

Executive Onboarding Package

Module 1: AI-CTO Leadership Foundations and Strategic Onboarding
NorthStar AI Hub Initiative | LiteLLM Enterprise | NorthStar Infrastructure (masked)

Submitted by	Luke Kilpatrick
Program	University of Waterloo, AI-CTO Program
Submission	Module 1 Executive Onboarding Package
Assignments Included	Assignment 1A: The Role of the AI-CTO Assignment 1B: First 90 Days Assignment 1C: Effective Communications and Executive Presence
Initiative	NorthStar AI Hub, built on LiteLLM Enterprise

AI-CTO Program | Module 1, Assignment 1A

The Role of the AI-CTO: Strategic Context and AI Initiative Definition

Submitted by: Luke Kilpatrick | University of Waterloo, AI-CTO Program

About this brief. NorthStar Infrastructure is a masked enterprise infrastructure customer based on patterns I have seen in the field. The governance challenge mirrors work I am currently helping organizations solve with LiteLLM Enterprise: giving the business one governed front door for AI across product, engineering, and internal use.

A1. The AI-CTO Context

NorthStar Infrastructure is a global infrastructure and hybrid cloud software company shifting from a hardware-led business to software and subscription revenue. Its customers are enterprise IT teams in financial services, government, healthcare, and manufacturing. AI adoption is complicated because the company historically sold to infrastructure administrators, not software architects or AI builders.

Organization Size. Approximately 6,000 employees, multi-billion dollar annual revenue, public, and operating across North America, EMEA, APJ, and India. The field organization includes more than 700 sales engineers.

My Role. Director, AI Engineering and Platform. I lead the team building NorthStar AI Hub, the governed front door for every LLM and agentic tool used internally or in product. The role is jointly accountable to the CIO for internal tooling, governance, cost, and productivity, and to the CTO for product AI, developer experience, and architecture standards.

Reporting Line. Solid line to the CIO, dotted line to the CTO. I manage a team of seven: platform lead, two platform engineers, AI security engineer, AI FinOps analyst, developer experience engineer, and program manager. I work closely with InfoSec, Legal, Procurement, CDO, Product, Field Enablement, and Developer Relations.

AI-CTO Archetype. Enterprise. NorthStar has enterprise scale, board scrutiny, mature governance, and an entrenched go-to-market motion. There is no formal AI-CTO today. My joint CIO/CTO mandate is the closest operating model, and one year-one outcome is to decide whether that should become a formal role or remain a joint structure.

Current AI Maturity. Fragmented. Teams use multiple LLM accounts, direct provider integrations, and reimbursed personal ChatGPT and Claude subscriptions. There is no shared evaluation framework, cost view, audit trail, approved model list, or consistent policy for what data may leave the company in a prompt.

A1. The AI Initiative: NorthStar AI Hub

Initiative Name. NorthStar AI Hub, the company's single governed front door for AI, built on LiteLLM Enterprise.

Initiative Type. Primary: AI risk and compliance. Secondary beneficiaries: customer operations, AI engineering productivity, and AI-assisted product features.

One-Sentence Description. NorthStar AI Hub is a self-hosted enterprise AI gateway that routes all internal and customer-facing LLM usage through one governed control plane for identity, model access, audit, cost, redaction, and incident response.

The Dual Mandate: One Gateway Platform, Two Audiences

NorthStar AI Hub serves both customer-facing product AI and internal employee AI on the same controls. In practice there will be two separate gateways one serving internal users and another serving our products. By using one platform for both we can save complexity and training. Product teams use AI Hub for embedded AI features, documentation assistance, code generation, and future agentic workflows in our products. Employees use it as the approved replacement for scattered personal ChatGPT and Claude accounts with controls to prevent data leaking into to public systems.

The model registry, audit log, guardrails, budgets, secret management, and compliance posture are similar for both systems. A single gateway software avoids duplicate platforms, reduces policy debate, and gives the company a common interface to govern AI.

What the Hub Actually Is

NorthStar AI Hub is a self-hosted LiteLLM Enterprise deployment in NorthStar's infrastructure. It provides one OpenAI-compatible endpoint for approved providers, including OpenAI, Anthropic, Azure OpenAI, Gemini, Bedrock, and internal models. It centralizes SSO, SCIM, scoped virtual keys, budgets, audit logs, redaction, moderation, request metadata, and observability. The operating rule is simple: no raw provider keys in code, no unapproved models, and no AI requests outside the governed path.

