Red Hat

Sustainable service providers

75% of a service provider's power consumption comes from their RAN.

Executive summary

Sustainable computing is a multidimensional challenge across a variety of industries. Service providers specifically are keen to adopt technology and techniques to reduce their power consumption and to support their sustainability goals, particularly within the radio access network (RAN) as this accounts for approximately 75% of their total power consumption. Energy and CO_2 optimizations in the carrier networks can be categorized within node, cluster, system, and domain dimensions from the granular control of hardware and its parameters, specialized schedulers and autoscalers, to holistic orchestration with smart workload placement using artificial intelligence and machine learning (AI/ML). All this takes into account unique system features and functionality applicable to a specific technology or domain.

Sustainability for net-zero operations

Service providers are under pressure to establish sustainability goals being demanded by different stakeholders, from generations more conscious of their impact on the environment to governments aligning to protect the planet for future generations.

National goals for reduction of CO_2 emissions and regulations mandate service providers move towards carbon neutral operations within the next few years, with net-zero operations as their end goal. New overarching regulations require public listed companies to have standardized environmental, social governance (ESG) reports. In the US, companies need to comply with new proposed rules enhancing and standardizing climate related disclosures by the security and exchange commission (SEC), and in the EU companies must comply with the reporting requirement set by the European Union commission's corporate sustainability reporting directive (CSRD).

Sustainability is the environmental pillar of an ESG report. ESG is based on standards set by policies, while sustainability uses a science-based approach to define standards for measuring environmental sustainability. Even though the two have some overlap, a good ESG performance does not guarantee sustainability, nor does having sustainability guarantee good ESG performance.

Technology and sustainability

Sustainability is a cross-industry concern where each organization faces its own specific challenges. When considering technology services, service provider goals tend to be aligned with CO_2 emissions and energy consumption. For the former, it is achieving carbon neutral and net-zero emission operations by prioritizing renewable energy, and for the latter, it is reducing power consumption and optimizing for energy efficiency.

CO_2 emissions

CO_2 emissions constitute 79% of greenhouse gasses.

The reduction in CO_2 emissions is a multifaceted challenge. The greenhouse gas protocol (GHG)[1] sets the standard for measuring and managing emissions. To reduce greenhouse gas, we can focus on reducing CO_2 emissions, which make up 79%[2] of the overall greenhouse gasses.

As shown in Figure 1, the GHG protocol defines three emissions scopes, each extending the dependencies beyond the direct organization:

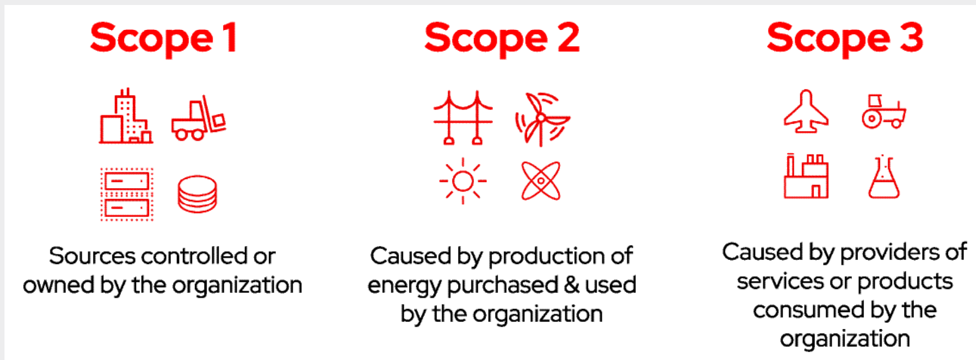

Figure 1. GHG protocol

- **Scope 1** emissions are direct greenhouse gas (GHG) emissions that occur from sources directly controlled or owned by the service provider. These include company facilities (e.g., datacenters, buildings, offices, remote locations) and company equipment (e.g., servers, generators, company vehicles).
- **Scope 2** emissions are indirect greenhouse gas emissions caused by the production of energy purchased and used by the service provider.
- **Scope 3** emissions refer to the emissions created during production and life cycle management of products (e.g., software, hardware) and services (e.g., transportation, cloud) consumed by the service provider.

