Syniti Knowledge Platform Connector and Associated Services

Security and Requirements

BackOffice Associates, LLC DBA Syniti
info@Syniti.com
www.Syniti.com

Contents

Overview

What Is the Syniti Knowledge Platform Connector?

The Syniti Knowledge Platform Connector is a lightweight, dynamic connection engine that allows the Syniti Knowledge Platform (SKP) to communicate with remote services via APIs and other TCP connections. The Knowledge Platform Connector uses software intelligence to automatically, seamlessly and securely establish and maintain connectivity between remote networks and SKP (see Figure 1).

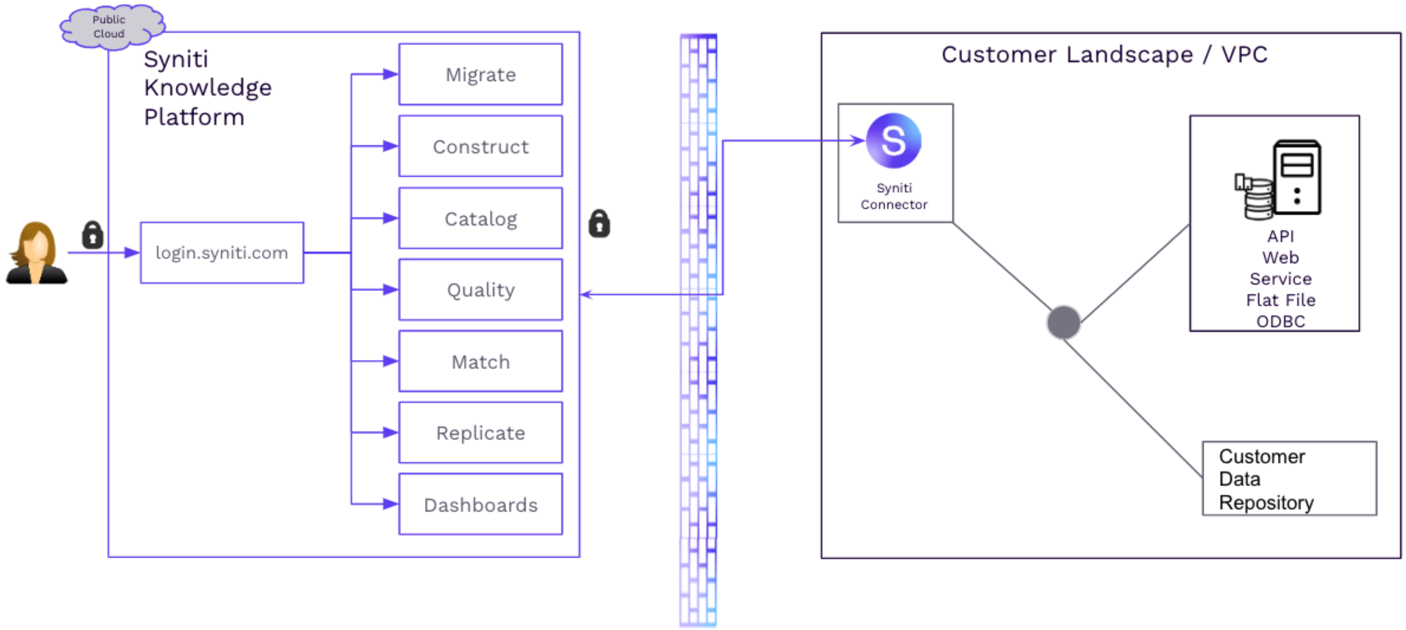

Figure 1: High Level Architecture

Why Did Syniti Create the Knowledge Platform Connector?

A core design principle is that customer master, transactional and operational data is not persisted within the Syniti Knowledge Platform itself. To provide core product capabilities such as data quality reporting, data profiling, metadata scanning, advanced deduplication, matching and other key functionality, SKP uses an Connector-based architecture to securely distribute execution outside of the SKP application environment.

