

# Microsoft SQL Server on SUSE Linux Enterprise Server

SUSE Linux Enterprise Server 15 SP6 and later  
Microsoft SQL Server 2022

Bernd Schubert, Software Engineer SAP (SUSE)  
Gereon Vey, Senior Manager SAP Emerging Technologies (SUSE)

**Date:** 2025-2-16

This guide helps users install and configure a basic Microsoft SQL Server deployment on SUSE Linux Enterprise Server.

**Disclaimer:** Documents published as part of the SUSE Best Practices and the Technical Reference Documentation series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

# Contents

- 1 Motivation 4
- 2 Installation 6
- 3 Configuration 9
- 4 Tools 25
- 5 Administration 28
- 6 Summary 32
- 7 Appendix 33
- 8 Legal notice 35
- 9 GNU Free Documentation License 37

# 1 Motivation

## 1.1 Background

Since Microsoft released SQL Server for Linux in 2017<sup>1</sup>, the feature gap<sup>2</sup> for SQL Server between Windows and Linux has been closing quickly with each update. With the 2022 release<sup>3</sup>, it is now extremely viable to consider running SQL Server workloads on Linux, aside from some niche features<sup>4</sup>.

## 1.2 Business case: Why SQL on SUSE Linux Enterprise Server

SUSE has a long history of been used as a host for commercial databases within the Linux world. Many of the initial developments and ports of well know commercial databases to Linux where done on and with SUSE. SUSE Linux Enterprise Server was first with many Linux features used by modern databases like memory handling, file systems and provides a rock solid base for business workloads.

Microsoft makes SQL Server a platform that gives you choices of development languages, data types, on-premises or cloud deployments, and operating systems by bringing the power of SQL Server to Linux, Linux-based containers, and Windows. Many organizations, particularly those that have used SQL Server on Windows, are moving to Linux and now have the advantage of having their database available there. SQL on SUSE is Enterprise-ready with great performance. The TPC-H benchmark results done by HP show that SQL Server on Linux delivers amazing performance. SUSE Linux Enterprise Server proved to be faster than Windows in that benchmark too. This really showcases that SQL Server's functionality, its' performance, and scalability are also there when deployed to SUSE Linux Enterprise Server. Second, there is no difference when

---

1 <https://cloudblogs.microsoft.com/sqlserver/2017/10/02/sql-server-2017-on-windows-linux-and-docker-is-now-generally-available> ↗

2 <https://learn.microsoft.com/en-us/sql/sql-server/what-s-new-in-sql-server-2022?view=sql-server-ver16&viewFallbackFrom=sql-server-linux-ver16> ↗

3 <https://learn.microsoft.com/en-us/sql/sql-server/sql-server-2022-release-notes?view=sql-server-ver16&preserve-view=true> ↗

4 <https://learn.microsoft.com/en-us/sql/sql-server/what-s-new-in-sql-server-2022?view=sql-server-ver16&viewFallbackFrom=sql-server-linux-ver16> ↗

running SQL Server on Linux. For example, the tools you use as developer on Windows are the same as on Linux. As Linux administrator you feel at home too, as SQL Server is installed and updated like any other Linux package using Zypper or YaST.

## 1.3 Audience

This guide is intended for SQL Server DBAs, developers, and DevOps/SRE engineers who are familiar with SQL Server on Windows and are looking to migrate to Linux. Operators who are adding an SQL Server requirement into a primarily Linux environment for consistency and simplicity may prefer tools that run only on Linux servers. Another reason may be lower negotiated pricing for Linux subscriptions to replace existing SQL Server on Windows machines.

## 1.4 Scope

The guide covers a basic installation of SQL Server on SUSE Linux Enterprise Server including the OS tuning specific for SQL Server. It is meant to be agnostic of underlying infrastructure excepting the nuance of registering your server discussed in [Section 2.2, "Server registration"](#) *Section 2.2, "Server registration"*.

## 1.5 Important notice

### 1.5.1 Multi Vendor solution

This section will provided some advise how to handle business critical solutions where multiple vendors are involved.



### 1.5.2 SUSE and Microsoft SQL Server

The solution in this guide requires a support contract for both products. With SUSE (valid and active subscription) for the operating system and with Microsoft a license for MS SQL Server.

Details about the SUSE subscriptions can be found here:

- <https://www.suse.com/support/> 


Details about the SQL Server licensing can be found here:

- <https://www.microsoft.com/en-us/sql-server/sql-server-2022-pricing> 
- <https://www.microsoft.com/licensing/docs/view/SQL-Server> 

In case your infrastructure is running in a public cloud environment, your CSP (cloud service provider) may have different support contract offerings.

- BYOS ( = SUSE Subscription + support for OS)
- PAYG ( = only CSP support for OS)

In case of support is required: Depending on the support contract, a decision must be made.

**Support from Microsoft.** A collection support channels can be found here: <https://learn.microsoft.com/en-us/troubleshoot/sql/database-engine/install/windows/support-policy-sql-server#obtain-support-from-microsoft> 

## 2 Installation

At first the manual installation for a standalone SQL Server is described by starting at [Section 2.1, "System requirements"](#).

### 2.1 System requirements

SQL SERVER HAS THE FOLLOWING SYSTEM REQUIREMENTS FOR LINUX:

- 2 GHz CPU with 2 cores
- x64-compatible only
- 2 GB<sup>1</sup> RAM
- XFS or Ext4 file system for the SQL DB files

- other file systems, such as Btrfs, are not supported
- 6 GB disk space (not including data)
- swap space  $\geq 2$  GB

<sup>1</sup> 2 GB is the minimum required memory to start SQL Server on Linux, which accommodates system threads and internal processes. You must take this amount into consideration when setting **max server memory** and **MemoryLimitMB**.

IF YOU USE NETWORK FILE SYSTEM (NFS) REMOTE SHARES IN PRODUCTION, NOTE THE FOLLOWING SUPPORT REQUIREMENTS:

- Use NFS version 4.2 or higher. Older versions of NFS do not support required features, such as fallocate and sparse file creation, common to modern file systems.
- Locate only the /var/opt/mssql directories on the NFS mount. Other files, such as the SQL Server system binaries, are not supported.

## 2.2 Server registration

To gain access to SUSE repositories, you first need to register your server with SUSEConnect. If you are launching an On-Demand (or Pay-As-You-Go) instance and not a BYOS (Bring Your Own Subscription) instance at a public cloud provider, skip this step.

Usage of SUSE SCC if you have Internet access.

- SUSE SCC (<https://scc.suse.com>)  (SUSE Customer Center)

```
sudo SUSEConnect --regcode ${REGISTRATION_CODE} --email ${EMAIL_ADDRESS}
```

Alternatively, if you have access to a

- SUSE RMT (<https://documentation.suse.com/sles/15-SP6/single-html/SLES-rmt/#book-rmt>)  (Repository Mirroring Tool)

server you want to use, use the --url option instead.

```
sudo SUSEConnect --url ${REGISTRATION_SERVER_URL}
```

The registration procedure for SUSE Multi-Linux Manager client's are different.

