Digital Innovation

Introduction

The Vision

Catalyze digital innovations that help **tackle urban challenges** and **establish a new standard** for the **responsible collection** and use of data in cities.

The ability to create the conditions for digital innovation is at the heart of Sidewalk Labs' vision for the city of the future. Digital innovation is the basis for many of the core planning initiatives that Sidewalk Labs has proposed throughout this Master Innovation and Development Plan to improve mobility, affordability, sustainability, and economic opportunity. It is also essential for catalyzing an ecosystem of new services and solutions by individuals, Canadian companies, local Toronto entrepreneurs, and other third parties from around the world.

That ecosystem is thriving in Toronto. Today, digital innovation is powering the region, from the cybersecurity and software startups in the Toronto-Waterloo corridor to local institutions like MaRS Discovery District, Communitech, the Vector Institute for Artificial Intelligence, and Civic Tech Toronto. Together these forces are driving Toronto's future: in 2015, the digital economy generated $117 billion nationwide,[1] supported 4,000 new Toronto businesses,[2] and provided 400,000 jobs for the city.[3]

But digital innovation raises a number of challenges that cities like Toronto are just starting to address. These include making sure basic digital infrastructure is affordable and open to everyone, making sure data is standardized and publicly accessible, and making sure there is a transparent process for protecting privacy and the good of the city.

These challenges are especially complicated for "urban data," which Sidewalk Labs defines as information gathered in the city's physical environment, including the public realm, publicly accessible spaces, and even some private buildings. While Canada has a strong foundation of privacy laws around personal information, and recognizes privacy as a fundamental human right, urban data creates a new set of questions that have surfaced during the Sidewalk Toronto public consultation process.

How can both cities and companies use data in a responsible way in the digital age?

How should the collection of data in public spaces evolve to match the speed of today's digital devices and the rapid development of artificial intelligence?

How can cities continue to engage in a meaningful public dialogue that addresses valid concerns about the impact on personal privacy, or about using urban data for the greater good?

Toronto and Ontario have taken some important initial strides to advance the conversation around data governance principles, including calling for public consultations to discuss how the digital economy can support business while protecting privacy. But while every city faces new barriers in the digital age, no place has yet adopted a comprehensive approach to address these challenges and create the conditions for digital innovation to flourish responsibly. The Sidewalk Toronto project presents a unique opportunity to do just that, and Sidewalk Labs proposes a holistic approach to digital innovation with four core components.

The innovation plan.
First, Sidewalk Labs proposes to establish open digital infrastructure that provides a shared foundation for using urban data to improve quality of life. This core infrastructure would be anchored by ubiquitous, affordable internet connectivity within the IDEA District, consistent with Waterfront Toronto's aspirations for closing the digital divide. It would also include physical mounts that can significantly reduce the cost of launching new digital innovations and help ensure that cities do not get locked into using proprietary solutions.

Second, Sidewalk Labs proposes to outline clear standards that make data publicly accessible, secure, and resilient. Today's urban data tends to be scattered across many owners, outdated, or

Key Term
Urban data
refers to information gathered in the city's public realm, its publicly accessible spaces, and even some private buildings.

stored in messy file formats, making it difficult for the community to use as a foundation for new ideas. Clear standards would make (properly protected) urban data accessible to researchers and the community in real time, and make it easy for third parties to build new services or competitive alternatives to existing ones.

Third, Sidewalk Labs proposes a trusted process for responsible data use that would apply to all parties (including Sidewalk Labs). This process would be anchored by a Responsible Data Use (RDU) Assessment — an in-depth review that is triggered by any proposal to collect or use urban data — and guided by a set of RDU Guidelines that incorporates globally recognized Privacy by Design principles. The process, including approvals, would be overseen by an independent Urban Data Trust created to be a steward of urban data and the public interest without stifling innovation.

Finally, Sidewalk Labs proposes to launch a minimal set of digital services that would catalyze this ecosystem of urban innovation. These services and applications — all of which would be open to competition and subject to the proposed responsible data use process — represent innovations currently not being pursued by the market but that remain essential to achieving Waterfront Toronto's quality-of-life objectives. Furthermore, the (properly protected) urban data generated by these launch services would be made publicly accessible (on a non-discriminatory basis), enabling companies, community members, and other third parties to use it as a foundation to build new tools.

Benefits of implementing the vision

Pilot new digital services that improve quality of life

Build fast, affordable digital infrastructure for residents and workers

Help make Toronto a global urban innovation hub

Establish a new standard for responsible data use

The impact.
At the small neighbourhood scale of Quayside, Sidewalk Labs' proposed approach would help pilot a range of services that improve daily life for neighbourhood residents, workers, and visitors across its core innovation pillars. These include a mobility management system that could use travel data to improve congestion and safety; an outdoor-comfort system that could use weather data to make the public realm more usable; a building-code system that could use structural and noise data to support a mix of residential and commercial uses; and energy management tools that could use data on energy demand and pricing to reduce peak-hour use, and thus greenhouse gas emissions.

Applied at the full scale of the IDEA District, the conditions of urban data, digital infrastructure, and core services would catalyze a new ecosystem for urban innovation, filled with technological advances by others that make urban challenges easier to tackle. That might include anything from a next-generation bike-share service, to small business tools that help retailers launch a successful pop-up, to civic tools that help families find an affordable home, to improved building designs that reduce energy use, to new apps that bring people together outdoors. The list would be bound only by imagination.

Sidewalk Labs' proposals for digital innovation would make it possible for the IDEA District to achieve key quality-of-life objectives. It would also serve as the cornerstone of a new global hub for urban innovation, estimated by Sidewalk Labs to generate $14.2 billion in annual economic activity and give rise to 93,000 total jobs, including nearly 10,500 jobs focused on urban innovation — attracting entrepreneurs from all over to the IDEA District.[4]

Above all, Sidewalk Labs' approach aims to demonstrate to Toronto, Ontario, Canada, and the rest of the world that cities do not need to sacrifice their values of inclusion and privacy for economic opportunity in the digital age.
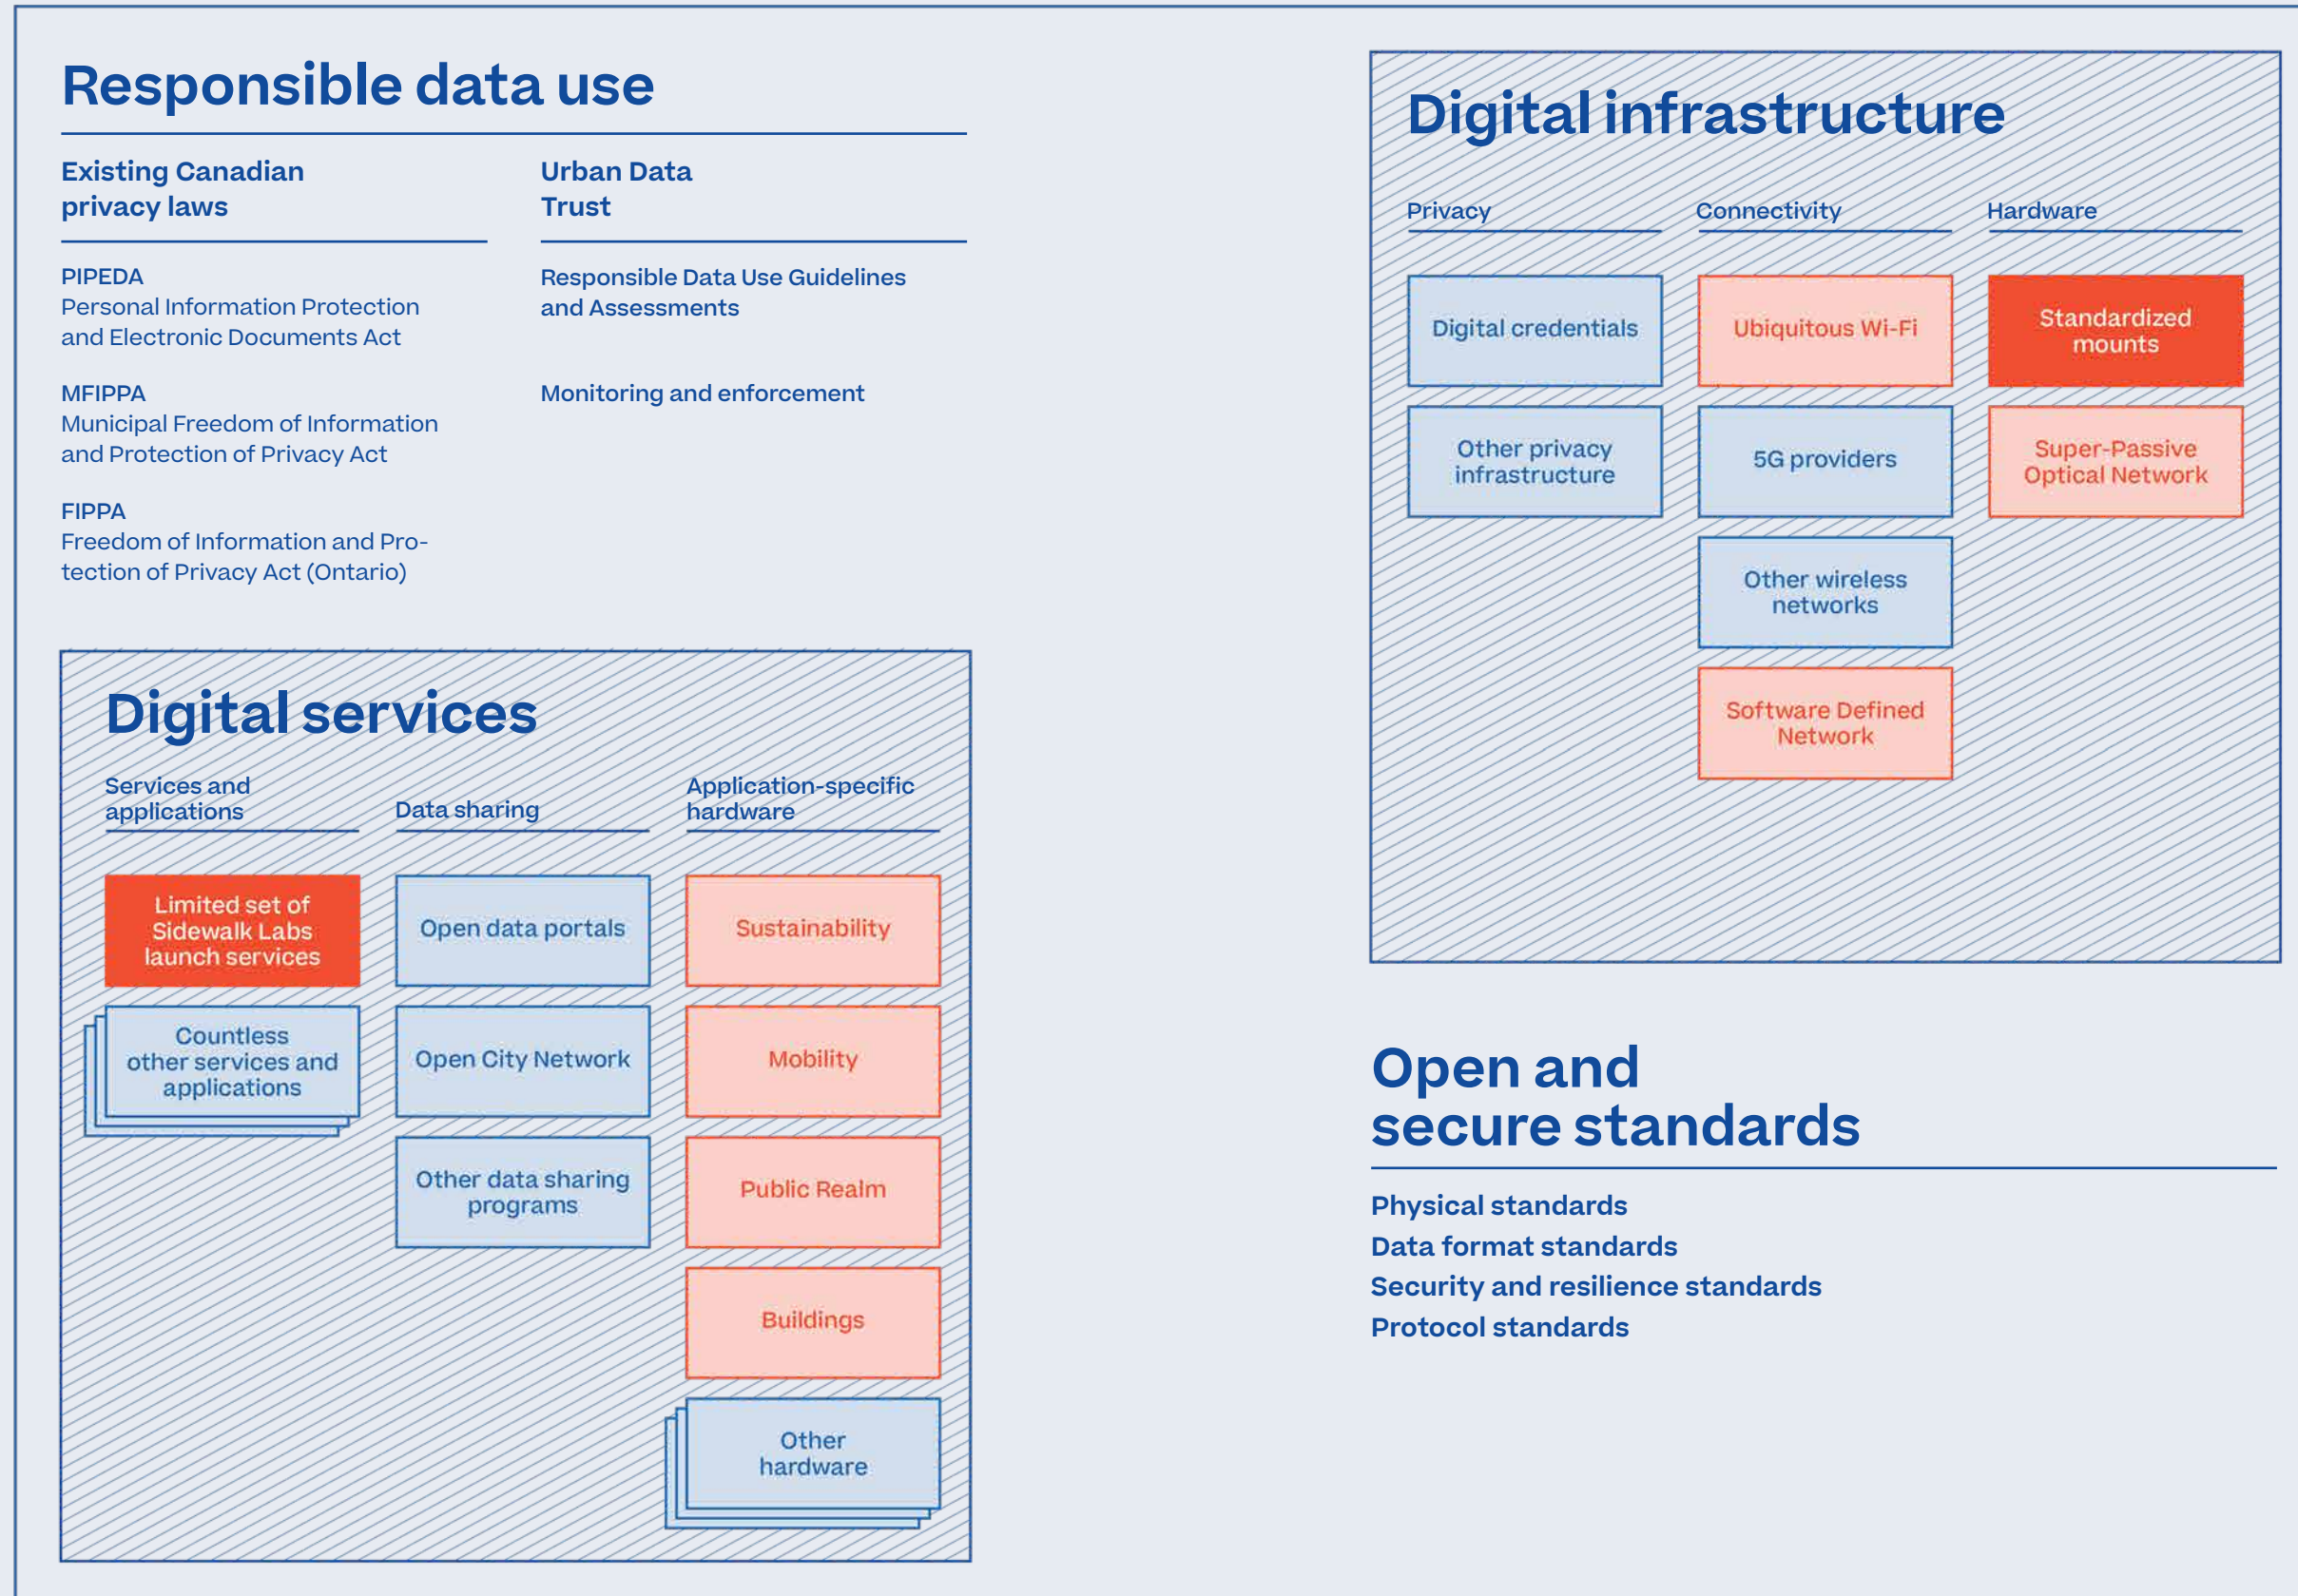
IDEA District

The 77-hectare Innovative Design and Economic Acceleration (IDEA) District, consisting of Quayside and the River District, provides sufficient geographic scale for innovations to maximize quality-of-life impact and to become financially viable.

Sidewalk Labs' role in creating the core conditions for digital innovation

Sidewalk Labs proposes to establish a set of core conditions that would catalyze an ecosystem of urban innovation along Toronto's eastern waterfront, consistent with Waterfront Toronto's objectives of improving quality of life and creating new economic opportunities in the digital age. These conditions include shared digital infrastructure, an open and secure approach to architecture and standards, a catalyzing set of digital services, and a trusted process for responsible data use.

As the diagram on this page shows, the role that Sidewalk Labs proposes to play would vary across these conditions and would follow a general approach of enabling innovation by others.

Responsible data use

Existing Canadian privacy laws

PIPEDA
Personal Information Protection and Electronic Documents Act

MFIPPA
Municipal Freedom of Information and Protection of Privacy Act

FIPPA
Freedom of Information and Protection of Privacy Act (Ontario)

Urban Data Trust

Responsible Data Use Guidelines and Assessments

Monitoring and enforcement

Digital services

Services and applications

Limited set of Sidewalk Labs launch services

Countless other services and applications

Data sharing

Open data portals

Open City Network

Other data sharing programs

Application-specific hardware

Sustainability

Mobility

Public Realm

Buildings

Other hardware

Digital infrastructure

Privacy

Digital credentials

Other privacy infrastructure

Connectivity

Ubiquitous Wi-Fi

5G providers

Other wireless networks

Software Defined Network

Hardware

Standardized mounts

Super-Passive Optical Network

Open and secure standards

Physical standards
Data format standards
Security and resilience standards
Protocol standards

Key

- Sidewalk Labs participation
- Some Sidewalk Labs participation but primarily third-party involvement
- Third parties play a leading role

General approach: Buy rather than build, wherever possible.

In keeping with its role as catalyst in the Sidewalk Toronto project, Sidewalk Labs prefers to purchase third-party technology — or partner with third parties to create (or enhance) it — whenever there are existing companies that have the capability and incentives to implement the systems required. Sidewalk Labs plans to give priority to technology that is local to Toronto, Ontario, or Canada.

In cases where technology does not currently exist, and where entrepreneurs or established companies are not building them, Sidewalk Labs plans to build the technology. These are likely to be cases that require significant up-front investment the market is not currently making, or where success focuses on longer-term objectives that other companies are designed to pursue.

In all cases, other entities would be free to develop and provide competing services to those offered by Sidewalk Labs.

Digital infrastructure role.

Sidewalk Labs plans to develop several components of digital infrastructure related to hardware, connectivity, and privacy, working alongside third parties to build out certain aspects of these systems.

For the proposed Wi-Fi network, Sidewalk Labs hopes to work with existing telecommunications companies with experience on the Toronto waterfront to build out infrastructure and conduct research and development of new technologies. Waterfront Toronto has worked for over a decade to eliminate the digital divide in their new communities, working with a local telecommunications provider to deliver gigabit service to every residential unit that gets built on public land, including in affordable housing.

For other infrastructure components, Sidewalk Labs expects to play a larger role that still involves others. These include standardized mounts that would reduce the cost of deploying digital innovations and an advanced optical network and software-defined network that makes connectivity faster and more secure. While Sidewalk Labs does not expect others to have sufficient incentives to create this infrastructure alone, it believes these components would play a critical role in boosting the success of digital innovations that address urban challenges.

Sidewalk Labs also expects third parties alone to provide other aspects of digital infrastructure that include 5G cellular connectivity (at much lower costs thanks to standardized mounts), other advanced communications networks, and additional privacy-enhancing infrastructure.

Digital services role.

To achieve fundamental quality-of-life goals through innovations the market has not pursued, Sidewalk Labs plans to offer a limited set of core digital services related to its essential programs for transportation, affordability, housing, energy, or public space. These services would rely on application-specific hardware devices created primarily by third parties but adapted or extended by Sidewalk Labs, working closely with these device manufacturers.

These launch services could still involve working with partners and buying existing technology. For example, the proposed mobility management system (see Page 452) could require computer-vision technology that performs de-identification at source, retaining an aggregate count of travellers but deleting any footage or images. Local companies are working on such technology, and Sidewalk Labs would explore options for purchasing those devices as this mobility system (or other proposed services) may require them.

Sidewalk Labs believes the urban data generated by these services would catalyze third parties to create countless other applications to improve quality of life, along with the application-specific hardware designed to support them.

For that to occur, this data must be shared publicly, and there are many companies and organizations in Toronto and beyond that specialize in making data available, such as ThinkData Works, the City of Toronto's Open Data Portal, and the Open City Network. Sidewalk Labs hopes to work with them to help provide the services necessary for the Sidewalk Toronto project.

Open and secure standards role.

Making data publicly available is necessary but not sufficient to catalyze digital innovation. That requires publishing the data in standard formats that third parties can easily build on, with good documentation for both the method of access and for interpreting the data format.

There are a small number of existing data formats for urban data, but Sidewalk Labs would focus on working with partners and standards bodies to develop, refine, and promulgate a much wider range of formats that support quality of life goals (see Page 403). Sidewalk Labs plans to take the same approach to standard communications protocols (such as software-defined networks), physical standards (such as standardized mounts), and security and resiliency standards (see Page 408).

Responsible data use role.

All digital innovations that propose to use or collect urban data in the IDEA District — whether developed by third parties or Sidewalk Labs — would be reviewed by and require approval from an independent Urban Data Trust (not controlled by Sidewalk Labs or Waterfront Toronto). These proposals would involve submitting an RDU Assessment to ensure that privacy and security are protected and that the innovations adhere to RDU Guidelines established by the Urban Data Trust. This proposed process would apply in addition to existing privacy laws.

Sidewalk Labs believes the Urban Data Trust could evolve into a public-sector or quasi-public agency over time.

By offering this unique set of catalyzing conditions in a defined geography, Sidewalk Labs hopes to encourage and invite countless urban innovators to view the IDEA District as a global launchpad for urban innovation.

Part 1

Providing More Affordable and Flexible Digital Infrastructure

Key Goals

1
Expand opportunity with ubiquitous connectivity

2
Reduce installation and maintenance costs with an "urban USB port"

3
Use distributed credential infrastructure to protect privacy

Digital infrastructure is a basic building block of the future city — the backbone of connectivity that helps residents, companies, organizations, and local agencies use data to launch new services that improve urban life. Many of the improvements to mobility, housing, energy use, and the public realm described throughout the MIDP are only possible today thanks to advances in digital infrastructure, such as fast internet connectivity and digital devices capable of collecting information.

Digital infrastructure is what enables an adaptive traffic light to prioritize a light rail vehicle that is running late, and what enables a heated bike lane to warm up in advance of a storm so a cyclist can get to work on a snow-free path. It is what enables an extendable awning to cover a ground-floor market space just before it rains, and what enables a small business to launch a pop-up at an affordable cost. It is what enables someone who suffers from asthma to request alerts whenever there is a decline in air quality, what enables a dishwasher to operate when energy is cleaner, and so much more.

Digital infrastructure is what unlocks these innovations, and more importantly, the significant leaps forward in affordability, mobility, sustainability, and opportunity that come with them. It is also the catalyst for new services or businesses no one has thought of yet, and the cornerstone of a digital economy. For the IDEA District to become both an inclusive neighbourhood that evolves over time and a hub for ongoing exploration into the next great idea for urban life, fast and low-cost connectivity should not be a luxury for the few — it should become the new standard.

But today's digital infrastructure can be expensive and difficult to replace. Too often, cities rely on proprietary hardware and software to collect data and connect people, locking them into using the same tools for years, even when better options become available. That makes it hard for residents, workers, and businesses to take advantage of the latest technologies that promise faster connections at lower costs.

