Reference: National Institute of Standards and Technology
[Docket No.: 110727437–1433–01]
Soliciting Input on Research and Development Priorities for Desirable Features of a Nationwide Public Safety Broadband Network

Requirements analysis.

The first consideration to the requirements should be to sort them[1]. The salient characteristic of the internet (including network segments that are sloppily labeled as 'broadband') is a modular separation between infrastructure and application. Some of the requirements listed deal with infrastructure, some with applications. And some are mutually contradictory. So the first step of analysis should be to sort them.

I. Infrastructure requirements:

(1) A dedicated high-quality network connection always available for sending and receiving continual data streams to support monitoring and resource tracking;
(3)[2] An infrastructure that is hardened and secure, providing a high level of system availability;
(5) Geographic coverage that has no limitations within the footprint of the National Public Safety Broadband Network;
(6) Dynamic management and control of the network;
(11) Access to and from external information sources[3];
(13) Automatic management and control of the network;

II. Application requirements:

(2) At a minimum, access to initial and updated basic incident information (voice- and text-based incident data);
(4) When voice is converged for normal operations and in the event the infrastructure is compromised, public safety communications must remain stable and with clear voice communications;

1 Readability. I've copied relevant parts of the solicitation into this response, all in 10-point font.
2 Maintaining your numbering unchanged. Security is an end-to-end function and therefore part of applications. 'Hardening' is an infrastructure issue.
3 Both infrastructure and application, but if the infrastructure is not interoperable, then the application interoperability is moot, so I sorted it to here.

• Infrastructure-less communications, with talk-around for the ability to talk one-to-one and one-to-many
• Optimal audio quality during adverse field conditions
• No latency on mission critical voice applications
(10) Single devices that support voice, video, and data;

#(x) Content security -- authenticity and confidentiality – are end-to-end requirements. Consequently they cannot be solved segment-by-segment within the infrastructure. First, emergency services need the content security (omitted from the list, in error). Secondly, requirement needs to be added to the applications list.

III. Hopelessly entangled:

(7) Interoperability, including with existing public safety-based systems;
(8) Ability to send and receive large amounts of information;
(9) A non-proprietary network based on industry standards;
(12) Easy integration with other technologies;
(14) Current and future enhancements available to commercial consumers are provided to public safety with no limitations; and
(15) Ability to send, receive, and process information from the public (citizens and media).

This last, #(15), statement is very important, and inadequate. The system should indeed allow citizen <--> government communications (e.g. 911). But it should also support citizen <--> citizen communications. This ability often decants a missing-persons problem into a self-solving one removing it from the emergency services to-do list.

Analysis of the Hopelessly Entangled category:

There are three motivations:

- life cycle maintainability
- system integration
- general purpose interoperability

which appear in this list. All three logic trains lead to a need for modular system design. In practice this turns into 1) the division between application and infrastructure observed above and 2) the division of the communications infrastructure into terrestrial-WAN, radio-WAN and LAN – all routable networks. If we are faithful to this modularization, then all three motivations are addressed.

Modularity payoff. Terrestrial-WAN, radio-WAN and LANs are all communications segments that have protocols and modular boundaries that make them routable networks. This means they can be snapped together into internetworks. And technology substitutions between the routers are both possible and transparent to the system as a whole. Further, the infrastructure is agnostic about the applications it supports, so we have no limitations in the infrastructure to application development.

#(7) contains a self-contradiction. The legacy communications systems we have, as your Notice explicitly states, are not interoperable. If we maintain backward compatibility, we cannot step forward into interoperable information systems. The priority requirement should be to move into an era of modular (aka routable) communications networks and infrastructure-agnostic applications. Backward compatability should be, at best, a secondary requirement.

This contradiction can be managed by focusing on a communications system of routable networks (an internetwork, looking forward) with a set of layer 7 gateways at the fringe that allow attachment of the mismodularity-limited legacy systems. This is how the 'telephone' system is organized today: the backbone is all routable network and the legacy circuit-switched (aka POTS) phone is only seen at the local loop fringe.

