

# Power11 Performance Best Practices

## A brief checklist

Description	Instructions
Ensure firmware is current	<a href="#">Fix Central</a> provides latest updates. Install the latest version of the system firmware. Use the FLRT tool to obtain the recommended levels for a given platform. <b>NOTE:</b> Ensure required HMC level is installed when updating F/W.
Memory DIMMs	Follow proper memory plug-in rules
Ensure OS level is current	Fix Central provides the latest updates for AIX, IBM i, VIOS, Linux, HMC and F/W. In addition to that, the FLRT tool provides the <b>recommended levels</b> for each H/W model. Use these tools to maintain your system up to date.
LPM	Use of 10Gb or faster dedicated network for LPM is preferred.
Sizing a system	When migrating to Power11, we recommend considering using SMT8, and size the LPARs based on the SMT8 rPerf values; in many instances, this will likely reduce the number of VPs required.
Right-size your Shared LPARs	<ul style="list-style-type: none"> <li>Assign entitled capacity (EC) to sustained peak utilization for LPARs with critical SLA requirements</li> <li>Assign EC to average utilization and number of virtual CPUs to peak utilization (physical core consumption) for LPARs with non-critical SLA</li> <li>Ensure the average LPAR utilization is equal or less than 75% of the entitled capacity</li> </ul>
Resource Groups	Consider use of Resource Groups and Resource Groups Advisor to isolate workloads.
Partition Placement	Current FW levels ensure optimal placement of the partitions. However, if constant DLPAR operations are executed on partitions on the CEC, it is recommended the use DPO to optimize placement (requires current level of F/W).
Compilers	Use latest compiler versions: <ul style="list-style-type: none"> <li><a href="#">IBM Open XL C/C++ for AIX 17.1.3</a></li> <li><a href="#">IBM Open XL Fortran for AIX 17.1.3</a></li> <li><a href="#">IBM Advance Toolchain for Linux on Power 16.0-5</a></li> <li><a href="#">Open SDK for Rust on AIX 1.84 (and above)</a>, GCC 14.3 (and above)</li> </ul> IBM i compilers in Rational Development Studio (5770-WDS)
Java	Java applications using Linux and AIX can seamlessly take advantage of new Power11 ISA features on operating systems running in Power11 mode by using the Java runtime versions listed below or newer: <ul style="list-style-type: none"> <li>Java 8               <ul style="list-style-type: none"> <li>IBM SDK 8 SR8 FP45</li> <li>IBM Semeru Runtime Open Edition 8u452: openj9-0.51</li> </ul> </li> <li>Java 11               <ul style="list-style-type: none"> <li>IBM Semeru Runtime Certified Edition 11.0.27: openj9-0.51</li> <li><a href="#">IBM Semeru Runtime Open Edition 11.0.27: openj9-0.51</a></li> </ul> </li> <li>Java 17               <ul style="list-style-type: none"> <li>IBM Semeru Runtime Certified Edition 17.0.15: openj9-0.51</li> <li>IBM Semeru Runtime Open Edition 17.0.15: openj9-0.51</li> <li><a href="#">OpenJDK 17</a></li> </ul> </li> <li>Java 21               <ul style="list-style-type: none"> <li>IBM Semeru Runtime Certified Edition 21.0.7: openj9-0.51</li> <li>IBM Semeru Runtime Open Edition 21.0.7: openj9-0.51</li> </ul> </li> </ul> Java runtime can be download at : <a href="https://developer.ibm.com/languages/java/semeru-runtimes/downloads/">https://developer.ibm.com/languages/java/semeru-runtimes/downloads/</a>
IBM i	Ensure IBM i is at a currently supported version with appropriate current Technology Refreshes applied. (updates can be found <a href="#">here</a> )

