Signpost: Trusted, Effectful Internet names

Jon Crowcroft from original slides by Anil Madhavapeddy,
University of Cambridge
Keynote for U-Net@ICC, Ottawa, 11.6.2012

BACKGROUND: 1980

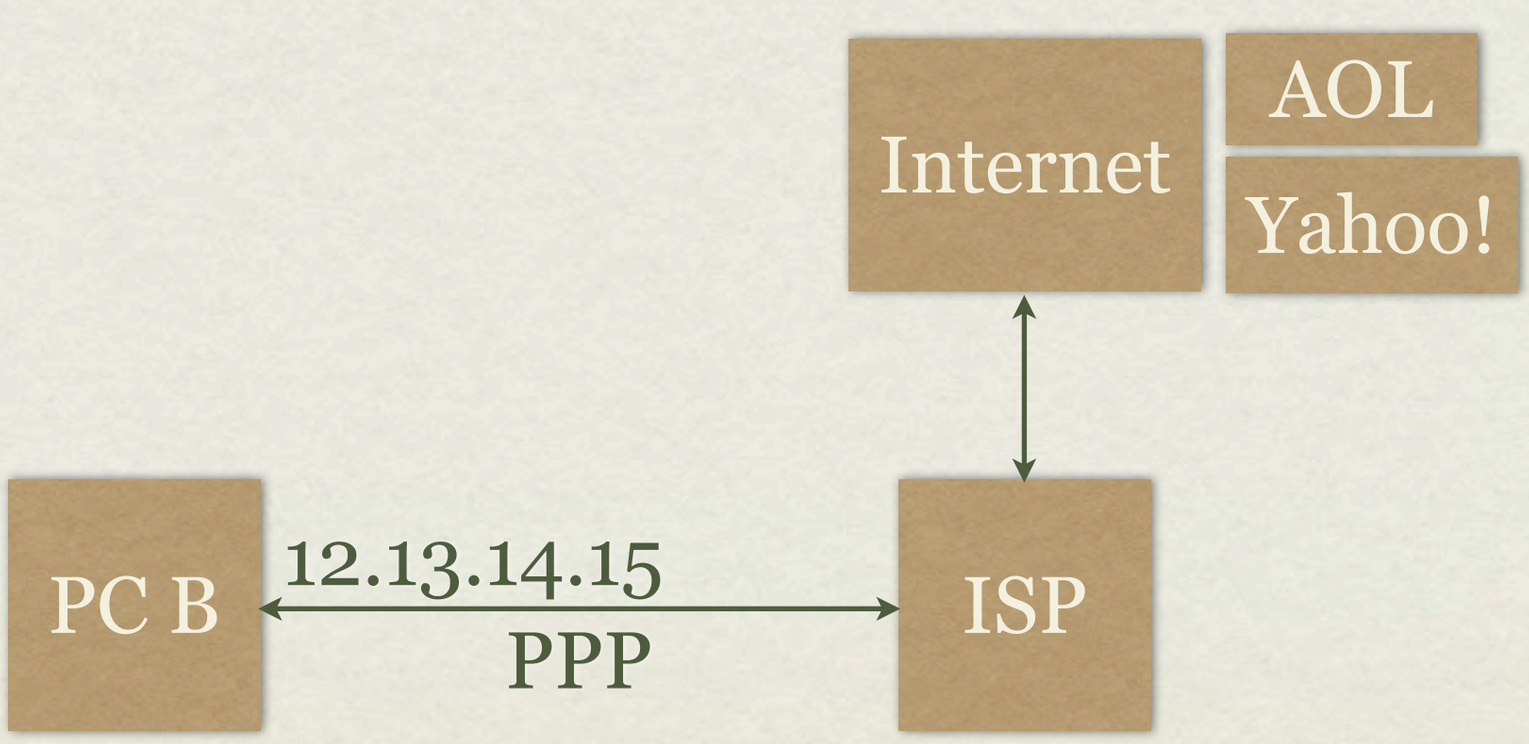

BACKGROUND: 1990

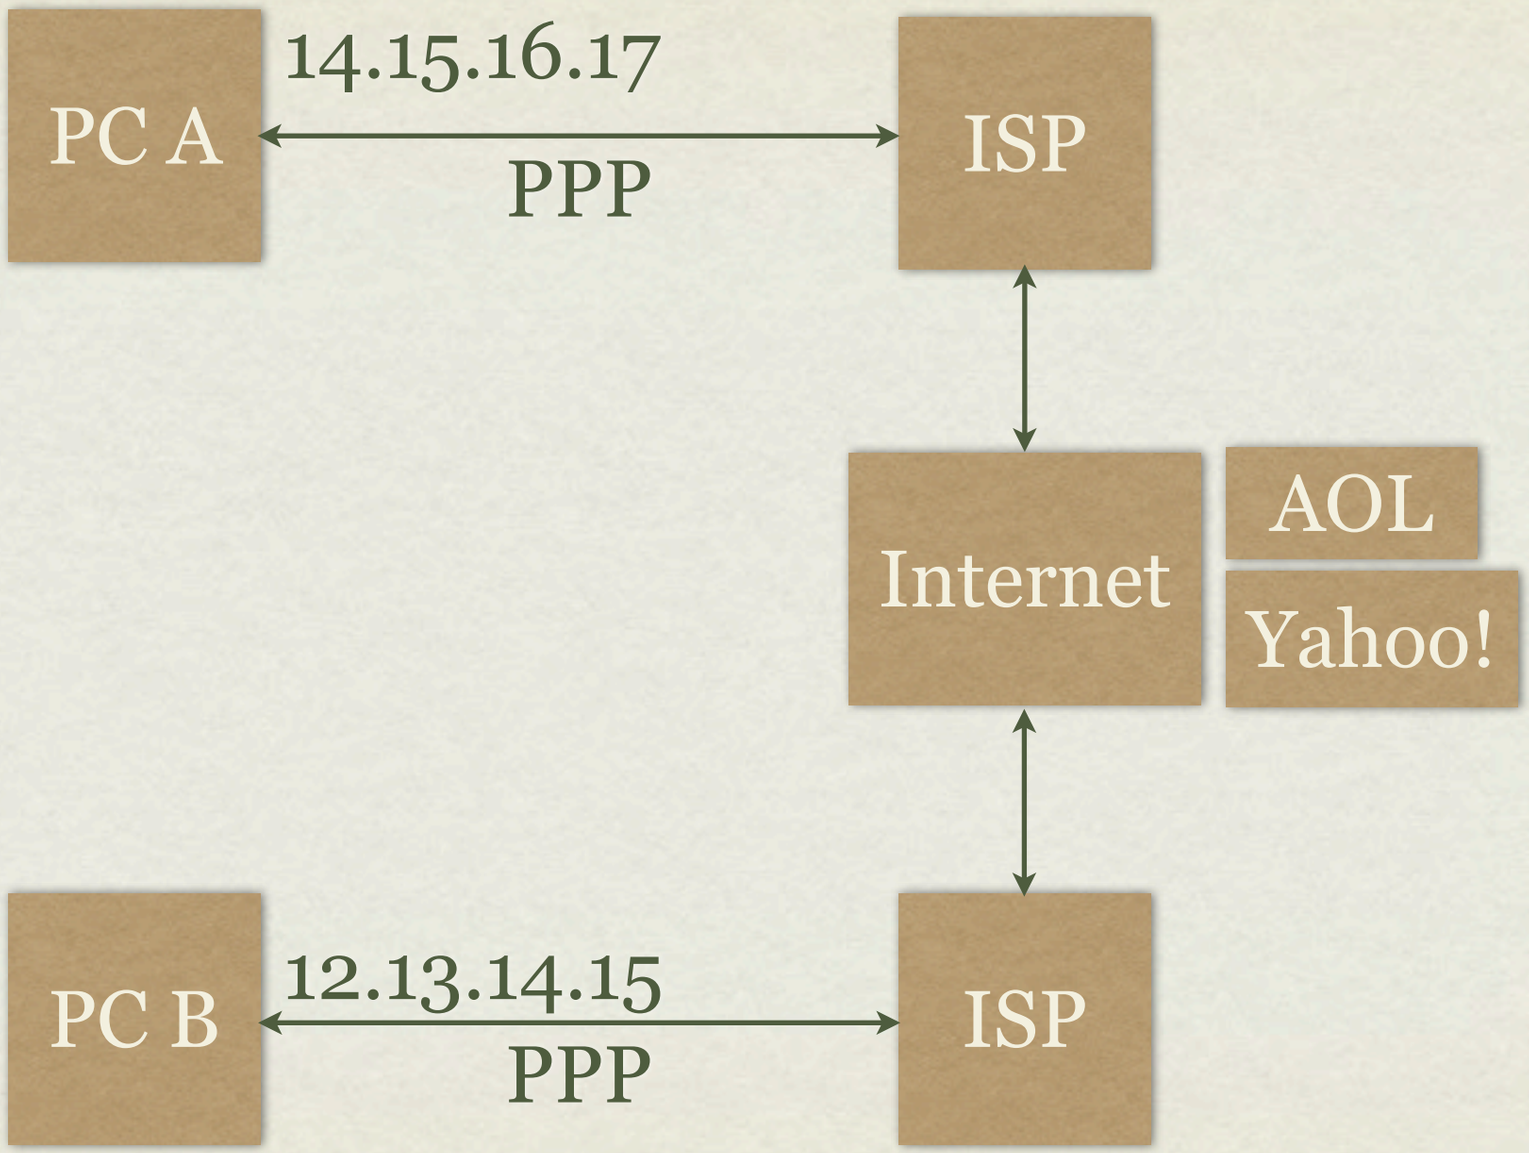

BACKGROUND: 2000

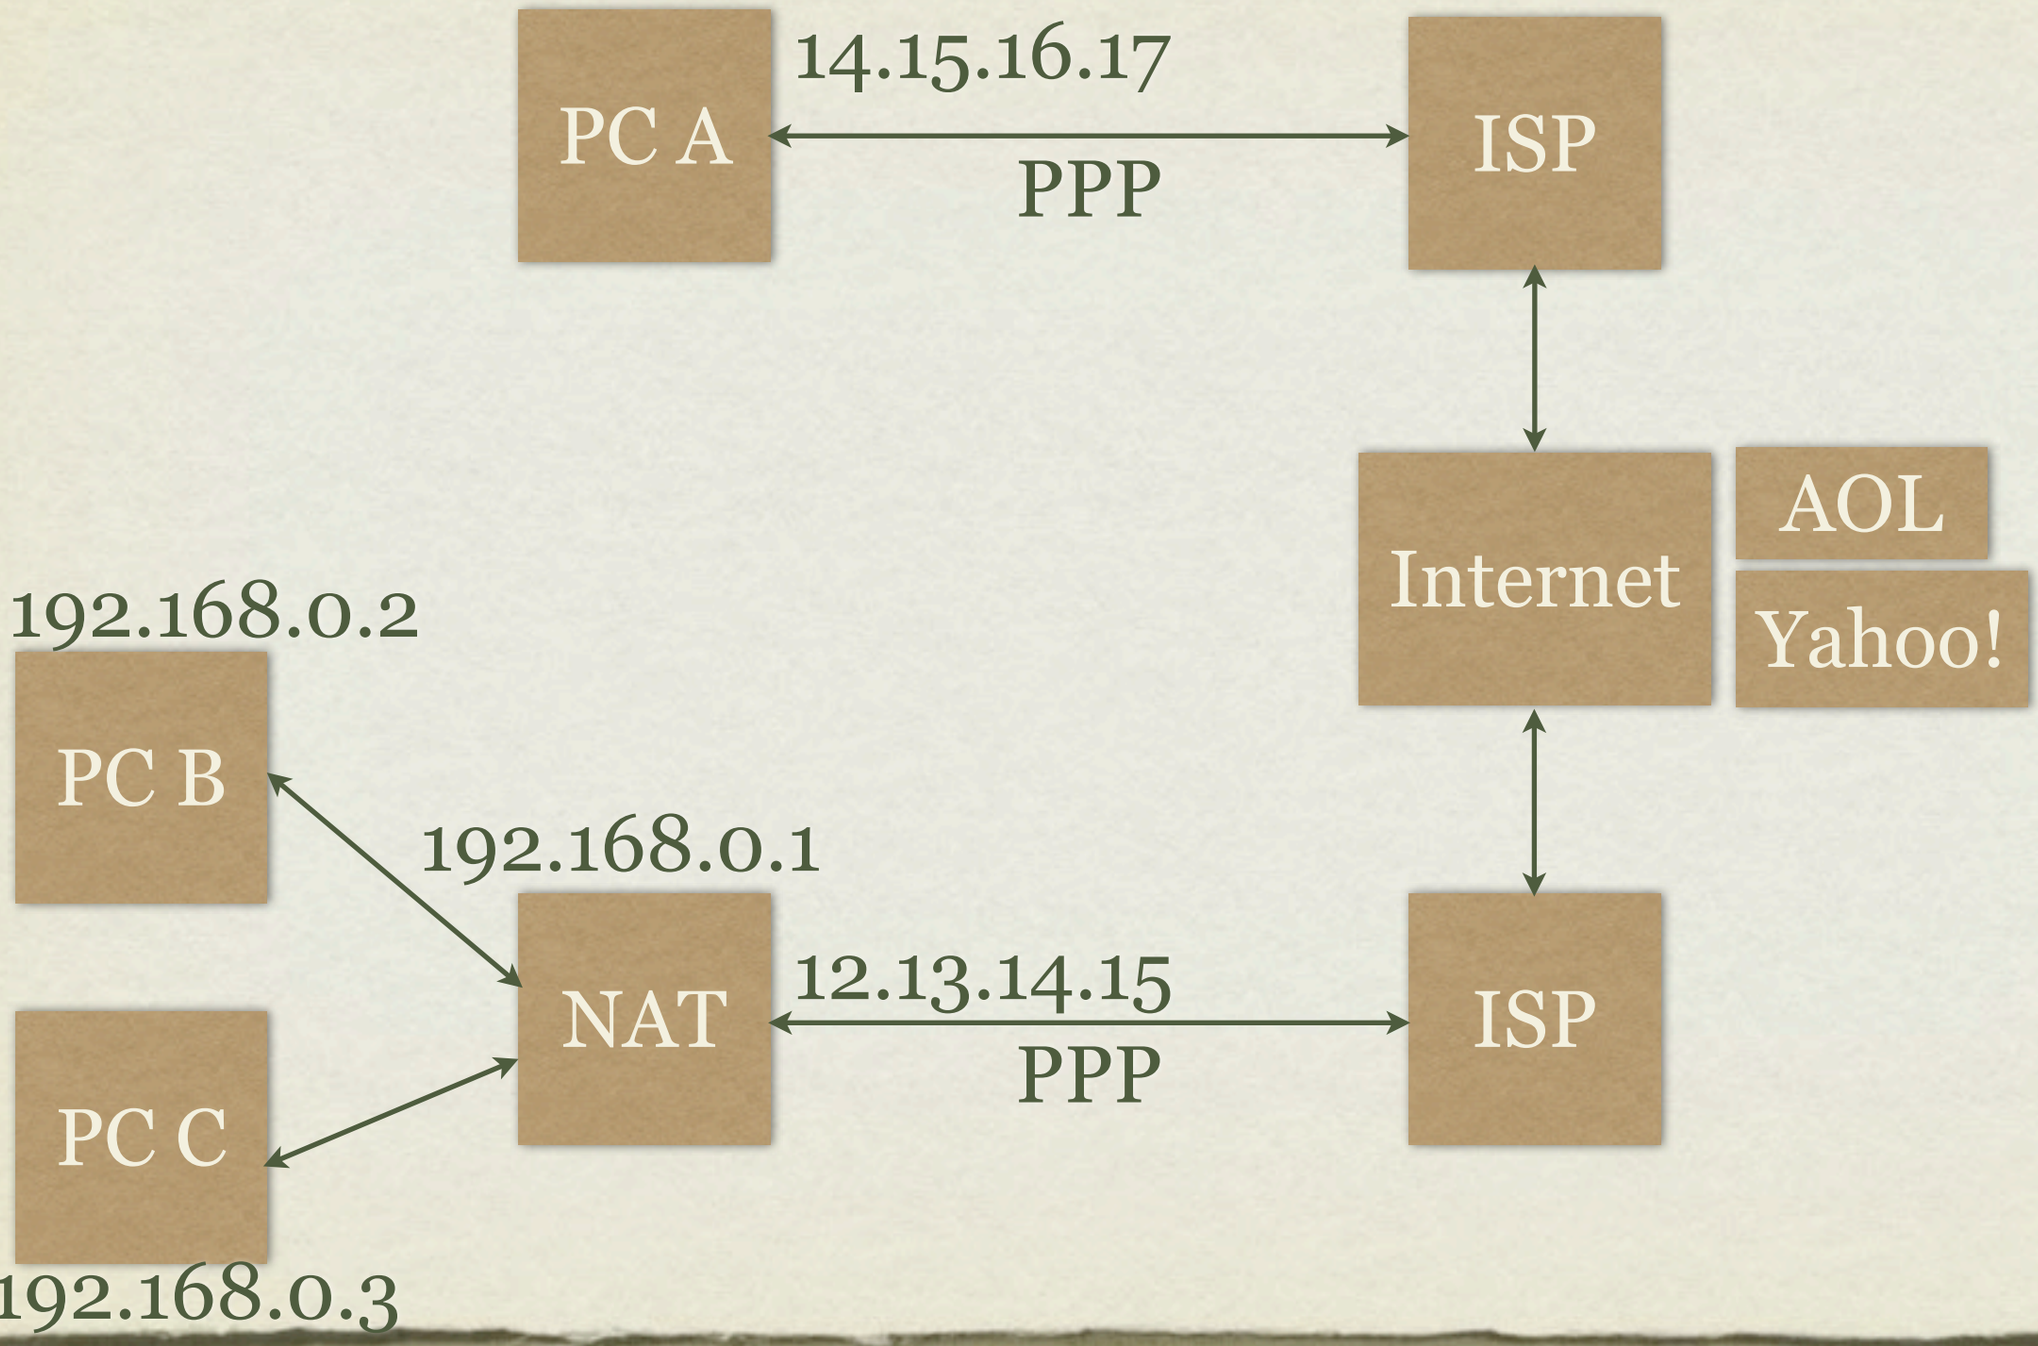

BACKGROUND: 2010

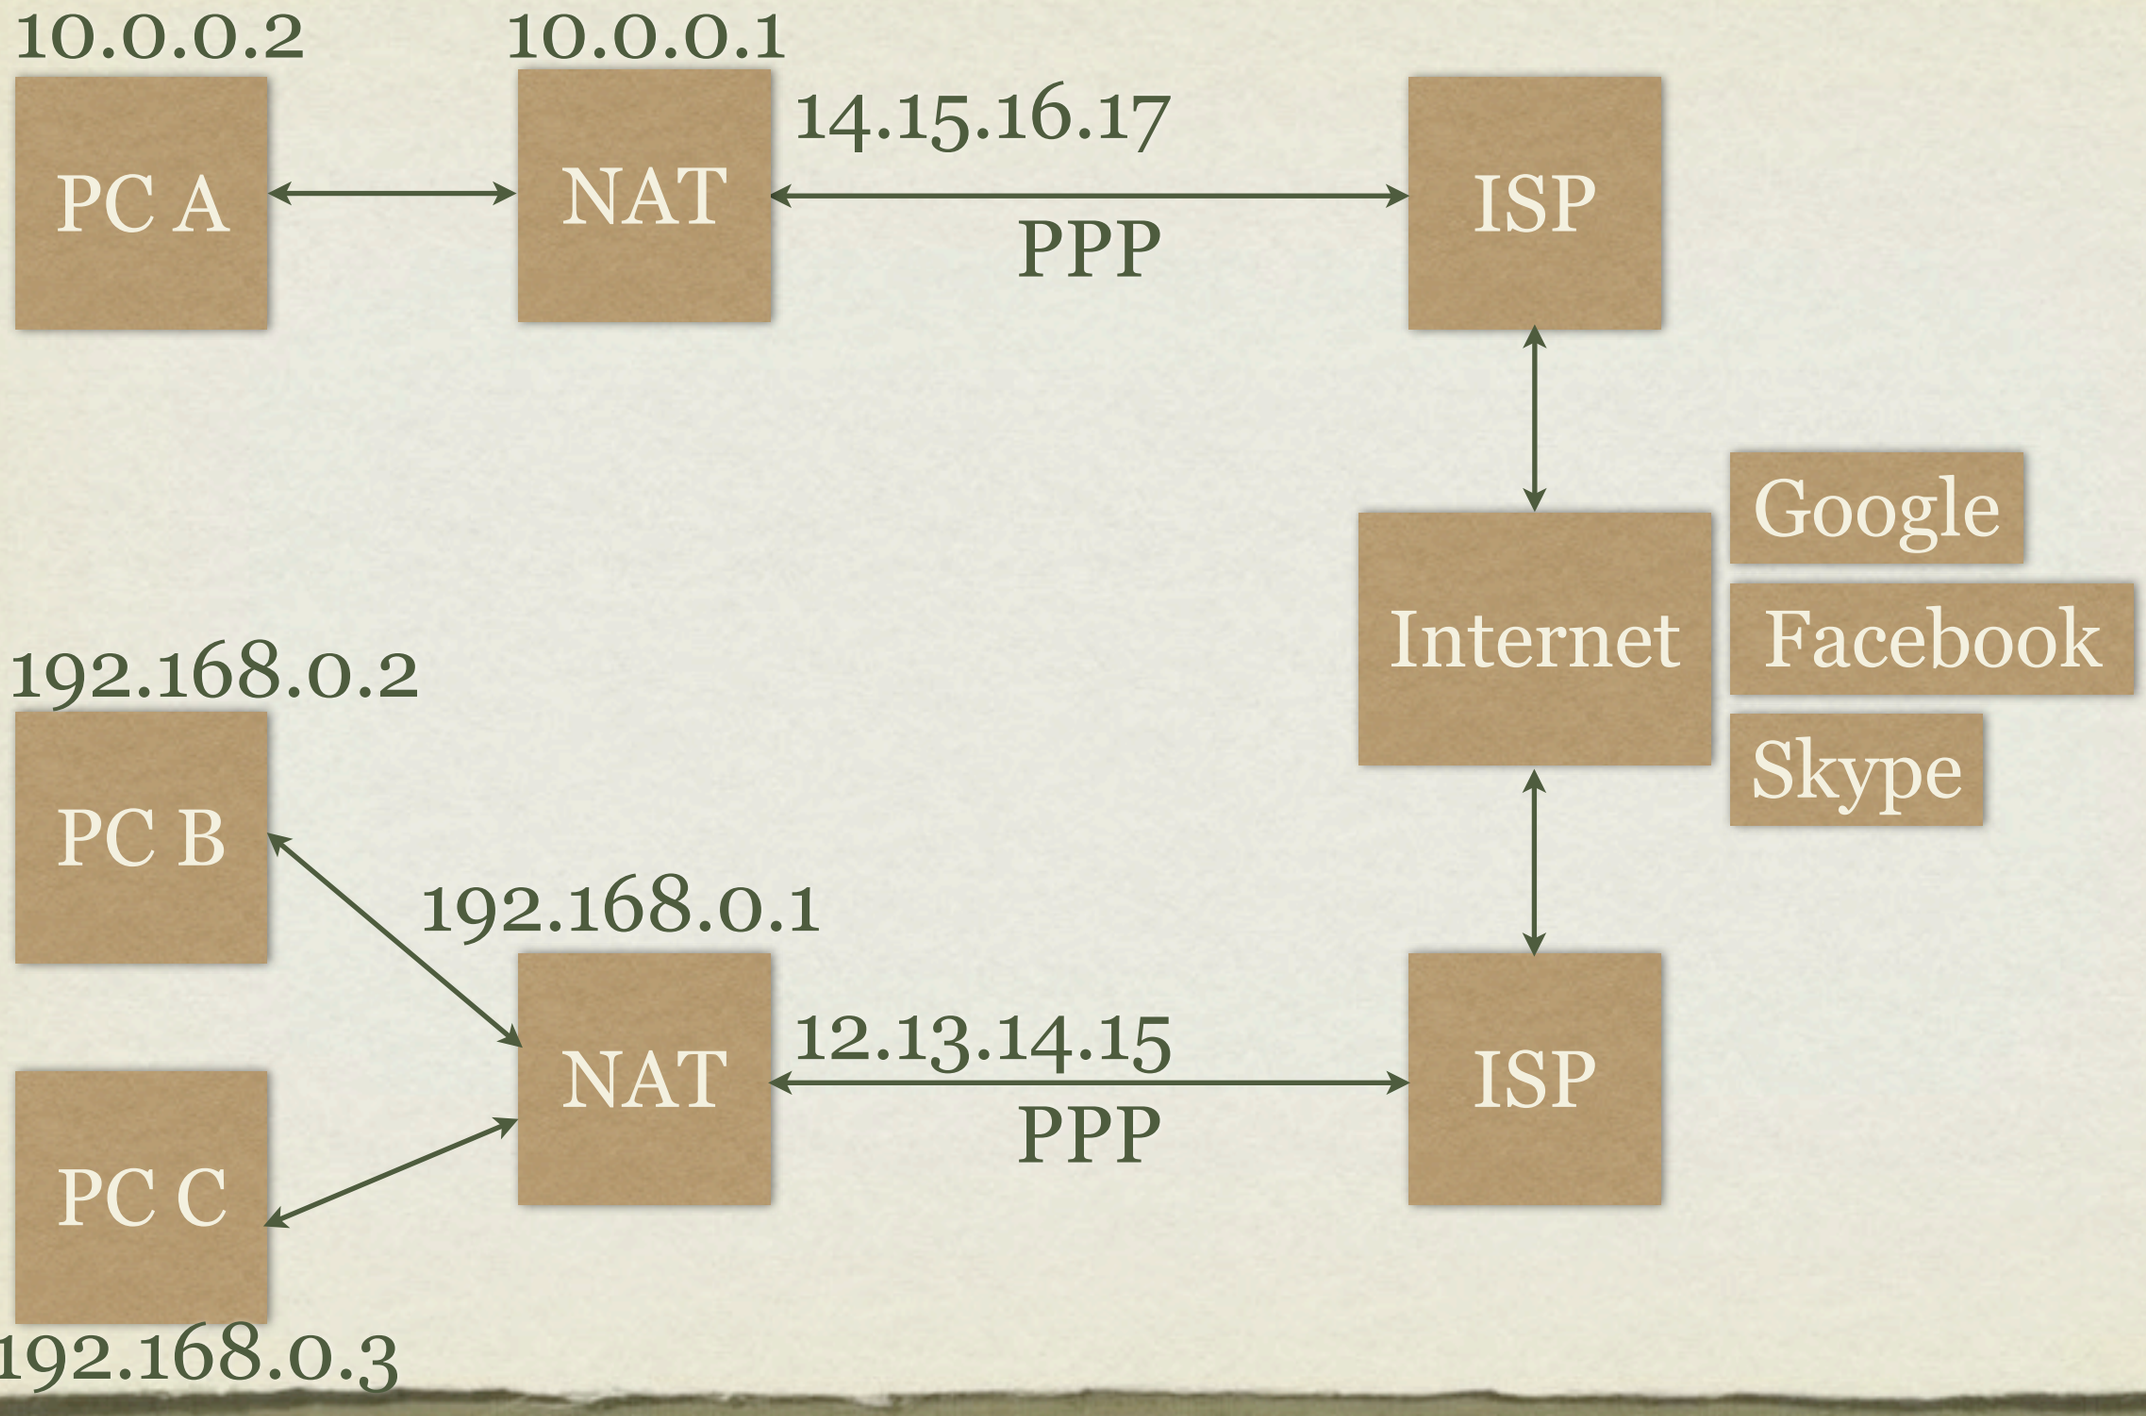

BACKGROUND: 2012

