Solution Brief

KDDI and Intracom Telecom, Collaborate on More Efficient 5G Network Power Usage Supporting Lower CO_2 Emissions

Intracom Telecom's NFV-RI™ Machine Learning (ML) based solution, leveraging Intel server telemetry and insights, can optimize KDDI's 5G Network power efficiency, lowering energy consumption, and enabling reduced CO_2 emissions.

Authors

KDDI Corporation
Masanori Harada
Tooru Yamamoto
Hajime Miyamoto
Hidenori Ishimura

Intracom Telecom
Nikos Anastopoulos
Evangelos Angelou
Christos Rizos

Intel Corporation
Wei Xian
Emma Collins
Takeru Tsuchiya
Masato Hotta

Executive Summary

KDDI Corporation (KDDI) is one of Japan's largest Tier-1 telecommunications operators. As a business entity, KDDI emits approximately one million tons of CO_2 annually through the consumption of energy such as electricity. For context, that's equivalent to the CO_2 emissions of approximately 500,000 households each year. By far, the largest portion of those emissions (98%) is due to the electricity consumed by mobile phone base stations, communication station premises, and data centers. Further increases are anticipated as the rollout of 5G continues and communications traffic volumes grow over time.

To move toward carbon neutrality, KDDI is firmly committed to reducing the power consumption of its mobile phone base stations and communications equipment and the use of renewable energy. To align with Japan's national Climate Action Plan, KDDI aims to achieve net zero CO_2 emissions from its business activities by fiscal 2030, ahead of the previous declaration. In addition, the entire KDDI Group aims to achieve net zero CO_2 emissions by fiscal year 2050.

To achieve this target, KDDI has several approaches to reducing emissions, including research into finding new ways to conserve energy in existing infrastructure.

Intracom Telecom is a global telecommunication systems and solutions vendor providing equipment and services to the industry for over 40 years. As virtualized infrastructures become increasingly critical in Communication Service Provider (CoSP) deployments, Intracom Telecom has developed solutions that employ Machine Learning (ML) to find the best power profile for optimizing power consumption. The Network Function Virtualization – Resource Intelligence (NFV-RI) solution can determine the ideal distribution and configuration of resources automatically, in closed-loop form, and dynamically, under any traffic.

KDDI, Intracom Telecom, and Intel partnered to explore possibilities for reducing power consumption for 5G core network workloads. The resulting Proof of Concept (PoC) is the subject of this paper.

The PoC demonstrates the ability to reduce power consumption by around 30 to 35W for a single data plane pod running on a single data plane server with a baseline power consumption of 360W, with possibilities to reduce power consumption further while maintaining service quality levels.

Problem Description

The benefits of Network Functions Virtualization (NFV) are well understood by CoSPs and are a key enabler as 5G networks expand. By virtualizing functions in the two fundamental parts of a mobile network, namely 1) the mobile core and 2) the Radio Access Network (RAN), many CoSPs have reported a significant Total Cost of Ownership (TCO) reduction.

The 5G data plane constitutes a significant part of a 5G mobile core. To meet stringent Key Performance Indicators (KPIs) such as deterministic performance, low latency, and zero packet drops, the 5G data plane relies on the Data Plane Development Kit (DPDK), the de-facto framework for accelerated packet processing. The DPDK employs polling to meet these KPIs, which is however, by design, oblivious to the load level of the user planes. As a result, CPU cores are always at their highest utilization, even when a user plane is idle or receiving the lowest amount of traffic. As a result, data plane servers always consume the maximum possible power, resulting in a higher Operational Expense (OpEx).

Figure 1 illustrates the problem in a graph that shows the notional network load over a 24-hour period. CPU cores are always 100% busy due to the polling mechanism even when the traffic load is light. The green bars show opportunities for power savings where CPU frequency can be turned down and still support the traffic load without packet drops. Conversely, CPU frequency can be turned up to handle the high traffic load during busy periods.

