That pesky critical infrastructure

Jason Larsen
First 2010

www.inl.gov

Idaho National Laboratory

What is Critical Infrastructure?

- *Critical Infrastructure [em por tant]*- 17 Industries necessary for the nation to function
  - Power
  - Water
  - Chemical
  - Manufacturing
  - ……

What is Critical Infrastructure?

- *Critical Infrastructure [no zee]*– Stuff private industry owns the government wants to “help” with

(Isn’t that bad for my bottom line?)

Boring Five Minutes

- Since this is a mixed audience, I’m going to spend 5 minutes on control systems 101
  - Feel free to check e-mail and take a power nap
- I’m going to cover high-level and low-level concepts in this presentation
  - I hope the mix comes out right

Typical Control System Layout

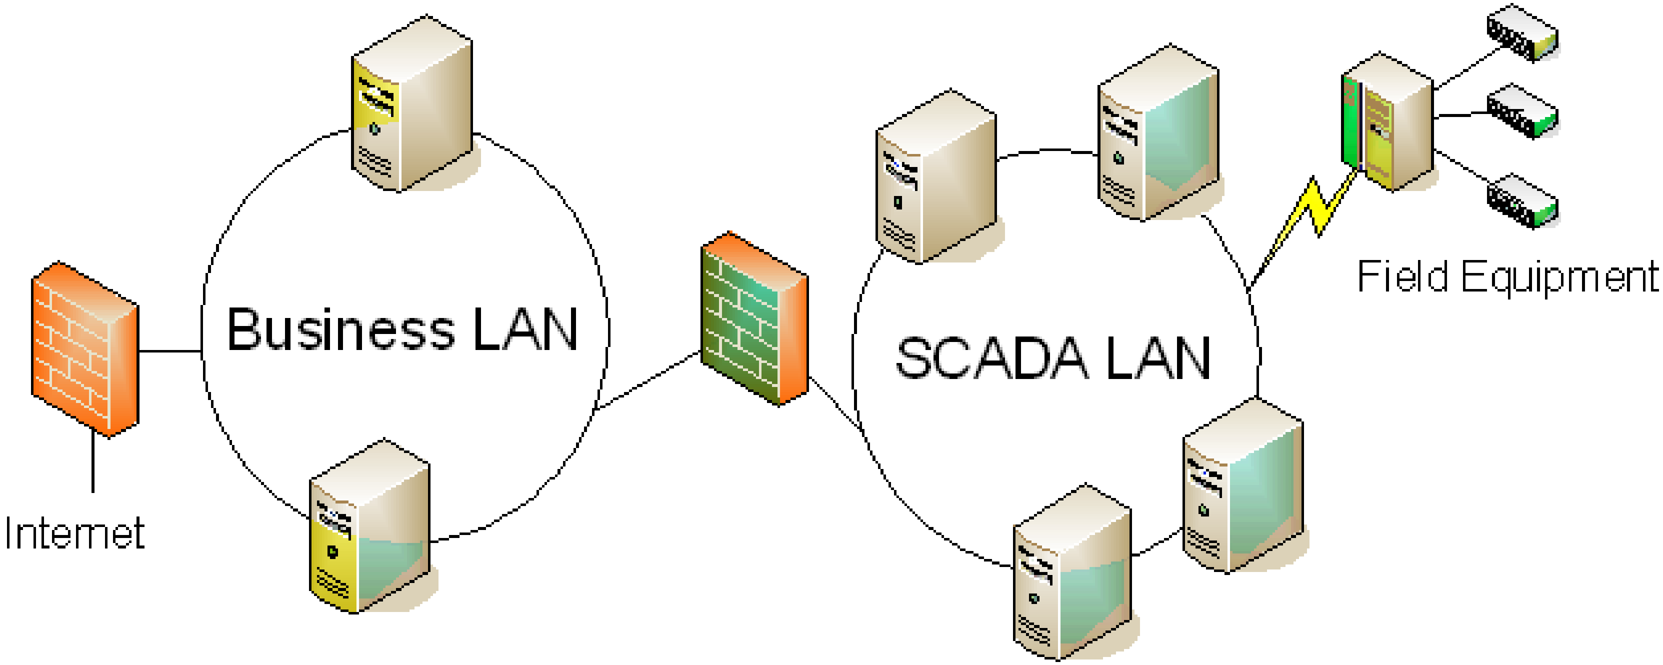

Where are the ticklish parts?

