

# lgtm enterprise



## LGTM Enterprise System Requirements

---

Release 1.19.2, February 2019

**Semmlé**<sup>TM</sup>

**Semmler Inc.**

44 Montgomery Street  
3rd Floor  
San Francisco, CA 94104

Copyright © 2019, Semmler Ltd. All rights reserved.

LGTM Enterprise release 1.19.2

Document published February 11, 2019

# Contents

<b>Introduction</b> .....	<b>4</b>
<b>Languages supported</b> .....	<b>5</b>
<b>System requirements</b> .....	<b>6</b>
<b>Work pool software requirements</b> .....	<b>9</b>
<b>Deployment</b> .....	<b>11</b>
<b>System integrations</b> .....	<b>14</b>
<b>Browser support</b> .....	<b>16</b>

# Introduction

LGTM Enterprise is built on Semmle’s tried and tested analysis technology. The product is the on-premises twin of Semmle’s [LGTM.com](#) platform—a publicly available, free to use, cloud-based service that continuously analyzes over 100,000 open source projects worked on by over 700,000 developers, and containing data for more than 39 million commits.

## Note

“LGTM” is the most commonly posted comment on code review platforms (such as GitHub). It means: “looks good to me.”

Most notably, LGTM Enterprise offers the full power of LGTM.com analysis to customers in an on-premises version. LGTM allows developers to efficiently browse alerts in their source code (using either a web browser or a plugin for their IDE), while at the same time incorporating broader insights. It seamlessly integrates with platforms like GitHub Enterprise and Bitbucket Server for code review, authorization, and authentication.

## About this document

This document describes the minimum system requirements for a deployment of LGTM Enterprise. For more detailed information about the services mentioned in this document, see the [administrator help](#).

## Related documentation

- [LGTM Enterprise System Architecture](#) (PDF)
- [LGTM Enterprise Installation Guide](#) (PDF)
- [LGTM Enterprise administrator help](#)

To access this, click **Admin help** at the top of the administration pages in LGTM Enterprise, or browse to [help.semmle.com/lgtm-enterprise/admin](https://help.semmle.com/lgtm-enterprise/admin).

# Languages supported

LGTM Enterprise currently supports analysis of code written using the following languages. Where there are several versions or dialects of a language, the supported variants are listed.

Language	Variants	Compilers
C/C++	C89, C99, C11 C++98, C++03, C++11, C++14, C++17	Clang extensions up to Clang 6.0 GNU extensions up to GCC 7.3 Microsoft extensions up to VS 2017
C#	C# up to 7.2 together with .NET versions up to 4.7.1	Microsoft Visual Studio up to 2017 .NET Core up to 2.1
COBOL*	ANSI 85 or newer	Not applicable
Java	Java 11 or lower	javac (OpenJDK and Oracle JDK) Eclipse compiler for Java (ECJ) batch compiler
JavaScript	ECMAScript 2018 or lower. In addition, the following are analyzed: <ul style="list-style-type: none"> <li>• JSX, Flow, and TypeScript code</li> <li>• YAML, JSON, and HTML files</li> </ul>	Not applicable
Python	2.7, 3.5, 3.6, 3.7	Not applicable
TypeScript	2.6, 2.7, 2.8, 2.9, 3.0, 3.1	Standard TypeScript compiler

\* Available on demand. Ask your support contact for details.

# System requirements

In some systems the control pool consists of a single machine—referred to as the "controller." This also hosts the web and data aggregation services so that there is no separate web pool. At larger scales a web pool of additional machines can be added to increase capacity for web and data aggregation services, with the "controller" remaining on a separate machine.

