AI AND THE FUTURE OF PHYSICAL SECURITY

Preparing for 2025 – 2030

JULY 20, 2025

Andrew Robinson,
City of Toronto
Manager, Corporate Security
Andrew.Robinson@toronto.ca

Contents

Executive Summary

Artificial intelligence is no longer just something we read about in trade journals or see in the private sector. It's already making its way into physical security work, and it's moving quickly. In 2025, we're seeing more and more municipalities across North America running AI pilots: from automating video surveillance reviews, to enabling touchless access controls, and deploying security robots in public spaces. These early tests are showing promise. In one example, the City of Gary, Indiana was able to dramatically reduce investigation times by using AI to scan through hours of CCTV footage and surface only the key moments worth reviewing.

Looking ahead to 2030, most experts expect AI to be deeply embedded in city security operations, helping us detect threats before they escalate, respond to incidents faster, and automate some of the more repetitive work we currently rely on human guards to do. The technology is advancing quickly, and the market is following. Just consider that security robotics alone are projected to top $39 billion globally by 2030.

Of course, this transformation isn't just about adopting new tools. It raises some real challenges: ethical, operational, and cultural. Issues like privacy, bias, and workforce disruption all need serious attention. But the opportunities are equally significant. AI can analyze surveillance footage 24/7 without fatigue, make sense of large volumes of data, and even help unify systems that used to operate in silos.

If we want to keep pace with evolving risks and stay ahead of them, we can't afford to wait. The next five to seven years will likely bring major shifts in how physical security and data intersect. Cities that invest now in pilot programs, infrastructure, and staff training will be better positioned to protect public spaces and make the most of constrained resources. This report aims to help municipal leaders, including security managers, directors, and finance decision-makers, understand where AI is heading, what's working today, and what steps we can take to prepare.

The bottom line is this: AI is already reshaping the security landscape. The sooner we lean in and shape that future ourselves, the better prepared we'll be to keep our communities safe, smartly, ethically, and responsibly.

Introduction

In a municipal setting, physical security means protecting the people, infrastructure, and public spaces that keep a city running. That includes everything from doors and access points at civic buildings, to perimeter cameras, water treatment plant patrols, and crowd safety at large events. Traditionally, we've relied on human presence and passive systems, such as CCTV monitored by staff, keycards and locks, and scheduled patrols. For many years, that approach was effective.

But today, AI is beginning to reshape physical security just as it is transforming many other industries. Its real strength lies in the ability to process large volumes of data quickly and identify

patterns that are too complex or too time-sensitive for people to catch on their own. That's especially important in our line of work. A single high-definition camera can generate three terabytes of video every month, and cities like Toronto have thousands of them. We've reached a point where it is no longer practical for operators to watch every feed or respond to every event manually.

With AI in the picture, our security tools are becoming smarter. Cameras can detect faces, recognize unusual movement, and trigger alerts automatically. Access control systems can learn how credentials are normally used and flag irregular behaviour. Sensors are starting to interpret activity on their own, using data rather than simple input-output logic.

This shift moves us away from reactive security, where we respond after something has already happened and toward a more proactive and preventative approach. Imagine a system that flags a potential issue as it develops, or one that limits door access when it detects a pattern that doesn't fit. These kinds of capabilities are no longer far-off ideas. They're already being tested and, in some cases, rolled out in North American cities.

AI is also helping connect systems that used to operate in silos. Video, access control, alarms, and sensor networks are being integrated into cohesive platforms that provide a real-time view of risks. This kind of unified situational awareness is a major step forward for municipal security.

In the sections ahead, we'll look at how AI is currently being used in municipal physical security, how we expect that use to evolve by 2030, and what steps city leaders can take today to prepare.

AI In Physical Security Today (2025)

As of 2025, we're seeing real momentum in how AI is being applied across different areas of physical security. Municipal governments, public agencies, and private companies are running pilots and, in some cases, rolling out full systems. The key areas where AI is showing the most promise include access control, video analytics, perimeter protection, and predictive maintenance for security equipment. Below is a breakdown of what that looks like in practice.

Access Control: Smarter Credentials and Biometric Systems

AI is making physical access to city buildings and facilities more intelligent and secure. Traditional keycards and PINs are being backed up, and in some places replaced, by biometric systems and behavior-based verification. We're seeing facial recognition checkpoints and touchless fingerprint scanners used in government buildings and some airports to verify identity. More advanced systems can now look at when, where, and how a credential is used, and flag anything unusual, such as a stolen badge or tailgating.

One example is Guardly, a Canadian company offering a mobile credential platform that uses real-time location tracking. City staff can unlock doors with their phones, and AI monitors usage

patterns to detect irregularities. These systems improve both security and user convenience. Another area to watch is behavioral biometrics, where AI learns how someone moves or types to verify their identity continuously. This is still emerging, but it shows promise in reducing manual spot checks and making unauthorized access much harder to pull off.

Video Analytics: Turning Cameras Into Real-Time Detection Tools

Out of all the use cases, AI-powered video analytics might be the most mature right now. Cities are using these tools to help law enforcement and security operators sift through massive amounts of footage faster and smarter. The City of Gary, Indiana, for instance, cut down investigation times by using AI to identify faces, vehicles, and movement patterns across multiple cameras. This helped their police focus their efforts instead of watching hours of video manually.

