Metering the smart way – a smartgrid for the datacenter

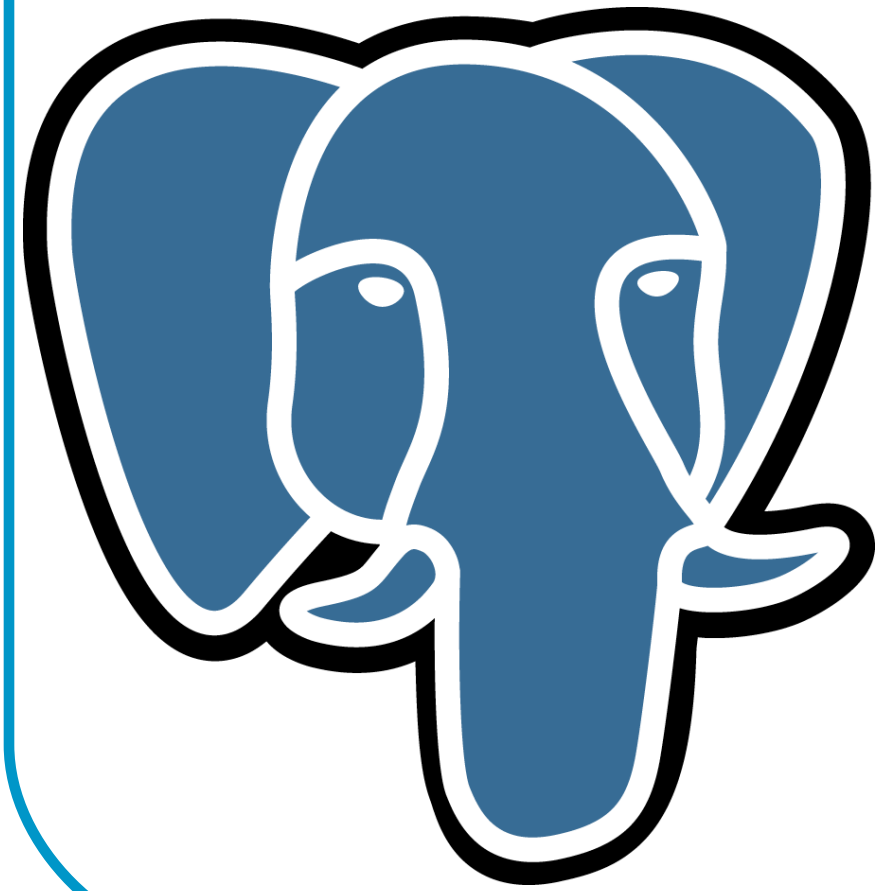

Stefan Kaltenbrunner
stefan@kaltenbrunner.cc
http://www.kaltenbrunner.cc/blog
Conova communications GmbH

PostgreSQL Conference Europe 2011

Amsterdam,
The Netherlands

Questions to the audience

How many of you have some sort of operational experience with relational databases?

Questions to the audience

How many of you can do coding or scripting at some level?

Questions to the audience

How many of you have experience in the field of building automation(SCADA) or experience with industrial bus systems?

In the beginning...

A clash of cultures...

