Best Practices
VMware VMotion
with
HyperIP

6420 Sycamore Lane N.
Maple Grove, MN 55369
http://www.netex.com

Adding HyperIP into your network

HyperIP improves the performance of backup and replication applications over your IP WAN. **HyperIP does not alter application protocols nor modify any file systems.** It efficiently moves block or file data over the IP WAN under any network conditions.

HyperIP also provides:

- support of WAN speeds scaling from 1-800 Mb/s
- virtual or physical appliance footprint
- adaptive lossless block level compression
- time of day rate controls for changing throughput requirements

HyperIP requires at least two appliances (virtual or physical), one residing on each side of the WAN, as shown in the figure below. Multiple servers and storage at each site can utilize the HyperIP data path. HyperIP can also be deployed in a hub or mesh configuration.

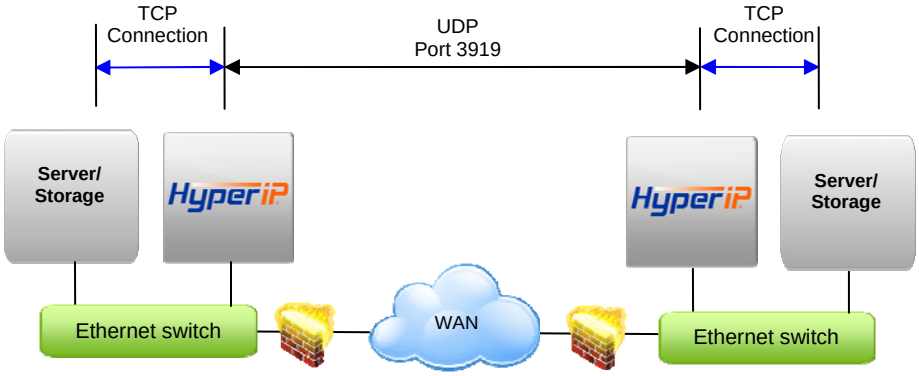

HyperIP terminates TCP connections locally and tunnels the data between HyperIPs using UDP port 3919. **Network devices filtering IP traffic in the data path between the HyperIPs must be configured to allow UDP port 3919.**

HyperIP must be *in* the data path to optimize the movement of data. HyperIP connects to a (virtual) LAN switch with a single Gigabit Ethernet NIC and has two modes of operation to facilitate being inserted into the data path:

- **Gateway Mode: User must add route statements in the *data movers* (application servers, storage devices, etc.)** defining HyperIP as the IP gateway for the destination IP addresses or networks. Alternatively, these IP route statements or redirect filters, may be configured in a router. Gateway mode **requires users to define HyperIP intercepts** based on IP addresses, TCP ports and/or protocols to determine what traffic to act on.

- **Proxy Mode: HyperIP requires additional local IP addresses (proxy) which represent remote IP addresses of the application servers or storage devices.** This local proxy IP address is then used to communicate with the remote application. HyperIP is configured with a 1:1 mapping in which each destination IP address requires an associated local proxy address. *Applications that do not support Network Address Translation (NAT) must use the HyperIP gateway mode.*

Each HyperIP requires its own key associated with the HyperIP fingerprint. You must connect to the user interface on each HyperIP to retrieve its fingerprint and forward this to support@netex.com to request the key.

For further explanation on the feature/functionality of HyperIP see the HyperIP User guide at: http://www.netex.com/support/product-support/hyperip

HyperIP in a VMotion Network

The VMotion Network is leveraged within VMware environments to move Virtual Machines (VMs) from one ESX server to another or to change storage for a VM. When running VMotion across a WAN, TCP protocol is a potential bottleneck and can affect the time to complete a migration. HyperIP shields TCP from WAN issues like latency and packet loss so the migration completes as quickly as possible.

HyperIP is installed on a virtual machine within an ESX server. HyperIP can tunnel traffic from the ESX server it resides on and/or from servers and storage outside its ESX server. Typically, only one HyperIP per site is required.

The two user selectable VMotion options in vSphere 4.0 that will be discussed in this document are:

Change Host – Can be performed for a live (running) or stopped (powered off) VM
Change Both Host and Datastore - Can be performed only for a stopped VM

The local ESX server requires routes to the remote ESX VMkernel and/or Service Console depending on which VMotion change option is selected.

