PLANETLAB

IPv4 Address Use in PlanetLab

Jeff Sedayao
Intel Corporation

PDN–04–020
April 2004

Status: Ongoing Draft.

IPv4 Address Use in PlanetLab

Jeff Sedayao

1. Introduction

The choice of IPv4 addresses used by a PlanetLab node has important implications for the usability and availability of that node. When assigning an IP address to a PlanetLab node, several design decisions that must be made. The IP address must be taken from some block of IP addresses. That address block needs to be chosen, along with the characteristics of the block, such as whether the block is directly reachable from the entire Internet and whether the address is contiguous with other PlanetLab nodes at the same site. Ownership for address space can be delegated to organizations such PlanetLab central. Address space decisions can affect how a node is accessed and even how security reporting concerning the node is done. This PDN describes these and other key IPv4 addressing issues regarding PlanetLab.

2. IP Addressing Design Issues and Choices

This section goes over three key design decisions when addressing PlanetLab nodes:

- what address space to use
- how much address space to use
- delegate address space authority

For each choice, we examine alternatives, make recommendations, and define requirements.

2.1. What Address Space to Use

A key decision when configuring a PlanetLab node is deciding what address space to use. There are three basic choices:

- Local address space
- PlanetLab owned address space
- Route Restricted IP Addresses

In addition, there is another kind of address space that can be used with all three:

- dynamically allocated address space.

A first option is use the address space that is local to the organization hosting the PlanetLab node. We can delegate portions of a large address space owned by PlanetLab to each site hosting a node. Alternatively, we can use space that is not easily reached by sites on the Internet, such as private network space. Another choice we can make is whether to use dynamically allocated address space, allocated by protocols such as DHCP, or use statically allocated address space. We will go over each choice in this section, covering plusses and minuses for each.

2.1.1. Local Address Space

With local addressing, the PlanetLab host uses IP address provided by the organization hosting it. For example, if an organization uses network 143.183.0.0/16, then the PlanetLab node hosted by that organization would use an IP address within 143.183.0.0/16. This approach has several advantages. Installation of nodes can be easier since the local organization already controls the address block. They can do not need to consult PlanetLab Central before setting up the node. There should be fewer issues with allocating and routing the address space since the address space is already used. The local organization will also be informed of any security incidents and complaints associated with PlanetLab. Some organizations will find that particular feature to be a drawback, as they would prefer that PlanetLab Central handle security incident handling.

2.1.2. PlanetLab Address Space

Instead of local address space, another option is to use parts of address space owned by PlanetLab. Blocks of this address space would be delegated to each site setting up PlanetLab hosts. The primary advantage of this is scheme is that the security incident reports could be channeled to PlanetLab administrators and optionally to the appropriate organization hosting the PlanetLab node. An additional advantage is that it would be easy to write policies for filter PlanetLab traffic, since all PlanetLab hosts would be from a limited set of networks.

There are many disadvantages to this scheme. First, PlanetLab Central would find it difficult to obtain a sufficient amount of address space to delegate to organizations hosting PlanetLab nodes. Typically, the smallest address space routed on the Internet today is a Class C (256 addresses), and these are only allowed from a specific range of addresses. As the number of PlanetLab sites increases, we would require a Class C per site. Scaling to 256 sites would then require that PlanetLab obtain a class B address space (a /16 or 64 K addresses). Even if were possible to obtain a class B address, much of the space required will not be used. Second, using non-local address spaces complicates routing and would make the task of landing much more difficult. Organizations landing PlanetLab nodes would need to advertise the appropriate parts of PlanetLab address space, and this typically would require changes at their ISP connectivity. Third, use of non-local address space can cause routing anomalies. Some geographies will route networks that it considers non-local in strange ways (e.g. via different continents) even if they are advertised locally. This is contrary to one of the goals of PlanetLab, to have a planetary scale test bed that experiences local conditions.

2.1.3. Using address space not directly routable

A significant portion of IP address space is not directly reachable from the general Internet. This address space includes, firewalled address space, private address space [1] and Internet2 networks. An advantage of using this space is that there is much more of these kinds of address space than directly reachable address space. Many organizations put significant amounts of address space behind firewalls. Also, there are communities of interest to experimenters that can be measured or serviced by PlanetLab experiments.

