

Talking about vs. talking to: effective networking in a natural disaster

Van Jacobson
Research fellow, PARC
van@parc.com

MIT Communications Futures Project
Cambridge, MA

23 October 2008

“Timely, accurate, impartial information is central to saving lives and strengthening recovery; the power lies in its effective management, analysis and application...”

– United Nations Office for the Coordination of Humanitarian Affairs (UN-OCHA), Global Symposium on Information for Humanitarian Action, Geneva, Switzerland, 22 October 2007

- Communication is crucial in a disaster but data communications didn't work during the 2004 tsunami and 2006 Pakistani earthquake.
 - Wide-area comm infrastructure was one of the first things to break.
 - Responders couldn't agree on who should be the communications hub and transit provider(s).
- I believe an architectural problem with our networking model is the root cause.

Inventory tracking today



Inventory tracking today



- Server has to be connected to every network 24/7
- Every plane, truck and clinic must be configured to talk to server(s)
- ... and must have trust relationship

Something that did work: voicemail on wheels



Objectives matter

“I need the medicine inventory”
can be *much* easier to accomplish than

“I need to talk to server X”

In general, the chance of success is maximized when applications and their networking component solve the same problem.

Intrinsic Incompatibility?

Applications require data.

Networks require data's location.

Is data location essential?

Core idea: communicate via named data not endpoint identifiers

The web gave us named data (URI/URL) so ...

- Receiver expresses interest in some collection of data.
- Sender(s) with matching data respond.
- Receiver validates relevance and integrity via information in the response.

CCN Inventory tracking



CCN Inventory tracking

d762/i/lv/oxfam/h9/123: #3456789



d762/i/tk/unhcr/t41/4317: #3456789



CCN Inventory tracking

d762/i/lv/oxfam/h9/123: #3456789



d762/i/tk/unhcr/t41/4317: #3456789



d762/i/lv/unhcr/t41/4317: #3456789



d762/i/tk/unocha/w31: #3456789

CCN Inventory tracking

d762/i/lv/oxfam/h9/123: #3456789



d762/i/tk/unhcr/t41/4317: #3456789



d762/i/lv/unhcr/t41/4317: #3456789



d762/i/tk/unocha/w31: #3456789

- Aggregators broadcast interest
- Planes, trucks, clinics, ..., broadcast transaction reports
- Reports follows 'interest' breadcrumbs to aggregators

- With CCN, data producers and consumers require no configuration yet data shows up everywhere there's interest (and nowhere else).
- Communication incentive is built-in -- entities are announcing that they're doing their job.
- Entire system is transparent and cross-verifiable.

- Cross-validation requires some verifiable notion of identity.
- ‘Identity’ usually involves hierarchical PKI which doesn’t work well in unstructured, heterogeneous, dynamic environments.
- Can we generate a useful, robust identity from intrinsic properties of the environment?

- The environment is highly diverse but communication intensive and richly interconnected.
- All CCN content is cryptographically signed and producer/consumer incentive structure rewards maximum information in signatures.

I.e., any consumer discards data that it does not believe is valid, relevant or trustworthy. Burden of proof is on the producer.

Content integrity incentive is tied directly to the act of communication

- Consumer wants content that is relevant (matches expressed interest) and valid.
 - ▶ has to announce its interests
- Producer wants its content consumed.
 - ▶ has to respond to interests
 - ▶ has to augment response so consumer(s) can prove relevance and validity.

User supplied name

Versioning &
segmentation

Content proxy
(e.g., SHA256
checksum)

