

Interconnection Schemes and
Strategies for ISPs Tutorial:
Internet Exchange Points

Monday Morning, May 26, 2008

In Conjunction with NAPLA / LACNIC


Bill Woodcock

Research Director

Packet Clearing House


ISP Lifecycle: Simple Aggregator

Single Transit Provider ——— IXPs




ISP Lifecycle: Redundancy and LCR

Redundant Transit Providers — IXPs




ISP Lifecycle: Local Peer

Redundant Transit Providers — IXP




ISP Lifecycle: Network Service Provider

Redundant Transit Providers — IXP




Hot Potato Routing




Hot Potato Routing

Red Customer sends to Green Customer via Red NSP




Hot Potato Routing

Red NSP delivers at *nearest* IXP




Hot Potato Routing

Green NSP backhauls from distant IXP




Hot Potato Routing

Green ISP delivers to Green Customer




Hot Potato Routing

Green Customer replies via Green NSP




Hot Potato Routing

Green NSP delivers at nearest IXP




Hot Potato Routing

Red NSP backhauls from distant IXP




Hot Potato Routing

Red NSP delivers to Red Customer




Hot Potato Routing

Red Network is responsible for its own costs




Hot Potato Routing

Green Network is responsible for its own costs



Hot Potato Routing

Symmetry: Fair sharing of costs



The old circuit-switched networks have dubbed our financial model “bill and keep”

Tools for thinking about Internet Exchanges in economic terms

What are we, as ISPs, selling?

The right to modulate bits.

That right is a perishable commodity.

Where do we get the potentially-modulatable bits?

The right to modulate bits

Any Internet connection is a serial stream of time-slices.

Each time-slice can be modulated with a binary one or zero, one bit.

Each customer purchases potentially-modulatable bits at some *rate*, for example, 2mbps, which is 5.27 trillion bits per monthly billing cycle.

That's a perishable commodity

The quality (as opposed to quantity-per-time) characteristics of an Internet connection are *loss, latency, jitter, and out-of-order delivery*.

Loss increases as a function of the number and reliability of components in the path, and the amount of contention for capacity.

Latency increases as a function of distance, and degree of utilization of transmission buffers by competing traffic sources.

Jitter is the degree of variability in loss and latency, which negatively affects the efficacy and efficiency of the encoding schemes which mitigate their effects. Jitter increases relative to the ratio of traffic burstiness to number of sources.

Out-of-order delivery is the portion of packets which arrive later than other, subsequently-transmitted packets. It increases as a function of the difference in queueing delay on parallel paths.

All of **these properties become worse with time and distance**, which is a reasonable definition of a perishable commodity.

So where do we get the bits?

The value of the Internet is communication.

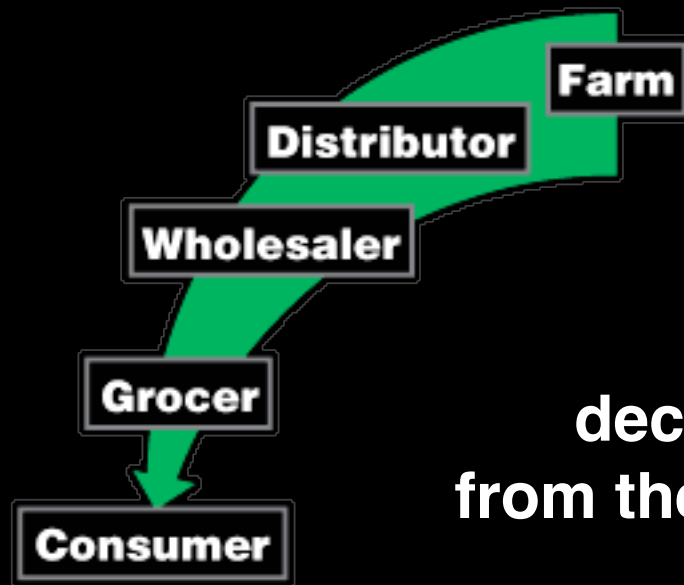
The value is produced at the point at which communication occurs between two ISPs, and it is transported to the customers who utilize it.

