



PROFILE



LINKEDIN



AGENTIC AI SYSTEMS

19-OCT-2026

Document Number: 202610-AASPUB-0001

Enterprise AI Compute Node: Local Qwen AI Integration with Microsoft 365 and the Future of Intelligent Business Automation

Applied Artificial Intelligence Engineering | Enterprise Digital Transformation | Autonomous Workflow Systems

Author: Reddy (Eddie) Jamisola

Organisation: Agentic AI Systems

Location: Darwin, Northern Territory, Australia

Publication: 19 October 2026

Website: <https://www.agenticaisystem.com.au/>

Intelligence. Automation. Advantage.

Abstract

Artificial intelligence is rapidly transforming how organisations communicate, process information, manage financial transactions and coordinate daily operations. Despite these advances, many businesses continue to rely on fragmented systems and labour-intensive administrative processes.

Agentic AI Systems has successfully developed and demonstrated an enterprise integration connecting a **locally hosted Qwen Large Language Model (LLM), operating through an AI Compute Node, with Microsoft 365 email services.**

The demonstrated system enables authenticated email retrieval, intelligent message analysis, structured summarisation and identification of actionable information using locally operated AI inference.

This achievement establishes a foundation for a broader enterprise automation platform designed to extend AI capabilities into financial document processing, accounts payable workflows, voice-controlled email management, calendar organisation, intelligent correspondence preparation and controlled business approvals.

Future development includes Optical Character Recognition (OCR), automated invoice and quotation identification, structured financial data extraction, CSV generation for accounting and finance systems, speech-to-text (STT), text-to-speech (TTS), AI-assisted email drafting, calendar scheduling and secure timesheet approval workflows using authorised PIN-based confirmation.

The initiative demonstrates how independently developed AI harnesses and locally operated computing infrastructure can complement established enterprise platforms while supporting greater flexibility, operational efficiency and human oversight.

This publication describes demonstrated capabilities and proposed development directions without disclosing proprietary software architecture, source code, model configurations, orchestration methods or internal security mechanisms.

Keywords: Agentic AI, Qwen, AI Compute Node, Microsoft 365, Enterprise AI, Intelligent Automation, OCR, Financial Processing, CSV Integration, Voice AI, STT, TTS, Calendar Automation, AI Harness, Human-in-the-Loop, Digital Transformation.

1. Introduction

Enterprise organisations rely on digital communication, financial applications, document management systems and collaboration platforms to conduct daily operations.

Although these technologies have improved business efficiency, many processes still require employees to manually interpret emails, review attachments, extract financial information, prepare correspondence, organise meetings and authorise routine transactions.

These activities are often repetitive, time-consuming and distributed across multiple software applications.

Artificial intelligence offers opportunities to reduce these administrative burdens by introducing intelligent assistance directly into existing enterprise workflows.

However, many AI solutions operate as standalone services or require organisations to depend on externally hosted inference platforms.

Agentic AI Systems has developed an alternative approach involving an independently managed AI Compute Node capable of supporting locally hosted language models and enterprise application integration.

The initial demonstrated capability connects a Qwen language model with Microsoft 365 email services.

The longer-term objective is to develop a unified enterprise AI assistant capable of interpreting information, preparing actions and coordinating business workflows while maintaining appropriate human authorisation.

2. AI Compute Node – Enterprise AI Infrastructure

An AI Compute Node is a dedicated computing environment designed to execute artificial intelligence workloads independently of conventional end-user applications.

Depending on operational requirements, the computing environment may incorporate CPUs, GPUs or combinations of hardware resources.

Agentic AI Systems has undertaken engineering involving locally hosted AI models, NVIDIA and AMD GPU infrastructure, distributed inference environments and custom AI orchestration tools.

The AI Compute Node provides a foundation for operating AI processing services under organisational control.

Potential advantages include:

- Local execution of selected AI workloads.
- Greater flexibility in model selection.
- Reduced exclusive dependence on commercial inference providers.
- Integration with existing cloud and enterprise applications.
- Support for specialised business automation.
- Adaptability to organisational infrastructure requirements.
- Opportunities to optimise computing resource utilisation.

