

Lenovo Campus Switches Quick Configuration Guide

This document will show some of the key tasks which will need to be performed on essentially all new installs and changes to existing installs of the campus switch products. It includes samples of configuration text to execute these tasks and notes on common errors or issues which have been experienced.

However, this document is not a substitute for the definitive documentation for the command line (CLI) or the browser interface (BBI) on the switch. When there is a conflict between this document and those definitive documents, the definitive documents should be presumed to be correct unless one is informed otherwise.

Contents

These common tasks are documented:

1. Common and basic command notes
2. Firmware and configuration management
3. Management port configuration and access (“service port”)
4. User access and security
5. VLAN definition and port VLAN assignment
6. Link Aggregations
7. Stacking
8. Layer-3 addressing configuration
9. Layer-2 redundancy configuration (failover)
10. Layer-3 redundancy configuration (VRRP)
11. IPv6
12. SNMP configuration
13. Miscellaneous commands

1. Common and basic commands: notes

There are some immediately noticeable differences between the Campus NOS CLI and others including CNOS and ENOS. They are listed below.

- The *configure* command does not have a *'terminal'* option; *config t* or *configure t* will not work.
- The *exit* command will take you to higher levels of configured mode and others, and will get you back to the CLI enabled mode, but will not exit from a telnet or ssh session. In order to exit your session you must use the *logout* command.
- *wr mem* will expand into *write memory* if typed in with a space after each word, and will work but if typed without the automatic expansion, it will not work.
- *copy run start* will not work. It is equivalent to *write memory*, but the full syntax is *copy system:running-config nvram:startup-config*.
 - o It is also possible to copy *nvram:startup-config* and *nvram:backup-config* to each other.
- There is no VRF support under Campus NOS, but one can force the use of the management (service) port with the *'source serviceport'* option on several commands. It is also possible to designate any other interface with an IP address with the *source* option on these commands.
- While in config mode or VLAN database mode, the *"do"* command will allow any enable mode CLI command to be executed. For example, *"do ping xx.xx.xx.xx"*.

Other differences will be discussed in the remaining sections of the document.

2. Firmware and configuration management

Firmware images and configuration files can be retrieved from and saved to external servers on Campus NOS just as they are on CNOS, but the command syntax is noticeably different.

Campus NOS uses an active and backup firmware image just as CNOS does. To upgrade the firmware, the new image can be copied with this command:

- `copy url://<IP-address>/filename active|backup`
 - *The url: can be ftp:, tftp:, scp:, http:, or https:*
 - *It is also possible to save an image to a server by copying from the active or backup image.*

To select which image to use when rebooting the switch, the command is shown below. The *unit-number* option is to boot only a single unit in a stack; by default, every unit in the stack is booted. There is no feature that allows a rolling reboot of a stack, and thus no way to avoid some amount of outage time.

- `boot system [unit-number] active|backup`

The current status of the image to be used on boot can be viewed with this command:

- `show bootvar`

Configuration text files can also be copied to and from a server. There are two copies of the configuration text stored in flash memory – *startup* and *backup* - which can be managed as follows:

- `copy system:running-config nvram:startup-config` – is the equivalent of “*copy run start*”.
 - *“write memory” is supported to perform the same function, but must be spelled out.*
- `copy url:// ... nvram:startup-config` loads a saved configuration into flash from an external server. A boot would be required to make the loaded configuration active on the switch.
 - *ftp, tftp, scp, sftp, http, and https are supported for loading a configuration.*
- `copy nvram:startup-config nvram:backup-config` and the reverse are supported.
- `copy nvram:startup-config url://` is supported, but not for http or https.

3. Management configuration and access

The physical management port on the campus switches is referred to as the *service port* in the documentation. It can be set to learn its address via DHCP or BOOTP, or it can be configured with a static address, using the commands shown below. Note that these commands are all entered from *enable* mode and not from *config* mode.

- Set IP address: *serviceport ip <IP address> <mask> <default gateway address>*
- Set DHCP or BOOTP: *serviceport protocol [DHCP|BOOTP|NONE]* – “NONE” uses the configured static address. This command should be entered before the ip address command; unlike CNOS entering an IP address will not automatically disable the other protocols.

It is possible to enable management access (CLI, browser) via the data ports if this is desired. The commands to do this are shown below:

- Set IP address: *network parms <IP address> <mask> <default gateway address> | NONE*
 - o NONE removes the address shown
- Set DHCP/BOOTP: *network protocol [DHCP|BOOTP|NONE] as above; for the network protocol the default is 'NONE' and so this command is not required before configuring an IP address.*
- Select VLAN for management access: *network mgmt._vlan <vlan-id>*
 - o Note that assigning an address to a VLAN this way does not make it routable to/from other VLANs on the switch.
 - o A routable address takes precedence over a management network address configured on the same network management VLAN. If there is a routable address then the switch will not respond to the address in the *network parms* command.

