

A GridRPC Model and API for End-User Applications - Summary for Submission to the OGSG -

Copyright Notice

Copyright © Global Grid Forum (2005). All Rights Reserved.
Copyright © Open Grid Forum (2007). All Rights Reserved.

Abstract

This document presents a model and API for GridRPC, i.e., a remote procedure call (RPC) mechanism for grid environments. Specifically this document is targeted for end-user applications, not middleware. That is to say, this document presents a simpler version of the GridRPC model and API that is completely sufficient for end-users and does not include the more complex features and capabilities required for building middleware. As a Recommendations track document in the Global Grid Forum, the goal of this document is to clearly and unambiguously define the syntax and semantics for GridRPC, thereby enabling a growing user base to take an advantage of multiple implementations. The motivation for this document is to provide an easy avenue of adoption for grid computing, since (1) RPC is an established distributed computing paradigm, and (2) there is a growing user-base for network-enabled services. By doing so, this document will also facilitate the development of multiple implementations.

1. Introduction

The goal of this document is to clearly and unambiguously define the syntax and semantics for GridRPC, a remote procedure call (RPC) mechanism for grid environments, thereby providing an avenue of easy access to grid computing. Specifically this document is targeted for end-user applications that do not require the more complex features and capabilities required for middleware packages. As such, it is outside the scope of this document to review or discuss those issues related to middleware, or the important issues related to network-enabled services or to provide any kind of tutorial information. Nonetheless, a Related Work section is provided to capture many references and pointers to relevant works that have lead up to this document. A preliminary version of this model and API appeared as [15]. A longer version of that paper is available as [16]. Comparison with CORBA is shown as [21].

2. The Basic GridRPC Model




Figure 1. The Basic GridRPC Model.

Figure 1 illustrates the basic GridRPC model. The functions shown here are very fundamental and, hence, appear in many other systems. A service registers with a registry. A client subsequently contacts the registry to look-up a desired service and the registry returns a handle to the client. The client then uses the handle to call the service which eventually returns the results.

In the GridRPC terminology adopted here, the service handle is a *function handle* which represents a mapping from a simple, flat function name string to an instance of that function on a particular server. Once a particular function-to-server mapping has been established by initializing a *function handle*, all RPC calls using that function handle will be executed on the server specified in that binding. A *session ID* is an identifier representing a particular non-blocking GridRPC call. The *session ID* is used throughout the API to allow users to obtain the status of a previously submitted non-blocking call, to wait for a call to complete, to cancel a call, or to check the error code of a call.

3. Document Scope

This simple, common model nonetheless represents multiple fundamental issues. It is clearly impossible to deal with them all at the same time. Hence, we now clarify what this document defines and does not define.

3.1 In Scope

This document focuses on just defining the API and the minimal programming model needed to understand and use the API for end-user applications. More specifically, it focuses on simple, client-server interaction since this comprises the majority of usage scenarios.

3.2 Out of Scope

The following topics are very important but are nonetheless out of the scope of this document:

- **Middleware.**

Middleware must be able to deal with situations that don't typically arise in end-user code, e.g., a variable number of arguments in a specific GridRPC call that is not known until call time.

- **Service Discovery.**

How the actual service registry or look-up is done is not addressed in this document. It is assumed that some type of registry or grid information service is available to accomplish this function.

- **Non-flat Service Names.**

The current API assumes simple name strings for GridRPC services. Describing and discovering GridRPC services by attributes or metadata schemas would certainly be very useful but is not addressed here.

- **General Workflow.**

Defining general mechanisms for managing grid workflows are not in the scope of this document. However, simple extensions to the API may be possible that allow the use of workflow management tools.

- **Interoperability between Implementations.**

Since this document focuses on the GridRPC API, it says nothing about the protocols used to communicate between clients, servers, and registries. Hence, it does not address interoperability.

4. The GridRPC API

We begin the presentation of the GridRPC API by defining the data types used. We then present the initialization/finalization calls, function handle management calls, the function calls

themselves, and the control and wait calls. Each call definition includes a table of possible error codes that it can return.

