

Power10 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. Install the latest version of FW1010 for the Power10 E1080 1-node or 2-node servers. Use the FLRT tool to obtain the recommended levels for a given platform. NOTE: 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 recommended levels for each H/W model. Use these tools to maintain your system up to date. |
| NX GZIP | To take advantage of NX GZIP H/W accelerator on Power10 processor-based systems the LPAR must be in POWER9 compatibility mode (not POWER9_base mode), or POWER10 mode. |
| LPM | Use of 10Gb or faster dedicated network for LPM is preferred. |
| Sizing a system | When migrating to Power10, 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"> Assign entitled capacity (EC) to sustained peak utilization for LPARs with critical SLA requirements Assign EC to average utilization and number of virtual CPUs to peak utilization (physical core consumption) for LPARs with non-critical SLA Ensure the average LPAR utilization is equal or less than 75% of the entitled capacity |
| 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 | Take advantage of new Power10 ISA features for operating systems running in Power10 mode by using the following compiler releases or later: <ul style="list-style-type: none"> IBM Open XL C/C++ for AIX 17.1.0 IBM Open XL Fortran for AIX 17.1.0 IBM Advance Toolchain for Linux on Power 15.0-0 GCC 11.2 |
| Java | Java applications can seamlessly take advantage of new Power10 ISA features on operating systems running in Power10 mode by using the Java runtime versions listed below or newer: <ul style="list-style-type: none"> Java 8 <ul style="list-style-type: none"> IBM SDK 8 SR6 FP36 IBM Semeru Runtime Open Edition 8u302: openj9-0.27.1 Java 11 <ul style="list-style-type: none"> IBM Semeru Runtime Certified Edition 11.0.12.1: openj9-0.27.1 IBM Semeru Runtime Open Edition 11.0.12.1: openj9-0.27.1 Java 17 (drivers may not be available yet) <ul style="list-style-type: none"> IBM Semeru Runtime Certified Edition 17: openj9-0.28 IBM Semeru Runtime Open Edition 17: openj9-0.28 OpenJDK 17 |
| IBM i | Ensure Technology Updates are current (updates can be found here) |

| Description | Instructions |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AIX | Tunables <ul style="list-style-type: none"> Tuning a VIOS is not recommended unless directed by VIOS/AIX support. Restricted tunables should not be modified (unless directed by AIX/VIOS development) Tunables should be at default unless otherwise recommended by AIX support. If tunables have been modified from default, they should be re-evaluated prior to moving to a different processor platform. CPU usage <ul style="list-style-type: none"> On POWER10, 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 customer 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"> VIOS 3.1.3 runs on Power9 compatibility mode 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 If configured with shared processors: 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 Assign uncapped mode and set variable weight capacity of VIOS partition higher than all client LPARs serviced by VIOS 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 link. For vFC, ensure the number of client connections does not exceed the limits of the physical adapter. 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. For vSCSI adapters, ensure you configure VTDs based on the following formula: Max VTDs = (512 - 2) / (virtual_q_depth + 3) |
| Virtual Ethernet adapters on AIX | <ul style="list-style-type: none"> 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 NOTE: For desired buffer size adjustments, refer to "AIX on Power – Performance FAQ" link below |
| vNIC | <ul style="list-style-type: none"> AIX - Change the following vNIC interface settings as follows: chdev -l ent# -a rx_que_num=16 -a tx_que_num=16 -a rx_que_elem=2048 -a tx_que_elem=1024 -a use_rec_q_val=no Linux distros - Update to latest kernel. Also, set rx/tx queues to maximum ethtool -L eth# rx 16 tx 16 |
| 100GbE adapter | <ul style="list-style-type: none"> RHEL8.4: For network bandwidth sensitive workloads, we recommend increase the receive queue size from 1024 to 4096 and increase the number of combined queues to 64. AIX 7.2: For network bandwidth sensitive workloads, we recommend increase the receive queue size from 1024 to 4096, increase the number of receive queues to 20 and transmit queues to 12. |

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](#)
- [Power system to AIX Maps](#)

Java / WebSphere

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

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