

Implementation of Resource Namespace Service

Masahiro Nakamura, Osamu Tatebe
University of Tsukuba


Resource Namespace Service (1)



- RNS lets you map any resource into single, hierarchical namespace
- Resources are referred to in a form of EndpointReference (WS-Addressing)
- RNS Specification is published as GFD-R-P.101

<http://www.ogf.org/documents/GFD.101.pdf>

Resource Namespace Service (2)



- Hierarchical namespace management that provides name-to-resource mapping
- Basic Namespace Component
 - **Virtual Directory**
 - Non-leaf node in hierarchical namespace tree
 - **Junction**
 - Name-to-resource mapping that interconnects a reference to **any existing resource** into hierarchical namespace

Operations

add (string: name, EPR: reference, . . .)

list ([string: regexp])

update ([EPR: parent, string: name,
EPR: reference, . . .])

query ()

remove ([string: name])


- All operations take implicit parameter in SOAP header: EPR of target entry

Defining Schema / WSDL

- Schema and WSDL are essential to define a web service
- In the schema, we define each entry type, list type of entries and fault types
- WSDL describes message types, etc.
- WSDL 2.0 is available, but no stable parser available yet. Use WSDL 1.1 instead.

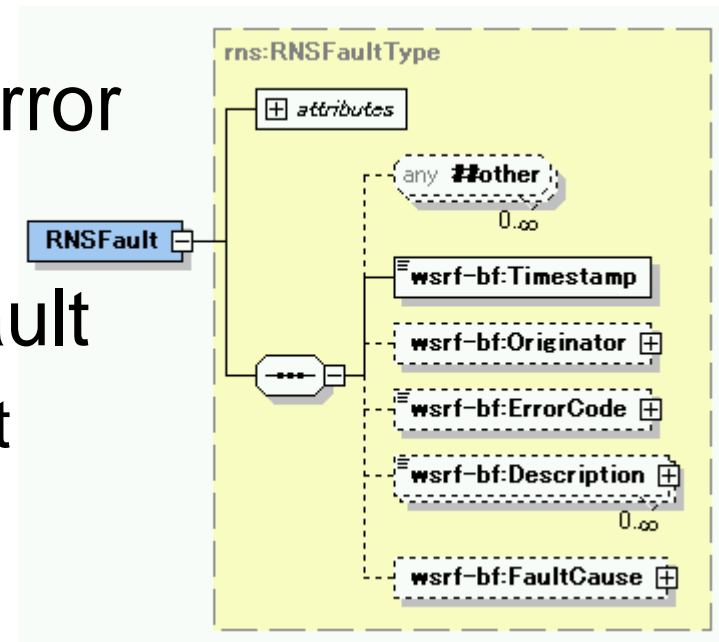
Entry

- Each entry has
 - one entry_name
 - any number of entry_references
 - any other properties (extensible elements)



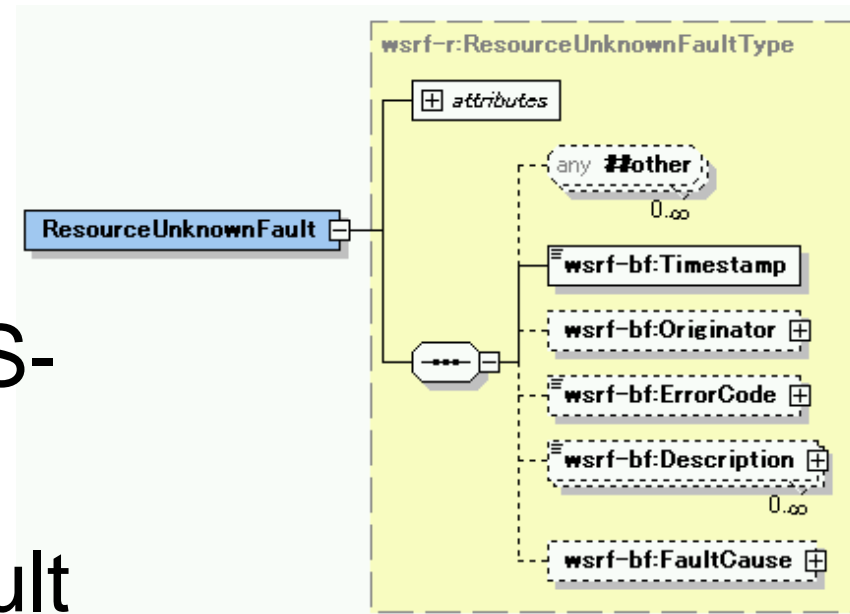
Faults (1) - RNSFault

- base fault class that extends WS-BaseFault
- used for internal server error
 - principally SQL errors
- “Path” element in RNSFault
 - specified in RNS spec, but considered not to be necessary



