

Lumen's Network Automation Journey

From Manual Workflows to Orchestrated Automation
to AI-Assisted Autonomy

LUMEN® ♦  Intentia® ♦  SELECTOR



Table of Contents

Executive Summary	2
What Changed and Why it Worked.....	2
What the Program Delivered.....	3
Results at a Glance	3
Lumen’s Network Automation Journey: From Manual Workflows to Orchestrated Automation to AI-Assisted Autonomy	5
1. Background and Business Context.....	5
2. The Cultural Foundation: People, Roles, and the “Automation-First” Mandate	6
3. Swim Lanes: Observability and Orchestration	8
4. Design Tenets: Modular Orchestrations that Anticipate Autonomy, and How the Work Got Done	11
5. Measuring What Matters: ROI, Reliability, Adoption	12
6. From 1,000 Identified Workflows to 700 Autonomous Candidates	15
7. ServiceNow as a “Front Door” for Internal Self-Service	16
8. On API Quality, MCP, and Agentic Patterns: What’s New and What’s Just Good Engineering.....	17
9. Review: What Key Success Factors Made This Work at Lumen?.....	18
10. A Repeatable Method for Operators.....	19
11. Is AI “a radical shift” from where Lumen started?.....	21
12. What to Expect Next	21
13. Conclusion	25
14. Call to Action.....	25
15. Acknowledgements.....	26
Appendix A: Glossary	27

Executive Summary

Lumen's network automation transformation combined Lumen's cultural and organizational pivot, Itential's automation and orchestration capabilities, and Selector's AI-driven observability, to move from fragmented scripts to a governed, API-first automation platform now delivering a tremendous amount of automation with safe autonomy in production.

The success stems not from a single breakthrough, but from leadership's framing of network automation as a skills evolution, enforcing engineering discipline, measuring outcomes in business terms, and moving toward autonomous AI only when proven safe and valuable.

What Changed and Why it Worked

Leadership and Culture First. Automation was treated as a new operating model, not a project. Leadership empowered engineers to automate routine tasks, showcase live demos on a regular basis, and measure tangible results. In the beginning, small, safe wins built trust and participation.

Modular Designs Anticipating Autonomy. Actions were decomposed into atomic utilities and then composed into reusable runbooks accessible via service portals. Selector aggregates hundreds of alerts into single, explainable incidents with audit trails.

"We aspire to have 80% of all our interactions with our network be machine-to-machine by 2025." – Greg Freeman, Vice President of Network and Customer Transformation at Lumen Technologies

Oversight and Controls. Every workflow required design reviews (PDD → SDD), pre- and post-checks, and rollback validation. The up-front design effort cut rework and made changes safe for headless execution, and the instantiation of CI/CD processes ensured continued discipline. Process governance and guardrails paved the way to build trust and a path toward safe autonomy.

Itential builds the orchestration layer that enterprises need to accelerate into the AI era, where their infrastructure must evolve from static systems into adaptive, intelligent platforms that move as fast as their business. Their vision is simple: Unify and automate every part of the infrastructure lifecycle – across data centers, clouds, and networks – so teams can deliver secure, compliant, and scalable operations at software speed.



Selector is the industry's leading AIOps platform, analyzing and correlating telemetry across metrics, logs, events, and topology to deliver real-time visibility into network and service health. The platform enables automated root cause analysis, reduces alert noise, and supports faster incident resolution across hybrid and multi-cloud environments, delivering insights that improve service reliability and reduce operational overhead.



Clear Swim Lanes Between Tools. Selector acts as the AI “brain,” consolidating signals and triggering actions; and Itential is the “hands,” executing audited workflows. Machine-to-machine control loops are primary, with human-in-the-loop reserved for exceptions.

What the Program Delivered

North Star. The goal: 80% machine-to-machine operations by 2025 – measured, governed, and then delivered as customer self-help with zero-touch reliability.

Measured Outcomes. Lumen tracked workflows, executions, and financial metrics (OpEx saved as well as value created) through Splunk dashboards.

Operational Efficiency and Reliability. Pre- and post-checks, along with intelligent scheduling, reduced then eliminated maintenance windows and customer notices with a belief that “the easiest problem to fix is the one you never get.”

AI Event Log Reduction. Early efforts cut more than one billion alerts to roughly 57,000 actionable via 2,700 ML patterns and expert rules, enabling safe closed-loop readiness.

Cross-Domain and Customer Value. Automation now spans IP, optical, and service assurance; and self-service portals deliver safe, automated services, and diagnostics to engineers and customers.

Financial Framing. Lumen distinguishes between *labor avoided* and *new value created*, explicitly recognizing avoided outages and risk reduction as quantifiable outcomes.

Results at a Glance

The combined Lumen + Itential + Selector team efforts have delivered the following as of November 2025.



Over **350** live workflows



Alerts reduced from over **1B** → **57k actionable**



Fewer than **2 minutes** Service Assurance Platform downtime YTD



Internal self-service live via ServiceNow with zero-touch guardrails

A Recipe for Success

Start with People and Incentives. Charter an automation-first culture, celebrate live demos, and provide low-code and high-code on-ramps.

Enforce Design Gates. Key processes and checkpoints are critical. These include: technical peer review/governance, workflow validation before production deployment, PDD and SDD review for explicit success and safety criteria, and more.

Decompose to Safe Atoms First. Then build modular, reusable actions, and later make them callable from existing portals.

Define Clean Swim Lanes. Use AI/ML for detection, and then automation for governed execution.

Measure What Matters. Track workflows, runs, errors, OpEx savings, and value creation.

