# DIGITAL TRUST

Building Public Confidence in Public Safety Technology

**George Perera**

Public Safety Technology Executive

# A Note to the Reader

This paper is not about a particular company or product. It is about a problem I have watched grow across public safety: our ability to deploy technology is moving faster than our ability to govern it. The framework applies to AI, Drone as First Responder, connected cameras, license plate recognition, real-time crime centers, digital evidence, cloud services, cybersecurity, emergency communications, and whatever comes next.

My purpose is practical. When an agency is considering a new capability, the conversation usually centers on what the technology can do and what it will cost. Those questions matter, but they are incomplete. Leaders also need to know whether the capability can be controlled, explained, measured, and defended over time.

That leads to the central argument of this paper: Digital Trust has to become infrastructure. In a connected environment, trust cannot sit inside one policy, one platform, or one vendor's assurance package. It has to remain attached to the information, the analysis, and the decisions as they move from system to system and from one organization to another.

> **Technology is transforming public safety. Digital Trust determines whether that transformation strengthens public confidence.**

Readers can use the Digital Trust framework as the foundation and the technology-specific chapters as stand-alone modules for training, discussion, publication, or agency review.

# Contents

# Executive Summary

I have spent most of my career watching public safety absorb new technology. This moment is different. We are not adding one more device or application to the stack. We are changing the environment in which calls are received, officers are deployed, evidence is created, and decisions are made. AI is already assisting with reports and analysis. Drones can arrive before the first unit. Cameras and license plate readers extend awareness across jurisdictions. Real-time centers combine information that used to sit in separate systems. Cloud platforms carry evidence and core operations. Next Generation 911 is bringing richer data into the first seconds of an emergency.

That creates enormous opportunity. It also creates a level of dependency and responsibility that agencies cannot treat as somebody else's problem. The same systems that improve speed, awareness, and safety can expose sensitive data, produce incorrect results, expand beyond their original purpose, or fail at the worst possible time.

Most agencies can buy advanced technology. Buying it is no longer the hard part. The hard part is knowing whether the agency can govern it, explain it, measure it, and remain accountable when it produces the wrong answer or stops working altogether.

I use the term Digital Trust to mean justified confidence that public safety technology is tied to a legitimate mission, protected, governed, transparent enough to be understood, and supported by evidence of how it actually performs. The word justified matters. Trust is not created by a press release, a respected brand, or an expensive contract. It is earned through controls that work, policies that are followed, decisions that can be reconstructed, and results that can be shown. At enterprise scale, those capabilities become Trust-as-Infrastructure: a common way to preserve provenance, integrity, verification, accountability, and control across otherwise separate systems.

The Digital Trust Compass turns that idea into five questions: Mission, Governance, Security and Privacy, Transparency, and Evidence. I would ask those questions before procurement, during implementation, after a major update, following an incident, and before retirement. A system does not remain trustworthy simply because it passed a review years ago.

The Compass is not enough by itself once information begins crossing the enterprise. Identity, provenance, policy, validation, human authority, and outcome evidence have to remain connected as a workflow moves through different platforms. Interoperability moves the data. Trust orchestration preserves the context needed to rely on it.

The chapters that follow apply the framework to AI, DFR, connected cameras and LPR, real-time crime centers, digital evidence, cloud and cybersecurity, and NG911. Each technology presents a different risk profile, but the leadership obligation is consistent: define the mission, establish limits before scale, protect the system and the data, keep people accountable for consequential decisions, explain the capability honestly, and prove that it is working as represented.

The paper also separates two responsibilities that leaders often blend together. Trust by Design applies to new capabilities, when requirements can still shape architecture, procurement, and implementation. Trust through Stewardship applies to the environment already in place. Most agencies inherit systems they did not choose and contracts they did not write. Those systems still have to be inventoried, constrained, improved, validated, or retired.

This is not an argument for slowing innovation. It is an argument for making innovation durable. I have seen technology deliver real operational value and still lose support because the agency could not

explain it, measure it, or show that it was under control. A disciplined program is far more likely to survive the first controversy, budget cycle, audit, or change in leadership.

> **The standard is not perfect technology. It is technology whose value, safeguards, and responsible use can be shown.**

PART I

# The Technology Revolution

> **Technology is no longer supporting the operating environment from the sidelines. It is becoming the operating environment.**

THE TECHNOLOGY REVOLUTION

## 1. Public Safety Has Entered a New Operating Environment

Public safety has always adapted to technology. Patrol cars extended an officer's reach. Radios changed command and coordination. CAD changed dispatch. Mobile data changed what personnel could know before arriving. DNA reshaped investigations, and body-worn cameras changed evidence and accountability. Those were major shifts. What is happening now is broader because the systems are no longer operating independently.

Consider a single 911 call. Before the first unit arrives, the event may generate location data, live video, sensor alerts, an automated summary, a drone launch, camera and plate searches, and analytical support. The information can then move through communications, a real-time center, patrol, investigators, prosecutors, and the courts. At each step, a different system, vendor, or organization may control part of the workflow.

When that environment works, response is faster, personnel know more, and evidence moves cleanly. When it fails, the impact is immediate. Calls stack up. Video disappears. Access is lost. Evidence is delayed. The public does not care which provider, integration, or contract clause caused the failure. The agency owns the consequence.

That is why a camera is no longer just a camera, a drone is no longer just an aircraft, and an AI assistant is no longer just software. Once those tools are connected to data, analytics, evidence, and operational decisions, the surrounding workflow matters as much as the product itself.

Integration creates advantage, but it also creates dependency. Data gathered for one purpose can quietly become available for another. A configuration change can disrupt several workflows. A vendor update can change the behavior of an AI model overnight. One compromised identity may reach far beyond the system where it began. Connected systems leave very little room for casual governance.

A modern technology decision therefore cannot end with a feature chart, a demonstration, and a price comparison. Leaders need to understand who owns the data, who can enter the environment, what decisions the system may influence, how it fails, how it changes, and what evidence will still be available years after the sales team has left.

### The operational promise

- Faster and more informed response to emergencies.
- Greater situational awareness before personnel enter uncertain or dangerous environments.
- More efficient review of large volumes of video, records, and digital evidence.

- Improved coordination among communications, patrol, investigations, intelligence, prosecutors, and partner agencies.
- Better documentation, auditability, and continuity of evidence.
- New opportunities to prevent harm rather than respond only after it occurs.

### The strategic reality

Technology will keep moving whether public institutions are ready or not. The choice is not modernization or no modernization. It is deliberate modernization or adoption by drift: features enabled because they are available, integrations built without clear ownership, and policy written only after a problem becomes public.

> **Technology is not the limiting factor. Responsible deployment - and the ability to sustain confidence - increasingly is.**



# 2. The Trust Gap

Most public safety technologies follow a familiar path. Someone identifies a real operational problem. A vendor demonstrates a capability that appears to solve it. Leaders can immediately see improvements in response time, coverage, staffing, safety, or investigative reach. Then the questions that determine whether the program can survive begin to surface.

Who can see the data? How long is it kept? Can it be used for a different purpose? How often is the alert wrong? Who verifies it before action is taken? Can a search be reconstructed months later? What is shared outside the agency? What changes when the vendor updates the product? These are not secondary questions. They are the program.

I call the distance between demonstrated capability and justified public confidence the Trust Gap. A product can prove what it does in a ten-minute demonstration. An agency may need months or years to prove that the same capability is being used within clear limits and under real accountability.

A drone can show that it reaches a scene faster than a patrol car. A camera network can locate a vehicle. AI can produce a polished summary in seconds. Cloud services can show impressive scale and uptime. None of those demonstrations answers whether the technology will stay within purpose, whether misuse will be detected, or whether the agency can explain a failure without pointing at the vendor.

The gap widens when deployment outruns policy, training, security, public explanation, and independent review. It also widens when leaders accept assurances that cannot be verified. A system may work exactly as advertised and still lose support because the institution has not shown that it understands or controls the consequences.

The gap closes through visible discipline. Name the purpose. Name the owner. Protect the data. Train the users. Audit meaningful activity. Acknowledge limitations. Show the operational result. When something goes wrong, correct it and say what changed.

