ENTERPRISE GUIDE

Claude Enterprise Deployment Guide

A practitioner's guide to workshopping, designing, deploying and configuring Anthropic's Claude across your organisation — from first discovery session to production at scale.

Audience
CIO / CTO
IT & Security Leads
AI Programme Champions
Department Heads

Covers
Requirements Workshop
Architecture Design
Implementation Guide
Cowork Configuration
As-Built Documentation

Products
Claude Enterprise
Claude Cowork
Claude Code
MCP Integration
Claude API / Excel / PPT

Produced by Fynaptix · Enterprise Claude AI Architecture & Deployment · fynaptix.com

NAVIGATION

Table of Contents

This guide covers the complete lifecycle of a Claude Enterprise deployment — from initial discovery workshops through to production governance. Each section is designed to be used independently or as part of a sequential engagement.

01 Introduction

The Enterprise AI Deployment Gap

Social media is full of Claude demos. What's missing is a practical end-to-end guide that takes an enterprise from "we should be using AI" to production systems that actually change how work gets done. This guide fills that gap.

Why Deployments Stall

The pattern is familiar: a team runs a successful proof-of-concept. Claude handles a legal document review, summarises a week of emails, or drafts a board report. Executives are impressed. Then nothing happens for six months.

The bottleneck is never capability. The bottleneck is the governance layer: identity integration, data controls, RBAC, spend management, observability, and the change management that turns a tool into a habit. These are solvable problems. This guide shows you how.

The Demo Gap
- Impressive POC with no follow-through
- IT Security not engaged
- No RBAC design completed
- No spend controls defined
- No governance framework

The Integration Gap
- SSO not configured
- Connector mapping undefined
- MCP servers not planned
- Data classification unclear
- No SIEM integration

The Adoption Gap
- No onboarding programme
- No use-case library built
- No champion network
- No success metrics defined
- Tool sits unused: shelfware

The Claude Enterprise Product Family

A complete deployment touches multiple Claude products. Understanding how they relate is essential before architecture design begins.

Product	Primary Users	Primary Use	Key Dependencies
Claude Enterprise (claude.ai)	All staff	Chat, Projects, deep research, document drafting	SSO, RBAC, data retention config
Claude Cowork	Knowledge workers	Agentic desktop tasks, file automation, recurring workflows	MDM, connectors, MCP servers, OTel/SIEM
Claude Code	Engineering teams	Agentic coding, codebase analysis, CI/CD workflows	CLAUDE.md, MCP servers, git integration
Claude API	Developers	Product integration, custom agentic systems, RAG pipelines	API keys, rate limits, prompt versioning strategy
Claude for Excel / PowerPoint / Word	Business users	In-app AI assistance for M365 workflows	M365 tenant, Enterprise licence

02 Requirements Workshop

Discovery & Scoping

A structured two-day discovery workshop is the foundation of every successful deployment. It aligns stakeholders, surfaces constraints, and produces the artefacts your architecture and implementation phases depend on.

Who Must Be in the Room

Role	Contribution	Day
CIO / CTO	Strategic priorities, budget envelope, success definition	Both
IT Security / CISO	Data classification, compliance constraints, SIEM requirements	Day 1
IT Infrastructure	Identity provider, MDM platform, network topology	Day 1
Legal / Compliance	Data residency, retention policy, regulatory obligations	Day 1
HR / People & Culture	Acceptable use policy, training requirements, change concerns	Day 1
Department Heads (2-4)	Use case identification, current pain points, tool landscape	Both
AI Champions (1-2)	Existing AI experiments, prompt experience, technical curiosity	Day 2

Day 1 Agenda — Current State Mapping

Morning (3 hours): Organisation & Constraints

09:00 — Scene setting: Why this deployment, what success looks like, executive mandate
09:30 — Current AI maturity audit: What's already in use, shadow AI inventory, existing tool landscape
10:30 — Data classification workshop: Identify what's PUBLIC, INTERNAL, CONFIDENTIAL, RESTRICTED
11:30 — Compliance constraints: Industry regulations, data residency, retention obligations
12:00 — Wrap morning, document open questions list