Keep Humans in the Loop Where Needed. Expand autonomy only when reliability data proves safety.

Together, these practices allowed Lumen to evolve from manual work to safe, explainable, and auditable AI autonomy. This created a cognitive NOC where machines handle routine operations, engineers focus on innovation, and customers experience fewer network events and fast, safe access to services.

Lumen's Network Automation Journey: From Manual Workflows to Orchestrated Automation to AI-Assisted Autonomy

1. Background and Business Context

[Lumen](#) is a large ISP that operates a large, heterogeneous network. The network was built through acquisition and merger, with legacy and modern networks spread across multiple domains.

Early automation and orchestration challenges centered on integration and scale. These included knitting together service activation and inventory stacks, internal portals and processes (e.g., Global Change Requests or GCRs), and service-assurance instrumentation so automation outputs could be controlled by consistent guardrails.

It was determined that more and better automation was needed and a leader was chosen.

- ✓ #1 - Lumen operates the most peered network*
- ✓ ~340,000 global fiber route miles
- ✓ ~163,000 on-net fiber buildings globally
- ✓ 350 +Tbps global backbone capacity
- ✓ 4 million Quantum Fiber locations in 16 US states

**CAIDA, AS Rank, January 2025*

Why This Matters

- ▶ Eliminated slow, human-gated change bottlenecks.
- ▶ Replaced fragile scripts with governed, observable automations.
- ▶ Exposed automations as API-based services for platform tools.


An Early Proof Point. At the very beginning of the project, Lumen's executive team challenged one of their vendors (Itential) to (1) integrate quickly across diverse systems, and (2) prove that orchestrations could be adapted when a system changed. Itential demonstrated both within hours, establishing platform credibility as well as opening the door to broader service-assurance use cases that the team later scaled internally.

Why This Matters: The team developed the vision in which agility eliminated a chronic bottleneck of slow, human-gated release and change pathways outside of the Operations team's control. This enabled Lumen to re-platform workflows from fragile scripts to governed, observable automations that could be exposed as services (and later, as API-based MCP tools).

2. The Cultural Foundation: People, Roles, and the “Automation-First” Mandate

Project Leadership: Greg Freeman

Greg Freeman is Vice President of Network and Customer Transformation at Lumen Technologies, where he leads the company’s network automation, orchestration, and AI strategy across the highest technical tier of network operations.



A veteran NetOps leader, Freeman’s remit at Lumen spans end-to-end operations (IP, optical/transport, metro, and voice) and championing NetDevOps practices that blend software engineering with operations. He has over 25 years in telecom and Internet engineering, and is guided by an automation-first mindset.

It’s worth noting that any network automation project begins with analysis of manual work being done, and then creating detailed workflow definitions. Freeman understands this and took great care to establish and enforce workflow validation and approval processes (with technical peer review and governance) before any subsequent automations were put into production deployment.

Regarding who could build automations, network engineers who do their network operations jobs day after day were deemed “citizen developers” and given the opportunity to contribute to the automation program via both low-code and high-code paths. We’ll expand on these activities below.

“No project was too small; everyone needs their ‘hello world.’”
– Jeff Torchia, Lumen

“These software fundamentals were all new to the network guys. There was heartburn in the beginning, but once we learned it, it got easier.”
– Marc Montez, Lumen

Freeman’s approach emphasizes measurable outcomes and repeatable workflows. He outlined how Lumen tracks (1) the number of workflows, (2) the number of workflow execution runs, and (3) two financial metrics (OpEx saved versus value created) streamed into Splunk for analysis. He also describes “ops reviews” of automation jobs that rebalance asynchronous jobs into optimal windows with the goal of preventing incidents before they occur, reinforcing that “the easiest problem to fix is one you never get.”



The easiest problem to fix is one you never get.

On scale and criticality, Freeman notes Lumen operates AS3356 – one of the world’s most interconnected backbones – so even a single high-impact workflow (e.g., one that prevents a route leak) can have much more value than many routine automations.

Day-to-day, he built teams that intentionally mix software developers with network engineers (noting that “unicorns” who do both are rare), reflecting his view that automation is a multidisciplinary sport spanning coding as well as IP and TL1/Optical domains.

Notably, network automation projects do not typically invoke images of dealing with the optical layer and interacting with network elements via TL1. But Freeman’s background helps bridge these network technology silos together and shows that automations can be orchestrated across domains.

It also shows that you can leverage predictive information from observability tools like Selector to make adjustments to traffic before components fail, or to trigger parts replacement workflows before traffic is ever impacted.

Why Leadership Matters: Freeman’s background and his VP-level mandate, in combination with outcome-driven metrics, cross-disciplinary teams, and community engagement, set the context for Lumen’s automation program design choices, as well as its trajectory toward increasingly autonomous operations.

“Nothing is possible without the leadership aligning to that vision. Greg has been absolutely vocal about it, and he’s been instrumental.”

– Varija Sriram, Selector

Why Leadership Matters

- ▶ Strong leadership sets clear mandates and outcome metrics that guide automation design.
- ▶ A leader unites teams and communities, enabling progress toward autonomous operations.



As long as you’re thinking “automation-first,” you belong here.

MANDATE

Freeman set an aggressive goal early in the project: Automate approximately 80% of manual activity within three to four years, which galvanized the team around the same scoreboard.

INCLUSIVITY

Freeman’s program then explicitly encouraged and welcomed both

low-code builders and Python-forward engineers to try their hand at the work. The message was simple: as long as you're thinking "automation-first," you belong here.

"If you get too restrictive, I'm going to do my own thing. If you get too loose, then it turns into the Wild West. There's a balance there. That's an art."

