

Report of the RIPE Enhanced Cooperation Task Force

RIPE Enhanced Cooperation Task Force

Document ID: ripe-464

Date: March 2009

Abstract

This document has been produced by the RIPE Enhanced Cooperation Task Force to explain:

- How the existing structures of Internet governance evolved
- Why these structures are uniquely suited to facilitate ongoing development and innovation in the Internet

With stakeholders outside the traditional Internet community, particularly in the public sector, taking an increasing interest in Internet governance, it is vital that these points be effectively communicated, and that we ensure that innovation and technical development continue. As the "Information Society" considers the future development of Internet governance, "enhanced cooperation" between the RIPE community and these new stakeholders must be a high priority for both the RIPE community and the RIPE NCC.

Contents

PART I: INTRODUCTION TO INTERNET ADDRESSING	3
Chapter 1: Internet Addressing Explained	3
What is an IP Address?	3
IP Addresses and the Domain Name System Distinguished	3
IP Addresses and Internet Routing	4
Chapter 2: Address Management Explained	5
Bottom-Up Coordination	5
Chapter 3: Mechanisms For Address Management	6
PART II: HOW INTERNET ADDRESS MANAGEMENT IS CONDUCTED IN THE RIPE NCC SERVICE REGION	7
Chapter 4: Organisational Structures	7
RIPE	8
RIPE Working Groups	8
RIPE Task Forces	8
RIPE NCC	8
Global Coordination	9
ICANN and IANA	9
Chapter 5: Policy Development And Implementation	10
Types of Policy	10
Address Allocation Policies	10
Database Policies	11
Best Practice Statements	11
The Limits of Addressing Policy	11
RIPE Policy Development Process	11
RIPE NCC Activity Plan	11
PART III: ENHANCED COOPERATION	12
Chapter 6: The Need For Enhanced Cooperation	12
Chapter 7: Consultation Processes	13
RIPE Meetings	13
RIPE NCC Government Relations	13
RIPE NCC Government Roundtables	13
Other Processes	14
Chapter 8: Learning Points	14
What Works	14
What Doesn't Work	14
What Is Not Yet Known	14
Recommendations	15

Part I: Introduction To Internet Addressing

Chapter 1: Internet Addressing Explained

What is an IP Address?

At the heart of how the Internet works is the Internet Protocol, the standard according to which different devices communicate. The genius of the Internet Protocol is that when devices use this standard they can connect to each other and exchange information, even if they are completely different kinds of device (such as a personal computer running Windows and a mobile phone), made by different manufacturers, owned by different users and are connected to different networks in different parts of the world. Using the Internet Protocol (also known as *IP*), truly global interoperability is achieved between highly diverse information processing systems.

In order for one device, such as a PC, to send data to another, it must be possible to identify the destination and distinguish from alternative possible destinations. The Internet Protocol achieves this by saying that each device must have an address, known as an *IP address*. As the public Internet seeks to achieve global interoperability, each device must have an IP address that is distinct from that for every other device; its address must be *globally unique*¹. If the addresses were not unique (that is, if two or more devices use the same address) then data intended for one physical destination could be mistakenly sent to another. This would not merely result in a loss of confidentiality: if data is sent randomly to several devices experiencing such an *address clash* it will generally result in a complete communications malfunction.

Ensuring that each Internet-connected device is capable of being assigned a globally unique IP address is therefore a primary requirement for the successful operation of the Internet.

IP Addresses and the Domain Name System Distinguished

IP addresses are the means by which machines identify themselves. In the case of IP version four (*IPv4*), each address represents a 32-bit number, which is actually displayed as a series of four numbers (each between zero and 255) separated by dots, such as 10.201.57.254. This format is not particularly memorable for humans.

In order to identify a particular location on the Internet (such as a website), humans generally use a different addressing scheme called the *Domain*

¹ The Internet Protocol can also be used to enable communications on private networks as well as the public Internet. The IP addresses of devices on a private network do not need to be globally unique, merely unique amongst the devices connected to that network. However, if such devices are also to communicate with the public Internet they must have a globally unique IP address, or else rely on some intermediary device that does. For details of strategies whereby many devices may share the use of such an intermediary, thereby reducing the number of IP addresses that are needed, see *Network Proxy* and *Network Address Translation*.

