Privacy Policy Expression Languages - An analysis

W3C Workshop,
Ispra, October 2006

Co-authors:

Paul Madsen, NTT
Marco Cassasa Mont , HP Labs
Robin Wilton, Sun Microsystems

Aims of the Paper

- Identify a set of actors and flows relating to online privacy;
- Classify the points at which privacy policy could apply in principle;
- Look at candidate policy expression languages and their scope;
- Match this against the policy enforcement points identified above;
- Look at Liberty's ID-WSF framework in this context.

We did *not* set out to:

- Perform an exhaustive or in-depth survey;
- Recommend one candidate over others;
- Design an end-to-end privacy preference solution.

The Basic Elements

- What are 'privacy' and 'policy' in this context?
- Who is involved in online privacy enforcement?
- What flows and interfaces does this imply?

- What conclusions can we draw about privacy as a set of relationships, roles, rights, obligations and expectations?
- Are there gaps in either the model or the current implementation picture?

High-level View of Actors and Flows

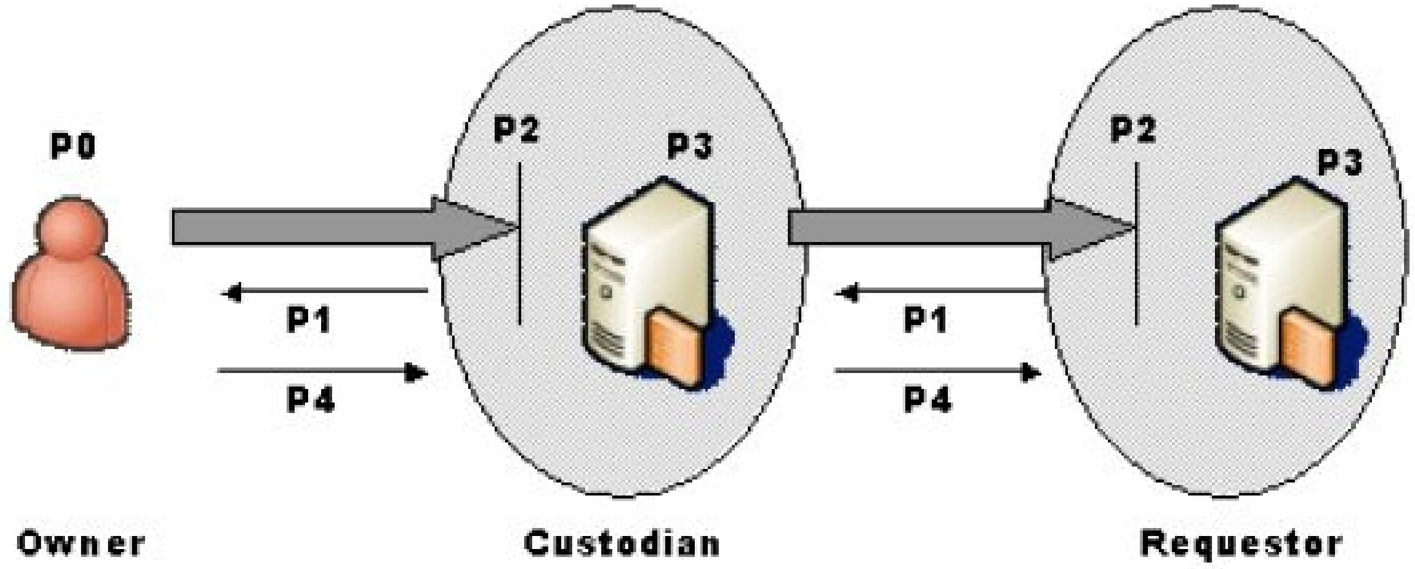

Implied policy control points:
P0: User expresses a preference
P1: Custodian expresses a 'promise'
P2: Custodian enforces acceptance conditions
P3: Custodian enforces usage conditions
P4: Custodian or owner asserts 'transitive' conditions

Note that P0, P2 and P3 have 'local' scope;
P1 and P4 are 'cross-domain'

Mapping current options onto this model

	P0	*P1*	*P2*	*P3*	*P4*
P3P	No	Yes	No	No	No
XACML	No	No	Yes	Yes	Yes
EPAL	No	No	No	Yes	No
ODRL	No	No	No	No	Yes

Notes:
1 – 'No' means 'not optimised for', rather than 'unsuitable for';
2 – Liberty's ID-WSF potentially supplements the existing options by adding the means to enforce policy for P1 and P4 flows, using a <UsageDirective> element in the SOAP header of a message;
3 – This mechanism can be used in support of a 'purpose of usage'/ 'conditions of usage' approach as described above;
4 – Such an implementation would still assume that some other component is enforcing P3 policy expressions.

