Implementing a Virtual MSAG

Jerry Steenson
Jerry Steenson GIS, Inc.
10/24/2011

Introduction

The National Emergency Number Association (NENA) Next Generation 9-1-1 i3 (NG911) standards, while not complete, are largely finished. In the realm of data management, one of the biggest departures from the past that is introduced by these standards is that they advocate a far greater role for Geographic Information Systems (GIS) as the center of 9-1-1 spatial data management. The Location Validation Function (LVF) and Emergency Call Routing Function (ECRF) are to be 'provisioned' from the GIS in a fully function i3 environment, and the MSAG and ALI will no longer exist.

For most organizations, this change will not happen overnight. As of this writing, the standards for provisioning the LVF and ECRF are not even finalized, though many of the 9-1-1 equipment vendors are moving forward as if it has been. 9-1-1 agencies looking to upgrade to a Next Generation standard, whether it is strictly i3-compliant or not, will need to sort through a bewildering choice of options, which may look radically different from vendor to vendor, before even making a decision. And, once a decision has been made, there are a large number of different systems that will need to be upgraded. For most agencies, all of the systems won't be upgraded at the same time, so even if an LVF/ECRF regime is implemented, the information embodied in the MSAG and ALI will still be needed.

Currently, virtually all 9-1-1 systems in the U.S. are dependent on the Master Street Address Guide (MSAG) and Automatic Location Information (ALI) databases. What will help many agencies is to have a roadmap for how to gradually transition from a pure MSAG/ALI-driven spatial data regime to one where the GIS, feeding the LVF/ECRF is what is done. What will be presented here is just such a roadmap, with enough flexibility to allow optional paths that should fit a large number of institutional arrangements and technical situations.

Overview

The underlying premise that will drive this discussion is that the MSAG and ALI, for many organizations, can be effectively eliminated in the near future, though their vestiges can be kept for as long as is needed to transition away from those legacy systems that still require them. But, before describing the details of that premise, it is useful to step back and take a look at the nature of the MSAG and ALI databases to see why it would be best if they were eliminated.

The MSAG and ALI can be thought of as spatial databases (or tables in a spatial database). The reason for this is that each piece of data in these two objects is tied to a specific location or spatial entity on earth. For an ALI record, this is pretty simple. Each ALI record, since it is tied to a customer record, is also tied to a service address for that customer. Optionally, there is location information, giving us things like apartment number, or whether the service location is in the basement. These elements define the inherent spatial nature of the ALI record.

The MSAG is also spatial data. Each MSAG record has a community, street name, high and low address range, and optionally, odd and even information, which can restrict it to one side or the other of the street range. Again, this places us at a pretty specific location for each MSAG record. Back in the 1970's, when the 9-1-1 system was being devised, the ALI and MSAG databases made pretty good sense as concise ways to locate people in need of emergency services.

A factor to consider is that, back in the 60's and 70's, cellular telephones did not exist. So, phone calls came exclusively from fixed landline locations. Now, roughly 2/3 to 3/4 (and sometimes more) of all 9-1-1 calls are coming from mobile phones, and not landlines. So, the MSAG and ALI aren't of any use for the majority of all 9-1-1 calls. And, as we all know, the use of landlines is rapidly declining. Additionally, VOIP technologies are also making traditional MSAG and ALI location methodologies obsolete. This rapid shift to non-landline phones has made using GIS data to locate callers much more important than in the past. Over the past 10 years, there has been a huge shift in this respect, so that a large majority of Public Safety Answering Points (PSAPs) now have some kind of computerized mapping in place to help locate callers, and determine routing and response. That mapping uses GIS data, whether it is maintained by the 9-1-1 agency, or others.

MSAG and ALI or GIS?

Since most 9-1-1 agencies already are already using GIS data in their operations, there is an opportunity and a need to compare that data to the existing MSAG and ALI data sources that all PSAPs have. This can be thought of as two operations. One is comparison, and the other is synchronization[1]. Comparison is the process of systematically checking the GIS data against the MSAG/ALI database, and identifying where the two disagree with each other. Synchronization is the process of taking what is learned from the comparison, and reconciling the differences between the two, with the desired result being that the two are logically identical.

[1] The NENA standard for this process (71-501 v1 Synchronizing GIS with MSAG & ALI) combines the two into a single process.

