

Performance and Scalability

LightWave Server™ was designed to be fully scalable using the Supervisor-Worker pattern. A LightWave Server *instance* consists of a LightWave Server SERVER process pair and optionally one or more LightWave Server SWORKER processes. The SERVER process listens on the instance's configured TCP/IP ports, hosts the LightWave Server Console, and services new connection request from API clients. If no SWORKER processes have been configured, the Server process will also perform API request processing. Any number of LightWave Server instances may be running on a single system, as long as they are configured to use different TCP/IP ports.

As application load increases, SWORKER processes may be started and attached to a SERVER process. When SWORKER processes are attached, the SERVER process will delegate new connections to the SWORKER processes in round-robin fashion. Any number of SWORKER processes may be started in any CPU, providing adequate capacity for any application.