Initial conclusions

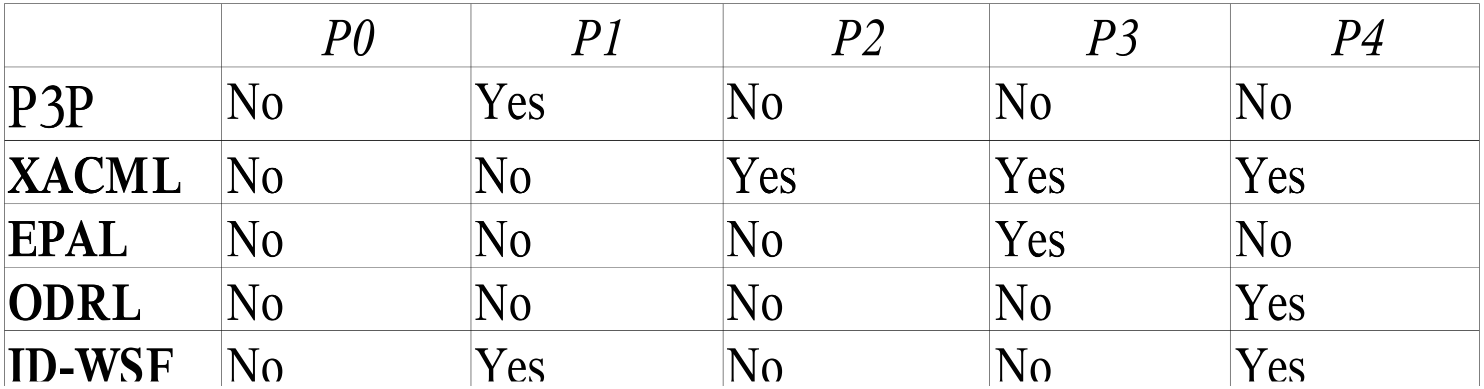

	P0	*P1*	*P2*	*P3*	*P4*
P3P	No	Yes	No	No	No
XACML	No	No	Yes	Yes	Yes
EPAL	No	No	No	Yes	No
ODRL	No	No	No	No	Yes
ID-WSF	No	Yes	No	No	Yes

A combination of XACML (for P3 enforcement) and ID-WSF carrying XACML payloads (for P1-P4 enforcement) appears viable, and usefully extends the coverage of the model.

Discussion Notes - p1

Because “privacy” is such a personal/introspective concept, we tend to start from the assumption that it relates to something entirely under our control. But of course, it only starts to have meaning as a concept from the moment we divulge (or expose) something to someone else.

At that point, we have to rely not just on our control of that information, but our ability to influence what someone else might do with it.

This suggests three possible 'points of control':

1 – Personal control (P0)

2 – Expression of preferences (P1)

3 – Enforcement of preferences (P2, P3, P4)

Discussion Notes - p2

At this point it is also worth asking why we habitually refer to 'preferences' rather than, say, 'requirements'...

“I would prefer that you don't tell anyone how old I am” does not mean the same as “I will only reveal my age to you on condition that you don't reveal it to anyone else”.

This also correctly introduces the idea that any 'control' we wish to retain over what someone else does with our information must cater for the possibility that that person could divulge it to a third party.

We then have to be able to decide whether the three 'control points' are necessarily best implemented in technology or by other means, for any given context of legislation, regulation, best practice and functional requirements.

Discussion Notes - p3

It is useful to construct a basic privacy model as follows:

Data Subject -> Data Custodian -> Data Consumer

while remembering that the roles of Data Custodian and Data Consumer often conflate into the same entity.

This model makes explicit the 'collection' and 'usage' steps.

'Benefit' flows around such a model in various ways, reflecting the fact that we rarely (intentionally) surrender privacy for absolutely no purpose; usually there is some *quid pro quo*, even if the cost and benefit are not always explicit or commensurate.

Is there a Privacy Preference 'Sweet Spot'?

