Open Connect

Overview

FAQ

Peering Information

Hardware Design

Software Design

Deployment Guide

ISP Inquiry

Overview

Netflix Open Connect Content Delivery Network

Around the world, people are enjoying nearly a billion hours per month of movies and TV shows from Netflix. Now, Netflix is enabling ISPs to get Netflix video data directly from Open Connect, a single-purpose content delivery network we've established.

ISPs can choose to have Open Connect Appliances within their datacenters, or to peer with the Open Connect network at common Internet Exchanges. Netflix will provide either form of access at no cost to the ISP.

ISPs interested in joining the Netflix Open Connect network can learn more from the Open Connect FAQ.

As part of Open Connect, we are also sharing our hardware design and the open source software components of the server. These designs are suitable for any other provider of large media files, and are very cost efficient. We welcome commentary and improvements, which will be shared with the community with the goal of a faster, less expensive Internet for all.

Open Connect

Overview

FAQ

Peering Information

Hardware Design

Software Design

Deployment Guide

ISP Inquiry

Open Connect FAQ

Where can I peer with Open Connect?

Open Connect currently offers settlement-free peering in the following locations:

- Ashburn, VA: Equinix Internet Exchange
- Atlanta, GA: Telx Internet Exchange
- Chicago, IL: Equinix Internet Exchange
- London, UK: LINX and LONAP
- Los Angeles, CA: Any2 Internet Exchange
- Miami, FL: NAP Of The Americas
- New York, NY: Telx Internet Exchange, Equinix Internet Exchange, NYIIX
- San Jose, CA: Equinix Internet Exchange

The latest peering information is available here.

How can I peer with Open Connect?

Review our peering guidelines and submit an ISP inquiry.

What if I want to peer at a different location?

Submit an ISP inquiry or contact peering@netflix.com and the Netflix peering team will assess the viability of peering in additional locations.

What if I have large concentrations of Netflix traffic in one or more of my metro regions?

Depending on your traffic profile, it may be more efficient to install Open Connect appliances in one or more metro network areas. Typically, this makes sense for individual markets serving a population of 100,000 or more broadband subscribers.The Open Connect appliance requires a 10 Gbps port. If your network fits this profile, you can find more details on deploying Open Connect Appliances in our Open Connect Appliance Deployment Guide.

Open Connect Peering Guidelines

General Overview

Netflix has an open peering policy and actively peers with networks that have end-users viewing Netflix content. We do not peer with networks that only have a homemarket of end-users in countries in which we do not presently operate.

General Requirements

1. The ISP network must be located in or connected to the same peering locations as the Netflix network (AS2906).
2. Both parties must provide a 24/7 contact who can escalate critical issues in a timely fashion.
3. Both parties are expected to register their routes in a public Internet Routing Registry (IRR) database, for the purpose of filtering. Both parties shall make good faith efforts tokeep this information up to date.
4. Either party may accept or reject peering requests as network management needs require.
5. Neither party shall establish a static route, a route of last resort, or otherwise send traffic to the other party for a route not announced via BGP.
6. Neither party shall announce to the other the more specific routes of prefixes learned via a third party transit customer.

Additional Private Interconnection Requirements

1. Interconnection(s) should occur in mutual locations with SMF 10Gb Ethernet interfaces.
2. At least 2 Gb/s of aggregate IPv4 traffic, measured using 95th Percentile in either direction, must be exchanged on an ongoing basis.
3. Each interconnection shall have no less than 10 Gb/s of capacity.

Join the Open Connect Network

To get the planning conversation started, submit an ISP Inquiry.

To request peering, submit an Open Connect peering request.

Questions

Contact the Open Connect team at peering@netflix.com

If you are interested in joining the Content Delivery or another team at Netflix, apply at www.netflix.com/jobs

Open Connect

Overview

FAQ

Peering Information

Hardware Design

Software Design

Deployment Guide

ISP Inquiry

Open Connect Appliance Hardware

Objectives

When designing the Open Connect Appliance Hardware, we focused on these fundamental design goals:

- Very high storage density without sacrificing space and power efficiency. Our target was fitting 100 terabytes into a 4u chassis that is less than 2' deep.
- High throughput: 10 Gbps throughput via an optical network connection.
- Very low field maintenance: the appliance must tolerate a variety of hardware failures including hard drives, network optics, and power supply units.
- Simple racking and installation. Front mounted power and network ports are the only things to connect at install time.

