

Wind River, Intel Solution Supports More Cloud Native RAN Services with Less Compute

Wind River Studio is distributed cloud software for virtualized and containerized workloads; the combination of 4th Gen Intel® Xeon® Scalable processor powered servers and Wind River Studio enables two-times capacity increase¹ while running on a single physical core



5G is entering a new era of performance, low latency and services as more mobile network operators (MNOs) support standalone 5G networks.² Standalone networks move away from 4G core networks to embrace a complete 5G network infrastructure from radio access network (RAN) to core. Investing in standalone 5G networks necessitates using telecom cloud and edge infrastructure to build out a nationwide network. Making the most of this new infrastructure will require a dramatic increase in small cell and macro base stations to deliver seamless coverage.



To support these base stations, MNOs will need to maximize coverage efficiency and simplify management to fully capitalize on the new standard and maintain a suitable level of cost efficiency. Some potential challenges with this roll out include:

- **High total cost of ownership:** With thousands of base stations needed due to reduced transmission radius of higher-speed wireless technologies, it's possible that deployment, maintenance, and power costs could increase significantly.
- **Coverage efficiency:** With more RAN nodes, fine tuning the coverage to deliver a seamless mobility experience to the subscriber while using the fewest possible hardware resources is a challenge.
- **Manageability:** Operators need automation and orchestration features to manage the extensive scale of a distributed cloud deployment.

Solving these challenges falls to the cloud platform that delivers service deployment, management, automation and analytics. Wind River, an Intel® Network Builders ecosystem Titanium member, offers its Wind River Studio that offers a comprehensive solution to deploy, manage and optimize 5G networks.

Wind River has optimized the Wind River Studio Cloud Platform control plane to run on the latest 4th Gen Intel® Xeon® Scalable processor. These optimizations include:

- Shrinking the platform footprint from an industry leading two cores down to a single core
- Leveraging Intel® Advanced Vector Extensions (AVX) instructions on 4th Gen Intel® Xeon® Scalable processors
- Addition of Kubernetes Operator for management of Intel® vRAN Boost integrated acceleration

The combined benefit of these optimizations include:

- Delivering two times increase¹ in capacity and,
- Approximately 20% power reduction compared to 3rd Gen Intel® Xeon® Scalable processors, and a
- 50% reduction in minimum platform core consumption for Wind River Studio.

Table of Contents

Wind River Studio Delivers Efficient Distributed Cloud Operations.....2

4th Gen Intel® Xeon® Scalable Processors.....3

Built-in Accelerators Improve Performance4

Test System Set Up4

Comparing CPU Performance Improvements5

Conclusion.....6

Edge Compute: A Limited Resource Drives Efficiency

Compute power at the edge is in increasingly high demand but unlike a data center with easy access to additional compute, it is a limited resource that must be allocated thoughtfully. For illustrative purposes, if an MNO decided to deploy multiple cells of a 200Mhz 64T64R massive MIMO RAN deployment on a 4th Gen Intel® Xeon® Scalable-based 32 core server, each core used for software infrastructure reduces the number of cores available for commercialization headroom and feature growth by 8.33%.

In this scenario, if the MNO considers L1, L2, OAM/fronthaul as relatively fixed and consuming 20 cores, this leaves 12 open cores. With Wind River Studio Cloud Platform, the MNO can consume as little as a single core for CAAS, leaving 11 (92%) to utilize to increase the capacity per cell site, including:

- Increasing the number of users served.
- Increasing the throughput or data speeds for existing users.

- Reducing additional hardware cost and power consumption by adding new applications such as a vRouter to the same server.

In terms of feature growth, the extra core capacity allows OEMs/RAN vendors to add new features without needing additional servers or optimizations.

In contrast, with this same scenario, a non-optimized CAAS solution that requires 4-6 cores reduces access to compute power by between 33%-50%.

As seen in Figure 1, with Wind River Studio, less compute power is needed to deliver more capacity with higher performance. The remaining compute power can be used for additional capacity, feature growth, and/or new revenue generating services, simultaneously reducing overhead and increasing revenue opportunities. These improvements are seen in tests of a demonstration platform that are documented in this paper (see tests details below and server specifications in Table 1).

