

The Single Source of Truth for Network Automation



Andy Davidson <andy@asteroidhq.com>

Friday 15th June

TREX Workshop 2018, Tampere Finland

Automation Journey



Reporting

Most network engineers begin their automation journey by producing some simple reporting software. It is low-risk, has a positive useful impact, and a good introduction to network scripting and the many libraries that support network automation.

Automation Journey



Reporting

```
1  
2 Dear support,  
3  
4 We are discarding above threshold to the following ports, please see  
5 if the customers can be persuaded to upgrade?  
6  
7 Discarding towards ports  
8  
9  
10 10.3.3.4 thing.example.com cherry-sw1 Ethernet3/4  
11 10.4.3.1 lovely.example.com apple-sw1 Etherent2/1  
12  
13 Lovingly,  
14 Your switch robot
```

Most network engineers begin their automation journey by producing some simple reporting software. It is low-risk, has a positive useful impact, and a good introduction to network scripting and the many libraries that support network automation.

Automation Journey



Reporting



Tooling

Eventually, tasks which are repetitive, and simple to automate start to look like great candidates to automate. Engineers discover that the great libraries that integrate with software tools can be used to write as well as read configuration, and simple standalone tools are created.

Automation Journey



Reporting



```
[dhcp-15:~ andy$ ./new_peer.py --asn 12345 --ipaddr 10.1.1.4  
[dhcp-15:~ andy$  
Session with AS12345 configured at EXAMPLE_IX, State: Active  
dhcp-15:~ andy$
```

Tooling

Eventually, tasks which are repetitive, and simple to automate start to look like great candidates to automate. Engineers discover that the great libraries that integrate with software tools can be used to write as well as read configuration, and simple standalone tools are created.

Automation Journey



Reporting



Tooling



Application

More complex tools are eventually produced. Engineers begin to “configure the network and not the device”, so state becomes a problem (I mean state becomes properly managed). This takes the look and feel of a proper application.

Automation Journey



Reporting



Tooling



Application



More complex tools are eventually produced. Engineers begin to “configure the network and not the device” (I mean state becomes properly managed). This takes the look and feel of a proper application.

Automation Journey



Reporting



Tooling



Application



Business

The ultimate place to reach is a fully automated and integrated business with a set of processes enforced and delivered by software. “Configure the product, not the network”.

Generally solved by businesses with scale challenges (mass access, hosting) but now a commonplace medium sized ISP/IXP requirement.

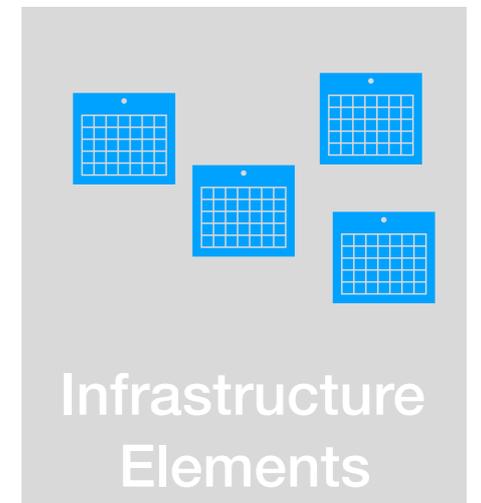
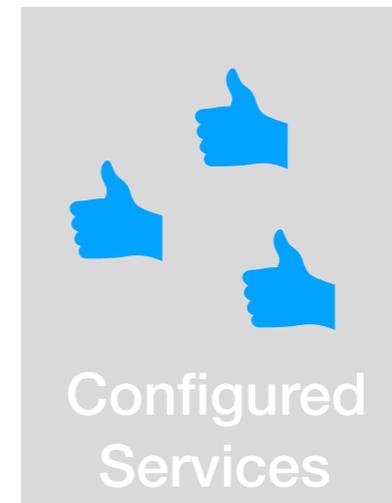
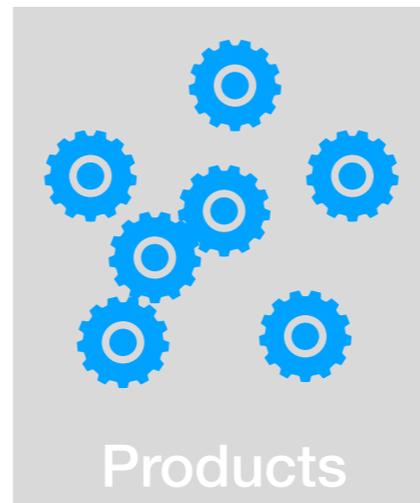
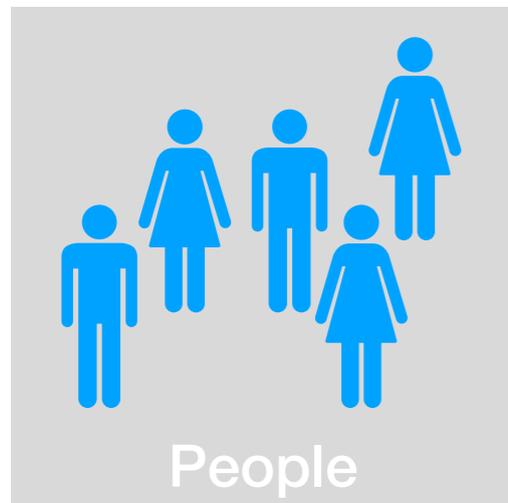
This presentation..

- Offer a technical perspective/thoughts on architecture on Greenfield deployment at the ‘automated business’ end of the spectrum
- What motivated this decision?
 - Replication - “as a service” product
 - Efficiency, leanness
 - Service assurance (rapid provisioning, ongoing high availability)
 - Integration with third party peering networks, Euro-IX, PeeringDB
 - Experience in this field, and frustration with traditional model
 - Chance to align business and technical process from the start - in our “DNA”

This presentation .. (2)

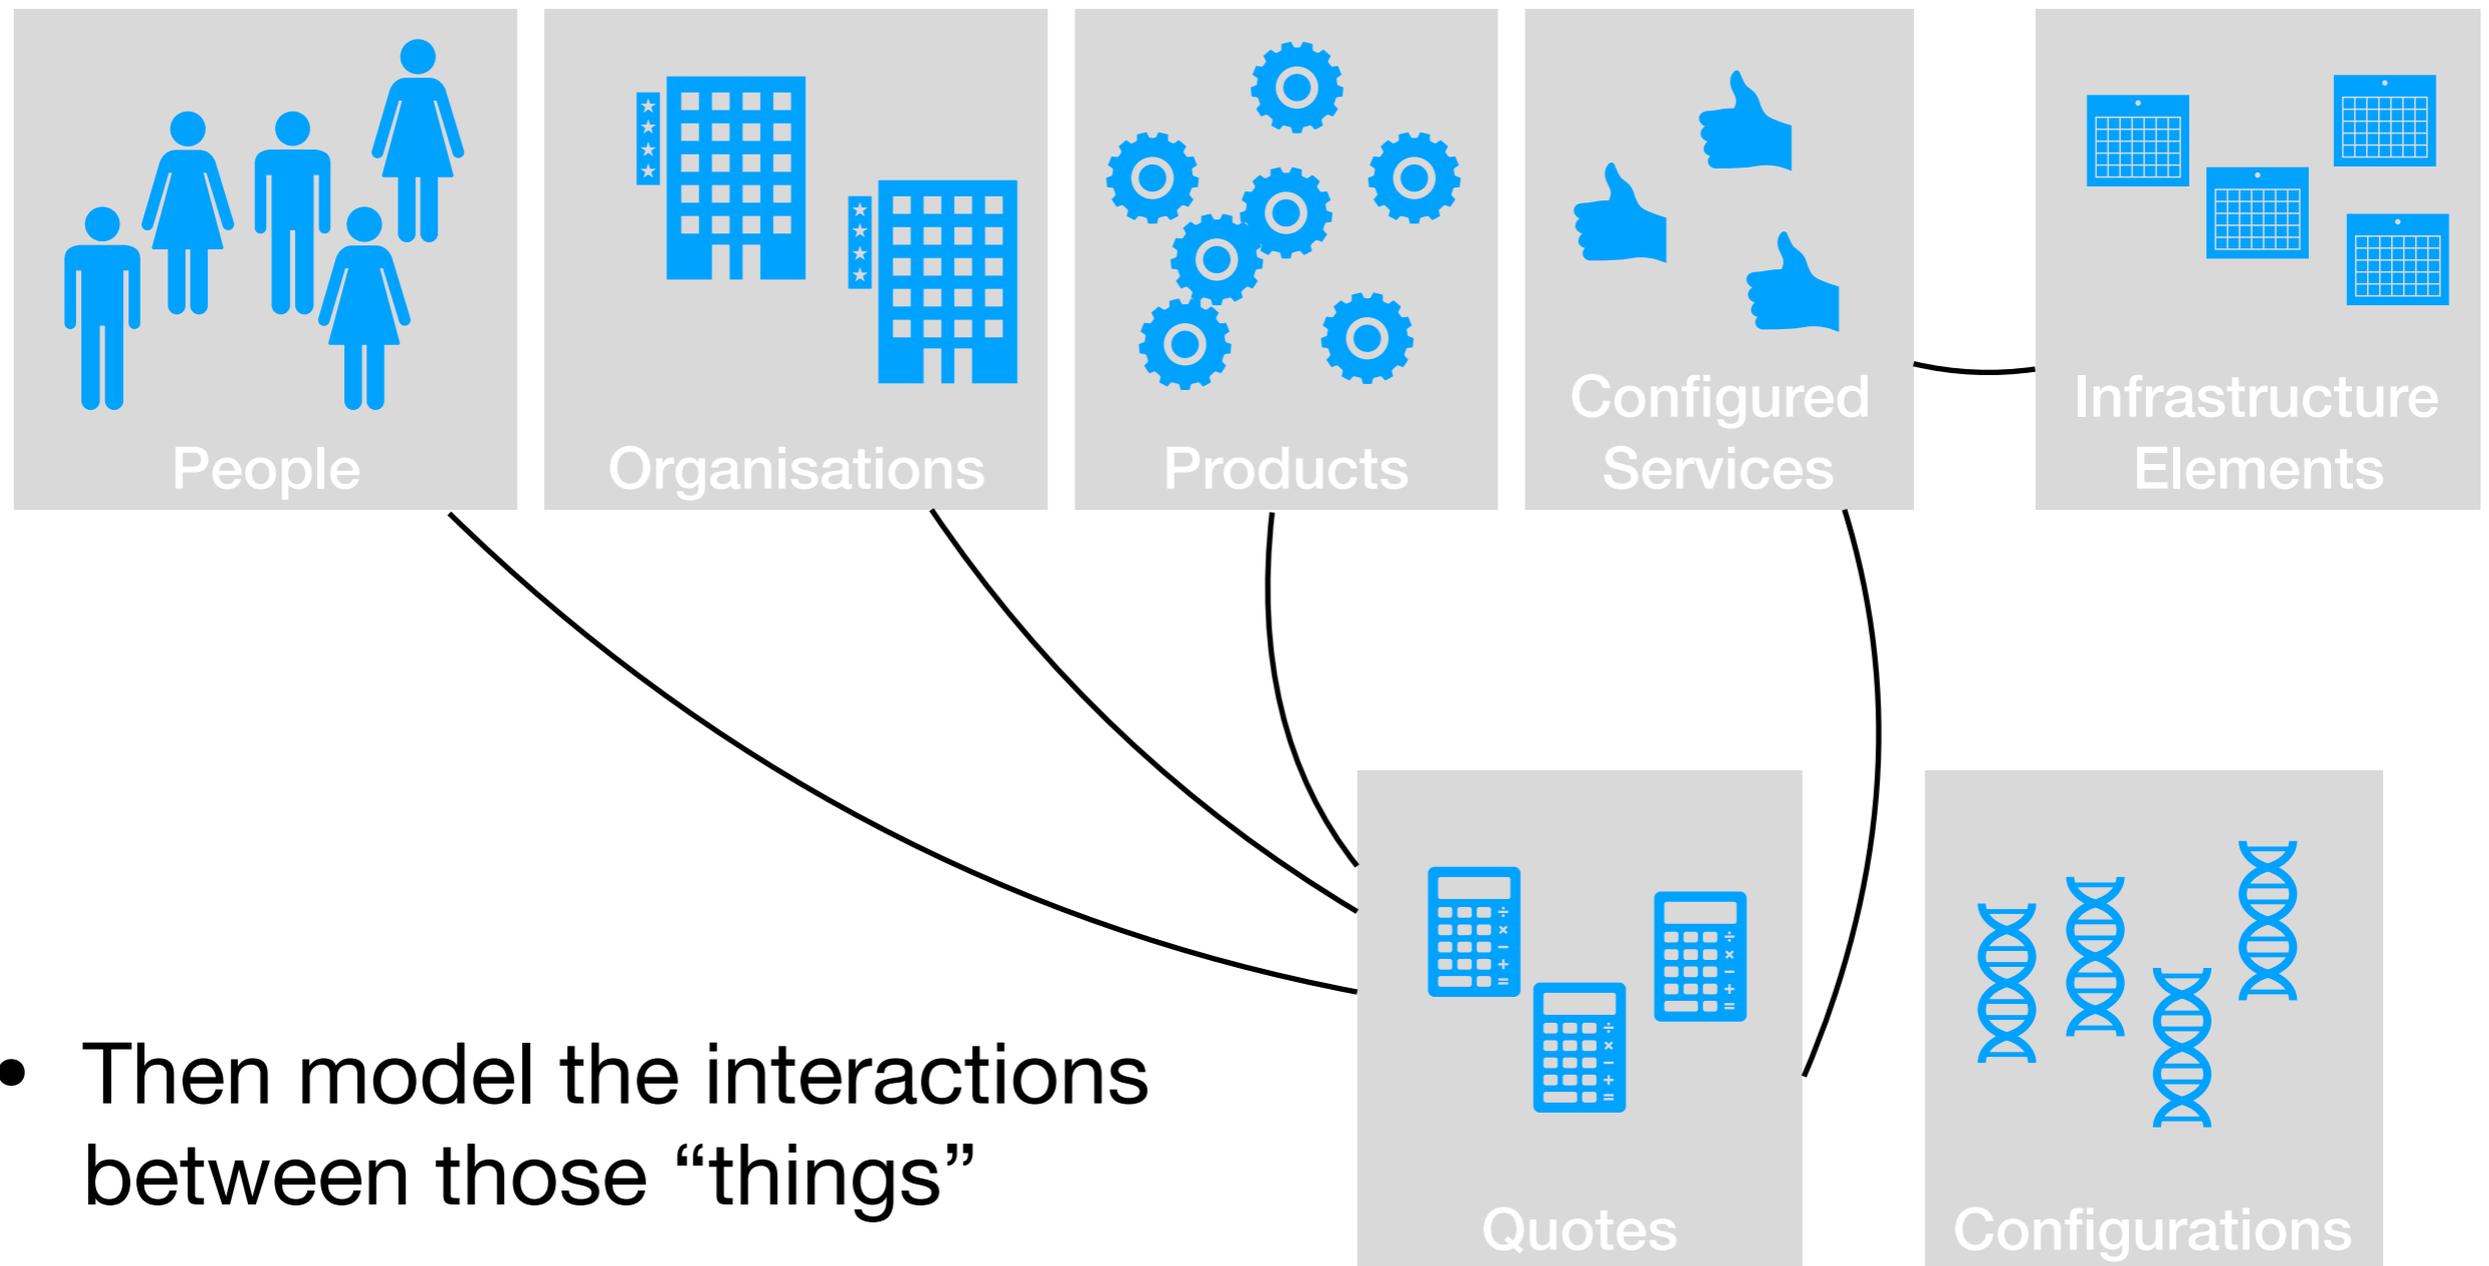
- Data model
 - Why and how to build a data model to support integrated automated business
- Software architecture for network centric businesses
 - Abstraction
 - APIs & API integration with customers
- Software testing
- Useful third party tools

