NGINX
Part of F5

PRODUCT DATASHEET

F5 NGINX SERVICE MESH

Enterprise-Class Connectivity for Kubernetes Services

WHY USE NGINX SERVICE MESH?

Increase Uptime
Ensure availability of backend services in scalable, dynamic environments, preventing connection timeouts and errors

Improve Security
Integrate strong centralized security controls within a Kubernetes cluster to protect distributed services at scale

Gain Insight
Achieve better visibility into service health and performance to reduce outages and simplify troubleshooting

Increase Uptime, Improve Security, and Gain Better Insight into Service Health

Kubernetes adoption is growing rapidly as organizations discover it's the most suitable way to deploy and run containerized microservices-based applications at scale.

However, organizations are facing challenges with security, reliability, observability, and scalability when they run Kubernetes in production:

- Connection timeouts and errors in scalable, dynamic environments lead to service interruptions
- Inadequate protection across distributed environments increases risk of exposure to cybersecurity threats
- Insufficient visibility into service health and performance causes outages and troubleshooting difficulties

NGINX Service Mesh is a part of NGINX Connectivity Stack for Kubernetes, designed to address service connectivity challenges in production environments – including on premises, in the cloud, and at the edge – with its enterprise-class availability, security, and visibility features:

- Ensures availability of business-critical backend services with advanced load balancing and connectivity patterns
- Improves protection with strong centralized security controls within a Kubernetes cluster
- Reduces outages and simplifies troubleshooting with granular real-time and historical metrics and dashboards

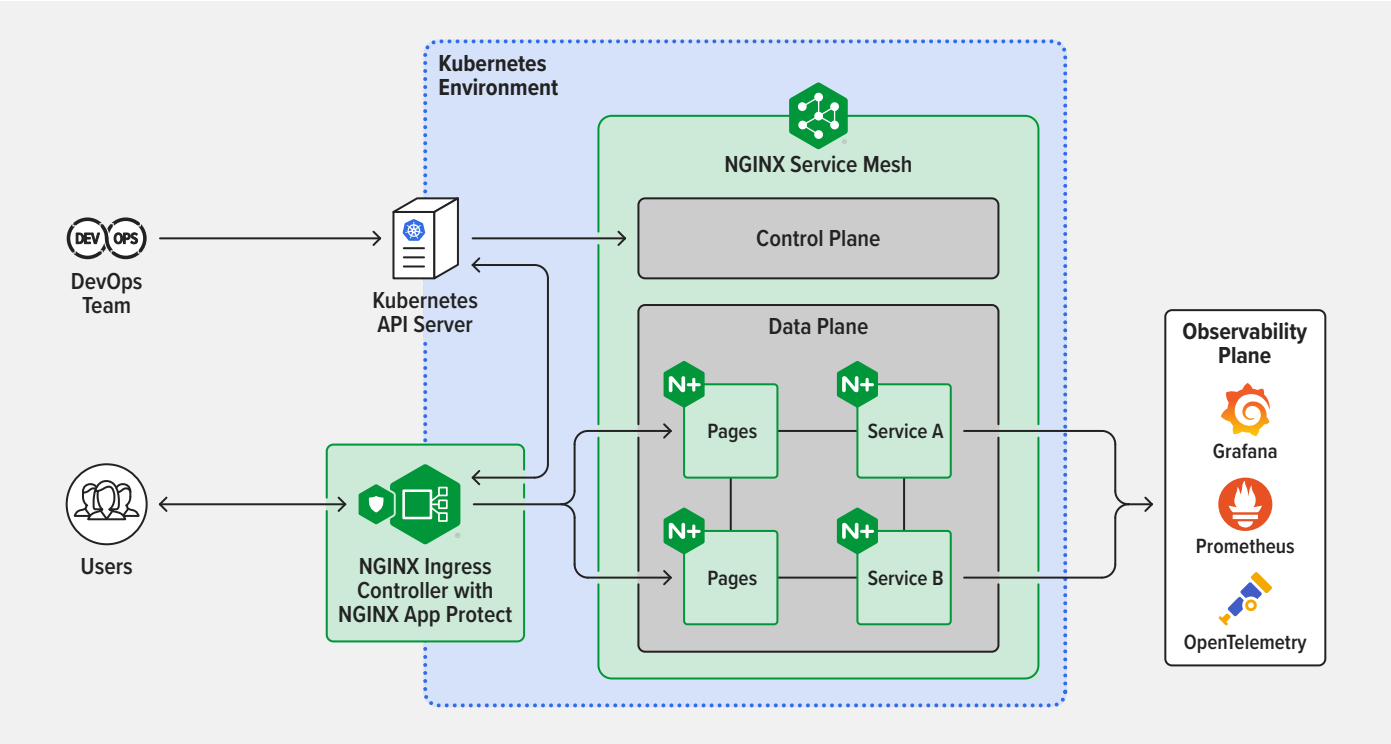

Benefits of NGINX Service Mesh

Simplify and streamline service connectivity in any Kubernetes environment. Reduce complexity with a lightweight, yet comprehensive, low-latency service connectivity fabric that's easy to deploy and use within a Kubernetes cluster.

Ensure Availability

Prevent connection timeouts and errors and avoid downtime when rolling out a new version of a backend service or during topology changes, extremely high request rates, or service failures.

- Advanced Layer 7 (HTTP, gRPC) and Layer 4 (TCP, UDP) load balancing with dynamic updates of target service instances
- Blue-green and canary deployments
- Rate limiting and circuit breaker connectivity patterns

Improve Visibility

Gain better insight into app health and performance with over 200 granular real-time and historical metrics to reduce outages and simplify troubleshooting.

- Discover problems before they impact your customers
- Find the root cause of app issues quickly
- Integrate data collection and representation with ecosystem tools, including OpenTelemetry, Grafana, Prometheus, and Jaeger

Strengthen Protection

Ensure secure service connectivity with enforced service identities, authorization, access control, and encrypted communications.

- SPIFFE and the SPIRE runtime with built-in or external certificate authority (CA) and automated certificate lifecycle management
- Access-control policies to allow communications to and from specific source and destination endpoints
- Authentication and encryption with mTLS

Simplify Operations

Reduce complexity and tool sprawl through technology consolidation for faster and easier app delivery

- Tight integration with NGINX Ingress Controller for unified app connectivity into, out of, and within the cluster
- Data and control planes are the same across all hybrid and multi-cloud environments
- Focus on core business functionality, offloading security and other non-functional requirements to the platform layer

To learn more, visit nginx.com/k8s

NGINX
Part of F5