ONE CONSORTIUM

Digital Identity Working Group

Digital Identity for Business Communications:

Demonstrating Legitimacy, Reducing Fraud, and Enabling Verifiable Communications

A Briefing for GIRAF

> ***Note on terminology:*** *Terms such as authenticated and verifiable carry specific technical and legal meanings that differ from everyday usage. The working group will define these terms precisely in subsequent documents. As used in this briefing, "verifiable" means that any receiving party can independently check the evidence of identity for themselves, without simply trusting that someone else has already done so.*

Executive Summary

Consumer trust in voice and business messaging is in structural decline. GASA estimates that USD442bn was lost to scams worldwide in 2025.[1][2] Similarly, Europol reports that spoofing, specifically caller ID spoofing, drives financial fraud and enables social engineering scams, resulting in substantial economic and societal damage, with an estimated EUR 850 million lost worldwide annually. Phone calls and texts are primary attack vectors, accounting for approximately 64% of reported cases; in this measure, Caller ID spoofing is the primary enabler of consumer-facing voice and SMS fraud. This is done by deliberately manipulating the information displayed on a user's caller ID, often using Voice over Internet Protocol (VoIP) services or [3] apps, to show a false name or number that appears legitimate and trustworthy.[4]

This briefing introduces the concept of digital identity for business communications: a framework in which a business entity can be properly identified, vetted, and issued with Portable credentials that can accompany—and be verified within—every communication it sends. Legitimate organisations benefit because their identity is now provable. Fraudsters lose their ability to impersonate them. Regulators and consumers gain tools to distinguish one from the other.

Using open standards, inclusive approach, with focus on interoperability across ecosystems and without reliance on any single central database or a centralized authority; our proposal puts the well-intentioned communications provider at the centre: driving a pro-competition environment and weeding out scam and fraud on voice and text communications.

The proposal

[1] https://www.feedzai.com/global-state-of-scams-report-2025/

[2] "Global State of Scams 2025" (GASA, 2025)

[4] "Position Paper on Caller ID Spoofing" (Europol, 2025)

This paper seeks to lay the foundations for an industry and regulator discussion of a pro competition, communications operator centric, non-proprietary system that does *not* require a single central database or central authority.

The express ambition is that the proposal can work with existing call and message authentication e.g. STIR/SHAKEN, RCD, Verified Origin adopted in Brazil, etc.

Therefore, we lay down an open and pro-competition framework.

The proposal centers around the following requirements:

- **Open and interoperable.** Any framework that depends on a single vendor, a proprietary protocol, or a closed registry will face adoption barriers and create new dependencies. The approach must be based on open, inspectable standards that can be independently implemented by any operator or service provider.
- **Inclusive and proportionate.** Standards must not inadvertently exclude smaller organisations, or organisations in jurisdictions where formal identity infrastructure is less developed. The framework should support a range of vetting approaches appropriate to different contexts and should focus on incentives for adoption rather than punitive exclusion.
- **Privacy-preserving.** Only the minimum necessary information should be revealed to any given verifier. An organisation's full identity record does not need to be disclosed to every party; what matters is that the relevant claims—'this is a legitimate business with the right to use this number'—can be confirmed.
- **Actionable without mandating wholesale change.** The framework should deliver value to early adopters without requiring universal adoption before any benefit is realised. It should be possible to add one organisation at a time, with verifiers immediately able to distinguish between communications with proven identity and those without. The framework is also designed to be proportionate: the burden of participation should scale with the level of participation, and enhanced services above the baseline are a matter of commercial choice for operators.

The delivery of the quadruple requirements is proposed to introduce a four-step digital identity framework for business-to-consumer communications. **First**, a business entity is vetted through a Know Your Business (KYB) process that confirms its legal existence and right to use the numbers, brand assets, and channels it deploys. **Second**, that verified identity is cryptographically linked to each individual communication — not stored in a registry to be queried but bound to the specific call or message. **Third**, any party in the delivery chain — terminating carrier, recipient device, analytics system, or regulator — can verify the identity evidence independently in real time or retrospectively. **Fourth**, a trust signal is presented to the end user, giving the consumer a visible, reliable indication that the communication has been verified.

The initial scope covers traditional voice and standard business messaging (SMS, RCS), with extension to OTT platforms and fixed-line endpoints anticipated. Whilst B2C represents the initial focus of this discussion, it does so merely for reasons of tackling large scale fraud (which tends to happen in a B2C calling set up) and in order to make the initial discussion more manageable. Personal subscriber identity verification is foreseen to be discussed in the next phase. Personal subscriber identity is explicitly out of scope in this phase.

Mandatory action, open standards, and governance

The framework is expressly not proposed as a mandate. It is designed to deliver value to early adopters without requiring universal adoption, and industry is explicit that it is not seeking regulatory compulsion. However, regulatory backing is treated as essential for the consumer-facing trust signal: a mark that any operator can display without consequence is one fraudsters can equally display. GIRAF members are asked to define outcome standards for verified communications, support conditions under which trust indicators may be displayed, and impose consequences for misuse — with the document acknowledging that NRA powers alone may be insufficient in some jurisdictions and that statutory backing may be required.

The framework is built entirely on open, interoperable standards; No single vendor, proprietary protocol, or central registry is mandated, required or desired. Multiple vetting bodies and credential sources can coexist, provided they meet applicable governance standards. Verification is decentralised: no single central authority needs to be queried. Any authorised party — including the communication recipient — can check the cryptographically linked identity evidence independently, without trusting that an intermediary has already done so. Each jurisdiction retains full authority over its own numbering policy and enforcement; the framework provides a common mechanism for expressing and transporting that evidence across borders, not a harmonisation of the underlying national rules.

Trust should not be siloed. Evidence of identity must be able to travel with every communication—call, message, or API—enabling verification across networks, borders, and channels.

This is not a proposal to replace existing systems, but to complement and uplift them in an operator centric and pro competition manner. Industry will develop the frameworks and mechanisms; the regulatory community has an essential role in providing the governance backing that makes those mechanisms trustworthy, and provably compliant.

This briefing is a call to engage—to define together what the next layer of the communications trust model should look like, openly, interoperably, and in a way that works across jurisdictions and channels. The specific actions this working group believes GIRAF should take are set out in Section 7.

1. The Problem: Trust in Business Communications Is Collapsing

1.1 Scope: business identity, not personal identity

The initial focus of this briefing is on business-to-consumer communications: calls and messages sent by organisations—companies, public bodies, service providers—to the individuals they serve. The identity question here is not about individuals, but about business entities.