What I mean, “data model”?

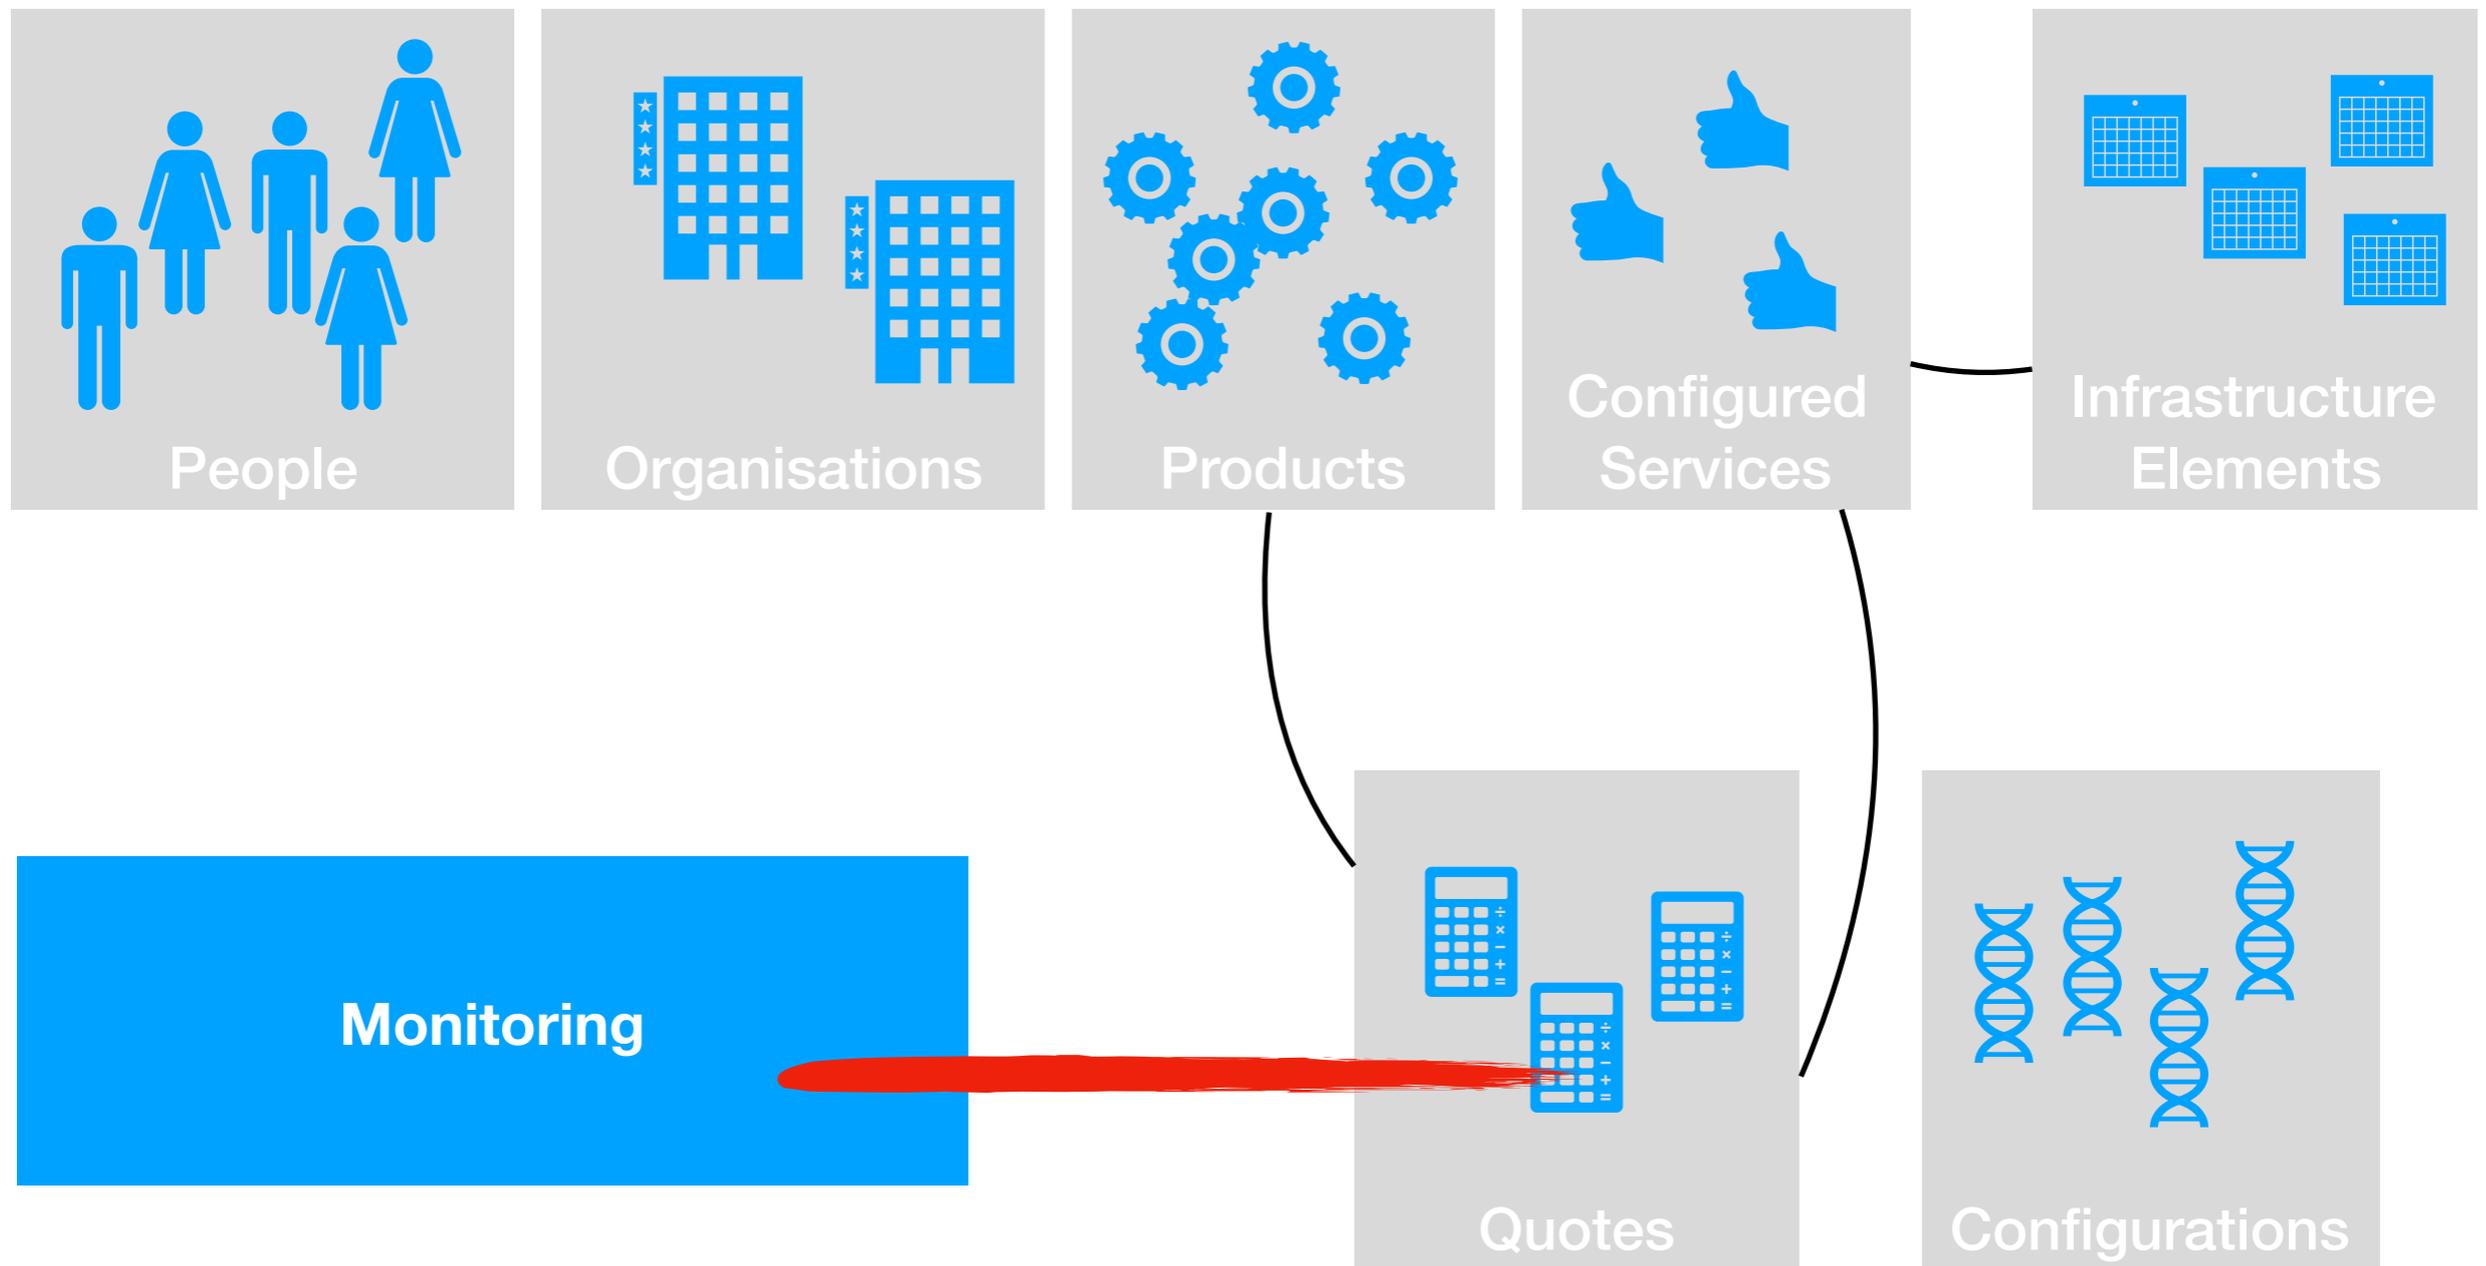


- A description of the things your business needs to ‘know’ in order to operate
- Start with the steady state of the business

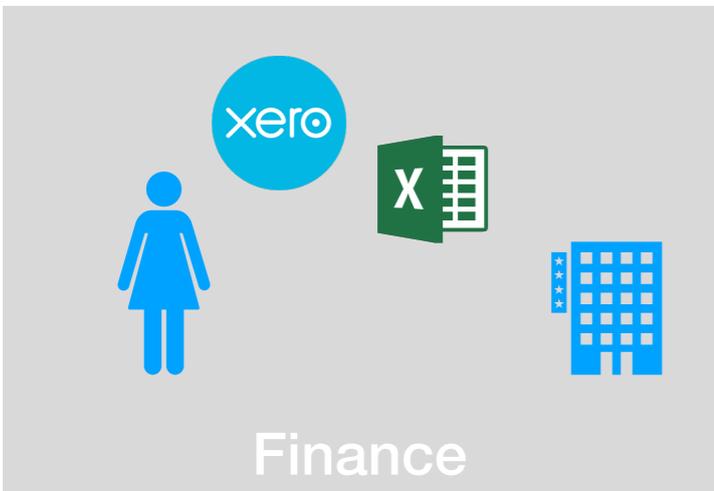
What I mean, “data model”?



