



PROACTIVEDATASYSTEMS · WHITEPAPER

AgenticOps

**What AI-run networks actually
mean for Indian enterprises.**

The pitch, translated. What is shipping, what is beta, what your telemetry has to look like first, and where your data goes when you ask the question.

AUGUST 2026

Author: Kunal Hans, Chief Technology Officer

THE ARGUMENT

AgenticOps Is Real. You Are Probably Not Ready for It.

Both halves of that sentence matter. The technology is not vapour, and the gap between it and most Indian enterprise networks is not about the AI at all.

Your CIO is being pitched this right now. Some version of the pitch is accurate. Cisco has built network-specific models, embedded them across the portfolio, and started shipping agents that read telemetry, reason about it and act. That is a real change in how networks are operated, and in five years the argument will look settled.

The problem is what sits underneath.

Agentic operations run on telemetry. Continuous, structured, complete telemetry, from devices that are running supported software, configured consistently, and named in a way a machine can reason about. An agent asked why the Pune branch is slow can only answer if the Pune branch has been reporting, in a format the model recognises, for long enough to establish what normal looks like.

Most Indian enterprise estates do not look like that. They have three generations of hardware, a management platform running two releases behind, a naming convention that changed when the previous network manager left, and assurance switched on at four sites out of forty because that was where the pilot happened.

None of that is a criticism. It is what a network looks like after fifteen years of doing more with less. But it means the honest answer to "should we buy the AI thing" is usually not yes or no. It is: not yet, and here is the specific list of what has to be true first.

The constraint is not the intelligence. It is the data you are feeding it, and most estates are not feeding it anything.

There is a second constraint, and it is particular to India. Cisco's documentation is clear that content and systems information associated with your use of Catalyst Center, including anything you type into the AI Assistant, is primarily processed and stored in the United States, regardless of where you are. For a regulated entity working under the DPDP Rules and the RBI's expectations, that is not a footnote. It is a conversation with your compliance function that has to happen before the technical one.

This paper takes the pitch seriously and takes it apart. What the components actually are. What has shipped, what is in beta, and what is a roadmap slide. What it genuinely requires. What changes for a team of six people running forty sites. And a candid readiness test, because the most useful thing anyone can tell you about this technology in 2026 is whether you are in a position to use it.

IF YOU READ ONE PAGE

The Paper in One Page

1. AgenticOps is an operating model, not a product.

Cisco's framing combines purpose-built models, agents that reason and act, and a workspace where humans and agents work on the same problem.

2. The models are the genuinely novel part.

A Deep Network Model, a Foundation Security Model and a Time Series Model, trained on Cisco's own telemetry corpus. A general-purpose language model does not know what a TLOC is.

3. Delivery is staged across 2026, not finished.

Agentic capabilities for campus, branch and industrial began rolling out in February 2026. Cisco Cloud Control entered controlled availability in the United States on 2 June 2026 with global availability planned for July. Several capabilities are beta, and Digital Twin went to alpha in July.

4. The entry point for most people is the Catalyst Center AI Assistant.

It is available at every licence tier, but what it can do follows the tier of the underlying feature, so Advantage-tier analytics stay Advantage-tier.

5. Cloud connectivity is not negotiable.

The AI Assistant requires an active connection to Cisco's cloud. An air-gapped Catalyst Center gets none of this.

6. Your telemetry goes to the United States.

Including your prompts. Settle that with compliance before you settle it with the network team.

7. The real prerequisite is boring.

Supported software, consistent configuration, assurance enabled everywhere rather than in a pilot, and naming a machine can reason about.

8. For a lean team the value is triage, not replacement.

The gain is in the first twenty minutes of an incident, which is where small teams lose most of their time.

9. The honest answer for most Indian estates in 2026 is "not yet".

Chapter 10 is the list of conditions that make that a yes.

10. Doing the prerequisites is worth it regardless.

Every item on the readiness list improves your network whether or not you ever switch an agent on.