We initiate the discussion about new market-place enhanced solutions tailored for large business communications (B2C) solely to give initial focus and by recognizing that the majority of scam, fraud and robocalling originates from impersonated ‘businesses’ to the detriment of Consumers. We do however not otherwise differentiate between businesses and consumers and C2C as well as C2B verification should follow naturally.

The question this framewok addresses is: is this communication genuinely from the organisation it claims to be from, and does that organisation have the right to communicate using this number, this brand, and this channel?

1.2 Consumers no longer trust what they receive

Numerous independent reviews confirms the death of the cold call and the decline of voice trust.

- **Unanswered calls**: "80% of calls from unrecognized numbers now go completely unanswered."[5]
- **The "Lifeline" Collapse:** In the FCC's public comment dockets, small business owners and medical professionals routinely submit evidence to the regulator stating that the voice network has been rendered "nearly useless." For example, the FCC's record includes filings from healthcare practices noting that **4 out of 5 outbound calls to patients go straight to voicemail** because consumers treat unknown local numbers as automated security risks by default.[6]
- **End-user suspiciousness:** "Over half of consumers (**57%**) report that suspicious, spoofed calls are actively becoming more frequent."[7]

The same evidence is recounted by Regulators; two examples:

In the USA, the FCC found that: The erosion of trust is so severe that the FCC's primary, top-listed consumer advice for avoiding phone fraud is a literal mandate of the statistic: *"Don't answer calls from numbers you don't know. Let unknown calls go to voicemail."* The FCC noted that the public's total refusal to answer the phone is what forced regulators to mandate STIR/SHAKEN protocols.

In the UK, Ofcom study shows that phone users heavily rely on network-level spam alerts or simply refuse the call because over half of all UK mobile users report receiving a suspected scam call or text on a weekly basis.

The consequences are severe for all parties. Citizens miss important calls from banks, healthcare providers, government agencies, and businesses they deal with every day. Businesses find their outreach campaigns failing not because their message is wrong, but because it cannot be distinguished from fraud. Public institutions discover that their communications are being impersonated, undermining public confidence in official channels.

1.3 Fraud and impersonation are escalating

Scams initiated through calls and messages represent significant and growing harm to consumers globally. Advances in generative artificial intelligence now allow fraudsters to clone voices and manufacture compelling impersonations at scale, targeting individuals with highly personalised attacks. This is illustrated by the following graph, which shows the value lost to scams instigated

[5] "Most Americans don't answer cellphone calls from unknown numbers" (Pew research, 2020)
[6] CG Docket No. 17-59 (Advanced Methods to Target and Eliminate Unlawful Robocalls) and WC Docket No. 17-97 (Call Authentication Trust Anchor)
[7] "[...]Brits lose confidence due to scam calls" (Nationwide, 2026)

by call or text message (in millions of US dollars) per complaints to the US Federal Trade Commission (FTC).

Figure: Value of Scam Losses Initiated by Call and Text Message per Complaints to the FTC

For comparison, the following graph shows the value of scam losses (in millions of Australian dollars) initiated by call and text messages per complaints to the Australian Communications and Media Authority (ACMA).

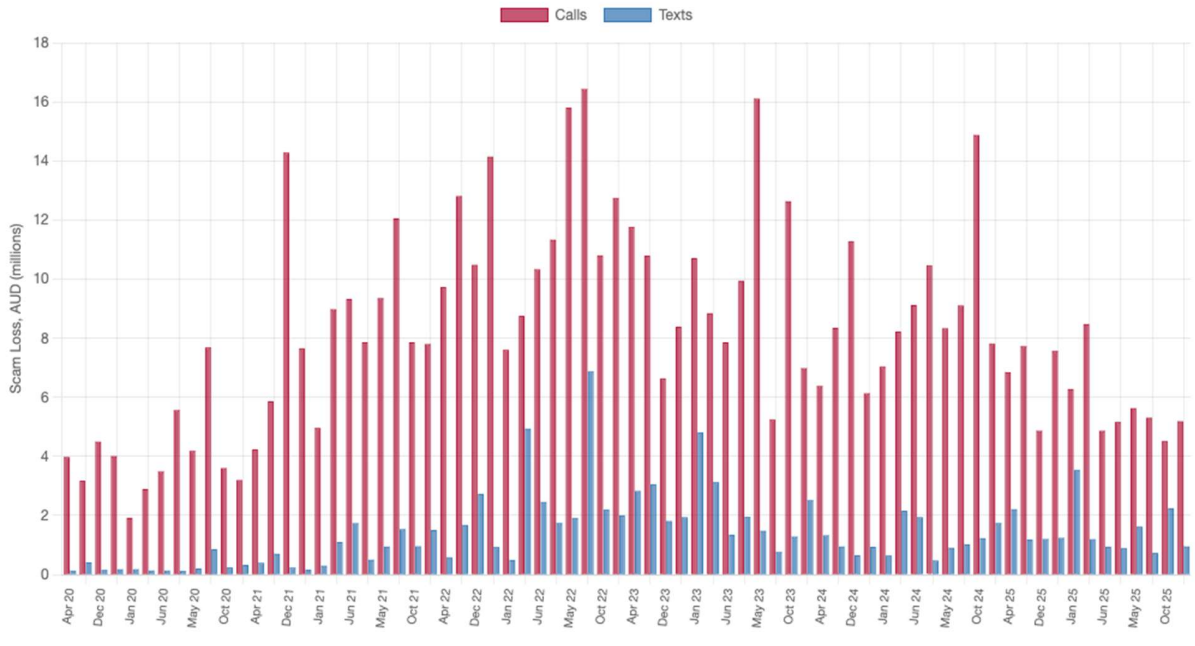

Figure: Value of Scam Losses Initiated by Call and Text Message per Complaints to the ACMA

Europol, in a 2025 position paper on caller-ID spoofing drawing on engagement with law enforcement agencies across 23 EU member states, found as follows:

Spoofing, specifically caller ID spoofing, drives financial fraud and enables social engineering scams, resulting in substantial economic and societal damage, with an estimated EUR 850 million lost worldwide annually. Phone calls and texts are primary attack vectors, accounting for

approximately 64% of reported cases. This is done by deliberately manipulating the information displayed on a user's caller ID, often using Voice over Internet Protocol (VoIP) services or specialised apps, to show a false name or number that appears legitimate and trustworthy. The ability of malicious actors to conceal their true identity and origin, severely impedes the capacity of law enforcement agencies (LEAs) to trace and prosecute cybercriminals. (underlining added)

In the UK, Ofcom's tracking shows that over **70% of UK landline and mobile users** report receiving a suspected scam call or text message every single week. Because of this saturation, voice fraud accounts for hundreds of millions of pounds in annual losses, particularly targeting vulnerable or elderly demographics who still rely heavily on landline communications. In the USA, the FTC received roughly 3 million fraud complaints, noting a staggering **$15.9 billion** in total consumer losses.

