

The Middleware Adoption Journey

The integration and management of enterprise middleware and observability platforms



Welcome aboard



Navdeep Sidhu

CEO, meshIQ

A technology Tower of Babel

The members of the IT team are the unsung heroes of every large organization. Even at technology companies, their colleagues have only a cursory understanding of what they do—something I've discovered and rediscovered during my thirty years in software engineering and IT management.

Everyone knows when the sales team lands a big account. After months of Scrum sprints result in a new product release, accolades pour in. A brand or website relaunch is lauded by colleagues and broadcast to customers over multiple marketing channels.

But the better IT professionals are at their jobs, the less likely the rest of the organization is to take note.

Much of a company's IT function is hidden away—and none more so than middleware, the digital glue that bonds together the mosaic of foundational systems at any large enterprise. Unlike the platforms and applications that it connects, most employees never interact with middleware directly. And like the IT team that oversees it, the better it functions, the less likely anyone is to notice it. But keeping middleware functioning at peak performance means keeping data flowing through it unerringly and uninterrupted. And that's no small feat.

A kind of universal translator, middleware was invented so applications and platforms programmed in different languages for myriad functions can communicate by exchanging data. Inevitably, the types of middleware became nearly as disparate as the systems they connect, each requiring specific knowledge to deploy and oversee effectively.

The road more or less traveled

This report maps out a journey: the route that many an IT team have taken to connect systems that were never designed to connect. As with any map, it's possible to trace alternate paths. While your organization may not have followed the trail we've chosen to highlight, chances are you'll recognize yourself and the challenges you face at one or more of the six way stations we've identified on the middleware adoption journey. Hopefully, the solutions we present for navigating detours and surmounting obstacles on such a voyage will prove useful.

Whether you're a bank connecting a decades-old mainframe with cloud-based, AI-driven fraud detection software; an insurance company that needs its policy administration, claims processing, underwriting, and billing systems to interact; a retailer who wants to connect inventory management to logistics and point-of-sales to marketing automation; or one of dozens of other enterprises whose middleware garden has grown organically, we believe this journey will be fruitful.



Keeping middleware functioning at peak performance means keeping data flowing through it unerringly and uninterrupted.

Key takeaways



The initial decisions on middleware integrations at large companies are often made by business units rather than IT, leading to organic middleware proliferation, software sprawl, and costly redundancies.



Implementing automated monitoring, remediation, and cloud-based compute scaling through an API abstraction layer further improves uptime and developers' productivity, expedites migrations and new integrations, and lowers cloud-related overhead.



The resulting inefficiencies include protracted rollouts, over-reliance on specialized staff, limited observability that impedes issue resolution, approval delays for developer access, and an overly complex technology architecture.



A unified observability platform equips a centralized support team to manage middleware performance more efficiently and offer developers self-service access to resources; its single API abstraction layer enables automations that improve mean time between failures/to recovery.



Centralizing middleware management reduces expense and technical debt while improving middleware performance, but supporting multiple middlewares may make it difficult to meet SLA requirements and does not fully address other pain points of organic growth.



AIOps technologies, adaptive middleware, and Model Context Protocol (MCP) servers connected by universal APIs are changing the nature of how middleware operates and can prepare enterprises to support more automated, AI-driven business processes.



Governing access privileges centrally while offering developers a self-service model for accessing resources improves their user experience, fosters creativity, encourages middleware adoption, and bolsters productivity while diminishing compliance issues and access-related ticketing debt.



01

Ad hoc middleware adoption



The computer was born to solve problems that didn't exist before.

Bill Gates

Middleware everywhere

Ideally, an organization's technology infrastructure would be built according to a rational plan. In reality, technology growth is often spontaneous or improvised. In pursuit of key performance indicators, business units implement new platforms without the input of a central IT team. To control costs, legacy technologies remain in place long after they should be upgraded or because no alternatives exist. Platforms with overlapping or duplicate functions reside in multiple departments, while needed functionality remains lacking elsewhere.

The results? Application islands. Data silos. Information deserts. Stand-alone IoT platforms. And clumsy, manual data transfers among systems and across departments that would benefit from one another's data but have no efficient means for exchanging it.