Name System (DNS). The DNS contains the addresses that are widely recognised, such as www.example.com. When a human uses a computer and seeks to contact a website using such an address, the computer must first translate the DNS address (which is understood by humans) into an IP address (which is understood by computers). www.example.com becomes 10.201.57.254.

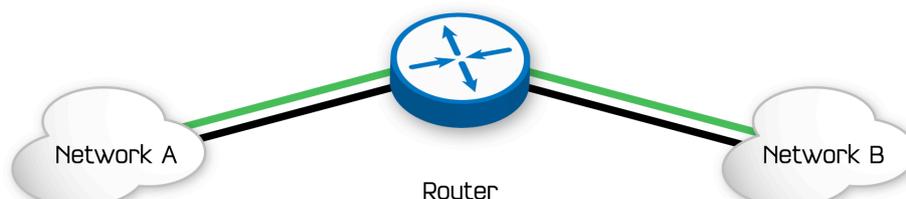
Both IP addressing and the Domain Name System are critical Internet systems and present important issues for Internet governance. They are, however, quite different systems; they are managed separately and face different issues. This paper is concerned with the management of IP addresses only.

IP Addresses and Internet Routing

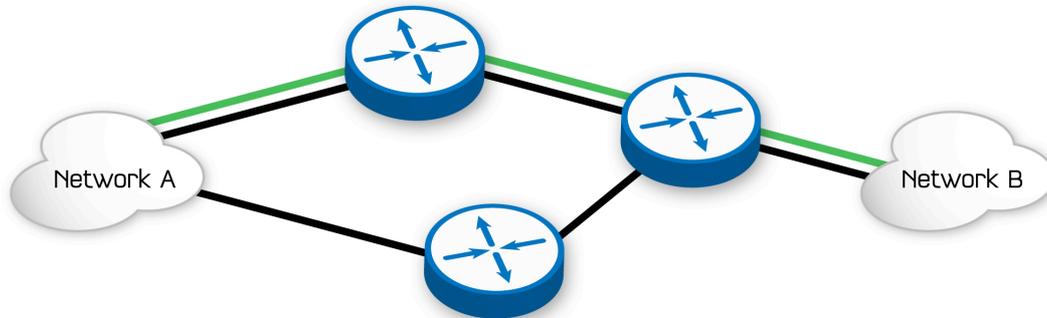
A machine that accepts data from another device and passes it on towards its ultimate destination using the Internet Protocol is called a *router*. Internet networks consist of routers, which select paths or *routes* over which they transmit messages toward their destination.

To transmit a communication across the Internet, the data is first split up into *packets*. This is true whether the communication is an e-mail or a web page, a voice call or a video stream, or any of the other myriad types of communication that occur on the Internet. Each packet is labelled with the destination IP address, and passed from one router to the next until it reaches its destination, at which point all the packets are reassembled into the format of the original communication. This means that routers do not need to understand the different kinds of higher-level communication that occur on the Internet – to a router, the packets that make up an e-mail appear no different from those that make up a video stream; routers simply need to know how to route Internet packets towards their destination.

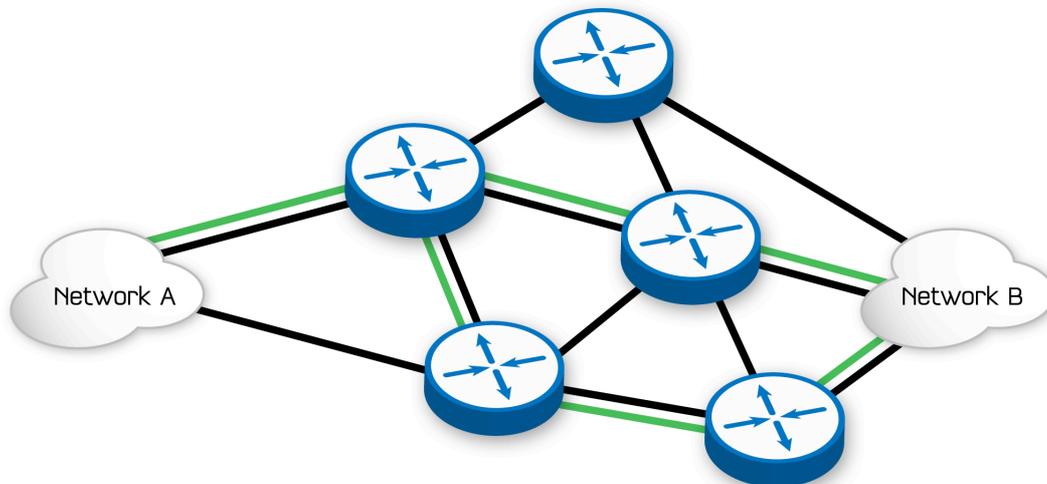
A router may be connected to a few or many other routers, but since the Internet spans the world, no router will be directly connected to every other router. It is not possible to be directly connected to more than a tiny proportion of all the routers that exist. An Internet packet will therefore commonly pass through a chain of routers as it moves progressively closer to its destination. The core function of a router is to choose the best nearby router to pass each packet to, so that the packet will reach its destination.



