Information Management

IBM

DB2 pureScale Feature Overview

Kelly Schlamb
WW IM Technical Sales Acceleration, DB2 pureScale Specialist
kschlamb@ca.ibm.com

November 22, 2012

DB2 pureScale

- **Unlimited Capacity**
  - Buy only what you need, add capacity as your needs grow
- **Application Transparency**
  - Avoid the risk and cost of application changes
- **Continuous Availability**
  - Deliver uninterrupted access to your data with consistent performance

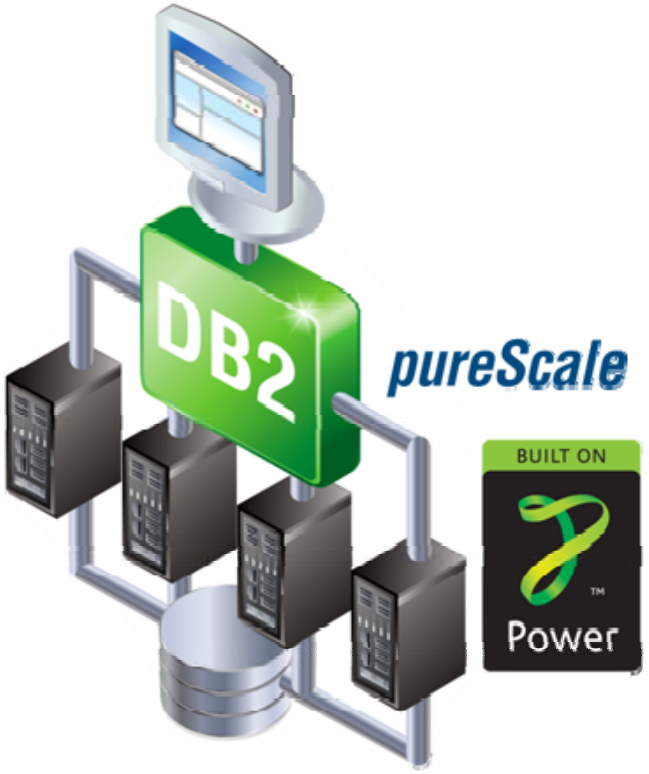

Learning from the undisputed Gold Standard... System z

DB2 pureScale Architecture

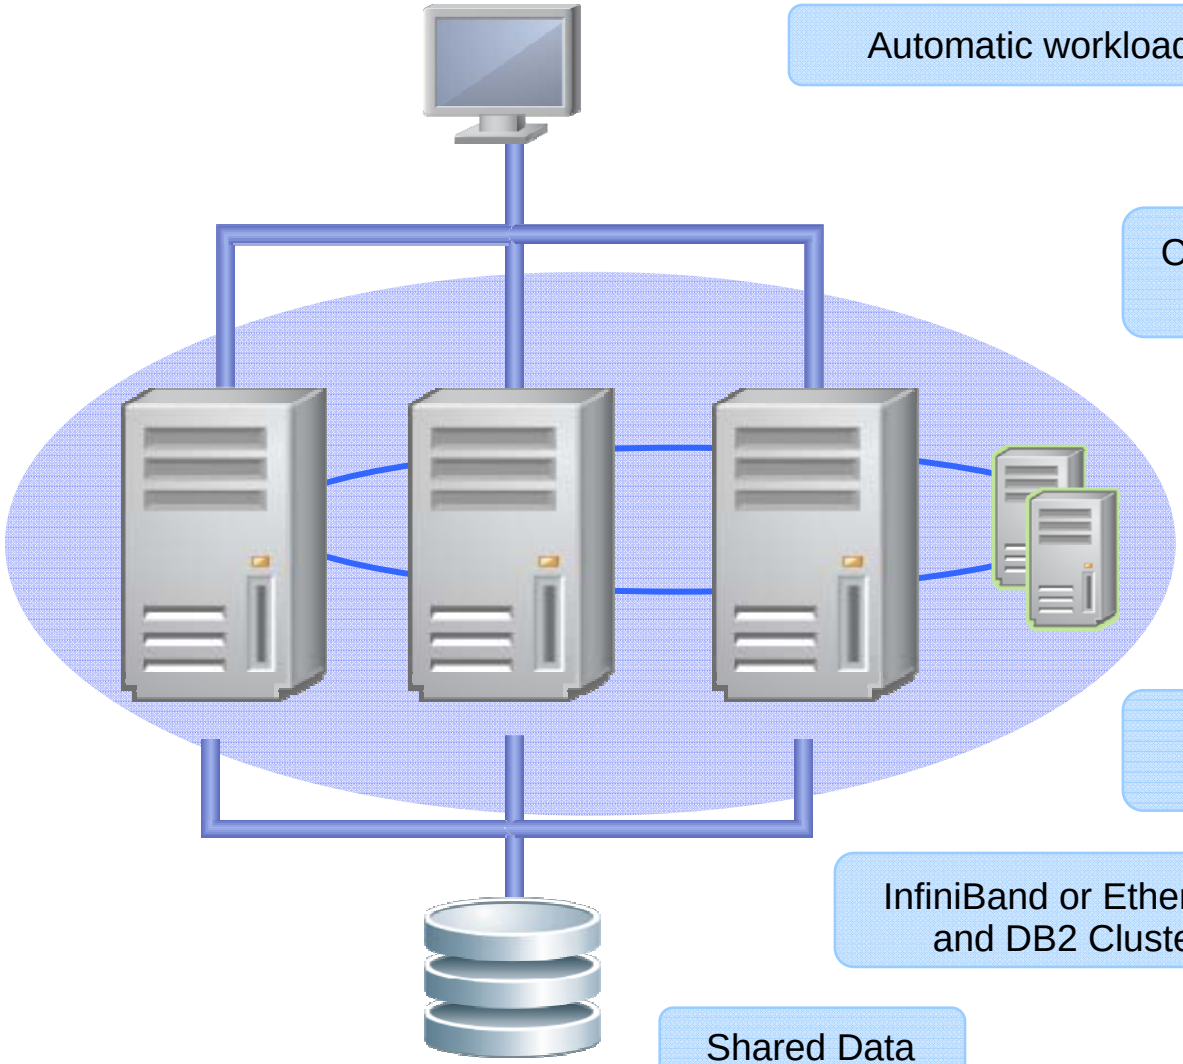

Automatic workload balancing

Cluster of DB2 nodes running on Power or Intel servers

Leverages the global lock and memory manager technology from z/OS

Integrated Cluster Manager

InfiniBand or Ethernet Networks and DB2 Cluster Services

Shared Data

DB2 pureScale : Technology Overview

Leverage IBM's System z Sysplex Experience and Know-How

