Programatic Networks - Autogen

"The history of your network"

or

"How we all got into this mess"

Vijay Gill, Michael Shields, Google Engineering

Agenda

What is Automatic Configuration Generation

Typical Operational Issues

Policy Enforcement

Case Study

Automatic Configuration Generation

Policy generation for the network

Audit for correctness and policy adherence

Ensure completeness of your architectural standards

Modeling

Typical Turn Up

Buy equipment

Provision power and space

Rack and stack

Interconnect

Activate

- Insert into protocols, meshes, monitoring
- Audit

Handoff

Forwarding

Policy And Scale

Implementing correct policy at scale is hard

If it's not automated, it will not scale

(most people will never need to scale)

Race to the bottom – we are in a commoditized business

OSS/NMS is a competitive advantage

“If your policy is in a wiki or a document, it doesn't exist” –Dan Cohn

Typical Errors

Do not adhere to policy

- Missing or Incorrect Security ACLs
- Incomplete BGP Meshes (mysterious blackholing)
- Incomplete MPLS Mesh
- Incomplete or Incorrect QoS configuration

Typical response – Add more procedures

- End up with a mass of MOPs and policies that look the same

Why

Muscle Memory

- People will cut/paste
- “know the correct configuration”

See what is not there

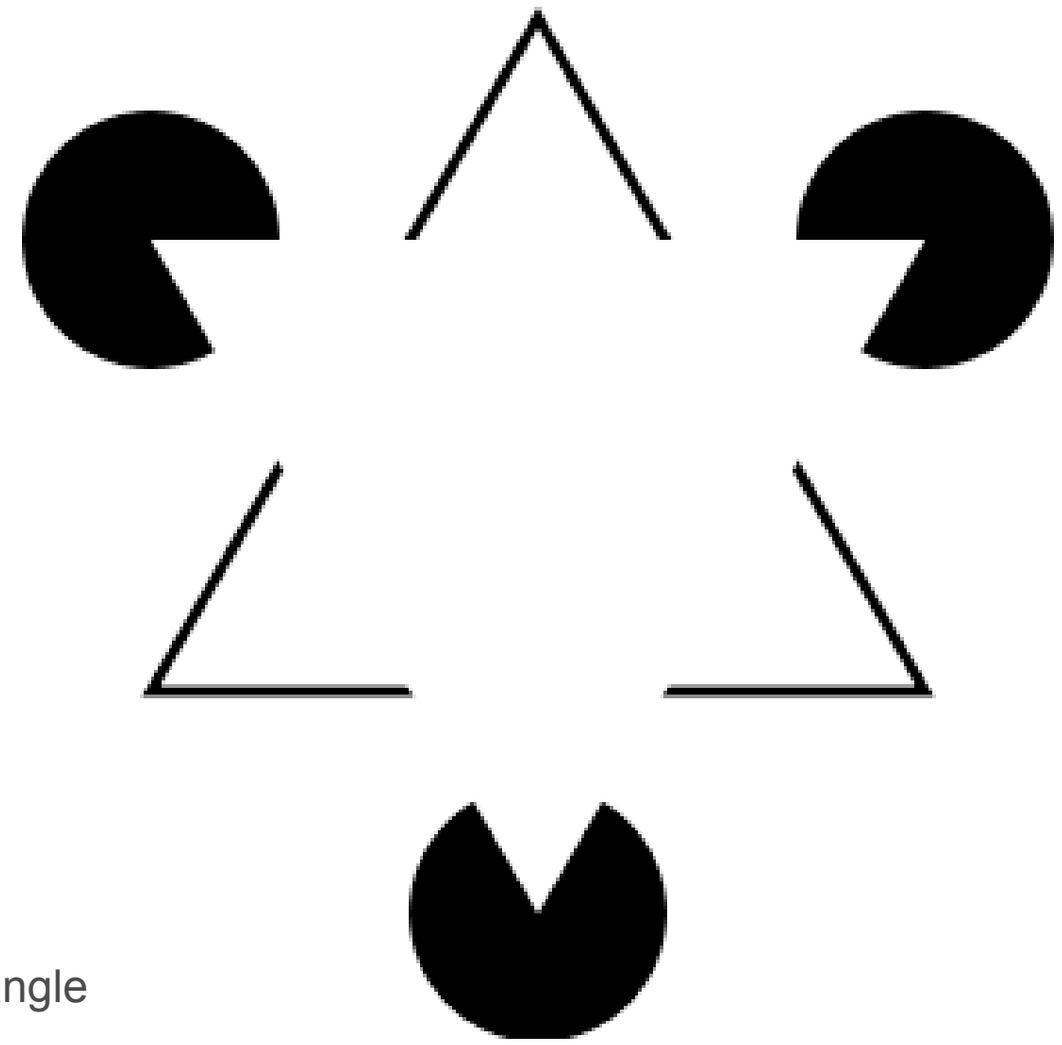

Kanizsa triangle

Transformation

Don’t get trapped by your special corner cases

- Reduces flexibility

Respond quickly to new service rollouts

Respond to security holes/bugs

Audits/Compliance/Planning

Consistent and quick configuration changes important

Canonical example - changing IGPs

- F(data) → old configuration
- F′(data) → new configuration

Data Dictionaries

Extensible Entity-Attribute-Relationship (EAR) Implementation

Core Data Model

- Set defined in order
- Share the definition of a record or field
- Allow tools to asses the impact of changing metadata on systems that use it

Dependencies

- Impact of change analysis possible
- Change the IP addr field from 4 bytes to 16

Relationships point to a specific version in the version stack

- Allow future state of the network to be described
- Automated tools can calculate the configuration changes required to morph the network

Policy Enforcement

Configurations are templates with variable substitution

Enforce Policy by tools, not by documentation

- Not that documentation shouldn't exist

Tools don't get tired or skip steps (bug free ones)

