

GWD-R (-00)
Job Submission Description Language (JSDL) Specification

<http://forge.gridforum.org/projects/ogsa-hpcp-wg>

Authors:
Blair Dillaway, Microsoft
Marty Humphrey, UVA
Chris Smith, Platform (Editor)
Glenn Wasson, UVA

6/11/2007

HPC Basic Profile, Version 0.3

Status of this Memo

This memo provides information to the Grid community regarding the specification of the HPC Basic Profile. Distribution is unlimited.

Copyright Notice

Copyright © Global Grid Forum (2003-2005). All Rights Reserved.

Abstract

This document defines the HPC Basic Profile, consisting of a set of non-proprietary specifications, along with clarifications, refinements, interpretations and amplifications of those specifications which promote interoperability. The single use-case addressed in this Profile is the “Base Case” (Section 2) of [HPC-U].

Contents

Abstract	1
1 Introduction	3
2 Notational Conventions	3
3 Job Description	4
3.1 JobDefinition	4
3.2 JobDescription	4
3.2.1 JobIdentification	4
3.2.2 JobName	4
3.2.3 JobProject.....	4
3.2.4 Application.....	4
3.2.5 Resources	4
4 Job Scheduling and Management Services	5
5 Security Considerations	6
5.1 Security Requirements of the HPC Basic Profile	6
5.1.1 Environment Assumptions.....	7
5.1.2 Securing the HPC Profile Messages	7
5.2 HPC Basic Profile Message Security	8
5.3 X.509 Certificate Based Mutual Authentication.....	8
5.3.1 Faults	Error! Bookmark not defined.
5.4 Username-Password Client Authentication	9
5.4.1 Faults	Error! Bookmark not defined.
6 Author Information	10
7 Contributors.....	11
8 Acknowledgements	11
Full Copyright Notice	11
Intellectual Property Statement	11
Normative References.....	12

1 Introduction

The HPC Basic Profile is a document that is used to describe how a particular set of specifications are composed in order to solve a basic use case around the use of HPC systems [refer to use case document]. The single use-case addressed in this Profile is the “Base Case” (Section 2) of [HPC-U].

The Profile consists of references to existing specifications, along with any clarifications of the contents of those specifications, restrictions on the use of those specifications, and references to any normative extensions to those specifications. While it is envisioned that many systems will have capabilities above and beyond those described in this profile, this profile describes a basic set of capabilities that can be used as the basis of interoperability testing between systems claiming compliance.

The document is structured as a set of sections, each of which is used to reference a particular aspect of an HPC Basic Profile compliant system. The first is that of job description, which references the Job Submission Description Language, version 1.0 [JSDL10] and the HPC Profile Application Extension [JSDLHPC]. The second is job scheduling and management, which references the OGSA Basic Execution Services specification [BES10].

It is worth noting that this profile is focused on describing the basic capabilities that must be supported by a compliant system. In many cases, the systems in question will support higher levels of functionality than described here, and many systems will support various extensions to the functionality described in the referenced specifications. It is not the goal of this profile to prohibit the use of such extensions, but to define a set of capabilities that can provide a basis for interoperability. As such, this profile may implicitly allow the use of various constructs, but not make any statement about the semantics of such use, and thus these constructs should not be used as the basis of any interoperability testing of HPC Basic Profile compliant systems.

2 Notational Conventions

The key words “MUST,” “MUST NOT,” “REQUIRED,” “SHALL,” “SHALL NOT,” “SHOULD,” “SHOULD NOT,” “RECOMMENDED,” “MAY,” and “OPTIONAL” are to be interpreted as described in RFC-2119 [RFC 2119].

The document refers to an “HPC Basic Profile compliant system” as a “Compliant system”.

This specification uses namespace prefixes throughout; they are listed in Table 2-1. Note that the choice of any namespace prefix is arbitrary and not semantically significant.

Table 2-1: Prefixes and namespaces used in this specification.