– Bill Terry, Lumen

PROJECT NAMES MATTER

That inclusivity converted skeptics into contributors, helped bring out the great ideas from people who had the vision but not the skills, and accelerated output. Because there were two on-ramps, network experts felt comfortable using low-code workflow definitions, and software developers were able to work in "high-code" alternatives (such as Bash, Perl, and Python), encapsulating scripts as governed actions within Itential.

Internally, this project was characterized as a "development program," not just a tool implementation. This was institutionalized by engineers documenting use cases, getting reviews on the high-level design, and then building and iterating on implementations. This created skill mobility and a growing bench of "citizen developers" and "automation architects."

"The whole goal is to keep people from imploding the network—that's why we review." – Marc Montez, Lumen

At every stage, the team measured visible ROI at the workflow level. Workflow onboarding required engineers to measure the manual time to do a workflow versus automated execution. The platform recorded runtime and fed an internal ROI calculator in Splunk, and this made savings and reliable repeatability concrete for executives.

3. Swim Lanes: Observability and Orchestration

HISTORY OF TOOLS USED

Before the current platform existed, engineers fixed pain with whatever was at hand: one-off Perl, bash, or expect scripts, or a PHP diagnostic here, and a cron job there. It "worked," but it was fragile. Every change to authentication, logging, or device syntax meant hunting through dozens of personal scripts. Pushing anything "real" depended on paperwork, weekly windows, and release processes outside operations.

EARLY WINS

When Freeman's Automation-First project began, the team's first wins were intentionally small and safe. They started by codifying common diagnostics (ping/trace, interface checks), wrapping them with pre- and post-checks, and making them callable from the tools people already used. This did two things fast: It removed tickets and handoffs for routine work, and it built trust as operators could see the same action run the same way every time, but with evidence and rollback baked in.

EARLY FORMALIZATION

From there, the group formalized how work gets done. They introduced lightweight but strict gates such as Process Design Documents (PDDs) which matured into Software/Solution Design Documents (SDDs) with reviews. SDD became builds. The defined process also included code reviews and a demo cadence so every team could see running software in the regular review meetings and not just slideware.

Any workflow that was determined not to be "revenue-impacting" was moved out of maintenance windows. Workflows that could be revenue-impacting required approvals, but gained standardized checks.

Crucially, they started measuring what mattered. Not just "how many workflows exist," but also "how often they run," "how often they succeed," "what toil they remove," and "what incidents they prevent." From the beginning, all of this data was streamed into shared dashboards so wins were visible.

As these patterns took hold, more ambitious use cases were attempted, such as:

- Closing the loop on noisy events by correlating them to likely causes,
- Attaching explain-your-work context to tickets, and
- Letting trusted automations be invoked headlessly by other systems (and later by natural-language agents with guardrails).


Selector brings extensive capability in this area by allowing users to interact not just with a network element, but with a large dataset of observability and related information via natural language.

SWIM LANES FOR TOOLS AND VENDORS

The early successes set up the next move: Make the swim lanes explicit. That way, each component of the stack gets excellent at its job without stepping on the others, with tight collaboration between them.

As the scale grew, the team codified swim lane boundaries as follows:

- **Selector would serve as a signal consolidator to provide AI-assisted decision support.** Selector digests high-volume telemetry and events, surfaces probable causes, and proposes or triggers the *right* remediation when confidence is high. That data ingested across IP and Optical domains provides Lumen with a “data quality” advantage versus raw, multi-layer ingestion across networking technology silos.



- **Itential would serve as the orchestration engine and policy enforcer.** As the “orchestration engine” and the “governed execution and policy enforcement” layer, Itential ensures AI-driven proposals from Selector translate into safe and auditable network actions. This applies to both workflows authored within Itential and scripts ingested into and managed by the Itential platform.

Giving each vendor a clear swim lane and reinforcing a “better together” stance reduced vendor friction. Selector surfaced insights, and Itential made them reliably actionable. This allowed Lumen to expose orchestration as a callable capability for decision systems on top, so the AI brain doesn’t have to be the hands – it just reliably calls them.

“We’re never going to do what [Selector is] doing. The customer wants orchestration through [Itential] because they’ve built guardrails so they can securely expose certain things to certain people and machines.”

– *Karan Munalingal, Itential*

Why Swim Lanes Matter: Defining clear roles for tools and vendors ensured reliable, safe network actions, and directly facilitated the transition from ad-hoc scripts to standardized, automated processes at scale. It played a key role in both reducing manual work and increasing trust between the tools and processes, which in turn led to safe autonomous operations.

Why Swim Lanes Matter

- ▶ Clear roles ensured safe network automation.
- ▶ Facilitated transition to standardized processes.
- ▶ Reduced manual work and increased trust.

4. Design Tenets: Modular Orchestrations that Anticipate Autonomy, and How the Work Got Done



Atomsics Compose. The team decomposed flows into atomic actions (e.g., read-only diagnostics like ping) and then composed them into larger runbooks.

Revenue-Impacting Changes Protected. Keep complex, potentially revenue-impacting changes behind approvals, while letting provably “safe” actions run hands-free. At each step, whenever there was a risk in uptime or availability (and therefore revenue), a “human-in-the-loop” failsafe was inserted into the runbook. This way, things like routing changes, changes to metro rings, and customer-impacting activities require human approval with enriched context provided by AI tooling. By comparison, routine maintenance actions could be exposed as APIs and run as scheduled jobs.

Careful Along the Way. Pre- and post-checks with error handling were baked in. Orchestrations were built *as if* no human would be present. They were instrumented to detect state, validate changes, and unwind safely when needed.

