

NAPLA
May 22, 2007

Bill Woodcock

Prefix-List Sanity-Checker

A Free Online Tool

by and for Network Operators

Version 1.2

March, 2007

Bill Woodcock

Packet Clearing House

Goals:

An interactive tool to help ISP provisioning techs qualify prefixes that customers want to BGP-advertise.

Generate properly-formatted Cisco and Juniper prefix-lists to cut-and-paste into configurations or provisioning systems.

Input Parsing

Pasted or uploaded text in many formats:

CIDR or classful

Address-ranges (including pre-CIDR)

Cisco or Juniper prefix or filter lists

Cisco route statements

sho ip bgp

sho ip route

RIR whois entries

IRR entries

Analysis

Flag bogons and critical infrastructure like IXP subnets and root / TLD nameservers.

Flag long prefixes.

Graphically map geographic distribution.

Malware incident to address ratio report

Table sortable by registrant, country, originating AS, malware reports, or status.

Output

Cisco prefix-lists

Juniper prefix-lists

Juniper set-commands



Live Demo!

Future Features

Custom acceptable prefix-length tables

Custom prefix-rejection tables

Custom AS-rejection and accept tables

Previous transit AS uniformity check

Suggest potential aggregates

XML API for provisioning system integration

IRR submission generation

IPv6

Others?

Known Issues

Performance.

The entire service is currently running on a single 1.66Ghz Mac Mini. As use picks up, we'll look for a grant to cover a stack of Xserves for it.

Thanks, and Questions?

Copies of this presentation can be found in PDF and QuickTime formats at:
[http:// www.pch.net / resources / papers / prefix-list-sanity-checker](http://www.pch.net/resources/papers/prefix-list-sanity-checker)

More importantly, the tool discussed in this presentation can be found at:

[https:// prefix.pch.net](https://prefix.pch.net)

Please give it a try, and let us know what features would make it more useful to you.

Bill Woodcock
Packet Clearing House
woody@pch.net

Netflow Routing Economics Analyzer

A Free Online Tool
by and for Network Operators

Version 0.9

May, 2007

Bill Woodcock

Packet Clearing House

Goals

Re-write our previous tool, which was too slow, too ugly to distribute, batch-oriented, and required ISPs to ship us hard disks each time they wanted a report.

Provide timely and easy-to-understand analysis of potential peering arrangements and potential IXP participation

Provide an overview of international Internet traffic for the research and development-aid communities.

Input

ISPs already peer with us in more than 1000 locations, which gives us an understanding of the topology of the Internet.

Anyone who wants to can direct netflow exports (sampled or full) to an adjacent hosted analyzer, or can run the analyzer themselves.

Output

Private per-ISP report tells you how much traffic (total and percentage) you could push/pull through another ASN (including all their down-stream customer ASes) if you peered with them.

Private per-IXP report tells you how much traffic you could push/pull through the open peers (including their customers) at each IXP, if you participated at it.

Output

Public country summary aggregates and anonymizes data to show what portions of each country's traffic are exchanged with each other country.

Public region summary shows the same thing between regions.

All reports updated daily, and summarized weekly, monthly, and annually. Available by PGP email or through an HTTPS web login.



Live Demo!

Current State of the Project

Code is about 85% complete.

Per-ISP report and most analysis code is done, Per-IXP and summary reports, and email delivery code not yet complete.

Internal alpha testing, nearly ready for external beta testing.

Privacy Policy

PCH does not (cannot, in fact, due to the volume) retain any netflow data.

Data belongs solely to the ISP that generates it, and only that ISP gets the analysis based on it.

PCH will anonymize and aggregate data in order to make public national-level and regional-level reports. These contain no individually-identifiable data.

It's open-source, so anyone who doesn't trust us or the privacy agreement is welcome to inspect the code and run it on their own servers.

Agreement with LACNIC

LACNIC is providing Spanish and Portuguese localization of all user interface and documentation.

Initial beta sites will be in São Paulo, Buenos Aires, and Miami.

At least three more locations in Latin America within the next year.

Good Practices in Internet Exchange Point Documentation

Version 1.4

May, 2007

Bill Woodcock

Packet Clearing House

Background: Motivation

Since 1994, PCH has been maintaining a canonical global directory of Internet exchange points, visible at www.pch.net/ixpdir. We are frequently approached by researchers, NGOs, and governmental and intergovernmental bodies, asking for more complete and more accurate numbers which can be used to characterize the size and rate of growth of the Internet.

