computers are bad

You are receiving this facsimile because you signed up for fax delivery of this newsletter. To stop delivery, contact Computers Are Bad by email or fax.

https://computer.rip - me@computer.rip - fax: +1 (505) 926-5492

2023-02-07 secret government telephone numbers

Very nearly a year ago, I wrote a popular article about secret military telephone buttons. To be clear, the "secret" here was a joke and these buttons are in fact well documented. The buttons I was talking about were the AUTOVON call precedence buttons, used for a five-level prioritization scheme within the AUTOVON military telephone network. The labels on these buttons, FO, F, I, P, for Flash Override, Flash, Immediate, and Priority, directly reflected the nuclear C2 scheme at the time. The AUTOVON telephone network is long retired, but military telephone systems continue to provide a call precedence scheme today, admittedly usually without dedicated buttons.

Well, the idea of call precedence without the priority buttons leads naturally to a followup that I promised: government and defense call prioritization schemes on the general, civilian telephone network. It has long been recognized that in the event of a national emergency, many people involved in the response would not have access to a dedicated government telephone network. This is particularly true when you view civil defense as a wider remit, beyond just military reprisal. Recovery from a disaster of any type will involve federal, state, and local government leaders, as well as staff of response organizations like critical utilities, hospitals, and disaster relief organizations. Not all of these people can realistically be furnished with a phone on a dedicated network. The only way to practically ensure prioritized communications in a disaster is to provide that capability as a feature of the public switched telephone network.

It is perhaps useful to provide a little bit of background on the need for call prioritization. In the landline phone network, the major limitation on capacity is long-distance trunks. Particularly in the era before the TDM digital telephone network, long-distance trunk capacity was relatively expensive. In the 1950s, a small city might have only one or two dozen trunks for outbound long distance calls. This means that only one or two dozen simultaneous calls outside of the city could be connected. By the mid-1960s, rapid expansion of the Long Lines microwave network had dramatically increased long-distance trunk capacity (in part through added routing flexibility), so it was no longer necessary to request most long-distance calls in advance. Still, into the ’90s it was not unusual for long-distance calls during peak periods like Christmas to result in an "all trunks are busy" intercept message.

The need for some sort of emergency prioritization of traffic on the landline network has long been known, but solutions have been uneven, at least in the US. In the UK, where telephone service was a state enterprise until 1984 and the first and second World Wars created a more expedient need, some type of basic telephone prioritization scheme has long been in place. Through WWII and the Cold War it was a simple one: in cases of emergency, all local loops not flagged as being required for emergency service would be

disconnected. This had the dual benefit of freeing capacity for government traffic and denying an invading enemy the use of the telephone network. A somewhat more sophisticated version of the same idea, called the Government Telephone Preference Scheme, remained in service until 2017.

In the US, though, there was little effort towards such a system during WWII and the Cold War. It's hard to know what exactly to make of this situation. First, the situation has always been a little different in the US than in the UK: during WWII, Great Britain was the target of ongoing bombing campaigns. Excepting mostly minor terrorist attacks (such as during the Escobar Rebellion in Mexico), the US has not suffered a military assault on its homeland since the Civil War. During the Cold War era, an attack on the US was assumed to be a nuclear one, and most likely an all-out nuclear campaign. The end result is that civil defense in the UK tended to be viewed as a practical system oriented around sustaining operations during an attack, while civil defense in the US was viewed as more of a theoretical exercise in preparation for reconstruction. This is all just background to explain that part of the reason the US did not have a telephone prioritization scheme through most of the Cold War is simply because the US is historically pretty bad at deploying civil defense infrastructure.

But that's not the whole story, there are other, more positive, reasons as well. For one, since WWII AT&T had operated a dedicated telephone network for military use (known as AUTOVON for most of the Cold War) and it was relatively large in scope and well-hardened against attack. For the most part, the assumption was that all emergency traffic would be on AUTOVON, not the PSTN. Another factor is the government's close relationship with AT&T. Despite AT&T being a private entity, it was exceptionally close to the federal government and provided many government services in secret. It is possible, even likely, that AT&T had arranged for some sort of emergency call prioritization scheme that wasn't discussed in public.

