CCNx Nameless objects

Marc Mosko (marc.mosko@parc.com)

ICNRG (San Francisco, USA) October 3, 2015

parc®
A Xerox Company

Outline

- Motivation

- Overview of Nameless Objects

- Use with Manifests

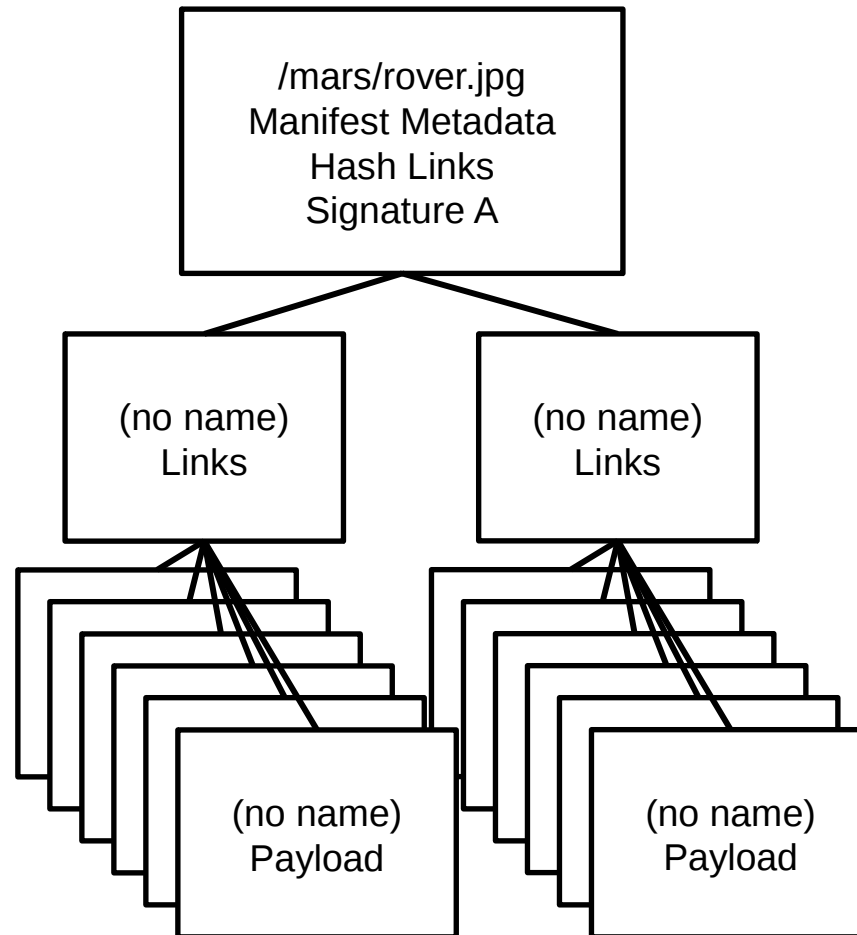- Conclusion

parc
A Xerox Company

Motivation

- How to position content in multiple locations
  - Without renaming it
    - Expensive to rename and re-sign. Requires original publisher to re-sign or some sort of trust delegation.
  - Without wrapping it
    - Similar to renaming.
  - Without changing routing
