DARKHORSE

EMERGENCY

EMERGENCY SERVICES STATION PLANNING GUIDE

Learn why station planning
is needed and how to do it well.

OVERVIEW

In this guide, we look broadly at *why* and *how* to do station planning.

In part one, we explore situations where stations are not actually helping your response performance, but are in fact hindering it. Sometimes, politics and short sightedness have saddled a community with a poor station scheme so much so that closing the underperforming station is the best course of action.

Next, we look at the context around neighborhood life cycles, demographics, and how they influence the need for stations. We show how neighborhoods age and renew and how road networks influence station coverage. If done well, a station plan should remain robust decades into the future.

Finally, we connect station planning to response performance and more tangible outputs like lives saved or property damage averted. We find that adding a single station to a typical community will prevent, on average, one cardiac death every other year and reduce property damage by about one quarter million per year.

Part two looks at how to do station planning well. First, we look at forecasting in the context of planning - how much detail is necessary, how far in the future should you look, and how to address both spatial and vertical growth. In general, spatial growth drives the majority of station needs, although if there are areas of intense densification, then you may need to reconfigure stations within your community boundaries.

Next, we look at the mathematical underpinnings of station modeling. The key point is that using traditional coverage models will result in a station plan that is too sparse. Probabilistic models are the most accurate and result in station schemes with the appropriate station density.

Finally, we look at what to build. Large, monolithic structures are becoming a thing of the past. The cost of land - particularly in urban centers - is prohibitive. Satellite stations, posts, and mixed-use structures are the wave of the future, and smaller is almost always better.

DARKHORSE EMERGENCY

Emergency service stations are a critical component of any response service. Stations are costly to build, and finding land in an ideal location is a challenging endeavour. If you are like most chiefs or directors, you will hoard them like pirate treasure. One never knows when budget for an additional station will be approved - let alone when the NIMBY and political hurdles can be overcome so you can start building it.

1.1 THE COST OF UNDERPERFORMING STATIONS

WHY WOULD YOU EVER WANT TO REDUCE YOUR STATION NUMBERS?

We've done station location projects for over a dozen multi-station cities, and in almost every case, there are stations that underperform - i.e., they bear a disproportionately small share of responses.

We all know these stations. Sometimes, they're called the "sleepy hollow". They're an attractive pre-retirement assignment. Sometimes, they're quiet because the neighborhood has changed character through gentrification, but just as often, it was a poorly located station to begin with. In some extreme cases, these underperforming stations act as a drag on overall performance.

In one fairly large city, we found two stations that were located in areas with very low call volume. In spite of the fact that this particular city was severely under-stationed, we found that removing these stations saw no noticeable decrease in response performance. These two sites were close to so few calls that they wouldn't be missed whatsoever.

BUT WHAT IS THE COST?

Keeping facilities like these can cost you in ways that aren't immediately obvious. First, there is the capital tied up in the building and land. In many cases, this is multiple millions of dollars. Secondly, there is the operational cost of cleaning, stocking, heating, and maintaining the facility. This typically runs in the hundreds of thousands of dollars per year. Both of these are "opportunity costs" where this funding could be better used elsewhere in the organization. As currently deployed, it offers no value to your system.

The inefficiency of these stations can result in other costs as well. First, crews that are deployed to an underperforming station are rarely if ever the closest responder. Their skills suffer attrition and they will only be called upon when the system is busy, or when they've been flexed to another location.

DARKHORSE EMERGENCY

Even if they're dispatched to a call, their response times will be much slower than busier crews who are responding to the majority of calls in your system. Secondly, even in a busier system, crews will need to start and end their shift in these quiet stations and will thus be effectively unavailable for these intervals, further hampering system performance. Finally, these stations become a source of conflict between crews since workloads are so unbalanced.

HOW WOULD YOU USE A MILLION DOLLARS TO IMPROVE YOUR OPERATION?