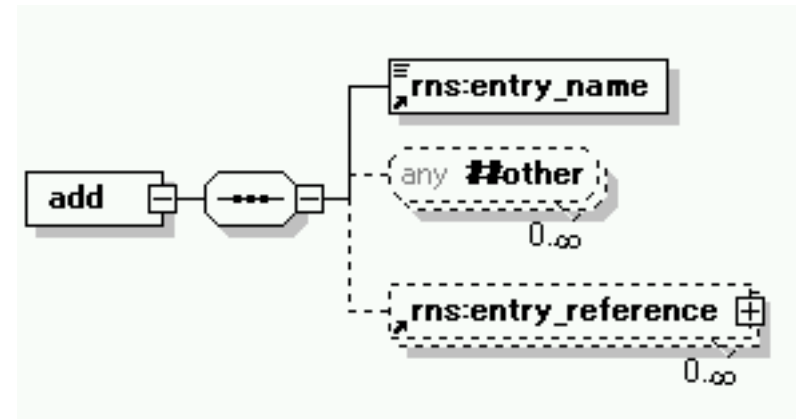
Faults (2) - ResourceUnknownFault

- The spec does not specify a “not found” fault
- Defined in OASIS WS-Resource
- extends WS-BaseFault



Operations (1) - add

- Adds an entry within a virtual directory
- Input parameters
 - (implicit) EPR of a parent virtual directory
 - entry_name
 - {any}^{*}
 - entry_reference^{*}
- Return value
 - entry_self



add(2)

- if no entry_reference is present in a request message, add a virtual directory
- otherwise, add a junction

```
<xsd:element name="add">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element ref="rns:entry_name" />
      <xsd:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded" />
      <xsd:element ref="rns:entry_reference" minOccurs="0" maxOccurs="unbounded" />
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>
```

Demonstration

- 5 operations of RNS
 - add, list, query, update, remove
- SOAP messages

RNS Implementation (1)

- Using JAX-WS (Java API for XML Web Services) as a web service framework
 - Apache Axis2 1.3 has a bug and unable to compile our WSDL
- JAX-WS lacks support of handling WS-Addressing on client side
 - Manually adding headers for now
 - Planned for the future releases of JAX-WS

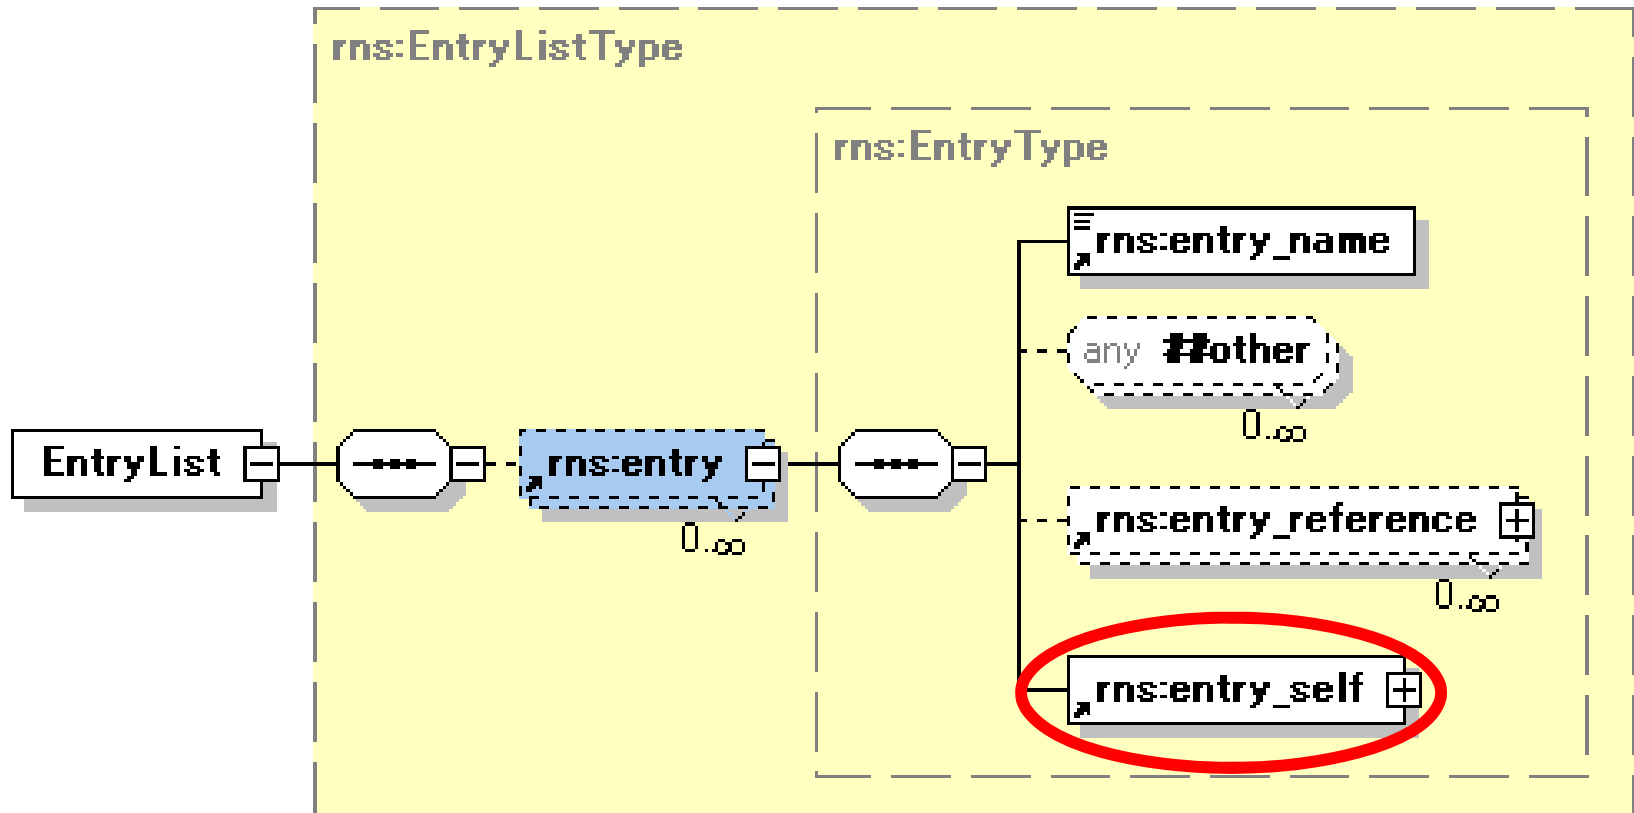
RNS Implementation(2)

