



# This is your operating system; Let me help you

**Rolf Riesen**

**24 March 2016**

**Copyright © 2016 Intel Corporation. All rights reserved.**



# Legal Disclaimer



- Intro
- Roles
- Interfaces
- Details
- Layers
- Functions
- Flexibility
- mOS
- Summary

Statements and opinions in this talk are those of the author and do not reflect on Intel product plans.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

\*Other names and brands may be claimed as the property of others.

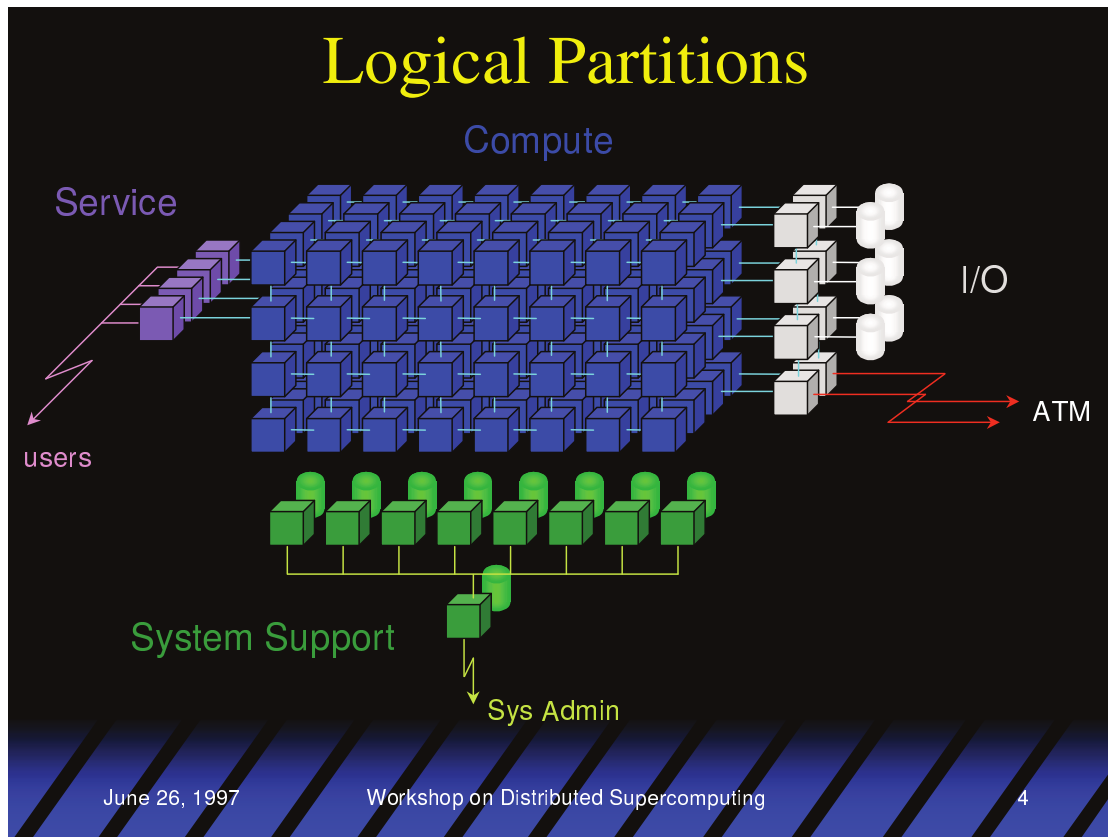
Copyright © 2016, Intel Corporation. All rights reserved.

# Introduction



- Showed this slide 19 years ago at SOS 1 in Santa Fe
- Partitioning and non-demand paged memory is still important
- Extracting a twenty-year old PowerPoint slide is not easy

- Intro
- Roles
- Interfaces
- Details
- Layers
- Functions
- Flexibility
- mOS
- Summary