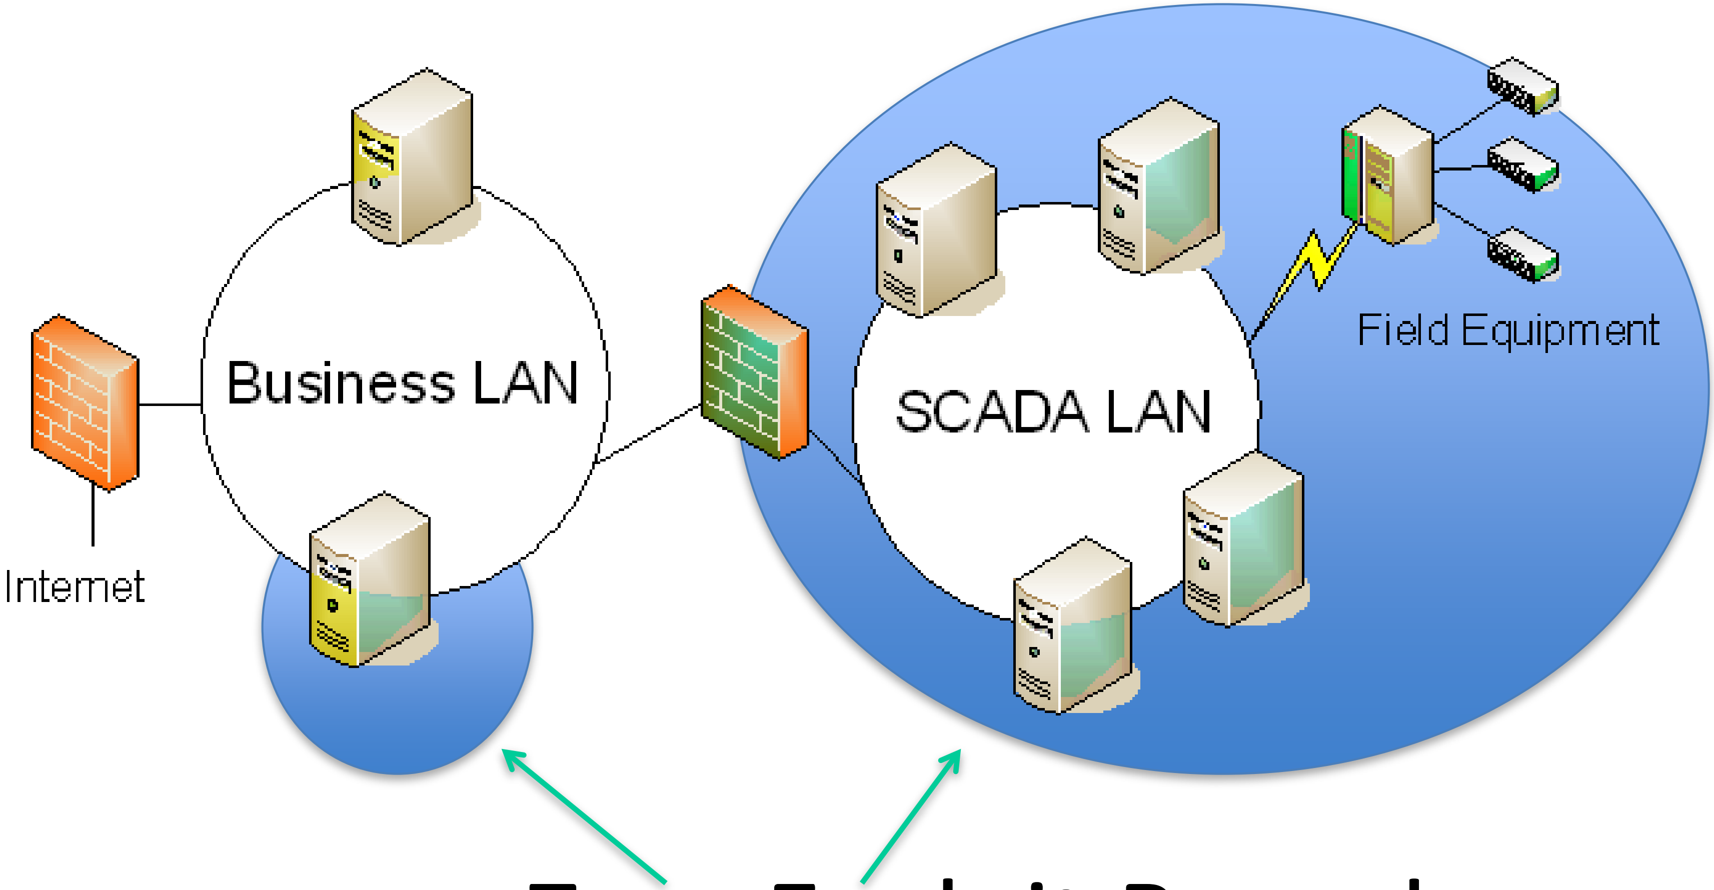

Zero Exploit Boundary

(If we get this far, we’ve already won)

Stages of a SCADA Attack

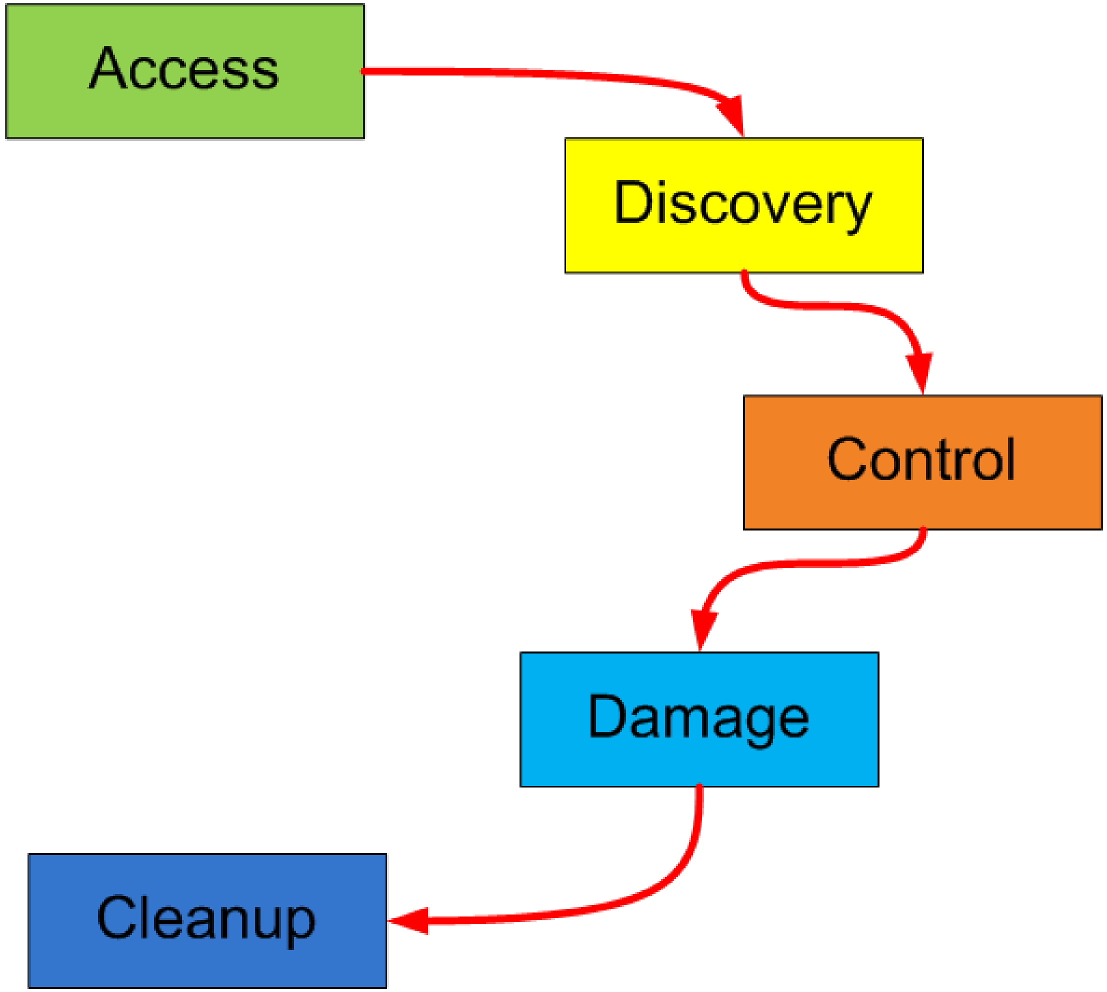

Regulation

- I may only have one outsourced firewall and one outsourced IDS, but I'm compliant
- Security standards have been insanely expensive
- They haven't changed the playing field much
  - Good companies still have good security
  - Bad companies still have bad security

Where are the Attackers Now?

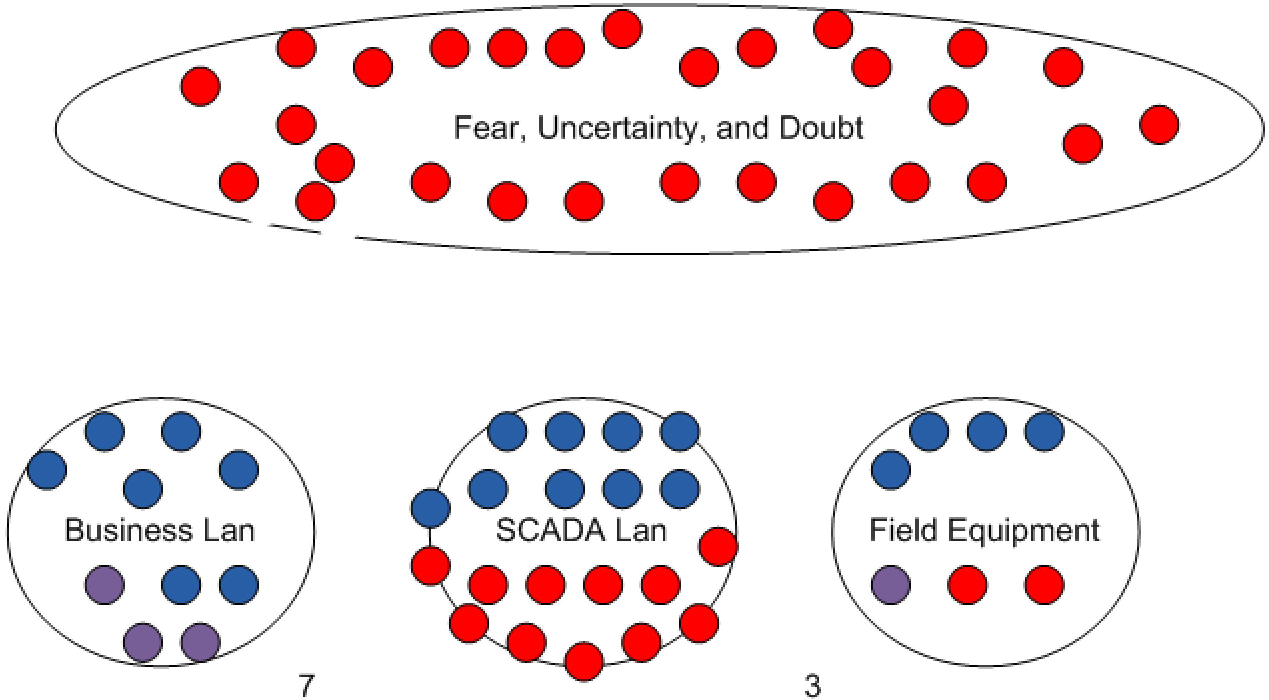

Fear, Uncertainty, and Doubt

(Plus lots of dots)

