

Evolution of OSCARS

Chin Guok, Network Engineer

ESnet Network Engineering Group

Winter 2012 Internet2 Joint Techs

Baton Rouge, LA

Jan 23, 2012





Outline

What was the motivation for OSCARS

History of OSCARS

What's new with OSCARS (v0.6)

Where is OSCARS (v0.6)

Who's using OSCARS

OSCARS supporting research

What have we learned

What's next for OSCARS

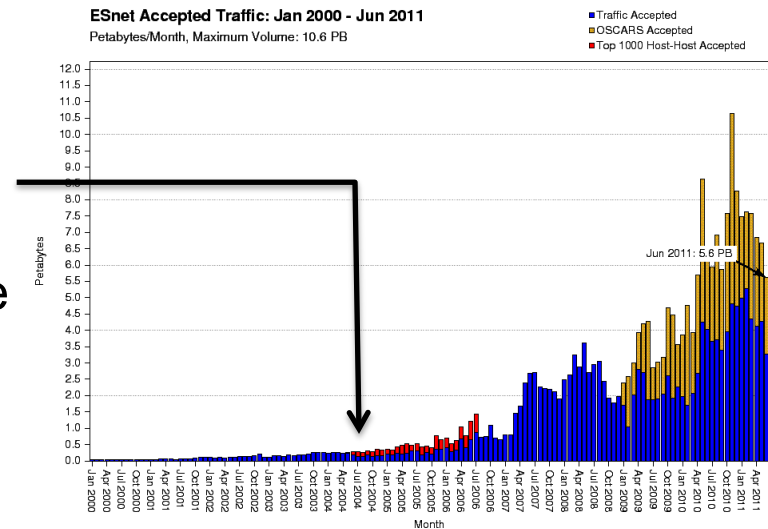
OSCARS in a nutshell



What was the motivation for OSCARS

- A 2002 DOE Office of Science High-Performance Network Planning Workshop identified bandwidth-on-demand as the most important new network service (for):
 - Massive data transfers for collaborative analysis of experiment data
 - Real-time data analysis for remote instruments
 - Control channels for remote instruments
 - Deadline scheduling for data transfers
 - “Smooth” interconnection for complex Grid workflows (e.g. LHC)

- Analysis of ESnet traffic done in early 2004 indicated that the top 100 host-to-host flows accounted for 50% of the utilized bandwidth