This is part of mission performance, not a public-relations exercise. Confidence affects funding, employee adoption, legislative support, information sharing, litigation exposure, and community cooperation. I have watched technically sound programs stall because nobody could clearly explain why the capability existed, who controlled it, or what value it actually produced.

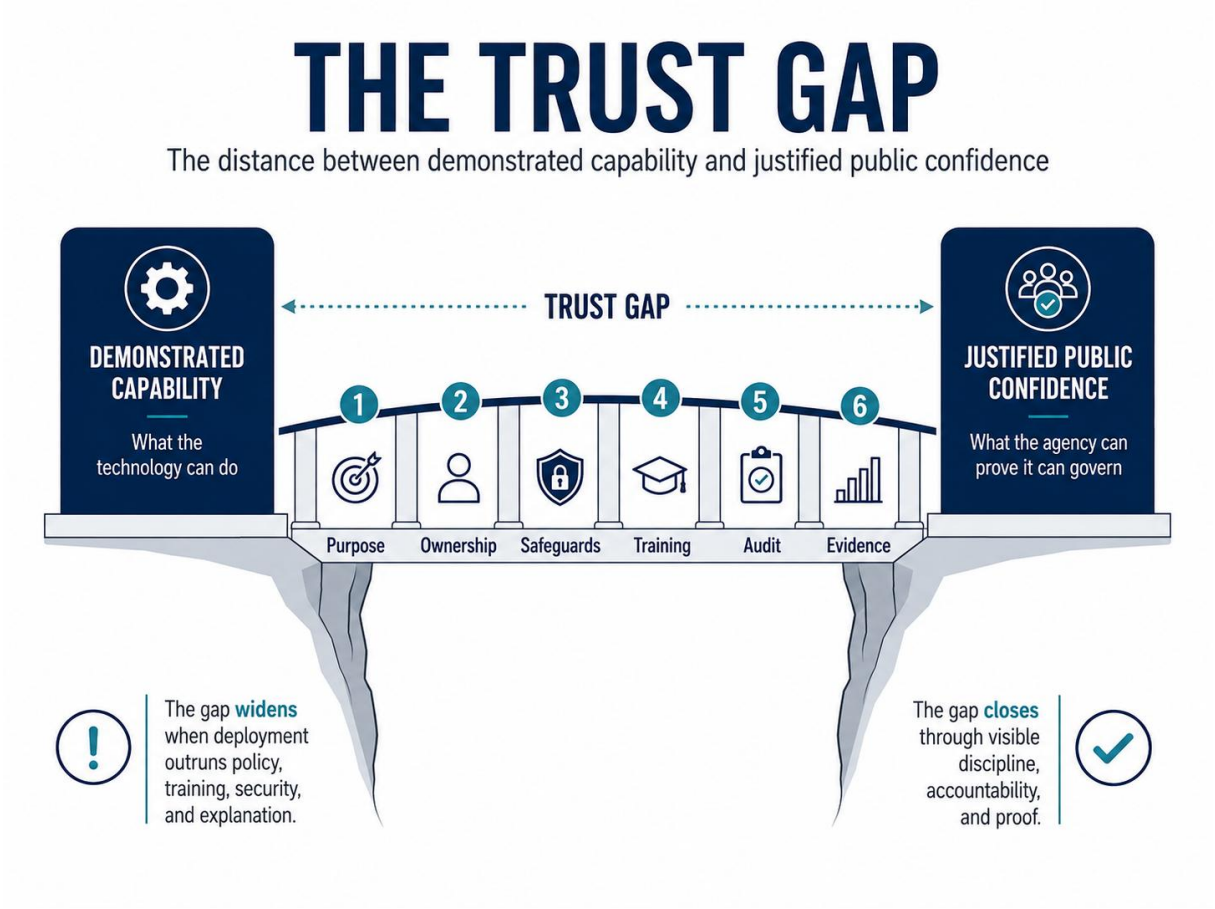


*Figure 1. The Trust Gap*

## Why the gap is growing

- Scale: a single platform can collect, search, connect, and retain information across an entire region.
- Speed: automated alerts and recommendations can influence action before traditional review occurs.
- Visibility: technology use, mistakes, and controversies can become public immediately.
- Opacity: advanced analytics and AI may be difficult for non-specialists to evaluate.
- Interdependence: agencies increasingly rely on vendors, cloud providers, data partners, and shared systems.
- Permanence: digital information can be copied, combined, and reused long after the original event.
- Expectation: communities and elected officials increasingly expect proof of safeguards, not general assurances.

## Public questions are governance signals

Questions about retention, access, accuracy, bias, sharing, privacy, security, and accountability are sometimes dismissed as objections from people who do not understand the technology. I view them differently. They are governance signals. They show leaders where the program is most likely to lose legitimacy, create legal exposure, or fail the people expected to operate it.

> **The public does not need every technical detail. It needs credible reasons for confidence.**

PART II

# The Digital Trust Framework

> **Digital Trust is not a philosophy added after deployment. It is an operating standard.**

THE DIGITAL TRUST FRAMEWORK

## 3. Defining Digital Trust

> **Digital Trust is justified confidence that public safety technology is mission-aligned, secure, governed, transparent, accountable, and supported by credible evidence.**

The definition begins with justified confidence. A public institution cannot reasonably ask people to trust a system because the company is well known, the demonstration was impressive, or the contract was expensive. Confidence has to rest on things that can be examined, tested, and corrected.

Mission alignment asks whether the technology solves a defined public safety problem. Security addresses systems, identities, data, integrations, and resilience. Governance establishes who may use the capability, for what purpose, under whose authority, and within what limits. Transparency requires an honest explanation of purpose, safeguards, limitations, and oversight without exposing sensitive tactics. Accountability means consequential actions can be attributed, reviewed, corrected, and stopped. Evidence means the agency can demonstrate its claims instead of repeating them.

Digital Trust is broader than cybersecurity. A system can be technically secure and still be used outside its stated purpose. A policy can exist while nobody follows it. A model can be accurate on average and still be impossible to defend in a particular case. A platform can deliver value while leaving the agency unable to explain who accessed the data or why a decision was made. Trust comes from the operating environment as a whole.

Digital Trust does not require universal agreement. Public agencies make difficult decisions, and some people will oppose them even when the process is sound. The objective is not unanimous approval. It is to show that the purpose was legitimate, the safeguards were reasonable, the use was disciplined, and the institution remained accountable.

The public does not need to understand source code, cryptography, cloud architecture, flight controls, or machine-learning models. It does need confidence that the people responsible for those systems understand them well enough to impose limits, detect misuse, recognize failure, and act honestly when the technology is wrong.

Responsibility is shared, but it is not equal. Providers must design products that support responsible public use and make their claims testable. Agencies remain responsible for the mission and for governing the capability over its full lifecycle. Elected officials should examine more than purchase price and promised features. Independent assessors have a role when neither the vendor nor the agency should be the only source of assurance.

Applied consistently, those responsibilities become more than a set of policies. They become infrastructure: a common way to preserve purpose, provenance, integrity, authority, and evidence when information crosses system, vendor, and organizational boundaries.

### Seven principles

1. Technology must begin with a clearly defined mission problem, not a search for a use case after purchase.
2. Human accountability remains even when software recommends, prioritizes, detects, or automates.
3. Trust requirements begin before procurement and continue through retirement.
4. Operational value and public confidence are complementary measures of success.
5. Claims about performance and safeguards should be supported by evidence appropriate to the level of impact.
6. Trust is shared: agencies and vendors have different responsibilities, but neither can create it alone.
7. Trust must travel with data, evidence, analytics, recommendations, and decisions as they move across systems and organizations.

### What Digital Trust is not

- It is not a substitute for operational performance.
- It is not a communications campaign designed to make an unacceptable practice appear acceptable.
- It is not a promise that technology will never fail.
- It is not a requirement that every technical detail be made public.
- It is not a checklist completed once and placed on a shelf.
- It is not opposition to innovation.

> **Public safety should not have to choose between technology and trust. It has to advance both.**

**THE DIGITAL TRUST FRAMEWORK**

# 4. The Digital Trust Compass

The Digital Trust Compass is deliberately simple. It gives an executive team, procurement committee, project manager, or community group five questions that force the right conversation. It does not eliminate judgment. It makes the basis for that judgment visible.


