Subject: Re: [PATCH 1/4] Virtualization/containers: introduction
Posted by Paul Jackson on Wed, 08 Feb 2006 04:21:37 GMT
View Forum Message <> Reply to Message

The driving force for cpusets are NUMA architectures.

Cpusets represent the topologies of NUMA systems, with hierarchies
of cabinets, drawers, boards, packages, cores, hyperthreads, and
with chunks of main memory associated usually with the board, but
sometimes a layer or two up or down.

Since not all cpus have the same access performance (delay and
bandwidth) to all memory chunks (nodes), for optimum performance one
wants to bind tasks, cpus and memory together, so as to run tasks on
sets of cpus and memory that are "near" to each other, and to size the
sets appropriately for the workload presented by the tasks.

Cpusets have no explicit awareness of topology; they just provides a
file system style hierarchy of named, permissioned sets, where each set
has:
  mems - the memory nodes in that set
  cpus - the cpus in that set
  tasks - the tasks running on these cpus and mems

For any cpuset, the 'cpus' and 'mems' are a subset of its parent in the
hierarchy, and the root of the hierarchy (usually mounted at /dev/cpuset)
contains all the online cpus and mems in the system.

--

I won't rest till it's the best ...
Programmer, Linux Scalability
Paul Jackson <pj@sgi.com> 1.925.600.0401