January 2020

TRUST IN SMART CITY SYSTEMS
Characteristics and Key Considerations

Executive Summary

The Smart Cities Council, a global advocate for smart city adoption, states that the term "smart cities" still lacks a universally agreed upon definition. [1] However, the term generally refers to the integration of information technology (IT) with the management and operation of civic functions. As these civic functions can include operational technology (OT) elements that monitor and operate physical systems, a smart city system can be seen as representing the intersection of the IT, OT, and public service domains of practice. All three domains are represented by mature fields of practice, but their combination in cross-domain projects can expose gaps within each domain, where key characteristics important to one domain might not be considered by the others. If not addressed, these gaps can introduce security, safety, and privacy risks, including risks to critical infrastructure and its underpinning technology.

These risks are not some "distant future" concern. In 2018, investment in smart city technologies is expected to reach $22 billion, and the investment is expected to grow in the future. [2] Hundreds of projects have been deployed or are in some stage of development within the United States. Since these systems are often intended to remain in use for years or decades, the impacts of failures to fully address risks introduced by these systems can be felt for years to come. Between the potential for lingering impact and the aforementioned connection to critical infrastructure, the relevance of smart city projects to the Cybersecurity and Infrastructure Security Agency (CISA) which is charged with securing this critical infrastructure, is straightforward. Products such as this paper are one way CISA can help key stakeholders ensure that smart city projects are better prepared to address the risks associated with such projects and that the result is part of a more secure and robust national infrastructure.

This paper presents a set of key "trust characteristics" that need to be considered when planning for a smart city project. A trust characteristic is an attribute or behavior of a smart city system that the users and operators of that system need to believe the system will provide or preserve. When these stakeholders are confident that the system will perform as expected with regard to all trust characteristics, then the system is said to be "trustworthy." [3] These characteristics are identified by looking at the key priorities of the domains of practice contributing to and touched by smart city projects. While all identified trust characteristics have been noted in prior work, those works tend to focus on characteristics important to one or two domains of practice; this paper attempts to identify in a single location the key characteristics important to all three contributing domains.

This paper is intended for stakeholders participating in the initial, high-level design of a smart city project and its objectives. For the purposes of this paper, "stakeholders" could be someone with a particular goal or vision they hope to achieve through the project, parties who will be impacted by the project, service providers who are interested in integrating with the project, and parties who are providing funding or other resources to support the project. This paper is intended for use early in the design process, since considering these characteristics early allows them to be part of the core smart city system design (rather than add-ons) and can help ensure that appropriate community members and experts are engaged in the design process. The key characteristics presented in this paper are intended to help guide these engagements and design decisions. It is hoped that identifying these trust characteristics in a single place, along with guidance on how to consider these characteristics, will better prepare smart city designers and implementers to create smart city systems that address the trust expectations of all parties impacted by the system.

The key trust characteristics are as follows:

- **Security** - Assurance that actions and data are available only to authorized parties.

- **Reliability** - The dependability of a process to operate correctly and within expected parameters.

- **Safety** - Avoiding injury or damage to persons, facilities, and the environment.

- **Resilience** - The ability to continue to operate under adverse conditions.

- **Privacy** - A party's ability to control the exposure of data associated with them.

- **Maintainability** - Assurance that the system will remain operational in the future and can adapt and grow as needed.

- **Compliance** - The ability to conform to requirements associated with a particular context.

- **Well-being** - Preservation of livelihood, quality of life and environs, and minimization of disruption.

- **Fairness** - Lack of bias, equal access, and transparency.

- **Integration** – The ability to work with or alongside existing infrastructure and processes without disrupting them.

- **Measuring Utility/Impact** – The ability to perform the needed job and/or create the desired change expected by the system's stakeholders.

All smart city projects need to consider each of these key trust characteristics. The degree to which each trust characteristic is important to a project can vary based on the purpose, behavior, and function of a system. Moreover, since every city and project is different, how each characteristic is addressed can vary widely as well. This paper provides additional detail on each trust characteristic as well as a list of key considerations for each characteristic. The latter helps project planners better understand how each characteristic fits with their particular project.

Acknowledgments

The Department of Homeland Security does not advocate or endorse any particular approach to Smart City Systems and readers will have to make the necessary decisions to determine what options work best for them.

Table of Contents

1   Introduction

The Smart Cities Council, a global advocate for smart city adoption, states that the term "smart cities" still lacks a universally agreed upon definition. [1] The term generally refers to the integration of information technologies (IT) with the management and operation of civic functions. "Smart city projects" are expected to have a large impact on U.S. citizens in the near future. Investment in smart city technology is expected to reach $22 billion in 2018 in the United States alone ($80 billion globally), and this investment is expected to grow in the future. [2] Because of the broad scope of smart city efforts, these projects could impact virtually every aspect of modern life, including communications, utilities such as water and power, transportation, and government services. Both the volume of investment and the scope of impact of these projects mean that U.S. citizens will be using and relying upon more smart city technologies.

Smart city projects are challenging. Despite massive investments, nearly one-third of smart city projects fail, and almost 80% of prototypes have not successfully scaled up to reach the desired scope. [4] Beyond the economic impact of such failures, these projects can inadvertently lead to security, safety, privacy, and infrastructure risks for communities, either by creating such issues themselves, or by providing attack vectors that allow malicious use of the services and components. One reason why these projects are so challenging is that they are almost always cross-cutting in nature. They often contain IT, operational technology (OT), and public service elements, and might overlap roles from different elements of city government and different activity sectors. IT, OT, and public service domains of practice often have different perspectives on what is important in a system. For example, IT systems often prioritize security over reliability, while OT systems often take the opposite stance. [3] The confluence of perspectives from these domains of practice can create severe problems for smart city projects if the differing, sometimes even contradictory, practices and priorities associated with those domains are not effectively reconciled.

The material in this paper is pulled from sources and experts from all of these perspectives. The result is the identification of a set of key characteristics for smart city projects. These key characteristics range from familiar risk contributors, such as security, safety, and privacy, to operational considerations, such as compliance, integration, and maintainability, to citizen concerns, such as well-being and fairness. The breadth of these considerations reflects the many perspectives and practices impacted by smart city projects. While many individual characteristics have been identified in prior efforts, such work tended to have a siloed perspective and addressed only specific characteristics of importance to certain areas of practice. A broader, more complete identification of key characteristics, and an understanding of the different perspectives that see them as important, is a necessary first step to the integration of solutions that address the broad scope of smart city stakeholder concerns. Thus, this paper seeks to provide a more comprehensive picture of key smart city characteristics in a way that is bound neither to a certain technology perspective (IT or OT) nor to any specific industry sector.

1.1   Audience

The audience for this paper is stakeholders participating in the initial, high-level design of a smart city project and its objectives. "Stakeholders" could be someone with a particular goal or vision they hope to achieve through the project, parties who will be impacted by the project, service providers who are interested in integrating with the project, and parties who are providing funding

or other resources to support the project. A project's set of stakeholders could include representatives from all these groups.

This document is best employed during the very early phases of a smart city project design, before specific technical requirements and solutions have been identified. There are multiple reasons for this:

- Early consideration of these characteristics allows them to be "designed in" rather than "bolted on."

- Early consideration of key characteristics means that they can be considered more holistically within a project, rather than applying them piecemeal to individual components.

- Some of these characteristics should be addressed during the design and deployment of a smart city project, rather than simply being elements of the final result. For example, holding open conversations and employing transparency in the design early in the project can be a useful technique for building community trust; waiting to have such conversations after project completion is far less effective.

- The characteristics might suggest additional parties who should be involved in the planning, requirements development, and implementation of a smart city project. In a project that crosses domains and perspectives, success will require diverse participation to reflect those perspectives. Bringing in diverse parties early in a project, even in the conceptual stage, allows their inputs to be addressed in the core design, and will increase their interest and investment in the project's success.

The information provided about each characteristic will help facilitate deliberate and intentional decisions on the part of stakeholders with regard to important aspects of a project's design. While this document does not identify specific solutions for each characteristic, largely because any solution will be highly dependent on the nature and context of the smart city project, by considering the key characteristics presented herein, project designers will have a better chance of avoiding many of the oversights that can plague projects that combine IT, OT, and public service domain elements.

1.2   Methodology

This paper identifies its list of key characteristics by comparing and combining the perspectives from domains of practice involved in smart city projects. These perspectives were identified using a variety of sources.

The authors conducted a review of articles from academic, government, and trade sources. Many of these articles dealt specifically with smart cities. Other articles focused on specific domains of practice to better understand their specific perspectives and priorities.

The authors also attended multiple conferences/events focused on smart cities. These conferences included numerous talks and panels by parties who were actively engaged in smart city projects. These events provided information about emerging practices, technology and social trends, risks to project success and how to deal with them, and case studies about existing efforts. They also provided networking opportunities that helped the authors identify individuals for follow-on interviews.

The final source employed in this paper consisted of direct interviews with smart city practitioners and thought leaders. The team conducted more than a dozen interviews over the course of several months. While not all interviewees are cited directly, all were very helpful in developing the authors' overall understanding of the smart city sector and the key characteristics important to smart city projects.

1.3 Relation to Cybersecurity and Infrastructure Security Agency (CISA) Responsibilities

This paper examines smart city projects through the lens of homeland security priorities. It was initiated and led by the Department of Homeland Security (DHS), Cybersecurity and Infrastructure Security Agency (CISA). The Homeland Security Systems Engineering and Development Institute (HSSEDI), a DHS-owned Federally Funded Research and Development Center (FFRDC), developed the paper.

Beyond the general impact smart cities are expected to have on the nation as noted above, many smart city projects support designated critical infrastructure sectors such as power or water [5], provide a service whose interruption could severely disrupt the lives of the impacted community, and/or involve information or activities that could be used to harm citizens if compromised by malicious parties. [5] As such, the security of such smart city projects and their potential impacts on citizens' safety and livelihood are often of great importance. These roles mean that the observations here are particularly relevant to CISA, which is tasked with protecting U.S. critical infrastructure.

Characteristics such as safety and security have a clear relationship to the CISA mission. One might therefore reasonably question why characteristics such as intergration and well-being are included and given equal attention in a paper that purports to reflect a homeland security perspective. There are multiple reasons for this more expansive treatment of smart city characteristics.

One such reason is that key characteristics do not fail independently—rather, they are interdependent. For example, if a smart city project cannot be readily maintained, then eventually the expense and difficulty of maintaining the system will have security impacts (e.g., due to failure to patch new vulnerabilities or address the latest security threats) and/or safety impacts (e.g., due to regular wear and tear of critical components, or failure to accommodate changing use patterns). Similarly, if a smart city project negatively impacts well-being, such as by creating delays or making routine activities harder, then citizens might seek to neutralize or bypass these disruptions, potentially in a way that compromises safety and security considerations. There are numerous examples in IT wherein users choose to circumvent security in the name of convenience and "getting the job done." [6] A system cannot be safe or secure in the long term unless all key characteristics are adequately considered.

Additionally, in many cases, addressing these key characteristics will involve trade-offs. A method of supporting one characteristic could potentially complicate support for a different characteristic. For example, techniques that make a system more resilient might make it harder to maintain. In any trade-off required in an implementation, it is necessary to understand the value of both sides of the exchange. Failure to understand both sides can lead to undervaluation and inadequate accommodation of characteristics important to some stakeholders. As noted earlier, failure to

account for stakeholder concerns and interests is a common contributor to the failure of projects that are inherently crosscutting, such as smart city efforts.

It is not the goal here to make every stakeholder in a smart city project an equal advocate for all characteristics. It is, in fact, highly beneficial to have stakeholders who focus on specific characteristics in order to ensure that they are adequately addressed and not undervalued by others. If a particular characteristic has no advocate among a project's stakeholders, then it is almost certain that it will fare poorly when trade-offs are made. Use of this paper in smart city projects will encourage early identification of stakeholders who can serve as advocates for each key characteristic. To be useful contributors, advocates must still be capable of understanding the importance of other characteristics. Advocates who are unwilling or unable to recognize the importance of other characteristics not only risk inadequate consideration of those other characteristics, thus putting the entire project at risk, but risk being seen as obstructionists whose views are to be avoided, since they cannot be integrated.

To summarize, while CISA's primary focus might be on a subset of the identified characteristics, it is impossible to adequately support any subset of the identified characteristics while ignoring or undervaluing others. To be an advocate for CISA's cybersecurity priorities within smart city projects requires an understanding of *all* characteristics that are important to the full range of smart city stakeholders.

1.4   Risk and Trust

All applications of new technology include an element of risk, if only due to the introduction of the novel into a system. However, this risk is amplified in smart city projects by the fact that much of the technology is designed to have a direct impact on a community, sometimes regarding elements considered "critical" to citizens [5], and thus the risks associated with that technology are experienced more directly and profoundly by citizens.