Schools are also investing. The Clovis School District in New Mexico signed a four-year contract with ZeroEyes, a company that uses AI to detect firearms on camera in real time. When a potential weapon is spotted, the system sends an image to a monitoring center staffed by trained analysts, who verify the alert and notify police within seconds. The goal is to act before a shot is ever fired. Transit systems are exploring similar tools. In 2025, the New York MTA began testing AI cameras to identify unsafe behavior like people entering tracks or crowds building up on platforms.

Other cities are using video analytics to manage everyday issues, counting people for crowd control, flagging abandoned packages, or alerting staff to fights or loitering. In all these examples, AI is helping turn camera networks from passive systems into active tools that improve situational awareness and help teams make faster, more informed decisions.

Perimeter Protection: AI Watching the Edges

Securing the perimeter of critical sites is another area where AI is making an impact. Cities are using smart sensors and cameras that can tell the difference between a real threat and a false alarm. For example, a smart fence system might recognize when a person is climbing over a gate compared to when a dog runs past or a tree branch blows in the wind.

Drones are an emerging concern, helpful in some cases but also a risk for unauthorized surveillance or disruptions. AI-powered systems like Dedrone are being used to detect, track, and even identify the location of rogue drones and their operators. Utilities like Con Edison in New York are already using this to protect energy infrastructure, and Barcelona's police have deployed similar tech to manage urban airspace.

On the ground, cities and campuses are testing autonomous security robots like those from Knightscope. These mobile units patrol outdoor areas, use 360-degree video and thermal sensors, and send real-time alerts to command centers when something unusual is detected. Some municipalities are starting to see these robots as a scalable, 24/7 addition to their frontline security, especially in places where adding human patrols is difficult or cost-prohibitive.

Predictive Maintenance: Keeping Systems Healthy Behind the Scenes

While not as visible as access or video analytics, AI is also helping to maintain the systems we already rely on. Cities have thousands of cameras, sensors, alarms, and card readers. If one of those devices fails, for example, if a door doesn't lock or a camera goes offline, it creates a blind spot or vulnerability. AI can monitor the health of these systems in real time and flag issues before they cause a problem.

One platform, Sentry AI, offers a camera "health check" that uses machine learning to detect when a camera is off-angle, obstructed, out of focus, or tampered with. Other platforms like Nanodems are providing full dashboards that track the live status of every security device, helping technicians stay ahead of failures. This kind of proactive maintenance helps reduce downtime, improves reliability, and lowers the cost of manual inspections, which can be hard to scale across hundreds of sites.

Where Things Stand in 2025

AI is no longer just a buzzword in physical security. Cities that have started experimenting with it are already seeing better response times, new capabilities like drone detection, and less time spent reviewing footage or performing routine checks. A City of Gary official described their upgraded system as a "force multiplier," giving them virtual eyes where officers can't always be.

At the same time, these tools come with new challenges, especially around privacy and accuracy. Some school boards and municipal councils are moving slowly, making sure any use of AI is backed by policy, community support, and safeguards to reduce bias. That's a healthy approach.

As we look toward the next phase, the pace of change is only going to accelerate. The next section explores what that could look like as we move toward 2030, and how cities can start preparing now to lead the way.

The Path to 2030: Projected Trends and Disruptions

AI in physical security is moving quickly, and the next five years will bring a wave of change. We are not just talking about upgrades in technology. We are looking at fundamental shifts in how security is delivered, how teams operate, and how cities prepare for what's coming. Below are five trends that, in my view, municipal leaders should be watching closely as we plan for 2030.

1. Rapid Acceleration of AI Capabilities

AI tools are improving fast. Features we think of as pilot projects today could become standard practice in just a few years. The development cycle is speeding up exponentially, especially in areas like computer vision and pattern recognition. As a result, many of the technologies we are testing now will be significantly more powerful, more affordable, and more widely available by the end of the decade.

Municipal cameras and sensors will likely come with AI capabilities already built in or easily connected through cloud platforms. Tasks like scanning footage for threats, verifying identities, or detecting suspicious behavior may become fully automated in real time. Cities that lock into rigid technology now may find themselves stuck in two or three years when something better becomes available.

The takeaway here is that AI planning needs to be agile. Contracts should allow for software upgrades and model retraining. Pilot programs should include room for iteration and learning. A good example would be using an anomaly detection model in a city park that keeps learning over time and eventually starts picking up subtle cues that suggest something is about to go wrong. The cities that stay flexible and build with scale in mind will be best positioned to adapt as capabilities evolve.

2. Physical and Cyber Security Are Merging

The old line between physical security and cybersecurity is disappearing. By 2030, we can expect to manage security more holistically, with integrated platforms that bring both worlds together. These new systems will combine feeds from access control, video surveillance, door alarms, and network security tools, making it easier to connect the dots.

This kind of convergence is already starting. Modern PSIM platforms now act like central nervous systems, pulling in information from physical devices and IT systems to detect coordinated threats. For example, a cyberattack that disables a camera at the same time as a door is forced open would be flagged immediately.