Being Careful Takes Work: Lumen developed a continuous improvement / continuous deployment (CI/CD) pipeline over five years, moving from manual enterprise admin deployments to fully-automated CI/CD for Itential content, Ansible, and other supporting applications. By 2023, this enabled moving from three-hour manual maintenance windows down to zero-downtime deployments. In 2024, they launched their 2.0 pipeline that enabled developers to deploy to dev/test independently. A two-approver requirement for test/production deployments was implemented, as well as other comprehensive guardrails such as naming standards checks, adapter validation, file naming validation, and more. In mid-2024, the 2.0 pipeline resulted in the elimination of deployment windows, resulting in a full CI/CD process with no separate rollback strategy. Currently, the team continues to work to identify errors and timeout issues so they can better understand operational conditions that cause them, and then put in the right guardrails and error handling so errors and timeouts are reduced.

“One of the first [AI projects] was ticket summarization. AI is great at that, [but] we have to have predictive workflows with very good guardrails. We can put it in a workflow, because we have the checks inside of that predictive workflow.” – Marc Montez, Lumen

Trust Grows with Experience and Data. The right for something to run autonomously was earned by setting thresholds of error-free operation over a defined time period or number of runs. While most operations were required to have human intervention at the beginning, the team logged inputs and outcomes across thousands of runs, and began to know which workflows “never fail in certain conditions.” Over time, they could elevate those to autonomous execution when trusted conditions were met

An Example of Light Touch Autonomy. A practical example of light-touch autonomy emerged around power management in optics: When link activity is low, the system can turn lasers off to save energy, and then bring them up on demand with protocol checks – no humans required. Small steps like this build trust after consistent execution, paving a path to potentially greater adoption of autonomous processes.

Why Earned Autonomy Matters

- ▶ Tools need a process to guide use.
- ▶ Repeatable runs without errors build confidence.
- ▶ Allows for variable conditions to earn autonomy.

Why Earned Autonomy Matters:

Tools alone don’t get the work done. The processes that invoke the tools need to guide how and when the tools get used. In an increasingly software-driven world, CI/CD provided the right framework for Lumen to develop, deploy, and orchestrate its automations. Trusted processes that evolve smartly over time, and with measured success, are necessary ingredients for autonomous behavior in network automation.

5. Measuring What Matters: ROI, Reliability, Adoption



From the start, Lumen treated both return on investment (ROI) and reliability as “first-class citizens,” not vague aspirations. At the same time, operational metrics guided what got built next.

The team defined “workflow” as the smallest unit of value that could be created. The team then tracked three things: how many discrete workflows exist, how often each individual workflow is run, and the tangible value each workflow either generates or saves: time or money.

Freeman describes how that matured into two complementary ROI metrics captured in a Splunk repository:

- **OpEx saved** (replacing a manual, once-per-period task), and
- **Value add** (running safe, necessary automations far more frequently than a human ever could).

This let leadership see both the labor avoided as well as the extra reliability and improved customer service gained when automations run 24x7.

As the program scaled, the metric stack evolved from simply “how many **workflows** did we ship?” to “how many transactions are we executing?” and “what are the **error rates per workflow**?” All this was streamed to and tracked via Splunk dashboards.

That evolution mattered: Transaction volume reveals which automations are doing real work, while error-rate trends expose where the team needed to harden guardrails before removing approvals for autonomous activity.

The measurement framework also acknowledged that *not all automation runs are equal*. Some automations trigger constantly and shave minutes of toil each time, while others might run once but avoid a headline-making incident.

For example, an orchestration that can prevent a BGP route leak on AS3356 illustrates the “value add” side of the ledger, but it’s hard to price an incident that didn’t happen – even though avoiding those events may be worth a few million dollars each time they are avoided. By capturing savings separately from OpEx, the team could defend investments that are also about reliability and risk reduction, not just labor substitution.

Reliability and Safety Required. Reliability and safety were treated as prerequisites for speed. The operating model institutionalized pre-checks and post-checks, code reviews, and an explicit AAA/permissions layer, so only approved personas (human or system) can execute specific actions.

The pattern is defined as the workflows (and the Python or Ansible they invoke) pass review, and then are permissioned via an internal AAA app. They also must include programmatic pre- and post-conditions that verify safety before and after any change. These same guardrails apply whether a human clicks “run” or an AI agent makes a call.

When things did go wrong, the post-mortems drove tighter controls. Klaus Bonapart and Lord Concepcion of Lumen recounted an early incident that was triggered “out of process” from User Acceptance Testing (UAT) against

production. This is the kind of event that led to stronger stage gates and protections “to protect ourselves from ourselves.”

Over time, the platform’s reliability picture became a strength. Jeff Torchia cites less than two minutes of downtime year-to-date (nine months) for the automation estate, even across data-center failovers – evidence that the plumbing enabling orchestration is itself highly available.

Those safety practices also enabled a pragmatic policy change: Where evidence showed a workflow was both (1) safe and repeatable with low/zero incident linkage, and (2) had high success rates, the team could justify removing some human approvals while keeping approvals in place for potentially revenue- or latency-impacting actions.

In parallel, operations began a regular **Ops Review** for automation. Using Grafana views of estate load, they shifted low-risk jobs to off-peak times and prioritized preventive workflows, further reducing maintenance windows and disruption.

Adoption by End Users. Adoption velocity accelerated when the automation team met users where they already worked, such as in familiar systems like ServiceNow. Selected automations were exposed externally to customers through Lumen’s customer portal, so customers could safely self-serve common actions.

