An NCC Group Publication

A Blueprint for Secure Smart Cities

Prepared by:
Matt Lewis

Contents

In this paper we present a high-level blueprint for secure smart cities which includes principles of security by design, threat modelling, secure architecture, strong governance with appropriate policies and processes and various security assurance activities that support testing of discreet IoT components, edge, cloud and backend systems and complete end-to-end systems.

The intended audience of this paper is town and city planners and municipalities who are involved in smart city visions, strategies and rollouts, in addition to end-system and device manufacturers and third party integrators and operators who will build, deploy and operate smart city applications through outsourced business models.

1. Overview

The aim of the paper is to provide all relevant parties with pointers to the correct questions to ask regarding security throughout the entire lifecycle of smart city applications and broader smart city operation.

Smart cities are underpinned by the capture and processing of vast amounts of data via a wealth of distributed sensors and systems to derive intelligence that can be used to improve the quality and provision of services and welfare to citizens.

Town planners, local authorities and municipalities are increasingly looking to smart technology to support their city-specific needs. These needs can differ radically between cities and thus there is no one-off, off-the-shelf smart city technology. Rather, smart cities will grow and develop over time, leveraging technology and innovative solutions that are commensurate with the needs of their respective cities, as dictated by city parameters such as geography, population and budgets amongst others.

For example, some cities may be under pressure to make cost savings, or have growing pressures on transport systems and congestion. Many cities are currently looking at solutions to help them reduce environmental impact and improve air quality due to emissions. Other cities may be looking to adapt to population growth, while others may be looking to lessen the impact on limited local healthcare services due to an ageing population.

Much of the marketing and public collateral around smart cities is expectedly optimistic but with little to no reference to security. Many cities around the world are already investing heavily in trials, testbeds and in many cases operational system deployments such as smart street lighting, smart parking sensors and smart waste management to name but a few.

From a security perspective, each new smart city application brings with it a new set of potential threats.

The evolutionary nature of smart cities means that different applications will interconnect over time as cities seek to generate new datasets and insights through continued innovation, thus creating an ever-growing complex web of interconnected systems that will require assurance of the underlying data confidentiality, integrity and availability in order to minimise the potential for hacking and cyber attacks.

In addition, many smart city applications may directly impact on citizen and visitor privacy, thus introducing the need for strong data protection principles that align with citizen consent and fair use.

Retro-fitting security to such applications may not be possible, or at least incredibly hard, with the cost of getting it wrong potentially at the detriment to city reputational loss, direct impact on health and safety (concerning smart city cyber-physical applications) or potential monetary fines as a result of failures or divergences from legislation and regulatory guidelines. As such, there is a strong need to render Smart Cities and their applications secure by design.

2. Smart city goals dictate security requirements

A key question for smart city planners to ask is what are the goals of a smart city or smart city application?

The goals of a smart city will dictate the underlying security requirements. A key question for smart city planners to ask is what are the goals of a smart city or smart city application? While a lot of smart city operation is about capture and mining of vast amounts of sensor data, the goals of an application that relies on such data could be quite different to other applications that leverage the same data sources.

For example, municipalities may wish to consume sensor data to use in AI model training for prediction of crowd gathering or patterns of crowd behaviour; or an application may simply be to monetise the data captured by sensor networks by selling it, or insights gained from it, to third parties.

The security requirements of these two different applications that might leverage the same underlying sources are quite different – the former might have stronger availability needs, especially if used in any real-time applications, while the latter may have stronger integrity needs if the data is being sold for profit (thus demanding accuracy and integrity in its manifestations), and potentially confidentiality aspects that need consideration should the data contain any personal information relating to city citizens and visitors.

The Importance of Strategies and Feasibility Studies

Smart city strategies and discreet feasibility studies on all proposed smart city applications are strongly recommended, with focus on privacy and security.

On the surface, many smart city visions and applications present compelling cases for progressing with their rollout due to perceived benefits, however as noted in [1]:

Implementing a Smart City initiative/ strategy consists of a complex set of tasks and politics that are difficult to resolve in practice and require multi-stakeholder negotiations, policy changes and investments to address.

This is in addition to most municipalities not typically possessing the in-house skills and personnel (e.g. cyber security specialists, data scientists, radio network engineers etc.) to properly strategise and plan secure smart city applications, rendering them reliant on third-party advice and guidance which may or may not be adequate with regards to a smart city's security requirements [2].

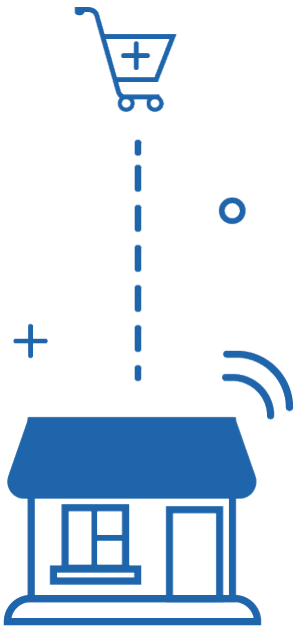

3. Smart city governance and risk ownership

Robust security begins with robust governance. Key to the success of the security of any project, smart city or otherwise, is a clear governance model with defined roles and responsibilities around risk management and ownership.

Robust security begins with robust governance. Key to the success of any project, Smart City or otherwise, is a clear governance model with defined roles and responsibilities around risk management and ownership.

In the corporate world governance structures are typically well-defined and understood – e.g. C-Suite members with a Chief Information Security Officer (CISO), Chief Technology Officer (CTO), Data Protection Officer (DPO) etc. aligned to an ongoing and maintained Information Security Management System (ISMS). In the world of smart cities where potentially many different applications are rolled out and maintained or operated by disparate third parties and system integrators, or outsourced entirely, the governance model is less clear, begging questions such as:

- Who's the CISO for the smart city?
- Who owns the risk if smart city applications become subject to cyberattack?
- Whose responsibility is it to respond to cyber incidents affecting smart cities?
- Who's the data controller for each smart city application?
- What types of personal data are being captured, transmitted and stored by smart city applications?
- What is the security policy for the smart city?
- Who sets the minimum standard for encryption, authentication and logging across smart city applications?

There are many more questions [3] – the aim of the select few above is to demonstrate how much there is to think about and establish by way of a robust smart city security governance model - much assumption may creep in as to who is responsible for security, when and where.

Even if municipalities expect that outsourced solutions will address the necessary security requirements, unless such requirements are captured and prescribed as contractual obligations, there is a risk that the security assurances needed and expected simply may not present themselves.

Establish a Smart City Security Working Group

As part of smart city governance, municipalities should establish internal security working groups, with routine meetings commensurate with the scale and security requirements of the underlying city and its smart applications [4]. The smart city security working group should establish a governance model, roles and responsibilities with regards to risk management and ownership and seek to understand the security and privacy implications of all smart city applications so as to understand the necessary security controls and risk management strategies needed per application.

Smart city security working groups should non-exhaustively include town planners, legal counsel, and departmental representatives where relevant to specific application rollouts (e.g. utilities, transport, sanitation etc.). CISO and DPO roles are naturally key to such working groups but as previously noted, may not exist in the context of smart cities. For those municipalities with ambitious and well-funded smart city visions, it is highly recommended that CISO and DPO-type roles are created to drive and maintain the required security governance.

Smart City Standards

Internationally recognised standards can help municipalities in their smart city security strategy and governance pursuits. For example, the International Organisation for Standardisation (ISO) provides frameworks to help cities structure what it means for them to be "smart" and how to go about achieving smart city goals [5]:

- ISO 37101:2016 is aimed at city leaders and is a management system concerning areas that need to be addressed in order to derive "smartness" through system and technology change, including environmental management, citizen health and well-being, governance, mobility etc.
- Specifically for security, ISO/IEC 27001 and ISO/IEC 27002 for information security management systems exist to support municipalities in addressing security and privacy issues
- ISO/IEC 30182 – the smart city concept model provides guidance for establishing a model for data interoperability which is particularly pertinent for applications that will likely capture and share a lot of data across smart cities
- ISO/IEC 21972 - Information technology provides an upper level ontology for smart city indicators
- ISO/IEC 27550 - Information technology – Security techniques supports concepts around privacy engineering
- ISO/IEC 27551 - Information technology – Security techniques sets out requirements for attribute-based unlinkable entity authentication which also supports concepts of privacy by design

Adoption of internationally-recognised standards is certainly recommended to support and guide smart city development

Adoption of internationally-recognised standards is certainly recommended to support and guide smart city development, but also to ensure consistency across different municipalities – the benefits of adopting standards early on will pay dividends at such a point that different municipalities may wish to share data and interconnect systems.

Monitoring and Adapting to Changes in Policy, Regulation and Legislation

Technology and threat landscapes are constantly changing; this means that policy, regulation and legislation regarding technology use is always at risk of change either in support of underlying systems, or to their detriment.

Smart city security governance therefore demands a routine monitoring of relevant policies, regulation and legislation; while municipalities may need to react quickly to any changes within those domains, particularly where privacy forms a pertinent factor of an underlying application or system. For example, in January 2020, the European Commission revealed that it is considering a ban on the use of facial recognition in public areas for up to five years [6].

Contemplation of such measures has certainly been a result of public mistrust and backlash in the rollout of pervasive CCTV and surveillance technologies across cities. The implications for affected municipalities within the EU could be severe - for example, a facial recognition monitoring solution currently being rolled out across a city at significant cost to the taxpayer might have to be stalled, while in another city, an operationally-effective facial recognition system may need to be halted, potentially impacting on law enforcement operations or functions such as lost or missing person detection across cities.

A robust smart city security governance model should help minimise the impact of any policy, regulation or legislation change.