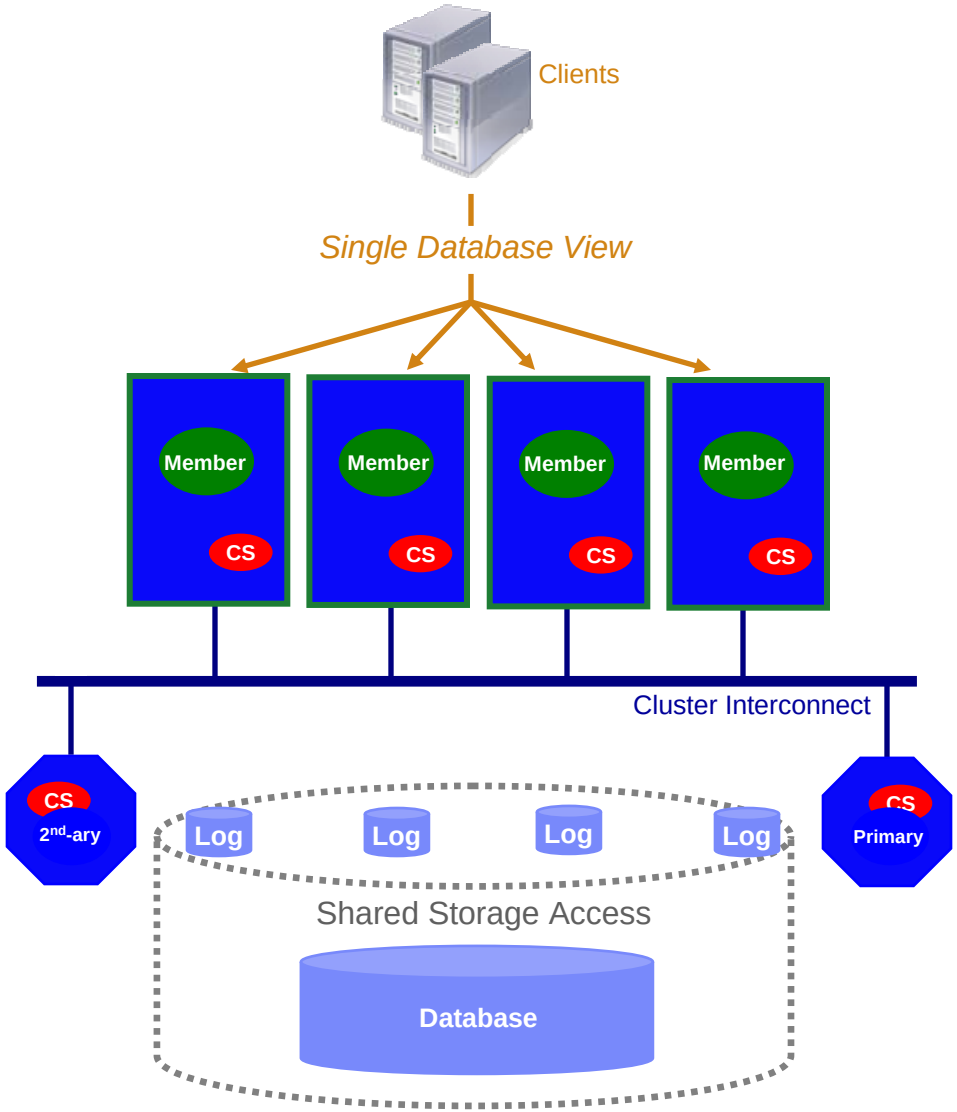

Clients connect anywhere,… … see single database

- Clients connect into any member
- Automatic load balancing and client reroute

DB2 engine runs on several hosts

- Co-operate with each other to provide coherent access to the database from any member

Integrated Cluster Services

- Failure detection, recovery automation

Low latency, high speed interconnect

- Special optimizations provide significant advantages on RDMA-capable interconnects (eg. Infiniband)

Cluster Caching Facility (CF)

- Efficient global locking and buffer management
- Synchronous duplexing to secondary for availability

Data sharing architecture

- Shared access to database
- Members write to their own logs
- Logs accessible from another host

Application Transparency

Take advantage of extra capacity instantly

- No need to modify your application code
- No need to tune your database infrastructure

