



Contribution ID: 19

Type: **not specified**

ML applications in Linux kernel

(1) A description of the overall topic

The number of areas of Machine Learning (ML) approaches application is growing with every day. There are already research works and industry efforts to employ ML approaches for configuration and optimization of Linux kernel. The using of ML approaches in Linux kernel is inevitable step. But there are multiple unanswered questions how ML can be used in Linux kernel efficiently. This micro-conference has goal to discuss:

1. How can we use ML approaches in Linux kernel?
2. Which infrastructure do we need to employ in Linux kernel for adoption of ML methods?
3. How can we use ML to automate testing, bugs detection, and bugs isolation in Linux kernel?
4. Can we use ML for automated refactoring and bug fix in Linux kernel code?
5. Can we use the whole fleet of Linux kernel deployments for massive and distributed testing of Linux kernel functionality?
6. How can we optimize the Linux kernel by using ML techniques?

The main discussion point is how to make ML techniques working in Linux kernel without affecting performance and efficiency the Linux kernel operations?

(2) The list of sub-topics that would be appropriate for the MC

- Using AI/ML for Linux kernel code management
- Using ML for kernel testing and bug fix
- Using ML techniques for efficient task scheduling policies elaboration
- ML infrastructure in Linux kernel
- Using ML Techniques + eBPF for achieving efficient kernel configuration
- Using ML methods for enhancing file systems reliability
- Using ML methods for prediction storage drives failures
- Using ML methods for memory access patterns analysis
- Using ML methods for delta-encoding and decreasing WAF in file systems
- Using ML methods for FSCCK and file systems synthesis
- Using ML methods for efficient data placement policy
- Using ML methods for CXL benchmarking
- Using ML for re-writing the whole Linux kernel in Rust
- Cognitive file systems: using ML for data classification and data organization
- Using hardware acceleration for employing ML methods on kernel side

(3) List of key attendees

- Sasha Levin
- Mcgrof Chamberlain
- Jim Huang
- SeongJae Park
- Cong Wang

- Jasmine Mou
- Anuj Gupta
- Neema Mashayekhi
- Wenhui Zhang
- Henry Huang

Primary author: DUBEYKO, Viacheslav (IBM)

Presenter: DUBEYKO, Viacheslav (IBM)

Session Classification: Birds of a Feather (BoF)

Track Classification: Birds of a Feather (BoF)