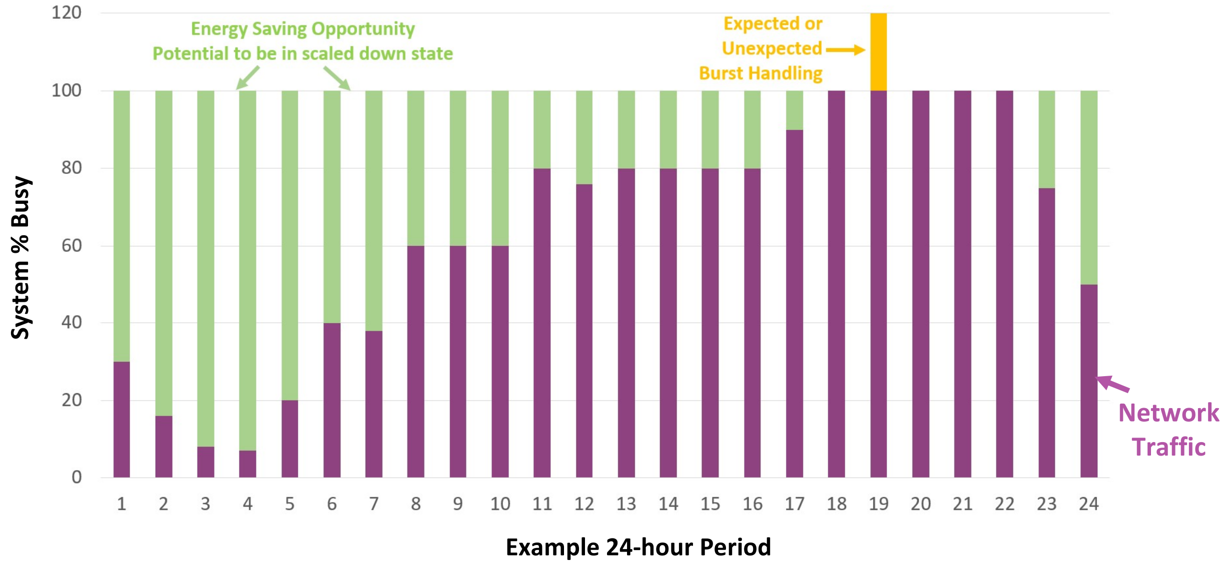

Figure 1. Notional Network Traffic Load Over a 24-Hour Period

Existing system software, such as the Linux operating system, include power management policies to dynamically manage power for CPUs when they are not busy. However, because the data planes running DPDK always appear to be at 100% CPU utilization (due to the polling scheme), the OS is unaware of their actual CPU demand at any given time.

Finding ways to leverage knowledge of the actual data plane demand facilitates the implementation of mechanisms to adapt CPU frequencies according to data plane load. Adjusting a CPU's frequency changes the power consumption leading to opportunities for power savings and consequently reduced CO_2 emissions.

Solution Description

To improve the energy footprint of their 5G mobile network, KDDI teamed up with Intracom Telecom to develop a PoC evaluating Intracom Telecom's NFV-RI solution to reduce the power consumption of KDDI's 5G data plane servers. NFV-RI leverages Intel's telemetry and power management technologies to retrieve the metrics necessary for network load determination and adjust CPU frequencies accordingly. The result can be significant energy savings.

High-Level Data Plane Structure

This PoC focuses on the 5G data plane that employs servers based on Intel's 2nd Generation Intel® Xeon Scalable processors.

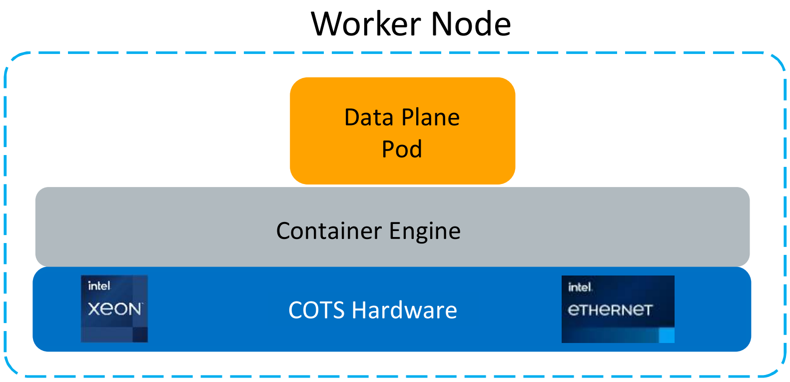

Figure 2. 5G Data Plane Setup

The setup for emulating the 5G core nodes around the 5G data plane server is the subject of the next section.