Thus, all the bits we sell come from an Internet exchange, whether nearby, or far away.

An analogy

Let's look at another perishable commodity with more readily observed economic properties... **Bananas.**


Value decreases with time & distance



The **value** of a banana decreases, the further it gets from the farm which produced it.


The shelf-life which the consumer can expect decreases, and eventually it becomes overripe, then rotten.


Cost increases with time & distance



The **cost** of a banana increases, the further it gets from the farm which produced it.


Salaries and hourly labor, warehouse leasing, diesel fuel, truck amortization, loss and spoilage, insurance, and other factors contribute additively.






In a competitive environment, retail price is limited by competition, so time and distance influence the price more than the number of middlemen.

The problem is the same:




ISPs form a delivery chain, bringing perishable bits to the consumers who purchase them.

So how do we improve things?



Bring the customer nearer an IX...



...or bring an IX nearer the customer.

So how do we recognize a successful exchange?

The purpose of an IX is to lower participating ISPs' average per bit delivery costs (APBDC).

A cheap IX is probably a successful one.
An expensive IX is always a failure.
Reliability is just hand-waving by salespeople.

The irony inherent in that

An efficient IX is an ISP's lowest-cost delivery method.

In order to shift latency-sensitive traffic toward the lowest-cost delivery method, it must also be the highest-capacity pipe.

Regardless of degree of utilization.

Thus many IX connections run at low utilization: apparent inefficiency.

Break

Determining Need

Sufficient end-user base?

No existing facility to build upon?

Sufficient degree of locally-destined traffic?

Geographic Location

User population

Fiber facilities or rights-of-way

Founding participants

Density

Centralized in one room

Campus of adjacent buildings

MAN

Frame or ATM cloud

Building Management

Telco hotel

University computing or
telecommunications facility

City emergency services facility

In-Building Facilities

Pathways

Power

Cooling

Access and security

Services

Switch fabric

Crossconnects

Route-server

Remote hands

NTP

Web caching

Business Structure

Incorporated or unincorporated?

Staffed or volunteer?

Non-profit or for-profit?

Cooperative or external ownership?

Cost-recovery (predictive or actuals), ad-hoc, or market pricing?

Policies

BLP, MLPA or MMPLA?

Mandatory looking-glass?

Routing and switch-port information
public or members-only?

Secrecy in the event of security
problems, failures, or mistakes

Extensible switch fabric?

Thanks, and Questions?

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

[http:// www.pch.net / resources / tutorials / ix-construction](http://www.pch.net/resources/tutorials/ix-construction)

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

Lessons Learned from the Russian-Estonian Cyber-Conflict

Version 1.0

June, 2007

Bill Woodcock

Packet Clearing House

What Was New?

Incidents beyond counting in the last twenty years, but...

What Was New?

Incidents beyond counting in the last twenty years, but very few have been state-actor...

What Was New?

Incidents beyond counting in the last twenty years, but very few have been state-actor, this one was relatively large (though not optimized for size or network impact)...

What Was New?

Incidents beyond counting in the last twenty years, but very few have been state-actor, this one was relatively large (though not optimized for size or network impact), and very few have been defended with such complete success.

What Was New?

Incidents beyond counting in the last twenty years, but very few have been state-actor, this one was relatively large (though not optimized for size or network impact), and very few have been defended with such complete success.

Thus, there are new lessons to be learned, and not just the usual NSP operators should be paying attention.

What Made Estonia so Successful?

Command structure was young and connected.

They skipped the whole “is the network worth protecting” step that older countries seem to invariably get hung up on.

They already had the necessary channels of communication established prior to the attack.

What Makes Other Countries More Vulnerable?

Lack of understanding and conviction.

Lack of commitment to funding defensive operations.

Scale. Many countries are too large for everyone who matters to already know everyone else who matters.