Over-emphasis on backward compatibility is very easy to do because we tend to consider our EMS infrastructure as enduring[4]. By contrast, how old is your cellphone? The legacy turns over faster than we estimate ... unless we've frozen it in amber due to poor modularity. It's worth observing that any information system worth anything outlives all the components in the information system.

Definitions. The Notice, and especially this section, uses several terms that are poorly defined. This is probably unavoidable in the IT business, but it needs to be recognized. Just one example, from wikipedia: 'Different criteria for “broad” have been applied in different context and at different times.' Sloppy definitions unlock a whole Pandora's box of evils, without any compensating hope.

The FCC was saddled with this same problem in the National Broadband Plan: in their Notice of Inquiry, the very first question FCC sought public comment on was what the definition of broadband should be[5]. The answers they got must have been unsatisfying since

4 And requirements-writers are describing shortcomings in what they have in hand.

5 FCC's GN Docket No. 09-51 , April 8, 2009, 15. “Broadband can be defined in myriad ways. In order to ensure that all people of the United States have access to broadband capability, we must make sure that the Commission appropriately identifies goals and benchmarks in this regard. Here, we seek comment on how the Commission should define “broadband capability.”

they left the issue open.

One term that you need to get straight, early, is 'dedicated high-quality network [6]'. While no one will argue that emergency services needs priority access into a highly robust communications system, that does not dictate a separate infrastructure from the civilian sector. Indeed a segregated infrastructure would negate #(15). DoD's experience might be helpful here. In the 1990s, DISA owned three terrestrial networks that were indeed segregated infrastructures: NIPRNET, SIPRNET and Defense Switched System (for unclassified, secret-system-high, and POTS voice respectively). Today, DISA has a single backbone: SIPRNET is tunneled through NIPRNET using virtual private networking tunneling. And DSN, except for the local PBXs, is all converted to VOIP and passed through the same infrastructure. The result has been higher quality by just about every metric you'd wish to apply. In many areas where we need emergency services communications, we will never be able to afford dual infrastructures. The FCC was on the right track with the 'public-private partnership' discussion, but didn't get very good at it.

6 #(1) sorted into infrastructure

Analysis of Infrastructure Requirements category:

The best way to organize this list is to ask the question: what makes emergency services communications different from out-of-box internet?

Those differences are four[7]:

- Higher availability requirements. The emergency services network needs to work when nothing else does.
- Security. Emergency services applications always need authenticity[8] and sometimes need confidentiality; the typical smartphone applications lack these end-to-end security features.
- Reach to mobile platforms. Understanding the differences here will unlock some of the QoS mysteries. There are a couple of 'hard problems' in here, from which you should derive the lesson that getting the terrestrial WAN portion of the network right FIRST will make this part of the problem easier. 'Backhaul' should not be an afterthought.
- Quality of Service control. Emergency services rightly have a need for priority access in emergent situations.

In many years of analysis and teaching, I've found that these four categories encompass the infrastructure problem, and therefore recommend the logic to you. The rest of this section expands each of these four differences. I believe that the result will be more comprehensive and balanced than the list you started with.

High availability. In my experience, the first/worst problem is inadequate expression of the need. I've seen consulting reports and requirements statement (Iike #(3) in your Notice) that fail to quantify the need. How many nines? This is a failure by the requirements sponsor (aka fire chief or police chief). The key to quantification is to transform the definition of availability (uptime/total time) into (total time - down time)/total time. And tie tolerable down time to real world requirements (what does 'continual' in #(1) mean?)

Once you have a real requirement stated, apply the three principles of high availability engineering. Which are:

- eliminate single points of failure (altroutes and backup power)
- reliable crossover (routable networks do this all day every day)
- detect failures as they occur (the real reason for #(6)).

The emergency services community frequently mouths terms like 'assured access'. The term expresses a general desire but lacks definition. The first step to assured access is a highly available and survivable infrastructure.

You should also note that technologies to implement high availability exist today. There is little need for either research or standards-writing. But there is a need for competent requirements analysis and systems engineering.