Often, the introduction of middleware is similarly decentralized. Middleware is chosen, deployed, and overseen by development teams or IT staff assigned to individual business units. The result is a heterogeneous landscape of middleware—sometimes with duplicative or overlapping functionality—selected to fulfill specific business or functional priorities such as speed, accuracy, and stability. And each type of middleware may require specialists to oversee it.

Middleware monitoring is likely dispersed as well; if a centralized monitoring team does exist, it, too, is composed of specialists trained to diagnose and remedy issues for a specific flavor of middleware.



The road to good intentions is paved with middleware

By avoiding the encumbrances of a multilayered approval process, an improvised approach to middleware adoption and monitoring accomplishes the short-term goal of connecting disparate systems quickly. But the lack of an overarching plan leads to new stumbling blocks related to redundant tooling, software sprawl, inefficient rollouts, and a fragile, overly complex technology architecture that's difficult to secure and maintain:

- ⚠ **Deploying the same middleware independently in different contexts** increases software costs and duplicates integration and management resources.
- ⚠ **Multiple observability tools that provide limited visibility** into system-wide messaging bottlenecks result in resource-intensive configuration modifications, delays in issue resolution, and costly outages.
- ⚠ **Employing multiple middleware specialists increases overhead,** and those who work with mainframes are increasingly difficult to find, making it challenging to scale operations; when such people move on, institutional memory and key expertise may depart with them.
- ⚠ **The absence of clear lines of accountability** leads to finger-pointing rather than problem-solving when issues arise.

Perhaps most importantly, decentralized middleware deployment **disperses the controls for managing access privileges**. This hinders developers' access to staging platforms and crucial resources like brokers, channels, and objects, and creates a disjointed user experience. The knock-on effects touch every aspect of application development and middleware management:

Stifled creativity and delayed product releases

Excessive ticket generation, increased technical debt, and a protracted access approval process

Increased data privacy risks that result from incorrectly configured access privileges

Inordinate time spent on routine, low-value tasks such as provisioning, basic remediation, and testing that keeps IT and DevOps teams from strategic initiatives

An improvised approach to middleware adoption and monitoring accomplishes the short-term goal of connecting disparate systems quickly.

02

Centralized middleware management



The most dangerous phrase in the English language is: 'We've always done it this way.'

Grace Hopper

Inventor of the code compiler and COBOL

The ABCs of SLAs

To address some of the challenges that arise from ad hoc middleware adoption as it proliferates, companies may adopt a shared services or “center of excellence” model for managing specific integrations and begin the process of standardizing the tools for middleware oversight.

By centralizing middleware support services like provisioning, organizations begin to leverage the economy of scale:

- ✓ **Eliminating duplicative resources** lowers the total cost of ownership for the middleware.
- ✓ **More efficient workflows** reduce ticketing debt and provide faster access to resources for developers.
- ✓ **Improved overall monitoring and observability** diminish missing transactions, bottlenecks, and outages.
- ✓ **Teams assembled purposefully and thoughtfully** can manage more rollouts and oversee more implementations with fewer people, reducing—or sometimes even halving—management costs.

As an internal service-provision group, the new center of excellence must, like all service-provision groups, meet service-level agreement (SLA) targets for its internal customers. These SLAs typically relate to uptime, transaction completion rates, cost savings, and response times. And therein lies a new set of challenges.

By centralizing middleware support services, organizations begin to leverage the economy of scale.



Less than excellent

Faced with clients across the entire organization, centralized teams can be overwhelmed by access requests from DevOps and developer teams. Delays in approvals can cause delays in new platform deployments, infrastructure improvements, and product releases. And managing an outsize workload can increase security and regulatory risks if access controls are misapplied.

In addition, the sprawling, ever-expanding, enterprise-wide infrastructure for which the center of excellence is now responsible has many more points of failure. Despite initiatives to standardize middleware across the enterprise, the logistical barriers to doing so can't always be overcome, and it's not uncommon for multiple middleware technologies to persist. This complexity can make it difficult for the new, streamlined centralized team to meet its SLAs.

