

Network performance degradation from 2.6.11.12 to 2.6.16.20

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-06/msg04579.html>

- *From:* Harry Edmon <harry@xxxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Fri, 16 Jun 2006 09:01:23 -0700
-

I have a system with a strange network performance degradation from 2.6.11.12 to most recent kernels including 2.6.16.20 and 2.6.17-rc6. The system is has Dual single core Xeons with hyperthreading on. The application is the LDM system from UCAR/Unidata (<http://www.unidata.ucar.edu/software/lDM>). This system requests weather data from a variety of systems using RPC calls over a reserved TCP port (388), puts them into a memory mapped queue file, and then sends the data out to a variety of downstream requesting systems, again using RPC calls. When the load is heavy, the 2.6.16.20 kernel falls way behind with the data ingestion. The 2.6.11.12 kernel does not. I have tried an experiment with a 2.6.17-rc6 system where it just does the ingestion, and not the downstream distribution, and it is able to keep up. I would really appreciate any pointers as to where the problem may be and how to diagnose it. I have attached the config files from both kernels and the sysctl.conf file I am using. I have also included the output from "netstat -s" on the 2.6.16.20 system during a time when it was having problems.

```
#
# Automatically generated make config: don't edit
# Linux kernel version: 2.6.11.12
# Wed Oct 19 11:16:32 2005
#
CONFIG_X86=y
CONFIG_MMU=y
CONFIG_UID16=y
CONFIG_GENERIC_ISA_DMA=y
CONFIG_GENERIC_IOMAP=y

#
# Code maturity level options
#
CONFIG_EXPERIMENTAL=y
CONFIG_CLEAN_COMPILE=y
CONFIG_LOCK_KERNEL=y

#
# General setup
#
CONFIG_LOCALVERSION=""
CONFIG_SWAP=y
```

Network performance degradation from 2.6.11.12 to 2.6.16.20

```
CONFIG_SYSVIPC=y
CONFIG_POSIX_QUEUE=y
CONFIG_BSD_PROCESS_ACCT=y
# CONFIG_BSD_PROCESS_ACCT_V3 is not set
CONFIG_SYSCTL=y
# CONFIG_AUDIT is not set
CONFIG_LOG_BUF_SHIFT=14
CONFIG_HOTPLUG=y
CONFIG_KOBJECT_UEVENT=y
# CONFIG_IKCONFIG is not set
# CONFIG_EMBEDDED is not set
CONFIG_KALLSYMS=y
# CONFIG_KALLSYMS_ALL is not set
# CONFIG_KALLSYMS_EXTRA_PASS is not set
CONFIG_FUTEX=y
CONFIG_EPOLL=y
# CONFIG_CC_OPTIMIZE_FOR_SIZE is not set
CONFIG_SHMEM=y
CONFIG_CC_ALIGN_FUNCTIONS=0
CONFIG_CC_ALIGN_LABELS=0
CONFIG_CC_ALIGN_LOOPS=0
CONFIG_CC_ALIGN_JUMPS=0
# CONFIG_TINY_SHMEM is not set

#
# Loadable module support
#
CONFIG_MODULES=y
CONFIG_MODULE_UNLOAD=y
CONFIG_MODULE_FORCE_UNLOAD=y
CONFIG_OBSOLETE_MODPARM=y
CONFIG_MODVERSIONS=y
# CONFIG_MODULE_SRCVERSION_ALL is not set
CONFIG_KMOD=y
CONFIG_STOP_MACHINE=y

#
# Processor type and features
#
CONFIG_X86_PC# CONFIG_DM_MULTIPATH is not set

#
# Fusion MPT device support
#
# CONFIG_FUSION is not set
# CONFIG_FUSION_SPI is not set
# CONFIG_FUSION_FC is not set
# CONFIG_FUSION_SAS is not set

#
# IEEE 1394 (FireWire) support
```

Network performance degradation from 2.6.11.12 to 2.6.16.20

```
#
# CONFIG_IEEE1394 is not set

#
# I2O device support
#
CONFIG_I2O=m
CONFIG_I2O_LCT_NOTIFY_ON_CHANGES=y
CONFIG_I2O_EXT_ADAPTEC=y
CONFIG_I2O_EXT_ADAPTEC_DMA64=y
CONFIG_I2O_CONFIG=m
CONFIG_I2O_CONFIG_OLD_IOCTL=y
# CONFIG_I2O_BUS is not set
CONFIG_I2O_BLOCK=m
CONFIG_I2O_SCSI=m
CONFIG_I2O_PROC=m

#
# Network device support
#
CONFIG_NETDEVICES=y
# CONFIG_DUMMY is not set
# CONFIG_BONDING is not set
# CONFIG_EQUALIZER is not set
# CONFIG_TUN is not set

#
# ARCnet devices
#
# CONFIG_ARCNET is not set

#
# PHY device support
#
# CONFIG_PHYLIB is not set

#
# Ethernet (10 or 100Mbit)
#
CONFIG_NET_ETHERNET=y
CONFIG_MII=m
CONFIG_HAPPYMEAL=m
CONFIG_SUNGEM=m
# CONFIG_CASSINI is not set
CONFIG_NET_VENDOR_3COM=y
CONFIG_VORTEX=m
CONFIG_TYPHOON=m

#
# Tulip family network device support
#
```

