

# **Development of OSIX – A Peering Point in a Box**

By

Aun Haji Iqbal

A thesis  
submitted to the Victoria University of Wellington  
in fulfilment of the  
requirements for the degree of  
Masters of Engineering  
in Network Engineering.

Victoria University of Wellington  
2013



## **Abstract**

Sending traffic over international communication links is much more expensive than sending traffic locally. Unfortunately, there are situations where two local networks end up using the international links because there are no national links between the two networks. To avoid this, a peering relationship should be established between these two local networks to allow them to directly exchange traffic. Peering relationships are implemented at Internet Exchange Points (IXPs).

There has been a significant increase in the number of IXPs in developed countries but take up in developing countries has been slow despite these countries having the most to gain due to the high prices that they pay for international bandwidth. Research has identified that lack of technical skills is a key barrier to the deployment of IXPs in these countries. In particular, although skills exist to maintain an IXP there is a lack of technical expertise to integrate individual tools together to implement an IXP.

The goal of this thesis is to develop an integrated IXP solution that could be easily deployed in developing countries. This content of this thesis includes analysis of the requirements for such a solution, development of a design, description of implementation trade-offs and evaluation of a final solution.



# Contents

<b>1. Introduction.....</b>	<b>1</b>
1.1 Contributions.....	4
1.2 Thesis Organization .....	4
<b>2. Background .....</b>	<b>7</b>
2.1 Principles of Peering .....	7
2.1.1 Why Peer? .....	8
2.1.2 Why some ISPs avoid Peering? .....	8
2.1.3 Peering Negotiation & Implementation .....	9
2.2 IXPs in Existence .....	10
2.3 Identified limitations of Peering at IXPs .....	11
2.4 Summary .....	12
<b>3. Analysis of New Zealand IXPs.....</b>	<b>15</b>
3.1 New Zealand Internet Setup.....	15
3.1.2 Formation of New Zealand Internet Infrastructure .....	16
3.1.3 Domestic Links .....	17
3.1.4 Communication with Rest-of-the-World .....	18
3.1.5 New Zealand IXPs .....	20
3.1.6 Why should non-ISP Peer? .....	22
3.1.7 The Current Peering Scenario .....	23
3.1.8 Domestic Transit .....	23
3.1.9 International Experience .....	24
3.1.10 Cost of Joining a Peering Point.....	25
3.2 Quantitative Evaluation .....	27
3.2.1 APBDC Model.....	27
3.2.2 Bandwidth Speed & Delay Comparisons .....	30
3.2.3 Bandwidth & Time Delay Graph Model .....	38
<b>4. Analysis and Design .....</b>	<b>41</b>
4.1 Requirements Gathering .....	41
4.2 Design Considerations .....	42
4.2.1 Layer 2 IXP.....	43
4.2.2 Layer 3 IXP.....	44
4.2.3 Layer 2 versus Layer 3 IXP .....	44
4.2.4 Physical Design Considerations.....	45
4.2.5 Peering Design Considerations .....	45
4.2.6 Routing Design considerations .....	46
4.2.7 IP Address Space Design considerations .....	46
4.2.8 Critical Infrastructure .....	47
4.3 IXP Structure Design .....	49
4.3.1 Basic Peering .....	50
4.3.2 Building OSIX .....	51
4.3.3 How it works at OSIX?.....	53
4.4 Route Server Design .....	54
4.5 Policy Design .....	57
4.6 Interface Design .....	60
<b>5. Implementation .....</b>	<b>63</b>
5.1 Implementation Based on Design .....	63
5.2 Operating System Selection and Implementation.....	66

5.3 Accessing OSIX.....	67
5.4 Route Server Implementation .....	67
5.5 Policy Implementation .....	68
5.6 Web Server Implementation .....	69
<b>6. OSIX Testing .....</b>	<b>75</b>
6.1 Test Infrastructure .....	75
6.2 Test Cases .....	75
6.3 Test Cases Summary .....	89
<b>7. Conclusions.....</b>	<b>91</b>
7.1 Major Contributions.....	91
7.2 Future Work.....	92
<b>8. Appendices.....</b>	<b>95</b>
Appendix A – IXP Fact Finding Survey.....	95
Appendix B – IXP Survey Email.....	99
Appendix C – IXP Survey Results.....	101
Appendix D – GPL License .....	113
Appendix E – VUW Human Ethics Policy.....	127
Appendix F – OSIX Terms & Conditions .....	129
Appendix G – Test Cases Results .....	131
Bibliography .....	148

## **Acknowledgments**

I like to thank my supervisors, Dr. Ian Welch and Andy Linton, for their guidance and support throughout this thesis.



## List of Figures

Figure 1 – Peering Breakeven Point & Cost Model [20].....	10
Figure 2 – Domestic Bandwidth Percentage (July 2011 – July 2012) [22].....	11
Figure 3 – NZ Internet users as percentage of population [24] .....	16
Figure 4 – Topological View of ISPs in New Zealand [25] .....	17
Figure 5 – SCCN Managed Fibre Connections [26].....	19
Figure 6 – SCCN Segments Length & Transmission Delays [26] .....	19
Figure 7 – Average Traffic Throughput at the DE-CIX [30].....	25
Figure 8 – Cisco Hardware Cost Reduction Comparison [7] .....	26
Figure 9 – Downloading file hosted at NZ Server from VUW.....	32
Figure 10 – Time Graph of file hosted at NZ Server from VUW.....	32
Figure 11 – Downloading file hosted at UK Server from VUW .....	33
Figure 12 – Time Graph of file hosted at UK Server from VUW .....	33
Figure 13 – Downloading file hosted at NZ Server from non-peering location.....	36
Figure 14 – Time Graph of file hosted at NZ Server from non-peering location.....	36
Figure 15 – Downloading file hosted at UK Server from non-peering location .....	37
Figure 16 – Time Graph of file hosted at UK Server from non-peering location .....	37
Figure 17 – Bandwidth Time and Delay Graph from peering location .....	38
Figure 18 – Bandwidth Time and Delay Graph from non-peering location.....	39
Figure 19 – Layer 2 IXP Model.....	43
Figure 20 – Layer 3 IXP Model.....	44
Figure 21 – Quagga System Architecture .....	54
Figure 22 – OSIX Physical Layout .....	56
Figure 23 – OSIX Logical Layout .....	57
Figure 24 – OSIX OS Layer Process View .....	59
Figure 25 – OSIX Process Flow Diagram .....	60



## List of Tables

Table 1 – IXP Growth [21, 22] .....	10
Table 2 – Accessing OSIX.....	67
Table 3 – Test Cases Summary.....	89



## Listings

Listing 1 – RTT Delay to NZ and UK Servers from VUW .....	31
Listing 2 – Hops to NZ and UK Servers from VUW .....	31
Listing 3 – RTT Delay to NZ and UK Servers from non-peering location .....	34
Listing 4 – Hops to NZ and UK Servers from non-peering location.....	35
Listing 5 – Linux Kernel Running on OSIX Servers.....	66



# 1. Introduction

At present most of the developing countries do not deploy domestic traffic peering, which means that traffic between the local Internet Service Providers (ISPs) is exchanged through international links. In the absence of the domestic peering in these countries, then an ISP must send all outbound traffic through its international links, most commonly via satellite and occasionally via submarine fibre [1].

International links entail both upstream and downstream packet traffic, the costs of which must be borne by either the sending or the receiving ISP. Here, we observe a troubling imbalance: Unlike in the telephony world, where ITU-mandated rules require that the costs of international calls be shared 50/50 between telecom operators, international Internet connectivity operates according to the peering/transit dichotomy. ISPs are not subject to the ITU's cost sharing rules; rather, connectivity costs are allocated according to bilateral contracts, which can generally be classified as either peering or transit agreements [15].

Peering is the direct exchange of mutual customer routing information and traffic between independent autonomous systems (AS) networks [16]. Peering involves two networks coming together to exchange traffic with each other freely, and for mutual benefit. This 'mutual benefit' is most often the motivation behind peering, which is often described solely by "reduced costs for transit services". Other less tangible motivations can include:

- Increased redundancy (by reducing dependence on one or more transit providers).
- Increased capacity for extremely large amounts of traffic (distributing traffic across many networks).
- Increased routing control over your traffic.
- Improved performance (attempting to bypass potential bottlenecks with a "direct" path).
- Improved perception of your network (being able to claim a "higher tier").

- Ease of requesting for emergency aid (from friendly peers).

Besides mutually benefiting the peering ISPs, peering has wider benefits. For example, peering encourage the development of local content and applications. Once peering is established, it becomes a natural location to host a variety of other services that reduce bandwidth requirements and improve the speed and reliability of Internet access for local users. Such improvement of the access speed for local content often results in incentives for local developers to produce local content and applications [2].

Also, it often encourages international content providers to establish themselves in the country. For example, after Argentina and Kenya started domestic peering, Google started hosting its services in those countries. That not only created employment opportunities in those countries but also improved access speed to Google's services. The development of local content and applications will also make the Internet more relevant to the local population, which makes the Internet more socially and economically beneficial to the country [14, 17].

Peering is implemented using Internet Exchange Points (IXPs). These components of Internet infrastructure that can increase the affordability and quality of the Internet by enabling local networks to efficiently exchange information at a common point within a country rather than needing to exchange local Internet traffic overseas. In many of the developing countries, for example, Internet messages need to be exchanged beyond their border, which adds significant costs because of lack of connectivity between domestic networks [17].

IXPs can be established with relatively minimal equipment and overhead costs. IXP are key elements of Internet Infrastructure which facilitates interconnection between Internet-based networks, which creates the potential for a range of technical and economic benefits for the local Internet community [3].

In a typical IXP-based network, the IXP is operated on a switching platform, which is used to interconnect the different service providers or subscribers. It is simply a case of using a standard layer 2-switch that can be shared among different subscribers for peering via bi-lateral agreement between their routers to exchange routes [4].

Autonomous systems (AS) in the Internet today exchange interdomain routing information through Border Gateway Protocol (BGP). Unlike interior gateway protocols, which periodically flood an intradomain network with all known topological information, BGP is an incremental protocol that sends update information only upon changes in network topology or routing policy. Routing information shared among BGP speaking peers has two forms announcements and withdrawals [9].

When two members decide to peer, they are only required to establish a BGP session between their routers. Since this requires both routers have interfaces in the same IP subnet, the IXP will assign an IP address to the IXP-facing router interface of each of its members from the IP prefix(es) allocated to the IXP by Internet Registry responsible for the IXP [18].

An Internet Exchange Point is usually a location operated by a single entity to allow the exchange of Internet traffic between three or more ISPs. An IXP is also known by its neutrality among all subscribers and it is preferably administered by a non-profit organization or association [5].

The barriers to establishing IXPs in countries where they do not yet exist are largely non-financial: there is often a lack of mutual appreciation of benefits among all stakeholders, as well as resistance from those providers with market dominance. In addition, limited technical skills and a lack of open competitive markets in telecommunication and Internet services make it more difficult to establish an IXP.

The problem of lack of open competitive markets is beyond the scope of this thesis instead the focus is on lowering the amount of technical knowledge required to implement an IXP.

## 1.1 Contributions

The main goal of this thesis is the design and implementation of an IXP, which can be easily deployed anywhere in the world without having any licensing or technical constraints. This thesis makes the following contributions:

- Evaluation of the need of peering and how it can benefit the ISP to peer locally where possible to avoid International transit cost and delays.
- Development of a set of requirements for an integrated solution for IXPs through review of international surveys on peering and our own survey of international IXPs.
- Design and implementation of an IXP integrated solution from components that meet the requirements identified from surveys and the literature.
- Evaluation of through testing of the solution using a laboratory test bed.

## 1.2 Thesis Organization

The following chapter describes the background of this project, providing an overview of IXP's peering issues, different kinds of peering and in what scenario peering may not be feasible.

- Chapter 2 talks about the background of IXP and its reasoning.
- Chapter 3 we talk about the New Zealand Peering Infrastructure and quantitative evaluation of Internet traffic on both domestic and International links. This analysis is intended to be used when promoting the benefits of peering and to show why it is worthwhile in a wider sense.
- Chapter 4 covers the project design and development phases of the project.
- Chapter 5 presents the steps followed during the implementation of the Internet Exchange Point.

- Chapter 6 provides verification and validation of test cases in building the OSIX.
- Finally, Chapter 7 provides the conclusion and suggests possible future improvements for the Internet Exchange Points.



## 2. Background

In this chapter, we will discuss the background of peering at IXPs.

### 2.1 Principles of Peering

The Internet is a collection of separate and distinct networks, each one operating under a common framework of globally unique IP addressing and global BGP routing.

The relationships between these networks are generally described by one of the following three categories:

- Transit (or pay) - You pay money (or settlement) to another network for Internet access (or transit).
- Peer (or swap) - Two networks exchange traffic between each other's customers freely, and for mutual benefit.
- Customer (or sell) - Another network pays you money to provide them with Internet access.

Furthermore, in order for a network to reach any specific other network on the Internet, it must either:

- Sell transit (or Internet access) service to that network (making them a 'customer'),
- Peer directly with that network, or with a network who sells transit service to that network, or
- Pay another network for transit service, where that other network must in turn also sell, peer, or pay for access.

The Internet is based on the principle of global reachability (sometimes called end-to-end reachability), which means that any Internet user can reach any other Internet user as though they were on the same network. Therefore, any Internet connected network must by definition either pay another network for transit, or peer with every other network that also does not purchase transit [6] [8].

### **2.1.1 Why Peer?**

There are many reasons for Tier 2 ISPs to peer together:

- The most important reason for ISPs to peer is to reduce the transit. It is very expensive to purchase transit from a Tier 1 ISP. On the other hand transferring data can come at no cost through peering.
- The two ISPs peering together do not need Upstream Transit to communicate and transfer data.
- Since an intermediary (Tier 1 ISP) is removed from the communication chain and the interconnection is more direct between the ISPs peering customers. This also reduces traffic latency and thus provides better performance to the peering ISPs customers.
- Several ISPs (both Tier 1 and Tier 2) tends to charge their customers based on the volume of data they use. Therefore, ISP interconnection through peering offers better performances including a diminution of packet loss. This is also a reason why occasionally some Tier 2 ISPs can provide a cheaper rate than tier 1 ISPs.
- By reducing the packet loss, it also enables ISPs to use more effectively their network and ensure as much bandwidth as possible to their customers.

### **2.1.2 Why some ISPs avoid Peering?**

As we discussed above that peering presents several advantages, but there are still many possible minor issues which can lead to conflict-of-interest and can stop peering agreement to be formed between ISPs.

- Lack of symmetry in the traffic and in the investments can be a major issue when working on a peering agreement. If one of the two ISPs is to send more traffic than the other or if the costs resulting from peering are mainly supported by one of the two parties, the ISPs will have difficulties in peering together.
- Some important resources are needed to apply and maintain peering between ISPs. Manpower is required to work on the peering agreement, the connection has to be established, and the routing tables need to be changed, it is a time-consuming process.

- Some equipment's costs cannot be avoided such as router or additional line-cards.
- Sometime depending on bilateral or multilateral peering, the time and financial resources needed to peer could be seen as more effective if traffic is simply transferred via Tier 1 ISP uplinks.
- As peering is a mutual agreement, it is difficult to give assurances in case of problems on the connection. Some ISPs would rather pay expensive transit costs and obtain guarantees about the reparation of outages than peering and relying on peering agreement clauses.
- Finally, ISPs do not want to commoditize data transit. Most of the major ISPs are competing on the quality of their service and the quality of their network interconnection. By providing access to other ISP through peering, major ISPs reduce their competitive advantage as they would improve competitor's performance by carrying their traffic [19].

### 2.1.3 Peering Negotiation & Implementation

The price of Internet Transit is relevant to the Peering Community since the alternative to peering is to simply send traffic to an upstream transit provider. So when doing the financial analysis of peering versus simply sending the traffic to an upstream ISP, Peering Coordinators do a side-by-side comparison between the two.

**Cost of Peering:** Peering costs generally include at least a fixed monthly fee for a peering port on a shared peering fabric, and in some cases a membership fee. In this first (simple) model we will assume that the ISP doing the analysis is already at a co-location site that has a public peering fabric. For the first model, let's assume the cost of a 10G port, membership fees, and other miscellaneous expenses add up to \$4,500 per month.

The figure below shows the costs of send traffic through peering or transit services in 2009, assuming you are already at a co-location with an IXP.

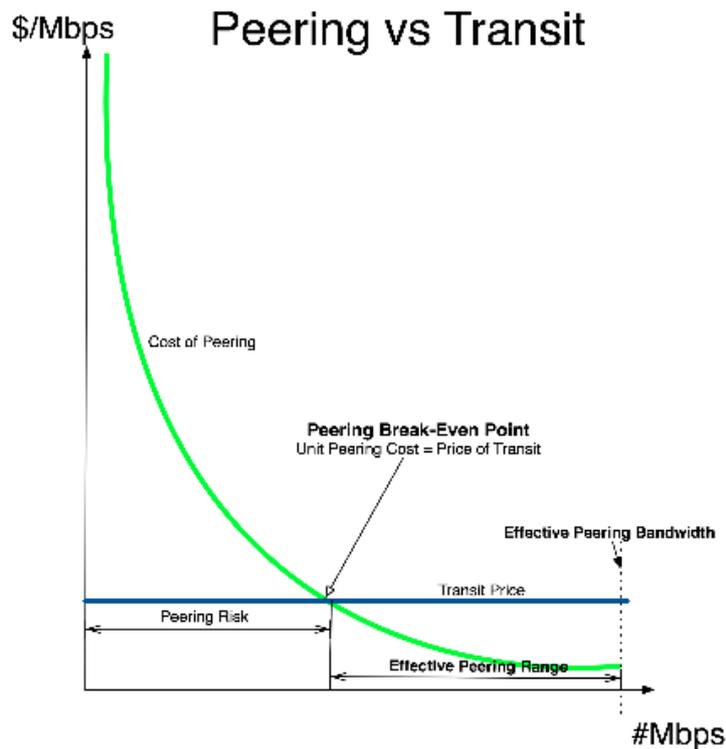


Figure 1 – Peering Breakeven Point & Cost Model [20]

The peering and transit lines intersect at a point called the Peering Breakeven Point, the point where an ISP is indifferent between peering and buying transit. Beyond this point, peering can save the ISP money, up until the capacity of the 10G port. Here we see the breakeven point about 500Mbps [20].

## 2.2 IXPs in Existence

As of 2012, there are than 361 IXPs have so far been set up worldwide, see Table 1. There is an increase of 17.45 per cent since 2007. Regionally, Latin America has experienced the fastest recent growth in numbers of IXPs, reaching 20 by the end of 2007 and the number incremented to 34 in 2012. However, developing countries have generally lagged behind the rest of the world in establishing IXPs. The Asia-Pacific region grew the slowest in 2007, bringing the total number of IXPs in that region to only 67. Africa has the fewest IXPs - only 17 nations had IXPs in 2007 - and growth reached to only 21 in 2012 [21, 22].

Region	No. of IXPs in 2007	No. of IXPs in 2009	No. of IXPs in 2010	No. of IXPs in 2012
Africa	17	22	22	21
Asia-Pacific	67	72	72	77
Europe	107	127	130	141
Latin America	20	29	31	34
North America	87	88	88	88
Total	298	338	343	361

Table 1 – IXP Growth [21, 22]

This project provides an introduction to IXPs by outlining their role as a key component of Internet infrastructure and covering all the essential elements of an open-source IXP solution like, the routing policy, routing database, backup mechanism of customer routers & IXP switch and graphical view of traffic, which must be considered in establishment of IXP Management server.

The following figure [22] gives a brief statistics on IXP and its domestic bandwidth production as a comparison between 2011 and 2012.

Region	Internet Exchange Points				Domestic Bandwidth Production			
	Jul 2011	Jul 2012	Net Change	Percent Change	Jul 2011	Jul 2012	Net Change	Percent Change
North America	86	88	+2	+2%	878G	984G	+106G	+12%
Asia-Pacific	76	77	+1	+1%	1.19T	1.39T	+193G	+16%
Europe	139	141	+2	+1%	5.85T	8.18T	+2.33T	+40%
Africa	21	21			4.29G	5.62G	+1.33G	+31%
Latin America	34	34			77.4G	130G	+52.1G	+67%
<b>Total</b>	<b>356</b>	<b>361</b>	<b>+5</b>	<b>+1%</b>	<b>8T</b>	<b>10.7T</b>	<b>+2.68T</b>	<b>+25%</b>

Figure 2 – Domestic Bandwidth Percentage (July 2011 – July 2012) [22]

The above figure show major differences in the growth of IXPs in developing countries and highly-developed countries. For example, Africa being a very large continent and having more than 53 independent nations have only 21 IXPs, while Asia being the most populated continent in the world have only 77 IXPs. On the other hand, Europe alone has more IXPs than combined numbers of Asia and Africa.

### 2.3 Identified limitations of Peering at IXPs

According to Internet Governance Forum, some of the challenges to IXP development are briefly explained below [23]:

**Lack of Trust between Service Providers:** IXPs are non-profit and they depend on their participants to cooperate and coordinate to be effective. Building trust and emphasizing neutrality and mutual benefits were underscored as essential in order to bringing parties together to establish an exchange point.

**Limited Technical Expertise:** The success of an IXP hinges on its ability to route traffic in an efficient, cost effective manner. This requires competent engineers to implement and support day-to-day operations at both the participating ISPs and the IXP switching facility.

**Cost of Network Infrastructure:** The absence of reliable and affordable local infrastructure can reduce the incentive and justification for operators to develop and connect to an IXP. In many countries, purchasing a domestic leased line across a city or region can be more expensive than sending traffic through an international link.

**Cost of Hosting an IXP in a Neutral Location:** The cost of operating IXP infrastructure in an appropriate, neutral facility can present challenges. In many countries, costs associated with leasing space, ensuring reliable power supply, providing adequate air-conditioning, security, and hiring IXP maintenance staff can outweigh the savings that participants might realize from its operation.

## 2.4 Summary

The case for IXPs is compelling, and the obstacles relatively clear and well-understood. In order to achieve wider IXP deployment in the developing world, the following issues must be addressed:

- Regulatory reform and liberalization,
- the overcoming of monopoly telecom resistance, and
- The organization of competitive ISPs into associations capable of neutrally administering shared facilities on behalf of their members.
- Technical expertise.

Both regulators and competitors need to be convinced and fully aware of the benefits of domestic Internet traffic exchange, and of the broader proposition that communication in a developing economy is very crucial. Lower costs for competitors can lead to greater revenues for all, stronger investment from abroad, and lower-cost, higher-quality services for all users [15].

My focus is on reducing the technical barriers. At present IXPs are built out of separate components and requires expertise to integrate together.

My solution is an integrated solution as so far there is no commercial or open-source solution available as of this writing.

It is very hard to build an IXP “in a box” from scratch as it involves many components and libraries to build a fully working solution.

The solution we proposed is completely open-source under GNU license and thus there will be no licensing overheads for the deployment.



### **3. Analysis of New Zealand IXPs**

This chapter analyses the New Zealand Internet Infrastructure and we are these analysis to prove a success-story of IXPs development in New Zealand. The initial analyses on New Zealand Internet were conducted by Neil Bertram [12]. Our analyses are an extension to what has been done before. We have chosen New Zealand as our case study because we have access to most of the information on domestic links and transit, even though our working solution, OSIX (Open Source Internet Exchange) can be deployed anywhere in the world.

#### **3.1 New Zealand Internet Setup**

New Zealand has one of the highest Internet access per head of population in the world and 83% [24] of the population (as of 2010) has some kind of access to the Internet. However, the continuous incrementing need to consume rich content applications like video-on-demand is adding additional burden on the International transit links that connect New Zealand to the rest of the world. In near future this likely to become an even bigger problem, as services such as Video on Demand, online music and movie purchases and majority of other high-bandwidth content are hosted overseas. New Zealand customers are already unhappy with the bandwidth provided for International links and they are also annoyed by the average speed and high prices due to the long-haul distance the content must be transferred.

Interestingly, it seems more appropriate for high-bandwidth-intense applications will have to be hosted within New Zealand's infrastructure to serve domestic clients, and possibly even hosted on regional IXPs.

The figure below shows the graph of increasing number of Internet users in New Zealand until 2010 [24].

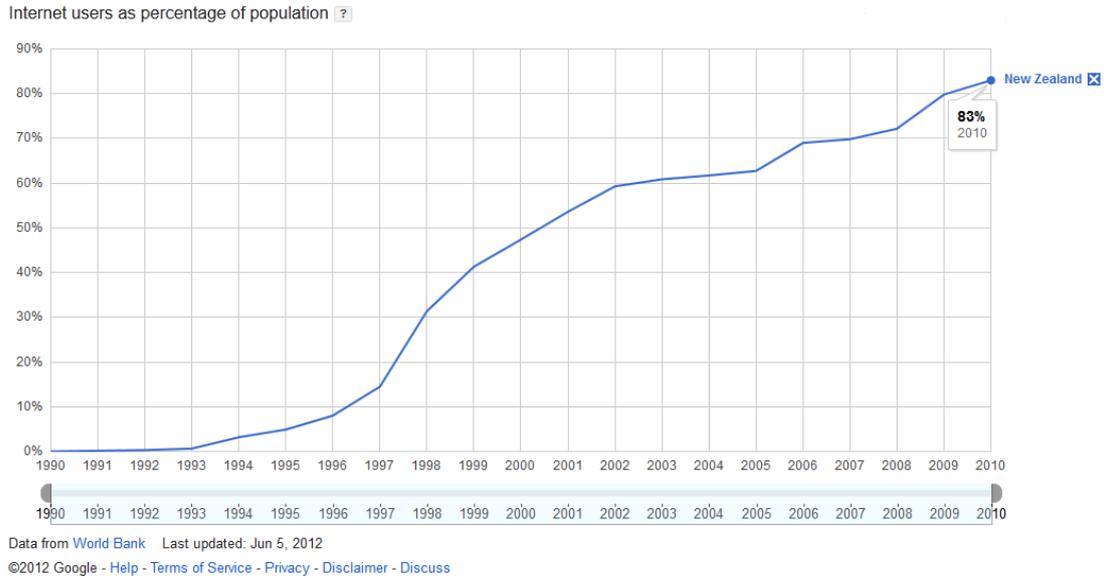


Figure 3 – NZ Internet users as percentage of population [24]

### 3.1.2 Formation of New Zealand Internet Infrastructure

At a basic level, New Zealand’s Internet infrastructure is divided into two groups of participating entities, those that peer and those that do not. The two main players in Internet Transit, or Tier 1 ISP are Telstraclear and Telecom NZ, also provide services to residential subscribers. The third emerging player, FX Networks only provides corporate services to Tier 2 ISPs and other corporations. The growth of non-peered providers is happened largely during 2004, and has given rise to a two-tier architecture that is similar to many other countries. The only thing that makes New Zealand slightly different from many other countries is that majority Internet users subscribe to an Internet services provided by the Tier-2 ISPs, rather than the transit providers purely providing transit.

