



NextGRID Data Transfer Profile

Editors:	Michel Drescher	FLE
	Steven Davey	UEDIN
	Ali Anjomshoaa	FLE

Date	Author	Comments	Version	Status
Dec 14, 2006	Steven Davey	Initial draft outline	v0.1	Draft
30 Mar 2007	Michel Drescher	Introduction section revised	v0.2	Draft
30 Mar 2007	Michel Drescher	Editorial changes	v0.3	Draft
04 Apr 2007	Michel Drescher	Added/started 3 more sections	v0.4	Draft
10 Apr 2007	Michel Drescher	Cleaned up document structure	v0.5	Draft
17 May 2007	Michel Drescher	Revised references, "TBD"s fixed	v0.6	Draft
13 Sep 2007	Michel Drescher	Added ByteIO as baseline transfer protocol	v0.7	Draft
31 Jan 2008	Michel Drescher	Update to the OGSA-DMI draft 50	v0.8	Draft
4 Feb 2008	Michel Drescher	Removed superfluous DMI WSRF rendering references, update to NextGRID BP 1.0	V0.9	Draft
19 Feb 2008	Ali Anjomshoaa	Final edits	V0.95	Draft
11 Mar 2008	D. Snelling	Final QA	V1.0	Final



- 1 INTRODUCTION..... 3**
 - 1.1 Profile Overview 3
 - 1.2 Relationship to Other Profiles 4
 - 1.3 Notational Conventions 4
 - 1.4 Profile Identification and Versioning 4
- 2 PROFILE CONFORMANCE 4**
 - 2.1 Conformance Targets 5
 - 2.2 Claiming Conformance 5
- 3 NEXTGRID BASIC PROFILE 5**
- 4 ADDRESSING 6**
- 5 RESOURCE PROPERTIES 7**
 - 5.1 Source Resource Properties 8
 - 5.1.1 Source Resource Properties Values 8
 - 5.2 Sink Resource Properties 9
 - 5.2.1 Sink Resource Properties Values 10
- 6 REFERENCES..... 12**



1 Introduction

This document defines the NextGRID Data Transfer Profile 1.0 (hereafter, "this Profile"), consisting of a set of de facto, institutional or evolving institutional Web services specifications, along with clarifications, refinements, interpretations and amplifications of those specifications that promote interoperability among implementations of those specifications.

Section 1 introduces this Profile, and explains its relationships to other profiles.

Section 2, "Profile Conformance," explains what it means to be conformant to this Profile.

Each subsequent section addresses a component of this Profile, and consists of two parts: an overview detailing the component specifications and their extensibility points, followed by subsections that address individual parts of the component specifications. Note that there is no relationship between the section numbers in this document and those in the referenced specifications.

This Profile forms part of the NextGRID Generalized Specifications, which aim at capturing NextGRID architectural concepts in a set of profiles that may be composed together. These profiles are specified in such a way that they could be implemented in terms of other well known specifications. While overall consistency is achieved at the conceptual level, and captured through the motivating Use Cases accompanying each Generalized Specification, system implementations, which may be based on other specifications, may not be consistent with this Profile. Thus, each profile defines a realisation of the underlying concept that can be implemented. However, implementers of the NextGRID architecture may need to support multiple underlying specifications.

1.1 Profile Overview

This Profile is intended for use when implementing Data Transfer services that are in line with the concepts of NextGRID [1]. It mandates the adherence to a certain set of specifications and clarifies their use. A service implementation that uses those specifications in a manner conformant with the Profile may be said to be an "implementation of the NextGRID Data Transfer Profile 1.0" or, informally, to be a "NextGRID Data Transfer Service".

The primary issues addressed in the Profile are as follows:

- *NextGRID Basic Profile.* This Profile mandates the use of aspects of the NextGRID Basic Profile 1.0 [2].
- *Addressing.* This Profile mandates the use of WS-Addressing endpoint references and places some constraints on their structure, thus enabling interoperability.
- *Resource Properties.* This Profile mandates that data transfer resources support certain resource properties, thus facilitating introspection, and also mandates support for certain operations that are optional in the WS-ResourceProperties specification. Both sets of requirements facilitate interoperability by allowing a client to that resource to make more assumptions about a service's interfaces.

