

The End To End Internet

Harald Alvestrand

Cisco Fellow

IETF Chair

Hat Identification

- **Not speaking for Cisco (I don't know what they are doing)**
- **Not speaking for the IETF (they don't know what I'm saying)**
- **Speaking on the basis of experience and thinking**
- **Speaking as myself**

The Original(?) Principle(s)


- **If you need to do something at the end of a connection, why bother with doing it in the middle too?**
- **If you don't attempt to do in the middle what you have to do at the end, the network is simpler, more open and better**
- **Having application state in the network is stupid – keep it at the endpoints**

Problem identification

- **”End” is a simple concept**
- **A process makes a connection, sends data, receives responses. The connection is end-to-end.**
- **Life isn't always that simple.....**

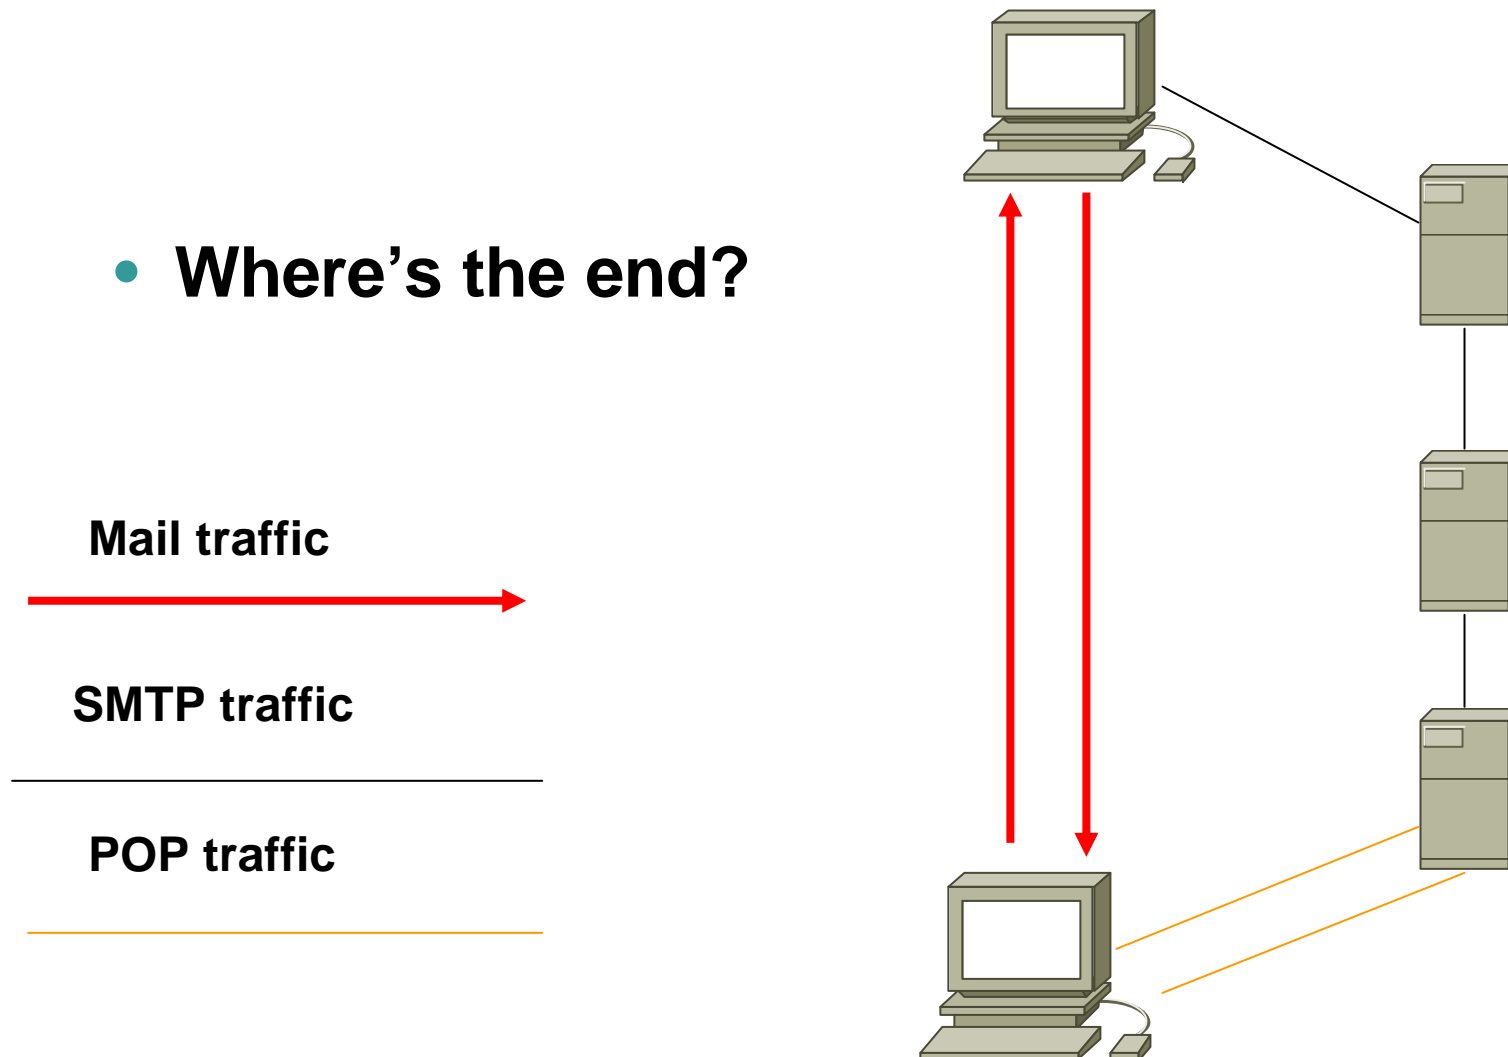
A more real-life example

- **To the outside: A single webserver**
- **On the inside: Duplicated everything, multilayer architectures.**
- **Where's the end?**



Another example: Mail

- Where's the end?



The two concepts of end-to-end (1)

Cisco.com

- **Technical: That which cannot be done well in the middle must be done at the edge.**

....functions placed at low levels of a system may be redundant or of little value when compared with the cost of providing them at that low level. (Chiappa)

The function in question can **completely and correctly** be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Saltzer/Reed/Clark)