Board Lens: Risk Reduction

The board's priority is not immediate revenue acceleration. It is defensible control of AI before shadow use creates legal, IP, security, or customer-data exposure. Today, employees paste customer names, draft contracts, roadmaps, support tickets, and pre-release code into consumer AI tools that are not auditable. The Hub brings that activity inside the front door, where it can be observed, governed, and explained to customers and auditors.

Why Now

Three forces make this urgent. First, the controls now exist in a self-hosted gateway: SSO, SCIM, audit logs, budgets, redaction, moderation, model allowlists, and required metadata. Second, the cost of inaction is visible in reimbursed ChatGPT and Claude accounts, scattered provider keys, and unexplained OpenAI spend. Third, customers in regulated industries now ask where their data goes when NorthStar employees use AI. Today the company cannot answer consistently. With the Hub, it can.

The Political Constraint

The single biggest barrier is human, not technical: the CRO and the field organization already rely on fast, unmanaged consumer AI to accelerate deal cycles. Because the Hub introduces visible IT governance like SSO, data redaction, and approved models, the field will view this as friction that jeopardizes their ability to hit quota. To overcome this resistance and ensure the CRO becomes a co-sponsor, we must prove the official Hub is actually faster and better by grounding it in proprietary playbooks and securing pilot endorsements from respected Sales Engineers. Adoption will only succeed if the governed path becomes the easiest path.

A1. AI-CTO Swim Lane Clarity

Decision Domain	Who Owns It at NorthStar	Where My Boundary Sits (LiteLLM Enterprise)	Unresolved Tension
AI model deployment decisions	CTO Office owns model strategy and architecture. My team owns the model registry. Product teams choose from approved models. Both local and commercial models are used.	I own the registry entry: provider, inference contract, budget owner, and guardrail profile. Nothing becomes broadly available without an approved registry entry.	The exception process is still unclear. Product and DevRel want fast frontier-model trials; the CTO Office wants a tight default list. We need to wait for liteLLM support to use the model that came out yesterday.
Data access and governance	CDO owns classification and policy. InfoSec owns controls and DLP. Legal owns customer-data and IP commitments. My team enforces policy at the gateway.	I own redaction, blocked users and keywords, and per-team logging. The Hub enforces the policies set by CDO, Legal, and InfoSec.	Use of anonymized support tickets and internal docs for grounded answers is unresolved. Legal sees risk; the field sees high productivity value.
AI vendor selection	Procurement owns commercial terms. InfoSec owns security review. CTO Office owns architectural fit. My team owns functional fit and production cost.	I recommend vendors that sit behind or integrate with the Hub, including model providers, LiteLLM Enterprise, observability, and secret management.	The standard MSA does not yet cover AI-specific terms such as training-data use, model versioning, output indemnity, and gateway-provider roles.
AI incident response	InfoSec owns security incidents. Legal and Communications own external response. Product owns in-product incidents. My team owns Hub-level response.	I can disable a model, revoke a key, or take a downstream AI feature offline at the gateway. Hub audit logs become the canonical incident record.	The company lacks an AI-specific incident playbook for hallucinated guidance, prompt injection, data leakage, model deprecation, or model behavior failures.

Distributed Mandate: Where the Gaps Are

NorthStar does not have a defined AI-CTO. The mandate is split across the CTO Office for platform strategy, CIO for internal AI, CDO for data policy, Product for in-product AI, and my team for the gateway and controls. My joint reporting line closes the most damaging seam between product AI and internal AI, but not the broader policy gap.

The biggest gap is ownership of cross-surface AI policy. The Hub gives NorthStar one technical enforcement point, but policies for data classes, approved models, guardrails, audit retention, and exceptions still require one accountable owner. My Week 2 ninety-day plan will use the Hub as the forcing function to establish either a formal AI-CTO role or a standing AI Governance Council with my team as the operating arm.

Appendix: Why LiteLLM Enterprise

LiteLLM Enterprise matters because it turns NorthStar’s AI policies into enforceable platform controls. The appendix below focuses only on capabilities that support the swim lane decisions above.

