# Common Service Definitions

Presented to the GLIF Tech working group

Jerry Sobieski
Director, Research Initiatives
Mid-Atlantic Crossroads

San Diego, CA

September 30, 2005

# What are the actual services we are building?

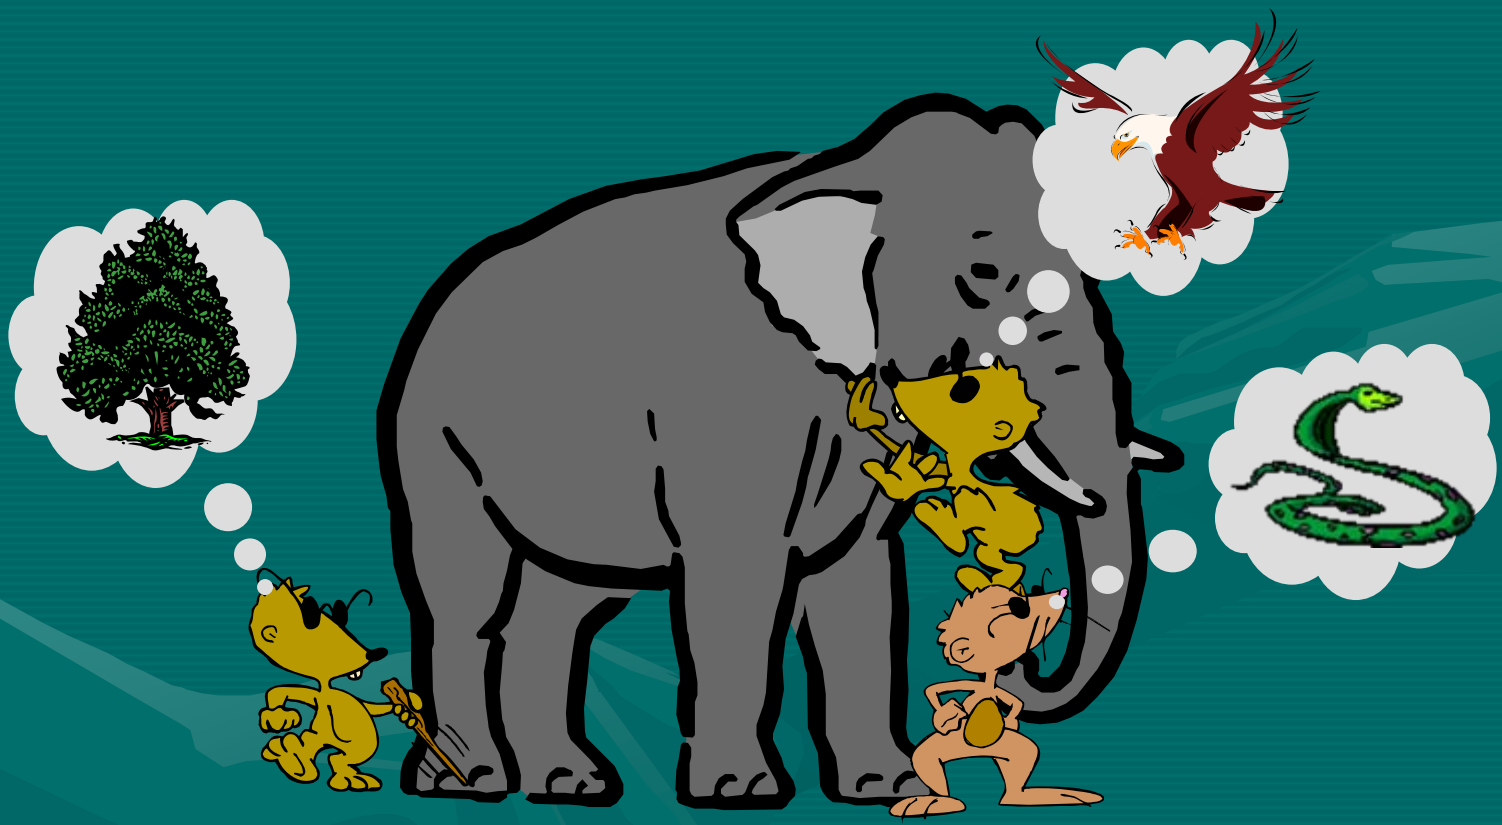

The three blind application developers and the elephant…

# The Problem:

- How do we build a global network of wavelength services that are:
  - Verifiable
  - Predictable
  - Repeatable
  - End-to-end

