

# POWER8 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 (link below) provides latest updates, which address optimal placement for partitions. Latest F/W levels as of this writing : FW860 for all POWER8 models 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	For optimal performance on workloads that are memory bandwidth sensitive follow these recommendations: <b>S814/S824/S822/S822L:</b> <ul style="list-style-type: none"> <li>Assign minimum 4 DIMMs per socket</li> <li>2 DIMMs behind each P8 chip with all same size DIMMs</li> </ul> <b>E870/E880:</b> <ul style="list-style-type: none"> <li>Assign a minimum of 4 DIMMs of same size per socket</li> </ul> <b>All POWER8 :</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.
40GbE adapter	<ul style="list-style-type: none"> <li>It is recommended that this adapter (FC EC3A &amp; EC3B) be installed on the internal CEC slots.</li> <li>RHEL7 : For network bandwidth sensitive workloads, we recommend increase the receive queue size from 1024 to 8192</li> </ul>
Java	<ul style="list-style-type: none"> <li>IBM Java 7.1 SR1 is the minimum level to exploit POWER8; recommend to be at least at IBM Java 8 SR1 FP10 or later.</li> <li>JDK 7.0 and below defaults to ppc-common instructions set</li> <li>Use of 64k size pages increases application performance</li> <li>If the application uses encryption, try H/W acceleration for encryption by setting <code>Dcom.ibm.crypto.provider.doAESInHardware=true</code></li> </ul>
WAS	<ul style="list-style-type: none"> <li>WAS 8.5.5.2 bundled with JDK 7.1 SR1 recommended level for P8</li> <li>Minimum supported level are 8.0.0.7 or 7.0.0.31</li> </ul>
Partition Placement	<ul style="list-style-type: none"> <li>Use DPO to optimize placement. Recommend to run DPO after DLPAR operations (memory or CPU) or LPM operations.</li> <li>Ensure that LPAR affinity score is above 90, for best performance.</li> </ul> <b>NOTE:</b> Additional memory or CPU resources may be required to achieve a score of 90 or above.
Sizing a system	<ul style="list-style-type: none"> <li>Use the Enhanced rPerf values in the System Performance Reports to correctly size P8 systems. Scale down both entitled capacity and virtual CPUs of a partition in accordance with relative rPerf value, for similar performance.</li> <li>Use Workload Estimator (WLE) rather than benchmark metrics 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 to sustained peak utilization for LPARs with critical SLA requirements</li> <li>Assign entitled capacity 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>

Description	Instructions
AIX Tunables/ VIOS Tunables	<ul style="list-style-type: none"> <li>Tunables <b>should not</b> be migrated across AIX levels</li> <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>All tunables (VMM, network, etc.) are configured for best performance. For desired adjustments, refer to "AIX on Power – Performance FAQ" link below</li> </ul>
AIX CPU utilization	The system is optimized for best raw throughput at higher CPU usage. If the customer requires to reduce CPU usage, use the <b>schedo</b> tunable <code>vpm_throughput_mode</code> to tune the workload and evaluate the benefits of raw throughput vs. CPU usage.
IBM i	<ul style="list-style-type: none"> <li>Ensure Technology Updates are current (link below)</li> </ul>
VIOS configuration	<ul style="list-style-type: none"> <li>Assign total entitlement of all VIOS partitions to be 10-15% of cores in shared pool. Refer to the PowerVM Best Practices for additional recommendations</li> <li>If configured as a shared partition, assign a CPU ratio of 2:1 (vCPU:ec) to the VIOS</li> <li>Assign uncapped mode and set variable weight capacity of VIOS partition higher than all client LPARs serviced by VIOS</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: <math>\text{Max VTDs} = (512 - 2) / (\text{virtual\_q\_depth} + 3)</math></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>
LPM	<ul style="list-style-type: none"> <li>Migration of LPARs (in VIO only mode) to P8 system (P6/P7 compat mode):               <ul style="list-style-type: none"> <li>Manual Migration (via backup/restore) should work as is</li> <li>LPM operations using FC NPIV may need LPM ifix based on migrating LPAR's AIX levels. Please use the FLRT tool to see the requirements (OS, VIOS, HMC and F/W) for a given migration.</li> </ul> </li> </ul>
Virtual Ethernet adapters on AIX	<ul style="list-style-type: none"> <li>Increase the virtual Ethernet device driver buffers if the partition is dropping packets on the virtual interface even when running with entitled CPU capacity. e.g., <code>chdev -l ent# -a max_buf_xxx=NNNN</code></li> <li><b>NOTE:</b> For desired buffer size adjustments, refer to "AIX on Power – Performance FAQ" link below</li> <li>Set largesend on virtual Ethernet adapter to improve performance: <code>chdev -l ent# -a mtu_bypass=on (or) ifconfig ent# largesend</code></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](#)
- [VIOS Sizing](#)
- [AIX Network Tuning for 10GigE and Virtual Network](#)
- [IBM AIX MPIO : Best practices and considerations](#)

### **Java / Websphere**

- [Best Practices for Java and IBM WebSphere Application Server \(WAS\) on IBM POWER9](#)

### **Databases**

- [AIX and Oracle Database Performance Considerations \(ICC\)](#)

## **Advisor Tools:**

- [IBM Systems Workload Estimator](#)
- [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:

<http://www.ibm.com/systems/support/p/bestpractices.html>

For suggestions/changes to this document, please contact [abraham1@us.ibm.com](mailto:abraham1@us.ibm.com)