Building automation

Building automation

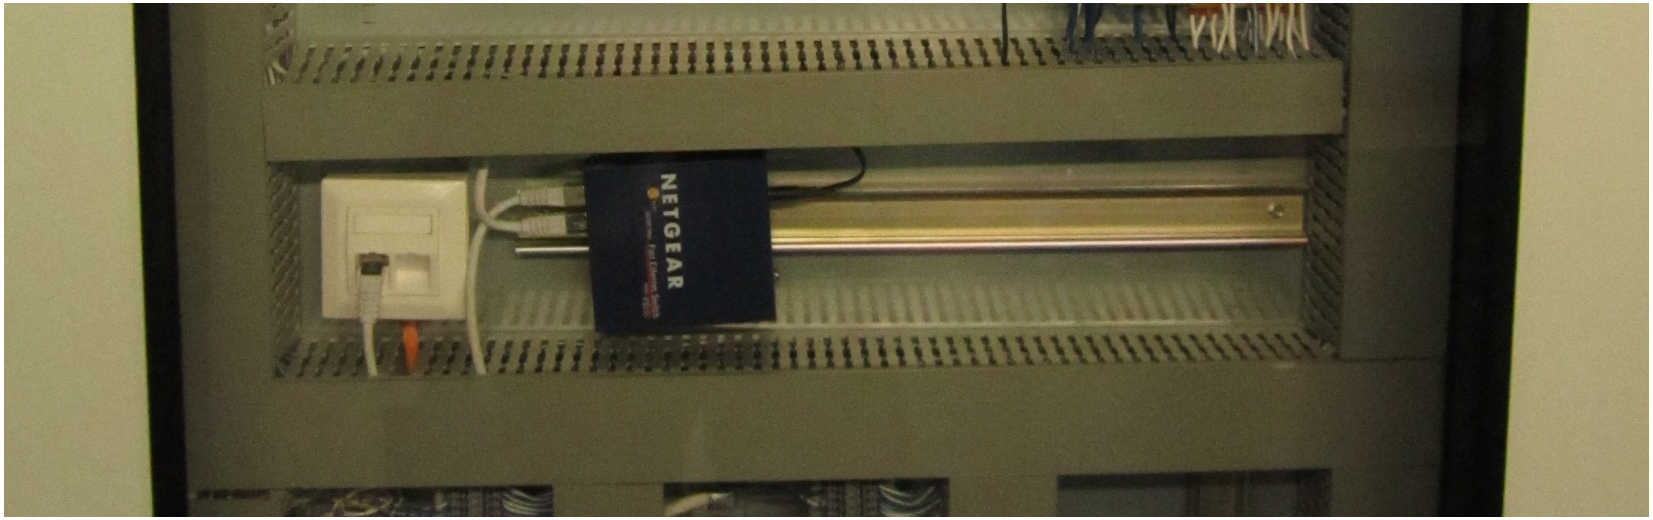

IT vs. building automation

A datacenter?

- A server?
- A rack?
- A bunch of racks?
- A serverroom?
- Something else?

A datacenter - really?

Only 30% of a large datacenter is the actual server room...

A datacenter - really?

- What about the rest?
  - Cooling (both inside and outside)
  - Electric systems (Switchboards, Transformers,...)
  - UPS(UPS itself, batteries, flywheel)
  - Emergency Power (genset)
  - Fire suppression systems
  - Security Systems (access control, surveillance)
  - Storage and Service areas

