Risk Managed Demand:
Phase One Technical Brief

The City of Seattle is committed to providing comprehensive public safety services. In response to Statement of Legislative Intent SPD-017-A-001, the Seattle Police Department (SPD) undertook a project to develop a model to predict outcomes of Calls for Service (CfS). This report documents Phase One of the Risk Managed Demand (RMD) project. Nearly 2 million records of police response between 2017 and 2022 were considered in the construct the Phase One risk matrix. The resulting risk scores and their related response tiers were used to generate staffing estimates from three years of normal operation (pre-pandemic)[1] response data (2017 – 2019)[2].

These analyses are intended to inform capacity building efforts in advance of implementation (Phase Two) of a call triage system sufficient to operationalize the response model. Notably however, and as has been presented previously, the purpose of RMD is not to identify a discrete list of call types that can be "diverted" from police. Beyond a basic list of low-level calls for which little to no risk can be assumed[3], it is not possible to predict the outcome based simply on how a Call for Service is initially categorized.

Rather, this phase establishes a framework for risk management. The risk matrix (risk assessment) and reclassification (risk management) of call response tiers in support of an ideal response ecosystem (alternate/co/differential response) is the first step toward an ongoing cycle of active risk management. SPD is the first department to engage in this more sophisticated form of analysis.

Importantly, this analysis focuses on risk: except for some manual reclassifications needed where the results were incomplete, necessary overlays of legal requirements (*e.g.*, mandatory arrest DV), practical considerations (such as how many resources the City might want to divert to marine responders), and labor obligations are not included. This matrix provides the context to explore what is possible.

Overview:
RMD is the natural evolution of the effort, in collaboration National Institute for Criminal Justice Reform (NICJR), to identify calls which can be safely handled by a police alternative. Since the early 20th century, American police operations have been driven by a need to respond to CfS. In the latter half of that century, a series of efforts to reform the Social Welfare State (SWS) (*e.g.*, deinstitutionalization of mental health), resulted in a dramatic expansion of the types of emergency requests made of municipal governments.

[1] Pre-pandemic response data was used to generate workload estimates on the assumption the response environment will return to baseline levels. Shifts in the use of public spaces and the effects of staffing on response capacity limit the utility of pandemic and post pandemic data for modeling.
[2] This was the same period assessed in the NICJR analysis.
[3] For example, Denver STAR responds to assist, intoxicated person, suicidal series (without weapon or self-harm), welfare check, indecent exposure, trespass unwanted person and syringe disposal.

Given the risks inherent in these requests, municipalities have elected to deploy police as *all hazards* responders, accepting the risk to responders and collateral harms resulting from an imperfect response. Although death and serious injury are relatively rare, the high profile and catastrophic nature of these outcomes has caused a broad reexamination of response to CfS, particularly in the wake of the murder of George Floyd. RMD seeks to establish an objective, transparent, and reliable framework for risk management, as well as a technical process for call triage—right sizing the response.

Since the NICJR report was produced, pursuant to Executive Order 2020-10, the SPD has 1) designed the RMD method,[4] 2) contracted a technical and consulting services vendor[5] to implement the call triage system ($750,000), 3) coordinated with partners (SFD[6], ITD[7]) to develop prototype data in an appropriately regulated environment,[8] 4) in coordination with Community Safety Communication Center (CSCC), identified a vendor who can support the development and operation of a call triage system sufficient to support the RMD tiered response model, and 5) generated staffing estimates (presented here for the purposes of informing capacity building).

In the coming months, the SPD will work with its partners to develop a first of its kind Natural Language Processing (NLP)-based call triage system. This Phase Two work will provide additional information that call takers at CSCC will use to right-size the response. While RMD will provide guidance, it WILL NOT override the human decision-making necessary to protect life and safety. The statistical models developed under RMD will be made available, publicly, as open-source coding for transparency and critical review. The collaborative work to develop and operationalize the technology represents a giant leap forward for the industry, meeting a critical gap in call triage identified by the Transform 911 project as a barrier to alternate response.

