

# 1. Infrastructure runtime cluster

## Introduction

This chapter describes what infrastructure is needed in order to run Blueriq Runtime in a clustered environment.

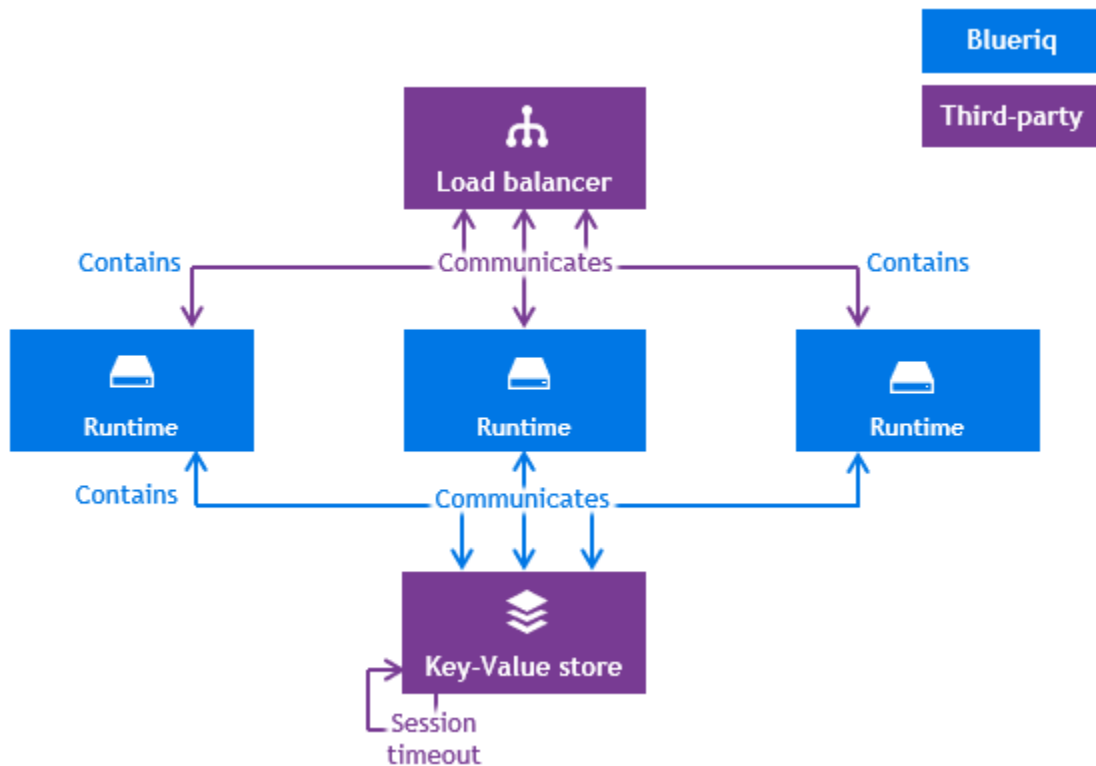
A **cluster** is a group of multiple server instances, spanning across more than one node, all running identical configuration. All instances in a cluster work together to provide high availability, reliability, and scalability.

- [Introduction](#)
- [Minimum Hardware and Software Requirements](#)
  - [Load balancer](#)
  - [Blueriq Runtime](#)
  - [Key-Value store server](#)

## Minimum Hardware and Software Requirements

Each paragraph described below is a software requirement that is assumed to be running in it's own node, or even on multiple nodes where applicable.

*Runtime cluster consists of the following entities:*



### Load balancer

A reverse proxy for load balancing HTTP requests for the Blueriq runtime is needed. It also runs in it's own separate node.

Blueriq does not specify a specific load balancer to be used, but for following chapters regarding clustering, nginx will be used as [example](#).

### Blueriq Runtime

Two or more Blueriq Runtimes running each of them on different nodes with the same configuration.

Documentation on how to set up the runtime for cluster configuration can be found [here](#).

### Key-Value store server

A Redis server (Key-Value store server) running in it's own node is required for running Blueriq in a clustered environment. It is used by the Blueriq Runtime for storing and retrieving data about the currently logged in user, and applications current state.



Redis is the default in-memory data structure store supported by Blueriq. For using other in-memory data store and how to add support for them in runtime can be found [here](#).

**Previous:** [Runtime cluster](#)

**Next:** [2. About the load balancer](#)