Agntic.ai

A PRACTICAL GUIDE
FOR AUSTRALIAN MSPS

The Managed Service Provider Guide to Agentic AI

From a service desk that drowns in tickets and after-hours calls to a digital teammate that answers every one, triages it cleanly, and knows exactly which action it must never take on its own.

PUBLISHED BY

Agntic.ai
AI voice agents & digital workers

EDITION

2026 · Australia
General information only

FIND US

Agntic.ai
Book a 20-minute walkthrough

FOREWORD

The pressure has moved to the service desk.

In the year to June 2025, the Australian Signals Directorate (ASD) received more than eighty-four thousand seven hundred cybercrime reports, around one every six minutes, and the average self-reported cost to a small business climbed fourteen per cent to $56,600.[1] Behind almost every one of those numbers is a phone call or a ticket, and for hundreds of thousands of Australian businesses the first place that call lands is a managed service provider (MSP).

Most MSPs are not short of demand. They are short of capacity at the one point where demand arrives: the service desk. Calls ring out at the morning peak. After-hours alerts pile up behind a voicemail or an outsourced answering line. Tickets sit unlogged while a senior engineer is heads-down on a project. The service level agreement (SLA) clock keeps running the whole time, and none of it shows cleanly on a profit and loss statement, which is exactly why it goes unaddressed for years.

This guide is about a specific, practical answer to that problem: an agentic artificial intelligence (AI) voice agent that works the phones, the inbound tickets and the routine triage around them, so your engineers can do the work that actually needs their hands. It is written for MSP owners, service delivery managers and service desk leads who want to understand what this technology does, what it must never do, and how to put it to work without becoming the soft spot in their own clients' security.

We have tried to be honest throughout. There is a clear line we will keep coming back to, the line between handling a request and acting on it, because the service desk is now a favourite target for attackers, and a good deal of this guide is spent making sure the technology stays firmly on the right side of it.

Brad Riley

CEO, Agntic.ai

1. Australian Signals Directorate (ASD), Annual Cyber Threat Report 2024–25 (released October 2025). ASD received over 84,700 cybercrime reports in FY2024–25, an average of one every six minutes, with the average self-reported cost of cybercrime per report for small business rising 14% to $56,600.

WHAT IS INSIDE

Contents

SECTION ONE

Agentic AI, in plain language

Before the benefits, the basics. What a digital teammate is, why it is different from the chatbots you have already met, and the single boundary that makes it safe to put on a service desk.

01

THE IDEA

A teammate, not a chatbot.

You have used generative AI already. You type a question, it writes you an answer, and the conversation ends there. It is a clever tool, but it waits for you and does nothing on its own.

Agentic AI is the next step. An **agent** does not just talk. It is given a goal, a set of rules and access to the tools it needs, and it carries the task through from start to finish. For an MSP the goal is usually simple to state: answer the call, work out what the user actually needs, and either resolve the routine request or log and triage it cleanly into your professional services automation (PSA) system, the same way a capable level-one dispatcher would.

That is why we call it a **digital teammate** rather than a chatbot. It speaks naturally on the phone, it listens, it asks the follow-up questions your desk would ask, and then it acts: it captures the detail, classifies the ticket, sets the priority against your SLA, and routes it to the right queue or engineer. When something falls outside its rules, it hands over to a person.

A chatbot answers. A digital teammate finishes the routine job, then knows exactly when to step back.

The difference matters most at the morning peak and after hours. A chatbot on your website might capture a message. A digital service desk agent actually picks up the phone at 7am, through the lunch crush and at 11pm, holds a real conversation, verifies who it is speaking to, and leaves a properly triaged ticket in the queue by the time your engineer logs in.

None of this replaces your team. It removes the repetitive, interruptive work that stops your engineers from doing the parts of the job that genuinely need a person: the diagnosis, the project work, the difficult client conversation, the judgement call.

THE ONE RULE

What it must never do.

A digital teammate on an MSP service desk is built around a boundary that matters more here than almost anywhere else. It handles the conversation and the administration around a request. It does not, on its own authority, take a privileged action that could open a door into a client's environment.

That means it does not reset a password or a multi-factor authentication (MFA) token on the strength of a phone call, it does not grant access, it does not disable security controls, and it does not start a remote session because someone on the line asked it to. If a caller requests any of those things, the agent's job is not to oblige. Its job is to verify identity to your policy, then route the request to an authorised human to action.

