Can the End-to-End Principle Survive?

NANOG, Feb 2004, Miami

Phil Karn

karn@qualcomm.com

Disclaimer

This talk includes my personal opinions.

I am not speaking for Qualcomm.

Qualcomm may or may not agree with me.

(But they should.)

Intro

- This talk is adapted and extended from my Mobicom '99 & AUUG 2003 talks
- Very little has changed
  - IPv6 and IPsec more widely deployed
  - worms, viruses, spam much worse

The End-to-End (E2E) Principle

- Seminal 1981 Saltzer, Reed & Clark paper: *End-to-End Arguments in System Design*
  - *IMHO, the most important network paper ever written*
- Many functions in a computer system are best done on an end-to-end basis
- A function can *sometimes* be justified at a lower layer as a performance enhancement
  - e.g., link level acks on a radio channel

Some Natural E2E Functions

- Reliability
  - end-to-end check still required even if subnet provides per-hop acks
- Security
  - end-to-end encryption protects the entire path
  - per-hop encryption can thwart traffic analysis
- Mobility
  - more flexible and efficient at application layer

Origins of the E2E Principle

- In the mid 1970s, the microprocessor created diseconomies of scale in computing that would clearly only grow
  - *distributed* became a buzzword
- The telephone system: unduly monolithic, complex, inflexible and expensive
  - precisely because it did too much; telcos still haven't learned from AIN fiasco
  - VoIP will be sweet revenge...someday...

E2E and the Internet

- The Internet architecture was originally conceived, designed, built, operated, tested and actually used *by the same people*, who were sponsored by other prospective end-users (the DoD)
  - "Every good work of software starts by scratching a developer's personal itch" (Eric S. Raymond)

Bogus Arguments Against E2E

- "How will we bill for our service?"
  - e.g., in VoIP; persistent "free Internet" myth
- "No *real* person will *ever* want/need to ----
  - run a server
  - have a home LAN
  - use the Internet
  - own a computer
  - (your excuse here – I've heard them all)

Some Real Threats to E2E

- IP address space exhaustion
  - more specifically, kludges like NAT
- Pervasive host security problems
  - thanks, Microsoft!
  - firewalls: packet filters, proxies, gateways, spam & virus filters, etc

More Threats

- Misguided performance concerns
  - "ack-spoofing" gateways (e.g., TCP over sat)
  - "lightweight" protocols
    - e.g., WAP, Unwired Planet (R.I.P)
- New layers on existing E2E mechanisms
  - no true E2E check in relayed email; TCP becomes by-hop between relays

A More Ominous Threat to E2E

- Carriers creeping up the stack
  - controlling address and name spaces
    - Cable modem, DSL providers charging for extra or static IP addresses "because they can"
  - restricting/modifying content
    - Port blocking
    - inserting ads, censoring content, transparent proxies
- Raw pipes aren't glamorous enough
- US regulation of wire owners has failed

Even More Ominous Threats

- Legal persecution of P2P networks
  - The Internet was *designed* to be "P2P"!
- Spam, worms, viruses, DoS attacks
  - endemic and rapidly getting worse
- Used to justify all sorts of anti-E2E violence:
  - outbound port 25 blocking
  - MAPS DUL blocking
  - AUP server prohibitions
  - mandatory spam/virus filters

The Real Issue

Who's in charge here? The end-user or the carriers?

Defending E2E

- Tunneling (e.g., IPv6 6to4 and IPsec)
  - Encryption nicely thwarts content restrictions
- QoS support
  - cleaner way to differentiate service offerings
  - a rare low-level feature that *should* exist, but doesn't
- Open source software
  - powerful way to meet *users'* (vs vendors') needs

IPv6

- 6to4 is excellent for NAT avoidance
  - will become very popular when implemented in consumer routers
- Hosts will be dual stack (many already are)
  - Non-global IPv4 address behind NAT
    - common current practice
    - fine for existing web & email clients
  - Global IPv6 address in 6to4 block
    - ideal for new P2P applications, e.g., VoIP

IPv6, contd

- Biggest myth about IPv6: “We can't use it until our carriers support it”.
- With 6to4 tunneling, you only need IPv6 support at the endpoints
  - and most already have it (XP, Linux, OS X)
- I actually *like that my carriers don't do IPv6*
  - *when they do, they'll arbitrarily filter, redirect, block, charge, spindle and mutilate IPv6 as they now do IPv4*

Will IPv6 succeed?

- Who will "own" the v6 address space?
  - Many complaints already about cost & difficulty of getting v6 address space
    - implicit /48 assignment big advantage of 6to4
- Requires host, app and router upgrades
  - already in most host OSes
  - older applications don't have to have it
  - not yet in consumer-grade gateways
- Will worsen host security problems

Security Threats

- Many distinct security problems, e.g.,
  - Spam
  - Worms/viruses
  - DoS attacks
- *Different resources being attacked*
  - *User eyeballs*
  - *Host resources*
  - *Network resources*
- *IMHO, Biggest single threat to E2E*

Preserving E2E Against Security Threats

- *Security Placement Principle: Place security mechanisms as close as possible to the resources being protected*
- *Ergo,*
  - *must* distinguish between host and network attacks
  - host attacks best prevented by host mechanisms
    - with net mechanisms as performance enhancement
  - net attacks only prevented by net mechanisms

Security Philosophy

- Humans are more valuable than machines
  - Primary goal of spam blocking is to save *my time; the network is secondary*
- End users *must* retain ultimate control
  - any filtering functions performed by ISPs as performance enhancement *must* be under end-user control

Our #1 Security Problem: Microsoft!

