

Description

Memory usage and settings are described here.

`memory` displays a report on Stata's current memory usage.

`query memory` displays the current values of Stata's memory settings.

`set maxvar`, `set niceness`, `set min_memory`, `set max_memory`, and `set segmentsize` change the values of the memory settings.

If you are a Unix user, see [Serious bug in Linux OS](#) under *Remarks and examples* below.

Quick start

Display memory usage report

```
memory
```

Display memory settings

```
query memory
```

Increase the maximum number of variables to 8,000 in Stata/MP or Stata/SE

```
set maxvar 8000
```

Set maximum memory allocation to avoid potential memory allocation bug in Linux

```
set max_memory 16g, permanently
```

Syntax

Display memory usage report

```
memory
```

Display memory settings

```
query memory
```

Modify memory settings

```
set maxvar      # [ , permanently ]
set niceness    # [ , permanently ]
set min_memory  amt [ , permanently ]
set max_memory  amt [ , permanently ]
set segmentsize amt [ , permanently ]
```

where *amt* is #[b | k | m | g], and the default unit is b.

Parameter	Default	Minimum	Maximum	
maxvar	5000	2048	120000	(MP)
	5000	2048	32767	(SE)
	2048	2048	2048	(BE)
niceness	5	0	10	
min_memory	0	0	max_memory	
max_memory	.	2×segmentsize	.	
segmentsize	32m	1m	32g	(64-bit)

Notes:

1. The maximum number of variables in your dataset is limited to `maxvar`. The default value of `maxvar` is 5,000 for Stata/MP and Stata/SE, and 2,048 for Stata/BE. With Stata/MP and Stata/SE, this default value may be increased by using `set maxvar`. The default value is fixed for Stata/BE.
2. Most users do not need to read beyond this point. Stata’s memory management is completely automatic. If, however, you are using the Linux operating system, see [Serious bug in Linux OS](#) under *Remarks and examples* below.
3. The maximum number of observations is fixed at 1,099,511,627,775 for Stata/MP and is fixed at 2,147,483,619 for Stata/SE and Stata/BE regardless of computer size or memory settings. Depending on the amount of memory on your computer, you may face a lower practical limit. See `help obs_advice`.

4. `max_memory` specifies the maximum amount of memory Stata can use to store your data. The default of missing (.) means all the memory the operating system is willing to supply. There are three reasons to change the value from missing to a finite number.
 1. You are a Linux user; see [Serious bug in Linux OS](#) under *Remarks and examples* below.
 2. You wish to reduce the chances of accidents, such as typing `expand 100000` with a large dataset in memory and actually having Stata do it. You would rather see an insufficient-memory error message. Set `max_memory` to the amount of physical memory on your computer or more than that if you are willing to use virtual memory.
 3. You are a system administrator; see [Notes for system administrators](#) under *Remarks and examples* below.
5. The remaining memory parameters—`nice ness`, `min_memory`, and `segment_size`—affect efficiency only; they do not affect the size of datasets you can analyze.
6. Memory amounts for `min_memory`, `max_memory`, and `segment size` may be specified in bytes, kilobytes, megabytes, or gigabytes; suffix b, k, m, or g to the end of the number. The following are equivalent ways of specifying 1 gigabyte:

```
1073741824
1048576k
1024m
1g
```

Suffix k is defined as (multiply by) 1024, m is defined as 1024^2 , and g is defined as 1024^3 .

7. 64-bit computers can theoretically provide up to 18,446,744,073,709,551,616 bytes of memory, equivalent to 17,179,869,184 gigabytes, 16,777,216 terabytes, 16,384 petabytes, or 16 exabytes. Real computers have less.
8. Stata allocates memory for data in units of `segment size`. Smaller values of `segment size` can result in more efficient use of available memory but require Stata to jump around more. The default provides a good balance. We recommend resetting `segment size` only if your computer has large amounts of memory.
9. If you have large amounts of memory and you use it to process large datasets, you may wish to increase `segment size`. Suggested values are

memory	segment size
32g	64m
64g	128m
128g	256m
256g	512m
512g	1g
1024g	2g

10. `nice ness` affects how soon Stata gives back unused segments to the operating system. If Stata releases them too soon, it often needs to turn around and get them right back. If Stata waits too long, Stata is consuming memory that it is not using. One reason to give memory back is to be nice to other users on multiuser systems or to be nice to yourself if you are running other processes.

The default value of 5 is defined to provide good performance. Waiting times are currently defined as

niceness	waiting time (m:s)
10	0:00.000
9	0:00.125
8	0:00.500
7	0:01
6	0:30
5	1:00
4	5:00
3	10:00
2	15:00
1	20:00
0	30:00

Niceness 10 corresponds to being totally nice. Niceness 0 corresponds to being an inconsiderate, self-centered, totally selfish jerk.

