/dev/!random

Profiling Entropy Collection in the Linux Random Number Generator

Richard Roberts

Justin MacIntosh

Advisor: Dr. Nadia Heninger

Random Numbers

Cryptography

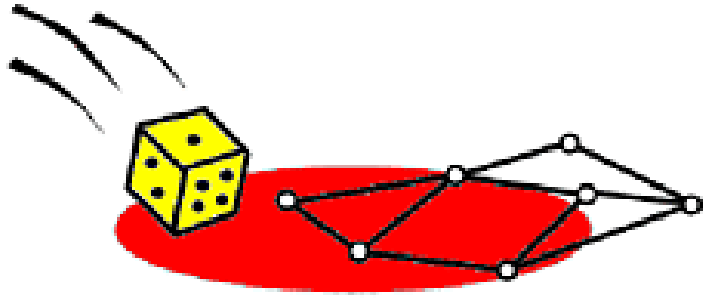

Randomized Algorithms

Web Browsing

Desktop Applications

Generating Random Numbers

Software Requiring Random Numbers

“Any one who considers arithmetical methods of producing random digits is, of course, in a state of sin.” –John von Neumann

Generating Random Numbers

Software Requiring Random Numbers

Operating System

Generating Random Numbers

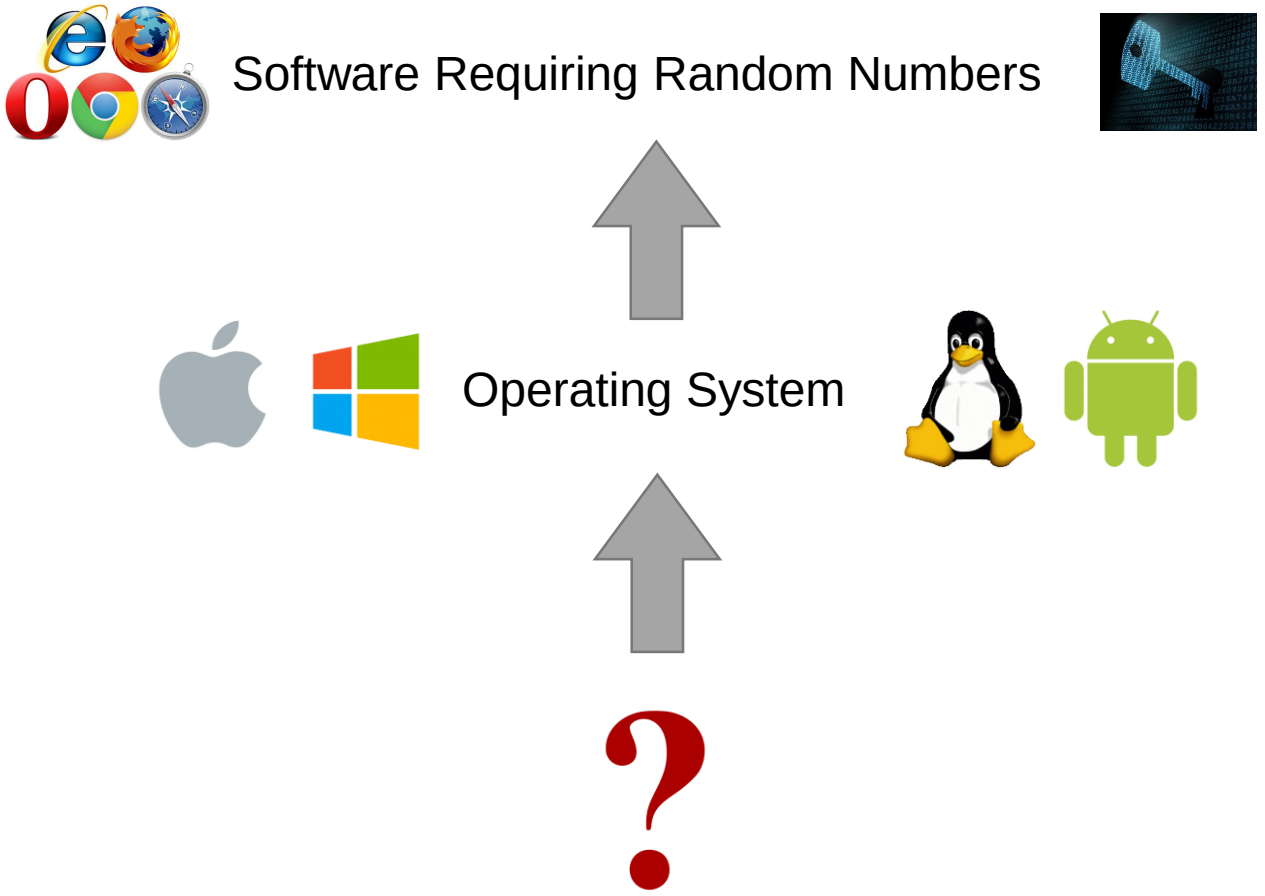

Using a Random Process

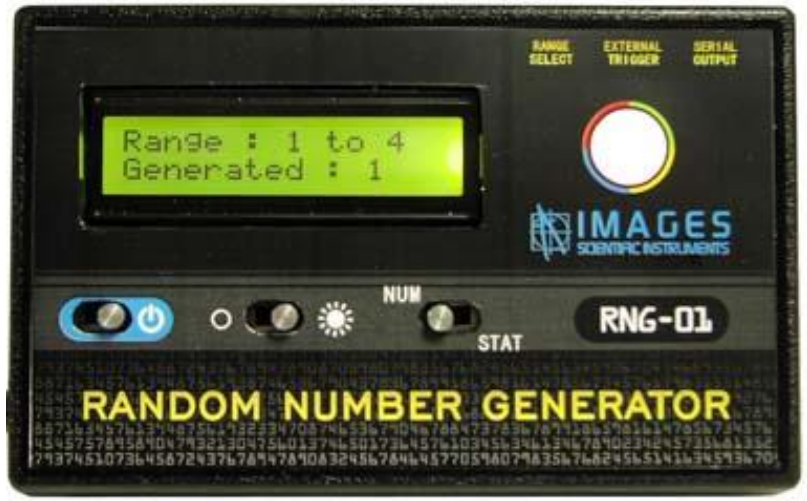

Radioactive Decay Hardware

“Lavarand” Generator

Linux Random Number Generator

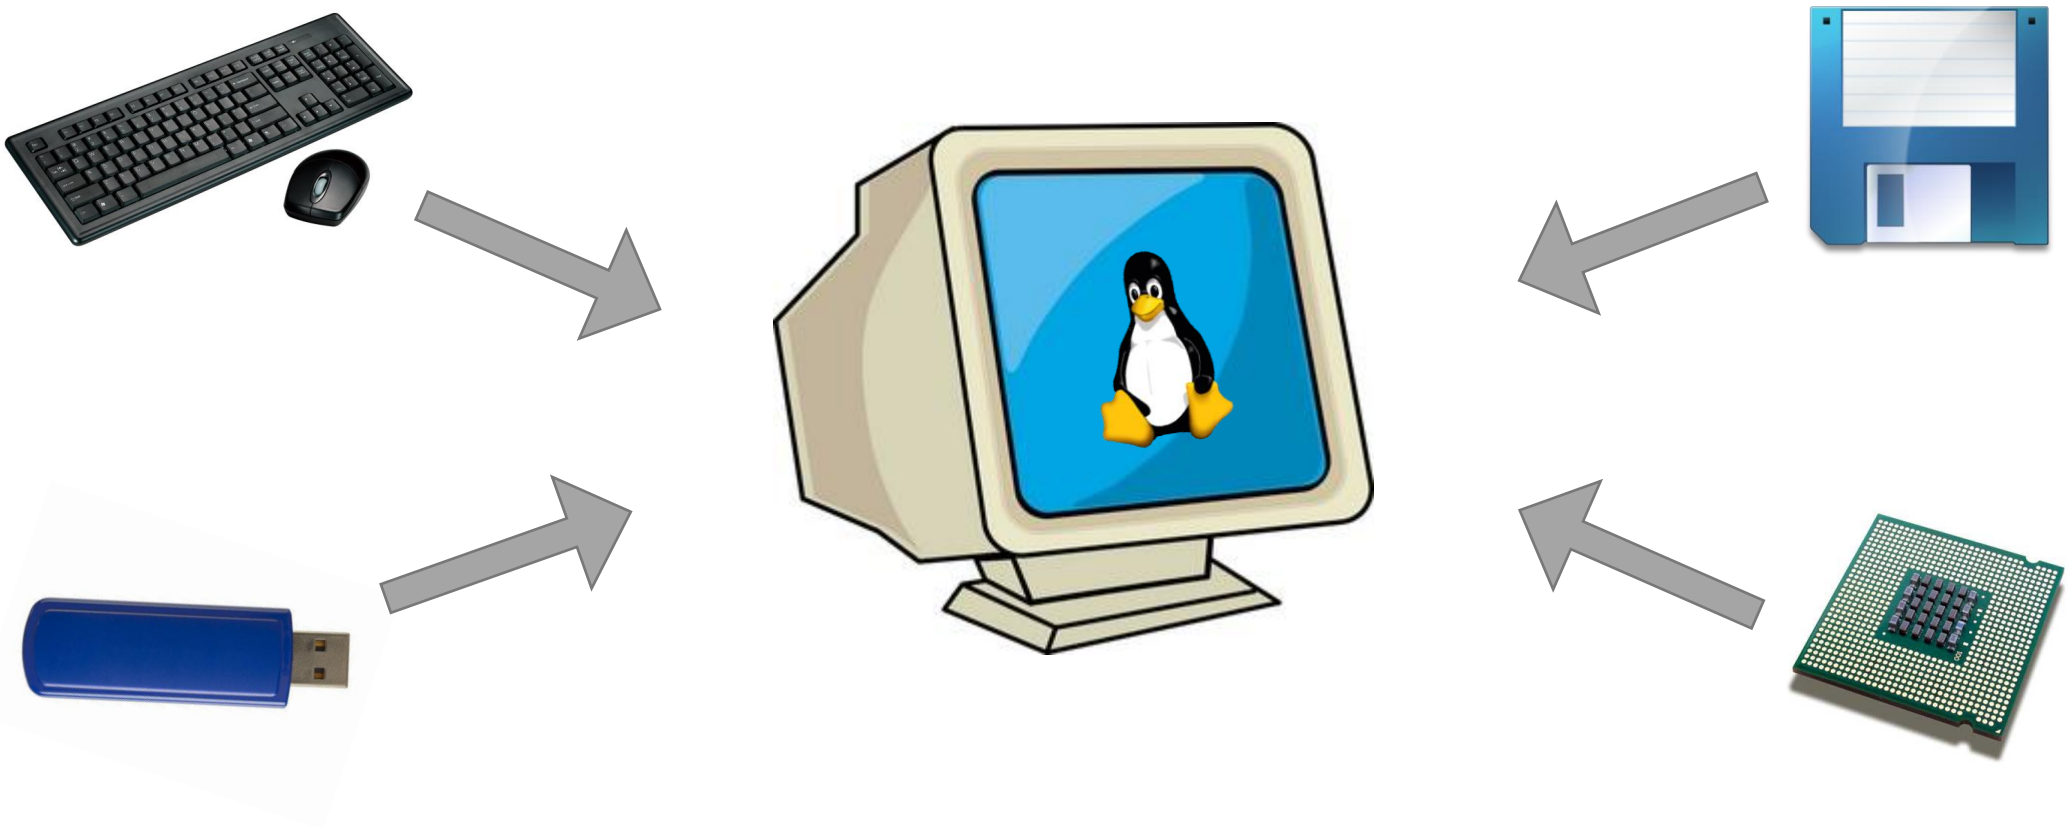

Problems

Tension between academics and developers

Designed Decades Ago

The design is known by academics to be theoretically flawed

