

How to Successfully Choose an IT Dependency Mapping Tool

3 September 2025 - 15 min read

By: Ankita Hundal, Roger Williams

Initiatives: [I&O Operations Management](#); [Evolve Service Management and Cloud Operations](#)

Heads of infrastructure and operations struggle to realize value from IT dependency mapping tools due to unclear goals, narrow vendor reviews and mismatched solutions. Our insights help set clear objectives, compare sourcing options and show how to validate tools with a focused proof of concept.

Overview

Key Findings

- The effectiveness of IT dependency mapping tools varies significantly because of protocol coverage, credentials, encrypted traffic, dynamic cloud and container topologies, and data access limits. This gap leads to inaccurate maps and low data trust.
- A single vendor or sourcing path often misses needs and inflates cost. Pricing models, data portability and integration depth differ between vendors and change total cost.
- Vague goals, weak sponsorship and insufficient testing frequently undermine dependency mapping tool deployments. This – combined with missing data owners and success metrics – leads to tools that perform well in demos, but fail to deliver in production.

Recommendations

- Establish a one-page charter for dependency mapping by defining expected outcomes, the specific data to be mapped, data owners and access owners to ensure clear goals, alignment, active sponsorship and measurable success.
- Evaluate suite, domain and stand-alone options by using a simple scorecard on protocol, coverage, data access and integration effort. Use this to select the best fit and reduce total cost of ownership.
- Conduct a rigorous proof of concept (POC) on critical and odd systems by testing critical data and organizational constraints to validate coverage, ensure accurate maps and reduce roll-out risk.

Introduction

Heads of I&O are faced with the challenge of ensuring infrastructure reliability and agility as digitization and competitive dynamics intensify. As organizations accelerate cloud and hybrid infrastructure adoption, the ability to map component relationships and dependencies accurately becomes essential to minimize unplanned outages, track availability and enable effective service delivery (see [Note 1](#) for a list of representative vendors of IT dependency mapping capabilities).

Mapping is inherently hard: results depend on vendor design, access to credentials, protocol handling and fit with your network, cloud and container stack. Poor dependency data, whether at application or service layer, can derail strategic initiatives, delay migrations and erode I&O's credibility. Frequent change and security controls add friction. The wrong tool selection can result in wasted investment, project delays and operational blind spots. Coverage needs to vary by environment – whether cloud, on-premises or hybrid – so global “best” vendor rankings are rarely useful.

Organizations cannot expect to simply buy one tool and have it meet their requirements. Instead, align tool selection with your organization's protocols, platforms, containers and data access requirements, and use a tailored, evidence-based process to meet business and technical needs. When approached this way, dependency mapping becomes a critical enabler of outcomes that matter most to the business, including reducing unplanned outages, accelerating migrations, improving change enablement, supporting audit and compliance and enabling faster root cause analysis.

Given these challenges, heads of I&O must ensure that dependency mapping investments deliver reliable, service-level, actionable insights that can be trusted across hybrid environments – from cloud migrations to on-premises improvements. This research provides a practical, three-step process to achieve that outcome, for dependency mapping success (see Figure 1).

Figure 1: Three Steps to Enable IT Dependency Mapping Success

Three Steps to Enable IT Dependency Mapping Success



Source: Gartner
832441

Analysis

Set Clear Project Goals and Ownership for Dependency Mapping

Dependency mapping supports two main categories of IT initiatives:

- **Projects** depend on accurate mapping to uncover hidden dependencies and reduce the risk of unforeseen disruptions. Examples of these types of projects include data center moves, on-premises-to-cloud migration, or an infrastructure upgrade
- **Service improvement programs (SIPs)** use the data from dependency mapping to streamline decision making and optimize ongoing IT service delivery. Examples of these types of SIPs include efforts to reduce downtime, decrease change failures or prevent false alerts.

Projects typically demand less process rigor than ongoing programs, as their defined timelines, focused stakeholder involvement and stable data requirements simplify management. In contrast, service improvement programs require sustained process maturity to ensure data remains current and relevant as environments evolve. Table 1 compares projects and ongoing service improvement programs.