11. `min_memory` specifies an amount of memory Stata will not fall below. For instance, you have a long do-file. You know that late in the do-file, you will need 8 gigabytes. You want to ensure that the memory will be available later. At the start of your do-file, you set `min_memory 8g`.
12. Concerning `min_memory` and `max_memory`, be aware that Stata allocates memory in `segmentsize` blocks. Both `min_memory` and `max_memory` are rounded down. Thus the actual minimum memory Stata will reserve will be

`segmentsize*trunc(min_memory/segmentsize)`

The effective maximum memory is calculated similarly. (Stata does not round up `min_memory` because some users set `min_memory` equal to `max_memory`.)

Options

`permanently` specifies that, in addition to making the change right now, the new limit be remembered and become the default setting when you invoke Stata.

`once` is not shown in the syntax diagram but is allowed with `set niceness`, `set min_memory`, `set max_memory`, and `set segmentsize`. It is for use by system administrators; see [Notes for system administrators](#) under *Remarks and examples* below.

Remarks and examples

Remarks are presented under the following headings:

Examples

Serious bug in Linux OS

Notes for system administrators

Examples

Here is our memory-usage report after we load `auto.dta` that comes with Stata using Stata/MP:

```
. sysuse auto
(1978 automobile data)

. memory
```

Memory usage	Used	Allocated
Data	3,182	100,663,296
strLs	0	0
Data & strLs	3,182	100,663,296
Data & strLs	3,182	100,663,296
Variable names, %fmts, ...	4,178	396,279
Overhead	1,081,344	1,082,136
Stata matrices	0	0
ado-files	53,718	53,718
Stored results	0	0
Mata matrices	10,880	10,880
Mata functions	2,720	2,720
set maxvar usage	4,636,521	4,636,521
Other	3,497	3,497
Total	5,773,999	106,849,047

We could then obtain the current memory-settings report by typing

```
. query memory
```

Memory settings		
set maxvar	5000	2048-120000; max. vars allowed
set niceness	5	0-10
set min_memory	0	0-1600g
set max_memory	.	32m-1600g or .
set segmentsize	32m	1m-32g
set adosize	1000	kilobytes
set max_preservemem	1g	0-1600g

Serious bug in Linux OS

If you use Linux OS, we strongly suggest that you set `max_memory`. Here's why:

“By default, Linux follows an optimistic memory allocation strategy. This means that when `malloc()` returns non-NULL there is no guarantee that the memory really is available. This is a really bad bug. In case it turns out that the system is out of memory, one or more processes will be killed by the infamous OOM killer. In case Linux is employed under circumstances where it would be less desirable to suddenly lose some randomly picked processes, and moreover the kernel version is sufficiently recent, one can switch off this overcommitting behavior using [...]”

— Output from Unix command `man malloc`.

What this means is that Stata requests memory from Linux, Linux says yes, and then later when Stata uses that memory, the memory might not be available and Linux crashes Stata, or worse. The Linux documentation writer exercised admirable restraint. This bug can cause Linux itself to crash. It is easy.

The proponents of this behavior call it “optimistic memory allocation”. We will, like the documentation writer, refer to it as a bug.

The bug is fixable. Type `man malloc` at the Unix prompt for instructions. Note that `man malloc` is an instruction of Unix, not Stata. If the bug is not mentioned, perhaps it has been fixed. Before assuming that, we suggest using a search engine to search for “linux optimistic memory allocation”.

Alternatively, Stata can live with the bug if you set `max_memory`. Find out how much physical memory is on your computer and set `max_memory` to that. If you want to use virtual memory, you might set it larger, just make sure your Linux system can provide the requested memory. Specify the option permanently so you only need to do this once. For example,

```
. set max_memory 16g, permanently
```

Doing this does not guarantee that the bug does not bite, but it makes it unlikely.

Notes for system administrators

System administrators can set `max_memory`, `min_memory`, and `nice` so that Stata users cannot change them. They can also do this with `max_preservemem` (see [P] [preserve](#)). You may want to do this on shared computers to prevent individual users from hogging resources.

There is no reason you would want to do this on users’ personal computers.

You can also set `segmentsize`, but there is no reason to do this even on shared systems.

The instructions are to create (or edit) the text file `sysprofile.do` in the directory where the Stata executable resides. Add the lines

```
set min_memory 0, once
set max_memory 16g, once
set nice 5, once
```

The file must be plain text, and there must be end-of-line characters at the end of each line, including the last line. Blank lines at the end are recommended.

The 16g on `set max_memory` is merely for example. Choose an appropriate number.

The values of 0 for `min_memory` and 5 for `nice` are recommended.

Stored results

`memory` stores all reported numbers in `r()`. StataCorp may change what memory reports, and you should not expect the same `r()` results to exist in future versions of Stata. To see the stored results from `memory`, type `return list, all`.

Also see

[P] [creturn](#) — Return c-class values

[R] [query](#) — Display system parameters

[U] 6 Managing memory

Stata, Stata Press, Mata, NetCourse, and NetCourseNow are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow is a trademark of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2025 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on [citing Stata documentation](#).