Hysteresis Histrionics

Mark R Lindsey, Jon Chleboun, James Puckett
{lindsey,jchleboun,jpuckett}@ecg.co

https://ecg.co/sipnoc

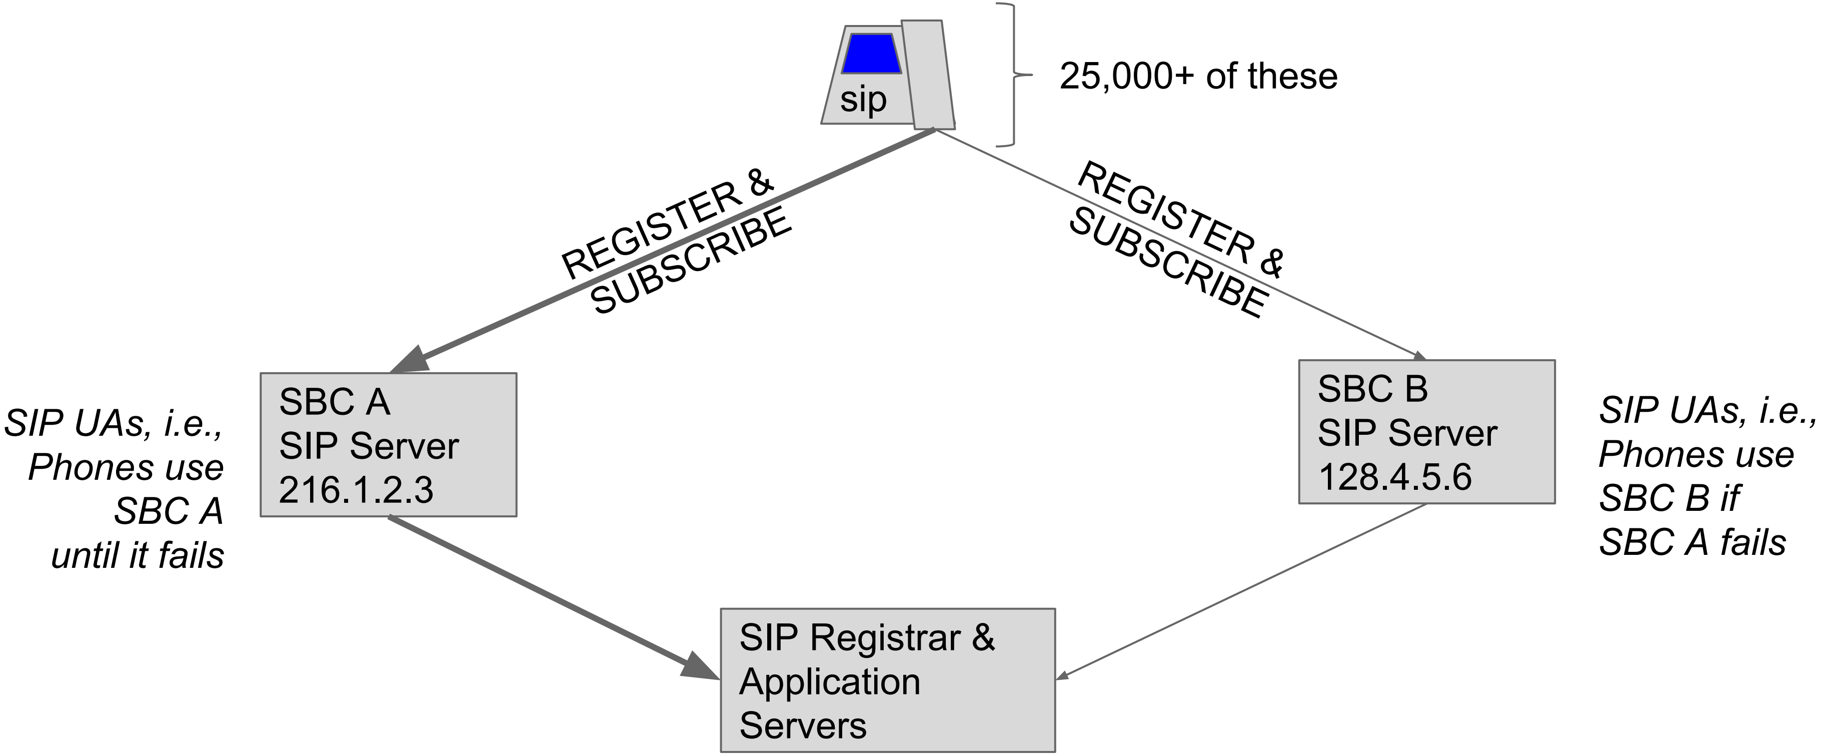

sip

25,000+ of these

REGISTER & SUBSCRIBE

REGISTER & SUBSCRIBE

SBC A
SIP Server
216.1.2.3

SIP UAs, i.e., Phones use SBC A until it fails

SBC B
SIP Server
128.4.5.6

SIP UAs, i.e., Phones use SBC B if SBC A fails

SIP Registrar & Application Servers

Context

When SIP failure occurs: *walk, don't run.*

ARPANET, October 1986, Daily Meltdown: "Congestion Collapse"

ARPANET: prototype of the Internet since 1969.

In the afternoons in 1986, the performance would drop: from 32 kbps to 40 bps net throughput… between machines in the same room.

Even though the network was completely busy sending data, nearly all of it was discarded because a few packets were lost.

1980s: Van Jacobson and the ARPANET

ARPANET nodes were attempting to establish connections *on startup.*

"If a node sent a dozen packets in trying to establish a connection and one didn't get through, it would resend all twelve." *

Jacobson: "It was a combination of really bad startup behavior and poor recovery code. We were just wasting all the bandwidth."

* *Cade Metz, Wired Magazine*

2.01

SIP messages per minute from typical SIP Phone.

Very Efficient!