Prefix	Namespace
xsd	http://www.w3.org/2001/XMLSchema
jsdl	http://schemas.ggf.org/jsdl/2005/11/jsdl
jsdl-hpcp	http://schemas.ggf.org/jsdl/2006/07/jsdl-hpcp
bes-factory	http://schemas.ggf.org/bes/2006/08/bes-factory
hpcp-bp	http://schemas.ggf.org/hpcp/2007/01/bp

3 Job Description

This section describes restrictions and clarifications to the Job Submission Description Language, version 1.0 [JSDL10] and the HPC Profile Application Extension [JSDLHPC] specifications.

The following elements within a JSDL document **MUST** be supported by a Compliant system. For the purposes of this document, supporting an element has a stronger meaning than with [JSDL10]. In order to support an element, a Compliant system must not only parse the element, but must accept the element as part of the JSDL job definition, and apply the semantics as indicated by the referenced specification with any clarifications or restrictions as described in this section.

JSDL documents **MAY** include additional elements from [JSDL10] beyond those listed in this section. A Compliant system **MAY** support any such additional elements should it encounter them in a submitted JSDL document. However, a Compliant system **MAY** also instead return a BES `UnsupportedFeatureFault` in response to encountering any such additional elements from [JSDL10].

3.1 JobDefinition

As in [JSDL10].

3.2 JobDescription

A Compliant system **MUST** support the `jsdl:JobIdentification`, `jsdl:Application`, and `jsdl:Resources` sub-elements.

3.2.1 JobIdentification

A Compliant system **MUST** support the `jsdl:JobName` and `jsdl:JobProject` sub-elements.

3.3 JobName

As in [JSDL10].

3.4 JobProject

As in [JSDL10].

3.5 Application

A Compliant system **MUST** support the `jsdl-hpcp:HPCProfileApplication` sub-element, as defined in [JSDLHPC].

3.6 Resources

A Compliant system **MUST** support the following sub-elements within the `jsdl:Resources` element: `jsdl:CandidateHosts`, `jsdl:ExclusiveExecution`, `jsdl:OperatingSystem`, `jsdl:CPUArchitecture`, and `jsdl:TotalCPUCount`.

3.7 CandidateHosts

The `jsdl:CandidateHosts` complex type will be supported as described in [JSDL10].

3.8 ExclusiveExecution

As in [JSDL10], with the clarification that the resources being allocated to the job are "hosts". That is, if a job runs exclusively on a host, then no other jobs may run concurrently on the same host.

3.9 OperatingSystem

The `jsdl:OperatingSystem` complex type will be supported as described in [JSDL10]. If the consuming system does not provide the requested operating system, or if the JSDL special token "other" is used as the content of the `jsdl:OperatingSystemName` sub-element, and if the consuming system does not understand the provided extension content, then the consuming system MAY return the BES `InvalidRequestMessageFault` to the requester.

3.10 CPUArchitecture

The `CPUArchitecture` complex type will be supported as described in [JSDL10]. If the consuming system does not provide the requested CPU architecture, or if the JSDL special token "other" is used as the content of the `jsdl:CPUArchitectureName` sub-element, and if the consuming system does not understand the provided extension content, then the consuming system MAY return the BES `InvalidRequestMessageFault` to the requester.

3.11 TotalCPUCount

The description is as in [JSDL10]. A Compliant system MUST support non-negative integer values of the `jsdl:Exact` element from the `jsdl:RangeValue_Type`. It MUST support non-negative integer values of the `jsdl:UpperBoundRange` and `jsdl:LowerBoundRange`, and MUST support the `exclusiveBound` attribute on these elements. It MAY support non-integer values, it MAY support the `epsilon` attribute of `jsdl:Exact`, and it MAY support the `jsdl:Range` element, but MAY instead return a `UnsupportedFeatureFault` in response to encountering such elements

4 Job Scheduling and Management Services

This section describes restrictions and clarifications to the OGSA Basic Execution Services specification [BES10].