Risk is usually expressed using some variation of the following shorthand:

$$\textbf{Risk = Consequence x Likelihood } [7]$$

One's risk is the combination of one's understanding of the severity of a set of consequences with the likelihood that those consequences will be experienced. While there are many variations in the details of this expression ([8]**,** [9]**,** [10])**,** the centrality of likelihood and consequence severity in understandings of risk appears to be nearly universal. A large body of work on risk management deals with estimating both likelihood and consequences to better prioritize risk mitigation efforts, and all projects are well advised to employ robust risk management strategies.

The focus here is treatment of smart cities at a very high, sector-agnostic level, making a focus on risk less useful. Risk assessments that are not grounded in specific details tend toward what-if scenarios that might make riveting reading, but that often are of limited use in directing an organization's limited resources to provide optimal risk mitigation.

The Industrial Internet Consortium (IIC) defines the trustworthiness of a system as "the degree of confidence one has that a system performs as expected in respect to all key system characteristics in the face of [disrupting factors]." [3] Trust and risk are often viewed as different sides of the same coin: while risk considers the undesired outcomes that could happen, trust focuses on the desirable traits that need to be preserved. Trust in the preservation of key characteristics of a system is a prerequisite for understanding risk in that risks cannot be identified and prioritized without

first understanding what needs to be protected and with what priority. To emphasize this focus on trust, and specifically the need to trust that the key characteristics are addressed and preserved, the phrase "key trust characteristics" is used going forward.

Looking at key trust characteristics is feasible even when considering systems in the abstract. This is because, while the nature of a system's risk profile can vary considerably based on implementation details, even at the most conceptual stage stakeholders will understand what they want their system to accomplish and what characteristics they want their system to preserve. It is important to consider the key trust characteristics that are commonly shared by smart city projects.

This definition of trust clearly incorporates the idea that key system characteristics behave as expected during operation. This definition can also be interpreted within a longer timeframe. Specifically, it is not sufficient for trust characteristics to be preserved only in the here and now; there must be confidence that they will continue to be preserved in the future. For example, in the case of privacy, there needs to be confidence that systems not only preserve individual privacy now, but that collected data will continue to be protected in the future. In the case of public services, trust includes an assurance that the system, and the services it provides, will remain available and be dependable. As a result, in the identification of key trust characteristics, the authors considered both short- and long-term perspectives on trust.

1.5   Perspectives on Trust Characteristics

This paper identifies the key trust characteristics that apply to smart city projects in a sector-agnostic manner, so that those considering implementing or supporting smart city projects will consider all of these characteristics in their plans. As noted above, the characteristics are best considered in the very early phases of a project as an aid to identification of design requirements. This allows the project design to incorporate elements that protect and preserve those characteristics from the beginning, rather than attempting to bolt on new elements to patch problems after they appear.

Smart city projects need to integrate perspectives from a broad set of stakeholders, including the governments that often implement and run the projects, the technologists who design and integrate the components, the individuals who oversee the ongoing operations and maintenance of the system, and the citizens whose lives are impacted by the system. Each of these perspectives often comes with its own understanding of which trust characteristics are key, but might not readily see the criticality of characteristics valued by other perspectives. The IIC "Industrial Internet of Things Security Framework" [3] observes that the Industrial Internet of Things (IIoT) represents the confluence of IT and OT perspectives. They note that the IT perspective tends to prioritize privacy, security, and reliability, while the OT perspective usually prioritizes reliability, resilience, and safety. A trustworthy IIoT system needs to address all of these characteristics and do so in a way that recognizes the legitimacy of both perspectives. There are multiple examples of what can happen when a device's software fails to adequately address safety considerations [11], or when connected physical devices fail to adequately address privacy [12] or security [13].

Smart cities add yet another perspective to this mix of elements. While smart cities make use of IT (always) and OT (often) elements, smart city systems include public service elements as well. Unlike industrial systems, which are the primary focus of the IIC, smart cities systems operate in proximity to, and directly impact, the public at large. This added public service perspective brings

with it several additional trust characteristics that might not be immediately considered by those whose focus is primarily on the technical aspects (IT and/or OT) of solution development.

1.6 A Complete Picture of Trust

Eleven key trust characteristics are identified in this paper. In discussions with city leaders, product vendors, and other stakeholders of smart city projects, no party independently identified more than a handful of characteristics. This is not to say that these parties did not feel the other characteristics were important—no one ever expressed disagreement with any of the characteristics identified. However, given the limited time and resources available to them, stakeholders often felt a need to focus on a small number of factors so that they would not become overwhelmed. While the number of top-level considerations identified varied, no one listed 11 priorities. How, then, are practitioners expected to deal with the 11 key trust characteristics identified here?

They key is to understand these characteristics not as a to-do list, but instead as a "map of the terrain" in which a particular smart city project is situated. When using a map, not every symbol and feature is going to be relevant, and even among the relevant features, some will be more important than others. This paper is intended to help smart city stakeholders better understand the "trust terrain" in which their project is situated and use this understanding to better prioritize the elements relevant to their project, as well as to identify and discard elements that are not relevant. While it is important that all trust characteristics are sufficiently addressed, what constitutes "sufficiency" can vary widely between projects. The goal of this paper is to help designers understand what constitutes "sufficiency" for their particular project with regard to each trust characteristic so that they may direct an appropriate amount of effort toward its achievement. Stakeholders will still need to determine how best to support each characteristic and to what degree, but this paper will help them make those decisions with a more complete understanding.

2   Trust Characteristics Overview

This document considers key trust characteristics that tend to be important to smart city systems. Briefly, the identified characteristics are:

- **Security** - Assurance that actions and data are only available to authorized parties.

- **Reliability** - The dependability of a process to operate correctly when it is expected.

- **Safety** - Avoiding injury or damage to persons, facilities, and the environment.

- **Resilience** - The ability to continue to operate under adverse conditions.

- **Privacy** - A party's ability to control the exposure of data associated with them.

- **Maintainability** - Assurance that the system will remain operational in the future and can grow as needed.

- **Compliance** - The ability to conform to requirements associated with a particular context.

- **Well-being** - Preservation of livelihood, quality of life and environs, and minimization of disruption.

- **Fairness** - Lack of bias, equal access, and transparency.

- **Integration** – The ability to work with or beside existing infrastructure and processes without disrupting them.

- **Measuring Utility/Impact** – The ability to perform the needed job and/or create the desired change expected by the system's stakeholders.

Their order is not intended to imply priority, and, in fact, different characteristics will be viewed as primary in different projects. Moreover, under some perspectives, some of these characteristics would be seen as means, while others might view the same characteristic as an end in itself. The perspective one brings is likely to play a significant role in one's view of the list elements.

Each of these characteristics is important to achieve broad trust in a smart city system. Any given project is likely to need to meet additional requirements to be operationally successful—the trust characteristics presented are necessary for the system to be trusted by key stakeholders.

The trust characteristics listed above are general concepts rather than disjointed categories. In some cases, boundaries between the characteristics can overlap. For example, reliability has significant overlap with the security concept of service availability, integration is closely related to the well-being element of being minimally disruptive, and just about any of these trust characteristics might have associated rules or regulations with which implementations must comply. The intent here is to ensure sufficient coverage of key trust elements, rather than define distinct silos that are addressed independently. Indeed, treating each of these characteristics as independent elements adopted in isolation is counterproductive (and sometimes simply impossible) because of the many ways they connect to each other.

Before looking at these characteristics in greater detail, the following subsections address some important points relevant to all characteristics. These ideas should be kept in mind when reading about any of the individual trust characteristics.

2.1 Characteristics are Non-Binary

Trust characteristics need to be managed, rather than simply having a binary state of "addressed" or "not addressed." In most cases, the "perfect" support of a characteristic (e.g., "perfect security" or "perfect safety") is not possible. Even approaching some hypothetical perfection might not be practical due to the costs required or adverse impacts to other elements of the system.

Instead of always trying to be as close to perfect as possible with regard to all characteristics, designers need to identify a "sufficient level" that supports each characteristic. Determining what qualifies as "sufficient" will depend on the nature of the project. A risk-based assessment can be useful in making this determination, as in any system the failure of some characteristics will be more consequential than others. A system on which public safety depends will need to be very resilient; a system that provides a minor convenience will only need to be resilient enough that it does not annoy people by being offline too often. The question is not whether to support a trust characteristic, but what degree of support makes sense.

2.2 Trade-Offs

Certain characteristics might be in tension with each other, and any solution needs to represent a compromise between competing objectives. This sort of trade-off is frequently seen in other endeavors. In IT, security is often prioritized over reliability, and it is not uncommon for a system

to be taken offline to contain a suspected compromise rather than risk exposure of confidential information. [14] By contrast, many industrial systems (examples of OT) prioritize reliability and have been known to resist patching or upgrading software products because doing so could potentially disrupt the existing system. Both cases represent a trade-off by operators: IT operators value the reliability of their systems, but may be willing to compromise it in the name of security; OT operators value security, but may be unwilling to risk reliability for it. In both cases, designers and operators made a trade-off based on system priorities. Smart city system stakeholders will have to make similar trade-offs, and understanding which characteristics to prioritize in the system is an important part of making good trade-offs.

2.3   The Importance of Perception in Trust

Perception can have critical implications for trust in a smart city system. A system that is safe (or secure, or reliable) will not be trusted if people believe it is unsafe (or insecure, or unreliable). In addition to the analysis of potential issues surrounding the characteristics described above, stakeholders should also consider elements related to the perception of these characteristics. User trust in the smart city system will certainly be influenced by its actual performance record, but that trust can also be influenced by issues that have more to do with belief and appearance.

Toward that end, how and when a project engages with the public can have a significant impact on perception. Actively listening to the concerns of the public not only helps identify ways to improve perceptions in the design and implementation of the system, but the act of engagement itself can be helpful in improving trust. Similarly, educating the public about the project, including its intended functions and benefits, can make the system less of a mysterious "black box" and improve public trust. A project that clearly articulates the safety (or other) aspects of the new system is likely to be more trusted than a system where the public is left to make its own determination in the absence of information.

Failure to anticipate and address perception can lead to the same loss of trust as an actual failure to address a given trust characteristic. This is not to say that stakeholders should engage in efforts that are designed to look impressive but have little practical impact. However, some consideration should be given to making the mechanisms that address these trust characteristics clear and visible.

2.4   The Primacy of Relationships

A key aspect of one's trust in a system will revolve around one's trust in those who designed, implemented, and manage the system. This point, and the need to develop and maintain relationships to create this trust, was emphasized repeatedly at conferences and in interviews. The quality of an organization's work on a system is unlikely to win over a public that does not trust that organization to act in their interests. This is another element where engagement during the initial design, development, and deployment is critical; by the time a system is operational, many opportunities to build trust relationships will be lost. As noted in the discussion of perception, transparent engagement is an important step in trust building. This engagement has the ability to create relationships on which trust can be developed. In particular, having members of the team— including engineers, maintainers, and operators—meet and talk with the other stakeholders can help the public build trust. It is often easier for community members to trust people than to trust organizations or complex systems. It can also be useful to identify trusted members of various

citizen groups and actively work with them to address community concerns. In such a situation, the community's trust in their representative can serve as a bridge to trust in the system itself.

2.5   Characteristic Importance Can be Emergent

As noted earlier, smart city leaders tend to focus on only a handful of characteristics. Citizens and other stakeholder groups can behave in the same way, and in any discussion with such parties, they are likely to focus on a subset of the characteristics listed in this paper. One interviewee who served as a state Secretary of Technology observed that citizens tended not to raise issues of safety and security in discussions of proposed smart city projects and suggested that this was because preserving these roles has traditionally fallen to government organizations (e.g., police and fire departments), and thus they assumed that the government would address those characteristics. [15] By contrast, citizens frequently raise concerns about privacy, cost and debt (aspects of maintainability) with regard to smart city projects, potentially because there is less trust that governments will address these characteristics. However, should trust in any of these characteristics be violated (e.g., the system proves to have safety problems), the reaction will be no less severe.

In other cases, citizens might not raise concerns regarding a particular project because they are not aware that the project touches on those concerns. For example, citizens might assume that a reactive traffic light system would not have privacy implications. However, if those lights do have privacy implications (e.g., the cameras are storing video of traffic), then those implications need to be adequately addressed or a public backlash is likely.

These are just a few examples of how issues that are ultimately important to the public might not be raised in public conversations. The point is that, while it is important to listen to input from citizens and other groups regarding a smart city project, this is only one data point to drive prioritization of the trust characteristics. Those developing smart city projects need not only to listen to this input, but also to use their knowledge of the planned system to anticipate what issues might be important to the public, even if unvoiced.

