Let's Make SIP Geographic Redundancy Actually Work.

Mark R Lindsey
mark@ecg.co

This Talk Could Save Your Life.

- VoIP is common for lifeline emergency services
- US FCC Rules, February 2012: You must report a 30-minute 911-affecting outage to the Feds.
- Any one site can be affected by environmental problem, *threat* of terrorism, network routing misconfiguration

Overview

- SIP Geographic Redundancy Stinks.
- The standards don’t matter if the vendors won’t implement them.
- We need some fully-defined mechanisms to support geographic fault-tolerance in *real-world* customer access networks.
- Every failover problem until this is resolved is a gift from your endpoint vendors.

SIP Has Helpful Mechanisms

- RFC 3263 (Locating SIP Servers, Rosenberg &al. 2002)
  - Explains how to identify alternate SIP Servers
  - Basic ground rules for transactions
- Basic SIP has mechanisms for redundant transactions, and handling duplicates

SIP Relies on Smart Endpoints

- Easy to concentrate all effort into the core servers and SBCs
- But the SIP Phones, IADs, ATAs have much of the responsibility
- To make SIP failover work properly, SIP CPE Endpoints must do it properly

Failover Requirements Unclear

- Key questions remain unanswered for SIP Phones & IADs:
  1. When should a phone re-REGISTER with the secondary?
  2. When should the phone re-attempt the primary?
  3. What happens to subscription state?
  4. What happens to the calls after the primary SBC fails?

Today: Sunny Day Scenario

- Customers use NAT Device to REGISTER with SBC
- SBC stores NAT external IP/Port for
- Core SIP Server stores SBC core IP/port

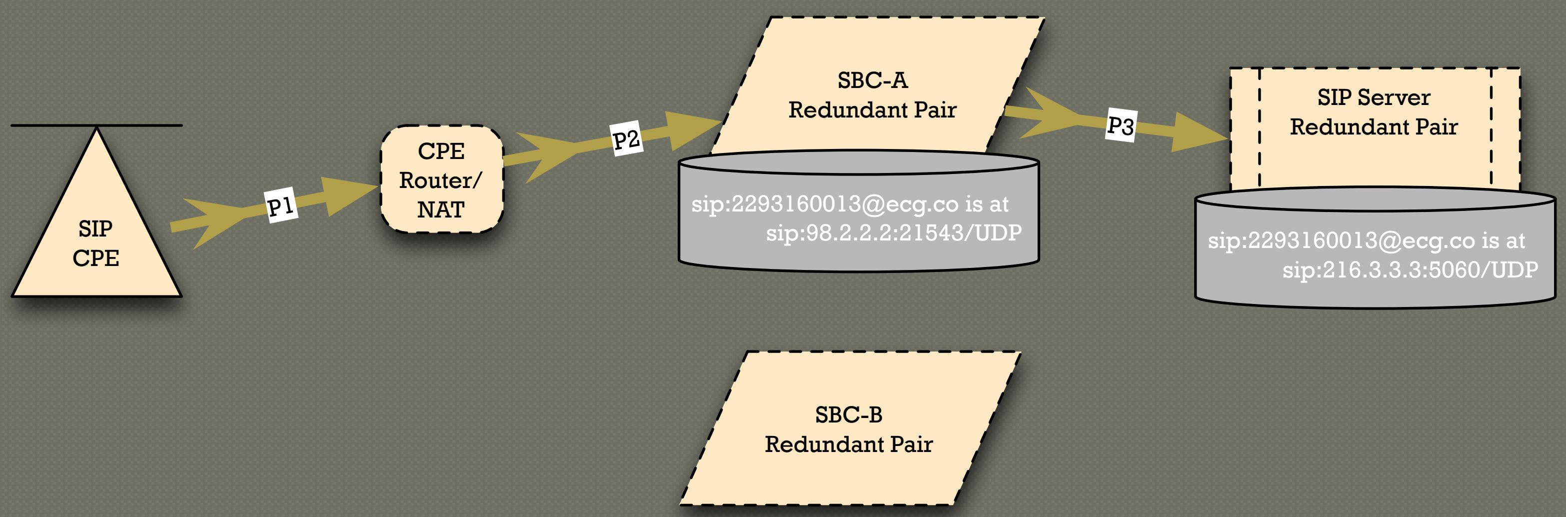

Geo-Redundant Networks = Multiple Access SIP IP Addresses

- We assume that the two sites fail independently (of course!)
- Most designers accomplish this with separate SBC IP addresses for each of the two sites. (i.e., VRRP not appropriate)
- Some systems (BroadWorks) use separate Core Server IP addresses, for each site