Although sophisticated station models or simulations can clearly identify these stations, often they're easy to spot with a quick look at your call data. **Ask your operational analyst to give you a count of responses originating from each run area over the last month.** You may be surprised to find one or two that hasn't had a single response.
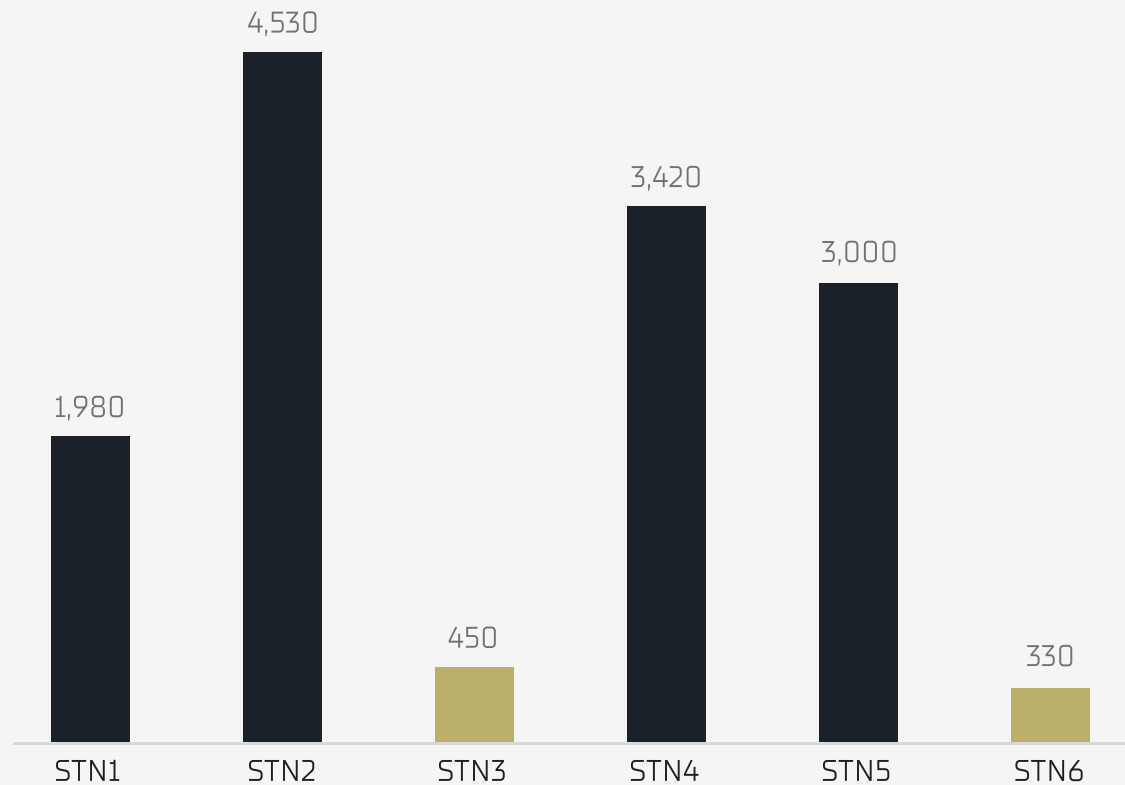
Figure 1: Comparison of existing stations.

DARKHORSE
EMERGENCY

Why Reconfigure Stations?

Call demand, road congestion, and urban sprawl are the ever-growing features of the modern emergency service system. They put pressure on your response performance and drive demand for capital and operating budgets.

Although there are many ways to address the supply and demand side, we will focus on station locations in this article. What causes station schemes to become sub-optimal?

1.2 STATION CONTEXT: WHAT DRIVES STATION PERFORMANCE

NEIGHBOURHOOD LIFE CYCLE

Neighborhood life cycle theory describes how a typical neighborhood progresses through development, maturity, decline, and renewal. The profile of the residents in a neighborhood can change substantially - especially in those areas entering renewal.

If the resident profile changes, then call volumes will change with it. Over a period of years or decades, a single neighborhood can see call volumes change by a factor of 10. Added up over a city or municipality, these changes can have a profound impact on call demand (demography has a strong effect on both fire and EMS calls).

Most importantly, the changes in demand are not uniform. In some areas, calls are increasing, while in others, they're declining. These "lumpy" changes can influence the performance (and viability) of station locations.

ROAD CONGESTION

Cities rarely grow as planned. Sometimes, economic downturns delay developments, while changes in public transit completely transform an area. Transportation planners are tasked with building roads and interchanges that will remain useful in the long-term. Sometimes, they fail. Traffic volumes grow more quickly than anticipated, or traffic patterns are different than expected.

The impact is felt on your response times, and the decisions can render your station locations almost useless. In communities with a number of bottlenecks (such as bridges or tunnels), this can be particularly problematic. Station plans may need to be revamped to adapt to the transportation reality.

LEGACY STATION ISSUES

Sometimes, stations are poorly located irrespective of demography or road networks. Some were built for political reasons. Others were added to spend excess budgets. Still others were chosen because an inexpensive property became available and a planner didn't want to lose out. As cities grow, these poor locations compound to the point where the entire system is inefficient and a hindrance to response performance.

In one jurisdiction, we found that three stations (one-fifth of the total) could be closed with virtually no performance impact. In fact, in a few cases, closing a station would actually increase response performance by ensuring crews were located closer to call hotspots.

HOW LONG DOES A CONFIGURATION LAST?

If your service has located stations in a well-planned manner, then chances are, even your oldest stations are reasonably located. We studied a quick-growing community of around one million people. Using a 10-year call data set, we analyzed how optimal station configuration changed over time.

What we found was surprising: the difference from the optimal plan 10 years ago to the optimal plan today was insignificant. This happened in spite of a heavily densifying core and rapid expansion of city suburbs. The new communities would need to add stations, but legacy stations were all still well-located even after a decade.

THE COST OF POOR LOCATIONS

There is an obvious performance cost to having poor station configuration. You're often too far away to respond in a timely manner and your response performance suffers. But there are other costs as well.

First, the workload between your stations will vary widely. Some crews will be busy their entire shift, while others won't respond to a single incident. This will lead to conflict and gaming as staff try to get the "plum assignment". Secondly, the crews at your less-busy station may not be busy enough to satisfy their training requirements. Finally, there is the opportunity cost of all the capital (and maintenance costs) tied up in underperforming stations.