The AI Compute Node concept supports hybrid enterprise environments in which locally operated AI services interact with established platforms such as Microsoft 365.

This does not remove the need for cloud security, identity management, data governance or appropriate information-handling controls.

3. Demonstrated Achievement – Qwen and Microsoft 365 Email

Agentic AI Systems has successfully implemented and tested an integration between a locally hosted Qwen language model and Microsoft 365 email services.

The demonstrated system retrieves authorised email information and uses local AI inference to analyse correspondence.

3.1 Current demonstrated capabilities

- Authenticated Microsoft 365 email access.
- Email message retrieval.
- Sender, subject and date identification.
- Individual message analysis.
- AI-generated email summaries.
- Identification of key information.
- Recognition of potential follow-up actions.
- Structured presentation of results.
- Integration through independently developed AI software.

The demonstrated implementation establishes functional interoperability between enterprise cloud communication services and locally hosted AI processing.

This capability provides the starting point for further development of intelligent enterprise workflows.

4. Future Development – AI Financial Email and Document Processing

One of the proposed extensions is an intelligent financial correspondence assistant.

Finance departments routinely receive invoices, quotations, purchase documentation, statements and other financial records through email.

These documents may arrive in different formats and require manual review before information can be entered into accounting systems.

The proposed AI financial processing capability aims to reduce repetitive data entry while preserving financial validation and approval requirements.

4.1 Automated identification of financial correspondence

The AI assistant would examine authorised incoming correspondence and identify messages likely to contain financial documents.

Potential document categories include invoices, supplier quotations, credit notes, purchase orders and statements.

The system would assist employees by distinguishing relevant financial correspondence from ordinary business messages.

4.2 OCR and intelligent data extraction

Where information is contained within PDF documents, scanned images or other supported attachments, OCR and AI-assisted document interpretation could be used to extract relevant fields.

Potential information includes:

- Supplier name and ABN.
- Invoice or quotation number.
- Document date and payment due date.

- Purchase order reference.
- Line-item descriptions, quantities and unit prices.
- GST and total amounts.
- Cost centre or account coding, where supported by approved business rules.

The objective is to convert unstructured financial documents into consistent, reviewable information.

4.3 Finance-ready CSV generation

Following extraction and validation, approved financial information could be converted into structured CSV files compatible with an organisation's accounting or finance import requirements.

<box> ****Illustrative workflow****

Financial email received → Document identified → OCR and AI extraction → Data validation → Human review → CSV export → Finance system import

</box>

This would support potential integration with enterprise finance applications, including systems such as TechnologyOne and other accounting platforms, subject to their supported import formats and validation requirements.

A CSV file should be treated as an import-ready candidate only after required checks are completed. AI extraction alone does not establish that an invoice is genuine, correctly coded or authorised for payment.

4.4 Expected operational benefits

The proposed capability could reduce manual transcription, improve consistency, accelerate document preparation and assist finance teams in managing high-volume correspondence.

Actual time savings and extraction accuracy would require testing against representative financial documents.

5. Future Development – Voice-Controlled Enterprise AI Assistant

Another proposed capability is natural-language voice interaction with enterprise applications.

The objective is to allow authorised employees to request routine administrative assistance through spoken instructions.

The proposed system would combine speech-to-text, language understanding and text-to-speech technologies.

5.1 Speech-to-text (STT)

Speech-to-text technology would convert spoken instructions into text for interpretation by the AI assistant.

Example commands include:

"Read my important emails."

"Summarise the messages received this morning."

"Write an email to the finance manager."

"Check my calendar for tomorrow."

"Prepare a meeting invitation for Friday."

The assistant would interpret the request and prepare the appropriate information or proposed action.

5.2 Text-to-speech (TTS)

Text-to-speech technology would allow the assistant to communicate information verbally.

Potential applications include reading email summaries, presenting calendar schedules, explaining proposed actions and requesting confirmation.

5.3 Voice-assisted email preparation

A user could instruct the assistant:

"Write an email to the project manager advising that the report will be ready tomorrow."

