



IBM[®] LinuxONE 4



The rules of business have fundamentally changed. Acceleration of digital transformation and adoption of hybrid cloud and AI are creating exciting opportunities to drive new innovations in your business. Sustainability, which was already high on the CEO agenda, is being given new emphasis by consumer and investor demands. Yet challenges remain, especially from the evolving landscape of cyber-threats. IBM's goal is to create innovative technologies that help you address these challenges and leverage opportunities for competitive advantage.

CEOs see sustainability as the number one challenge over the next two to three years as their businesses continue their digital transformation. One way to reduce energy consumption is improving IT efficiency, as it is estimated that data centers globally already use 200 to 250 TW of electricity -- about 1% of all electricity consumed. As compute needs grow to address growth in data, transaction volumes and cyber threat, so does energy consumption by the data center.

IBM® LinuxONE 4 delivers an optimized architecture built to handle the multi-workload needs of a modern scalable digital business, while delivering the highest qualities of service and the utmost economic and energy efficiency.

That's why many businesses that are known for sustainable practices, from large banks such as Citi to startups such as [Plastic Bank](#) and [Newlight Technologies](#), leverage IBM® LinuxONE at the core of their businesses.

Helps you turn your sustainability strategy into action

Powered by high-performance IBM Telum™ cores, IBM® LinuxONE 4 delivers high levels of scalability to support hundreds of workloads in a single system. The Telum processor also features additional compute resources for dedicated workloads such as on-chip AI processing and dedicated crypto processors that deliver protection at speed and scale.

The optimized architecture in IBM® LinuxONE enables you to grow workloads and add advanced IT functionality with a negligible increase in energy usage. It provides better performance with fast cores (4.6 -5.2 GHz), larger caches and more of them. IBM® LinuxONE 4 has on-chip accelerators for data compression, encryption, and AI. IBM® LinuxONE's hardware optimizations for Java™ and specialized I/O processors make it ideal for mixed workloads. IBM® LinuxONE helps resolve complexity and optimize queries by enabling big databases to be held on a single system.

IBM® LinuxONE Rockhopper 4 scales up to 68 4.6 GHz configurable Telum cores in either a tower (the Model LA2) or rack mount (Model AGL) configuration, which can be installed in either an IBM rack or a customer-supplied rack. Either way, IBM® LinuxONE Rockhopper 4 is available with four options for core capacity (Max5, Max16, and Max32 in a single processor drawer, and Max68 with two processor drawers) and up to 16 TB of RAIM, at up to 8 TB per processor drawer.

Highlights

- Integrate AI into business processes and IT operations to increase decision velocity and customer value. Protect data now and in the future with quantum-safe cryptography
- Build a cyber resilient environment
- Modernize for hybrid cloud to deliver new value

With IBM® LinuxONE Emperor 4, process up to 300 billion inference requests per day with 1ms response time using a Credit Card Fraud Detection model.

IBM® LinuxONE Emperor 4 systems, with GDPS®, IBM DS8000® series with Hyper Swap® and running a Red Hat® OpenShift® Container Platform environment, are designed to deliver 99.999999% availability.²



The IBM® LinuxONE 4 Integrated Accelerator for AI

IBM® LinuxONE 4 systems integrate new artificial intelligence (AI) acceleration via an on-chip AI coprocessor to reduce latency and deliver outstanding performance for in-transaction inferencing.

Now organizations can embed AI directly into business processes and existing IBM® LinuxONE applications to improve business outcomes and deliver customer value in each interaction at unprecedented scale and speed within stringent service level agreement (SLA) response time guidelines.

The on-chip compression accelerator

The on-chip compression accelerator function can reduce data storage, communications requirements, and costs, as well as increase data transfer rates to boost throughput without adversely impacting response times. The on-chip Compression accelerator improves systems performance for pervasive encryption, so that customers can encrypt 100% of their data, 100% of the time.