A datacenter

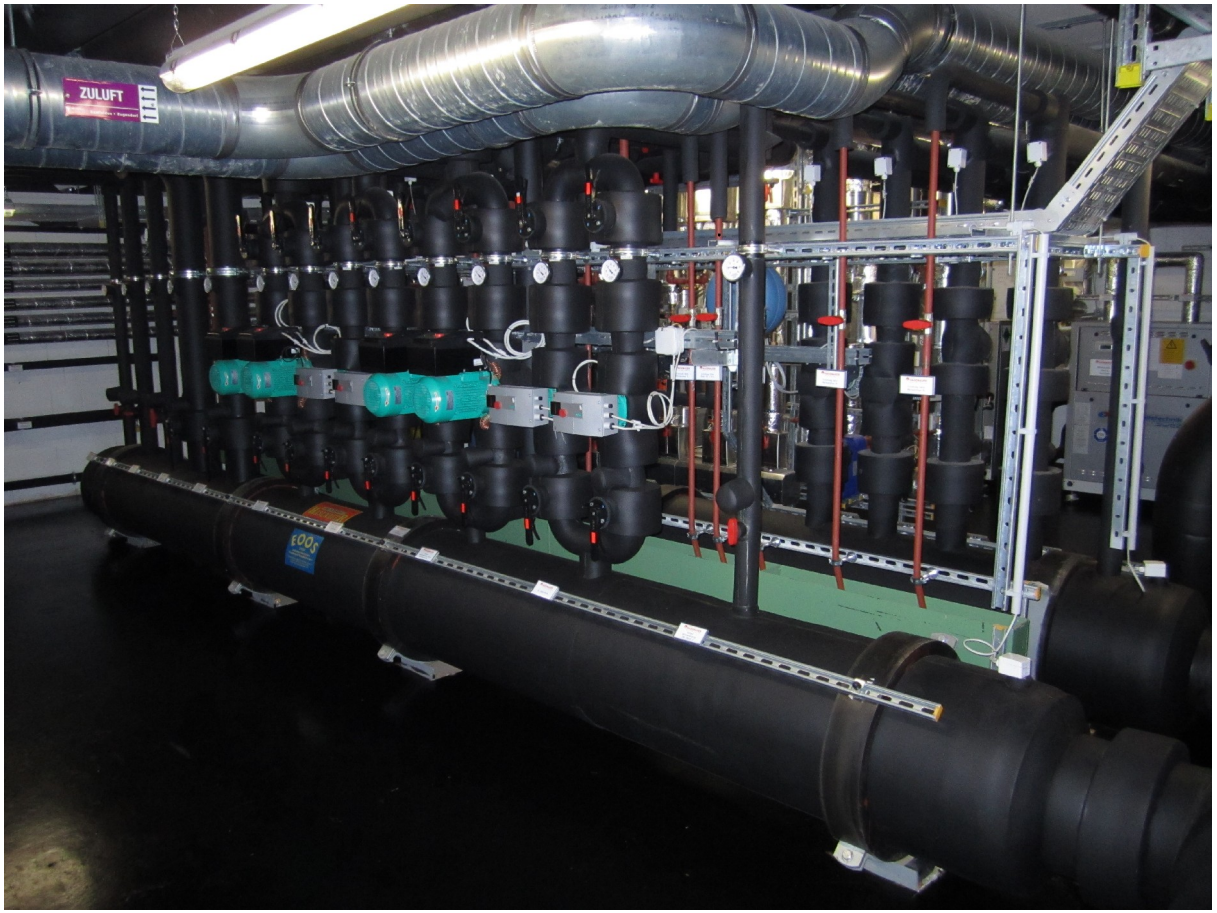

A datacenter

How does it work?

- A datacenter really is an industrial complex
- A datacenter may contain servers but those are nothing without the supporting infrastructure
- A datacenter consists of industrial and commercial bus systems
- Industrial systems are 20 years behind IT

IT vs. Industrial systems

- Pros
  - Industrial bus systems are simple
    - 2-wire
    - Long distances
    - Cheap
    - A bus (10-250 devices)
  - Industrial systems are field proven
    - Systems run for 30-40 years
    - Low requirements on environment

IT vs. Industrial systems

- Cons
  - Slow/parallelism
  - No (usable) APIs
  - Limited data collection capabilities
  - Requires ”dedicated PC”
  - Custom development environment
  - Reliability/Redundancy

    → *Very expensive to buy and maintain!*

What now?

- Design your own...
  - Use tools you know
    - PostgreSQL
    - Perl
  - Use tools everybody uses
    - Google
    - Wikipedia
  - Use free tools
    - Munin
    - nagios

Sensors – the usual ones!

