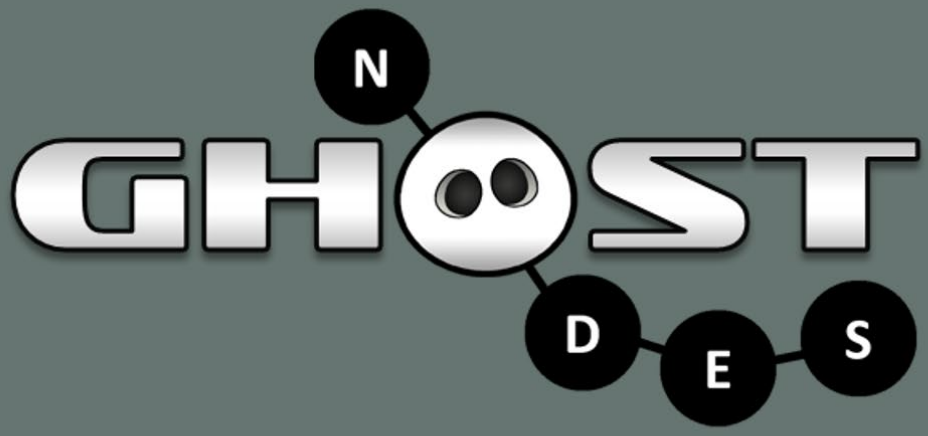
SEP
2026



GHOST LABS

WHITE PAPER

Distributed Infrastructure for Integration, Automation and AI

www.ghostnodes.com

1 The integration problem has changed

Enterprise environments are no longer centralized. Integration often still is.

Applications and workloads now live across SaaS, public cloud, private cloud, on-premises infrastructure, data centers, operational environments and the edge. AI adds yet another dimension, with data, models and execution subject to different requirements for security, latency, sovereignty and control.

The integration challenge has changed with them.

It is no longer simply about connecting system A to system B. It is about creating reliable flows across environments that evolve independently — without redesigning the surrounding architecture every time a system moves, a workload changes or a new environment is introduced.

Yet much of enterprise integration still relies on a centralized execution model.

That creates an increasingly awkward mismatch:

Ghost Nodes takes a different approach.
Ghost Nodes is distributed infrastructure for integration, automation and AI. It keeps management centralized, while execution takes place across the environments where systems, data and workloads actually live.

The architecture should follow the environment — not force the environment to follow the platform.

THE PROBLEM

Distributed systems forced through one **central runtime**

2 Why centralized execution becomes a bottleneck

Centralized integration platforms solved an important problem. They brought governance, visibility and structure to increasingly complex system landscapes.

Those things still matters.

What has changed is the assumption that governance and execution need to happen in the same place.

When every integration depends on one shared runtime, data may have to travel through infrastructure far from the systems exchanging it.

Capacity has to be scaled centrally even when the workload exists somewhere else. Maintenance or failure in one shared runtime can affect flows with little relation to each other.

And systems that could communicate locally become dependent on connectivity to a central execution point.

Moving that runtime to the cloud does not make the architecture distributed.

