

# Network Cloud Family Blueprints Installation Overview

PLEASE REFER TO R1 NETWORK CLOUD RELEASE DOCUMENTATION

[NC Family Documentation - Release 1](#)

THIS DOCUMENTATION WILL BE ARCHIVED

## Contents

- [Prerequisites](#)
- [Steps](#)
- [Deployment Components](#)
- [High Level Requirements](#)
  - [Compute Node Details](#)
    - [SR-IOV](#)
    - [BDF6 Addresses](#)
  - [Network](#)
  - [Storage](#)
  - [Redfish](#)

## Prerequisites

1. Internal and external network connectivity on all target hardware.

## Steps

1. Ensure all [High Level Requirements](#) are met.
2. Clone and download repositories and packages for the appropriate Akraino release. (Linux Foundation credentials required.)
  - a. [Akraino Gerrit](#): From the [list of projects](#), clone all relevant repositories.
  - b. [Akraino Nexus 3](#): Download all relevant packages.
3. Install the [Regional Controller Node](#):
  - a. Bootstrap the bare metal regional server node from the central node.
  - b. Run installation scripts to launch the Portal, Camunda Workflow, and Database components.
4. Login to the Akraino Portal UI.
5. Install the [Edge Node](#) via the Portal UI:
  - a. Complete the appropriate YAML template according to site requirements:
    - i. Site name
    - ii. Username and ssh key(s) for node access
    - iii. [Server names and hardware details](#)
    - iv. PXE, Storage, Public, and IPMI/iDrac [network details](#)
    - v. [SR-IOV](#) interface details, including the number of virtual functions and [BDF6 addresses](#)
    - vi. Ceph [storage](#) configuration
  - b. Choose the site to build, choose the required Blueprint, and select Build.
  - c. Upon successful build, select Deploy. The following scripts will be run, with status conveyed to the UI:
    - i. 1promgen.sh
    - ii. 2genesis.sh (invokes genesis.sh)
    - iii. 3deploy.sh

## Deployment Components

The following components are deployed in automated sequential fashion:

- [Genesis Host](#)
  - This is the first control node. Genesis serves as the seed node for the control cluster deployed on Edge sites.
  - Genesis contains a standalone Kubernetes instance with undercloud components (e.g., Airship) deployed via Armada.
  - Once the Undercloud is deployed, Ceph is deployed via Armada.
  - Remaining cluster control nodes are deployed next from bare metal, using MaaS. This requires an available PXE network. The Genesis host will provide a MaaS controller.
- [Control Hosts](#)
- [Compute Hosts](#)
- [Airship](#)
- [Apache Traffic Server](#) (VNF)
- [Ceph](#)
- [Calico](#)
- [ONAP](#)
- [OpenStack](#)
- [SR-IOV](#)

# High Level Requirements

Review requirements in the following order:

## Compute Node Details

Herewith are three methods to locate sufficient hardware details:

```
$ sudo dmidecode -s system-manufacturer
HP
$ sudo dmidecode -s system-version
Not Specified
$ sudo dmidecode -s system-product-name
ProLiant DL380 Gen9

$ sudo dmidecode | grep -A3 '^System Information'
System Information
    Manufacturer: HP
    Product Name: ProLiant DL380 Gen9
    Version: Not Specified

$ sudo apt-get install -y inxi
[ ... ]
$ sudo inxi -Fx
System:      Host: mtxnjrsvl24 Kernel: 4.4.0-101-generic x86_64 (64 bit gcc: 5.4.0) Console: tty 10
              Distro: Ubuntu 16.04 xenial
Machine:     Mobo: HP model: ProLiant DL380 Gen9 serial: MXQ604036H Bios: HP v: P89 date: 07/18/2016
CPU(s):      2 Multi core Intel Xeon E5-2680 v3s (-HT-MCP-SMP-) cache: 61440 KB
              flags: (lm nx sse sse2 sse3 sse4_1 sse4_2 ssse3 vmx) bmips: 119857
              clock speeds: [ ... ]
Graphics:    Card: Failed to Detect Video Card!
              Display Server: X.org 1.18.4 drivers: fbdev (unloaded: vesa)
              tty size: 103x37 Advanced Data: N/A for root out of X
Network:     Card-1: Broadcom NetXtreme BCM5719 Gigabit Ethernet PCIe
              driver: tg3 v: 3.137 bus-ID: 02:00.0
              IF: eno1 state: down mac: 14:02:ec:36:52:c4
              [ ... ]
Drives:      HDD Total Size: 1320.2GB (16.2% used)
              ID-1: /dev/sda model: LOGICAL_VOLUME size: 120.0GB temp: 0C
              ID-2: /dev/sdb model: LOGICAL_VOLUME size: 1200.2GB temp: 0C
Partition:   ID-1: / size: 28G used: 17G (66%) fs: ext4 dev: /dev/dm-0
              ID-2: /boot size: 472M used: 155M (35%) fs: ext2 dev: /dev/sdal
              ID-3: /home size: 80G used: 21G (28%) fs: ext4 dev: /dev/dm-2
RAID:        No RAID devices: /proc/mdstat, md_mod kernel module present
Sensors:     System Temperatures: cpu: 48.0C mobo: N/A
              Fan Speeds (in rpm): cpu: N/A
Info:        Processes: 397 Uptime: 39 days Memory: 41943.1/257903.7MB
              Init: systemd runlevel: 5 Gcc sys: 5.4.0 Client: Shell (sudo) inxi: 2.2.35
```

## SR-IOV

Configure and determine the SR-IOV NIC as follows:

```
$ # update /etc/default/grub with this line
$ export GRUB_CMDLINE_LINUX_DEFAULT="intel_iommu=on iommu=pt"
$ sudo -E update-grub
$ sudo reboot now
$ cat /proc/cmdline
$ sudo echo '32' > /sys/class/net/ens3f0/device/sriov_numvfs
$ sudo ip link show ens3f0 # to verify it worked
$ # add line to /etc/rc.local so it does this on reboot
$ sudo echo '32' > /sys/class/net/ens3f0/device/sriov_numvfs
```

## BDF6 Addresses

Intel provides a [script to locate BDF6 addresses from their NICs](#). Learn more about [Bus:Device:Function \(BDF\) Notation](#).

## Network

This Network Cloud blueprint requires:

1. A network that can be PXE booted with appropriate network topology and bonding settings (e.g., a dedicated PXE interface on an untagged /native VLAN)
2. A segmented VLAN with all nodes bearing routes to the following network types:
  - a. Management: Kubernetes (K8s) control channel
  - b. Calico
  - c. Storage
  - d. Overlay
  - e. Public

## Storage

This Network Cloud blueprint requires:

1. Control plane server disks:
  - a. Two disk RAID-1 mirror for the operating system.
  - b. Configure remaining disks as JBOD for Ceph, with Ceph journals preferentially deployed to SSDs where available.
2. Data plane server disks:
  - a. Two disk RAID-1 mirror for the operating system.
  - b. Configure remaining disks per the host profile target for each server (e.g., RAID-6; no Ceph).

## Redfish

This Network Cloud blueprint requires:

1. Configuring BIOS with HTTP boot as a primary device.
2. Adding MAC address of the card to Switch and DHCP server for traffic to flow.
3. Creating the configuration file for pre-seed on DHCP server.
4. Rebooting the server to boot on HTTP device.
5. Getting IP and the related package of OS to install Operating System.