CONTENTS

What Is Inside

OPENING

[The Argument: AgenticOps Is Real. You Are Probably Not Ready for It.](#)

[The Paper in One Page](#)

PART ONE · WHAT IT ACTUALLY IS

[1. The Pitch, Translated](#)

[2. The Three Layers: Models, Agents, Workspace](#)

[3. Shipping, Beta, or Slide: A Status Table](#)

PART TWO · WHAT IT ACTUALLY REQUIRES

[4. Telemetry, and the Data Hygiene Nobody Costs](#)

[5. Licences, Cloud Connection and Catalyst Center](#)

[6. The India Question: Where Does Your Telemetry Go?](#)

PART THREE · WHAT IT CHANGES

[7. What a Team of Six Actually Gets](#)

[8. What It Does Not Replace](#)

[9. The Economics, Honestly](#)

PART FOUR · ARE YOU READY?

[10. You Are Not Ready for This Yet If...](#)

[11. The Readiness Scorecard](#)

[12. What to Do in the Next Twelve Months Instead](#)

REFERENCE

[Frequently Asked Questions](#)

[Sources, Disclaimer and About](#)

01

PART ONE

What It Actually Is

The pitch translated into plain description, the three technical layers underneath it, and an honest status table separating what has shipped from what has not.

- [Chapter 1 · The Pitch, Translated](#)
- [Chapter 2 · The Three Layers](#)
- [Chapter 3 · Shipping, Beta, or Slide](#)

CHAPTER ONE

The Pitch, Translated

AgenticOps is Cisco's name for an operating model in which software agents read telemetry, reason about what it means, take action, and learn from the result, with humans supervising rather than performing each step.

Strip the language back and there are three claims.

The first is that a model trained specifically on network data understands networks better than a general-purpose one. That is straightforwardly true. A general language model can explain what BGP is. It cannot look at four hours of telemetry from your campus and tell you which of nine correlated symptoms is the cause.

The second is that once a system can reason about network state, it can also act on it: raise the change, validate it against the live topology, work out the blast radius, and either execute or hand a decision to a human with the evidence attached. Cisco describes this as risk-aware agentic assessment, and it is the part that separates an assistant from an agent.

The third claim is about people. Networking, security and observability have historically been three teams looking at three consoles during the same incident. A shared workspace where all three, plus the agents, work on one problem is a change in how outages get resolved rather than a change in the technology being resolved.

What is genuinely new here

Automation in networking is not new. Templates, scripts, intent-based configuration and assurance dashboards have existed for years, and most of them under-delivered for a reason worth naming: they automated the doing, not the deciding. You still had to know what to do.

The agentic claim is that the deciding is now partly automatable. That is a real distinction, and it is why this wave deserves more attention than the previous one, even from people who found the previous one disappointing.

What the pitch tends to leave out

Three things, consistently.

That the models need telemetry, and that telemetry is a project. That most of the interesting capability requires a cloud connection and a licence tier. And that the data leaves the country.

None of those are hidden. All of them are in Cisco's documentation. They are simply not what gets presented in a forty-minute briefing to a CIO.

CHAPTER TWO

The Three Layers: Models, Agents, Workspace

It helps to separate what is a model, what is an agent and what is a user interface. Vendors and press coverage use the names interchangeably, and they are not the same thing.

Layer one: the models

Cisco has embedded three purpose-built models: a Deep Network Model, a Foundation Security Model and a Time Series Model, trained on its own telemetry corpus. This is the genuinely differentiated part of the story, and it is difficult for anyone without a large installed base to replicate, because the moat is the data rather than the algorithm.

A time series model matters more than it sounds. Most network problems are deviations from a pattern rather than absolute failures. Knowing that this access point's client count is unusual for a Tuesday afternoon in this building is a different capability from knowing it is above a threshold, and thresholds are how most monitoring still works.

Layer two: the agents

Agents are what use the models. Cisco has described risk-aware assessments that validate proposed changes against live topology, configuration and telemetry, including the blast radius of a change, and agentic workflow creation that produces deterministic, repeatable automations rather than one-off answers.

That word deterministic is doing important work. An agent that produces a different answer each time it is asked is a demonstration. An agent that produces a verifiable workflow you can inspect, approve and run repeatedly is an operations tool. The distinction is the difference between something you show a board and something you let near production.

