People Verification vs Photo ID, Video, Notary, and KBA: What Still Holds Up in the Deepfake Era

Buyer's Guide / Last updated 2026-06-11 / https://www.scrambleid.com/learn/people-verification-vs-traditional-methods

> **Status (June 2026):** Early access. The People verification family is live with early-access customers and isn't generally available yet. The comparison below describes the shipped design as it stands today; talk to your ScrambleID account team about access and timelines.

In one sentence: Photo ID, video calls, remote notary apps, knowledge-based questions, and "call them back to verify" are all probabilistic signals that AI can increasingly defeat at scale; person-to-person cryptographic verification produces a deterministic answer in a few seconds (versus the 30 to 90 seconds typical of knowledge-based questions) and is the only method on this list whose security properties are not eroded by AI capability progression.

TL;DR (canonical)

- **Photo ID + visual match:** designed before commodity deepfakes existed; defeated at scale by high-quality forged IDs and face-augmentation models.
- **Video conferencing as identity verification:** broken by real-time deepfake video. The Arup Hong Kong $25.6M loss in early 2024 made the failure mode public.
- **Remote online notarization (RON):** appropriate for notarized documents, too heavy and too slow for routine identity verification (wire approvals, helpdesk).
- **Knowledge-based questions (KBA):** defeated by public records and breached data; NIST SP 800-63A discourages KBA as primary proofing.
- **"Call them back to verify":** valid only if the callback number is known-good and the callback channel is itself authenticated; voice cloning has eroded callback alone.
- **People verification (cryptographic round trip):** deterministic, completes in a few seconds, AI-quality-independent. The signature either verifies or it does not.

How to read this comparison

Every method below has legitimate uses. The question is not "which method is best?" The question is "which methods still work for high-trust decisions in a world where AI can clone any voice from 30

seconds of audio and produce real-time deepfake video on a laptop?"

The dimensions that matter:

1. **Determinism.** Does the method produce a binary yes/no, or a probability score?
2. **AI-resistance.** Does AI capability progression erode the method's effectiveness?
3. **Latency.** How long does a verification take?
4. **Audit quality.** What does the post-event audit trail prove?
5. **Operational reach.** Does the method work in person, over voice, and remotely?
6. **Coercion resistance.** Does the method require explicit, observable action by both parties?

People verification scores well on all six. Most traditional methods score well on at most two or three.

Side-by-side matrix

Property	Photo ID + visual	Video call as ID	Remote notary (RON)	KBA between humans	Callback to known-good	ScrambleID People (cryptographic)
Determinism	Probabilistic	Probabilistic	Probabilistic + procedural	Probabilistic	Depends on channel	Deterministic
AI-resistance	Eroding fast	Broken at scale	Eroding (RON video portion)	Long-defeated by breached data	Eroding (voice clone)	Cryptographic, AI-independent
Median latency	5-30 seconds	Multi-minute	Multi-minute	30-90 seconds	Multi-minute	A few seconds
Audit trail	"Looked OK to me" note	Recording, if retained	Notarized record + video	Question/answer log	Call log	Signed cryptographic event with TTL, device IDs, attribute set
Phone (voice) channel	Not applicable	Not applicable	Not applicable	Common (and weak)	Native	Type Code variant
In-person channel	Native	Awkward	Not applicable	Awkward	Not applicable	QR scan native
Remote / distant	Not applicable	Native	Native	Common	Native	SMS deep link variant
Coercion resistance	Low	Low	Procedural (notary)	Low	Medium	High (both sides explicitly act)
Cost per verification	Low (staff time)	Medium	High (notary fee)	Low	Low (staff time)	Sub-cent (cryptographic)

Property	Photo ID + visual	Video call as ID	Remote notary (RON)	KBA between humans	Callback to known-good	ScrambleID People (cryptographic)
Standards posture	Vendor-specific	Vendor-specific	State law, MISMO	NIST 800-63A discourages	None	NIST 800-63B / 207, FIDO2/WebAuthn

The pattern: traditional methods score well on familiarity and existing process integration; they score poorly on the dimensions that matter most when AI is the adversary (determinism, AI-resistance, audit quality, latency).

Where each method belongs in 2026

Photo ID + visual match

Where it works: Low-stakes identity confirmation. State ID at a bar. Boarding pass at a gate. The human is mostly checking that the ID looks plausible and the photo roughly resembles the person. The verification is socially scaffolded ("you'd have to be brazen to use a fake here").

Where it fails: High-value transactions. Bank branches handling six-figure withdrawals. Corporate facilities with sensitive areas. Any context where a forged ID + a face the verifier has not previously seen is a path to large-loss compromise.

The deepfake-era posture: Photo ID remains useful as a low-assurance signal alongside cryptographic verification for high-trust decisions. It is not sufficient alone.

Video conferencing as identity verification

Where it works: Established relationships. You've worked with this person for years; the video call confirms presence and continuity. The video does not verify identity; it confirms an existing trust relationship is being exercised by the expected parties.

