

The History and Evolution of Voice Recognition Technology

By ClearlyIP | Published April 18, 2025 | 15 min read



The History and Future of Voice Recognition

Speech recognition has evolved over the past seven decades from primitive analog devices to sophisticated AI-driven systems. Early research in the 1950s–1970s laid the groundwork for today’s accurate, end-to-end deep learning models. We begin with a chronological overview, then examine key technological milestones (algorithms and hardware), major contributors (industry and academia), current state-of-the-art technologies, commercial applications, ongoing challenges, and future directions of voice recognition.

Historical Development of Voice Recognition

Early voice recognition was very limited. In 1952 [Bell Labs](#) built **Audrey**, a system that recognized spoken digits [computerhistory.org](#). A decade later IBM demonstrated the **"Shoebbox"** at the 1962 World's Fair – a machine that understood 16 spoken words (digits 0–9 plus commands) [computerhistory.org](#). These analog-era devices could only handle tiny vocabularies, but proved that computers could process [human speech](#).

Figure: Bell Labs' "Audrey" (1952) – one of the first speech recognizers for digits [computerhistory.org](#).

In the 1970s, government research accelerated. DARPA funded speech projects (1971–1976) aiming for 1,000-word recognition [computerhistory.org](#). This led to CMU's **HARPY** system (1976) using beam search, which achieved DARPA's target by recognizing about 1,000 words [computerhistory.orgcomputerhistory.org](#). Around the same time, work on **Hearsay** (also at CMU) and other rule-based systems pioneered **continuous speech** (speaking without pauses) and early "blackboard" architectures, though these were hand-engineered by linguistic rules and templates [computerhistory.orgcomputerhistory.org](#).

The 1980s saw the introduction of statistical methods. At CMU, James Baker and others applied *hidden Markov models* (HMMs) to speech [en.wikipedia.org](#), replacing older pattern-matching (DTW) approaches. In fact, "the HMM proved to be a highly useful way for modeling speech and replaced dynamic time warping to become the dominant speech recognition algorithm in the 1980s" [en.wikipedia.org](#). IBM's **Tangora** system (mid-1980s) was one of the first to use HMMs to recognize up to 20,000 words [en.wikipedia.orgcomputerhistory.org](#). Simultaneously, Hidden Markov and n-gram language models became standard, enabling statistical recognition of larger vocabularies [en.wikipedia.orgcomputerhistory.org](#).

Figure: IBM's Tangora project (circa 1985) – led by Fred Jelinek, it used HMMs for large-vocabulary dictation (≈20,000 words) [en.wikipedia.orgcomputerhistory.org](#).

Commercialization began in the 1980s. In 1974, James and Janet Baker developed **DRAGON** at Carnegie Mellon using HMMs (later forming Dragon Systems) [computerhistory.org](#). In 1982 Dragon Systems was founded to commercialize speech software [computerhistory.org](#). By 1984 they licensed their recognizer to the Apricot PC – *the first personal computer with built-in speech recognition* [computerhistory.org](#). IBM also rejoined the field with Tangora and later ViaVoice. DARPA-sponsored competitions (e.g. through NIST) and the rise of Digital Signal Processing hardware drove rapid accuracy gains by decade's end [computerhistory.orgcomputerhistory.org](#).

Figure: ACT Apricot personal computer (1980s) – Dragon Systems' software in 1984 gave it built-in voice recognition computerhistory.org.

By the late 1980s and early 1990s, **large-vocabulary continuous speech recognition (LVCSR)** became practical. CMU's Kai-Fu Lee (a former Baker student) combined Dragon's HMMs with HARPY's beam search to create **SPHINX-1 (1987)** computerhistory.org. SPHINX-I was the first speaker-independent system capable of continuous recognition on large vocabularies computerhistory.org. In 1990 Dragon released **Dragon Dictate**, the first consumer speech product (discrete, "talk-and-pause" dictation). Dragon's major breakthrough came in **1997** with **Dragon NaturallySpeaking**, the first truly continuous dictation system for PCs (recognizing ≈100 words/minute, albeit requiring some user training) computerhistory.org. Microsoft licensed Dragon/Nuance technology for Windows and Office (e.g. Speech API, Office Dictation), and IBM's ViaVoice appeared in the late 1990s. By 2000, systems like these could handle vocabularies in the tens of thousands, though in quiet environments computerhistory.orgwikipedia.org.

