



The Db2Night Show

# Building Blocks of Db2

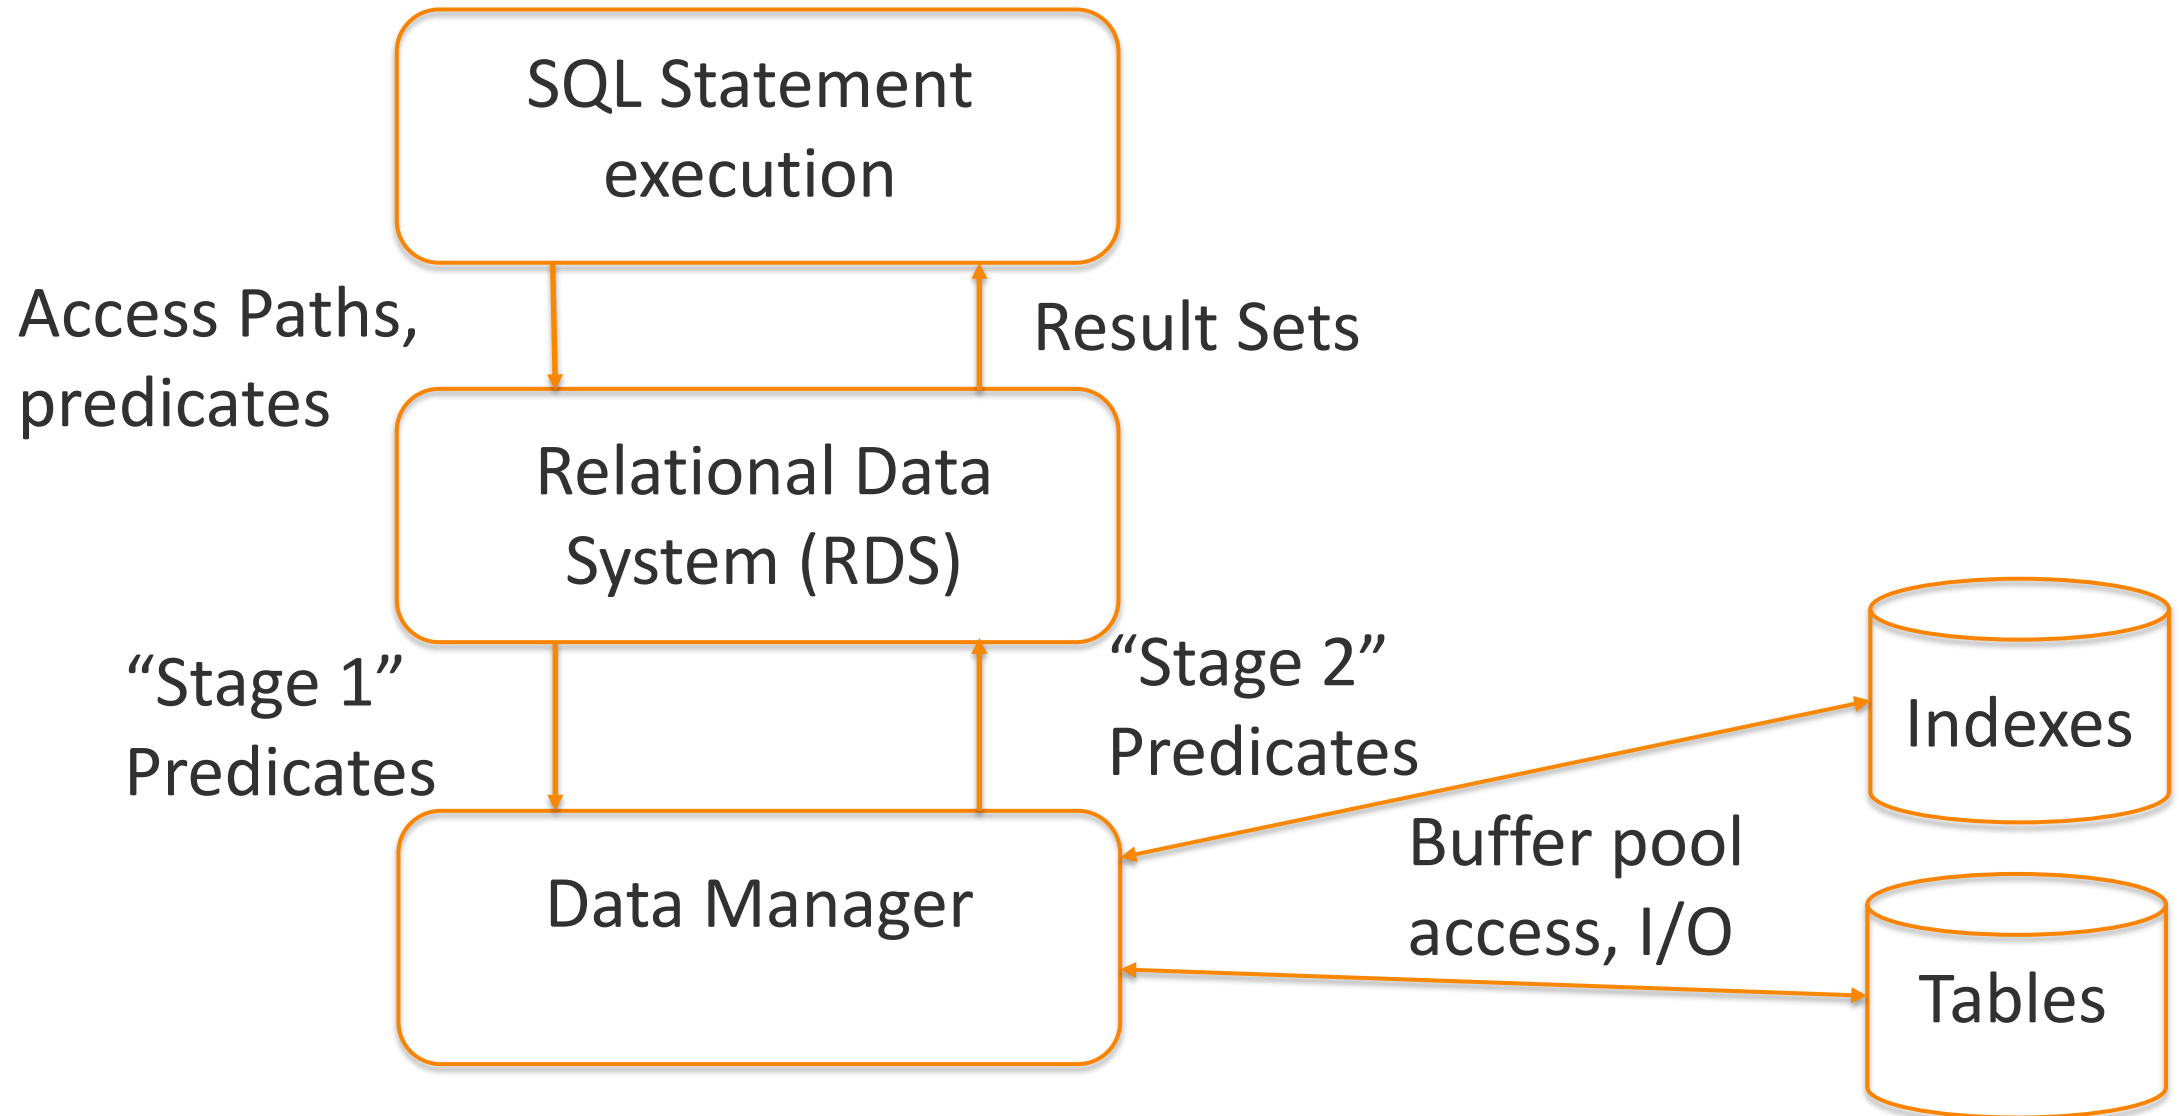
**Jim Dee**

Corporate Architect, BMC Software

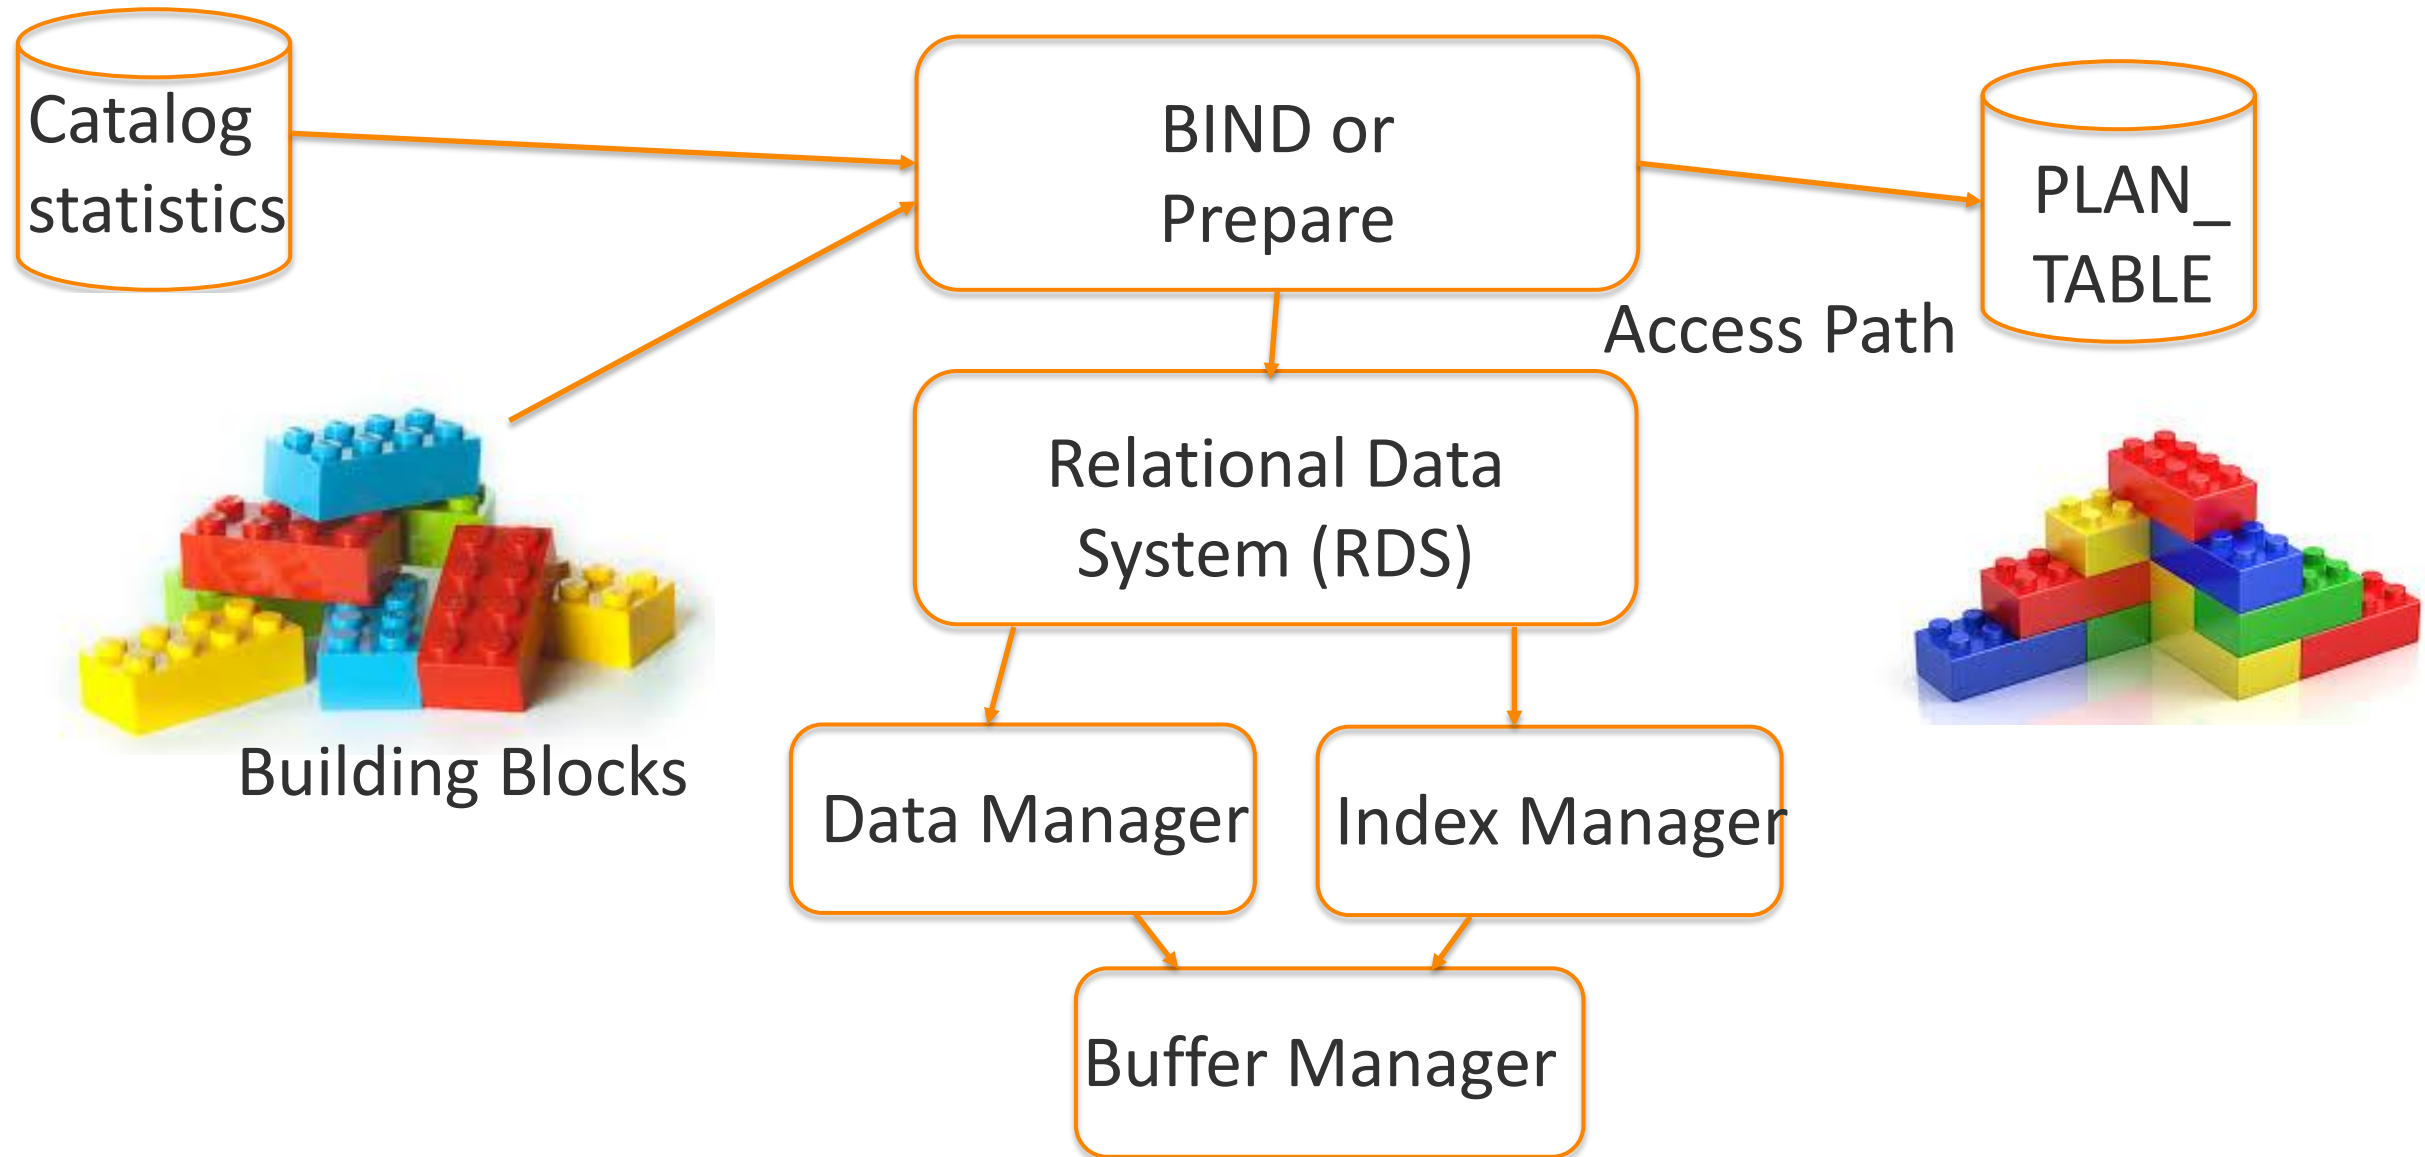
October 22, 2021



# DB2 SQL Execution Architecture

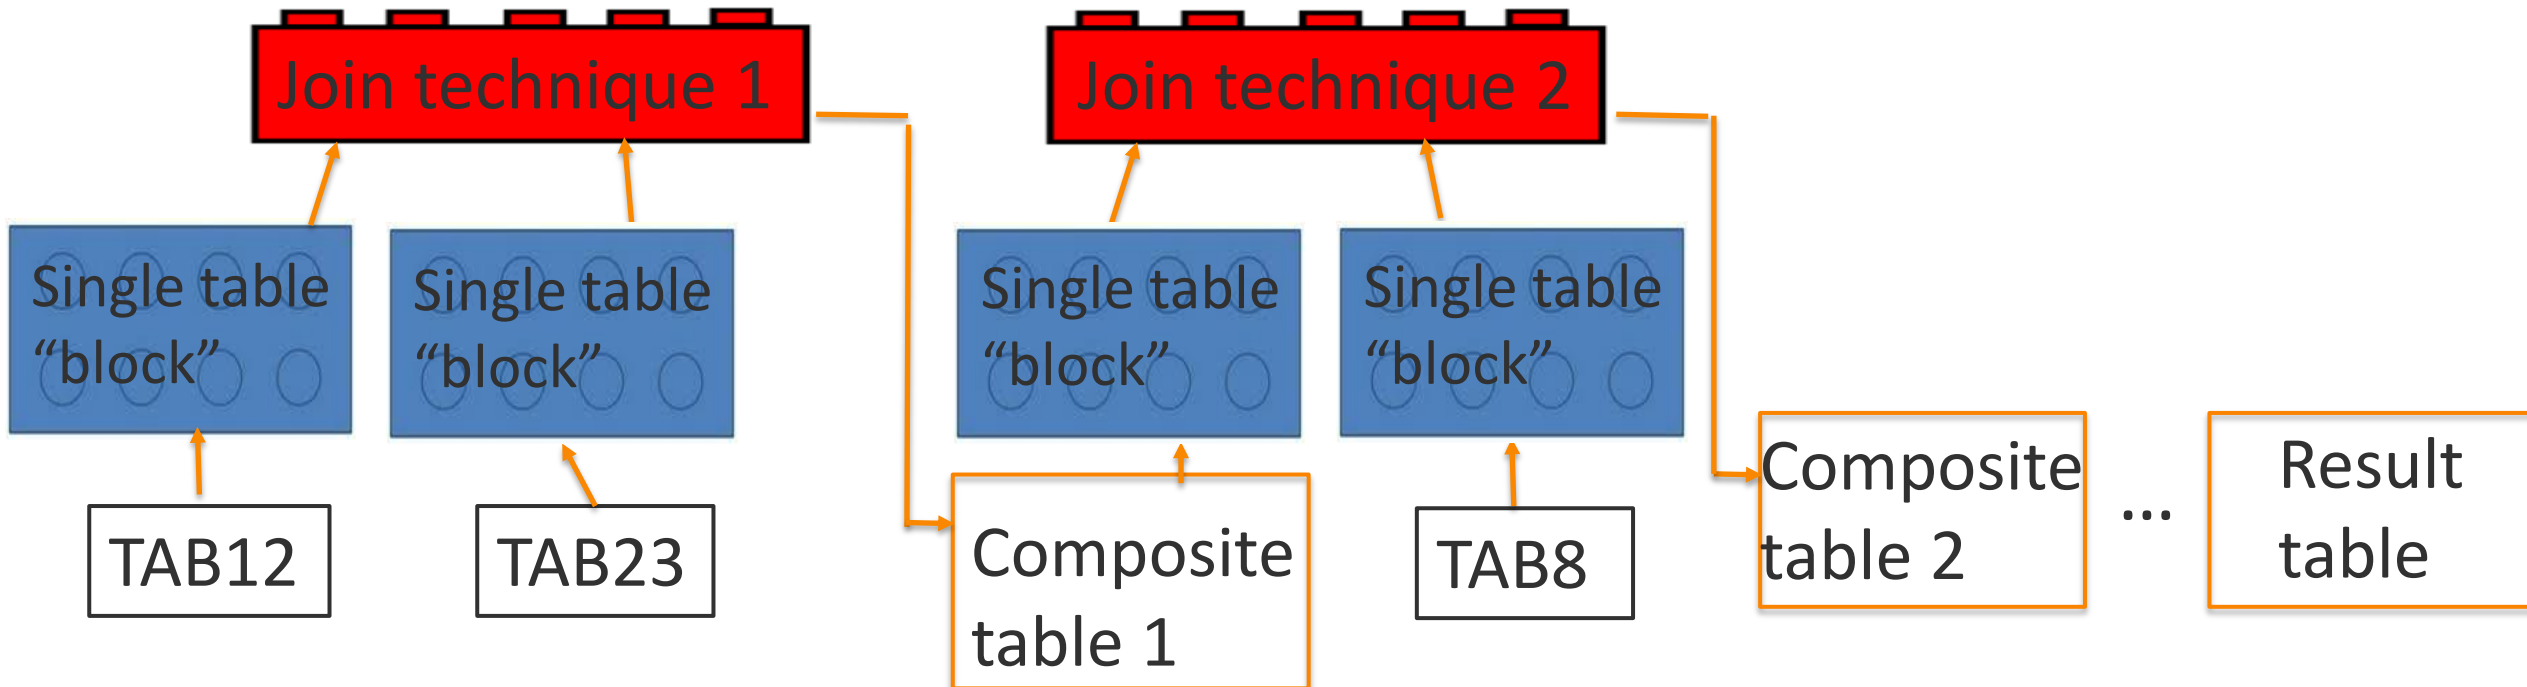


# The DB2 Optimizer and RDS



# Joins

```
SELECT A.COL1, B.COL2, ... Z.COL26
FROM TAB1 A, TAB2 B, ... TAB26 Z
WHERE A.JCOL1 = B.JCOL2
      AND B.JCOL3 = C.JCOL4
      AND ...
```