Encode Tribal Knowledge in a backed up format

Google™

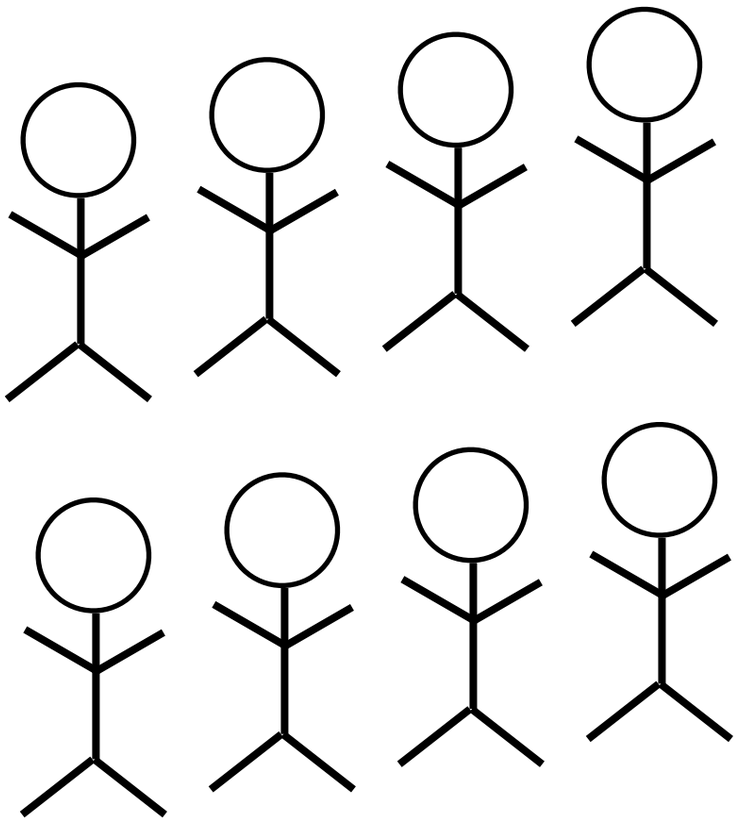

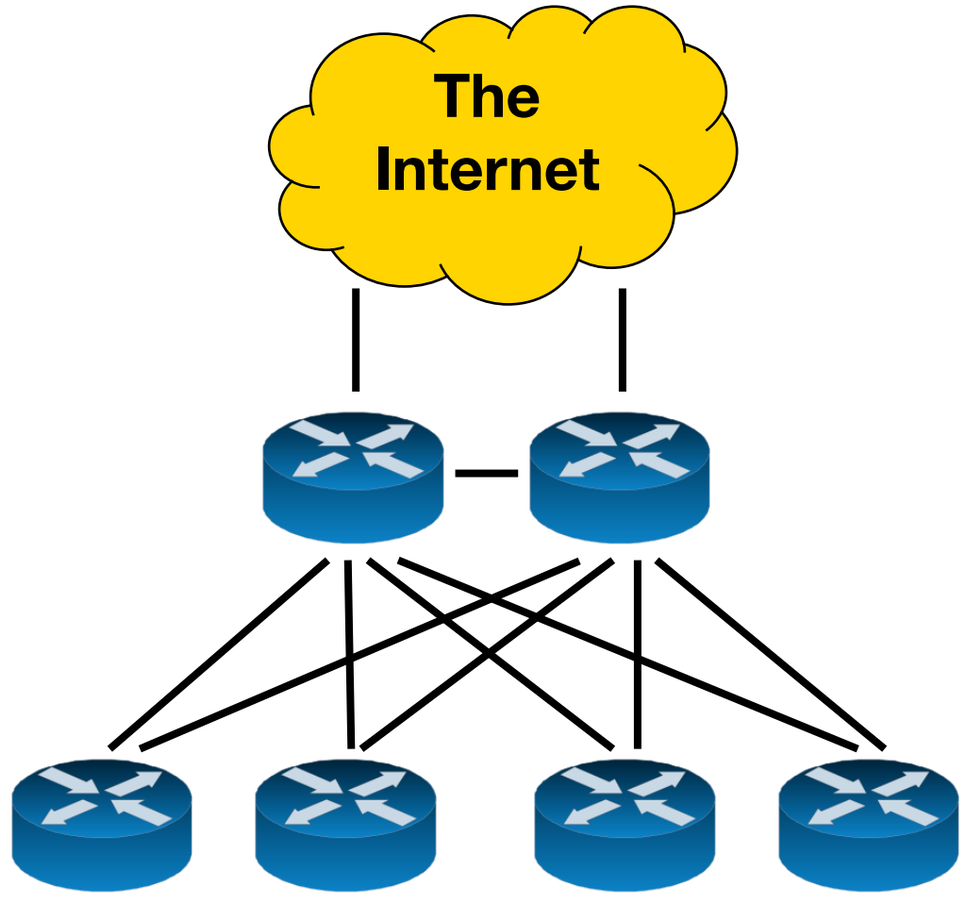

Google™

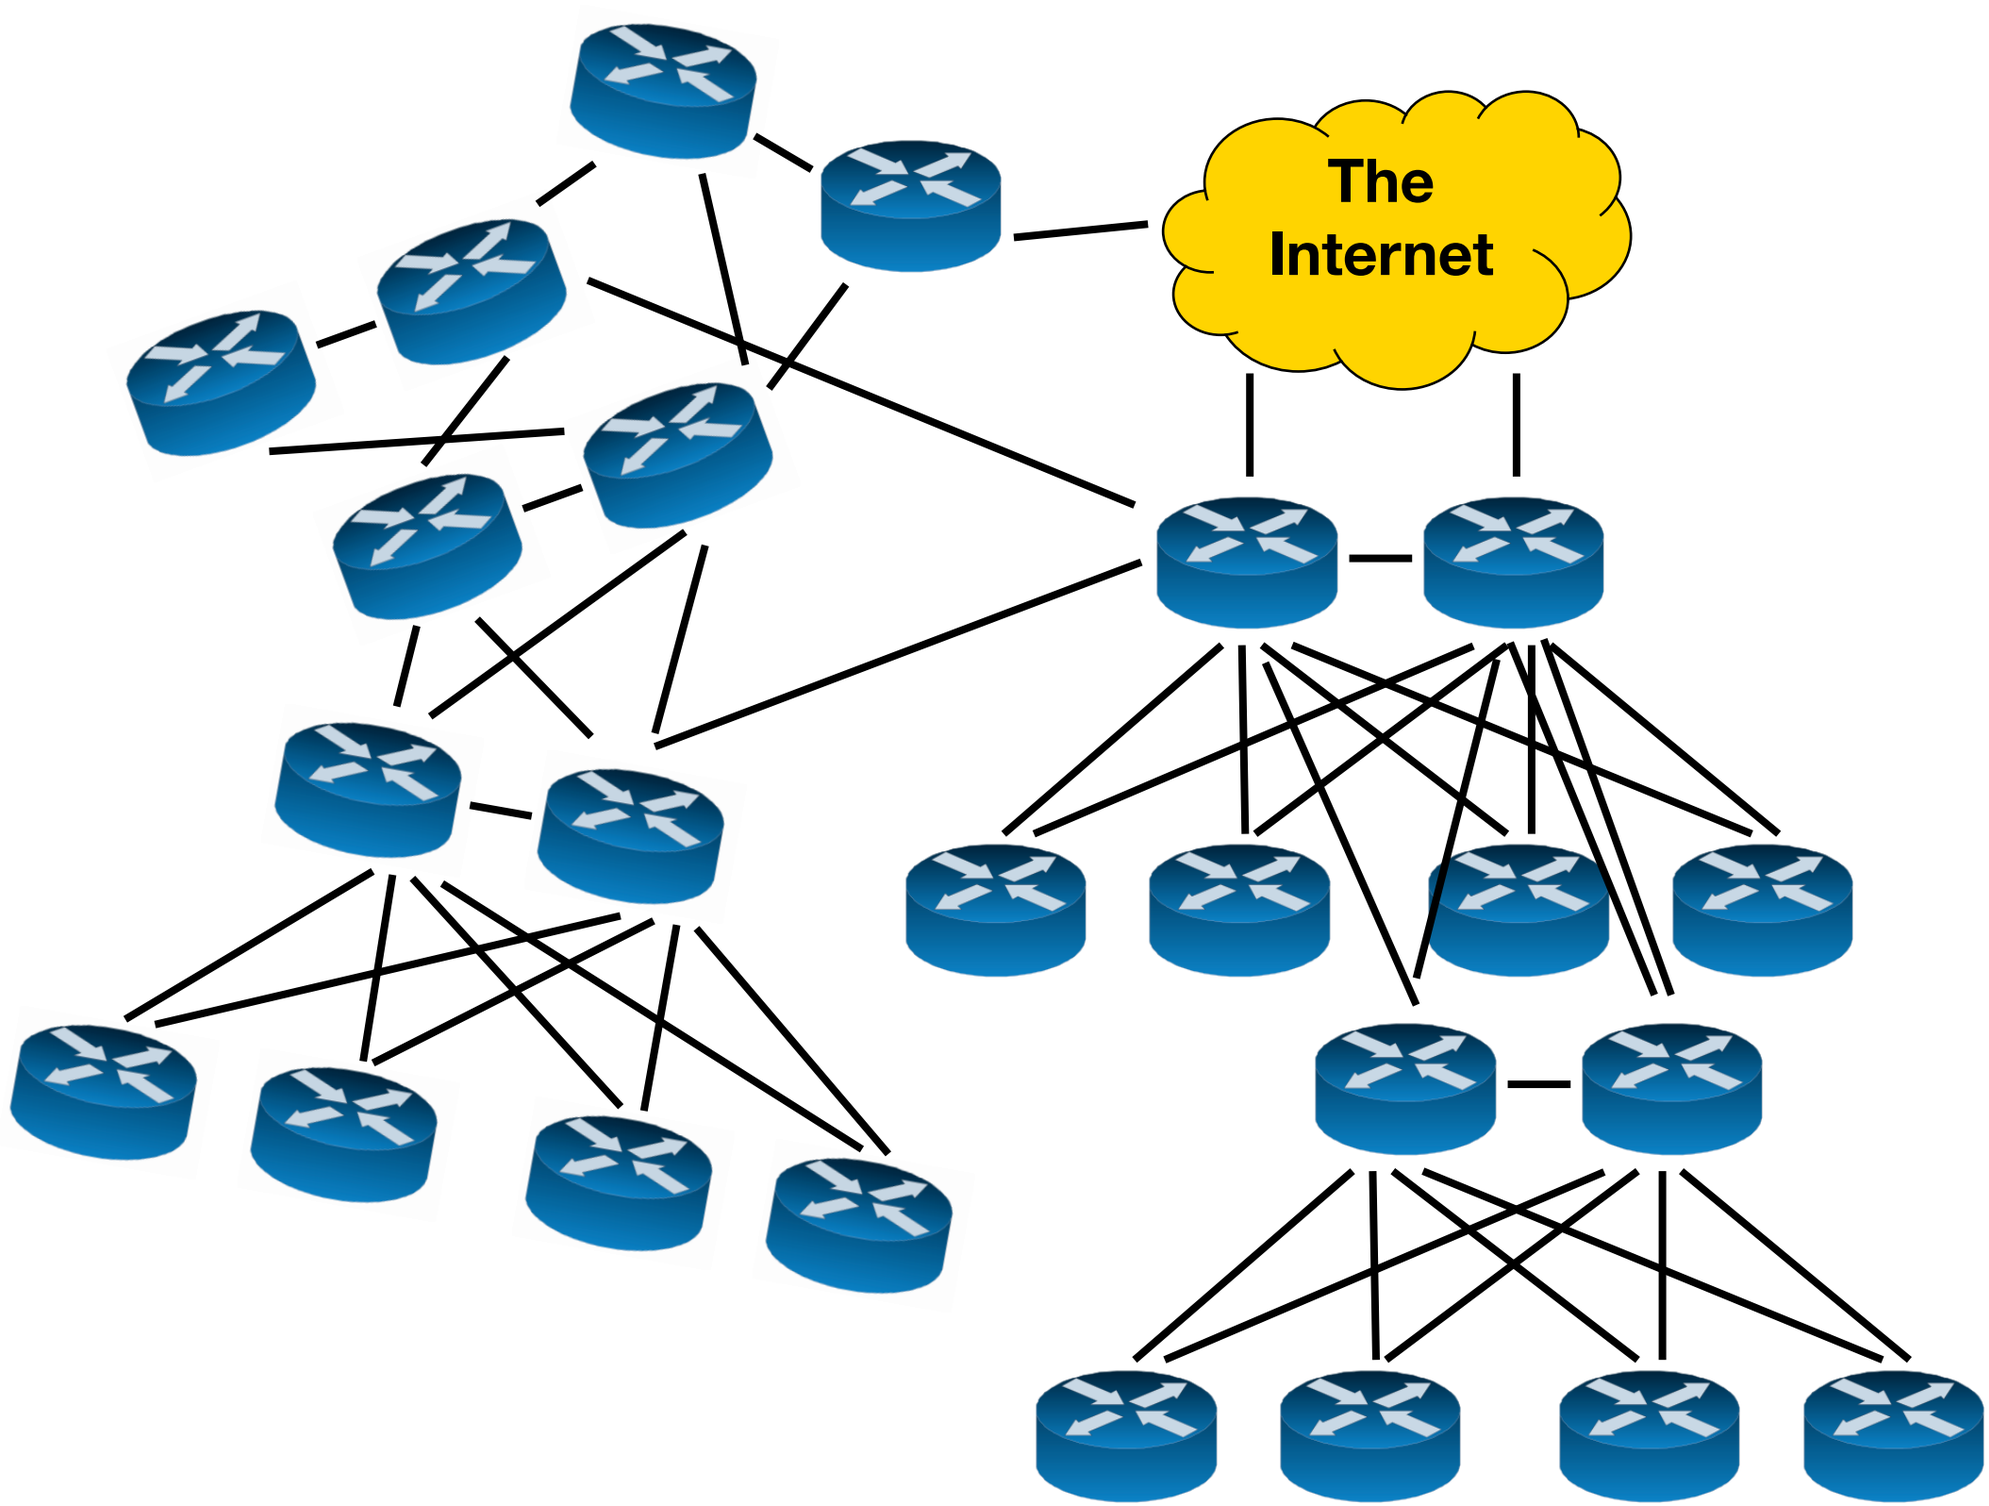

Google™

CC licensed: http://flickr.com/photos/toddneville/2231683148/

Google™

1997–2000

1998–2002

2002–2006

2006–

2007–

```
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
 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