Figure: CMU's Kai-Fu Lee (center) developed SPHINX-I (1987) – the first speaker-independent large-vocabulary speech recognizer computerhistory.org.

The 2000s saw voice recognition go mobile and cloud. Google introduced **Voice Search** (e.g. phone-based GOOG-411 in 2007, and iPhone app in 2008) and later built it into Android. In **2007–2011**, Microsoft integrated speech into Windows Vista and Office, and launched **Cortana** (2014). Apple debuted **Siri** in 2011 (licensed from Nuance) en.wikipedia.org, pioneering smartphone voice assistants. In **2014**, Amazon's **Echo/Alexa** brought always-on voice control to homes en.wikipedia.org. These assistants used cloud-based ASR to enable voice commands, search, and smart-home control. Over this period, advancements in machine learning (detailed below) steadily improved accuracy. By 2016–2017, deep neural network-based systems achieved **human-level accuracy** on conversational English (on par with professional transcribers at ~5–6% word error rate) computerhistory.orgmicrosoft.com.

For reference, Table 1 summarizes key milestones:

YEAR	EVENT/TECHNOLOGY	NOTES/CITATION
1952	Bell Labs <i>Audrey</i> (digits)	Recognized spoken digits computerhistory.org
1962	IBM <i>Shoebbox</i>	Understood 16 words at Seattle World's Fair computerhistory.org
1976	CMU <i>Harpy</i>	Recognized 1,000 words (DARPA goal met) computerhistory.org computerhistory.org
1987	CMU <i>SPHINX-I</i>	Speaker-independent LVCSR; won DARPA eval computerhistory.org
1997	Dragon <i>NaturallySpeaking</i>	First consumer continuous speech product (100 wpm) computerhistory.org
2007	Google Voice Search (mobile)	Early mobile speech search
2011	Apple <i>Siri</i>	iPhone voice assistant (speech + understanding) en.wikipedia.org
2014	Amazon <i>Alexa</i> (Echo)	Home smart speaker voice assistant en.wikipedia.org
2016	Deep Learning ASR (Microsoft et al.)	Surpassed human parity on Switchboard (~5.9% WER) computerhistory.org microsoft.com
2022	OpenAI <i>Whisper</i>	Multilingual end-to-end Transformer ASR (68k hrs data) openai.com openai.com

Key Technological Milestones

Statistical Models (1980s–2000s): The shift from rule-based templates to statistical methods (HMMs and n-gram language models) in the 1980s dramatically improved accuracy en.wikipedia.orgcomputerhistory.org. By modeling phoneme probabilities and word sequences, HMM-based systems achieved much higher vocabularies. In the 1990s, large corpora and benchmarking (DARPA/NIST evaluations) drove rapid progress: for example, CMU's Sphinx system won the 1992 LVCSR DARPA eval with continuous speech en.wikipedia.org. Accuracy steadily

improved: early 1990s systems had word-error-rates (WER) in the tens of percent, but by decade's end – after optimizations like discriminative training and larger language models – WER on tasks like broadcast news or Wall Street Journal reading fell to ~10–15%. Users nonetheless still had to “train” systems to their voice and speak carefully.

Deep Learning Revolution (2010s): Around 2010, deep neural networks (DNNs) began supplanting Gaussian mixture models. Hinton et al. showed that DNN acoustic models trained on large data gave *orders-of-magnitude* reductions in WER en.wikipedia.org. Google reported a 49% error reduction in 2015 by switching to LSTM-based models on voice search en.wikipedia.org. The 2016 Microsoft “human parity” system used CNN+LSTM, reaching ~5.9% WER on conversational English microsoft.com/computerhistory.org. In practice, DNNs (and later sequence models) became ubiquitous: nearly all state-of-the-art ASR now use deep nets with end-to-end training.

