

Jay Falck, CISSP, HCISPP



# Linux for SQL Server Professionals

# Please Support Our Sponsors



# Local User Groups

## SQL Malibu User Group

3<sup>rd</sup> Wednesday of each month (remote)  
[sqlmalibu.pass.org](http://sqlmalibu.pass.org)

## Los Angeles User Group

3<sup>rd</sup> Thursday of each odd month  
[sqlla.pass.org](http://sqlla.pass.org)

## Los Angeles – Korean only

Every Other Tuesday  
[sqlangeles.pass.org](http://sqlangeles.pass.org)

## Orange County User Group

2<sup>nd</sup> Thursday of each month  
[bigpass.pass.org](http://bigpass.pass.org)

## San Diego User Group

1<sup>st</sup> & 3<sup>rd</sup> Thursday of each month  
[meetup.com/sdsqplug](http://meetup.com/sdsqplug)  
[meetup.com/sdsqbig](http://meetup.com/sdsqbig)

SQLSaturday Sacramento

SQLSaturday San Diego

#SQLSatLA

July 28

September 22



# PASS Summit

Annual International Conference

November 6 -9 | Seattle, WA

2 Days of Pre-Cons

200+ sessions over 3 days

Over 5,000 SQL Professionals

Evening Networking Activities

**Discount Code: LSDISEI2J**

#SQLSatLA



# Agenda

- A little about me
- \*nix versions (a few anyway)
- Why SQL Server on Linux
- Where do I get Linux?
- Where do I get SQL Server for Linux?
- Introduction to the BASH Shell
- Common Command Equivalents between Windows and Linux

#SQLSatLA



## Agenda (continued)

- User and group setup and associated security
- Applying updates
- Location of database, transaction log, and server log files
- Task scheduling using cron
- Backup and restore
- Accessing files on a Windows System from Linux

#SQLSatLA



## Agenda (continued)

- Accessing Linux files from Windows
- PowerShell and Linux
- Commands for Tuning
- Advanced and/or not ready for Prime Time Topics
- After thoughts
- Online resources
- Questions

#SQLSatLA



## A little about me

- Worked as night shift computer operator in high school
- Application developer for a large state agency
- Mainframe system programmer at two large state agencies
- DBA and developer for State Parks Reservation System
- Transitioned into Healthcare IT in 1996
- Two main IT passions
  - Make it secure
  - Make it fast



#SQLSatLA



## \*nix versions (a few anyway)



#SQLSatLA



- These are a few different version of \*nix over the years.

- Coherent was the first version I used on Intel processors
- I used IBM AIX and HP-UX in the mid-90's to early 2000's for healthcare claims processing
- Solaris was used in the Public Lands division at TPWD for geographic mapping and CAD for state park enhancements
- Fedora, Red Hat (enterprise), and CentOS (development, free) are one of three basic distribution groups of Linux
- Debian and Ubuntu are another. Ubuntu is free, published by Canonical Ltd from the UK with an office in Austin. Canonical provides paid support.
- SUSE and OpenSUSE are another. This was popular in mid '90's because it was relatively low cost. In 2003 SUSE was acquired by Novell. In 2011 Novell and

SUSE were acquired by The Attachmate Group who allowed SUSE to operate as an independent company. In October 2014, the Attachmate portfolio was acquired by Micro Focus International (yes, the COBOL folks) with SUSE still operating as an independent.

## Why SQL Server on Linux?

- Linux in business world is growing
- Linux is leading the charge in some key areas
  - Big Data
  - Containers
  - DevOps
- There are shops that only have Windows for SQL Server. Everything else they do is some form of \*nix.

#SQLSatLA



## Where do I get Linux?

- [Red Hat Enterprise - Developer Edition](#)
- [SUSE](#)
- [Ubuntu \(server\)](#)
  - This link is for most current version. You may need to get an older version until Microsoft certifies latest.