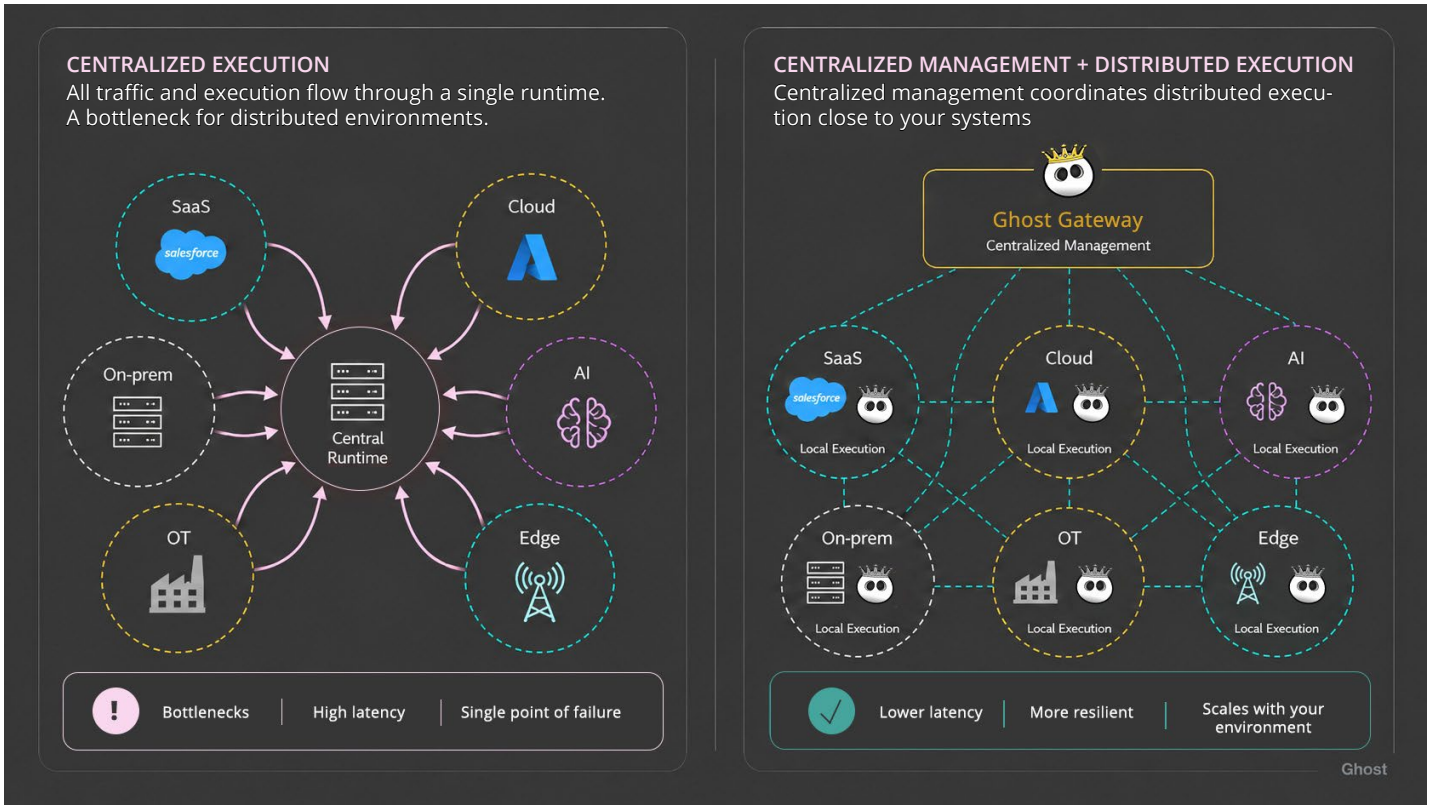
It simply changes where the center is.

Centralized execution also creates architectural coupling. As more integrations depend on the same runtime, infrastructure and operational model, changing one part of the landscape can have consequences far beyond the system being changed.

The issue is therefore not central control.

Central governance is valuable.
Central execution is not always necessary.

The two can be separated.



3 A different architecture

Centralized management. Distributed execution.

Ghost Nodes separates centralized control from distributed execution.

Integrations are designed, versioned, deployed and managed as part of one platform, while execution is distributed across Ghost Agents close to the systems, data and environments they serve.

There is no requirement for every integration to pass through one shared execution engine.

A single Ghost graph can span multiple Agents and multiple environments. One part of a workflow can execute on-premises, another in cloud infrastructure and another at the edge — while remaining part of the same managed flow.

That distinction matters.

It means organizations can keep central visibility, governance and lifecycle management without making every workload dependent on one central runtime.

Execution can happen where it makes architectural sense.

Close to the application.
Close to the data.
Inside an operational environment.
At the edge.

Or inside an isolated environment when required. Central control where it creates value. Distributed execution where it creates resilience, performance and freedom.

4 How Ghost Nodes works

One platform to design and manage. A distributed fabric to execute.