# Define the Service

- The Common Service Definition specifies and describes the service characteristics experienced by the consumer.
  - It does not specify how the service is engineered or provisioned, it simply specifies what is delivered to the end user. How the service provider decides to construct the service infrastructure is not part of the Common Service Definition.
- A well defined service can be ***predicted*** to have some set of performance characteristics
  - The user knows in advance what to expect of the service in terms of performance or other characteristics
- A well defined service is ***verifiable***:
  - It can be measured/tested at the service delivery points and found to conform to some set of prescribed service characteristics
  - This also allows for important pragmatic processes related to fault isolation
- A well defined service is ***repeatable***
  - It performs the exact same way every time it is invoked
- A well defined service covers the service ***end to end***
  - It should not matter which provider(s) or which networks are involved in delivering the service for the end users – they will all result in similar and conforming capabilities.

# Why must the GLIF formally define our services…

- **GLIF users need to know exactly what they can expect from GLIF services**
- **GLIF network engineers need to know exactly what capabilities they need design and deploy within their networks – and how they need to interoperate with their peers**
- GLIF is an integrated facility
  - It represents the sum of many different networks around the world, employing many different types of hardware, with varying capabilities and service models.
  - In order to deliver wavelength services across these heterogeneous and globally distributed facilities, we must be able to engineer and concatenate the services provisioned across any combination of these networks into *deterministic* data paths that are
    - Predictable
    - Verifiable
    - Repeatable
    - End to end