---

7 Only!

8 There are no exceptions.

Security. The security problem divides into two parts:

- End-to-end security (aka content security) spans multiple network segments and can therefore not be solved in the infrastructure – content security can only be addressed in the design of applications. Language like 'secure VOIP' or 'S/MIME' (secured e-mail) are in the right territory. What's important is to not let the end-to-end content security requirements leak into the infrastructure design requirements list; else you present yourself with an impossible problem.
- Infrastructure protection tends to become a secondary problem once the content protection part is decanted off. Here terms like denial of service, jamming, traffic and traffic flow analysis, interceptability ... appear[9]. These issues pertain to a particular network segment (with the pain usually in the radio-WAN).

Reach to mobile platforms. In the civilian sector, 'mobile' usually generates discussions of WiFi, cellphones, and the corpus of [smartphones, PDAs, laptops etc] that attach. Aside from the gadgets instead of infrastructure myopia, this is not the emergency services problem.

- Topology. Emergency services ubiquitously connote mobile platforms. An ambulance may have several end systems that need to connect, not just one. So the first difference to recognize is that we have a radio-WAN problem, not a LAN one (wired or wireless). We need to reach to the mobile platform, and that is a router-to-router interconnect problem.
- Shared media. Most of the commercial internet is made up of point-to-point pipes (fiber optic or copper). But radio is a shared media. This surfaces two issues:
  - multicast. Defined as delivery to more than one destination for the price of a single transfer. The proportion of multicast (#(4) in the applications list hints) is much higher than in civilian world. While multicast protocols exist in the commercial internet, the uptake, particularly in applications, has been poor because there's little payoff in the all-terrestrial, civilian internet.
  - MAC. With a shared medium, you need a Media Access Controller to take turns on that media. There are two kinds of MACs out there (contention and contention-free[10]).
- Four orders of magnitude less capacity. Today, both the terrestrial-WAN and the LAN portions of the internet can be easily provisioned at 10Gbit/sec capacities. By contrast, radio 'broadband' network segments have 10Mbit/sec capacities, pro-rated across all the users of a network segment. This capacity limitation, while getting better, is approaching the bits/Hertz limits where we can't get better without more spectrum.

The deckplate systems engineering lesson here is that the radio-WAN problem will be the hardest and most expensive part of the overall comms system. Overdo the terrestrial-WAN part and the LAN-in-mobile-platform part in order to keep the radio-WAN part as small/simple as possible. For example, keep end systems off the radio-WAN (other side of the router, on the LAN), thereby keeping the radio-WAN problem purely a router-router interconnect one.

---

9 In " Why (Special Agent) Johnny (Still) Can't Encrypt: A Security Analysis of the APCO Project 25 Two-Way Radio System " by Sandy Clark , Travis Goodspeed , Perry Metzger , Zachary Wasserman , Kevin Xu , Matt Blaze, presented at USENIX 2011, we have a perfect illustration. The confidentiality vulnerabilities are properly application shortcomings; the passive and active traffic analysis and jamming are infrastructure shortcomings.

10 Known in the voice world as free- and directed- nets respectively.

Quality of Service Control[11]. The QoS problem requires understanding of the reach-to-mobile situation. We have three components worth discussing, one in the terrestrial-WAN and two in the radio-WAN.

- in the civilian internet, qos[12] has been routinely addressed by overprovisioning for forty years now. If the capacity is larger than the demand, there's no congestion. No QoS measure can possibly improve qos. One of the side effects of the first principle of high availability is overprovisioning. Do it whenever you can. The surfeit of capacity is the primary reason why the QoS control standards, while existing, are not widely implemented – no payoff. But in the radio-WAN, we must accommodate the four orders of magnitude observation.
- The congestion in an emergency services internetwork will always appear in the radio-WAN. This means we must have contention-free MACs because they are stable under overload. All known radio-WAN contention-free MACs use some kind of TDMA mechanism meaning that each subscriber station gets a guaranteed transmit slot every frame (analogous to determinism in computer operating systems).
- The congestion in the radio-WAN will be manifest at the routers that surround it. Each subscriber station will be running multiple applications with different users and different urgency requirements. So its the router's job to stack up the most urgent traffic at the top of the queue so it gets sent first. Differential Services is the technology and standards to accomplish.

