Common Service Definition

Version 0.6
March 6, 2005

Jerry Sobieski
Mid-Atlantic Crossroads

Tom Lehman
University of Southern California
Information Sciences Institute East

1 Introduction

One of the basic obstacles to provisioning end-to-end "light path" services on a global basis is the lack of consistent and measurable performance definitions for these light paths.

Put another way, there is no mechanism to insure that the transport services instantiated by assembling resources through multiple administrative domains and across different network technologies will actually create an end to end path that is predictable, repeatable, and consistent with the user's request. Fundamentally, there is no clear "service definition" adopted by the network service providers that clearly describes the capabilities and performance characteristics of their services.

This creates a situation whereby the "service" offered or presented at one end of the light path may not be completely satisfied by the network technologies and capabilities across the intermediate networks. These service discrepancies may not be obvious or intentional, or the providers may honestly believe that some minor service differences will not be noticeable or are not significant. Indeed, it may be that the service provided actually exceeds the service requested by the user. But this service variability makes the service non-deterministic, unpredictable, and unrepeatable, and often causes significant performance impairments for the end user/application. There is no objective means of measuring and verifying the service.

In order to enable truly deterministic and consistent network service end-to-end, the users requesting those services and all the networks providing those services must adopt a **common service definition.** By doing so, the networks providing the services have a standard by which the service can be engineered and constructed, and operationally concatenated and tangibly measured and validated. The users will also have a means by which they can request services that are predictable, deterministic, and again measurable and verifiable.

This paper attempts to define the following:

i) **Common Services Paradigm**. Define the philosophy, need, and paradigm of a "Common Service Definition" and explain how it can form the basis for global network service provisioning

ii) **Common Service Definitions**. Define an example common network service. This example service is presented in an appendix and the primary intent is to provide a starting point for community discussion and iteration of specific network service definitions. It is hoped that this can provide a place for multiple common network services to be defined based on various transport services such as Ethernet, SONET/SDH, Fiber Channel, Infinband, and others.

In general, this paper is addressing the concept of a "service" for the purpose of furthering dynamic allocation of end-to-end data paths with dedicated network resources. In particular, we are looking at "light path" services - services that bear a strong resemblance to traditional connection oriented services or circuits, but which incorporate more recent developments and aspirations for optical and all-photonic networks.

The concept of a "service" is as viewed from an "end-system" which is connected to (and making a request from) a network. In this context, "end-system" is a very general term and may mean a host, computational cluster, data storage cluster, radio telescope, router, or any device which is connected to the network.

As noted, this "service" may traverse multiple administrative domains and network technologies. It is not the intent of this paper to address inter-domain routing, peering, or signaling issues. Likewise, the issue of how a single end-to-end path constructed from multiple heterogeneous network technologies might satisfy a specific service request is not addressed. The service described in this paper is as would be defined, measured, and verified at the end points of the provisioned path. Discussions regarding how to provide these services across multi-domain networks and heterogeneous network technologies are anticipated to be topics for future community discussions. It is expected that the common service definitions will establish an initial set of driving requirements for these other issues.

The remainder of this paper is organized as follows: Section 2 provides a rationale for the need for a Common Network Service Definition. Section 3 provides a framework for a Common Network Service. Section 4 discusses issues with service verification. Section 5 provides an overview of other related issues. Finally, Appendix A provides an example for a basic Ethernet service definition.

2 The Need for a Common Service Definition

Networks provide "services". These may be for example "best effort IP" service, or "sonet" services, or "transparent LAN service", or even "domain name service", etc. These services can be defined by a set of very specific service characteristics (or parameters) and allowable values for those parameters.

For instance, an "ethernet" service may have as a parameter the "Protocol Bandwidth" with a defined value of 1 Gigabit/second (Gbs).

Another less obvious "ethernet" parameter might be "Maximum Transmission Unit (MTU Size)". Unless specifically defined, there is no basis to assume a specific size. A knowledgeable engineer might reasonably assume the MTU cannot exceed 1500 bytes, but this is only an assumption and without an explicit specification, there is no guaranty that the maximum MTU might be larger (or smaller.).

