

# Product Usage Schemas

Jorge Luis Williams, Principal Architect



Created by: Jorge L. Williams  
Modified Date: 6/7/2014  
Classification: Public

# Background


# What is Rackspace?

---


- A Hosting Company

# What is Rackspace?

- “The Open Cloud Company”




**200,000+**  
**CUSTOMERS**  
90,000+ SERVERS  
26,000+ VM  
≈ 70 PB STORED



**GLOBAL**  
**FOOTPRINT**  
CUSTOMERS IN  
120+ COUNTRIES



**9 WORLDWIDE**  
**DATA CENTERS**



# Cloud Computing...

---

- A major shift in IT...
  - Enables you to provision (or deprovision) computing infrastructure and software platforms dynamically
  - Computing resources can be hosted...
    - On premise (private cloud)
    - Off premise (public cloud)
    - Both (hybrid cloud)



# Cloud Bursting...

---

- The ability to scale work to a public cloud when a private cloud's capacity is exceeded...



# Cloud Bursting...

---

- We can achieve bursting because cloud platforms make extensive use concepts expressed by Service Oriented Architecture (SOA)
  - Loosely coupled cloud services written using the REST architectural style
  - **Resources:** Compute Nodes, Load Balancers, Databases, Storage Volumes etc...
    - Accessed via URIs
    - Via uniform interface provided by the HTTP protocol
    - This simplifies dynamic provisioning, deprovisioning, configuration etc...

# Utility Computing

- Customer pays only for resources consumed
- Providers implement metering in order to track usage
  - Services sold at different pricing tiers known as **flavors**

<i>Flavor</i>	<i>Disk Space</i>	<i>VCPUs</i>	<i>Memory</i>	<i>Price</i>
<i>512 MB Standard</i>	<i>20 GB</i>	<i>1</i>	<i>512 MB</i>	<i>\$0.022/Hr</i>
<i>2 GB Performance</i>	<i>20 GB (SSD)</i>	<i>2</i>	<i>2048 MB</i>	<i>\$0.08/Hr</i>
<i>90 GB Performance</i>	<i>40 GB (SSD)</i>	<i>24</i>	<i>92160 MB</i>	<i>\$5.44/Hr</i>



# Utility Computing

- A Customer is allowed to change a resource's flavor at will, this is known as **resizing**
- A resize operation may upgrade or downgrade a resource

<i>Flavor</i>	<i>Disk Space</i>	<i>VCPUs</i>	<i>Memory</i>	<i>Price</i>
<i>512 MB Standard</i>	<i>20 GB</i>	<i>1</i>	<i>512 MB</i>	<i>\$0.022/Hr</i>
<i>2 GB Performance</i>	<i>20 GB (SSD)</i>	<i>2</i>	<i>2048 MB</i>	<i>\$0.08/Hr</i>
<i>90 GB Performance</i>	<i>40 GB (SSD)</i>	<i>24</i>	<i>92160 MB</i>	<i>\$5.44/Hr</i>



# Datacenters and Regions

---

- Customers may provision resources at different geographical regions in order to...
  - Control Network Latency
  - Achieve high availability and fail over
  - Abide by governmental restrictions
- A region may encompass more than one data center
  - Customer controls region...provider controls data center
- At Rackspace we use Airport codes to identify regions
  - DFW, LON, HKG, etc...

# Creating an Invoice

---

- Usage data, must be collected for
  - Individual Products:
    - Compute
    - Load Balancers
    - Object Storage
    - Databases, ...
  - Across all regions
    - DFW, ORD,LON, HKG, ...

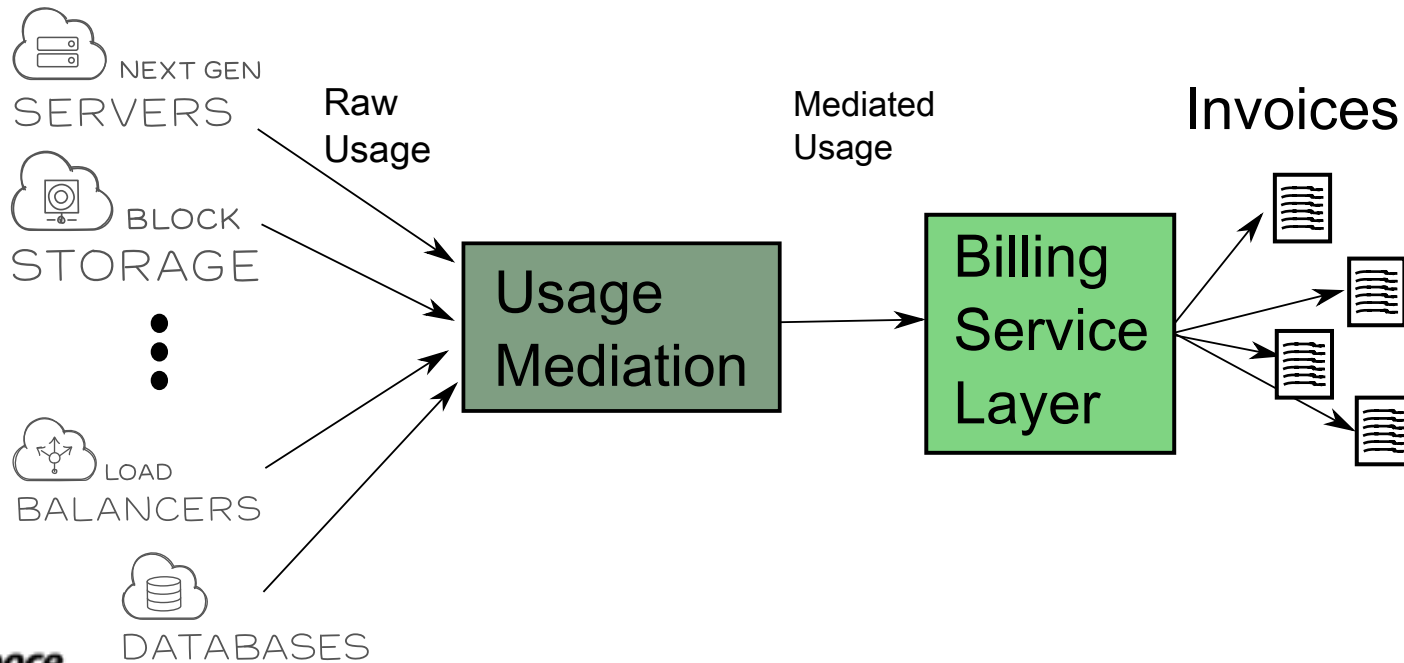
# Mediation

---

- Usage data is aggregated along a number of dimensions:
  - The owner of the resource (the **tenant**)
  - The resource itself the (individual compute node, load balancer, database)
  - The billable usage created by the resource (CPU cycles, bandwidth, etc.)
- The data is enriched with additional info such as the billable account number
- Daily summaries are produced
- The whole process is known as **Mediation**

# Mediation

- Products generate **raw usage**
- Mediation takes **raw usage** and creates **mediated usage**
- Mediated usage is then consumed by a number of billing services that rate the usage and produce invoices...



