# Cache Fusion: Demystified Arup Nanda

Longtime Oracle DBA

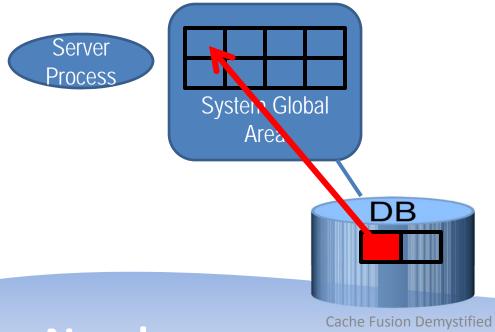
Blog: arup.blogspot.com Twitter: @ArupNanda

# Why this Session?

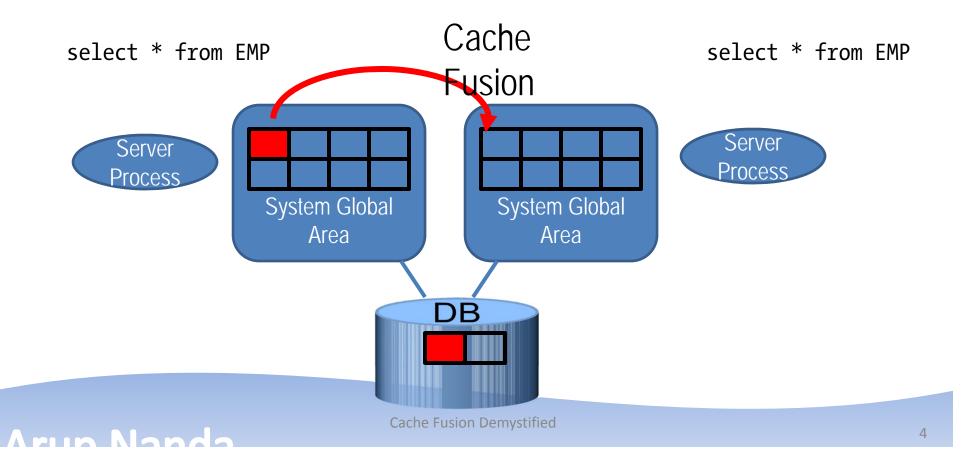
- How does Cache Fusion know where to get the block from?
- How are block locks vary from row locks?
- I'm confused about Global Cache Service (GCS), Global Resource Directory (GRD) and Global Enqueue Service (GES)
- We will understand how all these actually work under the covers.

### **Buffer Cache**

select \* from EMP



#### RAC - More than 1 Buffer Cache

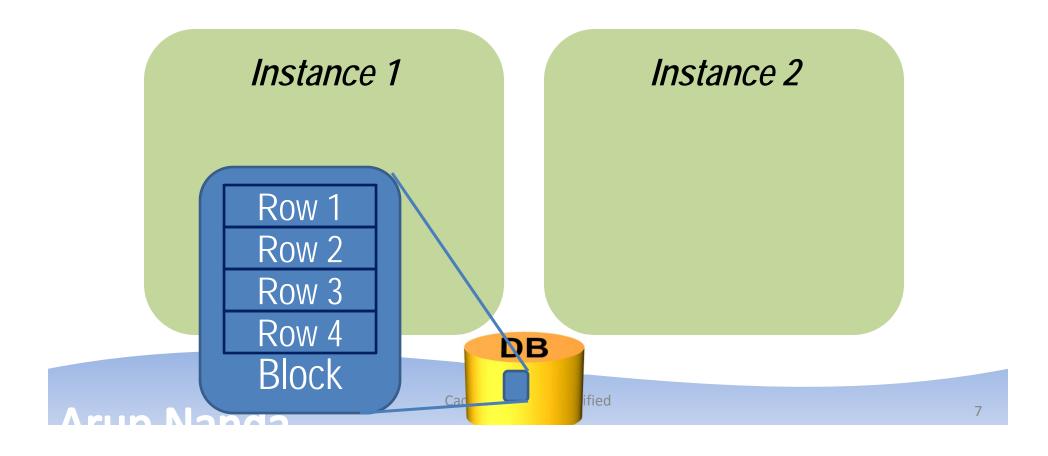


#### To Cache Fusion or Not?

- When a block is requested, the buffer cache is searched
- If not found, there are two options
  - Get from disk
  - Get from the other cache
- If found, there are three options:
  - Send the buffer to the user
  - Examine other caches for the presence of this buffer
  - Get from the disk
- How does it decide which option to take?