I have had substantial experience over the past 10 years in performing these functions, and authored a leading suite of GIS software utilities for MSAG comparison and synchronization[2]. A surprising fact emerged during this time working with the comparison and synchronization processes. When there are disagreements between the locally-maintained GIS and the MSAG/ALI databases, in most instances (I would guess over 80% of the time), the disagreements are resolved in favor of the GIS, and it is the MSAG and/or ALI that must be changed. Not only has this been true in the aggregate, but in every single comparison and synchronization project that I have been involved with, this has been true. This statement is truer with respect to the MSAG than it is with respect to the ALI. Because of the way that the two databases are managed, the ALI tends to be less error-prone than the MSAG. However, it is not unreasonable to say that the MSAG, as a starting point, can always be assumed to be less accurate than the GIS. This is true if there is locally managed GIS data, such as might be maintained by a county or city GIS department. If the only data available is public, such as from Census TIGER (Topologically Integrated Geographic Encoding and Referencing system) or Open Street Map, this might not be true. Also, this depends somewhat on the appropriate definition of error, at least for the MSAG.

Having made the blanket statement that GIS data is always better than the MSAG and ALI, it is worthwhile considering why this might be the case. The next few paragraphs summarize why I think this is true.

First, the GIS, by the nature of how it is used and edited, is less error-prone. This is true in part because the GIS user can see data in a visual, spatial relationship to other data. If the user is looking at a set of street segments on the map, he is also seeing the structures along that street, with their addresses, intersecting streets, and parallel streets that might have similar address ranges. All of this contextual information makes it much easier for the user to detect errors, even without specialized software to assist this process. When editing the MSAG or ALI, the user only sees a single record or a small group of records at a time. There is no way that I have seen in any of the commercially available MSAG/ALI management systems to get any real sense of spatial context, except for systems that show MSAG records in relation to the corresponding street records on the map. There are comparisons of the MSAG to the ALI data, so that the MSAG has to include the range of the ALI records, but that is of a much more limited nature than what can be done in the GIS environment. If the user has good data validation tools, the likelihood of error will be reduced. This is especially true if those tools are capable of automatically checking the integrity of the data in context with surrounding data.

Second, there is the fact that the MSAG is usually edited and managed by a single person. In most jurisdictions around the country, a single MSAG Coordinator keeps track of all the edits and updates to the database. But, even in many small counties, where this is true, the GIS data is managed and edited by multiple users. This introduces some cross-checking of errors, just in the normal operation of the system.

Third, there is the narrow user-base of the MSAG and ALI. The MSAG Coordinator is frequently the only person who directly interacts with the MSAG database. The telephone companies (telcos) typically do

[2] Those utilities are collectively known as MapSAG, which I developed as Chief Technologist for Contact One, which is now a part of Intrado.

all the direct entry work in the ALI, which does sometimes cause errors in the MSAG to be uncovered, and then fixed. But, there really are not many people who directly work with these databases. The GIS, on the other hand, has become a central functional component in the day-to-day operations of almost all levels of government. As such, many users directly work with and view this data all the time. So, when any of these users find an error in the data, there's a good chance that they will report that back to the people who manage the GIS, who can then rectify the situations.

Fourth, I would point to certain habits of how the MSAG in particular has been managed in most jurisdictions. Here is how some MSAGs look, even today:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	0	1000000	B
1st Ave	0	1000000	B
Hideaway Ln	0	1000000	B

In other words, all the ranges are 0 to 1 million, because it's just too much work to get the actual ranges in. And, from the standpoint of the harried MSAG Coordinator, this is just fine, because with the MSAG set up this way, they don't get as many ALI fallouts reported from the telcos. So, in effect, what is being maintained isn't an MSAG at all, but really a Master Street Name Guide (MSNaG for short). This egregious mismanagement of the MSAG isn't encountered often, but instances of this kind still exist. There are also cases where some of the streets have fairly accurate ranges, and others have these kinds of huge ranges that we see in this example. This is usually the case when the MSAG was originally set up with big ranges, just to get it working, with the intent of fixing the ranges, and the process of fixing the ranges is either occurring very slowly, or simply got dropped.

