

Dist ZTP

ZTP for CORE and DIST type devices is supported from NMS v1.2

The first switch might always have to be configured manually, unless some other network element can be configured to forward DHCP requests.

Assuming there is a "parent" switch already configured and managed by CNaaS, a second switch could be initialized using ZTP using this process:

1. Parent switch has some interfaces configured as "fabric" interfaces, but no linknets are specified on the interface (using linknets API calls)
In this case, the "fabric" interface is configured as access vlan 1, meaning untagged packets are sent to the ZTP DHCP server
2. The new device goes through DHCP_BOOT and DISCOVERED states the same as access switches
3. The administrator has to check in a new device in the settings repository with a new hostname and interfaces.yml configuration, or use a general interfaces.yml config for this device model
4. The device_init API-call is extended to allow for ZTP of type DIST and CORE, in addition to hostname this API-call can also take a list of expected neighbors as an argument
 - a) device_init will check LLDP neighbors to see which interfaces are connected to other fabric devices and make sure both ends are configured as ifclass fabric
 - b) device_init will check that the list of neighbors hostnames exists and are of the correct device_type (neighbors to a DIST device should be CORE, and neighbor to CORE should be DIST)
 - c) device_init will check that all peers are synchronized and config hash check passes
 - d) device_init will create new linknets in database using interface information gathered from LLDP and assign IPv4 linknets from block configured in settings
 - e) device_init will push new configuration to the new device, at this point contact is lost to the device, state changes to INIT
 - f) device_init will run syncto on peer devices so they re-configure their "fabric" interfaces and apply the new linknet configuration instead of ZTP config, and add BGP-peers etc
 - g) device_init will check that the new device is now reachable via loopback, and change state to MANAGED if successful

Special cases:

New device is not directly connected so LLDP information is not seen, in this case peer hostnames list can be set as an empty list and all interfaces and BGP peers have to be configured manually? Also manual reconfig of peer devices might be needed at step 4f

Caveats

It might not be able to fully ZTP devices that should go in the "evpn_peers" list (CORE devices), since it's probably not possible to populate "evpn_peers" list before the devices exist, and routing connectivity to the management loopback will not work until after the evpn_peers has been established. Solved by using all CORE type devices as evpn_peers if no peers are explicitly set?

Troubleshooting

If the subtask push_base_management fails (for example because a bad template) the neighbor devices might already have had their linknets reconfigured to fabric linknets. In this case the device that failed ZTP will not be reachable via DHCP IP nor MGMT IP until the device is deleted from the database and the neighbor devices are re-synchronized so the ZTP device can reacquire a DHCP lease.