The Figure below show the topological view of ISPs in New Zealand [25]:

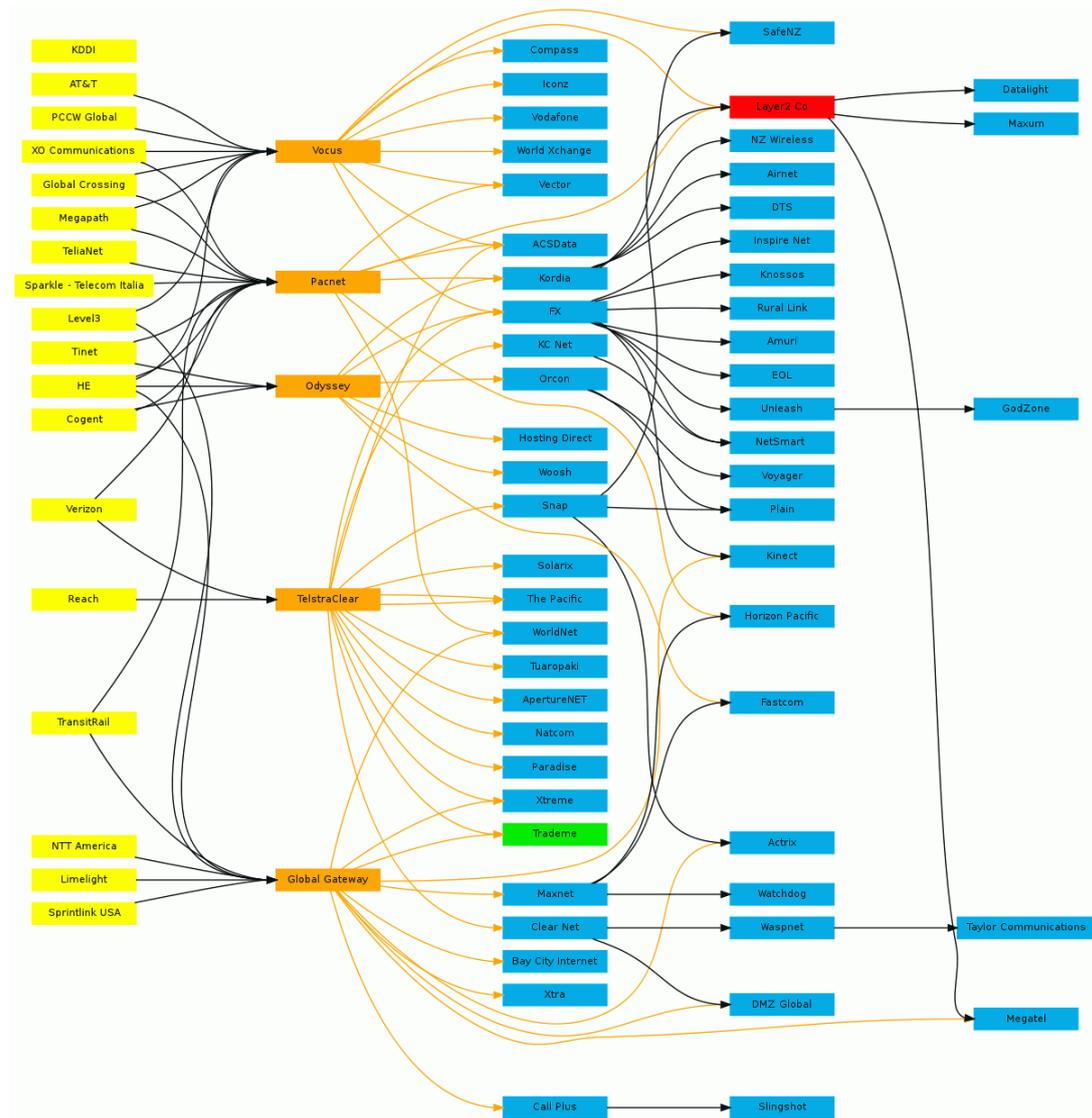


Figure 4 – Topological View of ISPs in New Zealand [25]

### 3.1.3 Domestic Links

In 90's, New Zealand had a massive deployment of data transit services that were IP-capable, mainly triggered by the introduction of competition between the NZ Telecom and newcomer Clear Communications, now owned by Telstraclear. Both providers built high-capacity nationwide fibre networks to provide both voice and data services.

As the general direction of telecommunications around the world has moved towards IP convergence and replacing legacy systems with those based solely on IP, much

investment has been poured into strengthening domestic backbones to handle large amounts of IP traffic over ATM and MPLS networks.

Several other providers and in particular FX Networks have also built national and regional backbones capable to transit IP data.

The big deployment of data networks has eventually provided free-of-cost national transit between any two (or more) points-of-presence. The Tier-2 ISP networks interconnected peering points as well as form the upstream link to Tier-1 ISP backbones.

The national transit networks in New Zealand have fibre networks capable of connecting businesses and ISPs to national backbones. In big cities, inexpensive metropolitan Ethernet service provisioned over fibre is available at speeds of up to 10Gbps from providers such as Vector and Citylink, providing an ideal environment for local peering points or transit services.

### **3.1.4 Communication with Rest-of-the-World**

Geographically New Zealand is remote to most of the popular content around the world and this cause International connectivity a major concern. Most data transit traffic out of the New Zealand is provided by international transit providers, such as Reach, Vocus and global Gateway, leasing bandwidth on the Southern Cross Cable Network (SCCN), which is currently delivering 295 Gigabit/s of fully protected bandwidth terminated in Takapuna (Auckland) and has the potential as demand growth requires to increase to 1.2 Tbps using the existing 10 Gbps technology or 4.8 Tbps using 40 Gbps. This scarcity of bandwidth, as well as the need to recoup the large investment that was required to build the network, prices for International data transit is relatively high.

Figure below shows the SCCN managed fibre connection from USA (Hawaii) to Australia via New Zealand [26].

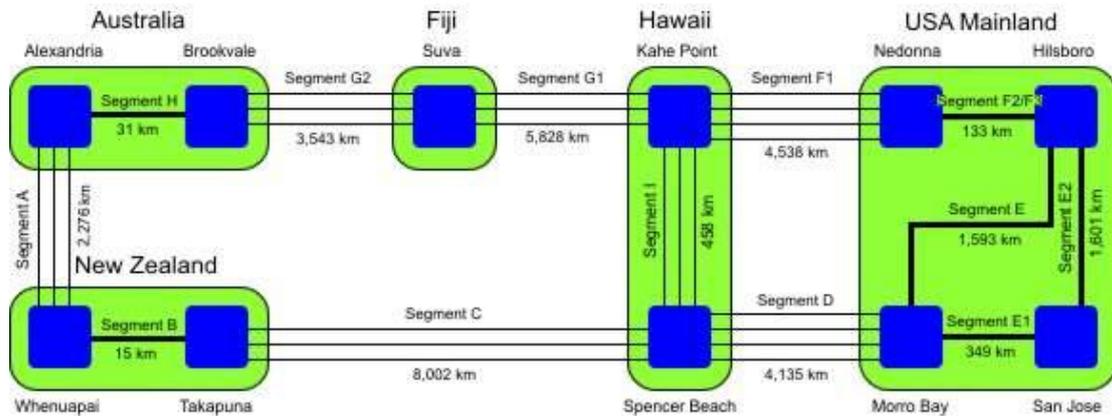


Figure 5 – SCCN Managed Fibre Connections [26]

Avoiding excessive use of international transit is a key priority for most ISPs, who employ measures such as transparent web caches to reduce their utilisation of the service. Due to its long-distance nature, the SCCN also incurs a minimum 140ms round-trip latency on communications, which is destructive to many latency-sensitive applications such as video conferencing – but unfortunately there is no solution to the speed of light problem.

The difference of cost between delivering traffic within the country and to the Internet at large is one of the main reasons why the domestic content delivery infrastructure is so important, and provides a solid reason to deploy more IXPs within New Zealand. The following Figure shows the propagation delay for each hop on SCCN passing through New Zealand [26].

Segment	From	To	Segment Type	Distance (Km)	Delay (ms)
A	Alexandria (AU)	Whenuapai (NZ)	Submarine	2,276	11.38
B	Takapuna (NZ)	Whenuapai (NZ)	Terrestrial	15	0.08
C	Spencer Beach (HW)	Takapuna (NZ)	Submarine	8,002	40.01
D	Morro Bay (US)	Spencer Beach (HW)	Submarine	4,135	20.68
E	Hillsboro (US)	Morro Bay (US)	Submarine	1,593	7.97
E1	Morro Bay (US)	San Jose (US)	Terrestrial	349	1.75
E2	Hillsboro (US)	San Jose (US)	Terrestrial	1,601	8.00
F1	Kahe Point (HW)	Nedonna (US)**	Submarine	4,538	22.71
F2 or F3##	Hillsboro (US)	Nedonna (US)**	Terrestrial	133	0.67
G1	Kahe Point (HW)	Suva (FJ)	Submarine	5,828	29.14
G2	Brookvale (AU)	Suva (FJ)	Submarine	3,543	17.72
H	Alexandria (AU)	Brookvale (AU)	Terrestrial	31	0.16
I	Kahe Point (HW)	Spencer Beach (HW)	Submarine	458	2.29

Figure 6 – SCCN Segments Length & Transmission Delays [26]

### 3.1.5 New Zealand IXPs

Historically, New Zealand Internet has always been based on the support of neutral peering of ISPs at exchanges. Initially, there were only few peering points at universities and an only major peering point and international gateway was setup in 1988 at the University of Waikato [27], which later became a separate entity and known as the New Zealand Internet Exchange (NZIX) when Telecom New Zealand took over operations in 1996.

Citylink has played an important part and are being credited the most for providing physical location for the high-capacity Internet exchanges that handled most domestic traffic through the late 90s. Their vision was based on setting up regional IXPs where ISPs, Universities and Content Providers could interconnect freely, while keeping the source traffic and destination traffic of large amounts of data as close to the IXP as possible. It was also envisioned that as the services of a domestic content provider becomes more popular and their business grows, a provider could scale it by replicating the servers at other IXPs to handle the load for that region only.

There are several open neutral Internet exchanges operating in New Zealand where participants can find others who wish to exchange IP traffic in a cost effective manner.

By peering at an Exchange, the party agree to provide routing information about their networks to other Exchange participants to optimise traffic flows across that Exchange.

In order to participate at an Exchange, one needs to be attached to that Exchange switch, have a router capable of running BGP4, and a block of IP numbers that they wish to route (at least a /24 subnet mask).

The first public regional Internet exchange in 1997 was opened by Citylink in Wellington and now it's known as the Wellington Internet Exchange (WIX). This exchange had no physical location at the time it was built; instead it was distributed

amongst all participants subscribing to Citylink's existing metropolitan fibre service. WIX is also known as the world's first IXP [28].

Since, most of the ISPs, Government Departments, Content Providers, and large businesses in Wellington already had Citylink's fibre service, joining the exchange to freely swap data amongst themselves made a lot of sense. In order to join the exchange, subscribers only need to advertised their address space using the BGP to a pair of route servers located at Citylink, and from that point on they could route data across WIX instead of it having to make the round trip to the NZIX in Hamilton.

The WIX eventually became a huge success and there are 163 subscribers as of June 2012, all can transfer data at up to 1Gbps.

After a huge success of WIX, Citylink decided to build additional IXPs in all major cities starting with Auckland because of its high volume of traffic demand. Auckland Peering Exchange (APE) was formed in 1999 and is located on the 48th floor of the Sky Tower in Auckland CBD. It did not adapt the exact same model as WIX as Citylink did not own a pervasive metropolitan fibre network in Auckland.

Sky Tower in Auckland is considered to be the hub for both national and international traffic and various international transit providers terminate their service on this collocation deck, so almost every ISP in New Zealand already buy some kind of presence to this location, even though they don't necessarily have to peer at APE.

Anyone who wishes to peer at APE can connect via local fibre, copper or wireless options from a number of providers that terminate on the Sky Tower's collocation deck. Even if APE a home to only 59 freely peering parties, more traffic flows over APE than WIX because of Auckland's greater population, and the exchange's significance as a gateway to the major ISPs as well as the rest of the world.

As of now APE is considered to be the biggest IXP by volume of traffic and WIX is known as the biggest IXP by number of participants. In New Zealand, as of to-date, the following neutral Internet Exchanges are in operation:

- Auckland Peering Exchange (APE)
- Hamilton Internet Exchange (HIX)
- Palmerston North Internet Exchange (PNIX)
- Wellington Internet Exchange (WIX)
- 3 Cities Internet Exchange (3CIX) - (Hutt City, Upper Hut, Porirua)
- NZ IPv6 Internet Exchange (V6IX) – [WIX6 and APE6]
- Christchurch Internet Exchange (CHIX)
- Dunedin Peering Exchange (DPE)

The following exchange is planned but not yet operational:

- Southland Internet Exchange

While the exchanges are not directly connected to each other, any participants that wish to are able to lease capacity on transit circuits from a number of providers which are connected to both exchanges, or lease dedicated fibre from national transit providers to connect them to a distant exchange.

### **3.1.6 Why should non-ISP Peer?**

Peering is principally a social, political and economic problem. The technical procedures needed to implement peering are well defined and relatively straightforward, when compared with the machinations required to get agreement to exchange routes with a prospective peer.

Policies towards peering in New Zealand vary widely amongst ISP's. The smaller ISP's (generally those that don't have direct Southern Cross capacity, and which are locally owned) will usually cheerfully peer with all and sundry. Larger organisations with aspirations toward 'Telco' status can be more difficult to convince to peer. In general, peering in New Zealand is no harder, and arguably easier to organise than elsewhere in the world. The high cost differential between international and national bandwidth provides a strong incentive for network providers to keep New Zealand

traffic within New Zealand, and that encourages those providers to look upon peering favourably.

### **3.1.7 The Current Peering Scenario**

Previously, almost every ISP in New Zealand was either directly or indirectly peered at either WIX or APE, or even both. This was an ideal situation for IXPs as it tends to serve the consumption of large volumes of data as cheap as the data circuit to reach the exchange.

Unfortunately in 2004, NZ Telecom and Telstraclear, the two key providers in New Zealand decided it was no longer in their best interests to peer without charge at the IXPs and instead access to their customers should be provided by means of a “domestic transit”, which could be delivered either by a dedicated connection from the Telco, or through a private peering agreement at any peering point.

Since access to these customers now only available over a metered connection, most content hosts were also forced to adopt this new service. A drawback of this was also that the amount of traffic handled by the peering exchanges dropped dramatically, and it seemed like the value of IXP has lost. So when the two key players decided to drop the peering, it also de-peered many other organisations as well as a number of organizations that were already using Telecom or Telstraclear as their upstream provider and relying on their peering to make their content available at IXPs.

### **3.1.8 Domestic Transit**

Domestic transit is a service offered by both NZ Telecom and Telstraclear, where they sell other Tier 2 ISPs to connect into their backbone network. This offer is also open to non-ISP organizations who buy the same service anywhere else in New Zealand.

As NZ Telecom and Telstraclear have a direct link between their backbones, it doesn't make any difference which ISP service is purchased from, as transit customers will still gain access to both network's customers.

This kind of domestic transit service is considered to be typical for tier-2 ISPs and traffic is only sent through the upstream link of Tier-1 providers when it is necessary to reach traffic at other IXPs that are not available at a peering exchange where they are connected. Since NZ Telecom and Telstraclear only connect their subsidiary ISPs to their respective domestic transit networks and not to IXPs as used to do previously, the domestic transit service is now also the primary means for all other Tier-2 ISPs to delivering data to their customers.

In New Zealand all the ISPs now has some path to a domestic transit service that belongs to either NZ Telecom or Telstraclear, most ISPs do not find it economic to use domestic transit as a mean of transferring data nationally, as peering to an IXP is a much cheaper but there is not much option and they have to use domestic transit mainly because some domestic content is not available on IXPs.

### **3.1.9 International Experience**

Majority of the developing countries are encouraged to setup an IXP by looking at the successful IXP Model in European exchanges and the constant growth in traffic in those exchanges are taken as a sign that IXP many ISPs are moving towards IXP.

Generally, the traffic volumes at the biggest European IXPs: DE-CIX, AMS-IX, and LINX are characterized by a constant growth.

For example, average traffic throughput at the DE-CIX reached approx. 500 GB/s at the end of 2010, approx. 800 GB/s at the end of 2011 and roughly 1.250 GB/s at the end of April 2012. See below figure [30]

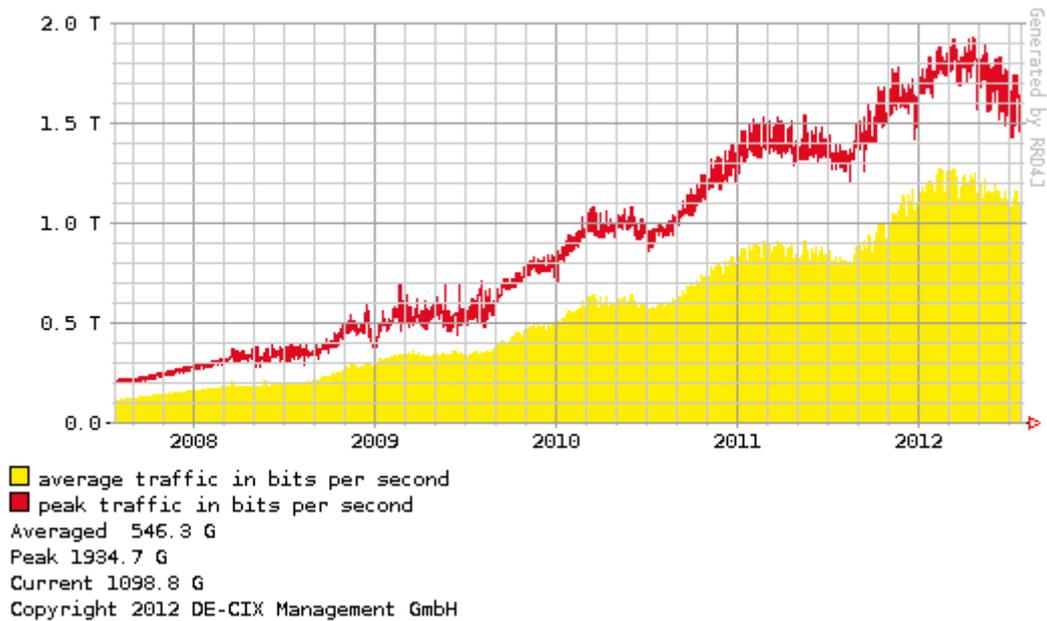


Figure 7 – Average Traffic Throughput at the DE-CIX [30]

### 3.1.10 Cost of Joining a Peering Point

The overall cost position of network operators is affected in particular by two aspects: the increase in overall traffic volumes and technological improvements. Generally, increasing traffic volumes – both in fixed and mobile networks – imply higher absolute costs for network operators if they have to invest in additional equipment. However, per unit costs may even fall depending on the economies of scale. Technological progress is the second factor impacting both, on overall and particular per-unit costs of a network operator. If technological progress leads to cost improvements (on a per unit basis) which outweigh the increase in traffic volumes then there would be no negative effect on the overall cost position of a network operator. Of course, in practice, a comprehensive assessment would require to also take into account the revenues generated by this operator.

Moore's law provides an interesting illustration that technological progress leads to significant performance improvements. In 1965 Gordon Moore stated that the number of integrated circuits on a computer chip was doubling every 18-24 month.

In order to assess the relevance of Moore's law on cost positions of network operators it is necessary to separately look at unit costs in the core network and in mobile networks.

In the core network costs for routers and optics showed significant declines over the years. The following figure from Cisco illustrates that the costs per Gbps for their routers decreased at an annual rate of 23% (1997-2012). See figure below [7]:

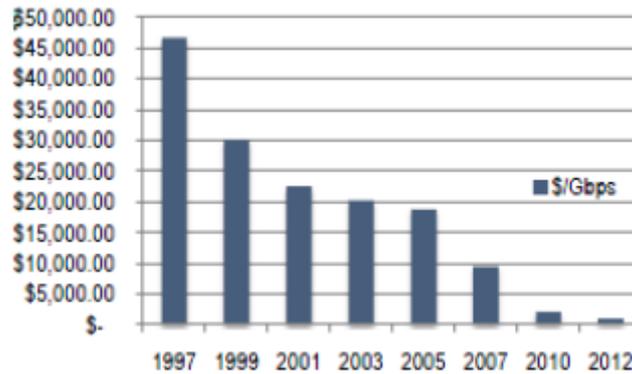


Figure 8 – Cisco Hardware Cost Reduction Comparison [7]

The above figure indicates that hardware cost should not be a significant barrier for an organisation joining a peering point.

## 3.2 Quantitative Evaluation

This section provides a quantitative evaluation on how ISPs or Content Providers can generate more revenues by peering with IXP

### 3.2.1 APBDC Model

#### APBDC Model

The following calculations are based on the of APBDC (Average Per-Bit Delivery Costs) Model developed by Packet Clearing House [13]. It shows the cost incurred on domestic peering links versus International transit links.

This model describes mathematically how an ISP can generate more revenues by peering to IXP on top of their existing E1 line to their carrier. In our context, it shows that an ISP will be more cost-effective by peering with an IXP.

#### Calculating APBDC

If an ISP has an E1 line of transit, which runs at an average of 40% utilization, that's 2.0Mbps of total potential capacity, but approximately 800Kbps of actual utilization.

800kbps (0.8Mbps)  
= 48Mbits/minute  
= 2.88Gbits/hour  
= 69.12Gbits/day  
= 2108.16Gbits/month (69.12GB/Day x 30.5)

If the ISP is spending NZ\$12,000 per month for that circuit, a simple interpretation of that price is that it's NZ\$6,000 megabit/second/month. However, that doesn't take utilization into account at all.

APBDC for 2,108.16 gigabits at NZ\$12,000 is:

$$\frac{\$12,000}{2,108.16Gb} = \$5.69 / Gb$$

Or NZ\$5.69 per Gigabit

By contrast, the ISP might also have a 100Mbps Ethernet link to an IXP. If they spend NZ\$3,000 on equipment, and they amortize the equipment over two years, they have a monthly cost of NZ\$125/month associated with the peering. If they run that connection at 15% utilization, that's 15Mbps.

15Mbps  
= 900Mbits/minute  
= 54Gbits/hour  
= 1,296Gbits/day  
= 39,528Gbits/month (1,296GB/Day x 30.5)

APBDC for 39,528 Gigabits at NZ\$125 is:

$$\frac{\$125.00}{39,528.00Gb} = \$0.00316/Gb$$

Or NZ\$0.003 per Gigabit

### **Aggregating APBDC**

If the ISP mentioned above can have only those two circuits as means of delivering packets, it's delivering a total of 15.8Mbps, and spending NZ\$12,125/month:

$\$12,000 + \$125 = \$12,125$   
 $800Kbps (0.8Mbps) + 15.0Mbps = 15.8Mbps$   
 $15.8Mbps * 2,635,200 \text{ seconds/month} = 41,636.16Gbits/month$

$$\frac{\$12,125.00}{41,636.16Gb} = \$0.291/Gb$$

So the overall APBDC for the ISP is \$0.291/Gigabit. There are two contributing cost-factors: transit at \$5.69/gigabit, and peering at \$0.003/gigabit. Obviously if you are

buying two products, and reselling them at the same price, you want to be buying as much of the cheap one as possible, and as little of the expensive one as possible, particularly when there's a three-order-of-magnitude difference between the two.

### **Shifting Traffic to Improve APBDC**

The inflexibility-with-respect-to-demand of circuit costs is a major problem for ISPs wishing to bring their contributing costs into line with their needs.

If the example ISP were able to shift 200Kbps of traffic (8.7%) from transit to peering, that would reduce their transit from 800Kbps to 600Kbps, and increase peering from 15Mbps to 17Mbps.

That would reduce average utilization of the 2Mbps transit circuit from 40% utilization to 30% utilization. If that allows the ISP to reduce the capacity of the circuit from 2Mbps to 1Mbps, which might reduce the price from \$12,000 to \$8,000, and yielding a new utilization of 60%, well within the commercially-reasonable envelope.

$$600\text{Kbps} * 2,635,200 \text{ seconds/month} = 1,581.12\text{Gbits/month}$$

$$\frac{\$8,000.00}{1,581.12\text{Gb}} = \$5.06 / \text{Gb}$$

$$17\text{Mbps} * 2,635,200 \text{ seconds/month} = 44,798.4\text{Gbits/month}$$

$$\frac{\$125.00}{4,4798.4\text{Gb}} = \$0.00279 / \text{Gb}$$

$$15.8\text{Mbps} * 2,635,200 \text{ seconds/month} = 41,636.16\text{Gbits/month}$$

$$\frac{\$8,125.00}{41,636.16\text{Gb}} = \$0.195 / \text{Gb}$$

So here, by shifting 200kbps of traffic from transit to peering, the ISP has reduced APBDC from \$0.291/GB to \$0.195/GB, and real-world costs from \$12,125/month to \$8,125/month.

### 3.2.2 Bandwidth Speed & Delay Comparisons

In this section we will evaluate the bandwidth speed and delay comparisons over local and international links.

#### Scenario 1

In this scenario, for New Zealand server, we have downloaded the file from a location peering with local IXP. For this test we are using Victoria University of Wellington (VUW) as it peers with WIX.

As a test we will download an ISO Image for SUSE Linux 12.1 from two servers and observed the speed and time-delay comparison. The follow are the two links where ISO Image is available to download.

#### UK Mirror Server

[http://mirror.rackspace.co.uk/openSUSE/distribution/12.1/iso/openSUSE-12.1-DVD-x86\\_64.iso](http://mirror.rackspace.co.uk/openSUSE/distribution/12.1/iso/openSUSE-12.1-DVD-x86_64.iso)

#### NZ Mirror Server

[http://mirror.xnet.co.nz/pub/opensuse/distribution/12.1/iso/openSUSE-12.1-DVD-x86\\_64.iso](http://mirror.xnet.co.nz/pub/opensuse/distribution/12.1/iso/openSUSE-12.1-DVD-x86_64.iso)

Before we begin the download, we will do a simple ping test to both mirrors and observe the Round-trip time (RTT). We can see from the results below that RTT for an ICMP packet is 19ms for server located in NZ and RTT for server located in UK has 303ms, which shows us that if the similar packet is retrieved from both servers, the UK mirror server will take additional 284ms than NZ mirror server.

```
C:\Windows\System32>ping mirror.xnet.co.nz
Pinging mirror.xnet.co.nz [58.28.25.150] with 32 bytes of data:
Reply from 58.28.25.150: bytes=32 time=26ms TTL=56
Reply from 58.28.25.150: bytes=32 time=16ms TTL=56
Reply from 58.28.25.150: bytes=32 time=18ms TTL=56
Reply from 58.28.25.150: bytes=32 time=16ms TTL=56
Ping statistics for 58.28.25.150:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 16ms, Maximum = 26ms, Average = 19ms
C:\Windows\System32>
```

```

C:\Windows\System32>ping mirror.rackspace.co.uk
Pinging mirror.rackspace.co.uk [92.52.126.40] with 32 bytes of data:
Reply from 92.52.126.40: bytes=32 time=306ms TTL=37
Reply from 92.52.126.40: bytes=32 time=303ms TTL=37
Reply from 92.52.126.40: bytes=32 time=292ms TTL=37
Reply from 92.52.126.40: bytes=32 time=312ms TTL=37
Ping statistics for 92.52.126.40:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 292ms, Maximum = 312ms, Average = 303ms
C:\Windows\System32>

```

Listing 1 – RTT Delay to NZ and UK Servers from VUW

We have done another test using TRACERT (Traceroute) which sends a sequence of ICMP echo request packets addressed to a destination host. Determining the intermediate routers traversed involves adjusting the time-to-live (TTL), aka hop limit. The following results will show you that it takes only 9 hops to reach the mirror server in New Zealand, while it takes 24 hops to reach the mirror server in UK.

```

C:\Windows\System32>tracert mirror.xnet.co.nz
Tracing route to mirror.xnet.co.nz [58.28.25.150]
over a maximum of 30 hops:
  0  168 ms  230 ms  293 ms  10.140.0.1
  1  12 ms   10 ms   6 ms   130.195.199.190
  2  7 ms    13 ms   8 ms   130.195.199.115
  3  10 ms   14 ms   5 ms   130.195.197.67
  4  12 ms   11 ms   12 ms  130.195.197.133
  5  5 ms    17 ms   5 ms   ntp.vuw.ac.nz [130.195.196.196]
  6  30 ms   16 ms   14 ms  wxnz1.wix.net.nz [202.7.1.22]
  7  23 ms   27 ms   31 ms  ip-58-28-25-149.wxnz.net [58.28.25.149]
  8  59 ms   17 ms   20 ms  mirror.xnet.co.nz [58.28.25.150]
Trace complete.
C:\Windows\System32>

C:\Windows\System32>tracert mirror.rackspace.co.uk
Tracing route to mirror.rackspace.co.uk [92.52.126.40]
over a maximum of 30 hops:
  0  333 ms  6 ms    5 ms   10.140.0.1
  1  9 ms    *       4 ms   130.195.199.190
  2  6 ms    7 ms    7 ms   130.195.199.115
  3  10 ms   8 ms    43 ms  130.195.197.67
  4  4 ms    4 ms    7 ms   130.195.197.133
  5  9 ms    5 ms    9 ms   ntp.vuw.ac.nz [130.195.196.196]
  6  8 ms    12 ms   8 ms   41.36.69.111.dynamic.snap.net.nz [111.69.36.41]
  7  *       *       *      Request timed out.
  8  *       *       *      Request timed out.
  9  31 ms   27 ms   25 ms  111-69-27-66.core.snap.net.nz [111.69.27.66]
 10 28 ms   28 ms   42 ms  111-69-27-65.core.snap.net.nz [111.69.27.65]
 11 161 ms  168 ms  156 ms te7-3.ccr01.sjc05.atlas.cogentco.com [38.122.92.105]
 12 169 ms  172 ms  170 ms te0-0-0-4.mpd22.sfo01.atlas.cogentco.com [66.28.4.149]
 13 192 ms  192 ms  *     te0-2-0-2.mpd22.mci01.atlas.cogentco.com [154.54.7.221]
 14 222 ms  202 ms  206 ms te0-0-0-3.mpd22.ord01.atlas.cogentco.com [154.54.25.78]
 15 311 ms  368 ms  310 ms te0-0-0-2.ccr21.bos01.atlas.cogentco.com [154.54.43.186]
 16 305 ms  360 ms  302 ms te0-1-0-2.mpd21.lon13.atlas.cogentco.com [154.54.30.130]
 17 305 ms  389 ms  305 ms te4-2.ccr02.lhr01.atlas.cogentco.com [130.117.50.218]
 18 334 ms  300 ms  322 ms 149.6.8.122
 19 301 ms  300 ms  300 ms coreb-edge5.lon3.rackspace.net [164.177.137.40]
 20 297 ms  315 ms  299 ms core3-coreb.lon3.rackspace.net [164.177.137.19]
 21 301 ms  300 ms  299 ms aggrit9a-core3.lon3.rackspace.net [92.52.76.69]
 22 298 ms  298 ms  303 ms seg-support-aggr-a.lon3.rackspace.net [92.52.76.165]
 23 294 ms  295 ms  297 ms 92.52.126.40
Trace complete.
C:\Windows\System32>

```

Listing 2 – Hops to NZ and UK Servers from VUW

As a test, we have downloaded the image from both UK and NZ mirror servers. The ISO Image size is 4.31 GB. We will observe the time and speed delay between the two downloads.