THE LINE, IN ONE SENTENCE

The agent handles requests and triage. A privileged action always passes through a verified identity and an authorised person. If a caller pushes for an immediate reset or access change, the agent is built to follow your verification steps, decline to act outside them, and escalate. It never lets urgency or pressure substitute for proof.

This is not a limitation we apologise for. It is the design, and in this sector it is the whole point. An MSP that adopts this technology should be able to say, hand on heart, that no credential was ever reset and no access was ever granted by a machine acting alone. Everything in the rest of this guide is built on top of that promise, and Section Five sets out exactly how it is enforced.

SECTION TWO

A shift on the service desk

The cost of an overloaded desk does not appear on any report. It shows up as a breached SLA, a senior engineer pulled off project work, and a client who quietly started shopping for a new provider. Here is the shift as your dispatcher actually lives it.

02

DANI'S MONDAY · A BUSY MANAGED SERVICE PROVIDER

Two queues, one phone, the clock always running.

Dani coordinates the service desk at a twelve-engineer MSP. Nothing here is unusual. That is the point.

MORNING

8:05	**The Monday flood.** Forty tickets banked overnight, the phone already ringing. Dani triages what lands while two engineers ask which job is theirs.
8:40	Three calls ring out during the rush. One was a director locked out before a board meeting. None leaves a voicemail; the SLA clock started anyway.
10:15	**A reset request.** A caller wants a password reset "right now, I'm travelling". Dani is sure it is fine, but the identity check takes eight careful minutes, and the queue grows while she does it.
12:30	Lunch. One person covers the desk. Nine calls come in over the hour. Four are answered. The rest become afternoon tickets, logged late and thin on detail.

AFTERNOON

2:00	**Project work, interrupted.** A senior engineer is pulled off a planned migration to take overflow calls. The migration slips to next week, again.
4:30	Two SLAs are at risk because the tickets were never properly categorised this morning, so nobody saw the priority until now.
6:05	Phones roll to the after-hours answering service. From now until 8am, a third party takes messages it cannot action. The serious ones wait until morning.

THE INVISIBLE COST

None of it was anyone's fault.

Dani is good at her job. The MSP is well run. And yet by the end of the day two SLAs were breached, a migration slipped, a senior engineer spent an afternoon on level-one calls, and a handful of tickets were logged so late and so thin that someone will have to ring the client back just to understand them.

This is the trap of service desk work. The losses are real but invisible. A missed call is not recorded as a missed call; it is simply a call that never reaches a person. A poorly triaged ticket does not announce itself; it just surfaces as a breach hours later. The senior engineer pulled onto the phones does not file a complaint; the project simply moves, and the client wonders why the roadmap keeps slipping.

Because nobody can see the cost, nobody can justify hiring against it, especially when level-one staff are hard to find and harder to keep. The work that gets dropped is always the same work: the calls at the edges of the day, the careful logging, the patient identity check. The work, in other words, that a digital teammate is built to pick up.

The rest of this guide is about handing that specific layer of work to an agent, so the next Monday looks different: every call answered, every ticket logged and triaged on contact, identity verified the same way every time, and your engineers left on the work only they can do.

SECTION THREE

The five jobs a digital service desk agent does best

Not everything should be automated, and a good deal of service delivery never will be. These five jobs are where an AI voice agent is genuinely strong, and where MSPs see the change first.

03

WHERE IT EARNS ITS PLACE

Five jobs, done properly, every time.

JOB 01

ANSWERED

Answering the phone and the overflow

Every call picked up on the first ring, at every hour, including the Monday peak, the lunch crush and after close. No hold queue, no voicemail, no caller left to an answering service that cannot help them.

JOB 02

LOGGED

Logging, classifying and triaging tickets

The agent captures the full detail on the first call, sets the category and priority against your SLA, writes it straight into your PSA, and routes it to the right queue. No thin tickets, no late triage, no missed clock.

JOB 03

VERIFIED

Handling reset and access requests, safely

It runs your identity verification steps the same way every time, then routes a verified request to an authorised engineer to action. It never performs the privileged step itself, which is exactly what keeps the desk secure.

JOB 04

RESOLVED

Known-fix and status enquiries

Routine, documented level-one fixes and the endless "any update on my ticket?" calls, answered from your own runbooks and live ticket status, with anything unusual passed to a person.

