



Netflix Open Connect Appliance Deployment Guide

This guide describes the deployment of embedded Open Connect Appliances. If you are interested in peering or an overview of the Open Connect program, see the [Open Connect web site](https://openconnect.netflix.com) (<https://openconnect.netflix.com>).

10 July 2018

Copyright © 2018 by Netflix, Inc. All rights reserved. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, without express permission from Netflix, Inc

[REQUIREMENTS FOR DEPLOYING EMBEDDED APPLIANCES >](#)

[NETWORK AND DATA CENTER CRITERIA >](#)

[TEAM ROLES >](#)

[TRAFFIC ANALYSIS AND SIZING >](#)

[SHIPPING, RECEIVING, AND INSTALLATION >](#)

[SHIPPING OCAS >](#)

[RECEIVING OCAS >](#)

[RACKING AND INSTALLATION >](#)

[OCA RETURNS AND REPLACEMENTS >](#)

[NEXT STEPS >](#)

[NETWORK CONFIGURATION >](#)

[OVERVIEW >](#)

[RECONFIGURING THE IP ADDRESS OF AN OCA >](#)

[ROUTER INTERFACE CONFIGURATION >](#)

[ROUTING AND CONTENT STEERING VIA BGP >](#)

[CLUSTERING ARCHITECTURES >](#)

[EMBEDDED OCAS COMBINED WITH PEERING SESSIONS >](#)

[FLASH-BASED APPLIANCES >](#)

[EXAMPLE ROUTER CONFIGURATIONS >](#)

[BROCADE LAG BGP >](#)

[CISCO IOS BGP >](#)

[CISCO IOS XR BGP >](#)

[CISCO IOS XR LACP >](#)

[CISCO IOS INTERFACE LACP >](#)

[CISCO IOS XR BRIDGED INTERFACE >](#)

[CISCO IOS SWITCH WITH VLAN >](#)

[JUNOS BGP >](#)

[JUNOS INTERFACE LACP >](#)

[JUNOS BRIDGED INTERFACE >](#)

[FILL, UPDATES, AND MAINTENANCE >](#)

[FILL WINDOWS >](#)

[FILL PATTERNS >](#)

[PEER FILLING >](#)

[TIER FILLING >](#)

[OUT-OF-CYCLE FILLS >](#)

[AUTOMATIC SOFTWARE UPDATES >](#)

[HEALTH MONITORING >](#)

[UNPLANNED OUTAGES OR SERVICE-IMPACTING EVENTS >](#)

[POTENTIAL SERVICE-IMPACTING MAINTENANCE >](#)

[PARTNER MAINTENANCE >](#)

[RETURNS AND REPLACEMENTS >](#)

[APPLIANCE RMA GUIDELINES >](#)

[PARTS RETURNS AND REPLACEMENTS >](#)

[CONTACT US >](#)

Requirements for Deploying Embedded Appliances

[NETWORK AND DATA CENTER CRITERIA >](#)

[TEAM ROLES >](#)

[TRAFFIC ANALYSIS AND SIZING >](#)

Netflix team members collaborate with you to tailor and optimize your implementation. Initially, Netflix will work with you to determine whether the combination of offload and fill traffic that is associated with an embedded OCA deployment will benefit your ISP. If embedded OCAs are warranted and you decide to deploy them in your network, the Netflix team will work with you to do capacity and network configuration planning as described in more detail [here](#).

To deploy embedded OCAs in your network, you should meet the following basic criteria. In certain non-standard deployments, your Netflix PEM can consider deviations from some of these requirements - however, deviations are analyzed on a case-by-case basis.

Network and data center criteria

Criteria	Description
Network	Your ISP should have an officially registered autonomous system number (ASN) that identifies your network.
Netflix traffic	Traffic requirements vary depending upon the content library that is being served, with the expected range being at least 5 Gbps of peak traffic . Lower traffic levels are typically served better by SFI peering alone.

Criteria	Description
Network capacity	To enable nightly fill and updates , you must have the capacity to handle 1.2 Gbps of inbound traffic for a 12-hour period per appliance. You can configure the fill window so that it occurs during low traffic periods.
Interconnection	You must connect to Netflix via SFI peering at one or more Netflix points of presence (PoPs) .
Installation	You must be able to physically install OCAs at your assigned site within 10 business days of receipt , or as soon as possible thereafter.
Physical connection	You must be able to provision 2-4 x 10 Gbps ethernet ports in a LACP LAG per OCA. The exact quantity depends on the OCA type.
Power supply maximum power output	750W

Team roles

To participate in the program, you will need to identify a person or a set of people who are available to perform the following roles. Your team members will work closely with the corresponding members of the Netflix team to facilitate the overall process.

Partner role	Description	Corresponding Netflix role(s)
Engagement manager	Serves as the main point of contact for the Netflix Open Connect partner engagement manager (PEM)	Netflix Open Connect Partner Engagement Manager (PEM)
Legal representative	Reviews and accepts the legal agreements that are required for OCA deployments	PEM, Open Connect legal representative, Open Connect business operations

Partner role	Description	Corresponding Netflix role(s)
Network engineer	Provides information about partner sites, OCA configurations, maintenance, and network routing	PEM, Open Connect Operations
Logistics representative	Directs the shipment of OCAs Your logistics representative should be involved at the beginning of the engagement process to avoid shipping delays at later stages of the process.	PEM, Open Connect Logistics, Open Connect Operations
Data center operations	Installs OCAs on the partner network	Open Connect Operations
Network operations	Works with Netflix to troubleshoot routing and other configuration issues that might arise	Open Connect Network Engineering, Open Connect Operations

Traffic analysis and sizing

After Netflix receives your OCA request, Netflix Open Connect PEMs will perform a traffic analysis on the autonomous system number (ASN) that you provided in the online appliance request form to determine whether your levels of Netflix traffic are currently (or projected to be) above the minimum threshold.

Assuming that you meet the minimum Netflix traffic requirements, additional capacity planning activities are initiated.