A Compliant system MUST support the BES base case specification. It MAY additionally support BES extension profiles.

4.1 BES Vector Operations

The BES `GetActivitiesStatus`, `TerminateActivities`, and `GetActivityDocuments` operations include a vector input parameter that specifies the set of activities that the operation should be applied to. A Compliant system MUST support a vector length of 1. A Compliant system SHOULD support input vector lengths greater than 1 but MAY return a BES `UnsupportedFeatureFault` in response to input vector lengths greater than 1.

4.2 FactoryResourceAttributesDocument contents

The `bes-factory:FactoryResourceAttributesDocument`, as returned by the `GetFactoryAttributesDocument` operation, includes a list of activities currently managed by the BES as well as a list of contained resources that are allocated for the use of these activities. If the numbers of activities or contained resources gets large, then the corresponding size of this document can also be quite large. Given that repeated requests for this document could incur a large cost for both clients and servers in transferring and parsing this document, the BES MAY choose not to return the `ActivityReference` or `ContainedResource` sub-elements of the `bes-factory:FactoryResourceAttributesDocument` on a request to request basis.

In order to distinguish between the absence of any activities being managed by the BES, and the BES implementation choosing not to return the `ActivityReference` sub-elements, the BES MUST provide the number of managed activities in the `TotalNumberOfActivities` sub-element of the `bes-factory:FactoryResourceAttributesDocument`.

In order to distinguish between the absence of any contained resources available to the BES, and the BES implementation choosing not to return the `ContainedResource` sub-elements, the BES

MUST provide the number of available contained resources in the TotalNumberOfContainedResources sub-element of the bes-factory:FactoryResourceAttributesDocument.

4.3 *BasicFilter extension*

Since there are cases when a client explicitly requires the complete list of both activities or contained resources, the BES MAY support the hpcp-bp:BasicFilter extension element within the content of the bes-factory:GetFactoryAttributesDocumentType. A BES that chooses to support this extension MUST return a BESExtension sub-element of bes-factory:FactoryResourceAttributesDocument containing the URI "<http://schemas.ogf.org/hpcp/2007/01/bp/BasicFilter>", and MUST provide the following semantics when encountering a hpcp-bp:BasicFilter element in the bes-factory:GetFactoryAttributesDocumentType.

The hpcp-bp:BasicFilter has the structure (the normative schema is provided in Appendix A):

```
<hpcp-bp:BasicFilter>
  <ActivityReferences> true|false </ActivityReferences>
  <ContainedResources> true|false </ContainedResources>
</hpcp-bp:BasicFilter>
```

There are four possible cases:

1. If both the ActivityReferences and the ContainedResources sub-elements are false in the BasicFilter, then the BES MUST NOT return either ActivityReference or ContainedResource sub-elements in the bes-factory:FactoryResourceAttributesDocument.
2. If the ActivityReferences sub-element is true and the ContainedResources sub-element is false in the BasicFilter, then the BES MUST return an ActivityReference sub-element for each activity managed by the BES, and the BES MUST NOT return any ContainedResource sub-elements in the bes-factory:FactoryResourceAttributesDocument.
3. If the ActivityReferences sub-element is false and the ContainedResources sub-element is true in the BasicFilter, then the BES MUST NOT return any ActivityReference sub-elements, and the BES MUST return a ContainedResource sub-element for each contained resource available to the BES in the bes-factory:FactoryResourceAttributesDocument.
4. If both the ActivityReferences and ContainedResources sub-elements are true in the BasicFilter, then the BES MUST return an ActivityReference sub-element for each activity managed by the BES, and the BES MUST return a ContainedResource sub-element for each contained resource available to the BES.

5 Security Considerations

In this section, interoperable security mechanisms which HPC Basic Profile compliant implementations must support is defined. These mechanisms are limited to those necessary to address the requirements of the "Base Case" (Section 2) of [HPC-U]. Compliant implementations may support additional security mechanisms required for extended functionality as discussed in Section 3 of [HPC-U].

