ITAM FORUM

The AI Ladder: A Strategic Roadmap to Augmented and Autonomous ITAM

Introduction

For decades, the IT Asset Management (ITAM) community has operated as the analytical engine of the enterprise – balancing cost, risk, and compliance. But the landscape is undergoing a fundamental transformation. The traditional, reactive administrative function of tracking hardware and software licenses is evolving into a proactive, strategic discipline powered by Artificial Intelligence.

The challenge we face is one of scale and speed. As organisations manage increasingly complex hybrid estates and decentralised SaaS portfolios, manual processes simply cannot keep pace. We are at an inflection point where responsibilities are increasing, yet headcounts remain flat. To navigate this shift, we need more than just new tools – we need a roadmap.

Introducing the AI Ladder for ITAM

The AI Ladder is a practical maturity model designed to help you locate your organisation's current standing and define the path forward. It moves from the foundational necessity of visibility to the ultimate vision of bounded autonomy.

05 **Autonomy (Bounded):**
Transition to agentic workflows where AI acts as a strategic team member within human-defined guardrails.

04 **Action (Close the Loop):**
Implement 'always-on' governance to manage the volatility of consumption-based software pricing.

03 **Insight (Spot Patterns Humans Miss):**
Use machine learning to identify anomalies, license shortfalls, and predictive maintenance needs.

02 **Trust (Make the Data Usable):**
Ensure AI-driven insights are reliable through rigorous data stewardship and governance.

01 **Visibility (Stop Flying Blind):**
Establish a near real-time source of truth across on-premises, cloud, and SaaS environments.

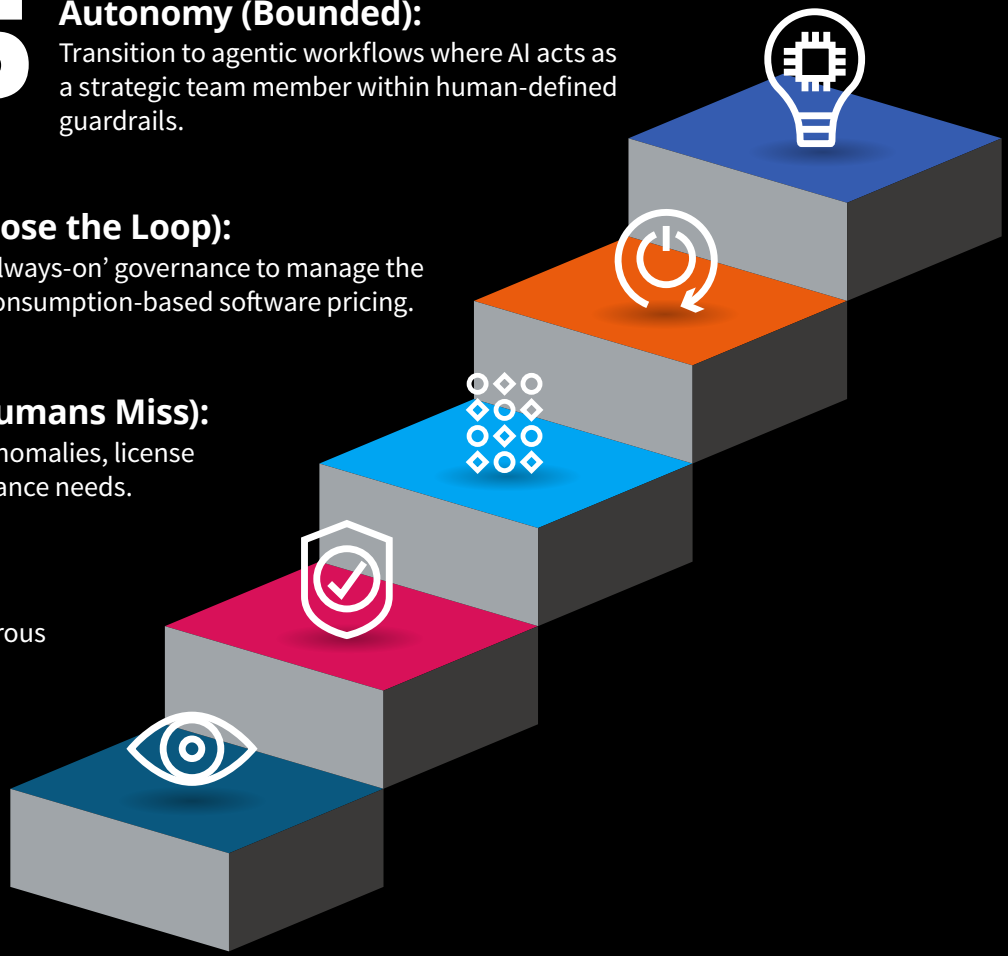

Our Contributors

To build this framework, I've gathered insights from eight, leading, industry voices. Each brings a unique perspective to a specific rung of the ladder.

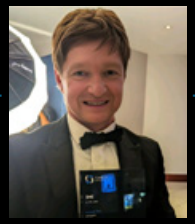

Gareth Davies
Director, bedigital

A champion for data stewardship as the practical foundation for improving ITAM outcomes.

Stephanie Day
VP, SaaS Management, Calero

An expert on how consumption-based pricing is breaking traditional governance models.

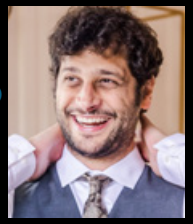

Phil Perfetti
Senior Product Marketing Manager, Flexera

A specialist in how AI introduces speed and precision to automated discovery and risk detection.

Ingrid Roodenburg
General Manager, Noventiq Benelux

A vocal advocate for the "Garbage In, Garbage Out" principle and the criticality of data foundations.

Victor Paredes-Colonia
Director of Product Management, ServiceNow

A thought leader on treating AI not as a tool but as a 'team member' to be onboarded and developed.

Tyrone Howland
Lead ITAM Consultant – UK & Europe, SHI

An authority on the partnership between machines for scale and humans for strategy.

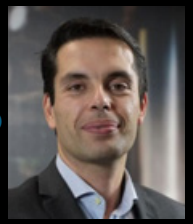

Luís Miguel Martins
Global Sales Strategy Director, SoftwareOne

A pioneer of 'augmented ITAM' and the rise of ITAM bots to support human judgment.