This QoS control discussion usually rounds out the issue of 'assured access'. What's important to grasp is that most of the internet and 'broadband radio' technologies do all of this already (for example, Differential Services only needs to be turned on in most routers). The missing part of the discussion is in applications: end systems need to set the Differential Services Control Point data correctly in order that routers can do their queue-sorting jobs.

11 Tarbaby warning.

12 Notation convention for the non-IETF-initiate. qos, lower case, is the quality of service being experienced in the default situation. QoS, upper case, means Quality of Service control or what you would do to change the defaults.

Analysis of Applications category:

There is one internet out there; but there are 1001 applications that run over the internet. With more every day. It is impossible to enumerate all the features wanted; your list is pretty minimal[13].

The most important requirement that all emergency services applications must support is end-to-end security[14]. Electronic mail has supported such (S/MIME) for over a decade; secure VOIP is a work in progress today.

Standards body shortcoming. In examining the S/MIME and secure VOIP technologies, we find no standards bodies that cover 'end-to-end'. IETF (author of most of these standards) is concerned with 'bits over the wire', but their scope does not include through the operating systems and computer storage. We have uncovered instances where the traffic is secure over the network, but that security stops immediately on data entry into the computer operating system where all kinds of potential malware can get at the content in unprotected form.

In the civilian sector there are several mechanisms to invent and improve applications over time. The open source operating system distribution is perhaps the most stunning example. Each release is better than the previous. There does not seem to be any analog in emergency services. A couple examples that would have payoff:

- a next-generation 911 application that would include identity, authenticity and position as part of the 'smartphone' application[15].
- A 'Common Alerting Protocol' editor that works on portable computers such as laptops and smartphones. A matching requirement would be a CAP feed display application. Again, authenticity is a requirement.

13 The NPSTC statement ”Mission Critical Voice Communications Requirements for Public Safety ” released Sep 2011 is NOT a treatment of 'mission critical' (a treatment of mission critical would be almost all infrastructure). But it is a decent start at features needed in a 'fire department VOIP application'.

14 Ironically omitted from the NPSTC statement.

15 This intersects the FCC's next-generation 911 inquiry where FCC acknowledges that the infrastructure no longer infers location, identity or authenticity – these features must be in the application

Feature List

You requested comments on this feature list. Herewith:

Feature List (organized around the four overarching themes noted above):
To ensure resiliency in an emergency:
• Resiliency: The ability of operable systems to recover from mishap, change, misfortune, or variation in mission or operating requirements.

See the comments on high availability above. The principles can be applied to anything; they fit routable internet technologies very well.

• Self-Organizing: Self-organizing networks dynamically manage their own configuration by automatically making changes to ensure messages reach their destinations.

The internet has been doing this ever since the introduction of Routing Information Protocol back in the ARPANET. Today, BGP and OSPF are commonly used in tandem.

• Meshing (ad-hoc device-to-device communication): A type of networking where each node must not only capture and disseminate its own data, but also serve as a relay for other sensor nodes, that is, it must collaborate to propagate the data in the network.

The term 'mesh' is undefined (or more correctly, swamped with too many sloppy definitions). Your statement does not differentiate between end systems and intermediate ones (like routers). Most computers' operating systems contain the routing protocols to allow them to route as well.

Most of the MANET (layer 3) and mesh (layer 2, see IEEE 802.16m) protocols have had little uptake in the last 15 years. Part of the reason is poor scaleability – the layer 2 meshing protocols require homogeneity of the nodes that is impractical. The layer 3 protocols (OLSR, DYMO) carry quite a bit of communications overhead, some of which increases proportional to the square of the number of nodes. Of course, if the comms system is made up of non-routable networks (e.g. P25), then the discussion is moot.

In practice, we've found little quarrel with existing routing information protocols (e.g. OSPF). They seem to do the job fine; but can be easily upgraded with a software download if necessary.