# The Problem

# Initial Implementation

---

- Mediation system consumed raw usage directly from products
- Products were in complete control of the data format
- There were different methods of obtaining the data
  - FTP
  - HTTP
  - Direct Database Access
  - ...

# Problems...

---

- Common cross-product attributes differed from one product to the next
  - Different names, timezones (GMT, CST), etc
- Attributes change as product evolves
  - No process to notify mediation team of changes
- Difficult to catch errors (even after extensive testing)
- Errors found in later stages of billing pipeline
  - Best case: Reprocess data
  - Worst case: Compensate customer or observe loss



# First Steps

# Standardization

---

- All products transmit messages to centralize AtomPub server using the Atom Syndication Protocol
- Standard event format for common cross-product attributes
  - Enforce GMT timezone
  - Validate against XSD
  - Fail fast if something validation fails
- Products specific attributes defined in their own XSD and also validated..

# Why Atom?

---

- Rackspace standardized on REST, AtomPub is an authoritative example of a RESTful protocol.
- Backed by a large number of implementations, Rackspace implemented its own AtomPub server: **Atom Hopper**
- Atom is extensible, easy to make it suite our needs
- Atom has a standard archiving extension (RFC 5005) — good model for providing access to long term archives for auditing purposes

# Raw Usage in Atom Feeds and Entries

---

- Raw usage is packaged into discrete events in Atom Entries
- These entries are collected in feeds, one feed per product
- The feeds are mapped to a specific URI at a shared endpoint
- In order to avoid network delays one instance of Atom Hopper is deployed per region...
- The mediation service extracts messages from all regions

# Atom Feeds...

DFW:



DATABASES



<http://feeds.dfw.rackspace.com/dbaas/events>



LOAD  
BALANCERS



<http://feeds.dfw.rackspace.com/lbaas/events>

LON:



DATABASES



<http://feeds.lon.rackspace.com/dbaas/events>




LOAD  
BALANCERS



<http://feeds.lon.rackspace.com/lbaas/events>

# USAGE Events

- Capture utility of a resource over a time duration
- Emitted at constant intervals throughout the day,
- after a resize operation, or when a resource is deleted



# USAGE\_SNAPSHOT events

---

- Captures cases that do not correspond directly to the utility model:
  - One Time Charges
  - Subscriptions
- Emitted on an ad hoc basis.

# The Usage Event Format


```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry xmlns="http://docs.rackspace.com/core/event"
  xmlns:atom="http://www.w3.org/2005/Atom"
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas">
  <atom:title type="text">LBAAS</atom:title>
  <atom:content type="application/xml">
    <event type="USAGE" version="1"
      tenantId="3737"
      resourceName="MyLoadBalancer"
      endTime="2012-06-15T10:19:52Z"
      startTime="2012-06-14T10:19:52Z"
      region="DFW" dataCenter="DFW1"
      id="b79cc3de-b399-3883-b555-61829bb7f966"
      resourceId="b79cc3de-b399-3883-b555-61829bbccd38">
      <lbaas:product serviceCode="CloudLoadBalancers"
        resourceType="LOADBALANCER" version="1"
        sslMode="MIXED" vipType="PUBLIC" numVips="44"
        numPolls="10"
        bandwidthOutSsl="345345346" bandwidthInSsl="364646770"
        bandwidthOut="3460346" bandwidthIn="43456346"
        avgConcurrentConnectionsSsl="4566.0"
        avgConcurrentConnections="30000.0"
        status="ACTIVE"/>
    </event>
  </atom:content>
</atom:entry>
```



# The Usage Event Format

```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry xmlns="http://docs.rackspace.com/core/event"
  xmlns:atom="http://www.w3.org/2005/Atom"
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas">
  <atom:title type="text">LBAAS</atom:title>
  <atom:content type="application/xml">
    <event
      tenantId
      resourceName
      endTime
      startTime
      region
      id
      resourceId

      resourceType
      sslMode
      numPolls
      bandwidthOutSsl
      bandwidthOut
      avgConcurrentConnectionsSsl
      avgConcurrentConnections
      status
    </event>
  </atom:content>
</atom:entry>
```



Atom Entry Envelope


# The Usage Event Format

```
<?xml version="1.0" encoding="UTF-8"?>  
<atom:entry
```

```
<atom:title  
<atom:content
```

```
<event type="USAGE" version="1"  
  tenantId="3737"  
  resourceName="MyLoadBalancer"  
  endTime="2012-06-15T10:19:52Z"  
  startTime="2012-06-14T10:19:52Z"  
  region="DFW" dataCenter="DFW1"  
  id="b79cc3de-b399-3883-b555-61829bb7f966"  
  resourceId="b79cc3de-b399-3883-b555-61829bbccd38">
```

Event, captures  
common cross product  
attributes



```
  resourceType  
  sslMode  
  numPolls  
  bandwidthOutSsl  
  bandwidthOut  
  avgConcurrentConnectionsSsl  
  avgConcurrentConnections  
  status
```

```
</event>  
</atom:content>  
</atom:entry>
```

# The Usage Event Format

```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry
```

```
  <atom:title
  <atom:content
    <event
      tenantId
      resourceName
      endTime
      startTime
      region
      id
      resourceId
```

```
    <baas:product serviceCode="CloudLoadBalancers"
      resourceType="LOADBALANCER" version="1"
      sslMode="MIXED" vipType="PUBLIC" numVips="44"
      numPolls="10"
      bandwidthOutSsl="345345346" bandwidthInSsl="364646770"
      bandwidthOut="3460346" bandwidthIn="43456346"
      avgConcurrentConnectionsSsl="4566.0"
      avgConcurrentConnections="30000.0"
      status="ACTIVE"/>
```

Product Specific  
Attributes



```
  </event>
</atom:content>
</atom:entry>
```

# The Usage Event Format

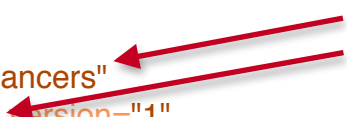
```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas"
  <atom:title
  <atom:content
    <event
      tenantId
      resourceName
      endTime
      startTime
      region
      id
      resourceId
      <lbaas:product serviceCode="CloudLoadBalancers"
        resourceType="LOADBALANCER" version="1"
        sslMode="MIXED" vipType="PUBLIC" numVips="44"
        numPolls="10"
        bandwidthOutSsl="345345346" bandwidthInSsl="364646770"
        bandwidthOut="3460346" bandwidthIn="43456346"
        avgConcurrentConnectionsSsl="4566.0"
        avgConcurrentConnections="30000.0"
        status="ACTIVE"/>
      </event>
    </atom:content>
  </atom:entry>
```

Note product attribute  
are in their own  
namespace.

