Making Smart Decisions about Smart Cities

ACLU
Northern
California

“Smart city” products are often touted as cost-efficient solutions to providing government services and addressing societal challenges, but in reality these technologies are no more intelligent than the data and algorithms behind them. Like other tools, their effectiveness depends on when, where and how they are used. And they often carry a wide range of risks or costs that may not be obvious on the surface, including real consequences for your residents and visitors.

In some cases, technological solutions to urban issues can reflect or even exacerbate racial or economic inequality rather than allocate city resources and services fairly. Technologies that collect information about residents and visitors can generate and expose deeply personal information or even be repurposed to function as surveillance tools. And smart initiatives built around public-private partnerships or otherwise reliant on third parties can lead to negative outcomes for your city if the relationships are not clearly defined or the incentives of everyone involved are insufficiently aligned.

The best approach to smart city technology is neither to blindly deploy it nor to reject it outright. Instead, your city needs to determine whether a technology-centric initiative or project is really the “smart” thing to do. This guide can help you do just that. It provides a step-by-step approach to evaluating data-driven urban technologies and the risks and costs they may entail, drafting policies and contracts that can mitigate potential harms, and collaborating with the community to decide whether and how to proceed with smart initiatives at all. It features real-life case studies that help highlight the benefits of making informed decisions about technology rather than trusting the technology to be smart for you.

Finally, note that this report focuses on technologies that are not designed specifically for law enforcement use. Please see our companion report, *Making Smart Decisions about Surveillance*, for information about those technologies.

We hope you find this report helpful in making truly smart decisions for your city.

Author: Chris Conley, ACLU of Northern California

Contributors: Nicole Ozer & Matt Cagle

Design & Layout: Gigi Harney & Jennifer Zhao

This publication was underwritten with support from the ACLU Foundation
and the ACLU's generous members and donors.

PUBLISHED BY THE ACLU OF CALIFORNIA
Technology & Civil Liberties Project
November 2017