DOES YOUR COMMUNITY HAVE A POOR STATION CONFIGURATION?

If you have stations that were located without any thought about response performance (whether for politics, economics, or some other reason), then your station configuration is probably dragging down your response performance. If your least busy station isn't pulling its weight, then your station configuration may be the culprit (particularly if the quiet station is not in a new suburb).

This, however, presents an intriguing proposition. You can perhaps close an underperforming station and free up the capital for a better location. This is even more compelling in cases where your deferred maintenance costs are substantial.

Why Reconfigure Stations?

If you have stations that were located without any thought about response performance (whether for politics, economics, or some other reason), then your station configuration is probably dragging down your response performance. If your least busy station isn't pulling its weight, then your station configuration may be the culprit (particularly if the quiet station is not in a new suburb).

This, however, presents an intriguing proposition. You can perhaps close an underperforming station and free up the capital for a better location. This is even more compelling in cases where your deferred maintenance costs are substantial.

1.3 THE VALUE OF RESPONSE TIME

BUT DOES IT ACTUALLY MAKE A DIFFERENCE?

It makes intuitive sense that faster is better. A quicker response will reduce the amount of structure damage in a fire. If a citizen isn't breathing or needs defibrillation, intervention within two minutes is surely better than 20. But keep in mind, these are a small subset of calls. The vast majority of fire and EMS calls do not benefit from a rapid response.

Public perception seems to go the same way. Patients believe they have received better service if there is a quicker response; homeowners count the seconds before fire suppression arrives. Even for cases where a quick response isn't strictly necessary, the stress associated with waiting for first responders is real.

SO WHY NINE MINUTES?

The original study of response times was conducted in Seattle in the 1970s by Alvarez and Cobbs. They found a lower mortality rate in patients reached in eight minutes or less than those in nine minutes or more. (Their research relied on punch clocks that didn't display seconds, so some services use 8 minutes, some use 8:59, and others now use 9 minutes). This research informed the decisions of municipalities around the world and the nine-minute standard is now nearly ubiquitous for both fire and EMS.

Unfortunately, subsequent research has called into question the importance of the nine-minute target. Several studies by Pons, Eisenberg, and others have found a more significant threshold of around five minutes for the most critical patients (cardiac arrest in particular). Studies of fire damage and response time have found mixed results. There is some evidence that flashover typically occurs around nine minutes after a fire starts, but other studies have found much faster or slower times.

DARKHORSE
EMERGENCY

Why Reconfigure Stations?

Although they don't agree on a common threshold time, they all agree that time is critical in a subset of emergency events. Faster is better.

The newer studies reveal something else. The outcomes are not binary at a particular time threshold. In other words, the house is not destroyed or the patient doesn't die once a magic response time is passed. Instead, the probability of death increases linearly and the amount of structure destroyed does the same. We can see this graphically in the next images.

SURVIVAL RATE DROPS ~1% EVERY ADDITIONAL MINUTE IN RESPONSE
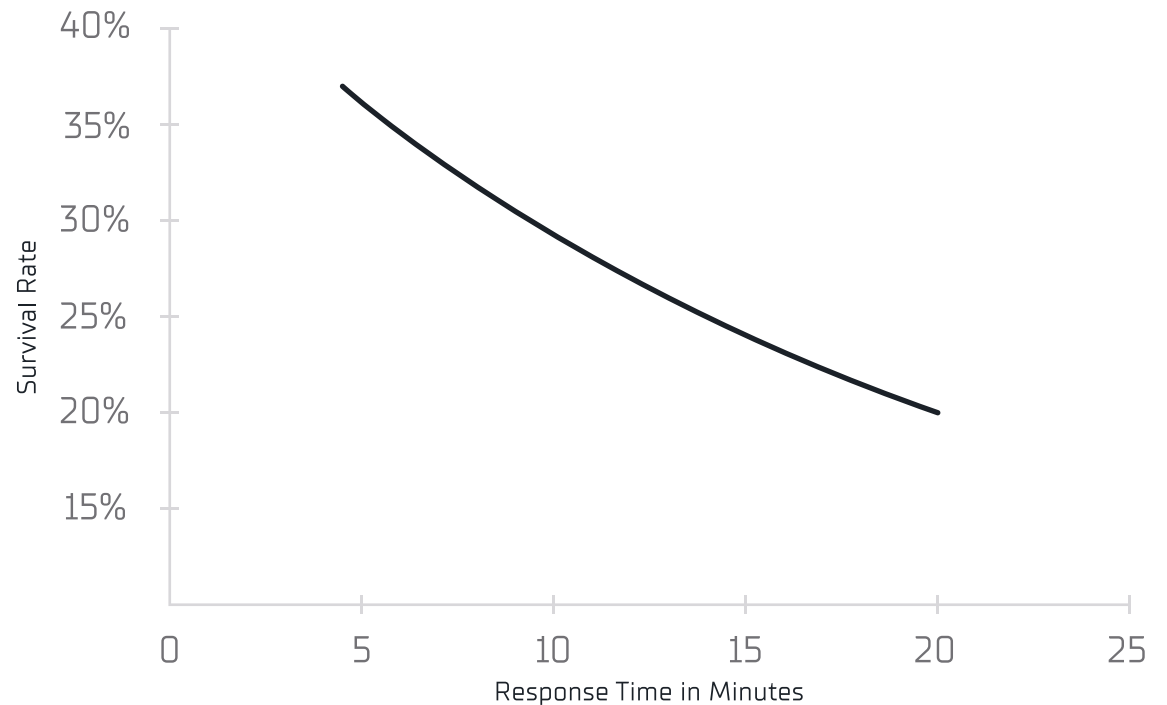
Figure 1: Cardiac survival from 911 to defibrillation taken from Karch et al study of cardiac arrests in Las Vegas casinos.