- SUSE Multi-Linux Manager (<https://www.suse.com/products/multi-linux-manager>) (former known as SUMA)
- Client Registration (<https://documentation.suse.com/suma/5.0/en/suse-manager/client-configuration/registration-overview.html>)

More information about registering can be found in the [SUSE Linux Enterprise Server 15 SP6 Deployment Guide \(https://documentation.suse.com/sles/15-SP6/single-html/SLES-deployment/#sec-yast-install-scc-registration\)](https://documentation.suse.com/sles/15-SP6/single-html/SLES-deployment/#sec-yast-install-scc-registration).

## 2.3 Repositories

Configure repositories for installation and upgrade of SQL Server 2022 on Linux. To verify packages from Microsoft's SQL Server repositories, first add their package signing key: As non privileged user the `sudo` must be added in front of each command.

```
sudo rpm --import https://packages.microsoft.com/keys/microsoft.asc
```

Then add the repository.

```
sudo zypper ar https://packages.microsoft.com/config/sles/15/mssql-server-2022.repo
sudo zypper ar https://packages.microsoft.com/config/sles/15/prod.repo
sudo zypper --gpg-auto-import-keys refresh
```

The `mssql-server` package requires `gdb` from the development tools module (which in turn needs desktop applications module):

```
sudo SUSEConnect -p sle-module-desktop-applications/15.6/x86_64
sudo SUSEConnect -p sle-module-development-tools/15.6/x86_64
```

Optional: the `sle-ha` module is needed for a High Availability (HA) setup:

```
sudo SUSEConnect -p sle-ha/15.6/x86_64
```

## 2.4 Package Installation

To install the SQL Server package non-interactively, and the add-ons, run the following command:

```
sudo ACCEPT_EULA=Y zypper install -y mssql-server mssql-tools18 unixODBC-devel glibc-locale-base sqlcmd tuned
```



## 3 Configuration

### 3.1 Initial configuration

This section is covering the OS modification, the NIC configuration, the recommended storage setup and the SQL Server configuration.

#### 3.1.1 OS configuration (CPU, Kernel, Memory)

##### CPU | sysctl | disk | memory setting

SUSE Linux Enterprise Server contains a Tuned profile for mssql (within the tuned package), but it's not according to the SQL Server best practices guide. The next steps describe and line out the changes recommended for SQL Server. Using Tuned it automatically configures CPU frequency governor, ENERGY\_PERF\_BIAS, and min\_perf\_pct settings appropriately due to the throughput-performance profile being used as base for the mssql profile. C-States parameter must be configured manually. The disk readahead section is also covered by including the file throughput-performance, please check the settings if they are equal and skip them if not needed.

```
mkdir -p /etc/tuned/mssql
cat >/etc/tuned/mssql/tuned.conf <<EOF
[main]
summary=Optimize for Microsoft SQL Server
include=throughput-performance

[cpu]
force_latency = 5

[vm]
transparent_hugepages = always

[sysctl]
vm.swappiness = 1
vm.dirty_background_ratio = 3
vm.dirty_ratio = 80
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs = 100
vm.max_map_count = 1600000
```

```

net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
kernel.numa_balancing = 0
[disk]
# Comma separated list of devices, all devices if commented out.
# devices = md125,md126,md127
readahead = 4096
EOF

chmod +x /etc/tuned/mssql/tuned.conf

```

Alternative settings are:

```

# For multi-instance SQL deployments, use
transparent_hugepages = madvise
# Multi-instance only. Can be configured by Tuned profile Set host variable
mssql_is_multi_instance: true

```

Depending on the infrastructure it might be necessary that the CPU setting must be disabled (e.g. log message: `cpu0: 'energy_perf_bias' = 'None', expected 'performance'`).

This can be achieved by copying the source file `throughput-performance` and modifying it.

Comment the line `energy_perf_bias=performance`, this will avoid error messages (like mentioned before) if `Tuned` is started and can't set this parameter.

```

mkdir -p /etc/tuned/throughput-performance
cp /usr/lib/tuned/throughput-performance/tuned.conf /etc/tuned/throughput-performance/
tuned.conf

vi /etc/tuned/throughput-performance/tuned.conf
...
[cpu]
governor=performance
# energy_perf_bias=performance
min_perf_pct=100
...

```



## Important

The final step to make the changes permanent is to disable `sysctl` re-execution after `Tuned` has modified the system. Otherwise, `sysctl` resets the settings to a value that is not expected in the above configuration. Disable `reapply_sysctl` by changing the 1 to 0.

```
vi /etc/tuned/tuned-main.conf
...
# Whether to reapply sysctl from the e.g /etc/sysctl.conf, /etc/sysctl.d, ...
# If enabled these sysctls will be re-applied after Tuned sysctls are
# applied, i.e. Tuned sysctls will not override system sysctls.
reapply_sysctl = 0
#reapply_sysctl = 1
...
```

Starting Tuned and activate the start during reboot.

```
systemctl enable --now tuned
```

Load and verify the new profile. Checking the log file will show settings which are not executed or have failed.

```
tuned-adm profile mssql
tuned-adm active
tuned-adm verify
cat /var/log/tuned/tuned.log
```

## Intel CPU C-state

Adapting the Intel C-state setting by appending `intel_idle.max_cstate` and `processor.max_cstate` to the kernel command line. The appropriate BIOS settings are essential.

```
vi /etc/default/grub
...
GRUB_CMDLINE_LINUX_DEFAULT="... intel_idle.max_cstate=1 processor.max_cstate=1"
```

Update the boot loader and reboot the system. (To avoid multiple reboots, a final reboot can be scheduled at this point [server reboot](#))

```
update-bootloader --refresh
reboot
```

After the reboot verify the setting is persistent.

```
cat /proc/cmdline
```

## Tuning storage IO

INFO: As multipath I/O is often used for SQL Server deployments, configure the device mapper (DM) multi-queue target to use the blk-mq infrastructure, by enabling the `dm_mod.use_blk_mq=y` kernel boot option. The default value is `n` (disabled). This setting, when the underlying SCSI devices are using blk-mq, reduces locking overhead at the DM layer. For more information on how to configure multipath I/O, can be found at <https://documentation.suse.com/sles/15-SP6/single-html/SLES-storage/#cha-multipath> or <https://www.suse.com/support/kb/doc/?id=000021020>

SCHEDULER is one of `bfq`, `none`, `kyber`, or `mq-deadline`. DEVICE is the block device (`sda` for example). For permanent I/O scheduler change (reboot persistent) for a particular device, copy `/usr/lib/udev/rules.d/60-io-scheduler.rules` to `/etc/udev/rules.d/60-io-scheduler.rules`, and edit the copied file to suit the needs.

Copy `60-io-scheduler.rules`.

```
cp /usr/lib/udev/rules.d/60-io-scheduler.rules /etc/udev/rules.d/60-io-scheduler.rules
```

Unchanged file:

```
...
# 1. BFQ scheduler for single-queue HDD
ATTR{queue/rotational}!="0", TEST!="%S%p/mq/1", ATTR{queue/scheduler}="bfq",
GOTO="scheduler_end"

# 2. BFQ scheduler for every HDD, including "real" multiqueue
# ATTR{queue/rotational}!="0", ATTR{queue/scheduler}="bfq", GOTO="scheduler_end"

# 3. For "real" multiqueue devices, the kernel defaults to no IO scheduling
# Uncomment this (and select your scheduler) if you need an IO scheduler for them
# TEST=="%S%p/mq/1", ATTR{queue/scheduler}="kyber", GOTO="scheduler_end"

# 4. BFQ scheduler for every device (uncomment if you need ionice or blk-cgroup features)
# ATTR{queue/scheduler}="bfq", GOTO="scheduler_end"

# 5. mq-deadline is the kernel default for devices with just one hardware queue
# ATTR{queue/scheduler}="mq-deadline"
```

The example above shows option 1 uncommented. Depending on the disk type, choose between option 1 to 5 or add a personal one. The old default must be deactivated by comment the line with a `#`.