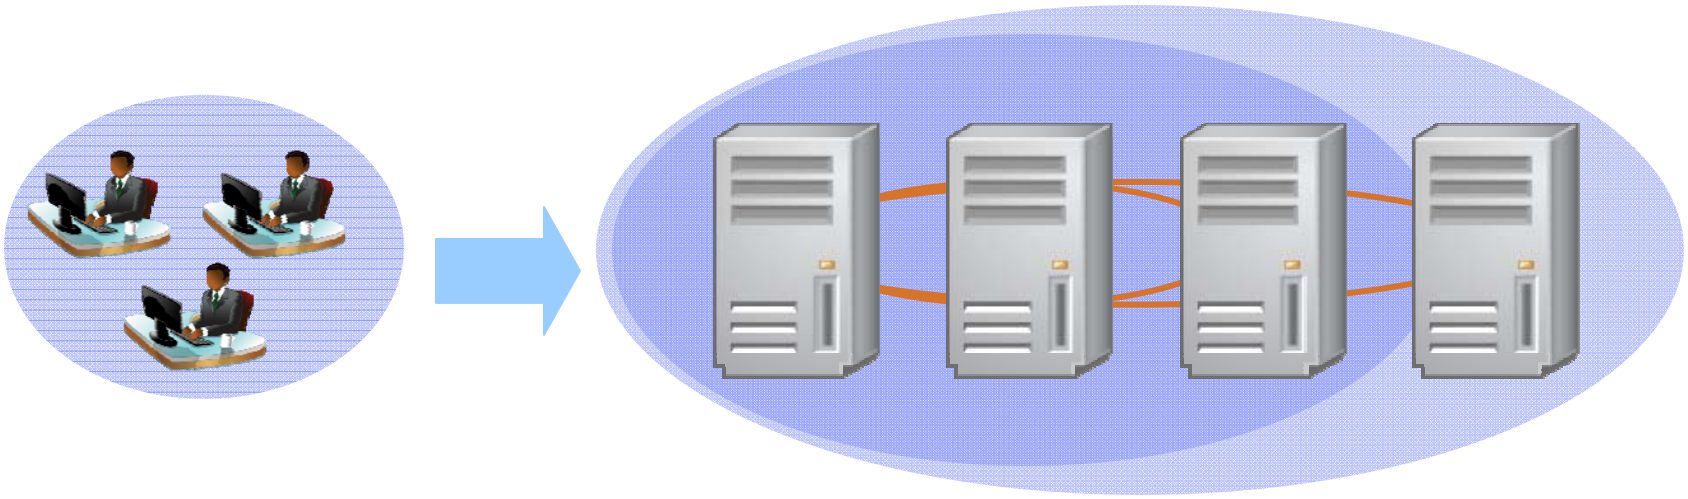

Your DBAs can add capacity without re-tuning or re-testing

Your developers don't even need to know more nodes are being added

The Key to Scalability and High Availability

- **Efficient Centralized Locking and Caching**
  - As the cluster grows, DB2 maintains one place to go for locking information and shared pages
  - Optimized for very high speed access
    - DB2 pureScale uses Remote Direct Memory Access (RDMA) to communicate with the powerHA pureScale server
    - No IP socket calls, no interrupts, no context switching

- **Results**
  - Near Linear Scalability to large numbers of servers
  - Constant awareness of what each member is doing
    - If one member fails, no need to block I/O from other members
    - Recovery runs at memory speeds

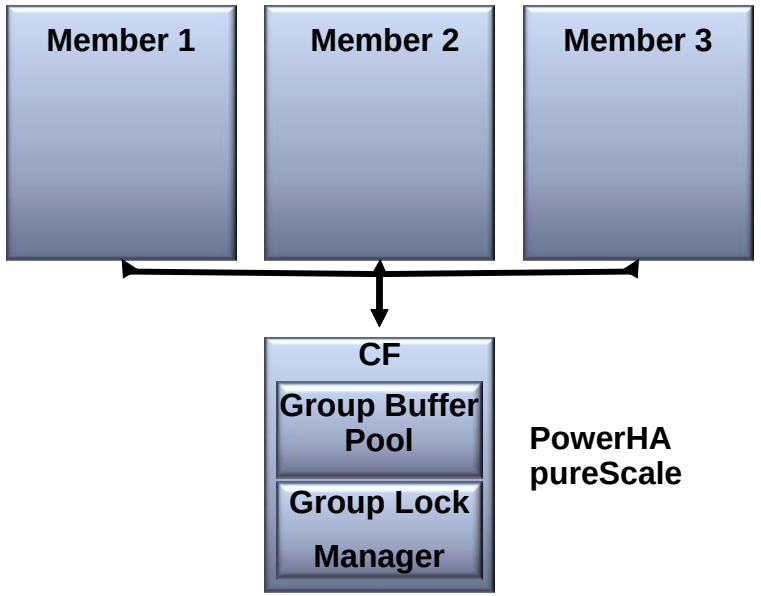

Online Recovery

- **DB2 pureScale design point is to maximize availability during failure recovery processing**
- **When a database member fails, only *in-flight* data remains locked until member recovery completes**
  - In-flight = data being updated on the failed member at the time it failed
- **Target time to row availability**
  - <20 seconds

