Solution Brief

Power Management – Leveraging AI for Smarter Data Center Power Efficiency

With QiO's AI-driven Foresight Optima DC+™ leveraging 3rd Gen Intel® Xeon® Scalable processors telemetry and insights, data center server workload efficiency can be determined and optimized to reduce energy consumption, leaving network operators in control of when and how to reduce energy costs and CO_2 emissions.

Executive Summary

Driven by the industry demand to reduce electricity consumption and carbon emissions, Intel and QiO Technologies have partnered to create a solution that identifies, recommends, and optimizes power consumption of Intel® Xeon® CPUs.

QiO is an industrial IoT AI software products company that delivers sustainability solutions for asset heavy and energy intensive sectors. QIO is leveraging the 3rd Gen Intel® Xeon® Scalable processor technology, to develop a novel solution for data centers in order to reduce power, emissions, and energy consumption. The QiO application learns how workloads operate on each server and uses Intel's power management technologies to avoid inefficient power use and reduce carbon emissions by using only the necessary energy required to accomplish the task.

QiO's AI driven Foresight Optima DC+™ application identifies power consumption inefficiencies and recommends actions to save energy. The end user can manually implement the recommendations (open loop) or fully automate the implementation of those recommendations (closed loop).

This solution brief provides an overview of the approach, describes the Proof of Concept (PoC) developed to showcase the solution, and explains the benefits.

This paper will show that up to 52% power saving can be achieved with idle, and up to 24% with a representative load, on a given CPU.

Note: In the graphs and results in this paper, the carbon emissions we refer to are CO2e (Carbon dioxide equivalent). "Carbon dioxide equivalent" or "CO2e" is a term for describing different greenhouse gases in a common unit. For any quantity and type of greenhouse gas, CO2e signifies the amount of CO2, which would have the equivalent global warming impact. For more information, refer to https://ecometrica.com/greenhouse-gases-co2-co2e-and-carbon-what-do-all-these-terms-mean/.

The CO2e calculations here are UK based. For US-based calculations, please see https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator#results

This document is part of the Network Transformation Experience Kit.

Solution Description

QiO Technologies suite of products – the Foresight Sustainability Suite™ (Foresight Optima™, Foresight Maintenance®, Foresight Service®) supports energy intensive and asset heavy industries to reduce energy consumption, increase energy savings, reduce carbon emissions, monitor the health of the asset, reduce unplanned downtime, enhance circularity of the asset, and automate ESG reporting.

QiO's Foresight Optima™ has been enhanced as Foresight Optima DC (data center)™ to continuously monitor and control HVAC and cooling systems to reduce energy costs, water usage, and improve overall power usage effectiveness (PUE). In addition to Foresight Optima DC™, QiO has developed Foresight Optima DC+™ in collaboration with Intel, for data center server optimization.

Foresight Optima DC+™ is the focus of this paper. It augments CPU-level optimization by leveraging Intel's platform telemetry and suite of CPU power management technologies and the Artificial Intelligence (AI) capabilities of Foresight Optima DC+™, specifically its multi-variate, multi-objective optimization algorithms, to intelligently learn server power usage patterns and propose (open loop) or automatically trigger (closed loop) actions to adjust power consumption to meet the needs of the server workloads more efficiently, without compromising operations or quality of service. See Figure 1.

In this solution, the Intel® processer based platform provides a rich suite of power specific telemetry. Foresight Optima DC+™ provides the data analysis and generates recommendations for improved power efficiency. This is then combined with extensive suite of Intel® Xeon® Scalable processor power management technologies to implement these recommendations to deliver additional energy savings and carbon emissions reduction.