Today: Primary Site Fault

- SIP CPE Detects Site A Fault and REGISTERs with SBC-B
- SBC-B learns NAT external IP/port
- Core Server stores new Contact from SBC-B

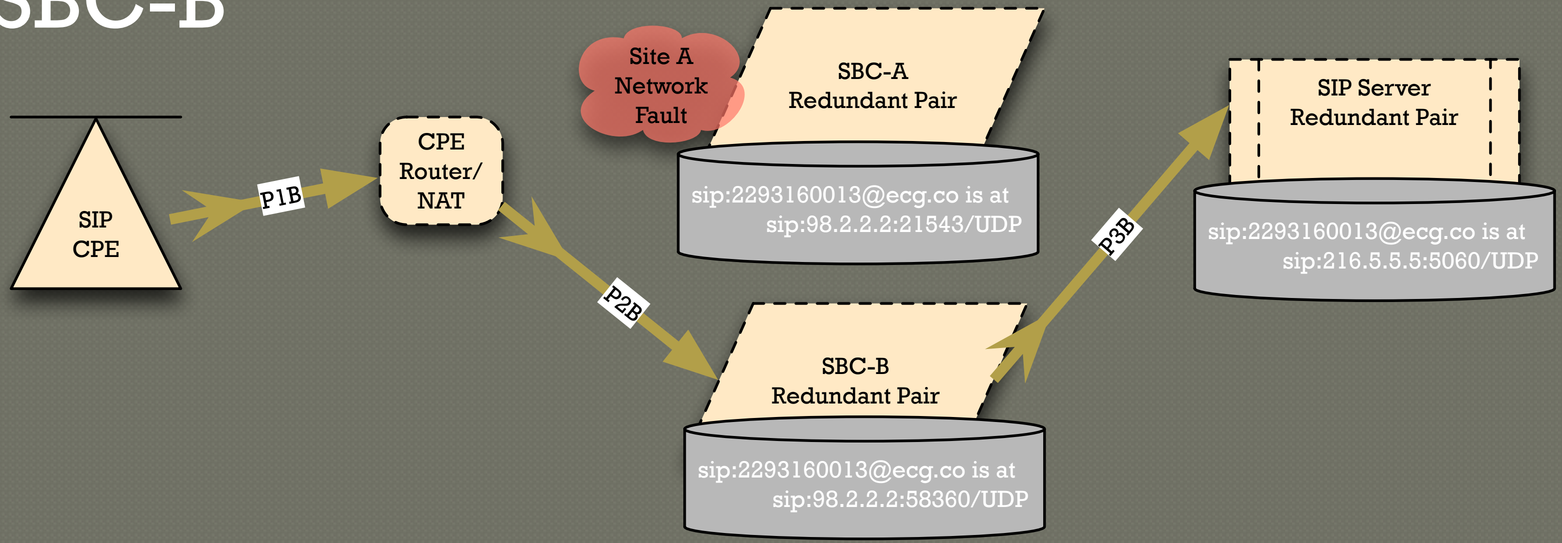

Today: Ugly Questions

- How long did it take for SIP Phones to detect the fault?
- How does the SBC / Core SIP Server handle the registration avalanche?
- How long was the outage?
  (i.e., What's the worst-case time to restore service for any one device?)
- When should the SIP phones re-REGISTER with the primary?

Today: Ugly Truth

- SIP Phones detect fault based on Registration expiration; typically <90 seconds
- SBC and Core server often drop some SIP messages to handle the re-registration storm
- Outages in the ballpark of 30 minutes
- SIP Phones vary widely in re-registration behavior

Could a smarter SBC help?