Telnet and SSH access to the switch can be enabled, disabled, and limited as shown below:

- To enable: *ip telnet server enable* and/or *ip ssh server enable*
- Maximum number of concurrent sessions can be set: *telnetcon | sshcon maxsessions <1-5>*
- Timeout can be increased up to 160 minutes: *telnetcon | sshcon timeout <1-160>*

Outbound telnet sessions are supported with the *telnet* command but there is no corresponding *ssh* command. Outbound telnet sessions must also be enabled as follows:

```
line console  
transport output telnet
```

4. User access and security

Campus NOS uses similar user access features to those in CNOS and ENOS. It can support remote authentication via RADIUS or TACACS, and supports user privilege levels 1 and 15 for TACACS. There are some differences in the required configuration for RADIUS and TACACS.

Local user credentials can also be defined, and the syntax for the commands to do this is noticeably different from both CNOS and ENOS.

Always verify that your security configuration is functioning properly before signing out of the session from which it was configured – open a new session to check that security is working as you intend. You can avoid locking yourself out of the switch by following this practice.

4.1 RADIUS

It is possible to use RADIUS for both authentication (access to the CLI and browser) and accounting (command logging). There are distinct commands so that these functions can be performed by separate servers if desired.

The configuration for a RADIUS server is as follows:

```
radius server [primary] host [acct|auth] <ip address>  
radius source-interface [serviceport|other IP interface] – identifies which port to send radius requests via  
radius server key [acct|auth] <ip address> -- a prompt for the key will be provided and it will be  
encrypted for use in the configuration
```

For use of RADIUS accounting, this command is also needed:

```
radius accounting mode
```

4.2 TACACS.

Configuration for a TACACS server is similar to RADIUS and is as follows:

```
tacacs-server host <ip-address>  
tacacs-server source interface [serviceport|other]  
tacacs-server key <key> -- there is no prompt for the key
```

4.3 Local user configuration

Local users can be configured and can be used to authenticate to gain access to the switch in the event of a failure on remote authentication servers.

The CLI commands for this are as follows:

```
username <name> level [0|1|15] [override-complexity-check] password
```

Level 0 is no access; level 1 is user access; level 15 is enable-mode access. The password is not entered on the line where the command is issued; instead, a prompt is provided.

For a newly created user to have the proper privileges, one additional command must be entered:

```
username <name> usergroup default-usergroup-name
```

The *default-usergroup-name* should be entered exactly as shown; a configured usergroup created by a customer can also be used.

There are numerous commands to manage password aging, complexity, and history under the *passwords* command family; these are for local credentials only. There is also a user groups function which allows users to be easily given a set of attributes.

4.4 Security with servers and local credentials

To enable the use of a remote security server the AAA command family is used, similarly to CNOS.

For sign-on to the switch (CLI or browser), the command is of the form:

```
aaa authentication login default [local|tacacs|radius|none]
```

Multiple options can be used and are tried in the order listed; common use would use RADIUS and then local or TACACS and then local.

802.1x security is supported for devices attached to data ports on the switch and the global command for this is:

```
aaa authentication dot1x default [local|tacacs|radius|none]
```

There are additional options for accounting support, some of which are only usable with RADIUS.

5. VLAN definition and port assignment

A quick sample portion of a configuration to create a VLAN and a routable interface for it; it can be repeated for additional VLANs as desired. For details, read the remainder of this section.

```
vlan database
    vlan 2
    name VLAN-2
    vlan routing 2 1
    ! creates routable interface #1 as interface vlan 2 and also 0/4/1
    exit
! config mode commands
config
interface vlan 2
ip address 192.168.1.1 /24
! note the space before /24; old style mask 255.255.255.0 can also be used
! for routing on the switch, the following is always required
ip routing
! default gateway, if desired
ip route default 192.168.1.254
! add members to the VLAN with ONE of the following
int 1/0/10
! note that the first number (1 above) identified the node number if a stack is used; the "10" identifies the
! physical port number
switchport mode access
switchport access vlan 2

! or
Int 1/0/11
switchport mode trunk
switchport trunk allowed vlan [add] 2
! optionally
switchport trunk native vlan 2
```

More details:

Campus NOS supports both trunk and access port modes, similarly to CNOS, ENOS, and other switch firmware. It can use the same commands to define a port status. However, defining VLANs and routable VLAN interfaces is very different, and will be discussed first.