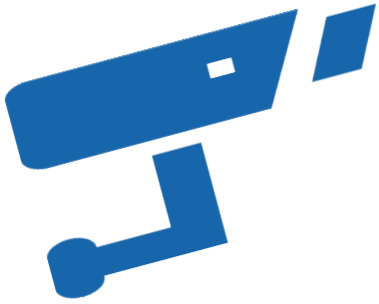

4. Smart city threat modelling

At the start of smart city application design, municipalities should threat model the application so as to enumerate vulnerabilities or absence of security controls such that suitable mitigations can be designed into solutions, or at least be risk-managed.

Smart city applications may be developed and/or make use of potentially vulnerable, legacy technologies, while new technologies and the nature of their use may present new vulnerabilities.

Additionally, the deployment of smart city systems in public spaces presents challenges around physical security of underlying hardware and components, while the nature of the smart city applications themselves, particularly if cyber-physical in the sense that they impact on the physical world, and could present risks to the health and well-being of citizens and visitors.

As such, the exercise of threat modelling [7] smart cities and their applications is paramount.

Threat modelling helps to reveal assumptions about security in systems. It helps to identify trust boundaries within systems and where vulnerability or weakness might manifest itself, and via which type of threat; be that denial of service, breach of confidentiality etc.

At the start of smart city application design, municipalities should threat model the application so as to enumerate vulnerabilities or absence of security controls such that suitable mitigations can be designed into solutions, or at least be risk-managed.

Smart City Threat Actors

When threat modelling it's useful to understand who the threat actors against a target system (or smart city in this instance) might be.

Understanding threat actors and their likely motivation will help inform the security requirements and controls needed to underpin smart city applications.

Here are a few examples of potential smart city threat actors and their motivations:

Citizens Committing Fraud

Where smart city applications may have some financial aspect, or provision of paid-for services, this may open up attempts by some citizens to bypass payments or to fraudulently gain access to services. For example, smart energy meters might be tampered with so as to avoid or minimise cost of energy bills.

Consider a smart parking application where parking bays inform a parking app which bays are free and available for use – a citizen who has a preferred, convenient parking spot, might tamper with the parking bay sensor to make it report that the bay is always occupied (when indeed it isn't), thus allowing the fraudster citizen to arbitrarily use their preferred parking spot whenever they wish.

Organised Crime

Organised criminal gangs may seek to exploit smart city applications for financial gain, or theft of goods and services. For example, a ransomware attack executed across a backend smart city command and control centre could cripple normal operations of a city and potentially impact on health and safety. Without backup systems, municipalities may be pressured into paying large ransom sums which may or may not result in resolution.

Understanding threat actors and their likely motivation will help inform the security requirements and controls needed to underpin smart city applications.

Consider also smart city applications where GPS might be used for driver, vehicle and goods tracking – cyber criminals may seek to spoof or fake coordinates to redirect vehicles to a convenient hijacking location for example.

Cyber Criminals and Hackers

Similar to organised crime but not necessarily financially driven, hackers for hire, or general curious hackers may seek to compromise smart city applications simply for the technical challenge and kudos. Perhaps they might want to play around with traffic light systems for amusement, or perhaps expose large datasets of citizens' personal information that they may have found following a successful breach of a backend smart city application database.

Hacktivists

Hacktivists are activists with a cause and with technical skills or access to those with the skills necessary to perform hacking activities which helps to raise awareness of their cause. For example, environmental activism is currently widespread globally owing to concerns around climate change, and specifically the impact that many large cities have on the environment.

Conceivably activist groups who already use disruptive tactics around cities to raise awareness of climate change, might look to use cyber attack methods against smart city applications for the same purpose, such as disrupting smart traffic light systems and dynamic routing applications to cause gridlock and impact on mobility during periods of activism.

Cyber Terrorists

Terrorist actions usually involve inflicting physical harm or death upon innocent civilians. Mostly this tends to be via physical actions involving violence and weaponry. However, where smart city applications have cyber-physical components that could (if abused) cause injury or death, then cyber terrorists may look to exploit these vectors as they could provide the means for performing high impact, large-scale terrorist acts.

Hostile Nation States

Depending on threat and geo-political landscapes, hostility might arise between nations, fuelling actions that seek to undermine a nation's economy and ability to operate effectively. Nation states will also possess large budgets and have access to highly-skilled individuals allowing for development of all manner of smart city exploitation and disruption techniques, whether large-scale jamming of wireless sensor networks to impair their operation, or disruption of entire electricity grids [8]; with the latter, the impact could be huge on those cities that have rolled out many smart city applications that are strictly dependent upon uninterrupted electricity supply for operation.

Industrial Espionage

The vast amount of technology distributed around smart cities, in addition to the vast amount of data and useful insights generated by that technology might increase the threat of industrial espionage, whether that be attempts at stealing and reverse engineering technology to copy or steal intellectual property, or attempts to gain unauthorised access to big data stores so as to leverage the data for some form of commercial gain.

Example Smart City Threats

The table below, borrowed and adapted from [1] shows just some of the threats that might exist against smart city technologies within different domains.

The aim of this table is simply to show the different types of threat that need to be considered and addressed, and how they will differ across technologies and use-case domains – certainly the number of threats per domain and technology are in no way exhaustive:

The aim of this table is simply to show the different types of threat that need to be considered and addressed.

Domain	Example Technologies	Potential Threats
Government	Urban Dashboards	Sensor networks are tampered with so as to generate inaccurate data which misinforms information displayed on urban dashboards.
	e-Voting	Vote data is tampered or votes are fraudulently generated so as to manipulate democratic decision making within cities.
Security and Emergency Services	Digital Surveillance	Hackers breach the back-end watch list database of a surveillance system in order to remove members from the watch list. Hackers place an innocent citizen's face onto a compromised backend watch list which results in a false arrest, resulting in privacy and data protection concerns, likely resulting in negative public reactions.
	Predictive Policing	Sensor networks are manipulated with fake events in order to make parts of a city appear overly busy, thus redirecting on-the-ground police away from areas that actually need them present.
Transport	Smart Travel Cards	Organised crime or hackers seek to generate fraudulent cards to avoid payment or to sell fake cards at cheaper prices on black markets. Hackers cause a denial of service to the ticketing infrastructure, meaning automatic barriers don't open which results in crowd build-up around city stations.

Domain	Example Technologies	Potential Threats
Transport	Dynamic Road Signs	Hackers or hacktivists manipulate messaging on dynamic road signs with false information, potentially resulting in traffic chaos on city roads.
	Adaptive Traffic Lights	Hackers or hacktivists gain unauthorised control over traffic lights and manipulate them to cause chaos on city roads – manipulation could cause harm or death through traffic collisions if lights are made to show green when they should show red, and vice-versa.
Energy	Smart Street Lighting	Hackers may gain unauthorised control over street lighting, allowing them to turn them off when they are needed at dark, potentially resulting in an impact on public safety due to poor visibility.
	Smart Electric Vehicle Meters	Attackers seek to manipulate charging meters and data to avoid paying for charging.
Waste and Environment	Smart Bins and Dynamic Waste Collection	Attackers block or manipulate the signals generated by smart bins to inform backend systems that they are full and need emptying, resulting in build-up of waste, full bins and likely increased littering as a result. Smart bins are compromised by hacktivists who manipulate the bin heat sensor to make it think it's on fire, which results in an automatic callout to the fire service, wasting the fire service's time and resource.
	Pollution Sensors	Organisations known to be heavy on emissions employ hackers to manipulate readings on nearby pollution sensors in order to avoid regulatory fines and penalties.
	Automated Flood Defences	Cyber terrorists gain command and control of automated flood defences and open or jam them in ways that result in actual flooding of urban areas.

For each smart city application under consideration, municipalities should engage cyber security specialists to assist in threat modelling workshops with all relevant stakeholders.

Domain	Example Technologies	Potential Threats
Health	Connected Hospitals	Cyber criminals launch a ransomware attack across connected hospitals, causing major disruption to critical administration of treatments and drugs, possibly resulting in loss of life.
	In-Home Medicine Dispensing and Telecare	Hackers compromise in-home medicine dispensing apparatus and force equipment to administer lethal doses of drugs.
Buildings	Building Management Systems (BMS)	The Heating, Ventilation and Air-Conditioning (HVAC) systems of smart buildings might be attacked and forced to make office environments too hot or cold, forcing evacuations and impacting negatively on business operations and overall city economy.
Homes	Smart Meters	Some citizens attempt to tamper with smart meters to manipulate their readings so as to reduce energy bills.
Citizens, Visitors and Tourists	Broadcast Messaging	A smart city application which has capability of pushing out SMS, Bluetooth and Wi-Fi messages to mobile handsets for citizen information purposes might be compromised by attackers, who might broadcast misinformation such as instructions to evacuate a city, resulting in pandemonium.

Assurance in the Data Supply Chain

Big data is the essence of smart cities. The accuracy and integrity of this data is paramount in terms of the quality and impact of delivery of services and insights that might build on such large data sources. In reality, municipalities may have little to no control over how and where data is captured, transmitted and stored, meaning that achieving high assurance in the so-called data supply chain could be difficult across an entire smart city and its associated applications.

For example, a third-party system integrator may be responsible for data captured from sensor networks; however, that data may then be passed to other third parties who perhaps use cloud providers for storage and processing of the data to mine it for insights and intelligence. Any insights might then be passed on to further third parties who consume those insights and present back to municipalities or citizens in some digestible form or dashboard reporting. In this albeit simple example, the potential for data corruption or manipulation throughout the entire data supply chain is high. Where municipalities have little to no control over their city data, assurances may therefore need to be sought from third parties, whether through demonstrable security testing and assurance activities performed by the third parties, and/or contractual obligations and Service Level Agreements (SLAs) associated with the security of the data and its supply chain.

Codes of Connection (CoCo) [9] should be considered as part of threat modelling and secure design of data supply chains. CoCos involve setting and establishing baseline security controls to be implemented when connecting (ideally accredited or assured) systems to form new pipelines of data flow. For example, two neighbouring smart cities may wish to exchange data to improve public services that traverse both cities - CoCos would help in this instance by ensuring that both parties properly plan for the nature of the connection of systems and the data that will be shared, so that adequate controls can be put in place to assure the data supply chain.