Whilst all losses experienced on the end-user / consumer side are deeply regrettable, it is particularly concerning and amoral that scammers and fraudsters specifically target the more harmful and vulnerable in society, including the elderly.

The harm extends beyond direct fraud losses. When consumers disengage from communications channels they no longer trust—declining calls, ignoring messages, abandoning services—new features and capabilities that would otherwise benefit them are either delayed or never reach them at all. Services that depend on voice or message-based engagement cannot achieve the utilisation levels needed to make investment sustainable. The damage is therefore not only to individual consumers who fall victim to fraud, but to the broader value that trustworthy communications can deliver to society—value that remains unrealised for as long as distrust persists.

Critically, for countries or companies that have not deployed STIR/SHAKEN, the methods used to commit this fraud exploit a basic structural weakness: there is no harmonised reliable, real-time way for a recipient—or the network delivering a call or message—to verify the accuracy of the presented telephone number and that the claimed sender is who they say they are. This weakness is not a technology failure. It is a governance and architectural gap that existing systems were not designed to close.

1.4 Existing defences are necessary but not sufficient

The communications industry and regulators have deployed significant resources to address this problem. Call authentication frameworks, sender registration systems, national CLI blocking rules, blacklists, analytics engines, and machine-learning classifiers all play an important role and have meaningfully reduced some categories of harm. These are valuable tools; the framework proposed in this briefing is designed to complement them, not replace them.

The most common include:

STIR/SHAKEN

Whilst STIR/SHAKEN enabled call authentication is robust, based on open industry standards and allows for provider centric implementation and operation, it is currently limited to three countries, though development to 'extend' reach via Cross Border Call Authentication (CBCA) is in progress. We do not know whether global NRAs, legislators or industry players will, and if so when, adopt STIR/SHAKEN.

Rich Call Data (RCD)

RCD (Rich Call Data) allows caller ID to carry verifiable (veracity of brand information and a brand information — the organisation's name, logo, and the stated purpose of the call — displayed on the recipient's screen before they answer. This makes it harder for fraudsters to spoof brand names and logos, and more informative for consumers deciding whether to respond.

RCD (Rich Call Data) allows caller ID to carry verified brand information. It is verifiable both in the sense of the veracity of the organisation's name, logo, and the stated purpose of the call, as well as cryptographically verifiable proving the integrity and authenticity of the verified information. It is then displayed on the recipient's screen before they answer.

STIR/SHAKEN currently provides that proof layer in the United States and several other jurisdictions. Outside those jurisdictions, RCD-equivalent branded calling can in principle be supported by any alternative cryptographic substrate that fulfils the same function — binding verified organisational identity to the communication with integrity. The framework proposed in this briefing is designed to serve exactly that role, whether as a complement to existing STIR/SHAKEN deployments or as a foundation where STIR/SHAKEN is not present.

BCID (Branded Calling ID)

BCID is a US industry-managed ecosystem built *on top of,* and *downstream to*, RCD to vet enterprises and securely deliver these branded details across major carrier networks.

Proprietary AI Solutions to enhance caller ID certainty

With numerous providers, regulators, industry forums and commentators seeking swift remedy against the global plague of communications scam, fraud and robocalling, we see proprietary and non-interoperable solutions being rolled out commercially, at an increasingly rapid pace.

Whilst there are certainly many positive uses of AI powered solutions, including detecting and stopping scam and fraud in the first place, as well as defence against Artificial Inflation of Traffic (AIT), preventing misuse of one-time-passwords and the deployment of hyper-realistic, conversational AI bots to intercept scammers in long and timewasting conversations, these do not seek to certify the caller's identity.

Where these AI solutions are deployed to prevent inbound fraud and scam calling, they often use heuristic and arithmetic based predictions of the caller and not actual verified identity. These solutions appear to provide some general protection for end-users, but they do not offer platform interoperability precisely because they are proprietary and non-interoperable. These inbound systems typically also require the establishment of a central database held by the proprietary owner, not the communications provider, which raises its own concerns about competition.

In short, whilst these proprietary offers *may* provide some protection for the end user against fraud, scam and robocalling, they re-introduce the shortcomings of a market based on an *essential facility* for terminating calls. Specifically, if an originating operator chooses not to participate in the AI system selected by the terminating operator, or has chosen another vendor for its own uses, their calls destined for termination may never pass over to terminate on the proprietary filtering provider's network.

The proposal of this paper is not to suggest a rejection the use of AI to combat fraud and scam, as said, there are many valuable uses. Our proposal should rather be seen as a valuable complement to AI solutions for termination networks. Rather than merely filtering calls based on proprietary AI assessments, our proposal – using open and transparent standards – can provide direct and verified identity assurance of the calling party which can be utilised by AI solution at the terminating end to enhance their anti-fraud and scam outcomes. As the AI solutions in this field are all proprietary and offer no interoperability, the reverse is not available.

Other

Additionally, it is worth noting that a small number of real-world systems already provide verified sender identity rather than probabilistic inference. In business messaging, RCS for Business (RBM) allows Google or the mobile network operator to verify and display the sender's identity in a form that cannot be spoofed. In voice, services such as BCID in the United States, and comparable schemes in Brazil and Egypt, incorporate Know Your Business (KYB) vetting and right-to-use (RTU) verification into branded calling.

These examples demonstrate that verified business identity in communications is technically achievable. However, they share a common limitation: each is siloed to a specific service or jurisdiction, requiring organisations to undergo separate vetting for every channel and country in which they operate. The framework proposed in this briefing addresses precisely this gap—making verified identity portable, once established, across channels and borders.