*Figure 2. The Digital Trust Compass*

| Dimension | Core Question | Evidence of Readiness |
|---|---|---|
| **MISSION** | What problem are we solving, and what measurable public safety outcome should improve? | Need statement, baseline measures, expected benefits, success criteria |
| **GOVERNANCE** | Who is authorized to use the technology, for what purposes, under whose oversight, and with what accountability? | Policy, roles, training, audit process, retention and sharing rules, exception handling |
| **SECURITY & PRIVACY** | How are systems, identities, data, integrations, and individual rights protected? | Architecture, access controls, encryption, logging, privacy review, resilience and incident plans |
| **TRANSPARENCY** | Can we explain the purpose, scope, safeguards, limitations, and oversight clearly and honestly? | Public fact sheet, executive briefing, policy summary, complaint and inquiry process |
| **EVIDENCE** | What proves the technology is secure, effective, accurate enough, compliant, and operating as represented? | Testing, audits, validation, performance metrics, error analysis, independent assessment |

## From a framework to trust infrastructure

The Compass can evaluate a single technology, but the larger problem is enterprise-wide. CAD, RMS, evidence platforms, cloud services, analytics, AI, and operational systems will not all be replaced by one product. What many agencies lack is a vendor-neutral, storage-agnostic way to connect the evidence of how those systems were used.

I call that function trust orchestration. It keeps purpose, identity, provenance, integrity, policy, access, validation, human authority, and outcomes connected across a workflow. The test is practical: can the agency show what entered the process, what changed it, who or what acted, what was recommended, what a person accepted or rejected, and what happened next?

**Interoperability moves information. Trust orchestration makes it understandable, verifiable, and accountable.**


*Figure 3. Trust Orchestration Across the Enterprise*

The Compass is technology-neutral on purpose. I would use the same five questions for an AI assistant, a camera network, a drone program, a digital evidence platform, a cloud migration, a 911 integration, or a legacy system that has operated for years without a serious governance review.

It is also not a one-time exercise. A system that was acceptable when purchased may look very different after new use cases, analytics, data sources, model updates, users, or integrations are added. Material changes, new uses, or expansion of authority should reopen the trust review. It should not slip into operations as a routine software update. Continuous review matters because trust can erode long after initial implementation.

The answers have to produce evidence. A policy on a shared drive is not the same as a policy that is trained, enforced, and audited. A vendor statement is not an independent assessment. A claim of operational value becomes credible when the agency can show response times, outcomes, errors, user behavior, and community impact.

Before full-scale deployment, I would add a formal evaluation gate. It should consider operational performance, user feedback, community perception where appropriate, readiness of policy and training, and whether the capability is delivering the benefit the agency promised.

## Five lifecycle gates

1. Need: Define the operational problem and why technology is an appropriate response.
2. Procurement: Convert trust requirements into specifications, contract terms, evaluation criteria, and proof obligations.
3. Implementation: Configure controls, train users, establish oversight, test integrations, and prepare public communication.
4. Operations: Monitor performance, audit use, review incidents, measure outcomes, review user feedback and community perception where appropriate, update policy, and validate material changes, new use cases, and expanded integrations.
5. Retirement: Preserve records, export data, revoke access, terminate integrations, document disposition, and confirm vendor obligations are complete.

## A decision standard

The Compass does not generate an automatic approval or rejection. It exposes uncertainty so leaders can make a deliberate decision. Some capabilities will offer high value with manageable risk. Others may need a limited pilot, narrower authority, stronger safeguards, or outside validation. A few will not be ready. That is not a failure of innovation. It is responsible leadership.

> **Capability alone is not a green light. The agency must be able to show how the technology will be governed and how its claims will be verified.**

PART III

# Applying Digital Trust

> **The value of a framework is demonstrated through use. The following chapters apply the Digital Trust Compass to the technologies reshaping public safety today.**

APPLYING DIGITAL TRUST

## 5. Artificial Intelligence

AI may become the most consequential technology public safety adopts in this generation. It can summarize reports, search enormous collections, translate and transcribe, identify patterns, assist investigators, automate routine work, and help personnel navigate systems that have become too complex. Agentic AI may eventually coordinate tasks across several applications and act more like an operational partner than a traditional software tool.

The attraction is obvious to anyone who has worked inside a public safety organization. We are overwhelmed with information and short on time. Used well, AI can return attention to judgment, investigation, communication, and service. It can also give smaller agencies analytical capability they could never staff on their own.

But AI concentrates several trust problems at once. It can be persuasive and wrong. It can carry limitations from its training data. Its behavior can change after an update. Sensitive information can leave the agency through careless use. People can defer to a recommendation because it arrived quickly and sounded certain. In some systems, the answer appears without enough provenance to explain where it came from.

The first governance mistake is treating every use of AI as equally risky. A tool that corrects grammar in an internal email is not the same as one that identifies a person, ranks investigative leads, assesses a threat, recommends enforcement action, or writes material into an official record. The controls should follow the consequence, not the label.

A low-impact use may require an approved platform, clear data rules, training, and ordinary review. Higher-impact uses demand more: a documented purpose, legal and privacy review, testing in the agency's actual environment, logging, model and version awareness, quality control, error monitoring, and meaningful human verification. Autonomous action, where allowed at all, should be narrow, explicit, and easy to stop.

The phrase human in the loop is not enough. I have seen workflows where the person is technically responsible but is given neither the time nor the information needed to challenge the system. Human review is meaningful only when the user can trace important claims, recognize uncertainty, reject the recommendation, and document why.

AI-generated content and AI-derived insight also require different controls. A generated report, image, or narrative needs labeling and review. An analytical linkage needs provenance and corroboration. An automated action needs defined authority, thresholds, monitoring, and an immediate override mechanism. Treating those uses as interchangeable creates false confidence.

For a consequential use, the agency should be able to reconstruct the decision path: source material, model and version, material prompts or rules, relevant configuration, output, human review, acceptance or rejection, and the action that followed. Not every AI interaction belongs in an evidence package. But when the output materially affects a person, an investigation, or an official record, the institution should be able to explain how it got there.

Good AI governance should control change without freezing the technology in place. New use cases, model updates, integrations, incidents, and material performance changes need a route back through review. The goal is not to keep personnel away from AI. It is to prevent invisible, untested, or unaccountable AI from becoming routine.

## A responsible AI baseline

- An inventory of approved AI systems and use cases.
- Impact classification that matches controls to consequence.
- Clear rules for sensitive, criminal justice, personnel, and public data.
- Meaningful human review for consequential outputs.
- Logging of material prompts, outputs, sources, model versions, and approvals where operationally appropriate.
- End-to-end decision lineage for consequential uses, including human acceptance, rejection, correction, or override.
- Testing for accuracy, failure modes, bias, security, and performance within the agency's actual use case.
- Visible labeling of AI-generated or materially AI-modified official content.
- A process to report errors, suspend use, correct records, and revalidate after significant changes.

**Questions leaders cannot delegate**


*Figure 4. Questions Leaders Cannot Delegate*

- What specific task is the AI performing, and what decision could its output influence?
- What information enters the system, where does it go, and can the provider use it to train or improve other models?
- What must a human verify before the output is acted upon, placed in a record, or shared?
- Can we identify sources, model versions, significant prompts, outputs, and changes?
- Can we reconstruct the path from source data through AI output, human judgment, and final action?
- How was the system tested for our actual use case, population, data, and operating environment?
- What happens when the AI is wrong, unavailable, compromised, or materially changed?
- Can we stop the system quickly and continue the mission without it?

> **AI should expand human capability without hiding human accountability.**

**APPLYING DIGITAL TRUST**

# 6. Drone as First Responder

DFR changes the order of response. An agency no longer has to wait for the first officer to arrive before seeing what is happening. In the right conditions, a drone can be overhead within minutes, confirm

whether the event is real, identify hazards, locate people, support de-escalation, and help supervisors send the right resources.

The strongest case for DFR is operational, not technological. I would judge the program by response time, coverage, the calls it can actually support, verified incidents, avoided or downgraded responses, safety outcomes, availability, weather, battery logistics, staffing, airspace, camera performance, integrations, and total cost. A remarkable aircraft does not make a successful program if the operating model cannot deliver value.

