Delivering Trusted Services at Global Scale

Managed identity verification to help critical services stay secure, resilient, and accessible

ENTRUST

SECURING A WORLD IN MOTION

Executive Summary

Digital services rely on trusted identities. Before a citizen can claim a benefit, a customer can open an account, or an employee can log in, their identity must be established and verified with confidence.

This paper is written for leaders in government, financial services, and regulated enterprises who are responsible for high-trust services. It makes two arguments. First, identity verification delivered as a managed service is often the most effective model when the service must be secure, resilient, accessible, and scalable. The implication is practical. Identity verification in this context is not a single check. It's one part of a live digital service that must be available, resilient, and explainable every day. Organizations should choose the technology and delivery model that match their requirements, operational constraints, and risk profile and require evidence that a provider can sustain performance in production. They should evaluate identity verification as a critical part of their overall system infrastructure.

Second, a provider should be chosen based on the operational requirements of the program, not against feature lists alone. Mature buyers are no longer only asking whether a tool can capture a document or match a face in a controlled demonstration. They are asking whether a provider can operate the identity layer under real-world pressure: shifting demand, new document formats, changing regulatory expectations, evolving fraud techniques, accessibility requirements, and the need for clear audit trails. The answer depends on accountability, operational depth, compliance posture, user experience, and continuous improvement.