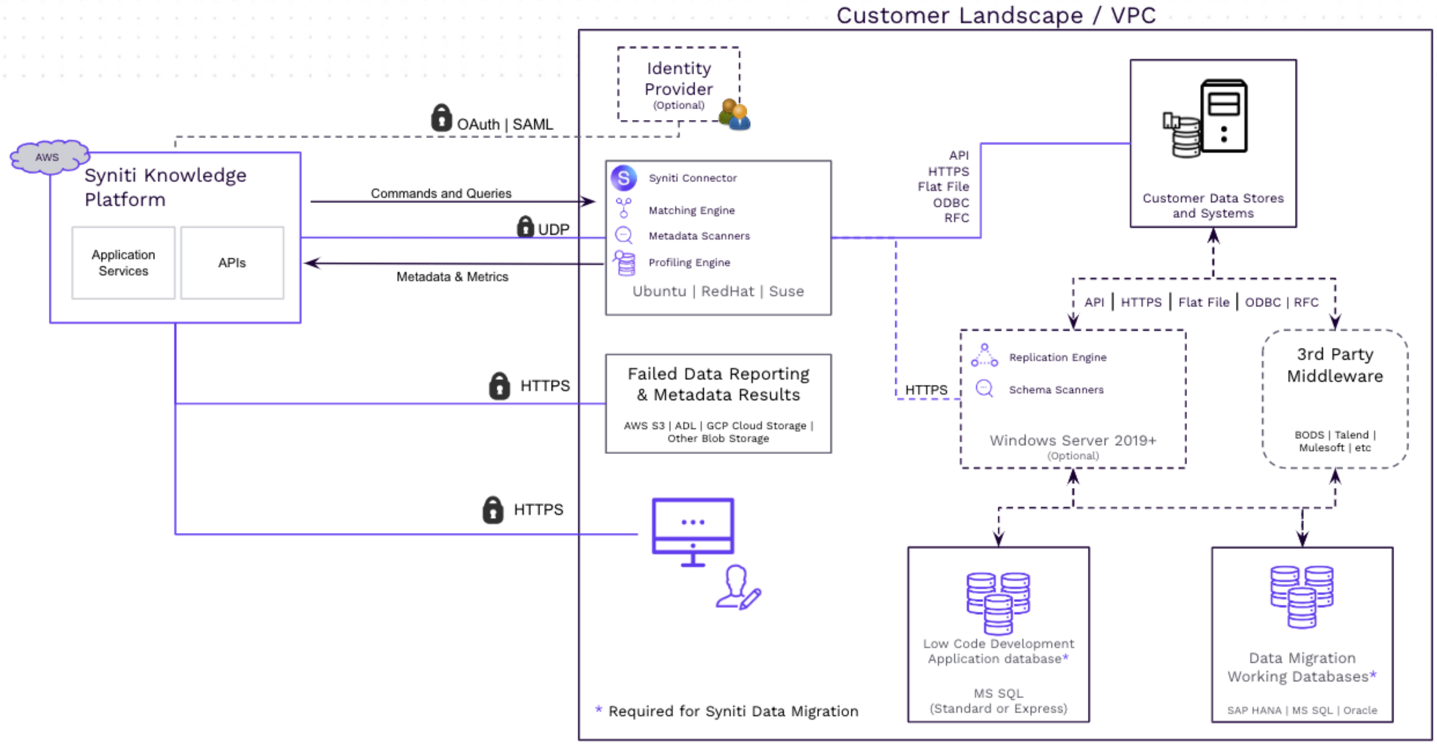

Figure 2: Detailed Protocol Diagram

As is shown in the above diagram, the SKP sends commands and queries to the Knowledge Platform Connector (referred to as the Syniti Connector in the diagram above). These commands include instructions such as RESTful API calls to local services and data quality queries. As a result of these commands, only metadata and metrics are sent to SKP for storage and processing.

The key to successful processing via the Knowledge Platform Connector is that, from the Knowledge Platform Connector instance itself, communication pathways are open to the data storage and processing systems that will be scanned for metadata, queried for data quality concerns, used to execute data migrations, and integrated with the Syniti Knowledge Platform.

Security

Security is a fundamental consideration for all software written at Syniti and the Knowledge Platform Connector is no exception. A core design principle of this system is that Syniti prioritizes confidentiality and integrity over availability when these tenets of the CIA triad come into conflict. If the system cannot operate securely, it must never revert to some less secure mode of operation to maintain availability; in such situations Syniti would expect to compromise availability to maintain confidentiality and integrity of the data and systems it can access.

Encrypted Communication via a Software Managed VPN

The Syniti Knowledge Platform Connector is based on the Wireguard VPN technology that is available in the Linux kernel. From this baseline, a modern VPN style connection using state of the art cryptography is established that is designed to be faster and more efficient, with a smaller attack surface than other encryption and VPN standards. Technical details on the encryption, standards and performance can be found here.

