

Reducing Energy Use and Carbon Footprint of Open RAN Networks

Aspire Technology profiles the energy reductions of power saving features using 3rd Gen Intel® Xeon® Scalable processors for Open RAN 5G centralized units and distributed units.



Overview

Energy efficiency within mobile telecommunications has been and continues to be a focal area of interest for all in the industry. The estimated electricity consumption of the global operator community in 2021¹ was around 293 million MWh representing approximately 20% of operating expenditures. Led by the GSMA, many operators are seeking ways to shrink their carbon footprint by reducing energy draw and adopting renewable energy.



Traditionally, mobile operators deploy RAN baseband functionality from a cell site using a tightly integrated baseband unit (BBU) comprised of proprietary hardware and software from telecom equipment manufacturers (TEM). In traditional RAN, a BBU vendor selects the 'best' or most suitable component combinations from various hardware manufacturers and builds their baseband units from a mix of those components. This hardware solution runs the baseband application software that has been designed specifically for that hardware. The resulting fixed-function solution is solely focused on acting as a BBU, the vendor has full control over this solution from the components to the RAN application software.

The maturation of Open RAN has helped mobile network operators (MNOs) to embrace a new model that can reduce costs and increase innovation.

Key to this evolution is the standardization of open RAN interfaces and decoupling of hardware and software, so that the traditional baseband application software can now be deployed on commercial off-the-shelf (COTS) hardware. The traditional BBU stack, comprising layer 1, layer 2 and layer 3 functionality has been redefined, and incorporates a new open fronthaul protocol called evolved Common Public Radio Interface (eCPRI), as well as disaggregating the BBU component software.

The disaggregation of BBU components - distributed units (DUs) and centralized units (CUs) - allows the centralization of various RAN functions in data centers and/or localized hubs. It is expected that these new deployment scenarios can leverage existing data center energy efficiency programs. The distribution of network functions across physical server pools will allow for additional potential energy efficiencies. Server pool sizes can be adjusted dynamically with no end user impact.

As MNOs move from traditional hardware to COTS, the industry is searching for the most energy-efficient deployments. COTS hardware is the opposite of bespoke baseband hardware - it is general-purpose infrastructure, built with high-capacity processors, RAM and hard drives without customization for the baseband role it is being used for in Open RAN.

Computational power and resources of centralized DUs required to manage distributed radio units (RUs) can be optimized as utilization fluctuates in the served locations. In low traffic hours, computational power can be considerably lowered if

Table of Contents

- Overview..... 1
- Aspire in Telecom’s Energy Efficiency2
- Finding the Right Processor for 5G Open RAN.....2
- Test Environment and Execution2
- Managing Idle CPU Behavior with C-States4
- Traffic Model5
- Metrics.....5
- Power Consumption Reduction Results5
- L1 BBU CPU Utilization %6
- CPU C-State Residency.....7
- Conclusion.....7

the server resources have the capacity to enter a low consumption mode. However, there will still be a substantial percentage of DUs deployed on site, serving very few cells, where the DU hardware efficiency becomes very important. Energy consumed within the RAN is usually considered across all components (BBU and the radio units), however this paper specifically focuses on the baseband savings.

Aspire Technologies, an Intel® Network Builders Program partner, have independently documented a level of increased energy efficiency on servers using 3rd Gen Intel® Xeon® Scalable processors over prior generation CPUs with a significant reduction in CPU execution for a modular RAN. The output of this project provides the additional energy savings achieved from the power state features that have been introduced.

Aspire in Telecom’s Energy Efficiency

Leaders in RAN Energy Efficiency:

Aspire Technologies has emerged as a pioneer in the domain of energy efficiency within radio access networks (RANs). Aspire’s extensive expertise in optimizing energy consumption within traditional RAN and Open RAN environments reflects the company’s commitment to sustainability while ensuring seamless network operations.

Innovative Solutions for RAN Optimization:

With a team deeply versed in telecommunications and radio software development, Aspire leverages cutting-edge technologies like artificial intelligence (AI), machine language (ML), and predictive analytics to create innovative solutions that not only reduce energy usage but also enhance the performance of RAN and Open RAN deployments. This proficiency uniquely positions Aspire to cater to operators seeking to achieve a balance between cost reduction, environmental responsibility and network efficiency.