JOB 05

COVERED

After-hours triage and on-call escalation

Overnight and on weekends, the agent answers, verifies, captures the detail and decides against your rules what waits for morning and what wakes the on-call engineer. Your people are roused for genuine priority-one events, not for a forgotten password. Every interaction is logged and consent to record is captured up front.

WHY THESE FIVE

Rising demand, rising stakes, a workforce you cannot simply hire.

The jobs worth handing over share a shape. They happen often, they follow rules you already have, and getting them wrong costs you a client or opens a door. That is precisely the shape an agent handles well, and where a scarce, expensive engineer is wasted.

1 every 6 min

A cybercrime report is made to the ASD on average every six minutes, more than 84,700 in the year, and an MSP desk is often the first call.

ASD ANNUAL CYBER THREAT REPORT 2024-25

$56,600

Average self-reported cost of cybercrime per report for a small business, up 14% on the year. Your clients feel this, and they call you first.

ASD ANNUAL CYBER THREAT REPORT 2024-25

1.3m by 2030

Tech workers Australia will need by 2030, against roughly 950,000 today. You cannot simply hire your way out of a busy desk.

ACS AUSTRALIA'S DIGITAL PULSE 2025

Read those numbers together and the case is straightforward. Demand on the desk is climbing, the cost of getting it wrong is climbing with it, and the workforce to absorb the load is not there to be hired. Hand the repetitive front-line moments to an agent and you keep your engineers for the work that needs them, while every call still gets answered.

Sources: ASD, Annual Cyber Threat Report 2024–25 (October 2025): over 84,700 cybercrime reports, an average of one every six minutes; average self-reported small-business cost per report $56,600, up 14%. Australian Computer Society (ACS), Australia's Digital Pulse 2025: around 1.3 million technology workers needed by 2030, against an estimated 950,000 in 2025.

SECTION FOUR

What an unanswered desk really costs

A missed call feels like nothing. A senior engineer on the phones feels like Tuesday. Stacked up over a year, they are a different story. Here are the four quiet leaks, and what they add up to.

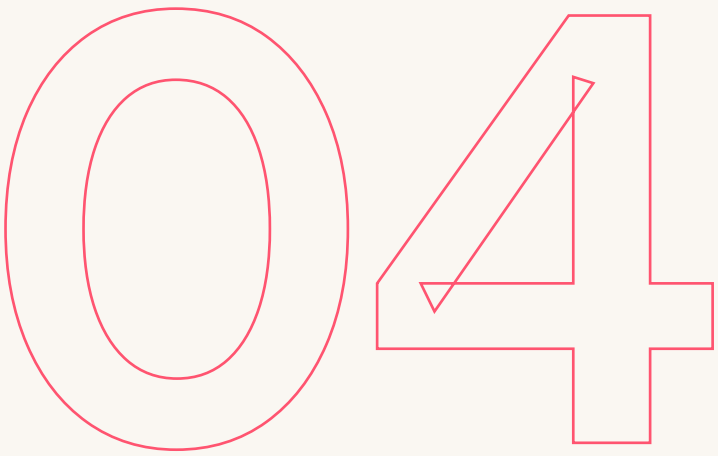

THE LEAKS LEDGER

Four leaks, one total.

The weekly figures below are illustrative drivers for a twelve-engineer MSP, not a quote. Every MSP's numbers differ. The value of laying them out is that the leaks stop being invisible.

Senior engineers on level-one calls ~30 hours a week of logging, triage and overflow soaked up by engineers who should be on project work, at a $55 fully-loaded rate.	$1,650 /wk
After-hours cover and SLA credits An outsourced answering service that cannot action anything, plus the service credits that slow nights and breached targets trigger.	$650 /wk
Slow first response, at-risk renewals One mid-size contract drifting toward a competitor over a year of missed calls, amortised across the weeks.	$750 /wk
Overflow calls lost and the rework Calls that ring out at peak, then come back as thin tickets that need a call-back just to understand.	$300 /wk
The combined leak	**$3,350 /wk**

That is roughly $160,000 a year leaking quietly through four holes, none of which shows up on an invoice. You do not need these exact numbers to act; even at half the assumptions the annual cost dwarfs the price of closing it. A digital teammate addresses all four at once: it answers the calls that ring out, it logs and triages on contact, it covers the nights, and it gives your engineers their project time back. Section Eight builds the full model with your own numbers.

Illustrative figures for a twelve-engineer MSP, shown to make the leaks visible. They are drivers, not a quote, and every MSP's volumes and rates differ.