History of OSCARS

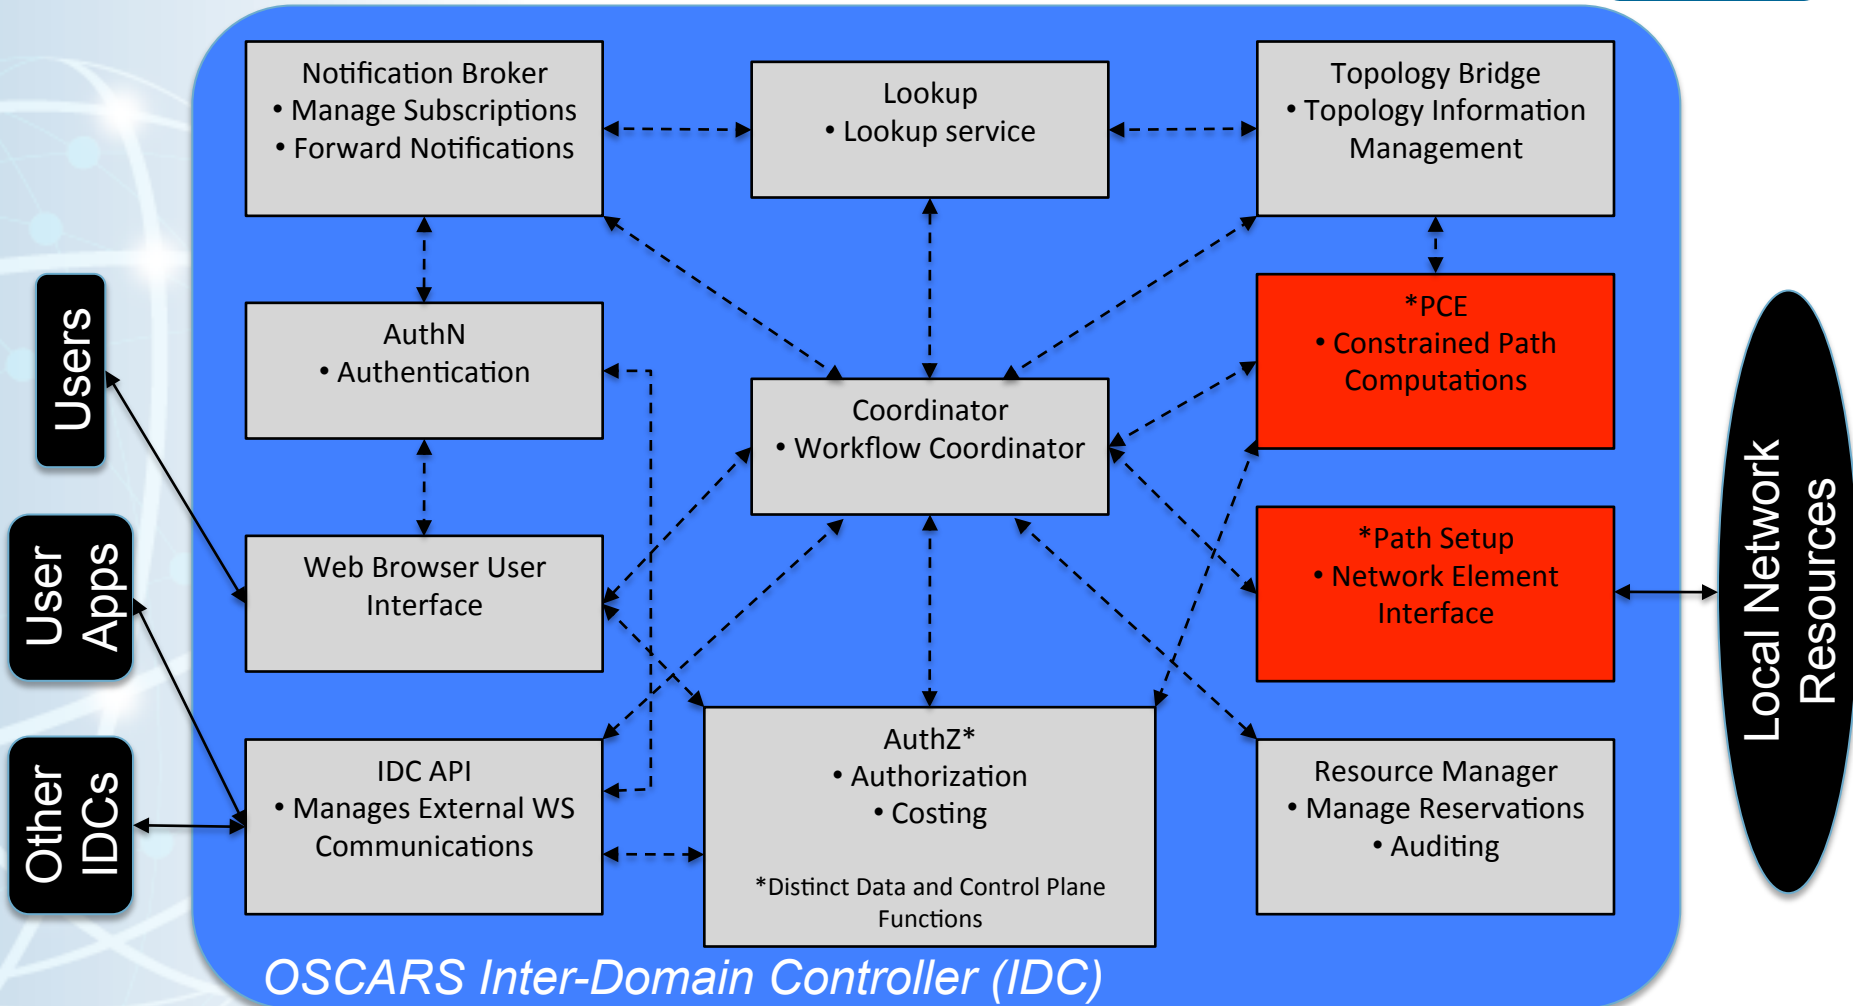
- OSCARS v0.6 RC1 released (Dec 2011)
 - OSCARS v0.6 is field tested in SCinet (SC11), and ESnet ANI 100G prototype network (Nov 2011)
 - **OSCARS interops with OGF NSI protocol v1 using an adapter at NSI Plugfest at GLIF Rio, participants include OpenDRAC (SURFnet), OpenNSA(NORDUnet), OSCARS(ESnet), G-lambda (AIST), G-lambda (KDDI Labs), AutoBAHN (GÉANT project), and dynamicKL (KISTI) (Sep 2011)**
 - OSCARS v0.6 SDK released allowing researchers to build and test PCEs within a flexible path computation framework (Jan 2011)
-
- **DICE IDCP v1.1 released to support brokered notification (Feb 2010)**
-
- First use of OSCARS by SCinet (SC09) to manage bandwidth challenges (Nov 2009)
 - Successful control and data plane interop between OSCARS, g-Lambda and Harmony using GLIF GNI-API Fenius (Nov 2009)
 - **Draft architecture design for OSCARS v0.6 (Jan 2009)**
-
