



## UDAP Use Cases and Requirements

|                 |                |      |
|-----------------|----------------|------|
| <b>Editors:</b> | A. Anjomshoaa  | EPCC |
|                 | Henning Mersch | FZJ  |
|                 | Ph. Wieder     | FZJ  |

| Date       | Author     | Comments           | Version | Status       |
|------------|------------|--------------------|---------|--------------|
| ...        | H. Mersch  | Initial version    | V0.1    | Draft        |
| 14/05/2007 | Ph. Wieder | Update             | V0.2    | Draft        |
| 09/2007    | Ph. Wieder | Response to review | V0.3    | Stable draft |
| 02/2008    | M. Memon   | Final version      | V1.0    | Final        |



- 1 INTRODUCTION ..... 3**
- 2 COMPONENTS ..... 4**
  - 2.1 Definition of Activity ..... 4**
    - 2.1.1 Activity Information ..... 4
  - 2.2 UDAP Concepts ..... 5**
    - 2.2.1 UDAP Document ..... 5
    - 2.2.2 UDAP Provider ..... 6
    - 2.2.3 UDAP Activity ..... 6
    - 2.2.4 Activity Registry ..... 6
    - 2.2.5 UDAP Advertiser ..... 7
    - 2.2.6 UDAP Client ..... 7
  - 2.3 UDAP as an Intermediary ..... 7**
- 3 USE CASES ..... 8**
  - 3.1 Job Submission ..... 8**
  - 3.2 Service and Resource Discovery ..... 8**
  - 3.3 Supporting an SLA Framework ..... 9**
  - 3.4 Supporting Data Centricity ..... 9**
  - 3.5 Accounting ..... 10**
  - 3.6 Archiving ..... 10**
- 4 UDAP REQUIREMENTS ..... 11**
  - 4.1 Scalability ..... 11**
  - 4.2 Security ..... 11**
- 5 REFERENCES ..... 12**



## 1 Introduction

In the current state of Grid architectures, information about an activity is fragmented and dispersed. Activity information, such as resource usage, security data, activity state, data requirements, et cetera, is currently captured using a variety of schemata and stored in different ways and by different logical components.

This dispersion of activity information leads to management, security and logistical overhead in discovering, accessing and using that information. This uncoordinated scheme for finding activity information assumes a lot of Grid components and Grid system coordinators. It results in an environment where activity information is managed by many systems. This makes it difficult for Grid components that rely on that information to know where the information is, leading to them having to search for and to “keep an eye” on many sources of activity information.

The Universal Dynamic Activity Package (UDAP) aims to bring all of the information fragments that are associated with an activity, regardless of the various schemata that are used to describe and capture these fragments, into one logical package. It then aims to specify i) how to write activity information to that package, ii) how to query that package for activity information, iii) how to extract information from that package for read and update operations, and iv) how to insert new activity information into that package.

Using UDAP, any Grid component requiring activity information has a single source for retrieving and using that information. Grid components can subscribe to the interface of an entity managing an activity, so that they may be notified of changes in activity information, thus allowing them to fulfil their role with respect to that activity. The same Grid components can poll the UDAP management interface for activity information in read, update and append operations, in addition to or instead of subscribing for notifications.

The subscribe-notify capability of the UDAP components leads to many interesting use cases, some of which are introduced later in this document (see Section 3). If applied as a component central to a system, the variety of UDAP use cases is not limited. In contrast, looking at UDAP from a technical angle, the number of use cases is very much limited to what has said before: write, query, read, update, and insert information about an activity through a subscribe-notify mechanism. In general, it means that Grid components can simply wait for activity information generated events that would trigger particular actions. This leads to an event based activity management environment. Nevertheless, to outline potential usage scenarios of UDAP in a NextGRID context, this document describes a number of activity-related use case and the resulting requirements.

## 2 Components

The Universal Dynamic Activity Package (UDAP) is a scheme for managing all the information associated with an activity on a Grid. This in turn allows any Grid component or resource which is concerned with that activity to access and use that information using the UDAP scheme to do its task.

### 2.1 Definition of Activity

An activity is a *unit of work* on a Grid. It can be a job; a task; a data processing operation; a data access operation; an application execution; a program execution; a Web Service invocation; something that a user or application needs to do, take care of, or execute.