The chief difficulty of using these kinds of address space concerns accessing PlanetLab nodes in those address space. While there are available methods to access such nodes through proxies (currently used with Internet2 PlanetLab nodes) or various network address translation (NAT) techniques, setting these up adds additional time and complexity. When there are problems, the additional indirection and address translations increases the time and coordination needed to debug and solve those problems.

In the near term, use of this kind of address space is discouraged except when there is significant research interest in the communities using those more isolated address spaces.

2.1.4. Dynamically Assigned Addresses

At the moment, PlanetLab recommends statically allocated IP addresses. IP addresses are currently configured on a file on a floppy disk, and this file must be reconfigured manually if this an IP address changes. Also, traffic accounting for incident handling is much simpler to do if a node IP address does not change. At some point in the future, however, we may consider using mobile PlanetLab nodes or nodes with intermittent connectivity. Access to pools of addresses (see below) would be simplified with access to dynamic address allocation mechanisms. In addition, use of dynamic allocation protocols such as DHCP may allow sites that are short on address space to conserve their addresses.

2.1.5. Recommendations

We recommend using local address spaces that are reachable from commercial Internet when addressing PlanetLab hosts using statically assigned addresses. In the near term, Addressing from route restricted spaces (such as on Internet2) or firewall limited space such only be used if there is sufficient interesting in studying or providing services in these address spaces and if there are readily available techniques for access them. Longer term, since there is a vast pool of hosts in address spaces that are in private address space or behind firewalls, we will develop techniques to land PlanetLab hosts in such spaces and access them from all of the existing Internet. PlanetLab should adapt to using dynamically allocated IP addresses to deal with mobile, servers with intermittent connectivity, address pool use, and to allow more efficient use of address space by hosting organizations.

2.2. Address space ownership

Another issue with IP address space is who will own the address space. “Ownership” has two key aspects:

- what organizations and people will designated as responsible for all technical and security matters regarding that piece of address space,
- who will perform and manage the IP address to name mapping.

Both of these have implications for security incident handling and management. We will discuss both of these issues in the following sections.

2.2.1. Registered Ownership of Address Spaces

Each block of IP Address space routed on the Internet has one or more owners associated with it, as registered with Internet numbering authorities like ARIN, APNIC, and RIPE. Owners are responsible for all issues associated with the block, and this information is publicly available on the Internet. Spam and security reporting software often use this information to inform the perceived source of an attack or spam.

There are two options for ownership:

- the local organization can own the space
- the space can be delegated to another organization.

If a local organization continues to own the space, spam and security reports will be sent to that organization. In some cases, this may be considered desirable, but in other cases, an organization's abuse desk or security personnel may not be sufficiently staffed to deal with a potential increase in incidents. An organization can delegate the address to PlanetLab central and let incidents be handled by PlanetLab central. This is has the advantage of reducing the burden of incident handling away from the local hosting organization. It does not eliminate the burden, as some incident response tools will copy not only the organization responsible for a block but also the organization that delegated it. Some organizations may want to know about all security incidents regarding their space, and this can be handled by setting up an e-mail alias for the addresses space abuse address that copies both the local organization and PlanetLab central. In some cases, this may not be possible, especially if the address space has been subdelegated already or if the IP addresses of the PlanetLab hosts are in addresses that cannot be delegated as a block.

We recommend that organizations delegate their address space if possible, but this not mandatory. This should reduce the workload of local organization while keeping them informed. We understand that this may not be possible or desirable by a hosting organization, so the ultimate choice is left to them.

2.3. How much Address Space to Use

The final question regarding IP address is what address space to use. There are a number of parts to this question:

- use of contiguous blocks
- support for multiple interfaces
- addressing of slices

We look at each of these areas in the next section and make recommendations on each.