Table 1: Project vs. Service Program Comparison

Aspect	Projects	Ongoing service improvement programs
Primary use of maps	Reduce migration and upgrade risk	Cut downtime and reduce change failures
Process rigor required	Typically less, due to defined timelines and stable data requirements	Higher – requires sustained process maturity as environments evolve
Stakeholder involvement	More focused and limited to project stakeholders	Broader, ongoing involvement across teams
Data requirements	Stable and well defined for project duration	Must remain current and relevant over time
Management complexity	Simpler, due to clear start/end and scope	More complex, due to continuous updates and integration with other processes

Source: Gartner

To ensure dependency mapping investments are aligned with business objectives, heads of I&O should establish a formal one-page charter for each project or program, specifying:

- Purpose
- Funding sources
- Stakeholders (including managed service partners, if relevant)
- Data scope

- Privacy and retention
- Measurable success metrics

For projects, define how maps will be used and what the post-cutover plans are. For programs, specify data ownership, integration with configuration management database (CMDB) and change management platforms, use cases, and the assets – or attributes – to be discovered to prevent scope drift and maximize value.

As part of goal setting, validate whether current tooling already meets these needs, or if gaps are due to people, process or other factors rather than technology alone. Clearly articulate must-have requirements for new tooling to address unmet needs. Heads of I&O should also engage the CISO early in the process to ensure that security policies, credential management and access approvals required for discovery and POC are identified and secured before moving forward. Without these, the initiative cannot proceed.

If a project will transition into a program, plan licenses, maintenance and resource needs early. Review existing tools for overlap, avoid redundant investments and ensure leadership commitment for both initial and ongoing support.

Compare Suite, Domain and Stand-Alone Options Using a Scorecard

No single tool can address every use case. To ensure the selected solution aligns with both technical requirements and budget constraints, heads of I&O should enable proactive evaluation of a range of vendors and sourcing paths, rather than defaulting to a single provider or existing partner.

Check how you will get data out, how the product roadmap aligns with your plans and what it will take to run at scale. This approach reduces the risk of vendor lock-in and increases the likelihood of finding a solution that best fits your use cases.

Understanding Sourcing Paths

The main sourcing paths for dependency mapping tools can be grouped into three categories:

- **Suite:** Tools embedded within larger ITSM/ITOM platforms. These tools are explicitly designed to support service modeling and tight integration with CMDBs, governance features and compliance controls. These tools go beyond raw discovery (or ADM) and can represent business services, CIs and relationships inside a governance framework. They typically require longer deployments and higher investment but are a strong choice for organizations prioritizing service governance and life cycle integration.
- **Domain:** These tools are designed to monitor application performance, transactions and telemetry and are delivered by observability or application portfolio management (APM) vendors. They build runtime service maps by tracing flows, API calls and dependencies across distributed systems, microservices, containers and cloud-native stacks. By themselves, they are more ADM-like, but can feed into service dependency mapping (SDM) by exporting data in a CMDB or providing freshness and accuracy that SDM programs often lack. They are typically great at accuracy in dynamic environments but lacking when it comes to governance and service-level modeling.
- **Stand-alone:** Specialized independent tools focused specifically on discovery and dependency mapping. These provide fast time-to-value, flexible data portability and independence from large suites. They often cover infrastructure and ADM, and have begun layering basic service modeling. These tools are well-suited for tactical projects (like migrations or data center exits), environments where exportable data can be linked to CMDBs, and organizations seeking affordable solutions for cost-sensitive environments. Though they may lack the governance depth of suite options.

Heads of I&O should consider the following key categories, among others, when comparing dependency mapping providers. While the categories below represent the most critical evaluation dimensions for dependency mapping tools, they are not exhaustive. Heads of I&O should tailor assessments to their own environment, adding organizational criteria such as organizational maturity, operating model or sector-specific regulatory demands. Along with an outline of key assessment categories, Table 2 also provides a comprehensive comparison of sourcing paths and testing methods, success signals, potential red flags and use-case fit to guide your selection process.

By mapping vendors against these criteria, leaders can align sourcing choice – suite, domain or stand-alone – with the outcomes their environment requires.