Motivation

The OECD and the FCC have asked PCH to work with the rest of the global Internet exchange community to prepare a good practices document which describes documentation practices for Internet exchange points and, once done, to promote it and help IXes implement it, globally.

Balance

On one hand, we're trying to solve a problem which exists due to present insufficient documentation. On the other hand, we don't want to claim a "best practice" that is not actually currently implemented.

Balance

Therefore, we've taken as our model the practice as currently extant in the IXes which have generally been most consistently diligent in documenting their operations. We have examined the many specific current practices and distilled what we believe to be their commonality, and about half a dozen IXes have now implemented it exactly as written in this document. We believe this is a reasonable and achievable goal, while being sufficient solution to be worth implementing.

Specifically

Publication of traffic statistics

Publication of membership contacts

Maintenance of the IN-ADDR zones

Publication of Traffic Statistics

On the IX web site, in addition to any other localized or translated URLs, a page shall be located at `/statistics` relative to the root of the web site, containing any statistical information which the IX chooses to publish, in the language of its choice. This shall include at a minimum, a graphical histogram of the past 24 hours' aggregate traffic, counting each bit which has flowed across the IX switch fabric exactly once, with care not to double-count at inter-switch connections. The time-base of the graph may be in the local time-zone of the IX. The graph image itself shall also be reachable at `/statistics/graph`. The same data shall also be provided in tabular form at `/statistics/table`, of at least 288 lines, where each consists of a Unix-style timestamp five-minute aligned UTC, a tab, and the number of bits passed in the preceding five-minute period, divided by 300 and rounded to the nearest integer (in other words, five-minute average bits per second), followed by a carriage return, and the next entry. In the event that data is simply unavailable, no value should follow the tab. Under no circumstances should a lack of data be represented by a zero or an interpolated value. The earliest entry shall be at the beginning (top) of the file, and the most recent entry shall be at the end (bottom) of the file.

Publication of Participant Contacts

On the IX web site, in addition to any other localized or translated URLs, a page shall be located at `/participants` relative to the root of the web site, containing any information about the IX's membership or participants which the IX chooses to publish, in the language of its choice. This shall include, at a minimum, the organization name and contact method, such as a peering information URL, for each participant. Participant data shall also be provided in tabular form at `/participants/table`, exactly one line per assigned IP address of each subnet used for public interconnection across the IX, where each line consists of the assigned IP address, a slash, the subnet mask length in bits, a tab, the organization name of the participant to which it has been assigned, a tab, the principal Autonomous System Number which the participant reports to be in use on that interface, a tab, and the preferred contact information for the responsible party for that participant organization, followed by a carriage return, and the next entry. Unassigned IP addresses shall not appear in the file. The file shall be sorted by IP address, with the lowest address at the beginning (top) of the file, and the highest address at the end (bottom) of the file. ISPs shall not use private Autonomous System Numbers for public peering.

Maintenance of the IN-ADDR

In the publicly-visible in-addr DNS, one PTR record shall be maintained for each assigned IP address in each subnet used for public interconnection across the IX, consisting of the fully-qualified domain name received from the participant to which that IP address is assigned, and each participant shall be encouraged to maintain a matching A record in their forward domain. Unassigned IP addresses shall have no associated PTR record in the publicly-visible DNS. Address ranges for which the IX does not control authoritative IN-ADDR delegation shall not be used for the exchange of Internet traffic.

Cautions to Client Implementors

Note that the canonical URLs specified in this document are likely to be implemented on the server side using either http redirects or html redirects, so clients must be prepared to follow such redirects to reach the data.

This Document Describes Minimums

It is important to understand that this document describes a least common denominator practice, intended to facilitate easy programmatic interaction, and is in no way intended to discourage the implementation of additional features or language localization for the convenience of the IX's first constituency, its participants.

Beta-Phase Implementors

Amsterdam Internet Exchange

Seattle Internet Exchange

Nepal Internet Exchange

Toronto Internet Exchange

West Australia Internet Exchange

Wellington Internet Exchange

Internet Exchange of Puerto Rico

Feedback, Please!

This is a living document, and although its utility will greatly benefit from stability, we should also expect the document to be updated periodically, as needs and capabilities change.

Thanks, and Questions?

Copies of this presentation can be found
in Keynote, PDF, and QuickTime formats at:

<http://www.pch.net/resources/papers/ix-documentation-bcp>

Bill Woodcock
Research Director
Packet Clearing House
woody@pch.net

How These Last Two Fit Together

The IXP traffic statistics and the netflow analyzer may at first appear to be unrelated, but actually, they're complementary in a way which will give us a lot more information about the size and shape of the Internet...

