OSPF to ISIS

Vijay Gill <vijaygill9@aol.com>
Jon Mitchell <jrmitche@aol.net>

RIPE 47, Amsterdam

January 2004

Why

"But in our enthusiasm, we could not resist a radical overhaul of the system, in which all of its major weaknesses have been exposed, analyzed, and replaced with new weaknesses."

-Bruce Leverett

- Features
  - Convergence
- Security
- Simplicity

ISIS is Hard

Lets go shopping

*but to learn ISIS you have to know the secret handshake and be a *%##%ing 33rd level mason*

-Chance Whaley

Security

- http://www.nanog.org/mtg-0006/katz.html
- Packet bombs
- Wasn’t as big of a deal for AOL
  - We have packet filters on most line cards
    - Most is not ALL
- Runs directly on L2
  - Harder to spoof or attack

Simplicity

- Found out we didn’t need areas
  - Added complexity
    - Configuration
    - Typos
- Slowed it down
  - DV
- Flat area easy to configure and maintain
    - Stupid, but no stupider (apologies to Einstein)

State of the Art

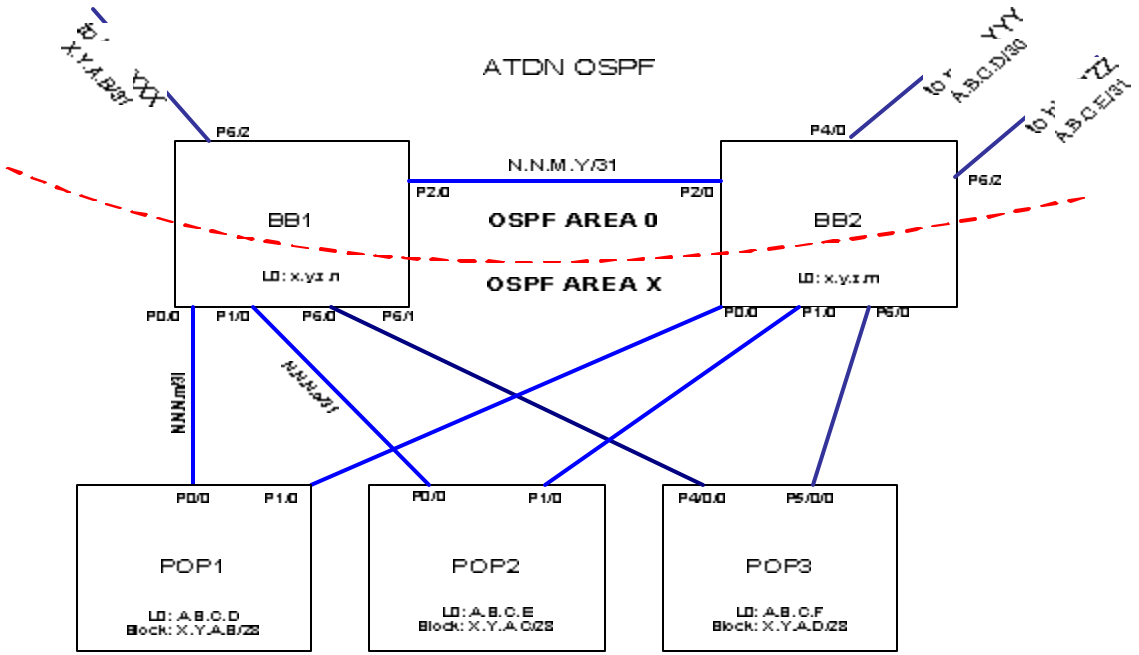

BB Sample Config

```
router ospf 1
log-adjacency-changes
area 0 authentication
area x authentication

passive-interface Loopback0
network A.B.C.0 0.0.3.255 area 0
network A.B.D.0 0.0.1.255 area X
network A.B.C.0 0.0.7.255 area X

maximum-paths 6

area X range A.B.C.x 255.255.255.240
area X range A.B.C.y 255.255.255.240
```

Note: Area X is the BGP cluster-ID of the site

POP Sample Config

```
router ospf 1
log-adjacency-changes
area X authentication
redistribute connected subnets
passive-interface Loopback0
network A.B.C.0 0.0.1.255 area X
....
Maximum-paths 6
```

Strategy

Ships in night

- Run parallel
- Verify routes

- Raise OSPF admin distance
- Verify network after change
- Remove OSPF

The plan is in the works, but we have not activated the implementation phase.
-Frank Caddeo