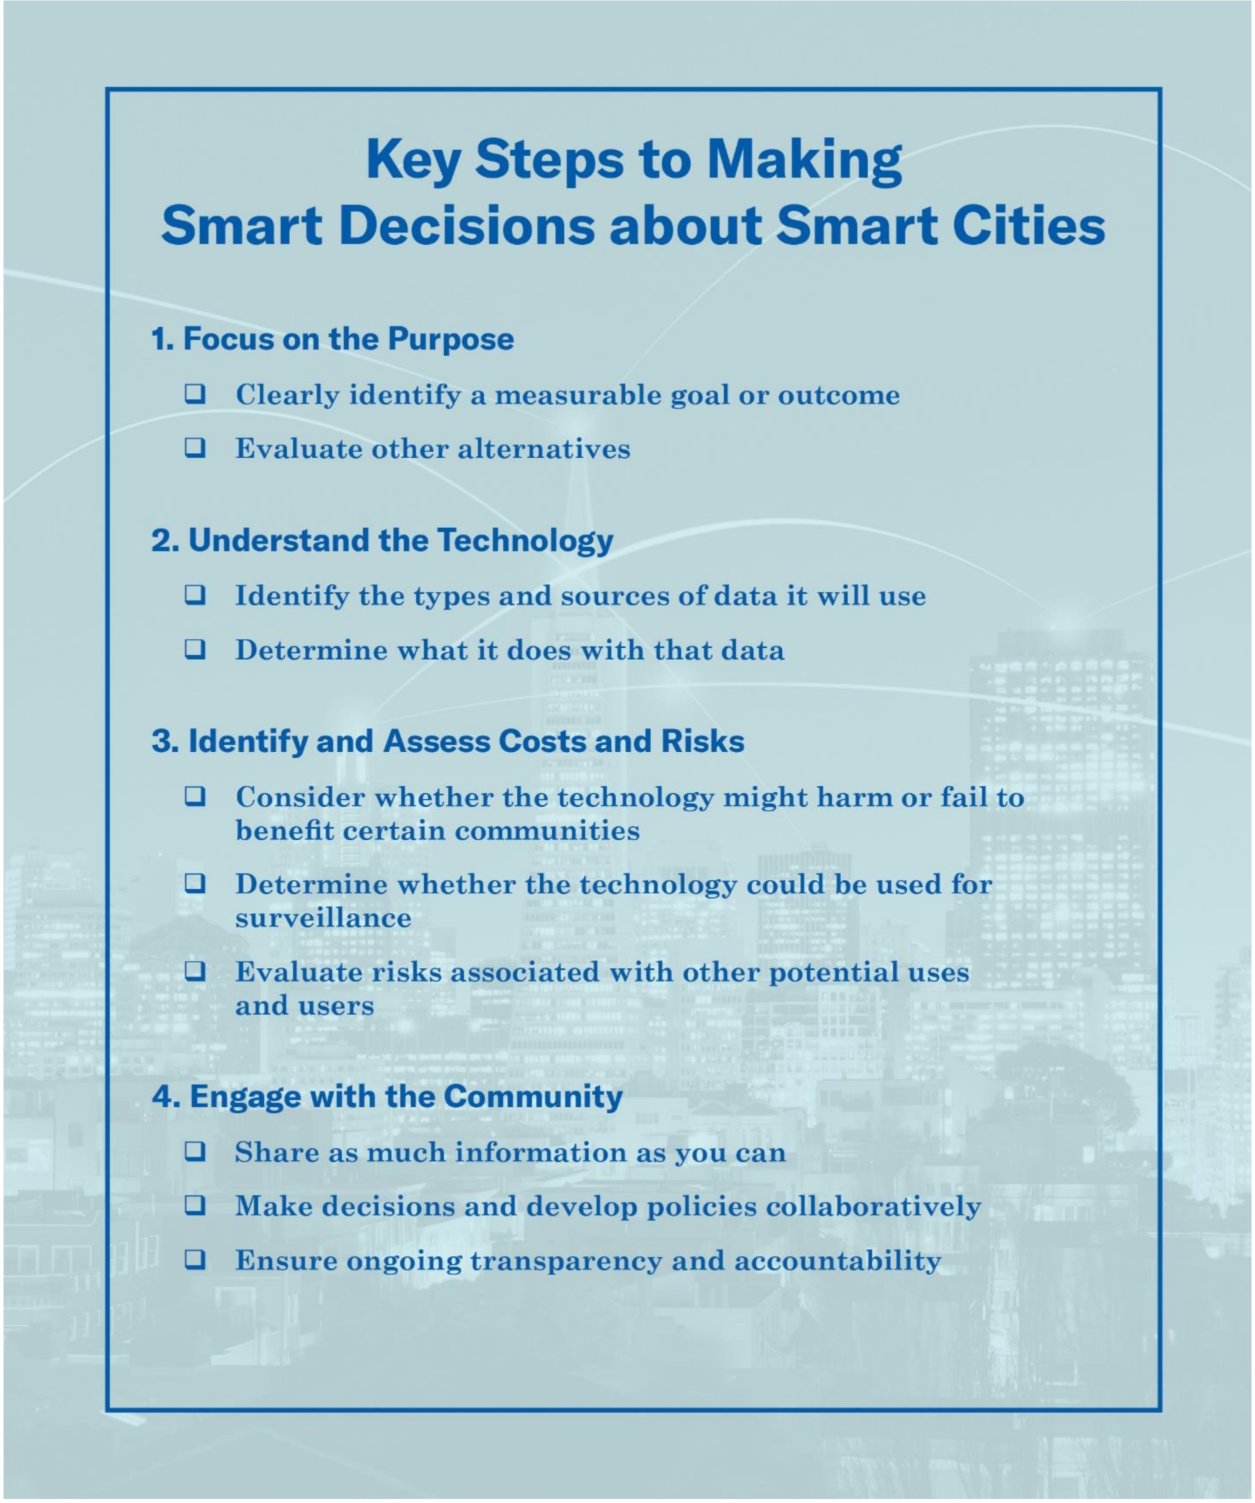

Key Steps to Making Smart Decisions about Smart Cities

1. Focus on the Purpose

- ❑ Clearly identify a measurable goal or outcome
- ❑ Evaluate other alternatives

2. Understand the Technology

- ❑ Identify the types and sources of data it will use
- ❑ Determine what it does with that data

3. Identify and Assess Costs and Risks

- ❑ Consider whether the technology might harm or fail to benefit certain communities
- ❑ Determine whether the technology could be used for surveillance
- ❑ Evaluate risks associated with other potential uses and users

4. Engage with the Community

- ❑ Share as much information as you can
- ❑ Make decisions and develop policies collaboratively
- ❑ Ensure ongoing transparency and accountability

Making Smart Decisions about Smart Cities

1. Focus on the Purpose

The first question to ask about any smart city product or initiative is perhaps the most basic: why? What is the goal you are trying to achieve or problem you want to solve? And will the technology you are considering actually accomplish that goal or solve that problem?