Fast and low-cost connectivity should not be a luxury for the few — it should become the new standard.

Sidewalk Labs' proposal for digital infrastructure centres on two core hardware components. One is ubiquitous connectivity that would offer residents, workers, and businesses access to their own secure, super-fast internet network no matter where they are, at an affordable cost. The other is a new type of "urban USB port" that would provide a physical mount, power, and connectivity to digital devices in the public realm — such as Wi-Fi antennae, traffic counters, or air-quality sensors fixed to street poles and traffic signals — at much lower cost than the connected mounts cities use today.

Additionally, Sidewalk Labs plans to explore the use of a new type of privacy-preserving software infrastructure that would enable people to share only the minimum amount of information necessary to complete a transaction with a digital service or app, with the person's full consent.

These proposed components would not be exclusive; on the contrary, any third party could provide a competing offering.

At the neighbourhood scale of Quayside, ubiquitous connectivity could draw people outdoors, further bridge the digital divide, and provide secure access across the entire neighbourhood. However, this type of network would only become financially sustainable at a larger service area, given the number of residents or businesses needed to recoup the initial investment. Deployed at the full scale of the IDEA District, this advanced connectivity would dramatically reduce the time and effort required to set up networks and enable residents to use their own network everywhere — from their couch to a park bench.

Similarly, in Quayside, the proposed urban USB port would make it much easier and less expensive to deploy technology in the service of improving a neighbourhood. But new hardware standards require significant geographic distribution to gain the wide adoption needed for device manufacturers to incorporate the standard into their own designs; for example, a Wi-Fi antenna producer would not change its design for a small handful of cases. Deployed across the IDEA District, however, this standardized mount would reduce the time needed to mount a device in the public realm by 92 percent over current infrastructure.

At the full scale of the IDEA District, this approach to digital infrastructure would enable the creation of many urban innovations described throughout the MIDP — as well as all those waiting to be invented in the future.

Sidewalk Labs' role in digital infrastructure.
As explained on Page 382, in keeping with its role as catalyst, Sidewalk Labs would first look to others to help deliver its digital infrastructure proposals, including the proposed connectivity network, standardized mounts, and privacy-preserving software. Other infrastructure components, such as 5G, could be provided entirely by third parties.

Expand opportunity with ubiquitous connectivity

The internet is essential to modern cities: it is needed at all corners of a community at all times. To provide ubiquitous connectivity, Sidewalk Labs proposes a secure, high-speed, uninterrupted network across the IDEA District, both indoors and outdoors, that can support the use of roughly 10 million simultaneous devices.

Toronto's waterfront currently incorporates world-leading internet speeds, thanks to the work of Waterfront Toronto with its telecommunications partners. For example, in places like the Bentway, Waterfront Toronto has collaborated with telecommunications partners to provide free Wi-Fi as a way to extend digital access into the public realm.

Sidewalk Labs proposes to push this work even further by taking advantage of recent advances in fibre-optic technology and new approaches to network management. Sidewalk Labs would provide technical guidance and requirements and work with Waterfront Toronto's procured telecommunications partner to build out the required physical infrastructure and operate the network.

At the core of Sidewalk Labs' proposed network is the belief that residents, workers, and visitors should have continuous access to their own secure Wi-Fi connection everywhere they go, from the basement of an office building to sidewalk underpasses connecting the IDEA District with the rest of Toronto. This ubiquity would mean residents and workers

can stay connected to their own home or office Wi-Fi network no matter where they are, without worrying about joining an insecure network.

This type of ubiquitous connectivity would also create new opportunities for small businesses and entrepreneurs to get up and running faster, and for residents and community groups to focus their energy in new directions, whether that means launching a pop-up retail shop, showing a digital media art installation, or finding a new job.

Advanced optical network

As part of its network planning, Sidewalk Labs is exploring a new technology called Super-PON (Passive Optical Network).

Conventional fibre-optic networks are constructed with a stranded fibre-optic cable running from the network provider's central office to the user's site, typically a single building. This type of system can reach 32 or 64 users per fibre strand,[5] with 20 kilometres of transmission reach.[6]

In contrast, Super-PON technology is capable of supporting 768 users per strand and extending the reach to 50 kilometres[7] — meaning that a single cable could now provide connectivity to multiple buildings across a neighbourhood or district. Super-PON achieves this improvement by splitting light into many different colours (or wavelengths) over a single strand of fibre-optic cable, with

How Super-PON technology outperforms traditional fibre-optics on seven key metrics

	Typical network approach	Super-PON approach
Users per fibre strand	32–64	768
Maximum transmission distance	20 km	50 km
Wi-Fi signal interference	Signal interference from neighbouring homes and businesses degrades Wi-Fi connectivity, especially during peak usage	A continuously managed Wi-Fi signal optimizes for speed and coverage to prevent slowdowns, even at periods of heavy usage
Router configuration	Users independently configure their own routers	Configuration is automated and secure to simplify setup and increase security
Security	Firewalls configured per router, making access difficult and often opening security holes	Holistically configured routes that allow access for authorized uses only — simultaneously more convenient and more secure
Wi-Fi availability	Few public Wi-Fi access points; most access points configured for private access only; difficult to connect devices like smart switches, thermostats, lighting	Wi-Fi access points situated throughout the neighbourhood, indoors and outdoors, for seamless connectivity and access while remaining secure
Access to home or networks	Difficult to access when elsewhere without complicated, insecure custom configuration	Allows people to connect directly to devices in their homes, schools, and offices easily and securely using software-defined networks

each colour serving as its own signal.[8] In one possible configuration, each light wavelength (for example, red, yellow, or blue) would provide connectivity to a specific building.

This technology infrastructure could result in a higher-bandwidth network with a number of additional benefits. The ability to split cables among more users means the network would require less fibre material and physical infrastructure than traditional networks, enabling it to be constructed faster and at lower cost. The network would also use less electrical power because its extended reach requires fewer "stops" for a signal (a traditional network could require rooms with electric boosters every 20 kilometres).

This Super-PON specification is now being studied by the IEEE Standards Association,[9] the world's largest technical professional organization, for possible inclusion in its 802.3 international standards for telecommunications. If applied in Quayside, Super-PON would make Toronto the first Canadian city with this technology (it currently exists in San Antonio, Texas),[10] and would help ensure fast connectivity throughout the IDEA District.

Extensive fibre-optic backbone

Beginning in Quayside, Sidewalk Labs' proposed design for a fibre-optic backbone would be connected to two major internet Points of Presence (POPs) in downtown Toronto. The proposed designs would support at least 10 times the amount of anticipated bandwidth needed. Sidewalk Labs plans to evaluate whether an additional POP is required to provide sufficient redundancy.

The proposed network could support

10 times

the bandwidth needed in Quayside.

In Quayside, Sidewalk Labs proposes that the conduits holding the fibre have express and local routes, as well as regular handholes (access points). Each building would serve as an aggregation point for outdoor fixtures capable of mounting digital devices, such as street lights or poles, and would have fibre-optic runs to provide additional access if needed.

At the proposed full scale of the IDEA District, further enhancements could be possible, including laying out the fibre-optic backbone as a loop so that a fault at any location would not disrupt access further along the fibre.

Flexible building connections

In Quayside, Sidewalk Labs plans to ensure that buildings conform to the following specifications that balance the goals of this Super-PON network with the ability for other providers to offer their own network services:

Conduits.
Sidewalk Labs proposes that incoming conduits meet a set of specifications provided to all developers, including buried depth, distance from water and sewer lines, slope from buildings, coating materials, size and amount, and duct plug features. These conduits should either run directly to a "Meet Me Room," or connect with the matching number of horizontal conduits that run to the Meet Me Room.

The proposed fibre-optic network would be designed to reach every building in Quayside

Key

— Connection to Toronto Point of Presence

— Proposed fibre-optic network
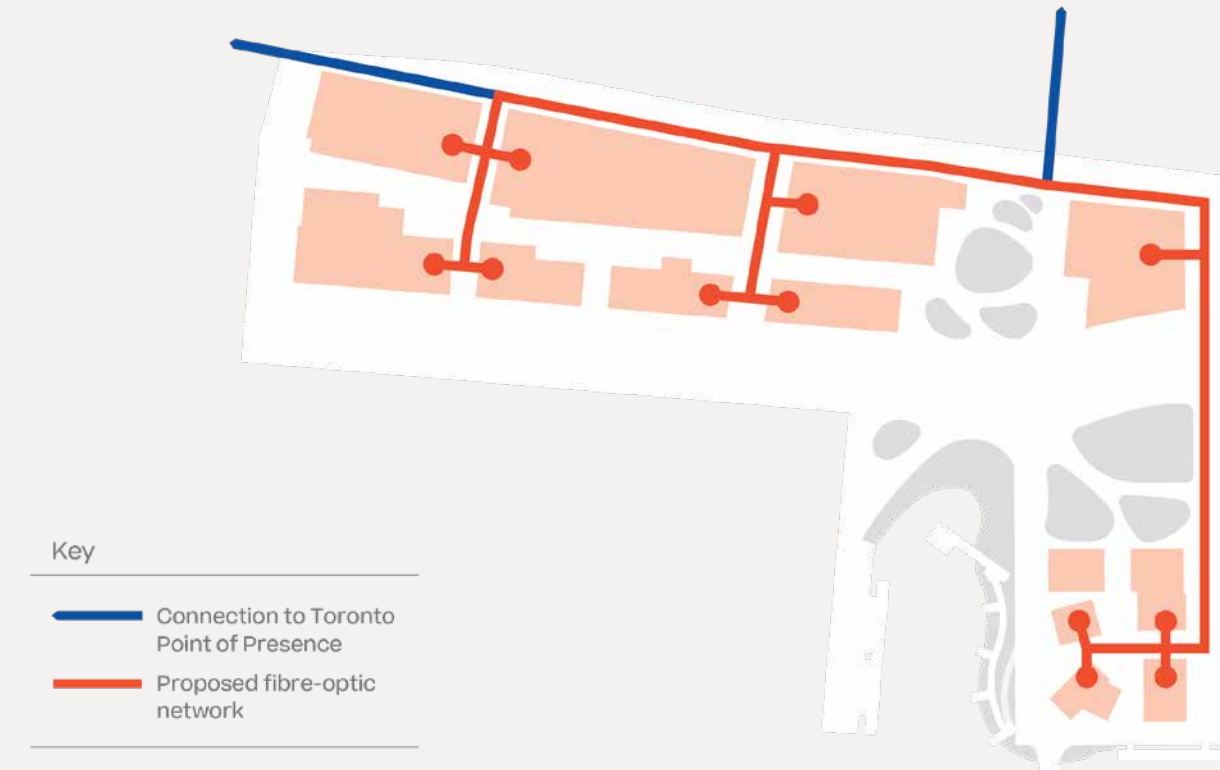
Meet Me Room.
This room would be a single location in the building where all communications-related equipment would be installed. It would be dedicated to communications use; other utilities should be located elsewhere to reduce risk of disruption of communications services. This room should have backup power and spare capacity for easy upgrades or new technologies.

Risers.
A vertical riser, dedicated to communications wiring, should be accessible on each floor and extend from base to the top floor and roof. The riser should be sized for future cabling. Ideally there would be two or more diverse risers that are separated by at least five metres for resiliency. Horizontal risers, on each floor, would connect each vertical riser to each individual unit.

Cabling.
Sidewalk Labs plans to implement Cat 6A wiring in each room for power-over-ethernet wireless access points, from a central point to form a local area network within the unit. This wiring would allow flexibility for installing additional radios — for example, the forthcoming 60 gigahertz products that offer multi-gigabit speed but whose signals cannot penetrate walls.
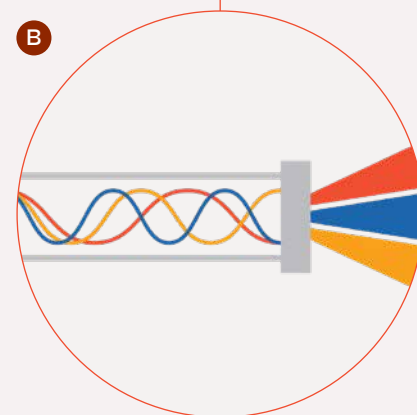
How it works: Super-PON connectivity

By splitting cables using new wavelength technology, Super-PON (Passive Optical Network) is capable of providing connectivity to multiple buildings across a neighbourhood or district.

Each building gets a dedicated wavelength (colour) on a single fibre stand, helping to reduce materials, reduce infrastructure, and increase speed.

A Third-party Point of Presence.
The fibre-optic backbone would be connected to two major internet Points of Presence in downtown Toronto.

B Super-PON fibre.
A single Super-PON fibre strand can serve multiple buildings in a neighbourhood.

C Meet Me Room.
A location in each building dedicated to communications utilities.

D Vertical riser.
A pipe or channel for communications wiring should be accessible on each floor and sized for future cabling.

E Loop return.
A circular structure ensures better access and fewer service disruptions.

Optimized wireless infrastructure

Next-generation wireless systems could offer amazing speeds, but they actually require significantly more antennae and wired backhaul connections than today's systems. Sidewalk Labs is working to determine the optimal location for antennae, both inside buildings and throughout the public realm, using software that automatically takes the site plans for Quayside and creates a predictive radio frequency study. This study includes locating Wi-Fi access points, mobile phone antennae (such as 4G, 5G, LTE, and 3.5 GHz CBRS), LoRaWAN gateways, and more.

A seamless and secure neighbourhood-wide network

When the internet was invented in the 1970s, every device could connect to every other device.[11] "Routers" performed the task of getting packets of information from the transmitting device to the receiving one, usually by taking multiple hops. Over time, the internet became less connected: for security purposes, some sub-networks (subnets) walled themselves off by having the router that connected them to the rest of the internet reject most incoming information packets. This was the origin of the internet "firewall" — a now-common feature of an internet router.

For this reason, it is very difficult for people to connect to a home device when they are not at home. Instead, they must engage with a home device (such as a smart thermostat or home-security camera) via a third-party website or app that this device contacts from time to time.

To help address this challenge, Sidewalk Labs proposes to take advantage of an emerging security approach called "software-defined networks."

As its name suggests, a software-defined network uses software to "define" the way that information travels through the network's hardware (its physical communications links and the routers that connect them). In such a system, users would not need to configure their own routers independently and have those routers reject all incoming communications using a firewall. Instead, the software-defined system would automatically configure the routers to create private networks that would remain available and secure across an entire neighbourhood — providing both greater convenience and heightened security.

Greater convenience.
In Quayside, these private networks would be available anywhere in the neighbourhood, including in parks and public spaces, using the ubiquitous Wi-Fi network. Using a neighbourhood software-defined network would enable people to connect to all of the same devices regardless of whether they are at home, in the office, in the park, in a light rail vehicle — anywhere. And nobody else (unless authorized) would have access to those devices. A neighbourhood-wide software-defined network could also make set-up easier than the current set of routers and firewalls that internet service providers use.

Consider, for example, a family that wants to check on their pet while they are out. Right now they would normally have to make sure their in-home video camera was cloud-connected, because otherwise they would lose contact with their camera as soon as they were out of range of their home Wi-Fi access point. A better approach would enable the family to access this video using data from their home directly, just as if they were at home, without that data having to be transferred or stored at any cloud provider. And just as some people use a virtual private network (or VPN) to connect to their office network, there would be a way to connect to the neighbourhood SDN when they are outside the neighbourhood to maintain the same access.

Heightened security.
A further advantage of software-defined networks is security. Because the software network would know what kind of data each device is supposed to be transmitting, it would be able to detect if any of them have been compromised. For example, if a thermostat that normally sends a few bytes every minute starts streaming megabytes per second, the software-defined network could quickly disconnect the device from the network — putting it in a kind of quarantine. This ability could help avoid "distributed denial of service" attacks and other exploits aimed at vulnerabilities in connected devices.

As with all digital infrastructure proposed by Sidewalk Labs, residents and businesses would not be required to use this network.

Digital infrastructure and inclusion

Building on the work of Waterfront Toronto to connect Toronto's waterfront communities, Sidewalk Labs plans to meet all the requirements for digital inclusion outlined by the National Digital Inclusion Alliance, a U.S.-based non-profit. Beyond affordable connectivity, these requirements include access to internet-enabled devices; quality technical support and digital literacy training; and applications designed to enable and encourage self-sufficiency, participation, and collaboration.

For those without smartphones or who require digital support, Sidewalk Labs plans to provide free-to-use devices, tech support staff, and digital literacy programming in the Civic Assembly and the Care Collective. This digital infrastructure would help the population seamlessly leverage digital tools for daily activities, advance in the digital jobs economy, and access critical services, such as government and healthcare support. It would also enable service providers to develop digital tools that they know can reach and support every community member.

To further encourage the development of truly inclusive tools, Sidewalk Labs is currently funding an inclusive usability testing program founded by Code for Canada called GRIT Toronto (see Page 443), working with local communities to develop a launch service aimed at participation in community decisions called Collab (see Page 446), and supporting Toronto-based service providers to develop technology solutions (see Page 382).

Providing More Affordable
and Flexible Digital Infrastructure

Reduce installation and maintenance costs with an "urban USB port"

Sidewalk Labs has designed a standardized mount called "Koala" that would make it fast, inexpensive, and safe to install a device on a light pole or other street fixture by providing a sturdy physical mount, power, and network connectivity. Just as USB ports made it easier to connect external devices with computers, this new type of urban USB port would create a standard connection point for cities that drives down the cost of installing and maintaining digital hardware.

Today, according to public records, Toronto has at least 11,000 devices mounted to public infrastructure, including Wi-Fi access points, cellular nodes, environmental sensors, and traffic or public safety cameras.[12] Installing these devices often requires significant disruption to street life, creates risks to workers in bucket trucks, and costs thousands of dollars, because light poles and other street fixtures were never designed to host digital hardware.

Adding a single car-counting device to an intersection requires the city to take the following steps:

Shut down a lane of traffic for hours or even days.

Send a bucket truck with several staff to the intersection.

Devise a creative mounting solution involving special clamps to adapt to the particular conditions of a traffic pole while maintaining safety standards.

Employ an electrician to shut down the supply to the pole and possibly run a network wire up the pole, a process that might involve digging a trench to the nearest connection point.

Repeat much of this labour-intensive process for repairs or upgrades.

Because this process of deploying digital hardware is so onerous, cities (and the private vendors they hire) tend to invest in high-priced, ultra-reliable devices that are expensive to repair and upgrade. If it were possible to deploy, maintain, and upgrade such devices in an inexpensive way, cities could buy much less expensive technology, replace the small fraction of devices that fail, and provide some redundancy of devices to improve reliability around things like Wi-Fi networks. They would also be able to upgrade technology on a much more rapid timeline and have more resources to conduct pilots or explorations for new services.
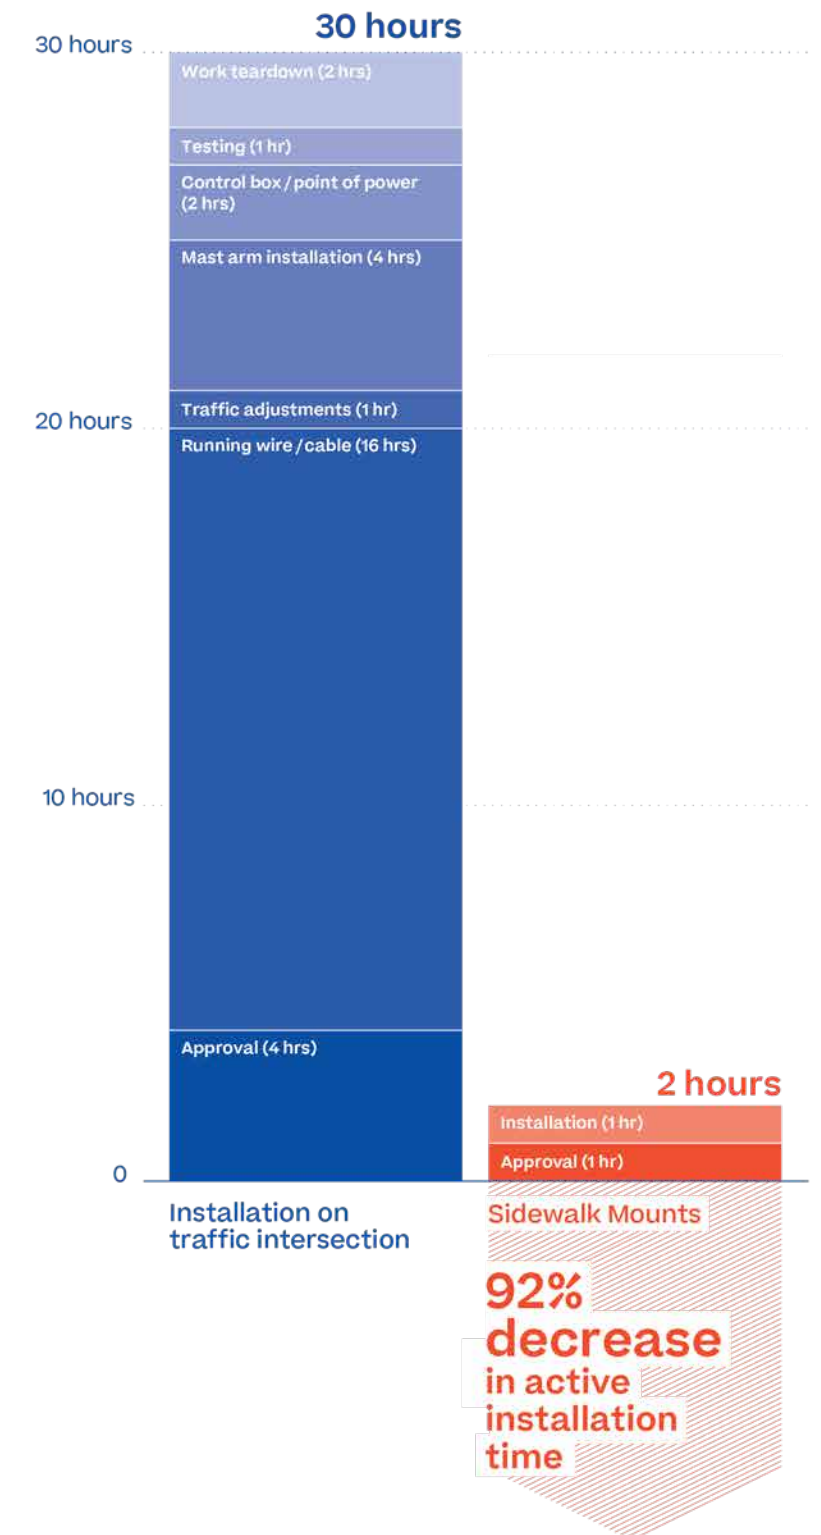
Sidewalk Labs' Koala mounts would provide a low-cost, low-fuss way for cities or third parties to improve urban life using urban data collected in the public realm. (All such data use would be subject to the proposed responsible data use process described on Page 414 of this chapter.) Koala mounts would be designed to provide power and connectivity to devices without the need to run new electric wires or close down streets. On the contrary, a device could be installed quickly using a common ladder or even a reacher grabber. Sidewalk Labs estimates its mounts would reduce the time of installation by roughly 92 percent — down from 30 hours today to two hours.

Koala mounts would be designed to work with any devices that meet its published standards, just like a USB port. As with Sidewalk Labs' ubiquitous connectivity network, companies would be free to use other mount offerings or stick with the traditional approach.

Device installation time savings of 92%

The proposed mount from Sidewalk Labs could dramatically reduce the amount of time it takes to install a device — down from 30 hours today to two hours. It could dramatically decrease costs, too. Assuming labour costs of $75 an hour, installing a device on a proposed mount would cost $150, compared with $1,980 for a standard traffic installation.