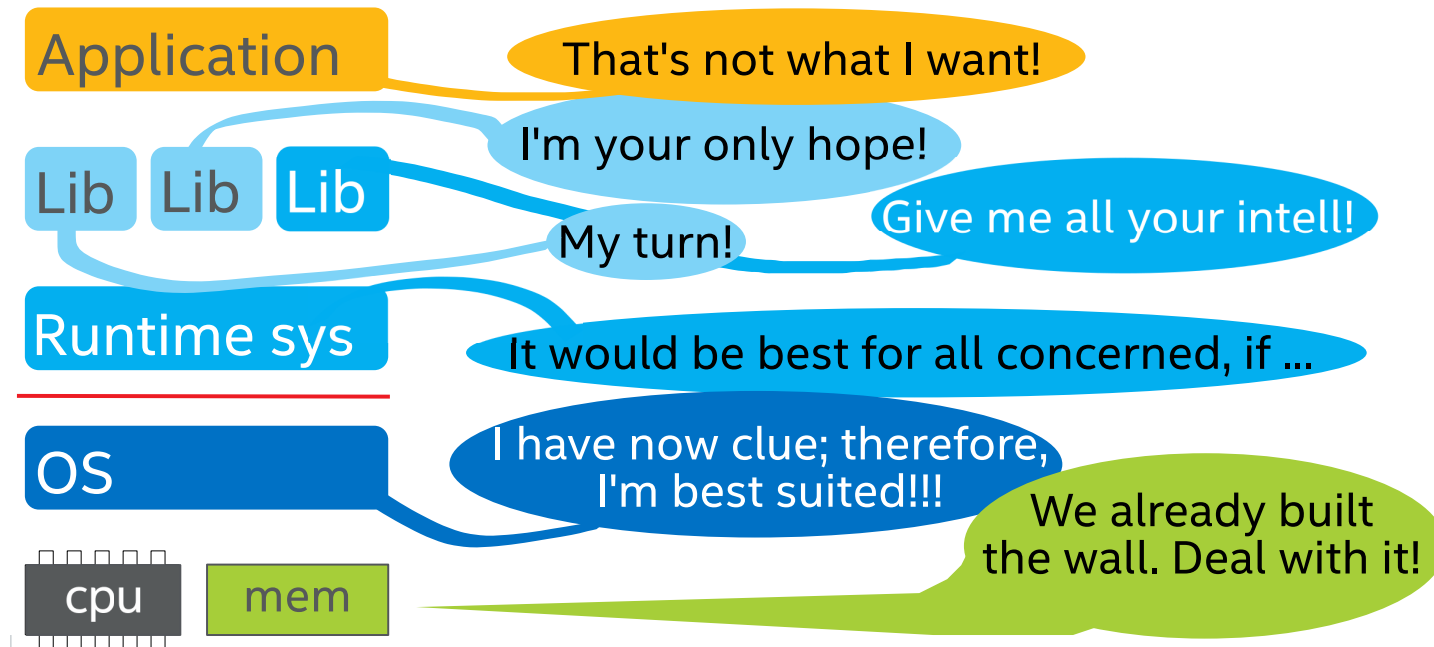


# The role of the OS



*What's the role of the OS and runtime system(s) in managing the memory hierarchy?*

- Intro
- Roles**
- Interfaces
- Details
- Layers
- Functions
- Flexibility
- mOS
- Summary



# Application interfaces



- Intro
- Roles
- Interfaces
- Details
- Layers
- Functions
- Flexibility
- mOS
- Summary

*What application interfaces are needed to help manage the memory hierarchy?*

- Back to the 90s: Socket API not a good match for message passing
  - ◆ Need to know where to put the data before it arrives
- `mbind()` before `mmap()` (or flags to `mmap()`)
- `madvise()` before `mmap()` (or flags to `mmap()`)
  - ◆ Need to know page size, kind, and total size wanted
- `sched_setaffinity()` before `clone()`, or
  - ◆ `clone()` with target CPU

# Level of detail



*What level of detail should the OS expose about the memory hierarchy?*

- Why would anything but full detail be a good option?
- Can always hide things or add a layer of abstraction above
  - ◆ Once it is hidden, it's gone. Just like scalability

- Intro
- Roles
- Interfaces
- Details**
- Layers
- Functions
- Flexibility
- mOS
- Summary

# Software stack



*To what level of the software stack should the OS expose details of the memory hierarchy?*

- I'll target the next level up. Let them take it from there

- Intro
- Roles
- Interfaces
- Details
- Layers**
- Functions
- Flexibility
- mOS
- Summary

# Memory mgmt functions



*What memory management functions should the run-time system contain?*

- Don't know. Maybe all?
- Thomas Sterling, Pavan, ..., have ideas and maybe answers

Intro  
Roles  
Interfaces  
Details  
Layers  
**Functions**  
Flexibility  
mOS  
Summary



# Flexibility and adaptability



*How flexible or adaptable do memory management policies need to be?*

- Very!