Stuart Moffat
Senior Product Manager, USU

A strategist mapping the evolution of intelligent support to autonomous ITAM agents.

Rung 1: Visibility (Stop Flying Blind)

The first rung is about ending the era of manual, retrospective tracking. As **Stephanie Day** rightly points out, “Most ITAM programs are still structured around static inventories and retrospective reviews,” but the “result is a governance model that no longer aligns with reality.” In today’s landscape, “Consumption-based pricing converts software from a known expense into a live meter.”

If you cannot see the meter in real-time, you cannot manage the cost.

Scaling Discovery through AI

The complexity of modern estates has moved beyond human reach. **Phil Perfetti** observes that, “The complexity of managing assets has grown exponentially, leaving traditional ITAM processes struggling to keep pace.” AI changes this because “AI-powered discovery tools can identify these assets in real time, normalise data and eliminate duplicates faster than any manual process.”

Deploying ITAM Bots

To achieve this level of oversight, **Luís Miguel Martins** advocates for the use of ITAM bots, which “can continuously discover assets across environments, normalise inconsistent data, [and] reconcile inventory sources in near real time.”

This visibility is further sharpened by “integrating data across systems such as IT Service Management (ITSM) and Identity and Access Management (IAM) to create a more accurate and consistent source of truth.”

Automating Software Recognition

A significant barrier to visibility has always been the manual drudgery of data entry. **Stuart Moffat** explains that “Machine learning models can now match raw traces to a standardised software catalogue,” which “improves consistency, enhances accuracy, and reduces manual effort.” This ensures your team starts with “clean, accurate data as the foundation for every strategic decision.”

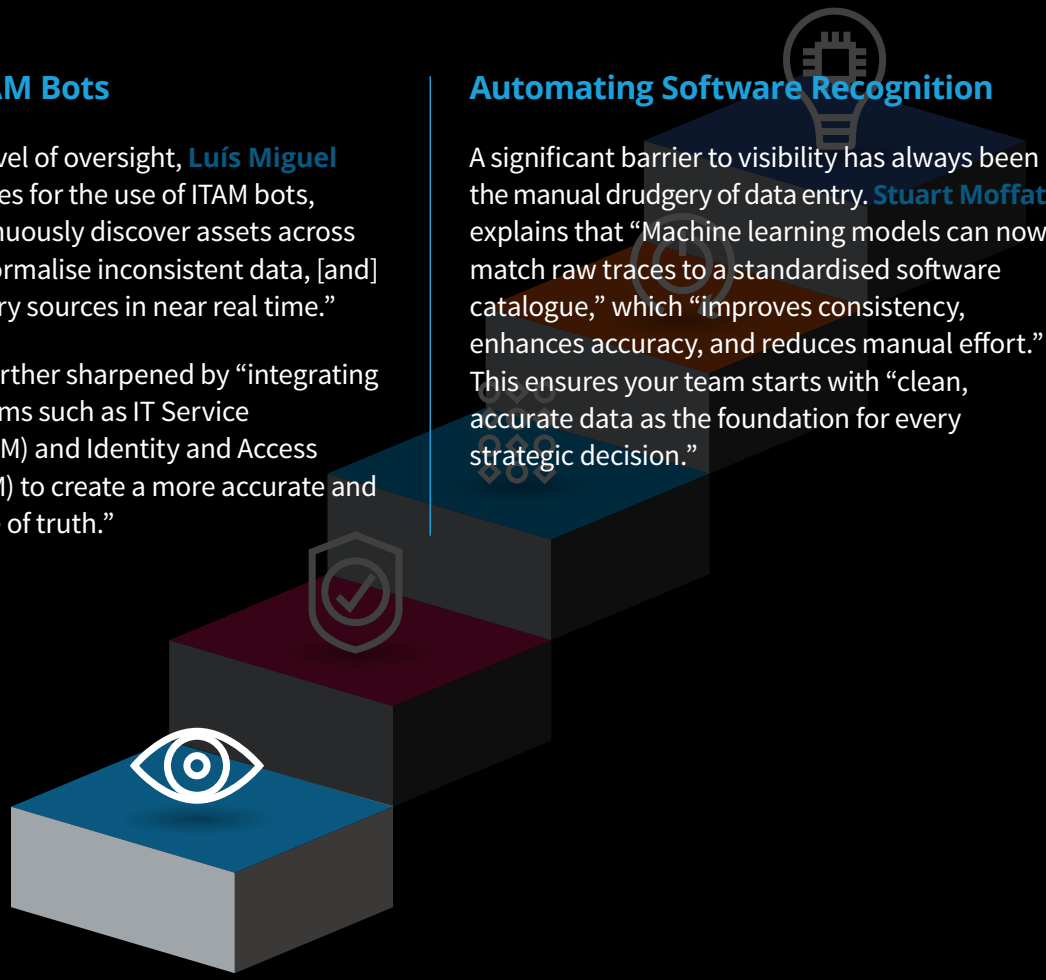

01 **Visibility (Stop Flying Blind):**
Establish a near real-time source of truth across on-premises, cloud, and SaaS environments.

What Good Looks Like: The Visibility Checklist

- **Real-Time Identification:** Your tools "identify these assets in real time" across all on-premises and cloud environments.
- **Unified Truth:** You have integrated "data across systems such as IT Service Management (ITSM) and Identity and Access Management (IAM)."
- **Automated Normalisation:** You use AI to "normalise data and eliminate duplicates faster than any manual process."
- **Live Metering:** You are tracking software that "prices itself continuously" rather than relying on "annual reconciliation".

If you only do one thing – Audit your current discovery coverage.

As **Ingrid Roodenburg** warns, "It can't optimise what it can't see." If you are ignoring "cloud workloads or remote endpoints," you are operating with "partial visibility" that creates "dangerous blind spots."