The Netflix Open Connect teams will work directly with your network planning and engineering teams to analyze traffic patterns within the various regions of your network, determine the appropriate type and number of appliance for each location, and determine the best resiliency options. The goal is to provide the best experience possible for your customers by architecting a design with maximum resilience and by serving Netflix content as close as possible to your end users.

The number of appliances that are deployed is determined by three general factors:

- The availability that is required at a given site
- The amount of Netflix traffic (bandwidth) that is delivered from that site
- The percentage of traffic that you want to originate from the site

The Netflix team will also work with you to map your current traffic volumes to our internal regional forecasts, to ensure that your OCA configuration will sustain the future peak of traffic in the coming year. After the initial deployment, appliance performance is monitored by the Netflix team and additional capacity planning is done periodically as needed.

Shipping, Receiving, and Installation

[SHIPPING OCAS >](#)

[RECEIVING OCAS >](#)

[RACKING AND INSTALLATION >](#)

[OCA RETURNS AND REPLACEMENTS >](#)

[NEXT STEPS >](#)

Netflix's logistics suppliers have different requirements for shipping, importing, and customs processes that are specific to each country in which an ISP might operate. During the engagement process with Netflix, the shipping specifics for your country will be reviewed with you and agreed upon.

Shipping OCAs

Before shipping OCAs to your facility, the Netflix team pre-configures the appliances and fills them with the latest content. To avoid having the pre-filled content become stale, thereby requiring an extended high-traffic initial fill period, it is highly recommended to get the appliances online as soon as possible. You will be asked to confirm site readiness before any OCAs are shipped to you.

To facilitate smooth delivery and installation, you will be directed to an online site survey form before shipment where you can specify the final configuration details of your appliances. In this form, you will provide information such as:

- Site addresses
- Shipment contacts
- IP configuration details (IPv4, IPv6, default gateway, and so on)
- BGP configuration details (IPv4, and IPv6 if applicable)
- Type of rack rails (2-post vs 4-post)

- Power supply types (AC vs DC, plug type)
- Optic types (LR-SM, SR-MM)

Each appliance is shipped with the rail kits and power supplies that you specify in the online form and fiber patch cables are provided upon request to connect the 10G interfaces from the OCA to your router.

Receiving OCAs

Requirements for receiving an appliance shipment are roughly as follows. Specific requirements will be communicated to your logistics team by the Netflix logistics team before the shipment.

Criteria	Requirement
Estimated appliance weight	100 to 125 pounds (45 to 57 kg)
Estimated weight of pallet and shipping materials	63 pounds (29 kg)
Loading dock and jack	Upon request, Netflix will ensure delivery trucks arrive with a pallet jack and lift gate if a loading dock and jack are not available at the destination.
Inside delivery	Inside delivery can be requested.
Consistent room temperatures	No higher than 78°F (26°C)
Maximum room temperatures	No higher than 104°F (40°C) for very short periods of time only

Racking and installation

Netflix ships each OCA to your requested location ready to deploy. As soon as you receive your OCA shipment, you are responsible for:

- Racking the appliance

- Providing the necessary 2-4 x 10Gbps optical network connections (exact quantity depends on the appliance type)

Reminder: As mentioned earlier, because the content that is pre-loaded onto the appliances can become stale quickly, you should complete the installation as soon as possible.

Airflow:

- For storage and global OCAs, optics face the cold aisle.
- For **flash-based appliances**, optics face the hot aisle.

If you are unsure about what type of OCA you are installing, contact [Netflix NOC](#) for assistance.

OCAs are either 2- or 4-post friendly. Rail kit instructions can be accessed online:

<http://openconnect.netflix.com/railkit-2-4.pdf> (<http://openconnect.netflix.com/railkit-2-4.pdf>)

If you are using DC power, follow these supplemental instructions:

<http://openconnect.netflix.com/dc-power-supply-adapter-kit.pdf>
(<http://openconnect.netflix.com/dc-power-supply-adapter-kit.pdf>)

Racking guidelines

- **Elevated Operating Ambient:** If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be

greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified by the manufacturer.

- **Reduced Air Flow:** Installation of the equipment in a rack should be such that the amount of airflow required for safe operation of the equipment is not compromised.
- **Mechanical Loading:** Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.
- **Circuit Overloading:** Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.
- **Reliable Earthing:** Reliable earthing of rack-mounted equipment should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power strips).

OCA returns and replacements

If an OCA is defective upon receipt, Netflix will send a Return Material Authorization (RMA) replacement appliance. For more information see: [Returns and replacements](https://openconnect.netflix.com/rma/) (https://openconnect.netflix.com/rma/)

In the event of a hardware failure, please contact [Netflix NOC](https://openconnect.netflix.com/contact) (https://openconnect.netflix.com/contact) for assistance.

Next Steps

After the initial installation, the next step is to establish a BGP session with the appliance, following the [network configuration instructions](#)

(<https://openconnect.netflix.com/network-configuration>). After the appliance is initially brought on-line, you notify the Open Connect Operations team by responding to the email notification that they send prior to each shipment. The Operations team will configure each OCA to receive fill traffic. Assuming that the appliance was installed within a [reasonable amount of time](#)

(<https://openconnect.netflix.com/requirements-for-deploying/#network-datacenter>) after receipt, the fill process generally takes about two days. If installation was delayed, the fill process can take longer.

The appliance will not start serving client traffic until both you and Netflix have agreed that the system is healthy and ready to serve, and when acceptable offload can be achieved. When the system is ready, the Netflix Open Connect operations team will enable the system to serve clients and will continue to monitor its performance.

Network Configuration

[OVERVIEW >](#)

[RECONFIGURING THE IP ADDRESS OF AN OCA >](#)

[ROUTER INTERFACE CONFIGURATION >](#)

[ROUTING AND CONTENT STEERING VIA BGP >](#)

[CLUSTERING ARCHITECTURES >](#)

[EMBEDDED OCAS COMBINED WITH PEERING SESSIONS >](#)

[FLASH-BASED APPLIANCES >](#)

This section describes how Open Connect Appliances are typically configured in a network. If you are an Open Connect ISP partner, Netflix works closely with you to determine the optimal configuration for your particular needs.

