

A Curmudgeonly Operator's View of Resiliency and Research

DARPA-PI

2002.01.16 San Diego

randy@psg.com

<<http://psg.com/~randy/020116.darpa-pi.pdf>>

The Internet **DOES** Work

- IP forwarding is just as fast as switching, they're all just one-table lookups
- Actual measurements show QoS is just fine due to proper provisioning, and proper provisioning is cheaper than adding a layer
- QoS is a decision of which packets to drop. Do we want to be dropping packets?
- There are reasons the Internet has taken over the data world and threatens voice
- Now we have to treat it as a mature system

The Internet **DOES** Work (cntd)

- Reliability and Resiliency are core strengths, the Internet was designed for them
- But they're being stretched
- IPv4 addresses keep going and going and ...
- BGP may have over a decade of life left
- Operators are used to keeping the net up
- It is not time to panic
- But it is time for prudent research and engineering

But Secretly, We're Scared to Death!

- DDoS attacks
- Routing attacks
- Other infrastructure attacks
- But we are **not** scared of
 - running out of address space or ASs
 - BGP collapse
- I am scared by complexity
 - MPLS 2547 VPNs
 - QoS

The Only Real Problem is

Scaling

All the others inherit from that one

If you can scale, everything else must be working.

-- Mike O'Dell - Chief Technologist UUNET

Complexity: the Arch-Enemy of Scaling

- The telephants glorify complexity
- But look what it has done to OpEx
- RFC 1925 section 2(3), "With sufficient thrust, pigs fly just fine."
- This does not mean we can afford the fuel costs

Where the Smarts are

- Traditional Voice has stupid edge devices, phone instruments, and a very smart core
- The Internet has smart edges, computers with operating systems, applications, ..., and a simple stupid core, which just does packet forwarding
- Adding an entirely new Internet service is just a matter of distributing an application to a few consenting desktops (until NATs)
- Compare that to adding a service to Voice

Where the Reliability is

- The Voice network has smart central organs which are heavily armored, have rooms of battery backup, etc.
- The Internet **assumes** major component **failure** and achieves reliability through redundancy in the protocol designs
- I.e. 2/3 of the DNS root servers could die and no one would notice except for loading
- The protocols simply find a working one and remember it until it fails

But MO is Wrong, Scaling is **not** the Only Problem

- Security
 - the internet was not designed for it
 - the world has changed - crackers
- Resiliency
 - the real world has changed too - Sep 11
 - and it's not a research network any more
- Reliable equipment == less features
- And basically we do not understand the dynamics of the network <blush>

What Operators Need from Researchers?

- Understanding of IP dynamics
- Understanding of protocol dynamics
- Understanding scaling issues
- Help with security architectures

What Operators Don't Need

- Pages of numbers with not one insight
- Pretty visualizations with no insight
- Sensationalist FUD about routing, DNS, address space, ...
- Protocol designs which are based on overly static DNS or routing models
- Complexity
- Useless features that give us reduced reliability

Insight into IP Dynamics

- TCP flows & Congestion
- Different types of flows & congestion
 - http
 - voip
 - streaming
- Application protocol designs: good/bad
- Recognizing bad flows
 - DoS & DDoS
 - flash events

Insight into Protocol Dynamics

- How is routing really performing?
- How much benefit if the redundant BGP announcements were eliminated?
- What redundant announcements are due to implementation and which to protocol?
- What is the performance and reliability of DNS protocol? DNS data?
- If DNS data were better, how is scaling?

Understanding Growth

- How long will I Pv4 address space last?
- What is the lifetime of BGP?
 - what if implementations are cleaned up?
 - and/or does the protocol need small hacks?
 - and/or are operational changes needed?
- Analysis of anomalies

Help with Security

- DDoS prevention, not just tracking (Push-Back is the only proposed **active** defense)
- Prevention/detection of routing attacks
- Protection of critical services
 - DNS & root/tld servers
 - address & routing registries

Ongoing Work - Reaching out toward Researchers

- We are working on a proof that operating the internet is NP hard
- We strongly suspect that we can operate an approximate internet in polynomial time and dollars