Each UA averages ~2 SIP requests/min

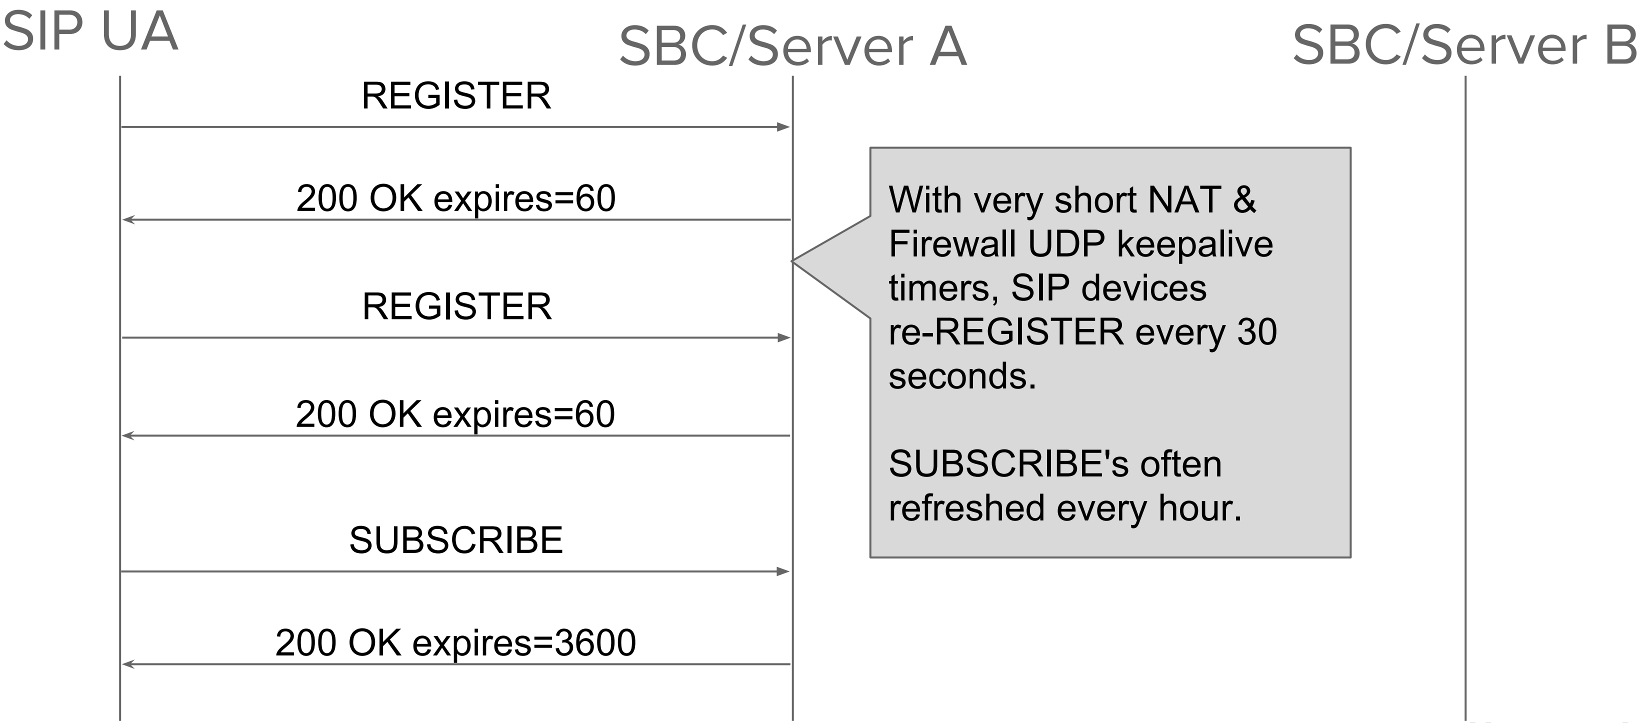

With very short NAT & Firewall UDP keepalive timers, SIP devices re-REGISTER every 30 seconds.

SUBSCRIBE's often refreshed every hour.

One Registration, Many Subscriptions

Many advanced Voice services registration, plus 1 to 8 SIP Subscriptions for features.

Busy-Lamp Field / Line State Monitoring

Message Waiting Indicator

Do-Not-Disturb

Call Forwarding

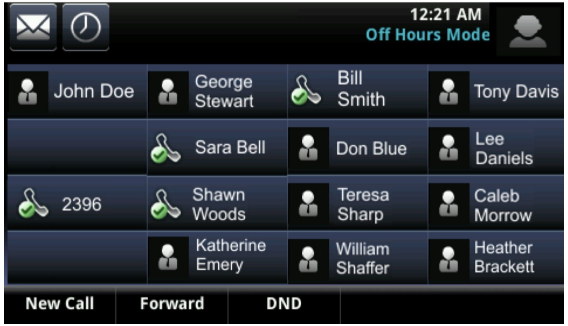

12:21 AM
Off Hours Mode

John Doe	George Stewart	Bill Smith	Tony Davis
	Sara Bell	Don Blue	Lee Daniels
2396	Shawn Woods	Teresa Sharp	Caleb Morrow
	Katherine Emery	William Shaffer	Heather Brackett

One transaction fails? Startup again!

If any with transaction fails (with 503 response from SBC),

1. The SIP UA fails over -- starts up -- on secondary server.
2. Triggers NOTIFY to SIP UA for each subscription.

> If a single SIP REGISTER or SUBSCRIBE transaction fails, *all subscriptions and registrations failover, and start over.*