Ghost Nodes is a distributed **infrastructure platform for integration, automation and AI.** It separates design, coordination and execution into distinct responsibilities, allowing management to remain centralized while workloads execute across the environments where they belong.

Integration, automation and AI workflows are built as visual graphs in **Ghost Studio**. **Ghost Gateways** coordinate the distributed environment, while **Ghost Agents** execute the actual workload close to the systems and data involved.

Together, the distributed components form the Ghost Fabric — the distributed infrastructure that connects and coordinates execution across cloud, on-premises, edge, operational and isolated environments as one managed platform.

In plain terms:

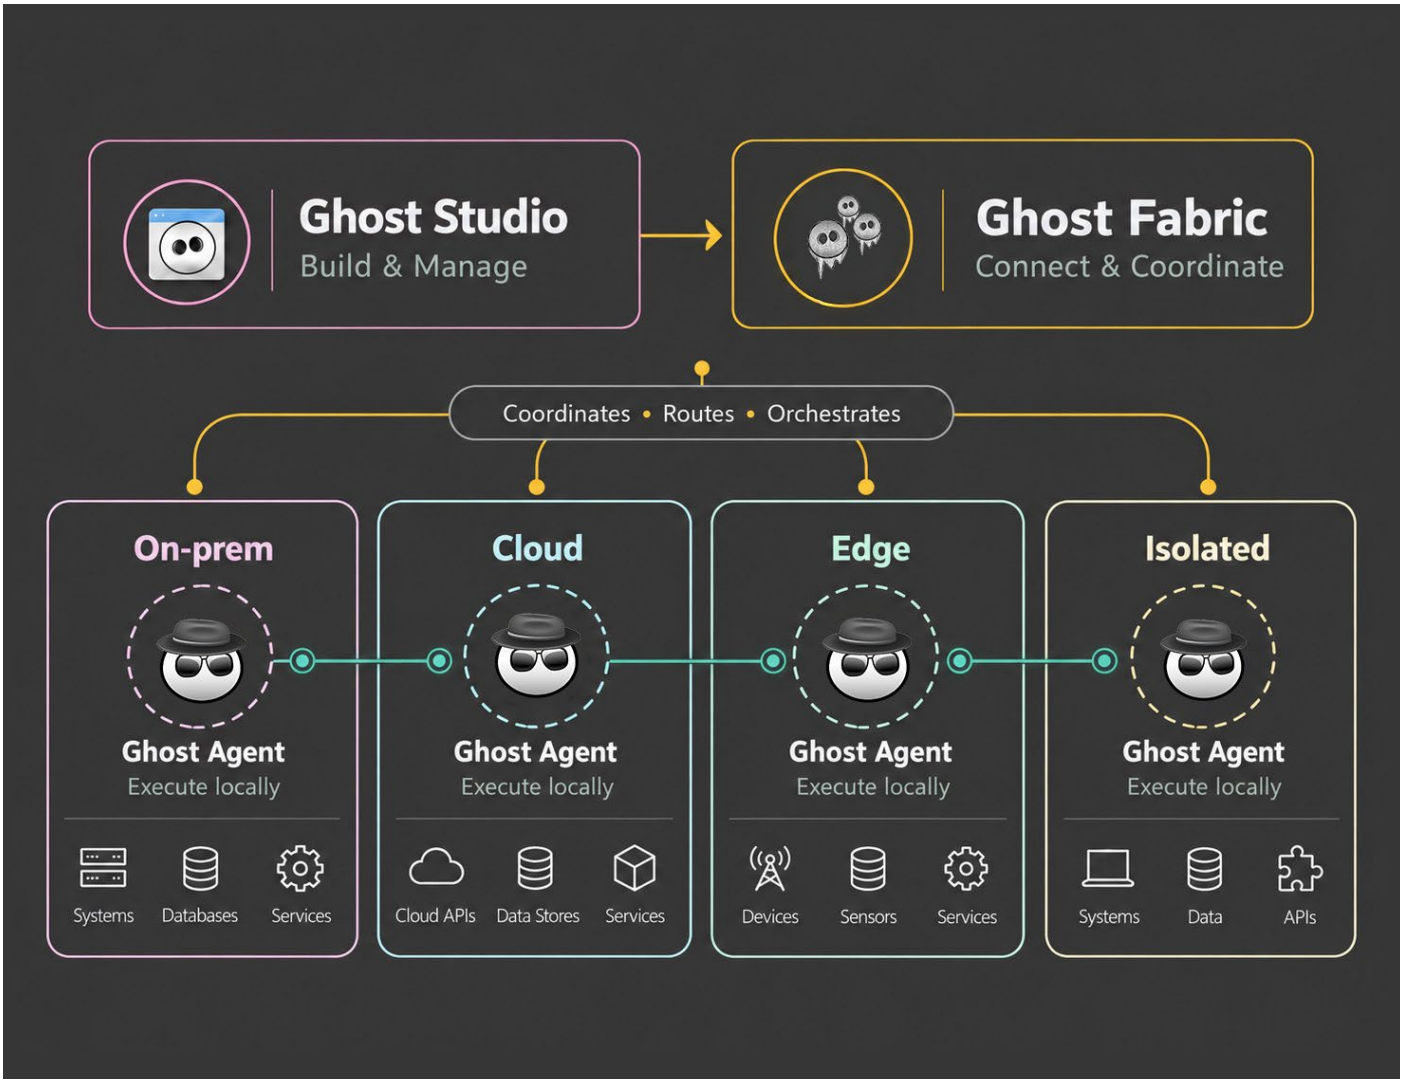
- Studio is where you build and manage
- Gateways coordinate
- Agents do the work
- Graphs define the flow
- Adapters connect systems, data and services
- The Fabric ties the distributed environment together

