

## Disk I/O Performance

Beginning with AirWave 7.0, support for polling of switches was expanded and we overhauled the way user data is stored. These changes have had a significant impact on disk I/O on some servers.

In AirWave 7.1 we made several improvements to increase the efficiency of disk I/O. Outside of that, there are several areas under your control where disk I/O can be improved.

### Checking I/O Performance in AirWave

On the System → Performance page there are three graphs that relate to disk performance:

- System Disk IOPs
- System Disk Outstanding I/O Requests
- System Disk Utilization

If disk utilization is frequently approaching 100% there are some things you can do to reduce the load on the hard drive system.

### Disk Drive Spindle Speed

The AirWave server is I/O intensive. For this reason we always require drives with 15,000 RPM spindle speeds. And on systems monitoring more than 500 devices we require multiple drives configured in a RAID-10 array. Please refer to the *AWMS Hardware Sizing Guide* on the Support Downloads page for details:

<http://www.airwave.com/download>

To see how to check the spindle speed of your server's drives follow the link below to a Knowledge Base article on the topic:

<http://kb.airwave.com/?sid=5014000000aF3U>

To verify the configuration of your hard drives, run the following command from the AirWave command line:

```
# df -h
```

If the total size of all partitions does not equal the combined capacity of the drives on the system then your drives may not be configured correctly.

### Router / Switch and Interface Polling

If the server is monitoring a large number of switches, consider backing off or disabling related polling to improve I/O performance. You can find these settings in the **SNMP Polling Periods** and **Routers and Switches** sections on the Groups → Basic page for group(s) containing routers and switches. The default periods are as follows:

(See screen shot on the following page.)

Home	Helpdesk	<b>Groups</b>	APs/Devices	Users	Reports	System	
List	Monitor	<b>Basic</b>	Templates	Security	SSIDs	AAA Servers	Radio
CDP Neighbor Data Polling Period:			30 minutes				
CDP Neighbor Data Polling Period:			30 minutes				
<b>Routers and Switches</b>							
Read ARP Table:			4 hours				
Read CDP Table for Device Discovery:			4 hours				
Read Bridge Forwarding Table:			4 hours				
Interface Up/Down Polling Period:			10 minutes				
Interface Bandwidth Polling Period:			15 minutes				
Interface Error Counter Polling Period:			30 minutes				
Poll 802.3 error counters:			<input type="radio"/> Yes <input checked="" type="radio"/> No				
Poll Cisco interface error counters:			<input type="radio"/> Yes <input checked="" type="radio"/> No				

## CDP-related Polling

If your network contains no Cisco IOS APs, or if it does and you've completed discovery of these devices, you should consider disabling CDP-related polling. On the Groups → Basic page under the **SNMP Polling Periods** section, disable the *CDP Neighbor Data Polling Period*, and under the **Routers and Switches** section, disable *Read CDP Table for Device Discovery*.

## I/O Scheduling

Another way you might be able to improve disk I/O performance is to change the I/O scheduling algorithm that's being used. Both the OS and the disk controller are capable of I/O scheduling. But when multiple disks are combined logically (via striping or a RAID setup) the OS cannot distinguish one physical disk in the stripe from another, it sees only one disk. The controller on the other hand is still aware of the individual disks and can potentially organize I/O more efficiently.

By default the OS is set up to use the Completely Fair Queuing or **CFQ** algorithm. By selecting the **NOOP** algorithm --which simply inserts all I/O requests into a FIFO queue, essentially disabling scheduling by the OS-- scheduling is deferred to the disk controller.

(For more information on the various scheduling algorithms refer to the links at the bottom of this article.)

## Changing the I/O Scheduling Algorithm

First check to see which I/O scheduling algorithm is currently selected:

```
# cat /sys/block/<BLOCK_DEVICE>/queue/scheduler
```

Ex.

```
# cat /sys/block/sda/queue/scheduler
```

(For more informaton on BLOCK\_DEVICE designations refer to the end of this article.)

The output will look something like this:

```
noop anticipatory deadline [cfq]
```

The algorithm in brackets [] is the one currently selected.

Change the algorithm to **NOOP**:

```
# echo noop > /sys/block/<BLOCK_DEVICE>/queue/scheduler
```

Ex.

```
# echo noop > /sys/block/sda/queue/scheduler
```

```
[noop] anticipatory deadline cfq
```

That's it.

---

## Additional Information

For more information on I/O scheduling:

[http://en.wikipedia.org/wiki/I/O\\_scheduling](http://en.wikipedia.org/wiki/I/O_scheduling)

<http://en.wikipedia.org/wiki/CFQ>

[http://en.wikipedia.org/wiki/Category:Disk\\_scheduling\\_algorithms](http://en.wikipedia.org/wiki/Category:Disk_scheduling_algorithms)

The following describes the BLOCK\_DEVICE designations:

**sd***xn*

**hd***xn*

Ex: sda1

**sd** : SCSI, SAS, SATA

**hd** : IDE

**x**- drive designation

sda : primary master

sdb : primary slave

sdc : secondary master

sdd : secondary slave

**n**- partition number

Ex: sda1 - first partition of the primary master

Primary partitions are generally assigned numbers 1-4.

Logical drives are generally assigned numbers 5-n.