BBNA Configuration Examples

Summary

We explain the configuration and application of BBNA (Broadband Bonding™ Network Appliance) devices through discussion of several examples.

The BBNA devices are highly versatile devices that can provide true bandwidth aggregation in a variety of settings. These include stand-alone operation, point-to-point communication between two sites, as well as supporting robust data communication across a large global enterprise. In all cases, the BBNA devices enable high performance, high reliability data communications by making concurrent use of all available communication resources. By using the most cost effective communication resources available at a site location, significant cost savings over traditional transport means may be achieved, while at the same time significantly increasing performance and reliability.

The BBNA devices come in different form factors. For example, the Truffle has 4 WAN ports and 2 LAN ports (upgradeable to 8 or 12 WAN ports), Truffle Lite has 2 WAN ports and 2 LAN ports, and the PortaBella 141 has 1 LAN port, and 4 USB ports (or 8 USB ports) for wireless WAN connections. The BBNA devices equipped with the VLL feature are interoperable can peer with Bonding Proxy Appliances by Mushroom Networks, as long as they are running the most recent software release of the BBNA system software. The examples discussed herein pertain to all BBNA models.

Stand-Alone Mode versus Peered Mode

A BBNA device can operate in stand-alone mode, or it can peer with BBNA that has a server license through the VLL feature. In order to configure a BBNA in peered mode it is first necessary to configure it in stand-alone mode.

In the stand-alone mode the BBNA aggregates bandwidth from multiple Internet connections in two ways. Data transfers over the HTTP protocol in the downlink direction, even for a single file transfer, can be concurrently spread over all available access lines so that the data transfer speed is close to the sum of the available data throughputs available from each Internet connection. Other traffic besides HTTP downlink traffic is aggregated using load balancing at the granularity of an IP flow.

In the peered mode, a BBNA may be configured to provide a data tunnel, called a VLL tunnel, to a Truffle with server license (which is yet another BBNA device). A VLL tunnel can be configured between two BBNA devices equipped with the VLL feature (one client and one server). Multiple VLL tunnels may also be established within a single BBNA device, so that a Truffle with server license may communicate directly with multiple BBNA devices. For each VLL tunnel, one BBNA device is configured as the "client" and the Truffle with server license is configured as the "server" for that VLL tunnel. An IP subnet is associated with each side (client, server) of a VLL tunnel. Local traffic on a subnet that is destined for the remote subnet will be forwarded through the VLL tunnel. A VLL tunnel can also be configured in a “Capture-All" mode in which *all* traffic at the client BBNA is forwarded through the VLL tunnel to the server BBNA, except for traffic that is destined for the designated subnet of the client BBNA.

In peered mode, the various features of stand-alone mode are still operational. In particular traffic that does not match the criteria for forwarding on the VLL tunnel(s) is processed as it would be in stand-alone mode. This provides reliability in case a remote BBNA peer becomes non-operational.

Pass Through mode versus non Pass Through mode

The BBNA features a "Pass Through" mode that is commonly used in configurations where the BBNA is installed as an upgrade to a legacy network that had a single Internet connection feeding a firewall/gateway. The advantage of using the Pass Through feature is that it does not require configuration changes (e.g. changes to IP addresses and firewall rules) in the existing network. In this mode WAN interface 1 is used to connect to the legacy Internet connection, and additional Internet connections are added on the other WAN interfaces. The WAN interface 1 is not assigned an IP address in Pass Through mode, while the other WAN interfaces are assigned IP addresses.

Note that Network Address Translation (NAT) is used for local devices on all WAN interfaces, except WAN interface 1 since it is in Pass Through mode. This means that traffic leaving WAN interface 1 will have the same source IP addresses as in the legacy network, but traffic leaving other WAN interfaces will have its source IP address set to the IP address of the corresponding WAN interface.

In some installations the legacy firewall/router will be discarded from the current network and will be replaced by the BBNA, which has firewall/router functionality. In this case, WAN interface 1 is not configured in the Pass Through mode. The Pass Through mode is usually not used when the entire network

is designed at the time of BBNA installation. When the Pass Through mode is not used, NAT is used for all WAN interfaces of the BBNA.

BBNA Installation in Branch or Satellite Office

A BBNA device can be installed in an existing branch or satellite office, in which case usually the BBNA device is providing aggregation of DSL, cable, T1, cellular, or satellite links. We consider an example where the existing network configuration is as illustrated here. The legacy site received Internet connectivity from a cable, DSL, or T1 modem, for example. The modem feeds a firewall or other layer 3 device that acts as a local gateway for the network. The firewall typically provides a local private subnet, which in this example is 192.168.1.0/24 in CIDR notation. The devices on the local private subnet are usually not visible from the public Internet, except if a firewall or port forwarding rule is configured in the firewall. In our example, the firewall in the legacy network has the IP address 75.76.77.78, and it belongs to the subnet 75.76.77.72/28. The gateway IP of the firewall is 75.76.77.76, and the IP of the primary DNS is 68.94.155.2.

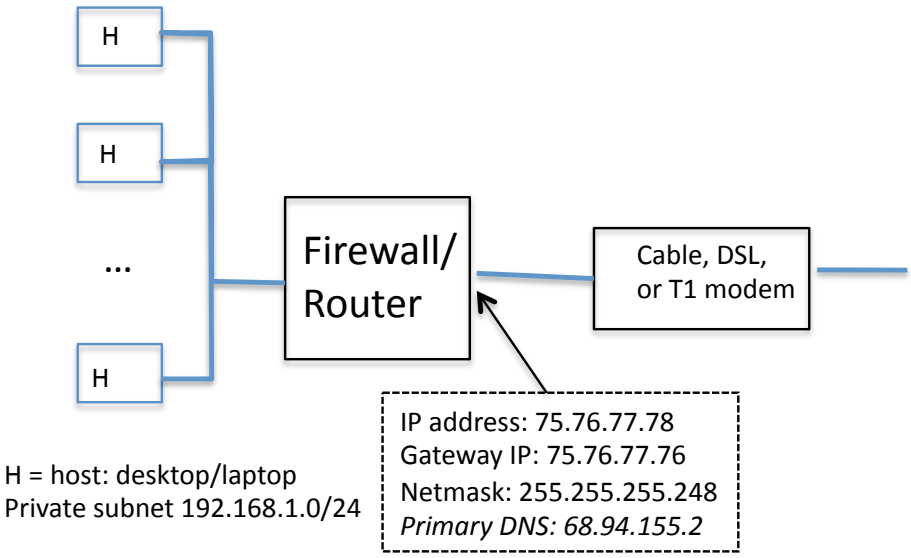

Example Configuration before BBNA installation

Pass Through Mode

If it is desired to retain the legacy firewall/router in the new network, then usually the Pass Through mode is used, since it does not require changing the IP address of the firewall or any firewall rules.