Lets look at an example of two users with 10gigabit LAN PHY ethernet interfaces. They may request that a "10 gig Ethernet" connection between them be established so that they can transfer important files quickly over the link. The first time the users request this light path be established, the end to end link may be provisioned over intermediate 10 gigabit ethernet LANPHY links. The users are able to successfully transfer their files at 10 Gbs. However, the next time the request is made, one of the intermediate networks provision the path across a 9.6 Gbs WANPHY ethernet link. These users are no longer able to get a full 10 Gbs of thruput. In fact, the users may experience a significant amount of packet loss (400 Mbs!) due to the capacity mismatch between the WANPHY and LANPHY media. Officially, the network was carrying ethernet frames, and offered what is generally considered as 10 Gbs gigabit of capacity, but the service experienced by the application was markedly different from one instance to the next. Here is an example of where a provider may naively believe that 4% service variance is unimportant. This would be incorrect since this can impact higher layer protocols. In reality, the users might be perfectly happy with a 9.6 Gbs path if only they knew in advance that this was the maximum sustainable capacity.

This is an example of the mis-match that can occur between the user's request (and expectations) and the network's delivered service parameters. Lets look at another potential aspect with this example.

Consider the same scenario, but this time the intermediate networks tried to load balance the 10Gbs path across some bonded 9.6Gbs links along the path. The full 10 Gbs capacity would be available to the end users, but the packets are reordered by merit of the load balancing process. Does the service meet the user's expectations? Does is meet the providers commitments? What exactly was promised by the provider? If the there was no specific sequential commitment in the service offering, then the users really had no right to expect that all packets would be delivered in order and the service would meet the definition. The point is that the service definition is required to clearly and formally define all aspects of the services – and by exclusion, all that is not part of the service.

Similar situations may occur between the intermediate transit domains as well. One domain may define a “gigabit Ethernet” service that is provisioned across a WDM system point to point that is transparent, i.e. will carry VLAN tags and Spanning Tree Protocol and all the 8B/10B encoding, flow control, etc; a neighboring domain may define a “gigabit Ethernet” service that is switched across a number of ethernet switches but carries untagged frames only and does not allow Spanning Tree. For many users, such service variance may go unnoticed, but unless these characteristics are clearly defined – or explicitly undefined, there will be confusion as these services are provisioned across multiple domains, over heterogeneous media, on a global end-to-end basis.

Service variance is not always detrimental to the applications. An application may actually get better performance than it had requested. In these cases, the user application must be aware that service variance experienced on one instantiation may not be provided on a subsequent instantiation, or indeed, it might actually vary within the existing instance. Even such positive service variance makes the service non-deterministic.

A good example of a non-deterministic service is Best Effort IP. While this paper is trying to establish a means for clearly defining services that can be predicted and consistent, it should **not** be construed that non-deterministic services are useless or detrimental. Best effort IP has more than proven itself as a very effective service for the overwhelming majority of current networked applications. It is left as an exercise for the reader to formally define a “Best Effort IP” service.

In order to ensure that the emerging optical networks can provision end-to-end “light path” services in a consistent, deterministic, and predictable manner, the network community needs to define a **common services definition** which can be utilized as the based for user requests and service verification. The following section discusses a framework for this service definition.

3 Common Network Service Definition Framework

At a very fundamental level, the "Service Definition" is a set of very specific and measurable service parameters and their allowable values. These values are defined by agreement between the interested parties – i.e. the service providers and/or the users. Only those service parameters/characteristics that are explicit can be relied upon to behave in any predictable manner. These services could be unique services offered by a single particular provider; or alternatively, several providers could collaborate and define a set of services that they all agree to provide - and having done so, they can peer these services and concatenate them for end-to-end delivery.

So how do we go about creating a formal and common *service definition*?

To start, the initial service concept should be as simple as possible. A well understood service should be selected, say for example, a "basic ethernet" service, and the essential qualities of that service should be identified and defined. This process could be a joint effort by multiple interested parties thus insuring that the resulting service definition could be broadly and easily adopted as the basis of service offerings among many interconnected networks. This process will quickly result in a "reasonably good" service definition – not necessarily the ideal or complete, but suitable to allow multi-domain services to at least talk apples-to-apples.

