

LiteSwitch is now available via on-premise installation, enabling the deployment and use of LiteSwitch services within your own technical infrastructure. Getting started with this new implementation method is simple – more information on this service and how to get started can be found below:


What's changing?

LiteSwitch was initially brought to market as a “virtual sidecar” implementation. This was a SaaS-based integration designed for quick onboarding and scalability but created potential exposure to issues related to latency and variable hardware cost.

The implementation method is migrating to an on-premise deployment, where almost all LiteSwitch services will be installed in the client's infrastructure, thereby eliminating risks related to latency and hardware cost.


Requirements

The LiteSwitch server image is transferred within a Docker container and does not impose specific requirements for infrastructure parameters. However, the following or similar configuration is recommended:



CPU:

- Type: Intel Icelake or AMD EPYC
- Cores: 8 to 16 cores (can vary based on trading volumes)
- Architectures: x86-64 or aarch64 (arm64)



RAM

- Capacity: 8 GB
- Type: DDR4 ECC



Storage

- Type: Enterprise-grade SSD (preferable)
- Capacity: from 100 GB (can vary based on trading volumes)

How does it work?

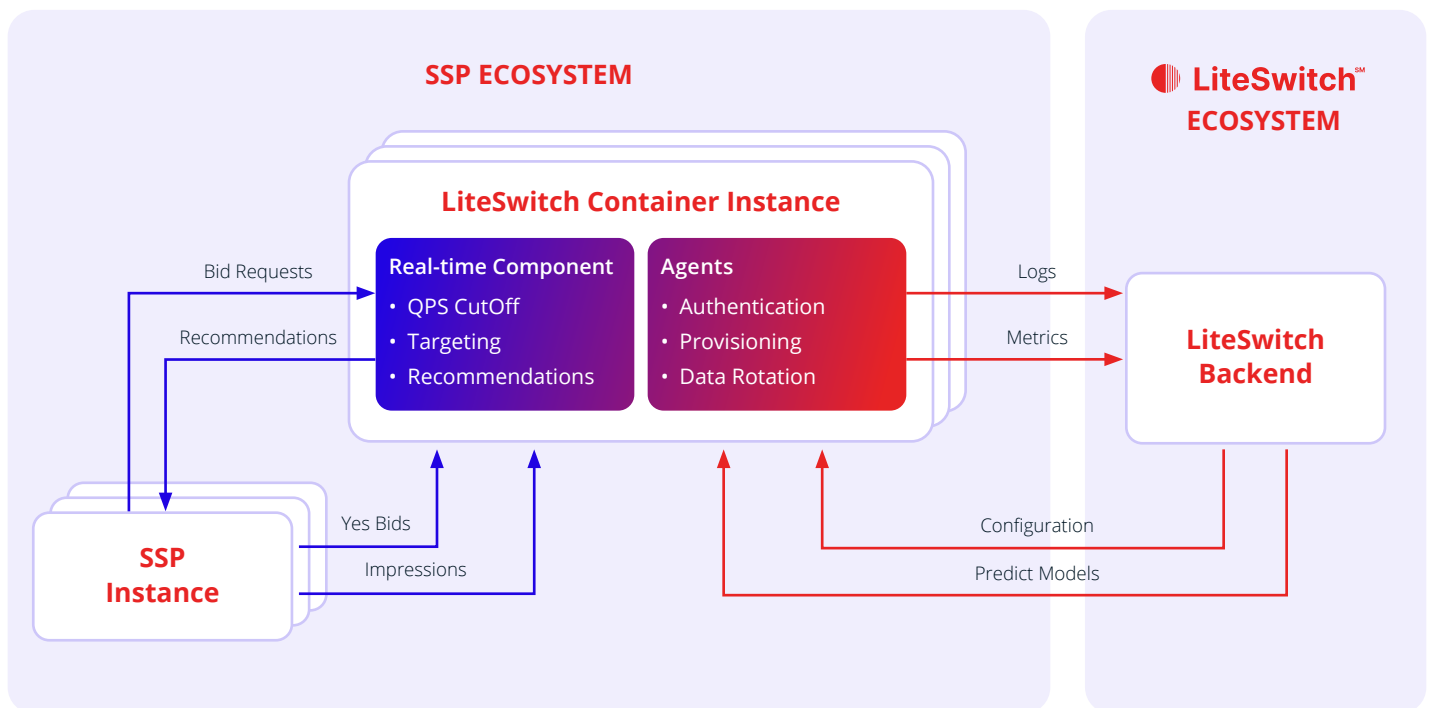
LiteSwitch will be distributed in a form of a Docker container, which consists of two distinct components:

- **Real-Time:** interacts with the client's supply-side platform (SSP) to process requests and provide recommendations, enforces configurations such as QPS caps and targeting, and processes yes-bid/impression data
- **Agents:** interacts with the LiteSwitch backend to authenticate the client-side instance, retrieves configuration metrics and predictive models, deliver logs and metrics for system monitoring, and provisions the Real-Time component

The container is only discoverable within your ecosystem, and each container will run in a separate instance.

Containers scale linearly, which means processing volumes are directly correlated to the number of required instances.

The instances require internet access to reach the LiteSwitch backend to perform actions such as fetching configuration updates and delivering logs/metrics, but they do not require external accessibility from the internet itself.



How to Get Started

To get started on implementations, clients need to complete the following steps:

- 1. Infrastructure Configuration:** Complete an [intake questionnaire](#) to define current infrastructure and setup configuration
- 2. Team Allocation:** Confirm dedicated team members responsible for installation
- 3. Timing:** Align on desired timeframe to run installation

Throughout this process, the LiteSwitch team will provide:

- Detailed documentation to guide installation
- Dedicated support throughout installation
- Continuous communication via email and dedicated Slack channel
- Post-installation support for a smooth transition