



HAL
open science

R'n'R building of an experimentation platform for the Edge

Pierre Neyron

► To cite this version:

Pierre Neyron. R'n'R building of an experimentation platform for the Edge. Journée Datamove 2024, Jun 2024, Chichillianne, France. pp.1-12. <hal-04604573>

HAL Id: hal-04604573

<https://inria.hal.science/hal-04604573v1>

Submitted on 7 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

R'n'R building of an experimentation platform for the Edge


Journée Datamove 2024

Pierre Neyron - Ingénieur de Recherche LIG/CNRS


6 juin 2024



Edge computing devices?



Cloud	Fog	Edge	IoT
Thousand of DC	Millions of nodes	Billion of devices	
Infrastructure		End-devices	
High bandwidth, e.g. wired		Low power, possibly wireless	
Linux or equiv.			RTOS/baremetal

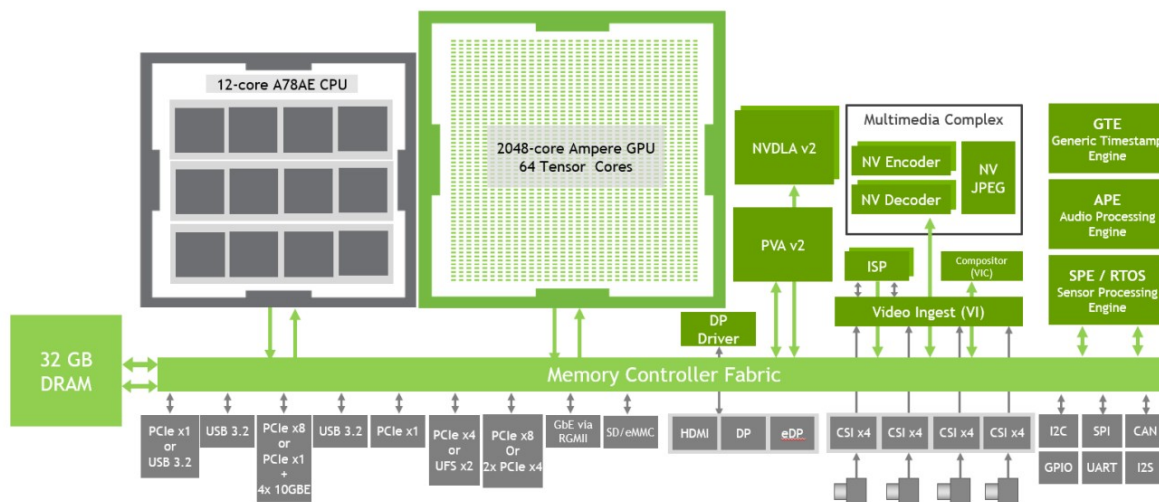


Edge computing devices?





- Typical edge computing device: **Nvidia Jetsons**

- ▣ System on Chip (SoC) – CPU+GPU+... all in one
- ▣ ARM CPU – Low power, fewer cores (4 → 16) than servers
- ▣ Nvidia embedded GPU – smaller/one gen behind servers
- ▣ Other accelerators (DL, Vision, ...)



Platforms for research 1/2

- Grid'5000
 - 2004...
 - Large scale experiments, distributed clusters : 750 nodes
 - Various hardware: CPUs, Accelerators, Memory/Disks, Network
 - Services for research: HPC, Grid, Cloud, P2P, Network, ..., IA, ...
→ an experimental data-center: **only server-class machines**
 - Form-factor: rackable
 - Server components: e.g. only server-class CPU, GPU, RAM, Network
 - Locality: centralized places (regional DC)
- Toward “Slices”
 - ... 2040
 - Continuum from Cloud to IoT
 - Work-in-progress



Platforms for research 2/2

- Slices builds from other platforms also:

- Fit-IoT-Lab

- The Grid'5000 of IoT
- Grid of sensors, mobile devices, ...


- Walt (Drakkar)

- Networked devices deployment

- Others...

- Integrating Edge computing devices is pretty new

- → *How does it fit in the Grid'5000 model ?*



Grid'5000 integration

Technical challenges


- Device constraints

- Exotic purchase
- Exotic form factor
- Exotic architecture
- Exotic/lack of components
- Exotic/lack of boot mechanism
- Exotic OS stack

- Grid'5000 constraints

- Datacenter hosting
- Control over network
- OS reconfiguration **X 1000 ?**
- Hardware validation
- Network topology reconfiguration
- Power meters

Grid'5000 integration make it to the data-center!