Lets explore the basic ethernet example:

We can informally define a basic ethernet service as having the following arbitrary characteristics:

- It is a point to point service
- It carries standard ethernet frames and delivers them in the same order received
- IT supports frames up to 9000 bytes in size
- It provides up to 1000megabits of bandwidth capacity

Using a simple syntax to formalize the definition, this basic service could be defined as:

```
Basic_EtherSegment := {
      Service_Type         := Point_to_Point;
      Framing              := IEEE_802.3;
      Protocol_Bandwidth   := 1000 Mbs;
      Maximum_MTU_Size     := 9000 Bytes;
      }
```

In this example, the parties defining the Basic_EtherSegment" service can decide what a Point_to_Point service type means – perhaps the idea in this case was to create a gigabit Ethernet link that felt like a simple back to back patch cable between two servers the implication being that no other traffic would be seen or felt on the link.

Once a service is deployed, experience will likely expose aspects of the service that need further definition.

One could imagine the Basic_EtherSegment example above needing refinement to address the needs of jitter sensitive applications. The parties reconvene and create a new service definition :

```
Basic_EtherSegment := {
      Version              := 1.1;
      Service_Type         := Point_to_Point;
      Framing              := IEEE_802.3;
      Protocol_Bandwidth   := 1000 Mbs;
      Maximum_MTU_Size     := 9180 Bytes;
      Maximum_Jitter       := .0001 seconds;
      }
```

The service definition can be amended to include additional service parameters or refinements to existing parameters, or an entirely new service definition can be created reflecting the differentiated service parameters. The overall process can be duplicated to define other services as well such as SONET/SDH, dim fiber service, Fiber Channel, and others. It is also anticipated that users will also be interested in other service parameters which will be harder to provide than the initial set envisioned. This may include parameters such as packet latency, jitter, loss, and reordering.

How a formal service definition is "published" is an open topic. The authors of this white paper believe it should be considered in a broader scope. The form of the service definition could be something as simple as a shared text document, or something more structure such as the pseudo code used in these examples. In eiterh case, the service definition will play a role in how the services are actually delivered. And these services may be many and varied. The definitions will likely lend themselves well to emerging standards such as Extensible Markup Language (XML). XML has potential benefits far beyond simply providing a formal syntax for describing a service. XML could conceivably be used in broader interactions such as user network requests, distribution of inter-domain transport service capabilities, describing applications requirements, grid services interactions, and many others. So while the recommendation is to start simply, with easily understood initial service descriptions such as the included examples, the longer term requirements may dictate a more powerful approach.

Appendix A presents an initial proposal for a Basic Ethernet Service definition. The intent is that once the community agrees on a general format for such service definitions, then additional services can be defined in a similar manner.

4 Service Verification

"Service variance" refers to the difference in service characteristics found by comparing the service definition with the service characteristics actually measured on the path. If a service parameter is specified, it should be measurable and verifiable at the user interface.

If a particular service characteristic is not explicitly specified as part of a formal definition, then there is nothing that can be said of that particular characteristic – it may be present, or it may not; it might be present for one request, and not for a subsequent request.

In general, service variance should be non-existent on any given end-to-end service. If service variance is zero, then the user is guaranteed to have a consistent set of network resources and performance characteristics at their disposal on an end-to-end path, every time the path is requested. Service verification can be accomplished by measuring network service characteristics at the user connection points and comparing to service request.

5 Related Issues (Routing, Signaling, AAA, and Scheduling)

The service definition described in this paper is one component required to develop the capability to provision end-to-end "light path" services on a global basis. Other important capabilities include mechanisms to actually provision the services (routing and signaling), scheduling, and AAA (Authentication, Authorization, and Accounting).

It is recognized that many discussions and agreements with respect to protocols, interfaces, and message formats will be required in order to realize these features across multiple networks. While this paper does not address these issues directly, it envisioned that these common services agreements will be a driving requirement for the other advanced capabilities and features. In this manner, the common service definitions are intended to be a building block service. That is, definition of the features, parameters, and requirements for the basic transport service is a good starting point for the requirements definition for routing, signaling and other mechanisms necessary to instantiate a requested service.

In general, for the purposes of this discussion, a separation is maintained between the data (or transport) plane and the control plane. The choice and implementation of a control plane should not be a driver for the nature and performance of the data plane. For this reason, detailed discussion of these issues is recommended for a time after an initial framework for the common services is defined.

