



# Serious About Software

## Kay's Dictum in Practice

Bryan Cantrill  
Oxide Computer Company

0x



# Serious about software?


- Ca. 2016 we were at a cloud computing pioneer (Joyent) that had been purchased by Samsung because it needed its own cloud infrastructure...
- The approach that we had taken at Joyent was to control the system software (our own operating system and cloud control plane), but run that on traditional (“commodity”) x86 server hardware
- Running that at much larger scale at Samsung, everything broke – and we found it was impossible to build a reliable system on unreliable hardware

# Alan Kay's dictum

**“People who are really serious about software  
should make their own hardware.”**

– Alan Kay

# Alan Kay's dictum: Origin?



# Oxide: Really serious about software

- In 2019, we started Oxide to deliberately **co-design** software and hardware into a single, unified system to deliver elastic infrastructure on-premises
- We have not developed any of our own silicon, but have developed more or less everything else: our own rack resign, our own compute sleds, our own switch (!) – and then all of our own software



# The scourge of reference designs

- The x86 server industry has been broadly propelled by reference designs: designs from microprocessor vendors that serve as a starting point for an OEM/ODM design
- These reference designs enshrine (ancient!) architectural decisions that impede the software one wants to build to support cloud computing!
- From the outset, we knew we wanted to do things in software that necessitated entirely different board design...

# Software-driven hardware design

- We eliminated the BMC, replacing it with a microcontroller service processor (STM32H753) running our own operating system (Hubris!)
- We eliminated the TPM, replacing it with a dedicated root-of-trust (NXP LPC55S69), which also runs Hubris and measures and controls the SP
- We eliminated the UEFI BIOS (and SMM!), replacing it with host software that boots holistically, bringing up lowest-level units on the CPU directly
- We eliminated the switch operating system, controlling the switch as a (large, complicated!) PCIe device from an adjacent sled

# Really serious software!




Oxide Gimlet compute sled (top view)

# Challenge: Fear of first principles?

- Part of the reason we were being steered away from doing original designs was the fear of implementing from first principles
- This was felt most acutely in the high-speed aspects of the system (DDR, PCIe), as well as the power distribution network
- But even in reference designs these aspects of the system are frequently brittle – it felt like there was an opportunity to do better
- For these aspects, **simulation was essential**


# Simulating DDR5

- For our AMD Turin-based compute sled, we have 12 DDR5 DIMMs with 1 DPC
- Used simulation to determine eye opening vs. reference board at 6000MT
- Simulation results were promising – and DIMMs trained successfully
- Our approach to system software allows us to measure the eye in situ



# Simulating PCIe Gen5

- For our AMD Turin-based compute sled, PCIe Gen5 can have a long path (e.g., for storage: socket, CEM, riser board, U.2)
- Simulation results showed adequately open eyes overall – but that some ports may be marginal at Gen5
- Results on actual boards showed simulation results conservative in some regards; drive firmware XXX



# Kay's Dictum in practice

- We knew that doing our own hardware would allow the software we wanted – but it's been a delightful surprise to see the reverse also true
- By designing our hardware around our software ambitions, we have been able to develop the software to support our hardware design
- The upshot is first-principles design that has allowed us to deliver a sophisticated system – one in which hardware and software have been truly co-designed

# More information

- We have a weekly(ish) podcast, *Oxide and Friends*, that goes into many of the technical details of what we have built, e.g.:
  - [A Measurement Two Years in the Making](#) featuring Larry Williams from Ansys discussing our simulation of our backplane
  - [Oxide and the Chamber of Mysteries](#) on EMI compliance
  - [Bringing up Cosmo](#) on our AMD Turin-based sled
  - [Raiding the Minibar](#) on a test fixture we built