Business Lan

SCADA Lan

Field Equipment

	7		3		
	3		1		
				1	
31	11	1	8	1	0
Not Hacking	Access	Discovery	Control	Impact	Cleanup

Attackers are finally here

- We've been waiting for years for the SCADA hackers to declare themselves
- We now have direct evidence of attackers on a control network that knew what they had hacked into
- We now have direct evidence of attackers interacting with a controller using its native control protocol

Attackers

- Wait. You promised me fireworks.
- Explosions. I want the explosions.
- If you’re waiting for them to wreck the place, you’re going to be waiting for a long time

Attackers

- Destroying a process isn't very profitable
- It's much more profitable to monitor and wait for the perfect opportunity
- Attackers know they've compromised control systems
  - The helpful notes left on the system told us so
  - They don't seem at all interested in controlling the process
    - At least not yet……

End of Boring 5 Minutes

- OK. On to the present.

“Only half the battle for control systems will only be fought IP space”

The future

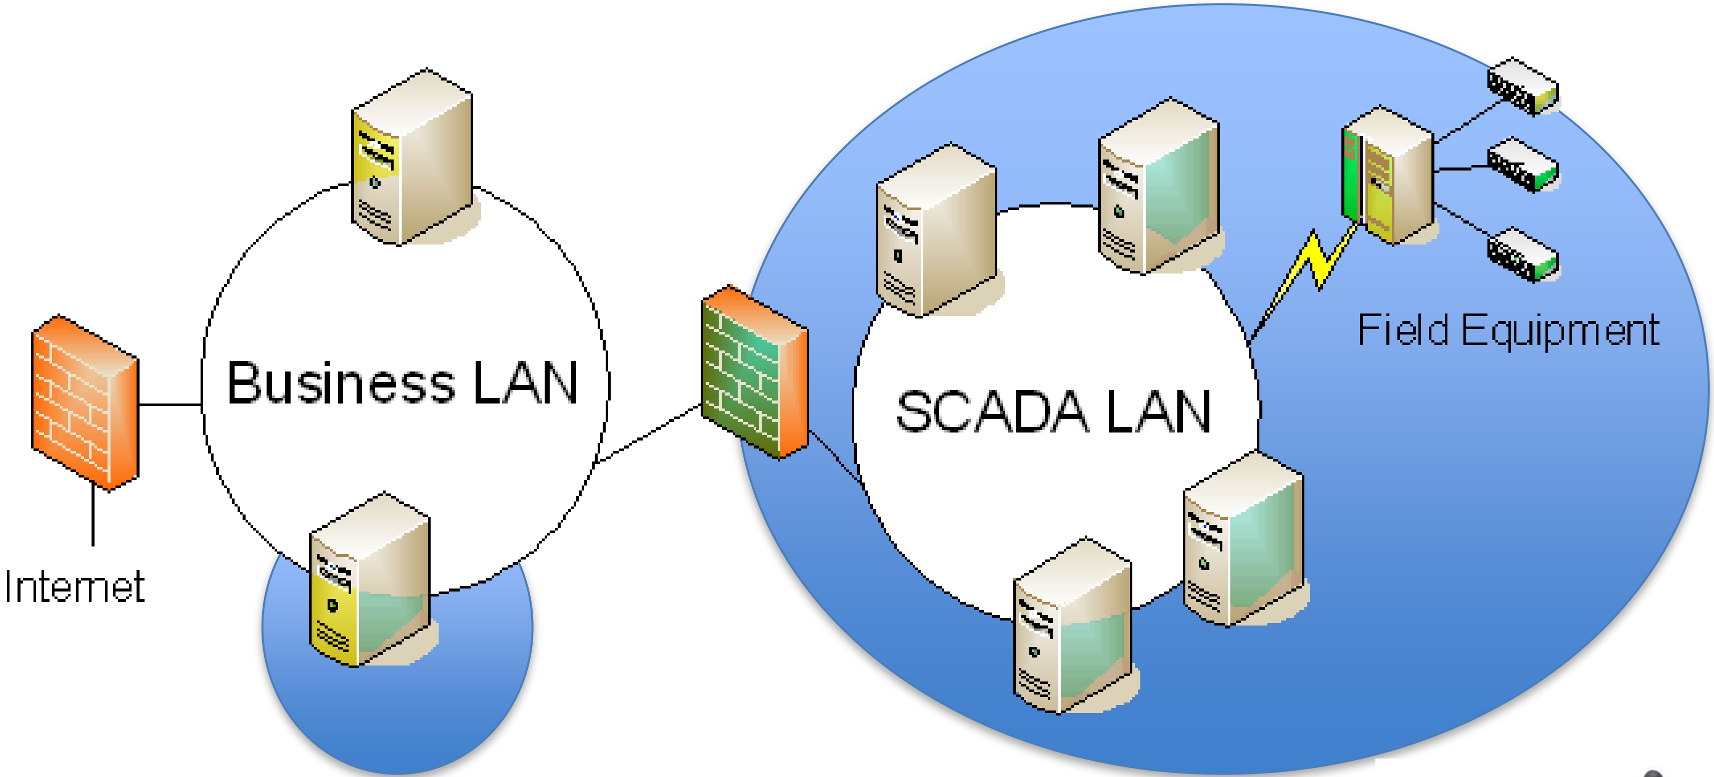

The ~~future~~ present

- There’s a guy running around with a Windows Mobile handheld that can operate breakers in a substation
- There’s a guy walking in a chemical plant right now controlling set points
  - His handheld doesn’t speak IP

The ~~future~~ present

- Wait! That’s a bad Idea.
- You’re 4 years late to the argument.
  - Go to the back of the line.
  - Wait for the next holy war. You lost this one.

The Smart Grid

- Nobody knows what the smart grid is
- It’s about more than meters
  - Power Modeling
  - Alternatives to spinning reserves
  - Solving the grid faster in case the windmills stop suddenly
- On the other hand, hacking meters is *really* fun

The Smart Grid

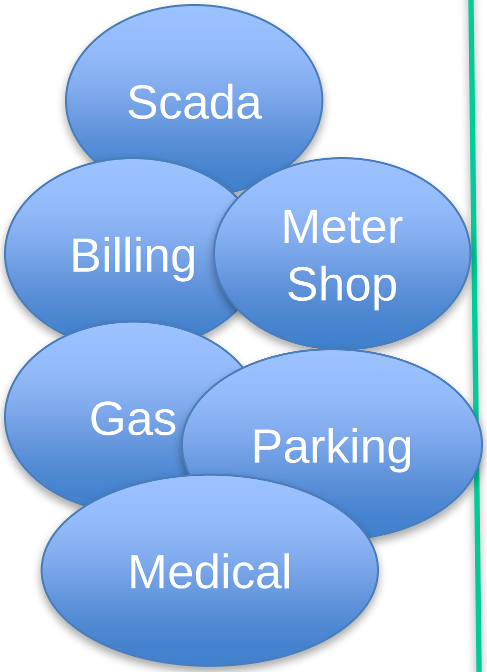

Scada

Billing

Meter Shop

Gas

Parking

Medical

Utility Interface

Backhaul Network

Home Area Network

Home Area Network

- Control of the HAN *does not* give access to or control of the backhaul network
- This is the most hackable surface
- It’s also the surface most exposed to the customer

Ti CC2x50 PRNG Problem

- Travis Goodspeed reported a problem in the ChipCon chips
- Basically, it only generated 2^{15}-1 keys
- It used a hardware pseudo-random number generator
  - (Also good for calculating a CRC-16)

ChipCon Problem

From the Documentation:

> The random number generator is a 16-bit Linear Feedback Shift Register (LFSR) with polynomial $X^{16} + X^{15} + X^{2} + 1$ (i.e. CRC16).
>
> It uses different levels of unrolling depending on the operation it performs. The basic version (no unrolling) is shown in Figure 27.