Nor does a centralized team solve all the pain points left over from an ad hoc middleware adoption strategy:



The management of multiple types of middleware still requires specialists—or falls to people who may lack the expertise to do so efficiently.



Configuration modifications and issue resolution therefore remain time- and resource-intensive.



Access approval delays continue to diminish user experience and productivity for developers.



Tool proliferation persists, as each type of middleware requires its own toolset for management and monitoring.

03

Centralized observability



It is a capital mistake to theorize before one has data.

Sherlock Holmes

“The Adventure of the Cardboard Box”

Unification, not proliferation

To combat tool proliferation, simplify middleware management, and improve SLA compliance, centers of excellence often adopt a unified observability platform specifically for middleware monitoring but linked to the rest of the IT landscape. Such a platform can provide real-time observability, remediation, end-to-end message tracking and tracing, and predictive analysis that averts performance problems.

Adopting a unified observability platform ameliorates many of the challenges faced by a center of excellence managing a legacy ad hoc middleware adoption strategy.

Such a platform improves SLA compliance by:

- ✓ **Flattening the learning curve for new technology**, which accelerates onboarding and time to value
- ✓ **Replacing middleware specialists with IT generalists**, thereby reducing overhead
- ✓ **Arresting tool proliferation and improving efficiency**
- ✓ **Reducing ticketing and technical debt** related to access and support requests
- ✓ **Providing developers with a better user experience**, which improves productivity

To combat tool proliferation, simplify middleware management, and improve SLA compliance, centers of excellence often adopt a unified observability platform.

- ✓ **Improving issue response time**—in some cases by as much as 90%
- ✓ **Decreasing mean time to repair (MTTR) and increasing mean time between failures (MTBF)**

The combination of improved response time, decreased MTTR, and increased MTBF is especially powerful: It's not uncommon for large enterprises to see as much as a 75% reduction in outages. However, as long as the DevOps team is handling access requests, developers remain vulnerable to delays related to approvals or request volume.



04

Democratizing access for developers

“

If we want users to like our software, we should design it to behave like a likeable person.

Alan Cooper

Inventor of Visual Basic

Turning an obstacle course into a superhighway

As any IT director knows, backlog can be a bigger impediment to progress than problem-solving. With a DevOps team snowed under by a blizzard of tickets, a developer might wait months for the credentials needed to address a problem that takes just minutes to fix.

Pressed for time and pressured for results, developers often find workarounds that circumvent the middleware altogether. These jury-rigged solutions exist outside the purview and governance of security and IT teams and can therefore introduce serious risks—including failed security audits, data loss, and extended downtime—whose magnitude increases with every new shortcut.

Adopting a self-service model for developer access to crucial resources mitigates those consequences. Self-service technology provides one normalized user interface for developers similar to the way a unified observability platform enables the management of multiple middlewares. After IT establishes granular, role-based permissions, developers can access the resources they need by learning a single tool, obviating specialized middleware knowledge.

With self-service technology, developers can access the resources they need by learning a single tool.



By removing obstacles to access and minimizing interruptions to their work, self-service developer access:

- ✓ **Improves middleware adoption rates**, thereby minimizing risk
- ✓ **Elevates user experience**, which fosters creativity and boosts productivity
- ✓ **Eliminates the need for access request tickets**, lowering overall support costs and minimizing charge-backs

A developer self-service platform also has benefits beyond the applications team:

- ✓ **Centralizing access management** automates and accelerates permissions-related compliance reporting, eliminating the need for oversight by specialized security experts.
- ✓ **Unifying control over permissions** diminishes errors, enhances overall security, and decreases the risk of financial penalties related to unauthorized access.

Middleware—and by extension, the need for rigorous middleware management—has now become foundational to the company's technology infrastructure and a key part of any future digital transformation projects, from automation to databases to artificial intelligence. Uptime has therefore become increasingly critical—but a wholly manual middleware management strategy puts uptime at risk.

And while the management and governance responsibilities of the center of excellence have been simplified, the performance of repetitive, low-value tasks like provisioning has been shifted to the applications teams.