SECTION FIVE

Identity, security and the line you never cross

This is the section that matters most for an MSP and the one most guides skip. Why the service desk is now a target, how the caller is verified, how client data is protected, and exactly how the agent is kept clear of any privileged action.

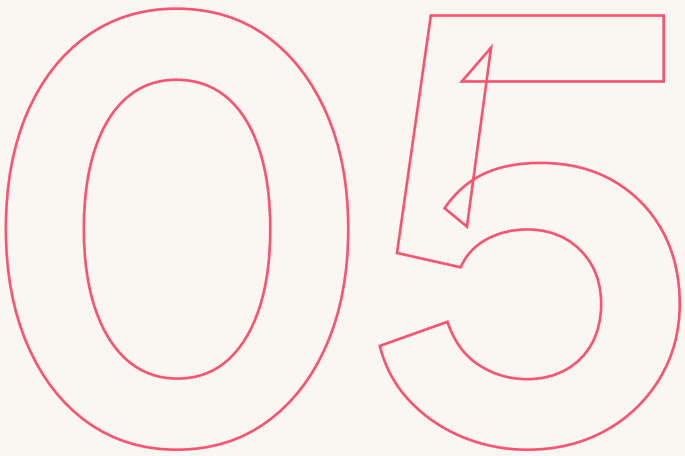

THE DESK IS NOW A TARGET

Why attackers call the help desk.

An MSP holds the keys to many businesses at once. That is its value to clients, and exactly why it is valuable to attackers. The cyber authorities of Australia, the United Kingdom, the United States, Canada and New Zealand have jointly warned that malicious actors increasingly target MSPs to exploit the trusted connection between a provider and its customers, where one compromise can cascade across an entire client base.[2]

The favourite way in is no longer a clever exploit. It is a phone call. In the well-documented MGM Resorts incident in 2023, attackers from the group known as Scattered Spider simply rang the help desk, impersonated an employee, and talked staff into resetting multi-factor authentication, a foothold that led to ransomware and reported losses exceeding one hundred million United States dollars. The same playbook, social engineering of the service desk to reset credentials or MFA, ran against major retailers through 2025.[3]

This changes what a service desk agent, human or digital, is for. The pressure to be helpful, fast and accommodating is precisely the pressure an attacker exploits. The defence is not rudeness. It is a verification process that runs the same way every time, that no amount of urgency or seniority on the line can shortcut, and a hard rule that the privileged action itself waits for an authorised person.

The most dangerous request on a service desk is the urgent one that asks you to skip a step.

A digital teammate is well suited to exactly this discipline. It does not get flustered, it is not socially pressured, it does not make an exception because the caller is angry or claims to be the chief executive. It runs your identity checks faithfully, every call, and it is built so that it cannot reset a credential or grant access on its own. The consistency that is hard for a tired human at 2am is, for the agent, simply how it works.

2. Joint advisory, "Protecting Against Cyber Threats to Managed Service Providers and their Customers" (AA22-131A), authored by CISA, NSA, FBI, ACSC (Australia), NCSC-UK, CCCS and NCSC-NZ, 2022. 3. CISA/FBI advisory on Scattered Spider (AA23-320A, updated 2025); MGM Resorts 2023 incident and 2025 retail campaigns widely reported.

THE BOUNDARY, ENFORCED

How the line is held.

Section One set the rule: the agent handles requests and triage, never a privileged action on its own authority. This is how that rule is enforced in the way the agent is actually built.

THE AGENT WILL	THE AGENT WILL NOT
+ Run your identity verification steps the same way on every call	– Reset a password or MFA token on its own authority
+ Log, classify and triage tickets against your SLA rules	– Grant access, change permissions or disable a security control
+ Answer documented known-fix and ticket-status questions	– Start a remote session because a caller asked it to
+ Route a verified reset or access request to an authorised engineer to action	– Skip a verification step for urgency, seniority or pressure
+ Decline politely and escalate the moment a caller pushes past the process	– Make a security judgement that belongs to your team

HUMAN IN THE LOOP

A person is never removed from the privileged action. The agent is a verified, well-governed layer in front of the desk, with clear escalation paths to your engineers and a standing instruction to hand over the moment a request needs human authority. You set the rules; the agent keeps to them, every time; you can see everything it did.

CLIENT DATA AND GOVERNANCE

Built to be inspected.