The BBNA device is installed in-line between the firewall and the legacy modem device. Specifically the legacy modem device is plugged into WAN interface 1 of the BBNA, and the firewall device is plugged into a LAN port of the BBNA. Additional Internet access lines, beyond the legacy Internet connection, are plugged into the remaining WAN interfaces. The additional Internet access lines can independently be configured either with a static IP address provided by the ISP, with a dynamic IP address from DHCP, or with an IP address from PPPoE. If it is desired that a WAN interface receive inbound traffic, then the WAN interface will typically be configured with a static IP address.

An illustration of an example configuration after BBNA installation is shown below.

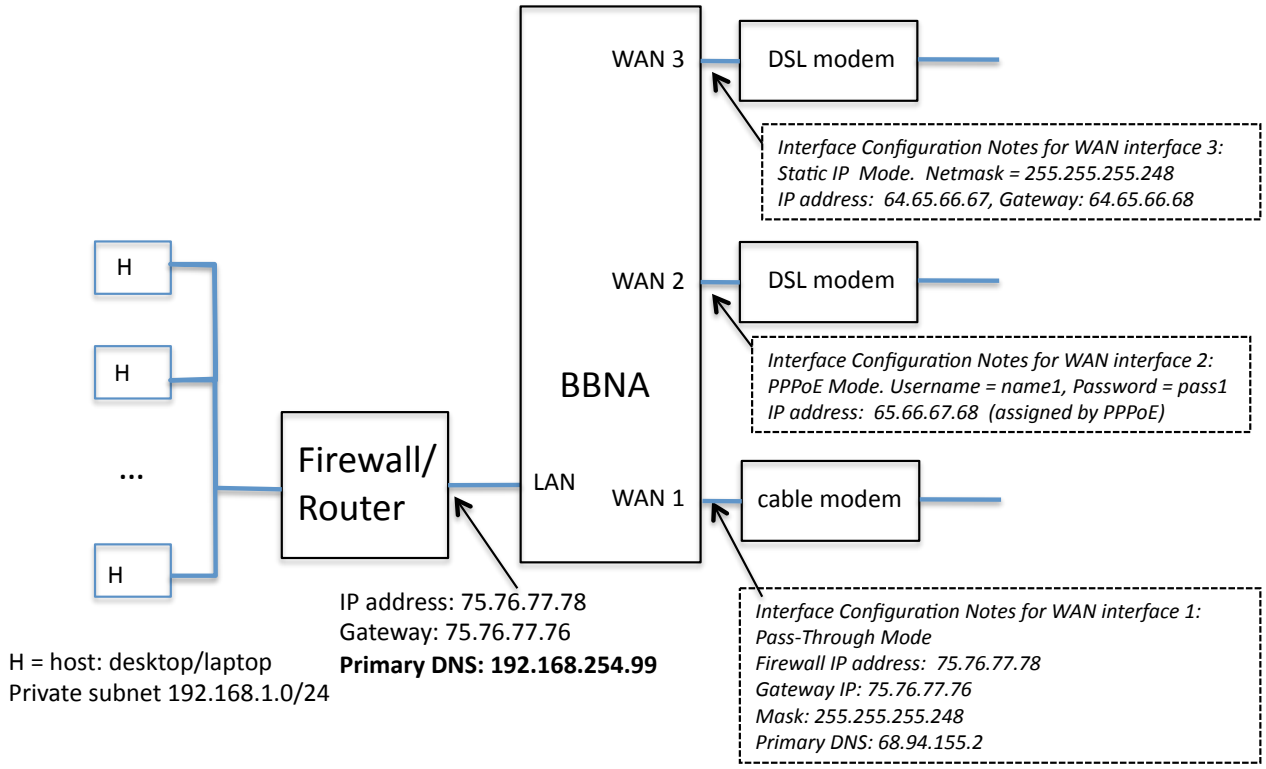

Example Configuration after BBNA installation: Pass-Through to Legacy Firewall/Router

Note that after installation of the BBNA, the public IP address of the firewall is unchanged, as well as its gateway IP address. The private IP addresses of all of the devices connected to the LAN side of the firewall are also unchanged, as well as any port forwarding and firewall rules that may have been set up. There is one configuration change that is recommended, however, and this is the primary DNS address for the firewall. In this example, before the installation, the primary DNS was 68.94.155.2. After the installation, it is recommended that the primary DNS of the firewall be set to the BBNA LAN address, which is 192.268.254.99 by default. This is recommended for better performance and reliability, since the DNS service provided by the BBNA aggregates DNS services across the different available WAN interfaces, including the legacy DNS service. The legacy DNS IP address, 68.94.155.2 in this example, should be entered in the Pass Through configuration for WAN interface 1.

In the example illustrated above, two additional Internet access lines are added, on WAN interfaces 2 and 3. WAN interface 3 is configured with a static IP address. The subnets on each WAN interface other than interface 1 must be non-overlapping with the subnet of the firewall (75.76.77.72/28 in this case).

Dynamic DNS can be used to provide load balancing of *inbound* traffic over the multiple WAN interfaces to the firewall and devices connected to the private LAN of the firewall. In the event of failure of one or more of the WAN connections, automated fail-over for inbound WAN traffic is also supported through Dynamic DNS.