Another development to watch is the use of digital twins. These are 3D virtual replicas of real buildings or areas. In 2025, vendors began offering digital twin systems that let you “walk through” a building remotely, viewing live camera feeds and alarms in the virtual space. By 2030, cities may use digital twins for major sites like city halls or transit hubs, especially for emergency response or training.

We should be thinking now about how to connect our systems, share data between departments, and ensure physical devices are treated as part of the broader city network. This kind of integration gives us better visibility and helps us manage risks more proactively.

3. The Security Workforce Will Shift

AI will change how our teams work. Some traditional duties will likely be reduced, especially jobs that focus on monitoring video or reviewing access logs. But this doesn’t mean security jobs will disappear. Instead, they will shift into more technical, more analytical, and more decision-oriented roles.

We may need fewer people watching camera feeds, but we will need more staff who can interpret alerts, tune AI systems, and investigate patterns the system uncovers. Roles like AI System Supervisor or Security Data Analyst could become common in our field by the end of

the decade. Front-line guards may begin using smart devices that feed them real-time information from AI tools, helping them make quicker, better-informed decisions on the ground.

Cities should start preparing for this now. That means updating job descriptions, investing in upskilling, and making sure our current staff have clear paths forward. We also need to communicate openly with unions and employees about how roles are evolving, not disappearing. If we do this well, we can actually improve job satisfaction by removing the most repetitive tasks and allowing staff to focus on what they do best, making judgment calls, engaging with people, and responding when it counts.

4. Policy, Ethics, and Public Trust Will Matter More Than Ever

As AI becomes more powerful, it will also come under more scrutiny. The public wants to know how these tools are being used, and why. Privacy, bias, and transparency will all be front and center. By 2030, we can expect stricter laws, clearer regulations, and higher expectations from residents.

Some cities have already started putting limits on technologies like facial recognition. Others are requiring clear policies about how long AI-generated images and data from surveillance footage is kept and who has access to it. These questions are not going away. In fact, they will only get louder.

Municipal leaders will need to be proactive. That means setting clear rules for how AI is used, conducting regular audits for fairness and accuracy, and making that information available to the public. It may also mean creating advisory boards or working with third parties to review how systems are performing.

Trust will be a deciding factor. If residents believe that AI is being used responsibly and can see the benefits in faster response times and safer spaces, they are more likely to support it. But if the system misidentifies someone or if policies are unclear, the backlash could slow or stop progress. Managing that balance is as important as managing the technology itself.

5. The Human Element Still Matters Most

With all the focus on automation and analytics, it is easy to forget this: municipal security is still about people protecting people. AI can process information, flag unusual behavior, and recommend action, but it cannot provide leadership, empathy, or judgment in the moment.

By 2030, we will see a shift in how human and AI roles complement each other. AI will handle the heavy lifting, scanning feeds, detecting anomalies, and filtering routine events. Humans will step in when things are unclear or when a decision needs to be made with context and care. We will also see new responsibilities emerge, like overseeing the AI tools themselves, managing system health, or responding when the system raises a concern.

In some cases, guard roles may shift into something closer to a “technology conductor,” managing alerts from multiple systems and stepping in when human intervention is required. IT

teams will likely be more involved in physical security than ever before, especially in supporting cloud-based tools, software updates, and integration with other city systems.

Done right, this shift can improve performance, reduce burnout, and help our staff focus on what matters most. We are not replacing people. We are freeing them up to do higher-value work that technology alone cannot handle.

In Summary

The road to 2030 is not just about adopting new tools. It is about rethinking how we deliver safety in a digital world. Cities that start planning now will be in a strong position to scale AI responsibly. That means budgeting with agility, removing silos between departments, helping the workforce evolve, and engaging the public with honesty and transparency.

The next sections of this report will explore who the major industry players are and highlight real-world use cases where cities are already seeing results.

Key Industry Players & Platforms

As AI becomes more common in physical security, several companies are stepping up with real solutions that are already making a difference in cities. Municipal leaders planning security upgrades or pilots should keep an eye on vendors with a proven track record in public-sector environments. Here are some of the main players shaping the landscape today.

1. BriefCam

BriefCam has been a frontrunner in AI video analytics. Their platform makes it possible to scan hours of security footage in just minutes using technology called “Video Synopsis.” It also supports live analytics like people counting, facial recognition, and license plate detection. BriefCam is already used in many police departments and city-wide safety programs to speed up investigations. Cities that already have a video system in place can easily integrate BriefCam’s software on top. Since Canon acquired BriefCam, it works well with a range of VMS platforms, which is useful for municipalities that want to add AI capabilities without replacing their infrastructure. The City of Toronto has purchased Briefcam but has not fully implemented it.

2. Evolv Technology

Evolv offers AI-powered screening systems designed to detect weapons at public entrances. Instead of traditional metal detectors, Evolv scanners use multiple sensors and AI to distinguish between everyday items and actual threats. People walk through at regular speed without stopping or emptying their pockets. The system flags only items of concern, which reduces false alarms and long lines. Baltimore’s school system recently committed over $5 million to roll Evolv

out in 28 schools. Cities are using Evolv in courthouses, city halls, and transit hubs to enhance safety while keeping operations moving smoothly. New York City even tested Evolv in its subway system.

