

Proactive Network Configuration Validation with Batfish

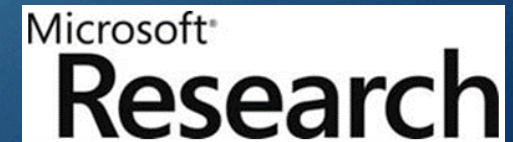
Ari Fogel
Todd Millstein



Luis Pedrosa
Ramesh Govindan



Meg Walraed-Sullivan
Ratul Mahajan
Jitu Padhye



Misconfigurations are common

Misconfigurations are expensive

2

Time Warner Cable Says Outages Largely Resolved

August 27, 2014 8:08 AM

Researchers say U.S. Internet traffic was rerouted through Belarus. That's a problem.

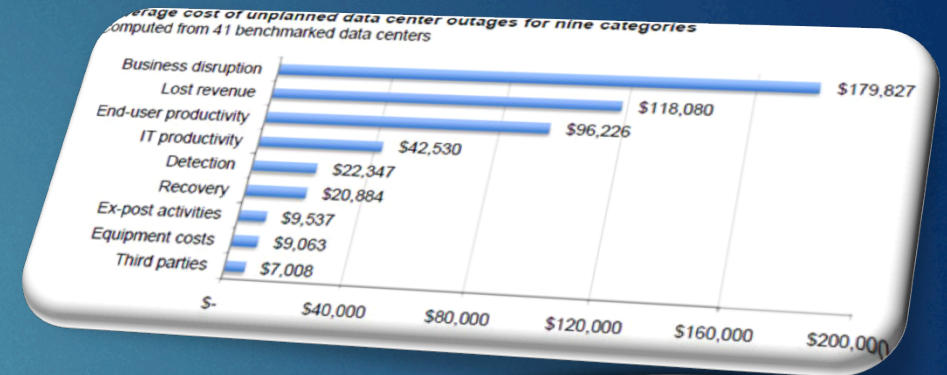


VMware Joins Cloud Outage Party
With Cloud Foundry Blackout

Intuit Service Outages Leave Frustrated
Verizon Customers In Their Wake

RIM outage costs could top \$100 million
Online Banking Upgrade contributed to Bank of America Outage

Yahoo Mail suffers outage; users react



informationWeek
THE BUSINESS VALUE OF TECHNOLOGY
IT Downtime Costs \$26.5 Billion In Losses

\$46 MM
lost from outages per year



Configuration is Hard

Low-Level Directives

- interface-level metrics
- protocol metrics
- per-network policy

Multiple Protocols:

- BGP
- IS-IS
- OSPF

Protocol Interactions:

