Centre de Calcul
de l'Institut National de Physique Nucléaire
et de Physique des Particules

IN2P3-CC ARM evaluation

nov 23 – LCG France

ARM evaluation : motivations

ARM widely recognized as Watts sparing (W/HS23)

- CC-IN2P3 is a huge consumer : financial and environmental impacts are quite large.
- Does that help solving more use cases at constant budget ?
- Does that help with AMD/Intel TDP’s explosion : power distribution, heat dissipation ?

ARM evaluation : motivations

ARM widely recognized as Watts sparing (W/HS23)

- CC-IN2P3 is a huge consumer : financial and environmental impacts are quite large.
- Does that help solving more use cases at constant budget ?
- Does that help with AMD/Intel TDP's explosion : power distribution, heat dissipation ?

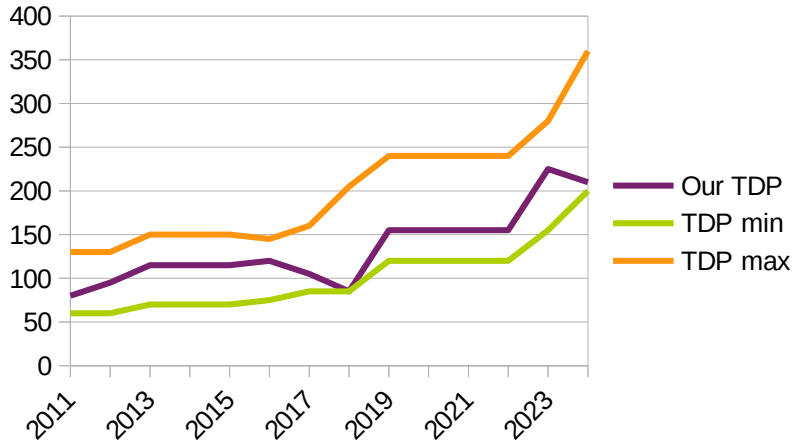

ARM evaluation : motivations

ARM widely recognized as Watts sparing (W/HS23)

- CC-IN2P3 is a huge consumer : financial and environmental impacts are quite large.
- Does that help solving more use cases at constant budget ?
- Does that help with AMD/Intel TDP's explosion : power distribution, heat dissipation ?

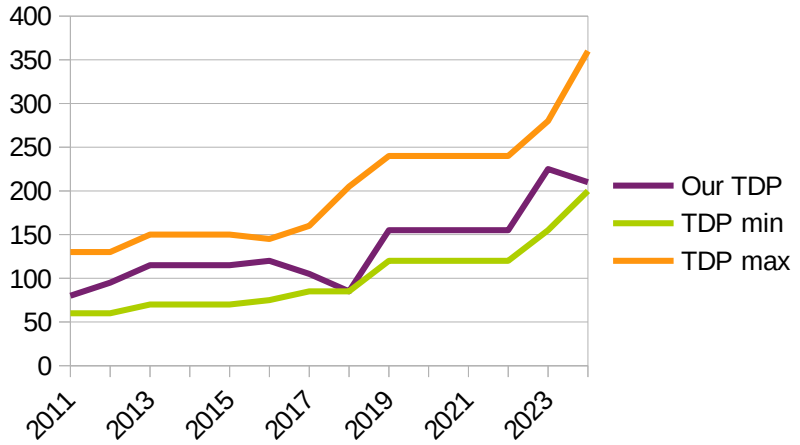

How is that true ?

ARM vs EPYC Milan & Genoa

Benchmarked platforms :

- HPE Proliant RL300 Gen11 with Ampère Altra Q80-30 and Ampère Altra Max Q128-30
- HPE Proliant DL385 Gen11 with AMD EPYC Genoa 9474F / 9354 / 9334
- HPE Apollo 2k Gen10+ with AMD EPYC Milan 7453 (last production batch)

CPU model	freq (GHz)	cores/ sock	thread/sock	sockets	TDP	Process (nm)	Release
Ampere Altra Max (M128-30)	3	128	128	1	250	7	2021
Ampere Altra (Q80-30)	3	80	80	1	210	7	2020
AMD EPYC Milan 7453	2.75	28	56	2	225	7	2021
AMD EPYC Genoa 9474F	3.6	48	96	2	360	5	2022
AMD EPYC Genoa 9354	3.25	32	64	2	280	5	2022
AMD EPYC Genoa 9334	2.7	32	64	2	210	5	2022

Gamme	min cores	max cores	min TDP	max TDP
Ampere Altra	32	128	45	230
EPYC Milan	8	64	155	280
EPYC Genoa	16	96	200	360

