

Configure Pair Bonding and Bridges for KVM Hypervisor

Configure Pair Bonding and Bridges for KVM Hypervisor

TUSK is normally installed on KVM hypervisors. These hypervisors should be configured with pair bonding in case one port fails, and KVM bridges are absolutely necessary for KVM guests to have IP addresses exposed for their web servers. Unfortunately, no Red Hat published GUI currently configures these, so it must be done manually.

These settings are for a pair of network ports in a simple failover configuration.

- Allocate the hostname and IP address for the TUSK server in DNS.
 - Always do this first, to avoid DNS conflicts or other surprises.
 - Make sure the reverse DNS matches the forward DNS.
 - If there will be a web proxy or load balancer in front of a TUSK server, allocate the DNS and IP address for that separately.
- Access the server with a remote console or keyboard and monitor.
 - Do not rely on an SSH login to make network changes. It's too easy to cut off access while reconfiguring a network.
- Do not do this on an active, production server, there can be very confusing service interruptions.
- Edit `/etc/modprobe.d/bond0.conf` to contain this. **Do not** set other bonding parameters here.

```
alias bond0 bonding
```

- Edit `/etc/sysconfig/network` to set the hostname and various other useful settings.
 - IPv6 **must, must, must** be enabled for pair bonding and other features to work reliably.
 - `NOZEROCONF=yes` avoids the unnecessary "169.254.*" fallback addresses and routing from being activated, which are purely confusing in normal environments.

```
HOSTNAME=hostname.example.com
# IPv4
NETWORKING=yes
NOZEROCONF=yes
# IPv6, necessary for bonding
NETWORKING_IPV6=yes
IPV6INIT=yes
```

- These notes are based on the network devices being called "eth0", "eth1", etc.
 - Edit as necessary for devices referred to as "em0", "em1", etc.
 - Run `"/sbin/ifconfig -a"` to get the list of all available devices and their MAC addresses.
- Edit `/etc/sysconfig/network-scripts/ifcfg-eth0`.
 - Make sure that HWADDR in this file the actual MAC address of the port.
 - Confirm this with `"/sbin/ifconfig eth0"`.
 - Notice the bonding settings for MASTER and SLAVE.
 - Notice the `NM_CONTROLLED=no`, to block NetworkManager from ever touching this port.
 - Notice the "MTU=9000" setting to enable jumbo frames. This is a common practice in server environments, and should be activated as needed.

```
BOOTPROTO=none
DEVICE=eth0
HWADDR=aa:bb:cc:dd:ee:ff
MASTER=bond0
#MTU=9000
NM_CONTROLLED=no
NOZEROCONF=yes
ONBOOT=yes
SLAVE=yes
TYPE=Ethernet
```

- Edit `/etc/sysconfig/network-scripts/ifcfg-eth1`.
 - Adjust the HWADDR to match the actual MAC address of the port.
 - Run `"/sbin/ifconfig eth1"` to report this.
 - Notice the "DEVICE" name is different than ifcfg-eth0.
 - Notice that all other settings are identical to ifcfg-eth0.

```
BOOTPROTO=none
DEVICE=eth1
HWADDR=aa:bb:cc:dd:ee:ff
MASTER=bond0
# Uncomment for jumbo frames
#MTU=9000
NM_CONTROLLED=no
NOZEROCONF=yes
ONBOOT=yes
SLAVE=yes
TYPE=Ethernet
```

- Edit **/etc/sysconfig/network-scripts/ifcfg-bond0**.

- Notice the BONDING_OPTS, suitable for fail-over bonding and 100 msec failover times.

```
ARP=no
BONDING_OPTS="mode=1 miimon=100"
BOOTPROTO=static
BRIDGE=br0
DEVICE=bond0
# Uncomment for jumbo frames
#MTU=9000
NM_CONTROLLED=no
ONBOOT=yes
```

- Edit **/etc/sysconfig/network-scripts/ifcfg-br0**.

- This selects the appropriate bridge for KVM based virtualization.

```
BOOTPROTO=static
DEVICE=br0
GATEWAY=192.168.1.xxx
IPADDR=192.168.1.yyy
NETMASK=255.255.255.0
NM_CONTROLLED=no
ONBOOT=yes
SLAVE=bond0
TYPE=Bridge
```