A Global Reach with Local Relevance:

Operating on a global scale, Aspire’s engagements with operators worldwide have enriched the company’s understanding of diverse RAN ecosystems. This global

footprint empowers Aspire to tailor solutions that accommodate regional regulations, network infrastructures, and operational requirements, making Aspire Technologies the partner of choice for driving energy efficiency in both traditional and Open RAN networks.

Finding the Right Processor for 5G Open RAN

The scope of this research is to profile and document the energy saving gains achieved with Intel’s power saving features which are available as part of 3rd Gen Intel® Xeon® Scalable processor-based servers when executing Open RAN software, specifically when executing the DU-like workloads of the 5G RAN (gNB).

Test Environment and Execution

The processor energy consumption test environment consisted of an Advantech carrier-grade edge server SKY-8132S-11 (see Figure 1 and Table 1), with a single Intel® Xeon® Gold 6338N processor running a DU-like workload, and an Intel reference board with Intel® Xeon® Platinum 8360Y processor as a load generator. The SKY-8132S-11 is a highly optimized, 1U, ultra-short depth, DU server that has been designed to meet 5G Open RAN performance, sustainability and TCO objectives. The server has been built for high reliability, to minimize costly service downtime and onsite interventions. Advanced features such as redundant BIOS and firmware images, fail-safe remote updates, support for single fan failure or the capability to withstand extreme operating temperatures (-40 to +65 degrees Celsius) make the SKY-8132S-11 a low-risk choice for wide-scale Open RAN deployment.

Advantech’s next-generation SKY-8134S-11 vRAN server based on 4th Gen Intel Xeon Scalable processor with Intel vRAN Boost is now available, keeping the same carrier-grade, ultra-compact, 1U, 300 mm depth design as SKY-8132S-11 but delivering unrivaled density, supporting up to 32 cores and sixteen onboard 25 GbE ports with PTP, SyncE and GNSS for high-precision time synchronization.



Figure 1. Front view of Advantech SKY-8132S-11 server.

This DU-like workload (Table 2) consisted of:

- **L1 workload:** provided by Intel® FlexRAN TestMAC running in timer mode. In this configuration, the L1 application is run in real-time, IQ samples are read in from reference files, loaded into DDR and the L1 application reads from and writes to memory.
- **L2 workload:** provided by Intel® FlexRAN UE and L2 simulators. Uplink and downlink traffic are generated by the Cisco TRex traffic generator and connected to those simulators on the server under test.

An external DC meter was connected between the Advantech server and its power supply unit (PSU) to gather power consumption metrics.