End-to-End & Transformers (2020s): More recently, end-to-end architectures (CTC, RNN-Transducer, sequence-to-sequence) have simplified pipelines by jointly learning acoustics and language. Self-attention **Transformer** models (originally from NLP) now dominate. Facebook/Meta's **wav2vec 2.0** (2020) showed that self-supervised pretraining on raw audio plus fine-tuning on limited data achieves state-of-the-art performance venturebeat.com. OpenAI's **Whisper** (2022) is an encoder–decoder Transformer trained on 680k hours of multilingual data; it matches or surpasses prior ASR robustness on accent and noise openai.com/openai.com. New architectures like Google's **Conformer** (combining convolutions with Transformers) also report top-tier accuracy. These models routinely handle noisy audio, far-field speech, and dozens of languages.

Accuracy Benchmarks: Systems' error rates have plummeted. For example, a benchmark task (Switchboard English conversation) saw WERs fall from ~20% in the 2000s to around **6% by 2016** (exceeding the ~5.9% human rate) microsoft.com/computerhistory.org. Similar progress occurred across tasks. Figure 1 (above) and Table 1 illustrate this timeline of error reduction. Today's research focuses on squeezing out the last errors (e.g. handling overlapping speech, rare words) and expanding to more languages.

Hardware and Computing: These advances were enabled by hardware. Early systems ran on mainframes (e.g. a 4 MB PDP-10 took 100 minutes to decode 30 seconds of speech in 1976 en.wikipedia.org). By the 2000s, GPUs and cloud compute allowed real-time DNN training. Today, dedicated AI accelerators (GPUs, TPUs, mobile DSPs/NPUs) allow on-device recognition for instant responses. For example, modern smartphones include neural engines to run ASR locally (e.g. wake-word detection), reducing latency and privacy risk.

Leading Players and Research Institutions

Research and development in speech recognition has been driven by a mix of industry labs and academia. Historic contributors include **IBM** (Shoebbox, Tangora, ViaVoice, Watson Speech-to-Text), **AT&T Bell Labs** (LPC coding, Dictation Lab; led by Lawrence Rabiner), **CMU** (Raj Reddy's Harpy and Sphinx, Kai-Fu Lee's SPHINX, grants), **BBN**, **SRI**, **MIT/LCS**, and others [microsoft.com/en-us/research/publications/](https://www.microsoft.com/en-us/research/publications/). Dragon Systems (later Nuance) was a key startup turned industry leader in dictation and medical ASR.

Today, major commercial players dominate: **Google** (Voice Search, Assistant, YouTube auto-caption, TensorFlow AI) is a leader in large-scale ASR; **Microsoft** offers Azure Speech and Cortana; **Amazon** created Alexa and AWS Transcribe; **Apple** runs Siri (with much on-device processing for privacy); **Nuance (Microsoft)** continues in niche dictation markets (medical/legal) [en.wikipedia.org](https://en.wikipedia.org/wiki/Nuance). Other notable contributors include **Facebook/Meta** (wav2vec, XLSR models), **Baidu** (Deep Speech, commercial voice apps), and **Xiaomi, Tencent, Alibaba** (voice in Chinese). Academic groups remain influential: CMU (Sphinx, CLSP), Stanford (Huang, Wocher), MIT, Tsinghua (speech labs), Johns Hopkins (CLSP), Cambridge (RWTH Aachen), etc., continually publish ASR research.

In summary, speech recognition has been advanced by a few large tech companies (IBM, Microsoft, Google, Amazon, Apple, Nuance) and many universities/institutes. These players drive both the core algorithms (machine learning, language modeling) and applications (assistants, accessibility tools).

State-of-the-Art Technologies

Modern voice recognition leverages end-to-end deep learning. In today's systems:

- **End-to-End Neural Models:** Speech is typically processed by a deep neural network that directly maps audio to text. Architectures include RNN-Transducers and encoder-decoder (seq2seq) models, which remove the need for separate acoustic/phonetic and language models. For example, Google's production ASR uses a streaming RNN-T to transcribe on the fly.
- **Self-Supervised Transformers:** Pretrained on massive unlabeled audio, transformer-based models like **wav2vec 2.0** and **Whisper** capture rich speech representations. These models handle noise, accents, and varied domains. OpenAI's Whisper (multilingual transformer) "approaches human-level robustness and accuracy" on English and supports transcription in multiple languages with good noise resilience openai.com. Facebook's wav2vec 2.0

achieved state-of-the-art with only 10 minutes of labeled data by leveraging 53k hours of unlabeled speech venturebeat.com. Such models have set new benchmarks (often achieving <5% WER on English benchmarks).