3. ZeroEyes

ZeroEyes is focused on a single problem: early detection of firearms. Their AI scans existing camera feeds for visible guns and sends an alert to a monitoring center, where a trained person confirms the threat before contacting local police. This human-in-the-loop model reduces false alarms and provides real-time alerts. It is gaining traction in schools and city buildings that are concerned about active shooter incidents. ZeroEyes integrates easily with existing VMS systems and offers a targeted upgrade path for cities wanting to strengthen response to gun violence without a full overhaul of their security setup.

4. Ava Security (Avigilon Alta)

Ava Security, now part of Motorola Solutions, builds AI-enabled cameras and a cloud-based platform under the Avigilon Alta name. Their cameras detect motion, count people, and even recognize sounds like glass breaking or gunshots. Their software, Ava Aware, lets security teams search video feeds using keywords and sends alerts for suspicious behavior, like someone loitering in a stairwell. What makes Ava stand out is its full integration with Motorola's Openpath access control system, creating an end-to-end solution for managing cameras and entry systems from one interface. For cities looking to move away from on-premises servers, Ava is a strong option that's already in multiple areas.

5. Rhombus Systems

Rhombus builds cloud-based security cameras with AI features like anomaly detection, unauthorized access alerts, and environmental sensors. These systems are easy to deploy and manage, making them ideal for municipalities without large IT teams. Cities can use Rhombus to monitor multiple facilities from a single dashboard, and the cameras learn over time. For example, if motion is detected in an unused building at 2 a.m., an alert can be sent automatically. While cloud systems require careful attention to data privacy rules, many cities appreciate how quickly and simply these tools can be rolled out.

6. Genetec and Milestone

Genetec (based in Canada) and Milestone (based in Denmark) are two of the most widely used video management platforms in the public sector. They are not AI companies, but both have added AI capabilities through partnerships and modules. Genetec's platform now includes object detection, license plate recognition, and integration with third-party AI tools. Milestone is known for its open architecture, which allows cities to plug in analytics platforms like BriefCam or Axis AI tools. Most cities using Genetec or Milestone are layering AI onto their existing

systems instead of starting from scratch. These platforms are trusted for their reliability and scale, and they play a central role in unified security operations.

Specialized and Emerging Vendors

There is also a growing group of niche vendors bringing AI into specific security areas:

- **Knightscope** builds autonomous patrol robots used in public areas, private campuses, and transit spaces. These robots use AI for navigation and threat detection. Knightscope is seeing increased adoption and has recently expanded its operations.
- **Dedrone** leads the market in drone detection and response systems. Their AI platform tracks drones across urban areas and is already in use in city drone defense networks. Dedrone has also partnered with Axon (maker of Tasers) to extend its capabilities.
- **Cognyte** and other analytics firms are offering Safe City platforms that combine AI video tools with social media monitoring and dispatch integration. These systems are more common internationally but are starting to show up in U.S. real-time crime centers.
- **Startups to watch** include:
    - **Everbridge** for AI-enhanced event management
    - **Omnilert** for instant lockdown and gun detection
    - **Spot AI** for budget-friendly AI video tools
    - **Ai-RGUS**, which monitors camera health and alerts teams if feeds go down

These smaller firms often fill gaps left by the larger vendors and may offer creative solutions for specific municipal needs.

What Cities Should Consider

When evaluating vendors, city security leaders should consider public-sector experience, cloud readiness, ease of integration, and compliance with Canadian privacy laws. It is also worth reviewing pilot programs and reference sites. Just because a platform works well in a corporate setting doesn't mean it will meet municipal standards. Some vendors offer flexible pricing or subscription models, which can help spread out costs. Cities that want to modernize without a full rebuild may benefit from tools that plug into existing infrastructure.

As of 2025, vendors like Genetec, Motorola, Axis, and Honeywell remain established players. At the same time, newer entrants like Evolv, ZeroEyes, and Dedrone are gaining ground fast with focused, high-impact tools. Partnerships between companies are also worth watching, as vendors combine strengths to deliver more complete packages.

In summary, there is no one-size-fits-all platform. Cities will need to match tools to their priorities, budgets, and infrastructure. Staying plugged into industry groups, attending trade

shows, and sharing lessons with other municipalities can go a long way in keeping security programs modern, informed, and ready for what is coming next.

Use Cases: Municipal AI in Action

To bring these ideas down to street level, here are a few short case studies showing how AI is already being used by cities and public agencies to improve physical security. Each example outlines the problem, the solution put in place, the results so far, and the vendor involved.

1. Weapons Detection in Schools (Baltimore)

Problem:
Baltimore City Public Schools had a serious issue with weapons showing up in schools. Traditional metal detectors created long lines, and manual bag checks were slow and sometimes missed key items. Staff wanted to improve safety without creating daily bottlenecks or putting too much pressure on security guards.

Solution:
In 2024, the district installed Evolv's AI-powered weapons detection systems in more than 20 high schools. Students now walk through at a normal pace. There's no need to stop or empty pockets. The scanners use sensors and AI to look for the shape of weapons, not just metal. Security guards monitor a tablet that shows an outline of anyone flagged, with a clear visual of where the item is on the body or in a backpack.