- Two kinds of worms & viruses endemic:
  - trojans (e.g., SoBig.F)
  - bug exploiters (e.g., Slammer, Blaster)
- Primary damage to infected hosts, but ubiquity clogs network
- Despite many promises, problem worsening rapidly

Fun & Games with SoBig.F

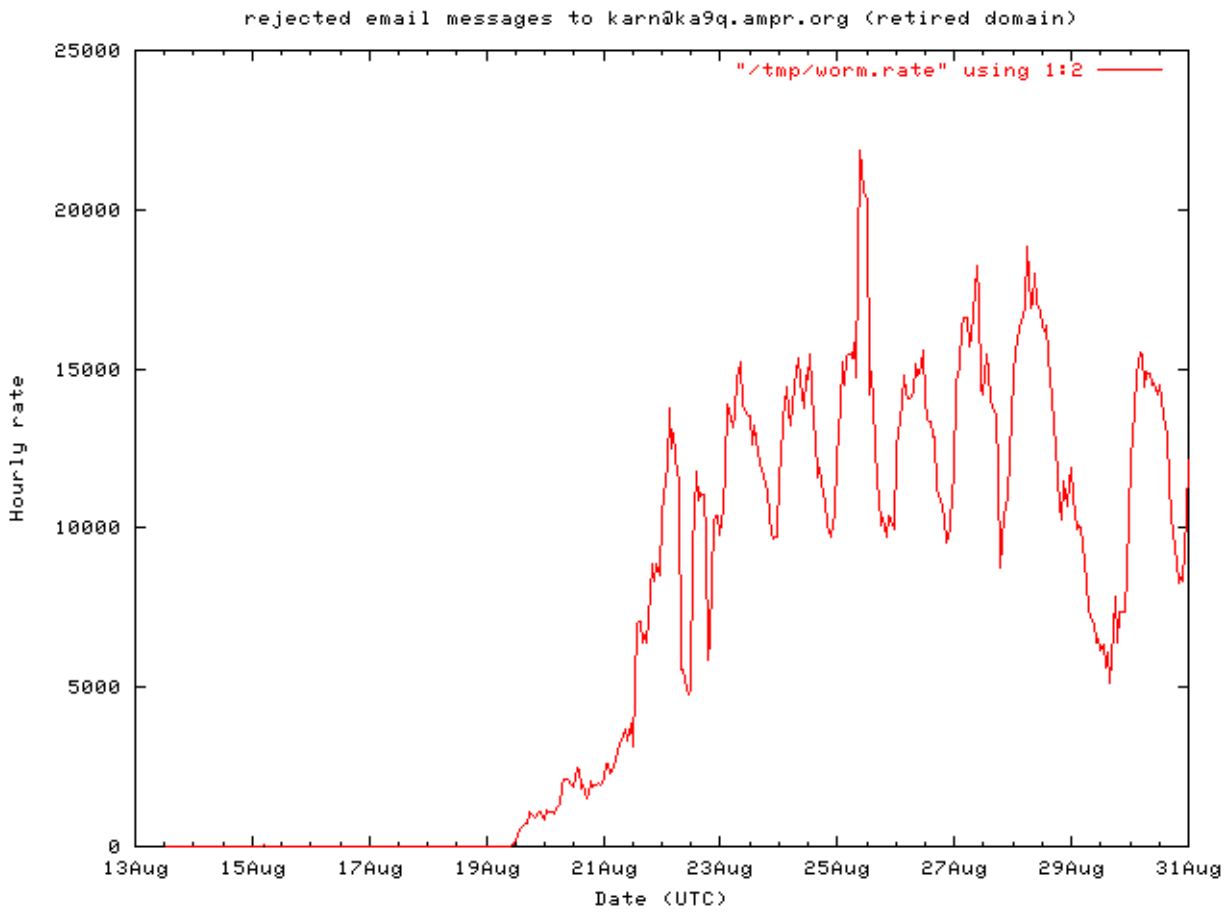

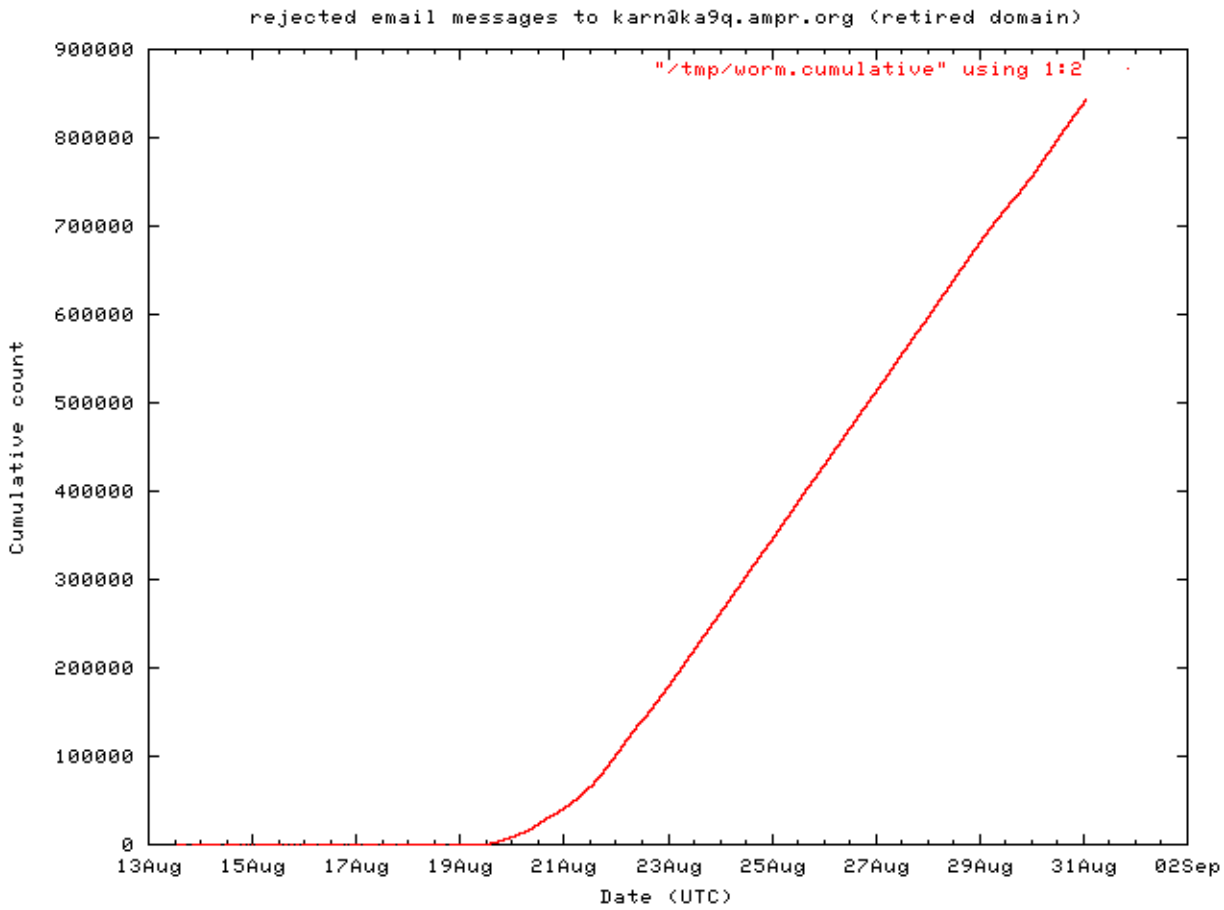
rejected email messages to karn@ka9q.ampr.org (retired domain)
"/tmp/worm.cumulative" using 1:2
Cumulative count
900000
800000
700000
600000
500000
400000
300000
200000
100000
0
13Aug
15Aug
17Aug
19Aug
21Aug
23Aug
25Aug
27Aug
29Aug
31Aug
02Sep
Date (UTC)

Example:
DoS in Cellular Networks

- We'd like to give every phone a global IPv6 address and make it a server (VoIP, text etc)
  - any host anywhere can send it packets
- Wireless is inherently slower than wired
- Denial-of-service attacks would be too easy
  - already pandemic in the wired Internet
  - excess capacity keeps them from being more destructive than they already are

Blocking DoS Attacks

- Filters in the phone won't work
  - the damage is to the wireless link, not the phone
- I.e., filters have to be in the network
- This problem isn't unique to wireless hosts
  - they are simply the most vulnerable
  - we need a general solution for all hosts if IPv6 is to restore the end-to-end model

Blocking Spam

- Special class of denial-of-service attack
  - attacked resource is user's eyes, not his link
  - already a serious problem with SMS in some areas
- Many ISP spam “solutions” are much worse than the disease (*e.g. dialup IP blocking)*
- Best solution so far: Bayesian analysis, performed upstream under user control

Conclusions

- Secure host software is mandatory
  - but Microsoft is highly problematic
- Will still need filtering to protect the net
- Challenge is to preserve the E2E model
- Needed: standard filtering mechanisms under end-user control

References

- Saltzer, Reed, Clark: *End to End Arguments in System Design*
  - http://people.qualcomm.com/karn/library.html
- Raymond, *The Cathedral and the Bazaar*
  - http://www.tuxedo.org/~esr/writings/cathedral-bazaar/
- Karn, *Why I Hate Microsoft: Part 1, Worms and Viruses*
  - *http://www.ka9q.net/worm*

More References

- Andrew Odlyzko, “Content is not King”, http://www.firstmonday.dk/issues/issue6_2/odlyzko/