Data:

The data underlying RMD was developed by relating nearly 2 million police calls to approximately 350,000 records of aid response by the SFD to generate a severity score.[9] Of note, the systems used to track police and fire responses do not relate to each other. When police require a fire response, a call must be made to the fire alarm center, and vice versa.

No native relationship between calls exists. To relate these calls, the SPD designed a novel process to identify responses occurring at the same place (within 200 ft) for overlapping periods of time (1/10th

[4] The NICJR analysis is based on a plan language reading of the description of the final call type of the call and does not accurate reflect the multidimensionality or risk of the event. On the advice of noted policing expert Dr. Geoffrey Alpert, the SPD designed the RMD based on objective criteria for severity and likelihood. Dr. Alpert's opinion was presented to City Council along with the NICJR report and is available upon request.

[5] Accenture was contracted under Scope of Work 10. Work began in July 2022 and is schedule to conclude in February 2023.

[6] Seattle Fire Department

[7] Information Technology Department

[8] The RMD data requires calculating a "severity score" based on the existence and type of aid response by the SFD. This score must be calculated in a secure and compartmentalized environment. ITD acts as an authorized third-party to calculate the score. SPD personnel do not have access to identified SFD aid response records. An MOU memorializing and authorizing this work, approved by all parties involved, is available upon request.

[9] Responses to events resulting a documented Use of Force were removed from the analysis and do not factor in the likelihood calculation. Severity scores are generated independent of the police response. Injuries documented in an aid run are not the result of the police response.

of an hour increments). From this relationship, the names of the people involved are matched using a probabilistic technique referred to as *fuzzy match*.

No known sources documenting the number of police responses involving a fire or aid response exists. To evaluate the performance of the matching technique, verifiable, high severity (death) outcomes were compared across systems. Independent analysis of the outcomes across both systems indicated an 85% match rate. Further examination of the nature of failed matches indicated not all these outcomes will be represented across both systems. Police and fire are not always required to respond to each death.

Estimates based on observation of failed match conditions suggest a match rate greater than 90% and were deemed sufficient to pass testing from the prototype phase. The computer code developed to produce this match will also be shared publicly, so that other jurisdictions might replicate this work. In total, 165,917 aid response records were related to the 1,991,498 police responses evaluated over the 5 1/3-year study period (2017 to May 2022).

***RECOMMENDATION: Fully** integrated (technical) police and fire call management systems (e.g., Computer Aided Dispatch) would eliminate the need for this imperfect process and further improve the accuracy of call triage.*

Risk Management:

Few enterprises exist without risk. Responsible management often comes down to a process of identifying, mitigating, and accepting those risks necessary to complete the mission of the organization. Beginning in 2006, for example, the International Civil Aviation Organization required most commercial operators to implement a Safety Management System (SMS), of which risk assessment/management is a critical function. Like commercial aviation, policing has a very low tolerance for failure. The death or injury of a community member or responder represents a significant event for the organization, much like a plane crash.

By evaluating the history of the organization, against some objective criteria, a confluence of severity and likelihood can help guide the organization to make risk management decisions. As proposed in the NICJR analysis, a dimensional scaling of the more than 300 call types currently in use can provide response protocol guidelines, informed by risk management (see Table 1, below). Based on whether the risk, whatever that may be, is acceptable to the organization, policy prescribes a course of action. Table 1 (below) demonstrates risk management policy, operationalized by a risk matrix. For events coded high risk by the organization, policy defines the risk to another than all hazards responder to be "unacceptable" and so a police response is necessary. Risk classified as Tier 2 or Tier 3 indicates the risk is acceptable, with mitigation.[10] Tier 4 responses pose no particular risk to responders or those involved. This area is safe for a deferred or alternative response (see Differential Police Response, below).