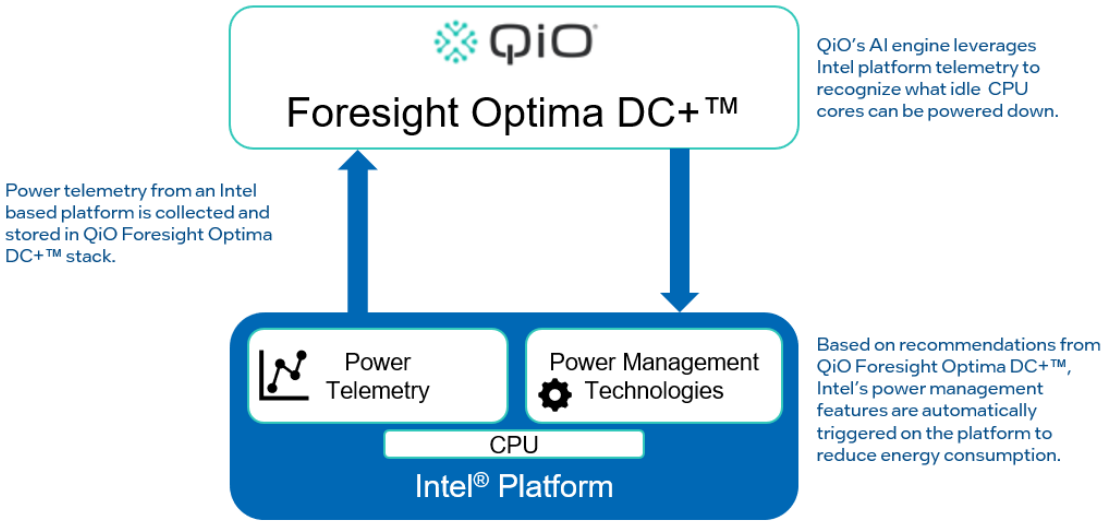

Figure 1. High-level Overview of Solution

Intel Power Management Technologies Used by Foresight Optima DC+™

Improved CPU Monitoring

With enhanced CPU power telemetry and insights, the end user can get a more fine-grained picture of how much power is being consumed, and where it is consumed on an Intel-based server. For example, busy, idle, and unallocated cores can be identified that will help to calculate an estimate of potential dollar and carbon savings. With this information, the end user can make informed decisions on the power saving technologies they wish to use and when.

CPU Power Saving Features

CPU power states or C-states are sleep states for the processor that allows cores to sleep and save power. With increasingly deeper levels of C-states, the energy efficiency gains are more significant. Workload performance and power savings can be balanced to give the most efficient solution by using C-states for use cases like the *time-of-day* scenario. This will be discussed in this document.

Other examples of using sleep states could be when unallocated compute, zombie, or stranded resources have been identified. If this compute is not active currently but, may be in the future, then why not enable the C-state capability. It will help reduce power consumption, which will also have a knock-on effect of energy saving for fans and external cooling.

For more information on other Intel Power Management Technologies, see
https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html

Use Case Examples Using C-States

In a large fleet of servers, the load on the server varies but the power consumption does not, which means there is an opportunity to match the load to the power consumption.

A simple use case that models the benefits of the improved CPU power monitoring and power saving features is the *Time-of-Day* scenario. In this scenario, system usage follows a particular pattern over a 24-hour period, which is well known to network operators.

Broadly speaking, the traffic load on the network starts low in the morning and gradually increases throughout the day (as more people wake up, logon, and go about their daily business), reaching a peak in the late evening (with video streaming for example). With this variation of network load, server usage increases and decreases accordingly. However, without power management in place, the power being consumed remains relatively high.

One of the key points that this paper highlights is the use of Intel's power management technology where the CPU cores can enter a deep sleep if the load on a server decreases, and CPU cores are idle. This, in turn helps to save power.

Figure 2 below shows a simple 24-hour period pattern, with CPU usage increasing from early morning throughout the day, reaching a peak in late evening (blue line). The energy consumption, however, remains relatively flat (red line).

By monitoring fine-grained CPU power telemetry from the Intel server, the end user can identify where the opportunity for power savings is (indicated by green arrows) and take action to ensure a more efficient use of system power, and reduce their carbon footprint.