Table 2: Dependency Mapping Sourcing Path Comparison Matrix
(Enlarged table in Appendix)

Category	Sourcing		
	Suite: Governance-driven, strong CMDB integration, regulated	Domain: Ops/dynamic environments, cloud-native stacks, real-time freshness is priority	Stand-alone: Tactical projects/migrations, quick visibility, portability
Strategic fit	Test: Map 2-3 business services into CMDB Pass: CI classes align cleanly, service models tender without workarounds, workflows intact. Red flag: Requires extensive customization or results in flat device lists.	Test: Push service-dependency data via API into ITSM or CMDB Pass: Relationships preserved, available in less than 24h. Red flag: Only raw flows export, CMDB context lost, throttled APIs.	Test: Export/import service maps into CMDB Pass: Relationships and meta-data preserved with minimal cleanup. Red flag: Only CSV/device lists exported, service-level context lost.
Tag & meta-tag management	Test: Bulk import tags into CMDB, validate ephemeral resource tagging Pass: Automated ingestion, tags linked to services. Red flag: Manual-only tagging, no ephemeral support.	Test: Ingest tags from cloud providers. Pass: Dynamic cloud tags consistently appear in service maps. Red flag: Partial or inconsistent tagging, manual reconciliation required.	Test: Apply tags to legacy and cloud assets, bulk update. Pass: Metatags preserved across environments. Red flag: Flat tagging (device-only) with no service-level link.
Coverage (protocols, cloud, containers)	Test: Discover infra across hybrid stack (servers, DBs, SaaS) Pass: All major protocols/services visible Red flag: Missing modern protocols or SaaS coverage	Test: Run discovery in container/serverless workloads. Pass: Real-time capture of microservices and PaaS. Red flag: Infra blind spots, fails on legacy protocols.	Test: Scan legacy servers, VMs, apps. Pass: Strong legacy/on-premises visibility, some hybrid cloud Red flag: No depth for PaaS, containerized workloads
Accuracy, freshness & change tracking	Test: Shut down/move a known service. Pass: Map updated in less than 6 hours, few false links. Red flag: Daily batch refresh only, stale dependencies.	Test: Run load test with ephemeral containers. Pass: Updates in less than 5 min, accurate removal of terminated nodes. Red flag: Missed ephemeral assets or persistent ghost links.	Test: Compare output with a golden truth set. Pass: >75-85% recall/precision, low false positives. Red flag: Frequent spurious or missing links.
Scale & reliability	Test: Deploy 3 collectors across sites, simulate failover. Pass: HA works, no map gaps, auto-resync in less than 10 min. Red flag: Collector outage breaks maps.	Test: Burst container count 5x. Pass: No discovery lag, PU/mem overhead less than 10%. Red flag: Collector bottlenecks, lag less than 15 min.	Test: Run 2-3 business services simultaneously. Pass: Stable outputs at project scale. Red flag: Fragile or fails beyond single-site deployments.
Data model & CMDB fit	Test: Import maps into CMDB staging Pass: Classes mapped correctly, merge rules apply Red flag: Flat device lists, duplicates, no soft delete	Test: Export via API to CMDB Pass: Stable mapping into CI tables. Red flag: Loss of service model fidelity.	Test: Export/import maps with tags Pass: Relationships and metadata intact. Red flag: CSV-only or incomplete export.
Integration & APIs	Test: Sync service to ITSM and back. Pass: APIs bulk-write, webhooks fire changes. Red flag: Rate limits or no bidirectional sync.	Test: Connect to AIOps/observability. Pass: Event-driven, API-first, low-latency updates. Red flag: Weak ITSM tie-in, monitoring-only.	Test: Bulk export/import via REST. Pass: Works with ITSM APIs, repeatable. Red flag: Manual flat-file transfer.
Portability & exit	Test: Export all maps from CMDB Pass: Full export, open formats, reimport possible. Red flag: Export blocked, paywalled, or lossy.	Test: API export of maps to JSON/YAML. Pass: High-fidelity export/import to other tools. Red flag: Proprietary formats, incomplete exports.	Test: Export/import with open schema. Pass: Preserves service relationships and history. Red flag: Only device/asset list exported.
Security & compliance	Test: Vault integration, RBAC, audit logging. Pass: Fine-grained access, auditable, region controls. Red flag: Shared admin accounts, uncontrolled residency.	Test: Cloud discovery scoped with least privilege. Pass: Scoped creds and audit logs visible. Red flag: Over-permissioned discovery, no logs.	Test: Review secret storage method. Pass: Read-only, audited. Red flag: Hardcoded or broad credentials.
Cost & TCO	Test: Build 3-year TCO model with services and support. Pass: Transparent licensing, roadmap aligned. Red flag: Hidden add-ons, pivot risk.	Test: Compare license vs. workloads monitored. Pass: Costs predictable and scalable. Red flag: Volatile SaaS costs or opaque tiers.	Test: Pilot 2-3 projects with pricing model. Pass: Low entry, portable, no hidden fees. Red flag: Weak roadmap or vendor instability.

