

# Building Impala

This page describes how to build Impala from source and how to configure and run Impala in a single node development environment.

## Prerequisites and Requirements

### Hardware Requirements

- CPU must support at least SSSE3
- Minimum memory: 16GB
- Hard disk space: 120GB (for the test data)

### Supported Operating Systems

#### Linux

- Ubuntu 14.04, 16.04, 18.04
- CentOS 7
- See `bootstrap_development.sh` for other supported versions (this wiki page may be stale).

#### OS X

Not supported

#### Windows

Not supported

### Option 1 - Building Impala (for developing Impala)

```
git clone https://gitbox.apache.org/repos/asf/impala.git ~/Impala
cd ~/Impala
export IMPALA_HOME=`pwd`
./bin/bootstrap_development.sh
```

### Option 2 - Building Impala without Test Data (for testing Impala)

```
git clone https://gitbox.apache.org/repos/asf/impala.git ~/Impala
cd ~/Impala
export IMPALA_HOME=`pwd`
./bin/bootstrap_system.sh
source ./bin/impala-config.sh
# Format the test cluster and start Impala and dependent services
./buildall.sh -noclean -notests -format -start_minicluster -start_impala_cluster
```

### Rebuilding after initial build

```
# Rebuild both backend and frontend
${IMPALA_HOME}/buildall.sh -skiptests -noclean

# Rebuild with optimized release binaries
${IMPALA_HOME}/buildall.sh -skiptests -noclean -release

# Incremental builds
source ${IMPALA_HOME}/bin/impala-config.sh # If you didn't already source impala-config.sh in this shell
# Optional: Rebuild the impala binary only
make -j${IMPALA_BUILD_THREADS} impalad
# Optional: Build the Java-side frontend only
make -j${IMPALA_BUILD_THREADS} java

# Restart the Impala cluster
${IMPALA_HOME}/bin/start-impala-cluster.py
```

See [Tips for Faster Impala Builds](#) for more tips on how to do incremental builds.

## FAQ

**Q: The build is stuck at 100% progress for a while (as the screenshot shows). What's going on there?**

```
[ 98%] Building CXX object be/src/util/CMakeFiles/impala-profile-tool.dir/impala-profile-tool.cc.o
[ 98%] Building CXX object be/src/service/CMakeFiles/impalad.dir/daemon-main.cc.o
[ 98%] Building CXX object be/src/util/CMakeFiles/loggingsupport.dir/logging-support.cc.o
[ 98%] Linking CXX executable ../../build/release/service/impalad
[ 98%] Linking CXX executable ../../build/release/util/impala-profile-tool
[ 98%] Linking CXX shared library ../../build/release/util/libloggingsupport.so
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
[ 98%] Built target impalad
[ 98%] Generating ../../build/release/service/admissiond
[100%] Generating ../../build/release/service/statestored
[100%] Generating ../../build/release/service/catalogd
[100%] Built target admissiond
[100%] Built target statestored
[100%] Built target catalogd
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
[100%] Built target impala-profile-tool
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
[100%] Built target loggingsupport
[100%] Linking CXX shared library ../../build/release/service/libfesupport.so
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
../../toolchain/toolchain-packages-gcc7.5.0/openldap-2.4.47/lib/libldap.a(os-ip.o):os-ip.c:functi
[100%] Built target fesupport
[100%] Built target notests_regular_targets
```

A: Usually the FE compilation is still running. Maven downloads a lot of artifacts. So you will see slow progress there if your internet connection is slow. You can check the logs at `$(IMPALA_HOME)/logs/mvn/mvn.log`