Layer three: the workspace

Cisco AI Canvas is the generative workspace built on the Deep Network Model, where network, security, cloud and observability people work on one incident with the agents in the room. Cisco Cloud Control is the wider platform bringing cross-domain telemetry, the models and the agents under one identity and control plane.

For most enterprises in India, though, the first contact with any of this will be neither of those. It will be the AI Assistant inside Catalyst Center, which is where AgenticOps reaches an on-premises network.

THE PRACTICAL POINT

If someone offers to show you AgenticOps, ask which layer they are demonstrating. A model demo is impressive and tells you nothing about your network. An agent acting on your own telemetry is the only demonstration worth scheduling a room for.

CHAPTER THREE

Shipping, Beta, or Slide: A Status Table

This is the table the pitch does not include. Dates and states as publicly described in 2026. Verify current status before you plan around any line of it, because this is moving quickly in both directions.

Capability	State in 2026	What it means for you
Catalyst Center AI Assistant	Available	The realistic starting point. Natural-language troubleshooting and root cause on an onpremises network.
AI Network Analytics	Available, Advantage tier	Baselining and anomaly detection. Requires cloud connectivity. Data is deidentified before it leaves.
Agentic capabilities for campus, branch, industrial	Rolling out from Feb 2026	Staged. Ask specifically which capabilities are live for your platforms and releases.
Cisco Cloud Control	Controlled availability, US, from 2 Jun 2026. Global availability planned Jul 2026	Not something to design an Indian operating model around today. Confirm regional availability.
Agentic Actions, Deep Reasoning, Expanded Experience Metrics, Multicloud Fabric	Beta from Jun 2026	Beta means beta. Suitable for evaluation, not for the runbook.
Digital Twin	Alpha from Jul 2026	Interesting, and years from being how you validate changes.
Marketplace and further AgenticOps capabilities	Later in 2026	Roadmap. Keep it out of any business case.

HOW TO READ A ROADMAP IN A PROPOSAL

If a capability you are being sold on is in beta or alpha, it should not appear in the justification for spending money this year. Ask for the proposal to be split: what you can use on the day of handover, and what you are hoping for later. Sign against the first list.

02

PART TWO

What It Actually Requires

Three chapters on the prerequisites: the telemetry and data hygiene nobody puts in the business case, the licence and connectivity position, and the question an Indian regulated entity has to answer first.

- Chapter 4 · Telemetry, and the Data Hygiene Nobody Costs
- Chapter 5 · Licences, Cloud Connection and Catalyst Center
- Chapter 6 · The India Question

CHAPTER FOUR

Telemetry, and the Data Hygiene Nobody Costs

An agent is only as good as what it can see. AgenticOps correlates telemetry across device, network and application layers to isolate an issue, which means the quality of the isolation is bounded by the completeness of the telemetry.

This is the part of the project that has no demo.

Ask yourself a simple question about your own estate. If a user in the Hyderabad office reports that video calls have been poor since Monday, is there a system that already knows what Monday looked like at that site, at that time of day, for that access point, compared with the previous six Mondays?

For most Indian enterprise networks the answer is no. Not because the capability does not exist, but because assurance was enabled at the sites in the pilot and never extended, or because half the estate runs software old enough that the telemetry it emits is thinner than the model expects.

The four things that have to be true

Coverage. Telemetry from all of the estate, not the part where the proof of concept ran. An agent reasoning across a partial view will reach confident conclusions about the wrong thing, which is worse than reaching none.

Continuity. Anomaly detection needs history. A model that has watched your network for two weeks does not know what your quarter-end looks like. Give it a baseline before you judge it.

Consistency. Devices configured differently produce telemetry that is difficult to compare. The value of standardising your configuration was always real. Agentic operations simply make the cost of not doing it visible.

Naming. This is the one people laugh at and then lose three months to. If your sites are named a mixture of city names, project codes, the name of the integrator who installed them and, in two cases, an engineer's initials, then a question phrased as "show me every branch in the west region" cannot be answered. Not because the model is weak, but because the answer does not exist in the data.

Nobody puts a naming convention in an AI business case. It is frequently the longest task in the project.

What this means for sequencing

The prerequisites are not a tax on the AI project. They are the network improvements you have been deferring, with a deadline attached.

