


The realization of virtual compute resources in a Utility Computing Data Center (UCDC) in the many core era

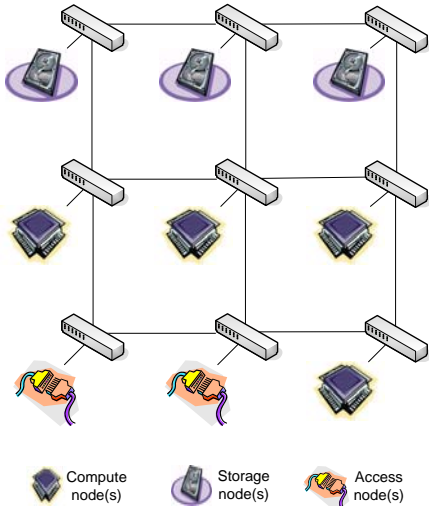
Many Core CPU

The development of CPUs that support the use of multiple cores trigger a paradigm shift of the traditional programming model. As the number of cores increases to hundreds, new architectural challenges emerge.



UCDC


In an abstract view a UCDC is a large collection of resources interconnected by an interconnection network where each resource falls into the category of compute node, storage node, or access node.



The jobs or services running in a UCDC typically have very diverse characteristics (different resource requirements, running time, quality of the software, security requirements, importance, etc). Virtual servers containing a subset of the available resources are dynamically created to fulfill user demands.

Flexible Partitioning

A typical job will come with a set of requirements relating to the number of compute nodes, storage nodes and access nodes it demands. This requires a partitioning of the network as a set of (possibly non-contiguous) partitions in order to maximize resource utilization.




If resource allocations must be of a rectangular shape, a request for 6 resource entities cannot be granted even if 6 entities are available (external fragmentation)

If the shape of the resource allocation may be irregular, a request for 6 resource entities can be granted (shown as a dashed line).

If irregular and overlapping resource allocations are allowed, a request for 7 resource entities may be granted (dashed line) if entity number 7 is lightly loaded or a multi-core processor etc.

Architectural Challenges for the Interconnection Network

Fault Tolerance



The effect of a faulty component in the interconnection network should be constrained to as few jobs as possible, and the set of jobs that are terminated by the fault should, to the greatest extent possible, be controlled from a specification of job importance. A solution to fault tolerance should allow unaffected jobs to run uninterrupted.

Predictable Service

Although in many cases it will be impractical to guarantee completely disjoint physical resources to different partitions, it should be possible to guarantee a specific portion of the interconnect network capacity to each partition or job. Further, it should be possible to differentiate between jobs based on importance. Three major topics that must be addressed are high-granularity service differentiation, congestion and admission control.