DFR also creates an immediate perception issue because the capability is visible. A resident sees a camera-equipped aircraft over the neighborhood, not a carefully designed workflow. The questions are predictable: Where does it fly? What does it record? Does it follow people? How long is video kept? Who can see the feed? What prevents the mission from expanding later?

Those questions should be answered in the design, not improvised after a complaint or news story. The agency needs clear launch criteria, authorized missions, geographic limits, prohibited uses, retention and access rules, evidence procedures, supervisor responsibilities, audit requirements, and a process for approving future expansion.

Transparency does not require publishing tactics that would endanger personnel or defeat the mission. It does require a plain explanation of why the program exists, what calls it supports, what it is not permitted to do, how performance is measured, and who is accountable.

DFR must be governed as a system, not merely as an aircraft. The dock, network, cloud platform, video stream, pilot interface, sensors, maps, evidence transfer, identities, vendor support, cybersecurity, and analytics are part of the same operating environment. A weakness in any one of them can affect the entire program.

Automation and remote supervision make exception handling even more important. Leaders need to know what happens when communications drop, GPS degrades, weather changes, a dock fails, the aircraft cannot return safely, the vendor service is unavailable, or a human operator has to take control without warning. When the aircraft is visible in the sky, operational failure becomes a public trust event very quickly.

## Operating DFR responsibly

- A clearly documented mission tied to call types and measurable response outcomes.
- Published high-level policy describing scope, safeguards, oversight, retention, and prohibited uses.
- Flight, access, and evidence logs sufficient to reconstruct significant events.
- Regular review of response time, availability, coverage, false or unnecessary deployments, safety events, complaints, and program outcomes.
- Cybersecurity and continuity planning for aircraft, docks, communications, cloud services, user identities, and evidence transfer.
- Defined rules for automated detection, follow functions, facial or object analytics, and any future expansion of capability.
- A process for community questions, elected-official oversight, audits, and policy revision.

## Questions to settle before scale

- What calls and mission outcomes justify the program?
- Where may the aircraft operate, and what uses are explicitly prohibited?

- When does video become evidence, how long is it retained, and who may access it?
- Can every launch, flight, view, download, and transfer be audited?
- What technology or policy changes would materially expand the program and require renewed approval?
- How are safety, cybersecurity, weather, communications failure, and vendor outage addressed?
- What performance evidence will be shared with leaders and the public?

> **DFR will endure only if the program can demonstrate purpose, discipline, safety, and accountability - not merely speed.**

**APPLYING DIGITAL TRUST**

# 7. Connected Cameras and License Plate Recognition

Connected cameras and license plate recognition can extend public safety awareness across time and geography. They help locate stolen vehicles, develop leads, support missing-person cases, identify movement connected to serious crime, and give responders useful information before they arrive. When jurisdictions and private partners connect their systems, the value can grow quickly.

The risk grows just as quickly. One camera may appear limited. A regional system that combines plate readers, fixed cameras, private feeds, alerts, hotlists, historical searches, and broad sharing is something very different. The public will judge the total capability, not the label attached to each device.

Purpose limitation has to come first. The agency should be able to say what the system is for, what it may not be used for, what authorization is required for different searches, and what must be confirmed before action is taken. An alert is a lead to check. It is not proof.

Hotlist quality deserves the same attention as camera accuracy because a bad entry can create an immediate real-world consequence. Who entered it? What supports it? When does it expire? Who reviews it? How is an urgent entry communicated? How is an error corrected across every partner that received it?

Retention and sharing should follow the mission. Keeping more data for longer may help an investigation, but it also increases cost, discovery burden, exposure, and the opportunity for misuse. Every retention period and every role with historical access should have an operational reason.

Private-camera integrations add another layer of responsibility. A business, school, neighborhood, or transportation system may voluntarily share a feed, but the agreement cannot blur ownership or accountability. Consent, access, logging, security, retention, secondary use, termination, and the limits of government control should be understood before the feed becomes routine.

A camera program also needs better measures than the number of devices installed or alerts generated. I would look at validated leads, recoveries, investigative contribution, response improvement, false alerts, audit findings, policy exceptions, and cases where the technology did not perform as expected. Volume is easy to report. Value is harder, and it matters more.

## Controls that matter

- Purpose-based policy for alerts, live viewing, historical search, sharing, and investigative use.
- Quality controls for hotlists, including source, expiration, review, correction, and partner responsibilities.
- Reasonable retention supported by operational need and periodic review.

- Confirmation requirements before stops, searches, or other consequential action.
- Role-based access, detailed audit logs, misuse detection, and regular supervisory review.
- Clear agreements for partner and private-camera access, security, ownership, and termination.
- Public explanation of the network's purpose, safeguards, and oversight.

### Questions the agency should answer

- What specific public safety purposes justify collection and historical search?
- Who may create alerts or hotlist entries, and how are accuracy and expiration controlled?
- What must personnel confirm before acting on an alert?
- How long is data retained, and what evidence supports that period?
- Which agencies, vendors, or private partners can access or receive data?
- Can we detect unusual searches, mass downloads, misuse, or access outside an assigned purpose?
- How will the agency measure both operational value and error?

> **Connected-camera trust depends on the rules, data quality, access, sharing, and accountability surrounding the network.**

**APPLYING DIGITAL TRUST**

# 8. Real-Time Crime Centers and Advanced Analytics

Real-time crime centers bring together information that used to be separated by system, unit, geography, and time. CAD, records, cameras, plate data, gunshot detection, drones, maps, corrections information, criminal intelligence, public sources, and partner data can now contribute to a common operating picture. Analytics help personnel find connections no individual could identify by manually reviewing every source.

Done well, that integration improves response and investigation. It gives field personnel context, locates relevant cameras, identifies vehicles or locations, connects related events, supports missing or endangered persons, and shortens the time between receiving information and acting on it. Done poorly, it creates a powerful environment with unclear authority and very little ability to distinguish fact from inference.

### From fusion to trusted intelligence

Fusion is only the first step. A mature intelligence environment connects information, establishes its context and provenance, reasons across sources, and recommends action with stated confidence and limitations. Only where authority is explicit should it execute a bounded task. Each step adds value, but it also raises the consequence of error and the need for traceability and human control.

> **Connect -> Understand -> Reason -> Recommend -> Execute. Authority and assurance must increase at every step.**

Aggregation creates a form of institutional power that is easy to underestimate. A record that appears routine in one system may become highly revealing when combined with several others. An analytical linkage can look authoritative even when it rests on stale, incomplete, or low-quality data. Governance

therefore has to address not only who can open each system, but how conclusions are created from the combination.

The center's role must be explicit. It may provide situational awareness, technical collection, analytical support, lead development, or referral. It should not quietly become the final authority for a criminal, intelligence, personnel, or administrative decision. The unit receiving the information remains responsible for understanding the source, limitations, and level of corroboration before acting.

Collection should be driven by Priority Information Requirements - the questions leaders and operators actually need answered - rather than by the simple fact that more data is available. PIRs impose discipline. They tie collection to a legitimate mission need, focus analysis, and connect dissemination to a decision someone is authorized to make.

Provenance is essential. A user should be able to tell the difference between a verified agency record, a partner report, an automated alert, open-source material, an analytical inference, and an unconfirmed lead. When those categories are blended, the confidence of the product can exceed the quality of the information beneath it.

At enterprise scale, the RTCC becomes part of the trust-orchestration layer. Its purpose is not to create another uncontrolled system of record or to make decisions silently. It should preserve source identity, confidence, transformations, validation, dissemination, and the handoff to the person or unit that remains accountable.

Broad connectivity does not require broad individual access. Users should see the sources and functions their role requires. Facial recognition, bulk analytics, personnel-related searches, sensitive exports, and external dissemination may warrant additional approval and documentation. Role-based access is not an inconvenience here. It is part of maintaining legitimacy.

Performance measures also need care. Counting searches, bulletins, and leads rewards activity without proving value. Better measures include time saved, response support, validated investigative contribution, serious incidents assisted, correction rates, referrals accepted, policy compliance, and cases advanced. The center should be able to show not only that it is busy, but that its work is accurate and useful.