HS06 correlation to HS23

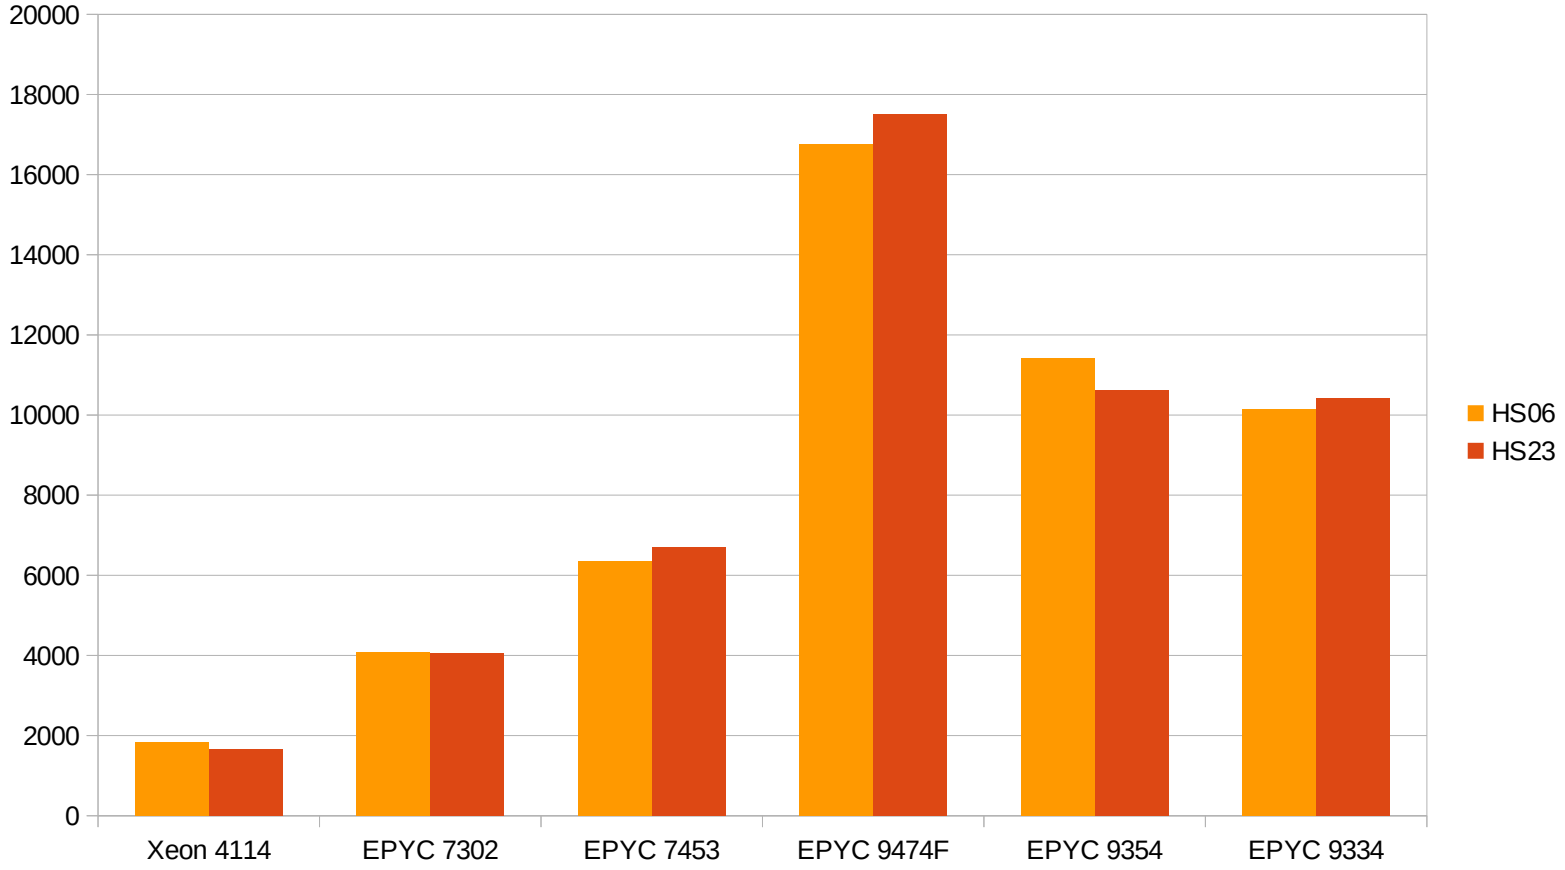

Correlation : 0,996177425123253

Computing power