Developers resist change on the grounds that no practical attack has been made

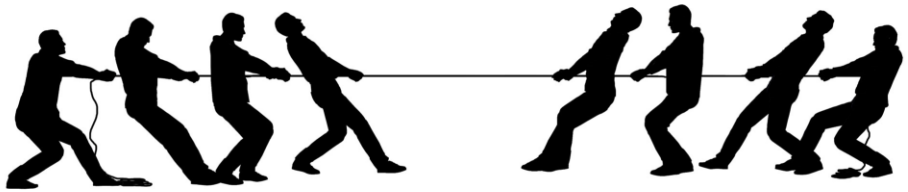

“Mining Your Ps and Qs”*

Flaw in design led to attack during computer startup

Fixed the startup issue, but did not rework overall design

*Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. Mining your Ps and Qs: Detection of widespread weak keys in network devices. In *Proceedings of the 21st USENIX Security Symposium*, August 2012.

Our Project

Measuring the LRNG in Real Use Cases

Linux source code is very complex

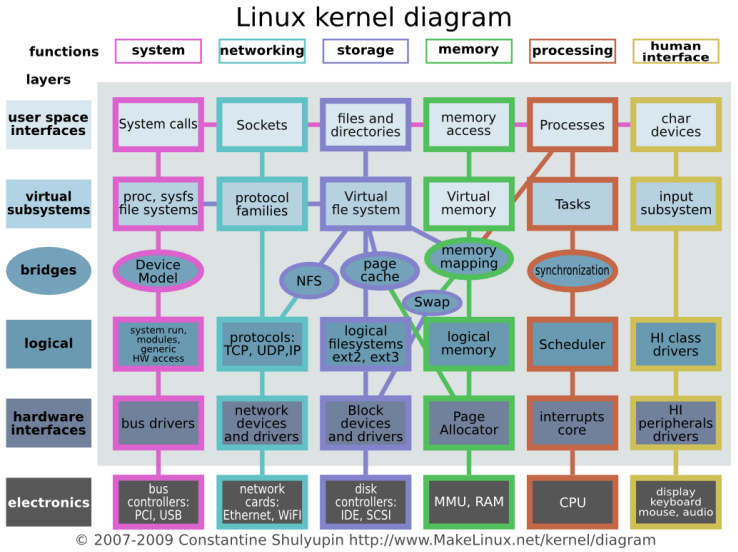

Create a standard testing framework that can be used without a thorough understanding of the source code

Generating Logs

Challenge: The LRNG source code is tightly woven into the operating system source code

Solution: We can create our own modified version of the operating system and install it on a computer

Storing Data

Challenge: We cannot save logs without making events that would need to send more logs