NorthStar Need	LiteLLM Enterprise Capability	Executive Value
One endpoint for every model	OpenAI-compatible proxy for OpenAI, Anthropic, Azure OpenAI, Gemini, Bedrock, and internal models.	Teams build once. Provider changes become configuration choices, not migrations.
Identity tied to corporate SSO	SSO and SCIM with Okta, Azure AD, Google Workspace, and OIDC/SAML providers.	Users get the same login as Salesforce. Offboarding follows corporate identity.
Multi-tenant product and internal use	Organizations, teams, projects, and virtual keys with budgets, model allowlists, rate limits, and isolated logs.	One control plane, separate tenants, separate budgets, and separate policies.
Cost control across providers	Per-team and per-model budgets, alerts, temporary increases, spend reports, and tag tracking.	Finance gets a consolidated AI spend view before overruns become surprises.
Governed data egress	Secret and PII redaction, moderation, banned keywords, blocked users, and required request parameters.	Sensitive customer, contract, roadmap, and code data can be stripped or blocked.
Traceability	Audit logs with retention and per-team logging to Langfuse, LangSmith, Arize, GCS, or Azure Blob.	Incidents can be tied to person, application, team, model, and cost center.
Secrets out of code	Integrations with AWS KMS, Secrets Manager, Azure Key Vault, Google KMS, HashiCorp Vault, and CyberArk.	Provider keys stay out of repos and laptops. Rotation follows InfoSec policy.
Self-hosted, in our cloud	Docker deployment in NorthStar cloud accounts, with AWS and Azure Marketplace procurement.	Data stays inside known controls. Procurement can land in weeks.
Production support	Enterprise support with Sev 0/Sev 1 response targets, shared Slack or Teams support, and patch timelines.	The support model is defensible when the Hub is on a critical workflow path.

AI-CTO Program | Module 1, Assignment 1B

First 90 Days — NorthStar AI Hub

Submitted by: Luke Kilpatrick | University of Waterloo, AI-CTO Program |

Building on Week 1. This plan extends Assignment 1A (The Role of the AI-CTO). The organization is NorthStar Infrastructure — a representative composite enterprise infrastructure customer. The initiative is NorthStar AI Hub, the company’s single front door for AI, built on LiteLLM Enterprise. My role is Director, AI Engineering and Platform, reporting jointly to the CIO and CTO. The political constraint identified in 1A — the CRO and the field organization — is the single biggest variable this 90-day plan is designed to navigate.

Executive Summary

This 90-day plan turns NorthStar AI Hub from concept into governed production. The arc is deliberate: in Month 1 we earn the right to ship by proving value with three internal pilot teams and an approved LiteLLM Enterprise PoC. In Month 2 we move from PoC to production, expand to a Field pilot co-sponsored by the CRO, and stand up the AI Governance Council that will own cross-surface AI policy from that point forward. In Month 3 we reach production scale with measurable shadow-AI displacement, the Council operating on a regular cadence, and the Hub embedded in NorthStar’s security and procurement playbooks. Every milestone in the table that follows is observable: a tenant in production, a contract signed, a percentage of spend on the Hub, a charter ratified.

Part A — The 90-Day Plan

The plan is organized along five tracks adapted from the Week 2 framework: People, Governance and Ethics, Infrastructure and Data, Quick Wins, and Strategic Projects. Each track has observable milestones at Day 30 (Understand), Day 60 (Initiate), and Day 90 (Perform), plus an Outcome and OKR. The table is the operational core of the brief and is meant to be referenced.