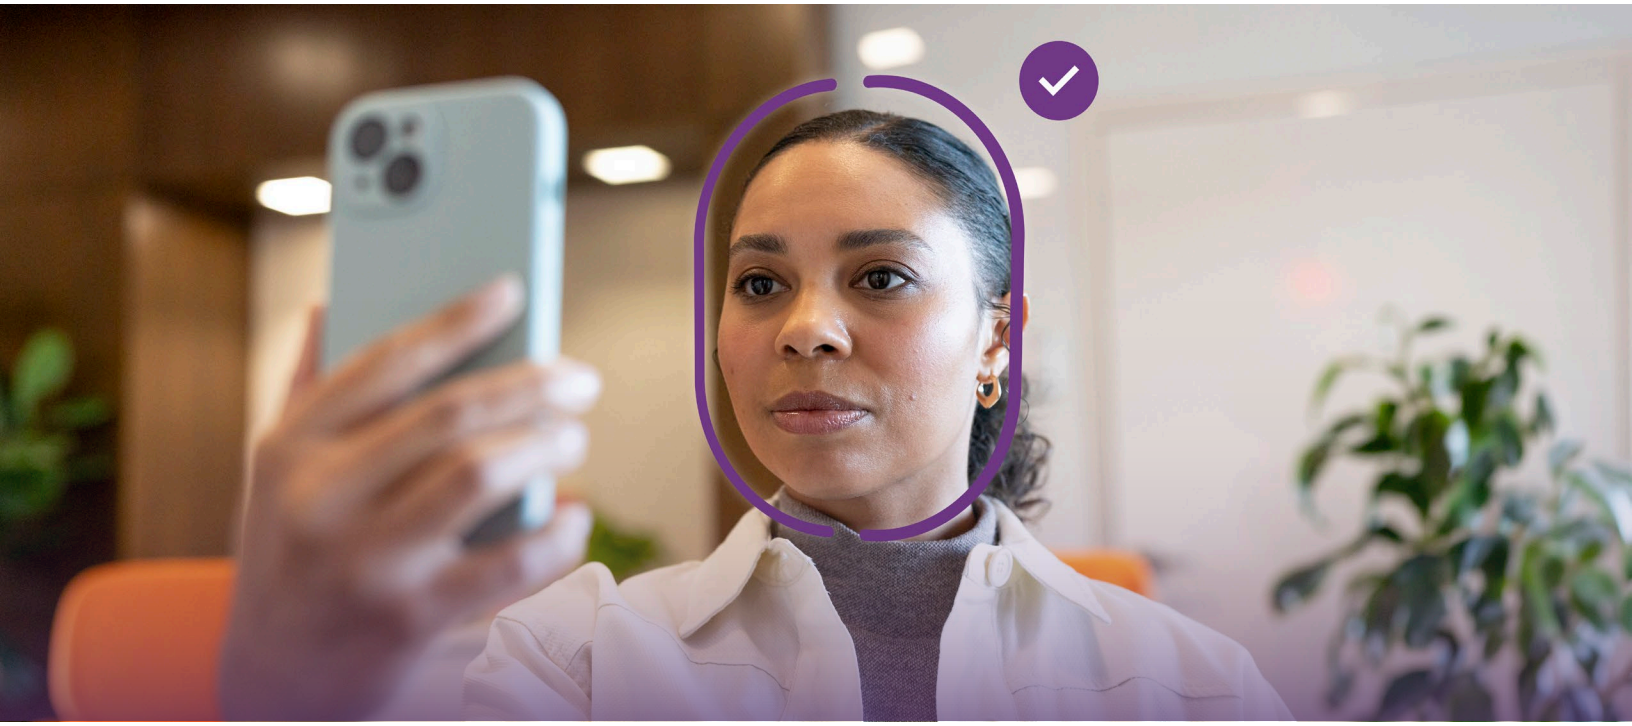

The Customer Mission

Organizations don't set out to buy identity technology. They set out to deliver safe, reliable, and scalable critical services. Identity verification is the control point that makes those services possible.

For government, the mission is to protect access to public services, reduce fraud, preserve trust, and keep essential programs available to legitimate users.

For financial institutions, the mission is to onboard and serve customers with enough assurance to satisfy regulators without creating unnecessary friction.

For enterprises, the mission is to give employees, customers, and partners trusted access to systems and services without increasing risk.

While the missions differ by sector, the operating requirement is the same – establish trust quickly, consistently, and defensibly.

The key is confidence. Organizations need to know that the person interacting with the service is who they claim to be. They also need to know that the decision can be explained if challenged, that exceptions can be handled safely, and that the process can keep running when demand spikes or the threat environment changes.

A weak identity layer does more than create technical defects. It can allow fraud to infiltrate a program, delay service delivery, exclude legitimate users, increase manual workload, and/or draw regulatory or public scrutiny. For a national- or global-scale service, even a seemingly small failure rate can become a large operational problem. A capture issue that affects a small percentage of users becomes thousands of retries. A manual review queue sized for normal traffic becomes a backlog under peak demand. A poorly governed exception process becomes an audit risk.

This is why identity verification should be framed around outcomes rather than inputs. The buyer is not simply procuring document capture, face matching, liveness checking, or data validation. The buyer is procuring confidence, continuity, and control. **The identity layer must support the service mission.**

This distinction changes the conversation for decision makers and buyers. The question becomes less about, "which tool has the strongest feature set?" and more about, "which operating model can keep this service trusted, available, accessible, and compliant over time?"

Operational Context

If the customer mission defines the outcome, the operational context defines the guardrails. Identity verification decisions are shaped by regulation, policy, risk appetite, sovereignty, privacy, auditability, accessibility, security, and the practical demands of running a live service. All of these requirements should be translated into procurement expectations before providers are compared.

In practice, the following questions need to be answered early on:

- Where can identity data be processed and stored?
- Who can access it?
- How are decisions recorded?
- What evidence is retained?
- How are exceptions handled?
- What availability is required?
- How will the service behave under peak demand?
- Can legitimate users complete the journey across different devices, languages, abilities, countries, and levels of digital confidence?

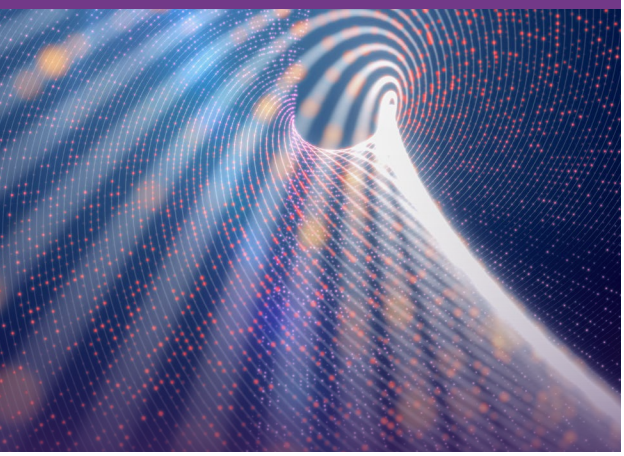

These are not secondary requirements. They determine whether the service can be trusted in production. A provider that performs well in a demonstration but cannot show auditability, resilience, accessibility, change control, and continuous service management is not a strong candidate for a critical program.