For cardiac arrests, there are dozens of studies on mortality and response times. Most report that mortality drops as response time improves, although the steepness of the relationship varies substantially. Nevertheless, every minute in response time means more lives saved.

As to fire loss and response times, there have been fewer attempts to prove the link, but the findings of the New Zealand paper are in line with the controlled experiments of fire spread and flashover rates. Every minute is worth thousands of dollars in property damage.

What is most interesting about these studies is to use the results as inputs in the station planning process. For example, it is possible to estimate how much property damage can be averted by adding or moving a fire station. Similarly, we can see how many lives could be saved by adding a single ambulance station.

PROPERTY DAMAGE INCREASES BY ~2% FOR EACH ADDITIONAL MINUTE OF RESPONSE TIME
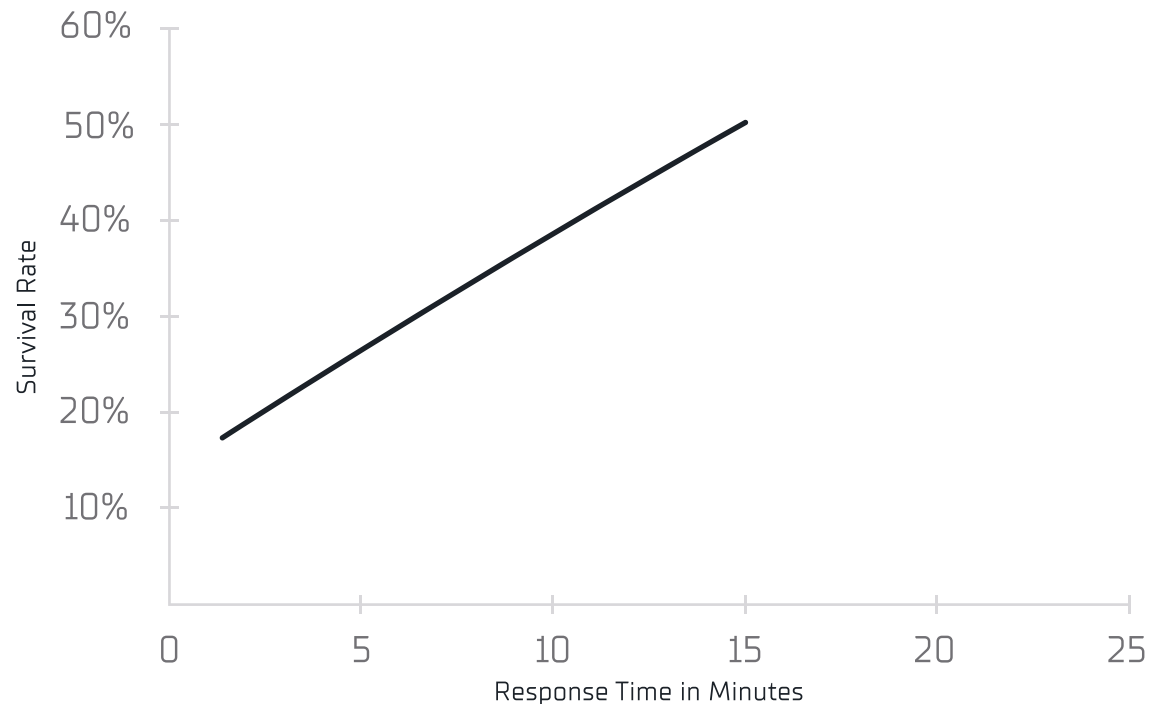
Figure 2: Value lost in structure fires as a function of response time. Adapted from Challands' study of New Zealand fires published in Fire Technology 2010.

Why Reconfigure Stations?

In a study we did for a moderate-sized suburban community (population ~50,000), we found that adding a third combined fire/EMS station would save $75,000 per year in property damage and save one additional cardiac patient's life every year.

The charts to the left show this effect through time.

Although the results will vary significantly depending on the size and type of community, it is clear that there is value in translating the impact of resourcing decisions into a currency that is readily understood.

So, what's the bottom line? **Improving response times will indeed move the dial on lives saved and/or property damage averted - even though this is driven by an extremely small subset of calls.** With a little work, you can approximate the relationship between response performance and more concrete measures like lives or dollars. While "an improvement from 87% response performance to 90%" may sound good, "25 lives and two million dollars" has a far more relatable impact.