2.6   Evidence and Verification

The adage "trust, but verify" holds true in the context of smart city trust. Even when a project can earn the trust of its stakeholders, it is important to provide stakeholders with the opportunities to validate that trust. Keeping records and audit trails regarding all key trust characteristics is one tool toward this end. Such records can cover the whole project, from design, to deployment, and then throughout operations, and can show the steps taken to ensure that key elements of each trust characteristic are considered and monitored. For example, audit trails of what information a system collects and how it is used can help enhance trust in the system's privacy, audit trails of security incidents and responses can enhance trust in the system's security, and audit trails of system uptime and proactive maintenance can enhance trust in the system's reliability. While most stakeholders will probably not actually review these records, simply knowing that this information is on record can enhance trust in the system.

Of course, one needs to be careful that such records do not end up undermining the very trust characteristics they are intended to support. For example, detailed information about security and resiliency capabilities could help attackers discover ways to subvert those capabilities, and any large data collection activity is going to be subject to privacy concerns. For these reasons, it is

important that auditing efforts be carefully designed and targeted, rather than simply collecting and recording every possible piece of data. Careful planning of auditing methods not only helps mitigate the risks, but is likely to make the records more usable as well by focusing on information of greater value without the need to filter out irrelevant data. For these reasons, project planners should give serious thought to what information is needed to demonstrate the trustworthiness of their system and how to gather and expose that information to best effect.

3 Trust Characteristic Details

This section provides an overview of each of the identified trust characteristics. Each of these characteristics is a large topic, often supported by a mature field of study. This paper can provide only a cursory overview of each topic to help the reader consider the topic in the context of smart cities. Development of any requirements or solutions to provide and protect these characteristics will require greater expertise on these topics than can be provided here. However, the sections on each of the key smart city trust characteristics should at least provide enough information to help readers understand how the characteristic relates to a given project and identify initial steps to address that characteristic.

Each characteristic is presented in general terms in order to be broadly applicable to a wide range of smart city projects. However, understanding the trust characteristic's relevance to a specific project, and in particular which aspects of the characteristic are applicable or inapplicable to the project's specific circumstances, requires considering the characteristics within the unique context of that project. To facilitate this, each trust characteristic section concludes with a list of key considerations to help stakeholders filter and focus the concepts discussed in that section. By doing this, they can identify those concepts most important to their project. As noted earlier, not all trust characteristics are equally important to all projects, and even when a characteristic is relevant, project designers need to determine what constitutes "sufficient" support for the characteristic. The key considerations for each trust characteristic are intended to be the first step in helping designers think about what constitutes sufficiency for their particular smart city effort. The list of considerations is certainly not complete, and stakeholders should treat it as a starting point in determining the role a given trust characteristic plays in a project.

3.1 Security

Security involves preventing unauthorized parties from using or disrupting a system. **A key element of security is the idea of "authorized parties"**—that the ability to perform an action is conditional on who or what is attempting to undertake that action. As a result, some way to measure authority is inherent in any security mechanism. Authority could be tied to identity, knowledge of some secret (e.g., a password), possession of some object (e.g., a key), or some other indicator. More sophisticated authority measures might include a combination of these, and potentially other attributes such as time and place. For example, police officers have different authority depending on whether they are on or off duty. Determining what activities require authorization, which authorities will be allowed to perform those actions, and how those authorities will be recognized are all key initial steps to any security solution.

A related, and critically important, element involves determining who/what can be trusted to be "authorized" for certain activities. For human actors, this can involve background checks and other tests to determine whether they are likely to abuse their authority. For machine actors, this

generally takes the form of testing to ensure that the component is sufficiently reliable for its role. The strength of a lock becomes moot if everyone is given access to the key. In the same way, the best technical security elements can be nullified by a weak vetting and access management process. **The security design of a system needs to consider how to determine if parties can be trusted with authorization.** Similarly, procedures need to be in place to review, modify, and revoke authorization as appropriate. Insider threats, where attacks are launched by current or former employees, are a significant security issue, and managing authorization (through vetting, monitoring, and revocation of unneeded access) is the key tool to combat it. [16] The bottom line is that the technical mechanisms that verify a party is authorized are only part of a security solution. These technical mechanisms need to be complemented by a sufficiently robust method of assigning, modifying, and revoking the authority of actors to take actions on the system.

Security is often described as having three key elements: [17]

- Confidentiality – Avoiding disclosure of information to unauthorized individuals

- Integrity – Avoiding changes (to processes or information) by unauthorized individuals

- Availability – Preventing unauthorized individuals from denying services to authorized individuals

In IT, this triad is often abbreviated as CIA, while in OT it is usually written as AIC. This reflects that IT tends to focus more on confidentiality, while OT focuses more on availability. [3] One important takeaway is that **not all security elements will necessarily be of equal importance in all circumstances** and that, as with the identification of trust characteristics, one's professional background can lead to bias as to which are most important.

Security protections can fail for multiple reasons. One common cause is a policy failure, where an undesired activity is allowed because the security policies and protections do not prevent it. This can happen due to a poorly designed or implemented security policy or because of failure to recognize that access to a particular resource needs to be controlled in the first place. In such a situation, the security mechanisms are performing nominally, but the outcome of that performance is not what was desired. Alternately, the security mechanisms themselves might be overcome. In this case an actor is not authorized to perform an action but is able to do so anyway. This can happen due to software vulnerabilities, physical alteration of the security mechanisms, failures by personnel to correctly perform security responsibilities, or some other procedural failure (e.g., an authorized action is initiated at the wrong time or in the wrong way).

Security crosses many potential segments of a smart city system. There is information security, physical security, operational security, institutional security, and potentially other security segments and sub-segments depending on the sector, activity, and how the proposed system is managed. **Each security segment defends against different types of threats and uses different mechanisms.** Information security deals with protection of data and data processes on computers and over communications media. Information security mechanisms include software and certain IT hardware elements that encrypt data, detect cyber intrusions, and preserve data against modification.

Physical security controls physical access to facilities and protects them from damage and vandalism. Physical security mechanisms include cameras and other sensors, locks, barricades, security guards, and other elements. Despite the vast differences in what is protected and how, **security segments are often co-dependent on each other**, and a security failure in one part can

lead to failures in other parts. For example, IT systems can be compromised through physical security failures (e.g., someone with unsupervised physical access to systems can plant bugs and other recording devices in computers) or through operational security failures (e.g., social engineering attacks). For this reason, security segments should not be designed in isolation, and there should be coordination between security operations.

Another important consideration for security is that it deals with stopping *actors*. In some cases, those actors might be benign—a user who accidentally double clicks on a file they are not authorized to read. However, **security often involves blocking actors seeking to deliberately circumvent those security features.** Given the active nature of this threat, preserving security requires adapting capabilities to meet and block new adversarial techniques. In other words, security needs to be maintained, rather than simply implemented. Because of this, it is useful to consider what objectives attackers might have. These objectives could relate to the smart city system itself, any capabilities that share infrastructure with that smart city system, or any capabilities that have dependencies on the behavior of that smart city system. An urban camera system could, for example, be attacked to disable it or to allow unauthorized parties to view video feeds. If the cameras were used to feed a road toll collection system, the system might be attacked to block collection of tolls. If the cameras shared network infrastructure with the credit card payment system used to process tolls, those cameras might be used as an entry point to compromise those payment systems. Each type of attack would serve a different end, which means that they would attract different types of attackers who bring with them different skills and resources. For example, an average motorist is unlikely to put much effort toward avoiding a few dollars in tolls, but a criminal organization might put significantly more resources toward compromising the credit card payment system that might share the same network as those cameras. While this is an IT-focused example, smart city systems might also have physical or operational overlaps and dependencies. In summary, both direct and indirect attack objectives should be considered when determining how much and what kind of security a smart city system requires.

Trust that a system is secure is often a prerequisite to having trust in other key trust characteristics. Confidentiality failures can expose private data (e.g., data breaches); availability failures can degrade resilience and reliability (e.g., denial of service); integrity failures (e.g., data corruption) can compromise the behavior of a system, harming impact and utility; and ultimately, a malicious party that gains the ability to manipulate mechanisms to preserve one or more trust characteristics can cause those mechanisms to be subverted. Thus, compromise of a system's security is almost universally a means to compromise the operation (and trust in) some other element of the system. As such, it is important to provide security controls not only surrounding the key functional elements of a system, but also around any other elements that are important to preserve.

3.1.1 Key Security Considerations

1. Consider what activities in the smart city system should be limited to authorized parties.

2. Consider the conditions under which a party would be authorized to perform a sensitive action or view sensitive data.

3. Consider how parties will be investigated/tested to determine whether they are worthy of being designated as "authorized" (trusted) for an action.

4. Consider what security requirements are most important (e.g., confidentiality, integrity, availability).

5. Consider what types of security are necessary (e.g., physical, information, operational).

6. Consider what objectives an attacker might have in attacking the smart city system.

7. Consider what systems share resources with the smart city system (and thus might be vectors to attack the smart city system, or the smart city system might be vectors to attack them).

8. Consider what systems are dependent upon the smart city system (thus making attacks against the smart city system a means to attack those other systems).

9. Consider how the security of the system will be vetted during system design.

10. Consider who will be responsible for maintaining the security of the smart city system.

11. Consider what resources will be needed to evolve and update the security of the system as threats change.

12. Consider what resources will be needed to review the trustworthiness of authorized parties and add new authorized parties, as necessary.

3.2 Reliability

Reliability is the dependability of a process to operate correctly and within expected parameters. Failure of a smart city system to be reliable will, at least, result in frustration among the system's users and, if too unreliable, can render the system unusable. In an industrial setting, reliability is generally measured as the fraction of a system's actual availability over its scheduled availability, with the latter taking into account things like turning off a system for planned maintenance. [3] However, this definition reflects that industrial facilities can, by definition, plan to reduce or eliminate demand for the system's services around planned downtime. Such planning is less practical when dealing with demand for utilities, whose need is constant; emergency services, whose use cannot be anticipated; or general public services, where it cannot be assumed that the users of the system will all be aware of planned downtime. As such, **in a smart city context, reliability will encompass not only minimization of unexpected downtime, but also minimization of the impact of planned downtime.**

Reducing unexpected downtime involves predicting and pre-empting likely failure scenarios. With regard to hardware systems, there is a mature field of practice associated with such predictions. Such systems look at statistics such as mean time to repair, mean time to failure, and mean time between failure metrics. These measures can be predicted by testing physical components to the point of failure in a lab environment. Reliability measures in IT are often expressed in terms of availability as measured by feature, load, and regression testing. [18] All these measures use time to indicate when predictive maintenance efforts should be performed to avoid unexpected failures. In some systems it is also possible to install sensors that measure signs of wear (in physical systems) or system instability (in IT systems) as a dynamic indicator of the need for predictive maintenance.

Predictive maintenance may necessitate downtime. IT systems might need to be patched and rebooted, while physical systems might need to be turned off while parts are replaced. **While**

predictive maintenance tends to be far less disruptive than reactive maintenance (which occurs only after the system has suffered a breakdown), it can still be disruptive to those who use and rely on the system. For this reason, support for reliability also involves minimizing the impact of such planned downtime. This is often accomplished through techniques such as phased maintenance, where only part of the system is down at a given time, thus allowing the system as a whole to continue operation, albeit at a degraded capacity. Maintenance can also be timed to occur during periods that have historically shown little demand for the service, such as timing IT maintenance to occur overnight. These and other techniques can help ensure that planned downtime is minimally disruptive.

The "operate correctly when expected" element of reliability can have multiple implications. For a web server, being reliable generally means being able to respond to a request for content in a timely manner. For a power generation service, it means providing a safe and uninterrupted supply of electricity to users. Obviously, there is a considerable range in the potential importance of this characteristic across different systems. Virtually any case where a system exhibits **poor reliability will lead to frustration on the part of users, but some reliability failures could have significantly broader implications.** Prioritization of a system's reliability, and thus the effort and expense that is sufficient to put into maintaining this characteristic, should reflect the probable impact of degraded reliability.

By the same token, "mere" user frustration with the system should not be dismissed. **Perception that a system is unreliable will erode confidence and can damage adoption**, potentially undermining the system's effectiveness. In fact, people's perceptions of the reliability of automated systems is, in general, lower than the actual rate of reliability. [19] That means that users may perceive the system as "down" more often than it actually is. For this reason, reliability efforts might also need to include mechanisms to specifically address perceptions of reliability—for example, by providing metrics that demonstrate the actual availability of the system.