Why to care from an engineering point of view?



Where does/shall data live?



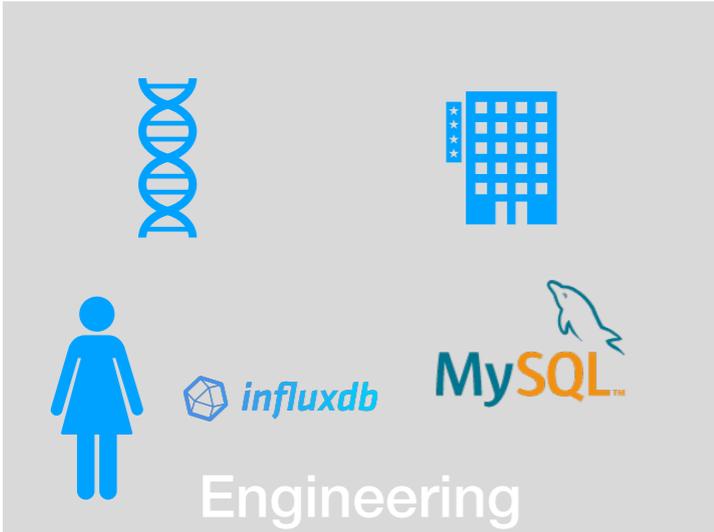
Finance

This block contains icons for a person, the Xero logo, an Excel spreadsheet icon, and a building icon, representing financial data storage.



Sales

This block contains icons for a thumbs up, an Excel spreadsheet icon, a building icon, a gear, a calculator, and a person icon, representing sales data storage. The SUGAR CRM logo is also present.



Engineering

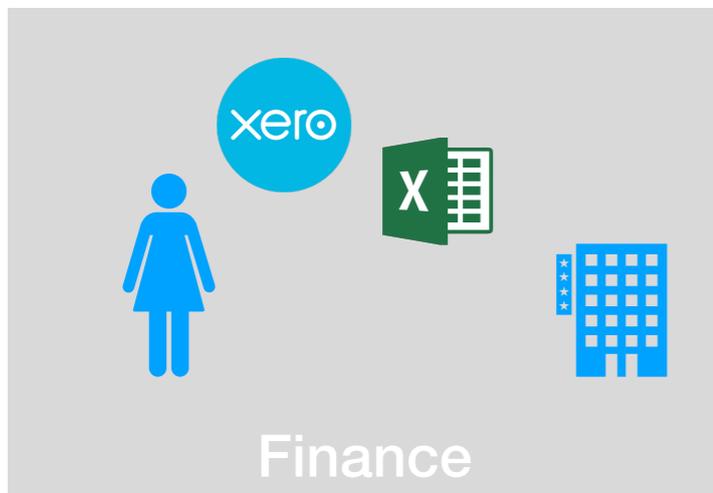
This block contains icons for a DNA helix, a building icon, a person icon, the InfluxDB logo, and the MySQL logo, representing engineering data storage.



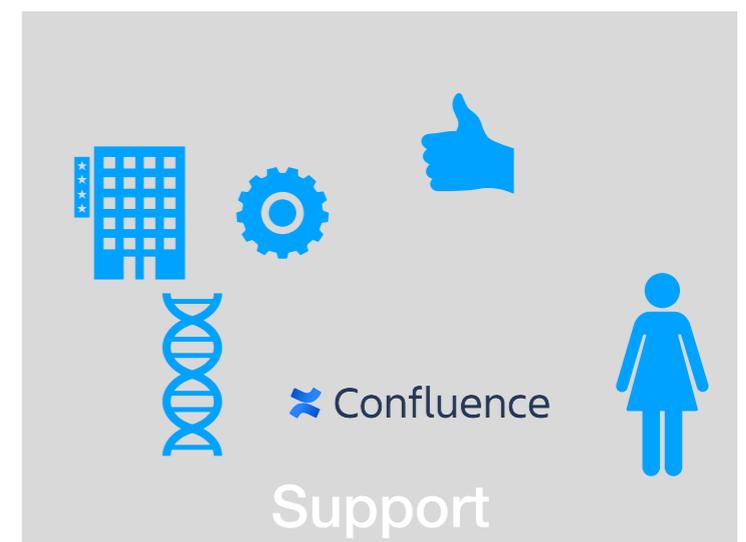
Support

This block contains icons for a building icon, a gear, a thumbs up, a DNA helix, the Confluence logo, and a person icon, representing support data storage.

Where does/shall data live?

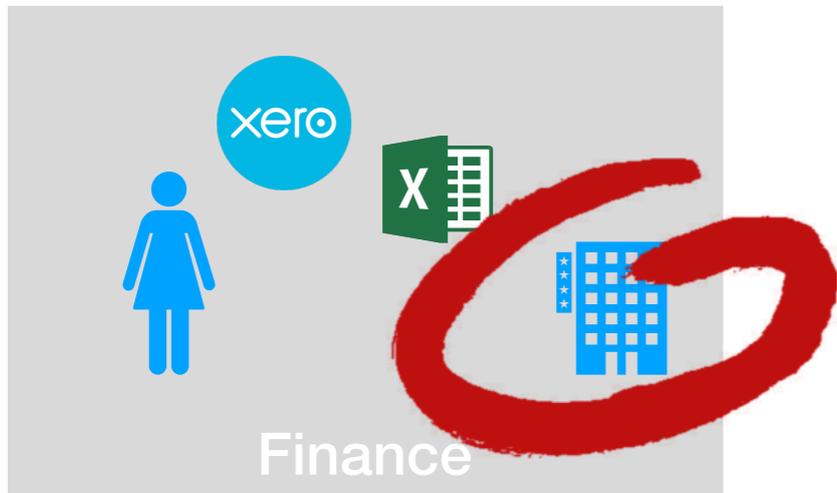


Fundamentally it is *fine* for data to “live” in different tools and databases

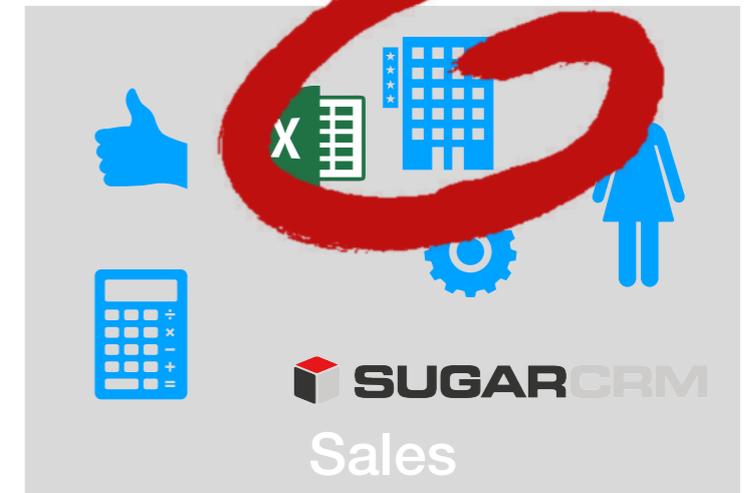


Where does/shall data live?

Search Engine, Inc.

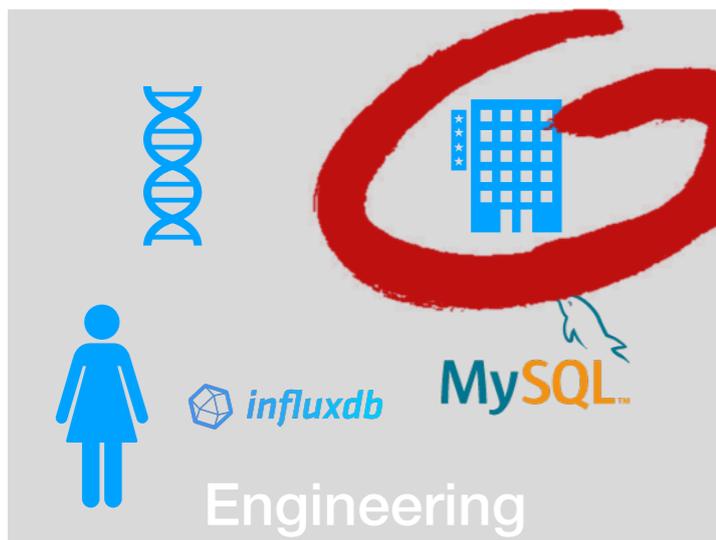


Finance

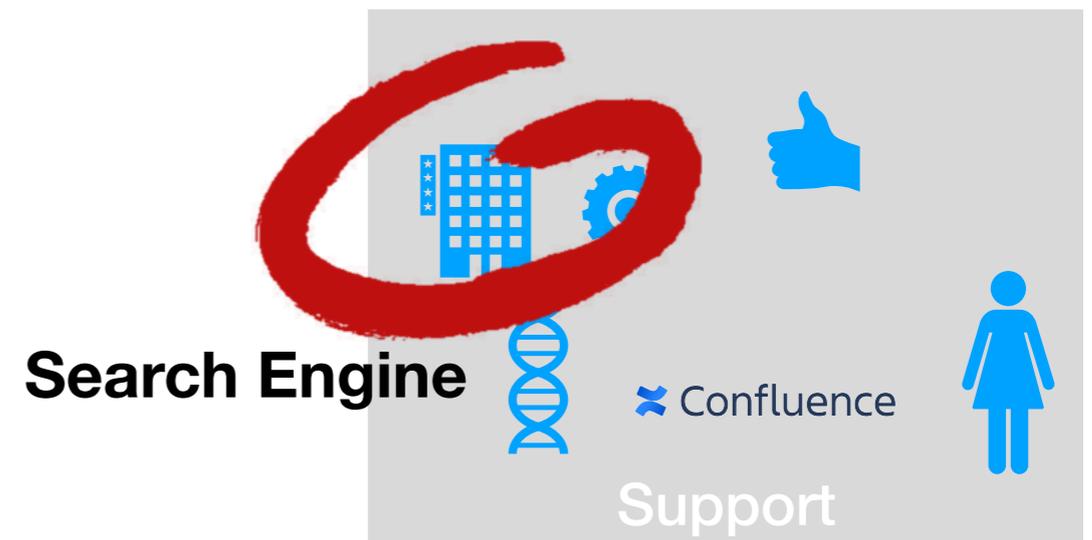


Sales

Search Engine (Netherlands) B.V.



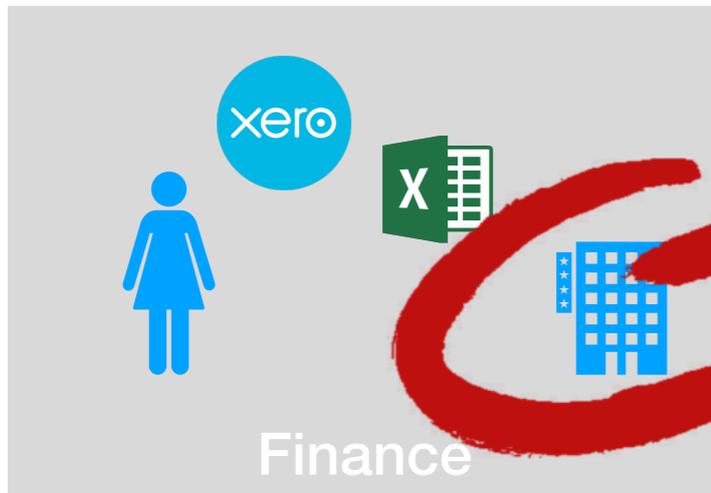
Engineering



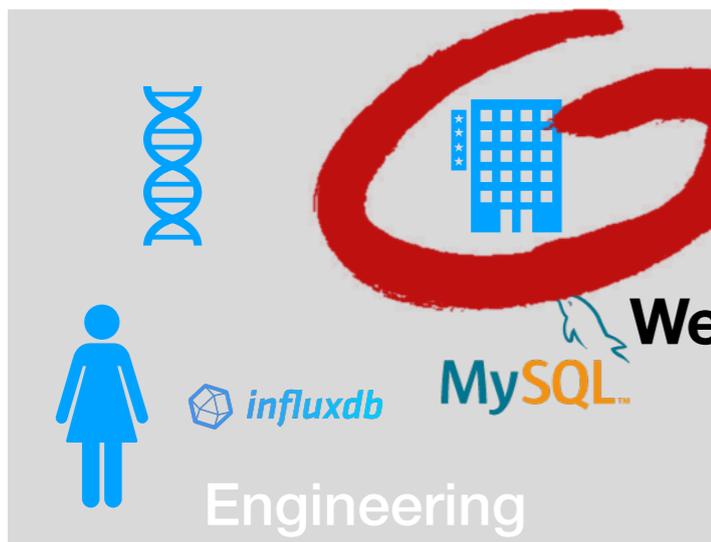
Support

Where does/shall data live?

Search Engine, Inc.



