

PROFILE

LINKEDIN

AGENTIC AI SYSTEMS

DF COMMANDER — EDGE-TO-SERVER VOICE EMAIL ASSISTANT

TUF Laptop Speech Recognition, Bonsai 1.7B, AIGENESIS Qwen3-30B-A3B, Microsoft 365 and Text-to-Speech

Publication ID: 202611-AASPUB-0012
Publisher: Agentic AI Systems
Publication Date: 14 November 2026
Category: Artificial Intelligence | Voice Interfaces | Systems Engineering
Classification: Technical Research & Engineering Publication
Evidence Basis: Historical project handover dated 4 November 2026; current service status not revalidated

Agentic AI Systems — Intelligence. Automation. Advantage.

Abstract

This publication documents a distributed, locally orchestrated voice-assistant workflow designed to convert spoken requests into Microsoft 365 email operations and return spoken responses. The system uses an ASUS TUF laptop as the edge interface, a small Bonsai language model for local command handling, an AIGENESIS AI server running Qwen3-30B-A3B for higher-capacity reasoning, a Microsoft 365 email connector, and Piper text-to-speech (TTS).

Historical engineering notes report successful retrieval and spoken playback of individual email bodies. They also identify incomplete dashboard synchronisation and intermittent command-routing failures. These observations are distinguished from planned improvements and do not constitute a fresh production acceptance test.

1. Engineering Objective

The design aims to provide hands-free email assistance while keeping speech capture and final audio playback close to the user. The laptop performs the immediate voice interaction; the server handles more complex language interpretation and authorised email operations. The architecture is intended to reduce repeated manual navigation through the mailbox while maintaining explicit boundaries around account access and message content.



PROFILE

LINKEDIN

AGENTIC AI SYSTEMS

2. End-to-End Processing Architecture

Stage	System / responsibility
1 — Voice capture	TUF-EDYBOY microphone; local speech-to-text transcription
2 — Edge interpretation	Bonsai 1.7B; classify and route email-related intent
3 — Server reasoning	AIGENESIS; Qwen3-30B-A3B interprets supported requests
4 — Mail operation	Microsoft 365 connector executes authorised email retrieval
5 — Response	Text returned to edge dashboard and Piper Alan TTS

3. Edge Laptop: Speech and Local AI

The historical configuration identifies an ASUS TUF laptop with an NVIDIA GTX 1650 Ti and 32 GB RAM. Whisper Tiny provides local speech recognition. Bonsai 1.7B Q1_0 is described as the edge language model, with a configured context window of approximately 32K tokens and a local inference listener at 127.0.0.1:8084. These are recorded configuration values, not current health-check results.

The edge model is intended to recognise commands such as reading a selected message, requesting a brief, and asking for recent mail. Routing a recognised email command to a connector is preferable to allowing a general-purpose model to improvise an answer about mailbox contents.

4. AIGENESIS: Server-Side Reasoning

The AIGENESIS server is recorded as hosting Qwen3-30B-A3B with an approximately 16K context configuration. A historical API endpoint is 192.168.0.123:8114. This server provides higher-capacity intent processing and language generation, while a separate connector is responsible for Microsoft 365 operations. The model must not be treated as having independent mailbox access merely because it can generate email-related text.

5. Microsoft 365 Connector

Historical records identify a connector at 192.168.0.123:8120, with POST /voice/email as an email operation endpoint. Supported command families were described as list, brief, priorities, today and read. The connector is the authorised boundary for fetching mailbox data; model-generated text alone is not evidence that an email was retrieved.

Access tokens, account consent, Microsoft Graph permissions and audit controls must be enforced by the connector. The records available for this publication do not establish the exact Graph permission scopes or the current authentication state.



PROFILE

LINKEDIN

AGENTIC AI SYSTEMS

6. Text-to-Speech and Dashboard

The response path returns text for display in the laptop dashboard and for playback through Piper Alan TTS. The historical dashboard endpoint is 127.0.0.1:8098, with an event receiver at 127.0.0.1:8099/event. The intended user experience combines an audible answer with a visible record of the content or summary that was spoken.