New setting:

```
...
```

```
# 1. BFQ scheduler for single-queue HDD
# ATTR{queue/rotational}!="0", TEST!="%S%p/mq/1", ATTR{queue/scheduler}="bfq",
  GOT0="scheduler_end"

# 2. BFQ scheduler for every HDD, including "real" multiqueue
# ATTR{queue/rotational}!="0", ATTR{queue/scheduler}="bfq", GOT0="scheduler_end"

# 3. For "real" multiqueue devices, the kernel defaults to no IO scheduling
# Uncomment this (and select your scheduler) if you need an IO scheduler for them
# TEST=="%S%p/mq/1", ATTR{queue/scheduler}="kyber", GOT0="scheduler_end"

# 4. BFQ scheduler for every device (uncomment if you need ionice or blk-cgroup features)
# ATTR{queue/scheduler}="bfq", GOT0="scheduler_end"

# 5. mq-deadline is the kernel default for devices with just one hardware queue
# ATTR{queue/scheduler}="mq-deadline"
### new entry
ATTR{queue/scheduler}="kyber"
```

Activate the setting during runtime and check the configuration for all devices starting with sd:  
**udevadm** usage:

```
udevadm control --reload
udevadm trigger
cat /sys/block/sd?/queue/scheduler
```

### 3.1.2 Network configuration

The network setting depend on the used NIC hardware. Some of the following configuration changes maybe cannot be made or have different possible values as shown. In the following step the execution as root user is described. If the root access is not possible the sudo command must be added in front of each command.

#### Configuring network port buffer size

In the example below, the NIC is named eth1, which is an Intel-based NIC. For Intel based NIC, the recommended buffer size is 4 KB (4096). The value to use here depends on the maximum values recommended for each NIC type and vendor. Make the ethtool option persistent by adding them to the ifcfg file for each interface (last line).

Read the current setting:

```
ethtool -g eth1
```

```
Ring parameters for eth1:
Pre-set maximums:
RX:    4096
RX Mini:  n/a
RX Jumbo: n/a
TX:    4096
TX push buff len: n/a
Current hardware settings:
RX:    256
RX Mini:  n/a
RX Jumbo: n/a
TX:    256
RX Buf Len:  n/a
CQE Size:  n/a
TX Push:  off
RX Push:  off
TX push buff len: n/a
TCP data split:  n/a
```

Set the new values:

```
vi /etc/sysconfig/network/ifcfg-eth1
IPADDR='192.168.1.11/24'
BOOTPROTO='static'
STARTMODE='auto'
ZONE=internal
ETHTOOL_OPTIONS='-G iface rx 4096 tx 4096'
```

Restart the NIC and validate the setting or do all the changes in once and restart the network later.

NIC restart:

```
ifdown eth1

ifup eth1

ethtool -?? eth1
```

For the concrete example it looks like this:

```
ethtool -g eth1
...
Current hardware settings:
RX:            4096
RX Mini:       n/a
RX Jumbo:      n/a
```

```
TX:      4096
...
```

Enable the network port buffer temporarily using the command below and check the settings:

```
ethtool -G eth1 rx 4096 tx 4096
ethtool -g eth1
```

## Enabling jumbo frames

Before enabling jumbo frames, verify that all the network switches, routers, and anything else essential in the network packet path between the clients and the SQL Server support jumbo frames. Add a line in the ifcfg file and set the MTU size.

```
vi /etc/sysconfig/network/ifcfg-eth1
...
MTU=9000
...
```

Restart your NIC, for example with [Section 3.1.2, “Network configuration” \(page 14\)](#) and validate the setting.

```
ip addr |grep 'mtu 9'
```

Enable jumbo frames temporarily using the command `ip`.

```
ip link set eth1 mtu 9000
ip addr |grep 'mtu 9'
```

If your SQL Server is running to this point in time, then the SQL Server needs to be configured for jumbo frames as well. After jumbo frames are enabled, connect to SQL Server and change the network packet size to 8060 using `sp_configure` as shown. (How to connect to the database? [Section 3.1.5, “SQL Server configuration” \(page 22\)](#))

Jumbo frames for SQL Server:

```
EXEC sp_configure 'network packet size', '8060';
GO
RECONFIGURE WITH OVERRIDE;
GO
```

## Setting the port for adaptive RX/TX IRQ coalescing

Meaning interrupt delivery is adjusted to improve latency when the packet rate is low and to improve throughput when the packet rate is high. This setting might not be available across all the different network infrastructures. Thus review the existing network infrastructure and confirm that this is supported. The example below is for the NIC named `eth1`, which is an Intel-based NIC.

```
vi /etc/sysconfig/network/ifcfg-eth1
...
ETHTOOL_OPTIONS_rx='-C iface adaptive-rx on adaptive-tx on'
```

Restart the NIC [Section 3.1.2, “Network configuration” \(page 14\)](#) and validate the settings.

```
ethtool -c eth1
```

Enable the RX/TX IRQ coalescing temporarily using the command `ethtool`.

```
ethtool -C eth1 adaptive-rx on adaptive-tx on
ethtool -c eth1
```

## Enabling RX and TX side of RSS queues

It is recommended to enable receive-side scaling (RSS) and combine the RX and TX sides of RSS queues by default. When working with Microsoft Support, there have been scenarios where disabling RSS has improved the performance as well. Test this setting in test environments before applying it on production environments. The following example is for Intel NICs.

```
vi /etc/sysconfig/network/ifcfg-eth1
...
ETHTOOL_OPTIONS_com='-L iface combined 8'
```

Restart the NIC [Section 3.1.2, “Network configuration” \(page 14\)](#) and validate the settings.

```
ethtool -l eth1
```

Enable combining the RX and TX side of RSS queues temporarily using the command `ethtool` and verify the setting.

```
ethtool -L eth1 combined 8
ethtool -l eth1
```



### 3.1.3 Firewall configuration

The connection to the SQL Server needs two ports opened in the firewall. \* 135/tcp MSDTC \* 1433/tcp SQL Server

Optional: \* 9100/tcp Prometheus node\_exporter \* 9664/tcp ha\_cluster\_exporter

```
firewall-cmd --permanent --add-port=135/tcp --add-port=1433/tcp && firewall-cmd --reload
```

Check the configuration

```
firewall-cmd --list-ports
```

### 3.1.4 Storage configuration

Use storage subsystem with appropriate IOPS, throughput, and redundancy. Based on the Microsoft SQL Server recommendation splitting the storage into 5 parts is the best approach.

- OS + swap (mount point / and /swap)
- data (mount point /data)
- transaction log (mount point /log)
- tempdb (mount point /tempdb)

The default filesystem for OS will be btrfs all others will have xfs. Based on the database sizing the required disk must be provided. The example will describe a Linux Software RAID based setup. The disks sdb - sdi are the partitions from the NVME storage underneath. The command `lsblk` helps to find the right partition name to build a reasonable RAID setup later and avoid having all partitions from only one NVME in the same RAID configuration.