Failover: burst of 17 SIP requests

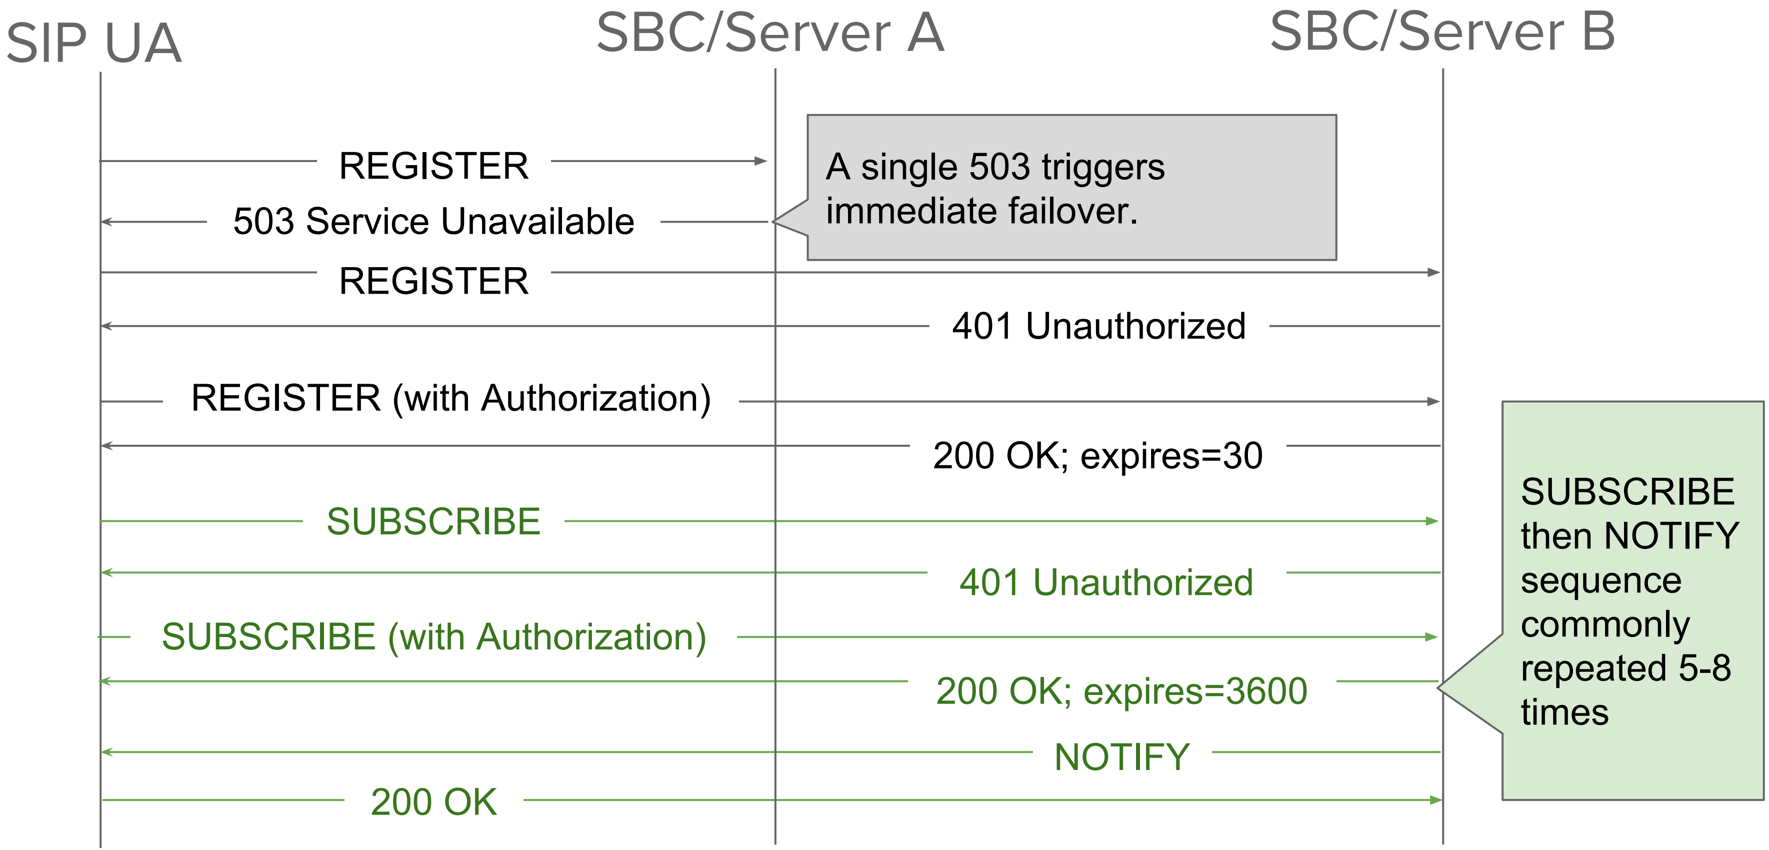

SIP UA

SBC/Server A

SBC/Server B

REGISTER

503 Service Unavailable

A single 503 triggers immediate failover.