Description	Instructions
AIX	<b>Tunables</b> <ul style="list-style-type: none"> <li>Tuning a VIOS is not recommended unless directed by VIOS/AIX support.</li> <li><b>Restricted tunables should not</b> be modified (unless directed by AIX/VIOS development)</li> <li>Tunables should be at default unless otherwise recommended by AIX support.</li> <li>If tunables have been modified from default, they should be re-evaluated prior to moving to a different processor platform.</li> </ul> <b>CPU usage</b> On Power11, the AIX OS system is optimized for best raw throughput at higher CPU usage when running with dedicated processors. When running with shared processors, the AIX OS system is optimized to reduce CPU usage (pc). If the Client requires to further reduce CPU usage (pc), use the schedo tunable vpm_throughput_mode to tune the workload and evaluate the benefits of raw throughput vs. CPU usage.
VIOS configuration	<ul style="list-style-type: none"> <li>Use latest VIOS versions: VIOS 3.1.4.60, or later, VIOS 4.1.1.10 or later, VIOS 4.1.0.40 or later</li> <li>Shared Ethernet adapters using a 10Gb, 40Gb or 100Gb Mellanox adapter as a backing device should enable the "flip_n_run" attribute, via chdev, on the network adapter port: chdev ent# -a flip_n_run=yes</li> <li>For vFC, ensure the number of client connections does not exceed the limits of the physical adapter.</li> <li>For vSCSI disks, ensure the queue_depth for virtual disks is less than or equal the queue_depth of the physical disk in the VIOS.</li> </ul> If configured with shared processors: <ul style="list-style-type: none"> <li>Assign total entitlement of all VIOS partitions to be 10-15% of cores in shared pool and assign CPU ratio of 2:1 (vCPUs:ec). Refer to the PowerVM Best Practices for additional recommendations.</li> </ul> Assign uncapped mode and set variable weight capacity of VIOS partition higher than all client LPARs serviced by VIOS
Virtual Ethernet adapters on AIX	On AIX LPARs, with very high network load, the following tuning can improve bandwidth. <b>Note:</b> CPU utilization may increase. chdev -l ent<#> -a queues_rx=4 -a queues_tx=12. <b>Note:</b> Setting rx_queue>0 disables SEA thread by default.
vNIC	On AIX LPARs with very high network load, enable the following settings for a better response: chdev -l ent<#> rx_que_num=8 -a rx_que_elem=2048 -a tx_que_elem=1024 This tuning will consume more pinned memory. Ensure VIOS has enough memory and CPU resources. On Linux LPARs with very high network load, enable the following settings for a better response: ethtool -L eth<#> rx 16 tx 16 ethtool -G eth<#> rx 4096 tx 2048
100GbE adapter	On AIX LPARs: For maximum bandwidth, incrementally increase the buffer size from 1024 to 4096 in steps of 1024, and double the number of transmit and receive queues. chdev -l ent<#> -a queues_rx=<size> -a queues_tx=<size> -a rx_max_pkts=<size> On RHEL8.4 and later: For network bandwidth sensitive workloads, we recommend increase the transmit and receive queue size to ethtool -L enP<#> combined 64 ethtool -G enP<#> rx 4096 tx 2048

## **Best Practices documents and References:**

### **POWER**

- [Power Virtualization Best Practices](#)
- [IBM Power Systems Performance Report \(Enhanced rPerf\)](#)

### **IBM i**

- [IBM Power Performance Capabilities Reference](#)
- [IBM i on Power – Performance FAQ](#)
- [IBM i Optimization Guide](#)

### **AIX and VIOS**

- [IBM AIX MPIO: Best practices and considerations](#)
- [Power system to AIX Maps](#)

### **Java / WebSphere**

- [IBM WebSphere Application Server Performance Cookbook](#)

### **Databases**

- [AIX and Oracle Database Performance Considerations](#)
- [The latest performance enhancements for Db2 for i](#)

### **Advisor Tools:**

- [VIOS Advisor](#)
- [Resource Groups Advisor](#)
- [IBM Systems Workload Estimator](#)

### **Redbooks:**

- [Power11 based systems](#)
- [Power Systems Private Cloud with Shared Utility Capacity](#)

### **Software Updates:**

- [IBM i Technology Refreshes](#)
- [IBM i Fixes](#)
- [Fix Central \(for Firmware, AIX and VIOS updates\)](#)
- [Fix Level Recommendation Tool \(FLRT\)](#)

The latest copy of this document can be found under the “Service and Support Best Practices” website:

<https://www.ibm.com/support/pages/node/883882>

For suggestions/changes to this document, please contact [quinnw@us.ibm.com](mailto:quinnw@us.ibm.com)