In order for the devices attached to the private LAN of the firewall to be accessible through WAN interfaces other than interface 1, an appropriate forwarding rule is necessary on the BBNA. For example, in this case, a forwarding rule might be desired so that all inbound traffic which is destined to 64.65.66.67 (traffic arriving on WAN interface 3) is forwarded to the firewall with IP address 74.75.76.78, as indicated in the graphic here. If it is desired that WAN interface 3 should not be used for inbound traffic, this forwarding rule is not necessary.

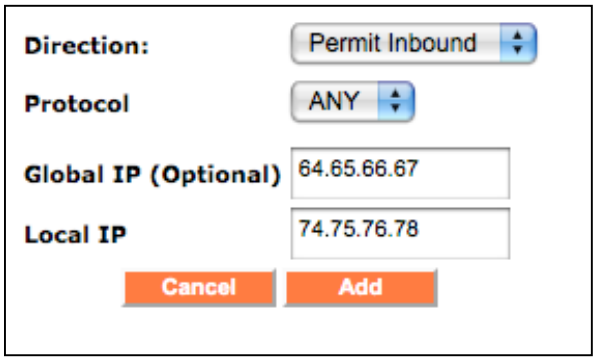

Installation at Branch/Satellite Office without Pass Through Mode

The BBNA device can function as a firewall/router and replace the firewall/router in a legacy installation. In our example, the legacy firewall provided a local private subnet of 192.168.1.0/24. An example of an installation without the Pass Through mode in this case is shown below. The LAN ports of a BBNA belong to the same switch, but an optional Ethernet switch, as illustrated here, can be plugged into the BBNA LAN port will provide additional Ethernet ports for local devices.

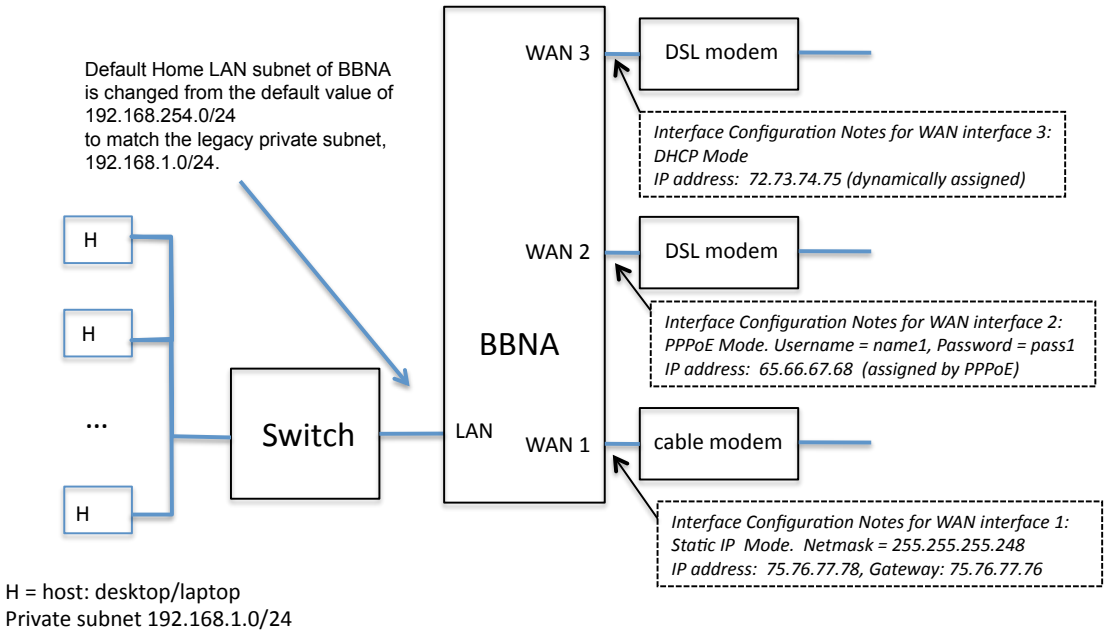

Example BBNA Configuration: BBNA and optional switch replace existing firewall

By default, the BBNA has a home LAN subnet of 192.168.254.0/24. This can be changed to match the private subnet of the legacy network if desired, so that the IP addresses in the legacy private subnet do not need to be reassigned. This is done through the HOME tab in the BBNA Management Interface.

In this example, there is a cable modem plugged into WAN interface 1, configured to support a static IP address. DSL modems are plugged into WAN interfaces 2 and 3. These could be from different providers,

but it is not necessary. In this example, WAN interface 2 is configured using PPPoE, which assigns an IP address to the WAN interface (usually this is a dynamic IP address). WAN interface 3 is configured using DHCP to obtain a dynamic IP address.

A discussed already, if the devices attached to the private LAN need to be accessible from the public Internet, this can be accommodated by configuration of firewall rules within the BBNA that implement "port forwarding". Specific ports of the WAN interfaces can be mapped to private LAN addresses through a firewall rule.

BBNA Installation at Headquarters or Data Center Site

Another BBNA device called BPA (Bonding Proxy Appliance) can be installed for termination purposes in a headquarters or other data center site, where there is typically a very high speed Internet connection available (eg. Fiber based OC-3, etc). A typical installation is illustrated below.

Due to the very high speed of the Internet connection available at the site, only one WAN interface is used on the BBNA. The WAN interface should be configured with a static IP address so that it can be externally addressed from the Internet. In this example the static IP address of the BBNA is set to 65.66.67.68. The gateway IP of the WAN interface, as well as the DNS IP should also be specified for the WAN interface, which is WAN interface 1 in this example.

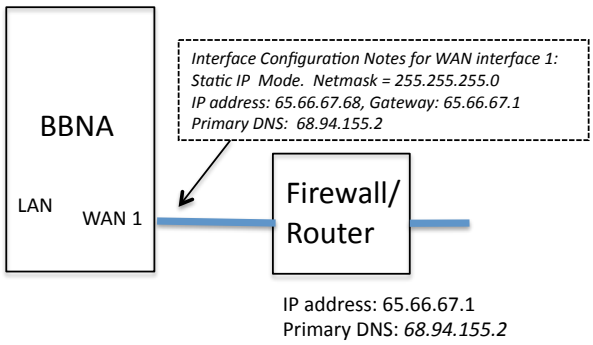

Example BBNA Configuration: Headquarters or Data Center site

The LAN port of the BBNA is used for configuration of the BBNA through the Management Interface, but after configuration it is not necessary to have any devices plugged into the LAN interface of the BBNA. The Management Interface can also be accessed through the WAN interface once it is so configured. If there are no devices plugged into the LAN port, as illustrated here, the BBNA is said to be in a stub configuration.