Main Backbone Nodes

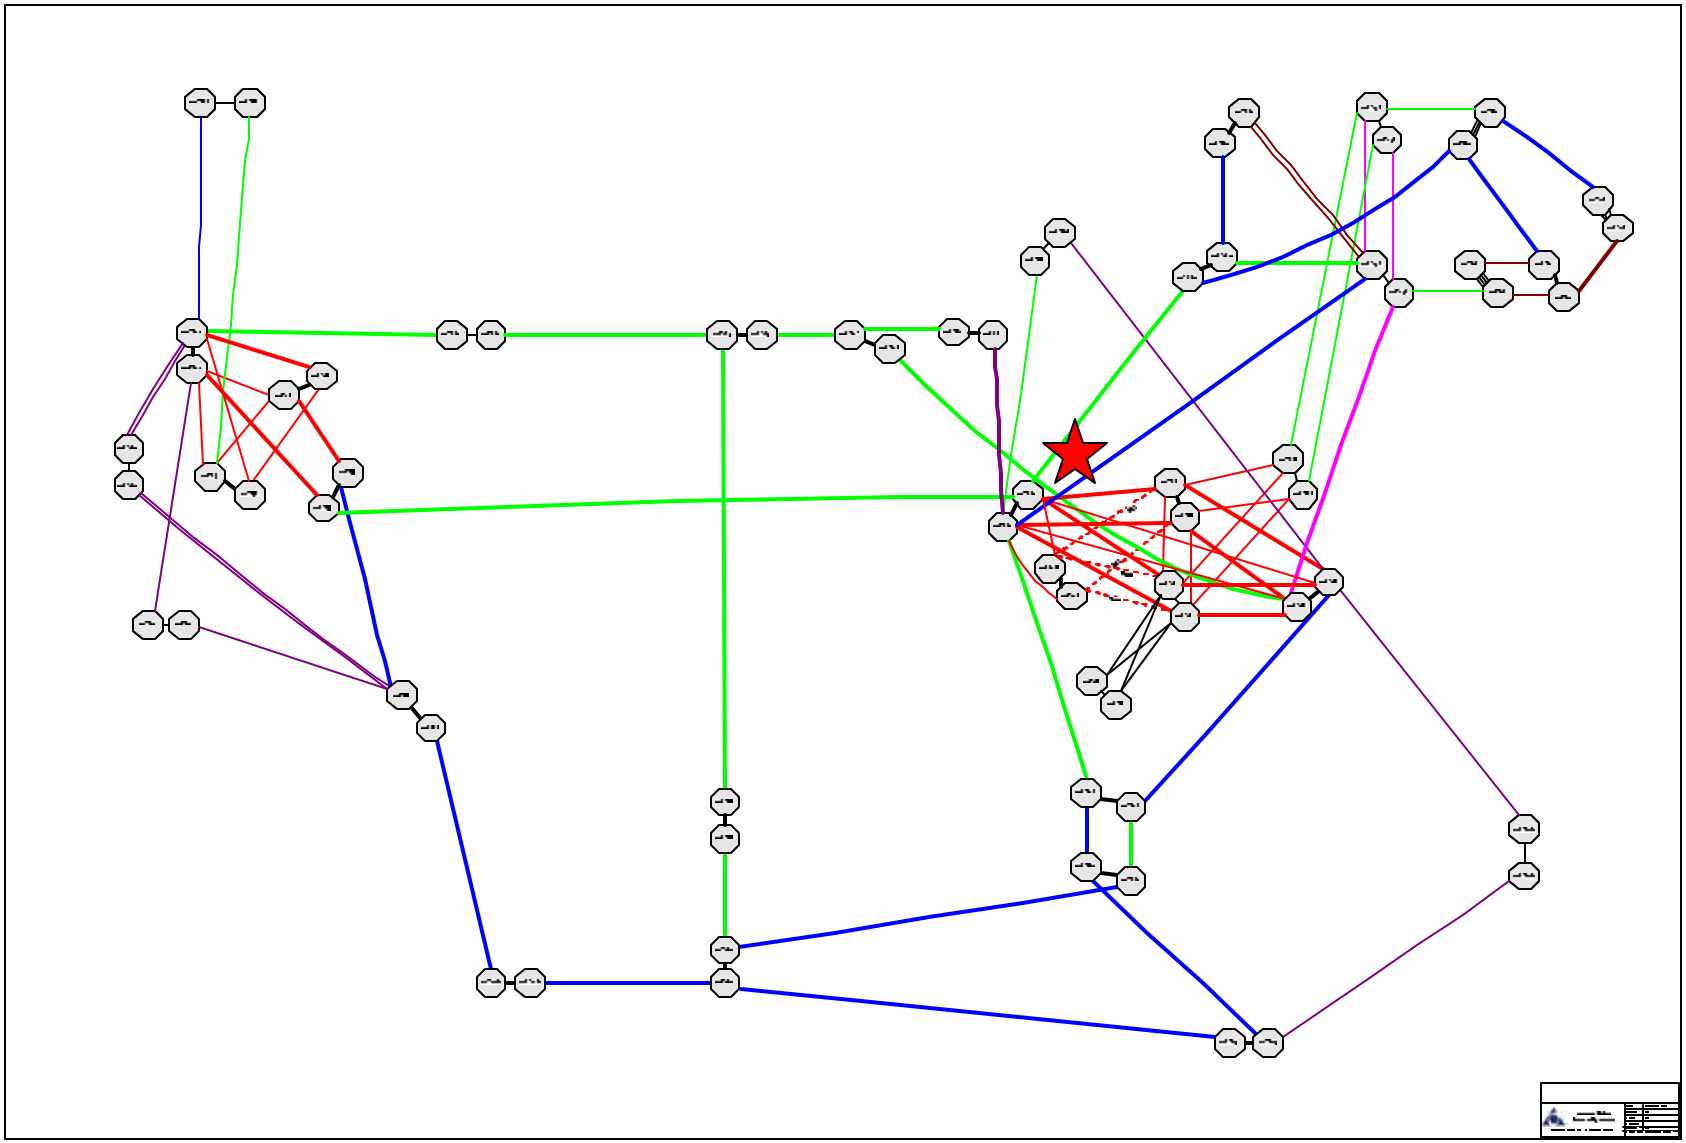

Out of Band

“OOB is the saving throw when you @#$% up”

-RS

- Verified OOB reachability to all POPs beforehand

IS-IS Migration Prep

Pre-Migration

- Load IS-IS configuration built with scripts on RTL routers
  - Non Customer PoP
- Develop/test scripts to check IS-IS neighbor relationships and route consistency

Migration Week

- Load IS-IS configuration
- Verify IS-IS neighbor relationships
- Verify LSPs in IS-IS database
- Change OSPF administrative distance to 254
  - On some edge routers

Some mornings, it's just not worth chewing through the leather straps.
-Emo Phillips

Migration Week (cont)

- Compare IS-IS and OSPF routes on pair of pop routers
- IS-IS vs. OSPF cost check on all interfaces in network
- Change OSPF administrative distance of all remaining routers to 254
- Verify no OSPF routes in fowarding table
- Basic network reachability
  - Ping all routers
  - Check connectivity to some external sites
    - Standard NOC monitoring

Your rules are really beginning to annoy me
-Snake Plissken

Post-Migration

- No verification
  - Verification done as part of migration
- Run a script to remove the OSPF configuration from all ATDN routers

We had more than enough genuine headaches as it was, and trivial aesthetic concerns weren't even close to making it onto our agenda.
-Geoff Miller