For more information, see the [FAQs](https://openconnect.netflix.com/faq/) (<https://openconnect.netflix.com/faq/>)

Overview

OCAs are *directed cache appliances*, meaning that the manner in which traffic is directed to the appliance is determined explicitly by you and by Netflix, not by the appliance itself.

An OCA only serves clients at IP addresses that you advertise to the OCA via a BGP session. In other words, traffic is only delivered from your embedded OCAs to the customer prefixes that you explicitly announce to them, as described in the following sections. Therefore, you as the ISP partner have full control over the networks that the appliances will serve. BGP sessions are established between appliance(s) and the closest connected router.

If content is requested that is not contained on an embedded OCA, the client request is directed to the closest Netflix content site via peering (if present) or via transit.

Reconfiguring the IP address of an OCA

Each appliance comes fully configured based on the IP address details that you provided to Netflix in your site survey before it was shipped.

For all appliances, the IP address can be updated via a keyboard and monitor. Interfaces are on the front of the chassis, but might be hidden behind a panel. The updated IP address will only take effect after a reboot, so it is important to drain the appliance two hours before the change by shutting down its BGP session (or sessions) to ensure that there is no traffic being served by the appliance.

For step by step instructions, see the article in our Partner Portal help center:

<https://openconnect.zendesk.com/hc/en-us/articles/360000921752>

(<https://openconnect.zendesk.com/hc/en-us/articles/360000921752>)

Router interface configuration

When you are connecting the appliances to your router, follow the guidelines in this section. See also [Example router configurations](#)

(<https://openconnect.netflix.com/sample-network/>).

- Each OCA requires one usable IPV4 address, and it is highly recommended to also assign one IPV6 address.
- You can assign an address to the appliance from an IPv4 subnet of /31 and larger, or an IPv6 subnet of /127 and larger.

- It is acceptable to assign an address to the appliance from a larger subnet (for example, a /24). However, because only one IPv4 address is required per appliance, a smaller subnet is typically used.
- The router interfaces must be configured for Link Aggregation Group (LAG) with LACP.
- A standard maximum transmission unit (MTU) must be configured on each router interface. Jumbo frames are not supported.
- If there are multiple routers available that can provide redundancy in a site, it is recommended to stagger appliances between routers. Appliances on the same router should be in the same subnet to optimize filling. Appliances on separate routers should be in separate subnets. Appliances are not designed to be connected to two separate routers.
- Each OCA is hardened against network attack and is designed to be directly connected to the internet. Filtering inbound or outbound traffic can cause operational issues, so we strongly recommend that you allow all traffic on all ports, do not use ACLs, and ensure that your router has a default route or full routing table. If you absolutely **must** filter, the current list of inbound and outbound usage follows. Please note that these can change at any time without prior notification.
 - **Traffic from OCA:** Allow all destination addresses and ports.
 - **Traffic to OCA:** Allow TCP 22, 53, 80, 179, 443, UDP 53 and 123 (source and destination), ICMP types 0, 3, 8, 11, and all ICMPv6 from any public IP/port. Allow all return traffic from any appliance-initiated connection (TCP established).
- Each network interface must be receiving between 0 dBm and -10 dBm of light to ensure good data throughput. The LCD panel on the front of the appliance displays the current light levels for each interface. If your appliance does not have an LCD panel, access the console with a keyboard and mouse,

then follow the console instructions to check light levels. If light levels are out of the acceptable range, clean the optics. If cleaning the optics does not bring them into the acceptable range, contact Netflix to have new optics shipped to you.

Routing and content steering via BGP

We steer clients to our OCAs based on an ISP's BGP advertisements, coupled with the routing and steering algorithms in the Open Connect control plane. ISP partners can control some aspects of content steering via the BGP routes that are announced to each OCA.

The control plane steers requests from end user clients to the best available appliance based on multiple factors. Assuming that the appliance has the requested title and available serving capacity, the control plane provides clients with a ranked list of appliances (typically 3 or more reliable sources) to stream from.

Appliance selection criteria

The following appliance selection criteria are considered, in order, by the Open Connect control plane services. If there is a tie for a given criterion, then the next criterion is considered. If there is a tie on all criteria, traffic is balanced between appliances.

1. The appliance that receives the most-specific route to the client's prefix.
2. The appliance that receives the route to the client's netblock with the shortest AS path.
3. The appliance that receives the route to the client's netblock with the lowest multi-exit discriminator (MED). (See the notes on MEDs below).

4. The geographically closest appliance. We geolocate based on client IPs, whose location is then compared to the latitude and longitude of nearby OCAs to determine the closest available system.

BGP notes

- Route announcements for Open Connect embedded appliances:
 - IPv4 prefixes between /8 and /31 (inclusive) are accepted.
 - IPv6 prefixes between /19 and /64 (inclusive) are accepted.
- Route announcements for Open Connect peering sessions:
 - IPv4 prefixes between /8 and /24 (inclusive) are accepted.
 - IPv6 prefixes between /19 and /48 (inclusive) are accepted.
- As an implicit requirement, **all** appliances must have a BGP session configured in order to correctly participate in Netflix content steering and delivery.
- Advertised routes that are received by an OCA are synchronized with Open Connect control plane services approximately every five minutes.
- To localize traffic, the best practice is to advertise the most specific routes to the appliance. For example, if you are announcing a /22 to the OCA, but a /24 is received from the same block over settlement-free interconnection (SFI) peering or transit, the /24 will be preferred, delivering content traffic from the remote source instead of the OCA.
- If you are deploying only one OCA in your network, you should advertise the most specific (shortest) prefix for that OCA over the peering session that you want the OCA to use for nightly filling purposes.
- If you are deploying multiple OCAs in your network:
 - To enable efficient nightly fill: If you have separate clusters that are located in two different sites, ensure that the appliances within one

cluster can hear the subnets for the appliances in the other cluster via the BGP connection that is established with your router. See the [Fill and updates](#) information for more details.