An example: Historically, “after repair actions” leave a service on “protect,” and customers had to call in and wait for an engineer to manually revert the service to “work.” The team turned that into a one-click, guardrail-checked orchestration that a customer schedules themselves with no human touch required by Lumen.

The platform architecture explicitly includes ServiceNow in the mediation/integration layer, reflecting how automations are consumed without forcing users into a new portal. In closed-loop scenarios with Selector, Itential not only executes remediations, but also updates the ServiceNow incident and sends confirmation back to Selector. The result is that the operator stays in one pane of glass and gets end-to-end evidence.

“If there’s a config event, we are able to then send that to Itential, ‘Can you check if this is compliant? Was there a change request? Does that conform to the compliance checks? If not, revert the change.’” – *Varija Sriram, Selector*

Feedback Loop Flywheel. The “measure → harden → expose” loop created the flywheel for the team as follows:

- Measurement made wins visible and defensible to the executives by showing OpEx savings, value creation, and error-rate improvement; and to the engineers through usage, merges, and demo-driven learning.
- Reliability engineering, such as guardrails, pre/post-checks, and stage gates, minimized blast radius and underwrote decisions to remove approvals in well-understood cases.
- Thoughtful distribution, such as putting automations behind ServiceNow and in the customer portal, removed friction from consumption, turning workflows into callable services for humans and systems alike.

The result is a program that is simultaneously governed and fast, and runs millions of executions a year, rising reliability, and has a growing set of use cases where autonomy can safely “just run” because the evidence has proven it should.

Automation as a Force Multiplier: We noted above that not all automation runs are equal. Lumen saw that once certain workflows were automated, the ability to invoke them much more frequently and at a greater scale can be a real force multiplier in an automated network.

Additionally, the process of carefully defining and understanding a workflow for the purposes of automation with advanced observability, automation, and orchestration tooling in mind can lead to insight on how to radically improve a workflow. Always be on the lookout for new workflow optimizations and think critically along the lines of “How can I use these new tools to do this job in a way I couldn’t as a human operator?”

Why ROI, Reliability, and Adoption Matter:

Automating networks should be done to produce provable ROI. Lumen kept this top of mind and took a sufficiently detailed approach (the number of discrete workflows, how often each one runs, and the tangible value created or OpEx saved) without being overly complex. This is key to showing the actual value of network automation, and to getting the necessary support for a large project that spans years.

Why ROI, Reliability, and Adoption Matter

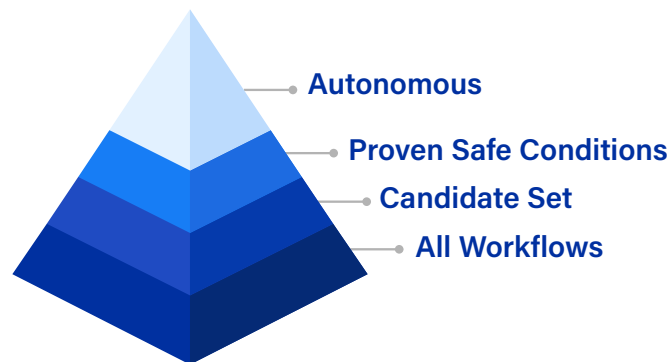
- ▶ Network automation must yield provable ROI.
- ▶ Measure success with detailed, but simple, metrics (workflows, frequency, value).
- ▶ This approach secures support for long-term projects.

6. From 1,000 Identified Workflows to 700 Autonomous Candidates

The team is now over five years into this project, and has identified over **1,000** workflows that are candidates for automation. They have over **350** built and

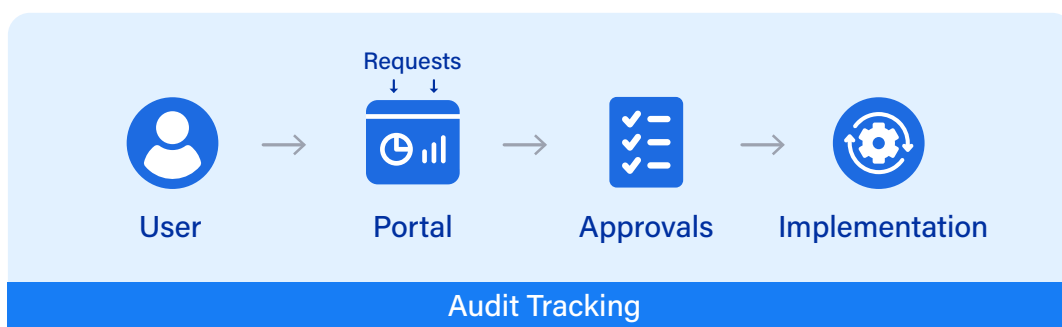
deployed as of October 2025, and are on a path to another **700** which could run autonomously once confidence criteria are met.

Lumen has learned and codified which inputs, thresholds, and contexts justify automatic action versus require approval. The criteria for autonomous mode are generally that if the workflow runs 1,000 times and they understand the conditions where the workflow “never fails,” then it becomes a candidate for Selector’s AI execution. At that point, the workflow can move into the category of “just let it fly.”



Any workflows that haven’t met those criteria require a human in the loop to complete. This approach treats moving a workflow to autonomy not as a switch, but as a **spectrum** governed by evidence and per-workflow service level objectives (SLOs).

7. ServiceNow as a “Front Door” for Internal Self-Service



A common friction in large enterprises is the “**Portal Gap.**” Engineers can build automations in weeks, but publishing them requires waiting for centralized portal/IT support teams to put the submitted automations into production to make them accessible and usable.