Usually the sole purpose of the BBNA device in this configuration is to provide termination of a VLL tunnel. If that is the case, since the VLL tunnel uses UDP port 9899, any upstream firewalls must be configured to pass traffic on port 9899. If the encapsulation mode for VLL is disabled, then any upstream firewalls must pass traffic on port 18925, since that is the port used in this case.

Peered Mode of Operation : VLL Tunnels

As mentioned earlier, BBNA client devices enabled with the VLL feature can be paired with another BBNA device called BPA, so that a high speed and reliable data pipe may be enabled between them. For each VLL tunnel, there is an associated "client" BBNA, and a "server" BBNA (BPA). Traffic is forwarded along the VLL tunnel from the client BBNA to the server BBNA and vice versa, i.e. it is a bi-directional data tunnel for IP packets. In order to configure a VLL tunnel it is first necessary to configure each BBNA in stand-

alone mode, as described above. In particular, each BBNA may be using the Pass Through feature on WAN interface 1, or not. We begin with the simplest possible case, where the Pass Through feature is not used on either BBNA device. Also, we first assume there are no other firewalls, except for the BBNA devices themselves.

Associated with each VLL tunnel there is a client IP subnet, and a server IP subnet. In the normal mode of operation, packets seen by the client BBNA that are destined for the server subnet are forwarded along the VLL tunnel to the server BBNA. In the "Capture-All" mode of operation for the VLL tunnel, however, *any* packet seen by the client BBNA that is not destined for the client subnet is forwarded along the VLL tunnel to the server BBNA. In either the normal mode of operation or the Capture-All mode, packets seen by the server BBNA that are destined to the client subnet are forwarded along the VLL tunnel to the client BBNA. Unless the Capture-All mode is used, there is no essential reason that one BBNA should be selected as a client for the VLL tunnel versus the server for the VLL tunnel.

In order to configure a VLL tunnel, a pop-up menu form must be filled out within the GUI for each of the two BBNA devices. On the HOME tab of the BBNA Management Interface, this form will appear by clicking on the "Add Remote LAN" button. Once the form has been completed, click ADD to complete the VLL configuration for that BBNA. The VLL parameters can be edited by clicking on the name of the VLL tunnel in the Remote LAN table of the HOME tab. The order in which the client and server BBNA are configured for the VLL tunnel is arbitrary. The status of a VLL tunnel is displayed in the Remote LAN table of the HOME page. As we will see later, a BBNA can terminate multiple VLL tunnels, so there can be multiple entries in the Remote LAN table.In the first case we consider, both the client subnet and the server subnet are private subnets that are accessible on the respective BBNA device, as illustrated below.

