10Gbps Test Lab Setup

Speedtest is capable of achieving 10Gbps speeds under ideal network conditions. This document outlines some considerations and recommendations for achieving multi-gigabit speeds, up to 10Gbps, in lab settings. In order to reach these peak speeds, the following dependencies must be met.

- The server utilized for the test must be 10Gpbs capable, and have sufficient capacity available to support the expected client load.
- The client devices from which the user is testing must be capable of reaching those speeds.
    - For example, a mobile device may not have the network or processing power required to reach 10Gbps, but a more powerful laptop or desktop setup may.
    - Client side software bottlenecks such as those created by the web browser must be eliminated by relying on one or more of our native Speedtest applications, rather than a web-based test client.
- All network components used in-path must be capable of multi-gigabit transport.
- Ensure that OoklaServer is updated or that you are using the newest version.
- Lastly, it is important to note that Speedtest utilizes a TCP-based protocol, which means that throughput performance is directly related to the latency of an end-to-end connection. In order to achieve high throughput performance, it is imperative to reduce latency and packet loss between the client and server.
    - This is very achievable in a lab situation, but remember that real speeds will decrease from what was tested in a lab environment.

Server Setup

Supporting multi gigabit testing requires careful planning and scoping of hardware and software configuration. While we can make suggestions about the capacity of the hardware and configuration that may be used, ***Ookla does not explicitly endorse a particular hardware solution over another***.

Operating System

Any up-to-date mainstream 64-bit base Linux system running an up-to-date kernel is recommended. Modern kernels tend to include optimizations to the network subsystem.

Server Hardware

While it is possible to run OoklaServer in a virtual environment, it is not recommended to do so when measuring high bandwidth connections, as they often require an additional layer of tuning and have a lower availability of serviceable platforms. Here are the key recommendations that we suggest for the physical server hardware:

1. CPU
    a. As a general rule, OoklaServer applications benefit from a CPU profile that includes fewer faster CPU cores, rather than more numerous, slower cores.
        i. Inter process communication, IPC, is the most important metric for performance with OoklaServer.
        ii. Turbo cores may result in unstable results.
        iii. Single socket board designs are easier to tune than multi socket
    b. Consider CPUs that have minimal DPC latency.
    c. Required CPU instructions and capabilities
        i. PCIe Generation 2 x4 at a minimum will be required for most 10Gbps Network Interface Cards (NICs). It is suggested that a system with PCIe Generation 3 be used due to the higher performance nic architecture afforded.
        ii. Direct PCIe lanes to the CPUare suggested, rather than going through a PCIe or PCH bridge
            1. Intel and AMD processor lines offer these direct connections, but not all system boards take advantage of this architecture.
        iii. Some workstation class systems have been observed to have bios limitations, only allowing PCIe x16 performance to GPUs while supporting other cards at x1 bandwidth.
            1. Confirm that non-GPU PCIe cards can run at full bandwidth in the slot that the NIC will be provisioned in.
2. Memory
    a. Memory size does not have a tangible impact on per-stream performance. However, it does directly correlate to the number of simultaneous clients it can serve.
        i. 16GB minimum is recommended in general, however, 32GB is recommended when serving a large quantity of clients.

    b. Modern memory bandwidth differences do not have a measurable impact on OoklaServer performance.
3. Storage
    a. Storage performance is not relevant to the performance of OoklaServer.
    b. Ensure that at least 1GB of disk is free for OoklaServer and future upgrades.
        i. Additional space is required for logging, depending on log level setting (configurable).
4. NIC
    a. A NIC that can support network connections of at least 10Gbps is required
    b. Chelsio and Mellanox produce high quality NICs that can offer improved performance, and have proven to be effective with high performance OoklaServer implementations.
        i. Intel NICs are a viable alternative.
    c. Newer generations of NICs generally offer improved efficiency
        i. these improvements are sometimes marginal relative their increased cost
    d. We suggest using fiber interfaces 10GBase-F if possible for all multi-gigabit implementations
        i. Copper 10Gbps NICs create marginal performance degradation as compared to lower latency fiber interfaces, but can provide adequate 10Gbps performance in most cases.

Example Build

A server setup that we have used with success to support single client 10Gbps testing is as follows. Many factors will impact the performance of your testing environment. This is not a guarantee that this hardware will work perfectly in your network.

- Intel® Xeon® CPU E3-1270 v3 @ 3.50GHz
- Intel® 82599ES 10 Gigabit Ethernet Controller
- 16GB RAM

To achieve higher speeds, or support multiple simultaneous 10Gbps tests, use a 20Gbit bonded or 40Gbit NIC or create a “cluster” of 10Gbps capable servers.

Server BIOS and OS Configuration

Here are the key recommendations that we suggest for server configuration:

1. CPU Tuning
    a. Hyperthreading or SMT should be disabled for peak performance. If hyperthreading must be enabled, assigning the process to non-conflicting cores will help.
        i. Example on 4 core 8 thread, assign core 0, 2, 4, 6.

    b. If NUMA is not avoidable such as on modern dual and quad socket server configurations, ensure that kernel networking threads are allocated to a single NUMA node. Interrupts will also have to be carefully tuned. The OoklaServer processes should also be restricted to the same NUMA node as the interrupts and memory
    c. Unless leveraging a tool such as FD.io, ensure that interrupts are tuned for the socket(s) in use for the OoklaServer processes
        i. IRQBalance
        ii. https://fasterdata.es.net/host-tuning/linux/100g-tuning/interrupt-binding/
    d. If possible, disable netfilter, ipset, cillium etc. and perform firewalling upstream
    e. Ensure that the intel_pstate governor is used with the HWP + performance flags
2. Network Hardware Tuning
    a. Ensure that NIC hardware offloading is enabled
    b. Ensure that Generic Receive Offload (GRO) is enabled.
    c. Ensure the NIC attaches at the correct bus bandwidth and signalling rate (this can be seen with lspci on Linux, the nic properties on Windows, etc.)
    d. Enable Direct Cache Access if the NIC and BIOS supports it
    e. Enable IO/AT if the NIC and BIOS support it
    f. If using Mellanox nics, follow their Linux tuning guidelines too: https://community.mellanox.com/s/article/performance-tuning-for-mellanox-adapters

Network Configuration

1. For lab purposes disabling firewalls can help prevent test interference.
    a. Any type of Deep Packet Inspection or network traffic filtering may impact and degrade performance in a testing scenario.
2. Increase TCP send/receive buffer sizes as outlined in our performance optimization guidelines.
3. Use congestion control algorithms that match your network. Most Linux distributions default to CUBIC which is loss based congestion control and may underperform in some cases where loss may not indicate congestion such as radio networks.