A graph is not tied to a single Agent, runtime or location. The same workflow can span multiple Agents and environments — processing one step on-premises, another in cloud infrastructure and a third at the edge — while remaining one managed graph.

Adapters connect graphs to applications, APIs, protocols, data sources and services. Processing and business logic can then execute where it makes the most sense, instead of forcing data through one shared central engine simply because that is where the integration platform happens to run.

The result is not a collection of disconnected runtimes. Configuration, deployment, visibility and lifecycle management remain coordinated across the platform while execution stays distributed.

One infrastructure platform.
One distributed
fabric. Execution where it belongs.



5 Run where the systems are

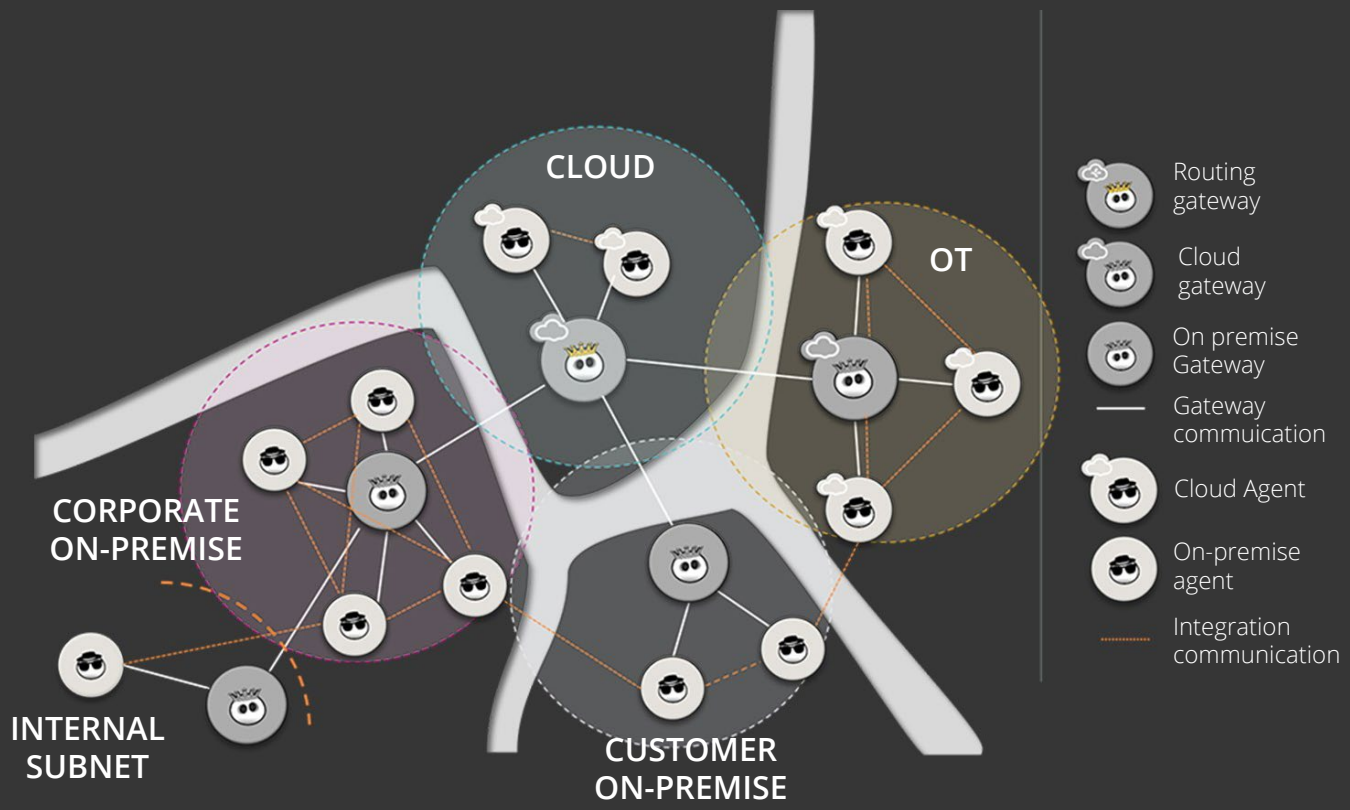
Most enterprise environments do not live in one place. They span public and private cloud, on-premises infrastructure, data centers, operational networks and edge locations — often at the same time.

Ghost Nodes is built for that reality.

Ghost Agents can run close to the systems and data they serve — inside traditional IT environments, operational technology, cloud infrastructure, edge locations or isolated networks. Different parts of the integration landscape can therefore execute in different environments while remaining part of the same managed platform.
That allows execution to stay where there is a reason for it to stay: close to a business system, inside a protected network, near machines and equipment, or alongside cloud-based applications and services.

Systems do not need to be moved, exposed or routed through a distant runtime simply to fit the integration platform.

Ghost Nodes adapts to the landscape already there — not the landscape you wish you had.