Afternoon (3 hours): Use Case Discovery

13:00 — Department-by-department time-sink mapping: What takes most time that Claude could handle?
14:00 — Use case scoring: Impact x Feasibility x Data Risk x Speed-to-value matrix
15:00 — Tool landscape mapping: Which connectors will Claude need? (Slack, GDrive, M365, Salesforce...)
15:45 — Identify pilot candidate teams: one technical + one non-technical team recommended
16:30 — Day 1 retro and homework: each attendee documents top 3 use cases overnight

Day 2 Agenda — Future State & Pilot Design

Morning (3 hours): Architecture Decisions

09:00 — Use case prioritisation: review overnight submissions, consensus scoring, select top 8-10
09:45 — Infrastructure decision: Anthropic-hosted vs AWS Bedrock vs Google Vertex vs Azure Foundry
10:30 — Governance model: RBAC role matrix draft, connector approval process, MCP allowlist ownership
11:15 — Observability requirements: SIEM platform, alert priorities, triage ownership
11:45 — Security review: data handling commitments, model training opt-out, retention configuration

Afternoon (3 hours): Pilot Planning

13:00 — Pilot team selection: define champion cohort (8-15 people across 2 teams)
13:30 — Pilot use case selection: 3 recurring tasks per person to delegate to Claude
14:00 — Success metrics: time saved, task completion rate, user satisfaction, error rate
14:45 — Communications plan: how will the organisation hear about this? Change management approach?
15:30 — Roadmap draft: 6-month phased rollout from pilot to organisation-wide
16:00 — Close: actions, owners, deadlines, next workshop date

Workshop Output Artefacts

Five artefacts produced during the workshop feed directly into architecture and implementation:

Use Case Register
Scored & prioritised list
Data classification per use case
Pilot assignment per use case
Success metric per use case

Data Classification Map
4-tier classification schema
System-to-classification mapping
What Claude can access per tier
Approved connector list per tier

RACI Matrix
Deployment programme RACI
Ongoing governance RACI
Incident response ownership
New connector approval chain

Pilot Brief
Team composition
Use cases selected
Duration & cadence
Measurement framework

Constraints Register
Compliance obligations
Data residency requirements
Blocked integrations
Security non-negotiables

6-Month Roadmap
Phase 1: Pilot (weeks 1-8)
Phase 2: Expand (weeks 9-16)
Phase 3: Scale (weeks 17-24)
Phase 4: Optimise (ongoing)

Key Decisions the Workshop Must Produce

Inference provider: Anthropic direct, AWS Bedrock, Google Vertex, or Azure Foundry?
Identity provider and SSO protocol: Okta, Azure AD, Google Workspace, or SAML 2.0?
MDM platform for desktop deployment: Intune, Jamf, or Kandji?
SIEM platform: how will Cowork OpenTelemetry data be routed to it?
MCP allowlist ownership: who approves new servers? What is the approval SLA?
Data retention period: configure in Enterprise admin before any users are provisioned
Acceptable use policy stance: does it cover Claude generating content for external use?
Minimum Cowork version: enforce >= 2.0.65 via MDM for current CVE patch coverage

03 Architecture & Design

Technical Blueprint

A well-designed Claude Enterprise architecture has four properties that ensure it ages well: provider-agnostic client, per-device/per-user identity, schema-owned audit trail, and per-call inference routing.

Architecture Decision Record

Record these four core decisions before any configuration begins. They determine every downstream technical choice.

Decision	Options	Recommendation	Key Consideration
Inference Provider	Anthropic Direct / AWS Bedrock / Google Vertex / Azure Foundry	Bedrock or Vertex for regulated industries	Data sovereignty & existing cloud commitments
Identity Integration	Okta / Azure AD / Google Workspace / SAML 2.0	Use existing IdP — no new identity system	Configure SCIM for automated lifecycle management
Desktop Deployment	Intune / Jamf / Kandji / Manual	MDM mandatory for Cowork enterprise governance	Managed settings enforce allowlists & version floors
Observability	Anthropic native / SIEM via OTel / Both	OTel to SIEM required — Cowork not in audit logs	Cowork activity excluded from all native compliance tools

Identity & Access Design

SSO integration is the first technical task. All Claude Enterprise access should flow through your existing identity provider — no local accounts, no shared credentials. Configure SCIM provisioning so that when someone joins or leaves, Claude access is managed automatically through your HR and identity systems.

