

# Automated network configuration

Big networks can be correct

Michael Shields, Google Engineering  
RIPE 57, October 2008

policy generation for the network

audit for correctness and policy adherence

ensure completeness of your architectural standards

modeling

y equipment

vision power and space

ck and stack

erconnect

**configure the device based on your documentation**

**have another engineer check your work**

profit!

**We've done it a  
thousand times.**

**What could go  
wrong?**



---

```
interface ethernet [x/y]
  ip address [address] [netmask]
  vrrp 1 priority [120, 100]
  vrrp 1 authentication cisco
  vrrp 1 timers advertise 3
  vrrp 1 timers learn
  vrrp 1 ip [address]
  no shutdown
```

```
interface ethernet 1/0
  ip address 10.1.0.2 255.255.255.0
  vrrp 1 priority 120
  vrrp 1 authentication cisco
  vrrp 1 timers advertise 3
  vrrp 1 timers learn
  vrrp 1 ip 10.1.0.10
  no shutdown
```

```
interface ethernet 1/0
  address 10.1.0.2 255.255.255.0
  vrrp 1 priority 100
  vrrp 1 authentication cisco
  vrrp 1 timers advertise 3
  vrrp 1 timers learn
  vrrp 1 ip 10.1.0.10
  shutdown
```

```
interface ethernet 1/0
  ip address 10.1.0.2 255.255.255.0
  vrrp 1 priority 120
  vrrp 1 authentication cisco
  vrrp 1 timers advertise 3
  vrrp 1 timers learn
  vrrp 1 ip 10.1.0.10
  no shutdown
```

```
interface ethernet 1/0
  address 10.1.0.2 255.255.255.0
  vrrp 1 priority 100
  vrrp 1 authentication cisco
  vrrp 1 timers advertise 3
  vrrp 1 timers learn
  vrrp 1 ip 10.1.0.10
  shutdown
```

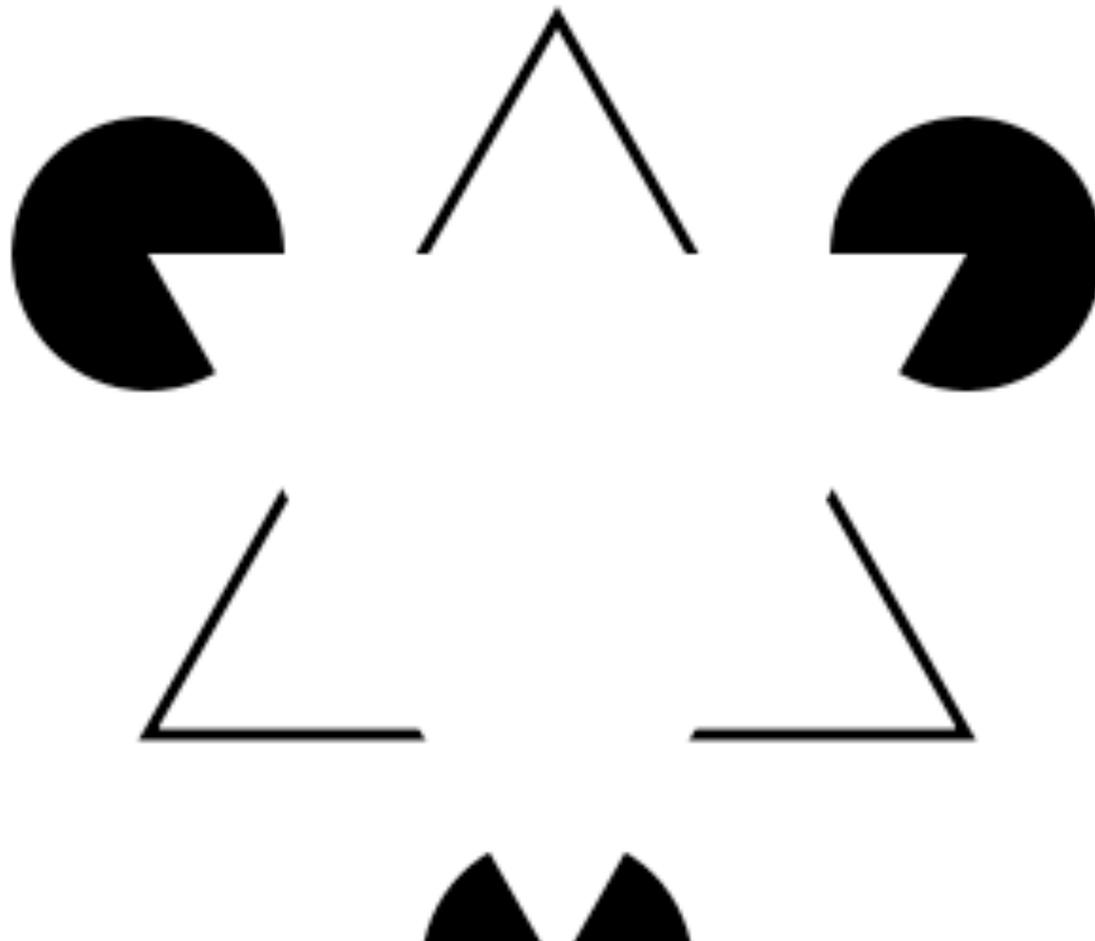
```
interface ethernet 1/0
  ip address 10.1.0.2 255.255.255.255
  vrrp 1 priority 120
  vrrp 1 authentication cisco
  vrrp 1 timers advertise 3
  vrrp 1 timers learn
  vrrp 1 ip 10.1.0.10
  no shutdown
```

people will cut and paste. Errors propagate.

humans expect to see patterns.