Lumen addressed this by exposing some of the “let it fly” Internal automations for internal use in ServiceNow. This cut the time-to-adoption and aligned with IT

governance while keeping infrastructure decisions with infrastructure leaders.

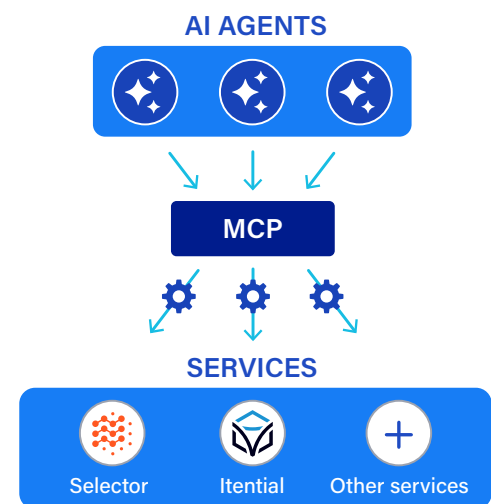
Why This Works: Infrastructure leaders own the orchestration and controller stack, while IT preserves governance via ServiceNow. Both charters are met without blocking delivery.

Result: Fewer “shadow IT” tools, faster consumption, better telemetry on what’s used, and where to invest next.

8. On API Quality, MCP, and Agentic Patterns: What’s New and What’s Just Good Engineering

API Reality. Lumen and Itential repeatedly ran into inconsistent or minimally “wrapped” APIs across vendors – evidence that programmability pressure can outpace quality. The market is improving API maturity and security investments, but variability remains a frustration for integration teams.

Enter MCP and Agentic Control Planes. To safely expose network capabilities to AI and their agents, Lumen prototyped an MCP server on Itential. This allowed them to preserve all the guardrails while enabling a tool-calling experience from the northbound AI layer.



In parallel, Itential and Selector have both productized their MCP servers so customers can onboard agents just like systems or APIs and keep execution governed.

“We’d rather support an elite [agentic] experience, but still relay that to network/security/infrastructure through governed actions.”

– Karan Munalingal, Itential

Lumen sees promise with MCP and will continue to investigate its use on top of the solid foundation built with observability, automation, and orchestration tooling and processes.

“We have a proof of concept using MCP tied to several LLMs. As of today, we have [over 90] tools tied to it. You talk to it in English, and it gives you a graph in seconds.” – Rabih Nahas, Lumen

Agent-to-Agent (A2A) Implications. As agents coordinate across domains (tickets, telemetry, config), MCP tooling and explicit swim lanes become more, not less, important. What changes is *who* initiates a call and *how* confidence is established, not the need for pre- and post-checks, predictability of success, and rollbacks.

Why Good API Engineering Matters: Lumen is not the only company to encounter issues with API security and quality, so driving API quality and consistency is good for Lumen’s needs and the entire networking industry. This is especially true as software-centric tools and processes become increasingly important for network operations.

AI agents are becoming more useful for NetOps tasks, and agent access to services via APIs and MCP as a proxy for APIs will also put more focus on this area. Intentional discovery of and experimentation with AI agents, MCP, and other rapidly evolving mechanisms is key to any organization determining how these tools could be useful for them.

Why Good API Engineering Matters

- ▶ Strong API quality and consistency matter for Lumen and the wider networking industry.
- ▶ Growing AI agent and MCP use makes intentional exploration and experimentation essential.

9. Review: What Key Success Factors Made This Work at Lumen?



Strong Product Leadership + a Foundation of Credibility. Leadership intentionally created the right team dynamics and established deliberate processes with small incremental steps that built trust.

Inclusive Contribution Model. Low-code and high-code tracks prevented a culture war, turned vision into action, and attracted and developed more builders.

Evidence Over Opinion. Per-workflow ROI, reliability, and repeatability metrics provided evidence for progress, support for a project spanning years, and laid a foundation for increased autonomy over time.

Swim Lanes by Design. Intential’s action layer and Selector’s decision support avoided overlap, sped collaboration, and sharpened accountability.

Governed Exposure of Capabilities. Deliberate, detailed internal process definition, along with portal-based access to automations, allows more people and machines to use automations safely.

10. A Repeatable Method for Operators

Below is a condensed, reusable framework with artifacts you can use as templates.



BEFORE YOU START WORK: CHARTER

- Define and agree on what your overall goals and metrics are as.
- **Artifacts:** Team Charter and North Star Metrics

PHASE 1: BASELINE

- Recognize and commit to detailed workflow analysis and documentation. There are no shortcuts to understanding what you do today and looking for opportunities to improve those workflows with new tools. Leverage and codify internal knowledge and investigate new tools that can help accelerate workflow documentation and optimization.
- Define swim lanes of observability, orchestration, source of truth, and other key functions up front.

- Commit to and budget time for per-workflow ROI and safety metrics from day one.
- **Artifacts:** Automation Opportunity Register (AOR). This is where you spend the time documenting manual tasks, volumes, time spent, and blast radius.

PHASE 2: PLATFORM AND GUARDRAILS

- Stand up the orchestration plane and re-platform existing scripts as governed actions.
- Implement pre- and post-checks, rollback paths, and audit logging at each step, not as afterthoughts.
- Expose trusted, callable actions through controlled access portals to speed and ease consumption and align with governance.
- Thoroughly define and communicate your security requirements, recognizing the incredible pace of evolution in AI tooling and potential security and identity exposures.
- **Artifacts:** Documentation of System and Security Requirements.