Current Setup

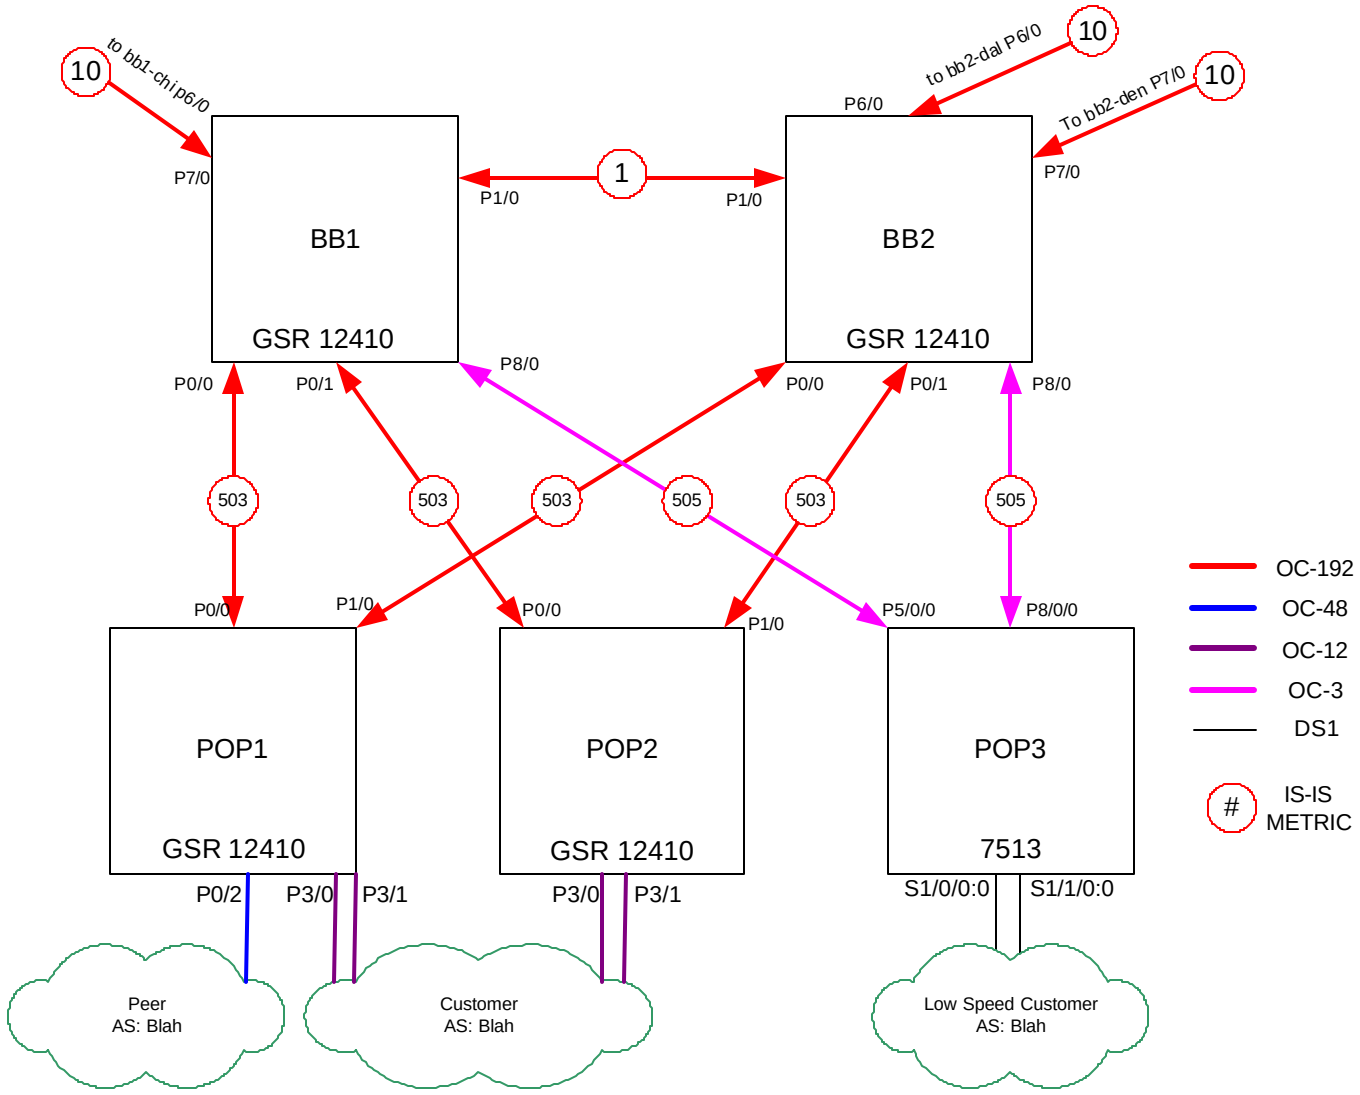

Config Bits

- *!*
- *interface Loopback0*
- *isis metric 1 level-2*
- *!*
- *interface POS5/0*
- *description P5/0: bb1-nye-P5-0-pop1-nye-P5-0 (66.p.x.y/31 direct-cabled)(T=pbNYE)*
- *ip router isis*
- *isis metric 503 level-2*
- *isis password ISISPASSWORD(hint, this isn't the real password) level-2*
- *!*
- *router isis*
- *passive-interface Loopback0*
- *maximum-paths 6*
- *net 39.752f.0100.0014.0000.5000.1668.router.id.inIPv4.00*
- *is-type level-2-only !**Why Level 2?***
- *domain-password this-isn't-the-real-password-either*
- *metric-style wide !*
- *external overload signalling !* **Ensure that IS-IS will tear down adjacencies when dCEF is disabled on an interface**
- *set-overload-bit on-startup wait-for-bgp !* ***Avoid placing router on IGP SPF before bgp***
- *max-lsp-lifetime 65535*
- *lsp-refresh-interval 65000*
- *no hello padding !* ***Hello padding to mtu is deprecated***
- *log-adjacency-changes all*
- *!*

Design

- All connected interfaces are redistributed into BGP
- IS-IS will be preferred
- Redistribution into BGP chosen to reduce the number of links in the SPF
  - Is it an issue in practice
    - Not really

cluelessness leads to flapping... flapping leads to dampening... dampening leads to suffering
-RS

Cost and RR Design

- Backbone links
  - Used OSPF metrics
  - BB-POP Interconnects
    - OSPF metric + 500
    - Avoids Inversion on BB-BB link failure
    - Mirrors OSPF w/ Areas behavior
  - MED oscillation issue
    - Full mesh of POP routers
    - No client-to-client reflection
    - Cost (InterPOP) > cost differences IntraPOP
- New cost out procedure
  - add 10000 to the interface