6 Sovereignty, resilience and control

A problem somewhere should not have to become a problem everywhere.

Where data and execution run — and who controls them — should be an architectural decision, not a platform limitation.

Ghost Nodes lets organizations decide where integration logic, data processing and AI workloads run: on-premises, in private or public cloud, at the edge or inside isolated environments.

For sensitive or regulated workloads, that means selected data and processing can remain within the organization's own infrastructure or chosen geographic boundaries instead of being forced through a central cloud service.

Regulatory frameworks such as NIS2 and DORA are placing greater emphasis on understanding and controlling ICT dependencies, third-party services and where critical data is processed [1].

But sovereignty is not only about geography.

It is also about **dependencies**: which infrastructure a workload relies on, which external services it needs and what happens when those services are unavailable.

Ghost Nodes is designed to give organizations control over where their integration infrastructure runs, where data is processed and which external services are part of the architecture.

The distributed execution model also improves resilience. Workloads running in one environment do not have to depend on the availability of unrelated environments. Network disruptions, maintenance or a problem elsewhere in the landscape can therefore remain more contained instead of automatically affecting every integration.

Central management still provides visibility and governance across the platform — without requiring every workload, every data flow or every environment to depend on one central runtime.

You decide what runs where, what depends on what — and what is allowed to leave the environment.