The two concepts of end-to-end (2)

- **Moral: Network bad. Endpoints good; middlebox state bad; soft state good; stateless even better.**
- **This is a design choice. And it has served us very well. It is not a natural law.**

the intermediate packet switching nodes, or gateways, must not have any essential state information about on-going connections. Instead, they are stateless packet switches, a class of network design sometimes called a "datagram" network. (Clark95)

Digression: The Phone System Illusion

- **The "end" being an apparatus, not a person (earpiece, microphone and dial)**
- **End was "obviously" stupid. All "smarts" in the middle.**
- **A billing relationship implied by the physical wire and logical number plan**
- **Additional functionality being built on the trust relationship that was based on the wire**
- **Complexity (services, mobility) was added while keeping the basic model**
- **As a single model, it worked fairly well**

Phone Systems in Trouble

- **The PABX: Groups of endpoints????**
- **Interconnection: do you trust your competitor/customer/provider?**
 - **See MCI local-call case**
- **Intelligent endpoints using the phone network as dumb carrier: Dialup ISPs**
 - **Lots of work to stop the network from being too smart**
 - **”To disable call waiting::”**
- **Telco bypass: Voice over the Internet**
 - **Within and outside the E.164 dialling plan**
- **All signs that the marriage of functionality and connectivity was artificial and ill-funded.**

Does end-to-end always work?