Timeline

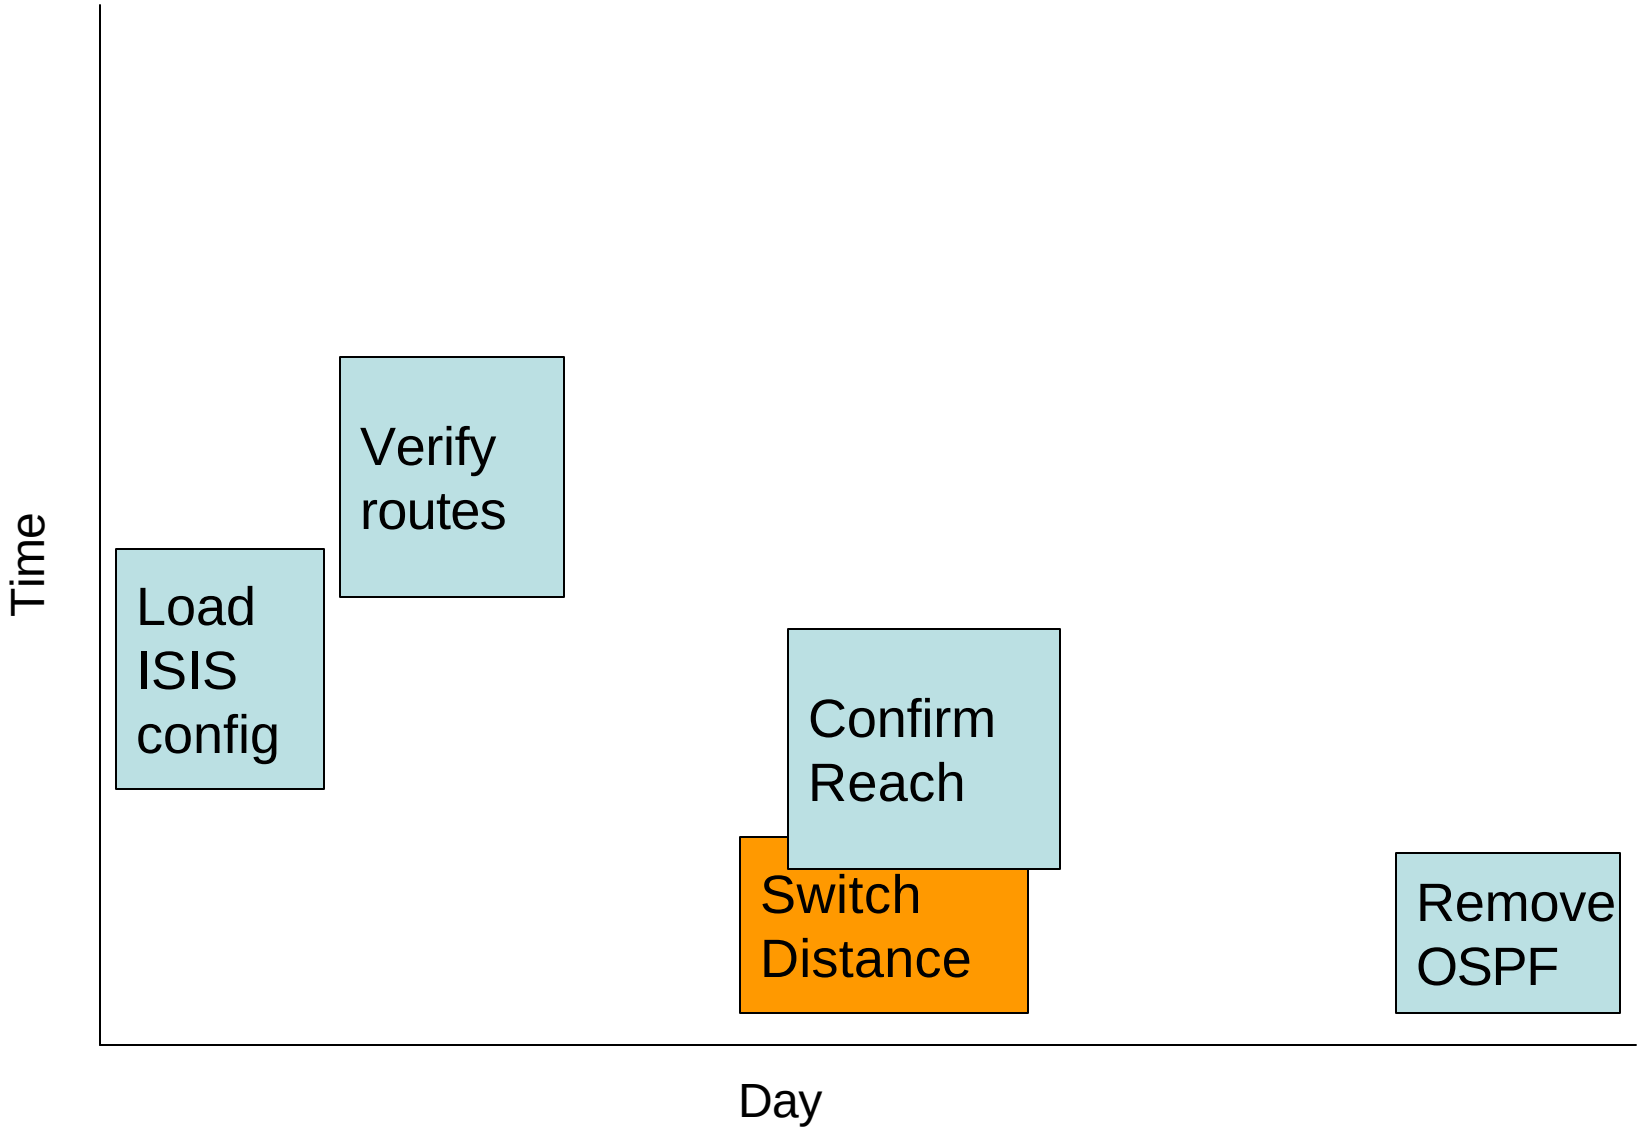
Time
Load ISIS config
Verify routes
Confirm Reach
Switch Distance
Remove OSPF
Day

Loading ISIS Config

- Non Disruptive
- Config was loaded in a three hour window, Monday 0600-0900 EDT
- Script (OSPF) -> IS-IS
- Output was copied to each router
- No IS-IS routes in use

If you can't remember, then the claymore is pointed at you

Route Verification

- Compare IS-IS neighbor topology with OSPF
  - show clns neighbor
  - show ip ospf neighbor
- Check IS-IS database on all routers
  - Ensure all other routers LSP’s installed in IS-IS database (sh isis database)

Route Verification

- On selected edge routers
  - Change OSPF admin distance to 254
  - Verify traffic to peers
  - Compare IS-IS and the OSPF routes
    - All routes in the network are correctly in IS-IS?
- Go or No Go

Great ideas, in theory, should not be hampered by pesky reality
-Dys

MED

- Changing metric affects MEDs
  - New metric in the BGP one minute after distance change
  - Ratchet down
    - Does not propagate for another 10 minutes
  - One Large Peer – LP
    - Listened to MEDs
    - Not enough capacity to fit all of traffic in one circuit
  - All routers connected to LP
    - Migrated at roughly the same time
    - Manually cleared soft out after the metric advertisement updated