PoC Setup

The PoC setup comprises a traffic generator to simulate and handle traffic on the following 5G interfaces as shown in *Figure 3*:

- N3 interface to a Radio Access Network (RAN), for example a radio base station (gNodeB) with connected User Equipment (UE) devices
- N4 interface to the Session Management Function (SMF)
- N6 interface to the Data Network that provides IP services

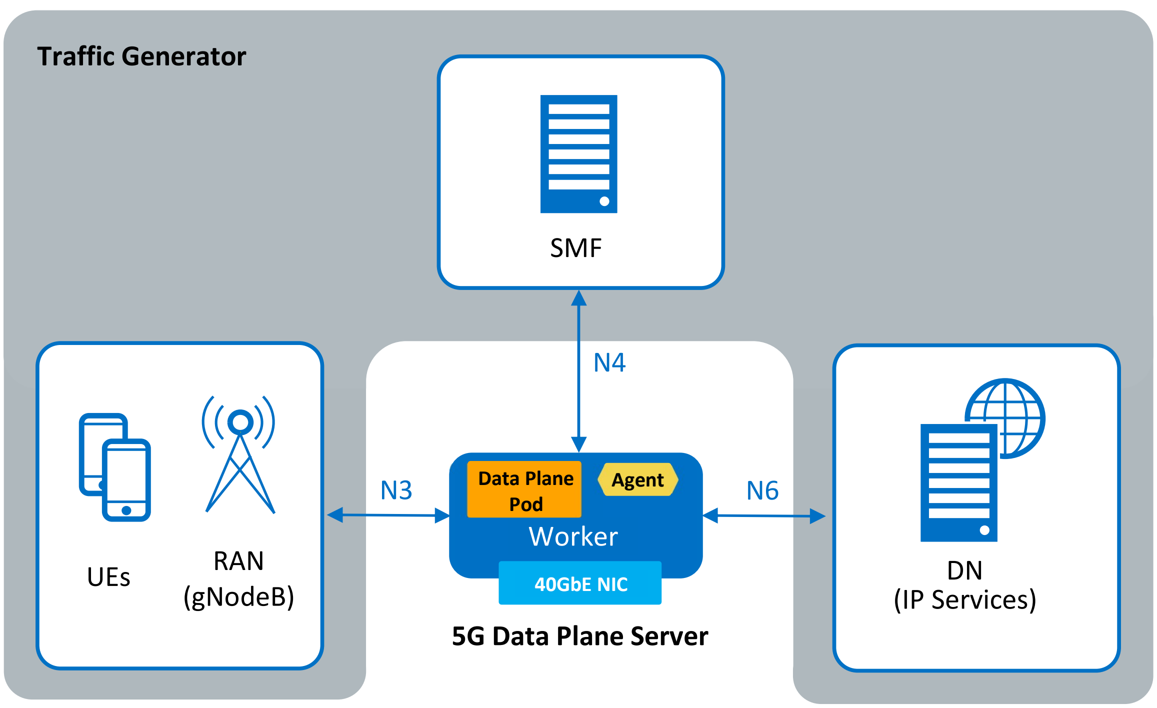

Figure 3. Data Plane Server in an Emulation Environment

The emulated 5G data plane server (Worker) has the following features:

- **CPU**: Intel® Xeon Gold 6230N 2.30 GHz 20C/40T, two socket
- **Memory**: 312 GB
- **Hard Disk Driver (HDD)**: 900 GB
- **Network Interface Card (NIC)**: 40 GbE

The software includes:

- Kubernetes bare metal platform software (cloud platform software)
- 5G data plane software
- Intracom Telecom NFV-RI with Frequency Feedback Loop (FFL)

To facilitate CPU performance scaling on the Worker, the `intel_pstate` driver is enabled.

The server runs multiple instances of a commercial 5G data plane Network Function (NF), which consists of complex traffic accelerations including a DPDK distribution from a leading industry vendor, implemented as Cloud-native Network Functions (CNFs) for Kubernetes-based environments.