30 hours	
Work teardown (2 hrs)	
Testing (1 hr)	
Control box / point of power (2 hrs)	
Mast arm installation (4 hrs)	
Traffic adjustments (1 hr)	
Running wire / cable (16 hrs)	
Approval (4 hrs)	

Installation on traffic intersection

2 hours
Installation (1 hr)
Approval (1 hr)

Sidewalk Mounts

**92% decrease** in active installation time

A standardized mount to reduce disruption

The proposal Koala mount would create a standard connection point for digital devices that drives down the cost of installing and maintaining digital hardware.
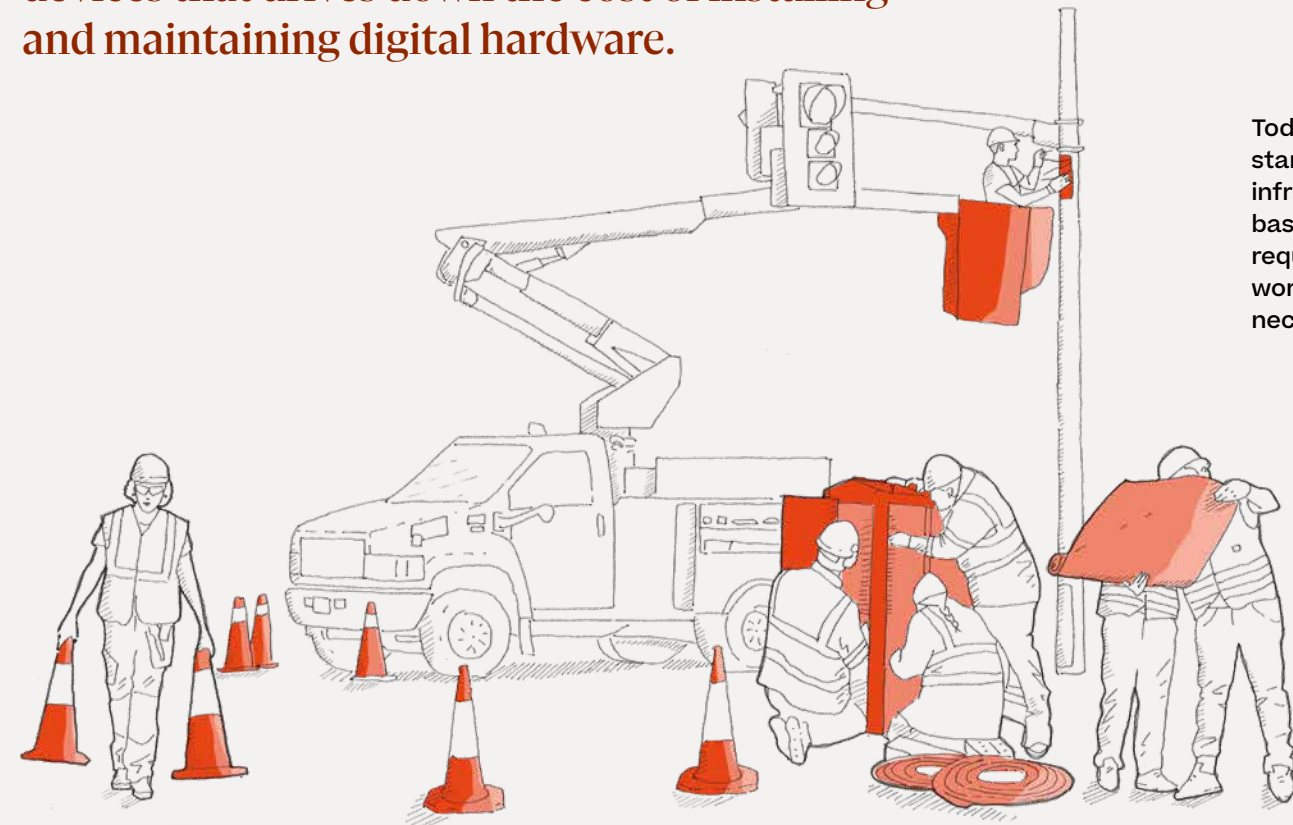
Today, without standardized digital infrastructure, even a basic traffic counter requires hours of work to mount, connect, and test.
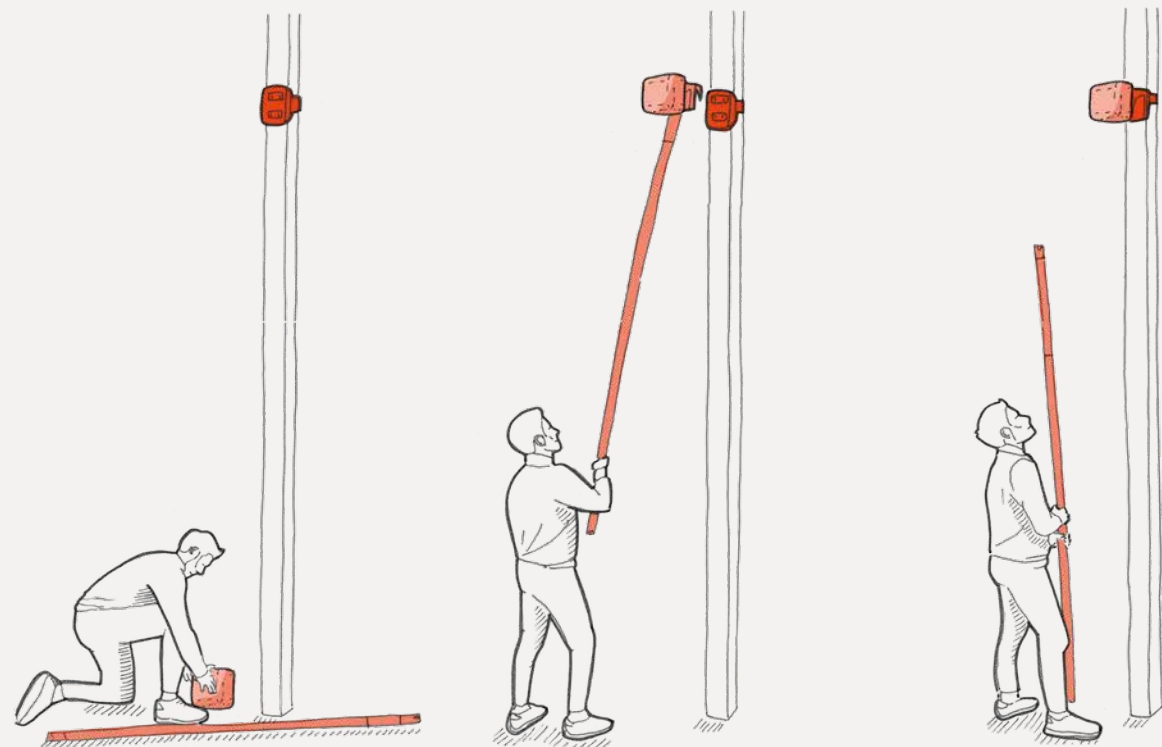
Koala mounts would make it easy and quick to connect to a ubiquitous network and collect urban data for a multitude of purposes, from bicycle counting to air-quality monitoring to interactive public art installations.

Koala mounts would provide a low-cost, low-fuss way for cities or third parties to improve urban life using urban data.

# Use distributed credential infrastructure to protect privacy

Many products and services in cities require some information about the people using them. But Sidewalk Labs believes that city residents, workers, and visitors should have to share no more information than absolutely necessary to use a digital service, receive a benefit, or conduct common personal or business transactions.

As an example, consider applying to rent an apartment. Potential tenants are often asked to reveal a lot of sensitive personal information as part of the rental application, such as their Social Insurance Number, driver's licence, tax history, and pay stubs.[13] But the minimum amount of essential information would likely include evidence of financial responsibility, such as recent credit history or score. It should not be necessary to include other information about the individual that could be used to discriminate against an applicant, such as their age or ethnicity.

To help tackle this challenge, Sidewalk Labs has been exploring the field of distributed digital credentials. This emerging approach uses privacy-preserving techniques to enable interactions such as the one described above in a way that provides only the minimal amount of information necessary, with a person's full consent over what information is shared.

Such privacy infrastructure is being developed by many groups around the world, including the open-source community, global organizations (such as the consortium piloting the DECODE project in Europe), startups, large financial institutions, and governments (for example, the Province of British Columbia). Sidewalk Labs plans to work with these types of groups to explore ways to incorporate this existing technology into many of its digital services that involve personal information, and to adopt a standard for handling personal data transactions in a trustworthy way.

This structure for digital services enables transactions between two parties that do not involve the creators of the digital services at all (whether Sidewalk Labs or another third party). Instead, credentials would be stored on user devices, not in the cloud (thus distributed, and not centralized), and the credential infrastructure would not act as an intermediary between the two parties. Continuing the rental application example, only the landlord and the rental applicant would ever have access to the information in their transaction.

In the rental application example, such a system could process a credential digitally signed by a trusted financial institution confirming the applicant's financial status without divulging further information that is not required for the application process — and with the applicant having full control over sharing this information.

Distributed credentials can ensure that people share the least information necessary to complete any digital transaction.

This interaction is enabled by technological advances in cryptography such as zero-knowledge proofs, digital signatures, and auditable data structures — which together make it possible for the applicant to prove their financial eligibility for an apartment without revealing data such as their name, address, or employer, all of which might bias a reviewer. In this case, zero-knowledge proofs allow the renter to prove their financial information is in an acceptable range without revealing exact values; the digital signature allows the reviewer to guarantee that the data is authentic and confirmed by a trusted counterparty like a bank; and auditable data structures give users the ability to make sure that no one has compromised their account or stolen their identity information.

In other words, only the people providing information about themselves and the service they are interacting with should know what is happening with the data involved — balancing the needs for privacy and authenticity for many types of urban interactions, both digital and physical.

Part 2

Setting Data Standards That Are Open and Secure

Key Goals

1
Enable third-party innovation with published standards

2
Use best-in-class resiliency and security

The ability to collect urban data is the first step to creating the conditions for digital innovation in the future city. But collection alone is not sufficient to use that information to create new services or tools that improve people's lives. To do that requires making the data publicly accessible to others in a way that encourages innovation but remains secure.

Perhaps the best example of a place catalyzing digital innovation via open standards is Estonia (see sidebar). The country's digital services platform, called "X-Road," makes it quick and easy for residents to do everything from apply for a bank loan to contest parking tickets to file their taxes.[14] And because the platform is publicly accessible through a published standard, the capital of Tallinn has become a hub of innovation in areas such as cybersecurity and blockchain technology.[15]

Standardized data formats, the kind that software developers can easily read and build on, are a key catalyst for digital innovation.

Of course, to create a vibrant ecosystem of new applications using data, that data must be provided in a standard format, with good documentation for both the method of access and for interpreting the data format. That is typically done through well-designed application programming interfaces, or APIs. APIs are standardized programming tools that enable computer systems to communicate; for example, when a Transit App shows bike-share availability at a nearby dock, it is using an API to connect with the bike-share system's real-time database, process that data, and display it on a phone.[16]

Currently, there is a gap between well-designed APIs and those of a typical open dataset. A well-designed API provides application developers with a clear description of the kind of data they can retrieve, the exact format the data will be provided in, sample code to access and use the data, and example applications that have been built using these same ingredients. That is not the way that the vast majority of open data is provided today. Making urban data available in ways that software developers can readily build on could provide the conditions for significantly increased innovation in city technology.

Key Term
APIs
are standardized programming tools that enable computer systems to communicate.

How Estonia's "X-Road" makes lives easier

At the start of the 21st century, only about one-third of Estonia's population had ever used the internet.[17] Less than 20 years later, this small Baltic nation of 1.3 million people is home to the most advanced civic data system in the world.

Estonia's residents go online to vote, file taxes, apply for bank loans, share education transcripts, view health records, contest parking tickets, and more. Estonians do not need to register their kids for kindergarten; the system does it for them, based on their child's date of birth and home address. The pet e-registry tells them when it is time for another round of vaccinations. Estonians do not even carry driver's licences or vehicle registration papers with them when they drive.

The only thing Estonians need is their e-ID card, which comes with two PINs to ensure security. The first PIN is for personal authentication when citizens log on; the second is for their digital signature, when they need to approve online transactions. And all those transactions take place on X-Road: the secure, government-run data exchange where residents interact with businesses and government.

Instead of notifying multiple government offices of a change of address, Estonians do it once, in the population registry, and give X-Road permission to share it with the voter registry, health ministry, banking institutions, and so on. X-Road shares only what it is instructed to share. And every time a third party views a person's information, it is traceable via a blockchain-style distributed ledger. Estonians can not only view their own health records, but also see which physicians and specialists have accessed them as part of their care.

X-Road processes half a billion queries annually, leading to substantial cost and time savings.[18] Transactions and verifications that used to take hours are completed in seconds. The process of registering a new business in Estonia takes 18 minutes;[19] by contrast, the same process in Ontario takes roughly 20 business days.[20] The country's courtrooms, once backlogged, are now remarkably efficient. Prescriptions flow from physician to pharmacist, and patients need not wait to get them written or filled. A 2015 World Bank report calculated that X-Road saved Estonians a total of 2.8 million annual hours — the equivalent of 3,225 people working around the clock for a full year.

The development of X-Road has given Estonia a competitive advantage in technology industries, helping to foster a robust startup ecosystem and giving the capital city of Tallinn a global reputation as a leading innovation centre. Estonia is also exporting X-Road to countries such as Finland, Moldova, Panama, and others.[21] As former Estonian President Toomas Hendrik Ilves told the *New Yorker*: "It's very popular in countries that want — and not all do — transparency against corruption."

Discussions of open data must also recognize the potential security risks that come with it. Addressing these risks begins with the network itself; as described on Page 392, a software-defined network could provide a heightened level of security by monitoring the amount of data that a device is transmitting and shutting off access if it detects anomalous behaviour. But security is not about implementing a single measure; rather, it best occurs with an established process for resiliency, transparency, and vigilance.

Sidewalk Labs proposes to catalyze innovation through the use of urban data that is both open and secure. First, Sidewalk Labs plans to develop and apply a set of published standards around open architecture, access, and sources that enable third parties to build on top of available information. Second, in support of that effort, Sidewalk Labs plans to use best-in-class security and resiliency techniques that aim to prevent disruptions, detect risks, and rapidly restore services.

Deployed at the full scale of the IDEA District, this plan for open and secure urban data would enable a vibrant ecosystem of urban innovation for startups, government agencies, researchers, civic organizations, and anyone else.

Sidewalk Labs' role in data standards. As explained on Page 382, in its role as project catalyst, Sidewalk Labs would aim to partner or rely on existing tools to achieve its goals for standards and security, including working with the many companies and organizations in Toronto that specialize in providing data in standard formats.

Setting Data Standards
That Are Open and Secure

Enable third-party innovation with published standards

At the core of Sidewalk Labs' approach to catalyze innovation is the belief in the importance of published standards for digital hardware and software, and public access to urban data that can reasonably be considered a public asset.

Openness is essential to provide new services that help improve quality of life and to inspire urban innovation by third parties. Just as no single company owns the web, no single company, organization, or agency should own the data or databases used by cities. They must be publicly accessible to improve upon, build on top of, or even replace.

Sidewalk Labs proposes a three-part plan to achieve its goal of a digitally open city. First, it proposes to provide data in standard formats and via well-defined, public APIs (open architecture), and where relevant standards do not exist, it would work with other companies, researchers, and standards bodies to create those standards. Second, it proposes to make this data publicly accessible by default (open access). Third, it proposes to make the software source code required for others to integrate with each of these systems publicly available under a free software licence (open source).

Open architecture avoids the

lock-in costs

of proprietary systems.

Open architecture: Public standards

All too often, today's cities buy bespoke, proprietary data systems from private vendors. The result is costly lock-in: the city must pay this provider forever for the use and support of the system or throw away the technology and pay a new provider for replacement.

For the Sidewalk Toronto project, any digital hardware and software that Sidewalk Labs creates would use public standards that make it possible not just to access data easily but also to replace aspects of the hardware or software itself, avoiding lock-in from a single technology provider and encouraging innovation.

This approach follows that of the World Wide Web. The reason that someone browsing the web can use any browser to view any web page, and that any web page could be served by any web server, is that the web is based on a collection of public, internationally recognized standards. These standards are a medley of letters: HTTP (how web pages can be requested), HTML (how text and images are specified), CSS (page formatting), SSL (security), and so on. Because these standards are universally followed, anyone with sufficient technical expertise can create a new version of any component of the web, including a new web server, a new web browser, or a new website.

Such standards have a number of advantages. First, they help ensure that no single company has a monopoly on providing a critical component. On the contrary, standards make it easy to improve — or even replace — any single component without throwing away the entire system.

Second, public standards inspire innovation. Web standards are now used for tasks that the creators never dreamed about. For example, standards originally designed for simple web pages are now used to support email, social networking, video-conferencing, virtual reality, and banking.

See the "Sustainability" chapter of Volume 2, on Page 296, for more details on the Brick standard.

Public data standards prevent any single company from monopolizing a critical digital system or component.

Where relevant standards exist, Sidewalk Labs plans to use them. These would likely include:

GTFS Realtime, a standard for reporting the location of public transit vehicles within the neighbourhood in real time (see sidebar)

General Bikeshare Feed Specification (GBFS), for reporting the availability of bike-share bikes and docks

Brick, a standard for describing building infrastructure, including HVAC systems

IFC, a standard for building information modelling, along with the Linked Data extensions

OpenStreetMap, a representation of roads and other public realm infrastructure

CityGML and CityJSON, standards for describing building shapes and sizes

OpenTraffic and OpenLR, emerging standards for describing traffic and street segments

Public Life Data Protocol, a standard from Gehl Institute on the use of public space

Sidewalk Labs commits to publishing an ongoing list of standards it uses, and proposes that the Urban Data Trust require other entities using urban data in the IDEA District to do the same.

Open architecture: APIs

Public data standards provide the lingua franca necessary for open architecture. Another important aspect is the methods by which data is exchanged via APIs.

As explained on Page 401, APIs provide a well-documented way for software developers to access public data. Too often today, even if a city makes its data publicly accessible, that data is too inconsistent and unpredictable to use without significant manual processing.

For example, if two entities collect the temperature in different parts of Toronto, an API would specify that both parties should use Celsius, collect the position of the data using latitude and longitude, and store the time in Coordinated Universal Time. If these parties did not agree to speak this common language before publishing their data, using that data correctly would be time-consuming and error-prone for software developers. The result would be that a startup or organization would have to invest a lot of money to standardize the data or, all too often, abandon an idea that might otherwise lead to a promising new service.

Sidewalk Labs plans to make its own APIs well-documented and publicly available, as well as to use public standards where they exist. Where public standards do not exist, Sidewalk Labs plans to work with others to define formats that could become standards in the future. Finally, Sidewalk Labs proposes that the Urban Data Trust ensure that other organizations and individual developers collecting and using urban data in the IDEA District do the same.

Innovation spotlight

GTFS: How transit riders get real-time trip data

Perhaps the best example of the power of open-data standards in an urban context is a format for transit data known as the General Transit Feed Specification, or GTFS. Its technical name notwithstanding, GTFS is easy to understand: it is what makes it possible for a navigation app to show users when the next streetcar, subway, or bus is scheduled to arrive.[22]

Not long ago, bus or subway riders standing on a street corner or platform had only the vaguest idea of when they would be on the move. The schedule posted in fine print on a pole offered no assurance. Their ride could be two, 20, or 200 minutes away.

Today, in most major North American cities, smartphone apps can tell riders when their transit vehicle is coming down to the minute, thanks in large part to GTFS. Initially developed in 2005 as a collaboration between Google and Portland, Oregon's TriMet transit agency, GTFS allows transit agencies and other developers to integrate static and real-time transit data into a wide variety of apps.[23]

GTFS has since served as the template for bike-share data (known as GBFS) and could do the same for everything from autonomous vehicle fleet movements to parking availability, allowing them to be integrated together. It is all part of a trend: providing better mobility not from more rail lines or asphalt, but from better and timelier information.

Open access

Publicly available data has enabled innovation across multiple industries by making it easy for students, researchers, and entrepreneurs to try out new ideas. To take one example, the openness of the web turbocharged research on information retrieval by providing access to public web pages. This research led to the creation of search engines, adding to the web ecosystem.

To take another example, in the late 1980s, the U.S. Census Bureau developed the Topologically Integrated Geographic Encoding and Referencing (TIGER) database to support the 1990 census.[24] The TIGER database describes land attributes, such as roads, buildings, rivers, and lakes. By releasing the data publicly, the census bureau enabled new services and products from digital mapping and navigation companies, such as NAVTEQ and TomTom, and eventually from online mapping services, such as MapQuest and Google Maps.

The time has come to prioritize not just the data that is easy to acquire and publish, but to gather and distribute data that will have the largest positive impact on quality of life. Sidewalk Labs believes that providing open access to data that has been expressly collected for the purpose of improving mobility, sustainability, accessibility, economic opportunity, and other aspects of urban life would have an even greater potential impact than much existing open data.

As described on Page 424, in the section on RDU Guidelines, Sidewalk Labs proposes that properly de-identified and non-personal urban data be made publicly accessible by default, enabling others to use it to create new services, tools, or products.

As an extension of this policy, Sidewalk Labs proposes that this information be integrated into existing open-data portals containing relevant urban data, including the Open Smart Cities Framework, the Toronto Open Data Portal, and the Ontario Open Data Catalogue — expanding access even further.

Open source

Once data is made publicly available in standardized formats through well-documented interfaces, anyone with sufficient expertise could, in principle, create innovations that integrate with urban infrastructure and digital services. But that does not make it easy. Parsing the standard formats, processing public data for particular common purposes, or communicating with APIs often takes a lot of time and effort — and reduces the likelihood that innovators will engage and succeed.

Where there are common tasks like these, Sidewalk Labs plans to share its software code publicly as "open source" — under licences like the Apache License (Version 2.0) or the MIT License — and encourage others to do the same. This approach has become common practice in the software industry, because it increases engagement with software systems. Over time, with contributions from software engineers across the world, this approach creates more robust and useful software.

In keeping with the belief that open-source tools inspire creative new uses, Sidewalk Labs has released several of its tools as open source, including the CommonSpace app for supporting public life studiesand the Toronto Transit Explorer prototype (available through the Sidewalk Toronto website). Sidewalk Labs plans to continue doing so in the future and to encourage others to do the same. 📖

See the "Public Realm" chapter of Volume 2, on page 118, for more details on CommonSpace.

Launching an open-source transit tool
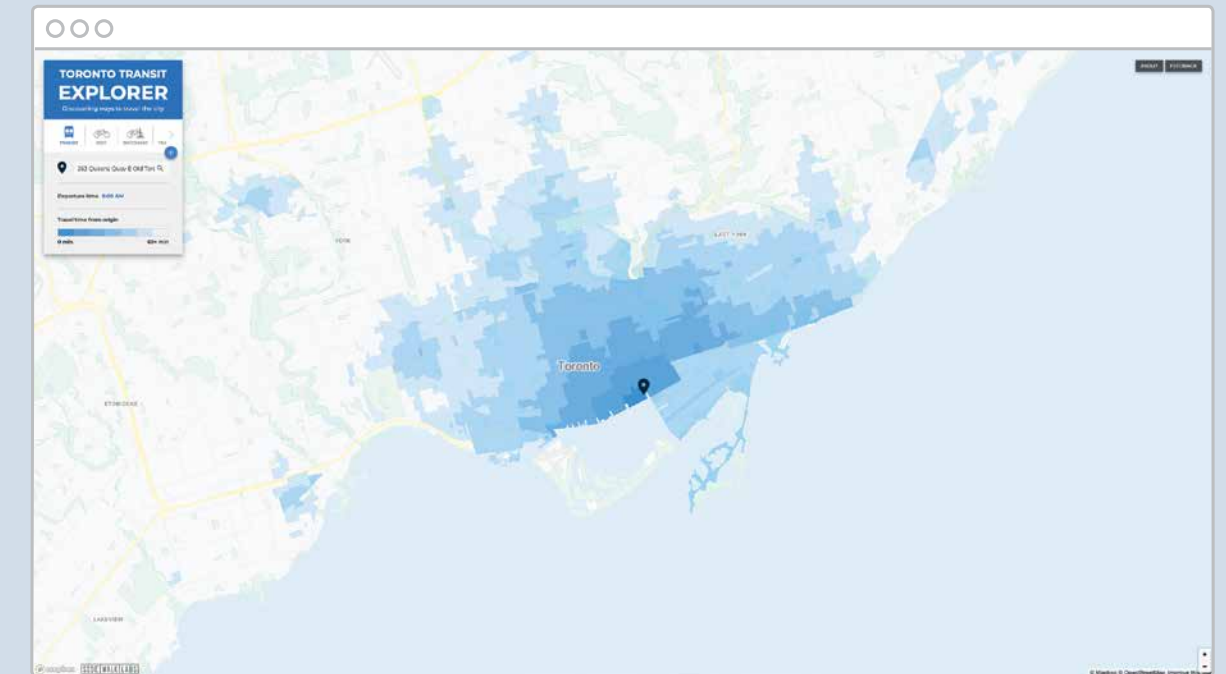
The Toronto Transit Explorer's open-sourced data format, front-end visualization, and server code enable others to improve the tool over time.

As an exercise in getting to know Toronto, while using open data and open-source software, Sidewalk Labs developed and launched a tool called the Toronto Transit Explorer in 2018.[25] The tool lets Torontonians explore how easy it is to get from any point in Toronto to any other using a range of travel modes.

To create this tool, Sidewalk Labs improved an existing open-source transit router called R5, adding features such as the ability to combine bike-share and transit into a single trip, as well as the ability to filter for wheelchair-accessible transit. Sidewalk Labs published these changes publicly so others could take advantage of these improvements in the future.

Sidewalk Labs then created a web application for exploring Toronto's transportation options and a server that used the improved R5 router to calculate data on the fly for the user interface.

Early iterations of the app were shared at the first two Sidewalk Toronto Public Roundtables and at a Civic Tech Toronto meetup. This important community feedback led to a redesign that made it easier for people to choose their origin and destination points.

To enable others to take this work and create new apps and variations along similar lines, Sidewalk Labs open-sourced the Toronto Transit Explorer front-end visualization as well as the server code under the Apache License (Version 2.0). Sidewalk Labs has since received feature requests, code contributions, and ideas for improving the tool from doctoral students, urban planners, software engineers, and members of the Toronto community who saw the potential for using the tool in their own work.