- Could you just replicate the learned NAT contacts from SBC-A to SBC-B?
- If so, complexity can be handled in Core/ SBC

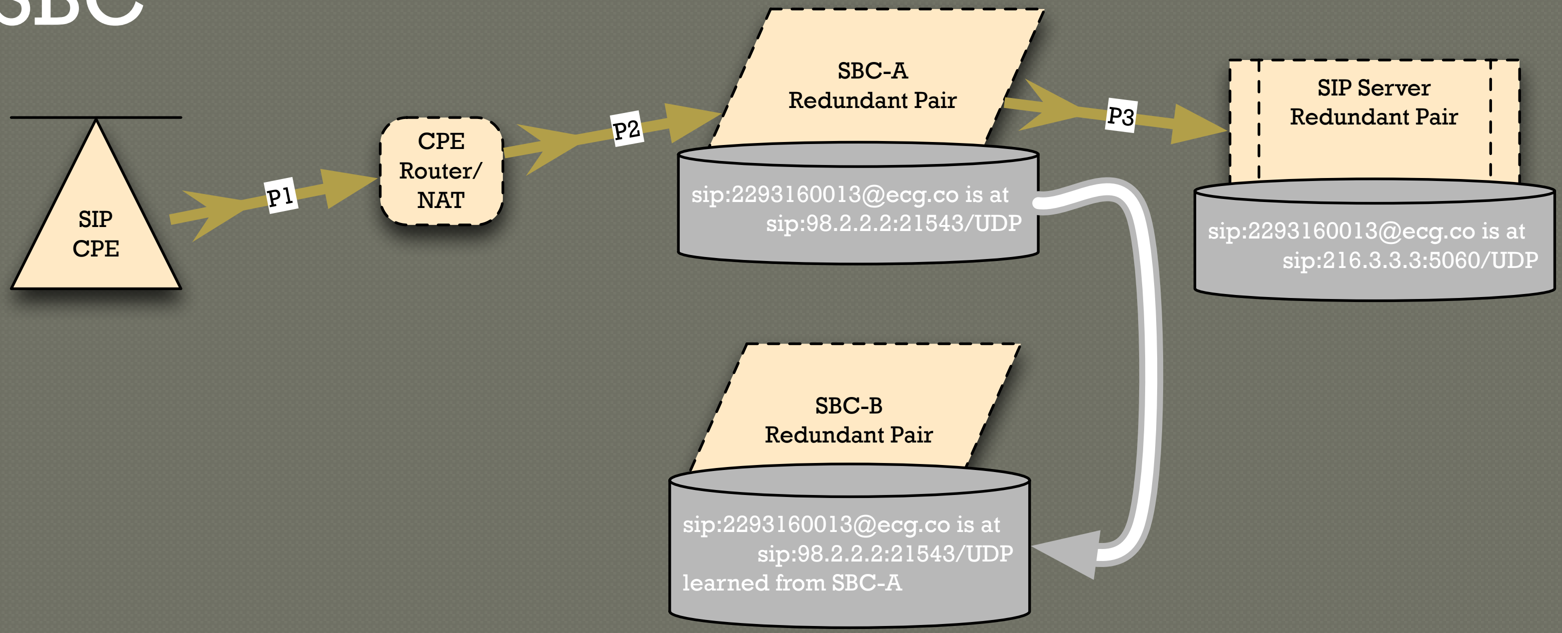

Replicating NAT Contacts Fails

- *Reminder:* Registration from SIP Phone opens path from (a) Core server, (b) via SBC, (c) via NAT device, to (d) phone
- NAT device or Firewall *should* block traffic from SBC-B.