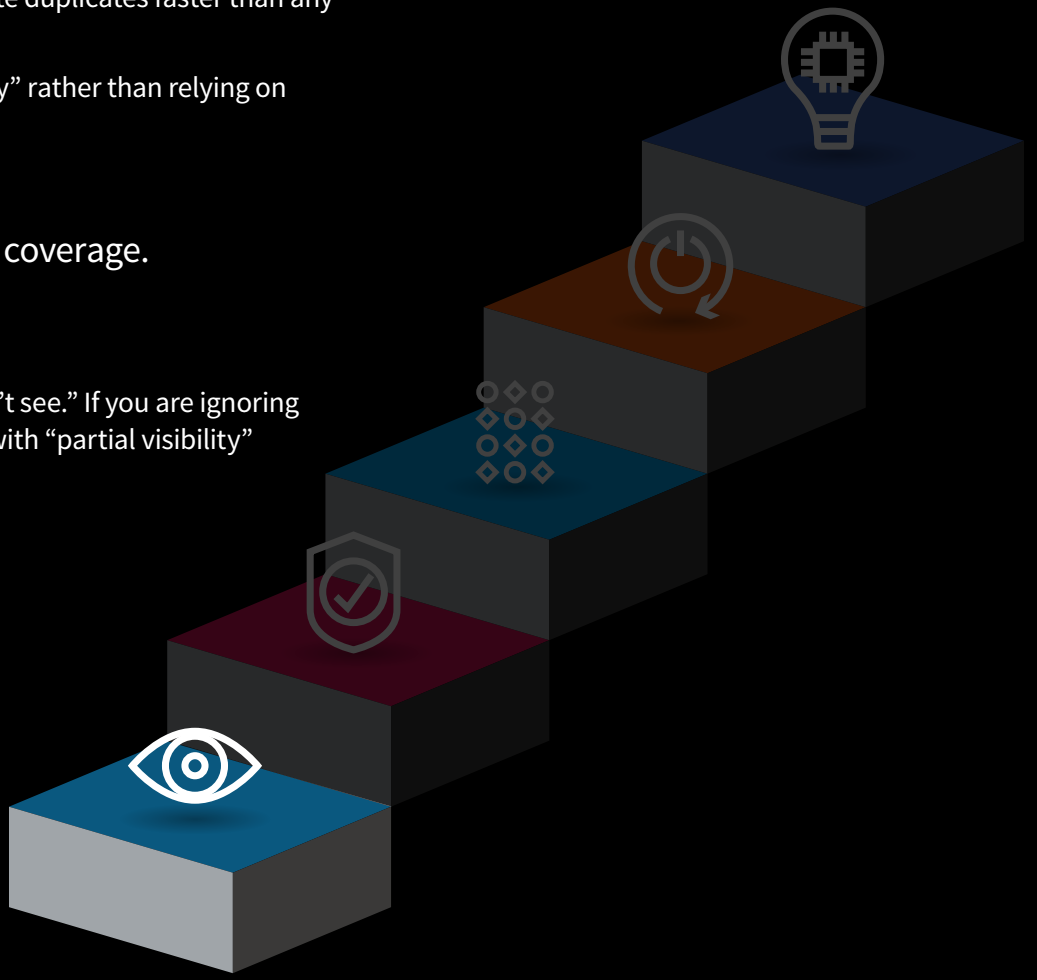

01 **Visibility (Stop Flying Blind):**
Establish a near real-time source of truth across on-premises, cloud, and SaaS environments.

Rung 2: Trust (Making the Data Usable)

The enthusiasm for AI in our field often leads organisations to skip a vital step. We want the magic of autonomous optimisation, but as **Ingrid** warns, "AI is not a magic solution, it's an enabler that demands strong foundations." Without these foundations, we risk moving from "manual errors" to "automated catastrophes."

Data Stewardship as the Backbone

To build trust, we must move beyond passive governance and into active stewardship. **Gareth Davies** defines this shift clearly: "As the operational arm of data governance, data stewardship ensures asset data is accurate, consistent, and usable throughout the asset lifecycle." He argues that stewardship is not a "standalone governance activity," but a "practical foundation for improving ITAM outcomes."

Without this backbone, the analytical power of AI remains a hollow promise. **Gareth** is blunt about the consequences: "Without accurate, well-governed asset data, AI-driven insights are unreliable". Furthermore, "Poor-quality or incomplete asset data leads directly to unreliable insights, suboptimal procurement decisions and missed optimisation opportunities."

The 'Garbage In, Garbage Out' Reality

The principle of 'rubbish in, rubbish out' takes on a new level of danger when applied to machine learning.

Ingrid captures the risk perfectly: "Without robust data foundations, AI in ITAM risks amplifying errors, creating blind spots, and leading organisations toward costly missteps." She reminds us that "AI is a magnifying glass. It amplifies the quality of your data. Clean, accurate data leads to valuable insights, while poor data magnifies errors, risks, and inefficiencies."

The Role of Validation

Trust is built through continuous validation. **Ingrid** explains that "Data validation ensures that discovered assets are accurate, normalised, and reconciled." She warns that if we neglect this, "AI may misclassify software, misinterpret usage metrics, or generate flawed forecasts."

Gareth suggests that the path to trust involves, "Defining who is responsible for the accuracy and security of specific data domains" and "appointing accountable data stewards to maintain standards across systems and teams."

02 **Trust (Make the Data Usable):**
Ensure AI-driven insights are reliable through rigorous data stewardship and governance.

What Good Looks Like: The Trust Checklist

- **Defined Ownership:** You have established clear ownership by defining who is responsible for specific data domains.
- **Standardised Governance:** “Clear data governance policies” are in place to ensure consistency across naming conventions and metadata.
- **Continuous Validation:** You perform “Continuous data validation” to maintain the integrity of AI-driven insights over time.
- **Reusable Data Products:** You are building “trustworthy data sets that can be used across multiple AI applications,” not just for ITAM.

If you only do one thing – Formalise the role of the Data Steward.

As **Gareth** suggests, don’t treat this as a sideline task. By assigning accountability for data quality, you ensure, “AI outputs can be trusted and acted upon” rather than ignored by skeptical stakeholders.