# Predicates

**After EXPLAIN, the following columns are of interest:**

**SEARCHARG in DSN\_PREDICAT\_TABLE ('Y' => Stage 1)**


**TYPE, LEFT\_HAND\_SIDE, RIGHT\_HAND\_SIDE**

**STAGE in DSN\_FILTER\_TABLE (joined to PREDICAT by PREDNO)**

- Has more detail about ordering in predicate resolution


# Filter Factor

```
SELECT A.CUSTNAME, B.ORDERNO  
      FROM PROD.CUSTOMER A, PROD.ORDER B  
      WHERE A.CUST_RATING > 8  
            AND B.ORDERDATE > TODAY - 5  
            AND A.CUSTID = B.CUSTID;
```



# Tablespace Scan


```
SELECT LASTNAME, SSN FROM CUSTOMER;
```



No index  
used!

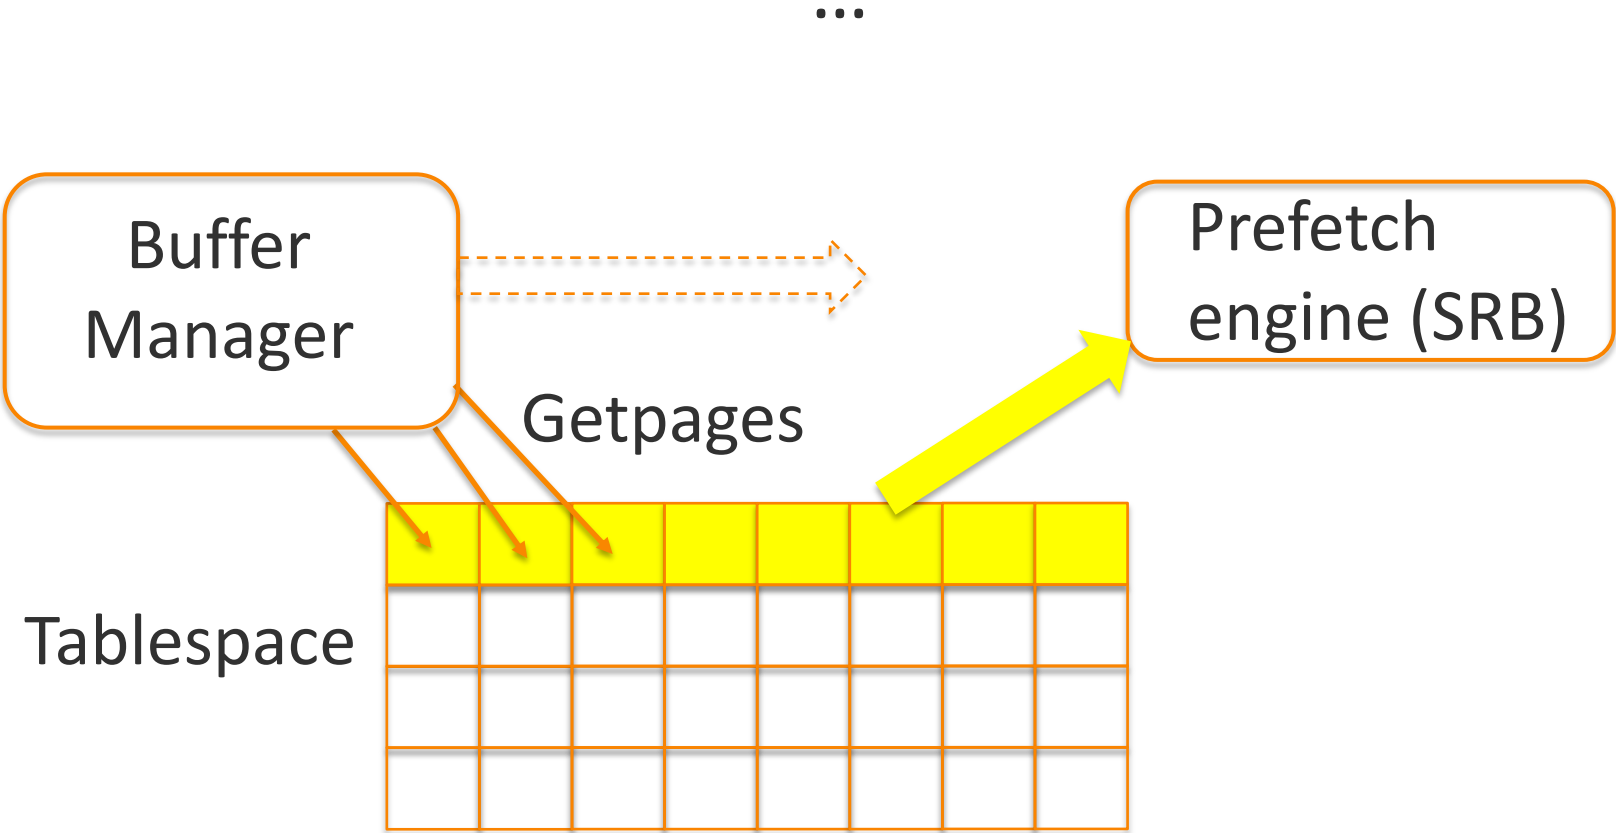


Getpages



Tablespace

# Sequential Prefetch





# Sequential Prefetch – More Detail

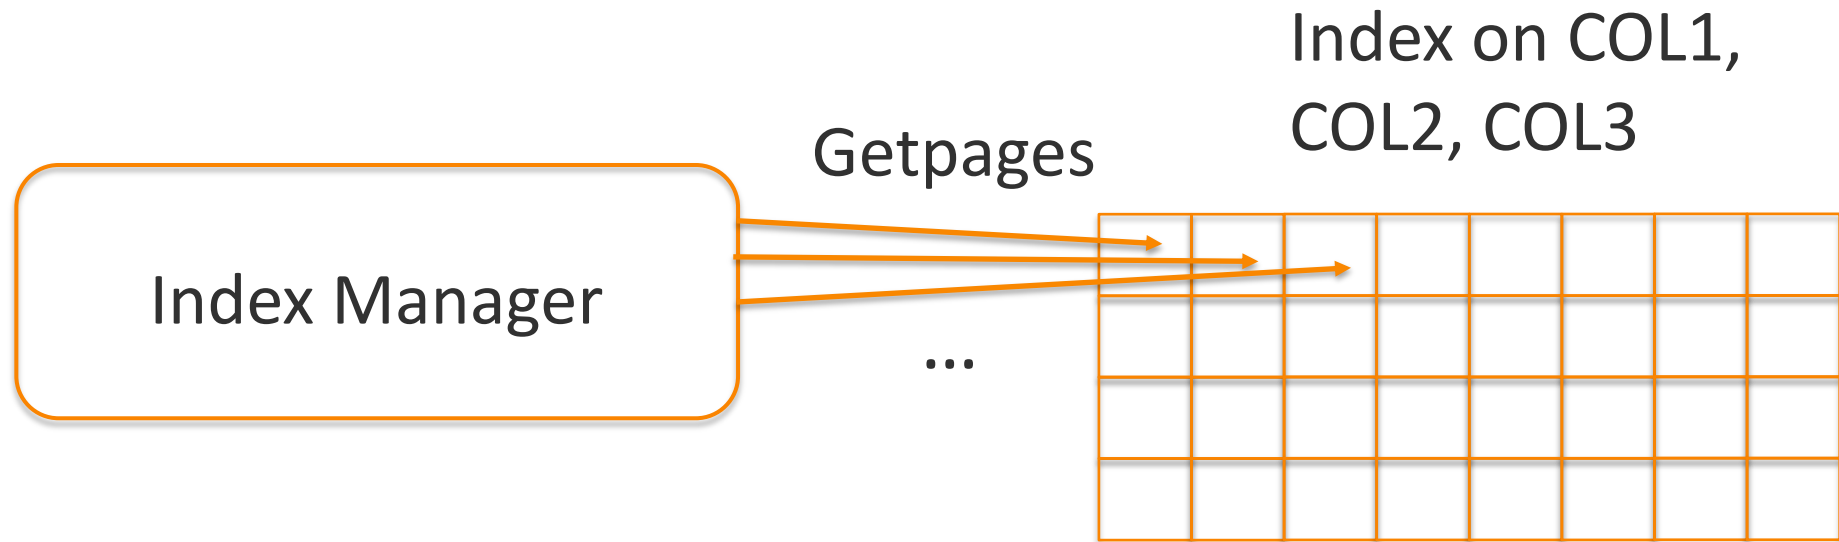
Data Manager	Buffer Manager	Prefetch Engines
Getpage page 89	Finds page 89 in buffer pool	
Getpage page 90	Starts prefetch engine	Issue reads for 90-127, 128-191
	Waits thread for page 90	
Getpage page 91	Finds page 91 in buffer pool	
...		
Getpage page 128	Triggers prefetch	Issue reads for 192-255
...		