Data security for today and tomorrow

IBM® LinuxONE 4 represents a breakthrough in data security. Quantum-safe cryptography is embedded in the system to improve the resiliency to

cyber-attacks from bad actors with future access to quantum computing resources. Today's cyber threats often involve harvesting encrypted data for decryption later when these resources can break today's encryption algorithms. IBM® LinuxONE 4 represents a step forward as customers have a safe and tested infrastructure that can deploy the more sophisticated and complex cryptography needed to protect today's sensitive data from cyber risks as they emerge.

IBM® LinuxONE 4 sustainability

IBM® LinuxONE 4 is designed to enable even greater sustainability success for businesses.

IBM® LinuxONE 4 systems are built for the modern data center to optimize flexibility and sustainability. With the introduction of on-chip AI acceleration, new capabilities for partition-level power monitoring, and the IBM® LinuxONE Rockhopper 4 ability to install into a client data center rack and power distribution, both IBM® LinuxONE Emperor 4 and Rockhopper 4 systems continue to deliver advantages to help reduce your carbon footprint. These key architectural advantages help distinguish IBM® LinuxONE 4 for sustainability in your data center, especially when consolidating workloads from x86.

To help our clients better understand the environmental impacts associated with IBM® LinuxONE 4 products, a product carbon footprint report is available for representative product configurations. This product carbon footprint uses the Product Attributes to Impact Algorithm (PAIA) to provide estimated greenhouse gas emissions in carbon dioxide equivalent over the product's lifecycle using average emissions factors. The product carbon footprint information can be used to identify the major drivers of impact, known as hotspots, within the materials acquisition, manufacturing, fulfillment, and use of the product.

Monitoring sustainability data is critical for companies to make progress on their net zero commitments. Beginning with IBM® LinuxONE 4, granularity of power efficiency and carbon footprint data is now available down to the logical partition (LPAR) level. This new telemetry information is supported by an enhanced hardware management console (HMC) Environmental Dashboard, aligning the available sustainability telemetry information in compliance with the latest Tier 1 Data Center Infrastructure Management (DCIM) standards by the American Society of Heating, Refrigerating and Air-conditioning Engineers (ASHRAE). This information is also accessible through enhancements to the HMC Web Service application programming interface (API), allowing for integration into modern DCIM systems through a set of secure, Representational State Transfer (REST) APIs. In addition, clients can track energy consumption with the IBM Instana™ Observability sensor for IBM® LinuxONE.


Offering deployment flexibility is important as clients look to optimize their data centers and reduce their carbon footprints. IBM has focused on the environmental attributes of its product packaging and shipping since the late 1980s. IBM directly procures source paper and paper-based or wood-based packaging from forests that are sustainably managed and certified as such. With the rack mount IBM® LinuxONE Rockhopper 4, clients can install IBM® LinuxONE hardware into their existing data center rack and power distribution unit (PDU) infrastructure, eliminating the need to ship an IBM frame and PDU hardware, thereby reducing the product's carbon footprint.

The combination of these IBM® LinuxONE 4 enhancements with existing IBM® LinuxONE sustainability benefits and IBM's overall commitment to sustainability means that having an IBM® LinuxONE 4 system in your data center will go a long way toward helping you meet your sustainability goals.

Flexible infrastructure

A new IBM offering, Flexible Capacity for Cyber Resiliency, is for clients who purchase a IBM® LinuxONE 4 system and want to transfer compute capacity easily and efficiently between different data centers. This can be highly valuable for disaster recovery, regulatory compliance, maintenance, and other business needs.

Remote Code Load (RCL) is available as an alternative to having the IBM Service Support Representative (SSR) come on site to update the IBM® LinuxONE firmware. RCL optimizes resilience and keeps your IBM® LinuxONE up to date with the latest features, fixes, and



maintenance without requiring someone to be physically in the data center to install and monitor planned updates to your system. This optional feature provides secure, remote installation and monitoring by IBM for planned updates to your IBM® LinuxONE.