1.2 Relationship to Other Profiles

This Profile extends the NextGRID Basic Profile Version 1.0 [2].

1.3 Notational Conventions

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [3].

Normative statements of requirements in the Profile are presented in the manner detailed in the WS-I Basic Profile 1.1 Conformance Requirements section [4].

Both requirement statements and extensibility statements can be considered namespace-qualified. This specification uses a number of namespace prefixes throughout; their associated URIs are listed below. Note that the choice of any namespace prefix is arbitrary and not semantically significant [5].

Table 1 Namespaces used by NextGRID Data Transfer Profile 1.0

Prefix	Namespace
wsdl	http://schemas.xmlsoap.org/wsdl
wsa	http://www.w3.org/2005/08/addressing
wsrf-rp	http://docs.oasis-open.org/wsrp/rp-2
dmi	http://schemas.ogf.org/dmi/2007/05/dmi
rbyteio	http://schemas.ggf.org/byteio/2005/10/random-access

1.4 Profile Identification and Versioning

Profile identification and versioning uses the style described in WS-I Basic Profile 1.1 [4] and abides by the normative descriptions contained therein. The name of this Profile is "NextGRID Data Transfer Profile" and version number is "1.0".

2 Profile Conformance

Conformance to the Profile is defined normatively in WS-I Basic Profile 1.1 [4]. This Profile abides by those definitions.

2.1 Conformance Targets

Since the Profile is an extension of the NextGRID Basic Profile 1.0 [2] it may place further restrictions on conformance targets defined therein.

The following conformance targets are used in this Profile:

- **ENDPOINTREFERENCE** – the serialization of the `wsa:EndpointReference` element and its content (from OGSA WSRF Basic Profile [6]).
- **DESCRIPTION** – descriptions of types, messages, interfaces and their concrete protocol and data format bindings, and the network access points associated with Web services (e.g., WSDL descriptions) (from WS-I Basic Profile 1.1 [4]).
- **INSTANCE** – software that implements a `wsdl:port` (from WS-I Basic Profile 1.1 [4], without “bindingTemplate” from the namespace `urn:uddi-org:api_v2`).
- **CONSUMER** – software that invokes an INSTANCE (from WS-I Basic Profile 1.1 [4]).
- **SENDER** – software that generates a particular message according to the protocol(s) associated with that message (from WS-I Basic Profile 1.1 [4]).
- **RECEIVER** – software that consumes a message according to the protocol(s) associated with that message (e.g., SOAP processors) (from WS-I Basic Profile 1.1 [4]).

2.2 Claiming Conformance

Claims of conformance to the Profile and the attachments mechanisms are the same as normatively described in WS-I Basic Profile 1.1 [4].

The conformance claim URI for this Profile is
“<http://www.nextgrid.org/datatransferbasicprofile/v1>”.

3 NextGRID Basic Profile

The following profile (or sections thereof) is referred to in this section of the Profile:

- NextGRID Basic Profile 1.0, Section 4 [2]
- NextGRID Basic Profile 1.0, Section 5 [2]
- NextGRID Basic Profile 1.0, Section 6 [2]

The NextGRID Basic Profile defines a set of conformance statements that are either mandatory or optional when designing and implementing Web Services that are intended to be conformant to the NextGRID Basic Profile. This profile mandates compliance with these requirements:

R0301 ENDPOINTREFERENCES, INSTANCES, DESCRIPTIONS of INSTANCES, CONSUMERS and RECEIVERS MUST conform to the NextGRID Basic Profile 1.0, Section 4, “OGSA WSRF Basic Profile”.

R0302 ENDPOINTREFERENCES MUST conform to the NextGRID Basic Profile 1.0, Section 5, “Basic Security”.

R0303 INSTANCES, CONSUMERS, SENDERS and RECEIVERS MUST conform to the NextGRID Basic Profile 1.0, Section 6, “Secure Channel”.

4 Addressing

This section of this Profile incorporates the following specification by reference:

- OGSA-DMI Functional Specification 1.0 draft 50 [7].