Complete assurance coverage improves your mean time to resolution on its own. Consistent configuration reduces outages on its own. A clean naming and site hierarchy makes every report and every audit easier on its own. If the agentic capability disappointed you entirely, you would still be ahead.

That is the argument to take to a finance director, and it is a considerably stronger one than a demonstration.

CHAPTER FIVE

Licences, Cloud Connection and Catalyst Center

Two practical gates sit between an Indian enterprise and the AI Assistant in Catalyst Center: what your licences entitle you to, and whether your management platform is allowed to talk to Cisco's cloud.

The licence position is more subtle than it looks

The AI Assistant itself is available across licence tiers. What it can actually do follows the tier of the underlying feature it is drawing on. So the Assistant is not gated, but the capability behind any given answer may be.

AI Network Analytics, which provides the baselining and anomaly detection that makes most of this useful, is part of the Advantage software licence. If your estate is on Essentials, the Assistant will still answer questions. It will not have the analytics underneath that make the answers interesting.

This is worth checking device by device rather than assuming, because mixed estates are normal. A campus with Advantage at the core and Essentials at the access layer produces exactly the partial-visibility problem described in the previous chapter.

The cloud connection is compulsory

Active cloud connectivity is a prerequisite for the AI Assistant to function. Catalyst Center is registered to Cisco's cloud through System, Settings, External Services, using a cisco.com account.

Read that sentence again if your Catalyst Center currently has no path to the internet, which is a common and deliberate design in Indian BFSI, defence-adjacent manufacturing and government estates. In that architecture, none of this is available to you today. Cisco has indicated that on-premises deployment options are coming for customers with data sovereignty requirements, which is the thing to track if this describes you.

Gate	What to check
Catalyst Center release	Whether your current release supports the capabilities being pitched. Several are release-gated.
Licence tier, per device	Advantage where the analytics matter. Audit the estate rather than assuming uniformity.
Cloud registration	Whether policy permits it. This is a security and compliance decision, not a network one.
Device software	Whether platforms are on releases that emit the telemetry the models expect.
Support status	Anything past its last date of support will not be getting there.

CHAPTER SIX

The India Question: Where Does Your Telemetry Go?

Cisco's documentation states that customer content and customer systems information associated with your use of Catalyst Center is primarily processed and stored in the United States, regardless of your physical location. That includes the prompts you type into the AI Assistant and the deployment data it analyses to answer them.

This is the paragraph that should be in every proposal you receive and is in none of them. It is not a scandal and it is not hidden. It is how a great deal of cloud-delivered software works, and Cisco documents it clearly. But it converts a technical evaluation into a compliance conversation, and in India that conversation now has statutory weight.

Why it matters more than it did two years ago

The DPDP Rules 2025 place obligations on data fiduciaries around reasonable security safeguards and the handling of personal data, with the compliance deadline of 13 May 2027 now inside most organisations' planning horizon. Network telemetry is not obviously personal data, and in many cases it is not. But it is not obviously not, either.

Consider what assurance data contains. Device identifiers, user names where identity services are integrated, MAC addresses, location by access point, connection times, application usage patterns. A record showing which employee connected to which access point at what time, for how long, is closer to personal data than most network teams instinctively assume.

Regulated entities have a second layer. The RBI's IT governance expectations for banks, NBFCs and other supervised entities cover outsourcing, data localisation and third-party risk, and telemetry leaving the country under a cloud registration is squarely within the scope of those questions.

THE QUESTION TO ASK FIRST

Before the technical evaluation, ask your compliance or legal function one question: are we permitted to send network telemetry and operator prompts to a service that processes and stores them in the United States?

If the answer is no, or is a six-month review, then the sequencing of your entire AgenticOps programme changes and it is far better to know that in week one.

What good practice looks like

Get the data flows documented before the pilot, not after. Which categories of data leave, in what form, whether deidentification is applied, where processing happens and under what contractual terms. Cisco applies deidentification to network event data sent to the cloud for AI Network Analytics, and knowing precisely what that covers is worth an hour of somebody's time.

Then have the conversation once, properly, and record the outcome. Organisations that skip this end up having it repeatedly, informally, in the middle of a deployment, with a different answer each time.