# The Usage Event Format

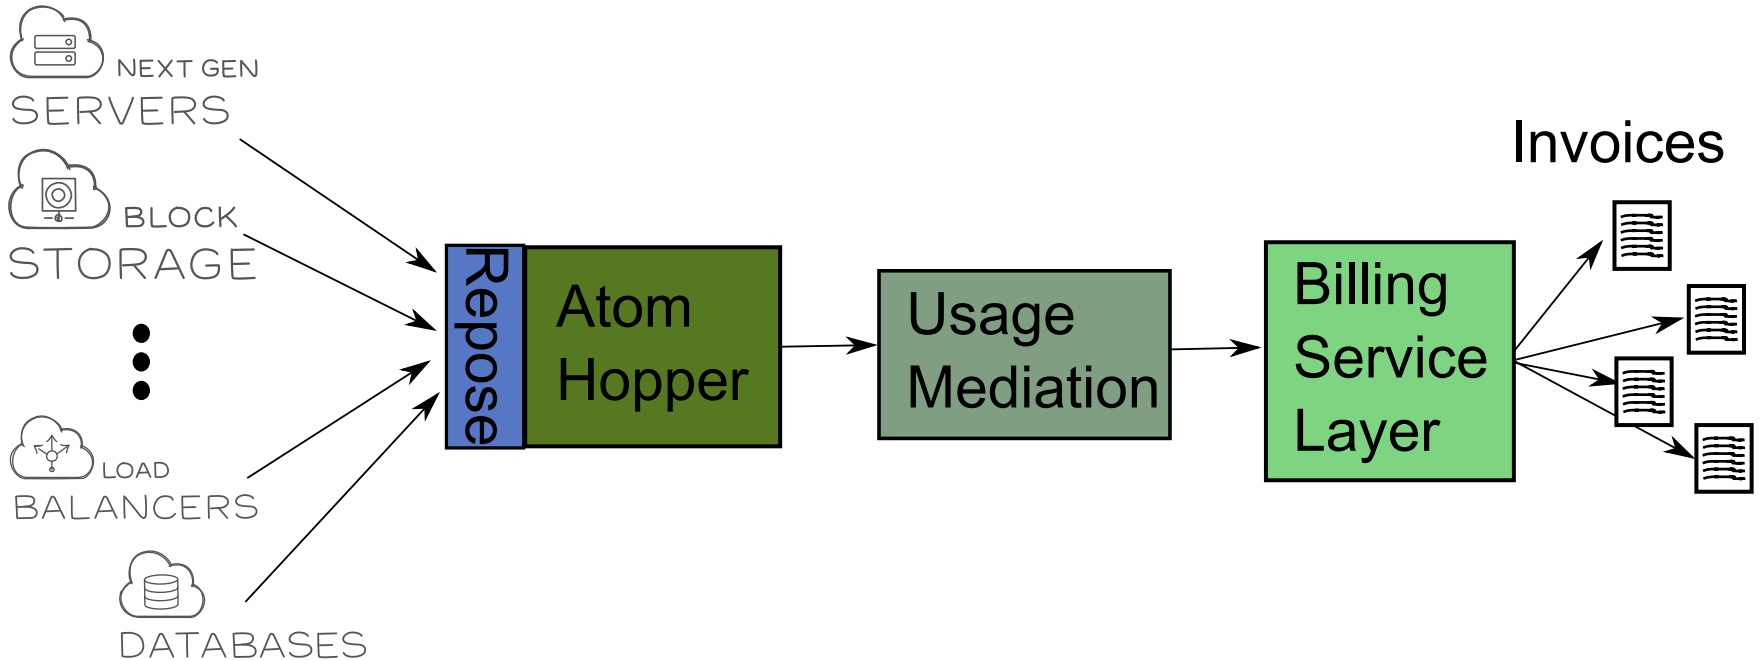
```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas"
  <atom:title
  <atom:content
    <event
      tenantId
      resourceName
      endTime
      startTime
      region
      id
      resourceId
      <lbaas:product serviceCode="CloudLoadBalancers"
        resourceType="LOADBALANCER" version="1"
        sslMode="MIXED" vipType="PUBLIC" numVips="44"
        numPolls="10"
        bandwidthOutSsl="345345346" bandwidthInSsl="364646770"
        bandwidthOut="3460346" bandwidthIn="43456346"
        avgConcurrentConnectionsSsl="4566.0"
        avgConcurrentConnections="30000.0"
        status="ACTIVE"/>
    </event>
  </atom:content>
</atom:entry>
```