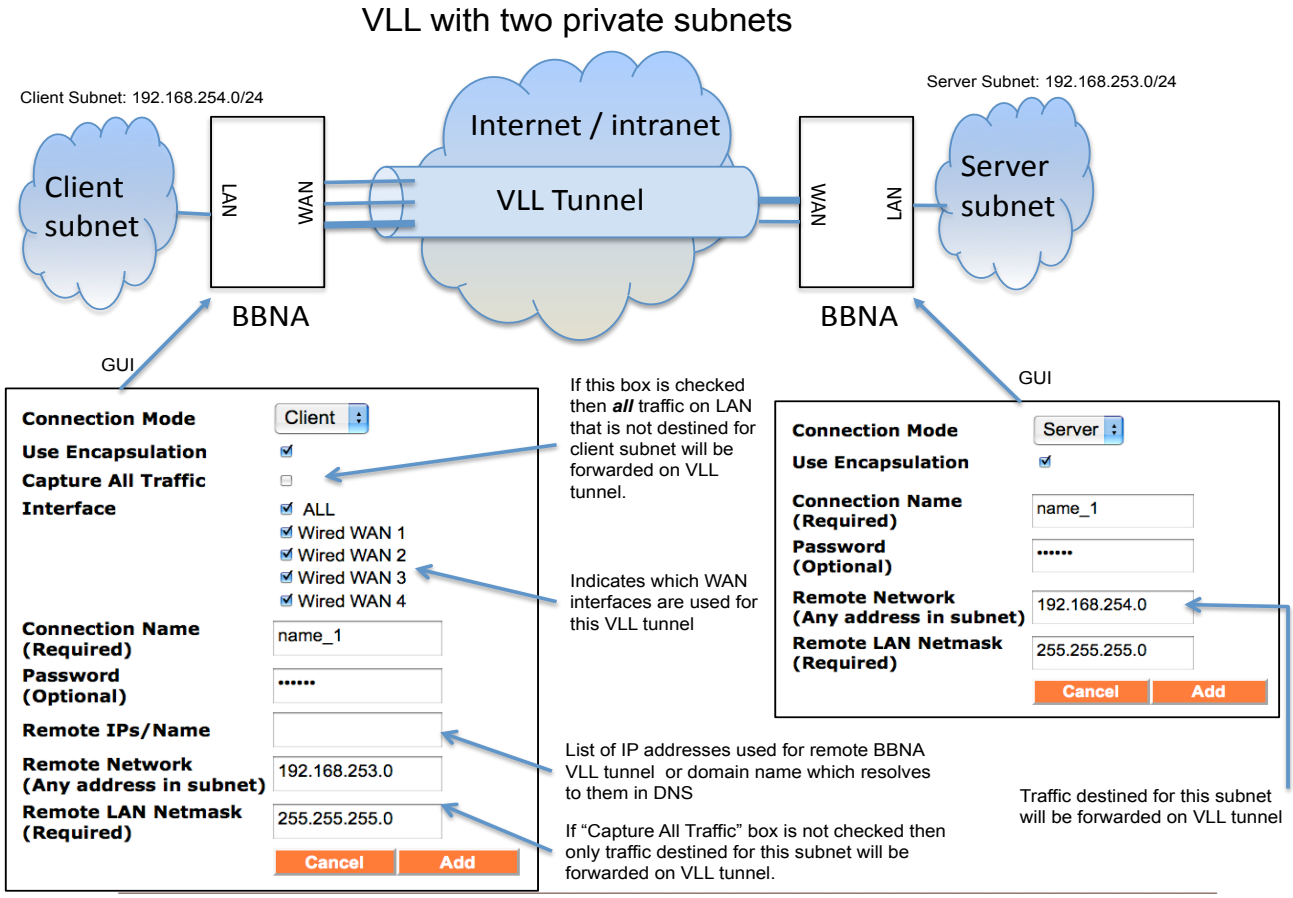

As illustrated, the client subnet in this example is set to the BBNA Home LAN subnet, i.e. the subnet which is associated with the BBNA IP address on a LAN interface. By default this is 192.168.254.0/24. The client subnet and the server subnet must be non-overlapping, so in order to use the BBNA Home LAN subnet on the client BBNA we need to change the default LAN subnet. In this example, the LAN subnet on the client BBNA is changed from the default of 192.168.254.0/24 to 192.168.253.0/24.

In the figure above, the pop-up menus associated with the depicted VLL tunnel are illustrated. As seen, the Client Connection Mode is selected on the client BBNA, and the Server Connection Mode is selected on the server BBNA.

On the associated pop-up menu for the client side, we can select whether the VLL tunnel uses the Capture-All mode or not, as discussed above. We can also select which WAN interfaces will be used on the client side of the tunnel by checking the appropriate boxes. The definition of which WAN interfaces will be used on the server side is made through the text field "Remote IPs/Name" on the client menu. This should be a list of routable IP addresses, separated by commas, associated with the WAN interfaces of the server BBNA. Alternatively a domain name can be specified for this text field, which will resolve to this set of IP addresses, through DNS. In this particular example there are two active WAN interfaces shown on the server BBNA, so in this case those two IP addresses, separated by a comma, should be entered in the "Remote IPs/Name" field unless a domain name is entered.

A "Connection Name" and Password are specified for each end of the VLL tunnel, which must match exactly at each BBNA. The server subnet is specified on the client menu set up and the client subnet is specified on the server menu (the "Remote Network" and Remote LAN Netmask" fields) set up.

In passing it is worthwhile to note that neither the client subnet nor the server subnet of a VLL tunnel needs to be a BBNA Home LAN subnet. The client subnet and the server subnet can be any non-overlapping IP subnets. Generally the VLL subnets should be private IP subnets, unless the Pass Through feature is used. If the client BBNA uses the Pass Through mode, however, then the client subnet needs to be set to the subnet of the firewall attached to WAN interface 1 of the client BBNA. Similarly, if the server BBNA uses the Pass Through mode, then the server subnet needs to be set to the firewall attached to WAN interface 1 of the server BBNA. We discuss the case of the VLL mode with the Pass Through feature through some examples.

In the next example below, the server BBNA is configured in Pass Through mode on WAN interface 1. In this example, the IP address of the firewall attached to the LAN port of the server BBNA is 74.75.76.77, and the associated subnet is 74.75.76.72/28. Therefore the client subnet must also be 74.75.76.72/28. This is done through the values 74.75.76.77 and 255.255.255.248 in the Remote Network and Remote LAN Netmask fields in the pop-up menu for the VLL tunnel on the client BBNA.