Each router *announces* or *advertises* a list of routes it can process, expressed as ranges of IP addresses for which it can provide routing service. Put simply, a router tells its peers, “If you have traffic intended for IP addresses in the range between 192.0.0.1 and 192.5.255.255 (for example), pass me those packets and I can route them onwards.”



This is something of an over-simplification: in practice, a router will often be directly connected to more than one other router capable of reaching a particular destination. There is therefore a communications protocol that routers use, called *Border Gateway Protocol* (or BGP), which assists routing decisions. This includes helping routers to determine which is the best (shortest) path for a packet to reach its destination.



Chapter 2: Address Management Explained

Bottom-Up Coordination

As discussed in Chapter 1, it is imperative that IP addresses are globally unique. However, IP addresses are just numbers that have been programmed into a computer or other network device as its address. How does the person configuring a new computer with an IP address know which address to use?

In principle, a new device could select any unused IP address and use that. This is slightly complicated by the fact that the device must persuade the router to which it is connected to honour that choice of address, and to announce to its neighbours that it can now route traffic to this new address. And how does the router, or chain of routers, know whether to honour such a new address? In one sense, it doesn't matter: one IP address is much like another; they're just numbers and no one address is superior to any other. Accordingly, one might suppose that any such announcement should be honoured.

However, if everyone simply selected their IP address at random, there would inevitably be clashes, and as explained earlier this would destroy the global interoperability of the Internet – data sent across the public Internet would no longer reliably reach its destination. There is therefore a need for IP address users to coordinate amongst themselves so that when configuring a new device they do not select an IP address that is already in use.

The design of the coordination process is called a “bottom-up” process because it is based on the concept that all IP address users can make an active contribution to that process. Every user of IP addresses shares the same interest in ensuring that address conflicts do not occur, and so the community of IP address users comes together in various groups to organise, discuss and make decisions on how the process will work. This is distinguished from “top-down” hierarchies, such as governments or corporations, where an established and recognised authority has the right to determine a policy and to instruct others to carry it out.

Within the geographic region comprising Europe, the Middle East and parts of Central Asia, the name given to the bottom-up community is RIPE (Réseaux IP Européens). The primary purpose of RIPE is to ensure the necessary administrative and technical coordination is achieved. The RIPE Network Coordination Centre (NCC) acts as secretariat to RIPE and carries out the administrative directives of the community.

The RIPE community has determined that it needs to act collectively to prevent IP address space conflicts. This requires mechanisms to facilitate coordination, develop *policies* for the operation of those mechanisms, institutions to operate those mechanisms in a neutral fashion in accordance with the policy, and processes for the development of new policy and the governance of the institutions. The next chapter discusses these mechanisms, policies, processes and institutions, and explains how these are all derived from and closely tailored to support the requirement to prevent address space conflict.

Chapter 3: Mechanisms For Address Management

To satisfy the requirement that IP addresses be globally unique, the mechanisms for IP address management must be global in scale. As noted in Chapter 2, bottom-up policy is coordinated in Europe, the Middle East and

parts of Central Asia by RIPE, but there are similar organisations in other geographic regions that play similar roles. These other organisations are:

- AfriNIC: Africa
- APNIC: Asia Pacific
- ARIN: North America
- LACNIC: Latin America and the Caribbean

As well as acting as a forum for policy decision-making, these organisations perform the function of *Regional Internet Registries* (RIRs). RIRs have two vital tasks:

1. To distribute IP addresses to their regional community according to the policies developed by that community
2. To maintain a publicly available registry of which addresses have been allocated for use in that region, and to whom

Under the first task, RIRs are designated as the "gatekeepers" of IP addressing, though they act on the instruction of the community members to whom they distribute addresses.