- See the additional information about [clustering architectures](#).
- MEDs can be injected and will be respected. You can use MEDs to de-preference routes between appliances or otherwise distribute traffic to meet your requirements. See additional notes below.
- Netflix does not use any BGP community information that is advertised by partners to OCAs or via Open Connect peering.

Additional notes on MEDs

- **Important:** Marking MEDs on already installed and working Open Connect Appliances can be hazardous, because it must be done on all BGP sessions for all appliances at the same time.
- There is no cap on the maximum MED value.
- A missing MED is treated the same as a MED of 0, and indicates that the appliance should receive all servable traffic for the associated prefixes (also often referred to as *MED-missing-as-best*). Remember, if multiple appliances receive the same prefix with the same metric, traffic is load-balanced across those appliances. Because a missing MED will be equivalent to 0, it is preferred over any >0 MED on other appliances.
- For information about MEDs with respect to peering only, see [Peering Locations](#).

Clustering architectures

Two or more OCAs that are intended to serve the same set of customers can be configured by the Open Connect operations team as a manifest cluster. OCAs in a

manifest cluster share content storage and function together as one logical server/storage unit. The Netflix team collaborates with you to determine whether clustering is warranted and how to set up optimal clusters, depending on your particular site and network configurations.

Although partners do not need to configure manifest clusters, it is important to understand some basic clustering concepts. In particular, there are implications to consider when OCAs in a cluster are taken down for maintenance or moved to a different site.

Clustering has the following potential benefits:

- **Greater offload for unique content**

In a typical two-OCA cluster, both appliances will use approximately 40% of their storage for the same popular content. This popular content typically represents roughly 60% of the OCA's total offload. The remaining 60% of storage space on each OCA is used to store a unique set of less-frequently-accessed content. Because we do not store the same exact set of content on each single OCA in a cluster, a cluster of OCAs provides greater total offload than an unclustered group of OCAs. This strategy helps the OCAs in a cluster function more efficiently.

For more information about storage strategies, see [this tech blog post](https://medium.com/netflix-techblog/distributing-content-to-open-connect-3e3e391d4dc9) (<https://medium.com/netflix-techblog/distributing-content-to-open-connect-3e3e391d4dc9>).

- **Better resiliency**

Redundancy is generally acceptable in a two-OCA cluster. In the event of a single OCA failure, the healthy appliance will take over the majority of the traffic that the failed unit was serving. See the failover scenarios in the [sample architectures](https://openconnect.netflix.com/sample-architectures/#failover) (https://openconnect.netflix.com/sample-architectures/#failover).

Notes for partners:

- OCAs intended to serve the same set of customers can be clustered if they are located at the same site, or if they are in close geographical or network proximity.
- OCAs cannot be clustered if they are not intended to serve the same set of customers.
- After a set of OCAs has been installed in a site and grouped together as a cluster by the Open Connect team, they should be thought of as one big server. Therefore, any changes you make to a single OCA in a site has the potential to negatively impact the serving efficiency and behavior of the group.
- If you need to make changes to the OCAs in an established site - for example, if you intend to relocate an OCA from one site to another or disable one or more OCAs for a significant period of time - it is important to notify the Open Connect team so that they can make the necessary changes to the cluster configuration. Failing to do so can cause undesired consequences. For example, you may see traffic being steered to the wrong site, fill patterns may become suboptimal, and hot spots might develop.
- To enable optimal and balanced traffic patterns, OCAs in a site and within a cluster must receive the exact same BGP route advertisements. Therefore, if you relocate an OCA you must revisit your BGP route announcements to ensure that traffic continues to be steered appropriately.

Embedded OCAs combined with peering sessions

The ideal Open Connect implementation is a mixture of both SFI peering and deployed embedded OCAs. Netflix uses two separate autonomous systems for peering:

- AS2906 is the AS number that Netflix uses for peering at its PoPs
- AS40027 is the AS number that embedded OCAs use to peer with ISP networks

See [BGP notes](#) for prefix announcements that are accepted on peering sessions.

The same prefix announced both to a private or public peering session (using AS2906) and to an OCA (using AS40027) will always be preferred on the appliance over peering, because the Open Connect control plane will have two BGP entries for that prefix:

- one with an AS PATH LENGTH of 1 (<AS_NUMBER>) from the appliance itself
- one with an AS PATH LENGTH of 2 (2906 <AS_NUMBER>) from an IX location

When OCAs and Open Connect SFI peering is combined, OCAs are nominal and peering is used primarily for backup, for filling, and for serving long-tail titles.

Flash-based appliances

If you are an ISP with very large amounts of Netflix traffic, we will likely include flash-based appliances in your OCA deployment architecture. Flash-based appliances are 1U flash memory-based servers that are deployed when you reach

a threshold number of OCAs, to augment the delivery capability of the main (storage) appliances.

Flash-based appliances have different fill window requirements, see [the fill information](https://openconnect.netflix.com/fill/) (https://openconnect.netflix.com/fill/) for more details.

Flash-based appliances are not clustered unless they are in the same site, and they should not be set up in the same manifest cluster as the main storage appliances.

Rail kit instructions for flash-based appliances can be accessed online:

<http://openconnect.netflix.com/railkit-1.pdf> (http://openconnect.netflix.com/railkit-1.pdf)

Note: Flash-based appliances, unlike storage and global appliances, have their ports on the back of the device facing the hot isle.

Example Router Configurations

[BROCADE LAG BGP >](#)

[CISCO IOS BGP >](#)

[CISCO IOS XR BGP >](#)

[CISCO IOS XR LACP >](#)

[CISCO IOS INTERFACE LACP >](#)

[CISCO IOS XR BRIDGED INTERFACE >](#)

[CISCO IOS SWITCH WITH VLAN >](#)

[JUNOS BGP >](#)

[JUNOS INTERFACE LACP >](#)

[JUNOS BRIDGED INTERFACE >](#)

Use the following examples to help you understand typical router configurations.

