

CPU Hot-Plug

Status And Challenges

Andreas Färber

andreas.farber@suse.com

SUSE Linux GmbH



Outline

Overview Of This Presentation

- Motivation
- History and status
- Challenges
- Discussion

Motivation For CPU Hot-Plug

Why Invest In CPU Hot-Plug?

- Scalability
- Testing

History

CPU “Hot-Plug” In qemu-kvm

CPU Scaling By Ballooning

- vCPUs pre-allocated
- cpu-set HMP command
 - Sets them online/offline
 - Conceptually similar to virtio-balloon for memory
- Not working in last qemu-kvm versions
- Never available in qemu

Status

CPU Hot-Plug

First Real CPU Hot-Add

- vCPUs dynamically added
- `cpu-add QMP/HMP` command
 - Creates a new `X86CPUState` and vCPU dynamically
 - Allows to add *maxcpus* minus initial number of vCPUs
- Windows Server 2012R2+: CPU online right away
- Linux: Need to explicitly set CPU online via `/sys`
- Credits:
 - Andreas Färber – CPUState preparations
 - Igor Mammedov, Eduardo Habkost – x86/ACPI/QMP parts

Challenges

CPU Hot-Unplug

Reverting The Effects Of CPU Hot-Plug

- vCPUs dynamically removed
- Add `cpu-del` QMP/HMP commands?
 - How does argument relate to CPU indices? (sparse?)
- Or reuse `device-del` infrastructure?

- Credits:
 - Fujitsu et al.

From Custom QMP to QOM Device?

- Originally for cpu-set CPU was on Object
- Now CPU is a Device → device-add?
 - Not every machine a PC
 - Allow to add different CPUs
 - Where do we add them to? (topology)
 - Reuse existing infrastructure
 - Which granularity? (ABI stability...)
- Credits:
 - Andreas Färber – RFC QOM'ifying underneath cpu-add
 - Bharata Rao – CPU index rework

CPU Hot-Plug For Power Architecture

- pseries machine is a virtual machine (sPAPR)
- Conclusion: Hot-add on CPU core level
 - SMT a fixed property of core
- Credits:
 - Bharata Rao

CPU Hot-Plug For System z

- There are only virtual machines (PR/SM, LPAR)
- Hot-add on CPU core or on thread level?
- Dispatching required
- NUMA?
- Credits:
 - Jason J. Herne et al.

More?

.

qemu-project.org

Thank you.







Corporate Headquarters
Maxfeldstrasse 5
90409 Nuremberg
Germany

+49 911 740 53 0 (Worldwide)
www.suse.com

Join us on:
www.opensuse.org