



Product Bulletin 50 (Revised)

Product	Any version of StorNext® that is running on Linux kernel 2.6.10 or higher
Summary	Performance for Linux systems on kernel 2.6.10 and higher can be optimized by changing the default I/O scheduler
Date	November 2011 (First revision published March 2010)

Overview

The I/O throughput of Linux Kernel 2.6.10 (SLES10 and later and RHEL5 and later) can be increased by adjusting the default I/O settings. The tuning procedures for your Linux kernel are available from your Linux distribution vendor, and are also summarized in this document.

For more information, contact the Quantum Technical Assistance Center and reference CR 29432.

Note: Compared to the first version of this bulletin published in March 2010, this revision updates the information in the Symptoms and Solution sections.

Symptoms

Using the default scheduler in the Linux Kernel 2.6.10 will result in a significant reduction in I/O performance when compared to older kernel releases. The default RHEL5 scheduler I/O read performance can be degraded by 50%, compared to that of a RHEL4 distribution.

Solution

Beginning with the 2.6 kernel, the Linux I/O scheduler can be changed to control how the kernel does reads and writes. There are four types of I/O scheduler available in Linux kernel 2.6.10 and higher:

- The completely fair queuing scheduler
- The no operation scheduler
- The deadline scheduler
- The anticipatory scheduler

The default scheduler in most distributions is the completely fair queuing (cfq). Experimentation shows that the deadline scheduler provides the best improvement. The default size of the Linux I/O scheduler's queue depth should also be changed from its default value of 128 to 4096. Refer to your Linux tuning guide for additional information on I/O scheduler tuning.

To optimize the I/O scheduler, create a file named `/usr/cvfs/config/deviceparams` and add the following lines:

```
scheduler=deadline  
nr_requests=4096
```

(The `nr_requests` value is the size of the Linux I/O scheduler queue depth.)

After adding these lines, restart the StorNext service.

You can verify the I/O scheduler changes were implemented by examining the `/var/log` messages.

```
Nov 16 10:06:44 rhe15sp6-1 fsmpm[1543]: deviceparams: scheduler=deadline  
Nov 16 10:06:44 rhe15sp6-1 fsmpm[1543]: deviceparams: nr_requests=4096  
Nov 16 10:06:44 rhe15sp6-1 fsmpm[1543]: deviceparams: set /sys/block/sdb/  
queue/scheduler to deadline  
Nov 16 10:06:44 rhe15sp6-1 fsmpm[1543]: deviceparams: set /sys/block/sdb/  
queue/nr_requests to 4096
```

(repeats for all block devices used by StorNext)

And by checking the settings for one of the above devices.

```
[root@rhe15sp6-1 queue]# pwd  
/sys/block/sdb/queue  
[root@rhe15sp6-1 queue]# cat scheduler  
noop anticipatory [deadline] cfq  
[root@rhe15sp6-1 queue]# cat nr_requests  
4096
```

Contacting Quantum

More information about StorNext is available on the Quantum Service and Support website at www.quantum.com/ServiceandSupport. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs). You can also access software, firmware, and drivers through this site.

For further assistance, contact the Quantum Technical Assistance Center:

North America	+1 800-284-5101 Option 5
EMEA	00800 9999 3822
Online Service and Support	www.quantum.com/OSR
World Wide Web	www.quantum.com/ServiceandSupport

(Local numbers for specific countries are listed on the Quantum Service and Support Website.)