Brocade LAG BGP example configuration

```
lag "NETFLIX" dynamic
ports ethernet 1/7 to 1/8
primary-port 1/7
deploy
port-name "NETFLIX-OCA" ethernet 1/7
port-name "NETFLIX-OCA" ethernet 1/8
interface ethernet 1/7
enable
ip address <interface address/mask>
no ip redirect
ipv6 address <interface address/mask>
```

```
sflow forwarding
link-aggregate active
interface ethernet 1/8
enable
link-aggregate active
```

```
router bgp
  neighbor Netflix peer-group
  neighbor Netflix remote-as 40027
  neighbor NetflixIPv6 peer-group
  neighbor NetflixIPv6 remote-as 40027
  neighbor <IPv4 address of Appliance> peer-group Netflix
  neighbor <IPv4 address of Appliance> description "Netflix Ope
  neighbor <IPv6 address of Appliance> peer-group NetflixIPv6
  neighbor <IPv6 address of Appliance> description "Netflix Ope
```

```
address-family ipv4 unicast
  neighbor Netflix prefix-list no-advertise in
  neighbor Netflix prefix-list NETFLIX-LOCAL out
  no neighbor NetflixIPv6 activate
  exit-address-family
```

```
address-family ipv6 unicast
  neighbor NetflixIPv6 activate
  neighbor NetflixIPv6 prefix-list no-advertise-ipv6 in
  neighbor NetflixIPv6 prefix-list NETFLIX-LOCAL-IPV6 out
  exit-address-family
```


Cisco IOS BGP example configuration

```
!  
router BGP <ASN>  
neighbor <IP address of Appliance> remote-as 40027  
neighbor <IP address of Appliance> route-map NETFLIX-OUT out  
!  
route-map NETFLIX-OUT permit 10  
match ip address prefix-list NETFLIX-LOCAL  
set metric 0  
!  
route-map NETFLIX-OUT permit 20  
match ip address prefix-list NETFLIX-BACKUP  
set metric 100  
!  
route-map NETFLIX-OUT deny 30  
!
```

Cisco IOS XR BGP example configuration

```
!  
router bgp <ASN>  
neighbor <IP address of Appliance>  
remote-as 40027  
address-family ipv4 unicast  
route-policy NETFLIX-OUT out  
!
```

```
route-policy NETFLIX-OUT
  if destination in NETFLIX-LOCAL then
    set med 0
  elseif destination in NETFLIX-BACKUP then
    set med 100
  else
    drop
  endif
end-policy
!
prefix-set NETFLIX-LOCAL
  192.0.2.0/25,
  192.168.100.0/25
end-set
!
prefix-set NETFLIX-BACKUP
  192.0.2.128/25,
  192.168.100.128/25
end-set
!
```

Cisco IOS XR LACP example interface configuration

```
!
interface Bundle-Ether1
  bundle minimum-active bandwidth 1
```

```
ipv4 address <address/mask>
ipv6 address <really-long-address/mask>
!
interface TenGigE0/0/2/0
  bundle id 1 mode active
  lacp period short
!
interface TenGigE0/0/2/2
  bundle id 1 mode active
  lacp period short
!
```

Cisco IOS Interface LACP example configuration

```
!
interface Port-channel1
  ip address <address> <mask>
!
interface GigabitEthernet1/1
  channel-group 1 mode active
  channel-protocol lacp
!
interface GigabitEthernet1/2
  channel-group 1 mode active
  channel-protocol lacp
```

!

Cisco IOS XR Bridged Interface example configuration

For multiple servers with LAG, using the same default gateway

```
interface BVI100
  description NETFLIX
  ipv4 address <address/mask>
  ipv6 address <really-long-address/mask>
```

```
l2vpn
load-balancing flow src-dst-ip
bridge group NETFLIX-OCAS
  bridge-domain NETFLIX-OCAS
    interface Bundle-Ether1
    interface Bundle-Ether2
    routed interface BVI100
```

```
interface Bundle-Ether1
  l2transport
interface Bundle-Ether2
  l2transport
```

Cisco IOS Switch with VLAN example configuration

For multiple servers in LAG groups, behind a switch

```
! SVI Configuration
!
interface Vlan30
  description **NETFLIX OCA SERVERS**
  ip address <address/mask>

! Port channel configuration
!
interface Port-channel31
  description **LINK TO NETFLIX-OCA-1**
  switchport
  switchport mode access
  switchport access vlan 30
!
interface Port-channel32
  description **LINK TO NETFLIX-OCA-2**
  switchport
  switchport mode access
  switchport access vlan 30

! Interfaces configuration
!
! OCA1
interface TenGigabitEthernet1/1/1
  description **LINK TO NETFLIX-OCA1-TEN0**
  switchport
```

```
switchport mode access
switchport access vlan 30
channel-group 31 mode active
!
interface TenGigabitEthernet1/1/2
description **LINK TO NETFLIX-OCA1-TEN1**
switchport
switchport mode access
switchport access vlan 30
channel-group 31 mode active

! OCA2
interface TenGigabitEthernet2/1/1
description **LINK TO NETFLIX-OCA2-TEN0**
switchport
switchport mode access
switchport access vlan 30
channel-group 32 mode active
!
interface TenGigabitEthernet2/1/2
description **LINK TO NETFLIX-OCA2-TEN1**
switchport
switchport mode access
switchport access vlan 30
channel-group 32 mode active
```

Junos BGP example configuration

```
protocols {
  bgp {
    group NETFLIX {
      neighbor <IP address of Appliance> {
        export NETFLIX-EXPORT;
        peer-as 40027;
      }
    }
  }
}
```

```
policy-options {
  policy-statement NETFLIX-EXPORT {
    term NETFLIX-LOCAL {
      from {
        prefix-list NETFLIX-LOCAL;
      }
      then {
        metric 0;
        accept;
      }
    }
    term NETFLIX-BACKUP {
      from {
        prefix-list NETFLIX-BACKUP;
      }
      then {
        metric 100;
      }
    }
  }
}
```

```
        accept;
    }
}
then reject;
}
}
```

