

## Factsheet

### Tests – Gaming

# Gaming

## Steam

Supported clients: Whiteboxes, Routers

The Steam client measures download speed of a game from Steam's game distribution network.

Measurement begins by fetching the manifest file for a freely available and publicly accessible game hosted by Steam. The manifest file provides information on the different "cells" (regions) that host the content and the chunks that are available in those cells (Steam files are split into chunks).

The client carries out the measurement phase by establishing one or more concurrent TCP connections to the user's local cell that hosts the chunks. These chunks are downloaded repeatedly for a configurable duration (defaulting to 10 seconds).

The client fetches the manifest itself, without the use of an offload server, so any geographic or ISP-specific logic implemented by Steam will equally affect our measurement client.

The client captures the following metrics:

1. The cell that was used
2. The average DNS lookup time
3. The average TCP connection time
4. The average download speed across all concurrent TCP connections

## DOTA 2

Supported clients: Whiteboxes, Routers

The DOTA2 measurement client measures round-trip latency and packet loss to the nearest DOTA2 game server. DOTA2 divides game servers into regions, with multiple game

servers (and relays on top of that) serving gameplay traffic amongst peers. Valve, the producer of DOTA2, publish a list of game servers and relays at [http://cdn.dota2.com/apps/sdr/network\\_config.json](http://cdn.dota2.com/apps/sdr/network_config.json), which is updated frequently. The game will periodically refresh its local cached copy of this server list.

The SamKnows DOTA2 measurement client performs its measurement in two phases. The first phase carries out a brief latency measurement to every DOTA2 server globally, in order to find the server with the lowest latency to the client. The second phase then carries out a fuller latency measurement to the server found in phase one. By default, only the results for the nearest server are reported.

The test uses ten ICMP echo packets to carry out its measurement by default. The test can also record the number of hops to the server.

The measurement client records the following:

- The IP address of the server.
- The region and datacenter name, if available, of the server.
- The average, minimum and maximum and standard deviation of round-trip time.
- The number of packets sent and received, allowing packet loss to be calculated.

## FIFA 2021

Supported clients: Whiteboxes, Routers

The FIFA 2021 measurement client measures round-trip latency and packet loss to the EA (Electronic Arts) game servers that are used

during gameplay. EA makes use of a variety of cloud hosting providers to host their game servers, including AWS.

The FIFA 2021 measurement client performs its measurement in two phases. The first phase carries out a brief latency measurement to every FIFA 2021 server globally, in order to find the server with the lowest latency to the client. The second phase then carries out a fuller latency measurement to the server found in phase one. By default, only the results for the nearest server are reported.

The test uses ten ICMP echo packets to carry out its measurement by default. The test can also record the number of hops to the server.

The measurement client records the following:

- The IP address of the server.
- The region and datacenter name, if available, of the server.
- The average, minimum and maximum and standard deviation of round-trip time.
- The number of packets sent and received, allowing packet loss to be calculated.

### League of Legends

Supported clients: Whiteboxes, Routers

The League of Legends measurement client measures round-trip latency and packet loss to the Riot network infrastructure, which serves gameplay traffic for League of Legends. Riot operates the following distinct regions that users can play in:

- North America
- Europe West
- Europe Nordic & East
- Oceania
- Russia
- Turkey
- Brazil
- Latin America North
- Latin America South

- Japan

In each of these regions, Riot operates a POP (Point of Presence) where traffic is exchanged with local ISPs and the game servers are hosted. These game servers utilise a modified version of the ENET protocol (based upon UDP) with supplemental encryption.

Whilst the game servers themselves cannot be measured reliably (due to the safeguards Riot employs to prevent cheating), the network infrastructure directly in front of the game servers can be measured.

The League of Legends measurement client performs its measurement in two phases. The first phase carries out a brief latency measurement to every League of Legends server globally, in order to find the server with the lowest latency to the client. The second phase then carries out a fuller latency measurement to the server found in phase one. By default, only the results for the nearest server are reported.

The test uses ten ICMP echo packets to carry out its measurement by default. The test can also record the number of hops to the server.

The measurement client records the following:

- The IP address of the server.
- The region and datacenter name, if available, of the server.
- The average, minimum and maximum and standard deviation of round-trip time.
- The number of packets sent and received, allowing packet loss to be calculated.

### Fortnite

Supported clients: Whiteboxes, Routers

The SamKnows Fortnite client measures round-trip latency and packet loss to the Epic Games infrastructure that hosts the Fortnite gameplay servers. These servers are segregated by region. The real Fortnite game

will automatically choose the region to exchange gameplay traffic with, whilst the SamKnows client will carry out measurements to all regions. Fortnite currently relies on Amazon's AWS for most of its gameplay infrastructure. It also makes use of some proxying infrastructure - most notably that of Subspace.com - to overcome regional peering issues.

Fortnite divides infrastructure into the following region definitions:

- North America - East
- North America - West
- Europe
- Oceania
- Brazil
- Asia
- Middle East

The Fortnite measurement client performs its measurement in two phases. The first phase carries out a brief latency measurement to every Fortnite server globally, in order to find the server with the lowest latency to the client. The second phase then carries out a fuller latency measurement to the server found in phase one. By default, only the results for the nearest server are reported.

The test uses ten UDP packets to carry out its measurement by default. The test can also record the number of hops to the server.

The measurement client records the following:

- The IP address of the server.
- The region and datacenter name, if available, of the server.
- The average, minimum and maximum and standard deviation of round-trip time.
- The number of packets sent and received, allowing packet loss to be calculated.

### Apex Legends

Supported clients: Whiteboxes, Routers

SamKnows has developed a measurement client that assesses round-trip latency and packet loss to the Apex Legends gameplay infrastructure. Like other gaming providers, Apex Legends segregates its infrastructure by region. Apex Legends hosts its infrastructure with multiple cloud platforms (Amazon AWS and Google GCE) and colocation providers.

Apex Legends divides infrastructure into the following region definitions:

- North America - East
- North America - Central
- North America - West
- Brazil - South
- Europe - West
- Asia
- South East Asia
- Japan
- Korea
- Australia

Each region is often served by multiple datacenters from different hosting/colocation providers. For example, "North America - West" is served by AWS, GCE and also by a third-party colocation operator.

The Apex Legends measurement client performs its measurement in two phases. The first phase carries out a brief latency measurement to every Apex Legends server globally, in order to find the server with the lowest latency to the client. The second phase then carries out a fuller latency measurement to the server found in phase one. By default, only the results for the nearest server are reported.

The test uses ten ICMP echo packets to carry out its measurement by default. The test can also record the number of hops to the server.

The measurement client records the following:

- The IP address of the server.
- The region and datacenter name, if available, of the server.

- The average, minimum and maximum and standard deviation of round-trip time.
- The number of packets sent and received, allowing packet loss to be calculated.