A variation on this fourth problem, which is in my experience the usual case, is that breaks in the ranges are not used to create separate MSAG records. So, say that Main St. has 0 to 400 blocks, then has two missing blocks where there is a park and no street. Then, Main St picks up at the 700 block and goes to 999. Typically, this is represented in the MSAG like this:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	1	999	B

More accurately, to capture the fact that those two blocks are missing, this should be represented in the MSAG like this:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	1	499	B
Main St	700	999	B

The problem with not having the range broken as shown is that an ALI record for, say, 585 Main St can be entered into the ALI database, and it would pass MSAG validation without any problems. When an emergency call for a heart attack comes in for 585 Main St, well, that's just too bad for the person who lives at 485 Main St that the mistaken ALI record actually belongs to.

If the MSAG is incorrect, it is to be expected that there may be errors in ALI records as a result. This is somewhat less likely, as noted above. This is because there are multiple indirect checks on the accuracy of the ALI database. One is that the customer himself is responsible for adding their data to this database, at the time they apply for new service. But, the MSAG only checks to see that an address falls somewhere in a valid range. So, if the telco customer rep types '114' instead of '113', the address will pass the MSAG check, and the error can be inserted into the database. Also, if there is an error in the address, either because the customer gave the wrong address, or because of a data-entry error, the service tech installing or repairing the customer's service will probably catch the problem. But, there's no guarantee that the tech will relay this back to whomever it is that has to make sure that the database is correct. And, again, note that there is a relatively small number of people interacting with this data, as opposed to GIS data, which is continually being checked, even if inadvertently, by all of the system's users.

So, there are multiple reasons why the MSAG in particular, and to a lesser degree the ALI, can be expected to be systematically less accurate than the GIS data. And, anecdotal (but very strong) empirical evidence has consistently confirmed this expectation. The conclusion that can be reached from this observation is that any 9-1-1 agency can expect to improve its spatial data, and therefore the accuracy of its emergency response, by simply replacing the MSAG with their locally produced and locally maintained GIS data. The next section will discuss some methodologies for accomplishing this goal.

The ALI database does tend to be less error-prone than the MSAG. If a jurisdiction has good point address GIS data, or perhaps parcel data with accurate situs addresses, it is usually a pretty simple process to check to see how well the two match. In such cases, the ALI database will be a subset of the GIS address data. This is because many addressed locations in the GIS address data do not have phones, and therefore are not in the telco's database. As cell phones become more popular, this is even more true. But, it is to be expected that each and every ALI record can be unambiguously matched to a GIS address[3]. Making sure that this is true is a good guarantee that the MSAG replacement will be a fairly painless process.

Replacement Methodologies

Once an agency has decided to go ahead and just replace the MSAG with GIS data, there will be different ways to do this, dependent on a number of factors. Once the specifics of the situation are understood, the appropriate methodology may be implemented.

[3] A significant qualifier here is that the Location field in the ALI database is important to the use of the ALI data, and frequently does not match to any known information in the GIS data. So, apartment and suite numbers may not be tracked in the GIS. But, if they are, the data can be validated down to the level of the Location as well as address.

How will the MSAG be structured?

There are a variety of ways to represent information in the MSAG database. Remembering that what is desired is to replace the MSAG with better data from the GIS, it will probably be advisable to change how the MSAG is structured. Here is a summary of the most common ways to structure the MSAG:

1) *Traditional aggregated ranges*. This is essentially the example from above. The only change suggested here is that breaks between blocks be eliminated. But, the data still is stored as full 100-block ranges, like this:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	1	499	B
Main St	700	999	B
Main St	1000	1098	E

   This is not what I would recommend. First, there is a problem because this does not give a 1-to-1 correspondence back to the GIS data. Though there will be fewer records in the MSAG this way, which MSAG Coordinators like, it makes the job of keeping the MSAG and GIS synchronized much more difficult. If the MSAG is going to be treated as a derived product of the GIS, anyway, the need to aggregate ranges is far less obvious, so this ought not to be an issue.
   The question of how to assign the odd/even/both values is somewhat difficult. In this type of storage, there will probably need to be software, or at least some sophisticated SQL queries, to generate the MSAG records. The question of whether a particular street range is Odd or Even, in particular, is not easy to determine, and will probably require specialized processes from case to case.