Open Connect Appliances are servers based on commodity PC components (similar to the model used by all large scale content delivery networks). We were influenced by the excellent write-ups from the Backblaze team, and use a custom chassis due to a lack of ready made options for a compact unit.

To achieve over 100 TB of storage, spinning hard drives provide the highest affordable density, in particular 36 3TB SATA units. The hard drives are not hot swappable, as we wish to avoid the operational burden of field service. For lower power utilization and simpler sourcing we select commodity units from two vendors and use software to manage failure modes and avoid field replacement. Dead drives reduce the total storage available for the system, but don't take it offline. We also add 1 TB of flash storage (2 solid state drives) for system files, logs and popular content. To augment the motherboard attached controller, we use two 16 port LSI SAS controller cards that connect directly to the SATA drives. This avoids I/O bottlenecks of SATA multipliers or SAS expanders, and also reduces system complexity.

From a compute point of view, the system has modest requirements moving bits from the storage to network packets on the interface. To reduce the power usage and hence also cooling requirement (which in turn reduces vibration from case fans) we use a single low power 4 core Intel Sandy Bridge CPU on a small form factor Supermicro mATX board with the full 32 GB of RAM installed.

We use redundant, hot swappable power supply units that have interchangeable AC and DC options for maximum installation flexibility. Zippy reversed the fan rotation of the units to allow mounting at the front of the case, and thus allow network and power connects to be positioned here.

The network card has two 10 Gbps modules, which can power a variety of SR and LR optic modules, for installation flexibility and scalable interconnection.

The following system was developed and first deployed at the end of 2011.

Bill of Materials

Description	Vendor & Model	Quantity
Chassis	TST custom	1
Motherboard	Supermicro X9SCM-F	1
Processor	Intel E3-1260L	1
Memory	8GB ECC 1333MHz	4
Hard Drive	Hitachi Deskstar 5K3000 3TB	36
Hard Drive (alternate)	Seagate Barracuda 7200.14 3TB	36
Controller	LSI SAS 92016-16i 16 port	2
Network card	Supermicro AOC-STGN-i2S	1
Redundant Power Supply Unit (AC/DC options)	Zippy MRW-5600V4V/DMRW-5600V4V	1
Misc.	2U active CPU Heatsink, SATA Cables, NIC optics	

Acknowledgements

Building this system we collaborated with a wide range of suppliers. In particular we would like to thank Netty, John, Tim, Dave and the team at Intequus, our system integrator. TST provided the custom chassis. For storage guidance and troubleshooting we would like to thank George and Andy at HGST, Ben, Steve and team at LSI, and John, Dave, Anik and team at Seagate.

Questions

Contact the Open Connect team at openconnectappliance@netflix.com

If you are interested in joining the Content Delivery or another team at Netflix, apply at www.netflix.com/jobs

Open Connect

Overview

FAQ

Peering Information

Hardware Design

Software Design

Deployment Guide

ISP Inquiry

Open Source Software

Open Connect Appliance Software

Netflix delivers streaming content using a combination of intelligent clients, a central control system, and a network of Open Connect appliances.

When designing the Open Connect Appliance Software, we focused on these fundamental design goals:

- Use of Open Source software
- Ability to efficiently read from disk and write to network sockets
- High-performance HTTP delivery
- Ability to gather routing information via BGP

Operating System

For the operating system, we use FreeBSD version 9.0. This was selected for its balance of stability and features, a strong development community and staff expertise. We will contribute changes we make as part of our project to the community through the FreeBSD committers on our team.

Web server

We use the nginx web server for its proven scalability and performance. Netflix audio and video is served via HTTP.

Routing intelligence proxy

We use the BIRD Internet routing daemon to enable the transfer of network topology from ISP networks to the Netflix control system that directs clients to sources of content.

Acknowledgements

We would would like to express our thanks to the FreeBSD community, the nginx community, and Ondrej and the BIRD team for providing excellent open source software. We also work directly with Igor, Maxim, Andrew, Sergey, Ruslan and the rest of the team at nginx.com, who provide superb development support for our project.

Questions

Contact the Open Connect team at openconnectappliance@netflix.com.

If you are interested in joining the Content Delivery or another team at Netflix, apply at www.netflix.com/jobs