

Contextualization of Resolution (c15n) BoF

49th IETF, San Diego

December 12, 2000

The problem

- (UR) identifier resolution infrastructure
- End users/clients have a stake in expressing the policies they want applied in resolution
- For example, the "appropriate copy" of a resource – preferably finding a "local" copy of a journal rather than (re)purchasing one from the authoritative publisher.
- Software systems being built to solve this in today's deployed systems are using specialized, non-interoperable, non-scalable approaches.

Resolution – 50k foot

- Implementation of a transaction between a resource provider and a content seeker
- Valid transactions:
 - repeatable end to end: every seeker gets same bits directly from provider (origin)
 - provider tailors (content/location) to individual seeker
 - seeker consciously(*) overrides provider's policies (proxy)
- Treacherous transaction:
 - neither seeker nor provider can explain what happened in the transaction – policy implemented by an unidentified, unknown (to endpoints) intermediary

(*) Grey area: may include service agreements.

How is c15n different than...

- CDN's are about letting providers preferentially distribute their content and manage access to it
- (Transparent) proxies implement (local) network policies for replication of content
- Returns pointer or pointers, not resource

Existing contextualization work

- The SFX project described at:
 - http://www.dlib.org/dlib/october99/van_de_sompel/10van_de_sompel.html
 - http://www.doi.org/workshop_19sep00/doi_wkshp_0900_11version.ppt
 - Uses HTTP cookies
 - (web) client asks the global resolution for an identifier
 - sends a cookie which is a key for the appropriate context
 - global system uses this key to redirect the query to the appropriate local knowledge server (address), which makes the judgment about where an appropriate copy of the resource shall be obtained.

Desired Solution

- part of (URI) resolution infrastructure, not middle-box support structure
- application-independent
- standardized

Top down vs. bottom up

- Top down:
 - assigner designates only authoritative result
 - everyone gets the same result, always, everywhere
- Bottom up:
 - user (agent) of resolver controls parameters and out come of resolution
 - what you get depends on where you are in the net

Contextualized

- Ability to respect the authoritative chain
 - Incorporate elements of “context” in resolution, eg:
 - where you are in the net
 - current hardware parameters
 - language
 - Uniform mechanism for expressing context and finding necessary local information
- N.B.:* may have multiple different “contexts” at different times in the same network.