The AI assistant would prepare a draft using the instruction and any relevant authorised context.

The draft would then be presented for review.

The system would not send the message merely because the user requested that it be written.

Sending would require a separate, authorised execution action.

This separation between preparation and execution is a central design objective.

6. Future Development – Intelligent Email Analysis and Prewritten Responses

The proposed email assistant would extend beyond summarisation into the preparation of suggested correspondence.

Many employees spend considerable time reading messages and composing routine replies.

An AI-assisted workflow could analyse incoming correspondence, identify its purpose and prepare an appropriate response for employee review.

6.1 Proposed capabilities

- Identify messages requiring responses.
- Classify correspondence by purpose and urgency.
- Extract questions and requested actions.
- Prepare contextual draft replies.
- Highlight missing or uncertain information.
- Present recipient, subject and proposed message.
- Maintain drafts pending authorised user action.
- Record approved sending actions where required.

6.2 Human-controlled execution

A central principle is that **AI may prepare an action without automatically receiving authority to execute it.**

For example, the system may identify a supplier enquiry and prepare a response.

The employee would review the proposed message, make any required changes and explicitly approve sending.

<box> ****Illustrative workflow****

Email received → AI analysis → Suggested reply generated → Draft awaiting review → User approves or edits → Authorised email transmission

This approach is intended to balance automation efficiency with accountability and user control.

7. Future Development – AI Calendar Organiser

Calendar management is another proposed enterprise integration capability.

Employees frequently coordinate meetings, appointments, deadlines and follow-up activities through email.

The AI calendar organiser would assist with identifying scheduling requirements and preparing calendar actions.

7.1 Proposed capabilities

- Read authorised calendar information.
- Summarise upcoming appointments.
- Identify potential scheduling conflicts.
- Suggest available meeting periods.
- Prepare appointment invitations.
- Extract meeting requests from correspondence.
- Organise reminders and follow-up activities.
- Prepare rescheduling recommendations.
- Coordinate email and calendar information.

7.2 Natural-language scheduling

A user could request:

"Check my calendar and arrange a meeting with the IT team next Wednesday afternoon."

The assistant would assess authorised availability information and prepare a proposed appointment.

The user would then review and approve the meeting details before invitations were issued.

The intention is to reduce administrative coordination while preserving control over calendar changes.

8. Future Development – Intelligent Timesheet Approval

Timesheet processing is a common administrative responsibility across government, councils, businesses and service organisations.

Traditional workflows may involve employees submitting timesheets, supervisors reviewing entries and finance or payroll teams processing approved records.

The proposed AI-assisted timesheet capability would support the review and approval process.

8.1 Timesheet analysis

The system could assist authorised supervisors by reviewing submitted timesheet information against defined organisational rules.

Potential functions include identifying incomplete records, unusual hours, missing approvals, duplicate entries and discrepancies requiring attention.

The AI would present findings and recommendations rather than independently determining payroll entitlements.

8.2 Approval PIN mechanism

A proposed feature is an approval process requiring an authorised supervisor to confirm selected actions using an approval PIN.

<box> ****Illustrative workflow****

Timesheet submitted → AI-assisted validation → Supervisor review → Approval PIN confirmation → Authorised approval recorded → Payroll processing
</box>

The PIN would provide an additional confirmation step within an authenticated approval process.

For enterprise deployment, PIN confirmation should be combined with appropriate identity verification, role-based permissions, secure credential handling and auditable approval records.

A PIN alone is not sufficient protection for sensitive payroll transactions.

8.3 Human accountability

The proposed design preserves the supervisor's responsibility for approving timesheets.

AI assistance would focus on highlighting exceptions, improving consistency and reducing repetitive administrative checking.

Final authorisation would remain with the designated employee.

9. Unified Enterprise AI Assistant

The longer-term development direction is to bring these capabilities together within a coordinated enterprise AI environment.

Rather than operating separate AI tools for email, finance, scheduling and approvals, organisations could use a common AI assistance framework.

The proposed environment would allow authorised users to interact with multiple enterprise functions through a consistent interface.

