White Paper

Data Center

Streamlining Server Power Management

Intel® Infrastructure Power Manager enables optimization and management of power settings across a compute cluster. Compatible with servers based on a range of Intel® processors, Intel Infrastructure Power Manager helps to cut energy consumption, costs, and carbon emissions.

Authors

Tim O'Driscoll
Software Product Manager, Intel

Chris MacNamara
Senior Principal Engineer, Intel

Mohammad Abdul Awal
Senior Engineering Manager, Intel

Executive Summary

Data centers present significant opportunities for energy savings, though unlocking this potential sometimes requires overcoming complexity challenges. The wide variation in the applications running, deployment models, and server vendor solutions underscores the need for a scalable, adaptive, and unified approach to optimize power usage across this infrastructure.

Intel® Infrastructure Power Manager addresses this challenge by delivering a comprehensive software solution for automatically managing and optimizing power settings across the compute infrastructure. This suite of software agents features integrated power management control loops and application programming interfaces (APIs), enabling intelligent energy optimization and enhanced performance-per-watt across a broad range of Intel® Xeon® Scalable processors and Intel® Core™ processors.

Beyond offering extensions to the Kubernetes API to provision fine-grained power settings across server platforms, the solution also includes enhanced algorithms for actively governing power at run time that significantly reduce power consumption. The algorithms have been optimized for Intel® processors. We measured average power savings of 34% on a general all-purpose workload.[1]

Intel Infrastructure Power Manager simplifies power management, empowering organizations to improve energy efficiency and reduce operational expenses. It can also support sustainability initiatives by helping to reduce carbon emissions. Intel Infrastructure Power Manager is a critical enabler for modern, energy- conscious data center operations.

Industry Challenges in Optimizing Processor Power

Enterprises face growing pressure to reduce energy consumption to lower operating costs and meet sustainability and carbon reduction goals. The data center has substantial potential for savings because the workloads on servers may be highly variable over time, and yet they typically run near full power regardless. For example, infrastructure is often overprovisioned: a financial trading application might be busy during stock market trading hours, but the servers that run those applications could have many idle processor cores on the weekend that are nevertheless fully powered up. Similarly, telecommunications network hardware is often sized for peak demand but may frequently operate at much lower utilization at off-peak times.

Even when CPU power features are enabled, managing them effectively remains complex. The following are examples of the types of challenges that data center operators face as they look to manage power.

Table of Contents

- It can be difficult to manage power settings at scale, choosing the right Basic Input/Output System (BIOS) settings or optimal operating system (OS) power governor per server, especially when hosting various types of applications.
- Out-of-the-box power settings available through the OS enable the frequency and sleep states of the processor cores to be changed in line with utilization. Still, depending on the scenario or use case, they may not always provide optimal control. It is possible to find extra power savings when frequencies are managed more effectively.
- In a cloud-native environment, applications may be instantiated or destroyed on any processor cores, so real-time monitoring and re-configuration of power settings would be required to continually optimize power savings while maintaining application performance.
- Packet-processing applications based on the Data Plane Development Kit (DPDK) or Vector Packet Processor (VPP) libraries make the cores they run on appear to be busy all the time because of the way they poll for packets, even though there may be no actual packets. As a result, operating system power governors cannot see any power savings opportunities, resulting in systems running at maximum frequencies unnecessarily.
- Virtual machines are abstracted from the underlying hardware, so they can't see or control the power management features. To ensure the virtual machines only consume the energy they require, there needs to be visibility of the load on the system and the ability to see and change power management settings.

Managers need scalable, adaptive solutions that intelligently manage power settings across diverse hardware and application environments, delivering efficiency gains without sacrificing performance. While performance has traditionally been a concern, especially for latency-sensitive workloads like packet processing in networking and telecommunications, advancements in power management in Intel® processors now make it possible for data center managers to enable energy-saving features without compromising performance.

Introducing Intel® Infrastructure Power Manager

Intel Infrastructure Power Manager simplifies power management across complex, multi-node environments. Intel Infrastructure Power Manager empowers IT teams to reduce operational costs, improve energy efficiency, and support sustainability goals without compromising performance. It enables precise and automatic/ dynamic control over processor power settings at scale.

The solution incorporates Intel-specific and Intel-optimized power governors that go beyond traditional hardware controls, enabling more intelligent and responsive power optimization. This allows organizations to fine-tune power settings at runtime.

The software is available to all Intel customers and is validated on a broad range of Intel® platforms, as described later in this paper.