Goal 2

Use best-in-class resiliency and security

The digital systems and services proposed in the MIDP would help improve street safety, clean energy use, construction efficiency, and more. But connecting these systems creates new risks; intentional actions, inadvertent disruptions, even weather-related or environmental events could have a negative impact on digital services or infrastructure.

Planning for these risks requires a high level of security and reliability. Technologists often focus on digital security to prevent intentional acts. Sidewalk Labs plans to build on that foundation to ensure that the digital technology used in the IDEA District is resilient as well as secure. Digital systems should not only be secure from hackers — they should also be reliable in the face of inadvertent actions or environmental effects and maintained in a way that keeps them functioning at a consistent level over time.

Sidewalk Labs' approach to digital reliability emphasizes three design goals. First, as much as possible, prevent disruptions and the loss of functionality. Second, rapidly detect any loss in functionality or increased risk of loss of functionality through audits and other approaches. And third, prepare to rapidly restore functionality to any service that experiences a disruption.

These priorities are modelled after the standard approach taken by government and municipal services to ensure the

resilience of critical systems, and are parallel to the software architecture concept "security by design." Security by design refers to the principle that rather than being an afterthought, security should be considered at the beginning of the systems design process. This approach avoids designing a system or service in a way that makes security less effective or more difficult to implement.

Preventing disruption

Digital systems should, wherever possible, use public standards and open-source software with strong institutional and community support. This approach includes using tools like OpenSSL and the Linux kernel, which large organizations and governments around the world already depend on.

By using these tools, if a potential failure mode is discovered, a significant global community with a shared sense of urgency can help to address the issue. If any participating member of the community discovers a problem, all members can contribute to and benefit from the fix. Sidewalk Labs plans to use the Common Vulnerabilities and Exposures system — a public catalogue of security threats used by many other public- and private-sector digital service providers — to learn about and mitigate potential problems.

Additionally, Sidewalk Labs plans to give preference to the modularity of systems whenever possible, making it easier to

isolate any component of a system that might experience a disruption and to replace any individual component with newer technology.

When open-source software is not available, Sidewalk Labs plans to develop tools in concert with the security community. This effort could include inviting security and reliability researchers to test various systems, following the industry practice of issuing "bug bounties" to researchers who responsibly disclose issues or help patch vulnerabilities. Sidewalk Labs plans to run regular tests with a "red team" to simulate security breaches and failures.

As new technology emerges, best practices change. That makes specific recommendations (such as using a certain encryption method) less appropriate, effective, and nimble than having a broad strategy to remain up-to-date with — and be able to adjust in response to — emerging recommendations by the security community. Sidewalk Labs plans to use this broader, more resilient approach for all the technologies it develops or maintains.

For example, when using cryptography, Sidewalk Labs would not develop its own methods of encryption, and instead would use algorithms certified by the Cryptographic Algorithm Validation Program, the cryptographic standards program run by the U.S. National Institute of Standards and Technology and the Canadian Communications Security Establishment. Similarly, Sidewalk Labs plans to follow security and reliability standards defined by the greater community, including two notable benchmark security standards, SOC2 and ISO27001, for applicable products and services.

Though best practices in cybersecurity are always evolving, there are a number that Sidewalk Labs follows today, including:

Encrypting as much data as possible in storage and in transit using AES keys of 256 or 512 bits

Storing keys in a key management system backed by FIPS 140-2 Level 3-certified hardware security modules

Enabling client-managed encryption keys running on top of the same modules for any storage or computing resources to third parties

Using HMAC to ensure message integrity with symmetric encryption

Preferring elliptic-curve-based approaches over RSA for asymmetric encryption and digital signatures

Using SHA-256 for general hashing and bcrypt for passwords

Preferring multi-factor authentication methods over passwords alone

Routing all traffic through TLS and, when that is not an option, physically partitioning devices from other networks

Detection and auditability

Ongoing auditability is an important way for the security community to confirm the integrity and reliability of a digital system. Sidewalk Labs plans to use auditing systems such as Trillian to achieve this objective and would closely follow the state of security research to maintain best-in-class approaches.

Additionally, Sidewalk Labs would have regular third-party audits of any platforms and code it maintains, not only to confirm that it is consistent in running the same software it shares but also to confirm that it meets the quality expected by the Urban Data Trust. As part of this effort, Sidewalk Labs plans to build both technical and policy-based controls to provide strong assurance to the community that the digital systems it implements are behaving consistently with the Urban Data Trust's expectations.

Another key approach to transparency and auditability is the use of modular systems. Modularity enables a high degree of transparency: even when data itself is encrypted, the amount of data being transferred between systems can be shared, when appropriate, to provide guarantees about what is being saved and transferred. For example, an auditor who sees a very low amount of data leaving a computer-vision camera would know that data is being processed on-site and that the raw video is being deleted — even while the data itself would not be visible to the auditing party.

Finally, Sidewalk Labs is eagerly evaluating the growing field of transparency and auditability for machine learning and artificial intelligence. As the field develops, Sidewalk Labs plans to synthesize findings and principles established as best practices in industry and academia. Broadly, Sidewalk Labs believes that machine learning should be as auditable and transparent in its decisions as traditional software and engineering are (see sidebar).

In the case of a disruption, practicality may require keeping information temporarily contained to the people managing the incident and relevant authorities; for example, security vulnerabilities need to be patched before they are shared. But Sidewalk Labs plans to give strong preference to publication, including regular external audits, and commits to sharing publicly full post-mortems of any incident or report once resolved or stabilized.

Sidewalk Labs' commitment to "Responsible AI"

Many Canadians interact with artificial intelligence systems on a daily basis. Some applications of AI are as benign as email spam filters. Others carry more significant impacts, such as how banks approve loan applications.

One very common example of AI exists in "recommender" systems, which try to predict the preference or rating an individual would give to an item. Recommender systems function by collecting and analyzing the behaviour or activity of individuals and by comparing individuals to others who are similar to them. Many common recommender systems are considered helpful — for example, they can pre-populate a music playlist based on listening history. But some recommender systems can impact individuals in more significant ways or reveal potentially sensitive information about that individual.

The continued development and use of AI systems raises digital governance challenges that go beyond privacy. It is possible for organizations to be in full compliance with privacy laws yet still use data in ways that could impact people in harmful or unexpected ways.

To help protect against these unexpected outcomes and guide its use of AI, Sidewalk Labs has developed a Responsible AI framework guided by six overarching principles that are contextual, progressive, and applicable to all types of technology (existing and future). This framework is inspired by leading international standards, such as the Declaration on Ethics and Data Protection in Artificial Intelligence, which was signed by the Privacy Commissioner of Canada.[26]

(These principles would work alongside the proposed RDU Guidelines described on Page 424.)

Fairness and equity.
All projects involving AI systems should be designed and developed responsibly from the start and should consider an individual's reasonable expectations and the original purposes of data collection.

Accountability.
Organizations should always remain accountable for the AI systems they create and deploy.

Transparency and explainability.
Individuals should be informed when they are interacting directly with an automated system and when their personal information is being used to make consequential decisions about them. When feasible, AI systems should be designed with the ability to be explained in terms people can understand. In addition, AI inputs (or training sets) and potential biases should be understandable and debuggable.

Relevance.
All AI systems should be developed and designed with high standards of scientific excellence and with a multi-disciplinary approach that includes sharing research and best practices with regard to AI.

Value alignment.
AI systems should be designed, developed, and used in line with international human rights and local community values.

Respect for human dignity.
Individual autonomy and agency should be upheld through a diverse and multi-disciplinary design process. AI systems should be used to empower individuals and communities and enhance public engagement.

Preparedness and response

Designing plans for detection of or response to incidents requires anticipating potential issues (a practice known as "threat modelling") and setting up processes for continuous readiness to respond to a service disruption.

Threat modelling is an iterative process that seeks to identify the assets of an application or service that are at risk of disruption. These assets are then reviewed for mitigations of potential issues (or "threats") against their integrity. The risks posed by these threats are evaluated by taking into account factors such as the likelihood of some external factor triggering a disruption.

Response readiness focuses not only on preparing plans for responding to the threats generated in the modelling exercise, but also on ongoing drills to practice the plan. In many cases, this readiness requires staff, drills, and ongoing collaboration with external stakeholders to ensure that there are clear lines of communication in the event of an incident.

Each digital system that Sidewalk Labs implements for the Sidewalk Toronto project would use a preparedness assessment (see Page 413) to provide clear answers to key questions on threat modelling and response readiness. These assessments would be reviewed by a Sidewalk Labs security team as well as by parties that operate or maintain relevant dependent systems; for example, the potential for a problem with a traffic management system (an upstream system) requires designing a strong line of communication with emergency services (a downstream dependent).

Prioritize data residency

The decision on where to store data (known as data residency) is based on many considerations, including whether there is sufficient technical and physical architecture to store the data securely, the cost of storing the data abroad versus in the organization's home country, and applicable laws.

As with all matters relating to data, Sidewalk Labs' approach begins with a baseline that abides by existing laws. Canada's federal private-sector privacy law does not require data to be stored or processed solely within Canada. Instead, it seeks to make organizations accountable by imposing obligations to ensure that data is properly safeguarded. Similarly, the federal and provincial public-sector privacy laws that may be applicable do not dictate data residency. Sidewalk Labs continues to monitor developments in this area, including the Office of the Privacy Commissioner of Canada's consultation on the transborder flow of data, initiated April 9, 2019.

During the development of the MIDP, Sidewalk Labs engaged with numerous stakeholders and community interest groups to guide its approach to data residency, and heard clearly the desire to store data in Canada. For that reason, Sidewalk Labs commits to using its best efforts at data localization — for storage, processing, and communication — as long as there are Canadian-based providers who offer appropriate levels of security, redundancy, and reliability. To the extent that it is deemed infeasible to store data solely in Canada, Sidewalk Labs would be transparent about such a decision.

Information about data residency would be part of the proposed RDU Assessment (see Page 429) required for all parties.

Preparedness assessments enable faster responses to security risks

To improve security and resiliency for digital systems, Sidewalk Labs plans to use a preparedness assessment. Such documents aim to identify security risks as well as mitigation approaches through questions around threat modelling and response readiness.

The questions on this page are included here for illustrative purposes only.

Threat modelling

What are the ways in which this service could be disrupted (such as partial outage, corrupted data, full outage, and illicit access or control)?

Assess the likelihood of each disruption and (if available) any potential known ways that each disruption could be triggered.

For each of these scenarios, will any systems external to the service be affected?

For each of these scenarios, how will the disruption be detected? Could the disruption avoid detection?

For each of these scenarios, are there up-front investments that can lessen their effect?

For each potentially affected service listed above, what is the escalation path for notifying that service of a disruption?

Response readiness

For each of the scenarios above, please provide a playbook describing a communication and mitigation plan.

How regularly will there be drills practicing the protocol outlined in the playbook?

Do these drills involve downstream and upstream stakeholders?

Will there be "on call" staff available for response?

If no, outline a response plan that obviates the necessity for staffing.

If yes, outline the responsibilities and training for this staff. Also outline a continuity plan for maintaining this staff.

Part 3

Creating a Trusted Process for Responsible Data Use

Key Goals

1
Implement the Urban Data Trust

2
Establish RDU Guidelines

3
Set a clear process for urban data use or collection

In addition to flexible digital infrastructure and published standards, a third core condition for digital innovation is instilling community trust that information collected in cities will preserve the privacy of individuals and be used for the greater good — while promoting the growth of new businesses and the rise of new tools to improve urban life.

The pace of change for digital technologies such as the internet, social networks, and artificial intelligence has accelerated globally. When Canada established its federal private-sector privacy law, known as the Personal Information Protection and Electronic Documents Act (PIPEDA), some 20 years back,[27] just 42 percent of the population owned a personal computing device and smartphones did not exist.[28]

Canada is poised to lead a change. Canada recognizes privacy as a fundamental human right, with the right to privacy rooted in the Canadian Charter of Rights and Freedoms.[29] On top of that foundation, recent conversations convened by federal, provincial, and municipal regula-

tors have called for stronger national and provincial data strategies that protect individual privacy while enabling companies to create valuable new services using data, rather than competing to own data outright.

All three levels of government are at various stages of consultations with the public. The Government of Canada launched national consultations on digital and data transformation in 2018.[30] Ontario launched its data strategy consultations in early 2019.[31] The City of Toronto also announced it would begin to develop a city-wide policy framework and governance model associated with digital infrastructure.[32]

The Sidewalk Toronto project itself has sparked significant conversations about a new approach to digital governance in cities, generating new ideas from Canadian experts, stakeholders, and the public. This ongoing, comprehensive engagement and consultation has shaped the ideas Sidewalk Labs is proposing in this MIDP and would continue to help them evolve with the project.

How public consultation shaped Sidewalk Labs' ideas

To receive guidance on a full range of issues relating to responsible data use, Sidewalk Labs convened a Data Governance Working Group made up of independent experts and community representatives. Sidewalk Labs and this group have benefited from other insights, including those of Waterfront Toronto's Digital Strategy Advisory Panel.[33] Sidewalk Labs also consulted with all levels of government, and met with the Office of the Privacy Commissioner of Canada, the Information and Privacy Commissioner of Ontario, and various departments within the City of Toronto.

Such collaboration has been critical, because there is no comprehensive and unified digital governance model in Canada for the type of community Sidewalk Labs hopes would emerge within the IDEA District. The aforementioned consultations being driven by the three levels of government represent important starts to this conversation, and Sidewalk Labs offers the proposal in this chapter for consideration.

Over the course of its own public consultation to date, Sidewalk Labs has heard three key themes that have helped shape its proposal.

Canada is poised to lead a global change when it comes to data governance strategies.

1

What we heard: Protect more data.

The first theme was a recognition that while it is paramount to protect personal information, as Canada's privacy laws currently do, individual privacy is only part of the discussion around responsible data use.

Existing privacy laws only apply to or protect "personal information," meaning information about an identifiable individual. Sidewalk Labs heard through its consultations that Torontonians are also concerned about the collection and use of data gathered in the city's public realm, publicly accessible spaces, and even some private spaces — whether or not that data identifies specific individuals.

This type of data collection merits special focus for a variety of reasons. Its collection in public spaces raises concerns about surveillance that are exacerbated by computer processing power and the proliferation of sophisticated digital tools, such as cameras and sensors. Certain types of this data might reasonably be considered a collective public asset. Individuals are also not always aware of either the collection or use of such data. For example, in the case of on-street pedestrian counters or lobby cameras, collection and use notices often lack adequate information to fully inform individuals, are not visible until the individual is within the field of view, do not consider language barriers, or are absent altogether.

Furthermore, Torontonians are concerned about how the collection and use of non-personal information could impact groups of people or the community. For example, federal privacy commissioner guidance encourages companies to consider the potential impacts that

aggregated or de-identified data can have on individuals or communities at large, but companies could benefit from further guidance and comprehensive standards.[34]

How we responded: A new category of "urban data."

For all these reasons, Sidewalk Labs proposes a new category of data called "urban data" that includes both personal information and information that is not connected to a particular individual. The term "urban data" nods to the fact that it is collected in a physical space in the city and may be associated with practical challenges in obtaining meaningful consent. Urban data therefore seems worthy of additional protections.

Urban data would be broader than the definition of personal information and include personal, non-personal, aggregate, or de-identified data (see sidebar) collected and used in physical or community spaces where meaningful consent prior to collection and use is hard, if not impossible, to obtain. In that sense, urban data would be distinct from more traditional forms of data, termed here "transaction data," in which individuals affirmatively — albeit with varying levels of understanding — provide information about themselves through websites, mobile phones, or paper documents.

The proposed responsible data use process would protect urban data while building on existing protections for personal information — knowing that both urban data and transaction data must be handled responsibly for a better city.
Of course, the creation of a new term creates positives and negatives for companies and regulators alike, and Sidewalk Labs welcomes additional discourse on this term and its use in the context of the Sidewalk Toronto project.

Four types of urban data

There are different ways urban data can be categorized, each with different impacts on individuals and groups of people.

Non-personal data is data that does not identify an individual and can include other types of non-identifying data that is not about people. Some examples of non-personal data are aggregated data sets, machine-generated data (such as weather and temperature data), or data on maintenance needs for industrial machines. There are many benefits for consumers and members of industry to processing this type of data. The European Union recently passed a regulation protecting the free flow of non-personal data.[35] Even though non-personal data is not about identifiable individuals, it can still have unintended harmful impacts on people — for example, if AI systems use aggregated data sets to make predictions or recommendations to individuals.

Aggregate data is data that is about people in the aggregate and not about a particular individual. Aggregate-level data is useful for answering research questions about populations or groups of people. For example, aggregate counts of people in an office space can be used in combination with other data, such as weather data, to create an energy-effi-

ciency program so consumption is controlled, with the goal of saving money and reducing energy use. As with other types of data, the use of this data can have bias and fairness consequences.

De-identified data is data about an individual that was identifiable when collected but has subsequently been made non-identifiable. Third-party apps and services may wish to use properly de-identified data for research purposes, such as comparing neighbourhood energy usage across a city. When data is de-identified correctly — using principles including k-anonymity, and frameworks such as differential privacy — it is no longer personal information. While de-identification of data may not completely eliminate the risk of the re-identification of a data set, when proper guidelines and techniques are followed, the process can produce data sets for which the risk of re-identification is very small. The Information and Privacy Commissioner of Ontario has released a set of De-identification Guidelines for Structured Data, which provide basic concepts of and techniques for de-identification. The guidelines highlight the key issues to consider when de-identifying personal information and provide a step-by-step process for removing personal information from data sets. The

biggest risk of using de-identified data is that it is sometimes possible to link pieces of information together to re-identify the individual.[36] This risk can be mitigated by having trusted external experts regularly attempting re-identification in a controlled environment, in order to harden the system.

Personal information has a legal definition in Canada and is the subject of privacy laws, including PIPEDA.[37] The broad legal definition of personal information includes any information that could be used, alone or in combination with other information, to identify an individual or that is associated with an identifiable individual. Individuals routinely share their personal information with governments and businesses, whether applying for a licence or business permit, shopping, or ordering a ride-hail service. In some cases, personal information has to be shared to receive the service; for example, when people order food for delivery, the restaurant needs to know where to deliver it. Individuals often receive benefits from sharing their personal information, but society has seen many of the harms from illegal or unethical uses of personal information.

2

What we heard:
Consider urban data a public asset.

A second big theme heard during public consultation was that, in addition to personal and collective privacy, Torontonians are concerned with the ownership and stewardship of urban data.

Increasingly, some types of urban data can be understood as a community or collective asset. Take the example of traffic data. Since that data originates on public streets paid for by taxpayers, and since the use of that data could have an impact on how those streets operate in the future, that data should become a public resource.

In its extensive consultations with the public, stakeholders, government, and expert advisors, Sidewalk Labs heard that data collected in the public realm or in publicly owned spaces should not solely benefit the private or public sector; instead, it should benefit multiple stakeholders, provided any privacy risks have been properly minimized.

Part of using data responsibly involves making sure that no one entity — Sidewalk Labs or another — controls urban data that could reasonably be considered a public asset. The opportunities to use urban data to create new digital innovations must be available to everyone, from the local startup to the global corporation.

How we responded:
An independent Urban Data Trust.

If urban data is a common good, it should not be exclusively "owned" in the traditional sense. The question then becomes: Who should be the steward of urban data? Sidewalk Labs proposes that an independent entity called the Urban Data Trust manage urban data and make it publicly accessible by default (if properly de-identified).

As described on Page 420, part of this entity's responsibilities would involve establishing an accountable and transparent process for approving the use or collection of urban data in the first place, given the potential of urban data to impact people's daily lives.

3

What we heard:
Apply consistent guidelines.

A third major theme emphasized by public consultation was that Sidewalk Labs should not have a special advantage in the development of urban innovations. Quayside and the IDEA District must welcome all kinds of local companies, entrepreneurs, researchers, and civic organizations using urban data to improve life.

How we responded:
A single process for all parties.

The process proposed applies to all entities that seek to collect urban data in the IDEA District, including Sidewalk Labs.

The result: A proposed process for using urban data managed by an independent entity

These insights formed the basis of Sidewalk Labs' proposal for responsible data use, which builds on the strong foundation established by privacy laws and aims to establish an enhanced privacy standard.

Provincial and federal privacy commissioners would continue to oversee compliance with all privacy laws. Additionally, this proposal calls for the establishment of an independent Urban Data Trust, tasked first with establishing a set of RDU Guidelines that would apply to all entities seeking to collect or use urban data in the IDEA District and, second, with implementing and managing a four-step process for approving the responsible collection and use of urban data:

① ② ③ ④ Step 1:
Classify the data.
Does the proposed data activity involve urban data, and if so, does it involve personal information?

① ② ③ ④ Step 2:
Submit an RDU Assessment.
How would the data be used and collected? What measures, such as consent or de-identification, would be taken to ensure privacy and avoid harm?

① ② ③ ④ Step 3:
Receive a decision.
Do the benefits outweigh the risks enough to merit approval by the Urban Data Trust?

① ② ③ ④ Step 4:
Meet post-approval conditions.
Have devices been registered? How would access be facilitated? How would auditing occur?

The following sections describe the proposed implementation of the Urban Data Trust in greater detail, propose initial RDU Guidelines for consideration, and describe each of the proposed steps required when applying to use or collect urban data. This description is followed by two examples of how the process could work for digital innovations.

(This particular proposal is just one of many that should be considered on this important topic. Sidewalk Labs also supports the consideration of other recent proposals, including from MaRS[38] and the Toronto Region Board of Trade,[39] calling for independent entities whose mandate could be to govern data collection and use, provide oversight of digital technologies, enhance radical transparency for the placement of sensors in the public realm, and encourage that standards are published to enable third-party innovation.)

Implement the Urban Data Trust

Goal 1

Key Term

An independent

Urban Data Trust

would oversee all requests to use or collect urban data.

Sidewalk Labs proposes that the Urban Data Trust oversee matters of the digital governance of urban data for the IDEA District, including the approval and management of data collection devices placed in the public realm, as well as addressing the challenges and opportunities arising from data use, particularly those involving algorithmic decision-making. (Note that this entity is not intended to be a "trust" in the legal sense; see sidebar on Page 423.)

Sidewalk Labs believes the Urban Data Trust should be managed through a democratic process, but also recognizes that the novelty, complexity, and scale of this approach means that it could take some time to figure out how to appropriately implement the entity. For these reasons, Sidewalk Labs proposes that the Urban Data Trust could be implemented in two phases.

A first phase would be focused on getting the entity up and running quickly to establish the rules and give it experience working through use cases, perhaps first working through Sidewalk Labs' proposed use cases in Quayside; a second phase would work towards a more long-term solution.

Initial implementation period

Sidewalk Labs proposes that initially the Urban Data Trust be implemented through the final agreement between Waterfront Toronto and Sidewalk Labs. The agreement would call for the creation of the Urban Data Trust as the independent digital governing entity for the Sidewalk Toronto project (not controlled by either Sidewalk Labs or Waterfront Toronto). A key component of the agreement would require any organization requiring a permit to build or operate in the IDEA District to consider whether they plan to engage in data-gathering activities. If those activities would involve the collection or use of urban data, the agreement would require that the organization apply to the Urban Data Trust and obtain its approval before urban data collection and use could occur.

The agreement would also set up the structure of this initial Urban Data Trust and authorize that a non-profit entity be created with the charter to address the digital governance challenges related to urban data while also promoting data-driven innovations that benefit individuals and society. Sidewalk Labs proposes that this entity would have a board consisting of five members. The board initially could include a data governance, privacy, or intellectual property expert; a community representative; a public-sector representative; an academic representative; and a Canadian business industry representative.

The board could act in ways similar to Internal Review Boards or Research Ethics Boards in academic institutions for research, or to content moderation boards set up in-house at social media companies. In these examples, a team of experts are assembled to review and assess whether certain decisions should be made while balancing different interests. The independence of the board would be ensured by the application of best practices such as diverse representation of interests, term limits, staggering term lengths to ensure balanced succession, maintaining appropriate boundaries with clear conflict of interest policies, and other measures.

The proposed board would also hire (as an employee of the Urban Data Trust) a Chief Data Officer to run the entity's daily operations. This position could be filled by a data governance and privacy expert, potentially similar to the type of experience a former privacy commissioner might have.

Under the direction of the board and requiring its approval, the Chief Data Officer would be responsible for developing the charter for the Urban Data Trust; promulgating RDU Guidelines that apply to all parties proposing to collect urban data, and that respect existing privacy laws and guidelines but also seek to apply additional guidelines for addressing the unique aspects of urban data (see Page 424); structuring oversight and review processes; determining how the entity would be staffed, operated, and funded; developing initial agreements that would govern the use and sharing of urban data; and coordinating with privacy regulators and other key stakeholders, as necessary.

Sidewalk Labs anticipates that the Chief Data Officer would use a number of resources to inform its decisions, including the RDU Guidelines, the RDU Assessments (see Page 426) completed by proposed data collectors, published guidance from privacy regulators, and input from the board. The Chief Data Officer's decisions would be made to ensure that all actors in the IDEA District comply with applicable laws, such as PIPEDA and provincial or municipal privacy laws. The Chief Data Officer and the board would also develop protocols on when and how data could be stored outside of Canada.