Achieving high assurance in the so-called data supply chain could be difficult across an entire smart city and its associated applications.

Data Poisoning and Second-Order Attacks

Relevant to data supply chain security is consideration of the potential attacks associated with data poisoning and second-order attacks.

Data poisoning attacks will involve deliberate acts by attackers in manipulating smart city data, either to cause general disruption, or to specifically influence decisions that might be made based on modified (poisoned) data. Encryption (see later) plays a strong part in minimising the potential for data poisoning attacks, however even encryption can only go so far in that at some point, data must be decrypted in order to perform operations upon it – attackers may be able to manipulate data at the point of decryption (e.g. if they have gained unauthorised access to backend systems), thus data poisoning attacks are not exclusive to the sensor network only.

Municipalities also need to be aware of the threats associated with second-order attacks against data. Such actions can be quite subtle and hard to detect, but essentially involve poisoning or corrupting data at some point within the data supply chain, so that when that data is used in a secondary (or beyond) process, the underlying poisoned data influences a decision or effect that an attacker intended. For example, big data collection across large sensor networks might be performed in order to create a large dataset of training data, to be used to train a machine learning model that can then be used to make future predictions or classifications on future data captures. If attackers could influence the training data with poisoned or manipulated inputs, then when that data is used to train new models, the model is created in a way that it perhaps misclassifies in ways that benefit attackers in some way, or simply causes disruption or provision of incorrect information to citizens.

Threat modelling the data supply chain is therefore a recommended key activity to support secure smart city design.

Threats Involving Real-Time Responses and Automated Decision Making

The 'smart' aspect of many smart city applications usually denotes some level of automatic decision making or real-time adaptive responses to certain events.

For example, there are various traffic light solutions that consume real-time data from vehicle detection sensors and automatically adapt the phasing of lights to accommodate for variances in traffic load.

When threat modelling systems that employ a level of real-time adaption or automated decision making, it's important to draw out what the consequences could be based on specific threats, since there may not be an easy human-based override mechanism upon cyber attacks against such systems that might be causing severe operational issues as a result of automatic decisions made on manipulated data.

The same principles apply to smart city applications that might leverage AI or machine learning in ways that implement algorithmic autonomy (e.g. through autonomous vehicles). Municipalities should understand who is responsible for any erroneous or negligent decisions made by autonomous systems as a result of manipulated or compromised input data.

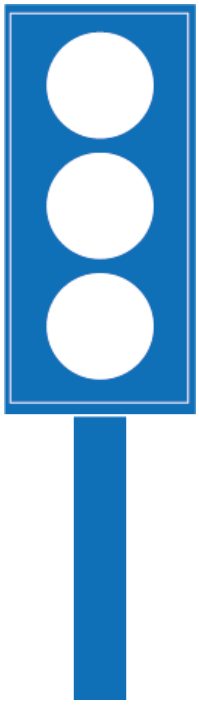

5. Smart City Secure Design

During or shortly after threat modelling exercises regarding smart city applications, principles of secure design should be followed so as to ensure that security requirements manifest themselves as in-built controls and features within production systems.

During or shortly after threat modelling exercises regarding smart city applications, principles of secure design should be followed so as to ensure that security requirements manifest themselves as in-built controls and features within production systems.

Engage with Citizens from the Outset

As part of smart city secure design, to allay citizen security concerns and to understand what citizens want by way of a secure smart city, it is highly recommended that citizens are engaged early on as part of consultations and oversight committees. Municipalities should encourage security-focussed town hall meetings and engage citizens of all ages to facilitate transparency around smart city application plans and rollouts. As written in [1], *The smart city needs to find an effective means to shift citizens from users and consumers to active stakeholders in order to become more democratic in nature.*

Certainly an example where inadequate citizen engagement has hampered smart city plans has been seen in Toronto, Canada [10], where various concerns around privacy and data in relation to Alphabet-owned Sidewalk Labs' proposed smart city research and initiatives has resulted in delays and citizen campaigns seeking to block related deals and developments.

Privacy

Smart city applications may consume data that is personal in nature by default, such as GPS location data relating to citizen whereabouts within cities (as consumed from their mobile devices), facial imagery captured by surveillance and facial recognition applications or device identifiers such as Bluetooth IDs of mobile handsets captured by crowd monitoring applications. Indirectly, there may be ways to derive personal information on citizens as a result of aggregation of multiple data sources, such as correlation of surveillance, GPS and device ID monitoring data. Much public concern around privacy in smart cities relates to what data is captured by what systems (specifically those captures that occur without notification or consent), and also who has access to all captured data and perhaps the capability to aggregate in ways that allow for entire tracking and surveillance of citizens throughout their city visits. Issues of function creep are also pertinent here, where citizens may have concerns that over time, municipalities, law enforcement and intelligence agencies may access and use the rich telemetry generated by smart cities in ways that potentially compromise their privacy.

Data Protection Impact Assessments (DPIAs)

Municipalities should perform DPIAs [11] on all proposed smart city applications, particularly those that are overtly privacy-impacting. The DPIA will assist in identifying and managing privacy risks arising from new projects and proposed systems, and can be used to demonstrate considered thought by authorities when engaging with citizens on smart city applications.

Consent

As noted in [1] with regards to smart cities, *"Issues of notice and consent are difficult to deal with in practice as people move through environments saturated with networked sensors, actuators and cameras that generate huge volumes of data about them".*

This poses a challenge for smart city applications that do overtly capture personal information, and perhaps large volumes of such data in real-time 24/7 applications.

Municipalities should therefore consider all options for obtaining and recording city citizen and visitor consent [12]. In some regards, citizen consent may be slightly easier to achieve through awareness campaigns and websites used to inform those citizens on, say the local electoral register, or those registered by local councils for council tax and use of council services. Citizens could be invited to visit websites whereby they are provided with information on proposed smart city applications and their potential impact on privacy vs. potential improvements in safety and/or delivery of services, such that they could then make informed decisions on whether or not they consent to use of their data. However, managing lack of consent may simply not be an option in terms of how some smart city applications operate and as such, methods for anonymising captured data will need to be investigated.

Where consent can be engineered into applications, it should always be opt-out by default, as opposed to opt-in by default.

Some technical solutions may be applicable for at least notifying city citizens and visitors of surrounding applications that might be capturing their data, if not obtaining their explicit consent. Examples might include 'push' technologies that send messages to passing mobile devices such as Bluetooth broadcasts – conceivably a message could be pushed to cell phones in specific city areas such as "You are entering a city zone which captures device identifiers to monitor crowd volumes. Do you consent to us using an anonymised version of your mobile device identifier during your visit?"

Where consent can be engineered into applications, it should always be opt-out by default, as opposed to opt-in by default.

Withdrawal of Consent

Under data protection regulation and legislation, applications that accept and process the consent of citizens and visitors will need to offer the ability to withdraw consent. Such withdrawals should not be onerous and easily actionable, allowing individuals to request that any personal data or identifiers captured during a consenting period can be easily identified and deleted from all applicable systems, and that future capturing of such data is ceased until such time that consent might be reissued by the individual.

State Powers

Depending on country, state and city, there may be applicable and usable special powers or legislation that allows for capture of specific data types (such as surveillance systems capturing facial images) to support aspects of law enforcement and national security. Such systems may be built and operated by law enforcement agencies themselves, while other systems may be owned and operated by municipalities where at a state and national level, special warrants might be used to access specific data sets captured by the city's applications. To earlier points on citizen engagement, they should at least be informed of any law enforcement or state-run systems within cities, and of the process that government agencies may be able to follow to access specific data sets, and what that data might be.

Physical Signage and Notices on Surveillance Applications

For surveillance systems that perhaps under state powers or law enforcement do not require consent, some consideration could at least be made around use of physical signage in specific city areas or zones, informing on the presence of surveillance cameras, such as "Facial recognition and CCTV applications operate in this area", or "Automatic Number Plate Recognition (ANPR) operates in this area".

Data Anonymisation/Pseudonymisation

Anonymisation or pseudonymisation of smart city data could mitigate many issues around consent both for citizens and for city visitors. Various techniques exist for anonymising or pseudonymising data in ways that don't affect the utility of the data but that preserve anonymity. For example, a crowd monitoring sensor might capture the Bluetooth IDs of passing mobile devices; given the application only needs to monitor volume, there should be no need to capture and store entire IDs; thus the IDs could be stripped of some of their values (such as the first three digits), or pseudonymised so as to replace all digits with different values.

Municipalities should however exercise caution with data anonymisation activities; particularly where data is stored and

what it might offer when aggregated or correlated with other data sets – there may exist methods for de-anonymising (or re-identification) of data when correlating with other data sets, thus rendering anonymisation methods ineffective. Data anonymisation specialists should be engaged from the outset when planning for data anonymisation and pseudonymisation operations on smart city data.

Subject Access Requests (SARs) and Freedom of Information (FoI) Requests

When designing smart city applications that capture and process personal data, municipalities should plan mechanisms for handling SARs [13] and Freedom of Information (FoI) requests. Under data protection legislation in some countries, SARs provide citizens with the right to ask an organisation whether or not they are using or storing personal information about them. Citizens can also request copies of their personal information.

On a broader scale, FoI [14] legislation can provide public access to information held by public authorities, who are obliged to publish certain information about their activities, and where members of the public are entitled to request information from those authorities. Citizens of smart cities, particularly those with concerns around privacy and surveillance, may therefore seek information on specific applications and the types of data that they capture, process and store. Local authorities will need to engineer mechanisms for dealing with FoI requests and should seek to provide transparency from the outset on smart city applications through provision of relevant information on public websites for example.