• Adaptability: The ability of the network and/or device to modify/ change behavior based upon external conditions.

One of the standard functions of a management protocol is 'configuration management[16]'. Simple Network Management Protocol provides an extensible framework to add whatever configuration variables are desired.

To ensure reliability and availability:
• Prioritization: The ability to prioritize network traffic based on assigned priority schemes.
• Quality of Service (QoS): The set of standards and mechanisms for ensuring high-quality performance for critical applications. By using QoS mechanisms, network administrators can use existing resources efficiently and ensure the required level of service without reactively expanding or over-provisioning their networks. The goal of QoS is to provide preferential delivery service for the applications that need it by ensuring sufficient bandwidth, controlling latency and jitter, and reducing data loss.

As stated, this reads like you want to be able to manufacture capacity out of thin air. This is not possible. What is possible is to prioritize traffic. As discussed above, this requires a combination of contention-free MAC and differential services.

To enable security:
• Strong, Dynamic Access Control: Access control lists can be configured to control both inbound and outbound traffic on networks and authentication/ verification of users/devices on the network. The level of access control should be sufficient to allow for entree into a broad set of systems and databases needed by public safety (e.g., criminal history databases, medical records, public work records, etc.).

There are two issues in here, for which two different solutions are needed. Access control (as implemented at routers) 'keeps the riffraff out' of a particular portion of the internet. This makes sense as an infrastructure protection measure.

But access control is not a satisfactory solution for protecting confidentiality[17]. For that, encrypt the data.

To ensure affordability/commercial alignment:
• Compatibility with Commercial Infrastructure: The utilization of a variety of commercial services when public safety is in areas not covered by

16 The other functions are fault, accountability, performance and security.

17 Terminology alert. Within operating systems (and implemented in the SELinux work in most Linux kernels today) there are terms like Mandatory Access Control and Discretionary Access Control. These terms have little relevance to the issues at hand, but the indefinition of the terminology will drag at least some of an audience down a dead end.

the public safety broadband network.

Terminology: do you not mean 'interoperability' rather than 'compatibility'?
This is important for several reasons.

- interoperability with the citizens that emergency services exist to serve,
- mainstream technology.

Compatibility in communications systems generally only means non-interference. You are inferring a desire for something more.

• Network sharing: The shared use of infrastructure between commercial and public safety users.

More of the same? In lightly populated parts of the US we have emergency services requirements just like urban areas. But we don't have economic justification for more than one communications infrastructure. Anchor-tenant agreements with commercial ISPs should be strongly encouraged.

We have several instances of anchor-tenant arrangements where a commercial carrier provides communications services for emergency services (the anchor tenant) and then sells any capacity excess to that for profit. There are two issues herein: one is that we need 'shared technology' which routable networks provide (extend the internet). The second is a business model that takes advantage of the technical interoperability (the anchor-tenant arrangements). We have cases in the internet domain where this has worked successfully. We have no cases where segregated comms systems have allowed infrastructure sharing.

• Multi-Modal: The ability of the network to support voice, video, data, and multimedia simultaneously.

What of this list does the internet not do today?

• Scalability: The ability of a system, network, or process to handle growing amounts of work in a graceful manner or its ability to be enlarged to accommodate that growth.

The internet design characteristics of

- separation of infrastructure from application
- communications as routable networks

has resulted in an network that has scaled, often at rates exceeding 100% per year.

At the design phase, this could include requirements to ensure that scalability can be achieved, to the extent possible, by software enhancements and upgrades as opposed to by hardware replacements. Scalability also includes the need, in the case of a large scale event, to accommodate a rapid increase in the number of users in a limited geographic area.

• Power Awareness: The ability of network/devices to control power functions.

Treat the radio-WAN function – reaching to mobile platforms – separately from the LAN function (reaching to end systems). The latter can be done by, for example, lighting up a WiFi hotspot on a fire truck or a wired ethernet LAN in an ambulance. This puts the radio-WAN power requirements in the vehicle-mounted radio-WAN subscriber station, not the end system. And because the end system only needs to be 'WiFi aware' or 'ethernet aware' it can be mainstream technology (e.g. smartphone or laptop, or adaptation from that).