An activity is atomic. This means that an activity is an indivisible unit of work from an activity management perspective. When you stop an activity, you stop all of it, not any one part of it.

And activities can be composed together to form chains of activities that may be managed conditionally, sequentially, or in parallel. Activities can be the atomic nodes in a workflow. Further, they can be conditionally used to process data in a data centric process.

Furthermore, an activity is considered to be a resource. The realised activity must implement the management interfaces of a resource so that it may be managed as such on a Grid. This leads to activities being dependent on and used by other resources. An activity may require another activity in its list of resource requirements. The statements: “depending on an activity” and “requiring an activity resource” are syntactically identical.

#### 2.1.1 Activity Information

We take a holistic view of an activity. We consider all that there is to know about an activity:

- all of its requirements,
- all of its dependencies (on data and other activities) for the composition and management of activities for workflow, scheduling and brokering processes,
- all of its contextual information, such as:
  - topical domain (financial markets, weather forecasting, etc.),
  - security (who owns the activity, who is allowed to run it, etc.),
  - SLAs, QoS and other related policies, and
- all of its monitoring information, such as:
  - state,
  - history,
  - resource information,
  - accounting,
  - policy conformance, et cetera.

Activity information therefore includes such things as:

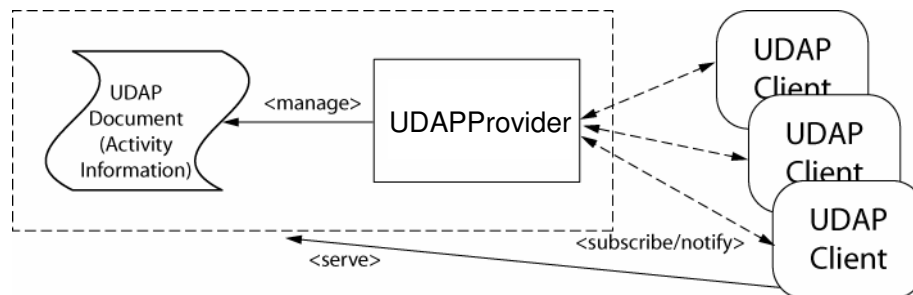
- resource requirements (hardware, software, services, data, etc.),
- dependencies on data and other activities,
- subsistence state,
- security and ownership context,
- migration rules and policies,
- SLA policies governing transactions,
- accounting and auditing,
- history and provenance, and
- any other information associated with the activity.

## 2.2 UDAP Concepts

The core of the UDAP model is the UDAP document and in terms of implementation the respective UDAP Activity resource. Its information is managed by a UDAP Provider and processed by UDAP clients. In the following, these concepts are described in detail. The respective implementations of the concepts are outlined in the “Resource Discovery using the UDAP System” recipe of the NextGRID cookbook [1]. Please note that the existing components are implemented according to the NextGRID Basic Profile [2] and the upcoming implementations have to follow the Basic Profile, too.

### 2.2.1 UDAP Document

At the heart of the UDAP model is the UDAP Document (the relation between UDAP concepts is pictured in Figure 1). This conceptual document is the package that holds all the information associated with an activity. It needs a schema structure that allows it to classify the information that it holds for an activity, and to reflect the state of that activity at any point in time. This schema



**Figure 1** The figure shows the UDAP Provider in charge of a UDAP Document for an activity (represented by the dashed box) on a Grid. UDAP Clients concerned with that activity may subscribe to the UDAP Provider so that they can be notified of changes in the activity.

is defined as an XML Schema named “UDAP Schema” which is described in the respective NextGRID Generalised Specification [3].

## 2.2.2 UDAP Provider

The UDAP Provider is a management component that deals with the management-oriented aspects of all other UDAP components. It is the single entity responsible for catering the management function of other entities. These functions include the creation and destruction of managed UDAP Documents.