For estates where the answer is a firm no, the honest position is that agentic operations are a 2027 or 2028 conversation, pending on-premises options. That is not a reason to do nothing. It is a reason to spend the intervening period on the prerequisites, which are valuable on their own.

03

PART THREE

What It Changes

The honest version of the benefit case for a small team running a large estate, the things this does not touch, and what the economics look like when the prerequisites are costed properly.

- [Chapter 7 · What a Team of Six Actually Gets](#)
- [Chapter 8 · What It Does Not Replace](#)
- [Chapter 9 · The Economics, Honestly](#)

CHAPTER SEVEN

What a Team of Six Actually Gets

The typical Indian enterprise network team is small relative to the estate it runs. Six to twelve people covering forty to two hundred sites is normal. For that team the value of agentic operations is concentrated in one place, and it is not where the marketing points.

It is the first twenty minutes of an incident.

Think about how those twenty minutes currently go. A ticket arrives saying the Chennai office is slow. Somebody opens the assurance dashboard. Somebody else checks whether a change went in last night. A third person asks the application team whether anything moved. Twenty minutes later the team has established what is not wrong, which is genuinely useful and enormously expensive in a team of six.

An agent that correlates across device, network and application telemetry and arrives with a ranked set of probable causes does not fix the problem. It removes the search. For a small team, the search is the cost.

The second gain: the question nobody has time to ask

Small teams operate reactively because they have no choice. Nobody has an afternoon to ask which access points have been quietly degrading for six weeks, or which branch has drifted furthest from the standard configuration, because there is always a ticket.

Natural-language interrogation of telemetry lowers the cost of asking that kind of question from an afternoon to a sentence. Whether your team then acts on the answer is a management question rather than a technology one, but the barrier moves.

The third gain: change confidence

Risk-aware assessment validates a proposed change against live topology, configuration and telemetry, and identifies the blast radius. In a small team, change confidence is usually a function of who is available to review it. Making that assessment consistent rather than dependent on which engineer is on shift is a real improvement in operational risk, and it is the capability least discussed in the pitch.

WHERE THE VALUE IS NOT

It is not in headcount reduction. Any proposal that justifies this spend by removing people from a team of six is describing an outcome nobody has achieved, on a network that will still need those six people to run the changes the agents recommend.

CHAPTER EIGHT

What It Does Not Replace

A short chapter, because the list is short and clear, and because being clear about it is how you keep the credibility to make the positive case.

Design. An agent will not tell you whether your campus should be a fabric or a traditional three-tier network, or whether a plant needs a separate industrial zone. Those are judgements about risk, cost, skills and the next seven years. They are made once and lived with.

Physical reality. No model resolves a fibre cut, a failed power supply in a plant room, or an ISP whose last mile floods every monsoon. The most common causes of downtime in Indian branch networks remain stubbornly physical.

Accountability. When a change goes wrong at two in the morning, somebody has to decide to roll back and somebody has to answer for it. An agent's recommendation does not carry accountability, and any operating model that pretends otherwise will fail its first serious test.

The skills of your team. This is the one that gets lost. Reviewing an agent's reasoning requires more expertise than following a runbook, not less. A team that cannot tell when the answer is wrong is in a more dangerous position with these tools than without them.

Getting the estate supportable. Nothing here fixes hardware past its last date of support, software two releases behind, or a licence position that does not cover what you have deployed.

CHAPTER NINE

The Economics, Honestly

Most business cases for this technology compare the licence cost with a projected saving in engineer hours. Both halves of that comparison are usually wrong, and the direction of the error is not the one you would expect.

The licence cost is understated, because it is quoted as the incremental cost of the AI capability rather than the cost of getting the estate to a state where the capability functions. The saving is overstated, because it is modelled on incidents resolved rather than on the portion of an incident that is actually search.

Here is a more honest frame.

