

Table of Contents

Equipment Recommendations	3
Minimum Requirements	3
Recommended Requirements for Every 10Gbps Peak Traffic on DPI	3
Detailed Recommendations	3
Tips from Yandex ClickHouse	4

Equipment Recommendations



Do not install the module on a server with a DPI platform!

Minimum Requirements

The component can be installed on a VM for testing purposes with the following minimum requirements:

1. Processor (CPU) 2.5 GHz - 1 core
2. RAM - at least 16 GB
3. Hard disk (SSD highly recommended) - at least 500 GB
4. Operating system - CentOS 8.x, [VEOS](#), CentOS Stream 8.x, Oracle Linux Server 8.x, AlmaLinux 8.x
5. Network interface card (NIC) - at least 1Gbps

Recommended Requirements for Every 10Gbps Peak Traffic on DPI

1. Processor (CPU) from 2.5 GHz - 6 cores
2. RAM - 64 GB
3. Hard disk (SSD highly recommended) - at least 500 GB, see below for storage volume calculation and storage organization recommendations
4. Operating system - CentOS 8.x, [VEOS](#), CentOS Stream 8.x, Oracle Linux Server 8.x, AlmaLinux 8.x
5. Network interface card (NIC) - 2x10Gbps. It should be noted that each DPI generates an IPFIX flow at a speed of 0.5% to 1% of the real traffic speed. It is also recommended to aggregate the ports on QoE into a LAG for fault tolerance.

Example of a QoE server receiving IPFIX from DPI for 100Gbps peak traffic (in+out): Server platform (2U, AMD EPYC 7713 processor with 64 cores, 512 GB RAM, HW RAID Controller, 2 x 960GB SSD RAID1 for OS, 4x3.84TB SSD NVME RAID0 stripe default disks + HDD/SSD RAID50 disks for storage of a specific volume, 2x network adapter 2x25GbE, 2xPSU)

Storage Volume Calculator Based on Average Traffic Speed

It is assumed that the average daily traffic is 60% of the total peak (in+out) traffic.

In the provided calculator, you need to change the traffic value to get the storage volumes.

Detailed Recommendations

CPU	<p>One processor supporting SSE 4.2 instructions starting from Intel Nehalem and AMD EPYC Zen2 with 4 or more cores, base clock speed of 2.5 GHz and higher. Choose processors with more cores. Clock speed is less important. For example, 16 cores at 2600 MHz is better than 8 cores at 3600 MHz.</p> <p>Do not disable Hyper-threading and Turbo-Boost.</p>
RAM	<p>At least 16 GB, memory modules must be installed in all processor channels on the motherboard. Memory should be no less than the volume of requested data. More memory improves performance when generating reports. More memory also reduces disk load.</p> <p>Always disable swap file.</p>
Disks	<p>To optimize storage costs, several types of disks are used: default — fast disks for data reception and aggregation process, it is recommended to use SSD NVMe in RAID0. hot — disks for storage during periods of higher likelihood of report requests on this data, typically up to 3 months, SSD disks in RAID-10, RAID-5, RAID-6, or RAID-50. cold — slow, large-volume disks for long-term storage, HDD disks in RAID-10, RAID-5, RAID-6, or RAID-50 are recommended.</p> <p>The storage duration for each level is set in the configuration via the GUI. Data migration between disks and data cleanup happens automatically based on the settings. A mechanism for overflow control is also provided to protect the database. The main volume of data is stored in the /var/lib/clickhouse directory. Temporary data (IPFIX dumps) are stored in the /var/qoestor/backend/dump directory. For better performance, it is important (recommended) that these directories are located on a separate disk or array. See Disk Configuration.</p> <p>For OS and QoE Stor software installation, use 2 disks of at least 256GB capacity, combined in RAID 1 (mirror). A hardware RAID controller is required.</p>
QoE Cluster (Sharding)	<p>It is better to create multiple nodes and combine them into a cluster: GUI optimizes queries so that all nodes generate reports in parallel. IPFIX-balancer is used to evenly distribute data across nodes (roundrobin), significantly improving system performance.</p> <p>In case of node failure, the balancer will automatically distribute data to the remaining nodes. General recommendation: as many nodes as possible and as little data per node as possible. This will provide:</p> <ol style="list-style-type: none"> 1. High performance 2. Good fault tolerance 3. Scalability (by adding nodes to the cluster)

Tips from Yandex ClickHouse

You can read Yandex ClickHouse operation tips at <https://clickhouse.yandex/docs/ru/operations/tips/>.