As a global industry forum with a sole focus of re-introducing trust in the communications market, One Consortium echoes the views of the European Commission that summarised the challenges seen from a global and cross border perspective:

End-users are increasingly vulnerable to fraudulent activities carried out through interpersonal communications services. Several Member States have introduced national measures to combat evolving fraud schemes, such as CLI spoofing, phishing, smishing, and vishing. However, those measures remain uncoordinated and reflect primarily national perspectives, enabling fraudsters to shift their activities to jurisdictions where such practices are not effectively addressed.

But the above means to combat scam calling and spoofing share common limitations: AI solutions work on probabilities, not proof. They identify likely bad actors based on patterns, history, and inference. They cannot positively identify good actors based on verifiable evidence. Others, even where a call authentication framework confirms that a calling line identifier has not been modified in transit, this does not confirm who is legitimately behind that number, nor whether the organisation using it has the right to do so. Knowing the CLI was not spoofed tells us nothing about whether the organisation presenting it is who it claims to be. As a result:

- Legitimate calls and messages are frequently blocked or labelled as suspicious, imposing real costs on businesses and citizens who rely on those channels.
- Fraud that is novel, low-volume, or well-disguised continues to reach its targets.
- The same number can be authenticated by one system and flagged by another, creating inconsistency and eroding confidence.
- Enforcement and traceback—locating the true origin of fraudulent traffic—remains inconsistent and globally incomplete, precisely because identity is not built into the communication itself.

The industry needs not only additional defences against bad actors, but a positive framework for demonstrating the legitimacy of good ones. It is useful to distinguish three distinct ways in which digital identity supports anti-fraud goals: punitive (strengthening traceback and enforcement against bad actors); protective (tagging verifiably well-identified traffic so networks and devices can treat it differently); and promotional (enabling branded communications that verified businesses can use to build engagement and trust with customers). Existing authentication and registration tools primarily support the first. The framework proposed in this briefing additionally enables the second and third.

2. The Opportunity: Proving Who Is Communicating

2.1 The positive case: demonstrating legitimacy

We currently spend a lot of time trying to detect who the bad actors are. Put differently: restore confidence in the good actors, and the bad actors are excluded by default. The most powerful change we can make would enable legitimate actors to prove their identity through every communication they send. We will then be able to progressively exclude the bad actors and restore confidence when members of the public answer their phone.

> *Legitimate businesses do not need to change their behaviour. They simply need a way to show, credibly, who they are. Fraudsters who impersonate them will not be able to do so.*

When a consumer receives a call from their bank, they should be able to know—not guess—that it genuinely originates from that bank. When a government agency sends a message, citizens should be able to verify it. This is what digital identity for communications makes possible: a world in which the question “who is really communicating with me?” has a provable answer.

This is fundamentally a positive innovation. It protects citizens. It builds trust in channels that businesses and public services depend on. And it creates accountability: where identity is required, bad actors lose the cloak of anonymity. It is equally important to be clear about what this framework does not claim: it does not guarantee that consumers will answer every verified call. The goal is to enable an informed choice — by equipping every party in the chain, from the originating network through to the end user, to make better-informed decisions about the communications they handle or receive.

A useful analogy is the transformation of the internet through HTTPS and the browser padlock. Before verified website identity became standard, the web was fundamentally untrustworthy for commerce: users had no reliable way to know whether the site they were visiting was genuine. The introduction of verified certificates—displayed as a padlock, and for high-assurance cases a green address bar—was the pivotal change that made e-commerce possible. It did not make the internet perfectly safe, but it provided a reliable, visible signal of verified identity that consumers could act on. Digital identity for communications proposes the same shift for voice and messaging: a reliable, standardised signal that a communication originates from a verified business entity—the equivalent of the padlock for calls and messages.

2.2 What digital identity for communications provides

A digital identity framework for communications addresses four connected needs:

- **Identity.** Identifying the business entity: reliably establishing, through a Know Your Business (KYB) vetting process, that an organisation is who it claims to be, that it exists as a legitimate legal entity, and that it has the right to use the telephone numbers, short codes, and other brand assets it is deploying.
- **Linking.** Connecting verified identity to each individual communication through cryptographic binding — so that the identity established through vetting is inseparably associated with the specific call or message being sent, not merely stored in a registry to be queried. Unlike a registry lookup against a CLI, which returns information about whoever is registered to that number, a linked credential proves who is actually behind this specific communication. Note that the term “Linking” as used throughout this framework refers specifically to cryptographically binding identity evidence to a unique, individual communication. This is entirely distinct from “linkability” in the data protection sense. The GDPR concept of linkability concerns the use of identifiers — including phone numbers — to correlate an individual’s activities across separate transactions or channels in order to construct a profile, which raises right-to-privacy concerns under Article 4(1). Cryptographic binding of evidence to a specific communication creates no such cross-transaction correlation and raises no equivalent privacy concern. The two concepts should not be conflated

- **Verification.** Enabling verification: making it possible for any party in the communication chain—the terminating network, the handset, the analytics system, the regulator—to check the identity evidence independently, without simply relying on the word of an intermediary. Crucially, a verifier in one country can examine credentials issued under another country’s governance, seeing exactly what was verified, by whom, and to what standard—without having to simply trust that the foreign vetting process was adequate.
- **Presentation.** Signalling trust to the end user: ensuring that verified identity is communicated to the recipient in a clear, reliable, and tamper-resistant way—whether through a brand display on a smartphone, a CLI prefix on a feature phone, or an audible indicator. The presentation layer is the moment at which all the technical and governance work becomes visible and meaningful to the citizen receiving the call or message.

These four elements together constitute what might be called an unbroken chain of custody: from the governance process that vets the organisation, through the credentials that are issued, to the communication that is sent and verified, to the signal received by the end user. A chain that is unbroken at every point is what separates this approach from the patchwork of partial solutions that exist today.

3. Why Existing Approaches Fall Short

3.1 Registries: useful but siloed

Many jurisdictions have established registries—databases that record which organisations are permitted to send particular types of communications. These registries serve an important function. But they face structural limitations that prevent them from solving the problem at scale:

- They are channel-specific: a registry that covers business messaging may not cover voice calls, and vice versa. Fraud simply migrates to the unregistered channel. Even

within a single channel such as business messaging, separate registries may exist for different identifier types—short codes, standard telephone numbers, and alphanumeric Sender IDs may each be governed by distinct registration systems, creating further silos within the same channel.