    - Requires changing global state
    - Who's going to authorize random cache to advertise someone else's name?
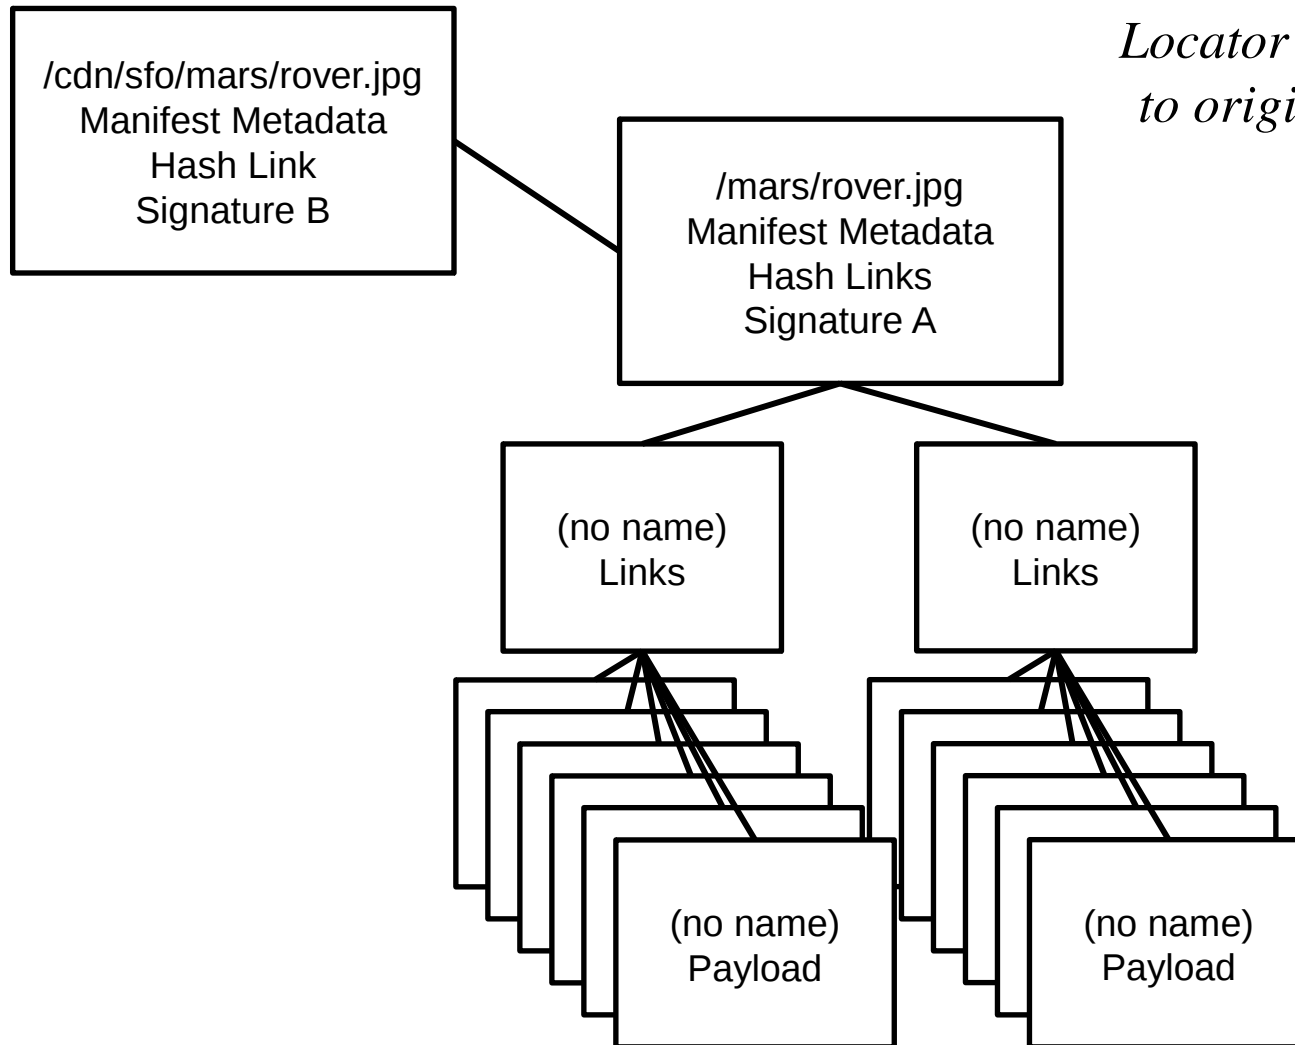
parc
A Xerox Company

Overview of nameless objects

- CCNx Name can be a locator or identifier
  - We want to use Identifiers for user content but Locators (to nearest replicas) to fetch it.

- CCNx Content Object bakes a name in to an object
  - Because it is either (a) signed or (b) accessed by hash based name.
  - So, it can have either an Identifier or a Locator, but not both.

