Help me optimize this code which enumerates all possible GUIDs

devblogs.microsoft.com/oldnewthing/20131029-00

October 29, 2013

Raymond Chen

Hi, I'm looking for help optimizing this code. It creates a file containing every possible GUID.

```
class Program {
 public static void Main(string[] args)
 {
  using (var w = new System.IO.StreamWriter("all_guids.txt")) {
   for (long l1 = 0; l1 <= 0xFFFFFFFFL; l1++)
   for (long l2 = 0; l2 <= 0xFFFFFFFFL; l2++)
   for (long l3 = 0; l3 <= 0xFFFFFFFFL; l3++)
   for (long l4 = 0; l4 <= 0xFFFFFFFFL; l4++)
    System.Console.WriteLine(
     new System.Guid(string.Format("{0:X8}{1:X8}{2:X8}{3:X8}",
                              l1, l2, l3, l4)));
  }
 }
}
```

I know this code will take a really long time to run, so any performance improvements would be appreciated.

Okay, wait a second. You do realize that even if you had a 100GHz machine that could generate a new GUID every cycle, it will still take $2^{128} \div 10^{11} = 3 \times 10^{27}$ seconds or 10^{19} years to complete.

Right, that's why I was looking for help making the code run faster.

Look at this way: Suppose you could somehow get this algorithm to run a quintillion times faster, so it finishes in under a year. Your output file is going to be $2^{128} \times 16 = 2^{132}$ bytes in size. That's around 10^{27} terabytes. One terabyte of SSD storage weighs around 100 grams. The mass of the earth is 10^{24} kilograms. Therefore, before you run this program, you will need to acquire 100 earth-sized planets and convert them all to SSDs.

What are you trying to do with all of these GUIDs?

There's a Web site that uses GUIDs to identify resources. I want to enumerate all the GUIDs so I can request each one from the Web site to see if it contains anything interesting.

Okay, so it doesn't matter how fast your GUID enumeration algorithm is, because that's not the rate-limiting factor in this entire exercise. The limiting factor is the speed of the Web server you are attacking. (And at this point, it's fair to say that what you're doing is *attacking* a Web server.) Even if the Web server can handle a billion hits per second, it will take 10^{22} years for it to respond to all your requests.

Looking at it another way: Suppose you could enlist the resources of *every computer on the Internet* to send a million requests per second. It would still take over 10^{13} centuries for them to try every GUID.

And the server administrators might suspect something is up after, oh, about seven centuries.

This problem requires a sociological, not engineering, solution. You need to contact the site administrators, describe what information you would like, and see if you can come to some sort of data-sharing agreement.

I've contacted the site owners, but they are not interested in helping me, so I have no choice but to try every GUID.

Your computer is already plenty powerful enough to generate GUIDs as fast as the Web site can handle inbound requests. Maybe you can upgrade your cable modem to a higher level of service.

Actually, what you need to do is upgrade the *other site's* hardware and Internet connection to a higher level of service. Because they are the bottleneck, not you. "Here's a half million dollars to get faster servers. Just a gift from me. No really, no strings attached. Have fun. Just get those faster servers online right away, okay?"

I wish you good luck.

Raymond Chen

Follow