- **Successful control plane interop between OSCARS and g-Lambda using GLIF GNI-API GUSI (GLIF Unified Service Interface) (Dec 2008)**
 - **DICE Inter-Domain Controller (IDC) Protocol v1.0 specification completed (May 2008)**
-
- Successful Layer-2 reservation between ESnet (OSCARS) and GÉANT2 (AutoBAHN), and Nortel (DRAC) (Nov 2007)
 - First dynamic Layer-2 interdomain VC between ESnet (OSCARS) and Internet2 (HOPI) (Oct 2007)
 - First ESnet Layer-2 VC configured by OSCARS (Aug 2007)
 - **Adoption of OGF NMWG topology schema in consensus with DICE Control Plane WG (May 2007)**
-
- First dynamic Layer-3 interdomain VC between ESnet (OSCARS) and Internet2 (BRUW) (Apr 2006)
 - **Formulation of DICE (DANTE, Internet2, CANARIE, ESnet) Control Plane WG (Mar 2006)**
 - **Collaboration with GÉANT AMPS project (Mar 2006)**
-
- First production use of OSCARS circuit to reroute LHC Service Challenge traffic due to transatlantic fiber cut (Apr 2005)
 - **Collaboration with Internet2 BRUW project (Feb 2005)**
-
- **Funded as a DOE project (Aug 2004)**