Public frameworks reinforce this direction. NIST SP 800-63A-4 defines technical requirements for proofing across identity assurance levels, reflecting the need to match identity proofing controls to the level of risk. The UK Electronic Travel Authorisation (ETA) factsheet requires applicants to provide biographic, biometric, and contact details, supporting border security by increasing knowledge about those seeking to come to the UK. These examples point to the same expectation: Identity systems must support assurance, usability, explainability, and operational control.

The operational context also forces a more disciplined view of user experience. **A high-trust service cannot simply remove friction. It must apply the right friction in the right places.** A routine, low-risk journey should be fast and accessible. A higher-risk or ambiguous case may require stronger controls, additional checks, or expert review. Good service design isn't about choosing assurance over experience or experience over assurance. It's about designing an identity journey that supports both, according to risk.

This context should shape the entire buying process. It defines the evidence a provider must bring, the operating commitments that matter, and the governance model the buyer will need after launch. It also defines the limits of a purely component-led view of identity verification. A component may be accurate in isolation, but critical services require a working system of policy, technology, operations, reporting, support, and accountability.

The Hidden Complexity of Identity Verification

Identity verification can often seem straightforward when described as a sequence: Capture a document, take a selfie, run checks, and return a decision. In reality, the user experience is just the visible layer – the tip of the iceberg – and sits atop a complex, invisible orchestration layer. This layer must absorb variation from users, documents, devices, networks, regulatory requirements, and fraud techniques.

The complexity begins at capture. Users apply in different lighting conditions, on different devices, with different camera quality, network strength, and levels of familiarity with digital journeys. A document image may be blurry, cropped, glared, expired, unsupported, or difficult to classify. A selfie may fail because of lighting, camera quality, accessories, or a genuine mismatch that requires further review. A network timeout may interrupt the session. A user asked to repeat steps too often may abandon the process.

The complexity continues in orchestration. Identity proofing may involve document authenticity checks, document data extraction, biometric comparison, liveness detection, device or session signals, data checks, sanctions or watchlist checks, risk scoring, decisioning, exception handling, and audit logging. These signals are not useful simply because they exist. They must be captured, interpreted, monitored, tuned, and governed as the program evolves.

Complexity introduces vulnerabilities that can plague large-scale programs in particular. Failure modes commonly observed during onboarding can be caused by manual data entry, repeated capture attempts, limited channel support, and slow review processes. These failures can weaken fraud controls, reduce conversions, or increase risk. Troubleshooting guidance for verification failures commonly points to low-quality images, glare, expired documents, biometric mismatches, liveness failures, unsupported documents, and technical timeouts. These issues may sound tactical, but at scale they become operational constraints.

The hidden complexity is that no single component has to fail for the service to fail. A document verification tool may perform well. A biometric tool may perform well. The integration may pass testing. Yet a surge in demand, a new document variant, a change in fraud patterns, or an unexpected manual review load can still create queues, abandonment, or service disruption.

A typical failure mode for large-scale programs looks like this: A service launches successfully and performs well in controlled testing. Volume grows. A subset of users begins failing capture because of device, document, or environmental conditions. Manual review volumes exceed the expected baseline. The provider responsible for one component points to the integrator. The integrator points to the customer's workflow. Internal teams struggle to diagnose the root cause because reporting is fragmented. The issue may start as a localized defect, but it becomes a service problem because ownership is distributed.

This is why mature buyers look past feature comparisons. They want to know who owns the journey when something breaks. They want to know how quickly documents can be updated, how exceptions are handled, how false rejects are monitored, how abandonment is measured, how fraud patterns are addressed, and how service changes are made without forcing the buyer to coordinate every moving part.

Identity verification is therefore not a one-time integration. It's a continuous operating capability. It must be designed, monitored, tuned, and improved in production. The more critical the service, the more important the operating capability becomes.