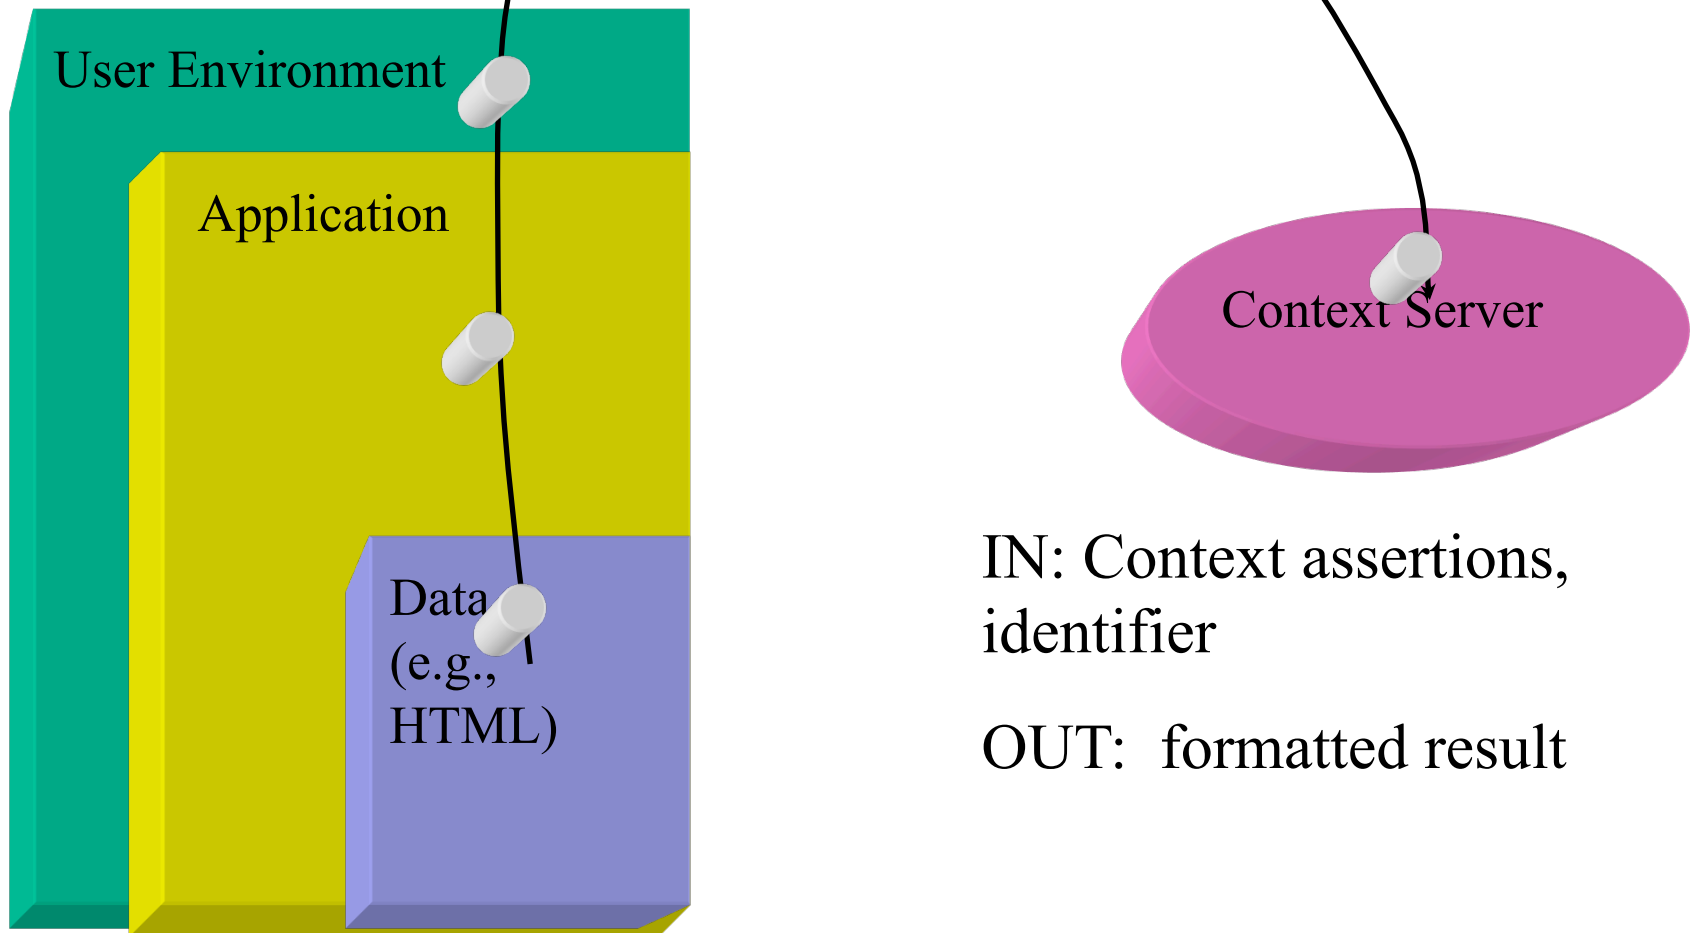
Scope and applicability

- At least all URI
 - beyond http, web
- Potentially elsewhere
 - e.g., looking up “human-friendly” names to map to DNS entries

Why standardize?

- Enable incorporation into all application-level infrastructure
- Within a given application, many different context services may be accessible (e.g., for library resources, localized identifiers, etc).

Graphically...



I/O

- In
 - context assertions
 - core fixed vocabulary
 - extensible
 - target identifier
 - URI
 - other
- Out
 - Formatted response
 - URI
 - URI set...

Examples of context assertion vocabulary

- location id
- user id
- language
- locale
- capabilities
- cost
- etc

The part we want to solve

- protocol to the context server
- CNRP-like, with more explicit match
- RESCAP?

Issues

- Fixed core vocabulary
- URI or more general
- Protocol vehicle
- Context server discovery
 - top-down (authoritative)
 - locally-configured
 - “yes”

Draft WG charter – Problem Statement

The URN WG's Dynamic Delegation Discovery System (DDDS) describes a generalized architecture for 'top down' resolution of identifiers such as URIs. This works well when a (software) client wants or needs to dynamically determine the explicit authoritative delegation of resolution. However, there are times when it is desirable to incorporate other elements of contextual control information in determining, for example, the "appropriate copy" of a resource -- preferentially finding a "local" copy of a journal rather than (re)purchasing one from the authoritative publisher. This is generally applicable to all URI resolution, but it is more specific than "web caching". Software systems being built to solve this in today's deployed systems are using specialized, non-interoperable, non-scalable approaches.

Draft WG charter – Problem Statement

Cont' d

This working group will develop a protocol for passing an identifier and context assertions to a context server, in order to determine an identifier (or a set of identifiers) that the context server deems a “best fit” for the expressed context.

Draft WG charter – Milestones

- Feb 2001 Goals & Requirements draft
- Mar 2001 (Context) Core vocabulary draft
- Mar 2001 Vehicle draft document
- Apr 2001 Sample application integration (e.g., HTTP) draft
- Apr 2001 Goals & Requirements to RFC
- May 2001 Core vocabulary doc to RFC
- Jun 2001 Sample application integration to RFC