An MSP handles personal information on behalf of its clients, and that brings obligations under the Privacy Act 1988 (Commonwealth) and the Australian Privacy Principles (APPs). A digital teammate has to be built to that standard, not retrofitted to it: callers are told plainly that they are speaking with an automated assistant, information is collected only for the purpose of the request, and recordings and transcripts are handled under your own privacy policy with access limited to the people who need it.

That framework has recently been strengthened. The Privacy and Other Legislation Amendment Act 2024 began the most significant overhaul of the Act in its history, and from 10 December 2026 organisations must disclose in their privacy policy where decisions affecting individuals are made by substantially automated means. An MSP using an AI agent should plan for that disclosure now rather than scramble for it later.

Because every interaction is logged, transcribed and summarised, you end up with a clearer record of front-line activity than most desks keep today, when a phone call leaves no trace at all. That record supports the obligations you already meet and the frameworks you may already align to, such as the Essential Eight or your own ISO 27001 controls. The agent operates inside those frameworks rather than around them.

Consistency becomes a setting rather than a hope, and every call can be reviewed.

It also makes review simple. You can listen back, read transcripts, and adjust the agent's rules and verification steps in plain language. If you decide a particular request must always escalate, that change is made once and applied to every call from then on. This guide is general information, not legal advice; before you go live, confirm your privacy and security obligations with the appropriate adviser. What the technology gives you is a desk that is easier to govern, not harder.

SECTION SIX

Under the bonnet

You do not need to be deeply technical to use this, but it helps to know what is happening when the phone rings. Here is the whole thing, in six parts.

06

HOW IT WORKS

Six parts, one conversation.

PART 01

It picks up and listens

The agent answers in a natural voice, on the first ring, and tells the caller they are speaking with your automated service desk assistant. It understands ordinary speech, accents and interruptions.

PART 02

It verifies who is calling

Before anything else, it runs your identity checks the same way every time. No verification, no privileged request proceeds. Pressure and urgency change nothing.

PART 03

It follows your rules

SLA priorities, ticket categories, known-fix runbooks, which client maps to which contract: it works from the rules you set, in plain language, not from guesses.

PART 04

It writes to your PSA

It connects to your professional services automation and ticketing tools through permissioned access, logs the full detail, sets the priority and routes to the right queue. No thin tickets, no double entry.

PART 05

It escalates cleanly

A privileged action, a priority-one event after hours, or any call that leaves its rules is handed to the right human, with the context already captured. It never actions the privileged step itself.

PART 06

It reports back

Every call is logged, transcribed and summarised. You see volumes, outcomes, verification results and anything escalated, and you tune the rules from there.

SECTION SEVEN

Before you switch it on

You can stand up a working agent quickly. The MSPs that get the most from it spend a little time first, getting the groundwork right. Here is what to map across two weeks.

07

THE GROUNDWORK

An hour of mapping saves a month of patching.

Getting an agent live is fast. The difference between a good launch and a frustrating one is whether you have written down the things your desk currently holds in its head. Work through this list before you go live.

- [ ] Your top ten reasons clients call, in order
- [ ] Your identity verification steps, written exactly
- [ ] The known-fix runbooks the agent may resolve from
- [ ] How clients and contacts map to contracts in your PSA
- [ ] Consent and recording wording for collecting information
- [ ] Ticket categories, priorities and SLA targets per client
- [ ] Which requests must always reach an authorised human
- [ ] Your after-hours rules: what waits, what wakes on-call
- [ ] Escalation paths and the engineers behind each queue
- [ ] Who owns the agent's rules and reviews its reports

This is a fortnight of light work, not a project. Most of it is writing down decisions you have already made informally. Once it is on paper, the agent can be configured to match exactly how your desk already runs, which is the whole point: it should sound and behave like your MSP, on its best day.

SECTION EIGHT

Building the business case

An illustrative model, not a promise. Plug in your own numbers and the shape of the return tends to hold: it pays for itself on returned engineer hours alone, before you count coverage and retention.