Finally, when estimating the reliability needs of a smart city system, it can be useful to **consider that dependency on a system can grow.** This can occur because increased adoption increases user dependency on the system, or because the system is integrated into another smart city project and thus becomes a technical dependency for the new system. For these reasons, it may be useful to consider a design that allows reliability mechanisms to be improved over time.

3.2.1   Key Reliability Considerations

1. Consider the potential impacts on users/citizens if the smart city system is unavailable.

2. Consider the potential impacts on other technical systems if the smart city system is unavailable.

3. Consider what methods are available to anticipate when predictive maintenance should be employed.

4. Consider what methods are available to reduce the impact of planned downtime of the system.

5. Consider how reliability of the system will be measured.

6. Consider whether it make sense to publish system measurements to help ensure a perception of reliability and how this might be done.

3.3 Safety

Safety means avoiding injury or damage to persons, facilities, and the environment. [3] Examples of safety issues using this broad definition include things like fall hazards and exposed sharp objects (personal safety), fire and explosion hazards (personal and property safety), pollution issues (all three), and even radio and light emissions (environmental safety). When considering a smart city system, it is important that all aspects of safety (personal, property, and environmental) be considered.

The first safety requirements date to ancient times (e.g., "When you build a new house, make a parapet around your roof so that you may not bring the guilt of bloodshed on your house if someone falls from the roof." Deuteronomy 22:8). Thus, it should be no surprise that **the field of safety analysis is quite mature, with established norms, techniques, standards, and specializations**. Involving experts in the field of safety in the design of a smart city system is not only highly advisable, but it could be a legal requirement. In many cases, security norms have been codified into regulations, such as city building codes and Federal Communications Commission radio emissions guidelines. Compliance with regulations like this is dealt with more directly in the section on compliance.

Addressing the safety trust characteristic involves looking at the specific role and context of a given smart city project and identifying safety issues that may arise. If regulations are well written, there should be significant overlap between compliance with safety regulations and the mitigation of actual safety issues. However, **even with the best regulations, there are likely to be safety issues that are not adequately addressed by mere compliance.** This is especially true for types of services and technologies that are relatively new and might not have sufficient established precedents to drive appropriate conventions and regulations. Given the relative novelty of the smart city sector, this can be a common characteristic of smart city projects. This makes it especially important that security in smart city projects be considered beyond the confines of mere regulatory compliance.

Safety is often a key consideration in OT projects. Hazard analysis of physical systems, such as those dealing with toxic or flammable materials, is quite mature. [20] [21] OT safety assessment techniques focus on physical processes and employ empirically derived component failure probabilities, such as assessment of the mean time to failure of a given part. [3] By contrast, safety is rarely a significant consideration in IT systems since safety impacts are all physical in nature, and thus rarely impacted in a pure IT system. Moreover, IT systems do not have the same failure modes as OT systems; where OT failure modes are generally associated with repeated use and associated wear, IT parts do not fail due to repetition, but instead due to unexpected input or state conditions. [3] Examples of software failures that resulted in injury or death, such as the case of the Therac-25 medical accelerator [11] or the Multidata Systems International therapy planning software [22], occurred because the machines were put into states that the designers never anticipated.

The fact that IT and OT systems designers have such differing understandings of the role of safety presents a challenge in smart city systems, which will often combine both perspectives. While the ultimate impact of an explosion or gas leak is not changed based on whether the system in question is IT-connected, the failure modes that can lead to such a situation can be impacted by the presence of IT elements, often in ways that are not well understood if safety analysis is conducted within narrow IT or OT-exclusive silos.

In addition, because smart city systems are designed to interact with, influence, and/or inform the public (through their public service elements), any public safety issue is likely to have a more direct impact on people than if the system existed in an isolated industrial area. **Patterns of human use, both of the smart city services themselves and in the vicinity of any physical elements of that service, can have a significant impact on the magnitude of a safety issue.** OT engineers might have more experience with industrial projects that operate within controlled environments and are surrounded by parties trained in their safe operation; different expertise will be necessary to address the safety concerns of projects that operate in public spaces.

This analysis should consider the safety of the system under nominal operations, as well as safety during potential failure situations. Nominal operation safety could consider things like protecting citizens from fall hazards or sharp edges. Safety under failure conditions would consider potential issues if parts of the system cease to function correctly. Such situations could range from gas leaks to the service being unavailable (such as if traffic lights are no longer operational). A good safety analysis will consider both aspects.

Perception of safety is another critical matter for consideration. **It is not enough for the system to be safe—people need to feel safe using the system.** Failure to address perceptions of safety can significantly inhibit adoption and can even lead to people subverting elements of the system in an attempt to make themselves feel safer. An example of this is the persistent myth that one is safer not wearing seatbelts. [23] Significant research has been done on methods to help people feel safer; designers of smart city systems that might have safety implications are advised to avail themselves of this research.

3.3.1 Key Safety Considerations

1. Consider whether the smart city system has the potential to impact the physical safety of people.

2. Consider whether the smart city system has the potential to physically damage property.

3. Consider whether the smart city system has the potential to damage its environment.

4. Consider which parties need to be engaged in the design of the system to ensure that safety concerns are addressed.

5. Consider whether there are any legal requirements for certain parties to be involved in the system design to ensure system safety.

6. Consider whether additional experts, beyond those legally required to certify the safety of the system, are needed to sufficiently minimize safety concerns.

7. Consider how perceptions of safety will be addressed and managed.

8. Consider how users and those in proximity to these systems will be made aware of how to act in a safe manner.

3.4 Resilience

Resilience is the ability to continue to operate under adverse conditions. Cyber resiliency is defined as "the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses,

attacks, or compromises on systems that use or are enabled by cyber resources regardless of the source." [24] This definition could reasonably extend to non-cyber elements of a system as well.

Engineering for resilience works on the assumption that bad things will happen, and so it is important to prepare for that eventuality. It focuses on minimizing and containing damage, restoring at least partial functionality as quickly as possible, and ultimately recovering from the event. Adverse conditions may include adversary actions, natural disasters, unexpected resource loss such as power outages, unexpected failure of part of the system, and human errors, among others. Since any condition by any cause is within scope for resiliency considerations, engineering for resiliency is less about prevention, like security, and more about contingencies and containment.

The following are key elements in developing a resilient system: [25]

- **Prepare** by anticipating potential failure scenarios. For this, it is important to understand the components and dependencies of the smart city system and the potential impacts of their failure.

- Develop designs to **protect** the system. This involves developing mechanisms to prevent failures, but even more important, it involves developing mechanisms to contain failures if they do occur. A resilient system might suffer the failure of a component, but it will include mechanisms that minimize the impact of that failure and reduce the chance that component failure leads to failure of the system as a whole.

- Develop mechanisms to **detect** failures in the system. Early detection helps speed recovery and is another way to prevent a component failure from turning into a system failure.

- Have a **response** mechanism so that component failures can be swiftly addressed. This involves having procedures (e.g., a "playbook") in place so that operators know what to do, as well as ensuring that the system design includes mechanisms that facilitate these responses. In the simplest case, the latter would involve ensuring that components are easy to replace, and a supply of spare parts is readily available. A more sophisticated resilience mechanism might have redundant components/processes pre-deployed for immediate use.

- Include **recovery** mechanisms to restore services. This is not necessarily the same as repairing the failed component (although such a repair would certainly be one way to recover services), but can instead mean ensuring that the full services of the smart city system are restored quickly, even if by alternate means.

Resiliency certainly has a technical dimension, and there are many ways a system can be engineered to make it more resilient. However, **human factors are especially important in creating resilient systems**, especially those that are expected to be resilient against more significant potential disruptions. Establishing important communications channels early (so they can be immediately employed during an emergency), ensuring that emergency roles are clear, and training personnel so they respond appropriately are all important tools for creating a resilient system.

A critical element to any resiliency preparation is the development of a plan to guide an organization's response to disruptions. Such plans, often called response plans or Continuity of Operations Plans (COOP), include guidance and procedures for organizations to use to prepare for and continue to operate under adverse conditions. The Federal Emergency Management Agency

(FEMA) has published a template that organizations can use to develop such plans. [26] While this template focuses on federal departments and agencies, the guidelines can be helpful for any organization seeking to develop plans to support continued operation during disruptions.

Although a resiliency plan is important in virtually all smart city systems, **making a system highly resilient can become expensive, particularly when protection methods such as redundancy are employed.** As a result, consider the criticality of the system when prioritizing costs associated with this characteristic. For example, if an app that helped people locate a free parking spot went down, it would cause citizen frustration and perhaps some additional road congestion, but not much else. However, smart city systems that support power distribution and emergency responders could have life-and-death consequences. These may be examples where protective resiliency measures such as redundancies (e.g., a hospital has a backup generator) are worth the expense.

3.4.1 Key Resilience Considerations

1. Consider how disruptive "failure" will be for a particular service (e.g., minor frustration or a matter of life and death).

2. Consider the potential ways in which the system could fail. This is less about enumerating possible causes of failure, and more about identifying dependencies of the system whose failure could be disruptive.

3. Consider what can be done to reduce the impact of these failures and whether it is possible to contain a failure so that it does not bring down the entire system.

4. Consider how the system will be monitored for failures.

5. Consider what parties will need to respond to failures and what resources these parties will need.

6. Consider the ways the service could be restored quickly, even if the original failure cannot be immediately fixed.

3.5 Privacy

Privacy is a party's ability to control the exposure of data associated with them. Privacy concerns about smart city plans were one of the most often raised constraints the authors heard at conferences and interviews. This concern is mirrored in numerous news reports and research papers concerning privacy and smart city systems. [27] [28] [29] For these reasons, it is especially crucial that smart city projects address, and be seen addressing, privacy concerns.

Privacy covers the collection, storage, and use of a person's information and the degree of control that person has over those activities. It is important to note that **privacy is far more than avoiding the exposure of private data.** While keeping collected information confidential is part of privacy, it is only one part, and privacy violations can (and do) occur even if the individual's information is never exposed to others. "Personal information" itself is a broad topic, including data about identity, an individual's physical body, location and movement, communications, transactions, and territory, including personal space, objects, and property. [30] As a result, serious thought needs to be given to how a smart city project might touch on privacy issues before adequate solutions can be identified.

Data collection considerations are a critical first part in any privacy solution. Most smart city systems generate data of some kind and, due to the proximity these systems can have with individuals, this information can be personal in nature. This can happen even if the system is not intended to collect personal information. Even though the personal information was not the target of the collection, it can still lead to a breach of individual privacy. For this reason, **it is important to know what information is collected, and not just what information is intended for collection.** Any data collected that could have privacy implications will put a privacy management debt on all systems receiving, transmitting, processing, and storing that data. It might be desirable to disable the collection of the relevant data, or pre-process the data (e.g., anonymize or aggregate) at the point of collection to reduce or eliminate privacy implications.

Data use is also an important privacy consideration. Community members might allow their personal data to be used for some purposes, but might consider it a privacy violation for that same information to be used for other purposes. For example, citizens might be willing to accept the presence of license plate readers that can automatically bill users for tolls and parking as a convenience, but if that same information were made available to retailers to target advertisements based on where they were parking, some might consider this an unacceptable privacy violation. And to be clear, some citizens might not be comfortable with the automated billing scheme in the first place. To help manage privacy, **it can be important for those whose information is being collected to understand how the data is being used**, and potentially be able to set controls on its use. Awareness of intended use is not only necessary for citizens to exert control over their personal information, but could also make them more accepting of certain data collection, especially if they see that it provides a benefit for them or if they feel it serves a good cause.

For IT systems, a typical privacy control mechanism consists of opt-in/opt-out controls. That is, an individual is given the option to consent to share their data at the outset of using an application. This might work with IT-based public service systems, but is less effective when data is collected by passive sensors, such as cameras. For example, there is no way to opt out of being recorded by security cameras beyond simply avoiding entering areas where they are present, which may not be a practical option if these cameras monitor public places. [30] **In cases where opting out of data collection is not possible, it is important to make sure the public is aware** of these sensors, how the data is used, how the data is protected after collection, and how protections are monitored and verified. So, while the public may not be able to control whether or not the data is gathered, they will at least be aware of the potential privacy impact of these systems.

Another important consideration is the degree to which data can be associated with an individual. **When information on an individual is collected, the individual can be explicitly identified or they can be coded, de-identified, or anonymized.** When identity information is coded, a combination of letters or symbols is used in place of a personal identifier, and information can be traced back to the source by someone with the code or key. [31] While a person's name or other information might not be present, it is important to note that coding still tracks people individually, and there may be ways to link these records back to identifying information.