The OGSA-DMI Functional Specification defines the following information element:

- DataEPR – Section 3.2.1

Furthermore, the OGSA-DMI Functional Specification defines the following two roles, among others:

- Source – Section 3.2.2
- Sink – Section 3.2.3

OGSA-DMI introduces two distinct instances of DataEPRs – the SourceDEPR and SinkDEPR – that are used to describe the existing data at the Source, and the data location at the Sink. OGSA-DMI further introduces a three-way matching process to select a concrete data transfer protocol such as FTP to transfer data from the Source to the Sink. To perform the matching a Data Transfer Factory requires certain metadata describing suitable data transfer protocols for the Source and for the Sink, respectively.

OGSA-DMI allows for three different usage scenarios related to the DataEPRs for the Source and Sink described in sub-sections of Section 3.3:

- Explicit Data Transfer – Section 3.3.1
- Defined Data Transfers – Section 3.3.2
- Logical Data Transfers – Section 3.3.3

The first two usage scenarios effectively reduce the DataEPR structures to data containers. The last scenario allows for more dynamic and composable Grid deployments, enabling an OGSA-DMI implementation to make timely queries for Source and Sink related metadata, generally enhancing accuracy and reliability in a Grid.

This profile mandates using the third usage scenario, “Logical Data Transfers”:

R0401 A SourceDEPR ENDPOINTREFERENCE MUST NOT conform to the requirements given in OGSA-DMI Functional Specification 1.0 draft 50, Section 3.3.1.

R0402 A SourceDEPR ENDPOINTREFERENCE MUST NOT conform to the requirements given in OGSA-DMI Functional Specification 1.0 draft 50, Section 3.3.2.

As well as:

R0403 A SinkDEPR ENDPOINTREFERENCE MUST NOT conform to the requirements given in OGSA-DMI Functional Specification 1.0 draft 50, Section 3.3.1.

R0404 A SinkDEPR ENDPOINTREFERENCE MUST NOT conform to the requirements given in OGSA-DMI Functional Specification 1.0 draft 50, Section 3.3.2.

Consequently:

R0405 A SourceDEPR ENDPOINTREFERENCE MUST conform to the requirements given in OGSA-DMI Functional Specification 1.0 draft 50, Section 3.3.3.

R0406 A SinkDEPR ENDPOINTREFERENCE MUST conform to the requirements given in OGSA-DMI Functional Specification 1.0 draft 50, Section 3.3.3.

Additional requirements on the Source and Sink are given in a later section of this Profile.

5 Resource Properties

The OGSA-DMI Functional Specification describes a set of entities that interact while performing a data transfer. These entities are named Source, Sink, Data Transfer Factory and Data Transfer Instance. Section 4 above mandates that these entities are implemented as WS Resources.

This section mandates the support of certain Resource Properties when implementing and publishing either of those entities.

5.1 Source Resource Properties

This section of this Profile incorporates the following specification by reference:

- OGSA-DMI Functional Specification 1.0 draft 50 [7]

OGSA-DMI Functional Specification defines the following Information Element

- [dmi:Protocol] – Section 4.2.1.2

that describes the availability of the described protocol for reading the data from the Source. This Profile mandates modelling the Source as WS-Resources in Section 4, and mandates using Logical Data Transfers, also in Section 4. Using Logical Data Transfers the Data Transfer Factory needs to learn about the available protocols for the Source before it can find a suitable data transfer protocol. Consequently:

R0501 A Source INSTANCE MUST contain in the Resource Property Document schema, as referred to by the `wsrf-rp:ResourceProperties` attribute on the `portType`, a Resource Property Element `dmi:Protocol` as defined in OGSA-DMI Functional Specification 1.0 draft 50, Section 12.

R0502 A Source INSTANCE, when acting as RECEIVER, MUST respond to a `wsrf-rp:GetResourceProperty` request where the child of the `wsrf-rp:GetResourceProperty` is the QName `dmi:Protocol`, with the complete list of supported source protocols each serialised into one `dmi:Protocol` element.