Key Capabilities

- Post-boot power provisioning: Configure power settings such as C-states, P-states, Turbo Boost, and uncore frequency from within the operating system post-boot. This enables consistent and scalable power profiles across nodes, tailored to workload demands. For example, users can define minimum/maximum frequency limits and core sleep depth and enable or disable Intel® Turbo Boost Technology per-core or system-wide. Custom resource definition (CRD) files are used to integrate with cloud infrastructure.
- Application-aware power optimization:
    - For high-performance packet-processing applications (using DPDK or VPP), an enhanced P-state governor dynamically adjusts the frequency based on the true workload intensity.
    - For general-purpose applications, built-in OS governors and an enhanced uncore frequency governor dynamically scale down power usage based on system activity.
- Idle core optimization: The software detects idle processor cores in real time and dynamically applies power-saving features, adjusting P-states, C-states, and uncore frequency to minimize energy consumption during low-utilization periods.
- Programmable REST API: A comprehensive REST API allows users to integrate power management into orchestration tools or custom applications, offering real-time policy enforcement and dynamic control.

Proven in Telecom, Built for All

While widely applicable across industries, Intel Infrastructure Power Manager has been extensively validated in the telecommunications sector, where it supports both packet and non-packet processing workloads within mission-critical environments. Network operators value the ability to manage diverse workload types through a single, unified interface.

Intel Infrastructure Power Manager Deployment Models

Intel Infrastructure Power Manager may be deployed as in-band agent(s) in the Linux OS. The Intel Infrastructure Power Manager software may be run on a host directly, run as a container on a host, be deployed as a container running in a Kubernetes Pod, or run in a virtual machine. Intel Infrastructure Power Manager may be deployed as part of the infrastructure software, part of a Container-as-a-Service (CaaS) software stack or included with a given application.

Some potential scenarios include:

- Intel Infrastructure Power Manager as CaaS; Intel Infrastructure Power Manager provisions power states and uncore governor in operation, managing system-wide uncore frequency.

- Intel Infrastructure Power Manager is included with the application; Intel Infrastructure Power Manager manages CPU core frequency for application-only cores.

Power Saving Test Results

The authors of this paper carried out extensive testing covering typical network conditions. Figure 1 shows the impact of using Intel Infrastructure Power Manager for a general all-purpose workload. The example is a control plane application in telecommunications. Control plane applications do not involve packet processing and can demonstrate the savings Intel Infrastructure Power Manager enables through using C-states and the enhanced management of uncore frequencies. The average power saving was 34% under moderate network traffic, with a peak power saving of 42% in a non-busy hour, and a 19% power saving in a busy hour.

Intel Infrastructure Power Manager was tested on a User Plane Function (UPF) packet-processing workload, as shown in Figure 2. This showed an average savings of 38.3% in power over 24 hours and a peak saving of up to 41% at idle, with no impact on I/O throughput.[2]

Validated Platforms

Intel Infrastructure Power Manager has been validated on servers based on 2nd Gen Intel® Xeon® Scalable processors through to Intel® Xeon® 6 processors and on 12th Generation Intel® Core™ processors through to Intel® Core™ processors (14th gen). The solution may run on other Intel-based systems, but similar results are not guaranteed.

The solution has been tested on Ubuntu, CentOS, Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server (SLES) operating systems at the time of publication.

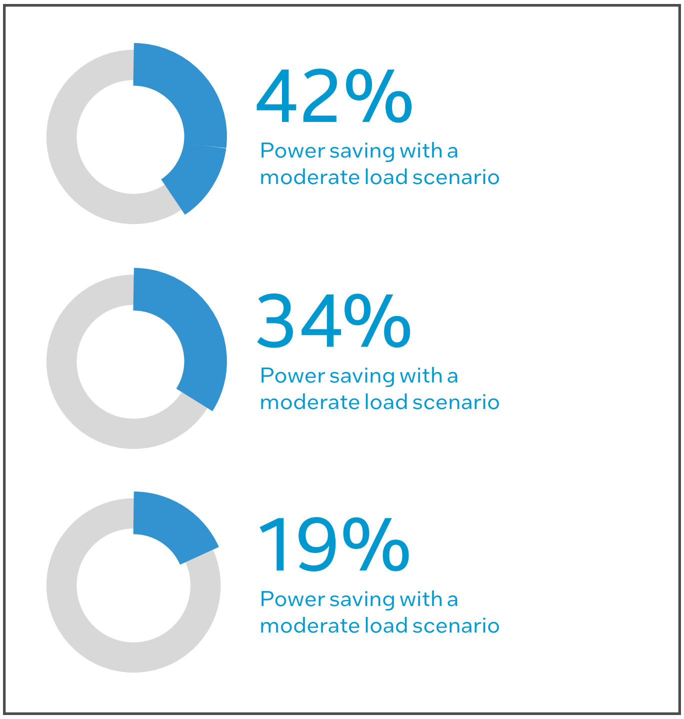

Figure 1. Using Intel Infrastructure Power Manager to manage a control plane application enables a 34% average power saving under moderate network traffic.[1]