5.1 *Security Requirements of the HPC Basic Profile*

The environment in which an HPC Basic Profile service/client will operate is described below along with the requirements for securing the HPC Basic Profile messages.

5.1.1 Environment Assumptions

In addressing the Base Case some common assumptions are made about the environment and relationships between the users and BES web service schedulers. The security mechanisms defined in this specification build on this environment.

1. There is an identity management infrastructure deployed for provisioning users and services with identity credentials.
 - Web services are provisioned with X.509 [RFC 3280] service certificates following industry standard practice.
 - It is required that users be provisioned with username-password credentials or X.509 certificates. If an organization uses X.509 client certificates, username-password credentials may also be utilized but are not required.
2. Trust relationships are pre-configured and uniform
 - Users trust the CA(s) issuing X.509 service certificates and services trust the authority provisioning username-password credentials or the CA(s) issuing X.509 user certificates.
 - All BES Web services are fully trusted with respect to managing and executing activities within the environment and safeguarding any confidential user and activity information.
 - Users may not fully trust each other. They may require their activities be free from tampering by other users, or in some cases that the details of their activities (job type, data source, ..) not be exposed to other users.
3. X.509 certificate revocation may be supported using industry standard mechanism such as CRLs [RFC3280] and OCSP [RFC 2560] responders. It is up to the relying party whether to take advantage of revocation information.
4. It is assumed BES services are well-known to users and other services and may be located using commonly deployed mechanisms such as DNS (Domain Name Service) or UDDI (Universal Description Discovery and Integration) look-ups.
5. Authorization is based on authenticated user/service identities and attributes carried in the provisioned identity credentials. The authorization mechanism employed is outside the scope of this specification.

5.1.2 Securing the HPC Profile Messages

There is a need to secure messages exchanged between users and BES scheduler services to support the Base Case. The security mechanisms must support required message sender authentication (BES requests and responses), integrity protection, and confidentiality. These are summarized below:

BES Request Message Authentication – BES services require authentication of clients (may be a user or other service) invoking their services to ensure only authorized actions are performed. This includes, limiting who may create an activity, cancel an activity, and query an activity's status.

BES Response Message Authentication – Entities requesting BES services will require authentication of the responding service. This is needed to ensure that returned status information or faults can be relied upon.

Integrity Protection – High assurance message integrity is necessary to prevent attackers from modifying activity definitions for purposes such as creating incorrect billing or denial of service.

Confidentiality - In some environments, activity details and status information will be considered confidential. As such, it will be mandatory to encrypt the BES messages to prevent disclosure to unauthorized entities. Confidentiality of this information may not be critical in other environments, though message encryption is still acceptable.

5.2 HPC Basic Profile Message Security

This specification takes the position that security interoperability for the Base Use case is best achieved through a few widely deployed, standards-based, technologies and vetted implementation guidance. It is not a goal of this specification to innovate in the security area or drive adoption of new technologies.

To that end, use of TLS/SSL transport layer security as the basis for interoperable secure messages is adopted. This provides greater functionality that absolutely required for some environments, but minimizes the number of mechanisms which must be supported. It is not believed the tools, and supporting infrastructure, for interoperable message-level security (based on the WS-* family of specifications) have reached the level of adoption and deployment needed to rely on their use as the primary security mechanism for this profile.

The HPC Basic Profile builds on the "WS-I Basic Security Profile" (WS-I BSP) [WS-I Basic Security Profile Version 1.0, Working Group Draft, 2006-08-17] as the foundation for interoperable message security. In particular, the transport layer security mechanisms identified in Section 4 of that specification are used. The more restrictive cipher suite guidelines specified in the "OGSA Basic Security Profile 1.0 - Secure Channel" (OGSA BSP-SC) are also adopted as described in Section 5.3.