The examples in this guide are valuable for determining information required to deploy HyperIP in the VMotion network. It also includes information required to complete the installation worksheets in the HyperIP HyperStart Guide. The subsequent sections describe the VMotion data flows to better understand how HyperIP is deployed with each type of migration.

VMotion Change Host Migration Option

When performing a *Change Host Migration* traffic flows between ESX VMkernel NICs. The data flow moving a VM from the ESX server "Canada" to the ESX server "Mexico" using HyperIP is shown in the drawing below. The actual IP addresses in your VMotion network should be used to replace the IP addresses in this example.

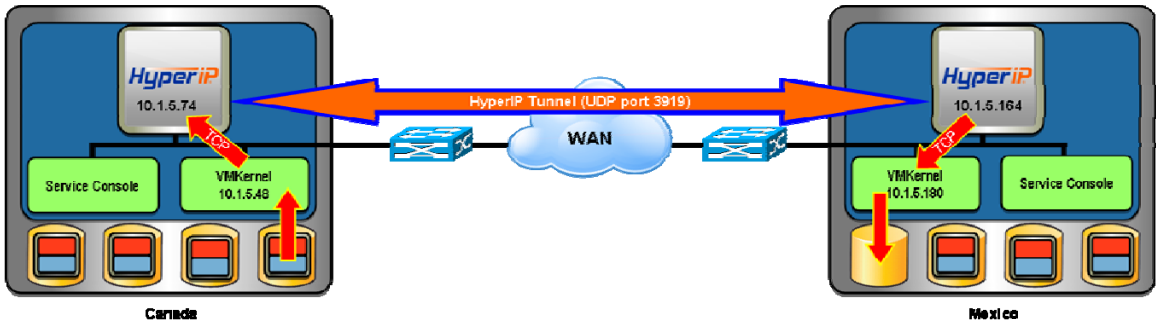

The HyperIPs must be configured to intercept VMotion traffic for these sites (see Appendix A for HyperIP configuration) and the VMkernel routing tables need to be modified to direct the traffic to the HyperIPs. Local route statements are added to the VMkernel using the local HyperIP as the gateway to the remote VMkernel. You must be logged into the ESX server with the authority to make VMkernel routing changes.

HyperIP must be on the same subnet and VLAN as the VMkernel for the gateway statement to be effective. If HyperIP can not be placed on the same network, contact support@netex.com for other options to direct traffic to HyperIP.

In this example use the following IP routing statements in the ESX VMkernel to direct the traffic to HyperIP during a VMotion *Change Host Migration*:

Use the following ESX CLI command to install a route on Mexico's VMkernel to route traffic to Canada's VMkernel using the HyperIP on Mexico as the IP gateway:

```
esxcfg-route –a 10.1.5.48/32 10.1.5.164
```

You should receive the following response to the command above:

```
Adding static route 10.1.5.48/32 to VMkernel
```

Verify the route was added correctly using the "esxcfg-route –l" command.

```
VMkernel Routes:
Network            Netmask              Gateway
10.1.5.0           255.255.255.0        Local Subnet
10.1.5.48          255.255.255.255      10.1.5.164
default            0.0.0.0              10.1.5.50
```

Use the following ESX CLI command to install a route on Canada's VMkernel to route traffic to Mexico's VMkernel using the HyperIP on Canada as the IP gateway:

```
esxcfg-route –a 10.1.5.180/32 10.1.5.74
```

You should receive the following response to the command above:

```
Adding static route 10.1.5.180/32 to VMkernel
```

Verify the route was added correctly using the "esxcfg-route –l" command.

```
VMkernel Routes:
Network            Netmask              Gateway
10.1.5.0           255.255.255.0        Local Subnet
10.1.5.180         255.255.255.255      10.1.5.74
default            0.0.0.0              10.1.5.50
```

VMotion Change Host and Datastore Migration Option

When performing a *Change Host and Datastore* Migration, traffic flows between the ESX servers' Service Console. The data flow for moving a VM from the ESX server "Canada" to the ESX server "Mexico" using HyperIP is shown in the drawing below. The actual IP addresses in your VMotion network should be used to replace the IP addresses in this example.