• Standardized Common Interfaces: Protocols, Application Program Interfaces, application platforms, radio capabilities, etc. that allow for competitive provisioning.

A lot of government treatment of standards attempt to standardize 'the box' rather than the interface to the box. And quotation of a laundry list of standards cannot add up to a modularization model. You will want a universally-observed modularization model to enable the focus on interfaces[18].

• Uniform, Universal Access: The ability to access the network and data anywhere at any time through any device.

Request for Comments
For each feature listed above, NIST is requesting input on the following:
• Your assessment of the importance of the feature in relation to a Nationwide Public Safety Broadband Network;

Without knowing what a Nationwide Public Safety Broadband Network is, it's difficult to address this question. I hope that the term does not connote a separate emergency-services-only infrastructure – that contradicts several of your interoperability and sharing desiderata. 'Assured access' and 'mission critical' are important, but they do not require infrastructure segregation – indeed, that would be counterproductive[19].

• Current gaps that exist preventing the realization of the full potential of the feature;
• Possible research and development that could take place to close any technical gaps;

18 API is the term for an interface definition in the computer world; in the communications domain, the analogous term is SAP: Service Access Point.

19 This always triggers emotional outbursts usually using the quoted terms. Emotions will get to the wrong engineering answers unless they are soberly analyzed.

Several comments at the end of this submission. But in general, the issues are systems engineering ones, and 'this is the way we've always done it' ones. There is very little missing R&D.

• Any challenges that public safety could face in realizing the full potential of these features given currently implemented solutions;

You have several bullets here that glance upon acquisition issues (multi-vendor, open standards, etc). Within the information technology world, from the government perspective, there is often great value in procurement vehicles like IDIQ contracts. But the trace needs to start with observations about budgets.

EMS budgets tend to 'roller coaster'. Between large scale disasters, they are minimal and flat; once Katrina comes to town, money magically appears. Whining about this characteristic probably won't be productive, but there are two valuable antidotes:

- One is to observe that the communications that the schools need is the same communications that our emergency services need. And both government activities are funded from the same tax base. In the past, because of the stovepipe characteristics of emergency services communications, a unified communications budget and plant was infeasible. No longer the case.
- The second is to place contracts for the 'right stuff' that can be exercised to 'bulk up' when needed. Do the planning during the budgetary flat spots, not in the heat of a crisis. This thread might drive you beyond just standards-making to the placement of IDIQ-type contracts on behalf of undefined (but EMS) users – making the supply system more 'shovel-ready'.

• Best practices from other industries that could be leveraged to expedite public safety's realization of these key features.

Additionally, NIST is requesting input on the following further considerations for the nationwide public safety network:
• What is the importance of employing open standards for the nationwide public safety network?

Well, yes. It's hard to imagine someone maintaining that open standards would not be important. But ...

First, standards, without a modularization model are akin to a cart without a horse. It's very easy to hopelessly mismodularize an information system using a set of open standards. The modularization model (in this case separation of application from infrastructure) is the first step; then you can crystalize the modular interfaces with standards.

Second, it depends on which open standards you select. P25 is purportedly an open standard, but it is for a trunked-radio, voice-only communications system; it is not a routable

network. IEEE 802.11 (WiFi) is certainly an open standard, and one for a routable network. But IEEE 802.11 contains a contention-MAC making it wholly unsuitable for a radio-WAN purpose (it's fine as a LAN technology). IEEE 802.16, on the other hand, is a standard for radio-WANs – we can extend the internet with equipment that adheres to the standard.