Solution: Send logs to another machine for persistent storage

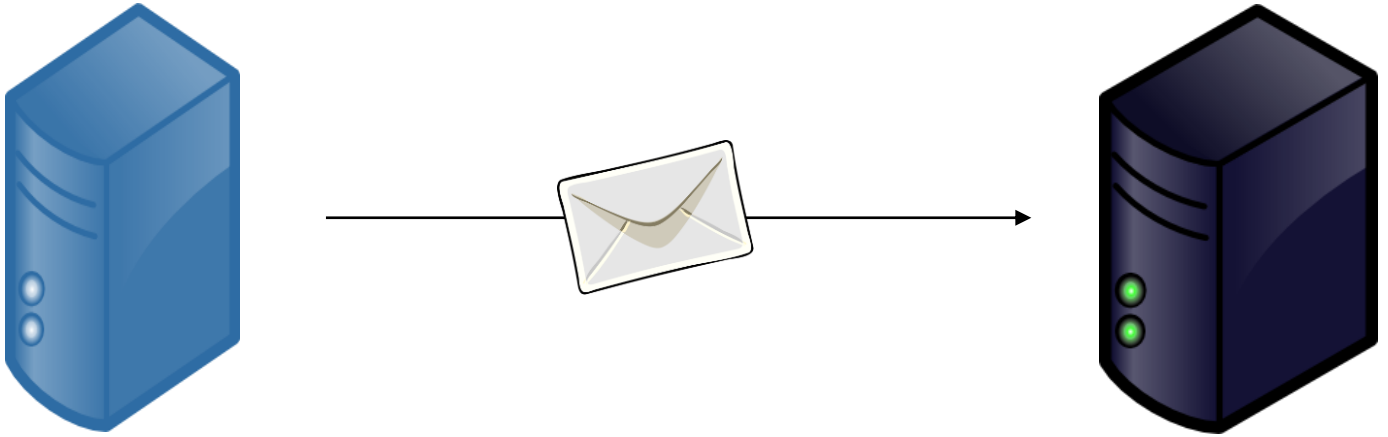

Generation

Collection

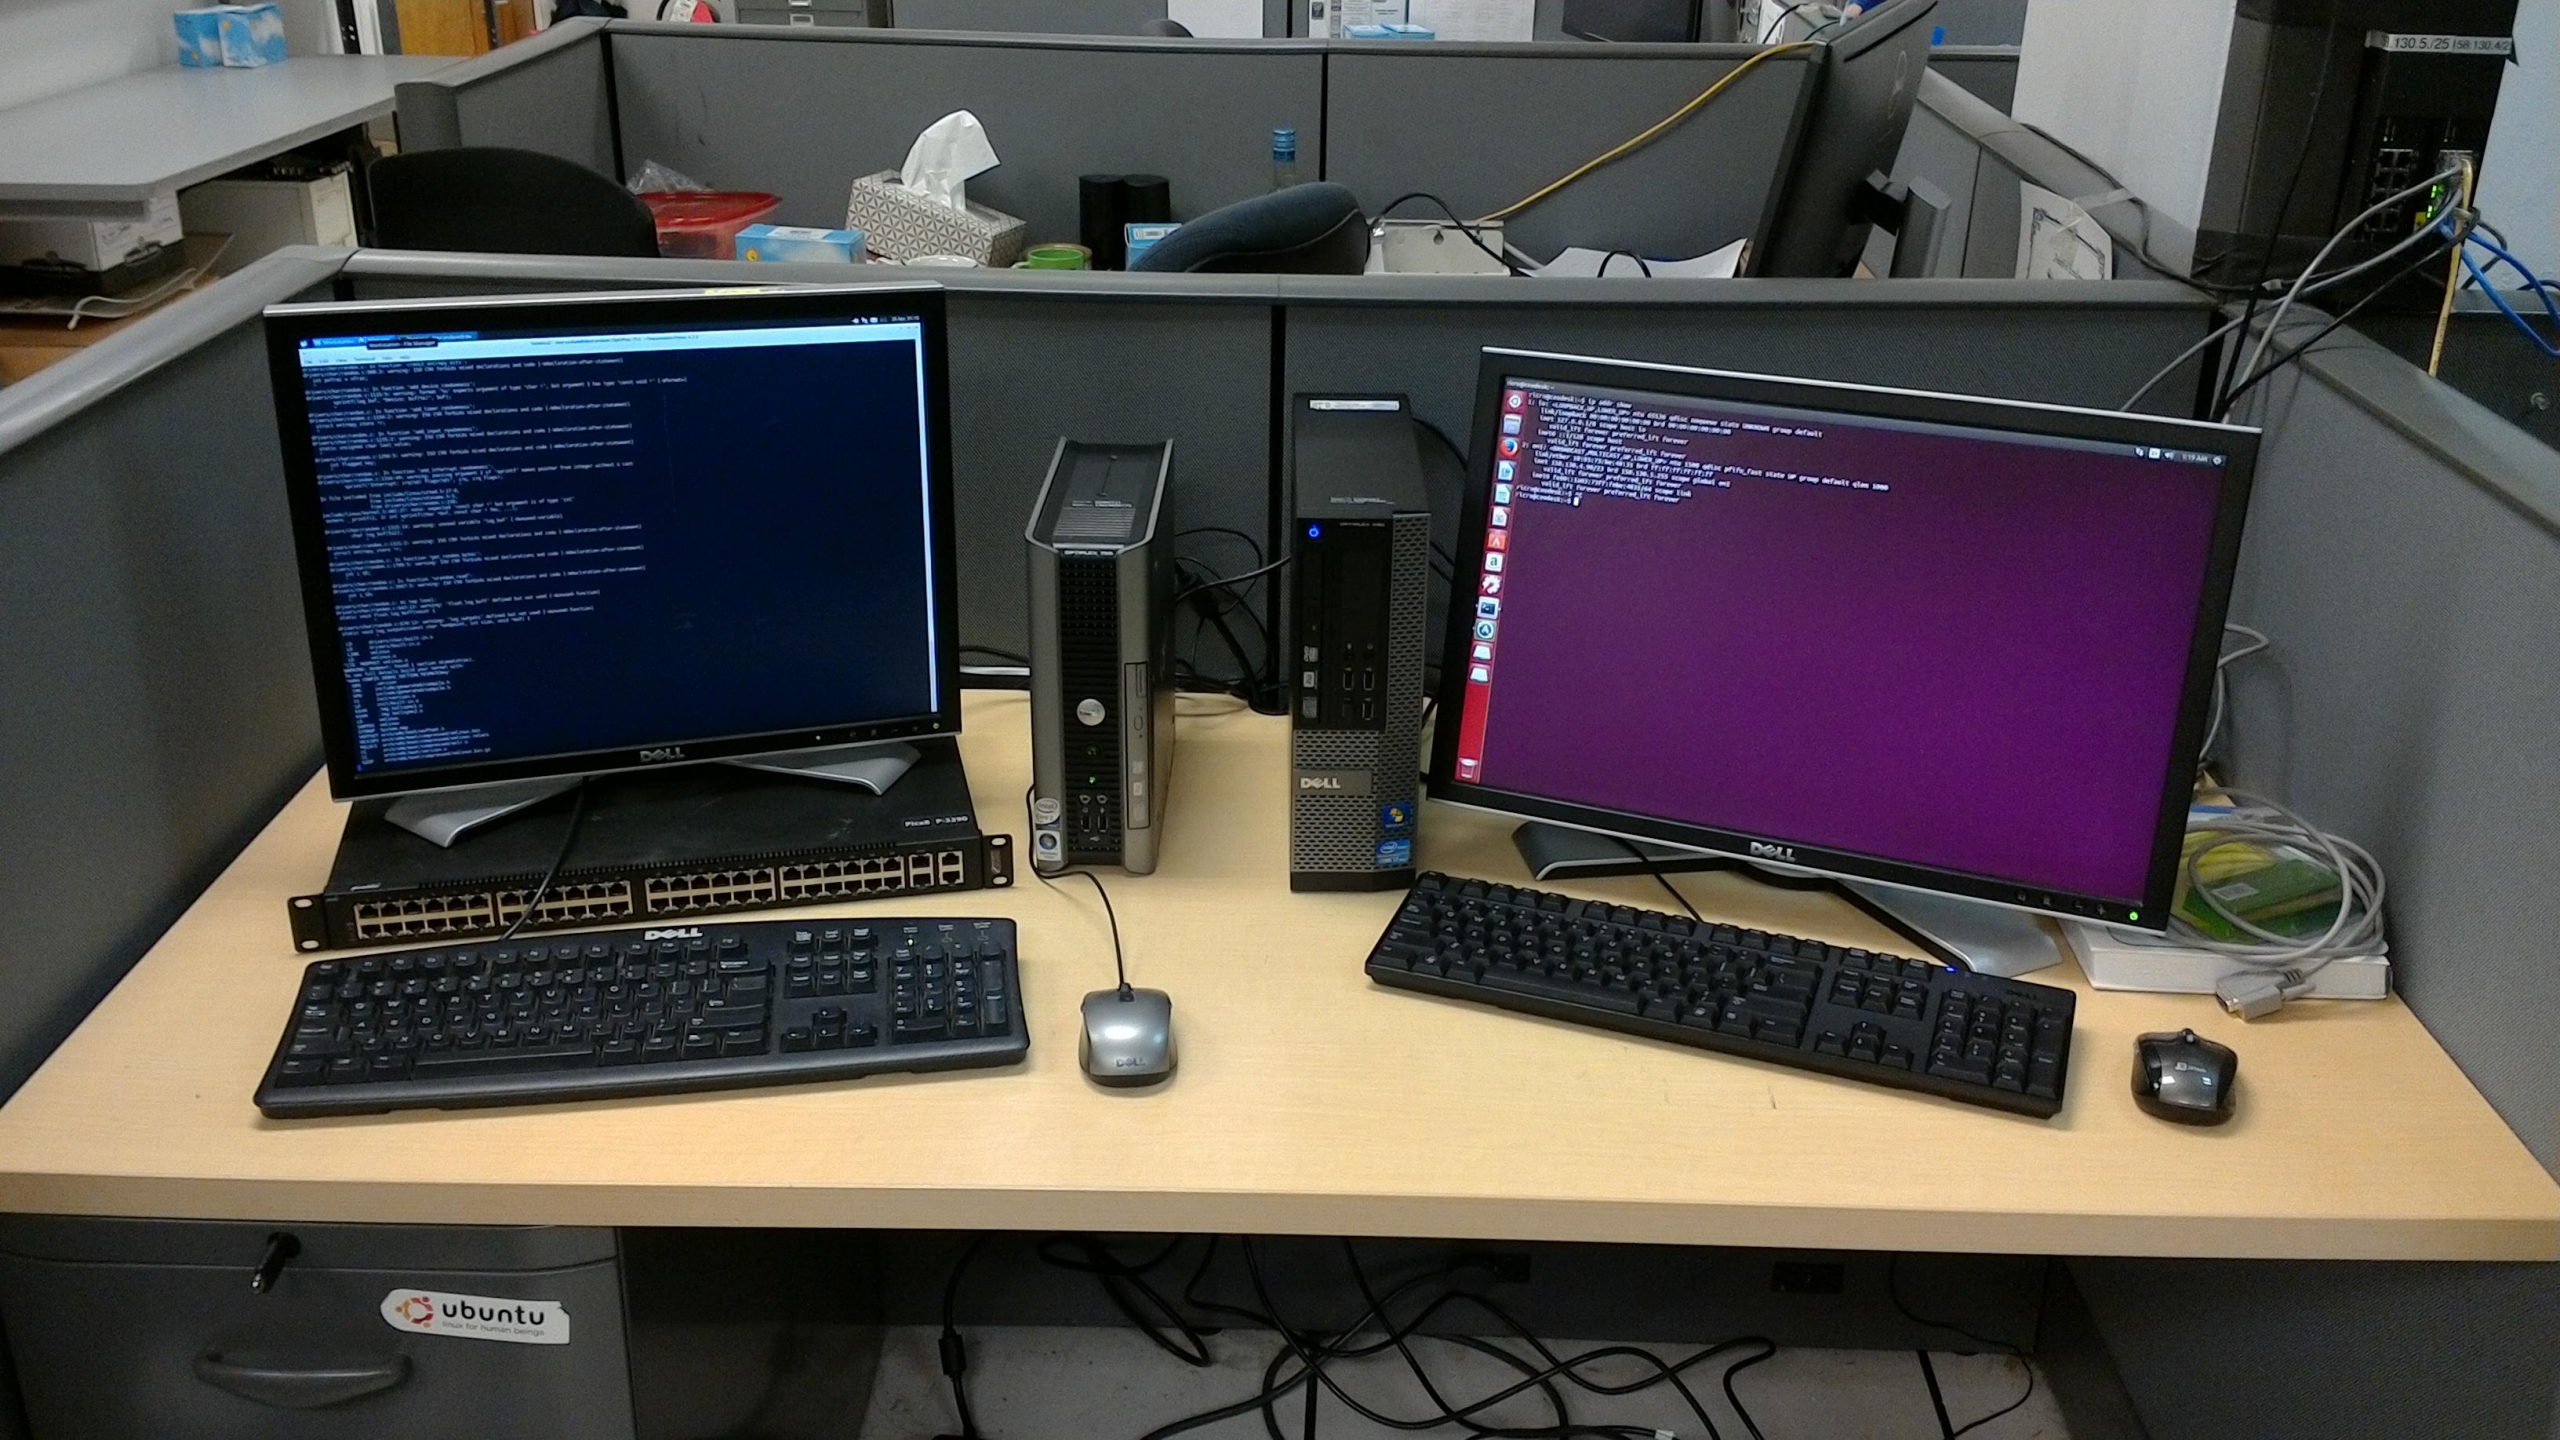
DELL
ubuntu
linux for human beings

```
ricro@ceodesk:~$ sudo tcpdump udp -i em1 port 2048 -Aqs 0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on em1, link-type EN10MB (Ethernet), capture size 65535 bytes
01:42:49.989533 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 274
E.......@.3........Z......./.Input Pool 1: fd03189a5903ae172d6c8110c5e265a4cf58af6131e7db1f922abbebb12cdbb4c6b7b2c5c304fd2f391ee7464bae7bb1d2387a23aa3052f8afa86e77fc0a75032e1beadb644d22bef53a612426dca5f390948c4d8e783d4d111f236e03322d4817243b8e324b6607876121ff27f67b70b35ef19b945f0897ae6d32301d528e5b END
01:42:49.989580 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 274
E.......@.3........Z........Input Pool 2: cc8fba68dd9c31b365446faabbdbcc62be4122b56925db59542d4b076b3c7114e11383ada1026b46a0c986046b450f78c9bbd454155a063fe88df6a89b63061176cbbf56356bd22748a7b94f97e4419eeabc47e1d245d11de46957a2d0de121c46f69123ed2049b997ae35eaa2dd359b0d2929ad2ed0dd1e511c00426dfaae36 END
01:42:49.989599 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 274
E....   ..@.3........Z........Input Pool 3: 55655581ab791b73796a40f7a5b9b90adb3618c57bed281b8d4878f96e3a90e38a8cf1ecc0f3c442697c657a104640c5bed475cac7c7e5c831313cd606a7ff968b63547c52f537409de3b763e0e77c1d4fce4a494bead607b22b1f30b989850fcc168b1c29771a3131ee66502529a0abc54ff3e7bf8fabca1dc5fb590c3ca076 END
01:42:49.989609 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 274
E....
..@.3........Z........Input Pool 4: 2dcb241102779e65190428619efc4485ec425305aa8df549274cd358d5f63049f851b25a40aa627ba2e568e7a3fc428c8bd676bc1fdb706ac009564d3c51e26b53adae995c3fb9139a91161e7c2e0d3df1aa80b4e30f270ef70c05eb7971d192ed0d6f47f7327b01fc8ee3b160cdfde31dd9165bf45c6ab7ce289b6b2d0431db END
01:42:49.989618 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 275
E../....@.3........Z......d.Blocking Pool: a59e89ef4dad679824c22bb729aa4ed15d783d5c9c921b5d2a5d5f0bfc006fd22ee5769f14dfcf6b1c6c5caa9025d4ade8f4965e4ef9e462c89c49a816a071ddf1ecf30a07261ce2b220d5236e61a728152f994204ade5f237ee6a41cb48c16a2ce2cc722860e14d45da6f8ecd3ba0d419a553394165145cd09aa2c1d5669762 END
