

The Plandek data collection process is separated into several stages to ensure that your source code does not leave your network and to put minimal strain on your services.

Gathering

Plandek's gatherers run on your servers and connect to the APIs exposed by your services. They fetch raw data, which is stripped down to the minimal amount of metadata required to power the product. This metadata is then sent over a secure TLS socket to Plandek's API to be processed.

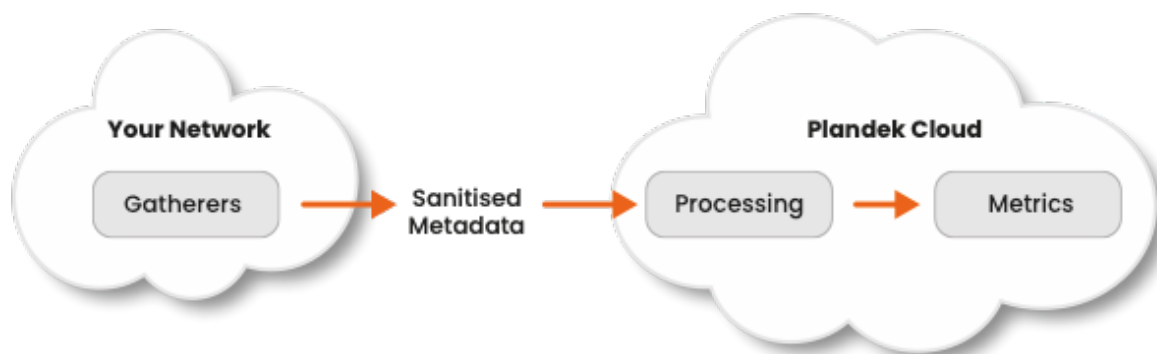
More information about the data that the gatherers capture and store can be found in the Plandek Gathering and Security document, however it is worth noting that your codebase never leaves your network.

Processing

Processing runs on Plandek's servers and converts the metadata into events which are used to generate the metrics that power the Plandek Web Application.

Metrics

Metrics runs on Plandek's servers and consumes the events from processing and generates metric data based on them.



How it is deployed and managed

The on-premise gatherers are designed to run on a kubernetes cluster and are packaged as a docker container. This is comprised of a long-running orchestrating component which communicates with Plandek's systems and launches short-lived gathering pods on the kubernetes cluster. We provide a Kubernetes skeleton and an API key to set up this component, and it should immediately begin to work. This component requires that Kubernetes Role Based Access Control be used to permit the podserver to launch other pods. There are some configuration options to control which nodes that pods can be launched on and the maximum capacity that it is able to deploy.