[10] Although the literature is insufficient, it is reasonable to assume that in some circumstances the presence of police may represent a risk of escalation. For this reason, alternative response such as STAR and CAHOOTS are believed to be effective risk mitigation. Under the tiered response model, STAR and CAHOOTS type resources would full under Tier 3. In these instances, the response ecosystem (e.g., police and fire) would be aware an alternate response was in progress (coordinated via a Common Operating Picture such as Computer Aided Dispatch) and may stage nearby, able to rapidly intervene, but would not be in attendance. Both CAHOOTS and STAR also utilize Tier 2 co-response options.

TIER	DESCRIPTION	POLICY
TIER 1	Police Response	High Risk – Police Response
TIER 2	Co-Responder Assisted Police Response	Medium High – Must Mitigate
TIER 3	Police Assisted Co-Response	Medium – Should Mitigate
TIER 4	Deferred Response	Low Medium / Low – No Mitigation

Table 1 – Tiered Response Model

For this project, the definition of severity is fixed and determined by objective criteria (see Table 2, below). The plain language reading of the CAD event type, used in the NICJR analysis, is limited by both the third-party rater, who may not be familiar with the use of the code, and the unverified nature of the coded outcomes (CAD final call types do not reflect the presence of a corroborating record, e.g., arrest report). The severity coding employed by RMD is dependent on an objective and verified record, outside the police data system.

SEVERITY	DESCRIPTION
SEVERITY 5	An event with an associated, unnatural, death.
SEVERITY 4	An event with an associated injury requiring the ALS[11] (*i.e.*, paramedic) standard of care and transfer to a medical professional.
SEVERITY 3	An event with an associated injury requiring the BLS[12] (*i.e.*, EMT) standard of care and transfer to a medical professional.
SEVERITY 2	An event with an associated injury requiring care but no transfer of the patient to another medical professional (*e.g.*, first aid).
SEVERITY 1	No physical injury to a person.

Table 2 – Severity

All severity coding is based on the fact of a fire aid response at a particular level (*e.g.*, ALS, BLS, etc.) and does not include injuries that are the result of the response. For the purposes of this analysis, all CAD events related to a police Use of Force (UoF) or where a police officer was reported as the injured party are not considered by the analysis. The severity coding here is not a perfect representation of what might happen absent a police response but is an acceptable proxy for the seriousness of the event.

While society is currently contemplating low-level enforcement as a strategy for social control (e.g., misdemeanors and infractions), there is little disagreement that government has an imperative to act in defense of life and safety.[13] For this reason, any CAD events related to a NIBRS reportable "crime

[11] Advanced Life Support
[12] Basic Life Support
[13] While an imperative to respond to reported crime and active threats to life/health/safety is critical to trust and legitimacy, many differential police or alternative responses have been proposed. This framework allows for a risk managed approach to implementation of differential/alternative response. https://www.transform911.org/blueprint/chapter-7-response/

against person" offense are likewise removed from the analysis. Response data for the same three-year period used by NICJR was used in this analysis. In total, 727,423 dispatched responses were considered for the years 2017 – 2019.

Calculating likelihood is less defined. Several complex statistical techniques might serve to reflect natural relationships in the data; however, they are likely to be overly complex, and challenge the transparency of the RMD program. For each final call type, the rate at which the most severe outcome occurred was calculated and expressed as a percentage. Likelihood is coded from Very Unlikely (1) to Very Likely (5) on a five-point scale (see Table 3, below).

LIKELIHOOD	DESCRIPTION
VERY UNLIKELY	Occurring less than or equal to 1% of the time.
UNLIKELY	Occurring less than or equal to 5% of the time.
POSSIBLE	Occurring less than or equal to 10% of the time.
LIKELY	Occurring less than or equal to 25% of the time.
VERY LIKELY	Occurring greater than or equal to 50% of the time.

Table 3 – Likelihood

Table 4 (below) reflects the RMD risk matrix in its pure analytical form. In this *risk assessment* each combination of severity and likelihood classifies the call types according to the tiered call response model (above). The number of call types classified is summed in each cell (e.g., Severity 5 and Very Unlikely = 68). Clusters of cells are color coded to define the response tier.