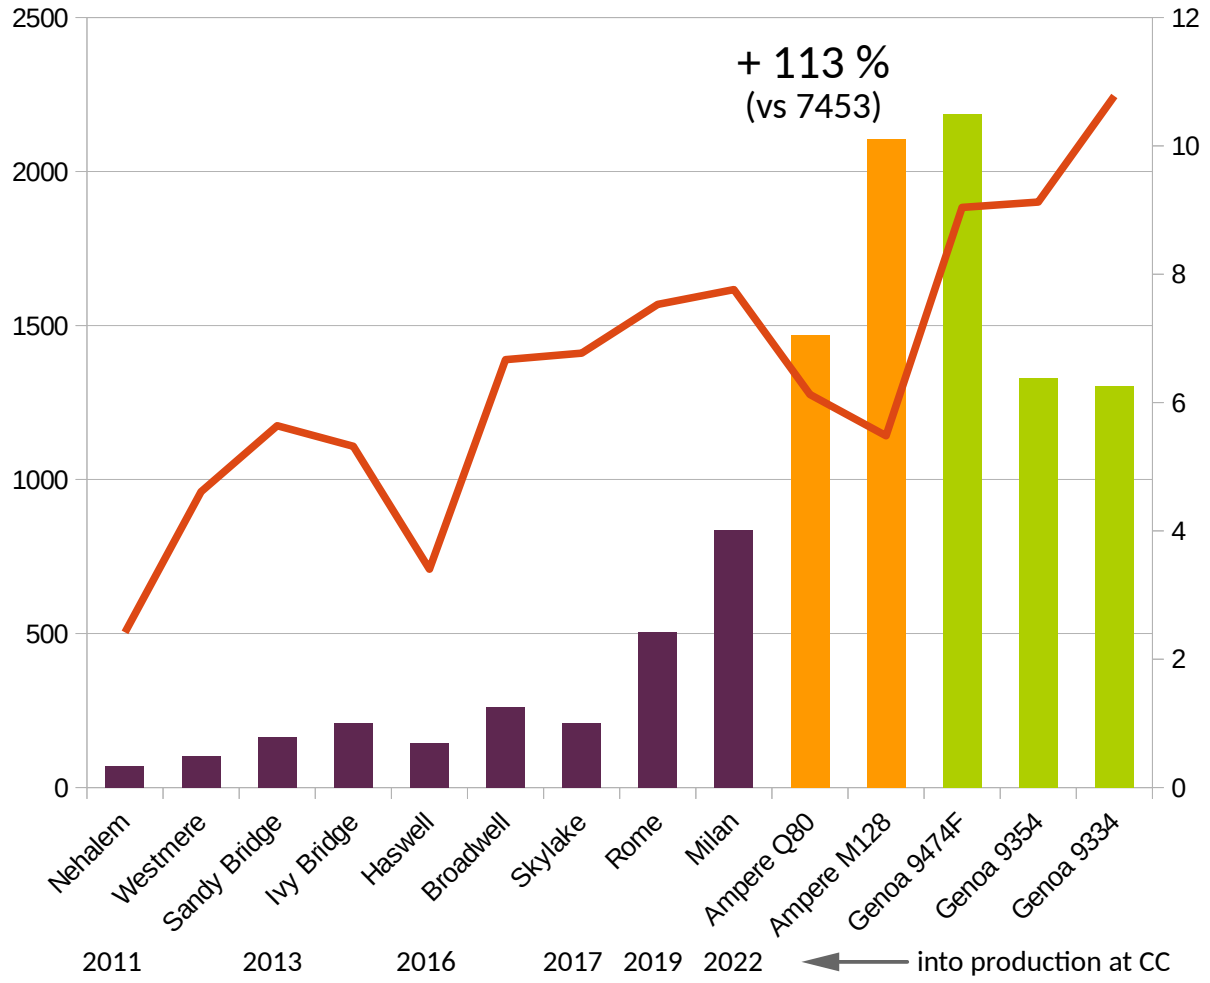

HS06 measurements up to Broadwell,
HS23 since
→ let call this HS0623

Power effectiveness (W/HS0623)