Besides managing activities, the UDAP Provider handles the interaction with Negotiator components, which leverage its use to allow the SLA artefacts in the form of activities in order to be managed and persistent for negotiation and discovery use cases. One of the use cases is to manage agreed SLAs while provider's negotiator reaching towards the SLA agreement. In this perspective, negotiation is not the only component which is interacting with UDAP Provider, also operational managements' service delegating some of its responsibilities to this component. In order to foster the SLA management, the useful functions are to add and update the agreed SLA properties via WSRF standards' based service interfaces.

## 2.2.3 UDAP Activity

As illustrated in Figure 2, the *UDAP Activity* component is the central component of the UDAP Provider. In each interaction with the activities, the UDAP Provider acts a factory interface for instantiating UDAP Activity resources. A UDAP Activity could be a stateful web service resource; thereby it exposes all of the attributes and lifetime characteristics in the form of interfaces. A UDAP Activity represents UDAP Document as a remote interface which can be represented as a *resource properties document* containing attributes related to activity type and states. A UDAP Activity is an abstract resource representation of an activity; its specializations can supplement the resource properties and functions and extended it with domain specific interfaces.

Since UDAP Activity contains UDAP Document, therefore the interface possesses the incoming and outgoing requests as instances of UDAP activities. Each activity is according to its history, state, and description associated with a particular provider. Therefore the main requirement concerning the Activity Registry is to provide tracking and archiving of information is to enable the logging, and monitoring, etc. of registered activities. The Activity Registry is an intrinsic part of UDAP and the underlying component for incorporating activity information persistently.

## 2.2.4 Activity Registry

The UDAP Model and Framework are activity-oriented and all the data records which are registered are considered as activities. Each activity is according to its history, state, and description associated with a particular provider. Therefore the main requirement concerning the ActivityRegistry is to provide tracking and archiving of information to enable logging, monitoring, etc. of registered activities. The Activity Registry is an intrinsic part of UDAP and the underlying component for incorporating activity information persistently. It is one of the essential components for executing the resource discovery use case.

## 2.2.5 UDAP Advertiser

The UDAP Advertiser is another integral concept of UDAP. It performs the match-making of computing resources which have already been published in to the registry. The registration is performed by the ActivityRegistry where all the resource-specific information with respect to the schema is exposed. The candidate resources are registered by sending a request to the UDAP Advertiser in the form of JSDL-compliant information, which is then transformed into the UDAP Document and other elements if they are also part of the request into the information schema instance understood by the ActivityRegistry. In case of discovering compute resources, the UDAP Advertiser accepts requests in the form of JSDL documents.

## 2.2.6 UDAP Client

A UDAP Client is any Grid component that uses and/or produces activity information. UDAP Clients invoke the public management interface of the UDAP Provider for read, update and append operations of activity information.

A UDAP Client can subscribe to the interface of a UDAP Activity Resource in order to receive notification of activity information based events. The subscription of a UDAP Client may be conditional, where the condition dictates the type of activity information based event that the client is interested in, e.g. “notify me if the state of the activity changes to running” or “notify me if the resource usage of the activity has exceeded the budget of the activity owner”.

## 2.3 UDAP as an Intermediary


An interesting feature of UDAP is that it acts as an intermediary<sup>1</sup> between activity owners and the service and resource providers that serve that activity (see Figure 2). This is as a result of the UDAP Activity Resource being at the centre of the interactions between the activity owner and the service providers.

The activity owner, who initiates the activity, must do so through a Grid component, such as a job submission client, or Grid portal. That Grid component, then, in turn is a UDAP Client and can subscribe to the UDAP Activity Resource that it has initiated for activity information based events.

This intermediary quality of UDAP leads to a symmetric, uniform interface for activity management, with the activity initiating Grid component on one side, and the service providers that serve that activity on the other, all interacting with and using the Activity Resource using the same interface.

---

<sup>1</sup> The idea of an UDAP Intermediary has been massively used within the first realisation of the UDAP system, mainly to de-couple resource provider and resource consumer concerning the advertisement and request of resources. Please refer to **Error! Reference source not found.** for a detailed description.



**Figure 2 UDAP as an intermediary between an activity owner/initiator and the service serving that activity. The “UDAP Activity” box represents the abstract activity entity.**

### 3 Use Cases

The potential use cases for UDAP are manifold and not limited by the model itself. In this section we list a number of these cases to give an idea of what can be achieved by the application of the UDAP model.