- They are jurisdiction-specific: a registry in one country provides no verification for a call crossing an international boundary.
- They are static: they record what was true at the time of registration, not what is true now. An organisation that has changed, been acquired, or whose licence has lapsed may remain in a registry as though nothing has changed.
- They rely on the registry operator as the point of trust: any party wishing to verify must query a central database. If that database is unavailable, compromised, or simply not accessible across borders, verification fails. Furthermore, registry lookups typically depend on the accuracy of the calling identifier itself—a lookup against a spoofed number returns information about whoever is legitimately assigned that number, not whoever is actually originating the call.
- These limitations apply whether the registry is operated under national regulatory mandate or by a commercial provider — but the nature of the problem differs. National registries tend to be siloed by jurisdiction and slow to update, while commercial registries vary widely in their coverage, standards, and interoperability. Neither model, as currently deployed, provides the portable, cryptographically linked identity evidence that this framework proposes.

3.2 Call authentication: necessary but incomplete

A range of approaches exist to validate the call setup path—including in-band authentication frameworks such as STIR/SHAKEN (widely deployed in North America), national CLI blocking rules (the primary approach used across Europe, Asia, and elsewhere, blocking international calls presenting a domestic CLI unless their legitimacy can be verified), and out-of-band parallel database matching. Each represents a genuine contribution to reducing CLI manipulation. But authentication of the call setup path is not the same as authentication of the business entity originating the call.

National CLI blocking rules illustrate the limitation particularly well. The approach works by blocking incoming international calls that present a domestic CLI unless their legitimacy can be verified — typically through direct inter-operator validation or a proxy model. This is a meaningful defence against simple spoofing. But it proves only that the CLI is nationally consistent; it does not prove who is behind it. And it creates a genuine problem for legitimate multinational enterprises that operate contact centres across borders and wish to present local numbers to customers in each market — so that those customers can reach the company at a familiar local number without incurring international charges. Today, such an enterprise has no reliable way to prove to a terminating network in a foreign market that it has the right to use those local numbers. The result is that legitimate cloud-based calling services face the same blocking rules as fraudsters, while fraudsters continue to find ways around them.

A call can be technically authenticated—meaning the number was not spoofed in transit—while still originating from an organisation that has no legitimate right to use that number, or while impersonating a brand. The identity of the communicating entity and their rights to use digital assets like phone numbers and logos, as distinct from the integrity of the call setup, remains unproven.

It is important to be clear about what the framework proposed in this briefing is not: it is not a proposal to replace any of these existing approaches. They represent genuine and substantial

investment by carriers and regulators, and the working group recognises their value. What this framework proposes is to offer the opportunity to add a digital identity layer which enhances existing call authentication infrastructure — adding the evidence, directly linked to each unique communication, that makes authentication meaningful. Either on top of, and complementing existing STIR/SHAKEN based RCD or offer it for countries and companies that do not yet have STIR/SHAKEN.

For the multinational enterprise example, this means that a verified credential establishing the right to use a local number in a given market can travel with the call itself, giving the terminating network the evidence it needs to deliver it confidently.

3.3 The missing piece: portable, linked identity evidence

What is missing from today's architecture is a way for identity evidence to travel with a digital communication—to be genuinely portable across channels and borders, and to be cryptographically linked to the specific call or message in a way that any verifier can check independently and immediately.

Without this, it is impossible to answer the fundamental question with certainty: not "was this call's identifier unmodified?" but "is this genuinely from the organisation it claims to be from, and does that organisation have the right to make this communication?"

4. The Framework: How Digital Identity Works in Communications

The following describes the high-level architecture of a digital identity framework for communications. This section is deliberately non-prescriptive about specific technologies; the goal is to describe the functional requirements that any viable approach must meet.

4.1 Step one: establishing identity through vetting

Before an organisation can demonstrate its identity in a communication, its identity must be established. This means going through a vetting process in which:

- The legal existence of the business entity is confirmed against authoritative sources.
- The organisation's right to use (RTU) the telephone numbers, short codes, and alphanumeric Sender IDs it deploys is verified with the carrier or range-holder that has assigned those identifiers.
- Any additional attributes relevant to the communication context—brand assets, licence status, accreditation—are documented and verified by an appropriate authority.

The result of this vetting process is not simply a record in a database, but a set of portable, re-usable, and verifiable credentials: digitally signed documents that capture what was verified, by whom, and when. These credentials can be updated or revoked if circumstances change. Importantly, the framework does not require organisations to be vetted from scratch: existing KYB data, business registrations, and carrier relationships can serve as inputs into the credential issuance process, reducing the burden of adoption.

The vetting process must be governed by clear rules for who may vet, what standards they must meet, and how disputes are resolved. The governance framework is as important as the technical mechanism. We propose that the framework is open to multiple identity verification mechanisms and credential sources—no single registry, authority, or provider is mandated as the sole root of trust. We have mentioned the HTTPS framework above and another framework worthy of mentioning is the long, varied and industry diverse list of certifying authorities for, the reasonably

similar, European electronic signatures (eIDAS) framework. The eIDAS certification framework is built up as a public / private partnership where the National Supervisory Bodies (public) oversees the National Accreditation Bodies (public) which legally accredits the Conformity Assessment Bodies (private); already today, NRAs and some communications providers are acting as CABs.

Similarly, there is no requirement for a single digital identity scheme: the framework is designed to work with multiple coexisting approaches to vetting and credential issuance, provided they meet the governance standards applicable in the relevant jurisdiction.

4.2 Step two: linking identity to each communication

A credential, or cryptographic token, on its own is not sufficient. It must be connected—linked—to the specific call or message being sent. This is the step that prevents a fraudster from obtaining or copying credentials legitimately belonging to another organisation and using them deceptively.

The mechanisms by which this linking is achieved may vary: it may be incorporated into the communication signalling itself, or it may be accessible through a reference that accompanies the communication. What matters is that the link is specific to that communication, and that it can be verified in real time.

This “linking” step—connecting the identity evidence to the individual communication—is what transforms a registry into a live, communication-level proof of identity. It is the critical architectural piece that most current systems lack.

There are broadly three approaches to achieving this link: 1) in-band (where the identity evidence is carried within the communication signalling itself), 2) out-of-band (where the evidence is carried via a separate channel and referenced by the communication), and 3) hybrid combinations of both. The framework must accommodate all three, since different network environments and deployment contexts will favour different approaches. By accommodating all three approaches, we are enabling interoperability with other valid methods such as STIR/SHAKEN and RCD, and hence allow countries and companies to choose their own path to combat fraud and scam due to without risking sunk costs, regulatory uncertainty or the path leading towards rejecting interoperability and the well-functioning of the competitive communications market.