PHASE 3: BUILDER ON-RAMPS AND PLAYBOOKS

- Offer low-code and high-code contribution pathways.
- Run weekly design reviews of PDDs and SDDs and adopt software-centric processes for developing and deploying automations, such as CI/CD.
- Run weekly demos of working automations, not slideware.
- **Artifacts:** HLD template, “Design Windows & Pre/Post-Checks” checklist, “Run-Outcome” telemetry schema (inputs, outputs, errors) for every workflow, schedules of weekly demonstrations.


PHASE 4: DECISION LAYER AND AUTONOMY READINESS

- Integrate observability and AI to propose actions, but maintain human approvals on revenue-impacting changes.
- Over time, track “never-fail” conditions as well as gray zones to determine when to graduate automation candidates to autonomy as evidence warrants.
- **Artifacts:** Autonomy Readiness Score (ARS) per workflow. Confidence/Trigger policy.

PHASE 5: FUTURES: MCP, AGENT ENABLEMENT, AND EMERGING TOOLS

- Stand up an MCP server (perhaps multiple, grouped by functions) to expose tools to agents without bypassing orchestration guardrails.
- Define agent-to-agent contracts (who can call what, with which proofs, under which triggers).
- Establish organizational processes for proactively evaluating rapidly evolving tools and technology, AI and otherwise; always be evaluating.
- Start with read-only or reversible actions and, over time, expand by evidence.
- **Artifacts:** MCP and Agent contracts. Org procedure for product/tech evaluation.

11. Is AI “a radical shift” from where Lumen started?



Yes and No. The *northbound* user experience has changed. Now agents are calling tools via MCP, but Lumen’s core tenets of modularity, guardrails, and human-in-the-loop approvals where it matters are the same design grammar they started with. What’s changed is how fast you can surface the right action and who/what initiates it.

Where There Is Continuity: Itential always modeled everything as actions. The addition of MCP and agents is a natural extension of “onboard a system, expose an action.”

What Has Changed: Selector’s AI has shifted more of the “should we act?” decision toward machines when confidence is warranted, and this shortens the loop between detection and remediation. Selector’s MCP and agentic operating model with clear roles drives more capability to machines – agents with agency.

Danger Will Robinson – This Is Not a Free Pass: The team still draws a hard line at autonomous changes that can impact revenue and customer traffic unless the changes meet very stringent evidence and guardrail criteria.

12. What to Expect Next

MORE WORKFLOWS MOVE TO AUTONOMY

Lumen’s next phase, in addition to another 700 workflows moving to automations, is about graduating more work into autonomy, but only where the data says it’s

safe. The playbook is clear: Prove reliability in context (device family, region, time window), then start to remove human approvals for that slice while keeping audit and rollback intact.

Two operational practices support this.

- First, the team’s recurring automation ops reviews look at estate load (using Grafana), and shift asynchronous jobs to optimal times, while retaining service windows only for truly notice-worthy changes. The motto “the easiest problem to fix is the one you never get” keeps preventive workflows at the front of the queue.
- Second, engineers expand autonomy piecemeal to build confidence. They do this by discovering inventory types, handling specific alarm classes, and then permitting small, reversible moves of traffic before tackling broader actions that could be revenue-impacting.

“Start with deterministic checks/actions (“power? ping? shut/no-shut?”) before riskier changes, collect feedback, then expand. First two steps which are very safe: **build trust that way to build a safe path to autonomy.**”

– Sachin Natu, Selector



As correlation quality rises and noise falls, more remediations can “graduate.” As an example, Lumen’s initial work (with Watson AIOps) collapsed >1B alerts down to ~57K actionable via >2,700 ML patterns plus 100+ expert rules. This is a prerequisite to a closed loop as it raises the signal-to-noise and explains the “why” behind a trigger.

Expect Lumen to keep using that evidence (success rates, error/rollback incidence, context boundaries) to decide where autonomy can simply run and *where* approvals should remain.

“I want the primary way to be that AI/ML interacts with the orchestrator to do closed loop automation. [Natural-language is] a secondary way to interact with the network.” – Greg Freeman, Lumen

MORE FOCUS ON API QUALITY AND SOURCE OF TRUTH

A second focus is API quality and source of truth (SoT) maturity, so that agents and decision systems act on high-quality, unambiguous data.

The architecture already routes external callers through a REST mediation layer (Auth/Apigee), with ServiceNow, Laser Portal, and other apps consuming orchestrations as services rather than scripting devices directly.

On the configuration side, Lumen is pulling standard configs from engineering's Basegen (their "source of truth") to avoid procedure documentation/wiki/MOP drift when building partner NNIs.

Inventory is being treated the same way: Selector-fronted MCPs now support inventory edits (e.g., moving a device to a new region) by writing to systems like NetBox through controlled tools rather than ad-hoc scripts.

Upstream of all that, data gathering is expanding beyond SNMP to controller APIs and time-series data lakes for both IP and optical, which are necessary if closed-loop actions are going to span domains.

Put simply, the more consistent the SoT and the more standards-compliant the APIs, the easier it is to let agents reason reliably and for the orchestrator to enforce policy.

MORE AUTOMATED CUSTOMER INTERACTION

Third, MCP will provide an option for customers to interact with Lumen infrastructure through natural language. An end customer could issue requests for info on circuit status, which could be derived from the Selector observability data stores, with the appropriate guardrails put in place by Lumen protection mechanisms.

AGENT-TO-AGENT

Fourth, Lumen is formalizing agent-to-agent (A2A) governance so multiple AI agents can coordinate safely without policy drift or "action collisions."

The principle of keeping Intent as the governed execution path is already a design tenet. Lumen frames AI/ML systems as the triggering brain and the orchestrator as the hands, and uses ChatOps/NL interfaces as a secondary path for humans. That separation becomes even more important as Lumen trials multi-agent environments.

