

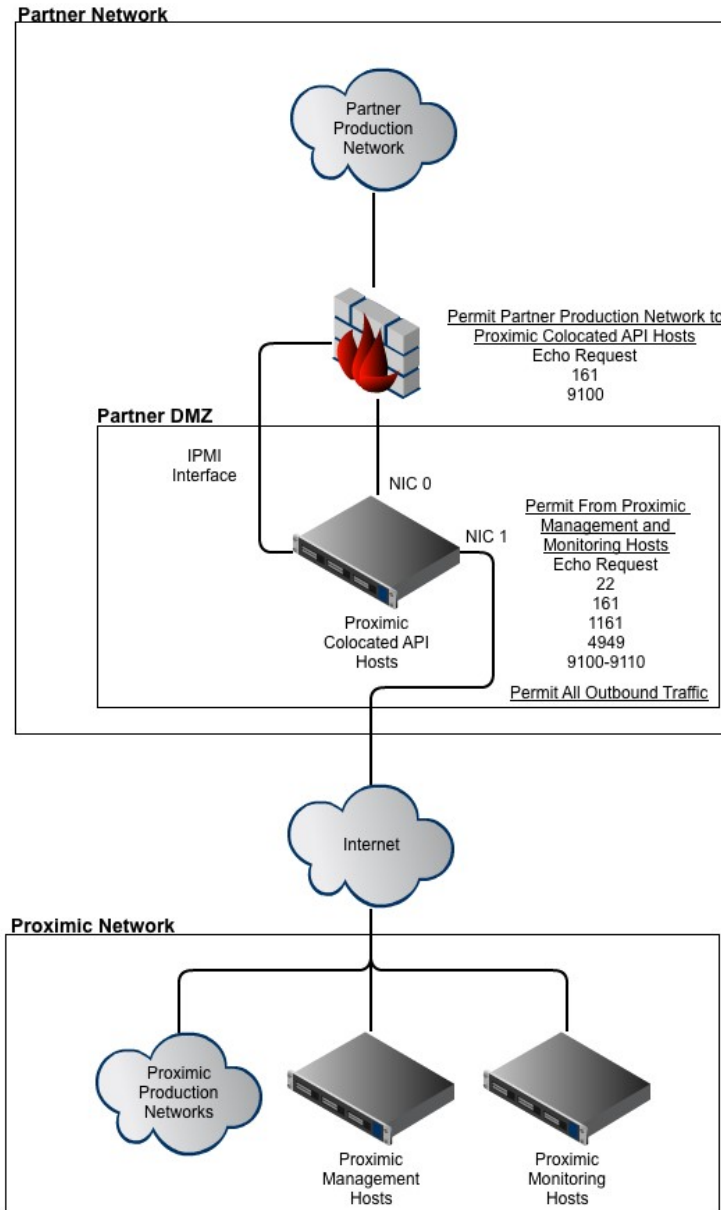
Colocated API

For customers with high volume output requirements (>1,000 requests per second) a colocated caching / API layer is highly recommended. The caching layer contains full data sets of the customer's inventory while the API layer provides access to the data and functionality. In a colocated setup the instance is queried with every impression.

The API returns responses at RTB compliant performance within 10ms (usually between 1-5ms). Using bare metal with 16 cores (E5620) and 96 GB RAM a single tier runs at about 75,000 requests per second. Virtualized environments run at about 20,000 requests per second depending on the provider / setup. Multiple instances can be used to cover the required throughput and provide additional reliability.

Best Practice

The following diagram shows how the colocated caching/API tier is implemented within the partner's infrastructure and how it communicates with Comscore's processing backend.



Comscore sets up, operates, and maintains a caching / API layer on the 3rd party hardware. The 3rd party hardware and network configuration is supposed to follow these requirements. Variations are most likely possible.

Partner will apply load balancing services in case more than one colocated instances is required. The number of nodes and type of load balancing depends on traffic volume and characteristics of partner.