The cellular network faces similar problems but at a different point. In the landline network, it is possible but rare for a local exchange switch to become overwhelmed by the number of phones off-hook. Conventional telephone switches use a relatively low "oversubscription" ratio from local loops to call-processing elements (the nature of which depends on the switch architecture) and are often designed so that most components of the switch are only in use during the setup stage of the call, and are disconnected once the call is established and made available to process additional calls. This means that overwhelming a telephone exchange to the point that it can no longer offer dial tone and accept dialing from telephones requires picking up a significant number of the connected telephones at the same time. In practice, the long distance trunks were virtually always the bottleneck. The cellular network is not so lucky.

Before data became the major driver of cellular network architecture, cellular carriers used a much higher oversubscription ratio. Cell phones spend only a very small portion of the time connected to calls, less than landlines early on due to the higher pricing, and so cellular base stations were designed to track associations with far more phones than they could actually handle call traffic with. In an emergency that affected even a small area, it was very easy for enough people to make cellphone calls that the tower hit its capacity limit and began rejecting additional calls. Since the limitation was often in the actual bandwidth on the radio side, this limit applied to all simultaneous calls, not just calls in setup. Some freeway accidents were dramatic enough emergencies to cause cellular calls to fail.

This problem, in the case of civil emergencies, is not at all far-fetched. We will refer to 9/11 repeatedly because it is main origin of modern emergency telecommunications programs. On 9/11, Verizon reported a 1.5x to 2x increase in national call volume.

Cingular a 20% increase. Nextel, a carrier that I will write a whole article on eventually because of the interesting technology they employed, saw an increase of several hundred percent in the hours immediately following the first attack. These statistics, though, reported at a national level, understate the impact. There was a problem nationwide, but there was a far greater problem in New York City, where local cellular base stations became so overwhelmed that the call completion rate in the Manhattan area is thought to have dropped to a few percent for some time. Normal cellular service to Manhattan was not restored for several days.

And during that same time period, basically the entire government and commercial disaster response capability was attempting to communicate plans... via cell phone. 9/11 is currently the most memorable event demonstrating the need for a call prioritization system, but it is only one on a string of incidents both before and after where emergency response was significantly impeded by capacity limitations in the PSTN.

As you would expect from a federal emergency preparedness initiative, the government set out to address the problem of limited telephone network capacity through a confusing maze of several separate programs, administered by separate organizations, and accessed by separate means. Without a clear place to start, I will choose the best-known of these services: GETS, the Government Emergency Telecommunications Service.

The history of GETS is surprisingly obscure, perhaps because it has been tossed like a hot potato between a good half dozen executive branch agencies over the last few decades. Here's what I have put together from various archives and a helpful former AT&T source: the story of GETS starts with the National Communications System, or NCS, established by Kennedy in response to the Cuban Missile Crisis. NCS had an objective of ensuring reliable communications for the civilian government during nuclear conflict, and seems to have started down the path of developing a civilian version of AUTOVON. This would be essentially an independent, parallel telephone network employing dedicated resources engineered for survivability.

This program, called the National Emergency Telecommunications System or NETS at least late in its life, faltered mostly due to cost. AUTOVON had been a tremendously expensive project that went through mostly because of the military's tremendous budget. To save money, NETS was re-imagined as an "overlay" network built on top of the PSTN. Conveniently, this later version of the proposed NETS is detailed in a 1987 National Research Council report evaluating the plan on the request of NCS.

The 20,000 federal government users of NETS would each be issued an Access Security Device, which would use some sort of digital signaling to authenticate itself to a device called a Call Controller integrated into (or attached as a peripheral to) a telephone switch. The NETS Call Controller would then communicate with other call controllers to attempt to establish a route to the dialed number. This is the real magic of NETS: the telephone network at the time, organized largely around hierarchical long-distance routing, had fixed, pre-planned routes in place between tandem switches. If the configured route to a destination was damaged, the call would fail---there was no self healing the sense we expect from modern networks. NETS implemented self-healing as a feature of the CC, which would initially attempt to establish the call via a default route, check for the ability to communicate with a CC at the other end, and give up and try a different route if the remote CC didn't respond. This way NETS CCs should be able to "discover" a working route even with appreciable damage to the telephone network.

NETS, according to the National Research Council report, had numerous shortcomings. For one, the proposal dealt almost purely with federal employees, even though it was clear from NCS's objectives that many non-government users would need access in an emergency

(hospitals and infrastructure operators, for example). This was an especially large problem because of the seemingly over-paranoid design of the "access security device," a piece of hardware which would apparently have to be stockpiled in various federal offices to be issued in the event of an attack. The report recommends eliminating the ASD and replacing it with a dialed PIN code, a far simpler arrangement that didn't seem much worse considering that the ASD didn't implement encryption of the call contents but instead just a challenge-response authentication process.