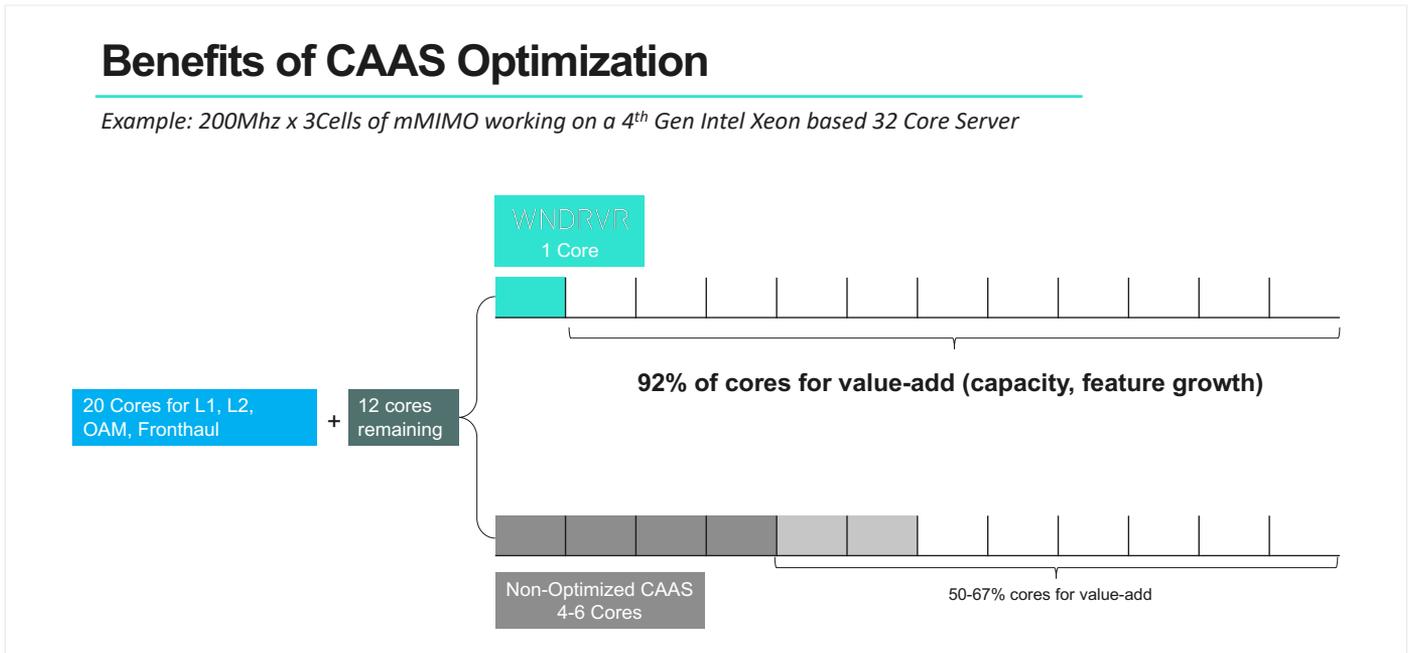


Figure 1. Illustration of benefit of Wind River Studio's reduced compute footprint.

Wind River Studio Delivers Efficient Distributed Cloud Operations

Wind River Studio is the first cloud-native software for the development, deployment, operation, and servicing of ORAN/vRAN systems. The software also works with other mission-critical intelligent edge systems that require security, safety, and reliability to execute on a single core (with 4th Gen Intel® Xeon® Scalable processor-based servers). Wind River Studio is architected to deliver digital scale across the full lifecycle through a single pane of glass to accelerate transformative business outcomes.

Wind River Studio delivers an integrated cloud platform that unifies infrastructure, orchestration, and analytics capabilities. This platform enables mobile network operators (MNOs) to deploy and manage globally distributed 5G edge networks, including remote servers that process vRAN traffic. The core capabilities of Wind River Studio include:

Cloud Platform: A carrier-grade Kubernetes cloud platform for managing distributed cloud infrastructure. Based on the open source StarlingX project, Wind River Studio compiles optimal open-source technology to deploy and manage distributed networks.

Conductor: Comprehensive orchestration and automation capabilities provide one platform to achieve multi-cloud automation and zero-touch operation. MNOs can use the app catalog to select applications, deploy them to a carrier-grade cloud platform, and orchestrate the resources needed for the applications at the edge network site. This orchestration allows scalability from a handful of nodes to thousands of nodes in a geographically dispersed, distributed environment.

Analytics: Using machine learning algorithms, Wind River Studio supports effective management of a distributed cloud system by consuming and processing data and producing meaningful insights for decision making. Wind River Studio uses full stack monitoring of the cloud infrastructure cluster to collect, analyze, and visualize cloud behavioral data to improve uptime and optimize operations.