Track	Tools	Understand — Day 30	Initiate — Day 60	Perform — Day 90	Outcome / OKR
People	Skills gap analysis. AI literacy survey. SE Champion identification. Org design review for AI Engineering & Platform team.	Skills gap analysis complete. AI literacy baseline measured across Field, Product, Marketing, Engineering. Three SE Champions identified for Field pilot.	Two AI Platform Engineers and one AI Security Engineer hired or backfilled. AI literacy training launched for first 200 employees. SE Champions on Hub.	Team of seven fully staffed. 600+ employees through AI literacy module. Champion network expanded to 12 SEs across regions, with measurable influence.	Outcome: AI capability lift across the org. OKR: 80% of Field, Product, Marketing, Engineering complete AI literacy by Day 90.
Governance & Ethics	AI Governance Council charter. Acceptable Use Policy. Model Risk Tiering. Ethics Charter aligned to NorthStar values.	Governance gaps documented. Draft AI Acceptable Use Policy circulated to CDO, Legal, InfoSec, CTO Office. Council membership proposed.	AI Governance Council constituted with first meeting held. Acceptable Use Policy ratified. Model Risk Tiering (High/Medium/Low) live in the Hub registry.	Council operating on monthly cadence with published minutes. Ethics Charter ratified and communicated. First quarterly governance review delivered.	Outcome: Single accountable owner for cross-surface AI policy. OKR: 100% of new model deployments pass through review.
Infrastructure & Data	LiteLLM Enterprise PoC. AI Maturity Model assessment. Data classification workshop. Secret manager integration plan.	LiteLLM Enterprise PoC approved. Three internal pilot teams identified (Marketing content, DevRel docs, Engineering productivity). Maturity baseline: Level 2 (Repeatable).	LiteLLM Enterprise contracted via AWS Marketplace. Hub in production for three internal tenants with SSO, virtual keys, audit logs, and PII redaction live. Vault-based secret rotation enabled.	Hub serving five internal tenants plus one Field pilot. 95%+ of internal LLM spend routed through the Hub. Maturity advanced to Level 3 (Defined).	Outcome: Governed AI infrastructure in production. OKR: 95% of internal LLM API calls route through the Hub by Day 90.
Quick Wins	Quick-win selection criteria. Shadow-AI expense audit. Single-tenant deployment template. Marketplace procurement template.	Shadow-AI expense audit complete: actual annual spend on personal ChatGPT/Claude accounts and uncoordinated API keys quantified to the dollar.	First quick win shipped: Marketing tenant live, replacing 40+ personal ChatGPT subscriptions with a governed Hub team. Soft-budget alerts active.	Three quick wins concluded: Marketing, DevRel, and Engineering tenants in production. Reimbursable personal-AI expenses down 70%+ in those orgs. Wins communicated company-wide.	Outcome: Visible, measurable cost-and-risk wins that build credibility for Strategic Projects. OKR: 70%+ reduction in reimbursable personal-AI expenses in pilot orgs.
Strategic Projects	Map of AI dependencies. Field pilot design. CRO co-sponsorship plan. Customer-facing product AI roadmap.	Map of AI dependencies delivered to CIO and CTO. CRO briefed twice; co-sponsorship plan drafted. Customer-facing product AI roadmap aligned with CTO Office.	Field pilot launched with CRO co-sponsor and three respected SEs. First customer-facing product AI feature in staging on the Hub. Standard AI MSA addendum drafted with Legal.	Field pilot expanded from three to twelve SEs across regions, with documented productivity lift. First product AI feature in limited GA on the Hub. Year-2 strategic plan approved.	Outcome: Hub adopted by the field; product AI shipping on governed infrastructure. OKR: Field pilot reports ≥1 hour/SE/week saved with 90%+ satisfaction.

Sequencing Logic — Why This Order

Three sequencing choices in this plan. First, internal use cases lead, customer-facing use cases in products follow. Marketing, DevRel, and Engineering pilots in Month 1 give us battle-tested controls before we put a single customer-facing feature on the Hub.

The reverse sequence — ship product AI first, govern internal use later, is how most enterprises end up with two parallel, half-governed systems on different platforms.

Second, the Field pilot waits until Day 60. Trying to convert the field on Day 1 is a losing fight. By Day 60 we have shipped governed Marketing and DevRel tenants, three SE Champions are using the Hub daily, and the CRO has been briefed twice with concrete data on shadow-AI spend in his own organization. The pitch is no longer abstract; it is “your SEs are already telling you this is faster than ChatGPT and has more flexibility in model choice.”

Third, the AI Governance Council is constituted at Day 60, not Day 30. Standing up a Council before the Hub has shipped anything produces a governance theater that the CRO and the field will rightly ignore. Standing it up the moment the first quick win lands means the Council has real decisions to make from its first meeting, model approvals, data classifications, incident response and earns its authority by exercising it.