Junos Interface LACP example configuration

```
interfaces {
    ae1 {
        description NETFLIX-OCA;
        aggregated-ether-options {
            lacp {
                active;
                periodic fast;
            }
        }
    }
    unit 0 {
        family inet {
            address <address/mask>;
        }
        family inet6 {
            address <really-long-address/mask>;
        }
    }
}
```



```
xe-0/0/0 {
    description NETFLIX-OCA-cxgbe0;
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/1 {
    description NETFLIX-OCA-cxgbe1;
    ether-options {
        802.3ad ae1;
    }
}
}
```

Junos Bridged Interface example configuration

For multiple servers with LAG, using the same default gateway

```
vllans {
    NETFLIX-OCAS {
        vlan-id 100;
        l3-interface irb.100;
    }
}
interfaces {
    irb {
        unit 100 {
```

```
    family inet {
        address <address/mask>;
    }
    family inet6 {
        address <really-long-address/mask>;
    }
}
ae0 {
    description NETFLIX-OCA-1;
    unit 0 {
        family ethernet-switching {
            vlan {
                members NETFLIX-OCAS;
            }
        }
    }
}
ae1 {
    description NETFLIX-OCA-2;
    unit 0 {
        family ethernet-switching {
            vlan {
                members NETFLIX-OCAS;
            }
        }
    }
}
xe-0/0/0 {
```

```
        description NETFLIX-OCA-1-PORT1;
        ether-options {
            802.3ad ae0;
        }
    }
xe-0/0/1 {
    description NETFLIX-OCA-1-PORT2;
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/2 {
    description NETFLIX-OCA-2-PORT1;
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/3 {
    description NETFLIX-OCA-2-PORT2;
    ether-options {
        802.3ad ae1;
    }
}
}
```

Fill, Updates, and Maintenance

[FILL WINDOWS >](#)

[FILL PATTERNS >](#)

[PEER FILLING >](#)

[TIER FILLING >](#)

[OUT-OF-CYCLE FILLS >](#)

[AUTOMATIC SOFTWARE UPDATES >](#)

[HEALTH MONITORING >](#)

[UNPLANNED OUTAGES OR SERVICE-IMPACTING EVENTS >](#)

[POTENTIAL SERVICE-IMPACTING MAINTENANCE >](#)

[PARTNER MAINTENANCE >](#)

Each Open Connect Appliance (OCA) stores a portion of the Netflix catalog, which in general is less than the complete content library for a given region. Popularity changes, new titles that are added to the service, re-encoded movies, and routine software enhancements are all part of the nightly updates, or fill, that each appliance must download to remain current.

Fill windows

By design, OCAs follow a “push fill” methodology. They fill every day during a window of time that corresponds to your off-peak hours. The timing of the fill window is defined in partnership with your network planning team.

The goal is to set the fill window such that:

- It occurs during the trough of your Netflix traffic
- It does not disrupt your inbound traffic peaks

Setting the fill window this way allows the system to continue to download updates and serve clients without negatively impacting playback performance.

Requirements for the fill window, by hardware type:

Setting	Storage	Global	Flash
Length of time	12 hours	12 hours	3 hours
Timeframe	2AM to 2PM (local time) (can be shifted +/- 2 hours)	2AM to 2PM (local time) (can be shifted +/- 2 hours)	3AM to 6AM (local time) or 6AM to 9AM (local time)
Bandwidth	Maximum rate of 1.2 Gbps per system	Maximum rate of 0.4 Gbps per system	Maximum rate of 1.2 Gbps per system
Total traffic during fill window	Up to 6.5 TB	Up to 2.1 TB	Up to 1.5 TB

The system might complete filling before the end of the fill window, and it might fill more than once if new content becomes available before the fill window ends.

If flash-based appliances are part of your network, they will be configured to fill as much as possible from the storage appliances that are on site. Because they are SSD-based, they are configured with a shorter staggered fill window, and they will not serve while they are filling. Because the flash-based appliances do not serve while they are filling, all Netflix traffic is served by the storage appliances on site during their fill period.

Fill patterns

In general, appliances determine where to receive fill using selection criteria that is similar to [the steering criteria that is used by Netflix client devices](#)

(<https://openconnect.netflix.com/network-configuration/#routing-steering>). The following scenarios provide more detail.

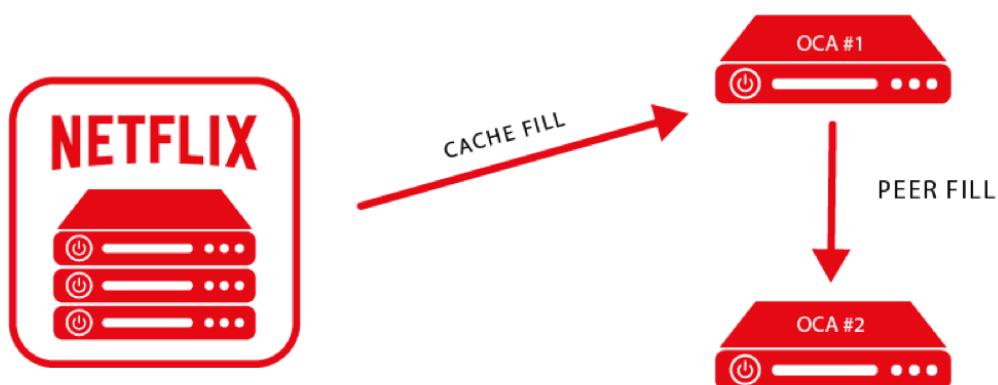
If you have specific fill requirements, discuss them with your partner engagement manager during the initial site design process. If you require changes to your appliance's fill behavior post-deployment, [contact the Open Connect operations team](https://openconnect.netflix.com/contact/). (<https://openconnect.netflix.com/contact/>)

Peer filling

OCA's will attempt to peer fill from each other in either of these two cases:

- If they are clustered
- If they are on the same subnet

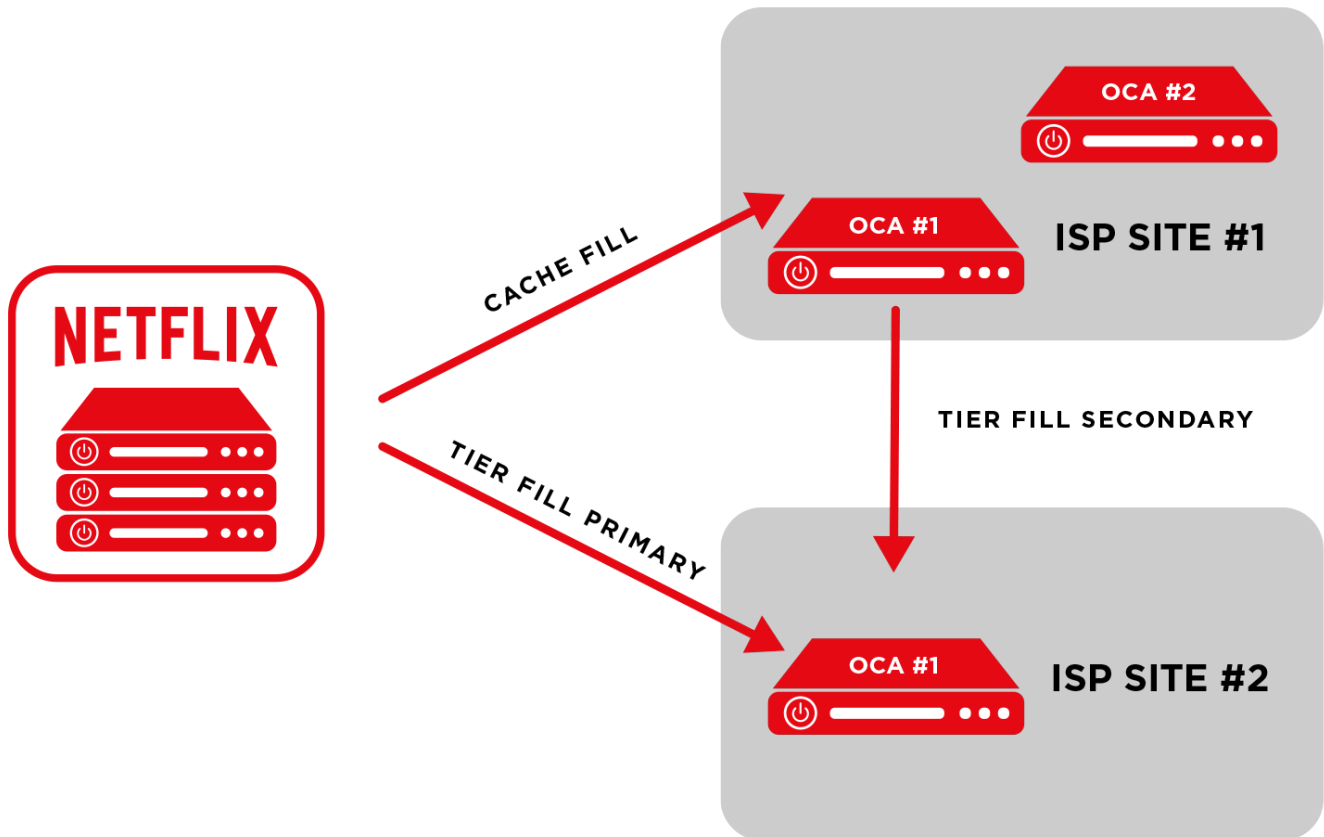
For example, assume that in the following diagram, OCA #1 and OCA #2 are in a cluster. If OCA #1 downloads a movie via a regular cache fill, OCA #2 will then seek to download the movie from OCA #1 instead of via transit or Open Connect peering. Routes to the OCA's must be advertised over Open Connect peering to ensure that the peering session is used for fill traffic.



If OCAs on multiple sites are configured within the same cluster, there is a strong chance that they will not be numbered within the same subnet. However, even if the cluster spans multiple sites, the “same-cluster” rule still applies - therefore the OCAs will attempt to peer fill as described.

Tier filling

Peer filling is the most efficient method, however as a second-best option tier filling is considered. OCAs that can see each other’s IP address in their BGP feed (but are not in the same cluster or subnet) will consider filling from each other rather than via a regular cache fill. To enable tier filling if you have separate clusters that are located in two different sites, ensure that the appliances within one cluster can hear the subnets from the other cluster via the BGP connection that is established with your router.



Out-of-cycle fills

There will be cases that cause an urgent need to deploy or update a small subset of titles on the OCAs outside of the normal fill window. This need can arise for various reasons, for example:

- If re-encoding of a popular title is required to repair viewing issues
- To make a new, “day-of-broadcast”, or “day-after-broadcast” title available

When these out-of-cycle fills occur, you will see a small portion of fill traffic for a limited period of time outside of the configured fill window. Doing an out-of-cycle fill instead of waiting until the normal fill window helps to ensure that popular titles can continue to be served from your OCAs (the preferred method) and not via peering or transit. This approach improves the overall efficiency and offload

potential of the system. Because out-of-cycle fills are limited to very specific popular titles, they do not negatively impact overall server efficiency - rather, they increase the local offload of popular content.

Automatic software updates

Netflix manages the OCA software, and we will periodically update or enhance the software. All non-emergency software updates are performed during the configured fill window, and you will be notified before and after they occur. All appliance software, including the operating system, web server, and control code, is updated as a single image. If a new image fails to properly load, the appliance will automatically fall back to the previous image. In either case, you will receive an email to let you know the status of the upgrade after it is attempted.

Health monitoring

As part of the Open Connect program, Netflix continuously monitors all deployed Netflix OCAs, including all aspects of the performance and availability of the appliance. In addition, you can monitor the basic health of the OCAs that are deployed on your network by using the Netflix Partner Portal at <https://my.oc.netflix.com> (<https://my.oc.netflix.com>). You will be invited to create an account in the Partner Portal as part of the OCA installation and turn-up process.

Unplanned Outages or Service-Impacting Events

For service-wide outages with broad impact to the service of multiple OCA partners, partners might receive an initial notification, status updates on a regular basis if the outage is prolonged, and a final resolution notification. If you want to filter these announcements, the sender address is: cdn-noc@mail.netflix.com.