IBM® LinuxONE Emperor 4 is a tower system built within a 19" factory rack, and flexibly scales from one frame to four, depending on the configuration. IBM® LinuxONE Rockhopper 4 is also built with a 19" housing that can be installed into either an IBM rack or a client's rack. The rack mount configurations enable clients to seamlessly integrate Rockhopper 4 into existing racks with other equipment in a hybrid cloud data center. Clients migrating to Rockhopper 4 from earlier IBM® LinuxONE systems or consolidating Linux® workloads from distributed systems will achieve significant sustainability improvements, including improved environmental efficiencies and floor space reduction for most clients. Changes to the footprint mean:

- The factory frame no longer requires the I/O drawers to be locked into fixed locations.
- Support continues for both raised and non-raised floors as well as top exit and bottom exit for I/O and power cabling. All cabling is routed to the back of the frame with new brackets to contain cables.
- Clients have embraced intelligent power distribution unit (iPDU) technology implemented on the IBM® LinuxONE Emperor 4 and the LinuxONE Rockhopper 4 factory frame. The doors are designed for acoustics and optimized for airflow. The IBM factory frame requires 3-phase power.
- The use of the iPDU power may improve power efficiency and reduce overall energy costs depending on the required configuration (IBM® LinuxONE Emperor 4 and Rockhopper 4 factory frame only).

The 19" frame technology supports the A3 operating class as defined by ASHRAE. The benefit of having an A3 class rating is being able to save on heat, ventilation, and air conditioning (HVAC) costs due to the wide range of operating conditions that will now be available.

The optional IBM® LinuxONE Hardware Management Appliance (HMA) can be ordered with IBM® LinuxONE 4 systems to provide HMC and Support Element (SE) functions within the 19" frame, eliminating the need for a separate HMC outside of the server.

Continuous Compliance

IBM® LinuxONE 4 systems contain new capabilities to make it easier and more productive to comply with the Payment Card Industry – Data Security Standard (PCI-DSS) regulatory guidelines. Audit preparation times can be significantly reduced and require less staff to complete. IBM® LinuxONE 4 systems are integrated with the IBM® LinuxONE Security and Compliance Center to monitor and record system, network, and application data for changes and adherence to PCI-DSS standards. A user-friendly dashboard enables infrastructure personnel to easily and quickly generate reports that auditors need and to ensure a continuous compliance posture that mitigates the potential for noncompliance regulatory fines.

IBM® LinuxONE 4 for Hybrid Cloud

IBM® LinuxONE is continuing to deliver new and improved cloud capabilities on the platform. IBM is empowering developers across the organization by adopting a broad set of open and industry standard tools including an agile DevOps methodology to accelerate modernization.

Now we have taken another big step forward in modernization technologies with IBM® LinuxONE 4.

IBM announced several open-source compilers to leverage the Integrated Accelerator for AI on the IBM Telum processor chip that powers IBM® LinuxONE 4. This enables programmers to embed inferencing in applications easily and execute at scale. Existing IBM compilers have been updated to exploit the latest IBM® LinuxONE 4 architecture. This capability allows them to deliver cross-platform development and integration, operate with Java, Swift or Node.JS and optimize enterprise workload performance without recompiling, as well as reduce the central processing unit (CPU) cycles needed to complete the job.

IBM® LinuxONE 4 provides more compute capacity, more cache per core, and more memory to Linux deployments. IBM® LinuxONE 4 systems provides the platform for modernizing, developing, and managing on-premises containerized applications. The new addition of rack mount options for IBM® LinuxONE Rockhopper 4 opens new opportunities within a data center.

Fast and secure access to data

High-speed connectivity to data is critical to achieve balanced performance with storage device and exceptional transaction throughput. IBM® LinuxONE 4 offers:

- Storage connectivity: A 2-port fibre connection IBM FICON® Express32S adapter that connects your IBM® LinuxONE 4 system to switches, directors, and storage devices at up to 32 Gbps. With support for native FICON, High Performance FICON for IBM Z (zHPF) and industry standard Fibre Channel Protocol (FCP), the adapter helps meet the low latency and increased bandwidth demands of applications. When the fibre channel connection endpoints use the FICON Express 32S adapter or FICON Express16SA adapters to the IBM DS8900F storage, authentication of the endpoints is enabled.