#SQLSatLA



## Where do I get SQL Server for Linux?

- The following distros are supported by Microsoft at this time:
  - [Red Hat](#)
  - [SUSE](#)
  - [Ubuntu](#)
  - [Docker](#)
  - [Azure VM](#)
- If you were testing the preview version, you will need to update to GA repositories
  - [Configure Repositories](#)

#SQLSatLA



# Introduction to the BASH shell

- What is BASH
- How to make a script executable
- Pipes and redirection
- Background tasks
- Useful commands available in BASH
  - grep – search a file based on regular expressions
  - awk – parse text files
  - find – find files in a directory (including subdirectories) and print the filenames, or perform another command

#SQLSatLA

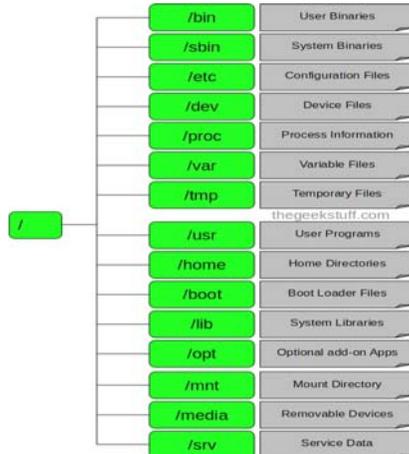


The BASH shell is the most common default shell distributed with Linux. The original shell in \*nix was called the bourne shell. BASH is the bourne-again shell as it replaced the bourne shell.

Default naming convention for BASH shell scripts is name.sh. This is just for human readability as the extension in \*nix, unlike DOS, doesn't cause the shell to do anything special. If you want the script to be executable, you give the file execute attribute using the chmod command. You can also place a header in the script that will tell the shell which command processor should be used to execute the script. This is sometimes referred to as "sh-bang" and the format is `#!/<pathto processor>`. For example, to execute using the default shell use `#!/bin/sh` or `#!/bin/bash`. If you've created a python script it would be `#!/bin/python`. Once this is added, you can invoke by simply typing `./scriptname`.

# Introduction to the BASH shell (continued)

- Typical Linux folder layout



#SQLSatLA



## 1. / – Root

Every single file and directory starts from the root directory.

Only root user has write privilege under this directory.

Please note that /root is root user's home directory, which is not same as /.

## 2. /bin – User Binaries

Contains binary executables.

Common linux commands you need to use in single-user modes are located under this directory.

Commands used by all the users of the system are located here.

For example: ps, ls, ping, grep, cp.

## 3. /sbin – System Binaries

Just like /bin, /sbin also contains binary executables.

But, the linux commands located under this directory are used typically by system administrator, for system maintenance purpose.

For example: iptables, reboot, fdisk, ifconfig, swapon

## 4. /etc – Configuration Files

Contains configuration files required by all programs.

This also contains startup and shutdown shell scripts used to start/stop individual programs.

For example: /etc/resolv.conf, /etc/logrotate.conf

## 5. /dev – Device Files

Contains device files.

These include terminal devices, usb, or any device attached to the system.

For example: /dev/tty1, /dev/usbmon0

## 6. /proc – Process Information

Contains information about system process.

This is a pseudo filesystem contains information about running process. For example:

/proc/{pid} directory contains information about the process with that particular pid.

This is a virtual filesystem with text information about system resources. For example:

/proc/uptime

## 7. /var – Variable Files

var stands for variable files.

Content of the files that are expected to grow can be found under this directory.

This includes — system log files (/var/log); packages and database files (/var/lib); emails (/var/mail); print queues (/var/spool); lock files (/var/lock); temp files needed across reboots (/var/tmp);

## 8. /tmp – Temporary Files

Directory that contains temporary files created by system and users.

Files under this directory are deleted when system is rebooted.

## 9. /usr – User Programs

Contains binaries, libraries, documentation, and source-code for second level programs.

/usr/bin contains binary files for user programs. If you can't find a user binary under /bin, look under /usr/bin. For example: at, awk, cc, less, scp

/usr/sbin contains binary files for system administrators. If you can't find a system binary under /sbin, look under /usr/sbin. For example: atd, cron, sshd, useradd, userdel

/usr/lib contains libraries for /usr/bin and /usr/sbin

/usr/local contains users programs that you install from source. For example, when you install apache from source, it goes under /usr/local/apache2

## 10. /home – Home Directories

Home directories for all users to store their personal files.

For example: /home/john, /home/nikita

## 11. /boot – Boot Loader Files

Contains boot loader related files.

Kernel initrd, vmlinux, grub files are located under /boot

For example: initrd.img-2.6.32-24-generic, vmlinuz-2.6.32-24-generic

## 12. /lib – System Libraries

Contains library files that supports the binaries located under /bin and /sbin

Library filenames are either ld\* or lib\*.so.\*

For example: ld-2.11.1.so, libncurses.so.5.7

## 13. /opt – Optional add-on Applications

opt stands for optional.