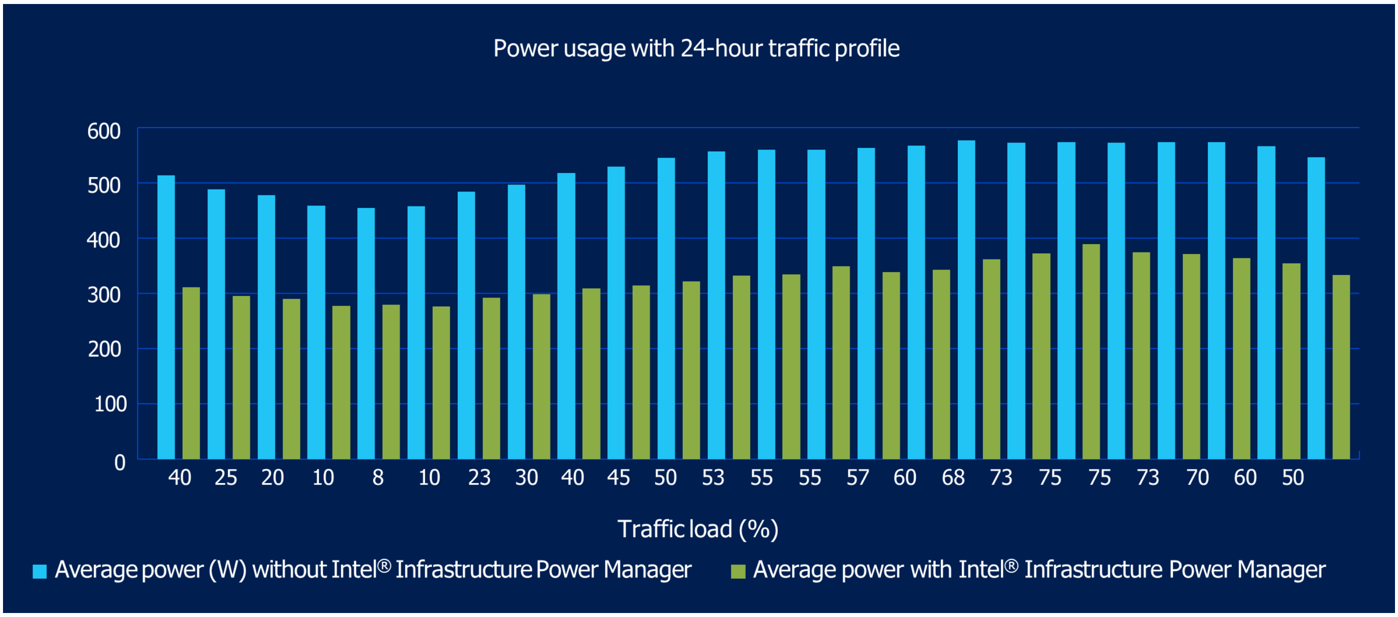

Figure 2. Using Intel Infrastructure Power Manager for a packet processing workload enables processor core frequencies (P-states) to be adjusted in line with network traffic and core utilization.[2]

Summary

Intel Infrastructure Power Manager streamlines the process of managing power features on Intel® processor platforms. It provides a single software stack for power management across various hardware platforms and applications.

It has three main use cases:

- Provisioning power configurations at scale: Intel Infrastructure Power Manager gives users a consistent API-based process to configure settings for C-states, P-states, Turbo Boost, and the uncore frequency, across server platforms from different vendors.
- Runtime power optimization for packet-processing applications: For these applications, based on DPDK or VPP, the solution adjusts core frequencies (P-states) in line with real processor core utilization. The OS power governors and optimal uncore frequency are set up at deployment time.
- Runtime power optimization for other applications: Intel Infrastructure Power Manager enables OS C-state power governors at deployment time to optimize power savings while protecting application performance. The solution then manages the uncore frequency in real time to match demand on the processor cores.

Find the solution that is right for your organization. Contact your Intel representative for further information or a demonstration of Intel® Infrastructure Power Manager.

Learn More

- Intel® Infrastructure Power Manager for 5G Core Solution Brief
- Intel Power Management Technology

---

[1] Configuration: Two 3rd Generation Intel® Xeon® Gold 6338 Processors (32 cores at 2 GHz), 512 Gbytes total memory, Intel® Ethernet Network Adapter E810-CQDA2 (qty=2). For full details see: Dynamic Power Savings in Cloud-Native 5G Wireless Infrastructure Network Functions.

[2] Configuration: Test by Intel as of 11/21/2024. 1-node, 2x Intel® Xeon® 6780E processors, 144 cores, 320W TDP, HT N/A, Turbo On, Total Memory 1024GB (16x64GB DDR5 6400 MT/s [6400 MT/s]), BIOS 1.0.1, microcode 0x30001b3, 10 x Ethernet Controller E810-C for QSFP, 1x 28.6G SanDisk 3.2Gen 1, 1x 894.2G Dell BOSS-N1, 4x 1.7T Dell NVMe PM1743 RI E3.S 1.92TB, Red Hat Enterprise Linux CoreOS 416.94.202408132101-0, 5.14.0-427.31.1.el9_4.x86_64. For more details, see Intel-Dell Verified Reference Configuration for 5G Core on Red Hat OCP.

Notices & Disclaimers

0925/CG/DJA/PDF Please Recycle 366871-001US