Network-level enforcement: inject the anthropic-allowed-org-ids HTTP header at your proxy/firewall. This prevents users on managed devices from authenticating to personal Claude accounts — a critical control for data handling compliance.

RBAC Role Matrix

Design your role matrix before any configuration. Apply most-restrictive defaults to all users and elevate specific groups based on demonstrated need. RBAC controls six Cowork capabilities but does not cover Chrome, plugins, MCP servers, or connectors — treat it as one layer of a defence-in-depth model.

Role	Claude.ai	Cowork	Connectors	Plugins	MCP	Code

Read-Only	Chat	—	—	—	—	—
Standard	Full	Basic	Tier 1 only	Curated	—	—
Power User	Full	Full + scheduled	Tier 1-2	All approved	Allowlisted	—
Developer	Full	Full	All approved	All	All approved	Full
Admin	Full + Mgmt	Full	All	All + publish	All + register	Full

Connector & MCP Governance Tiers

Every connector and MCP server is an attack surface. Define governance tiers before deployment and enforce them through MDM managed settings.

Tier 1 — Pre-Approved
Google Workspace, Microsoft 365
Slack, standard productivity tools
Internal read-only MCP servers
Auto-enabled for Standard+ users
No additional approval required

Tier 2 — Manager Approval
Salesforce, HubSpot, DocuSign
FactSet, S&P Global / Kensho
Read-write internal MCP servers
Manager approves + IT Security logs
SLA: 5 business days

Tier 3 — CISO Approval
Custom external MCP servers
Third-party unvetted plugins
Computer Use capability
CISO sign-off + pen test required
Quarterly re-approval

Inference Gateway Architecture

For regulated industries or data sovereignty requirements, configure an inference gateway rather than routing to api.anthropic.com directly. Set inferenceProvider: gateway and inferenceGatewayBaseUrl in MDM managed preferences. The wire format is the Anthropic Messages API — AWS Bedrock, Google Vertex, and any OpenAI-compatible endpoint via conversion layer are all compatible. This provides provider portability as the model capability landscape evolves.

Observability Architecture

Cowork activity is NOT captured in Anthropic's native Audit Logs, Compliance API, or Data Exports on any plan tier — including Enterprise. You must route OpenTelemetry telemetry to your own SIEM. Cowork streams tool calls, file access events, and approval states via OTel. Map these to your SIEM schema and define alert rules before go-live.

Minimum six alert rules to deploy on day one:

- Bulk file access: >50 files in a single Cowork session — possible data exfiltration
- Off-hours scheduled tasks: new recurring task created outside 07:00-20:00 local
- MCP tool calls to sensitive systems (HR/Finance/Legal) outside business hours
- Cowork client version below minimum floor (< 2.0.65) on any managed device
- Unauthorised OAuth grant: user self-approved a connector scope
- Inference gateway error rate: >5% request failure over 10-minute window

04 Implementation

Step-by-Step Deployment

Implementation follows a strict sequence. Each step has dependencies on the previous. Do not run steps in parallel — governance must be in place before users are onboarded.

Phase 1 — Foundation (Weeks 1-2)

1

Claude Enterprise Account Setup

Contact Anthropic Sales at anthropic.com/contact-sales. Enterprise covers Claude Code and Cowork under one agreement. Configure your data retention policy and model training opt-out in the admin console BEFORE any users are provisioned. Note your data residency configuration if applicable.

2

SSO & Identity Provider Integration

Connect your IdP (Okta, Azure AD, Google Workspace) via SAML 2.0 or OIDC. Configure SCIM for automated user provisioning and deprovisioning. Test the full JIT provisioning flow with a test user before proceeding. Inject the anthropic-allowed-org-ids header at your proxy/firewall layer.

3

Network & Inference Gateway Configuration

If using Bedrock/Vertex: configure your inference gateway, set managed preferences via MDM (inferenceProvider, inferenceGatewayBaseUrl). Test end-to-end inference from a managed device through the gateway. Confirm TLS inspection policy allows Claude traffic without breaking certificate validation.

4

RBAC Configuration & Spend Limits

Create role groups in the Claude admin console mapped to your IdP groups. Apply most-restrictive defaults to all users. Create Power User and Developer elevation groups. Set group-based spend limits — most-restrictive precedence applies. Document every role-to-capability mapping in your RBAC register before proceeding.