ADDING STATIONS REDUCES EXPECTED CARDIAC FATALITIES
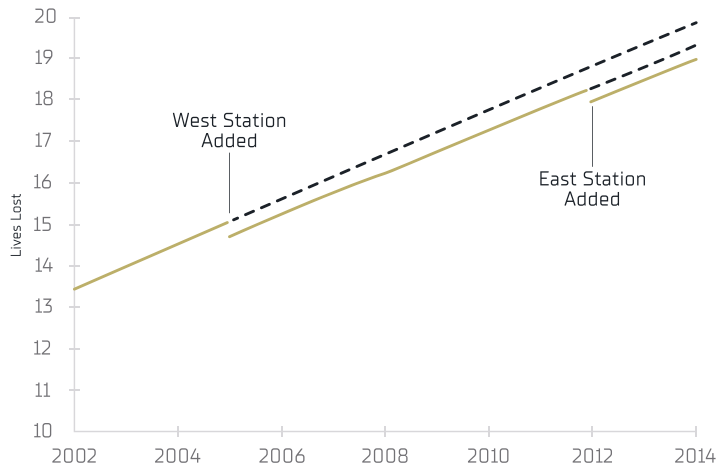
Figure 3: Impact on cardiac fatalities of adding additional stations. Adapted from a study for the City of St. Albert, Canada.

ADDING STATIONS REDUCES FIRE LOSS
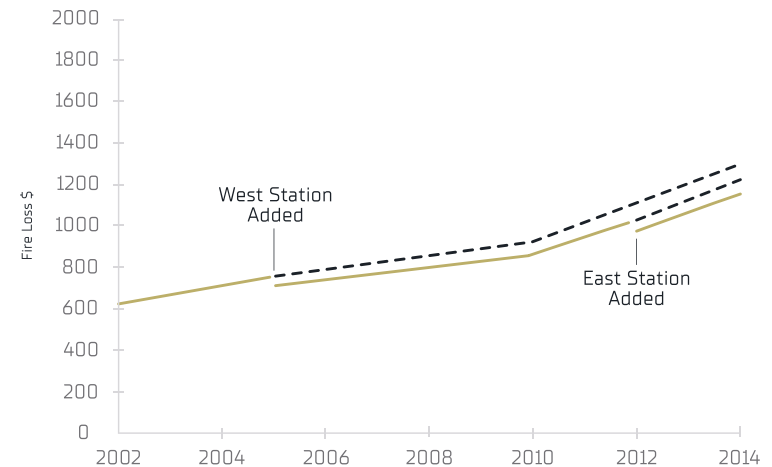
Figure 4: Impact on fire loss of adding stations. City of St. Albert study.

To learn more about how to define your "why", explore Darkhorse Diagnostics.

[SEE PLATFORM]

DARKHORSE
EMERGENCY

How To Do Station Planning Well?

Forecasting is a time-honored tradition in emergency services. It is used to support strategic plans, justify budgets, and prove to citizens that you're not just reacting to emergencies. Forecasts can also be used to adjust resources in the short-term, or to target more effective inspection efforts.
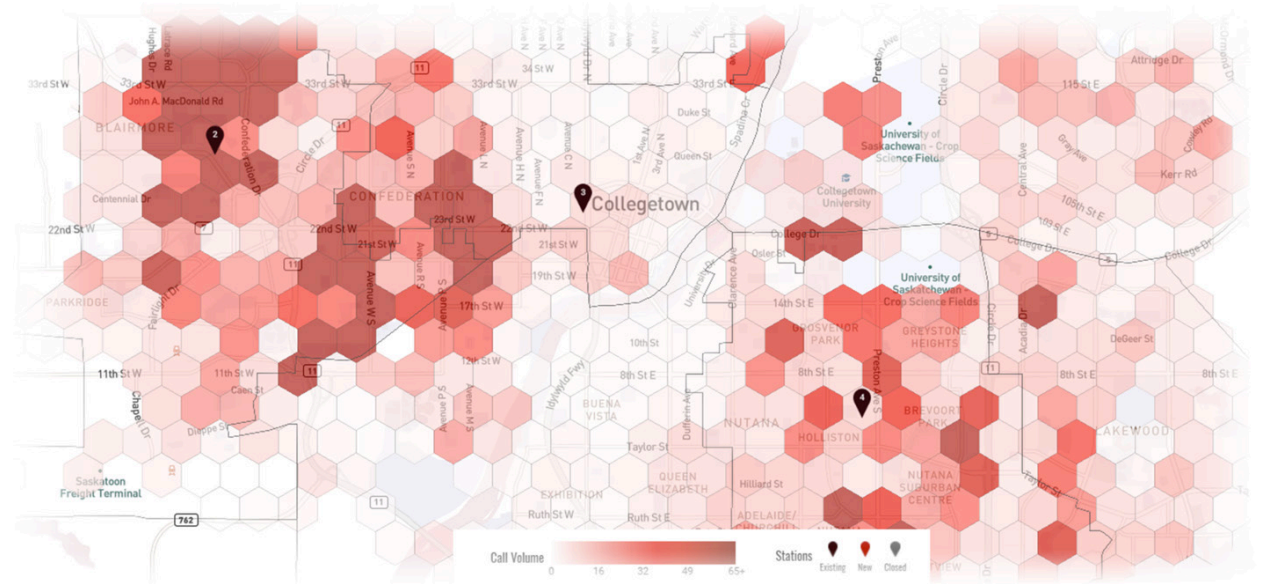
Forecasts are always wrong, but that doesn't mean they aren't useful.