Physical Security and use of Street Furniture

Much of smart city technology and components will be deployed as physical hardware in public spaces; be these small sensors, wireless gateways that consume sensor data and route to the Internet, telecoms infrastructures such as Radio Access Networks (RANs) and cell towers, and edge-computing devices such as high-end servers performing high-intensity computation on sensor network data, to name but a few.

Attacks against physical technology components need to be considered, since the ability to physically access such components could provide mechanisms for threat actors to manipulate sensor data on the devices themselves. Such attacks would not require wireless network data interception and decryption of data for manipulation, since the data could be manipulated directly on the hardware before transmission.

Components may either be directly integrated into a city's physical infrastructure, such as smart parking sensors embedded into concrete in the ground, or connected to street furniture such as streetlamps, bus stops, and traffic lights for example. In addition, the street furniture components themselves might be the 'smart' technology component, such as smart bins and smart bollards.

As part of smart city secure design it is important to understand the implications of smart technology being accessible in public spaces. This increases the potential for physical tampering of technology for example, which may not be easy to monitor in the case of an application utilising hundreds or even thousands of smart sensors distributed around a city; it would not be possible to use CCTV for example to monitor the physical security of all such components.

Physical security principles or controls should therefore be designed and built into solutions where possible.

These might include:

- Use of unattainable heights – e.g. placing sensors or technology components high up on street furniture such as streetlamps provides a barrier to easy tampering compared to say sensors embedded at street level

- Strong casing with locks – use of robust weatherproof cases with strong physical locks can help secure critical components such as gateways, though the associated physical key management processes could add a layer of complexity (where are physical keys stored, and who has access to them?)

- Electronic tamper detection mechanisms – there may be ways to engineer tamper detection into

technology, such as automatic alerting upon detection of case opening. This however might not be feasible due to increased cost in technology and monitoring process required to render the control effective, but should at least be considered for critical components

» Anti-tamper labelling (stickers) – while fairly crude in nature, and demanding of human inspection which could be onerous due to the need for routine inspection, the use of tamper stickers on devices might at least provide a retrospective view of potential tampering of devices, allowing for any onwards investigation around tamper-based activities on affected devices

Network Architectures and Topologies

Different network architectures and topologies may be available for different types of application [15]. Different topologies will present different types of threat, and understanding this at design stage is key so that any necessary controls or redundancy to support availability can be architected into solutions [16].

Hub and Spoke

Many sensor-based networks will operate under a hub-and-spoke architecture, meaning that multiple sensors will wirelessly connect to a central hub (gateway) which will route data onwards accordingly, such as to Internet-facing cloud applications for processing. A key security consideration for hub-and-spoke architectures is availability – if there is only one hub and its availability becomes disrupted (whether through active attack or inadvertently through power failure), then the entire system experiences an outage. The criticality of sensor-based networks should dictate the availability requirements on those networks – where some applications may be able to tolerate occasional downtime, other more real-time critical systems may not be able to afford even seconds of unavailability, thus demanding at design stage, consideration for architecting redundancy and failover controls into proposed systems.

Peer-to-Peer (P2P)

Some smart city applications may operate on a P2P basis, whereby ad-hoc or meshed networks may dynamically create themselves upon certain conditions. Conceivably smart city mobile device applications might be developed with P2P functionality, whereby citizen mobile devices auto-connect to fellow citizen's devices that were opted into the same smart city application in order share and route data between them.

P2P adds a layer of complexity to management of data flows, while it might also increase the potential for manipulation (or exposure) of data as it traverses devices and technology that are not under the control of the city.

Edge Computing

Edge computing provides high performance computation power closer to where it needs to happen, as opposed to relying on cloud or backend systems. For example, a smart traffic light system may need to process large volumes of camera imagery and sensing data from around a city, in real-time, in order to dynamically change the phasing of traffic lights so as to maintain optimum traffic flow. Using edge computing reduces latency in these situations, but presents issues around potential physical access or manipulation or disruption of devices that are attached to street furniture.

Cloud

Cloud computing will play a huge role in smart cities. Much of the data captured by the myriad of sensors around a city will typically find its way to cloud systems whereupon it can be processed and mined for new insights, and stored in large volumes accordingly. Cloud also offers a number of features such as availability which is a common key requirement for smart city applications. A key consideration at design stage for cloud components of smart city applications is how will data be ingested, and how will it be accessed?

Likely various Application Programming Interfaces (APIs) will be exposed on the cloud which will provide programmatic methods to consume and access data. Data access requirements might solely be for authorised city officials, legitimate third parties (e.g. where data access and insights may be monetised in some way) or open [17] to anyone as part of open data strategies, aimed at allowing citizens and entrepreneurs to access and innovate over the data captured by the underlying system.

The criticality of the application and the nature of any monetised aspect will dictate the security assurances needed on access methods such as APIs.

Static vs. Mobile Sensors

While many smart city sensor-based applications will be static, in that the sensors will be deployed and remain in situ, some applications may make use of mobile sensors, whereby the geo-location of sensors changes over time. An example here could simply be citizen mobile devices that move around with their owners and capture and transmit various telemetry data via their underlying technologies (cameras, microphones, accelerometers etc.). On a different scale, mobile sensors could relate to tracking sensors in freight movements around a city, or even drone or autonomous ground vehicles capturing some sort of data for future processing.

The threat model of applications will change significantly depending on whether underlying sensors and technologies are static or mobile.

Mobile Applications

Many smart city applications will make use of citizen mobile devices for the purpose of presenting data to citizens, and/or retrieving data and telemetry from them as part of some broader smart city application. The presentation and/or extraction of data from mobile devices may occur via an app, which would communicate with backend systems over cellular or Wi-Fi networks.

Municipalities should ensure that the security requirements of mobile applications are captured and addressed during the design stage, particularly where the data captured, stored and processed by mobile applications may be personal in nature regarding the owner of the underlying mobile device.

LPWAN Technologies

Many of the sensor-based wireless networks that will underpin smart city applications will make use of Low-Power, Wide Area Network (LPWAN) technologies. There are a number of LPWAN technology alternatives available (LoRaWAN, NB-IoT and Sigfox to name just three), each with their respective advantages and limitations with regards to security. There are therefore fundamental design considerations around LPWAN and security.

Battery Life

LPWAN sensors will be low cost, and likely will not have access to a permanent power source meaning they will operate on batteries. LPWAN concerns the use of discrete, fairly infrequent data transmissions regarding sensing inputs, such that the battery life of the sensors can last (perhaps up to ten years) a long time before needing replacement. Where gateways or backend applications need to communicate with sensors, an increase in the consumption of finite battery power on the sensor will be experienced. This means that heavy computation operations such as software updates may not be feasible on sensor-based networks with limited battery life. It's possible that in some applications, a sudden need to update thousands of sensors with a security patch for a reported critical vulnerability may not be possible.

Unlicensed Spectrum and Line of Sight Needs

Some LPWAN technologies such as LoRaWAN operate in the unlicensed wireless spectrum, meaning that there's little governance over activities within those spectrums which could open the potential for wireless jamming to cause disruption. LPWAN systems may require line of sight and strong signal strength between sensor and gateway. In cities that are densely built and populated, achieving good line of sight and avoiding overlapping communications in unlicensed spectrums could prove problematic for availability. In addition, some LPWAN technologies by virtue of being unidirectional from sensor to gateway do not present any guarantee of transmission or receipt of messages, meaning data loss is a risk.

Range also needs to be considered for LPWAN technologies. It is important to understand what distance ranges will need to be handled for specific technologies as this will drive the decision for specific technology use. For example, if only meters are required between sensors and a gateway then perhaps Bluetooth would be appropriate, as opposed to the need to transmit data from sensors to gateways across tens of kilometres of a city.

Optimised Deployment

Related to line of sight considerations, the actual placement of sensors and their ability to maximise broadcast range will need to be considered. Poor placement of sensors (and gateways) may severely impact on the performance of a system. Optimised deployment strategies will therefore need to be determined and will be unique per application – such strategies need to include aspects of ease of access for maintenance, in addition to physical security and secure street furniture deployment.

Sensor and Gateway Authentication and Authorisation

Different LPWAN technologies will offer different mechanisms for sensor and gateway authentication and authorisation to the network – i.e. there will likely need to be controls configured to deny rogue or ad-hoc sensors the ability to arbitrarily join sensor networks, otherwise the ability to do so could result in sensors injecting erroneous or malicious data into smart city applications.

Encryption

Lack of encryption, or weak encryption mechanisms open up the potential for data interception and replay attacks against LPWAN systems. Encryption will therefore be required to minimise the potential for attackers to intercept and/or manipulate LPWAN communications. However, small form-factor sensors manufactured to low cost will not likely possess the capability to perform high-end encryption. Certain concessions may therefore need to be made regarding LPWAN encryption such as the strength of algorithm and length of encryption keys used.

With encryption also comes the need for secure key management – who creates the keys, how are they configured on devices, and who has access to the keys (are they backed up somewhere)? Municipalities will need to ensure strong governance around encryption and key management, whether that be handled by local authorities themselves, or outsourced to third party system operators.

Some open source software libraries do exist to support encryption in embedded devices with limited computational resources. LibDisco [18] for example condenses state of the art cryptographic protocols and primitives into an extremely compact cryptosystem making it easy to fit into embedded devices.

Community vs. Closed Networks

There are various community-based LPWAN networks that exist around cities, whereby enthusiasts and general contributors offer up gateways into community networks. The Things Network [19] is one such example community built around LoRaWAN.

Municipalities will need to understand at design stage whether they intend to use community networks, or closed networks built and operated by commercial third parties. Where community networks are used, awareness is needed around lack of control over what other applications are built and operated on the same networks, and thus what other data might be passing though shared gateways and the potential impact that data load might have on performance and availability. Similarly, community networks may lack robust SLAs on availability and aspects such as maintenance (e.g. software updates for security issues), rendering roles and responsibilities around performance and maintenance unclear.