The linking architecture should support a decentralised and federated model—avoiding a single point of control or failure—consistent with the openness and interoperability principles set out in Section 5.

4.3 Step three: Verification by any party

The final step is verification: the ability of any party in the communication chain to check the identity evidence independently. This includes:

- The terminating carrier, which can confirm the communication is from a vetted and authorised sender before delivering it.
- The recipient’s device, which can present verified identity information to the user.
- Analytics systems and fraud detection tools, which can incorporate verified identity as a high-quality signal.
- Regulators and enforcement agencies, which can use the evidence for traceback, compliance assessment, and enforcement.

Critically, the verification that needs to be done should not require querying a single central database operated by one provider. The evidence should be independently checkable by any authorised party—including the recipient of the communication—at any point in the delivery chain, and at any time, including retrospectively for audit or enforcement purposes.

4.4 Step four: Call presentation

Verification alone does not restore consumer trust—that requires a visible, reliable signal that the communication has been verified. The presentation layer is where the technical and governance infrastructure becomes tangible to the citizen receiving the call or message.

For smartphones, this may take the form of a branded display—showing the organisation's name, logo, and verified status—that the recipient can see before answering. For feature phones and fixed-line endpoints that cannot render rich displays, alternative mechanisms are needed. One example already in use in certain Asian markets is prepending a "#" character to the displayed CLI to indicate that the calling line identity has been verified as non-spoofed. This approach requires no handset upgrade and is visible on even the most basic phone display. Its effectiveness, however, depends entirely on regulatory enforcement: if any originating party can prepend the symbol without consequence, the signal becomes meaningless.

The critical insight is that no presentation mechanism is effective without governance backing. A trust indicator that any operator can display without regulatory oversight is a trust indicator that fraudsters can spoof. The presentation layer is therefore inseparable from the governance framework: its value depends entirely on the rules that govern who may display it, and the consequences for misuse.

Accessibility is a core requirement of the presentation layer. The framework must ensure that trust signals are expressible in forms accessible to all users — including those who cannot see a screen, those using screen readers or assistive technologies, and those using feature phones or fixed-line endpoints. The specific mechanisms will vary by device type and deployment context; the principle is that no user should be excluded from the trust signal by virtue of the device or assistive technology they use.

4.5 Scope: traditional telephony first, emerging platforms to follow

This working group's initial focus is on traditional voice communications and business messaging over standard telephony infrastructure (SMS, RCS), where the trust crisis is most acute and the regulatory pressure most immediate. The framework is designed to apply subsequently to over-the-top and emerging communications platforms as they develop.

A shared identity framework across channels is not merely convenient; it is necessary. There are two reasons for this. The first is defensive: fraud does not respect channel boundaries. If one channel is secured while another remains open, fraudsters will migrate. A framework that works only for one voice standard, or only for one messaging standard, will not solve the problem.

The second reason is economic. The benefits of restoring trust in communications are not limited to fraud reduction. Across the channels this framework is designed to address, new services and features are being held back by consumer reluctance to engage. If people do not use these services because they cannot trust them, the macroeconomic gains those services are positioned to deliver will not be realised. Investment cases become unsustainable when the volumes that would justify them fail to materialise.

The scope of this framework must also extend to fixed-line and non-smartphone endpoints. Wireline carriers participate in the communications ecosystem as both originating and terminating parties, and cannot be excluded from a framework that claims to work at scale. The presentation mechanisms may differ by device type, but the underlying vetting, linking, and verification architecture applies consistently across all endpoints.

Beyond these initial channels, the same vetted identity credential is applicable wherever verified business identity is required — for example, in authorising access to operator APIs such as those being developed under the GSMA Open Gateway initiative. A related capability on the roadmap

is enabling consumers to verify the identity of an enterprise they are calling, not just the enterprise calling them: when a customer calls their bank, they could see the bank's verified name and brand on their screen, providing the same assurance in the reverse direction. This is a subsequent phase of development.

5. Principles for a Workable Framework

For a digital identity framework to be adopted at scale, it must meet a set of practical requirements that go beyond the technical. The following principles guide the work of this group.

- **Open and interoperable.** Any framework that depends on a single vendor, a proprietary protocol, or a closed registry will face adoption barriers and create new dependencies. The approach must be based on open, inspectable standards that can be independently implemented by any operator or service provider. No single proprietary provider should be able to control access to the attestation, signing, or verification infrastructure. The framework is open to multiple identity verification mechanisms and credential sources, allowing operators to participate at the level appropriate to their context and to differentiate competitively in the services they offer on top of the baseline standard.
- **Portable across channels and borders.** Identity evidence must work across channels and, ultimately, across national boundaries. Achieving adoption at national level is itself a substantial and meaningful goal; cross-border portability—where a credential established in one jurisdiction is recognised in another—is the structured evolution that follows. The framework is designed to support this progression: national implementations can be built in a way that is compatible with future cross-border recognition, without requiring global agreement before any benefit is realised.
- **Governed transparently.** The rules for who may vet, what standards apply, how credentials are issued and revoked, and how disputes are resolved must be clear and auditable. Governance is not a detail; it is the foundation of trust.
- **Transnational by design, sovereign in implementation.** The framework should be designed to operate across sovereign jurisdictions without requiring identical national rules or globally centralised enforcement. Its purpose should be to provide an internationally recognised mechanism for creating, transporting, and verifying evidence of business identity. Each jurisdiction would retain the right to determine the evidence it requires, the vetting standards it recognises, and the privileges or access rights that verified identity confers within its borders. The objective is therefore not necessarily full harmonisation of national rules, but structured interoperability: enabling identity evidence produced under one jurisdiction's governance to be understood, assessed, and relied upon by another.
- Where possible, the framework should support an internationally accepted mechanism for issuing and transporting a vetted organisational identity token. Such a token should include appropriate transparency regarding the underlying vetting process, including which entity or entities performed the verification, which organisational attributes were verified, what access privileges or use cases the identity supports, and the applicable chain of trust behind the vetting authority.
- A global framework should nevertheless seek to maximise harmonisation wherever practicable, particularly around core vetting principles, organisational identity attributes, assurance levels, and token formats. Recent work by i3Forum, One Consortium and GIRAFF on global "KYC" standards illustrates the industry's growing recognition that

trusted communications require common foundations for organisational identity verification.