How Delivery Models Shape Outcomes

As illustrated previously, identity verification does not usually fail because a single algorithm underperforms. It fails because the service is unable to absorb change. Whether it's new document types, new fraud patterns, shifts in demand, evolving regulatory requirements, or the realities of how users behave in live environments – what happens next determines how quickly the system is able to recover from a service disruption.

The delivery model determines who owns the complexity, how quickly it's managed, and whether the service holds up under pressure. Most organizations evaluating identity verification encounter four delivery models. Each reflects a different approach to control, capability, and accountability. None are inherently wrong. Each performs differently once the service is live.

Delivery Model	Best Suited For	Strengths	Common Constraints	Primary Ownership Burden
In-house Orchestration With Specialist Components	Organizations with mature platform teams that can integrate and manage multiple identity capabilities	Flexibility, vendor choice, component-level optimization, ability to swap tools over time	Integration fragility, inconsistent reporting, shared accountability, root-cause diagnosis across vendors	Buyer owns orchestration, service management, and the middleware between components
Systems Integrator-Led Build and In-House Operations (With Specialist Components)	Organizations that need high control, have strong internal engineering and operations teams, and intend to own the service long term	Tailored design, alignment to internal standards, strong control over architecture and roadmap	Handover risk, ongoing tuning burden, internal skills dependency, slower response if identity competes with other priorities	Buyer owns live service performance after implementation.
Systems Integrator-Led Build and Managed IDV Service (With Specialist Components)	Programs where identity is part of a broader transformation and the buyer wants one program-level delivery partner	Reduced internal burden, integration expertise, structured delivery and governance	Change cycles can be slower, domain depth varies, cost can expand as requirements evolve	The SI owns delivery and operation, but identity-specialist performance still needs close oversight
Identity Verification as a Managed Service From a Single Provider	Critical services with high or variable volume, strict availability expectations, regulatory scrutiny, fraud pressure, and limited appetite to run specialist operations in-house	Clear end-to-end accountability, operational resilience, continuous tuning, faster adaptation, service-level reporting	Requires strong governance, trust in provider model, and clear evidence of production performance	The provider owns service outcomes, with the buyer governing performance, risk, and compliance

The table is useful as a starting point, but the real decision sits beneath it. Buyers should understand how each model behaves when the service moves from design to operation.

In-House Orchestration With Specialist Components

In this model, the organization assembles its own identity stack, integrating specialist providers for document verification, biometrics, data checks, fraud controls, and orchestration. It's typically chosen for flexibility: the ability to combine best-of-breed components, change vendors over time, and align identity flows to internal platform strategy.

In controlled environments, this can produce strong results. Each component may perform well in isolation, and the organization can tailor the flow to its specific requirements. The complexity emerges in orchestration. Identity verification is not a set of independent checks. It's a coordinated system where failure at one point affects outcomes across the journey.

When responsibility is distributed, failure modes tend to concentrate in the gaps between components. Integration changes can affect downstream outcomes. Vendors may evolve at different speeds. Reporting may not pinpoint the root cause of abandonment, false rejects, or manual review spikes. Accountability can become shared precisely when the buyer needs it to be clear.

On a small scale, these issues are manageable. But at a national or global scale, they become systemic risks. The model remains viable where an organization has engineering depth and operational maturity to manage these dependencies actively and continuously.

Systems Integrator-Led Build With In-House Operation

This model is often chosen when control is a priority. A systems integrator (SI) designs and implements the architecture, typically combining multiple technologies, before handing the service over to internal teams to operate.

It can work well when the organization has strong internal engineering, security, and operations capability, identity is one part of a broader transformation program, and long-term ownership of the platform is a deliberate strategic choice. The design can be tightly aligned to internal standards, and the organization retains full control over how identity is implemented and evolved.

Where this model becomes difficult is after go-live. Identity verification is not static. Documents change. Fraud patterns shift. User behavior introduces variability. Manual review volumes fluctuate. What begins as a delivery program becomes a continuous operating burden. Common failure patterns include handover gaps, operational overload, slow response to change, and fragmented ownership when multiple vendors are embedded in the solution.

Under steady conditions, the service may perform well. Under sustained pressure, the burden of ownership often becomes the limiting factor.