Designing for Availability

System up-time will be paramount for many smart city applications; specifically those that operate in near real-time and that perhaps operate cyber physical systems such as traffic lights. As such, availability options will need to be explored and designed to cater for the up-time and performance requirements of underlying applications and systems [20].

Availability considerations should include:

- Redundancy – this is where additional, duplicate critical components may need to be deployed to provide failover options in the event of a device failure. For hub and spoke architectures for example, redundancy can provide assurance at the hub layer. Understanding the criticality of a smart city application will help determine the redundancy requirements upon it.

- Capacity Planning – thought should be given to likely future extensions and additions to deployed architectures, and what effect those might have on underlying device performance. For example, an LPWAN gateway deployed to support arbitrary LPWAN applications will have some limitation in terms of the number of connections that can be supported. If no thought around increase in capacity is made at design stage, then there is risk of deploying systems that cannot easily adapt to increase capacity without being replaced by newer, more powerful options which would involve system downtime for upgrade.

- Backup Power – more specifically for critical systems, there may be a need for backup power sources in the event of power loss in order to maintain availability. Options might include solar PVC panels, backup battery packs or diesel generators. The feasibility of backup power options will be dictated by the physical, geographic placement of powered components.

- Out of Band Communications – in the event of disruption to communication technologies, out of band access and communication methods can be considered to maintain remote management and operations. For example, an LPWAN gateway might include an embedded mobile SIM card so as to provide a backup mechanism to connect it to the Internet via the cellular network. Additionally the same LPWAN gateway might expose an authenticated Wi-Fi access point to allow for proximity wireless connections for management and configuration in the event of an impact on other supported communication technologies and protocols.

Designing for Data Confidentiality and Integrity

Data confidentiality and integrity is achieved through encryption and other cryptographic methods. Municipalities need to understand early on what data is captured, transmitted, processed and stored by specific smart applications. This information should present itself from PIAs as discussed earlier. PIA outputs in addition to data mapping exercises will help understand the entire data journey and lifecycle across systems, thus allowing for appropriate planning of encryption of data in transit (when being transmitted), and at rest (when stored).

As mentioned earlier, some technologies such as low-cost sensors will offer limited encryption capabilities, thus some concessions may need to be made due to technical and hardware limitations. Layers of encryption may also need to be designed – this is common in some LPWAN technologies such as LoRaWAN, which includes encryption mechanisms for data between sensors and gateways, with additional encryption layers for application data that is relayed to backend systems. Consideration is needed at this point on where encryption termination points exist, and thus at what point data is available in an unencrypted state within systems.

Credential and Key Management

Encryption brings with it additional security considerations in order to preserve assurance of encrypted systems – mostly this relates to encryption key and credential (e.g. password) management:

- Encryption Method – static key encryption involves the use of the same key for encrypting and decrypting data, whereas asymmetric encryption (or public key encryption) involves use of multiple keys. Each method has its own advantages and disadvantages and depending on underlying technology, may or may not be possible to implement.

- Key Generation – the method for generating encryption keys needs to be secure and not predictable or guessable, or easily crackable through offline brute-force password attacks. How and where keys are generated needs to be understood early on in design.

- Key Storage and Access – understanding everywhere where encryption keys will be stored, and how they will be accessed and by whom is paramount. The more opportunity there is for keys to be exposed, the more likely the potential for unauthorised use of those keys for decryption of data.

Earlier phases of establishing governance across smart cities should ensure baseline secure policies and procedures are defined for credential and key management. Setting a minimum standard will help provide assurance on the minimum level of confidentiality and integrity that is present across a smart city.

Long-Term Operation

The long-term operational requirements for smart city applications need to be understood as part of the design stage.

The fact that many components will be rolled out across cities and installed through a number of potentially disruptive methods (e.g. digging up roads to install infrastructure components such as smart parking sensors, or installing gateways and antennas on the roofs of tall buildings) means that the cost of rectifying errors in the system design and underlying choice of components could be significant.

> For example, a rollout of say 10,000 parking sensors across a city that fail to properly communicate their state to distributed gateways could be an extremely costly problem to rectify, whereupon revelation of such a mishap could generate citizen and tax payer anger and resentment and reluctance toward future smart city technology adoption.

Smart City Simulation Tools

Town and city planners may already be familiar with various simulation tools concerning traffic, people flow and general urban design. The use of sensor network and smart city application simulation tools is highly recommended at design stage. Use of such tools can help in understanding feasibility of system deployment, including aspects such as optimal wireless sensor and gateway placement (for coverage and reliability) for example.

A number of smart city and sensor network simulation tools exist [21]; both commercial and open source. CupCarbon [22] for example is an open source smart city and IoT wireless sensor network simulator which has a number of features driven through an easy-to-navigate GUI, including the ability to take into account the topology of city buildings as well as the radio visibility and radio propagation within different environments.

Google Earth [23] also offers a number of different ways to asset track IoT devices at specific GPS coordinates and can also be used to examine Radio Frequency (RF) overlays to help plan optimal deployment of wireless sensor networks for maximum availability.

Simulation software might also assist in playing out specific cyber security incidents or scenarios in a safe, non-disruptive environment, such as analysing how simulated data flows react to component failures or compromises.

6. Smart city secure build

There are a number of considerations regarding construction or build of smart city applications in order to ensure that secure design principles manifest themselves as real-world controls and mitigations, rather than remaining as design ideals.

Smart city secure design, threat modelling and simulation helps stakeholders understand the security requirements and controls necessary to realise smart city systems and applications. Security does not however end at design – there are a number of considerations regarding construction or build of smart city applications in order to ensure that secure design principles manifest themselves as real-world controls and mitigations, rather than remaining as design ideals. A few key areas therefore demand attention during the build phase of smart cities.

Test and Lab Environments

By their nature, smart city applications will be large and comprise multiple components and sub-systems and may be dispersed across large distances of large urban areas or entire cities. As such, the production of a representative test or lab environment may not be feasible, meaning that many systems may be built directly into cities without prior lab testing. In this situation, the "city becomes a living lab in which experimentation is practised as systems are developed and refined" [1].

Where possible, test and lab environment (or digital twin) testing is always recommended to support security assurance; however, municipalities may not have this luxury due to time and budgetary constraints. Without the use of a lab or test environments, previous aspects discussed such as public consent need to be considered since it may not be possible to truly simulate or implement a representative test environment.

As such, where possible, simulation software should be used to model sensor networks and interactions, and certainly at the product level (e.g. sensor, gateway), test lab facilities should be used to understand the security of embedded systems, before they are rolled out at scale across entire cities.

Suppliers and Third Party Due Diligence

Municipalities will likely engage third parties throughout the design, implementation and maintenance of underlying smart applications. This could include a number of technology and platform providers and system integrators for example. Ideally any necessary third party due diligence activities will have been performed long before the point of smart city application construction, so as to provide municipalities with assurance in the security posture and practices of their contracted third parties.

When there are multiple parties involved, it is important to clarify roles and responsibilities with regards to security assurance, otherwise there is a danger that all parties assume security is or has been handled elsewhere or by others, or if security hasn't been prescribed by procurement, then it may simply not find its way into architected and constructed systems.

Relevant here is also consideration around aspects such as encryption and key management – as smart city systems that employ encryption are constructed, who manages the encryption keys, where are they stored, and how/when can they be refreshed or modified?

Third party due diligence is an entire subject in itself, however some guidance towards pragmatic approaches in this domain has been authored by NCC Group [24].

Asset Management

The number of assets demanding tracking across smart cities could be vast. Such assets could range from static sensors as part of a large sensor-based network, mobile sensors used in transportation applications, or static critical components such as LPWAN gateways.

Asset management is crucial for security, and particularly when needing to manage security incidents. Not having a handle on what assets exist, where, and associated metadata regarding those assets such as last time updated with software or batteries replaced, could render smart city operations incredibly ineffective.

There are many commercial and open source software offerings to support asset management. The choice of asset management software doesn't need to be expensive or bespoke. Consider Google Earth for example and the screenshot below of a fictitious smart bin LPWAN system. The image shows how a LoRaWAN Gateway and three associated smart bins could be asset tracked by their actual GPS coordinates. A suitable naming convention is used to uniquely identify each asset (bins and gateways), while any amount of associated metadata for each asset could also be tracked within this Google Earth application, such as the last time and date of asset maintenance (software or battery upgrade).

Asset management is crucial for security, and particularly when needing to manage security incidents.

For sensor networks, battery life tracking is quite critical as part of asset management. Consider a sensor network comprising thousands of parking sensors across a city that operate on a 5-year battery lifespan – municipalities would need to ensure a phased and timely approach to replacing batteries to ensure continued availability and operation of the underlying smart parking system. Failure to track such aspects could result in thousands of sensors quickly running out of power in succession, rendering the replacement process chaotic while in the meantime, the underlying smart parking application is rendered ineffective or simply unavailable.

Other relevant data points that could be captured with such asset management processes include which devices connect to which gateways, and which systems interconnect with each other [25].