## From NZ Server

The following figure shows the download over a 1GB Ethernet connection from VUW:

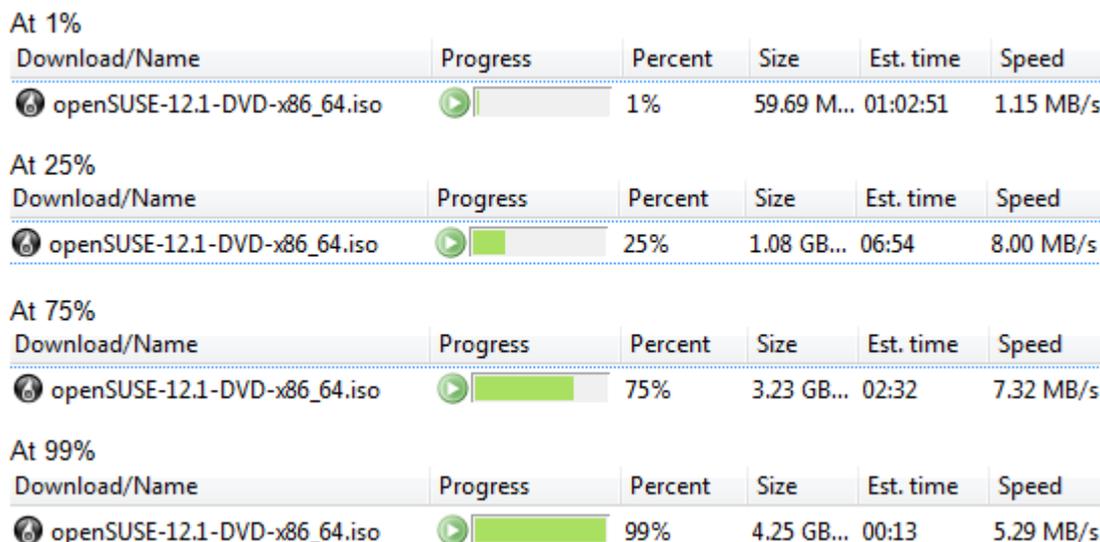


Figure 9 – Downloading file hosted at NZ Server from VUW

The following figure displays the Average Download speed was 9.25Mbps over the total duration of 09:32 minutes

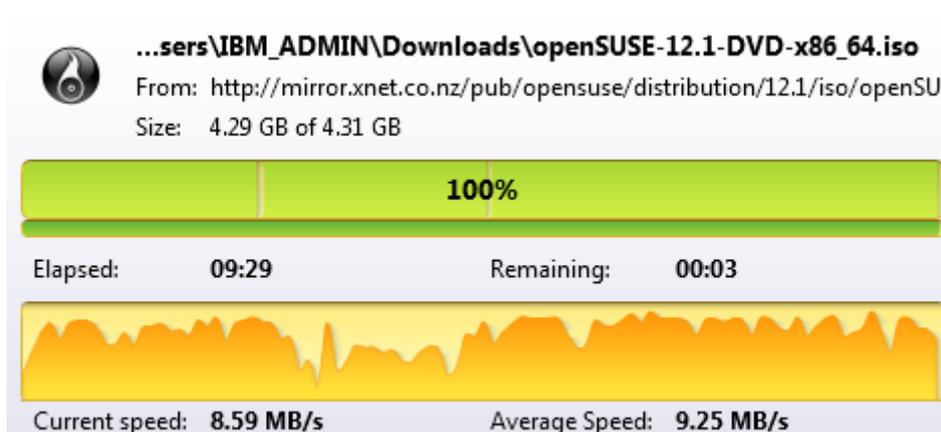


Figure 10 – Time Graph of file hosted at NZ Server from VUW

### From UK Server

The following figure shows the download over a 1GB Ethernet connection from VUW:

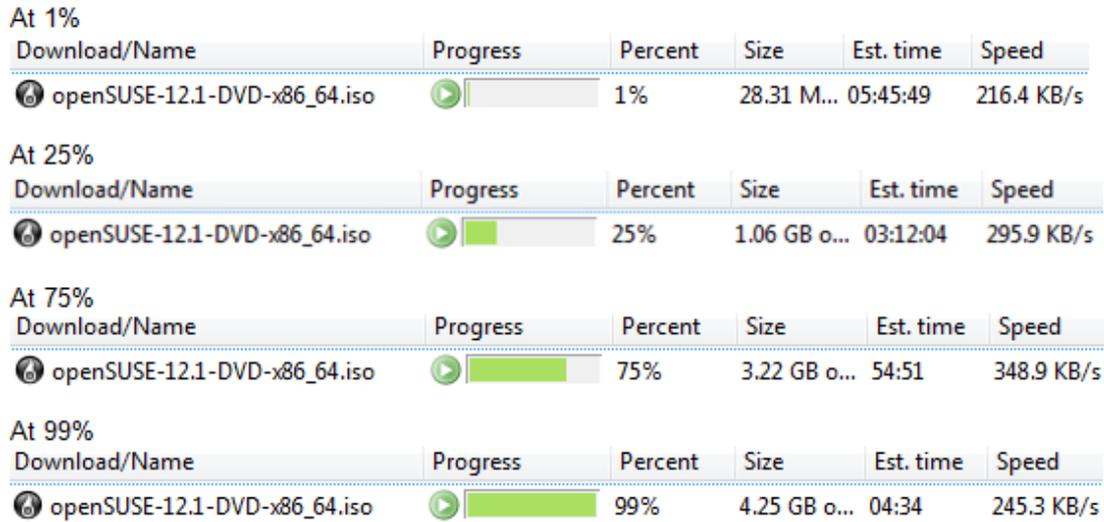


Figure 11 – Downloading file hosted at UK Server from VUW

The following figure displays the Average Download speed was 192.5Kbps over the total duration of 04:06:28 hours.

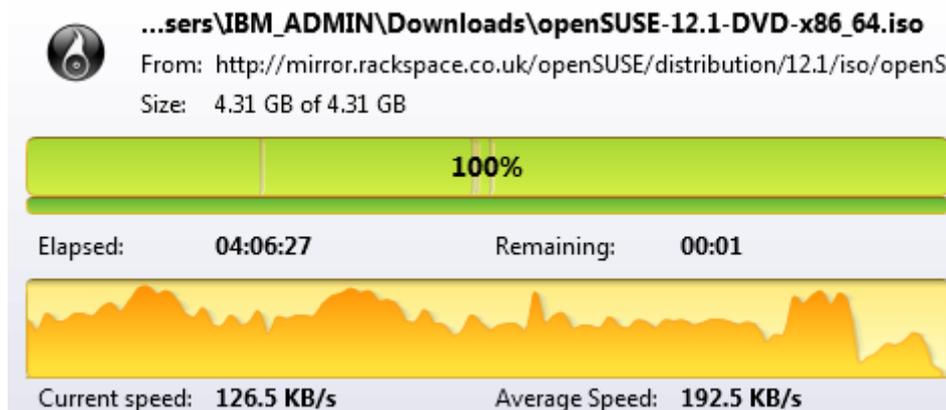


Figure 12 – Time Graph of file hosted at UK Server from VUW

To verify our understanding, we have now done a second test and this time we have decided to download an ISO file from a different domestic and International location. As a test we will download an ISO Image for Ubuntu Linux 12.04 from two servers and observed the speed and time-delay comparison. The follow are the two links where ISO Image is available to download.

## UK Mirror Server

<http://releases.ubuntu.mirrors.uk2.net//precise/ubuntu-12.04-server-amd64.iso>

## NZ Mirror Server

<http://ftp.citylink.co.nz/ubuntu-releases/precise/ubuntu-12.04-server-amd64.iso>

Before we begin the download, we will do a simple ping test to both mirrors and observe the Round-trip time (RTT). We can see from the results below that RTT for an ICMP packet is 157ms for server located in NZ and RTT for server located in UK has 293ms, which shows us that if the similar packet is retrieved from both servers, the UK mirror server will take additional 284ms than NZ mirror server.

```
C:\Windows\System32>ping ftp.citylink.co.nz
Pinging ftp.citylink.co.nz [202.7.6.9] with 32 bytes of data:
Reply from 202.7.6.9: bytes=32 time=166ms TTL=46
Reply from 202.7.6.9: bytes=32 time=155ms TTL=46
Reply from 202.7.6.9: bytes=32 time=155ms TTL=46
Reply from 202.7.6.9: bytes=32 time=155ms TTL=46
Ping statistics for 202.7.6.9:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 155ms, Maximum = 166ms, Average = 157ms

C:\Windows\System32>
C:\Windows\System32>ping releases.ubuntu.mirrors.uk2.net
Pinging releases.ubuntu.mirrors.uk2.net [83.170.94.3] with 32 bytes of data:
Reply from 83.170.94.3: bytes=32 time=304ms TTL=44
Reply from 83.170.94.3: bytes=32 time=292ms TTL=44
Reply from 83.170.94.3: bytes=32 time=297ms TTL=44
Reply from 83.170.94.3: bytes=32 time=282ms TTL=44
Ping statistics for 83.170.94.3:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 282ms, Maximum = 304ms, Average = 293ms
C:\Windows\System32>
```

Listing 3 – RTT Delay to NZ and UK Servers from non-peering location

## Scenario 2

In this scenario, for New Zealand server, we have downloaded the file from a location that does not peer with local IXP.

We have done another test using TRACERT (Traceroute) which sends a sequence of ICMP echo request packets addressed to a destination host. Determining the intermediate routers traversed involves adjusting the time-to-live (TTL), aka hop limit. The following results will show you that it takes 17 hops to reach the mirror server in New Zealand, while it takes 19 hops to reach the mirror server in UK.

```

C:\Windows\System32>tracert ftp.citylink.co.nz
Tracing route to ftp.citylink.co.nz [202.7.6.9]
over a maximum of 30 hops:
  1  17 ms  8 ms  1 ms
  2  2 ms  3 ms  3 ms
  3  2 ms  2 ms  2 ms
  4  3 ms  2 ms  2 ms
  5  2 ms  2 ms  2 ms
  6  3 ms  3 ms  3 ms
  7  12 ms 13 ms 12 ms
  8  24 ms 13 ms 13 ms
  9  13 ms 12 ms 13 ms
 10 138 ms 138 ms 138 ms i-0-6-4-0.tlot-core01.bx.telstraglobal.net [202.84.142.110]
 11 142 ms 142 ms 142 ms i-4-2.eq1a01.bi.telstraglobal.net [202.40.149.134]
 12 144 ms 144 ms 144 ms gblx-peer.eq1a01.pr.telstraglobal.net [134.159.61.22]
 13 146 ms 149 ms 149 ms e5-1-40G.ar6.LAX1.gblx.net [67.17.111.65]
 14 302 ms 208 ms 208 ms united-layer.ethernet3-15.ar6.lax1.gblx.net [67.17.157.86]
 15 156 ms 155 ms 155 ms v1an802.br2.sf7.unitedlayer.com [209.237.224.137]
 16 155 ms 155 ms 155 ms v170.br01-eqx-ash.unitedlayer.com [207.7.146.2]
 17 156 ms 155 ms 155 ms ftp.citylink.co.nz [202.7.6.9]
Trace complete.
C:\Windows\System32>

C:\Windows\System32>tracert releases.ubuntu.mirrors.uk2.net
Tracing route to releases.ubuntu.mirrors.uk2.net [83.170.94.3]
over a maximum of 30 hops:
  1  13 ms  3 ms  3 ms
  2  2 ms  3 ms  3 ms
  3  2 ms  2 ms  2 ms
  4  2 ms  2 ms  2 ms
  5  2 ms  2 ms  2 ms
  6  3 ms  3 ms  3 ms
  7  12 ms 12 ms 12 ms
  8  17 ms 12 ms 12 ms
  9 138 ms 137 ms 137 ms te7-3.ccr01.sjc05.atlas.cogentco.com [38.122.92.105]
 10 139 ms 139 ms 139 ms te0-1-0-5.mpd22.sfo01.atlas.cogentco.com [66.28.4.181]
 11 177 ms 177 ms 176 ms te0-0-0-2.mpd22.mci01.atlas.cogentco.com [154.54.6.33]
 12 188 ms 188 ms 189 ms te0-5-0-4.mpd22.ord01.atlas.cogentco.com [154.54.45.158]
 13 208 ms 208 ms 208 ms te0-7-0-23.mpd22.jfk02.atlas.cogentco.com [154.54.43.77]
 14 292 ms 298 ms 284 ms te0-2-0-5.mpd22.lon13.atlas.cogentco.com [154.54.84.142]
 15 290 ms 297 ms 293 ms te0-4-0-1.ccr22.lon01.atlas.cogentco.com [130.117.0.245]
 16 300 ms 301 ms 294 ms 149.6.3.158
 17 300 ms 301 ms 291 ms 83.170.70.234
 18 301 ms 292 ms 294 ms dc5.as13213.net [83.170.70.138]
 19 296 ms 291 ms 299 ms mirrors.uk2.net [83.170.94.3]
Trace complete.
C:\Windows\System32>

```

Listing 4 – Hops to NZ and UK Servers from non-peering location

As a test, we have downloaded the image from both UK and NZ mirror servers. The ISO Image size is 684.29 MB. We will observe the time and speed delay between the two downloads.

**From NZ Server**

The following figure shows the download over a 1GB Ethernet connection from a different location and not VUW:

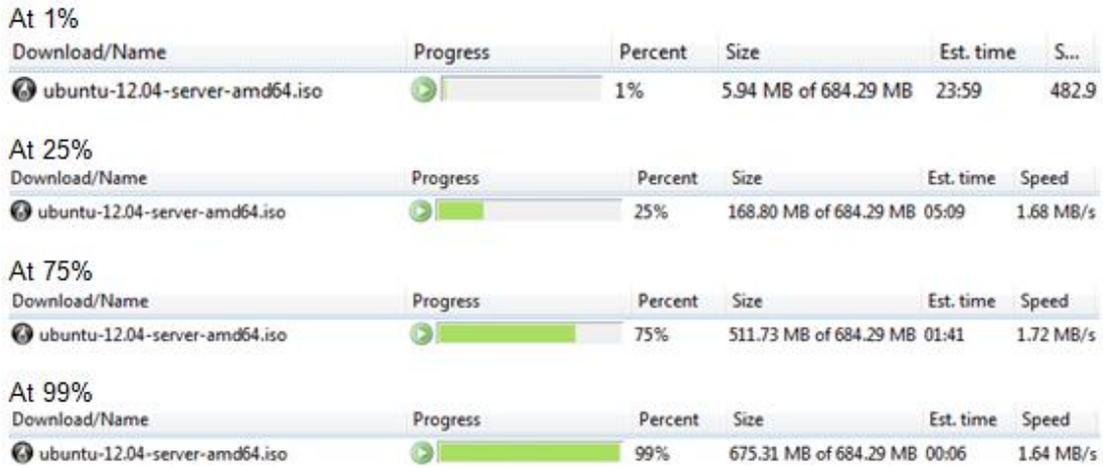


Figure 13 – Downloading file hosted at NZ Server from non-peering location

The following figure displays the Average Download speed was 1.64Mbps over the total duration of 00:06:57 minutes

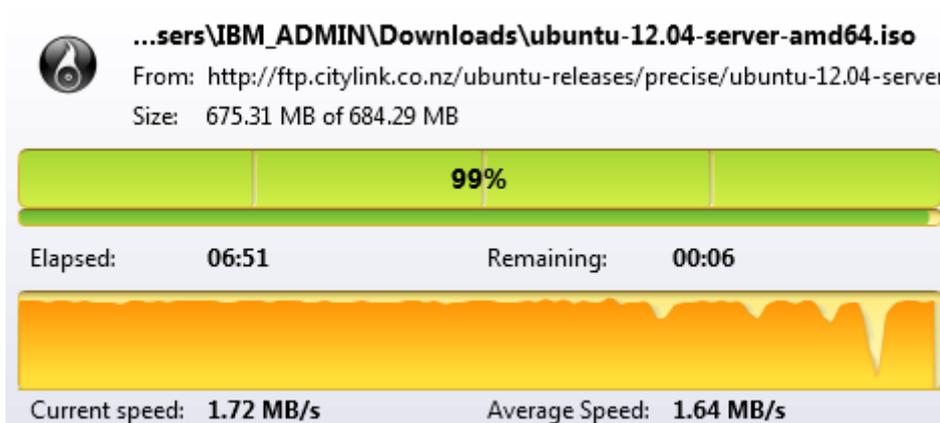


Figure 14 – Time Graph of file hosted at NZ Server from non-peering location

### From UK Server

The following figure shows the download over a 1GB Ethernet connection from a different location and not VUW:

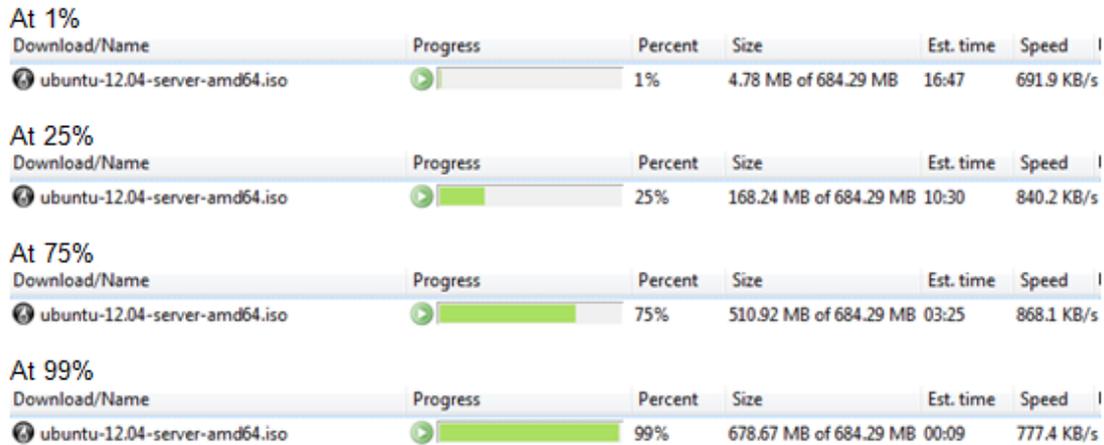


Figure 15 – Downloading file hosted at UK Server from non-peering location

The following figure displays the Average Download speed was 777.4Kbps over the total duration of 00:15:14 minutes.

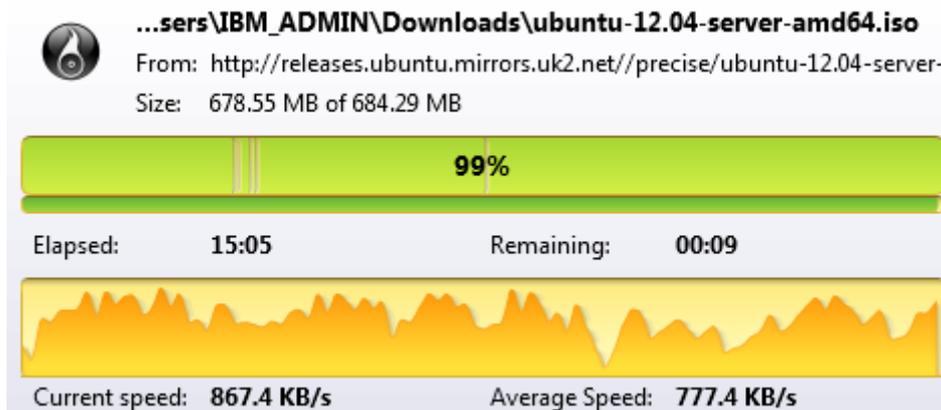


Figure 16 – Time Graph of file hosted at UK Server from non-peering location

### Scenario 1 vs. Scenario 2

Here we will discuss the difference between downloading the file from a local content provider, hosting its services at local IXP. We have downloaded the file from two different locations. One of them peers with IXP and the other one does not.

In Scenario 1, we noticed that we received a much higher bandwidth for downloading a large file size of 4.29GB from a local IXP and the average downloading speed went as high as 9.25MB/s as we were downloading the file from VUW, which directly peers with local IXP.

In Scenario 2, we downloaded a small file size of 684.29MB from a local IXP, but the average downloading speed only went up to 1.64MB/s as we downloaded the file from a location that does not peer with local IXP.

### 3.2.3 Bandwidth & Time Delay Graph Model

This section explains the difference between scenario 1 and 2 using linear graphs.

#### Linear graph for Scenario 1:

File Size: 4.31 GB

Average Download speed = 192.5Kbps. Time = 04:06:28 hours for server in UK

Average Download speed = 9.25Mbps. Time = 00:09:32 minutes for server in NZ

The following figure displays that Bandwidth is high while time utilized is very low for NZ sever (the higher-end of graph), while bandwidth is very low and time utilized is very high (the lower-end of graph).

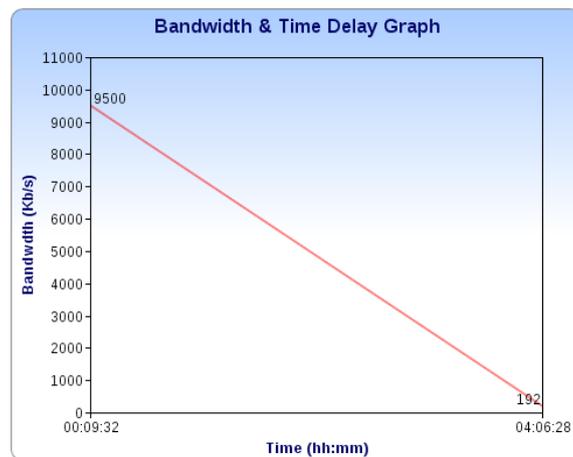


Figure 17 – Bandwidth Time and Delay Graph from peering location

#### Linear graph for Scenario 2:

File Size: 684.29 MB

Average Download speed = 192.5Kbps. Time = 00:15:14 minutes for server in UK

Average Download speed = 9.25Mbps. Time = 00:06:57 minutes for server in NZ

The following figure displays that Bandwidth is high while time utilized is very low for NZ sever (the higher-end of graph), while bandwidth is very low and time utilized is very high (the lower-end of graph).

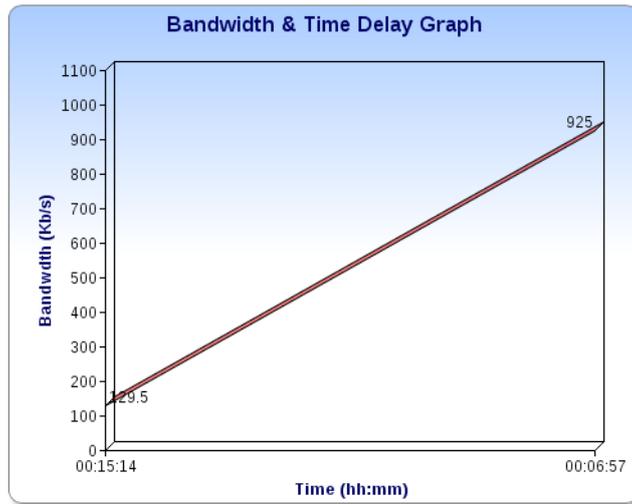


Figure 18 – Bandwidth Time and Delay Graph from non-peering location



## 4. Analysis and Design

In this chapter we will discuss the analysis and design of Open Source Internet Exchange (OSIX).

### 4.1 Requirements Gathering

In this chapter we will introduce the rationale for the design of OSIX, our peering point in-a-box and we will present a set of tactics and different approaches to setup the IXP. Finally we will present the design document for our IXP and our implementation will be based on our design.

We have conducted an online ‘IXP Fact Finding Survey’ with IXP stake-holders across the globe. This was conducted in accordance with the VUW Human Ethics Policy. See Appendix E. We shared our survey with more than 50 participants around the world, but unfortunately the response was not very high. We selected all participants randomly from the list of IXP provided on PCH and Wikipedia website. We contacted the IXP administrators via email and requested them to respond to the Survey. See Appendix A, Appendix B and Appendix C for the IXP Survey and the results.

The survey was published and participants were notified on 31<sup>st</sup> May 2012, we kept it open until 30<sup>th</sup> July 2012 but only four participants chose to respond (one participant only completed half of the survey).

We have also used the results from another survey by Packet Clearing House (PCH) as it is relevant to our research and it was conducted recently in between October 2010 and March 2011. The results were published on May 2, 2011 at PCH website.

[29]

Based on our survey results, we have identified expectedly that majority of the IXPs in developing countries like to have an easy-to-run and simple-to-manage IXP which can handle large volume of traffic.

Most of the participants have also chosen to adapt Layer 2 IXP (we will discuss in later sections) and like to see the IXP to support both Multilateral and Bilateral peering.

The following is the breakdown in percentile of the key responses in our survey. See Appendix C for complete results and statistics:

- 75% prefer Layer 2 IXP
- 75% prefer to have multilateral peering.
- 100% have preferred to have a dedicated management server at their IXP
- 50% prefer to have both CLI and Web access to the management server.

## **4.2 Design Considerations**

In this section we will discuss some of the considerations in IXP Design process in Layer 2 IXP and what sort of peering should be used.

Internet Exchange Point (IXPs) is a vital part of Internet, without IXP the Internet could not function at its best because the different networks that make up the Internet would not be able to exchange local traffic with each other and hence they have to rely on International bandwidth. The simplest form of an exchange point is a direct connection between two Internet Service Providers (ISPs). [21]

When more than two providers operate in the same area, an independent switch operates more efficiently as a common interconnection point at which to exchange traffic between the local networks. [21]

The design of IXP is the most critical part of the entire process as the right approach cannot only save from a lot of grief in further steps but also ensures that any future upgrades or changes can be implemented tracking the design. In this scenario of designing the IXP solution at a minimum apply the following principles:

- All peering shall take place using BGP 4

- It supports both bilateral and multilateral peering.
- Two route servers must be running in the IXP to provide resiliency and failover in case if one of the server loses connectivity or have been taken down for maintenance.
- The design allows yearly, monthly, weekly and daily graphical statistics in graphical web-interface.

Peering means that the ISPs have to buy circuits between each other. Each participant has to buy one whole circuit from their premises to the IXP neighbours. In IXP, rather than N-1 half circuits to connect to the N-1 other ISPs. Since each ISP participates in IXP, the cost is minimal – one local circuit covers all domestic traffic and International circuits are used just for international traffic and backing up domestic links in case the IXP fails.

In layer 2 IXP, only one switch is required but two switches can also be used for redundancy. From design perspective, it also requires neutral ISP management usually funded equally by IXP participants. It provides 24x7 cover, support and value-added services. IXP should also be located on secure neutral location.

#### 4.2.1 Layer 2 IXP

All traffic is exchanged outside routers that are connected to a shared media (i.e. Ethernet 10/100/1000BaseTX)

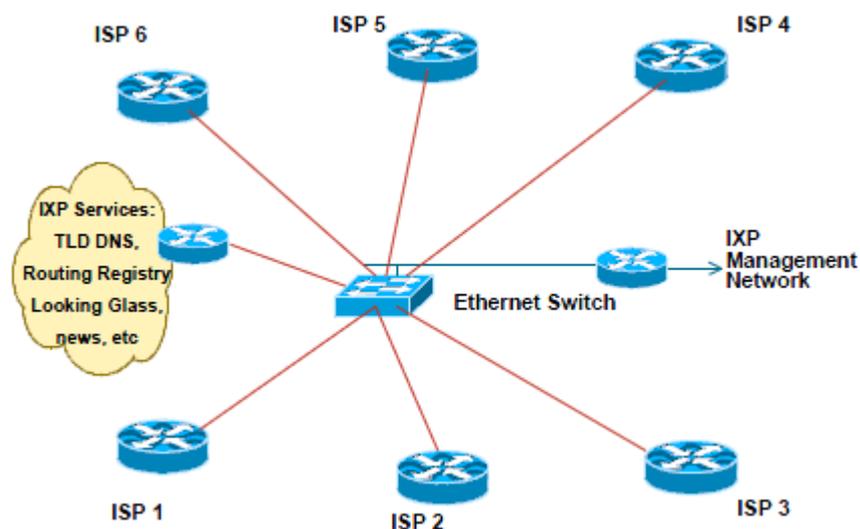


Figure 19 – Layer 2 IXP Model

### 4.2.2 Layer 3 IXP

All traffic is exchanged inside a router and Layer 3 IXPs limit the autonomy of the members. Someone has to manage the router in the middle. Create business issues, as ISPs don't have control with whom they can peer with.