Urban data agreements.
During the initial implementation period, the Urban Data Trust entity would enter into contracts with all entities, institutions, and organizations that are approved to collect or use urban data in the IDEA District. The contracts ("urban data agreements") could be similar to data sharing agreements or data licence agreements and include parameters that govern the collection, disclosure, storage, security, analysis, use, and destruction of urban data. Since these terms would be stipulated in the contracts, the breach of any term would be legally enforceable, with breaches actionable in court by the Urban Data Trust entity. The Urban Data Trust could also publish breach notifications about data collectors who fail to comply with the contract, and the contracts could potentially provide the entity with the right to enter onto property and remove sensors and other recording devices if breaches are identified.

Funding.

While the details on funding the initial implementation of the Urban Data Trust would need to be worked out in a consultation process, Sidewalk Labs proposes that as part of each contract, each party that desires to collect and use data in the designated geography pay a data collection and use administration fee to cover the costs of the Urban Data Trust. These costs would include salaries for the Chief Data Officer and the staff to manage applications, reviews, audits, and enforcement, as well as honoraria and other customary expenses for the board.

Longer-term options

After a certain period of time — once the Urban Data Trust has overseen the collection and use of data in the IDEA District and has gone through multiple use cases with provincial and federal privacy regulators — it is possible that other, more enduring arrangements could be implemented.

Looking long-term, Sidewalk Labs puts forth that the Urban Data Trust could be transformed into a public-sector agency or a quasi-public agency, either of which could give it more long-term viability or broader coverage.

Public-sector agencies receive their mandate from enabling legislation, are responsible for performing a public function or service, and are accountable to the minister responsible for that legislation. An advantage of transforming the Urban Data Trust into a public-sector agency is that the concept and process could then be applied to a wider group of organizations and places where similar technologies are being deployed. A disadvantage is that housing the Urban Data Trust in a public-sector entity would require new or amended legislation, and the passage of legislation can take time and would need to account for emerging technologies.

Sidewalk Labs notes that the Toronto Region Board of Trade recommended that the Toronto Public Library (a public-sector agency) be charged with the responsibility and authority for a Toronto Data Hub, citing the library's expertise in managing data and its credibility and trustworthiness to put the public interest first.[40] Sidewalk Labs supports a further review of this proposal.

Quasi-public bodies include entities that have been granted authority to act in the public interest, but that are at arm's length from government. For example, in Ontario, certain professions are governed by self-regulatory colleges, which regulate those professions in the public interest.[41] These colleges are responsible for ensuring that their regulated professionals act in a safe, professional, and ethical manner. They have the power to set practice and competency standards, investigate complaints about members, and, where appropriate, discipline members. The advantages of a quasi-public body include that it can act independently of government and that its reason for existence is to protect the public interest. A disadvantage is that these agencies are usually publicly funded until they can be fully self-funded.

Sidewalk Labs believes each of these options to be credible and worthy of further discussion in collaboration with Waterfront Toronto's Digital Strategy Advisory Panel, government, the community, academia, and industry.

Why the "Civic Data Trust" became the "Urban Data Trust"

One of Sidewalk Labs' initial proposals for responsible data use called for an independent Civic Data Trust to be the steward of urban data.[42] Sidewalk Labs heard consistent feedback from many advisors and critics who felt that calling this entity a "trust" raised questions such as: "Who would be the trustee, and who are the beneficiaries?"

Sidewalk Labs notes that this entity is not intended to be a "trust" in the legal sense — legal trusts are not designed to benefit the general public. Instead, Sidewalk Labs aligns with the definition of a data trust from the Open Data Institute, a U.K. non-profit, as "a legal structure that provides for independent stewardship of data," as articulated in the institute's 2019 report, "Data trusts: lessons from three pilots."

While Sidewalk Labs proposes a non-profit entity, the final legal structure (and name) would be determined based on input from government, the community, researchers, and industry. Sidewalk Labs also now calls this entity the "Urban Data Trust" to clarify the proposed responsibilities.

Additionally, Sidewalk Labs heard that some people prefer to use the term "digital" rather than "data," as the considerations of an entity like the trust extend beyond data to all digital matters. Sidewalk Labs agrees and believes that the proposed RDU Guidelines and Assessment embrace this concept by assessing the broader issues arising from digital innovations and data ethics.

Creating a Trusted Process
for Responsible Data Use

Establish RDU Guidelines

Sidewalk Labs believes that an essential early step for the Chief Data Officer would be to create a set of RDU Guidelines that establish clear, common standards for responsible data use and can be applied consistently to all parties engaged in the collection and use of urban data.

The RDU Guidelines should address the concerns around privacy and data ownership that have been raised about the Sidewalk Toronto project, recognizing that similar concerns apply to other entities engaging in similar work. Rather than being constrictive, these rules should provide greater clarity and transparency to all innovators who want to set up shop and use data in a responsible way.

Sidewalk Labs believes the RDU Guidelines should build on the world-renowned approach to privacy called Privacy by Design, which outlines principles that should be implemented from the very beginning of a data activity to embed privacy protections into the design, operation, and management of a product, project, operation, or service.[43] But the proposed RDU Guidelines should go beyond privacy to address key areas of digital governance, ethics, and open access to information, as well as the ways in which aggregate or de-identified data can impact individuals and groups of people through the use of advanced analytics, such as artificial intelligence.

Key Term

Privacy by design

is a world-renowned approach to privacy that outlines principles that should be implemented from the very beginning of a data activity.

Sidewalk Labs believes the Urban Data Trust would be in a position to determine the most appropriate RDU Guidelines. For consideration as an initial set, however, Sidewalk Labs submits the following guidelines, which it has implemented internally for pilots that undergo privacy assessments:

Beneficial purpose.

All proposed uses of urban data must incorporate Canadian values of diversity, inclusion, and privacy as a fundamental human right. To meet this standard, there must be a clear purpose and value to any proposed use of urban data, as well as a clear, direct connection to the ways in which the project and proposed data collection activity would benefit individuals or the community. A proposal or project should not be collecting data for the sake of having data.

Transparency and clarity.

Organizations should inform individuals of how and why data would be collected and used, and should do so in a way that is proactive, clear, and easy to understand. Organizations should provide examples of how they plan to inform individuals about the data-collection activity.

Data minimization, security, and de-identification by default.

Organizations should collect the minimum amount of data needed to achieve the beneficial purpose and use the least invasive technology available to achieve the beneficial purpose. Organizations should seek to use up-to-date de-identification techniques to reduce the amount of personal information that they collect and use. Organizations should demonstrate the need for the amount of data to be collected and should be prepared to detail what, if any, personal information is desired; what they are planning to do with it; what safety and security safeguards would be used to protect individuals; and how these efforts would be audited.

Publicly accessible by default.

Organizations should make properly de-identified or non-personal data that they have collected publicly accessible to third parties by default, formatted according to open standards. This approach would help to ensure that individual privacy is preserved while also enabling data and source code to be accessible by others to catalyze innovation. Organizations should be prepared to detail their methods for making such data publicly accessible, and to justify any plans to restrict data access.

No selling or advertising without explicit consent.

While there would not be proposed prohibitions placed on data collectors who would like to sell data containing personal information or to use such data for advertising, a higher level of scrutiny should be placed on projects that want to use personal information for these purposes. Organizations that want to engage in this activity have an obligation to follow all applicable privacy laws; they should also provide clear justifications for this activity and demonstrate (with examples) how they plan to obtain explicit consent from the affected individuals. Such precautions are necessary because individuals often do not know when their personal information is being sold or used for such purposes.

(Sidewalk Labs has already committed publicly that it would not sell personal information to third parties or use it for advertising purposes. It also commits to not share personal information with third parties, including other Alphabet companies, without explicit consent.)

Responsible AI principles required.

To ensure that issues around the use of artificial intelligence systems are being considered and addressed by data collectors and developers, organizations should be required to detail if they are going to be developing AI systems. If so, they should be required to show how they have incorporated Responsible AI principles into their development and decision-making to reduce the likelihood of biased and unethical outcomes. (See Page 411 for more information.)

Goal 3

Set a clear process for urban data use or collection

Sidewalk Labs proposes that once the Urban Data Trust and RDU Guidelines have been established, a transparent, four-step process should be created for any proposals seeking to collect or use urban data in the IDEA District.

① ② Step 1:
③ ④ Classify the data

Step 1 would involve the person or entity determining whether or not its proposal involves urban data, transaction data, or both types.

Urban data.
If the data activity involves the collection or use of urban data, then Sidewalk Labs proposes that the data collector must move on to Step 2 of the process, which calls for submitting an RDU Assessment to the Urban Data Trust (see sidebar on Page 428).

Urban data can include information collected in the public realm — defined as commonly shared spaces not owned by a private entity, such as streets, squares, plazas, parks, and open spaces — by devices such as pedestrian counters or traffic cameras. It can include information collected in privately owned but publicly accessible spaces, such as building lobbies, courtyards, some parks, ground-floor markets, and retail stores. And it can include information collected by a third party in private spaces, such as data on tenant or building noise, air quality, and energy use.

Transaction data.
If the data activity solely involves the collection and use of transaction data, then no assessment is required.

Transaction data is information that individuals consent to providing for commercial or government-operated services through a direct interaction, such as apps, websites, and product or service delivery. This data includes things like the credit card information a customer provides when signing up for a home delivery, an email address given to sign up for a local business's e-newsletter, or a phone number submitted to a banking app for text updates.

Sidewalk Labs believes that transaction data should not be under the Urban Data Trust's purview for several reasons. First, the data collector is already accountable under applicable privacy laws either to obtain consent to the collection and use of such data if the data is personal information or, if it is a public-sector entity, to ensure they have the proper legislated authority. Second, this type of data arguably is not uniquely connected to public spaces, nor is it generally considered a public asset requiring additional protections within the public interest.

This proposal to remove transaction data from the purview of the Urban Data Trust does not dismiss any ongoing concerns or questions that people have about the collection and use of transaction data in the areas of consent, transparency, and accountability, among others. Instead, it reflects the belief that incorporating transaction data into a governance model for the Sidewalk Toronto project would be unworkable given the lack of a relationship between this kind of data collection and a specific geography.

Sidewalk Labs appreciates that there would be ongoing dialogue about the scope of data collection and use under the Urban Data Trust's purview, and welcomes that dialogue.

(Even though this proposal does not place transaction data under the purview of the Urban Data Trust, Sidewalk Labs commits to applying the RDU Guidelines to any of its own commercially launched products and services that involve transaction data.)

Both types of data.
If the data activity involves the collection and use of both types of data, such activity would fall under the stewardship of the Urban Data Trust. One realistic example is an app-based ride-hail service whose vehicles are equipped with sensors or cameras capable of collecting data on passengers or the environment. While this organization's collection and use of data through the app would not fall under the jurisdiction of the Urban Data Trust, its collection and use of urban data through sensors and cameras would fall under that jurisdiction, thus requiring an RDU Assessment to be filed.

Is it urban data?

The following questions can be used by public- or private-sector entities to ascertain whether the data they want to collect and use is subject to the Urban Data Trust process.

1 Is the data solely transaction data?

↓ **No** — Please continue to Question 2.

⊗ **Yes** — This is not urban data. This is a traditional form of data that Sidewalk Labs calls "transaction data," which does not fall under the stewardship of the Urban Data Trust.

2 Is the data proposed to be collected within the IDEA District?

↓ **Yes** — Please continue to Questions 3–5. If you answer yes to any of these questions, then the data is urban data and subject to the stewardship of the Urban Data Trust.

⊗ **No** — Your data will not be governed by the Urban Data Trust.

3 Is the data proposed to be collected in the public realm — on the street, in public squares, at plazas, in parks, or in open spaces?

✓ **Yes** — Urban data.

↓ **No** — Continue to next question.

4 Is the data proposed to be collected in privately owned spaces commonly used or accessed by the public — including building lobbies, privately owned but publicly operated parks, ground-floor markets, retail stores, or ride-hail vehicles?

✓ **Yes** — Urban data.

↓ **No** — Continue to next question.

5 Is the data proposed to be collected by a third party in an individual's private spaces or about an individual in their private spaces? (Examples include a building owner collecting noise, air quality, or energy-usage data on a tenant; a utility collecting data on a tenant's water consumption; or a building collecting information on tenant waste.)

✓ **Yes** — Urban data.

⊗ **No.** → This data is not urban data and will not be overseen by the Urban Data Trust.

① **②** Step 2:
③ ④

Submit an RDU Assessment

As a second step in the process, Sidewalk Labs proposes that entities, both public and private, seeking to collect or use urban data complete an RDU Assessment — an in-depth review outlining the purpose of the digital proposal, the type of urban data it aims to collect, its potential impact on the community, and its risks and benefits. This step would also apply to entities proposing to use urban data collected by an existing device for a new purpose. RDU Assessments would be conducted during the design phase, prior to urban data collection or use.

(Sidewalk Labs has been developing an RDU Assessment template since the summer of 2018, and it is currently used internally to assess the privacy compliance and responsible data use of pilots, projects, services, and products. This process requires collaboration from different teams to ensure that privacy is not just a compliance exercise and that privacy is truly done "by design" from the start.)

The entity applying for data collection would submit the RDU Assessment along with an application to the Urban Data Trust for review and approval. The Urban Data Trust would use the RDU Assessment to assess how the proposal conforms to the RDU Guidelines, privacy laws, Privacy by Design principles, and any other relevant factors or applicable laws. If necessary, the Urban Data Trust should help startups, companies, and organizations understand these factors when preparing the RDU Assessment.

The RDU Assessment would incorporate and build on one of the strongest existing data governance tools for protecting individual privacy: the "privacy impact assessment." A privacy impact assessment identifies any privacy and security risks associated with new digital technologies or data-related services, as well as how they are mitigated in the design of the project. All three orders of government currently require or encourage privacy impact assessments. Similar assessments are also a cornerstone of the General Data Protection Regulation, Europe's 2018 privacy initiative, which has raised the bar on responsible data use.

The proposed RDU Assessment would follow the same guidelines as a privacy impact assessment, attempting to identify potential privacy risks of new programs or services, to begin such an analysis at the outset of development, and to be adjusted and refined through stakeholder feedback. The RDU Assessment would exceed current privacy compliance requirements because it would consider the broader social and ethical considerations of new and existing technologies and their potential impact on people.

How the RDU Assessment relates to the RDU Guidelines

When assessing whether to approve a digital proposal, the Urban Data Trust would review an RDU Assessment and consider many factors, including how well the proposal conforms to the RDU Guidelines. Many of the example questions on this page have a close tie back to the guidelines.

Beneficial purpose

Transparency and clarity

Data minimization, security, and de-identification by default

Publicly accessible by default

No selling or advertising without explicit consent

Responsible AI principles required

Sidewalk Labs' proposed RDU Assessment includes four primary components:

Purpose.
The first section of the RDU Assessment would ask for a description of the purpose of the project, service, or product, including its objectives and goals, as well as the urban challenges it hopes to address. Examples of questions that might be asked in this part of the RDU Assessment might include:

What is the objective for this project? Clearly state the problem that is being solved.

Clearly state the measurable goal or outcome of the project.

How likely are the proposed technology and collection and use of data to solve the problem as described?

What are the alternatives to the technology or method of collection? Why are they not sufficient?

Data sources.
The second section of the RDU Assessment would require a description of the technology or data-collection methods, the data sources or types, and the parties who have access to the data. Some of the questions asked in this section might include:

What are all the sources of the data, internal and external?

Does the data activity involve personal information?

Does this project involve the collection or use of data about people?

Is the data stored in Canada? If not, is there a reason beyond business case or financial considerations that the data would not be stored in Canada?

Is the data, or a subset of data, going to be used for advertising purposes?

Is the data going to be sold to third parties?

Will the data be matched against, combined with, or augmented by other data sets?

Legal compliance.
The third section of the RDU Assessment would capture conformance to applicable privacy laws. Examples of questions asked in this section might include:

Have individuals been given choices about the collection of their personal information?

Describe how the data activity complies with applicable privacy laws.

If the data activity involves personal information, there must be explicit, express consent for collections, uses, or disclosures that: (i) involve sensitive information; (ii) are outside the reasonable expectations of the individual; and/or (iii) create a meaningful residual risk of significant harm. Please explain how you have achieved this requirement.

Does the data activity include mechanisms that explain how data is used, how benefits and risks to individuals are associated with the processing, and how individuals may participate and object where appropriate?

If the data activity includes personal information, how has it been de-identified?

Is there a less privacy-invasive way to achieve the goals of the data activity (including potential insights)?

What are the safety and security safeguards (such as encryption or internal access controls)? Is internal access audited?

Risk-benefit analysis.
The fourth section of the RDU Assessment would ask the proposing entity to detail and rate the risks and benefits associated with the project and data collection activity, and how any risks have been mitigated. Example questions might include:

Could the anticipated use of technology harm or benefit certain individuals, groups of people, or communities in unintended or unexpected ways?

What are the benefits to the individual or groups of individuals?

How will this data-collection activity impact the community?

Will the de-identified or non-personal data be made publicly accessible? If not, why?

If personal data is being de-identified, when in its lifecycle is this done? How long is identifiable data retained on devices?

Explain your external threat model and countermeasures.

What format will the data be made available in? Is this format a public standard? If there is no relevant standard currently available, where is the documentation for the format that you will use? What partners or standards bodies do you plan to work with to promulgate this format?

In this project, is the project owner using analytics-driven models, insights, or algorithmic decision-making that could impact individuals?

Receive a decision

Once the RDU Assessment is completed, the proposed data collector would submit it to the Urban Data Trust for review, assessment, and decision by the Chief Data Officer.

Balance benefits and risks.
Sidewalk Labs proposes that the Chief Data Officer look at all of the information the data collector provided in the RDU Assessment and determine whether the data activity should proceed based on the organization's attestation of compliance with applicable laws, as well as a subjective and objective assessment of the RDU Assessment that takes into account the appropriateness of the proposed data collection and uses and the resulting net balance of impact.

The Urban Data Trust would assess the balance of the proposed benefits and the potential harms, weighing their significance and likelihood of occurring against any mitigation efforts. The entity could also make use of published guidelines from the privacy commissioners regarding personal information; for example, if a data collector indicates that it plans to receive consent for the collection of personal information, the Urban Data Trust could look to the Office of the Privacy Commissioner of Canada's guidelines on meaningful consent to determine how closely they align with the data collector's proposed methods.

Similarly, if the data collector indicated that it plans to de-identify the data, the Urban Data Trust could look at the Information and Privacy Commissioner of Ontario's guidelines on de-identification for structured data, among other industry standards, to assess the techniques used by the data collector, as well as any standards established by the entity.

The Urban Data Trust could also interact with the data collector in a consultative process to the extent that additional information is needed to make the assessment or to assist the data collector in improving its data activity.

Final decision.
Sidewalk Labs proposes that a final decision be issued as "denied," "approved," or "approved with conditions."

Because the RDU Assessment is highly contextual and does not lend itself to black-and-white rules, several case studies have been included on Pages 436-440 to help readers understand how approval decisions could work in practice. Ultimately, the decision-making standards would be set by the Urban Data Trust.

A note on legal compliance.
An organization's approach to legal compliance would be part of the Urban Data Trust's decision-making process, but the organization itself would ultimately be responsible for legal compliance. Failure to abide by relevant privacy laws could result in enforcement action by the appropriate regulator and legal remedies imposed by the Urban Data Trust.

Of note: if personal information (as defined by PIPEDA) is involved in a proposal, the "legal compliance" section of the RDU Assessment would collect information detailing how the data is in compliance with privacy laws. The Urban Data Trust would not assess whether the organization is in compliance with Canadian laws, because under PIPEDA, organizations must remain accountable for the personal information they collect, use, and disclose. There are also practical reasons involving accountability and liability that account for why the Urban Data Trust should not be responsible for this compliance.

The Urban Data Trust could deny applications based on overt or apparent non-compliance. But the Urban Data Trust's opinion on legal compliance — for example, through the acceptance or rejection of an RDU Assessment based on PIPEDA compliance — should not be taken as validating compliance or as evidence or a ruling on legal compliance.

Meet post-approval conditions

As a final step in the process, Sidewalk Labs proposes that, once an entity or organization receives approval to collect or use urban data in the IDEA District, the Urban Data Trust should meet a set of post-approval conditions around transparency, device registration, data access, data sharing and licencing agreements, and auditing.

RDU Assessment transparency.
Sidewalk Labs proposes that the summaries of approved RDU Assessments be made publicly available by the Urban Data Trust to ensure transparency and encourage accountability by the public, privacy advocates, and regulators alike. Proprietary or confidential information, such as intellectual property or trade secrets, would not be published.

Device registry.
Sidewalk Labs proposes that, as part of the RDU Assessment filing and application process, entities must submit a map with the proposed locations of all data-collection devices, such as sensors or cameras. (This requirement would not apply to private owners or tenants of residential units or houses, such as those installing home security cameras for personal safety reasons.) Once the application including these locations has been approved, the entity must register these devices with the Urban Data Trust, which would upload the devices' locations and fields of view to an interactive map that would be publicly accessible. This registry would provide the public with a real-time inventory of information on what kind of data is being collected, as well as why, how, where, and by whom.

Facilitating access.

Sidewalk Labs believes that, in line with its proposed RDU Guidelines, properly de-identified, aggregate, or non-personal urban data should be made publicly accessible by default. Public access to urban data is crucial to innovation, equity, and the provision of digital services that improve quality of life.

If the data or source code were to be made publicly available, the Urban Data Trust would manage this access through data sharing agreements and facilitate integration with existing open-data portals and tools.

Facilitating access could be accomplished in a variety of ways, from having the Urban Data Trust actually hold the data to having it set rules that require collectors to publish de-identified, aggregate, or non-personal data in real time. This access should be free for basic use, but reasonable fees could be applied for commercial purposes or heavy use.

Access restrictions.

Data sharing agreements would also include information about any access restrictions approved by the Urban Data Trust. There could be cases when urban data cannot be released publicly for a variety of reasons. These cases could involve data that contains personal information — for example, a government organization that collects transponder data or images of licence plate numbers for enforcement.

Other cases could involve proprietary data collected at great cost to a company. The public release of such data would undermine investment and competitive advantage, discouraging businesses from locating within the IDEA District.

For example, consider a company building an alternative robotic delivery system for transporting packages and items to and from a storage facility. For robots to be able to navigate tunnels, sidewalks, building entrances, lobbies, elevators, and hallways, they would need to know where they are at any given moment with a high level of precision. Existing positioning technology like GPS or Wi-Fi triangulation would be too coarse — especially in urban environments, where GPS signals are often blocked by buildings. Recent developments in positioning technology can provide accuracy within a few millimetres, but significant investment would be required to deploy transmitters throughout the neighbourhood.

While this type of location data would technically occur within the public realm, the considerable cost of compiling it — and the likelihood that the company would either choose to pursue the project elsewhere, or not at all, if forced to make the data available, in real time, to its competitors — could merit a proprietary restriction in the view of the Urban Data Trust. The entity would still be able to audit the data collection and use, and the RDU Assessment summary would be publicly accessible.

Data sharing and licencing agreements.

As described on Page 421, Sidewalk Labs proposes that the Urban Data Trust facilitate access to urban data via data sharing agreements, including the terms of any potential restrictions or licencing fees.

In these cases, the Urban Data Trust would first make a determination about whether or not access to the data should be restricted, and then negotiate the terms of this restriction with the company or entity. These terms might include making the data accessible through an agreed-upon licencing fee, endowing the Urban Data Trust with rights to facilitate access based on certain specifications, requiring permission from the original entity for another party to access the data, or potentially even prohibiting access.

From that point forward, any entity seeking access to this data would have to apply for approval through an RDU Assessment, agreeing to abide by the negotiated access or licencing terms.

Data sharing agreements would also include a copy of the RDU Assessment and application, fees payable to the Urban Data Trust, the rationale for retaining any data in an identifiable manner, details on how the organization or entity would be audited, details on any certification marks the organization has obtained for its practices or project, and a limitation of liability and indemnification to the Urban Data Trust.

Auditing and enforcement.

The Urban Data Trust should retain the authority to audit all collections and uses as needed and order the removal of digital devices in the event it discovers a violation. The terms of auditing would depend on factors such as the sensitivity of the data, the track record of the organization, and the uses of the data, including whether advanced data analytics would be run on the data and whether the organization plans to use the data for ads based on consent obtained.

The Urban Data Trust would be able to seek legal remedies for violation of agreed-to conditions of data collection and data use.

The question of more traditional enforcement authority should be considered as part of the ongoing consultation for this work — for example, auditing could occur with the assistance of privacy regulators or via contractual agreements.

How it works: RDU Assessment case studies

It can be hard to talk about digital governance in the abstract. While the proposed Urban Data Trust would ultimately create its own governance standards and guidelines, the following illustrative examples are presented here to help guide readers through the responsible data use process and to give a broad sense of how decisions around responsible data use could be made. The process described here would apply to any public or private entity proposing to collect or use urban data in the IDEA District, including Sidewalk Labs.