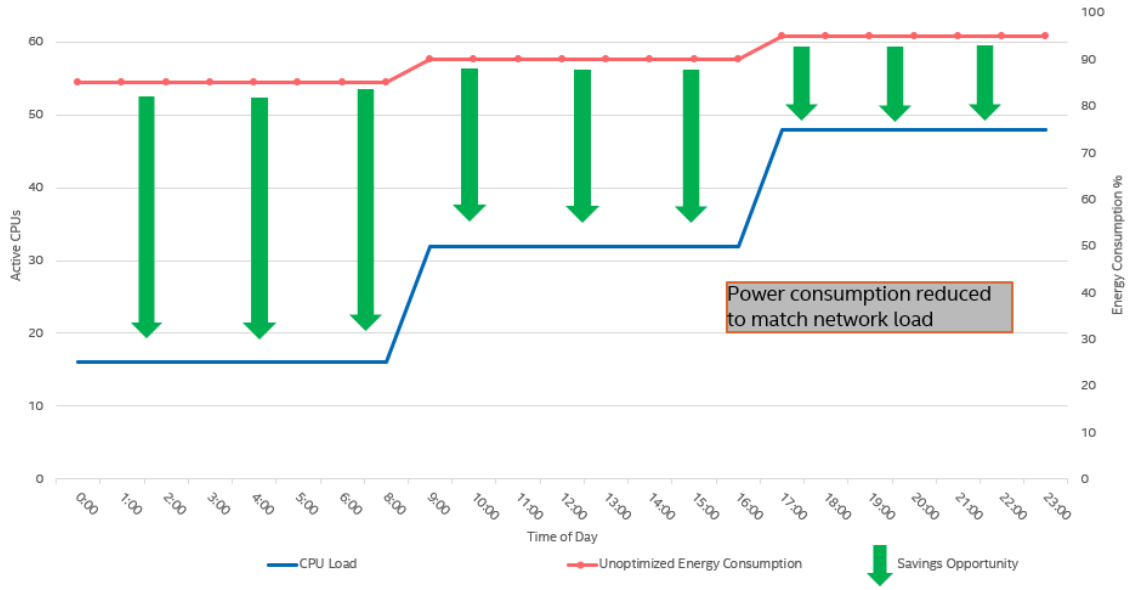

Figure 2. Illustrating the Savings Opportunity when Power Management Enabled

Proof of Concept (PoC)

QiO's Foresight Optima DC+™ has been tested in a cluster of multiple servers and demonstrated that it is possible to monitor and provide power savings recommendations for a fleet of servers. The results that follow however focus on the savings that can be achieved at a single node level to highlight potential power savings.

The workload server uses stress-ng to simulate a synthetic workload. Power-related telemetry is collected via Collectd and ingested by Foresight Optima DC+™ as depicted in Figure 3.

Collectd is an open-source telemetry collection daemon that runs on Unix-like OS's. It has a pluggable architecture and uses the concept of south bound plugins, which will read metrics from a platform feature, and north bound plugins, which will make this information available in various formats, so that it can be consumed as necessary (such as through monitoring or analytics solutions). The frequency of metrics gathering from the platform and/or processes is fully configurable, ranging from seconds to minutes, depending on the granularity required.

