

ICN - KuD addons

Kud addon plugins will be invoked by Kud (Kubernetes deployer) in R2 release to deploy K8s addons in edge location, Kud addon plugins include:

- Prometheus
- Rook: Deploy containerized Ceph cluster by Rook operator and provide storage service in edge location node.
- Virtlet
- NFD
- OVN
- Intel device plugins: deploy Intel device plugin for K8s to enable the use of Intel devices such as SRIOV, QAT, GPU & FPGA in edge location node

Node Feature Discovery detects hardware features and advertises them through a Kubernetes concept called node labels which, in turn, can be used to control workload placement in a Kubernetes cluster. NFD runs as a separate container on each individual node of the cluster, discovers capabilities of the node, and finally, publishes these as node labels using the Kubernetes API.

Some of the features that NFD currently detects:

CPUID: Intel® processors have a special CPUID instruction for determining the CPU features, including the model and support for instruction set extensions, such as Intel® Advanced Vector Extensions (Intel® AVX). Certain workloads, such as machine learning, may gain a significant performance improvement from these extensions (e.g. AVX-512). NFD advertises all CPU features obtained from the CPUID information.

SR-IOV : networking: Single Root I/O Virtualization (SR-IOV) is a technology for isolating PCI Express* resources. It allows multiple virtual environments to share a single PCI Express hardware device (physical function, PF) by offering multiple virtual functions (VF) that appear as separate PCI Express interfaces. In the case of network interface cards (NICs), SR-IOV VFs allow direct hardware access from multiple Kubernetes pods, increasing network I/O performance, and making it possible to run fast user-space packet processing workloads (for example, based in Data Plane Development Kit). NFD detects the presence of SR-IOV-enabled NICs, allowing optimized scheduling of network-intensive workloads.

Intel® RDT: Intel® Resource Director Technology (Intel® RDT) allows visibility and control over the usage of last-level cache (LLC) and memory bandwidth between co-running workloads. By allowing allocation and isolation of these shared resources, and thus reducing contention, RDT helps in mitigating the effects of noisy neighbours. This provides more consistent and predictable performance which may be essential in meeting Service Level Agreements (SLA), for example. NFD detects the different RDT technologies supported by the underlying hardware platform.

Intel® Turbo Boost Technology: Intel® Turbo Boost Technology accelerates processor performance for peak loads, dynamically overclocking processor cores if they are operating within the power, current, and temperature limits of the processor. This can provide significant performance benefits for CPU-bound workloads. On the other hand, some workloads behave better when this technology has been disabled. NFD detects the state of Intel® Turbo Boost Technology, allowing optimal scheduling of workloads that have a well-understood dependency on this technology.

IOMMU: An input/output memory management unit (IOMMU), such as Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) technology,

allows isolation and restriction of device accesses. This enables direct hardware access in virtualized environments, highly accelerating I/O performance by removing the need for device emulation and bounce buffers. This can be crucial for I/O heavy workloads in Kubernetes deployments using hypervisor-based container runtimes, such as Kata* Containers. NFD detects if an IOMMU is supported by the host hardware platform and enabled in the kernel of the host operating system.

Linux* kernel: Some specific workloads may be highly dependent on the kernel version of the underlying host operating system. For example, some kernel features may be required to be able to run an application, or, they provide measurable performance benefits. NFD detects the kernel version and advertises it through multiple labels, allowing the deployment of workloads with different granularity of kernel version dependency.

References

Intel Node Feature Discovery Document Number: 606833-001 7