- **Multilingual and Zero-Shot Models:** Systems now aim to serve dozens of languages. Training on multilingual data lets one model generalize cross-lingually. For instance, Whisper can transcribe 100+ languages and even translate speech to English openai.com. Cutting-edge research explores **zero-shot ASR**, where a single acoustic model outputs language-independent units. Meta AI's recent MMS Zero-Shot framework (trained on 1,078 languages) reduced character-error by ~46% on 100 unseen languages, with no language-specific training arxiv.org. This suggests ASR can soon handle low-resource languages without transcripts.
- **Speaker Diarization and Context:** Many systems now include **diarization** ("who spoke when") to segment multi-speaker audio. Recent advances in neural diarization allow reasonably accurate speaker labeling in meetings or broadcasts. End-to-end models are starting to jointly optimize transcription and speaker ID. Similarly, models incorporate context via attention to improve punctuation and phrasing in transcription.
- **Real-Time and On-Device ASR:** Modern voice assistants and mobile apps require streaming, low-latency recognition. On-device keyword spotting and first-pass recognition is common (e.g. Alexa's wake-word runs locally). GPUs and NPUs on phones enable running neural ASR models offline. Meanwhile, cloud services (e.g. Azure Speech, Google Cloud Speech-to-Text) can handle heavier loads. The trend is "edge intelligence" – running as much as possible on-device for responsiveness and privacy.
- **Deep Learning Integration:** All current production ASR systems are deep learning-based. From Google's massive re-training of voice models to Apple's use of DNNs on-device, end-to-end neural modeling is ubiquitous en.wikipedia.org/computerhistory.org. Even hybrid systems (neural acoustic + neural LM) dominate.

Applications of Voice Recognition

Voice recognition is now ubiquitous across domains:

- **Digital Assistants:** Apple Siri, Google Assistant, Amazon Alexa, Microsoft Cortana, Samsung Bixby and others use ASR to accept voice commands for search, navigation, scheduling, smart-home control, etc. These platforms (available on smartphones, speakers, cars) have hundreds

of millions of users. For example, Google Assistant runs on over 500 million devices and supports voice queries in 30+ languages [blog.google](#). Amazon Echo (Alexa) has similarly popularized voice-driven IoT control [en.wikipedia.org](#).

- **Dictation & Transcription:** Consumer and enterprise applications use ASR to convert speech to text. Dragon NaturallySpeaking (Nuance) dominates professional dictation (medical/legal). General-purpose dictation is built into OSs and applications (e.g. Windows Speech Recognition, Google Docs Voice Typing). Video platforms (YouTube, Zoom, Teams) and broadcast media auto-generate captions using ASR. Call centers and voicemail systems use automated transcription for record-keeping and analytics.
- **Accessibility:** ASR provides accessibility for the deaf or motor-impaired. Live captioning services (e.g. Google Live Transcribe, Apple Live Caption) use speech recognition to display spoken words in real time. Voice control enables hands-free device use for mobility-impaired users.
- **Automotive:** Cars increasingly offer voice interfaces (navigation, climate control). Examples include Siri eyes-free mode (CarPlay), Android Auto, and proprietary systems (Mercedes MBUX, Ford SYNC). These rely on robust far-field ASR that works in noisy, moving environments.
- **Customer Service & IVR:** Automated speech recognition powers phone menus and chatbots. Many companies deploy ASR in call centers for interactive voice response (IVR) systems. Newer AI-driven “voicebots” can understand natural queries (e.g. “I want to change my billing address”) and route calls or even handle simple dialogs.
- **Industry-Specific Uses:** ASR is specialized for sectors like healthcare (medical transcription), legal depositions, banking (voice biometrics for authentication), and media (subtitle generation). For instance, Nuance’s Dragon Medical One is a cloud-based ASR platform for clinical notes.

[computerhistory.org](#) By the mid-2010s, voice assistants (Alexa, Google Home, Siri, Cortana) became commonplace consumer products [computerhistory.org](#), and today ASR underpins search, smart devices, dictation, captioning, automotive interfaces, and much more.