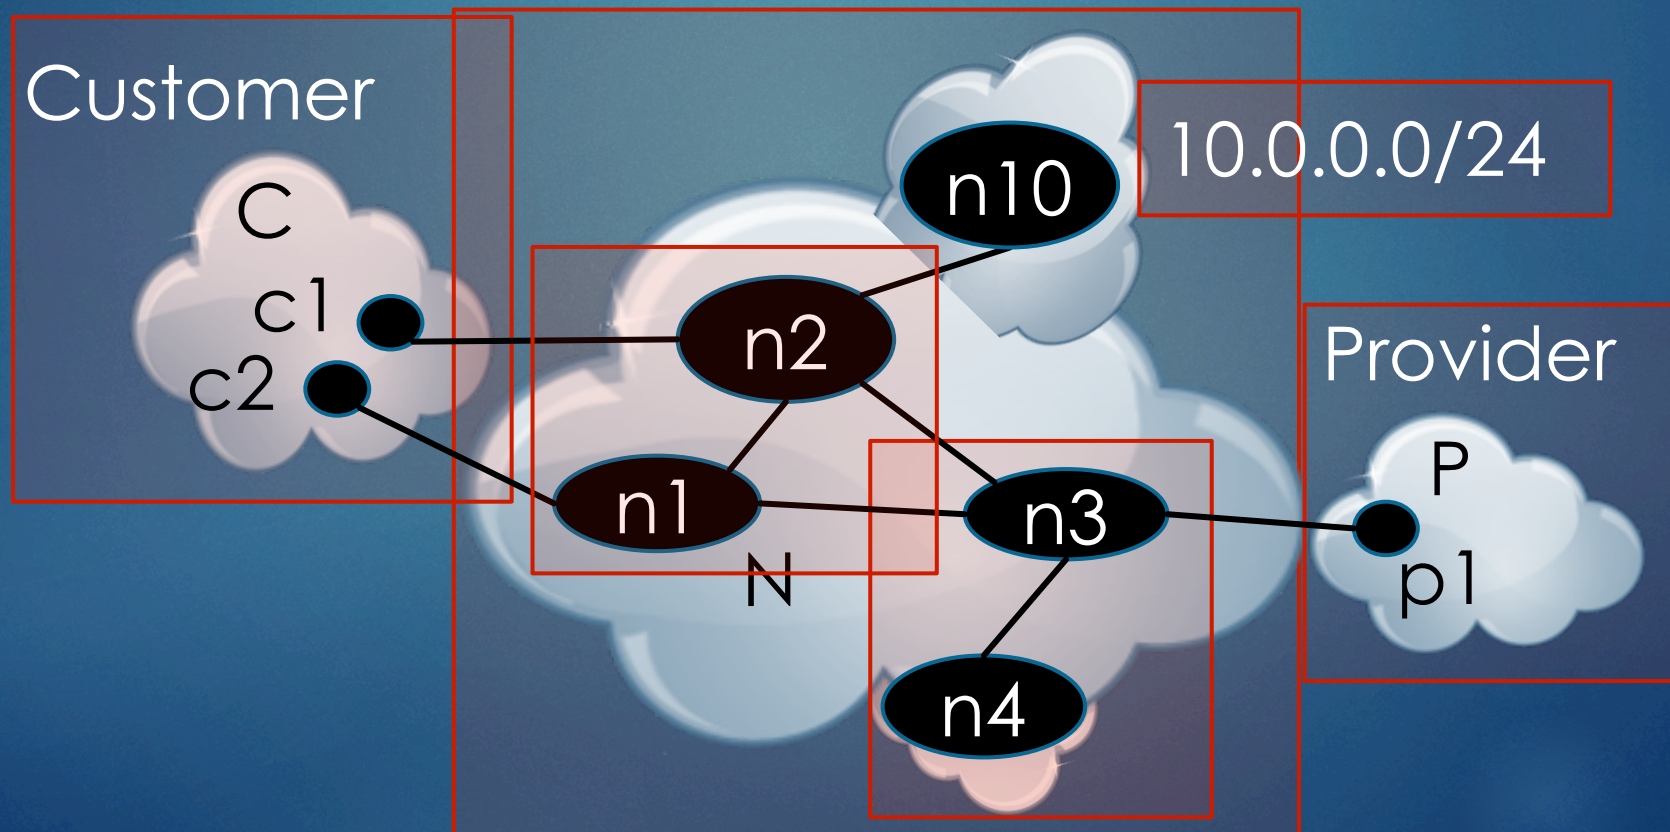
- Route Redistribution
- Protocol Preference
- Re-advertisement

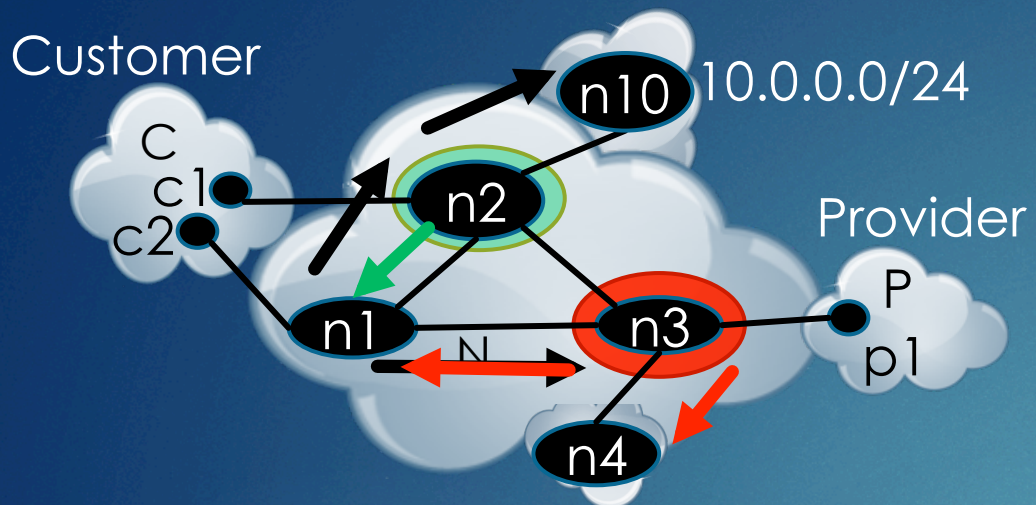
```
ospf interface int3_1 metric 1
ospf redistribute static metric 10
bgp neighbor p1 AS P Accept ALL
static route 10.0.0.0/24 drop, log
```

Example

4

- ▶ 10.0.0.0/24 should be:
 - ▶ Reachable from C
 - ▶ Unreachable from P, n4





► 10.0.0.0/24 should be:

► Reachable from C

► Unreachable from P, n4

5

3 interface int2_10 ip 10.0.0.1/24
4 ospf redistribute connected metric 10

//-----Configuration of n2-----

1 ospf interface int2_1 metric 1

2 ospf interface int2_3 metric 1

3 interface int2_10 ip 10.0.0.1/24

4 ospf redistribute connected metric 10

5 prefix-list PL_C 10.0.0.0/24

6 bgp neighbor c1 AS C apply PL_C out

4 static route 10.0.0.0/24 drop
5 ospf redistribute static metric 10

//-----Configuration of n3-----

1 ospf interface int3_1 metric 1

2 ospf interface int3_2 metric 1

3 ospf interface int3_4 metric 1

4 static route 10.0.0.0/24 drop

5 ospf redistribute static metric 10

6 bgp neighbor p1 AS P Accept ALL

Batfish

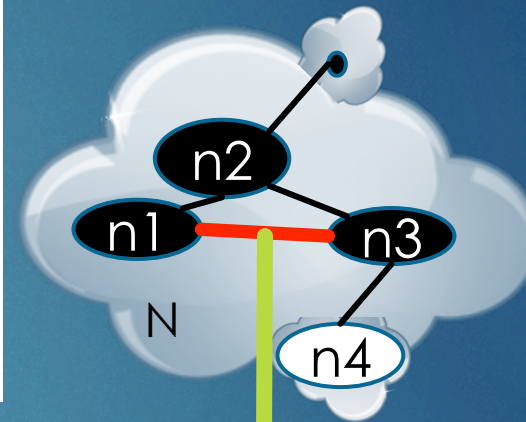
6

- ▶ Offline configuration safety checker
- ▶ Available at <http://www.batfish.org>
- ▶ Has found real bugs in real networks
- ▶ 4 stages:
 - ▶ Configuration processing
 - ▶ Configuration analysis
 - ▶ Forwarding table generation
 - ▶ Forwarding table analysis

Stage 1: Process router configurations

7

```
//-----Configuration of n3-----  
1 ospf interface int3_1 metric 1  
2 ospf interface int3_2 metric 1  
3 ospf interface int3_4 metric 1  
  
4 static route 10.0.0.0/24 drop  
5 ospf redistribute static metric 10  
  
6 bgp neighbor p1 AS P Accept ALL
```



Fact about OSPF
interface costs

OspfCost(
 node:n3,
 interface:int3_1,
 cost:1).
)

Fact about topology
LanNeighbors(

node1:n3
 interface1:int3_1,
 node2:n1,
 interface2:int1_3).
)

Stage 2: Analyze configurations

8

//-----Parsing-----

No parsing errors

//-----Basic checks-----

Undefined reference to route-map '**loch_ness_policy**'

//-----Custom checks-----

// No IP reuse

IP '**192.168.1.13**' assigned to both **rtr1:int5** and **rtr3:int6**

// All loopback networks exported into OSPF

rtr5:loopback0 neither active nor passive for any OSPF process

Stage 3: Compute forwarding tables

9

OspfExport(
 node=**n2**,
 network=**10.0.0.0/24**,
 cost=**10**,
 type=**ospfE2**).

InstalledRoute(*route*=
 node=**n1**,
 network=**10.0.0.0/24**,
 nextHop=**n2**
 administrativeCost=**110**,
 protocolCost=**10**,
 protocol=**ospfE2**}).

Fib(
 node=**n1**,
 network=**10.0.0.0/24**,
 egressInterface=**int1_2**).

Stage 4a: Identify forwarding violations

10

Counterexample of multipath consistency

```
{  
    IngressNode=n1,  
    SrcIp=0.0.0.0,  
    DstIp=10.0.0.2,  
    IpProtocol=0  
}
```

Stage 4b: Explain forwarding violations

11

Counterexample packet traces

ViolationTraceRoute(

flow={ *node*=**n1**, ... ,*dstIp*=**10.0.0.2** },
 1st hop:[n1:int1_2 -> n2:int2_1]
 2nd hop:[n2:int2_10 -> n10:int10_2]
 fate=**accepted**).