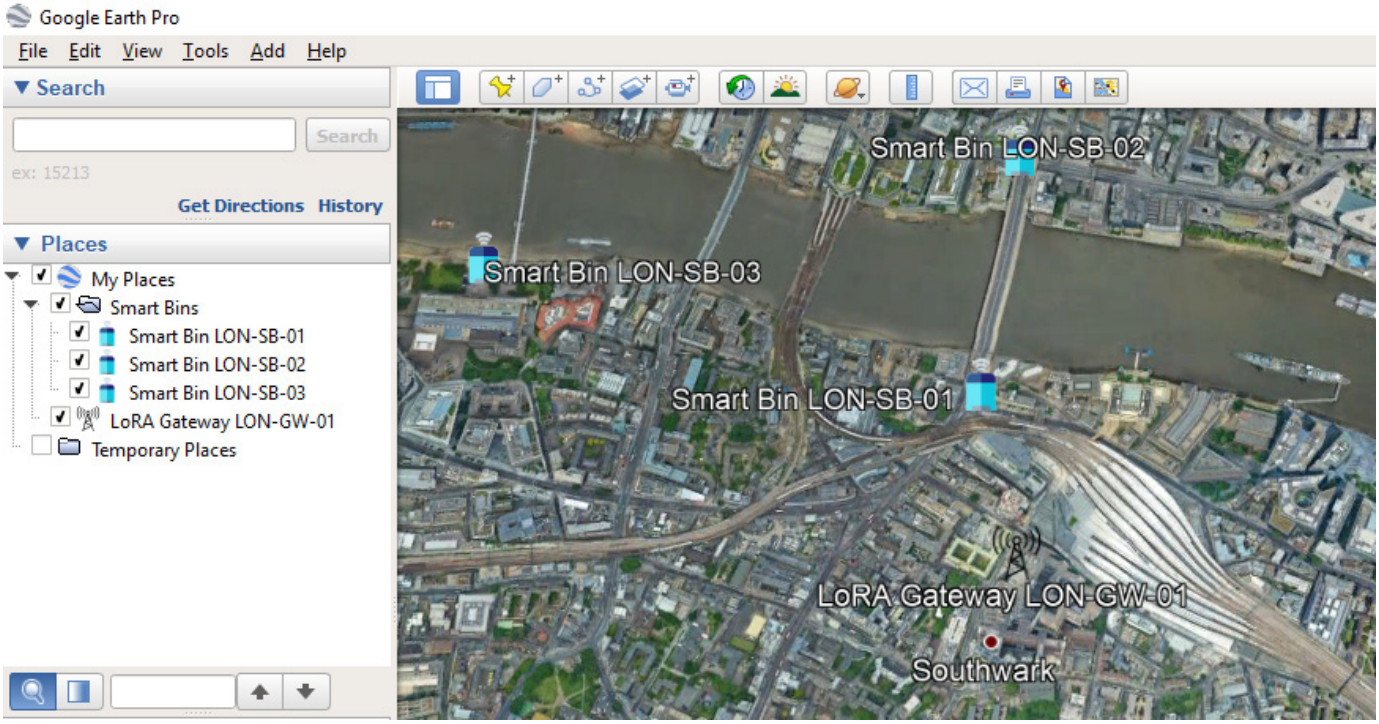

Figure 1 - Using Google Earth as a smart city asset tracking mechanism

Digital Disruptors

Most of this paper has assumed a position where municipalities and town planners are those who are the innovators, designers and implementers of smart city applications. In reality, many smart city applications will likely present themselves without any input or governance by underlying city councils.

The presence of many community networks such as The Things Network for LoRaWAN [19], and the fact that technologies such as LoRaWAN operate in the unlicensed spectrum means that individuals and organisations are free to build and host entire smart city infrastructures at their physical premises (such as LPWAN gateways), as opposed to being deployed on local council assets such as street furniture.

In reality, many smart city applications will likely present themselves without any input or governance by underlying city councils.

Not all smart city applications will necessarily be built and operated under the auspices of city councils. There is therefore much potential for innovators, start-ups and general digital disruptors to surface and deploy smart city systems and applications that provide some benefit or utility to city visitors and citizens.

The ridesharing company Uber [26] is a prime example of a digital disruptor and smart city application that has revolutionised mobility in many major cities around the world, without needing much, if any (in some cases), interaction with or approvals from city authorities.

Likely, many commercial enterprises will seek to deploy their smart city solutions, possibly without any oversight or consultation with local authorities – municipalities may wish to ensure suitable smart city application registration and licensing processes, in order to maintain a handle on different applications being rolled out by third parties across cities, and to ensure some level of governance or oversight over what digital disruptors may be seeking to achieve and whether the data that they are capturing and processing is being done in alignment with relevant data protection regulations, and not in ways that potentially compromise the privacy or safety of city citizens and visitors.

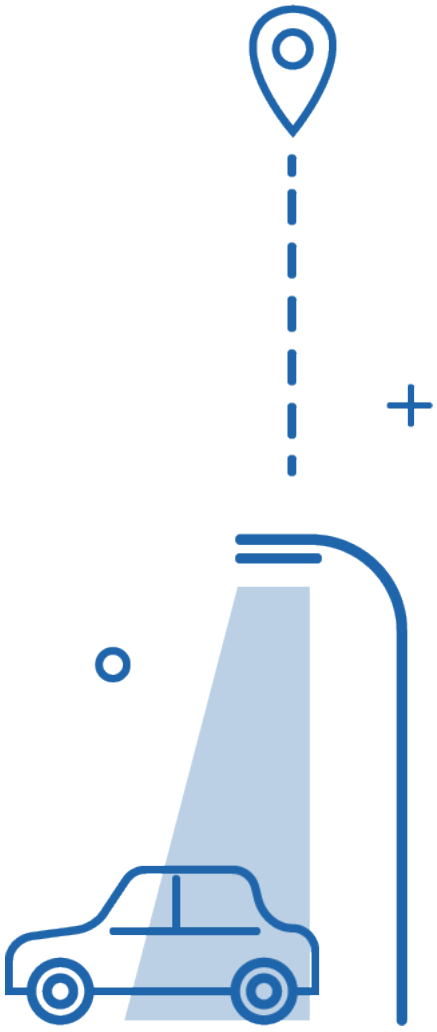

7. Smart city secure maintenance and operation

A successful ransomware attack against a smart city's command and control centre could cripple effective operation of the city until such time that the affected systems could be recovered from such an attack

Once smart city systems and applications have been built and deployed, they enter operational and maintenance phases which demand operational security processes and procedures.

Smart City Orchestration

Smart cities comprising multiple smart applications and systems will likely require some level of centralised orchestration. This would typically take the form of a command and control centre, which ingests data from discrete systems, possibly presenting it to smart city operators in some digested or dashboard form. In addition, such orchestration might include remote control or issuing of commands to remote systems, such as remotely controlling motorised CCTV cameras, or performing some sort of override on a smart traffic light system in order to manage unexpected traffic flows in a city. Orchestration dashboards might be used to simply provide real-time data on the health, or uptime of applications, while in addition, evidence of potential cyber-attacks might be detected by firewall or intrusion detection systems, and presented to orchestration dashboards, allowing operators to enact relevant incident response procedures.

How smart city orchestration will look and operate will likely be radically different and bespoke per city, and will somewhat be dependent on how much of smart city operation is outsourced to third parties. Some municipalities may elect to build and operate their own control rooms, while others may be satisfied with such operations forming part of contracted managed security services.

There are a number of different smart city orchestration offerings available – whether outsourced or deployed in-house, municipalities should satisfy themselves with the security of those platforms since an exploitable vulnerability in smart city orchestration software could potentially provide attackers with full control over a city and its underlying applications.

On a related note – smart city control centres will present a target of great interest to attackers. If such centres are operated by systems that have Internet access for operator email for example, then the usual risks associated with phishing attacks become pertinent. I.e. motivated attackers may be able to socially engineer smart city operators through phishing attacks in order to gain a foothold on the actual command and control system of the underlying smart city. The secure architecture of smart city command and control and orchestration systems therefore demands close scrutiny and strong security governance and controls to minimise the potential for its compromise. Equally troublesome could be the impact of a ransomware attack against control centres.

A successful ransomware attack against a smart city's command and control centre could cripple effective operation of the city until such time that the affected systems could be recovered from such an attack – this highlights the importance of system backups to support smart city availability and disaster recovery processes.

Smart City Secure Operating Centre (SOC) and Monitoring

In addition to general monitoring of smart city application operation and system health, municipalities should ideally monitor specifically for security incidents. Such monitoring might for example employ solutions that use machine learning to detect anomalies across sensor networks for example, or ingest intrusion detection logs concerning any exposed web-based APIs that provide access to smart city data.

Other relevant considerations here include any SOC monitoring that might impact on privacy. For example, a smart city application using facial recognition may use a city's CCTV infrastructure to check for images of individuals on a maintained watch list. In order to preserve the privacy of citizens who aren't on such watch lists, solutions might need to be considered that when presenting real-time video feeds to operators, the facial images of non-watch list individuals are blurred in real-time for example. Similarly, any applications that use or track GPS coordinates of citizen mobile devices should not be easily searchable by operators in ways that might allow for specific individual tracking across a city.

In reality, many smart city applications will likely present themselves without any input or governance by underlying city councils.

Command and control systems should therefore be built and configured with appropriate levels of authorisation (audited actions) for performing smart city data query and viewing operations that could result in some level of citizen or visitor privacy infringement.

Software Patching and Firmware Updates

Secure maintenance of smart city technology demands good patch management. The amount of software that might underpin smart city applications could be vast, spanning embedded sensors and gateways, edge computers, cloud web services and mobile apps to name but a few. This is where asset tracking of smart city components becomes important, in addition to incident response procedures that help municipalities act in the event of a critical security vulnerability being identified and disclosed concerning some underlying component of a smart city.

Much effort in this domain might involve pure risk management as opposed to technical procedures. For example, suppose an LPWAN network comprising thousands of sensors is reported as vulnerable due to some software vulnerability in the sensors; patching thousands of sensors around a city, whether remotely or in situ just may not be logistically feasible. Similarly, the low power aspect of some sensor networks might mean that pushing out a patch wirelessly to sensors would not be feasible, or would at least drain the finite battery power source of the underlying sensors thus reducing their operational lifespan.

Other key aspects to consider regarding patching include the impact on downtime. If smart city applications are performing real-time critical operations such as dynamic traffic light signalling, then taking those systems offline for even just minutes in order to apply security patches may not be feasible, or at least not feasible during busy periods. This would demand suitable consideration for system outage time, and the invocation of any necessary backup controls during such downtime, such as manual traffic operators and signs being deployed to affected areas while systems are updated.