VLANs (other than VLAN 1) must be explicitly created in VLAN configuration mode, which is different from *config* mode and is reached from enabled mode on the CLI by entering *vlan database*.

Once in database mode, VLANs are created with the *VLAN <id>* command and can be given a name with the *VLAN name <id> <text name>* command.

- Campus NOS also supports dynamic creation of VLANs learned from a neighboring switch via the GVRP protocol. These can be made permanent with the *makestatic* option.

Routable interfaces for VLANs of the form *interface VLAN <x>* are not created by default for each defined VLAN. Instead, they must be explicitly created in VLAN database mode with this command:

- *VLAN routing <vlan-id> [<routable-interface-number>]*

The created interface can be configured as *interface 0/4/<routable-interface-number>* as well as by the more common *interface VLAN <vlan-id>*.

Once VLANs are created, ports can be assigned to them with the familiar *switchport* commands. Note that ports default to mode '*general*' which can lead to unexpected behaviors and should be avoided. , as follows:

- Access ports (one VLAN only, no tagging): *switchport mode access* and then *switchport access vlan <x>*.
- Trunk ports (multiple VLANs, mostly with tagging): *switchport mode trunk*, *switchport trunk native VLAN (optional)*, and *switchport trunk allowed VLAN* can be used to control tagging and VLAN membership.

There are options to enable multiple VLANs to be received untagged and assigned to VLANs based upon the subnet of the destination IP address; this is usually seen with older models of IP telephone instruments.

6. Link Aggregation configuration

A quick sample config is shown; discussion and additional details follow.

To create a link aggregation (aka port-channel in CNOS):

```
config
int 1/0/1-1/0/2,2/0/1-2/0/2
! or other range/list of ports
addport 0/3/1
! defines interface lag 1; 0/3/1 will always be a valid alias
!
! the below is always required for LACP – not for static LAG
interface lag 1
no port-channel static
lacp admin key 10
! if individual interfaces are configured with lacp actor admin key, it must use the same value
```

Campus NOS allows for the creation of link aggregation groups on a single switch much as other firmware does. We don't support a multi-switch option such as vLAG, vPC, or MC-Lag.

LAG interfaces are identified in the configuration as *LAG <x>* or also as *<0/3/x>*. Up to 64 LAGs can be created in the current firmware.

To create a LAG, the individual ports need to be added to the LAG using the *addport <LAG-ID>* command under the physical interface, where the LAG-ID is one of the identifiers shown in the previous paragraph. Port(s) can be removed from a LAG using the *deleteport* command in the same way.

In a stack, any two or more ports from any member of the stack can be part of the same aggregation. This substantially mitigates the lack of a function such as vLAG in Campus NOS.

One item to note is that LAGs are static by default; they function as if they had been configured with *channel-group <x> mode ON*. To set a LAG to use LACP, the minimum requirement is to use this command under the configuration of the LAG interface:

- *interface lag 15*
- *no port-channel static*
- *lacp admin-key 10 (optional)*

If an admin-key is not configured, one will be generated automatically.

It is also possible to configure individual ports with LACP attributes, including active or passive state, individual state, and LACP admin key. Also, configuring multiple individual ports with the same LACP key is not sufficient to aggregate the ports together; they must be added to a lag interface.

In general, attributes for a *lag* should be configured on the *lag* interface; it is not recommended to configure them on the physical ports except for the *individual* option, which can be used when needed as below:

- *interface 1/0/2*
- *[no] lacp actor admin state individual*

7. Stacking

Campus NOS supports stacking using the 10Gb ports which are on all of the campus switch models. There are four ports on each switch; typically two of them would be used to provide a stack which can survive a link or port failure. The members of the stack would be cabled in a ring topology.

In a stack of campus switches, one switch is the management switch and all of the CLI and browser functions are performed on it. Another switch can be designated as a standby which would take on the role of the master in the event of a master switch failure.

In order to build a stack, the *stacking* command needs to be part of the configuration, and two of the four 10Gb ports need to be configured as stack ports.

Ports are designated as stacking ports or changed back to standard Ethernet ports with this command:

- `stack-port <id> ethernet/stack`

The stack ports are the 10Gb ports and are physically ports 25-28 on a -0128 switch model or 49-52 on a -0152 model.

By default, all of the 10Gb ports are configured as stack ports; it is necessary to manually change them to Ethernet mode if it is desired to use them to uplink to another portion of a customer's network. Changing the mode of the stack ports required a reboot to become effective.