R0503 A Source INSTANCE, when acting as RECEIVER, MUST respond to a `wsrf-rp:GetMultipleResourceProperties` request where the child of the `wsrf-rp:GetMultipleResourceProperties` contains the QName `dmi:Protocol`, with the complete list of supported source protocols each serialised into one `dmi:Protocol` element, among the values requested for other Resource Properties.

5.1.1 Source Resource Properties Values

This section of this Profile incorporates the following specification by reference:

- OGSA-DMI Functional Specification 1.0 draft 50 [7]
- OGSA-ByteIO Specification 1.0 [8]

The OGSA-DMI Functional Specification defines the following Information Elements and Data Transfer Protocol identifiers:

- [dmi:Protocol] – Section 4.2.1.2
- Data Transfer Protocol Identifiers – Section 4.1.1

that together describe the support of the identified data transfer protocol at the data movement Source.

OGSA-ByteIO Specification 1.0 defines the following ByteIO property (amongst others):

- “Readable” – Section 4.1

The OGSA-ByteIO specification mandates the support of this property. However, the specification allows this property to have the Boolean value “false” which would be incorrect for the Source of a DMI-governed data movement (the defining characteristic of a Source is that data can be read from it).

This Profile mandates that the Source of a DMI-governed data movement supports the OGSA-ByteIO “Random-Access” protocol as a base line data transfer to read data from the Source. It further mandates that the Source must allow read access to the data if it is invoked using the ByteIO Random-Access protocol.

R0504 A Source INSTANCE MUST support the OGSA-ByteIO Random-Access protocol as a data transfer protocol. A Source INSTANCE MUST provide one dmi:Protocol element that identifies OGSA-ByteIO using the URI <http://www.nextgrid.org/datatransferbasicprofile/protocols/ogsa-byteio/v1.0/random-access>”.

R0505 A Source INSTANCE, when acting as RECEIVER, MUST respond to a wsrf-rp:GetResourceProperty request where the child of the wsrf-rp:GetResourceProperty is the QName rbyteio:Readable, with the rbyteio:Readable element “<rbyteio:Readable>true</rbyteio:Readable>”.

R0506 A Source INSTANCE, when acting as RECEIVER, MUST respond to a wsrf-rp:GetMultipleResourceProperties request where the child of the wsrf-rp:GetMultipleResourceProperties contains the QName rbyteio:Readable, with the rbyteio:Readable element “<rbyteio:Readable>true</rbyteio:Readable>”, among the values requested for other Resource Properties.

5.2 Sink Resource Properties

This section of this Profile incorporates the following specification by reference:

- OGSA-DMI Functional Specification 1.0 draft 50 [7]

OGSA-DMI Functional Specification defines the following Information Element

- [dmi:Protocol] – Section 4.2.1.2

that describes the availability of the described protocol for writing the data to the Sink. This Profile mandates modelling the Sink as WS-Resources in Section 4, and mandates using Logical Data Transfers, also in Section 4. Using Logical Data Transfers the Data Transfer Factory needs to learn about the available protocols for the Sink before it can find a suitable data transfer protocol. Consequently:

R0507 A Sink INSTANCE MUST contain in the Resource Property Document schema, as referred to by the `wsrf-rp:ResourceProperties` attribute on the `portType`, a Resource Property Element `dmi:Protocol` as defined in OGSA-DMI Functional Specification 1.0 draft 50, Section 12.

R0508 A Sink INSTANCE, when acting as RECEIVER, MUST respond to a `wsrf-rp:GetResourceProperty` request where the child of the `wsrf-rp:GetResourceProperty` is the QName `dmi:Protocol`, with the complete list of supported source protocols each serialised into one `dmi:Protocol` element.

R0509 A Sink INSTANCE, when acting as RECEIVER, MUST respond to a `wsrf-rp:GetMultipleResourceProperties` request where the child of the `wsrf-rp:GetMultipleResourceProperties` contains the QName `dmi:Protocol`, with the complete list of supported source protocols each serialised into one `dmi:Protocol` element, among the values requested for other Resource Properties.