Our brains see things that aren't there.

errors happen, and you won't find all of them just by looking.



## Most common types of errors:

- Missing or incorrect security ACLs
- Incomplete BGP meshes (mysterious blackholing)
- Incomplete MPLS mesh
- Incomplete or incorrect QoS configuration
- IP address confusion

## Typical response: Add more procedures

- End up with a mass of procedures and policies that look the same
- Can't ever keep the corpus of documentation self-consistent
- Can't ever keep the network consistent with documentation
- You'll go blind trying. Need to create abstractions so we can think big.

**the more actions you take, the more mistakes you'll make.**

**if it's not automated, it will not scale.**

[Corollary: if your network can be managed by hand, it is small.]

**correct networks scale better.** This is a competitive advantage

But you already have a network. It's up and running.  
You're good at your job, so your network is pretty good.

It's good enough. But how do you get to better?

What should you do instead?

and how do you get there?

**Configurations are templates with variable substitution**

enforce policy by tools, not only by documentation

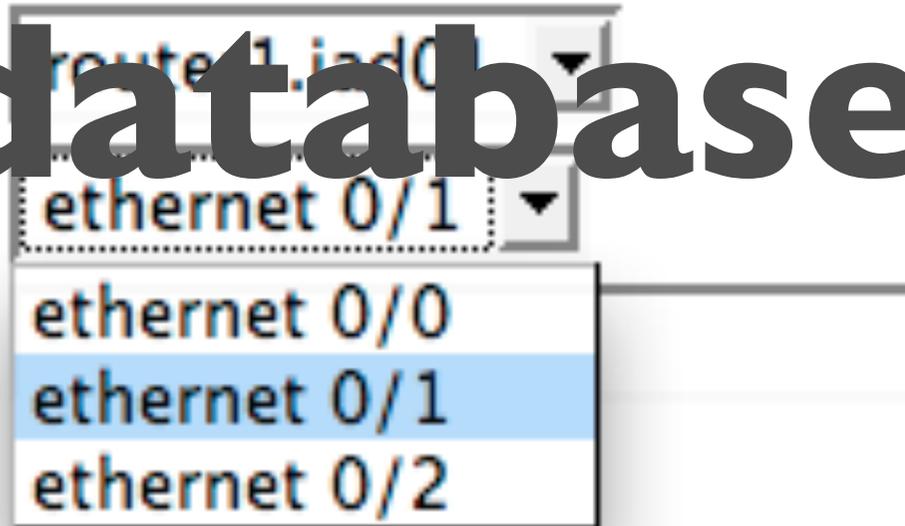
forces you to find your exceptions

**Don't touch the  
router until it's in  
the database**

**Router:**

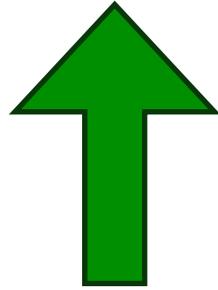
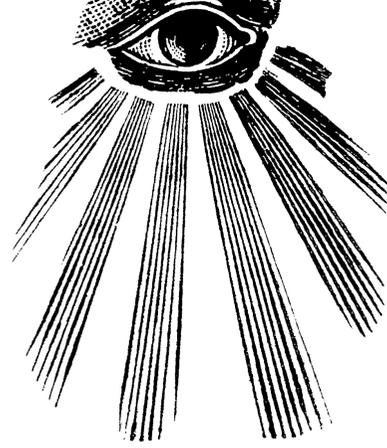
**Interface:**

**IP address:**



The image shows a screenshot of a network configuration interface. On the left, there are three labels: "Router:", "Interface:", and "IP address:". To the right, there is a dropdown menu for selecting an interface. The dropdown menu is currently open, showing a list of interface options: "ethernet 0/0", "ethernet 0/1", and "ethernet 0/2". The "ethernet 0/1" option is highlighted in blue. The text "route 1.iad0" is visible above the dropdown menu, and "ethernet 0/1" is visible in the dropdown menu's header area.