# Sequential Prefetch

	PLANNO		METHOD		TABNO		ACCESSTYPE		MATCHCOLS		PREFETCH	
	2		2		2		R		0		S	

# Index Scan

```
SELECT COL1, COL2 FROM TAB;
```

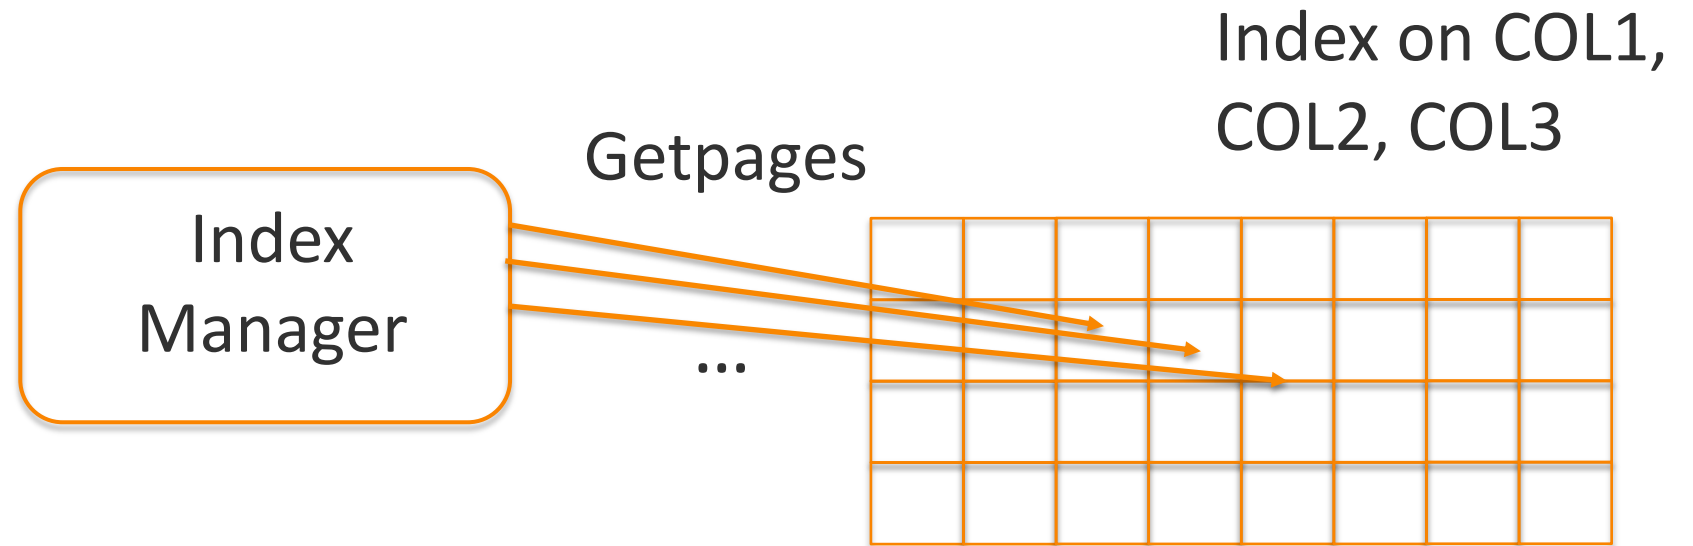


# Index Scan

PLANNO	METHOD	ACCESSTYPE	MATCHCOLS	PREFETCH	CREATOR	TNAME	TABNO
1	0	I	0	S	SYSIBM	SYSCOLUM	1

# Matching Index Scan

```
SELECT COL1, COL3, COL4 FROM TAB  
WHERE COL1 BETWEEN 3 AND 5  
AND COL3 = 'ABC' ;
```

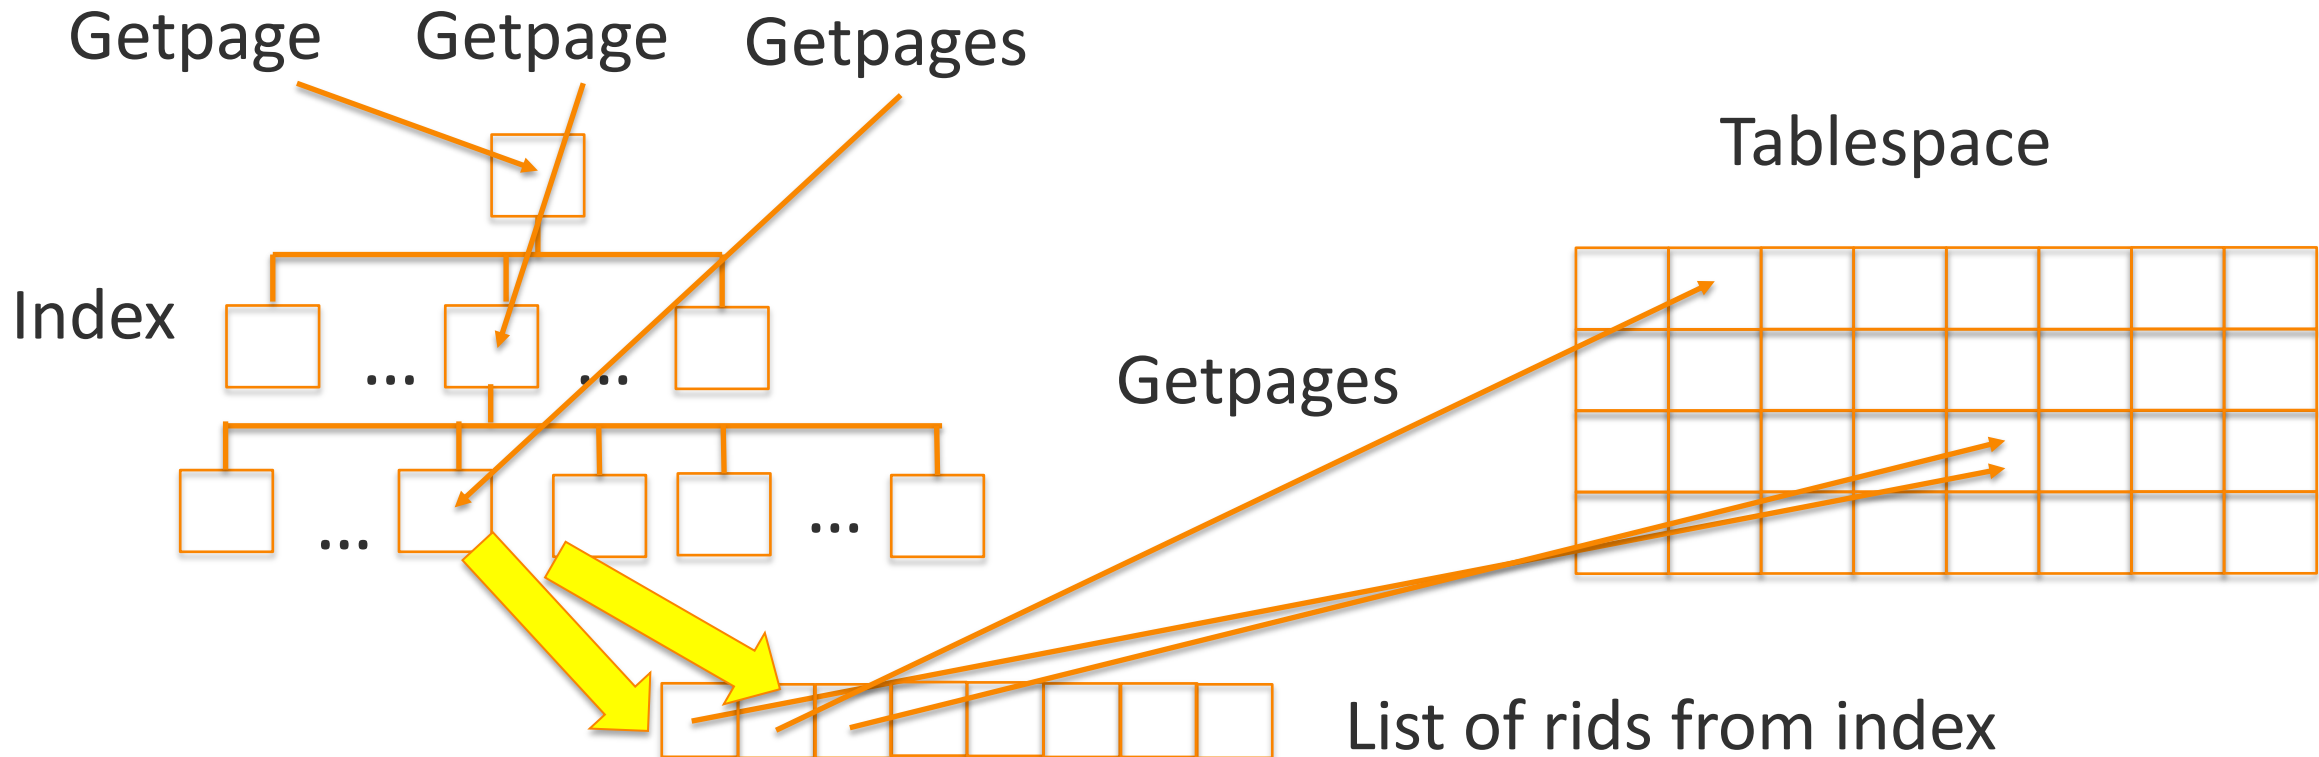


# Matching Index Scan

	PLANNO		METHOD		TABNO		ACCESSTYPE		MATCHCOLS		PREFETCH	
	1		0		1		I		1		S	

# Range Retrieval and List Prefetch

```
SELECT ORDERNO, ORDERDATE, AMOUNT FROM ORDER  
WHERE CUSTID = :Customer-id  
AND ORDERDATE BETWEEN :start and :end;
```