This means that reduction in CO_2 emissions across the three scopes has direct dependencies on the ability to track energy sources and location closely as well as in close collaboration with the suppliers of services and products.

Reduction in CO_2 emissions can be achieved within a service provider's network at the node, cluster, system, and domain level.

Reduction in CO_2 emissions can be analyzed as a byproduct of the four dimensions discussed below (node, cluster, system, and domain level) augmented by relevant metrics. Achieving a carbon neutral operation or a net-zero emissions operation, the power or energy optimizations techniques should be enhanced with metrics identifying the energy sources, for example, a percentage of energy from renewable versus fossil sources for each node.

1 *Greenhouse Gas Protocol. "We set the standards to measure and manage emissions." accessed 7 Nov. 2022.*

2 *EPA. "Overview of greenhouse gasses." 16 May 2022.*

By augmenting the metrics considered by the optimizers with information pertaining to CO_2 emissions, the same techniques can be used within four dimensions, considering the CO_2 emissions information to determine the optimization to apply.

A key hurdle is the lack of standardized application programmable interfaces (APIs) or data sources from which to obtain the information of Scope 2 CO_2 emissions from the power consumed by the nodes at a particular time in a specific location. While there are several providers of these metrics with various degrees of accuracy and there are even some providers with certain metrics available at no cost, they typically have limited geographical coverage and resolution.

Open source communities are working together on the sustainable computing protocol to create carbon awareness metrics.

There are also efforts underway in the community, such as the sustainable computing project, on a protocol to get this information in a cloud-native way or via the green software foundation with their carbon aware software development kit (SDK), which helps service providers bring carbon emission awareness to their software. Even with these community efforts, there is still a lack of nationwide or multinational metric providers or a single standard to follow. This is an area the industry should work on together to have standardized APIs, metrics sources, and metrics resolution.

Energy consumption

Service providers strive to reduce their power consumption and are regularly looking for ways to improve energy efficiency within their network.

According to the TMForum, service providers account for about 3% of the total power consumption of humanity.[3] As shown in Figure 2, approximately 75%[4] of that is consumed by RAN. This represents a significant portion of the operational expenditure (OPEX) of service provider networks.

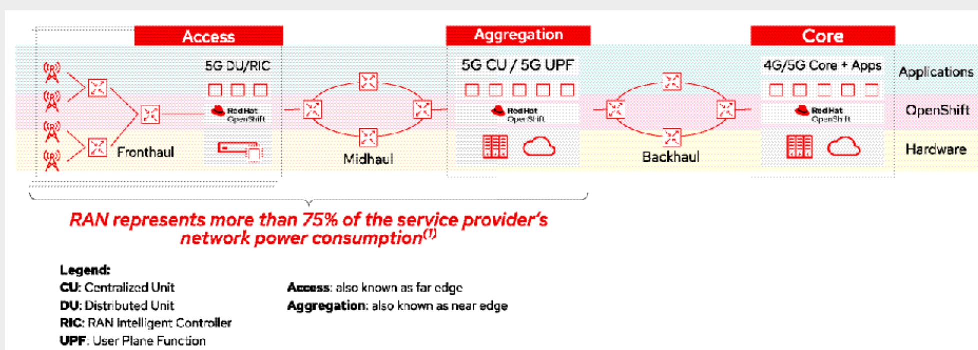

Figure 2. Typical service provider network power consumption

Service providers have a vested interest in adopting technology and techniques that support their sustainability initiatives and to drive the reduction of their energy cost while maintaining regulatory compliance.