Contains add-on applications from individual vendors.

add-on applications should be installed under either /opt/ or /opt/ sub-directory.

## 14. /mnt – Mount Directory

Temporary mount directory where sysadmins can mount filesystems.

15. /media – Removable Media Devices

Temporary mount directory for removable devices.

For examples, /media/cdrom for CD-ROM; /media/floppy for floppy drives;

/media/cdrecorder for CD writer

16. /srv – Service Data

srv stands for service.

Contains server specific services related data.

For example, /srv/cvs contains CVS related data.

## Introduction to the BASH shell (continued)

- File permissions

```
drwxr-xr-x 7  root  root  4096  Nov 28  11:54  ..
drwxrwxr-x 2  root  root  4096  Nov 20  15:38  .
-rw-rw-r-- 1  root  root  1289  Nov 20  15:38  restore_aw2017.sql
-rw-rw-r-- 1  root  root   180  Nov 20  15:37  update_aw2017.sql
-rw-rw-r-- 1  root  root   875  Nov 20  15:37  backup_aw2017.sql
-rw-rw-r-- 1  root  root   327  Nov 20  15:34  restore_aw.sql
-rw-rw-r-- 1  root  root    31  Nov 11  17:24  helloworld_2.sh
-rw-rw-r-- 1  root  root    23  Nov 11  17:22  helloworld_1.py
-rw-rw-r-- 1  root  root   214  Nov 10  15:11  backup_aw.sql
-rw-rw-r-- 1  root  root    85  Nov 10  14:51  awk.sh
-rw-rw-r-- 1  root  root   599  Nov 10  01:22  chineseyear.ps1
-rw-rw-r-- 1  root  root    42  Nov  9  18:34  helloworld_2.py
-rw-rw-r-- 1  root  root    19  Nov  9  18:26  helloworld_1.sh
root@ucdevlinuxsql:~/scripts$ _
```

#SQLSatLA



## Common Command Equivalents between Windows and Linux

Description	Windows	Linux
Copy a file	copy	cp
Move a file	move	mv
Rename a file	ren	mv
Delete a file	del	rm
Edit a file	notepad	vi, nano, pico, vim
Run elevated	runas	su sudo
Change file attributes	attrib	chmod
Change group attributes	attrib	chgrp

#SQLSatLA



There are three big differences between Windows and \*nix that will take some getting used to:

- 1) \*nix is case sensitive
- 2) Where Windows (usually) requires a backslash, \*nix always uses forward slash. In \*nix the backslash is an escape character
- 3) Line terminators are different. Linux uses newline (1 ascii character), Windows/DOS use CR/LF (2 ascii characters)

## Common Command Equivalents between Windows and Linux (continued)

Description	Windows	Linux
Display file contents	type	cat less more head tail
Change directory	cd	cd
Create directory	md	mkdir
Delete directory	rd	rmdir
List directory contents	dir	ls
Clear screen	cls	clear

#SQLSatLA



## Common Command Equivalents between Windows and Linux (continued)

Description	Windows	Linux
Archive files	7zip, RAR, WinZIP	tar
Compress files	7zip, RAR, WinZIP	gzip
Check network config	ipconfig	ifconfig ip
Shutdown/restart	shutdown [/r]	shutdown [-r] [-t n][[now]
Get help	Command /? Maybe	man

#SQLSatLA



## User and group setup and associated security

- Create a new user
  - useradd
- Delete a user
  - userdel
- Modify a user
  - usermod
- Create a group
  - groupadd groupname

#SQLSatLA



Add a user with a home directory: `sudo useradd -m nebula -p IHateMyFather2!`

Groups are just that, groups. They have no inherent meaning (other than root)

`sudo groupadd guardians`

Add a user to a group: `usermod -a -G nebula guardians`

`usermod -a -G groot guardians`

`usermod -a -G groot mssql`

## User and group setup and associated security (continued)

- Location of log files
  - /var/log

#SQLSatLA



# Applying updates

- Ubuntu
  - apt-get
  - apt
- Red Hat
  - yum (yellow dog update manager)
- SUSE
  - zypper

#SQLSatLA



## Location of database, transaction log, and server log files

- By default, SQL Server files are stored in /var/opt/mssql
  - log
  - data
  - secrets
  - .system
  - mssql.conf

#SQLSatLA



You must be root to get to this folder as it is owned by mssql user.

log – location of sql log files

data – default location for system and user databases. With CU1, this can be customized by mssql-conf

secrets –

.system

mssql.conf – Since there is no registry in Linux, settings you would expect to find in the registry are stored here.