## Guardrails for an integrated center

- A written charter defining mission, authority, services, limits, referral pathways, and accountable leadership.
- Data-source inventory showing ownership, quality, access, permitted use, retention, and sharing restrictions.
- Priority Information Requirements that drive collection, analysis, and dissemination.
- Clear labeling of source, confidence, inference, and corroboration status.
- Tiered access and elevated approval for high-impact searches or tools.
- Audit and supervisory review that examine both misuse and patterns of overreach.
- Documented handoff from analytical support to the authority responsible for investigation or action.
- Decision lineage that records how significant recommendations moved from source information to accountable action.
- Outcome measures that value validated contribution rather than activity alone.

### Questions for accountable intelligence

- What decisions is the center authorized to support, and what decisions remain elsewhere?
- Which data sources are integrated, and what quality or use limitations accompany each?
- Can users distinguish facts, automated alerts, analytical inferences, and unconfirmed information?
- What searches or tools require elevated approval or documentation?
- How are leads corroborated before enforcement, investigative, intelligence, or personnel action?
- How are errors corrected in downstream reports and shared systems?

> **The value of an RTCC is not simply that it sees more. It is that it helps the agency understand more while preserving the distinction between information, inference, and accountable decision-making.**

**APPLYING DIGITAL TRUST**

# 9. Digital Evidence

Digital evidence is no longer a specialized category. In many investigations, it is the case. Body-worn and in-car video, drone footage, fixed cameras, phones, cloud records, social media, messages, documents, photographs, audio, sensor data, and system logs can create volumes that were unimaginable a generation ago. Collection is only the beginning. The harder work is proving integrity, controlling access, finding what matters, meeting discovery obligations, and maintaining continuity over years.

Digital evidence platforms solve real problems. They centralize storage, automate upload, manage permissions, support redaction, connect prosecutors, and preserve audit trails. AI and analytics can help locate relevant material within collections no team could review manually. Those capabilities are necessary, but every transformation, derivative product, and automated action adds another question the agency may later have to answer in court.

Integrity remains the foundation. The agency has to show that evidence was lawfully collected, preserved from alteration, linked to reliable metadata, transferred through controlled processes, and accessed only by authorized users. The audit history must capture meaningful events - viewing, sharing, editing, redaction, download, deletion, and administrative action - rather than produce a thin log that cannot answer a real question later.

That integrity should be independently verifiable across the lifecycle. The storage platform may protect the file, but the institution benefits from a vendor-neutral, storage-agnostic method to verify identity, cryptographic integrity, custody events, transformations, access, sharing, and disposition across platforms. The trust layer does not need to become another repository. Its job is to preserve proof.

Synthetic and manipulated media make authenticity harder. Agencies need a disciplined way to preserve originals, document transformations, evaluate provenance, distinguish source evidence from derivative products, and identify AI-generated or AI-enhanced material. A transcript, translation, enhanced image, summary, or selected clip may help the investigation. It should never silently replace the original.

Machine-generated actions and AI artifacts should be recorded as meaningful evidentiary events when they materially affect an investigation or official record. Inputs, model and version, prompts or rules,

configuration, outputs, transformations, and human acceptance or rejection should remain connected to the source. That is decision lineage in an evidentiary context.

Discovery and public-records obligations belong in system design, not in a cleanup project after deployment. Search, export, metadata preservation, legal hold, redaction, format, chain of custody, and portability determine whether investigators, prosecutors, defense counsel, courts, oversight bodies, and the public can rely on the evidence.

Cloud hosting does not transfer the agency's responsibility to the provider. Leaders still need to understand ownership, location, encryption, administrative and subcontractor access, support practices, backup, disaster recovery, availability, exit rights, and how evidence will be retrieved if the relationship changes.

Retention requires judgment. Keeping everything forever may feel safe, but it creates cost, exposure, privacy concerns, legal complexity, and operational clutter. Retention should follow law, case status, evidentiary need, policy, and record type. Deletion must be controlled and documented, with legal holds and investigative need taking precedence.

## A defensible evidence environment

- Preservation of original evidence and metadata with verifiable integrity controls.
- Independent, storage-agnostic integrity verification that survives platform changes and interagency sharing.
- Complete, reviewable audit trails for significant access and changes.
- Clear distinction between original evidence, redacted copies, enhanced media, transcripts, summaries, and AI-derived products.
- Decision lineage for AI-generated, AI-enhanced, and machine-created evidentiary products.
- Role-based access, legal hold, retention, disposition, export, and chain-of-custody procedures.
- Validated redaction and disclosure workflows that preserve context and accountability.
- Continuity plans for vendor outage, cyber incident, litigation, merger, contract termination, or platform migration.
- Periodic testing that evidence can actually be located, exported, authenticated, and presented when required.

## Questions that expose weak links

- Can we prove the integrity and history of evidence from collection through final disposition?
- Can integrity and chain of custody be independently verified without relying solely on the system that stores the evidence?
- Who can view, share, alter, redact, export, delete, or administer evidence, and can each action be audited?
- How are AI-generated summaries, transcripts, enhancements, and selections labeled and validated?
- Can the agency meet discovery, public records, legal hold, and court requirements at scale?
- What happens to evidence during a prolonged outage, cyber incident, vendor failure, or contract termination?
- Can data be exported with complete metadata and usable chain-of-custody information?
- Are retention and deletion rules lawful, purposeful, and consistently enforced?

> **Digital evidence earns confidence when the agency can show where it came from, what happened to it, and why it remains reliable.**

**APPLYING DIGITAL TRUST**

# 10. Cloud, Cybersecurity, and Operational Resilience

Cloud computing changed what public safety agencies can deploy and how quickly they can deploy it. Capabilities that once required substantial local infrastructure and specialized staff can now be delivered as managed services for evidence, analytics, AI, collaboration, backup, and regional information sharing.

But cloud is not a destination, and it is not a guarantee of security. It is a responsibility model. The provider may secure the underlying service, but the agency still owns the mission, identities, access, data, configuration, legal compliance, continuity, and consequences of failure. In a real incident, nobody outside the contract team will care where that boundary was supposed to be.

No single provider's control plane gives leadership the whole picture in a multi-platform environment. The agency needs an enterprise view of identities, data flows, integrations, configuration changes, support access, provenance, and assurance that remains meaningful across cloud and on-premises boundaries.

In practice, the most common cloud failures are not exotic. They are excessive privilege, weak identity controls, poor configuration, unmanaged integrations, unmonitored support access, incomplete logging, unclear data classification, and the assumption that someone else is handling it. Easy connectivity can turn a small mistake into broad exposure.

Identity is the control plane. Multi-factor authentication, least privilege, privileged-access management, service-account control, conditional access, timely provisioning, and rapid revocation are foundational. The agency also needs to know when provider or subcontractor personnel entered the environment, who approved the access, how long it lasted, and what they did.

Security also has to be operationally aware. A control that blocks unauthorized access is valuable, but if it locks out dispatchers, investigators, or responders at the wrong moment, it has created a mission failure of its own.

Resilience matters as much as prevention. A platform may be secure and still fail the mission if it becomes unavailable at the wrong time. Backup, recovery, geographic redundancy, offline capability, internet and identity dependencies, vendor recovery commitments, and degraded operations have to be tested rather than assumed.

Cybersecurity incidents are trust incidents. A breach, ransomware event, outage, or compromised account can affect evidence, investigations, communications, employee safety, privacy, and the credibility of the institution. Technical containment is only one part of the response. The agency must preserve evidence, meet legal obligations, continue operations, communicate verified facts, correct weaknesses, and prove that recovery is complete.

The level of independent assurance should match the consequence. A questionnaire and a certification logo may be enough for a low-impact service. Mission-critical systems and sensitive data environments require more: architecture review, vulnerability and configuration evidence, penetration testing where

appropriate, incident history, supply-chain visibility, recovery exercises, and contractual rights to obtain continuing proof.

An exit strategy is not procurement paperwork. It is a control against dependency. Before the agency becomes reliant on a service, it should know how data, evidence, logs, configurations, and operating processes will be transferred or preserved. Portability, deletion certification, transition support, continued access during a dispute, and surviving obligations should be settled while the agency still has leverage.