Network performance degradation from 2.6.11.12 to 2.6.16.20

```
CONFIG_NET_TULIP=y
CONFIG_DE2104X=m
CONFIG_TULIP=m
# CONFIG_TULIP_MWI is not set
# CONFIG_TULIP_MMIO is not set
# CONFIG_TULIP_NAPI is not set
CONFIG_DE4X5=m
CONFIG_WINBOND_840=m
CONFIG_DM9102=m
# CONFIG_ULI526X is not set
CONFIG_HP100=m
CONFIG_NET_PCI=y
CONFIG_PCNET32=m
CONFIG_AMD8111_ETH=m
# CONFIG_AMD8111E_NAPI is not set
CONFIG_ADAPTEC_STARFIRE=m
# CONFIG_ADAPTEC_STARFIRE_NAPI is not set
CONFIG_B44=m
CONFIG_FORCEDETH=m
# CONFIG_DGRS is not set
CONFIG_EEPRO100=m
CONFIG_E100=m
CONFIG_FEALNX=m
CONFIG_NATSEMI=m
CONFIG_NE2K_PCI=m
CONFIG_8139CP=m
CONFIG_8139TOO=m
CONFIG_8139TOO_PIO=y
CONFIG_8139TOO_TUNE_TWISTER=y
CONFIG_8139TOO_8129=y
# CONFIG_8139_OLD_RX_RESET is not set
CONFIG_SIS900=m
CONFIG_EPIC100=m
CONFIG_SUNDANCE=m
# CONFIG_SUNDANCE_MMIO is not set
CONFIG_TLAN=m
CONFIG_VIA_RHINE=m
# CONFIG_VIA_RHINE_MMIO is not set

#
# Ethernet (1000 Mbit)
#
# CONFIG_ACENIC is not set
CONFIG_DL2K=m
CONFIG_E1000=m
CONFIG_E1000_NAPI=y
# CONFIG_E1000_DISABLE_PACKET_SPLIT is not set
CONFIG_NS83820=m
CONFIG_HAMACHI=m
CONFIG_YELLOWFIN=m
CONFIG_R8169=m
```

Network performance degradation from 2.6.11.12 to 2.6.16.20

```
# CONFIG_R8169_NAPI is not set
# CONFIG_SIS190 is not set
# CONFIG_SKGE is not set
# CONFIG_SKY2 is not set
CONFIG_SK98LIN=m
CONFIG_VIA_VELOCITY=m
CONFIG_TIGON3=m
# CONFIG_BNX2 is not set

#
# Ethernet (10000 Mbit)
#
# CONFIG_CHELSIO_T1 is not set
CONFIG_IXGB=m
# CONFIG_IXGB_NAPI is not set
CONFIG_S2IO=m
# CONFIG_S2IO_NAPI is not set

#
# Token Ring devices
#
# CONFIG_TR is not set

#
# Wireless LAN (non-hamradio)
#
# CONFIG_NET_RADIO is not set

#
# Wan interfaces
#
# CONFIG_WAN is not set
# CONFIG_FDDI is not set
# CONFIG_HIPPI is not set
# CONFIG_PPP is not set
# CONFIG_SLIP is not set
# CONFIG_NET_FC is not set
CONFIG_SHAPER=m
CONFIG_NETCONSOLE=m
CONFIG_NETPOLL=y
# CONFIG_NETPOLL_RX is not set
# CONFIG_NETPOLL_TRAP is not set
CONFIG_NET_POLL_CONTROLLER=y

#
# ISDN subsystem
#
# CONFIG_ISDN is not set

#
# Telephony Support
```

Network performance degradation from 2.6.11.12 to 2.6.16.20

```
#
# CONFIG_PHONE is not set

#
# Input device support
#
CONFIG_INPUT=y

#
# Userland interfaces
#
CONFIG_INPUT_MOUSEDEV=y
CONFIG_INPUT_MOUSEDEV_PSAUX=y
CONFIG_INPUT_MOUSEDEV_SCREEN_X=1024
CONFIG_INPUT_MOUSEDEV_SCREEN_Y=768
CONFIG_INPUT_JOYDEV=m
CONFIG_INPUT_TSDEV=m
CONFIG_INPUT_TSDEV_SCREEN_X=240
CONFIG_INPUT_TSDEV_SCREEN_Y=320
CONFIG_INPUT_EVDEV=m
CONFIG_INPUT_EVBUG=m

#
# Input Device Drivers
#
CONFIG_INPUT_KEYBOARD=y
CONFIG_KEYBOARD
```