1

Example #1: A mobility management system

A private company proposes to launch a mobility management system, working in collaboration with the city's transportation department.

The proposed mobility management system could help coordinate all the roads, traffic signals, curbside loading zones, and trip options, ensuring a safe and efficient travel experience for residents, workers, and visitors. To work properly, such a system would need to collect real-time information on mobility-related measures such as traffic volume (for pedestrians, cyclists, transit riders, and cars alike), transit delays, curb demand, parking demand, route closures, emergency dispatches, weather patterns, and more. This information would help the system do things like set prices for pick-up and drop-off zones to reduce congestion, or hold traffic signals for pedestrians who need more time to cross the street.

① ② ③ ④ **Step 1:**
Classify the data

The proposed mobility management system would operate in Quayside. It would require the placement of sensors and devices in public spaces, including on traffic signals, such that individuals would not have the practical opportunity to provide prior meaningful consent for the collection and use of this data.

For these reasons, the data collected would be considered "urban data." The proposal should advance to Step 2.

① ② ③ ④ **Step 2:**
Submit an RDU Assessment

Because the mobility management system seeks to collect and use urban data, it must complete an RDU Assessment. This assessment, plus an application, must be filed with the Urban Data Trust and approved before the service can launch.

The RDU Assessment would help the Urban Data Trust assess how well the proposed mobility management system conforms to relevant decision factors, such as the RDU Guidelines, applicable privacy laws, and Privacy by Design principles. Some of the relevant details from the assessment could include:

> The proposed system has a clear beneficial purpose, with an aim toward improving public safety, traffic congestion, and travel times.

> Much of the data required to run the system is non-personal, such as sensors to detect available curb spaces. The system also uses de-identified data by computing aggregate counts of pedestrians, cyclists, and vehicles directly on the camera and immediately deleting any raw video footage, safeguarding the privacy of individuals who might be visible in the raw footage. Together these efforts reflect Privacy by Design principles and data minimization.

> The city also proposes to collect some personal information (such as transponder information or licence plate images) for enforcement of curb rules; the city would attest to compliance with the applicable laws, including the Municipal Freedom of Information and Protection of Privacy Act.

The information collected by the system would not be sold for advertising purposes or used for behavioural tracking purposes.

While direct consent would not be possible for traffic signal information, the system would submit a map with the proposed placement of all mobility-related sensors to the Urban Data Trust so people could know the locations and purposes of the devices, improving transparency.

Non-personal data would be made publicly accessible to others. Some access to de-identified data is proposed to be restricted as the system trains and tests its algorithm, to safeguard privacy and security.

The system's cameras would use computer vision to de-identify pedestrians, cyclists, and vehicles at the source. Some de-identified information would be kept for an indefinite period to help train the algorithm to properly de-identify images. The data would only be accessible by key personnel with valid reasons to access the data for quality assurance and security purposes. Because data would be used by an algorithm and to influence decisions, Responsible AI guidelines should be considered in the assessment of this technology and proposed data use.

① ② **Step 3:**
③ ④ # Receive a decision

As a next step, the Urban Data Trust would review the RDU Assessment and the application. Again, the Urban Data Trust should establish its own decision-making guidelines, but based on the proposed RDU Guidelines, this particular proposal would seem to meet criteria for approval, given the balance of benefits to risks.

Benefits: The system proposes to help achieve a reduction in traffic congestion, an increase in public transit ridership, and reductions in carbon emissions related to driving. The resulting accessibility of aggregate, non-personal, and de-identified data made publicly available would ease traffic and provide new opportunities to develop safety devices and applications. The data controllers would plan to store data in Canada.

Risks: The personal information collected as part of the system could be used to identify location patterns and schedules, including access by law enforcement and civil discovery. Other risks could include the de-identification process and the retention period of some of the images for calibration.

Decision: Given the proposed RDU Guidelines, the Urban Data Trust would likely approve this data activity, given its clear benefits and its proposals to effectively manage risks, which would include using the minimum amount of data, de-identifying data at the source, and ensuring any personal information collected by the city is secured and encrypted. The data controllers would also attest that the data activities are in conformance with applicable privacy laws.

① ② **Step 4:**
③ ④ # Meet post-approval conditions

Once approved, the data collectors would register the data-collection devices to the publicly accessible device registry. The data collectors would still work with the Urban Data Trust to meet post-approval conditions around transparency, data access, and auditing.

Transparency: The summary RDU Assessment would be made publicly available.

Device registration: All devices would be registered with the Urban Data Trust and placed on a publicly accessible map.

Data access: Non-personal and aggregate data is made publicly accessible via the city's open-data portal. For example, a researcher could study this data to detect near misses between cars and pedestrians, and evaluate the performance of intersection designs on street safety.

Data sharing agreements: While access to properly de-identified data would be restricted to train the algorithm, the Urban Data Trust recommends that once testing is complete, the data and source code be made open so the benefits can spread. For example, a self-driving technology startup could use the same type of insights to create an improved pedestrian detection system. Personal information that would be collected and used by the city would not be made publicly accessible.

Auditing: The Urban Data Trust could decide that it would audit the system's de-identification techniques once in the next year. The Urban Data Trust could also recommend that the company retain an external auditing company to assess its de-identification techniques.

2

Example #2: An automated parking payment system

A private parking garage owner proposes to install CCTV cameras for security purposes, and to use the data to create an automated payment system as drivers enter and leave the garage. The cameras are capable of reading licence plates and capturing images of drivers and passengers. The garage owner does not plan to de-identify these images. The garage owner also plans to share the data with a data broker for a fee.

Individuals who are regular users of the parking garage could opt in to this system for automatic payment. Individuals who use the garage as one-offs and who do not opt in to (or even know about) this service would also have their licence plates captured, although these customers must pay for parking using a parking app or with cash.

① ② **Step 1:**
③ ④ # Classify the data

The proposed parking payment system would operate within the IDEA District. The placement of cameras would be in a privately owned public space, and individuals would not have the opportunity to provide explicit consent for the collection and use of their data. Additionally, the payment system would be linked to an individual's credit card or parking app account.

For these reasons, the data collected would be considered "urban data" as well as "transaction data," and the proposal should advance to Step 2.

① ② **Step 2:**
③ ④ # Submit an RDU Assessment

Because the proposal seeks to collect and use urban data, the parking garage owner must file an RDU Assessment and an application with the Urban Data Trust for approval before the service can launch.

For this illustrative example, some of the relevant details from the assessment could include:

The garage owner claims a beneficial purpose for the proposal related to security and automated billing for customers. The garage owner would like to sell the data to a data broker, claiming this would benefit customers by offsetting fees to help keep parking prices low. However, selling data to third parties without explicit consent from the individual is in violation of RDU Guidelines.

The garage owner intends to provide notice of the cameras with "CCTV signs" posted around the garage, achieving some transparency. There would also be information printed on the back of the parking garage ticket on how the data is used and directing the user to the garage website, where a more complete description of the data practice would be available.

The garage owner attests compliance with PIPEDA and any other applicable law on the application form accompanying the RDU Assessment.

The video stream would be available to the parking lot attendant when in the office and would be kept in the case of an incident and subsequent examination by authorities for a period of two weeks. Because the purpose for data collection is to deter or investigate safety and security incidents, there would be no obligation to de-identify the footage, and this use would be permissible by Canadian laws, as long as the Office of the Privacy Commissioner's guidance on video surveillance is followed. But the parking garage owner also proposes to use the video footage for another purpose (selling to data brokers) without obtaining consent and would not de-identify this data.

While the parking garage owner acknowledges that sharing personal information with a data broker would likely be surprising to individuals, the owner does not detail any risk mitigation efforts, claiming that the risks would be necessary and justified by the benefits.

① ② **Step 3:**
③ ❹ ## Receive a decision
As a next step, the Urban Data Trust would review the RDU Assessment and the application. Once again, the entity should establish its own decision-making guidelines, but based on the proposed guidelines, this particular proposal would likely be denied, given that its risks outweigh its benefits and that the data activity does not comply with RDU Guidelines.

Reasons: The data activity, as a whole, would stand in violation of the RDU Guidelines by selling data for advertising purposes or to third parties without consent and not de-identifying the data used for this purpose by default. The rationale for not de-identifying by default would likely not be compelling, as there were no actions taken to mitigate the risk.

The Chief Data Officer would likely consider the data activity, as a whole, in violation of PIPEDA, as the garage owner did not specify in the legal compliance law section of the RDU Assessment that they had obtained consent from the vehicles' owners, and also proposes to sell personal information without consent.

Conditions: The garage owner would have the opportunity to resubmit the RDU Assessment and application after consultation with the Urban Data Trust. Unless and until the RDU Assessment and application gains approval, the garage owner would not be able to install the CCTV cameras and begin collecting data. If an audit discovered that CCTV cameras had been placed in the garage and had started to collect data, the garage owner could be sued for breach of the contract entered into upon leasing the garage in the IDEA District.

① ② **Step 4:**
③ ❹ ## Meet post-approval conditions
In this case, failure to gain approval would mean the proposal would not advance to Step 4.

The Urban Data Trust would help ensure privacy protections, make urban data a public asset, apply consistent and transparent guidelines, and be publicly accountable to all Torontonians.

Part 4

Launching Core Digital Services That Others Can Build On

Spotlights

1
An outcome-based building code system to enable a safe, vibrant mix of uses

2
An Office Scheduler to optimize energy use

3
A mobility management system to reduce congestion and improve safety

Digital infrastructure, published standards, and a trusted responsible data use process together set the foundation for digital innovation. But a true ecosystem of urban innovation requires a catalyst that makes it possible for third parties to build new digital applications, services, products, or tools that improve people's lives.

To serve as that catalyst, Sidewalk Labs proposes to launch core digital services that are essential to achieving quality-of-life objectives from Day One in Quayside (see table on Page 444). These launch services would not only deliver improvements to affordability, mobility, sustainability, and economic opportunity, but also would make the urban data they generate accessible to others — enabling countless subsequent innovations to emerge from local companies, entrepreneurs, startups, researchers, agencies, civic groups, and others.

These proposed core digital services would have a multiplier effect, since making their non-personal, aggregate, or de-identified urban data publicly acces-

sible would catalyze digital innovations by a wide and growing range of third parties, inspiring a new generation of tools for city living:

The shipping company that uses micro-location data to develop a robot that can deliver packages straight to a person's door

The mobility entrepreneur who uses trip data on shared rides to launch a shuttle service with on-demand beach chairs and umbrellas

The retailer who pairs foot-traffic data with weather information to identify the best locations or times for pop-up vendors to set up shop

The environmental researcher who uses building data to recognize common recycling mistakes and teams up with a digital fabrication studio to design a more sustainable coffee-cup lid piloted by local restaurants

The list is truly endless. Just as no one could have expected that a satellite-positioning system would eventually change the way people hail a cab, ride a bike, order food, meet with friends, take pictures, or even find romance — digital services have the power to enable new ideas no one can imagine.

The following pages provide an overview of several core services proposed by Sidewalk Labs, as well as a description of the urban data they use, an illustrative sense of what their RDU Assessments could emphasize, and the types of third-party innovations that they might make possible.

Merely collecting urban data is not an end to itself. Urban data should only be gathered as a means of creating a new application, use, service, or product that can improve the lives of city residents, workers, visitors, and businesses.

Sidewalk Labs' role in digital services.
As explained on Page 382, Sidewalk Labs plans to offer this limited set of core digital services in cases where achieving fundamental project goals around transportation, affordability, housing, energy, public space, and other areas would require an innovation the market has not pursued.

Some of these launch services could still involve working with partners or buying existing technology, and other entities would be free to develop competing services. All proposed digital services would be subject to the proposed responsible data use approval process overseen by the Urban Data Trust, which would include completing RDU Assessments to ensure privacy is protected.

GRIT Toronto: Involving the community in digital tool development

Traditionally, user testing has taken the form of market research: a small group of people is recruited to come to an office during working hours to give feedback on a new technology. This method can result in narrow or even biased feedback.

To explore a more inclusive kind of user testing, Sidewalk Labs is currently funding GRIT Toronto (Gathering Residents to Improve Technology), a program founded by Code for Canada. The program meets people of all digital skill levels, cultures, ages, and backgrounds where they are — in community spaces outside of working hours, for example — and incorporates their feedback into the creation of new digital services and products, helping to ensure these tools reflect the needs of the populations they are intended to support.

Launched in late 2018, the GRIT Toronto pilot has recruited over 350 residents from Toronto's 25 wards, representing a diversity of backgrounds, lived experiences and technical skill levels. What unites them is a desire to shape the digital products and services that could impact their lives and their city. This initiative could help software developers in Quayside collaborate with a broad range of community members and ensure that their digital solutions truly have neighbourhood needs in mind.

Sidewalk Labs' proposed launch services

This table seeks to provide an overview of the initial digital services proposed by Sidewalk Labs as part of the Sidewalk Toronto project, including a sense of their purpose, data sources and access, and potential to catalyze third-party innovation. All digital innovations (whether created by Sidewalk Labs or others) would be subject to the independent responsible data use approval process described on Page 424, as well as applicable privacy laws. The information here should be viewed as illustrative but not necessarily exhaustive.

Sidewalk Labs' proposed service or application	What urban data it proposes to use and/or publish	Possible third-party applications that could build on this data	What existing ecosystem the innovation supports (Names are illustrative only.)
Mobility management system To reduce congestion and encourage shared trips, this proposed mobility management system would coordinate all travel modes, traffic signals, and street infrastructure, and apply demand-based pricing to curb and parking spaces.	**Non-personal:** Curb space availability (e.g., occupancy sensors) **Non-personal and/or de-identified at the source:** Pedestrian and cyclist detection and counts; vehicle detection, counts, speed **Restricted data (not published for privacy reasons):** Vehicle identification data, such as license plates or transponders, collected and used directly by the city for parking enforcement	A **policymaker** could create more informed policy decisions around parking availability and transit service. A **self-driving technology** startup could improve its pedestrian-detection system. A **researcher** could detect pedestrian near misses and evaluate the performance of intersection designs on street safety. **Employers** could start programs that encourage workers to shift commute times to decrease congestion.	**Self-driving vehicles:** Aptiv, Cruz, Lyft, Uber, Waymo **Sensor and traffic management:** Axilion, Brisk Synergies, GRIDSMART, LeddarTech, Miovision, NoTraffic, Numina, P3Mobility, RapidFlow, SMATS Traffic Solutions **Parking:** Cloudpark, Curbway, Jrop, Passport, Pay by Phone, Sensys **Routing apps:** Apple/Bing/Google Maps, Transit App, Waze
Outdoor comfort system A proposed system of outdoor-comfort tools, deployed in real time, could dramatically increase the amount of time it is comfortable outside, including building "raincoats" to block rain, awnings to provide shade, and fanshells to provide group cover.	**Aggregated and/or non-personal:** Hyper-local temperature, humidity, wind speed, rainfall, and sunshine levels **Non-personal:** Raincoats and fanshell status	A **retail startup** could build an app that identifies the best locations or times for a pop-up store based on weather patterns. **Health organizations** could build apps that show residents a jogging route that avoids wind and snow and maximizes sun and interesting views. (These apps could also draw from the mobility sensors to avoid congested areas.)	**Weather data:** Ambience Data, Earth Networks, IBM, The Climate Corporation **People flow:** Ecocounter, Numina, PeopleFlow
Flexible retail platform (Seed Space) A proposed leasing platform called Seed Space would help small businesses and other retailers book a wide range of ground-floor space sizes, from anchor-tenant spaces to micro stalls, for short- or long-term uses.	**Aggregated and/or de-identified:** Footfall and rate data, aggregated tenant turnover rates **Non-personal:** Space size, availability **Restricted data (not published for privacy reasons):** Leasing, rent, or transactional data collected with clear consent	A **retail startup** could create an app that determines the best times of the year or day for an entrepreneur to set up in the area. (This use could also draw on hyperlocal weather data from the outdoor comfort system.) An **economic development firm** could conduct (or have a startup create an app to conduct) retail industry analyses of neighbourhood turnover rates by size of space. **Business Improvement Areas** could use this data to understand the economic impact of events or policy decisions.	**Location mapping:** InnerSpace, MappedIn **Space mapping:** A Retail Space, Chatter Research, POTLOC **Space availability:** Booqd, Breather, Harbr, PiinPoint
Open space usage and management (CommonSpace) A proposed digital application called CommonSpace (created with the local organization Park People and the Gehl Institute) would make it substantially easier, faster, and less expensive to collect more reliable data on how people use public spaces — helping park operators better respond to community needs.	**Aggregated and/or non-personal:** Gehl public realm activity categories, usage counts **Non-personal:** Extremely high-level demographic details	**City planners, community groups, and others** could use this information to research park spaces and equipment that show the highest use in different parks throughout the city. **Community-based groups** could develop planning apps and tools that allow community members to better suggest park uses for all ages and abilities in their neighbourhoods.	**Open space management:** Range of government, non-profit, and community groups **Park operations:** Gehl Institute and other urban planning and design groups **City operations:** mySidewalk, Namara, Stae, and other platforms supporting city operations insights
Public realm maintenance map A proposed real-time map of public realm assets — from park benches to drinking fountains to landscaped gardens — would enable proactive maintenance and keep spaces in good condition.	**Non-personal and/or aggregated:** Evapotranspiration, plant health, moisture, waste bin volume, air quality **Non-personal and/or de-identified:** Public realm asset location, usage, damage detection; decibel meter (e.g. only volume level, not recording audio)	**Software developers** could use this information to create automated maintenance services, such as precision agriculture systems or landscaping bots. **Industrial manufacturers** could use data on utility maintenance to identify more durable materials or component designs. **City officials, business improvement districts, and others** could use this information to better schedule core operations, such as waste collection or green-space watering, to lower costs and improve quality of life.	**Physical asset location:** Bench Mark, BeWhere Inc., Estimote, Tekt **People flow:** Eco-Counter, Numina, PeopleFlow **Autonomous equipment:** BigMow, Husqvarna, Kobi **Predictive maintenance:** AI Incorporated, Arable, Mero Technologies, Nanophyll, Opti, Plantix, Sensoterra

Sidewalk Labs' proposed service or application	What urban data it proposes to use and/or publish	Possible third-party applications that could build on this data	What existing ecosystem the innovation supports (Names are illustrative only.)
Civic engagement (Collab) A proposed digital application called Collab (prototyped with local communities and Digital Public Square, a non-profit spun-out of the University of Toronto) would aim to engage community members in local decisions that could shape their neighbourhood, such as programming in a central public space, through a transparent process that reveals the decision-making framework and all community inputs. (Try the prototype at collab.sidewalklabs.com.)	Non-personal: Program choice selections, pre-populated and user-generated options Aggregated and/or de-identified: Broad demographic information (only upon clear opt-in / consent)	A **neighbourhood association** could clearly explain the tradeoffs associated with a decision about public space programming: for example, a farmers market provides fresh produce and draws a lot of foot traffic, but the space may feel too congested for a community picnic. A **research team** could analyze data to see if inputs are inclusive and representative of the community. A **community group** could evaluate user-generated inputs without revealing personal information.	Public input support: Range of government, non-profit and community groups such as neighbourhood associations, business improvement areas, public realm management organizations, and planning departments Community engagement and decision making: Decidem, Neighborland, Ethelo, and other platforms
Outcome-based building code This proposed real-time building code system could monitor noise, nuisances, and structural integrity to help a mix of uses thrive without sacrificing public safety or comfort.	Non-personal, aggregated, and/or de-identified: Strain gauges, vibration, odour, sound pressure, decibel meter (e.g. only volume level, not recording audio) Aggregated and/or non-personal: Safety sensors (e.g. sprinkler pipe pressure, fire pump diagnostics, heat, smoke, CO_2, CO PM 2.5, PM10, VOC, lead detection) Restricted data (not published for privacy reasons): Individual measurement data for the safety metrics above	**City government** could use this information to develop new outcome-based regulatory systems for code compliance. **Planning researchers** could use this information to study the relationship between mixed-use development and local economic growth. **City agencies or architectural groups** could create apps to visualize building structural integrity issues.	Environmental collection: Aclima, AQMesh, Awair, Concrete Sensor, Fibos, Koto Labs, NoiseAware, Safehub Building outcomes mapping: The Black Arcs, Map Your Property, RATIO.CITY
Active stormwater management A proposed active stormwater system would rely on green infrastructure and digital sensors to retain stormwater, reuse it for irrigation, and empty storage containers in advance of a storm to avoid combined sewer overflow.	Non-personal and/or aggregated: Stormwater tank level, stormwater flow meter, total suspended solids, valve and gate status, underwater water quality near shore	**Environmental researchers** could design an app to determine the number of plantings and amount of greenery needed to reduce stormwater flows and the need for secondary treatment. **City planners** could use this information to better plan (and minimize) hard infrastructure needs for stormwater, such as tanks and treatment facilities.	Digital management: Aquatic Informatics, IBM, Innovyze, Opti, Parjana, RainGrid, SUEZ, Veolia North America Water quality: Acoubit, FREDsense, Orb, Xylem, ZwitterCo

Sidewalk Labs' proposed service or application	What urban data it proposes to use and/or publish	Possible third-party applications that could build on this data	What existing ecosystem the innovation supports (Names are illustrative only.)
Energy management system (Schedulers) This proposed system of Home, Office, and Building Operator Schedulers would automate energy use to optimize residential, commercial, and building heating, cooling, and electricity systems — reducing energy waste and relying on clean energy while increasing tenant comfort.	Non-personal: Outdoor weather Aggregate and/or de-identified: Data on room temperature and humidity; energy use by type (e.g., from plug loads, lighting, HVAC); motion or occupancy; ambient light; comfort levels / complaints Restricted data (not published for privacy reasons): Individual measurement data for the metrics above (e.g. timestamped data about particular plug loads, occupancy detection for particular rooms) and any data about individual residential units	**Energy researchers** could use this data to compare neighbourhood energy usage across a city. **Architects and designers** could use this information to improve building designs. **Regulators** could use this information to create a dynamic energy code system based on actual operators instead of design-based models. **Climate organizations** could create apps to help individuals or households gamify their energy savings (provided users consent to share their data).	Building management systems: Automated Logic Controls, Johnson Controls, Schneider, Siemens Niche building analytics providers: Basking Automation, Comfy, eleven-x, Encycle, Parity, Peak Power, Cortex, Raybased, SensorSuite, SimpTek, SHIFT Energy, Thoughtwire, Density, InnerSpace Energy use measurement: VoltServer, Enertiv, Sense, Wemo, Currant Thermostats: Ecobee, Honeywell, Google Nest, Samsung Smart switches, lighting, appliances, and other hardware: Lutron, Enlighted, LG, TZOA
Building waste management systems To help divert landfill waste, a proposed program of responsive digital signage would help residents and businesses sort their trash, recyclables, and organics (foods) by illustrating common sorting mistakes. "Pay-as-you-throw" waste chutes would support this recycling program while helping to reduce overall waste.	Aggregated and/or de-identified: Trash volume, pressure scales (weight), waste classification for sorting using computer vision, contamination issues	An **environmental researcher** could team up with a fabrication studio to design a more sustainable coffee-cup lid based on disposal habits. **City planners** could use this information to understand best practices in buildings and to test new systems and strategies to scale to other buildings. **Computer-vision startups** could use information on common recycling errors to design augmented-reality apps that could help people classify waste. **Environmental groups** could design an app that provides feedback to consumers, both residential and commercial, encouraging higher recycling rates.	Smart waste: AMP Robotics, Anaconda, CleanRobotics, Compology, Enevo, Recycle Track Systems, Rubicon Global, Zerocycle

An outcome-based building code system to enable a safe, vibrant mix of uses

For most of the 20th century, cities separated residential, commercial, and industrial uses geographically to protect homes from noise, air pollution, and other nuisances.[44] This approach made sense in a world without reliable tools to monitor the environmental nuisances of commerce and industry. But it also discouraged an active mix of home, work, and retail into the same neighbourhood — let alone the same building.

Working alongside local government, Sidewalk Labs proposes to create a real-time building code system designed around the premise that buildings should be able to house a diverse range of tenants — residential, commercial, and light industrial alike — so long as everyone adheres to agreed-upon "outcomes," such as minimizing noise, air pollution, and other public nuisances.

What urban data it proposes to use.

The proposed outcome-based building code system would monitor several types of building regulations on an ongoing, real-time basis via environmental sensors that collect non-personal data. The environmental information collected is considered "urban data," because it would be data collected in a privately owned common space in the IDEA District.

Devices would be placed in building hallways to collect information on structural integrity and vibration, odours, interior air quality, and noise levels. This system

would be designed to collect only the specific data pertaining to building codes. Additionally, buildings would implement non-personal safety sensors to measure things like sprinkler pipe pressure, fire pump diagnostics, heat and smoke, and particulate matter.

This information would be provided from the third-party owners of these devices to an outcome-based code datastore. Any violation detected in this datastore would be sent to building managers for next steps and resolution.