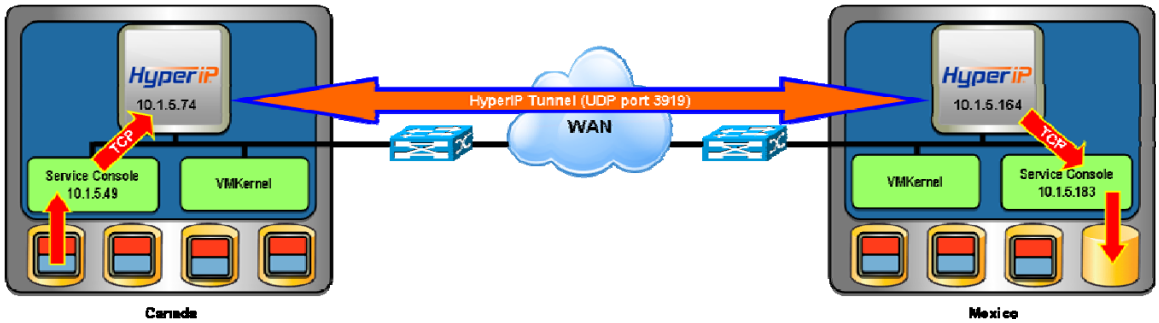

The HyperIPs must be configured to intercept VMotion traffic for these sites (see Appendix A for HyperIP configuration) and the ESX routing tables need to be modified to direct the traffic to the HyperIPs. Local route statements are added to each Service Console using the local HyperIP as the gateway to the other Service Console. You must be logged into the ESX server with the authority to make routing changes.

HyperIP must be on the same subnet and VLAN as the Service Console for the gateway statement to be effective. If HyperIP can not be placed on the same network, contact NetEx support at support@netex.com for other options to direct traffic to HyperIP.

In this example use the following IP routing statements to direct the traffic to HyperIP during a VMotion *Change Host and Datastore Migration*:

Use the following ESX CLI command to install a route on Canada's Service Console to route to Mexico's Service Console using HyperIP as the IP gateway:

```
route add -host 10.1.5.183/32 gw 10.1.5.74
```

Verify the route was added correctly using the "route" command.

```
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
10.1.5.183      10.1.5.74       255.255.255.255 UGH   0      0        0 vswif1
10.1.5.0        0.0.0.0         255.255.255.0   U     0      0        0 vswif1
0.0.0.0         10.1.5.50       0.0.0.0         UG    0      0        0 vswif1
```

Use the following ESX CLI command to install a route on Mexico's Service Console to route to Canada's Service Console using HyperIP as the IP gateway:

```
route add -host 10.1.5.49/32 gw 10.1.5.164
```

Verify the route was added correctly using the "route" command.

```
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
10.1.5.49       10.1.5.164      255.255.255.255 UGH   0      0        0 vswif1
10.1.5.0        0.0.0.0         255.255.255.0   U     0      0        0 vswif1
0.0.0.0         10.1.5.50       0.0.0.0         UG    0      0        0 vswif1
```

VMotion – Both Options of Migration

The most likely scenario is for HyperIP to be configured for both migration options. The complete IP routing required for this scenario is noted below:

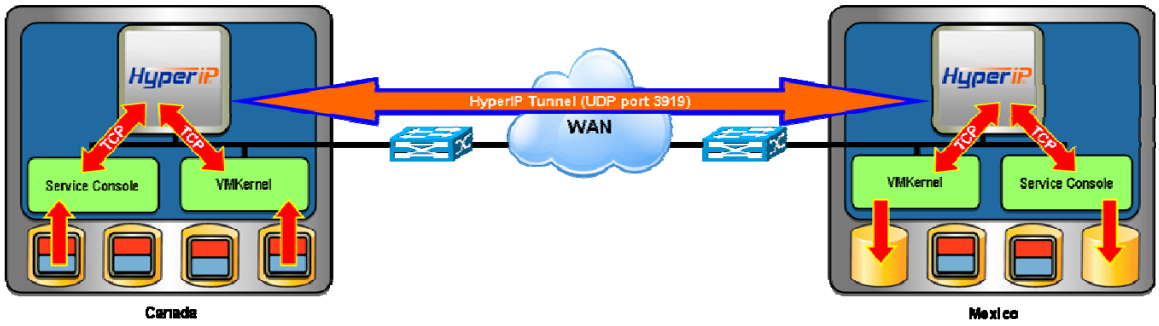

The VMkernel and Service Console routing information is shown below. The commands to modify the routing tables are as described in the previous sections.