For those systems that cannot be turned off or patched for various reasons, then upon revelation of any security vulnerabilities in affected components, the situation becomes more about managing legacy and vulnerable equipment – this may demand additional controls to be put in place around the affected devices, such as alerting and firewalling for example.

Asset Tracking and Inspection

The importance of asset management has already been discussed, however in the context of smart city operation, this includes consideration around tracking and inspection of those assets. For example, where assets might be deployed on street furniture and may potentially be easily accessible by the general public, processes and procedures should be defined and followed for routine physical inspection – this might be as simple as checking anti-tamper stickers so as to obtain assurance that devices in easily-accessible places haven't obviously been tampered with. Certainly if solutions are used within control centres to alert on anomalous behaviour of smart city assets, then physical inspection of those assets should be invoked as and when the alerting arises.

Smart City Cyber Incident Response

As part of overall governance, municipalities will need to define robust cyber incident response plans. These will need to clearly identify roles and responsibilities during incidents, such as when third party operators are responsible and/or the local authorities themselves.

Procedures need to be defined for triaging potential incidents, while for cyber-physical systems that may be affected, appropriate and timely engagement with the emergency services should be initiated.

Indeed, some smart city applications may implement some form of auto-emergency service communication, or could be configured to do so, such as smart bins automatically calling out to the fire service (via in-built SIM card and 4/5G telecoms capability) if their heat sensors detect high temperatures that might be indicative of the bin being on fire; or smart traffic light systems automatically calling out to traffic enforcement so as to swiftly deploy manual traffic direction to maintain safe traffic flows, while any underlying cyber-attack or system downtime is investigated.

Disaster Recovery (DR) procedures should also be considered as part of incident response.

Disaster Recovery (DR) procedures should also be considered as part of incident response. Municipalities should plan for DR across all applications, understanding the worst-case scenario for if each application failed and what processes would need to be followed to recover, or at least ensure continued safe and secure operation of a city. This might include establishing DR backup processes and systems.

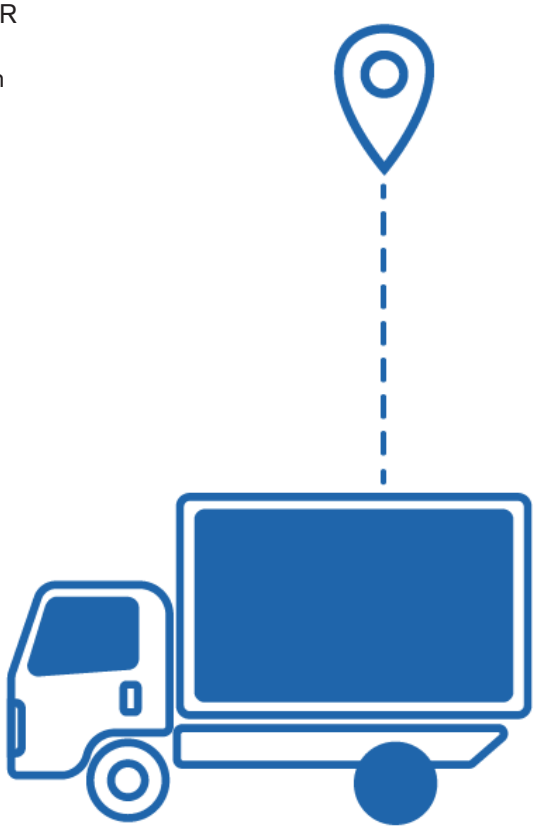

8. Security testing of smart cities

Even if testing had occurred in test or lab environments, depending on the criticality of certain smart applications, there may still be merit in production system security testing in order to ensure that security issues, vulnerabilities or misconfigurations haven't manifested themselves between design, construction and operation.

Security testing of smart city components and applications will ideally have occurred before rollout. Even if testing had occurred in test or lab environments, depending on the criticality of certain smart applications, there may still be merit in production system security testing in order to ensure that security issues, vulnerabilities or misconfigurations haven't manifested themselves between design, construction and operation.

Security testing of vendor devices and applications is certainly recommended in order to validate any vendor claims on the security of their products; in addition, the general modus operandi of some systems, once fully deployed, may be exhibiting privacy-impacting behaviours that were otherwise unaccounted for during design and deployment phases.

For example, consider a simple crowd monitoring system which uses Bluetooth listeners distributed around urban areas; the listeners capture the Bluetooth ID (address) of passing cell phones and use these unique identifiers as a simple counter regarding the number of people passing by and present within a specific area.

Security testing of such a solution might uncover misconfigurations, such as the system capturing the full device ID (a personal identifier, without anonymisation) and potentially storing all of those IDs (without consent) in a cloud-based platform, despite the system offering a configuration option to pseudonymise the identifiers, and system integrators believing they had configured the correct security setting.

With the same system, security testing might identify an implementation flaw whereby use of specialised attacker equipment could create invalid or corrupted device IDs that when processed by the Bluetooth listeners, causes them to crash due to ineffective error handling, thus affecting the availability and operation of the crowd monitoring system.

Tooling and Capability

Specialised tooling will likely be required to perform specific technical security tests against smart city applications.

For example, LPWAN technologies such as LoRaWAN and Sigfox operate on specific radio frequencies and implement specific protocols – tooling to test the security of such technologies may not be readily available, thus demanding specialised services or devoted effort to develop the tooling necessary.

At a fundamental level, owing to the various wireless-based systems that underpin smart city applications, smart city security testing tooling will need to offer the capability to intercept, replay, stress-test and fuzz [27] various wireless protocols.

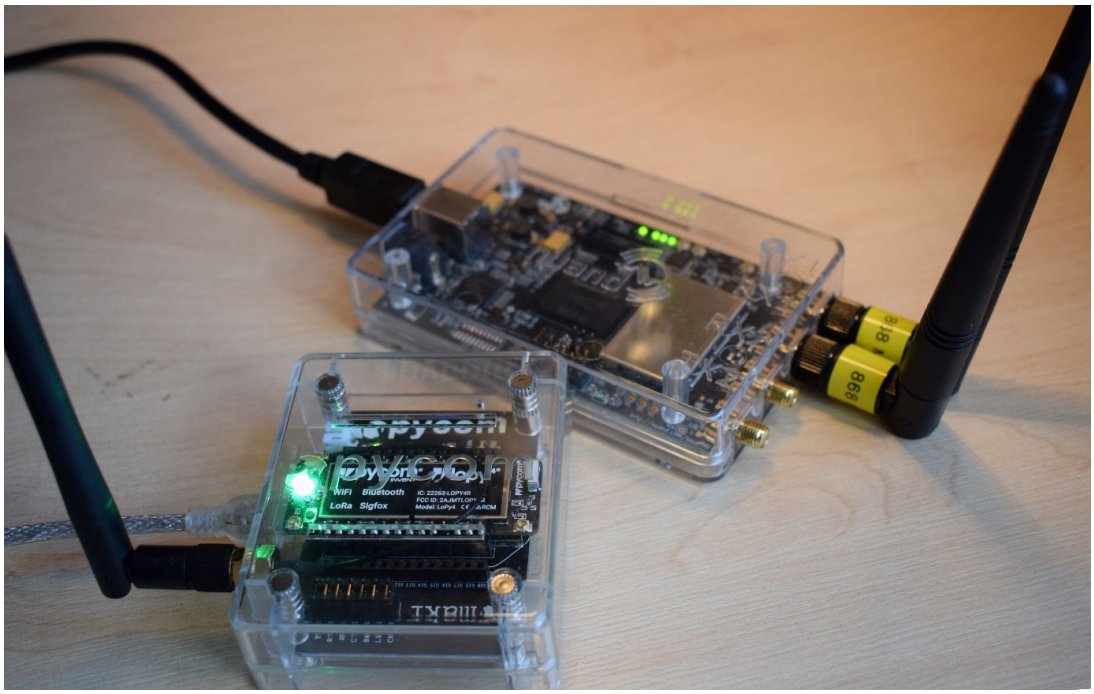

Figure 2 - NCC Group's R&D on LPWAN security tool development using COTS products and Software Defined Radio (SDR

Component Testing – Hardware and Embedded Security

Security testing of hardware components (sensors, gateways, edge computers etc.) will help understand any implementation or configuration flaws that need to be managed. In addition to uncovering security flaws, hardware and embedded security activities can also help municipalities understand:

- Component specifications – there may be scope to use different (more secure) sub-components or modules within devices; security testing of different components can help identify their respective security capabilities and limitations.
- Potential for intellectual property theft - reverse engineering activities on hardware components can help understand the level of difficulty involved in potentially copying and reproducing technologies. This may be of particular interest to those applications that rely on installation of hardware components in easily accessible public spaces. Also relevant here is understanding the level of complexity involved in retrieving sensitive data from embedded devices, such as hard-coded encryption keys or passwords. If attackers can easily retrieve such credentials from devices deployed in public spaces, then those credentials may allow for unauthorised connection to closed networks and services.

- Supply chain security – understanding where specific components are manufactured, and the supply chain of those components can help identify potential vulnerability or additional assurances needed around the secure supply of hardware components. In addition, understanding the security around product repair activities (where do devices go for repair, and who performs repair?) is important so as to identify any potential for tampering of equipment during repair processes.

- Forensics – if a smart city application or hardware component is assumed or revealed to be compromised in some way, there may be a need to perform technical forensics on the device in order to help triage the source or nature of attack and its manifestation. Security testing can help identify methods for accessing and retrieving logs and telemetry from hardware devices that can help in forensic investigations on those devices, such as understanding how or when a smart city sensor may have been tampered with, and via what mechanism.

Edge and Web Services Security Assessment

Edge computing may be used in smart city applications in order to provide low latency, high compute power capabilities, particularly for those systems that require real-time operation and that perhaps are critical in terms of their cyber-physical components. Given edge computers may be present in public places, physical and electronic security testing may be required in order to derive assurances on their secure, non-tampered operation.