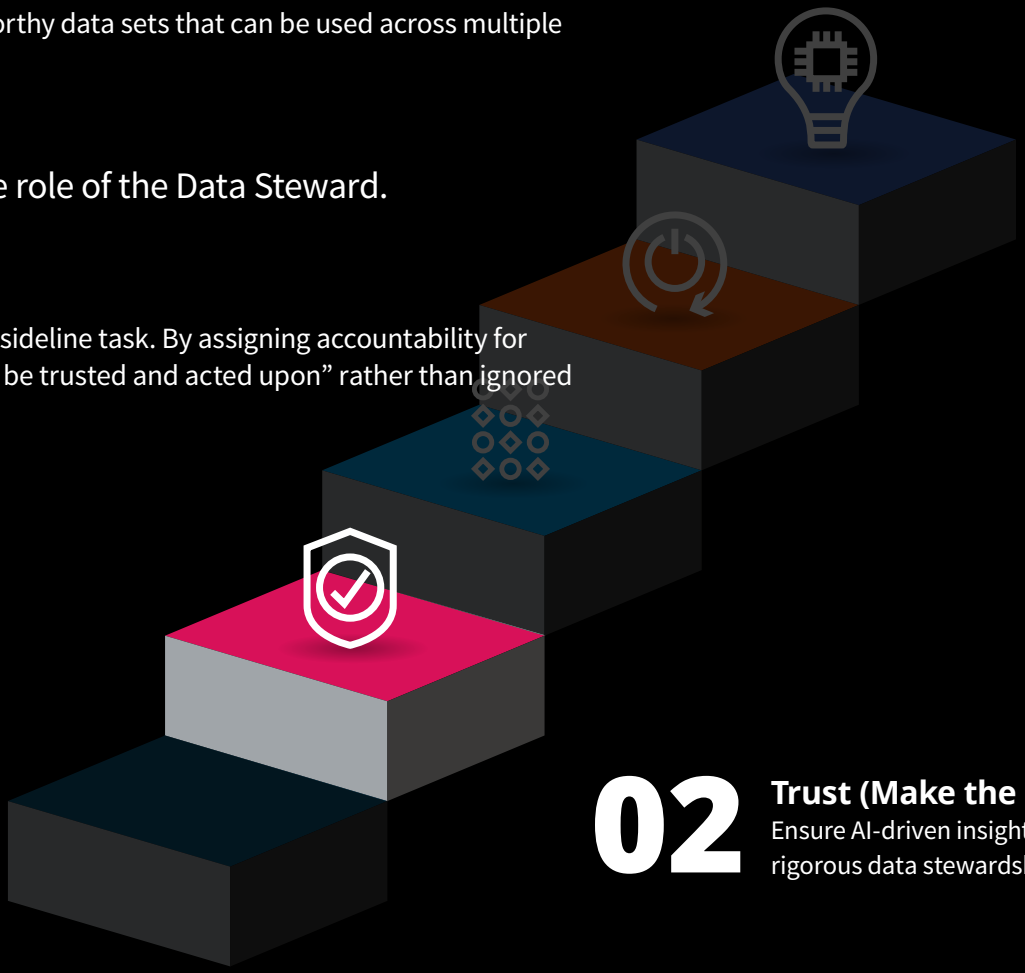

02 **Trust (Make the Data Usable):**
Ensure AI-driven insights are reliable through rigorous data stewardship and governance.

Rung 3: Insight (Spotting Patterns Humans Miss)

At this stage of the ladder, we move beyond simple record-keeping. As **Phil Perfetti** notes, “AI changes that equation by introducing speed, precision, and scalability at a level previously unimaginable.” This rung is where, “AI processes data at machine speed, identifies patterns humans would miss, and operates 24/7 without fatigue.”

Predictive Intelligence

The real power of Rung 3 lies in shifting from reactive firefighting to proactive management. **Tyrone Howland** explains, “AI can analyse consumption trends, seasonal patterns, and organisational behavior to predict future SaaS usage, renewal requirements, and cloud cost movements.” This allows us to look ahead; or as **Phil Perfetti** puts it, “Instead of reacting to compliance risks or overspending, ITAM teams can proactively optimise costs and negotiate smarter contracts.”

This predictive capability extends to the very hardware we manage. **Gareth** highlights, “AI also supports predictive asset lifecycle management by analysing historical and usage data to forecast device failures, renewal requirements, or software obsolescence.”

Uncovering Hidden Waste

Finally, AI excels at finding the leaks in your budget invisible to the naked eye. **Tyrone** states, “AI can detect dormant accounts, underused features, and opportunities for licence downgrades.” In complex cloud environments, it “identifies oversized instances, unused reservations, and misconfigurations that drive unnecessary costs.”

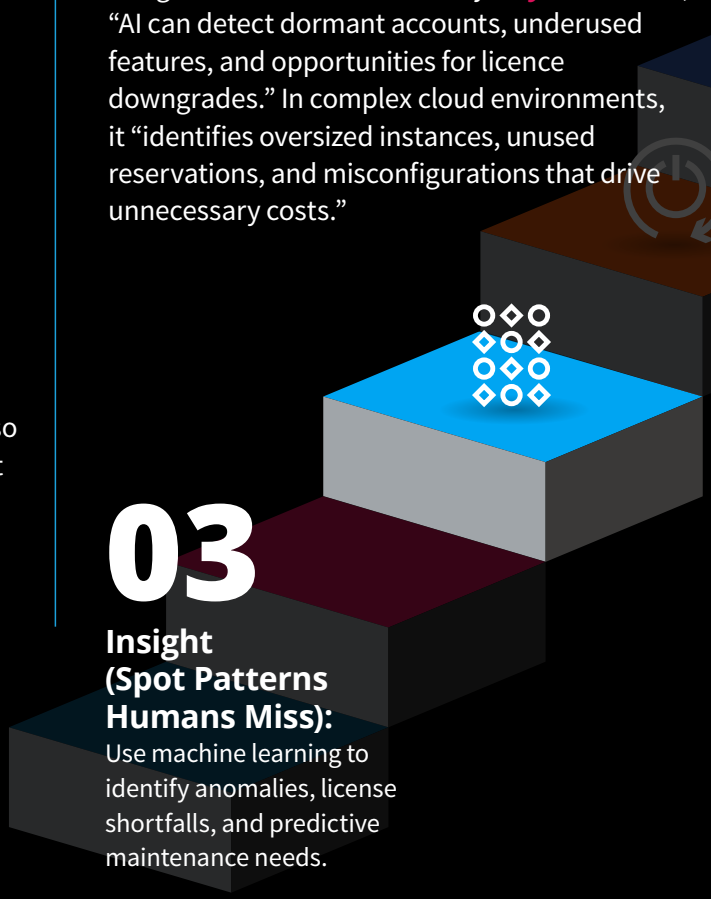

03

Insight (Spot Patterns Humans Miss):
Use machine learning to identify anomalies, license shortfalls, and predictive maintenance needs.