Further, unlike AUTOVON which provided dedicated service all the way to the user, NETS was essentially a feature of the long-distance network only. In an emergency when a local exchange switch was overwhelmed and unable to service a newly off-hook like, NETS would be of no hope. Still, the report finds that this is a relatively small problem, and although users might have to wait a while it is expected that after taking their phone off-hook they would get a dial tone eventually. Because the telephone switches of the time worked by fabric-switching digit receivers to off-hook phone lines, each person picking up a phone was essentially placed in a queue to receive dial tone. Normally the queue was empty and so this happened instantaneously, but during especially busy calling periods it was not all that unusual for their to be a very brief delay.

Finally, the NETS proposal included not only voice but fax and data. Because the "damage-avoidance" routes developed by the NETS CCs would often be longer than typical long-distance phone routes, the connections would be of relatively poor quality. This required additional electronics at the ends of the call, and potentially integrated into user PBXs, for line conditioning. It seems like the use-case for the data mode wasn't very well established, and so the report recommended dropping the requirement for additional line conditioning.

The conclusions of the report seem to have been influential, because most of the features the report identified as problematic disappeared from later consideration. Unfortunately, some of the beneficial features did as well. NETS had included a call preemption system, for example, which did not make it past the end of the NETS project. The use of additional dedicated controllers for routing also seems to have been abandoned, but on this issue we must consider that in the late '80s computer-controlled switches were becoming the norm and were capable of much more complex routing logic. This may have just eliminated the need for an additional control system dedicated to routing emergency calls.

The evolution of NETS, now called GETS, launched in 1994. Contemporaneous reporting supports the odd lack of historical information on GETS. A 1996 Chicago Tribune article about the 710 area code (see, people have been bringing up 710 as a weird area code for a long time!) says that while GETS was never a secret, it also received very little promotion even within the government. As of that article, the director of the NCS said that GETS would be "fully operational" in 2001... an auspicious date in emergency telecommunications.

Indeed, the events of September 2001 lead to a massive federal reconsideration of continuity of government, emergency communications, and emergency preparedness in general. To make a long story short, the attacks of 9/11 immediately prompted an almost complete failure of federal emergency communications systems. Everything from complex continuity of government plans to the phones in the White House basement were found to be completely nonfunctional, the result of decades of under-investment in an increasingly incompetent civil defense apparatus [1]. In the year after 9/11, federal bureaucrats found themselves blowing the metaphorical dust off of a number of half-finished or half-forgotten communications programs, GETS included. It seems that in the 2001-2002 time period, NCS launched sort of a public relations program to promote GETS not just in

the federal government but across state and local governments and industry. While GETS is technically a Gulf War-era program, for most purposes it is a post-9/11 program.

GETS is perennially discussed in telephone arcana circles, including *Computers Are Bad,* because it has the distinction of being the sole use of the "federal government" 710 area code. GETS is accessed by dialing 1-710-NCS-GETS, where NCS is left over from the former National Communications Service. Amusingly, the special nature of the 710 area code poses a challenge: 710 is so rarely used (to the degree that it is sometimes listed as a "reserved" or unused area code) that some phone systems with trunk routing logic, like larger corporate PBXs, will not route it correctly.

The decision to put GETS in the little-used (at the time, entirely unused) 710 area code was probably more practical than whimsical. GETS is expected to work across the entirety of the landline telephone network, which in 1994 contained a lot of legacy switching systems. Trunk selection based on area code would have been nearly universal at the time (I say "nearly" because dial service was not yet universal in 1994), so it would have been relatively easy to configure older switches to direct calls to the 710 area code via dedicated trunks. Basically, the prioritization of GETS could be implemented as a special case in the "LERG" routing table used to connect long-distance calls. On the other hand, as a workaround for phone systems that do not handle the 710 area code correctly, GETS also offers toll-free access numbers as an alternate. These were likely put in place later on, when more flexible 4ESS/5ESS or DMS-100 switches had become the strong majority of the telephone network.

