




Smart AI Setup Guide



Help you
to choose
the **best**
AI
hardware
setup
without
needing
to be a
tech
expert.

Prepared by : Team D – AI Hardware Scoping Project | May 2025

For Business Owners, Startup Founders, and Enterprise Decision-makers



Dear Business Owner,

We know AI can feel overwhelming. From cloud costs to technical jargon, it's tough to know where to start - especially if you're not a tech expert.

That's why we created this guide.



Our goal is simple:

“To help you understand what AI can actually do for your business and show you which type of hardware setup will get the job done, without overpaying or overbuilding.”

You'll find:

Real-world examples of AI use in business



Setup options across different budgets



Clear answers to: “Can I do what I need to do with this setup?”



Whether you want to add a chatbot, summarize customer feedback, or train your own AI tools — we've got you covered

Let's get started!

-Team D, Federation University-

What Can AI Actually Do for You?

Here are just a few things you can do with your own AI setup:



Answer customer questions instantly (via chatbot)



Summarize long reports, emails, or legal documents



Translate communications across multiple languages



Generate responses, marketing drafts, or memos



Train the AI to understand your tone, policies, or product details



Automate HR, support, or data review tasks

**No cloud subscription.
No privacy worries.**

Everything runs on your own hardware.

FAST, SECURE, AND SCALABLE.



Do I Need a Big AI Model or a Small One?

You don't need a massive model for most tasks.

Here's how to decide what's "enough":




Business Goal	Recommended Model Type	Why?
Internal chatbot, FAQs, summarization	LLaMA 7B or similar (small)	Fast, private, low-cost
HR support, compliance emails, internal reports	Falcon 7B–40B (medium)	Handles longer text and context
Deep training, multi-lingual support, large data	BLOOM 176B or Falcon 180B (large)	Advanced AI, high-power tasks



TIP:

If your goal is “auto-reply to customers,” you probably don’t need to spend more than \$2–5K.



How to Use This Catalogue

We've broken this guide into five price tiers, each with three ready-to-go hardware options:

Option Type

- Inference-Optimized
- Balanced
- Training-Optimized

What It Does

Best for answering questions and summaries

Can both serve and train smaller models

Built for learning and customizing your AI

Each setup is written in plain language:



What it includes



What you can do



Why it's useful

\$2K Tier Overview: Smart Start for Small Business



Perfect For

- ✓ Solo entrepreneurs, consultants, and local retailers
- ✓ Businesses just starting with AI
- ✓ Tasks like chatbots, simple content automation, and document summarization

What You Can Do at This Tier

- Run small AI models (e.g. LLaMA 7B) on your own computer
- Build a chatbot to answer customer queries
- Summarize documents and automate routine paperwork
- Save money on cloud costs by running AI in-house

Quick Benefits

- Budget-friendly starting point
- Low power use – no need for special cooling or setup
- Can be upgraded later as needs grow
- Fits in any small office or even a home setup



Option 1: Inference-Optimized (\$2K)

“The Smart Start Setup”

Ideal For:

Small businesses needing a fast, quiet system for handling customer questions, automating emails, or document summaries.

Motherboard

MSI B650 Gaming Plus WiFi

RAM

32GB (2×16GB) DDR5
6000 CL36

OS Drive

1TB Samsung 990 Pro
NVMe M.2 SSD

Case

Corsair 4000D Airflow
Mid-Tower

Power Supply

MSI A1000G PCIE5 1000W
80+ Gold PSU

CPU

AMD Ryzen 7 7700,
8-Core 5.3GHz Turbo (65W)

GPU

NVIDIA GeForce
RTX 4060 Ti 16GB

Data Storage

2TB Samsung 970 EVO
Plus NVMe

Cooling

Noctua NH-D15 chromax.
black CPU Air Cooler

Other

Estimated Total - \$2,300 AUD

What You Can Do:

- Run your chatbot 24/7 without paying cloud fees
- Instantly summarize HR reports or emails
- Deploy small AI models without technical help

Benefits:

- Easy to maintain and upgrade
- Runs quietly in any office
- Handles AI tasks without delays

Option 2:

Balanced Setup (\$2K)

"The Flex Performer"

Ideal For:

Business owners who want flexibility, some training, some live AI use. Great for development teams or startups.

