

# Power9 Performance Best Practices

## A brief checklist

This document is intended as a short summary for customers on key items that should be looked at when planning a migration. For a more in-depth and more complete set of recommendations, please refer to the document links provided on the second page

Description	Instructions
Ensure firmware is current	Fix Central provides latest updates. Latest F/W levels as of this writing : FW950 for ALL Power9 systems. 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	<b>All POWER:</b> 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.
SMT8	To take full advantage of the performance of Power9 CPUs, we recommend clients evaluate the use of SMT8. Proper sizing is also recommended to maximize the Power9 improvements. We recommend when moving to SMT8 to reboot the partition to get the best performance of this change.
40GbE adapter	RHEL7: For network bandwidth sensitive workloads, we recommend increase the receive queue size from 1024 to 8192.
Sizing a system	<ul style="list-style-type: none"> <li>When migrating to Power9, 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.</li> <li>Use Workload Estimator (WLE) for sizing LPARs for CPU consumption as it provides better sizing results.</li> </ul>
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>
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	<ul style="list-style-type: none"> <li>IBM XL C/C++ for AIX V16.1 and XL Fortran V16.1 added support for Power9. Also adds support for C++11 and C++14.</li> <li>IBM XL C/C++ for Linux V16.1.1 &amp; XL Fortran V16.1.1 support for P9 ISA</li> <li>Advanced Toolchain: 11.0-3 and later</li> <li>gcc: Version 7 of gcc is recommended for P9 ISA support. Also includes support for "-mtune=power9"</li> </ul>
vNIC	<ul style="list-style-type: none"> <li>AIX - Change the following vNIC interface settings as follows: chdev -l ent# -a rx_que_num=8 -a tx_que_num=8 -a rx_que_elem=2048 -a tx_que_elem=1024 -a use_rec_q_val=no</li> <li>Linux distros - Update to latest kernel. Also, set rx/tx queues to maximum ethtool -L eth# rx 8 tx 8</li> </ul>
Java	<ul style="list-style-type: none"> <li>IBM JDK8 SR5 is the minimum level to exploit Power9</li> <li>Open JDK 1.8 provides partial support for P9 ISA</li> <li>Use of 64k size pages normally increases application performance</li> </ul>
IBMi	Ensure Technology Updates are current (see link below)

Description	Instructions
LPM	<ul style="list-style-type: none"> <li>VIOS 3.1 runs on native POWER9 mode</li> <li>VIOS 2.2.5 and above provide performance improvements over prior versions.</li> <li>Use of 10Gb dedicated network for LPM is preferred.</li> <li>Use of dual MSPs can improve performance (minimum requirements: VIOS 2.2.5 and FW860)</li> </ul>
AIX Tunables / VIOS Tunables	<ul style="list-style-type: none"> <li>Tuning a VIOS is not recommended unless directed by VIOS/AIX support.</li> <li><b>Restricted</b> tunables <b>should not</b> be modified (unless directed by AIX/VIOS development)</li> <li>Tunables <b>should not</b> be migrated across H/W or AIX levels.</li> </ul>
AIX CPU utilization	The AIX OS system is optimized for best raw throughput at higher CPU usage. If the customer requires to reduce CPU usage (pc), use the <b>schedo</b> 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>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>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> <li>Assign uncapped mode and set variable weight capacity of VIOS partition higher than all client LPARs serviced by VIOS</li> </ul> </li> <li>For performance and flexibility, it is recommended to use IBM i to virtualize internal storage to IBM i. If you must use VIOS, follow the wiki at the following <a href="#">link</a>.</li> <li>For vFC, ensure no more than 64 client connections total per physical fcs adapter on the VIOS. Also, ensure no more than 64 storage ports configured per vFC adapter on the client. These are physical limits; practical limits may differ based on workload.</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> <li>For vSCSI adapters, ensure you configure VTDs based on the following formula: Max VTDs = ( 512 - 2 ) / ( virtual_q_depth + 3 )</li> <li>Only enable the largesend attribute on the SEA (physical adapter backing the SEA) if all LPARs serviced by the VIOS are AIX partitions.</li> </ul>
Virtual Ethernet adapters on AIX	<ul style="list-style-type: none"> <li>Increase the virtual Ethernet (vETH) device driver buffers if the partition is dropping packets on the virtual interface even when running with entitled CPU capacity. e.g., chdev -l ent# -a max_buf_xxx=NNNN</li> <li><b>NOTE:</b> For desired buffer size adjustments, refer to "AIX on Power – Performance FAQ" link below</li> <li>Set largesend on vETH adapter to improve performance (AIX): chdev -l en# -a mtu_bypass=on (or) ifconfig en# largesend</li> </ul>

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

### **POWER**

- [Power Virtualization Best Practices](#)
- [IBM Power Systems Performance Report \(Enhanced rPerf\)](#)
- [Power9 Migration hints and tips](#)

### **IBM i**

- [IBM i on Power – Performance FAQ](#)

### **AIX and VIOS**

- [AIX on Power – Performance FAQ](#)
- [IBM AIX MPIO : Best practices and considerations](#)

### **Java / WebSphere**

- [Best Practices for Java and IBM WebSphere Application Server \(WAS\) on IBM Power9](#)

### **Databases**

- [IBM Power9 and SMT performance for Db2](#)
- [Oracle Database 11g and 12c on IBM Power Systems S924, S922 and S914 with Power9 Processors](#)
- [AIX and Oracle Database Performance Considerations \(ICC\)](#)

### **Advisor Tools:**

- [VIOS Advisor](#)

### **Redbooks:**

- [PowerVM Best Practices](#)
- [PowerVM Managing and Monitoring](#)
- [PowerVM Virtualization Introduction and Configuration](#)
- [POWER Optimization and Tuning Guide](#)

### **Software Updates:**

- [IBM i Technology Updates](#)
- [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 [abraham1@us.ibm.com](mailto:abraham1@us.ibm.com)