 ip 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
 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
 ip 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
 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
 ip 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
 no shutdown
```

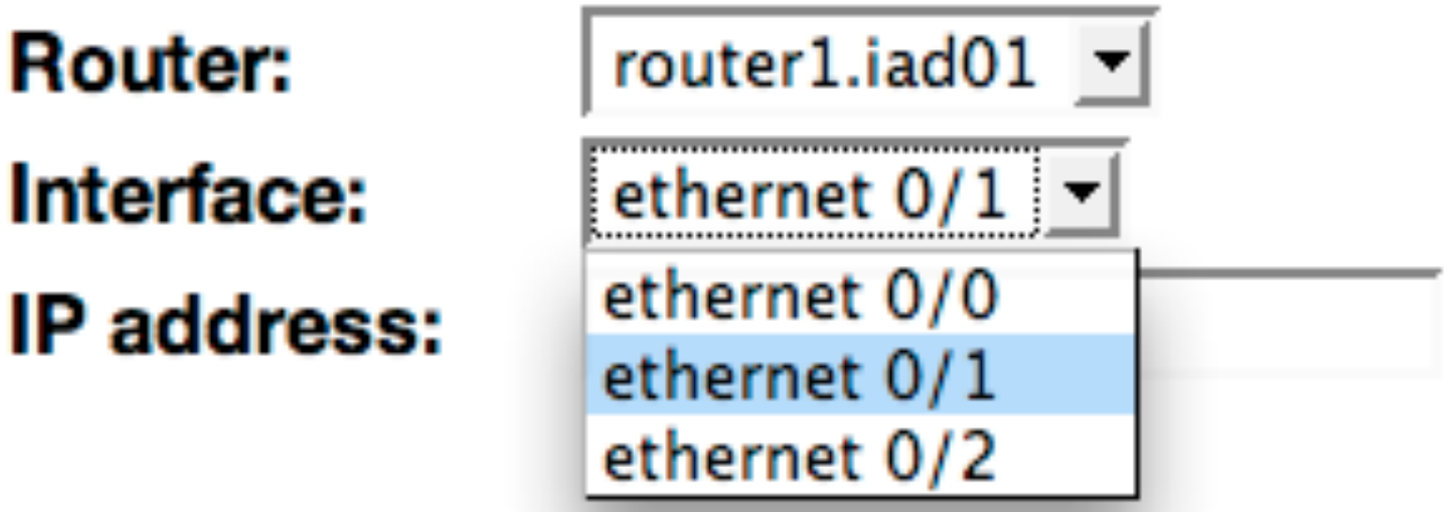
Router:
router1.iad01
Interface:
ethernet 0/1
IP address:
ethernet 0/0
ethernet 0/1
ethernet 0/2