ChipCon Problem

From the Code:

```
 *  The seed value must not be zero or 0x0380 (0x8003 in the polynomial).  If it is, the psuedo
     *  random sequence won’t be random.  There is an extremely small chance this seed could randomly
     *  be zero or 0x0380.  The following check makes sure this does not happen.
     */
    if (rndSeed == 0x0000 || rndSeed == 0x0380)
    {
      rndSeed = 0xBABE; /* completely arbitrary "random" value */
    }
```

Seeding is only bad in a certain case

Another Key Problem

- Unnamed vendor (Until they fix it)
- Key generated by reading the least significant bit of the onboard temperature sensor
  - Crypto Generates Heat
  - A hot chip returns 0xFFFF for the temperature
  - Hhhhmmmm……..

Home Area Network Worms

- Even though you can’t shut off the power, a HAN device exploit can still be a problem
- In densely populated areas, the radios of one HAN will be within transmission distance of the neighbor’s HAN
- HAN worms have been shown to be possible
- What happens when my water heater attacks your refrigerator?

Backhaul Network

- Control of the backhaul network *might* give control of billing and remote disconnect
- Some vendors have fixed this problem

Zigbee Buffer Overflows

- Buffer overflows in Zigbee stacks have been shown to give access to the backhaul network
- This is not a given