Example for `lsblk`.

```
# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sda          8:0    0   50G  0 disk
├─sda1       8:1    0    8M  0 part
├─sda2       8:2    0   48G  0 part /var
|
|                               /opt
|
|                               /home
|
|                               /srv
|
|                               /root
|
|                               /usr/local
17          /tmp
```

```
|
|                                     /boot/grub2/x86_64-efi
|
|                                     /boot/grub2/i386-pc
|                                     /.snapshots
|                                     /
└─sda3   8:3    0    2G    0 part [SWAP]
sdb      8:16   0   10G    0 disk
sdc      8:32   0   10G    0 disk
sdd      8:48   0   10G    0 disk
sde      8:64   0   10G    0 disk
sdf      8:80   0   11G    0 disk
sdg      8:96   0   11G    0 disk
sdh      8:112  0   12G    0 disk
sdi      8:128  0   12G    0 disk
```

INFO: The following chapter is done as root user, if this is not possible the sudo must be put in front of each command.

Create the database storage RAID:

```
# For Data volume, using 4 devices, in RAID 5 configuration with 8KB stripes
mdadm --create --verbose /dev/md0 --level=raid5 --chunk=8K --raid-devices=4 /dev/sdb /
dev/sdc /dev/sdd /dev/sde

# For Log volume, using 2 devices in RAID 10 configuration with 64KB stripes
mdadm --create --verbose /dev/md1 --level=raid10 --chunk=64K --raid-devices=2 /dev/sdg /
dev/sdf

# For tempdb volume, using 2 devices in RAID 0 configuration with 64KB stripes
mdadm --create --verbose /dev/md2 --level=raid0 --chunk=64K --raid-devices=2 /dev/sdi /
dev/sdh
```

The Ext4 and XFS file systems are supported, therefore format the volumes with XFS (case-sensitive).

```
mkfs.xfs /dev/md0 -f -n version=ci -L datavolume
mkfs.xfs /dev/md1 -f -n version=ci -L logvolume
mkfs.xfs /dev/md2 -f -n version=ci -L tempdb
```

Use the noatime attribute with any file system that stores SQL Server data and log files. Creating the mount point:

```
mkdir -p /data /log /tempdb
```

Mounting the disks during boot requires an entry in the /etc/fstab. The blkid helps to find the right UUID for each disk.

```
10 blkid |grep md
Microsoft SQL Server on SUSE Linux Enterprise Server

vi /etc/fstab
```

```
...
UUID="xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx" /data xfs rw,attr2,noatime 0 0
UUID="xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx" /log xfs rw,attr2,noatime 0 0
UUID="xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx" /tempdb xfs rw,attr2,noatime 0 0

mount -a
```

Create subdirectories for the SQL Server:

```
mkdir -p /data/masterdata /log/masterlog
```

Prepare file system permission for the dedicated storage partition:

```
chown -R mssql.mssql /data /tempdb /log
```

Check the disk are mounted and have the right permissions:

```
df -hT
ls -l /
```

INFO: This would be a good point in time to reboot your system to activate all the configuration changes and settings before the SQL Server is started the first time. This would also help to check if all settings made before are reboot persistent.

server reboot.

```
reboot
```

### 3.1.5 SQL Server configuration

To configure and start SQL Server, `mssql-conf` can be used to accept the EULA, set the SQL Server Edition, the SA password and many more details.

Add `mssql-tools18` to your `PATH` environment variable in a Bash shell:

```
echo 'export PATH="$PATH:/opt/mssql-tools18/bin"' >> ~/.bash_profile
echo 'source ~/.bash_profile' >> ~/.bashrc
source ~/.bash_profile
```

Running the command `/opt/mssql/bin/mssql-conf list` will a list of all possible parameter what can be set.

<code>/opt/mssql/bin/mssql-conf list</code>	
<code>control.alternatewritethrough</code>	Enable optimized write through
<code>flush for O_DSYNC requests</code>	
<code>control.hestacksize</code>	Host extension stack size in KB
<code>control.stoponquestprocessfault</code>	Stops the process if any guest
<code>process reports unhandled exception</code>	

<code>control.writethrough</code>	Use <code>O_DSYNC</code> for file flag write
<code>through requests</code>	
<code>coredump.captureminiandfull</code>	Capture both mini and full core
<code>dumps</code>	
<code>coredump.coredumptype</code>	Core dump type to capture: mini,
<code>miniplus, filtered, full</code>	
<code>coredump.disablecoredump</code>	SQL Server disable core dump
<code>distributedtransaction.allowonlysecurerpcalls</code>	Configure secure only rpc calls
<code>for distributed transactions</code>	
<code>distributedtransaction.fallbacktounsecurerpcifnecessary</code>	Configure security only rpc calls
<code>for distributed transactions</code>	
<code>distributedtransaction.maxlogsize</code>	DTC log file size in MB. Default
<code>is 64MB</code>	
<code>distributedtransaction.memorybuffersize</code>	Circular buffer size in which
<code>traces are stored. This size is in MB and default is 10MB</code>	
<code>distributedtransaction.servertcpport</code>	MSDTC rpc server port
<code>distributedtransaction.trace_cm</code>	Traces in the connection manager
<code>distributedtransaction.trace_contact</code>	Traces the contact pool and
<code>contacts</code>	
<code>distributedtransaction.trace_gateway</code>	Traces Gateway source
<code>...</code>	

The following options are used to do a simple setup for a host with 1TB RAM.

```
/opt/mssql/bin/mssql-conf set sqlagent.enabled true
/opt/mssql/bin/mssql-conf set sqlagent.errorlogfile /log/sqlagent.log
/opt/mssql/bin/mssql-conf set sqlagent.errorlogginglevel 2
/opt/mssql/bin/mssql-conf set telemetry.customerfeedback false
/opt/mssql/bin/mssql-conf set filelocation.defaultdatadir /data
/opt/mssql/bin/mssql-conf set filelocation.defaultlogdir /log
/opt/mssql/bin/mssql-conf set filelocation.masterdatafile /data/masterdata/master.mdf
/opt/mssql/bin/mssql-conf set filelocation.masterlogfile /log/masterlog/mastlog.ldf
/opt/mssql/bin/mssql-conf set network.rpcport 135
/opt/mssql/bin/mssql-conf set control.writethrough 1
/opt/mssql/bin/mssql-conf set control.alternatewritethrough 0
# for example using 950GB of 1048576MB available RAM for MSSQL DB instead of 80% default,
# example for 1TB
/opt/mssql/bin/mssql-conf set memory.memorylimitmb 950000
```

## Initial setup and start of SQL Server

Finally configure and start `mssql - server` (`mssql - conf` starts the `mssql - server` immediately after configuring), run the following command:

```
ACCEPT_EULA='Y' MSSQL_PID='Developer' MSSQL_SA_PASSWORD='Strong!Passw0rd' /opt/mssql/bin/
mssql-conf setup
```

```
#alternatively
MSSQL_MEMORY_LIMIT_MB='<some value here>' ACCEPT_EULA='Y' MSSQL_PID='Developer'
MSSQL_SA_PASSWORD='Strong!Passw0rd' /opt/mssql/bin/mssql-conf setup
```

- ACCEPT\_EULA accepts the SQL Server EULA
- MSSQL\_SA\_PASSWORD sets the SA user password. Ensure password requirements as outlined in *Section 7.1, “SQL Server password requirements”* are followed.
- MSSQL\_PID sets the SQL Server edition, acceptable values are:
  - Evaluation
  - Developer
  - Express
  - Web
  - Standard
  - Enterprise
  - EnterpriseCore
  - A product key



## Warning

It is recommended to change the SA password later with `mssql-conf set-sa-password` or disable the history prior to configuring SQL Server with `set +o history`, and re-enabling it afterward with `set -o history` (Bash).