The combination of the automation features in Conductor and Analytics for decision making allows for very efficient operations in terms of fast response to network issues without waiting for technician intervention. The integration of Wind River Studio with Intel® architecture CPUs results in a tightly optimized and built-for-the-edge solution that delivers efficiencies, lowers risks, enhances agility, and expedites time to market.

In a vRAN/ORAN deployment with an Intel® architecture compute infrastructure at each site in the network (Figure 2), Wind River Studio offers a single geo-distributed cloud operating multiple, individual distributed clouds with centralized management.

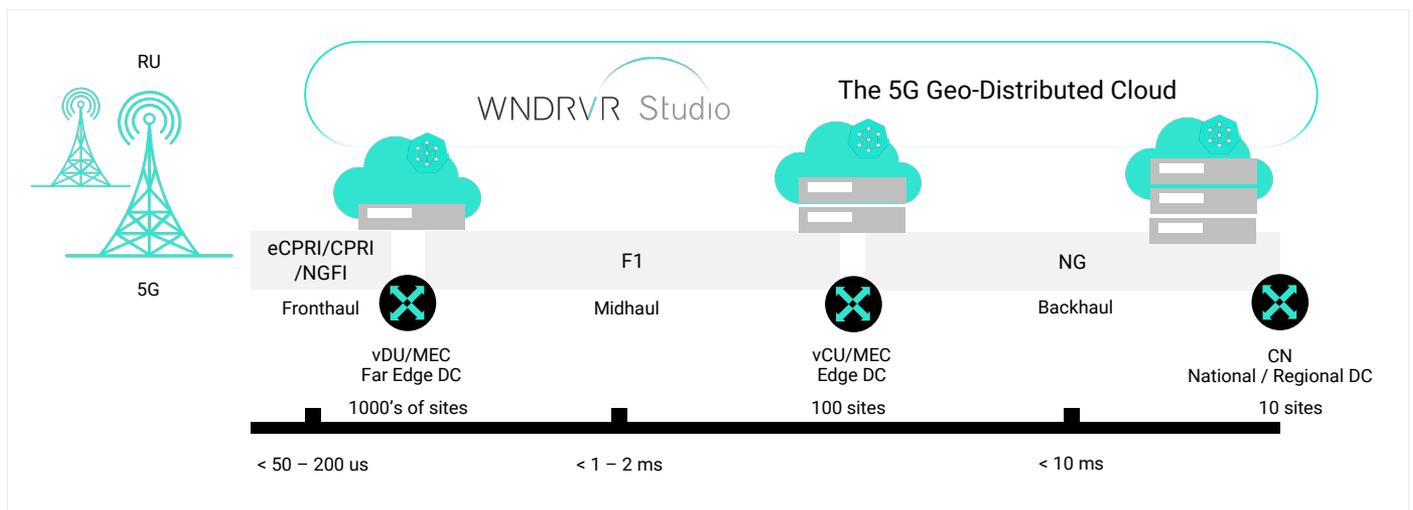


Figure 2. Wind River Studio Cloud Platform deployed as Individual distributed clouds with centralized management.

4th Gen Intel® Xeon® Scalable Processors

The 4th Gen Intel® Xeon® Scalable processor is designed to accelerate performance across the most demanding workloads. The new processor architecture combines high-performance processor cores with up to eight built-in accelerators (see Figure 2)³ to help maximize performance efficiency for demanding workloads.

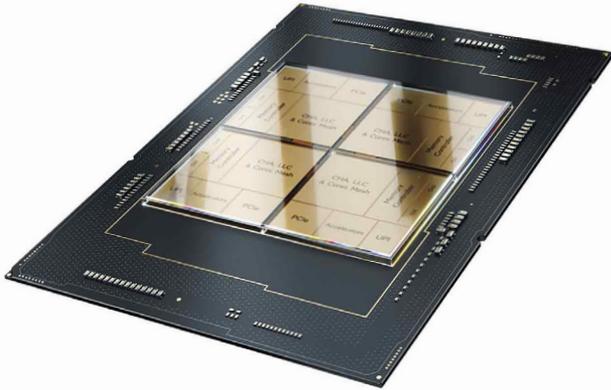
These latest processors offer a breakthrough in I/O capabilities, with up to 80 lanes of PCIe 5.0 connectivity and Intel® Scalable I/O Virtualization for more virtual machines

per server. The CPUs also support Compute Express Link (CXL), a cache-coherent interconnect for processors, memory expansion, and accelerators. Figure 3 summarizes the features of these CPUs.