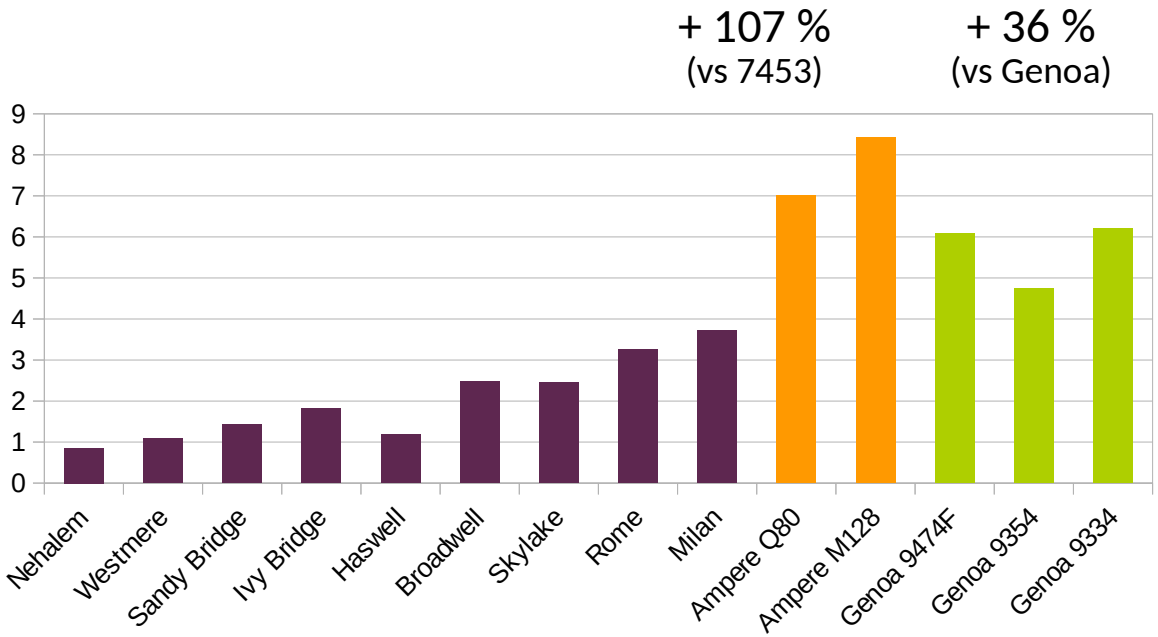

Processor HS0623 / TDP

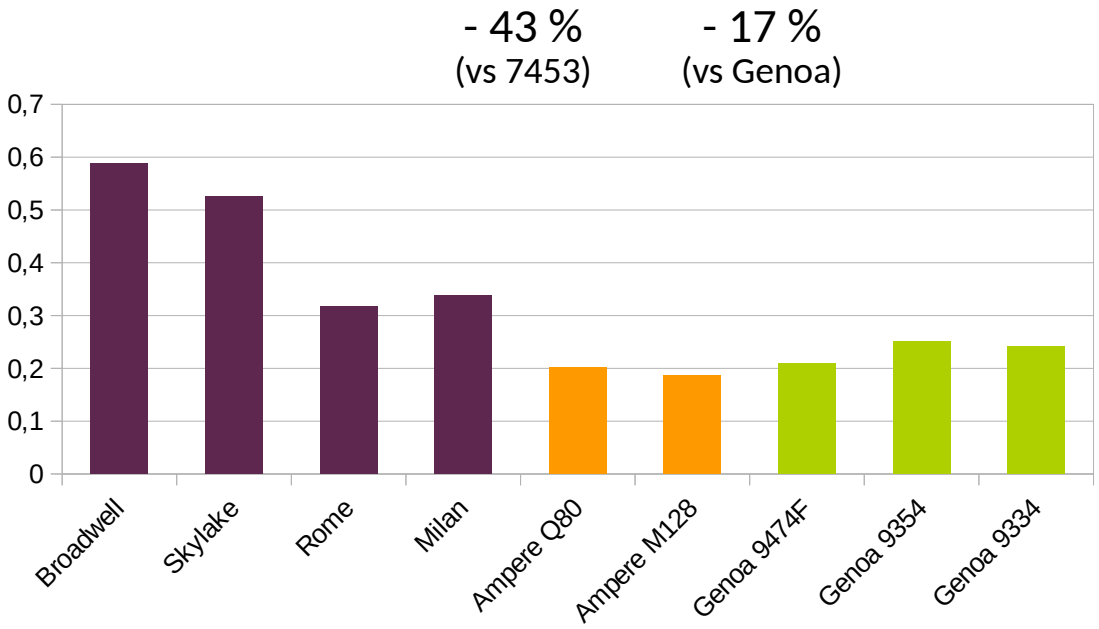

Chassis W / HS0623

Caution : Genoa power measurements extrapolated to 2U4N BP from a HPE DL380.

Final gain is undermined by server form factor (1U1N UP).

Density (HS0623/U, W/U)