➢ *Clearly identify a measurable goal or outcome*

Smart city initiatives should be focused on improving the lives of the people in your community. There are many ways technology might help you do that: make your city more efficient by leveraging data, attracting investment by building a reputation as a technology center or tackling environmental issues like air quality. But whatever your goal may be, it needs to be clearly defined.

Beyond a clear purpose, it is also important to establish concrete markers of success for your project. If you are aiming to reduce citywide electricity consumption or transit delays, you can establish a specific percentage target and determine whether or not you hit it. Vaguer "purposes" without a ready measuring stick make it harder to evaluate your initiative's actual success, and much easier to lose sight of your true objective as a result.

Even if you are soliciting ideas for potential initiatives, you can still prioritize ideas and proposals that address specific problems and have measurable outcomes. Anyone proposing an initiative should be able to identify the actual ways it will benefit your city — and provide evidence to back up these claims. Models like Boston's Smart City Request for Information can help you evaluate the tangible benefits of potential smart initiatives.

On the other hand, just "test driving" new technologies without a specific purpose can be risky and problematic. At best, a "pilot program" that does not address a clearly-defined problem may result in nothing more than "a glossy presentation, and a collective shrug" after drawing attention and resources from other options. At worst, the lack of a defined purpose can lead to real harms that were not identified because the details of the project were not fully developed. This applies both to brand-new initiatives and to unanticipated expansions of existing programs into new areas.

➢ *Evaluate other alternatives*

Technology can make some intractable problems easier to address — but it can also distract attention from underlying issues and divert funds from other effective strategies. As a result, smart city proposals should be evaluated on the same basis as other proposals: are they likely to actually solve the problem at hand? At what cost? Otherwise, they risk drawing attention and resources from alternatives that might be more effective.

2. Understand the Technology

In order to identify and address potential costs and risks, you need to know not only what a smart technology is intended to do but how it actually works. Most importantly, you need to understand the data that drives the technology: what it is, where it comes from, how accurate it is and how it is used. Only then can you really understand the impact it might have on your community and identify potential harms and solutions.

➢ *Identify the types and sources of data it uses*

The data that smart technologies use can come in many forms and from many sources, even for the same purpose. Data can be generated from government-operated sensors, culled from historical records, voluntarily disclosed by residents or visitors, or obtained through a contract or agreement with a third party. The data can be identifiable to a specific person or (at least nominally) anonymous — or it may not be linked to specific people at all. It might be gathered very transparently, like the increasing number of toll bridges and roads that clearly indicate the use of license plate readers, or obtained without any meaningful notice to the affected individuals, like cameras unobtrusively mounted on light poles. And many technologies incorporate multiple types or sources of data.

Sensor Data

Many technologies generate their own data from sensors, often widely dispersed throughout the city. These sensors may capture or generate various types of information, including:

- Video
- Audio
- Location, e.g. GPS devices that track the location of busses or license plate readers that assign geospatial coordinates to captured images
- Environmental, e.g., temperature, lighting, air quality, etc.
- Consumption of water, electricity, or other utilities, i.e. smart meters
- Pressure, used to determine the presence or passage of a vehicle or pedestrian
- And many more

Even within a category, sensors can vary widely. Video sensors, for example, range from simple motion detectors to sophisticated cameras capable of capturing, analyzing and storing high-definition images. Both the value and the potential risks may depend on the quality as well as the amount of information generated.

San Diego's "Shotspotter" Captures More than It Aims For

San Diego's initial decision to install smart street lights invited little controversy, but that changed when the city decided to attach "Shotspotter" microphones to the lights. Activists pointed out that while the product is billed as a gunshot detector, it apparently is capable of recording any kind of sound, including conversations. Because the decision to install Shotspotter was made with "no public input and only limited input from city council, the city has faced demands for public hearings and greater transparency going forward.

Other City-Collected Data

Cities also gather data from more traditional interactions with community members and the area: tax assessments, employee-reported damage to recreational spaces or equipment, and so on. This can include historical records that may not reflect the current demographic makeup of the city or that may be biased by historical segregation or structural racism. Relying on this data without identifying and correcting for its weaknesses may lead to inequitable treatment of communities of color and other marginalized groups.

