Agreement Signalling and Network Service Provisioning for Grids

Tiziana Ferrari*, Kostas Kavoussanakis†, Charaka Palansuriya†, Anand Patil‡, Elisabetta Ronchieri*

*Istituto Nazionale di Fisica Nucleare CNAF, IT (Email: {Tiziana.Ferrari, Elisabetta.Ronchieri}@cnaf.infn.iti)
†University of Edinburgh, UK (Email: {kavousan, charaka}@epcc.ed.ac.uk)
‡DANTE, UK (Email: anand.patil@dante.org.uk)

***Abstract*— Resource-sharing is one of the Grid intrinsic properties enabling efficient usage of services by clients from many Virtual Organizations. This is particularly true for network infrastructures, which are typically designed to serve a large number of heterogeneous users. In such an environment, the support of network Quality of Service requires the establishment of Service Level Agreements between users and network service providers, admission control as well as network control planes for automatic or semi-automatic service provisioning. The availability of guarantees allows flexible usage of the distributed infrastructure and enables application-controllable performance. Nevertheless, many network providers today only offer best-effort services and do not support Service Level Agreement signalling. In this paper we analyze requirements and introduce two bandwidth-oriented services for enhanced Grid data transfer: the *Guaranteed Delivery File Transfer* and *Virtual Leased Line* services. We propose a complete architecture – fully integrated with Grid middleware – for network Service Level Agreement management and inter-domain network service provisioning. This is based on the functional division into agreement and service provider layers proposed by the Grid Resource Allocation and Agreement Protocol working group of the Global Grid Forum.**

I. Introduction

Advance reservation is the mechanism that allows a user to request exclusive access, for a specific future time interval, to a set of resources that satisfy specific requirements. This mechanism is useful to various mission-critical applications, for example to Grid applications to reserve network bandwidth for bulk transfers; to Data Scheduling services requiring storage space reservation prior to the start of a data transfer session; and to Workload Management services to allocate computing power (e.g., to high-priority execution tasks). An advance reservation request obtains the full specification of the resources needed through a set of resource-specific attributes. It may supply run-time information at a later stage through a binding operation. Reservation allows the exclusive access to a resource pool, while allocation enforces the Quality of Service (QoS) guarantees.

The application of a specific QoS level to a Grid activity – regardless of its nature – requires the negotiation and establishment of a contract, named *Service Level Agreement* (SLA), between the customer or proxy, and the service provider. An SLA quantitatively defines the performance level for the service requested and it defines the obligations for the parties involved in the contract. The SLA is typically specified through a template containing both quantitative and qualitative information. In particular, an SLA specification supplies administrative information (e.g., the identity of the entities involved in the contract, and the penalties applicable to the parties when the SLA guarantees are not honored), whereas the Service Level Specification (SLS) is a set of attributes and values describing the profile of the requested performance level [1].

The Grid Resource Allocation and Agreement Protocol working group (GRAAP-WG) [2] of the Global Grid Forum (GGF) proposes an architecture for Grid SLA signalling comprising two functional levels: the *agreement layer* and the *service provider layer* [3]. The agreement layer implements the communication protocol used to exchange information about SLAs between the customer and the service provider, and it is responsible for ensuring that the SLA guarantees are enforced by a suitable service provider. In addition, the agreement layer exposes information about types of service and the related SLA templates offered, processes agreement creation requests to check the compliance of the so-called agreement offers (i.e. the SLA requests), and handles well-formed requests to the service provider layer. An agreement is successfully created if the corresponding guarantees can be enforced by one or more service providers, which are also responsible for supervising the status of their resource pools.

We propose a resource-independent architecture for SLA management, modelled according to the two-layer service architecture described above, and based on a novel agreement-layer service that we call *Agreement Service*. We show how the proposed service interacts with and relies on complementary Grid services. We then instantiate this architecture, by focusing on the problem of reservation management and provisioning the network context.

We initially identify a set of Grid application and middleware requirements, and we propose two services for bandwidth management specifically designed to assist Grid file transfer. We then detail the functionality needed for the negotiation of the corresponding reservations, relying on a novel service called *Bandwidth Allocation and Reservation*, which instantiates the Agreement Service from a functional point of view, and supports operations for reservation creation and deletion. We complete this architecture by proposing a service-provisioning model for inter-domain network environments and we specify the interfaces between the various components. The overall architecture proposed is fully integrated with Grid middleware developed in the framework of the EGEE project [4], in particular with the Grid Workload and Resource Management services (WMS), and the Data Management services. We conclude the paper with a survey of related work, some closing remarks and a description of future work.