While a switch is a member of a stack, some commands have an operand to identify which member of the stack is to be effected, such as the following: *reload [unit]*, which by default would reload every switch in the stack at the same time.

Port names in a stack use the unit number as the left-most part of the interface identifier; physical interfaces are of the form <unit>/0/<port>. VLAN interfaces, routable interfaces, and LAGs are unchanged and apply across all of the members of the stack. Physical ports from any stack member can be in a LAG group, assigned to VLAN membership, etc. (Note that a single switch always has ports numbered 1/0/<port>).

8. Layer-3 addressing and routing protocol configuration

An IP address and similar parameters (RIP, OSPF, VRRP, etc.) can be applied to a routable interface; this includes individual physical ports as well as VLAN interfaces. LAG interfaces can not be configured for these functions.

IP addresses are configured similarly to other familiar firmware, except that two formats are supported:

- `ip address <x.x.x.x> <m.m.m.m>` for the older style subnet mask, and
- `ip address <x.x.x.x> /<prefix>` where `/24` and similar options must be preceded by a space.

IP routing must be explicitly configured with the `ip routing` command; by default it is disabled. Similarly, the `routing` keyword must be part of the configuration for a physical port; it is the equivalent of the `no switchport` command on CNOS.

RIP and OSPF can be configured with routing disabled but will not function.

Routable interfaces can be configured to use RIP and/or OSPF as follows:

- `interface <1/0/x> or VLAN <x>`
- `routing`
- `ip ospf area <area>`
- `ip rip [receive version <x>] [send version <x>]`

Additional options are available for both RIP and OSPF. Note that BGP is not supported.

9. Layer-2 redundancy configuration (failover)

This feature is equivalent to similar *failover* features in CNOS, ENOS, and elsewhere. All of these features allow the monitoring of port(s) or aggregation(s), typically those that connect to an upstream network, and automatically disabling designated downstream or server-facing ports if the upstream ports fail.

The use case for this is to support a server with two paths to the upstream network, so that if there is a failure along one path, it is disabled, and a NIC teaming/bonding driver will forward all outbound traffic from the server along the presumably surviving path.

Configuration of this feature in Campus NOS involved creation of “*link state groups*”.

The groups are created with this command: *link state group <number> <action>*, where the *action* is either up or down. Almost all of the time, the desired action will be *down*, which will mirror similar functionality in ENOS, CNOS, and elsewhere.

Individual interfaces are added to a link state group with this command as part of the interface configuration:

link state group <number> [upstream/downstream], where downstream ports are the ones brought up or down, and the upstream ports are the ones whose state is monitored.

10. Layer-3 redundancy configuration: VRRP

VRRP is an IETF standard protocol which is used to allow seamless and rapid failover of routed traffic in the event of a link or device failure. It implements a shared IP address and shared MAC address which can be used as the next hop in a routed network. At least two switches are enabled to share these addresses, but only one – the current *master* – normally carries traffic. In the event the *master* fails, one of the *standby* switches will take over as *master*. All of the switches typically have a configured priority which determines which will be master and which will take on the role of master should the current master fail.

VRRP switches send out periodic messages which include the identifying number of the VRRP instance as well as their current priority. Failures are detected when these messages cease to be received by a standby switch, at which point an election process takes place and the surviving switch with the highest priority becomes the master.

The configuration as shown below needs to be applied on each of the two (or more) switches, but the priority values should designate which is the preferred master switch. If a server is connected downstream of the switches and the teaming/bonding driver designates a preferred NIC port, then the switch connected to that port should have higher priority.

Note that only VLAN and physical port configurations can be used; LAG interfaces can not be given an IP address in the current Campus NOS firmware. If necessary, select an unused VLAN number and assign the members of the LAG (and no other ports) to that VLAN.

A sample configuration for VRRP v2 follows. A different set of commands, beginning with *vrrp* is available for VRRP v3, which can support IPv6 as well as IPv4. Only one version of the VRRP protocol can be active at any given time. There are also distinct commands to display *vrrp*: v2 uses *show ip vrrp* and its operands; v3 uses *show vrrp* and operands.

