



Escala M8-1100

The most resilient Escala server ever built to support large-scale, complex IT environments



The Escala M8-1100 high-end enterprise server, based on the Power11 processor, remains the flagship of the Power Systems lineup. It is built to support large-scale, complex IT environments with up to 256 Power11 processor cores and 64 TB of DDR5 memory.

<p>99.9999% uptime the most resilient Escala server ever¹</p>	<p>28% better server efficiency with the new Energy Efficient mode compared to Maximum Performance mode²</p>	<p>5-8 hours saved Of IT staff time per ticket</p>
---	--	---

Increased performance and efficiency

The Escala M8-1100 delivers outstanding performance, robust security, and high reliability while integrating effortlessly into existing infrastructures with minimal complexity and cost.

Designed for organisations seeking secure and efficient scalability of core operational workloads, such as databases, and a cloud infrastructure for business-critical applications the M8-1100 offers:

Expanded core capacity:

now supporting from 10 to 16 cores per socket, enabling configurations of up to 256 cores and 2,048 simultaneous threads

Performance boost:

achieves a 7–10% increase in rPerf, thanks to the innovative 2.5D ISC (Integrated Stacked Capacitor) Interposer, which delivers a 50–100 MHz frequency uplift to the Power11 processor.

Ensuring business continuity

With a **99.9999% uptime**, the **Escala M8-1100** is the most resilient server in the Escala lineup providing a solution for mission-critical workloads deployments with uncompromising reliability.

Business continuity is reinforced by the market-leading RAS (Reliability, Availability, Serviceability) features of Power and AIX, ensuring uninterrupted operations even under demanding conditions.

M8-1100 can offer spare core capabilities at the silicon level, with a pool of **idle processor cores** that can be activated in response to detected faults. This design is intended to help sustain computing

capacity and reduce the risk of processing interruptions in mission-critical environments.

The Power11 generation introduces capabilities that allow **planned maintenance without taking critical workloads offline**, ensuring **continuous availability**.

Sustainability and energy efficiency

The Escala M8-1100 delivers more performance with less energy consumption to achieve sustainability goals:

- 7% More Energy Efficient than the previous generation (M8-1000) in terms of rPerf per Watt, reinforcing its commitment to greener IT operations.
- Up to 33% Higher rPerf per Watt with Power11 compared to Power10 (maximum configuration at 100% utilisation under typical operating conditions)
- Energy Efficient Mode offers up to 28% better energy efficiency compared to Maximum Performance mode ideal for optimising power usage without compromising reliability.
- 2x more performance per watt at maximum input power on M8-1100 compared to x86 system³

Security

The Escala M8-1100 strengthens security by simplifying end-to-end encryption thanks to the hardware crypto-engines embedded in the Power11 processor. The Power11 processor's Matrix Math Accelerator architecture brings AI inferencing right where your data resides for faster insights.

The server also features quantum safe encryption for secure boot and LPM.

¹ Based upon unplanned downtime of a single Power E1180 system as calculated in the POWER11 Processor-Based Systems RAS

² Based upon internal measurements of performance per watt on servers comparing Maximum Performance Mode to Energy-Efficient Mode while running compute-, disk-, and memory-based workloads on Power11 systems with fully configured sockets and memory as follows: M8-1100 with 4x10c / 64x64GB DDIMM

³ Performance is based on QPI data as of May 15, 2025 from IDC available at <https://www.idc.com/about/qpi>

