

Miért más az IBM Power, mint az x86?

Angyal László
Infrastructure Specialist
laszlo.angyal@ibm.com




1B


*new applications are expected by 2028¹
...and enterprises are drowning in complexity*



Silos



Data



Outages



*Cost
Optimizations*



*Skills
Shortages*

*...which makes keeping critical
systems continually running a
near impossible task...*

*...and these challenges will
only intensify as enterprises
strive to remain competitive in
the AI era*

POWER világ



(jellemzően)

- Kevesebb, „nagyobb” rendszer
- Inkább vertikális skálázás (scale-up)
- Konszolidált workloadok
- Integrált, nagyvállalati rendszer

x86 világ




(jellemzően)

- Sok kicsi szerver
- Horizontális skálázás (scale-out)
- Elosztott rendszerek
- „Commodity” infrastruktúra


Nem minden CPU core egyforma! Miért erősebb egy POWER core?

POWER világ



- Akár SMT8 (értsd: 8 thread / core)
- Akár 2x per-core teljesítmény
- Kevesebb core >>> nagyobb hatékonysággal


x86 világ



- Max. SMT2 (értsd: 2 thread / core)
- Kisebb per-core teljesítmény
- Skálázási lehetőség: sok-sok core / socket


A teljesítmény közvetlen hatása a költségekre

POWER példa



- 4 core (SMT8)
- 4 „per-core” licenc szükséges

x86 példa




- 8 core (SMT2)
- 8 „per-core” licenc szükséges

Ugyanaz a workload → kevesebb core → alacsonyabb licenc költség


Virtualizáció másképp: erőforrás-alapú megközelítés

POWER világ



- Dinamikus erőforráskiosztás (pl. I/O, uncapped vCPU, Active Memory Sharing)
- CPU/memória „okos” megosztása
- Magasabb kihasználtság

x86 világ



(jellemzően)

- VM-ek kiosztása fix erőforrásokkal
- CPU és memória „előre lefoglalva”, sokszor „üresjáratban”
- Gyakan alacsony kihasználtság

Nem pusztán VM-eket adunk, hanem biztosítjuk a szükséges kapacitást

Oracle consolidation: Hardware vs. Software virtualization



POWER

Sharing CPU

VM's idle CPU is used



Bare metal Virtualization



PowerVM Virtualization is integrated by design

No CPU over-provisioning → better CPU utilization (EC,VP, capped/uncapped)
Higher sustained Utilization
Better resource control and management

Exadata with Virtualization

- Weak and poor virtualization features and resiliency
- Maintenance requires all VMs shut down on same DB Node
- Limited resource sharing → more licenses
- Scale memory and local storage **require** VM Restart
- No VM Mobility (cold or live) capabilities

PowerVM Virtualization


- Workloads isolated by firmware
- Mixed environments, prod, dev, test, ..
- Run DBs and Applications
- Flexible, on the fly resource management
- Easier maintenance
- Simple operations
- No overhead, VMs scalability

$$\text{Peak} \sum (\text{VMs}) < \sum \text{Peak}(\text{VMs})$$


x86

Mapping CPU

VM's idle CPU is wasted



Additional x86 SW Virtualization



1 to 1 Mapping
2 Virtual CPUs : 1 Physical Core, Idle CPU in VMs is wasted

"CPU over-provisioning is possible, but workload performance conflicts can arise if all guests become fully active"
<https://www.oracle.com/a/tech/docs/exadata-kvm-overview.pdf#page=13>


Klasszikus x86 „kisszerveres” környezet

Átlag <20% erőforráskihasználtság

- Configuration planned for peaks (50% unused?)
- Configuration planned for growth (20% unused?)
- System waits for I/O and memory access even when it is working (20% unused?)

x86 Linux	Avg CPU%	Peak CPU%
app	8%	10%
db	15%	25%
email	10%	20%
net	7%	10%
SysMgt	7%	10%
web	10%	15%
f&p	7%	10%
oos	5%	8%
Overall	8%	15%

