



TELECOM RELIABILITY ISSUES, AS APPLIED TO 9-1-1

Brian Rosen

Principal

Brian Rosen Technologies, LLC



Who am I and why should you care?

- I'm the tech weenie here
- Systems architect on high availability systems
- Principal Architect of Next Generation 9-1-1
 - *Co-chair of NENA i3 architecture working group*
- 30+ year background in telecom and IP networking
- Consultant to State government on
 - *NG9-1-1 RFPs, vendor selection, contract monitoring*
 - *9-1-1 Outage investigations*

Disclaimer

- Personal opinions, not necessarily backed up by a whole lot of facts
 - *But its my job to understand how reliable systems are built, and what goes wrong with them*
 - *Been around a long time, like when there was "Ma Bell". Worked for Tier 1 suppliers to telcos.*
 - *I've been spending a fair amount of time lately looking deeply into failures of the 911 system*
 - It's not pretty

The landscape today in telecom

- Plain Old Telephone systems ("wireline") has close to zero investment, but still a viable business, and the carrier of last resort for many rural citizens
- Wireless is the flash, big money maker, large investment, if you are ATT, TM or VZW
- IP networking and its various delivery platforms, including fiber to the curb are healthy businesses
- Underlying all of these are the optical networks that actually carry all the traffic

A primer on reliability

- What we actually care most about is “availability” = percentage of time a system is available to serve its customers. Usually this is expressed in “nines” availability %
Downtime per year
 - 90% - *"one nine"* 36.5 days/yr downtime
 - 99% - *"two nines"* 3.65 days/yr downtime
 - 99.9% - *"three nines"* 8.76 hours/yr downtime
 - 99.99% - *"four nines"* 52.56 minutes/yr downtime
 - 99.999% - *"five nines"* 5.26 minutes/yr downtime
- Availability can be measured (on a very mature, stable system) or predicted based on statistics for the components of the system and the arrangement of those components
- Predictions are based on “Mean Time Between Failures” and “Mean Time To Repair”

What we see happening (specifically in 9-1-1)

- Systems using somewhat older tech are failing more often, and taking longer to repair
 - *One reason for this is that telcos have lost their ability to track how the networks are actually deployed*
- Newer systems are failing mostly by either fiber cuts or software bugs
 - *Fiber cuts are more common now, probably because the construction industry is better at digging holes*
 - *Software bugs are really hard on availability: to raise availability, you install redundant elements, so if one fails, another is available to take the load. But software bugs affect all the replicas, so that entire mechanism doesn't work.*

Some issues are just poor execution

- We have seen a significant number of failures of the 911 system that were entirely avoidable “dumb” people mistakes
 - *My current worst case was a single fiber cut brought down a significant part of the Nebraska 911 system because an earlier cut resulted in a temporary work around which persisted for **11 years!***
- We see examples of insufficient redundancy that is completely obvious but either someone thought it was too expensive to do it right, or no one verified that the redundant path/circuit/equipment was actually ready to serve if needed, and of course it wasn't
- We also see somewhat deliberate poor design where the carrier is arguing the regulations that otherwise would require a better design don't apply

We have some regulatory issues

- Next Generation 9-1-1 is substantially different from E9-1-1
 - *All new protocols*
 - *All new databases*
 - *All new functions*
- Regulations at both state and FCC level were written for E911 and need to be updated
- Carriers are using their own interpretation of the regs, and it's not great
- Latest FCC R&O on NG9-1-1 is a great start, but more needed

There are some bright spots (again, in 9-1-1)

- Internet Protocol networks really are inherently more reliable than older technology.
- One reason is that an IP network will find a path, if it exists, between two points even if there is major damage to the network
- Older technology relies on pre-determined alternate routes
- The latest standards incorporate more/easier ways to make redundancy easier to deploy and manage
- We're being somewhat successful at getting vendors to show their system designs can meet a 5 nines predicted availability

One of my latest attempts to improve

- The next version of the base technical standard for NG9-1-1 describes a mechanism for one state to back up another state's NG9-1-1 system
- If the entire state system goes down, the backup system takes over
- This is specifically aimed at dealing with software bugs. The two systems don't have to, and should not, use the same software
- But the mechanism is (currently) optional, so we have to convince RFP writers to require it, and insist vendors implement and deploy it

When Failures Occur, Call in the Experts

- Shameless plug for consultants like me ☹️
- Don't rely on what the utility tells you what happened
- Its usually truthful, but not the whole truth
- You need to ask the right questions and analyze the response to understand what really happened, and what needs to be done to avoid recurrence
- I'm big on the true root cause
 - *Usually, what you get is a symptom of a bigger problem*
 - *We want to avoid the whole class of problems, not just this specific one*
 - *The Five Whys*