	Sev 1	Sev 2	Sev 3	Sev 4	Sev 5
Very Likely	114	0	2	1	0
Likely	0	0	1	1	3
Possible	0	0	5	4	1
Unlikely	0	0	14	2	0
Very Unlikely	0	2	51	87	68

Table 1 – RMD Risk Matrix v1[14][15]

[14] Color-coding of the risk matrix includes four major categories of response. Within the fourth classification (Tier 4 – green), a subclass of risk can be accommodated by incorporating a distinction between high and low harm offenses.

[15] This matrix reflects the original response tier coding, without recode (see below). Counts per cell are distinct call types.

Recoding:

Risk management is based on the initial risk assessment. In this operation, the original risk assessment is managed to produce an ideal form. In commercial applications, risk management is used to mitigate unacceptable risks blocking actions or behaviors necessary to the business. In service applications the ideal form is risk managed. Where the service prefers a more appropriate service provider but cannot accept the unmitigated risk to that provider, a mitigation is necessary to achieve that end state. For example, where a Mental Health Professional (MHP) is the ideal responder (*e.g.*, crisis events), risk to the responder can be mitigated by a co-response. Under this risk managed response, a resource capable of intervening to protect life might be assessed as a reasonable mitigation, sufficient to render a lower risk classification.

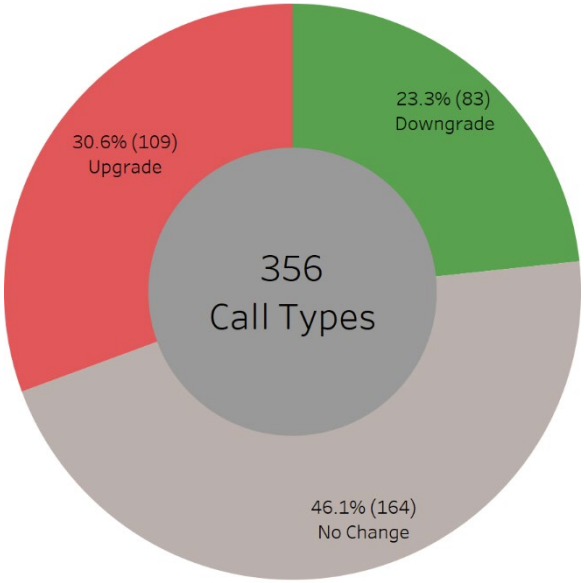

The Blueprint proposed by the Transform 911 project outlines a "significant investment in a diverse ecosystem of response options so that callers can be met with the right response at the right time." The initial risk assessment (see Table 4, above), represents the demonstrated risk in a police only response ecosystem. The act of risk managing an alternative or co-response examines the assessed risk against a reassessed state given a reasonable mitigation (*e.g.*, co-responder assisted police response). The recoded response tier represents this mitigated state.

Of the 356 call types evaluated, approximately 54% were recoded because the analysis did not necessarily reflect the desired or ideal state (mitigated). Of the 54% that were recoded, 23% were downgraded to a lower response tier. A reading of the call type suggested the subject or type of service may benefit from a right sized response. In nearly 31% of call types, the response tier was upgraded because they were either a specific request for a police response (*e.g.*, "SFD – ASSIST ON BLS FIRE OR MEDICAL RESPONSE" Tier 1) or were a call type requiring enforcement action, up to and including arrest, to resolve (*e.g.*, "LIQUOR VIOLATIONS – MINOR" Tier 2).