* IBM Scorpion study aggregate results




Single workload on a single system

- Average Utilization: 20.7%
- Peak: 79%


Result is that 80% of the hardware, software, maintenance, floor space, and energy that you pay for, is wasted

For A Single Workload, We Require A Machine Capacity 6.0x The Average Demand




Assumes standard deviation or sigma = 2.5

When We Consolidate 4 Workloads We Only Require 3.5x Average Demand




Assumes standard deviation or sigma = 2.5

When We Consolidate 16 Workloads We Only Require 2.25x Average Demand



Assumes standard deviation or sigma = 2.5

When We Consolidate 144 Workloads We Only Require 1.42x Average Demand



Assumes standard deviation or sigma = 2.5

Szerverkonszolidáció virtualizált, nagy (POWER) szerverre

Klasszikus x86 forgatókönyv esetében: 64 workload fut 64 különálló szerveren

- Átlagos kihasználtság: 20.7% >>> sokat „löttyög” üresjáratban
- Csúcskihasználtság: 79%

POWER forgatókönyv esetében: 64 workload fut egyetlen POWER szerveren*



- Átlagos kihasználtság: 61% >>> „okos” kihasználtság
- Csúcskihasználtság: 78%

64x4 = 256 CPU helyett 72 CPU elegendő (3.5:1 arány)

A POWER szerverek jellemzően magasabb kihasználtsággal működnek.

* Egymástól független workloadokat feltételezve

**64 to 1 Systems Consolidation
(72 CPUs)**



RAS

(R)esiliency/Reliability – Minimize potential errors

- Minimizes outage opportunities through hardening in manufacturing
- Minimize total components and use higher grade components
- Error Detection and Self-Healing



(A)vailability – Minimize unplanned downtime

- Predict, fix and keep running despite problems
- Self diagnosing – work around faulty components/self heal



(S)erviceability – Minimize outage/downtime

- Accurate problem diagnostics
- Concurrent repair, guided repair
- End to end service – from problem determination to call home to closing service action

Why Power has Better RAS

End-to-End RAS with rich history and capabilities built into ...

- Hardware components
- Interfaces
- Firmware
- Hypervisor
- Software (OS)
- Management stack

IBM Power proactively develops tests for end-to-end RAS ...

- Everyday use through extensive modeling and simulation
- Resolving unexpected errors through error injection to validate end-to-end error resolution
- Natural disasters through earthquake, shock and vibration, and connector wear
- Environmental issues with Temperature and Humidity chamber testing

Support & Service

- IBM Call Home Support through [IBM Electronic Service Agent™](#)
- [IBM Clientvantage](#)
- [IBM Support Insights](#) delivers actionable indicators for preventive maintenance



T&H smaller chamber



4inch free fall drop
non-operational shipping
simulation



T&H walk-in chamber



Earthquake simulating
Richter 8.3 scale
with the workload running
un-interrupted




Operational S&V
with the workload
running un-
interrupted

Power RAS Involves More Than Just the Processor


■ IBM develops, tests, integrates the entire stack for RAS

- ▶ I/O drawers / memory management unit
- ▶ Processors and all chips in CEC
- ▶ Hypervisor (PowerVM) and VIOS
- ▶ Device drivers, PCI adapters
- ▶ Operating system (AIX, System i, Linux)
- ▶ Middleware and Clustering software

Power Stack - Integrated RAS



Intel Processor RAS



RAS – Power11

- Standard
- Optional
- Not Available

	S1122/S1124	E1150	E1180
Redundant / Hot Swap Power Supplies	●	●	●
Pluggable VRMs with redundant/spare voltage phases	—	● N+1	● N+2
OCC error handling w/ power safe mode	●	●	●
EPOW reporting	●	●	●
TOD Battery Concurrent Maintenance	—	●	●
Redundant Fans/Blowers	●	●	●
Hot Swap Fans / Blowers	●	●	●
NVMe drives and/or Media Concurrent Maintenance	●	●	●
PCI Adapter Concurrent Maintenance	●	●	●
EEH (Enhanced Error Handling) Support	●	●	●
Concurrent Op Panel Repair	LCD only	●	●
Redundant Service Processors	N/A	N/A	●
Dynamic Service Processor Failover	N/A	N/A	●
Auto Service Proc. Reset/Reload	●	●	●
Redundant System Clocks	—	●	●
Dynamic System Clock Failover	—	● Int. Spare	●
Redundant VPD – system VPD	● FW Copy	● FW Copy	●

