


The Future of Mission-Critical Voice: Looking ahead at the next 10 years

By: Antoine Stephen
Head of Product, Voice & Cloud9



Looking ahead at the next 10 years

The next decade of voice communication will look radically different than the past 10 years.

Today's (and certainly tomorrow's) trader demands technology that is fundamentally incompatible with the rigidity of legacy systems - workers are increasingly reliant on mobile, flexible workplaces and AI. The end of the on-premise era brings an imperative for critical trading systems to deliver not only zero-downtime reliability but also unprecedented flexibility across multiple endpoints with integrated intelligence.

We are committed to consistent innovation and questioning the status quo. This is part of our legacy, which Leo Papadopoulos discusses in his piece looking at the founding innovations of Cloud9. As we look ahead to the next 10 years, the team is working hard to build upon this legacy.



Looking ahead at the next 10 years

The below are some of the areas our engineering team is focused on building:

Public Cloud: redefining mission-critical resilience

Being on the public cloud is not merely a hosting decision; it is an architectural foundation for resilience. This is one of the early decisions that influences how we approach product design today. Cloud9 is engineered to fully embrace the benefits of hyperscale infrastructure, a contrast that highlights the fragility of legacy private cloud setups.

The core engineering challenge with a cloud-native approach is to leverage the hyperscale power and redundancy of the public cloud while maintaining the low latency and security essential for mission-critical trading. Over the past few years, this philosophy has enabled us to expand the number of endpoints to access our software - on mobile, on tablet, beyond the trading floor.



The key to our higher reliability is the powerful combination of the public cloud and modern RTC technologies, such as WebRTC. This pairing allows the control plane to be massively redundant while the voice plane maintains dynamic, low-latency connections.

A traditional system or even a private cloud deployment typically operates with limited redundancy, often relying on just two data centers. This limits geographic BCP scope and introduces critical latency during failover.

Cloud9 is built to maximize the cloud's native capabilities, providing multi-region and multi-zone deployment. This distributed model is a key differentiator:

- **Superior redundancy:** We deliver multi-region and multi-zone redundancy for all control plane services, and continue to expand to more regions as our customer base grows. Even without relying on the cloud for core P2P voice traffic, this architecture ensures exceptional resilience, far surpassing the two-data-center model.
- **Inherent BCP:** Our design ensures inherent Business Continuity Planning (BCP) by dispersing services globally, making us resilient to regional outages, something private-cloud legacy competitors cannot match.
- **Scalability and elasticity:** The public cloud provides the necessary elasticity to scale resources instantly during peak market volumes or for rapid global deployment, bypassing the capital expenditure and provisioning delays inherent in private infrastructure.

Directory: from disconnected silos to a community platform

The true value of a communication platform lies in its ability to facilitate connections. Legacy systems are fundamentally disconnected and even those with large user bases often run multiple, non-interoperable platforms, creating silos of identity that prevent seamless counterparty discovery and connectivity.



Cloud9 is designed as a community platform. The core of its design is a unified, open community – the backbone to the directory.

This community model is built for the future:

Unified discovery: It allows users to easily find, connect and call counterparties across firms, a seamless experience impossible with siloed legacy platforms.

Open architecture and interoperability potential: Our architecture is designed to be a connective layer. By opening the Cloud9 Directory to the broader Symphony platforms, we not only build a professional community but also position ourselves to facilitate communication across other mediums (i.e. Symphony Messaging, WhatsApp, Microsoft Teams). This expands the scope of our directory capabilities.

Crucially, our underlying technology possesses the engineering capability to bridge diverse audio streams – meaning we could, hypothetically, connect a secure trader voice audio directly with a compliance enabled third-party application like WhatsApp voice. This inherent flexibility transforms the Cloud9 Directory into a unified professional identity, future-proofing firms for global, compliant communication regardless of the channel.

By 2035, we predict that the industry will have moved towards a unified directory model where traders don't know what tools others are using. We're working hard to prepare for this evolution by expanding our directory scope and federating with other communities.

End Points: the multi-device future

The reliance of legacy systems on proprietary, fixed hardware is a critical vulnerability. Our SaaS model and public cloud architecture enable a modern, flexible multi-endpoint story, which we believe is the future. In the next decade, we'll see traders moving out of the trading floor to an increasingly mobile workflow.