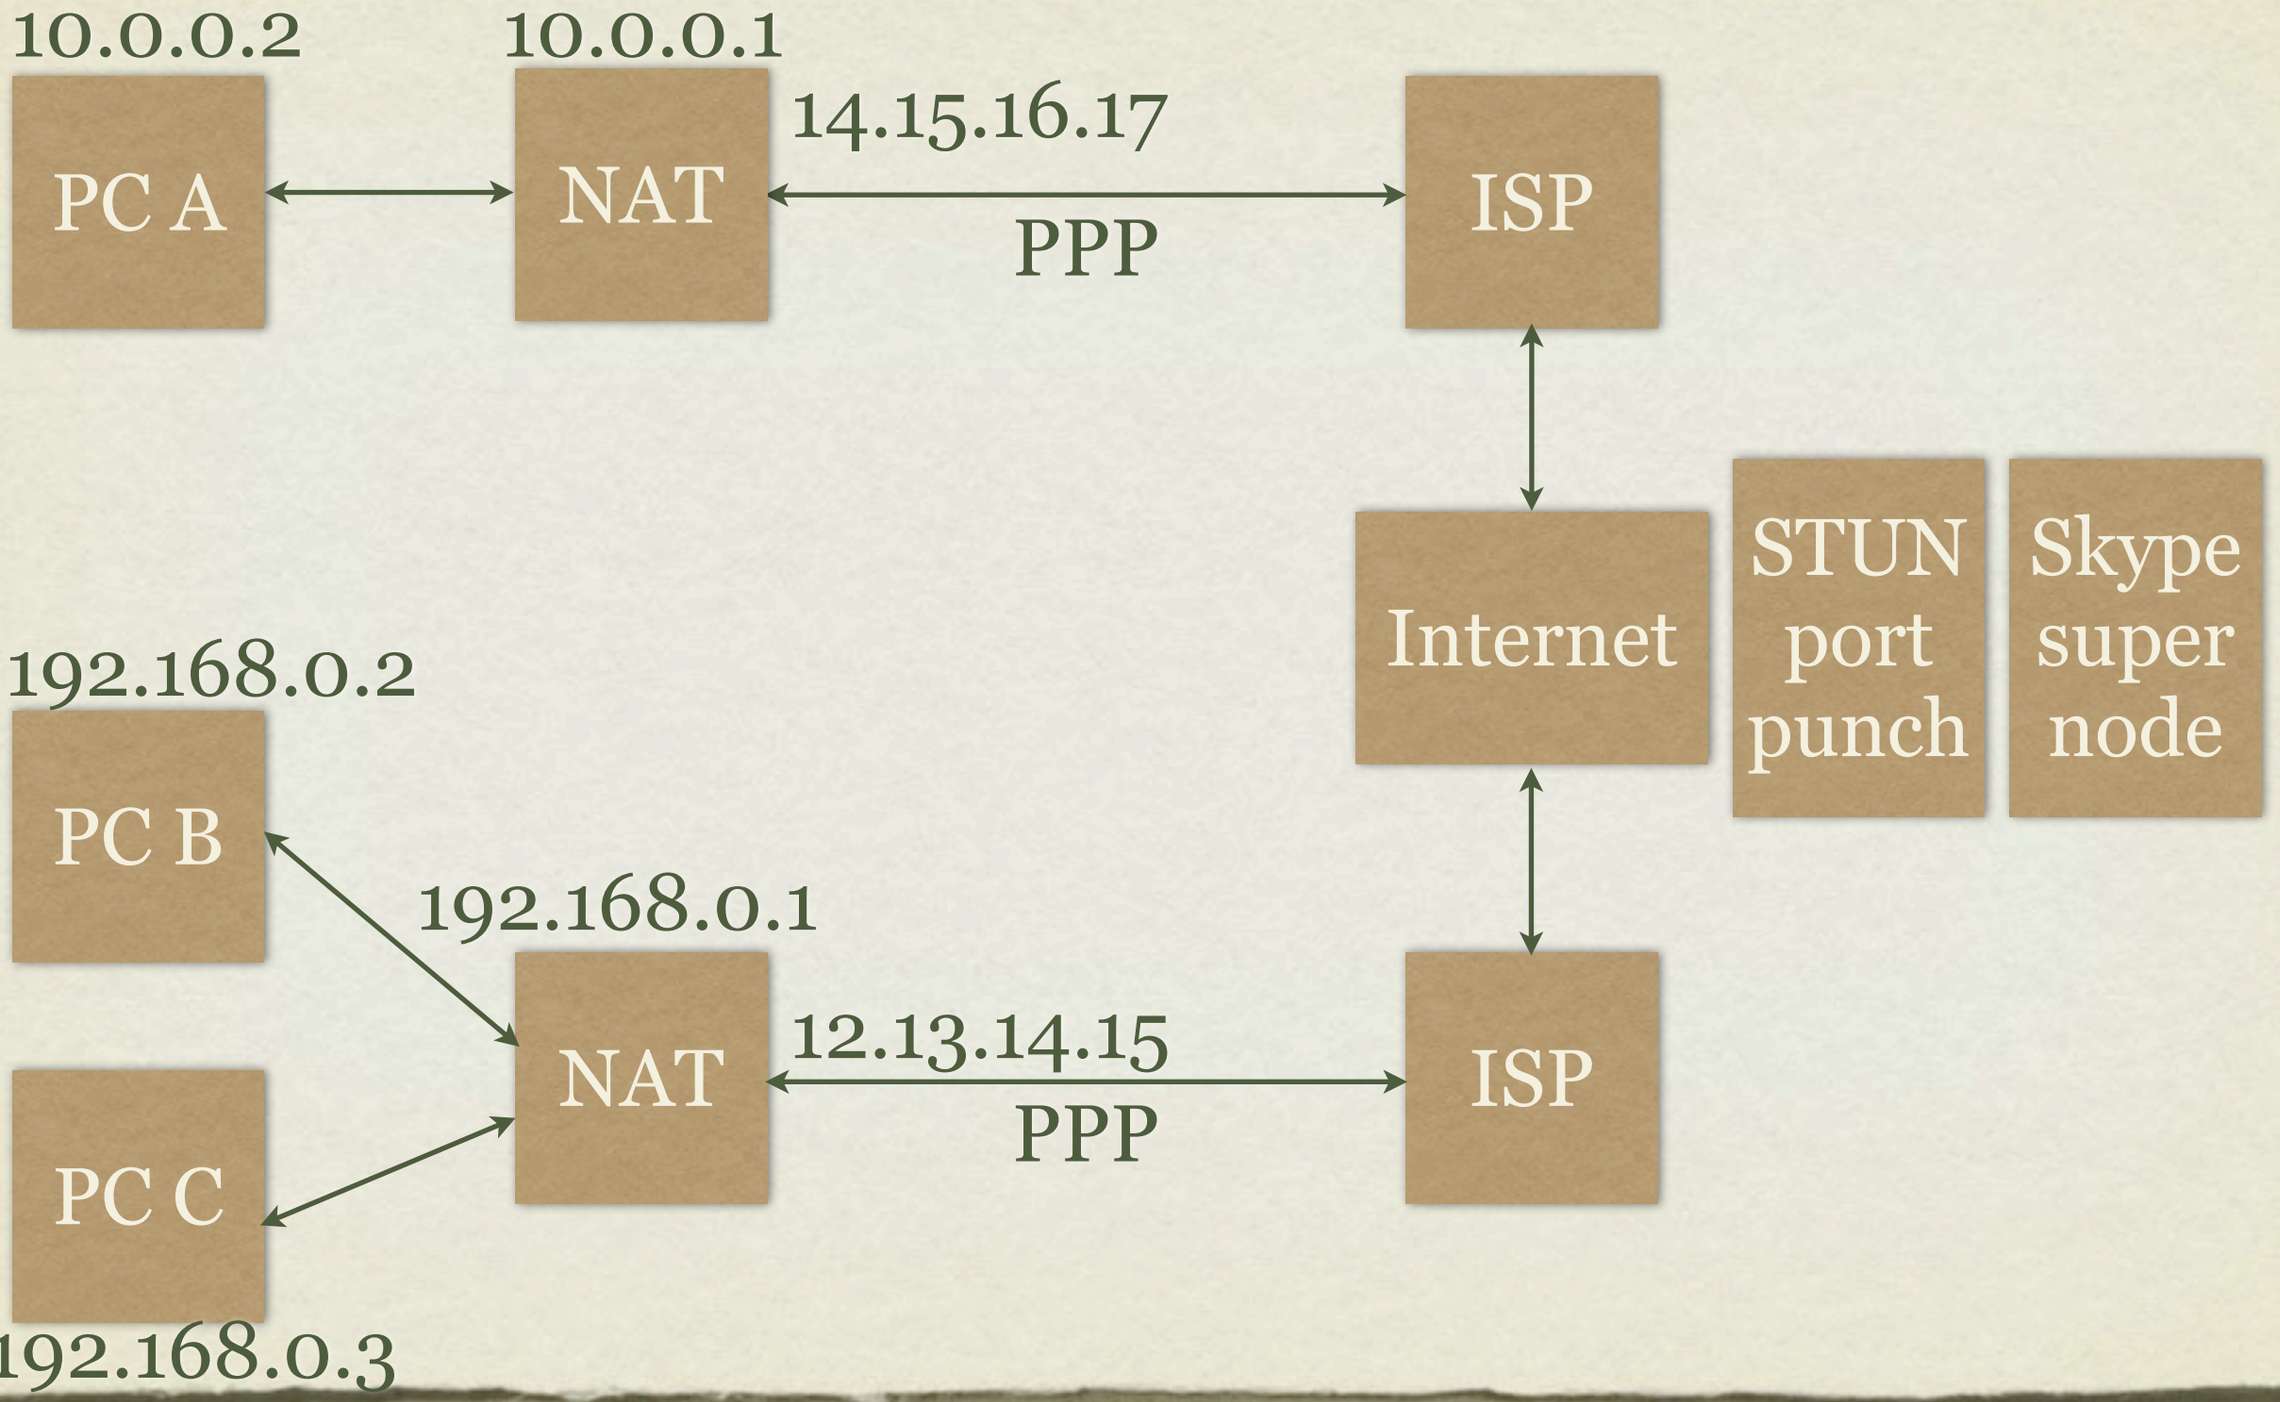

BACKGROUND: EDGE COMPLEXITY

home PC

work PC

iPhone

iPad

BACKGROUND: EDGE COMPLEXITY

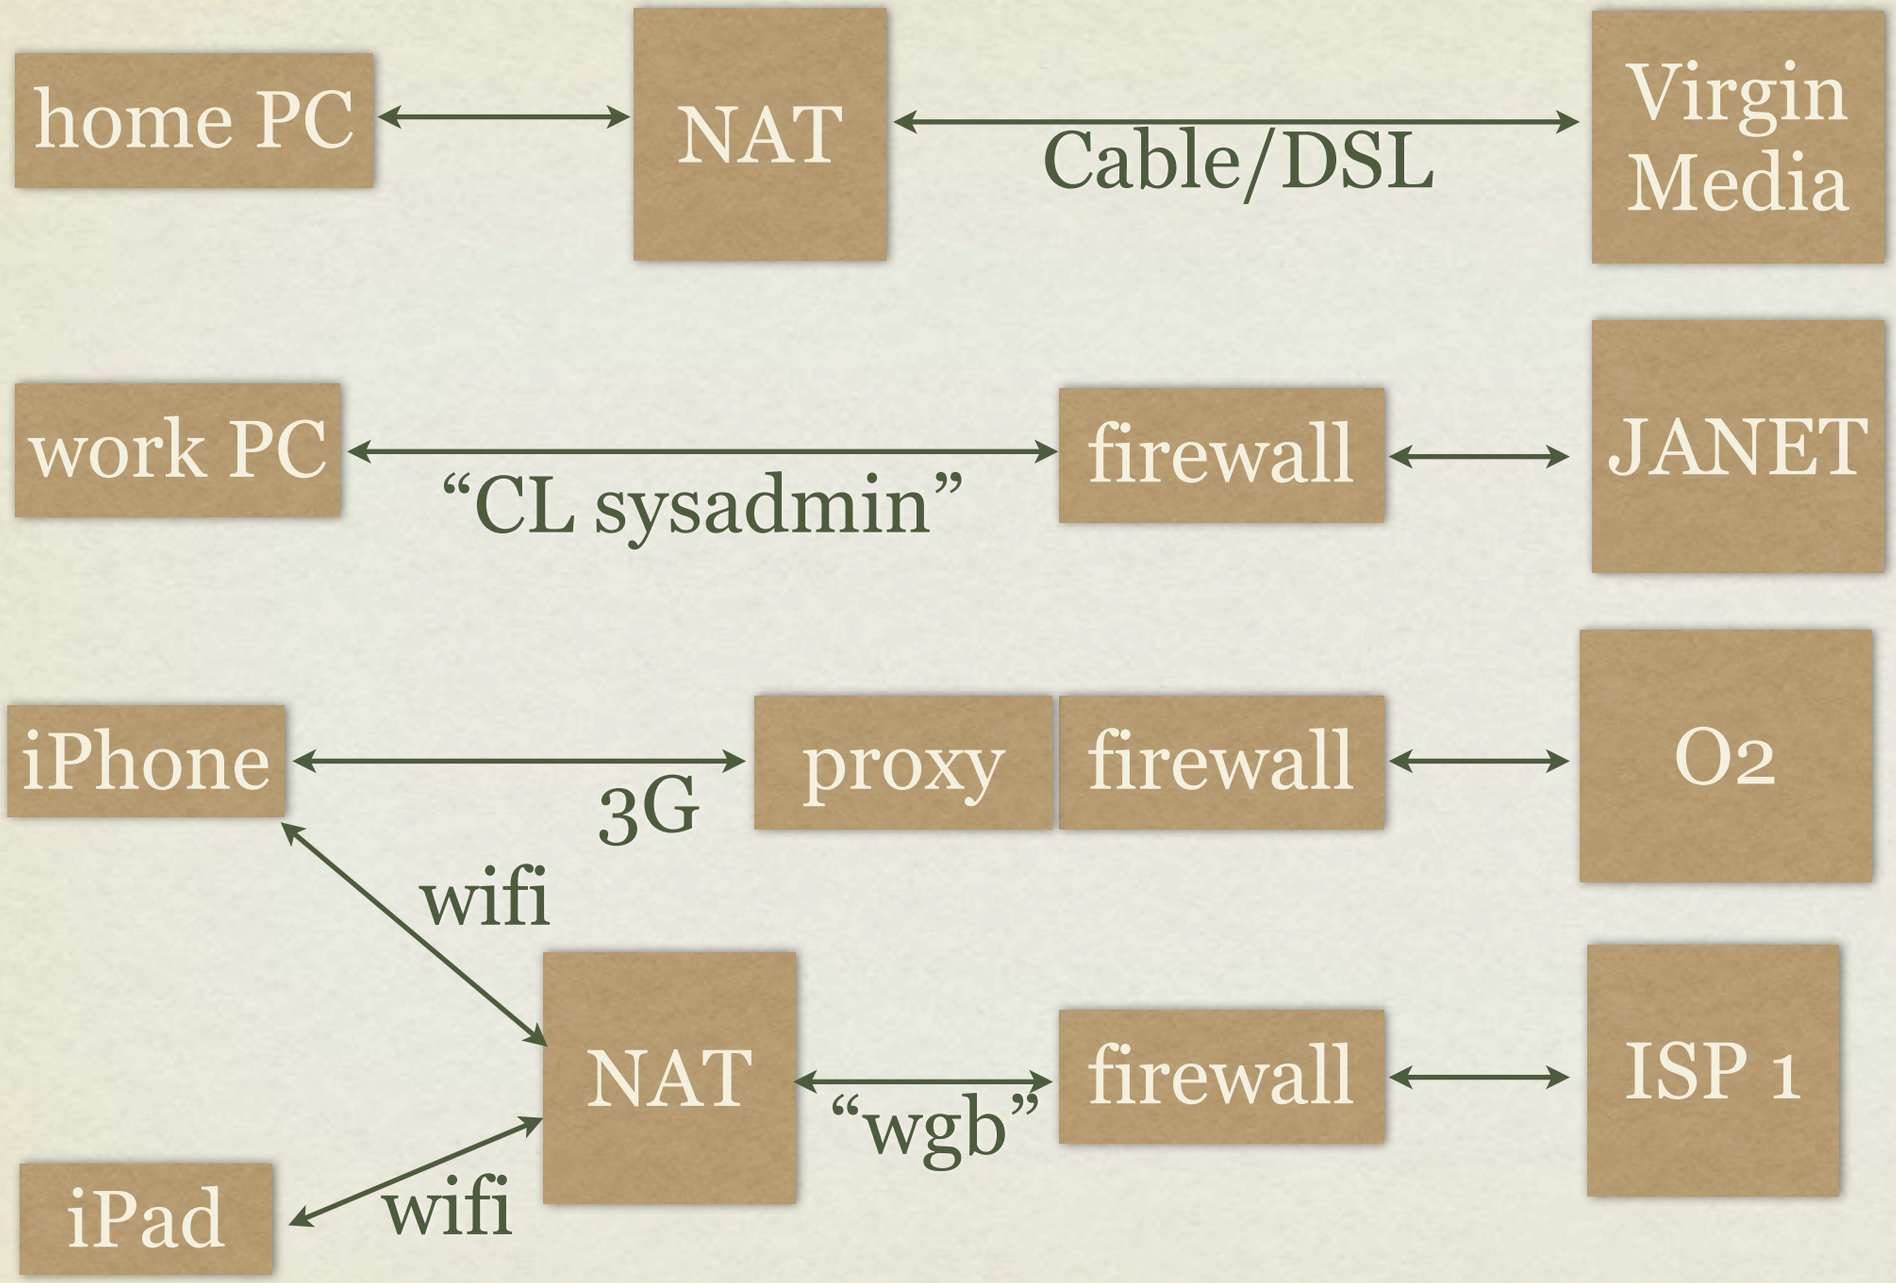

BACKGROUND: THE CLOUD

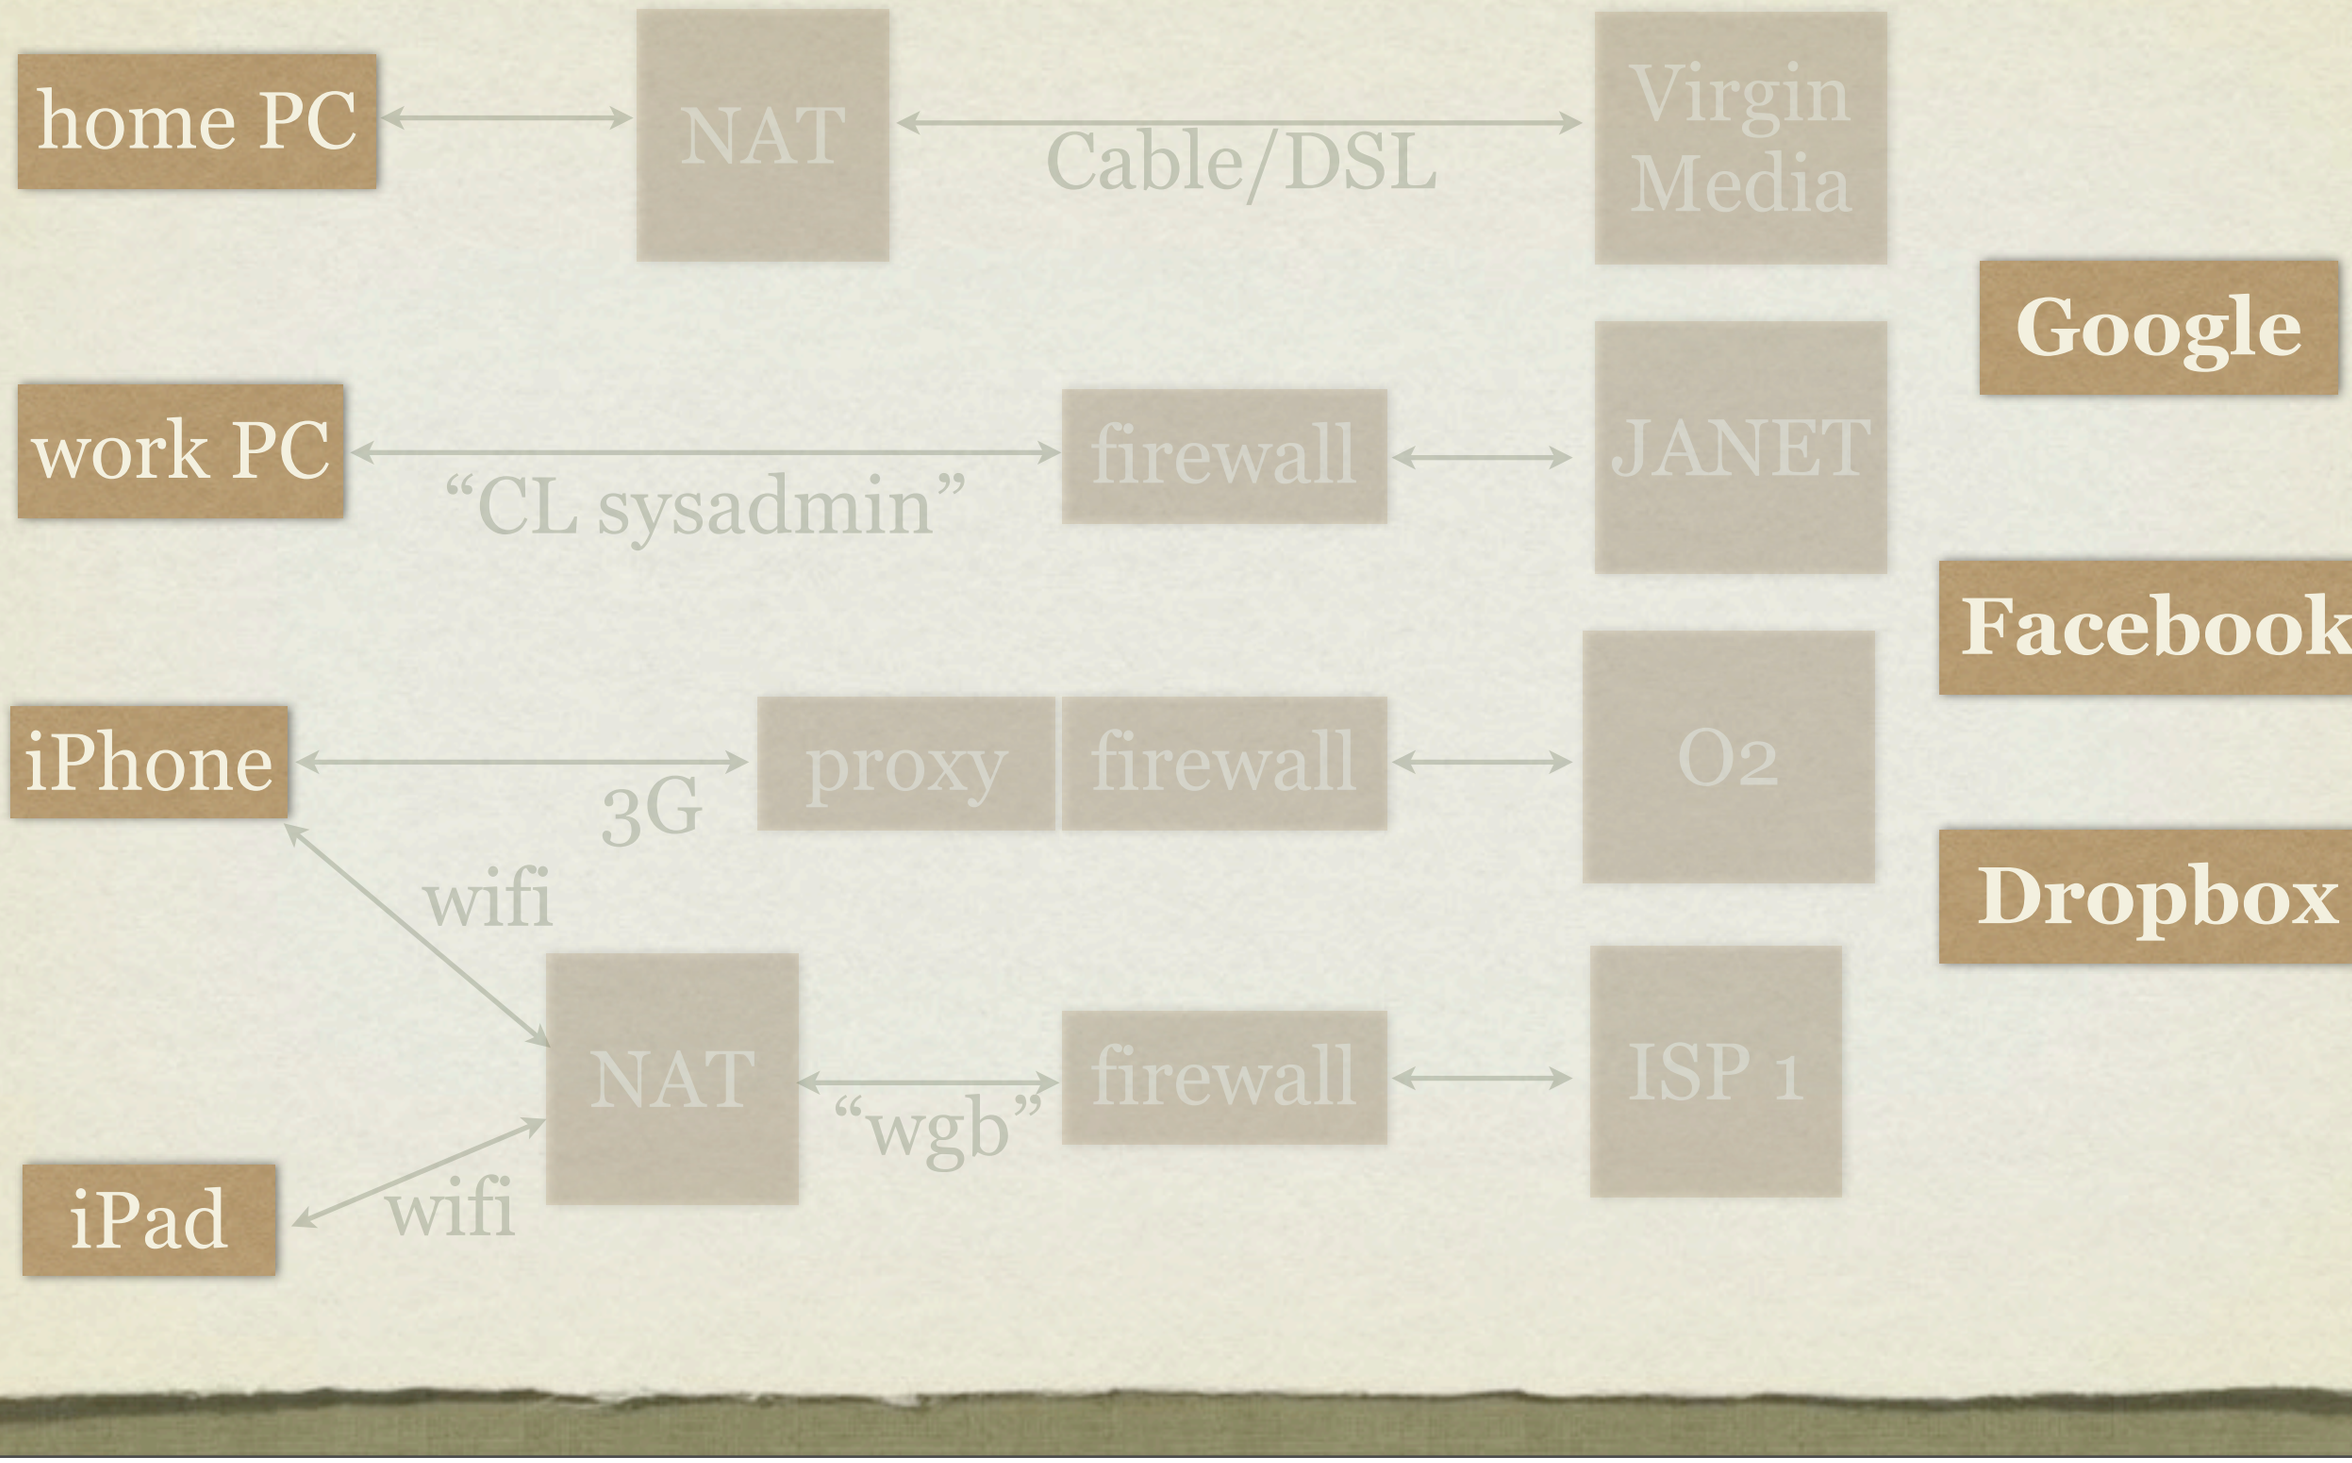

BACKGROUND: THE CLOUD

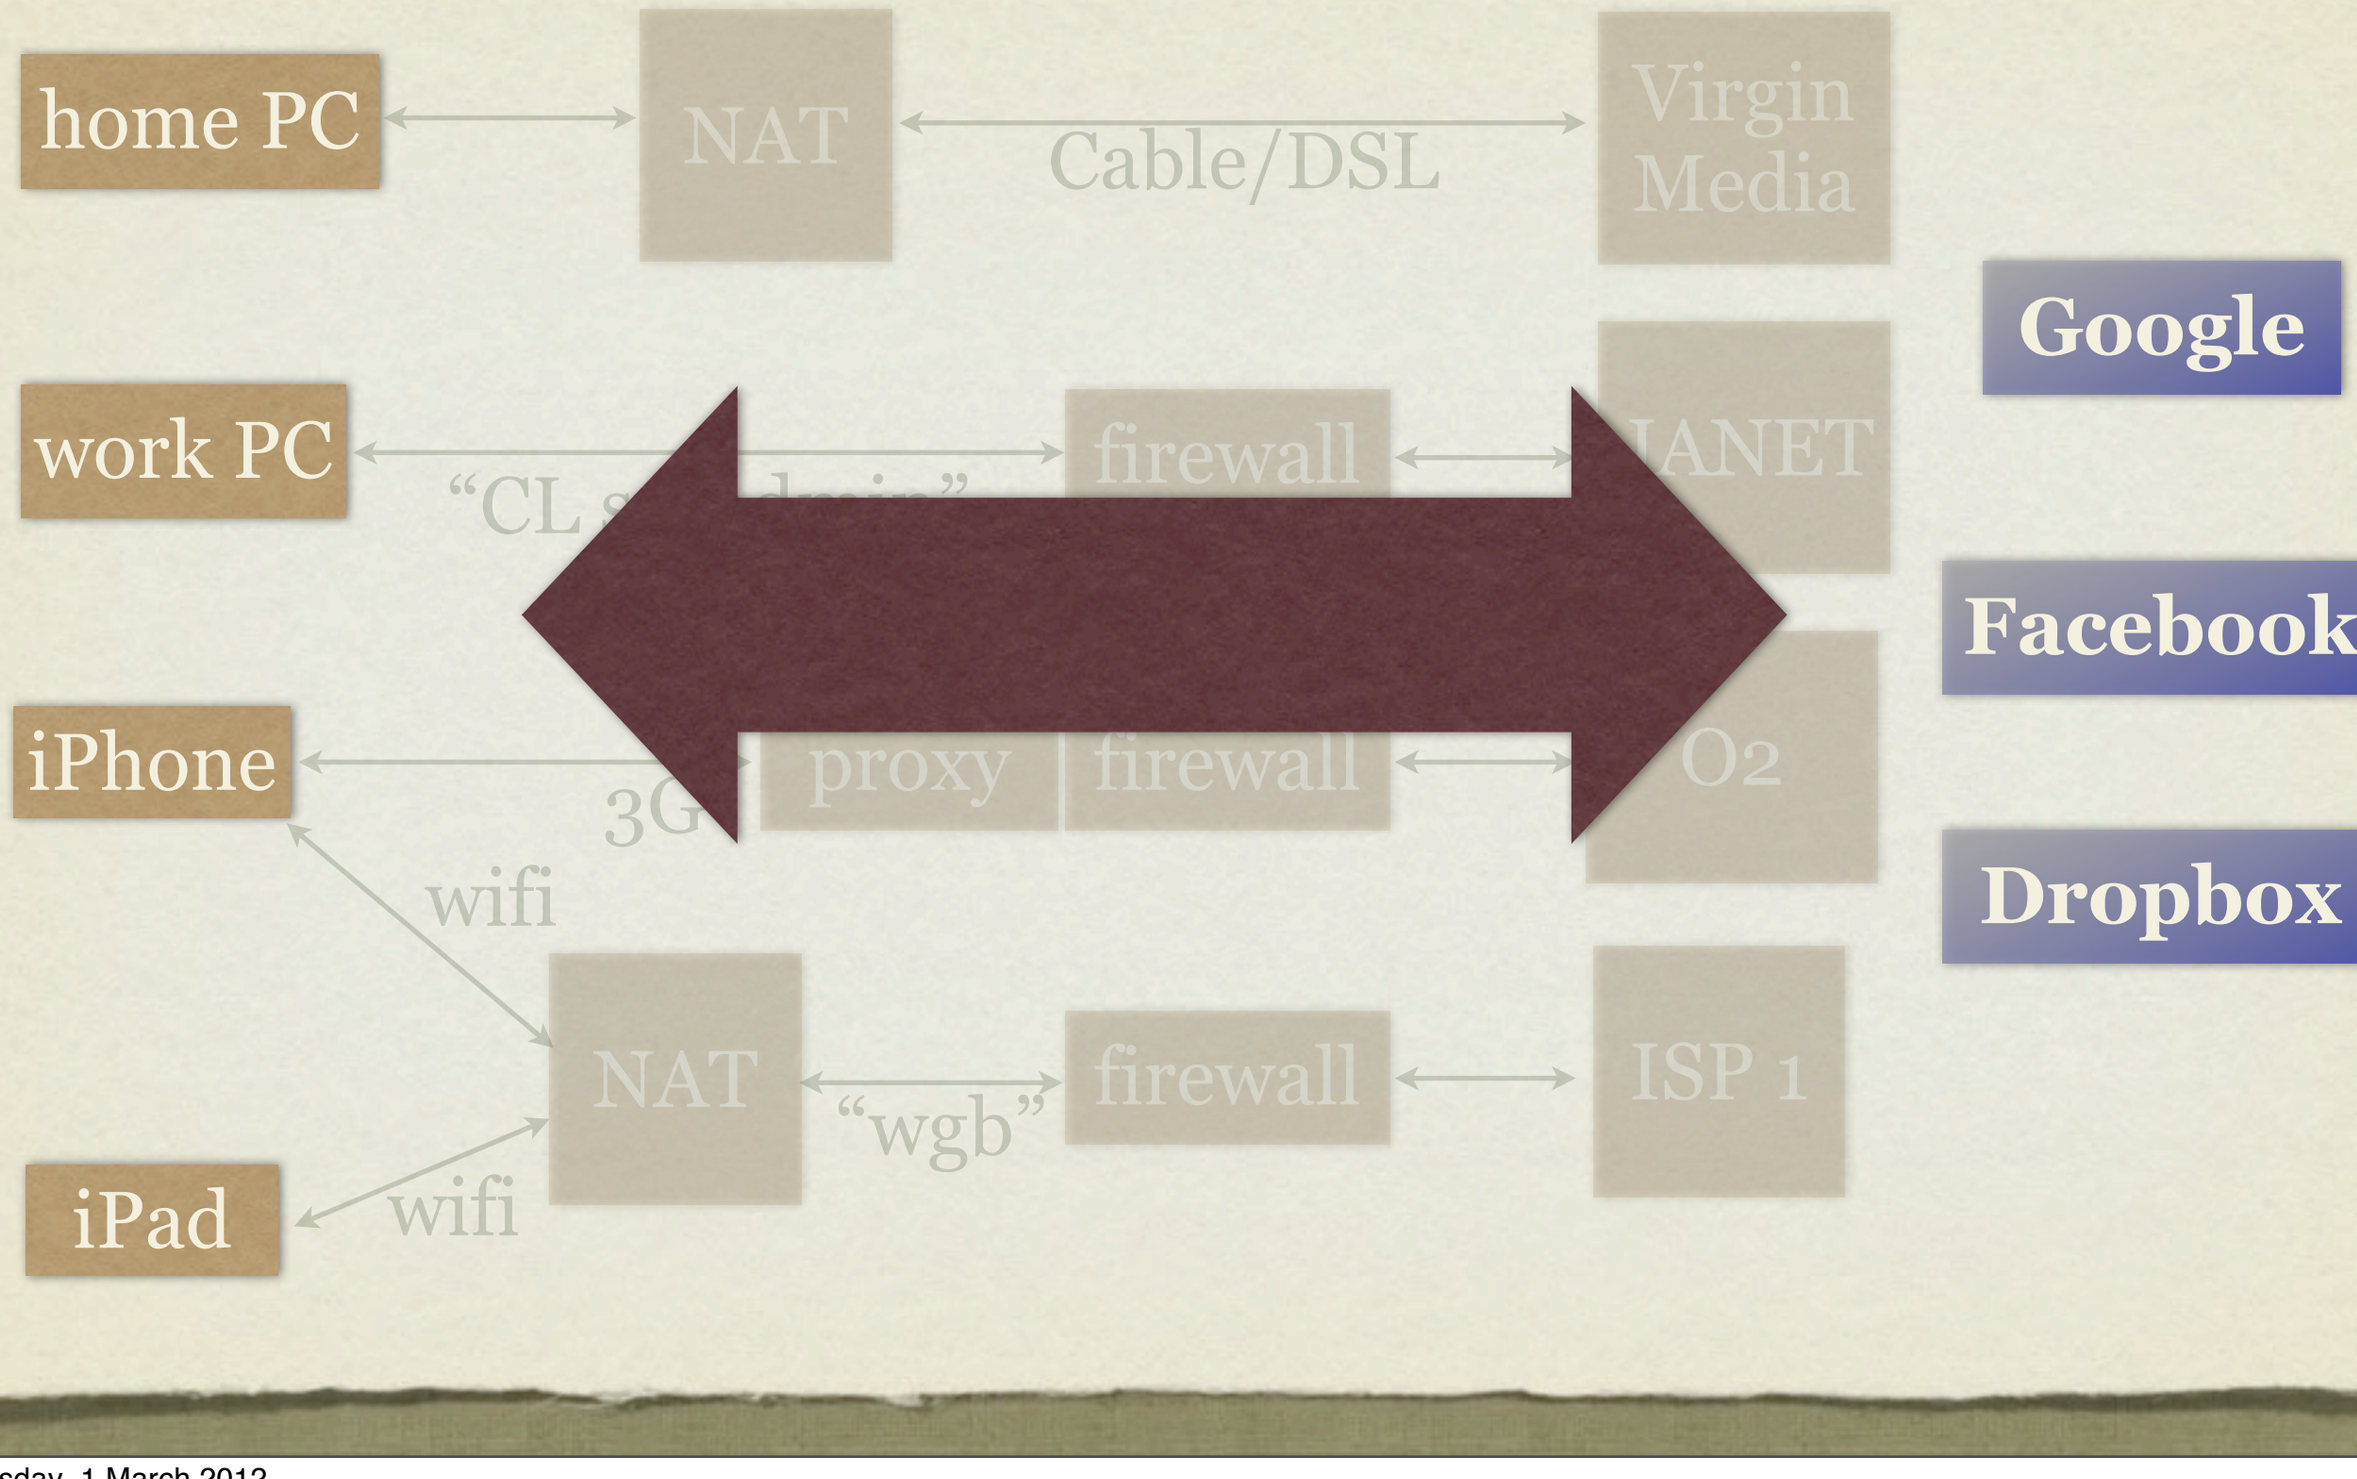

BACKGROUND: CLOUDS ROCK

- **Identity:** high-level, easy-to-use device registration (“my iPhone”, “work computer”).
- **Visibility:** only outbound connections required.
- **Reliability:** an army of professional sysadmins to worry.
- **Social:** cloud services can connect to each other.

BACKGROUND: CLOUDS SUCK

- **Privacy:** all data controlled by third-party, with their own policies (Google real name!).
- **Security:** one leak is all it takes. Irrevocable loss.
- **Cost:** orders of magnitude more resources on edge networks (e.g. bandwidth/latency).
- **Availability:** what if your house is disconnected?
- **Energy:** cost of moving data to/from edge and cloud.

LET A MILLION CLOUDS BLOOM!

- **Why can't we all have our own cloud between our devices and networks?**
  - #1: we have no **identity** online.
  - #2: lack of end-to-end **connectivity** limits visibility**.**
  - #3: who hosts our stuff **reliably**?
  - #4: why bother? What **new services** does this enable?