These are common cross-product attributes with product specific validation rules



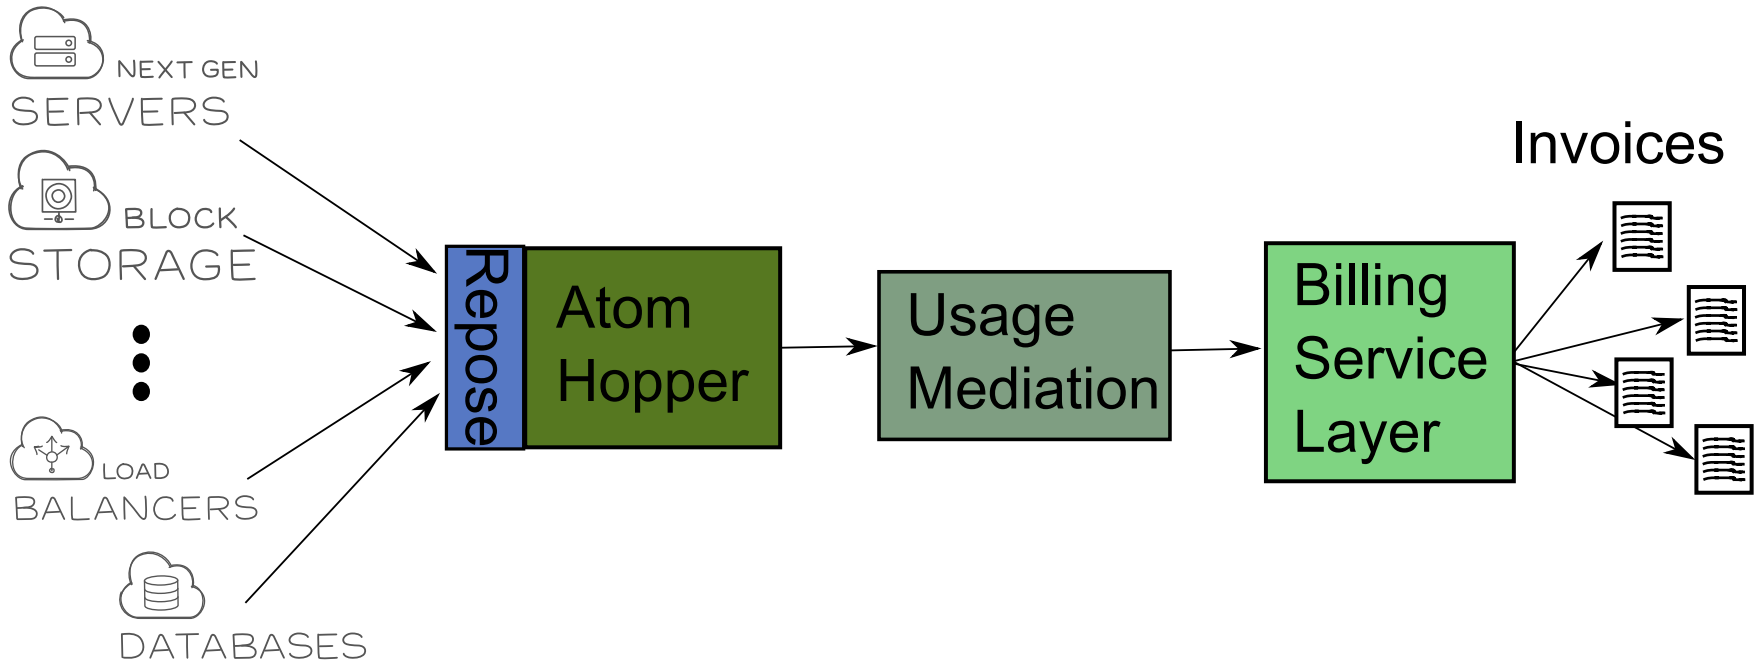
# Validating Events

- Validation achieved via Repose
  - Programmable HTTP proxy
  - Capable of validating REST APIs described by a WADL



# Validating Events

- All Rackspace Specific Business Logic is encapsulated by **Repose**
- Atom Hopper (AtomPub server) is unaware of validation rules.



# Atom Hopper WADL

```
<?xml version="1.0"?>
<application xmlns="http://wadl.dev.java.net/2009/02">
  <grammars>
    <include href="core_xsd/entry.xsd"/>
  </grammars>
  <resources base="http://localhost">
    <resource path="autoscale/events" type="wadl/feed.wadl#AtomFeed wadl/feed.wadl#Unvalidated"/>
    <resource path="backup/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudBackup"/>
    <resource path="bigdata/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#BigData"/>
    <resource path="cbs/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudBlockStorage"/>
    <resource path="files/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudFiles"/>
    <resource path="glance/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#Glance"/>
    <resource path="lbaas/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudLoadBalancers"/>
    <resource path="meta/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#AtomHopper"/>
    <resource path="monitoring/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudMonitoring"/>
    <resource path="netdevice/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#NetDevice"/>
    <resource path="nova/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudServersOpenStack
      wadl/product.wadl#RHEL"/>
    <resource path="queues/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudQueues"/>
    <resource path="servers/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudServers
      wadl/product.wadl#RHEL"/>
    <resource path="sites/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudSites"/>
    <resource path="ssl/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#Ssl"/>
    <resource path="support/events" type="wadl/feed.wadl#AtomFeed wadl/product.wadl#Support"/>
  </resources>
</application>
```



# Atom Hopper WADL

```
<?xml version="1.0"?>
<application
  <grammars>
    <include
  </grammars>
  <resources
    <resource path="autoscale/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="backup/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="bigdata/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="cbs/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="files/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="glance/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="lbaas/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="meta/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="monitoring/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="netdevice/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="nova/events" type="wadl/feed.wadl#AtomFeed
      wadl/product.wadl#RHEL"
    <resource path="queues/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="servers/events" type="wadl/feed.wadl#AtomFeed
      wadl/product.wadl#RHEL"
    <resource path="sites/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="ssl/events" type="wadl/feed.wadl#AtomFeed"
    <resource path="support/events" type="wadl/feed.wadl#AtomFeed"
  </resources>
</application>
```

Common cross feed properties  
(GET operations)



# Atom Hopper WADL

```
<?xml version="1.0"?>
<application
  <grammars>
    <include
  </grammars>
  <resources
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource path="lbaas/events"
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    </resources>
  </application>
```

Product Specific properties  
prevent cross posting of events  
assign product specific assertions



```
type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudLoadBalancers"/>
```

```
wadl/product.wadl#RHEL"
```

```
wadl/product.wadl#RHEL"
```

# Atom Hopper WADL

```
<?xml version="1.0"?>
```

```
<application
```

```
  <grammars>
```

```
    <include
```

```
  </grammars>
```

```
  <resources
```

```
    <resource
```

```
    <resource
```

```
    <resource
```

```
    <resource
```

```
    <resource
```

```
    <resource
```

```
    <resource path="/baas/events"
```

```
    <resource
```

```
    <resource
```

```
    <resource
```

```
    <resource
```

```
  <resource
```

```
  <resource
```

```
  <resource
```

```
  <resource
```

```
  <resource
```

```
</resources>
```

```
</application>
```

Product Specific properties  
named after the service code  
which uniquely ID the product



```
type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudLoadBalancers"/>
```

```
wadl/product.wadl#RHEL"
```

```
wadl/product.wadl#RHEL"
```

# Atom Hopper WADL

```
<?xml version="1.0"?>
<application
  <grammars>
    <include
  </grammars>
  <resources
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource path="nova/events"
      type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudServersOpenStack
        wadl/product.wadl#RHEL"/>
    <resource
    <resource path="servers/events"
      type="wadl/feed.wadl#AtomFeed wadl/product.wadl#CloudServers
        wadl/product.wadl#RHEL"/>
    <resource
    <resource
    <resource
  </resources>
</application>
```

Properties  
can be shared  
across feeds

# Atom Hopper WADL

```
<?xml version="1.0"?>
<application
  <grammars>
    <include
  </grammars>
  <resources
    <resource path="autoscale/events" type="wadl/feed.wadl#AtomFeed wadl/feed.wadl#Unvalidated"/>
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource

    <resource
    <resource

    <resource
    <resource
    <resource
  </resources>
</application>
```

An unvalidated feed.  
Useful for onboarding



wadl/product.wadl#RHEL"

wadl/product.wadl#RHEL"

# Atom Hopper WADL

```
<?xml version="1.0"?>
<application
  <grammars>
    <include href="core_xsd/entry.xsd"/>
  </grammars>
  <resources
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
    <resource
  </resources>
</application>
```



XSD grammar used when validating usage messages...defines Atom entries that contain usage data

wadl/product.wadl#RHEL"

wadl/product.wadl#RHEL"

# The Atom XSD

- We provide our own schema for atom that strictly requires the event element in the content...

```
<complexType name="UsageContent">
  <complexContent>
    <extension base="atom:BaseContent">
      <sequence>
        <choice>
          <element ref="event:event"/>
        </choice>
      </sequence>
    </extension>
  </complexContent>
</complexType>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry xmlns="http://docs.rackspace.com/core/event"
  xmlns:atom="http://www.w3.org/2005/Atom"
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas">
  <atom:title type="text">LBAAS</atom:title>
  <atom:content type="application/xml">
    <event type="USAGE" version="1"
      tenantId="3737"
      resourceName="MyLoadBalancer"
      endTime="2012-06-15T10:19:52Z"
      startTime="2012-06-14T10:19:52Z"
      region="DFW" dataCenter="DFW1"
      id="b79cc3de-b399-3883-b555-61829bb7f966"
      resourceId="b79cc3de-b399-3883-b555-61829bbccd38">
      <lbaas:product serviceCode="CloudLoadBalancers"
        resourceType="LOADBALANCER" version="1"
        sslMode="MIXED" vipType="PUBLIC" numVips="44"
        numPolls="10"
        bandwidthOutSsl="345345346" bandwidthInSsl="364646770"
        bandwidthOut="3460346" bandwidthIn="43456346"
        avgConcurrentConnectionsSsl="4566.0"
        avgConcurrentConnections="30000.0"
        status="ACTIVE"/>
    </event>
  </atom:content>
</atom:entry>
```