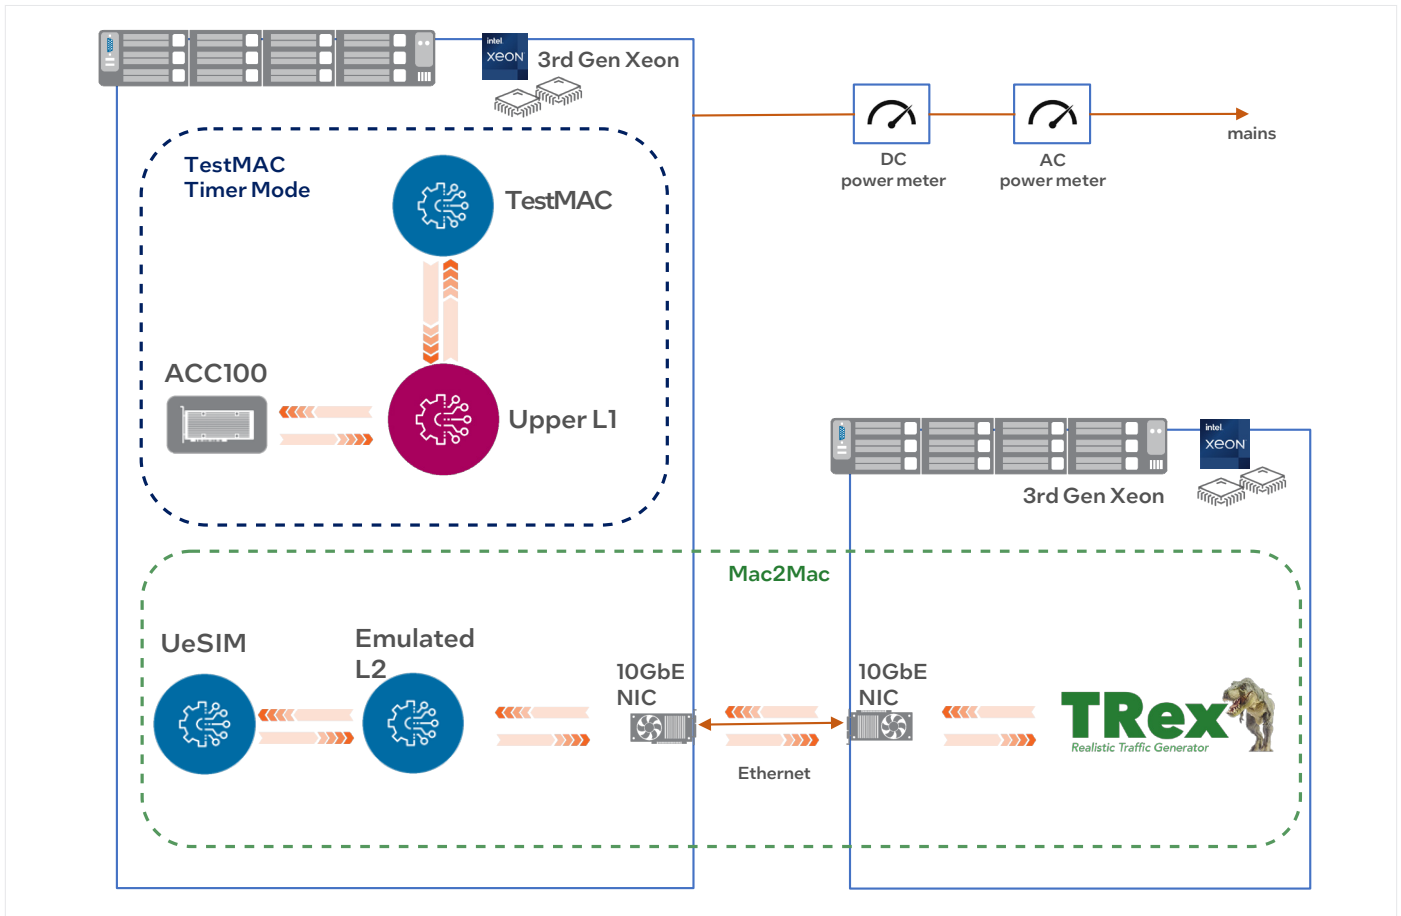


Figure 2. Test environment for L1 and L2 workloads.

System	Hardware	
Server Under Test 1 (SUT1)	Server	Advantech SKY-8132S-11
	CPU	1x Intel® Xeon® Gold 6338N Processor
	FEC accelerator	Silicom Lisbon P1 accelerator card
	NIC	Intel® XXV710 25GbE
Traffic Generator Server	Server	Intel® Customer Reference Board
	CPU	2x Intel® Xeon® Platinum 8360Y Processor
	NIC	Intel® XXV710 25GbE
Measurement	DC Power Meter	Elektro-Automatik PSI 9000
	AC Power Meter	Shelly Pro 4PM

Table 1. Test environment hardware.

Software	Version
CentOS	7.9.2009
DPDK	21.11
FlexRAN	22.03

Table 2. Test environment software.



Figure 3. C-State characteristics.

Managing Idle CPU Behavior with C-States

C-States are CPU-idle power saving states (see Figure 3) that work by adjusting the (internal) clock signal and power from within the CPU in response to traffic levels. The amount of energy saved depends on the reduction in the voltage and the clock signal. The first C-State is C0, which is when the CPU is at a normal CPU operational or active state, where the clock signal is not stopped, and voltage is not reduced. All other C-States (C1-Cn) represent idle sleep states where the processor clock is inactive (cannot execute instructions) and different parts of the processor are powered down. The higher the C-State number, the lower the core voltage and the longer the wakeup time. These states can be applied at a CPU package or individual core level.