In some cases, not enough observations were present in the data to render a reasonable assessment of likelihood. This is known as "out of sample" and is a common limitation cited as a root cause underlying *why some predictions fail.*[16] For this reason, call types with only one or no maximum-severity outcome during the study period were evaluated manually.[17] Where the call type was an obvious police matter (a specific call for a police officer or unique police asset), the response tier was recoded. After this initial manual adjustment, the remaining call types were reviewed and where the response could conceivably benefit from a co-responder present (*e.g.*, taking a report of a completed

[16] Silver, N. (2012). *The signal and the noise: Why so many predictions fail-but some don't*. Penguin.
[17] The *washing out* effect observed here is a limitation of the over-granulation of call types which has occurred, organically, over the more than 14 years the system has been in service. A lack of data governance (gap resolved pursuant to the consent decree) and strategic guidance on the use and maintenance of call types has resulted in an interrater reliability problem between call takers and across response resources. Fidelity to the call classification model cannot be assumed with a broad and overly diffused coding scheme. See recommendation for call type cleanup, below.

crime), or an alternative response (*e.g.*, dead animal disposal, down powerlines) the response tier was manually recoded (see Figure 1, below).

RECOMMENDATION: *Throughout this process, the scale of call type classifications has been noted as unreasonably detailed and not an effective means of call classification. SPD is working with CSCC to reduce the scale and sometimes redundant nature of these call types.*

Call Types:

Why not isolate those call types, by description, which do not require a police response? Put simply, managing alternative or co-response is a function of risk mitigation. As outlined in the tiered response model proposed by NICJR, while there may be significant opportunity to provide the right resource for a person in crisis, the potential for harm to the responder or others must be mitigated.

The call type is insufficient to represent the total risk of the event, as evidenced by the limitations of the current call classification scheme, the discretionary and often one-to-many relationships that can occur with call outcomes, and the nearly 41,000 permutations of calls observed across initial call type, final call type, call clearance and how it was received. To leverage a more reliable data source, a label must be applied from a dimensionally scaled view of the call type (*i.e.*, call response tier). Phase Two will leverage this label to forecast risk from the words spoken by the person requesting assistance. While the final call type, in this analysis, represents critical context, it is not the complete analysis.

Compounding Disparity:

The academic community is rightly concerned about the use of police data in the development of advanced analytics and the extent to which inequalities present in the criminal justice system generally may be perpetuated in analyses generated from criminal justice data. The use of data to direct future delivery of police service thus comes with very real risk of compounding disparate effects by inappropriately focusing police resources on a continuous feedback loop. RMD mitigates this risk by relying in part on outcome measures independent of the delivery of police service (*i.e.*, SFD Aid Runs). Similarly, while death or injury are objective and independently verifiable outcomes, sometimes a police response can result in one of these outcomes. For this reason, all CAD events including use of force are eliminated from the count of adverse events. As such, the calculated "likelihood" of an adverse event (severity levels 2 – 5) is decoupled from the biasing effect of the police response.

It must be acknowledged that this analysis can only represent what has happened. When the intention is to radically disrupt the emergency response ecosystem, a direct measure of the risk of physical harm cannot be achieved. Instead, the outcomes utilized by this analysis are proxies for dangerousness. This initial estimate is assumed to be imperfect. As RMD is implemented and begins to inform the deployment of co/alternate response resources, the resulting outcomes will be tracked as a measure of system efficacy, in keeping with the "triple – T's" of Evidence-Based Policing (EBP).

Staffing:

Phase One of the RMD analysis suggests (see Figure 2, next page) 182 call types (approximately 51% of call types, or 73.16% of service time) are appropriately classified as a Tier 1 police response. That is, even though the report of a crime may not accompany this response, a significant risk to the people involved was identified.

The 38 call types (10.7%) identified for Tier 2 co-response represent an opportunity to provide a more tailored resource (*e.g.*, Mental Health Professional, see Blueprint), within a risk-mitigated response (*i.e.*, where an officer is available to immediately intervene, if necessary, to prevent physical harm to the co-responder). The RMD analysis suggests approximately 31 FTE[18] are necessary to build capacity for this type of response. However, while Mental Health Professionals are the obvious resource, several professional disciplines may be viable in this role.