Benefits:
Throughput improved right away. The system can screen over 3,000 people an hour, which is a major time-saver. More importantly, it catches real threats while reducing false alarms. Early feedback from students showed they felt safer, and less bothered by the process. Baltimore also took the time to explain the system to the public, which helped build trust. This is a good example of how to modernize screening without slowing everything down.
Vendor: Evolv Technology

2. AI-Assisted Investigations (City of Gary, Indiana)

Problem:
The City of Gary was dealing with high crime and limited police resources. Officers often spent hours scrubbing through video footage from city and private cameras to identify suspects or vehicles. The process was slow, and critical leads were often delayed.

Solution:
In 2022, Gary Police launched a Real-Time Crime Center (RTCC) that uses AI to analyze video across the city. The system includes facial recognition, license plate readers, and object

tracking. Instead of manually reviewing footage, analysts enter details like a suspect's appearance or vehicle type, and the AI does the scanning. Alerts for hot-listed vehicles or persons are pushed directly to patrol officers. Gary also partnered with Fūsus, which connects private cameras into the city's network and applies analytics across both.

Benefits:

The AI tools have already reduced response time and improved arrests. In one case, a robbery suspect was identified and arrested the same day after AI matched his face with a mugshot. In another, stolen vehicles were located and recovered using plate readers. Officers say they now spend less time reviewing footage and more time responding strategically. The City of Gary's approach shows that even a mid-sized city can stretch its staffing resources by letting AI handle the heavy lifting.

Partner: Fūsus

3. Gun Detection and Emergency Response (Clovis, New Mexico)

Problem:

Clovis, New Mexico, schools had cameras, but no way to detect weapons in real time. Monitoring was manual, and false alarms from fake threats disrupted the school day. They needed something proactive that would catch a gun before shots were fired.

Solution:

In 2023, the district installed ZeroEyes, an AI platform that monitors existing cameras for firearms. When the system sees what it thinks is a gun, it sends an image to ZeroEyes' operations team, where trained military veterans verify the sighting. If it's real, they alert school security and local police within 10 seconds, often before anyone calls 911.

Benefits:

There haven't been any active shooter incidents since the system went live, but even in drills and accidental cases involving toy or replica weapons, the system worked. It provided fast, clear alerts to the right people. The big win here is peace of mind. Staff know every camera is now a smart sensor. Response time is faster, and lockdowns are less frequent because the human review layer reduces false alarms. Clovis now has a 24/7 digital guard on duty, built into their existing system.

Vendor: ZeroEyes

4. Real-Time Watchlist (Lakeland, Florida)

Problem:

Lakeland had a recurring problem downtown with a few people who had been legally trespassed from public spaces like libraries and markets. With a small police force, there wasn't always someone around to catch them if they came back.

Solution:
In 2024, the city installed 14 high-resolution cameras in the downtown core and used Verkada's facial recognition tools to build a watchlist. Only three individuals were included, based on past incidents. When one of them enters the area, an alert goes out to police and downtown safety staff in real time, allowing officers to respond quickly.

Benefits:
Since deployment, the problem people have mostly stayed away. Businesses and the public say the area feels safer, and police can enforce trespass orders without being everywhere at once. Transparency was key: the city explained the limited use and who was on the list. That helped ease privacy concerns, though advocacy groups still raised flags. Lakeland's case shows how facial recognition can work as a narrow, well-communicated tool for managing repeat offenders.
Technology: Verkada Face Recognition

Takeaways for Other Cities

These examples offer several lessons:

- Baltimore showed the value of strong communication and training when rolling out new tech
- Gary's success hinged on public-private integration and real-time analytics
- Clovis benefited from a clear use case and human-in-the-loop verification
- Lakeland's case stressed the importance of transparency and scope control

If your city is considering AI, start by defining a clear problem. Then look at tools that fit that need, not the other way around. Consider pilot programs before going all-in. These stories prove that when done right, AI can strengthen city safety, improve efficiency, and do it in a way that works within tight municipal budgets.

Phased Implementation Roadmap (2025–2030)

For municipal security leaders, rolling out AI is not a single project. It's a multi-year journey that works best when implemented in phases. Each phase allows time for testing, learning, and scaling. Below is a practical roadmap broken into three stages: early groundwork in 2025 to 2026, broader deployment in 2027 to 2028, and full integration in 2029 to 2030.