II. Use Cases

Grid applications and middleware rely on the availability of a robust and reliable communication infrastructure, enabling the timely and efficient handling of heterogeneous data streams, which typically share common general-purpose links. Networks need to offer a service portfolio to address this large variety of requirements. In what follows, we describe some of the most interesting Grid network use cases. For a comprehensive analysis, see [5] and [6].

Applications in high-energy physics and astronomy require distributed analysis of enormous amounts of data. In the former case, data are generated by one central experimental installation (the so-called *Tier-0* site) and are distributed hierarchically to *Tier-1* computing centres and smaller *Tier-2* and *Tier-3* sites.

Similarly, in the *Electronic-Very Long Baseline Interferometry* (E-VLBI) experimental environment [7], data are gathered from a large number of peripheral nodes to a single, central computing site, where data are finally correlated and processed. These two examples illustrate how Grid application traffic patterns can easily generate bottlenecks. The problem is typically alleviated by replicating data to many different resources.

Bulk data replication is characterised by flows that can deliver data in a large time frame by deadline. Conversely, the exchange of high-priority traffic bulk data replication (such as the one generated by database queries during job execution), can be more sensitive to latency and it can require completion in a shorter time scale.

While replication typically involves the exchange of large amounts of traffic, mission critical control traffic, such as the flows generated by Grid control operations (for example queries, lookups, notification messages, etc.), typically produce smaller data volumes, but in this case the individual messages are sensitive to one-way delay and network reliability. Such traffic, which can be bursty in nature, is not time-windowed, and flows over a long period of time.

Visualization applications applied to biomedicine represent another important application category. They produce different flows, such as voice, video and remote-steering type of flows. Each stream requires slightly different bandwidth, one-way delay, IP packet delay variation (IPDV) and packet loss guarantees, but all of them need a robust network. Consequently, point-to-point, time-windowed telemedicine applications typically need real-time guarantees.

Similarly, the E-VLBI astronomy applications need a very high bandwidth with relatively strict bounds on IPDV, being typically composed by several flows that must all converge, in a synchronised fashion, to the same point for correlation. This gives rise to the need for guarantees on one-way delay and IPDV.

Finally, Grid applications, such as in biomedicine, can require a secure connection for confidentiality. Usually networks do not offer encrypted services but may use the confidentiality attribute, for example, as an indication to drive routing decisions.

III. Bandwidth Allocation and Reservation Services

Several are the services that can be identified to address the requirements illustrated in Section II. In this paper, we just focus on the support of file transfer enhanced performance based on *Bandwidth Allocation and Reservation* (BAR) services. In particular, we discuss two types: the *Guaranteed Deadline File Transfer* service (GDFT) and the *Virtual Leased Line* service (VLL). These services, together with *Video Quality* service and *Visualisation* service, were identified in [8] to address our use cases by taking into consideration the capabilities of the underlying network.

Guaranteed Deadline File Transfer: provides network transport between two end-points/sites guaranteeing the movement of a specified volume of data within a given period of time. This implies that only an average bandwidth is guaranteed over time, while the instantaneous bandwidth may vary. GDFT is a service designed to assist data replication.

Virtual Leased Line: provides a dedicated amount of capacity on the network path between two servers or network domain end-sites over a period of time. The actual bandwidth remains constant.

From an implementation point of view, the proposed BAR services can rely on different implementation techniques. For example, they can rely on the IP *Differentiated Services* traffic conditioning techniques [9], and, in particular, on the *Premium IP* Per-Domain Behaviour [10], which is currently supported by the GÉANT network and a number of European National Research and Education Networks (NRENs). Premium IP requires that [11]: *for the selected packets, capacity is conserved and hence packet loss is null or negligible, apart from bit error rate and other similar causes, but that loss is never due to congestion. For the selected packets, the one-way delay and IPDV along a path are independent of the load of the path and are very similar to the values obtained on an empty network. The IP Premium Service aims at providing users with the equivalent of a leased line*. Thanks to the quantitative guarantees on bandwidth offered, Premium IP lends itself well to the bulk data replication use case described in Section II.

The usage of the proposed services requires the full integration with the user application and Grid middleware. In order to do so, Grid infrastructures need to support mechanisms for automatic SLA signalling, management and monitoring.

In what follows we detail the architecture of the entire stack of protocols and agents that are needed to accomplish the tasks of reservation and allocation. In particular, the proposed SLA signalling approach is resource-independent, i.e. is applicable to many different contexts regardless of the specific content and purpose of a given SLA. The mechanisms – detailed in Section V – for provisioning of the BAR services are a good example of a resource-specific approach. We believe the proposed model to be applicable to any general-purpose network inter-domain scenario.

