

ASCE NETWORKS
Omnirange for Multi-homed Network

Technical White Paper

ASCE
NETWORKS[®]
SUPERIOR IN IP TECHNOLOGY.

OMNIRANGE™

A Traffic Manager for Multi-homed Networks

Scope

The Internet continues to grow, with seemingly no end in sight. As more data traverses the global IP backbone, the need for continuous “up time” and reliable data access becomes more crucial. To ensure constant Internet access, networks ranging from a corporate entity to e-commerce sites are deploying multiple Internet connections, usually through multiple Internet Service Providers (ISPs). This network architecture is commonly referred to as MultiHoming. Although this design creates a more resilient network and Internet access model, it also introduces certain complexities that are not always easily resolved. To help manage traffic in such a network design, Asce Networks introduces the OMNIRANGE, an Internet Traffic Management product specifically designed for simplifying networks that utilize a multi-homed scheme. This document discusses the need for the OMNIRANGE and highlights key features of the product that can prove beneficial for any multi-homed network.

The Need

The term “multihoming” generally refers to a network that utilizes multiple connections to the Internet, usually through multiple ISPs (Note that multihoming also refers to a host with multiple network interface cards). For the purposes of this document, we use multihoming as a term to describe a network connected to the Internet through multiple connections. Multi-homed networks are increasing in popularity, because they provide networks with better reliability and performance. Better reliability comes from the fact that the network is protected in case one of the Internet links or access routers fails.

The performance gain comes from the fact that the network’s bandwidth to the Internet is the sum of the bandwidths available through each of the access links. It should be noted, though, that better performance is only achieved if all the links are used collectively. Figure 1 below shows a typical multi-homed network:




Figure 1: A Typical Multihomed Network

In the network depicted by Figure 1, the network has 2 connections to the Internet, one through ISP1 and one through ISP2.

Although at first glance such a network design may seem simple, further analysis proves otherwise. Several complications are introduced with a design such as the one shown in Figure 1. First and foremost is the IP addressing scheme. There are 2 common possibilities that can be deployed in regards to the IP scheme that the internal network in Figure 1 uses:

- ✍ A single IP network number is assigned to the internal network. This will require communication and cooperation between the 2 ISPs in order to advertise proper routes for this single IP network to the rest of the Internet.
- ✍ Each ISP assigns the internal network a different IP address range. Therefore, 2 IP ranges will be active at the same time for the internal network.

Each of the 2 schemes above presents a unique set of challenges. In the first case, where a single IP address is used, the 2 ISPs must coordinate and work together in order for a proper route for the single subnet to be advertised to the rest of the Internet.

Also, care must be taken to assure that both links are used for incoming traffic. If only a single ISP was used to deliver inbound traffic to the network, then part of the motivation and benefits for multihoming would not be realized.

In the second case (which is more popular for average multi-homed networks and the primary focus of this document), where distinct IP ranges from each ISP are assigned to the internal network, there is the issue of what range to use for outbound traffic. If range1 (assigned to the network by ISP1) is used and the link to ISP1 fails, there is no way for the response

traffic to return to the network, since the world knows range1 to be accessible only through ISP1. Furthermore, if only range1 is used, the ISP2 link will never be used for inbound traffic, again since the world knows range1 as accessible through ISP1. Then, there is the issue of what IP addresses to advertise to the world for inbound traffic. If for example, the network had a web server that needed to be accessed from the world, which IP range would the web server belong to? If it were only one of the ranges, the web server would be inaccessible if the ISP responsible for that range lost its link to the network. If addresses from both ranges were advertised, then DNS failover and resiliency would become additional factors that would need to be addressed. All this illustrates that the deployment of a multi-homed network brings with it a set of complexities that need to be addressed carefully.

Address management is only part of the problem. In both scenarios described above, a Truly robust system should have an addressing scheme, which is well known to the Internet, to compensate for link failures. This may involve the deployment of complex routing protocols such as BGP, which would try to assure that a valid path to the inside network (be it a single IP range or multiple ranges) would always be available to the Internet.

Aside from the complexities, a multi-homed network also has some benefits that are typically never fully realized. Consider the following points:

- ✍ The network has multiple links to the Internet. Even with the most sophisticated routing protocols, true load balancing will never be achieved through the multiple links for outbound traffic. Any load balancing decisions that a routing protocol makes will be crude at best, providing perhaps “load sharing”, but nothing more.
- ✍ Some Internet resources are better accessible through one ISP rather than the other. Routing protocols may know basic proximity information, but they generally have no knowledge of dynamic link conditions.
- ✍ For inbound traffic (e.g. Internet hosts trying to access a web server on the multi-homed network), one ISP may provide a better path into the network than another ISP. Again, there is no way to factor in dynamic link conditions for choosing the best path into the network at any given time.