The second task is equally important, however, in providing the coordination to prevent address space conflict. Registries, known as "whois" databases, are maintained by each of the RIRs. The whois database maintained by the RIPE NCC is called the RIPE Network Management Database, which is generally shortened to the *RIPE Database*.

As well as IP addresses, a whois database contains the names of organisations or customers to whom the addresses have been allocated, and related Points of Contact (POC). They also contain registration information for Autonomous System (AS) Numbers, a separate numbering system used in network-to-network routing.

Coordinated IP address distribution and publicly available databases of address registration information are important factors in the effort to ensure global uniqueness of addresses. The bottom-up coordination of such infrastructure, however, involves a variety of organisational structures, which vary from region to region. Part II of this document looks at these structures and how they operate in the RIPE NCC service region.

Part II: How Internet Address Management Is Conducted in the RIPE NCC Service Region

Chapter 4: Organisational Structures

Implicit in the adoption of bottom-up coordination is the understanding that all members of the Internet community must have the opportunity to participate in the development of IP address management policies.

RIPE

As discussed in Chapter 2, the Internet community in Europe, the Middle East and parts of Central Asia is represented by RIPE, a collaborative forum open to all parties interested in wide area IP networks in the region and beyond. Importantly, there are no membership requirements for participation in RIPE, and no membership fees; its activities are performed on a voluntary basis, with decisions made by community consensus.

RIPE conducts its coordination activities through a number of smaller bodies.

RIPE Working Groups

The majority of work done under the auspices of the RIPE community is carried out in *working groups*. Each working group focuses on a specific topic or area of interest, and each has one or more mailing lists where relevant topics and questions are discussed. They also hold face-to-face meetings twice a year, as part of RIPE Meetings.

As with RIPE generally, membership of these working groups is open and, for the most part, non-formal (though there are formally appointed Chairs and Co-chairs). New working groups can be formed with the consensus agreement of the RIPE community (or can close down if their area of interest ceases to be of relevance to the community), but they are designed to be ongoing forums for discussion.

RIPE Task Forces

RIPE Task Forces, on the other hand, are designed to complete a specific task or set of tasks. A task force is established by the RIPE community and is given a specific directive and timeframe. At the conclusion of its task, a task force will generally disband.

Anyone with an interest can volunteer, but the number of participants is usually limited, depending on the nature of the task. The outcome of a task force generally takes the form of a report with specific recommendations. These recommendations will be discussed by the RIPE community, and implemented if consensus can be reached.

RIPE NCC

The decisions and policies of the RIPE community are enacted through its secretariat, the RIPE NCC, which is an independent, not-for-profit organisation. As discussed in Chapter 3, the RIPE NCC also acts as the Regional Internet Registry (RIR), providing Internet resources (IPv4 and IPv6 addresses and AS Numbers) and related services to members in the RIPE NCC service region.

The RIPE NCC is funded by its membership, which is a subset of the RIPE community, specifically those community members who have obtained resources from the RIPE NCC. These members are referred to as *Local Internet Registries* (LIRs) and they are generally Internet Service Providers

(ISPs), telecommunication organisations and large corporations with significant presence in the region.

Global Coordination

The nature of the Internet means that there are many challenges whose solutions require global coordination. The RIPE community and the RIPE NCC offer a channel through which global discussion can be conducted, and global agreements reached when necessary.

These discussions happen in many ways, both formal and informal, and there are well-established processes within the RIR system for the development and implementation of global policies. Such policies require discussion and consensus agreement in all RIR service regions.

All of the RIRs are also members of the *Number Resource Organization* (NRO), a body through which they formalise their cooperative efforts. This can be useful when communicating and coordinating with other Internet organisations, or when engaging parties outside the traditional Internet community (for instance, the NRO has participated in the *World Summit on the Information Society* (WSIS) and *Internet Governance Forum* (IGF) processes).

