	RID
Jansen	Jan	1,18,1
Rekels	Els	1,31,2
Rekels	Els	1,18,4
Visser	Annet	1,31,4

Page type: Index		
prev	next	
LName	FName	RID
Visser	Klaas	1,18,3
Zandt, 't	Willen	1,95,1
Zwart	Sonja	1,95,4

Leaf level
(Level 0)

Summary

Storage structures (for on-disk rowstore data)

Heap

- Unorganized

- Forwarding pointers

Clustered index

- Index columns

- Rest of data

Nonclustered index

- To clustered index, or to heap

Next chapters

Chapter 2: Scan operators

- Table Scan

- Index Scan

- Clustered Index Scan

Chapter 3: Seek operators

Chapter 4: Lookup operators

Chapter 5: Special scans