Technical specifications

System package per node	<ul style="list-style-type: none"> Up to four 5U nodes + 2U System Control Unit, 19" rack 	
System nodes	<ul style="list-style-type: none"> 1 - 4 	
Configuration options	<ul style="list-style-type: none"> 1 system node 	<ul style="list-style-type: none"> System maximum (4 system nodes)
# socket	<ul style="list-style-type: none"> 4 	<ul style="list-style-type: none"> 16
Processor module offerings (SMT8)	<ul style="list-style-type: none"> SCM 10-cores + 1 spare – 3.9 to 4.2 GHz SCM 12-cores + 1 spare – 3.9 to 4.4 GHz SCM 16-cores – 3.8 to 4.3 GHz 	
# Cores (+ spare)	<ul style="list-style-type: none"> 40 (+4), 48 (+4) or 64 	<ul style="list-style-type: none"> 160 (+16), 192 (+16), or 256
L2 cache / core	<ul style="list-style-type: none"> 2 MB 	
L3 cache / processor	<ul style="list-style-type: none"> 8 MB of L3 cache/core, total of 128MB shared each core can access to the full 128 MB 	
Performance Mode	<ul style="list-style-type: none"> Dynamic mode by default Save/Max is set in Power Management 	
System memory (RAM)	<ul style="list-style-type: none"> 64 DDIMM slots Up to 16 TB DDR5 DDIMMs 4000 MHz 	<ul style="list-style-type: none"> 256 DDIMM slots Up to 64 TB DDR5 DDIMMs 4000 MHz
Processor-to-memory band-width	<ul style="list-style-type: none"> 512 GB/s per socket w/ 4x 32GB DDIMM, 1024 GB/s per socket otherwise 	
System memory min / max / (min % active)	<ul style="list-style-type: none"> 256 GB / 16 TB / (50%) 	<ul style="list-style-type: none"> 1 TB / 64 TB / (50%)
Integrated PCIe adapter slots	<ul style="list-style-type: none"> 8 PCIe Gen5 	<ul style="list-style-type: none"> 32 PCIe Gen5
PCIe I/O Expansion Drawers	<ul style="list-style-type: none"> Up to 4 (12 PCIe adapter slots each) 	<ul style="list-style-type: none"> Up to 16 (12 PCIe adapter slots each)
Internal storage	<ul style="list-style-type: none"> 4 slots for NVMe U.2 	<ul style="list-style-type: none"> 16 slots for NVMe U.2
Flexible Service Processors	<ul style="list-style-type: none"> 2 	
HMC ports	<ul style="list-style-type: none"> 2 	
USB ports	<ul style="list-style-type: none"> USB PCIe adapter must be used for enabling USB access 1 x USB 3.0 in System Control Unit 	
rPerf/core (# of cores): rPerf	<ul style="list-style-type: none"> 37.5 (40): 1,499 38.5 (48): 1,848 35.7 (64): 2,285 	<ul style="list-style-type: none"> 37.5 (160): 5,999 38.5 (192): 7,395 35.7 (256): 9,141 and 0.506 rPerf/Watts

PowerVM Enterprise – Hypervisor (included)	<ul style="list-style-type: none"> • LPAR, Dynamic LPAR; Virtual LAN (Memory to memory inter-partition communication) • Up to 20 Partitions per core; Virtual I/O Server; Shared Processor Pools; Shared or Dedicated Capacity; Live Partition Mobility (LPM) and Active Memory Sharing (AMS); NovaLink (virtualisation management agent). • Resource Groups which can be combined with shared processor pools (MSPPs) and also offers LPM.
Reliability, Availability, and Serviceability (RAS) features	<ul style="list-style-type: none"> • Six 9's availability (99.9999% uptime). Spare Cores for increased availability (4 spare cores per drawer). Quantum safe encryption for secure boot and LPM. First failure data capture (FFDC). Processor instruction retry. L2 and L3 cache ECC protection with cache line-delete. Core checkstops. Dynamic processor deallocation. Chip kill protection for x4 DDIMMs, with DRAM sparing. Processor fabric and memory buses retry with data lane sparing and ½ bandwidth mode. High-speed inter node cables, with passive components and advanced fault isolation diagnostic capabilities. Guided FSP and SMP cable installation. Concurrent repair of the external SMP cable. Redundant phase and spare phase for VRMs supplying processors. Spare PMIC for DDIMM power regulation. Redundant system clocks with dynamic failover. Redundant, hit-swappable power supplies and cooling fans. Concurrent add or repair of I/O drawers. Extended error handling on PCIe slots. Hot-plug and blind-swap PCIe adapter slots. Concurrent repair of the Op-Panel. Concurrent repair of the Time of Day Battery. Selective dynamic firmware updates.
Capacity on Demand	<ul style="list-style-type: none"> • Mobile CoD, Power Enterprise Pool as an option
Software support	<ul style="list-style-type: none"> • AIX: 7.2, 7.3 or later • VIOS 4.1.0, VIOS 4.1.1 • RHEL: 8.6, 9.4, 9.6 and 10 • SLES: 15 SP6 or later • Red Hat OpenShift Container Platform: 4.19 or later • IBMi: 7.4, 7.5, 7.6 or later - Contact Bull representative for more details
Warranty	<ul style="list-style-type: none"> • 1 year

Bull and IBM: a perfect fit

Since 1992, Bull and IBM have built a unique relationship, with IBM leading to a highly productive technological cooperation. This has fundamentally strengthened the AIX ecosystem, by regularly generating innovations, in areas such as scalability, RAS, virtualisation and cloud enablement.

Connect with us
bull.com



Bull is a registered trademark © Copyright 2026, Bull SAS – All rights reserved.