So, not only does a multi-homed network create various design complexities that involve addressing schemes, routing protocols, and DNS, but it also provides for some benefits that are never fully utilized. The combination of all these factors has created the need for Asce Networks' OMNIRANGE. OMNIRANGE eliminates all complexities inherent in the multihoming design, providing a single, easy to manage, “appliance” that intelligently optimizes and utilizes all Internet links. The OMNIRANGE also enables the benefits of a multi-homed network that are never fully taken advantage of with a

traditional design.

OMNIRANGE provides the following advantages for a multi-homed network:

- ✍ OMNIRANGE intelligently manages the IP address ranges assigned to the network from the various ISPs.
- ✍ OMNIRANGE ensures that all ISP links are optimized by intelligently load balancing all outgoing traffic through the available links, while at the same time managing the address spaces used for the outgoing traffic.
- ✍ OMNIRANGE uses Asce Networks' proven detection algorithms to choose the best ISP for outbound traffic.
- ✍ OMNIRANGE assures that both ISP links are used for all incoming traffic and no address from a failed ISP link is ever advertised to the Internet.

In essence, the OMNIRANGE becomes a single, easy to administer, traffic manager for the multi-homed network, eliminating the complexities of routing protocols and uncertain traffic patterns. It also optimizes the multiple ISP connections of the multi-homed network to ensure that all links are used to the best of their potential, thereby making the entire network more efficient.

The remainder of this document will discuss the major mechanisms the OMNIRANGE uses to accomplish the tasks outlined above.

OMNIRANGE for Outbound Internet Traffic

Let's consider a multi-homed network using a OMNIRANGE for traffic management. A typical OMNIRANGE network is shown in Figure 2 below:

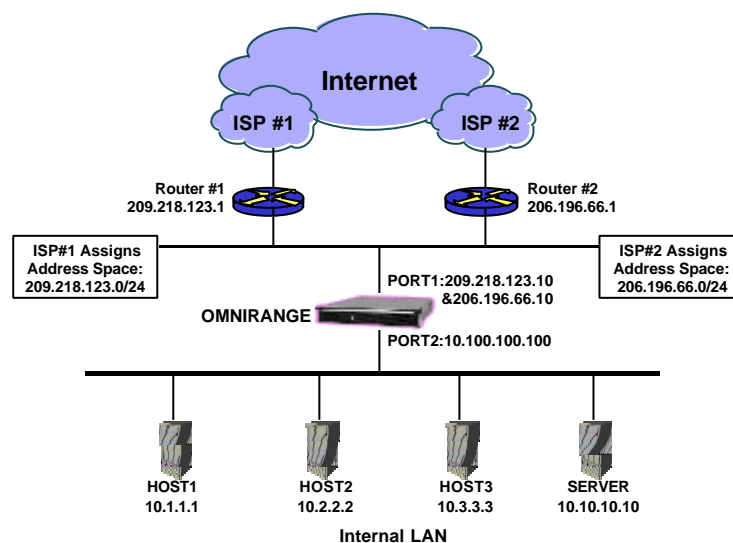


Figure 2: A OMNIRANGE Network

In Figure 2, the multi-homed network is connected to the Internet through ISP1 and ISP2. Each ISP has assigned the network an IP address range: 209.218.123.0/24 from ISP1 and 206.196.66.0/24 from ISP2 (Note: The IP scheme presented here is in the form of “network-IP/bits-in-network-mask.” 206.196.66.0/24 represents the IP network 206.196.66.0 with a mask of 255.255.255.0). Likewise, the Internet knows to get to the 209.218.123.0/24 space through ISP1 and to 206.196.66.0/24 through ISP2. The actual hosts on the network are placed on a private 10.0.0.0/8 network, in order to delegate all address management issues to the OMNIRANGE. Below are highlights of OMNIRANGE features that can optimize traffic management.

IntelligentNAT™

For intelligent address management of outbound traffic, the OMNIRANGE utilizes an algorithm called IntelligentNAT. As mentioned above, the problem with simple NAT

(Network Address Translation) is that it can cause return delivery issues. Let's say, for example, that the OMNIRANGE translated all 10.0.0.0 addresses to 209.218.123.0 (ISP1's address range) addresses. If this were the case, every packet leaving the network would carry a source IP address of 209.218.123.X. Now, let's say the network's link to ISP1 has failed. The OMNIRANGE recognizes this, and now solely uses Router2 and ISP2 to send packets to the Internet. However, if the OMNIRANGE continues to use the 209.218.123.X addresses as source addresses for outbound packets, the packets will get to the destination host, but will encounter return delivery problems. On the response, the Internet host will respond to 209.218.123.X, and the Internet will attempt to deliver the packets through ISP1 (because 209.218.123.X belongs to ISP1's address range and is accessible from the Internet through ISP1, according to global routing tables).