IV. Agreement Signalling

We propose an approach to SLA signalling based on the instrumentation of the WMS and on the availability of a new collective service called the *Agreement Service*.

The WMS [12], [13] comprises a set of middleware components responsible for the distribution and management of tasks across Grid resources and services. In this context the term *resource* refers to fabric-layer facilities, for example computing farms, data storage systems and networking infrastructures. Conversely, we use the term *service* to refer to a general service instance regardless of the provided functionality level with respect to the Grid service architecture.

Scheduling policies define the internal strategy which aims at maximizing the individual task execution efficiency, and, at the same time, at optimizing resource and service usage. WMS greatly simplifies the usage of Grids by handling execution errors, by supporting the discovery of the services satisfying the user and application requirements, by allowing re-submission in case of failure and by offering logging and bookkeeping facilities which can increase the internal robustness of the WMS.

The WMS instrumentation is needed in order to provide the capability of handling the allocation and reservation of Grid resources, network bandwidth being a specific example. The WMS supports SLA signalling, discovery of service instances satisfying the user's requirements, and monitoring.

We borrow the notations of *Agreement Initiator*, *Agreement Offer* and *Service Provider* from the WS-Agreement proposed specification of the GRAAP-WG of the GGF [3]. In particular, as we focus on reservation and allocation, in this context the provider is called *Reservation and Allocation Service Provider* (RASP). The proposed architecture is illustrated in Figure 1, which details the interactions among the above components. Three types of resource are shown: the *Network Service Access Point* (NSAP) – presented in Section V – the Computing Element (CE) and the Storage Element (SE). In what follows we focus our attention on the NSAP.

The agreement initiator is the entity which triggers the SLA signalling process. In our scenario, users, applications or Grid middleware components can be initiators provided that the identity of the resources to be involved in the SLA signalling is known. If available, we require such information to be conveyed by the SLA request, in order to allow the allocation and reservation activity to be submitted directly to the AS. If users have no a-priori knowledge of the Grid infrastructure, we propose an approach where the allocation and reservation task is submitted through the WMS in order to take advantage of its discovery capabilities. After the submission, the WMS determines both the RASPs meeting the task requirements, and the ASs which are capable of supporting the corresponding RASP interfaces. In this scenario, the WMS has the agreement initiator role and acts on behalf of the user or application.

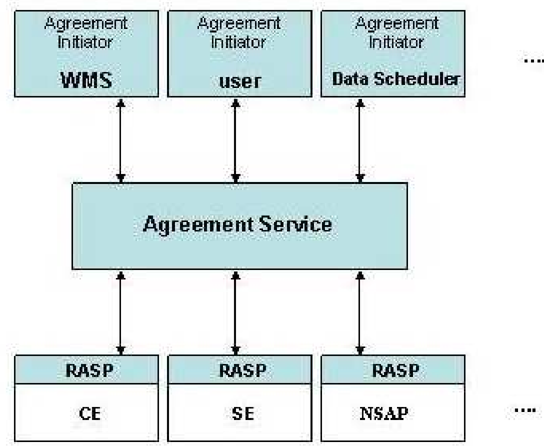

Fig. 1. Agreement Service and the other components of the proposed architecture

The AS can receive agreement offers from various types of initiators, such as the user, the WMS, or other Grid middleware components (e.g. the Data Manager, which is responsible for moving data across Grid space storage devices). The AS initially checks SLA compliance to the appropriate corresponding template, it extracts attributes from the SLA request, and, finally, it dispatches them to a number of RASPs by invoking the appropriate interface.

RASPs are fabric-layer resource-specific services that are responsible for the management of the availability of the local resource over time, for admission control, for the control of the satisfiability of each authorized incoming requests and, finally, for the enforcement of the guarantees. We propose each reservation-capable resource to be associated to an individual RASP instance. A computing farm, a data storage system or a network domain are examples of resources supervised by a RASP.

A. Agreement Initiator

The initiator is responsible for triggering the negotiation of a new agreement by interacting with the AS, as defined in [3]. Initiators can be the reservation consumer itself or a proxy. The former case is applicable when the consumer has a-priori knowledge of the resources to be reserved. Alternatively, the consumer can decide to submit the request to the WMS, which performs discovery.

We request the agreement initiator to provide information to the AS about the candidate RASPs, as this approach gives the opportunity to use existing WMS capabilities and it reduces the overall AS complexity.

Requests originate from the user or application and are submitted to the WMS. We propose them to be modelled following the agreement offer structure defined in [3]. Resource requirements and preferences are part of the SLS and they consequently need to be specified as a list of *Service Description* and *Guarantee Terms*.