ICANN and IANA

It is also through the RIRs (and sometimes the NRO) that the Internet community communicates with and participates in the "top level" of Internet coordination.

The *Internet Assigned Numbers Authority* (IANA) is a high-level body responsible for coordinating some of the key elements of the global Internet. Its responsibilities relate to some aspects of domain names, number resources and protocol assignments. The IANA is currently operated as one of the activities of the *Internet Corporation for Assigned Names and Numbers* (ICANN), an internationally-organised non-profit organisation mainly responsible for the stability of the DNS.

The IANA is responsible for the stewardship of the remaining pool of unallocated addresses. As such, the IANA distributes IP addresses to the RIRs. The global policies under which this distribution occurs are decided within the policy setting forums provided by ICANN, with input from a wide range of stakeholders. The RIRs in turn distribute those addresses to their communities.

The RIRs contribute to these processes both through direct communication with ICANN and IANA representatives at meetings around the world and through more formal relationships. The *NRO Number Council*, a body made up of three community members from each RIR service region, also fulfils the role of the *Address Supporting Organization* (ASO) *Address Council*, whose main tasks include overseeing recommendations on global IP address policy and the appointment of several Directors to the ICANN Board of Directors.

Chapter 5: Policy Development And Implementation

While policies guiding the management of IP addressing are vital to the successful operation and development of the Internet, it does not follow that the same policies will be appropriate in all regions. Different economic, geographic and historical factors mean that the requirements of the Internet community in each region can be very different.

For this reason, Internet policies of various kinds are agreed on at a regional level by the Internet community, and implemented by the RIRs. The broad definition of "Internet community", which includes ISPs, governments, regulatory bodies, network engineers, end users and anyone else with an interest, means that this system allows for all concerns to be addressed before a policy is agreed on.

Types of Policy

The policies developed and agreed on by the RIPE community are expressed in RIPE Documents, and fall into several broad categories (though some policies may fall under more than one of these):

- Documents relating to address policy and address management (including IPv6 documents)
- RIPE Database documents
- RIPE NCC organisational documents
- Information Services documents
- Request forms and supporting notes

All RIPE Documents are accessible to anyone on the RIPE website:

<http://www.ripe.net/ripe/docs/index.html>

Address Allocation Policies

Policies governing the allocation and assignment of IP addresses and AS Numbers are central to the RIPE NCC's role as the Regional Internet Registry, and central to the overall goal of global uniqueness in addressing. These policies are created and implemented to fulfil four major requirements:

- Uniqueness: All public IP addresses worldwide must be unique. This is an absolute requirement guaranteeing that every host on the Internet can be uniquely identified.
- Aggregation: The distribution of IP addresses in a hierarchical manner permits the aggregation of routing information, which helps to ensure proper operation of Internet routing.
- Conservation: Public IP address space must be fairly distributed to organisations that operate networks and can demonstrate a legitimate need.
- Registration: The provision of a public registry documenting resource allocations and assignments must exist to ensure uniqueness and to provide information for Internet troubleshooting at all levels.

These goals will be represented variously in policies covering IPv4 addresses, IPv6 addresses and AS Numbers. Separate policies are necessary to meet the varying technical and operational needs of the different protocols.

Database Policies

Responsibility for the RIPE Database is another central role of the RIPE NCC, and the community must agree on any substantive decisions about the operation of this database. A Database Working Group exists for discussion of database-related issues, and a RIPE Document relating to the database can be viewed at:

<http://www.ripe.net/ripe/docs/ripe-419.html>

Best Practice Statements

Some RIPE Documents look at Current Best Practices in relation to various aspects of network operation. By helping to foster practices that are recognised as efficient and useful for network trouble-shooting, these statements are useful to the whole Internet community.

The Limits of Addressing Policy

It is important to note the limits of IP addressing policy. As noted earlier, the fundamental design of the Internet means that anyone can at any point "announce" any address they choose to. IP addressing policies cannot change this fact, but at the same time it is in no one's (legitimate) interest to make the Internet unusable.

IP addressing policy, as it is currently developed, can take into account the needs and concerns of all sectors of the Internet community. This results in policies that produce the best result for the Internet itself, and for the community at large.