Motherboard

ASRock B650 PRO RS
(AM5, ATX)

CPU

AMD Ryzen 9 5900X,
12-Core 4.8GHz Turbo

RAM

64GB (2×32GB) DDR4-
3200

GPU

NVIDIA GeForce
RTX 4070 Ti 16GB

OS Drive

1TB WD Black SN770
NVMe SSD

Data Storage

4TB Seagate Barracuda
HDD

Case

Phanteks Eclipse P400A
Mid-Tower

Cooling

Cooler Master Hyper 212
Evo V2 Air Cooler

Power Supply

Corsair RM850e 850W 80+
Gold PSU

Other

Estimated Total - \$2,800 AUD

What You Can Do:

- Do basic training on your own product or customer data
- Try custom AI responses or tone control
- Run models for customer segmentation or lead scoring

Benefits:

- Versatile for both development and delivery
- Handles large datasets easily
- Great choice for growing teams

Option 3: Training-Optimized (\$2K)

“The Mini AI Lab”

Ideal For:

Tech-savvy founders, researchers, or students who want to start training models themselves.

Motherboard

MSI PRO Z790-P WiFi
(LGA1700, ATX)

CPU

Intel Core i5-13400F,
10-Core

RAM

32GB (2×16GB) DDR5-5200

GPU

2× NVIDIA GeForce
RTX 3060 12GB

OS Drive

500GB Samsung 970 Evo
Plus NVMe

Data Storage

2TB SATA SSD (Crucial
MX500)

Case

Cooler Master - MasterBox
TD500

Cooling

DeepCool AK400 CPU
Cooler

Power Supply

SilverStone DA750 750W
80+ Gold PSU

Other

High-performance
thermal paste

Estimated Total - \$2,300 AUD

What You Can Do:

- RTrain smaller language models on your own customer data
- Test how model tone or accuracy changes with training
- Run multiple experiments side-by-side

Benefits:

- Best training performance at this budget
- Compact dual-GPU power
- Ideal for universities, small labs, or curious developers

\$5K Tier Overview: Smart Start for Small Business



Perfect For

- ✓ Medium-sized teams, tech-forward startups, research groups
- ✓ Teams ready to move beyond AI experimentation into AI operations
- ✓ AI product development, support automation, training your own models

What You Can Do at This Tier

- Fine-tune models like LLaMA 13B or Mistral 7B
- Automate internal support or HR questions
- Train AI to “speak your brand voice”
- Start building AI-powered services (e.g., legal, policy, or product AI)



Quick Benefits

- Much more power without jumping to enterprise costs
- Run multiple models or workflows simultaneously
- More memory, faster storage, and better cooling
- Great middle ground before scaling to \$10K–\$20K setups



Option 1: Inference-Optimized (\$5K)

“The Reliable AI Workhorse”

Ideal For:

Teams running multiple chatbots, internal Q&A tools, or customer-facing automation at scale.

Motherboard

ASUS Pro WS B650M-SYS
MicroATX

CPU

AMD Ryzen 9 7900
(12-core, 5.4GHz boost, 65W)

RAM

64GB DDR5-5200 ECC
(4×16GB UDIMM ECC)

GPU

NVIDIA RTX A5000 24GB
(Workstation GPU, Ampere)

OS Drive

2TB Crucial P3 Plus M.2
NVMe SSD

Data Storage

4TB Samsung 870 QVO
SATA SSD

Case

Fractal Design Define 7 (Mid
Tower, sound-dampened)

Cooling

Noctua NH-U12A chromax.
black (high-end air)

Power Supply

Seasonic Focus PX-850
850W 80+ Platinum PSU

Other

Extra case fans (2× Noctua
NF-A14 for airflow)

Estimated Total - \$5,300 AUD

What You Can Do:

- Run LLaMA 13B, Falcon 7B for customer support, helpdesk, and internal queries
- Serve multiple AI models 24/7
- Handle legal, compliance, or IT documentation Q&A tools

Benefits:

- Built for reliability can run nonstop
- Super quiet - suitable for shared workspaces
- Easy to manage and upgrade

Option 2:

Balanced Setup (\$5K)

“The All-Rounder AI Studio”

Ideal For:

Startups or teams building and using AI at the same time train, test, and deploy from one machine.