User-Submitted Data

Some technologies rely on users to submit information, either manually (e.g. by filling out an online form or sending a text) or automatically (e.g. by downloading a mobile app that uses a smartphone's sensors to detect potholes while driving). This data often comes at low cost to the city and can foster feelings of community and public-private collaboration. However, the data may over- or under-represent certain communities based on historical relationships with the government, economic inequality, or other factors that can create bias or outright error in the data set. It can also raise privacy concerns if, for example, the mobile app collects data even when not in use.

Metadata

Many devices and sensors attach all kinds of additional information, called metadata, to any content they collect or create. For example, a captured image may be tagged with the time and location where it was taken, or an online form could record information about the submitter's device and Internet connection. This information can be highly identifiable and sensitive, both independently and in the aggregate, even when the "content" itself is not.

Third Party Data

Cities are increasingly obtaining data through third parties, either via formal public-private partnerships that install and manage sensors or simply by purchasing or otherwise acquiring existing data about the city and its communities. The risk of bias or inequity may depend not only on what types of data it at issue but how it is collected, used and stored at its source and whether it is provided to the city in raw form or subject to the private party's own analysis or editing. Your city needs to understand the context in which third party data is collected and acquired in order to assess its value and risks.

➢ *Determine what it does with that data*

In addition to understanding what data a technology uses, it is important to determine how that data is used, stored and transmitted or otherwise exposed. Relevant questions include:

- How accurate is the data? Is that accuracy sufficient for its purpose? Is it more accurate than necessary? What happens when data is inaccurate?
- Does the technology do anything with the data it collects or receives? What inferences or analysis does it generate?
- How long does the technology store any data it collects or generates?
- Is the data remotely accessible automatically transmitted in any way? If so, how and to whom?
- Does the technology use the data for any other reason other than the assigned purpose?

3. Identify and Assess Costs and Risks

Unfortunately, technology often comes with hidden costs and risks, and that is no less true for smart city technologies. Because they rely on data, they can present significant threats to individual rights, including not only invasions of privacy but also the potential to reflect or even exacerbate inequalities arising from racism and other societal biases. Thoroughly evaluating the potential costs and how or whether they can be avoided or mitigated is an essential step in determining whether a given smart technology is the right choice for your community.

➢ *Consider whether the technology might harm or benefit certain communities*

Any data-driven system is vulnerable to biases. The data that it collects may reflect historical or present inequality, including those driven by discrimination. Its placement and use can reduce or conversely worsen these injustices. Concerns about bias cannot be simply brushed aside with the assumption that computers are inherently neutral or that your "smart" system will mitigate them. Instead, as New York City noted, it is important to make equity "an explicit guiding principle in [your] work."

Data Biases

Data-driven technologies are only as good as the data that goes into them. If your data is biased, it can lead your system to erroneous and inequitable conclusions.

The types of bias depend largely on the particular data your system will use:

- **Sensor data**: Are the sensors equitably deployed? Wil they work equally well everywhere? Are there differences between areas or communities that need to be taken into account when interpreting the data?

- **Participatory data**: Does everyone participate equally in generating data, or are there barriers to participation for some groups (particularly those with reduced access to technology, reluctance to communicate with the government or lowered expectations due to historical mistreatment) that causes the data to be biased? How can you avoid favoring "squeaky wheels" or users from groups with greater access to power? Are participants adding their own biases to the mix?

BART Watch App Allegedly "Promotes Racial Profiling"

The Bay Area Rapid Transit agency released its "BART Watch" app with the goal of empowering riders to keep BART safe by reporting suspicious activity on the trains. But an investigation showed that the app was disproportionately used to target Black and homeless riders, whether or not the target was actually committing a crime. The app also sparked a lawsuit for allegedly collecting information about its users, including location information, even when the app was not being used.

- **Historical data**: What historical biases are present in the data? Are some communities over- or under-represented in the data set, or is data for some communities less accurate? How can these be accounted for to ensure that the biases are not perpetuated?

Algorithmic Biases and Equitable Outcomes

Smart city technology can also help or hinder the cause of social justice in the way that it allocates scarce resources or provides equitable or inequitable services to all members of a community. These systems can perpetuate or even strengthen existing inequalities within the community rather than addressing them. For example, congestion pricing systems and similar systems that increase prices when demand exceeds supply can disproportionately impact lower-income residents and neighborhoods, which can lead to "pricing out the poor" for whom a marginal increase in costs may be unaffordable. Even attempts to specifically serve marginalized groups can fail if those groups lack the resources to take advantage.