6 Summary

In closing, the Common Service Definition Paradigm has the following features:

It is simple. It allows the process of defining a service to start very simply. A single domain can unilaterally define an initial version of a newly conceived service with as little as a single characteristic.

It is extensible. Additional service characteristics can be defined – either as a global service characteristic whose meaning is standardized, or a private service characteristic whose meaning is only important to a single domain (or set of domains.)

It is iterative. As a service offering is deployed, and networks and applications gain experience with its provisioning and utilization, additions or modifications can be easily incorporated into newer versions of the service definition.

It is open and verifiable. The service characteristics are public and well defined and can be measured objectively at the service access points.

It can support a hierarchy of service characteristics. As in the example above, hierarchical inheritance can be asserted from basic service characteristics into more tailored service offerings for a more targeted user community or market segment.

As the research community evolves towards multi-service infrastructure, and more advanced applications require a broader portfolio of network services, the ability to formally define network capabilities and to dynamically integrate them into the applications environment will require an approach such as was described in this paper. We hope this paper forms a basis for discussion and provides at least an initial means of addressing the issue and moving forward.

References

2nd Int'l Workshop, Optical Control Planes For the Grid Community, MCNC Workshop Organizer, Pittsburgh PA, November 12, 2004, http://www.mcnc.org/mcncopticalworkshop/nov04

Workshop, Optical Control Planes For the Grid Community, MCNC Workshop Organizer, Chicago, ILL, April 23-24, 2004, http://www.mcnc.org/mcncopticalworkshop

Workshop, Optical Network Testbeds, NASA Ames Organizer, Moffet Field, CA, August 9-11, 2004, http://www.nren.nasa.gov/workshop7

4th Annual Global LambdaGrid Workshop, Nottingham, United Kingdom, September 2-3, 2004, http://www.glif.is/meetings/04.html

Control Plane Conference Calls, DRAGON, CHETTAH, OMNInet, UltraScience Net, CY 2004-2005, hosted from USC/ISI East, Arlington VA.

DRAGON, University of Maryland/Mid-Atlantic Crossroads, University of Southern California/Information Sciences Institute, George Mason University, National Science Foundation (NSF), http://dragon.maxgigapop.net

CHETTAH, University of Virginia, http://cheetah.cs.virginia.edu/

OMNInet, Northwestern University, International Center for Advanced Internet Research, http://www.icair.org/omninet

UltraScience Net, Department of Energy, http://www.csm.ornl.gov/ultranet/

CANARIE, Canada's Advanced Internet Development, http://www.canarie.ca

Appendix A

Common Network Service Definitions

DRAFT

A.1 Basic Ethernet Service

General Description:
The Basic Ethernet Service is meant to provide a network service similar to that of an ethernet patch cable between two hosts. Such a link would be point-to-point and would be sequential in nature with no possibility of packet reordering or other packets being inserted along the path. This service definition is intended to cover 1 Gigabit/s, 10 Gigabit LANPHY, and 10 Gigabit WANPHY service offerings. A distinction is made between the connection interface speed and end-to-end bandwidth available.

Service Definition:

```
Basic_Ethernet_Service_Definition :=
{
      Connection_Profile      := Point_to_Point,
                                 Sequential,
                                 Bi-Directional;
      /*This implies there is no other source of traffic that should be
      seen or impact this service path.  Sequential implies no packet
      reordering, and bi-directional means the service assumes both
      directions are equally provisioned.*/

      Protocol_Bandwidth      := 100 Mb/s--->10 Gb/s;
      /*Bandwidth can be specified 100 Mb/s increments.  Must be
      compatible with connection interface type.  Bandwidth available
      from preamble to tail*/

      Protocol_Framing        := 1G | 10GLAN | 10GWAN;
      /*802.3 frames delineated by IFG & preamble bytes.



      Maximum_Transmission_Unit     := 9180 Bytes;
      /*this service allows jumbo frames*/

      Maximum_Latency         := number-in-milliseconds;
      /*this parameter defines maximum latency which will be observed
      based on endpoints defined for a specific service request.*/
```

```
Maximum_Jitter          := number-in-microseconds;
/*this parameter defines maximum jitter which will be observed
based on endpoints defined for a specific service request.
Standard packet traffic profile need to be defined as part of
this specification.*/

Maximum_Loss            := lossrate;
/*this parameter defines maximum packet loss which will be
observed based on endpoints defined for a specific service
request.  Standard packet traffic profile need to be defined as
part of this specification.*/

Service_Locations_Type  := IPv4_ADDR | IPv6_ADDR ;
/*Define how endpoints are identified.  Options include IPv4 or
IPv6 addresses.*/

Service_Locations       := IPv4_ADDR+ | IPv6_ADDR+;
/*A list of served end points.  For this point to point service,
only two endpoint will be included.*/

}
```

Note: Several other common network services can be envisioned. Without attempting to define them at this a time, a list is provided below for future consideration.