4.1 GridRPC Data Types

grpc_function_handle_t

Variables of this data type represent a specific remote function that has been bound to a specific server which might be chosen by underlying GridRPC system. They are allocated by the user. After a *function handle* is initialized, it may be used to invoke the associated remote function as many times as desired. The lifetime of a *function handle* is determined when the user invalidates the *function handle* with a *handle destruct* call.

grpc_sessionid_t

Variables of this data type represent a specific non-blocking GridRPC call. *Session IDs* are used to probe or wait for call completion, to cancel a call, or to check the error status of a call. *Session IDs* are also allocated by the user but their lifetime is determined automatically. A *session ID* is initialized when a non-blocking GridRPC call is made. It is invalidated, or destroyed, when (1) all return arguments have been received, and (2) a wait function has returned a "call complete" status to the application. If an invalid *session ID* is passed to any GridRPC call, an error will result.

grpc_error_t

This data type is used for all error and return status codes from GridRPC functions.

4.2 Initializing and Finalizing Functions

The initialize and finalize functions are similar to the MPI initialize and finalize calls. Client GridRPC calls before initialization or after finalization will fail.

grpc_error_t grpc_initialize(char *config_file_name)

This function reads the configuration file and initializes the required modules. After this function is called once, subsequent call will return with GRPC_ALREADY_INITIALIZED.

grpc_error_t grpc_finalize(void)

This function releases any resources being used by GridRPC, canceling all the unfinished asynchronous calls.

4.3 Remote Function Handle Management Functions

The *function handle management* group of functions allows the creation and destruction of function handles.

grpc_error_t grpc_function_handle_default(grpc_function_handle_t *handle, char *func_name)

This creates a new function handle using a default server associated with the given function name. This default could be a pre-determined server or it could be a server that is dynamically chosen by the resource discovery mechanisms of the underlying GridRPC implementation. The server selection process could be postponed until the actual call is made on the handle. Once selection is made, all the calls through the handle must go to the server.

grpc_error_t grpc_function_handle_init(grpc_function_handle_t *handle, char *server_name, char *func_name)

This creates a new function handle with a server explicitly specified by the user.

grpc_error_t grpc_function_handle_destruct(grpc_function_handle_t *handle)

This releases all information and resources associated with the specified function handle. It also cancels a running session bound to the handle, if exists, before releasing the handle itself.

```
grpc_error_t grpc_get_handle(
    grpc_function_handle_t **handle,
    grpc_sessionid_t sessionID)
```

This returns the function handle corresponding to the given **session ID** (that is, corresponding to that particular non-blocking request).

4.4 GridRPC Call Functions

Two GridRPC call functions are available for end-users. These two calls are similar but provide either blocking (synchronous) or non-blocking (asynchronous) behavior. In the non-blocking case, a *session ID* is returned that is subsequently used to test for completion.

```
grpc_error_t grpc_call( grpc_function_handle_t *handle, <varargs> )
```

This makes a blocking remote procedure call with a variable number of arguments.

```
grpc_error_t grpc_call_async(
    grpc_function_handle_t *handle,
    grpc_sessionid_t *sessionID,
    <varargs> )
```

This makes a non-blocking remote procedure call with a variable number of arguments. A *session ID* is returned that can be used to probe or wait for completion, cancel the call, and check for the error status of a call.

The GridRPC Recommendation does not define which implementation-related operations may be assumed to be complete when an asynchronous call returns. However, all asynchronous GridRPC calls must return as soon as possible after it is safe for a user to modify any input argument buffers.

4.5 Asynchronous GridRPC Control Functions

The following functions apply only to previously submitted non-blocking requests.

```
grpc_error_t grpc_probe( grpc_sessionid_t sessionID)
```

This call checks whether the asynchronous GridRPC call represented by the *session ID* **sessionID** has completed. If it has completed, **GRPC_NO_ERROR** is returned. Otherwise, **GRPC_NOT_COMPLETED** is returned.

```
grpc_error_t grpc_probe_or(
    grpc_sessionid_t *idArray,
    size_t length,
    grpc_sessionid_t *idPtr )
```

This call checks whether the asynchronous GridRPC calls represented by the array of *session IDs* in **idArray** have completed. If any calls have completed, the function return value is **GRPC_NO_ERROR** and the **grpc_sessionid_t** pointed to by ***idPtr** contains exactly one valid, completed call. If no call has completed, the function return value is **GRPC_NONE_COMPLETED** and the **grpc_sessionid_t** pointed to by ***idPtr** is undefined. If any of the *session IDs* in **idArray** are invalid, no operations will occur and an **GRPC_INVALID_SESSION_ID** error will be returned. However, the array of *session IDs* may contain completed *session IDs* without causing an error.

```
grpc_error_t grpc_cancel( grpc_sessionid_t sessionID )
```

This cancels the specified asynchronous GridRPC call.

```
grpc_error_t grpc_cancel_all( void )
```

This cancels all outstanding asynchronous GridRPC calls.