We identify four term categories:

1) *General* attributes: they describe the type of task requested and the corresponding functionality.
2) *Time* attributes: they provide information about the start time, end time and duration of the reservation.
3) *Functionality* attributes: they supply reservation and resource-specific information.
4) *Guarantee* attributes: they define the Quality of Service profile.

B. Agreement Service

The AS is the novel service component that provides SLA management capabilities to the WMS. We propose the AS to be responsible of interacting with the agreement initiator (from the server side) and with the RASPs (from the client side) which can satisfy the initiator's request. In particular, the AS accepts agreement offers from initiators and is responsible for compliance checking as well as for the interaction with one or more RASPs in case of success.

The AS exposes RASP capabilities through agreement templates [3]. The template is the SLA skeleton and includes a list of creation constraints applicable to the SLS. The templates advertised by one AS explicitly define the types of service the agreement service can handle, and consequently depend on the service providers the AS interacts with. Consequently, different templates are applicable to SLAs for different network services. The AS we propose can offer different templates for the same type of service, in order to offer various types of agreement to different user groups according to their Virtual Organization (VO) membership. Membership, role and capability information is provided by the VO Membership Services (VOMS) [14].

The AS can perform negotiation in order to assist initiators during the establishment of a SLS whose characteristics depend on the current status of the RASP. For every well-formed offer, SLS attributes are translated to low-level RASP-specific terms according to the RASP interface.

The operations exposed by the AS are: `createAgreement`, `agreementKill` and `agreementUpdate` for agreement management (creation, deletion and update respectively), `getAgreementStatus` and `getAgreementProperties` for monitoring purposes and `getProperties` for retrieval of the templates handled by a given AS.

Finally, the AS supplies information about the status of agreements under negotiation and the attributes of the agreements successfully established. We proposed these information elements to be managed by the Logging and Bookkeeping component of the WMS.

C. Reservation and Allocation Service Provider

The RASP interface determines the set of agreement services that can invoke the service provider and this type of information is used by the WMS in order to select a list of candidate AS's. The service provider interface can be of any type, and it is resource-specific. We propose a NSAP interface for the GDFT and VLL services in Section V-B.

V. Network Service Provisioning

We propose an instantiation of the architecture described above specific to the provisioning of BAR services. The entities involved in this scenario are the *High Layer Middleware* – at the agreement initiator layer; the *Bandwidth Allocation and Reservation service* (BAR) – loosely mapping[1] to the agreement layer; and the *Network Service Access Points* (NSAP) – at the service provider layer. The hierarchy is illustrated in Figure 2 while the layer mapping is shown in Figure 3.

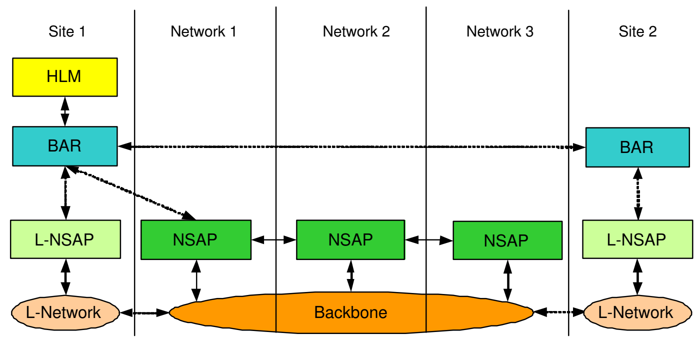

Fig. 2. BAR Architecture outline

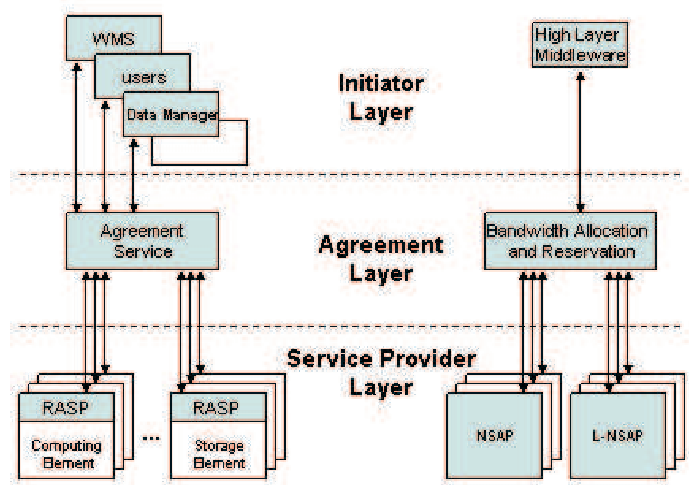

Fig. 3. Mapping of services to layers

In what follows we describe the network-specific functionality of the individual components.