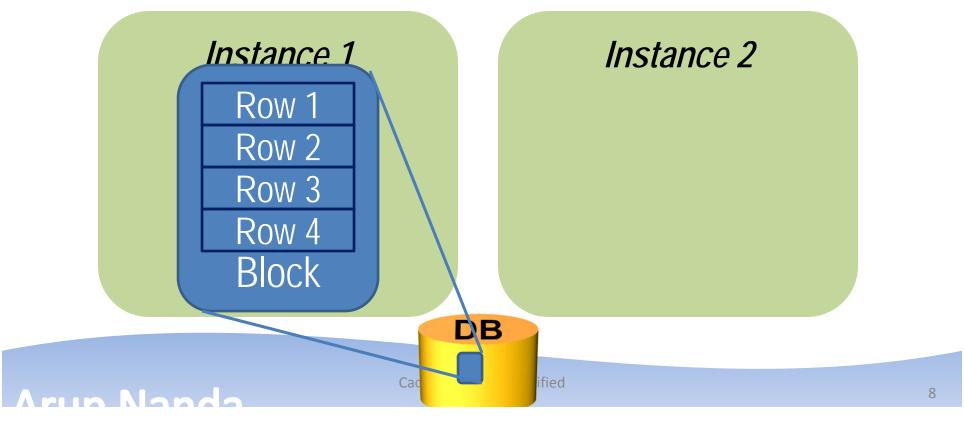
# Poly Symm Transience

# Block - Row Relationship

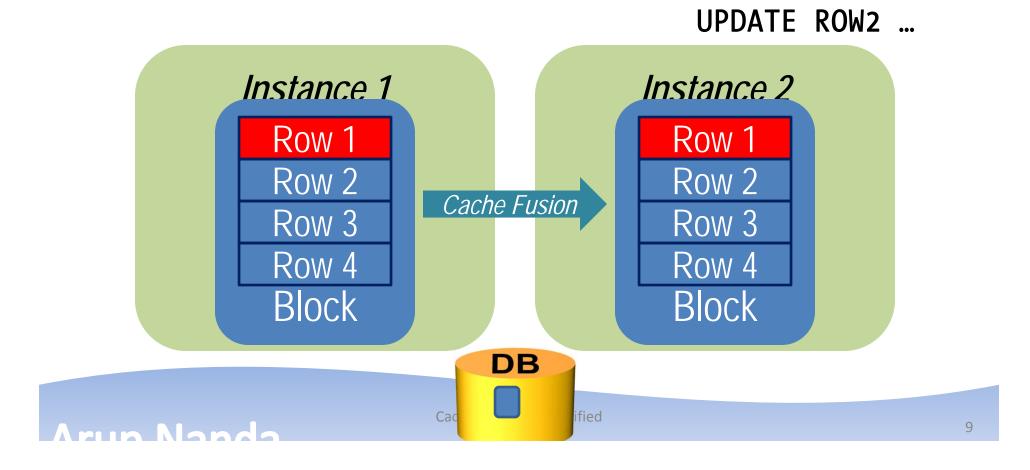


# Update on One Instance

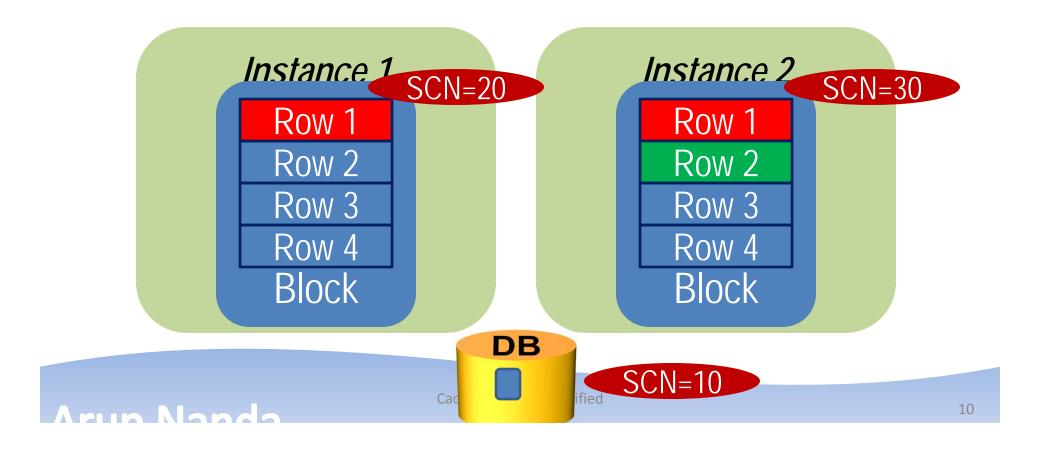
#### UPDATE ROW1 ...



