


Virtualisation Best Practices – A Technical Point of View



Agenda

- Introduction
- Logical Partition Configuration and Considerations.
- Virtual I/O Server Sizing, Configuration and Considerations.
- MPIO Recommended Settings
- Firmware, software patching and maintenance strategy recommendations.
- Documentation - System Planning Tool, Estate-wide documentation.

Logical Partition Configuration and Considerations.

- When defining the values for processor units and virtual processors always ensure that for the minimum values set them to 0.1 and 1, for desired set the units to the estimated average work load and the virtual processors to the next whole number (e.g. 5.5 processor units and 6 virtual processors), for maximum values a good default is to set double the desired values.
- When defining the values for memory, firstly do not set the maximum memory value too high – this is to minimise the amount of hypervisor memory used to maintain the potential maximum memory values. Set the minimum value to around 1GB, you do not need to have larger values.
- Set ALL virtual and (if used) physical adapters to desired, do not set to required as this will prevent the use of Live Partition Mobility and stop DLPAR operations from working. The goal is to allow adapter changes to happen dynamically with no downtime required on the partition.
- Always use separate VSCSI adapters for the AIX root volume group and data volume groups, for IBM i keep the load source disk on a separate VSCSI adapter from other disks.

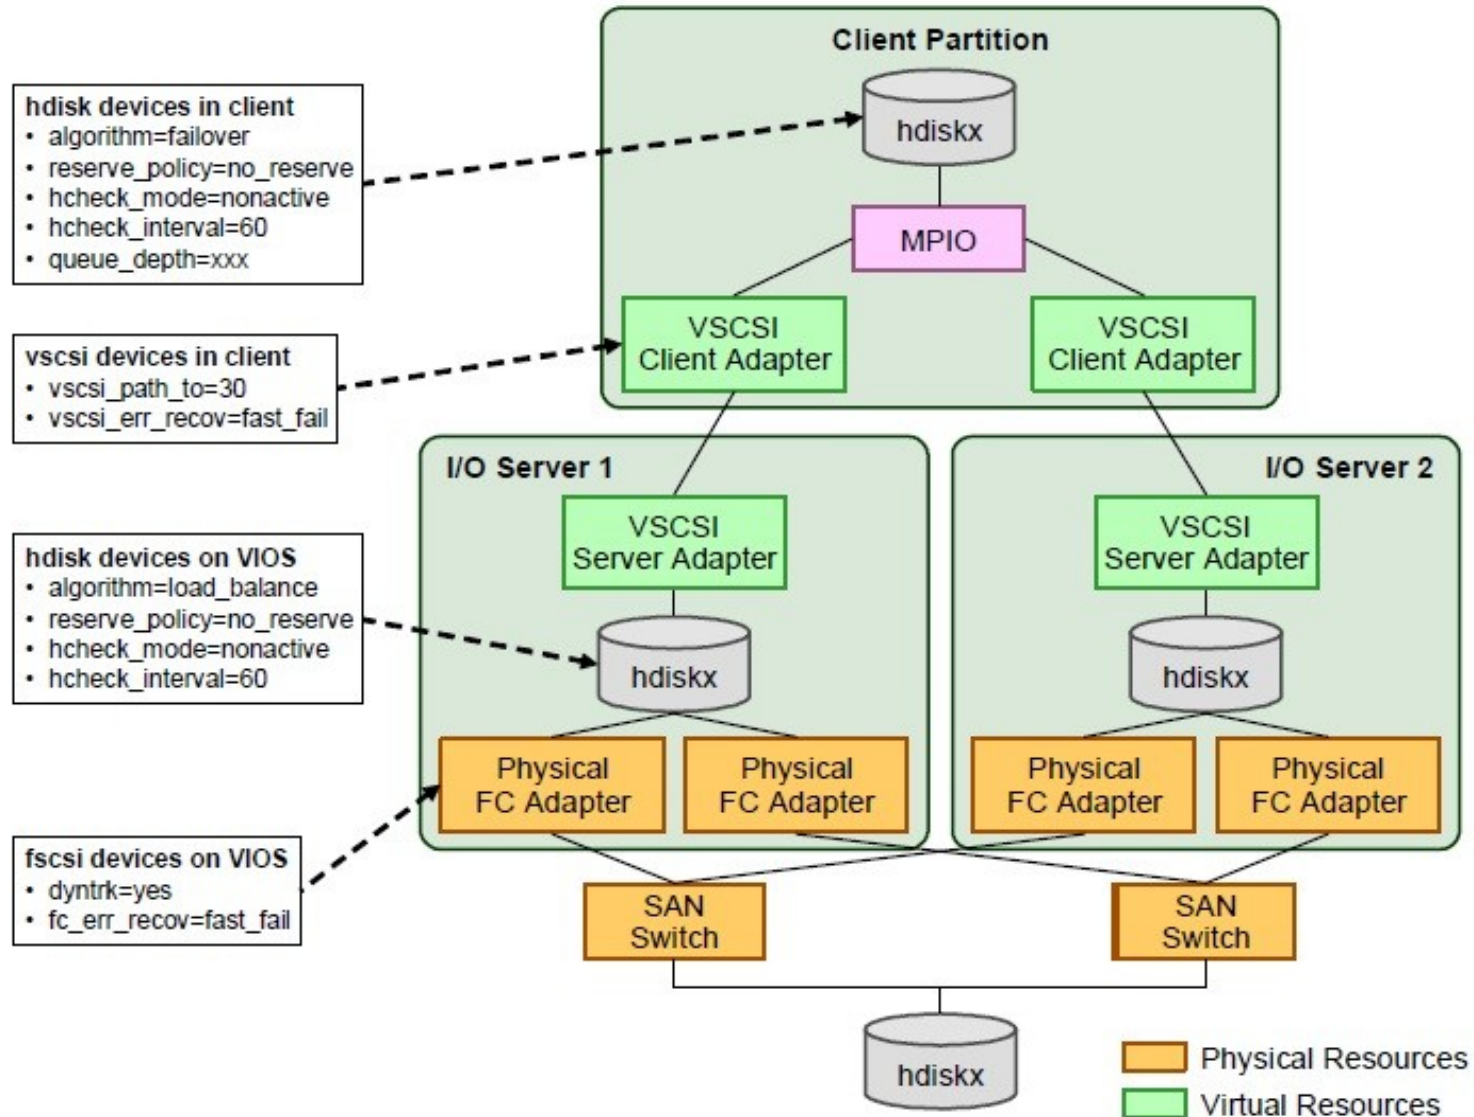
Virtual I/O Server Sizing, Configuration and Considerations.