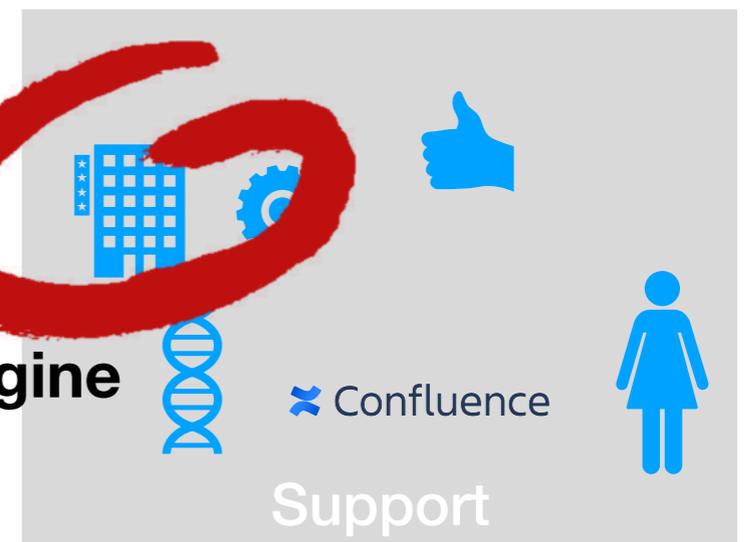
Search Engine (Netherlands) B.V.



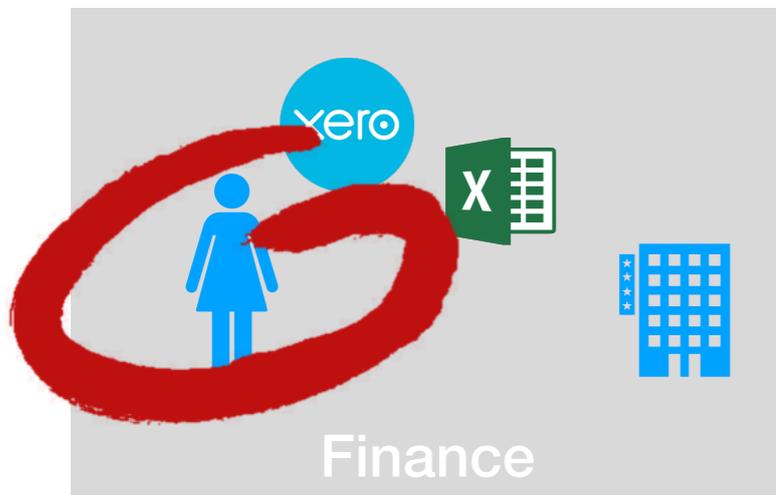
We just deal with Fred



Search Engine



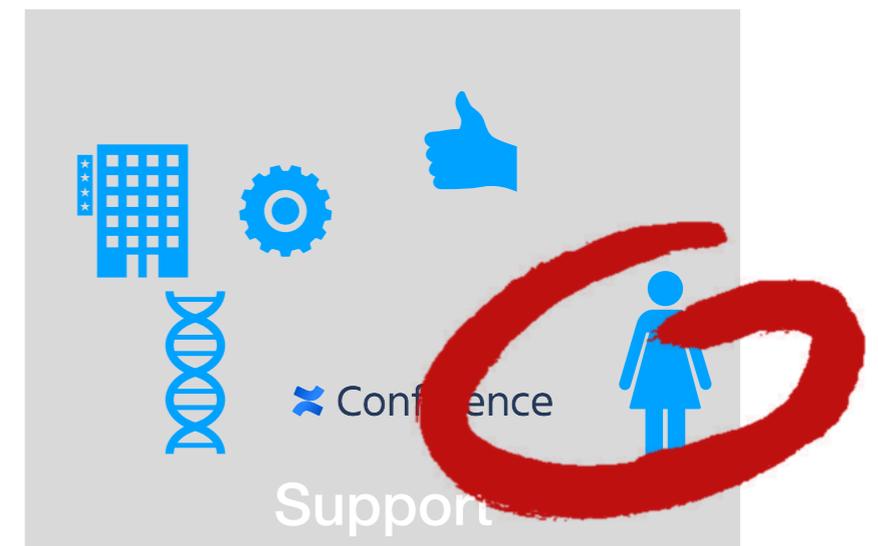
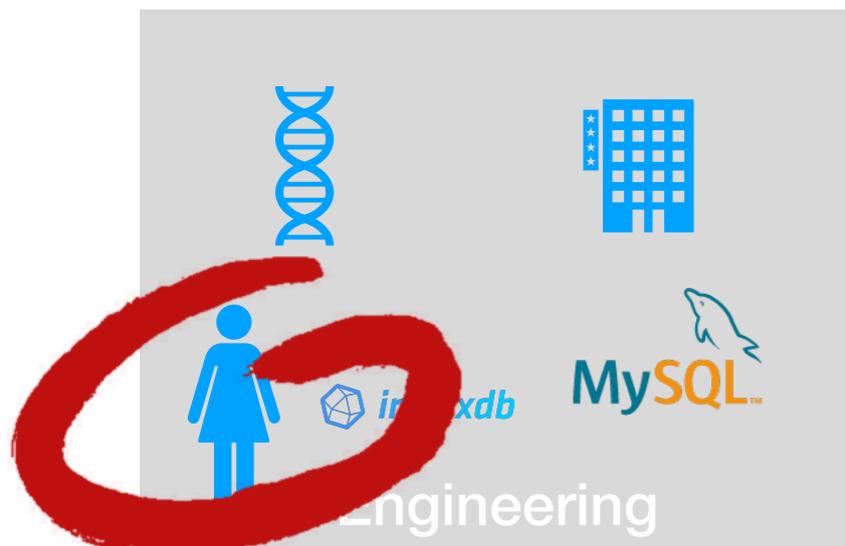
Where does/shall data live?



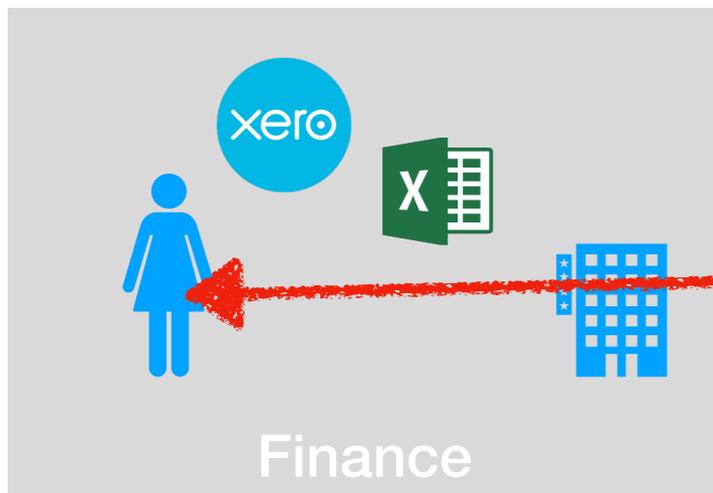
Fundamentally it is *not fine* for more than one data place to be authoritative for any single type of record



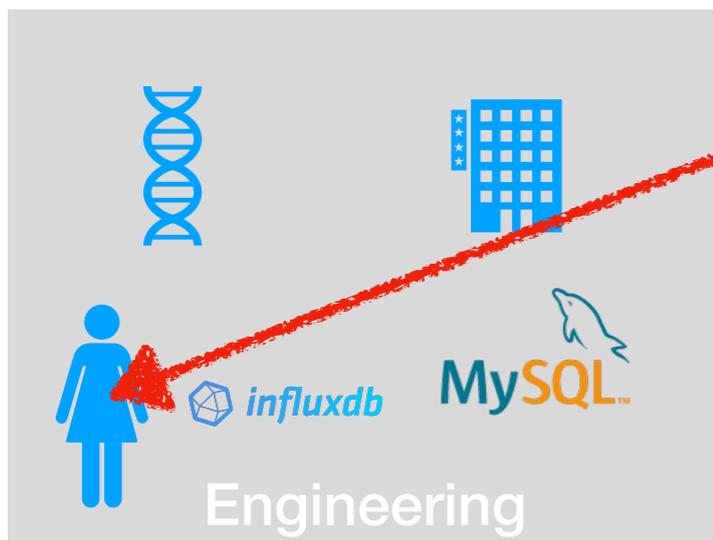
The other databases must refer to the key (id) of a single authoritative source



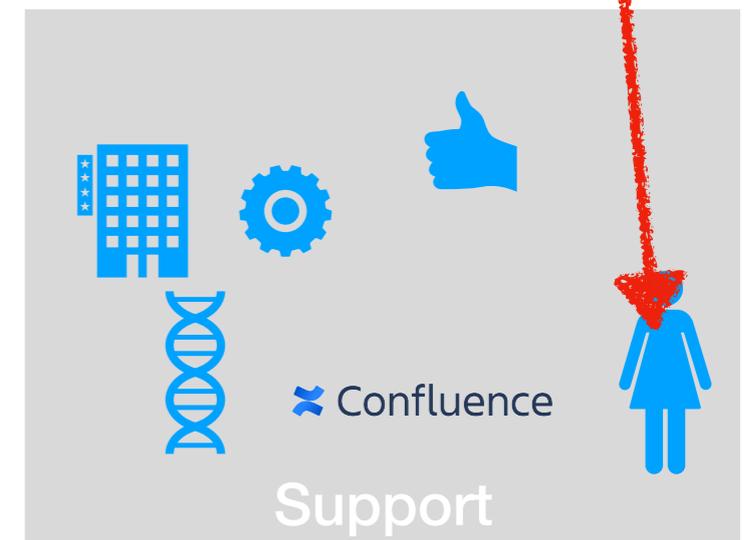
Where does/shall data live?



Fundamentally it is *not fine* for more than one data place to be authoritative for any single type of record



The other databases must refer to the key (id) of a single authoritative source



We will talk about how to configure and enforce that shortly.

Rules of Engagement

- Store any item of data **ONCE**
 - Easy to ensure that it is correct
 - “Third normal form”
- Give every record a **unique** ID which has nothing to do with the record
 - (ASN is not to be used as ID!)
- Decide **where** it will be authoritative
- Requires buy in and planning from across the business.

Separate your customer/ infrastructure data

Port

port_id

customer_id

bridge_id

port_name

Service

service_id

port_id

service_item1

product_id

Ensure infrastructure centric and customer-centric data is not in the same table
This will make your data substantially easier to maintain in terms of portability