SIGNPOSTS

- The **minimum coordination infrastructure** required to establish routes between edge devices.
  - DNS is woefully under-used to date. (ab)use it for global signalling through middleboxes.
  - Work offline and support lazy synchronisation
  - Support confidential lookups
- **Desired user experience:** when I address a device by its hostname, the result should just work (e.g. iphone.anil)

PROBLEM #1: IDENTITY

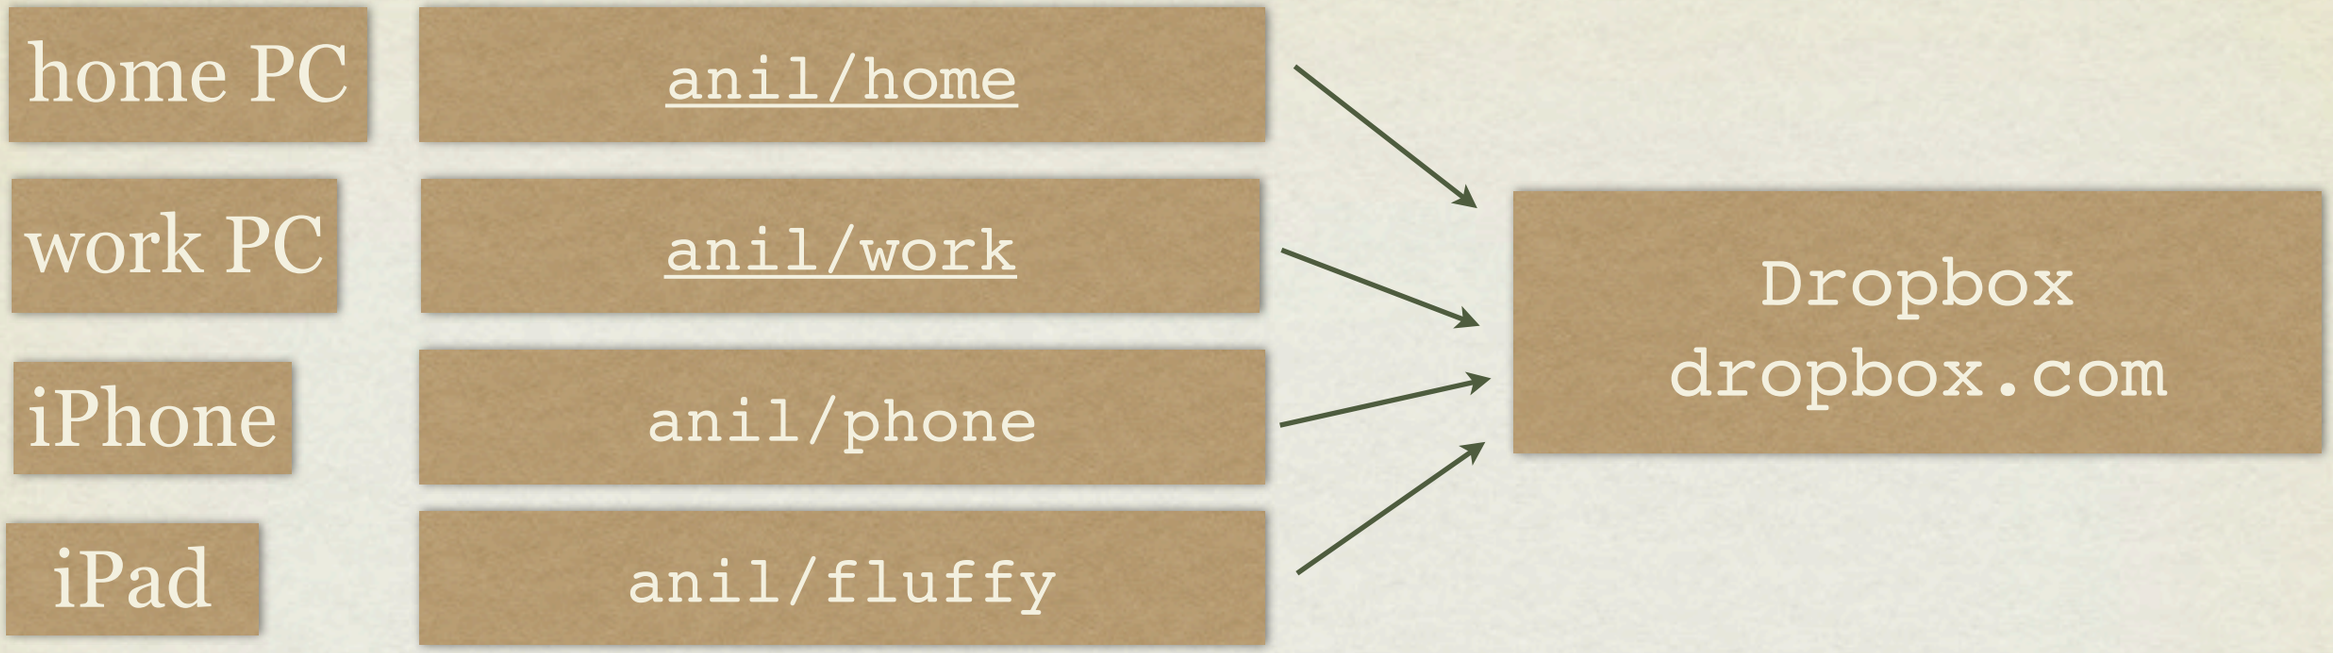

PROBLEM #1: IDENTITY

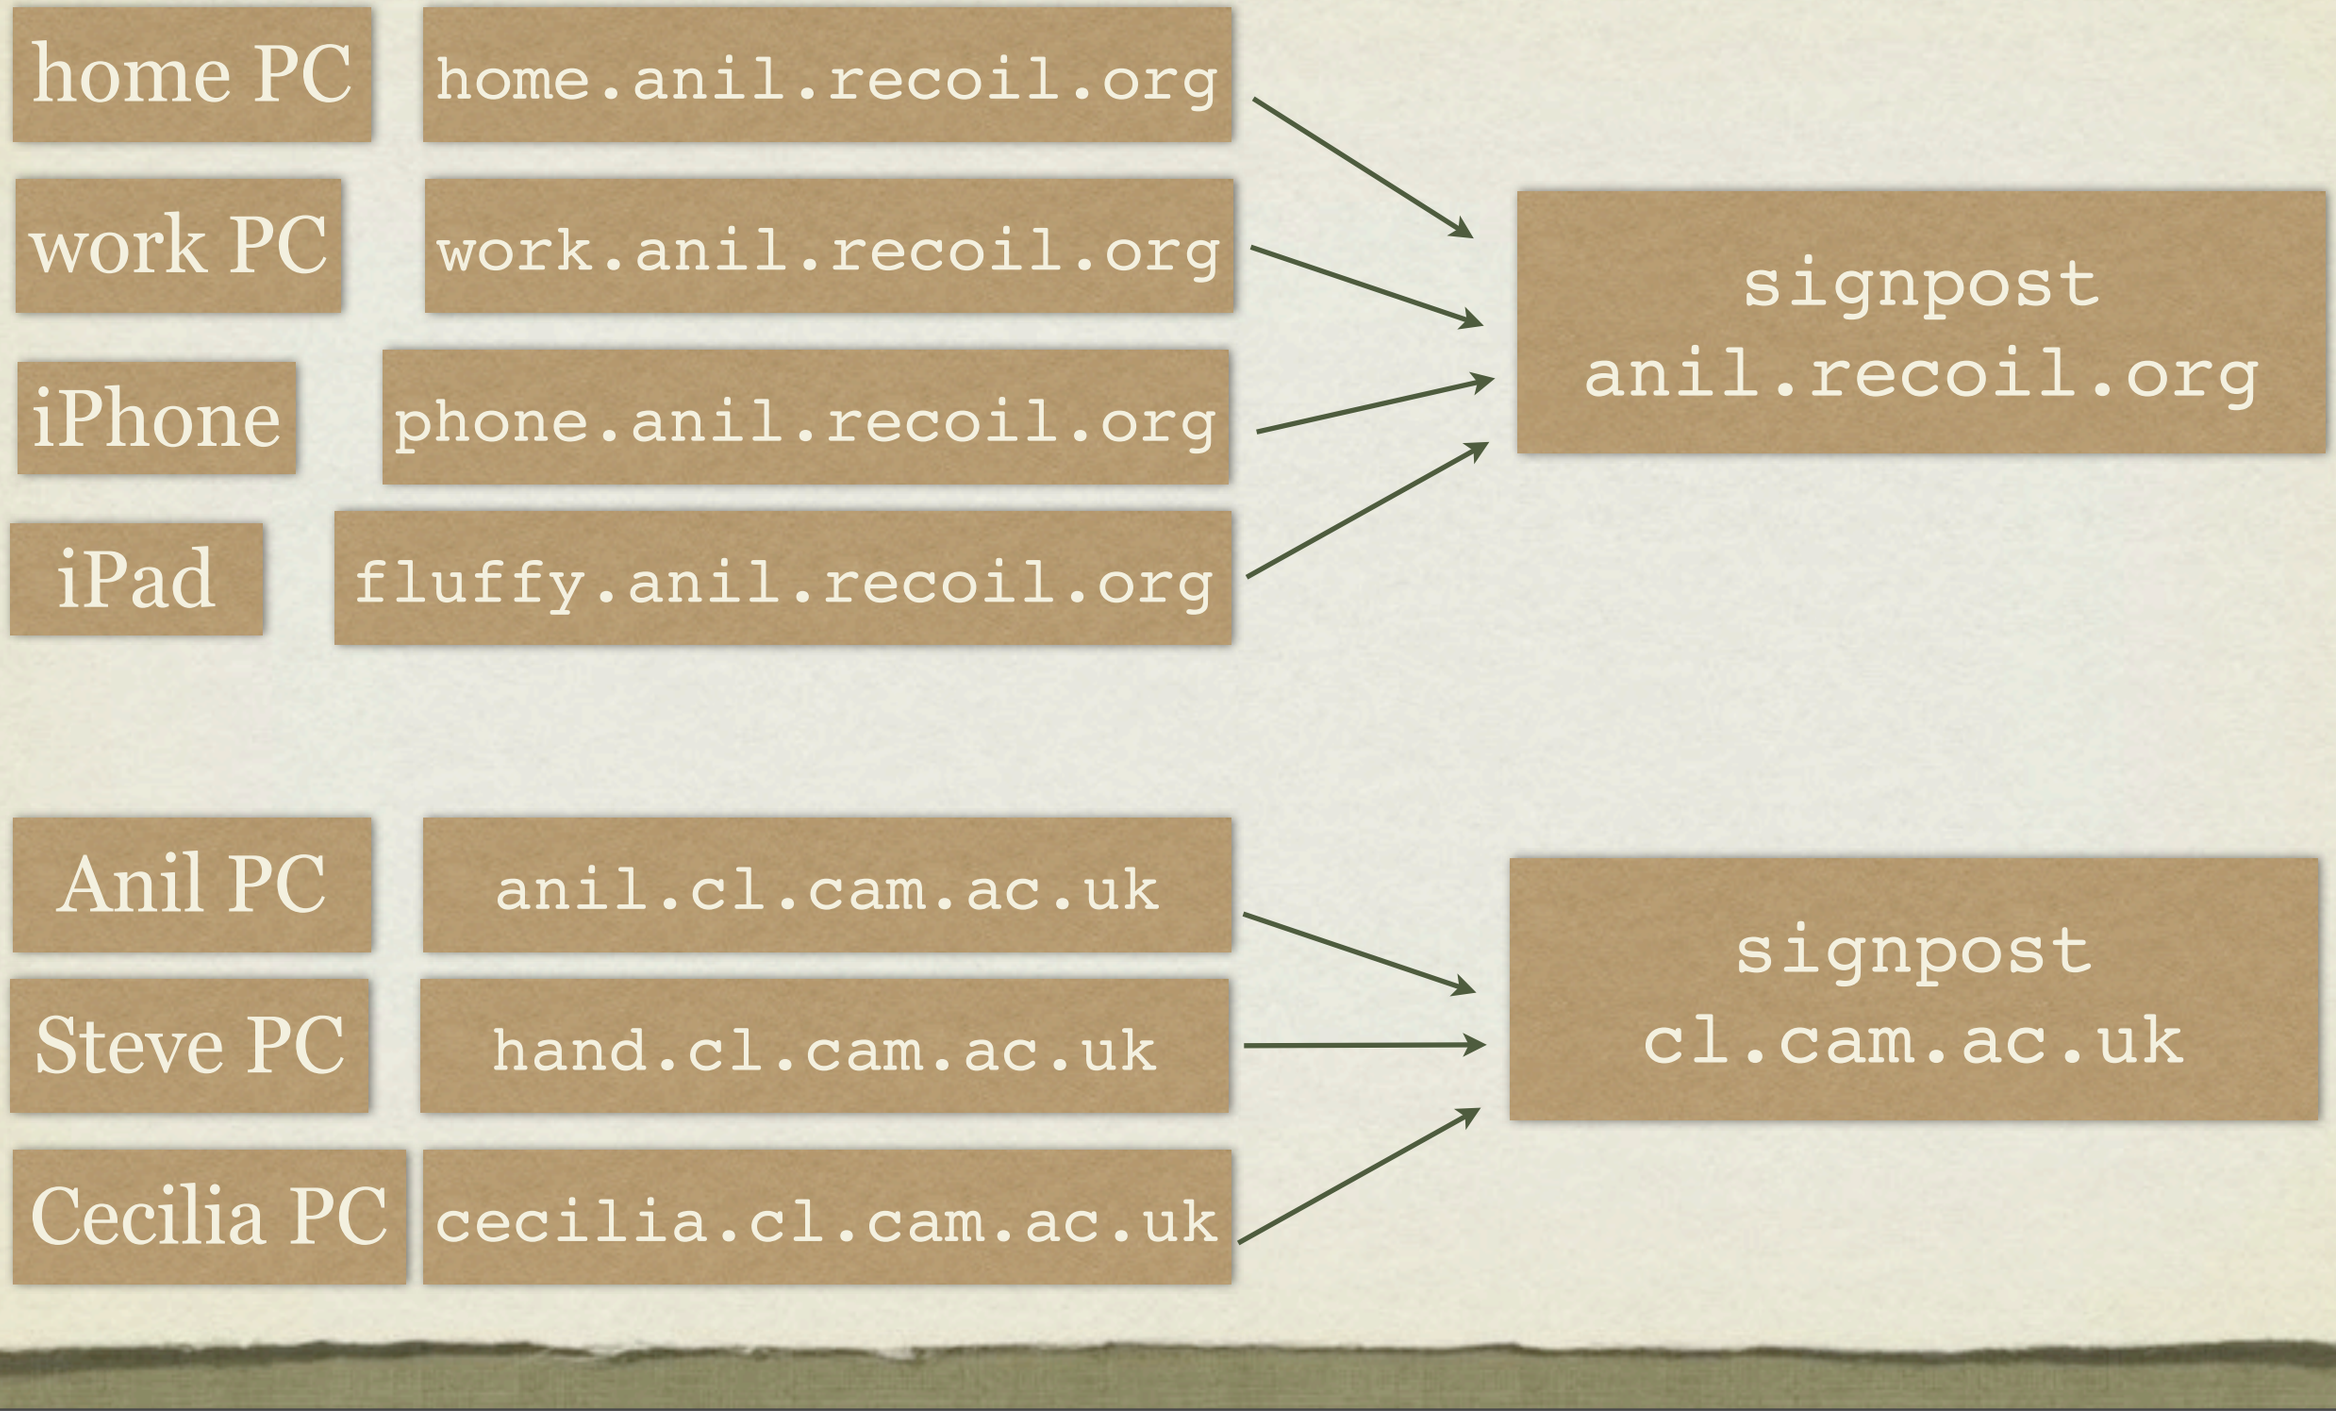

PROBLEM #1: IDENTITY

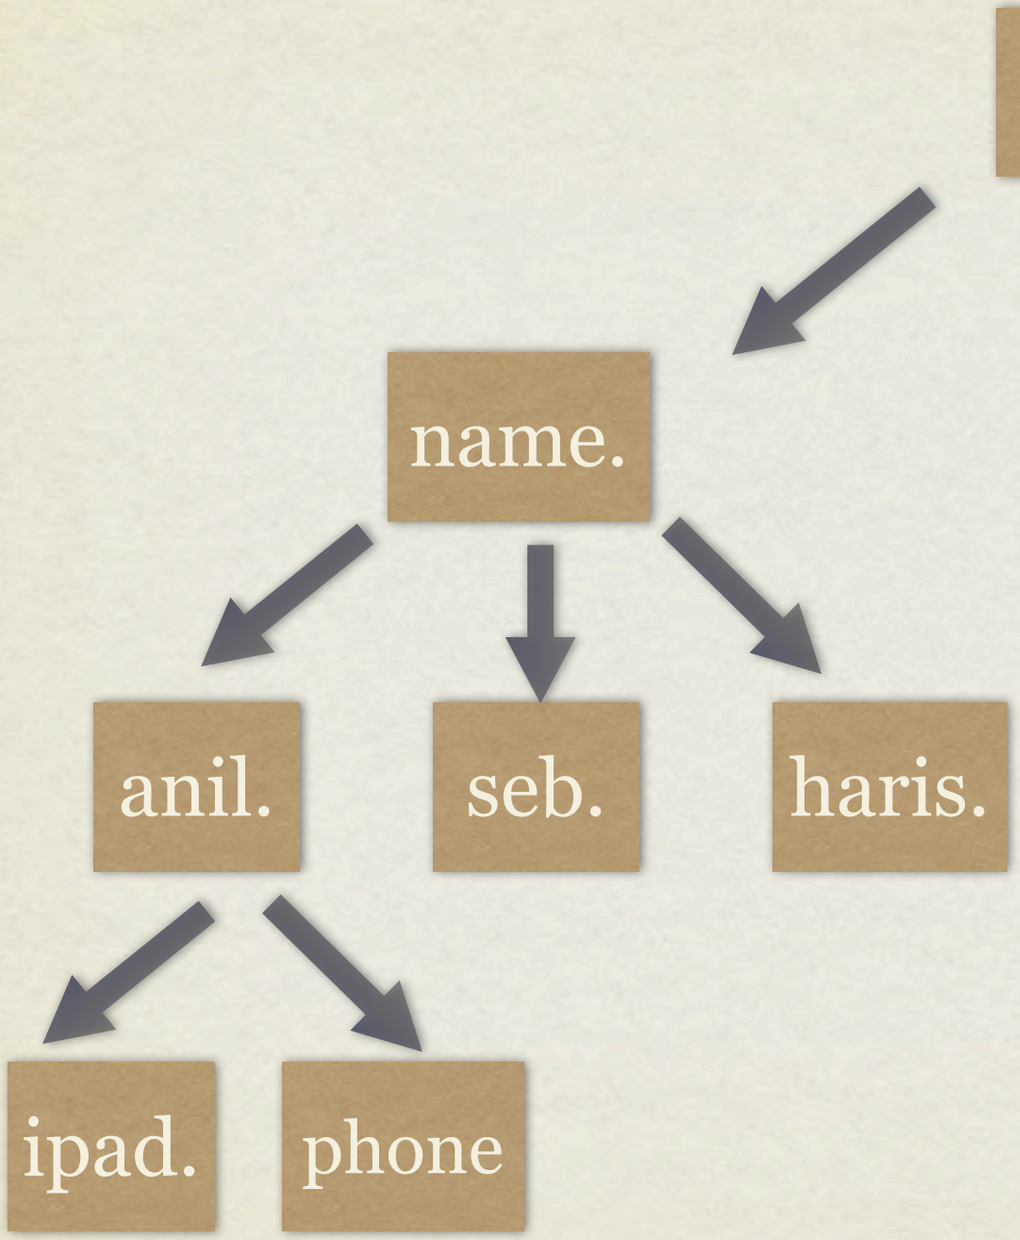

PROBLEM #1: IDENTITY

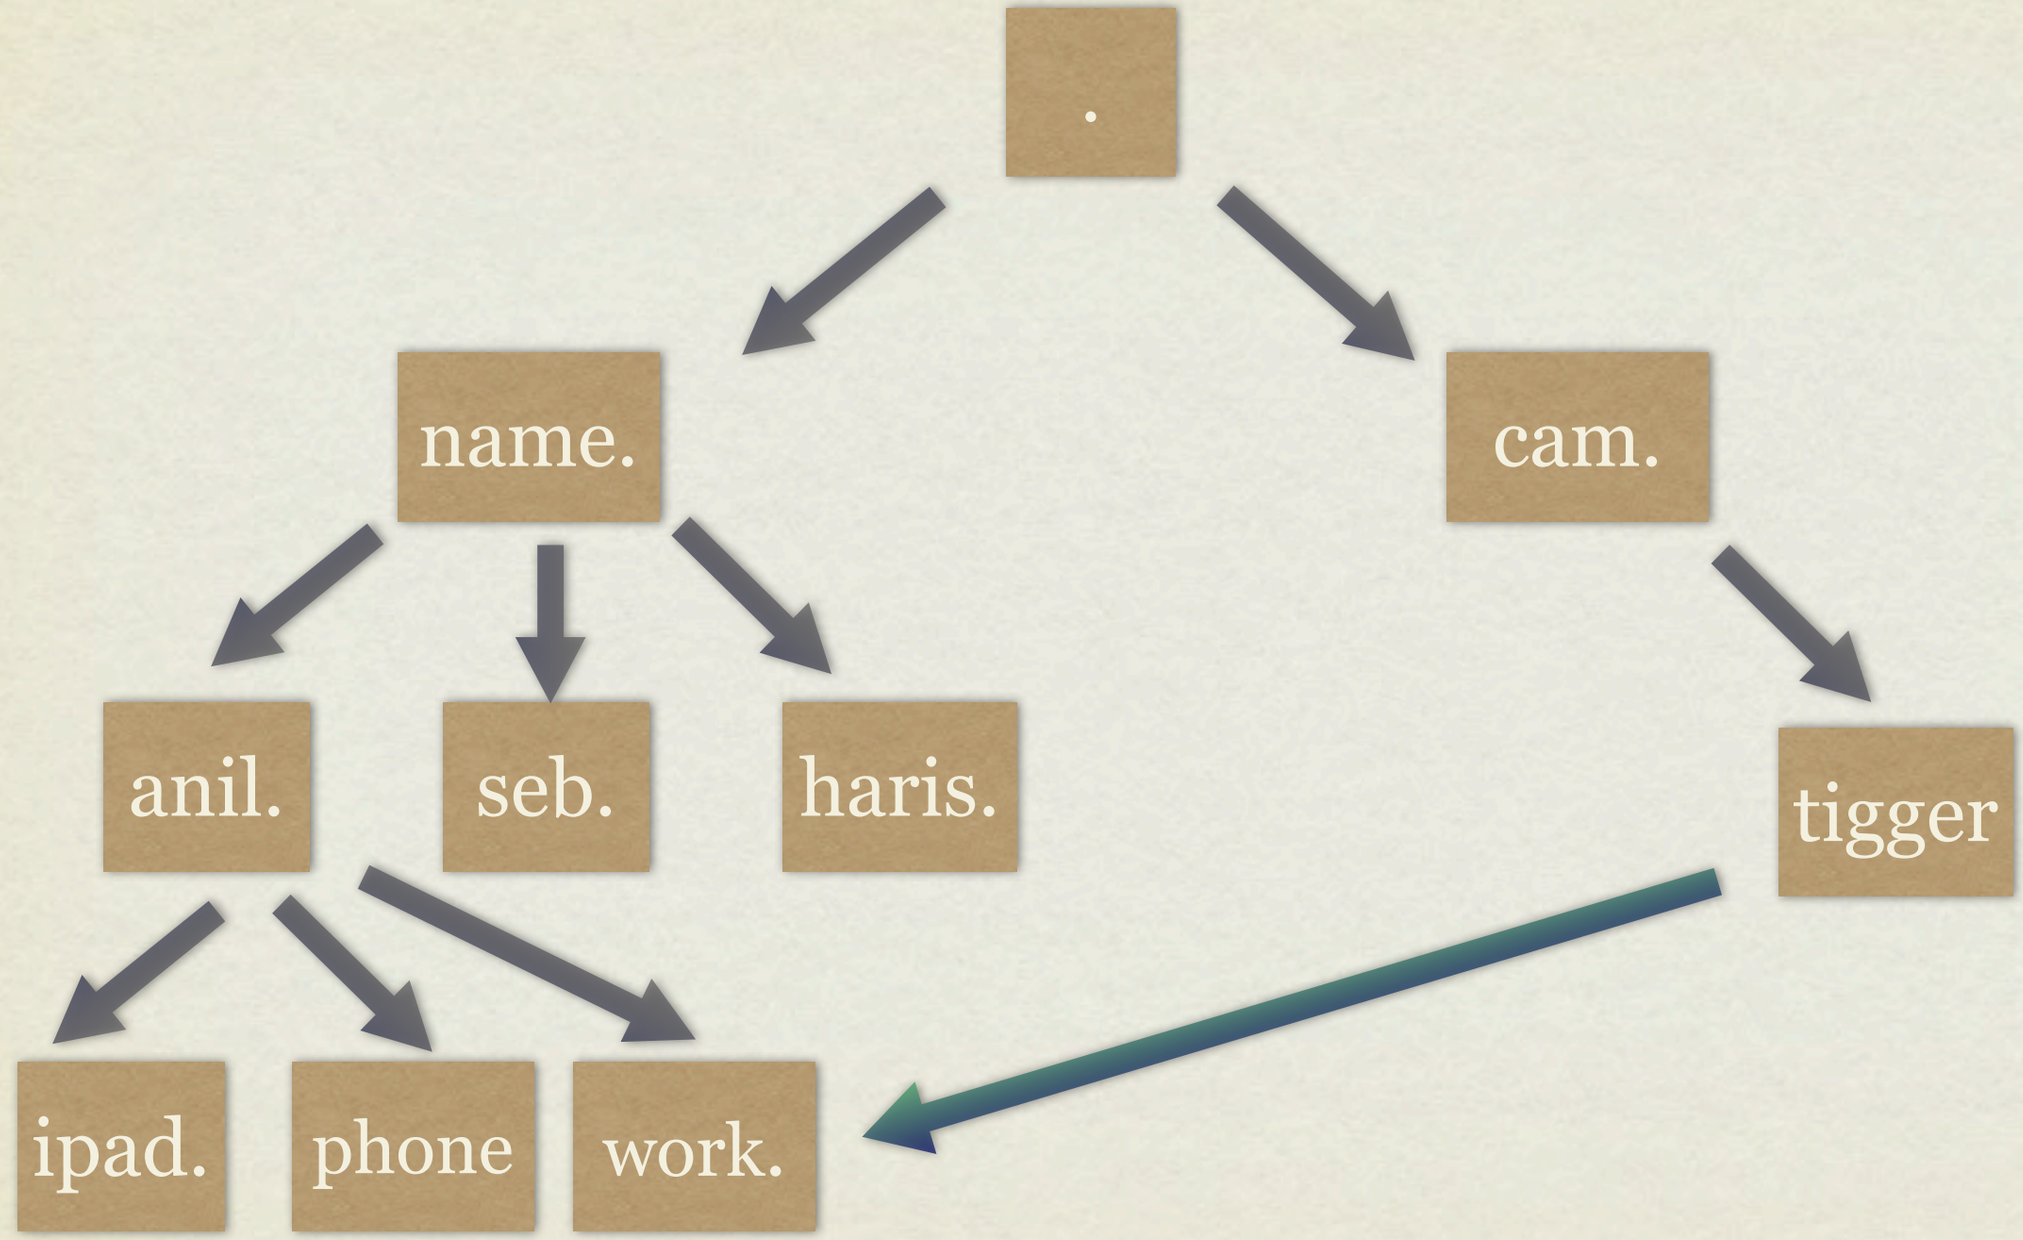

PROBLEM #1: IDENTITY

- **Identity:** every individual has a domain name hierarchy
  - DNSSEC means you register a single public/private key (`anil.recoil.org`)
- Proxy identity to social networks (`anilmadhavapeddy.facebook.com`)
- Use address book to invisibly associate names to DNSSEC domain keys.
- Bind devices to your domain (“resurrected duckling”)

LET A MILLION CLOUDS BLOOM!

- **Why can't we all have our own cloud between our devices and networks?**

- #1: we have no **identity** online.
- #2: lack of end-to-end **connectivity** limits visibility**.**
- #3: who hosts our stuff **reliably**?
- #4: why bother? What **new services** does this enable?

