

# A Simple XML Producer-Consumer Protocol

Warren Smith, Dan Gunter, Darcy Quesnel  
[wwsmith@nas.nasa.gov](mailto:wwsmith@nas.nasa.gov), [dkgunter@lbl.gov](mailto:dkgunter@lbl.gov), [quesnel@mcs.anl.gov](mailto:quesnel@mcs.anl.gov)

Grid Forum Performance Working Group  
November 30, 2000

## 1. Introduction

This document describes a simple XML-based protocol that can be used for producers of events to communicate with consumers of events. The protocol described here is not meant to be the most efficient protocol, the most logical protocol, or the best protocol in any way. This protocol was defined quickly and its intent is to give us a reasonable protocol that we can implement relatively easily and then use to gain experience in distributed event services. This experience will help us evaluate proposals for event representations, XML-based encoding of information, and communication protocols.

The next section of this document describes how we represent events in this protocol and then defines the two events that we choose to use for our initial experiments. These definitions are made by example so that they are informal and easy to understand. The following section then proceeds to define the producer-consumer protocol we have agreed upon for our initial experiments.

## 2. Events

There is an entire Performance Working Group document on event schemas. Since two of the authors of this document are authors of the event schema document, the events defined here closely relate to the events defined in that document. The difference is that the events defined here are meant to be very simple and easy to write code to translate from programming language structures to XML and from XML back to programming language structures. In this document, we use the following simplifications:

- Implementers are not required to validate the information in events against the information defined in event types.
- Element values are always strings (follows from the previous simplification).
- Values in events do not have any accuracy or unit attributes.
- We will not support sending context information outside of events. For example, a stream of events will not begin with a description of common information in those events (their context).

We assume that an event consists of a type and a set of elements where each element is a <name, value> pair and the names and values are strings. Each element can be either implicitly required or optional. We say implicitly because there is not any explicit type checking being performed at this time.

To further simplify things, we initially only handle 2 different types of events. The next two subsections describe these two types.

### 2.1. CPU Load Event

The CPU load event type is a simple event for containing the load information returned by the Unix uptime command. We therefore use the event type "UptimeCPULoad" for this event to differentiate it from other means of measuring CPU load. This event must contain the following elements:

- **TimeStamp.** The time at which the CPU load event was generated. The format of the time contained in the value of this element is a restricted version of the ASCII format proposed by Dan Gunter in his document "A Standard Timestamp for Grid Computing." This standard proposed by Dan is an extension of one of the variations of the ISO 8601 time format. To restrict Dan's time format, we assume that all times are reported in GMT or zulu time and that precision is not represented. An example time with milliseconds would look like: 2000-11-09T11:12:00.000Z.
- **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.
- **HostName.** The name of the host the load measurement is made on.

Here is an example UptimeCPULoad event in our XML encoding:

```
<UptimeCPULoad>
  <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>
```

Note that the indenting and new lines are for readability purposes only. We assume that the events will not contain any new lines or spaces around element tags when the events are transmitted.

When asking for a CPU load event, the following input parameters can be specified:

- **Period.** The number of seconds between each uptime measurement and event generation. This parameter is only used when a subscription is performed. If this parameter is specified for a query, it is ignored.

An example of a description of which events a consumer desires is:

```
<UptimeCPULoad>
  <Period>600</Period>
</UptimeCPULoad>
```

## 2.2. Round Trip Time Event

The second event we define here is a latency event with data produced by the Unix ping command. We simply call this event a "Ping" event. The event must contain the following elements:

- **TimeStamp.** The time at which the ping was performed. The format is the same as the time stamp above.
- **SourceHostName.** The host name or IP address of the host that is performing the ping command.
- **TargetHostName.** The host name or IP address of the host that the source host is pinging. Note that this name or IP address can indicate 1 of several network interfaces on the target host.
- **RoundTripTime.** The round-trip time reported by the ping command in milliseconds.

Here is an example Ping event in our XML encoding:

```
<Ping>
  <SourceHostName>foo.nas.nasa.gov</SourceHostName>
  <TargetHostName>bar.lbl.gov</TargetHostName>
  <RoundTripTime>7</RoundTripTime>
  <TimeStamp>2000-11-09T21:53:45Z</TimeStamp>
</Ping>
```

When asking for a ping event, the following input parameters can be specified:

- **Period.** The number of seconds between each uptime measurement and event generation. This parameter is only used when a subscription is performed. If this parameter is specified for a query, it is ignored.
- **TargetHostName.** The name or IP address of the host that will be pinged.

An example of a description of which events a consumer desires is:

```
<Ping>
  <Period>600</Period>
  <TargetHostName>bar.lbl.gov</TargetHostName>
</Ping>
```

### 3. XML Protocol

This section describes the XML protocol we use for communication between producers and consumers. We require that our protocol allow:

- Consumers to subscribe for events
- Consumers to unsubscribe from events
- Producers to send events to consumers
- Consumers to query for a single event

We assume that:

- TCP sockets are used for communication.
- No authentication, authorization, or security is supported.
- XML is used to represent control and data messages.
- One socket between a producer and a consumer is used for both control and data messages.
- Multiple subscriptions can be active over a socket.
- Multiple requests can be outstanding on the same connection.