Cost line	What people forget
Licence uplift to Advantage	Across every device that needs to be visible, not the ones already there.
Software upgrades	Platforms have to be on releases that emit what the models expect. Some of that is a project.
Hardware replacement	Anything past support is not going anywhere. This is often the largest line and it is rarely in the AI business case.
Assurance rollout	Extending coverage from the pilot sites to the whole estate.
Standardisation and naming	Weeks of unglamorous work with no demo at the end.
Compliance review	Real elapsed time, particularly for regulated entities.
Training	Reviewing agent reasoning is a skill. Budget for it.

Now the other side of the ledger, and this is the part that rescues the case. Every one of those cost lines buys you something whether or not the agents ever earn their keep. Supported hardware, current software, full assurance coverage, standard configuration and a documented data-flow position are the definition of a well-run network. You are not buying AI readiness. You are buying the thing you should have been buying anyway, with a reason to get it approved.

Build the case on the network you get, not on the agents you hope for. Then the agents are upside rather than justification.

04

PART FOUR

Are You Ready?

The candid version. Eleven conditions that mean the answer is not yet, a scorecard you can run in an afternoon, and the twelve-month programme that turns a no into a yes.

- [Chapter 10 · You Are Not Ready for This Yet If...](#)
- [Chapter 11 · The Readiness Scorecard](#)
- [Chapter 12 · What to Do in the Next Twelve Months Instead](#)

CHAPTER TEN

You Are Not Ready for This Yet If...

Any one of these makes an agentic operations programme premature. Two or more makes it a distraction. This is the chapter your vendor will not write for you.

1. Part of your estate is past its last date of support.

It will not run current software, will not emit modern telemetry, and cannot be patched. Fix this first. It is an audit finding independent of any AI question.

2. Assurance is enabled at some sites and not others.

An agent reasoning over a partial estate gives confident answers about an incomplete picture, which is worse than none.

3. Your licence position is mixed and undocumented.

If you cannot say which devices carry Advantage today, you cannot predict what the Assistant will be able to do.

4. Catalyst Center has no path to Cisco's cloud, and policy will not permit one.

This is a hard stop for the current generation. Track the on-premises options rather than fighting it.

5. Compliance has not ruled on United States processing.

Ask before the pilot. A negative answer found in month four wastes a quarter.

6. Your naming and site hierarchy are inconsistent.

Questions phrased by region or function cannot be answered if the data does not carry those attributes.

7. Configuration has drifted and nobody has a standard.

Comparing sites is how anomaly detection works. If every site is bespoke, everything looks anomalous.

8. You have less than a quarter of telemetry history.

Baselines take time. Two weeks of data judges the technology unfairly and expensively.

9. Nobody can name your top five recurring incidents.

If you do not know where the time goes, you cannot evaluate whether anything saved it.

10. Change management is informal.

Risk-aware change assessment is worth little where changes go in without records. Fix the process before automating around it.

11. The business case depends on capabilities in beta.

Build the case on what works on handover day. Everything else is upside.

WHY WE ARE TELLING YOU THIS

Because selling you something you cannot use costs us more than the sale is worth. The prerequisites in this list are work we would rather do with you now than explain to you in month six of a programme that is not going to deliver.

CHAPTER ELEVEN

The Readiness Scorecard

Score each row honestly. Two points where the statement is fully true across the whole estate, one where it is partly true, zero where it is not. An afternoon with your team is enough to complete it.

Statement	Fully true (2)	Partly (1)	Not yet (0)
Every device in the estate is within its supported life			
Platforms run software releases current enough to emit modern telemetry			
Assurance is enabled across the whole estate, not a pilot			
We hold at least one quarter of continuous telemetry history			
Licence tier is documented per device and covers what we need			
Catalyst Center is on a current release			
Cloud registration is permitted by our security policy			
Compliance has ruled on US processing of telemetry and prompts			
Site naming and hierarchy are consistent and machine-readable			

Statement	Fully true (2)	Partly (1)	Not yet (0)
Configuration is standardised, with drift measured			
Change management is recorded and auditable			
We can name our top five recurring incident types			

Reading your score

20 to 24. You are ready. Run a scoped pilot against a real operational problem, with a measurement you agreed before you started.

13 to 19. You are close. Identify the two lowest-scoring rows, fix those, and revisit in two quarters. Do not start a pilot on the assumption that the gaps will close during it.