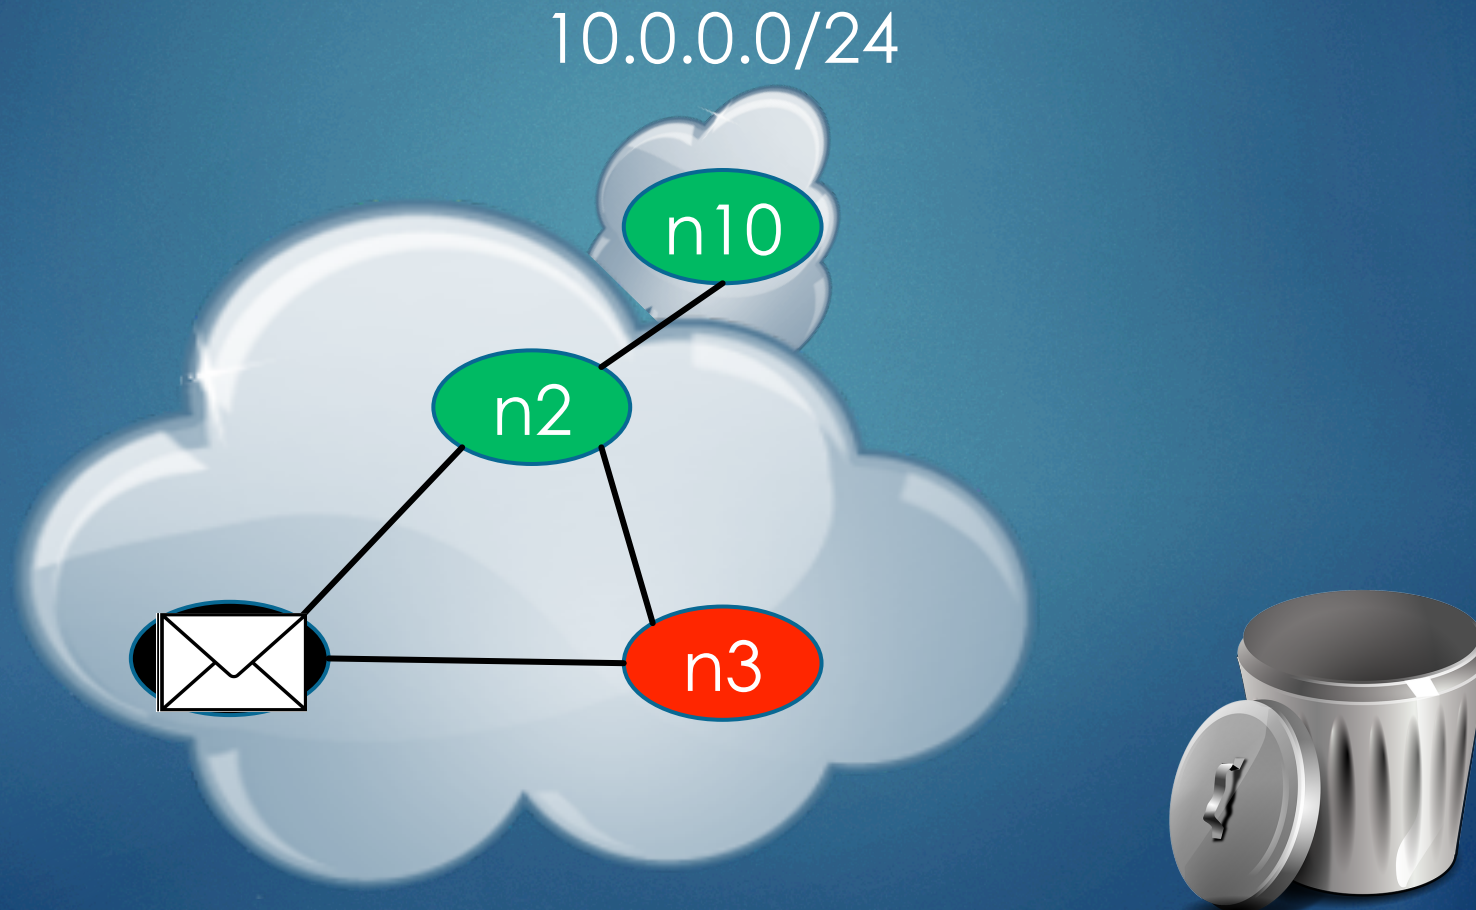
ViolationTraceRoute(

flow={ *node*=**n1**, ... ,*dstIp*=**10.0.0.2** },
 1st hop:[n1:int1_3 -> n3:int3_1]
 fate=**nullRouted by n3**).