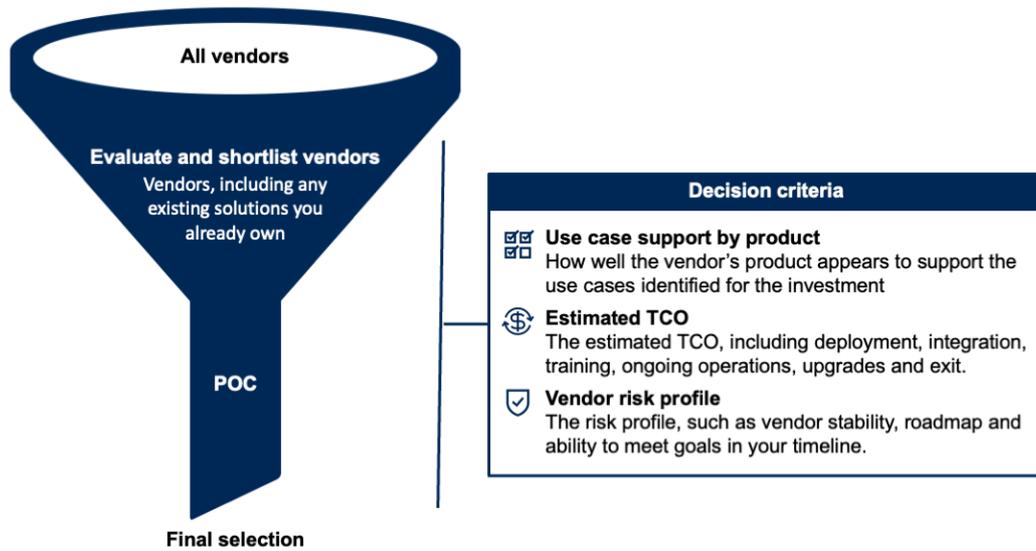
Source: Gartner

After conducting the dependency mapping, shortlist two to four vendors, including any existing solutions you already own (see Figure 2). Assess each vendor based on:

- How well the vendor’s product appears to support the use cases identified for the investment.
- The estimated TCO, including deployment, integration, training, ongoing operations, upgrades and exit.
- The risk profile, such as vendor stability, roadmap and ability to meet goals in your timeline.

Figure 2: Vendor Shortlisting Funnel

Vendor Shortlisting Funnel



Source: Gartner
832441

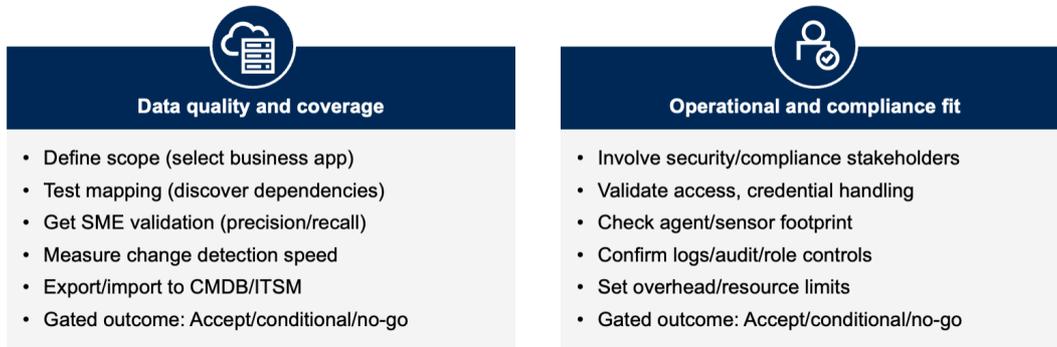
Conduct a Rigorous POC Against Critical Data and Organizational Constraints

After shortlisting dependency mapping solutions, heads of I&O should require a written POC in the target environment before making a purchase decision. Ensure contractual terms allow for comprehensive testing and clear data handling terms clarifying expectations around liability and indemnification. Avoid committing to any tool that does not permit in-environment validation.