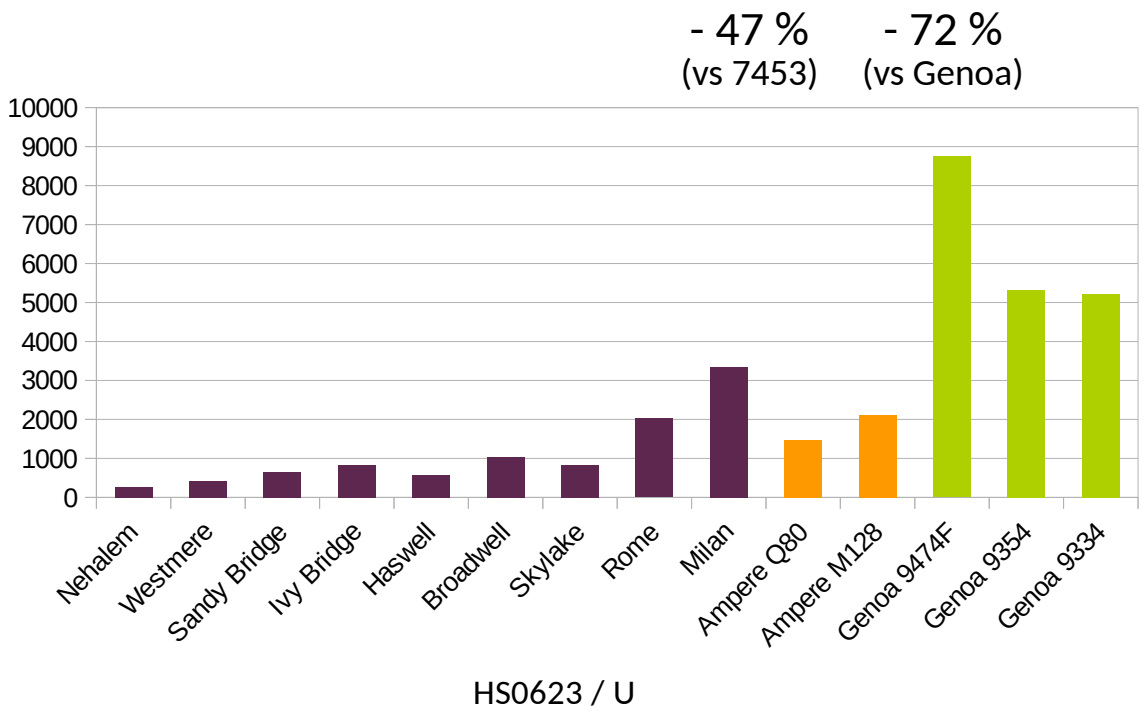

HS0623 / U

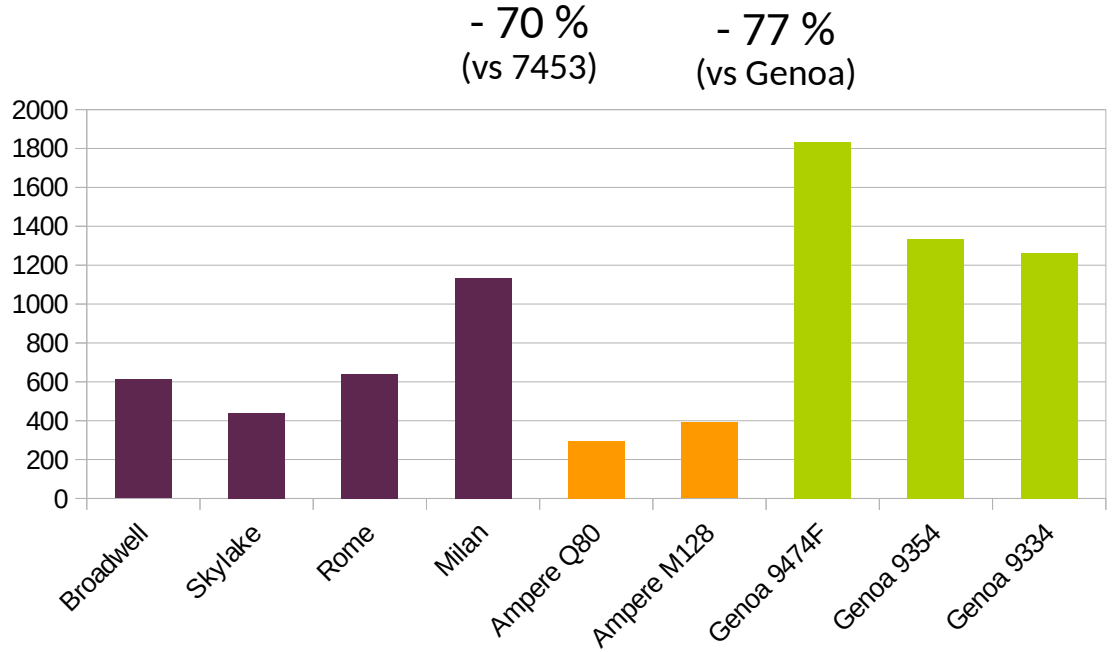

W / U

Note : again, density is largely the 1U1N UP integration consequence.

