

SnapLogic Ultra Pipelines

Enable real-time data delivery at scale

Lag times, glitches, and outages are unacceptable when it comes to your mission-critical data delivery processes. Across all industries, our customers are building applications that demand low latency, resilient services.

- Banks need digital transactions to clear quickly and securely
- Retailers need e-commerce sites to process orders in real-time
- Advertisers need to handle the web traffic spikes that coincide with ad placements
- Identity management providers need to verify people's credentials instantly

Ultra Pipelines, an extension of the SnapLogic Intelligent Integration Platform (IIP), ensure ultra-fast, high throughput, and reliable data transfer for scenarios that are essential to your business. If your business depends on resilient and low-latency processes available 24 hours a day, 7 days a week then you need Ultra Pipelines.

Enjoy 'Ultra'-High Speed

Historically, running low-latency pipelines has been a complex and code-heavy process. Ultra Pipelines rapidly deliver time-sensitive data to cloud and on-premises applications and data stores. Those requesting data get answers under a second to their requests, up to 10x faster, than without Ultra Pipelines. Our "always-on, always-active" integrations eliminates the need for connection setup when a new request comes in. Ultra Pipelines enable real-time operations and are perfect for streaming data from Kafka producers or from IoT endpoints.

Benefits of Ultra Pipelines

Low latency - Provides sub-second response time with 'always-on' connection. Ensures data is passed between applications in the fastest manner possible

High throughput - Because pipeline setup is not needed, you can run more instances of Ultra pipelines on a given node, giving you 4x as much throughput.

Greater resiliency - Even if there are failures of execution or queuing nodes built in redundancy of Ultra delivers data where it's needed

Process More Requests

With SnapLogic Ultra, you can run more instances of a pipeline on a SnapLogic node. Because the Ultra connection is always on, computing resources are not wasted on setting up and tearing down the connection for every event, freeing up processing power for executing the requests themselves.

Get Fault Tolerant Data Processing

Get guaranteed delivery and security against failures without losing scalability. Ultra Pipelines have redundancy built into the Execution nodes as well as in the FeedMaster nodes, that feed requests to Execution nodes. So even if a node fails, your real-time requests will continue to get served.

How Does It Work?

A standard configuration for Ultra is as shown in Figure 1. It consists of 2 execution nodes that execute multiple instances of a pipeline in parallel. Two feedmaster nodes distribute incoming requests to the execution nodes and are fed these requests by a load balancer. Figure 2 shows the second configuration that can be used with Ultra. Here the requests come from a Kafka streaming service or equivalent or from a message queue such as Java Messaging Service (JMS), RabbitMQ, Amazon SQS, Google Pub/Sub and are fed to one of the two execution nodes.

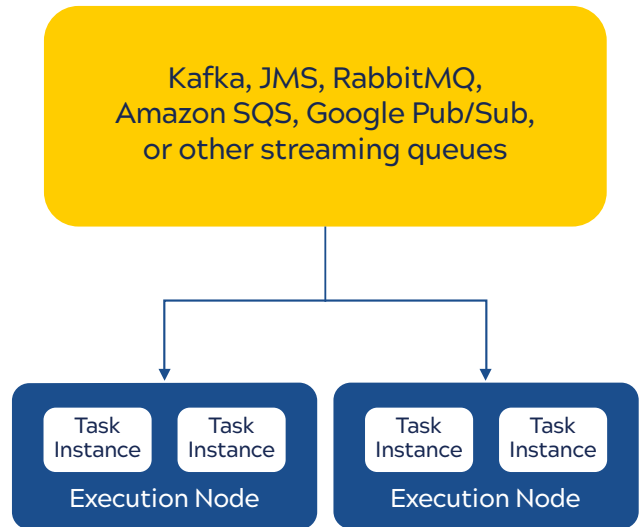


Figure 2: Ultra Configuration for Kafka and other streaming applications and message queues

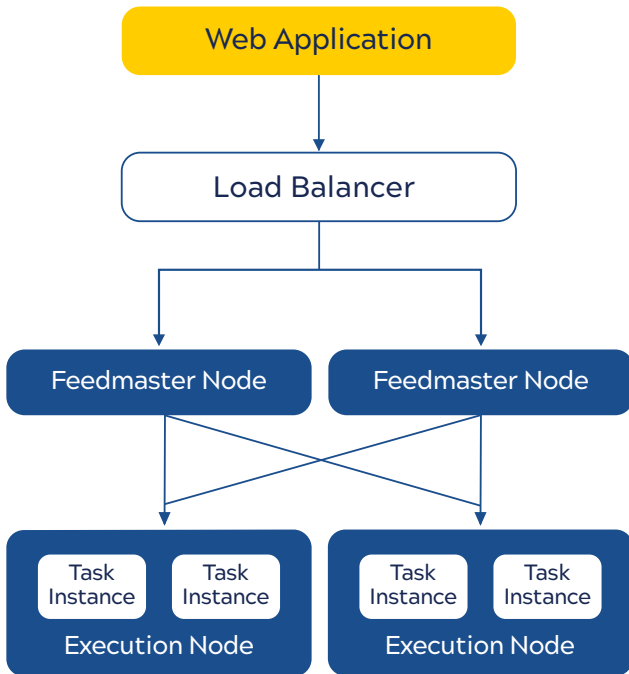


Figure 1: Standard Ultra Configuration with 2 Feedmaster nodes and 2 Execution nodes.

SnapLogic powers the automated enterprise. The company’s self-service, AI-powered integration platform helps organizations connect applications and data sources, automate common workflows and business processes, and deliver exceptional experiences for customers, partners, and employees. Thousands of enterprises around the world rely on the SnapLogic platform to integrate, automate, and transform their business. Learn more at snaplogic.com.