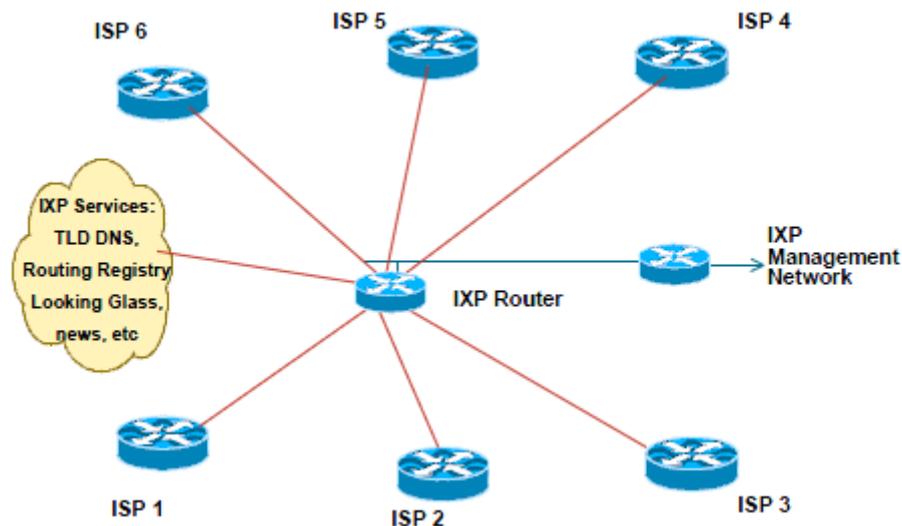


Figure 20 – Layer 3 IXP Model

### 4.2.3 Layer 2 versus Layer 3 IXP

The following section shows the key differentiates between Layer 2 and Layer 3 IXPs.

#### Layer 3 IXP

- IXP team requires good BGP knowledge
- Rely on 3rd party for BGP configuration
- Less freedom on who peers with whom
- Usually competes with IXP membership
- Tends to be distributed over wide area

#### Layer 2 IXP

- IXP team does not need routing knowledge
- Easy to get started
- More complicated to distribute over wide area

- ISPs free to set up peering agreements with each other as they wish

Layer 2 is considered as 'real' IXP whereas Layer 3 is mostly known as a concept used by Transit ISPs.

#### **4.2.4 Physical Design Considerations**

The IXP Core is an Ethernet switch and it can be from the cheapest and smallest 12 or 24 port 10/100 switch to the largest 32 port 10GigEthernet switch. Each ISP participating in the IXP brings a router to the IXP location. One Ethernet port on the router connects to IXP switch and one WAN port to connect to the WAN media leading back to the ISP backbone. Router must be able to run BGP

IXP switch is located in one equipment rack dedicated to IXP and it also includes other IXP operational equipment such as Route servers and IXP Management server. Routers from participant ISPs are located in neighbouring/adjacent rack(s) and physical connections are usually made using Copper (UTP) connections made for 10Mbps, 100Mbps or 1Gbps connections. Fibre can also be used for 10Gbps and 40Gbps.

#### **4.2.5 Peering Design Considerations**

Each participant in IXP needs to run BGP and they need their own Public AS number and not private ASN. Each participant configures external BGP directly with the other participants in the IXP. This setup is known as multilateral peering. In multilateral peering, each participant peers with every other participant. However, mandatory multilateral peering is discouraged as history suggests that they have not been very successful.

In bilateral peering, participants set up peering with each other according to their own requirements and business relationships; this is the most common situation at IXPs today.

#### **4.2.6 Routing Design considerations**

ISP border routers at the IXP generally should not be configured with a default route to carry the full Internet routing table, carrying default or full table means that this router and the ISP network is open to abuse by non-peering IXP members. Therefore the correct configuration is only to carry routes offered to IXP peers on the IXP peering router.

Preferably participant routers at the IXP should also not be configured to carry the IXP LAN network within the Interior Gateway Protocol (IGP) or Internal Border Gateway Protocol (iBGP) and should use next-hop-self BGP concept.

#### **4.2.7 IP Address Space Design considerations**

The following are some of the recommended considerations that should be followed when selecting the IP address scheme:

- Critical IP address range should be used in IXP environment
- Public address space means IXP network could be leaked to Internet which may be undesirable
- Because most ISPs filter RFC1918 address space, this avoids the problem
- IXP terms of participation often forbid the Internet Exchange LAN to be carried in the ISP member backbone

Our design supports both IPv4 and IPv6 deployment. Globally, the sooner ISPs transition over to IPv6, the quicker they will be able to achieve greater efficiencies and develop more dynamic applications. These efficiencies and applications may have the potential of reduce costs and additional revenue generation.

One issue that might motivate Internet companies to transition more quickly to IPv6 is having easy IPv6 peering support and information from IXs. This will create a critical mass of IPv6-enabled Internet traffic and applications. In sum, these changes enabled at the exchange level, can shape the growth of IPv6 in the world. [11]

## **IXP Address Assignment**

Critical Infrastructure addresses can be provided by the Regional Internet Registries (RIRs). There are total of five RIRs:

- AfriNIC - African Network Information Centre
- ARIN - American Registry for Internet Numbers
- APNIC – Asia-Pacific Network Information Centre
- LACNIC - Latin America and Caribbean Network Information Centre
- RIPE NCC for Europe, the Middle East, and Central Asia

APNIC (in case of Asia Pacific) will assign a minimum /24 for the IXP transit LAN. Multiples of /24 can be requested if there are multiple IP sites and if it is needed.

### **4.2.8 Critical Infrastructure**

Assignments to critical infrastructure are available only to the actual network operators of the network infrastructure performing such functions. Registrar organizations that do not actually host the network housing the registry infrastructure will not be eligible for an assignment.

#### Domain registry infrastructure

- Root domain name system (DNS) server
- Global top level domain (gTLD) nameservers
- Country code TLD (ccTLD) nameservers

#### Address registry infrastructure

- Internet Assigned Numbers Authority (IANA)
- Regional Internet Registries (RIRs)
- National Internet Registries (NIRs)

### **Services that can be offered at IXP and its considerations**

Services offered should not compete with member ISPs (basically at IXP) e.g. web hosting at an IXP is a bad idea unless all members agree to it [10].

IXP operations should make performance and throughput statistics available to members. For Example, tools such as MRTG can be used to produce throughput graphs for member (or public) information.

- ccTLD DNS
  - The country IXP could host the country's top level DNS e.g. "SE." TLD is hosted at Netnod IXP in Sweden
  - Offer back up of other country ccTLD DNS
- Root server
  - Anycast instances of I.root-servers.net, F.root-servers.net etc are present at many IXPs
- Route Collector
  - Route collector shows the reachability information available at the exchange
- Looking Glass
  - One way of making the Route Collector routes available for global view (e.g. [www.traceroute.org](http://www.traceroute.org))
  - Graphs can be for both public or members only access
- Content Redistribution/Caching
  - For example, Akamai update distribution service
- Network Time Protocol
  - Locate a stratum 1 time source (GPS receiver, atomic clock, etc) at IXP
- Routing Registry
- Used to register the routing policy of the IXP membership (more later)

### 4.3 IXP Structure Design

This section discusses the flow of information from OSIX to the subscriber.

In order for a subscriber to participate at OSIX, they will need:

- a router that supports BGP-4
- a block of IP addresses to advertise - OSIX cannot supply these addresses. They will either have them already or need to get a block address from their ISP or directly from APNIC.
- A public Autonomous System (AS) number, they need to get one from APNIC if they don't already have one. We will not assign one.
- to complete the application form.

Once the subscriber has completed the application form we'll send them some details on the IP address they will use on their router attached to the OSIX and a sample configuration for Quagga on a Linux-based router. This configuration will also be very similar to the configuration need for a Cisco router. If they are using another type of router you'll need to reconfigure it accordingly.

In basic operation, OSIX requires BGP to communicate with subscribers and that all the peers be on the same IP subnet and within the same broadcast domain as it is not easy to multipartite peering if there are routers in the way of some of the peers. Therefore large Metro-Ethernets are good choice for any exchange points. On OSIX we will have two route Linux-based route servers running the Quagga BGP listener and these servers will have identical configuration, but they will be located in different buildings.

Almost anyone can join OSIX, the relatively low cost of entry means there are many peers of all sorts of sizes, as the cost of participating in OSIX is a negligible increase over the basic costs of getting a connection.

The higher cost involved in getting a circuit to OSIX has meant that there is a smaller number of more committed peers.

The OSIX support a couple of modes of operation:

- bilateral peering between consenting parties.
- multilateral peering via the route servers at each exchange.

Bilateral peering tends to suit organisations with legal departments who have time and energy to spend on organising and managing their peering agreements.

Multilateral peering tends to appeal to organisations that don't have a dedicated peering department, especially as the number of peers increases (the n-squared nature of bilateral arrangements bites pretty quickly as the number of peer's increases).

Different organisations therefore have different peering policies. Small ISP's and non ISP organisations tend to favour relying on the route servers (especially in Wellington, where there are now over 80 non-ISP peers) - you expend 5% of the effort, for 90% of the value.

Medium sized ISP's use the route servers, but also use bilateral sessions for peers of extra importance. Small and medium ISP's tend to be fairly promiscuous - they'll peer with anybody who wants to.

Large ISP's prefer to only enter into bilateral sessions with other large ISP's, but with a bit of effort it is possible to get them to bend their definition of **large** so that they'll peer more widely.

### **4.3.1 Basic Peering**

The aim of peering is to get rid of as much traffic as possible, before the subscriber give the remainder to their ISP, and likewise to have as much of their traffic arrived without it needing to pass via their ISP. To achieve this, they need as many peers as possible to listen to their route advertisements, and they need to learn as many routes as possible from other peers.

We will define basic peering as "announcing some routes to the route servers, and listening to everything the route servers send". It does not get the subscriber every possible peer on OSIX, there are some ISP's that only advertise to the route servers, and do not listen to announcements from the route servers, and others that do not peer at all with the route servers.

### **Setting up basic peering**

There are a few things that subscribers need to assemble, before they can start peering on OSIX:

- A connection of reasonable bandwidth of at least 100MB Ethernet, or more preferably 1GB Ethernet to OSIX.
- A router than can handle BGP4, and with enough throughput for the traffic load you expect to put on it. Most peers use Juniper or Cisco routers or even Quagga based routers built on Linux OS Servers.
- An AS number. These can be obtained from APNIC (as our OSIX is in New Zealand), directly through APNIC. On OSIX we do not support the use of private AS numbers. Besides if you plan to do any kind of multi-homing between multiple ISP's, then you definitely want the public AS number.
- Some IP addresses, the more the better. We encourage the /29 subnet and below are accepted by the route servers, but some ISP's only accept /24 and shorter.

Assuming that the above is sorted, then subscriber is ready for basic peering on OSIX. The route servers employ strict filtering on inbound announcements from peers, so subscriber will have to let OSIX know what IP networks they plan to advertise to the route servers. The route servers are configured via RPSL, which allows for automated generation of the BGP configurations from a database driven policy specification.

### **4.3.2 Building OSIX**

Once the subscriber has got basic peering with the route servers going, they will notice fairly quickly that they get some oddities with traffic - depending on who is

listening and who is advertising, traffic can take asymmetric routes, and subscriber will also notice some holes in their route table where they might think certain ISP's should be. The solution is to encourage those ISP's who are not listening to their announcements into the route server to do so.

Often this is relatively straightforward, an email to any local networking user group from OSIX staff asking everybody to listen to the new prefix announcement, followed with a few phone calls can be all it takes to get a new subscribers routes out to many of the ISP's who don't have a default "listen-all" stance. However there are going to be some ISP's who just won't listen to their announcements and unfortunately, these tend to be the biggest ISP's.

In real world, non ISP's that are net sinks of traffic requesting peering are not likely to get much favour by big ISPs as they have little bargaining power, since they basically want something for free, and give little value to the ISP in return for the peering. The lack of favour increases as the disparity between source and sink increases. As rule of thumb, the more of a leech subscribers are, the less likely they are to get peering.

However, if subscribers are a net source of traffic, then they are more likely to get favours in terms of peering. As an example if subscriber is a provider of services that NZ Internet users want to see, then that gives them significant bargaining power over the ISP's. If that traffic is primarily of interest to local users, which means it has little international component to cloud the issue, and especially if that traffic is timing sensitive, meaning it is intolerant of packet reordering or latency variation, then so much the better.

The ideal situation is to offer a high bandwidth service that is performance sensitive to a bunch of users who will put more pressure at their ISP if it does not work well. For example, video conferencing or media streaming is ideal, but sizeable http or ftp traffic will also add value to the case. The content does not need to be available 24x7, but needs to be in demand for enough time for the ISPs to notice that are being affected.

Once the basic peering is established, then subscribers negotiate access to content providers. Subscribers can announce to their region via local user groups, for example, in New Zealand its known as New Zealand Network Operators Group - NZNOG and even possibly via direct contact with ISPs that they expect significant local demand for their content, and that they should listen to the announcements for their prefix on OSIX.

Once the subscribers are on the exchange, they are in a good position to organise peering sessions with ISPs where they deliver subscriber a full route table, or at least a default route. This provides redundant routes in the event of a failure.

### **4.3.3 How it works at OSIX?**

There are two route servers at OSIX which hold details of the routes exchanged. These servers filter routes in two ways:

- Incoming announcements are checked against individual filter lists for each peer.
- Outgoing announcements are filtered to make sure that routes are not announced for RFC1918 addresses and other bogon or martian addresses.

The filter lists and other details of the route server configuration are built automatically from an RPSL database using the RtConfig tool from the IRRToolset. In case if subscriber need to change the prefixes they want to announce, OSIX has a web page where they can do this online.

The OSIX Routing Registry data is available either by using a standard “whois” client connected to: `whois.osix.net:43` or using a web based interface.

For example, subscriber can examine the OSIX routing policy by looking up the details for AS7675 for route server 1 and AS8285 for route server 2.

Each OSIX peer can check the list of networks that they are allowed to announce to the OSIX by looking up the relevant route-set. The peers at OSIX should comply with the terms and conditions in Appendix F.

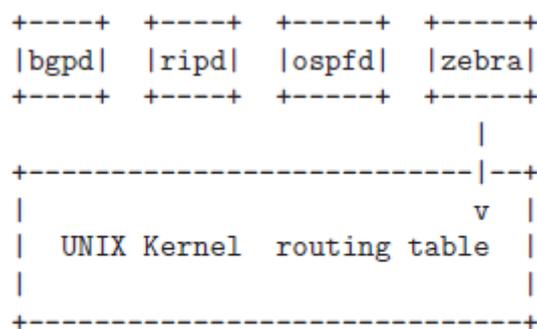
## 4.4 Route Server Design

We have built our route server using Quagga on Ubuntu Linux server. Quagga is a routing software package that provides TCP/IP based routing services with many routing protocols support, especially BGP.

There are many other open-source routing daemons available under GNU license, but we have decided to use Quagga, because of its easy-to-install and easy-to-administer in complex environments.

Another reason for choosing Quagga is that it uses advanced software architecture to provide a high quality, multi-server routing engine. Quagga has an interactive user interface for each routing protocol and supports common client commands. Due to this design, we can add new protocol daemons to Quagga easily. We can also use Quagga library as our program's client user interface. Quagga is distributed under the General Public License GNU.

Traditionally routing software is made as a one process program which provides all of the routing protocol functionalities. Quagga takes a different approach. It is made from a collection of several daemons that work together to build the routing table. There may be several protocol-specific routing daemons and zebra the kernel routing manager. Though in our design, will only need BGPd daemon, the following figure shows how the different routing daemon runs on top of Linux kernel.



Quagga System Architecture  
Figure 21 – Quagga System Architecture

Zebra is an IP routing manager. It provides kernel routing table updates, interface lookups, and redistribution of routes between different routing protocols.

BGP stands for a Border Gateway Protocol. The latest BGP version is 4. It is referred as BGP4. BGP4 is one of the Exterior Gateway Protocols and de-fact standard of Inter-Domain routing protocol. BGP4 is described in RFC1771, A Border Gateway Protocol 4 (BGP-4). Many extensions have been added to RFC1771. RFC2858, Multiprotocol Extensions for BGP-4 provides multiprotocol support to BGP-4.

The AS (Autonomous System) number is one of the essential element of BGP. BGP is a distance vector routing protocol, and the AS-Path framework provides distance vector metric and loop detection to BGP. RFC1930, Guidelines for creation, selection, and registration of an Autonomous System (AS) provides some background on the concepts of an AS.

The AS number is a two octet value, ranging in value from 1 to 65535. The AS numbers 64512 through 65535 are defined as private AS numbers. Private AS numbers must not to be advertised in the global Internet.

At an Internet Exchange point, many ISPs are connected to each other by external BGP peering. Normally these external BGP connection are done by 'full mesh' method. As with internal BGP full mesh formation, this method has a scaling problem.

This scaling problem is well known. Route Server is a method to resolve the problem. Each ISP's BGP router only peers to Route Server. Route Server serves as BGP information exchange to other BGP routers. By applying this method, numbers of BGP connections is reduced from  $O(n*(n-1)/2)$  to  $O(n)$ .

Unlike normal BGP router, Route Server must have several routing tables for managing different routing policies for each BGP speaker. We call the routing tables as different views. bgpd can work as normal BGP router or Route Server or both at the same time.



The following public IP subnets are taken from VUW as part of thesis to provide connectivity to the outside world:

VUW IPv4 Subnet: 130.195.10.0/24

VUW IPv6 Subnet: 2404:2000:2000:10::/64

The following diagram shows the logical flow of peering at OSIX:

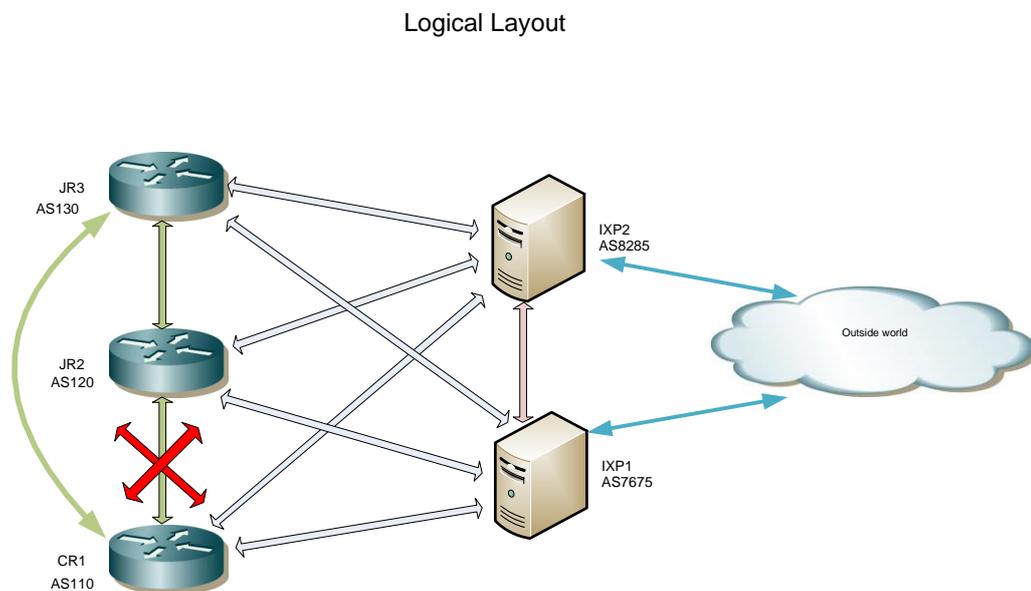


Figure 23 – OSIX Logical Layout

As per above diagram, you can see that following AS numbers are allocated to route servers and the three peers:

IXP1 (Route Server1) – AS 7675

IXP2 (Route Server2) – AS 8285

CR1 – AS 110

JR2 – AS 120

JR3 – AS 130

## 4.5 Policy Design

In order to build our policy design at OSIX, we are using the combination of IRRd and RtConfig (a tool in IRRToolset). IRRd is a stand-alone Internet Routing Registry

database server. IRRd can store information and answer queries about local network, campus and ISP backbone topology, address allocation and routing policies.

IRRD can be used as an independent local database server, or as part of the global Internet Routing Registry (IRR). IRRd supports the Routing Registry Specification Language (RPSL) routing registry syntax.

As of version 2.2.0, IRRd also supports the RPSLng IPv6 and Multicast extensions to RPSL. The IRRd distribution includes all needed IRR support services, including: automated real-time mirroring of other IRR databases, update syntax checking, update security checking, and update notification. The current version of IRRd also supports several RIPEdb whois flags.

When IRRd is used in conjunction with policy tool RtConfig, it allows:

- Automated generation of router configuration files and access-lists
- Internet topology visualization
- Network trouble-shooting and debugging

In addition to user-oriented whois queries, the IRRd Server also provides several query commands for performing RPSL set expansions and AS number to route prefix mappings which are useful for automated tools (such as IRRToolSet). The IRRd distribution also includes the `irr_rpsl_submit` e-mail/TCP front-end update program which performs RPSL syntax and authentication checking.

IRRD is also running as a daemon on our IXP server at OSIX along with Quagga and many other applications. See the figure below for the complete integration of the process environment at OSIX Servers.

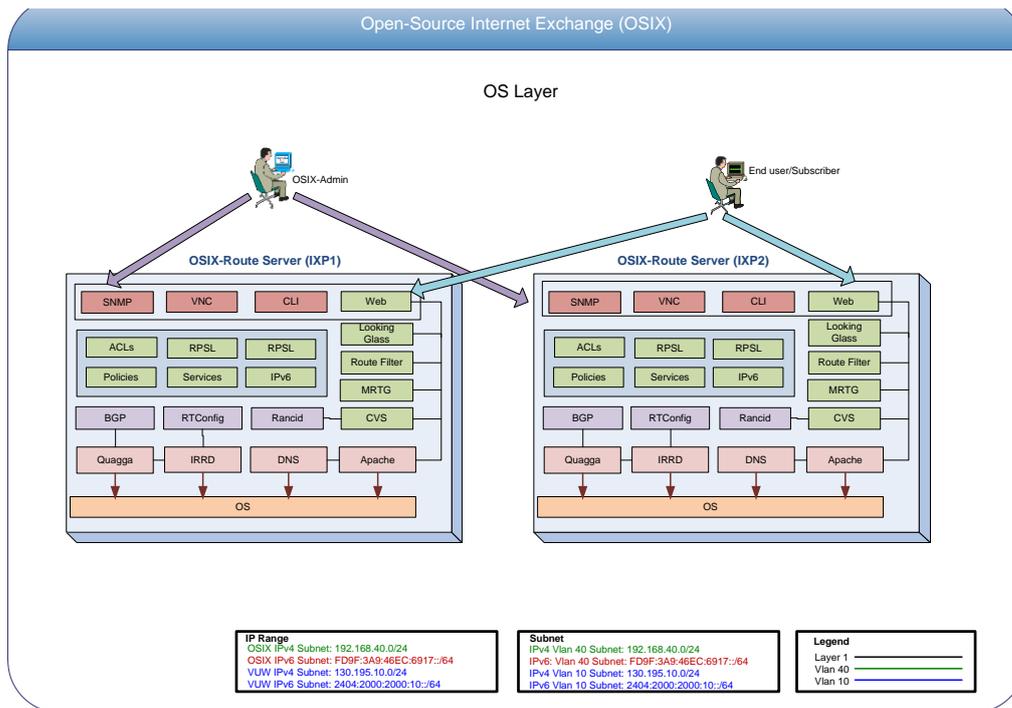


Figure 24 – OSIX OS Layer Process View

The following figure represents on how the OSIX configuration system interacts with all the other processes and entities.

As the OSIX configuration is the central part of our system, it is based on the standards set by ‘routing policy’. Routing policy also generates the IRRD database using RPSL script.

IRRD then generate a complete configuration for OSIX using RtConfig. RtConfig analyses the routing policies registered in the Internet Routing Registry (IRR) and produces router configuration files. It currently supports Cisco and Junos router configuration file formats.

RtConfig reads lines form the standard input, and prints them to the standard output, except for the lines that start with "@rtconfig" which instruct rtconfig to perform special operation. RtConfig establishes a whois connection to query IRR.

The OSIX Config System then generates a ‘bgp.conf’ file with the information for all the peers and it can be used by Quagga. See figure below for complete process flow:

RPKI – Resource Public Key Infrastructure is a security framework for verifying the association between resource holders and their Internet resources via resource certification. In this context, 'resource holders' are organizations such as Regional Internet Registries (RIRs), Internet Service Providers (ISPs), or end-user organizations, while 'Internet resources' are IPv4 and IPv6 address blocks and Autonomous System Numbers (ASNs). We have not implemented RPKI in our system, but this is something we can look at as part of our future enhancements.

Process Flow Diagram

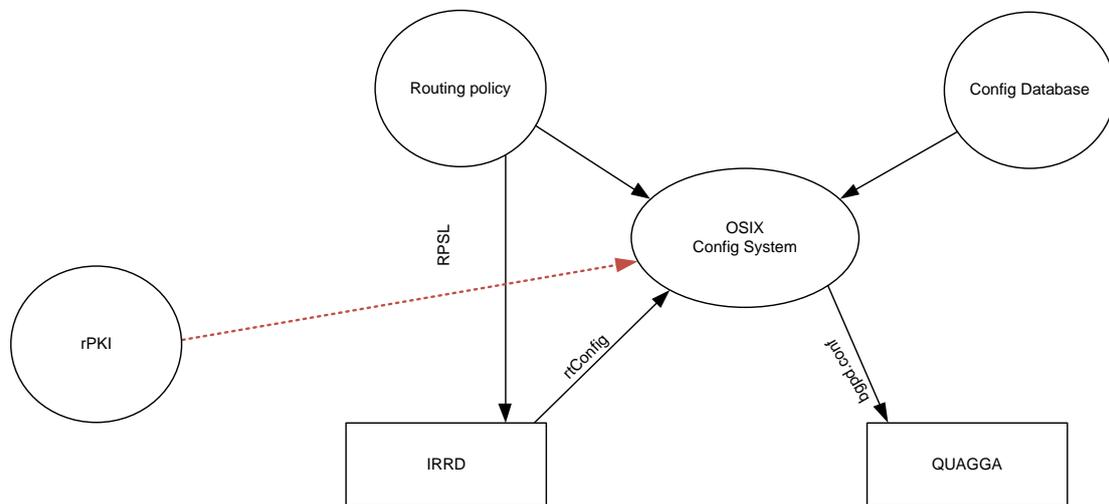


Figure 25 – OSIX Process Flow Diagram

## 4.6 Interface Design

Based on our survey results and our experience, we believe that users have different preferences in accessing the information they are looking for. In order to design the user interface for OSIX Management servers, we have made it accessible to administrators via Command Line Interface (CLI), Linux desktop accessible via console and VNC – Virtual Network Computing, SSH with tunnelling or local console and via website.

The OSIX peers can only access some features via public website which can help them in peering. For example, they can run database queries using Perl script on website and they can also “show routes” script on website as well.

OSIX peers can also generate configurations for their peering routers for both Cisco and Juniper routers.

There are different needs for all three interfaces and they are all justified as they can all serve very useful in different scenarios.

### **CLI**

- This interface is only available to OSIX Administrator
- It is only accessible via SSH (Secure Shell Console)
- It is needed to access Quagga Router Daemon
- It is also need access IRRd Daemon
- Several commands and scripts can only be run via CLI, so accessibility to this interface is easily justified.

### **VNC/Desktop**

- This interface is also available only to OSIX Administrators
- This is a handy feature, in case if Administrator needs to run utilities like, Wireshark to capture packets on network or need access to Firefox to test scripts.
- Since the interface is only available to Administrator, it has unlimited access, just like shell.

### **Website**

- This interface is available to OSIX peers and also to general public.
- OSIX website will also have all the features and information that OSIX peers will need in order to peer with OSIX.
- Features like, show-routes, who-is queries, generation of prefix-lists will be available on website for general public.

We will design OSIX website in simple html and it will be mostly text-based. We will not be using flash or any other heavyweight tools.

Besides this website is only targeting ISPs and Network Engineers, who are only here for information on how to peer with us and they prefer to get information in simple-text based style.

## 5. Implementation

This chapter describes the implementation of OSIX given the design described in the previous chapter. As part of our implementation of the OSIX, we have used the design as we have discussed earlier and we will discuss with reasoning on implementing each component of implementation.

### 5.1 Implementation Based on Design

In order to build the OSIX System, the most difficult part was to choose the platform where we can build the entire system and that can be easily re-deployed anywhere in the world with minimum changes once it is built. Therefore, we decided to use a virtual machine (VMware) to build our operating system. We have built OSIX system on VMware as it supports our selected operating system, Ubuntu Linux, which we have modified structurally and installed all the required applications. The other reason to use VMware is because of its portability and backward compatibility with all versions and different brands. The VMware image is also compatible with other freely available virtualisation software like Virtual Box.