Type	Hardware	Software
Control pool*	<b>Controller</b> <ul style="list-style-type: none"> <li>• 8 cores</li> <li>• 16 GB RAM</li> <li>• Solid-state drive (SSD)<sup>†</sup></li> <li>• 1 TB free disk space</li> </ul>	<b>Supported Linux distributions/versions:</b> <ul style="list-style-type: none"> <li>• CentOS 7</li> <li>• Debian 9</li> <li>• Red Hat Enterprise Linux 7</li> <li>• Ubuntu 16.04 LTS</li> <li>• Ubuntu 18.04 LTS<sup>‡</sup></li> </ul>
	<b>Web pool*</b> <ul style="list-style-type: none"> <li>• 2 cores</li> <li>• 8 GB RAM</li> <li>• SSD recommended<sup>†</sup></li> <li>• 50 GB free disk space</li> </ul>	
Work pool*	<ul style="list-style-type: none"> <li>• 2 cores</li> <li>• 8 GB RAM</li> <li>• SSD recommended<sup>†</sup></li> <li>• 50 GB free disk space (per worker assigned)</li> </ul>	<ul style="list-style-type: none"> <li>• Linux versions as above, or Windows 7 or higher.</li> <li>• The 64-bit version of the most recent release of OpenJDK 8, or Oracle JDK 8. Note: IBM Java, and other versions of Java are not supported.</li> <li>• Python 2.7</li> </ul>

\* For more information, see the *LGTM Enterprise System Architecture* PDF, or the "[Server overview](#)" topic in the LGTM Enterprise administrator help.

† Semmler recommends that you use solid-state drives (SSDs) rather than hard disk drives (HDDs) for all of the machines, in both the control pool and the work pool. It is strongly recommended that you use an SSD for the "controller" machine, as files are served from here to all other cluster machines.

‡ OpenJDK 8 or Oracle JDK 8 must be installed if you use Ubuntu 18.04.

### Note

During installation, upgrade and configuration processes, root access is required for each machine.

Virtual (for example, VMWare, Xen, KVM) or physical machines are both suitable for the pools of machines, but the term “virtual machines” is used exclusively below.

The computational power required for the work pool depends on the number of projects to analyze, their size, build times (for compiled languages such as Java), and the average commit frequency.

As this information is usually difficult to ascertain in advance, it makes sense to plan for an incremental deployment based on the following information:

- Minimum deployment size for LGTM Enterprise
- Hardware requirements at scale
- Incrementally adding capacity

## Third-party packages required for installation

LGTM Enterprise also requires access to the following third-party packages.

- Java JRE 8
- PostgreSQL 9.5 or higher
- RabbitMQ Server
- Nginx
- Minio release 2018-07-23T18-34-49 (bundled with LGTM)
- Solr 7.1.0 (bundled with LGTM)
- Git
- Subversion
- Python 2 (version 2.7 or higher will be installed)

By default, these are installed automatically on the relevant machines during the installation of LGTM Enterprise. With the exception of Minio and Solr, the exact versions that are installed depends on the latest version available for your Linux distribution. Any dependencies of these packages are also installed.

## Access to third-party packages during installation

The simplest way to install LGTM Enterprise is to enable internet access and to run the installation script supplied with the application. This enables the installer to download and

install the required packages from their official online package repositories using the standard utilities available in your Linux distribution. If you're using Red Hat Enterprise or CentOS, this will require access to the EPEL (Extra Packages for Enterprise Linux) repository. Once LGTM has been deployed, internet access can be closed down.

If internet access is not available, you can do either of the following:

- Configure the machine so that it can access an internally hosted mirror that contains these package repositories and their dependencies. The installer will download and install the required packages from your internal mirror.
- *For installations on Red Hat or CentOS:* Extract the third-party dependencies pack before you run the installer. Details for doing this are provided in the installation instructions.

# Work pool software requirements

The LGTM work pool must be set up with machines that include the right software and the correct environment to enable checkout and analysis of the code bases you want to analyze.