<box gap={2}> ****Proposed enterprise AI capabilities****

<grid columns={2} gap={2}> <grid-item> <box border radius="lg" padding={3} gap={1}> <icon name="mail" color="secondary" size="lg"/> ****Email Assistant**** <text color="secondary" size="xs">Analyse, summarise and prepare replies</text> </box> </grid-item> <grid-item> <box border radius="lg" padding={3} gap={1}> <icon name="file-spreadsheet" color="secondary" size="lg"/> ****Finance Assistant**** <text color="secondary" size="xs">OCR, invoices, quotations and CSV exports</text> </box> </grid-item> <grid-item> <box border radius="lg" padding={3} gap={1}> <icon name="mic" color="secondary" size="lg"/> ****Voice Assistant**** <text color="secondary" size="xs">STT, TTS and spoken commands</text> </box> </grid-item> <grid-item> <box border radius="lg" padding={3} gap={1}> <icon name="calendar-days" color="secondary" size="lg"/> ****Calendar Assistant**** <text color="secondary" size="xs">Scheduling, reminders and appointments</text> </box> </grid-item> <grid-item> <box border radius="lg" padding={3} gap={1}> <icon name="clipboard-check" color="secondary" size="lg"/> ****Approval Assistant**** <text color="secondary" size="xs">Timesheets, validation and PIN confirmation</text> </box> </grid-item> <grid-item> <box border radius="lg" padding={3} gap={1}> <icon name="shield-check" color="secondary" size="lg"/> ****Governance**** <text color="secondary" size="xs">Access control, auditing and human oversight</text> </box> </grid-item> </grid> </box>

The objective is to create an intelligent enterprise assistant that can interpret requests, prepare information and coordinate proposed actions while respecting organisational permissions and approval requirements.

10. Custom AI Harnesses and Multi-Agent Engineering

Agentic AI Systems independently develops AI harnesses and specialised software tools to support its broader engineering programme.

An AI harness provides the surrounding software environment required to coordinate language models, business information and operational activities.

The development programme explores multi-agent orchestration, specialised AI functions, persistent operational context, automated quality assurance and intelligent task coordination.

These concepts support the potential development of enterprise assistants that can handle different categories of work without relying on a single model to perform every function.

For example, document interpretation, language analysis, information validation and workflow preparation may require different capabilities.

The broader engineering objective is to coordinate these functions within a controlled operational environment.

Specific orchestration mechanisms, software architecture and proprietary implementation methods remain confidential.

11. Cybersecurity, Privacy and Responsible AI

Enterprise AI integration introduces security and governance requirements that must be addressed before production deployment.

These requirements become especially important when AI systems interact with email, financial records, employee information, payroll processes and business approvals.

The proposed enterprise AI environment is intended to support the following principles.

Least-privilege access: AI services should access only the information and functions required for authorised tasks.

Human-controlled execution: High-impact actions, including sending correspondence, approving timesheets and importing financial transactions, should require appropriate authorisation.

Data protection: Sensitive information must be handled according to organisational security and privacy requirements.

Auditability: Significant AI-assisted actions and approvals should be recorded to support accountability.

Output validation: AI-generated information must be checked before being relied upon for financial, legal or operational decisions.

Secure approval mechanisms: Authentication, authorisation and confirmation controls must protect business transactions.

Operational reliability: Systems should support monitoring, exception handling and recovery from processing failures.

Locally hosted AI offers greater control over selected processing resources but does not automatically guarantee compliance or eliminate cybersecurity risks.

12. Enterprise and Government Relevance

The proposed capabilities have potential applications across government agencies, regional councils, commercial enterprises and other organisations that rely on Microsoft 365 and structured business processes.

Finance departments could benefit from assisted invoice extraction and preparation of finance-ready data.

Administrative teams could use intelligent email summaries and suggested responses.

Executives could receive consolidated correspondence and scheduling assistance.

Supervisors could use AI-supported timesheet validation while retaining final approval authority.

Employees could interact with authorised enterprise functions through natural-language voice commands.

These applications may be particularly relevant to organisations managing distributed teams, regional offices and high volumes of administrative correspondence.