Today, the team has an MCP-based, multi-agent, proof-of-concept with over 90 tools connected. It's powerful but not yet in production precisely because scope, determinism, and QA need stronger guardrails.

Lumen implemented the following three-phase approach to make agents safer.

- First, limit what LLMs decide: Selector performs the heavy correlation on raw telemetry.
- Second, use LLMs narrowly to translate results to English and drive user experience. This reduces hallucination risk while preserving explainability (“show your work,” including the alerts collapsed into a single ticket).
- Third, permission and review stages remain in place: Itential runs with AAA-scoped permissions, code/design reviews, and stage environments; those same controls apply whether a human presses “run” or an agent calls the API.

These near-term moves sit inside a North Star that has guided Lumen throughout the project: 80% machine-to-machine operations with customer-facing self-service wherever safe.

“On a network you’ll see hundreds of alerts. In the ServiceNow database, you’ll see one ticket with all those hundreds of alerts presented [by Selector] as one ticket, so [you can see] why they were collapsed, how they were connected.” – Sachin Natu, Selector

CROSS-DOMAIN AUTOMATIONS

Finally, expect more cross-domain automations. Lumen is deliberately wiring triggers from Optical into actions at IP and vice-versa so that an event in one layer prompts the right change in another through the orchestrator’s policy checks.

The organization already teams up software developers with network engineers to make that practical across TL1/Optical and IP stacks, and the automation ops reviews will keep fine-tuning when and where those automations run.

SUMMARY OF WHAT’S NEXT

- More workflows earn autonomy as their telemetry proves safety and reliability.
- APIs and SoT keep tightening, so agents act on truth.
- A2A governance ensures only one set of hands (Itential) executes changes under policy while multiple brains (Selector and Itential AI/ML, with natural language agents) decide when and why. This way, Lumen will expand safe autonomy without losing control of safety, auditability, or customer experience.

13. Conclusion

Lumen's journey is a great example of "you don't leap to autonomy, you earn it."

- By having a mandate and then investing in culture, integration, and guardrails, Lumen was able to define team and vendor swim lanes.
- By investing in instrumentation and measurement, they were able to measure outcomes per workflow, have a durable runway where AI can safely speed decisions, and MCP can safely expose tools without compromising control or reliability.



*You don't leap to autonomy,
you earn it.*

We believe this is an operating model that other large network operators can adopt:

- Have bold North Star metrics and executive support
- Spend time documenting and rethinking workflows
- Be inclusive with low and high-code onramps
- Start small, compose, and iterate
- Instrument everything
- Let evidence move the line between approval and autonomy
- Keep the "hands" governed even as the "brain" gets smarter
- Make safe automations available through existing portals

14. Call to Action

Solutional helps NetOps teams to analyze AI and other new technologies for adoption into their tech stacks, and how organizations need to adapt with those new tools and technologies.

Would you like to move your NetOps team toward Lumen's model for network automation, orchestration, and autonomy (e.g., the [Repeatable Method](#) in Section 10)? [Contact us](#) for a no-commitment conversation.

15. Acknowledgements

We want to express our gratitude to all the individuals that contributed to and supported this effort.

From Lumen

Greg Freeman

Jeff Torchia

Bill Terry

Marc Montez

Lord Concepcion

Klaus Bonaparte

Rabih Nahas

From Itential

Karan Munalingal

Kristen Rachels

Chris Wade

From Selector

John Capobianco

Sachin Natu

Varija Sriram

Stephen Ochs

Packet Pushers for hosting and their support of the *Total Network Operations* podcast.

Appendix A: Glossary

A2A: Agent2Agent Protocol

API: Application Programming Interface

CI/CD: Continuous Improvement / Continuous Deployment

MCP: Model Context Protocol

PDD: Process Design Document

ROI: Return on Investment

SDD: Software (Solution) Design Document

SoT: Source of Truth

TL1: Transaction Language 1, a standardized, ASCII-based language used for managing network elements, particularly in older or legacy telecommunications and optical equipment.

UAT: User Acceptance Testing

ABOUT ITENTIAL

Itential is the infrastructure and network orchestration platform for the AI era. Our cloud-native solution unifies lifecycle operations across hybrid network and cloud environments, empowering enterprises to automate everything from configuration management and compliance to service delivery. With support for both low-code orchestration and high-code extensibility, Itential enables teams to build, adapt, and scale automations faster. Through agentic orchestration via MCP and intelligent workflows, Itential connects IT systems, CI/CD pipelines, and network infrastructure to deliver secure, end-to-end automation that's ready for the next generation of digital and AI-driven operations. Learn more at itential.com.

ABOUT SELECTOR

Selector helps enterprises prevent and resolve network and service issues faster by unifying observability, AI, and automation in a single platform. By correlating telemetry across metrics, logs, events, and topology, Selector delivers real-time visibility into service health, dependencies, and root causes across hybrid and multi-domain environments. Purpose-built AI and network-trained large language models accelerate detection, diagnosis, and remediation, enabling IT teams to reduce downtime, increase reliability, and operate with greater efficiency. Selector empowers organizations to shift from reactive troubleshooting to proactive, data-driven network operations. Learn more at Selector.ai.

ABOUT SOLUTIONAL

Solutional is a technology and marketing consulting firm helping IT teams, marketing departments, and business leaders bridge the gap between cutting-edge technology and real-world outcomes. With deep expertise in networking, automation, modern infrastructure, and AI for network operations, Solutional delivers strategic guidance in consulting services, technical marketing, and education that drives innovation, from the data center to the boardroom. Visit solutional.com to learn more.