Under this frame of reference, we identify four dimensions for optimizations of energy consumption and CO_2 emissions:

3 *"Can the telecoms industry power down its impact on the environment?" Inform Tmforum, 19 Aug. 2021.*

4 *"A holistic approach to address RAN energy efficiency." Ericsson, 16 Dec. 2021.*

- **Node level:** refers to the optimizations that can be achieved through innovations of central processing units (CPUs) and their architectures, types of hardware, for example, power consumption by smart network interface cards (SmartNICs), hardware accelerators and graphic processing units (GPUs), or the fine tuning of node level parameters, for example, core and memory frequencies, and disabling cores.
- **Cluster level:** refers to optimizations that are available or applied holistically to a cluster. For example, energy aware schedulers and de-schedulers, energy aware cluster, or pod auto scalers.
- **System level:** represents the entire clusters and elements involved in delivering a service. For example, all the clusters, switches, routers, and antennas that comprise the mobile network.
- **Domain level:** refers to a specialized service or system functionality, for example, the RAN domain, the multiaccess edge computing (MEC) domain, or a content delivery network (CDN) domain. Each one of these services has optimizations that are unique to the domain they operate. For example, in the case of RAN, the energy consumption of antennas can be reduced by adjusting the energy to individual sectors, or enabling and disabling them based on utilization characteristics. In the case of MEC, the optimization can focus on the location of the workload based on the type of energy source, or the overall utilization of a node.

These four dimensions act as a reference and can be used by service providers to explore options available today, the market reality, and what they can do to improve their options in the future.

Node level optimizations

Node level optimizations can be grouped into two types:

Disabling unused components within a service provider's network will result in noticeable reductions in power consumption.

- Hardware optimizations
- Operating system (OS) optimizations

As shown in Figure 3, optimizations range from granular control of hardware components, for example, disks, NICs, and GPUs to identifying and disabling idle components, for example, turning off unused cores or hardware accelerators to adopting alternate hardware architectures, for example, SmartNICs or Advanced RISC Machine (ARM) CPUs.

Available OS optimizations include power efficient profiles, for example, tuned server-power-save profiles and granular selection of OS settings for energy efficiency, for example, selection of CPU governors, core frequency settings, and adjusting p- or c-state.

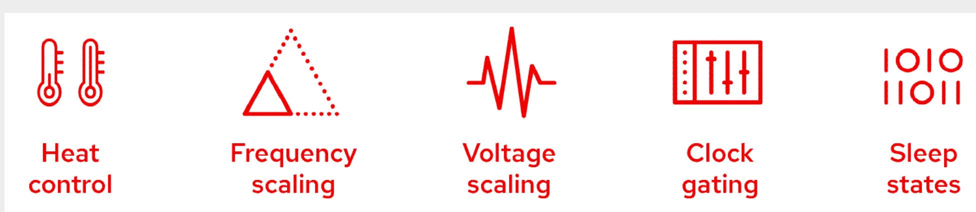

Figure 3. Node level optimizations

Red Hat

Disabling or turning off unused hardware components or having the ability to disable unused CPU cores or vGPUs adds up to a noticeable reduction in power consumption. Consider hundreds of thousands of nodes that are part of a service provider network—the power reduction achieved by the combined nodes could be significant.

Granular control of OS settings

Experience suggests the use of power efficient profiles, controlling frequency settings, CPU governors, and p-state and c-state can lead to reduced power consumption. However, the reduction of power consumption may have a negative effect on energy efficiency, and the efficiency of the power supply unit (PSU) influences the power consumption of the node.

There is not a direct correlation or linear relationship between reducing power consumption and improving energy efficiency. Power reduction can result in a less efficient execution of tasks, such as packet processing, that affects the performance of low latency applications.

There are limitations on the controllability of the PSU. Even on an idle node, the PSU can be consuming approximately 45% of its power rating, for example, a 600 watts PSU will always consume a minimum of approximately 270 watts. Even with significant reduction in power consumption due to OS optimizations, the power consumed may remain at approximately 45% or higher. OS optimizations also cannot control penalties resulting from the efficiency rating of the PSU. When consuming, for example, 600 watts, even a PSU with a 95% efficiency rating will be consuming approximately 632 watts.