- Temperature
- Humidity
- Powermeters
  - Really power analysers
  - Real, reactive and appearant loads
  - Peak/burst draws
  - Frequency, voltage
  - Power factor

Power meters

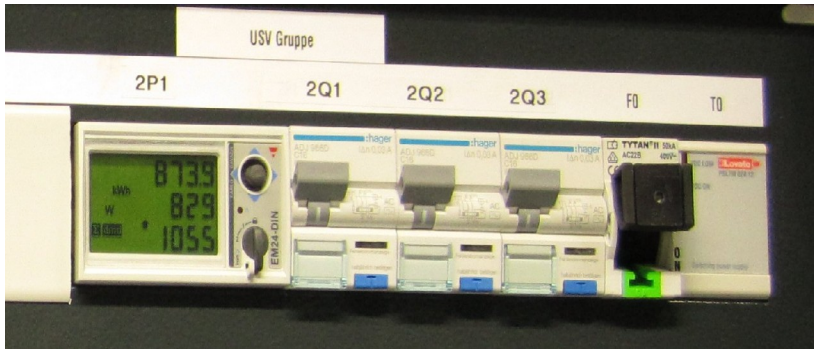

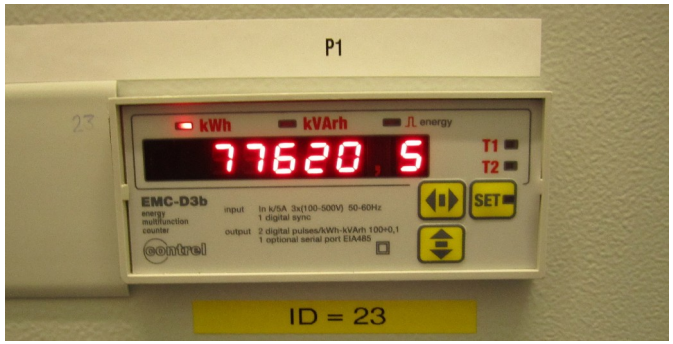

Sensors – not so usual...

- Water flow
- Status of valves
- Cold water chiller states
- Fire control system states
- Intrusion detection sensors
- Rack status sensors
- Video surveillance gear

Various ”Sensors”

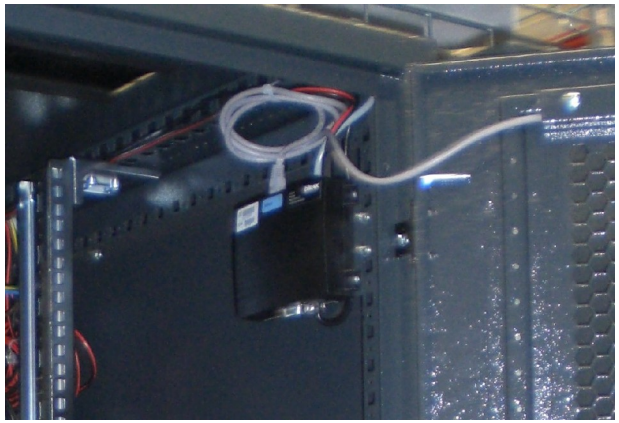

What is a sensor?

- A sensor is part of something bigger, usually a sensor gateway
- A sensor can be directly or indirectly assigned to a customer
- A sensor has metadata
  - Resolution
  - Scaling
  - SI unit
  - Multiplication factor

Use a database!

- PostgreSQL
  - ~412M rows right now
  - Partitioning
    - Per month/day
    - ”current value”
    - Per sensortype (more and less interesting sensors)
  - Preaggregation tables
    - 15min intervals
    - 24h intervals

Use a database!

- Database metadata
  - Business dependencies/rules
    - SLA limits (humidity, power usage, pricing)
    - Access control information (tokens, biometric access)
    - Basic contract data (end of contract,...)
    - Sensor thresholds
  - ”Mapping of objects”
    - Room names
    - Rack/Room names (customers != internal)

Use a database!

- Metadata
  - Logical dependencies
    - Recursive Structures (phase L1 is in compartment 41 of rack DC4-1F, that rack is in Row 1, that row is in Datacenter DC4, and that one in building Greencube
    - Logical groups - (powermeter 17 is in the feed for AC unit 1-4 and powermeter 213 is in the feed for the chiller → power usage for ”cooling system 1”)
  - technical dependencies
    - Breaker 1 is handling powerfeed XYZ

Feeding the System!