# Demos

#SQLSatLA



Show how you have to elevate to superuser to get to the directory and then look at some files.

# Task scheduling using cron

- Why cron
  - Original task scheduler for \*nix
  - Linux admins are familiar with it
  - Early release of SQL Server for Linux did not have SQL Agent
- Format of crontab
  - mm hh dd MM dow <command to execute>
    - \* in any position says to run based on other parameters but any for this one are OK
    - mm – 0-59
    - hh – 0-23
    - dd – 1-31
    - MM – 1-12
    - dow – 0-6 (Sunday is 0)

#SQLSatLA



The crond daemon runs jobs on the system according to information stored in `/etc/crontab`. You can also have jobs run hourly, daily, weekly, and monthly by simply placing your desired script in a special folder named `/etc/cron.<period>`.

Use the `crontab` command to modify the crontab. Arguments are `-e` (edit) `-l` (list) `-r` (remove the file) `-v` (when was crontab last modified)

You can set a default editor for crontab by issuing command `export EDITOR=vi` (or editor of your choice)

If you want a task to run every 2 minutes use `*/2` for the mm value

## Backup and Restore

- Guess what, it's the same as windows
- Principal difference is direction of slashes

#SQLSatLA



## Accessing files on a Windows System from Linux

- The following steps will allow access to file shares on remote windows machines
- Get the cifs utilities
- Create credentials file in home directory and protect it
- Add entry to /etc/fstab (for persistence)
- Now mount the share that was added to /etc/fstab
- It is possible to mount without adding to /etc/fstab

#SQLSatLA



As previously mentioned, most \*nix commands and parameters are either acronyms or the name of the original developers pet. In this case, cifs is an acronym which stands for Common Internet File System.

Adding the entry to /etc/fstab will cause the share to be mounted on each restart.

### /etc/fstab entry

- Note: be careful editing this file as it could corrupt the system. Need to run editor in elevated mode.
- `//servernameoraddress/sharename /mnt/mountname cifs  
credentials=/home/username/.smbcredentials,icharset=utf8,sec=ntlm,vers=3.0,uid=username,gid=groupid 0 0`

Mount the share

```
sudo mount -a
```

Doing a non-persistent mount:

```
sudo mount -t cifs //servernameoraddress /sharename /mnt/mountname cifs -o
credentials=/home/username/.smbcredentials,icharset=utf8,sec=ntlm,vers=3.0,ui
d=username,gid=groupid
```

Get cifs:

- 1) sudo apt-get update
- 2) sudo apt-get install cifs-utils -y

Create credentials:

- 1) cd \$HOME
- vi .smbcredentials
  - Insert username=somewindowsuser
  - Insert password=somewindowspassword
  - Insert (optional) domain=domainnameofauthuser
- chmod 600 ~/.smbcredentials

## Accessing Linux files from a Windows System (Ubuntu)

- Install samba
- Create a samba user and password
- If it doesn't exist, create a folder you wish to share
- Add an entry to the samba configuration file for the new share
- Restart samba daemon
- Validate changes
- Map network drive from windows

#SQLSatLA



To install samba on Ubuntu:

```
sudo apt-get update
sudo apt-get install samba
```

To create a user in samba. Note: The user, if it doesn't exist in the normal passwd, will be added. Recommend using an existing user. Password can be different.

```
sudo smbpasswd -a <username>
```

To add entry to samba configuration

Make a backup copy of the existing file: `sudo cp /etc/samba/smb.conf ~`

Edit the file using editor of choice. Ex: `sudo vi /etc/samba/smb.conf`

Add the following information to the end of the file:

```
[<sharename>]
  path = <path to folder to share>
  valid users = <sambausername>
  read only = no
```

Note: Spaces before and after equal sign is required. Also indent each line under the share name tag.

When creating the samba user, you will be prompted for a password and verification. If you

use a username other than your own, you will need to be sure to grant access to any shares you create.

# Demos

#SQLSatLA



Show that we really can access files across systems.