The Big One

- Flip Admin Distance
  - IS-IS routes are preferred
- Current network metrics are consistent with config files?
- Slow Start
  - Manually change admin distance to 254 on more edge POPS
- Go No-Go?
- Script to flip the rest
  - From the edge to the center (with respect to ops2)
  - In order - LP, europe, asia, brazil, us-pop, us-bb, and dc
- External routes in OSPF now in iBGP

Main Backbone Nodes

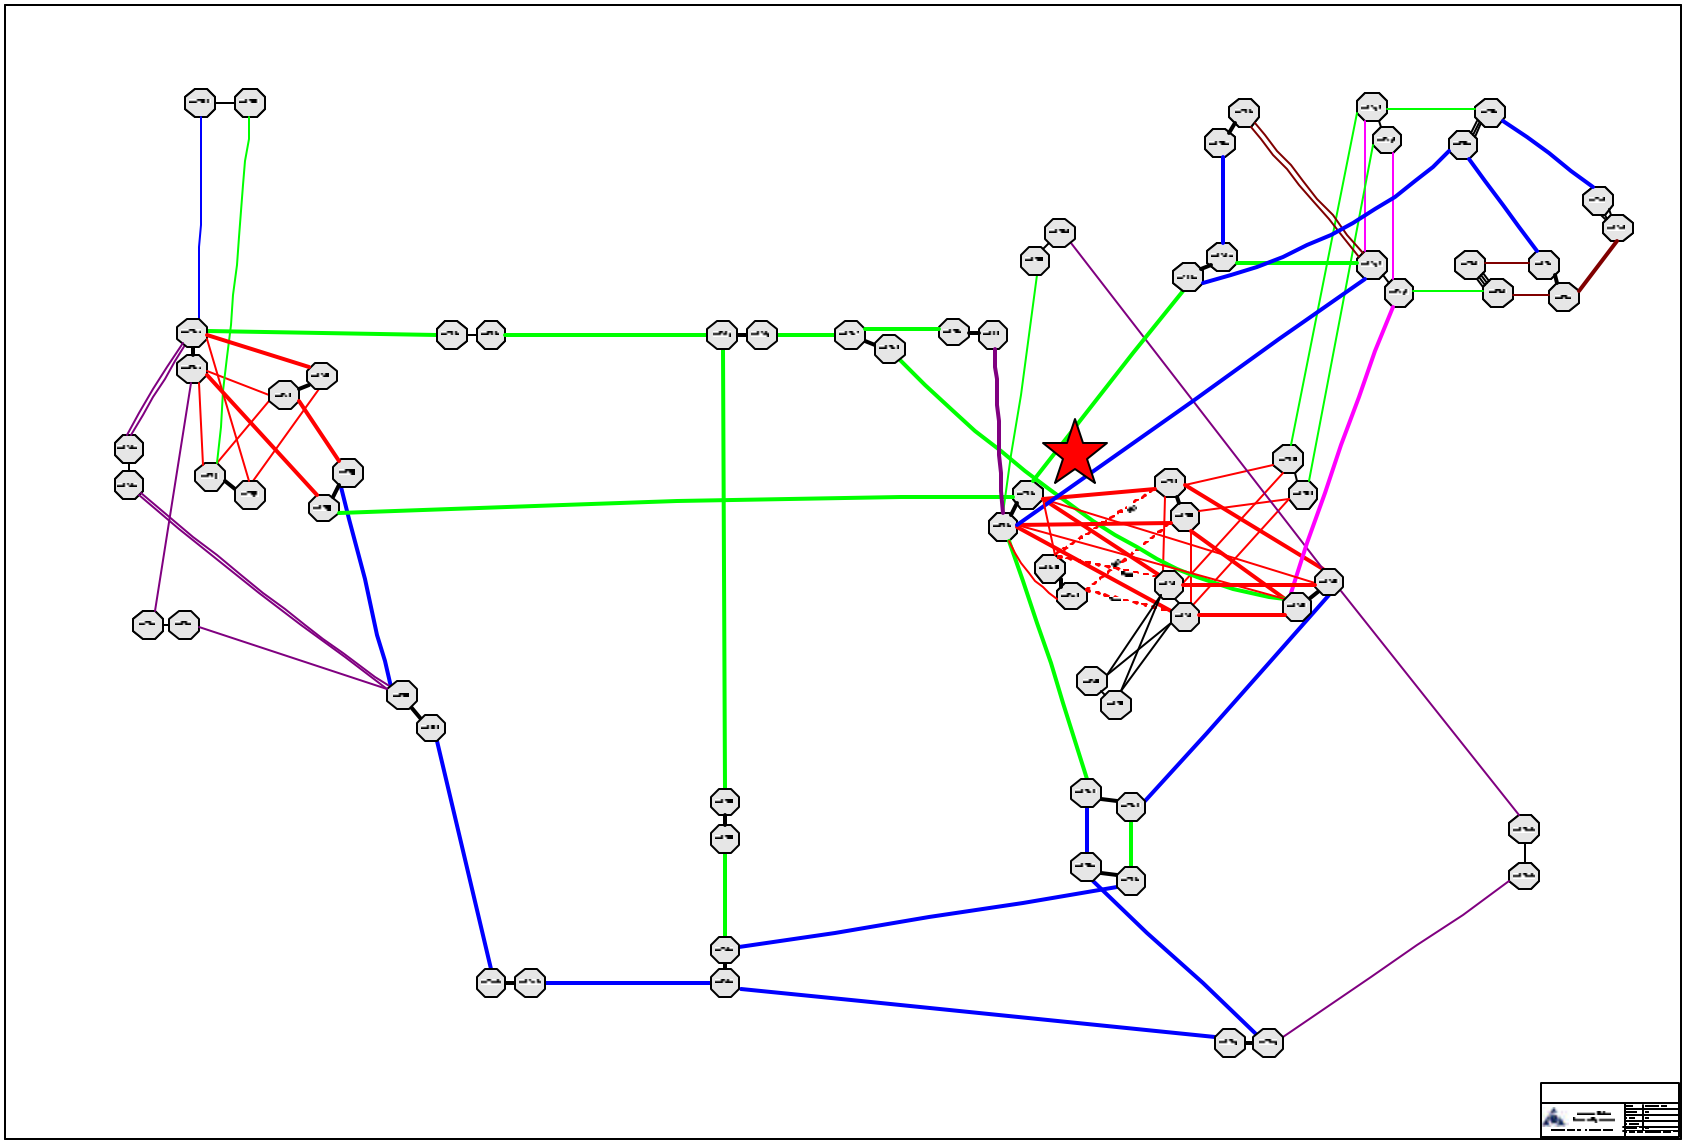

Routing

- Convergence time for the installation
  - <1 second
- No CEF updates
  - Costs changed but PATHS didn’t
- All production traffic is routed to Edge router loopbacks (n-h-s)
- Rollback
    - Remove admin distance command
    - Pre-written script