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

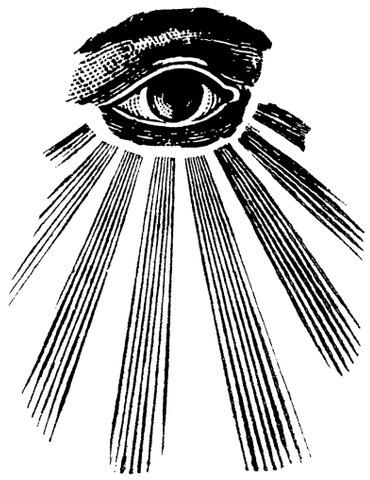

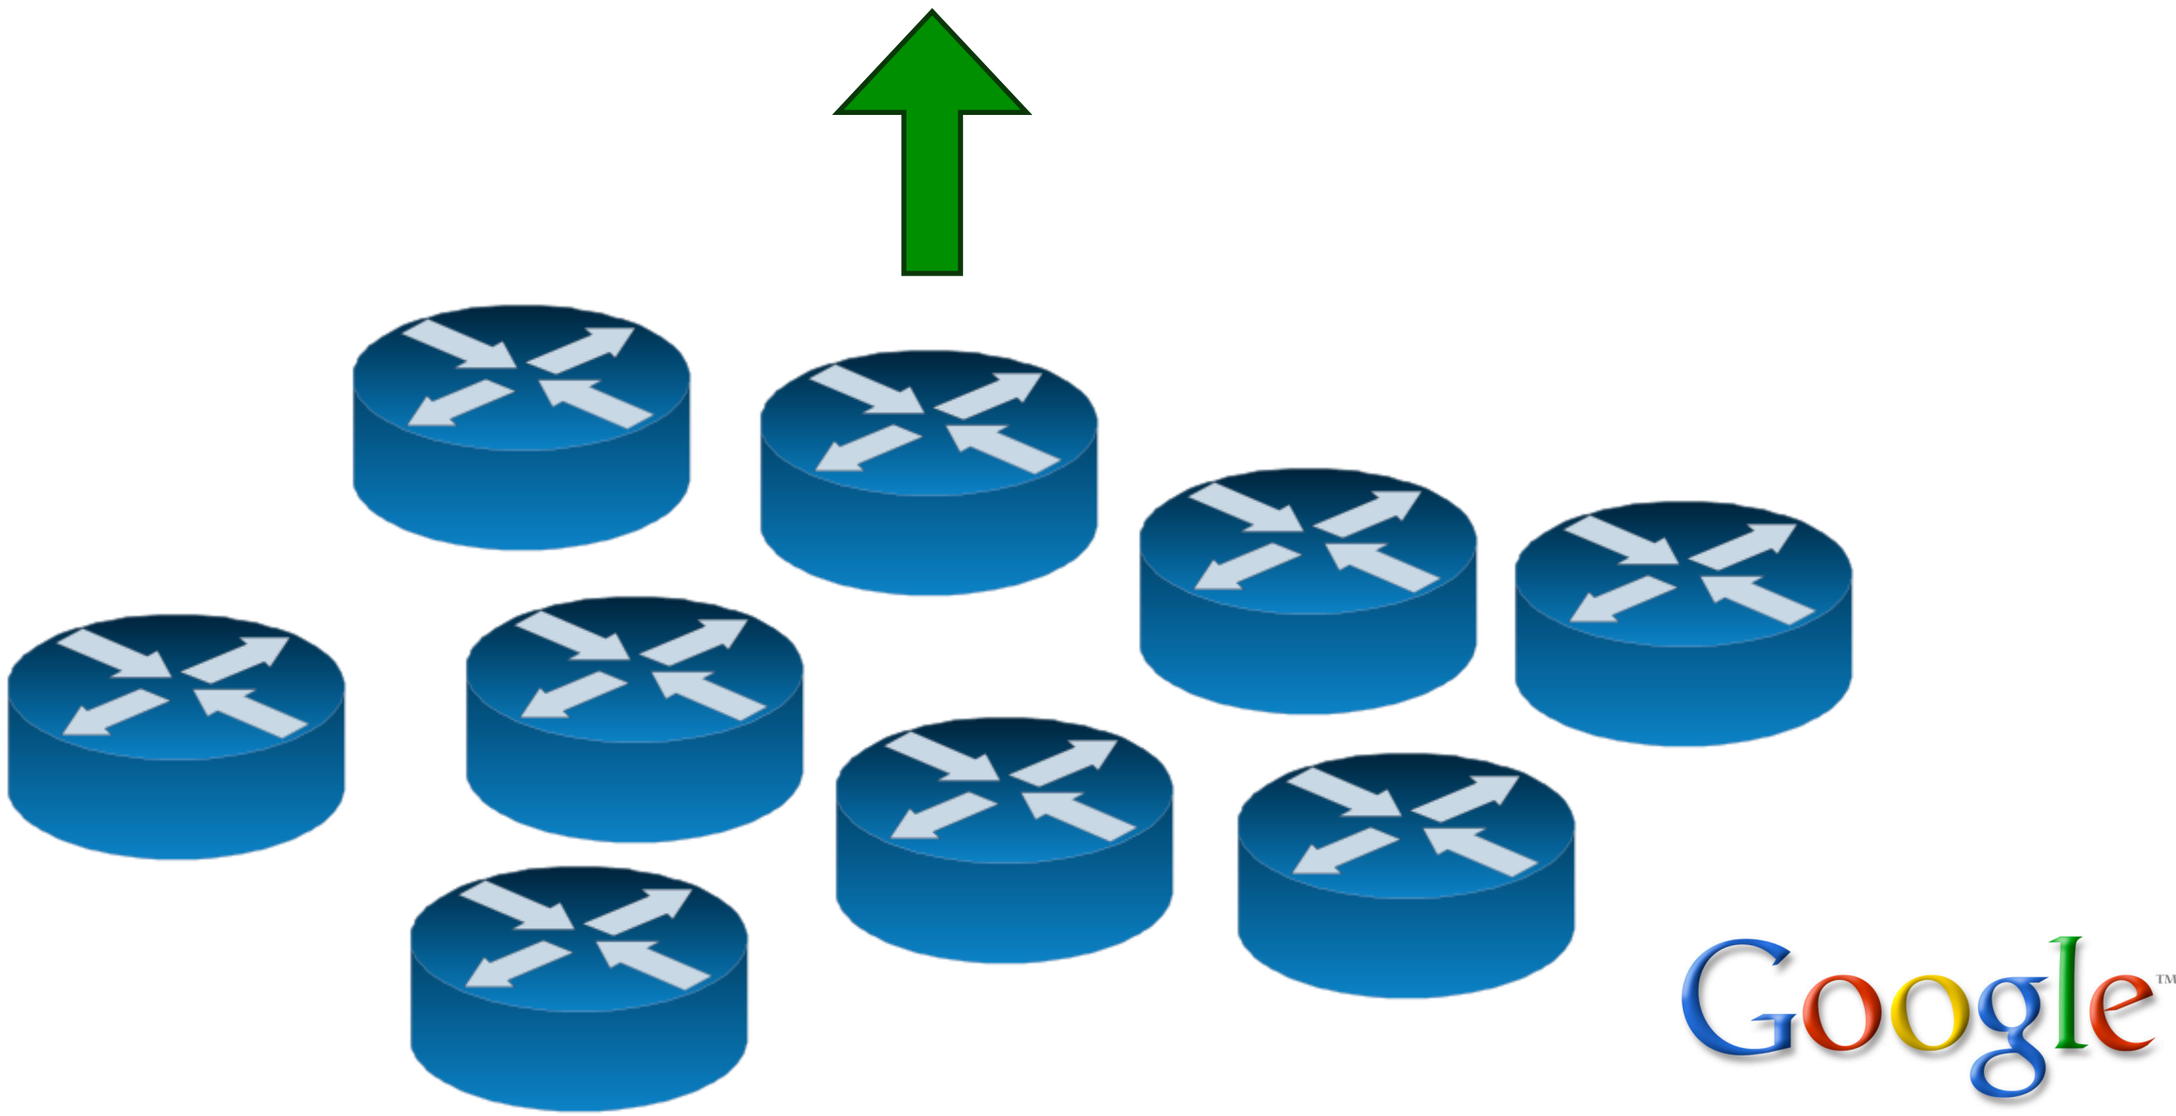

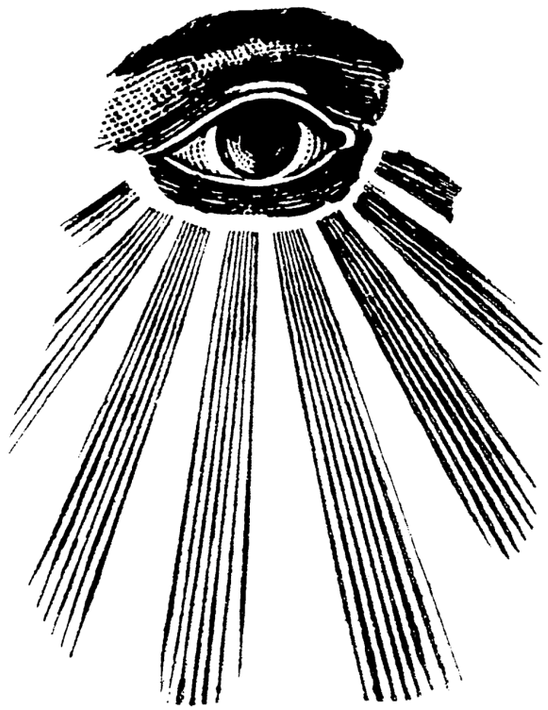

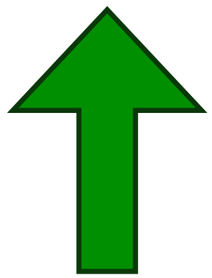

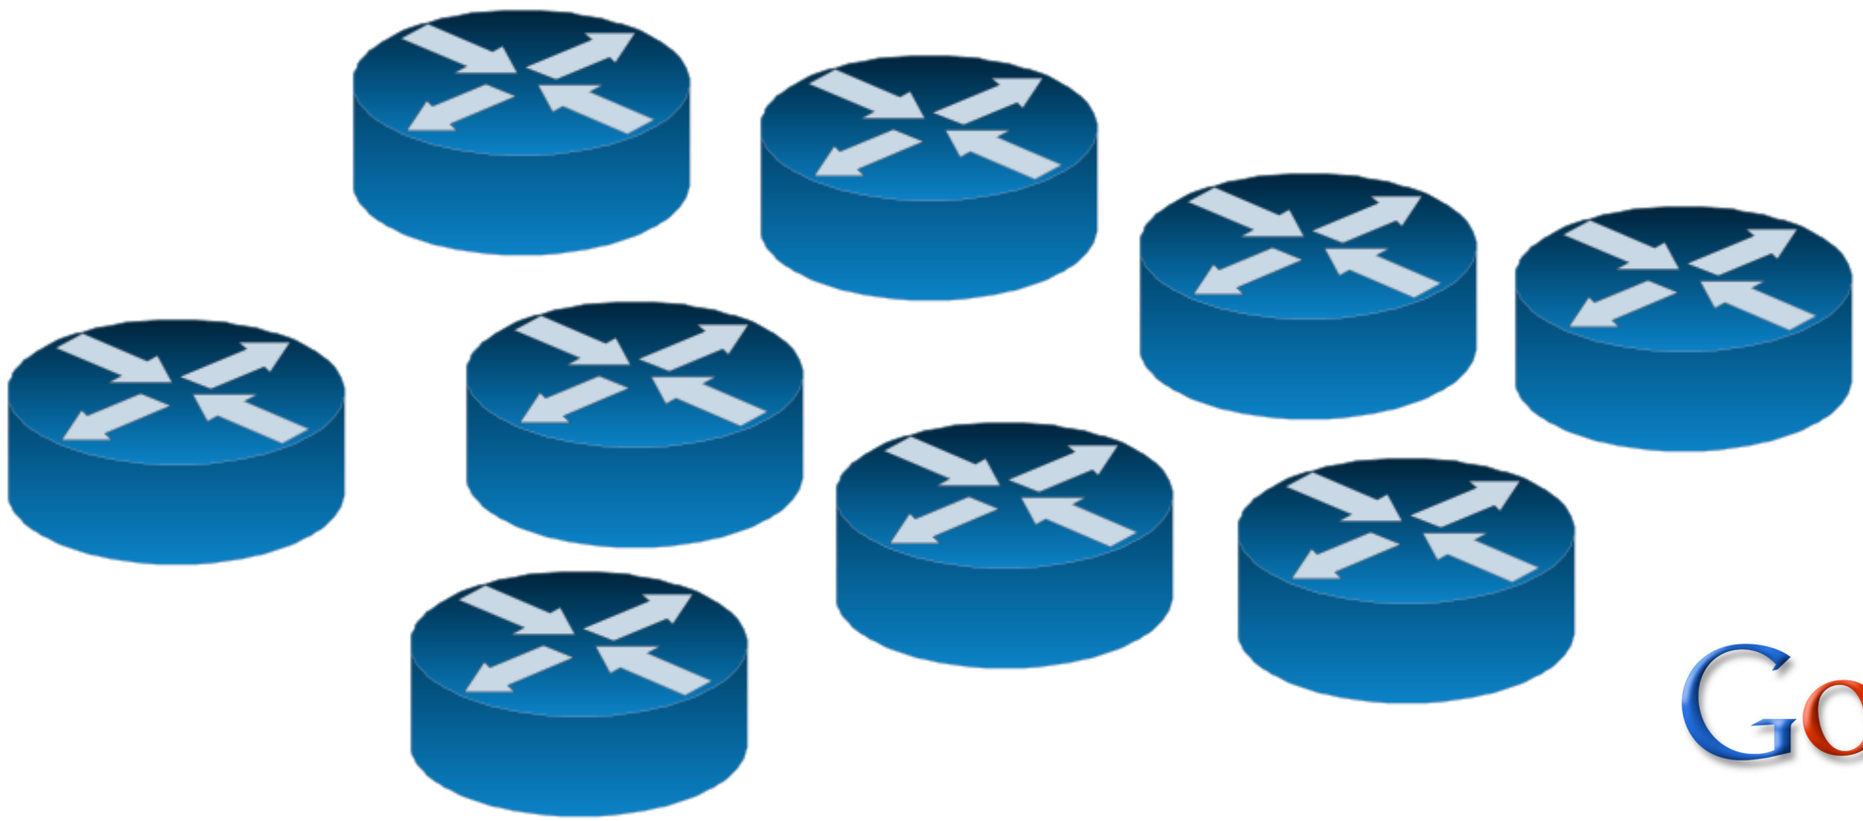