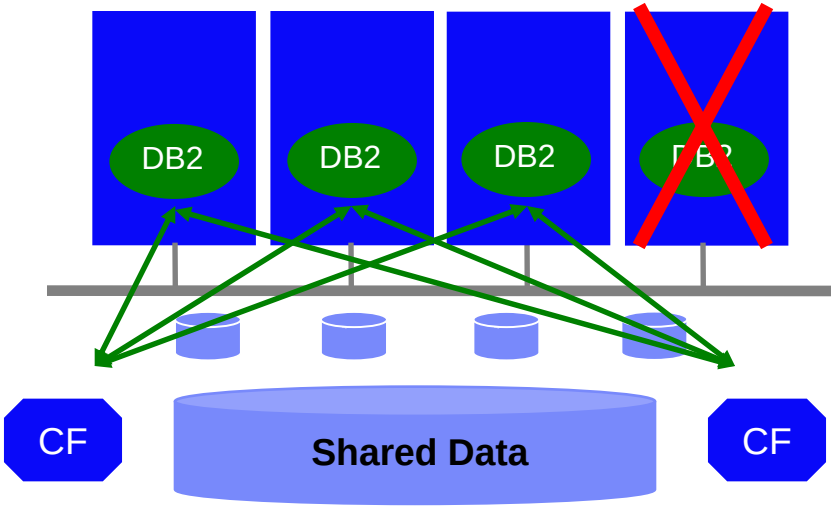

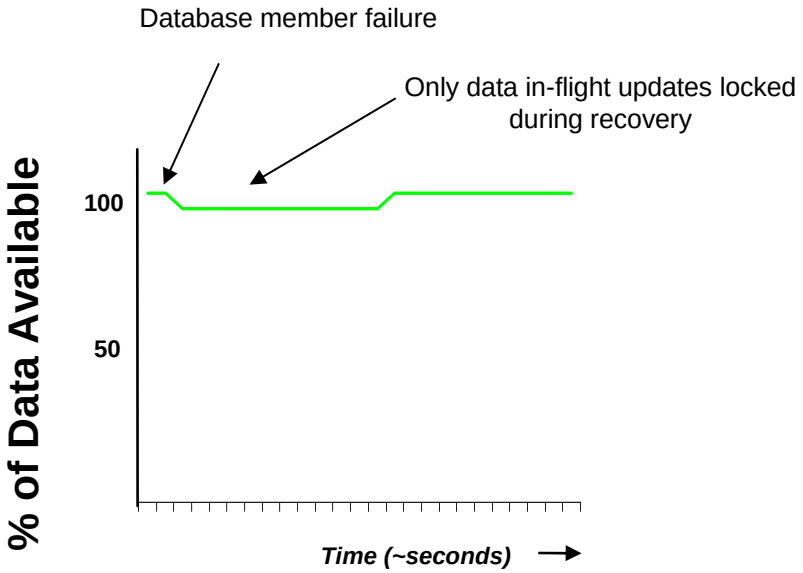

Competitive Advantages

- **Ease of Use Advantage**
  - Entire stack has integrated install, deployment and maintenance
  - DB2 Data Sharing automatically installs, configures, and patches
    - Database engine
    - Integrated cluster manager and availability monitor
    - Integrated shared disk file system
  - No need to partition your database or application to scale

- **Availability Advantage**
  - Centralized locking and a true global caching results in higher availability in the event of software or hardware failures

- **Scalability Advantage**
  - DB2 pureScale uses unique technology to deliver near linear scaling without the need to partition the application or database
    - No need for applications to be cluster aware in order to scale to dozens of members