# An Ethernet Light Path (on Tuesday)

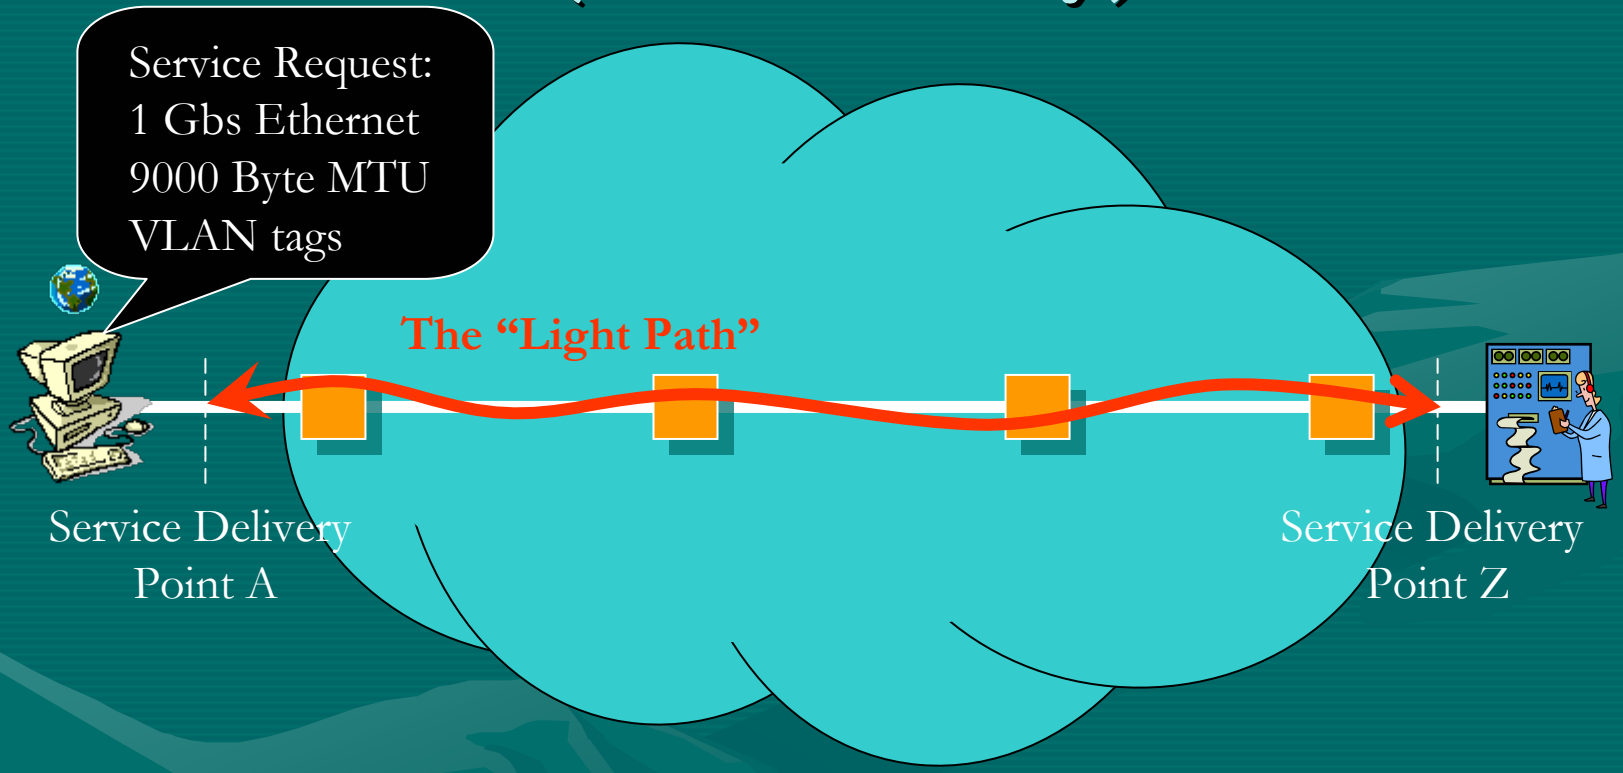

# An Ethernet Light Path (on Thursday)

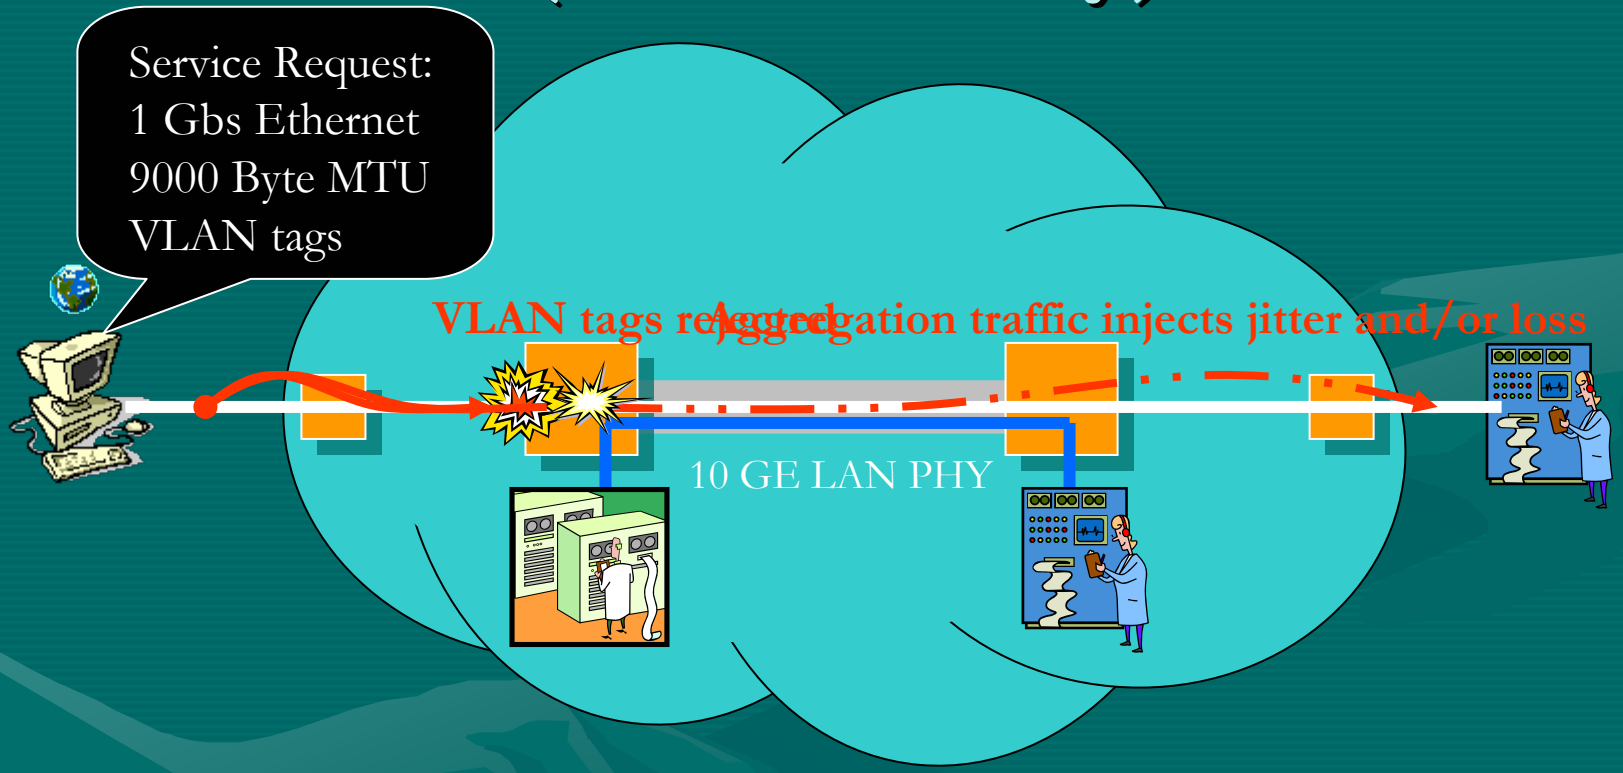

# An Ethernet Light Path (on Friday)

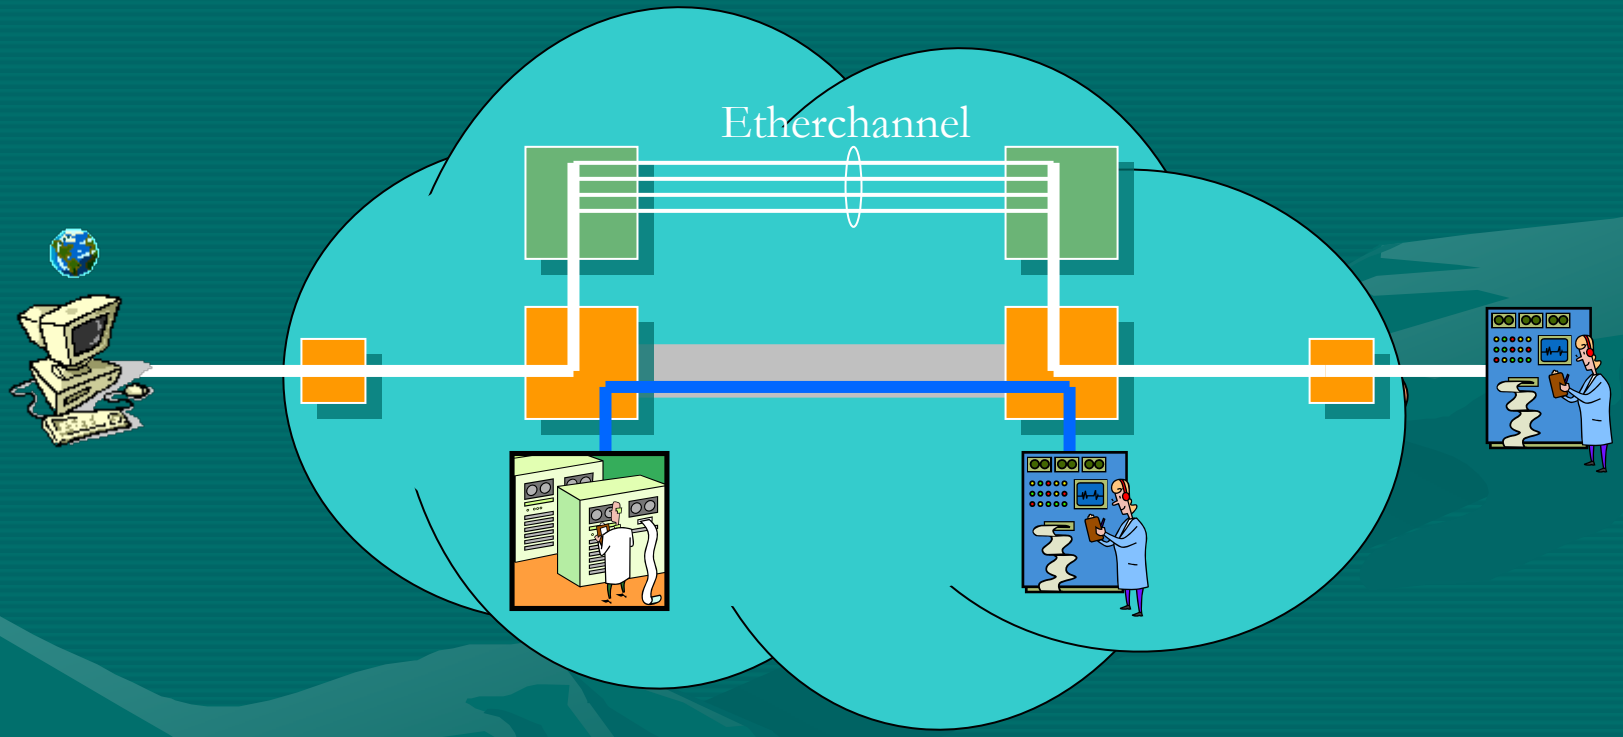

# Examples of common [mis-]assumptions: "Ethernet Light Path"

- **An Ethernet Light Path has only two end points (i.e. only two sources of traffic)**
  - What happens if it does not? (e.g. VLAN linking, broadcast storms, etc.)
- **Large MTUs are supported on all big fat pipes**
  - Not a safe assumption (even today) – often limited to 1500 Byte MTUs
- **"10Gigabit" means 10,000,000,000 bits/sec …always…end to end**
  - Transoceanic links only support 9.4 Gbs Sonet/SDH WANPHY
- **A point to point Light Path will not drop packets**
  - Almost all transport gear incorporates some type of buffering – particularly ethernet switch buffers may be overrun by bursty traffic…
  - Does the path infrastructure support flow control?
- **Ethernet frames will be delivered in order and transparently**
  - VLAN tags are often not carried end to end. Intermediate Ethernet infrastructure often uses VLANs themselves for provisioning Light Paths through the intermediate
  - Bonded etherchannel configurations reorder packets and/or constrain bandwidth
  - Is Spanning Tree supported? Is it intercepted by transit switches?
- **An ethernet light path does not introduce jitter**
  - Ethernet aggregation almost always causes queueing and introduces jitter – which can be compounded by multi-hop paths with raw aggregation provisioning…
  - Flow control/shaping introduces jitter end-to-end

# Moving from 1000 emails to 1000 milliseconds…

- Current processes for creating end-to-end light paths is fraught with extreme communication challenges between network operators, and between operators and applications developers/scientists. Service instantiations are slow and often variable.
- Early adopter tools for service establishment and agreement on service definitions
- Fully automated provisioning and a wide range of service capabilities…in seconds (or less)