New Consistency Properties

12

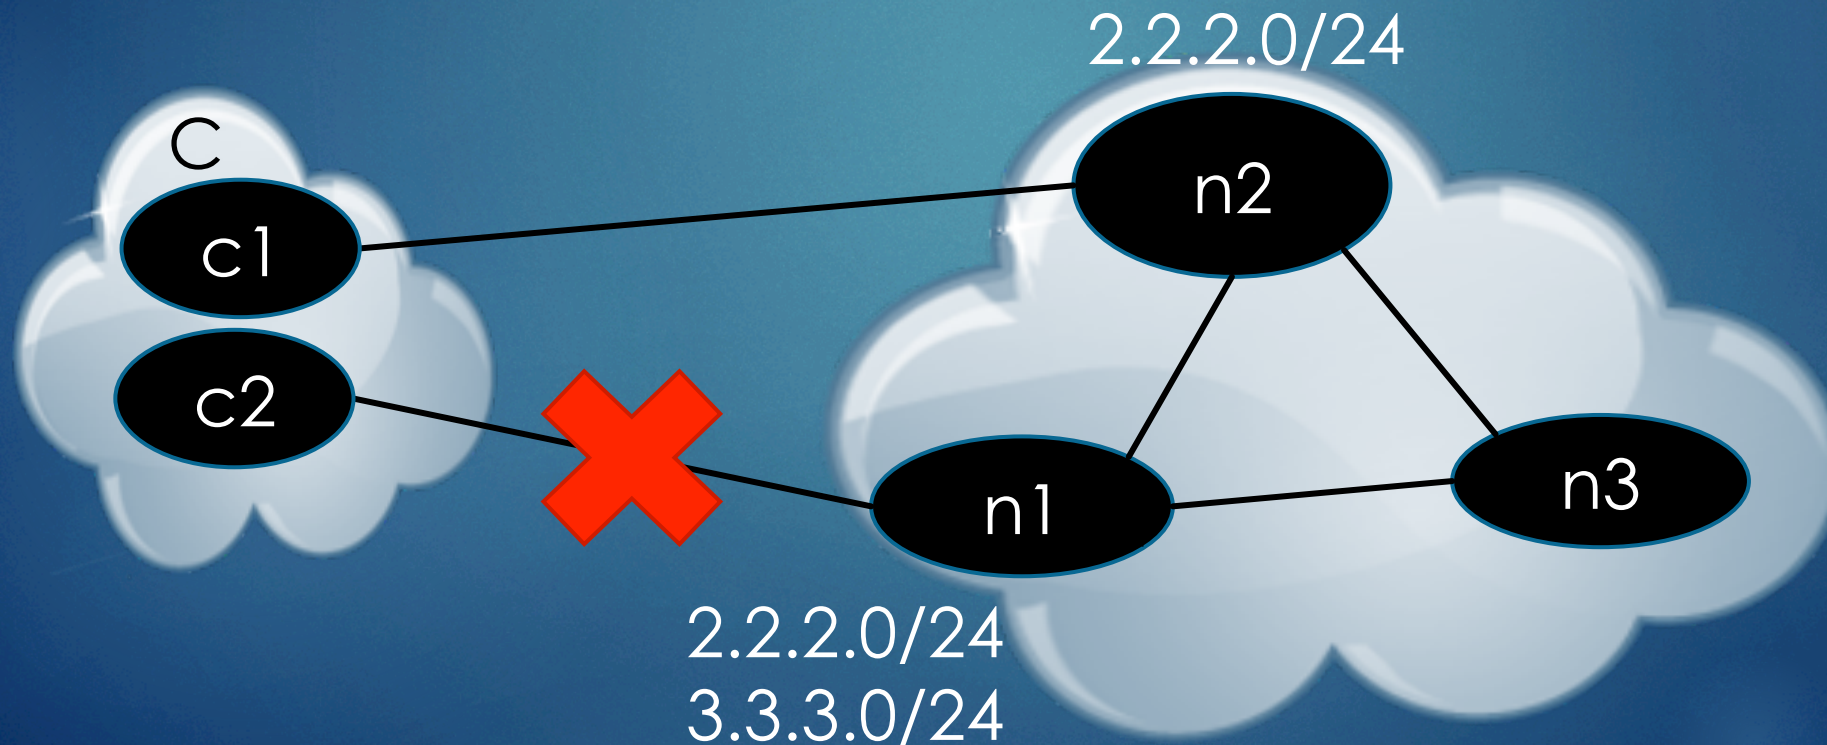
- Multipath – disposition consistent on all paths