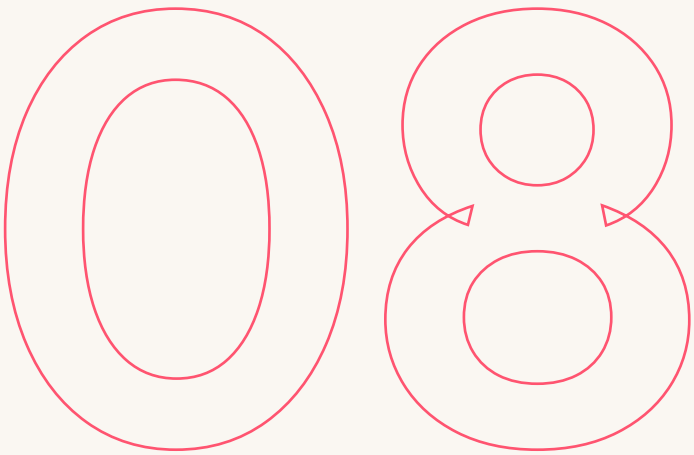

AN ILLUSTRATIVE MODEL

Where the return comes from.

A worked example to show the mechanism, not a quote. The numbers are illustrative and rounded; replace them with your own. The point is that the three returns stack, and the first one usually covers the cost by itself.

THE SCENARIO · A TWELVE-ENGINEER MSP

Supports about 40 client businesses and roughly 1,800 endpoints. The service desk fields around 500 calls and logs about 600 tickets a week. Around 40 calls a week currently ring out at peak or hit an after-hours answering service. Engineer time is costed at a $55 fully-loaded hourly rate.

Where the return comes from	Illustrative annual figure
Engineer hours returned ~30 hours/wk of first-line logging, triage and overflow handled by the agent × $55 × 48 weeks. The biggest line, and it frees project capacity you can bill.	**$79,000**
After-hours and overflow cover Replaces a ~$1,400/month outsourced answering service ($16,800) plus the ~$5,200 of SLA credits that slow nights trigger across a year.	**$22,000**
Retained client revenue Reliable, fast first response protects renewals; conservatively, one mid-size contract held that would otherwise have drifted to a competitor.	**$36,000**
Cost of the agent Indicative annual platform cost for a desk of this size, plus the groundwork time in your first fortnight.	**($18,000)**
Net illustrative return Returned hours and protected revenue, less the cost of running it.	**$119,000**

Read it conservatively and the case still holds. Halve every assumption and the model clears the cost of the agent several times over, on returned engineer hours alone, before the night cover and the steadier renewals are counted. We will build this with your real numbers in a short call rather than ask you to take a generic figure on faith.

SECTION NINE

What you don't need

Some of what holds MSPs back is not cost or risk, but a set of assumptions that are simply not true. Here are the ones worth retiring before you start.

09

MYTHS WORTH RETIRING

Less than you think.

YOU DON'T NEED

- To replace your engineers. The agent takes the repetitive layer, not the people
- To rip out your PSA or RMM. It connects to what you already run
- To weaken security to move faster. The opposite: it verifies the same way, every time
- A long integration project. A working agent is a matter of days, not quarters
- To let it touch privileged actions. By design, it never does

YOU DO NEED

+ A clear picture of why clients call, and your SLA rules
+ Your identity verification steps written down exactly
+ Agreement on what always reaches an authorised human
+ One owner inside the MSP who watches the reports
+ A willingness to start with one job and grow from there

The honest summary is that the barrier is smaller than the reputation of "AI on the service desk" suggests. You are not rebuilding your MSP. You are adding a reliable, well-governed layer to the front line and keeping a firm hand on what it is allowed to do, which in this sector is the entire value.

SECTION TEN

In practice

Three composite MSPs, drawn from the kinds of deployments this technology suits. The names are illustrative; the situations are not.

10

COMPOSITE CASE STUDIES

What changed, and how fast.

COMPOSITE · GENERAL IT

A twelve-engineer MSP

High call volume, Monday floods, senior engineers pulled onto the phones.

0

calls to voicemail after going live, day or night

The Monday flood, absorbed.

The desk's worst hour was the start of the week, when overnight tickets and a ringing phone collided and engineers were pulled off project work to cope. The agent now answers every call at the peak and after close, verifies the caller, and logs a fully triaged ticket straight into the PSA.

Within the first fortnight the team noticed project work stopped slipping, because the level-one interruptions that used to scatter the day were being handled at the front line.

COMPOSITE · SECURITY-FOCUSED

A managed security provider

Acutely aware of help-desk social engineering after industry incidents.

100%

of reset requests verified to policy before any action

The reset request, never rushed.

The provider's concern was not call volume but consistency: a tired human at 2am is exactly who an attacker calls. The agent now runs the identity checks the same way on every reset request, declines anything that fails them, and routes only verified requests to an authorised engineer to action.