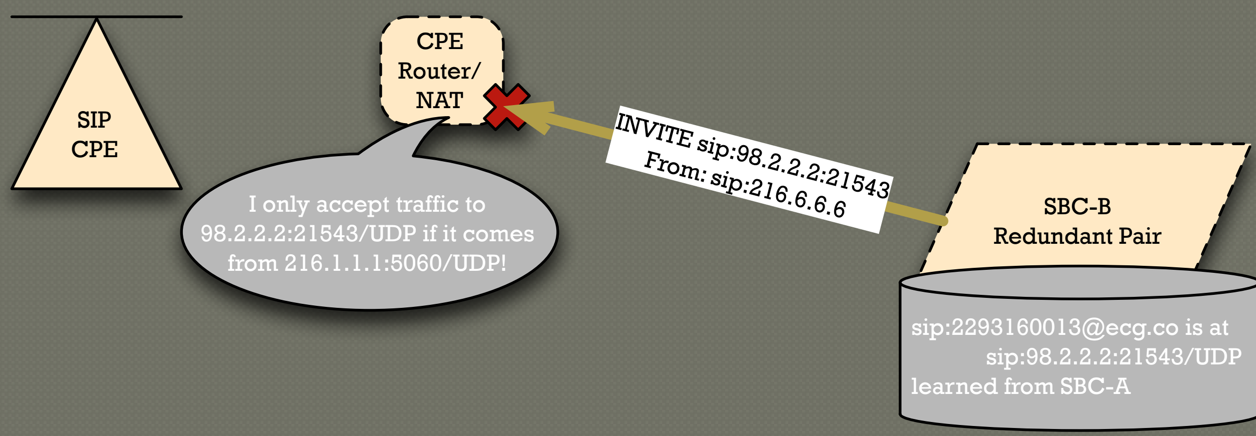

Parallel Registration Redundancy

- Proposal: SIP-PRR, a common-Sense mechanism for improved redundancy
- Key idea:
  - ALL users
  - Register with ALL sites (SBCs)
  - ALL the time
- Core system ALWAYS has a path back to the SIP endpoint

SIP-PRR Clarifies Failover

- SIP-PRR Premise: All SIP users are registered Continuously through Multiple NAT-SBC Paths
- Core SIP Server knows about all paths
- Therefore Core SIP Server can send SIP through any path
- The job of the SIP phone is to maintain registration with all possible paths

SIP-PRR: Sunny Day Scenario

- SIP CPE registers through all paths
- Therefore, NAT device is ready to receive calls from all SBCs
- Core SIP Server has all contacts concurrently

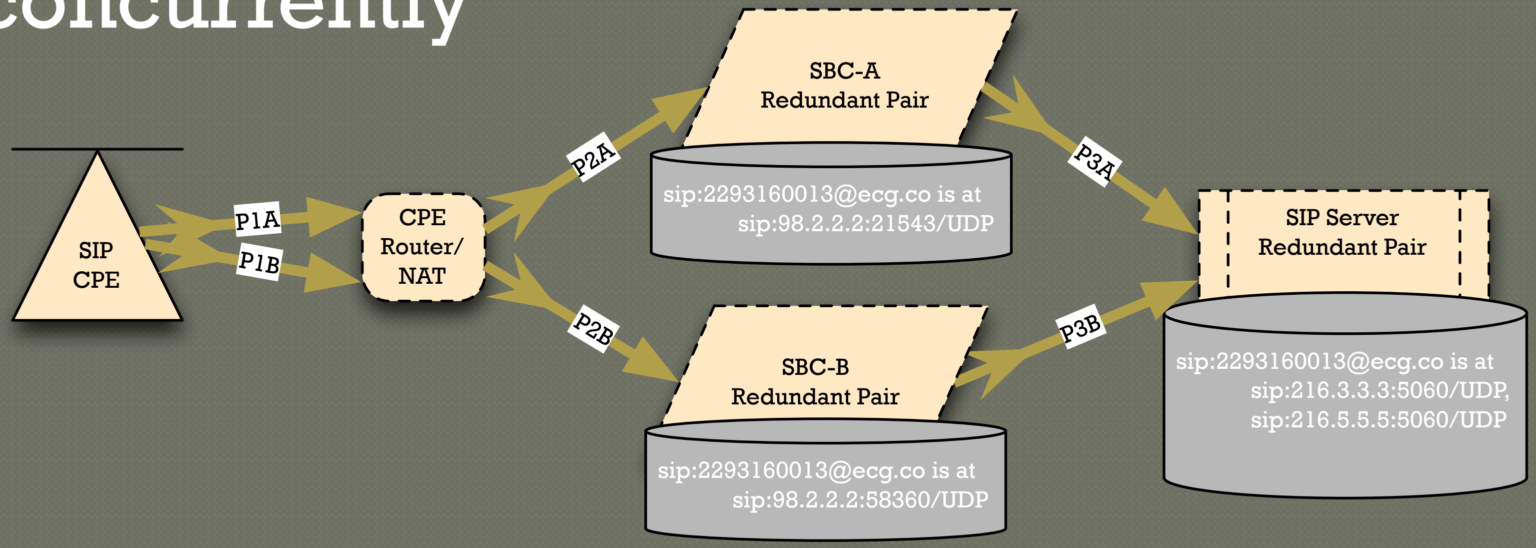

Fault Under SIP-PRR Primary Mode

- Core SIP Server uses one path as "Primary"
- SIP Server attempts Primary path first; then after timeout attempts other paths