RAS – Power11

- Standard
- ◻ Optional
- Not Available

	S1122/S1124	E1150	E1180
Concurrent FW Update – service pack	w/ Opt HMC	w/ Opt HMC	●
Memory Preserving IPL	●	●	●
HW and OS Dump	●	●	●
Processor Instruction Retry	●	●	●
Spare core	● (1)	●	● (2)
Predictive De-allocation (such as, core)	●	●	●
Core check stop	●	●	●
Dynamic Core Init. Changes	●	●	●
L2/L3 Line Delete	●	●	●
Fabric Bus CRC/Retry	●	●	●
SMP Cable Fault Isolation	N/A	N/A	●
Ext SMP cable concurrent maintenance	N/A	N/A	●

(1) Except 4 & 10 core eSCM

(2) Except 16 core option

RAS – Power11

- Standard
- Optional
- Not Available

	S1122/S1124	E1150	E1180
Memory ECC/Chip Kill Handling	● (1) (2)	● (2)	● (2)
Spare DRAM	4U DDIMMs – 2 spare DRAMs per rank	4U DDIMMs – 2 spare DRAMs per rank	4U DDIMMs – 2 spare DRAMs per rank
Integrated N+1 Power regulator	4U DDIMMs only	■	■
Memory Scrubbing	●	●	●
Processor to Memory Buffer Recovery	● (3) (4)	● (3) (4)	● (3) (4)
Memory (DRAM) Row Repair	● dynamic (detect and deploy at runtime)	● dynamic (detect and deploy at runtime)	● dynamic (detect and deploy at runtime)
Predictive Dynamic Memory De-allocation	●	●	●
Active Memory Mirroring for Hypervisor	■ Priced option	■ Priced option	●

- (1) DDIMM 2U: 1st chip kill fixed with ECC, DDIMM
- (2) 4U : 1st & 2nd chip kill fixed w/spare, 3rd fixed w/ECC
- (3) Protected by CRC with retry for soft bit flips
- (4) ½ bandwidth mode (from 8 to 4 lanes) for hard lane fails.

Power11: New Horizons of Reliability

Six 9's
of uptime

Power11 is designed to be the most reliable server in the history of the IBM Power platform with 99.9999% of uptime.

A magas rendelkezésre állás nem extra, hanem alapértelmezett.

1. Based upon unplanned downtime of a single Power E1180 system as calculated in the IBM Power RAS Whitepaper.

Amikor a leállítás nem opció...

NASA, Kennedy Space Center, Artemis II kilövés – a Hold felé:

- Kritikus rendszerek valós idejű működése
- Nincs lehetőség kiesésre
- Maximális megbízhatóság szükséges

Vannak környezetek, ahol a hiba egyszerűen nem megengedhető:

- POWER szervereken futó kilövés irányító rendszer
- Több százezer adatpont valós idejű feldolgozása (telemetria, szenzorok)
- Kritikus rendszerek folyamatos monitorozása a teszteltől a kilövésig
- Azonnali eltérésfelismerés és döntéstámogatás



Hatékonyság: kevesebb infrastruktúra, nagyobb teljesítménnyel

POWER világ



- Kevesebb szerver
- Magasabb kihasználtság
- Alacsonyabb energiaigény
- Hatékonyabb működés

Tipikus x86 környezet



- Több szerver
- Több energiafogyasztás
- Több szoftverlicenc
- Magasabb üzemeltetési költség