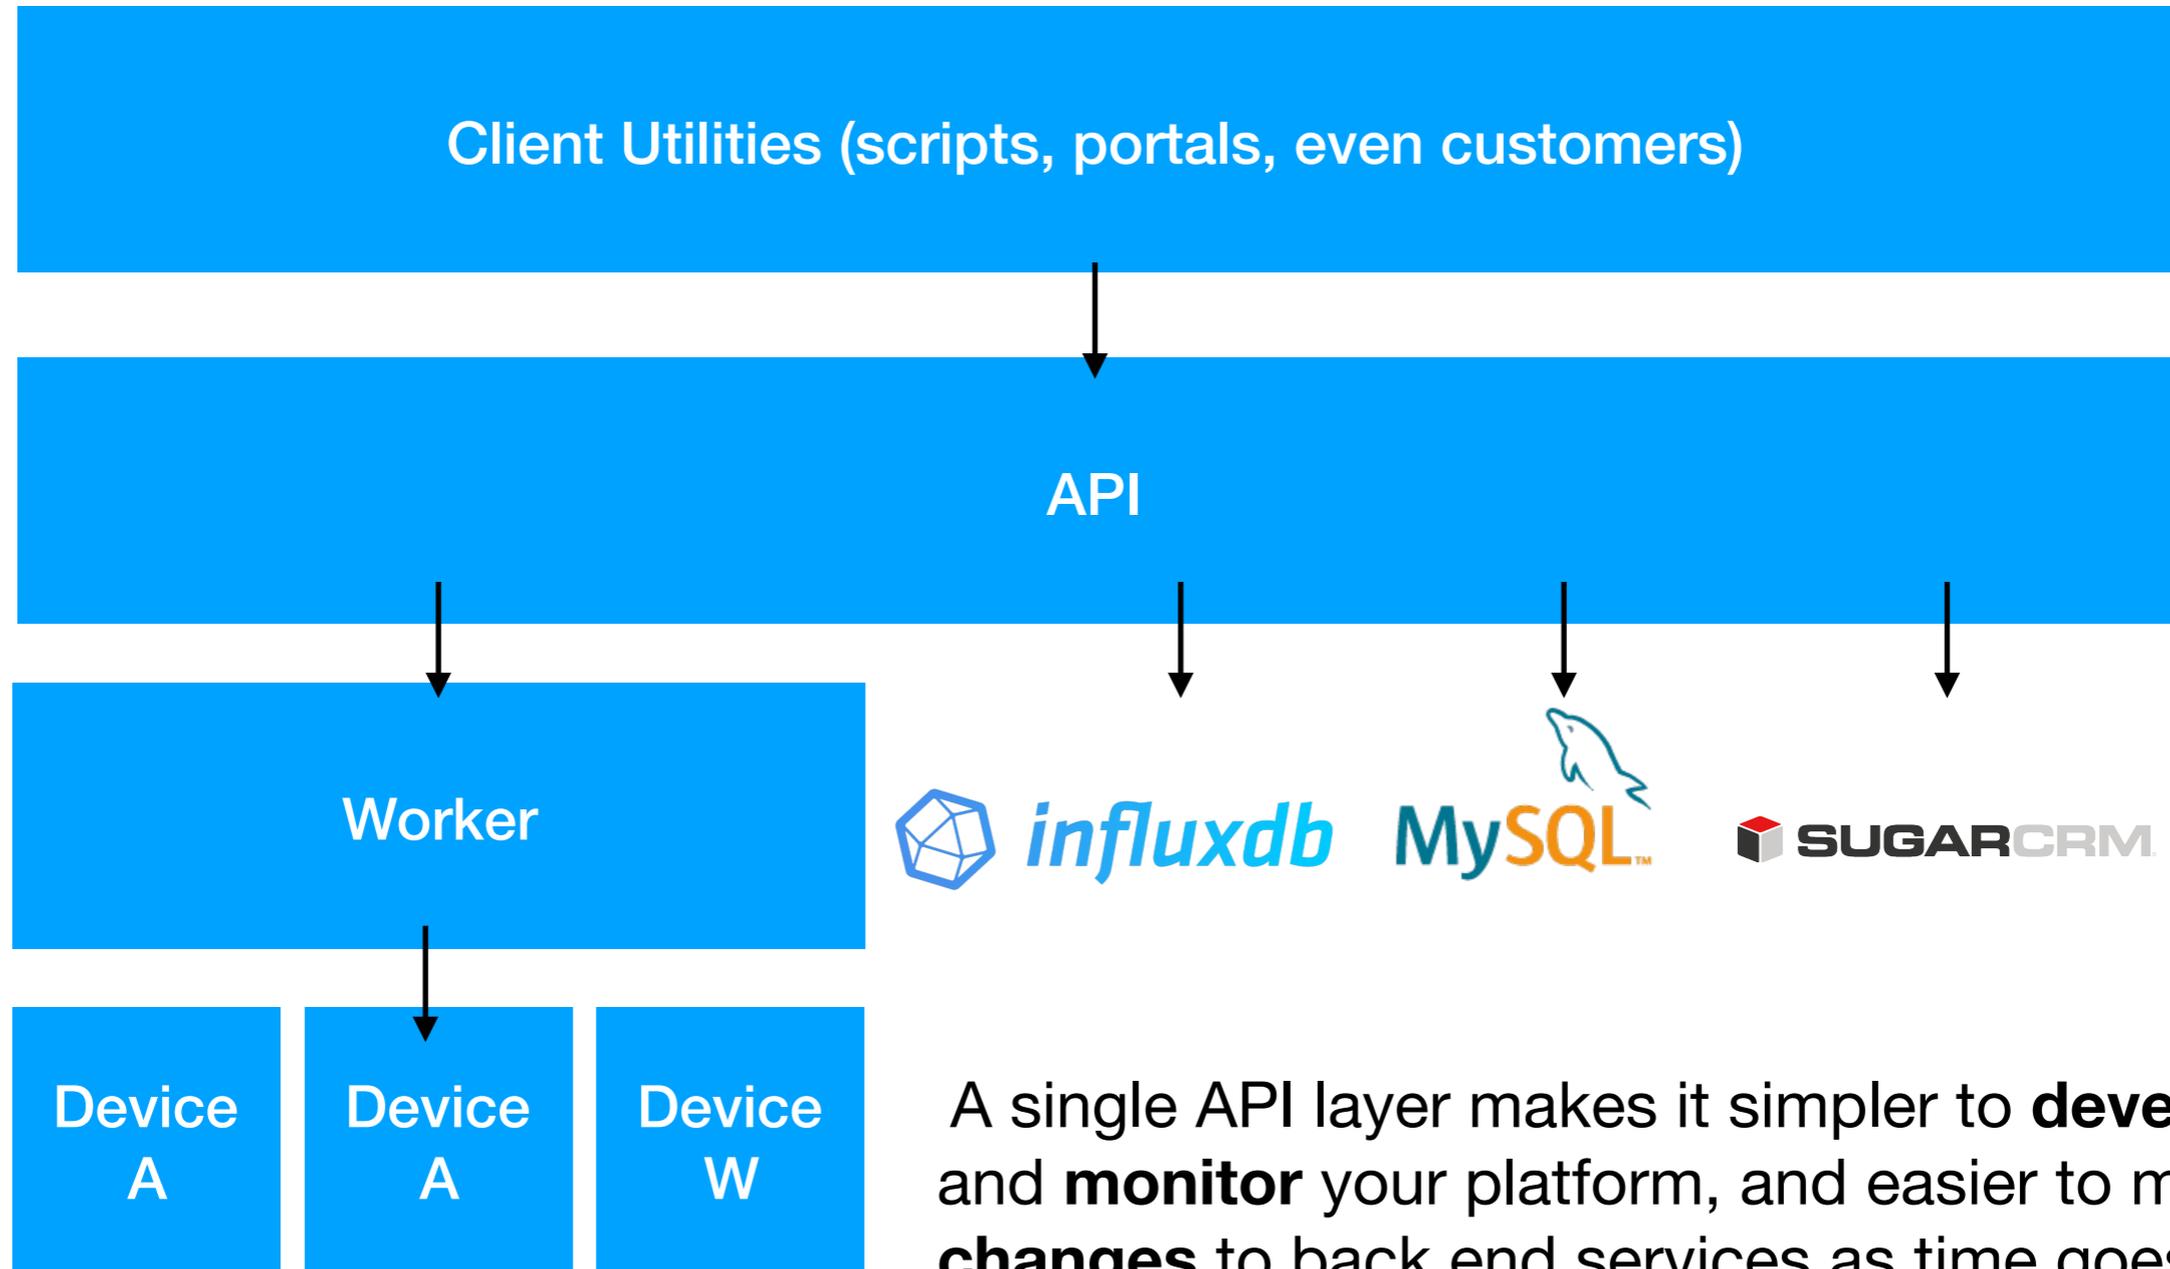
Database Fashions

- Document store -vs- RDBMS
 - Developers like document stores because they are very extensible and less strict
 - “Storage” cost reduced, so now we can be lazy
- Strict is a benefit / feature

Common Data Stores in Engineering

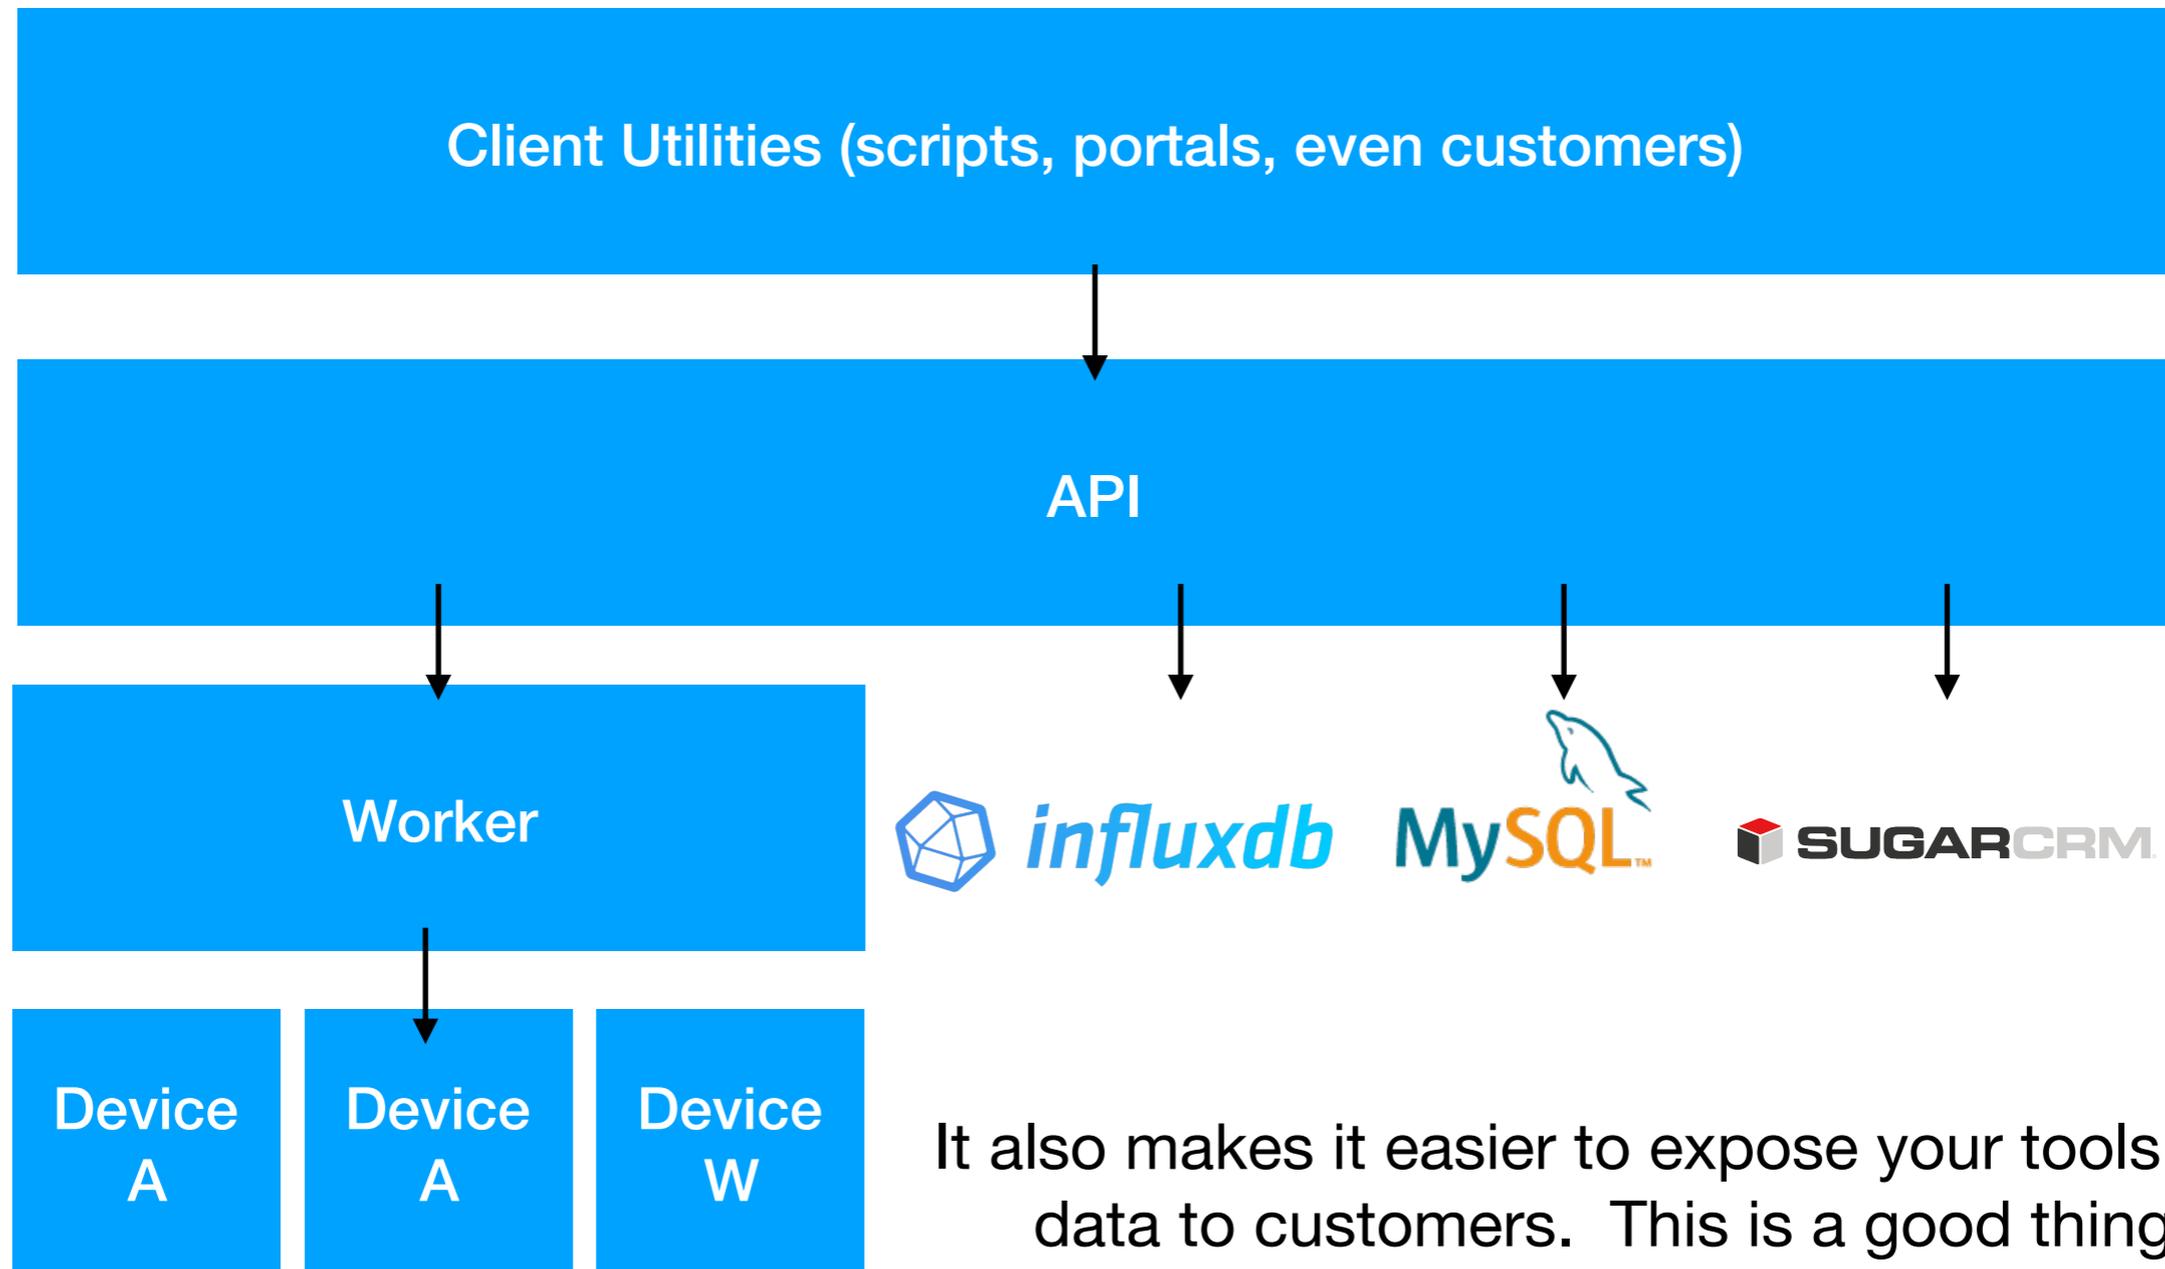
- SQL - Truths about users, ports, services, 'state', e.g. MySQL
- Time Series - e.g. Port utilisation, light level, error count, e.g. InfluxDB
- Third Party - Someone else's sorted data, e.g. CRM, e.g. EuroIX/PeeringDB