Google™

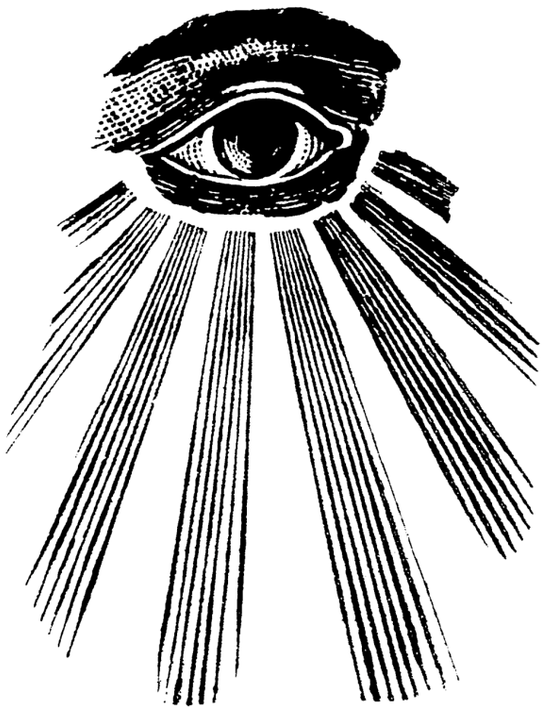

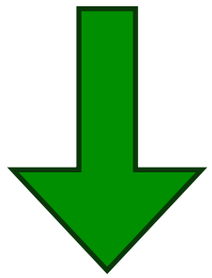

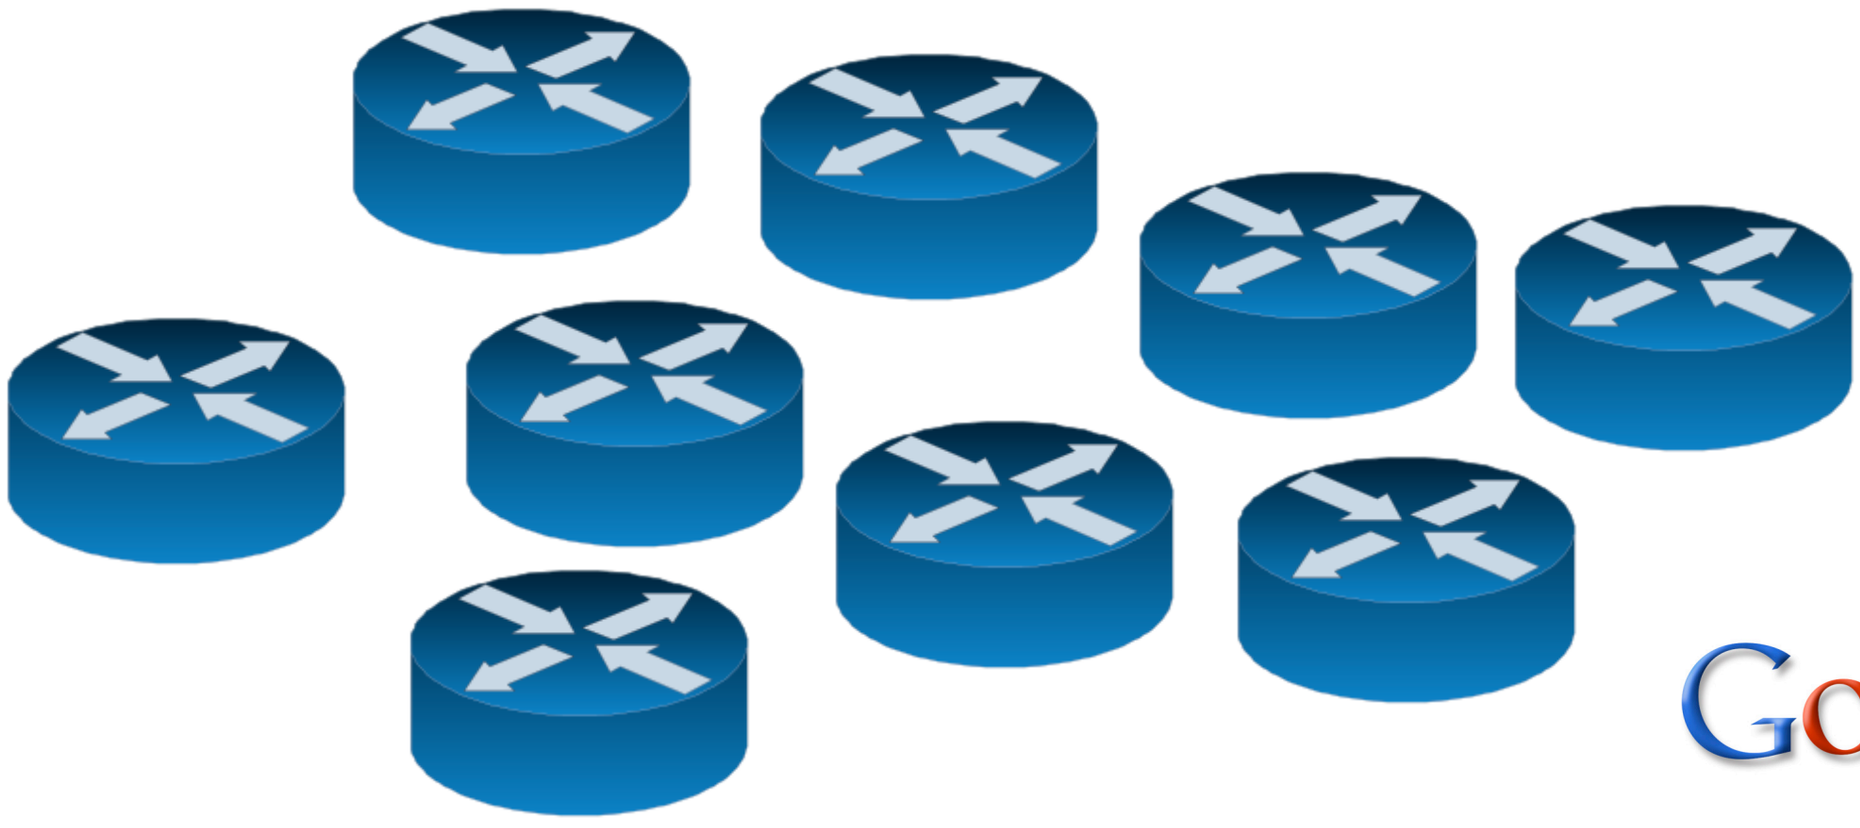

Google™

```
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
```

Error Reasons

1 - Bugs in code for initial population of DB

2 - Actual Configuration Errors

3 - Valid deviation for business reasons

Router	Type	Loopback
router1.iad01	Cisco AGS+	192.0.2.38
router2.iad01	Cisco AGS+	192.0.2.39
router1.lhr07	Cisco 4500M	192.0.2.207