Part B — Stakeholder Matrix

Twelve stakeholders matter for the success of NorthStar AI Hub. They are listed by the role as it drives the position regardless of who currently holds it. The four-quadrant matrix that follows positions each stakeholder by power (ability to influence the Hub’s success) and interest (degree to which the Hub affects their work). The detail table after that explains, for each, what they care about, what they fear, and what would move them from neutral to supportive.

Stakeholder Grid

Stakeholder	Role	Interest in NorthStar AI Hub
Chief Information Officer	Co-sponsor; owns internal AI tooling, employee productivity, governance, cost.	Wants a governed, defensible internal AI footprint and a unified cost view across the company. Sees the Hub as the cleanest path to both.
Chief Technology Officer	Co-sponsor; owns customer-facing product AI, architectural standards, foundation model strategy.	Wants product AI features to ship on infrastructure that does not create lock-in or audit risk. Cares about which models enter the registry and on what terms.
Chief Revenue Officer	The single biggest political variable; owns the field organization and its comp plan.	Wants quota attainment this quarter. Currently sees the Hub as friction. Will become a co-sponsor only when his SEs tell him the Hub is faster and better than what they have.
Chief Data Officer	Owns data classification and policy that the Hub enforces.	Wants explicit, defensible answers to “what data may leave the company in a prompt.” Currently providing those answers case-by-case rather than systemically.
Chief Information Security Officer	Owns access controls, DLP, secret management, incident response.	Wants no surprises. The Hub is welcome if it reduces the InfoSec attack surface (no raw API keys in code, audit logs, SSO/SCIM); resisted if it adds one.
General Counsel	Owns customer data commitments, IP exposure, and the AI vendor contract addendum.	Wants contracts that address training-data use, model versioning, and indemnification.
Chief Financial Officer	Owns budget, FinOps, and the answer to “how much are we spending on AI.”	Wants a single line item for AI spend, not a hundred reimbursed personal subscriptions. The Hub’s per-team and per-model budgets are the answer.
VP of Field Engineering / SE Leadership	The CRO’s technical lieutenant; manages the SE community.	Wants his SEs effective in front of customers. The most influential individual for Field pilot success — not the CRO, but the leader of the SE community.
SE Champions (3–12 named SEs)	Individual contributors with disproportionate informal influence over peers.	Want tools that make them faster and better in front of customers. Their endorsement of the Hub will move the field.
Procurement	Owns commercial terms and the standard MSA.	Wants standard terms and predictable cycle times. Will be neutral-to-supportive if AWS Marketplace procurement is on the table; resistant if asked to write a custom MSA.
Product VPs (3–4 individuals)	Own roadmap for product features that depend on the Hub for AI.	Want freedom to evaluate frontier models the day they ship. Will resist any registry process they perceive as slowing them down; supportive if exception path is fast.
Board AI Risk Committee	Receives quarterly governance reporting.	Wants confidence that NorthStar can answer customer security questionnaires about AI data handling. The Hub is the answer; their interest is in evidence it is working.

Power – Interest Matrix

The matrix below positions each stakeholder by their power to influence the Hub’s success and their interest in its outcome. The four-quadrant model dictates the engagement strategy: Manage Closely (highest priority), Keep Satisfied (active engagement, less frequent), Keep Informed (regular updates, no negotiation), Monitor (light-touch awareness).

Keep Satisfied (High Power, Low Interest) Chief Financial Officer General Counsel Procurement Chief Data Officer Board AI Risk Committee	Manage Closely (High Power, High Interest) Chief Information Officer ★ Chief Technology Officer ★ Chief Revenue Officer Chief Information Security Officer
Monitor (Low Power, Low Interest) Product Management Software Developers	Keep Informed (Low Power, High Interest) VP Field Engineering / SE Leadership SE Champions Product VPs

★ Co-sponsors.

Stakeholder Detail — Priorities, Concerns, and What Moves Them

The five stakeholders below — the full Manage Closely quadrant plus the most influential individual in Keep Informed — are the ones whose movement determines whether this plan succeeds. Each is profiled against three questions: what they prioritize, what they fear, and what concrete action moves them from neutral to supportive.