# Service Definitions

- **Service Definitions are arbitrary but globally accepted descriptions of network capabilities.**
  - Each Service Definition specifies a fixed set of service parameters, and associated values/defaults for those parameters, that form the complete set of measurable service characteristics.
    - E.g. MTU >= 9000 Bytes; or
    - SpanningTreeProtocol=notTransported;
  - Any characteristic of a service that is not explicitly defined in the service definition is explicitly undefined
    - i.e. the user can neither expect it to be present nor expect it to be absent.
  - Different "services" should reflect fundamentally differnet network transport capabilities. Service characteristics should be used to indicate minor variances in particular instances, or to reflect availability of resources.
- **Service Definitions are living specifications**
  - They can be modified, augmented, enhanced, refined, etc from time to time as the GLIF community sees fit.
    - In the near term, the GLIF should adopt several basic broad and inclusive definitions that allow the process of define-deploy-review-refine to begin.
- **There may be multiple Service Definitions** (e.g. Ethernet, Sonet, Infiniband, Packet)
  - If properly designed, some services may be subsets of others, or offer a base set of characteristics inherited by an entire family of service dialects.

# Beginning the Process

- **At this Fall2005 GLIF meeting, we are submitting skeleton/draft descriptions of three initial Common Service Definitions**
  - We would like to see comment and discussion on these specifications
- **We would like to see a general adoption of the resulting initial rudimentary specifications by Dec 31st 2005.** (Perhaps an ad hoc committee of GLIF networks/user can meet at SC05 to adopt the 1.0 Specs…)
  - This will provide a guideline for service deployment, Open Xchange Point engineering, and conformance testing by the GLIF networking community over the coming year(s)
  - And it will allow GLIF users to begin architecting and re-engineering their applications to leverage these resources and capabilities.
- **We will translate these specifications to textual white papers and XML documents over the spring of 2006, and place these results on the GLIF web site** (or other such repositories as the community sees fit) for public review.
- **We propose a CSD Review at the next GLIF meeting in Tokyo in September 2006 to begin the iteration process.**

# Proposed Initial GLIF Services:

- **EtherPipeBasic** – Simple point to point transport of ethernet frames.
- **EtherPipeUltra** – Transport of high capacity flows, and transparent to [most] layer2 services.
- **TDMBasic** – Sonet/SDH service. Deterministic latency/jitter transport, intercontinental and commercial interworking
- LightWaveBasic – All-Photonic transport of ITU compliant waves. (work in progress…)
- PacketLSP – Half-baked idea for MPLS/IP service to complement the other service definitions. (Given commercial sector activity in layer3 VPNs and Ethernet VPLS, this may not be on the GLIF critical path)

# EtherPipeBasic 1.0beta

- **Service_Name** := EtherPipeBasic;
- **Version** := 1.0 beta 2005.09.30;
- **Description** := The “EtherPipeBasic” service is a point to point service that transports ethernet frames from one end to the other. This service is intended to be a useable ethernet connection for moving ethernet framed payloads – not for supporting ethernet networking (i.e. VLAN tags and Spanning Tree may or may not be transported by instances of this service.);
- **Framing** := IEEE 802; // Standard Ethernet frames…
- **MTU** := >= 1500 Bytes;
- **Bandwidth** := 100 Mbs ... 1 Gbs; // Continuous BW range 100-1000
- **Frame_Order** := In_Order; // Re-ordering not allowed
- **BER** := <= 10**-12; // One BE per thousand seconds

# EtherPipeUltra 1.0 beta

- **Service_Name** := EtherPipeUltra;
- **Description** := The “EtherPipeUltra” service is a point to point service that is to be used by large capacity users with certain performance criteria. The –Ultra service is compatible with the EtherPipeBasic services thus enabling a –Basic service request to be satisfied with –Ultra infrastructure;
- **Include** := EtherPipeBasic;
- **Framing** := IEEE 802;
- **MTU** := >= 9000 Bytes;
- **Bandwidth** := { 1|2|3|4|5|6|7|8|9|10 } Gbs; /* User can request BW in 1 Gbs granularity, or in the granularity specified in the –Basic service
- **Frame_Order** := In_Order;
- **BER** := <= 10**-12;
- **VLAN_Transport** := { True | False }; // User must explicitly indicate whether they want VLAN information carried uninterrupted E2E.
- **SPT_Transport** := { True | False }; // User must explicitly indicate whether they want SPT information carried uninterrupted E2E.