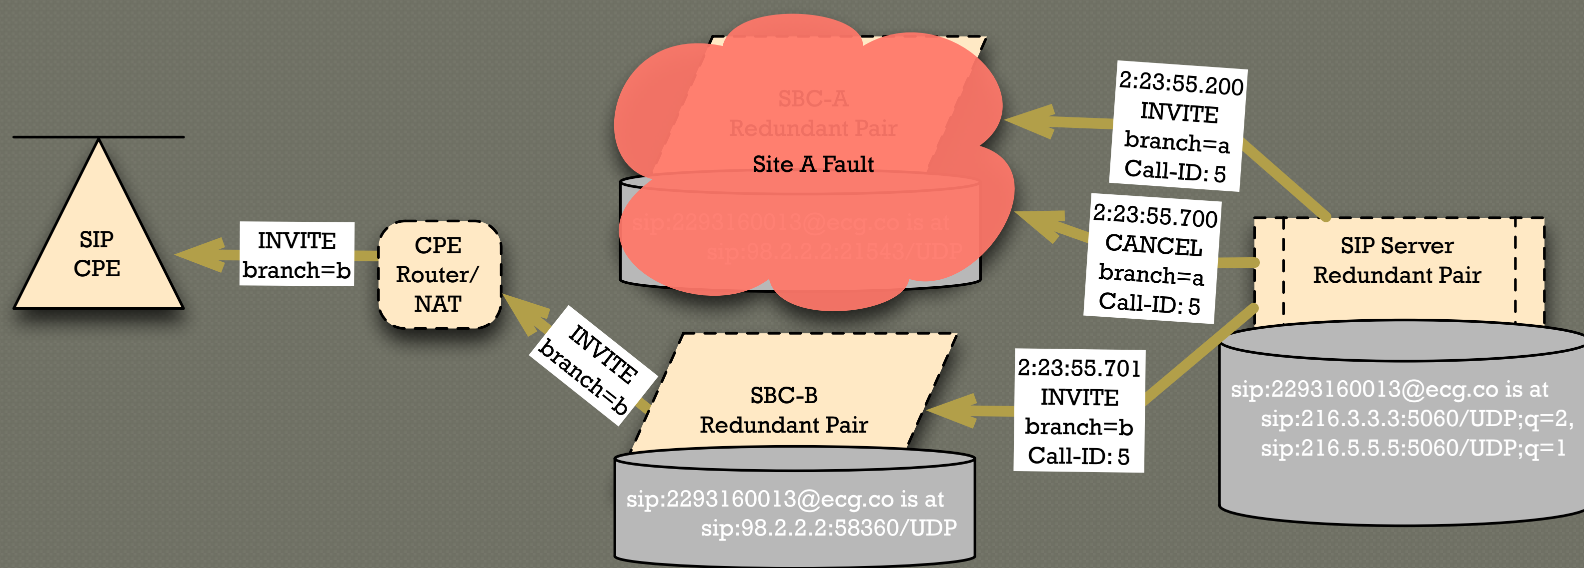

Fault Under SIP-PRR Simultaneous Mode

- All requests are sent through all paths simultaneously
- No failover time at all