TCO 5y (€/HS23)

10 cts / kW.h

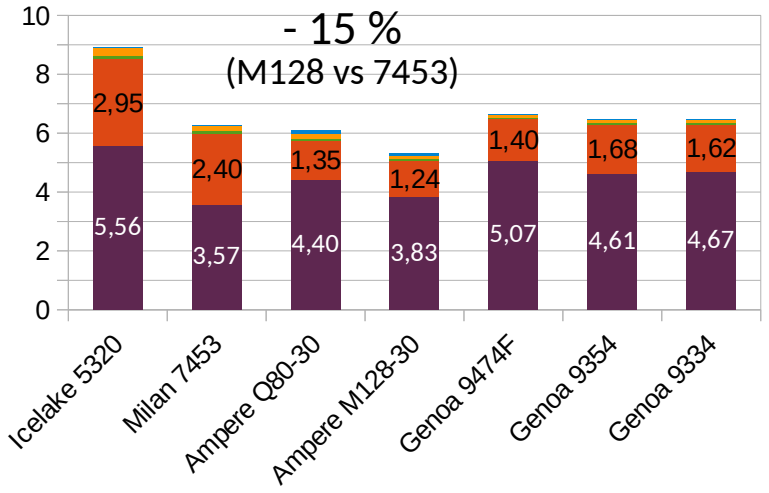

20 cts / kW.h

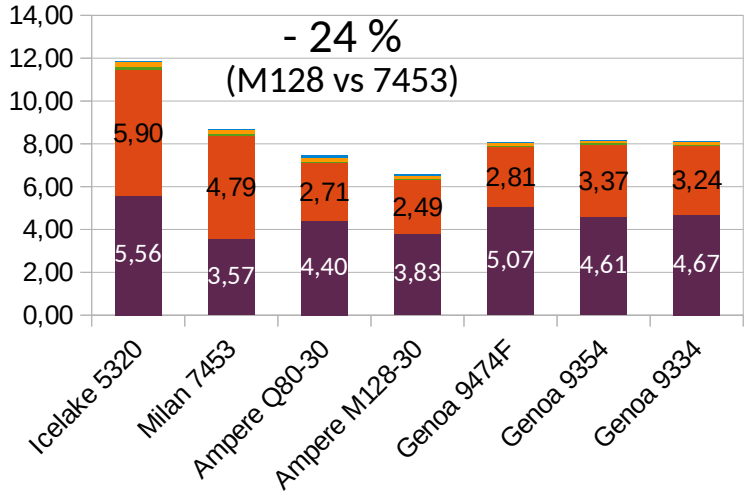

30 cts / kW.h

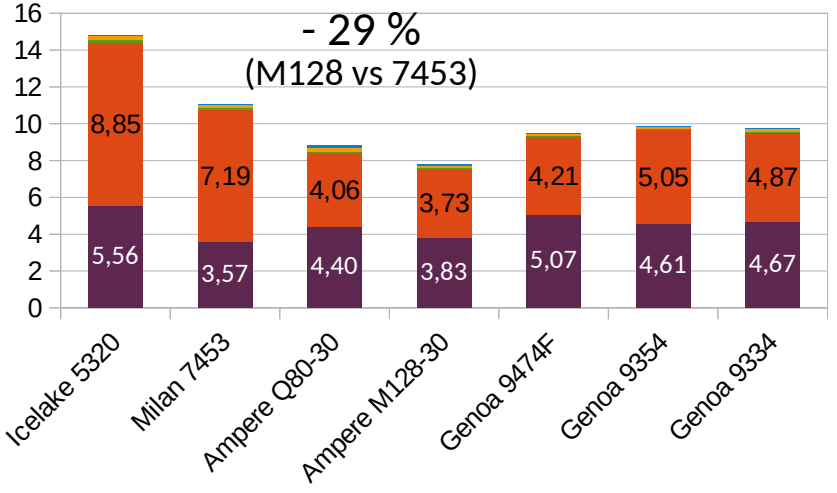

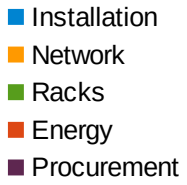

Caution : Genoa & Icelake procurement costs are not final

Procurement is a bit more expensive, compensated by power effectiveness.

Conclusion & follow up

Early conclusions

- x86 comparison : more, tinier, no SMT cores.
- RL300 is not the way :
  - W / HS23 43 % inferior to Milan 7453 but we may expect more.
  - RL300 chassis not dense enough (-47 % HS23 / U vs 7453).
- Reduced PCI peripherals offer compared to x86_64.

What's next ?

- use cases qualification :
  - LSST first player
  - users enrollment, first farm commissionning, software limitations…
- Integration :
  - CI
  - environment (installation, configuration, authn/z, filesystems…).
  - batch systems (Slurm/HTCondor ?).
- MATINFO limited to HPE RL300, follow up with GigaByte coming (day after tomorrow).
  - H262-P60 (2U4N).
  - R183 (AmpereOne 5nm 1U1N DP).
  - Others (SuperMicro, Foxconn, Wiwynn…) ?
- Shall we test NVIDIA's Grace chip ?