Smart cities will employ a wealth of web services and likely in different combinations of open (publicly accessible) and closed (restricted to specific devices or networks). Open web services may exist so as to provide access to smart city data to citizens or innovators, typically through APIs [27]. Closed systems may employ whitelisting or firewalling in order to restrict access to web services. Security testing of production web services will help municipalities understand whether the desired data access permissions have been properly configured. Similarly, stress-testing exposed APIs and web services may be necessary in order to understand whether systems can withstand unexpected loads in network traffic, whether those loads be benign (e.g. legitimate yet sudden/unexpected large volume access) or as a result of intentional Distributed Denial of Service (DDoS) attacks.

Mobile Apps

Smart city systems may require citizens or visitors to install a mobile device app in order to make use of the underlying system, whether that be through presentation of information to citizens, or allowing them to query smart city data in backend systems and retrieve results. Such mobile apps may be developed by third parties, and it is crucial that they undergo security testing in order to ensure that they do not expose security flaws that might compromise the safety, security or privacy of citizens and visitors. Mobile app testing will typically also include testing of remote web services with which the apps communicate. Any flaws in those services could also result in exposure or compromise of citizen or visitor data. For example, suppose a smart city app captures and transmits the GPS location of citizens to a remote web service. A security flaw in that web service could provide attackers with access to all GPS locations of citizens; such information could be used to expose the privacy of citizens, or potentially to identify sources of mobile devices for physical theft (muggings).

City-Wide Security Assessment

While smart city security might be tested through component or sub-system security testing activities, as smart cities grow in size and complexity, the feasibility of component testing may diminish.

At such times, municipalities may wish to consider city-wide security assessments, leveraging outputs from threat models of the city to understand those areas that are priority and require particular focus on security testing.

Factors to consider around city-wide security assessments include:

- Routine - how often should such assessments be conducted? Once a year, or perhaps upon any major, significant addition or change to the city infrastructure?

- Physical vs. Electronic – should city-wide assessments include both physical security testing (of sensors, gateways, edge computers and other such systems attached to street furniture), and/or simply electronic components such as wireless networks and cloud/web service APIs?

- War-Driving – traditionally, war-driving has been the act of moving around an area (e.g. a city) to enumerate Wi-Fi access points and their associated security settings (e.g. are they open or secured?). For smart cities, war-driving is a relevant enumeration activity as a first phase of security testing, and includes the need to enumerate not just Wi-Fi, but other smart city wireless technologies and systems such as LPWAN, ZigBee, and Bluetooth. Items enumerated during war-driving could be correlated with asset management systems concerning the underlying smart city, and this activity might help identify unauthorised or rogue wireless components.

- War-Parking (or in situ interception and enumeration) – when technologies such as LPWAN transmit small amounts of data at infrequent intervals, the act of intercepting such wireless communications for enumeration makes more sense when performed in a static location, as opposed to moving around a la war-driving. War-parking might therefore need to be performed as part of city-wide security assessments, whereby security testers use technologies to intercept and enumerate various wireless protocols and signals, perhaps deployed in interception mode for days or weeks on end. Such activities would be best performed at physical locations with good line of sight across cities, such as at decent heights within city centres. War-parking would also help understand what (if any) types of data leakage occur across smart cities – i.e. what types of data from what applications are transmitted and potentially unencrypted or easily decipherable?

Full Spectrum Smart City Attack Simulation

Full spectrum smart city attack simulation could provide municipalities with the full end-to-end assurances needed on the security of their smart cities. This involves emulating the tactics and techniques used by real-world adversaries. Such end-to-end assessment could help identify weaknesses in system configurations, staff training and awareness, and operational response.

A combination of different teams would be engaged during a full spectrum simulation:

- Black Team – with the aim of identifying weaknesses in physical controls and staff awareness (social engineering) that facilitates physical access to smart city components and associated premises.

- Red Team – assesses cyber preventative controls, staff security awareness and challenges any Blue Team (system operators and analysts) detection and response processes.

- Purple Team – combining the Red and Blue Team activity which sees attack and response experts embedded within a smart city SOC (Blue Team) during a Red Team engagement.

- Gold Team - identifies improvements in internal and external communications, crisis management procedures and decision making – this would include understanding aspects such as how to communicate and engage with citizens, regulators and the media when faced with publicised smart city security incidents.

For any security testing on production systems, municipalities will need to properly scope and plan such testing with appreciation for potential impact on those systems. For example, testing on critical cyber-physical systems may need to be performed out of hours, during quiet periods of system operation.

6. Conclusions

Security is a process and not a solution, and as such the topic of smart city security is not to be onerous or exasperating owing to the number of different elements that need to be considered.

The topics in this paper are by no means exhaustive. The aim of the paper was to at least introduce some of the core security areas that should be considered as part of smart city design, implementation and operation.

Security is a process and not a solution, and as such the topic of smart city security is not to be onerous or exasperating owing to the number of different elements that need to be considered. Rather, just knowing the security considerations and questions to ask is useful to focus mind-sets and to support principles of security by design, where hopefully municipalities can see many instances of security being an enabler, providing city citizens and visitors with adequate assurances around privacy and safety.

Given the gradual deployment and interconnectivity of disparate smart city applications over time, there is great scope to produce secure design patterns and repeatable governance models and policies that can be reused or adapted accordingly for new systems, thereby minimising the need to start from scratch upon inception of each new system, and also to ensure consistency in security considerations for each new smart city application.

"It is not feasible to halt the smart city agenda, and much of the adoption of networked technologies and software systems by municipal authorities across the world cannot simply be removed.

However, it is not too late to recognize the extent of the new cybersecurity vulnerabilities and risks posed by these technologies and to put in place strategies and approaches to mitigate and prevent them." [1]

7. References and further reading

[1] Creating Smart Cities, Edited by Claudio Coletta, Leighton Evans, Liam Heaphy & Rob Kitchin (Routledge, 2019)
[2] https://repository.library.georgetown.edu/bitstream/handle/10822/1053223/Bellefleur_Wang_Internet%20of%20Things%20Security%20Considerations%20and%20Solutions.pdf
[3] https://discovery.ucl.ac.uk/id/eprint/1388243/
[4] https://www.bristolonecity.com/about-the-one-city-plan/
[5] https://www.iso.org/publication/PUB100423.html
[6] https://www.bbc.co.uk/news/technology-51148501
[7] https://insights.sei.cmu.edu/sei_blog/2018/12/threat-modeling-12-available-methods.html
[8] https://www.energy.gov/sites/prod/files/2017/01/f34/Cyber%20Threat%20and%20Vulnerability%20Analysis%20of%20the%20U.S.%20Electric%20Sector.pdf
[9] https://whatis.techtarget.com/definition/Code-of-Connection-CoCo
[10] https://www.smartcitiesworld.net/news/news/waterfront-toronto-delays-sidewalk-labs-decision-day-4975
[11] https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/data-protection-impact-assessments/
[12] https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/consent/what-is-valid-consent/
[13] https://ico.org.uk/your-data-matters/your-right-to-get-copies-of-your-data/
[14] https://ico.org.uk/for-organisations/guide-to-freedom-of-information/what-is-the-foi-act/
[15] https://link.springer.com/article/10.1186/s13174-018-0097-0
[16] https://www.researchgate.net/publication/326062191_An_Information_Security_Architecture_for_Smart_Cities
[17] https://ieeexplore.ieee.org/document/7802649/
[18] https://www.discocrypto.com/disco.html
[19] https://www.thethingsnetwork.org/
[20] https://www.al-enterprise.com/-/media/assets/internet/documents/smart-city-network-architecture-guide-en.pdf
[21] https://www.researchgate.net/publication/309187620_SCSimulator_An_OpenSource_Scalable_Smart_City_Simulator
[22] http://cupcarbon.com/
[23] https://www.google.co.uk/intl/en_uk/earth/
[24] https://www.nccgroup.trust/globalassets/our-research/uk/whitepapers/2018/11/third-party-assurance_final.pdf
[25] https://www.nccgroup.trust/uk/our-research/using-graph-databases-to-assess-the-security-of-thingernets-based-on-the-thingabilities-and-thingertivity-of-things/
[26] https://www.uber.com/gb/en/
[27] https://owasp.org/www-community/Fuzzing
[28] https://ieeexplore.ieee.org/document/7802649/

Image credits

[1] Front page: Anastasiya Bleskina/Shutterstock.com

[2] Page 5: Irina Kostyuk/Shutterstock.com

[3} Page 12: Irina Kostyuk/Shutterstock.com

[4] Page 15: Viktoriya/Shutterstock.com

[5] Page 22: Irina Kostyuk/Shutterstock.com

[6] Page 25: Irina Kostyuk/Shutterstock.com

[7] Page 28: Irina Kostyuk/Shutterstock.com

11. About the Author

Matt Lewis is an experienced Technical Research Director. His specialisms include general security consultancy, scenario-based penetration testing, vulnerability research and development of security testing tools. He has over seventeen year's cybersecurity experience covering all manner of technologies across most sectors and industries.

12. Acknowledgements

Many thanks to Stephen Bailey, Head of Privacy and Cyber Consulting in the Risk Management and Governance Practice at NCC Group, for his review, insights and invaluable input.

13. About NCC Group

NCC Group is a global expert in cyber security and risk mitigation, working with businesses to protect their brand, value and reputation against the ever-evolving threat landscape. Through an unrivalled suite of services, we provide organisations with peace of mind that their most important assets are protected, available and operating as they should be at all times.

With our knowledge, experience and global footprint, we are best placed to help businesses identify, assess, mitigate and respond to the risks they face. We are passionate about making the Internet safer and revolutionising the way in which organisations think about cyber security.

Headquartered in Manchester, UK, with over 35 offices across the world, NCC Group employs more than 2,000 people and is a trusted advisor to 15,000 clients worldwide.