Centralized vs. Distributed Lock And Cache Management

DB2 pureScale

Centralized CF Design

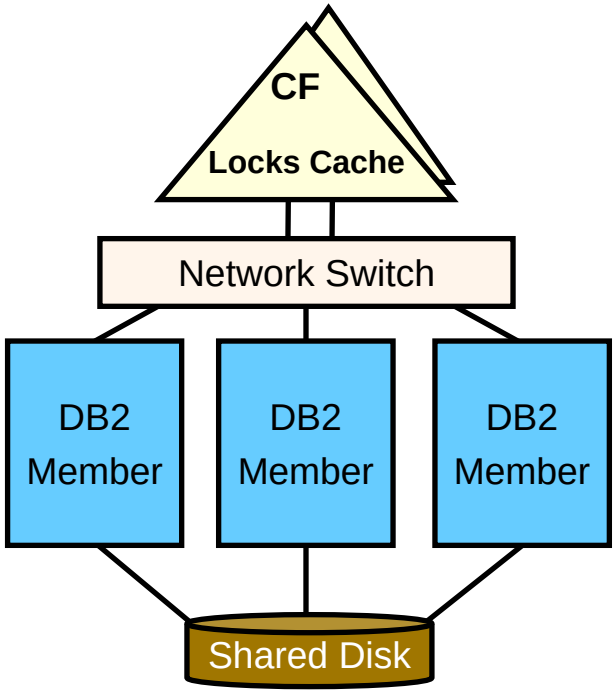

Centralized lock manager and cache in CF

Oracle RAC

Distributed Design

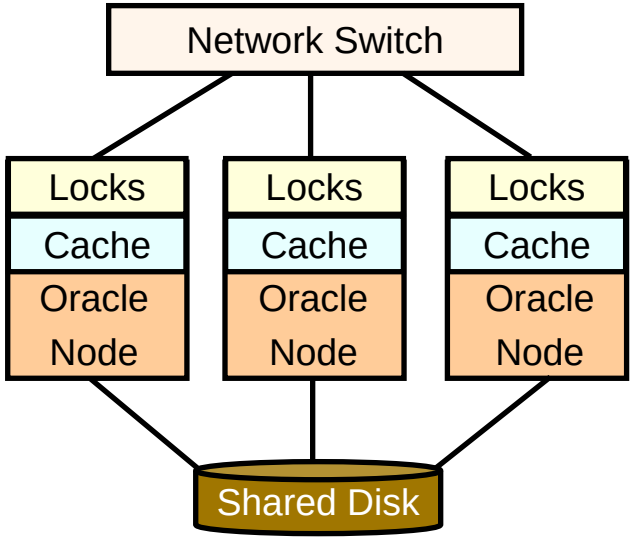

Distributed lock and cache management in each node

DB2 pureScale – No Freeze at All