When done well, they support planning and prevent costly mistakes on capital investments.

2.1 SPATIAL CALL FORECASTING

CALL VOLUME FORECASTS

Creating a long-term call volume forecast for a jurisdiction is reasonably straightforward. Call volumes are highly correlated with population levels and historic patterns. Over the long haul, call growth will stabilize around a rate per population. Simply extrapolating population growth and multiplying by the current call rate per population gives a sensible result for a call forecast.

It is possible to refine this further by considering daytime population, income levels, expected age, and transit population (vehicles passing through a jurisdiction), but the improvement in accuracy typically doesn't justify the effort. Benchmarking against similar communities to compare calls per station or calls per staffed vehicle can also provide a rough guide for capital planning.

How To Do Station Planning Well

This tells you if the number of stations you have is reasonable and how many you should expect to build as the system grows. What it doesn't tell you is where you need to build them. To have a forecast that's useful for station location planning requires much more effort. Communities don't grow in concentric circles from the downtown core. Land is developed in clumps and dependant on job prospects, utility access, transportation, and other factors. What's needed is a spatial forecast. You want to know where all those calls are going to originate.

Fire, ambulance, and police call drivers are persistent. The same neighborhoods and even the same buildings continue to drive call volumes year after year. (In fact, in one large city, a single bar accounted for 1% of all ambulance calls. The city eventually forced the bar to close). This is good news for planning because there is some stability to the location and volume of calls and thus stability to the station plan.

As a city grows - either vertically or horizontally - you will have to account for these new potential locations. Here the development plans and zoning will become your friend. Commercial and industrial developments tend to have similar call rates per square foot of area (a proxy for population visits or employees).

Residential areas are driven off of population. **Look at the expected population for a neighborhood when fully built out.** Your station plan is for the long-term, not immediate needs. **We recommend using a rate based on your current residential calls.** Even though new housing developments will have a certain profile of residents (perhaps young families), neighborhoods change over time. This approach works in new neighborhoods or in densifying neighborhoods.

The one wrinkle to consider is retirement centers. Although a typical neighborhood will start out with young families, move to empty nesters, and then return to young families, clusters of retirement centers may form. **These tend to drive calls at a much higher level than the typical population, and they are persistent.**

The best approach looks at specific drivers of calls (retirement centers and large commercial developments), call growth, and existing hot spots. Together, these will paint a pretty clear picture of the future. Just make sure to update it periodically.

See how **Darkhorse Emergency** helped Strathcona County implement use volume forecasts to select new station locations.

READ CASE STUDY

How To Do Station Planning Well?

Traditionally, station planners have relied on the concept of "coverage" to determine the best locations. In simple terms, each geographic area of a certain size needs to have a station in order to be considered covered. This works reasonably well from a resourcing standpoint (having enough response vehicles to address an area), but it introduces some problems.

2.2 STATION MODELING: COVERAGE OR RESPONSE PERFORMANCE

WHAT DOES COVERAGE MEAN?

If a call location can be reached within a target response time, the demand point is considered covered. For example, if a call is 7:59 minutes away on average, it is covered, but if it is 8:01 minutes away, it is uncovered. Unfortunately, this can bias your station spacing - typically encouraging a station plan which is insufficiently dense. Your stations end up too far apart.

The other issue with using coverage is that it has no connection to the actual fractile performance that you report. An area may be "covered", but only have a response performance of 75%. If your plan promises a coverage of 95% but your operations only deliver 75%, you may have to answer some tough questions. As we saw in the section on response performance, real life is not binary.
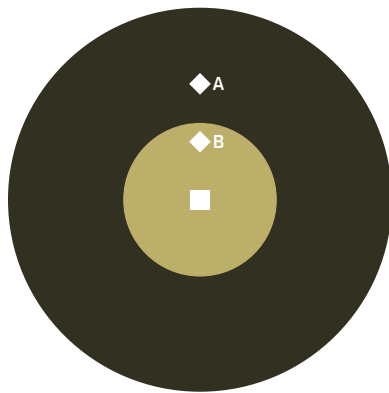
THE POWER OF VARIABILITY

Consider, for example, figures 1 and 2. In both figures, the centre of the circle is the location of the emergency services station and the two diamonds are call locations A and B.
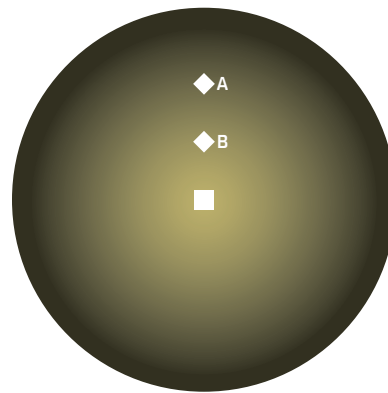
DARKHORSE EMERGENCY

How To Do Station Planning Well

Figure 1 is a representation of a Binary Coverage Model. The dark circle represents the area around the station that is "covered" by that station in 4 minutes. The light grey circle is the area around the station that is "uncovered". As you can see, with the Binary Coverage Model, all the calls from B would be considered "covered," whereas all calls from A would be considered "uncovered".