### Cloud controls that have to work

- Explicit mapping of agency and provider responsibilities.
- Strong identity, least privilege, privileged access, logging, and lifecycle controls.
- Data classification and rules for storage, processing, sharing, geographic location, and support access.
- Independent security review appropriate to mission impact and sensitivity.
- Tested backup, recovery, failover, degraded operations, and incident response.
- Continuous monitoring of configuration, vulnerabilities, accounts, integrations, and provider changes.
- Contractual portability, transition, evidence preservation, and secure deletion obligations.

### Questions for continuity and control

- Which mission functions and data depend on the service, and what happens when it is unavailable?
- Who administers the environment, including provider personnel, and how is privileged access controlled and reviewed?
- Where is data stored and processed, and what subcontractors or services can access it?
- What evidence demonstrates security beyond vendor self-attestation?
- How quickly can the agency detect, contain, recover from, and communicate a significant incident?
- Can the agency continue critical operations during identity, internet, cloud, or vendor disruption?
- How will data, logs, configurations, and operational capability be transferred or preserved at exit?

> **Cybersecurity protects the continuity, evidence, privacy, and credibility on which public safety operations depend.**

**APPLYING DIGITAL TRUST**

# 11. Next Generation 911 and Connected Response

Next Generation 911 is changing the first point of contact between the public and emergency services. A call may now include precise location, text, images, video, telematics, medical information, sensor data, automated transcription, translation, and links to other public safety systems. The communications center is becoming a data-rich operational hub rather than a voice-only gateway.

The potential benefit is immediate. Better information can improve call classification, accessibility, responder awareness, resource selection, and safety. It can also connect 911 to DFR, cameras, mapping, responder devices, records, and real-time operations during the first critical minutes of an incident.

Reliability comes first because the public assumes 911 will work. When it does not, the caller does not distinguish among the carrier, cloud provider, integrator, and agency. New features cannot weaken the availability, clarity, accessibility, and resilience of the core service. Failover, redundancy, routing accuracy, location quality, cybersecurity, backup communications, and surge performance are operational requirements.

Richer data also creates new responsibilities. A caller may send video or medical information without understanding how it will be retained, shared, or used as evidence. Automated transcription, translation, prioritization, and decision support can help, but a mistake can alter the response. Call-takers and dispatchers must recognize those limitations and retain clear authority to correct the technology.

Accessibility and equity are mission issues, not secondary design considerations. The system has to work for people with disabilities, different languages, limited connectivity, older devices, and varying levels of digital literacy. A platform can be technically advanced and still fail if a meaningful part of the community cannot use it.

Every connection from 911 into another system expands the trust boundary. Each integration needs a defined purpose, owner, security model, data flow, retention rule, monitoring plan, and fallback process. More connections create value only when the environment remains understandable and supportable.

## Requirements for connected response

- Reliability and continuity requirements that protect the core emergency communications mission.
- Validated location, routing, transcription, translation, and decision-support performance.
- Human authority to correct automated classifications, recommendations, and data.
- Clear handling of caller-provided media, sensitive information, evidence, retention, and sharing.
- Accessibility testing across language, disability, device, network, and user conditions.
- Mapped integrations with defined ownership, security, monitoring, and fallback procedures.
- Regular exercises for cyberattack, carrier failure, cloud outage, surge, disaster, and degraded operations.

## Questions for the communications mission

- Does each new capability improve response without reducing reliability of the core 911 mission?
- How accurate are automated location, transcription, translation, prioritization, and routing functions?
- What information is retained, shared, or treated as evidence, and how is the caller informed where appropriate?
- Can call-takers and dispatchers recognize and override automation errors?
- How does the system perform for people with disabilities, different languages, limited connectivity, or older devices?
- What dependencies could interrupt service, and when were failover and degraded operations last tested?
- Who owns each integration and is accountable when data or service fails between systems?

> **The most advanced emergency communications system earns trust only when it remains available, understandable, accessible, and accountable under the worst conditions.**

PART IV

# Building and Sustaining Trusted Technology

> **New technology should be built and procured with trust requirements from the start. Existing technology must be brought under disciplined stewardship. Public safety needs both approaches.**

BUILDING AND SUSTAINING TRUSTED TECHNOLOGY

## 12. Trust by Design and Trust through Stewardship

Trust by Design means treating mission, governance, security, privacy, transparency, accountability, and evidence as requirements from the beginning. A system is not ready merely because it performs its technical function. It also has to support the agency's ability to control, audit, explain, validate, correct, and eventually retire it.

That requirement should shape architecture and procurement. Logs need enough detail to detect misuse. Permissions should support least privilege. Retention should be configurable. Data must be portable. Automated outputs should carry provenance. Material changes should be visible. Human review, incident response, outside assessment, continuity, and exit cannot be added after the agency has become dependent.

Most leaders do not inherit a clean environment. They inherit years of systems, contracts, data, workarounds, integrations, and vendor relationships. Some were purchased before today's security, privacy, AI, and transparency expectations existed. Others expanded gradually until actual use no longer matched the original policy. That is the reality of public safety modernization.

Trust through Stewardship is the discipline of taking that inherited environment as it exists and bringing it under control. Identify the owner, purpose, data flows, users, safeguards, limitations, and evidence of performance. Correct the urgent gaps. Constrain uses that cannot be supported. Renegotiate where possible. Retire what cannot be responsibly sustained.

Stewardship avoids two bad choices. One is complacency: assuming a long-running system is acceptable because it has not yet produced a public controversy. The other is paralysis: believing nothing can improve without replacing everything. Most agencies need a practical, prioritized program that addresses the greatest consequence first while keeping the mission running.

A single capability may require both approaches. An existing camera network may need stronger stewardship while a new analytics feature goes through Trust by Design. An evidence platform may remain in service while export, retention, access review, and AI labeling improve. An inherited cloud environment can be brought under stronger identity, logging, configuration, and recovery controls without waiting for wholesale replacement.

### A practical lifecycle

1. Define the mission: identify the problem, affected people, expected benefit, and measures of success.
2. Classify impact: determine the consequence of error, misuse, unavailability, bias, exposure, or unauthorized expansion.

3. Establish authority: assign an accountable owner and define permitted, prohibited, and exceptional uses.
4. Configure controls: implement security, privacy, access, logging, retention, human review, and continuity requirements.
5. Validate readiness: test technical performance, operational workflow, failure modes, legal compliance, and public explanation.
6. Train and communicate: prepare users, supervisors, leaders, partners, and appropriate public-facing information.
7. Monitor and improve: measure outcomes, audit use, review complaints and incidents, validate changes, and correct weaknesses.
8. Retire responsibly: preserve required records, export data, remove access, terminate integrations, and verify deletion or disposition.

### Prioritizing an inherited environment

- Start with technologies that affect constitutional rights, enforcement decisions, public-facing operations, sensitive data, evidence, or mission continuity.
- Identify capabilities that expanded beyond their original purpose or policy.
- Address missing ownership, unsupported systems, weak identity controls, unavailable logs, uncontrolled data sharing, and unclear retention first.
- Create a documented remediation decision: correct, constrain, replace, or retire.
- Review material vendor changes and new features before they become operational defaults.
- Make progress visible through recurring executive review and evidence, not one-time declarations.

### Questions for the inherited environment

- Which technologies would create the greatest consequence if wrong, misused, unavailable, or publicly misunderstood?
- Who is the accountable owner for each high-impact capability?
- Which systems lack current policy, reliable logs, access review, retention rules, validation, or public explanation?
- What can be corrected immediately through configuration or policy, and what requires contract or replacement?
- Which new features or integrations are being activated without a renewed trust review?
- What evidence will show that remediation is complete and operating effectively?

> **Trust by Design prepares the next generation. Trust through Stewardship brings today's environment under responsible control.**

**BUILDING AND SUSTAINING TRUSTED TECHNOLOGY**

# 13. Procurement, Vendors, and Evidence-Based Assurance

Procurement is often treated as a commercial process that ends with selection and signature. In public safety technology, it is one of the most important governance events. This is where an agency either

preserves leverage or gives it away. Requirements, evaluation criteria, contract language, and implementation decisions determine how much control will remain after the vendor has been selected.

Trust requirements have to be written before leverage is lost. If the agency will need usable audit logs, data portability, retention control, breach notification, restrictions on AI training, visibility into subcontractors, model-change notice, decision lineage, independent integrity verification, continuity, or secure deletion, those items belong in the requirements and the contract.