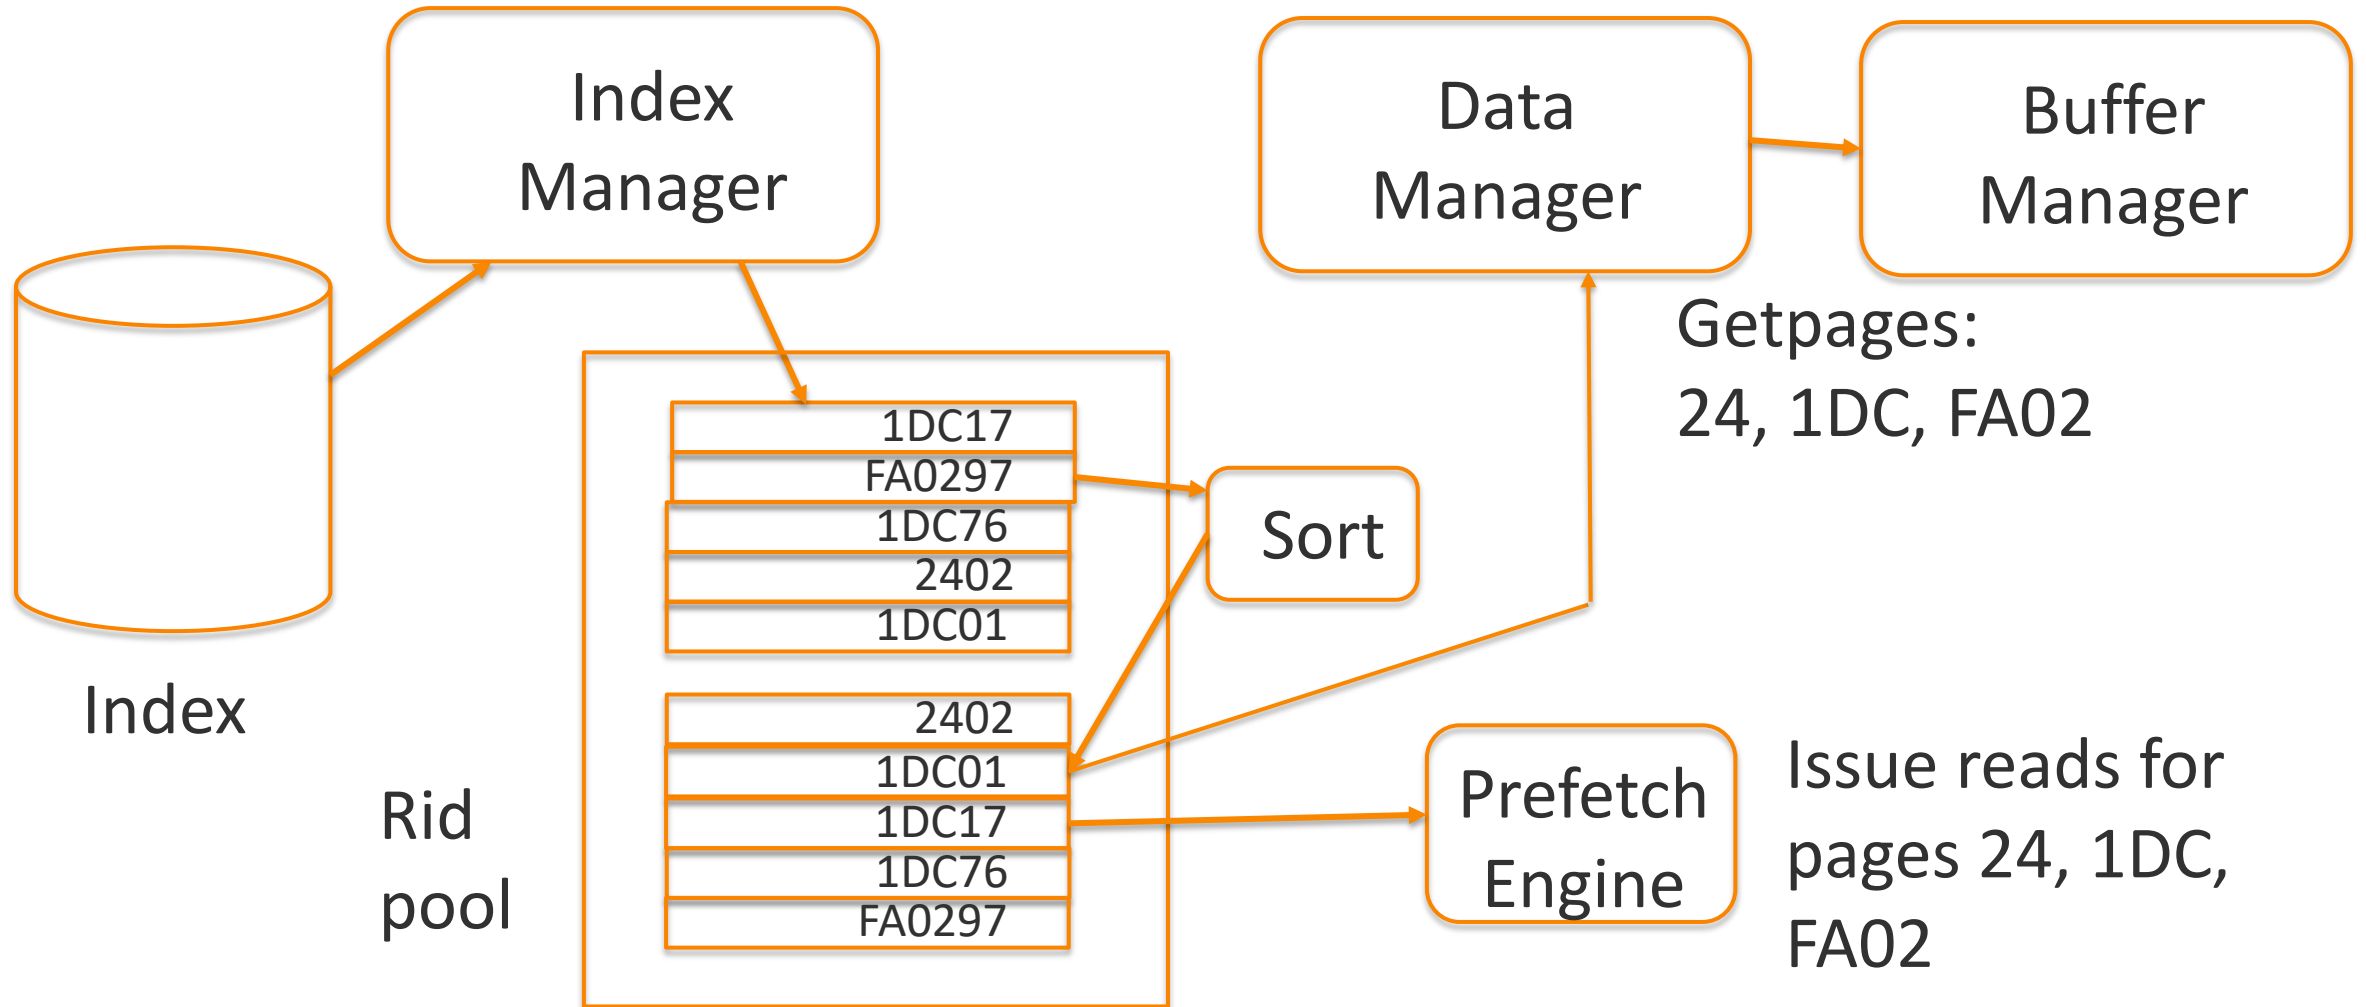


# List Prefetch

PLANNO	METHOD	ACCESSTYPE	MATCHCOLS	PREFETCH	TABNO
1	0	I	1	L	1

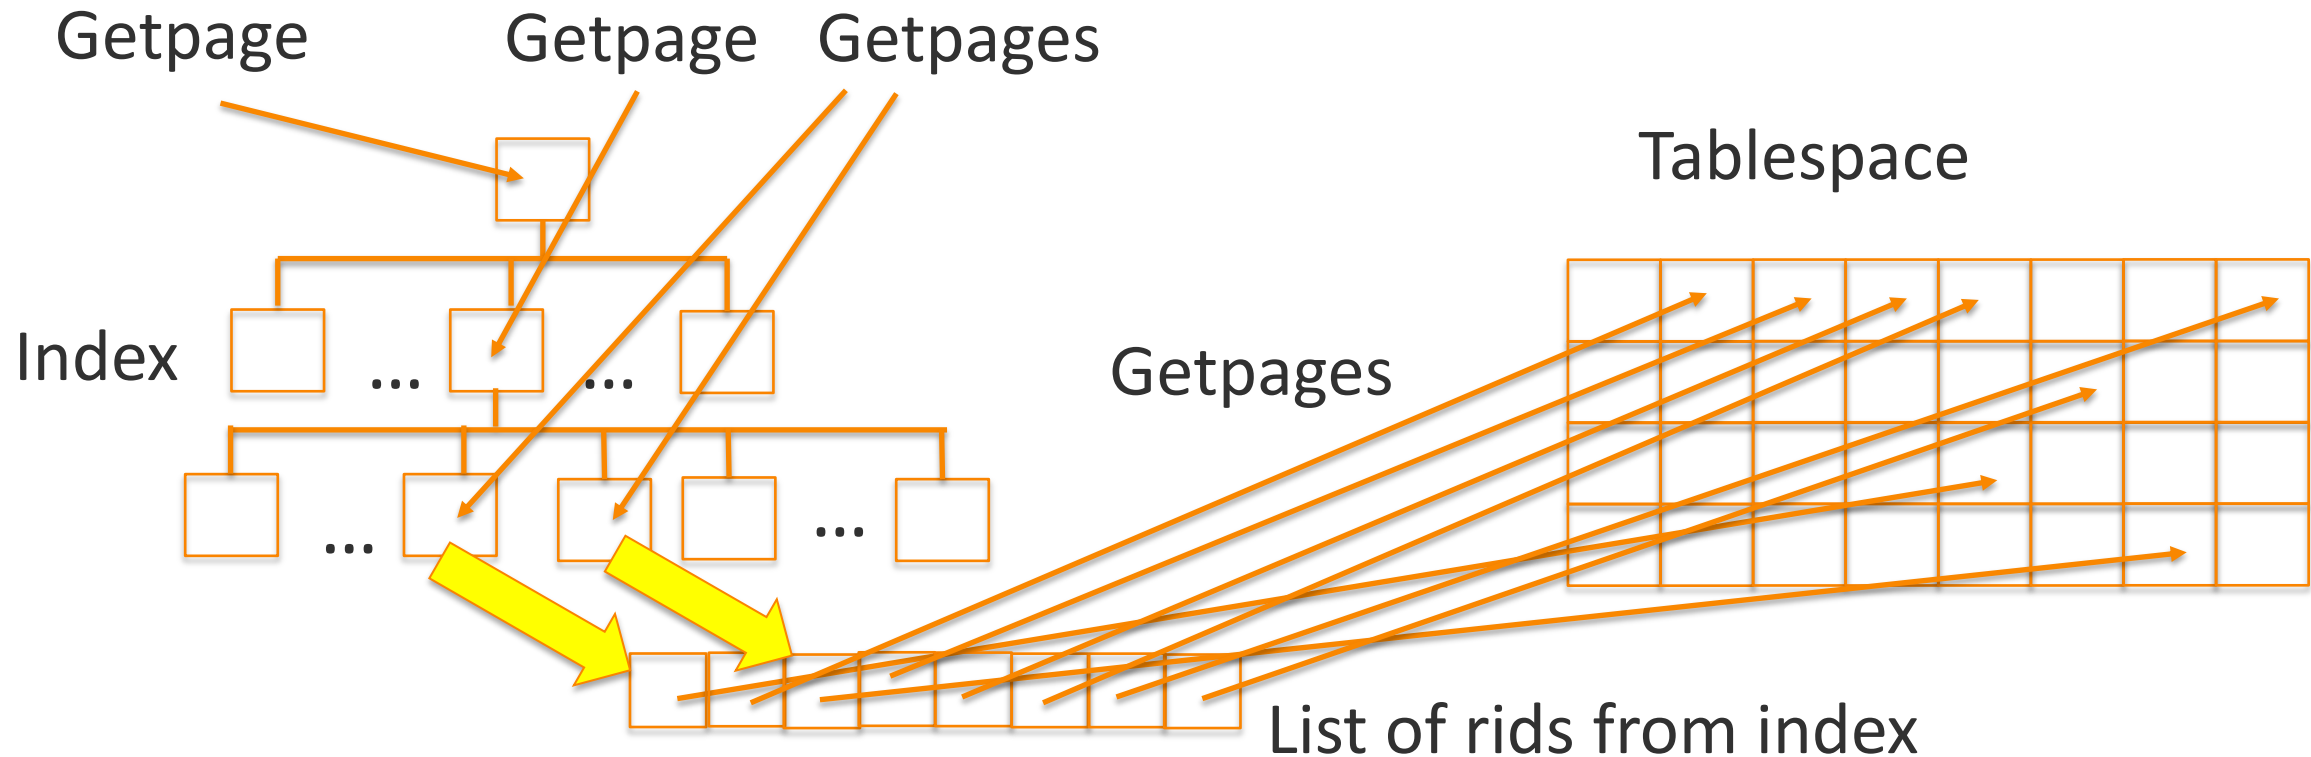


# List Prefetch – More Detail



# Dynamic Prefetch

```
SELECT ORDERNO, ORDERDATE, AMOUNT FROM ORDER
WHERE CUSTID = :Customer-id
AND ORDERDATE BETWEEN :start and :end;
```





# Dynamic Prefetch

PLANNO	METHOD	ACCESSTYPE	MATCHCOLS	PREFETCH	CREATOR	TNAME	TABNO
1	0	I	0	S	SYSIBM	SYSCOLUM	1

# Single Row Retrieval

```
SELECT LASTNAME, SSN FROM CUSTOMER
WHERE CUSTID = :Customer-id;
```



# Single Row Retrieval

PLANNO	METHOD	ACCESSTYPE	MATCHCOLS	PREFETCH	CREATOR	TNAME	TABNO