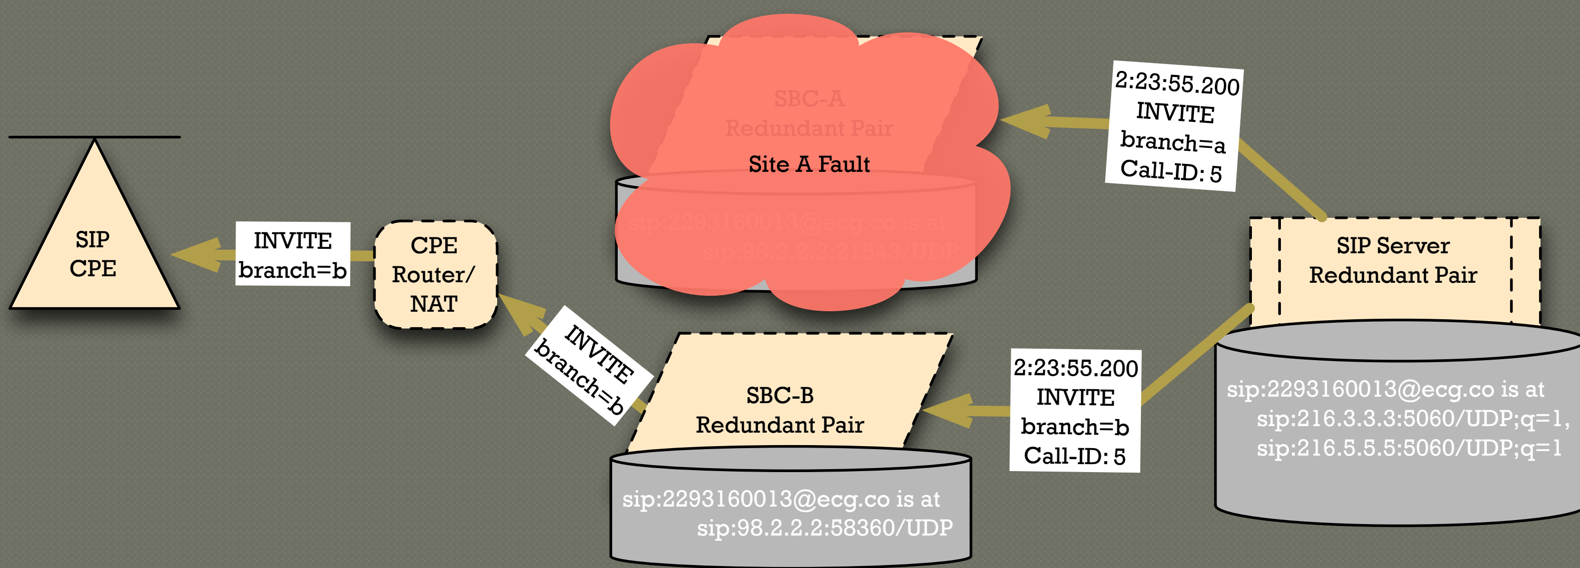

Standing on the Shoulders of Jonathan and Henning?

- This talk does not show all important details
- We're just asking vendors for a specific form of SIP Forking *to actually be implemented!*
- Could SIP-PRR be used to recover a call that was originally established via a different site? Yes!