Canada's routing tables should look like this to utilize HyperIP for both options:

```
[root@canada ~]# esxcfg-route -l
VMkernel Routes:
Network             Netmask             Gateway
10.1.5.0            255.255.255.0       Local Subnet
10.1.5.180          255.255.255.255     10.1.5.74
10.1.5.190          255.255.255.255     10.1.5.74
default             0.0.0.0             10.1.5.50

[root@canada ~]# route -n
Kernel IP routing table
Destination     Gateway    Genmask         Flags Metric Ref    Use Iface
10.1.5.183      10.1.5.74  255.255.255.255 UGH   0      0        0 vswif1
10.1.5.0        0.0.0.0    255.255.255.0   U     0      0        0 vswif1
0.0.0.0         10.1.5.50  0.0.0.0         UG    0      0        0 vswif1
```

Mexico's routing tables should look like this to utilize HyperIP for both VMotion options:

```
[root@mexico ~]# esxcfg-route -l
VMkernel Routes:
Network             Netmask             Gateway
10.1.5.0            255.255.255.0       Local Subnet
10.1.5.20           255.255.255.255     10.1.5.164
10.1.5.48           255.255.255.255     10.1.5.164
default             0.0.0.0             10.1.5.50

[root@mexico ~]# route -n
Kernel IP routing table
Destination     Gateway    Genmask         Flags Metric Ref    Use Iface
10.1.5.49       10.1.5.164 255.255.255.255 UGH   0      0        0 vswif1
10.1.5.0        0.0.0.0    255.255.255.0   U     0      0        0 vswif1
0.0.0.0         10.1.5.50  0.0.0.0         UG    0      0        0 vswif1
```

Appendix A-HyperIP Configuration

HyperIP may reside on its own internal switch and traffic sent outside the ESX server, through an external switch and back to the HyperIP or HyperIP could reside on the same virtual switch as the VMkernel and/or Service Console NIC. For unrestricted performance, HyperIP should be configured with a dedicated network interface.

The following information is required when configuring HyperIP:

- Interface IP address and network mask.
- HyperIP hostname
- HyperIP default gateway
- HyperIP Domain name
- IP addresses or networks utilizing HyperIP (required to configure intercepts)

Using this information follow the instructions in the HyperStart for HyperIP on VMware guide to configure HyperIP for the VMotion Network. This document is found on the NetEx support website in the docs section for your version of HyperIP:

http://www.netex.com/support/product-support/hyperip

Example screen shots of the configuration webpages are included on the following pages.

This System webpage is used to configure or verify the basic system information and access settings:

Local IP Hostname
Default Gateway
Local Domain
Data and Mgmt IP addresses and Network mask