If you specify a product key, it must be in the format of `#####-#####-#####-#####-#####`, where `#` is a number or a letter in summary 25.

- SQL Server should be started at this point. You can verify this with `systemctl status mssql-server.service`.
- By default, SQL Server listens for connections on port `1433`, providing a second method to verify the SQL Server is up and running.
- SQL Server should be accessible with `sqlcmd`.

## systemd check:

```
systemctl status mssql-server.service
```

```
● mssql-server.service - Microsoft SQL Server Database Engine
   Loaded: loaded (/usr/lib/systemd/system/mssql-server.service; enabled; preset:
disabled)
   Active: active (running) since Thu 2024-10-17 17:21:22 CEST; 7min ago
     Docs: https://docs.microsoft.com/en-us/sql/linux
    Main PID: 12965 (sqlservr)
      Tasks: 180
         CPU: 16.947s
    CGroup: /system.slice/mssql-server.service
            └─12965 /opt/mssql/bin/sqlservr
            └─12970 /opt/mssql/bin/sqlservr
```

## Open port:

```
ss -tulpan |grep 1433
tcp    LISTEN 0      128          0.0.0.0:1433      0.0.0.0:*
      users: (("sqlservr",pid=23066,fd=102))
tcp    ESTAB  0      0          127.0.0.1:53189    127.0.0.1:1433
      users: (("sqlservr",pid=23066,fd=137))
tcp    ESTAB  0      0          127.0.0.1:35775    127.0.0.1:1433
      users: (("sqlservr",pid=23066,fd=122))
tcp    ESTAB  0      0          127.0.0.1:1433
127.0.0.1:53189 users: (("sqlservr",pid=23066,fd=135))
tcp    ESTAB  0      0          127.0.0.1:1433
127.0.0.1:35775 users: (("sqlservr",pid=23066,fd=134))
tcp    ESTAB  0      0          127.0.0.1:33679    127.0.0.1:1433
      users: (("sqlservr",pid=23066,fd=108))
tcp    ESTAB  0      0          127.0.0.1:1433
127.0.0.1:33679 users: (("sqlservr",pid=23066,fd=139))
tcp    LISTEN 0      128          *:1433           *:.*
      users: (("sqlservr",pid=23066,fd=105))
```

## sqlcmd login with `sqlcmd -S <hostname> -U SA -P <password>`:

```
sqlcmd -S mssql -U SA -P Strong\!Passw0rd -Q "SELECT @@VERSION" 2>/dev/null
```

```
-----
Microsoft SQL Server 2022 (RTM-CU15-GDR) (KB5046059) - 16.0.4150.1 (X64)
Sep 25 2024 17:34:41
Copyright (C) 2022 Microsoft Corporation
Developer Edition (64-bit) on Linux (SUSE Linux Enterprise Server 15 SP6) <X64>

(1 row affected)
```

## Moving tempdb files to a dedicate storage area

By default, a new installation of SQL Server on Linux creates multiple tempdb data files, based on the number of logical cores (with up to eight data files). One of our preparation steps for the storage component involved setting up a dedicated RAID array for the tempdb files. This setting can only be modified from inside the SQL Server. The following example moves tempdb from its current default location on the disk to another disk location. Because tempdb is re-created each time the MSSQLSERVER service is started, you do not need to physically move the data and log files. The files are created when the service is restarted. Until the service is restarted, tempdb continues to function in its existing location. Determine the logical file names of the tempdb database and their current location on the disk.

Check the current location with `sqlcmd -S <hostname> -U SA -P <password>`. (decide if a move is required or not)

```
sqlcmd -S mssql -U SA -P Strong\!Passw0rd
SELECT name, physical_name
FROM sys.master_files
WHERE database_id = DB_ID('tempdb');
GO
```

INFO: If this is the first time the database is up and running enable the jumbo frames for the SQL Server if needed. [Section 3.1.2, "Network configuration" \(page 15\)](#)

If all the preparation are made in the right order the `/tempdb` directory should already be in use and set as default location. In that situation the next step can be skipped.

Change the location of each file by using `ALTER DATABASE` with `sqlcmd -S <hostname> -U SA -P <password>`:

```
sqlcmd -S mssql -U SA -P Strong\!Passw0rd
USE master;
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = tempdev, FILENAME = '/tempdb/tempdb.mdf');
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = templog, FILENAME = '/tempdb/templog.ldf');
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = tempdev2, FILENAME = '/tempdb/tempdb2.ndf');
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = tempdev3, FILENAME = '/tempdb/tempdb3.ndf');
```

```
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = tempdev4, FILENAME = '/tempdb/tempdb4.ndf');
GO
```

Stop and restart the instance of SQL Server.

Verify the file change with `sqlcmd -S <hostname> -U SA -P <password>`:

```
sqlcmd -S mssql -U SA -P Strong\!Passw0rd
SELECT name, physical_name
FROM sys.master_files
WHERE database_id = DB_ID('tempdb');
GO
```

```
ls -lh /tempdb/*
-rw-rw---- 1 mssql mssql 8,0M 11. Nov 11:05 /tempdb/tempdb2.ndf
-rw-rw---- 1 mssql mssql 8,0M 11. Nov 11:05 /tempdb/tempdb3.ndf
-rw-rw---- 1 mssql mssql 8,0M 11. Nov 11:05 /tempdb/tempdb4.ndf
-rw-rw---- 1 mssql mssql 8,0M 11. Nov 11:05 /tempdb/tempdb.mdf
-rw-rw---- 1 mssql mssql 8,0M 11. Nov 11:09 /tempdb/templog.ldf
```

Delete the `tempdb.mdf` and `templog.ldf` files from their original location.

With SQL Server 16 more memory specific improvements are made with can be enabled. For more details and SQL Server specific internal tunings follow the official documentation from Microsoft. <https://learn.microsoft.com/en-us/sql/relational-databases/in-memory-data-base?view=sql-server-linux-ver16#memory-optimized-tempdb-metadata> ↗

## Further SQL Server configuration modification

For further configuration, use `mssql-conf` to set additional parameters. Changes will take effect after a restart:

```
mssql-conf set ${parameter}
systemctl restart mssql-server
```

Available `mssql-conf` options are described in [Configure SQL Server on Linux with the mssql-conf tool](https://learn.microsoft.com/en-us/sql/linux/sql-server-linux-configure-mssql-conf?view=sql-server-ver16) (https://learn.microsoft.com/en-us/sql/linux/sql-server-linux-configure-mssql-conf?view=sql-server-ver16) ↗.

An alternative way to configure SQL Server is using the `/var/opt/mssql/mssql.conf` file. Settings are stored in the `mssql.conf` format `mssql.conf` (https://learn.microsoft.com/en-us/sql/linux/sql-server-linux-configure-mssql-conf?view=sql-server-ver16#mssqlconf-format) ↗.