NFV-RI uses a Frequency Feedback Loop (FLL) workflow to predict a DPDK-based CNF's traffic levels and dynamically adjust its CPU cores' frequencies according to its incoming load, while promoting zero packet drops. The FFL uses Machine Learning (ML) to predict a CNF's behavior in upcoming traffic changes and adjust its CPU cores frequency accordingly, to enable reduced power usage during light or off-peak traffic periods, without compromising performance.

Figure 4 is a view of the setup from the Kubernetes perspective. The NFV-RI Agent pod describes the resource requirements of an NFV-RI Agent container, and the Data Plane pod describes the resource requirements of the 5G data plane server.

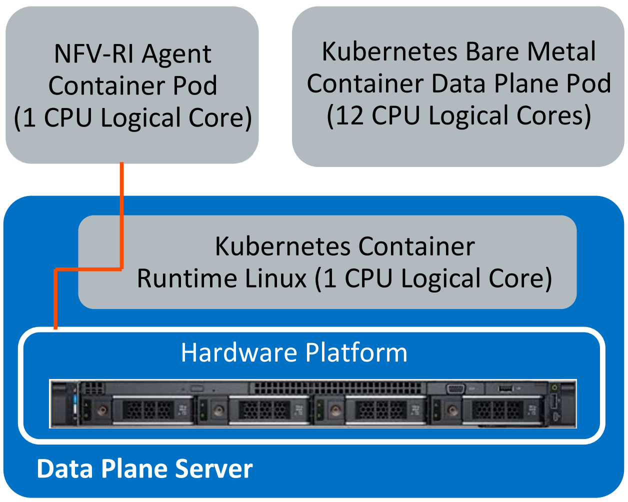

Figure 4. Kubernetes View

Emulated Traffic

Based on KDDI's own traffic model, we emulate the traffic load for a 24-hour period. This traffic is representative of typical data plane traffic, which has the lightest traffic load in the early hours of the morning (for example, 3 to 4 a.m.) and the heaviest traffic in the evening (for example, 6 to 8 p.m.) with varying traffic load levels in between. The maximum traffic load is set at six times the minimum traffic load.

Since by default, a DPDK application always utilizes its CPU cores at 100% due to polling, any time the traffic is not at its highest presents an opportunity to turn down the frequency (and consequently reduce power consumption). This is a key principle that the PoC leverages.

Server CPU Logical Core Allocation

In the PoC, only 12 of the available 80 CPU logical cores (40 per socket) are reserved (or pinned) for use by the 5G data plane. Intracom NFV-RI controls the frequencies of these 12 CPU logical cores. *Figure 5* shows the CPU logical core allocations for the data plane in the PoC.

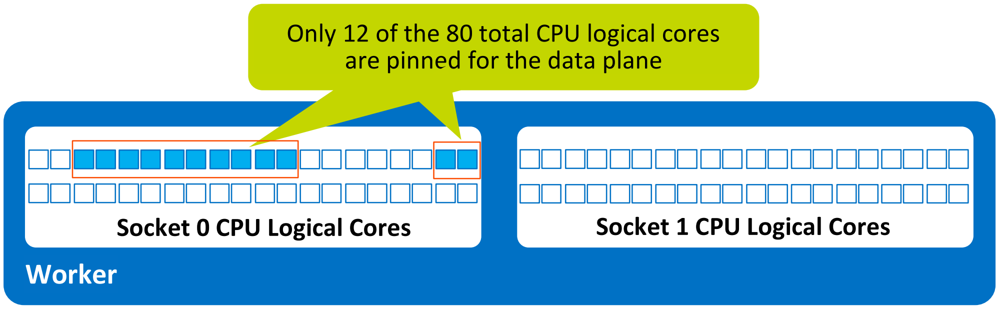

Figure 5. Server CPU Logical Core Allocation for 5G Data Plane

FFL Phases of Operation

The Intracom Telecom NFV-RI uses a Frequency Feedback Loop (FLL) workflow to predict CNF traffic levels and dynamically adjust the frequencies of each CPU core used by DPDK CNFs. The three phases of FFL operation are:

- **Configuration** - The user provides all the parameters needed for configuring FFL in its subsequent phases, like the suitable data sources or the mapping of CNFs to FFL instances.
- **Training** - In the training phase, FFL is trained to find ideal frequencies for different traffic levels. The ideal CPU frequency for a certain traffic level is the minimum CPU frequency that yields zero packet drops. In general, at higher traffic levels, the more CPU frequency is required to keep packet drops at zero. In this phase, the user is initially asked to specify the maximum known capacity of the CNF (Cmax), that is, the maximum amount of traffic that the CNF can sustain in production without errors. Subsequently, FFL attempts to identify the ideal frequencies for a certain number (N) of traffic steps, uniformly distributed in the range between 0 and Cmax, asking the user to feed CNFs with the corresponding traffic rate in each step. The outcome of this phase is a trained model that can predict how traffic is going to vary in the subsequent short-term period and decide the best CPU frequency for the predicted level that avoids packet drops and maintains Service Level Agreements (SLAs).
- **Operation** – In the operation phase, the user launches one or more FFL instances on the machine, one for each different CNF the user wants to control, specifying which trained model should be used for each. This phase may run indefinitely. Each FFL instance dynamically detects changes in traffic load and decides automatically which CPU frequency should be applied to the CNFs as shown in *Figure 6*.

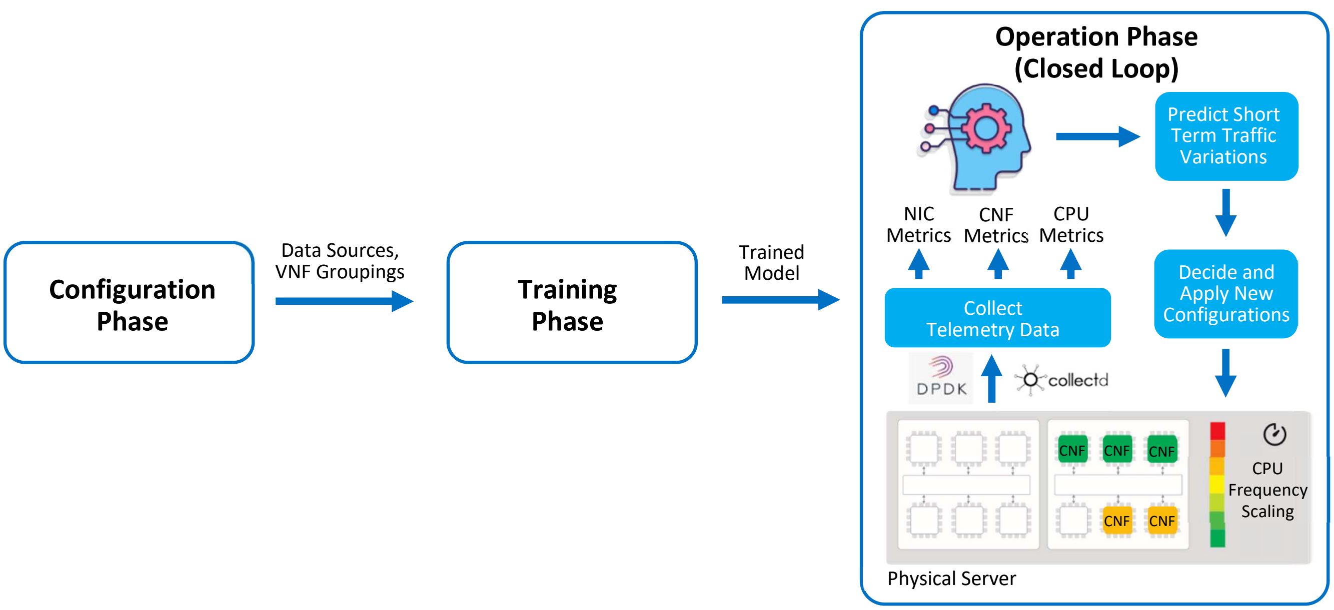

Figure 6. NFV-RI FFL Phases

NFV-RI FFL Agent

A key component in this PoC is the NFV-RI Agent. The Agent monitors:

- Low-level Performance Monitoring Unit (PMU) metrics on each CPU core, which provide an indication of the actual traffic load handled by that CPU core
- A packet drop counter from Prometheus running on the Kubernetes platform

The Agent uses this information to determine the P-state (and accordingly the frequency) to which the CPU core can be set while ensuring zero packet drops.