General Architecture



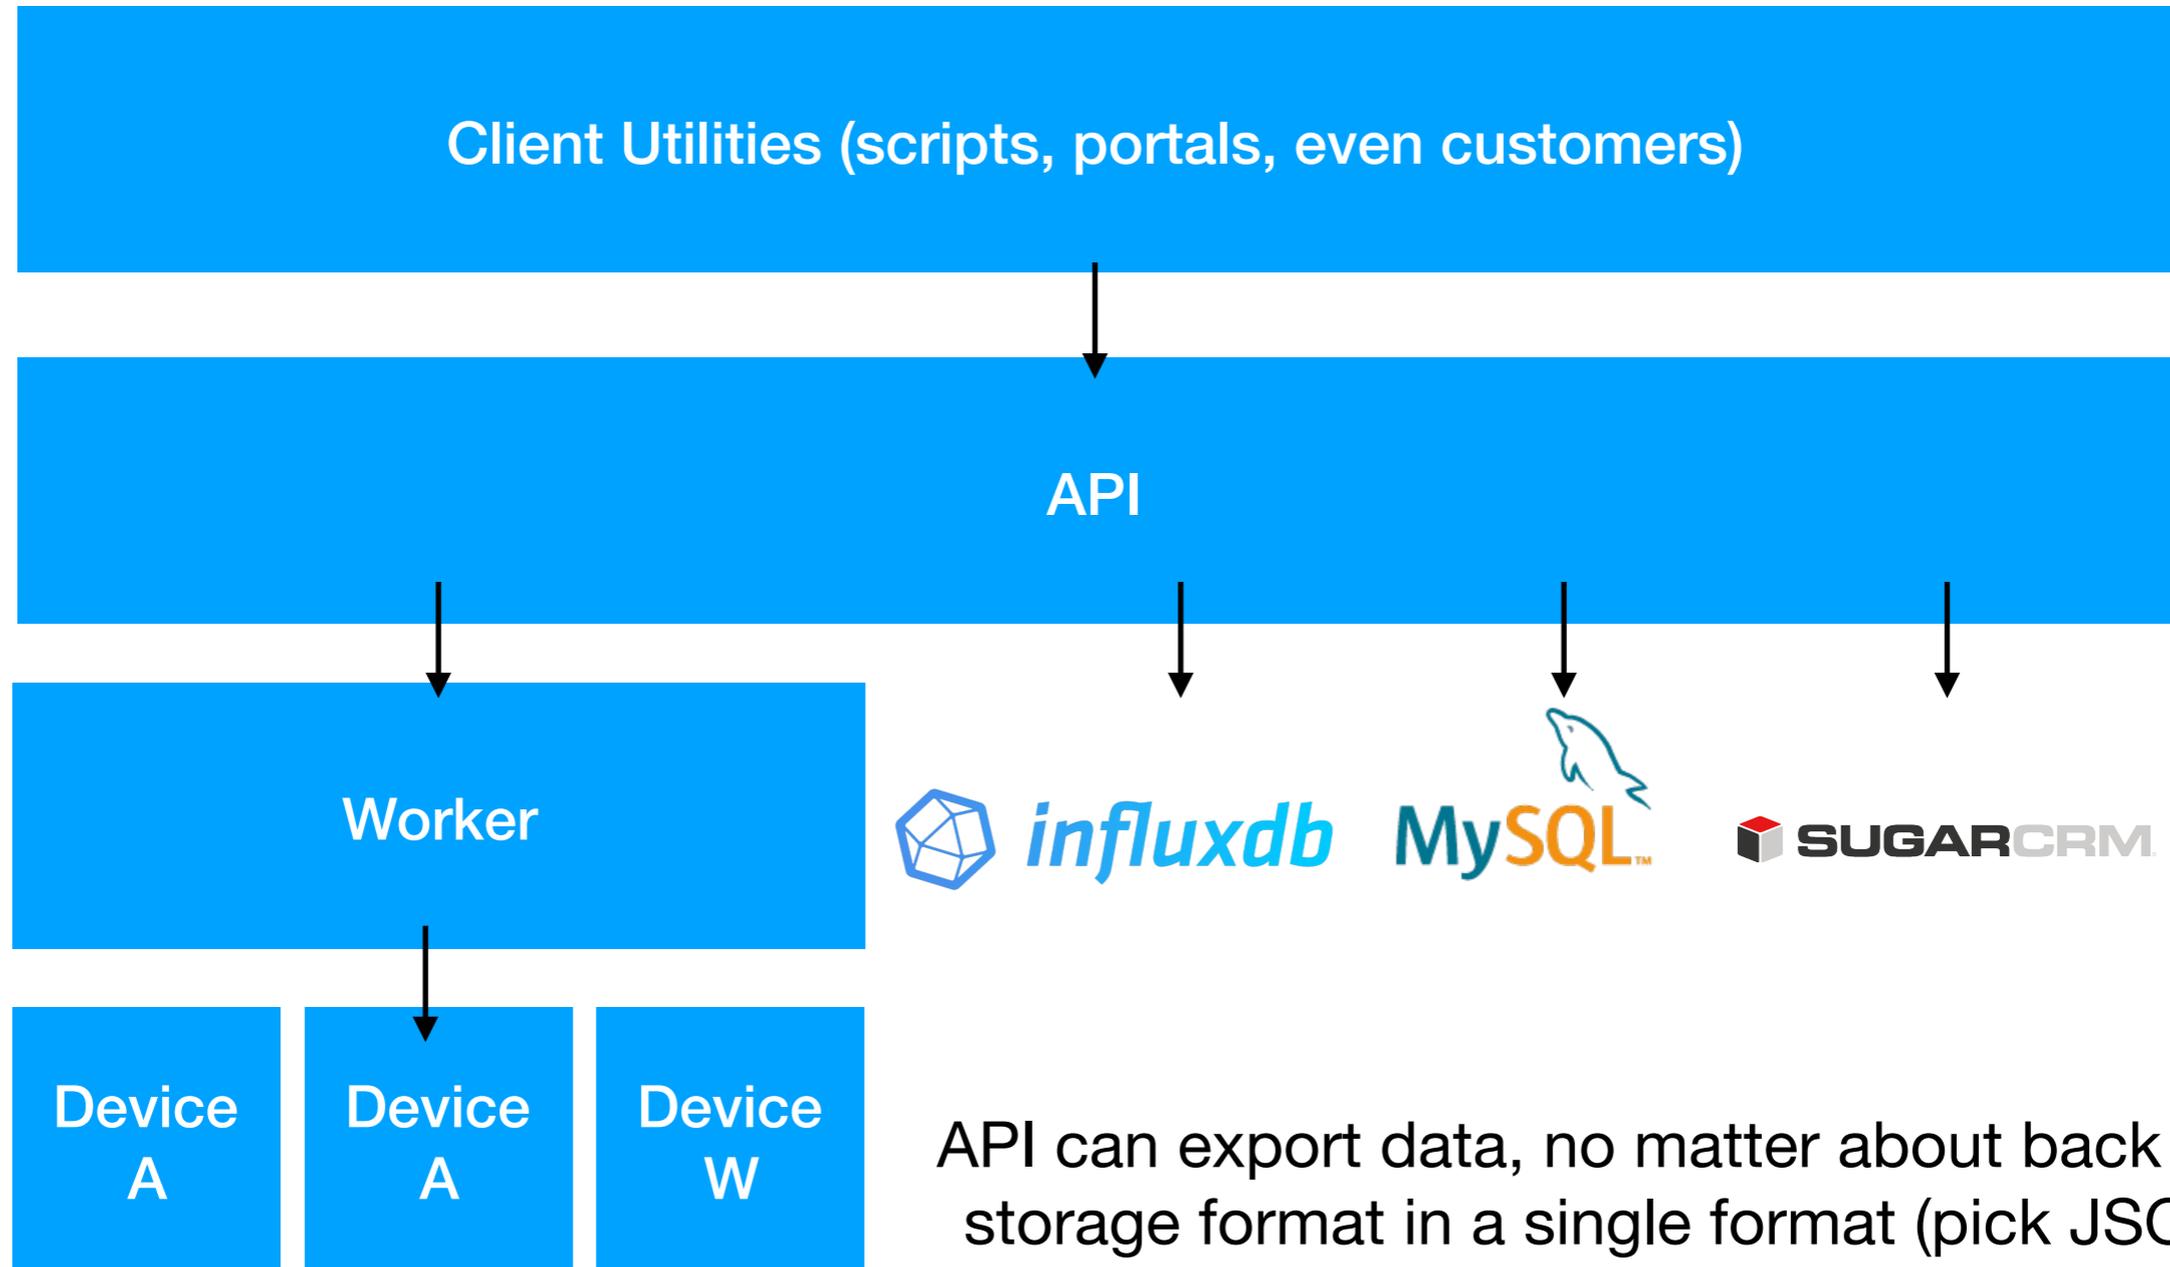
A single API layer makes it simpler to **develop** and **monitor** your platform, and easier to make **changes** to back end services as time goes by

General Architecture



It also makes it easier to expose your tools and data to customers. This is a good thing!

General Architecture



Worker, BIRD

Internal SQL

```
1 {
2   "asn": ,
3   "collector_bgp4_last_changed": "2018-01-07 12:39:39",
4   "collector_bgp4_observed_time": "2018-01-08 21:10:27",
5   "collector_bgp4_routes": 11,
6   "collector_bgp4_state": "Established",
7   "collector_bgp6_last_changed": "2018-01-07 12:39:40",
8   "collector_bgp6_observed_time": "2018-01-08 21:10:27",
9   "collector_bgp6_routes": 5,
10  "collector_bgp6_state": "Established",
11  "exchange_id": 1,
12  "exchange_name": "Asteroid Amsterdam Internet Exchange",
13  "exchange_shortcode": "ams",
14  "mac_address": ,
15  "network_id": 7,
16  "organisation_id": ,
17  "peering_dnsname": ,
18  "peering_ip6addr": "2001:7f8:b6::",
19  "peering_ipaddr": "185.1.",
20  "port_id": 7,
21  "route_server_config": 1,
22  "routeserver_bgp4_last_changed": "2018-01-08 11:02:08",
23  "routeserver_bgp4_observed_time": "2018-01-08 20:45:58",
24  "routeserver_bgp4_routes_filtered": 3,
25  "routeserver_bgp4_routes_imported": 7,
26  "routeserver_bgp4_state": "Established",
27  "routeserver_bgp6_last_changed": "2018-01-08 20:34:06",
28  "routeserver_bgp6_observed_time": "2018-01-08 20:45:58",
29  "routeserver_bgp6_routes_filtered": 0,
30  "routeserver_bgp6_routes_imported": 5,
31  "routeserver_bgp6_state": "Established",
32  "service_id": ,
33  "vlan_8021q_tag": "101",
34  "vlan_id": 1,
35  "vlan_name": "AMS_PEERING"
36 }
```

GET https://sputnik.asteroidhq.com/port/2

Pretty Raw Preview JSON

```
38   "port_enabled": 1,
39   "port_first_link_enable_time": "6th-February-2018 17:29:53",
40   "port_id": 2,
41   "port_identifier": "Ethernet1/2",
42   "port_original_lineitem_id": 1,
43   "port_original_quote_id": 1,
44   "port_repeatbill_id": 1,
45   "port_repeatbill_mrc": 0,
46   "port_speed_name": "10000",
47   "port_supervisor_lock": 0,
48   "port_switch_admin_enable": 1,
49   "port_switch_last_flapped_time": "13th-January-2018 00:01:49",
50   "port_switch_link_enable": 1,
51   "port_switch_observed_time": "2nd-March-2018 14:51:06",
52   "port_type": "physical",
53   "stats": {
54     "backtime": "1w",
55     "backtime_days": "7",
56     "bps_in": [
57       [
58         "2018-02-23T14:30:00Z",
59         1198.0533333333333,
60         1519396200
61       ],
62       [
63         "2018-02-23T15:00:00Z",
64         1336.2933333333333,
65         1519398000
66       ],
67       [
68         "2018-02-23T15:30:00Z",
69         1231.8666666666666,
70         1519399800
71       ],
72       [
73         "2018-02-23T16:00:00Z",
```

Worker, Arista

InfluxDB

Models <> Templates

- Once you have confidence in your data model you can harness the power of templated configuration
- Once your data model extends across the business you can do that with greater accuracy and devolved control
- e.g. at Asteroid, our sales people can deliver exchange ports directly from the **quotation**
 - So can customers
 - Simultaneous delivery of monitoring from the quotation

Automation fire triangle

Structured Data

Automation

Templates

Templates - Jinja

```
28 {% for network in exchange_data.json.member_list %}
29 ### AS{{ network.asnum }} - {{ network.name }}
30
31 protocol bgp pb_as{{ network.asnum }} {
32     description "Collector for AS{{ network.asnum }} - {{ network.name }} ";
33     local as collectorasn;
34     source address collectoraddress;
35     neighbor {{ network["connection_list"][0]["vlan_list"][0]["ipv4"]["address"] }} as {{ network.asnum }};
36     import all;
37     export none;
38 }
39
40 {% endfor %}
```