The previously recorded defect was that some retrieved email bodies were spoken successfully but were not displayed in the dashboard. Audio playback success therefore cannot be used as a substitute for dashboard end-to-end validation.

7. Recorded Implementation Paths

Component	Historical path / endpoint
Laptop listener	C:\AI\DF-Voice-V2.9\voice_listener_bonsai_alan.py
M365 connector	/home/edyboy/aas-m365-agent/aas_m365_voice_connector.py
Bonsai local API	127.0.0.1:8084
Qwen3 server API	192.168.0.123:8114
M365 connector API	192.168.0.123:8120
Dashboard / event	127.0.0.1:8098 / 127.0.0.1:8099/event

8. Historical Validation and Outstanding Issues

Area	Recorded result / qualification
Individual email read	Email body retrieval and spoken TTS output reported successful
Number-based selection	Recognition and read_number routing repairs implemented
Dashboard synchronisation	Outstanding: spoken message text sometimes absent from dashboard
Temporal email requests	Outstanding: “read yesterday’s emails” could bypass connector
Fallback model response	Qwen could incorrectly state it lacked email access
Current availability	Not verified by this publication



PROFILE

LINKEDIN

AGENTIC AI SYSTEMS

The evidence supports a partial functional baseline, not a complete production-grade voice/email acceptance. No measured speech-recognition accuracy, end-to-end latency, token throughput or connector reliability statistics are available in the supplied historical summary.

9. Failure Modes and Engineering Controls

Routing integrity: Treat recognised email commands as structured operations. Validate intent, date ranges and message identifiers before dispatch. When connector invocation fails, report that failure rather than asking the language model to invent mailbox state.

Prompt-injection resistance: Email bodies are untrusted content. Instructions embedded inside retrieved messages must remain data and must not override the assistant’s tool policy, user authorisation or routing rules.

Privacy and credentials: Keep Graph credentials and access tokens out of model prompts, logs and spoken output. Use least-privilege delegated permissions, protect transport between nodes, and limit retention of transcriptions and message bodies.

Action confirmation: Reading and summarising mail are distinct from sending, deleting or forwarding. Any future write action should require explicit user confirmation and server-side permission checks.

Observability: Associate each voice request with a trace identifier covering transcription, intent classification, connector execution, dashboard update and TTS completion. Avoid logging raw message contents unless explicitly required and protected.

10. Recommended Acceptance Test Plan

A repeatable acceptance suite should verify: microphone transcription; local Bonsai classification; remote Qwen connectivity; connector authentication; read-by-number; date-scoped listing; dashboard/event delivery; TTS playback; and recovery from network or token failures. Every test should record request IDs, response codes, timing and whether the answer was grounded in an actual connector result.

Negative tests should include an email containing malicious instructions, ambiguous message numbers, revoked Graph consent, unavailable AIGENESIS, and a disconnected dashboard. The system should fail safely and communicate uncertainty without fabricated mailbox content.

11. Conclusion

DF Commander illustrates a practical division of labour between edge speech processing, small-model intent handling, server-side reasoning, a dedicated Microsoft 365 connector and local speech synthesis. The recorded work demonstrates that individual email retrieval and spoken playback were achievable within this architecture. Reliable dashboard propagation, deterministic connector routing and security acceptance remain necessary before a broader operational claim can be made.



PROFILE

LINKEDIN

AGENTIC AI SYSTEMS

Publication and Attribution

Agentic AI Systems
Intelligence. Automation. Advantage.
Website: https://www.agenticaisystem.com.au/
Technical Publications: https://www.agenticaisystem.com.au/publications/

Publication Reference: 202611-AASPUB-0012
Published: 14 November 2026
© 2026 Agentic AI Systems. All rights reserved.

Technical disclosure: This report is based on a historical configuration and handover summary, not a live audit of endpoints or Microsoft 365 permissions. Documented implementation outcomes are separated from proposed controls and future acceptance criteria.

Video Demonstration — DF Commander V2

https://www.agenticaisystem.com.au/videos/watch/df-commander-v2/

Watch the distributed local AI and Microsoft 365 voice-assistant demonstration.