5.2.1 Sink Resource Properties Values

This section of this Profile incorporates the following specification by reference:

- OGSA-DMI Functional Specification 1.0 draft 50 [7]
- OGSA-ByteIO Specification 1.0 [8]

The OGSA-DMI Functional Specification defines the following Information Elements and Data Transfer Protocol identifiers:

- [dmi:Protocol] – Section 4.2.1.2
- Data Transfer Protocol Identifiers – section 4.1.1

that together describe the support of the identified data transfer protocol at the data transfer Sink.

OGSA-ByteIO Specification 1.0 defines the following ByteIO property (amongst others):

- “Writeable” – Section 4.1

The OGSA-ByteIO specification mandates the support of this property. However, the specification allows this property to have the Boolean value “false” which would be incorrect for the Sink of a DMI-governed data transfer (the defining characteristic of a Sink is that data can be written to it).

This Profile mandates that the Sink of a DMI-governed data transfer supports OGSA-ByteIO “Random-Access” protocol as a base line data transfer to write the data to the Sink. It further mandates that the Sink must allow write access to the data if it is invoked using the ByteIO Random-Access protocol.

R0510 A Sink INSTANCE MUST support the OGSA-ByteIO Random-Access protocol as a data transfer protocol. A Sink INSTANCE MUST provide one dmi:Protocol element that identifies OGSA-ByteIO using the URI <http://www.nextgrid.org/datatransferbasicprofile/protocols/ogsa-byteio/v1.0/random-access>”.

R0511 A Sink INSTANCE, when acting as RECEIVER, MUST respond to a wsrf-rp:GetResourceProperty request where the child of the wsrf-rp:GetResourceProperty is the QName rbyteio:Writeable, with the rbyteio:Writeable element “<rbyteio:Writeable>true</rbyteio:Writeable>”.

R0512 A Sink INSTANCE, when acting as RECEIVER, MUST respond to a wsrf-rp:GetMultipleResourceProperties request where the child of the wsrf-rp:GetMultipleResourceProperties contains the QName rbyteio:Writeable, with the rbyteio:Writeable element “<rbyteio:Writeable>true</rbyteio:Writeable>”, among the values requested for other Resource Properties.

6 References

- [1] D. Snelling, M. Fisher, A. Basermann, F. Wray, P. Wieder and M. Surridge, *NextGRID Vision and Architecture White Paper V5*, v18, NextGRID Project, 17 May 2007, http://www.nextgrid.org/download/publications/NextGRID_Architecture_White_Paper.pdf
- [2] V. Li, A. Anjomshoaa and D. Snelling, *NextGRID Basic Profile V1.0*, NextGRID Project, 2 January 2008.
- [3] S. Bradner, RFC 2119: *Key words for use in RFCs to Indicate Requirement Levels* (the Internet Engineering Task Force Best Current Practice), March 1997. <http://www.ietf.org/rfc/rfc2119>
- [4] K. Ballinger, D. Ehnebuske, C. Ferris, M. Gudgin, C. K. Liu, M. Nottingham, and P. Yendluri (eds.), *Basic Profile Version 1.1*, Web Services Interoperability Organization Final Material, 10 April 2006. <http://www.ws-i.org/Profiles/BasicProfile-1.1.html>
- [5] Bray, T., Hollander, D., Laymann, A., Tobin, R.: *Namespaces in XML 1.0*, W3C Recommendation, 16 August 2004. <http://www.w3.org/TR/REC-xml-names>
- [6] I. Foster, T. Maguire, D. Snelling, GFD-R-P.72: *OGSA® WSRF Basic Profile 1.0*, Open Grid Forum, 1 May 2006, <http://www.ogf.org/documents/GFD.72.pdf>
- [7] M. Antonioletti, M. Drescher, A. Luniewski, S. Newhouse and R. Madduri, *OGSA-DMI Functional Specification 1.0*, Public Review Draft 50, OGSA-DMI Working Group, Open Grid Forum, 1 February 2008, http://www.ogf.org/Public_Comment_Docs/Documents/2008-02/ogsa-dmi-spec-editor-submission.pdf
- [8] M. Morgan (Ed.), *ByteIO Specification 1.0*, GFD.87, Open Grid Forum, 31 October 2006.