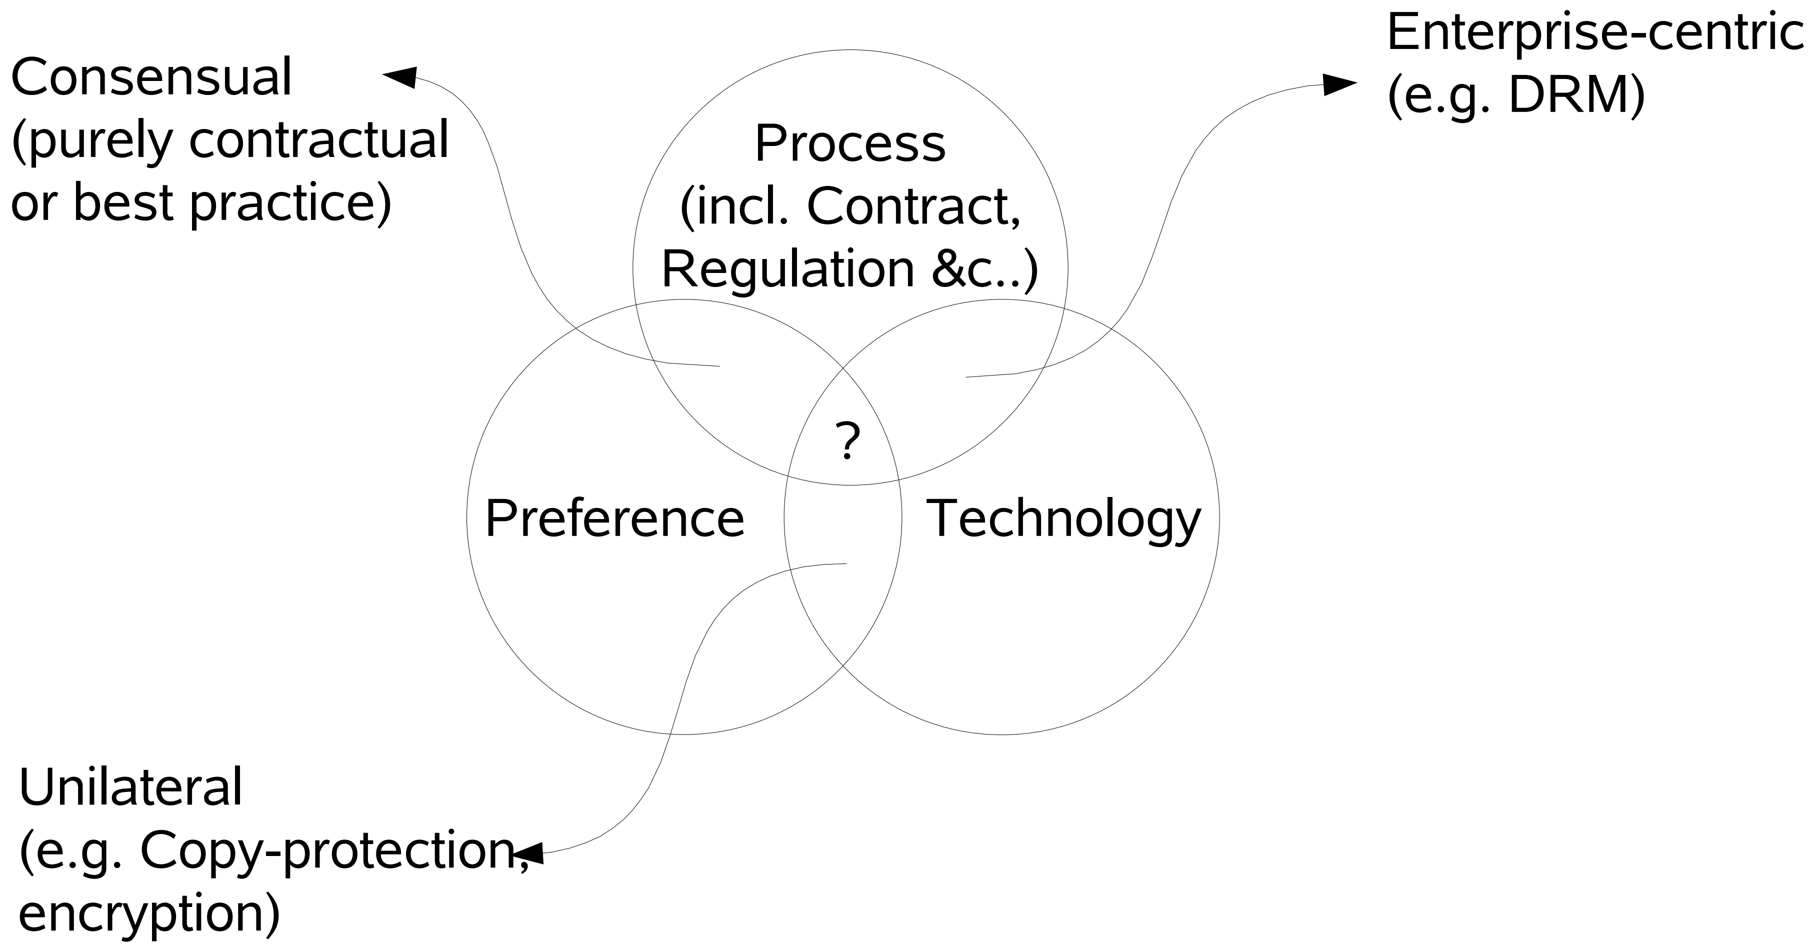

Thank you

Robin Wilton, Sun Microsystems
robin.wilton@sun.com
+44 705 005 2931
http://blogs.sun.com/racingsnake

A recent article...

“Liberty Alliance enables consumer privacy controls”

The Liberty Alliance, whose members include service providers like AOL and American Express, has unveiled the final version of the Identity Web Services Framework (ID-WSF) 2.0, InternetNews reports.

ID-WSF 2.0, a schema for building secure, interoperable Web services that can be piped over the Internet, was originally developed to transmit secure transactions between large businesses, the report says.

The schema now takes social-networking sites into consideration, allowing consumers and organisations to establish privacy controls, enabling them to guard against phishing, pharming and other forms of online identity theft while they share information on social networking sites, it says.

Source: SecurityNewsPortal.com (12 Oct 2006)