RIPE Policy Development Process

The RIPE community creates policy according to a well-defined process. This *Policy Development Process* (PDP) is laid out in a RIPE Document, and is available at:

<http://www.ripe.net/ripe/docs/pdp.html>

The RIPE PDP is designed to allow all interested parties to have input into the decision making process, whether through the RIPE Meetings or working group mailing lists. A policy that has come through the PDP will have had a chance to be considered by all interested parties, meaning that the final result is widely regarded as being in the best interests of the broad Internet community.

RIPE NCC Activity Plan

In its role of implementing the RIPE community policies, it is important that the RIPE NCC be seen to be fair, unbiased and transparent. The RIPE NCC Activity Plan is an important way in which this is demonstrated. It publicly lays out the plans and priorities for the RIPE NCC over the coming year, and all

members of the RIPE community have the chance to comment on it. At the end of this process, the RIPE NCC Board, who are elected by RIPE NCC's membership, will decide whether or not to approve the Activity Plan for the coming year.

Part III: Enhanced Cooperation

Chapter 6: The Need For Enhanced Cooperation

"Enhanced cooperation" is a term that was coined during the World Summit on the Information Society (WSIS) in 2005. It refers to the developing relationships between all stakeholders in the "Information Society", particularly between the private and public sectors. It is an issue that has taken on increased prominence in recent years, as the wider Internet community strives to ensure that all stakeholders' voices are heard and all interests served.

As detailed in the first two parts of this document, the existing structures of Internet governance (including RIPE and the global RIR system) have evolved over the past decades to meet the very specific challenges that the Internet poses. While the RIPE community therefore welcomes the increased participation of all Information Society stakeholders, it is important that these new stakeholders participate with an understanding of the ramifications that even local IP addressing policy can have on the global Internet.

For the RIPE community then, enhanced cooperation is an important opportunity to improve communication with those stakeholders in the Internet community who do not tend to participate in the standard RIPE forums. This especially includes government and regulatory bodies (the public sector), who in recent years have taken an increasing interest in issues of Internet governance.

It is also a chance to educate these other stakeholders on the existing systems, and why we believe they should be maintained and strengthened. The reasons for this belief include the following points:

Provide Assurance of Operational Stability

The RIR system emerged in part to address the need for a stable, reliable means of controlling IP address distribution and management globally. The structure of the Internet itself relies heavily on such stability and certainty. Factors such as the depletion of unallocated IPv4 address space and the increasing involvement of public sector organisations, however, mean that the Internet industry is facing a period of change.

In this environment, the need for open and clear dialogue between all stakeholders, and especially between the public and private sectors, is vital to ensuring the ongoing operational stability of the Internet's infrastructure.

Provide Assurance of Good Policy

While the RIR policy development system is designed to facilitate input from all interested parties, this goal can only be achieved when all parties are

aware of the process. In the past, this has meant that some parties with an interest in Internet addressing policy have not taken part in policy development.

One of the goals of enhanced cooperation is to increase awareness of and participation in the existing policy development processes, which is vital to the final goal of creating policies that are in the best interests of all parties. These processes allow for input from various stakeholders through a range of channels, not just RIR meetings.

Maintain Confidence in Institutions and Processes

The RIPE community believes that the existing system of institutions and processes for Internet policy development has proved itself to be fair, open, flexible and efficient. As such, it is well positioned to meet the challenges of policy development in a rapidly changing industry.

Enhanced cooperation is an important means of disseminating knowledge of the existing system, its institutions and processes, and of building confidence in that system's ability to meet the needs of all stakeholders. Confidence in these institutions and processes is vital to ensuring wider participation in the development of policy, and general respect for the policies produced.

Chapter 7: Consultation Processes

The RIPE NCC is already involved in a range of activities that fall under the description of enhanced cooperation. These include activities within the official policy development process, as well as outreach activities described in the Activity Plan.

RIPE Meetings

RIPE Meetings are some of the biggest undertakings of the RIPE NCC, and happen twice a year. As well as playing a central role in the policy making process, the meetings are a chance for members of the Internet community to gather and meet. Most importantly, RIPE Meetings (like all RIR meetings) are open to everyone, and therefore offer a unique opportunity to further enhanced cooperation through face-to-face contact between all sectors of the Information Society.