Contract and Risk Intelligence

One of the most tedious tasks in ITAM is the review of dense legal documents. AI is uniquely suited to solve this. **Tyrone** notes, “Contract intelligence tools can review licensing terms and highlight restrictive clauses or risk areas in moments.” This is echoed by **Stuart Moffat** who points out, “AI can highlight key clauses such as usage rights or privacy terms and link users directly to the relevant sections.”

Beyond contracts, AI acts as an early warning system for risk. **Phil Perfetti** observes, “AI can analyse patterns across millions of data points and flag irregularities that indicate potential breaches or non-compliance.” Similarly, **Tyrone** mentions, “Reconciliation engines can match entitlements to deployments continuously, identifying issues before they escalate.”

What Good Looks Like: The Insight Checklist

- **Trend Forecasting:** You are using “AI-driven analytics” to “forecast usage trends” and “predict license shortfalls” before they occur.
- **Anomaly Detection:** You have systems that “flag irregularities” in usage patterns that might indicate security breaches or non-compliance.
- **Automated Contract Review:** You use “contract intelligence” to identify “obligations, usage restrictions, and potential risks” hidden in legal documents.
- **Lifecycle Forecasting:** Your team can “forecast device failures” and “software obsolescence” based on historical usage data.

If you only do one thing – Apply AI-driven analytics to your most complex SaaS or Cloud contract.

As **Phil** notes, machine learning models can, “analyse historical usage patterns and predict when a subscription will become underutilised,” allowing you to intervene months before a renewal rather than days.

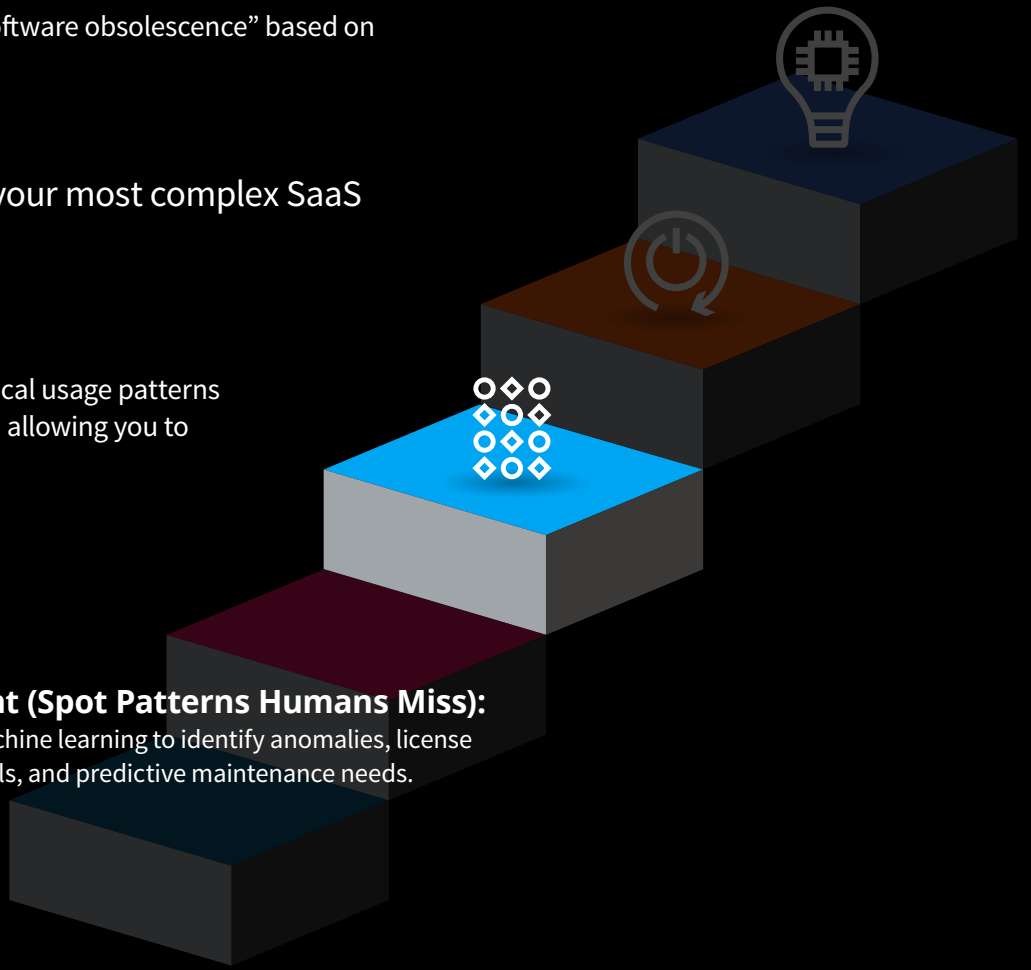

03 **Insight (Spot Patterns Humans Miss):**
Use machine learning to identify anomalies, license shortfalls, and predictive maintenance needs.

Rung 4: Action (Closing the Loop)

The transition to this rung is driven by a fundamental change in software economics. As I look at the landscape today, we can no longer rely on static reviews because the technology itself is no longer static. We must move toward a model where governance operates at the same speed as the software it manages.

The Reality of the Live Meter

The shift to consumption-based models has changed the stakes of ITAM. **Stephanie Day** observes, "Consumption-based pricing converts software from a known expense into a live meter." In this environment, "Usage accumulates quietly, minute by minute, across users, integrations, and automations." Because of this, she argues, "The goal is not to slow teams down, but to give the business a shared, real-time understanding of usage, cost, and risk, along with guardrails that operate at the speed of modern software."