2) *Single range per street segment, 'theoretical' ranges*. This is one way to get a 1-to-1 correspondence to the GIS data, and would look more like this:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	1	199	B
Main St	200	299	B
Main St	300	499	B
Main St	700	799	B
Main St	800	899	B
Main St	900	999	B
Main St	1000	1098	E

   In this instance, some of the street segments did not have breaks, presumably because there were no cross streets. The point is, however, that each individual street segment in the GIS (i.e., each row in the table in the GIS that represents the streets) has exactly 1 matching row in the MSAG. This does make the MSAG larger in the database, but with modern storage technology and processing speeds, this should not be a technical barrier.
   Getting to the O/E/B value is easier in this case, but still a little tricky. But, if street segments have 4-number address ranges with a standard for defining that a side of the street is not used,

making it either an odd or even segment, this can be managed. One solution to this that simplifies the extraction of the MSAG is to expand the listing above to look like this:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	1	199	O
Main St	2	198	E
Main St	201	299	O
Main St	200	298	E
Main St	301	499	O
Main St	300	498	E
Main St	701	799	O
Main St	700	798	E
Main St	801	899	O
Main St	800	898	E
Main St	901	999	O
Main St	900	998	E
Main St	1000	1098	E

The idea here is to have no 'B' ranges, and have only Odd and Even records. Thus, most street segments will generate 2 MSAG records, unless they are unaddressed on one side. That is, if a segment has 100-198 for ranges on one side, and 0-0 ranges on the other side, it would be treated as an Even record.

3) *Single range per street segment, 'actual' ranges.* This will be very similar to the previous example, except that what is sometimes called the actual range will be used. The actual range is the range of currently assigned addresses on the street. Here is what this MSAG table might look like for this case:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	15	190	B
Main St	205	273	B
Main St	311	455	B
Main St	700	798	B
Main St	832	865	B
Main St	901	994	B
Main St	1044	1088	E

Each of the numbers in this arrangement actually corresponds to a known physical address of a structure. All addresses in between the high and low actual address are considered possible addresses, though very few of the addresses between the high and low are usually assigned and in use.

An advantage of this arrangement is that fewer erroneous ALI records can be entered by the telco. In the previous example (2), the address 5 Main St could be entered. But, we know that there is no structure at 5 Main St, because the lowest address on Main St is 15. However, lots of incorrect addresses could still be entered, like 45 Main St, if it didn't exist as an actual structure.

Note that the parity (odd/even) values of the streets do not necessarily match up to the values in the Odd/Even/Both column, unless the segment is an Odd or Even segment. The Even segment at the end of the list has to have even addresses in the Hi and Lo ranges, because if those are actual assigned addresses, and we know it's an Even street, the physical addresses must always be even. Determining whether a particular segment is Odd, Even, or Both will depend on how the street data is derived and stored. If only 2-number address ranges are stored this way, some other information, such as an OEB attribute on the street segment, will need to be used to determine parity. More often, addresses are stored as 4-value ranges, and the values are something like:

Street Name	Left Low	Left High	Right Low	Right High
Main St	34	86	21	81

In this case the segment's low and high values are 21 and 86, and we know that it is a Both record, because we can see high and low parity values.

On the other hand, if the row looks like this, it is a little trickier:

Street Name	Left Low	Left High	Right Low	Right High
Main St	34	86	43	81

In this case the low and high values are 34 and 86. That might lead one to conclude that it is an Even record, but there are odd addresses, 43 and 81, in between the even high and low. Only if the record looked like this would we treat it as an Even record (assuming that 0 is not a valid address in this database):

Street Name	Left Low	Left High	Right Low	Right High
Main St	34	86	0	0

4) *Point-based.* This way of structuring the MSAG gives the agency the maximum control over the entry of erroneous addresses. In this scheme, each address point or parcel in the GIS data is used to generate an MSAG record, with a range of 1. Here is how that would look:

Street Name	Lo Range	Hi Range	Odd/Even/Both
Main St	15	18	B
Main St	87	87	B
Main St	190	190	B
Main St	205	205	B
Main St	246	246	B
Main St	273	273	B

Main St	311	311	B
Main St	402	402	B
Main St	455	455	B
Main St	700	700	B
Main St	722	722	B
Main St	798	798	B
Main St	832	832	B
Main St	865	865	B
Main St	901	901	B
Main St	977	977	B
Main St	994	994	B
Main St	1044	1044	B
Main St	1088	1088	B

In practice, there would be more rows than this in the table, but the point is that in this scenario, there is an MSAG record for each addressed structure. Note that this is not the same as the ALI, because there will be many structures in this table that do not have phone service. Therefore, this table will have more records than the ALI database. Note that in this example, the Odd/Even/Both column is simply defaulted to B, though it could be automatically set to O or E if desired. This might not be possible with whatever database management system is in use by the 9-1-1 agency, but is simple to change if needed.

