



Large Volume Logging

High-Performance Scalable Log Collection

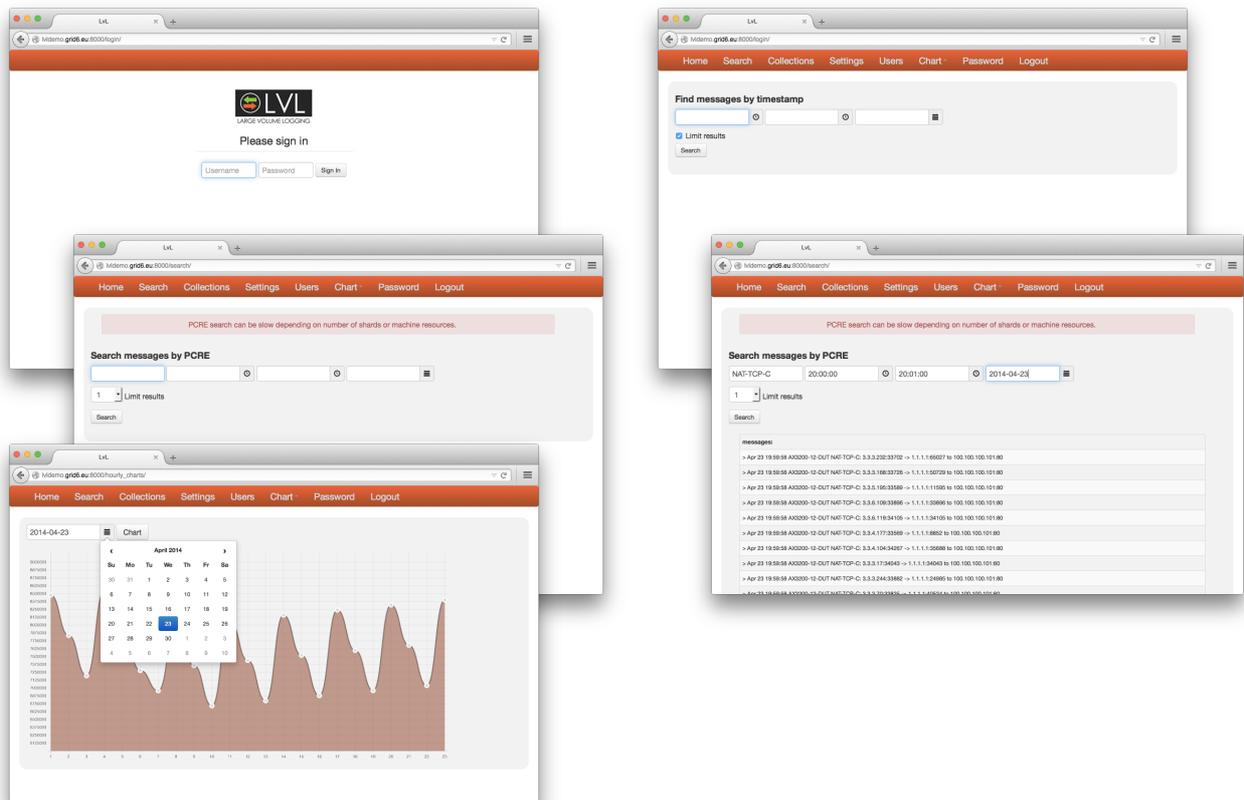
LVL provides high-performance log collection, indexing and easy searching. LVL is specifically designed for organisations that are mandated by data retention policy to collect and keep logging data.

Historical Auditing

When dealing with data retention policies and regulatory compliance, being able to audit historical data is a challenge. LVL provides indexing and easy search capabilities on log data collected months before.

Data Retention Policy

To comply with data retention policies, organisation must be able to collect, maintain and keep significant amounts of logging data. Being able to store this data in an efficient manor is critical. LVL uses intelligent indexing mechanisms to store logging data while maintaining easy access for auditing purposes.



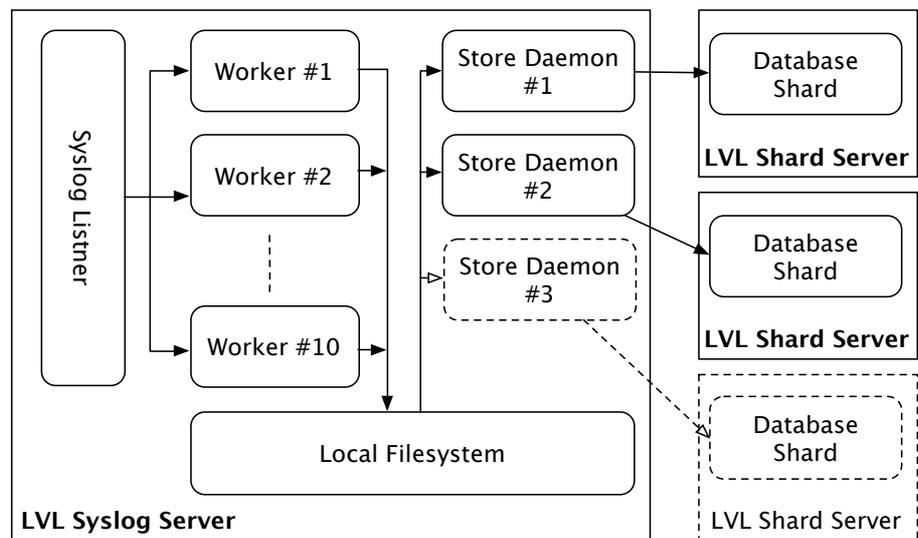


Datasheet

Highly Scalable Log Collector

When dealing with hundreds of thousands messages per second it's imperative that every single message is collected. LVL is designed to be highly scalable for this purpose. It can grow with the log collection needs by moving processes to dedicated hardware and scale horizontal.

Linear scaling can be achieved by load balancing multiple instances (on commodity hardware) of LVL. This will scale LVL beyond 400.000 messages per second.



Reference Architecture

An Intel Quad Core i7 (3Ghz) with 32G of memory, processes up to 400.000 messages per second with a message size of 300 bytes.

Storage per month	100 bytes	200 bytes	300 bytes
100.000 m/s	26 TB	52 TB	78 TB
200.000 m/s	52 TB	104 TB	156 TB
300.000 m/s	78 TB	156 TB	234 TB
400.000 m/s	104 TB	208 TB	311 TB

To find out more about LVL and how to get control of your logging data, you can contact Grid6 at <http://grid6.io> or call and talk to a Grid6 partner.