From the devkit to a ***data-center compatible*** chassis

- With out-of-band control (on/off, serial console)
- With out-of-band power meters

2 considered options:

- fablab made chassis → more work
- industrial solution → less flexible



Grid'5000 Integration



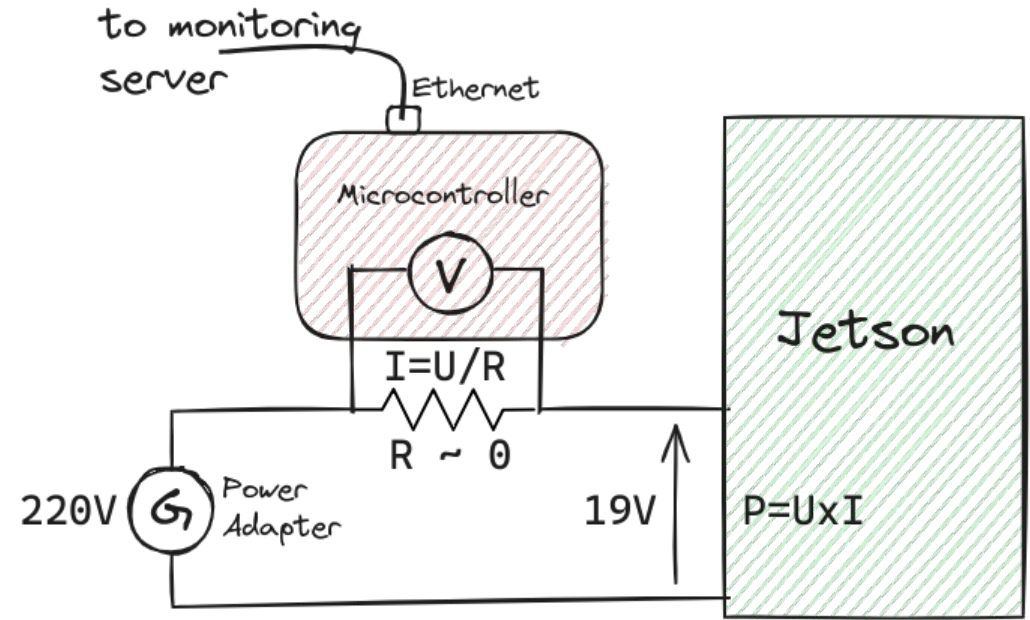
Out-of-band control

Data-center servers come with a BMC. What about Edge devices ?

Required functions:

- On/Off
- Serial console (if no network)
- Recovery
- **Power meter**

No on-the-shelf solution ?
DiY with microcontroller ?



Grid'5000 Integration OS swing – Boot switch

Datacenter servers come with PXE network boot


→ allow to remotely choose what to boot (**key to kadeploy**)

Embedded devices boot mechanism varies...

- Some have U-boot with network support: **good!**
- Nvidia happened to have a vendor bootloader without network: **bad**
- UEFI without working PXE: **bad**

Boot switch over USB ? Programmable USB device ? **wow**

Post Linux kernel swing ? (giving up on allowing to change kernel)... **mmm**



Grid'5000 Integration

OS reconfiguration

Edge == Embedded Linux devices come with exotic software stacks.

E.g. for Nvidia Jetson: ***Complex stacking of software layers***

- Base OS (e.g. Ubuntu)
- Nvidia layer, with specific kernel, libs (cuda, ...), services
- + Nvidia Partner layer: custom device description, additional drivers, ...

Non-standard machinery for installation (flash)

Required to ***reverse-engineer*** the installation process.

Need to automate: ***Kameleon*** to the rescue with build recipes!





Grid'5000 Integration

Network topology emulation

User control of the physical network topology of hundreds of devices is out of the scope of Grid'5000.

→ Emulation with the G5K network reconfiguration tool (kavlan) + Linux tools



Take away message

- Hard technical work (years!) but:
 - 12 Edge nodes in “production”, available in G5K Toulouse
 - A few services missing:
 - Many platform tests → need reworking the whole Grid’5000 test infra
 - Network topology emulation → kavlan-ng in dev, should support it
 - Out-of-band power meters → go custom design if still not on the shelf?
 - Used for Mathilde & Tuan’s work: federated learning
 - 100 nodes better → what about 1000 ?
- More to come in Slices...