Structure the POC around two key validation areas (see Figure 3):

Figure 3: Structure the POC around two key validation areas

Structure the POC Around Two Key Validation Areas



Source: Gartner
832441



Data Quality and Coverage

Define a granular scope for your mapping initiative by selecting specific business applications to map to their underlying infrastructure. IT teams often start with their own systems or partner with a product owner to focus on a particular application, ensuring the scope is manageable and relevant.

- Prioritize testing on critical and high-risk systems, including those that are unique, legacy or industry-specific.
- Engage technical subject matter experts to assess whether the tool accurately discovers and maps dependencies for these assets, ensuring the data meets the organization’s operational goals.
- Measure precision and recall against a truth set.
- Track time to the first useful map.
- Track time to detect a change.
- Aim for zero false links on crown jewels.

Validate that the tool:

- Discovers all in-scope components and maps true dependencies (including ephemeral and managed cloud services).
- Produces decision-grade outputs (service views usable by change, incident, migration teams).
- Detects topology changes within the time window your process actually requires (e.g., before CAB, during incident triage).
- Can export/import full-fidelity data into your CMDB/ITSM.

Gated Outcomes:

- **Accept (Go):** Crown-jewel services correctly mapped per SME validation; residual gaps are low-risk with workarounds; change updates arrive within your process window; export/import validated.
- **Conditional accept (Go with mitigations):** One or two material gaps but with compensating controls and a vendor-committed remediation path on an agreed timeline.
- **No-go:** Misleading links on crown-jewels, inability to generate service-level outputs, or lack of data portability.

Operational and Compliance Fit

Evaluate the tool's performance under real world organizational constraints.

- Involve security and compliance stakeholders to confirm that required access levels, credential storage and handling and agent deployments are compatible with enterprise policies and do not introduce unacceptable risk.
- Validate that the agent or sensor footprint does not negatively impact system performance or business operations. Involve security and compliance.
- Confirm logs, audit and role controls meet policy.
- Set overhead limits.

Stress test under operational realities with security, networking and platform stakeholders:

- **Access pattern:** Prefer outbound-only collectors via approved proxies/brokers. If inbound is required, document exceptions.

- Secrets and RBAC: Use your enterprise vault with rotation, least-privilege scopes and auditable usage.
- Footprint and reliability: Validate agent/sensor/collector resource budgets; prove no single point of failure via controlled failover; ensure logs/events integrate with SIEM.

Gated Outcomes:

- **Accept (Go):** Operates within policy for access, secrets, RBAC and resource budgets; HA collectors and failover succeed; compliance review passed with evidence.
- **Conditional accept:** Minor variances mitigated with tuning, dedicated nodes or maintenance windows; residual risk accepted formally.
- **No-go:** Unapproved inbound holes, secrets in-app without audit or material performance impact.

Decision Rule

Proceed if both areas are “accept” or one is “accept” and the other is “conditional” with signed mitigations and timelines. If the tool falls short but delivers relative improvement (e.g., faster triage, fewer failed changes, credible migration cutovers), document the gaps and proceed only with explicit compensations. **Do not buy on roadmap promises.**

Deliverables

Require a POC dossier that includes:

- SME-validated service maps
- Truth-set comparisons
- Change-detection evidence
- Export/import artifacts
- Resource-budget observations
- Access/secret/RBAC audit logs
- An integration demo into CMDB/ITSM

By insisting on a rigorous, stakeholder-driven POC, heads of I&O can ensure the selected dependency mapping solution is technically sound, operationally viable and compliant with organizational requirements maximizing the likelihood of a successful and valuable deployment.

