

# Best Practices for Connections between ESnet and Regional Optical Networks

Michael Collins, Joseph Burrescia, William Johnston ([wej@es.net](mailto:wej@es.net)), Chin Guok, Joseph Metzger, Michael O'Connor

## Introduction

---

Regional Optical Networks (RONs) play an important role providing networking services to the U.S. R&E community. Typically, a RON is a facility based optical network running on owned assets that can span geographic regions as large as one or more states. RONs serve as aggregators for connected sites and/or other local R&E networks. Many DOE sponsored scientists have a RON as their primary high bandwidth connection to external facilities and wish to use this path for collaboration with researchers at DOE Laboratories.

ESnet may opportunistically form interconnections with RONs. There are three basic types of possible connections between ESnet and a RON. Only two of the types are available at this time:

- o ESnet and a RON may peer at Layer3
  - This involves a traditional BGP based exchange of routes and traffic.
- o ESnet and a RON may set up a layer2 or layer 3 virtual circuit between sites or between experimental facilities within the sites.
- o ESnet cannot currently accept direct layer1 connections.

This document discusses the interconnections and ESnet's expectations.

## 1 Interconnection between ESnet and a RON

---

ESnet may opportunistically form interconnections with a RON. Usually, the connections are formed where the interconnection does not require new long-haul communication circuits. In order to protect network resources, ESnet maintains a formal boundary at the edges of its network. The demarcation point between ESnet and the RON is effectively the link between the ESnet switch/router and the RON connection device.

### 1.1 BGP peering

---

#### *Description*

A BGP peering is set up at layer 3 and permits the exchange of traffic between ESnet and RON customers. ESnet can accept both IPv4 and IPv6 BGP sessions. IPv6 traffic is transported natively by the ESnet backbone.

#### *Routing Plan*

ESnet and a RON will jointly establish a routing plan to specify prefix announcements. Common elements of a routing plan include the following, any of which are modifiable in the final plan:

- o The RON will only announce its customer prefixes to ESnet. In this context customer includes a leaf R&E network.
  - ESnet will announce the prefixes of sites that may access the RON using ESnet.
- o The RON announces ESnet prefixes only to those sites that reach ESnet via the RON
  - ESnet will announce the RON prefixes to sites that reach the RON through ESnet.

The routing plan will include such special issues as a RON

- o Supplying backup connectivity to another R&E network
- o With international connectivity

### **Expectations**

There are several expectations involving BPG peering.

- o The connection is established for the settlement free exchange of traffic
- o The RON must not set a default route pointing to ESnet
  - ESnet will not set a default route pointing to the RON.
- o The RON will supply ESnet with a list of prefixes the RON will announce and update the list as required.
  - Alternatively, and preferably, the RON should register its announcements in a Routing Registry.
  - ESnet registers its announcements *in the Internet Routing Registry (IRR)* and ESnet filters all BGP sessions.
- o Prefixes and AS numbers in BGP paths announced by the RON must be registered in the appropriate regional Registry
  - Prefixes and AS numbers announced by ESnet are registered *in the IRR*.
- o The RON may request that the BGP session be protected with a MD5 password.
  - ESnet does not routinely use MD5 passwords.
- o The RON may collect and publish statistics pertaining to the connection
  - ESnet will collect and publish statistics pertaining to the connection

## **1.2 Multicast Peering**

---

### **Description**

Interdomain multicast peerings are set up at layer 3 and permit the creation of dynamic PIM source/group (S,G) state and the exchange of traffic between ESnet and RON customers. ESnet supports both IPv4 and IPv6 multicast packet forwarding. IPv6 multicast traffic is transported natively by the ESnet backbone.

### **Routing Plan**

The routing plan is the same as that for BGP peering with the additional requirement that Multiprotocol BGP (MBGP) is required.

### **Expectations**

ESnet expects that multicast enabled RONs implement the following multicast protocols to exchange multicast traffic with ESnet.