## Checkout requirements

The essential requirements are:

1. **Appropriate client software** to allow the worker to check out code from the repository—that is, Git, Subversion or Team Foundation Version Control (TFVC). The LGTM Enterprise installer automatically installs Git and Subversion clients. For TFVC, you must install the Team Explorer Everywhere command-line client (TEE-CLC) on each worker host machine.
2. **Authentication method** to allow the worker host machine to authenticate with the repository host and have read access to the code. For Subversion and TFVC repositories, the credentials you specify when you create a repository provider (see "[Adding repository providers](#)" in the administrator help) are used. For Git repositories, the simplest method is use an SSH configuration, but other options are also available.

For more information, see "[Configuring hosts to checkout code](#)" in the LGTM Enterprise administrator help.

## Analysis requirements

During the analysis of compiled languages, LGTM builds the code. This means that there must be at least one machine in the work pool that is configured with the environment needed to build each code base that you want to analyze. Interpreted languages often also depend on the presence of other software in the environment, and these must be made available too.

### Core requirements

The core requirements listed below are the minimum required to analyze code written in each language. Your projects may have additional requirements. Compiled languages have further requirements (detailed in the following section).

Language	Essential
C/C++	Method of building the code (see next table).

<b>C#</b>	Method of building the code (see next table).
<b>COBOL</b>	No further requirements. Analysis is available on demand. Ask your support contact for details.
<b>Java</b>	Method of building the code (see next table).
<b>JavaScript</b>	No further requirements unless any TypeScript files are present. If the repository also includes TypeScript files, the TypeScript requirements must be met, or TypeScript analysis disabled.
<b>Python</b>	Access to the pip packaging management system. Installation of the packaging and virtualenv pip modules. Ability to download any packages that the Python code bases you intend to analyze depend on.
<b>TypeScript</b>	Node.js version 6.x or later.

## Compiled languages

For compiled code that is normally built using one of the most commonly used build systems (listed in the following table), LGTM Enterprise has an "autobuild" process which is designed to automatically build the code. If your organization uses these build systems, install them on the relevant worker machines. If some projects use other build systems, you also need to make these systems available. These projects will require a custom build configuration.

<b>Language</b>	<b>Autobuild available for</b>
<b>C/C++</b>	Linux workers: Autoconf, CMake, QMake, Meson, WAF, SCons, and Linux KBuild Windows workers: MSBuild
<b>C#</b>	MSBuild, .NET Core, Mono (Linux), and Visual Studio (Windows)
<b>Java</b>	Gradle, Maven and Ant

For more information, see "[Configuring hosts to analyze code](#)" in the administrator help.

# Deployment

The software is deployed to Linux machines using packages. Semmler provides packages in common formats (RPMs and Debian packages) for the LGTM Enterprise software and for those dependencies that are not available by default through recent Linux distributions.

Machines in the work pool should be configured to checkout and build the software that they will analyze and tested before LGTM is deployed to them. LGTM is deployed using LGTM RPM or Debian packages on Linux machines, and an MSI installer on Windows machines.

## Minimum deployment

For an initial trial, a suitable cluster of machines is:

- Control pool: 1 virtual machine (VM)
- Work pool: 5–20 virtual machines

The hardware requirements defined in the [system requirements](#) are enough to support one worker per work pool VM. If you want to run more than one general worker per work pool VM, multiply the disk space, RAM and CPU cores accordingly.

### Note

When you add projects with over 1 million lines of code to LGTM, you may need to increase the resources available on the work pool machines. More RAM and disk space may be required to build large projects. If the builds are parallelized, increasing the number of cores available on each machine may also be beneficial.

In addition to the hardware and software requirements above, each VM in the work pool must be set up with the build environment that is required to build the projects that will be analyzed. For example:

- For C# analysis the correct versions of the C# compiler must be installed.
- For Python analysis pip and the packaging module must be installed. In addition, to analyze code for Python 2 projects, the virtualenv module must be installed.

## Hardware requirements at scale

LGTM Enterprise is designed to scale successfully from analyzing the work of a few hundred developers to analyzing the work of many thousands of developers. The resources required

change according to the nature of the repositories analyzed, size of the projects, number of developers and commit frequency. We recommend that you start with a minimum deployment and scale this as required.