SIP-PRR: Network Load

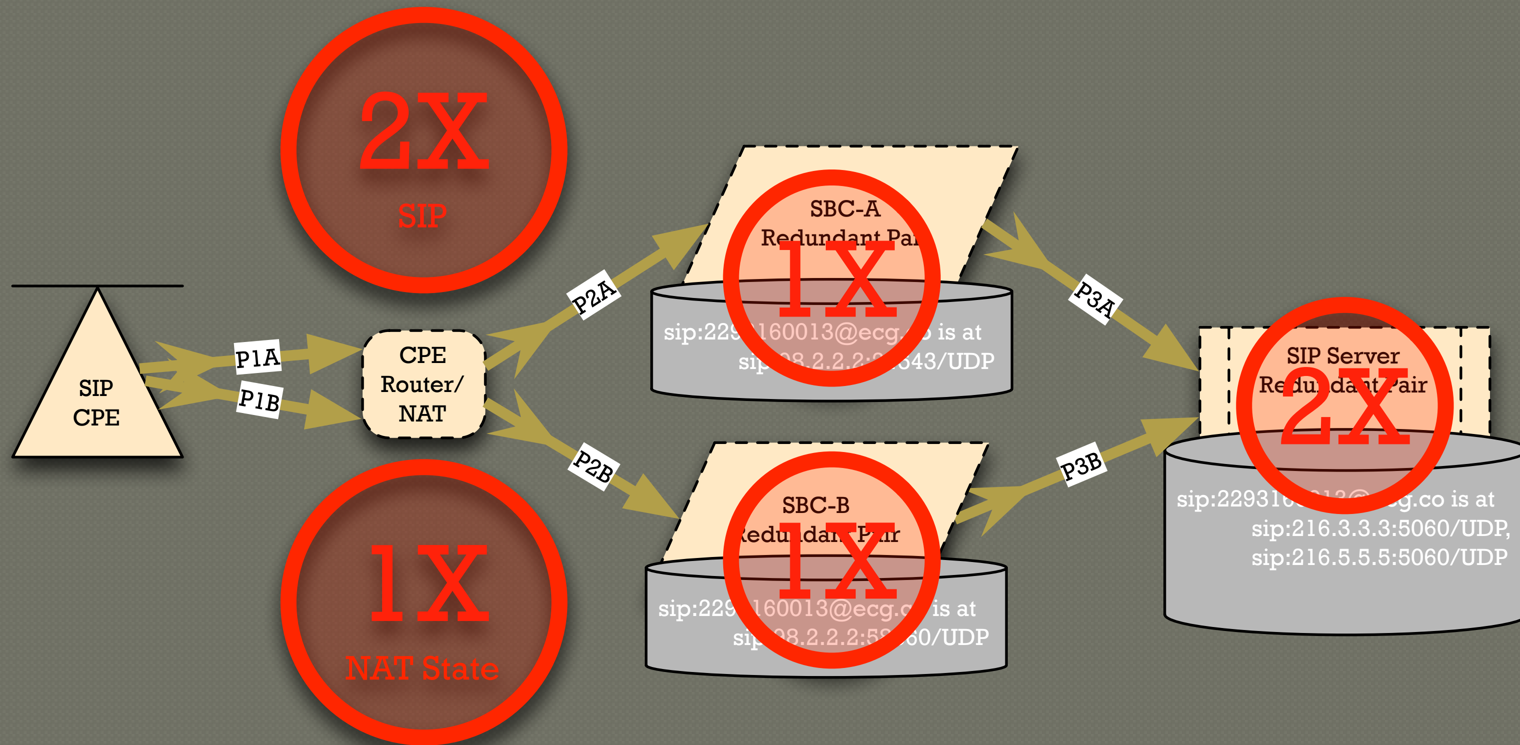

SIP-PRR: Network Load

- CPE Router:
  - Double the SIP Traffic for REGISTER and potentially for all, under SIP-PRR Simultaneous Mode
  - Unchanged State: *for today's standby redundancy, the NAT device must support SIP flows to each SBC*
- SBC:
  - Same load as today in each SBC
  - *Reduced peak-hour load: no register avalanche!*
- Core VoIP Servers:
  - Double SIP REGISTERs, plus SIP-PRR Simultaneous Mode
  - Double the state to track SIP Contacts (trivial increase – maybe 16 MB RAM for 100k subs)
  - *Reduced peak-hour load: no register avalanche!*

Deployment Considerations

- Likely the biggest impact: double SIP REGISTER traffic at CPE Edge Router (CE)
- After core ready: partial CPE support is OK

SIP-PRR for End Users

- Practically no delay for failover
- Every call works – inbound and outbound
- SIP-PRR makes Geographic Failover actually a substantial benefit work for end users, even if it occurs every day.

SIP-PRR for Network Engineers

- Make better use of all that geo-redundancy investment already incurred!
- SIP-PRR is the missing link for true geographic redundancy
- With SIP-PRR, maintenance / configuration / faults for an SBC pair or entire geographic site do not turn into a major event
- And there was much rejoicing

You Can Make SIP-PRR Reality

- A simple idea that could make an enormous difference for network reliability.
- But we need to flesh out all the details: help us complete the standard proposal.

Software Vendors REALLY Make SIP-PRR Reality

- SIP-PRR will require real work for your vendors:
  - SIP Phones, IADs, ATAs must support it: Polycom, Aastra, Adtran, Linksys, Audiocodes
  - Core SIP servers (but not SBC) must support it: BroadSoft, Metaswitch, Sonus, Taqua
- Vendors will *only* work on this problem if *you* make it important to them.

Contact me

- Mark R Lindsey
  +1-229-316-0013
  mark@ecg.co
  http://ecg.co/lindsey