GETS users, after completing an application, are issued a card with the GETS access number and a frustratingly long 12-digit PIN. They call the GETS access number, enter the PIN, and then dial a phone number. While GETS looks a lot like a prepaid calling card to the user, internally the telephone network prioritizes calls from and to GETS to ensure a higher chance of completion. Perhaps more importantly, at least early on GETS calls were routed using more complex logic than normal calls so that---similar to the original goals of NETS---GETS calls should automatically route around unavailable trunks. Today, with a higher degree of automation in traffic engineering, it's unclear how much of a difference that functionality would make.

GETS achieves the goal of call prioritization on the landline network, but it hits its limit when it comes to cellphones. The typical capacity problem with cellphones is not long distance trunks, but local tower capacity. GETS won't help with completing a call when your phone isn't able to set up a voice session in the first place. Well, in practice, some cellular carriers do seem to use more flexible modern GSM baseband capabilities to automatically prioritize GETS calls at the cellular network level as well, but this isn't universal and wasn't even really possible with earlier cellular standards.

So, for cellphones, we have a separate system: WPS, the Wireless Priority Service.

To use WPS, you have to apply to DHS to get an authorization letter and then set up WPS service with your cellular carrier. Once enabled for your line, the dialing prefix *272 will cause your call to be prioritized within the cellular network. WPS explicitly does not provide preemption, so you might still have to try multiple times, but you will have a higher chance of getting through eventually.

Internally, WPS is implemented as a vertical service code, a generalized capability of phone switches that allows codes prefixed with * and # to operate special features. The GSM specification also has extensive treatment of vertical service codes, making WPS prioritization easier to implement in the cellular network layer.

Like how GETS doesn't (necessarily) affect the cellular network, WPS doesn't affect the landline network. To ensure end-to-end prioritization of a call made on a cellphone, it's necessary to use both. For example, if I needed to reach Comcast in a genuine internet emergency, I would dial *272171062743871234567890121800934648 9. This is so incredibly convenient that the government offers a phone app for WPS/GETS users just to do the dialing for you.

Neither GETS nor WPS addresses the actual availability of a working telephone, though. For this, there is yet another program, but one that is less technical and more bureaucratic. The Telecommunications Service Priority program, or TSP, is an FCC program that allows government and infrastructure users to receive higher priority on telecommunications requests. This mostly means that TSP telephone services will be the top priority for repair after damage to the telephone network, but it also means that TSP users can order new service with higher priority as well. Like GETS and WPS, TSP starts out with an application form to DHS, and it's available to all government agencies and to industry in recognized critical sectors. TSP seems to have been formally created in 1988, indicating that it may be a result of the same NCS efforts that lead to GETS.

All of these services might seem somewhat antiquated and, well, they are, with even 2001 being quite a ways in the past now. As an effort to modernize emergency telecommunications, the First Responder Network Authority was established in 2012. Ultimately part of the National Telecommunications and Information Administration (NTIA) rather than DHS, the Authority is commonly referred to as FirstNet. This can be a bit confusing since the term FirstNet also refers to the actual service, delivered by AT&T.

This public-private situation can be a bit confusing.

While FirstNet is a government agency that includes a technical board and policy efforts in multiple areas, the bulk of FirstNet was contracted to AT&T. The objective of FirstNet in this sense is to develop a nationwide broadband network for emergency communications. AT&T, being a cellular carrier, is currently implementing this vision using cellular technology. In many practical senses, FirstNet is just a tier of AT&T cellular service that comes with prioritization.

It didn't totally start out this way. The original goal on FirstNet was more like a unification of the existing public safety communications systems across the US. These systems, based primarily on the APCO P25 trunking radio protocol but with increasing use of private LTE, are mostly operated by states and municipalities and could not easily be counted. FirstNet proposed to build new state-level unified networks or to integrate networks developed by the states. NTIA reserved a substantial federal radio spectrum allocation for this purpose, one that was suitable for broadband radio protocols like LTE. States were given the option to develop their own FirstNet access layer for integration into the nationwide system, or to allow the national FirstNet operator to do so. In practice, AT&T proposed to build out the new state network in each state and all 50 states accepted this proposal.

AT&T developed their access network mostly by addition to their existing cellular sites. FirstNet deployment started properly in 2017, and FirstNet now covers pretty much AT&T's entire network footprint with LTE service. FirstNet is being upgraded to 5G along with the broader AT&T network. AT&T handles the entirety of FirstNet administration, and you get FirstNet service simply by applying to AT&T for a FirstNet cellular plan. Many FirstNet users have their personal phones enrolled in FirstNet for the cost savings compared to a dedicated phone, so it really does look and feel like normal cellular service.