Phase 2 — Cowork Deployment (Weeks 2-3)

5

Desktop App Deployment via MDM

Package the Claude Desktop app for deployment via Intune, Jamf, or Kandji. Set minimum version requirement >= 2.0.65. Push managed preferences with: inferenceProvider, approved connector list, file system mount points, and approved plugin marketplace URL. Verify managed settings applied on a test device before fleet deployment.

6

Connector Provisioning

Enable Tier 1 pre-approved connectors for Standard+ users. Configure OAuth scopes at admin level — do not allow users to self-approve new OAuth grants. For each connector: document data classification touched, roles permitted, and add to Connector Register. Test each connector from a standard user account.

7

MCP Server Allowlist

Register all approved MCP servers in MDM managed settings. For each server: document tools exposed, data accessible, capability grants, and roles permitted. Internal read-only servers are Tier 1. Read-write servers require Tier 2 approval. Test every server from a managed device. Log all tool invocations to SIEM.

8

OTel / SIEM Integration

Configure OpenTelemetry export from Cowork to your SIEM (Splunk, Datadog, Microsoft Sentinel). Map OTel schema to your SIEM data model. Deploy the six minimum alert rules. Run a tabletop exercise: simulate bulk file access, verify the alert fires within SLA. Document SIEM integration in as-built documentation.

Phase 3 — Claude Code (Weeks 3-4, Engineering Teams)

9

Claude Code Installation & CLAUDE.md

Install Claude Code via npm (requires Node.js). Create an organisation-level CLAUDE.md template with: coding standards, approved MCP servers, security constraints, and escalation procedures. Distribute via repository template. Configure git integration. Set API key management via your secrets manager (Vault, AWS Secrets Manager).

10

Engineering MCP Server Network

Deploy MCP servers for engineering toolchain: Jira/Linear (issues), GitHub/GitLab (code), internal documentation, deployment pipelines. Each server: least-privilege tool grants, scoped to repos and projects the engineer can access. Do not grant Claude write access to production systems without an explicit human approval gate.

Phase 4 — Pilot Onboarding (Weeks 4-5)

11

Pilot User Provisioning

Provision pilot cohort (8-15 users across 2 teams) via SSO. Assign to Power User RBAC group. Enable all Tier 1 connectors. Brief IT helpdesk on expected ticket types. Issue welcome pack: what Cowork can do, what it cannot, how to flag issues, who to contact. Schedule weekly check-ins for the 6-week pilot period.

12

Baseline Measurement & Analytics

Before pilots begin using Claude, measure baseline on 3 target tasks per person: time to complete, self-assessed quality score. Set up usage analytics dashboard in the Claude admin console. Configure weekly automated report: active users, tasks completed, connectors used, spend by group.

Go / No-Go Checklist

Confirm all items before any user other than admins accesses the production environment:

- SSO tested end-to-end — JIT provisioning works; deprovisioning removes access within 1 hour
- Data retention policy configured — training opt-out confirmed active in admin console
- RBAC matrix applied and tested — verified standard user cannot access Power User features
- Spend limits set for all groups — tested that limit enforcement triggers correctly
- anthropic-allowed-org-ids enforced at network layer — tested on a managed device
- OTel/SIEM active — minimum 6 alert rules deployed and successfully tested via tabletop
- Cowork minimum version enforced via MDM — all managed devices on >= 2.0.65
- Connector OAuth scopes reviewed and locked — users cannot self-approve new grants
- MCP allowlist deployed via MDM — non-allowlisted servers blocked on managed devices
- Helpdesk briefed — triage runbook distributed to IT support team
- Acceptable Use Policy updated to include Claude — staff notification sent
- Incident response plan updated — Claude-specific scenarios documented and tested

05 Cowork Configuration by Department

Automation & Workflows

Cowork's value is highest when configured for how a specific team actually works. Each department section covers: recommended starting tasks, connector requirements, the plugin to build, and example recurring automation task descriptions to configure on day one.

A key principle: teams do not hand Claude their core work. They hand it the work that surrounds their most critical tasks — project updates, research sprints, collaboration decks, status digests. Start there and expand from demonstrated success.