COMPOSITE · GROWING MSP

A two-office regional MSP

Expanding client base, no after-hours cover, on-call engineers burning out.

P1 only

events that now wake the on-call engineer overnight

The nights, finally quiet.

On-call engineers were being woken for forgotten passwords and trivial questions, and the rota was wearing people down in a market where replacing them is hard. The agent now answers overnight, verifies and triages, and against the MSP's own rules decides what waits for morning and what is a genuine priority-one event.

The on-call engineer is now roused only for events that truly need a human at night, and the morning queue arrives already logged and prioritised.

A NOTE ON THESE EXAMPLES

These are composites built to illustrate common patterns, not named clients. Your results depend on your call volume, your SLA rules and how you choose to use the agent. We are happy to talk through a realistic picture for your specific desk.

SECTION ELEVEN

Questions MSPs ask

The questions that come up in almost every first conversation, with straight answers.

11

FREQUENTLY ASKED

The honest answers.

Will callers know they are talking to an AI?

Yes, always. The agent tells callers plainly that they are speaking with your automated service desk assistant. Transparency is a requirement, not an option, and a caller can ask for a person at any time.

Can it reset a password or grant access if someone insists?

No. It does not perform privileged actions on its own authority, by design. It verifies identity to your policy and routes a verified request to an authorised engineer. Urgency and seniority on the line change nothing. The whole guide is built on that boundary.

Does it work with our PSA and ticketing tools?

It connects to common professional services automation and ticketing systems through permissioned access, reading and writing the fields you allow. We confirm compatibility with your specific stack before you commit to anything.

Where does our clients' data go?

It is handled under your privacy policy and the Australian Privacy Principles, collected only for the request, encrypted, and held within boundaries you control. Every interaction is logged, and access is limited to the people who need it.

Will it replace our service desk staff?

No. It takes the repetitive, interruptive layer of the work so your people can focus on diagnosis, project delivery and the calls that need a human. MSPs generally redeploy their engineers onto billable work rather than reduce them.

How is it different from the chatbot in our ticketing tool?

A chatbot waits on a web form and captures text. A digital teammate answers the phone, holds a real conversation, verifies the caller, and completes the routine job end to end, including writing the triaged ticket into your PSA.

Could an attacker social-engineer the agent?

That is the threat it is designed against. It runs your verification the same way every call, cannot be hurried or flattered into skipping a step, and cannot perform the privileged action itself. Consistency at 2am is exactly where humans are most vulnerable and the agent is strongest.

How long until it is live?

Days, not quarters. The groundwork in Section Seven is the main task, and most of it is writing down how your desk already runs. After that, configuration and testing are quick.

What if we want to change how it behaves?

You change the rules and verification steps in plain language and the change applies to every call from then on. Consistency becomes a setting. You can also listen back and read transcripts whenever you like.

SECTION TWELVE

Your first seven days

A short, concrete path from reading this guide to a live agent answering your desk. Three steps, one week.

12

FROM HERE TO LIVE

Start with one job. Grow from there.

DAYS 1–2

Map the basics

Work through the Section Seven checklist. Pin down your top call reasons, SLA rules, verification steps and escalation points. This is the real work.

DAYS 3–5

Configure and test

We set the agent up to match your rules, connect it to your PSA in a controlled way, and test it together against real call scenarios, including the awkward ones, until it sounds like your desk.

DAYS 6–7

Go live on one job

Start with a single job, often after-hours triage or overflow calls, watch the reports, then widen its remit once you trust it. Small start, fast confidence.

THE ONE DECISION TO MAKE THIS WEEK

You do not need to commit to a full rollout. Pick the single job that hurts most right now, the Monday flood, the after-hours nights, or the senior engineers stuck on level-one calls, and let an agent take just that. The rest follows from what you learn.

When you are ready, the best next step is a short walkthrough where we build a realistic picture for your MSP: your numbers, your rules, and a clear view of what the agent would and would not do. No generic figures, no pressure.

Answer every ticket. Verify before you act.

A digital teammate that works your phones, your queue and your nights, runs your identity checks the same way every time, and never performs a privileged action on its own. Built for Australian MSPs.

TALK TO US

Agntic.ai
Book a 20-minute walkthrough for your desk

ABOUT THIS GUIDE

General information only.
Not legal, privacy or security advice.
Confirm your obligations with the appropriate adviser.