

Table of Contents

Hardware requirements	3
<i>Minimum requirements</i>	3
<i>Recommended requirements per each 10Gbps of peak traffic on DPI</i>	3
Storage size calculator based on average daily traffic	3
<i>Detailed recommendations</i>	4
<i>Operation tips from Yandex ClickHouse</i>	4

Hardware requirements



Do not install the QoE module on the SSG server (BRAS, NAT, DPI)!

Minimum requirements

The component can be installed on a VM for testing with minimal requirements:

1. Processor (CPU) 2.5 GHz - 1 core
2. RAM - from 16 GB
3. Hard drive (SSD strongly recommended) - from 500 GB
4. Operating system - CentOS 8.x, [VEOS](#), CentOS Stream 8.x, Oracle Linux Server 8.x, AlmaLinux 8.x.



When using VEOS, take into account the information in [VEOS Update Notes](#) when selecting hardware.

5. Network interface card (NIC) - from 1Gbps

Recommended requirements per each 10Gbps of peak traffic on DPI

1. Processor (CPU) 2.5 GHz or higher - 6 cores
2. RAM - 64 GB
3. Hard drive (SSD strongly recommended) - from 500 GB; see details about storage capacity calculation and recommendations below
4. Operating system - CentOS 8.x, [VEOS](#), CentOS Stream 8.x, Oracle Linux Server 8.x, AlmaLinux 8.x.



When using VEOS, take into account the information in [VEOS Update Notes](#) when selecting hardware.

5. Network interface card (NIC) - 2x10Gbps. It is important to consider that each DPI generates an IPFIX stream at a rate of 0.5% to 1% of real traffic speed. It is also recommended to combine ports on the QoE in a LAG for fault tolerance.

Example of a QoE server receiving IPFIX from DPI for 100Gbps of peak traffic (in+out): Server platform (2U, AMD EPYC 7713 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 for storage according to volume, 2x network adapters 2x25GbE, 2xPSUs)

Storage size calculator based on average daily traffic

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

In the provided calculator, change the traffic value to calculate storage size.

Detailed recommendations

CPU	<p>Single processor supporting SSE 4.2 instructions starting from Intel Nehalem and AMD EPYC Zen2 with 4 or more cores and a base clock speed of 2.5 GHz or higher. Choose CPUs 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>From 16 GB; memory modules must be installed in all CPU channels on the motherboard. The memory size should not be less than the volume of queried data. The more memory, the better performance when generating reports and the lower the disk load.</p> <p>Always disable swap file.</p>
Disks	<p>File system type: ext4.</p> <p>To optimize storage cost, multiple types of disks are used:</p> <ol style="list-style-type: none">default — fast disks for data ingestion and aggregation processes, SSD NVMe in RAID0 recommended.hot — disks for storing data likely to be queried (usually up to 3 months), SSDs in RAID-10, RAID-5, RAID-6, or RAID-50.cold — high-capacity slow disks for long-term storage, HDDs in RAID-10, RAID-5, RAID-6, or RAID-50 recommended. <p>Retention period at each level is configured via GUI. Data migration and cleanup occur automatically according to settings. A mechanism for overflow protection is also provided. The main data volume is stored in /var/lib/clickhouse. Temporary data (IPFIX dumps) are stored in /var/questor/backend/dump. For best performance, these directories should be located on a separate disk or array. See Disk space configuration. For OS and QoE Stor software installation, use two drives of at least 256GB combined in RAID1 (mirror). A hardware RAID controller is required.</p>
QoE Cluster (Sharding)	<p>It is better to create several nodes and combine them into a cluster: GUI can optimize queries so that all nodes build reports in parallel. IPFIX-balancer is used for even data distribution across nodes (round-robin), significantly improving performance.</p> <p>If a node fails, the balancer automatically directs data to remaining nodes. General recommendation: more nodes and smaller data portions per node. This ensures:</p> <ol style="list-style-type: none">High performanceFault toleranceScalability (by adding nodes to the cluster)

Operation tips from Yandex ClickHouse

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