PIM V2 – Protocol Independent Multicast Sparse Mode

- o PIM adjacency between ESnet and a RON is required. PIM is essential on all segments of the multicast shortest path distribution tree, from a source to all its intended receivers.
- o At its edge, ESnet filters PIM groups intended for internal use in both directions.

MSDP – Multicast Source Discovery Protocol

- o MSDP describes a mechanism to connect multiple PIM-SM domains together.
- o At its edge, ESnet filters MSDP advertisements intended for internal use.

MBGP - Multiprotocol Border Gateway Protocol

- o Is an extension to BGP that enables BGP to carry routing information for multiple network layers and address families.

SAP - ESnet does not forward the Session Advertisement Protocol group 224.2.127.254

### 1.3 Virtual circuits between ESnet and RON sites.

---

#### *Description*

If the volume, duration, or science requirements of a flow of science data between a RON site and an ESnet site warrants, ESnet and the RON may create a virtual circuit between the sites (or the science experiment facilities at the sites). The circuit consists of least five independently managed parts:

- o The portion of the circuit internal to the ESnet site up to the ESnet edge device
- o The ESnet portion of the circuit
- o The ESnet/RON interconnect
- o The RON portion of the circuit
- o The portion of the circuit internal to the RON site up to the RON edge device

The ESnet portion of the circuit may be managed statically or by the On-demand Secure Circuits and Advance Reservation System ([OSCARS](#))[1]. A complication of the ESnet/RON interconnection is that there is no standard network-network interface (NNI) for inter-domain virtual circuit handoff. The issue has been pragmatically addressed by developing the virtual circuit services as a collaboration among a number of the R&E serving networks, optical testbeds, and organizations (OSCARS, Internet2[2], GEANT[3], and the European regional networks (NRENs), Brookhaven National Laboratory[4], General Atomics[5], SLAC[6], DRAGON[7], and Ultra Science Network[8]). Currently the agreed data plane handoffs are IP packets for layer 3 virtual circuits, and VLAN packets for layer 2 virtual circuits.

#### *Expectations*

There are several expectations regarding Virtual Circuits:

- o The RON can deliver the reliability necessary to insure that programmatic needs are satisfied. For example, currently we understand the LHC circuit requirements as having a minimum reliability of 99.95%.
- o The routing policy of the RON excludes other traffic sources from using the circuit.
  - The routing policy of ESnet excludes other traffic sources from using the circuit.
- o If dynamically provisioned, the RON circuit setup will require an authentication and authorization process.
  - OSCARS will require an authentication and authorization process to reserve a virtual circuit in ESnet.

### 1.4 Common expectations for connections between ESnet and a RON

---

- o The RON will supply contact information as requested by ESnet
  - ESnet will supply contact information as requested by the RON
- o The RON will maintain at least a 8hX5d NOC contact that will answer the phone
  - Outside of normal working hours, the RON will supply a contact which can initiate a trouble ticket and contact trouble shooters
  - ESnet maintains a 24X7 NOC
- o The RON will maintain a trouble ticket system for tracking problems and assign a trouble ticket number to problems reported by ESnet.
  - ESnet maintains a trouble ticket system and assigns ticket numbers to anything affecting the network
- o The RON must have a documented escalation procedure to ensure the timely resolution of problem reports.
  - ESnet has such an escalation procedure.

### 1.5 Contact information

---

Contact information is available at <http://www.es.net/hypertext/welcome/pr/contactus.html>

## 1.6 References

---

- [1] ESnet OSCARS webpage: <http://www.es.net/oscars>
- [2] Internet2 BRUW Project: <http://discvenue.internet2.edu/wordpress>
- [3] GEANT PACE Project: <http://pace.geant2.net>
- [4] BNL TeraPaths Project: <http://www.atlasgrid.bnl.gov/terapaths>
- [5] General Atomics QoS Project: <http://www.fusiongrid.org/network>
- [6] SLAC IEPM Project: <http://www-iepm.slac.stanford.edu>
- [7] <http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/WebHome>
- [8] UltraScienceNet Testbed: <http://www.usn.ornl.gov>