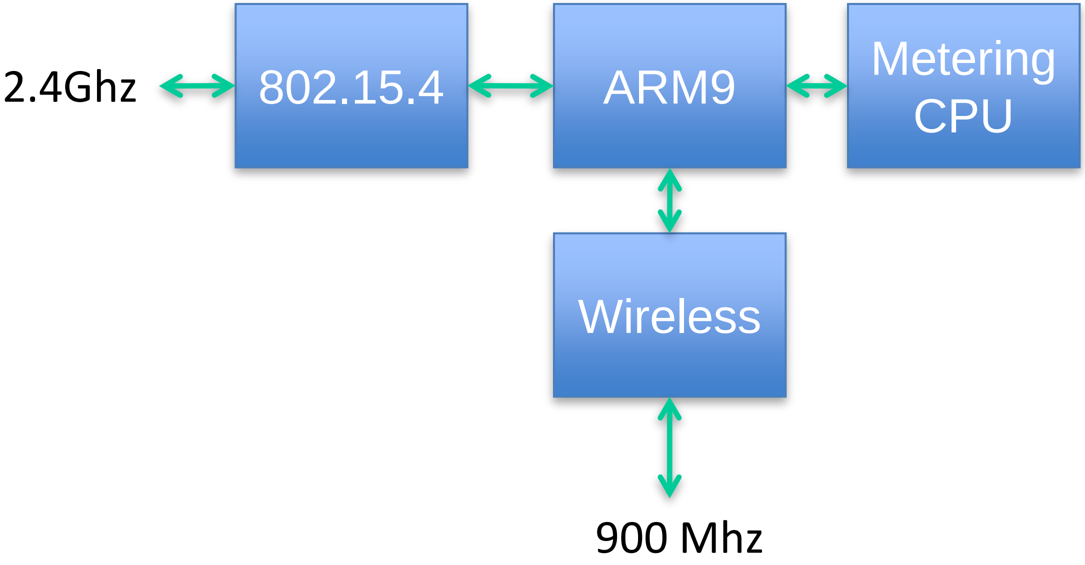

Meter Backhaul Worms

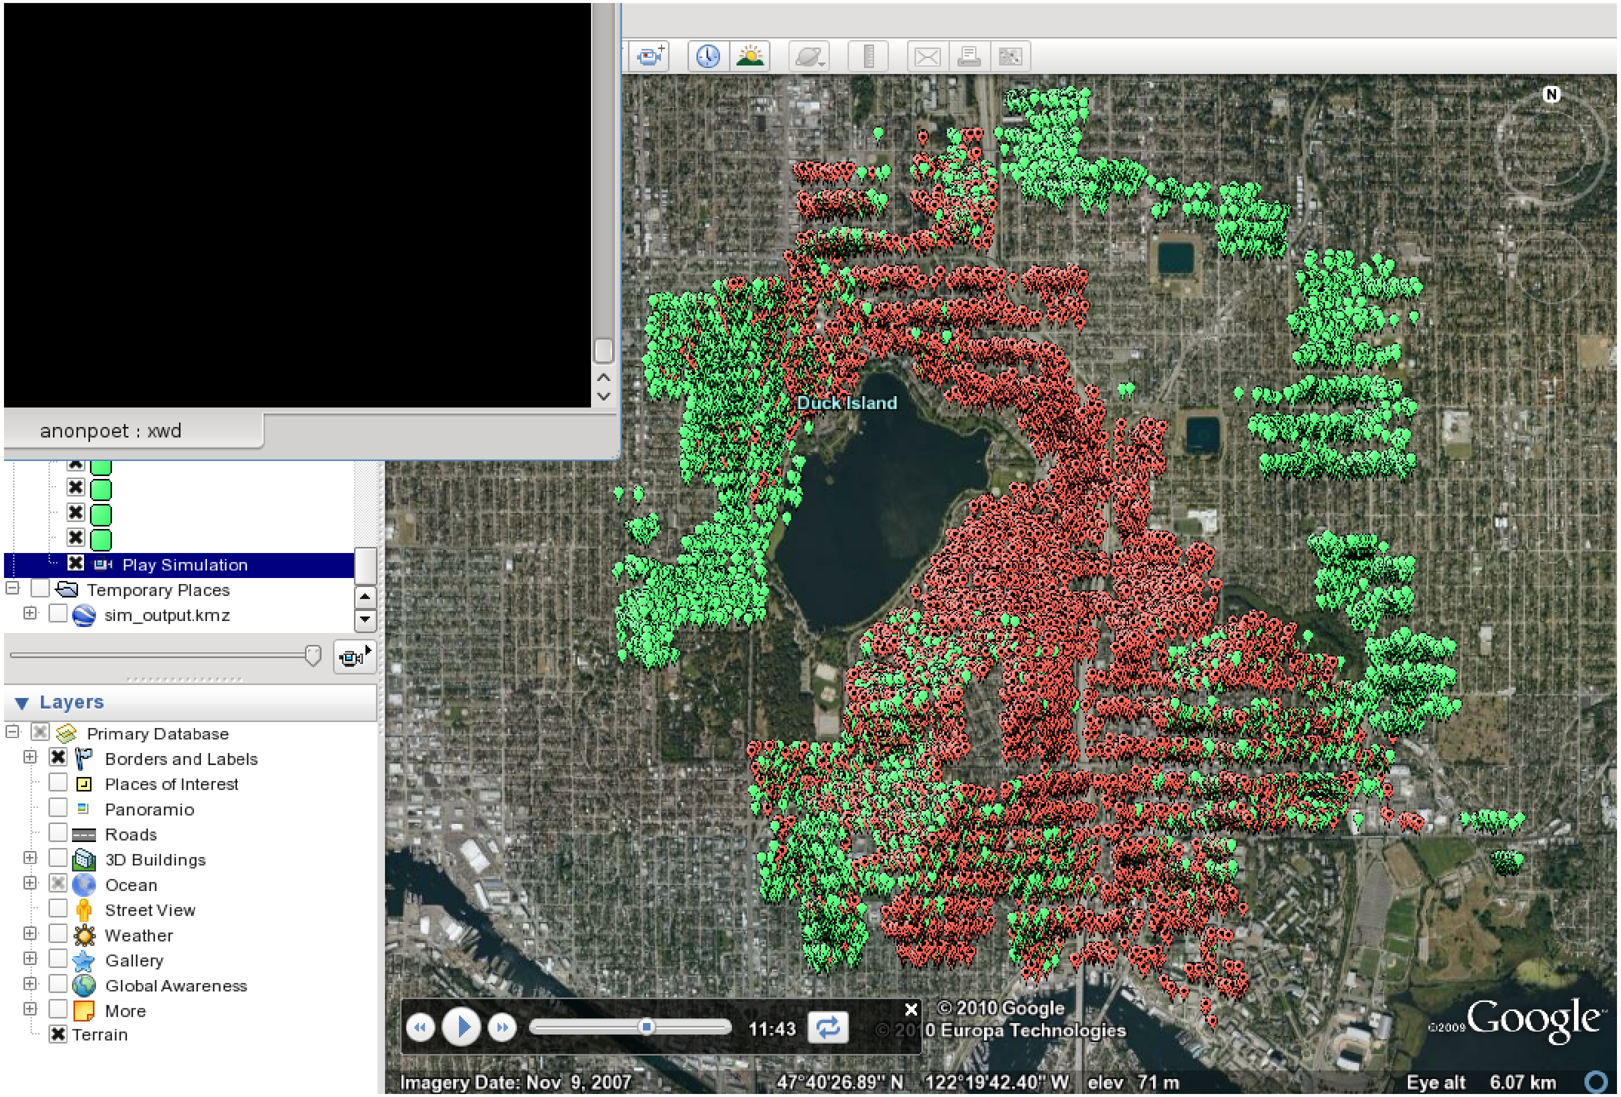

Meter Backhaul Worms

- Takeover of a city 24 hours +/- 2 hours
- Takeover of a state 24 hours +/- 2 hours
- Payloads can be interesting
  - Change Billing IDs
  - Remote Disconnect
  - Move 3 million meters to cell phone frequencies
- May have to touch every meter to clean up

Hacking Upstream

- So far no one has been able to hack from a meter into a control network
- We may not need to hack from the meters to get full control
- Lots of backend networks tie into an AMI system

Other Networks

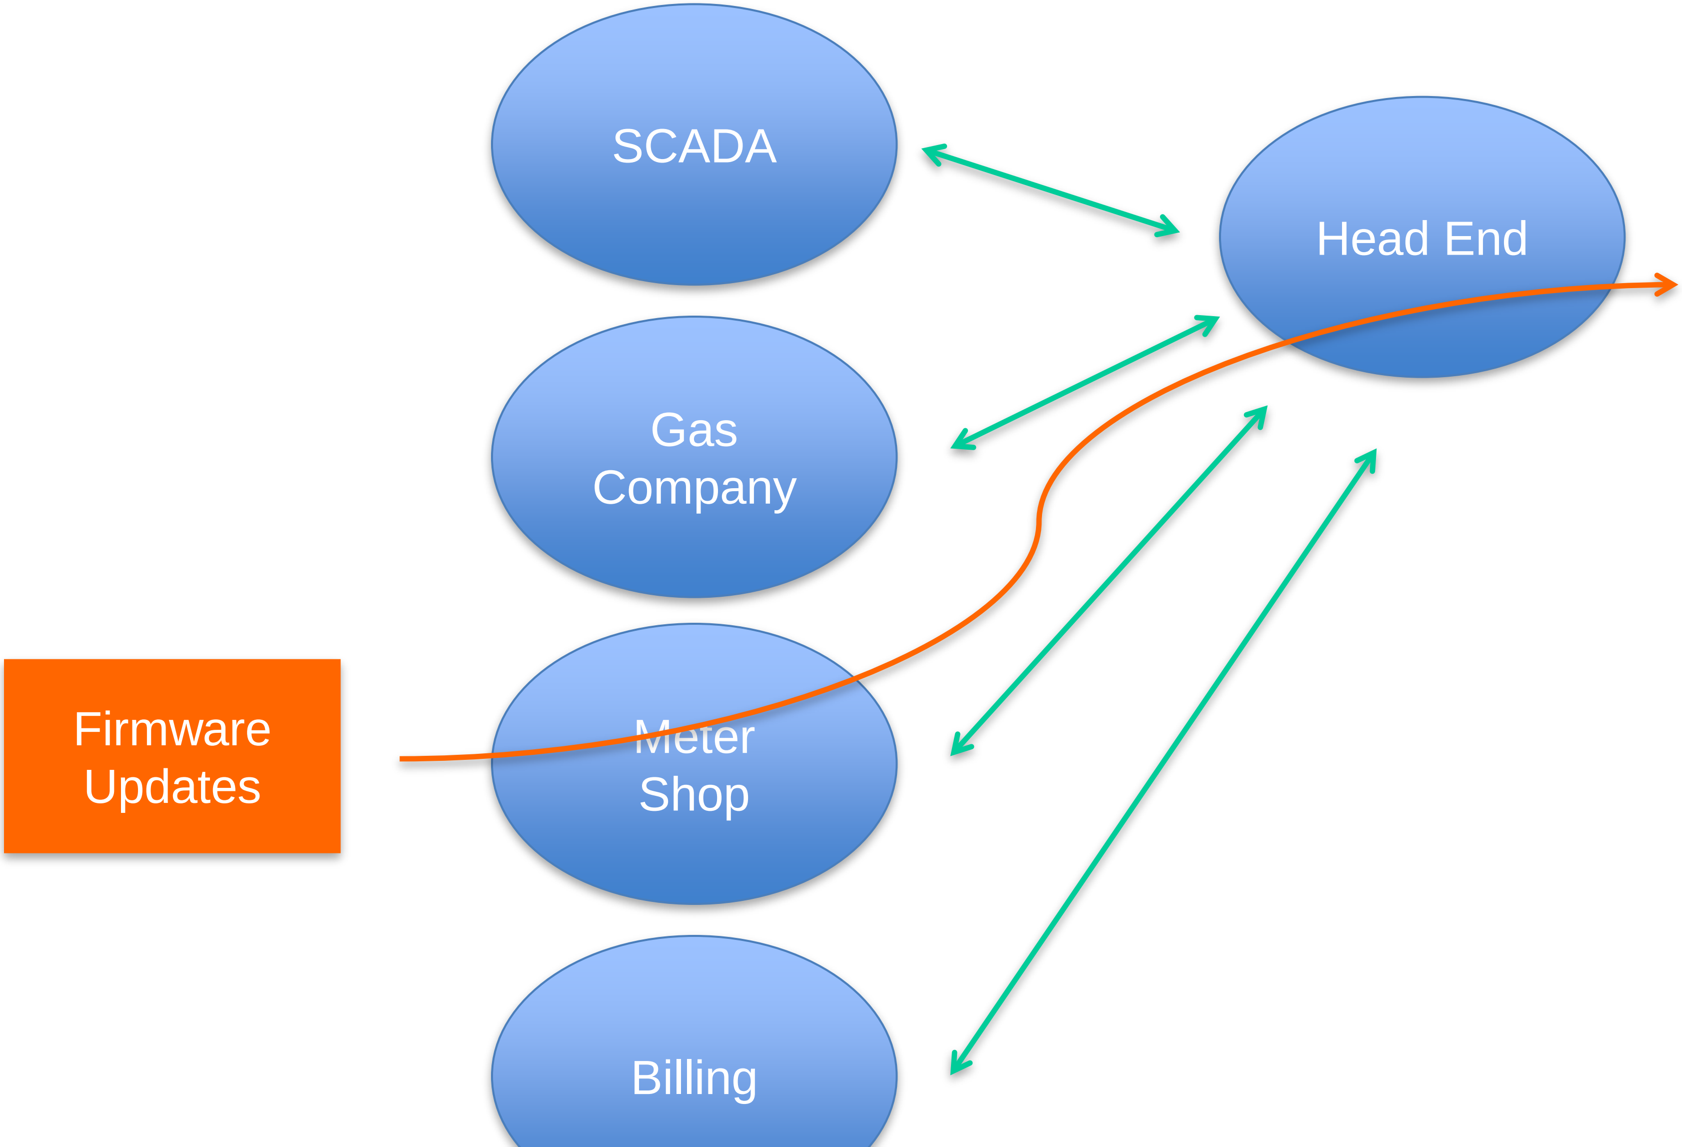

Other Industries

- Power gets most of the attention
- Other industries have also gone wireless

PLC Rootkit PoC

- A Proof-of-concept rootkit was presented at S3 over a year ago
- Attacker is able to install code on the embedded device and have a persistent presence