## PowerShell and Linux

- Don't want to learn BASH or Korn?  
PowerShell works in Linux
- Instructions for installing are available at  
GitHub
  - [Installing PowerShell Core on Linux](#)

#SQLSatLA



- Summarized here
  - # Import the public repository GPG keys
  - `curl https://packages.microsoft.com/keys/microsoft.asc | sudo apt-key add -`
  - # Register the Microsoft Ubuntu repository
  - `curl https://packages.microsoft.com/config/ubuntu/16.04/prod.list | sudo tee /etc/apt/sources.list.d/microsoft.list`
  - # Update the list of products
  - `sudo apt-get update`
  - # Install PowerShell
  - `sudo apt-get install -y powershell`

## Commands for Tuning

- The following commands can be used to check performance of Linux in general
  - ps – list running processes (point in time)
  - top – display running process (real time)
  - kill – kill a running process (be careful)
  - iostat – CPU and IO statistics
  - sar – collect, report, save system activity information
  - pidstat – statistics for Linux tasks
  - mpstat – statistics for each running processor
  - vmstat – virtual memory statistics
  - cifsioat – CIFS statistics (if in use)

#SQLSatLA



Some commands are not installed by default. To install: `sudo apt install sysstat -y`

## Advanced and/or not ready for Prime Time Topics

- [High Availability - RedHat Specific](#)
- [SQL Server Agent](#)
- [SQL Server Full-Text Search](#)
- [Active Directory Integration](#)
- [SQL Server Integration Services](#)

#SQLSatLA



## After thoughts

- Some things I thought of after original outline
  - If using the Ubuntu Desktop, there is a nice Notepad++ clone available
    - `sudo add-apt-repository ppa:notepadqq-team/notepadqq`
    - `sudo apt-get update`
    - `sudo apt-get install notepadqq`
  - Some interesting Linux commands you may or not normally use in DOS
    - `cal` – display calendar for current or specific month
    - `date` – display system date and time
    - `df` – show disk space free
    - `free` – displays memory usage
    - `netstat` – network usage info
    - `sort/uniq` – sort text files and (alternatively) remove duplicate lines
    - `tac` – same as `cat` but in reverse order
    - `touch` – create a new file or change the date/time stamp of existing
    - `wc` – produce count of words, lines, and characters in a file

#SQLSatLA



Cal options: -h turns off highlighting of current date, -j shows julian dates, -3 shows prev and next month, -y specific year, -m specific month

Df -h shows a more friendly display

Wc -l produces line count only, -c counts bytes, -m counts characters, -w counts words, -L length of longest line

## Online Resources

- Demo Code
  - [SQL Saturday 717 Demo Code](#)
- Microsoft Docs
  - [SQL Server on Linux](#)
  - [Performance Best Practices](#)
  - [Quick Start Guide - Registration Required](#)
- Adventure Works Database
  - [AdventureWorks sample databases - All Versions](#)
- Pluralsight
  - Anthony Nocentino
    - SQL Server on Linux Administration Fundamentals
    - Play by Play: Microsoft Open Source Powershell on Linux and Mac

#SQLSatLA



Some MSDN subscriptions include up to 6 months of access to Pluralsight.

## Online Resources (continued)

- Amazon Kindle (Inexpensive but no longer free)
  - [Linux for Beginners](#)
  - [Linux Administration](#)
  - [Command Line Kung Fu](#)
  - [Shell Scripting](#)
  - [Python Programming for Beginners](#)
- 7 Free eBooks for Noobs and Admins
  - <https://goo.gl/vtYtxY>
- Microsoft Virtual Academy
  - <https://mva.microsoft.com>
- [24 Hours of PASS](#)

#SQLSatLA



## Online Resources (continued)

- More free books at GoalKicker
  - [Bash](#)
  - [Git](#)
  - [Linux](#)
  - [Perl](#)
  - [PowerShell](#)
  - [Python](#)
  - [R](#)
  - [SQL](#)
- [We Speak Linux](#)
- [Performance Best Practices](#)

#SQLSatLA



## Questions

- Please fill out the evaluation form

#SQLSatLA



## Contact

- Email: [sql@unicorncomputing.com](mailto:sql@unicorncomputing.com)
- Twitter: @jayfalck
- LinkedIn:  
<https://www.linkedin.com/in/jayfalck>

#SQLSatLA