This thing severely violated the Rule of Complexity as applied to the problem. The Rule of Complexity states that if an answer seems too complicated to be the right answer, it is the wrong answer.
-Steve Cutchen

Removal of OSPF configuration

- After burn in
  - 0300 EDT
  - OSPF configuration removed
    - Non-disruptive change
  - Old OSPF configs archived via RANCID

*We **are** jolly green giants, walking the earth with routers.*
-Christopher Morgan (after no router ospf 10 at MFN)

Subject: From the install file
Date: 6/25/2003
To: network-ops-list@aol.net
CC: John

Network Install Doc for Non-Bounce June 25, 2003
General Maintenance (times noted with attribution):

c) Switching ATDN backbone from OSPF to ISIS as the igp. 0300
Expected Impact: None

Line of Truth

Dog will hunt
/vijay

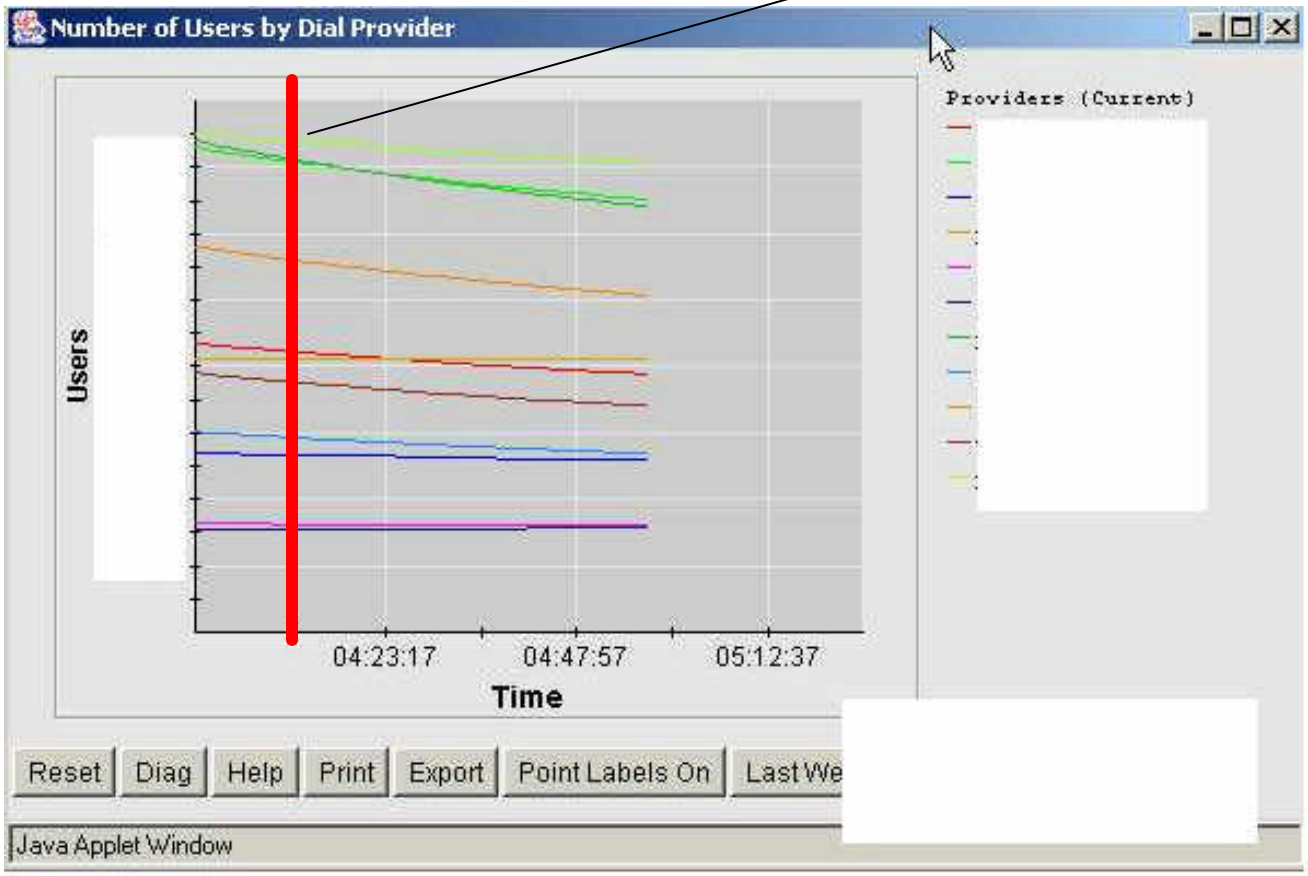

Traffic

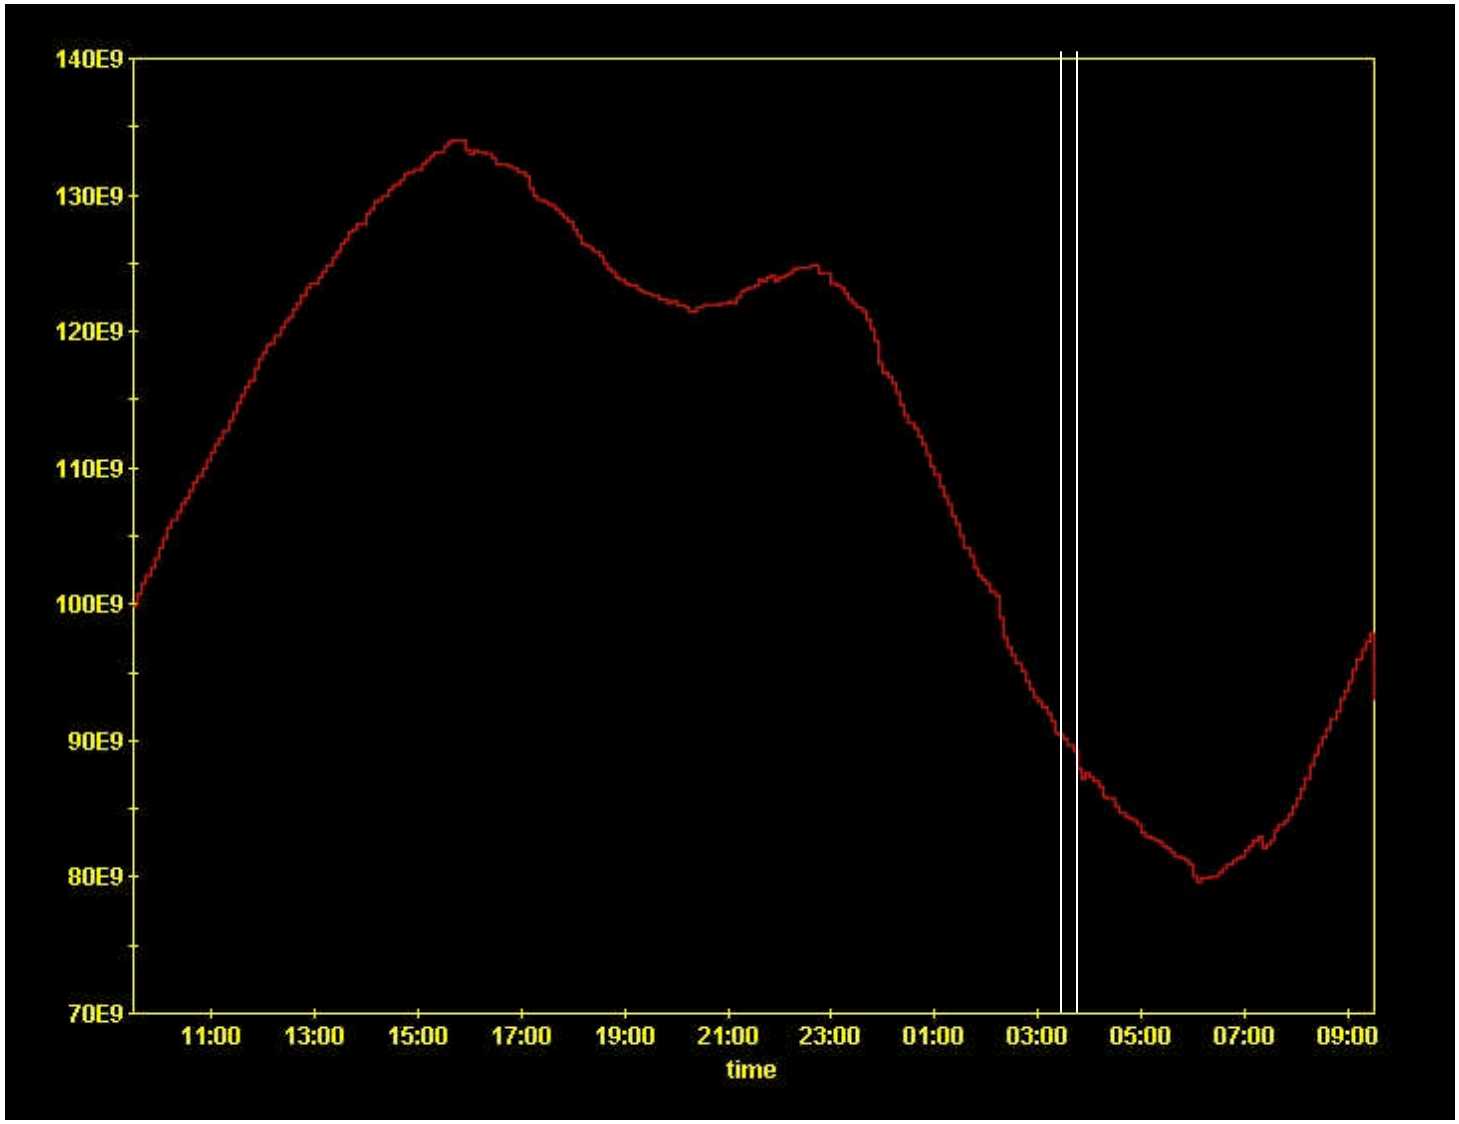
140E9
130E9
120E9
110E9
100E9
90E9
80E9
70E9
11:00
13:00
15:00
17:00
19:00
21:00
23:00
01:00
03:00
05:00
07:00
09:00
time