# The Atom XSD

- In the definition of the event itself we allow an element from a foreign namespace that must be strictly valid against an XSD

```
<any namespace="##other" processContents="strict"  
minOccurs="0" maxOccurs="unbounded"/>
```



```
<?xml version="1.0" encoding="UTF-8"?>  
<atom:entry xmlns="http://docs.rackspace.com/core/event"  
  xmlns:atom="http://www.w3.org/2005/Atom"  
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas">  
  <atom:title type="text">LBAAS</atom:title>  
  <atom:content type="application/xml">  
    <event type="USAGE" version="1"  
      tenantId="3737"  
      resourceName="MyLoadBalancer"  
      endTime="2012-06-15T10:19:52Z"  
      startTime="2012-06-14T10:19:52Z"  
      region="DFW" dataCenter="DFW1"  
      id="b79cc3de-b399-3883-b555-61829bb7f966"  
      resourceId="b79cc3de-b399-3883-b555-61829bbccd38">  
        <lbaas:product serviceCode="CloudLoadBalancers"  
          resourceType="LOADBALANCER" version="1"  
          sslMode="MIXED" vipType="PUBLIC" numVips="44"  
          numPolls="10"  
          bandWidthOutSsl="345345346" bandWidthInSsl="364646770"  
          bandWidthOut="3460346" bandWidthIn="43456346"  
          avgConcurrentConnectionsSsl="4566.0"  
          avgConcurrentConnections="30000.0"  
          status="ACTIVE"/>  
        </event>  
      </atom:content>  
    </atom:entry>
```



# The Product XSDs

- We initially required products to define their own elements and attributes in product specific XSDs

```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry xmlns="http://docs.rackspace.com/core/event"
  xmlns:atom="http://www.w3.org/2005/Atom"
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas">
  <atom:title type="text">LBAAS</atom:title>
  <atom:content type="application/xml">
    <event type="USAGE" version="1"
      tenantId="3737"
      resourceName="MyLoadBalancer"
      endTime="2012-06-15T10:19:52Z"
      startTime="2012-06-14T10:19:52Z"
      region="DFW" dataCenter="DFW1"
      id="b79cc3de-b399-3883-b555-61829bb7f966"
      resourceId="b79cc3de-b399-3883-b555-61829bbccd38">
      <lbaas:product serviceCode="CloudLoadBalancers"
        resourceType="LOADBALANCER" version="1"
        sslMode="MIXED" vipType="PUBLIC" numVips="44"
        numPolls="10"
        bandwidthOutSsl="345345346" bandwidthInSsl="364646770"
        bandwidthOut="3460346" bandwidthIn="43456346"
        avgConcurrentConnectionsSsl="4566.0"
        avgConcurrentConnections="30000.0"
        status="ACTIVE"/>
    </event>
  </atom:content>
</atom:entry>
```

# The Product XSDs

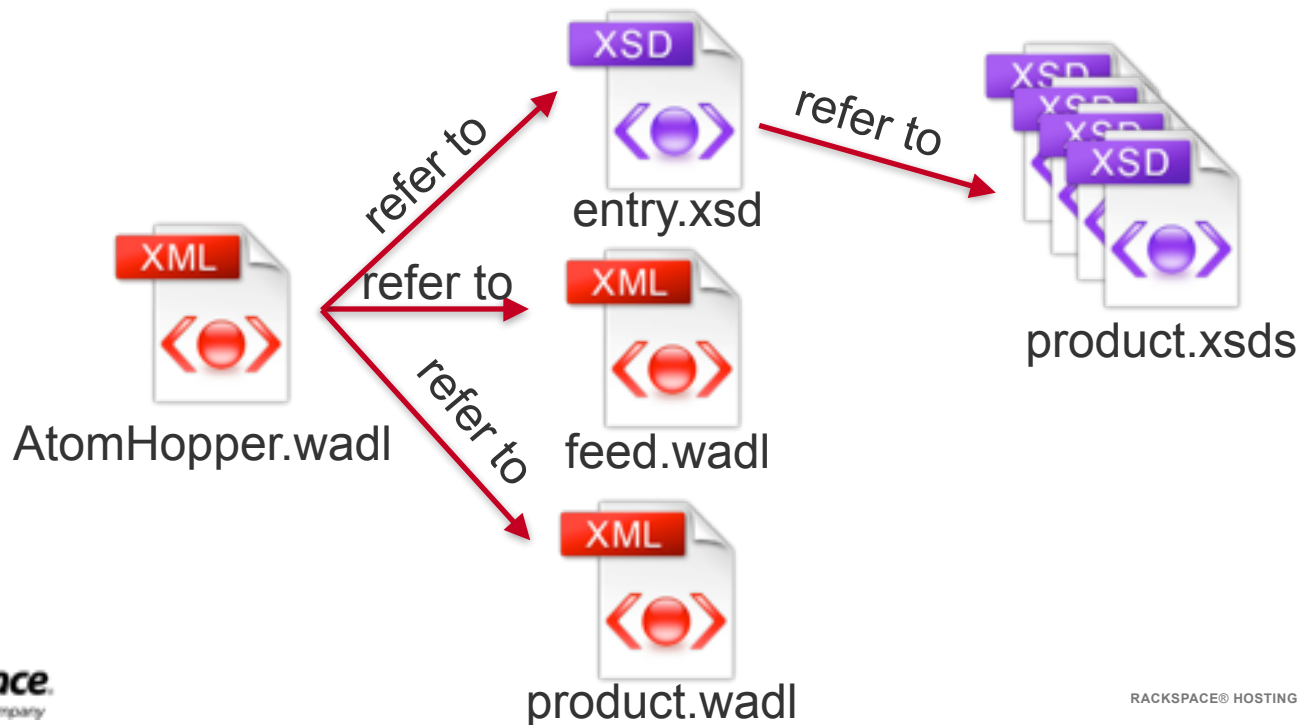
- Additionally we required products to annotate product attributes with instructions to help guide the mediation process

```
<attribute name="avgConcurrentConnections"
  use="required"
  type="p:avgConcurrentConnections1 Type">
  <annotation>
    <documentation>
      <html:p>The amount of conncurrent connections.</html:p>
    </documentation>
    <appinfo>
      <usage:attributes aggregateFunction="WEIGHTED_AVG"
        unitOfMeasure="COUNT"
        groupBy="false"/>
    </appinfo>
  </annotation>
</attribute>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<atom:entry xmlns="http://docs.rackspace.com/core/event"
  xmlns:atom="http://www.w3.org/2005/Atom"
  xmlns:lbaas="http://docs.rackspace.com/usage/lbaas">
  <atom:title type="text">LBAAS</atom:title>
  <atom:content type="application/xml">
    <event type="USAGE" version="1"
      tenantId="3737"
      resourceName="MyLoadBalancer"
      endTime="2012-06-15T10:19:52Z"
      startTime="2012-06-14T10:19:52Z"
      region="DFW" dataCenter="DFW1"
      id="b79cc3de-b399-3883-b555-61829bb7f966"
      resourceId="b79cc3de-b399-3883-b555-61829bbccd38">
      <lbaas:product serviceCode="CloudLoadBalancers"
        resourceType="LOADBALANCER" version="1"
        sslMode="MIXED" vipType="PUBLIC" numVips="44"
        numPolls="10"
        bandWidthOutSsl="345345346" bandWidthInSsl="364646770"
        bandWidthOut="3460346" bandWidthIn="43456346"
        avgConcurrentConnectionsSsl="4566.0"
        avgConcurrentConnections="30000.0"
        status="ACTIVE"/>
    </event>
  </atom:content>
</atom:entry>
```