Evidence

Based on analysis of Gartner client inquiries on IT dependency mapping.

Note 1: Representative Vendors That Offer IT Dependency Mapping Capabilities

List of representative vendors that offer IT SDM capabilities (by sourcing path):

- Suite vendors offering IT dependency mapping capabilities:
 - [Atlassian](#)
 - [BMC](#)
 - [EasyVista](#)
 - [Freshworks](#)
 - [IBM](#)
 - [IFS](#)
 - [Ivanti](#)
 - [ManageEngine](#)
 - [Matrix42](#)
 - [OMNINET](#)
 - [OpenText](#)
 - [ServiceNow](#)
 - [SolarWinds](#)
 - [SymphonyAI](#)
 - [USU](#)

- Domain vendors offering IT dependency mapping capabilities:
 - Amazon Web Services
 - Cisco
 - CloudSphere
 - Dynatrace
 - Datadog
 - Eracent
 - Flexera
 - Microsoft Azure
 - New Relic
 - Qualys
 - ReadyWorks
 - Resolve Systems
 - ScienceLogic
 - Splunk
 - Tanium

- Stand-alone vendors offering dependency mapping capabilities:
 - Faddom
 - Device42
 - Lansweeper
 - NetBrain

Document Revision History

How to Successfully Choose an IT Asset Discovery Tool for Service Dependency Mapping
- 6 July 2021

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[IT Score for Infrastructure and Operations](#)

[How to Build a Cloud Migration Cost Estimate](#)

© 2025 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

Table 1: Project vs. Service Program Comparison

Aspect	Projects	Ongoing service improvement programs
Primary use of maps	Reduce migration and upgrade risk	Cut downtime and reduce change failures
Process rigor required	Typically less, due to defined timelines and stable data requirements	Higher – requires sustained process maturity as environments evolve
Stakeholder involvement	More focused and limited to project stakeholders	Broader, ongoing involvement across teams
Data requirements	Stable and well defined for project duration	Must remain current and relevant over time
Management complexity	Simpler, due to clear start/end and scope	More complex, due to continuous updates and integration with other processes

Source: Gartner

Table 2: Dependency Mapping Sourcing Path Comparison Matrix

	Sourcing		
Category	Suite: Governance-driven, strong CMDB integration, regulated	Domain: Ops/dynamic environments, cloud-native stacks, real-time freshness is priority	Stand-alone: Tactical projects/migrations, quick visibility, portability
Strategic fit	<p>Test: Map 2-3 business services into CMDB.</p> <p>Pass: CI classes align cleanly, service models render without workarounds, workflows intact.</p> <p>Red flag: Requires extensive customization or results in flat device lists.</p>	<p>Test: Push service-dependency data via API into ITSM or CMDB.</p> <p>Pass: Relationships preserved, available in less than 24h.</p> <p>Red flag: Only raw flows export; CMDB context lost; throttled APIs.</p>	<p>Test: Export/import service maps into CMDB.</p> <p>Pass: Relationships and metadata preserved with minimal cleanup.</p> <p>Red flag: Only CSV/device lists exported, service-level context lost.</p>
Tag & meta-tag management	<p>Test: Bulk import tags into CMDB; validate ephemeral resource tagging.</p> <p>Pass: Automated ingestion, tags linked to services.</p> <p>Red flag: Manual-only tagging, no ephemeral support.</p>	<p>Test: Ingest tags from cloud providers.</p> <p>Pass: Dynamic cloud tags consistently appear in service maps.</p> <p>Red flag: Partial or inconsistent tagging, manual reconciliation required.</p>	<p>Test: Apply tags to legacy and cloud assets, bulk update.</p> <p>Pass: Metatags preserved across environments.</p> <p>Red flag: Flat tagging (device-only) with no service-level link.</p>
Coverage (protocols, cloud, containers)	<p>Test: Discover infra across hybrid stack (servers, DBs, SaaS).</p> <p>Pass: All major protocols/services visible.</p>	<p>Test: Run discovery in container/serverless workloads.</p> <p>Pass: Real-time capture of microservices and PaaS.</p>	<p>Test: Scan legacy servers, VMs, apps.</p> <p>Pass: Strong legacy/on-premises visibility, some hybrid cloud.</p>