Boston Pothole Detection App Forced to Patch Holes in Data

In 2012, Boston announced "Street Bump," a smartphone app designed to automatically detect and report potholes as drivers passed over them. The app was intended to help direct the city's resources for both short-term fixes and long-term projects. However, researchers pointed out that the app's data set largely excluded "[p]eople in lower income groups and the elderly who are less likely to have smartphones, which "result[ed] in an unequal allocation of funds." The city responded by "work[ing] with a range of academics to take into account issues of equitable access and digital divides."

In addition, if technology is going to truly improve your community, it must be deployed in a way that ensures that lower-income individuals, communities of color, and other disadvantaged residents reap their share of the benefits. Among other things, this means that technologies that leverage existing infrastructure or resources, ranging from public transit to available high-speed Internet, cannot be trusted to automatically deliver services equally. Your community instead needs to identify existing inequities and incorporate those into the plan for smart technologies.

Washington Bike-Share Program Struggles to Attract Low-Income and Minority Riders

Washington, D.C.'s "Capital Bikeshare," like many other such programs, was intended to provide all of its residents access to bikes – but has found that the vast majority of its riders are white and wealthy. Even though the program placed bike stations in a wide range of locations, low-income communities and people of color have been deterred from participating due to a combination of factors, including fear of liability and a higher risk of traffic accidents, harassment and crime. Even dedicated efforts to "recruit minorities and low-income residents" have thus far failed to ensure that the service provides value to the entire community.

➢ *Determine whether the technology could be used for surveillance*

Any smart technology that collects data about residents or visitors presents a risk that the technology will be used or perceived as surveillance. The risk of surveillance can be most acute with technologies that:

- generate data about individuals, including video or audio recordings, especially if that data is easily linked to a specific person;
- collect data that contains or can expose sensitive information, including information related to location or associations; or
- lack access controls, auditing and other safeguards to ensure the data is not used inappropriately.

Even a small amount of information can be profoundly revealing, identifying an individual as a participant in a political rally or a patient at a mental health clinic. But the risks and costs of surveillance are both heightened as the number of sensors and amount of data increase. Aggregating large data sets can permit both "connecting the dots" to tie a seemingly-anonymous bit of information to something that is associated with a specific person and data mining to drill into seemingly innocuous information to derive "hidden" inferences that can reveal sensitive information such as an individual's sexual orientation, relationship status or political activism.

For more information, please see our guide on surveillance technology, *Making Smart Decisions about Surveillance*.

Surveillance poses risks both to individual privacy and to civil rights and equitable treatment of all community members and visitors. Before adopting any technology that collects or produces data, cities need to identify those risks and take steps to mitigate them. This should include evaluating alternatives that might accomplish the same aims with less risk by reducing data collection in the first place, as well as establishing and publishing a use policy that will clearly spell out the protections for any data that is collected.

Identifiable Data

One simply way to avoid many (though certainly not all) risks associated with sensor data is to make sure the data cannot be linked to an individual person, either immediately (because the data is tagged with the person's name, email, etc.) or indirectly (e.g. by capturing a high-quality image of the person that could be recognized or identified by facial recognition). In many cases, identifiable or overly precise data does not forward the purpose of the technology at all. Simpler approaches can often accomplish the same purpose with both a lower financial cost and a reduced risk of harm.

Palo Alto Systems Generate Results, Not Data

In order to predict and plan for congestion and traffic, the city of Palo Alto turned to technology. To protect privacy at the same time, the city chose systems including parking sensors and low-resolution cameras that are sufficient to detect the presence of a car, bike or pedestrian but do not generate enough information to identify anyone. In addition, the city deletes all raw data as soon as it is processed. In doing so, the system minimizes any potential privacy risks that might be associated with higher-resolution cameras and images that are retained or shared.

Data Retention and Deletion

Another effective way to reduce the risk of surveillance or misuse is to eliminate data once its value is gone. Automated systems that purge old data once it has been processed or is no longer needed can eliminate the possibility of future misuse. And even data that continues to have value can often be aggregated or otherwise altered after a set period to make it less identifiable or sensitive.

➢ *Evaluate risks associated with other potential uses and users*

Many cities look at smart technologies as a way to act collaboratively, both with specific vendors and industry partners and with their community as a whole. While this can be both admirable and effective, it does raise specific concerns that need to be taken into account. A thorough plan for smart technologies needs to specify legitimate users and uses (rather than just users) and ensure that other access and use is effectively prevented.