It's at this point that enterprises look for ways to automate middleware management tasks.

Self-service developer access improves middleware adoption, elevates user experience, fosters creativity, and boosts productivity.

05

Intelligent automation

“

Software is the language of automation.

Jensen Huang

CEO, Nvidia

Making abstraction a reality

While self-service developer access promotes productivity, automated provisioning and troubleshooting can accelerate ticket resolution by as much as 90%, leaving just 10% of all tickets to be handled manually. But manual operations do more than impede recovery efforts and erode the time developers and IT have for transformative work. Fortunately, automation can do more than plow through a ticket blizzard.

A lack of automation makes it more difficult to assess things like usage requirements in hybrid on-prem/private cloud environments where capacity does not scale dynamically. This results in excessive compute consumption costs. Automated monitoring, self-executing remediation based on observability triggers, and predictive usage analytics improve uptime while optimizing resource usage and minimizing cost.

This kind of infrastructure-related automation requires a single, consistent API abstraction layer across all middleware. Much like an observability platform's self-service layer, API abstraction software provides a discrete, standardized interface that developers can use with multiple APIs. This makes it possible for developers to work with any API, and even to create new ones, without altering the code of the application itself.

With an API abstraction layer, developers can work with any API without altering the underlying application code.



Future-proofing middleware

Introducing this API abstraction layer empowers developers to create myriad efficiencies for themselves, administrators, DevOps teams, and centers of excellence:




Freeing developers from repetitive, low-value tasks like creating provisioning assets enables them to focus on application development or revenue-generating projects.



Automated, trigger-based alerts and issue resolution achieved with reusable scripts improve uptime and decrease the need for hands-on remediation while reducing service issues, MTTR, MTBF, high-level support requests, and technical debt.



Real-time data-integrity checks minimize costly data errors.



Automation fills the gaps in a workforce where middleware experts are retiring faster than their positions can be filled by new staff with the same expertise.

Best of all, automation developed using a single API abstraction layer can be applied not only to existing middleware but also to future middleware, expediting new integrations and migrations from old to new middleware. And the benefits go beyond automation: by isolating middleware and applications, an API abstraction layer enables middleware to be switched out without modifying the connecting applications or platforms.

Automation developed using an API abstraction layer can expedite future integrations and migrations.

06

The future of middleware adoption

“

*You can't connect the dots looking forward;
you can only connect them looking backwards.
So you have to trust that the dots will
somehow connect in your future.*

Steve Jobs

Are we there yet?

Your enterprise middleware architecture supports your existing business models beautifully, and data flows efficiently in streamlined schema from platform to platform. Old technologies are deprecated as new, more agile ones emerge, and integrating new middleware technologies is not the heavy lift it used to be. An API abstraction layer lets your team “develop once and deploy many,” reusing the same code for connecting any application to any API.

A unified management plane provides a single interface for observability, management, and tracking for messaging, events, and streaming. Developer provisioning and access are automated. Alerts to messaging issues are triggered automatically and errors resolved quickly—often without human intervention. Cloud compute usage scales automatically as well, reducing overhead. Ticketing and technical debt wither away. Failovers are failsafe. SLAs are exceeded.

Middleware nirvana? Perhaps. But new AIOps technologies and adaptive middleware that modifies its behavior in real time to handle changes in environment or application function are a reality. Combined with the development of a universal set of backward-compatible APIs, they are narrowing the distance between where we are and where we want to go.

And in some ways, we've already arrived....

New AIOps technologies and adaptive middleware that modifies its behavior in real time to handle changes in environment or application function are a reality.



AI and the road ahead

Middleware has long played a role in the field of artificial intelligence, ensuring that the interactions between the forerunners to large language models (LLMs) and the datasets that go into developing and training them are secure, compliant, and efficient.

In addition to performing typical middleware tasks such as load balancing, dataflow rate regulation, data encryption, message caching, and resource orchestration, AI-related middleware integrates with systems that play a role in LLM-specific tasks. These include de-risking, drift detection, performance and bias monitoring, and adaptive governance—the real-time evolution of rules, content filters, and logic that militate how models learn and what data they're permitted to learn from.