01:42:49.989622 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 277
E..1....@.3........Z....... Nonlocking Pool: e326ce82479b7354ebd6f2e367285509d1e3b14347bc6ad581f5fb1dbf8748ae05c9a010e475221571e9d41ee1a9f96b9a223952286247bdbc694ead5ceaea3f2ae348164970a6d05c6ce76393542422ee7c2da26afe03b84ecfaa88f917768aa40712fa1d2552a6f229962862784b45eb18f9de4c0c72823a82b525172dd157 END
01:42:56.172341 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
..@.4........Z.....,`.Input: Type(4) Code(4) Value(458842)
01:42:58.212264 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,^~Input: Type(4) Code(4) Value(458775)
01:42:58.284423 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,^~Input: Type(4) Code(4) Value(458775)
01:42:58.404351 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,`.Input: Type(4) Code(4) Value(458763)
01:42:58.508349 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,`.Input: Type(4) Code(4) Value(458763)
01:42:58.588407 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_.Input: Type(4) Code(4) Value(458764)
01:42:58.708390 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_.Input: Type(4) Code(4) Value(458764)
01:42:58.748391 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_~Input: Type(4) Code(4) Value(458774)
01:42:58.828348 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_~Input: Type(4) Code(4) Value(458774)
01:43:00.596172 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_~Input: Type(4) Code(4) Value(458774)
01:43:00.620344 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_.Input: Type(4) Code(4) Value(458764)
01:43:00.676378 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_~Input: Type(4) Code(4) Value(458774)
01:43:00.740336 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,_.Input: Type(4) Code(4) Value(458764)
01:43:02.756368 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,].Input: Type(4) Code(4) Value(458756)
01:43:02.820302 IP devrandom-OptiPlex-755.cis.upenn.edu.2048 > 158.130.4.90.2048: UDP, length 36
E..@....@.4........Z.....,].Input: Type(4) Code(4) Value(458756)
```