Challenges and Limitations

Despite advances, speech recognition faces challenges:

- **Noise and Acoustic Variability:** Background noise, reverberation, and distance (far-field speech) degrade performance. Multi-microphone beamforming helps, but real-world robustness is still imperfect. ASR accuracy drops with noisy ambient conditions, requiring noise-robust features and data augmentation.
- **Accents and Dialects:** Systems are often trained on standard accents. Non-native or regional accents can dramatically increase error rates. For example, a Stanford study found leading ASR systems make *twice as many errors* on African-American speakers' voices versus White speakers news.stanford.edu. Similar disparities exist for other non-standard dialects. Users must "speak clearly" or adapt to the system's expected accent.
- **Bias and Fairness:** Along with accent bias, gender and age can affect recognition. Past research shows ASR trained on imbalanced datasets may favor certain groups. These biases are an active area of concern, as voice tech becomes widespread.
- **Privacy and Security:** Always-on microphones and cloud processing raise privacy worries. Critics note that ASR "censors" speakers outside the trained accents or languages scientificamerican.com. There have been lawsuits alleging devices record private conversations without consent reuters.com. Regulations (GDPR, COPPA) require firms to protect voice data; companies increasingly process voice on-device to mitigate risk.
- **Contextual Understanding:** Transcribing speech doesn't guarantee understanding intent. Ambiguities (homophones, sarcasm) require natural language understanding on top of ASR. Current ASR passes raw text to NLP modules, but deeper contextual comprehension (semantics, pragmatics) remains limited.
- **Code-Switching and Multilingual Speech:** In many regions, speakers mix languages ("code-switching") mid-utterance. ASR typically assumes a single language context and thus fails on mixed utterances. As Scientific American notes, code-switching is effectively not handled by current ASR – "no code-switching ... either you assimilate, or you are not understood" scientificamerican.com. Low-resource languages also pose problems due to scarce training data.
- **Adversarial Inputs:** Like other AI, ASR can be fooled by adversarial audio or mimicry, raising security issues (e.g. hidden commands).

These challenges mean that despite high average accuracy, real-world ASR systems must be carefully tuned and supplemented with failsafes.

Future Trends in Voice Recognition

Looking ahead, research and product development point to several trends:

- **Zero-Shot and Few-Shot Recognition:** Models that generalize to new words, accents, or languages without retraining. Recent work trains one multilingual acoustic model on hundreds of languages, then transcribes unseen ones ("zero-shot") [arxiv.org](#). Such approaches will enable ASR for underrepresented languages by sharing phonetic knowledge.
- **Brain-Computer Speech Interfaces:** The cutting edge is moving beyond ear/microphone interfaces. In March 2025, UC Berkeley/Stanford researchers reported a real-time "brain-to-voice" neuroprosthesis: decoding neural signals directly into intelligible synthesized speech [engineering.berkeley.edu](#). This uses deep learning to stream brain activity into vocal output, promising communication for paralyzed patients. While still experimental, it illustrates a future where "speech" may bypass the vocal tract entirely.
- **Embedded and Edge AI Chips:** Future devices will have ever more powerful on-chip AI processors (NPUs, DSPs, neuromorphic chips) dedicated to speech. Google's Edge TPU, Apple's Neural Engine and Silicon, and Qualcomm's AI Engine already accelerate ASR. Next-generation hardware will enable always-on, ultra-low-power speech recognition, improving privacy and responsiveness.
- **Federated and Privacy-Preserving Learning:** To address privacy, companies are exploring on-device model updates. Google uses federated learning for voice data: the model improves by learning on-device user speech without uploading raw audio [support.google.com](#). This trend will grow, allowing personalization while keeping data local.
- **Human-Like Conversational Agents:** Voice recognition will merge with large language models for more natural dialog. Systems like Google Duplex (2018 demo) show near-human telephone conversations [wired.com](#). Future assistants (powered by GPT-style models) will have realistic prosody, context awareness and humor. OpenAI and others are already prototyping voice-chat versions of chatbots. We may see virtual agents that not only transcribe speech but hold extended, nuanced spoken dialogues.
- **Multilingual Synthesis and Translation:** Advances in text-to-speech (parallel to ASR) will produce natural multilingual voices and real-time speech-to-speech translation. For example, Whisper already translates speech to English. In future, one might speak any language and immediately hear it rendered in another by real-time neural systems.