The supply chain

- Critical infrastructure in general, but especially electric power evolves at a maddeningly slow pace
    - They think everything through before implementation
- That's changing with smart grid technologies
    - The availability of money has successfully accelerated the pace of adoption
    - New technologies built by small firms are being deployed quickly in the market

The supply chain

- Right now smart grid and green energy are not a threat to the bulk electric system
  - There aren’t enough remote disconnect meters
  - There isn’t enough generation
- That is changing
  - Some wind projects will generate as much power as a nuke

The supply chain

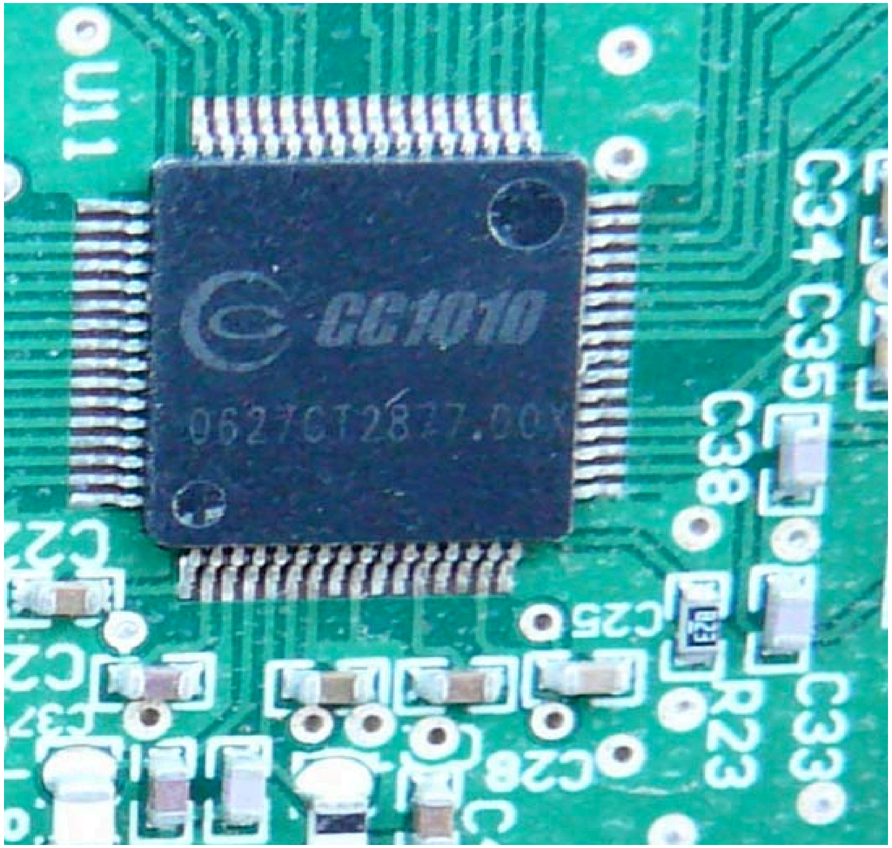

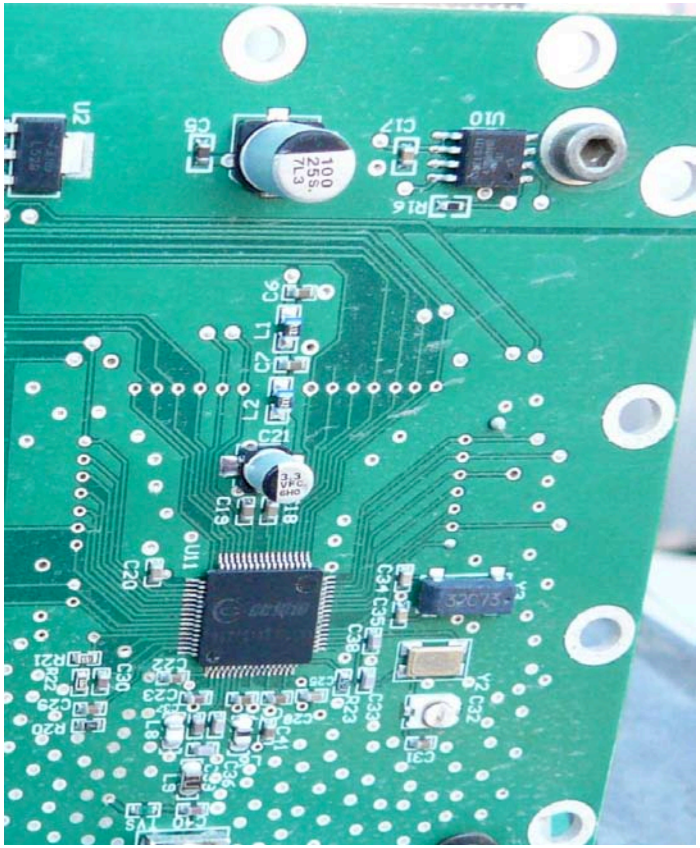

The firmware this board is based on was written by a college student as part of his degree.

Post Exploitation Research

- Defending the perimeter is becoming a well known problem
    - In many cases it’s also a lost cause
- After breaching the defenses, attackers still have tons of work ahead of them
- Post exploitation methodologies in SCADA have largely been unexplored
- This may represent the best chance defenders have of catching the attackers

Post Exploitation Research

- In order to be effective an attacker must figure out the constants used in the protocols

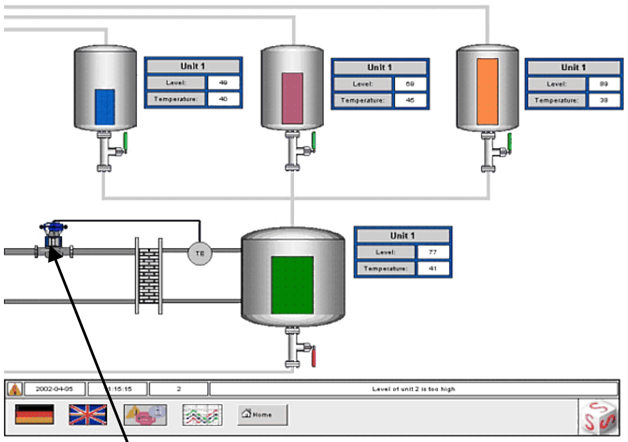

Unit 1	
Level:	49
Temperature:	40

Unit 1	
Level:	69
Temperature:	45

Unit 1	
Level:	89
Temperature:	38

Unit 1	
Level:	77
Temperature:	41

2002-04-05	1:15:15	2	Level of unit 2 is too high

Interesting Feedback Loop

```
0001a860  02 00 00 00 00 00 00 00  80 18 40 00 00 00 00 00  |..........@.....|
0001a870  80 18 00 00 00 00 00 00  a0 08 00 00 00 00 00 00  |................|
0001a880  05 00 00 00 0c 00 00 00  08 00 00 00 00 00 00 00  |................|
0001a890  18 00 00 00 00 00 00 00  6b 00 00 00 01 00 00 00  |........k.......|
```