• What is the need, if any, for commonality of functions across the system?
• What is the importance of a multi-vendor environment for the network and what are the lessons learned in deploying a multi-vendor environment from the cellular and other industries?
• What can be done to ensure both short- and long-term affordability of the network for all types of public safety agencies?
• In a recent report, the President's Council of Advisors on Science and Technology suggested the need to develop methods for implementing a "survivable core" of cyber-infrastructure that would be relied upon to provide truly essential services in the event of a catastrophic cyber-attack. Please comment on how NIST should pursue this recommendation. Among other things, commenters should address whether the goal should be to design a separate survivable core that is integrated and interoperable with the primary public safety network, or instead to design the primary network such that it can reconstitute rapidly—following a catastrophic event—to achieve some "core" level of service.

Almost every commentary that uses the term 'cyber' tends to get into a 'sky is falling' mentality that lacks much analysis. The most important tool, both for analysis and for designing against the threat lies in the distinction between end-to-end security (focus on the content) and infrastructure security. We habitually use infrastructure protection methods (for example, anything that secures WiFi or a USB memory stick) and then wonder why we get end-to-end security breaches.

Survivable core, yes. Separate, no. Very few attacks have targeted the internet itself; rather they use the internet as a conduit to attack end systems attached to it. If you segregate the network infrastructure you violate the first two principles of high availability engineering.

Reconstitution is a different issue – it's an availability/survivability one rather than a security issue. But it is important. If the pre-disaster network is an internet – made up of routable networks – then it is very straightforward to use rapidly deployed routable networks to reconnect the residual intact internetwork.

• What is the marginal cost of the feature/functionality versus equipment

available today?
• What network features or requirements have not been identified above, the lack of which may impair the network's ability to adequately serve the needs of public safety?
• How should NIST engage public safety practitioners and technologists as part of the planned R&D projects to ensure proper prioritization of efforts and effectiveness of developed solutions?

Internet architecture. The internet has grown up largely without government technical involvement:

- FCC has managed spectrum, but until National Broadband Plan, has not involved itself with protocol issues.
- DHS has not been involved with protocol issues either. Recently the SAFECOM treatment on the DHS grant guidance web page has acquired a single reference to 'broadband'.
- There has been previous NIST involvement. For some time, the chair of the IEEE 802.16 committee was a NIST employee[20]. And NIST sponsored the abortive, ill-conceived Government Open Systems Interconnect Profile (GOSIP, written up in FIPS 146).

It appears that this state of affairs is changing as we increasingly shift from circuit to packet switched systems in the radio environment – the government is getting more involved. Recommend some thought regarding how this is to be managed:

Is federal government involvement in internet architecture a good thing? To date, the reputation of the federal government dealing with things-internet has been almost all bad.

How many voices in the government will be speaking?

Modularization model for entire emergency services information system. Standards are not enough; a modularization model is needed to support systems integration (incl multiple vendors), life cycle maintainability, and overall interoperability and end-to-end security. None of the existing federal government practices or recommendations (e.g. in the DHS grant guidance web pages) have a modularization model.

Red team. P25 has been pushed to emergency services for well over a decade and contains egregious security concerns that have only recently been uncovered[21]. This is unsatisfactory federal government performance.

Information to EMS personnel and administrators. The government needs some means of reacting to discoveries like the egregious shortcomings of P25[22]. There appears to be no effective way to apprise emergency service personnel about the security vulnerabilities in the radios they use. Do we wish to live with informal conduits like Matt Blaze's web site?[23] This is also unsatisfactory federal government performance.

---

20 IEEE 802.16 defines specifications for a radio network that certainly qualifies as routable and also fits most definitions of broadband.

21 See Footnote 6.

22 DHS CERT has been wholly unresponsive.

23 http://www.crypto.com/blog/p25 and http://www.crypto.com/p25/

Catalyze open source solutions[24]. End systems applications thrive in the open source community. Directly applicable is the Linux / Android blossoming we are seeing now. Commission some emergency services applications; insist that the products be GPL and get them into the distributions.

But there are no open source applications that meet emergency services needs[25].

Vintage open source development, as the current open source operating system distributions work, does need some small adaptations (but not enough to kill off the ethic). These adaptations should consist of a 'certification review' to assure that end-to-end security needs are met. Most of this can consist of red teaming.