There is a lower limit for which any OS optimization will not have a noticeable impact on power consumption. There is also an upper limit or maximum utilization on the node, after which factors, including temperature of the CPUs and associated electronics, have a direct effect on the energy efficiency. This can impact attributes like the performance and latency of the workload.

By understanding the node-specific minimum and maximum thresholds of utilization, the application of various optimizations can cause optimal node sustainability. Beyond the node-level benefits, the ideal scenario would be for higher levels, for example, cluster level or system level, to account for these thresholds in their placement rules for the workloads.

Alternative hardware architectures and intelligent workload scheduling are being explored to reduce power consumption.

Adoption of alternative hardware architectures

Service providers are exploring moving from x86 to ARM-based processors as they offer built-in power consumption savings and can reduce overall cost. However, there are certain hurdles to be overcome. Most networking applications were not designed to run on ARM and would require major refactoring and extensive testing before deployment in a service provider network, so it remains to be seen if savings would meet expectations.

Off-the-shelf x86 and ARM processors simply cannot meet the demands of being able to process terabytes of traffic in a single rack footprint efficiently compared to application-specific integrated circuit (ASIC)-based processors, which is why internet core routers still run on custom silicon. Maintaining mixed x86 and ARM infrastructure with intelligent workload scheduling may be the answer. However, to avoid the operational complexity penalties of mixed architectures, service providers may need to invest in new and smarter operations.

Cluster-level optimizations

Moving from node to cluster-level optimizations will create new opportunities primarily due to scale.

Cluster-level optimizations include power aware schedulers and energy aware autoscalers. Using specialized schedulers allows the system to run workloads on nodes with the best characteristics for the desired optimization outcome. By considering the ideal utilization threshold of a hardware profile optimized for power consumption or performance, a specialized scheduler can determine if spreading the workload across multiple nodes, or running more applications per node leads to the best outcome based on the desired optimization goal, for example, energy efficiency versus power consumption versus CO_2 emissions.

In multicluster environments, each individual cluster may be deployed following a specific standard blueprint. Identifying when workloads hosted in the cluster are not using optional platform features or capabilities, such as cluster-wide metrics collections or Kubernetes operators whose functionality is not being used by any of the applications, and dynamically disabling or enabling those features or capabilities will take the burden off resources that allow for better distribution or aggregation of the active workloads.

Calculation of ideal network paths with lower power consumption and autoscaling clusters based on energy consumption are examples of system-level optimizations.

System-level optimizations

System-level optimizations focus more on techniques that drive smart workload placement and scheduling, as shown in Figure 4, across a multicluster environment. System-level optimizations use advanced techniques, including determining the ideal network path with lower power consumption for service intercommunication. Autoscaling clusters based on the energy requirements of the applications and services being consumed at a particular time is another example of advanced techniques that can be employed by system-level optimizations.

These types of optimizations would depend on sophisticated AI/ML-based orchestration mechanisms capable of performing advanced analysis and deciding and predictions based on the current and expected state of the system. Introduction of intelligent orchestration could be challenging and expensive, including from the power consumption point of view, so the gains would need to be carefully evaluated prior to deciding.