Phase 1 (2025 to 2026): Build Awareness and Run Pilot Projects

Goal: Learn what AI can offer and start small

- ***Educate the Team:*** Start by briefing your senior leadership, IT staff, legal advisors, and security personnel on how AI is being used in other cities. Consider running a few internal sessions that focus on realistic use cases and limitations. Setting the right expectations early will help reduce resistance later.
- ***Identify Opportunities:*** Take a good look at your current pain points. Are your operators overwhelmed by camera feeds? Are your access systems vulnerable to tailgating? Use input from security staff, facility managers, and even frontline workers to pinpoint areas where AI could help.
- ***Start Pilots:*** Choose one to three use cases to test. Maybe try AI-powered video analytics on a few cameras or a weapons detection unit at a public-facing facility like City Hall. Work with your procurement team to secure short-term agreements for these trials. Be sure to set clear success metrics up front, such as improved response time or fewer missed incidents.
- ***Check Infrastructure:*** Before deploying any tech, make sure your cameras, servers, and networks are ready. Some upgrades might be needed, such as switching to a cloud-based VMS or boosting bandwidth. You'll also want to improve data hygiene by updating floorplans, labelling camera feeds, and cleaning up your device inventory.
- ***Lay Policy Foundations:*** Even during a pilot, get ahead of privacy and ethical concerns. Work with your legal team to draft initial guidance on how AI-generated data will be stored, who can access it, and under what conditions it can be used. For example, if testing facial recognition, limit it to a specific use case and ensure council or community oversight.
- ***Engage Vendors and Train Staff:*** Take time to meet with multiple vendors, see demos, and gather references. Begin basic training with your internal staff, even if it's just an intro to AI terms and how alerts work. Start building a small internal group that can become your AI champions going forward.

Outcome by End of Phase 1:

You should have at least one working AI use case producing real results. Your leadership team will better understand what's possible, your staff will be more comfortable with the tech, and you'll be ready to plan for a more structured rollout.

Phase 2 (2027 to 2028): Expand Solutions and Formalize Policy

Goal: Move from pilot projects to full-scale implementation

- ***Scale Up What Works***: If your pilot showed value, expand it. For example, if a single building had AI-supported access control, add it to all high-risk sites. Use the lessons learned to craft better RFPs and long-term contracts. Focus on platforms that combine

multiple tools into one interface where possible, like integrating video analytics, access control, and incident management.

- ***Adjust Staff Roles:*** AI changes how teams work. Develop a staffing model that accounts for the new workflow. You may need to shift some positions or introduce new ones, such as “AI monitoring lead” or “security systems analyst.” Work with HR and unions early to map out training or reassignments. Make it clear that existing staff will be supported and trained, not left behind.
- ***Formalize Governance:*** Set up clear processes for how AI tools are used, maintained, and evaluated. Create operating procedures that outline what happens when an AI alert is triggered, how false positives are handled, and what oversight is in place. This is also when you update privacy policies to cover AI data and system logs. Share these publicly to build transparency.
- ***Check for Bias and Accuracy:*** Now that the systems are running more broadly, bring in an independent check. You can work with third-party experts or university researchers to audit performance. Make sure the tools are working fairly across demographics and that alerts are consistent and reliable. These audits should happen regularly and results should be documented.
- ***Collaborate Across Departments:*** Start thinking beyond security. Connect your AI systems with other departments. For example, integrate badge access logs with cybersecurity monitoring or allow your RTCC to share alerts with emergency management. These partnerships will multiply the value of your investment.
- ***Refine and Optimize:*** Not every vendor or tool will scale perfectly. Be prepared to adjust. This is the phase where you lock in your core platforms and eliminate anything that’s not delivering value. Keep collecting data that shows performance and return on investment.

Outcome by End of Phase 2:

You’ll have mature AI systems operating in key parts of your security program. Your team will know how to use them, your governance structure will be formalized, and your policies will be in place. AI will no longer be “new” but part of how your city manages security day to day.

Phase 3 (2029–2030): Full Integration and Optimization

Goal: Reach a point where AI is fully embedded in your municipal security operations. The focus shifts to performance, expansion, and long-term sustainability.

- ***City-Wide Integration Across Operations:*** By now, AI should be part of every major area of your security ecosystem. This is the time to bring together everything that’s been built over the years. Connect your security AI systems with other smart city tools. This could include tying in environmental sensors for air quality or hazardous materials with your real-time security platforms. You may also have digital twins running which are

live digital models of critical infrastructure or downtown zones where all your security data flows for real-time drills and emergency response. By 2030, your city should be operating a centralized security command center that handles everything from a door alarm to a multi-site incident. If different tools have been added over time, now is the time to make sure they talk to each other. Even if you can't merge into one platform, get them linked using APIs or middleware.

- ***Ongoing AI Training and System Learning***: AI tools are not fire-and-forget. They require regular updates and improvements. Set up a process to retrain models using new data and real-world feedback. If your video analytics system is producing false positives in certain lighting conditions or environments, feed that back in. Vendors may release updates, but your local context matters and training with your city's data can make a difference. If resources allow, consider hiring a part-time data scientist or partnering with a local university to help you fine-tune AI models for specific use cases. This helps improve accuracy and lets you build AI capabilities that reflect the unique needs of your city.

- ***Complete the Workforce Transition***: By this point, your security team should be fully transitioned into a tech-enabled model. Most retraining will be complete, and any job changes should have happened through attrition or planned redeployment. Now the focus is on keeping your staff skilled and confident. Make sure all front-line staff are comfortable interpreting AI alerts and understand when to rely on them and when not to. Offer regular upskilling opportunities, especially in areas like technology troubleshooting and AI ethics. This is also the time to clearly define roles. Who is responsible for maintaining AI systems? Who is accountable when an alert turns out to be incorrect? Make those responsibilities part of formal job descriptions and internal governance documents.