- **Immersive Voice Environments:** With AR/VR and the metaverse, voice interfaces will extend into 3D worlds. Recognizers will need to handle spatial audio cues and multiple virtual speakers seamlessly.

In summary, voice recognition is heading towards more universal, on-device, and human-like systems. Models are growing to cover every language (7,000+ spoken), dialect, and speaking style with minimal explicit training. The integration of speech with brain-computer interfaces and continual learning will open new frontiers.

Conclusion: Voice recognition has progressed from niche digit-understanding machines to pervasive AI systems. Through decades of advances in modeling (HMM→DNN→Transformers) and hardware (mainframe→GPU→Edge TPU), ASR now serves billions of users. Continued research in robustness, fairness, and self-supervised learning will shape its future. The next frontier blends voice with cutting-edge AI – from zero-shot language generalization to mind-to-speech interfaces – promising richer, more natural human-computer speech interaction openai.com/engineering/berkeley.edu.

Sources: Authoritative histories and research reports provide the above information computerhistory.org computerhistory.org openai.com/scientificamerican.com news.stanford.edu arxiv.org microsoft.com. Tables and figures summarize the timeline and key milestones for quick reference.

Tags: voice recognition, speech recognition, ai, deep learning, technology history, algorithms, human computer interaction, technological evolution

About ClearlyIP

ClearlyIP Inc. — Company Profile (June 2025)

1. Who they are

ClearlyIP is a privately-held unified-communications (UC) vendor headquartered in Appleton, Wisconsin, with additional offices in Canada and a globally distributed workforce. Founded in 2019 by veteran FreePBX/Asterisk contributors, the firm follows a "build-and-buy" growth strategy, combining in-house R&D with targeted acquisitions (e.g., the 2023 purchase of Voneto's EPlatform UCaaS). Its mission is to "design and develop the world's most respected VoIP brand" by delivering secure, modern, cloud-first

communications that reduce cost and boost collaboration, while its vision focuses on unlocking the full potential of open-source VoIP for organisations of every size. The leadership team collectively brings more than 300 years of telecom experience.

2. Product portfolio

- **Cloud Solutions** – Including *Clearly Cloud* (flagship UCaaS), **SIP Trunking**, **SendFax.to** cloud fax, **ClusterPBX OEM**, **Business Connect** managed cloud PBX, and **EPlatform** multitenant UCaaS. These provide fully hosted voice, video, chat and collaboration with 100+ features, per-seat licensing, geo-redundant PoPs, built-in call-recording and mobile/desktop apps.
 - **On-Site Phone Systems** – Including CIP PBX appliances (FreePBX pre-installed), ClusterPBX Enterprise, and Business Connect (on-prem variant). These offer local survivability for compliance-sensitive sites; appliances start at 25 extensions and scale into HA clusters.
 - **IP Phones & Softphones** – Including CIP SIP Desk-phone Series (CIP-25x/27x/28x), fully white-label branding kit, and *Clearly Anywhere* softphone (iOS, Android, desktop). Features zero-touch provisioning via Cloud Device Manager or FreePBX "Clearly Devices" module; Opus, HD-voice, BLF-rich colour LCDs.
 - **VoIP Gateways** – Including Analog FXS/FXO models, VoIP Fail-Over Gateway, POTS Replacement (for copper sun-set), and 2-port T1/E1 digital gateway. These bridge legacy endpoints or PSTN circuits to SIP; fail-over models keep 911 active during WAN outages.
 - **Emergency Alert Systems** – Including **CodeX** room-status dashboard, **Panic Button**, and **Silent Intercom**. This K-12-focused mass-notification suite integrates with CIP PBX or third-party FreePBX for Alyssa's-Law compliance.
 - **Hospitality** – Including **ComXchange** PBX plus PMS integrations, hardware & software assurance plans. Replaces aging Mitel/NEC hotel PBXs; supports guest-room phones, 911 localisation, check-in/out APIs.
 - **Device & System Management** – Including **Cloud Device Manager** and **Update Control (Mirror)**. Provides multi-vendor auto-provisioning, firmware management, and secure FreePBX mirror updates.
 - **XCast Suite** – Including Hosted PBX, SIP trunking, carrier/call-centre solutions, SOHO plans, and XCL mobile app. Delivers value-oriented, high-volume VoIP from ClearlyIP's carrier network.
-