One nice thing about this method is that, once the agency has made sure that its ALI database really is just a subset of the GIS records in its address point or parcel data, switching to this point-based MSAG can be extremely simple.

This way of structuring the MSAG can only work if the GIS department is entering all new addresses before they will be expected to apply for phone service. However, this is not all that uncommon these days, because of how many GIS departments are structured. In fact, with the increasing trend to have actual address assignment handled by GIS departments (or 9-1-1), it is not unusual for address points to be created well in advance of any possibility of applying for phone service, so the appropriate MSAG record for each new phone service request can safely be assumed to exist before the service is requested.

The choice of how to store the MSAG is partly driven by the agency's desire to have what it considers to be the best results. That means having the fewest number of incorrect ALI records as possible, and the lowest probability of incorrect ALI records being created, with a reasonable cost. The scenarios above, from 1 to 4, are arranged in increasing order of data quality. Budget considerations and the known quality of the GIS data (see below) will be big factors in making this decision. Note that the cost of maintaining each system is not necessarily greater for the different scenarios. Scenario 2 should be less expensive than scenario 1. And, in the right circumstances, scenario 4 might be the least expensive way to maintain the MSAG and ALI of any possible choices.

How good is the GIS data?

As pointed out above, a good starting assumption is always that the GIS data, assuming it is locally managed, is better than the MSAG. But, in order to convince the internal powers-that-be that this is the case might still require some preliminary work. If this is true, a rigorously controlled test of this assumption might be needed. To do this, and probably to carry out the replacement, an agency might need to modify its GIS data somewhat. The MSAG has a slightly unusual structure, which is generally not the same as what the GIS data will look like.

One defensible way to test would be to follow these steps:

1) Select perhaps 50 streets at random from the GIS.
2) For each street, using an SQL query or software for this purpose:
    a. Find the other streets that share the same name and community and ESN number. These street segments, taken together, define one or more MSAG records.
    b. Find the corresponding MSAG records in the existing system. For example, if one of the selected streets was 100-199 Main St. in Anytown, ESN# 111, search the MSAG to get the records that correspond to Main St. in Anytown, ESN #111. The system might return an MSAG record for 0 – 1000 in this case.
    c. Next, get all the streets in the GIS that match that ESN/community/name combination.
    d. Get all of the ALI records for the MSAG records.
    e. Do queries to see if the ALI records from d. all match to a street segment from c.
3) Ideally, each ALI record will match a single street segment. If there are mismatches:
    a. If the number of mismatches is unreasonably large, replacing the MSAG with street-generated MSAG records may not be feasible. The large number of errors may indicate that the GIS data just isn't ready for this project. But, a random check of the errors will further illuminate this issue. There is an outside, unlikely possibility that the errors are all in the ALI records.
    b. If the number seems OK, check the records that don't match, and see why. There might be a fairly simple explanation. Also, it's possible that the unmatched records represent errors in the ALI database. These might be records that correspond to no known address, and might be inactive at the telco, or they might be records with incorrect addresses. Either way, they are errors that should be fixed, which is a good reason to proceed with the MSAG replacement project.

Replacing the MSAG

Once the decision has been made to replace the MSAG with a version derived from the GIS, a number of questions must be answered before continuing.

The ability to transform the MSAG and ALI will be in part a function of what arrangements are already in place for MSAG and ALI database management. There are generally three ways that this is done, and each has specific implications for whether replacing the MSAG is possible, and if so, how it can be managed.

1) Service Provider

This is the most common model in use across the country. In this model, the MSAG and ALI data are stored remotely by a service provider. The agency has limited access to the data, but can submit updates to the MSAG, and view the ALI data in order to identify problems and suggest fixes. Since the ALI data is actually managed by the telcos, they also make almost all of the changes, though to a limited extent, changes may be made by the service provider when needed.

The format of the new MSAG will also have to be considered. It might not, for instance, be feasible to go to a point-based MSAG, as suggested above, due to various contractual or design issues. After all, a point-based MSAG will increase the number of records in the MSAG by at least a factor of 10. So, the decision of exactly what the new MSAG will look like might not be entirely up to the agency, and will need to be collaboratively determined with the service provider.