Stakeholder	Priorities	Concerns	What Moves Them to Supportive
Chief Information Officer (Co-sponsor)	A governed internal AI footprint. A unified cost view. A defensible answer to security questionnaires. A working partnership with the CTO Office on cross-surface AI.	That the Hub becomes shelfware if the CRO does not adopt it. That the team is under-staffed for the governance load it is taking on.	Already supportive. Maintain by delivering visible, measurable wins on the published cadence and never surprising him in a board setting.
Chief Technology Officer (Co-sponsor)	Product AI features shipping on infrastructure that does not create vendor lock-in or audit risk. Tight control over what enters the model registry.	That a permissive registry creates architectural debt. That the joint reporting line dilutes his control over product AI.	Already supportive. Maintain by giving the CTO Office final sign-off on registry entries and bringing the first product AI feature to limited GA on the Hub by Day 90.
Chief Revenue Officer	Quota attainment this quarter. SE productivity. A field organization that is faster, not slower, than competitors’.	That the Hub introduces friction the field cannot absorb. That IT is taking away tools letting his team hit their number.	Three respected SEs telling him, unprompted, that the Hub is faster than ChatGPT for their workflow. Plus a Day 60 Field pilot with measurable time-saved data and a co-sponsor credit on the rollout announcement.
Chief Information Security Officer	Reduced attack surface. No surprises. Audit-ready logs. Managed secret rotation.	That the Hub becomes a new attack surface (gateway compromise, prompt injection, exfiltration via tool calls). That AI-specific incidents have no playbook.	A Day 60 demonstration that the Hub eliminates raw API keys from code, routes everything through SSO/SCIM, and produces audit logs InfoSec can ingest. Plus a jointly authored AI incident playbook by Day 90.
VP Field Engineering / SE Leadership	SEs effective in front of customers. Tools that work the way SEs actually work, not the way IT imagines.	That the Hub forces SEs through workflows they will route around. That a centrally mandated tool will degrade rather than improve customer-facing demos.	Co-design the Field pilot with him as the named lead, not a recipient. Let him pick the three SE Champions. Give his organization first access to the customer-grounding feature — the highest-leverage capability the Hub can offer the field.

The Move That Matters Most

If I get one thing right in 90 days, it is the Field pilot. The CIO partnership is already in place; the CTO partnership follows from architecture choices I control; the CFO, Legal, and InfoSec move with the wins. The CRO is the only stakeholder whose movement is genuinely uncertain, and the CRO does not move on the basis of an executive briefing. He moves when his own SEs tell him something is faster and better than what they have. The entire 90-day plan is structured to put a tool in the hands of three respected SEs by Day 60 and let them build the case for me. Everything else is in service of that single move.

Effective Communications and Executive Presence

NorthStar AI Hub Initiative | LiteLLM Enterprise | NorthStar Infrastructure (masked)

Submitted by: Luke Kilpatrick | University of Waterloo, AI-CTO Program | Class of 2026

Context Recap

This is the third progressive deliverable for Module 1. Week 1 defined NorthStar AI Hub, built on LiteLLM Enterprise, as the governed front door for internal and customer-facing AI. Week 2 translated that initiative into a 90-day plan and stakeholder map. Week 3 turns the same case into an executive communication plan for the single audience whose support determines field adoption: the Chief Revenue Officer.

The communication challenge is not to explain LiteLLM. It is to show the CRO that a governed AI Hub protects sales velocity rather than slowing it down. The field already uses consumer AI because it helps them move faster. The executive message must therefore position the Hub as the safer and faster official path, not as an IT control layer imposed after the fact.

Part A: Audience Alignment and Message Positioning

A1. Executive Audience and Trust Positioning

The selected audience is the Chief Revenue Officer. The CIO and CTO are aligned on the need for governance. Legal and the Chief Data Officer are gating stakeholders. The CRO is the political fulcrum because the field organization is where shadow AI is already useful, risky, and culturally difficult to displace.