A Way Forward

Much, perhaps most, of the Internet's traffic flows through Internet exchange points. This traffic can be *quantified*, perhaps even to the extent of coming to a relatively accurate total, but it is very difficult to characterize *qualitatively*, due to performance and privacy considerations.

A Way Forward

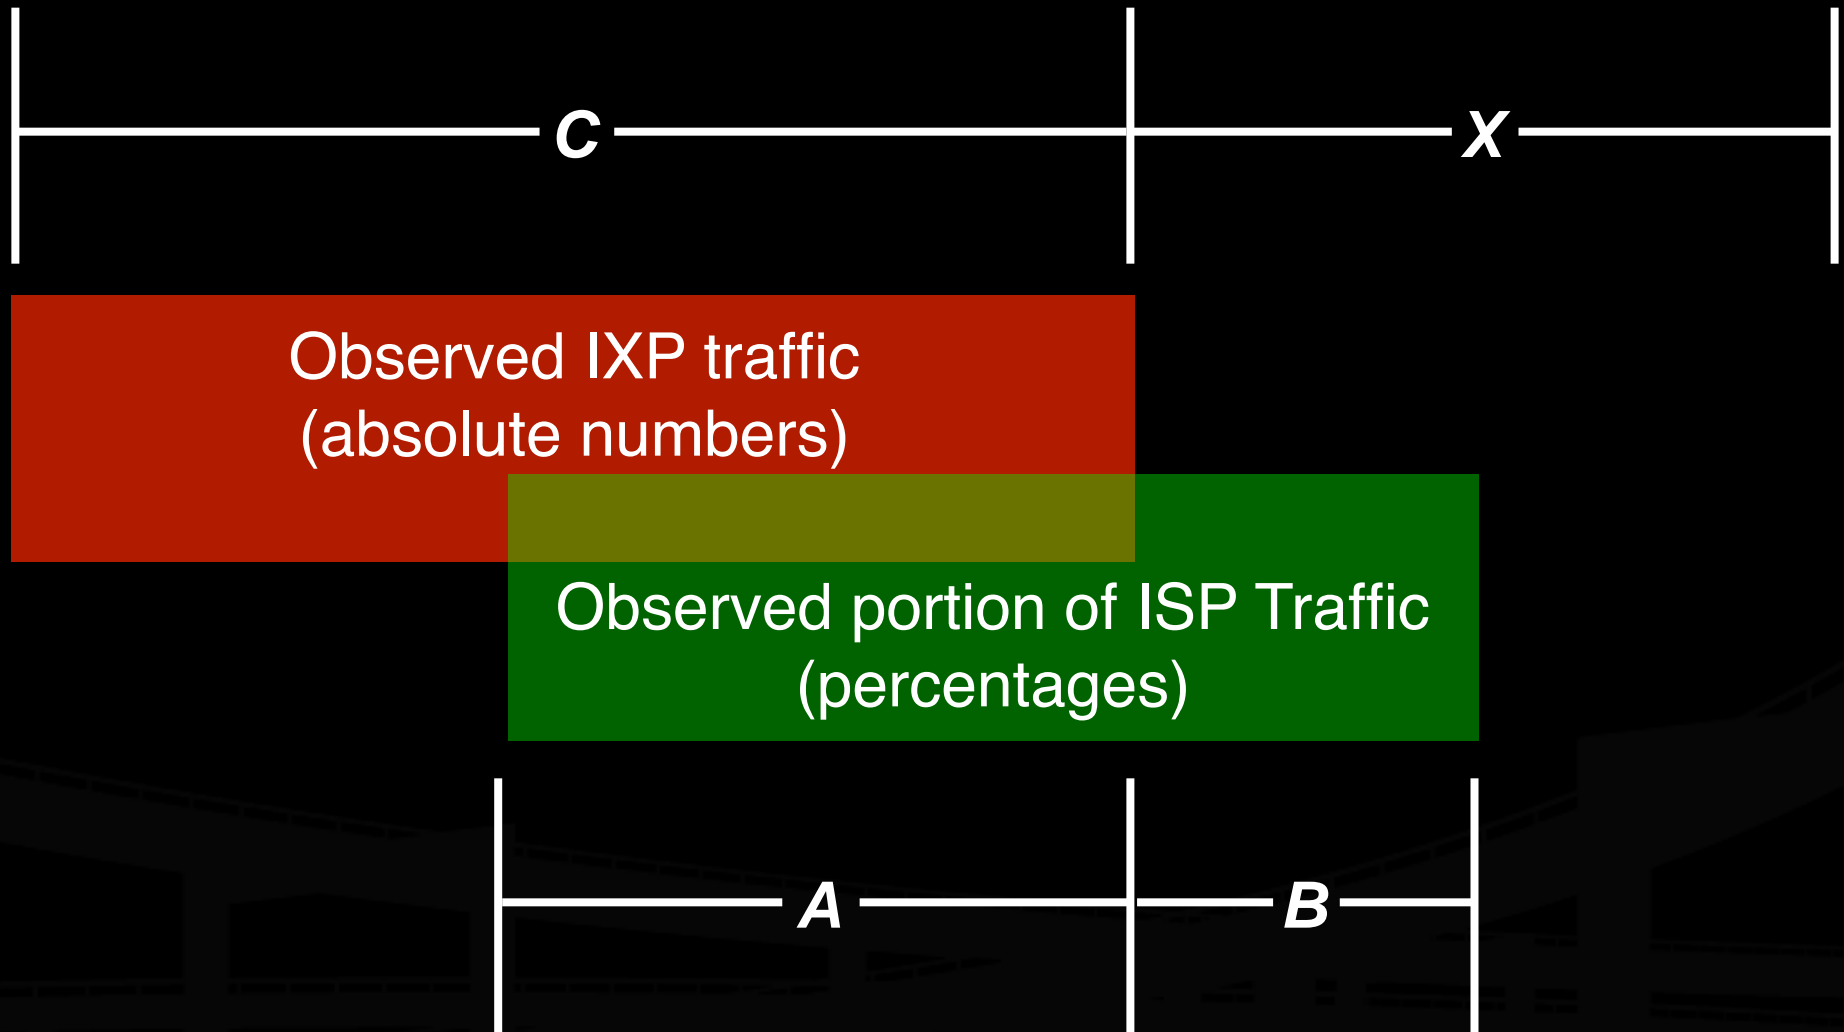
The Internet's traffic all flows through Internet service providers. Because these networks are numerous and are privately and diversely held, it is impossible to *quantify* their size or to directly measure what portion of the total one is observing. But one can accurately *qualitatively* characterize the portion that one does see.

