IP addressing

This White Paper is only applicable for Prelude and Nemesis series. For Foundation series, please refer to the Quick Start Guide.

The ISAAC® Platform is built from the ground up to enable users and administrators the flexibility to use and manage their systems without worrying about physical access.

ISAAC® is logically separated into 3 operational levels: the physical hardware, a virtualization solution, and the virtual machines that run the system's workload. Each level has at least one point where it is connected to a network for access and management.

What's on the network?

Unlike most traditional audiovisual equipment or appliances, the virtualization features of the ISAAC® Platform allow some of these components to dynamically move and balance across physical connections and hardware. Where that more traditional approach requires a one-to-one mapping of physical connections (ethernet patch points) to logical connections (IP addresses) the ISAAC® system is more flexible to changing demands in the virtualized components.

Level 1 - Hardware

All ISAAC® hardware is configured, managed, and monitored over the network.

Servers

ISAAC® servers have a discrete baseboard management controller (BMC) that can interact with and monitor the main system regardless of its current running state. This controller has a dedicated physical connection on all current ISAAC® servers. From the BMC's webpage you can power the server on and off, access the built-in network KVM for troubleshooting, and change hardware configuration.

Nemesis Quorum Device

The Quorum Appliance that ships as part of an ISAAC® Nemesis system is only used to form a quorum, so a vote can be done in case of one of the servers fails. The Quorum Appliance is connected to each host using dedicated physical network connections, using preconfigured IP addresses, and does not have any user interface.

Note: The Quorum Appliance is only considered during failover and unneeded in normal operation. If it is not operational during a failure event, automatic failover recovery will not occur.

Between the physical hardware and virtual machines running on an ISAAC® server there is a thin Hypervisor

Level 2 - Virtualization

layer. A Hypervisor is a very minimal operating system, like GNU/Linux or Windows, but is purpose-built to run and manage virtual machines. Each server runs a Hypervisor which has a logical network interface serving management web pages and a control API. In an ISAAC® system direct access to this level is rare, usually reserved for troubleshooting and configuration. Most control and management of the system will be done through the ISAAC Workspace, which in turn interacts with the virtualization layer on the user's behalf.

This is the first level at which logical connections may start to differ from physical connections. Every ISAAC® server has multiple general-purpose network ports that can be configured to conform with the greater network. The most common configuration is as a pool[1], where all ports are configured identically with any VLANs[2] that may be required, and any logical connections are dynamically assigned to some combination of physical connections.

In a Nemesis system there are additional physical interfaces and logical interfaces per server. These are dedicated links for High Availability features, and links to the Quorum Appliance, which do not interact with the outside network. The two hosts are directly connected to each other, and directly connected to the Quorum Appliance. These are assigned by default as shipped and don't need addresses assigned.

All the real work and user interaction with an ISAAC®

Level 3 - Virtual Machines

system is done in Virtual Machines (VM). These are the self-contained operating systems and applications used to control and manage your environment. All ISAAC® systems will have at least a VM that hosts the ISAAC® Workspace itself and may be configured with additional VMs for specific tasks (show or lighting control, programming tools, etc.).

As fully virtual parts of the system, the logical connections for these VMs will dynamically assign and migrate to the subset of physical ports available depending on which hardware they are running on, and balancing decisions made by the virtualization system. AV system designs using the ISAAC® Platform cannot assume that any VM will exclusively use any single host or physical network connection as that would prevent the redundancy and high availability the platform is designed for[3].

Network Default Configuration

ISAAC® systems (Prelude and Nemesis) come preconfigured in DHCP by default. The virtual switch[4] used for general network is labeled **vmbr0**. In some models such as Prelude 232, Prelude 432 and Nemesis 632, an additional high speed **vmbr3** can be present.

These two virtual switches can be changed from DHCP to static IP if required for integration in the greater network.

ISAAC® Nemesis systems also come pre-configured with two virtual switches **vmbr1** and **vmbr2** for Hosts and Quorum Appliance interlink communications. These links are not connected to the external world and **must not be changed**.

Switching from DHCP to static IP

This operation can be done via Proxmox UI.

On Nemesis systems, this operation must be done on both Hosts.

- Login into Proxmox UI[5]
- Select the Host > System > Network tab (screenshot of a Nemesis).

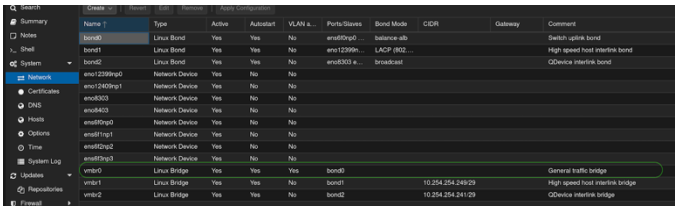

- Select **vmbr0** (Linux Bridge) and click Edit.

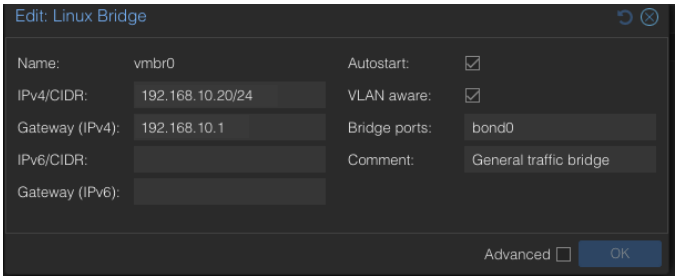

  - Enter the IPv4/CIDR, ex 192.168.10.20/24
  - Enter the Gateway (IPv4) address, ex: 192.168.10.1
  - Click OK.
- Changes must be applied, click on Apply Configuration.