Inter-Government Sharing

It can be tempting to multiply the benefits of the data you collect by sharing it with other government agencies. But doing so without foresight can create risks that the data may be misused or disclosed further, or that communities will simply refuse to engage with the smart technology. These risks are heightened when law enforcement is involved, particularly in communities with a history of tense relationships with the police. You need a clear (and public) policy spelling out whether and how data from any given technology is accessible to other agencies for other purposes.

Corporate Partners

Public-private smart city partnerships can raise significant concerns about aligned incentives. Private entities are likely to be driven to maximize their own profit rather than provide equitable service to all of your residents and visitors. This can impact not only how companies use data provided by the city but what data they provide themselves, which can create additional risks of bias or error. And, of course, your city should be sure that you will not be left with nothing if your partnership does not work out in the long run.

NYPD Loses Access to Data After Parting Ways with Palantir

After using Palantir's software to compile and analyze information about criminal and civil matters for several years, the NY Police Department decided to end its relationship with the company – and was surprised to learn that it would lose a significant amount of the information it had generated at the same time. The Department asked Palantir to provide both a data file with all of the generated analysis and a "translation key" so it could use that analysis in a new system. Palantir agreed to create the file but refused to provide a way to translate it, rendering the data largely useless. The two parties are still fighting it out.

Any partnership agreement needs to ensure that smart city data is used for the benefit of all. This requires understanding and taking into account

- data ownership and access;
- your partner's business model and long-term viability;
- any applicable legal and regulatory framework; and
- plans for auditing and evaluation.

as well as any other factors that might affect the partnership.

Finally, it is important to ensure that "trade secrets" or other demands for confidentiality do not supersede the need to thoroughly understand your partner's business model and data handling practices. This is especially true if your city is outsourcing the provision of services to third parties. Any partnership should ensure that the city is positioned to further and protect the interest of its residents and visitors, which is simply impossible without full disclosure of potential conflicts of interests or functions of the technology.

Toronto Urged Not to "Lose Sight of Personal Privacy"

Toronto has agreed to partner with Sidewalk Labs, a subsidiary of Google's parent company Alphabet, to rebuild a section of the city "from the Internet up." But this proposal has raised alarms from residents about the risks of a private party collecting vast amounts of information and handling other traditional government roles without electoral accountability. Activists and the local press have pushed for transparency, strong data use policies and other safeguards to ensure the city doesn't "lose sight of personal privacy" or allow Sidewalk Labs to put its own interests first.

Data Security

Even data that is not intended to be preserved or disclosed in any fashion may be subject to hacking or otherwise compromised. Cities need to ensure that any kind of data is properly protected, something that all too frequently is not the case for commercial "Internet of Things" devices that lack even basic security mechanisms.

Redland Camera Networks Hacked Due to Security Flaws

The city of Redland, CA deployed a network of "smart" cameras intended to allow the city to monitor traffic and other activities. However, a security researcher quickly demonstrated that the cameras lacked basic security precautions. As a result, the cameras and the footage they collected were readily available to hackers.

Open Data

Many smart city programs eagerly share information with third parties or the public as a form of transparency. While this can have benefits, it also poses risks, especially if the exposed data consists of individual records that have any chance of being linked to specific persons. Make sure you take steps to protect the privacy of individuals in any data set you make public.

Alternate Uses

Evaluating the potential impact of smart city programs means looking beyond the primary purpose and expected use case to explore how else the technology and related data could be used. A parking technology that relies on pressure sensors generates data that may have few uses outside of its primary purpose. But a parking technology that instead relies on cameras could easily be used for many other purposes. Potentially harmful uses might be foreclosed by technical or policy means, but first they must be identified.

San Jose Backs Off Proposal to Add Cameras to Existing "Smart Lights"

The smart street lights installed by San Jose were not designed exclusively as lights; instead, they were built with additional "ports" where new components could be added. However, both residents and city officials protested when the city proposed adding surveillance cameras to the street lights. The city has since decided to limit the cameras to recording "vehicle traffic" until a privacy and use policy is developed.

4. Engage with the Community

The previous sections spelled out questions you need to answer to determine whether any given smart initiative is the right answer for your city. But you do not need to answer these questions alone, nor should you. The best way to answer all of these questions is to bring the wisdom and expertise of the entire community to bear, both when you make the decision whether to use smart tech in the first place and on an ongoing basis through measures designed to ensure transparency and oversight. Doing so can help you ensure that your smart tech accomplishes its intended purpose and benefits your entire community.