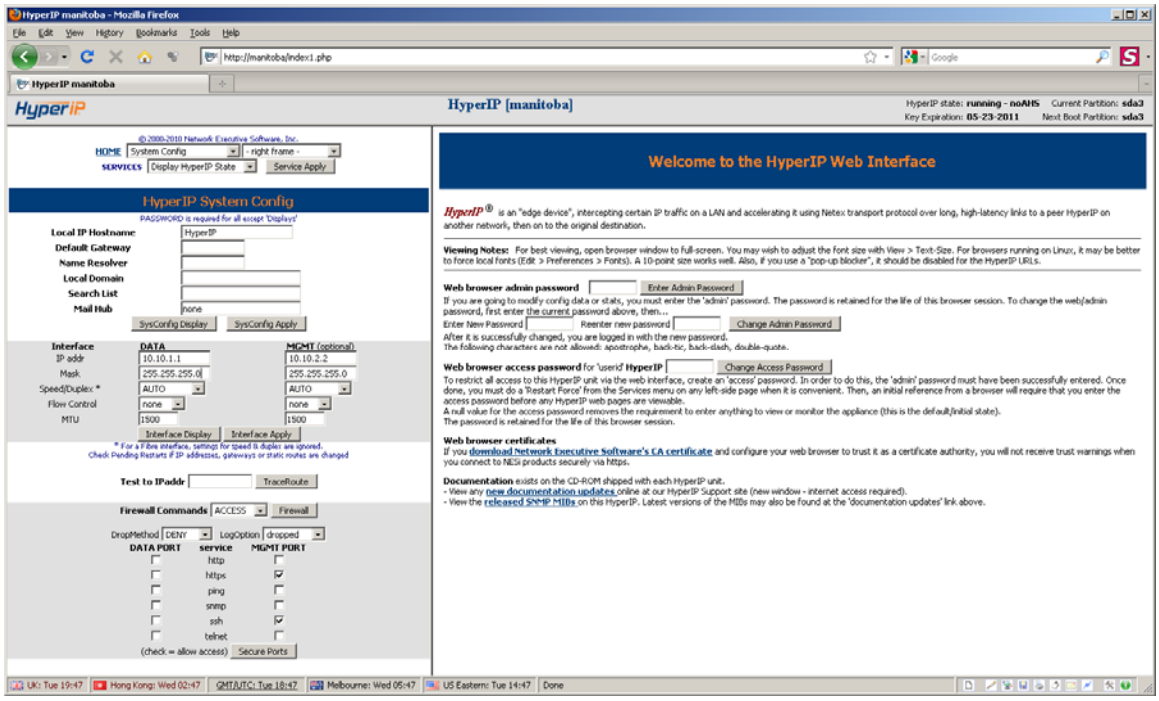

The HyperIP “Configure NxN” frame is launched from the HyperIP Configuration webpage and is used to add information about the sites. The first site definition on each HyperIP must define itself and you must put this site number into the “Configure this unit as site #” window.

Site number/Site name (user-defined and unique within the HyperIP environment)
Primary IPaddr - HyperIP data interface IP address
Segsize (default is 32768 – may modify after setup tests have been run)
MaxRate (required if there are multiple remote HyperIPs configured)

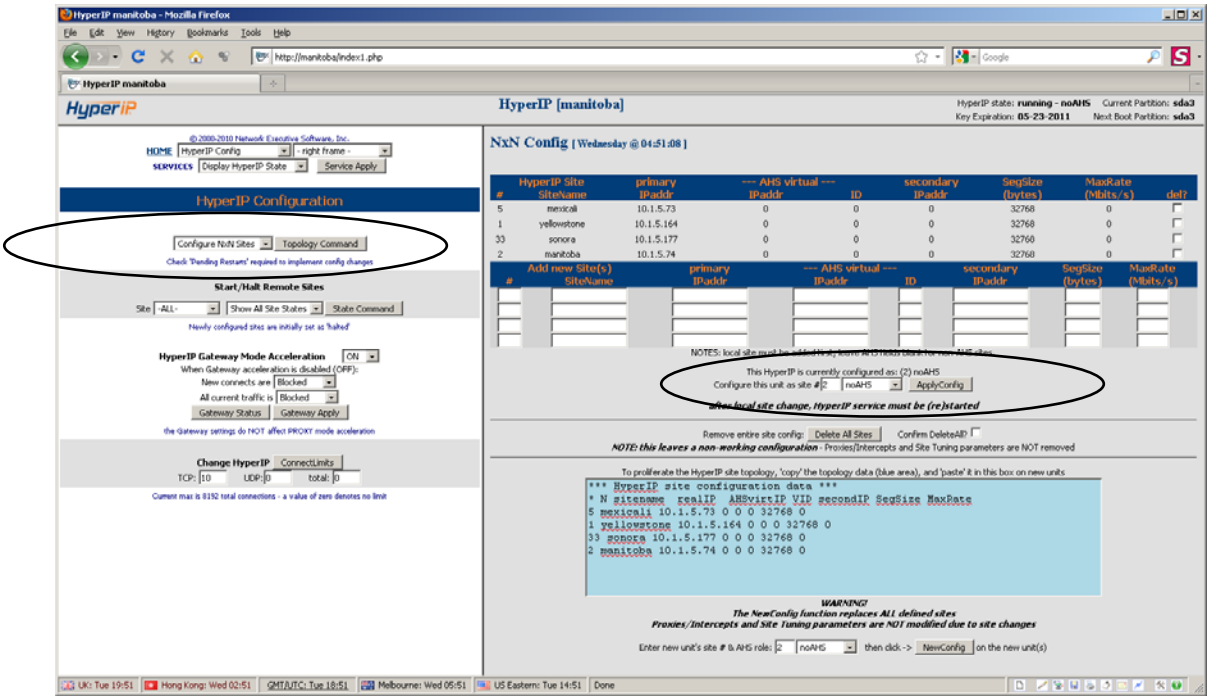

The HyperIP “Proxies & Intercepts” frame is launched from the HyperIP Configuration webpage and is used to configure what traffic destined to the remote site HyperIP will intercept:

IP addresses or networks utilizing HyperIP
Protocols and ports can be used to further qualify intercepted traffic

Note: Intercepts may use an asterisk wildcard. This wildcard can only be defined on a byte boundary. (Using 10.1.5.* will match IP addresses 10.1.5.0-10.1.5.255)