Automating the Response

To handle this "live meter," we must move beyond manual intervention. **Luís Miguel Martins** describes the use of 'ITAM bots' that "automate license assignment and harvesting, generate ELPs, flag anomalies in cost or risk, and prepare foundational renewal documentation."

This automation allows for immediate correction of issues. **Victor Paredes-Colonia** suggests that teams should, "Auto-remediate low-risk compliance issues within thresholds." **Stuart** adds that at this level of maturity, "Agents perform tasks like updating records or classifying software."

Orchestrated Workflows

The final piece of this rung is the transition to automated execution. **Tyrone** explains that in advanced functions, "AI insights automatically trigger operational actions, such as reclaiming unused licences or correcting cloud tags." This ensures that optimisation isn't just a recommendation, it's a result.

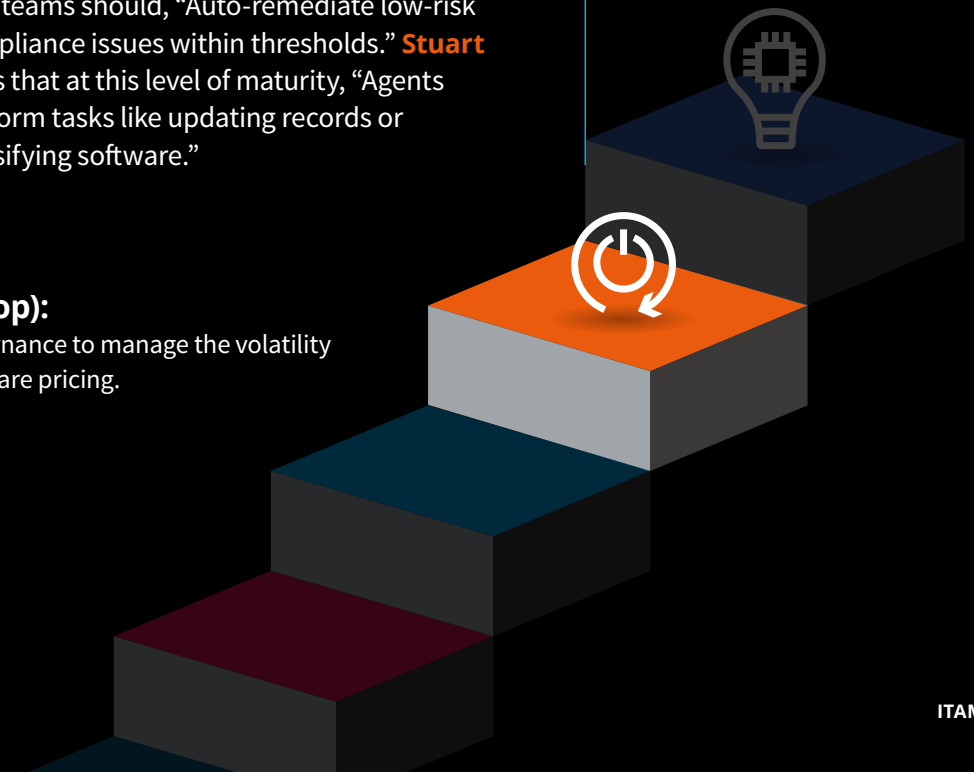

04 **Action (Close the Loop):**
Implement 'always-on' governance to manage the volatility of consumption-based software pricing.

What Good Looks Like: The Action Checklist

- **Continuous Oversight:** You have replaced “snapshots with live data” to match software that “prices itself continuously”.
- **Automated Reclaiming:** You utilise ‘ITAM bots’ to “automate license assignment and harvesting”.
- **Auto-Remediation:** You “auto-remediate low-risk compliance issues” based on defined “thresholds”.
- **Orchestrated Action:** AI-driven “insights automatically trigger operational actions” such as “correcting cloud tags”.

If you only do one thing – Review your “low-risk compliance issues”, and identify what can be auto-remediated.

As **Victor** notes, moving toward “Auto-remediate low-risk compliance issues within thresholds” is a key step in scaling your capacity without increasing headcount.

04 Action (Close the Loop):

Implement ‘always-on’ governance to manage the volatility of consumption-based software pricing.

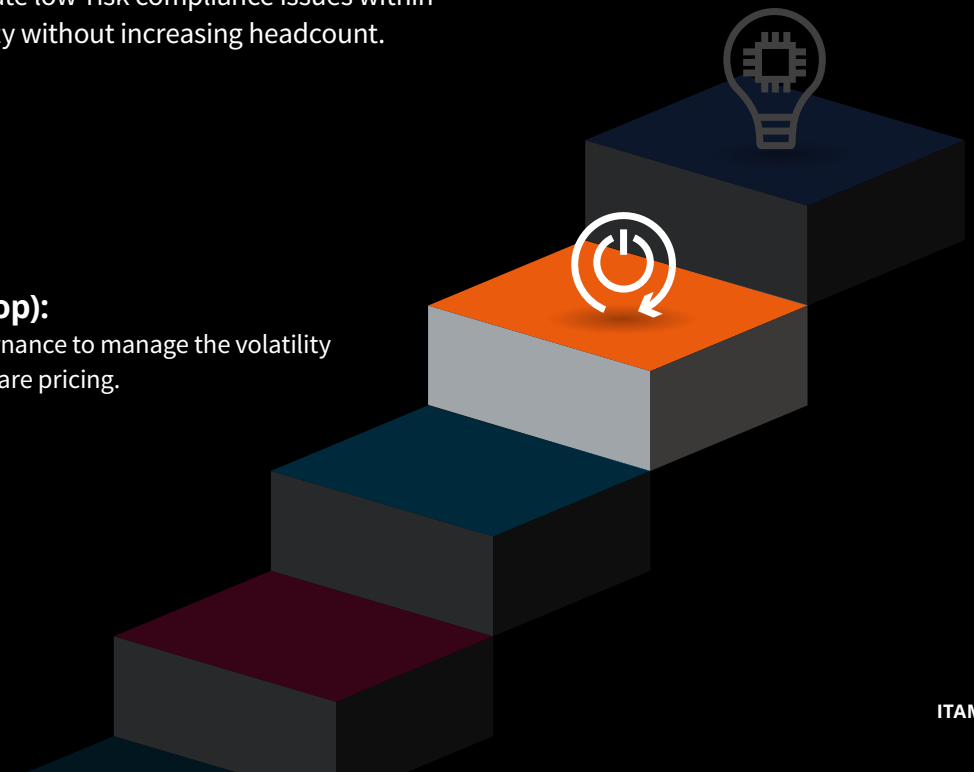

Rung 5: Autonomy (Bounded)

The final rung of the ladder is defined by the transition from tools to "team members." As **Victor** provocatively states, "The real opportunity lies in treating AI as a team member you develop over time." This requires a mindset shift: "Define clear guardrails, then let AI operate independently within them."