- Network connectivity: A new set of IBM Open Systems Adapters (OSAs) – OSA-Express7S 1.2 – meet the increased networking bandwidth demands driven by high-speed processors and faster network-attached storage devices.
- Support for IBM zHyperLink™ 1.1, a direct-connect, short-distance I/O adapter offering extremely low latency connectivity to FICON storage systems. The IBM Washington Systems Center offers the zBNA tool to help determine workload candidates that can benefit from this adapter.
- Shared memory communications (SMC) that are used for either direct memory placement of data (SMC-D) within the IBM® LinuxONE 4, or host-to-host memory communications, using Remote Direct Memory Access over Converged Ethernet (RoCE) Express adapters, without significant TCP/IP processing overhead.
- Ability to participate in a Server Time Protocol (STP) coordinated timing network using IBM Coupling cards.
- The IBM Adapter for NVMe is an interface card that allows installing a solid state (flash) drive (SSD) in the IBM® LinuxONE I/O subsystem, connecting to the processors at the system bus speed for minimizing latency. Up to 16 adapters – each with one client-supplied SSD – can be installed in a IBM® LinuxONE system.

Why IBM®?

Position for today and tomorrow

Success in the digital economy is contingent on making IT a creator of value internally and externally. Fundamental to this is a flexible infrastructure that positions organizations strategically, leveraging AI and hybrid cloud, while protecting existing investments and improving sustainability.

IBM® LinuxONE 4 delivers this with scale, agility, resiliency, performance, a security-rich environment, and a lower overall total cost of ownership (TCO). IBM® LinuxONE 4 provides confidence in meeting the future, in a world of uncertainty.

For more information

Specifications for IBM® LinuxONE 4:

	IBM® LinuxONE Emperor 4	IBM® LinuxONE Rockhopper 4 factory frame	IBM® LinuxONE Rockhopper 4 rack mount	IBM® LinuxONE 4 Express*
Maximum number of configurable cores	200	68	68	16
Maximum numbers of drawers	4	2	2	1
Maximum number of IO drawers	12	3	3	1
Number of factory frames	1 to 4	1	No frame	No Frame
Can co-locate with storage/switch...	No	No	Yes	Yes
Frequency	5.2 GHz	4.6 GHz	4.6 GHz	4.6 GHz
Telum Chip	Yes	Yes	Yes	Yes
Max Memory	48 TB	16 TB	16 TB	1 TB
Sizes	39, 82, 125, 168 and 200	5, 16, 32 and 68	5, 32 and 68	16

*Additional feature codes can be purchased and shipped separately from the IBM® LinuxONE 4 Express base configuration

Carbon Footprint Report:

IBM® LinuxONE Emperor 4:

<https://www.ibm.com/downloads/cas/2JBPXBMK>

IBM® LinuxONE Rockhopper 4:

<https://www.ibm.com/downloads/cas/OOVKYPZM>

IBM® LinuxONE rack mount:

<https://www.ibm.com/downloads/cas/YXBRG7NB>

IBM® LinuxONE 4 Express:

<https://www.ibm.com/downloads/cas/DBOKJNEZ>

Spec Sheets

IBM® LinuxONE Emperor 4

Machine Type 3931 Model LA1:

<https://www.ibm.com/downloads/cas/DVZEL8BR>

IBM® LinuxONE Rockhopper 4

Tower (factory frame): Machine Type 3932 Model LA2:

<https://www.ibm.com/downloads/cas/OK6DL6WL>

Rack mount: Machine Type 3932 AGL:

<https://www.ibm.com/downloads/cas/9JORJOV5>

IBM® LinuxONE 4 Express:

