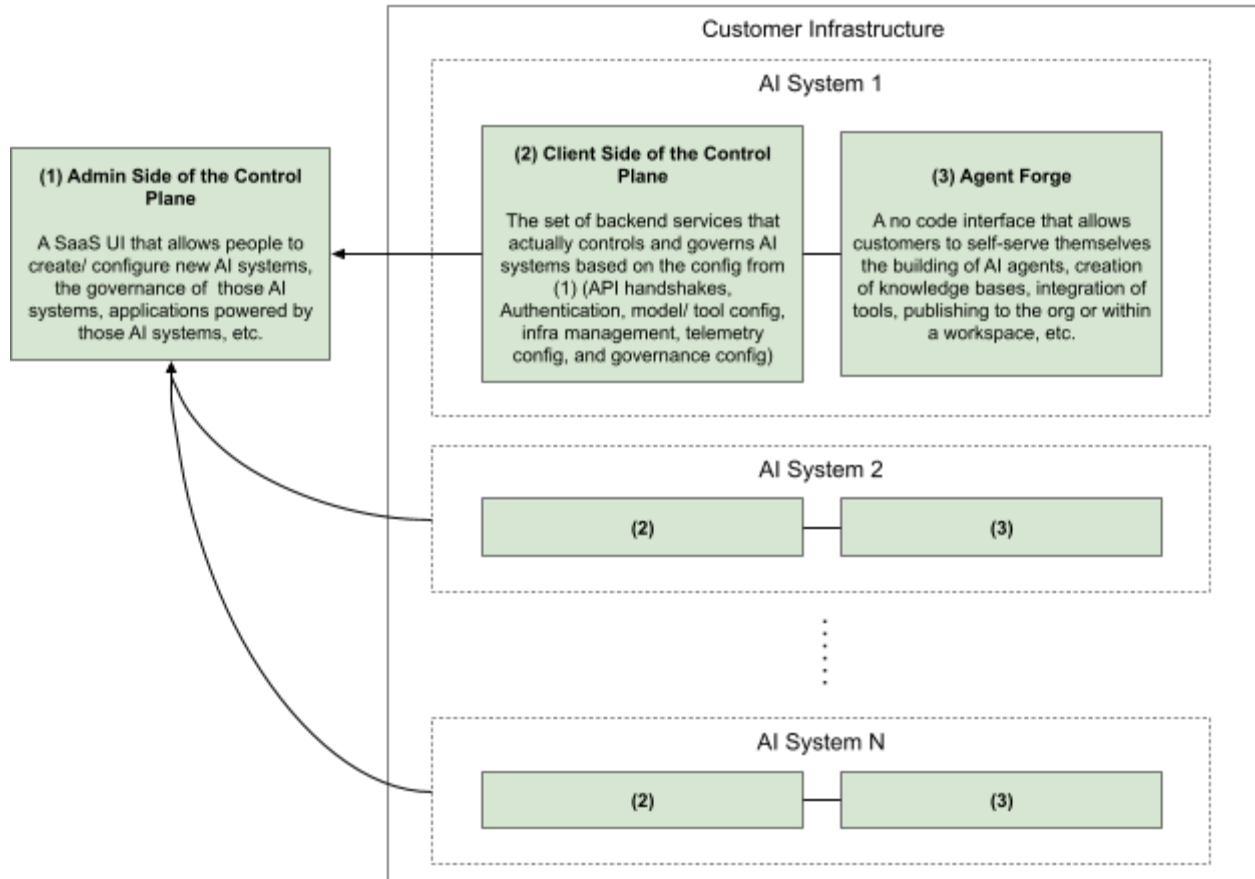


Prediction Guard Architecture



- The SaaS portion of our control plane is (1) pictured above.
- That portion of the control plane is currently hosted in GCP, but, it is containerized, and could be hosted in any k8s control plane (e.g., AKS)
- To use (1) (the SaaS side of the control plane), the customer logs in via username/ password or SSO. They create an AI system by deploying the client side of the Prediction Guard control plane in their infra (e.g., Azure). That client side of the PG control plane connects to (1) via outbound connection. The customer continues to manage, configure and govern AI systems via (1), with the client side components (2) applying those changes.
- The customers received access to the SaaS product via us provisioning them an login on the (1) UI.
- See the above diagram and description.
- In terms of what is running in and outside of Azure
 - (1) lives outside of Azure
 - (2) and (3) are deployed inside of Azure (or the relevant customer environments, such as on-prem or other cloud VPCs)