The Bandwidth Allocation and Reservation Service: can be invoked by High Layer Middleware and/or user applications. Reservation requests are expressed in a network-neutral language. The role of the BAR is to pass the request on to

[1]The BAR service currently implements a subset of the agreement layer functionalities described in [3].

three services: the local L-NSAP, the designated backbone NSAP and also to the remote L-NSAP (via the remote BAR) at the target destination. When interacting with NSAPs and L-NSAPs, BAR translates a high-level SLA request into a network-oriented one. This is one of the general roles of the AS, as explained in Section IV-B.

The Network Service Access Point: is a network entity concerned with admission control and configuration of network equipment on the backbone. We assume NSAPs to be present in transport networks, such as GÉANT and NRENs, thus forming an extended QoS-enabled network[2]. Its functionality includes issuing a response message to the BAR entity making the request, which contains information about the success of the operation. As explained in [8], the NSAP abstracts the network-specific services but still speaks a language that is network-oriented. A NSAP reference implementation is under development at the GÉANT2 project [15].

Local Network Service Access Point: for an end-to-end reservation, local networks outside the extended QoS-enabled network also need to be configured. We call the entities effecting this "last-mile" configuration *Local NSAPs* (L-NSAP). The L-NSAPs do not partake to the chaining activity of the NSAPs, as discussed in Section V-A. They only interact with their local BAR services, whether configuring the source or destination local network.

A. Interactions

The various BAR interactions are depicted in Figure 4. When invoked by a client for example HLM (step 1), the BAR service translates the application-oriented information conveyed in the request to a corresponding network-oriented request, structured according to the type of service specified. The BAR delegates it to the nearest designated NSAP at the edge of the backbone (step 2). We are currently reviewing options for the discovery of NSAPs.

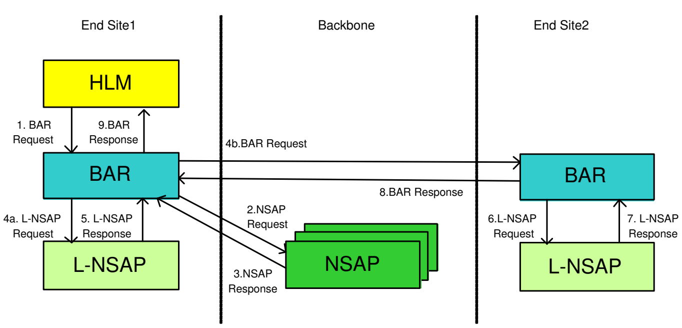

Fig. 4. Sequence of Interactions for BAR

Figure 5 is an elaboration of what happens at the network between steps 2 and 3 of Figure 4. In order to proceed with the enforcement of the guarantees, admission control is performed. In order to achieve this, the NSAPs need to interact with each other (steps 2a-2d). In particular, the originating NSAP coordinates with its peers along the path to the destination, and a collective decision is made as to whether the reservation is in principle approved.

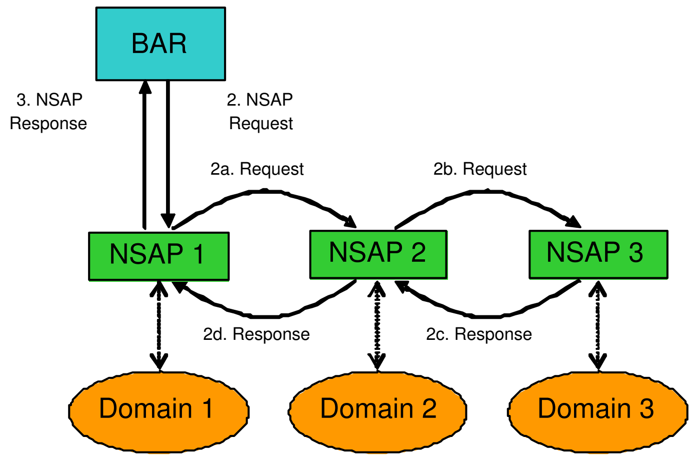

Fig. 5. Sequence of chaining between NSAPs for BAR

The designated NSAP conveys the output of the coordination to the invoking BAR (step 3). At this point, the BAR service responds to the client if the backbone is not able to honour the SLA requested. Otherwise, the BAR delegates it both to the L-NSAP at its end-site (step 4a) and to the peer BAR service at the target end-site (step 4b). The L-NSAP responds to its BAR service if it can in principle effect the reservation (step 5). Finally, the remote BAR service delegates the request to the remote L-NSAP (step 6) and returns the response to the originating BAR service (steps 7 and 8), which responds to the client (e.g., HLM) whether the reservation is successful (step 9).