1	0	I	3		SYSIBM	SYSCOLUM	1

PREDNO	TYPE	LHS	RHS	FILTER_FACTOR	TEXT
1	AND			+0.100000000000000000E+01	(( "SYSIBM"."SYS
2	EQUAL	TBCREATO	VALUE	+0.2034902200102806E-01	"SYSIBM"."SYSCO
3	EQUAL	TBNAME	VALUE	+0.1644736767048016E-03	"SYSIBM"."SYSCO
4	EQUAL	NAME	VALUE	+0.9412650251761079E-04	"SYSIBM"."SYSCO
5	COMPOUND			+0.8438107033725828E-04	COMPOUND: 2, 3
6	COMPOUND			+0.5818994395667687E-05	COMPOUND: 2, 3, 4

# Nested Loop Join


Outer  
table

	8	
	2	
	4	
	2	
	5	
	9	
	3	
	8	
	4	

Join columns

Inner  
table

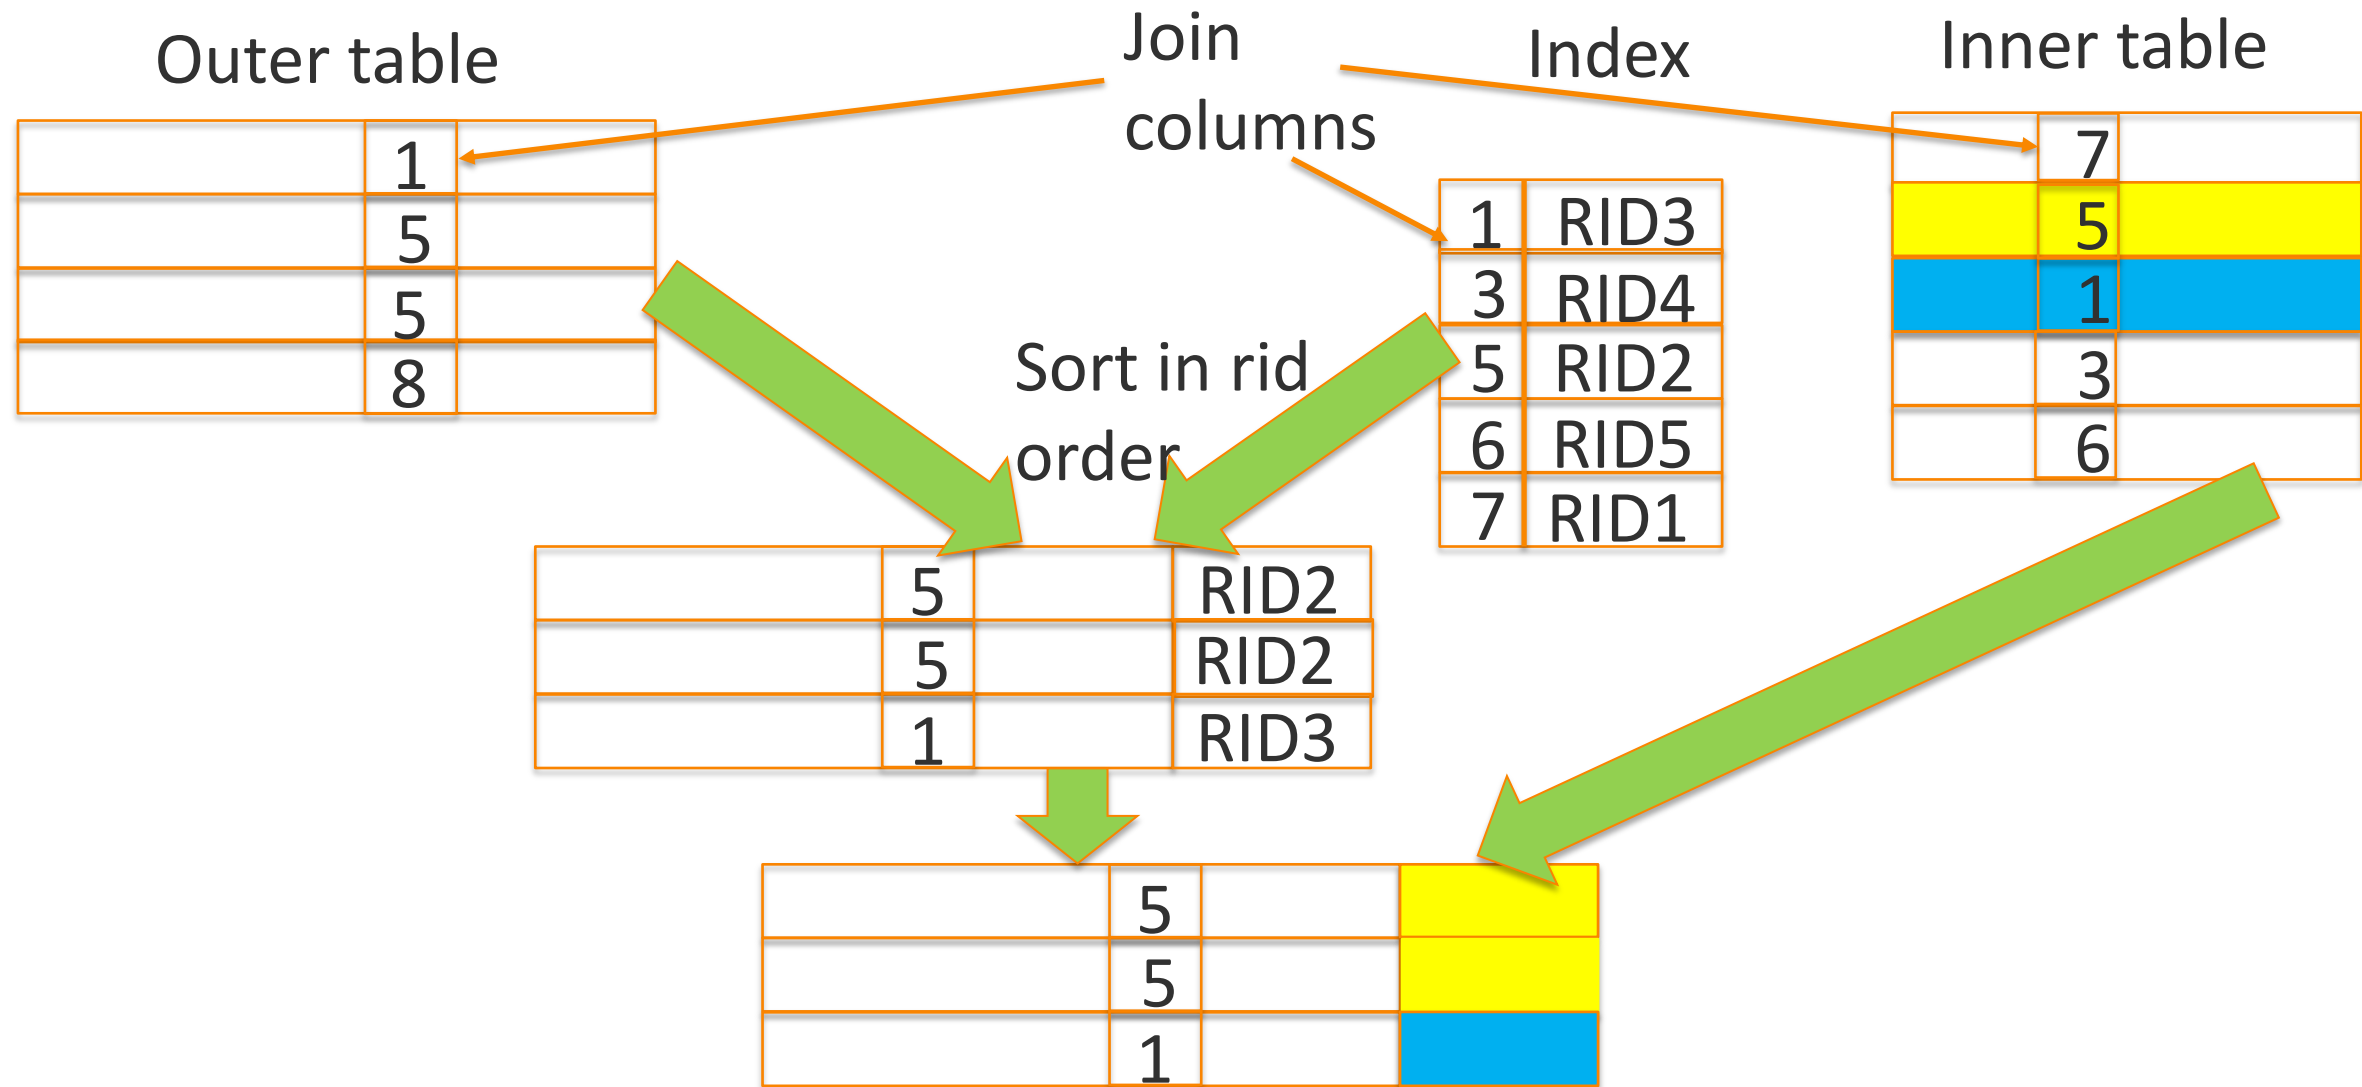
	8	
	1	
	2	
	1	
	5	
	9	
	6	
	7	
	5	
	7	
	4	
	3	