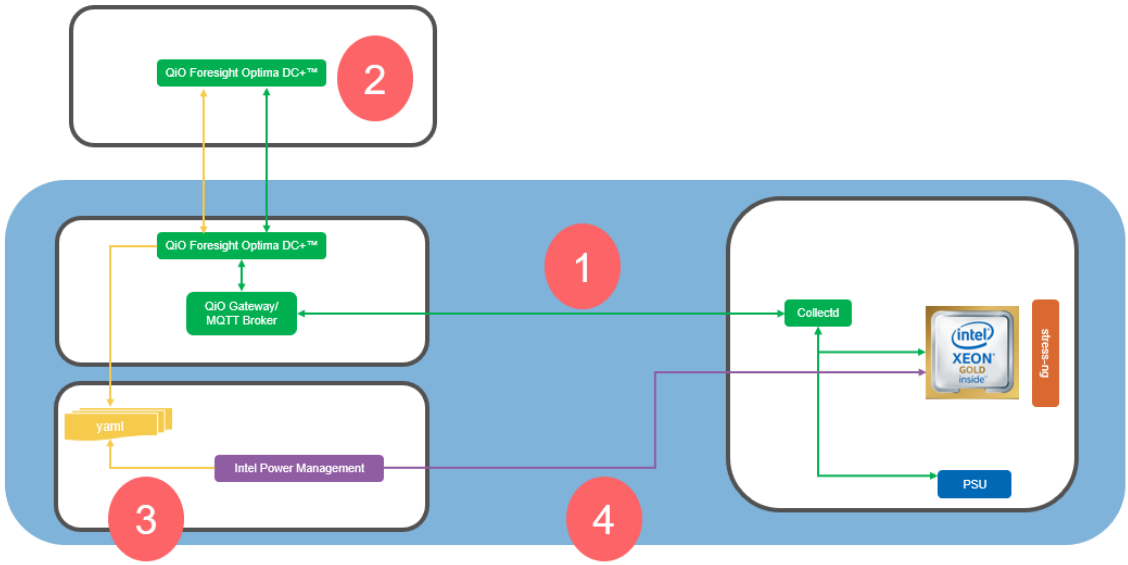

Figure 3. PoC Setup Showing a Single Server

The data flow is as follows:

1. CPU and system power are collected from the SUT via collectd Turbostat and IPMI plugins. They are then sent to the QiO gateway.
2. Foresight Optima DC+™ uses its Artificial Intelligence (AI) capabilities to automatically generate an Energy Efficiency Index (EEI) model based on the power-related telemetry received from the Intel® platform. The EEI is a relative indicator, in the range 0 to 1, of the power efficiency of a server. A value closer to 0 highlights inefficiencies in server power consumption versus usage and so, identifies the opportunities for power saving. A value of 1 indicates that a server is operating at its most efficient.
3. Based on this EEI, Foresight Optima DC+™ learns the usage pattern of the server over time and recommends the most appropriate Intel power management techniques to apply to save power (open loop), with the option of automating the implementation of those recommendations (closed loop).
4. In this case, the recommendations from Foresight Optima DC+™ are created in a YAML format and automatically applied to the SUT with the Intel Power Management scripts in a closed loop scenario.

The PoC was setup in this case to identify:

1. Idle servers (64 cores idle) and the power savings that could be achieved if power management features enabled
2. Idle cores while running a workload. In this case 50% of the cores (32 cores) were 30% utilized for a given time and the rest were idle. (In this case we used the stress-ng workload generator to provide the equivalent network traffic load).

By applying the power savings and efficiency recommendations suggested by Foresight Optima DC+™, the result is a significantly more efficient use of system resources, ensuring that workloads maintain their performance requirements while also achieving lower overall power consumption.

PoC Results

Scenario 1: Stress-ng workload

Stress-ng workload deployed on the SUT (Dual Socket Server with Intel® Xeon® Gold 6338 Processor) on 32 out of 64 cores @ 30% load - 16 out of 32 cores on each socket @ 30% load.

Power Management Inactive:

SUT running stress-ng workload. Intel CPU Power Management features are disabled.

Power Management Active:

SUT running stress-ng workload. Foresight Optima DC+™ is calculating an Energy Efficiency Index and is making **recommendations** as to when and where CPU C-states can be enabled on the SUT. These C-state recommendations are automatically enabled/disabled on the SUT.

Results: SUT (Dual Socket Server with Intel Xeon Gold 6338 Processor) running Stress-ng workload uses up to 24.78% less power when Foresight Optima DC+™ 22.0 is enabled.