Conclusion & follow up

Bigger picture, questionings & future

- ARM ecosystem looks not as mature as x86_64 : how do we manage rolling procurements, architectures compatibilities ?
- Limited to local Slurm farm for now (broader testings, less gains to expect)
  → looking forward WLCG's moves about pledges.
- How are power gains intrinsincs to ARM, how shall we forecast superiority against x86_64 ?
  (AmpereOne generation TDPs range from 200 to 350W)

Questions

Appendix : Ampere's roadmap

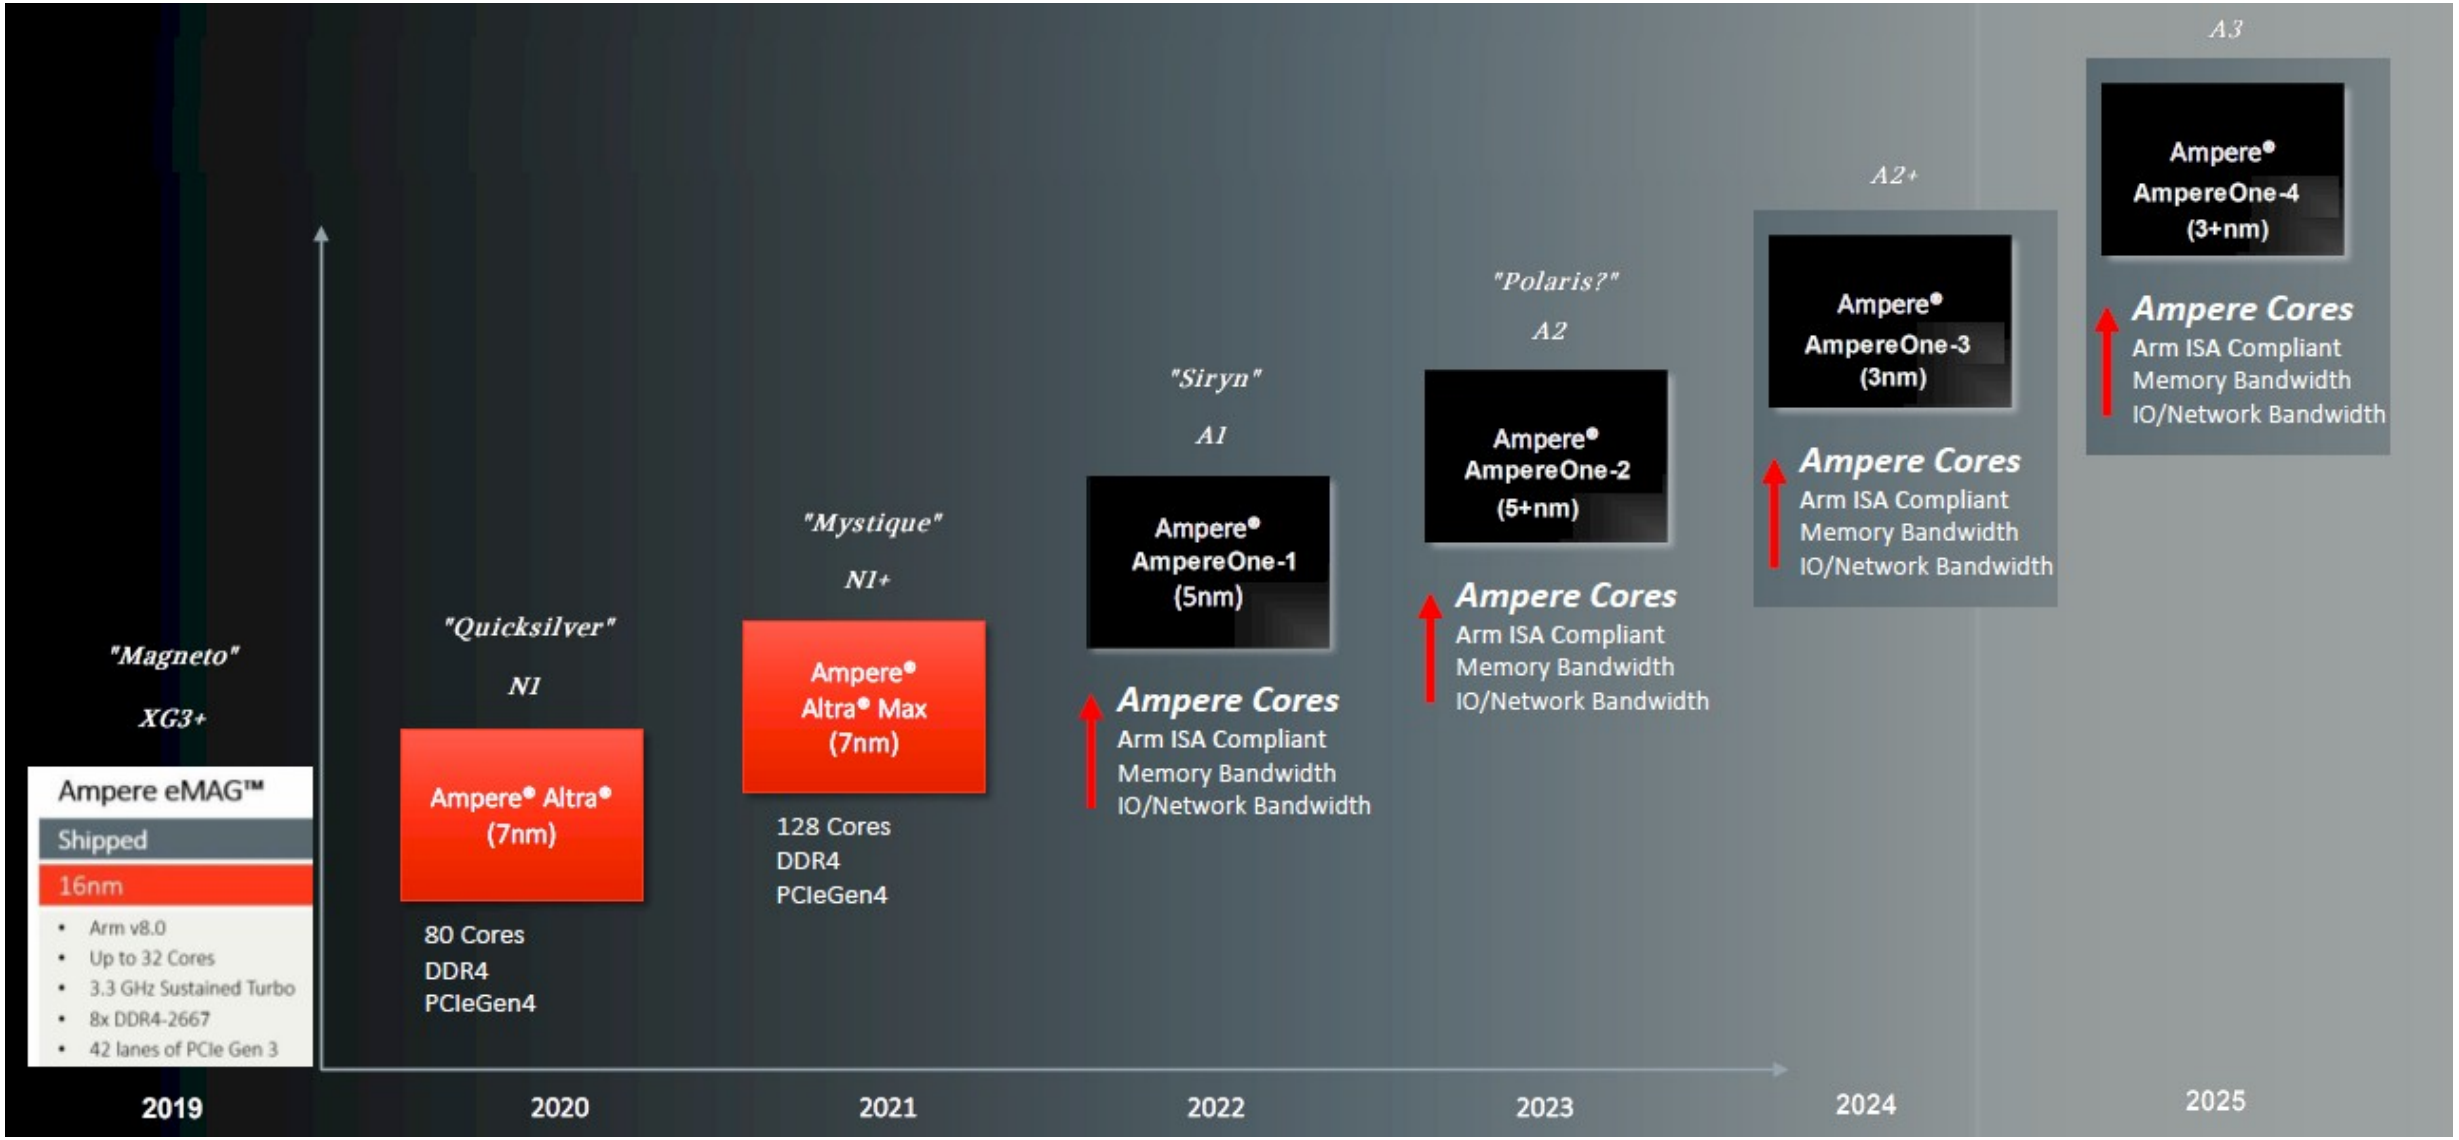

Appendix : AMD’s roadmap