3. Services

- **Telecom Consulting & Custom Development** – FreePBX/Asterisk architecture reviews, mergers & acquisitions diligence, bespoke application builds and Tier-3 support.
- **Regulatory Compliance** – E911 planning plus **Kari's Law**, **Ray Baum's Act** and **Alyssa's Law** solutions; automated dispatchable location tagging.

- **STIR/SHAKEN Certificate Management** – Signing services for Originating Service Providers, helping customers combat robocalling and maintain full attestation.
 - **Attestation Lookup Tool** – Free web utility to identify a telephone number's service-provider code and SHAKEN attestation rating.
 - **FreePBX® Training** – Three-day administrator boot camps (remote or on-site) covering installation, security hardening and troubleshooting.
 - **Partner & OEM Programs** – Wholesale SIP trunk bundles, white-label device programs, and ClusterPBX OEM licensing.
-

4. Executive management (June 2025)

- **CEO & Co-Founder: Tony Lewis** – Former CEO of Schmooze Com (FreePBX sponsor); drives vision, acquisitions and channel network.
 - **CFO & Co-Founder: Luke Duquaine** – Ex-Sangoma software engineer; oversees finance, international operations and supply-chain.
 - **CTO & Co-Founder: Bryan Walters** – Long-time Asterisk contributor; leads product security and cloud architecture.
 - **Chief Revenue Officer: Preston McNair** – 25+ years in channel development at Sangoma & Hargray; owns sales, marketing and partner success.
 - **Chief Hospitality Strategist: Doug Schwartz** – Former 360 Networks CEO; guides hotel vertical strategy and PMS integrations.
 - **Chief Business Development Officer: Bob Webb** – 30+ years telco experience (Nsight/Cellcom); cultivates ILEC/CLEC alliances for Clearly Cloud.
 - **Chief Product Officer: Corey McFadden** – Founder of Voneto; architect of EPlatform UCaaS, now shapes ClearlyIP product roadmap.
 - **VP Support Services: Lorne Gaetz** (appointed Jul 2024) – Former Sangoma FreePBX lead; builds 24x7 global support organisation.
 - **VP Channel Sales: Tracy Liu** (appointed Jun 2024) – Channel-program veteran; expands MSP/VAR ecosystem worldwide.
-

5. Differentiators

- **Open-Source DNA:** Deep roots in the FreePBX/Asterisk community allow rapid feature releases and robust interoperability.
- **White-Label Flexibility:** Brandable phones and ClusterPBX OEM let carriers and MSPs present a fully bespoke UCaaS stack.

- **End-to-End Stack:** From hardware endpoints to cloud, gateways and compliance services, ClearlyIP owns every layer, simplifying procurement and support.
 - **Education & Safety Focus:** Panic Button, CodeX and e911 tool-sets position the firm strongly in K-12 and public-sector markets.
-

In summary

ClearlyIP delivers a comprehensive, modular UC ecosystem—cloud, on-prem and hybrid—backed by a management team with decades of open-source telephony pedigree. Its blend of carrier-grade infrastructure, white-label flexibility and vertical-specific solutions (hospitality, education, emergency-compliance) makes it a compelling option for ITSPs, MSPs and multi-site enterprises seeking modern, secure and cost-effective communications.

DISCLAIMER

This document is provided for informational purposes only. No representations or warranties are made regarding the accuracy, completeness, or reliability of its contents. Any use of this information is at your own risk. ClearlyIP shall not be liable for any damages arising from the use of this document. This content may include material generated with assistance from artificial intelligence tools, which may contain errors or inaccuracies. Readers should verify critical information independently. All product names, trademarks, and registered trademarks mentioned are property of their respective owners and are used for identification purposes only. Use of these names does not imply endorsement. This document does not constitute professional or legal advice. For specific guidance related to your needs, please consult qualified professionals.