# TDMBasic 1.0beta

- **Service_Name** := TDMBasic
- **Version** := 1.0 beta 2005.09.30
- **Description** := The TDMBasic service is intended to be a Sonet/SDH point to point capability compatible with Next Gen Sonet/SDH features.
- **Framing** := ITU G.707;
- **Bandwidth** := {51.84; 9953.28; 51.84} Mbs // Any multiple STS1
- **MAX_Latency** := Minimize;

# PacketLSP 1.0 beta

- **Service_Name** := PacketLSP;
- **Version** := 1.0beta 2005.09.30
- **Description** := Packet LSPs are intended for relatively small flows that can be handled in large numbers over an MPLS/IP backbone, or for larger flows (~1 Gbs) required to access or egress other lower layer Light Path infrastructure.
- **Bandwidth** :=
- **Jitter** :=
- **BER** :=
- **Packet_Order** := Maintained;

# LightWaveBasic 1.0beta

- **Service_Name** := LightWaveBasic 1.0beta
- **Version** := 0.1 alpha 2005.09.30
- **Description** := This service is a true ITU compliant photonic wave transport. This service is framing agnostic but assumes certain basic encoding and modulation techniques that allow for engineering of service infrastructure. This service does not do wave translation or regeneration.
- **Wavelength_Spacing** := 100 GHz;
- **Maximum_Modulation_Rate** := 13 Ghz;

# Notes and Issues on Provisioning

- **The Common Service Definitions say *nothing* about how these services are or should be engineered or architected or provisioned inside the intermediate networks.**
  - The service provisioning between the end points may take a number of different paths over potentially many different technologies (e.g. Ethernet over GFP encapsulation over a SDH link…) Indeed, the intermediate networks may request lower layer Light Paths between themselves to support upper layer Light Path service requests from other users.
  - This engineering/provisioning is left to the individual service domains to implement as they see fit – as long as the service characteristics are maintained end-to-end.
- **A service domain may reject a service request if they do not have the resources available to support that request.**
  - Corollary: A domain should not advertise availability of a service unless they can provide that service in "most" instances.
- **Where CSD parameter offers the user a choice or range, the user must explicitly select one (or alternatively, the GLIF community can adopt an explicit default value.)**
  - This is important in that the first hop provider may need to assert the request upstream to other providers who may have a different default.
- **Looking forward, the GLIF Common Services Definitions and the iterative process of refining them should enable automated service routing techniques, scheduling, authorization/accounting, etc.**

# Notes and Issues on Provisioning

- The user should only have to ask their first hop service provider to establish a service instance.
  - The provisioning process downstream should be opaque to the end user so that the NSP has full flexibility in how they negotiate and fulfill the service request end-to-end.
  - This downstream service routing process is a Hard Thing and an open research and experimental networking topic. Keeping this process opaque to the user will allow the GLIF the greatest flexibility to implement and evaluate many different management models and user interfaces.
- Note: The "user" in the CSD model is simply the requester at the end of a service instance…
  - This does not break or impede novel service provisioning or management concepts, it simply defines the service characteristics experienced in the data plane at the service delivery points.

# Next Steps

- Timeline of milestones:
  - Now: Need ~four SMEs (people) and ~four Appl scientists (also people:-) to review, comment, and finalize v1.0
  - SC05 (Nov 05) -> Final comments and GLIF adoption
  - Apr 06 – Document and XML templates available on web site
  - Sep 06 – Community review and evolution
  - SC06 (Nov 06) Shake down with SCinet
- Open Issues:
  - Who will whip the slaves?
  - Who are the slaves?
  - How do we store and disseminate the resulting specifications and documents?

# Thanks!

- The "Common Services Definition" white paper (Sobieski/Lehman April 05) can be found at:

  dragon.east.isi.edu

Or contact

Jerry Sobieski (MAX)

or

Tom Lehman (ISI East)