To learn more about this server requirements please refer to the *Hardware and Network Requirements* section below.

To learn more about the functionality of the API please refer to [API Basics](#).

Hardware and Network Requirements

[System Requirements](#)
[Amazon Web Services](#)
[Google Cloud Platform](#)
[Operating System](#)
[NetworkingConfiguration Management Initial Access](#)
[Load Balancers](#)
[Client Connections](#)
[VLANs and Security Groups](#)
[Administrative Ports](#)
[Application Ports](#)

System Requirements

- **RAM:** Minimum: 32 GiB, recommended: at least 48 GiB for higher volumes. May be slightly less if required.
- **Disk:** Minimum: 100 GiB, at least 75GiB of which should be unallocated or available as /data (may be the root partition)
- **NIC:** Minimum: 1 Gb/s
- **CPU:** Minimum: 4 cores, recommended: 8 cores
- **OS:** Ubuntu 16.04 LTS

Amazon Web Services

- Instance Type
 - For low volume we recommend r5.xlarge instances.
 - For high volume we recommend r5.2xlarge instances.
 - r5d, r4, and r3 versions of these instances are also acceptable.
 - <http://aws.amazon.com/ec2/instance-types/>
- Disk Space
 - r3 or r5d require no additional storage
 - r4 or r5 (no local disk) require a 100GB EBS partition attached using any valid device name
- AMI
 - We use Ubuntu 16.04 LTS AMIs, a list of the current AMIs may be found at <https://cloud-images.ubuntu.com/locator/ec2/>
 - The following customized AMIs may also be used:
 - r3 instances:
 - ap-southeast-1: ami-0ddeaea373e1c04d3
 - ap-southeast-2: ami-09f200421b17d60ba
 - eu-west-1: ami-0667edf63084992b4
 - us-east-1: ami-09a6428d3133a8dad
 - us-west-1: ami-029973014e2238acb
 - r4, r5, and r5d instances:
 - ap-southeast-1: ami-0a2db599a6397fcd5
 - ap-southeast-2: ami-026f2b227386860ed
 - eu-west-1: ami-08beeeade3b8a94a6
 - us-east-1: ami-02826d25a445313b7
 - us-west-1: ami-0ba5825aa3265b8c3

To get started please provide your AWS account number and Comscore will make the images available.

These AMIs are subject to change **without notice**. Every time an instance is booted it is highly recommend to check whether it is performing as expected. The latest stock Ubuntu 16.04 LTS AMIs are always acceptable.

Google Cloud Platform

- Instance Type
 - For low volume we recommend n1-highmem-4 instances
 - For high volume we recommend n1-highmem-8 instances
 - <https://cloud.google.com/compute/docs/machine-types>
- Disk Space
 - 100GB minimum must be provided; this may be either the root partition or a persistent disk with at least 75GB.
- OS Image
 - We use Ubuntu 16.04 LTS images; the GCE Image family is ubuntu-1604-lts
 - <https://cloud.google.com/compute/docs/images#os-compute-support>

Operating System

Ubuntu 16.04 LTS

Configuration Management

In order to provide the best service possible, comScore needs to be able to have full control over the machine and integrate it into its configuration management and monitoring systems.

This means that root access to the machine is required to install puppet and snmpd for control. If required the partners' community string can be used to monitor the machine via SNMP.

Initial Access

In order to connect to the servers initially, the attached ssh public keys must be added to the ubuntu user's authorized key file (/home/ubuntu/.ssh/authorized_keys); no user accounts should be predefined. They may also be configured for the root user if remote root access is enabled initially (/root/.ssh/authorized_keys)

Attachment for /home/ubuntu/.ssh/authorized_keys

Please verify that the **permissions** of the folder and file are correct and there is no line wrapping in the authorized keys.

```
ls -ld /home/ubuntu/.ssh
drwx----- 12 ubuntu group 384 Apr 27 2019 /home/ubuntu/.ssh

ls -l /home/ubuntu/.ssh/authorized_keys
-r----- 1 ubuntu group 401 Oct 2 2018 authorized_keys

wc -l /home/ubuntu/.ssh/authorized_keys
4
```

Networking

Load Balancers

Instances may, but do not need to be, placed behind a load balancer.