Commands to configure *vrrp* instance 1 for v2 are below. IP addresses are included in this sample but would need to be changed to match a customer environment.

```
interface vlan 10
!alternatively for a single physical port
!interface 1/0/20 for port 20
routing
ip address 192.168.99.2 /24
ip vrrp 9 (must match on all switches using the shared address)
ip vrrp 9 mode (enable the VRID number)
ip vrrp 9 accept-mode (allows incoming traffic to a non-master port; all outgoing traffic will still use the master)
ip vrrp 9 priority 101 (100 is the default; master should have at least 101 – others can have lower prio)
ip vrrp 9 preempt (allows priority changes to cause a change in which switch is master)
ip vrrp 9 ip 192.168.99.1 (same subnet as interface address)
no shutdown (if required for the interface)
exit
```


The corresponding VRRP v3 configuration under the same interface is:

```
[no] fhrp version vrrp v3 (changed between v2 and v3)
vrrp 1 address-family ipv4 (or v6)
    accept-mode
    preempt
    priority 101
    address 192.168.99.1
    exit
```

11. IPv6

Campus NOS has rich support for IPv6, and it can be used in many of the applications shown above in place of IPv4. However, there is not support for “4 in 6” or “6 in 4” tunneling protocols nor for forwarding traffic from a v4 network to a v6 network.

Support for IPv6 includes the following:

- Interfaces can have v6 addresses either alone or along with v4 addresses. The interface must be configured with the *ipv6 enable* command before an address can be set.
- OSPFv3 for IPv6 is supported; it is configured much the same as OSPFv2 but commands are prefixed by IPv6 rather than IP. For example:
 - o *ipv6 router ospf* – begins configuration for OSPFv3
 - o *ipv6 router ospf area <x>* - as part of interface configuration for each interface with a v6 address
 - o The counterpart of the “*ip routing*” command is “*ipv6 unicast-routing*” and is required to enable routing and the routing protocols.
- RIPng for IPv6 is not supported.

V6 addresses can be assigned to interfaces (physical ports, VLANs, loopbacks) using similar syntax as v4 addresses, as follows:

- *ipv6 enable* – required on each interface
- *ipv6 address <prefix/length>* - set the address

12. SNMP

Support is available for all versions of SNMP; this document will show common commands for v1 and v2c. SNMP v3 commands are far more complex but are available.

Commands on the switch are available to configure various aspects of SNMP functionality.

a. *snmp-server commands*

The server address, community string, and source port for traps can be configured with the following commands.

- *snmp-server community <string> [ro|rw|su] [ipaddress <ip-address> [ipmask <mask>]]*
This establishes a community string, its level of access (su = superuser), and optionally the ip address or range of addresses for which it is valid. Typically one or more read-only and one or more read-write strings would be configured. No community strings are defaulted.
 - *snmp-server enable traps* – enables traps to be sent to a server. Additional keywords enable or disable various categories of traps if desired.
 - *snmp-server host <ip-address> [traps version [1|2] <community-string>]* – configures an SNMP server to be a trap receiver with the given community string.
 - *snmp-server sysname "name"* – this command is generated automatically if the *hostname* command is configured.
 - *snmp-server enable traps violation* – enables port security traps on a port; entered in interface configuration mode.
- b. *snmp-trap link-status [all]* – can be used globally or per interface to enable link traps.
- c. *snmptrap source-interface [serviceport|vlan<x>|physical port|loopback]* – designates source for traps. A routable interface must be selected. Note different syntax between this command and the one above.

13. Miscellaneous commands

13.1 Spanning tree

Campus NOS supports several STP standards including some which are not commonly in use today, such as original STP (802.1D) and rapid STP or RSTP (802.1w). Original Cisco PVST – which was replaced by PVRST – is also supported.

The STP protocols that are commonly used today are supported: MSTP (802.1s) and Cisco's Rapid-PVST (PVRST).

The default protocol for STP on Campus NOS is RSTP, which is not VLAN-aware. PVRST or MSTP are strongly recommended.

The command to set the spanning-tree protocol is:

```
spanning-tree mode [mst|rapid-pvst|pvst|stp|rstp]
```

The first two options are the recommended protocols.

MSTP requires additional configuration to set a region name, revision number, and to assign VLANs to the MSTP instances, and the commands to do so are similar to CNOS and other firmware. Many options are under the *spanning-tree configuration* command. Rapid-PVST (PVRST) does not require this additional configuration.

13.2 LLDP

LLDP is supported by Campus NOS but is not enabled by default, and a set of commands need to be entered to cause the switch to send and receive LLDP similarly to what is done by other devices. These commands are listed below, and are applied at the individual interface level:

```
lldp transmit – sends a minimum set of data
```

```
lldp transmit-tlv port-desc sys-cap sys-desc sys-name – transmits the listed information
```

```
lldp transmit-mgmt – sends information on the management IP address of the switch
```

Note that for aggregated links, the LLDP commands can only be applied to individual members of an aggregation, not to the *lag* interface itself.