If the agency is using a service provider, changes to the MSAG will likely be incremental in nature. It is commonly the case that only 50 updates per day will be processed, so changing the entire MSAG at once may not be feasible. But, with planning and patience, it can be done. And, an advantage of this methodology is that there will be a second set of eyes checking out the problems and changes, which might make issue resolution simpler.

Sometimes, a complete bulk update of the MSAG might be possible. All of this is highly dependent on local factors, such as the nature of the agency's contract with the provider, or PUC regulations governing just how such changes might be made. This must all be checked out in advance, because there may be charges from the service provider to do a bulk update.

If the process is done incrementally, one advantage is that ALI errors will be identified immediately. When an existing MSAG record is replaced with a new record or records, if there are ALI fallouts, there will be immediate feedback. Then, the resulting ALI errors can be worked and fixed before continuing. Having good MSAG/GIS editing and validation software can significantly improve and expedite this process. With such software, it should be possible to identify these potential ALI fallouts before submitting the new records, so that the errors can be resolved before submitting to the database. This will lead to a smoother transition with fewer headaches.

2) In House Storage

This is the scenario where the agency hosts at least its own MSAG database, and probably the ALI as well. If the system allows, the MSAG will be able to be changed out by the agency. Many of the considerations from the service provider model still pertain here. However, the agency may find it easier to manage the process, and be able to turn it around more quickly if it controls its own MSAG and ALI.

Some of the self-hosted systems will allow the bulk replacement of the MSAG, and then the incremental follow-up process of fixing any orphaned ALI records. An advantage to this is that the MSAG will immediately reflect the current state of the GIS data, and new ALI records submitted by the telcos will be checked against this data instead of the legacy MSAG. But, the agency must still be committed to fixing all the anomalous ALI records as quickly as possible, to avoid any potential issues.

3) Regional Storage

With a regional storage system, an agency might be a subscriber or co-owner of a system hosted at a different location, and may not have full control over how data is added. But, still, in most cases, this option will allow the agency very similar options to the in house storage option. There might be some issues that are handled more like in the service provider scenario, but the specifics of that will depend on how the local system is configured and managed.

A Virtual MSAG

Having converted to a GIS-based MSAG, in whatever form it happens to take, the next step is to begin implementing maintenance. If the MSAG is truly a direct reflection of the GIS data, I like to think of it as a Virtual MSAG. I say that because it performs the same functions as the legacy MSAG, but should not require the same kind of maintenance that the legacy MSAG did. It is Virtual in the sense that it's not directly managed, but is merely an artifact of the GIS data. In fact, the Virtual MSAG should require very little direct maintenance, and with the proper procedures and/or software solutions, can be driven directly by the ongoing maintenance of the GIS data.

So, the Virtual MSAG in this new world has the advantage of still being an MSAG, and being able to connect to, be connected to, and feed all of the legacy processes that the MSAG has traditionally served in the agency. This will allow the agency to reduce data maintenance while improving data quality, and reducing the probability of ALI errors.

Maintenance of the Virtual MSAG

Having replaced the MSAG with one that is derived directly from the GIS, it is almost certainly advisable to have in place a plan and methodology for maintaining this Virtual MSAG with what comes from the GIS data. The complexity of doing so will be dependent on a number of factors, including the style of MSAG to be adopted, how it relates back to the source GIS data, and how the MSAG is stored. Though it is unlikely, it is possible that the difficulty of maintaining the new MSAG may be such that it does not make sense to do so. However, even if a well-automated maintenance method is impractical, the alternative of using traditional MSAG maintenance methods, while consulting the GIS data to determine what edits to make, is likely still preferable to not replacing the MSAG at all. However, this topic is too complex to address in this paper. I intend to create a separate white paper to address the issue of maintenance.

Conclusion

In this paper, we've explored some of the implications of coming changes to 9-1-1 database management that are implied by the new i3 standards. In particular, strategies for dealing with transitional periods that might last years were discussed. The argument has been made that in most cases, locally-maintained GIS data is more accurate than legacy MSAG data. Because of this, it makes sense to do a complete replacement of the MSAG with a Virtual MSAG, derived from and driven by the GIS. Different scenarios and strategies for accomplishing this transition have been discussed, along with possible pitfalls that might arise. Overall, I believe that for many agencies, a transition to a GIS-driven Virtual MSAG is a wise, cost-effective solution that will move them well down the road toward handling the GIS data issues that will become so central to implementing i3.