Finance & Accounting

Starting Tasks
Weekly budget vs actuals narrative
Variance analysis write-up
Journal entry preparation notes
Invoice & PO reconciliation summary
Board report financial section draft

Connectors Required
Microsoft Excel (native integration)
SharePoint / Google Drive
Slack (output delivery channel)
ERP read-only MCP (NetSuite / SAP)
Email for approvals workflow

"Finance Analyst" Plugin
Pre-loads: chart of accounts
Pre-loads: approval thresholds
Pre-loads: period close calendar
Pre-loads: reporting templates
Skill: variance analysis workflow

Recurring automation tasks to configure on day one:

- Monday 08:00 — Pull prior week actuals from SharePoint, compare to budget, write 3-paragraph variance narrative, post to #finance-weekly Slack channel
- Last business day of month — Summarise month-end close checklist status, flag outstanding items, draft CFO summary email
- Daily 09:00 — Check invoice approval queue in ERP, list items pending >3 days, draft personalised reminder messages to approvers
- Weekly Friday 16:00 — Compile expense report submissions, flag policy exceptions, generate summary for Finance Manager review

Legal & Compliance

Starting Tasks
NDA triage and risk flagging
Contract clause extraction
Compliance checklist completion
Regulatory update summaries
Matter status report drafts

Connectors Required
DocuSign connector
Document management MCP (iManage / NetDocuments)
SharePoint / Google Drive
Regulatory feed (read-only web)

"Legal Analyst" Plugin
Pre-loads: approved clause library
Pre-loads: standard contract terms
Pre-loads: jurisdiction map
Pre-loads: escalation thresholds
Skill: NDA triage workflow

Recurring automation tasks:

On NDA receipt — Extract party names, term, governing law, obligations, non-standard clauses. Flag deviations from template. Route to assigned lawyer with structured summary.

Monday weekly — Scan regulatory update feed for jurisdiction-relevant changes, summarise implications, post to #legal-updates channel

Monthly — Generate matter status report: active matters, billing status, upcoming deadlines, dormant files >60 days

Operations & Project Management

Starting Tasks
Meeting notes to action items
Project status update drafts
Stakeholder communication drafts
Vendor research summaries
Process documentation drafts

Connectors Required
Slack (input & output)
Google Drive / SharePoint
Project management MCP (Asana / Monday.com / Jira)
Calendar connector + Email

"Ops Coordinator" Plugin
Pre-loads: project templates
Pre-loads: stakeholder map
Pre-loads: escalation matrix
Skill: status report workflow
Skill: meeting-to-actions workflow

Recurring automation tasks:

After each project meeting — Extract transcript, identify action items with assignees and due dates, update project tracker, post summary to project Slack channel

Friday 15:00 — Pull project status from tracker, generate RAG (Red/Amber/Green) summary for each active project, draft weekly ops report for leadership

Weekly — Scan trackers for tasks overdue >5 days, draft personalised reminder messages to task owners

Marketing & Communications

Starting Tasks
Campaign brief first drafts
Content calendar population
Competitive research summaries
Social media draft batches
Email copy variations by persona

Connectors Required
Google Drive / SharePoint
Slack
Chrome (web research)
CRM read-only MCP
Analytics platform MCP (read)

"Brand Writer" Plugin
Pre-loads: brand voice guide
Pre-loads: messaging framework
Pre-loads: approved claims list
Pre-loads: style guide
Skill: brief-to-copy workflow

Recurring automation tasks:

Monday weekly — Research top 5 competitor content pieces from prior week, summarise themes and gaps, identify content opportunities, post to #marketing-intel

Monthly — Pull campaign performance metrics, write performance narrative against objectives, draft recommendations for next month, prepare slide content

Weekly — Generate 10 social post drafts from content calendar topics using brand voice guide, save to drafts folder for human review and scheduling

People & Culture / HR

Starting Tasks
Job description first drafts
Onboarding documentation
Policy document updates
Employee survey analysis
Training material drafts

Connectors Required
HRIS read-only MCP
(Workday / BambooHR)
SharePoint / Google Drive
Slack
Email

"HR Partner" Plugin
Pre-loads: org chart snapshot
Pre-loads: policy library
Pre-loads: approved JD templates
Pre-loads: onboarding checklist
Skill: JD generation workflow

Recurring automation tasks:

On new role request — Generate first-draft JD from template, role requirements, and similar roles in library. Flag if responsibilities overlap with existing roles.