Where it fails: Initial verification. Trust establishment. High-value asks from someone you've never met or rarely interact with. The Arup Hong Kong loss involved a multi-participant video call where every participant except the victim was AI-generated. Real-time deepfake video on commodity hardware is now production capability.

The deepfake-era posture: Video calls are for human collaboration, not identity verification. High-trust asks require an out-of-band cryptographic verification regardless of who the video call appears to show.

Remote online notarization (RON)

Where it works: Documents requiring notarization (real-estate transactions, certain affidavits, regulated business filings). RON has its own legal framework, state-by-state rules, and the multi-step

procedure that includes KBA, document analysis, video, and a notary in the loop produces a notarized record.

Where it fails as routine identity verification: Latency, cost, and procedural overhead make RON impractical for high-volume verification (wire approvals, helpdesk, branch transactions). The KBA component of RON inherits the same weaknesses KBA has elsewhere.

The deepfake-era posture: RON for the documents it was designed for. Cryptographic people verification for routine identity verification of the parties involved.

Knowledge-based questions (KBA) between humans

Where it works: Almost nowhere as primary verification. KBA can serve as a weak supplementary signal or as a fallback for unauthenticated read-only inquiries.

Where it fails: Anywhere a determined attacker can spend 10 minutes on a search engine, dark-web data marketplace, or breached-data aggregator. NIST SP 800-63A discourages KBA as a primary identity-proofing signal.

The deepfake-era posture: KBA remains a fraud-tax indicator (the absence of basic identity facts is a signal that the caller is not the legitimate user). Its presence is not evidence of legitimacy because the data is widely available.

"Call them back to verify"

Where it works: Calling a known-good number from organizational records (not a number provided by the suspicious caller) plus a cryptographic verification on the callback. The callback channel is itself authenticated.

Where it fails: Calling alone, without cryptographic verification. Voice cloning from 30 seconds of audio is commodity capability; the legitimate-sounding voice on the callback is no longer evidence of legitimacy.

The deepfake-era posture: Callback to a known-good number remains a valid procedural step; it must be paired with cryptographic verification to provide assurance for high-value transactions.

People verification (cryptographic round trip)

Where it works: Any high-trust verification between two humans where the goal is a deterministic, audit-rich answer in a few seconds. Wire approvals, vendor banking changes, IT helpdesk credential requests, branch high-value transactions, executive-to-finance asks, contractor verification at physical sites.

Where it does not apply: Anonymous interactions, contexts where neither party has a ScrambleID-bound identity, edge cases where the verification is one-sided (the verifier confirms the presenter's identity without exposing their own).

The deepfake-era posture: People verification is the deterministic anchor for high-trust verification. Other methods (badge, photo ID, video) layer on top as social and operational signals; people verification provides the cryptographic ground truth.

A concrete walkthrough: vendor banking change

The threat: an attacker emails Accounts Payable claiming to be a long-standing vendor and requests a banking-detail update before the next payment cycle.

Method 1, photo ID emailed in. Defeated by a forged ID image attached to the email. AP has no way to verify the ID is real or that the email sender is the legitimate vendor.

Method 2, video call to "verify." The attacker schedules a Zoom call. The video shows a person matching the vendor contact's LinkedIn photo. The voice matches what AP remembers. The participants on the call could be anyone; the video stream could be a deepfake. AP would not know.

Method 3, RON. Overkill and slow. Vendor banking changes happen often; routing each through a notary process is operationally untenable.

Method 4, KBA. AP asks the caller for the vendor's tax ID, principal contact name, and last invoice number. All of these are typically retrievable from breached data, public records, or even the legitimate vendor's prior emails (if a vendor email account has been compromised, all of these are in the inbox).

Method 5, callback alone. AP calls the vendor's known-good number. The legitimate vendor's account has been compromised and the call is forwarded; or the legitimate voice has been cloned; or the legitimate person is on PTO and a delegated AP-receiver picks up.

Method 6, people verification. AP requests a people verification through the legitimate vendor contact's enrolled identity. The verification either completes (the legitimate vendor's hardware-bound private key signs the challenge) or it does not (the attacker has no private key to sign with). The result is deterministic. The audit trail records the device IDs, the attribute set shared, and the timestamp. The whole exchange takes a few seconds, versus the 30 to 90 seconds typical of knowledge-based questions.

The pattern repeats across wire approvals, IT helpdesk requests, executive sign-offs, branch high-value transactions, and contractor verification.

When traditional methods are still right

This article is not a prescription that people verification replaces everything. It is a prescription that People verification is the deterministic anchor for high-trust verification, and that traditional methods remain useful in specific contexts:

- **Photo ID at low-stakes interactions** (bars, gates, basic facility entry) remains operationally appropriate.
- **Video calls for established relationships** continue to be the right communication channel; just stop treating them as identity verification.
- **Remote online notarization for the documents it was designed for** is the right tool.
- **Callback to known-good numbers** remains a valid procedural step, paired with cryptographic verification for high-value asks.
- **KBA as a low-assurance fallback** for unauthenticated read-only inquiries can stay; just stop treating it as primary verification for high-trust decisions.

The change is at the high-trust end of the spectrum, where probabilistic methods now produce confident-looking false positives that AI can engineer at scale.