Network-optimized 4th Gen Intel® Xeon® processors are the next step in accelerating load balancing workloads while increasing energy efficiency, with a high-throughput, low-latency platform engineered for data centers, network core, and scalable to the edge for on-prem or cloud deployments.



Purpose Built for Accelerated Network Workloads



4th Gen Intel® Xeon® Scalable processors² offer:

Advanced execution resources in a range of core counts and feature sets, with improved per-core performance compared to the previous generation, enhanced by the most built-in accelerators in the industry.

Enhanced memory subsystem with up to eight DDR5 channels operating at up to 4800 MT/s, a 1.5x improvement in memory bandwidth and speed compared to the predecessor platform.

Faster, higher capacity I/O based on up to 80 lanes of PCIe 5.0 per socket, compared to 64 lanes of PCIe 4.0 per socket in the prior generation.

Figure 3. 4th Gen Intel® Xeon® Scalable processor networking features.

Built-in Accelerators Improve Performance

Integration of accelerators into the processor redefines CPU architecture. Using accelerators provides a more efficient way to achieve higher performance than relying solely on increasing the CPU core count for workload processing.

With all-new accelerated matrix multiply operations, 4th Gen Intel® Xeon® Scalable processors have exceptional AI training and inference performance. Other seamlessly integrated accelerators speed up data movement and compression for faster networking, boost query throughput for more responsive analytics, and offload scheduling and queue management to dynamically balance loads across multiple cores. To enable new built-in accelerator features, Intel supports the ecosystem with OS-level software, libraries, and APIs. This architecture allows users to run cloud and networking workloads using fewer cores with faster cryptography.

Test System Set Up

Performance testing was done on two systems under test (SUT) that were configured as an open RAN system to showcase the optimizations made in Wind River Studio.

The company worked closely with Intel to optimize resource consumption in Wind River Studio running on 4th Gen Intel® Xeon® Scalable processors, which resulted in the ability to run the entire cloud software platform in a single core – half the core count required for software to run on 3rd Gen Intel® Xeon® Scalable processor. This frees up a core for other revenue generating applications.



Figure 4. Key accelerators built into the 4th Gen Intel® Xeon® Scalable processor.

Comparing CPU Performance Improvements

For the tests, conducted by Intel in March 2023, two servers were configured to provide vRAN distributed unit (DU) functionality which includes antennas that supported massive multiple input, multiple output (mMIMO) transmission capabilities including 64 transmit/receive (64T/64R) capabilities.

One server (3rdGenSUT) was based on 3rd Gen Intel® Xeon® Scalable CPU with 32 cores (30 free cores after loading Wind River Studio) and featured an Intel® eASIC™ structured ASIC (Intel® vRAN Accelerator ACC100) to accelerate forward error correction (FEC) performance.

The performance from this server was compared to a second server (4thGenSUT) based on 4th Gen Intel® Xeon® Scalable Processor with 32 cores (31 free cores after loading Wind River Studio). This SUT used the new 4th Generation Intel® Xeon® platform with Intel® vRAN Boost, which eliminates the need for an external accelerator card by integrating vRAN acceleration into the CPU, thereby reducing power consumption and cost.

Both servers used Intel’s FlexRAN™ reference software as the workload. FlexRAN™ is 4G and 5G PHY reference software used for product development and testing.

The tests utilized a distributed cloud built using Wind River Studio Cloud platform 22.12 with two single-server sub-clouds comprised of the two SUTs.

The central region controller (CaaS Manager), Wind River Studio Analytics and Studio Conductor were also a part of the test and ran on two additional 3rd Gen Intel® Xeon® Scalable processor-based servers.

Studio Conductor was used to deploy the workload that ran on each of the servers. The workload was run on each of the servers in parallel which then measured the performance. The platform and workload performance KPIs were collected by Wind River Studio Analytics in real time from the two sub-clouds, brought to the central region and displayed side by side on a Kibana dashboard showing the performance gain and platform overhead reduction running the exact same workload on both SUTs.



Figure 5. Performance results from the Wind River Studio tests (higher is better) The 3rd Gen Intel® Xeon® Scalable processors support 20 radios on two cores, whereas the 4th Gen Intel® Xeon® Scalable processor supports 32 radios on one core.