(Note: The "OGSA Basic Security Profile 1.0 - Core" specification is not used as that addresses the binding of key information to an endpoint reference [in WS-Addressing], which is not relevant when using transport layer security.)

The HPC Basic Profile message security mechanisms and requirements are defined in Section 5.3 and 5.4. Compliant implementations are required to fully implement one of these mechanisms, though they may support both. The terminology of the WS-I BSP is used to define compliant implementations. Specifically, a conforming INSTANCE is "software that implements a wsd!port or a uddi:bindingTemplate".

5.3 TLS/SSL using X.509 Certificate Based Mutual Authentication

This specification supports use of the Transport Layer Security (TLS 1.0 and TLS 1.1)[RFC2246 and RFC 4346] or Secure Sockets Layer (SSL 3.0) protocol for BES message security with mutual authentication of the sender and receiver based on X.509 v3 certificates. This is done in accordance with the recommendations of WS-I BSP and the more restrictive cipher suite guidance of the OGSA BSP-SC specification. Faults shall be handled in accordance with the TLS/SSL specifications.

Specific requirements of this specification are:

R0501: An INSTANCE MUST support TLS 1.0, SHOULD support SSL 3.0, and SHOULD support TLS 1.1.

R0502: An INSTANCE MUST support the FIPS-140 compliant Ciphersuites TLS_RSA_FIPS-WITH_3DES_EDE_CBC_SHA and TLS_RSA_FIPS_WITH_3DES_EDE_CBC_SHA

R0503: An INSTANCE SHOULD support TLS_RSA_WITH_AES_128_CBC_SHA and TLS_RSA_WITH_AES_128_CBC_SHA.

R0504: An INSTANCE MUST support X.509 v3 certificates using RSA cryptographic keys and RSA/SHA-1 (<http://www.w3.org/2000/09/xmlsig#rsa-sha1>) digital signatures.

R0505 An INSTANCE SHOULD support X.509 v3 certificates using

RSA cryptographic keys and RSA/SHA-256 (<http://www.w3.org/2001/04/xmldsig-more#rsa-sha256>) digital signatures.

R0506: An INSTANCE must use TLS/SSL encryption key agreement based on the RSA algorithm. Diffie-Helman key agreement shall not be used.

R0507 An INSTANCE MUST support server authentication using X.509 v3 certificates.

R0508 An INSTANCE MUST support client authentication using X.509 v3 certificates.

5.3 TLS/SSL with Username-Password Client Authentication

This specification supports use of the TLS or SSL protocol for BES message security with x.509 server authentication and username-password based client authentication. When using this mechanism, a secure TLS/SSL session with the BES service must be first established. This is done in conformance with the recommendations contained in the WS-I BSP and requirements R0501 through R0507 above. That is, service authentication is done using an X.509 service certificate and a channel encryption key negotiated using RSA key transport.

Once an encrypted and integrity protected transport layer channel has been established, the client may transmit an HPC Basic Profile supported request messages, including their username-password authentication information as specified in the Username Token Profile 1.1 specification [Web Services Security UsernameToken Profile, Working Draft 2, OASIS, 23 Feb 2003].

Specific requirements of this specification are:

R0509: An INSTANCE MUST support client authentication using username/password credentials with cleartext (<http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-username-token-profile-1.0#PasswordText>) type encoding.

R0510: An INSTANCE MAY support client authentication using username/password credentials with digest (<http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-username-token-profile-1.0#PasswordDigest>) type encoding.

Since all password information is communicated within a secure transport layer by compliant implementations, this specification does not specify use of message-level encryption. Also, use of nonces or creation times to prevent replay attacks is not required by this specification and these may be omitted from a password digest calculation.

Faults occurring during TLS/SSL negotiation shall be handled in accordance with the TLS/SSL specifications. If faults arise based on processing of the clients username-password credential by the service, the service may silently drop the request message or respond with a SOAP fault message. When responding with a fault message, if the service is unable to validate the supplied credentials a SOAP fault with faultcode 'Client' should be returned otherwise a fault with faultcode 'Server' shall be returned. Compliant BES service implementations may wish to implement mechanisms to limit the number of invalid authentication attempts for a given username to prevent password guessing attacks.