Figure 2 is a representation of a Probabilistic Coverage Model. In this figure, we represent response time as a gradient. Call locations closer to the station have improved coverage compared to those that are farther away. Under this model, location B is still mostly covered (but not entirely covered) whereas location A is somewhat covered (as opposed to not at all in the binary model).

So which approach best approximates reality? Consider the following figure. Over a five year period, there were over 300 responses to a single address (a large retirement centre) from a single station. The travel time of each vehicle run is plotted in **Figure 3.**

Figure 1: Binary Coverage

Figure 2: Probabilistic Coverage
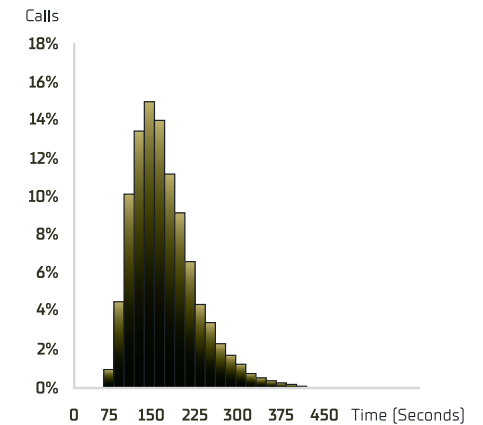
Figure 3: Travel time of each vehicle run.

Right away you can see the wide range of response times. Sometimes the vehicle reaches it in just over a minute (75 seconds), sometimes in just over six minutes (375 seconds). Although the average may be three minutes, for any single response, the drive time can vary immensely.

Probabilistic coverage models recognize this variability around drive times. And because of this, you can use these models to calculate an expected response performance before you build a station. The benefit of this approach is a more accurate prediction of performance, and better spacing of resources.

DARKHORSE EMERGENCY

How To Do Station Planning Well?

There is a conundrum in station planning. The largest stations with the most space for crews and response vehicles are actually needed in the place least able to accommodate them. Land in the busiest parts of the city (typically the downtown core) is at a premium. Innovative services have begun to use creative means to both improve performance and contain costs.

2.3 TOWARDS SMALLER STATIONS

A WELL-DESIGNED AND LOCATED STATION IS A THING OF BEAUTY

Many of these become iconic structures, integral to a vibrant community. Fire and EMS halls are celebrated as one of the foundational pieces of a successful municipality. They're a tangible reminder of how the community serves its citizens.

For fire and ambulance crews, stations are a place to eat, rest, exercise, and connect with colleagues. They're in many cases a home away from home. In fact, many members will spend more time at their station than they do at home.

A station is a place to store equipment, house vehicles, and train teams. More importantly, it's a base from which to respond to emergencies in a timely manner. It's an asset that is crucial to meeting a service's mandate. Unfortunately, fire and EMS stations are the most capital-intensive component of the budget. They're difficult to site, expensive to build, and costly to maintain.

The good news is there are alternatives that have been used by services to maintain the benefits of traditional stations without all the cost. The secret is in becoming smaller.

SATELLITE STATIONS

Fast growing communities are a challenge. When partially developed, they may require an emergency presence, but the available land may be far from the population base. What's needed is an interim location to house fire and EMS personnel while the area develops and grows.

Some services have had success in using "satellite" stations to bridge this gap. A satellite station is a small, first response base located in a residential area. Usually, it's a house on an undesirable and busy intersection. The service purchases and retrofits the house to meet the needs of the fire or EMS crews that will be staffed there and uses it as a forward base embedded deep in the community.

This approach provides a number of benefits. One is the greater variety of locations to choose from. The service is no longer limited to municipally-owned land or appropriately zoned locations. Planners can choose better locations with better call proximity and associated response performance. The cost, even after the retrofit, is a fraction of a typical station build. Finally, as the community grows and the station location is no longer ideally situated, the structure can be sold and reverted to a residential property.

There are challenges though. Zoning hurdles and community opposition should be expected, but the cost and performance differences are real enough to convince many.

Image 2: Single start station in Auckland, NZ.

There are a number of benefits to this model:

- Cleaning, maintenance, and inventory are centralized.
- The service can reduce costs and improve service from the economies of scale.
- The capital costs - especially when using posts - are much lower: a single building and a single piece of property.
- In theory, it should also reduce overtime or delays because vehicles and equipment are pooled centrally.

In practice, however, there are some shortcomings:

- In pure post models (with only a single central station), crews don't have a "place" to work at. They're always in their vehicle. This can affect health, morale, and the sense of camaraderie that develops at a fixed location. It can also be uncomfortable in particularly cool or hot climates.