A sample `mssql.conf` file is shown below. Edit the `mssql.conf` file and restart `mssql-server` to apply changes.

```
[sqlagent]
enabled = true
errorlogginglevel = 2
errorlogfile = /log/sqlagent.log

[telemetry]
customerfeedback = false

[filelocation]
defaultdatadir = /data
defaultlogdir = /log
masterdatafile = /data/masterdata/master.mdf
masterlogfile = /log/masterlog/mastlog.ldf

[network]
rpcport = 135

[EULA]
accepteula = Y

[control]
writethrough = 1
alternatewritethrough = 0
```

## 4 Tools

Now that SQL Server is running, you can query it. The `mssql-tools` package includes `sqlcmd`, which is a shell to query SQL Server. Install it similarly to the `mssql-server` package. (skip the step if it was done in the chapter Repositories)

Add the repository:

```
zypper addrepo --refresh --check https://packages.microsoft.com/config/sles/15/prod.repo
```

Install the `mssql-tools` package:

```
ACCEPT_EULA=Y zypper install --no-confirm mssql-tools
```

Add the tools to your `PATH`:

```
vi ~/.bashrc
```

```
echo 'export PATH="$PATH:/opt/mssql-tools18/bin"' >> ~/.bash_profile
source ~/.bash_profile
```

Then start `sqlcmd` and make a query. The `-S` option designates the server. The `-U` option specifies the user. Available options can be found at [sqlcmd Utility Syntax \(https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver16#syntax\)](https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver16#syntax).

Log in with `sqlcmd`

```
sqlcmd -S localhost -U SA
```

```
1> SELECT name from sys.databases
2> GO
name
-----
master
tempdb
model
msdb
tpcc

(5 rows affected)
```

`GO` is required here to execute the previous statements.

To exit `sqlcmd`, input `quit` or `exit`:

The full `sqlcmd` documentation can be found at [sqlcmd Utility \(https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver16\)](https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver16).

## 4.1 Azure Data Studio

Azure Data Studio is a lightweight, cross-platform data management and development tool with connectivity to popular cloud and on-premises databases. Azure Data Studio supports Linux, with immediate capability to connect to Azure SQL and SQL Server. [What is Azure Data Studio? \(https://learn.microsoft.com/en-us/azure-data-studio/what-is-azure-data-studio\)](https://learn.microsoft.com/en-us/azure-data-studio/what-is-azure-data-studio).

Download the **Azure Data Studio** RPM from [Download Azure Data Studio \(https://learn.microsoft.com/en-us/azure-data-studio/download-azure-data-studio?view=sql-server-ver16&tabs=linux-install%2Cwin-user-install%2Csuse-install%2Cwindows-uninstall%2Credhat-uninstall#download-azure-data-studio\)](https://learn.microsoft.com/en-us/azure-data-studio/download-azure-data-studio?view=sql-server-ver16&tabs=linux-install%2Cwin-user-install%2Csuse-install%2Cwindows-uninstall%2Credhat-uninstall#download-azure-data-studio).

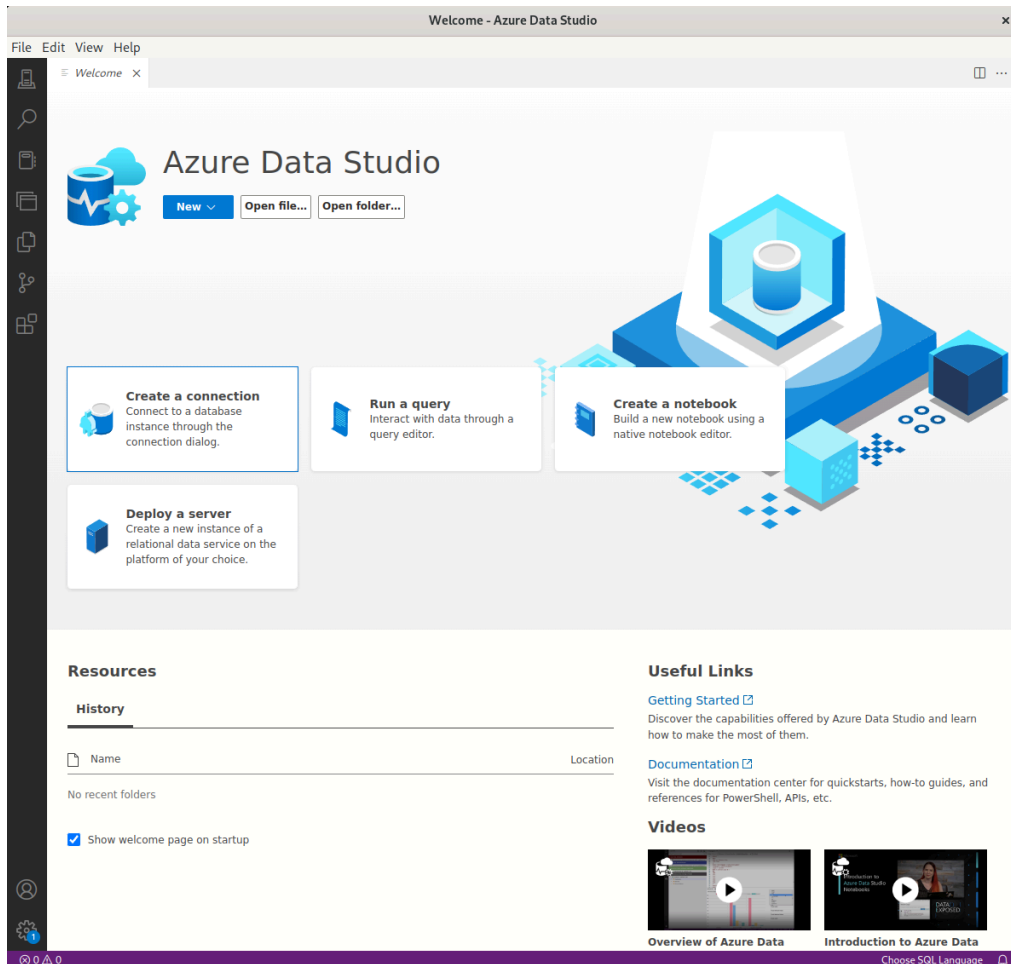
```
curl --location --request GET 'https://go.microsoft.com/fwlink/?linkid=2282380' --output
azure-data-studio.rpm
```

Install Azure Data Studio using .rpm:

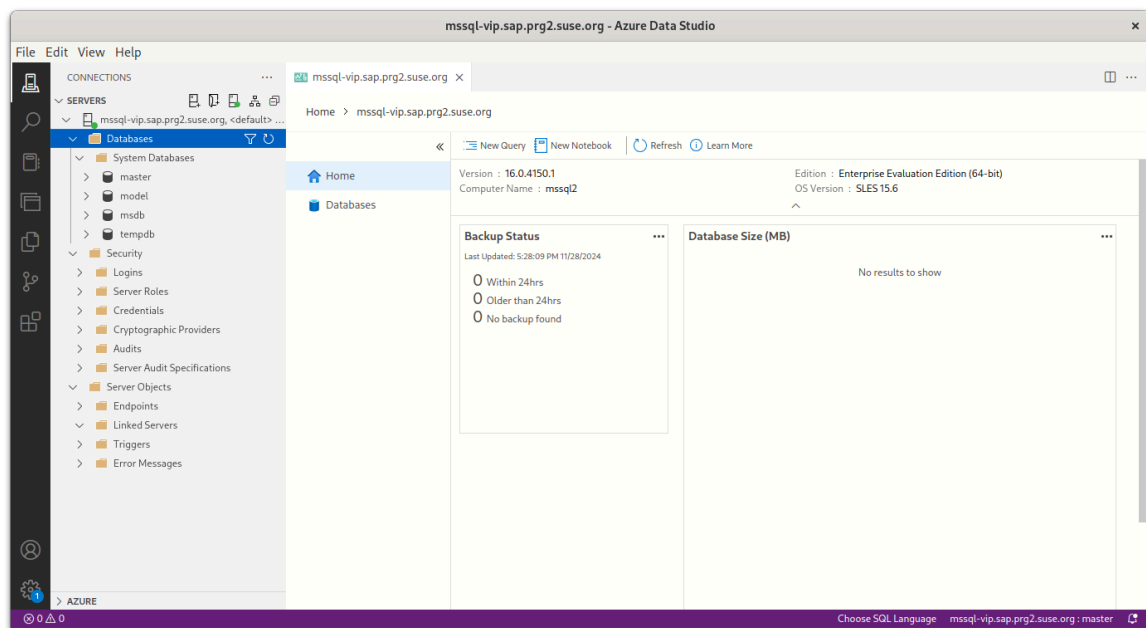
```
zypper install <name given before>.rpm
```

To start Azure Data Studio run the command on your shell:

```
azuredatstudio
```



With Create a connection the SQL Server can be added. A pop-up window is asking for the required parameter. The connection can be established to a single node or the virtual IP of a cluster setup.



## 5 Administration

### 5.1 systemd

The `mssql-server` package installs and configures SQL Server as a `systemd` (<https://systemd.io/>) service. `systemd` provides a framework for managing services, mounts, and system states. You can find more details about `systemd` unit files at `systemd.unit — Unit configuration` (<https://www.freedesktop.org/software/systemd/man/systemd.unit.html>). To control the `mssql-server` service, use `systemctl` to retrieve the `status`, `start`, `stop`, `restart`, `enable`, and `disable` the service.

Display `mssql-server` status as root user:

```
systemctl status mssql-server.service
● mssql-server.service - Microsoft SQL Server Database Engine
   Loaded: loaded (/usr/lib/systemd/system/mssql-server.service; enabled; preset: disabled)
   Active: active (running) since Mon 2024-11-11 11:05:43 CET; 1 week 0 days ago
     Docs: https://docs.microsoft.com/en-us/sql/linux
  Main PID: 23061 (sqlservr)
    Tasks: 195
          CPU: 3h 18min 15.202s
    CGroup: /system.slice/mssql-server.service
```

```
└─23061 /opt/mssql/bin/sqlservr
└─23066 /opt/mssql/bin/sqlservr
```

```
Nov 11 11:05:58 mssql sqlservr[23066]: [75B blob data]
Nov 11 11:05:58 mssql sqlservr[23066]: [96B blob data]
Nov 11 11:05:58 mssql sqlservr[23066]: [100B blob data]
Nov 11 11:05:58 mssql sqlservr[23066]: [71B blob data]
Nov 11 11:05:58 mssql sqlservr[23066]: [124B blob data]
Nov 11 11:05:59 mssql sqlservr[23066]: [157B blob data]
Nov 11 11:05:59 mssql sqlservr[23066]: [193B blob data]
Nov 11 11:05:59 mssql sqlservr[23066]: [155B blob data]
Nov 11 11:05:59 mssql sqlservr[23066]: [204B blob data]
Nov 11 12:05:49 mssql sqlservr[23066]: [97B blob data]
```

Start mssql-server as root user:

```
systemctl start mssql-server
```

Stop mssql-server as root user:

```
systemctl stop mssql-server
```

Restart mssql-server as root user:

```
systemctl restart mssql-server
```

Enable mssql-server to start on boot (mssql-server is enabled by default on installation):

```
systemctl enable mssql-server
```

```
Created a symlink from `/etc/systemd/system/multi-user.target.wants/mssql-server.service`
to `/usr/lib/systemd/system/mssql-server.service`.
```

Disable mssql-server to prevent starting on boot:

```
systemctl disable mssql-server
```

```
Removed symlink /etc/systemd/system/multi-user.target.wants/mssql-server.service.
```

Check the journalctl log for mssql-server.service:

```
journalctl -fu mssql-server.service
```

The parameter -f is to follow the log entries. If a static non updates list is preferred, use journalctl -u mssql-server.service

Check if the network port is active as root user:

```
ss -tulpan |grep 1433
```

## 5.2 Logs

For troubleshooting, the logs and crash dumps are written to `/var/opt/mssql/log` by default. Notable logs are the **errorlogs** (errorlog\*), **trace logs** (\*.trc), **sqlagent logs** (sqlagent\*), and the **extended events logs** (\*.xel). Core dumps are written with the `.tar.gz2` extension and SQL dumps with the `.mdmp` extension. To view these resources, you need root or the mssql user access.

```
ls /var/opt/mssql/log
errorlog      errorlog.6
  HkEngineEventFile_0_133741652559190000.xel  log_5.trc
  system_health_0_133736403350200000.xel  system_health_0_133752031587310000.xel
errorlog.1    FDLAUNCHERRORLOG
  HkEngineEventFile_0_133742310585270000.xel  log_6.trc
  system_health_0_133736403551080000.xel  system_health_0_133753552568770000.xel
errorlog.2    HkEngineEventFile_0_133736403215340000.xel
  HkEngineEventFile_0_133752031521550000.xel  log_7.trc
  system_health_0_133740616123760000.xel  system_health_0_133757931534220000.xel
errorlog.3    HkEngineEventFile_0_133736403466300000.xel
  HkEngineEventFile_0_133753552495030000.xel  log_8.trc
  system_health_0_133741585411750000.xel
errorlog.4    HkEngineEventFile_0_133740616015550000.xel
  HkEngineEventFile_0_133757931472470000.xel  mssql-conf
  system_health_0_133741652635900000.xel
errorlog.5    HkEngineEventFile_0_133741585328790000.xel  log_4.trc
               sqlagentstartup.log  system_health_0_133742310655250000.xel
```

## 5.3 SQL Server starting issues

If the database does not start, check the `journalctl -u mssql-server.service` logs. If the example below is shown, verify the right ownership for the directories and files.

```
systemd[1]: mssql-server.service: Main process exited, code=exited, status=229/
SELINUX_CONTEXT
...
systemd[1]: mssql-server.service: Main process exited, code=exited, status=229
```

Check the ownership for the files in `/data`, `/tempdb` and `/log`. All of them must be owned by `mssql` (UID and GID).

## 5.4 Loading sample data

Microsoft has provided some [sample databases \(https://github.com/microsoft/sql-server-samples\)](https://github.com/microsoft/sql-server-samples) you can use to seed your `mssql-server` instance with some data.

Here is an example of loading our SQL server instance with the sample database `WideWorldImporters`.