PROBLEM #2: CONNECTIVITY

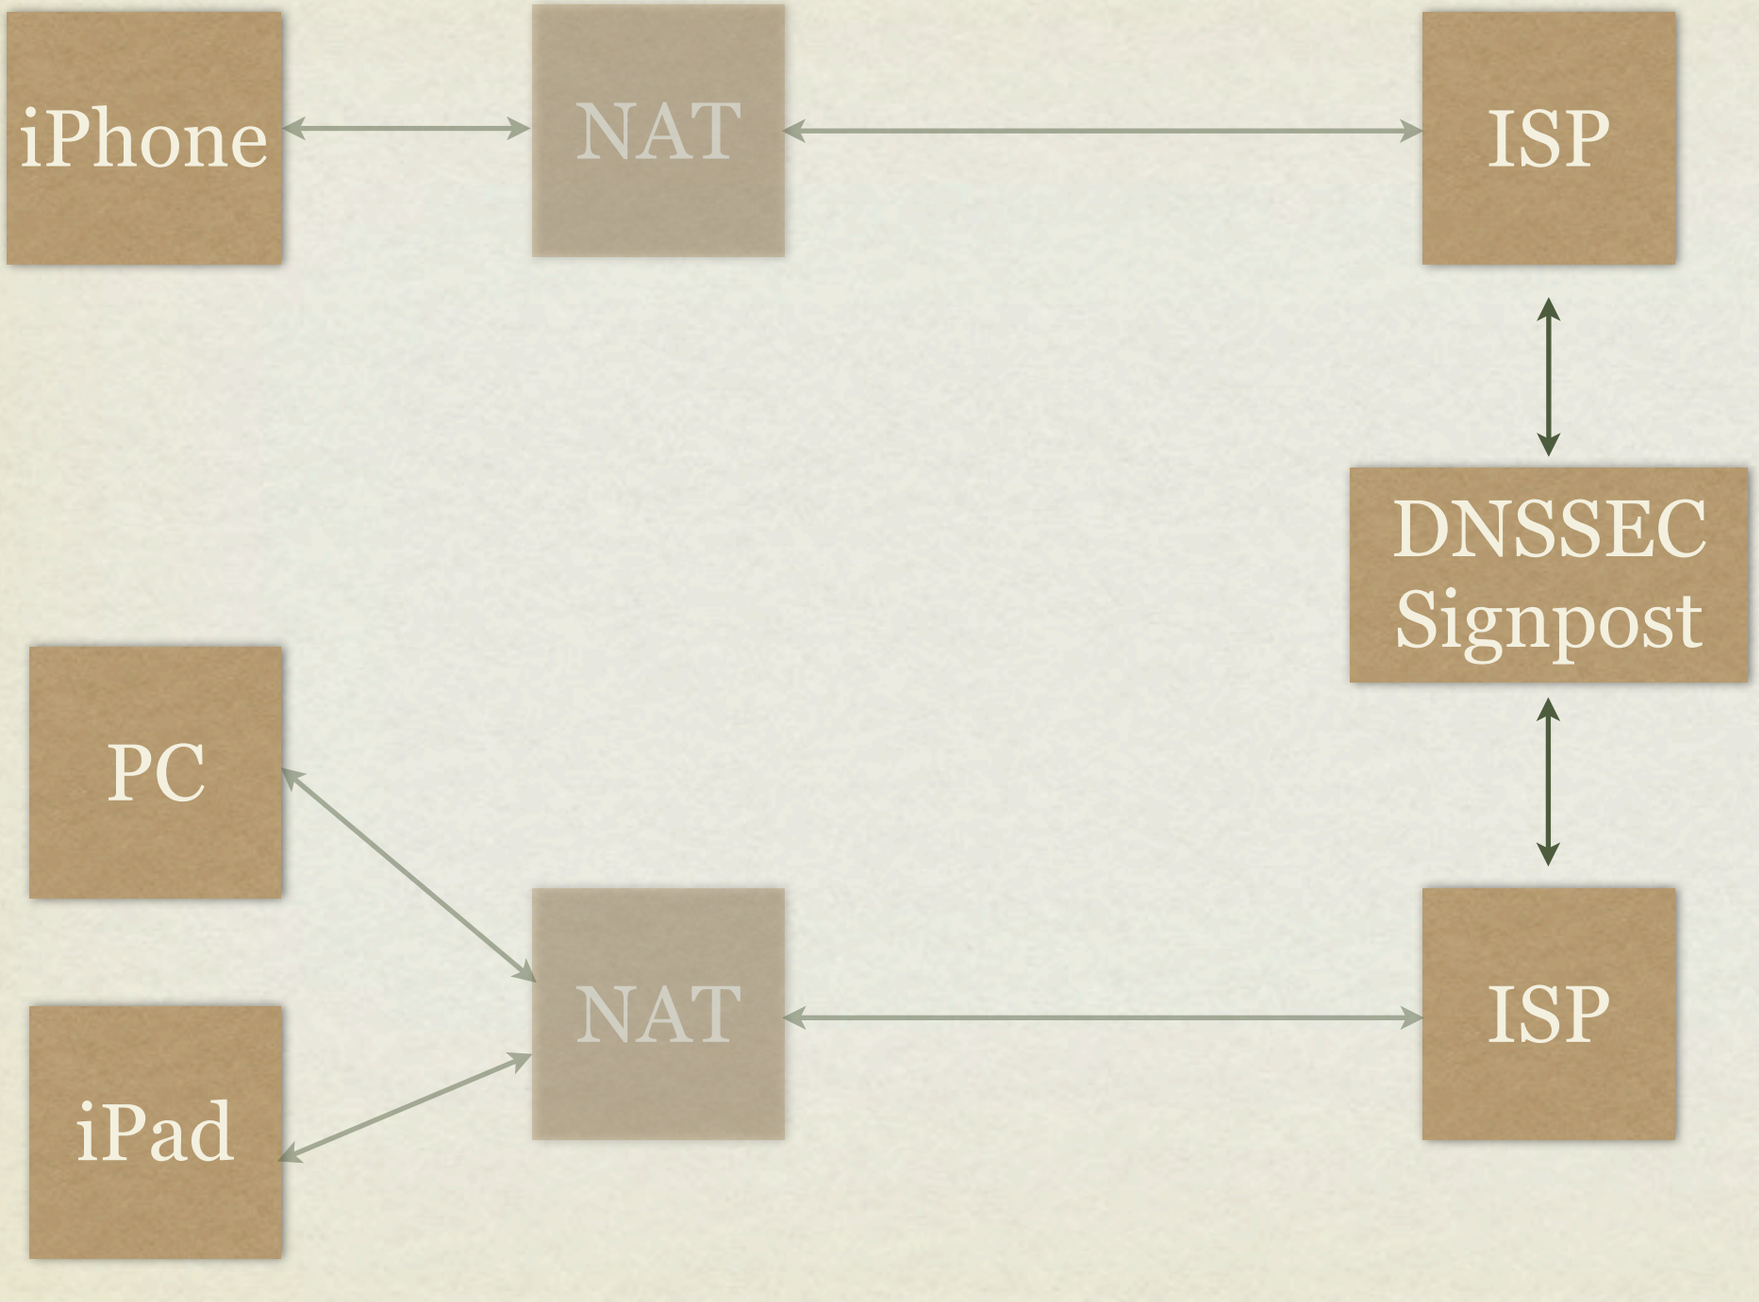

PROBLEM #2: CONNECTIVITY

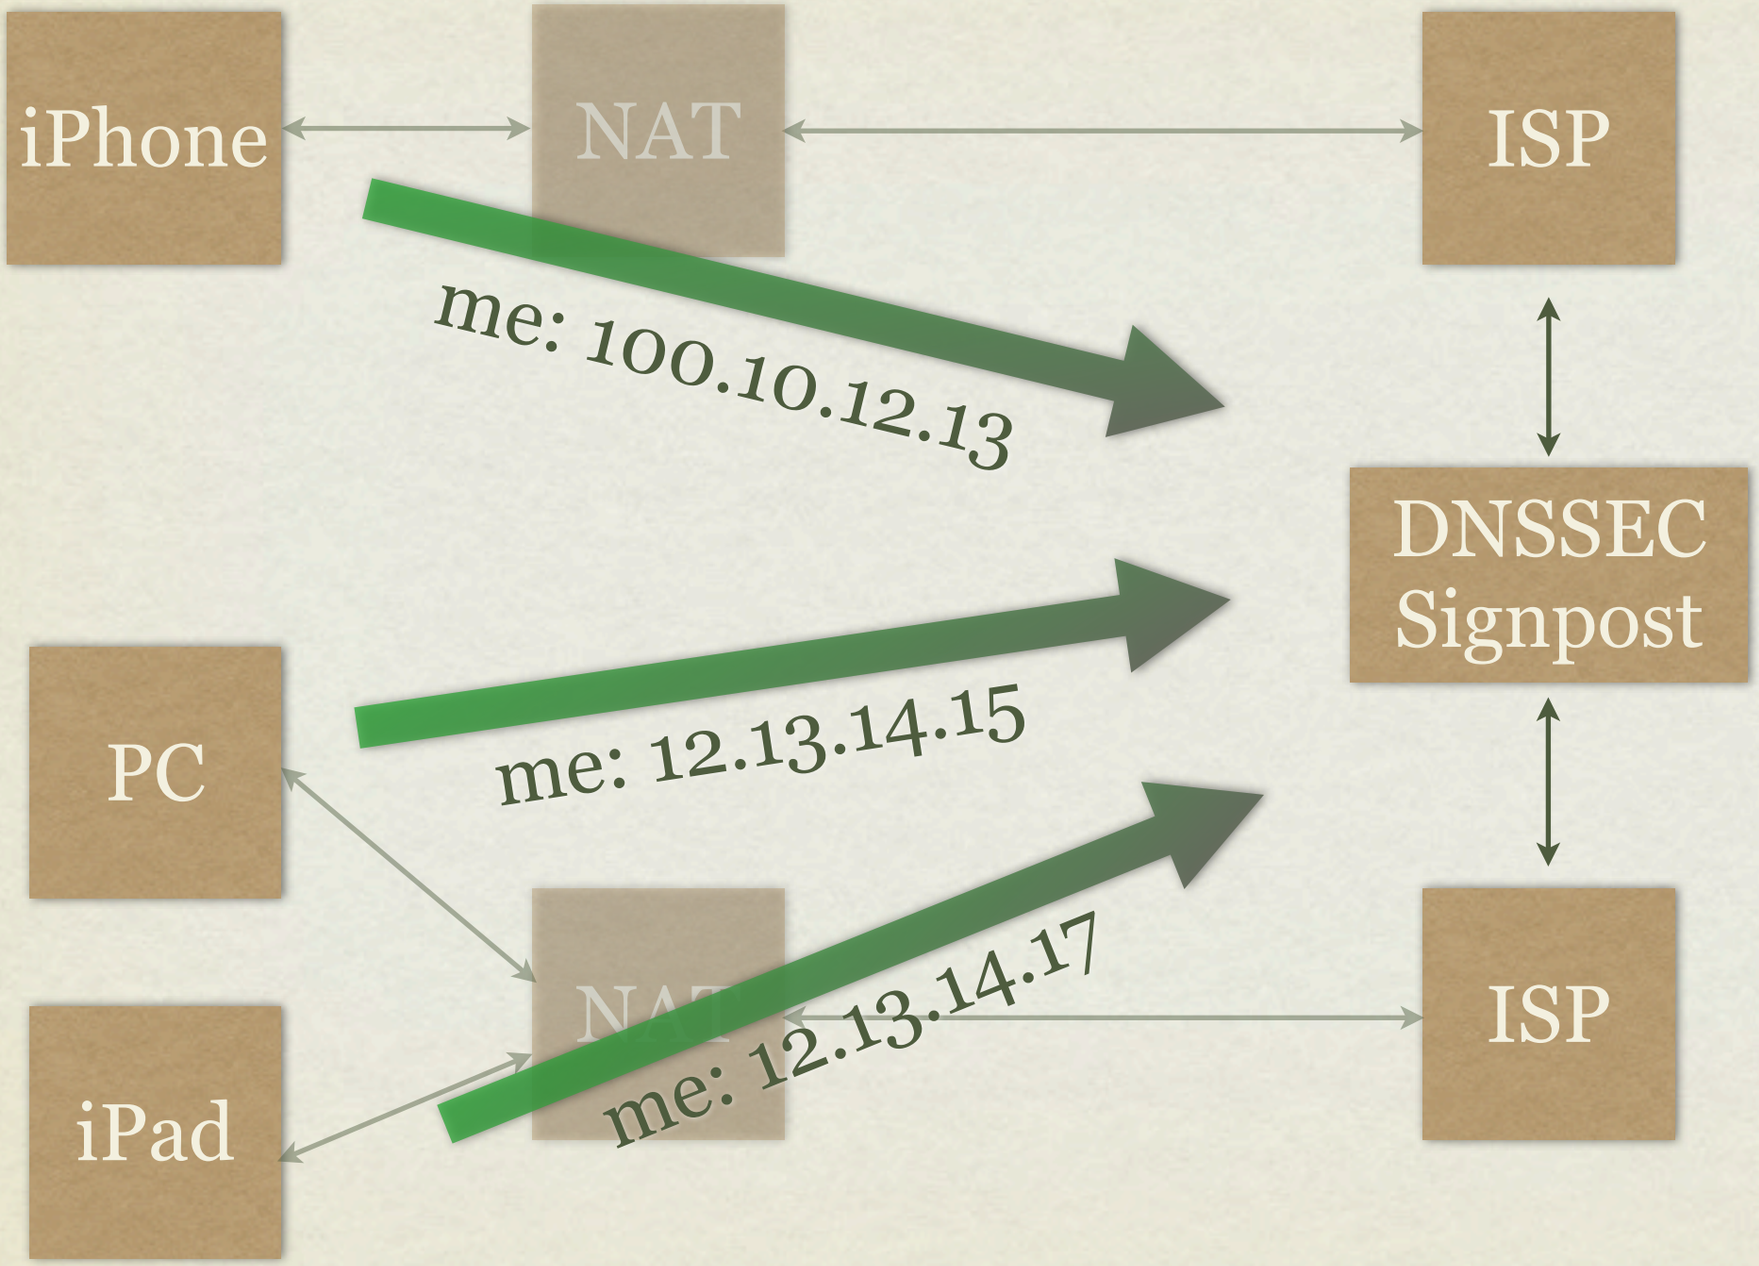

PROBLEM #2: CONNECTIVITY

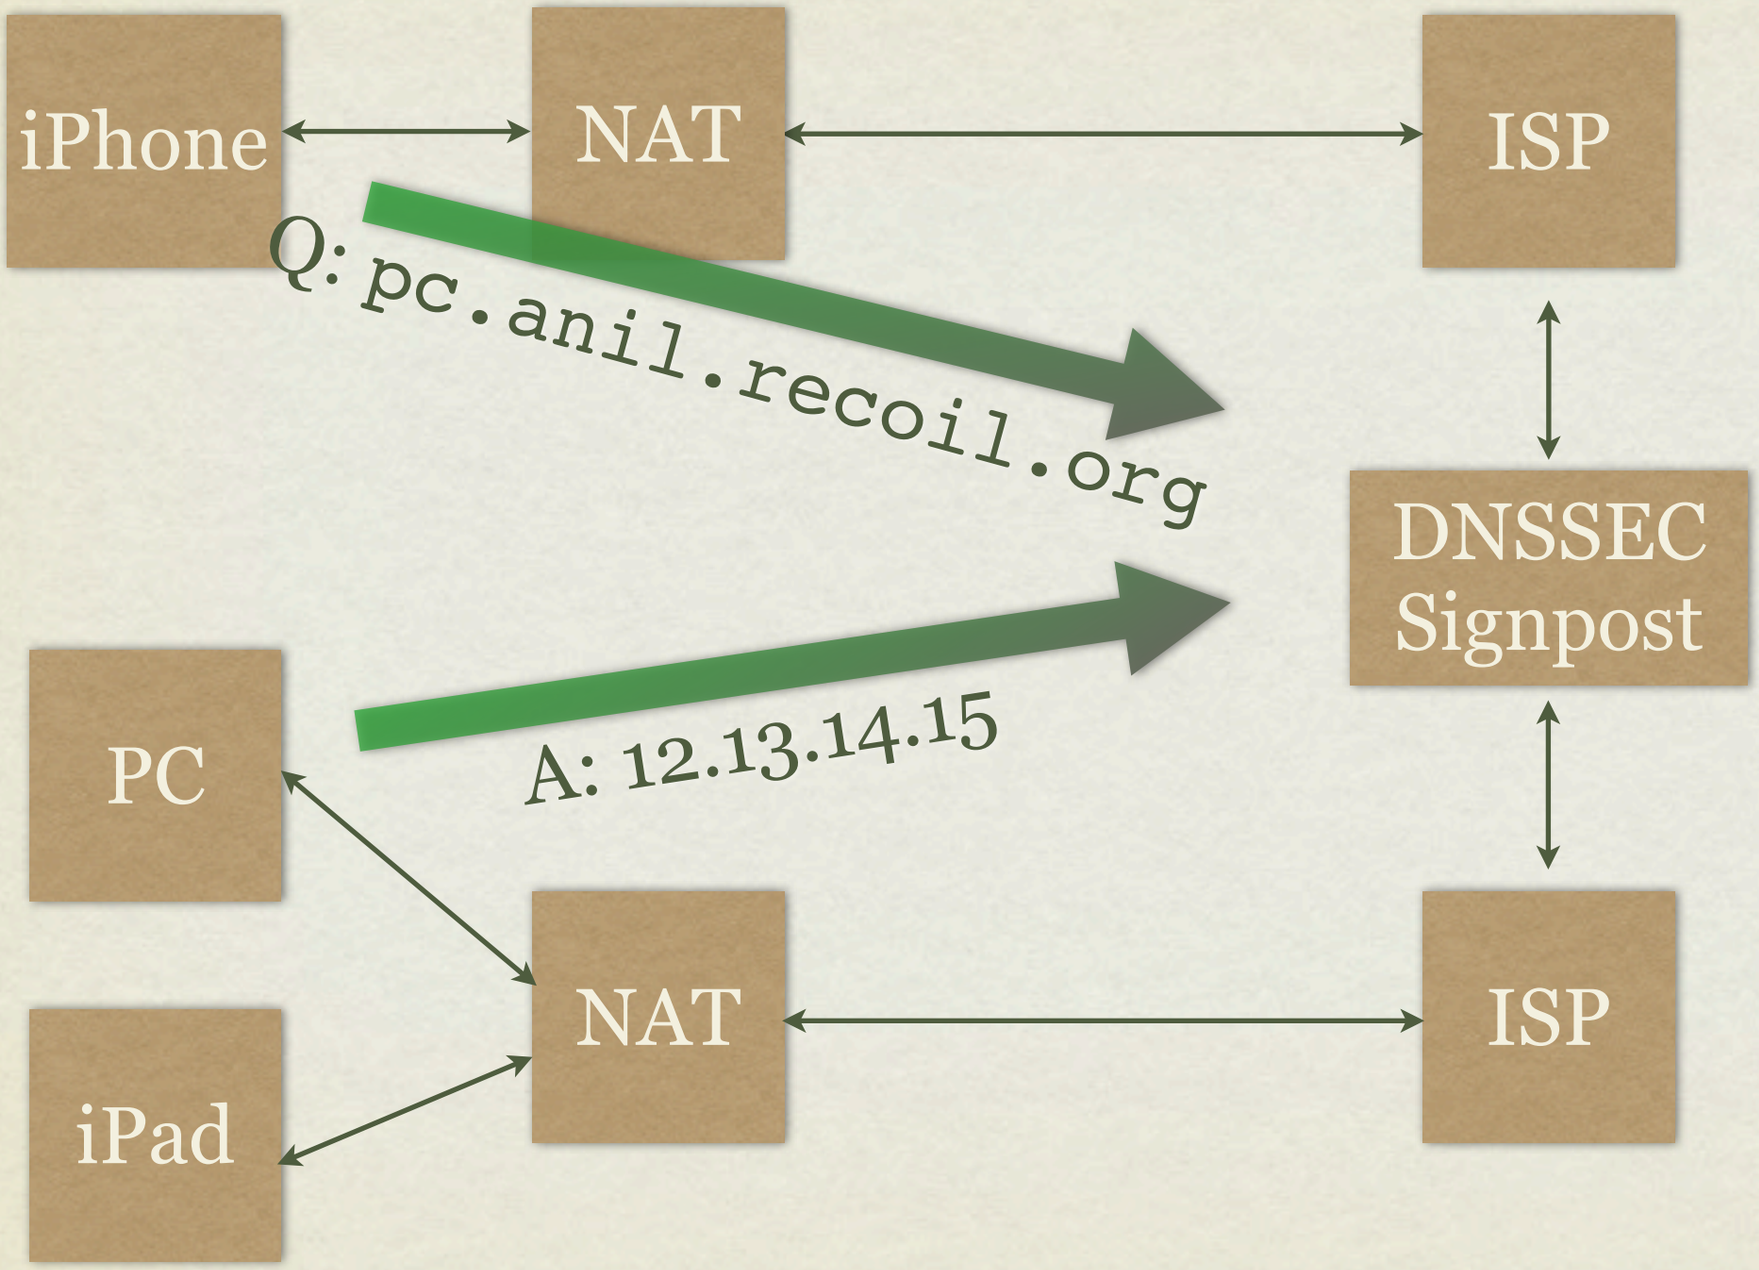

PROBLEM #2: CONNECTIVITY

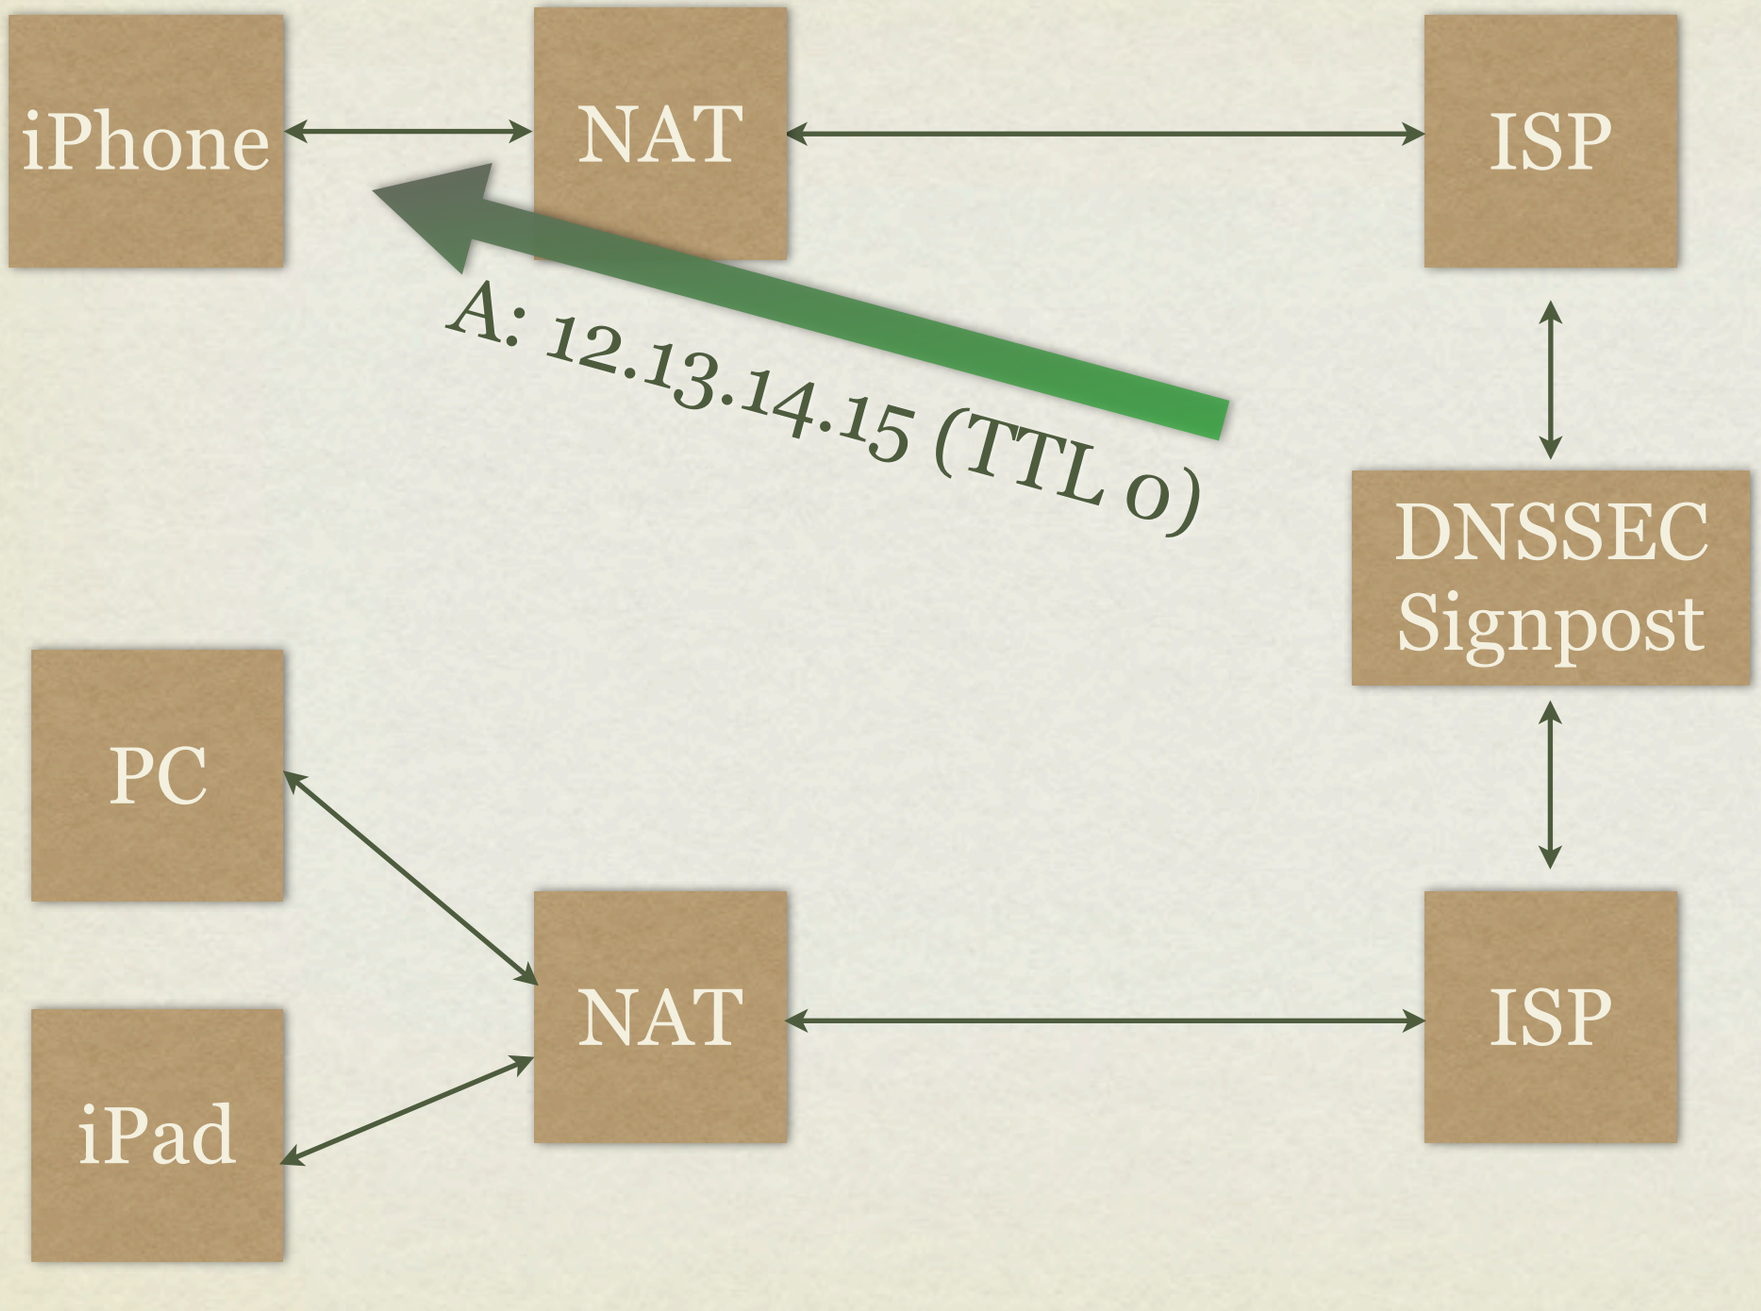

PROBLEM #2: CONNECTIVITY

- **Parallel Routing Tactics for p2p:**
  - **NAT punching:** act as a 3rd party STUN server
  - **UPnP** or other NAT control protocols
  - **Rendezvous** zeroconf discovery of peers
  - **IPSec setup**: VPN (great for “dumb” devices)
  - **HTTP/SMTP proxy:** corporate networks
  - **Wifi hotspot?** IP-over-DNS works (iodine)
  - Last resort: tunnel traffic to the cloud
- Your signpost is the ultimate dirty fighting middleboxer!

PROBLEM #2: CONNECTIVITY

- **“Effectful” name lookups**
  - When a name is looked up, the Signpost executes tactics to discover and establish routes.
  - Tactics form a simple dataflow graph of goals. E.g.:
    - “*ipad* wants to connect to *iphone*”
    - “*iphone*” requires a VPN tunnel **or** a NAT punch
    - attempt NAT punch - FAIL
    - attempt VPN setup - SUCCESS. Return IP to “*iPad*”.
- Tactics are composed via *functional reactive programming*. Lets us inspect *why* a route exists based on successful tactics.

PROBLEM #2: CONNECTIVITY

- **Routing tactics can auto derive other security keys from global public key!**
- **L2:** Ethernet authentication (802.1X), WPA certificates
- **L3:** IPSec, L2TP, OpenVPN
- **L4:** SSL (Notaries), TCPcrypt
- **L7:** HTTPS (Google Chrome), SSH (RFC4255), IMAP, CalDAV, WebDAV
- **“L8”:** Browser passwords, file encryption

PROBLEM #2: CONNECTIVITY

anil
iPhone

anil.recoil.org
Signpost

seb
PC

seb
iPad

seb.eide.name
Signpost

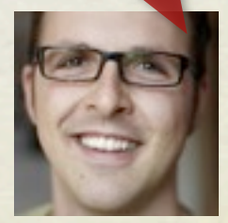

PROBLEM #2: CONNECTIVITY

anil
iPhone

Q: `ipad.seb.eide.name`

anil.recoil.org
Signpost

seb.eide.name
Signpost

seb
PC

seb
iPad

PROBLEM #2: CONNECTIVITY

anil
iPhone

A: 190.10.23.45

anil.recoil.org
Signpost

190.10.23.45

seb.eide.name
Signpost

seb
PC

seb
iPad

VPN

PROBLEM #2: CONNECTIVITY

anil
iPhone

anil.recoil.org
Signpost

A: 190.10.23.45

anil.recoil.org
Signpost

VPN

seb
PC

seb
iPad

190.10.23.45

seb.eide.name
Signpost

seb.eide.name
Signpost

LET A MILLION CLOUDS BLOOM!

- **Why can't we all have our own cloud between our devices and networks?**

- #1: we have no **identity** online.
- #2: lack of end-to-end **connectivity** limits visibility**.**
- #3: who hosts our stuff **reliably**?
- #4: why bother? What **new services** does this enable?

