

# Methodology

## Methodology

### Off-net and on-net test servers

It is important that measurements collected by the test architecture support the comparison of ISP performance in an unbiased manner. To this end, SamKnows maintains a fleet of "off-net" test servers, deployed at major internet exchanges and peering locations globally. Test servers are carefully sited to ensure optimal connectivity on a market-by-market basis.

Off-net test servers are located outside of the bounds of any one ISP's network, and seek to provide a neutral and consistent measurement target. Measurements made against off-net servers represent "real world" end-user experience, and off-net servers can also help bring to light peering problems or routing issues.

Regulator reports use data collected from measurements against off-net servers, and SamKnows recommends that all customers use off-net servers to ensure that they are getting a realistic overview of their customers' experience. Testing to multiple servers in a variety of locations will help build a picture of network health beyond your immediate borders, allowing you to view local, national, and international performance to better understand the service you are providing to your customers.

Test servers may also be deployed inside an ISP's network ("on-net" test servers). This allows them to segment end-to-end network performance and determine the performance of their own network versus third party networks. For example, an ISP can see what impact third party networks have on their end-users Quality of Experience ('QoE') by

placing test servers within their own network and at major national and international peering locations.

### Test server selection

The SamKnows measurement agents (Whiteboxes, routers, smartphone apps and browser-based tests) select the nearest test server by running round-trip latency checks to all nearby test servers before measurement begins. Note that the term "nearest" refers to the test server nearest to the agent from the point of view of network delay, which may not necessarily always be the one nearest geographically.

Alternatively, it is possible to override test server selection based on latency and implement a static configuration so that the agent will only test against test server chosen by the administrator. This is so that the administrator can choose to test to any test server that is of interest to the specific project and to maintain configuration consistency.

### Data stored on test servers

No measurement data is stored on test servers. The test servers provide a 'dumb' endpoint for the measurement agents to test against. All measurement performance results are recorded by the measurement agents, which are then transmitted from the agent to data collection servers managed by SamKnows.

### Test server management and monitoring

SamKnows controls its fleet of test servers via a popular open-source management tool called Puppet (<https://puppetlabs.com>). Puppet allows the SamKnows Operations team to easily manage hundreds of test servers and ensure that each group of test

servers is configured properly as per each project requirement. Coded in Python, Puppet uses a low-overhead agent installed on each test server that regularly communicates with the controlling SamKnows server to check for updates and ensure the integrity of the configuration.

This method of managing our test servers allows us to deal with the large number of test servers without affecting the user's performance in any way. We are also able to quickly and safely make changes to large parts of our test server fleet while ensuring that only the relevant test servers are updated. This also allows us to keep a record of changes and rapidly troubleshoot any potential problems.

While Puppet handles the configuration and management of the test servers; Nagios, (the most popular online monitoring application) is used by SamKnows to monitor the test servers. Each test server is configured to send Nagios regular status updates on core metrics such as CPU usage, disk space, free memory, and SamKnows-specific applications. Nagios will also perform active checks of each test server where possible, providing us with connectivity information - both via "ping" and connections to any webserver that may be running on the target host.