Decision criteria for buyers

Use this short list to score traditional and emerging verification methods against the threats your enterprise actually faces:

1. **What's the loss if a single verification produces a false positive?** Six-figure wire? Vendor banking change? Helpdesk credential reset? The higher the value, the more important deterministic verification becomes.
2. **Is the adversary capable of voice cloning, deepfake video, or AI-driven social engineering?** If your threat model includes any modern fraud actor (and it does), the answer is yes.
3. **What's your latency budget?** Cryptographic verification in a few seconds is achievable. Anything multi-minute is going to be operationally avoided.
4. **What audit trail does compliance require?** Signed cryptographic events with device IDs and TTL data hold up to forensics; "the caller answered the security question correctly" does not.
5. **What channels do your verifications happen across?** In person, over voice, distant. Cryptographic methods that adapt to all three (QR/QID, Type Code, SMS deep link) avoid per-channel point solutions.
6. **What does the recovery path look like if the primary verification fails?** A weak recovery path becomes the new attack surface (see Recovery and Fallback Playbook).

Key Takeaway

Traditional human-to-human verification methods (photo ID + visual match, video calls, remote online notarization, knowledge-based questions, callback to known-good numbers) were designed in a pre-AI era and are increasingly probabilistic and defeatable as AI-generated voice cloning, real-time deepfake video, and breached-data aggregation become commodity. Person-to-person cryptographic verification produces a deterministic answer in a few seconds (versus the 30 to 90

seconds typical of knowledge-based questions), is independent of AI capability progression because no AI can produce a signature without the matching hardware-bound private key, and produces a signed audit trail that holds up to forensics. People verification does not replace photo ID at low-stakes interactions, video for human collaboration, or RON for documents that require notarization; it replaces probabilistic verification at the high-trust end of the spectrum, where loss is large and probabilistic false-positives are now engineered by attackers.

FAQ

Can a deepfake defeat photo-ID verification?

Yes, in the directions that matter. A high-quality forged ID is hard to detect by visual inspection, and the visual face match between the ID photo and the in-person face has always been probabilistic. AI-generated face augmentation can defeat both human and many automated face-matching systems. Photo-ID verification was designed in an era before commodity deepfake generation; its assumptions no longer hold for high-value transactions.

Is calling someone back to verify still effective?

It depends on which number you call. Calling a number provided by the suspicious caller defeats the purpose. Calling a known-good number from your records is better but still relies on the assumption that voice authentication is meaningful. Voice cloning from 30 seconds of audio is now commodity capability. Callback to a known-good number plus a cryptographic verification (such as a cryptographic round trip) is the current standard; callback alone is not sufficient for high-value transactions.

How do remote notary apps compare to people verification?

Remote online notarization (RON) is a regulated identity-proofing process for legal documents, with state-by-state rules. It typically combines KBA, document upload, video conferencing, and a notary-in-the-loop. RON is appropriate for documents requiring notarization. It is too heavy and too slow for routine identity verification (wire approvals, vendor changes, helpdesk verification) where the goal is a cryptographic round trip that takes a few seconds rather than a notarized record.

What's wrong with knowledge-based questions (KBA) between humans?

Public records and the cumulative leakage from a decade of breaches have made knowledge-based questions trivially defeatable. SSN, date of birth, prior addresses, last four of a card, mother's maiden name, recent transactions: all available to motivated attackers. NIST SP 800-63A discourages KBA as a primary identity-proofing signal. KBA between humans (over phone) is the same problem with worse audit.

Is people verification slower than just looking at someone's badge?

No. End-to-end verification takes a few seconds after the verifier triggers consume, and most of that is the network round trip and the mobile UI render rather than the cryptography. A badge inspection takes about as long and produces a probabilistic answer; people verification produces a deterministic one.

Can I use people verification alongside other verification methods?

Yes, and the deployment model is designed for layering. People verification is layered: photo ID and badge inspection still happen at physical sites; KBA can still serve as a low-assurance fallback for unauthenticated read-only inquiries; video calls still exist for relationship work. The change is that high-trust decisions (wire approvals, vendor banking changes, IT helpdesk credential requests, in-person high-value transactions) require a people verification cryptographic verification rather than relying on the older signals alone.

References (public)

- FBI IC3 Annual Internet Crime Report: https://www.ic3.gov/AnnualReport/Reports
- FinCEN Alert on Deepfake Media for Identity Fraud: https://www.fincen.gov/sites/default/files/2024-11/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf
- CISA, Implementing Phishing-Resistant MFA: https://www.cisa.gov/sites/default/files/publications/fact-sheet-implementing-phishing-resistant-mfa-508c.pdf
- NIST SP 800-63A: https://csrc.nist.gov/pubs/sp/800/63a/4/final
- NIST SP 800-63B: https://csrc.nist.gov/pubs/sp/800/63b/4/final

---

Related reading

- What Is People Verification?
- Deepfake-Resistant Identity Verification
- Stopping Help-Desk Impersonation with People Verification
- People Trust Checks
- People Verification Implementation Guide
- What Is Identity Proofing?