- C0: Operating State: the CPU is fully turned on.
- C1: Halt: Stops CPU main internal clock and reduces core voltage.
- C6: Deep Power Down: Reduces the CPU internal voltage further and increases the wakeup time.

A series of C-State configurations were initially executed to choose the optimal test scenario. As expected, C6 state demonstrated the best efficiency in terms of power consumption in those preliminary tests. The C6 state is the deepest possible sleep state configurable on 3rd Gen Intel® Xeon® Scalable processors. In the C6 state, the CPU’s internal voltage is reduced to as low as 0 Volts. The two C-States configurations chosen for this test comparison were as follows:

Baseline = CPU cores associated with FlexRAN L1, L2 and UeSim applications have a max C-State set to C0. All other cores have a max C-State of C6.

SUT1 = All cores have a max C-State of C6. Therefore, in this configuration, cores associated with L1, L2 and UeSim can achieve deeper sleep states than with the baseline configuration.

For this test environment, CPU cores were allocated to the components as shown in Table 3.

Software Components	Number of Logical Cores	CoreIDs
OS/Kernel processes	4	0, 31, 32, 63
FlexRAN layer 1 BBU	10	4-8, 36-40
TestMAC	2	9, 10
FlexRAN layer 1 Others	3	1-3
Layer 2	8	12-18, 20
UE Simulator	6	19, 21-25
Unassigned	31	-

Table 3. CPU core affinities and configuration on the Intel® Xeon® Gold 6338N processor, 64 logical cores.

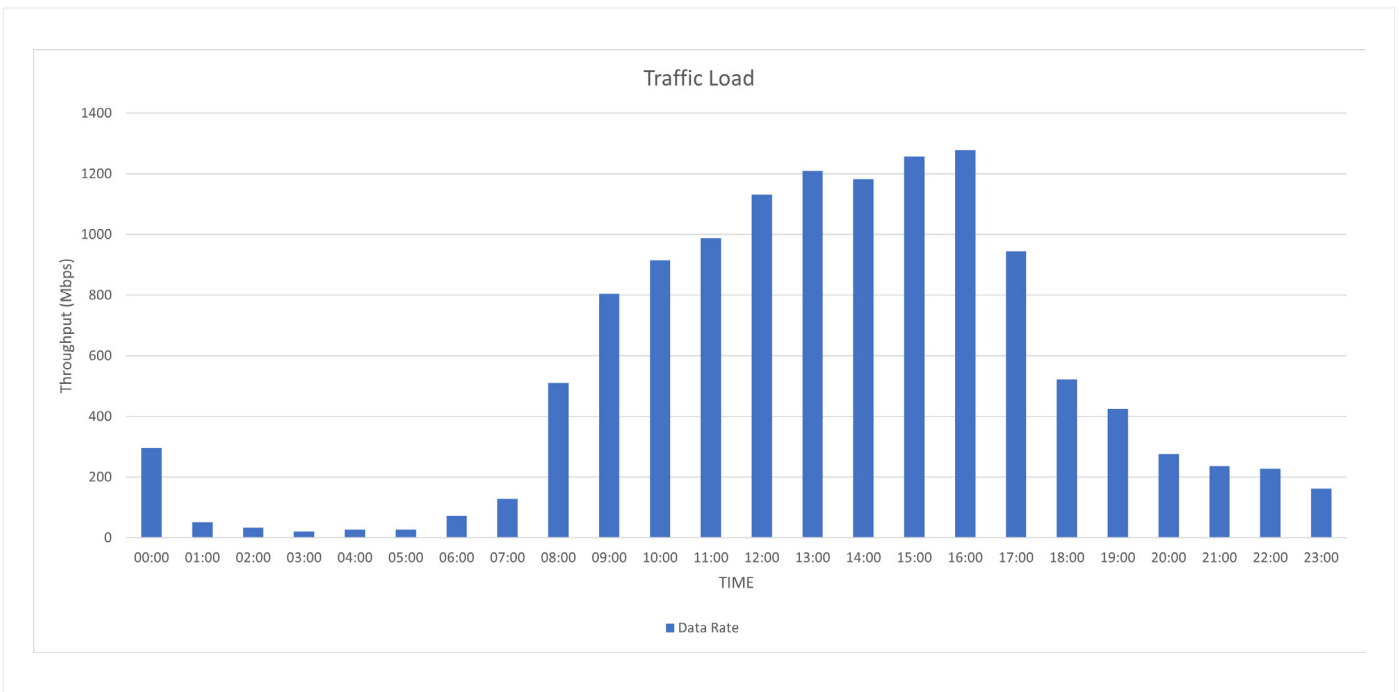


Figure 4. Hourly traffic profile.