- 1 core and 4GB memory per VIOS (minimum based on AIX clients)
- 1.5 cores and 4GB memory per VIOS (minimum based on IBM i clients)
- Dual VIOS servers recommended for redundancy
- Place into a CPU pool with maximum of 3/4/5 cores and un cap the partitions
- Above has been used for implementations with 20+ client partitions for disk and network virtualisation
- For a larger implementation, some customers choose to have separate pair of VIOS servers for production and DR workloads or for disk and network virtualisation
- **When using 10Gb Ethernet or 8Gb Fibre adapters the amount of memory required by the VIO Server increases dramatically. If using any of these high speed adapters then increase the desired memory from 4GB to 6GB or preferably 8GB.**
- For Blade Servers using 0.5 CPU and at least 2GB RAM as a starting point, for larger PS702/3/4 Blades use the normal guidelines above.

Virtual I/O Server Sizing, Configuration and Considerations (cont.).

- Make a virtual slot numbering scheme and stick to it! Pick a number range for virtual scsi adapters and another one for virtual networking. In a dual Virtual I/O Server configuration you could use even numbering for one and odd numbering for another as one common example.
- When using VSCSI attached disks make sure you name the disk mapping (the virtual target device name) using a naming scheme that reflects the target client partition and, if possible, the disk use. For example a root disk for an AIX client partition could be called *lpar1_rootvg* and an AIX data disk could be called *lpar1_datavg1*.
- Dont forget to record ALL disks, names, slot numbers and target client information in a document, database or spreadsheet for future reference. Do this while building the partitions and not afterwards.
- Where practical always use a redundant dual Virtual I/O Server configuration, this allows for patching and maintenance to be performed with one partition up and running live (supporting the client partitions) and the other down for patching.

MPIO Recommended Settings



Firmware, software patching and maintenance strategy recommendations.

- Set a date, every six months, to review your current firmware and software patch levels and compare them to the latest stable versions.
- Ensure a regular maintenance window is available to perform firmware updates and patching, once a year is the recommended time frame to perform the updates.
- When doing system firmware updates from one major release to another ALWAYS update the HMC to the latest available version first along with any mandatory HMC patches then doing the firmware. If the operating system is being updated as well, update the operating system first, then HMC code and lastly the system firmware.
- In a dual HMC configuration always update both HMC's in a single maintenance slot or disconnect one HMC until it is updated to the same level as the other HMC.
- Verify the recommended code levels using the Fix Level Recommendation Tool (FLRT) on the public IBM Support Site.
- The IBM recommended strategy documentation is available from here:
<http://www14.software.ibm.com/webapp/set2/sas/f/best/home.html>

Documentation - System Planning Tool, Estate-wide documentation.

- Documentation of the system configuration, naming conventions, slot allocations, firmware and software levels and maintaining that document is extremely important.
- Create a System Plan of any new system before any configuration work is performed and then create another System Plan after all partition configuration work has been completed.
- Use the first System Plan from the HMC to load into the System Planning Tool and perform configuration work in that tool first, this will allow you to set up different scenarios and review slot allocations and adapter placements before committing anything back to the hardware.
- Use the second System Plan in case of an unlikely DR event and the configuration needs to be re-created. Also, make sure partition profile data is saved on a regular basis in the HMC so profile information can be restored very quickly.
- Having everything documented will allow you to provide IBM Support with vital information very quickly, along with other data they request (snap files etc.), and reduce the time it takes to move a support request from information gathering onto the resolution stage.

Copyright and Notices

© IBM Corporation 2011
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, AIX, AIX 5L, BladeCenter, DB2, eServer, Lotus, Notes, Power, POWER, POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER7, POWER PC, Power Architecture, Power Systems, PowerPC 601, PowerPC 604, pSeries, RS/6000, SP, System I, System p, System p5 and WebSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of US trademarks owned by IBM may be found at <http://www.ibm.com/legal/copytrade.shtml>.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

TPC-C is a trademark of the Transaction Processing Council (TPPC).

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPECjpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

NetBench is a registered of Ziff Davis Media in the United States, other countries or both.

AltiVec is a trademark of Freestyle Semiconductor, Inc.

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

The IBM home page on the Internet can be found at: <http://www.ibm.com>.

The Power Systems home page on the Internet can be found at: <http://www.ibm.com/systems/power>.

The BladeCenter home page on the Internet can be found at: <http://www.ibm.com/ystems/bladecenter>.