What's new with OSCARS (v0.6)

- Up till OSCARS v0.5 the code was tailored specifically to production deployment requirements
- Monolithic structure with distinct functional modules
- In OSCARS v0.6 the entire code base was re-factored to focus on enabling research and production customization
- Distinct functions are now individual processes with distinct web-services interfaces
- Flexible PCE framework architecture to allow “modular” PCEs to be configured into the path computation workflow
- Extensible PSS module allows for multi-layer, multi-technology, multi-point circuit provisioning
- Protocol used to make requests to OSCARS (IDC protocol) was modified to include an “optional constrains field” to allow testing of augmented (research) features without disrupting production service model

Modularization in OSCARS v0.6

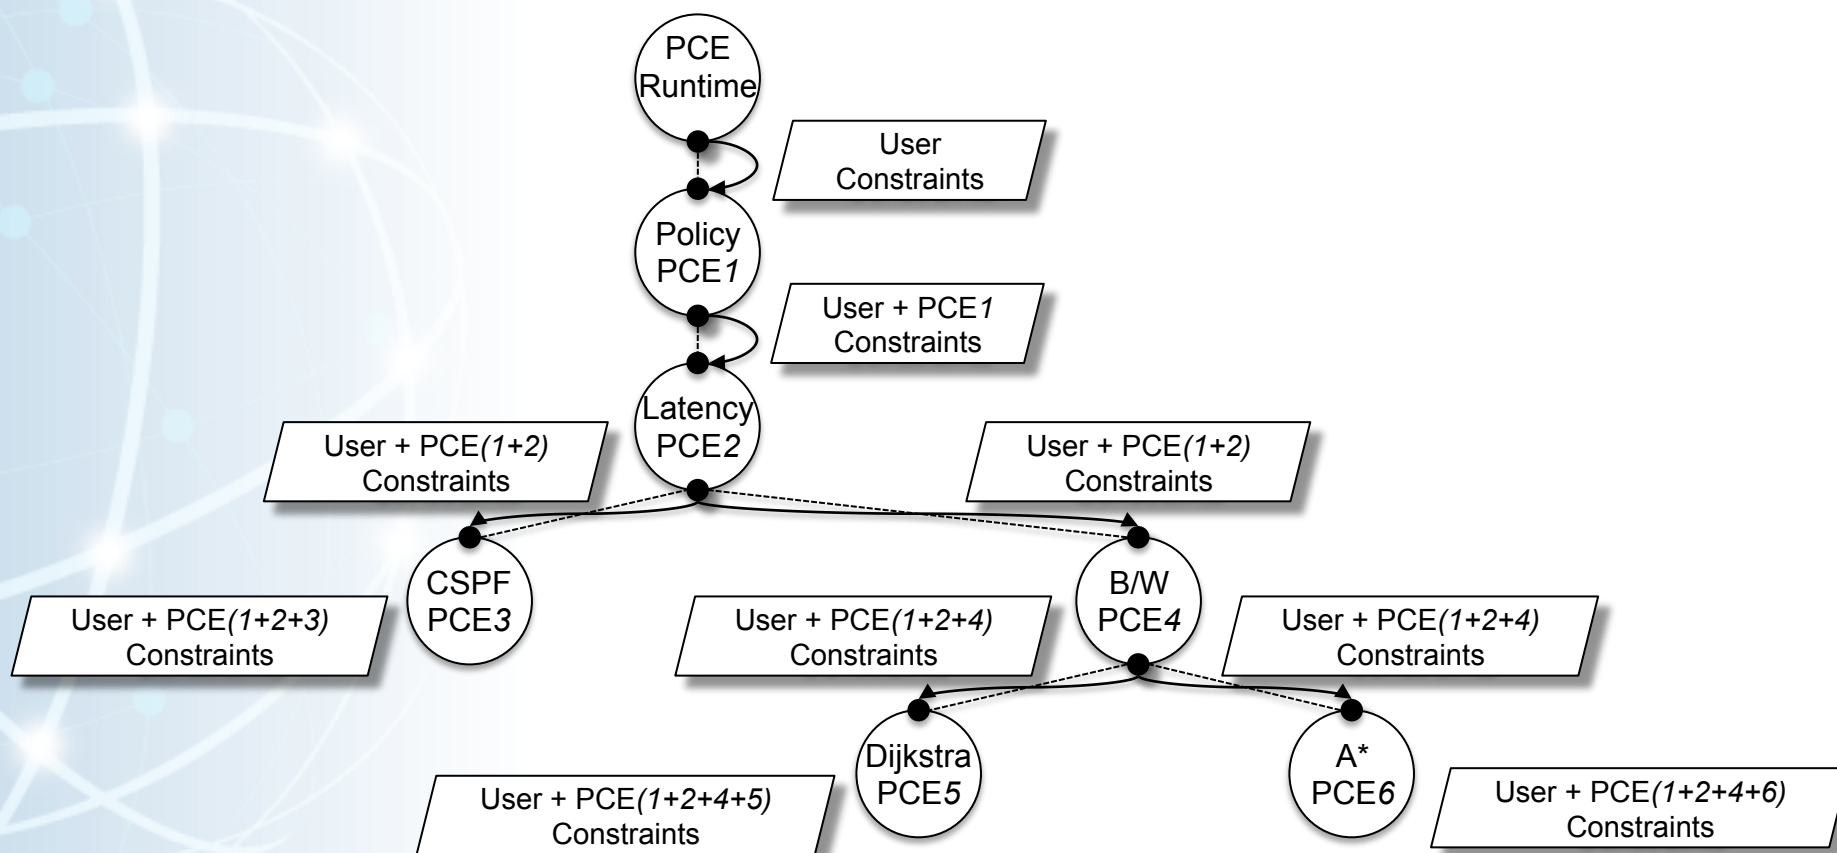


** Current focus of research projects*

Flexible PCE Framework

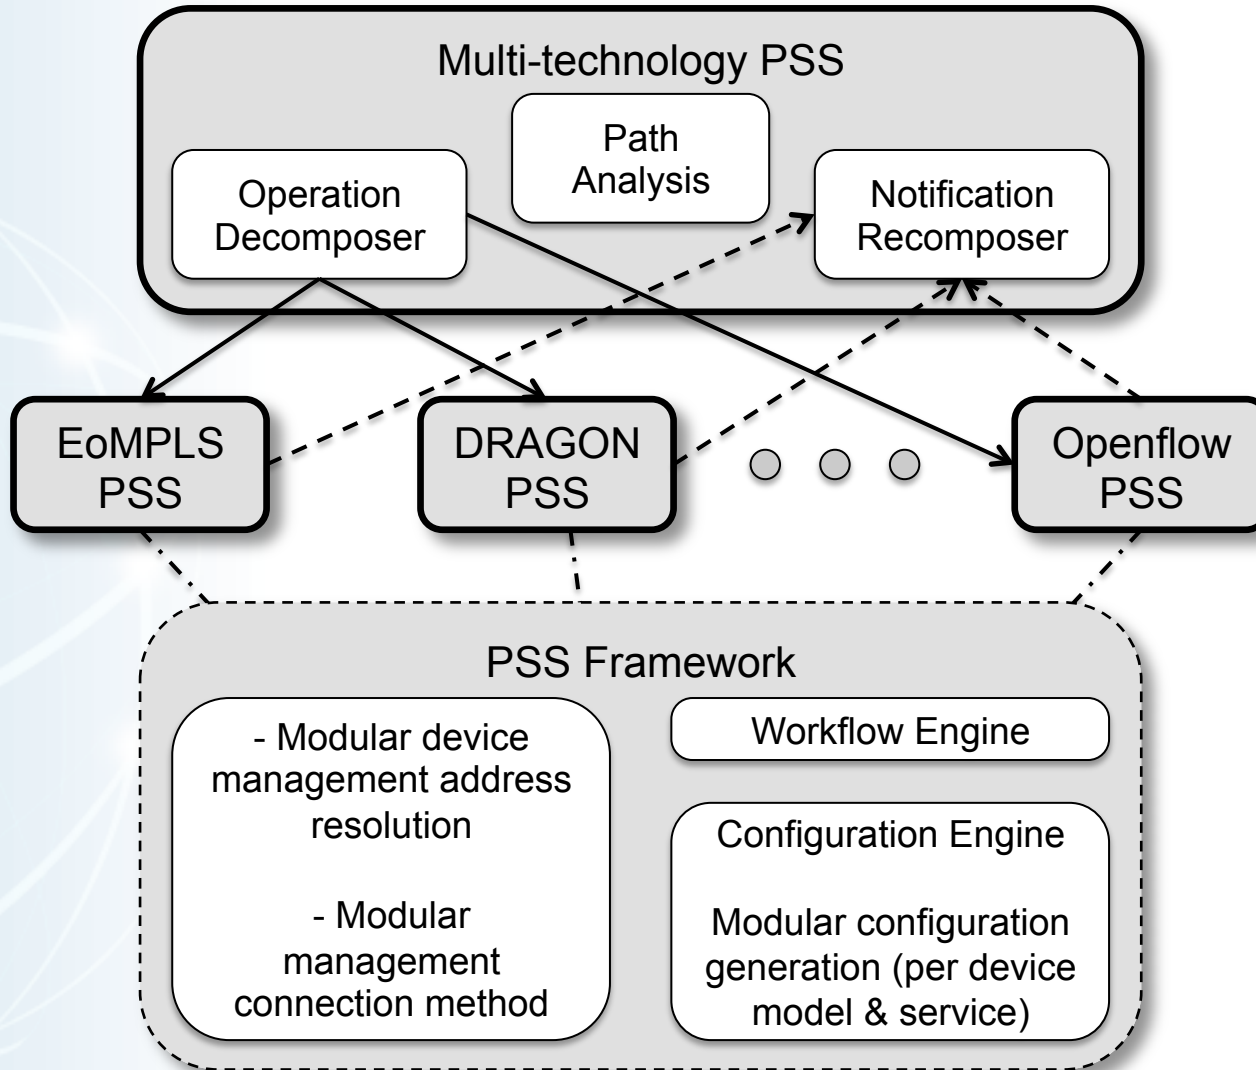
Supports arbitrary execution of distinct PCEs