Motherboard

Asus ROG STRIX Z790-F
Gaming WiFi (LGA1700)

CPU

Intel Core i9-14900KF (
24-core, 6.0GHz boost)

RAM

64GB DDR5-6400
(2×32GB) CL32

GPU

ASUS ROG Strix GeForce
RTX 4090 24GB

OS Drive

2TB Samsung 990 Pro
NVMe SSD

Data Storage

4TB Seagate IronWolf
Pro HDD 7200RPM

Case

HYTE Y70 Touch ATX Mid
Tower (LCD panel)

Cooling

Corsair iCUE H150i ELITE
LCD XT 360mm AIO Liquid

Power Supply

Corsair RM1000x SHIFT
1000W 80+ Gold PSU

Other

Estimated Total - \$5,500 AUD

What You Can Do:

- Run multiple apps — chatbot, summarizer, translator, sentiment analyzer
- Train mid-size models up to 13B parameters
- Serve AI + build product features on one device

Benefits:

- Massive power, compact setup
- Handles all stages: training, testing, and use
- Smart choice for growing AI product companies

Option 3: Training-Optimized (\$5K)

“The Dual-Engine Trainer”

Ideal For:

AI developers or researchers who need fast training cycles and room for experimentation.

Motherboard

Gigabyte X670 AORUS Elite AX (AM5)

CPU

AMD Ryzen 9 7950X3D (16-core, 5.7GHz, 3D V-Cache)

RAM

128GB DDR5-6000 (4×32GB) CL36

GPU

2× NVIDIA GeForce RTX 4080 16GB

OS Drive

2TB WD Black SN850X NVMe SSD

Data Storage

8TB Seagate Exos X18 7200RPM enterprise HDD

Case

Phanteks Enthoo Pro 2 (SSI-EEB Full Tower, high airflow)

Cooling

Arctic Liquid Freezer II 420 AIO (420mm rad)

Power Supply

EVGA SuperNova 1300 GT 1300W 80+ Gold PSU

Other

Powered riser cables for GPUs

Estimated Total - \$7,500 AUD

What You Can Do:

- Train LLaMA 13B, fine-tune BERT, and explore vision models
- Run multiple training jobs or model tests in parallel
- Build AI features faster with short iteration loops

Benefits:

- Extreme training performance under \$10K
- Ready for research, R&D, and development teams
- Built for the future upgrade path to even larger workloads

\$10K Tier Overview:

High-Performance AI for Builders

Perfect For

- ✓ Startups building AI-powered products
- ✓ Teams training models with large datasets
- ✓ Businesses moving beyond experiments into production-scale tools



What You Can Do at This Tier



- Train and fine-tune models like LLaMA 65B (quantized)
- Handle large document processing tasks
- Run AI services for legal, marketing, HR, or product teams
- Support multiple AI workflows (training + deployment + analytics)

Quick Benefits

- Long-lasting investment for scaling operations
- Much higher performance without jumping to enterprise-level cost
- Great for simultaneous training and live inference
- Flexible for building AI tools, customer apps, or automation workflows



Option 1: Inference-Optimized (\$10K)

“The Enterprise AI Base Station”

Ideal For:

Organizations running several AI services (e.g., legal, HR, support) that must stay fast and reliable around the clock.

Motherboard

Supermicro H12SSL-i ATX
(SP3 socket, server board)

CPU

AMD EPYC 7452 32-Core
@ 2.35GHz

RAM

256GB DDR4-3200 ECC
REG (8×32GB RDIMM)

GPU

4× NVIDIA RTX 3060 12GB

OS Drive

1TB Intel D7-P5510 U.2
NVMe (enterprise SSD)

Data Storage

8TB NVMe U.2 SSD
(2×4TB in RAID0)

Case

4U Rackmount Server
Chassis (Rails)

Cooling

Dynatron SP3 active cooler
+ 4× GPU blower fans

Power Supply

2× 800W 80+ Platinum
Redundant PSUs

Other

IPMI module for remote
management

Estimated Total - \$7,500 AUD

What You Can Do:

- Run multiple AI models at once across different teams
- Automate document workflows for legal, HR, compliance
- Support multi-user access to internal AI tools

Benefits:

- Rackmount-ready for integration into server rooms
- Enterprise-grade stability
- Designed to run 24/7 at scale

Option 2:

Balanced Setup (\$10K)

“The Creative AI Studio”

Ideal For:

Startups or creative teams who want one machine that can train, test, and run AI-powered apps all in one place.