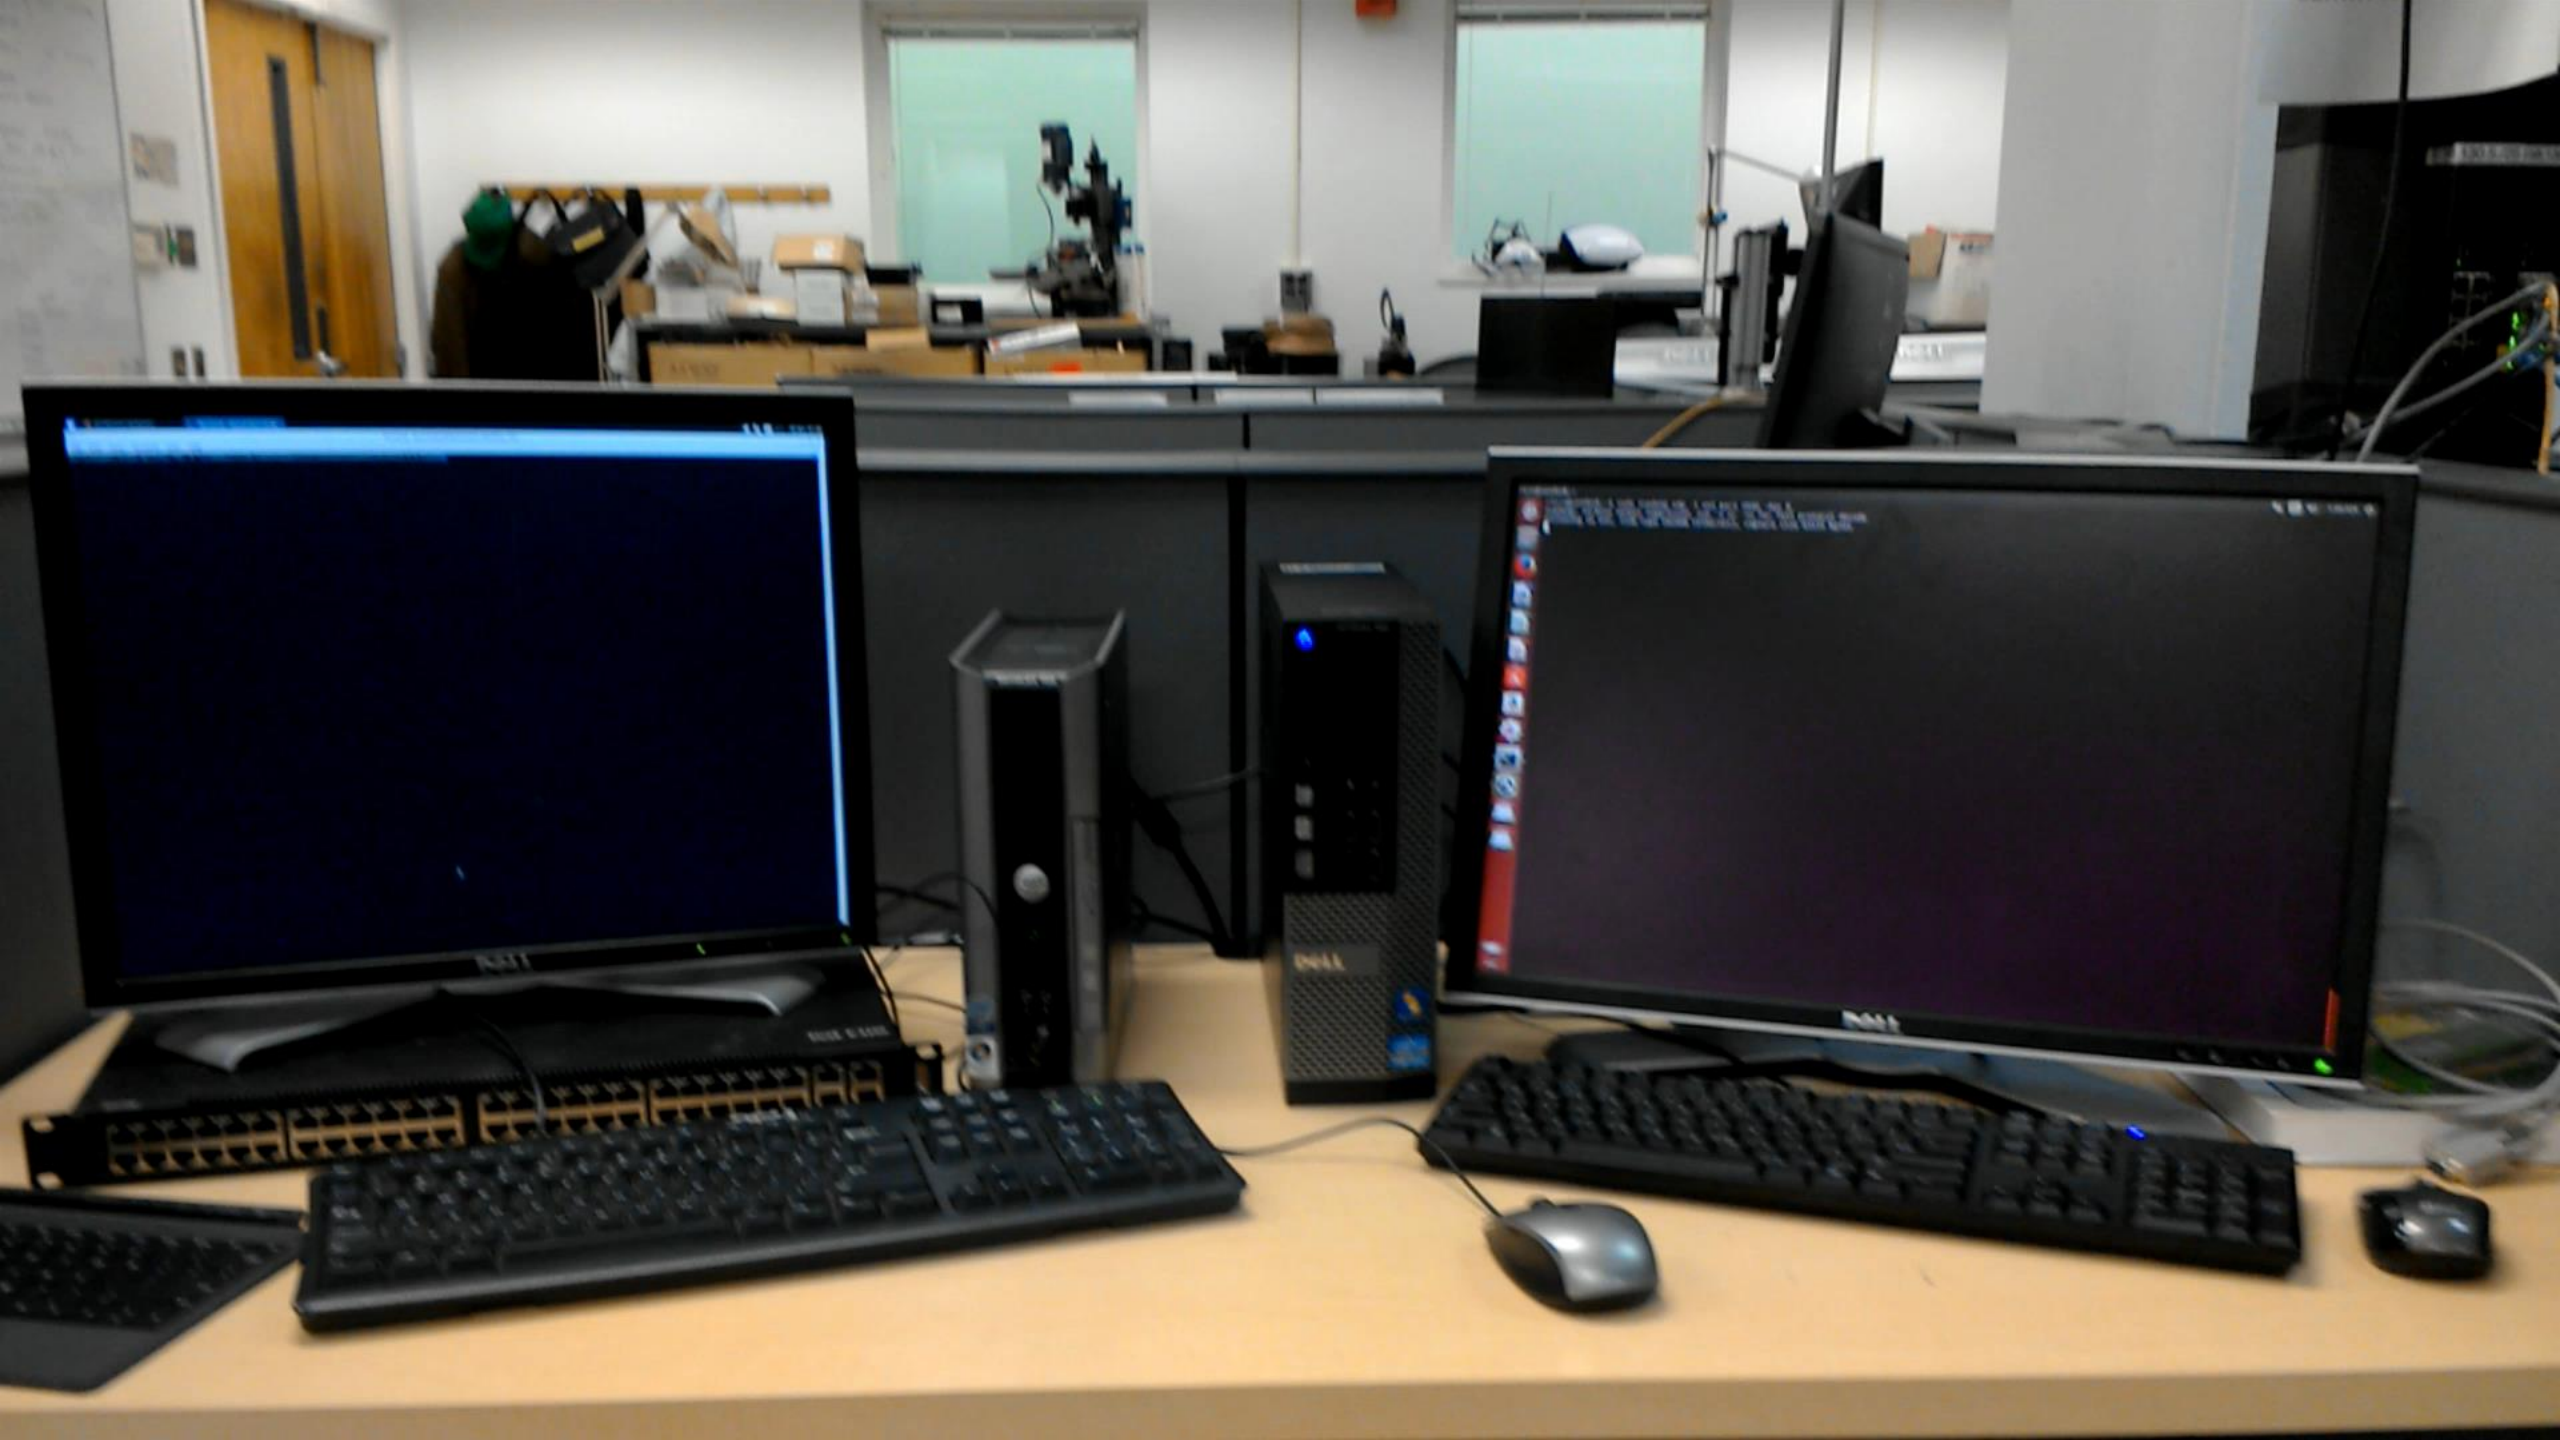

Analysis and Results

Experimentation

Run experiment with “typical computer usage”

Web browsing, desktop applications, etc.