We have used the following open-source applications under GNU Public License to build the OSIX system

- |                             |                                  |
|-----------------------------|----------------------------------|
| 1. Operating System:        | Ubuntu Linux 11.04 32-bit server |
| 2. Route Server:            | Quagga                           |
| 3. Web Server:              | Apache 2.2                       |
| 4. Configuration Server:    | Rancid/CVS                       |
| 5. Graph Tool:              | MRTG/RRDTool                     |
| 6. Configuration Generator: | RtConfig (IRRToolset)            |
| 7. Route Registry Daemon:   | IRRD                             |
| 8. Scripts:                 | OSIX RPSL Scripts                |

The following is the detailed description of each tool and its purpose in our system:

### **Ubuntu Linux 11.04 Server:**

We have decided to use Ubuntu Linux 11.04 because of its rich repository and its easy-to-install features. This is very important as the final by-product of our OSIX will be available to anyone who wishes to use at their exchange and having a friendly OS which can be configured and managed with minimum efforts is a big advantage. The version 11.04 is the most stable version at the time of deployment.

### **Quagga 0.99.20.1**

Quagga is built on the old routing daemon known as Zebra. Quagga is routing software that provides TCP/IP based routing services with routing protocols support such as RIP, OSPF and BGP4. Quagga also supports BGP Route Reflector and Route Server. In addition to traditional IPv4 routing protocols, Zebra also supports IPv6 routing protocols as well as SNMP.

### **Apache 2.22**

Apache is an open-source web server that can be easily installed on Linux OS via its repository and it can host web applications and scripts on the server. In OSIX, it is used to host the intranet webpage as well as public facing Internet website for anyone to retrieve peering info for OSIX.

### **Rancid 2.3.6/CVS**

Rancid (Really Awesome New Cisco confIg Differ) is a tool monitors a router or switch configuration, including software and hardware (cards, serial numbers, etc) and uses CVS (Concurrent Version System) or Subversion to maintain history of changes. CVS can also reside on a web server and its configuration repository can be accessible via Webpage.

### **MRTG 2.17.4/RRDTool**

MRTG (The Multi Router Traffic Grapher) is a tool to monitor the traffic load on network links. MRTG runs a web script which can reside on Web Server and generates HTML pages containing PNG images which provide a real-time visual representation of network traffic. This is a very useful tool as it provided the traffic

through put from peers connected at OSIX. MRTG access the devices using SNMP. RRDtool is the OpenSource industry standard, high performance data logging and graphing system for time series data.

### **RtConfig/IRRToolset 5.0.0:**

RtConfig is a part of IRRToolset (Internet Routing Registry Toolset) which is a suite of policy analysis tools to operate with routing policies in RPSL [RFC 2622] format, registered in Internet Routing Registry (IRR). RtConfig is used to make router configuration files, based on policies registered in Internet Routing Registries (IRR).

### **IRRD 2.3.10**

IRRD is a stand-alone Internet Routing Registry database server. IRRd can store information and answer queries about local network, campus and ISP backbone topology, address allocation and routing policies. IRRd can be used as an independent local database server, or as part of the global Internet Routing Registry (IRR).

IRRD supports the Routing Registry Specification Language (RPSL) routing registry syntax. As of version 2.2.0, IRRd also supports the RPSLng IPv6 and Multicast extensions to RPSL. The IRRd distribution includes all needed IRR support services, including: automated real-time mirroring of other IRR databases, update syntax checking, update security checking, and update notification. The current version of IRRd also supports several RIPEdb whois flags.

When used in conjunction with policy tools such as RtConfig, Roe, and Aoe, the IRRd server allows:

- Automated generation of router configuration files and access-lists
- Internet topology visualization
- Network trouble-shooting and debugging

In addition to user-oriented whois queries, the IRRd Server also provides several query commands for performing RPSL set expansions and AS number to route prefix mappings which are useful for automated tools (such as IRRToolSet). The IRRd

distribution also includes the irr\_rpsl\_submit e-mail/TCP front-end update program which performs RPSL syntax and authentication checking.

### **OSIX Scripts**

These scripts are used to communicate between RtConfig and IRRd. They also store the prefix-lists and generate configuration file based on policies defined for both IPv4 and IPv6.

All the components in our implementation plan are tested successfully in chapter ^ - OSIX Testing and all the results are recorded in Appendix G – Test Cases.

## **5.2 Operating System Selection and Implementation**

We have two OSIX Management Servers that run all the applications required to run the Exchange Point. Both the servers are configured in exactly the same way and each subscriber needs to peer with each server to provide full resiliency and redundancy.

The selection of OS was a difficult choice as we wanted to make sure its compatible with all the applications we need to run for OSIX and we also wanted to make sure it's easy to manage and its widely available.

We have decided to use Ubuntu server as it meets our requirement in terms of its adaptability to install most of the applications through its repository, which makes the installation not only easier but also easy to upgrade and migrate to newer versions.

### **Kernel Version:**

We are running the following kernel on our OS

```
iqbalaun@ixp1:~$ uname -a
Linux ixp1 2.6.38-15-generic-pae #61-Ubuntu SMP Tue Jun 12 19:32:42 UTC 2012 i686 i686
i386 GNU/Linux
iqbalaun@ixp1:~$
```

Listing 5 – Linux Kernel Running on OSIX Servers

### **User Authentication:**

The system is only accessible via local users created on the server and they are the “admin accounts” of OSIX Administrators. No peers or third party shall be allowed to access servers.

**Services:**

We have also taken measures to ensure that only the required services are running on server and any unwanted services are disabled to minimise the usage of memory and CPU.

**Port Accessible:**

We have run the “nmap portscan” to ensure that all the unwanted ports are blocked and only the desired ports are allowed to run on the server.

### 5.3 Accessing OSIX

The operating system on both IXP1 and IXP2 is only accessible via following three methods:

Service	Port	Description
SSH	22	Access to CI
VNC	5900	Access to Desktop
Web	10000	Access via Website

Table 2 – Accessing OSIX

In order to access OSIX via website, we have installed and configured a third-party package, named WebMin which allows the Ubuntu Linux OS accessible via web interface.

The accessibility of the system to desktop is also available via VNC. This feature is enabled as we are running several web scripts on the server and from the administration point-of-view; we may need to test several applications via GUI applications, such as Firefox or Wireshark.

### 5.4 Route Server Implementation

The Route Servers are built in Quagga version 0.99.20.1. We have built two route servers, one on each IXP Management Server

Both Route Servers are accessible via Telnet from the IXP Management Servers. OSIX Admin can telnet to BGP Daemon on port 2605 and Zebra Daemon is accessible via Telnet to port 2601

The Route Server on IXP1 has AS#: 7675

The Route Server on IXP2 has AS#: 8285

Each subscriber at OSIX must peer with both Route Servers using both IPv4 and IPv6 subnets using BGP4.

Route Servers do not communicate each other via BGP and they do not have any load-balancing configured, but they can be redundant to each other as they both peer with all the subscribers at OSIX.

As per our network design, we have also connected three routers as subscribers. All three routers are connected to both route servers and they all have a unique AS number. See the Network Diagram in Design Section for the details.

### **IXP Management Switch**

The IXP Management switch is configured as a standard Layer 2 switch and it does not perform any routing. The switch we have used in OSIX setup is a Cisco WS-C2950T with 24 Ports capable of speed up to 100Mbps.

## **5.5 Policy Implementation**

As per our design, we are storing the policies at the IRRd server, which is accessible via telnet to port 43.

By running the “sudo make” in osix or osix6 folder, it can generate the “bgp.conf.loaded” file in both IXP1 (m91) and IXP2 (m92) directories.

The lists of all the IPv4 peers are stored under osix.peers flat-file and the lists of all the IPv6 peers are stored under osix6.peers in flat file

Once we execute “sudo make” command from parent OSIX or OSIX6 folder it looks for all listed peers in osix.peers file and generates “.rpsl” file for both in IXP1 and IXP2 Route servers.

The script then tells the bgp.conf.loaded to read the respective “.rpsl” file. The RPSL file is built on the policy set by RtConfig.

RtConfig then uses the RPSL file and based on the IRR Policy stored in IRRd, it generates the bgpd.conf file which is the final version of the file ready to be used for the OSIX route servers.

## **5.6 Web Server Implementation**

We are running Apache2 as our web server and the following applications are hosted on the server and can be accessed by any web browser.

- CVSweb/Rancid
- Looking Glass
- MRTG Graph tool
- OSIX Route Queries

### **CVSweb/Rancid**

RANCID (Really Awesome New Cisco config Differ) monitors a router's or generally a device's configuration, including software and hardware (cards, serial numbers, etc) and uses CVS (Concurrent Version System) or Subversion to maintain history of changes. Rancid supports Cisco routers, Juniper routers, and many others.

CVS is a version control system designed for software projects. The CVS can have multiple users simultaneously online and working on a project, also in a file. The role of the CVS is to make the changes in the source code (including bugs) traceable to make documented. At the same time, older versions are saved and restored.

RANCID does this by the very simple process summarized here:

- login to each device in the router table (router.db),
- run various commands to get the information that will be saved,
- cook the output; re-format, remove oscillating or incrementing data,
- email any differences (sample) from the previous collection to a mail list,
- and finally commit those changes to the revision control system

Rancid is a Perl script which uses the login scripts (clogin) to login to a device, execute commands to display the configuration, etc, then filters the output for formatting and security by omitting any confidential data such as passwords.

We are using Rancid to take the backup of configurations from both Cisco and Juniper routers as they are the two most popular brands used by peers nowadays and we have used both brands as peers in OSIX.

Rancid can be executed using a command “rancid-run”. This command will look for all the devices listed in “router.db” file in both Juniper and Cisco directories and pull the configurations for each device. The login information for each device is stored in .cloginrc file in home folder. Once the configurations have been retrieved, we can run rancid-cvs command to update the CVS file.

The above process of taking backups using Rancid can be scheduled to run automatically at weekly or monthly frequency as well. The configuration files can be accessed on both IXP1 URL at <http://130.195.10.91/cgi-bin/cvsweb> and IXP2 using a web URL at <http://130.195.10.91/cgi-bin/cvsweb>

### **Looking Glass**

Looking Glass (LG) is a term used to describe servers on the Internet running publicly available Looking Glass software implementations in BGP. LG is accessed remotely for the purpose of viewing routing information. This is a read-only portal to routers of any organization running the LG server. This feature does not give access to the routers and it only shows the IPv4 and IPv6 BGP routes within OSIX. The other

features like, trace and ping can be disabled by simply removing it from the lg.cgi script.

We have implemented LG in both IXP1 and IXP2 and it will be allowed public access for viewing BGP routes.

### **MRTG Graph tool:**

Multi Router Traffic Grapher (MRTG) is a tool to monitor the traffic load on network links. MRTG generates HTML pages containing PNG images which provide a LIVE visual representation of this traffic.

MRTG consists of a Perl script which uses SNMP to read the traffic counters of your routers and a fast C program which logs the traffic data and creates beautiful graphs representing the traffic on the monitored network connection. These graphs are embedded into Webpages which can be viewed from any modern Web-browser.

In addition to a detailed daily view, MRTG also creates visual representations of the traffic seen during the last seven days, the last five weeks and the last twelve months. This is possible because MRTG keeps a log of all the data it has pulled from the router.

This log is automatically consolidated so that it does not grow over time, but still contains all the relevant data for all the traffic seen over the last two years. This is all performed in an efficient manner. Therefore you can monitor 200 or more network links from any halfway decent UNIX box.

MRTG is not limited to monitoring traffic, though. It is possible to monitor any SNMP variable you choose. You can even use an external program to gather the data which should be monitored via MRTG. People are using MRTG, to monitor things such as System Load, Login Sessions, Modem availability and more. MRTG even allows you to accumulate two or more data sources into a single graph.

## **OSIX Route Queries**

This feature allows users to request route queries from web and it is also available to anyone on Internet. We have allowed users to perform the following three queries.

1. Who Is (whois) query
2. Show Route query
3. Change Route query

### **Query: whois**

This query searches the OSIX Route Registry Database for any AS number submitted by user at the link: <http://130.195.10.91/route/whois.pl>

This feature is very simple to use and user can simply include the "AS" in front of an AS number. Example: AS110. Users are also allowed to use parameters in their search to get more precise results. For example, by including "-r" (no recursion) in front of the search key is to show only the object that matches the search key.

Additional parameters like "-t" (template only) or "-v" (template and description) in front of an object name will display the template

Example: -t inetnum

### **Query: show route**

This feature uses the RtConfig tool from the Internet Routing Registry Toolset to display the filters for the routes which can be imported to and exported from the two OSIX route servers (AS7675 and AS8285).

Users can also choose to display a list of routes which the Route Servers will accept in a BGP announcement from any of the ASes connected to OSIX route servers. The default format is 'Cisco prefix-list' but you can choose to see them in 'Cisco access list' or JunOS format if you prefer. This page can be accessed at the URL <http://130.195.10.91/route/ShowRoutes.cgi>

**Query: change route**

This query takes a number of parameters including your AS number and a list of ALL the routes that user want to export to the OSIX Route Servers from that AS. It creates a set of RPSL objects for submission to the NZRR Route Registry via email.

If the user have more than one router peering with the OSIX Route Servers, than they have to make sure the list of routes includes ALL the routes their AS will announce from ALL of their routers. This page can be accessed at the URL

<http://130.195.10.91/route/ChangeRoutes.cgi>



## **6. OSIX Testing**

We have performed testing on both OSIX servers (IXP1 and IXP2), since both servers are exact replica of each other and we have tested each feature in on the two servers, but in our testing we have supplied evidence from either one of the two servers.

### **6.1 Test Infrastructure**

The testing for OSIX infrastructure consists of Linux Servers, Cisco and Juniper Routers and Cisco Switch.

The following equipment is used in the test environment:

2x Ubuntu Linux 11.04 Servers

2x Juniper Routers

1x Cisco 2950T-24 Switch

The following equipment's are needed to perform the tests:

- PC/server with at least 100MB Ethernet NIC

### **6.2 Test Cases**

The following test cases have been used for the testing of OSIX functionality to ensure that all components are implemented as per design. All testing was completed successfully. See Appendix G for the Test case results.

The following test cases have been used for the testing of OSIX functionality to ensure that all components are implemented as per the design.

Test Case 1: VMware Image Replication for the two OSIX Servers

Test Case 2: Cold Start to servers and ensure all services come back up smoothly

Test Case 3: Warm Start to servers and ensure all services come back up smoothly

Test Case 4: Secure Shell Login to IXP1 and IXP2 on Port 22

Test Case 5: VNC Login to IXP1 and IXP2 on port 5900

Test Case 6: Webpage Login to IXP1 and IXP2 on port 10000  
Test Case 7: Telnet Authentication to Quagga routing Daemon on both Route Servers  
Test Case 8: BGP Establishment with the three peers from IXP1 and IXP2  
Test Case 9: Authentication to IRRd daemon via telnet to port 43  
Test Case 10: IPv4 Ping test to CR1, JR2 and JR3 to ensure end-to-end connectivity  
Test Case 11: IPv6 Ping test to CR1, JR2 and JR3 to ensure end-to-end connectivity  
Test Case 12: SNMP Walk to all the devices.  
Test Case 13: Generate Backup for both Cisco and Juniper using Rancid  
Test Case 14: CVS Repository is displaying all the updated configuration files  
Test Case 15: Apache Web Server is UP and hosting all required applications  
Test Case 16: MRTG real-time graphs are generated for all three peers at OSIX  
Test Case 17: OSIX RPSL scripts can generate configuration based on RPSL policy  
Test Case 18: Looking Glass can display IPv4 BGP Summary for all peers  
Test Case 19: Looking Glass can display IPv6 BGP Summary for all peers  
Test Case 20: Users can request “whois” query from OSIX webpage  
Test Case 21: Users can request “show route” query from OSIX webpage  
Test Case 22: Users can request “change route” query from OSIX webpage  
Test Case 23: Verify all logs are being stored at OSIX Log Server  
Test Case 24: Ensure only required services are running

### **Test Case 1: VMware Image Replication for the two OSIX Servers**

#### **Description:**

VMware Image portability for the two OSIX Servers

#### **Pre-requisites:**

- Fully configured PCs with VMware servers installed.
- Console Login to VMware servers.

#### **Procedure:**

- Taking snapshot of VMware image on VMware servers.

- Copy the VMware image on a different VMware server location and load the VMware image.

**Acceptance Criteria:**

- The VMware image files are copied at different locations with no errors.
- The OS installed Ubuntu 11.04 server loads up with no errors.
- The Ethernet interface on server's status is up.
- The line interface status is up.
- Ideally no input, CRC, framing, or output errors are seen on interfaces.

**Test Case 2: Cold Start to servers and ensure all services come back up smoothly**

**Description:**

Warm Start to servers and ensure all services come back up smoothly

**Pre-requisites:**

- Console access to the OSIX Servers.

**Procedure:**

- Take snapshot of the VMware image before rebooting the server
- System to be rebooted after graceful shutdown.

**Acceptance Criteria:**

- System returns back to power with all configurations.
- All the interfaces stand UP.
- All the daemons and services are resumed without manual start.
- Ideally no input, CRC, framing, or output errors are seen on interfaces.

**Test Case 3: Warm Start to servers and ensure all services come back up smoothly**

**Description:**

Cold Start to servers and ensure all services come back up smoothly

**Pre-requisites:**

- Take snapshot of the VMware image before rebooting the server
- System to be rebooted after graceful shutdown.
- End of test.

**Procedure:**

- System to be shut down by unplugging the power chord.
- Enter the **show interface** command on the switches to verify the vPC Peer Link interfaces status.

**Acceptance Criteria:**

- System returns back to power with all configurations.
- All the interfaces stand UP.
- All the daemons and services are resumed without manual start.
- Ideally no input, CRC, framing, or output errors are seen on interfaces.

**Test Case 4: Secure Shell Login to IXP1 and IXP2 on Port 22****Description:**

Secure Shell Login to IXP1 and IXP2 on Port 22

**Pre-requisites:**

- System is up and running.
- Any remote client like, putty or SecureCRT
- Login account to OSIX servers
- Remote access to servers

**Procedure:**

- Login to the OSIX servers remotely via SSH on port 22.
- Make sure users cannot login to the OSIX servers via Telnet.

**Acceptance Criteria:**

- The correct cable types and transceivers are installed.
- The physical interface status is up.
- The line interface status is up.
- Ideally no input, CRC, framing, or output errors or those errors are less than 1% of traffic.

**Test Case 5: VNC Login to IXP1 and IXP2 on port 5900****Description:**

VNC Login to IXP1 and IXP2 on port 5900

**Pre-requisites:**

- VNC Client
- Remote connection to OSIX Servers

**Procedure:**

- Make sure that OSIX servers allow remote login to Ubuntu Desktop.

**Acceptance Criteria:**

- Remote connection is established.
- User can browse and make changes to the desktop.

**Test Case 6: Webpage Login to IXP1 and IXP2 on port 10000****Description:**

Webpage Login to IXP1 and IXP2 on port 10000

**Pre-requisites:**

- Remote access to OSIX servers
- Client must have a web browser, like IE or Firefox.

**Procedure:**

- Open browser from the client and login via https to port 10000 on both OSIX servers.
- End of test.

**Acceptance Criteria:**

- User is able to login using their same login ID as they used to login via SSH.

**Test Case 7: Telnet Authentication to Quagga routing Daemon on both Route Servers****Description:**

Telnet Authentication to Quagga routing Daemon on both Route Servers

**Pre-requisites:**

- Quagga services must be running on OSIX servers.

**Procedure:**

- Login to Quagga server via telnet to localhost on port 2605.

**Acceptance Criteria:**

- User can login in successfully..

**Test Case 8: BGP Establishment with the three peers from IXP1 and IXP2****Description:**

BGP Establishment with the three peers from IXP1 and IXP2

**Pre-requisites:**

- BGP Configuration has been applied on both OSIX Route servers and peers.

**Procedure:**

- Run “sh ip bgp summary” command on Route server.

**Acceptance Criteria:**

- All the peers are established with OSIX route servers.

**Test Case 9: Authentication to IRRd daemon via telnet to port 43****Description:**

Authentication to IRRd daemon via telnet to localhost on both OSIX servers

**Pre-requisites:**

- IRRD must be installed and fully configured.
- Port must be enabled and defined in “services”

**Procedure:**

- Telnet to localhost using command “telnet localhost irrd”.

**Acceptance Criteria:**

- User can login to the telnet to the IRRd

**Test Case 10: IPv4 Ping test to CR1, JR2 and JR3 to ensure end-to-end connectivity****Description:**

IPv4 Ping test to CR1, JR2 and JR3 to ensure end-to-end connectivity

**Pre-requisites:**

- SSH to the OSIX server
- Run the ping test from the server to all three peers.

**Procedure:**

- Use the command “ping <ip address> from CLI

**Acceptance Criteria:**

- Ensure that all ICMP packets are sent to subscribers with no packet loss.

**Test Case 11: IPv6 Ping test to CR1, JR2 and JR3 to ensure end-to-end connectivity****Description:**

IPv6 Ping test to CR1, JR2 and JR3 to ensure end-to-end connectivity

**Pre-requisites:**

- SSH to the OSIX server
- Run the ping test from the server to all three peers.

**Procedure:**

- Use the command “ping6 <ipv6 address> from CLI

**Acceptance Criteria:**

- Ensure that all ICMP packets are sent to subscribers with no packet loss.

**Test Case 12: SNMP Walk to all the devices.****Description:**

SNMP Walk to all the peers at OSIX.

**Pre-requisites:**

- Ensure SNMP utility is installed on OSIX servers.
- SNMP Servers is defined at subscribers.

**Procedure:**

- Run the command “snmpwalk -v2c <ip address> -c public” from servers to all three subscribers

**Acceptance Criteria:**

- OSIX servers must be able to snmpwalk to all subscribers.

**Test Case 13: Generate Backup for both Cisco and Juniper using Rancid****Description:**

Generate Backup for both Cisco and Juniper Routers using Rancid

**Pre-requisites:**

- Rancid must be installed and configured on the OSIX servers
- Users must have the permissions to run the Rancid commands

**Procedure:**

- Run the command 'bin/rancid-run' from the /var/lib/rancid folder
- Make sure the new config file has been updated in their respective folders

**Acceptance Criteria:**

- The command should be executed with errors.
- Configuration file must be backed up from peers.

**Test Case 14: CVS Repository is displaying all the updated configuration files****Description:**

Apache Web Server is UP and hosting all required applications

**Pre-requisites:**

- Apache2 web server must be installed on OSIX servers.

**Procedure:**

- Run command "ps aux | grep apache" to ensure services is running.
- Restart Apache server by using command "sudo /etc/init.d/apache2 restart"

**Acceptance Criteria:**

- The Apache server should be running by default.
- The server should come back up after restarting it.

**Test Case 15: Apache Web Server is UP and hosting all required applications****Description:**

CVS Repository is displaying all the updated configuration files

**Pre-requisites:**

- Backup has been taken using Rancid
- CVSWEB must be running and hosted on Apache Web server

**Procedure:**

- Access the CVSWEB URL
- Access the configuration files in CVS Repository and make sure the new version is updated.

**Acceptance Criteria:**

- The new version of configuration file is created
- CVS Repository must show the differences between the configurations.

**Test Case 16: MRTG real-time graphs are generated for all three peers at OSIX****Description:**

MRTG real-time graphs are generated for all three peers at OSIX

**Pre-requisites:**

- MRTG has been installed and running on Apache Server
- OSIX subscribers are accessible via SNMP

**Procedure:**

- Access the URL for MRTG on OSIX servers

**Acceptance Criteria:**

- The daily, monthly, weekly and yearly graphs should give the real-time statistics

**Test Case 17: OSIX RPSL scripts can generate configuration based on RPSL policy**

**Description:**

OSIX RPSL scripts can generate configuration based on RPSL policy

**Pre-requisites:**

- RPSL Scripts are generated and assigned the correct permissions.
- All required Perl libraries are installed.

**Procedure:**

- Run the make command from “/usr/local/src/rpsl/osix” folder.

**Acceptance Criteria:**

- The “bgpd.conf” file is generated for both OSIX servers

**Test Case 18: Looking Glass can display IPv4 BGP Summary for all peers**

**Description:**

Looking Glass can display IPv4 BGP Summary for all peers

**Pre-requisites:**

- Looking Glass scripts are installed and correct permissions assigned.
- Scripts are modified to be able to login to OSIX peers

**Procedure:**

- Access LG from URL and ensure it comes up properly.

**Acceptance Criteria:**

- User must be able to see BGP Summary for all peers at OSIX

**Test Case 19: Looking Glass can display IPv6 BGP Summary for all peers****Description:**

Looking Glass can display IPv6 BGP Summary for all peers

**Pre-requisites:**

- Looking Glass scripts are installed and correct permissions assigned.
- Scripts are modified to be able to login to OSIX peers

**Procedure:**

- Access LG from URL and ensure it comes up properly.

**Acceptance Criteria:**

- User must be able to see only IPv6 BGP Summary for all peers at OSIX

**Test Case 20: Users can request “whois” query from OSIX webpage****Description:**

Users can request “whois” query from OSIX webpage

**Pre-requisites:**

- “Whois” web scripts are installed and accessible via webpage.

**Procedure:**

- Access the whois.pl via webpage
- User should be able to perform query to OSIX subscribers

**Acceptance Criteria:**

- Users can successfully query the OSIX peers.

## **Test Case 21: Users can request “show route” query from OSIX webpage**

### **Description:**

Users can request “show route” query from OSIX webpage

### **Pre-requisites:**

- “show route” web scripts are installed and accessible via webpage.

### **Procedure:**

- Access the “showroute.cgi” via webpage
- Users should be able to generate prefix-lists to all OSIX peers in both Juniper and Cisco format.

### **Acceptance Criteria:**

- Users can successfully import and export prefix-lists in both Cisco and Juniper format.

## **Test Case 22: Users can request “change route” query from OSIX webpage**

### **Description:**

Users can request “change route” query from OSIX webpage

### **Pre-requisites:**

- “Change route” web scripts are installed and accessible via webpage.

### **Procedure:**

- Access the “changeroute.cgi” via webpage
- Users should be able request OSIX to update or change any routing information.

### **Acceptance Criteria:**

- The page should allow users to submit request for changing routes.

### **Test Case 23: Verify all logs are being stored at OSIX Log Server**

**Description:**

Verify all logs are being stored at OSIX Log Server

**Pre-requisites:**

- All possible loggings must be enabled to save logs at “/var/logs” directory.

**Procedure:**

- Go to the “/var/logs” directory and ensure that logs are being generated.

**Acceptance Criteria:**

- Logs are being generated for all installed services.

### **Test Case 24: Ensure only required services are running**

**Description:**

Ensure only required services are running

**Pre-requisites:**

- Fully configured Linux with all services installed.

**Procedure:**

- Run the command “service --status-all” from CLI.
- Run NMAP Port Scan using command “nmap localhost”.

**Acceptance Criteria:**

- Only the required services should be running.

## 6.3 Test Cases Summary

The following table shows the result and present status of each application deployed in IXP server and its evaluation result if it has been successfully tested.

Test Case#	Test Result	Present status
Test Case 1	Successful	Completed
Test Case 2	Successful	Completed
Test Case 3	Successful	Completed
Test Case 4	Successful	Completed
Test Case 5	Successful	Completed
Test Case 6	Successful	Completed
Test Case 7	Successful	Completed
Test Case 8	Successful	Completed
Test Case 9	Successful	Completed
Test Case 10	Successful	Completed
Test Case 11	Successful	Completed
Test Case 12	Successful	Completed
Test Case 13	Successful	Completed
Test Case 14	Successful	Completed
Test Case 15	Successful	Completed
Test Case 16	Successful	Completed
Test Case 17	Successful	Completed
Test Case 18	Successful	Completed
Test Case 19	Successful	Completed
Test Case 20	Successful	Completed
Test Case 21	Successful	Completed
Test Case 22	Successful	Completed
Test Case 23	Successful	Completed
Test Case 24	Successful	Completed

Table 3 – Test Cases Summary



## 7. Conclusions

This thesis has presented an overview why it is important to keep the local traffic local and transfer the data locally where possible. The motivations and key challenges faced when working on a peering agreement were highlighted. The OSIX is built with a focus on developing countries but it can be deployed anywhere in the world. The IXP peering discussed and its limitations were identified.

The OSIX was built to overcome the technical barriers in developing countries for not having enough skills to build a complete system. The system was developed on a continually running environment bringing several constraints that had to be dealt with in order to pursue the development.

The IXP Fact finding survey was conducted and shared with selected IXP administrators worldwide and their responses were noted in developing the system.

During the testing we selected the test cases and we used them to identify any loopholes in the system, for any new features were identified and their resulting implications were shown in test cases.

### 7.1 Major Contributions

In summary this thesis has made the following contributions:

- *Quantitative evaluation of the need and benefits of peering.* We have done quantitative evaluation of traffic and bandwidth delays from both domestic and International servers and identified how much affective it is to peer with IXPs. We have also used APBDC model to prove that it is also more cost-effective for ISPs to peer with exchange points.
- *IXP Survey.* We published a fact-finding survey for the IXP administrators and stakeholders to share their experience and advise how a better IXP can help to reduce their shortcomings. Although we did not get the number of responses that we had hoped, nonetheless we were able to use responses and other surveys to inform our development of requirements for the solution.