[1] See Directive (EU) 2022/2555 (NIS2), Art. 21, on cybersecurity risk-management and supply-chain security; and Regulation (EU) 2022/2554 (DORA), Arts. 28–30, on ICT third-party risk and contractual requirements, including where services are provided and data is processed.

What changes when integration becomes infrastructure

Distributed execution changes more than where integrations run. It changes how much of the landscape has to scale, fail or change together.

SMALLER FAILURE DOMAINS

A problem in one Agent, flow or environment does not automatically have to become a platform-wide problem.

Failures can stay closer to where they occur.

SCALE WHERE THE LOAD EXISTS

Capacity can be added where processing actually happens instead of scaling one shared runtime for the entire integration landscape.

GREATER OPERATIONAL INDEPENDENCE

Local workloads can operate with fewer unnecessary dependencies on distant infrastructure or connectivity.

Different environments can retain more autonomy while still being part of the same manged platform.

CHANGE IN PARTS

A system can be replaced, a workload moved or a new environment introduced without forcing the surrounding integration architecture to move with it.

Change one part without unnecessarily disturbing the rest.

LESS ARCHITECTURAL LOCK-IN

Integration logic becomes less dependent on one execution location, cloud, infrastructure model or central runtime.

That gives organizations more freedom to change the technology around the integration layer over time.
Enterprise change rarely happens everywhere at once. The integration architecture should not require it to.

Integration, automation and AI — one integration model

Integration, automation and AI increasingly depend on the same systems, data, events and business processes.

Treating them as separate technology stacks often means adding another platform, another set of dependencies and another layer of orchestration.

Ghost Nodes takes a different approach.

A Ghost graph can move data, apply business logic, trigger actions and invoke AI models as part of the same distributed workflow. AI therefore becomes another workload that can be orchestrated alongside the rest of the business process — not a separate island beside it.

Where that AI runs is also part of the architecture.

.

Cloud-based models can be used where they make sense. Local or self-hosted models can be used when data sensitivity, sovereignty, latency or operational requirements make that the better choice.

AI becomes part of the flow — not another platform the flow has to work around.

AI-ASSISTED ENGINEERING

Ghost AI Studio applies AI to the integration lifecycle itself.

It can help teams create Ghost graphs and integrations from prompts, understand and document existing graphs, analyze and migrate integrations from platforms such as BizTalk and SnapLogic, optimize existing flows and generate new Ghost adapters.

That matters because modernization rarely starts with a blank page.

Most organizations already have years — sometimes decades — of integration logic, dependencies and platform-specific knowledge embedded in existing solutions. Before that landscape can be modernized, someone has to understand what it actually does.

Ghost AI Studio is designed to help with that work as well as the building of what comes next.



AI does not only run inside the integration.
It helps engineers understand, build and evolve it.

9 Modernize without the big bang

Replacing an integration platform does not have to mean replacing everything at once.