New Consistency Properties

13

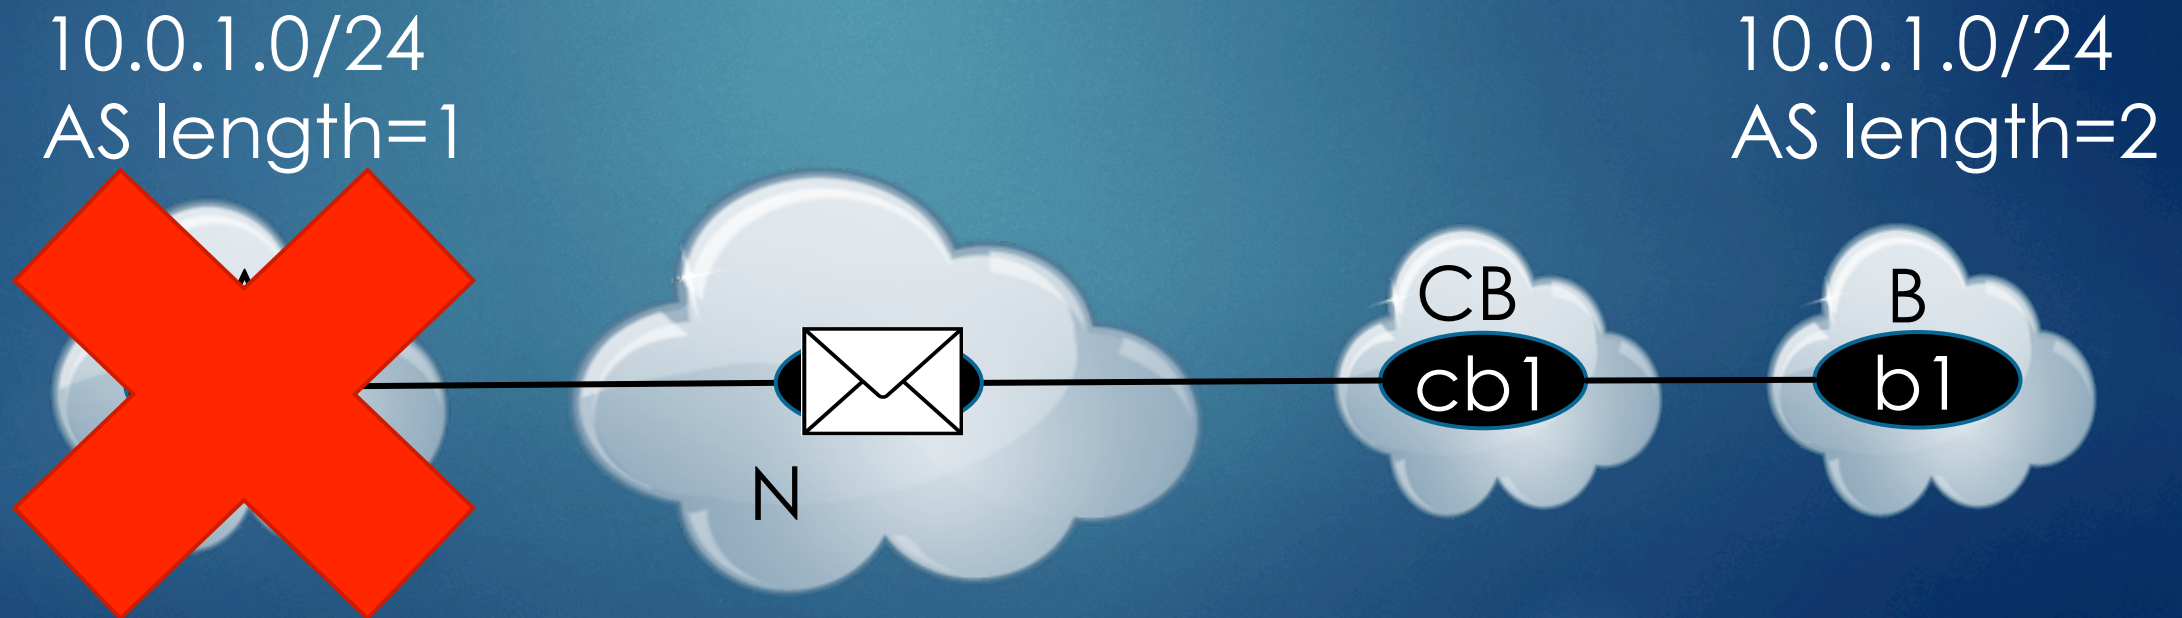
- ▶ Multipath – disposition consistent on all paths
- ▶ Differential reachability – reachability unaffected by change



New Consistency Properties

14

- ▶ Multipath – disposition consistent on all paths
- ▶ Differential reachability – reachability unaffected by change
- ▶ Destination – at most one customer per delegated address



