
Swappiness and Limits for SystemD Oses

Swappiness

Synchronization between memory cache and hard disc is accomplished at the Operating System level using two methods. First, the synchronization option for the file system is set in `/etc/fstab`.

```
sudo vi /etc/fstab
UUID=foo /var ext2 noatime,async,errors=continue 1 1
```

The option `sync` means that all changes to the according filesystem are immediately flushed to disk; the respective write operations are being waited for. For mechanical drives that means a huge slow down since the system has to move the disk heads to the right position; with `sync` the process has to wait for the operation to complete. In contrast, with `async` the system buffers the write operation and optimizes the actual writes; meanwhile, instead of being blocked the process continues to run. The default is `async`.

When `sync` mode is not enabled, the complicated writeback algorithm comes into play. The writeback algorithm is designed to limit IO operations. It assumes that the user prefers the system perform flushes to disk only occasionally: either after a deadline or when there are a threshold of "dirty" pages to flush. The second and recommended method is to tune the OS memory cache to disk write settings for best performance (Linux Oses only).

```
sudo vi /etc/sysctl.conf
vm.dirty_background_ratio = 50
vm.dirty_ratio = 90
vm.swappiness = 0
vm.dirty_writeback_centisecs = 500
vm.dirty_expire_centisecs = 60000
```

For more information on Disk caching and Performance tuning, please see the following:

https://lonesysadmin.net/2013/12/22/better-linux-disk-caching-performance-vm-dirty_ratio/
<https://www.kernel.org/doc/Documentation/sysctl/vm.txt>

Limits

Warning: RedHat/CentOS 7+, Ubuntu 16+, Debian 8+, SLES 12+ ignore `/etc/security/limits.conf` and `/etc/security/limits.d/*` settings.

Core File Limits

To permit core files without size limits on systemd-based systems use the following process. Note that these steps must be done as root:

```
sudo -s

cd /etc/systemd/system

mkdir solserver.service.d

cd solserver.service.d

printf "[Service]\nLimitCORE=infinity\n" > override.conf
```



Symas OpenLDAP

How-To Guides

```
systemctl daemon-reload
```

Now you can verify the change has taken effect for future slapd startup:

```
systemctl cat solserver
```

This should show the updated limits:

```
# /etc/systemd/system/solserver.service.d/override.conf
[Service]
LimitCORE=infinity
```

Once Symas OpenLDAP is installed you can view the process's limits file to see the Max open files limit:

```
systemctl show solserver | grep LimitCORE
```

```
LimitCORE=infinity
LimitCORESoft=infinity
```

If adjusting these settings after Symas OpenLDAP is installed and slapd is running, a restart of solserver is required to pick up the change.

```
systemctl restart solserver
```

NoFile Limits

The default number of files (nofile) allowed to be opened by any specific process can vary depending on the operating system. Defaults are typically 1024 - 4096. For large customers with high traffic volume on OpenLDAP servers these limits can cause slapd to return PANIC errors rather than completing the requested operations.

To increase nofile limits on systemd-based systems use the following process. Note that these steps must be done as root:

```
sudo -s

cd /etc/systemd/system/solserver.service.d

printf "LimitNOFILE=524288\n" >> override.conf

systemctl daemon-reload
```

Now you can verify the change has taken effect for future slapd startup:

```
systemctl cat solserver
```

This should show the updated limits:

```
# /etc/systemd/system/solserver.service.d/override.conf
[Service]
LimitCORE=infinity
LimitNOFILE=524288
```

Once Symas OpenLDAP is installed you can view the process's limits file to see the Max open files limit:

```
systemctl show solserver | grep LimitNOFILE
```

```
LimitNOFILE=524288
LimitNOFILESoft=524288
```

If adjusting these settings after Symas OpenLDAP is installed and slapd is running, a restart of solserver is required to pick up the change.

```
systemctl restart solserver
```



Symas OpenLDAP

How-To Guides

Virtual Memory (AS) Limits

While most Virtual Machine managers do a decent job allocating memory to their respective VMs as needed, the process of doing so can wreak havoc on a memory-reliant application such as OpenLDAP. This behavior can be overridden from within the VM itself.

To configure your VM to utilize all available virtual memory on systemd-based systems use the following process. Note that these steps must be done as root:

```
sudo -s

cd /etc/systemd/system/solserver.service.d

printf "LimitAS=infinity\n" >> override.conf

systemctl daemon-reload
```

Now you can verify the change has taken effect for future slapd startup:

```
systemctl cat solserver
```

This should show the updated limits:

```
# /etc/systemd/system/solserver.service.d/override.conf
[Service]
LimitCORE=infinity
LimitNOFILE=524288
LimitAS=infinity
```

Once Symas OpenLDAP is installed you can view the process's limits file to see the Max open files limit:

```
systemctl show solserver | grep LimitAS
LimitAS=infinity
LimitASSoft=infinity
```

If adjusting these settings after Symas OpenLDAP is installed and slapd is running, a restart of solserver is required to pick up the change.

```
systemctl restart solserver
```