#### 3.1 Job Submission

A Grid user, application or job submission client discovers a UDAP Provider from a registry. It then invokes that service to instantiate a new UDAP Activity Resource Service by submitting to it a description of the activity which includes the activity’s requirements and any contextual and dependency information about that activity.

The UDAP Provider performs a search on its list of subscribers (subscribed UDAP Clients) to find suitable ones in order to fulfil the activity. If it finds suitable UDAP Clients that match the requested activity’s requirements, it then instantiates a WS-RF UDAP Activity Resource.

The UDAP Provider then notifies the suitable UDAP Clients it has found of the activity’s existence, sending them the Activity Resource Service’s EPR. The UDAP Clients identified can then contact the new activity resource instance to find out further details of its requirements. If they are able to fulfil those requirements, an SLA will be agreed between each UDAP Client and the activity resource. The activity can then be served by those UDAP Clients that have an SLA in place with it.

#### 3.2 Service and Resource Discovery

Taking up a common scenario in Grids, a client-side agent requests compute resources and expects a list of potential service providers back from the UDAP system.

Within the setup procedure, the UDAP Provider is started and waits for incoming subscription requests. UDAP clients can subscribe themselves to the UDAP Provider, passing as arguments their EPR address [4] and an XPATH [5] query expression about their capabilities for processing UDAP Documents, i.e. the service or resource capabilities they provide.

Following this, the UDAP system is set up and can be used by UDAP clients generating activities. This client retrieves information about the resources to query from some other component like an end user interface or a broker. The client then generates, based on the information it has retrieved, a UDAP Document and submits this to the UDAP Provider, which implicitly means it will subscribe to changes on that UDAP Document. The UDAP Provider will now execute all capability descriptions of the resource providers to the UDAP Document and – in case they match – inform the resource providers with a reference to the UDAP Document.

The UDAP Subscriber could now retrieve this UDAP Document and process it by interpreting the content. Afterwards, the UDAP clients will modify the UDAP Document at the UDAP Provider. In a common scenario, where services or resources are to be discovered based on their capabilities, this would mean to insert one or more service provider EPRs.

Every time a modification is made to the UDAP Document, the (implicitly subscribed) UDAP client will be informed about the changes, which will be service provider EPRs for our concrete scenario. At any time new UDAP clients acting as subscribers could be added to the UDAP Providers, which enhances the facilities of the overall system.

### **3.3 Supporting an SLA Framework**

The NextGRID SLA Framework [6] requires that each pair of interacting Web Services must have in place an SLA that governs their interaction in terms of guaranteed quality of service.

In the UDAP scheme, this requirement leads to there being an agreed SLA in place between each UDAP Client-UDAP Activity Resource Service pair. This SLA would set out the levels of service and quality expected from a UDAP Client when serving a UDAP activity. It will also set out the conditions on the activity under which it will be served, such as the required security conditions being fulfilled and an appropriate price being paid for the service provided.

An activity being served by many UDAP Clients would have an agreed SLA in place with each UDAP Client.

### **3.4 Supporting Data Centricity**

Data centricity may be defined as the set of operations that need to be conditionally performed on some data of interest. These operations include:

- producing,
- storing,
- transporting,
- filtering,
- transforming,
- processing,
- accessing,
- reading,
- writing,

- updating,
- archiving,
- et cetera.

When user or application requirements are data centric, “activity operations”, such as those listed above, must be performed on the specified data as a result of events generated due to previous operations performed on that data, or due to the state of that data.

This leads to a workflow of activity nodes, that are conditionally executed based on events generated as a result of previous operations.

The UDAP scheme enables an environment to be built, where activities may be composed together and conditionally executed by a data centric job manager (a workflow engine) based on events, in order to meet the needs of a data centric process.

### **3.5 Accounting**

When an activity is being fulfilled by one or more UDAP Clients, there are accounting requirements that govern the use of the resources.

Accounting information may be kept and updated in the activity’s information directory (the WS-Resource Properties (WS-RF) Document [7]). When new accounting information is available, events can be generated and sent to subscribe UDAP Clients which are interested in that information.