Subnet	Interface	Customer
10.1.0.2/24	ethernet 0/1	6829 — E. Blofeld, Inc.
10.1.0.3/24	ethernet 0/2	3189 — Disco Volante
10.1.0.4/23	ethernet 0/3	17942 — Thanet Alloy



# What's up in the

# attic?



```
1 interface ethernet 1/0
2 ip address 10.1.0.2 255.255.255.0
3 vrrp 1 priority 100
4 vrrp 1 authentication cisco
5 vrrp 1 timers advertise 3
6 vrrp 1 timers learn
7 vrrp 1 ip 10.1.0.10
8 no shutdown
```

```
1 interface ethernet 1/0
2 ip address 10.1.0.1 255.255.255.0
3 vrrp 1 priority 100
4 vrrp 1 authentication cisco
5 vrrp 1 timers advertise 3
6 vrrp 1 timers learn
7 vrrp 1 ip 10.1.0.10
8 no shutdown
```

Why are things wrong?  
(in order of how easy to fix)

1. Bugs in code for initial population of database
2. Actual configuration errors
3. Deliberately unusual configuration

**s not enough to generate your configs once.**

ou should be able to recheck the actual network state against the  
generated network state at any time.

# Actual Google audit

```
protocols {  
  ...  
  ospf {  
    ...  
    area 0.0.0.0 {  
      {% for interface in dev.physicalinterface_set.all %}  
      {% for unit in interface.logicalinterface_set.all %}  
      {% if unit.ospf_metric %}  
      interface {{ unit.name }} {  
        metric {{ unit.ospf_metric }};  
      }  
      {% endif %}  
      {% endfor %}  
      {% endfor %}  
    }  
    ...  
  }  
}
```

Router	Type	Loopback	IS-IS NET
router1.iad01	Cisco AGS+	192.0.2.38	49.0001.0000.00 00.000a.00
router2.iad01	Cisco AGS+	192.0.2.39	49.0001.0000.00 00.000b.00
router1.lhr07	Cisco 4500M	192.0.2.207	49.0001.0000.00 00.000c.00

```
interface serial 1
 ip address 10.0.0.2 255.0.0.0
 ip ospf network point-to-multipoint
```

```
encapsulation frame-relay
 frame-relay map ip 10.0.0.1 201 broadcast
 frame-relay map ip 10.0.0.3 202 broadcast
 frame-relay map ip 10.0.0.4 203 broadcast
```

```
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
```

```
1 interface serial 1
2 ip address 10.0.0.2 255.0.0.0
3 ip ospf network point-to-multipoint
4 ip router isis
5 isis metric 503 level-2
6 isis password ISISPASSWORD level-2
7 encapsulation frame-relay
8 frame-relay map ip 10.0.0.1 201 broad
9 frame-relay map ip 10.0.0.3 202 broad
10 frame-relay map ip 10.0.0.4 203 broad
11 !
12 router ospf 1
13 network 10.0.0.0 0.0.0.255 area 0
14 !
15 router isis
16 passive-interface serial 1
17 maximum-paths 6
18 net 49.0001.0000.0000.000a.00
19 is-type level-2-only
20 metric-style wide
21 ...
```

(data) → old configuration

(data) → new configuration

**You can work at the level of the forest, not the trees**

# Allow for one-otts with jails



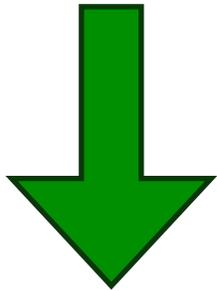
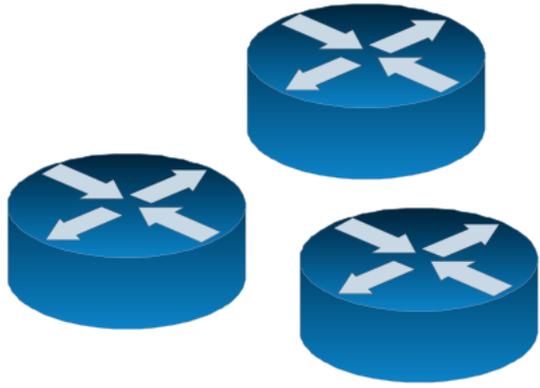
Ad hoc  
isolation jail

Configuration collector — RANCID is a good start

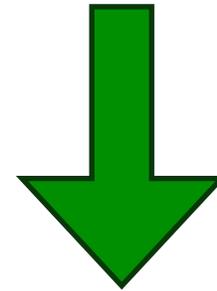
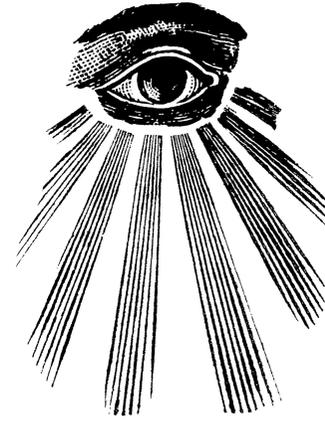
Database

Configlet generator

Comparison engine



configs  
(as-built)



configs  
(generated)



**enter all routine operations around an audit.**

**fix one class of problems at a time, networkwide.**

BGP mesh is a good place to start.

**continuously compare** your generated configs  
against actual configs, and get diffs to zero.

**use jails** to isolate nonstandardness.

Small networks don't need this.

Big ones do.

The transition is the hard part.

You can only get there incrementally.

**Out of the crooked timber of humanity,  
no straight thing was ever made.” —Kant**

# Questions?

Michael Shields

[mshields@google.com](mailto:mshields@google.com)

P.S. We're still hiring.