2.3.1. Address Block Boundaries and Contiguousness

There are a number of design choices regarding address block boundaries that need to be made. First is whether addresses allocated should be contiguous or not. A contiguous block of addresses makes it easier for local administrators to set appropriate firewall filters as needed. Of course, since address space may be in short supply, this may not always possible. For example, in some data centers where address space is tight, PlanetLab nodes may be directly connected to router interfaces. In other cases, PlanetLab nodes may be landed on segments with existing non-PlanetLab nodes.

Another issue is whether the size of the block should fall on CIDR boundaries. There are a number of advantages of having the block of addresses used for PlanetLab hosts fall on CIDR boundaries. Dedicating a CIDR block makes it possible to delegate address space ownership and makes it much easier to delegate reverse DNS. It also makes it easier to create firewall filters, particularly with filtering routers. A drawback to this approach is that potentially scarce address space is unused with this approach.

Since conditions vary between organizations, we recommend but do not mandate dedicated CIDR blocks to a segment where PlanetLab hosts live. It does can make maintenance, security, and other tasks easier for hosting organizations, but the address space demands of using CIDR blocks dedicated to PlanetLab may not be tolerable.

2.3.2. Multiple network interfaces

At this point, there is no compelling reason to use multiple interfaces, and several reasons not to use them. There is no foreseeable benefit to using multiple interfaces that land on the same segment. When a node has interfaces on multiple interfaces, this complicates routing and adds complexity when debugging problems. Multiple interfaces increase the possibility of Martian responses, the case when applications do not function properly when responses to service requests come from different IP addresses. Moreover, multiple interfaces can consume additional IP address space.

Despite these drawbacks, it is likely that a compelling reason for using multiple interfaces may emerge. While we recommend that for now, multiple network interfaces not be used, it would be best to prepare for the case where multiple interfaces are used.

2.3.3. Slice Addressing

A final question is how should slices be addressed. There are two extremes regarding this:

- each slice has its own address
- all slices share one IP address

Having each slice use an IP address has a number of advantages. It simplifies tracking of traffic for security purposes (since each slice uses a distinct IP address), and will allow multiple slices to use the same TCP or UDP port without the need for multiplexing schemes. The main disadvantage is the consumption of address space, as there may be tens if not hundreds of slices on a node.

The other option, one IP address for all slices, is currently used on PlanetLab nodes. A major disadvantage of this approach is that should more than one slice want to run a service on a port, a multiplexing scheme would be required. Another option that straddles the two extremes is for all slices is to have a pool of addresses used for each node. Should services need to share a port, an IP address would be allocated and a port would be allocated to slices requesting access to a port already in use. This approach has the advantages of simplifying sharing of TCP and UDP [2] ports, but has the disadvantage of consuming more IP addresses, although far fewer than the address per slice approach. One proposed use of an address pool would be for the construction of honey farms. Multiple addresses could be assigned to hosts, and use of an IP address other than the main address would be indications of a port scan or a scanning worm or virus.

PlanetLab currently uses the one address for all slices method, and for most slices, this should be adequate. PlanetLab nodes should evolve, however, toward the address pool concept of slice addressing. In addition, when there is a significant contention for a port, a port multiplexing scheme would be extremely useful. Both of these two methods should be available for use.

3. Summary

Below is a table of IP addressing recommendations.

Address Characteristic	**Current Recommendation/Requirement**	**Long-term Recommendation/Requirement**
Type of IP Address	Recommend using fully routable addresses ; other address space (NAT, route restricted or firewalled space) possible if there is sufficient interest	Proxy/NAT/tunneling techniques developed to traverse this space
Address Space Ownership	Use local address space is required.	Develop techniques for effectively using firewalled, NATted, and route limited space.
Support for Multiple interfaces	No support	Support multiple interfaces
IP to Name Mapping	Optional delegation to PlanetLab central	
Address Space delegation	Optional delegation to PlanetLab central	
Address Block	CIDR block recommended	
Static Addressing	Required	Work with Dynamically

		assigned addresses
Slice Addressing	Use one IP address for all slices	Use a pool of addresses for slices, develop multiplexing scheme for port usage

4. References

[1] Y. Richter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear. “Address Allocation for Private Internets.” RFC 1981. February, 1996.

[2] Jeff Sedayao and David Mazieres. PDN-03-016. Port Use and Contention in PlanetLab. November 2003