- Backend DBMS: MySQL
 - 2 InnoDB tables (entries / references)
- Application Server: Sun Java System Application Server 9.1
 - Works on Tomcat too

Further refinement needed for spec

- In examples of the spec, EPRs consist of the server address and entry pathnames. This convention cannot be assumed.
 - EPR to “/some/entry” can be “http://example.net/RNS?id=12345” instead of “http://example.net/RNS/some/entry”
- In order to traverse entries
 - “add” should return an EPR to a newly created entry
 - Response of “list” should include EPRs of listed entries

Response message of “list”



(mini) Benchmark results

Measured the time that a client takes to perform operations on junctions with random names.

Client execution time per entry (ms)

#junctions	add	list	query	update	remove
100	50.0	4.8	12.6	94.5	5.3
200	50.5	4.4	12.4	110.6	4.6

Environment: Pentium M 1.1GHz, 1.5GB Mem, Windows XP

Each entry has one entry reference, an XML element (<test>string</test>)

All junctions are stored in a single directory.

Performance Considerations



- The JAX-WS framework takes 6~8ms
- The performance is limited by HDD accesses on this benchmark (CPU is not 100% used during “add”/”update”)
 - It’s a laptop anyway...
- “Update” is getting slower as more entries reside in a directory
 - Merging two tables into one may improve the performance.

Several observations by Andrew

A screenshot of a Mozilla Firefox browser window. The title bar reads "[ogsa-wg] RNS experiences document - Mozilla Firefox". The address bar shows "http://www.ogf.org/pipermail/". The page content is an email message from Andrew Grimshaw, dated Tue Aug 21 11:44:02 CDT 2007. The message includes a list of links for previous and next messages, and sorting options. The main body of the email starts with "All," followed by a paragraph: "We have done an implementation of RNS and have several observations we write our experiences document." Below this is another paragraph: "First, the version of the document I am referring to is the September 2006 version." The message ends with "p.4 Add function".

[ogsa-wg] RNS experiences document

Andrew Grimshaw [grimshaw at virginia.edu](mailto:grimshaw@virginia.edu)
Tue Aug 21 11:44:02 CDT 2007

- Previous message: [\[ogsa-wg\] Slides presentation and minutes from OGF20 Service Level terms BoF](#)
- Next message: [\[ogsa-wg\] RNS experiences document](#)
- Messages sorted by: [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#)

All,

We have done an implementation of RNS and have several observations we write our experiences document.