- *Design and implementation of an IXP integrated solution (OSIX).* The solution has the following features:
  - *Implemented and integrated the OSIX Model for peering.* We have built the IXP which has all the components required to run the IXP successfully anywhere in the world with minimum efforts and technical knowledge.
  - *Tools for Management.* We have built a complete OSIX (in-the-box) solution, which has all the necessary components, needed to run an exchange point. All the components identified are discussed in detail in Chapter 4 and implemented in chapter 5.
  - *Used Virtual Machines (VM) to build OSIX.* We have built our complete system using virtual machines. This has the OSIX system highly portable and can now be re-distributed easily anywhere. The virtual image can be used in other freely available virtualisation software like, Virtual Box.
  - *Open Source Solution.* Our solution is fully open-sourced and available under GNU GPL Licence that avoids the licensing issues regarding free redistribution of our solutions.
- *Evaluation of the solution.* The solution's functionality has been evaluated using a laboratory test bed.

## 7.2 Future Work

The work on OSIX in this thesis demonstrates a complete working solution to run and manage the IXP. More work is required to make an ISP more resilient and reliable around the world depending on how much traffic its handling and the number of peers are subscribed to an Exchange Point.

- *RPKI in Exchange Points.* Resource Public Key Infrastructure (RPKI), also known as Resource Certification, is a specialized PKI framework designed to secure the Internet's routing infrastructure, specifically the BGP. RPKI provides a way to connect Internet number resource information (such as Autonomous System numbers and IP Addresses) to a trust party. The certificate structure mirrors the way in which Internet number resources are distributed. RPKI can be used by the legitimate holders of the resources to control the operation of Internet routing protocols to prevent route hijacking and other attacks.

- *MPLS in Exchange Points*. As part of future work, we will address the architecture model and the functionality of the Generalized Multi-Protocol Label Switching Exchange Point (GMPLS-XP), which is the equivalent of the Internet Exchange Point (IXP). Considering the importance of the role that the IXPs today have in supporting the global Internet operation we believe that GMPLS-XP is a necessary building block for the future multi-provider heterogeneous transport networks.

- *Types of Traffic Analysis*. In order to better understand what kind of services must be hosted on Exchange Points to provide a faster bandwidth, we need to perform an additional test to identify what are the requirements. This testing is also very important as it identifies what kind of IXP is needed in a particular region where traffic requirements may be different.

- *Real-time Testing of OSIX system*. In future we should also re-distribute OSIX system to Exchange Points in developing countries and run it parallel with their existing system and identify if there are any missing loopholes that can be fixed and prevented and it will also demonstrate how OSIX system handles traffic in production.

- *User Manual*. Even though the OSIX system is very easy to manage and it has all the components pre-installed, but having a user manual is always helpful. Due to time constraints, we couldn't complete a user manual, but as part of future work we should complete a detailed user manual for OSIX.



## 8. Appendices

### Appendix A – IXP Fact Finding Survey

Dale Carnegie  
Head of School of Engineering and Computer Science  
Cotton Building  
Victoria University of Wellington

16 March 2012

Dear HoS

I am a post-graduate research student at VUW and doing a Masters of Engineering with majors in Network Engineering with the topic of “Improving Tools for Development of Internet Exchange Points to Reduce Barriers to Internet Growth in Developing Countries”.

I intend to do a questionnaire and am seeking approval under section 4.7 of VUW HEC policy. This meets the criteria for approval by HoS (with copy send to HEC convener) for the following reasons (per section 4.7(b)):

- (i) The questionnaire is totally anonymous and no personally identifying information is collected. Additionally, the questionnaire is provided via a server so we do not observe the collection of the information and the server does not track IP addresses or other identifying information.
- (ii) Not contain questions on sensitive topics (e.g. sexual practices, drug taking, illegal activities). The questions are restricted to information about peering architectures and are factual in nature.;
- (iii) Questions are designed around the research goals of improving tools. The questions are aimed at discovering requirements from the community for tools for building better Internet Exchange Points
- (iv) This is a student project supervised by Ian Welch and Andy Linton.
- (v) The questionnaire states the purpose of the questionnaire, the use to which the results will be put, the disposal of the questionnaire forms, and the fact that the questionnaire is anonymous.

The questionnaire states this (see attached).

yours sincerely,

Aun Iqbal  
(Student ID: 300208565)

APPROVED BY:

DATE:

Please return this letter to Ian Welch.

## IXP Survey

You are being invited to participate in IXP Fact Finding survey. This survey is part of thesis from Masters of Engineering at VUW.

There are no known risks if you decide to participate in this research study. There are no costs to you for participating > in the study. The information you provide will be anonymous and will only be used for statistical purpose for the thesis. The questionnaire will take about 5-10 minutes to complete.

The information collected may not benefit you directly, but the information learned in this study should provide more general benefits. (Here you could say “In particular, the results will be used to guide the development of a peering toolkit (IXPsolutions) for making it easier to deploy a peering solution.”)

This survey is anonymous. Do not write your name on the survey. We do not collect IP addresses. No one will be able to identify you or your answers, and no one will know whether or not you participated in the study. Individuals from Victoria university of Wellington may inspect these records. Should the data be published, no individual information will be disclosed.

Your participation in this study is voluntary. By completing this survey you are voluntarily agreeing to participate.

The information from the survey will be kept on a secure server and deleted two years after the end of the study.

To the best of our knowledge, no survey of systems administrators has ever been conducted in an academic setting, so the information gathered can help focus future research efforts. We expect to release generalised results within one month at this address:

<http://ecs.victoria.ac.nz/Main/GradAunIqbal>

If you have any questions about the study, please contact Aun Iqbal at [aun.iqbal@ecs.vuw.ac.nz](mailto:aun.iqbal@ecs.vuw.ac.nz).

Please tick this box to show that you have read and understood the following notice:  (Tick here)

“I understand that the generalised results from this survey will be used to advance academic research, and that they may be published to disseminate the findings. No personal information has been or will be collected, and therefore cannot and will not be released. The form data is stored on a private server and protected by passwords and will be securely disposed of a year after the completion of the thesis.””

### **Your background**

1. What is your role in dealing with Internet Exchange Point (IXP)?
  - a. Network Administrator
  - b. Sales & Marketing
  - c. Manager
  - d. Other
  
2. Which of the following RIR group does your IXP operate in?
  - a. African Network Information Centre (AfriNIC)
  - b. American Registry for Internet Numbers (ARIN)
  - c. Asia-Pacific Network Information Centre (APNIC)
  - d. Latin America and Caribbean Network Information Centre (LACNIC)
  - e. Réseaux IP Européens Network Coordination Centre (RIPE)
  
3. How many years have you been using or operating IXPs?

### **Your IXP**

4. When was your IXP established?
  
5. How many subscribers do you have in your IXP?
  
6. Which of the following 'peering techniques' do you prefer to deploy?
  - a. Multilateral
  - b. Bilateral
  - c. Both
  
7. What volume (in GB) of traffic is your Exchange Point is handling per day?
  
8. Is your IXP Layer 2 or Layer 3?

### **IXP Management**

9. Do you have a dedicated management server at IXP?
  - a. Yes
  - b. No
  
10. Briefly explain how do you currently manage your exchange point?
  
11. Do you have one or more 'Route Servers' or 'Route Reflectors'?
  
12. What tools do you use to generate bandwidth and traffic utilization at the IXP?
  
13. Do you have any tools to take backups of the router and switch configurations at Exchange Point?
  
14. Are you satisfied with your existing infrastructure in Exchange Point?
  
15. Would you like to have a management server at Exchange Point?
  
16. Do you have a DNS Root Server at your IXP?  
Yes  
No
  
17. What user-interface is more convenient for your needs?

- a. CLI
- b. Web (https)
- c. GUI Desktop

### **Challenges**

- 18. What is your biggest concern when deploying a new feature?
- 19. Do you have a dedicated person(s) to manage the IXP?
- 20. Does your IXP support IPv6?

### **Future Development**

- 21. What would you do differently if you decided to redesign the Exchange Point?
- 22. Which of the following features would add value to the Exchange Point?
  - a. Subscriber Database
  - b. Backup of configuration
  - c. Route Server
  - d. Auto-generation of 'filter lists'
  - e. Bandwidth Utilization Graphs
  - f. Looking Glass
  - g. Root Server
  - h. NTP Server
  - i. Others
- 23. Any suggestions?

## Appendix B – IXP Survey Email

The following email was sent out to all participants requesting them to participate in survey.

Dear Sir/Madam,

I am a student at Victoria University of Wellington, New Zealand and currently studying for Masters of Engineering (ME) in Network Engineering.

As part of my ME Thesis on “Improving Tools for Development of Internet Exchange Points to Reduce Barriers to Internet Growth in Developing Countries”, I am required to do a survey to investigate how integrated tools could be used to improve the deployment and management of Internet Exchange Points.

I have retrieved your contact details via your Exchange Point's Website and I would like to request you to participate in this short survey.

- The survey is focused on Exchange Point Administrators/Operators.
- It is an on-line survey and should take between 5 and 10 minutes to complete.
- Responses are anonymous and we are not requesting or monitoring any identifying information.

Please find the link to the survey at: <http://ixpsolutions.org/survey>

Please note that this survey is purely for academic purpose and we are not collecting any personal information which identifies you.

The results will not be used for any marketing or promotional purposes.

I appreciate your help with this project.

Kind regards, Aun

Email: [aun.iqbal@ecs.vuw.ac.nz](mailto:aun.iqbal@ecs.vuw.ac.nz)



## Appendix C – IXP Survey Results

### Results

Number of records in this query: 4

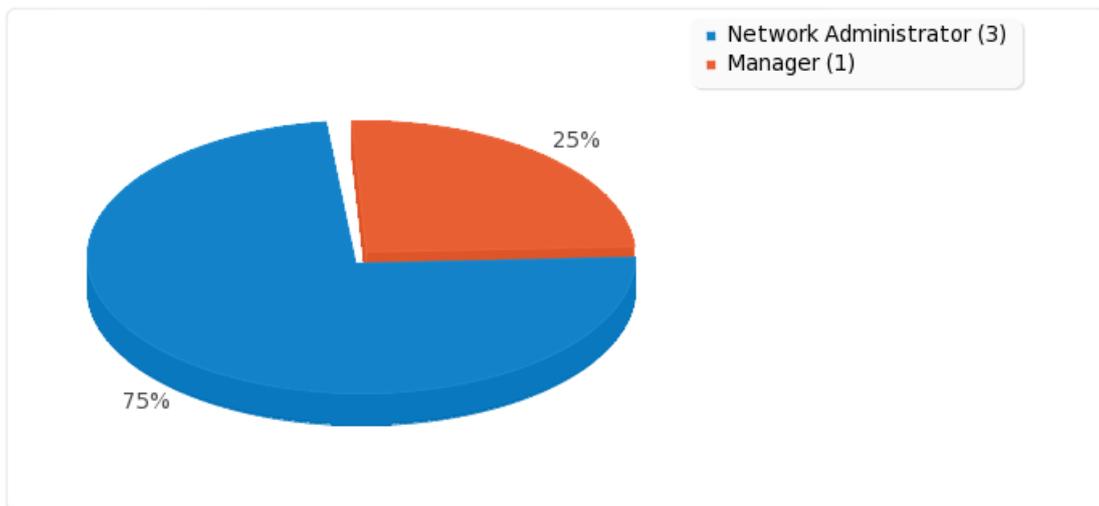
Total records in survey: 4

Percentage of total: 100.00%

#### Field summary for 1

##### What is your role in dealing with Exchange Points?

Answer	Count	Percentage
Network Administrator (A1)	3	75.00%
Sales & Marketing (A2)	0	0.00%
Manager (A3)	1	25.00%
Other	0	0.00%
No answer	0	0.00%



#### Field summary for 2

##### Which of the following RIR group does your IXP operate in?

Answer	Count	Percentage
African Network Information Centre (AfriNIC) (A1)	1	25.00%
American Registry for Internet Numbers (ARIN) (A2)	0	0.00%
Asia-Pacific Network Information Centre (APNIC) (A3)	3	75.00%
Latin America and Caribbean Network Information Centre (LACNIC) (A4)	0	0.00%
Réseaux IP Européens Network Coordination Centre (RIPE) (A5)	0	0.00%
No answer	0	0.00%



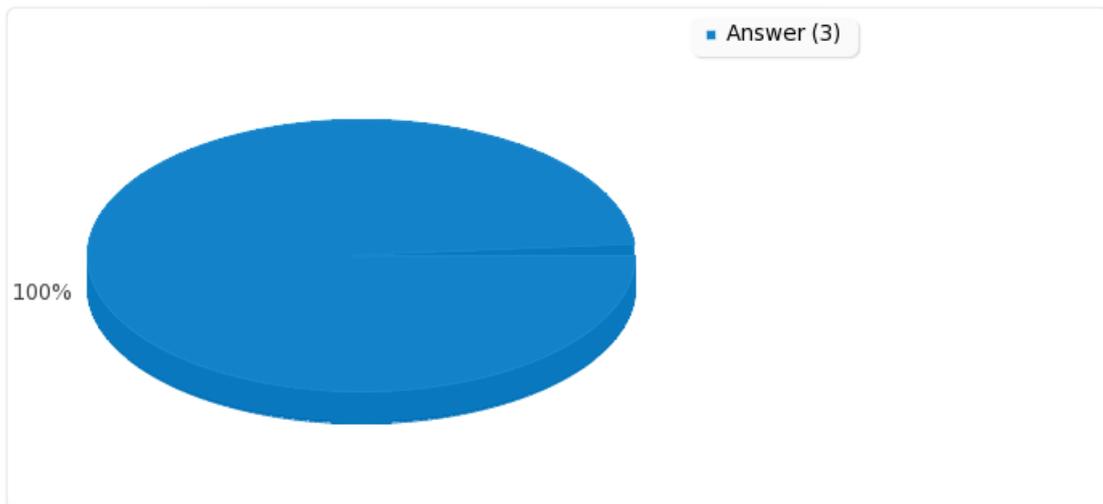


Answers:

9	2003
11	1996
12	2003

**Field summary for 5**  
**How many subscribers do you have in your IXP?**

Answer	3	100.00%
No answer	0	0.00%



Answers:

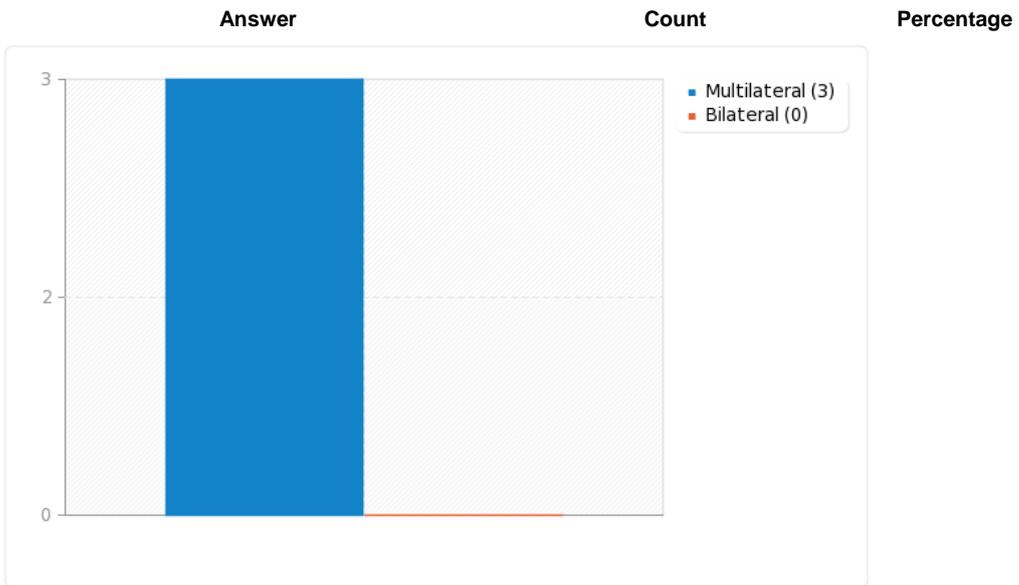
9	22
11	approx 100
12	101

**Field summary for 6**  
**Which of the following 'peering techniques' do you prefer to deploy?**

Answer	Count	Percentage
Multilateral (A6)	3	75.00%
Bilateral (A7)	0	0.00%

Field summary for 6

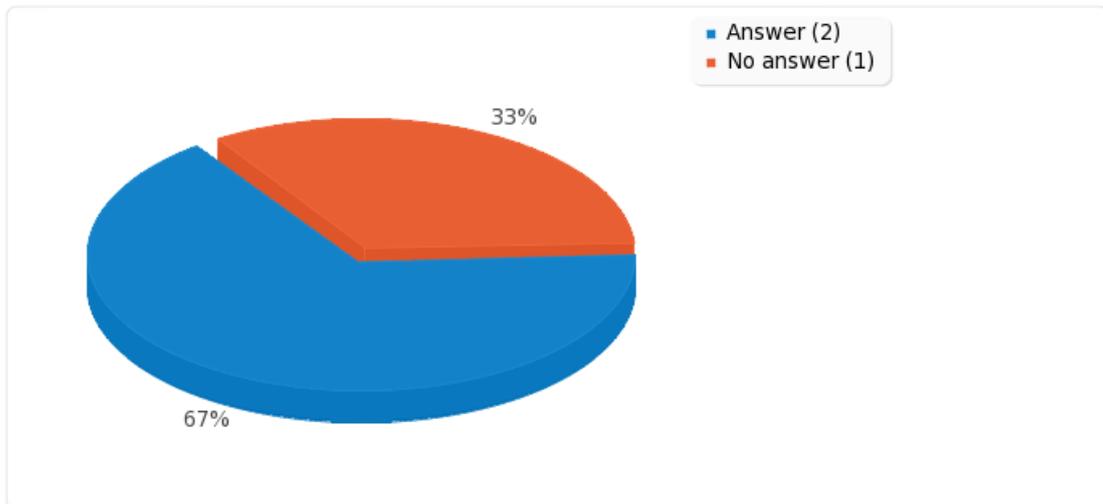
Which of the following 'peering techniques' do you prefer to deploy?



Field summary for 7

What volume (in GB) of traffic is your Exchange Point is handling per day?

Answer	Count	Percentage
Answer	2	66.67%
No answer	1	33.33%



Answers:

9	120
11	2.5Gbps peak, no idea how many actual bytes

Field summary for 8

Your IXP is:

Answer	Count	Percentage
Layer 2 (SQ001)	3	75.00%
Layer 3 (SQ002)	2	50.00%

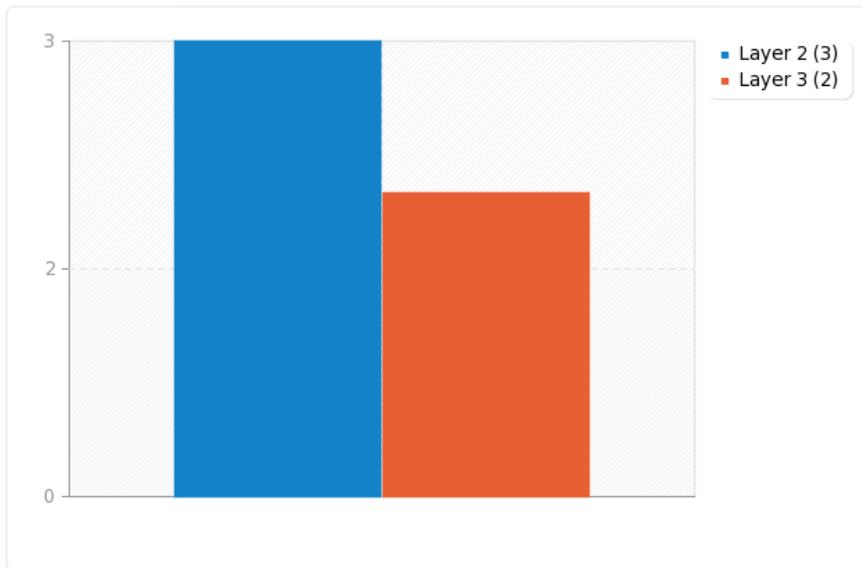
**Field summary for 8**

**Your IXP is:**

**Answer**

**Count**

**Percentage**



**Field summary for 9**

**Do you have a dedicated management server at your Exchange?**

**Answer**

**Count**

**Percentage**

Yes (A1)

3

100.00%

No (A2)

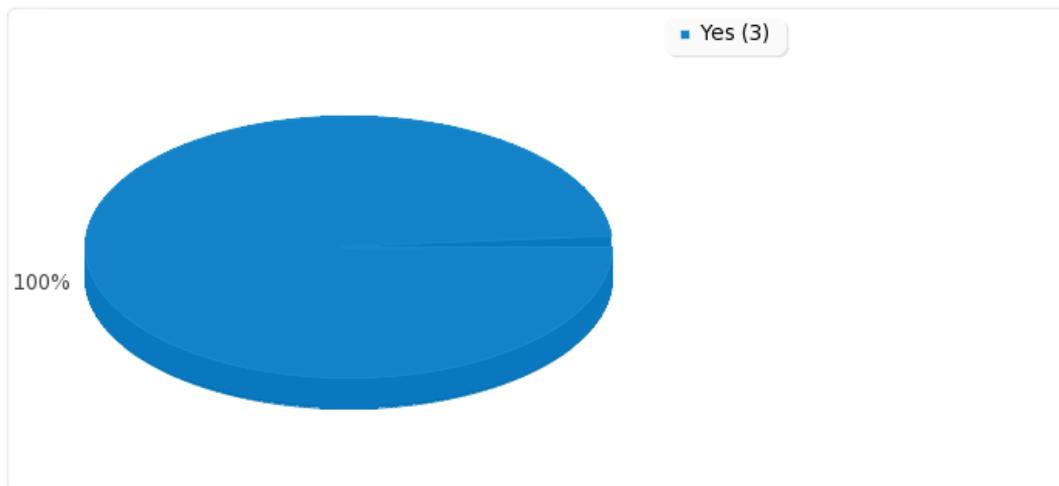
0

0.00%

No answer

0

0.00%



**Field summary for 10**

**Briefly explain how do you currently manage your exchange point?**

**Answer**

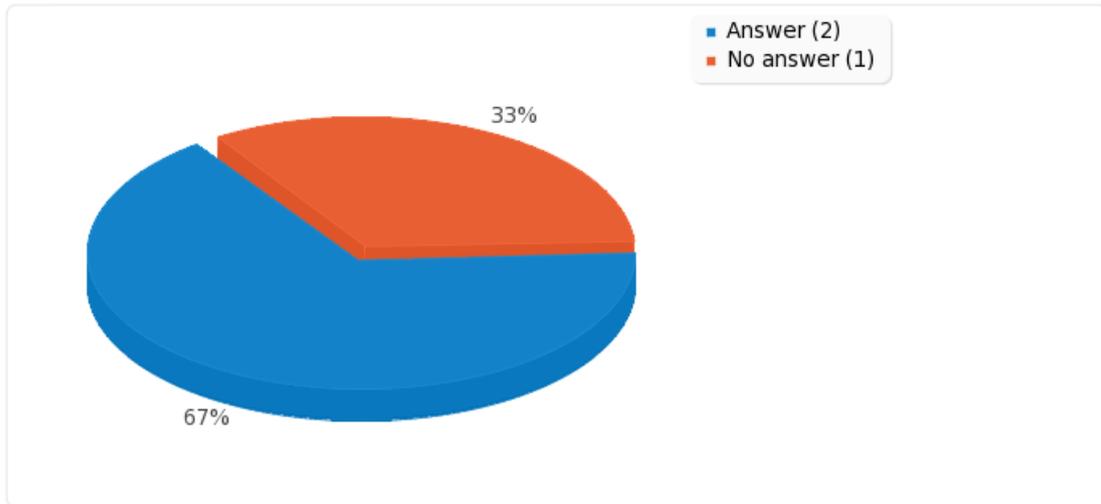
**Count**

**Percentage**

No answer

1

33.33%



**Answers:**

- 9 | 2x cisco eBGP speakers as route server, with manual incoming prefix filters manual config, ssh, vi etc
- 11 | automated scripts that build quagga configs from RPSL and static data.

**Field summary for 11**

**Do you have one or more 'Route Servers' or 'Route Reflectors'?**

Answer	3	100.00%
No answer	0	0.00%



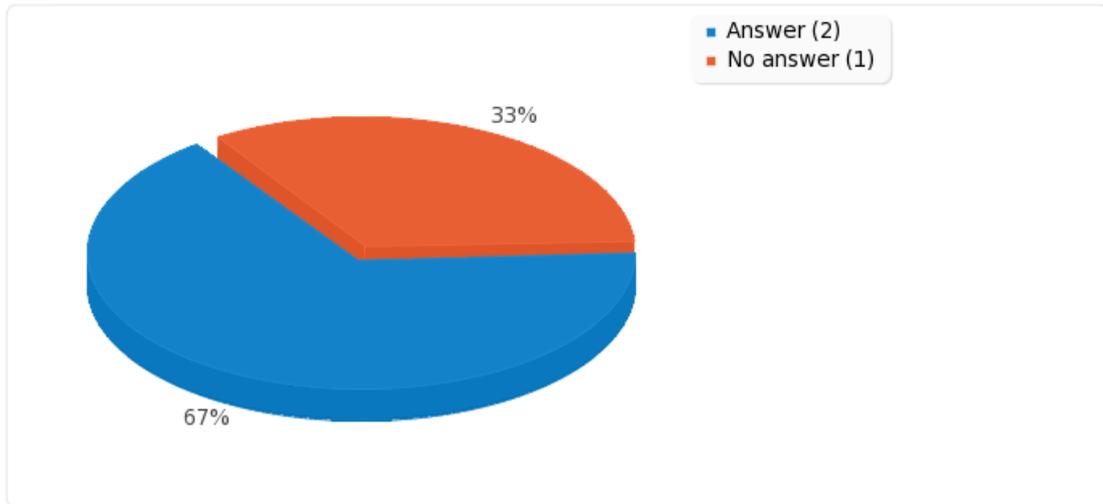
**Answers:**

- 9 | yes, 2 route servers
- 11 | Yes
- 12 | yes

**Field summary for 12**

**What tools do you use to generate bandwidth and traffic utilization at Exchange Point?**

Answer	2	66.67%
No answer	1	33.33%



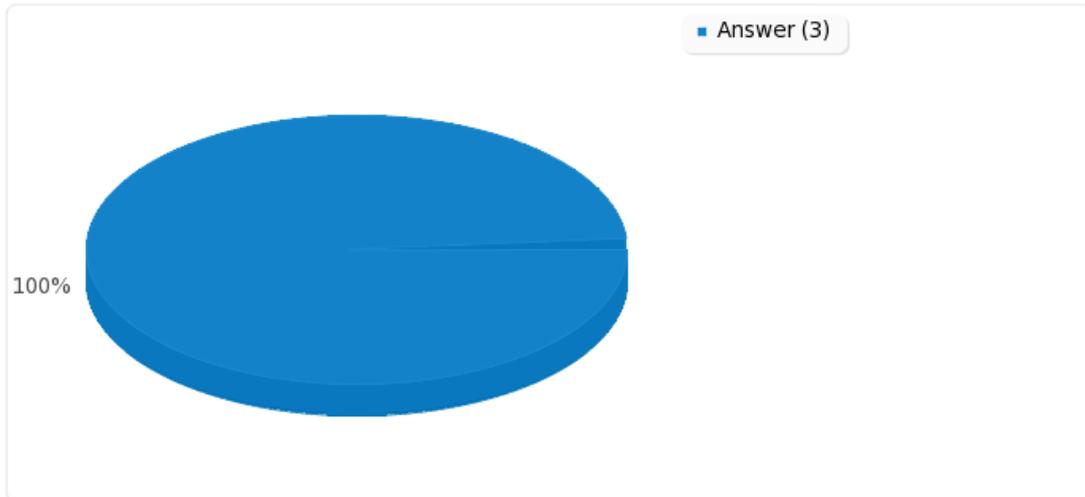
**Answers:**

9	mrtg + rrdtool
11	SNMP/Collectd/Munin

**Field summary for 13**

**Do you have any tools to take backups of the router and switch configurations at Exchange Point?**

Answer	3	100.00%
No answer	0	0.00%



**Answers:**

9	tftp / rancid
11	Yes
12	yes

**Field summary for 14**

**Are you satisfied with your existing infrastructure in Exchange Point?**

Answer	3	100.00%
No answer	0	0.00%



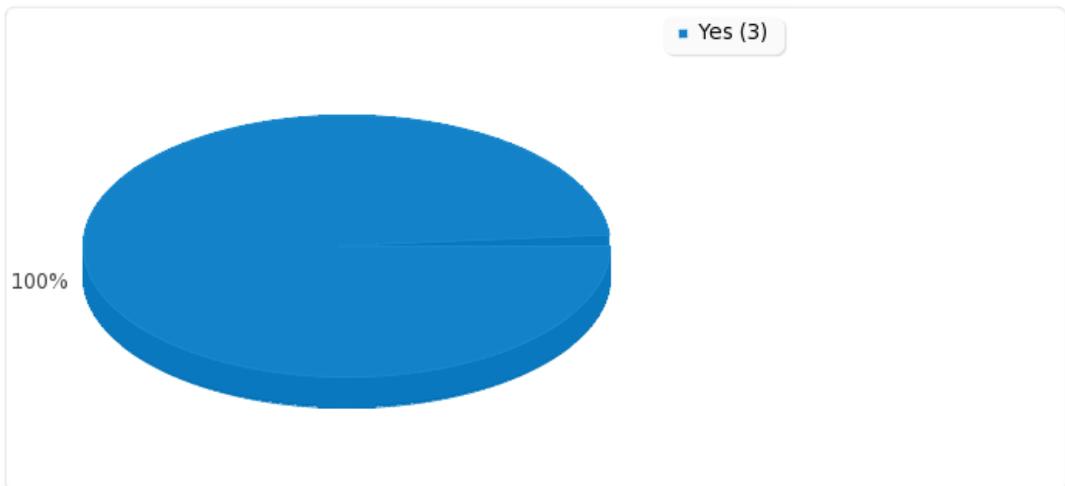
Answers:

9	almost
11	Yes
12	no

Field summary for 15

Would you like to have a management server at Exchange Point?