Kevesebb rendszer → alacsonyabb költség → egyszerűbb üzemeltetés

<https://powertco.precisionitinc.com/>

Power11: More performance with less energy consumption

2X

More performance per watt^{1,2} at maximum input power on IBM Power E1150 vs compared x86 system

1. Performance is based on Quantitative Performance Index (QPI) data as of May 15, 2025 from IDC available at <https://www.idc.com/about/qpi>. IBM Power E1150 (4x30c Power11) QPI of 241,000E versus HPE Compute Scale-up Server 3200 (4x60-core Intel cores at 1.9GHz) QPI of 208,898 and utilizations of 75% for E1050 based on IBM Power Performance Utilization Guarantee and 40% for x86 based on IT Economics assessments of customer environments running a total of 13,800 x86 cores where average measured IBM peak utilization is 16%. Peak utilization of 16% is derived from a weighted average of x86 workloads in four large enterprise client IT environments using a total of the compared 13,861 x86 cores in production and test environments. For additional information on x86 workload analysis contact the IBM IT Economics team, IT.Economics@us.ibm.com
2. Energy consumption is based on maximum input power :IBM Power E1050 with maximum power of 5,200 W <https://www.redbooks.ibm.com/redpapers/pdfs/redp5684.pdf> Superdome Flex 280 with maximum power of 4,740 W https://www.hpe.com/psnow/doc/a50004268enw.html?jumpid=in_pdp-psnow-q5

Pools

Erőforrások megosztása több rendszer között

POWER világ



Power Enterprise Pool

- Több szerver az erőforrás poolban
- Dinamikus kapacitás átcsoportosítás ill. bekapcsolás
- Hatékonyabb kihasználás

Hagyományos szerverek




- Szerverenként elkülönített CPU és memória erőforrások
- Statikus kapacitás
- Nehéz átcsoportosítás

Nem különálló szerverekben, hanem összesített kapacitásban gondolkozunk.

Power Enterprise Pools 2.0 with *Shared Utility Capacity*


Increase responsiveness and deliver cloud economics on-premises across a collection of Power11 and/or Power10 systems

Traditional systems




- Repeated procurement cycles increasing time to value
- Sub-optimal resource utilization with over or under-provisioned capacity
- High upfront investment
- Limited flexibility to scale up and down over time

Power9 & Power10 systems



- Offer seamless growth and instant, non-disruptive response to business changes
- *Capacity Upgrade on Demand* – permanent activations
- *Elastic Capacity* – temporary activations by the day
- *Mobile Capacity* – share permanent activations within a pool

Power Private Cloud w/ Shared Utility Capacity




- *Shared Utility Capacity* takes Mobile and Elastic Capacity to the next level
- Optimize resource utilization to reduce compute capacity costs by 20-30% with seamless, multi-system resource sharing
- Lower initial system price by over 50% and access fully-active standby resource as pay-per-use capacity, by the minute
- Real time & historical usage monitoring and analysis via the Cloud Management Console

Power Enterprise Pools 2.0 delivers *Shared Utility Capacity*

Deploy a Power Enterprise Pool across a collection of Power11 systems

- Purchase Power11 systems with Power Enterprise Pools Subscriptions
- *Base Capacity* from each system is aggregated and effectively shared across the pool
- Remaining capacity on each server is activated when a pool is started and made seamlessly available by the minute as *Metered Capacity*
- IBM Cloud Management Console with HMC monitors minutes of resource use that exceed a pool's Base capacity
- Minutes of Metered Capacity consumed may be invoiced at the end of each month* or debited against pre-paid Capacity Credits



Base and Metered Capacity elements

- Processor activations
- Memory activations*
- AIX and IBM i entitlements
- SLES subscriptions from IBM***
- RHEL subscriptions from IBM***
- OpenShift Container Platform subscriptions from IBM ***
- PowerHA SystemMirror for AIX & IBM i