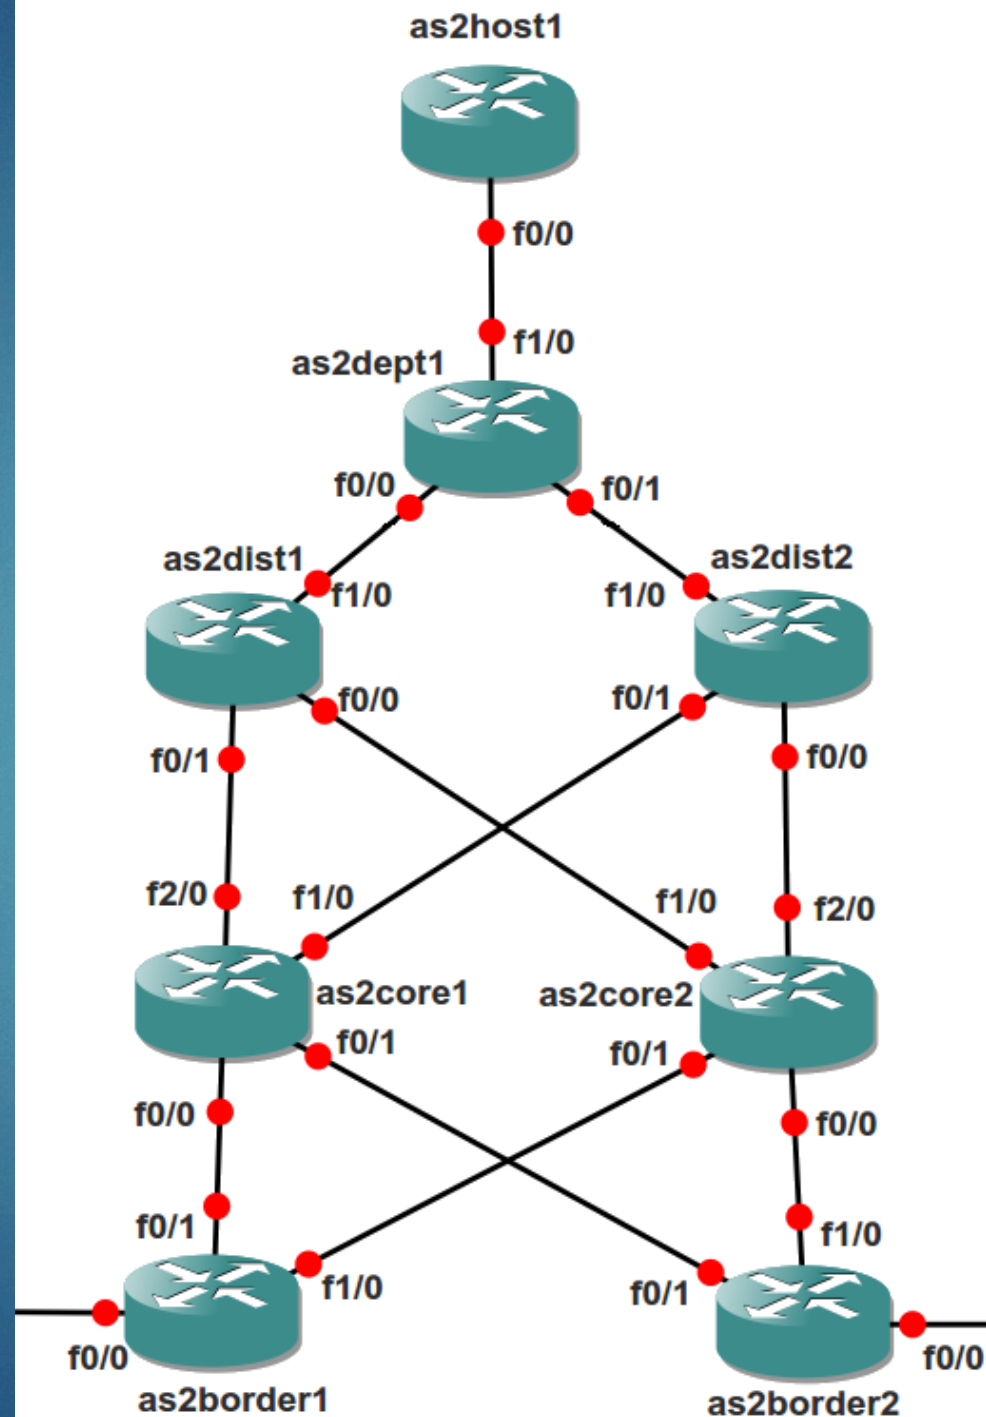
Implementation

15

- ▶ Support multiple configuration languages
 - ▶ IOS, NX-OS, Juniper, Arista, ...
- ▶ Broad feature support
 - ▶ Route redistribution, OSPF internal/external, BGP communities...
- ▶ Unified, vendor-neutral intermediate representation

