Logstat

SEASON 9 – FOURTH ROUND

The latest project in which Musala got involved (thanks GDPR) is a tool, that without remembering or using any kind of stored information about our site's users, calculates with good precision their count.

We will give you a giant list of IPv6 IPs, but you will only have **4 MiB memory limit**. With at most 10% difference from the correct answer, you must **find how many distinct IPs** there are among them. That means that if the correct answer is **X**, any number between **0.9*X** and **1.1*X** would be considered good enough and pass the test.

Input

From the input file `logstat.in`, IPv6 IPs are entered, one per line. A valid IPv6 consists of 8 blocks made of 2 bytes written in hex, separated by "`:`". The end of the input is marked by the line "`xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`".

Output

In the output file `logstat.out` print the number of distinct IPs in the input file.

Constraints

$$10\,000 \leq unique\ IPs \leq total\ number\ of\ IPs \leq 2\,000\,000$$

Time limit: 2.0 seconds
Memory limit: 4 MiB

Logstat

SEASON 9 – FOURTH ROUND

Example

Both tests are **not** complying with the constraints. They are just for demonstration purposes.

Input (logstat.in)	Output (logstat.out)
6e0a:c679:7155:ea15:8eac:dcf2:a18c:b7f3 7899:3b9c:8652:68be:a4c0:8150:a348:4c2d 42ab:a8b6:f093:66b6:7427:a38f:c105:c6ba 6e0a:c679:7155:ea15:8eac:dcf2:a18c:b7f3 2b73:1aca:7e4b:7833:c095:cc4c:4fa1:f59b 665a:01be:b643:ddb4:2bc8:ecc2:111c:dcd7 2b73:1aca:7e4b:7833:c095:cc4c:4fa1:f59b 080b:e1a8:c23a:1895:9485:3515:02c8:f7c3 b1d7:4f04:60c8:7bb0:cbe6:a7ab:453b:a61d 2b73:1aca:7e4b:7833:c095:cc4c:4fa1:f59b 665a:01be:b643:ddb4:2bc8:ecc2:111c:dcd7 42ab:a8b6:f093:66b6:7427:a38f:c105:c6ba 665a:01be:b643:ddb4:2bc8:ecc2:111c:dcd7 9d4a:945e:3700:4f08:71c6:09a4:c38a:3d71 b1d7:4f04:60c8:7bb0:cbe6:a7ab:453b:a61d 8038:73ff:522e:59d5:9cf3:feb1:53ca:436b 5730:02fd:4433:6bfb:c435:bbfe:2422:7ff4 8038:73ff:522e:59d5:9cf3:feb1:53ca:436b 665a:01be:b643:ddb4:2bc8:ecc2:111c:dcd7 8038:73ff:522e:59d5:9cf3:feb1:53ca:436b xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx	10

The correct answer is 10.

10% leeway allows you to output 9 or 11 and still pass the test.

Input (logstat.in)	Output (logstat.out)
a9a3:8967:94bd:7c45:06f0:9c62:4a0a:c445 a9a3:8967:94b9:7c45:16f0:fcb2:4a0a:c445 a9a3:8967:94bd:7d45:06f5:9cd2:f70a:c445 a9b3:8d67:94bb:7c45:06f0:9c64:4a0a:1445 a9a3:8967:94bc:79e5:06f0:9c62:4a3c:c445 a9a3:99a7:94bd:7345:16f0:9c62:4a0a:c545 a9a3:596f:947d:7c45:06d0:9c62:4ada:c445 afa3:8967:94bd:7c45:a6f0:4c62:4aba:c44d a9c3:8967:98b5:7c4f:06f0:9cf2:4a0a:c445 49aa:8967:94bd:7c47:06f0:9662:4a03:c445 a9a3:8967:9ebd:7c45:0650:ac62:4a21:c445 a9a3:8967:94fd:7c45:06f0:df68:4a0a:c245 a9a3:8967:94cd:7c45:09f1:9c62:4a0a:ce45 afa3:8967:94bd:7c45:a6f0:4c62:4aba:c44d a9a3:8967:94bd:7c45:06f0:9c62:4a0a:c445 a9a3:8967:94cd:7c45:09f1:9c62:4a0a:ce45 a9a3:8967:94bc:79e5:06f0:9c62:4a3c:c445 49aa:8967:94bd:7c47:06f0:9662:4a03:c445 a9a3:8967:94cd:7c45:09f1:9c62:4a0a:ce45 a9a3:8967:9ebd:7c45:0650:ac62:4a21:c445 xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx	13