REGISTER

401 Unauthorized

REGISTER (with Authorization)

200 OK; expires=30

SUBSCRIBE

401 Unauthorized

SUBSCRIBE (with Authorization)

200 OK; expires=3600

NOTIFY

200 OK

SUBSCRIBE then NOTIFY sequence commonly repeated 5-8 times

8.5x

SIP workload growth for failover over a single minute

Ex. 1: October 2018: Stabilized by *Removing* Failover

In October 2018, US-Nationwide recent failover scenario only stabilized by removing the failover configuration directly.

> *The workaround employed: Secondary SBC SRV records were removed from DNS.*

Ex. 1: OUTAGE: Failover outage, 2.5x recorded SIP load

SBCs could not handle failover.

Resolved by removing backup SBCs from DNS

Approx 2 hours to recover after manual intervention.

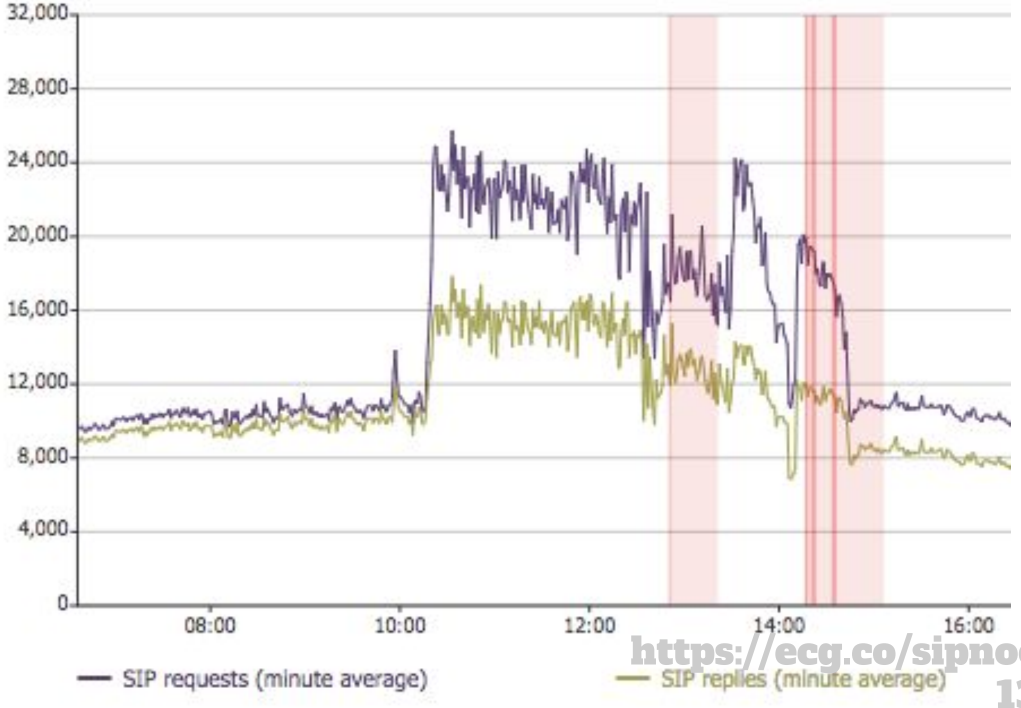

https://ecg.co/sipnoc