Router	Type	Loopback	IS-IS NET
router1.iad01	Cisco AGS+	192.0.2.38	49.0001.0000.0000.000a.00
router2.iad01	Cisco AGS+	192.0.2.39	49.0001.0000.0000.000b.00
router1.lhr07	Cisco 4500M	192.0.2.207	49.0001.0000.0000.000c.00

```
 1 interface serial 1                                    1 interface serial 1
 2  ip address 10.0.0.2 255.0.0.0                        2  ip address 10.0.0.2 255.0.0.0
 3  ip ospf network point-to-multipoint                   3  ip ospf network point-to-multipoint
                                                          4  ip router isis
                                                          5  isis metric 503 level-2
                                                          6  isis password ISISPASSWORD level-2
 4  encapsulation frame-relay                             7  encapsulation frame-relay
 5  frame-relay map ip 10.0.0.1 201 broadcast             8  frame-relay map ip 10.0.0.1 201 broadcast
 6  frame-relay map ip 10.0.0.3 202 broadcast             9  frame-relay map ip 10.0.0.3 202 broadcast
 7  frame-relay map ip 10.0.0.4 203 broadcast            10  frame-relay map ip 10.0.0.4 203 broadcast
 8 !                                                     11 !
 9 router ospf 1                                         12 router ospf 1
10  network 10.0.0.0 0.0.0.255 area 0                    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 ...
```

