



Jena rules experiences and implications for rule use cases

Dave Reynolds
HP Laboratories

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Framing

- Focus on rule languages for the semantic web
- Based on experience with Jena rules system look at:
 - spectrum of use cases noticed
 - inferred need for standards-based interoperability
 - implications for expressivity of rules language

Jena rules experiences

Jena rules:

- forward chaining (RETE) and back chaining (tabled LP) rules engines, targeted at RDF processing
- internal usage (RDFS, partial OWL) and user tool
- Jena user base ~20k, rules introduced mid 2003

Observations from user group traffic and HP usage:

- significant usage, split between standalone RDF processing and rules layered on top of RDFS/OWL
- use of procedural extensions
- little evidence of rule interchange other than reuse of supplied RDFS/OWL rules
- wide mix of apparent use cases ...

Use cases noted:

1. Deductive rule usage



<u>Use case</u>	<u>Interoperability?</u>	<u>Expressivity implications</u>
view transformation, abstraction and integration	<ul style="list-style-type: none"> • share transform libraries 	<ul style="list-style-type: none"> • semi-positive • object introduction • property variables • (recursive data structures)
implementing semantics of shared vocabularies e.g. SKOS	<ul style="list-style-type: none"> • reuse of implementations • definition of agreed subsets? 	<ul style="list-style-type: none"> • equality
application specific semantics	<ul style="list-style-type: none"> • tool independence • endorsement of approach 	
expressive KR on top of/along side RDFS/OWL	<ul style="list-style-type: none"> • reuse of expressive KBs 	<ul style="list-style-type: none"> • full OWL semantics or layering? • wf negation?

Use cases noted:

2. Integrity rule usage



<u>Use case</u>	<u>Interoperability?</u>	<u>Expressivity implications</u>
validation of dataset prior to further processing	<ul style="list-style-type: none">•publish agent preconditions	<ul style="list-style-type: none">•semi-positive•denial rules or reporting actions
Policy expression (e.g. access control)	<ul style="list-style-type: none">•unclear	<ul style="list-style-type: none">•richer reporting actions•non-monotonicity?

Use cases noted:

3. Reactive rule usage



<u>Use case</u>	<u>Interoperability?</u>	<u>Expressivity implications</u>
Event triggering (e.g. fault detection, SLA violation)	<ul style="list-style-type: none">• unclear	
Notification	<ul style="list-style-type: none">• public subscription interface	<ul style="list-style-type: none">• non-recursive could simply be a persistent SPARQL query

Conclusions

- Significant usage of rule processing already in semantic web apps:
 - indirect reason for standardisation (tool independence)
- Example areas with particular interop needs:
 - transformation libraries
 - shared implementation of vocabularies (OWL fragments, SKOS)
 - pre/post condition for remote processors (e.g. semweb services)
 - (more expressive KR)
- Expressivity observations
 - need for RDF level processing (no DL restrictions, property variables, bNode introduction)
 - several interop use cases exploit closed-world negation
is semi-positive sufficient or do we need full non-mon NAF?



i n v e n t