Mobile and Browser-Enabled Flexibility

Our use of WebRTC technology is foundational to this flexibility. It proves the point that a more open, public cloud platform allows us to move seamlessly into other form factors. WebRTC enables us to implement a Trader Voice client on any browser-enabled platform, including mobile devices, but not limited to them. This means a user could sign in and run essential voice flows on virtually any device with a browser, providing compliant, persistent connectivity to traders regardless of their location, a capability that legacy hardware-based systems can only imitate awkwardly.

Hardware

Crucially, being a SaaS solution allows us to support various devices - desktop application, mobile and even traditional hardware where required. This flexibility ensures that firms can adopt the system without being locked into a single proprietary device, future-proofing their investment against evolving workplace needs. location, a capability that legacy hardware-based systems can only imitate awkwardly.

Reliability & resilience

We plan to utilize AWS/GCP infrastructure to reach geographies that would be harder to be in if you run your own private cloud. Plus being distributed to more locations means lower latency and less exposure to failures (higher resilience).

AI: integrating intelligence

While the voice system's core architecture must ensure performance and resilience, the next layer of innovation is Artificial Intelligence. This application, particularly in voice transcription and analytics, is best delivered via a SaaS model.

Moving from legacy on-premise recording to a SaaS model provides two crucial advantages:

- **Continuous improvement:** Being cloud-native allows us to continuously collect the data necessary to fine-tune and improve our models, as well as fundamentals required for data capture like audio quality and noise canceling (especially important as mobile usage grows). This feedback loop is essential for enhancing Trader Voice Analytics, improving the accuracy of transcription for financial jargon and overlapping conversations and building out new features.
- **Compliance and insight:** Our AI capabilities transform unstructured voice data into actionable intelligence. This includes high-fidelity, individual call transcription and the ability to extract key insights for trade reconstruction, compliance surveillance and risk management. This data layer is decoupled from the core voice architecture, allowing us to rapidly iterate on AI capabilities without impacting voice performance.



Agentic Workflows: Automating Operations and Provisioning

The next frontier of operational efficiency lies in shifting from manual administration to agentic workflows. While the possibilities are endless, the initial use case we see is automating the full user provisioning lifecycle.

By leveraging Symphony's Agent Studio and Model Context Protocols (MCPs), this approach replaces manual configuration with agent-driven onboarding, reducing the risk of human error and significantly compressing the time-to-value for new deployments and environment setups.

Furthermore, we are extending this agentic capability into operations and compliance, transforming how we monitor and maintain communication quality. Agents are being deployed to proactively track telemetry, reconcile data across fragmented legacy silos and provide real-time assurance. Looking ahead, agent-to-agent workflows will operate within the same secure, auditable frameworks as human interactions. With built-in "human-in-the-loop" validation and granular access controls, these autonomous agents will be able to handle complex trade reconstruction, compliance surveillance, and risk management at scale, helping to ensure every automated action is as secure and compliant as our core voice architecture.



The next decade of voice trading innovation

The evolution of mission-critical voice is fundamentally an architectural challenge. It is the shift from managing the risk of proprietary hardware and limited failover to embracing the inherent resilience, elasticity and intelligence of the public cloud. Symphony Voice, anchored by Cloud9, provides the engineering blueprint for the next decade.

By architecting for multi-region redundancy, utilizing a P2P community model, integrating AI-driven intelligence and supporting multi-device endpoints, we are delivering a solution that is not just a compliant voice system, but a core piece of financial infrastructure designed for continuous operation and future innovation. For engineering and IT leaders, this means moving beyond the constraints of the past to implement a scalable, secure and truly resilient foundation for global trading.

The past 10 years were about revolutionizing the voice market - the next decade will be focused on refinement. First, our team is laser focused on ensuring unparalleled quality and constant refinement of our voice platform. This commitment extends beyond proactive service monitoring to perfecting the quality of archiving and data extraction. We leverage a continuous data feedback loop that fine-tunes models. This cloud-native approach enhances Trader Voice Analytics, dramatically improving transcription accuracy for financial jargon and overlapping conversations, thus transforming unstructured voice data into actionable intelligence for compliance surveillance, trade reconstruction, and risk management.

Second, we are establishing the core building blocks to support the next frontier: the integration of autonomous agents into the voice trading lifecycle. This future, which will include human-to-agent or even agent-to-agent workflows over Cloud9, necessitates a higher level of compliant data capture and harvesting than ever before. We are actively developing the architectural layers required to ensure these advanced, intelligent workflows are secure, auditable and seamlessly integrated.