# Nested Loop Join

PLANNO	METHOD	ACCESSTYPE	MATCHCOLS	PREFETCH	CREATOR	TNAME	TABNO
1	0	I	2		SYSIBM	SYSINDEX	1
2	1	I	2		SYSIBM	SYSKEYS	2
PREDNO	TYPE	LHS	RHS	FILTER_FACTOR		TEXT	
4	EQUAL	NAME	IXNAME	+0.1420454500475898E-03		"IX"."NAME"="KY"."IXNAME"	
5	EQUAL	CREATOR	IXCREATO	+0.1592356478795409E-02		"IX"."CREATOR"="KY"."IXCR	
15	COMPOUND			+0.6462451710831374E-04		COMPOUND: 4, 5	

# Hybrid Join





# Hybrid Join

PLANNO	METHOD	ACCESSTYPE	MATCHCOLS	PREFETCH	CREATOR	TNAME	TABNO
1	0	I	1	L	SJDID001	CUSTOMER	1
2	4	I	1	L	SJDID001	ORDER	2

1. Scan the index on CUSTOMER (CUST\_COUNTRY, CUST\_ADDRESS) for all entries for one value of COUNTRY.
2. Sort the RIDS from (1).
3. List prefetch the rows from CUSTOMER. As each row is retrieved, read the index entries in ORDER (CUSTID, ORDNO) for all entries for the value of CUSTID from the CUSTOMER row.
4. Sort the workfile built in (3) in RID order.
5. List prefetch the rows from ORDER, retrieving the other requested columns.