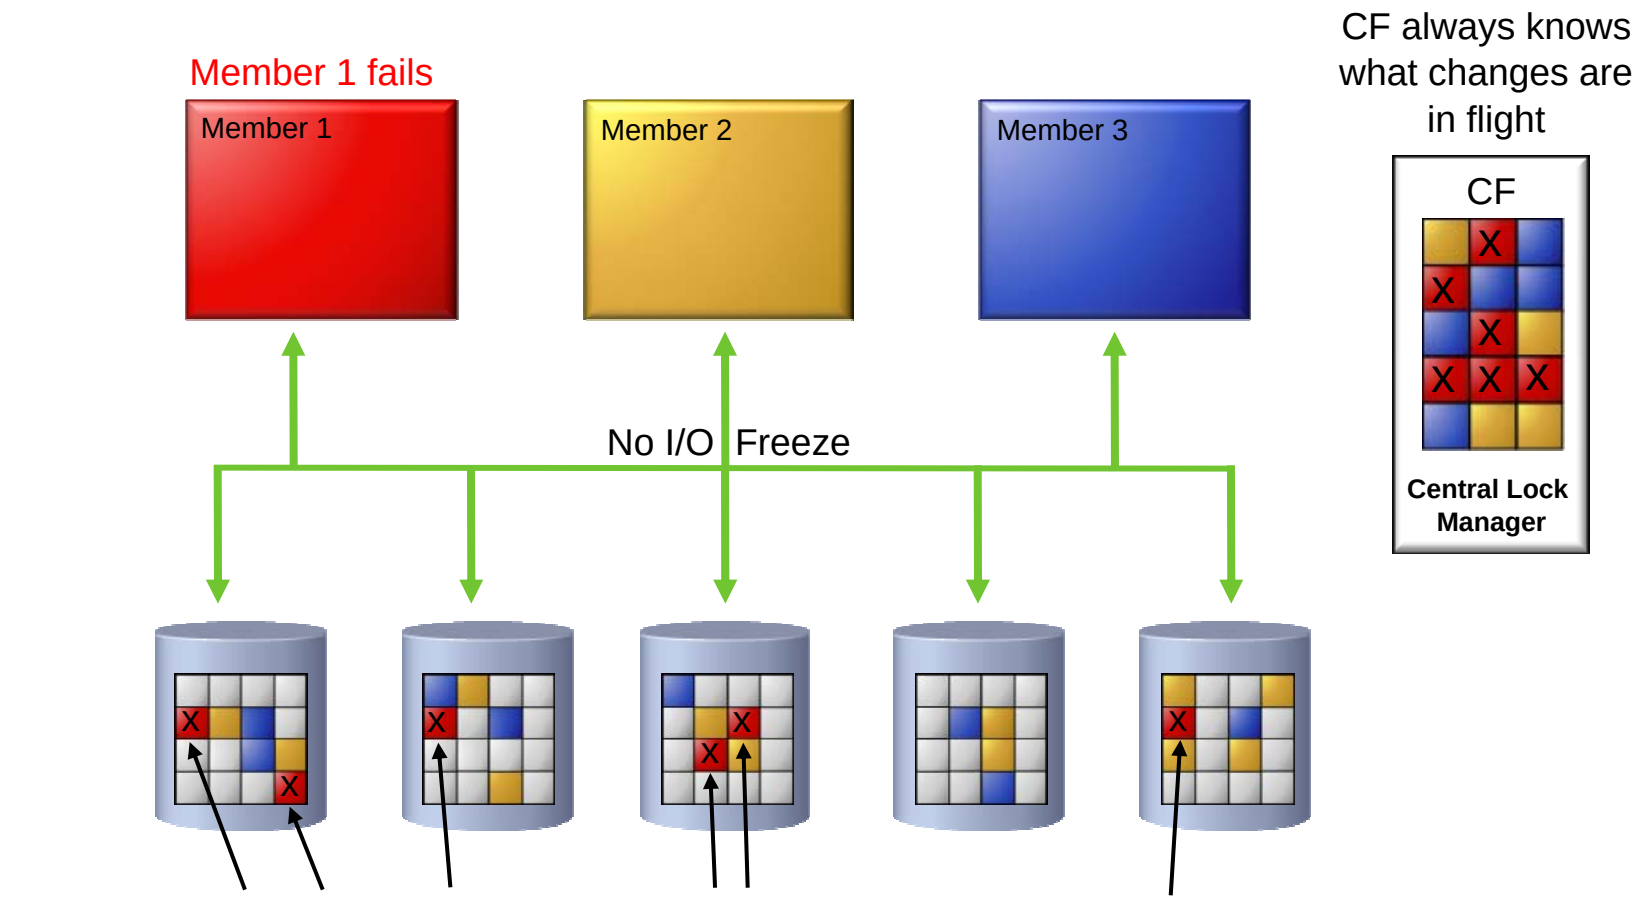

With RAC – Access to GRD and Disks are Frozen

- **Global Resource Directory (GRD) Redistribution**

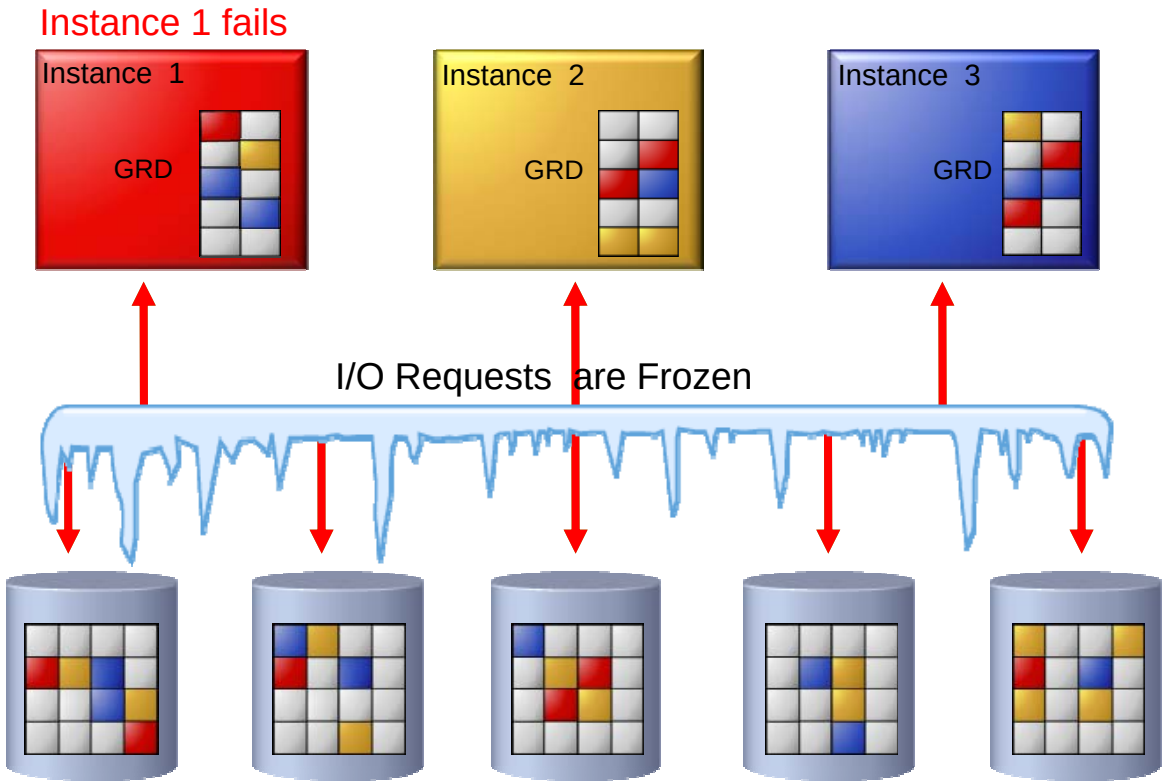

No more I/O until pages that need recovery are locked

No Lock Updates

With RAC – Pages that Need Recovery are Locked

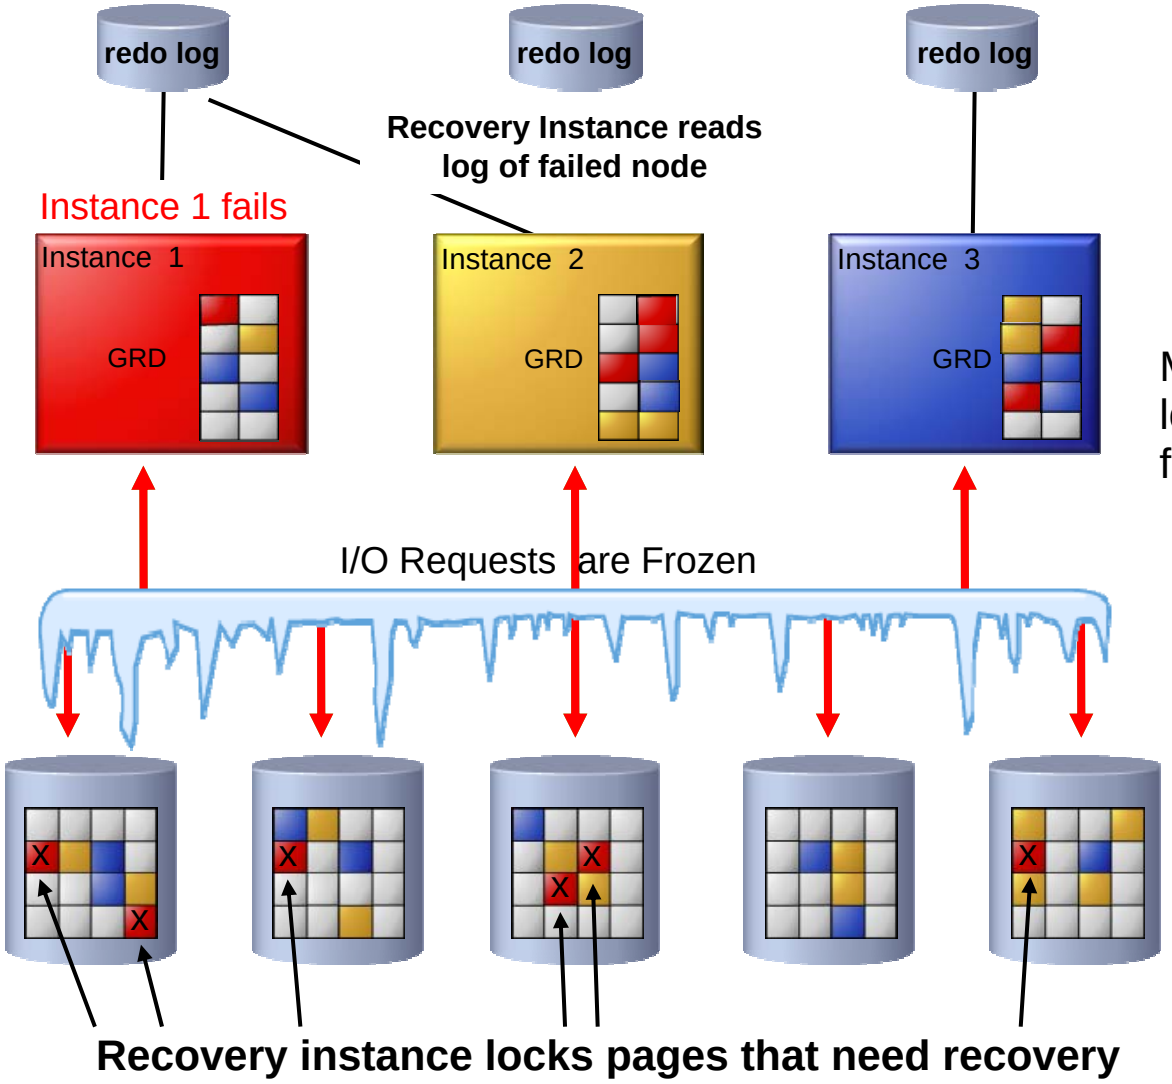

Must read log and lock pages before freeze is lifted.

Scalability Differences

- Distributed locking requires that you lock a page whenever there is the intent to update that page
- DB2 pureScale must lock a page whenever rows are actually being changed on that page
- DB2 pureScale improves concurrency between members in a cluster which results in better scalability and less of a need for locality of data

Summary – What can DB2 pureScale Do For You?

- Deliver higher levels of scalability and superior availability
- Better concurrency during regular operations
- Better availability during member failure
- Result in less application design and rework for scalability
- Improved SLA attainment