Weekly — Monitor probation tracker, identify employees entering final month, generate personalised check-in prompts for line managers

Monthly — Analyse survey responses, identify recurring themes by department, generate sentiment summary and recommended actions for People team

Engineering & Technology

Claude Code (Dev Work)
Feature development across services
Legacy code modernisation
PR review and code quality
Test generation and coverage
Documentation generation

Cowork (Surrounding Work)
Sprint planning doc drafts
Technical spec write-ups
Incident postmortem reports
Runbook generation from code
Architecture decision records

Engineering Plugin
Pre-loads: tech stack documentation
Pre-loads: coding standards
Pre-loads: CLAUDE.md template
MCP: Jira + GitHub + internal docs
Skill: ADR generation workflow

Key CLAUDE.md sections for engineering teams:

APPROVED_MCP_SERVERS — list approved servers with tool-level capability grants

CODING_STANDARDS — link to style guide, linting rules, test coverage requirements

SECURITY_CONSTRAINTS — no secret commits, no direct prod DB writes, PR required for main

ESCALATION — when to involve a human before proceeding with a significant action

CONTEXT — architecture overview, service dependencies, deployment pipeline summary

06 Adoption & Change Management

From Pilot to Organisation-Wide

Technical deployment is only half the job. Organisations that get the most from Claude invest as much in the human adoption layer as the technical one.

Five-Level Maturity Model

Every team moves through five levels of Cowork adoption. Understanding where each team sits determines what to focus on next.

Level	Name	What's Happening	Focus Next
1	Chat Q&A;	One-off questions, drafts, research via claude.ai chat	Identify 3 recurring tasks per person to delegate
2	Task Delegation	Individuals use Cowork to delegate multi-step tasks with connectors	Build shared prompt library, introduce team-wide use cases
3	Recurring Automation	Scheduled tasks run automatically without user initiation	Design first department plugin, measure ROI, brief leadership
4	Plugin Specialisation	Department has a private plugin with bundled skills, connectors and knowledge	Extend to adjacent teams, connect to enterprise systems
5	Org-Wide Platform	Claude embedded in how the organisation operates — not bolted on	Governance review cycle, advanced agentic systems design

6-Month Rollout Roadmap

Phase	Weeks	Activities	Success Metric
1 — Foundation	1-2	Technical deployment, SSO, RBAC, SIEM, admin training	Go/No-Go checklist 100% complete
2 — Pilot	3-8	Champion cohort onboarded, weekly check-ins, prompt library building	70%+ weekly active rate, 3 tasks/user delegated
3 — Expand	9-16	Pilot learnings applied, 2nd cohort onboarded, first department plugin	50%+ of target dept active, first plugin in production
4 — Scale	17-24	Org-wide rollout, all department plugins, manager enablement	60%+ org-wide active within 2.5 weeks of broad launch

Champion Network Design

The champion network is the single most important adoption lever. Identify 1-2 champions per department — people who are enthusiastic about Claude, credible with peers, and willing to be first point of contact for questions. Champions share what's working, curate the prompt library, and surface friction points back to the programme team. They are not IT support. Monthly champion sync, standing Slack channel, small recognition budget for sharing wins.

Training Programme

Tier	Audience	Format	Topics
Foundations	All staff	1-hour live + self-paced video	What Claude is, acceptable use policy, how to get help
Practitioner	All Cowork users	Half-day workshop per department	Cowork basics, connectors, task prompting, recurring tasks, use case library
Advanced	Power users & champions	Full-day intensive + follow-up	Plugin design, MCP usage, complex workflow design, prompt engineering
Technical	Developers	2-day developer workshop	Claude Code, API integration, MCP server development, agentic patterns

07 As-Built Documentation

Templates & Registers

Production Claude systems require documentation that survives staff turnover, model updates, and audit cycles. These templates define the minimum acceptable as-built documentation for a Claude Enterprise deployment.

1. Architecture Overview Document

- System diagram: Identity flow (User → IdP → Claude), network topology (managed device → MDM → inference gateway → provider), SIEM integration path
- Inference provider: endpoint URL, authentication method, model versions approved, update policy and change management process
- Data retention: retention period configured, training opt-out status confirmed, deletion SLA
- Network controls: proxy configuration, anthropic-allowed-org-ids enforcement point, TLS inspection policy
- Disaster recovery: inference gateway unavailability procedure, fallback path, RTO/RPO commitments

2. RBAC Register

One row per role. Must be reviewed and signed off quarterly.