Download the `WideWorldImporters` database:

```
curl --location https://github.com/Microsoft/sql-server-samples/releases/download/wide-world-importers-v1.0/WideWorldImporters-Full.bak \
--output /tmp/WideWorldImporters-Full.bak
```

Restore full backup into `mssql-server` with `sqlcmd` while updating paths for the data, user data, transaction log, and in-memory data: (Adapt the password and storage location to your setup)

```
sqlcmd -S localhost \
-U sa \
-P Strong\!Passw0rd \
-Q "RESTORE DATABASE WideWorldImporters \
    FROM DISK = '/tmp/WideWorldImporters-Full.bak' WITH \
    MOVE 'WWI_Primary' TO '/data/WideWorldImporters.mdf', \
    MOVE 'WWI_UserData' TO '/data/WideWorldImporters_UserData.ndf', \
    MOVE 'WWI_Log' TO '/data/WideWorldImporters.ldf', \
    MOVE 'WWI_InMemory_Data_1' TO '/data/WideWorldImporters_InMemory_Data_1'"

Processed 1464 pages for database 'WideWorldImporters', file 'WWI_Primary' on file 1.
Processed 53096 pages for database 'WideWorldImporters', file 'WWI_UserData' on file 1.
Processed 33 pages for database 'WideWorldImporters', file 'WWI_Log' on file 1.
Processed 3862 pages for database 'WideWorldImporters', file 'WWI_InMemory_Data_1' on
file 1.
Converting database 'WideWorldImporters' from version 852 to the current version 904.
Database 'WideWorldImporters' running the upgrade step from version 852 to version 853.
Database 'WideWorldImporters' running the upgrade step from version 853 to version 854.
...
Database 'WideWorldImporters' running the upgrade step from version 902 to version 903.
Database 'WideWorldImporters' running the upgrade step from version 903 to version 904.
RESTORE DATABASE successfully processed 58455 pages in 37.388 seconds (12.214 MB/sec).
```

When loaded, project ten table names from the `WideWorldImporters` database to test it:

```
sqlcmd -S localhost \
```

```
-U sa \  
-P Strong\!Passw0rd \  
-Q "SELECT TOP(10) table_name FROM \  
    WideWorldImporters.information_schema.tables \  
    WHERE table_type = 'BASE TABLE'"
```

table\_name

---

Colors  
Colors\_Archive  
OrderLines  
PackageTypes  
PackageTypes\_Archive  
StockGroups  
StockItemStockGroups  
StockGroups\_Archive  
StateProvinces  
CustomerTransactions

(10 rows affected)

## 6 Summary

Businesses around the world look to SUSE to help them simplify and optimize their IT environments, modernize their applications and infrastructure, and accelerate innovation on-premises, in the cloud, and at the edge. With SUSE Linux Enterprise Server support for Microsoft SQL Server, businesses can streamline their IT landscape and operations without changing their preferred enterprise database management system.

At this point, you should have a fair understanding of how to install SQL Server on SUSE Linux Enterprise Server, install SQL Server tools, query SQL Server and perform basic administration. To stay up to date on the latest SQL Server on Linux features, bookmark [Release notes for SQL Server 2022 on Linux \(https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-release-notes-2022?view=sql-server-ver16\)](https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-release-notes-2022?view=sql-server-ver16).



## 7 Appendix

### 7.1 SQL Server password requirements

Password complexity policies are designed to thwart brute-force attacks by expanding the range of potential password combinations. When the password complexity policy is enforced, new passwords must adhere to the following guidelines:

- The password does not contain the account name of the user.
- The password is at least eight characters long.
- The password contains characters from three of the following four categories:
  - Latin uppercase letters (A through Z)
  - Latin lowercase letters (a through z)
  - Base 10 digits (0 through 9)
  - Non-alphanumeric characters such as: exclamation point (!), dollar sign (\$), number sign (#), or percent (%).

Passwords can be up to 128 characters long. Use passwords that are as long and complex as possible.

### 7.2 Security limitations of SQL Server on Linux

SQL Server on Linux currently has the following limitations:

- A standard password policy is provided. MUST\_CHANGE is the only option you might configure. The CHECK\_POLICY option is not supported.
- Extensible Key Management is not supported.
- SQL Server authentication mode cannot be disabled.
- Password expiration is hard-coded to 90 days if you use SQL Server authentication.
- Using keys stored in the Azure Key Vault is not supported.
- SQL Server generates its own self-signed certificate for encrypting connections. SQL Server can be configured to use a user-provided certificate for TLS.

Find more detailed information at <https://learn.microsoft.com/en-us/sql/linux/sql-server-linux-security-overview?view=sql-server-ver16> and at <https://learn.microsoft.com/en-us/sql/t-sql/statements/alter-login-transact-sql?view=sql-server-ver16>.

## 7.3 Disabling the sa account as a best practice

Follow the instructions from Microsoft: [Disable SA user \(https://learn.microsoft.com/en-us/sql/linux/quickstart-install-connect-suse?view=sql-server-linux-ver16#disable-the-sa-account-as-a-best-practice\)](https://learn.microsoft.com/en-us/sql/linux/quickstart-install-connect-suse?view=sql-server-linux-ver16#disable-the-sa-account-as-a-best-practice)

## 7.4 Virtual machines and dynamic memory

If you are running SQL Server on Linux in a virtual machine, make sure you select options to fix the amount of memory reserved for the virtual machine. Do not use features like Hyper-V Dynamic Memory. The feature for KVM is called memballoon and can be set to none in the VM configuration. The default is virtio.

**KVM memballoon.**

```
virsh edit --domain mssql

...
    <memballoon model='none'>
        <address type='pci' domain='0x0000' bus='0x05' slot='0x00' function='0x0' />
    </memballoon>
...
```

The VM must be powered off to activate the changes. A reboot is not enough.

**before.**

```
lspci |grep balloon

05:00.0 Unclassified device [00ff]: Red Hat, Inc. Virtio 1.0 memory balloon (rev 01)
```

**after.**

```
lspci |grep balloon
```

## 7.5 References

- SUSE Linux Enterprise Server 15 SP6 (<https://documentation.suse.com/sles/15-SP6/>) ↗
- SUSE Multi-Linux Manager (<https://www.suse.com/products/multi-linux-manager>) ↗
- Repository Mirroring Tool Guide (<https://documentation.suse.com/sles/15-SP6/single-html/SLES-rmt/#book-rmt>) ↗
- Installation guidance for SQL Server on Linux (<https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-setup?view=sql-server-ver16>) ↗
- Quickstart: Install SQL Server and create a database on SUSE Linux Enterprise Server (<https://docs.microsoft.com/en-us/sql/linux/quickstart-install-connect-suse?view=sql-server-ver16>) ↗
- Configure SQL Server on Linux with the mssql-conf tool (<https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-configure-mssql-conf?view=sql-server-ver16>) ↗
- sqlcmd Utility (<https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver16>) ↗
- Password Policy (<https://learn.microsoft.com/en-us/sql/relational-databases/security/password-policy?view=sql-server-ver16>) ↗
- Security limitations for SQL Server on Linux (<https://learn.microsoft.com/en-us/sql/linux/sql-server-linux-security-overview?view=sql-server-ver16>) ↗
- SQL DB samples (<https://learn.microsoft.com/en-us/sql/samples/sql-samples-where-are?view=sql-server-ver16>) ↗
  - sql-server-samples (<https://github.com/microsoft/sql-server-samples>) ↗

## 8 Legal notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/> ↗.

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Below we draw your attention to the license under which the articles are published.

## 9 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

### 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.



- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

## ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.