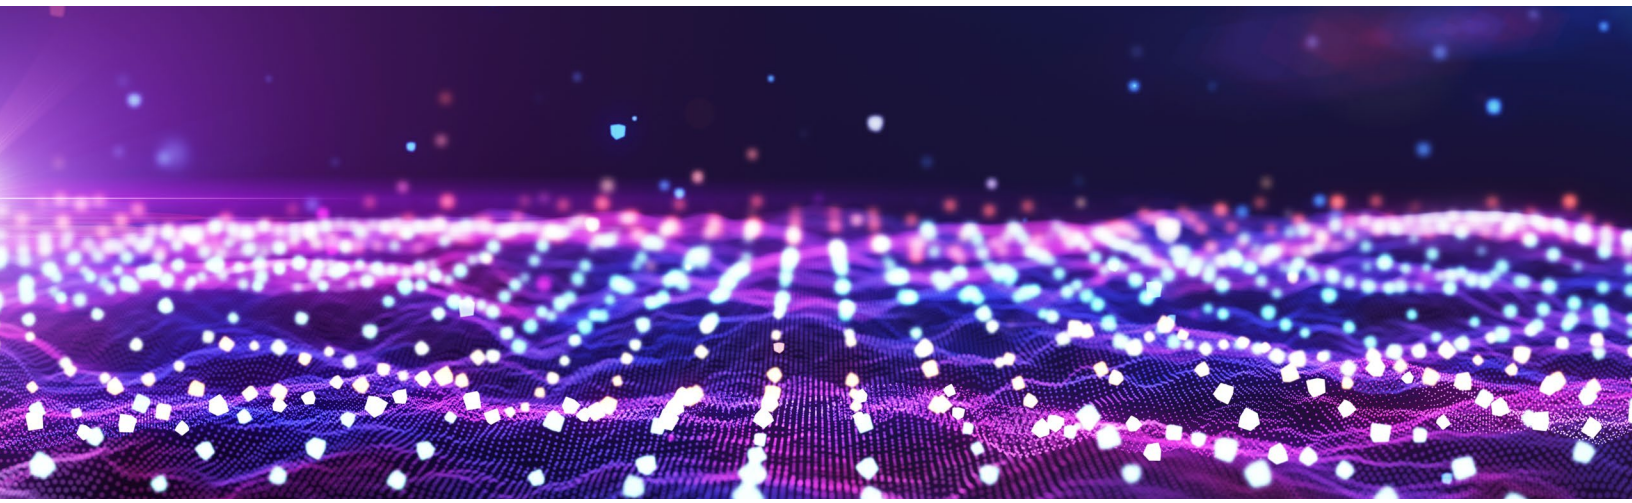

Systems Integrator-Led Build and Managed IDV Service

Here, the systems integrator retains responsibility not only for delivery but also for ongoing operation of the service, typically under a managed contract. This model is attractive when the organization wants to reduce internal burden, identity is part of a wider transformation, and there is a need for a single program-level partner.

It can provide a coherent delivery experience, especially where integration across multiple systems is required. The distinction lies in the nature of identity verification as a domain. Identity isn't just another integration component. It's a live control point, continuously influenced by fraud trends, regulatory changes, user behavior, and service demand.

Where integrators often excel in architecture and delivery, the ongoing operation of identity services introduces different demands: continuous tuning of verification performance, rapid response to emerging fraud techniques, management of exception handling at scale, and maintenance of audit-ready processes. Challenges often emerge in change velocity, domain depth, and cost predictability as ongoing changes and extensions accumulate.

This model can succeed when identity requirements remain relatively stable or when the integrator brings strong domain capability. It becomes more constrained where rapid adaptation is required.

Identity Verification as a Managed Service From a Single Provider

In this model, a single provider is responsible for designing, operating, and continuously improving the identity verification service. The defining characteristic is the operating model.

The provider is accountable for service performance and availability, verification outcomes, monitoring, tuning, optimization, fraud adaptation, compliance, and audit readiness. This model is typically adopted when identity is critical to service delivery, volume is high and variable, availability expectations are strict, regulatory scrutiny is high, and internal teams are not structured to operate identity as a continuous service.