Motherboard

ASUS Pro WS TRX50-SAGE
WIFI (sWRX8)

CPU

AMD Threadripper 3960X
24-Core @ 3.8GHz

RAM

128GB DDR4-3200
(4×32GB)

GPU

2× NVIDIA RTX 3090
Ti 24GB

OS Drive

2TB Corsair MP600 PRO
NVMe

Data Storage

4× 4TB WD Gold 7200RPM
HDD (RAID10)

Case

Fractal Define XL full
tower

Cooling

Custom water loop
(CPU + 2 GPUs)

Power Supply

EVGA 1600W T2 80+
Titanium PSU

Other

High-end thermal
compound + extra radiators

Estimated Total - \$10,000 AUD

What You Can Do:

- Train large models and use them in real time
- Run creative tools like AI image/video generators
- Manage long workflows with lots of input/output (e.g., content or legal teams)

Benefits:

- Smooth, fast, and flexible across multiple AI workloads
- Perfect mix of power and usability
- Stylish and quiet — fits creative offices

Option 3: Training-Optimized (\$10K)

“The AI Powerhouse Workstation”

Ideal For:

Developers or researchers training large language models in-house, or engineering teams powering their own AI stack.

Motherboard

Asus ProArt Z790-Creator
WIFI (LGA1700)

CPU

Intel Core i9-13900K
24-Core @ 5.8GHz

RAM

128GB DDR5-6000
(4×32GB)

GPU

2× NVIDIA RTX 4090
24GB (Founders Edition)

OS Drive

2TB WD Black SN850X
NVMe

Data Storage

1× 8TB Samsung 870 QVO
SATA SSD

Case



Cooling

2× 360mm AIO (one for
CPU, one shared by GPUs
with AIO bracket)

Power Supply

Corsair AX1600i 1600W
80+ Titanium PSU

Estimated Total - \$11,000 AUD

What You Can Do:

- Train LLaMA 65B (quantized) and similar high-load models
- Speed up development with fast iterations compliance
- Test, train, and evaluate models entirely in-house

Benefits:

- Best deep learning performance under \$15K
- Quiet, upgradeable, and office-friendly
- Future-proof setup for serious AI development

\$20K Tier Overview:

AI Infrastructure for Enterprise Teams

Perfect For

- ✓ Medium to large companies with internal AI teams
- ✓ Organizations running AI across multiple departments
- ✓ Those ready to move from cloud reliance to local, private AI hosting

What You Can Do at This Tier

- Train and deploy large-scale AI models (e.g., LLaMA 65B, Mistral 7B, Falcon 40B)
- Support real-time AI services for multiple teams (e.g., finance + legal + support)
- Store, fine-tune, and serve AI safely in-house
- Eliminate most monthly cloud AI costs over time



Quick Benefits

- High memory and GPU capacity
- Supports multiple concurrent users or services
- Can run AI day and night without downtime
- Suitable for secure or regulated environments



Option 1: Inference-Optimized (\$20K)

“The AI Service Hub”

Ideal For:

Enterprises hosting multiple internal chatbots or live model services for staff or clients.

Motherboard

Asus ProArt Z790-CREATOR
(LGA1700)

CPU

Intel Core i7-14700K
(20-core)

RAM

64GB DDR5-5200 (2×32GB)

GPU

2× RTX A6000 48GB

OS Drive

Corsair MP600 Pro NH
2TB NVMe

Data Storage

8TB Seagate FireCuda
530 NVMe

Case

Corsair Obsidian 1000D
Super-Tower

Cooling

2× Custom loops (one for
2 GPUs, one for CPU)

Power Supply

2× EVGA 1600W P+ (parallel
via dual-PSU adapter)

Other

NVLink Bridge for
2× A6000

Estimated Total - \$19,500 AUD

What You Can Do:

- Host multiple live AI applications at once
- Deliver AI-powered services securely in-house
- Serve legal, HR, finance, and helpdesk content in real-time

Benefits:

- Enterprise-grade reliability
- Quiet, office-safe setup
- Fast, consistent AI response with no outside dependencies

Option 2: Balanced Setup (\$20K)

"The VersaGrid AI Tower"

Ideal For:

Organizations who want one machine to serve, train, and experiment — from prototypes to deployment.

