Clustering with Oracle (Real Application Cluster)

Montse Collados Polidura,
Eric Grancher
IT-DB

DataBase Workshop January 29-31, 2002

Clustering with Oracle

- Shared disk (see later) infrastructure, all disk devices have to be accessible from **all** servers

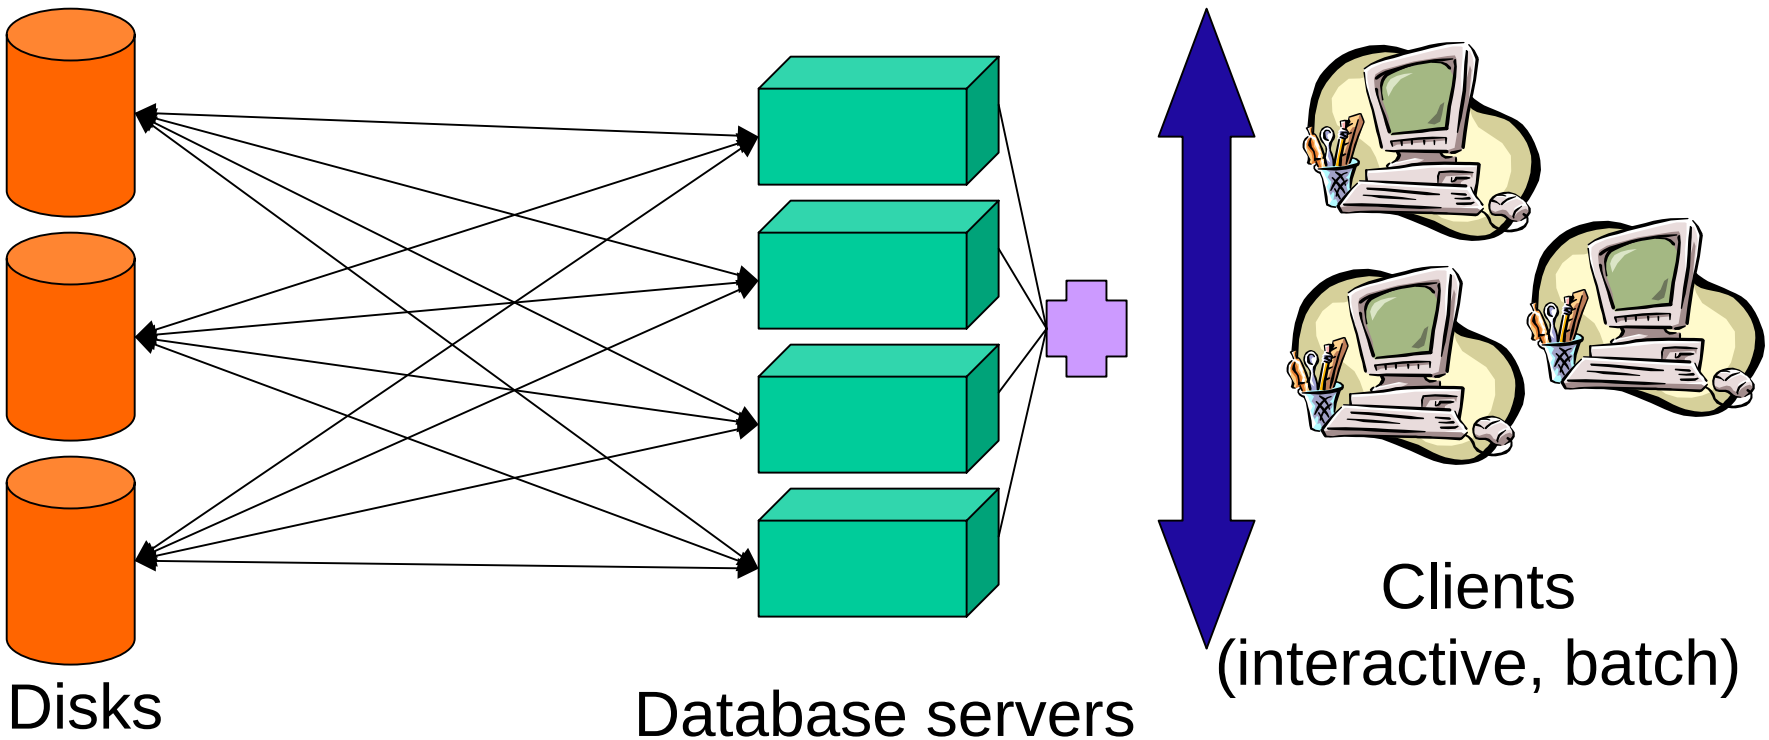

A clustered database

- It can be accessed from all of the nodes
- Functionality is the same as a database "served" by a single database (same statements –or OCCI calls –, transaction...)
- The nodes are using the interconnect (network between the nodes) to transfer the cached data

Pros

- High availability: one HW resource (CPU, power, OS corruption...) may fail
- Scalability (more power to achieve the task)
- Have several hosts working on the same data -> eases the data management
- Cost facility (can be used with several small machines instead of a large expensive one) -> commodity HW ?

Cons

- Cache coherency, implies design and configuration issues
- Complexity of the administration
- Cost of software (Enterprise Edition + RAC option)
- HW requirements (certified HW, today shared disk is not commodity)

History and experience

- Long clustering history at Oracle (since version 7, hasn't been a world wide success)
- CERN is using the product since 1996, mostly for the highly availability feature
- Major new releases Oracle8i OPS and Oracle9i RAC (shared cache)

Usage at CERN

- (Since 1996)
- 2 nodes “central service” (Sun, 4 CPU, 2.5GB memory, Sun shared disk, small amount of data), used mostly as a high available solution
- ~1000 accounts, ~ 200 concurrent
- We have gained experience from the day-to-day operation and the migrations

Tests on Linux

- (Autumn 2001, Montse C. P., Catherine D., Dirk G., George S.)
- Servers: 9 dual PIII CPUs, 512MB
- Storage: use 1 node that serves data to the others (instead of disks accessible from all nodes).
- Suse 7.2, Oracle 9.0.1

Tests on Linux

- Goal:
  - ✓ Test that it works with “cheap HW” + Linux
  - ✓ better understand the configuration issues
  - Check how it scales
    - Number of nodes
    - Network interconnect
    - CPU used for the cache coherency
    - Identify bottlenecks
- To be done
  - Work with faster + more storage

Direction

- Oracle seems to be highly committed to it (simpler, more efficient), new features 9iR2 (2002 Q2?)
- Commodity servers -> ok, Linux
- Commodity storage -> open issue (InfiniBand?)
- The scalability depends on the application / design, to be tested with a “typical” application
- Larger scale tests