Demonstrations are useful, but they are designed to show the product under favorable conditions. The agency needs evidence from the environment in which the system will actually operate. That may require a controlled pilot, representative data, references from comparable agencies, error and false-positive information, architecture and security review, recovery exercises, and validation of critical integrations.

The proof should match the consequence. A low-impact administrative tool may need basic security and performance evidence. A system that affects identification, enforcement, emergency response, evidence, or sensitive data should face a much higher standard and continuing review. The more consequential the output, the less acceptable vendor self-attestation becomes.

Contracts also need to account for change. Modern platforms do not remain static. New models, analytics, data sources, integrations, and user interfaces may alter the trust profile without triggering a new procurement. The agency needs notice of material changes, access to release information, testing rights, configuration control, and the ability to approve or disable a feature before it becomes operational.

Vendor access deserves the same specificity as employee access. Support personnel may need temporary entry into sensitive environments, but the agency should know who connected, from where, for what purpose, under whose approval, for how long, and with what logging. Subcontractors and offshore support arrangements should never be invisible.

The strongest vendor relationship is not based on an assumption that the provider is trustworthy. It is one in which trust can be demonstrated. A capable provider should be able to explain architecture, limitations, data use, security responsibilities, auditability, testing, failure handling, and what happens when a promised control or performance level is not met.

## Requirements that should be considered before signature

- Agency ownership and permitted vendor use of data, metadata, prompts, outputs, images, and derived information.
- Location of data and processing, subcontractors, support access, and cross-border considerations.
- Identity, logging, audit, retention, legal hold, export, and secure deletion capabilities.
- Security standards, vulnerability management, incident notification, evidence preservation, and cooperation obligations.
- Performance measures, service levels, accuracy information, error reporting, and remedies.
- AI model use, training restrictions, provenance, version changes, material feature notification, and human-control requirements.
- Interoperability, interfaces, data portability, configuration export, transition assistance, and exit rights.
- Vendor-neutral provenance, integrity verification, and decision lineage that remain usable across system and organizational boundaries.
- Independent assessment, documentation access, audit rights, and continuing proof obligations.

### Proof before promises

Procurement teams should separate three things that are often blended together: a claim, an artifact, and a demonstrated result. A claim is what the provider says. An artifact is a report, certification, test result, or reference offered in support. A demonstrated result shows how the capability performs in the agency's real use case. As the consequence rises, the decision should rely less on the first category and more on the third.

### Questions before signature

- Which trust requirements are mandatory, and are they written into evaluation criteria and the contract?
- What evidence supports claims about security, accuracy, resilience, privacy, interoperability, and operational performance?
- Can the agency control retention, access, audit, export, and deletion without provider intervention?
- Can the agency independently verify critical events and lineage without relying solely on the provider's own records?
- How will material product, model, data, or subcontractor changes be disclosed and approved?
- What access can provider personnel obtain, and how is it approved, monitored, and revoked?
- What happens if the service fails, the company is acquired, the contract is disputed, or the agency must exit?
- Which obligations survive termination and how will completion be verified?

> **Public confidence cannot rest on vendor assurance alone. The agency must preserve control and obtain proof.**



# 14. Explaining Technology to the Public

Public communication about technology usually starts too late. A capability is deployed, a controversy or media inquiry arises, and leaders try to explain the program under pressure. Even a correct answer sounds defensive when the public is learning about the technology for the first time through a conflict.

Transparency should begin earlier and should match the impact of the capability. The public does not need a technical manual. It needs a clear explanation of the problem, what the system can and cannot do, what information it collects, how it is used, what safeguards apply, who is accountable, how success is measured, and where concerns can be raised.

For most residents, the core question is simpler than the agency's: What is the benefit to me? A credible answer should explain how the capability improves response, safety, service, accuracy, or accountability for the community, not just how it helps the department.

The purpose is not to win unanimous approval. It is to reduce avoidable uncertainty and show that the institution considered both benefit and consequence. In my experience, acknowledging a limitation builds more credibility than making an absolute claim that a system is safe, unbiased, accurate, or incapable of misuse.

The explanation also has to remain consistent. Operational personnel, executives, legal counsel, procurement, public information staff, elected officials, and the provider should not describe the same

capability in materially different ways. When they do, the communication problem is usually revealing an unsettled governance problem.

Elected officials need concise, decision-oriented information: the mission problem, alternatives considered, expected outcomes, total cost, principal risks, safeguards, policy, measurement plan, and the conditions that would trigger reassessment. A highly technical briefing can obscure accountability just as easily as an incomplete one.

Communities are better served by a plain-language fact sheet that is available before anyone has to request it. One or two pages may be enough if the document explains purpose, scope, retention, sharing, safeguards, oversight, meaningful outcomes, and a responsible point of contact. It should change when the program changes.

Transparency has limits. Agencies must protect sensitive tactics, personal information, system vulnerabilities, and active investigations. The obligation is not to disclose everything. It is to explain the rules and accountability surrounding what cannot responsibly be released.

## The seven-part public explanation

1. Benefit: how the technology improves safety, service, response, or accountability for the public.
2. Purpose: the public safety problem and why technology is being used.
3. Capability: what the system can and cannot do.
4. Data: what is collected, where it comes from, how long it is kept, and who may receive it.
5. Safeguards: policy, access, security, human review, audits, and prohibited uses.
6. Accountability: who owns the program, who reviews use, and how errors or misuse are addressed.
7. Evidence: what results, audits, testing, incidents, limitations, and improvements can be demonstrated.

## Language that builds credibility

- Say what is known, what is not yet known, and how uncertainty will be resolved.
- Distinguish current capability from features that may be added later.
- Explain limitations and error, not only success.
- Avoid absolute statements that cannot be sustained under real-world conditions.
- Use plain language before technical terminology.
- Commit to a recurring update, not a one-time announcement.

## Questions before the press conference

- Can a resident understand in two minutes why the technology exists and what limits apply?
- Do agency leaders, users, legal counsel, and the provider describe the capability consistently?
- What information can be published proactively without compromising safety or investigations?
- Are limitations, errors, and future expansion addressed honestly?
- Who answers public questions and how are complaints, corrections, or requests handled?
- What outcomes and audit information will be reported over time?

> **Transparency is the ability to explain purpose, limits, safeguards, accountability, and evidence before a crisis demands it.**



# 15. When Technology Fails

No complex technology environment operates without failure. Systems go down. People make mistakes. Vendors release flawed updates. Credentials are compromised. Alerts are wrong. Data is incomplete. A policy meets a situation nobody anticipated. Digital Trust does not promise perfection. It requires preparation, integrity, and correction.

The first few hours often determine whether a technical event becomes a lasting trust failure. Minimizing the problem, delaying acknowledgment, speculating, or issuing inconsistent accounts adds damage. An agency that contains the harm, preserves evidence, establishes facts, communicates responsibly, and shows what it is correcting has a better chance of retaining confidence.

Incident planning has to bring the right functions together before the event: operations, cybersecurity, legal, privacy, evidence, public information, procurement, executive leadership, the provider, and internal affairs where appropriate. A cyber event may also be an evidence event, a privacy event, a continuity problem, and an officer-safety issue. Separate plans that meet for the first time during the incident will fail slowly and publicly.

Authority should be settled in advance. Who can suspend the system, isolate an integration, revoke vendor access, roll back a model or update, preserve logs, notify affected people, approve public communication, and move the mission into a degraded mode? Those decisions should not be invented during the first conference call.

Communication should be factual and iterative. Early information will be incomplete. Say so. Leaders can state what is known, what remains unknown, what action has been taken, what the operational impact is, and when the next verified update will be provided. Silence leaves the field to speculation.

Restoring service is not the end of recovery. The agency still has to understand root cause, affected data and decisions, downstream consequences, corrective action, vendor responsibility, policy or training changes, and whether anyone was harmed by inaccurate output or lost service. A significant event should trigger revalidation of both the technology and its governance.

Lessons also have to leave the technical team. A correction buried in an internal ticket does not improve the institution. Findings should shape architecture, policy, contracts, training, communication, oversight, and future procurements. When appropriate, a public after-action summary can show accountability without exposing sensitive details.