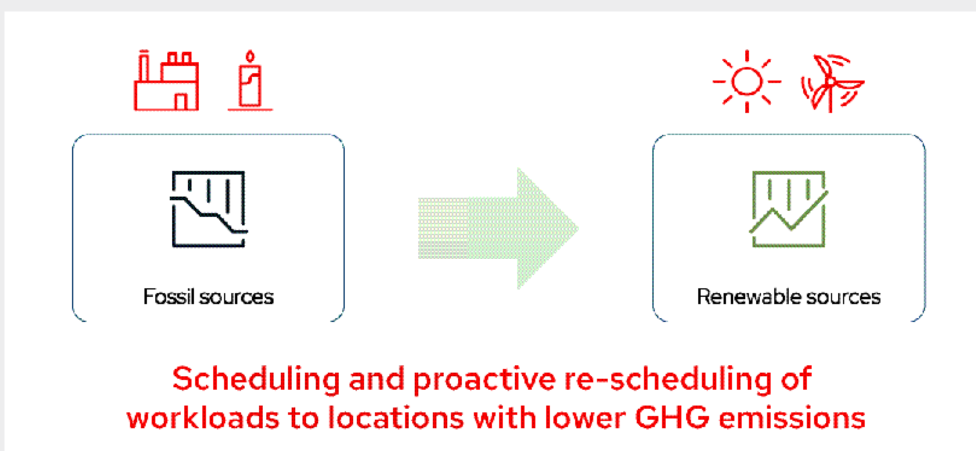

Figure 4. Workload scheduling to improve power consumption

Various RAN domain-level optimizations are achievable using AI and the RAN intelligent controller (RIC).

Domain-level optimizations

Domain-level optimizations are specially designed for the environment in which they operate. For the previously described dimensions, service providers have a relatively high control and influence over them. With domain-level optimization, service providers are relatively constrained by the options exposed by the provider of the solution of the particular domain.

The following are three domain-level optimization examples:

- Power reduction can be achieved on a service provider network using the AI and the RIC to control the power of each cell sector are well-known and documented by organizations like the open-RAN (O-RAN) alliance, European telecommunication standards institute (ETSI), Green G, Next G Alliance, and others. Unfortunately, these optimizations are only possible if the RAN vendor supports it. For example, for a programmable antenna, the particular RIC solution must have implemented the vendor-specific integrations with the antenna provider, and for the specific RAN software.
- Capgemini Engineering's whitepaper, Project Bose: A smart way to enable sustainable 5G networks, highlights the use of UE mobility prediction data by AMF to help RAN to assist in directional paging, which leads to energy savings on the 5G RAN side.[5]
- Similarly, for the CDN domain, enhancing the service with energy aware capabilities highly depends on the implementation and capabilities exposed by the CDN software.

Each service provider will have different goals for energy efficiency and power reduction, and will identify priorities for their specific organization.

The domain-level optimizations may have the most significant effect in reducing energy consumption for the RAN, which represents 75% of the power consumption of a service provider network. To make these a reality, service providers should be willing to enforce their influence on the RAN software vendors to make these available within their desired RAN solution.

Balancing competing goals

The combination of optimizations is driven by a service provider's goals. Energy efficiency improves the return on investment (ROI), reduction of power consumption improves their OPEX, and the reduction of CO_2 emissions improves ESG performance. Even when related, these can be competing goals and the priority of each one is specific to the service provider's needs. While reduction of CO_2 emissions can be achieved by the use of renewable energy, in a RAN environment and in areas where this type of energy requires the service provider to set up and maintain new power infrastructure, for example, solar panels, wind turbines, batteries, or inverters, it could cause a significant negative impact on CAPEX and the OPEX of the service provider's network. Similarly, reducing power consumption by lowering frequencies or applying other constraints to the network may affect its energy efficiency, tasks may take more time for completion and affect the user experience, or investments may be needed in more sophisticated and intelligent orchestration tools.

Trying to adopt all possible optimization strategies at the same time might not be an achievable goal. Some optimizations apply for niche workloads or configurations and will have minimal or no effect on other areas. The focus should be on incremental improvements towards the service provider's goals.

Service providers should identify the priorities of their organization that could include reduction of CO_2 emissions for achieving carbon neutrality or net-zero emissions, achieving energy efficiency to improve ROI, or reduction of power consumption to reduce OPEX. Service providers should sort

5 *"Project bose: A smart way to enable sustainable 5G networks." Capgemini, July 2022.*

relevant goals based on their priority. Once goals and priorities are identified, the right optimizations for each dimension will become clearer, and service providers will speed up the adoption of the strategies that have a direct impact on their goals.

How Red Hat can help

Sustainable computing is a multidimensional challenge and there is no single solution to fit all scenarios. Service providers have to consider options to optimize each one of the previously described dimensions.