7 to 12. Not yet. The work in Chapter 12 is worth more to you over the next year than any agentic capability, and it is what makes the capability usable later.

Below 7. The conversation to have is about network modernisation and supportability. Agentic operations are not the problem you have.

CHAPTER TWELVE

What to Do in the Next Twelve Months Instead

If the scorecard put you below twenty, here is the programme. None of it depends on any vendor's roadmap, all of it improves your network on its own terms, and at the end of it the agentic question answers itself.

Quarter	Do this	Why it pays regardless
Q1	Full installed-base audit by serial number. Map support status, end-of-life dates, software release and licence tier per device.	You cannot plan anything without it, and it closes an audit exposure immediately.
Q1	Ask compliance the United States processing question, in writing, and record the answer.	Determines whether the rest of the programme is a 2026 or a 2028 conversation.
Q2	Replace or upgrade anything past support. Bring platforms onto current releases.	Security patching, and the only route to modern telemetry.
Q2	Agree a naming and site hierarchy standard. Apply it everywhere.	Every report, audit and query gets easier. Unglamorous and permanent.
Q3	Standardise configuration by site type. Measure drift and close it.	Fewer outages, faster onboarding, comparable sites.
Q3	Extend assurance to the entire estate and leave it running.	Improves mean time to resolution on its own, and starts the telemetry history clock.
Q4	Formalise change management, with records and rollback plans.	Operational risk, and a prerequisite for any automated change assessment.
Q4	Baseline your top five recurring incident types and the time each consumes.	Without it you will never be able to prove whether anything helped.

Run that programme and one of two things happens. Either the agentic capabilities mature into something you can adopt on a network genuinely ready for them, or they do not, and you are left with a supported, standardised, fully instrumented estate with clean change control.

There is no version of that where you lose.

THE PARTNER TEST

Who Should Be Telling You This

The pitch you are receiving is not wrong. It is incomplete, and the incomplete parts are the ones that determine whether you get value.

The test to apply to anyone selling you agentic operations is simple. Ask them what has to be true about your network before it works, and then notice whether the answer is specific to your estate or general to the brochure.

Proactive Data Systems is a Cisco Preferred Networking Partner with thirty-five years of experience and more than 1,500 customers across India. Our CCIE-led team runs the readiness assessment before the conversation about licences: installed base by serial number, support and end-of-life status, software releases, licence tier per device, assurance coverage, configuration drift, and the data-flow position your compliance function will ask about.

Sometimes that assessment concludes that you should not buy this yet. We would rather tell you that now.

START HERE

Send us your installed-base list and your Catalyst Center release. We will complete the readiness scorecard with your team and tell you honestly which of the twelve rows you would score, and what the two-quarter path looks like.

proactive.co.in

REFERENCE

Frequently Asked Questions

What is AgenticOps in simple terms?

AgenticOps is Cisco's operating model in which software agents read network telemetry, reason about what it means, act on it and learn from the outcome, with people supervising rather than performing each step. It is built on purpose-built models trained on network data, and it is delivered progressively across Cisco's portfolio rather than as a single product you switch on.

Do I need a Network Advantage licence to use the Catalyst Center AI Assistant?

The AI Assistant is available across licence tiers, but what it can do follows the tier of the underlying feature it draws on. AI Network Analytics, which supplies the baselining and anomaly detection behind most useful answers, is part of the Advantage licence. Audit your estate device by device rather than assuming a single tier.

Does the AI Assistant work on an air-gapped Catalyst Center?

No. Active cloud connectivity is a prerequisite, with Catalyst Center registered to Cisco's cloud using a cisco.com account. For estates where policy forbids that connection, this generation of capability is unavailable. Cisco has indicated that on-premises options are coming for customers with data sovereignty requirements.

Where is our network telemetry processed?

Cisco documents that customer content and customer systems information associated with Catalyst Center, including prompts entered into the AI Assistant, is primarily processed and stored in the United States regardless of your physical location. Network event data for AI Network Analytics is deidentified before it is sent. Confirm the current position and the precise scope of deidentification for your deployment.

Is agentic operations going to reduce our network headcount?