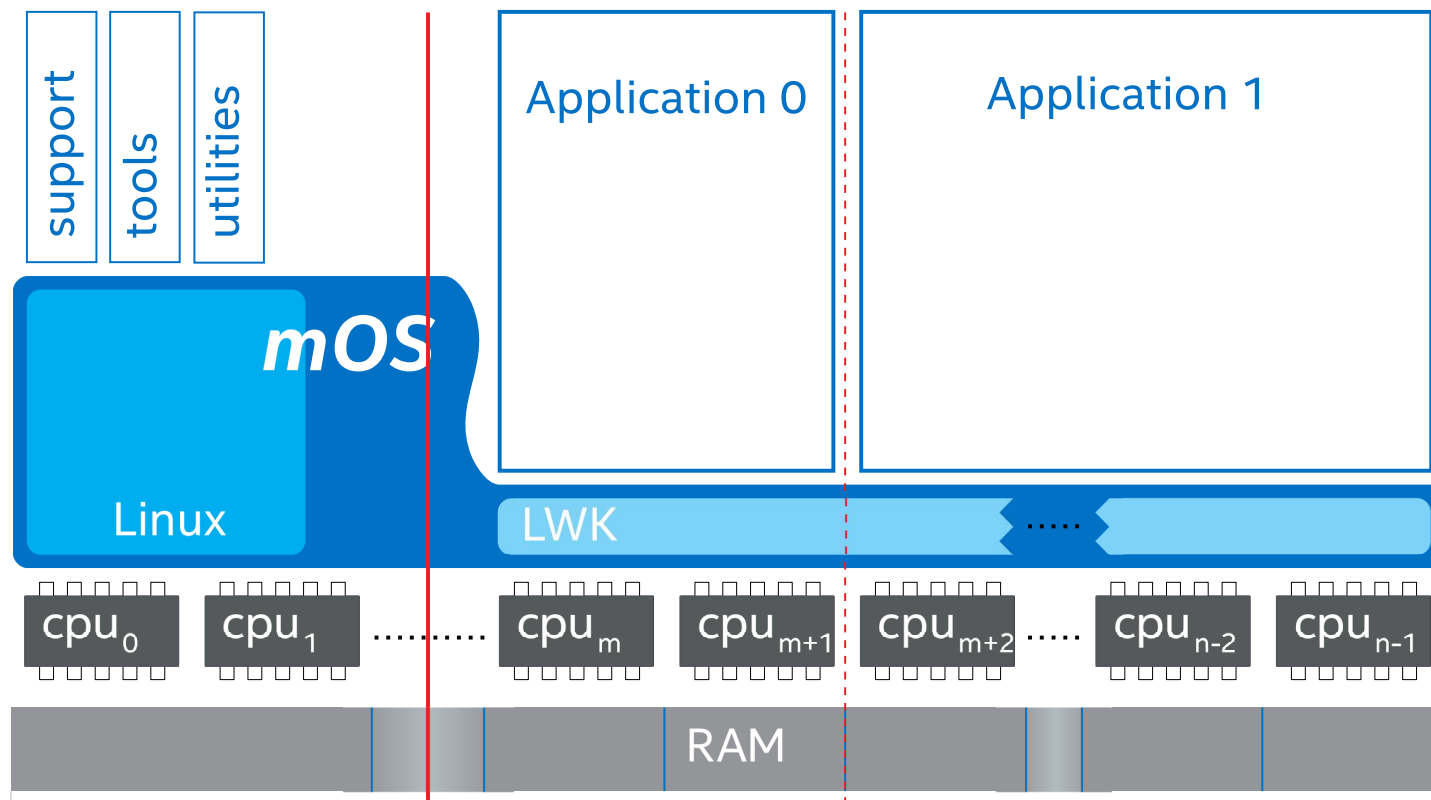
- Intro
- Roles
- Interfaces
- Details
- Layers
- Functions
- Flexibility**
- mOS
- Summary

# mOS partitioning



- Everything old is new again

- Intro
- Roles
- Interfaces
- Details
- Layers
- Functions
- Flexibility
- mOS
- Summary



# Summary



- Intro
- Roles
- Interfaces
- Details
- Layers
- Functions
- Flexibility
- mOS
- Summary

- I'd like you (runtime, lib, app, whoever) to do it
  - ◆ Less work for me, but ...
  - ◆ Some wont do it (right)
    - What works in Linux may not be ideal for mOS
  - ◆ Will get conflicting requests from each layer
- Have to provide defaults anyway
  - ◆ Have to decide at launch and alloc time what to do
- mOS goal: Provide good (sane) defaults (for high-end HPC apps)
  - ◆ Give lots of control to upper layers and user
  - ◆ Have an override switch when they mess up ;-)
- Work with us! Tells us what you need/want, don't want