Since ISP1 is down, the response packet will never reach the internal network.

To alleviate this problem, the OMNIRANGE will perform “smart” dynamic NAT. With this feature, the OMNIRANGE will have addresses from both ISPs' address ranges available for translation. Then, when a router is selected to carry an outbound session, the OMNIRANGE will choose an IP address that is associated with that router/ISP. Therefore, in the example of Figure 2, if the OMNIRANGE chooses Router1 as the router to deliver a session to the Internet, it will use an IP address of 209.218.123.X as the translated source address. Likewise, if it chooses Router2 as the router to deliver a session to the Internet, it will use a source IP address of 206.196.66.X. By choosing translated source IP addresses according to the chosen router, return delivery issues will not be encountered.

Content Routing™

In order to optimize outbound traffic, the OMNIRANGE can also optionally perform proximity calculations for outbound traffic. Let's consider the network in Figure 2 again. If an internal host wants to access a specific web site, it's possible that the route through one ISP is more efficient than the route through the other ISP for that specific content. So, the OMNIRANGE performs proximity calculations through all available ISPs (two, in case of Figure 2) to the destination. For future traffic to this destination, the OMNIRANGE will choose the best ISP connection, according to the results derived from these proximity calculations.

To achieve this, the OMNIRANGE will utilize Asce Networks' proven algorithms to efficiently and dynamically calculate the best ISP link per destination. This dynamic calculation optimizes all available ISP links in ways that were never before exercised.

Combining proximity and load balancing assures that ISP links are always optimized, according to traditional load measurements and destination resource based routing decisions.

OMNIRANGE for Inbound Internet Traffic

Not all traffic managed by the Omnirange may necessarily be outbound traffic. There may also be internal resources that need to be accessed from the Internet. Let's continue with our analysis of the network shown in Figure 2, where an internal server (10.10.10.10) is a web server that needs to be accessed from the outside world. In such cases, the Omnirange can be used to overcome some of the obstacles that were defined in "The Need" section of this document, for inbound traffic. Below are highlights of features that can help overcome some of the complications.

IntelligentNAT

IntelligentNAT not only encompasses dynamic IP address allocation and translation, but it also includes the ability to statically map internal resources to external IP addresses. Individual internal resources (such as servers) are mapped to multiple outside IP addresses (one from each ISP). For inbound traffic, the statically mapped IP address from the best available ISP is used. The static mapping of IntelligentNAT also compensated transparently for ISP link failure. If an ISP link is down, only available IP addresses are used for inbound traffic. By making an inside resource available through all available ISPs, uptime is guaranteed for that internal resource. Permanent access to the resource is available through the best and/or most available ISP link.

Proximity

For inbound traffic, the Omnirange utilizes the same proximity mechanisms that it uses for outbound traffic. The reason proximity is an issue for inbound traffic is once again a matter of optimization. If an Internet host needs to access an internal resource (like web server 10.10.10.10 from Figure 2), then it's likely that this Internet host can get to the multi-homed network more efficiently through one ISP versus the other.

To accomplish this, the Omnirange calculates proximity from its network to all networks with hosts trying to access internal resources. Then, it directs the Internet users into the network through the best-suited ISP link, therefore intelligently distributing the load between the available ISP links. In essence, each Internet client will access the Omnirange network through the ISP link that best meets that client's needs.

So, the job of the proximity mechanism in the Omnirange is two-fold. First, it's used to calculate proximity for outbound traffic. This results in Internet traffic traversing the most optimal ISP link, as far as proximity and load are concerned. Secondly, it's used to direct inbound traffic into the network through the most optimal ISP link. The combination of outbound and inbound proximity utilization allows the Omnirange to optimize the multiple ISP links in ways that were never truly realized before.