# Merge Scan Join

Outer  
table

	8	
	1	
	5	
	5	
	3	
	1	
	2	
	6	

Select and sort

	1	
	3	
	6	


Join columns

	7	
	5	
	1	
	3	
	6	

Inner  
table

Select and sort

	3	
	5	
	6	

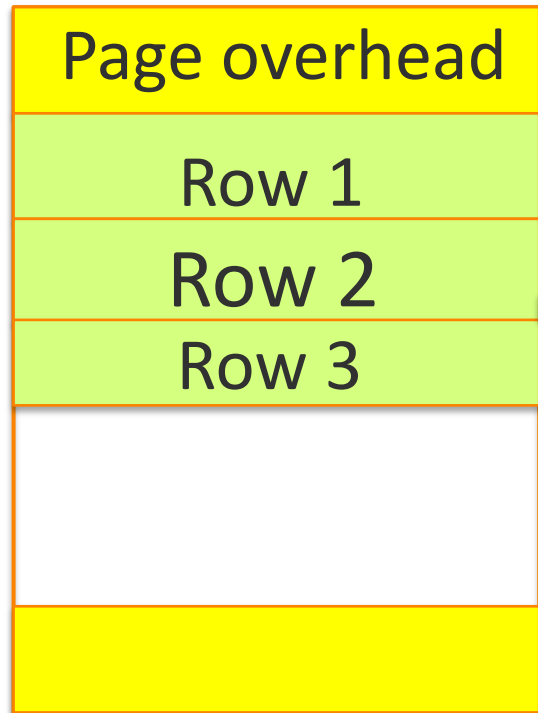


# Merge Scan Join


PLANNO		METHOD		TABNO		ACCESSTYPE		MATCHCOLS		PREFETCH	
1		0		1		I		0		S	
2		2		2		R		0		S	

# Indirection


Tablespace page



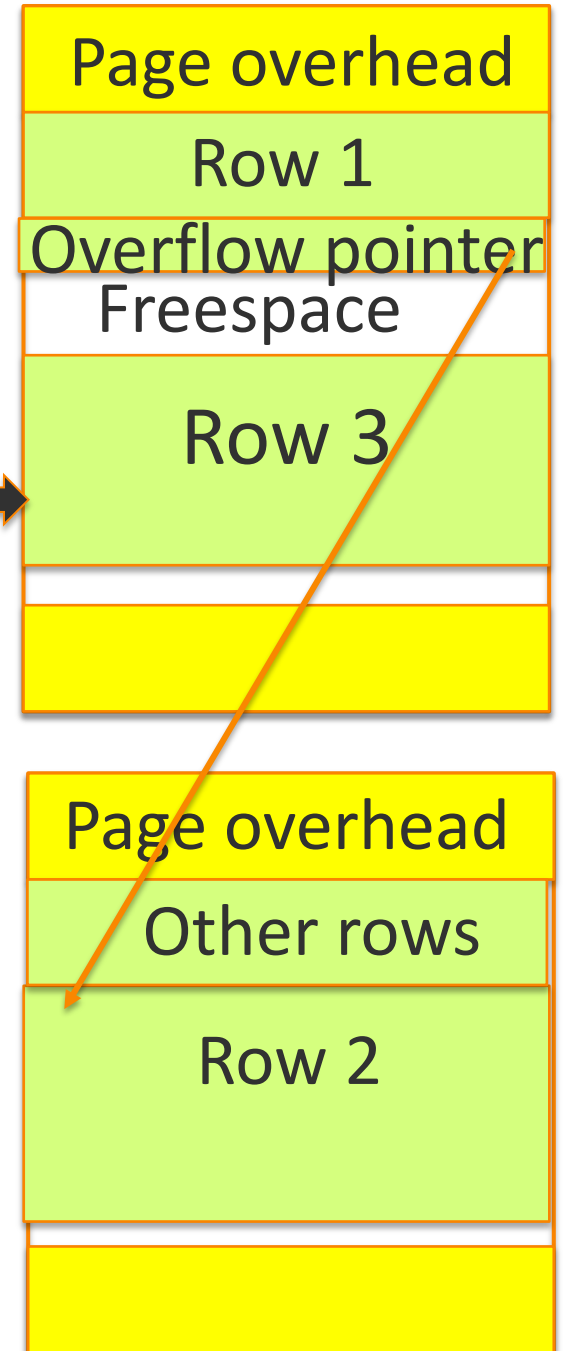
Row 2 grows



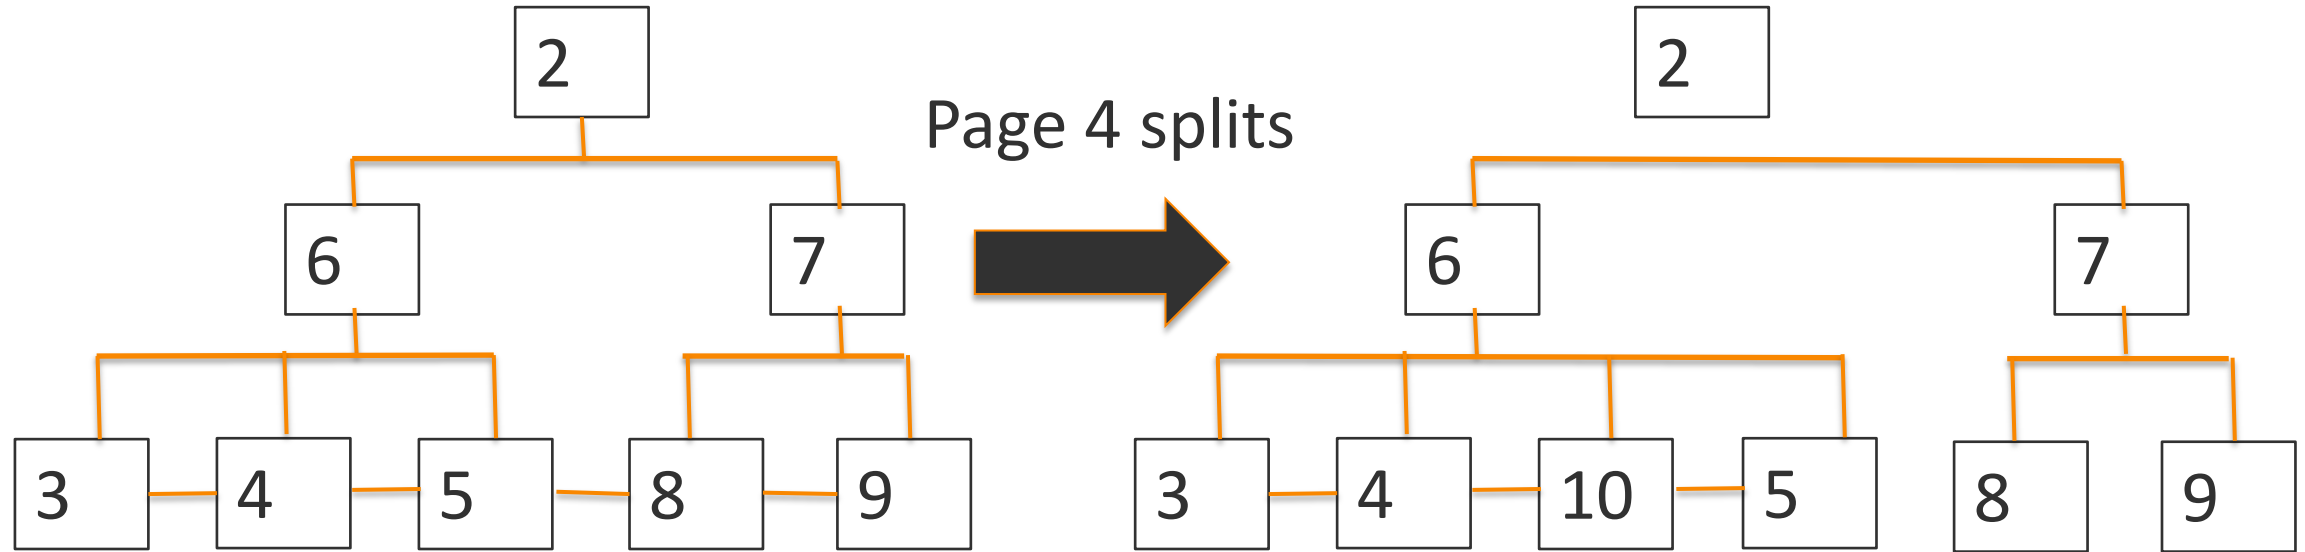
Row 3 grows



Row 2 grows again



# Index Structure and Updates

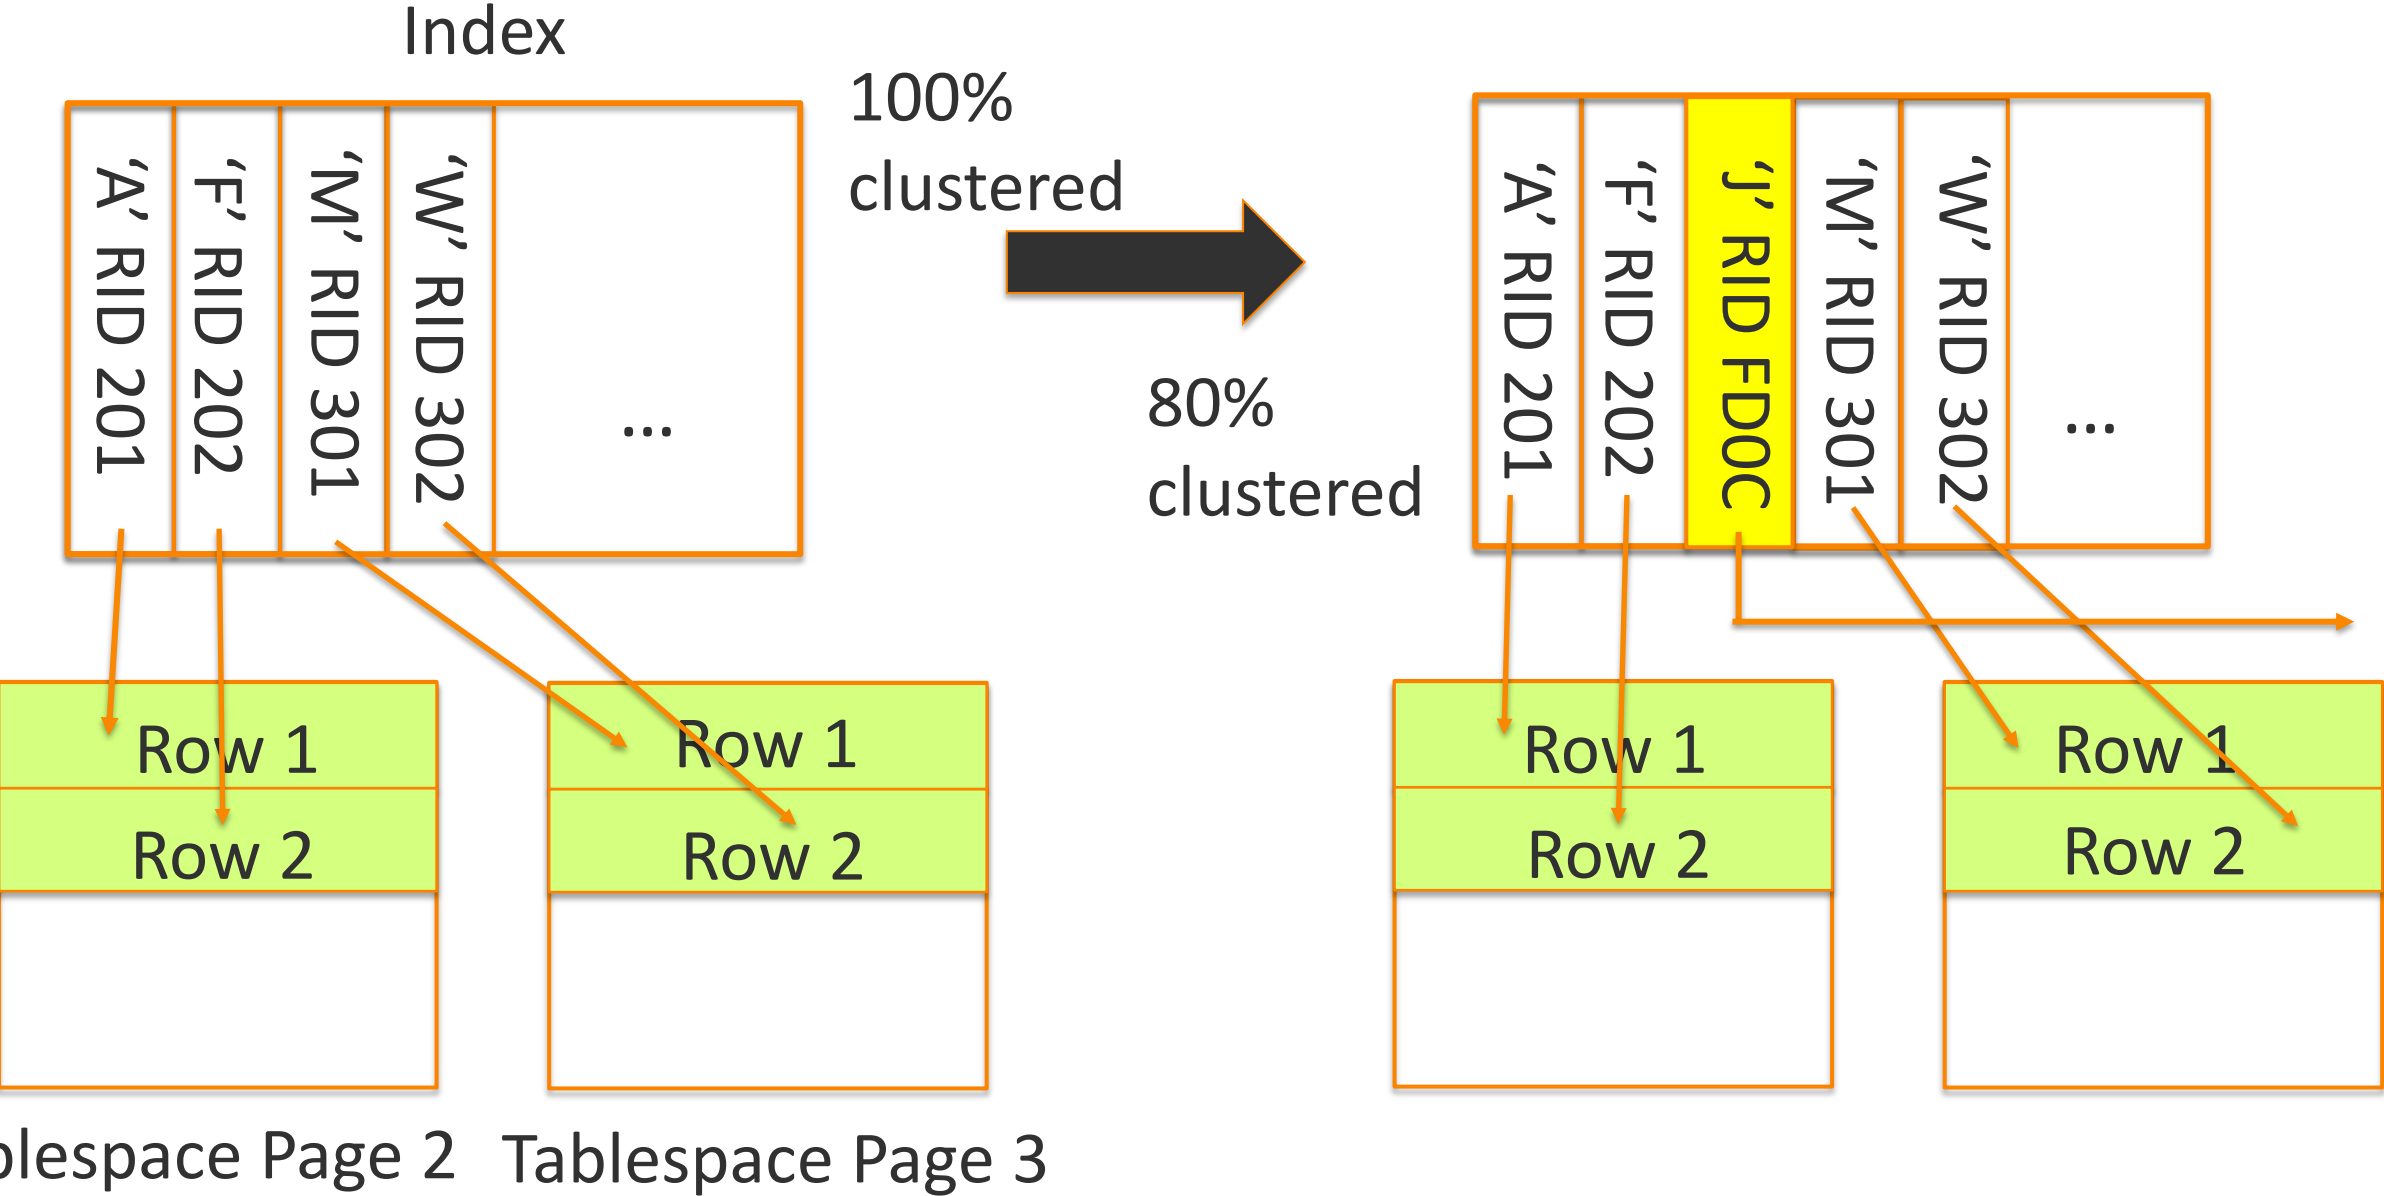


LEAFDIST = 150

$$((4-3) + (5-4) + (8-5) + (9-8) * 100) / 4$$

LEAFDIST = 320

# Clustering



# Thank You

---

**Bring IT to Life.™**