An XML-Based Protocol for Distributed Event Services

Warren Smith
NASA Ames Research Center

Dan Gunter
Lawrence Berkeley National Laboratory

Darcy Quesnel
Argonne National Laboratory

Motivation

- Government agencies and academia developing computational grids
  - Large-scale distributed systems with high-performance resources
- Requires resource, service, and application monitoring
  - Load on a computer, status of a service, performance of an application
- Different requirements for different types of monitoring
- Several projects addressing performance monitoring in grids
- Define standard protocols and data representations for interoperability in the Grid Forum

Approach

- Define a relatively simple event protocol
  - TCP
  - XML
- Implement the protocol
- Provide to the community
- Gain experience
- Improve the protocol
- In this talk: architecture, events, and event protocol

Grid Monitoring Architecture

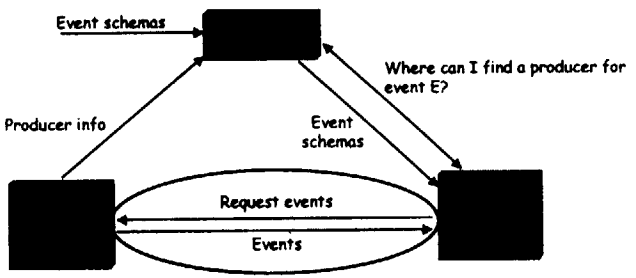

Events

- An event consists of:
  - Name
  - A set of elements
    - Name
    - Value
    - Optional attributes (units and accuracy)
- Has associated input parameters
  - Name
  - A set of elements
- Defined using XML Schema
- Define a base Event and EventParameters type and elements in separate namespaces
- Anyone can define their own namespace and events

CPU Load Event

- Event elements
  - HostName: The name of the host the measurement is made on.
  - Load1: The 1 minute CPU load reported by uptime.
  - Load5: The 5 minute CPU load reported by uptime.
  - Load15: The 15 minute CPU load reported by uptime.
  - TimeStamp: The time the measurement was made in a standard format.
- Event parameters
  - Period: How often to perform the measurement.

CPU Load Event in XML

Event Parameters

```
<UptimeCPULoad
xmlns="http://www.gridforum.org/Performance/EventParameters">
    <Period units="min">2</Period>
</UptimeCPULoad>
```

Event

```
<UptimeCPULoad
xmlns="http://www.gridforum.org/Performance/Event">
    <Load1>1.5</Load1>
    <Load5>1.6</Load5>
    <Load15>1.3</Load15>
    <HostName>foo.nas.nasa.gov</HostName>
    <TimeStamp>2000-11-09T21:51:45Z</TimeStamp>
</UptimeCPULoad>
```


Round Trip Time Event

- Event elements
  - SourceHostName: The name of the host performing the ping command.
  - TargetHostName: The name of the host being pinged.
  - RoundTripTime: The round trip time reported by the ping command in milliseconds.
  - TimeStamp: The time the measurement was made.
- Event parameters
  - TargetHostName: The name of the host being pinged.
  - Period: How often to perform the measurement.

Round Trip Time Event in XML

Event Parameters

```
<Ping
xmlns="http://www.gridforum.org/Performance/EventParameters">
    <TargetHostName>bar.lbl.gov</TargetHostName>
    <Period units="min">1.5</Period>
</Ping>
```

Event

```
<Ping
xmlns="http://www.gridforum.org/Performance/Events">
    <SourceHostName>foo.nas.nasa.gov</SourceHostName>
    <TargetHostName>bar.lbl.gov</TargetHostName>
    <RoundTripTime units="ms">7</RoundTripTime>
    <TimeStamp>2000-11-09T21:53:45Z</TimeStamp>
</Ping>
```


Protocol Interactions

- Subscription: A request that initiates the transmission of a series of events
- Consumer subscription
  - Consumer subscribes for events from a producer
  - Producer sends events to the consumer
  - Consumer unsubscribes for events
- Producer subscription
  - Producer initiates a subscription for a consumer
  - Producer sends events to the consumer
  - Producer or consumer unsubscribes
- Consumer query
  - Consumer asks the producer for an event and producer replies with the event

General Message Format

- Length (in bytes) as a 32-bit integer in network byte order
- XML tags that indicate the message type: <SubscribeRequest> ... <SubscribeRequest>
- Message specific data inside the tags

Example Interaction I

- Consumer->Producer

```
<SubscribeRequest xmlns="http://www.gridforum.org/Performance/Protocol"
                  requestID="2">
  <Ping xmlns="http://www.gridforum.org/Performance/EventParameters">
    <Period units="min">5</Period>
    <TargetHostName>bar.lbl.gov</TargetHostName>
  </Ping>
</SubscribeRequest>
```

- Producer->Consumer

```
<SubscribeReply xmlns="http://www.gridforum.org/Performance/Protocol"
                requestID="2">
  <Return>Success</Return>
  <SubscriptionID>1234</SubscriptionID>
</SubscribeReply>
```

Example Interaction II

- Producer->Consumer

```
<Event xmlns="http://www.gridforum.org/Performance/Protocol"
       subscriptionID="1234">
  <Ping xmlns="http://www.gridforum.org/Performance/Events">
    <SourceHostName>foo.nas.nasa.gov</SourceHostName>
    <TargetHostName>bar.lbl.gov</TargetHostName>
    <RoundTripTime units="ms" accuracy="0.1">7</RoundTripTime>
    <TimeStamp>2000-11-09T21:53:45Z</TimeStamp>
  </Ping>
</Event>
```


Example Interaction III

- Consumer->Producer

```
<UnsubscribeRequest xmlns="http://www.gridforum.org/Performance/Protocol"
                    requestID="15">
  <SubscriptionID>1234</SubscriptionID>
</SubscribeRequest>
```

- Producer->Consumer

```
<UnsubscribeReply xmlns="http://www.gridforum.org/Performance/Protocol"
                  requestID="15">
  <Return>Success</Return>
</SubscribeReply>
```


Event Protocol Performance

- 2 independent implementations of protocol
  - Java by Dan Gunter
  - C++ by Warren Smith
- Encoding/decoding an event message on a 933MHz PIII running Redhat 7.1 with JDK 1.3
  - Java
    - Encode 21,900 events/sec
    - Decode 600 events/sec
  - C++
    - Encode 28,100 events/sec
    - Decode 4,300 events/sec

Why didn't you use...?

- XML-RPC
  - May be superceded by other protocols
- SOAP
  - Were no good implementations available
  - Performance concerns
- Some other text protocol
  - XML has nice features, even if it is verbose
- CORBA event service
  - CORBA is not popular in grids
- Binary
  - We will probably define one at some point

Future Work

- Use our implementations to gain experience
- Extend the protocol
- Experiment with other protocols
  - SOAP
  - Binary
- Standardize the most appropriate protocols

Questions?

- Global Grid Forum
  - http://www.gridforum.org
- Performance Working Group
  - http://www-didc.lbl.gov/GridPerf