# Update a Different Row on Node 2

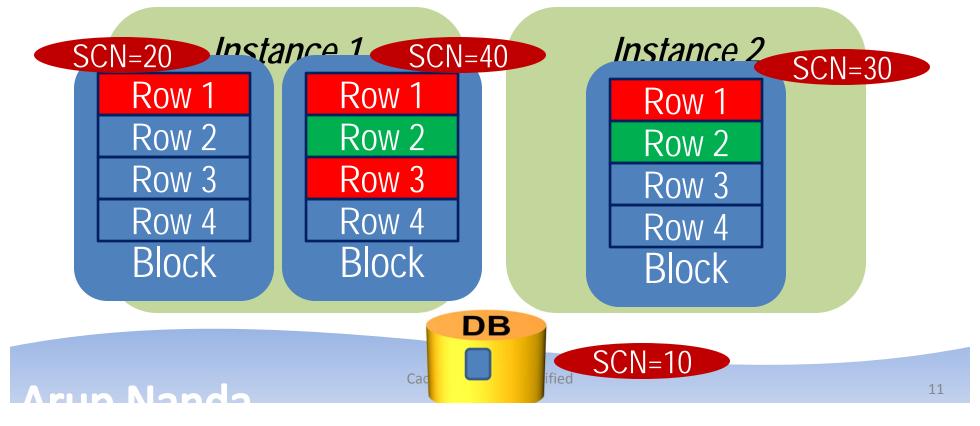


#### **Buffer Versions**



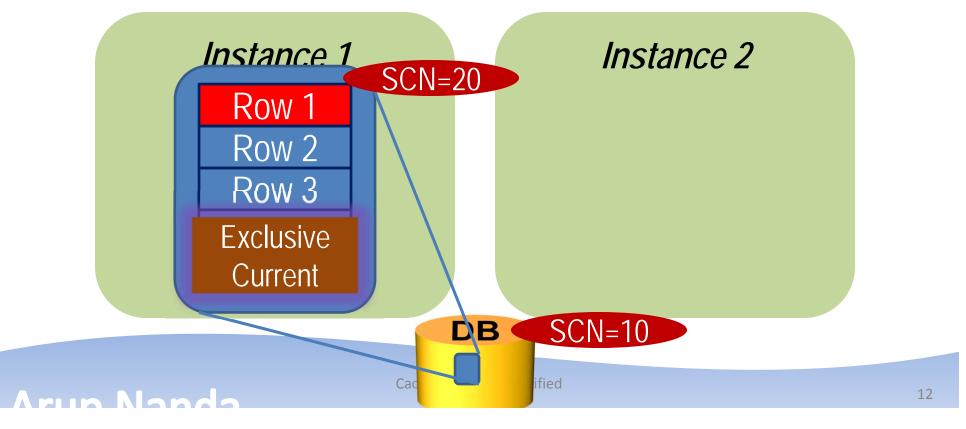
#### **Buffer Versions**

#### UPDATE ROW3 ...

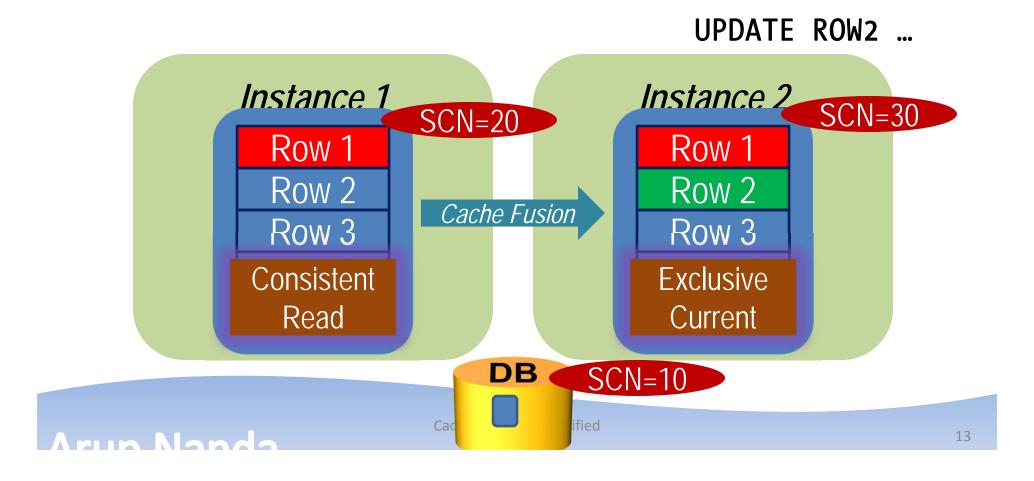


#### **Buffer State 1**

#### UPDATE ROW1 ...

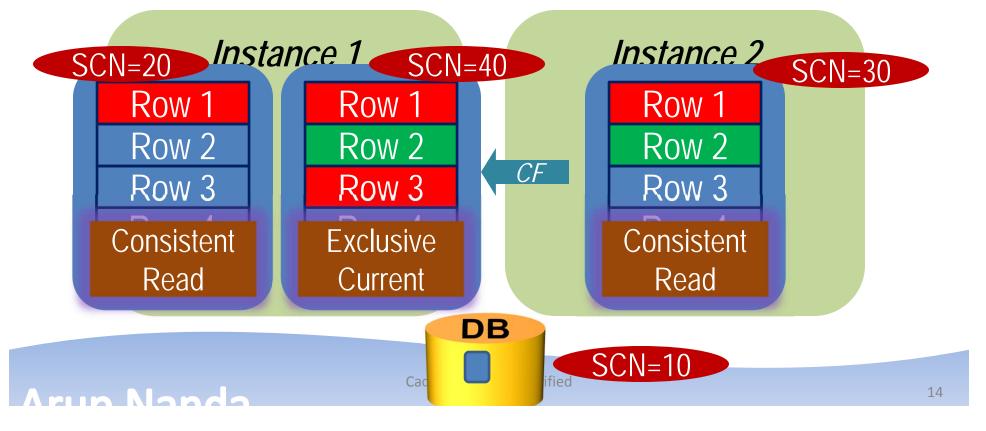


## Update a Different Row on Node 2



#### **Buffer Versions**

#### UPDATE ROW3 ...

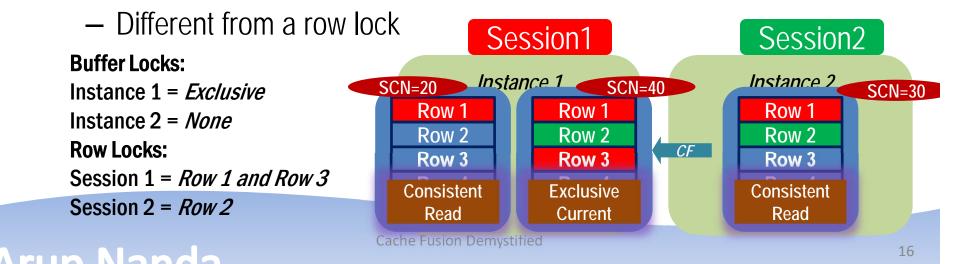


#### **Buffer States**

- The buffer can be retrieved in two modes
  - Consistent Read (CR)
  - Current
- There can be several CR copies of a buffer
   Upto 7. http://arup.blogspot.com/2011/04/can-i-fit-80mb-database-completely-in.html
- There can be only one current mode
  - For an instance
- Each current buffer is Shared Current
- Only one buffer in the entire cluster can be Exclusive Current

#### **Buffer Lock**

- When an instance wants to change the state of the buffer from CR to Exclusive Current
  - It must get a lock on that buffer
  - This is called a Buffer Lock



#### Global Cache Service

- Provides buffer from one instance to the other
  - But does not know who has what type of buffer lock

# Global Enqueue Service

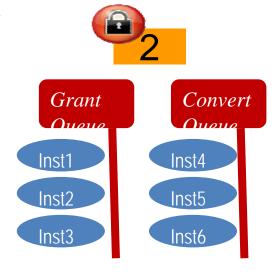
- Used to be called Dynamic lock Manager (DLM)
- Holds the information on the locks on the buffers.
- Each lock has a name shown in V\$LOCK\_ELEMENT (or X\$LE)
- This is different from row locking, which is on a specific row
- If a buffer is locked, the lock element name is shown in V\$BH.LOCK\_ELEMENT