Traffic Model

To test a traffic load consistent with a real world environment, a traffic profile representing a dense urban 5G NR environment was chosen (see Figure 4). A modelling interval of one hour was chosen, giving 24 individual traffic profiles per day.

During testing of each one-hour profile, the following was executed:

- Downlink Mac2Mac testcase, providing the L2 workload on the server. This involved running a Cisco TRex traffic generator providing load to a simulated L2/UeSim application on SUT1. (See figure 2) DL throughput was set based on the chosen traffic profile.
- Full-duplex TestMAC timer mode testcase. This provided the L1 workload. The TestMAC testcase was matched as closely as possible to the L2 test in terms of configuration and throughput speeds.

Each of the testcases was run for five minutes. This was sufficient as the power requirement remained constant throughout the test.

Metrics

The following metrics were collected while executing the testcases:

- AC and DC power (W) – using external meters
- L1 BBU CPU utilization (%) – using *mpstat* tool
- C-State residency (%) – using *turbostat* tool

Power Consumption Reduction Results

The DC power consumption was consistently lower on SUT1 than the baseline.

- The average power required on SUT1 was 27.9W lower than the baseline (a 9.2% reduction). This corresponds to a reduction in power consumption of 0.67kWh for a 24-hour period. The 24-hour DC power consumption for the baseline was 7.31kWh, while SUT1 consumed 6.64kWh. This represents substantial power consumption savings achieved by applying the C-State energy efficiency feature in an COTS deployment (see Figure 5).
- The largest reductions occurred during periods of low traffic as expected, as the CPU cores spend more time in an idle state and thus have more opportunity to reach deeper sleep states.

With no L1, L2, or UeSim applications running on the server, the DC power requirement was 205.6W on baseline and 162.5W on SUT1, a reduction of 21%.

The AC power consumption is dependent on the PSU device. In our testing, the difference between DC and AC power consumption remained relatively stable, with AC consumption being 0.07kWh higher than DC, on average.