The advantage is clarity of ownership. The complexity of the system does not disappear, but it sits with the party best positioned to manage it. New document types can be absorbed without requiring buyer-led redesign. Fraud patterns can be addressed through ongoing updates and operational controls. Exception handling can be integrated rather than bolted on. Performance can be measured against outcomes, not component activity.

This model does require buyer discipline. Governance shifts from building and managing components to setting service expectations, evaluating evidence of performance, maintaining oversight of compliance and risk, and ensuring commercial transparency. The trade-off is between control of inputs and ownership of outcomes.

There is no universal best approach. The right model depends on the nature of the service being delivered. A useful decision lens is to consider scale and variability of demand, criticality of the service, internal capability and appetite for ownership, and the rate of external change. When high scale, high criticality, limited internal capacity, and rapid change converge, the delivery model becomes a defining factor in a program's success.

An organization should worry less about which technology performs best in isolation and instead explore which model can sustain identity verification as a reliable, accountable, continuously operating service.

What Good Looks Like In Practice

A strong identity service is visible in production, not just in procurement. The test is whether it can deliver every day under real operating conditions. This is where buyers can separate technology that works in controlled situations from a service that holds up as part of critical digital infrastructure. The following signals help buyers make that distinction.

Outcome Accountability

The provider commits to measurable service levels, uptime expectations, reporting, and performance indicators that connect to the buyer's mission. Metrics should not sit in monthly exports alone. They should give operational teams timely visibility into trends, anomalies, completion rates, exception patterns, and potential fraud signals.

The purpose of measurement is not simply to prove that work occurred. It's to show whether the service is producing the outcome the buyer needs. This includes legitimate users completing the journey, risk being identified, exceptions being resolved, and the service remaining available under expected operating conditions.

Operating Resilience

A critical identity layer must handle volume spikes, document changes, release cycles, incident response, and support needs without turning every change into a buyer-led project. It should have tested runbooks, monitoring, service management discipline, and clear escalation paths.

Resilience is often most visible when conditions are imperfect. For example, let's say a country introduces a new identity document format for its citizens. Demand for the service peaks during certain seasons. A new release introduces an unexpected issue. Fraud evolves and becomes more sophisticated. The strongest signal of a good operating model is that it can detect the issue quickly and contain it. The provider communicates the issue with clarity and restores expected service levels without ever creating confusion over accountability of that issue.

Controlled Friction

A service that works only in ideal conditions is not enough. It has to be battle-tested in the real world. Users apply from different devices, networks, lighting conditions, countries, and levels of digital confidence. The identity journey should be accessible, quick, and low friction for legitimate users, while still applying stronger controls where risk warrants them.

Low friction should not mean weak assurance. It should mean that friction is purposeful. The journey should remove avoidable effort, such as repeated manual entry or unnecessary retries, while retaining the controls needed to protect the service. This balance is especially important in public services and regulated sectors, where exclusion and fraud are both unacceptable outcomes.

High Assurance With Practical Exception Handling

Automation should carry routine volume, but higher-risk or ambiguous cases need structured review, clear decisioning, and audit trails. The service must explain what happened, why it happened, and what was done next.

Exception handling is not a side process. It's part of the trust model. When ambiguous cases are not handled well, they can create manual bottlenecks, inconsistent decisions, user frustration, or compliance exposure. On the other hand, a strong service defines how exceptions are routed, reviewed, resolved, and learned from.

Compliance and Audit Readiness

Compliance should be built into the service. Organizations need evidence of what happened, when it happened, which controls were applied, how decisions were made, and how exceptions were resolved. Auditability cannot be reconstructed after the fact. It must be designed into the service from the start.

This is especially important where identity decisions affect access to public services, financial products, travel permissions, employment systems, or other regulated environments. The service must not only make decisions. It must preserve the evidence needed to defend those decisions.

Continuous Adaptation

Fraud techniques, biometric methods, document standards, user expectations, and digital identity ecosystems will keep evolving. A resilient service needs technical authority, composable architecture, and the ability to improve components and controls without forcing a full rebuild.

This is where a managed service can become more than an operating model. It can become a way to keep the identity layer current with best-of-breed technology without getting locked into a single vendor or microservices provider. Continuous monitoring, analytics, fraud intelligence, and controlled service evolution help prevent a program from aging in place between procurement cycles.

