


Part V – Applications

Cloud Computing: Evaluation

Assessment of cloud computing offerings

- How well do Amazon, Google, and Microsoft's cloud computing offerings scale with an increasing workload? Can indeed a (virtually) infinite throughput be achieved?
- How expensive are these cloud computing offerings and how does their cost / performance ratio (i.e., bang for the buck) compare?
- How predictable is the cost with regard to changes in the workload?

Distributed database architectures: Classic






Distributed database architectures: Classic

- Classical DB architecture
 - Advantages
 - Allows the use of "best-of-breed" components at all layers
 - Allows scalability and elasticity at the storage and web/application server layers
 - Limitation
 - Database server is a potential bottleneck. If it is overloaded, the only solution is to buy a bigger machine, which has to be provisioned for peak workloads.



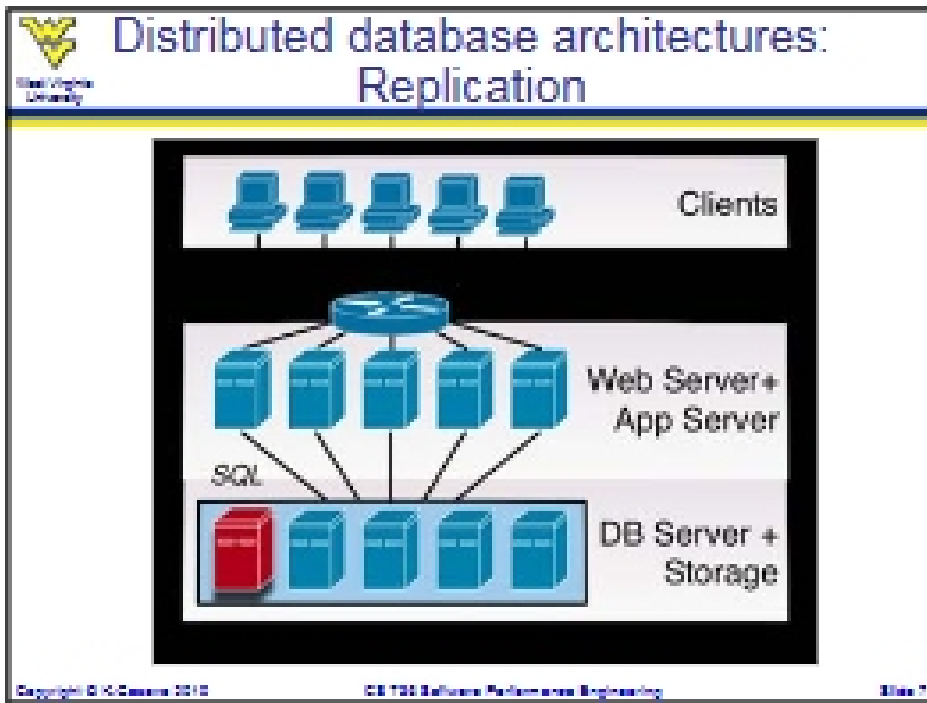
Distributed database architectures: Partitioning





Distributed database architectures: Partitioning

- The database is logically partitioned and each partition is controlled by a separate DB server
 - Partitioning can be transparent or visible to the application programmer
 - Storage can be attached to machines that run the DB servers or can be dissociated in SAN
- Limitation
 - Scalability with regard to dealing with a fluctuating workload
 - Adding (or removing) machines in order to deal with higher (or lower) workload involves repartitioning the data and, therefore, moving data between machines



- West Virginia University
- ## Distributed database architectures: Replication
- Each database server controls a copy of the whole database (or partition of the database, if combined with partitioning)
 - Replication increases both scalability and reliability
 - It is important to keep the replicas consistent
 - The most prominent protocol is ROWA (read-once, write all) based on a Master copy
 - Master server propagates all committed updates to the satellites when these updates have been successfully committed
 - If the workload is read-mostly it scales-out and down nicely with the workload
 - Otherwise the Master can become a bottleneck
- Copyright © K. Casares 2010 CE 722 Software Performance Engineering Slide 8