Motherboard

ASUS ROG Maximus Z790
Dark Hero (LGA1700)

CPU

Intel Core i9-14900KS
(24-core, 6.2GHz turbo)

RAM

128GB DDR5-6400
(4×32GB)

GPU

4× ASUS ROG Strix
RTX 4090 24GB

OS Drive

WD Black SN850X 2TB
NVMe

Data Storage

2× 4TB Samsung 990 Pro
NVMe (RAID0)

Case

HYTE Y70 (or Corsair
7000D) full tower

Cooling

EKWB Custom Loop for
4× GPUs + CPU

Power Supply

2× Corsair AX1600i 1600W

Other

High-amperage circuit
installation for 15A draw

Estimated Total - \$20,000 AUD

What You Can Do:

- Train large models, deploy them, and analyze results on the same system
- Support multiple departments or teams in parallel
- Move faster with high-speed storage + GPUs

Benefits:

- Versatile for R&D and production
- Designed for teams sharing AI infrastructure
- Balanced between inference, training, and experimentation

Option 3: Training-Optimized (\$20K)

“The Deep Learning Engine”

Ideal For:

Teams focused entirely on training powerful models, experimenting, or conducting AI research.

Motherboard

Supermicro H13SSL-N
(SP5 socket)

CPU

AMD EPYC 9554P
(64-core, 3.7GHz boost)

RAM

512GB DDR5-4800 ECC
RDIMM (8×64GB)

GPU

2× NVIDIA RTX 4090 24GB +
1× NVIDIA RTX 6000 Ada 48GB

OS Drive

2TB Samsung 990 Pro NVMe
(or 4TB variant if needed)

Data Storage

1× 16TB NVMe U.2 SSD
(OWC Accelsior 8M2)

Case



Cooling

4× 4U passive heatsinks +
high-CFM fans.

Power Supply

3kW Rack PSU system
(3×1000W N+1)

Other

Infiniband or 100Gb Ethernet

Estimated Total - \$40,000 AUD

What You Can Do:

- Train large models or multiple smaller ones in parallel
- Run deep learning workloads across long durations
- Build AI capabilities internally from the ground up

Benefits:

- Built for training faster results, more control
- Supports high-demand developers or researchers
- Perfect for teams building their own AI stack

\$50K+ Tier Overview: Supercomputing Power for AI Leaders



Perfect For

- ✓ AI companies, research labs, and enterprise innovation hubs
- ✓ Building custom AI models (e.g., LLaMA 65B+, Falcon 180B, BLOOM 176B)
- ✓ Supporting high-volume AI workloads or developing new AI products

What You Can Do at This Tier



- Train and serve full-scale language models
- Power multiple AI projects (chatbots, translation, research) at once
- Eliminate recurring cloud costs for AI training and usage
- Establish complete in-house AI infrastructure

Quick Benefits

- Designed to match or exceed cloud-scale compute
- Enterprise-class reliability and performance
- Supports large team access and multiple services
- Fully owned, private, and secure



Option 1: Inference-Optimized (\$50K+)

“The Enterprise AI Gateway”

Ideal For:

Organizations needing high-throughput, high-availability AI inference across multiple services and departments.

Motherboard

Supermicro X13SWA-TF (SSI-EEB, supports 4 GPUs)

CPU

AMD Threadripper Pro 7995WX, 96-core

RAM

512GB DDR5 ECC RDIMM (8×64GB)

GPU

4× NVIDIA H100 80GB PCIe

OS Drive

4TB Samsung 990 Pro NVMe

Data Storage

16TB OWC U.2 NVMe (as in Opt3 above)

Case

Phanteks Enthoo 719 (Super Tower, SSI-EEB)

Cooling

Custom loop with 3× 480mm radiators

Power Supply

2× 2000W 80+ Platinum (Super Flower Leadex)

Networking

Mellanox ConnectX-7 200Gb Infiniband + fiber cables

Estimated Total - \$270,000 AUD

What You Can Do:

- Power 24/7 AI services for HR, finance, analytics, customer engagement
- Host internal GPT-style models company-wide
- Deliver AI tools to hundreds of users at once

Benefits:

- Cloud-like scale, on your own terms
- Built for live use with high uptime
- Secure, private deployment without vendor lock-in

Option 2: Balanced (\$50K+)

“The Core AI Infrastructure”

Ideal For:

Tech companies and universities needing both training and inference in one scalable, powerful system.