- Example graph of PCE Modules



Constraints = Network Element Topology Data

Extensible PSS Module





Protocol Extension to IDCP

- “optionalConstraint” added to support research features without constant need to change base protocol
- Enhancements prototyped in “optionalConstraint” will migrate to base protocol once they have been baked

```
<xsd:complexType name="resCreateContent">
  <xsd:sequence>
    <xsd:element name="messageProperties" type="authP:messagePropertiesType" maxOccurs="1" minOccurs="0"/>
    <xsd:element name="globalReservationId" type="xsd:string" maxOccurs="1" minOccurs="0"/>
    <xsd:element name="description" type="xsd:string" />
    <xsd:element name="userRequestConstraint" type="tns:userRequestConstraintType" maxOccurs="1" minOccurs="1" />
    <xsd:element name="reservedConstraint" type="tns:reservedConstraintType" maxOccurs="1" minOccurs="0" />
    <xsd:element name="optionalConstraint" type="tns:optionalConstraintType" maxOccurs="unbounded" minOccurs="0"/>
  </xsd:sequence>
</xsd:complexType>
...
<xsd:complexType name="optionalConstraintType">
  <xsd:sequence>
    <xsd:element name="value" type="tns:optionalConstraintValue"/>
  </xsd:sequence>
  <xsd:attribute name="category" type="xsd:string" use="required"/>
</xsd:complexType>
<xsd:complexType name="optionalConstraintValue">
  <xsd:sequence >
    <xsd:any maxOccurs="unbounded" namespace="##other" processContents="lax"/>
  </xsd:sequence>
</xsd:complexType>
```



Where is OSCARS (v0.6)

OSCARS v0.6 is gaining adoption and seeing production deployments

- Field tested at SC11
 - Deployed by SCinet to manage bandwidth/demo bandwidth on show floor
 - Modified (PSS) by Internet2 to manage Openflow switches
 - Modified (Coordinator and PSS) by ESnet to broker bandwidth and coordinate workflow
- Currently deployed in ESnet 100G Prototype Network
 - Modified (PSS) to support ALU devices and “multi-point” circuits
- Adopted by Internet2 for NDDI and DYNES
 - IU GRNOC has modified OSCARS v0.6 (PSS and PCE) to support NDDI OS3E

OSCARS v0.6 RC1 is now available (open-source BDS license)

- https://code.google.com/p/oscars-idc/wiki/OSCARS_06_VM_Installation

Who's using OSCARS

- Currently deployed in 21 networks including wide-area backbones, regional networks, exchange points, local-area networks, and testbeds.
- Under consideration for an additional 26 green field deployments in 2012



Deployments using OSCARS v0.6 above dotted line



CALTECH



AMPATH™
INTERNATIONAL EXCHANGE POINT IN MIAMI

NORTHROP GRUMMAN



GpENI



VANDERBILT





OSCARS Supporting Research

Changes made to OSCARS v0.6 has been extremely useful in supporting research projects, e.g.:

- Advance Resource Computation for Hybrid Service and Topology Networks (ARCHSTONE)
 - <http://archstone.east.isi.edu>
- Coordinated Multi-layer Multi-domain Optical Network Framework for Large-scale Science Applications (COMMON)
 - <http://www.cis.umassd.edu/~vvokkarane/common/>
- End Site Control Plane System (ESCPS)
 - <https://plone3.fnal.gov/P0/ESCPS>
- Network-Aware Data Movement Advisor (NADMA)
 - <http://www2.cs.siu.edu/~mengxia/NADMA/>
- Virtual Network On-Demand (VNOD)
 - <https://wiki.bnl.gov/CSC/index.php/VNOD>