Constants in Control Protocol

Post Exploitation Research

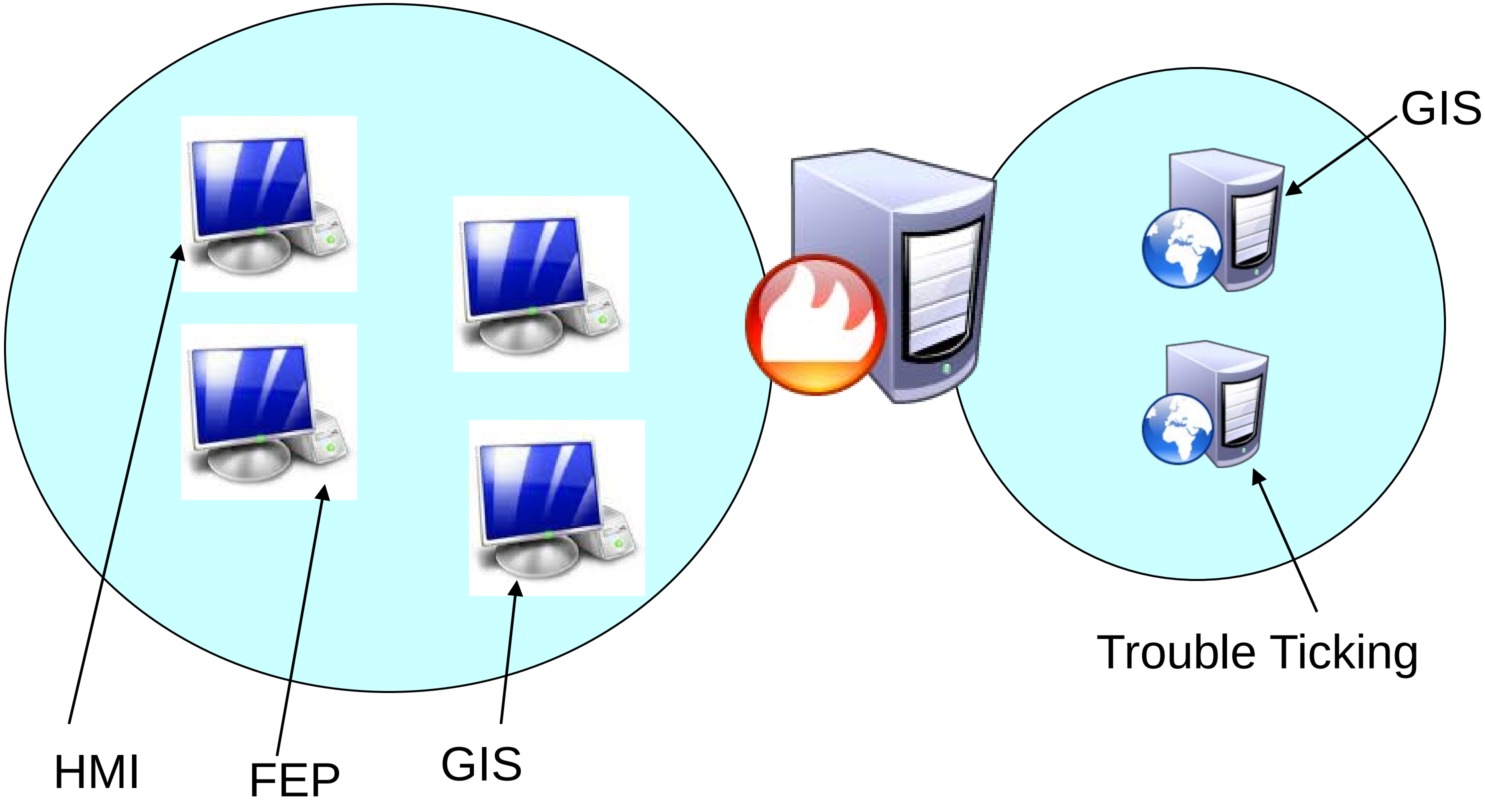

Information Leakage

The end of tools

Profibus

HART

IEC-870

Custom Serial Protocols

Modbus

Foundation Fieldbus

ZigBee

The end of tools

- After an attacker leaves the IP network, we have no tools to detect or do forensics
- A rootkit on a PLC is very different than a rootkit on a Windows machine