Motherboard

2× ASUS WS WRX80 SAGE
(for dual TR Pro builds)

CPU

2× AMD Threadripper Pro
5975WX, 32-core

RAM

2×256GB (each node)
DDR4 ECC

GPU

Node1: 4× RTX 4090;
Node2: 4× RTX 4090

OS Drive

Each node: 2TB NVMe

Data Storage

Shared NAS: 100TB RAID
via 12× HDD + flash cache

Case

2× Phanteks Enthoo Pro
(full tower for each node)

Cooling

Each node: 2×360 AIO +
additional fans

Power Supply

Each node: 1× 1600W PSU

Networking

100GbE Switch + 3× NICs
(50K cluster needs fast net)

Estimated Total - \$50,000 AUD

What You Can Do:

- Train and serve models like Falcon 180B (quantized) or BLOOM 176B
- Power multiple AI platforms (chat, genAI, predictive analytics)
- Allow multiple teams to run AI jobs at once

Benefits:

- Flexible for both R&D and product delivery
- Ready for next-gen workloads and tools
- Long-term investment that removes cloud dependency

Option 3: Training-Optimized (\$50K+)

“The AI Lab-in-a-Box”

Ideal For:

R&D labs, universities, or AI companies training their own large language models in-house.

Motherboard

Supermicro H13SSL-N
(or H13DCi for dual CPU)

CPU

2× AMD EPYC 9654,
96-core each.

RAM

1TB DDR5 ECC (16×64GB)

GPU

4× NVIDIA H100 SXM5
80GB (with NVSwitch)

OS Drive

2× 4TB NVMe U.2 (mirror)

Data Storage

8× 16TB NVMe (RAID0)
for 128TB ultra-fast

Case

Supermicro 8U chassis with
liquid cooling backplane

Cooling

Direct hot water liquid cooling
plates for CPUs/GPUs

Power Supply

4× 3000W PDUs (for redundant
supply, 3-phase feed)

Networking

Infiniband HDR 200Gb/s
full mesh (switch + NICs)

Estimated Total - \$300,000 AUD

What You Can Do:

- Train large foundation models from scratch
- Run full-lifecycle AI research in-house
- Support 10+ researchers or engineers simultaneously

Benefits:

- Matches academic and enterprise research centers
- Eliminates cloud training costs for large AI projects
- Expandable — perfect for pilot AI clusters



Tier Comparison Summary

Tier	Budget Range	Best For	Max Model Size	What You Can Do
\$2K	\$2000-2800	Solopreneurs, first-time users	LLaMA 7B (basic)	Run a chatbot, summarize docs
\$5K	\$5000-7500	Startups, mid-size teams	LLaMA 13B / Falcon 7B	Train light models, automate workflows
\$10K	\$7500-11000	Product dev teams, labs	LLaMA 65B (quantized)	Train, test, and run custom tools
\$20K	\$19500-20500	Enterprise teams	Falcon 40B / Mistral 7B	Parallel AI use across departments
\$50K+	\$55000-300000	Universities, AI companies	BLOOM 176B / Falcon 180B	Build your own GPT-like infrastructure



TIP:

The sweet spot for most businesses is \$5K–\$10K — powerful enough to train your own tools, affordable enough for fast ROI.

WHICH TIER IS RIGHT FOR YOU?

Ask yourself these 3 questions:

1

What do I need AI to do today?

- Answer emails?
- Train a chatbot?
- Write reports?

2

Do I want to train models or just use them?

- If just using
start with \$2K–\$5K
- If training your own
look at \$5K+ tiers

3

How much will I use it?

- Light daily tasks
Inference builds are enough
- Heavy usage or large data
Training-ready builds are better

Visual Decision Guide:

Your Goal

Just need a chatbot for customer queries

Want to train a basic model

Run AI for multiple departments

Serve and train models organization-wide

Replace cloud AI with your own lab

Start Here

\$2K Inference Option

\$5K Training Option

\$10K Balanced Option

\$20K+ Tier

\$50K+ Tier





Final Thoughts & Getting Started



AI is no longer just for tech giants.

With the right setup, you can bring AI in-house, protect your data, save money, and serve your customers faster.



Next Steps: ↘

1. Identify what your business needs AI to do.
2. Pick the tier that gives you enough power - not more than you need.
3. Talk to your IT provider, integrator, or local supplier with this catalogue in hand.

Need help building or choosing?



teamd.aihardware@feduniversity.edu.au



Federation University Australia – AI Research Group



Version: May 2025