Figure 5 shows the performance results. The 3rdGenSUT was able to deliver up to 3x100MHz 64T64R and support for up to 20 radios and the 4thGenSUT doubles that performance with up to 6x100MHz (2pk, 4avg) with support for up to 32 radios. These tests demonstrate that Wind River Studio Cloud Platform enables key performance differentiators in 4th Gen Intel® Xeon® Scalable processors, including the new 5G instruction set, integrated accelerators, and increased per core-performance.

The combination of Wind River Studio Cloud Platform and 4th Gen Intel® Xeon® Scalable processors enables MNOs to achieve up to two times capacity increase for vRAN workloads while at the same time reducing the software's footprint from a minimum of two physical cores to one physical core, an additional performance increase enabling support for a larger number of radio units or other applications.

Conclusion

MNOs will need to build out edge networks to support 5G standalone networks driving an increased demand for an edge cloud platform that is light weight, full featured with high throughput. Wind River has optimized its Wind River Studio Cloud Platform to meet these needs with the entire platform running in a single core of the 4th Gen Intel® Xeon® Scalable processor. In that configuration, the software ran up to three times faster than on a server based on 3rd Gen Intel® Xeon® Scalable processors. These performance and footprint improvements allow MNOs to have more capacity on their servers for revenue generating services.

Learn More

[Wind River in Telecommunications](#)

[Wind River Studio for Operators](#)

[4th Gen Intel® Xeon® Scalable processor](#)

[Intel Network Builders](#)



¹ Tests conducted by Intel in March 2023 using system under test defined in Table 1.

² <https://www.fiercewireless.com/5g/t-mobile-rolls-out-5g-standalone-25-ghz-spectrum>

³ <https://www.intel.com/content/www/us/en/products/details/processors/xeon/scalable.html>

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.

Your costs and results may vary.

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

© 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.

4Gen SUT Server Configuration

The server configuration is presented below:

- Processor:
 - Model: Intel® Xeon® Gold 6421N
 - Number of processors / Physical cores per processor: 1 / 32
- Memory:
 - Type / Amount: DDR5 / 128GB (8x16GB)
- Disk Size / Type: 1TB / SSD
- Network Cards:
 - Intel Corporation Ethernet Controller X710 for 10GBASE-T
 - Intel Corporation Ethernet Controller E810-C
- Accelerator (for FlexRAN environment):
 - Maclaren Summit Intel® vRAN Accelerator ACC100
- BIOS configuration:
 - Hyper-Threading: Enable
 - SpeedStep (Pstates): Enable
 - Boot Performance Mode: Max Performance
 - Energy Efficient Turbo: Disable
 - Turbo Mode: Enable
 - Hardware P-States: Disable
 - Enable Monitor Mwait: Enable
 - CPU C1 Auto Demotion: Disable
 - CPU C1 Auto unDemotion: Disable
 - Enhanced Halt State (C1E) Disable
 - CPU C6: Enable
 - Package C State: C0/C1 state
 - Dynamic L1: Disable
 - Power Performance Tuning: OS Controls EPB
- WRCP Configuration:
 - WRCP 22.12 GA release
 - Number of Physical Cores/Logical Cores for platform function: 1 / 2

3Gen SUT Server Configuration

The server configuration is presented below:

- Processor:
 - Model: Intel® Xeon® Gold 6338 CPU;
 - Number of processors / Physical cores per processor: 2 / 32
- Memory:
 - Type / Amount: DDR4 / 192GB (6x32GB)
- Disk Size / Type: 2TB / SSD
- Network Cards:
 - Intel Ethernet Controller X710 for 10GBASE-T
 - Intel Ethernet Controller X520-2
- Accelerator (for FlexRAN environment):
 - Intel® vRAN Accelerator ACC100
- BIOS configuration:
 - Hyper-Threading: Enabled
 - SpeedStep (Pstates): Enabled
 - Boot Performance Mode: Max Performance
 - Energy Efficient Turbo: Disabled
 - Turbo Mode: Enabled
 - Hardware P-States: Disabled
 - Enable Monitor Mwait: Enabled
 - CPU C1 Auto Demotion: Disabled
 - CPU C1 Auto unDemotion: Disabled
 - Enhanced Halt State (C1E) Disabled
 - CPU C6: Enabled
 - Package C State: C0/C1 state
 - Dynamic L1: Disabled
 - Power Performance Tuning: OS Controls EPB
- WRCP Configuration:
 - Software Version: WRCP 22.12 GA release
 - Number of Physical Cores/Logical Cores for platform function: 1 / 2

Table 1. SUT Specifications