

Performance Tuning DSpace

- 1 [Give Tomcat \(DSpace UIs\) More Memory](#)
 - 1.1 [Give Tomcat More Java Heap Memory](#)
 - 1.2 [Give Tomcat More Java PermGen Memory](#)
- 2 [Use the Parallel Garbage Collector](#)
- 3 [Give the Command Line Tools More Memory](#)
 - 3.1 [Give the Command Line Tools More Java Heap Memory](#)
 - 3.2 [Give the Command Line Tools More Java PermGen Space Memory](#)
- 4 [Give PostgreSQL Database More Memory](#)
- 5 [SOLR Statistics Performance Tuning](#)

The software DSpace relies on does not come out of the box optimized for large repositories. Here are some tips to make it all run faster.

Give Tomcat (DSpace UIs) More Memory

Give Tomcat More Java Heap Memory

Java Heap Memory Recommendations



At the time of writing, DSpace recommends you should give Tomcat \geq 1024MB of Java Heap Memory to ensure optimal DSpace operation. Most larger sized or highly active DSpace installations however tend to allocate more like 1536MB to 2048MB of Java Heap Memory.

Performance tuning in Java basically boils down to memory. If you are seeing "java.lang.OutOfMemoryError: Java heap space" errors, this is a sure sign that Tomcat isn't being provided with enough Heap Memory.

Tomcat is especially memory hungry, and will benefit from being given lots of RAM. To set the amount of memory available to Tomcat, use either the `JAVA_OPTS` or `CATALINA_OPTS` environment variable, e.g:

```
CATALINA_OPTS=-Xmx1024m -Xms256m
```

OR

```
JAVA_OPTS=-Xmx1024m -Xms256m
```

The above example sets the maximum Java Heap memory to 1024MB. They also allow the memory allocation to reset to a smaller number, if Tomcat determines that is necessary. In our experience Tomcat makes good memory allocation decisions, if you allow it to do so.

Difference between `JAVA_OPTS` and `CATALINA_OPTS`



You can use either environment variable. `JAVA_OPTS` is also used by other Java programs (besides just Tomcat). `CATALINA_OPTS` is *only used* by Tomcat. So, if you only want to tweak the memory available to Tomcat, it is recommended that you use `CATALINA_OPTS`. If you set **both** `CATALINA_OPTS` and `JAVA_OPTS`, Tomcat will default to using the settings in `CATALINA_OPTS`.

If the machine is dedicated to DSpace a decent rule of thumb is to give tomcat half of the memory on your machine. **At a minimum, you should give Tomcat \geq 512MB of memory for optimal DSpace operation.** (NOTE: As your DSpace instance gets larger in size, you may need to increase this number to the several GB range.) The latest guidance is to also set `-Xms` to the same value as `-Xmx` for server applications such as Tomcat.

Give Tomcat More Java PermGen Memory

Java PermGen Memory Recommendations



At the time of writing, DSpace recommends you should give Tomcat \geq 192MB of PermGen Space to ensure optimal DSpace operation.

If you are seeing "java.lang.OutOfMemoryError: PermGen space" errors, this is a sure sign that Tomcat is running out PermGen Memory. (More info on PermGen Space: http://blogs.sun.com/fkieviet/entry/classloader_leaks_the_dreaded_java)

To increase the amount of PermGen memory available to Tomcat (default=64MB), use either the `JAVA_OPTS` or `CATALINA_OPTS` environment variable, e.g:


```
CATALINA_OPTS=-XX:MaxPermSize=192m
```

OR


```
JAVA_OPTS=-XX:MaxPermSize=192m
```

The above example sets the maximum PermGen memory to 192MB.

Difference between JAVA_OPTS and CATALINA_OPTS

 You can use either environment variable. JAVA_OPTS is also used by other Java programs (besides just Tomcat). CATALINA_OPTS is *only used* by Tomcat. So, if you only want to tweak the memory available to Tomcat, it is recommended that you use CATALINA_OPTS. If you set **both** CATALINA_OPTS and JAVA_OPTS, Tomcat will default to using the settings in CATALINA_OPTS.

Please note that you can obviously set **both** Tomcat's Heap space and PermGen Space together similar to:

 CATALINA_OPTS=-Xmx1536m -Xms256m -XX:MaxPermSize=192m

On an Ubuntu machine (10.04) at least, the file /etc/default/tomcat6 appears to be the best place to put these environmental variables.

Use the Parallel Garbage Collector

This is a fairly complicated performance tuning subject, for more information, see this article on [performance tuning Tomcat](#). However, a number of large repositories have enabled the parallel garbage collector, and it seems to help with throughput. If you find you have to increase your heap memory allocation very high to maintain stability, you may wish to consider enabling parallel garbage collection as well.

```
JAVA_OPTS="$JAVA_OPTS -XX:+UseParallelGC"
```

Give the Command Line Tools More Memory

Give the Command Line Tools More Java Heap Memory

Similar to Tomcat, you may also need to give the DSpace Java-based command-line tools more Java Heap memory. If you are seeing "java.lang.OutOfMemoryError: Java heap space" errors, when running a command-line tool, this is a sure sign that it isn't being provided with enough Heap Memory.

By default, DSpace only provides 256MB of maximum heap memory to its command-line tools.

If you'd like to provide **more** memory to command-line tools, you can do so via the JAVA_OPTS environment variable (which is used by the [dspace]/bin/dspace script). Again, it's the same syntax as above:

```
JAVA_OPTS=-Xmx512m -Xms512m
```

This is especially useful for big batch jobs, which may require additional memory.

You can also edit the [dspace]/bin/dspace script and add the environmental variables to the script directly.

Give the Command Line Tools More Java PermGen Space Memory

Similar to Tomcat, you may also need to give the DSpace Java-based command-line tools more PermGen Space. If you are seeing "java.lang.OutOfMemoryError: PermGen space" errors, when running a command-line tool, this is a sure sign that it isn't being provided with enough PermGen Space.


By default, Java only provides 64MB of maximum PermGen space.

If you'd like to provide **more** PermGen Space to command-line tools, you can do so via the JAVA_OPTS environment variable (which is used by the [dspace]/bin/dspace script). Again, it's the same syntax as above:

```
JAVA_OPTS=-XX:MaxPermSize=128m
```

This is especially useful for big batch jobs, which may require additional memory.


Please note that you can obviously set **both** Java's Heap space and PermGen Space together similar to:

 JAVA_OPTS=-Xmx512m -Xms512m -XX:MaxPermSize=128m

Give PostgreSQL Database More Memory

On many linux distros PostgreSQL comes out of the box with an incredibly conservative configuration - it uses only 8Mb of memory! To put some more fire in its belly edit the shared_buffers parameter in postgresql.conf. The memory usage is 8KB multiplied by this value. The advice in the Postgres docs is not to increase it above 1/3 of the memory on your machine.

For More PostgreSQL Tips

 For more hints/tips with PostgreSQL configurations and performance tuning, see also:

- [PostgresPerformanceTuning](#)
- [PostgresqlConfiguration](#)

SOLR Statistics Performance Tuning

[This @mire article](#) covers two different methods to enhance performance for the SOLR statistics, that are part of DSpace 1.6 and newer versions.

Note that the **Auto Commit** method is already integrated in DSpace 1.7 and above.