The distinction is simple: Some technology works well in controlled situations. **A managed service for critical programs has to deliver in ordinary conditions, peak conditions, and adverse conditions.** That is the difference between a capability and an operating service.

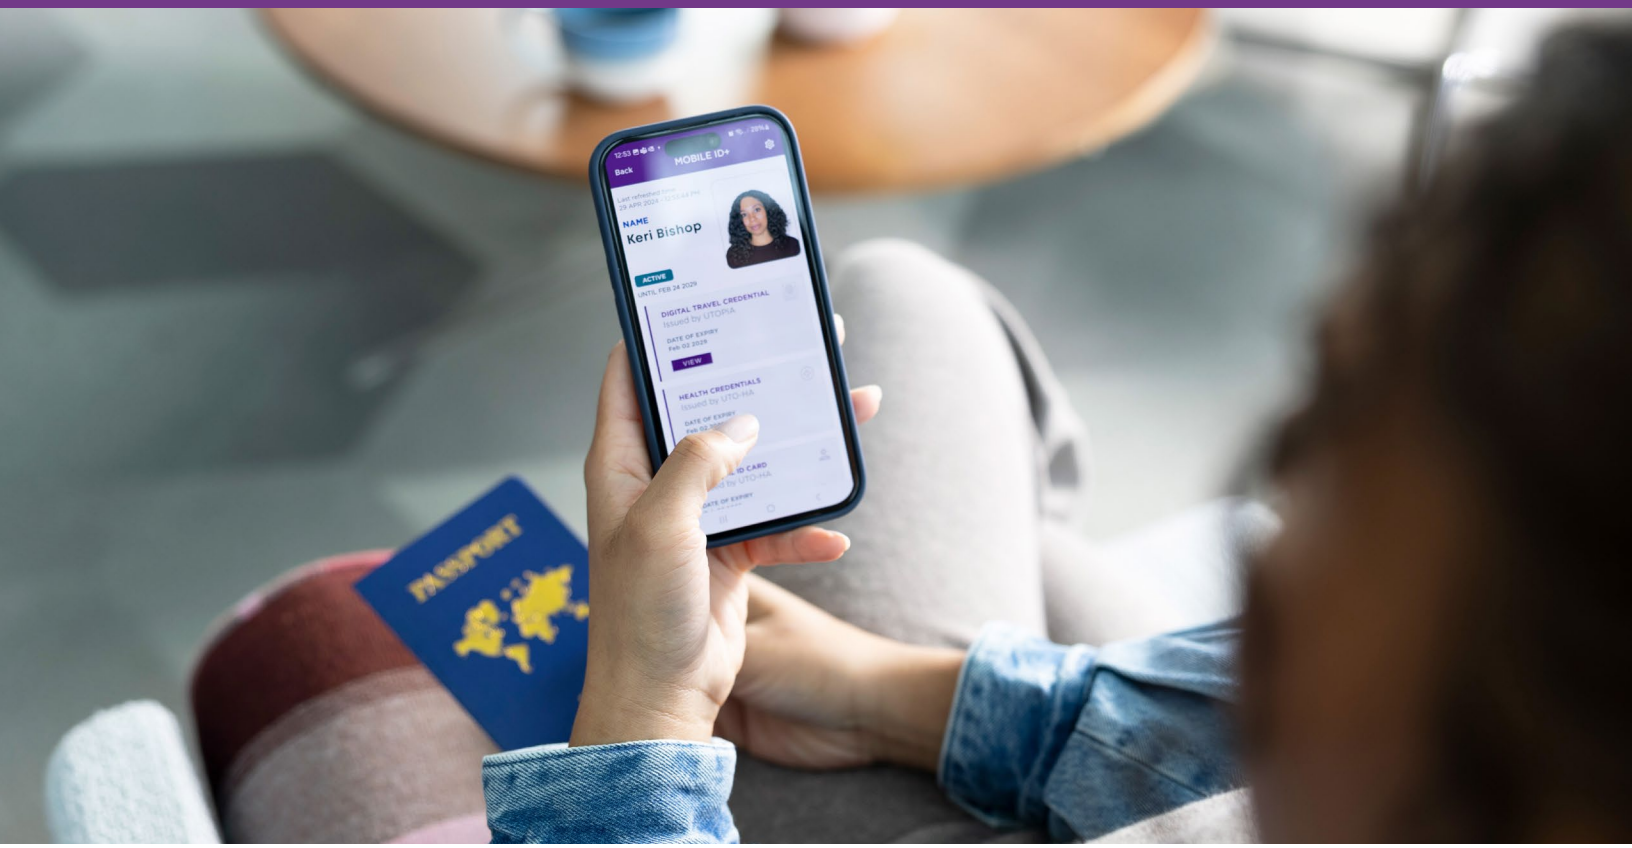

Case in Point: UK Home Office ETA Rollout for Visa-Exempt Nationals

Objective

The UK Home Office set out to introduce a mandatory digital permission to travel to the UK while preserving a seamless experience for legitimate visitors. The goal was to strengthen border security without disrupting travel for millions who previously required no advance clearance.

Requirements

The program demanded continuous availability, a simple and intuitive application journey, and near real-time decisions for most applicants. Automation was essential to prevent operational bottlenecks. Identity verification needed to be integrated into the journey in a way that supported completion. The rollout followed a phased approach, beginning with Qatar and the Gulf, expanding to around 50 countries including the United States, Canada, and Australia, and the European Union within just over a year.

Result

The Home Office approached ETA as a live service from the outset. Performance expectations for availability, speed, resilience, and response were defined early and supported by continuous monitoring and service management. The system now operates at sustained scale, handling about 75,000 applications daily, with a new application every 1.2 seconds. By the end of 2025, 24.8 million ETAs had been issued.

The Impact

When identity verification is operated as a managed service, it can help critical programs move from policy ambition to global-scale delivery – supporting high application volumes, sustained availability, and trusted digital access for millions of users.

The UK ETA example is useful because it shows why identity verification at scale must be treated as part of service delivery. The objective was not simply to verify identity. It was to support a policy outcome, preserve the travel experience for legitimate visitors, reduce pressure on operations, and maintain confidence in a mandatory digital permission process.

It also shows why the operating model matters. The requirements combined availability, speed, automation, assurance, and operational resilience. Those requirements can't be met by feature capability alone. They require service design, monitoring, incident response, change management, and cross-functional alignment across policy, digital, operations, and suppliers.