The request is booked if and only if all three responses above are positive. If the local networks reject the reservation, BAR needs to trigger a compensation task and cancel the reservation at the backbone. We require the reservation response to be immediate, regardless of the actual lead-time needed for service provisioning.

B. Interfaces

Specification of input and output information is important to guarantee interoperability [8]. We propose two interfaces: the HLM-BAR and BAR-NSAP interfaces, which conceptually structure the interaction between the agreement initiator and agreement provider layer, and between the agreement and service provider layer respectively.

The HLM-BAR Interface: is the interface exposed by the BAR to the HLM, as shown in Figure 6, and supports operations to effect both the GDFT and VLL reservations; these are equivalent to the `createAgreement` operation detailed in Section IV-B. Different input information is provided depending on the specific bandwidth reservation service type requested. Other operations exposed are *Query Reservation* and *Cancel Reservation* for reservation management. These are respectively equivalent to the `getAgreementStatus` and

[2]The term backbone in the rest of this paper is assumed to mean extended QoS-enabled network.

`agreementKill` operations detailed in Section IV-B. The proposed framework is extensible, as extra operations can be added to support additional service types and new types can be supported by adding the service-specific operations required.

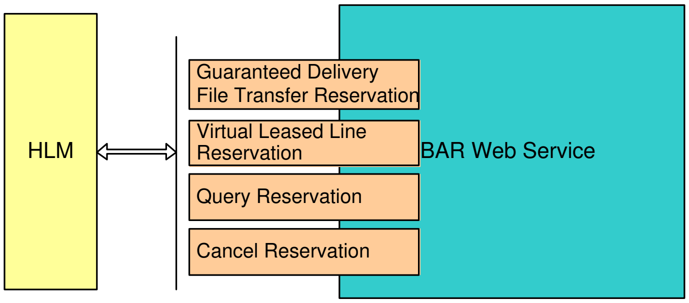

Fig. 6. HLM-BAR interface

The BAR-NSAP Interface: is the interface needed to pass on the SLS attributes to the appropriate NSAP. The operation *Request Network Reservation* requests a service by specifying values which depend on the network technology supporting the NSAP, for example Premium IP. The translation from high-level attributes characterizing the BAR services to low-layer parameters is performed by BAR. *Query Network Reservation* and *Cancel Network Reservation* are the other two operation exposed, as shown in Figure 7. They are needed for the network reservation management. Figure 7 shows the interface exposed by the NSAP web service [8].

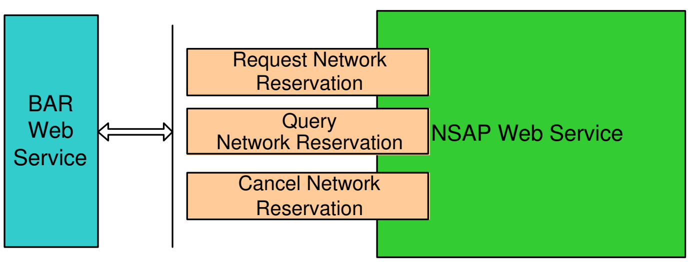

Fig. 7. BAR-NSAP interface

C. Security

Integrity, confidentiality, authorization and authentication are required at any level of the proposed architecture: between the SLA initiators and the BAR service, between peer BAR service instances (during the set-up of a bandwidth reservation at each end-site) and between individual NSAP services.

VI. Prototype

A prototype of the Agreement Service has been implemented to demonstrate the feasibility of the proposed architecture. As described in Section IV-B, the operations supported by the AS are resource-independent, and hence the AS is applicable to resources other than network bandwidth. The resource reservation capabilities are entirely dependent on the type of templates supported by the AS.

Two components have been developed: the WMS Resource Manager and a storage space AS, which is the client of a SE exposing a SRM 2.1 interface [16], and supports the signalling of agreements for storage space reservation. The realization of the web-service part of the AS was based on GSOAP v.2.7.0. The implemented AS (client and server side) are entirely based on the WS-Agreement XML schema definitions. The operations exposed by the AS are similar to the Agreement Factory ones defined in [3], however the prototype WSDL exposed by the AS does not fully comply with the standard, as it is not currently based on the WS-Resource Framework [17]. The AS prototype implementation will be revised to fully comply with the new WS-Agreement specification as soon as it is released.

The corresponding RASP offering space reservation enforcement to the prototype AS, is based on the StoRM [18] Service Provider exposing a SRM 2.1 interface which supports immediate space reservation, file copy and file creation. StoRM enforces space reservation by allocating space directly into the file system. This is accomplished by relying on the GPFS file system [19] pre-allocation capabilities.