Roles in Cyber-Conflict

User population

Network service providers

CERT

Law enforcement

Ministry of Foreign Affairs

Military

Roles in Cyber-Conflict

User population	Intelligence
Network service providers	Defense
CERT	Analysis & coordination
Law enforcement	Domestic prevention
Ministry of Foreign	International prevention
Military	Credible threat of offense

User Population

The population of users – people in their homes and offices – are essentially the only ones who can determine authoritatively that an attack is taking place and provide the intelligence that differentiates an attack from productive traffic.

What differentiates a DDoS from a slashdotting? Only end-user expectations.

Network Service Providers

NSPs have absolute control over the field of action when they choose to exercise it, but are essentially neutral until called into play by their user constituencies or attacked directly.

The cost of exercising control can be immense and unrecompensed.

CERTs

Computer Emergency Response Teams provide the dedicated intelligence analysis, real-time forensic capabilities, and specialized channels of coordination that end-users cannot afford to each maintain individually.

CERTs are the primary site of national shared investment in common defense.

Law Enforcement

A coherent and comprehensive system of laws which prohibit cyber-offense, and effective and obvious enforcement of those laws, are the mechanism whereby domestic attacks are forestalled.

Ministry of Foreign Affairs


Diplomacy, on the part of the department of state or foreign affairs or its equivalent, is the mechanism whereby a nation forestalls attacks on the part of other state actors or forces under their control.

Military

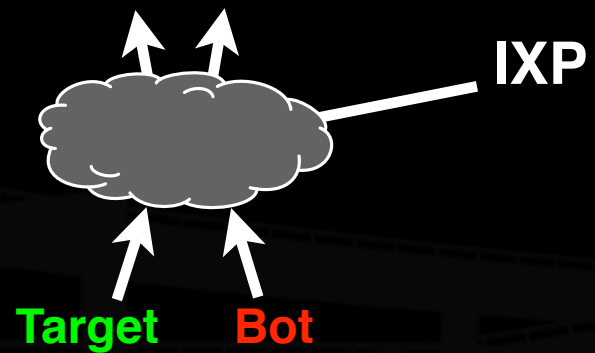
In countries other than China, militaries do not have privileged access to or control over the field, so militaries are reduced to solely offensive roles in cyber-conflict.

In this regard, they provide the credible threat behind diplomatic negotiation.