- Multiple bus systems
  - Usually wired serial Modbus (1979)
  - One gateway for every 10-100 devices
    - A device can have multiple sensors (coils/registers)
  - A Gateway bridges serial Modbus to Modbus/TCP
    - Somewhat ”documented”
    - Actual implementations very vendor specific
      - Endianess of encoded floats
      - Registers starting at 0 vs 1

Feeding the System!

- Perl modbus client
  - Custom code (libmodbus hard to use)
  - Reads metadata from databases/files
  - Quirks workarounds
  - Heavy serialization to avoid overloading gateways
  - ~250 lines of code, vs 10000€+ for commercial libraries

Use the data - part #1!

- Billing (the important one!)
  - Variable usage fee based on power draw
- Monitoring/Reporting
  - Munin
  - Nagios (passive and active checks)
  - Sales/Product planning
  - Other databases
  - Dashboards

Use the data - part #2!

- Phase balance (L1/L2/L3)
- Power usage patterns
  - Sudden surges or drops in power draw
  - close to fuse capacity
- Tell the vendor about reality (vs his ”we are green IT” marketing docs)
- Optimized the whole system (partial loads are tough!)
- Use the data to fix the data...

Example graphs...

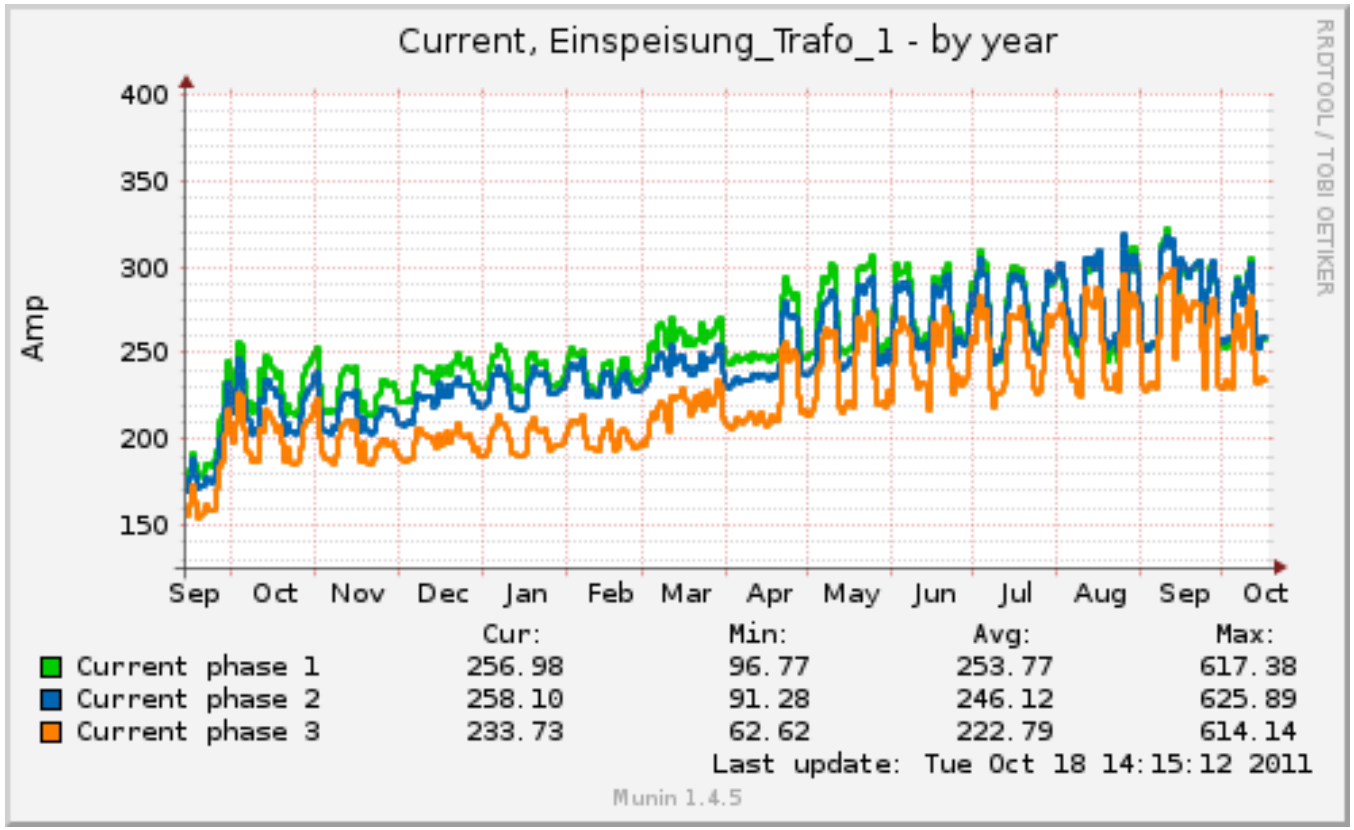

Example graphs...