However, each deployment would require assessment of operational needs, data quality, integration requirements, accessibility, security and governance.

13. Demonstrated Capabilities and Development Roadmap

The distinction between implemented functionality and proposed extensions is essential.

Capability	Development status
Local Qwen inference through AI Compute Node	Demonstrated
Microsoft 365 email access	Demonstrated
Individual email analysis	Demonstrated
AI-generated email summaries	Demonstrated
Identification of potential email actions	Demonstrated
Custom AI harness engineering	Existing development programme
Automated invoice and quotation detection	Proposed
OCR-based financial extraction	Proposed integration
Finance-ready CSV generation	Proposed integration
Voice-controlled enterprise commands	Proposed
STT and TTS enterprise assistant	Proposed integration
AI-prepared email responses	Proposed extension
Approval-controlled email sending	Proposed
AI calendar organiser	Proposed
AI-assisted timesheet validation	Proposed
PIN-confirmed timesheet approvals	Proposed
Unified enterprise AI assistant	Long-term development objective

The roadmap is subject to engineering validation, security testing, enterprise requirements and operational priorities.

14. Business Value and Operational Benefits

The proposed enterprise AI capabilities are intended to address several common business challenges.

Administrative productivity: Reduce time spent on repetitive correspondence and information-processing tasks.

Financial efficiency: Assist with preparing structured invoice and quotation information for accounting workflows.

Communication efficiency: Reduce the effort required to interpret and prepare routine email correspondence.

Scheduling efficiency: Simplify calendar coordination and appointment preparation.

Management oversight: Provide structured assistance with timesheet review and approval processes.

Technology flexibility: Allow organisations to evaluate locally hosted AI alongside existing enterprise cloud investments.

Operational control: Preserve human authorisation for actions involving financial records, external communications and employee approvals.

These benefits represent intended outcomes. Quantitative improvements would need to be established through controlled testing and operational measurement.

15. Intellectual Property and Proprietary Technology

The custom AI harnesses, integration software and associated engineering components developed by Agentic AI Systems form part of its proprietary technology portfolio.

This publication intentionally provides a descriptive overview rather than an implementation guide.

Source code, detailed architecture, orchestration logic, authentication configurations, model-serving parameters, internal communication mechanisms and security implementation details are not disclosed.

The publication is intended to communicate the demonstrated achievement, engineering direction and potential enterprise applications while protecting proprietary development work.

Third-party models, applications and technologies remain subject to their respective licensing and ownership arrangements.

16. Conclusion

Agentic AI Systems has successfully demonstrated the integration of a locally hosted Qwen language model with Microsoft 365 email services through an independently operated AI Compute Node.

The demonstrated system provides AI-assisted email analysis, structured summarisation and identification of relevant information using locally hosted inference.

This achievement establishes a foundation for a broader enterprise AI assistant capable of supporting financial document processing, OCR-based invoice extraction, CSV preparation, voice interaction, intelligent email drafting, calendar organisation and approval-controlled business workflows.

The proposed development direction places particular emphasis on maintaining human control over important actions, including external email transmission, financial data processing and timesheet approvals.

The broader significance is the potential to transform enterprise AI from a standalone conversational tool into a coordinated operational assistant integrated with existing business systems.

The objective is not simply to make AI capable of answering questions, but to develop AI systems that can interpret information, prepare useful work, coordinate business processes and support authorised human decision-making.

Through independent AI harness development, locally operated computing infrastructure and practical enterprise integration, Agentic AI Systems continues to explore the next generation of intelligent business automation.

Author and Organisation

Reddy (Eddie) Jamisola

Founder | AI Systems Engineer

Agentic AI Systems

Darwin, Northern Territory, Australia

Website: <https://www.agenticaisystem.com.au/>

Intelligence. Automation. Advantage.

© 2026 Agentic AI Systems. All rights reserved.

Publication Classification: Applied AI Engineering and Enterprise Digital Transformation White Paper.

This publication distinguishes demonstrated capabilities from future development objectives. Proprietary implementation details are intentionally withheld.