An example CreateActivity message, including a username and digest password is shown below.

```
<s11:Envelope
  xmlns:s11="http://schemas.xmlsoap.org/soap/envelope"
```

```

xmlns:wsa="http://www.w3.org/2005/08/addressing"
xmlns:bes-"factory="http://schemas.ggf.org/bes/2006/08/bes-factory"
xmlns:wss="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-
wssecurity-secext-1.0.xsd"
xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-
wssecurity-utility-1.0.xsd" >
  <s11:Header>
    <wsse:Security>
      <wsse:UsernameToken xmlns:wsse='http://docs.oasis-
open.org/wss/2004/01/oasis-200401-wss-wssecurity-secext-1.0.xsd' >
        <wsse:Username>Bert</wsse:Username>
        <wsse:Password Type='http://docs.oasis-
open.org/wss/2004/01/oasis-200401-wss-username-token-profile-
1.0#PasswordText'>Ernie</wsse:Password>
        </wsse:UsernameToken>
      </wsse:Security>
    <wsa:Action>
      http://schemas.ggf.org/bes/2006/08/bes-
factory/GetActivitiesStatus
    </wsa:Action>
    <wsa:To s11:mustUnderstand=1>
      http://www.bes.org/BESFactory
    </wsa:To>
  </s11:Header>
  <s11:Body wsu:Id='TheBody'>
    <bes-factory:CreateActivity>
      <bes-factory:activityDescriptionDocument>
        <bes-factory:ActivityDocument>
          {Any valid JSDL document}
        </bes-factory:ActivityDocument>
      </bes-factory:activityDescriptionDocument>
    </bes-factory:CreateActivity>
  </s11:Body>
</s11:Envelope>

```

6 Author Information

Blair Dillaway
Microsoft Corp.

Marty Humphrey
University of Virginia

Chris Smith
Platform Computing, Inc.

Marvin Theimer
Microsoft Corp.

Glenn Wasson
University of Virginia

7 Contributors

We gratefully acknowledge the contributions made to this specification by [insert names].

8 Acknowledgements

We are grateful to numerous colleagues for discussions on the topics covered in this document, in particular (in alphabetical order, with apologies to anybody we've missed) [insert names].

We would like to thank the people who took the time to read and comment on earlier drafts. Their comments were valuable in helping us improve the readability and accuracy of this document.

9 Full Copyright Notice

Copyright © Global Grid Forum (2003-2005). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE GLOBAL GRID FORUM DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

10 Intellectual Property Statement

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director (see contact information at GGF website).

11 Normative References

[RFC 2119] Bradner, S. *Key words for use in RFCs to Indicate Requirement Levels*. Internet Engineering Task Force, RFC 2119, March 1997. Available at <http://www.ietf.org/rfc/rfc2119.txt>

[JSDL10] Available at <http://www.ggf.org/documents/GFD.56.pdf>

Appendix 1 HPC Basic Profile XML Schema

```
<xsd:schema
  targetNamespace="http://schemas.ogf.org/hpcp/2007/01/bp"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:hpcp-bp="http://schemas.ogf.org/hpcp/2007/01/bp"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <!-- Filter Types -->
  <xsd:complexType name="BasicFilterType">
    <xsd:sequence>
      <xsd:element name="ActivityReferences" type="xsd:boolean"/>
      <xsd:element name="ContainedResources" type="xsd:boolean"/>
      <xsd:any namespace="##other" processContents="lax"
        minOccurs="0" maxOccurs="unbounded"/>
    </xsd:sequence>
    <xsd:anyAttribute namespace="##other" processContents="lax"/>
  </xsd:complexType>

  <xsd:element name="BasicFilter" type="hpcp-bp:BasicFilterType"/>
</xsd:schema>
```