In the case of an emergency (e.g., fire) or non-compliance, municipal officials could query the database directly.
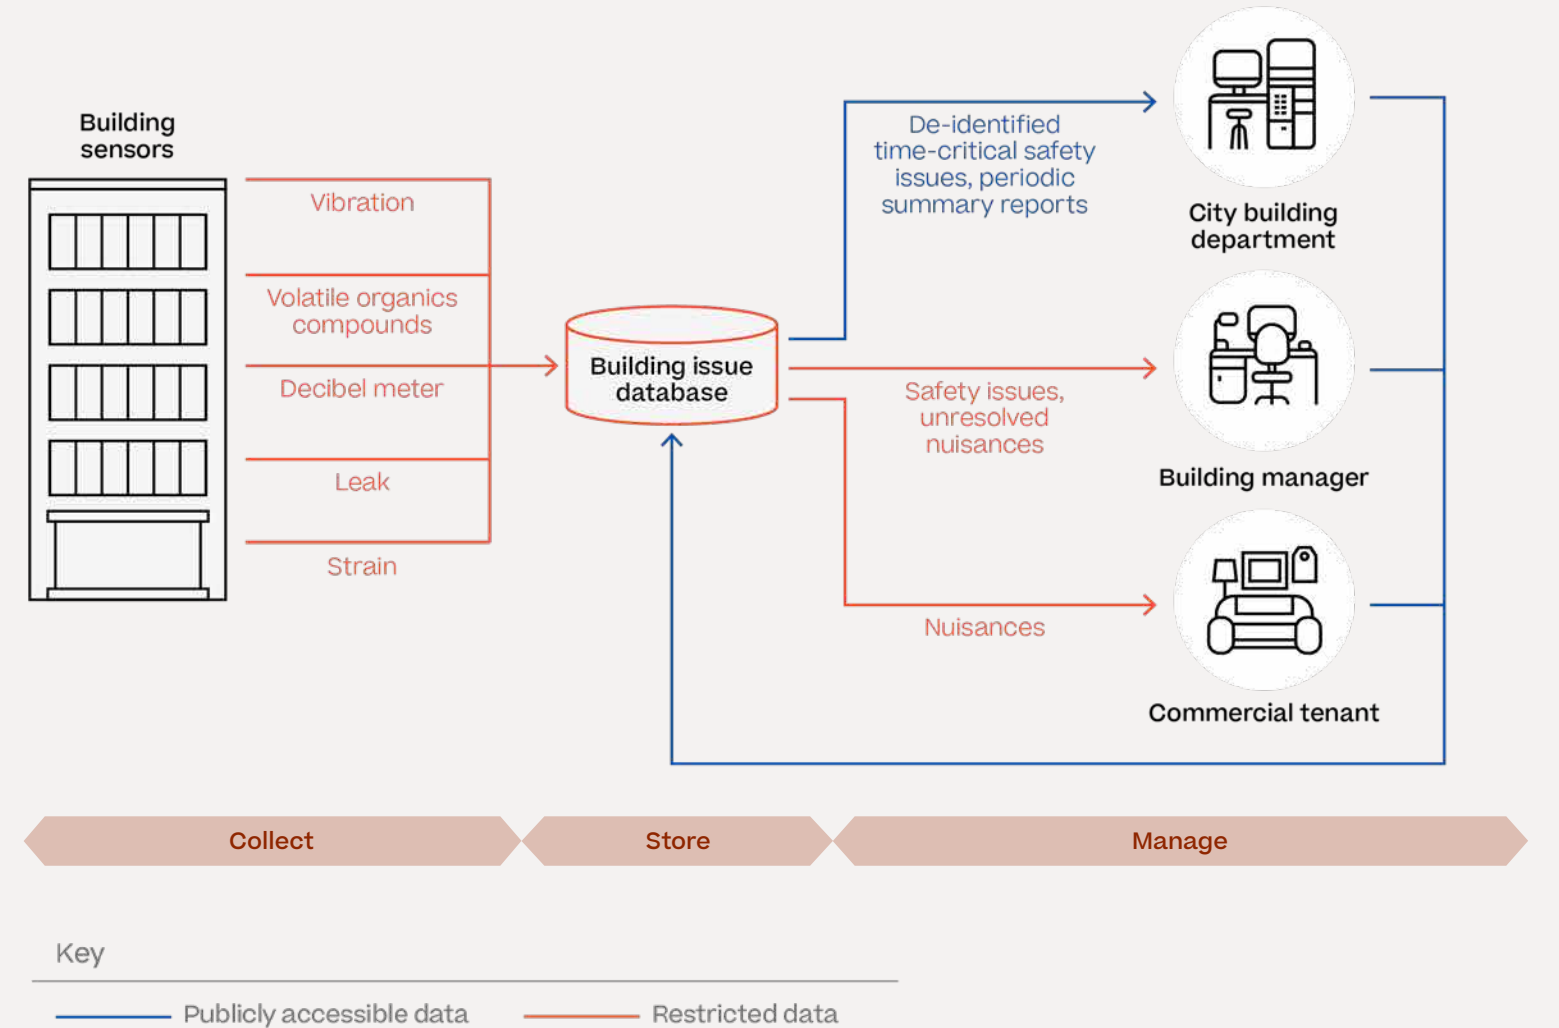
What the RDU Assessment could consider.

The beneficial purpose of this proposed innovation would be to enable a greater mix of residential, commercial, and light industrial spaces, helping to create a lively local economy and achieve Waterfront Toronto's goals for complete communities. The collection of urban data would be necessary to ensure the industrial spaces would comply with regulatory conditions, such as noise and odour requirements, thus enabling both commercial and residential tenants to coexist safely.

The proposal would be developed in accordance with the RDU Guidelines. The expected impact on people would be small, given that the sensors involved

How it works: Outcome-based code

Building sensors that detect code violations could send these issues to a restricted database accessible by the city, building managers, and tenants, with only aggregated data publicly accessible to third parties.

Building sensors
- Vibration
- Volatile organics compounds
- Decibel meter
- Leak
- Strain

Building issue database

De-identified time-critical safety issues, periodic summary reports → **City building department**

Safety issues, unresolved nuisances → **Building manager**

Nuisances → **Commercial tenant**

Collect — Store — Manage

Key
— Publicly accessible data   — Restricted data

in this initiative would collect non-personal information related to building codes. Because this data could be linked to individual building hallways, however, this data would be considered restricted and not publicly accessible. For these reasons, Sidewalk Labs believes the balance of impact of collecting the environmental data would weigh in favour of the proposal.

What it makes possible by others.

The non-personal data collected by the outcome-based code system, as well as information aggregated by neighbourhood level, would be shared with a pub-

licly accessible API, enabling third parties to build on top of it.

A potential future innovation could include the adoption by city government of a new system for code compliance or zoning based not on pre-existing, rigid standards but rather on real-time performance to help Toronto achieve its goals for mixed-use development. Additionally, city agencies or their private vendors might create an app to visualize a building's structural-integrity issues in real time. Such a tool could save money by efficiently identifying problems and catalyzing proactive maintenance. 📖

See the "Buildings and Housing" chapter of Volume 2, on Page 202, for more on outcome-based building codes.

Launching Core Digital Services
That Others Can Build On

An Office Scheduler to optimize energy use

Today, no one is focused on saving energy in commercial tenant spaces, such as offices. Existing energy management programs that could optimize thermostats and ventilation systems in commercial spaces are under the control of the building operator, not the tenant.[45] The result is that offices often operate based on default system schedules that do not match the tenant's needs.

To help commercial tenants manage energy consumption and costs, Sidewalk Labs proposes to use a tool called the Office Scheduler that would optimize all the systems under tenant control, based on factors such as energy prices. This tool is part of a suite of Scheduler tools that together would reduce greenhouse gases compared with standard downtown buildings, consistent with Waterfront Toronto's ambitions for achieving a climate-positive community.

What urban data it proposes to use.
To achieve this goal, the Office Scheduler would need visibility into electricity usage and cost, as well as real-time metering of all building energy systems, such as heating, cooling, lighting, and equipment. An encrypted building-energy datastore would aggregate information and automatically determine any optimization steps across systems for both occupant comfort and energy savings.

The proposed Office Schedulers would incorporate data from a set of energy management sensors (such as ambient

lights, motion sensors, plug load monitors, room temperature gauges, and digital thermostats) as well as from computer systems (such as calendar notifications) to reduce energy use when rooms are unoccupied or already comfortable. This information would be provided from the third-party owners of these devices to a data format translator.

To register requests for temperature changes from workers, the Office Scheduler would use some personal information by direct consent through an app (making this transaction data). This information could be used to respond to worker complaints, and if a change could not be accommodated due to competing requests, it could be used to guide workers to areas of the office that might be more comfortable.

What the RDU Assessment could consider.
The beneficial purpose of the Office Scheduler is to help achieve a climate-positive community through reducing energy consumption in commercial spaces and to optimize for clean energy use. Other benefits include a 20 percent reduction in building energy operating costs (when used in concert with the other Scheduler tools) and greater comfort for workers.

The expected negative impact on people would be small, given that minimal personal information is required and would be de-identified or aggregated for its

---

How it works: Office Schedulers
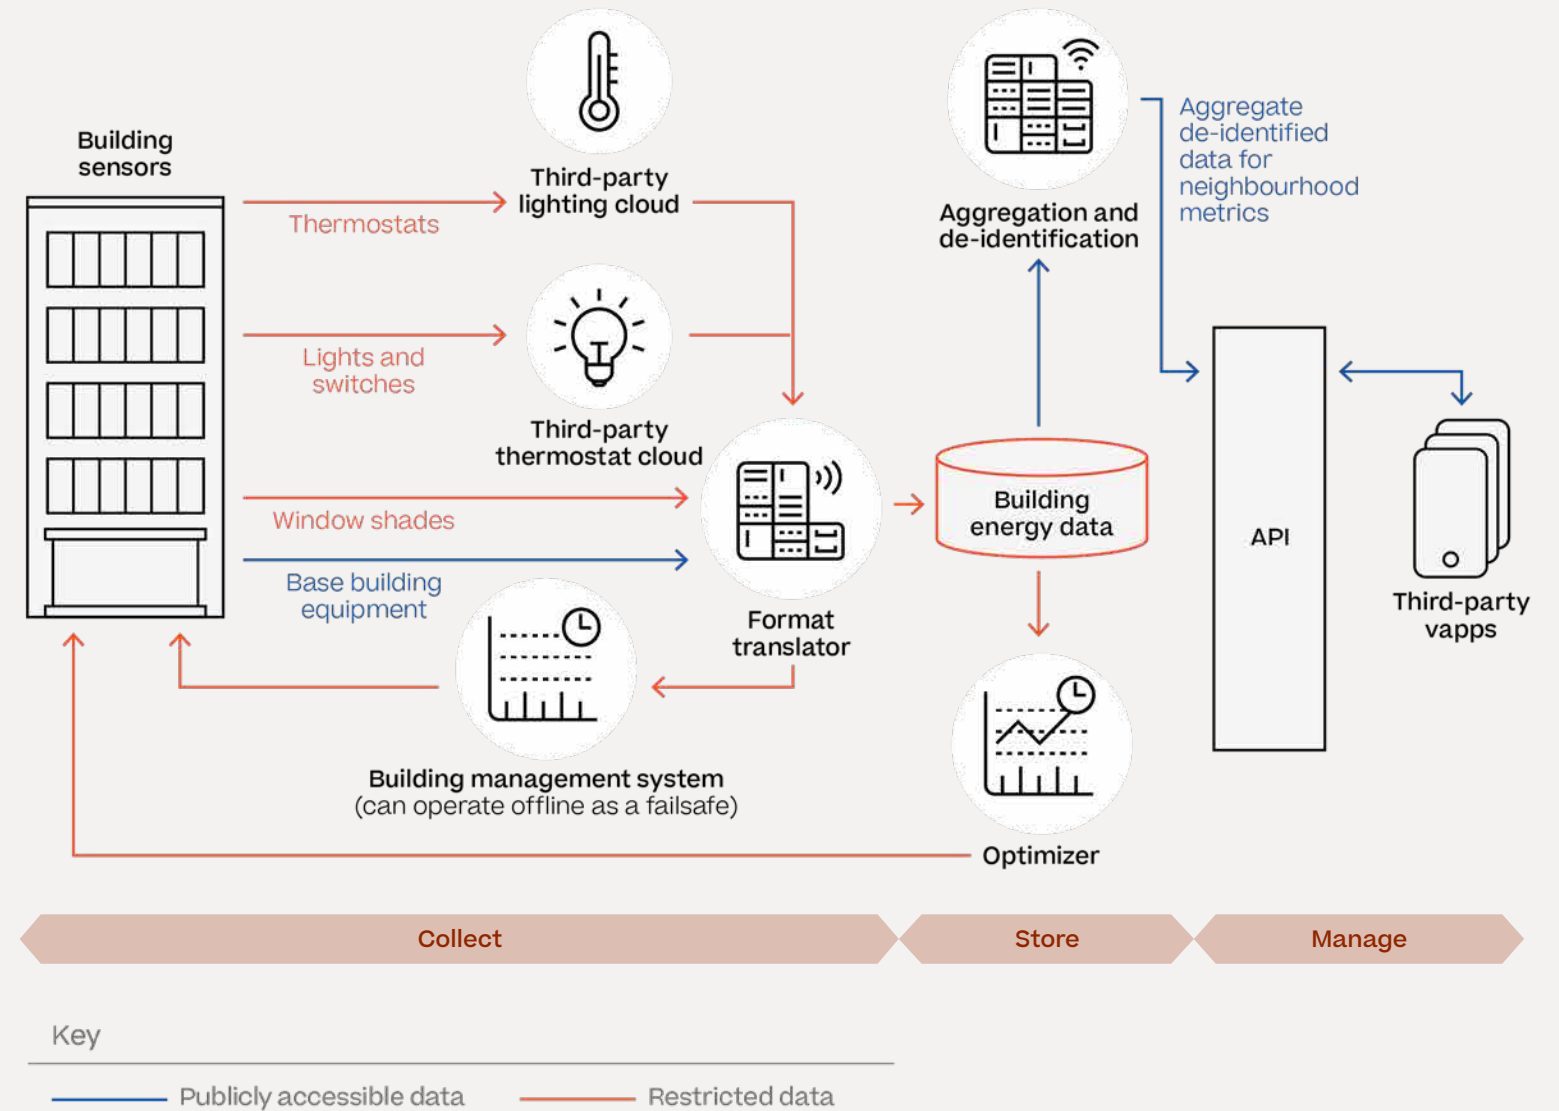
Information from energy-related sensors would help the Office Scheduler tool optimize building energy use, with aggregated and de-identified data made publicly accessible to third parties.

Key

— Publicly accessible data   — Restricted data

---

intended use. Non-personal and de-identified data, including neighbourhood-level metrics, would be made publicly accessible so that others could use this data. Personal information (which is subject to Canadian privacy laws) would be stored in a secure database with access restricted to certain employees and agents and only be kept as long as necessary to fulfill the original purpose.

While the Office Scheduler proposes to automate some energy actions, tenants would have the ability to override the automated system, and the algo-

rithm would also undergo a Responsible AI assessment. Sidewalk Labs believes the balancing of the risks of collecting the data in offices would weigh in favour of the data collection activity.

What it makes possible by others.
Third-party apps and services would be able to use de-identified and aggregated data for research purposes, such as comparing neighbourhood energy usage across a city to improve building designs or evaluate energy policies, or to create new tools, such as behavioural apps that help families gamify their energy savings.

See the "Sustainability" chapter of Volume 2, on Page 296, for more on the proposed Office Scheduler.

Spotlight 3

A mobility management system to reduce congestion and improve safety

Sidewalk Labs' proposed mobility management system would use non-personal and de-identified urban data (such as trip counts, traffic congestion measures, and curbside availability information) to help manage the transportation network in line with objectives around street safety, shared trips, and travel times. This tool would be able to understand how people are using the entire system (including all trip modes), analyze these travel patterns, and encourage trip choices that do not rely on private cars — all in real time.

What urban data it proposes to use.
To estimate traffic flows or prioritize pedestrian safety, lidar, radar, and cameras would need to be able to detect all travellers and vehicles at an intersection, de-identifying that information on the device and providing only an aggregate count. To manage congestion around curb spaces, in-pavement occupancy sensors would need to detect the presence of vehicles without identifying specific vehicles. A separate licence plate reader could capture parking data about vehicles violating parking rules to send it directly to the city for municipal enforcement.

Municipal enforcement could be performed via traditional methods used by the City of Toronto today, or improved by providing enforcement agencies with better information and tools (such as recommended areas where violations are more likely) or systems that enable the

city to perform automated enforcement (such as vehicle transponders or license-plate readers).

The data collected by the mobility system could flow to two key databases. All non-personal and de-identified information could flow to an open datastore, publicly accessible via an API. Private data could flow to an enforcement datastore, with access restricted to municipal officials only.

What the RDU Assessment could consider.
This mobility management system formed the basis for the illustrative RDU Assessment case study on Page 436. As noted there, Sidewalk Labs believes that under the proposed RDU Guidelines, this proposal would gain approval for having a beneficial purpose related to travel time and increased public transit use, helping to achieve Waterfront Toronto's objective for sustainable transportation. Privacy risks would be mitigated through de-identification.

If necessary, some of this data could be collected by a public entity that is authorized to enforce relevant bylaws and regulations. In these cases, only the city would have access to this data. As such, this collection and use would be governed by the Municipal Freedom of Information and Protection of Privacy Act, and the city would follow its own privacy practices.

How it works: Mobility Management System

To operate a "dynamic curb," a mobility management system collects information about curb availability, stores that information in databases, and makes non-restricted data publicly accessible to third parties.
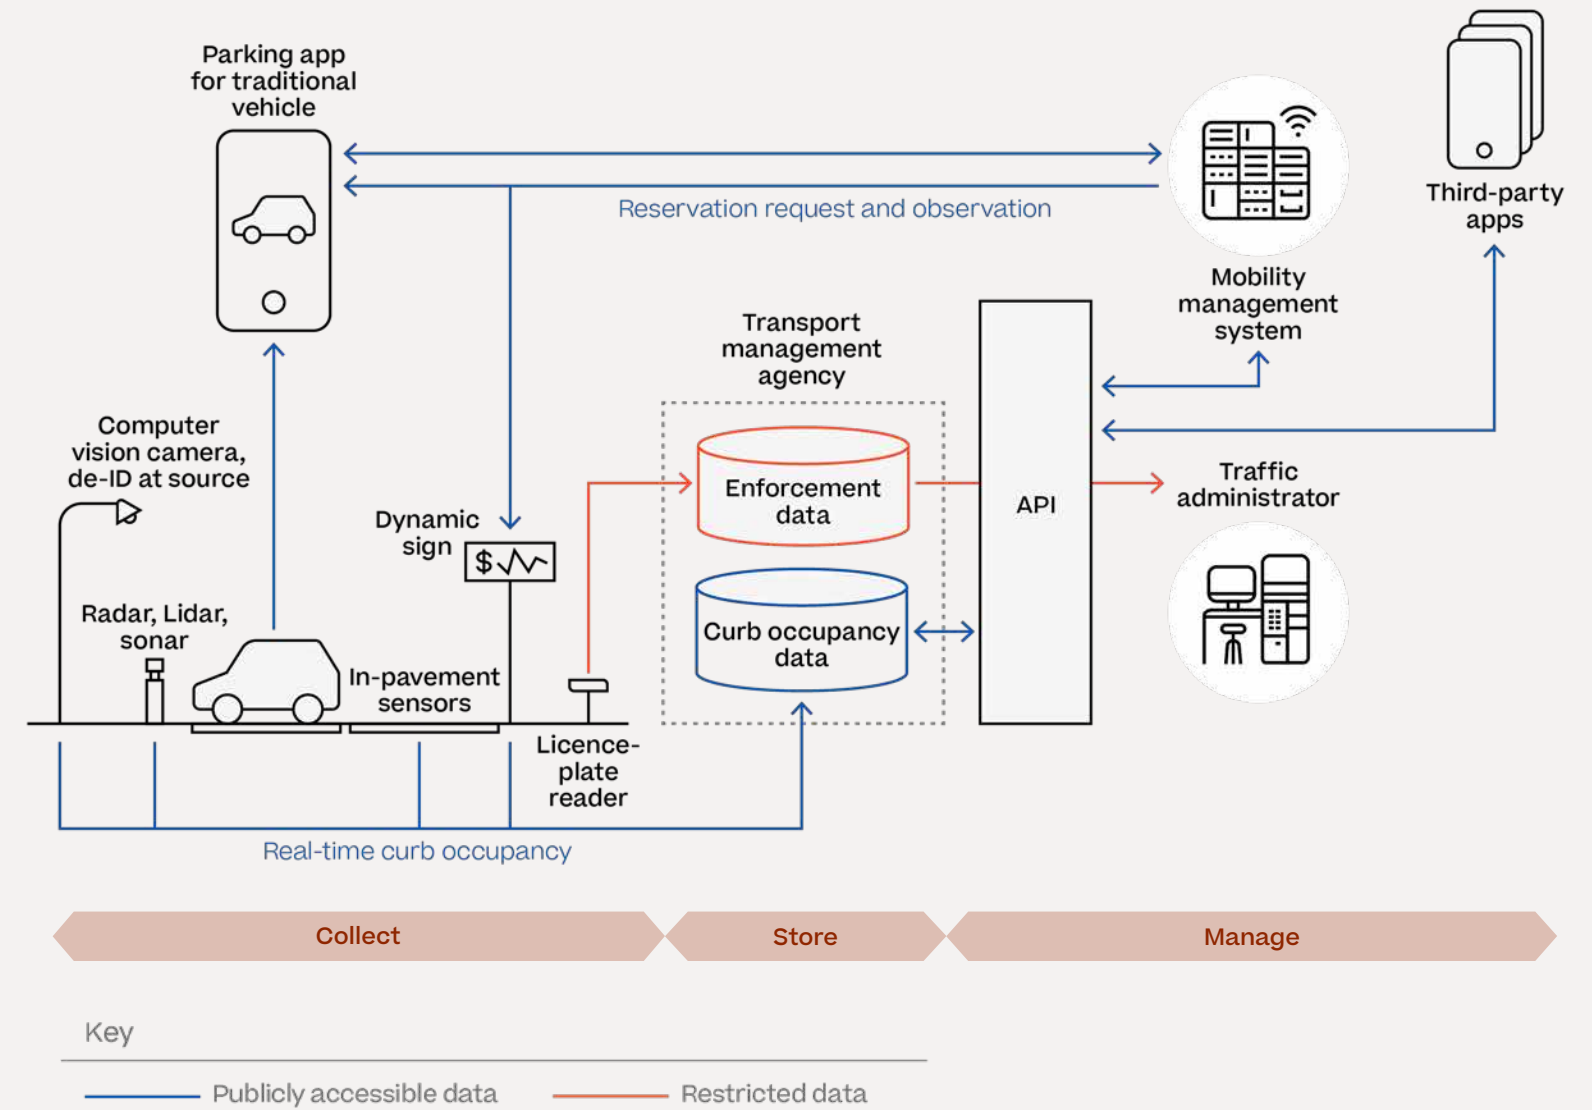
What it makes possible by others.
This mobility management system — along with third-party developers who create navigation apps or ride services — would be able to pull publicly accessible data from the API to provide travellers with information that helps them make trip choices, such as public transit arrival times, bike-share availability, or prices for curb space. Such publicly accessible data would also enable third parties to create new services in the future.

For example, a navigation app might use the aggregate trip patterns and available mode options to provide users with the fastest, cheapest, or greenest routes from A to B. Self-driving vehicle companies could use the information on intersection movement to improve technology that detects pedestrians or cyclists. Local officials would be able to use the curbside availability data to propose new guidelines for ride-hail services.

See the "Mobility" chapter of Volume 2, on Page 22, for more details on the proposed mobility management system.

PGM Note:
~12 lines/~70 words over

Public Engagement

The following summary describes feedback related to **digital innovations**, and how Sidewalk Labs has responded in its proposed plans.

As part of its public engagement process, members of Sidewalk Labs' planning and innovation teams talked to thousands of Torontonians — including members of the public, expert advisors, civic organizations, and local leaders — about their thoughts, ideas, and needs across a number of topics.

1 Protect people's privacy and use data to serve the public good

What we heard

Throughout the public engagement process, Torontonians were loud and clear: data privacy matters. Residents were wary about third-party access to data collection and the commercial sale of data. The Data Governance Advisory Working Group recommended that "Privacy by Design" principles be incorporated into the project. The Sidewalk Toronto Fellows advised Sidewalk Labs to ensure that, as a first principle, data be collected and used with the public good in mind.

Public Roundtable 4 participants who took part in a data-focused discussion were particularly helpful in defining the use cases they were comfortable with. For example, as long as data was de-identified, residents felt comfortable with data being collected and used for transit and mobility purposes. As one Reference Panel resident said: "Cities need aggregate data. … They need to know which modes of transportation people take when it's raining. They need to know how many people went through an intersection, not who went through it. And if they can legitimately anonymize the data they collect then I would accept that."

The Residents Reference Panel had many data-related concerns, including the need to ensure that algorithms would not perpetuate existing biases. They also wanted to ensure the cyber-security of this tech-enabled neighbourhood would be state of the art.