### **3.6 Archiving**

When an activity is being fulfilled by one or more UDAP Clients, there may be provenance and archiving requirements for that activity.

The UDAP schema allows for provenance and archiving of activity information to be captured in the schema. However, this may not be practical. As activities progress in their life, the provenance and archiving data that they carry may blow up and hinder performance of the UDAP Activity Resource Service. This should be borne in mind when implementing the UDAP Provider and UDAP Document concepts.

A good way to capture provenance and archiving information about an activity is to have a service that is dedicated to undertaking provenance and archiving. The implementation of this service may be specialised to do provenance and archiving in whatever way or form and using whichever schema. The important point to bear in mind is that this provenance and archiving service can be a UDAP Client which can subscribe to a UDAP Activity Resource to receive notification every time a piece of information about the activity changes (that information being captured in the Resource Properties Document of a WS-RF UDAP Activity Resource Service).

Provenance and archiving have a number of use-cases in themselves. For example, it may be necessary to undertake provenance and archiving for fault tolerance, in case of catastrophic failure



of services or resources that are serving an activity. Under such a circumstance, the provenance and archiving data will act as a snapshot of the “memory footprint” of the activity’s information and state and can be used for the activity to be migrated to other services or resources.

## 4 UDAP Requirements

UDAP manages various information belonging to (as prior defined) Activities, thus it has to be informed and used to manage this information from various components. Nevertheless, all components interacting with UDAP will directly benefit from its capabilities like acting as intermediary. Thus we expect UDAP to be one of the central components of a NextGRID implementation.

### 4.1 Scalability

For the UDAP Framework, the crucial component when it comes to scalability, is the UDAP Provider. Due to the fact it keeps track of the inserted UDAP Activities, we currently think about a scalable system of UDAP Providers. Another UDAP Provider could subscribe for newly and/or otherwise un-processable UDAP Activities and will be notified to process a UDAP Activity from the first UDAP Provider. This would lead to a kind of a *Peer-to-Peer*-Network of UDAP Providers, which is envisaged to scale very well. Due to the subscription/notification mechanism used, such a UDAP Provider network should be not too difficult to realise. Please note that this feature is not yet planned to be implemented to the UDAP Framework, but initial design considerations are under way.

### 4.2 Security

UDAP Clients that are able to fulfil the requirements of an activity, must match the requirements stated in the activity’s security context. The security requirements will most likely be bipartite, and so the activity must match the requirements of the UDAP Clients.

The ability to satisfy bipartite security requirements between the activity and a UDAP Client should be captured in the SLA between them.

Furthermore, details of the fulfilment of an activity’s security requirements by UDAP Clients must be put into that activity’s information set (the UDAP Document rendered as the WS-RF Resource Properties Document).



## 5 References

- [1] K. Tserpes, NextGRID Cookbook, NextGRID, March, 2007.
- [2] V. Li and D. Snelling, NextGRID Basic Profile, NextGRID, June, 2006.
- [3] A. Anjomshoaa and Ph. Wieder, UDAP Schema, NextGRID Generalised Specification, NextGRID, May, 2007.
- [4] M. Gudgin and M. Hadley (eds.), Web Services Addressing 1.0 – Core, W3C Candidate Recommendation, W3C, 17 August 2005. <http://www.w3.org/TR/2005/CR-ws-addr-core-20050817/>.
- [5] A. Berglund, S. Boag, D. Chamberlin, M. F. Fernández, M. Kay, J. Robie, and J. Siméon, XML Path Language (XPath) 2.0, W3C Candidate Recommendation, W3C, June, 2006. <<http://www.w3.org/TR/xpath20/>>.
- [6] B. Koller, Ph. Masche-Pakkala, B. Mitchell, and H. Mizani, NextGRID SLA Framework, NextGRID Project Output, P4.5.6, September, 2006.
- [7] S. Graham and J. Treadwell (eds.) Web Services Resource Properties 1.2 (WS-ResourceProperties), OASIS Standard, 1 April 2006. [http://docs.oasis-open.org/wsrp/wsrp-ws\\_resource\\_properties-1.2-spec-os.pdf](http://docs.oasis-open.org/wsrp/wsrp-ws_resource_properties-1.2-spec-os.pdf)