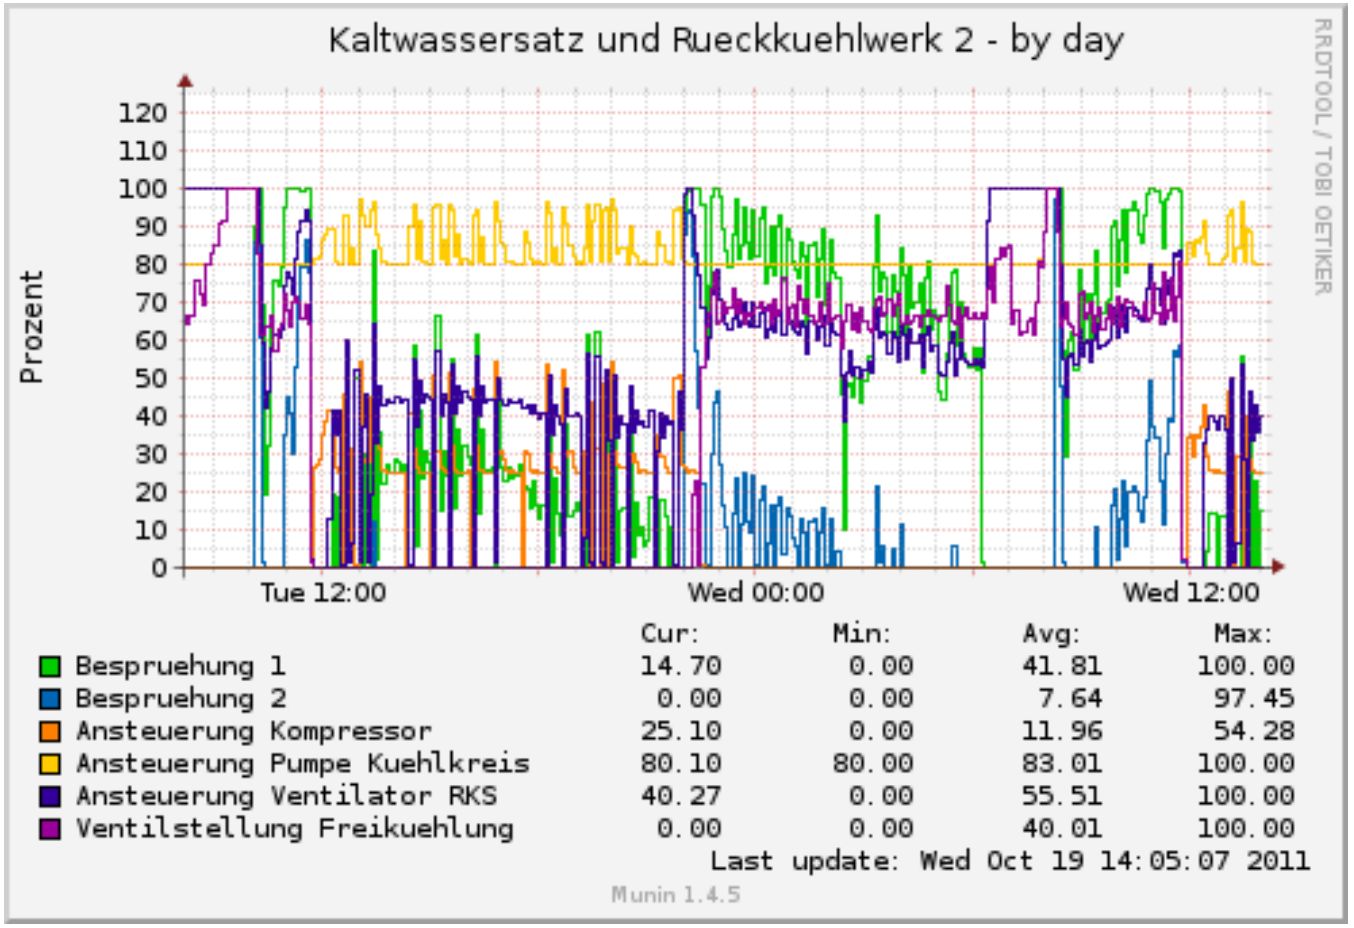

Example graphs...

Host	Service	Status	Last Check	Duration	Attempt	Status Information
Rueckkuehler_2	Ansteuerung_Befeuchter_1	OK	14:15:08	16d 1h 59m 4s	1/3	[0-100%] [11.4]
	Ansteuerung_Befeuchter_2	OK	14:15:08	2d 11h 34m 8s	1/3	[0-100%] [0]
	Ansteuerung_Ventilator	OK	14:15:08	1d 16h 49m 8s	1/3	[0-100%] [39.0]
	Aussentemp	OK	14:15:08	5d 4h 49m 4s	1/3	Temp OK [11.1C]
	Innentemp_Stufe_1	OK	14:15:08	11d 5h 9m 8s	1/3	Temp OK [30.3C]
	Innentemp_Stufe_2	OK	14:15:08	24d 10h 9m 8s	1/3	Temp OK [31.2C]
	Leckwaechter_HD-Pumpen	OK	14:15:08	13d 4h 39m 8s	1/3	No alarm [0]
	RL_Temp_Aussenbereich	OK	14:15:08	17d 16h 24m 5s	1/3	Temp OK [32.2]
	Registerheizung	OK	14:15:08	18d 10h 49m 8s	1/3	State: [0]
	Rueckkuehler_Alarm	OK	14:15:08	6d 8h 54m 6s	1/3	No alarm [0]
	Rueckkuehler_BM	OK	14:15:08	25d 3h 44m 8s	1/3	State: [1]
	Rueckkuehler_Frostwaechter	OK	14:15:08	30d 13h 39m 9s	1/3	No Alarm [0]
	Rueckkuehler_SM	OK	14:15:08	11d 0h 24m 9s	1/3	No failure [0]
	Ueberspannungsschutz	OK	14:15:08	15d 11h 59m 9s	1/3	No alarm [0]
	VL_Temp_Aussenbereich	OK	14:15:08	3d 9h 19m 8s	1/3	Temp OK [29.7]
	Ventilator_BM	OK	14:15:08	28d 19h 24m 9s	1/3	State: [1]

Data issues...

- Lots of sensor – (some) wrong readings
  - Powermeters suddenly report 0
  - Powermeters suddenly report random values

select max(kwh)-min(kwh) from datacenter.measurements_last_month where customer='foo'; → *(sometimes) wrong results*

  - Sensors do not report values at all
  - Scaling might change dynamically
    - Kwh/10 in one moment and kWh the next one...

Future direction

- Realtime display of data
  - Displays in the datacenter
  - Monthly reports for customers^managers
- External DW
  - Trending
  - Capacity planning/Power purchase
- A better dynamic graphing system
  - Visage/collectd

Thank you!

Questions?