De-identification, per the St*andards for Privacy of Individually Identifiable Health Information* by Health and Human Services in response to the Health Insurance Portability and Accountability Act of 1996 (HIPAA), records data without any individually assigned code. [32] De-identified data removes information about the individual, relatives, etc., that could be used alone or in combination with other data to identify an individual. [32] This means that an individual record is still associated with a single person, but it will generally not be possible to tell, given a collection

of records, whether that collection represents a single person or a group. In other words, the records lack a way of grouping by individual. Because each record is still associated with an individual, it remains possible for analytics to re-identify the subject of the data through correlation of specific data elements. However, it will not be as easy to do this as with coded data.

Anonymous data means that individuals cannot be identified within the data. Generally, this occurs either because the data represents an aggregate over time (e.g., counts of how many people entered an area during a given day) or because the data simply lacks the granularity to identify individuals (e.g., a pressure plate that can distinguish between some and no pressure on it, but cannot measure differences in weight).

It is important to note that these methods of obfuscation are not infallible. **Inference and analytics make it possible to combine data sets to derive identities.** Smart city technologies such as cameras and monitors, geo-location tracking enabled on smart devices, and commonplace technologies such as credit card transactions can be combined to identify the location and identity of an individual, even when privacy protections are in place. [30] Because of this, it is important to consider the privacy implications of data collection in aggregate, considering how multiple sources might be combined in a way that potentially creates privacy issues greater than those of any individual data source.

Control of data that has been collected is another important privacy element. Clearly, the security of this data (specifically its confidentiality and integrity) needs to be preserved (these elements are discussed more under the security trust characteristic). Many of the recent headlines concerning privacy breaches were due to groups failing to securely store collected data. [33] [34] **It is also important to consider who will have access to the data, as well as which parties will have control over the data.** For example, some products send data to the product's vendor rather than, or in addition to, those who purchased and operate those products. Some vendors collect the data from their products and then provide access to this data set to customers as part of their service. [35] While the collection and retention of data by third parties can be a source of concern, it can also provide benefits. It might be useful for the city to keep certain data "at arm's length" to avoid potential privacy concerns. For example, in a survey of citizens' opinions, the city would want to see the aggregate responses, but there might be concerns about allowing the city to attach individual responses to persons or areas. Having a third party handle the raw data and then deliver a summarized report can enhance the overall privacy characteristics of the process. In such cases, having third parties manage the data is not necessarily a bad thing. However, how these commercial entities secure this information and the uses to which this data is put should be clearly understood, as should the terms and expectations of all parties involved.

If cities store data themselves, this data might be subject to regulations that dictate how collected data must be handled. For example, the privacy policies of Seattle [36] and Kansas City [37] both include requirements that citizens be allowed to view and correct information collected about themselves, which has implications for the costs of operating and maintaining the data collection systems. Similarly, once data is in the government's possession, there may be a requirement to release the data to the public, which can have privacy implications. For example, traffic camera data released by a city could be used by regular citizens to track others. [38] In fact, multiple city leaders to whom the authors spoke noted that some of their systems did not retain collected data in order to avoid the added complexity this would entail. As a result, understanding the potential implications associated with managing one's own data collections is another privacy-relevant concern.

The significant concerns citizens have voiced regarding their privacy have, unsurprisingly, translated into a number of regulations. Conformance to regulations is discussed in greater detail in the section on compliance, but some observations are worth emphasizing. **Be sure to consider the scope of the parties potentially impacted by data collection, as this can have implications for the applicability of certain regulations.** Data collected via IT sources, such as over websites, could come from users from anywhere in the world. For example, the European Union's General Data Protection Regulation (GDPR) would apply to data collected in the United States on EU citizens. [39] Moreover, some vendors whose products collect and report data might be located in different jurisdictions, or even in other countries. All of these circumstances can have implications as to which regulations are applicable and what steps need to be taken to conform with them. Altering a system to become compliant with regulations after deployment is often more expensive than implementing such compliance at the outset of the project, so it is usually worth the cost to try to be thorough in identifying relevant regulations early in the project.

Given the degree to which users are constantly exposing private information (intentionally or otherwise) to commercial companies through social media, linked apps, rewards clubs, and other mechanisms, it can be tempting to adopt a stance of "moral relativism" and simply note that one's own collections are far less intrusive than those to which many citizens seem to have become accustomed in dealing with commercial entities. However, doing this ignores two important points. First, the behavior of these commercial companies is not something to be emulated, since many people are not happy with the volume of data collected. [40] [41] Even those who willingly surrender some of their privacy as a cost of access to a particular service may not be happy doing so. [42] The second key difference lies in the different roles, and particularly the power differential, between commercial enterprises and governments. Both might collect the same information, but governments have the power to use that information to deprive people of their freedom and property, through arrest or fines. [43] As a result, data that governments collect and use has the potential for significantly greater impact on people's lives. For this reason, **actions by a commercial company that might be allowed (or at least tolerated) might be seen as an unacceptable privacy violation if done by a government entity**.

3.5.1   Key Privacy Considerations

1. Consider what data is being collected (intentionally or unintentionally).
2. Consider how collected data will be used.
3. Consider what mechanisms can be put in place to keep collected data from being put to new uses without explicit consent.
4. Consider whether the data needs to be de-identified or anonymized.
5. Consider whether there are ways that de-identified data might be re-identified based on other collected information.
6. Consider how those whose data is collected will be made aware of what was collected and how it is to be used.
7. Consider whether there will be a way for system users to opt out of having their data collected or used for a particular purpose.
8. Consider whether there will be a way for system users to revoke consent after opting in.

9. Consider whether there will be a way for system users to view the data that has been collected about them.

10. Consider whether there will be a way for system users to correct information that has been collected about them that they feel to be erroneous.

11. Consider what entities will see the collected data (e.g., third parties such as product vendors, city governments, law enforcement).

12. Consider whether there are reasons that the city governments should not be allowed to see raw collected data.

13. Consider what controls are needed to ensure that all those who have access to the data respect the privacy of users and relevant privacy regulations.

14. Consider what data will be transmitted and stored, and how the security of the data will be preserved.

15. Consider whether there are any regulatory/process implications created by storing the collected data.

16. Consider the scope of parties whose data will be collected, and the implications of this scope on the set of regulations with which the system must comply.

3.6 Maintainability

Maintainability is the assurance that the system will remain operational in the future and can adapt and grow as needed. From an IT perspective, software maintenance may require patching and upgrading. If the component resides on a network, this might be done remotely. If not, it may be necessary to have physical access to perform such updates, which can add considerably to the time and expense of maintenance. From an OT perspective, physical components might need to be replaced or modified (either for predictive maintenance or to repair breaks). In both cases, maintenance costs can be reduced by ensuring that updates/replacements are easily performed.

It is important to note that **maintainability is not just about preserving the status quo, but also encompasses system growth and evolution.** Most smart city systems deploy first as pilots, with limited scope. It is important that these systems be able to grow beyond their pilot phase. Such growth might simply be a matter of scaling up the project, but will more likely involve other adaptations as well. Such adaptations might be necessary to address discovered issues, or to fit the system into differing demographic, geographic, or infrastructure contexts. The design of a system should also include the ability to respond if the system sees greater adoption/demand than anticipated. Additionally, given that the goal of the system is not just to operate but to provide a service, changes in use patterns, technology, or other factors could require that the system adapt in order to continue to fulfill its intent. Finally, smart city projects can often build on each other. For instance, installation of fiber connectivity can be used by multiple services, kiosks and poles created for one service often can be used by other services, and many data-gathering projects could find their data used by other smart city capabilities. Ideally, smart city projects should make such reuse and extension possible.

One of the key elements of a maintenance strategy is identification of funding. **Any project needs to identify sufficient funding not just to support design and deployment, but to maintain that**

system after deployment has completed. There are several ways to support a system's funding beyond depending upon yearly funding out of a city's budget. These include:

- Monetization of infrastructure – Find ways to generate revenue from the infrastructure the project needs to build. Examples include selling use of bandwidth on fiber cables or leasing space on poles.

- Revenue sharing opportunities – This involves partnering with private companies, and potentially creating revenue generating opportunities in return for having them help support ongoing maintenance of the system.

- Monetization of data – Many systems involve collection of data that can be monetized. However, great care needs to be taken to ensure that this does not violate expectations of privacy, security, or open data policies.

- Fee and fare collection – In these cases, users pay a fee to use the system's service.

- Cost savings over prior behaviors – In some cases, a system might present a significant cost savings over prior practice. The system still requires an ongoing stream of funding from the city, but it can be justified by demonstrating that this funding is significantly less than for the system that was replaced. For example, significant repair costs are curbed by early detection and incremental maintenance. [44]

In short, while smart cities will often require a significant capital outlay to develop and deploy a new system, such projects can also represent revenue opportunities for cities. Project stakeholders should carefully consider possibilities regarding the latter.

In addition to developing a funding plan, **it is important to provide solid measures of costs versus incomes related to the system.** Many smart city leaders noted that concerns about financing were frequently raised, both by citizens and by governments. Providing reasonable estimates of expected costs and incomes ahead of a project, as well as detailed tracking of these values after its deployment, will be necessary to win over these groups.

Beyond a secure funding stream, another requirement is access to the parts and services on which the system depends. There needs to be a way to acquire replacement parts and components for repairs, patches for software issues, upgrades to support hardware and software, and the means to extend the scope of a deployment (e.g., expanding the range of a project through deployment of new components in new locations). With regard to services, some smart city systems might use third parties to support data processing and collection. In all cases, the system needs to be designed to avoid overdependence on a single source. Companies can go out of business or cease support for a particular product or service. **If the smart city system is dependent upon some component that is no longer available, the survival of that system can be jeopardized.** One way to mitigate this risk is to employ products and services that conform to technical standards. Conformance to standards increases the chance that other, compatible components can be found.

In addition to needing technical parts and services, smart city systems also require personnel with the skill and expertise to operate, maintain, and repair the system. This means that **there must be mechanisms to retain institutional knowledge in the (expected) case that members of the support team retire or switch jobs.** It also requires people who can step in to fill the vacuum created when others leave the team. For the former, documentation of procedures plays a critical role. In the latter case, it can be used to support training to produce "bench depth" in the

maintenance team so there is more than one person who can do any given job. Multiple smart city leaders also cited the importance of establishing links with local educational institutions that could be sources of new hires to replace outgoing staff.

Finally, an important consideration for maintainability is that, while every city and smart city system is ultimately unique due to differing needs and context, too much uniqueness can be problematic. Cary, North Carolina's Chief Innovation Officer observed, "**The key to longevity is repeatability.**" [45] Elements of a system that are common across multiple cities will prove easier to maintain, both because cities will be able to learn from each other and because commercial companies will have a greater incentive and ability to continue support for those elements. By contrast, a system composed almost entirely of custom components will be especially vulnerable to supply chain issues if a necessary vendor ceases support (or significantly increases prices) for those components. For this reason, it can be useful to consider what other, similar projects are doing. While every smart city project needs to adapt to its local context, this should be balanced by the need to consider the potential cost impact that adaptation could have on future maintenance.

3.6.1 Key Maintainability Considerations

1. Consider how long components of the smart city system are expected to be operational, if properly maintained, before needing to be replaced.

2. Consider whether physical components of the smart city system can be readily replaced/upgraded.

3. Consider whether software components of the smart city system can be remotely patched/updated.

4. Consider probable ways in which it might be desirable to grow the smart city system (e.g., expanding its range or adding new capabilities).

5. Consider how ongoing maintenance of the smart city system will be financed.

6. Consider how costs and incomes related to ongoing maintenance of the smart city system will be measured.

7. Consider what vendors the smart city system depends upon and whether there are alternatives to these vendors if the vendor is no longer able to adequately support their part of the system.

8. Consider how institutional knowledge regarding the operation and maintenance of the system will be preserved.

9. Consider how the team responsible for maintenance and operation of the system will respond to the loss of a member.

3.7 Compliance

Compliance is the ability to conform to requirements associated with a particular context. Smart city systems may need to comply with a variety of legal, regulatory, and, potentially, integration requirements or standards, which may come from the federal, state, or local level. **The requirements with which smart city systems are obligated to comply may vary depending on who is acquiring the system, funding the system, or operating the system.** For example, a

private sector entity is not subject to Freedom of Information Act requests like its public sector counterparts. Even within a single city government, different departments might be subject to different compliance requirements. Thus, one of the first challenges in a smart city project is simply identifying all the sources of requirements with which the system must comply.

In addition, **stakeholders in a smart city system could impose their own requirements.** Vendors might require that certain components be interoperable with their products (potentially, exclusively with their products). Investors and community leaders might add requirements that correspond to their values (e.g., ensuring that certain communities are guaranteed access to the system) or that address concerns they fear could compromise the project (e.g., requiring conformance to certain security practices that might not otherwise be required). Insurers might impose requirements for the system to be covered by their policies. These requirements need to be clearly spelled out so that there is clear agreement when they have been met and to ensure that they are not overlooked.