➢ *Share as much information as you can*

Your community members need information both to understand how smart technology might affect their own lives and to reflect that understanding back to you to help you make decisions about the initiative. While not every community member cares about every last detail, providing as much information as possible helps experts and concerned residents alike participate in the process.

One of the most important pieces of information to share is your assessment of the technology, including an evaluation of any risks of bias or surveillance. Explaining to community members how a technology can serve them can help them embrace smart initiatives, while having a sincere discussion about costs and risk can help you not only prioritize and weigh issues but also identify potential solutions.

Bounder Smart Grid Struggles After "Leav[ing] Citizens in the Dark"

Boulder, CO was one of the first cities in the U.S. to install smart meters through a partnership with private utility Xcel Energy. However, residents were widely confused about the project, with many uncertain whether they had a smart meter at all. As the city evaluated its options, it recognized the need to dramatically improve communications with electric customers, making "consumer information, consumer choice, and consumer communication … a huge part of what we would do."

In addition, your community will better understand the system you are proposing if you give them as much information as possible. Technical specifications, vendor contracts, proposed use policies and other details can help your community truly understand and engage with the proposal.

➢ *Make decisions and develop policies collaboratively*

Once you have given your community members the information they need to make a decision (and the time they need to process it), you also need to provide an opportunity to participate in the decision-making process. There are many different ways that you might reach out to community members for feedback. These include

- **holding public hearings** to discuss the smart tech, or
- **forming citizen working groups** to evaluate new proposals.

In addition, you may want to reach out directly to community leaders or experts to participate in these events or simply spread the word and generate feedback.

Oakland Forms Privacy Advisory Commission to Solicit Feedback on City Technology

After triggering outrage from community activists over its use of surveillance technology, Oakland changed its approach and began to formally engage with informed citizens. Its recently-formed Privacy Advisory Commission, which includes both academics and advocates, is charged with “provid[ing] advice to the City of Oakland on best practices to protect Oaklanders' privacy rights in connection with the City's purchase and use of surveillance equipment and other technology that collects or stores our data.”

➢ *Ensure ongoing transparency and accountability*

Even if your community has already deployed smart tech, you still have an obligation to make sure it is being used appropriately and effectively. This includes ongoing conversations about whether the smart tech is actually fulfilling its primary purpose as well as transparency and oversight to identify and address any harms or abuse that have been identified.

Public Policy

A public use policy — crafted in consultation with your community —can make it clear how any technology or data associated with a smart initiative is to be used: who can access the system or its data (including how and whether the public can gain access), how it will be protected, when it should be destroyed, and many other details. A robust policy also can clarify the intended purpose of the initiative and expectations for anyone who uses it. This policy should, in most cases, include:

- **the specific purpose(s)** that the project is intended to advance;
- **the technology** that will be used in the project;
- **the authorized user(s)** of the technology or any data it collects;
- **the uses that are authorized** and the uses that are prohibited;
- **required safeguards** to protect the technology or data from misuse, including encryption or access controls;
- **the data retention period**, the reason for selecting that time period and the process by which data is deleted after the retention period expires;
- **data disclosure rules** that describe if and how any other government agencies or non-government entities, including law enforcement and members of the public, may use or access the technology or any data it collects; and
- **required auditing or other mechanisms** designed to measure the effectiveness of the initiative and compliance with the policy.

Ongoing Transparency

Your obligation to keep your city's residents informed and engaged does not end once you launch a smart initiative. To maximize the benefits and minimize the risks, you need to keep a running dialog with community members. This means being transparent about the effectiveness and newly-discovered costs, benefits or risks of smart initiatives, as well as any violations of your use policy or other relevant information, so that the community can make informed decisions about its continued use. It also means restarting the process if necessary when a new opportunity or proposed change calls into question the results of your initial collaborative decision.

Conclusion

Smart city technologies can benefit your community in various ways — but they also present the risk of very real harms. Technologies that allocate resources or drive decision-making can perpetuate or even exacerbate inequality and the effects of discrimination, and technologies that collect data can be used for surveillance or other purposes that may not align with the community's values or respect individual rights. Avoiding these risks requires identify your goals, understanding your technology, identifying costs and risks, and collaborating with your community. Taking these key steps will help you make smart decisions about smart cities.