# Overview...

- Repose reads AtomHopper.WADL which refers to:
  - Feed.wadl (Common cross feed WADL resource types)
  - Product.wadl (Product specific WADL resource types)
  - entry.xsd (Describes Atom Entries that contain Usage Data)
  - product.xsds (Describe product specific elements)



# Problems with the initial approach...

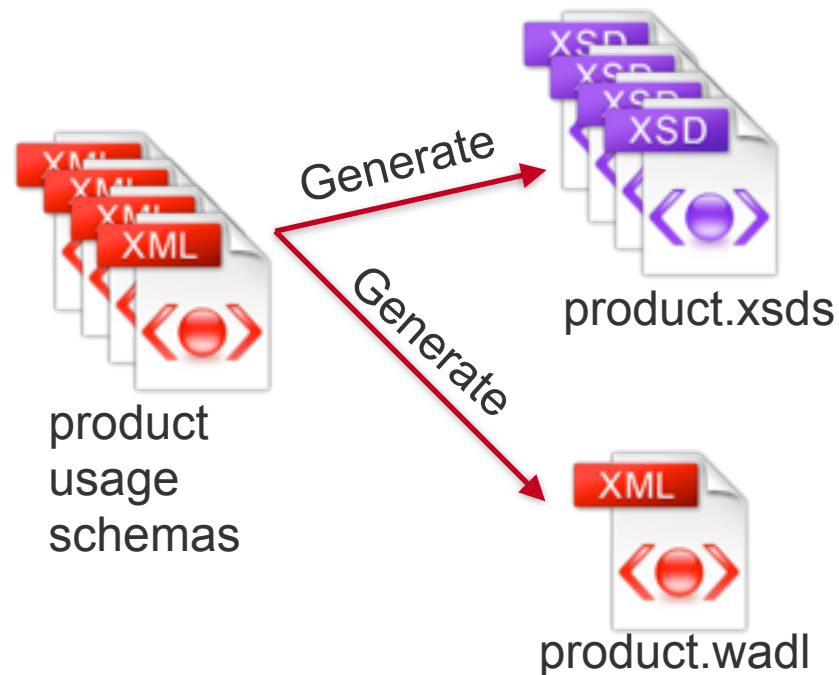
---

- Having product teams define their own XSDs meant that they had to knowledge of XML — something a lot of product teams don't have...
- The XSDs produced by products needed to conform to a number rules (element named product, with service code attribute, with attribute annotations etc). These rules had to be checked.
- The resource types in product.wadl were manually maintained which was tedious and error prone.

# Product Usage Schema

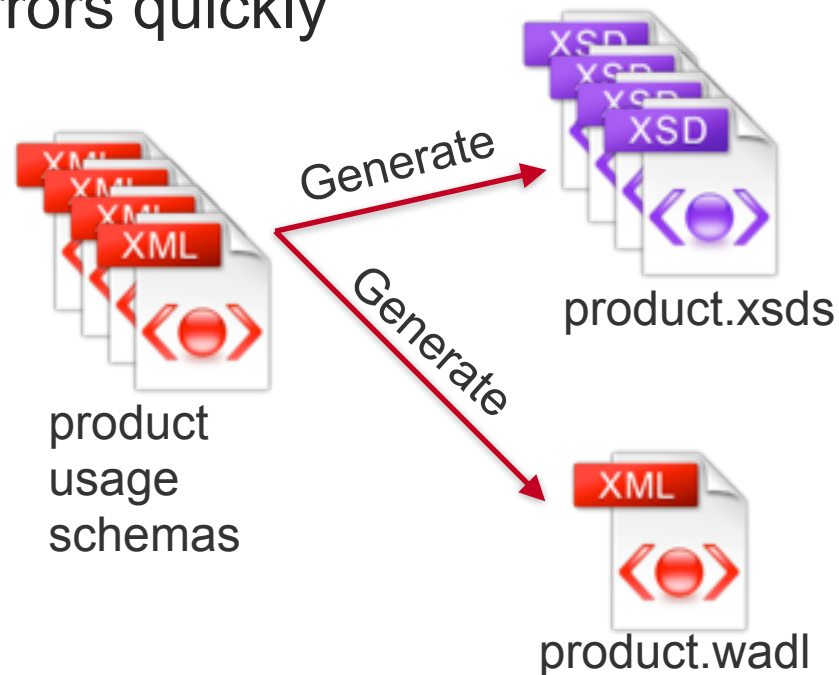
# Generate Artifacts from a Simplified Description: Product Usage Schema

- XProc Pipeline is used to generate product.xsds and the product.wadl from a simplified schema description
- Service Teams are required to use that description instead of an XSD



# Generate Artifacts from a Simplified Description: Product Usage Schema

- Lowers barriers for service teams: removes idiosyncrasies of XSD...is customize to our use case...
- Helps automate the process: eliminates manual checks, helps catch errors quickly



# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas"
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
  <description>
    Lbaas load balancer usage fields.
  </description>
  <attribute name="avgConcurrentConnections" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent connections.
  </attribute>
  <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent ssl connections.
  </attribute>
  <attribute name="bandWidthIn" type="unsignedLong" use="required"
    unitOfMeasure="B" aggregateFunction="SUM" min="0"
    max="10995116277760">
    The amount of bandwidth in, in bytes.
  </attribute>
  <attribute name="vipType" type="string" use="required"
    allowedValues="PUBLIC SERVICENET">
    The vip type associated with the load balancer.
  </attribute>
```



# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas" ← Namespace
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
  <description>
    Lbaas load balancer usage fields.
  </description>
  <attribute name="avgConcurrentConnections" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent connections.
  </attribute>
  <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent ssl connections.
  </attribute>
  <attribute name="bandWidthIn" type="unsignedLong" use="required"
    unitOfMeasure="B" aggregateFunction="SUM" min="0"
    max="1099511627760">
    The amount of bandwidth in, in bytes.
  </attribute>
  <attribute name="vipType" type="string" use="required"
    allowedValues="PUBLIC SERVICENET">
    The vip type associated with the load balancer.
  </attribute>
```

# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas"
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
  <description>
    Lbaas load balancer usage fields.
  </description>
  <attribute name="avgConcurrentConnections" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent connections.
  </attribute>
  <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent ssl connections.
  </attribute>
  <attribute name="bandWidthIn" type="unsignedLong" use="required"
    unitOfMeasure="B" aggregateFunction="SUM" min="0"
    max="1099511627760">
    The amount of bandwidth in, in bytes.
  </attribute>
  <attribute name="vipType" type="string" use="required"
    allowedValues="PUBLIC SERVICENET">
    The vip type associated with the load balancer.
  </attribute>
```

Common cross product attributes are required

# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas"
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
  <description>
    Lbaas load balancer usage fields.
  </description>
  <attribute name="avgConcurrentConnections" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent connections.
  </attribute>
  <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent ssl connections.
  </attribute>
  <attribute name="bandWidthIn" type="unsignedLong" use="required"
    unitOfMeasure="B" aggregateFunction="SUM" min="0"
    max="10995116277760">
    The amount of bandwidth in, in bytes.
  </attribute>
  <attribute name="vipType" type="string" use="required"
    allowedValues="PUBLIC SERVICENET">
    The vip type associated with the load balancer.
  </attribute>
```

Human readable  
documentation is required  
for each attribute...

# Product Schema


```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas"
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
  <description>
    Lbaas load balancer usage fields.
  </description>
  <attribute name="avgConcurrentConnections" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent connections.
  </attribute>
  <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent ssl connections.
  </attribute>
  <attribute name="bandWidthIn" type="unsignedLong" use="required"
    unitOfMeasure="B" aggregateFunction="SUM" min="0"
    max="1099511627760">
    The amount of bandwidth in, in bytes.
  </attribute>
  <attribute name="vipType" type="string" use="required"
    allowedValues="PUBLIC SERVICENET">
    The vip type associated with the load balancer.
  </attribute>
```

Mediation annotations are simply part of the attribute description

# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas"
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
  <description>
    Lbaas load balancer usage fields.
  </description>
  <attribute name="avgConcurrentConnections" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent connections.
  </attribute>
  <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
    aggregateFunction="WEIGHTED_AVG"
    unitOfMeasure="COUNT" min="0" max="1000000">
    The amount of concurrent ssl connections.
  </attribute>
  <attribute name="bandWidthIn" type="unsignedLong" use="required"
    unitOfMeasure="B" aggregateFunction="SUM" min="0"
    max="1099511627760">
    The amount of bandwidth in, in bytes.
  </attribute>
  <attribute name="vipType" type="string" use="required"
    allowedValues="PUBLIC SERVICENET">
    The vip type associated with the load balancer.
  </attribute>
```

Some attributes correspond directly to their equivalent in XSD



# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
  namespace="http://docs.rackspace.com/usage/lbaas"
  serviceCode="CloudLoadBalancers"
  version="1"
  resourceTypes="LOADBALANCER">
```

```
<description>
```

Lbaas load balancer usage fields.

```
</description>
```

```
<attribute name="avgConcurrentConnections" type="double" use="required"
  aggregateFunction="WEIGHTED_AVG"
  unitOfMeasure="COUNT" min="0" max="1000000">
```

The amount of concurrent connections.

```
</attribute>
```

```
<attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
  aggregateFunction="WEIGHTED_AVG"
  unitOfMeasure="COUNT" min="0" max="1000000">
```

The amount of concurrent ssl connections.

```
</attribute>
```

```
<attribute name="bandwidthIn" type="unsignedLong" use="required"
  unitOfMeasure="B" aggregateFunction="SUM" min="0"
  max="10995116277760">
```

The amount of bandwidth in, in bytes.

```
</attribute>
```

```
<attribute name="vipType" type="string" use="required"
  allowedValues="PUBLIC SERVICENET">
```

The vip type associated with the load balancer.

```
</attribute>
```

Similarly enumerations are defined directly...

# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
```

```
  namespace="http://docs.rackspace.com/usage/lbaas"
```

```
  serviceCode="CloudLoadBalancers"
```

```
  version="1"
```

```
  resourceTypes="LOADBALANCER">
```

```
    <description>
```

```
      Lbaas load balancer usage fields.
```

```
    </description>
```

```
    <attribute name="avgConcurrentConnections" type="double" use="required"
```

```
      aggregateFunction="WEIGHTED_AVG"
```

```
      unitOfMeasure="COUNT" min="0" max="1000000">
```

```
      The amount of concurrent connections.
```

```
    </attribute>
```

```
    <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
```

```
      aggregateFunction="WEIGHTED_AVG"
```

```
      unitOfMeasure="COUNT" min="0" max="1000000">
```

```
      The amount of concurrent ssl connections.
```

```
    </attribute>
```

```
    <attribute name="bandWidthIn" type="unsignedLong" use="required"
```

```
      unitOfMeasure="B" aggregateFunction="SUM" min="0"
```

```
      max="1099511627760">
```

```
      The amount of bandwidth in, in bytes.
```

```
    </attribute>
```

```
    <attribute name="vipType" type="string" use="required"
```

```
      allowedValues="PUBLIC SERVICENET">
```

```
      The vip type associated with the load balancer.
```

```
    </attribute>
```

There are only a finite set of simple types. Most borrowed from XSD... but we have specific types such as UUID

# Product Schema

```
<productSchema xmlns="http://docs.rackspace.com/core/usage/schema"
```

```
  namespace="http://docs.rackspace.com/usage/lbaas"
```

```
  serviceCode="CloudLoadBalancers"
```

```
  version="1"
```

```
  resourceTypes="LOADBALANCER">
```

```
    <description>
```

```
      Lbaas load balancer usage fields.
```

```
    </description>
```

```
    <attribute name="avgConcurrentConnections" type="double" use="required"
```

```
      aggregateFunction="WEIGHTED_AVG"
```

```
      unitOfMeasure="COUNT" min="0" max="1000000">
```

```
      The amount of concurrent connections.
```

```
    </attribute>
```

```
    <attribute name="avgConcurrentConnectionsSsl" type="double" use="required"
```

```
      aggregateFunction="WEIGHTED_AVG"
```

```
      unitOfMeasure="COUNT" min="0" max="1000000">
```

```
      The amount of concurrent ssl connections.
```

```
    </attribute>
```

```
    <attribute name="bandWidthIn" type="unsignedLong" use="required"
```

```
      unitOfMeasure="B" aggregateFunction="SUM" min="0"
```

```
      max="10995116277760">
```

```
      The amount of bandwidth in, in bytes.
```

```
    </attribute>
```

```
    <attribute name="vipType" type="string" use="required"
```

```
      allowedValues="PUBLIC SERVICENET">
```

```
      The vip type associated with the load balancer.
```

```
    </attribute>
```

Dates are always in GMT...

Whitespace separated lists denoted with a star(\*)


int\* string\*





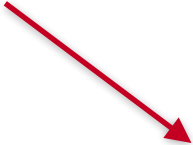
# XSD Translation...