Under the Tier 2 response, a police officer would be primarily responsible for the contact, with the co-responder available in a support role as necessary or until any potential threat to responders or bystanders is resolved. At this point, every Tier 2 response could revert to a Tier 3, where the co-responder takes the lead.[19] There is some overlap in response between Tier 2 and Tier 3. Phase One analysis indicates the need for approximately 8 FTE in this role. In addition to low-risk crisis contacts, Tier 3 responses include 49 call types (13.8% of responses), including low level calls for assistance, mediation of disputes, and initial investigation of noise complaints.

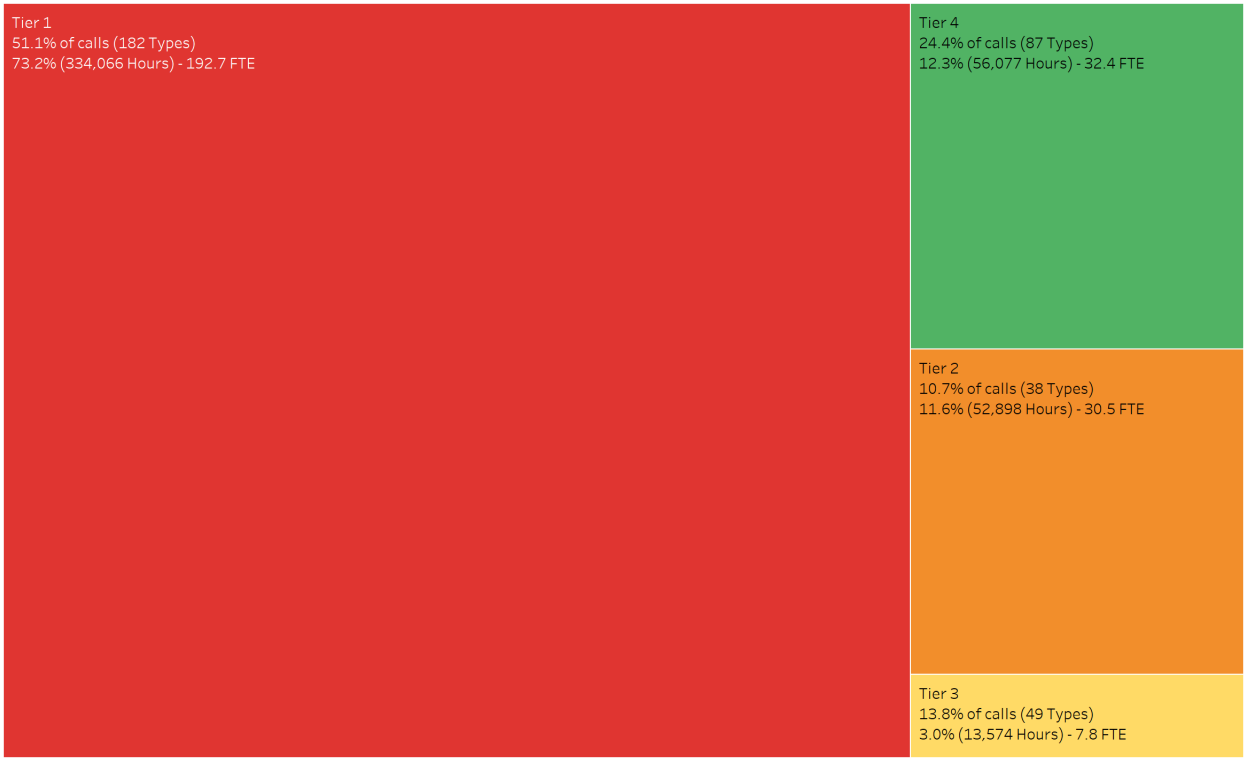

Figure 2 - RMD Staffing Analysis

Tier 4 is an opportunity for the City to improve customer satisfaction and realize some resource efficiencies. Represented in the 87 call types (24%) are crime reports necessary to file insurance claims and less well-defined requests for service. Some of these calls are for "hazards" or "traffic control", roles that can be safely filled by civilian responders. A combination of deferred response, online reporting, and alternate responders (*e.g.*, animal control) can fill this need.