The Evolution of the Agent

This stage is the culmination of the journey from simple support to proactive partnership. **Stuart** identifies this as the final stage of adoption: "Autonomy is the final stage. Agents act proactively, make recommendations, and execute tasks within defined limits." He explains that these agents "adapt in real time and escalate only when human input is needed, becoming strategic partners in ITAM."

The result of this autonomy is a fundamental change in our daily work. **Victor** notes, "ITAM roles evolve from firefighting toward strategic advising and proactive insights." In this environment, systems can "auto-allocate licenses for standard requests but escalate unusual patterns" and "recommend lifecycle decisions for commodity hardware, but flag anomalies."

The Human Anchor

Even at the peak of the ladder, human judgment remains the indispensable anchor. **Stuart** reminds us, "Automation does not mean loss of control. Human oversight remains essential." He emphasises, "Users must be able to review, approve, or adjust actions before they are executed."

Phil Perfetti reinforces this, stating, "The future of ITAM will be defined by how well we combine machine intelligence with human insight." He notes that human expertise is irreplaceable in areas like strategic decision making, ethical oversight, and vendor relationships.

As we reach this level of maturity, the partnership becomes the prize. **Victor** concludes those who "treat AI as a team member to develop rather than a script to run, will define the next era of IT Asset Management."

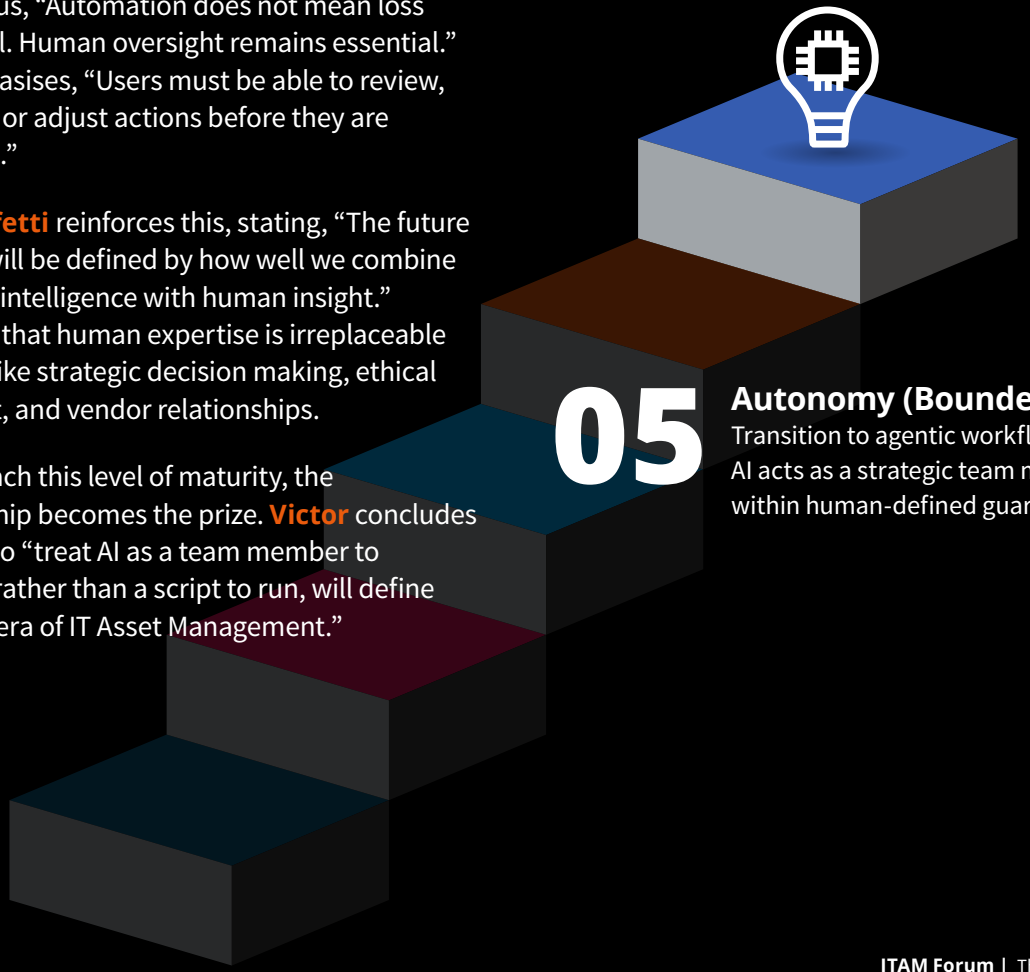

05

Autonomy (Bounded):
Transition to agentic workflows where AI acts as a strategic team member within human-defined guardrails.

What Good Looks Like: The Autonomy Checklist

- **Proactive Collaboration:** AI agents "anticipate needs, suggest actions, and execute them with human oversight."
- **Bounded Independence:** Systems "act proactively, make recommendations, and execute tasks within defined limits."
- **Strategic Shift:** The "new ITAM professional spends less time on operations, more time influencing technology investments."
- **Shared Accountability:** Humans "define strategy, set policies, and own ultimate accountability, even when AI executes decisions within those boundaries."

If you only do one thing – Begin the "onboarding" process for your AI.

As **Victor** suggests, treat it like a new hire. "Teach them business context, watch early work closely, and gradually give them autonomy within defined boundaries."