No credible evidence supports that today, and any proposal justified on headcount reduction should be treated with suspicion. The realistic gain for a small team is in triage: removing the search phase at the start of an incident, which is where lean teams lose most of their time. The people are still needed to act on what the agents find.

How much telemetry history do the models need?

Enough to establish what normal looks like for your network, including its weekly and monthly cycles. A quarter is a reasonable minimum before you judge the output. Evaluating anomaly detection on two weeks of data will produce a misleading result in either direction and waste the evaluation.

Should we wait for the technology to mature?

Waiting to buy is reasonable. Waiting to prepare is not. The prerequisites, supported hardware, current software, full assurance coverage, standardised configuration and clean change control, improve your network on their own terms and take about a year. Do those and the buying decision becomes straightforward whenever you choose to make it.

REFERENCE

Sources, Disclaimer and About

Sources

Cisco Newsroom, "Cisco Expands AgenticOps Innovations Across Portfolio", February 2026.

Cisco Newsroom, "Cisco Unveils Agentic Platform for Operating and Defending Critical IT Infrastructure", June 2026.

Cisco, AgenticOps for Enterprise Networking, cisco.com.

Cisco, Catalyst Center AI Assistant overview and documentation, revised May 2026, including licence applicability, cloud connectivity prerequisite and data processing location.

Cisco, Catalyst Assurance User Guide, Cisco AI Network Analytics overview, including Advantage licence requirement and deidentification of network event data.

Cisco Live 2026 sessions and partner commentary on Cisco Cloud Control and AI Canvas availability, including controlled availability from 2 June 2026 and planned global availability in July 2026.

Digital Personal Data Protection Rules 2025, notified November 2025, compliance date 13 May 2027.

Reserve Bank of India, Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices.

Disclaimer

This whitepaper is general information for enterprise technology buyers in India. It is not legal, regulatory, tax or compliance advice, and it is not a product specification or a price quotation. Product availability, feature states, licensing terms and data processing arrangements described here reflect publicly available information at the time of writing and change frequently, in some cases monthly. Verify anything you intend to rely on against current Cisco documentation and with your own legal and compliance advisers before acting. Cisco, Catalyst, Catalyst Center and related marks are trademarks of Cisco Systems, Inc. Proactive Data Systems is an independent Cisco partner and this document is not published or endorsed by Cisco Systems, Inc. Any examples are illustrative and do not describe an identifiable customer.

About Proactive Data Systems

Proactive Data Systems is a Cisco Preferred Networking Partner with thirty-five years of experience and more than 1,500 customers across India. The company designs, deploys and supports enterprise networking, security, collaboration, cloud and AI infrastructure, with CCIE-led design and a 24x7 network operations centre in India.

proactive.co.in

Proactive Data Systems is an enterprise IT infrastructure and services company. For more than three decades, the company has designed, deployed, and operated the technology foundations behind some of India's most demanding enterprise environments, from regulated industries to large distributed operations.

Founded in 1991, Proactive serves more than 1,500 customers, from fast-growing businesses to Fortune 500 enterprises, across banking and financial services, manufacturing, healthcare, public sector, IT and ITES, retail, education, and global capability centres. Our work spans enterprise networking, cybersecurity, cloud and AI-ready data centres, collaboration platforms, structured cabling, and managed services.

A Cisco Preferred Partner under the Cisco 360 Partner Program, Proactive is recognised across the Networking, Security, Collaboration, Cloud & AI, and Services portfolios, and is a Cisco Advanced Customer Experience Specialised Partner. The company is ISO 9001:2015 and ISO/IEC 27001 certified, and engineers solutions in partnership with Dell, Nutanix, NetApp, Palo Alto Networks, Fortinet, VMware, Veeam, and CommScope.

Headquartered in Delhi NCR, with offices in Mumbai, Bengaluru, Pune, Hyderabad, Indore, and Singapore.

To learn more about our solutions, write to us
at pro@proactive.co.in

 www.proactive.co.in

PROACTIVE[™]
Enhancing Digital Experience

© 2026 Proactive Data Systems Pvt Ltd. All rights reserved.

This report is for informational purposes only. It does not constitute legal or regulatory advice. Statistical claims are attributed to named sources; readers should refer to original sources for complete methodology and context.