Answer	Count	Percentage
Yes (Y)	3	100.00%
No (N)	0	0.00%
No answer	0	0.00%



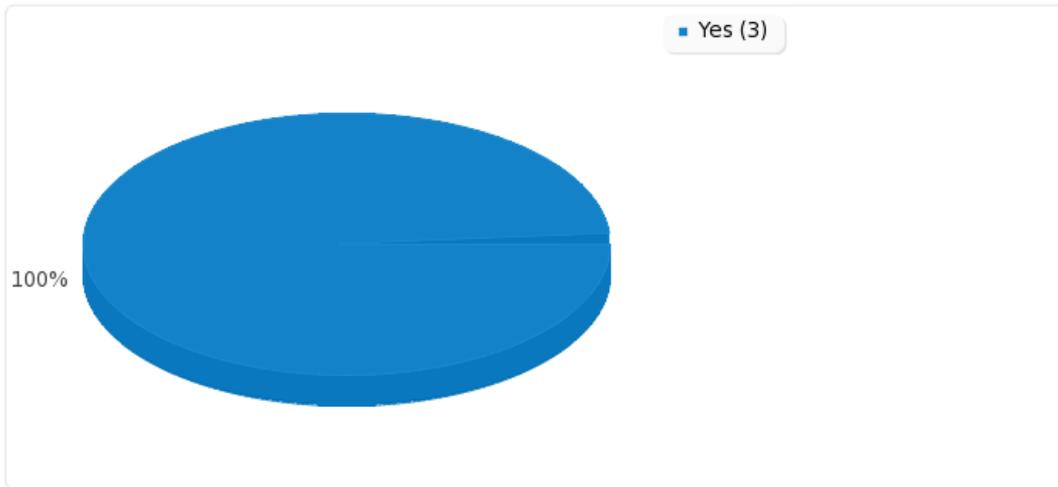
Field summary for 16

Do you have a DNS Root Server at your IXP?

Answer	Count	Percentage
Yes (Y)	3	100.00%
No (N)	0	0.00%
No answer	0	0.00%

**Field summary for 16**  
**Do you have a DNS Root Server at your IXP?**

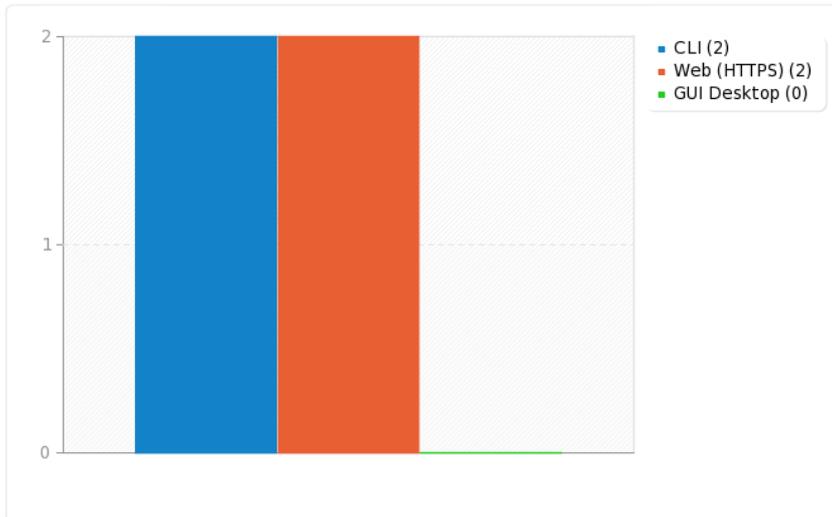
Answer	Count	Percentage
--------	-------	------------



**Field summary for 17**

**What user-interface is more convenient for your needs?**

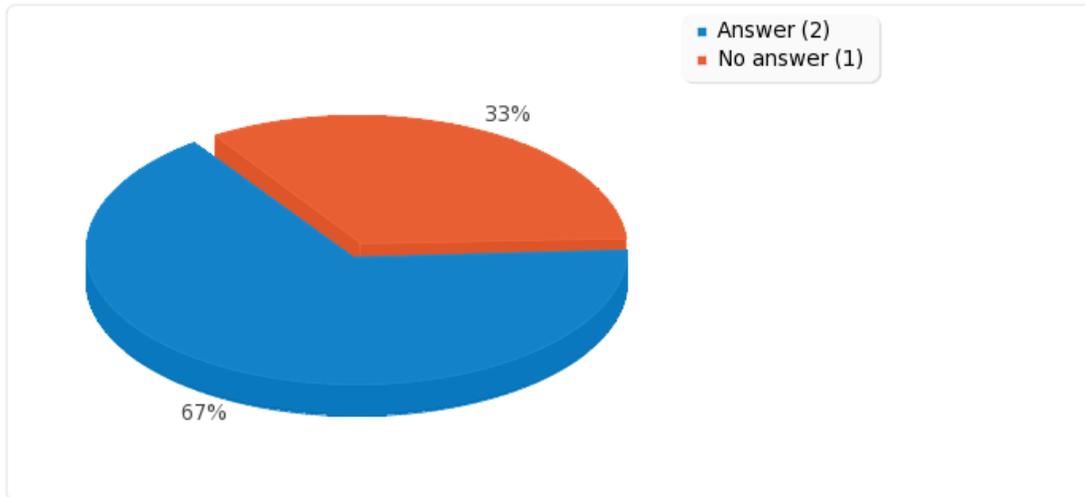
Answer	Count	Percentage
CLI (SQ001)	2	50.00%
Web (HTTPS) (SQ002)	2	50.00%
GUI Desktop (SQ003)	0	0.00%



**Field summary for 18**

**What is your biggest concern when deploying a new feature?**

Answer	2	66.67%
No answer	1	33.33%



**Answers:**

11	Causing instability.
12	Low downtime.

**Field summary for 19**  
**Do you have a dedicated person(s) to manage the IXP?**

Answer	3	100.00%
No answer	0	0.00%



**Answers:**

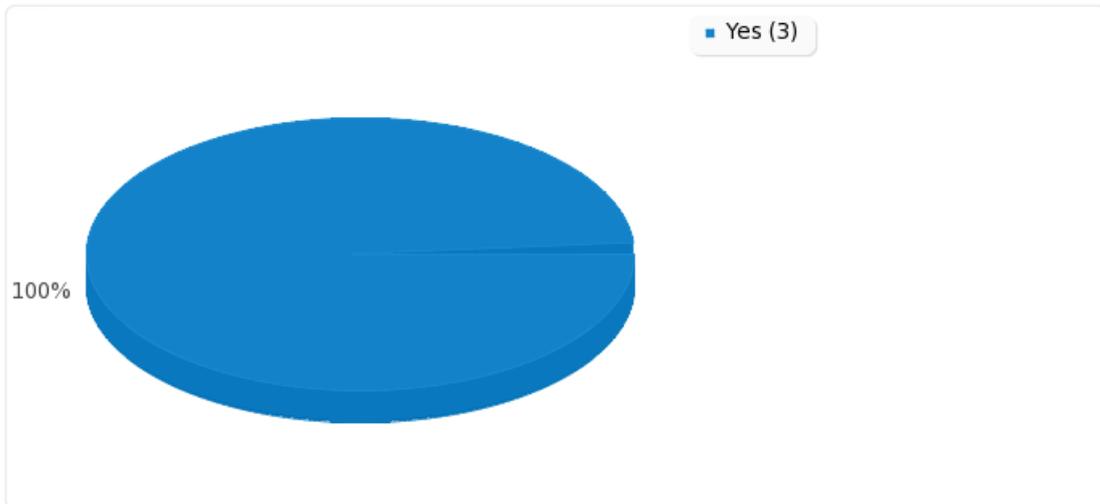
9	yes, voluntary part-time
11	No
12	yes

**Field summary for 20**  
**Does your IXP support IPv6?**

Answer	Count	Percentage
Yes (Y)	3	100.00%
No (N)	0	0.00%
No answer	0	0.00%

**Field summary for 20**  
**Does your IXP support IPv6?**

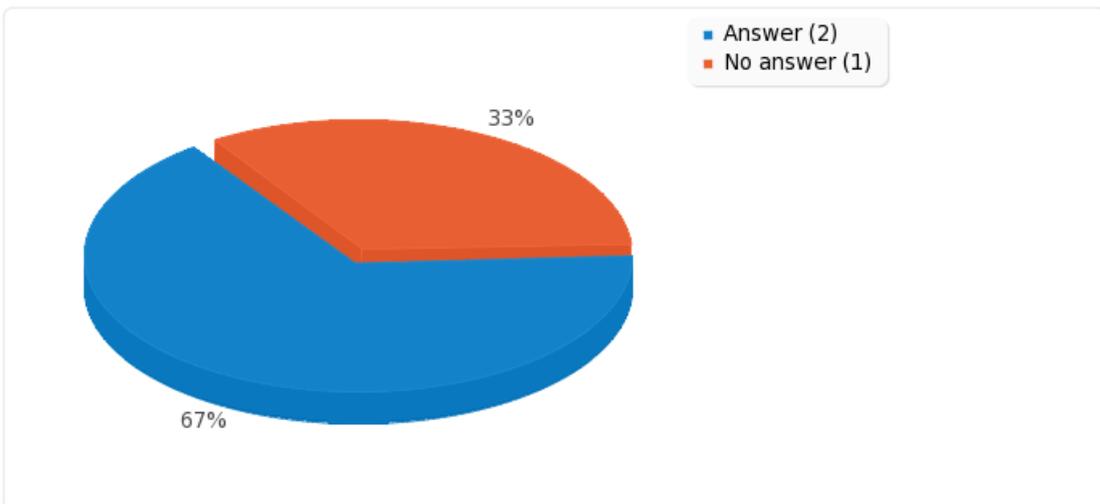
**Answer** **Count** **Percentage**



**Field summary for 21**

**What would you do differently if you decided to redesign the Exchange Point?**

Answer	Count	Percentage
Answer	2	66.67%
No answer	1	33.33%



**Answers:**

- 11 | More powerful servers to allow faster convergence. More automation.
- 12 | Allowing all companies who have AS Numbers. May be allowing Bilateral peering.

**Field summary for 22**

**Which of the following features would add value to your Exchange Point?**

Answer	Count	Percentage
Subscriber Database (SQ001)	0	0.00%
Backup of Configuration (SQ002)	1	25.00%
Route Server (SQ003)	1	25.00%
Auto-generation of 'filter-lists' (SQ004)	2	50.00%
Bandwidth Utilization Graphs (SQ005)	1	25.00%
Looking Glass (SQ006)	1	25.00%
DNS Root Server (SQ007)	1	25.00%
NTP Server (SQ008)	1	25.00%

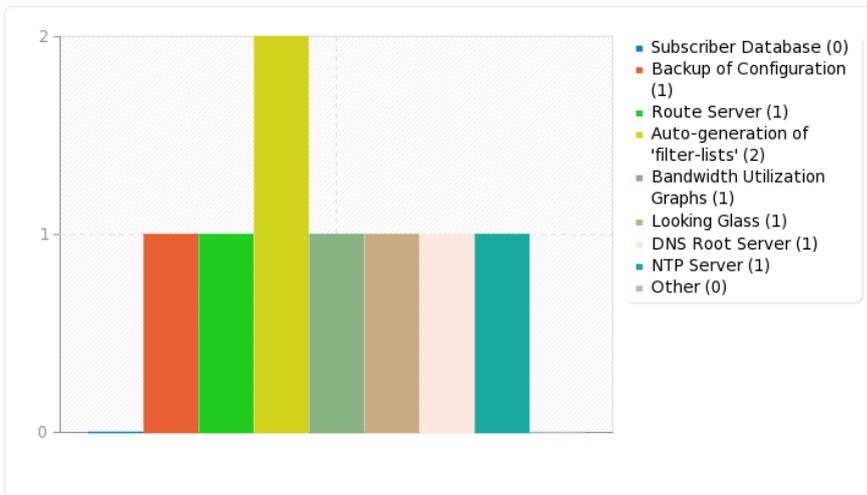
**Field summary for 22**

**Which of the following features would add value to your Exchange Point?**

**Answer**  
Other

**Count**  
0

**Percentage**  
0.00%



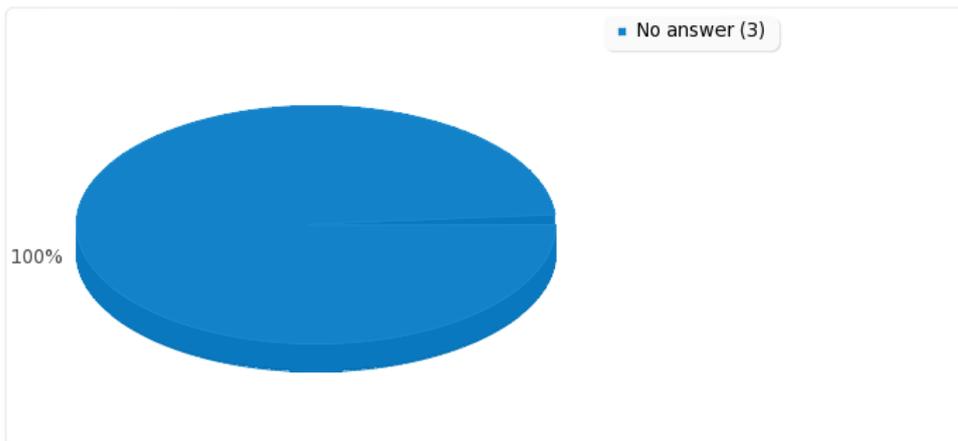
**Field summary for 23**

**Any suggestion?**

**Answer**  
No answer

**Count**  
0  
3

**Percentage**  
0.00%  
100.00%



## Appendix D – GPL License

This license is included here as a legal requirement, given that GPL-licensed code has been included in this document.

### GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

#### Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires

that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

## **TERMS AND CONDITIONS**

### **0. Definitions.**

“This License” refers to version 3 of the GNU General Public License.

“Copyright” also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

“The Program” refers to any copyrightable work licensed under this License. Each licensee is addressed as “you”. “Licensees” and “recipients” may be individuals or organizations.

To “modify” a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a “modified version” of the earlier work or a work “based on” the earlier work.

A “covered work” means either the unmodified Program or a work based on the Program.

To “propagate” a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To “convey” a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays “Appropriate Legal Notices” to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

## **1. Source Code.**

The “source code” for a work means the preferred form of the work for making modifications to it. “Object code” means any non-source form of a work.

A “Standard Interface” means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The “System Libraries” of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A “Major Component”, in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The “Corresponding Source” for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work’s System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

## **2. Basic Permissions.**

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output,

given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

### **3. Protecting Users' Legal Rights From Anti-Circumvention Law.**

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

### **4. Conveying Verbatim Copies.**

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

### **5. Conveying Modified Source Versions.**

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to “keep intact all notices”.
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an “aggregate” if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

## **6. Conveying Non-Source Forms.**

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object

code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

- e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

## **7. Additional Terms.**

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

## **8. Termination.**

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

## **9. Acceptance Not Required for Having Copies.**

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

## **10. Automatic Licensing of Downstream Recipients.**

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An “entity transaction” is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

## **11. Patents.**

A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's “contributor version”.

A contributor's “essential patent claims” are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, “control” includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a “patent license” is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To “grant” such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the

covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

## **12. No Surrender of Others' Freedom.**

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

## **13. Use with the GNU Affero General Public License.**

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

## **14. Revised Versions of this License.**

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

## **15. Disclaimer of Warranty.**

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

## **16. Limitation of Liability.**

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

## 17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it
does.>
```

```
Copyright (C) <year> <name of author>
```

```
This program is free software: you can redistribute it and/or
modify
it under the terms of the GNU General Public License as published
by
the Free Software Foundation, either version 3 of the License, or
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program. If not, see
<http://www.gnu.org/licenses/>.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author>
This program comes with ABSOLUTELY NO WARRANTY; for details type
`show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands ``show w'` and ``show c'` should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an “about box”.

You should also get your employer (if you work as a programmer) or school, if any, to sign a “copyright disclaimer” for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <http://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <http://www.gnu.org/philosophy/why-not-lgpl.html>.



## Appendix E – VUW Human Ethics Policy

This is an excerpt from Section 4.7 of the VUW HEP [94], under which the survey was approved without a formal Ethics Committee hearing.

“(a) Research in which the subject's participation is restricted to the completion of a written questionnaire in a manner not requiring the disclosure of the subject's identity, and which meets the criteria for questionnaires in section 4.7(b), may be approved in writing by the Head of the School...

(b) The questionnaire must:

- (i) Be totally anonymous (responses should be returned anonymously and there should be no coding or other means of identifying respondents from the response);
- (ii) Not contain questions on sensitive topics (e.g. sexual practices, drug taking, illegal activities);
- (iii) Be designed to meet the research goals set;
- (iv) In the case of student projects, be subject to appropriate supervision;
- (v) Normally state the purpose of the questionnaire, the use to which the results will be put, the disposal of the questionnaire forms, and the fact that the questionnaire is anonymous.”



## **Appendix F – OSIX Terms & Conditions**

### **Definitions**

"Subscriber(s)" means you and other users with networks connected to OSIX.

"OSIX" stands for Open Source Internet Exchange and is the general name to describe our Exchange Point. In this document OSIX refers to the exchange point you are planning to connect your network to.

"We" or "us" means OSIX.

"You" means the subscriber under this agreement and "your" has a corresponding meaning.

### **Agreement**

This agreement covers your connection to a particular OSIX. This is not to be confused with the set of bilateral OSIX peering agreements which are set up between your network and other networks on that OSIX.

This agreement begins when you correctly complete the on-line OSIX application form. The agreement will last for a period of one year, after which it will be extended automatically by periods of one year unless terminated by either party (you or us) giving three months notice in writing prior to the expiry of a one-year period.

### **Object of the agreement**

The general purpose of each OSIX is to promote the use of data communication and to provide a framework for cooperation in a particular geographic area that will allow the users of the participating networks to communicate with each other in the easiest way possible. The cooperation will, by implication, follow the prevailing rules and practices of the worldwide IP-network, the Internet, which each of the participating networks are a part of.

### **Connecting Your Network to an OSIX**

#### **General provisions**

To connect to an OSIX you require an IP router. The router will be connected to your network and to the particular OSIX which provides a neutral Public LAN network.

The following principles apply to the neutral Public LAN network:

- You are free to determine your policy towards the other Subscribers connected to an OSIX.
- You are free to set up bilateral network connections independently across an OSIX because it is a neutral Public LAN infrastructure.
- We will operate route servers on OSIX which you can choose to peer with. Peering connections are implemented using the BGP4 protocol. Any change to this will be made in consultation with all Subscribers connected to the OSIX.
- IGP protocols (RIP, IGRP, ISIS, OSPF, VRRP etc.) must not be used or transmitted on the neutral Public LAN network.

- Your connection to an OSIX must be implemented so that you are able to control your own internal traffic. Consequently, traffic being addressed within your network will remain on your network.

You will pay all costs related to connecting your network to the OSIX.

Correspondingly, we will cooperate with you in respect of the function of the peering point and also to provide the necessary assistance to you on normal terms.

### **Exclusion and disruption**

You must ensure that your usage of an OSIX is not detrimental to the Public LAN Network or to the usage of it by other Subscribers.

Under certain circumstances, we or you may disrupt the connection between the neutral Public LAN network and your network. These circumstances are:

- if either we or you deliberately misuse the networks in any way, or
- if either we or you cause an unreasonable load to the networks of individual Subscribers or in any other way cause a threat to the function and applicability of the network.

Each party (Us or You) should try to contact the other before such a disruption is made. In matters of urgency, however, disruption may be made immediately.

### **Liability**

No party (you or us) to this agreement can be held liable for indirect or consequential damage.

All liability rests with you with respect to the correctness and suitability of any and all information transmitted and any and all results thereby obtained.

### **Updating of equipment**

Each party (you or us) are entitled to install new versions of equipment and software.

### **Other provisions**

#### **Transfer**

No party (you or us) to this agreement may transfer its part of this agreement to any third party without the approval of the other party.

#### **Legal disputes**

Any dispute and controversy arising out of or in connection with this Agreement will be referred to arbitration according to New Zealand law.

## Appendix G – Test Cases Results

The following are the results for each test case discussed in Chapter 6.

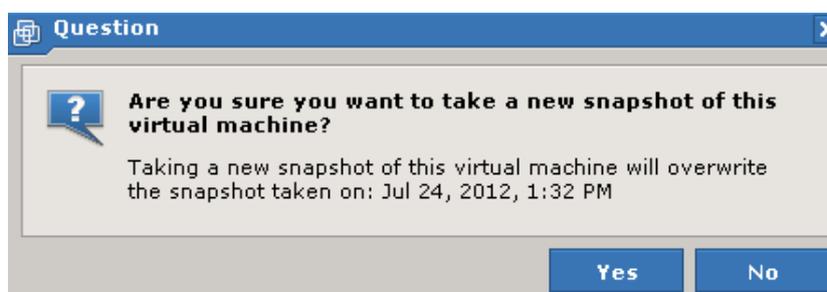
### Test Case 1:

#### IDENTIFICATION:

TEST ID: 01  
TITLE: VMware Image Portability Test  
PLATFORM: VMware 2.02 Server  
COMPONENT: Virtual Image  
DATE OF TEST: 24.07.2012

S: Success

F: Failure



Task	Target	Status	Triggered At	Triggered by	Completed At
Power On Virtual Machine	ixp1	Success	07/24/12 3:02:42 PM	iqbalaun	07/24/12 3:02:48 PM
Create Virtual Machine Snapshot	ixp1	Success	07/24/12 2:58:57 PM	iqbalaun	07/24/12 2:59:02 PM

#### Comments:

Test completed successfully.

### Test Case 2:

#### IDENTIFICATION:

TEST ID: 02  
TITLE: Warm Start of OSIX Servers  
PLATFORM: Ubuntu Server 11.04  
COMPONENT: Linux Server  
DATE OF TEST: 24.07.2012

S: Success

F: Failure

```
barretts: [~] % ssh 130.195.10.91
iqbalaun@130.195.10.91's password:
welcome to Ubuntu 11.04 (GNU/Linux 2.6.38-15-generic-pae i686)

* Documentation:  https://help.ubuntu.com/

system information disabled due to load higher than 1.0

Last login: Sat Jul 21 03:22:12 2012 from m167.dsrg.ecs.vuw.ac.nz
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$ uptime
03:01:01 up 2 min,  1 user,  load average: 1.00, 0.67, 0.27
iqbalaun@ixp1:~$
```

Comments: Test completed successfully.

**Test Case 3:**

IDENTIFICATION:	
TEST ID:	03
TITLE:	Cold Start of OSIX Servers
PLATFORM:	Ubuntu 11.04
COMPONENT:	Linux Server
DATE OF TEST:	24.07.2012

- S: Success
- F: Failure

```
Last login: Fri Jul 20 15:34:06 2012 from m167.dsrg.ecs.vuw.ac.nz
iqbalaun@ixp2:~$
iqbalaun@ixp2:~$ uptime
15:29:10 up 5 min,  2 users,  load average: 0.48, 1.05, 0.57
iqbalaun@ixp2:~$
```

Comments: Test completed successfully.

**Test Case 4:**

IDENTIFICATION:	
TEST ID:	04
TITLE:	SSH Authentication to OSIX Servers
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Linux Server
DATE OF TEST:	24.07.2012

- S: Success
- F: Failure

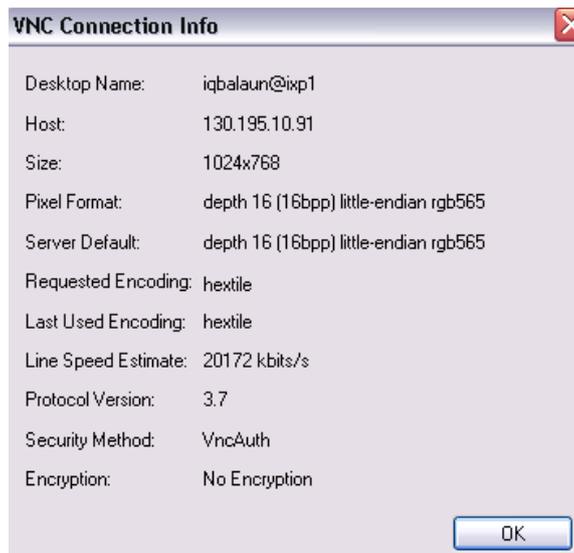
```
barretts: [~] % telnet 130.195.10.91
Trying 130.195.10.91...
telnet: Unable to connect to remote host: Connection refused
barretts: [~] %
barretts: [~] % ssh 130.195.10.91
iqbalaun@130.195.10.91's password:
welcome to Ubuntu 11.04 (GNU/Linux 2.6.38-15-generic-pae i686)
```

Comments: Test completed successfully.

**Test Case 5:**

<b>IDENTIFICATION:</b>	
TEST ID:	05
TITLE:	Remote Login to OSIX Servers Desktop
PLATFORM:	Ubuntu 11.04 Servers
COMPONENT:	Linux Servers
DATE OF TEST:	24.07.2012

- S: Success
- F: Failure



Comments: Test completed successfully.

### Test Case 6:

**IDENTIFICATION:**

TEST ID:	06
TITLE:	Web Authentication to OSIX Servers
PLATFORM:	Ubuntu 11.04 Servers
COMPONENT:	Linux Servers
DATE OF TEST:	24.07.2012

S: Success  
 F: Failure



System hostname	ixp1.dsrg.ecs.vuw.ac.nz (192.168.40.91)
Operating system	Ubuntu Linux 11.04
Webmin version	1.590
Time on system	Wed Jul 25 10:26:49 2012
Kernel and CPU	Linux 2.6.38-15-generic-pae on i686
Processor information	Intel(R) Pentium(R) D CPU 2.80GHz, 1 cores
System uptime	7 hours, 28 minutes
Running processes	148
CPU load averages	0.00 (1 min) 0.01 (5 mins) 0.05 (15 mins)
CPU usage	0% user, 0% kernel, 0% IO, 100% idle
Real memory	495.81 MB total, 190.04 MB used
Virtual memory	508 MB total, 1.57 MB used
Local disk space	48.68 GB total, 6.52 GB used
Package updates	All installed packages are up to date

**Comments:**

Test completed successfully.

### Test Case 7:

**IDENTIFICATION:**

TEST ID:	07
TITLE:	Telnet Authentication to OSIX Route Servers
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Quagga
DATE OF TEST:	24.07.2012

S: Success  
 F: Failure

```

iqbalaun@ixp1:~$ telnet localhost 2605
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.

Hello, this is Quagga (version 0.99.20.1).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

User Access Verification

Password:
ixp1> en
Password:
ixp1#
ixp1# sh ver
Quagga 0.99.20.1 ().
Copyright 1996-2005 kunihiro Ishiguro, et al.
ixp1#

```

Comments:

Test completed successfully.

### Test Case 8:

#### IDENTIFICATION:

TEST ID: 08  
TITLE: BGP Peering Establishment on OSIX Servers  
PLATFORM: Ubuntu 11.04 Server  
COMPONENT: Quagga  
DATE OF TEST: 24.07.2012

- S: Success  
 F: Failure

```

ixp1# sh ip bgp summary
BGP router identifier 192.168.40.91, local AS number 7675
RIB entries 0, using 0 bytes of memory
Peers 6, using 15 KiB of memory

Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ  OutQ  Up/Down  State/PfxRcd
192.168.40.110 4   110    468    468       0     0     0 07:45:49      0
192.168.40.120 4   120   1022    935       0     0     0 07:45:46      0
192.168.40.130 4   130   1016    935       0     0     0 07:45:50      0
fd9f:3a9:46ec:6917::110
  4   110    470    470       0     0     0 07:45:48      0
fd9f:3a9:46ec:6917::120
  4   120   1024    937       0     0     0 07:45:47      0
fd9f:3a9:46ec:6917::130
  4   130   1016    937       0     0     0 07:45:50      0
Total number of neighbors 6
ixp1#

```

Comments:

Test completed successfully.