A functionally limited prototype of the BAR service is also available. The prototype implements the BAR Request Service function for the GDFT service type and Cancel Service functions mentioned in V-B. It demonstrates the interactions between a BAR client, a BAR web service and a NSAP web service. It investigates how the interfaces defined can be used to perform a GDFT reservation request and a Cancel request. It also includes a basic security mechanism between a client and a BAR web service.

VII. Related work

Several approaches to resource reservation and allocation for Grids have been proposed, among which Cremona [20], the General-purpose Architecture for Reservation and Allocation (GARA) [21], [22], the User Controlled Lightpath Provisioning (UCLP) System [23] and the DataTAG framework for the advance reservation of heterogeneous network paths [24].

Cremona (Creation and Monitoring of Agreements) is an IBM initiative related to the Web Service Level Agreement (WSLA) project, which proposes an architecture for middleware implementing WS-Agreement and provides managed functionality for agreement templates and instances. Both Cremona and the architecture we propose in this paper, are based on WS-Agreement. However, while Cremona is intentionally designed in order to apply to general service environments, the approach we propose aims at defining a framework for SLA management which relies of existing general-purpose functionalities supported by Grid resource management and scheduling services, such as resource and service discovery, handling of communication and scheduling errors, resubmission, logging and bookkeeping.

GARA addresses the issue of advance reservation by proposing a common framework that can be applied to different resource types such as computing, storage and networking.

Resource management is integrated with a variety of general-purpose Grid services (e.g., for authentication and authorization).

UCLP focuses on the problem of dynamic provisioning of network paths called Lightpath. The Lightpath is a concatenation of individual links: the Root Lightpath Objects. The Lightpath is an end-to-end bandwidth-guaranteed tunnel that provides connectivity between two physical resources (e.g, Optical Cross-Connects).

DataTAG approach Paths are advertised by the Grid Information Service following the same approach adopted for any other Grid resource type. Only static attributes are exposed, while status information is only accessible to the corresponding resource manager for greater scalability.

GARA, UCLP and the DataTAG framework are prior to the WS-Agreement standardization effort and propose SLA signalling solutions only suitable to specific deployment scenarios.

For a comprehensive survey of network resource management please see [25]. While most of the systems only support a single or few network technologies, the approach we proposed is control-plane independent and exposes a common set of operations which are independent of the technology used in specific network domains.

In addition, some of these systems assume that admission configuration and control at only ingress and egress is required, the intervening network being *overprovisioned*. This is not generally true, as the network domain peering often constraints the applicability premium services. In this paper we have propose a minimal set standard operations to ensure interoperability between different network domains in a inter-domain scenario.

VIII. Conclusion and Future Work

We propose an approach for resource-independent Grid agreement signalling, and we show its integration with an existing middleware platform.

The general-purpose SLA signalling solution described in this paper is based on a novel component dubbed AS, which allows agreement management by relying on a number of WMS capabilities, such as service discovery, logging and bookkeeping and error handling through re-submission. The AS is complemented by the RASPs, which enforce the requested guarantees, as defined in the WS-Agreement architecture of the GGF.

We then focus on the problem of network QoS and introduce two network services for Grid file transfer based on bandwidth-management techniques: the GDFT and the VLL. We detail the functionality of the BAR Service that addresses the specific requirements of network reservation management. It relies on a general service inter-domain provisioning model based on the peering of various distributed NSAPs.

We plan to complete the instrumentation of the WMS in order to integrate it with a number of AS prototypes, which will be refined to improve error handling and to support authorization, authentication and agreement monitoring. In addition, the AS implementation will be adapted to fully comply with the WS-Agreement specification when available.

With regards to BAR, we intend to add the currently missing functionality (Request Service function for VLL Service Type and Query Service) and we plan to experiment with options available for the last-mile problem. We also aim to integrate the GÉANT NSAP implementation when available.

Acknowledgements

Antonia Ghiselli, Francesco Giacomini, Erwin Laure and Francesco Prelz provided valuable comments on the agreement architecture illustrated in this document. Fabrizio Pacini, Luigi Zangrando, Riaccardo Zappi, Luca Magnoni and many other colleagues from the EGEE JRA1 Workload Management team, assisted us during the prototype development phase. We are grateful to Maarten Buchli, Alistair Phipps, Florian Scharinger and Gloria Vuagnin who have worked on the design, development and deployment of the BAR service.

This work is supported by the EU project EGEE, sponsored by the European Union under contract number INFSO 508833; by the INFN project INFN-GRID; and by the UK Joint Information Systems Committee (JISC).

References

[1] *Institution of SLAs and appropriate policies*; EGEE Deliverable DSA2.2, May 2005 (`http://edms.cern.ch/document/565447`).