- ***Fine-Tune and Measure What Matters***: Now that you're fully operational, look at what's working and what needs adjustment. Are you getting too many false alerts? Work with your vendor to adjust the thresholds. Are some features going unused? Maybe your team needs more training or reminders about their value. You should also be collecting and reporting ROI. Show how AI contributed to fewer incidents, quicker response times, or cost savings on overtime. This kind of data is important for justifying the program during future budget discussions. Compare yourself to other cities where possible. See where you are ahead and where you might need to invest further. The goal is not just to be functional, but to be high-performing.

- ***Plan for the Future and Stay Flexible***: Technology will keep changing. Start exploring what's next. This might include edge AI, where sensors make decisions locally without sending data to the cloud or smarter integrations between physical and cyber security. Maintain some funding or partnerships to pilot new tech in small ways. Make

sure your existing systems and contracts give you room to grow or switch platforms when needed. Avoid getting locked into one vendor or architecture. Scalability should be part of your strategy now, especially as your city grows or adds new facilities.

- ***Maintain Public Trust and Transparency***: As AI becomes part of everyday operations, public trust matters even more. Publish regular reports explaining what AI is used for, what outcomes it has delivered, and how privacy is protected. Open the door to feedback, use surveys, host town halls, or create online spaces for residents to share input. Be ready for new laws and regulations that may come in by 2030, including those requiring third-party audits or more formal oversight. Consider becoming a knowledge source for other municipalities. Share your experiences through presentations, white papers, or conferences. Help set the tone for responsible AI use in public safety across the country.

Outcome of Phase 3:

By the end of 2030, your city should be running a fully integrated, AI-powered security program. Staff will be reskilled and confident. Systems will be working together across departments. Incidents will be responded to faster and more strategically. And employees and the public will feel that safety is being enhanced without sacrificing accountability or transparency. This phase isn't a finish line. It's the start of a continuous cycle of maintenance, learning, and refinement, where AI becomes a core part of how you keep the city safe.

This roadmap is designed to keep adoption paced with readiness. Jumping straight to city-wide AI without earlier groundwork can lead to failure or pushback. But moving too slowly risks missing out on the benefits. The goal is to build momentum, stay intentional, and make sure your rollout is aligned with your city's needs, resources, and values.

AI Readiness Checklist

Integrating AI into municipal security takes more than just buying software. It requires careful planning across both technology and the people who use it. This checklist is designed to help municipal security leaders get ahead of potential issues before rollout.

Technology and Data Readiness

Infrastructure Checkup

Do we have the right hardware, software, and network strength to support AI tools? Think about camera resolution, server processing power, or cloud connectivity. If we're using cloud services, is our bandwidth strong enough and is cybersecurity in place?

System Integration

Can our current systems like CCTV, access control, and alarms, connect to AI platforms? Start

by inventorying your systems and checking whether they have APIs or upgrade options. Aim to bring all data streams into one spot for the AI to analyze. A data lake or central dashboard might help.

Data Quality and Housekeeping

Are our data inputs accurate and useful? That means things like camera views covering the right areas, sensors working as intended, and logs properly time-stamped. Label your camera feeds properly, update site plans, and clean up outdated or mislabeled data. Develop a data retention plan that outlines what video or sensor logs to keep and for how long, with privacy in mind.

Security and Privacy Built In

AI systems must meet security standards. Require encryption at rest and in transit, set user access controls, and log who accesses what. If you're using sensitive tools like facial recognition, build in safeguards like face blurring for individuals not relevant to the alert. Ask vendors how they handle privacy and check that they're aligned with any city policies.

Vendor Selection

Don't just buy off a brochure. Check references with other cities or organizations. Look for third-party evaluations on accuracy and bias. When issuing an RFP, ask for explanations on how their AI makes decisions. Also ask about long-term support. Will they still be around and updating this system in five years?

Test Before Launch

Set up a sandbox. Before going live, test the AI system on recorded data to see how it performs. Build a rollout plan with fallback procedures. If the AI fails, can your team switch back to manual operations quickly and cleanly?

Room to Grow

Make sure your system can scale to support more sensors or higher workloads. Cloud systems are easier to scale but be mindful of contract flexibility. If you are on-premises, invest in additional capacity now if you expect to grow. Keep tabs on future tech like edge devices or 5G. Your setup should be able to handle these when they arrive.

Workforce and Organizational Readiness

Basic AI Literacy

Have we trained the team on what AI is and what it isn't? Set up workshops or short briefings. Staff should understand that AI helps them but doesn't replace them. Make sure they know the limits of AI and that it can make mistakes or "hallucinate." Also, train them on any new dashboards or tools the system introduces.

Managing the Change

Develop a plan for how jobs and routines will evolve. Talk to the team early and explain why AI is being introduced. Involve them in pilots so their feedback shapes how the tech is used.

Identify where resistance might come from. Maybe staff worry about job loss, and address it clearly. Share how you'll support retraining and offer growth paths.

Clarify Roles and SOPs

Review your procedures. If an AI tool triggers an alert, who responds? How? Make escalation paths clear. If one person sees the alert and assumes someone else will act on it, that's a problem. Consider whether to create a dedicated AI console role per shift or add the responsibility to your dispatchers. Update job descriptions if needed.