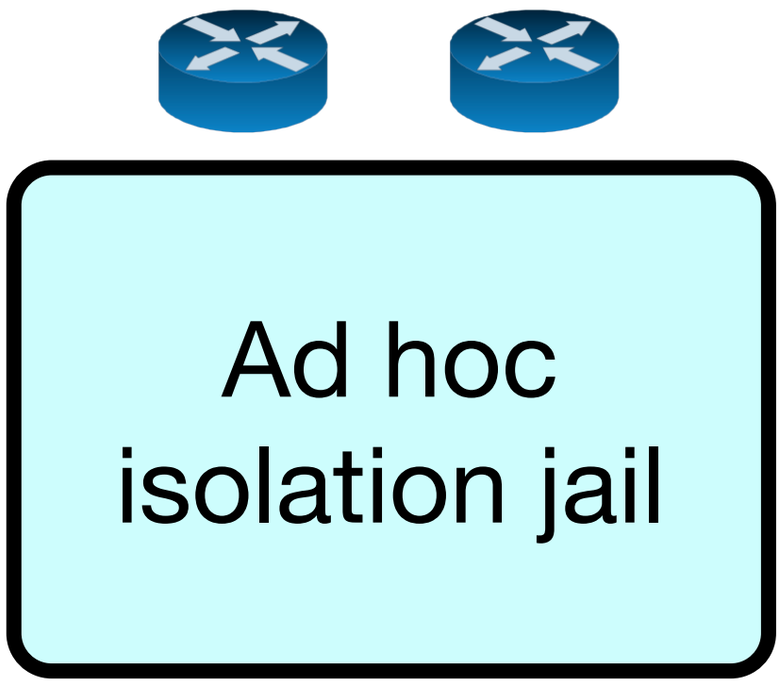
Ad hoc
isolation jail

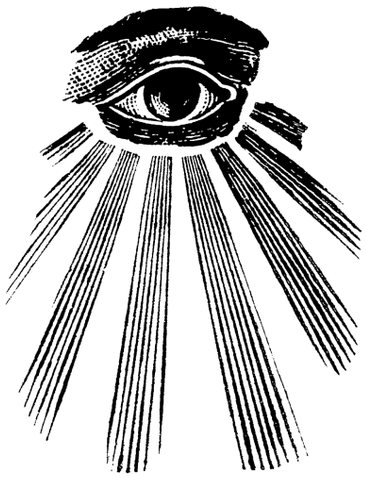

The Internet

Google™

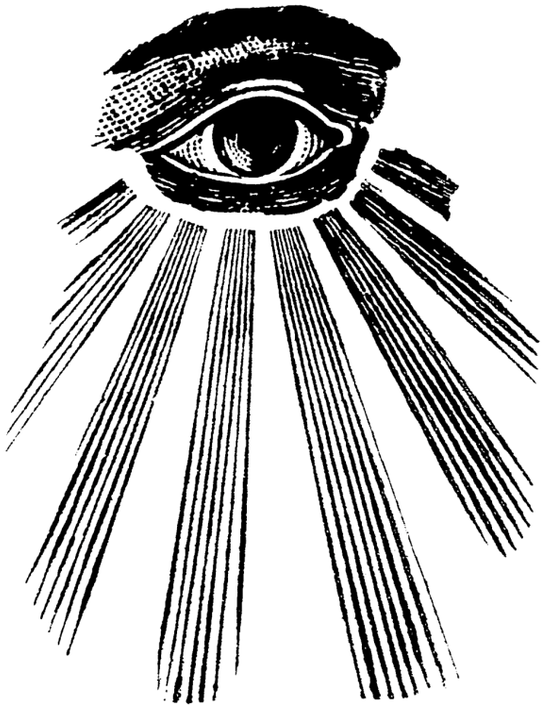

The Internet

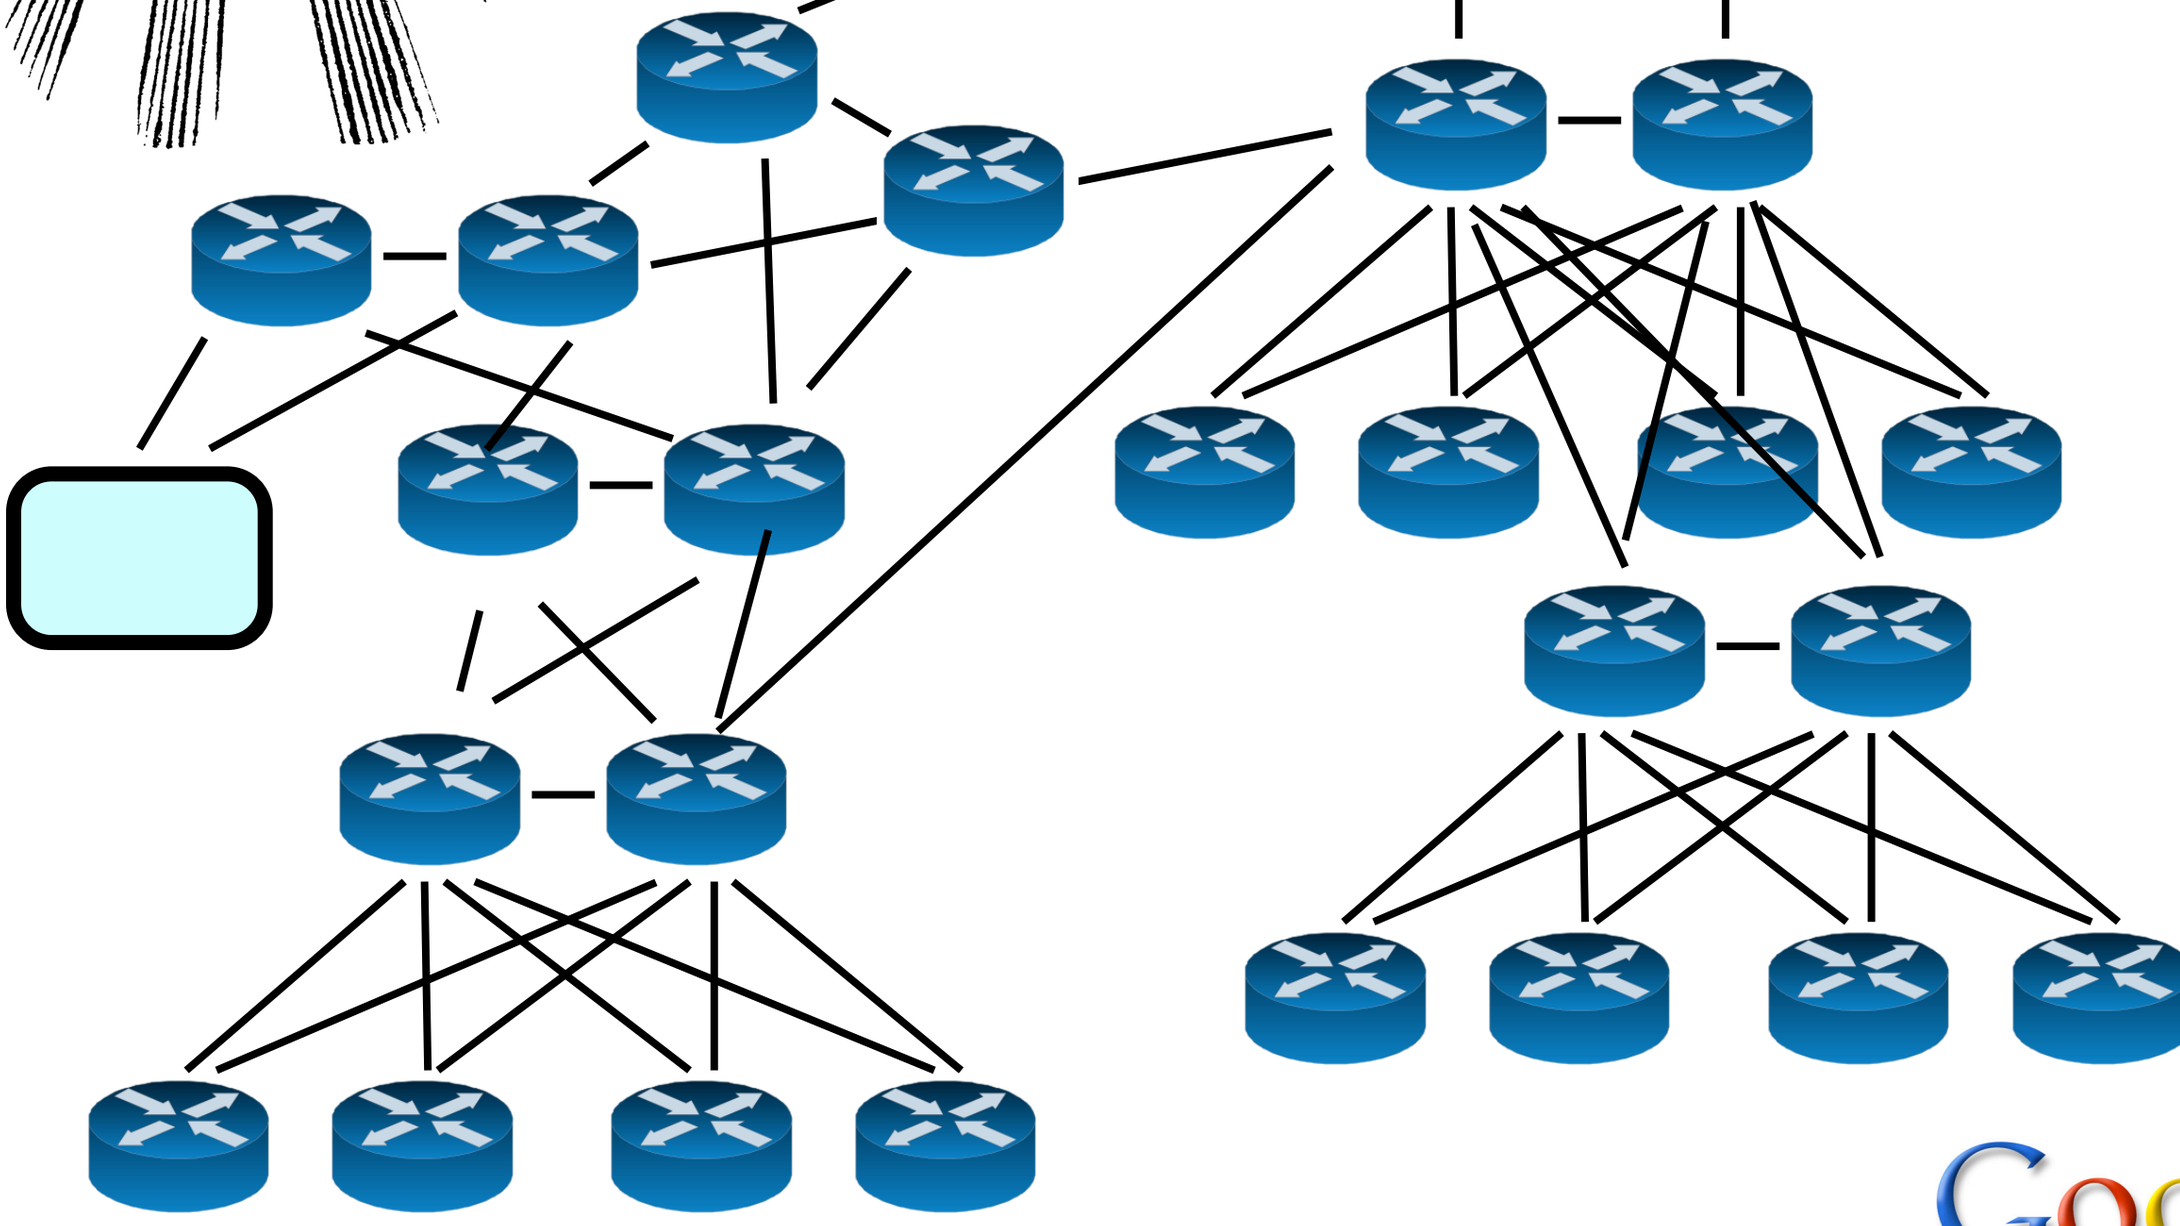

Google™

Summary

- If it isn't automated, it's wrong
- Compare your generated configs with actual configs and get diffs to zero
- Make jails to isolate nonstandardness
- Allows you to have metadata around the network

Summary

- Traffic engineering Databases (Load)
- Those data can be exported and utilized by the fleet mgmt software
- Integrate the fleet resource allocators with the real time network
- Programmatic Control
- Need incredible will to make it happen

Thanks

There is a difference between making something fool-proof and reducing the number of fools -Bill Barns

Questions/email vijay.gill@gmail.com