Demo

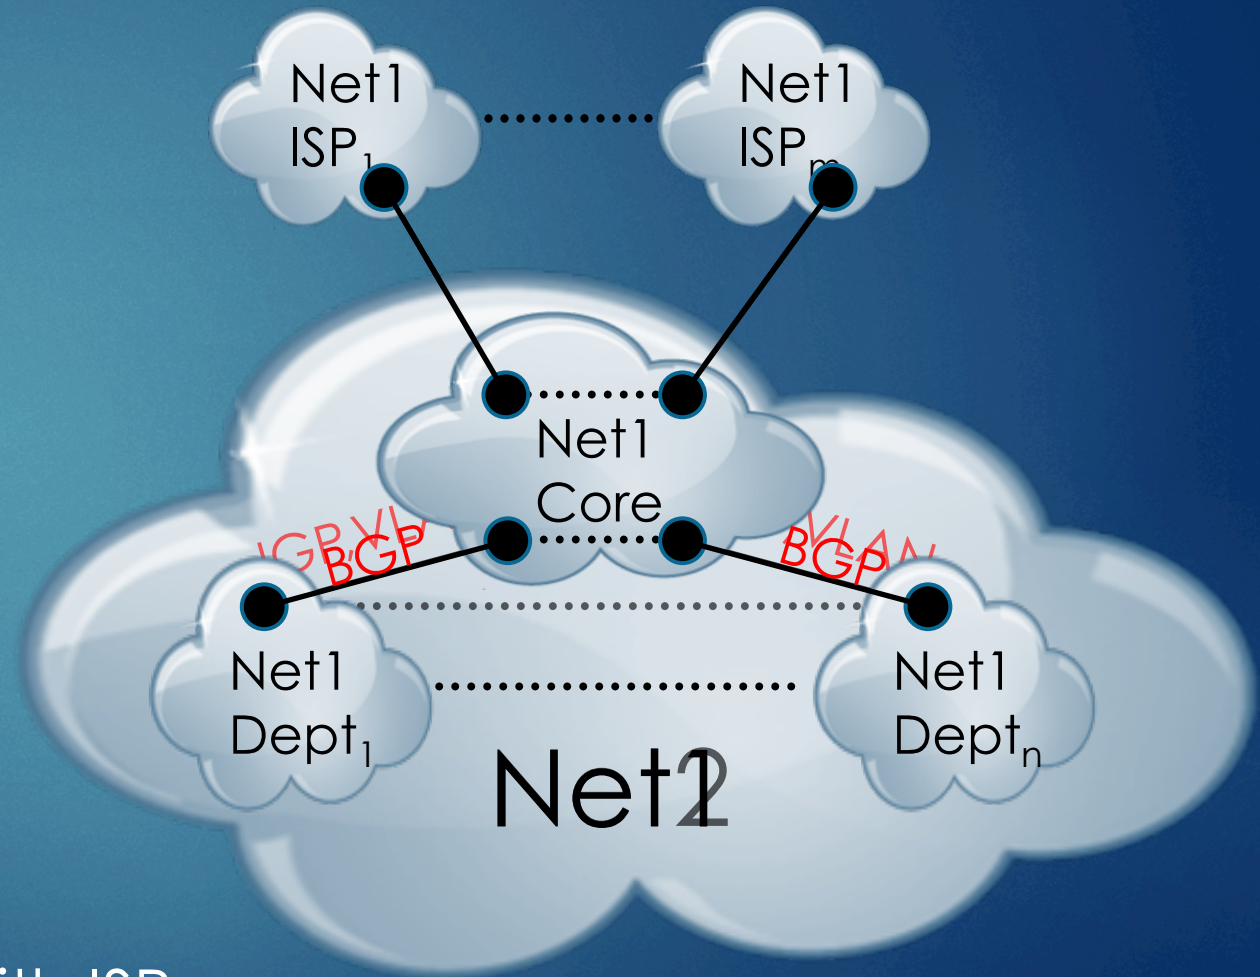
- ▶ Simplified version of Net1
- ▶ Cisco configuration files
- ▶ Multiple seeded bugs



Evaluation

17

- ▶ Two large university networks
- ▶ Net1 – 21 core routers
 - ▶ Federated network
 - ▶ Each department is own AS
 - ▶ Heavy use of BGP
- ▶ Net2 – 17 core routers
 - ▶ Centrally controlled
 - ▶ Heavy use of VLANs
 - ▶ Single AS
 - ▶ BGP communication only with ISPs



“P.S. WRT the prefix that was dual assigned from yesterday, one of my NOC [network operations center] guys stopped by today to ask what voodoo I was using to find such things :)” [emphasis added]
– email from the head of the Net1 NOC

Results

19

	Invariant	Total Violations
Net1	Multipath	32
	Diff.Reach.	16
	Destination	55
Net2	Multipath	11
	Diff.Reach.	77

Selected Violations

20

- ▶ (Multipath) Black-hole route cost too low (equal)
- ▶ (Diff.Reach.) Only one interface underlying VLAN
- ▶ (Destination) Prefix assigned to multiple departments

Conclusion

21

Take survey so we can support *your* network features and requirements in forthcoming versions:

<http://www.batfish.org/survey>

Send feedback/questions to:
arifogel@ucla.edu