Role	IdP Group	Capabilities	Spend Limit	Approval Path	Last Reviewed
Standard User	[idp-group]	Chat, Projects, Cowork basic	$XX/mo	Auto via SSO	DD/MM/YYYY
Power User	[idp-group]	Full Cowork + Tier 1-2 connectors + allowlisted MCP	$XX/mo	Manager + IT	DD/MM/YYYY
Developer	[idp-group]	All above + Claude Code + all approved MCP	$XX/mo	Eng Lead	DD/MM/YYYY
Admin	[idp-group]	Full management access	Unlimited	CIO approval	DD/MM/YYYY

3. Connector & MCP Register

One row per connector/server. Reviewed on any new addition and quarterly.

Name	Type	Tier	Data Class	Roles	Capabilities	Approved By
Google Workspace	Connector	1	Internal	Standard+	Read/write Drive, Gmail, Calendar	IT Security

Slack	Connector	1	Internal	Standard+	Read channels, post messages	IT Security
DocuSign	Connector	2	Confidential	Power+	Send, track, sign documents	Legal + IT
[ERP MCP]	MCP Server	2	Confidential	Power+	Read-only: GL, AR, AP	CFO + CISO
[Custom MCP]	MCP Server	3	Restricted	Power+ w/ approval	[Define tool grants]	CISO + pen test

4. SIEM Alert Runbook

One row per alert. Triage owner must be a named individual, not a team.

Alert	Trigger	Severity	Triage Owner	Escalation
Bulk File Access	>50 files in single session	HIGH	[Security Analyst]	CISO within 1hr
Off-Hours Task	New recurring task outside 07:00-20:00	MEDIUM	[IT Security]	Manager review
Version Below Floor	Cowork < 2.0.65 on managed device	HIGH	[IT Ops]	Force update via MDM
Unauthorised OAuth	User self-approved connector scope	HIGH	[Security Analyst]	Revoke + CISO notify
Sensitive MCP Access	HR/Finance/Legal MCP call after hours	MEDIUM	[Data Owner]	Business justification
Gateway Error Rate	>5% failure rate over 10 min	LOW	[IT Ops]	SLA clock starts

5. BAU Handover Checklist

Daily — Review SIEM alert queue; triage and close or escalate within SLA
Weekly — Review usage analytics: active users, spend by group, connector utilisation
Weekly — Champion network sync: collect friction, wins, and training gaps
Monthly — RBAC review: remove stale elevations, process new role requests
Monthly — Connector/MCP register: audit new additions, verify approvals documented
Monthly — Spot check deprovisioning: confirm departed staff lost access within 1 hour
Quarterly — Full security review: pen test any new Tier 3 MCP servers added in quarter
Quarterly — Model update review: when Anthropic updates model, test top-10 prompts for consistency
Annually — Full compliance review: refresh data retention, retrain staff, update acceptable use policy

08 Security, Governance & Compliance

Reference Guide

Cowork is not a chatbot with extra features. It is a full workstation agent with code execution, file access, browser automation, and scheduled task capabilities running on employee machines. Your security model must reflect this reality.

Understanding Cowork's Threat Surface

A successful prompt injection against Claude.ai leaks conversation context. A successful injection against Cowork can exfiltrate local files, execute shell commands, send messages as the user, and persist via scheduled tasks. Evaluate Cowork security accordingly — not as a chatbot risk, but as a workstation agent risk.

Code Execution
Runs in sandboxed VM on device
Sandbox escape possible for specific tasl
Constrain via MDM mount points
Monitor execution events via OTel

File System Access
Read and write to configured mount poin
Scope mount points to minimum require
Monitor file access volume via OTel
No access to system directories by defau

Browser Automation
Uses employee's authenticated session
Access = what the user can access
Prompt injection via web page content
Scope Chrome connector carefully

Scheduled Tasks
Run unattended — user not present
Can be persisted by malicious instructior
Alert on off-hours task creation
Monthly scheduled task audit minimum

MCP Server Risk
Every server is an attack surface
Third-party code in execution path
Least-privilege tool grants essential
Follow CIS MCP Companion Guide Apr 20

Compliance Gap
Cowork NOT in Anthropic audit logs
Cowork NOT in Compliance API
OTel to SIEM is your only audit trail
Applies to ALL plan tiers incl Enterprise

CVE Reference

Two notable CVEs as of May 2026. Enforce version floors via MDM.