Build New Skills

Identify what skills your team is missing. Maybe they need more experience reading data trends or working with new interfaces. Consider certifications in areas like "AI in Physical Security." If you're hiring, look for candidates who understand both security and tech. The long-term goal is a balanced team that knows security operations and can work with advanced tools.

Work Closely with IT and Data Teams

AI in security touches multiple departments. Your physical security team needs to partner with IT, CISO and data experts. Create a joint working group that meets regularly. Make sure data storage, software updates, and incident response plans cover both physical and digital aspects. For example, a badge misuse alert might need both the physical team and the cybersecurity group.

Ethics and Privacy Training

Anyone who manages or uses AI tools should be trained on privacy and ethics. Everyone should understand how to avoid bias, prevent misuse, and respect people's rights. Reinforce policies, for example, facial recognition matches are investigative leads, not final evidence. Promote a culture that uses AI responsibly and transparently.

Community Engagement

Get ready to talk to the public. Plan to host at least one public meeting or information session before deploying AI systems across the city. Give your staff key talking points. Consider creating a simple website that explains what tech is in use, why, and how the public can ask questions. A little transparency goes a long way in building trust.

Track and Review AI Performance

Set up regular monthly or quarterly review to look at how the AI is performing. Go over any incidents triggered by the AI. Were they accurate? Did the response make sense? Use these sessions to refine both the tool and how the team uses it. Assign someone to keep detailed logs of alerts, responses, and any issues. These logs may be needed later for audits or public requests.

Risk and Backup Plans

Include AI-related risks in your emergency planning. What happens if your AI system goes offline during a major event? Do you have manual procedures in place? Also, prepare for cases where the AI gets something wrong. For example, if someone is wrongly flagged, do you have a clear protocol for reviewing and resolving it?

This checklist helps ensure both your tools and your team are ready. A lot of cities focus heavily on buying the software and forget the people and processes around it. Reviewing this list helps you catch gaps early. Maybe your camera system can't connect to an AI solution or your staff haven't been fully trained yet. Catching that now saves major problems later.

Real readiness means having the right gear and the right mindset. The best software in the world won't deliver if your team isn't ready to use it or doesn't trust it. And the most motivated team can't succeed without proper tools. Revisit this checklist regularly as your program grows, some boxes will need to be checked again. Taking the time to do this properly will give your AI rollout the best chance of success.

Final Thoughts

AI in physical security isn't something coming in the distant future. It's already here, and it's moving exponentially fast. The question for municipal leaders is no longer whether to act, but how quickly. The tools are evolving, the threats are evolving, and the cities that move early will be the ones best prepared.

As this report has shown, early adopters are already seeing results, faster response times, smarter use of surveillance, and better deployment of staff. Waiting for AI tools to be perfect or for all the regulations to be finalized means falling behind. The world will keep moving, whether we're ready or not.

This is a call to action. Start by building awareness within your team and with your senior decision-makers. Use the information here to brief others and start the conversation. Then create a practical vision. For example, your goal might be to have AI-assisted surveillance at every high-risk facility by 2028, and to reduce security incidents by 40 percent by 2030 while making better use of staffing. Setting a clear goal makes it easier to focus your strategy and measure progress.

From there, take the first step. If you haven't already, launch a pilot. Work with a local university or tech company to test tools in a controlled environment. Talk to other cities. Security isn't a competition, and the more we share lessons and successes, the faster we all get better.

Most importantly, treat AI like a transformational shift, not just another procurement. That means investing in your people. Train them, support them, and involve them. An AI tool is only as strong as the people who trust it, use it, and know when not to. Likewise, never lose sight of the community. Public trust matters. If residents feel they've been left out of the conversation or are unsure how the tech works, that can slow or even stop your efforts. Be transparent, answer questions, and show the public how this is making the city safer for everyone.

Work across departments, too. Bring in IT, cybersecurity, legal, emergency management, and anyone else who will be affected. The more connected your planning is, the fewer surprises you'll face.

And throughout all of this, stay grounded in ethics. Just because AI can do something doesn't mean we always should. Build oversight into your process early. Even if it's not required by law yet, doing so builds credibility and gives the public confidence that this is being handled responsibly. AI will only become more powerful, so the guardrails you set now will matter more down the road.

Finally, stay flexible. Between now and 2030, there will be surprises, new technologies, new regulations, and new challenges (while writing this, OpenAI released ChatGPT Agents, which can understand goals and take actions on behalf of users, like booking a hotel). Some pilots might fail, and that's okay. Treat those as learning moments. If something works exceptionally well, don't be afraid to double down. Keep adapting.

AI gives us the chance to reimagine how public safety works. We can use it to spot risks earlier, respond faster, and make better use of our teams' time and energy. That doesn't mean removing people, it means freeing them up to focus on what matters most.

If we lead this transition with vision, integrity, and a willingness to learn, we can build safer, smarter cities by 2030. Cities that not only handle today's risks better but are ready for what's next.

So start now. Plan with intention. Pilot with purpose. Train your teams. And build trust every step of the way. That's how we lead the future of municipal security, one decision at a time.