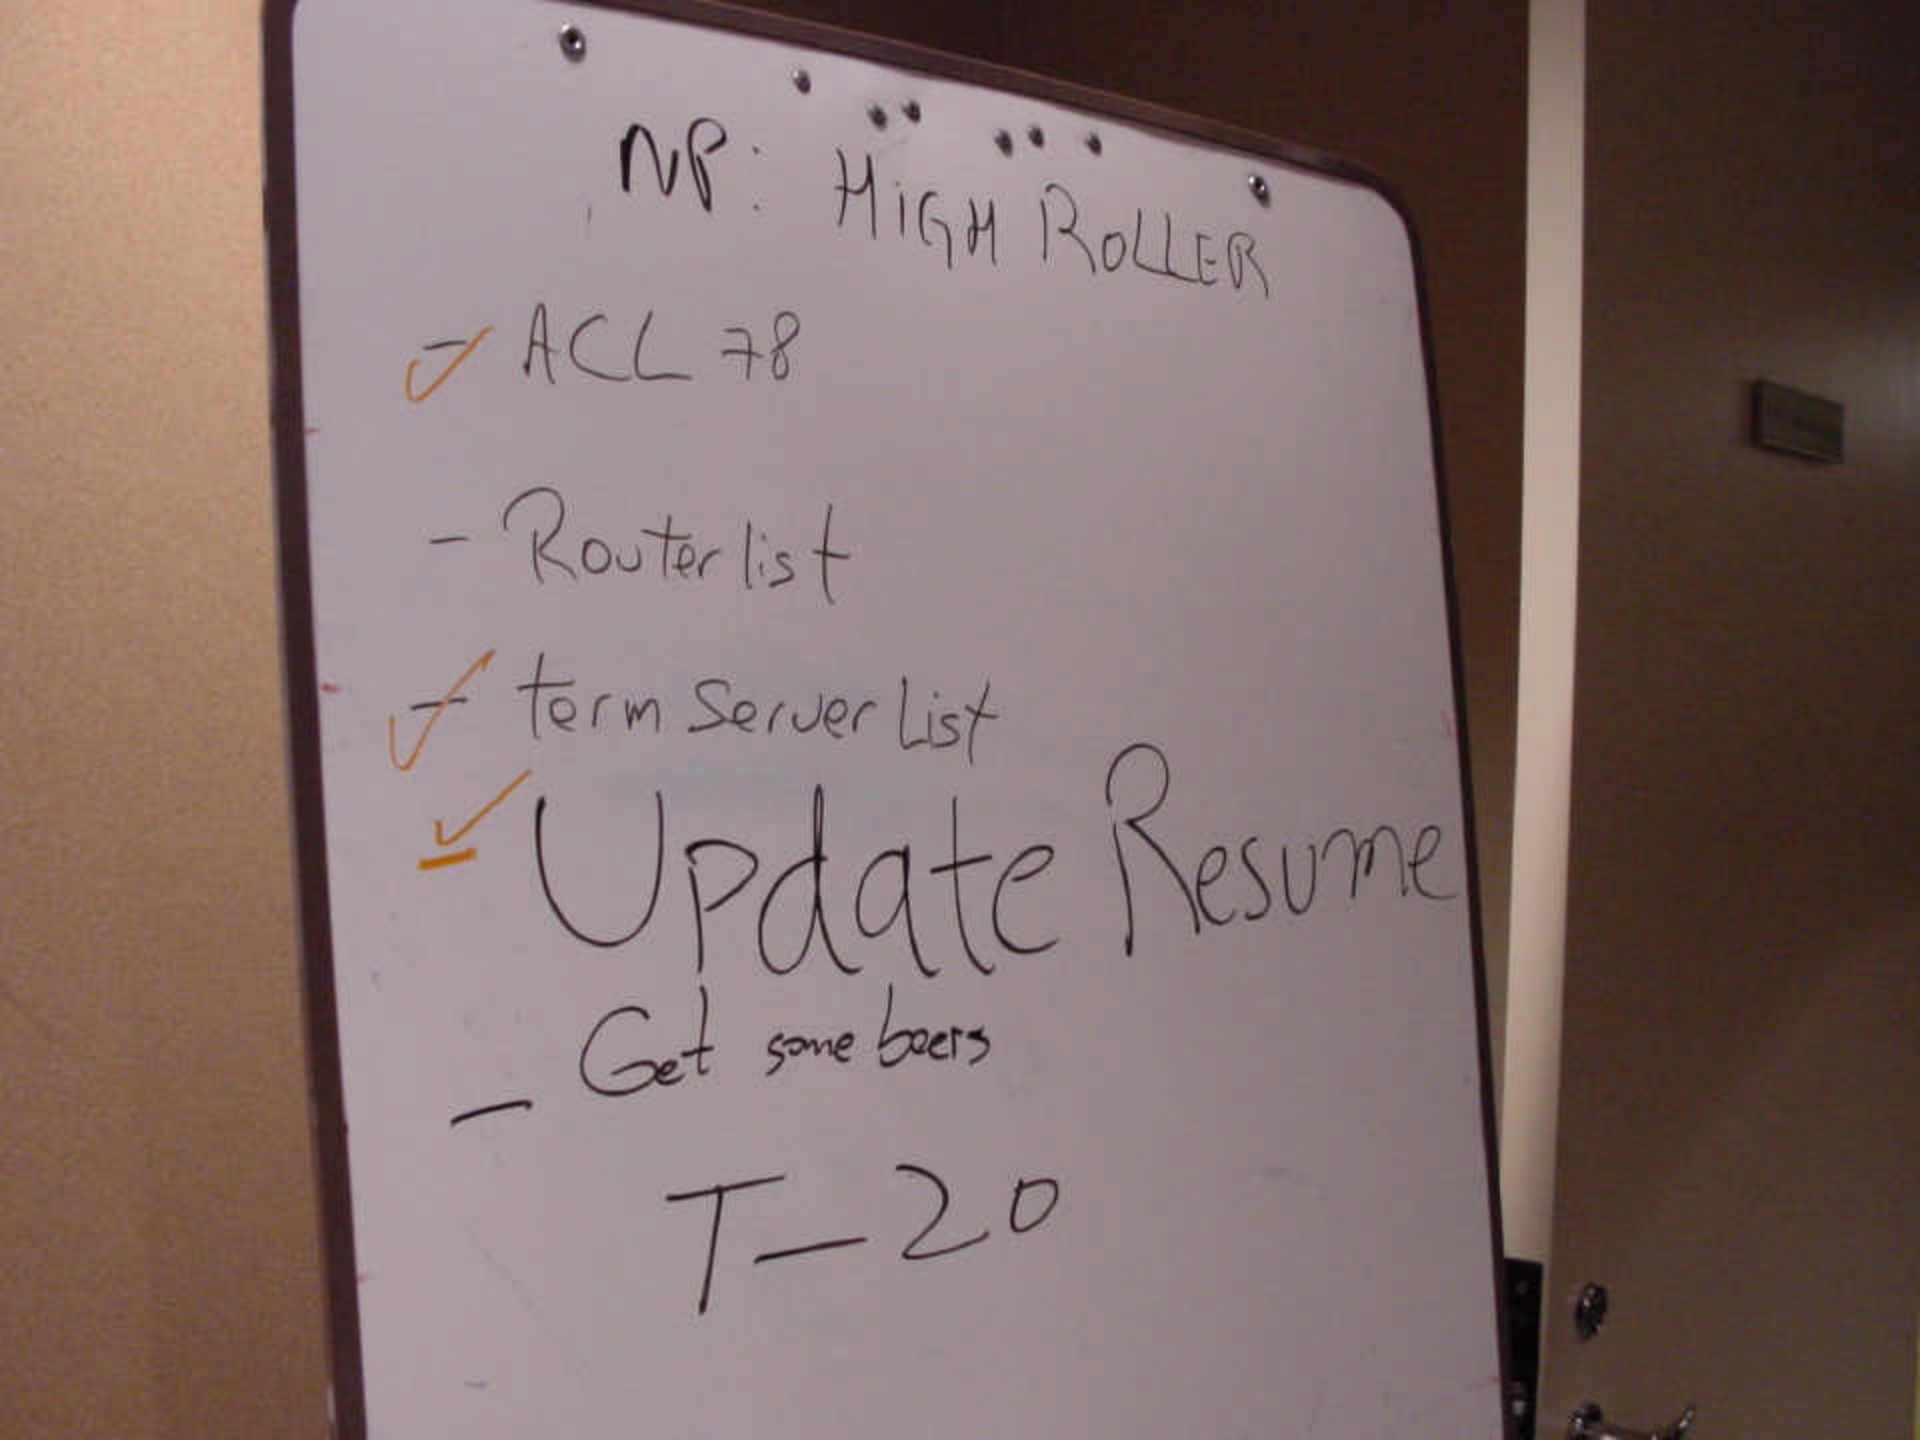
NP: HIGH ROLLER
ACL 78
- Router list
term Server List
Update Resume
- Get some beers
T-20

Questions?

You thinking about smoking off the MPLS hookah?

-Brook Bailey

There is a difference between making something foolproof and reducing the number of fools

-Bill Barns