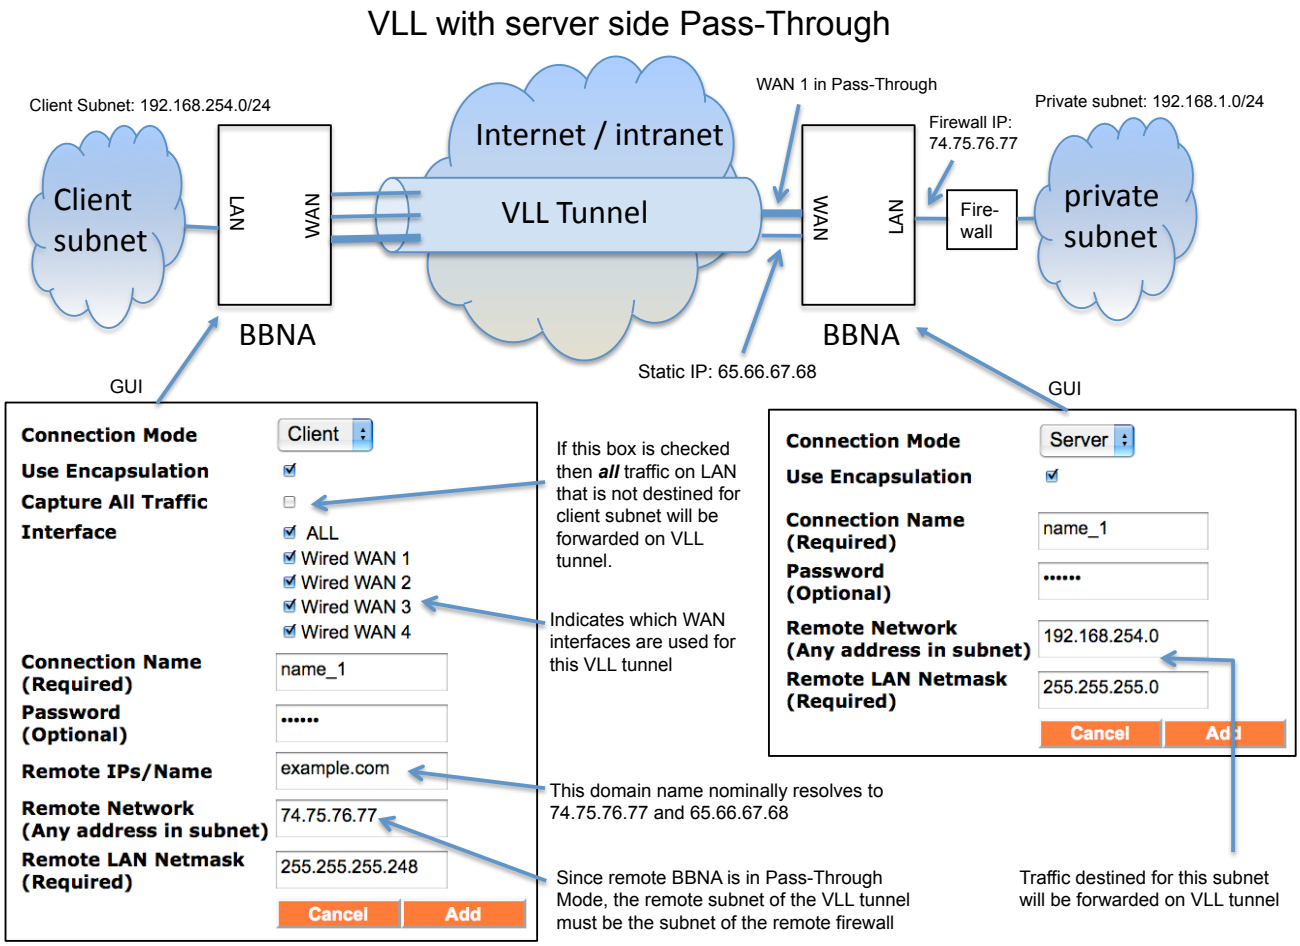

Note that in this case the client subnet is a private subnet and the server subnet is a public subnet, i.e. it is a subnet reachable from the Internet. Due to the VLL tunnel, the private client subnet at the client site is visible to devices within the private subnet behind the firewall at the server site. However, note that the reverse is not true. Namely, devices in the private subnet at the server site will not be visible to devices in the client subnet at the client site. This is due to the firewall at the server site.

It is also worth noting that the home LAN subnets at each BBNA, which are set to 192.168.254.0/24 by default, are separate subnets that can be used for servicing, or to connect other devices.

In the next example, the Pass Through mode is used on the client side, and the server BBNA sits *behind* a firewall. The Capture-All mode is used, and there are no devices present in the server subnet, so it is a "server stub" configuration as discussed previously. In this example the server site has ample bandwidth resources, so only one WAN interface is used on the server BBNA, which is configured with a static IP address. Since there are no devices on the server subnet, there are no devices attached to the LAN ports of the server BBNA. The server site profile matches a "headquarters" for an enterprise, or perhaps a data center that is used to provide Internet connectivity for the client site.

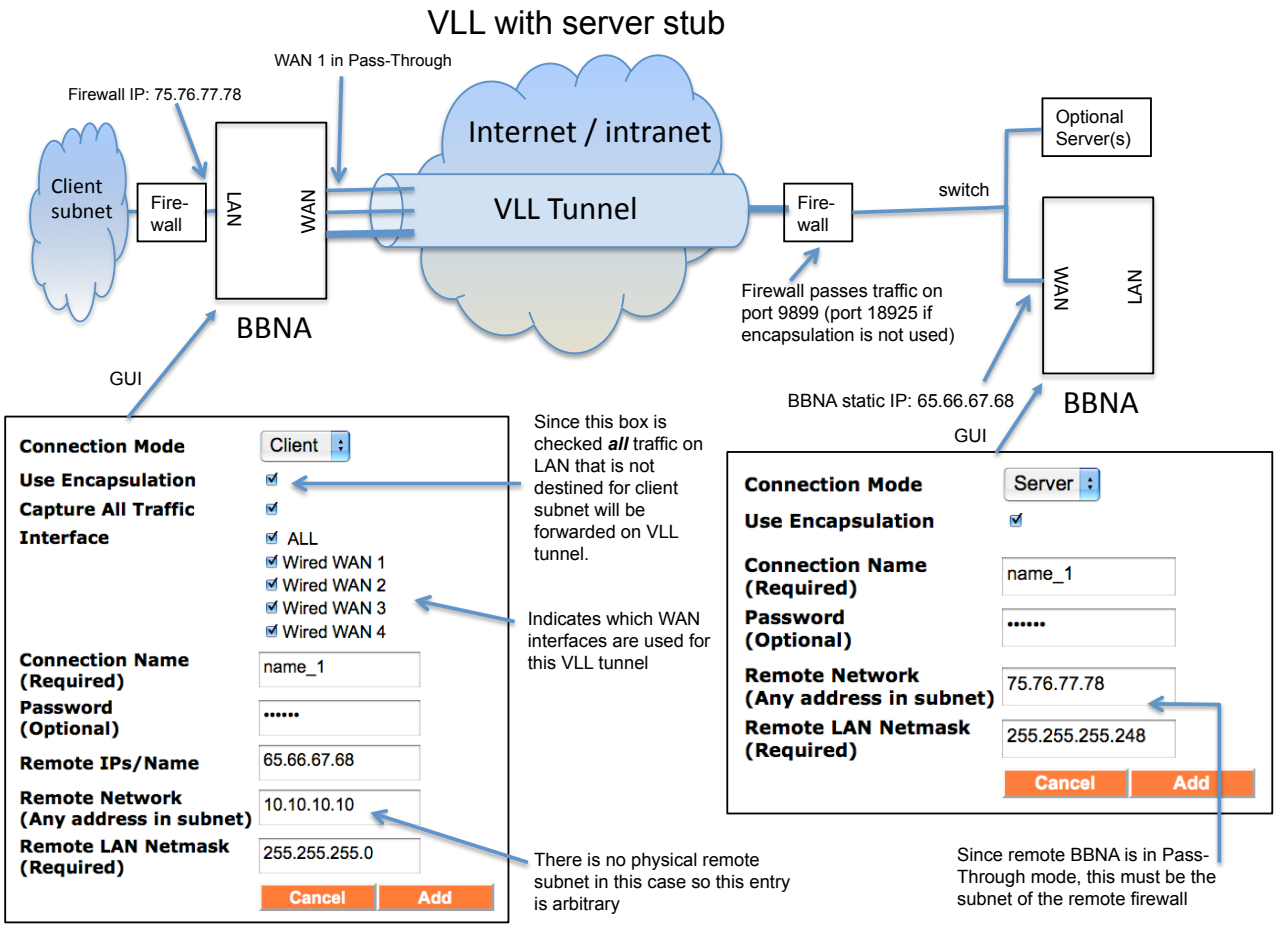

Since the client BBNA uses the Pass Through feature, the client subnet for the VLL tunnel must be the same as the subnet for the firewall at the client site. In this case, this is the subnet 75.76.77.72/28 which is equivalent to the specification of 75.76.77.78 as an IP with a subnet of 255.255.255.248 on the pop-up menu for the VLL tunnel on the server BBNA.