Given the requirements of our protocol, we define the following messages:

1. SubscribeRequest
2. SubscribeReply
3. UnsubscribeRequest
4. UnsubscribeReply
5. Event
6. QueryRequest

## 7. QueryReply

Note that many of the messages are pairs of request/reply messages.

### 3.1. General Message Format

In general, each message consists of:

1. Length: the number of bytes in the message as a 32-bit integer in network byte order (C functions: htonl, ntohl)
2. The XML tags that indicate the message type. For example: <SubscribeRequest> ... </SubscribeRequest> or <QueryReply> ... </QueryReply>
  - 2.1. Request messages always have a requester-unique request ID chosen by the requestor. This request ID is an attribute of the message tag. For example: <SubscribeRequest requestID="1"> ... </SubscribeRequest>.
  - 2.2. Reply messages always have a request ID. The request ID for a reply matches the request ID given by the requestor for the request that is being replied to. For example, a reply to the above subscribe request would be: <SubscribeReply requestID="1"> ... </SubscribeReply>.
  - 2.3. Reply messages always have a return code and may have a detailed return message. The Return element indicates if an operation was successful (Success) or a failure (Failure). These return codes will most likely be expanded later to contain more detailed error codes. The ReturnDetail element contains a text message that contains detailed user-readable information about what the status of a request.
3. The message-specific data inside the XML tags that identify the message. For example: <SubscribeRequest requestID="1"> <UptimeCPULoad> ... </UptimeCPULoad> </SubscribeRequest>

### 3.2. Subscribing for Events

When a consumer wants to receive a stream of events from a producer, it subscribes for events. For now we are assuming that after a subscription, a producer pushes events to the consumer that has subscribed for them. This subscription takes the form of a SubscribeRequest message from the consumer to the producer, followed by a SubscribeReply message from the producer to the consumer.

#### 3.2.1. SubscribeRequest Message

The subscribe request message consists of:

- Event type (required).
- Any input parameters needed to generate events (optional).
- A <requestor, socket>-unique request ID that the reply to this message refers to.

Here are two examples of SubscribeRequest messages:

```
<SubscribeRequest requestID="1">
  <UptimeCPULoad>
    <Period>600</Period>
  </UptimeCPULoad>
</SubscribeRequest>
```

```
<SubscribeRequest requestID="2">
  <Ping>
    <Period>300</Period>
    <TargetHostName>bar.lbl.gov</TargetHostName>
  </Ping>
</SubscribeRequest>
```

Note that the name of the element containing the input parameters is the same as the name of the event that will be generated. See Section 3.4.

In the future, we will probably want an event filter. After an event is generated by a producer, a filter is used to determine if the event should be sent to a consumer. For example, a filter may indicate that an event should be sent only if the load is greater than 5.0. The filter is a boolean expression of event element names and constants that is applied to an event. We think the filter should go inside the PingParameters.

### 3.2.2. SubscribeReply Message

The subscribe reply message consists of:

- Return (required). Success means the request the reply is for was successfully completed, Failure means the request failed. Other non-zero return codes to represent more detailed failures will most likely be added in the future.
- ReturnDetail (optional). Text giving further information about the successful or unsuccessful subscribe.
- Requested (required) of the request that this message is in reply to.
- An optional producer-unique SubscriptionID that identifies the subscription that was successfully made by the consumer (if one was). The subscription ID should be present if the subscription was successful and should not be present if the subscription was not successful. The subscription ID is used so that a consumer can easily route an event to the handler for events from that subscription.

Two examples of SubscribeReply messages are:

```
<SubscribeReply requestID="1">
  <Return>Failure</Return>
  <ReturnDetail>The period specified is too small.</ReturnDetail>
</SubscribeReply>
```

```
<SubscribeReply requestID="2">
  <Return>Success</Return>
  <SubscriptionID>1234</SubscriptionID>
</SubscribeReply>
```

### 3.3. Unsubscribing from Events

When a consumer wishes to stop receiving a stream of events from a producer, it unsubscribes from those events. This unsubscription is accomplished by the consumer sending an UnsubscribeRequest message to the producer and the producer sending an UnsubscribeReply message back to the consumer in response to the request.

### 3.3.1. UnsubscribeRequest Message

The unsubscribe request message consists of:

- The requested (required) of this request so that the reply to this request can be identified.
- The SubscriptionID (required) that identifies the subscription to be unsubscribed from.