[18] This calculation is a rough estimate of the FTE needed to address the workload (in hours). The calculation assumes 2,080-hours +20% per FTE (for modest shift relief) and would need to be adjusted for 24/7 scheduling. For reference, the shift relief factor for sworn ranges from 1.7-2.1 under the current 4-2 shift schedule. Shift relief factors vary widely based on the actual shift schedule of the deployable resource.

[19] The tiered response model proposed by RMD is identical to NICJR, except for coding. RMD *reverse codes* the response tiers to reflect current practice in call prioritization. Under RMD, a Tier 1 response represents the highest level of risk. A Tier 4 response represents almost no risk.

Operationalization:

As indicated previously, Phase Two (in progress) seeks to forecast the risk of a response. This information will be presented to call takers as they are listening to the request for service. Common in modern call center operation is on-the-fly NLP. This technology listens to the call and suggests insights to the call taker (customer service representative) to provide a more efficient and ultimately more satisfying experience for the caller.

From the call tiers defined by this analysis, preliminarily, bag-of-words models will be generated to train the intelligent call center models. Using a combination of cloud computing and web-based applications, this technology will inform CSCC as to the safety and appropriateness of an alternative response. The SPD is working to mature this technology, in collaboration with CSCC and Amazon Web Services. All technology will comply with relevant security and privacy best practices, as well as the City-specific requirements associated with the City's Privacy Policy and Surveillance Ordinance.

This effort is timed to coincide with response capacity coming online. Until capacity can be built, the SPD and CSCC will continue to mature the technology to become more accurate. The RMD project, as indicated in footnotes above, is designed to evolve with the City. As we more accurately represent the risks and benefits to responders and our community, the system will adapt to reflect those realities. Further, RMD assures the City of Seattle continues to set the standard for the future of police and emergency first response, nationally.

A Differential Police Response:

The overdiversification of police response has been contemplated since the 1960's. Experimentation with a Differential Police Response (DPR) began in the late 1970's but stumbled on the issue of call triage. RMD identifies opportunities to realize a DPR (Tier 2, 3 and 4). Co-response or an altogether alternative to police responders in these instances may have service or efficiency benefits. Within these strata, the City can deploy a resource capable of bringing a trauma informed response or a professional with a tailored specialty suited to the specific service being requested.

In a 2021 publication resulting from research supported by the SPD, researchers at John Jay College of Criminal Justice found, among other things, the availability of services to be predictive of a decision to arrest (community caretaking) a person in crisis. Absent a service to refer a subject or a Mental Health Professional (MHP) capable of providing the standard of care, police are often without an alternative to emergent detention. Similarly, in service to members of our community in critical need of housing or other services may be rendered by a person credentialed with a Master's in Social Work (MSW) or an outreach professional with lived experience.

In addition to providing a tailored service, RMD enables an efficient response that is both more satisfying to the community member requesting assistance (customer) and makes the best use of limited police resources. Completed crimes reported for the purpose of filing an insurance claim, for example, can be routed to the proposed *portal*—a modern, online reporting platform. Requests to retrieve found property can be routed to Community Service Officer's (CSO's), without compromising the safety of those involved.

Next Steps:

This technical brief is provided in the spirit of collaborative development. The SPD recognizes emergency response and more general requests for service from our community as a City responsibility. Within the areas of risk and police service, the SPD is uniquely suited to answer some questions. Some questions are better answered by stakeholders and members of our community, with their own unique perspective and values.

As we work to incorporate feedback and learn more about these data, the risk matrix and tiered response design will change. The insights presented here should be taken as preliminary estimates. In the coming weeks, the SPD intends to provide a dashboard that can be engaged by stakeholders to interact with the matrix, alter some of the design parameters, and re-evaluate the model.