Firmware rootkits

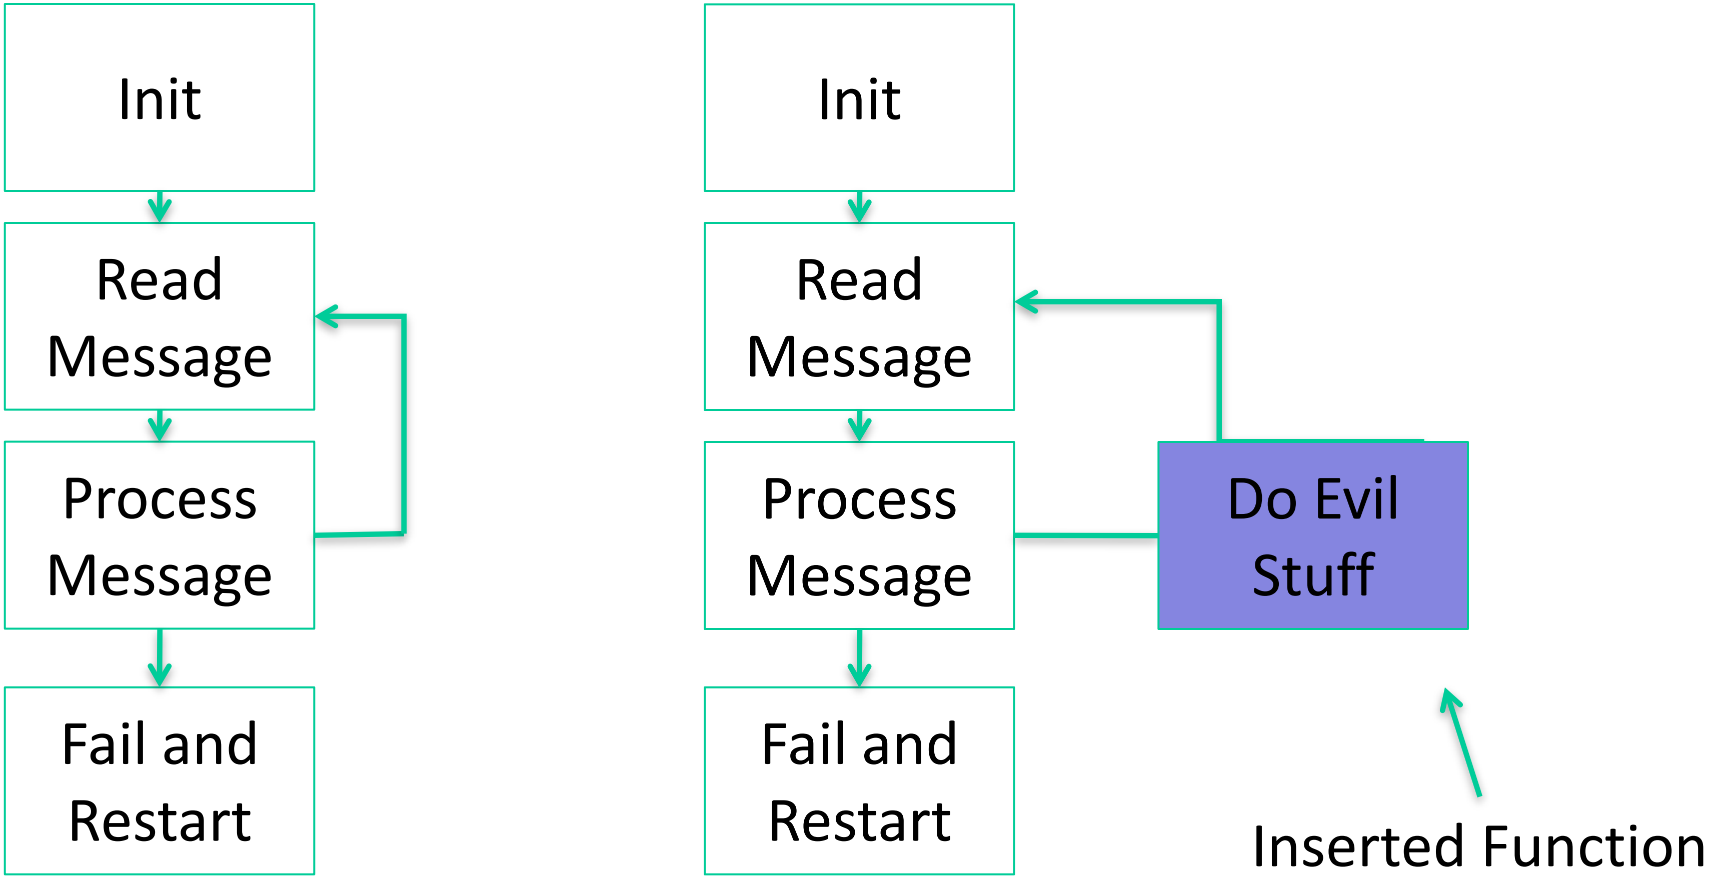

Forensics

- Here’s a meter.   Find the hacker.
- Looking to see if an embedded device has been compromised isn’t straightforward
  - Get the firmware
  - Understand the environment
  - Compare it to a known good
  - Reverse engineer the differences

Getting the Firmware

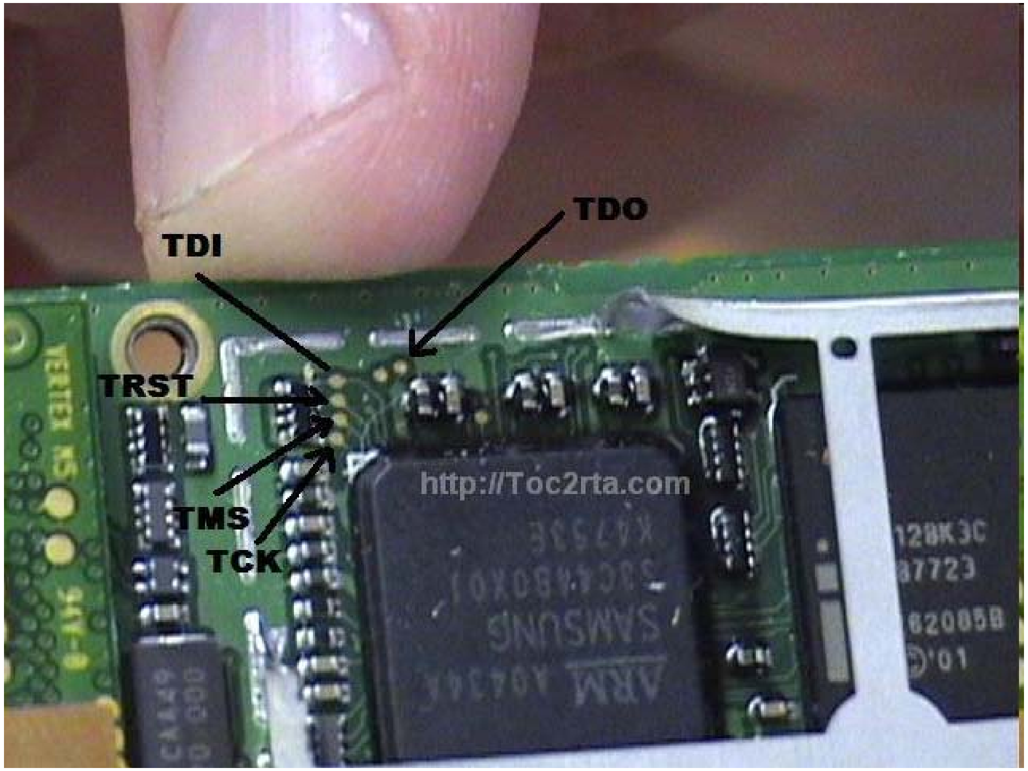

In some cases we may be forced to exploit the device just like the attacker

mCode

- Reverse engineering firmware isn’t too bad
  - Unfortunately there’s a bunch of them
  - There’s no way an analyst can learn them all
- Tools aren’t portable across microcontrollers
- I’m working on this

Assemblies

- ARM

```
 LDR    R3, [R11,#-8]
CMP     R3, #3
BGT     loc_8474
LDR     R3, [R11,#-8]
ADD     R3, R3, #4
MOV     R0, R3
BL      #0x8444
B       loc_847C
LDR     R0, [R11,#-8]
BL      TestFunc
```

- MSP430

```
cmp.w #0x4,R12
jge 0x801A
mov.w #0x4, R12
br #0x8010
add.w #0x4, R12
br #0x8010
```

- 8051

```
mov   r2,dpl
mov   r3,dph
clr   c
mov   a,r2
subb  a,#0x04
mov   a,r3
xrl   a,#0x80
subb  a,#0x80
jnc   00102$
mov   dpl,r2
mov   dph,r3
inc   dptr
inc   dptr
inc   dptr
inc   dptr
ljmp  100$
mov   dpl,r2
mov   dph,r3
ljmp  100$
```

Assemblies

- Each assembly has its own idiosyncrasies
- AVR uses the Z register like a stack
- ARM has the funky 16-bit Thumb instructions
- Inline indirect jumps
- It gets worse with all the ways to interact with the I/O

The experiment with mCode

- It may be possible to convert each assembly to a standard format
- The CS student’s motto:

  “There is no problem so complex that it can’t be solved with one more layer of indirection”

mCode

- It’s possible to represent each assembly in a standard form
- Most opcodes are common, but they have side effects
  - Add (ARM)
  - Add (AVR)
  - Add (X86)

Side Effects

- X86 Add eax, ebx

```
eax:=eax+ebx
If eax+ebx>0xFFFFFFFF:
   c:=1
If eax+ebx>0x7FFFFFFF:
  0:=1
If eax+ebx==0:
  z:=1
```

- ARM add r1,r2

```
r1:=r1+r2
if r1+r2>0xFFFFFFFF:
   c:=1
if r1+r2>0x7FFFFFFF:
  n:=1
if r1+r1==0:
  z:=1
```

Side Effects

- Most side effects don’t influence code execution
  - They can be culled from the instruction list

```
eax:=eax+ebx
if eax+ebx>0xFFFFFFFF:
  c:=1
ecx:=ecx+edx
if ecx+edx>0xFFFFFFFF:
  c:=1
jc 0x804855
```

Aggregation

- Instructions can then be combined into operations

```
if r1==0:
  z:=1
if z==1:
  pc:=label1
```

```
if r1==0:
  pc:=label1
```

Pseudo-C

```
if r1==0:
  z:=1
if z==1:
  pc:=label1
```

```
if r1==0:
  pc:=label1
```

```
if (r1==0){
    Label1();
}
```

Forensics

- The goal is that in the future we will be able to quickly analyze meters and other embedded devices for malware
- This is only part of the problem
  - Bad revision control on the part of the vendors
  - Board level environment
    - Where does output 4 go?
    - What does it turn on?

Interesting Times

- It’s pretty much the wild west of control systems hacking
- After years of slow adoption, new technologies are rapidly being deployed
- Only half the battle will be fought in IP-space
- We don’t really understand all the side effects of the what we’re deploying
  - This gives the advantage to the attackers

Questions??

Jason Larsen

Jason.larsen@inl.gov