The metamorphosis of artificial intelligence from sci-fi fantasy to amusing toy to a manifold tool with real-world business applications is transforming the way we work. It's also transforming middleware itself, from how it's coded to the use cases it can now solve for—and even how it's defined.

Model Context Protocol (MCP) is an open-source standard or framework developed to help LLMs connect to external platforms, data sources, development environments, and applications. MCP servers, while not technically middleware, perform a similar function, acting as a new kind of universal translator tuned to the proprietary data formats of various LLMs, data sources, and software tools.

But whether it's the more familiar middleware that bridges the rivers separating database from AI model or MCP servers that cause those rivers to evaporate, observability is and will remain essential to the human oversight required for the responsible development of trustworthy generative and agentic AI, and to govern how it's applied IRL.

The metamorphosis of AI from sci-fi fantasy to a tool with real-world business applications in transforming middleware itself.

07

The future of middleware observability

“

A journey of a thousand miles begins with a single step.

Lao Tzu

The shortest distance between two points?

The middleware adoption journey isn't a race to the finish line. For most IT teams, it's not even a linear run: Because most rollouts are handled sequentially, companies remain vulnerable to the inefficiencies and pain points encountered during the journey we've just taken with each new middleware integration. But even if the distance between Middleware Past and Middleware Future were a straight line, the end point moves with each new generation of middleware, integration, and application development platform.

Whether you're at one of the stages we've described, somewhere in between, or have different projects at different locations, understanding how you've arrived there and the potential benefits ahead can help you address the performance, staffing, productivity, and budgetary challenges every IT department faces. And wherever you are on your middleware journey, a unified observability and management platform is key to turning those challenges into successes.

Wherever you are on your middleware journey, a unified observability and management platform is key to your success.



meshIQ: The middleware management platform for your entire infrastructure



The meshIQ platform is a scalable management plane that sits atop any type of message-oriented and streaming middleware. It provides a centralized interface and easy-to-use tools for observing, managing, tracking, and remediating all middleware—on-prem and cloud-based—throughout your enterprise.



Self-service technology and granular, role-based access privileges democratize secure provisioning and access for developers, streamlining workflows, improving productivity, and reducing security risks, IT ticketing, and technical debt.



meshIQ's transaction and performance reports and visualizations, together with auto-generated root-cause and trend analyses, facilitate triage for fixes that improve mean time between failures, mean time to recovery, uptime, and SLA compliance.



The platform's single, consistent API abstraction layer automates repetitive tasks, simplifies code-driven deployments, and allows developers to work with any current or future API.



meshIQ's automated preemptive alerts to messaging anomalies and potential bottlenecks permit you to take preventative measures rather than fight fires, minimizing disruptions and improving uptime.



Object- and user-level audit trails consolidate compliance reporting. Straightforward database-enabled rollback capabilities for configuration changes simplify the correction of coding errors. And meshIQ's customizable dashboards let you monitor message flows, queues, and transactions in real time, providing observability across all middleware platforms in a single interface.

meshIQ provides real-time observability, management, and tracking for more than 100 enterprise customers that use middleware by IBM, Apache, TIBCO, RabbitMQ, Solace, RedHat, Confluent, Oracle, and Microsoft, as well as every flavor of Kafka.

Whatever middleware may be in your company's technology constellation, reach out today to set up a demo and learn how meshIQ can streamline your middleware operations.

sales@meshiq.com | meshiq.com





The **meshIQ** platform is a scalable management and observability plane that sits atop any type of on-prem or cloud-based middleware. Its centralized interface enables end-to-end message, streaming, and events tracking in real time.

meshIQ provides easy-to-use tools and automated features for remediating traffic and performance issues, and predictive analytics for avoiding future issues. It offers developers privilege-based self-service provisioning and a unified API abstraction layer to simplify configuration and deployment.

The meshIQ platform helps large enterprises improve uptime, productivity, and SLA compliance, lower management costs, decrease meantime to repair, and increase mean time between failures.

Written by Andrew Miller
Designed by Matteo Becucci
Produced by Jennifer Knutel

sales@meshiq.com