What have we learned

- Solve a real problem
 - Understanding the practical problem made it easy to judge if OSCARS met the original design requirements
- Be practical
 - It was necessary to “punt” or simplify specific designs in OSCARS in order to scope the problem and solve the core issues
- Balance optimization and customizability
 - As OSCARS started gaining adoption in various networks, it became evident that the ability to customize OSCARS was more important than optimization. *This was the driver for OSCARS v0.6.*
- Engage the right community
 - For OSCARS to be successful in an end-to-end and multi-domain environment, it was necessary to engage communities (e.g. DICE, GLIF) with similar goals and requirements
 - **OSCARS has been and is a collaborative effort!**



What's next for OSCARS

- Short-Term
 - Evaluate contributions from collaborative research projects (e.g. ARCHSTONE, COMMON, etc), and certify code for production deployment
 - Support standards protocol (e.g. OGF NSI) natively into OSCARS
 - Evaluate and develop solutions driven by production requirements (e.g. multi-point overlay networks, Nagios plug-ins, packaging)
- Long-Term
 - Explore co-scheduling and composable services framework to advance “Intelligent networking”
 - Develop strategies for OSCARS to mature in a larger community base (e.g. FreeBSD)

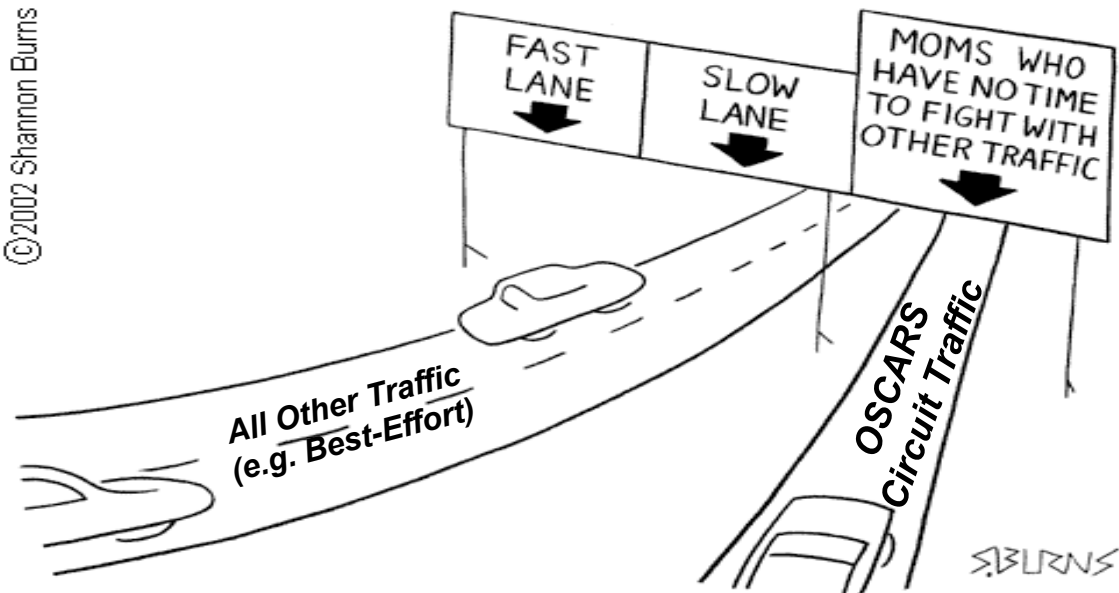
OSCARS in a nutshell



- OSCARS is an open-source collaborative project
- OSCARS is gaining critical mass in the deployment in networks
- OSCARS is support new exciting research in the area of intelligent networking
- OSCARS supports network resource management
...
and “MOMS” (Movement Of Massive Science)

www.shannonburns.com

©2002 Shannon Burns



For more info: <https://code.google.com/p/oscars-idc/>

Questions?



chin@es.net | Chin Guok