```
<attribute name="avgConcurrentConnections" use="required"
  type="p:avgConcurrentConnections1Type">
  <annotation>
    <documentation>
      <html:p>The amount of conncurrent connections.</html:p>
    </documentation>
    <appinfo>
      <usage:attributes aggregateFunction="WEIGHTED_AVG"
        unitOfMeasure="COUNT"
        groupBy="false"/>
    </appinfo>
  </annotation>
</attribute>
```




```
<attribute name="avgConcurrentConnections" type="double" use="required"
  aggregateFunction="WEIGHTED_AVG"
  unitOfMeasure="COUNT" min="0" max="1000000">
  The amount of conncurrent connections.
</attribute>
```

```
<simpleType name="avgConcurrentConnections1Type">
  <restriction base="xsd:double">
    <minInclusive value="0"/>
    <maxInclusive value="1000000"/>
  </restriction>
</simpleType>
```




# XSD Translation...

```
<attribute name="vipType" use="required" type="p:vipType1Enum">  
  <annotation>  
    <documentation>  
      <html:p>The vip type associated with the load balancer.</html:p>  
    </documentation>  
    <appinfo>  
      <usage:attributes aggregateFunction="NONE" groupBy="false"/>  
    </appinfo>  
  </annotation>  
</attribute>
```



```
<attribute name="vipType" type="string" use="required"  
  allowedValues="PUBLIC SERVICENET">  
  The vip type associated with the load balancer.  
</attribute>
```



```
<simpleType name="vipType1Enum">  
  <restriction base="xsd:string">  
    <enumeration value="PUBLIC"/>  
    <enumeration value="SERVICENET"/>  
  </restriction>  
</simpleType>
```

# Assertions

- It became evident that we couldn't catch all errors unless we add support for assertions...
- Assertions have what we call a scope (product, entry)
- Determines the current node is the XPath expression

```
<xpathAssertion test="if (@sslMode = 'OFF') then  
  ((@bandWidthInSsl = 0) and (@bandWidthOutSsl = 0) and  
  (@avgConcurrentConnectionsSsl = 0)) else true()">  
  If SslMode is OFF then bandWidthInSsl, bandWidthOutSsl, and  
  avgConcurrentConnectionsSsl should all be 0.  
</xpathAssertion>
```

```
<xpathAssertion test="$event/@resourceId castable as xs:integer" scope="entry">  
  The resourceId for a VIP should be an integer.  
</xpathAssertion>
```

# Assertions

- When the scope is entry we define: the variables \$entry, \$event, and \$product to reach into different parts of the message.

```
<xpathAssertion test="if (@sslMode = 'OFF') then  
  ((@bandWidthInSsl = 0) and (@bandWidthOutSsl = 0) and  
  (@avgConcurrentConnectionsSsl = 0)) else true()">  
  If SslMode is OFF then bandWidthInSsl, bandWidthOutSsl, and  
  avgConcurrentConnectionsSsl should all be 0.  
</xpathAssertion>
```

```
<xpathAssertion test="$event/@resourceId castable as xs:integer" scope="entry">  
  The resourceId for a VIP should be an integer.  
</xpathAssertion>
```

# Assertions

- How the assertion get's generated depends on the scope.
- Product scope assertions are generated directly into the XSD...(we use XSD 1.1)

```
<assert vc:minVersion="1.1"
  test="if (@sslMode = 'OFF') then ((@bandWidthInSsl = 0) and (@bandWidthOutSsl = 0) and
(@avgConcurrentConnectionsSsl = 0)) else true()"
  xerces:message="If SslMode is OFF then bandWidthInSsl, bandWidthOutSsl, and avgConcurrentConnectionsSsl should all be 0."
  saxon:message="If SslMode is OFF then bandWidthInSsl, bandWidthOutSsl, and avgConcurrentConnectionsSsl should all be 0.">
  <annotation>
    <documentation>
      <html:p>Assertion: If SslMode is OFF then bandWidthInSsl, bandWidthOutSsl, and avgConcurrentConnectionsSsl should
all be 0.</html:p>
    </documentation>
  </annotation>
</assert>
```

# Assertions

- How the assertion get's generated depends on the scope.
  - Entry scope assertions are generated directly into the product WADL via XSLT preprocess extension...
  - Identity transform with a choose statement...
  - WADL does offer support for XPath assertions in the spec, but we'd lose the ability to set variables like \$event

```
<xsl:choose>
  <xsl:when test="$event/@resourceId castable as xs:integer"/>
  <xsl:otherwise>
    <xsl:message terminate="yes">The resourceId for a load balancer should be an integer.</xsl:message>
  </xsl:otherwise>
</xsl:choose>
```

# Versioning

---

- Our versioning strategy requires that versions remain backwards compatible as much as possible
  - Clients don't have to change
  - We don't have to POST multiple versions of the same message
  - Version numbers are there to set different validation rules

# Versioning

---

- You can create multiple product schemas with the same namespace and the same service code...but different versions
- These multiple schemas generate a single XSD with an alternative element that selects a different complex type

```
<element name="product" vc:minVersion="1.1" type="p:BaseBigDataType">  
  <alternative test="(@version eq '1')" type="p:BigData1Type"/>  
  <alternative test="(@version eq '2')" type="p:BigData2Type"/>  
</element>
```



# Future Work

---

- Better XSD 1.0 support
  - Teams use JAXB etc. We take advantage of take advantage of the conditional inclusion support XSD 1.1, but the generated XSDs still require some manual tweaking...
- JSON Support
  - Generating JSON Schema, allowing JSON messages
- GUIs and Wizards for producing product schema...to lower barriers even more.
- Improve load time. Compiling WADLs, XSLTs, and XSDs is currently taking about 5 mins!

# Conclusion

---

- We transform our existing billing pipeline from one in which usage data was defined in irregular and inconsistent ways to one in which strict standards were followed.
- Product schema allowed us to leverage advanced features in XSD and WADL without having product teams deal with the idiosyncrasies of these languages.
- We can now produce product specific usage data in a highly consistent manner and to communicate requirements to the usage mediation team accurately.

# THANK YOU

jorge.williams@rackspace.com



**RACKSPACE® HOSTING** | 5 MILLINGTON ROAD | HAYES, UNITED KINGDOM UB3 4AZ  
**UK SALES:** +44 (0)20 8712 6507 | **UK SUPPORT:** 0800 988 0300 | **WWW.RACKSPACE.CO.UK**