- Where harmonisation is not achievable, interoperability between national or regional trust frameworks should be actively pursued. This would allow different domestic regimes to interwork while preserving national sovereignty over policy, regulation and enforcement.
- A globally coordinated approach — or, failing that, robust interoperability — would benefit citizens, businesses and public authorities by enabling trusted communications both domestically and internationally. A purely national solution addresses only part of the challenge. Business communications are increasingly cross-border, and fraudsters readily shift activity across jurisdictions to exploit gaps between national regimes. Trusted communications therefore requires a framework that can operate beyond national boundaries while respecting local legal and regulatory requirements.
- The experience of certificate-based secure web access, using X.509 certificates for HTTPS, is instructive. That framework has successfully enabled secure, domain-level trust for billions of national and international web transactions. However, it has been much less effective as a mechanism for reliable organisational identity or KYC. This demonstrates both the feasibility of global technical trust infrastructure at scale, and the need for any trusted communications framework to place stronger emphasis on verified business identity, assurance transparency and governance.

Top of Form

Bottom of Form

- **Inclusive and proportionate.** Standards must not inadvertently exclude smaller organisations, or organisations in jurisdictions where formal identity infrastructure is less developed. The framework should support a range of vetting approaches appropriate to different contexts, and should focus on incentives for adoption rather than punitive exclusion.
- **Privacy-preserving.** Only the minimum necessary information should be revealed to any given verifier. An organisation's full identity record does not need to be disclosed to every party; what matters is that the relevant claims—'this is a legitimate business with the right to use this number'—can be confirmed.
- **Actionable without mandating wholesale change.** The framework should deliver value to early adopters without requiring universal adoption before any benefit is realised. It should be possible to add one organisation at a time, with verifiers immediately able to distinguish between communications with proven identity and those without. The framework is also designed to be proportionate: the burden of participation should scale with the level of participation, and enhanced services above the baseline are a matter of commercial choice for operators.

6. Benefits by Stakeholder

For citizens and consumers

Scams initiated through calls and messages represent a significant and growing harm to consumers globally—estimated to affect more than one in five people and costing hundreds of billions of dollars annually. Consumers gain the ability to distinguish communications from genuine organisations from those that are impersonating them. Over time, as adoption grows, a communication accompanied by verified identity becomes a meaningful signal of trustworthiness.

Consumers can begin to engage with unknown-number calls and messages with greater confidence when that signal is present—and to treat its absence as a reason for caution.

For businesses and service providers

Legitimate organisations can finally demonstrate their identity affirmatively. The bank making a fraud alert call, the healthcare provider confirming an appointment, the delivery company advising on a parcel—all can communicate with confidence that their identity can be verified by the recipient. Call completion rates improve. Customer trust is rebuilt. The investment in compliant, high-quality communications pays off.

For carriers and network operators

Carriers gain a consistent, cross-network basis for identity verification that reduces their exposure to regulatory risk from delivering fraudulent traffic. They can offer verifiably authenticated calls and messages as a differentiating service, and gain a foundation on which enhanced services—branded calling, rich business messaging—can be built with genuine trust. Improved call answer rates benefit all parties. It is worth emphasising that participation in this framework does not require carriers to offer any particular set of value-added services: the framework sets a baseline of open, interoperable standards, and what operators build on top of that baseline remains a matter of commercial choice. This is by design: a genuinely open framework creates conditions for competitive differentiation rather than concentrating identity assurance in the hands of a single provider.

For regulators and enforcement agencies

When identity evidence travels with a communication, the chain of custody from originating organisation through to delivery is visible and auditable. Traceback becomes substantially more tractable, grounded in verifiable attribution of who was responsible for the content of the communication rather than just which provider carried it. Regulators in different jurisdictions can work from a common evidence base rather than having to reconcile incompatible national systems.

7. The Role of GIRAF and the Path Forward

The problem described in this briefing is not unique to any one country or regulatory jurisdiction. It is a global challenge that requires a global response—or at least a framework that can operate consistently across borders, even where its implementation varies by market. Industry can design the architecture, develop the standards, and build the tools. But without regulatory engagement, the framework will not reach its potential. There are specific points at which industry cannot succeed alone.

The urgency of this work is independently recognised at the law enforcement level. A 2025 Europol position paper on caller-ID spoofing, drawing on engagement with law enforcement agencies across 23 EU member states, called for harmonised technical standards, enhanced cross-border collaboration, and aligned regulatory frameworks as essential responses to spoofing-enabled fraud. Europol's paper also references One Consortium's work on traceback mechanisms as part of the solution landscape. The framework proposed in this briefing directly addresses the architectural gap Europol identifies. GIRAF and Europol are natural partners in translating that framework into enforceable outcomes. Secondly and remaining on the European continent, the proposed Digital Networks Act (DNA) which is, subject to final approval, slated for commencement as early as 2028, states that "The [Office for Digital Networks, i.e. the new EU

‘communications think tank’] shall collect information from national regulatory or other competent authorities, relevant European and global organisations and stakeholders, including EUROPOL and ENISA, on the fraudulent activities perpetrated by making use of [publicly available interpersonal communications services], in particular on those that are of cross-border nature.” It then follows: “BEREC, after consulting the European Data Protection Board and ENISA, shall issue guidelines on technical and legal measures that could effectively protect end-users against fraudulent activities in the Union[...]”. Therefore, especially with regards to providers active within the EU27, now is the time to act, and now is the time to agree an open and pro-competition global approach.

Critically, the framework is designed to respect national sovereignty: each jurisdiction retains full authority over its own numbering policy, vetting standards, and enforcement mechanisms. What the framework provides is a common mechanism for expressing, transporting, and verifying evidence of compliance with those sovereign requirements — enabling international communications to flow with confidence across jurisdictional boundaries.

- **Backing the presentation layer.** The moment at which all the technical and governance work becomes meaningful to a citizen is when they receive a signal that the communication has been verified. Industry can design the signal; only regulatory support can give it the authority that makes it meaningful. GIRAF members are well placed to define the outcome standard—“beyond reasonable doubt, verified”—and to support the conditions under which operators may display a trust signal and face consequences when they misuse it. This need not mean mandating a specific technology; it means providing the backing that makes the signal trustworthy.