We are among the first researchers to gather this data

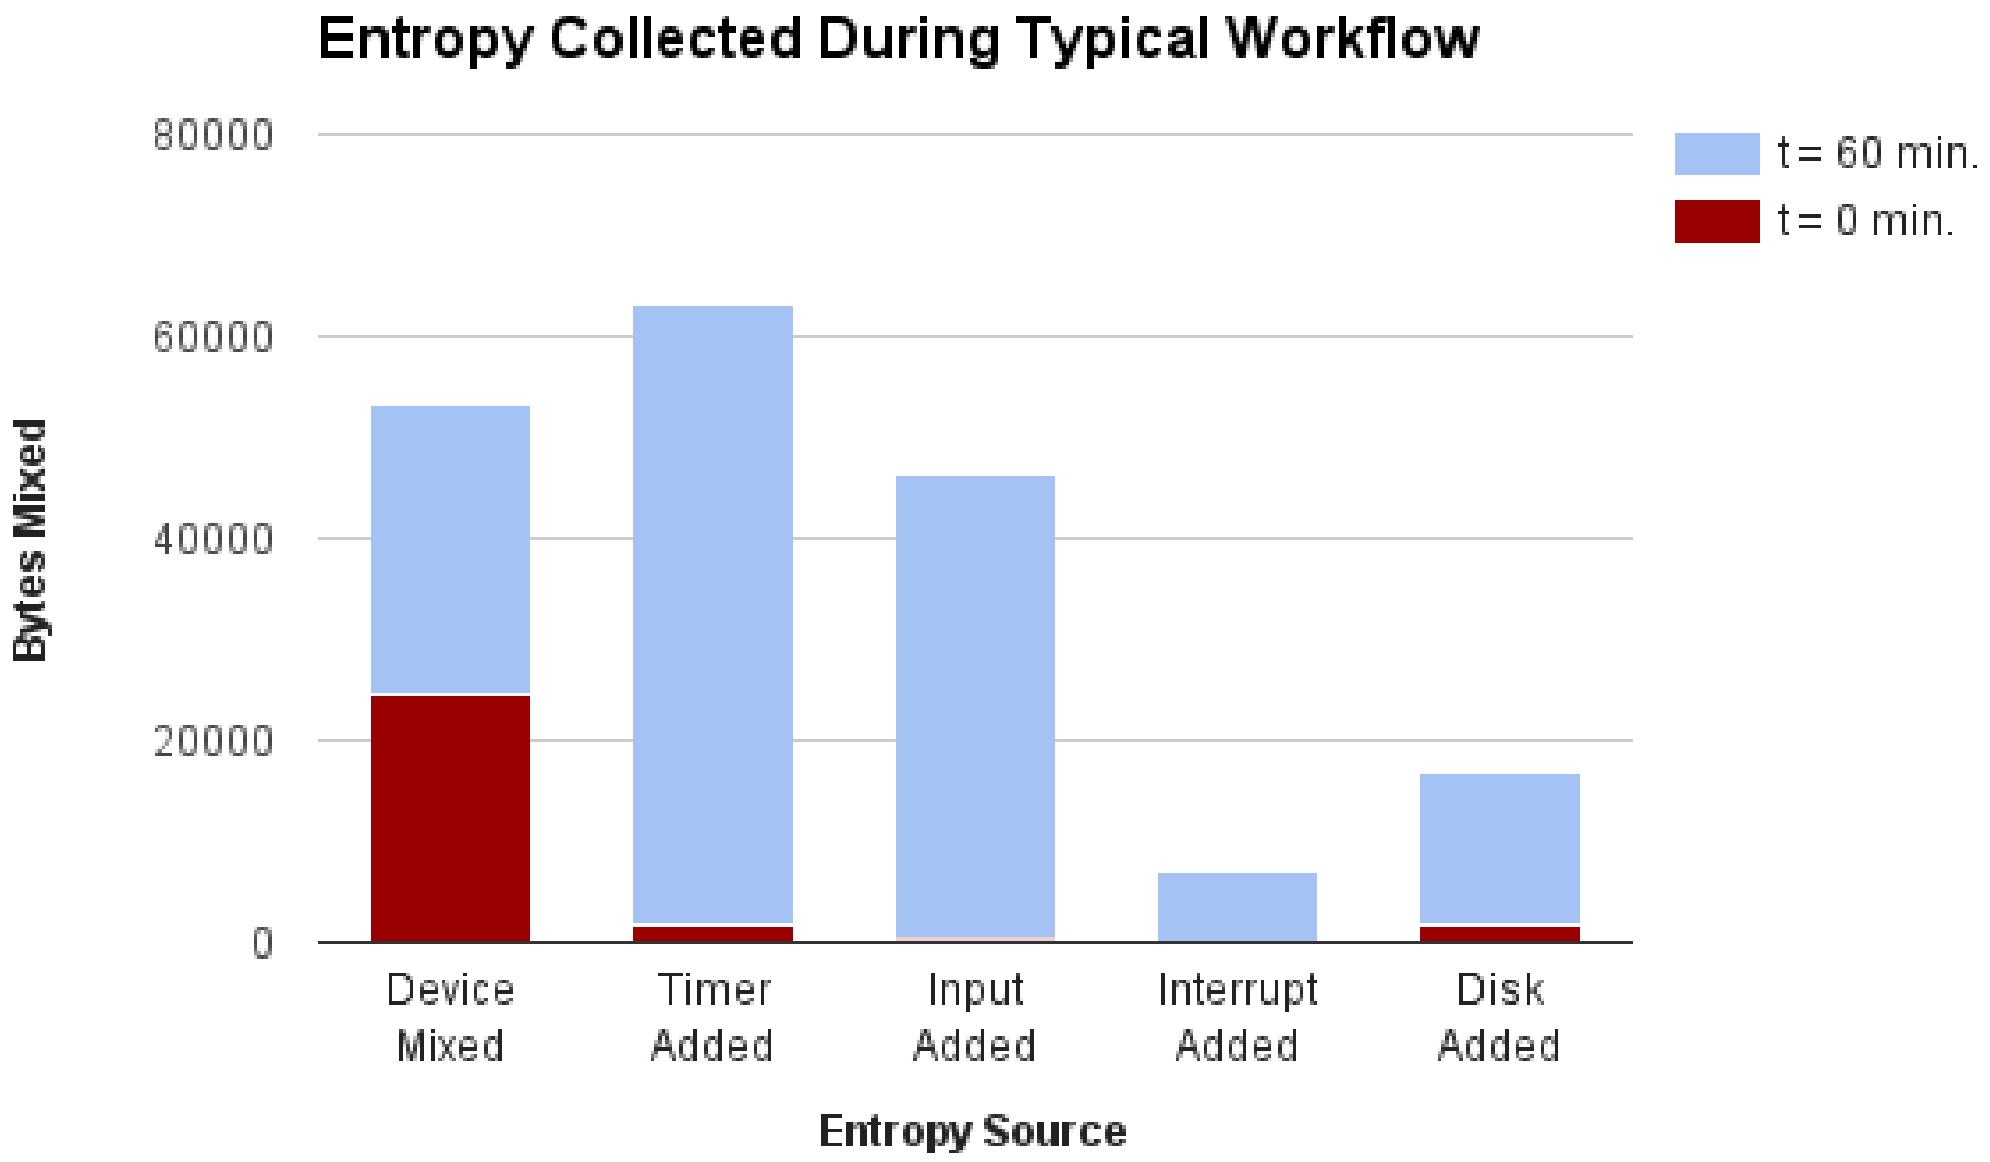
Entropy Collected During Typical Workflow
80000
60000
40000
20000
0
Bytes Mixed
t = 60 min.
t = 0 min.
Device Mixed
Timer Added
Input Added
Interrupt Added
Disk Added
Entropy Source

Ignoring Best Practices

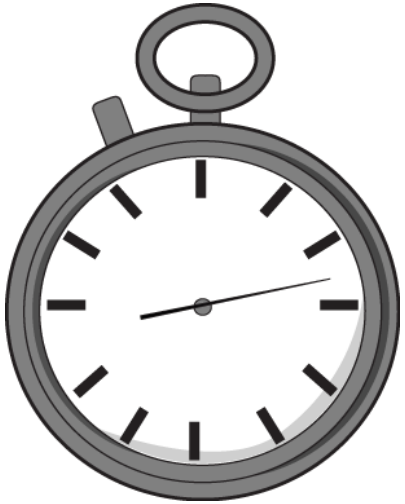

There are 2 ways to generate values

Performance vs. Stronger Security Guarantees

Empirical observation: application developers almost always choose the more performant option