You can monitor the free capacity of the work pool using the Infrastructure administration page. If you need to add extra capacity, you should provision additional machines and register them as part of the work pool.

In rare cases where there is a particularly large amount of data, the services on the controller can be hosted on different machines.

Other bottlenecks are usually best addressed by changing the control pool from a single machine to a cluster. For example:

- Create a web pool of several machines to load balance the web application and web proxies. These machines also process the results returned by the work pool.
- Host the database on a dedicated virtual machine with substantial storage.
- Host the file store on a dedicated virtual machine with substantial storage.

### Examples

There's an enormous variation in how companies organize their source code and development workflow. This makes it difficult to predict the hardware requirements for LGTM Enterprise in advance of installation, which is why LGTM is designed to scale smoothly. To help you get an idea of the likely requirements, here are some examples of the resources used by production instances of LGTM.

#### LGTM.com

This instance analyzes the work of over half a million developers working on code in over eighty thousand repositories. The following resources are used on Google Cloud Platform:

- A dynamic control pool that scales up and down automatically according to the current workload.
- A dynamic work pool that typically runs between 600 and 1250 workers. Each worker runs on a separate, [preemptible](#), virtual machine with 2 vCPUs, 7.5 GB of system memory, and 80 GB of disk space.

#### LGTM Enterprise

In this example, the instance of LGTM Enterprise analyzes the work of approximately 2,500 developers working on code in over 850 repositories. The following internal virtual machines are used:

- One controller running on a virtual machine with 6 vCPUs 2.6 GHz, 32 GB of system memory, and 1 TB disk space.

- A work pool containing 15 general workers and 1 query worker. These are spread across 8 virtual machines. Each of these virtual machines has 6 vCPUs 2.6 GHz, and 32 GB of system memory.

## Incrementally adding and removing capacity

As new projects are added to LGTM Enterprise, the work pool (and occasionally the control pool) will need to be expanded. Machines in the work pool are stateless and can all be configured identically. This makes it easy to scale the pool up and down elastically in response to demand, by provisioning virtual machines from a common image.

# System integrations

Integration	User authentication	User authorization	Code review integration	Repository access only
<b>Azure DevOps</b> (Git and TFVC projects)	✓	✓	Git projects only	
<b>Bitbucket Cloud</b> (Git projects only)	✓	✓	✓	
<b>Bitbucket Server</b> 4.5.0 and above (Git projects only)	✓	✓	✓ *	
<b>Git</b>				✓
<b>GitHub.com</b> (public projects only)	✓		✓	
<b>GitHub Enterprise</b> 2.9 and above	✓	✓	✓	
<b>GitLab.com</b>	✓		✓	
<b>GitLab Enterprise and Community Edition</b> 9.0 and above	✓	✓	✓	
<b>OneLogin</b> (SAML authentication)	✓			
<b>Subversion</b> 1.8 and above				✓

<b>Team Foundation Server</b> 2017 and above (Git and TFVC projects)	✓			
--	---	--	--	--

\* Requires that the Bitbucket Server trusts the LGTM instance's SSL certificate, and has the Post Webhooks plugin installed. For the latest information about which versions of Bitbucket the plugin supports, see the [plugin's documentation](#).

## Issue tracker integration

LGTM Enterprise integrates with version 7.6.9 and above of Atlassian Jira. For more information, see the topic "[Adding Jira integration](#)" in the administrator help.

Other issue tracker systems are supported via webhook integration—see the topic "[Adding webhook-based issue tracker integration](#)" in the administrator help.

## Browser support

The following web browsers are supported for accessing the web pages of an LGTM Enterprise deployment:

- Google Chrome
- Mozilla Firefox
- Apple Safari
- Microsoft Edge

Semmler aims to provide an optimal user experience for the most recent version of all of the browsers mentioned above, and will support browser versions for up to three years after their release.