****Base Memory activations are available on high-end and midrange systems**

***Subscriptions must be purchased as IBM Program Products with supported features included.

Power Enterprise Pools 2.0

Autonomous, multi-system resource sharing and fully-active standby resource, available by the minute



500+

Pools in production across the world


6 years

of proven performance in mission-critical environments

Up to 96

Power11 and/or Power10 systems supported in each pool

Power11 systems are *primed* for Power Enterprise Pools



Power11 offers a [simplified Processor and Memory Activation structure](#) to streamline configuration, deployment and use

- One Processor Activation price per model
- One Memory Activation price per model
- [Reduced prices per-core & per-GB on Power11 vs. Base/Mobile Activations on Power10](#)

A new Power Enterprise Pools Subscription is now used to enable Mobile Capacity or Shared Utility Capacity on a Power11 system


- 1-5 Year Power Enterprise Pools Subscription PIDs are priced per-socket by model, required when #EB35 (Mobile enablement) or #EP20 (Pools 2.0 enablement) added to a system or MES order
- No premium per-core or per-GB activation charges required to convert to Power Enterprise Pools

Up to 96 Power systems and up to 4,000 VMs may be supported in a single Power Enterprise Pool

Power11 systems may co-exist and share Mobile or Base Capacity resources with Power10 systems in the same Power Enterprise Pool

Power E1080 and E1050 systems may upgrade to Power11 processor technology via a Model Upgrade (same serial number), during which their Mobile or Base Capacity remain available as pool resources

Power Enterprise Pools 2.0 with Shared Utility Capacity



How it works

Purchase or lease one or more eligible 2-16 socket Power11 and/or Power10 systems, each with at least 1 Processor Activated and 256GB Memory Activated

Acquire a corresponding Power Enterprise Pools subscription to enable Shared Utility Capacity on each machine to be added to a pool.

Generate a new Pool ID via the IBM Entitled Systems Support portal, designating an initial system to associate with the new Pool ID. Then start the pool via the Cloud Management Console and add additional eligible systems to the pool.

Once a pool is started, each associated system's Processor Activations, Memory Activations and Operating System license entitlements become the pool's Base Capacity, which is aggregated and effectively shared across the pool

Remaining physical capacity on each server is activated automatically when a pool is started and becomes seamlessly available as Metered Capacity by the minute


IBM Cloud Management Console (CMC) with HMC monitors minutes of resource use that exceed the pool's Base capacity. A pool's resource consumption may be displayed in the CM real-time and historically by system, by resource and/or by VM.

Minutes of Metered Capacity consumed may be debited against pre-paid Capacity Credits on account or may be invoiced at the end of each month*

*Monthly invoicing of Metered Capacity is available in North America, on-shore Europe, and select Latin America, Middle East and African countries

NDA Required - Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.


In a traditional system infrastructure, each system must be individually provisioned and maintained for peak capacity



Power Enterprise Pools 2.0



IBM Cloud Management Console




When a Pool is started, all capacity above Base is activated and monitored by the CMC.

Minutes of Metered Capacity consumed above a Pool's aggregate Base are charged by the minute

Base Capacity is aggregated and effectively shared across the Pool


In a Power Private Cloud environment, each system may be configured optimally to support its expected average utilization, based upon the system's role in the infrastructure

Power Enterprise Pools 2.0




In a Power Private Cloud environment, each system may be flexibly configured with minimal Base Capacity to significantly reduce the initial system price, and additional capacity may be consumed by the minute or added as Base at a later time

Potential Metered Capacity consumption charges on one system may be offset by available (idle) Base Capacity on another system



Shared Utility Capacity provides unique innovation designed to optimize resource utilization and reduce the expense of handling volatile/peak demand across a collection of Power servers, while maintaining service level



Deploy a private cloud infrastructure on-premises using Shared Utility Capacity across a pool of Power11 systems



Optimize resource utilization to reduce compute capacity costs by 20-30% with seamless, multi-system resource sharing



Lower initial system price by over 50% and access fully-active standby resource as pay-per-use capacity, by the minute



Leverage Real time & historical usage monitoring and analysis via the Cloud Management Console

Cloud-like agility and economics with leadership business continuity and security

Miért választják a Power platformot?