For per-ISP partner incidents, including routing and connectivity anomalies that persist over a significant period of time, the Netflix operations team will be alerted by our internal monitoring tools. We will resolve any issues under our control as soon as reasonably possible. Depending on the severity and impact to traffic, we might ask relevant ISP partners for additional comments on the incident.

Potential Service-impacting Maintenance

Whenever possible, for any planned maintenance or activity that might adversely impact an OCA or the services that it provides, Netflix will communicate according to the following schedule:

1. Netflix will send a “Scheduled Maintenance” email notification to the partner NOC contact e-mail address in our system two (2) business days prior to the activity.
2. Netflix will send a “Confirm Maintenance” email two hours (02:00) before the start of the planned maintenance activity, to confirm that the maintenance will be performed as expected.
3. Netflix will send a “Maintenance Completed” email when the maintenance activity is completed, to signal the return to a normal operational state.

For emergency service activities that do not permit our standard advance period, we will notify as early as possible.

During scheduled and emergency Netflix maintenance windows, ISP partners might see their BGP session flap, notice a link-state change with their cache-connected interfaces, and see traffic served from the appliance decrease or go to zero. All of these events are expected.

Under normal conditions, software updates that are not service-impacting will occur during the fill window at a time of low traffic, thereby minimizing potential user impact.

Partner Maintenance

The only OCA components that Netflix will request partners to replace in the field are:

- Power supplies
- SFP+ optics

We will ship these components to you in case of failure. Netflix will not request any ISP partner to install or replace any internal OCA hardware components such as memory, hard drives, or fans. Note that OCAs are designed to operate in data center environments with high-capacity fans that are required to cool the dense disk and memory installed. You might notice loud fan noise from OCAs, which is normal.

For ISP partners who are engaging in maintenance that might impact Netflix OCA availability, please observe the following guidelines:

- Plan your maintenance or downtime to occur during the cache fill window in order to minimize service disruption to our shared subscribers.
- If your maintenance has the potential to impact Netflix customer service, notify the Netflix team by [opening a maintenance ticket](#) via the Partner Portal.
- Plan to drain appliances of traffic 2 to 3 hours in advance of the maintenance. This will serve to direct any new Netflix viewing activity to an alternate source, and sessions that are currently associated with the appliance should finish by

the start of the maintenance period. You can disable and drain appliances from the Open Connect Partner Portal following these instructions: [Draining and disabling appliances](#)

- After completion of maintenance, please ensure that the OCAs are re-enabled to serve traffic. This will redirect new session starts away from the alternate sources.
- If you intend to make any changes to a site that contains multiple OCAs in a cluster, you must notify the Open Connect team so that the cluster can be reconfigured properly. For more details, see the information about [clustering architectures](#).

See the [Operating Level Agreement](https://openconnect.netflix.com/ola.pdf) (<https://openconnect.netflix.com/ola.pdf>) for a PDF version of the maintenance information in this section.

Returns and replacements

[APPLIANCE RMA GUIDELINES >](#)

[PARTS RETURNS AND REPLACEMENTS >](#)

Appliance RMA guidelines

OCA's are designed for resiliency, including the ability to have multiple hard drives fail without a noticeable impact on delivery. At times, however, hardware performance might degrade to a point where replacement of the entire OCA is necessary to maintain the ability to offload traffic in a meaningful manner.

In general, RMA requests are initiated by the Netflix NOC team, who will reach out to ISP partners via email and request for them to fill out a short online form to start the RMA process.

Our standard process is to ship a replacement pre-loaded OCA within three business days. The partner returns the impacted hardware in the same shipping box, and Netflix pays shipping both ways, at no shipping cost to the partner. Instructions for how to return a defective appliance are included in the shipping box when you receive the replacement appliance.

In some instances, we might need additional time in order to fully pre-load content for specific geographic regions. The total time until shipment from our US facility is likely to be 10 US business days at most, however customs clearance can vary according to region and country and might delay final receipt.

Parts returns and replacements

As discussed in the [maintenance](https://openconnect.netflix.com/fill/#maintenance) section, we may ask partners to replace PSU or optics on an appliance. In such cases we will ship the replacement part to the partner, and we might also request that they return the defective part.

To return a defective part, follow these instructions:

1. Keep the packaging box for the RMA return. Please reuse the same packaging material.
2. Remove or black out old shipping label including the barcode(s).
3. Place the return shipping label on the most visible part of the packaging box.
4. Ship your package using one of the following methods:
 - Via Fedex:
 - Use your regular scheduled Fedex pickup.
 - Drop off at FedEx. Find the closest location at <https://www.fedex.com/locate>.
 - Schedule a pickup online. Go to <https://www.fedex.com/grd/rpp/>.
 - If you do not have a Fedex Center close by:
 - Contact Open Connect Logistics. Please send an email to CDNL@netflix.com requesting the pickup. Include the following information:
 - Pickup address
 - Contact name and phone number
 - Netflix NCO ticket number

Contact Us

Prospective partner inquiries

Peering inquiries: Contact the Open Connect team at peering@netflix.com (<https://openconnect.netflix.com>).

Embedded appliance inquiries: Submit an [online Open Connect Appliance request](https://openconnect.netflix.com/request/) (<https://openconnect.netflix.com/request/>), and a Netflix Open Connect Partner Engagement Manager (PEM) will contact you.

General questions about Open Connect appliances

Contact the Open Connect Appliance team at openconnectappliance@netflix.com (<https://openconnect.netflix.com>)

Current partners with embedded Open Connect appliances

For most non-urgent requests or questions, open a ticket via the Partner Portal following these instructions: [Managing support tickets](https://openconnect.zendesk.com/hc/articles/360000911791) (<https://openconnect.zendesk.com/hc/articles/360000911791>)

Netflix Network Operations Center (NOC) - For emergency support only

E-mail: cdn-noc@netflix.com

Availability: 7x24x365

Netflix Open Connect Operations

E-mail: cdn-ops@netflix.com

Availability: M-F 9AM – 5PM Pacific Time