

Procuring Open Compute Hardware

Olof Barring CERN IT



Agenda

Open compute

- What is it?
- Why is it interesting?
- How to buy?
- Lessons learned

History

Started by Facebook in 2011

“...most efficient server, storage and data centre hardware designs for scalable computing”

Foundation board

Facebook, Intel, Microsoft
Rackspace, Arista Networks
Goldman Sachs

opencompute.org

Why

- Hardware is designed for better
 - ✓ Interoperability
 - ✓ Energy efficiency
 - ✓ Cost
 - ✓ Serviceability
- Open
 - ✓ Not only hardware*

* <https://code.facebook.com/posts/1601610310055392/introducing-openbmc-an-open-software-framework-for-next-generation-system-management/>

First try

- In 2013
 - Two Hyve OCP 1500 twins
 - ~25% power gain




Further motivation

More power gains

- Rack with DC distribution
- Centralized AC/DC conversion

Probe market

- Supplier ecosystem
- Large volume price comparison



Who can supply?

“Solution providers”



Gold Membership + fee*

@CERN : “country of origin” is member-state

➤ Integrator + solution provider

* opencompute.org/community/get-involved/tiered-membership/

What to buy?

- Open design \neq unique design
- Can't just limit to design XYZ
- Reconcile needs & wants with reality
- Certification



→ 29 pg tender spec 😊

How much

- ❑ 40kHS06*
- ❑ 12x JBODs (~2 petabytes)
- ❑ 6x racks
 - <7kW / rack
 - Assembled and cabled
 - Our network switches

The rack is the unit!

*HS06 = HEP-SPEC06
2x E5-2630v3 ~360 HS06

Relative price

\$/HS06

- Compare to 250kHS06 contracts in 2014
(E5-26xx v3 in both)
- Compute alone: x1.06
- Compute in assembled rack: x1.30

\$/TB

- Compare to 25PB contracts in 2014: ~x1.14
(6TB vs 4TB)

Apples to apples:

Compute: +6% for 6x less volume

Storage: +14% for 12x less volume

Lessons

- Open \neq unique
- Qualifying suppliers
- Certification - early days

At the end of the tunnel

- Saving cost & power
- Standardization