- Generate any kind of configuration
- Takes variables from your JSON API
- Facilitates programatic methods in configuration strophe
 - Loops
 - Conditionals

Automation - Ansible

```
2
3 - name: get route-server settings for this ixp
4   uri: url="https://sputnik.asteroidhq.com/export/euroix/{{ exchange_id }}/participants.json" body=yes
5   register: exchange_data
6   tags:
7     - softconfig
8
9 - name: install bird IPv4 config
10  template: >
11    src=templates/collector4.conf.j2
12    dest=/etc/bird/bird.conf
13    owner=root
14    group=root
15    mode=0640
16    vars:
17      exchange_data: {{ exchange_data }},
18  register: bird4_changed
19  tags:
20    - softconfig
21
22 - name: restart bird (IPv4)
```

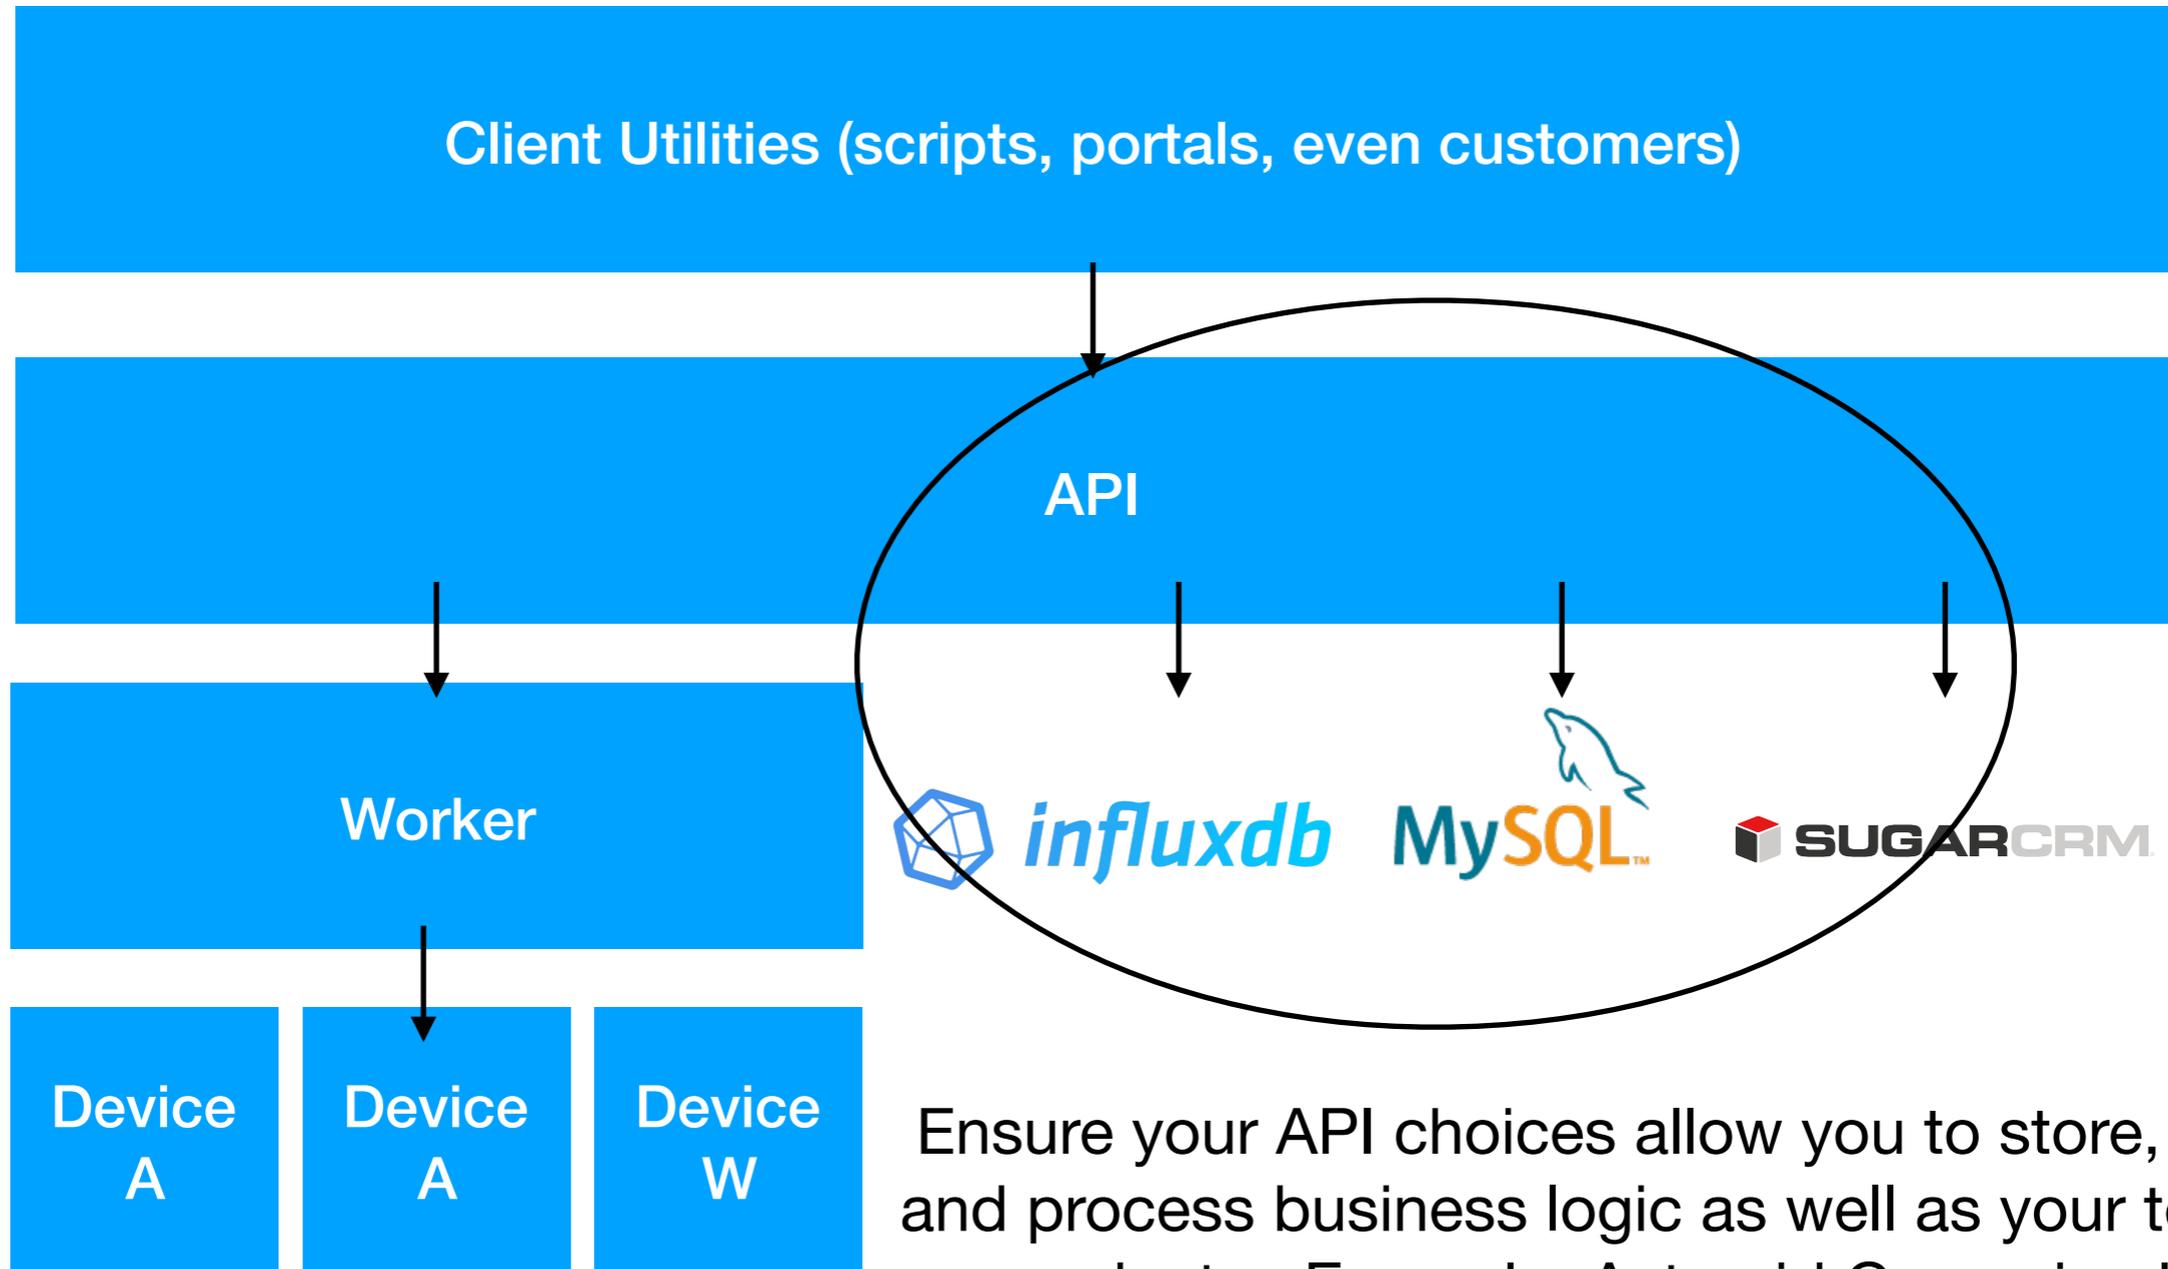
Conditional Logic without script

```
3 - name: Collect Prefix information
4   uri: url='http://46.51.199.18/prefix/{{ item["connection_list"][0]["vlan_list"][0]["ipv4"]["as_macro"] }}' body=yes
5   register: prefix
6   with_items: "{{ exchange_data['json']['member_list'] }}"
7   when: item['connection_list'][0]['vlan_list'][0]['ipv4']['routeserver'] == true
8   tags:
9     - softconfig
```