Even though there are no devices on the server subnet, the server subnet must still be specified. In this example, the server subnet is set to 10.10.10.10, which is an arbitrary private IP address. The server BBNA is set up with a static IP address in this case. Since there is only one active WAN interface on the server BBNA, only a single IP address needs to be specified for the "Remote IPs/Name" field in the menu for the VLL tunnel on the client side. In this example, that IP address is 65.66.67.68.

In this configuration, in effect the server BBNA device acts like a proxy for devices within the client subnet in order to reach devices outside the client subnet. The server BBNA performs NAT for devices within the client subnet, such that traffic from the client subnet will appear to other devices as if it is coming from the IP address of the BBNA, which is 65.66.67.68 in this case. This is related to why there are no devices on the server subnet. When traffic arrives to the server BBNA from the client subnet, it gets immediately forwarded out the WAN interface. From there it could go to a server device at the server site, or some other server in the Internet at large.

In this example, the VLL tunnel passes through a firewall at the server site. If the server BBNA is not within a DMZ, reconfiguration of this firewall may be necessary. By default, the VLL tunnel uses UDP encapsulation over port 9899, so the firewall at the server site may need to be configured such that traffic from this port is passed through. The UDP encapsulation is recommended for robustness, but if it is not used, then the firewall should pass all traffic on port 18925.

Before leaving this example, note that a similar configuration to the above, which is slightly simpler, is where the client side is not in Pass Through mode. In this case the client subnet is a private subnet, which is specified in the "Remote Network" and "Remote LAN Netmask" fields in the pop-up menu for the VLL tunnel at the server BBNA.

VLL Enterprise Networks

Thus far we have discussed only pairs of BBNA devices. The BBNA devices may be used to form networks for enterprises, which support global routing over the enterprise. This can be based on a private IP address space that is defined globally over the enterprise. Alternatively, it can be based on a publically routable IP address space, or a combination of publically routable IP address space and private IP address space. Large networks of BBNA devices can use hierarchical routing to ease network management, as we will explain below.

In most cases, each branch office of the enterprise would be equipped with a BBNA device. The headquarters and data center locations would also be each equipped with a server BBNA device, i.e. BPA. The resulting BBNA network will combine available Internet resources for each site for high performance and high reliability data communications across the enterprise, allowing use of the most cost effective transport technologies. For example, DSL is frequently more cost effective than synchronous transport based on T1 circuits, and may be used at most branch office locations. The headquarters and data center sites typically have access to very high-speed fiber-based Internet connections. The BBNA devices within these locations are typically used for VLL termination purposes rather than for bandwidth aggregation, i.e. the BBNA devices at these locations will usually only have one WAN connection.

Each BBNA device that supports the VLL feature can be configured with multiple VLL tunnels. There can be multiple tunnels where the BBNA is the server, as well as multiple tunnels where the BBNA is a client. In reference to any single BBNA device, the collection of server subnets for all the VLL tunnels where the BBNA is a client should be non-overlapping with each other, so that routing of traffic is non-ambiguous. For the same reason, the collection of client subnets for all the VLL tunnels where the BBNA is a server should also be non-overlapping with each other, as well as non-overlapping with the collection of server subnets referred to above. Also, there should be at most one VLL tunnel that is in the Capture-All mode where the BBNA is a client.

Each site where a BBNA is installed may use the Pass Through mode, in which case the local subnet for VLL termination is the subnet of the firewall router attached to a LAN port of the BBNA device. This is a publically addressable subnet. In general, the forwarding rules for handling traffic over VLL take precedence over the normal rules for handling IP traffic. In other words, if there is a conflict between the routing for a packet through a VLL tunnel versus the route taken by normal IP routing, the rule for VLL will take precedence. As a result of this, normal IP routing can be used in this case as a failover mechanism when one or more VLL tunnels go down due to a failure.

For sufficiently small networks it is manageable to have fully connected mesh between all BBNA devices. The case of 3 BBNA sites is illustrated here, where the 3 sites are labeled A, B, and C. In the fully connected mesh architecture, there is a VLL tunnel defined between each pair of BBNA devices. In this case, that translates to 3 VLL tunnels. The location of the server and client for each VLL tunnel is arbitrary, and the Capture-All mode is not used for any VLL tunnel. There is an associated local subnet at each site, which is the client subnet for each VLL tunnel where the BBNA is a client, and the server subnet for each VLL tunnel where the BBNA is a server. These three subnets must not be overlapping with each other.

Fully Connected Mesh Example

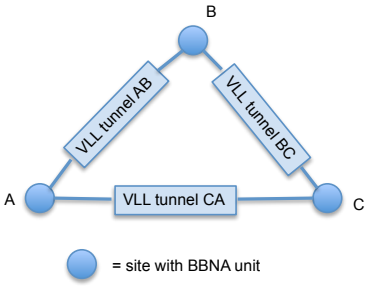

This is a general approach to providing global routing across the enterprise. It can be used for networks of more than 3 BBNA devices, though it may become un-manageable for large networks.

Another approach is appropriate where one "master" site has access to very high-speed Internet connections and is used as a "hub". The hub location will be equipped with the BPA. The BBNA device in each of the remaining sites has a VLL connection to the BBNA device at the hub site, where it is a client and the VLL tunnel is in Capture-All mode. An example with 5 sites total is illustrated here. The hub site is labeled HQ and terminates 4 VLL tunnels as a server. There are 4 "satellite" sites labeled B1 through B4. Each of the BBNA devices at these sites terminates a single VLL tunnel as a client, using Capture-All mode, and has an associated local subnet for the VLL tunnel. Theses subnets must be non-overlapping, so that the rules for forwarding packets at the HQ site are non-ambiguous. For example, traffic originating at the site B1 that is destined for site B4 first travels to the BBNA device at the HQ site since the VLL tunnel for site B1 is configured in Capture-All mode. When the traffic reaches the termination point of the VLL tunnel within the BBNA device at the HQ site, it is forwarded on the tunnel for site B4 to the local subnet at site B4.