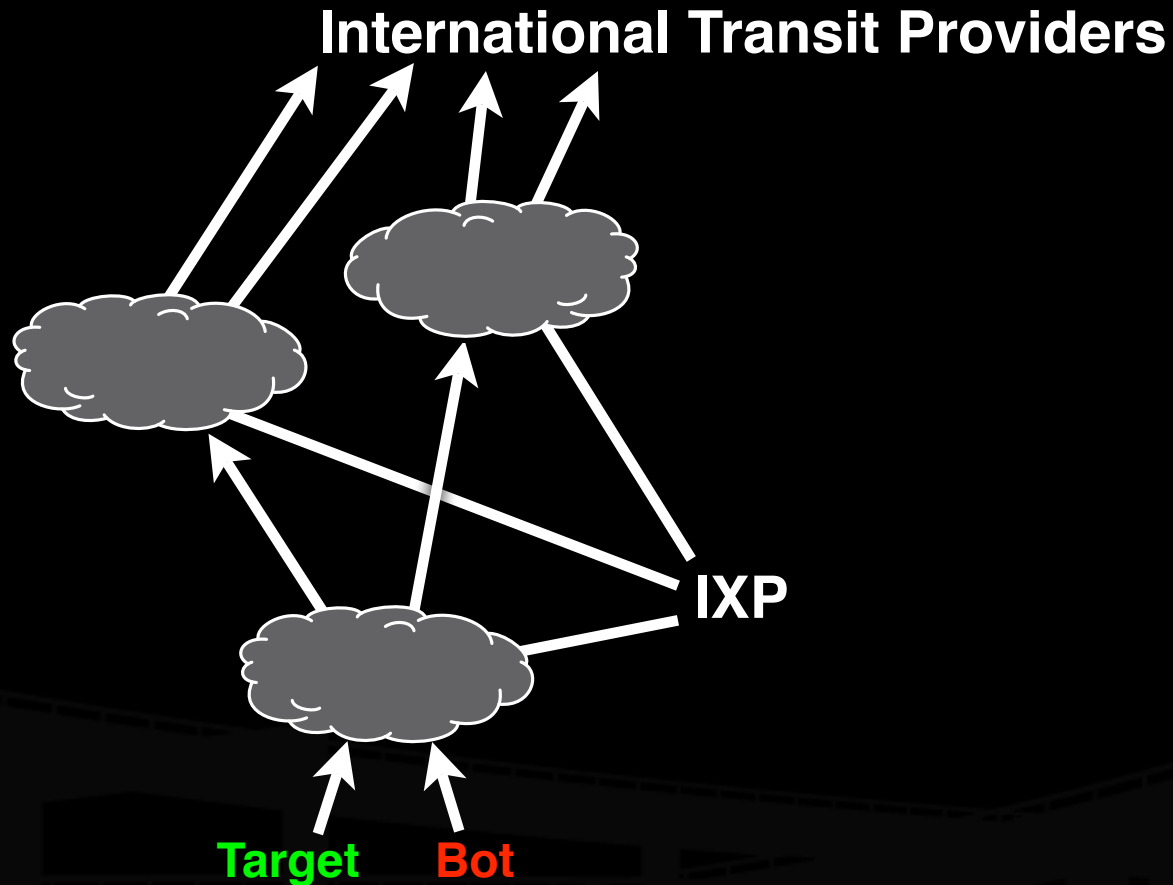
Topology of the Field



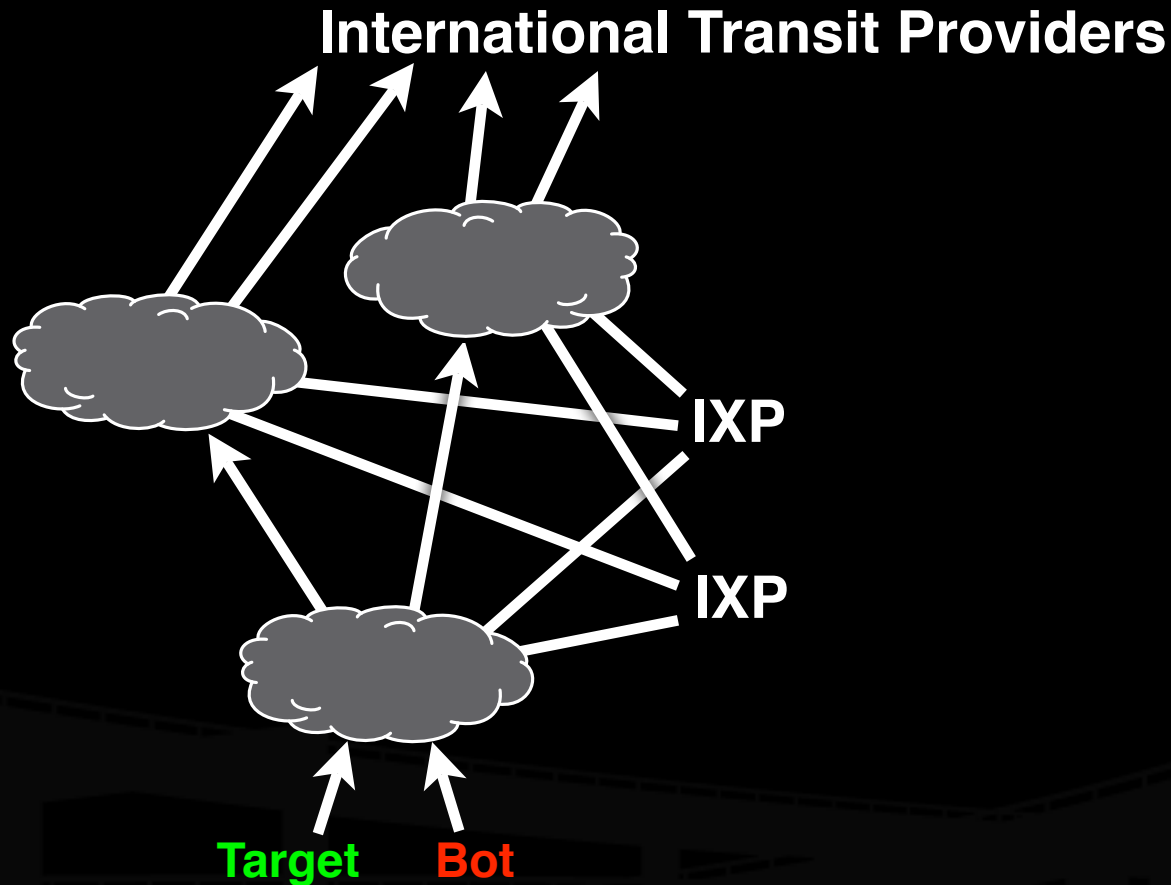
Topology of the Field



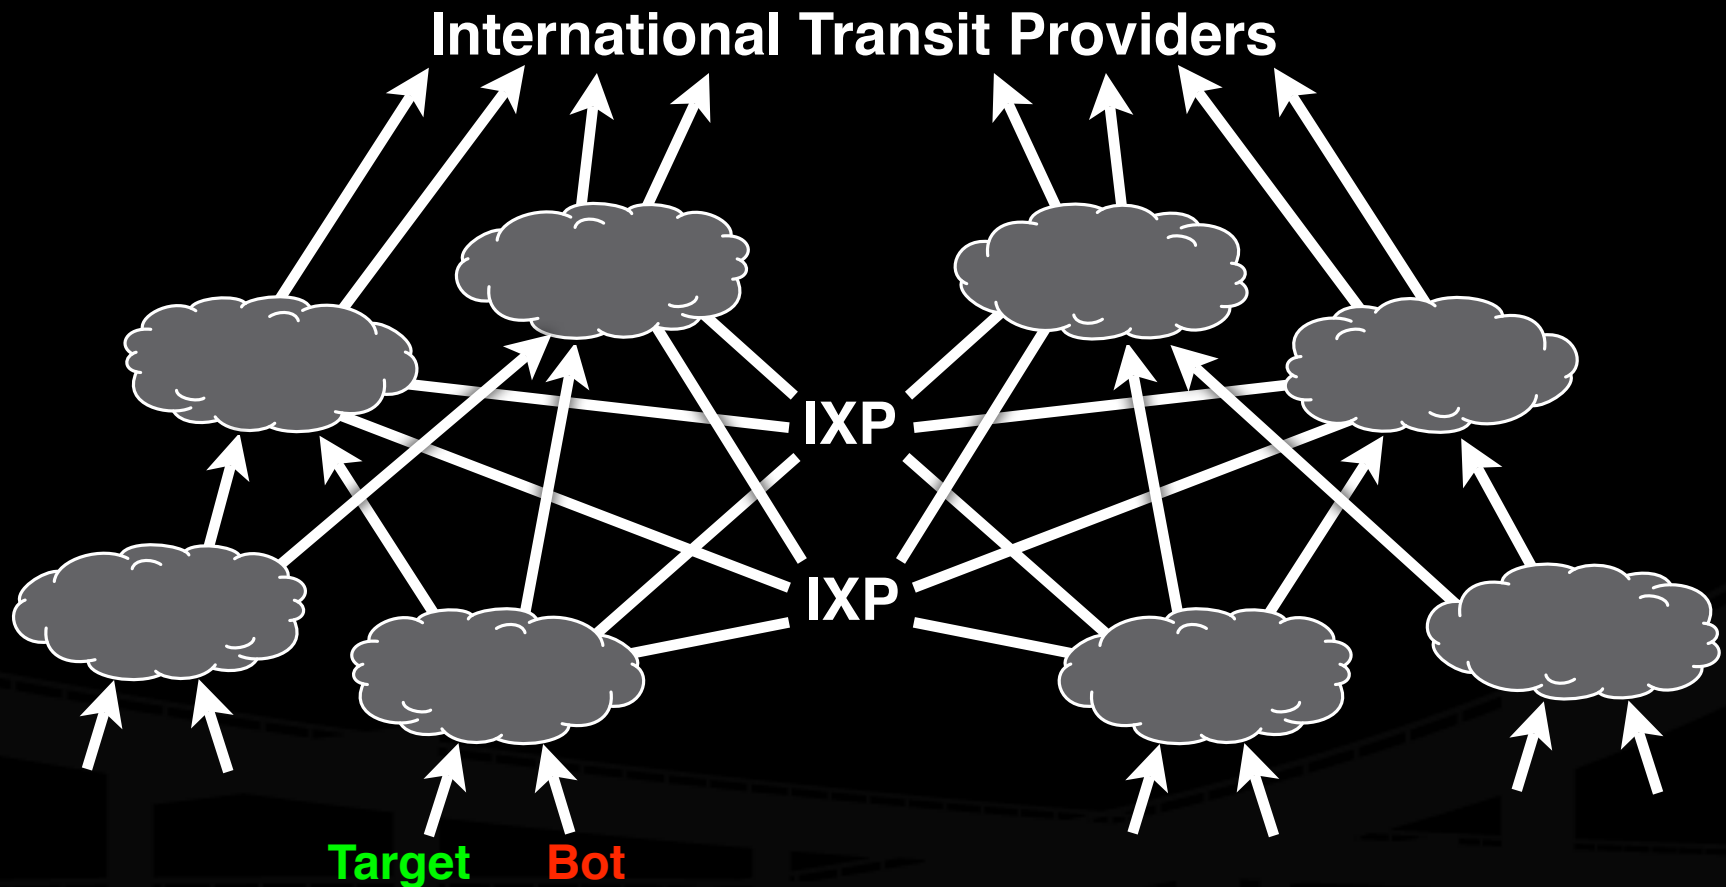
Topology of the Field



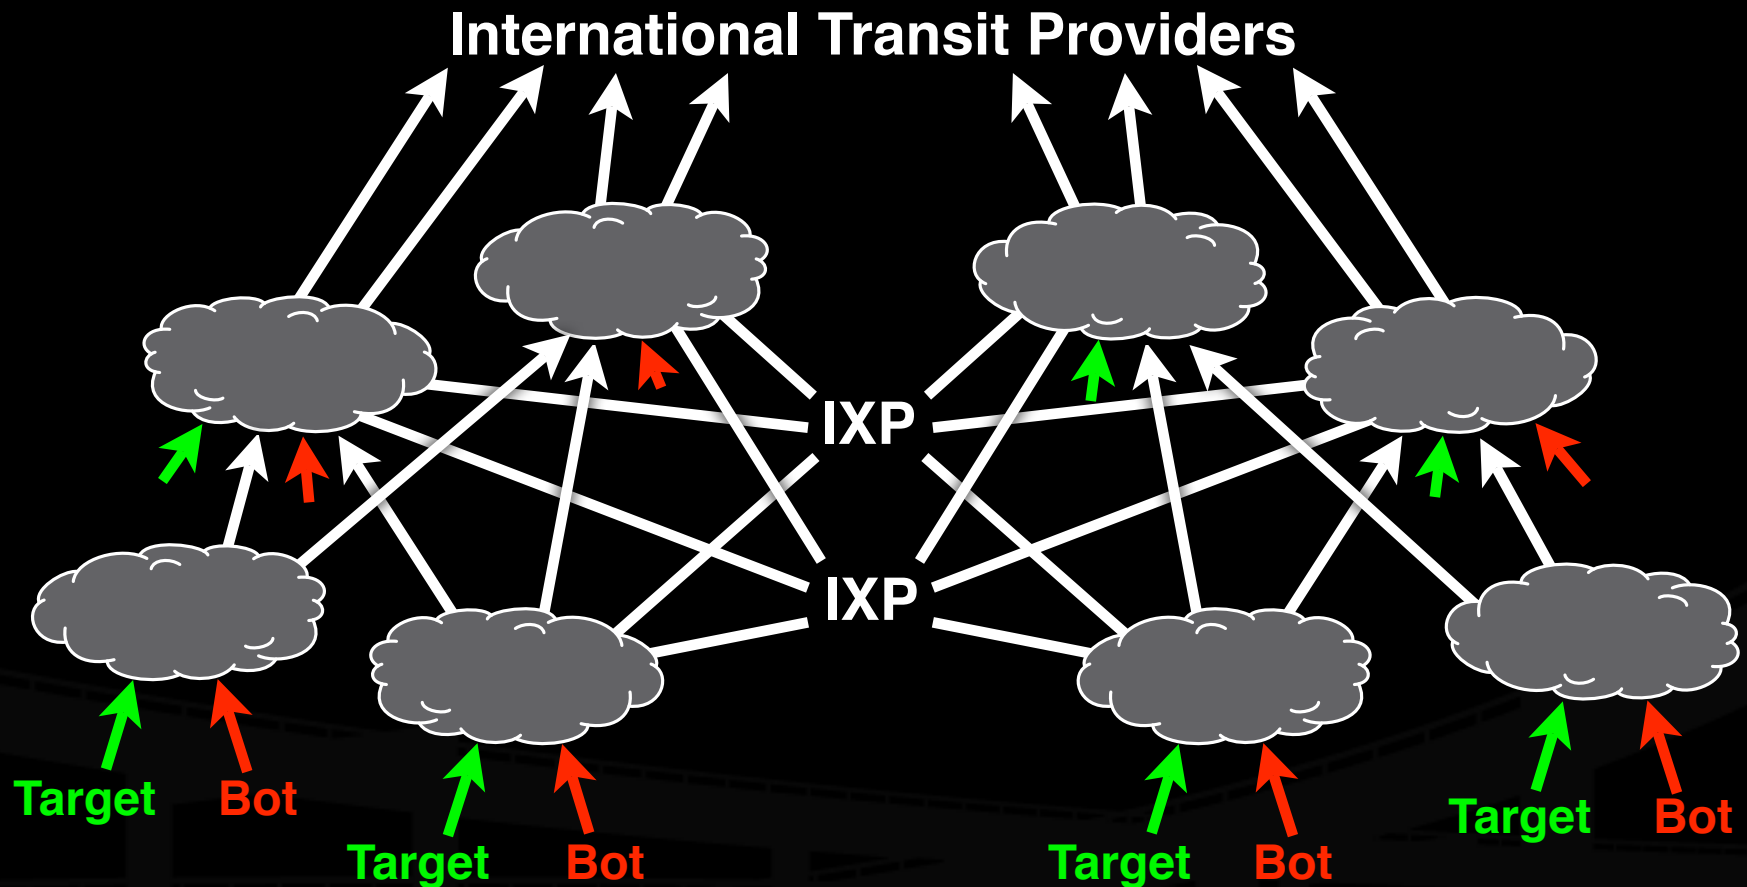
Topology of the Field




Topology of the Field




Topology of the Field




Topology of the Field




Topology of the Field




NSPs and Law Enforcement




NSPs and Law Enforcement




NSPs and CERT




NSPs and CERT



Diplomacy and International Law Enforcement



International Capacity





Impact on International Capacity

Approximately 4.1mpps at peak

Mitigated to 30kpps by 7am by diligent work through the night on the part of the CERT-EE and many NSPs.

Domestic Capacity



Impact on Domestic Capacity

None.

Entirely prevented by timely and decisive action by law enforcement.

Vizualization Tools in Use

MRTG / RRDTool

NfSen

Anomaly Detection Systems in Use

Cisco Detector (was Riverhead)

Panoptis

Arbor Peakflow SP

Narus Insight Manager

Lancope Stealthwatch XE

Q1 Labs Q1 Radar

Mitigation Systems in Use

Cisco Guard (was Riverhead)
Source-based uRPF filtering
Snort Inline

Thanks, and Questions?

Copies of this presentation can be had
in Keynote or PDF on request.

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