

Network architecture

Choosing a network card

Depending on your needs, the network card will have to support several features

- PTP
- PCI-passthrough
- SR-IOV

None of these features are required by SEAPATH, however, PTP is required by the IEC 61850 standard.

If your machine doesn't have enough interfaces, you can use an usb-ethernet interface. For example, the TRENDnet TU3-ETG was tested and does work natively on SEAPATH. These added interfaces will not be compatible with any of the features described above.

SEAPATH network on one hypervisor

The following description fits the case of a standalone hypervisor, but is also applicable to a machine in the cluster.

Physical interfaces

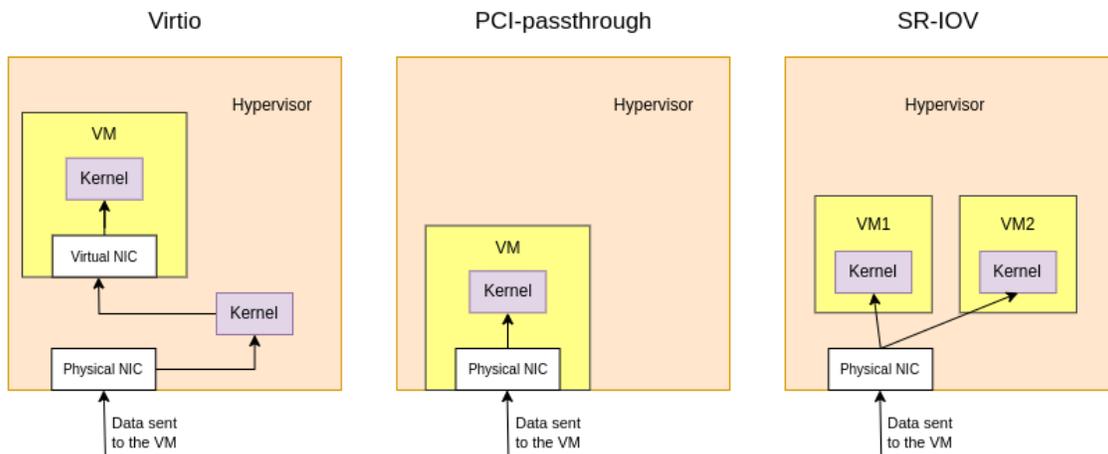
There are four different types of interfaces to consider on a SEAPATH hypervisor:

- Administration: This interface is used to access the administration user with ssh and configure the hypervisor. It is also used by ansible.
- VM administration: This interface has the same role of the previous one, but for the virtual machines.
- PTP interface: This interface is used to transmit PTP frames and synchronize the time on the machine.
- IEC61850 interface: This interface will receive the sample values and GOOSE messages.

VM interfaces

There are three different ways to transmit data from a physical network interface to a virtual machine:

- virtio: A virtual NIC (Network Interface Card) is created inside the virtual machine. The data received on the hypervisor will be treated by the physical network card and then pass to the VM virtual card. This is the simplest method, but also the slowest.
- PCI passthrough: The control of the entire physical network card is given to the VM. The received data is then directly treated by the virtual machine and does not go through the hypervisor. This method is much faster, but has two drawbacks:
 - The entire physical NIC is given to one VM. Other VMs or the hypervisor can't use it anymore
 - It requires a compatible NIC
- SR-IOV: Single Root I/O Virtualization is a PCI Express Extended capability which makes one physical device appear as multiple virtual devices (more details [here](#)). Multiple VMs can now use the same physical NIC. This is the fastest method, but it requires a compatible NIC.



Recommendations

Many network configurations are possible using SEAPATH. Here are our recommendations:

Administration

To avoid using too many interfaces, the administration of all VMs can use the same physical NIC. Hypervisor administration can also be done on this interface. This is possible by using a Linux or OVS bridge and connecting all VMs and the hypervisor to it.

i This behavior is achieved by the br0 bridge, preconfigured by Ansible (see Ansible inventories examples)

IEC 61-850 traffic

The Sample values and GOOSE messages should be received on an interface using PCI-passthrough or SR-IOV, the classic virtio interface is not fast enough. In this situation, all VMs receiving IEC 61-850 data must have one dedicated interface.

For testing purposes, to avoid using one interface per VM, these data can be received on classic interfaces using virtio. This is done by using a Linux or OVS bridge connected both to the physical NIC and to all the virtual machines.

i See the variables `bridges` and `ovs_bridges` in Ansible inventories examples.

PTP

PTP synchronization does not generate much data. The synchronization of the hypervisor can thus be used:

- Alone on a specific interface
- On the administration interface

Please remember that PTP requires specific NIC support.