Advantages

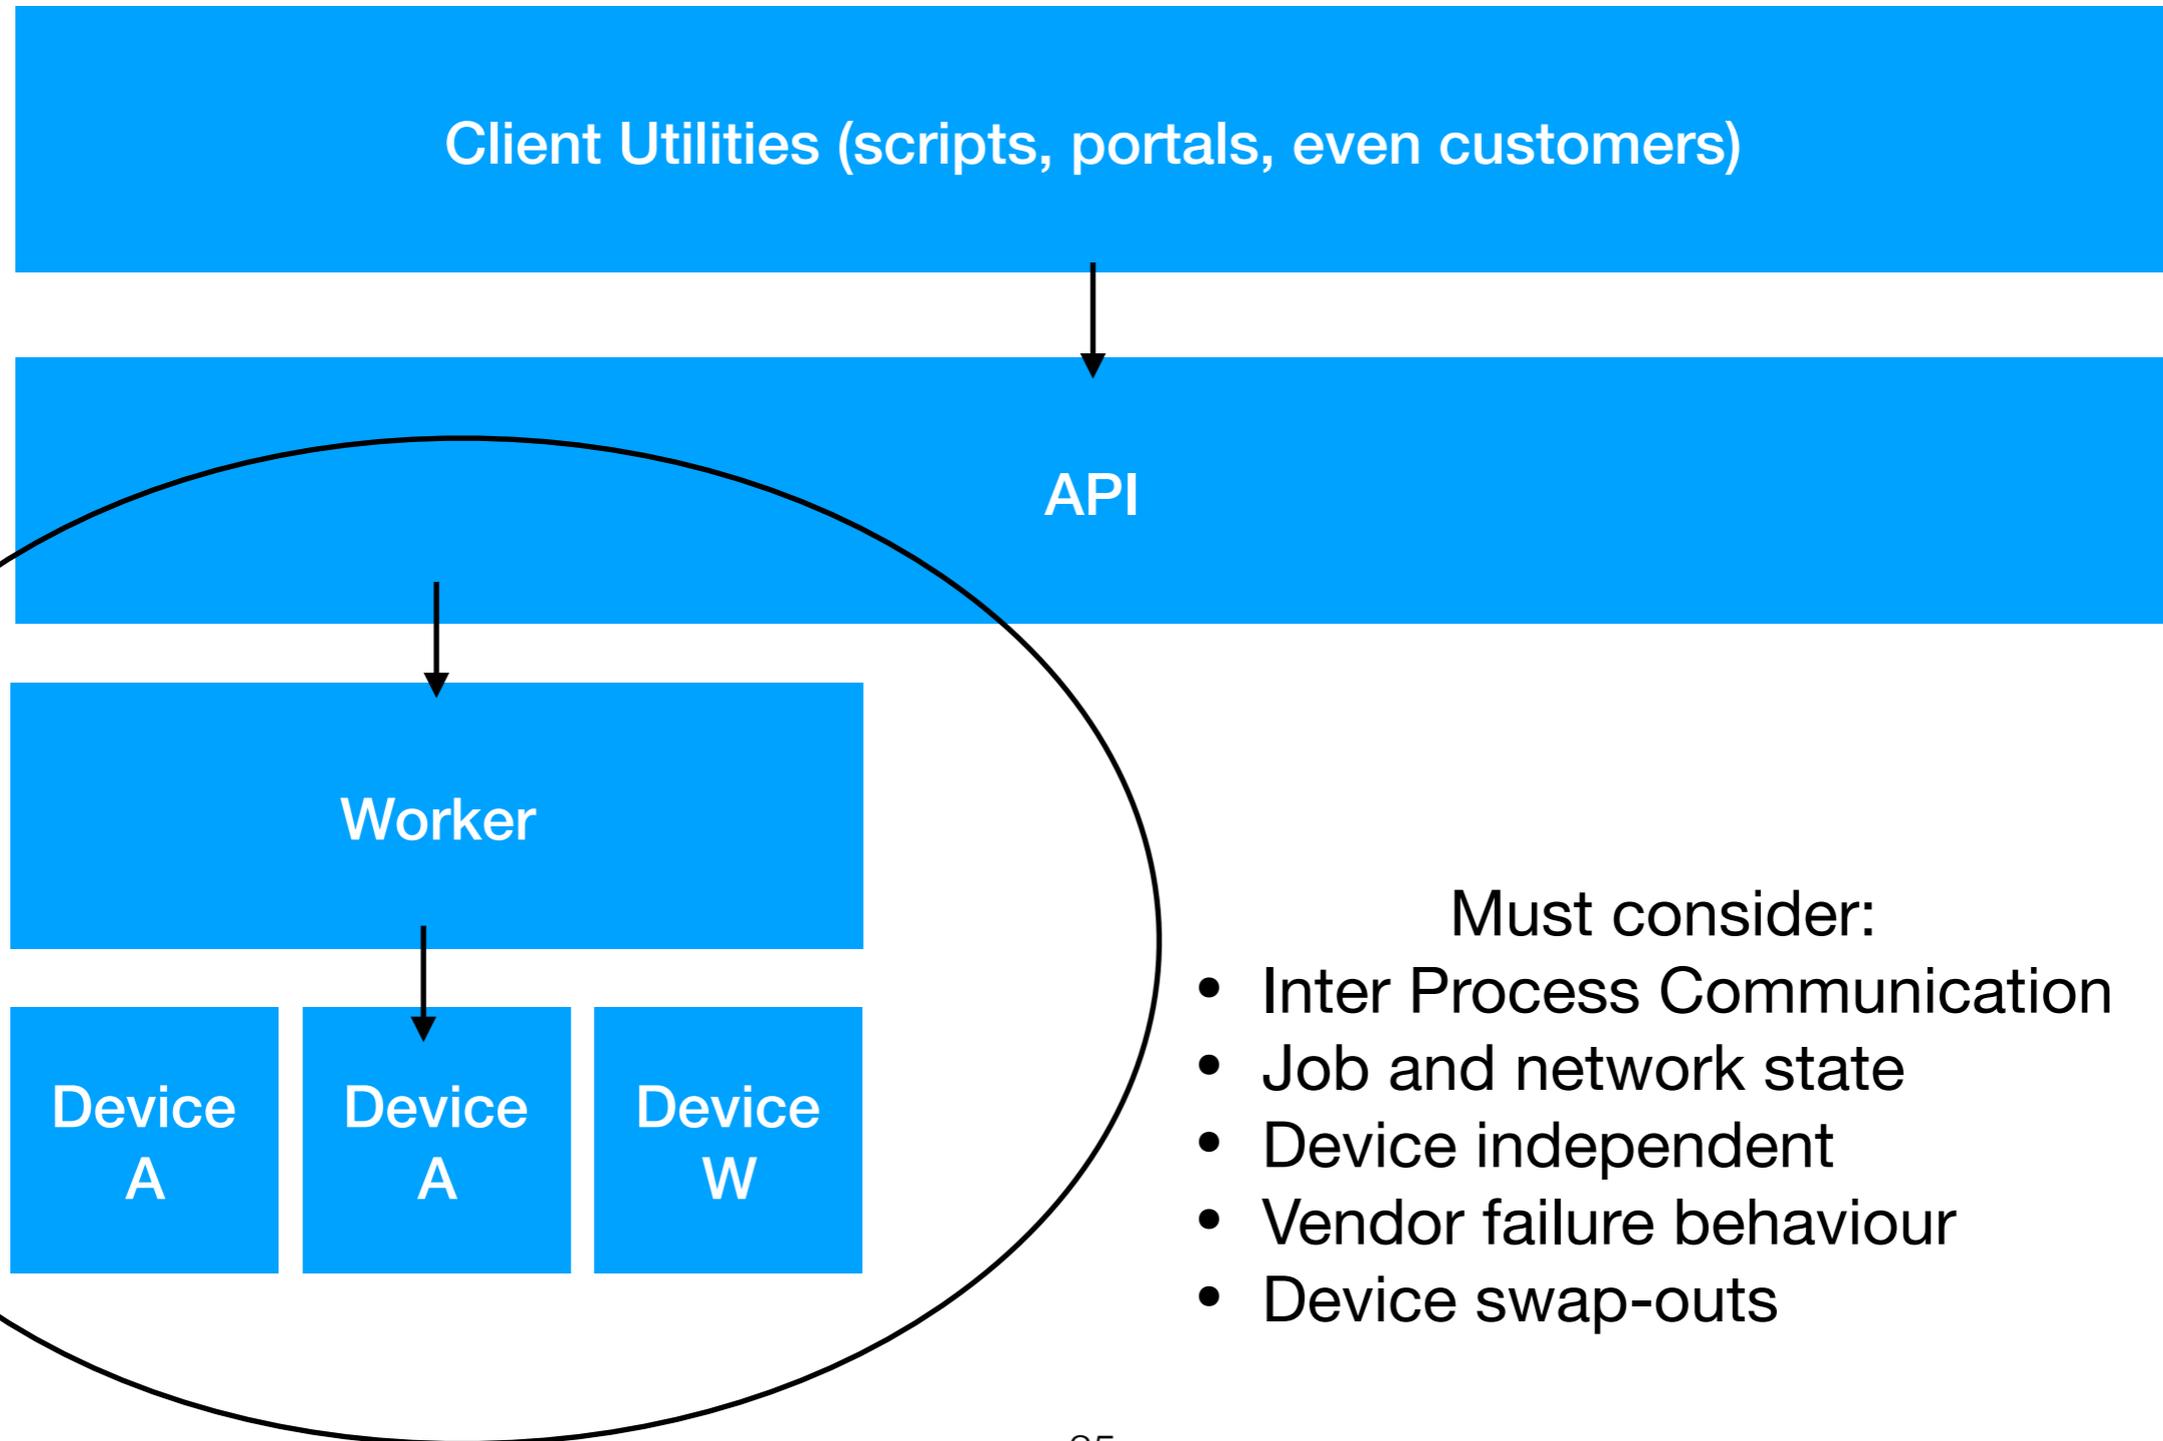
- API layer lightweight
 - Retrieve and update database records
 - Write in a familiar type-safe language (I chose Python)
- Automation layer lightweight
 - Essentially Ansible configuration files
 - Configuration “easier” than coding?

Business Logic



Ensure your API choices allow you to store, retrieve and process business logic as well as your technical products. Example: Asteroid Campaign logic.

Worker Architecture



Client Utilities (scripts, portals, even customers)

API

Worker

Device
A

Device
A

Device
W

Must consider:

- Inter Process Communication
- Job and network state
- Device independent
- Vendor failure behaviour
- Device swap-outs

Inter-Process communication

- Message Queue based?
 - e.g. RabbitMQ
 - 👍 Quite good support in major scripting languages
 - 👍 Fault tolerant, order matters, guaranteed delivery, HA
 - 👎 Extra software to support & Centralised
- Web Services
 - 👍 Same technology stack as central API
 - 👍 Inherently extensible
 - 👍 Decentralised
 - 👎 Extra software to write and more state to manage

Device Independence

- I chose to write a different worker per back end technology
- 🙅 A bit of copy/paste code, which is an anti-pattern
- 👍 No stress trying to treat different vendors generically
- 👍 NAPALM allows me to continue with Ansible
- 👍 Can swap out a switch/server architecture for sure

Device Swap-outs

- Using Ansible/NAPALM for switch configuration allows a process for rolling full configuration in event of device failure
- No need for specific software feature, an operational process is ok

Software Testing

```
[tatou:test_example andy$ ls
__pycache__  test_example.py
[tatou:test_example andy$ cat test_example.py
#!/usr/bin/env python

import unittest

class UnitTests(unittest.TestCase):
    def test_can_add_oneone(self):
        sum = 1 + 1
        self.assertTrue(sum == 2)
[tatou:test_example andy$ pytest
===== test session starts =====
platform darwin -- Python 2.7.10, pytest-3.0.5, py-1.4.32, pluggy-0.4.0
rootdir: /Users/andy/test_example, inifile:
plugins: cov-2.4.0
collected 1 items

test_example.py .

===== 1 passed in 0.01 seconds =====
tatou:test_example andy$
```

- Write the test first
- Red, Green, Refactor mantra

Integration vs Unit Testing

```
class UnitTests(LoginTests):
    def test_regular_user_can_view_own_contact(self):
        result_json = sputnik_get(self, "https://127.0.0.1:5001/contact/6")
        self.assertTrue(result_json["contact_firstname"] == 'Regular')

    def test_regular_user_cannot_view_others_contact(self):
        with self.assertRaises(AuthException):
            result_json = sputnik_get(self, "https://127.0.0.1:5001/contact/3")

    def test_cannot_load_administrator_systems_list_if_not_auth(self):
        with self.assertRaises(AuthException):
            result_json = sputnik_get(self, "https://127.0.0.1:5001/exchange/2/systems")
```

- If you are like me, you will prefer Integration tests
- Write lots, and remember to cover desired exceptions
- Run on your development instance after every change
- “Back to Zero” testing catches unexpected failures

The Joy of Errors

```
tatou:sputnik andy$ pytest
===== test session starts =====
platform darwin -- Python 2.7.10, pytest-3.0.5, py-1.4.32, pluggy-0.4.0
rootdir: /Users/andy/src/sputnik, inifile:
plugins: cov-2.4.0
collected 188 items

tests/test_api_asteroid.py .....
tests/test_api_whitelabel.py .....
tests/test_regular.py .....
tests/test_unit.py .....F.....

===== FAILURES =====
_____ UnitTests.test_can_load_organisations_quotes _____

self = <test_unit.UnitTests testMethod=test_can_load_organisations_quotes>

    def test_can_load_organisations_quotes(self):
        ql = QuoteList()
        testquote = ql.list_quotes_organisation(2)[0]
        self.assertTrue(testquote["quote_currency"] == "EUR")
> self.assertTrue(testquote["quote_lineitems"][0]["quote_lineitem_desc"] == "10Gbit Peering Port, Stellar Internet Exchange")
E       IndexError: tuple index out of range

tests/test_unit.py:391: IndexError
===== 1 failed, 187 passed in 5.82 seconds =====
```

Ubiquity of JSON for an ISP

- Especially in Peering!
 - PeeringDB
 - Euro-IX IXF-DB
 - Asteroid JSON

```
1 - {
2   "bridge_hostname": "ams-sw1",
3   "demarcation": {
4     "crossconnect_paying_organisation_id": null,
5     "crossconnect_ref": null,
6     "demarcation_id": 2,
7     "demarcation_label": "1/1/3 and 4",
8     "port_id": 2,
9     "rack_location": "Nikhef rack 108"
10  },
11  "exchange_id": 1,
12  "exchange_main_peering_vlan_id": 1,
13  "exchange_name": "Asteroid Amsterdam Internet Exchange",
14  "exchange_product_id": 2,
15  "exchange_product_name": "10Gbit Peering Port",
16  "exchange_tag": "ams",
17  "light_level": {
18    "latest_power_db": -2.8216313251307437
19  },
20  "organisation_id": 3,
21  "organisation_name": "Slashme BV",
22  "outstanding_work": [],
23  "peering_service": {
24    "mac_address": null,
25    "peering_dnsname": "asteroid-amsterdam.slashme.org",
26    "peering_ip6addr": "2001:7f8:b6::63fb:1",
27    "peering_ipaddr": "185.1.94.10",
28    "route_server_config": 1,
29    "service_id": 4,
30    "vlan_id": 1,
31    "vlan_number": "101"
32  },
33  "permit_config_change": true,
34  "port_bundle_group": null,
35  "port_bundle_type": "none",
36  "port_customer_status": "Live",
37  "port_description": "CUST::Slashme BV",
38  "port_enabled": 1,
39  "port_first_link_enable_time": "6th-February-2018 17:29:53",
40  "port_id": 2,
41  "port_identifier": "Ethernet1/2",
42  "port_original_lineitem_id": 1,
43  "port_original_quote_id": 1,
44  "port_repeatbill_id": 1,
45  "port_repeatbill_mrc": 0,
46  "port_speed_name": "10000",
47  "port_supervisor_lock": 0,
48  "port_switch_admin_enable": 1,
49  "port_switch_last_flapped_time": "13th-January-2018 00:01:49",
50  "port_switch_link_enable": 1,
51  "port_switch_observed_time": "2nd-March-2018 14:51:06",
52  "port_type": "physical",
53  "stats": {
54    "backtime": "1w",
55    "backtime_days": "7",
56    "bps_in": [
57      [
58        "2018-02-23T14:30:00Z",
59        1198.0533333333333,
60        1519396200
61      ],
62      [
63        "2018-02-23T15:00:00Z"
```

Summary

- Single source of truth under the control of all departments
- Which is used to configure services and network
- Accessible to all departments
 - Customer self service
 - Provision from quote
 - “Information in one place and tool”
 - Account Managers can do troubleshooting

Any Questions?



Andy Davidson <andy@asteroidhq.com>
www.asteroidhq.com