## A trust-preserving response

1. Acknowledge: recognize the event and activate accountable leadership.
2. Contain: stop ongoing harm, preserve mission-critical operations, and control access or automation.
3. Preserve: secure logs, evidence, configurations, communications, model versions, and relevant records.
4. Inform: communicate verified facts, operational impact, protective action, uncertainty, and next steps.
5. Correct: repair the technical issue and any affected records, decisions, notifications, or people.

6. Validate: independently confirm recovery where appropriate and review whether safeguards remain sufficient.
7. Learn: revise policy, training, contracts, architecture, oversight, and public information.

## Questions to answer before an incident

- Who can immediately suspend or constrain a technology when harm is possible?
- Can the agency preserve the evidence needed to understand what happened?
- How will critical operations continue during outage or isolation?
- Who coordinates operational, technical, legal, privacy, evidence, vendor, and communications decisions?
- How will affected records, cases, alerts, or people be identified and corrected?
- What independent validation is required before full return to service?
- How will lessons be documented and incorporated beyond the technical team?

> **A failure does not automatically destroy trust. Evasion, inconsistency, and the absence of correction often do.**

PART V

# The Executive Action Agenda

> **Digital Trust becomes real when it changes ownership, decisions, contracts, controls, communication, measurement, and the way an agency responds when technology does not perform as expected.**

THE EXECUTIVE ACTION AGENDA

## 16. A 90-Day Digital Trust Agenda

An agency does not need a new bureaucracy before it can improve Digital Trust. It needs visible ownership, a small number of disciplined decisions, and evidence that those decisions are being followed. Ninety days is enough to establish the operating rhythm, expose the largest gaps, and show leadership where intervention is required.

### Days 1-30: Establish visibility

1. Name an executive owner for Digital Trust and identify legal, operational, technology, cybersecurity, privacy, procurement, communications, evidence, and community-facing participants.
2. Create a high-level inventory of public-facing, high-impact, AI-enabled, evidence-related, surveillance, emergency communications, and mission-critical technologies.
3. For each priority system, identify purpose, owner, data, users, vendor, integrations, policy, retention, logs, validation, public information, and known incidents.
4. Select three technologies for immediate Compass review based on consequence, public visibility, rapid growth, weak governance, or known uncertainty.
5. Identify urgent control gaps that require immediate action, such as unsupported access, missing logs, unclear retention, uncontrolled sharing, absent policy, or no continuity plan.
6. Map one high-impact workflow end to end - from collection through analysis, recommendation, human decision, action, evidence, sharing, and retention - and identify where identity, provenance, or accountability is lost between systems.

### Days 31-60: Establish control

1. Complete the Digital Trust Compass for the three selected technologies and document decisions, evidence, gaps, owners, and deadlines.
2. Confirm accountable program ownership and update or issue interim policy where authority, permitted use, prohibited use, access, retention, sharing, human review, or incident response is unclear.
3. Review vendor obligations for data use, support access, security, incident notification, audits, material changes, portability, and exit.
4. Create or update a one-page public fact sheet and an executive briefing for each selected technology.
5. Establish performance measures that include operational value, error, misuse, availability, audit findings, complaints, and corrective action.
6. Define an initial trust-orchestration architecture: canonical identity, provenance, integrity, decision lineage, validation, and human-authority requirements that can operate across existing platforms.

## Days 61-90: Establish assurance

1. Conduct a targeted audit, test, tabletop, or independent review appropriate to the highest-risk gaps.
2. Exercise the incident process, including authority to suspend the system, continue operations, preserve evidence, coordinate with the provider, and communicate.
3. Brief executive leadership and appropriate elected or oversight officials on findings, decisions, unresolved risks, and the remediation plan.
4. Create a recurring quarterly review for high-impact technology and a trigger-based review for material changes, incidents, new data uses, or expanded integrations.
5. Publish or internally record the evidence of improvement and identify the next three technologies for review.
6. Pilot independent integrity or decision-lineage verification on one high-impact workflow without replacing the underlying systems of record.

## The executive dashboard

A Digital Trust dashboard should be simple enough to use in an executive meeting. It should show whether high-impact technologies have named owners, current policy, completed access reviews, tested continuity, usable logs, defined retention, validated performance, known incidents, user feedback, community concerns, open corrective actions, public-facing information, and upcoming material changes. The point is not another score. It is to keep important uncertainty from remaining invisible.

| Area | What leadership should see |
|---|---|
| **Ownership** | Named executive and operational owner |
| **Governance** | Current policy, training, and prohibited uses |
| **Security** | Access review, monitoring, vulnerability and continuity status |
| **Transparency** | Current executive and public-facing explanation |
| **Evidence** | Performance, error, audits, incidents, and validation |
| **Lineage** | Source, transformations, recommendations, human decisions, and action are reconstructable |
| **Change** | Pending features, integrations, contract changes, or expansion |

| Area | What leadership should see |
| --- | --- |
| Action | Open gaps, accountable owner, deadline, and disposition |

## Ten questions for every major technology decision

**1.** What public safety problem are we solving?

**2.** What measurable outcome should improve?

**3.** Who is accountable for the technology and its use?

**4.** What data is collected, created, shared, retained, or inferred?

**5.** What could go wrong, who could be harmed, and how would we know?

**6.** What human judgment or verification remains necessary, and can the decision lineage be reconstructed?

**7.** What safeguards prevent misuse, error, unauthorized expansion, and loss of service?

**8.** What evidence supports claims about security, integrity, provenance, performance, accuracy, and value?

**9.** Can we explain the technology and its limitations clearly to the public and elected officials?

**10.** How will we monitor, verify, change, suspend, correct, and eventually retire it?

> **The 90-day agenda does not declare Digital Trust complete. It establishes the ownership, evidence, and operating rhythm needed to keep earning it.**

# Conclusion: Technology That Deserves Confidence

Public safety is not going back to a simpler technological era. AI will become more capable. Autonomous systems will take on more tasks. Cameras and sensors will become more connected. Digital evidence will keep growing. Cloud platforms will deepen their role in operations, and emergency communications will carry richer information. Capabilities that are difficult to imagine today will eventually become routine.

The profession should welcome that opportunity. Technology can save time, improve awareness, strengthen safety, expand investigative capacity, preserve evidence, and help agencies serve their communities more effectively. Innovation belongs in public safety. But it has to arrive with discipline.

The enduring test is whether an institution can show that the innovation serves a legitimate mission, operates within clear limits, is protected against misuse and failure, is explained honestly, and is supported by evidence. That is the work of Digital Trust.

Trust cannot be treated as a communications problem to manage after a controversy. It has to shape requirements, architecture, procurement, configuration, policy, training, measurement, incident response, vendor relationships, public explanation, and retirement. New systems should be designed for it. Existing systems should be brought under it.

At scale, that work must mature from a collection of governance activities into infrastructure. Agencies need a common trust layer that preserves provenance, verifies integrity, connects machine and human actions, and maintains decision lineage as information moves across systems. Without it, accountability fragments at the same boundaries where technology is becoming more connected.

The public does not expect flawless technology. It does expect competence, responsibility, accountability, and honesty from public institutions. Confidence grows when an agency can explain why a capability exists, what value it produces, what limits apply, how people and data are protected, who remains accountable, what happens when the system is wrong, and what evidence supports those claims.

That standard protects more than reputation. It protects funding, continuity, employee confidence, community cooperation, evidence, privacy, legitimacy, and the ability to continue innovating. Trust is not separate from operational performance. It is one of the conditions that allows performance to endure.

The technologies shaping the future of public safety will not be judged only by speed, reach, intelligence, or scale. They will also be judged by whether the institutions using them have earned the confidence necessary to exercise that power in service to the public.

That is the responsibility before us: not simply to deploy more advanced technology, but to build and sustain technology that deserves confidence.

> **Technology is transforming public safety. Digital Trust determines whether that transformation deserves public confidence.**

# About the Author

George Perera is a public safety technology and cybersecurity executive with decades of experience leading digital transformation, cyber investigations, digital evidence, public safety communications,

emerging technology, and enterprise modernization. His work has focused on turning advanced technology into practical operational capability while preserving security, accountability, constitutional safeguards, and public confidence.

He has held national law enforcement technology leadership roles, advised public safety agencies and industry partners, taught digital crime and cybersecurity, and contributed to the profession's development of technology, evidence, cybersecurity, and governance practices.