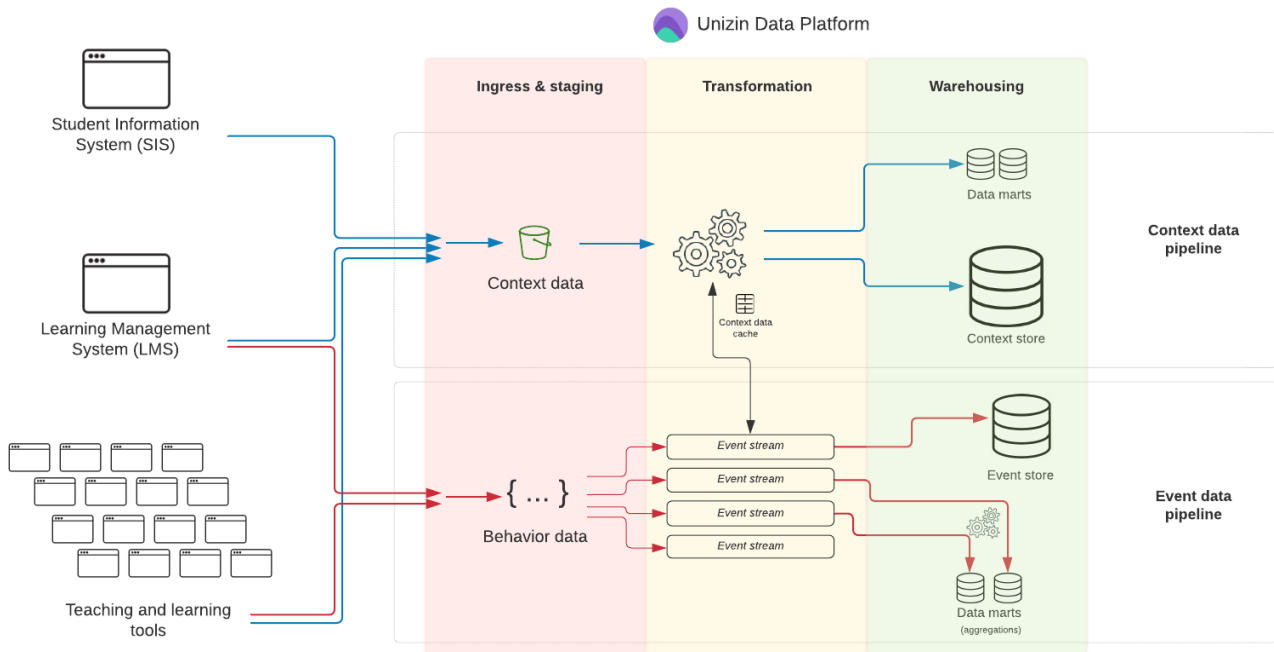


System overview



The Unizin Data Platform (UDP) is primarily composed of two data pipelines. Each data pipeline creates and maintains the data lakes and data marts that undergird the [UDP's data services](#). There exists one data pipeline for each [learning data category](#) integrated by the Unizin Data Platform (context data and behavior data). While each data pipeline is architecturally isolated from the other, there do exist points of interaction.

Each UDP data pipeline is composed of loosely-coupled applications/services and infrastructure. Broadly speaking, these loosely-coupled components map onto three separate phases of each pipeline: (1) data ingress & staging, (2) data transformation, (3) data warehousing.

Each pipeline maintains what's called a "foundational store" for each data category. The foundational stores are comprehensive/archival stores that capture all data integrated and stored into the UDP for all time. For context data, the UDP maintains the UDP Context store. For behavioral data, the UDP maintains the UDP Event store. Alongside the foundational stores are created and maintained a variety of data marts. The UDP's data marts are domain-specific data aggregations and derivations intended to serve related use-cases.

The UDP's applications/services and infrastructure are created during the UDP installation process. However, UDP customers may need to further configure certain aspects of these applications. In some cases, ongoing maintenance to, for example, add new data integrations is required.