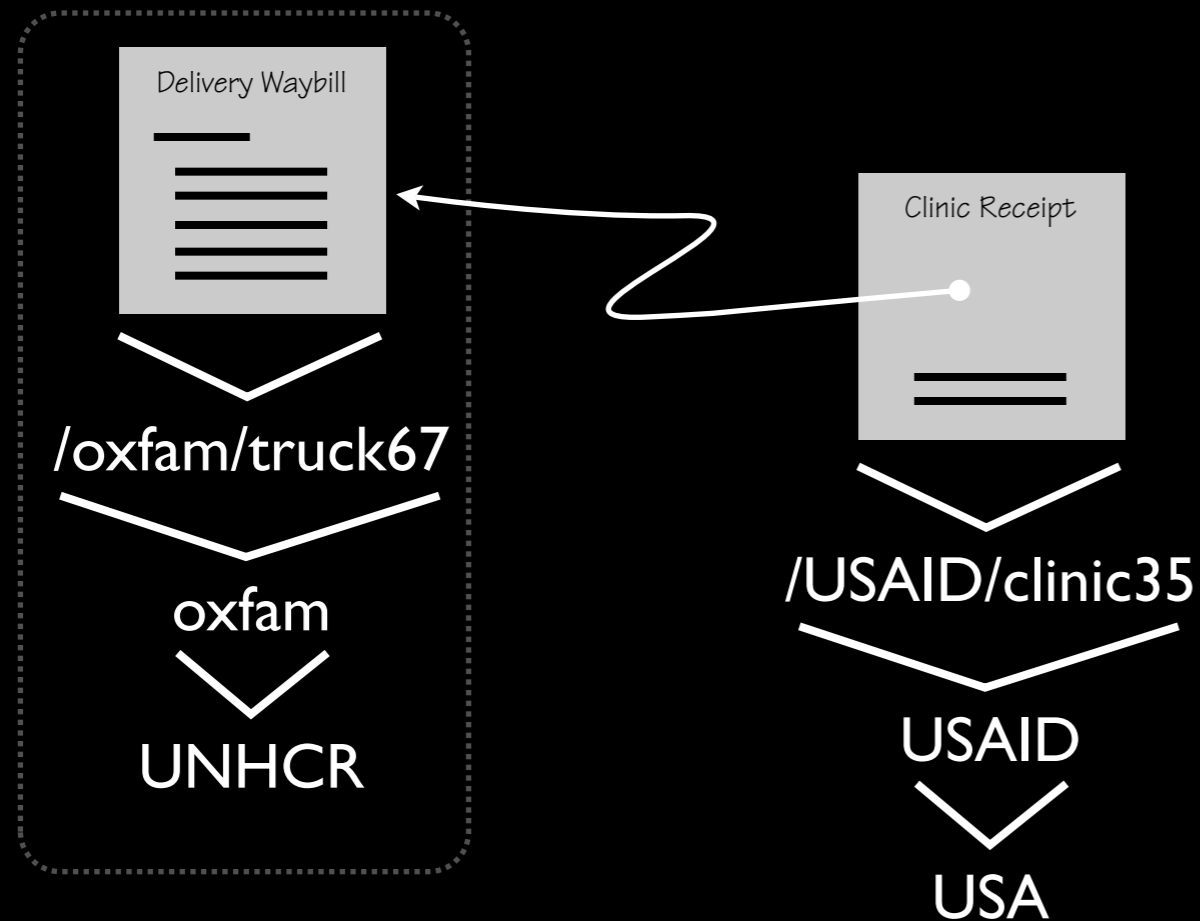
/d762/i/oxfam/h9/#101/v0/s0/0x3fdc96a4...

Signed by /oxfam/h9

Signed by oxfam

This binding is *immutable*
(the data associated with the name can't change)

Evidentiary identity



- Content you generate attests to your affiliation and contacts.
- Content that refers to your content attests to your productivity, reliability, longevity, ...

Like Google Page Rank, your value is not based on your content but on others' references to it.

Information sharing

- An organization expresses sharing policy as *“information X cannot be disclosed to Y ”*
- A network expresses it as *“host p cannot talk to host q ”*
- This only works if information can be made isomorphic to host identity
- Nearly impossible to do this within an organization; impossible in a coalition.

- Encryption allows location (where the bits are stored) to be decoupled from privacy (who has the key to decrypt them).
- But since conventional networking doesn't have named data it's *extremely* difficult to associate keys with content.
- CCN naming hierarchy makes it easy to algorithmically associate key with content.
- Since data can be in multiple hierarchies, a CCN name can collect exactly the data to be shared with some coalition partner.