The Path to Trusted Services, at Scale

Identity verification has become part of the operating fabric of critical digital services. The strongest programs treat it accordingly. They define the mission, understand the guardrails, identify the failure modes that matter, and choose a delivery model that fits the service risk.

For services delivered at national or global scale, managed identity verification offers a practical answer to a hard problem: how to combine assurance, accessibility, resilience, compliance, fraud adaptation, and operational accountability in one model. The decisive issue is whether the identity layer can be operated dependably over time. Components can be replaced. Integrations can be rebuilt. But once a critical service depends on identity, the cost of instability is paid by users, operations teams, regulators, and the organization's reputation.

When identity holds, critical services can scale with confidence. That is why identity should be evaluated as critical infrastructure, and why the operating model matters as much as the technology.

Make identity the service layer your critical program can trust. Assess your current model against the requirements that matter in production: accountability, resilience, assurance, accessibility, compliance, and scale. Then choose a partner who can demonstrate they can deliver identity assurance, availability, and accountability in production.

Entrust's specialists work alongside your teams to assess current initiatives against the criteria set out in this paper. We can help you create your digital transformation path with identity as the foundation.

Contact us today to begin your identity-centric journey

ABOUT ENTRUST

Entrust fights fraud and cyber threats with identity-centric security that protects people, devices, and data. Our comprehensive solutions help organizations secure every step of the identity lifecycle, from verifying identity at onboarding to securing connections and fighting fraud in everyday transactions. Ongoing monitoring supports compliance and safeguards keys, secrets, and certificates. With a foundation of identity-centric security, our customers can transact and grow with confidence. Entrust has a global partner network and supports customers in over 150 countries.

entrust.com Toll-Free: 888.690.2424 | International: +1.952.933.1223 | sales@entrust.com

ENTRUST
SECURING A WORLD IN MOTION