Client Connections

1,200 max total number of concurrent connections with **keep-alive enabled** per caching/API layer.

VLANs and Security Groups

It is strongly recommended to put machines used for this service into separate VLANs (or security groups if using cloud services such as AWS) which have no outbound access to other machines on the partner network.

Administrative Ports

Inbound to the host:

Port	Protocol	Description	Allow IP addresses
Echo Request	ICMP	Ping	107.20.163.99, 174.129.234.205, 206.121.247.54, 107.20.167.30, 54.84.180.254
22	TCP	SSH	107.20.163.99, 174.129.234.205, 206.121.247.54, 107.20.167.30, 54.84.180.254
161 and 1161	TCP and UDP	SNMP	107.20.163.99, 174.129.234.205, 206.121.247.54, 107.20.167.30, 54.84.180.254
4949	TCP	Munin	107.20.163.99, 174.129.234.205, 206.121.247.54, 107.20.167.30, 54.84.180.254

Outbound from the host:

Port	Protocol	Description	Comments
53	TCP and UDP	DNS	
80	TCP	HTTP	Package updates
123	UDP	NTP	Time Synchronization
443	TCP	HTTPS	Package updates
514	TCP and UDP	Syslog	
2003	TCP and UDP	Graphite	Graphing
5672	TCP	MQ	Monitoring
8140	TCP	Puppet	Configuration management

Application Ports

Inbound to the host:

Port	Protocol	Description	Comments
9110	TCP	HTTP	Connection to upstream doorman
9106	TCP	HTTP	Log traffic
9100	TCP	HTTP	Connection for requests to proxy doorman

Outbound from the host:

Port	Protocol	Description	Comments
9110	TCP	HTTP	Connection to upstream doorman
9106	TCP	HTTP	Log traffic
9100	TCP	HTTP	Connection for requests to proxy doorman

If you are using AWS and only using inbound ACLs, the following script will configure your group using the AWS CLI:

```

#!/bin/bash
set -u
if [[ $# != 2 ]]
then
  cat <<EOF
  USAGE: $0 proximic_instances_security_group comma_separated_list_of_internal_groups
  e.g.
  # sg-00ab12cd: contains all colocated API hosts
  # sg-99zz88bb,sg-zz00yy11: All hosts which will need to connect to the colocated APIs are in these groups
  $0 sg-00ab12cd sg-99zz88bb,sg-zz00yy11
  EOF
  exit 1
fi
security_group=$1
internal_groups=$(echo $2 | tr ',' ' ')
for ip in 107.20.163.99 174.129.234.205 206.121.247.54 107.20.167.30 54.84.180.254
do
  for port in 22 161 1161 4949 9100-9110
  do
    aws ec2 authorize-security-group-ingress --group-id "${security_group}" --protocol tcp --port "${port}" --
cidr "${ip}/32"
  done
  for port in 161 1161
  do
    aws ec2 authorize-security-group-ingress --group-id "${security_group}" --protocol udp --port "${port}" --
cidr "${ip}/32"
  done
  aws ec2 authorize-security-group-ingress --group-id "${security_group}" --protocol icmp --port 8 --cidr "${ip}
/32"
done
for group in ${internal_groups}
do
  aws ec2 authorize-security-group-ingress --group-id "${security_group}" --protocol tcp --port 9100 --
sourcegroup "${group}"
done

```

Verification

Please send the output of the following lines:

```

curl -s -v https://ipinfo.io/json
curl -s -v http://apt.prod.proximic.com/apt.gpg

```