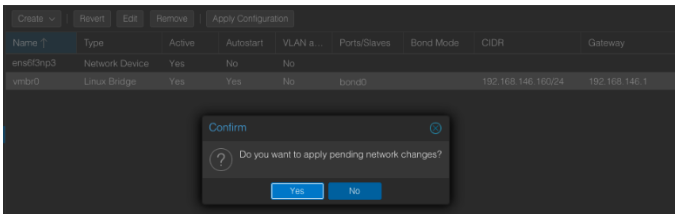

- The following dialog will show

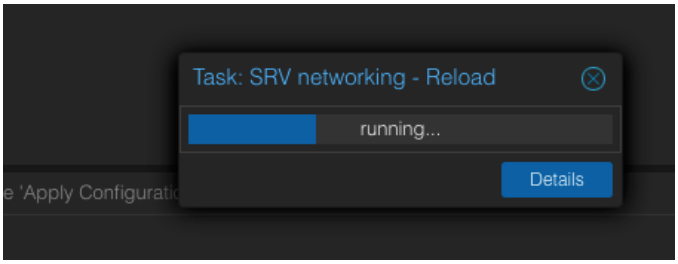

- At this point, the UI won't be responsive anymore, since the host will have changed its IP address. Use the new assigned IP address to login back into Proxmox UI.

The same procedure can be applied for **vmbr3** when present.

Switching from Static IP to DHCP

This operation can only be done via the host console.

On Nemesis systems, this operation must be done on both Hosts.

This process is at risk, since it doesn't have any guard rails, so only proceed if you're comfortable with Linux network settings.

- Login into the host console.
- Edit the **/etc/network/interfaces** file (nano is available to do so)
- Locate **vmbr0** virtual switch section

```
auto vmbr0
iface vmbr0 inet static
#General traffic bridge
    address 192.168.146.161/24
    gateway 192.168.146.1
    bridge-ports bond0
    bridge-stp off
    bridge-vlan-aware yes
    bridge-vids 2-4094
```

- Change **static** for **dhcp** and remove **address** and **gateway** lines

```
auto vmbr0
iface vmbr0 inet dhcp
#General traffic bridge
    bridge-ports bond0
    bridge-stp off
    bridge-vlan-aware yes
    bridge-vids 2-4094
```

- Save the changes.
- Reboot the Host.

The same procedure can be applied for **vmbr3** when present.

Switching from DHCP to static IP (alternative)

Although this operation can be done via Proxmox UI (this is the recommended way), it can also be done via the host console.

This process is at risk, since it doesn't have any guard rails, so only proceed if you're comfortable with Linux network settings.

- Login into the host console.
- Edit the **/etc/network/interfaces** file (nano is available to do so).
- Locate **vmbr0** virtual switch section.

```
auto vmbr0
iface vmbr0 inet dhcp
#General traffic bridge
    bridge-ports bond0
    bridge-stp off
    bridge-vlan-aware yes
    bridge-vids 2-4094
```

- Change **dhcp** for **static** and add **address** and **gateway** lines.

```
auto vmbr0
iface vmbr0 inet static
#General traffic bridge
    address 192.168.146.161/24
    gateway 192.168.146.1
    bridge-ports bond0
    bridge-stp off
    bridge-vlan-aware yes
    bridge-vids 2-4094
```

    - address uses CIDR notation.
- Save the changes.
- Reboot the Host.

The same procedure can be applied for **vmbr3** when present.

Summary

In general, to provide the remote management and robustness required of an ISAAC® system, each host will present multiple logical and physical network connections. Logical connections may be distributed across physical connections to maximize availability and will dynamically shift to balance workloads.

When designing an ISAAC® system into a network, keep the following addressing requirements in mind:

Prelude Systems (3+ IPs)

- 1 for hardware BMC[6]
- 1 for virtualization hypervisor[6]
- 1 for ISAAC® VM (Workspace)

Depending on specific configuration and usage:

- 1 or more for any additional VMs

Nemesis Systems[7] (5+ IPs)

Each system will require IP addresses for:

- 2 for hardware BMC (1 per server)[6]
- 2 for virtualization hypervisor (1 per server)[6]
- 1 for ISAAC® VM (Workspace)

Depending on specific configuration and usage:

- 1 or more for any additional VMs

---

[1] *The hypervisor balances across connections based on MAC address. Switch connections do not need to be configured for LACP/LAG bonding.*

[2] *ISAAC® Gen5 only supports VLAN untagged traffic on these ports.*

[3] *This does not require any special support by other networking equipment and will be compatible with all modern infrastructure.*

[4] *Virtual switches in Proxmox are technically Linux Bridges.*

[5] *The user must have at least the Sys.Modify permission, which is by default present in Administrator and PVEAdmin default Promox roles, to be able to proceed.*

[6] *In an environment with a dedicated management network, we suggest designating these items as management.*

[7] *Nemesis 900 systems may have different requirements.*