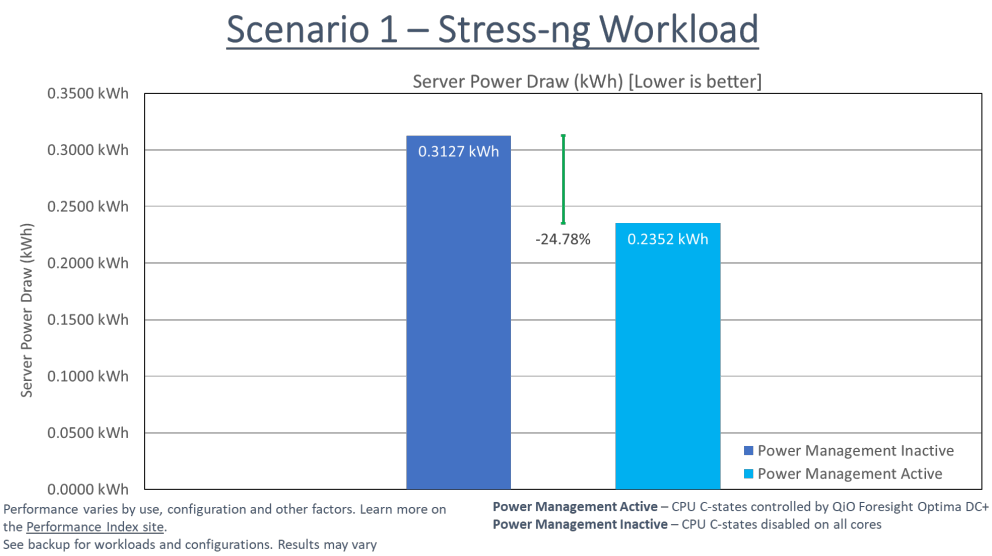

Figure 4. Stress-ng Workload Results

Scenario 2: Idle server

SUT (Dual Socket Server with Intel Xeon Gold 6338 Processor) at idle is tested to show the power savings available for enabling C-states on an idle server

Power Management Inactive:

SUT at idle. Intel CPU Power Management features are disabled. Idle means OS is running, cores in C0 and P1 frequency, uncore frequency scaling disabled and turbo disabled, with no workload running.

Power Management Active:

SUT at idle. QiO Foresight Optima DC+™ is calculating an Energy Efficiency Index and is making recommendations as to when and where CPU C-states can be enabled on the SUT. These C-state recommendations are automatically enabled/disabled on the SUT.

Results: SUT (Dual Socket Server with Intel Xeon Gold 6338 Processor) at idle uses up to 52.61% less power when Foresight Optima DC+™ 22.0 is enabled.

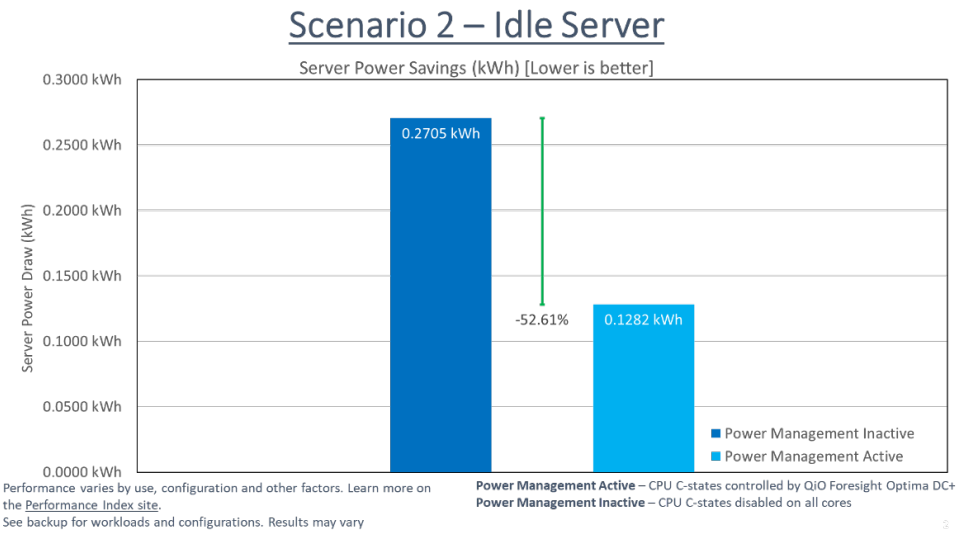

Figure 5. Idle Server Results

Foresight Optima DC+™ Dashboard

Foresight Optima DC+™ provides a web based dashboard to provide an overview and operational information on the server cluster. For each server you can view the Energy Efficiency Index, Carbon Emissions, and Energy Costs for a period of time.