A Way Forward


Observed IXP traffic
(absolute numbers)

Observed portion of ISP Traffic
(percentages)

A Way Forward



A Way Forward



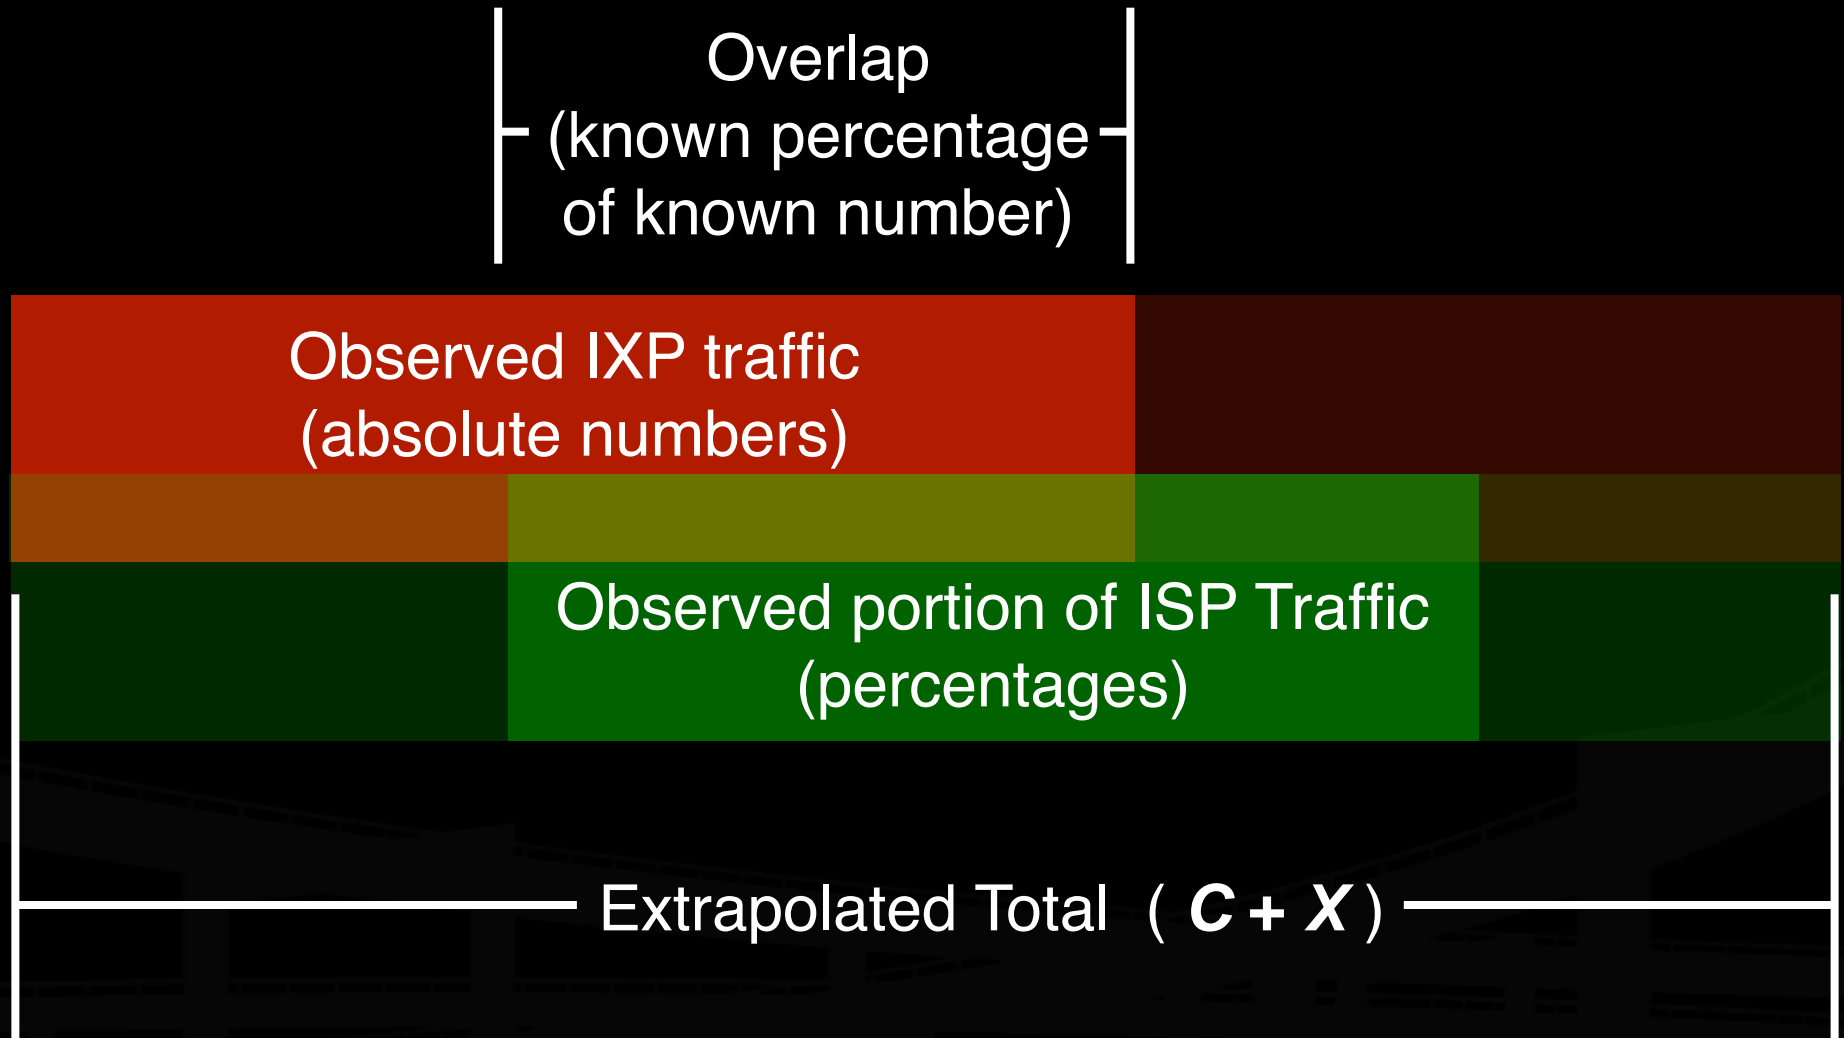
Observed IXP traffic
(absolute numbers)

$$\frac{A}{B} \approx \frac{C}{X}$$

Observed portion of ISP Traffic
(percentages)



A Way Forward




So where do we get the numbers?



ISP Netflow Exports

So where do we get the numbers?

IXP Aggregate Bit-Counts



Okay, now I'm *really* done.

Muchas gracias.