- **Stupid Endpoints**
PDA's, mobile phones, appliances
- **Security imposed from the middle**
Firewalls, authenticating proxies
- **Performance issues at special points**
Wireless links, fast long-delay links, congestion points
- **Money/control issues**
Those who can deliver a more complex function can demand more money for it (both groups!)
- **NAT!**

Stupid Endpoints

- **”Lightweight” devices connecting to a network**
Can’t authenticate, can’t encrypt, can’t remember
- **”Solution”**: Tie them to a larger system
- **Force all communication through this intermediary**
- **Works for GSM signalling**
- **Functionality totally dependent on intermediary**
- **In many cases, example of ”distributed endpoint”**

Security in the middle

- **Network owners want control**
- **An open architecture loosens control**
 - Ethernet jacks**
 - Standard PC architectures**
 - Windows**
- **If you cannot control the endpoint – control a chokepoint**
- **Problem for security: Security!**
 - Encrypted sessions are a major problem**
 - End-driven key management is even more of a problem: No way to check what's going on!**

Performance issues

- **Long-delay pipes**
 - Sending (or pacing) ACKs can increase performance, without upgrading end systems**
 - Content distribution networks move content "closer" to the consumer than the originator can**
- **High-loss, low-bandwidth networks ("wireless")**
 - Content adapters ("HTML downgrade" ++)**
 - Defeat TCP's "all loss is congestion"**
 - (Un)fairness based on identity**
- **There are other solutions to most issues**
 - Diffserv/RSVP, Fast-TCP,**
- **Content adaption can be seen as "split endpoint"**
- **Lack of model clarity is one of the bigger dangers here**

Money/Control issues

- **Desire of service provider: Get more money.**
- **Desire of equipment provider: Get more money.**
- **Method often used: Sell more complexity.**
 - Service provider: Fancier services (in network)**
 - Diffserv, Provider VPNs, voice-over-IP, managed services....**
 - User side: Fancier services (at endpoint)**
 - User VPNs, user voice-over-Internet, firewalls, scanners....**
- **If offering is unique, and customer buys it, customer can't escape. Price no longer decides.**
- **This is, long-run, a false way**
 - Competition will take care of those who spend too much**
 - Useful services will be used more, expanding the market**
- **In the short run, it works.**

So are there answers?

- **Valid architectures are those that provide value to the end-user**

- **Some services require middlemen**

The DNS service is a middleman

So is the telephone numbering system

- **Open architectures encourage the pieces to be considered separately**

Cross-subsidy is not a stable model – see DNS, voice-over-IP services

Attempting to block services is Not Good

So what is an "end"?

- **Dependent on context**
- **Related to function**

"end" of reliable byte delivery is not the "end" of database transaction processing, which is not the "end" of an email conversation

Clark et al knew this in 1981. We sometimes forget – and our systems work well enough that we often get away with it.

- **Needs clear thinking to identify**

Trust and the Principle

- **End users have to extend trust in order to get work done**
- **Middle-box security tries to force the issue, but can't know what the end-users are doing**
- **The Right Way is for the ends to extend trust, and tell the middle that they are doing so**
 - With the required authorization**
- **This is the end-to-end principle again!**

What does an "end" look like?

- **Commonly under a single administration**
 - Or seen by others as a single entity**
 - This is core to keeping innovation running**
- **Contains all the parts needed to perform its part of an useful communication function**
 - May have internal structure**
 - May rely on other services**
 - Does not force others to know about this structure**
- **Depends on your layer of abstraction!**

The End of the End Is Not Near

Cisco.com

- **We build services at many different levels**
- **We need to be aware of the "ends" we create in making those services**
- **The end-to-end principle is a design tool**
- **Do the right thing. Once. In the end.**

References

- **Noel Chiappa: “Would the real end-to-end principle please stand up?”**
http://users.exis.net/~jnc/tech/end_end.html
- **Saltzer/Reed/Clark paper (1981)**
<http://www.reed.com/Papers/EndtoEnd.html>
- **Dave Clark: “Design principles for the Darpa” (1995)**
<http://www.acm.org/sigcomm/ccr/archive/1995/jan95/ccr-9501-clark.pdf>
- **James Kempf and Rob Austein: “The Rise of the Middle and the Future of End to End: Reflections on the Evolution of the Internet Architecture” (2003)**
draft-iab-e2e-futures-03.txt