Prometheus @ Datacenters

Why Modbus Is Even Worse than SNMP

Richard Hartmann,
RichiH@{freenode,OFTC,IRCnet},
richih@{debian,fosdem,richih}.org,
@TwitchiH

2019-11-07

whoami

- Richard ”RichiH” Hartmann
- Swiss army chainsaw at SpaceNet
- Project lead for building one of the most modern datacenters in Europe
  - First one world-wide to be certified under the new revision of EN 50600
  - One of less than a dozen with both security and availability class 4
  - There’s no harsher non-military standard on Earth
- FOSDEM, DENOGx, PromCon staff
- Prometheus team member

Show of hands

- Who has heard of SNMP?
- Who has heard of ModBus?

ModBus

ModBus is worse

Relation

- SNMPv1: 1988
- ModBus release: 1979 (!)

SNMP

Without SNMP, the Internet would break down within hours

ModBus

Without ModBus, society would break down within hours

ModBus

- Without ModBus, you would have no power
- Without ModBus, you would have no water
- Without ModBus, you would have no ports, or trade
- Without ModBus, you would have no processed food
- Without ModBus, you would have no clothes

ModBus

Of course, ModBus has zero security built in

Flavours of ModBus

- Modbus RTU: Serial bus with binary data, most common. Hard real time
- ModBus ASCII: Serial bus with ASCII. Just don't. Hard real time
- Modbus TCP: Binary over TCP/IP. No hard real time requirements
- Modbus over TCP: Slight differences, not commonly used
- ModBus UDP, Modbus Plus, Pemex Modbus, Enron Modbus: Ignored

Which to use?

You want to use ModBus TCP

What if I can't?

Bridging RTU into TCP is common and you can buy ”master” units off the shelf

Master & slave

References to master & slave in modbus_exporter have been removed even though they are still part of the official standard

Addressing scheme

- 00001-09999: Read-Write, Discrete Output Coils
- 10001-19999: Read-Only, Discrete Input Contacts
- 30001-39999: Read-Only, Analog Input Registers
- 40001-49999: Read-Write, Analog Output Holding Registers

Addressing scheme

- 00001-19999: Bit-wise addressing into a 2-byte block. So you need sub-addressing
- 30001-49999: 2-byte block. Unless you need 16 bits, you need suba-ddressing or combination
- You always get 2-byte blocks back

Wat?

- No other data types defined
- Four ways to clobber a Float32 together:
  - Big endian (1 2 3 4)
  - Little endian (4 3 2 1)
  - Mixed endian (2 1 4 3)
  - YOLO endian (3 4 1 2)

Waat?

At least I have not seen YOLO endian yet

Waaat?

- Yes, ”Input” and ”Output” are from the perspective of the sender, not the actual device
- Yes, x0000 is skipped
- Yes, the binary 0x0000 maps to decimal 00001
- No, there’s no rule if you start counting with 0 or 1, it’s free for all
- Addresses up to 65536, or 105536, is ”extended range”

Waaaat?

This standard is enforced by devices simply stopping to work

Easy, reliable, horrible

Reminder

- Without ModBus, you would have no power
- Without ModBus, you would have no water
- Without ModBus, you would have no ports, or trade
- Without ModBus, you would have no processed food
- Without ModBus, you would have no clothes

Maps

ModBus maps are roughly what SNMP MIBs are

Maps

Only you can't unit test them and your production might stop working if you do something wrong

Maps

I have seen maps which are scans of photocopied paper

How do you work with that?

Industry standard is to have a hex viewer, a map, an Excel sheet, and strong nerves

What do I use this in datacenters for?

Everything

What do I use this in datacenters for?

Everything, except the cameras

What do I use this in datacenters for?

Access control, intruder detection, glass breakage, fire detection, fire suppression, cooling set points, groundwater pump, groundwater filters, ion exchange pump, reverse osmosis system, water leakage, fan speed, doors opening and closing, fence gates, lighting, MCCB & status, diesel engine status, diesel fuel tank levels, battery runtime, battery health, elevator access, elevator position, movement in secure areas, potential to ground, lightning strikes, microsecond events on power distribution, medium voltage, transformer load, transformer heat, floodlights, pressure release valves, airflow in office, temperature in office, temperature/humidity/pressure in data halls, smoke extraction fans, emergency exit status, LASER fence scanners, conductivity of cooling water, bullet-proof glass being shot at

What do I use this in datacenters for?

Not a complete list

Why?

Why?

Why?

I like pain

Why?

ModBus is the one standard supported by ALL industrial equipment

Why?

ModBus is horrible, but it's also extremely reliable

Why?

Because countless people would die if it wasn't

How?

`https://github.com/RichiH/modbus_exporter`

Max Inden did tons of work during a one-month networking & ModBus stint at SpaceNet

Caveats

If you have ModBus RTU, use a PLC as a gateway to expose ModBus TCP

Caveats

Reading out ModBus registers takes several seconds

Future work

Currently having my PLCs reprogrammed to expose seconds spent and might adapt exporter to calculate correct time

Future work

There is a semi-standard way to write a ModBus map and I want to have a generator like snmp_exporter's

Reminder

- Without ModBus, you would have no power
- Without ModBus, you would have no water
- Without ModBus, you would have no ports, or trade
- Without ModBus, you would have no processed food
- Without ModBus, you would have no clothes

ModBus

Without ModBus, society would break down within hours

Thanks!

Thanks for listening!

Questions?

Twitter: @TwitchiH