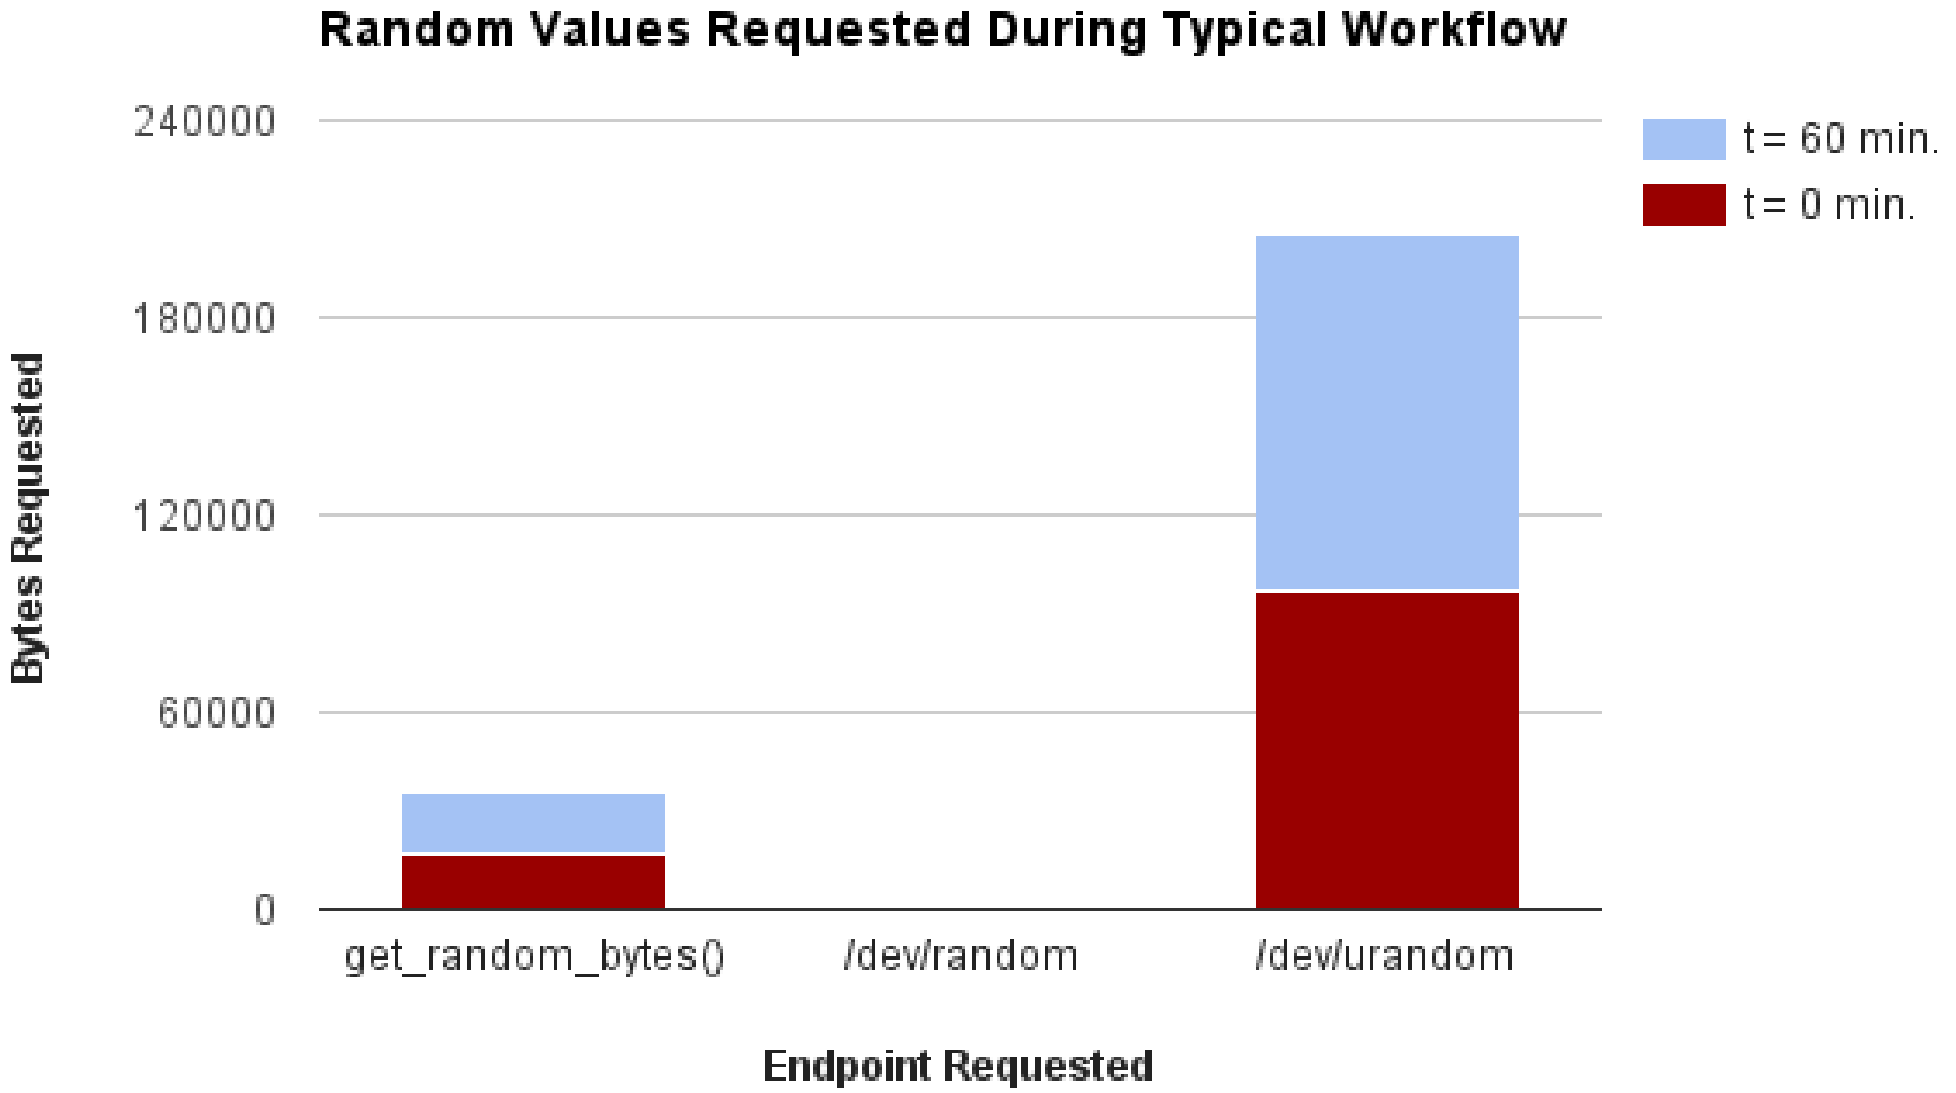
Random Values Requested During Typical Workflow
240000
180000
120000
60000
0
Bytes Requested
t = 60 min.
t = 0 min.
get_random_bytes()
/dev/random
/dev/urandom
Endpoint Requested

Firefox vs. Tor: Logging into Facebook

~1500 Bytes

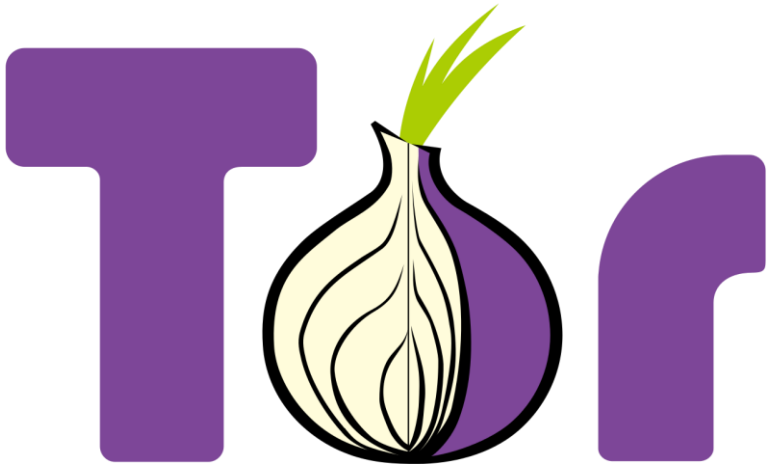

~1600 Bytes

What’s Next

Ongoing Research

Use our tool to discover if theoretical vulnerabilities* result in real-world attacks

Demonstrate that manipulating inputs can lead to biased outputs

*Dodis, Y., Pointcheval, D., Ruhault, S., Vergniaud, D., Wichs, D. Security analysis of pseudo-random number generators with input: /dev/random is not robust. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer Communications Security, CCS 2013,* November 2013.

Open Source

Plan to release source code to the community

Enable theory-minded researchers to conduct practical analysis

More people = More convincing argument

Underlying Motivations

Linux is ubiquitous in the world of computing

Design of a core component is known to be theoretically flawed

Discover vulnerabilities before they can be used

Acknowledgements

Dr. Nadia Heninger & SecLab

Dr. Ani Nenkova & Dr. Jonathan Smith

SEAS & the University of Pennsylvania’s CIS Department