	Red flag: Missing modern protocols or SaaS coverage.	Red flag: Infra blind spots; fails on legacy protocols.	Red flag: No depth for PaaS, containerized workloads.
Accuracy, freshness & change tracking	<p>Test: Shut down/move a known service.</p> <p>Pass: Map updated in less than 6 hours; few false links.</p> <p>Red flag: Daily batch refresh only; stale dependencies.</p>	<p>Test: Run load test with ephemeral containers.</p> <p>Pass: Updates in less than 5 min; accurate removal of terminated nodes.</p> <p>Red flag: Missed ephemeral assets or persistent ghost links.</p>	<p>Test: Compare output with a golden truth set.</p> <p>Pass: ≥75-85% recall/precision; low false positives.</p> <p>Red flag: Frequent spurious or missing links.</p>
Scale & reliability	<p>Test: Deploy 3 collectors across sites, simulate failover.</p> <p>Pass: HA works, no map gaps, auto-resync in less than 10 min.</p> <p>Red flag: Collector outage breaks maps.</p>	<p>Test: Burst container count 5x.</p> <p>Pass: No discovery lag; PU/mem overhead less than 10%.</p> <p>Red flag: Collector bottlenecks, lag less than 15 min.</p>	<p>Test: Run 2-3 business services simultaneously.</p> <p>Pass: Stable outputs at project scale.</p> <p>Red flag: Fragile or fails beyond single-site deployments.</p>
Data model & CMDB fit	<p>Test: Import maps into CMDB staging.</p> <p>Pass: Classes mapped correctly, merge rules apply.</p> <p>Red flag: Flat device lists, duplicates, no soft delete.</p>	<p>Test: Export via API to CMDB.</p> <p>Pass: Stable mapping into CI tables.</p> <p>Red flag: Loss of service model fidelity.</p>	<p>Test: Export/import maps with tags.</p> <p>Pass: Relationships and metadata intact.</p> <p>Red flag: CSV-only or incomplete export.</p>
Integration & APIs	<p>Test: Sync service to ITSM and back.</p> <p>Pass: APIs bulk-write, webhooks fire changes.</p> <p>Red flag: Rate limits or no bidirectional sync.</p>	<p>Test: Connect to AIOps/observability.</p> <p>Pass: Event-driven, API-first, low-latency updates.</p> <p>Red flag: Weak ITSM tie-in, monitoring-only.</p>	<p>Test: Bulk export/import via REST.</p> <p>Pass: Works with ITSM APIs, repeatable.</p> <p>Red flag: Manual flat-file transfer.</p>

<p>Portability & exit</p>	<p>Test: Export all maps from CMDB. Pass: Full export, open formats, reimport possible. Red flag: Export blocked, paywalled, or lossy.</p>	<p>Test: API export of maps to JSON/YAML. Pass: High-fidelity export/import to other tools. Red flag: Proprietary formats, incomplete exports.</p>	<p>Test: Export/import with open schema. Pass: Preserves service relationships and history. Red flag: Only device/asset list exported.</p>
<p>Security & compliance</p>	<p>Test: Vault integration, RBAC, audit logging. Pass: Fine-grained access, auditable, region controls. Red flag: Shared admin accounts, uncontrolled residency.</p>	<p>Test: Cloud discovery scoped with least privilege. Pass: Scoped creds and audit logs visible. Red flag: Over-permissioned discovery, no logs.</p>	<p>Test: Review secret storage method. Pass: Read-only, audited. Red flag: Hardcoded or broad credentials.</p>
<p>Cost & TCO</p>	<p>Test: Build 3-year TCO model with services and support. Pass: Transparent licensing, roadmap aligned. Red flag: Hidden add-ons, pivot risk.</p>	<p>Test: Compare license vs. workloads monitored. Pass: Costs predictable and scalable. Red flag: Volatile SaaS costs or opaque tiers.</p>	<p>Test: Pilot 2-3 projects with pricing model. Pass: Low entry, portable, no hidden fees. Red flag: Weak roadmap or vendor instability.</p>

Source: Gartner