Much of the FirstNet product is really just a matter of configuration (FirstNet provides prioritization to users an will even preempt non-FirstNet traffic when needed), but the dedicated bandwidth set aside by NTIA is still around in the form of band 14. Even this detail of FirstNet is confusing when it comes to the lines of AT&T vs. government services. Band 14 is dedicated to emergency communications, but as part of its deal with the NTIA, AT&T is allowed to use band 14 for their other customers as well, as a means of defraying their cost in deploying band 14 access points. So band 14 really consists mostly of plain old AT&T customers, but with the expectation that they will all be booted off the band when FirstNet traffic requires the bandwidth.

FirstNet is sort of a complex creation, heavily promoted by AT&T for reasons of their own profit motive, and supported by billions in government funding. As you would imagine, it has not been without controversy. The program has been expensive and slow to roll out, even with the significant advantage of mostly just using AT&T's existing network. Even after all that time and money, the original vision of unifying first responder communications hasn't really been achieved. While FirstNet is broadly used by groups like firefighters for cellular service, states and municipalities continue to operate their existing radio networks alongside FirstNet. Most of this is attributable to the high cost of FirstNet devices with ergonomics similar to existing systems (e.g. PTT) and the complexity of deployment relative to the land-mobile radio (LMR) systems public safety agencies are familiar with.

The issue of PTT is worth dwelling on. Users with a great degree of telecom nostalgia probably fondly remember Nextel, the short-lived cellular service provider that used the Motorola trunking-radio protocol iDEN. Nextel aggressively advertised the key advantage that iDEN held over the CDMA and 3GPP standards: with roots as an LMR protocol, it offered excellent support for PTT conversation between groups of phones. This made Nextel extremely popular with businesses like towing companies and the trades, where easy PTT communications between the office and employees in the field was convenient but did not justify the cost of an LMR system. Unfortunately, cellular PTT technology largely died with Nextel, replaced by IP-based systems that failed to offer the low latency and reliability of iDEN.

The concept of cellular PTT is not dead forever, though. The lack of good PTT support has long been seen as a critical deficiency of FirstNet and probably a complete blocker on its regular use by most first responders. Fortunately, AT&T has been an ongoing proponent of the 3GPP MCPTT or Mission-Critical PTT standard. MCPTT employees IP quality of service and multicast technology in the access layer to provide IP-based PTT that still mostly performs as well as radio-based systems. MCPTT in the form of AT&T FirstNet PTT (registered trademark) has been included in the FirstNet offering for just a couple of years now, and MCPTT-capable "LTE radios" like the Sonim XP5plus are now available (at a steep cost) to FirstNet users. These devices are essentially phones, some feature-phones and some just Android devices, but have a physical form factor more like a handheld radio including a large PTT button and top volume/channel controls. They may change the fortune of FirstNet in years to come.

On a more personal level, there seems to be a (probably justified) degree of mistrust when it comes to AT&T's performance in a severe emergency. Much of the US is now covered by a state-run shared public safety radio system. These systems usually use well-established technology (APCO P25) over state-owned microwave and fiber networks. Since they are relatively isolated from non-government traffic, they don't require any potentially complex prioritization schemes to ensure reliability in an emergency. Because they're state-owned, they're often easier for states to expand and modify. This gets at perhaps the most significant criticism of FirstNet: it's just so complicated. It's technically complicated, but moreover it's bureaucratically complicated, offered as

a joint venture of a private company and a semi-independent government authority to user through several degrees of contractual separation.

Today, between GETS, WPS, PTS, and FirstNet, it is now hopefully possible for a government or critical infrastructure user to reliably make a phone call in some cases some of the time. It took decades to get here, and there remain questions about the actual reliability of the service, but massive contracts with AT&T to deliver critical services of questionable quality has become a fine American tradition since divestiture. I saw post-divestiture because I feel like prior to that point even more massive contracts with AT&T usually delivered services that actually worked, and on schedule even, but maybe I'm just being nostalgic.

[1] I may be inserting some amount of opinion here, but honestly, the situation does not allow much room for debate. 9/11 made it extremely clear that the federal government had no meaningful civil defense capability.