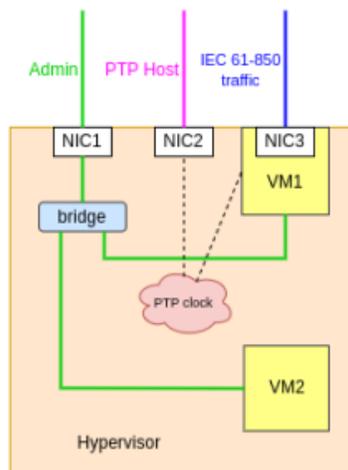
In order to synchronize the virtual machines, the PTP clock of the host can be used (ptp_kvm). It is also possible to pass the PTP clock on the PCI-passthrough/SR-IOV interface given to the VM.

See the [Time synchronization](#) page for more details.

Example of configuration

Below is an example of a SEAPATH machine with two VMs. VM1 receives SV/GOOSE and must thus be synchronized with PTP and use PCI-passthrough. VM2 is just for monitoring, and does not require either PTP or PCI-passthrough.

- One interface is used for both VM and hypervisor administration
- One interface to synchronize the host and VM1 in PTP
- One interface for PCI-passthrough in VM1 handling IEC 61-850



Connecting machines in a cluster

SEAPATH cluster is used to ensure redundancy of the machines (if one hypervisor fails) and of the network (if one link fails).

In order to connect all the machines together, two options are available:

- Connecting all machines to two switches (the second is needed to ensure redundancy if the first switch fails)
- Connecting all machines in a ring architecture

The second method is recommended on SEAPATH because it doesn't use an additional device (the switch). In that case, every machine is connected to its neighbors, forming a triangle. If one link breaks, the packets still have another route to reach the targeted machine. See on [Github](#) for more information.

In that case, two more interfaces per machines are needed, which leads to a minimum of four interfaces per machine (two for the cluster, one for the administration, and one for IEC 61-850 traffic)

Observer

An observer machine doesn't run any virtual machines. It is only used in the cluster to discriminate if a problem occurs between the two hypervisors.

On this machine, no IEC 61-850 packets will be received. It can be synchronized in PTP but it is not needed; it can stick to the simple NTP synchronization.



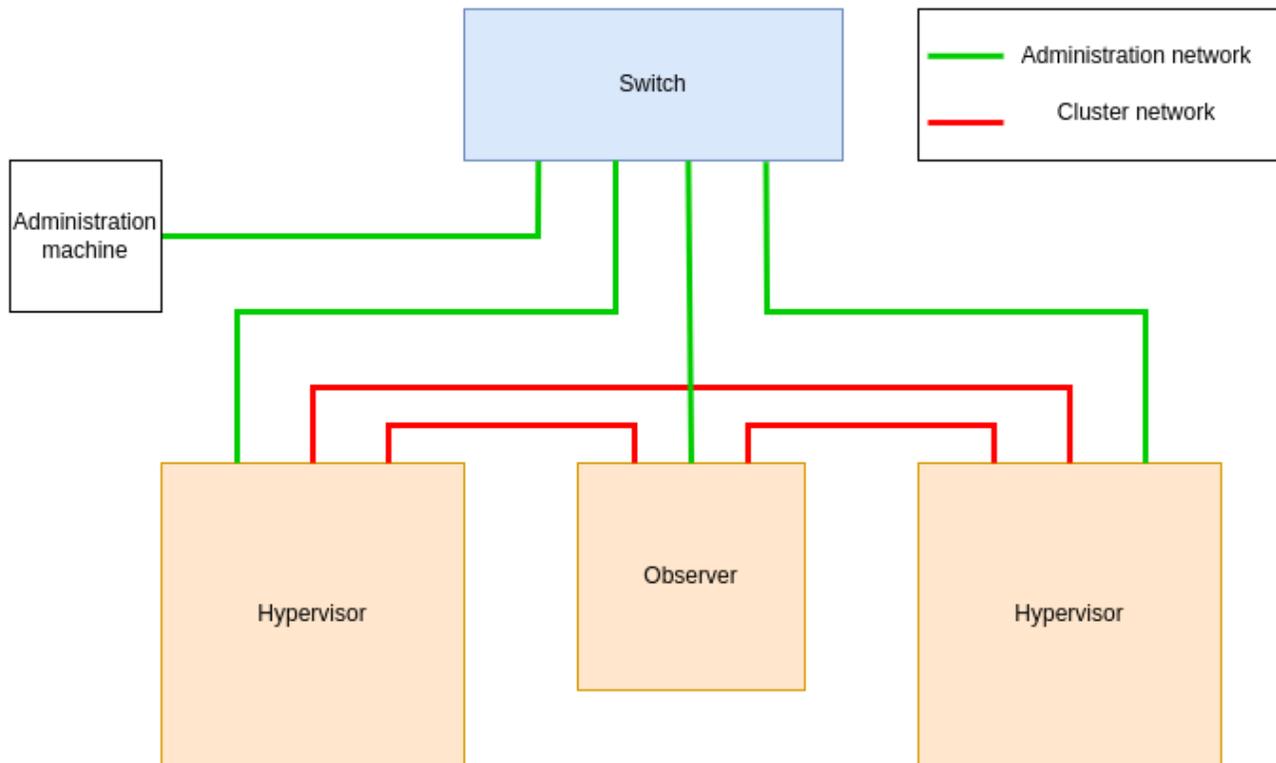
Even if PTP is not required, time synchronization must be ensured, at least in NTP. Otherwise, the machines will not be able to form the SEAPATH cluster.