Image 1: Temporary station in North Lenoir, SC.

SINGLE START STATIONS

On the extreme end is the single start station model. There are a couple permutations, but in general, staff come to a central location to gear up and collect their vehicle. They then drive to the depot - either a scaled-down station or a simple street address. They respond from that location until the end of their shift.

- Secondly, the drive from the central station to the post now happens during the crew shift. For larger cities, it can be over an hour of transit time from depot to central station and back.

- Coverage may also be exacerbated by locating the central depot far from the city center and its associated higher call volumes. Crews will spend a lot of time driving through areas with very little response demand. As a result, a typical ambulance service might see a 10% reduction in their effective unit hours when moving to single start.

STOREFRONT STATIONS

The monolithic fire hall or ambulance station is becoming increasingly rare. The cost to build and maintain a large station is prohibitive. And from a pure performance standpoint, you're better off having many small stations as opposed to a single large one.

Image 3: Alexandria, Virginia. Station 209.

The storefront or mixed-use model is becoming increasingly popular. In Alexandria, VA, the Station 209 fire hall includes a typical station layout, but it also contains commercial space and four storeys of residential apartments.

There are many examples of stations sitting on land which could be sold for tens or even hundreds of millions of dollars. Mixed use allows this capital to be used more efficiently with no degradation in service.

CONSOLIDATION VS. PATIENT-CENTRIC

To combat the ballooning capital costs of traditional stations, many communities are consolidating their services into large, mixed-service facilities. They're co-locating ambulance, police, and fire in large, public structures.

Image 4: Hastings, MN combined fire and EMS station.

On the surface, this seems like a good idea. The services have similar characteristics (lights and sirens), they share a dispatch center (or should) and they do similar training. In theory, there are economies of scale that come from putting similar services in the same location.
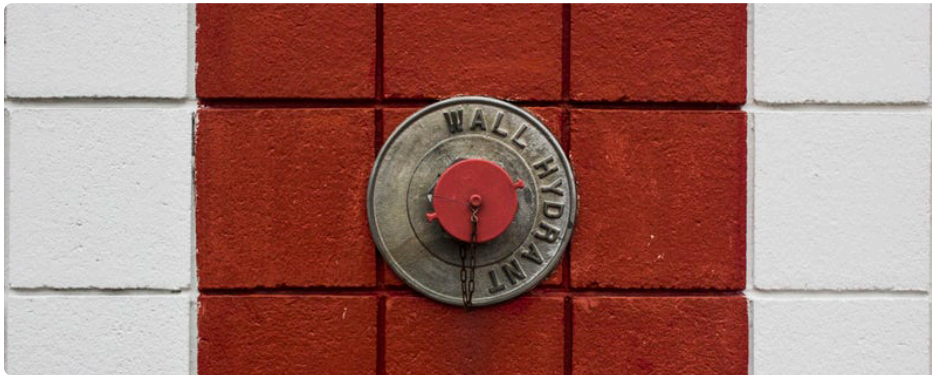
How To Do Station Planning Well

In practice, the results vary. Police, fire, and EMS have their cultural differences and may not "play" well together. Arguably, the biggest issue is that consolidation gets rid of the main benefit of having embedded stations: proximity.

Consider the following scenario: An elderly patient is unconscious and not breathing at the grocery store. She doesn't care if a fireman, a police officer, a paramedic, or a first-aid-trained bystander arrives first. She only cares how quickly life-saving measures can be administered. **Each second reduces her chance of survival.**

If each public safety agency is housed in a different location, one of them will almost certainly be closer to her than if they're all in the same facility. In fact, from her perspective, having only a single vehicle in each station with lots of stations is ideal. A patient-centric system would do precisely this - spread resources as widely as possible.

The cost of such a system is untenable with the current large station model, but hopefully, the examples above provide a way toward the ideal.

Darkhorse Deployment is designed to help emergency services do station planning well. See how.

EXPLORE PLATFORM

DARKHORSE
EMERGENCY

Station planning draws on a wide set of skills and resources. You must navigate the politics, wrestle with the budgets, and anticipate the future. And there's still no guarantee your plan will be approved.

DARKHORSE
EMERGENCY

CONCLUSION

From our perspective, the key points are as follow

- Poor locations cost you in more ways than you probably realize
- Well located stations tend to remain well-located for decades
- Although changes in a neighborhood's lifecycle affect performance in the short term, build for the long term
- Expected performance is a reasonable metric to evaluate station performance
- Projecting call volumes and road networks is important to keep you from painting yourself in a corner
- When using models, account for the variability of travel times or your station spacing will be wrong
- Think small and mixed use to make the most of your scarce capital
- Avoid co-location to be more patient-centric

We hope this e-book has given you some of the tools you will need to succeed at this process. At the very least, we hope you now have a deeper understanding of best practices and the breadth of considerations that go into doing it well.

To learn more about station planning or to see how the Darkhorse Emergency Decision Analytics platform can help navigate your planning needs, contact us today.

SCHEDULE A DEMO