



Low Latency/Jitter Optimization

User's Guide

Revision 2.1

The information in this USER'S MANUAL has been carefully reviewed and is believed to be accurate. The vendor assumes no responsibility for any inaccuracies that may be contained in this document, makes no commitment to update or to keep current the information in this manual, or to notify any person organization of the updates. Please Note: For the most up-to-date version of this manual, please see our web site at www.supermicro.com.

Super Micro Computer, Inc. ("Supermicro") reserves the right to make changes to the product described in this manual at any time and without notice. This product, including software, if any, and documentation may not, in whole or in part, be copied, photocopied, reproduced, translated or reduced to any medium or machine without prior written consent.

IN NO EVENT WILL SUPERMICRO BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, SPECULATIVE OR CONSEQUENTIAL DAMAGES ARISING FROM THE USE OR INABILITY TO USE THIS PRODUCT OR DOCUMENTATION, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN PARTICULAR, SUPERMICRO SHALL NOT HAVE LIABILITY FOR ANY HARDWARE, SOFTWARE, OR DATA STORED OR USED WITH THE PRODUCT, INCLUDING THE COSTS OF REPAIRING, REPLACING, INTEGRATING, INSTALLING OR RECOVERING SUCH HARDWARE, SOFTWARE, OR DATA.

Any disputes arising between manufacturer and customer shall be governed by the laws of Santa Clara County in the State of California, USA. The State of California, County of Santa Clara shall be the exclusive venue for the resolution of any such disputes. Super Micro's total liability for all claims will not exceed the price paid for the hardware product.

FCC Statement: This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the manufacturer's instruction manual, may cause harmful interference with radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case you will be required to correct the interference at your own expense.

California Best Management Practices Regulations for Perchlorate Materials: This Perchlorate warning applies only to products containing CR (Manganese Dioxide) Lithium coin cells. Perchlorate Material-special handling may apply. See <http://www.dtsc.ca.gov/hazardouswaste/perchlorate/> for further details.

Manual Revision 2.1

Release Date: April 1, 2016

Unless you request and receive written permission from Super Micro Computer, Inc., you may not copy any part of this document.

Information in this document is subject to change without notice. Other products and companies referred to herein are trademarks or registered trademarks of their respective companies or mark holders.

Copyright © 2016 by Super Micro Computer, Inc.

All rights reserved.

Printed in the United States of America

Introduction

These instructions are optimized for the SYS-1028UX-LL1/2/3-B8, SYS-1028-UX-CR-LL1/2 and SYS-6027AX-72RF/TRF- HFT1/2/3 systems. Steps may vary depending on SW application requirements and operating system.

1.1 OS Optimizations

1. Minimal OS installation of RHEL.
2. Kernel boot command line options: `intel_idle.max_cstate=0 processor.max_cstate=0 idle=poll pcie_aspm=performance mce=ignore_ce ipmi_si.force_kipmi=0 nmi_watchdog=0 hpet=disabled noht nohz=on nohalt nosoftlockup isolcpus=x,x (isolcpus option not required if using tuna to isolate CPU cores)`
3. Disable all unnecessary services using `CHKCONFIG` command.
4. Install `cpupowerutils`, `tuned-adm`, `tuna`, and `numactl`.
5. Run the following commands to tune your system before running your application. It is best to script the following commands for ease of execution.
 - a. `Modprobe acpi-cpufreq`
 - b. `cpupower set -b 0`
 - c. `cpupower frequency-set --governor performance`
 - d. `tuned-adm profile latency-performance`
 - e. `tuna -S0 -i; tuna -S1 -i`
 - f. `service irqbalance stop`

1.2 BIOS/IPMI Optimizations

1. Set “SuperMicro Hyper-Speed” at the highest, appropriate level
3. Set “SuperMicro Hyper-Turbo” as “Enabled”
4. Set IPMI fan speed to FULL.
5. Disabling unused devices such as USB controllers and SCU controller (PCH chipset’s storage controller) can help reduce system interrupts.

1.3 NIC Optimizations

1. To locate which NUMA node the NIC is directly connected to:
`#cat/sys/class/net/<nic port>/device/numa_node`
2. To ascertain which cores are located on the local NUMA node:
`#numactl --hardware`
3. Finally, use `taskset` or `numactl` to bind application to specific CPU cores

Mellanox NIC Acceleration:

Prefix application with: `LD_PRELOAD=libvma.so`

SolarFlare NIC Acceleration:

Configure the firmware:

- i. Download SolarFlare Utilities
- ii. `#sfupdate -write`
- iii. `#sfboot firmware-variant=ultra-low-latency`

Prefix application with: `onload --profile=latency`

Exablaze NIC Acceleration:

Prefix application with: `exasock`

1.5 Reading the CPU Frequency in Linux

To check the CPU frequency consistency in a Linux environment, use the “`turbostat -i 1`” command.