The fact that **smart city systems tend to have significant public visibility adds an additional layer of compliance: compliance with cultural expectations.** Those expectations will vary depending on the norms of the community. For example, different communities might have different expectations regarding the role of commercial entities in the operation of public services or the extent to which government should be involved in daily life, or might expect certain elements of local character to be preserved. While local leaders of a smart city project might be aware of these expectations, external companies and consultants might not. Identifying these requirements will require careful listening by outsiders and careful review by locals.

Conformance with technical standards is another important aspect of the compliance characteristic. While closely related to the integration characteristic, integration focuses on making components work together, while **compliance also covers conformance with standards or frameworks beyond what is needed for interoperability with the other elements of the system.** While compliance with these requirements might not be strictly necessary for the system to work, this added compliance conveys a couple of advantages. First, it affords greater integration opportunities in the future facilitated by these additional standards. Second, it means that the system is aligning with best practices, as identified by those standards. Useful guidance for identifying points where standards might be employed can be found in the National Institute of Standards and Technology (NIST) *Internet-of-things Enabled Smart (IES)-City Framework*, which provides guidance on identifying pivotal points of interoperability. [46] By conforming to standard practices at these pivotal points of interoperability, the system is positioned to integrate with new components and systems in the future.

3.7.1   Key Compliance Considerations

1. Consider the likely sources of regulations/laws with which the smart city project must comply.

2. Consider who the smart city project's stakeholders are (e.g., citizens, underwriters, commercial partners) and whether they have implicit or explicit compliance requirements.

3. Consider who can review the smart city system design and speak to its alignment with cultural expectations for its expected set of users.

4. Consider whether relevant technical standards can be employed by components within the smart city system.

3.8 Well-being

Well-being is preservation of livelihood, quality of life and environs, and minimization of personal disruption. Well-being, though often viewed as abstract and subjective, is very important from the perspective of public service systems. To thrive, cities must be desirable places to live, work, and play. Many smart city projects explicitly cite the objective of drawing businesses to the city through making the city a more attractive place to live for employees. Even when this is not an explicit goal of a project, smart city efforts that alienate citizens or make it harder for companies to attract employees will almost certainly be viewed as failures.

There is often a perception that qualities like well-being and livability are too "squishy" and abstract to be measured, much less managed. However, this is not the case. Many commercial companies put considerable investment into measuring the "intangible" aspects of customer experiences and drive their business decisions based on these measurements. [47] Companies also invest in measuring the happiness of their own employees, believing that happier employees are more productive. [48] There are tools available and companies that can assist cities in measuring the well-being of their citizens. A key element in understanding well-being is focusing on the experiences of those using or otherwise impacted by the smart city system. In technical projects there can be a tendency to focus on performance measures of the system, but such measurements are inward looking, rather than outward looking. [47] **Measures of well-being need to be focused on the experiences of those impacted by the system, rather than simply measures of the system itself.**

Preserving and improving well-being often means ensuring that the smart city project is a good fit within its local context. Even when smart city efforts are being undertaken at a regional level, it is important not to assume that a one-size-fits-all approach will work in terms of a solution or engagement plan. It is important to be aware of culture, history, and citizen concerns. Consider the following examples of two communities that would like to install kiosks to provide easy access to information. The first community would like the kiosks to provide access to community resources for citizens (e.g., e-gov kiosk [49]). The culture in this community dictates that blending in with the aesthetic of the community is very important. The kiosks may need to be unobtrusive (e.g., inside public buildings) to be accepted culturally. The second community is economically dependent on tourism and would like the kiosks to provide guidance to visitors and recommend local businesses. For the second community, the kiosks may need to be eye-catching and situated in view.

Maintaining well-being often requires performing research into the history of a location, as well as identifying important business and cultural considerations. Citizens' concerns should be heard, because even mistaken perceptions related to a technology cannot simply be dismissed. These **perceptions must be managed, because implementation of technologies that ignore citizen concerns, even if unfounded, can lead to pushback and, potentially, project failure.** For example, installation of new cell towers (i.e., base stations) are sometimes met with concerns from citizens about radiofrequency (RF) waves and the potential for cancer. Studies by the American Cancer Society, Federal Communications Commission, International Agency for Research on Cancer, Environmental Protection Agency, and World Health Organization (WHO) have each explored the issue and have all determined that RF exposure at the base of cell towers is well within (potentially thousands of times below) safety limits. [50] [51] However, rather than simply dismissing these concerns as false, planners need to actively engage the community segment to

foster better understanding. According to WHO, "education programs as well as effective communications and involvement of the public and other stakeholders at appropriate stages of the decision process before installing RF sources can enhance public confidence and acceptability." [51] Other smart city efforts might also require education or adjustment to manage the perception of risk, even though the actual risk is mitigated or negligible.

It is important not to underestimate the importance of well-being as a key trust characteristic of smart city systems. The well-being characteristics of a project can impact the degree of adoption a smart city system enjoys and can also have an economic impact by making a city more or less desirable for businesses. Beyond this, **well-being is the characteristic most directly related to trust in the system as a whole.** If citizens' well-being is addressed and supported, then citizens are given the impression that the system and its stakeholders are aware of and considering their interests. This, in turn, can help improve trust in other elements of the system, including faith in and support for other trust characteristics. The Chief of Civic Wellbeing for the City of Santa Monica observed, "Trust is at the core of well-being." [52] For this reason, it is critical that smart city systems address and support well-being and do so using relevant and effective measures.

3.8.1 Key Well-being Considerations

1. Consider the ways that the smart city system and its components will impact citizens.
2. Consider what parts of the smart city system will be visible/audible to citizens.
3. Consider what aspects of the smart city system will force/encourage changes to behavior by citizens.
4. Consider whether the smart city system will change the property value of certain community areas and whether the change will have secondary effects.
5. Consider likely citizen concerns with regard to the smart city system and how those concerns can be managed.
6. Consider how necessary changes in citizen behavior driven by the smart city system can be managed. Even if the goal is to change behavior, then the question becomes how to do this while minimizing disruption.
7. Consider how well-being of citizens will be measured in relation to the impact of the smart city system.

3.9 Fairness

Fairness encompasses concepts such as a lack of bias, equal access, and transparency. Smart city systems are generally understood to be implemented for the benefit of all citizens within a community, and some are expressly built to work toward closing opportunity gaps, such as the digital divide. [53] In these cases, communities are attempting to provide or improve equitable access to resources. From a public service system perspective, there are many reasons for the implementation of "fair" smart city systems, such as legal or policy requirements, improvement of the overall health of the community, as well as moral and ethical considerations.

There can be a perception that, by making a system "data-driven," one is creating a system that is inherently unbiased and fair. The fact that humans are not involved in the decision-making process can be seen as a protection against inherent human biases. Data-driven systems are *consistent*,

which is part of being fair, but this is not sufficient to create a system that is unbiased and fair. In fact, **algorithms can often exhibit a bias, even when developed with the best of intentions.** "Data is frequently imperfect in ways that allow these algorithms to inherit the prejudices of prior decision makers. In other cases, data may simply reflect the widespread biases that persist in society at large. In still others, data mining can discover surprisingly useful regularities that are really just preexisting patterns of exclusion and inequality." [54] Historical data going into a system, which is often how systems are "trained," may be biased (e.g., reflecting historical segregation) or inaccurate (e.g., based on out-of-date demographic or infrastructure information). [55] Limitations in available data can force a correlation between weak indicators or fail to account for nuances a human would recognize, leading to questionable conclusions. All of these factors can cause algorithms to behave in a manner that creates bias in the results. For this reason, data-driven processes need to be reviewed and monitored for bias in the same way that a human-managed process would be.

Ensuring fairness in smart city systems is made more difficult by the fact that many of the decision-making algorithms in these systems are "black boxes," meaning that it can be difficult to see how a particular conclusion was reached. For maintainers of these systems, this can make tracking and correcting bias challenging. This can also make it difficult to justify decisions coming from the systems, since the reasoning cannot be shared. This inability to explain reasons for decisions can reinforce citizens' perceptions of unfairness if the algorithm makes decisions with which they disagree.

It also needs to be observed that **being "right" might not always be adequate for a decision to be "fair."** In some cases, while a strictly objective view of the data might make a decision seem justified, when that decision is placed in a broader context, it could be seen as unfair and biased. For example, in 2016, online retailer Amazon decided to offer same-day delivery service to some of its customers. It based the decision of where to offer this service on the concentration of existing members of its Prime service within regions. However, after the service was rolled out, mapping the regions of same-day delivery against local demographics revealed the service was frequently offered to predominantly white neighborhoods while excluding predominantly black neighborhoods. Though it was true that more Prime members were in the designated one-day delivery zones, the reality of which groups were included and excluded from the service produced public backlash. In some cases, Amazon overrode the data-driven decision and expanded same-day delivery across cities. [56] The bottom line is that **bias can be an emergent quality of a system that is only visible when that system is placed in a larger context.** For this reason, it is important to review decisions within that broader context to determine if they create any unintended outcomes.

Measuring the impact of a system is an important step in detecting and addressing bias. If a particular group is not utilizing a service (or represents the predominant users of the service), it is useful to consider why this might be the case. The section on measuring utility/impact discusses metrics and measurement in more detail, including the importance of establishing a baseline for those measurements. In the case of detecting bias, however, be cautious about baselines, which can reflect long-standing inequalities in a location. Measurements that indicate that the baseline is preserved or only moderately improved could still indicate bias in the algorithm (potentially due to the use of historical data in training the algorithm). **In terms of searching for bias, looking at the impact of the systems using demographics and other factors can be helpful for service**

operators. If disparities in use are identified, the system operators will need to figure out why this is happening and what, if any, steps should be taken in response.

3.9.1   Key Fairness Considerations

1. Consider how the smart city system will be reviewed for bias before deployment.

2. Consider how the smart city system will be monitored for bias after deployment.

3. Consider whether there any specific fairness measures to which the smart city system must adhere (e.g., law, regulation).

4. Consider whether the smart city system can be "tuned" to address emergent bias.

5. Consider whether measurements will be shared with citizens to help "convince" them of the fairness of the smart city system.

3.10 Integration

Integration refers to the ability to work with or alongside existing infrastructure and processes. Integrating with infrastructure involves working with any tools, components, or devices that might already be in place. Integrating with processes involves aligning inputs and outputs between old and new systems and avoiding undesired disruption of existing practices and procedures. If the new smart city system is replacing an existing system, integration will entail making sure that all activities that feed into or use data from the old system are able to continue. Even if the new smart city system is not replacing existing systems, it is still going to need to integrate with processes and infrastructure for other systems. As such, integration is a universal requirement for any smart city system.

Technical integration (i.e., making sure components are able to interoperate) is a mature field, and professional "integrators" exist to help facilitate the process. When developing the technical integration plan for a project, bringing in these experts can be very beneficial. Two additional elements of technical integration should be considered. First, virtually any two software components can be made to work together given enough "glue code" that translates between them. However, **highly customized integration interfaces can be brittle and expensive to maintain.** For example, if one software component undergoes an upgrade or revision, it can change the component's inputs and outputs, requiring that the integration connector be rewritten. Instead, if possible, try to employ software components that use standards in their interfaces. This can simplify integration and make it more robust against changes in the connected components; it can also make it easier to replace a software component from one vendor with a different vendor's component if necessary. Note that the use of standards seldom completely eliminates the need for custom integration, since most vendors will want to add custom features for market differentiation. However, standards should at least reduce the need for custom software integration connectors.

Standards can also provide a similar benefit when integrating hardware components. In this case, standards would apply to types and sizes of cables, power connectors, and other types of "connectors" for inputs and outputs to the hardware component. Vendors that employ connectors with a customized size, shape, or other input (e.g., voltage) will complicate integration with hardware components from other vendors. Similarly, it can be useful to consider how modular a complex hardware component is as a whole. Specifically, if a multi-part component fails, is it necessary to replace the entire unit, or can the replacement be limited to only the broken part? If

the latter, can the replacement be performed "in house" or will an outside technician be required to make the repair? Understanding such issues before selecting the system's hardware components can have a significant impact on managing future maintenance costs.

A second important consideration for technical integration is potential indirect use of the new system. A system that generates data (e.g., embedded sensors) could see its data used by many downstream systems, and it might be that not all of those systems are immediately known. This is especially true if data is published using an open interface (described later in this section), as such an interface could encourage customized development of tools using the system's information. **When replacing an existing system in a way that alters the data that is produced (including but not limited to alteration of the data's structure, frequency of delivery, data semantics, or sensor precision), some effort should be taken to identify downstream users of the data** or at least publicize the upcoming change. Failing to do this can create unexpected disruptions when the transition occurs.