Star Topology Example

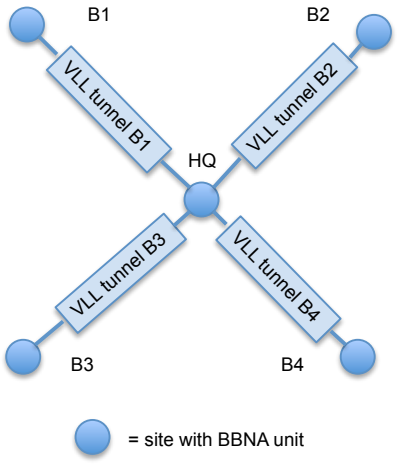

An advantage of this approach is that the BBNA devices within each satellite site have only one VLL tunnel, which is easier to manage. The HQ site has a VLL tunnel for each of the satellite sites. If all of the local subnets at each site are publically addressable via the Internet, then a failure of the HQ site will not preclude data communications between satellite sites, since normal IP routing will be used in that case.

For large and geographically wide networks, it may be desirable to combine the fully connected mesh approach and the star topology approach. For example, suppose there is a headquarters office in each region. The headquarters in each region can be configured in a fully connected mesh between them. The headquarters site within each region then serves as a hub for all the sites in that region. All sites except headquarters sites have a single VLL tunnel operating in Capture-All mode to the local headquarters office.

None of the VLL tunnels between headquarter sites use Capture-All mode. A single VLL tunnel between two headquarters sites may aggregate traffic for different subnets. For example, the local subnet at a headquarters site may be partitioned into smaller subnets, one for each branch site in the region. Therefore traffic incoming to the local subnet at the headquarter sites from a VLL tunnel from another headquarters site is further routed to the appropriate sub-subnet at the associated satellite site through another VLL tunnel. This hierarchical routing approach has the advantage of scalability, since the headquarters site in each region can be independently managed.

An example of this approach with three headquarters sites, labeled A, B, and C, is illustrated below.

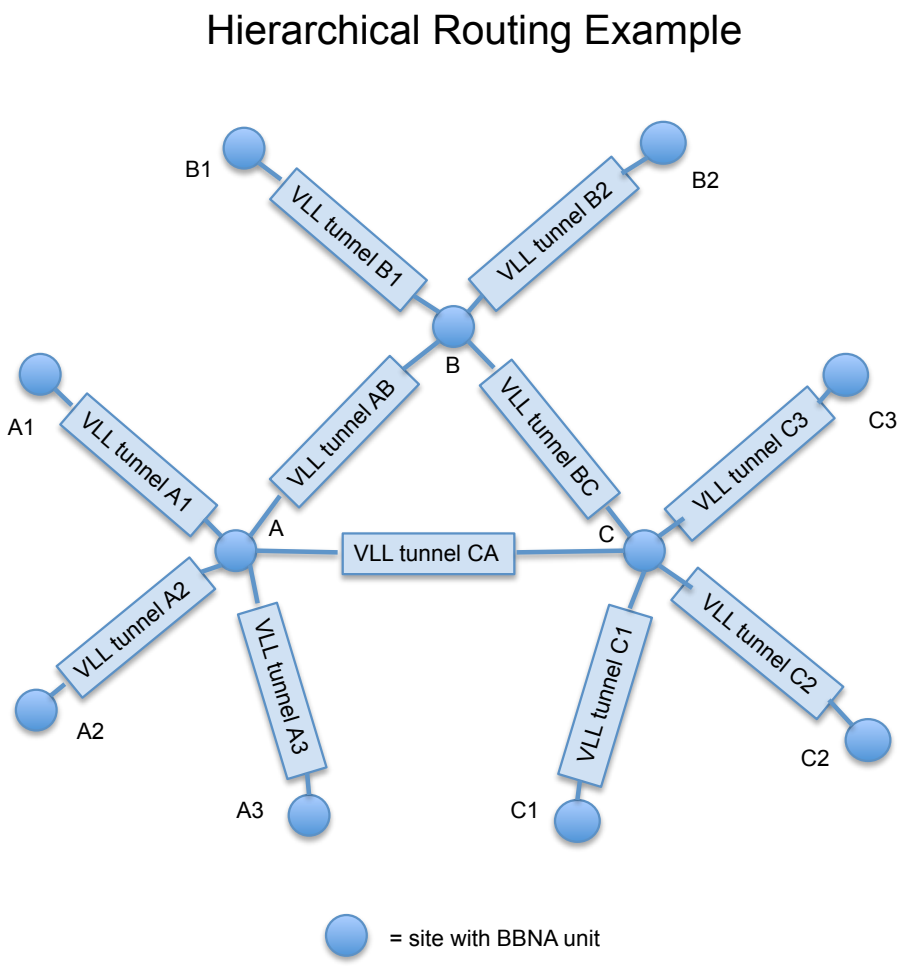

In this example, headquarters office A has 3 regional satellite offices A1, A2, and A3. Similarly, office C has 3 satellite offices C1, C2, and C3, and office B has satellite offices B1 and B2. Each of the satellite office has a BBNA device that is configured with a VLL tunnel as a client in Capture-All mode.

Note that if the subnets for the satellite offices in a region are not contiguous, it may be necessary to install multiple VLL tunnels between a pair of headquarters sites with this approach. Each one of those VLL tunnels between two headquarters sites may carry only a subset of all the traffic that is destined for the regional headquarters office if it is not possible to specify the union of all the subnets in the region as one large contiguous subnet.

As discussed before, if all of the subnets associated with the VLL tunnels are publically routable IP subnets then the architecture is resilient to failures. For example, even if the headquarters site in a region goes down, inter-site communication not involving the failed regional headquarters can still take place over normal IP routing instead of over VLL tunnels.