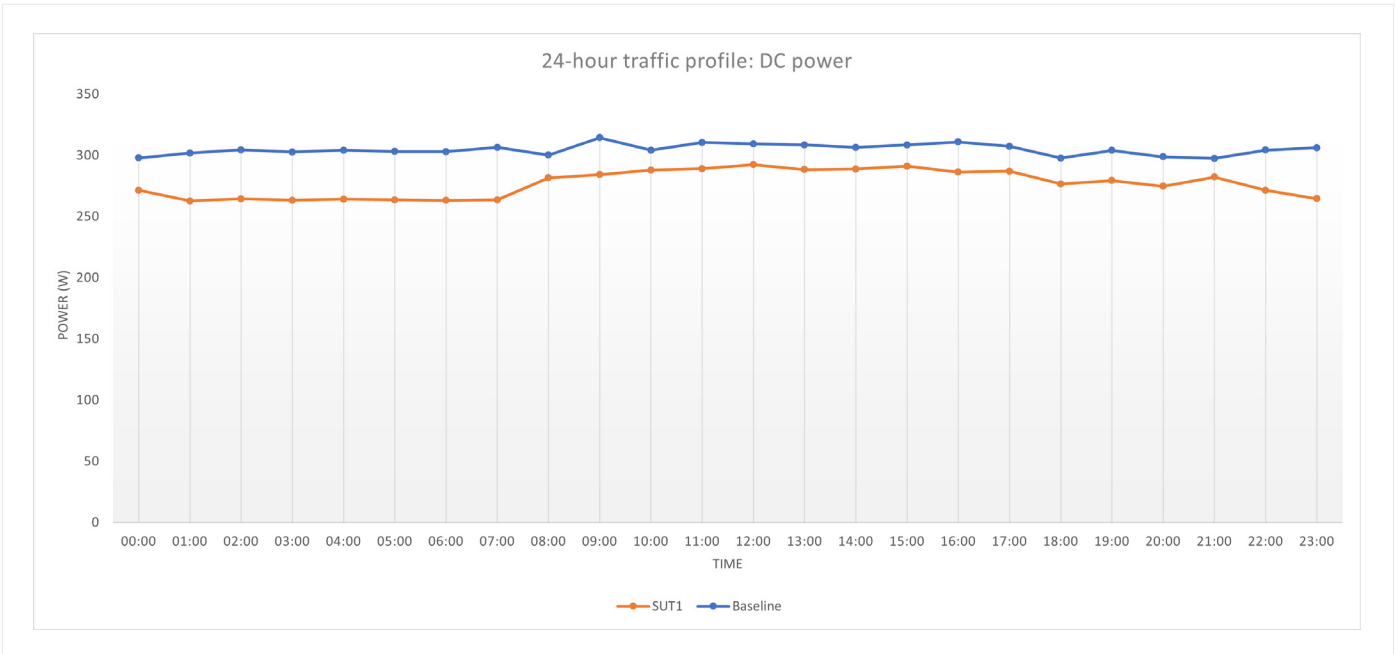


Figure 5. DC power.

L1 BBU CPU Utilization %

The FlexRAN layer 1 BBU threads were run on cores 4-8 and 36-40. The cores corresponding to those BBU threads on SUT1 had an average 10.4% lower CPU utilization than on baseline. The range of reduced utilization was between 2% and 28%. The periods of lower traffic had the largest reduction in CPU utilization. This similar behavior was observed in DC power consumption (see Figure 6).

SUT1 always had a lower usr% utilization than in the baseline configuration for the L1 BBU threads. This indicates that the use of deeper sleep states gives a reduced CPU utilization within this application, regardless of traffic load. No degradation in L1 performance was observed when running in deeper sleep states (SUT1) compared to the baseline configuration.