Ex. 2: 3.5x CPU Load on Startup

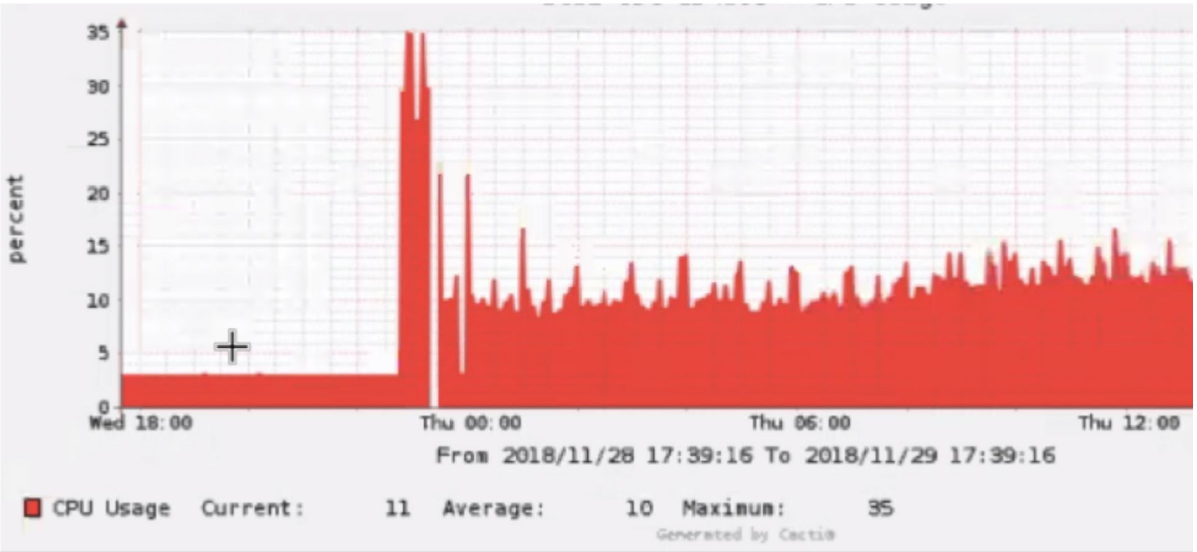

40,000 subscribers; Capacity sufficient. ~10% CPU load normal for maintenance. 1-2 Subscriptions/user.

Startup Consumes Huge Capacity

Regular maintenance is only a few packets.

When devices failover, UA's REGISTER & SUBSCRIBE as if they were entirely new to the network.

Backup servers don't trust the *same* devices from the *same* sources because they don't share Contact database, and don't reuse Authorization header data.

Ex 3: Failover 8.5x Authentication

Subscriber location is not usually shared between SBCs.

So failover causes a burst of authentication activity.

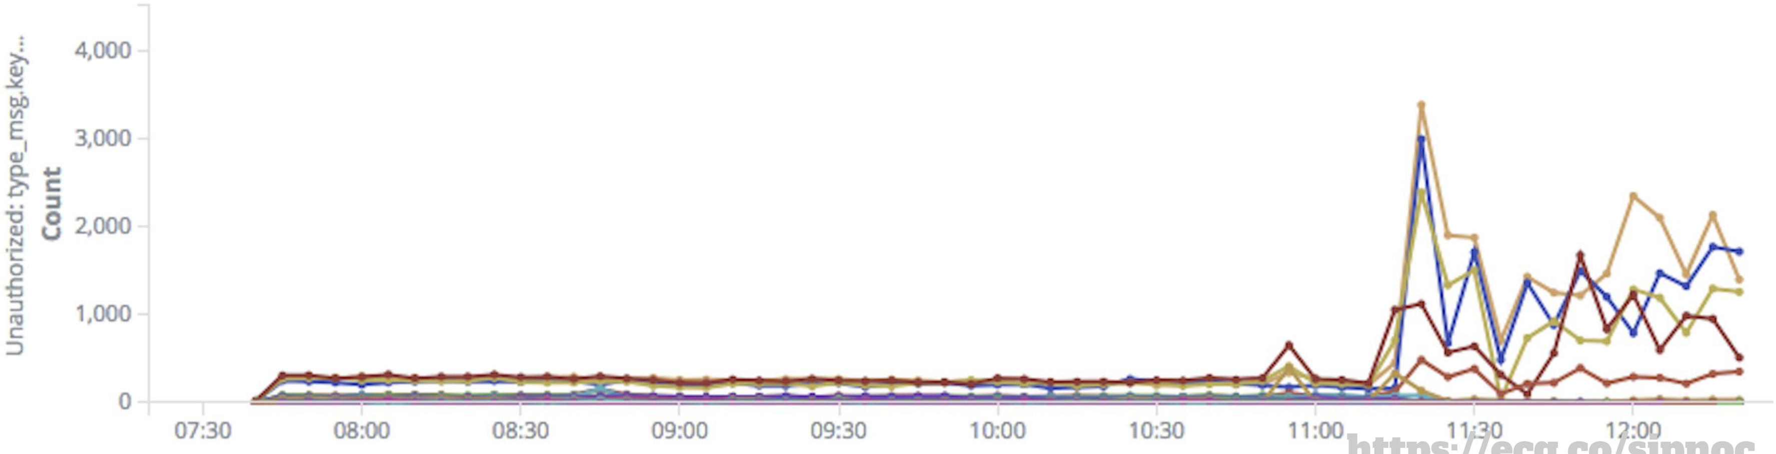

SIP UA's should retry failed requests

Today, SIP UAs *abandon the specific server* at the first 503 to a request.