A Power platformot kritikus üzleti rendszerek futtatására optimalizálták.

Megbízhatóság	Teljesítmény	Hatékonyság	Rugalmasság
▪ Minimális downtime ▪ Kritikus rendszerekhez tervezve	▪ Magas „per core” teljesítmény ▪ Hatékony párhuzamos feldolgozás	▪ Okosabb gazdálkodás az erőforrásokkal ▪ Alacsonyabb licenc- és energiaigény	▪ Dinamikus erőforráskezelés ▪ Cloud-szerű működés on-prem környezetben

Key competitive advantages of Power11 servers vs. the latest large x86-based systems

 Memory Power11 uses DDIMM highly reliable memory Security by design, Transparent Memory Encryption x86 servers use industry-standard memory	 Flexibility and Virtualization All Power11 servers feature highly flexible PowerVM Hot memory resizing (VMware can't do that with SAP) E1180 enjoys 32 production LPARs per box, better density, etc.	 Persistent Memory Native feature of Power11, key for SAP HANA startup time Intel discontinued the development of Optane persistent memory No persistent memory is available on Sapphire Rapids and later
 Memory Bandwidth Power11 features 16 channels per socket whereas Sapphire Rapids has 8 and Granite Rapids SP have 8 Therefore, Power11 easily retains the bandwidth advantage As per ITIC 2023 Global Server Hardware, Server OS Reliability Survey, IBM Power Servers have 100x better application uptime compared to nearest x86 vendor	 Scaling SAP HANA scales up to 32TB on Power11 Only 16TB in virtualized VMware environment, with many restrictions	 TCO and Sustainability IBM Power advanced virtualization + Tetris sizing methodology yields better density and usage, and therefore lower TCO and energy consumption Comprehensive set of customer references and case studies

Milyen operációs rendszerek futtathatók IBM Power szervereken?

Támogatott, natívan
futtatható operációs
rendszerek

- IBM i (OS/400) 
- AIX 
- VIOS
- Red Hat Enterprise Linux (RHEL) 
- SUSE Linux Enterprise Server 
- Ubuntu LTS 

Container platform:

- Red Hat OpenShift Container Platform 

Nem minden elterjedt OS
érhető el Poweren

- Van, ami fut natívan vagy nyakatekert emulációban, de nincs hozzá komolyabb gyártói támogatás
- A terméktámogatás megléte kulcsfontosságú éles üzemeltetési környezetben

Total Cost of Ownership (TCO) vs. Total Cost of Acquisition (TCA)

IBM Power hatása


- Kevesebb core >>> alacsonyabb licencdíj
- Konszolidáció >>> kevesebb rendszer
- Jobb hatékonyság >>> alacsonyabb energiaigény
- Egyszerűbb működés
- Minimális kiesés
- Kevesebb sérülékenység >>> nagyobb biztonság

Költséget meghatározó tényezők

- CPU alapú licencelés
- Infrastrukturális méret (pl. szerverzetek száma)
- Energiafogyasztás
- Üzemeltetési komplexitás
- Downtime költsége

A különbség nem csak technikai, hanem teljes életciklus költségben jelentkezik.

<https://powertco.precisionitinc.com/>



Thank you

© Copyright IBM Corporation 2026

IBM and the IBM logo are trademarks of IBM Corporation, 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 ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time.

Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.

Examples presented are illustrative only. Actual results will vary based on client configurations and conditions and, therefore, generally expected results cannot be provided.

Not all offerings are available in every country in which IBM operates.

It is the user's responsibility to verify the operation of any non-IBM products or programs with IBM products and programs. IBM is not responsible for non-IBM products and programs.

The client is responsible for ensuring compliance with all applicable laws and regulations. IBM does not provide legal advice nor represent or warrant that its services or products will ensure that the client