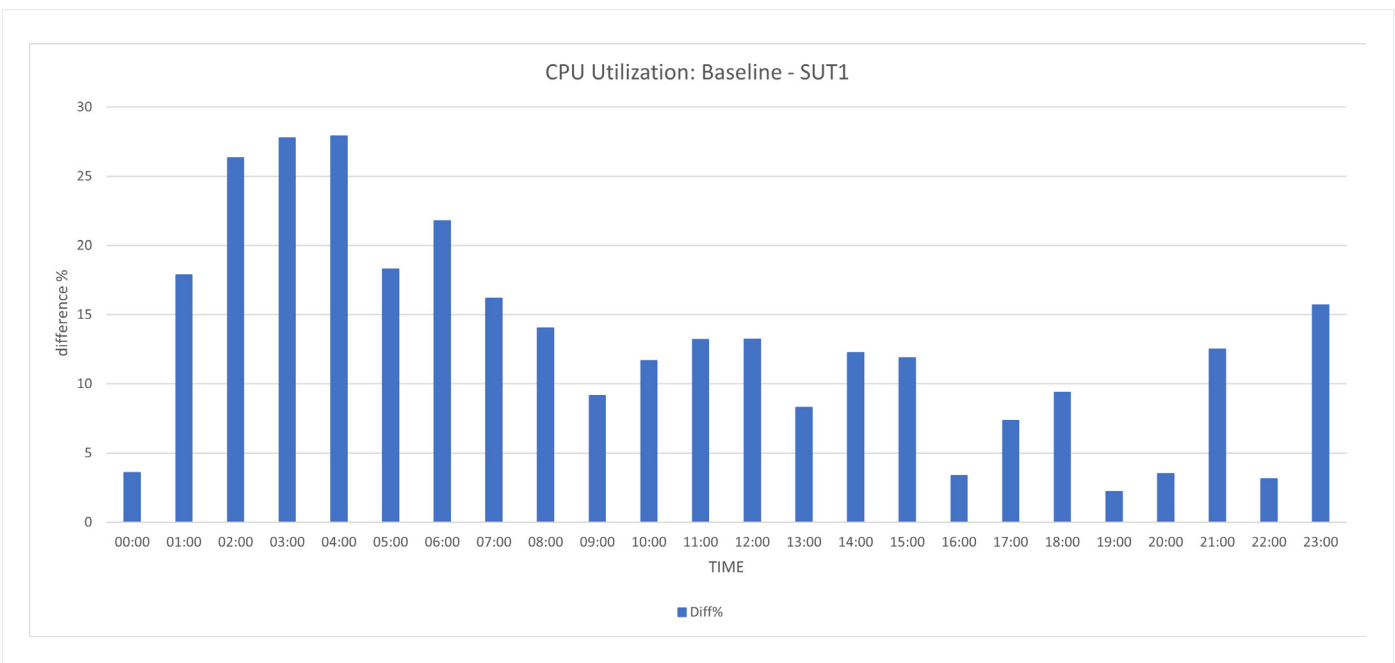


Figure 6. L1 BBU CPU utilization percentage.

CPU C-State Residency

CPU C-State residency shows what sleep states a given CPU core spends time in when idle. In this investigation, only C-States C0, C1 and C6 were possible to achieve. The analysis was focused on the **SUT1** scenario, where CPU cores assigned to applications could achieve a max C-State of C6. The **Baseline** scenario limited application cores to a max C-State of C0.

The power consumption was lower while in C6 state. There was substantially more usage of C1 state than C6 state, leaving considerable scope for potential optimization to push cores from C1 to C6 state more frequently, reducing the power consumption even further.

For FlexRAN L1 BBU assigned cores, we observed a 15%-31% C6 usage and a 31%-51% C1 usage. The DC power reduction in these same test cases (comparing **baseline** with **SUT1**) was in the range 5.6-13.2W. Unassigned cores demonstrated the highest C6 percent usage as expected:

- C6% = 40-50%
- C1% = 50-60%.

Conclusion

The use of deeper sleep states (C-states) does provide considerable energy consumption savings while running a DU-like workload on an 3rd Gen Intel® Xeon® based server.

In **SUT1**, there was less CPU usr% utilization running the same workload as **baseline**, indicating compute efficiency was improved with the use of C-States.

Low traffic periods provided the best energy savings (and CPU utilization). This is possibly due to the fact that as CPU utilization running the L1 and L2 workloads is reduced, there is more scope for idle cores to achieve deeper sleep states. Understanding this behavior fully will require further investigation.

Further optimization of application software to utilize deeper sleep state functionality could bring even more energy efficiencies. A study of additional configurations and traffic profiles would provide useful input for this and other optimization activities.

Aspire remains dedicated to advancing energy efficiency in the realm of telecommunications. The ongoing efforts focus on harnessing emerging technologies like CPU deep sleep states, refining optimization techniques, and collaborating closely with industry leaders to continuously push the boundaries of energy savings.

Aspire Technology delivers network solutions, open networks and consulting globally to fixed and mobile operators, vendors, system integrators and technology partners across the whole network lifecycle. Aspire started as specialized company in traditional mobile network technologies following the strong legacy of initial staff in software engineering and associated services for radio access network (RAN) including software development and support. Over time, Aspire built expertise in other network domains and technologies including Open RAN and Telco cloud for various vendors, and is currently delivering projects globally, with an international expert team of network and software engineers all around the world. Aspire built Open RAN expertise over the past four years with its Open Networks Lab in Dublin featuring over 30 vendors across core, CU/DU, RU, NFV, IT and test. For further information please visit www.aspiretechnology.com.

Learn More

[Aspire Technologies](#)

[Intel® Xeon® Scalable processors](#)

[3GPP](#)

[O-RAN Alliance](#)

[Advantech SKY-8132S-11](#)

[Intel® Network Builders](#)



¹https://www.gsma.com/betterfuture/wp-content/uploads/2022/11/Mobile_Industry_Position_Paper_Access_to_Renewable_Electricity_Nov22.pdf

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more on the [Performance Index site](#).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel technologies may require enabled hardware, software or service activation.

Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

0923/DC/H09/PDF

Please Recycle

356948-001US