Red Hat is working with node-level optimizations for specialized use cases including the RAN distributed unit (DU) and holistic approaches that include energy aware schedulers, energy aware workload autoscalers, and other ways to provide consistent and reliable energy consumption metrics for containerized applications on bare metal, virtualization, and cloud environments.

Some of the native features of the Red Hat® OpenShift® platform include:

Red Hat OpenShift features (subset)	
Tuned operator	Tuned helps manage node-level tuning. In Red Hat OpenShift, it defaults to high-performance profiles. The same functionality can be used to prioritize for power efficient profiles, or the creation of custom Tuned profiles to achieve service provider's goals.
Workload hints and runtime optimizations	The CRI-O runtime is optimized to detect annotations in the pod to optimize for Linux® Kernel control group (cgroup) settings towards the intended use. Workload hints capability allows the nodes to identify whether it should be tuned for very low latency at the cost of increasing power consumption, or if the more power efficient tunings can be applied to the node.
Node network latency profiles	Ability to assign the nodes to a profile that will react more aggressively or less aggressively to a change of node status (ready, unhealthy, unknown) based on the network latency that can be tolerated by a node and its workloads. This feature can be used to reduce the rescheduling of applications on environments experiencing a high network latency, which could impact the ideal power utilization thresholds of other nodes in the cluster.

Red Hat OpenShift features (subset)	
Idling applications	Red Hat OpenShift has native capabilities for the administrator to idle applications not being used. The application becomes active again when it receives network traffic or can also be manually activated. This capability allows applications to reduce resource consumption when not in use.
Cluster hibernation	When using Red Hat Advanced Cluster Management for Kubernetes to deploy and manage clusters on public cloud environments including Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP), the platform administrator can choose to hibernate clusters to reduce resource consumptions of clusters during times they are not being actively used, for example, a development or testing cluster during the night or weekend.
Scale to zero: Serverless autoscaling	Red Hat OpenShift Serverless provides automatic scaling or autoscaling for applications to match incoming demand. An application that is not receiving traffic can be scaled to zero when the scale-to-zero capability is enabled. The OpenShift Serverless service will scale up the application to meet demand if traffic to the application increases.
Autoscaling and rightsizing of applications	Horizontal pod autoscaling (HPA), vertical pod autoscaler (VPA), Kubernetes-based event-driven autoscaler (KEDA) is available in Red Hat OpenShift as mechanisms to adjust the number of replicas or the resource levels of a pod—based on general metrics or custom metrics. When CO_2 emission or power consumption metrics are available in the clusters, some of these capabilities can be used to scale workloads based on service provider's goals.

Some of the projects Red Hat is working on for sustainable computing include:

Upstream projects (subset of ongoing projects)	
KEPLER	The Kubernetes-based efficient power-level exporter (KEPLER) is one of the tools Red Hat is sponsoring for the accurate identification of energy consumption on pods.

Upstream projects (subset of ongoing projects)

KEDA	Red Hat is working with the community to bring native CO_2 and energy aware autoscaling capabilities to Kubernetes event-driven autoscaling (KEDA).
CLEVER	Red Hat is working with the container-level energy-efficient VPA recommender to allow energy consumption metrics to be used for VPA with an application.
PEAKS	Red Hat is working with the power efficiency-aware Kubernetes scheduler (PEAKS) on a Kubernetes scheduler that will take power metrics into consideration for the scheduling of an application.

Recognizing these as cross-industry challenges, Red Hat is actively participating in communities and standard development organizations (SDOs) connected to these topics. Some examples of these communities include the Cloud-Native Computing Foundation (CNCF) environmental sustainability Technical Advisory Group (TAG), OS-Climate, O-RAN Alliance, and TM Forum.

Learn more

Learn more about Red Hat's telco solutions, and find out more about Red Hat OpenShift, Red Hat Advanced Cluster Management for Kubernetes, and Red Hat OpenShift Serverless.

About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. A trusted adviser to the Fortune 500, Red Hat provides award-winning support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

facebook.com/redhatinc
@RedHat
linkedin.com/company/red-hat

North America
1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com