RIPE NCC Government Relations

The RIPE NCC is also taking an active role in reaching out to the public sector, including specific governments. With governments taking an increasing interest in Internet policy, this can often mean simply acting in an advisory role at government-organised events.

RIPE NCC Government Roundtables

The RIPE NCC is also pro-actively engaging the public sector through Government Roundtable meetings, at which are discussed Internet management issues relevant to governments, regulators and industry partners. These events provide a chance for attendees to learn more about how to participate in IP address management policy-making. High-level

discussions of IPv4/IPv6 address space and root name servers (one of which is operated by RIPE NCC) also provide attendees with an overview of the main elements involved in the technical coordination of the Internet.

Other Processes

The process of enhanced cooperation is ongoing, and will include activities beyond those discussed here. Such activities are often spearheaded by the RIPE NCC, but may involve contributions from other members of the RIPE community.

It is also important that the RIPE community be aware of the activities being undertaken in its name, and it is for this purpose that this Task Force recommends the formation of a Cooperation Working Group.

Chapter 8: Learning Points

What Works

Enhanced cooperation with the public sector has been a priority for the RIPE NCC since the lead-up to the WSIS events. This priority has been reflected in events such as RIPE NCC Government Roundtables and an increased presence at various government-organised events.

From this experience, it is clear that the best results in this area are produced by targeted activities. This can mean engaging the public sector participants in their own forums, or it can mean organising specific events beyond the traditional RIR event schedule.

What Doesn't Work

It is clear that the conventional channels and processes of the Internet community are not, of themselves, sufficient to meet the demands of enhanced cooperation. There is a range of reasons why interested parties outside the traditional RIPE community have not taken the opportunity to participate in forums such as the RIPE Meetings or mailing lists, but it is clear that if Internet policy is to have any authority, the policy development process must engage with these parties.

"Business as usual" in this case will not work. With the addition of targeted outreach activities and a flexible approach, however, the existing processes and institutions can still meet the needs of enhanced cooperation.

What Is Not Yet Known

At this point, the future of Internet addressing policy remains unclear. Enhanced cooperation is a broad strategy to ensure that as the Internet community adapts to the changes ahead, the needs of all stakeholders, both private and public sector, are recognised and reflected in Internet policy.

At the same time, it is also vital that the evolving policy development processes not hinder the ongoing operation of the Internet and development and innovation in the Internet industry.

Recommendations

This Task Force notes the ongoing importance of enhanced cooperation to the RIPE community and recommends that the RIPE community form a Cooperation Working Group, with the following charter:

The Cooperation Working Group is a forum for discussion focusing on cooperation between the private and public sectors on Internet matters. This kind of cooperation has taken on increased prominence in recent years, as the wider Internet community strives to ensure that all voices are heard and the interests of all parties are considered. Fostering more open dialogue between all stakeholders is vital to ensuring the continued stability of the Internet.

The working group discusses the following:

- 1. The working group will primarily discuss outreach from the traditional RIPE community to everyone else, especially governments, regulators and NGOs, all of whom we are trying to bring into our community. Topics are not to duplicate issues discussed in other working groups. This working group should complement the other working groups and help participants engage in appropriate work.*
- 2. The RIPE NCC's current outreach activities will be reported, and the RIPE NCC will seek advice and guidance on future activities. This is to make the discussions more focused - currently the only forum for these discussions is the ripe-list mailing list.*
- 3. The working group will develop and clarify the RIPE community's position on issues that are of relevance to the public sector or on which a community position has been sought.*
- 4. The working group will be responsible for maintenance of the RIPE Document produced by the Enhanced Cooperation Task Force, describing the RIPE community, existing policy development processes and outreach programs. The working group explicitly does not have change control over the RIPE Policy Development Process (PDP) itself.*

The working group is also an important channel through which the RIPE community can communicate with others in the Information Society. The Chairs are not to become special Ambassadors for RIPE. Their role is the same as other RIPE Working Group chairs, which implies they of course could be asked now and then what the status of the working group is. The process by which RIPE and RIPE NCC respectively coordinate with other bodies (such as the NRO) and communicate (mostly via RIPE NCC or the chair of RIPE) is not changed by creation and existence of this working group.