Smart city systems are often intended to create efficiencies or promote interoperability between previously stand-alone systems. For public services, this may mean government departments that had previously had independent processes will now be working together. In these instances, **it may be necessary to integrate the processes prior to implementing the smart city system.** The departments/disciplines involved will need to examine their processes and identify opportunities for integration and potential streamlining. For example, if a city would like to automate the process for applying for a restaurant license, their building department (e.g., building permit, occupancy permit), health and safety (e.g., inspection, food service license), fire department (e.g., fire protection permit, commercial kitchen license), etc., could all participate in clarifying their own processes and working together to develop the unified process that can be automated within the smart city system.

Some citizens may be unwilling or unable to transition to a new system. The ability to use a new system can be dependent on resources (e.g., access to the internet, smartphone, computer), skills (e.g., computer literacy), and willingness to change. **To avoid excluding citizens from public services when a smart city system is implemented, old ways of providing the service may have to coexist with the new system for a period of time.** For example, if a smart city system is put into place to provide easy access online to city government resources, this system can coexist with other methods of communication such as providing public service information in citizen water bills, thus providing an opportunity for those without online access to receive the same information. [57] Alternately, it may be necessary to maintain the old system in parallel to the new system to continue serving citizens who cannot make the transition. While the new system might have been intended to cut costs, the need to maintain both new and old systems in parallel could have the opposite effect in the short term. Effort should be taken to identify potential challenges for citizens transitioning from old to new systems, and identify steps to ensure that those who do not make the transition remain served.

Finally, **it can be useful to consider forward integration—that is, supporting integration with tools and components that do not yet exist.** One way of doing this is by supporting open interfaces that allow new components to use data generated by other components. Multiple cities have created such open interfaces and data portals, and encouraged developers to develop new tools using this information. [58] [59] Such efforts can be highly valuable because they can foster new and innovative uses of smart city data at little cost to the city.

3.10.1 Key Integration Considerations

1. Consider the other components with which the new smart city system will need to integrate (e.g., use data produced by the other system, or provide data to that other system).

2. Consider whether other systems are dependent on the existing system that the new smart city system will replace, whether the new system will introduce changes that those other systems must account for, and how downstream systems will be alerted to changes.

3. Consider whether technical standards can be employed at any of the integration points with the smart city system for hardware or software components.

4. Consider whether repairs to hardware components can be made piecemeal, or whether the whole component must be replaced in the case of component failure.

5. Consider what government departments and other institutions the new smart city system will touch upon and whether they need to adapt their processes to work with the new smart city system and/or to work with each other.

6. Consider the potential barriers that would make adoption of the new smart city system challenging or undesirable to users.

7. Consider what can be done to facilitate the transition by citizens to the use of the new smart city system with minimal disruption and how to accommodate those unable to make the transition.

8. Consider whether the new smart city system and an existing system will need to operate in parallel for a period and whether there are sufficient resources to support parallel operation during the transition period.

9. Consider whether there is a possibility to support open interfaces or other mechanisms in the smart city system to allow forward interoperability.

3.11 Measuring Utility/Impact

Utility and impact are the ability to perform the needed job and/or create the change expected by the system's stakeholders. All smart city projects have goals, and meeting those goals tends to be the focus of most of the system's design and implementation efforts. The need for a project to meet its intended objectives is obvious. However, potentially less obvious is the need to establish ways to determine the degree to which these goals are met. As such, this trust characteristic is concerned not only with designing the system to meet its goals, but with metrics that can measure the degree to which these goals are met, and the ability to expose this information to appropriate parties, including citizens.

Measuring utility and impact requires creating a robust baseline, defining metrics, and gathering appropriate data. [60] For example, if smart trash cans are expected to reduce the cost of sanitation services by reducing the number of trash pick-ups and smartly routing trucks, a baseline should be established by measuring the distances traveled, time spent, and costs of sanitation services before the smart trash cans are deployed. After deployment, measurements of those same items can be used to determine whether the desired efficiencies have been achieved.

Metrics, and especially the display of these metrics, can be powerful tools in helping to justify a smart city project after its deployment. This justification may be necessary to ensure sufficient

funding for ongoing maintenance and/or to support expansion of a project beyond an initial, trial deployment. Whether the system has demonstrable utility, supported by real measurements, could also have political and public perception implications. **Consideration should be given both to what information is collected and to how it is exposed.** In some cases, it might make sense to provide a web portal that allows citizens to monitor the measures directly; in other cases, it might be more effective (or necessary for security or privacy reasons) to provide only summary information. In either case, effort should be taken to make this information easy to understand—a giant spreadsheet of numbers is likely to be less helpful than a graph that provides a visual representation. Having a consultant experienced in data visualization can be very worthwhile, especially given that decisions made by citizens and city leaders based on this data could impact the system's longevity.

In addition, it can be useful to consider what parts of a system are visible to citizens compared to the parts that are "behind the scenes." For example, some smart city projects involve scanning roads for signs that potholes are developing and implementing a back-end infrastructure that routes this information to the city department that can act on it. However, from a citizen's perspective, the only part of this system that is likely to be noticeable is how long it takes for potholes to be repaired. While the system might quickly detect potholes and automatically assign a work item to the responsible city department, if this does not translate to a quick repair (for example, because the parties responsible for repairing the potholes lack the resources to respond quickly), then citizens are likely to be unhappy with the service. While metrics might show that the smart city system has led to faster pothole detections and work assignments, those metrics will likely be less important than measures of the time it takes for potholes to be repaired. Therefore, **it is helpful to consider a range of metrics, and be sure to give consideration to metrics that will help expose what citizens are seeing with regard to the smart city service's outcomes.**

It is also worth considering indirect measures associated with a smart city project. For example, the goal of a project might be to reduce traffic congestion through adaptive traffic light timing. Direct metrics could be measures of how long it takes to get from one end of the route to another. However, other measures might be of interest as well. It might be worth measuring traffic on surrounding roads that were not fitted with adaptive traffic lights. Perhaps these metrics would reveal that congestion on these roads was reduced as well (perhaps due to fewer cars being re-routed to side streets to avoid traffic), which would be another way to justify the efficacy of the project. **Consider the ways that a project might have indirect or secondary impacts, and consider establishing ways to measure these as well.**

Metrics also have a role in supporting maintenance. Fluctuations in measurements may indicate a problem and/or a need to make changes to the system. For example, metrics monitoring use of bike sharing that identify a decline in use may point to a need to redistribute those bikes to different locations or to assess the bikes for maintenance. In this way, **measures of impact can help identify opportunities to increase impact.** In short, metrics should not be "passive," but part of a feedback loop that can drive improvements.

It is important not to only consider impact, and measures thereof, as an "after deployment" issue. **Measures of how a project's deployment is proceeding can also be useful tools to allay concerns by citizens and city governments.** Consider what information might be tracked and how it might be shared. Similarly, identifying the specific metrics a project will use to measure its impact can be a useful part of the process to "sell" the project to potentially skeptical citizens and city governments. Citizens need to be convinced that a project has relevance to them if they are

going to adopt new services, and failure to convince citizens that a service will ultimately benefit them was cited as a historical technology adoption barrier. [61] By calling out measures that are aligned with real citizen needs, one can often make a stronger case as to why a project is relevant. In all these cases, these steps take place prior to deployment of the system.

Finally, projects should be sure not to neglect measures of trust. This includes measuring the behavior of system components that are included to support or preserve trust characteristics. This also includes measures of the public's trust in the system. The latter measures will reveal whether the efforts to support trust characteristics were effective and may identify opportunities to add or modify system components to enhance public trust. Measures of public trust can be direct, such as polls asking users how they feel about the system or monitoring public discussions about the system in social media. Measures of public trust can also be indirect, such as looking at changes in usage patterns of the system to reveal whether the public is embracing or avoiding the new system. **System designers should consider how they will measure the public's trust in the system as part of their overall measurement strategy as a means to validate and improve support for the trust characteristics.**

3.11.1   Key Measuring Utility/Impact Considerations

1. Consider what can be measured to better understand the user experience with regard to the smart city system.

2. Consider what can be measured to better understand the effectiveness and utility of the smart city system.

3. Consider what metrics can best demonstrate the relevance/importance of the smart city system to stakeholders and how such measurements will be shared with stakeholders.

4. Consider what can be measured that can feed back into the operations/maintenance processes of the smart city system to improve performance.

5. Consider whether there are indirect impacts of the system that might be useful to measure.

6. Consider what baselines need to be gathered ahead to prepare for measurements of the system.

7. Consider how measurements will be gathered.

4   Conclusion and Use of This Paper

This paper presents 11 key trust characteristics that need to be considered in any smart city project. They represent aspects of a smart city system that must be supported and preserved for users of the system to trust that it can be depended upon to maintain benign behavior. These trust characteristics come from the union of the three key domains of practice that contribute to smart city projects: IT, OT, and public service. Each of these domains brings important perspectives and expertise, but domain practitioners can prioritize trust characteristics with which they are familiar over characteristics that are not typically encountered in their regular practice. While this might not be problematic when each domain is acting in isolation, in projects that cross domain boundaries (such as most smart city projects), failure to account for trust characteristics can be a source of risk. This paper presents a more holistic list of characteristics built upon the cumulative set of priorities of IT, OT, and public service domains of practice. The result is intended to be a

comprehensive treatment of key characteristics, without the biases that might be present in a treatment tied to a single domain of practice.

The primary audience for this paper are those in the early stages of developing a smart city project. This is because, in many cases, the best way to address a given trust characteristic will involve adding a stakeholder who can serve as an advocate of that characteristic. These advocates can help the stakeholder team as a whole understand and make decisions regarding a sufficient level of support for each characteristic and make appropriate trade-offs between characteristics when necessary. In addition, because trust in a project's stakeholders and the importance of relationships are all key to developing trust in the developed system, transparency and ongoing engagement with regard to characteristics throughout the design and implementation of the project are important. All of these decisions and actions need to be taken well before the final implementation of the project, so considering these factors at the very beginning of the smart city project is critical.

Each of the characteristics is presented in detail, including multiple observations regarding its scope and impact. In addition, each characteristic concludes with a list of key considerations that can be useful to understand the role the characteristic plays in a given system and the elements that might need to be included in the system in order to support that characteristic. Not all characteristics are equally important to all projects. As such, the goal of these sections is not to present a list of tasks for each project to accomplish, but to help readers understand the ways in which a characteristic is important, as well as the different ways it can manifest in a project. Smart city stakeholders can then use this understanding to identify the specific ways in which each characteristic aligns with and impacts their project's goals and activities. The specific mechanisms used to preserve trust characteristics will vary significantly based on the individual needs of each project and the types of technology and processes it employs. As all the listed trust characteristics have mature fields of practice associated with them, bringing experts in those fields onto a project is likely to be an important part of solution development.

In summary, use of this paper facilitates an early, but important, step in the process of developing a smart city system. It is hoped that readers will use the information about trust characteristics to consider the several ways in which users need to have confidence in a smart city system and ways in which their smart city project can meet them. Smart city projects are challenging endeavors and can fail or underperform for many reasons. As such, addressing all trust characteristics does not guarantee a project's success. However, by incorporating appropriate levels of support for all trust characteristics, projects mitigate many factors that can lead to failure. More important, they help a project's stakeholders create a smart city system that is more likely to be embraced by those it serves because those parties have faith in the system's behavior and dependability.

References

[1]   Smart Cities Council, "Definitions and Overviews," Smart Cities Council, [Online].
       Available: https://smartcitiescouncil.com/smart-cities-information-center/definitions-and-
       overviews. [Accessed 26 March 2018].

[2]   T. Maddox, "Smart cities expected to invest $80B in technologies in 2018," TechRepublic,
       2 February 2018. [Online]. Available: https://www.techrepublic.com/article/smart-cities-
       expected-to-invest-80b-in-technologies-in-2018/. [Accessed 4 April 2018].

[3]   Industrial Internet Consortium, "Industrial Internet of Things Volume G4: Security
       Framework," 2016. [Online]. Available:
       https://www.iiconsortium.org/pdf/IIC_PUB_G4_V1.00_PB-3.pdf. [Accessed 26 April
       2018].

[4]   I. Rook, Interviewee, *Advanced Wireless Networks: The Foundation for Smarter Cities.*
       [Interview]. 28 March 2018.

[5]   United States Department of Homeland Security, "Critical Infrastructure Sectors," United
       States Department of Homeland Security, 11 July 2017. [Online]. Available:
       https://www.dhs.gov/critical-infrastructure-sectors. [Accessed 10 April 2018].