The Agent reads packet drop counts from Prometheus to close the loop as shown in *Figure 7*. Prometheus gathers packet drops either from the CNF directly, or from the NIC statistics. In our case, we have chosen to use the metric reported by the data plane itself. This closed-loop scheme ensures that the CPU core continually operates at the lowest possible frequency required to handle the predicted traffic load without incurring packet drops.

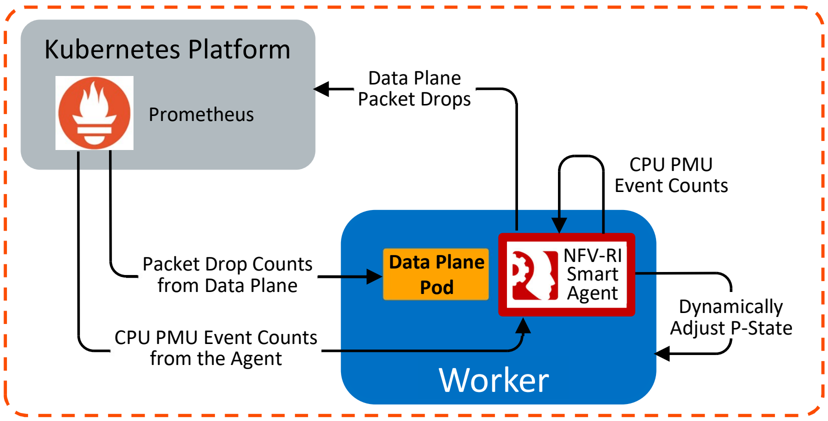

Figure 7. NFV-RI FFL Agent

Results

This PoC demonstrated initial savings of around 30 to 35W with 12 CPU logical cores. However, in a production environment, there are many variables that contribute to overall power consumption, for example:

- The High Availability (HA) and Cluster setup - Greater power savings can be expected in such setups because the traffic load is shared across multiple Data Planes.
- The Average Traffic or Data Plane Load - The higher this value, the less savings can be achieved because there are fewer opportunities to reduce core frequencies.
- The number of vCPUs used by Data Plane and managed by NFV-RI - The higher this number, the greater the power savings because the frequency/power optimization is applied to more vCPU cores.

Conclusion

Intracom Telecom's NFV-RI solution, leveraging Intel server telemetry and power management capabilities, provides significant opportunities for power consumption reduction in KDDI 5G network. This solution can contribute significantly to KDDI's quest to become carbon neutral and meet its CO_2 emission targets by 2030.

KDDI is encouraged by the potential for even greater savings and have expressed interest in exploring this solution further. For example, greater power savings can be expected if more than 12 CPU logical cores are allocated to the 5G data plane. It may also be possible to apply other power saving technologies to unused data plane cores, or unused control plane cores, or both, for even greater power consumption reductions.

References

The following table provides links to information related to the technologies described in this solution brief.

Table 1. References

	TITLE	SOURCE
[1]	Accelerating net zero CO_2 emissions to FY2030	https://news.kddi.com/kddi/corporate/newsrelease/2022/04/07/5984.html
[2]	Intra Telecom NFV Resource Intelligence (NFV-RI)	https://www.intracom-telecom.com/en/products/telco_software/sdn_nfv/nfvRI.htm
[3]	Intracom Telecom Machine Learning Boosts Energy Efficiency of Red Hat OpenShift NFV Workloads	https://networkbuilders.intel.com/solutionslibrary/intracom-telecom-machine-learning-boosts-energy-efficiency-of-red-hat-openshift-nfv-workloads
[4]	Intel® Xeon® Scalable Processors	https://www.intel.com/content/www/us/en/products/processors/xeon/scalable.html
[5]	Intel® 64 and IA-32 Architectures Software Developer's Manual Combined Volumes: 1, 2A, 2B, 2C, 2D, 3A, 3B, 3C, 3D, and 4 – Volume 3A, Chapter 14, Power and Thermal Management, Enhanced Intel SpeedStep® Technology	https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html

1222/CR/DF