Necessary?

> ***No.*** The 503 doesn't invalidate the Subscription.
>
> UA should retry later.

"If at first you don't succeed,

Backoff Exponentially."

Dan Sandler in

Google "Site Reliability Engineering"

Use *Entire* Failover Budget

Some Service Providers love with idea of instant recovery.
...but instant recovery is expensive!

Do you really need to failover
within 10 seconds? 30 seconds?

Many networks can endure minutes of failover with practically no customer impact.

Chains of Failover can crash whole networks

Suppose you have a fault on Server A that causes an overload on B...

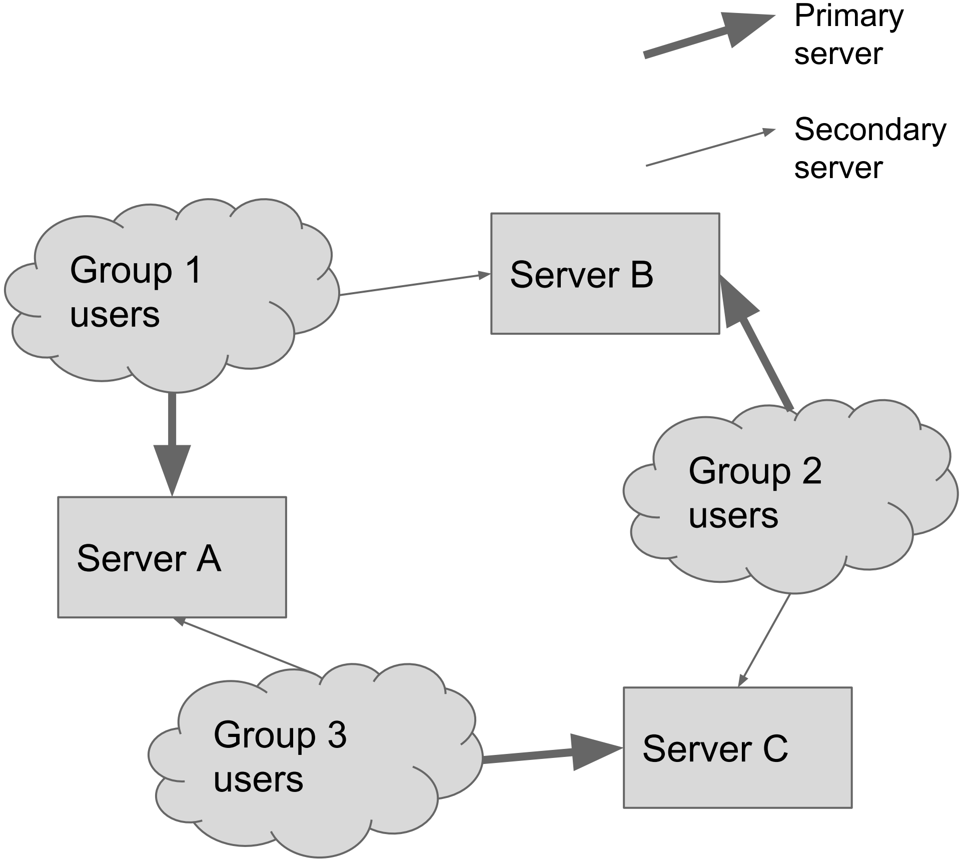

RFC 3263 says how to failover, not when

RFC 3263 on "Locating SIP Servers" grants some flexibility:

> 4.3 Details of RFC 2782 Process
>
> If a failure occurs, the client SHOULD create a new request ... That request is sent to the next element in the list as specified by RFC 2782.

***Maximize refresh delay:* Slow re-registration**

Retry with Exponential Backoff on 503 Failure.

Randomize delay on failover.

ECG: Voice that works.
All the time.

mark@ecg.co
+1-229-316-0013
https://www.ecg.co