[6]   R. Condon, "Users bypass security to get their jobs done," ComputerWeekly.com, 3 April
       2008. [Online]. Available: https://www.computerweekly.com/news/1308245/Users-bypass-
       security-to-get-their-jobs-done. [Accessed 27 April 2018].

[7]   M. Kassner, "Former NSA and CIA director recommends managing consequences instead
       of vulnerabilities," TechRepublic, 2 June 2016. [Online]. Available:
       http://www.techrepublic.com/article/former-nsa-and-cia-director-recommends-managing-
       consequences-instead-of-vulnerabilities/. [Accessed 23 February 2017].

[8]   Amey VECTRA Limited, "A Simplified Approach to Estimating Individual Risk,"
       [Online]. Available: http://www.hse.gov.uk/research/misc/vectra300-2017-r03.pdf.
       [Accessed 23 February 2017].

[9]   The.Project.Management.Hut, "Managing Project Risks," 23 January 2010. [Online].
       Available: https://pmhut.com/managing-project-risks. [Accessed 23 February 2017].

[10]   Cornell University and Penn State University, "Environmental Inquiry: Risk,"
       Environmental Inquiry, 2009. [Online]. Available: http://ei.cornell.edu/toxicology/risk/.
       [Accessed 23 February 2017].

[11]   A. Fabio, "Killed By a Machine: The THERAC-25," Hackaday, 26 October 2015. [Online].
       Available: https://hackaday.com/2015/10/26/killed-by-a-machine-the-therac-25/.
       [Accessed 12 April 2018].

[12]   B. Naylor, "This Doll May Be Recording What Children Say, Privacy Groups Charge,"
       National Public Radio, 20 December 2016. [Online]. Available:
       https://www.npr.org/sections/alltechconsidered/2016/12/20/506208146/this-doll-may-be-
       recording-what-children-say-privacy-groups-charge. [Accessed 12 April 2018].

[13] T. Brewster, "Warning: 50,000 Mi-Cam Baby Monitors Can Be Spied On With Ease," Forbes, 21 February 2018. [Online]. Available: https://www.forbes.com/sites/thomasbrewster/2018/02/21/50000-mi-cam-baby-cams-vulnerable-to-simple-spy-attacks/#47160eec1c7e. [Accessed 12 April 2018].

[14] P. Kral and C. Wright, "Incident Handler's Handbook," SANS Institute, 2011.

[15] K. Jackson, Interviewee, [Interview]. 16 February 2018.

[16] M. Gogan, "Insider Threats as the Main Security Threat in 2017," Tripwire, 11 April 2017. [Online]. Available: https://www.tripwire.com/state-of-security/security-data-protection/insider-threats-main-security-threat-2017/. [Accessed 26 July 2018].

[17] M. Nieles, K. Dempsey and V. Y. Pillitteri, "An Introduction to Information Security," National Institute of Standards and Technology, Gaithersburg, MD, 2017.

[18] J. D. Musa, Software Reliability Engineering: More Reliable Software, Faster and Cheaper, 2nd ed., AuthorHouse, Inc., 2004.

[19] L. H. Barg-Walkow, "Understanding the Role of Expectations on Human Responses to an Automated System," 2013. [Online]. Available: https://smartech.gatech.edu/bitstream/handle/1853/52909/BARG-WALKOW-THESIS-2013.pdf?sequence=1&isAllowed=y. [Accessed 18 April 2018].

[20] Environmental Compliance and Enforcement Network for Accession (ECENA) and The Regional Environmental Center for Central and Eastern Europe, *Overview of Industrial Risk Assessment.*

[21] P. Palmberg and N. Prophet, *Quantitative Risk Analyses in the Process Industries: Methodology, Case Studies, and Cost Benefit Analysis.*

[22] S. Garfinkle, "History's Worst Software Bugs," WIRED, 8 November 2005. [Online]. Available: https://www.wired.com/2005/11/historys-worst-software-bugs/. [Accessed 19 April 2018].

[23] Michigan State Police, "Myths and Facts About Seat Belts," State of Michigan, 2018. [Online]. Available: https://www.michigan.gov/msp/0,4643,7-123-1878_1711-13689--,00.html. [Accessed 30 June 2018].

[24] R. Ross, R. Graubart, D. Bodeau and R. McQuaid, "Systems Security Engineering: Cyber Resiliency Considerations for the Engineering of Trustworthy Secure Systems," 2018. [Online]. Available: https://csrc.nist.gov/CSRC/media/Publications/sp/800-160/vol-2/draft/documents/sp800-160-vol2-draft.pdf. [Accessed 5 July 2018].

[25] Symantec, "The Cyber Resilience Blueprint: A New Perspective on Security," 2014. [Online]. Available: https://www.symantec.com/content/en/us/enterprise/white_papers/b-cyber-resilience-blueprint-wp-0814.pdf. [Accessed 18 April 2018].

[26] Federal Emergency Management Agency, "Continuity of Operations Plan Template for Federal Departments and Agencies," April 2013. [Online]. Available: https://www.fema.gov/media-library-data/1386609058805-

b084a7230663249ab1d6da4b6472e691/COOP-Planning-Template.pdf. [Accessed 7 August 2018].

[27] L. van Zoonen, "Privacy concerns in smart cities," *Government Information Quarterly,* vol. 33, no. 3, pp. 472-480, 2016.

[28] R. Billings, "Smart Cities Come with Inherent Privacy Risks, ACLU Says," Portland Press Herald, 26 February 2018. [Online]. Available: http://www.govtech.com/fs/Smart-Cities-Come-with-Inherent-Privacy-Risks-ACLU-Says.html. [Accessed 2 July 2018].

[29] D. Curry, "Privacy a Key Issue Stalling Major Smart City Projects," RTInsights.com, 30 January 2018. [Online]. Available: https://www.rtinsights.com/smart-city-sidewalk-labs-privacy/. [Accessed 2 July 2018].

[30] R. Kitchen, ""The Ethics of Smart Cities and Urban Science"," *Philos Trans A Math Phys Eng Sci,* vol. 374, no. 2083, 28 December 2016.

[31] The Ohio State University Office of Research Office of Responsible Research Practices, "What is the difference between the terms coded, de-identified, and anonymous?," [Online]. Available: http://orrp.osu.edu/knowledge-base/what-is-the-difference-between-the-terms-coded-de-identified-and-anonymous/. [Accessed 21 April 2018].

[32] US Department of Health and Human Services, National Institutes of Health, "Institutional Review Boards and the HIPAA Privacy Rule," [Online]. Available: https://privacyruleandresearch.nih.gov/irbandprivacyrule.asp. [Accessed 16 April 2018].

[33] U.S. Office of Personnel Management, "Cybersecurity Resource Center: Cybersecurity Incidents," U.S. Office of Personnel Management, [Online]. Available: https://www.opm.gov/cybersecurity/cybersecurity-incidents/. [Accessed 2 July 2018].

[34] R. Abrams, "Target to Pay $18.5 Million to 47 States in Security Breach Settlement," The New York Times, 23 May 2017. [Online]. Available: https://www.nytimes.com/2017/05/23/business/target-security-breach-settlement.html. [Accessed 2 July 2018].

[35] Philips Lighting, "Smart city applications are best managed in the cloud," Philips, 2018. [Online]. Available: http://www.lighting.philips.com/main/inspiration/smart-cities/smart-city-trends/cloud#. [Accessed 2 July 2018].

[36] City of Seattle, "Privacy Principals," [Online]. Available: https://www.seattle.gov/Documents/Departments/InformationTechnology/City-of-Seattle-Privacy-Principles-FINAL.pdf. [Accessed 2 July 2018].

[37] City of Kansas City, "Privacy Policy," [Online]. Available: http://kcmo.gov/privacy-policy/. [Accessed 2 July 2018].

[38] D. Paredes and L. Wagner, "Orinda Surveillance Cameras Violate Privacy, Critics Say," 8 August 2016. [Online]. Available: https://www.nbcbayarea.com/news/local/Orinda-Surveillance-Cameras-Violates-Privacy-Critics-Say-389333452.html. [Accessed 23 April 2018].

[39] European Commission, "2018 reform of EU data protection rules," European Commission, 2018. [Online]. Available: https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en. [Accessed 6 August 2018].

[40] T. Fox-Brewster, "Facebook Is Playing Games With Your Privacy And There's Nothing You Can Do About It," Forbes, 29 June 2016. [Online]. Available: https://www.forbes.com/sites/thomasbrewster/2016/06/29/facebook-location-tracking-friend-games/#5632eb235f9c. [Accessed 6 August 2018].

[41] J. Turow, "Google Still Doesn't Care About Your Privacy," Fortune, 28 June 2017. [Online]. Available: http://fortune.com/2017/06/28/gmail-google-account-ads-privacy-concerns-home-settings-policy/. [Accessed 6 August 2018].

[42] W. Oremus, "Americans are losing trust in Facebook — here's why they'll keep using it anyway," Slate, 26 March 2018. [Online]. Available: https://slate.com/technology/2018/03/reuters-ipsos-poll-shows-americans-dont-trust-facebook-will-stay-on-it-anyway.html. [Accessed 6 August 2018].

[43] K. Harkness, Interviewee, *Panel: Challenges and Solutions in Developing a Smart City.* [Interview]. 1 June 2018.

[44] L. Miller, Interviewee, *Financing Smart Cities.* [Interview]. 8 May 2018.

[45] D. Ault, Interviewee, *Panel: City Spotlights: Citizen Life and Governance.* [Interview]. 29 March 2018.

[46] National Institute of Standards and Technology, "Internet-of-things Enabled Smart City Framework: A Consensus Framework for Smart City Architectures [Draft Release v2]," 2018. [Online]. Available: https://s3.amazonaws.com/nist-sgcps/smartcityframework/files/ies-city_framework/IES-CityFrameworkdraft_20180207.pdf. [Accessed 18 April 2018].

[47] A. Sardella, Interviewee, *Panel: Using the Science of Wellbeing to Guide Your Smart City.* [Interview]. 9 May 2018.

[48] M. J. Rowley, "New tech tools to track employee happiness," Cisco Systems, 7 August 2017. [Online]. Available: https://newsroom.cisco.com/feature-content?type=webcontent&articleId=1854540. [Accessed 1 July 2018].

[49] J. Bickers, "Self-service City Hall: The Benefits of E-gov Kiosks," [Online]. Available: https://www.kioskmarketplace.com/articles/self-service-city-hall-the-benefits-of-e-gov-kiosks-3/. [Accessed 25 April 2018].

[50] American Cancer Society, "Cellular Phone Towers," [Online]. Available: https://www.cancer.org/cancer/cancer-causes/radiation-exposure/cellular-phone-towers.html. [Accessed 23 April 2018].

[51] World Health Organization (WHO), "Electromagnetic Fields and Public Health: Base Stations and Wireless Technologies," May 2006. [Online]. Available: http://www.who.int/peh-emf/publications/facts/fs304/en/. [Accessed 23 April 2018].

[52] J. Rusk Chief of Civic Wellbeing for the City of Santa Monica, Interviewee, *Panel: Using the Science of Wellbeing to Guide Your Smart City.* [Interview]. 9 May 2018.

[53] Office of the Mayor News Release, "Mayor Mark Farrell Announces Plans for Universal Fiber Network that Mandates Net Neutrality and Privacy Protections," 31 January 2018. [Online]. Available: https://sfmayor.org/article/mayor-mark-farrell-announces-plans-universal-fiber-network-mandates-net-neutrality-and. [Accessed 20 July 2018].

[54] S. Barocas and A. Selbst, "Big Data's Disparate Impact," *California Law Review,* vol. 104, no. 3, p. 671, June 2016.

[55] C. Conley, "Making Smart Decisions about Smart Cities," November 2017. [Online]. Available: https://www.aclunc.org/sites/default/files/20171115-Making_Smart_Decisions_About_Smart_Cities.pdf. [Accessed 24 April 2018].

[56] D. Ingold and S. Soper, "Amazon Doesn't Consider the Race of Its Customers. Should It?," *Bloomberg,* 21 April 2016.

[57] B. Bennett, Interviewee, *Global Cities Team Challenge Kickoff -.* [Interview]. 7 February 2018.

[58] Chi Hack Night, "Chicago's weekly event to build, share & learn about civic tech," Chi Hack Night, 2018. [Online]. Available: https://chihacknight.org/. [Accessed 2 July 2018].

[59] Hack Portland, "Hack Portland," Meetup.com, 2018. [Online]. Available: https://www.meetup.com/Hack-Portland/. [Accessed 2 July 2018].

[60] S. Ladin-Sienne, "Measuring Community Vitality Through the 'Great Streets' of Los Angeles," Ash Center at Harvard Kennedy School, 2017.

[61] L. Miller, Interviewee, *Financing Smart Cities.* [Interview]. 8 May 2018.