Dimension	Response
Primary Audience	Chief Revenue Officer. Owns the global field organization, including Account Executives, approximately 700 Sales Engineers, field enablement priorities, and the operating rhythm that determines whether new tools are adopted or ignored.
What this audience prioritises	Quota attainment, deal velocity, and Sales Engineering capacity. The CRO will support tools that save SE time and increase confidence in customer-facing work. He will reject anything that feels like a governance tax, a training burden, or a workflow interruption.
Likely concerns or objections	First, IT friction: SSO, model allowlists, required metadata, and audit logs may sound slower than personal ChatGPT or Claude. Second, model quality: the field will worry the approved path is less capable. Third, autonomy: the CRO will resist any mandate that appears to take working tools away from his team.
Current stance	Resistant, but rationally so. He is not against AI. He is against an IT-led rollout that could slow his organization in the current quarter. His stance can shift if the pilot proves that the governed path improves speed, consistency, and customer-data safety.
What would make this audience approve the initiative	Three conditions must be true. His SE leaders co-design the pilot. The Hub is demonstrably faster because it is grounded in NorthStar playbooks, case studies, and approved patterns. The rollout starts small, with a clear Day 60 success metric and an explicit pause condition if the field does not prefer it.
Competence	I will show operational readiness: shadow-AI expense telemetry, the LiteLLM deployment model, provider and cost controls, initial guardrail design, and a pilot plan small enough to manage but meaningful enough to prove adoption.
Integrity	I will name the core objection directly: the field fears an IT friction tax. I will not overpromise adoption or revenue impact. I will commit that the pilot does not expand until the named SE cohort says the Hub saves time and improves their work.
Benevolence	I will frame the initiative as protection for the CRO's team first. The Hub reduces the risk that an SE causes a customer-data incident while trying to move quickly. It gives the field a sanctioned tool that works with their workflow rather than asking them to stop using AI.

A2. Audience Analysis: Six Questions

Question	Response
1. How much does my audience know about my topic?	The CRO understands AI at a board and vendor-pitch level. He does not need a technical lesson on gateways, model routing, or LiteLLM internals. He does need to understand one fact clearly: unmanaged AI use is already happening inside his organization, and the company currently lacks visibility into customer data, cost, model choice, and output quality.
2. How does my audience feel about my topic?	Wary, not hostile. He has seen productivity tools land as extra administration for the field. He will listen if the opening frames the issue as an existing field risk and a sales-productivity opportunity. He will tune out if the opening sounds like architecture, compliance, or platform standardization.
3. Will my audience be against me?	Initially yes, because the default assumption is that governance slows sales teams down. The path to persuasion is to agree with the concern, then reverse it: unmanaged AI is fast but unsafe and inconsistent; the Hub must be safer while also being faster because it is grounded in NorthStar-specific data.
4. What past events shape how this audience hears the message?	Past IT-led rollouts asked the field to change behavior without giving them a better workflow. That history matters. The presentation must show that the Hub is being launched like a product for the field, with SE champions, a clear success metric, and a commitment to pause if it does not create value.
5. What pre-presentation groundwork should happen?	Pre-wire the VP of Field Engineering. Identify three respected Senior SEs who can serve as pilot champions. Pull six months of expense data for reimbursed ChatGPT and Claude accounts. Confirm with Legal, CDO, InfoSec, and the CTO Office that the risk, data, and architecture framing will not be relitigated live.
6. Why am I making this presentation?	Because the AI mandate crosses CIO, CTO, Legal, Data, Product, and Field boundaries, but the adoption risk sits in the field. My role is to translate platform governance into a field adoption plan the CRO can support, sponsor, and defend.

A3. Purpose and Concise Executive Summary

Purpose. After hearing my presentation, the CRO will co-sign the NorthStar AI Hub charter as a joint sponsor, name three Senior Sales Engineers for the Month 1 pilot, and approve the Q1 pilot budget against a field-owned success metric.

Concise Executive Summary. NorthStar’s field teams are already using unmanaged consumer AI to draft customer emails, summarize calls, and generate technical material. The choice is not whether AI enters the sales workflow. It already has. The decision is whether NorthStar governs that footprint or leaves it scattered across personal accounts, corporate-card reimbursements, and unobservable model interactions.

NorthStar AI Hub, built on LiteLLM Enterprise, brings that activity inside a governed, SSO-enabled control plane with approved model access, audit logs, budget controls, redaction, and per-team visibility. The field benefit is not compliance for its own sake. The field benefit is a better AI tool: one connected to NorthStar’s approved playbooks, product patterns, and case studies, while reducing the risk that customer data leaks through a personal account.

The ask is deliberately narrow. Do not mandate global rollout. Start with three respected Senior SEs in a 60-day pilot. Measure demo-prep time, answer quality, and preference against personal tools. If the Hub is not faster and preferred, pause. If it is, the field becomes the on-ramp for governed AI adoption across NorthStar.