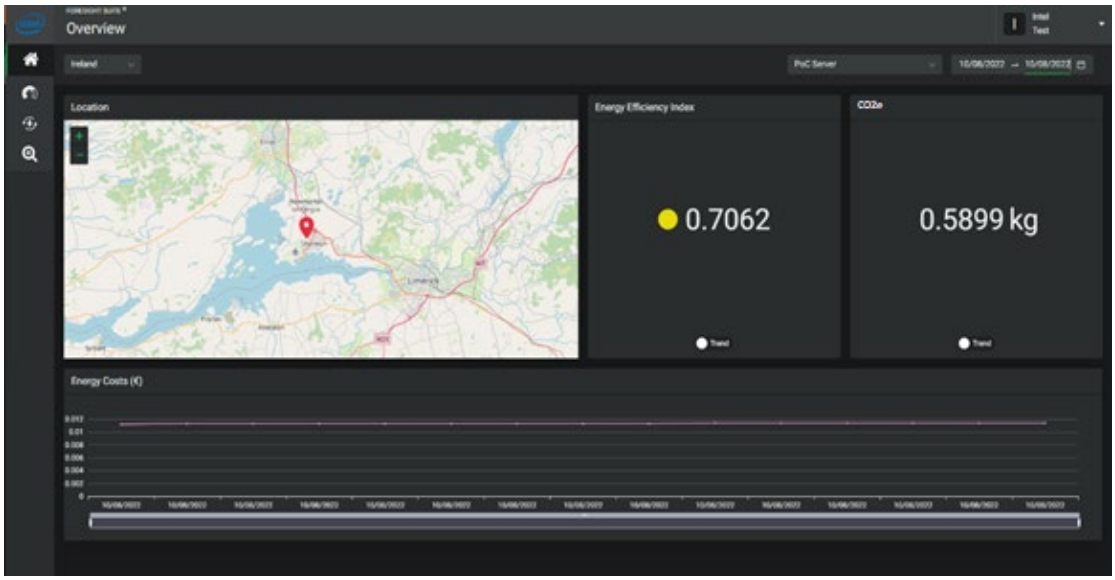

Figure 6. QiO Foresight Optima DC+™ Dashboard Overview

The Energy Efficiency Index and Carbon Emissions graphs for the period where power management was inactive (26/06), show that there are some efficiency improvements that can be made. During this period Foresight Optima DC+™ is recommending when and where to enable and disable C-states, but no action is being taken. By enabling Foresight Optima DC+™ to act on these C-state recommendations (27/06), there is an immediate improvement in the Energy Efficiency Index, and also significant reduction in carbon emissions. See Figure 7.

Figure 7. Stress-ng Workload Server EEI and Carbon Emissions in Foresight Optima DC+™ Dashboard

Similarly, with the SUT running completely idle. According to the EEI, the system is already at a good level of efficiency. However, there are still some efficiency gains to be made. This is made clear when Foresight Optima DC+™ acts on the C-state recommendations (15/07), and we see a significant reduction in Carbon Emissions from the server, as well as an improvement in EEI as a result. With the completely idle server there is opportunity for all cores (64 on this SUT in this case) to enter the power-saving C-states, which results in a significant reduction in power consumption from the server, and carbon emissions also as a result. See Figure 8.

Figure 8. Idle Workload Server EEI and Carbon Emissions in Foresight Optima DC+™ Dashboard

Benefits of Solution

This paper describes how it is possible to match the power consumption more closely to the demands of the workload. Adopting such a technique leads to:

- Improved power efficiency, translating to significant operational cost reductions
- A reduced carbon footprint, promoting a greener and more environmentally friendly approach

Our Partner – QiO Technologies

QiO Technologies (www.qio.in) is an Industrial IOT AI driven software products company that helps energy intensive and asset heavy industries make a faster transition to an ESG compliant, carbon neutral, sustainable based world economy. QiO was established in 2015 and has invested over 100+ person years in building, world leading AI driven, Industrial IoT suite of applications (Foresight Optima™, Foresight Maintenance®, Foresight Service®) that makes a measurable difference to the sustainability agenda by delivering costs/productivity improvements, energy savings, carbon emissions reduction, enhanced asset life, and automated ESG reporting to help companies accelerate their Net Zero carbon agenda.