[2] *The Grid Resource Allocation and Agreement Protocol Working Group*; Global Grid Forum (`https://forge.gridforum.org/projects/graap-wg`).

[3] Andrieux, A.; Czajkowski, Dan, A. et al.; *Web Services Agreement Specification (WS-Agreement)*; GGF Grid Resource Allocation Agreement Protocol draft, work in progress (`https://forge.gridforum.org/projects/graap-wg`).

[4] *The EGEE project home page*; (`http://www.eu-egee.org`).

[5] EGEE MJRA4.2 *Requirements and Uses cases from network applications*, EGEE Milestone MJA4.2, Jan 2005 (`https://edms.cern.ch/document/476742/1`).

[6] Ferrari, T. et al.; *Grid Network Services Use Cases*; Grid High Performance Networking Working Group, GGF, work in progress, Jun 2005.

[7] *International VLBI Service for Geodesy and Astrometry* (`http://ivscc.gsfc.nasa.gov/`).

[8] *Specification of interfaces for Bandwidth Reservation*; EGEE Project Deliverable DJRA4.1, May 2005 (`https://edms.cern.ch/document/501154`).

[9] Blake, S.; Black, D.; Carlson, M et al.; *An Architecture for Differentiated Services*, RFC 2475, Dec 1998.

[10] Nichols, K.; Carpenter, B.; *Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification*, RFC 3086, Apr 2001.

[11] Campanella, M.; Ferrari, T.; Leinen, S. et al.; *Specification and Implementation Plan for a Premium IP Service*; GÉANT project IST-2000-26417, Apr 2001.

[12] Andreetto, P.; Borgia, S.; Dorigo, A. et al.; *Practical Approaches to Grid Workload and Resource Management in the EGEE Project*; in Proc. of the Conference on Computing in High Energy and Nuclear Physics (CHE0P 2004), Vol. 2, pp. 899-902, Interlaken (CH), Sep 2004.

[13] *EGEE Middleware Architecture*; EGEE Deliverable JRA1.1, Aug 2004 (`https://edms.cern.ch/document/476451`).

[14] Alfieri, R.; Cecchini, R.; Ciaschini, V. et al.; *From gridmap-file to VOMS: managing authorization in a Grid environment*; Future Generation Computer Systems Journal, Vol 21 (2005), pp. 549-558.

[15] *The GÉANT2 Project Web Site* (`http://www.geant2.net/`).

[16] *The Storage Resource Manager Interface Specification*; Grid Storage Management Working Group, GGF (`http://sdm.lbl.gov/gsm/`).

[17] *The WS-Resource Framework* (`http://www.globus.org/wsrf/`).

[18] Zappi, R.; Magnoni, L.; Donno, F; Ghiselli, A.; *StoRM: grid middleware for disk resource management*; in Proc. of the Conference on Computing in High Energy and Nuclear Physics (CHE0P 2004), Interlaken, CH, Sep 2004, Vol. 2, pp. 1238-1242.

[19] *General Parallel File System* (http://www-1.ibm.com/servers/eserver/clusters/software/gpfs.html).

[20] Ludwig, H.; Dan, A.; Kearney, B.; *Cremona: An Architecture and Library for Creation and Monitoring of WS-Agreements*; in Proc. of the 2nd International Conference on Service Oriented Computing (ICSOC 2004), pp. 65-74, ACM Press, New York, 2004.

[21] Roy, A.; Sander, V.; *GARA: A Uniform Quality of Service Architecture*; published in Nabrzyski, J.; Schopf, J. M.; Weglarz, J. (Eds.), Grid Resource Management: State of the Art and Future Trends, Kluwer, 2003, pp. 377-394.

[22] Foster, I.; Roy, A.; Sander, V.; *A Quality of Service Architecture that Combines Resource Reservation and Application Adaptation*; in Proc. of the 8th International Workshop on Quality of Service; Pittsburgh, PA, Jun 2000.

[23] Wu, J.; Campbell, S.; Savoie, J., M. et al.; *User-managed end-to-end lightpath provisioning over CA*net 4*, in Proc. of the National Fiber Optic Engineers Conference (NFOEC), Orlando, FL, USA, Sep 2003, pp. 275-282.

[24] Curti, C.; Ferrari, T.; Gommans, L. et al.; *On Advance Reservation of Heterogeneous Network Paths*; Future Generation Computer Systems Journal, Vol. 21 (2005), pp. 525-538.

[25] *Network Resource Reservation Software and Interfaces, A State of the Art Survey*, EGEE Project Deliverable DJRA4.1 – Addendum, Nov 2004 (https://edms.cern.ch/file/501154/1/EGEE-DJRA41-Addendum-501154-v1.1.doc),