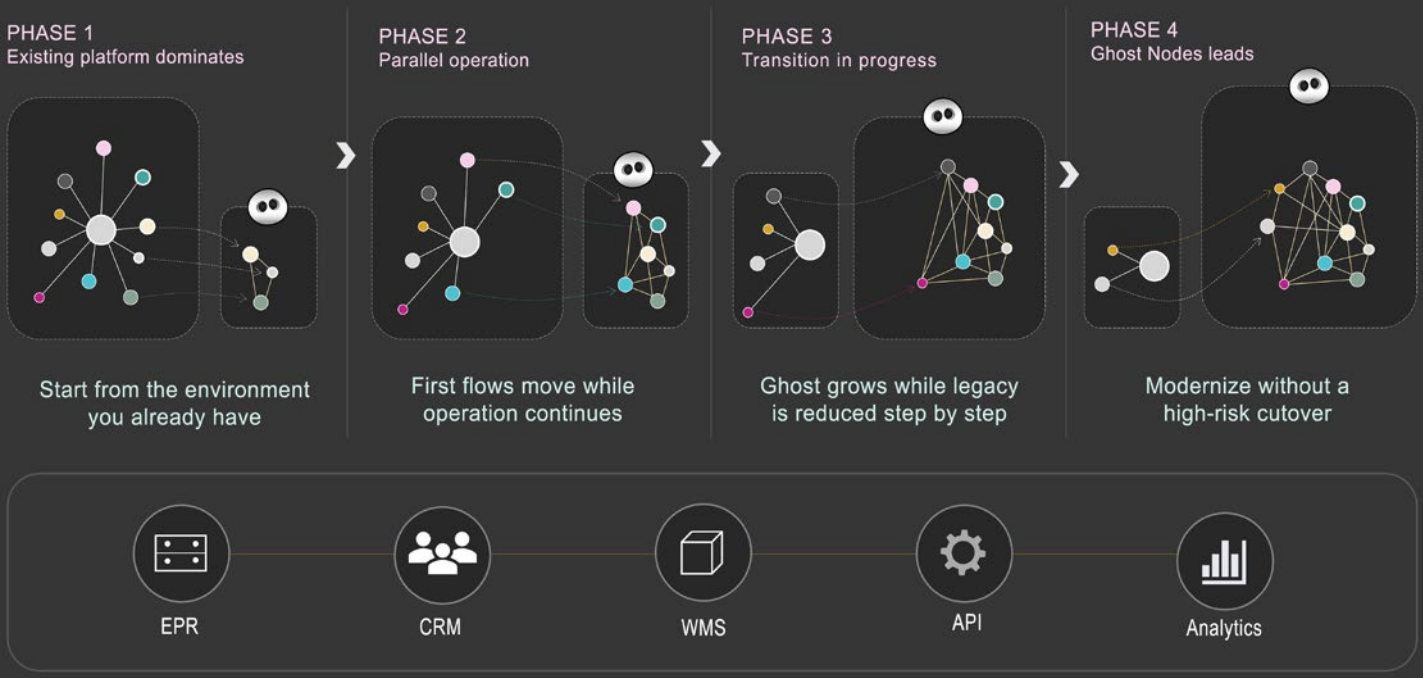
Ghost Nodes can run alongside existing integration platforms, allowing organizations to move flows gradually while production systems and business operations continue as normal. New integrations can be built in Ghost, existing ones transitioned when it makes sense, and legacy platforms reduced step by step rather than through one high-risk cutover.

That changes the nature of modernization.

Instead of treating migration as a single project with one critical deadline, organizations can modernize in controlled stges. Each flow can be moved, validated and put into production before the next one is touched.

That matters particularly in large integration landscapes, where hundreds of flows may connect business-critical systems and where the cost of getting a cutover wrong is far greater than the cost of taking a little more time.

Move flows gradually from existing integration platforms to Ghost Nodes - while operation continues as normal.



300+ BizTalk integrations transitioned while 40+ connected systems remained operational throughout the migration.

The same approach applies to BizTalk and other established integration platforms. Existing environments can remain in place while Ghost Nodes is introduced gradually alongside them.

> *Modernization becomes a transition — not an event.*

Proven in production

Ghost Nodes is not a reference architecture or a lab concept. It is running in live enterprise environments across legacy platforms, cloud services, operational systems and real-time workloads.

400+

BUSES CONNECTED
within one real-time operational environment

575+

ACTIVE INTEGRATIONS
running in one enterprise environment

500+

ON-PREMISES INTEGRATIONS
running in one customer environment

300+

BIZTALK INTEGRATIONS TRANSITIONED
for a single enterprise customer

50M+

TRANSACTIONS PER DAY
handled in one live customer environment

11

Built for environments that keep changing

Enterprise architecture is never finished.

Systems are replaced. Workloads move. Requirements change. New technologies create new possibilities — and new dependencies.

The integration layer should absorb that change, not turn every change around it into another integration project.

Ghost Nodes is built for exactly that: an integration architecture that can evolve with the environment around it.

> *Because the integration layer should not become the next thing holding you back.*

Ghost Nodes —
Distributed infrastructure for integration, automation and AI.

www.ghostnodes.com

by Ghost Labs