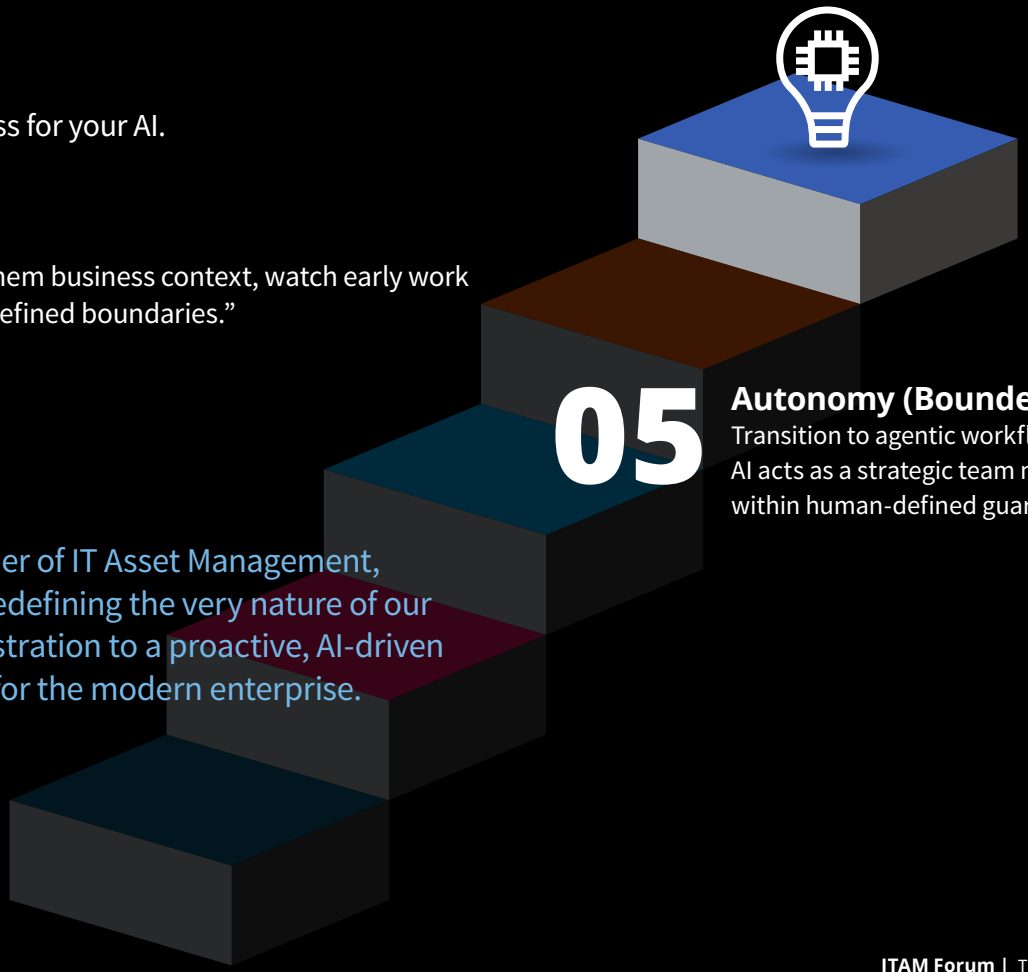

05

Autonomy (Bounded):
Transition to agentic workflows where AI acts as a strategic team member within human-defined guardrails.

As we conclude this exploration into the intelligent frontier of IT Asset Management, it's clear we are not merely adopting new tools; we are redefining the very nature of our profession. The transition from manual, reactive administration to a proactive, AI-driven discipline is no longer a choice. It's a strategic necessity for the modern enterprise.

The Executive Summary: A Strategic Roadmap

By climbing the AI Ladder, organisations can move from the chaos of fragmented data to the clarity of autonomous governance. This executive summary synthesises the collective wisdom of our eight contributors into a final roadmap for the future.

1. The Foundation: Visibility and Trust

The journey begins with a commitment to data integrity.

As **Ingrid** warns, "Without robust data foundations, AI in ITAM risks amplifying errors, creating blind spots, and leading organisations toward costly missteps."

Gareth reinforces this, noting, "Without accurate, well-governed asset data, AI-driven insights are unreliable," making data stewardship the "practical foundation for improving ITAM outcomes."

We must move toward "near real-time asset discovery" to manage what **Stephanie** describes as the "live meter" of consumption-based pricing.

2. The Engine: Insight and Action

With a trusted foundation, AI introduces "speed, precision, and scalability at a level previously unimaginable."

Phil highlights that AI-driven analytics can now "forecast usage trends, predict license shortfalls and highlight opportunities for rightsizing." These insights must lead directly to execution.

Luís Miguel introduces the concept of "ITAM bots" that "automate license assignment and harvesting, generate ELPs, flag anomalies in cost or risk." This ensures optimisation happens at the speed of the software itself, rather than during annual reviews.

3. The Vision: Bounded Autonomy and Partnership

The summit of the ladder is a new operating model: Augmented ITAM.

Victor challenges us to treat AI as "a team member you develop over time," moving toward a state where we, "Define clear guardrails, then let AI operate independently within them."

Stuart reminds us that this journey from "support" to "autonomy" is a deliberate progression where agents eventually "act proactively, make recommendations, and execute tasks within defined limits."

Final Takeaways for ITAM Leaders

1. **Prioritise Stewardship:** Treating data as an asset involves establishing clear ownership and recognising that high-quality data is an insurance policy against AI failure.
2. **Balance Machines and Minds:** While machines provide speed and scale, human expertise is irreplaceable for strategic decisions, ethical oversight, and vendor negotiations.
3. **Embrace Continuous Governance:** Traditional ITAM assumptions are broken. In an environment where applications are always on, governance must be as well.
4. **Scale Without Headcount:** By offloading repetitive tasks to AI, teams can scale capacity without scaling headcount and pivot their focus toward strategy, governance, and innovation.

If you only do one thing – Begin the "onboarding" process for your AI.

As **Tyrone** concludes, the future of ITAM is a partnership where, "machines handle the scale, and humans apply interpretation and direction."

The question is no longer if AI will transform our field, but, how fast we are willing to climb.