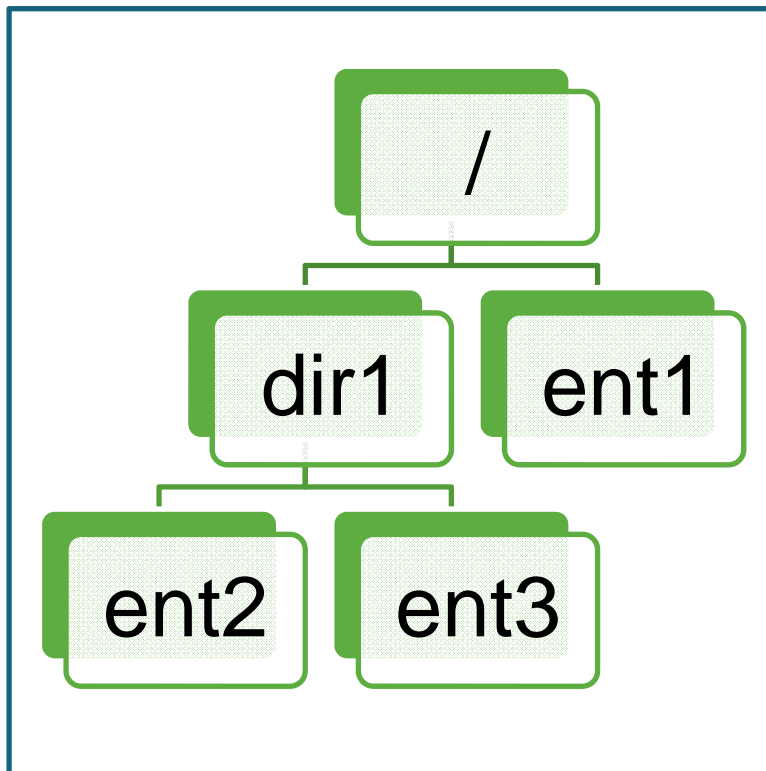
First, the version of the document I am referring to is the September 2006 version.

p.4 Add function

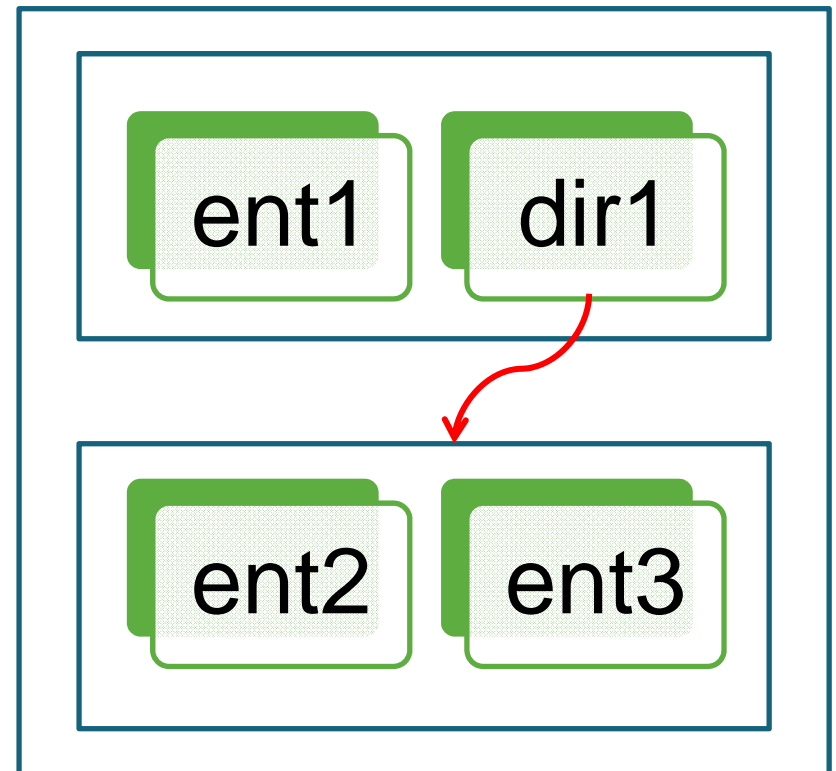
Difference from spec (1)

- Our implementation maintains a directory hierarchy (directories and junctions) within a single resource.
- Implementation of UVa (Genesis II) maintains junctions in a directory in a single resource.
 - Adding a virtual directory means to create a resource and add a junction to the resource.

Our namespace



UVa implementation



Difference from spec (2)

- Junctions cannot be referred to by EPRs
 - “add” cannot return an EPR of a newly created junction entry
 - To “update” a junction, additional entry_name required
 - “query” cannot be implemented
- No virtual directory entry
 - A virtual directory is a WS-resource without RNS metadata entry
 - Metadata cannot be attached to a virtual directory itself

Discussion about query

- Is “query” necessary?
 - “list” returns all information
 - you should know the parent directory when you perform query
- “list” cannot be used for root directory
 - RNS does not provide .(dot) entries for current directories while list can only be performed to directories and takes filename as a parameter
 - root directory may also have properties

Summary

- Overview of RNS
- Schema and WSDL definitions
- Implementation detail and performance results
- Needed refinement of the spec
- Some discussion with UVa team

Next Step

- Further performance and scalability tests
- We need to test interoperability with other implementations
- Make the implementation available
- Implement support for RNS on NAREGI and other middle-ware platforms