An example of an UnsubscribeRequest message is:

```
<UnsubscribeRequest requestID="9">
  <SubscriptionID>1234</SubscriptionID>
</UnsubscribeRequest>
```

### 3.3.2. UnsubscribeReply Message

The unsubscribe reply message consists of:

- Return (required). Success means the request the reply is for was successfully completed, Failure means the request failed. Other non-zero return codes to represent more detailed failures will most likely be added in the future.
- ReturnDetail (optional). Text giving further information about the successful or unsuccessful unsubscribe.
- requestID (required) of the request that this message is in reply to.

Example of an UnsubscribeReply message are:

```
<UnsubscribeReply requestID="9">
  <Return>Success</Return>
</UnsubscribeReply>
```

```
<UnsubscribeReply requestID="9">
  <Return>Failure</Return>
  <ReturnDetail>Unknown subscription ID.</ReturnDetail>
</UnsubscribeReply>
```

## 3.4. Event Message

An event message is sent from the producer to the consumer after a consumer has subscribed for events. An event message consists of:

- A subscription ID to identify which subscription the event is for
- The event data in the format described in Section 2.
- Optional Error and ErrorDetail elements to provide information to the consumer if an error occurs when generating an event. At the current time, we have not defined any error codes. The presence of an Error element indicates that an error has occurred and the user should examine the ErrorDetail element for further information.

Example Event messages are:

```
<Event subscriptionID="1234">
  <UptimeCPULoad>
    <Load1>1.5</Load1>
    <Load5>1.6</Load5>
```

```

    <Load15>1.3</Load15>
    <TimeStamp>2000-11-09T21:51:45Z</TimeStamp>
  </UptimeCPULoad>
</Event>

```

```

<Event subscriptionID="1234">
  <UptimeCPULoad>
    <Error>Authorization</Error>
    <ErrorDetail>You are no longer authorized to receive this
      information.</ErrorDetail>
  </UptimeCPULoad>
</Event>

```

```

<Event subscriptionID="1235">
  <Ping>
    <SourceHostName>foo.nas.nasa.gov</SourceHostName>
    <TargetHostName>bar.lbl.gov</TargetHostName>
    <RoundTripTime>7</RoundTripTime>
    <TimeStamp>2000-11-09T21:53:45Z</TimeStamp>
  </Ping>
</Event>

```

### 3.5. Querying for an Event

We believe there will be many cases when a consumer wants to ask for just 1 event from a producer. Instead of having a consumer subscribe, receive 1 event, and then unsubscribe, we allow a consumer to query a producer for an event. This query consists of a `QueryRequest` message that a consumer sends to the producer and a `QueryReply` message that the producer sends to the consumer in response to the `QueryRequest` message.

#### 3.5.1. QueryRequest Message

The query request message is very similar to the subscribe request message and consists of:

- Event type (required).
- Any input parameters needed to generate events (optional).
- Request ID so that a reply can be matched with a request.

Here are two examples of `QueryRequest` messages:

```

<QueryRequest requestID="15">
  <UptimeCPULoad>
  </UptimeCPULoad>
</QueryRequest>

```

```

<QueryRequest requestID="20">
  <Ping>
    <TargetHostName>bar.lbl.gov</TargetHostName>
  </Ping>
</QueryRequest>

```

### 3.5.2. QueryReply Message

The QueryReply messages are similar to the event messages and consist of:

- Return (required). Success means the request the reply is for was successfully completed, Failure means the request failed.
- ReturnDetail (optional). Text giving further information about the successful or unsuccessful unsubscribe.
- A request ID (required) to identify which QueryRequest this reply is for
- The event data in the format described in Section 2.

Example QueryReply messages are:

```
<QueryReply requestID="15">
  <Return>Success</Return>
  <UptimeCPULoad>
    <Load1>1.5</Load1>
    <Load5>1.6</Load5>
    <Load15>1.3</Load15>
    <TimeStamp>2000-11-09T21:51:45Z</TimeStamp>
  </UptimeCPULoad>
</QueryReply>
```

```
<QueryReply requestID="20">
  <Return>Success</Return>
  <Ping>
    <SourceHostName>foo.nas.nasa.gov</SourceHostName>
    <TargetHostName>bar.lbl.gov</TargetHostName>
    <RoundTripTime>7</RoundTripTime>
    <TimeStamp>2000-11-09T21:53:45Z </TimeStamp>
  </Ping>
</QueryReply>
```

## 4. Summary

This document describes a simple XML-based protocol that can be used for producers of events to communicate with consumers of events. The protocol described here is not meant to be the most efficient protocol, the most logical protocol, or the best protocol in any way. This protocol was defined quickly and it's intent is to give us a reasonable protocol that we can implement relatively easily and then use to gain experience in distributed event services. This experience will help us evaluate proposals for event representations, XML-based encoding of information, and communication protocols.

If you have comments on this paper they can be addressed to the Grid Forum Performance Working Group at [perf-wg@gridforum.org](mailto:perf-wg@gridforum.org) or to the authors individually. The Grid Forum Performance Working Group web page is located at <http://www-didc.lbl.gov/GridPerf> and contains links to email archives, documents written for the working group, and other useful information.