So, on an observer, only three interfaces are needed : two for the cluster and one for administration.

Administration machine

In order to configure all the machines in the cluster, we advise using a switch connected to the administration machine.

Below is an example with two hypervisors and one observer, all connected together in a triangle cluster. The administration machine is connected to them using a switch.



Receiving Sample values and PTP frames

Machines

The standard way to generate sample values is to use a merging unit. For PTP, it is a Grand master clock.

However, in a test environment, these two machines can be simulated:

- See this page to [simulate IEC 61-850 traffic](#)
- See this page to [simulate a grand master clock](#)

For this section example, we will use only one machine to simulate both PTP frames and IEC 61-850 traffic. This machine will be named "Publisher".

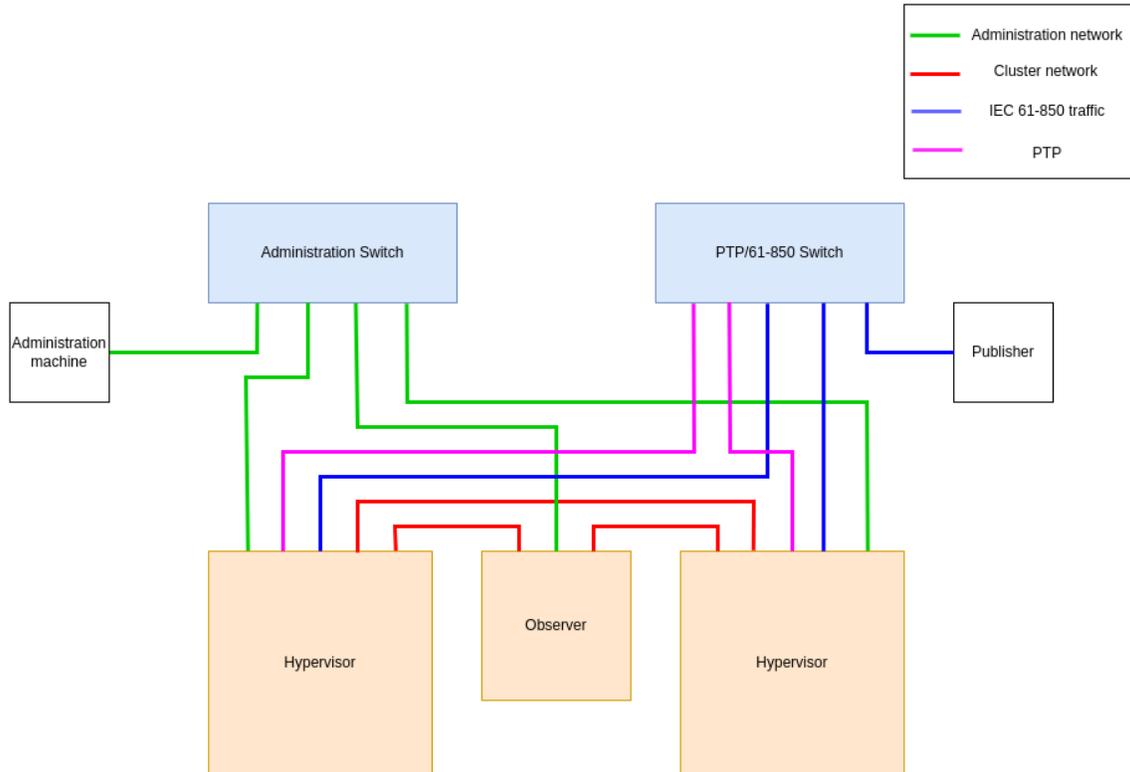
Connections

In order to connect the publisher (or the merging unit/Grand master clock), a direct connection, with only one cable is the best. It avoids unnecessary latency, but necessitates having a machine with many interfaces.

A switch can then be used between the publisher and the cluster machine. In this situation, we advise you to use a separate switch from the one used for administration. It is possible to use only one switch for both, but in that case, the two networks must be separated using VLAN.

Note: To manage PTP, the best is to use a PTP-compliant switch. See more information on the [Time synchronization](#) page

Below is a schema of a cluster with two hypervisors and an observer. The administration network is isolated on his switch. Another switch is used for PTP and IEC 61-850 traffic. Remember that if you use PCI-passthrough, you will need as many interfaces as you have Virtual machines.



i To avoid PTP frames interrupting with IEC 61-850 traffic, the best is to isolate PTP frames on a VLAN. This is done with the variable `'ptp_vlanid'` in the ansible inventories.