  In some jurisdictions, NRA powers alone may not be sufficient to achieve this. Some authorities lack the legal powers, or the enforcement credibility, to police trust indicator standards effectively. In others, there may be no telecommunications-specific regulator at all. And in many cases, the bodies best placed to act against fraud and scam at scale sit outside the telecoms field—as illustrated by the EU’s call for evidence on fraud and scam, which was issued by DG Home rather than the telecoms-relevant directorate. GIRAF members should therefore consider whether advocating for domestic legislative measures—giving trust indicator protections a statutory basis that goes beyond NRA powers where necessary, and engaging the full range of domestic authorities with a stake in the outcome—is part of the response this framework requires.
- **Supporting cross-border recognition.** For identity evidence to travel with communications across national boundaries, there must be a basis for mutual recognition—whether through harmonised standards, bilateral or multilateral agreements, or common governance frameworks. Industry can design credentials that are technically portable; whether they are accepted at termination in a foreign jurisdiction is a matter of regulatory agreement. GIRAF is the right forum to begin that conversation.
- **Shaping the governance of governance authorities.** The framework is deliberately agnostic about which bodies may act as governance authorities, recognising that rules differ across jurisdictions and that multiple authorities may coexist. But this agnosticism defers a question regulators must ultimately answer: what standards must a governance authority meet for its credentials to be recognised? And what happens when an authority’s rules are broken or its credentials are misused? Strong attribution—the ability to trace a fraudulent communication back to the party responsible—is a core property of this architecture. The consequences for misuse, and the bodies empowered to impose them, remain matters of regulatory design. GIRAF is well placed to lead that conversation.

- **Aligning on incentive structures.** Adoption is unlikely to be mandated everywhere at once, and industry is not asking for that. But regulators can create conditions that reward adoption and make the absence of verifiable identity increasingly visible—through labelling, through liability frameworks, or through conditions attached to number assignments. Industry will bring the solutions; regulatory frameworks that recognise and reward verified identity will accelerate adoption.
- **Engaging in standards evolution.** The technical standards underpinning this framework are being developed in relevant standards bodies, including work at GSMA and in 3GPP and IETF; and as stated above, ENISA, the European Union Agency for Cybersecurity, is expected to make known its views. Regulatory input into those processes ensures the resulting standards are fit for purpose from a governance and consumer-protection perspective. Equally important is who governs the evolution of those standards over time—the communications industry is highly dynamic, and new use cases will emerge faster than any fixed standard can anticipate. GIRAF members have a role not only in shaping standards at inception, but in ensuring that governance processes for standards evolution remain transparent, inclusive, and responsive to regulatory and consumer interests as the framework matures.

> *The question for GIRAF is not whether to engage with digital identity for communications, but how—and how quickly—to shape it in a direction that serves the interests of citizens and regulators, rather than simply inheriting the direction set by commercial interests alone.*

8. Next Steps and Further Work

This briefing is the first in a series of documents the One Consortium Digital Identity Working Group will produce to support informed engagement by GIRAF members. Our next publication in this series will contain further information on:

- Business vetting: how organisations can be identified and credentialled, what governance structures are needed, and how vetting can work across jurisdictions with different identity infrastructure.
- Linking identity to communications: the technical options for connecting identity evidence to individual calls and messages, and the trade-offs between different approaches.
- Incentives and adoption: how adoption can be encouraged without mandating it, how smaller organisations and those in less-developed identity ecosystems can participate, and what role regulators can play in creating the right conditions.

One Consortium is well placed to support this work by providing practical frameworks and templates—covering technical architecture, governance structures, and policy principles—for the key functions of a verified identity model. This includes supporting the structured evolution from national implementations toward international and multinational coordination, ensuring that early deployments are built in ways that can grow into cross-border interoperability rather than creating new silos.

The working group welcomes engagement from GIRAF members at every stage. The design choices involved in building this framework—on governance, on cross-border recognition, on the relationship between national and international standards—are not purely technical. They are policy questions that require regulatory decisions.

We look forward to working together on a framework that can restore trust in the communications that citizens and businesses depend upon every day.

Glossary of Key Terms

The following terms are used in this briefing with the meanings set out below. More precise technical definitions will be developed in subsequent working group documents.

Authenticated (communication): A communication in which certain properties—such as the integrity of the call setup, or the signing of the communication by the originating party—have been confirmed by a technical mechanism. Note: authentication of the call setup path does not by itself confirm that the communicating organisation has been vetted or has the right to use the identifier it is presenting. The distinction between signing and vetting is important: a communication can be technically authenticated while still originating from an entity that has no legitimate right to use the identifier in question.

Verifiable: Capable of being independently checked by any relying party, without requiring them to trust that someone else has already done so. A verifiable credential or claim is one whose evidence can be examined directly.

Business entity: An organisation (company, public body, service provider, or similar) that communicates with consumers. The identity framework described in this briefing focuses on business entities, not individual persons.

Credential: A digitally signed document that records the result of a vetting process: what was verified, by whom, and when. Credentials can be updated or revoked in real time.

KYB (Know Your Business): The process of verifying that a business entity is who it claims to be—confirming legal existence, ownership, and other relevant attributes—as the foundation for issuing credentials.

Linking (of identity to a communication): The process of cryptographically connecting identity evidence to a specific call or message, so that any verifier can confirm the evidence relates to that particular communication, not just to the organisation in general. Sometimes referred to as “binding” in technical contexts. This use of the term should not be confused with “linkability” in data protection law (as used in GDPR), which concerns the use of identifiers to correlate an individual’s activities across separate transactions or channels. Cryptographic linking of identity evidence to a specific communication is a distinct concept that does not involve cross-transaction correlation of personal data.

Portable (identity evidence): Identity evidence that can be used across channels (voice, messaging) and across borders, without needing to be re-issued for each context.

RTU (Right to Use): The verified entitlement of a business entity to use a specific identifier—telephone number, short code, or alphanumeric Sender ID—in communications. RTU verification confirms that the entity has been assigned or licensed that identifier by the appropriate authority.

Vetting: The process of verifying that a business entity is who it claims to be, that it has the right to use the numbers and brand assets it is deploying, and that it meets any other relevant standards for the communications context in which it is operating.

About this document

This document is a draft produced by the One Consortium Digital Identity Working Group for review and comment by GIRAF members. It represents the views of working group participants and does not constitute a formal position of One Consortium. Comments and contributions are welcome.

Authors: Randy Warshaw (Provenant), Eli Katz (XConnect), and the OC-DIWG.