PROBLEM #3: RELIABILITY

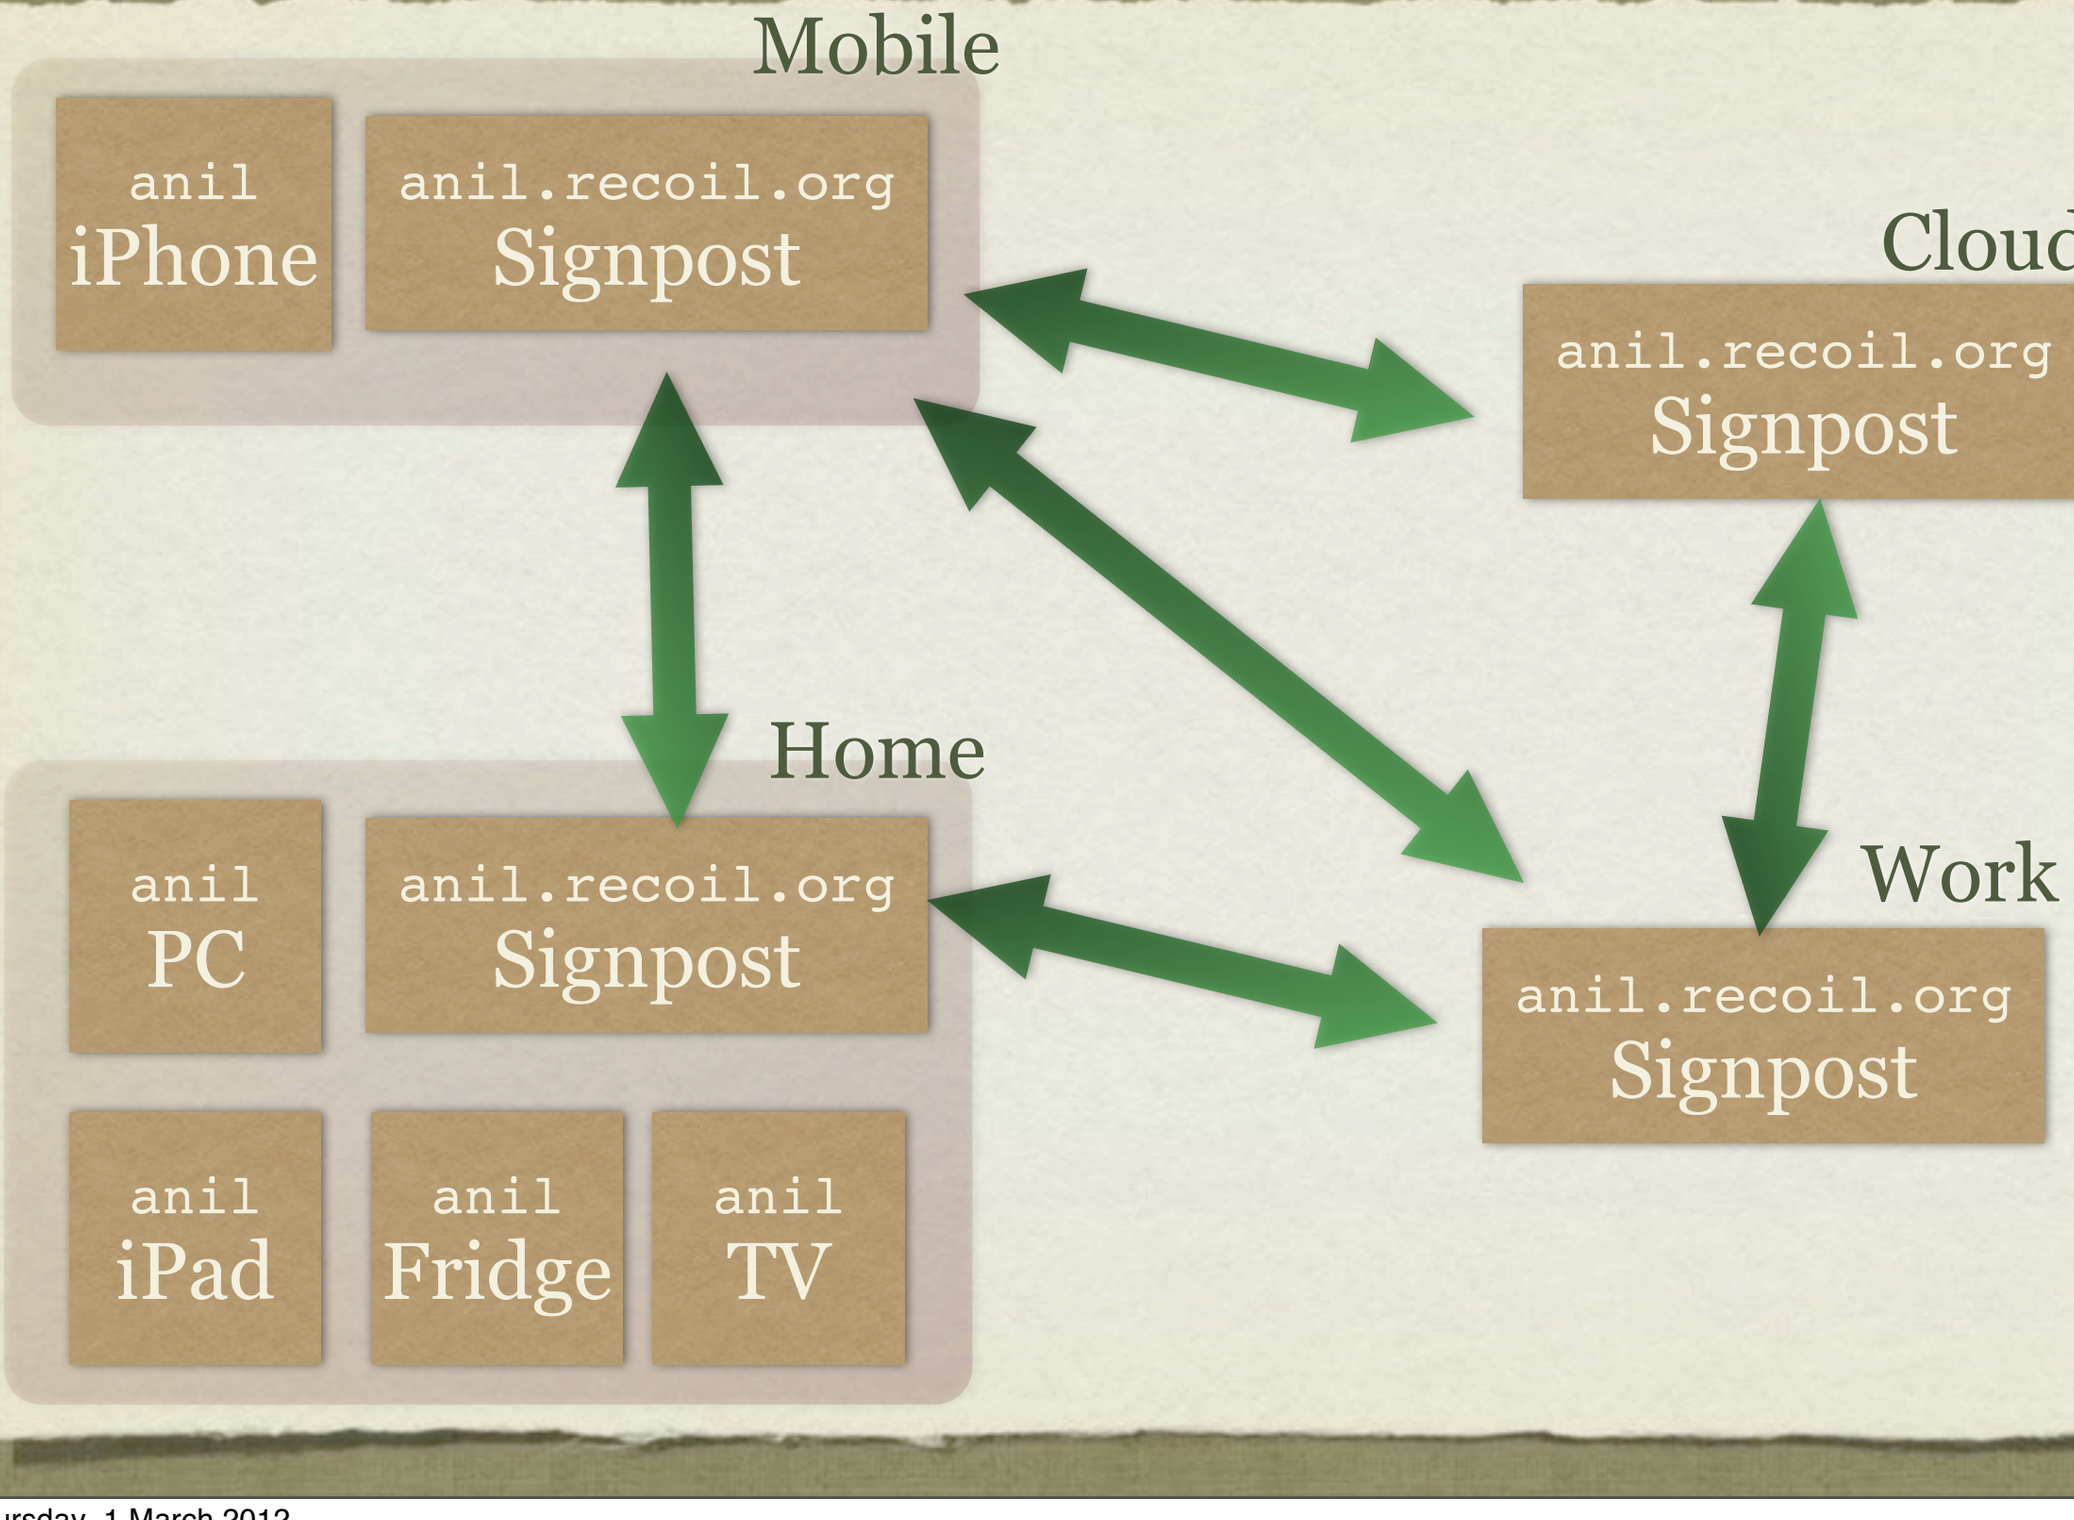

PROBLEM #3: RELIABILITY

- All signposts stay in communication and sync data
  - Eventually consistent lazy replication (Bayou)
    - Natural fit to DNS consistency model
    - Coordination data only: very low bandwidth
- Request resolution is a dataflow graph, where the nodes represent possible tactics (e.g. STUN or route setup).

LET A MILLION CLOUDS BLOOM!

- **Why can't we all have our own cloud between our devices and networks?**

- #1: we have no **identity** online.
- #2: lack of end-to-end **connectivity** limits visibility**.**
- #3: who hosts our stuff **reliably**?
- #4: why bother? What **new services** does this enable?

PROBLEM #4: WHY BOTHER?

- **Efficiency:** Apple devices support “sleep proxies” and multicast DNS *http://en.wikipedia.org/wiki/Bonjour_Sleep_Proxy
  - Devices register DNS services records (e.g. iTunes sharing or website) and go to sleep.
  - Router proxy wakes them up (Time Capsule or Airport Extreme).

- **Evaluation #1:** run Bittorrent to share files between two phones. Cycle between two spots in Cambridge: we hit eduroam, 3G, wgb wifi.

PROBLEM #4: WHY BOTHER?

- **Low latency services**, simply not possible with cloud.
  - Sub-millisecond image processing
  - Real-time video stitching (concerts, Olympics)

PROBLEM #4: WHY BOTHER?

- **Low latency services**, simply not possible with cloud.

PROBLEM #4: WHY BOTHER?

- **Low latency services**, simply not possible with cloud.
  - Sub-millisecond image processing
  - Real-time video stitching (concerts, Olympics)
- **Evaluation #2:** multipath video streaming is trivial with Signposts, as they take care of route setup and failover.

PROBLEM #4: WHY BOTHER?

- **Democratise our infrastructure**!
  - Hardware printing now possible (diydrones.com), Arduino, Raspberry Pi.
  - Not practical to hook things up to Twitter and Facebook at scale.
  - Machine-to-machine trust via Signpost gets more secure as it grows *(see Perspectives, USENIX Security)*

- **Evaluation #3:** middlebox probing and enable most efficient path security (TCPcrypt, IPSec).
- *"Policies in the ends, middlebox probing in the middle"*

SUMMARY

- **“An architecture for dynamic routing across distributed clouds via middlebox-controlled context-dependent naming”**

- or: **Network names that “just work”!**

- *Coming soon:* http://signpo.st/ http://github.com/avsm/
- *Related work:* Intentional names (MIT), Named Data Networking, Perspectives, Internet Indirection Infrastructure (I3)