How we responded

Designing for privacy.
For all its projects, Sidewalk Labs plans to incorporate Privacy by Design, an approach that requires thinking about potential privacy impacts at the very start of a project lifecycle and proactively embedding privacy measures into the design of a project (see Page 424).

Creating a steward.
To protect personal privacy and the public good, Sidewalk Labs proposes the creation of an independent entity called the Urban Data Trust to oversee digital matters and approve (or deny) proposals to collect or use urban data in the IDEA District (see Page 420).

Establishing guidelines.
Sidewalk Labs proposes that the Urban Data Trust establish a set of RDU Guidelines that apply to all parties engaged in the collection and use of urban data in the IDEA District. These guidelines would build on the strong existing framework of Canadian privacy laws (see Page 424).

Increasing transparency.
Sidewalk Labs proposes that all entities complete RDU Assessments with any proposal to collect or use urban data to ensure that digital services abide by the RDU Guidelines. RDU Assessments would be filed and publicly registered with the Urban Data Trust before a project or service could launch (see Page 429).

2 Earn public support through transparent policy, clear language, and data education

What we heard

Participants were concerned that Torontonians needed more education to advance their data literacy and that companies and organizations needed to be more transparent in the ways they collect data. They wanted to know more about how data collection would happen in a place like Quayside.

The Sidewalk Toronto Fellows, Reference Panel residents, and Roundtable participants urged Sidewalk Labs to proactively disclose when (and what kind of) data is being collected and used in clear language. As one roundtable participant noted: "Data privacy and responsible data use needs genuine commitment — that includes being specific and transparent about how it will be used."

Participants also wanted to ensure ways to consent or opt-out of data collection and use, especially in public spaces, where meaningful consent is a challenge. The Data Governance Advisory Working Group suggested that signage alerting the public to what data is being collected and how it is being used could be helpful.

Benefiting people.
Sidewalk Labs commits to applying Canadian values of diversity, inclusion, and privacy as a fundamental human right to its digital projects, providing a clear purpose and benefit to any proposed collection and use of urban data. No data for data's sake (see Page 424).

De-identifying by default.
Sidewalk Labs proposes that one of the RDU Guidelines state that personal information must be de-identified by default at first use, so it cannot be traced back to any individual (see Page 424).

Enhancing security.
Sidewalk Labs proposes to deploy a new security approach called "software-defined networks" capable of detecting security compromises and isolating impacted devices from the network (see Page 392). Sidewalk Labs also proposes to base all security and reliability standards on best practices and to emphasize resiliency across its systems (see Page 408).

Being proactive.
To establish a proactive approach to security, each digital system Sidewalk Labs proposes would use a preparedness assessment to provide clear answers to key questions on threat modelling and response readiness (see Page 412).

Protecting from ads.
Sidewalk Labs commits that it would not sell personal information to third parties or use it for advertising purposes. To encourage such behaviour from other companies or entities operating in the IDEA District, Sidewalk Labs proposes that the Urban Data Trust place greater levels of scrutiny on projects wishing to use personal information for ad purposes, including the need to justify this decision and to obtain explicit consent from users (see Page 425).

How we responded

Being transparent.
Sidewalk Labs proposes that all projects aiming to collect or use urban data must inform individuals of how and why their information is being collected and used, and do so in a way that is proactive, clear, and easy to understand — not written in legalese (see Page 424).

Providing clarity.
For the collection of urban data in public spaces, where meaningful consent cannot reasonably or reliably be achieved, Sidewalk Labs proposes that entities provide clarity of usage through efforts such as physical signs notifying people of a data device or informational websites describing a service or program in greater detail (see Page 424).

Improving design.
Sidewalk Labs released via Github a draft of new design patterns co-created with more than 100 participants from several cities worldwide. The goal of the new patterns was to build on the consent and notice requirements that exist under current privacy laws in a way that increases digital transparency and helps people quickly get a sense of the privacy implications associated with responsible urban data collection.

Registering devices.
Sidewalk Labs proposes that the Urban Data Trust not only approve the placement of data-collection devices but also publish and maintain an online registry and map of device locations, with easily accessible information on what kind of data is being collected, why, how, where, and by whom (see Page 433).

Supporting literacy.
In Quayside, Sidewalk Labs proposes to establish a Tech Bar that would provide community members with small-group or one-on-one assistance with digital tools, with the goal of improving digital literacy among the local community.

Attendees of the "Digital Transparency in the Public Realm" workshop are hard at work. Credit: Sidewalk Labs

3 Tech should be an enabler and an accessible amenity

What we heard

Residents were excited about the opportunity for Quayside to be a world leader in urban technology and to encourage and enable future tech innovations.

Torontonians hoped the Sidewalk Toronto project would improve existing public services, potentially by leveraging technology. As one Reference Panel resident explained: "The challenge is to find ways for technology to help foster a sense of community. That seems utopian but it's possible… I think Toronto can be a global model for a new kind of technology that helps keep us human." Participants were also open to new tools or options that would give community members more of a voice in decisions on programming and services.

Other residents were excited by new potential services, such as enhanced Wi-Fi connectivity. Still others wanted to see technology that would make Quayside more accessible, such as customizable tech that could be experienced in multiple ways.

The Data Governance Advisory Working Group encouraged Sidewalk Labs to pursue open data whenever possible, and the Sidewalk Toronto Fellows recommended that Sidewalk Labs develop an open data portal to encourage innovation for the public good.

How we responded

Connecting people.
Sidewalk Labs proposes to create a super-fast, ubiquitous connectivity network that would provide residents, workers, and businesses access to their own secure, personal high-speed network — no matter where they are — at an affordable cost (see Page 384). For people without smartphones or computers, devices and Wi-Fi kiosks would be available and free to use in communal spaces.

Standardizing data.
Sidewalk Labs plans to publish data in standard formats and via well-defined, public APIs. Where standards do not exist, Sidewalk Labs plans to work with companies, researchers, and standards bodies to create those standards (see Page 405).

Opening data.
To encourage innovation, Sidewalk Labs plans to make publicly accessible all urban data that could reasonably be considered a public asset. Sidewalk Labs plans to work with organizations and companies that are already building open data portals to provide access to this data, and also proposes that the Urban Data Trust facilitate integration with existing open data portals and tools (see Page 406).

Opening code.
Sidewalk Labs plans to make software source code public under free software licences and to encourage other entities creating services in the IDEA District to do the same (see Page 406).

Avoiding lock-in.
Sidewalk Labs proposes that any digital infrastructure it deploys be open to competition and alternatives. As one example, it proposes to deploy a new type of standardized mount that would make it easier for cities to swap in new digital tools and avoid relying on proprietary services (see Page 380).

Prioritizing accessibility.
In keeping with its accessibility principles, Sidewalk Labs commits to offering technology in multiple modes and maintaining best accessibility practices. (For further reading on accessibility, see Volume 1.)

Supporting inclusive usability testing.
Sidewalk Labs is currently funding GRIT Toronto, a program founded by Code for Canada that incorporates community feedback into the creation of new digital services and products, helping to ensure these tools reflect the needs of the populations they are intended to support (see Page 443).

Enabling civic engagement.
Sidewalk Labs is developing a prototype with Digital Public Square called Collab that would allow community members to propose ideas for events in their neighbourhood. The tool is designed to walk users through the tradeoffs associated with various proposals, including how their individual choice would impact the community (see Page 446).

Sidewalk Labs' Director of Design Michelle Ha Tucker describes the co-design process during a "Digital Transparency in the Public Realm" workshop at 307. Credit: Sidewalk Labs

4 Establish an ethical data governance model for the long-term

What we heard

The Sidewalk Toronto Fellows recommended that Sidewalk Labs establish an independent entity to ensure data stewardship, and the Residents Reference Panel suggested that, when possible, data be stored, regulated, and analyzed in Canada.

Residents wanted to know more about the Civic Data Trust initially proposed by Sidewalk Labs in 2018, including how the trust would integrate into existing legal and regulatory frameworks and ensure compliance for all. (The entity has now become the Urban Data Trust; see Page 423 for details on this shift.)

Residents also wanted to better understand the data-governance model overall — including how long-term data management and storage would work — and how the government could provide appropriate oversight over the project.

How we responded

Implementing an entity.
As noted earlier, Sidewalk Labs proposes the creation of an independent entity called the Urban Data Trust with the capacity to approve all proposals for use and collection of urban data and with a mandate to balance the public interest and the need for innovation (see Page 420).

Building on laws.
Sidewalk Labs proposes that the Urban Data Trust coordinate with privacy regulators and that the responsible data use process build on (not replace) existing privacy laws (see Page 419).

Ensuring accountability.
Sidewalk Labs proposes that the Urban Data Trust uphold data agreements through contracts that are legally enforceable and actionable (see Page 421).

Thinking long-term.
Looking long-term, Sidewalk Labs puts forth that the Urban Data Trust could be ultimately transformed into a public-sector agency or a quasi-public agency, either of which could give it more long-term viability or broader coverage (see Page 422).

Localizing data.
Sidewalk Labs commits to using its best efforts at data localization, as long as there are Canadian-based providers who offer appropriate levels of security, redundancy, and reliability. To the extent that it is deemed infeasible to store data solely in Canada, Sidewalk Labs would be transparent about such a decision (see Page 412).

Engagement spotlight

Attendees talk during the first "Digital Transparency in the Public Realm" workshop in Toronto. Credit: Sidewalk Labs

Alyssa Harvey Dawson heads privacy and data governance for Sidewalk Labs. When she first started at the company, she knew that the challenges facing a company whose mission is radically improving urban life through the use of technology would be unique. This realization came into greater focus in conversations with the Data Governance Advisory Working Group.

The working group pushed Alyssa and her team to consider how data privacy, use, and management take on new meanings when the source of that data is the public realm. "You can't just focus on personal information, which is where most privacy laws begin and end," says Alyssa. "The scope of data that could be collected from a private actor in public spaces, where you don't have all the usual protections, makes the concerns much more heightened. You have to think more broadly about the impact on people."

In response, Alyssa and her team coined a term, "urban data," that refers to aggregate, non-personal, de-identified, or personal data gathered in the physical spaces of a city, including its public realm, its publicly accessible spaces, and even some private spaces. They then proposed the creation of an independent entity that would represent the public interest and serve as the steward for the collection and use of all urban data across the IDEA District.

With these proposed initiatives, Alyssa and her team hope to advance the conversation about responsible data use in cities in new directions and inspire local solutions to this critical — and growing — challenge.

Toronto can demonstrate to the world that cities do not need to sacrifice their values of inclusion and privacy for economic opportunity in the digital age.

Acknowledgements

Endnotes

General note: Unless otherwise noted, all calculations that refer to the full proposed IDEA District scale are inclusive of the entirety of its proposed geography, including all currently privately held parcels (such as Keating West). Unless otherwise noted, all currency figures are in Canadian dollars.

Charts note: Sources for the charts and figures in this chapter can be found in the accompanying copy for a given section; otherwise, the numbers reflect a Sidewalk Labs internal analysis. Additional information can be found in the MIDP Technical Appendix documents, available at www.sidewalktoronto.ca/midp-appendix.

1. The Brookfield Institute for Innovation + Entrepreneurship, *The State of Canada's Tech Sector, 2016*. July, 2016.

2. PlanetWeb, *Toronto is Among the Fastest Growing Tech Innovators*. March 31, 2018.

3. Tech Toronto, *How Technology is Changing Toronto Employment*. 2016.

4. For more information on economic projections and job creation, please see the "Economic Development" chapter in Volume 1.

5. ScienceDirect, *Passive Optical Network*. Elsevier, 2019.

6. Photonics Media, *Fiber Optics, Understanding the Basics*. 2019.

7. Robbert van der Linden, *Adaptive Modulation Techniques for Passive Optical Networks*. Eindhoven University of Technology, 2018.

8. Craig Nevill-Manning and Prem Ramaswami. *DSAP Technology Update*. January 17, 2019.

9. *IEEE, The world's largest technical professional organization dedicated to advancing technology for the benefit of humanity*. IEEE, 2019.

10. Phillip Dampier, *AT&T and Comcast Successfully Slow Google Fiber's Expansion to a Crawl*. Stop the Cap, September 4, 2018.

11. LiveScience, *Internet History Timeline: Arpanet and the World Wide Web*. June 27, 2017.

12. *Privacy and Video Surveillance in Mass Transit Systems: A Special Investigation Report*. Information and Privacy Commissioner of Ontario, Privacy Investigation Report MC 07-68, March 3 2008; City of Toronto, *Red Light Cameras*. www.toronto.ca (accessed March 5, 2019); City of Toronto Transportation Data Catalogue, *Traffic Cameras*. opendata.toronto.ca (accessed March 5, 2019); Wireless Toronto, *Hotspot Map and List*. Brochure, December 2007. See also, "How many cameras are watching you? Toronto professor concerned about privacy." *CTV News Toronto*, February 26, 2015.

13. Office of the Privacy Commissioner of Canada, *Privacy in the Landlord and Tenant Relationship*. Reviewed 2018.

14. Petteri Kivimaki, "X-Road as a Platform to Exchange MyData." *Medium*, August 31, 2018.

15. *e-Estonian Guide*. e-Estonia, 2018.

16. Michael Bock, "Everything you need to know about APIs." *B2B News Network*, January 2015.

17. Kristjan Vassil, *Estonian e-Government Ecosystem: Foundation, Applications, Outcomes*. World Bank, June 2015.

18. e-Estonia, *Interoperability Services*. 2018.

19. Adam Rang, "Estonia (again) ranks 1st for tax competitiveness and 12th for ease of doing business." *Medium*, November 1, 2017.

20. Government of Ontario, *Register a Business Name or Limited Partnership*. Service Ontario, updated March 4, 2019.

21. Nathan Heller, "Estonia, the Digital Republic." *The New Yorker*, December 18 & 25, 2017.

22. GTFS, *General Transit Feed Specification*. 2018.

23. Bibiana McHugh, "Chapter 10, Pioneering Open Data Standards: the GTFS Story." *Beyond Transparency*, 2013.

24. GISGeography, *TIGER GIS Data (Topologically Integrated Geographic Encoding & Referencing)*. February 18, 2017.

25. Brad Bennett, *Sidewalk Labs' Toronto Transit Explorer Shows the Fastest Way to Get Around the City*. Mobile Syrup, May 4, 2018.

26. Lisa R. Lifshitz, *Ethics by Design: Canada Adopts AI Ethics and Data Declaration*. Canadian Lawyer, December 10, 2018.

27. PIPEDA was first implemented on January 1, 2000. For more information about PIPEDA, including subsequent amendments and applicable regulations, see Office of the Privacy Commissioner of Canada, *Privacy Laws in Canada*. www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/ (accessed February 27, 2019).

28. Econstats, *Personal computers per 100 people*. http://www.econstats.com/wdi/wdiv_597.htm (accessed February 27, 2019).

29. *Canadian Charter of Rights and Freedoms*. Part 1 of the *Constitution Act, 1982*, being Schedule B to the *Canada Act 1982 (UK)*, 1982.

30. Innovation, Science and Economic Development Canada, *Government of Canada launches national consultations on digital and data transformation*. News release, June 19, 2018.

31. Ontario Ministry of Government and Consumer Services, *Ontario's Government Launches Data Strategy Consultations*. News release, February 5, 2019.

32. Councillor Joe Cressy, seconded by Councillor Paul Ainslie, *Data Governance and Smart Cities*. Toronto City Council Notice of Motion, adopted February 26, 2019.

33. For further information on the panel's activities, consult the Waterfront Toronto Document Library, *Digital Strategy Advisory Panel Meeting Materials*, updated February 2019.

34. *Guidance on inappropriate data practices: Interpretation and application of subsection 5(3)*. Office of the Privacy Commissioner of Canada, modified January 3, 2019.

35. European Commission, *Free Flow of Non-Personal Data*. Digital Single Market Policy, updated January 9, 2019.

36. For more details, consult the report *De-identification Guidelines for Structured Data*. Information and Privacy Commissioner of Ontario, June 2016.

37. For the precise definition of "personal information" under Canada's *Personal Information Protection and Electronic Documents Act (PIPEDA)*, see Information and Privacy Commissioner of Canada, *PIPEDA in Brief*. January 2018.

38. Alex Ryan and Joe Greenwood, "Secure Toronto's smart data neighbourhood in a trust." Opinion, *The Toronto Star*, February 2019. See also MaRSDD, *A Primer on Civic Digital Trusts*. Gitbook, 2018.

39. *BiblioTech: Beyond Quayside: A City-Building Proposal for the Toronto Public Library to Establish a Digital Data Hub*. Toronto Region Board of Trade, January 2019.

40. *BiblioTech: Beyond Quayside*. TRBOT, January 2019.

41. Fairness Commissioner of Ontario, *Professions and Trades*. http://www.fairnesscommissioner.ca/index_en.php?page=professions/index (accessed March 4, 2019).

42. *Digital Governance Proposals for DSAP Consultation*. Sidewalk Labs, October 2018.

43. Privacy By Design Centre of Excellence, *The Seven Foundational Principles*. Ryerson University (accessed March 4, 2019).

44. William A. Fischel, "An Economic History of Zoning and a Cure for its Exclusionary Effects." *Urban Studies* Volume 1 Issue 2, February 2004.

45. U.S. Department of Energy, *Energy Efficiency in Separate Tenant Spaces - A Feasibility Study*. April 2016.

MIDP Acknowledgements

Sidewalk Labs Team

Rohit Aggarwala
Habon Ali
Jack Amadeo
Christopher Anderson
Brian Barlow
Mark Bauernhuber
Catherine Benz
Maya Borgenicht
Sarah Ruth Boyer
Simon Brandler
Matthew Breuer
John Brodhead
Marie Buckingham
Kia Burke
Laura Capucilli
Difei Betty Chen
Leo Chen
Chelsey Colbert
Reid Crombie
Colin Curzi
Chrystal Dean
Michael Delucia
Carrie Denning Jackson
Pino Di Mascio
Dasola Dina
Shaina Doar
Dan Doctoroff
Cara Eckholm
Andrew Edwards
Valerie Eisen
Robie Evangelista
Jan Fiegel
William Fields
Laura Fox
Krystina Francis
Benjamin Funk
Taylor Gesner
Johanna Greenbaum
Noah Greenbaum
Ryan Guptill
Abby Harrill
Shaun Harris
Alyssa Harvey Dawson
Marie Hlavaty
Brian Ho
Okalo Ikhena
Eric Jaffe

Nate Jenkins
Justin Joachimiak
Bronson Johnson
Nicholas Jonas
Mya Jones
Michael Kalt
David Katz
Tim Kau
Patrick Keenan
Ariel Kennan
Thomas Kennedy
Karim Khalifa
Emily Kildow
Laura Kilian
Eugene Kim
Annie Koo
Samridhi Kundra
Joanna Lack
Micah Lasher
Jacob Lazarus
Nicole LeBlanc
Lana Lee
Sara Lewis
Corinna Li
Jacqueline Lu
Mark Luckhardt
Christopher Macies
Alison Maddox
Drew Majkut
Jiten Manglani
Jorge Manrique Charro
Charlotte Matthews
Pierre Mayence
Davina Mazaroli
Mary-Margaret McMahon
Ayesha Menon
James Merkin
Amanda Meurer
Andrew Miller
Trista Miller
Amina Mohamed
Nerissa Moray
Nina Nappa
Neda Navab
Jennifer Neilands
Craig Nevill-Manning
Willa Ng

Moyosore Odubanjo
Douwe Osinga
Michelle Pacheco
Ananta Pandey
Marla Pardee
Julie Pedtke
Veronica Pinchin
Vanessa Quirk
Zahra Rajabi
Prem Ramaswami
Keerthana Rang
Lauren Reid
Evan Reidel
Marc Ricks
Daniel Riegel
Ninon Rogers
Sandra Rothbard
Jesse Shapins
Michael Shapiro
Joshua Sirefman
Chris Sitzenstock
Lauren Skelly
Kabir Soorya
Eric Sebastián Soto
Snezhana Stadnik
Andrew Staniforth
David Stein
Rachel Steinberg
Jeff Tarr
Ebony Thomas
Bradley Tran
Samara Trilling
Michelle Ha Tucker
Steven Turell
Daniel Vanderkam
Catherine Vargas
Michelle Velez
Ryan Vilim
Megan Wald
Leigh Whiting
Violet Whitney
Landry Doyle Wiese
Jonathan Winer
Andrew Winters
Alexis Wise
John Wittrock

Design and Technical Consultants

Master Plan Design and Engineering Team
Beyer Blinder Belle Architects + Planners
Stantec Consulting Ltd.
EllisDon
Golder Associates Ltd.
Greenberg Consultants
Integral Group
Kerr Wood Leidel Associates Ltd.
MGA | Michael Green Architecture
Mulvey & Banani Lighting Inc.
PUBLIC WORK office for urban design & landscape architecture
RWDI
Turncraft Advisors
W.F. Baird & Associates Coastal Engineers, Ltd.
WSP

Economic Analysis Team
Deloitte
Canadian Centre for Economic Analysis
urbanMetrics

Research and Development
Agritecture Consulting
The Altus Group
Aspect Structural Engineers
Bentall Kennedy (Canada) LP
Big Spaceship
Indigenous Design Studio/Brook McIlroy Inc.
BuildingGreen, Inc.
CadMakers
CHM Fire Consultants LTD
Cushman & Wakefield ULC
Digital Public Square
Doblin, a Deloitte business
Diller Scofidio + Renfro
Dubbeldam Architecture + Design
Energy Profiles Limited
Environics Analytics
EQ Building Performance
Equilibrium Consulting Inc. / Structural Engineer

Eric Miller, University of Toronto, Transportation Research Institute
Forsyth Street
gh3*
Graser & Co.
Heatherwick Studio
HR&A Advisors, Inc.
Idea Couture, part of Cognizant Interactive
Interface Engineering, Inc.
Jackman Reinvention, Inc.
James Urban, Urban Trees + Soils
JE Dunn Construction
Jensen Hughes
JLL Consulting
JLL Research
Peter Kiewit Sons ULC
Michael Wetter - Staff Scientist, Lawrence Berkeley National Laboratory
Lion Advisors for Community and Environment
Maffeis Structural Engineering
MASS LBP
Maximum City
Medow Consulting Inc.
Milborne Real Estate Inc – Residential Sales and Leasing
Mitch Stambler
nARCHITECTS
Nordic Structures
OptiRTC, Inc.
PARTISANS
Reshape Infrastructure Strategies
Ryerson City Building Institute
SE Health
SHS Consulting
Smarter Grid Solutions Inc.
Snøhetta
SPIN
Porous Technologies (Stormcrete®) LLC
Studio F Minus
Studio Ludo
Teeple Architects Inc.
Spanier Group
TWG

Professor Liat Margolis, University of Toronto
Urban Strategies Inc.
Urbanation Inc.
Urban Equation
Vanessa Pfaff, Diversity, Equity and Inclusion Consultant

With special thanks to:
Adam Barker, Charlotte Bovis, Naydia Chantarasompoth, Rebecca Craft, Steven Desrocher, Dina Graser, Elliot Jefferies, Neil Kittredge, Alex Laing, Dan Levitan, Chad Markel, Nancy MacDonald, Peter Macleod, Jon Medow, Richard Moore, Eddie Opara, Vanessa Pfaff, David Sauve, Karl Schaefer, Jenin Shah, Natalie Telewiak

MIDP Acknowledgements

Proposal Production

Project Leadership
Andrew Winters
Laura Fox
Laura Capucilli
Nerissa Moray

Editorial
Eric Jaffe (Editorial Director
    and lead writer)
Sophia Hollander (editor, writer)
Alex Marshall (writer, researcher)
Philip Preville (writer, researcher,
    fact-checker)
Vanessa Quirk (writer)
Anna Zappia (fact-checker)
N2 Communications (copyedit)

Creative Direction
Shaun Harris
Megan Wald

Book Design and Layout:
Pentagram Design Inc.
Chase Booker
Jack Collins
Brankica Harvey
Ryan Hewlett
Xinle Huang
Chantal Jahchan
Eddie Opara
Dana Reginiano
Yo-E Ryou

Book Design and Layout:
Sidewalk Labs
Sam Aitkenhead
Zainab Almehdi
Irina Koryagina
Leigh Whiting

Book Typefaces
Beatrice (Sharp Type)
Gliko Modern (R-Typography)
Neue Droschke (David Einwaller)

Book Paper
Interior: International Paper,
    70lb Accent Opaque Smooth Text
Cover: International Paper,
    120lb Accent Opaque Smooth Cover

Architectural Sketches & Renderings
Beyer Blinder Belle
    Architects + Planners
Picture Plane for Heatherwick Studio
    (Overview pages 116, 118, 120, 122, 124,
    140, 148, 242; Volume 1 pages 58, 60,
    62, 64, 66, 300, 308, 372, 438, 514)
Heatherwick Studio

Illustration and Cover Art
Hedof

Illustration
Emily Taylor

Eyes Closed Portrait Photography
Rich Gilligan

Map Graphics
Ben Oldenburg

Chart Graphics
MGMT.
Pentagram Design Inc.

Photography
David Pike
Jenna Wakani
Mark Wickens