<https://www.ibm.com/downloads/cas/5K6WYJXG>

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We

1. Performance result is extrapolated from IBM internal tests running local inference operations in an IBM® LinuxONE Emperor 4 LPAR with 48 cores and 128 GB memory on Ubuntu 20.04 (SMT mode) using a synthetic credit card fraud detection model (<https://github.com/IBM/ai-on-z-fraud-detection>) exploiting the Integrated Accelerator for AI. The benchmark was running with 8 parallel threads each pinned to the first core of a different chip. The `lscpu` command was used to identify the core-chip topology. A batch size of 128 inference operations was used. Results may vary.
2. IBM internal data based on measurements and projections was used in calculating the expected value. Necessary components include IBM LinuxONE Emperor 4, IBM z/VM V7.2 systems collected in a Single System Image, each running RHOCP 4.10 or above; IBM Operations Manager; GDPS 4.5 for management of data recovery and virtual machine recovery across metro distance systems and storage, including Metro Multi-site workload and GDPS Global; and IBM DS8000 series storage with IBM HyperSwap. A MongoDB v4.2 workload was used. Necessary resiliency technology must be enabled, including z/VM Single System Image clustering, GDPS xDR Proxy for z/VM, and Red Hat OpenShift Data Foundation (ODF) 4.10 for management of local storage devices. Application-induced outages are not included in the above measurements. Other configurations (hardware or software) may provide different availability characteristics.
3. Compared IBM Machine Type 3932 Max 68 model consisting of a CPC drawer and an I/O drawer to support network and external storage with 68 IFLs and 7 TB of memory in 1 frame versus compared 36 x86 servers (2 Skylake Xeon Gold Chips, 40 Cores) with a total of 1440 cores. IBM Machine Type 3932 Max 68 model power consumption was measured on systems and confirmed using the IBM Power estimator for the IBM Machine Type 3932 Max 68 model configuration. x86 power values were based on Feb. 2023 IDC QPI power values and reduced to 55% based on measurements of x86 servers by IBM and observed values in the field. The x86 server compared to uses approximately .6083 KWhr, 55% of IDC QPI system watts value. Savings assumes the Worldwide Data Center Power Utilization Effectiveness (PUE) factor of 1.55 to calculate the additional power needed for cooling. PUE is based on Uptime Institute 2022 Global Data Center Survey (<https://uptimeinstitute.com/resources/research-and-reports/uptime-institute-global-data-center-survey-results-2022>). x86 system space calculations require 3 racks. Results may vary based on client-specific usage and location.
4. Compared IBM Machine Type 3932 Max 68 model consisting of a CPC drawer and an I/O drawer to support network and external storage with 68 IFLs and 7 TB of memory in 1 frame versus compared 36 x86 servers (2 Skylake Xeon Gold Chips, 40 Cores) with a total of 1440 cores. IBM Machine Type 3932 Max 68 model power consumption was measured on systems and confirmed using the IBM Power estimator for the IBM Machine Type 3932 Max 68 model configuration. x86 power values were based on Feb. 2023 IDC QPI power values and reduced to 55% based on measurements of x86 servers by IBM and observed values in the field. The x86 server compared to uses approximately .6083 KWhr, 55% of IDC QPI system watts value. Savings assumes the Worldwide Data Center Power Utilization Effectiveness (PUE) factor of 1.55 to calculate the additional power needed for cooling. PUE is based on Uptime Institute 2022 Global Data Center Survey (<https://uptimeinstitute.com/resources/research-and-reports/uptime-institute-global-data-center-survey-results-2022>). Results may vary based on client-specific usage and location.

Learn more:

<https://www.ibm.com/products/linuxone-4>

© Copyright IBM Corporation 2024
IBM Corporation
New Orchard Road
Armonk, NY 10504

IBM, the IBM logo, ibm.com, IBM Z, DS8000, FICON, GDPS, HyperSwap, Instana, Telum and zHyperlink are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis.

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. The client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.