Rationale:

A “cancel array” call was considered but dismissed since it would cause difficult error handling.
End of Rationale.

4.6 Synchronous GridRPC Wait Functions

The following five functions apply only to previously submitted non-blocking requests. These calls allow an application to express desired non-deterministic completion semantics to the underlying system, rather than repeatedly polling on a set of *session IDs*.

grpc_error_t grpc_wait(grpc_sessionid_t sessionID)

This blocks until the specified non-blocking requests to complete.

grpc_error_t grpc_wait_and(
 grpc_sessionid_t *idArray,
 size_t length)

This blocks until *all* of the specified non-blocking requests in **idArray** have completed.

grpc_error_t grpc_wait_or(
 grpc_sessionid_t *idArray,
 size_t length,
 grpc_sessionid_t *idPtr)

This blocks until *any* of the specified non-blocking requests in **idArray** has completed. On a successful return, **idPtr** points to a completed request.

grpc_error_t grpc_wait_all(void)

This blocks until *all* previously issued non-blocking requests have completed.

grpc_error_t grpc_wait_any(grpc_sessionid_t *idPtr)

This blocks until *any* previously issued non-blocking requests has completed. On a successful return, **idPtr** points to a completed request.

4.7 Error codes and Error Reporting Functions

When a GridRPC call fails, an error code is returned. The table is omitted here for brevity.

These error codes satisfy:

0 = GRPC_NO_ERROR < GRPC_... < GRPC_LAST_ERROR_CODE

This specifies a useful numerical ordering of the error codes based on the set of integers without specifying a specific implementation.

The ability to check the error code of previously submitted requests is provided. The following error reporting functions provide error codes and human-readable error descriptions. These error descriptions can be more informative about the actual cause of the error.

char *grpc_error_string(grpc_error_t error_code)

This returns the error description string, given a GridRPC error code. If the error code is unrecognized for any reason, the string **GRPC_UNKNOWN_ERROR_CODE** is returned.

grpc_error_t grpc_get_error(grpc_sessionid_t sessionID)

This returns the error code associated with a given non-blocking request.

grpc_error_t grpc_get_failed_sessionid(grpc_sessionid_t *idPtr)

This returns the *session ID* associated with the most recent **GRPC_SESSION_FAILED** error. This provides additional error information on a specific *session ID* that failed for calls that deal with sets of *session IDs*, either implicitly, such as **grpc_wait_all()**, or explicitly, such as **grpc_wait_and()**. When there are more than two failed sessions, this function will return the session ID one by one. To make sure that all the failed sessions are handled, users have to call this function repeatedly until it returns **GRPC_SESSIONID_VOID**.

5. Security Considerations

Security issues of GridRPC are implementation-dependent and this document does not specifically address security in the API. For reference, security mechanisms of Ninf-G and NetSolve are described in this section. Security infrastructure of Ninf-G is based on GSI which is based on public key encryption, X.509 certificates, and the Secure Sockets Layer (SSL) communication protocol. This means that not only all the components are protected properly, but they can also utilize other Globus components, such as GridFTP servers, seamlessly and securely. NetSolve's current security is based on the ability to generate access control lists that are used to grant and deny access to the NetSolve servers. NetSolve uses Kerberos V5 services for authentication. The Kerberos extensions of NetSolve provide it with trusted mechanisms by which to control access to computational resources. At this time, the Kerberized version of NetSolve performs no encryption of the data exchanged among NetSolve clients, servers, or agents, nor is there any integrity protection for the data stream.

Author Contact Information

Hidemoto Nakada
National Institute of Advanced Industrial Science and Technology
hide-nakada@aist.go.jp

Satoshi Matsuoka
Tokyo Institute of Technology
National Institute of Informatics
matsu@is.titech.ac.jp

Keith Seymour
University of Tennessee, Knoxville
Seymour@cs.utk.edu

Jack Dongarra
University of Tennessee, Knoxville
dongarra@cs.utk.edu

Craig A. Lee
The Aerospace Corporation, M1-102
2350 E. El Segundo Blvd.
El Segundo, CA 90245
lee@aero.org

Henri Casanova
University of California, San Diego
San Diego Supercomputer Center
Casanova@cs.ucsd.edu

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.

Full Copyright Notice

Copyright (C) Global Grid Forum (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."