Multicast. One of the salient differences between civilian and emergency services is that emergency services have a much higher incidence of data that is multicast in nature (multi-party voice conversations, asset tracking (in military, known as common operational picture) and, for example, see your #(4)). Most of the internet infrastructure handles multicast quite well today (in ISO Reference Model terms, from layer 3 down). But because of dearth of demand, multicast transport protocols and muliticast applications have not had much uptake in the market.

This is a specific instance of the 'catalyze open source' issue. Multicast transport protocols have been around for 15 years (see RFC 5740) but neither NORM nor applications that use it appear in the open source (or even closed-source) distributions.

Another multi-cast related problem is confidentiality. The existing public key mechanisms rely on a publicly-known key for authenticity, so this matches a multicast environment. But confidentiality relies on a targeting each user's private key – which will all be different. This constrains the confidentiality mechanism to unicast. A feasible, scaleable solution is needed; it needs to work with many applications (e.g. both messaging and VOIP) and it needs to be standardized.

Concluding perspective.

Over the past forty years, we've watched packet switching technology take over various parts of the communications infrastructure. These are colloquially known as the bellhead/nethead wars.

Today the 'telephone' backbone is packet switched; ATM is a forgotten fad. DOCSIS-compliant internet extensions that use the cable TV plant abound (and several now-forgotten circuit-switched wannabes foundered here). Office automation has been taken over by ethernet-based packet technologies making circuit-switched fax, for example, a 'has been' ... Because of the advantages of packet switching, the technology has supplanted older circuit-switched (or altogether unswitched) technologies.

And today, the focus of the bellhead/nethead wars has shifted to what is perhaps the only remaining bastion: radio. We are in the middle of this transition, so a little perspective about what's going on should be useful to keep us out of dead ends.

First, the campaign is well-along. WiFi – an visible artifact of nethead – is ubiquitous. And what the cellphone companies call '4G' are all packet-switched technologies. Indeed, in the bellhead/nethead campaign, almost the last holdout is our emergency services

24 I belive this is the answer to your marginal cost query above, but I don't understand the question.

25 Leading to sober but laughable proposals to use twitter or SMS.

communications.

We can draw some generalizations by looking at various campaigns in the Forty Years War between bellheads and netheads. The first is that packet-switching technologies – extending the internet – have always won out. There aren't any exceptions, only holdouts. Once the emotional is decanted off, there are solid engineering reasons for this: bandwidth economy, ability to build high availability infrastructure, adaptability of applications, and an inherent modularity that separates end system (aka application) design from infrastructure. So there is almost an aura of inevitability: packet switching (labeled 'broadband' in this particular nethead campaign) is the direction to go.

The second observation is that we get a lot of terribly confused requirements statements. They don't 'get' the packet-switching features. A shining example of this is the above noted NPSTC statement that purports to define 'mission critical' but is in fact a list of features for an application ... which has nothing to do with, for example, high availability. This 'standing on fantail watching wake' phenomenon is very common and requires care. When a requirement is stated, you usually need to ask 'why?' a lot to relate it to real needs[26]. Looking at a packet-switched communications system through circuit-switched experience introduces a lot of (well meaning) distortions.

We see this confusion surrounding the term 'interoperability'. I've heard some of the silliest interpretations and definitions of the term in emergency services ... overlaying a very important real need. The internet delivers communications interoperability; so we should be extending it to our emergency services ... you have several contributing reasons sprinkled around in this Inquiry.

But emergency services has particular requirements that are often met only in potential in the vanilla internet. Most of the internet that we touch every day is at the fringe where it's almost always single-threaded (and therefore of inadequate Ao for EMS). This does not negate the packet-switching technology, but it should remind us that 'out of commercial box' solutions need some adaptation and some 'systems' thinking and balance when making plans and investments. Thus the plowshares --> swords analysis above. Virtually none of these adaptations require any research; and most require minimal development. Sinking budgets into federally funded R&D will not yield much return.

26 The requirement for full duplex in the NPSTC mission critical statement is one of these. Fulldux is a nonsense term in a packet-switching environment and relates to physical media dependent needs, not applications.