# **Lock Queuing**

- Each Buffer in a RAC instance has two queues
  - Grant Oueue

the queue where the requesters are queued for the locks to be granted in a certain mode

- Convert Queue
   the queue where the granted requests are queued to be notified to the requesters
- The queues for a specific buffer are placed in a single instance



#### Master Instance

- The SGA where the queues of a specific buffer are located
  - A Buffer has only one Master Instance
- The Master may change
  - Manually
  - By a process known as Dynamic Resource Mastering
- When an instance wants to get a lock on that buffer, it has to update the queue in the SGA of the master instance

# **Global Resource Directory**

- A list master instances of all buffers
- GRD is present on all the instances of the cluster
- To find out the master:

```
select b.dbablk, r.kjblmaster master_node
from x$le l, x$kjbl r, x$bh b
where b.obj = <DataObjectId>
and b.le_addr = l.le_addr
and l.le_kjbl = r.kjbllockp
```

# In Summary

- Buffers are gotten in 2 modes
  - CURRENT is need to be modified
  - CR if selected only for reading
- Every time other node wants the buffer
  - it is copied to a new buffer and sent (CR processing)
- There can be only one current state of the buffer in an instance in Shared Mode
- Only one Exclusive Current in the Cluster
- Each buffer has a master node that holds the lock Grant and Convert Queues
- GRD maintains information on the buffers' masters

# Thank You!

Blog: arup.blogspot.com Facebook.com/ArupKNanda Tweeter: @ArupNanda