QiO's Foresight Optima™ reduces energy costs and carbon emissions, optimizes renewable energy consumption and avoids tariff increases, delivering sustainability and energy efficiency. Foresight Optima DC™ and Foresight Optima DC+™ were designed to help data center operators accelerate delivery of their net-zero targets.

Summary

As a result of applying QiO Foresight Optima DC+™ savings of 24-52% can be seen by optimizing power states in Intel based servers. This lays the foundation for energy efficiency and carbon reduction.

By integrating the fine-grained CPU power telemetry from Intel-based servers with the intelligent algorithms and recommendations of QiO Foresight Optima DC+™, service providers can gain beneficial insights into the power usage and efficiencies of their platforms and applications. This allows strategic optimizations on power consumption and carbon emission targets to be met, while maintaining system performance and QoS.

To learn more about Intel's power management technologies and platform telemetry, see: https://networkbuilders.intel.com/intel-technologies/network-transformation-exp-kits

To learn more about Foresight Optima DC™ and Foresight Optima DC+™, see: www.qio.io

Backup

Server Configuration:

Stress-ng workload: 1-node, 2x Intel Xeon Gold 6338 on M50CYP2SBSTD with 256GB (16x16GB DDR4 3200MT/s [3200MT/s]) total DDR4 memory, microcode 0xd000332, HT on, Turbo on, Ubuntu 20.04.3 LTS, 5.4.0-99-generic, 1x 372.6G INTEL_SSDSC2BA40, stress-ng version 0.11.07 (gcc 9.4, x86_64 Linux 5.4.0-99-generic), Collectd 5.12, CommsPowerManagement commit: 87ab9d1167bd893819dbb720f8a68301dd5bb5e1, QiO Foresight Optima 22.0, test by Intel on 06/26/2022

Idle server: 1-node, 2x Intel Xeon Gold 6338 on M50CYP2SBSTD with 256GB (16x16GB DDR4 3200MT/s [3200MT/s]) total DDR4 memory, microcode 0xd000332, HT on, Turbo on, Ubuntu 20.04.3 LTS, 5.4.0-99-generic, 1x 372.6G INTEL_SSDSC2BA40, stress-ng version 0.11.07 (gcc 9.4, x86_64 Linux 5.4.0-99-generic), Collectd 5.12, CommsPowerManagement commit: 87ab9d1167bd893819dbb720f8a68301dd5bb5e1, QiO Foresight Optima 22.0, test by Intel on 07/14/2022

Terminology

Table 1. Terminology

Abbreviation	Description
AI	Artificial Intelligence
DC	Data Center
EEI	Energy Efficiency Index
ESG	Environmental, Social, and Governance
HVAC	Heating, ventilation, and air conditioning
PoC	Proof of Concept
PUE	Power Usage Effectiveness

References

The following table provides links for the references mentioned in this solution brief.

Table 2. References

Reference	Source
IEA Data Center and Data Transmission Networks, November 2021	https://www.iea.org/reports/data-centres-and-data-transmission-networks
	www.qio.io
Experience the How Wonderful Gets Done Event	https://www.intel.com/content/www/us/en/events/how-wonderful-gets-done.html
Network Transformation Experience Kits	https://networkbuilders.intel.com/network-technologies/network-transformation-exp-kits
stress-ng	https://manpages.ubuntu.com/manpages/artful/man1/stress-ng.1.html

Document Revision History

Revision	Date	Description
001	September 2022	Initial release.
002	September 2022	Updates made to third-party logo and trademarks

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. See backup for configuration details. 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 disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

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

© 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.

0922/DN/WIPRO/PDF 742231-002US