- Nameless objects
  - Allows a publisher to publish majority of content without a Name.
  - Publisher must provide a list of hash-based names.
  - Interests carry a locator name plus the hash, will match the nameless object.
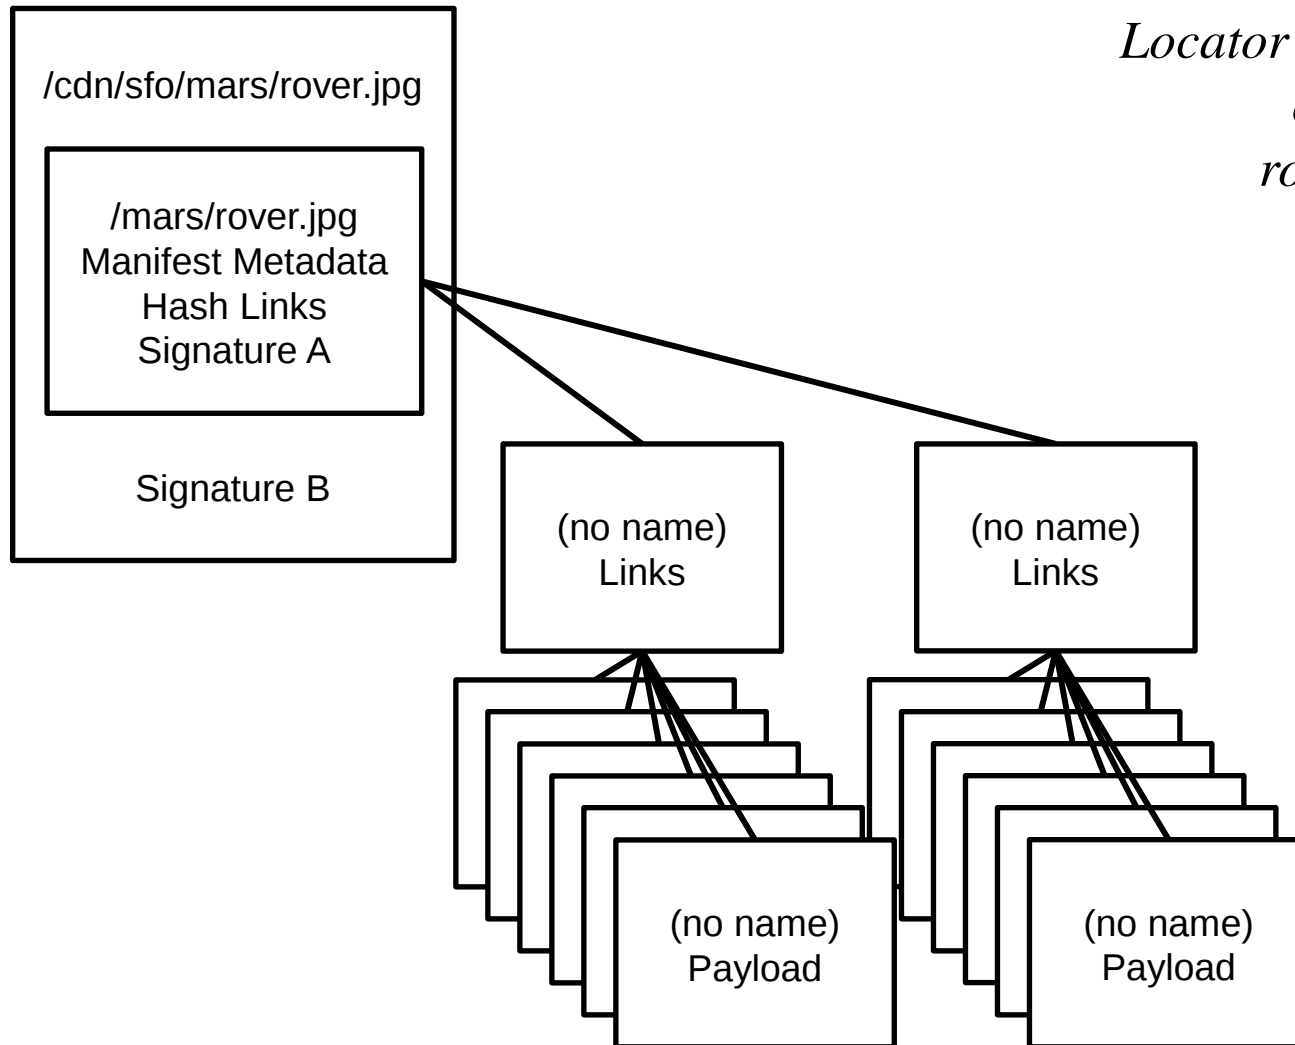
4

Usage with manifests

/mars/rover.jpg
Manifest Metadata
Hash Links
Signature A

(no name)
Links

(no name)
Links

(no name)
Payload

(no name)
Payload

parc
A Xerox Company

Usage with manifests

/cdn/sfo/mars/rover.jpg
Manifest Metadata
Hash Link
Signature B

/mars/rover.jpg
Manifest Metadata
Hash Links
Signature A

(no name)
Links

(no name)
Links

(no name)
Payload

(no name)
Payload

Locator name by reference to original root manifest.

6

parc
A Xerox Company

Usage with manifests

/cdn/sfo/mars/rover.jpg

/mars/rover.jpg
Manifest Metadata
Hash Links
Signature A

Signature B

*Locator name by inclusion
of original
root manifest.*

(no name)
Links

(no name)
Links

(no name)
Payload

(no name)
Payload

7

parc
A Xerox Company

Rules

1) Receive Interest

- The Interest must have a name.

- Try to satisfy from CS.

- Create PIT entry.

  – If Interest has a HashRestriction, must be able to match PIT entry without a Name.

- Aggregation rules apply as normal and now include additional key by only HashRestriction.

- Apply normal forwarding rules.

2) Receive ContentObject

- ContentObject has a Name.

  - Apply normal forwarding rules.

- Otherwise

  - Satisfy any PIT entry that matches on Hash alone.

  - Satisfy any PIT entry that matches on (KeyId, Hash) alone.

3) Satisy From CS

- Nameless object can only satisfy an Interest with a HashRestriction.

- Otherwise, normal rules apply.

8

conclusion

- Nameless Objects provide

    - A way to place content on any replica with minimum renaming/resigning.

    - The consumer still sees the original publisher signature on the hash tree, so will trust those hashes from that trust anchor.

- Processing Changes

    - Interest path must have way to save PIT entries (Hash) and (KeyId, Hash) without a name.

    - ContentStore only responds with a Nameless object if the Interest has a Hash restriction.

    - ContentObject path must be able to lookup in PIT without a Name.

parc®
A Xerox Company