General OMNIRANGE Highlights

Although special provisions for traffic management are available in the OMNIRANGE, there are also some core features that benefit traffic. The following lists essential OMNIRANGE features vital to its overall role in the network:

- ✍ Link failure detection: One of the primary roles of the OMNIRANGE in a multi-homed network is its ability to detect a failed ISP link. In order to achieve this, the OMNIRANGE continuously monitors the health of the routers it's load balancing to. Furthermore, the OMNIRANGE is capable of monitoring the path through each router, into the backbone of each ISP, if necessary. This allows the OMNIRANGE to provide full path failure detection, where any failure in the path from the network to the world through one ISP renders that ISP "out of service".
- ✍ Redundancy: The OMNIRANGE is deployed in a multi-homed network to provide redundancy for the links, which connect the network to the world. Therefore, it's only natural that the OMNIRANGE feature set includes a mechanism that allows two OMNIRANGES to be installed in parallel, with one backing up the other. OMNIRANGE uses Asce Networks' proven redundancy mechanism where unit health monitoring is done through the

network to compensate for unit failure as well as network failure. A backup OMNIRANGE can easily assume responsibilities for a primary OMNIRANGE if the primary unit, or one of its network connections, were to fail.

- ✍ SNMP polling: If the routers the OMNIRANGE is managing are SNMP-manageable, the OMNIRANGE can poll MIB variables from each router in order to gauge the router's health. Variables such as CPU utilization or bandwidth usage can be useful in such cases to dynamically adjust the amount of traffic sent to each ISP link.
- ✍ Recovery and warm-up timers: If an ISP link is intermittently going up and down, it's best not to send it any traffic until it is continuously stable for a predefined amount of time. The OMNIRANGE has the ability to delay directing session to an ISP link if it detects that a failed ISP link is now active again. This delay is user configurable and can ensure that a stable link is established before any sessions are directed to the ISP. Furthermore, the OMNIRANGE can gradually increase the amount of sessions sent to the ISP, once the ISP link has become active. This can be helpful in order not to overwhelm a sensitive link that has just become active.
- ✍ Backup routers: It's possible to configure a OMNIRANGE not to use an ISP link until all other links have failed. Such a configuration can be useful for small networks that have a low bandwidth backup link to an ISP but still require 7x24 operation.

Conclusion

Multihoming is an increasingly popular way of assuring a network is always connected to the Internet. However, a multi-homed network design brings with it complexities that need special care and attention. The OMNIRANGE alleviates these complexities by taking responsibility for link failure detection, IP address management, and DNS support for internal resources that need to be accessed from the Internet. These elements combined with the OMNIRANGE's ability to optimize the ISP links through detection, make the OMNIRANGE an ideal traffic manager for a multi-homed network. Not only does the OMNIRANGE make traffic management simpler, but it also optimizes the available resources in ways that traditional multi-homed network never utilized before.

Appendix: Related Technologies

Technologies other than OMNIRANGE can be used to support multiple ISP connections, although they fall short of the performance you can expect from OMNIRANGE. For comparison, this section gives an overview of some such technologies and their features. For instance, Border Gateway Patrol (BGP) routes connections using an algorithm that determines the shortest path, calculated by the number of hops (routers) between source and destination. BGP does not interfere with Multi-Link functionality. Virtual Router Redundancy Protocol (VRRP) and Hot Standby Router Protocol (HSRP) are used to make routers highly available. These specialized protocols, used for router redundancy, are not required but can coexist on your OMNIRANGE network with your MultiHoming implementation.

Border Gateway Protocol (BGP)

Independent organizations that maintain multiple Internet links to ensure high Internet availability often implement Border Gateway Protocol (BGP).

- BGP is routing technology that selects packet routes from all available ISPs.
- BGP shares the load but because it does not measure performance it does not perform true load balancing.
- BGP offers high availability for outbound packets but cannot manage problems routing inbound packets.
- BGP chooses carriers without measuring their performance. When BGP chooses slow or congested carriers, network performance suffers.
- When BGP coexists with OMNIRANGE, each BGP router (and each ISP link with redundant routers) is considered a single link.

Limitations

BGP is an ISP-level solution. It is not designed for implementation by end users so it requires specialized ISP resources and equipment. For instance, implementing BGP requires an ISP-independent IP address range. This poses significant risk of service failures leading to incorrect routing unless

the end user successfully negotiates dedicated cooperation between rival ISPs. The implementation is itself a multi-step process with several activities that fall well beyond the normal bounds of software configuration. The implementation team must negotiate agreements between rival ISPs, acquire and configure sophisticated hardware and routing schemes, and must possess advanced BGP programming expertise.

In comparison, OMNIRANGE is part of a managed product designed for end users that requires no additional or specialized hardware or software. This significantly reduces comparable implementation and maintenance costs. OMNIRANGE selects the connection with the fastest throughput, while BGP cannot tell whether a path with more hops is faster than a congested path with fewer hops. Finally, OMNIRANGE does not require additional processing capacity or hardware, while BGP resides on the router and requires extra processing capacity to calculate the shortest path, which is an added expense.

ASCE
NETWORKS[®]
SUPERIOR IN IP TECHNOLOGY.