### Test Case 9:

**IDENTIFICATION:**

TEST ID:	09
TITLE:	Telnet Authentication to IRRd Daemon
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	IRRD
DATE OF TEST:	24.07.2012

S: Success

F: Failure

```
iqbalaun@ixp1:~$ telnet localhost irrd
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
IRRD version 2.3.10 [17Nov2010]

User Access Verification

ixp1 Password: *****
ixp1 IRRd#
ixp1 IRRd# sh ver
2.3.10 [17Nov2010]
  ixp1
  Linux 2.6.38-15-generic-pae #61-Ubuntu SMP Tue Jun 12 19:32:42 UTC 2012 i686
  Compiled on Aug  5 2011
  UP for 7.82 hours
ixp1 IRRd#
```

**Comments:**

Test completed successfully.

### Test Case 10:

**IDENTIFICATION:**

TEST ID:	10
TITLE:	IPv4 connectivity test to OSIX peers
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Juniper and Cisco Routers
DATE OF TEST:	24.07.2012

S: Success

F: Failure

```
iqbalaun@ixp1:~$ ping 192.168.40.110
PING 192.168.40.110 (192.168.40.110) 56(84) bytes of data.
64 bytes from 192.168.40.110: icmp_req=1 ttl=255 time=0.828 ms
64 bytes from 192.168.40.110: icmp_req=2 ttl=255 time=0.798 ms
64 bytes from 192.168.40.110: icmp_req=3 ttl=255 time=0.804 ms
^C
--- 192.168.40.110 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2004ms
rtt min/avg/max/mdev = 0.798/0.810/0.828/0.013 ms
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$ ping 192.168.40.120
PING 192.168.40.120 (192.168.40.120) 56(84) bytes of data.
64 bytes from 192.168.40.120: icmp_req=1 ttl=64 time=0.323 ms
64 bytes from 192.168.40.120: icmp_req=2 ttl=64 time=0.553 ms
64 bytes from 192.168.40.120: icmp_req=3 ttl=64 time=0.274 ms
^C
--- 192.168.40.120 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2000ms
rtt min/avg/max/mdev = 0.274/0.383/0.553/0.122 ms
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$ ping 192.168.40.130
PING 192.168.40.130 (192.168.40.130) 56(84) bytes of data.
64 bytes from 192.168.40.130: icmp_req=1 ttl=64 time=0.531 ms
64 bytes from 192.168.40.130: icmp_req=2 ttl=64 time=0.493 ms
64 bytes from 192.168.40.130: icmp_req=3 ttl=64 time=0.452 ms
^C
--- 192.168.40.130 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 0.452/0.492/0.531/0.032 ms
iqbalaun@ixp1:~$
```

Comments: Test completed successfully.

**Test Case 11:**

IDENTIFICATION:	
TEST ID:	11
TITLE:	IPv6 connectivity test to OSIX peers
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Juniper and Cisco Routers
DATE OF TEST:	24.07.2012

- S: Success
- F: Failure

```

iqbalaun@ixp1:~$ ping6 fd9f:3a9:46ec:6917::110
PING fd9f:3a9:46ec:6917::110(fd9f:3a9:46ec:6917::110) 56 data bytes
64 bytes from fd9f:3a9:46ec:6917::110: icmp_seq=1 ttl=64 time=3.32 ms
64 bytes from fd9f:3a9:46ec:6917::110: icmp_seq=2 ttl=64 time=0.736 ms
64 bytes from fd9f:3a9:46ec:6917::110: icmp_seq=3 ttl=64 time=0.664 ms
^C
--- fd9f:3a9:46ec:6917::110 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 0.664/1.573/3.321/1.236 ms
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$ ping6 fd9f:3a9:46ec:6917::120
PING fd9f:3a9:46ec:6917::120(fd9f:3a9:46ec:6917::120) 56 data bytes
64 bytes from fd9f:3a9:46ec:6917::120: icmp_seq=1 ttl=64 time=1.58 ms
64 bytes from fd9f:3a9:46ec:6917::120: icmp_seq=2 ttl=64 time=1.41 ms
64 bytes from fd9f:3a9:46ec:6917::120: icmp_seq=3 ttl=64 time=2.03 ms
^C
--- fd9f:3a9:46ec:6917::120 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2004ms
rtt min/avg/max/mdev = 1.414/1.678/2.039/0.268 ms
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$ ping6 fd9f:3a9:46ec:6917::130
PING fd9f:3a9:46ec:6917::130(fd9f:3a9:46ec:6917::130) 56 data bytes
64 bytes from fd9f:3a9:46ec:6917::130: icmp_seq=1 ttl=64 time=1.50 ms
64 bytes from fd9f:3a9:46ec:6917::130: icmp_seq=2 ttl=64 time=2.06 ms
64 bytes from fd9f:3a9:46ec:6917::130: icmp_seq=3 ttl=64 time=1.39 ms
^C
--- fd9f:3a9:46ec:6917::130 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2005ms
rtt min/avg/max/mdev = 1.391/1.654/2.066/0.296 ms
iqbalaun@ixp1:~$

```

Comments:

Test completed successfully.

### Test Case 12:

#### IDENTIFICATION:

TEST ID:	12
TITLE:	SNMP Walk to all OSIX peers
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	SNMP
DATE OF TEST:	24.07.2012

- S: Success  
 F: Failure

```

iqbalaun@ixp1:~$ snmpwalk -v2c 192.168.40.110 -c public
iso.3.6.1.2.1.1.1.0 = STRING: "Cisco IOS Software, 1841 Software (C1841-
ADVENTERPRISEK9-M), Version 12.4(25a), RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled Fri 22-May-09 20:49 by prod_rel_team"
iso.3.6.1.2.1.1.2.0 = OID: iso.3.6.1.4.1.9.1.620
iso.3.6.1.2.1.1.3.0 = Timeticks: (2638718339) 305 days, 9:46:23.39
iso.3.6.1.2.1.1.4.0 = ""

iqbalaun@ixp1:~$ snmpwalk -v2c 192.168.40.120 -c public
iso.3.6.1.2.1.1.1.0 = STRING: "Juniper Networks, Inc. j2320 internet router, kernel
JUNOS 9.2R3.5 #0: 2009-01-15 04:29:52 UTC
builder@amalath.juniper.net:/volume/build/junos/9.2/release/9.2R3.5/obj-
i386/sys/compile/JSERIES Build date: 2009-01-15 04:55:15 UTC Copyright (c) 1996-200"
iso.3.6.1.2.1.1.2.0 = OID: iso.3.6.1.4.1.2636.1.1.1.2.23
iso.3.6.1.2.1.1.3.0 = Timeticks: (1012812943) 117 days, 5:22:09.43
iso.3.6.1.2.1.1.4.0 = STRING: "aunqbal@hotmail.com"

iqbalaun@ixp1:~$ snmpwalk -v2c 192.168.40.130 -c public
iso.3.6.1.2.1.1.1.0 = STRING: "Juniper Networks, Inc. j2320 internet router, kernel
JUNOS 9.2R3.5 #0: 2009-01-15 04:29:52 UTC
builder@amalath.juniper.net:/volume/build/junos/9.2/release/9.2R3.5/obj-
i386/sys/compile/JSERIES Build date: 2009-01-15 04:55:15 UTC Copyright (c) 1996-200"
iso.3.6.1.2.1.1.2.0 = OID: iso.3.6.1.4.1.2636.1.1.1.2.23
iso.3.6.1.2.1.1.3.0 = Timeticks: (1012896147) 117 days, 5:36:01.47
iso.3.6.1.2.1.1.4.0 = STRING: "aunqbal@hotmail.com"
iso.3.6.1.2.1.1.5.0 = STRING: "JR3"
iso.3.6.1.2.1.1.6.0 = STRING: "www.lab"
iso.3.6.1.2.1.1.7.0 = INTEGER: 4

```

Comments:

Test completed successfully.

### Test Case 13:

#### IDENTIFICATION:

TEST ID:	13
TITLE:	Configuration Backup of OSIX peers
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	RANCID
DATE OF TEST:	24.07.2012

- S: Success  
 F: Failure

```

iqbalaun@ixp1:/var/lib/rancid$
iqbalaun@ixp1:/var/lib/rancid$ bin/rancid-run
iqbalaun@ixp1:/var/lib/rancid$
iqbalaun@ixp1:/var/lib/rancid$ less Cisco/configs/
192.168.40.110 192.168.40.2 CVS/
iqbalaun@ixp1:/var/lib/rancid$ less Cisco/configs/
192.168.40.110 192.168.40.2 CVS/

```

Comments:

Test completed successfully.

### Test Case 14:

**IDENTIFICATION:**

TEST ID: 14  
TITLE: Apache Web Server is running with all services  
PLATFORM: Ubuntu 11.04 Server  
COMPONENT: Apache2 Web Server  
DATE OF TEST: 25.07.2012

- S: Success  
 F: Failure

```
iqbalaun@ixp1:~$ ps aux | grep apache
root      1405  0.0  0.5  5544  2556 ?        Ss   02:59   0:03 /usr/sbin/apache2 -k start
www-data  1406  0.0  0.3  5316  1784 ?        S    02:59   0:00 /usr/sbin/apache2 -k start
www-data  1409  0.0  0.5 227376  2940 ?        Sl   02:59   0:00 /usr/sbin/apache2 -k start
www-data  1412  0.0  0.5 227368  2668 ?        Sl   02:59   0:00 /usr/sbin/apache2 -k start
iqbalaun  8334  0.0  0.1   3616   880 pts/1    S+   11:51   0:00 grep --color=auto apache
iqbalaun@ixp1:~$
iqbalaun@ixp1:~$ sudo /etc/init.d/apache2 restart
[sudo] password for iqbalaun:
* Restarting web server apache2
... waiting ...done.
iqbalaun@ixp1:~$
```

**Comments:**

Test completed successfully.

### Test Case 15:

**IDENTIFICATION:**

TEST ID: 15  
TITLE: Updated configurations on CVS repository  
PLATFORM: Apache2 Server  
COMPONENT: CVSWEB  
DATE OF TEST: 25.07.2012

- S: Success  
 F: Failure

```
Juniper/configs/192.168.40.120 - diff - 1.3
130.195.10.91/cgi-bin/cvsweb/Juniper/configs/192.168.40.120.diff?r=H&only_with_tag=MAIN&r1=text&tr1=1.3&r2=text&tr2=1.2
[BACK] Return to 192.168.40.120 CVS log [TXT]
[DIR] Up to [Local Repository] / Juniper / configs
Diff for /Juniper/configs/192.168.40.120 between versions 1.2 and 1.3
version 1.2, 2012/07/18 11:28:32
Line 187
version 1.3, 2012/07/19 11:11:35
Line 187
```

**Comments:**

Test completed successfully.

### Test Case 16:

#### IDENTIFICATION:

TEST ID: 16  
TITLE: Graph to monitor bandwidth on all peers  
PLATFORM: Apache2 Web Server  
COMPONENT: MRTG  
DATE OF TEST: 25.07.2012

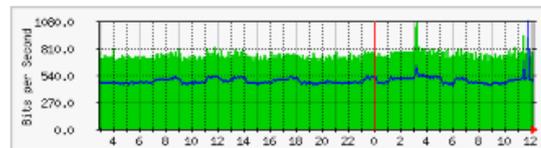
- S: Success  
 F: Failure

### Traffic Analysis for 192.168.40.130 -- JR3

System: JR3 in vuv\_lab  
Maintainer: aurigbal@hotmail.com  
Description: ge-0/0/0.0  
ifType: propVirtual (53)  
ifName: ge-0/0/0.0  
Max Speed: 1000.0 Mbits/s  
Ip: 192.168.40.130 ()

The statistics were last updated **Wednesday, 25 July 2012 at 12:10**,  
at which time 'JR3' had been up for **117 days, 6:23:07**.

#### Daily' Graph (5 Minute Average)



	Max	Average	Current
In	1080.0 b/s (0.0%)	728.0 b/s (0.0%)	680.0 b/s (0.0%)
Out	1080.0 b/s (0.0%)	480.0 b/s (0.0%)	448.0 b/s (0.0%)

#### Comments:

Test completed successfully.

### Test Case 17:

**IDENTIFICATION:**

TEST ID: 17  
TITLE: RPSL Scripts to generate RPSL based configuration  
PLATFORM: Ubuntu 11.04 Linux Server  
COMPONENT: RtConfig and Perl  
DATE OF TEST: 25.07.2012

S: Success

F: Failure

```
iqbalaun@ixp1:~$ cd /usr/local/src/rpsl/
iqbalaun@ixp1:/usr/local/src/rpsl$ cd osix
iqbalaun@ixp1:/usr/local/src/rpsl/osix$
iqbalaun@ixp1:/usr/local/src/rpsl/osix$ make
cd m91.dsrg.ecs.vuw.ac.nz; make
make[1]: Entering directory `/usr/local/src/rpsl/osix/m91.dsrg.ecs.vuw.ac.nz'
make[1]: `bgpd.conf' is up to date.
make[1]: Leaving directory `/usr/local/src/rpsl/osix/m91.dsrg.ecs.vuw.ac.nz'
cd m92.dsrg.ecs.vuw.ac.nz; make
make[1]: Entering directory `/usr/local/src/rpsl/osix/m92.dsrg.ecs.vuw.ac.nz'
make[1]: `bgpd.conf' is up to date.
make[1]: Leaving directory `/usr/local/src/rpsl/osix/m92.dsrg.ecs.vuw.ac.nz'
iqbalaun@ixp1:/usr/local/src/rpsl/osix$
iqbalaun@ixp1:/usr/local/src/rpsl/osix$
iqbalaun@ixp1:/usr/local/src/rpsl/osix$ less m91.dsrg.ecs.vuw.ac.nz/bgpd.conf
-----
! $Source: /usr/local/src/rpsl/osix/m91.dsrg.ecs.vuw.ac.nz/RCS/bgpd.conf,v $
! $Log: bgpd.conf,v $
! Revision 1.516 2012/07/18 22:06:47 root
! Backed up by make
```

**Comments:**

Test completed successfully.

### Test Case 18:

**IDENTIFICATION:**

TEST ID: 18  
TITLE: Looking Glass for IPv4 BGP Summary  
PLATFORM: Apache2 Web Server  
COMPONENT: Looking Glass  
DATE OF TEST: 25.07.2012

S: Success

F: Failure

Router: JR3 Junos  
 Command: show bgp summary

```

Groups: 4 Peers: 8 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0         0          0          0          0        0      0       0
inet6.0        7          1          0          0        0      0       0
Peer
192.168.40.91 7675      AS      InPkt   OutPkt   OutQ    Flaps Last Up/Dwn State|#Active/Received/Accepted/Damped...
inet.0: 0/0/0/0
192.168.40.92 8285      1114    1213    0        4      9:16:08 Establ
inet.0: 0/0/0/0
192.168.40.110 110      139764  152880  0        4      6w6d12h Establ
inet.0: 0/0/0/0
192.168.40.120 120      152937  152880  0        3      6w6d12h Establ
inet.0: 0/0/0/0
fd9f:3a9:46ec:6917::91 7675    1142    1242    0        9      9:29:05 Establ
inet6.0: 1/2/2/0
fd9f:3a9:46ec:6917::92 8285      1116    1214    0        4      9:16:07 Establ
inet6.0: 0/2/2/0
fd9f:3a9:46ec:6917::110 110     139817  152922  0        2      6w6d12h Establ
inet6.0: 0/2/2/0
fd9f:3a9:46ec:6917::120 120     152968  152911  0        3      6w6d12h Establ
inet6.0: 0/1/1/0
  
```

Comments:

Test completed successfully.

### Test Case 19:

#### IDENTIFICATION:

TEST ID: 19  
 TITLE: Looking Glass for IPv6 BGP Summary  
 PLATFORM: Apache2 Web Server  
 COMPONENT: Looking Glass  
 DATE OF TEST: 25.07.2012

- S: Success  
 F: Failure

Router: CR1

Command: show bgp ipv6 summary

```

BGP router identifier 2.2.2.2, local AS number 110
BGP table version is 66, main routing table version 66
2 network entries using 298 bytes of memory
5 path entries using 380 bytes of memory
6/2 BGP path/bestpath attribute entries using 744 bytes of memory
3 BGP AS-PATH entries using 72 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1494 total bytes of memory
BGP activity 15/13 prefixes, 54/49 paths, scan interval 60 secs

Neighbor      V   AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
FD9F:3A9:46EC:6917::91
4 7675  62369  62458    66     0    0 09:30:28    2
FD9F:3A9:46EC:6917::92
4 8285  10693  10702    66     0    0 09:17:32    2
FD9F:3A9:46EC:6917::130
4 130  154642 141386    66     0    0 6w6d        1
% NOTE: This command is deprecated. Please use 'show bgp ipv6 unicast'
  
```

Comments:

Test completed successfully.

**Test Case 20:**

IDENTIFICATION:

TEST ID:	20
TITLE:	WHOIS query on OSIX server
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Apache2 Server
DATE OF TEST:	25.07.2012

S: Success

F: Failure

Comments:

Test completed successfully.

**Test Case 21:**

IDENTIFICATION:

TEST ID:	21
TITLE:	“Show Route” query on OSIX server
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Apache2 Server
DATE OF TEST:	25.07.2012

S: Success

F: Failure

Comments:

Test completed successfully.

**Test Case 22:**

**IDENTIFICATION:**

TEST ID:	22
TITLE:	“Change Route” query on OSIX server
PLATFORM:	Ubuntu 11.04 Server
COMPONENT:	Apache2 Server
DATE OF TEST:	25.07.2012

- S: Success  
 F: Failure

**Comments:**

Test completed successfully.

**Test Case 23:**

**IDENTIFICATION:**

TEST ID:	23
TITLE:	OSIX Log Server is running
PLATFORM:	Ubuntu 11.04 Linux
COMPONENT:	Log Server
DATE OF TEST:	25.07.2012

- S: Success  
 F: Failure

**Comments:**

Test completed successfully.

## Test Case 24:

### IDENTIFICATION:

TEST ID:	24
TITLE:	Ensure all required services are running
PLATFORM:	Ubuntu Linux 11.04
COMPONENT:	Services
DATE OF TEST:	16.08.2012

S: Success

F: Failure

```
iqbal@ixp1:~$ service --status-all
[ ? ] acpi-support
[ ? ] acpid
[ ? ] alsa-restore
[ ? ] alsa-store
[ ? ] anacron
[ + ] apache2
[ - ] apparmor
[ ? ] appport
[ ? ] atd
[ ? ] avahi-daemon
[ + ] bind9
[ ? ] binfmt-support
[ ? ] bird
[ ? ] bird6
[ - ] bluetooth
[ - ] bootlogd
[ - ] brltty
[ ? ] console-setup
[ ? ] cron
[ ? ] cups
[ ? ] dbus
[ ? ] dmesg
[ ? ] dns-clean
[ ? ] failsafe-x
[ ? ] gdm
[ - ] grub-common
[ ? ] hostname
[ ? ] hwclock
[ ? ] hwclock-save
[ ? ] irqbalance
[ - ] kerneloops
[ ? ] killprocs
[ ? ] module-init-tools
[ ? ] network-interface
[ ? ] network-interface-security
[ ? ] network-manager
[ ? ] networking
[ ? ] ondemand
[ - ] openbsd-inetd
[ ? ] pcmciautils
[ ? ] plymouth
[ ? ] plymouth-log
[ ? ] plymouth-splash
[ ? ] plymouth-stop
[ ? ] plymouth-upstart-bridge
[ ? ] pppd-dns
[ ? ] procps
[ + ] pulseaudio
[ ? ] quagga
[ ? ] rc.local
[ - ] rsync
[ ? ] rsyslog
[ + ] saned
[ ? ] screen-cleanup
[ ? ] sendsigs
[ ? ] setvtrgb
[ ? ] speech-dispatcher
[ + ] ssh
[ ? ] stop-bootlogd
[ ? ] stop-bootlogd-single
[ ? ] sudo
```

```
[ ? ] udev
[ ? ] udev-fallback-graphics
[ ? ] udev-finish
[ ? ] udevmonitor
[ ? ] udevtrigger
[ ? ] ufw
[ ? ] umountfs
[ ? ] umountnfs.sh
[ ? ] umountroot
[ ? ] unattended-upgrades
[ - ] urandom
[ - ] x11-common
iqbalaun@ixpl:~$
```

```
iqbalaun@ixpl:~$ nmap localhost

Starting Nmap 5.21 ( http://nmap.org ) at 2012-07-17 11:56 NZST
Nmap scan report for localhost (127.0.0.1)
Host is up (0.0011s latency).
Not shown: 991 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
43/tcp    open  whois
53/tcp    open  domain
80/tcp    open  http
179/tcp   open  bgp
631/tcp   open  ipp
2601/tcp  open  zebra
2605/tcp  open  bgpd
5900/tcp  open  vnc

Nmap done: 1 IP address (1 host up) scanned in 0.20 seconds
iqbalaun@ixpl:~$
```

Comments:

Test completed successfully.

## Bibliography

- [1] Mohammad Z. Ahmad and Ratan Guha. 2010. Understanding the impact of internet exchange points on internet topology and routing performance. In *Proceedings of the ACM CoNEXT Student Workshop (CoNEXT '10 Student Workshop)*. ACM, New York, NY, USA, , Article 18. DOI=10.1145/1921206.1921226 <http://doi.acm.org/10.1145/1921206.1921226>
- [2] Ahmad, Mohammad Zubair; Guha, Ratan; , "Impact of Internet exchange points on Internet topology evolution," *Local Computer Networks (LCN)*, 2010 IEEE 35th Conference on , vol., no., pp.332-335, 10-14 Oct. 2010 DOI: 10.1109/LCN.2010.5735736 URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5735736&isnumber=5735635>
- [3] Ahmad, Mohammad Zubair; Guha, Ratan; , "A measurement study determining the effect of Internet eXchange points on popular webservers," *Local Computer Networks (LCN)*, 2010 IEEE 35th Conference on , vol., no., pp.976-982, 10-14 Oct. 2010 DOI: 10.1109/LCN.2010.5735844 URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5735844&isnumber=5735635>
- [4] Yihua He, Georgos Siganos, Michalis Faloutsos, and Srikanth Krishnamurthy. 2009. Lord of the links: a framework for discovering missing links in the internet topology. *IEEE/ACM Trans. Netw.* 17, 2 (April 2009), 391-404. DOI=10.1109/TNET.2008.926512 <http://dx.doi.org/10.1109/TNET.2008.926512>
- [5] Brice Augustin, Balachander Krishnamurthy, and Walter Willinger. 2009. IXPs: mapped?. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference (IMC '09)*. ACM, New York, NY, USA, 336-349. DOI=10.1145/1644893.1644934 <http://doi.acm.org/10.1145/1644893.1644934>
- [6] Matthias Wahlisch and Thomas C. Schmidt. 2009. Peer the peers: an overlay id assignment service at internet exchange points. In *Proceedings of the 5th international student workshop on Emerging networking experiments and technologies (Co-Next Student Workshop '09)*. ACM, New York, NY, USA, 45-46. DOI=10.1145/1658997.1659022 <http://doi.acm.org/10.1145/1658997.1659022>
- [7] BEREC Conference (29th May 20120 - An assessment of IP-interconnection in the context of Net Neutrality - [http://berec.europa.eu/doc/consult/bor\\_12\\_33\\_ip\\_ic\\_assessment.pdf](http://berec.europa.eu/doc/consult/bor_12_33_ip_ic_assessment.pdf). [Online accessed 23-July-2012]
- [8] Zhuoqing Morley Mao, Jennifer Rexford, Jia Wang, and Randy H. Katz. 2003. Towards an accurate AS-level traceroute tool. In *Proceedings of the 2003*

*conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM '03)*. ACM, New York, NY, USA, 365-378. DOI=10.1145/863955.863996  
<http://doi.acm.org/10.1145/863955.863996>

- [9] Labovitz Craig, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. 2000. Delayed Internet routing convergence. *SIGCOMM Comput. Commun. Rev.* 30, 4 (August 2000), 175-187. DOI=10.1145/347057.347428  
<http://doi.acm.org/10.1145/347057.347428>
- [10] Cisco Systems (2010) – Internet Exchange Point Design – ISP/IXP Workshops. URL: [www.pacnog.org/pacnog6/IXP/IXP-design.pdf](http://www.pacnog.org/pacnog6/IXP/IXP-design.pdf) [Online accessed 23-August-2012]
- [11] Domingues, M.; Friacas, C.; , "IPv6 in European Internet eXchange Points," Networking and Services, 2007. ICNS. Third International Conference on , vol., no., pp.109, 19-25 June 2007 DOI: 10.1109/ICNS.2007.77  
URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4438358&isnumber=4438243>
- [12] Bertram Neil - New Zealand's Internet Landscape: An analysis of peering, content, and scalability (VUW Honours Project)
- [13] Woodcock Bill (PCH) [July 2004] - APBDC Model v 0.9 -  
<http://www.pch.net/docs/tutorials/average-per-bit-delivery-cost/APBDC-Tutorial-v09.txt> [Online accessed 16-August-2012]
- [14] Rose, K.; , "Africa Shifts Focus from Infrastructure to Interconnection," Internet Computing, IEEE , vol.14, no.6, pp.56-58, Nov.-Dec. 2010  
DOI: 10.1109/MIC.2010.132  
URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5617058&isnumber=5617049>
- [15] McLaughlin A. (May 2004). Internet Exchange Points Their Importance to Development of the Internet and Strategies for their Deployment – The African Example. URL:  
<http://www.isoc.org/educpillar/resources/docs/promote-ixp-guide.pdf> [Online accessed 23-August-2012]
- [16] Open Peering initiative - Peering & Transit - Public Peering Link:  
<http://www.openpeering.nl/publicpeering.shtml> [Online accessed 23-August-2012]
- [17] Internet Exchange Points (IXPs). (Briefing Paper, Internet Society, 2009).  
URL: <http://www.internetsociety.org/fr/node/648>. [Online accessed 23-August-2012]
- [18] Augustin Brice, Balachander Krishnamurthy, Walter Willinger (November 2009). IXPs: mapped? IMC '09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference

- [19] Norton, W. B. 2000, Interconnection Strategies for ISPs, Equinix, <http://www.equinix.com/pdf/whitepapers/ISPInterconnectionStrategies2.pdf>. [Online; accessed 23 - July-2012]
- [20] DrPeering Press – Norton, W. B. (2011) – The Internet Peering Playbook – Connecting to the Core of the Internet. 2<sup>nd</sup> Edition. ISBN: eBook PDF ISBN: 978-1-937451-01-1
- [21] Jensen M. - A Summary Report Promoting the Use of Internet Exchange Points: A Guide to Policy, Management, and Technical Issues - <http://www.isoc.org/educpillar/resources/docs/promote-ixp-summary.pdf>. [Online; accessed 23 - July-2012]
- [22] PCH IXP Growth real-time Graph - [https://prefix.pch.net/applications/ixpdir/summary/growth-region/?sort1=ixp&sort2=\\_percent\\_change&order=desc](https://prefix.pch.net/applications/ixpdir/summary/growth-region/?sort1=ixp&sort2=_percent_change&order=desc). [Online; accessed 16-August-2012]
- [23] Report from the IGF Rio Best Practices Session Internet Traffic Exchange in Less Developed Internet: Markets and the Role of Internet Exchange Points. The Internet Society, <http://www.isoc.org/educpillar/resources/igf-ixp-report-2007.shtml>. [Online; accessed 23 - July-2012]
- [24] Internet Users as percentage of population - [http://www.google.co.nz/publicdata/explore?ds=d5bncppjof8f9\\_&met\\_y=it\\_net\\_user\\_p2&idim=country:NZL&dl=en&hl=en&q=new+zealand+internet+usage](http://www.google.co.nz/publicdata/explore?ds=d5bncppjof8f9_&met_y=it_net_user_p2&idim=country:NZL&dl=en&hl=en&q=new+zealand+internet+usage). [Online; accessed 23 - July-2012]
- [25] Topological Map of New Zealand Internet Service providers <http://www.ispmap.co.nz/topmap.html>. [Online; accessed 23 - July-2012]
- [26] SCCN Fibre Connections - <http://www.southerncrosscables.com/public/Network/default.cfm?PageID=42&MenuID=41>. [Online; accessed 23 - July-2012]
- [27] Waikato Linux Users Group (and contributors). New Zealand internet history. <http://www.wlug.org.nz/NewZealandInternetHistory>. [Online; accessed 23 - July-2012]
- [28] WIX: a Distributed Internet Exchange - May 26, 2005 By Richard Hulse - <http://www.linuxjournal.com/article/8073>. [Online; accessed 23 - July-2012]
- [29] Bill Woodcock & Vijay Adhikari1 [May 2 2011] - Packet Clearing House: Survey of Characteristics of Internet Carrier Interconnection Agreements
- [30] DE-CIX, 5 year graph Traffic Statistics - <http://www.de-cix.net/about/statistics>. [Online; accessed 23 - July-2012]