CVE	CVSS	Description	Fixed In	Required Action
CVE-2025-59536	8.7 HIGH	RCE via malicious hooks in project .claude/settings.json — code ran before trust dialogs appeared	v1.0.111 +	MDM version floor >= 1.0.111
CVE-2026-21852	5.3 MEDIUM	API key exfiltration via ANTHROPIC_BASE_URL override in managed settings	v2.0.65 +	MDM version floor >= 2.0.65 (current minimum)

Anthropic's Data Commitments

Customer prompts and responses are not used to train Anthropic's models by default. Retention is configurable. Anthropic's Trust Center (trust.anthropic.com) is the authoritative source for certifications

(SOC 2 Type II, ISO 27001), sub-processor list, and data handling policies. Direct your CISO and procurement team there — do not rely on secondary sources.

Six Governance Pillars

1. Identity Governance
SSO mandatory — no local accounts
SCIM for lifecycle management
anthropic-allowed-org-ids at network
MFA enforced at IdP level

2. Data Governance
Training opt-out confirmed active
Retention period configured
Data classification per connector
No RESTRICTED data in connectors

3. Access Governance
RBAC matrix documented + applied
Spend limits by group
Quarterly RBAC review cycle
Stale elevation removal process

4. Integration Governance
Connector register maintained
MCP allowlist via MDM
No self-service OAuth grants
Tier 3 servers require pen test

5. Observability Governance
OTel to SIEM for Cowork activity
Six minimum alert rules active
Named triage owners per alert
Monthly alert rule review

6. Operational Governance
BAU checklist with named owners
Incident response plan updated
Model update change management
Annual compliance review cycle

09 Next Steps with Fynaptix

Start Your Deployment

Every organisation's starting point is different. Whether you're scoping a first deployment, rescuing a stalled pilot, or scaling to organisation-wide adoption — Fynaptix has a structured path for you.

How Fynaptix Engages

AI Strategy & Readiness Assessment
For: Organisations at the start of their Claude journey
What: AI maturity audit, use case scoring, build-vs-buy matrix, compliance alignment
Output: Prioritised roadmap, architecture rec, investment case for internal stakeholders
Duration: 2-3 weeks

CoWork Enterprise Deployment
For: Organisations ready to deploy Cowork at scale
What: Full technical deployment — SSO, RBAC, MDM, connectors, MCP, SIEM + adoption programme
Output: Production deployment, as-built docs, 90-day adoption programme, team training
Duration: 6-10 weeks

Agentic System Design & Build
For: Teams with specific automation or AI product goals
What: Production agentic workflows, RAG, MCP server development, Claude API integration
Output: Working production system, runbooks, handover docs, BAU training
Duration: 6-16 weeks

Proof-of-Concept Programme
For: Teams needing executive buy-in first
What: 4-8 week time-boxed engagement, one specific use case end-to-end
Output: Working prototype, architecture decisions, cost projections, scaling roadmap
Duration: 4-8 weeks

Why Fynaptix

What We Bring	What This Means for You
Deep Claude expertise — practitioners building production systems daily	No learning curve on your time. Architecture decisions from experience, not theory.
Full-stack coverage — API, Cowork, Code, MCP, RAG, governance, training	One engagement partner for the entire Claude stack. No handoffs between specialists.
Governance from day one — every system built with the compliance layer	Your CISO and risk committee get documentation they can actually work with.
Adoption-focused delivery — technical build plus the onboarding that sticks	Deployments that drive real utilisation, not expensive shelfware.
Production-grade output — working infrastructure, not recommendations	Every engagement ends with deployed, documented systems your team can operate.

Ready to build the AI layer your business will run on?

hello@fynaptix.com · fynaptix.com · linkedin.com/company/fynaptix

This guide is produced by Fynaptix — an Australian enterprise Claude AI specialist based in Sydney. Fynaptix designs, builds, and operationalises production-grade Claude systems for enterprises across financial services, legal, healthcare, technology, and government.