A.2 Ethernet VLAN Service

A.3 Ethernet LAN Service

A.4 Basic SONET Service

A.5 Basic Fiber Channel Service

B Common Service Definition using XML

The prospect of a global multi-service inter-network is an enthralling concept. However, the ability to create such an environment is not separate from the user community that will employ it. Emerging technologies such as computational grids, application specific network topologies, security and survivability concerns, are all part of the same fabric.

For instance, the manner in which a user describes a distributed application - its nodal resource requirements and the dataflow requirements - must be translated into resource requests that the network can process, AAA keys that can open global resources, and application specific information for actual execution. The network domains that must provide the services must know of each other and their capabilities, there must be shared authentication and authorization architecture, common scheduling and reservation frameworks, just to name a few of the issues. These aspects will be dynamic – new networks and resources coming online, new services being defined, etc. A mechanism to elicit and disseminate this information is required, and this mechanism must be both widely accepted and easily modified to meet a broad set of requirements. The Extensible Markup Language (XML) and its associated tools and web based mechanisms provides one approach that appears to be very promising.

XML is a meta-language that can be used to structure and tag items in documents of many different types. These documents are text based (i.e. human readable), they can be validated in a number of ways; they can be distributed via web based protocols; they can be parsed and fed into applications, and the structures can be easily modified for different purposes. XML interfaces are already finding their ways into network configuration processes.

The following snippet shows how XML might be used to describe the Basic_EtherSegment service example:

```
<?xml version=”1.0”?>
<Service_Definition
xmlns=http://www.maxgigapop.net/research/DRAGON/csd.xml >
        <Service_Name> Basic_EtherSegment</Service_Name>
        <Service_Description>
                The Basic Ether Segment service is meant to provide a point-to-point
                Ethernet connection between two end systems that mimics a back to
                back patch cable between gigabit Ethernet nics on each end system.  This
                performance is measured with respect to data transport, not physical
                layer transparent Ethernet.  The implications here are that the Ethernet
                service provided must be able to transport 1000 mbs of user data,
                bidirectionally, between two points.  The data packets are standard
                Ethernet frames, up to 9180 bytes in size, and may contain any valid
                Ethernet header, and are not subject to any peripheral traffic.
                </Service_Description>
        <Service_Characteristics>
                <Profile   flags=”Point_to_Point   Sequential   Bi-Directional” >
                    <Profile_Desc>
                    This implies there is no other source of traffic that should be seen or
                    impact this service path.  Sequential implies no packet reordering, and
```

```
            bi-directional means the service assumes both directions are equally
            provisioned.
            </Profile_Desc>
        </Profile>
        <Framing>
            <Framing_Type>IEEE802.3</Framing_Type>
            <Framing_Desc>
            802.3 frames delineated by IFG & preamble bytes.  Std Ethernet.
            </Framing_Desc>
        </Framing>
        <Bandwidth>1000 Mbs</Bandwidth>
        <Bandwidth>10000 Mbs</Bandwidth>
        <MTU>
            <MTU_Size>9180</MTU_Size>
            <MTU_Desc>
            Maximum Transmission Unit: number of bytes allowed in the Ethernet
            frame.
            </MTU_Desc>
        </MTU>
    </Service_Characteristics>
</Service_Definition>
```


This is a very cursory example of how a common service definition could be described in XML. This document could be fed into service routing agents and user interface agents to provide a global path selection process for user requests. The basic document description could be extended to include authorization descriptors, or to include additional service descriptions (e.g. G.709 transport service), etc. It is beyond to scope of this paper to define the specific format for service definitions, but the authors felt the prospect of XML based service definition and dissemination poses enough promise to explore the concept it a very basic level.