Data Transit and Storage

All communications between the SKP and the Knowledge Platform Connector are performed over secure channels which provide authenticated encryption. There are two channels that are used for two distinct purposes:

1. **Control Plane:** VPN configuration details are passed from SKP to the Knowledge Platform Connector over TLS. See the section “Threat Scenarios” below for more information about specific measures employed to authenticate the configuration parameters.
2. **Data Plane:** All other communication between the Knowledge Platform Connector and the SKP occurs over a Wireguard VPN tunnel, which provides authenticated encryption.

The following Knowledge Platform Connector-related data is stored in the SKP:

- An association between tenants and the public keys of their registered Knowledge Platform Connectors. Syniti does not store any VPN private keys centrally. Both the customer premises Connector and the SKP peers generate encryption keypairs on-device when started, and the private key never leaves the device on which it was generated.
- Credentials stored in SKP by users for the purpose of connecting to customer resources. Syniti immediately encrypts the credentials using FIPS-validated hardware security modules. They are never written to persistent storage unencrypted. A full audit log of each decryption is kept.

Threat Scenarios

When deploying Syniti Knowledge Platform Connector, there are several common threat scenarios that are typically considered. What follows is a summary of these common scenario questions and the Syniti approach to protect against these threat vectors.

Threat Scenario: Attacker causes Knowledge Platform Connector to connect to malicious VPN endpoint

When the Knowledge Platform Connector starts up, it fetches its VPN configuration from an SKP cloud endpoint. Any attacker who can intercept or modify the response received by the Knowledge Platform Connector could cause the Knowledge Platform Connector to connect to an attacker-controlled Wireguard peer, which would in turn give the attacker access through the Knowledge Platform Connector to additional resources on the customer's network.

To mitigate this risk:

- The Knowledge Platform Connector requests the configuration over an encrypted and authenticated TLS connection. To inject a malicious configuration, attackers would need to compromise web PKI or install a malicious root certificate on the Knowledge Platform Connector.

- For situations where customers employ TLS man-in-the-middle technology, Syniti only relies on TLS to provide encryption (privacy) in transit. Coming soon, the configuration sent to the Knowledge Platform Connector will include a cryptographic signature to attest that it was legitimately created by Syniti. The Knowledge Platform Connector includes a public key with which the signature can be validated.
- To prevent an attacker from gaining control of old VPN endpoints once used by Syniti, replaying an old configuration and causing the Knowledge Platform Connector to connect to the former-Syniti-now-attacker-controlled VPN peers, a replay resistance mechanism is employed. The Knowledge Platform Connector sends its configuration requests with a nonce which SKP signs as part of its response. The Knowledge Platform Connector verifies the nonce in the response to ensure that it has received a fresh configuration.

Threat Scenario: Attacker causes customer to install malicious Knowledge Platform Connector software

The Knowledge Platform Connector is packaged as a standard Linux package for several distributions. The distribution package manager is responsible for cryptographically verifying the package before installing and running it. Syniti signs its Knowledge Platform Connector software with keys appropriate to the distribution mechanism, and publishes these keys for customers to verify themselves if desired. This mechanism is used both for initial Knowledge Platform Connector installation as well as package upgrades.

An attacker who gains access to Syniti internal release infrastructure could cause malicious Knowledge Platform Connector software to be published and signed with Syniti's keys. This is known as a supply chain attack. To mitigate this risk:

- Syniti employs the principle of least privilege to limit the number of people who have access to the signing machinery.
- Access to systems where the signing machinery runs requires multi-factor authentication.

Threat Scenario: Attacker gains access to SKP infrastructure and data

The Syniti Knowledge Platform is a cloud-based service available over the public Internet. Despite our best efforts to develop secure software and run it on secure infrastructure, there are inherent risks involved due to software bugs or design flaws both in Syniti code as well as in our underlying dependencies. Thus it is useful to speak in terms of a "trust boundary" that separates components of the Connector system with the underlying systems upon which it depends. If components beyond the trust boundary are compromised, this could lead to downstream compromises in the Knowledge Platform Connector system.

Our infrastructure provider, Amazon Web Services. lies outside the trust boundary for this model. If an attacker is able to compromise our infrastructure provider, then customer data could be at risk.

Attackers could take the form of:

- External entities who discover vulnerabilities in our software or employ social engineering techniques to gain elevated access
- Internal Syniti employees acting maliciously ("insider threats")

An attacker who gains access to the SKP control plane could use the Knowledge Platform Connector to gain access to customer networks. To mitigate these risks:

- The Knowledge Platform Connector can be manually configured by the customer with a list of IP addresses and ports it is allowed to provide connections to. If a request is made to the Knowledge Platform Connector from the cloud to connect to a host which is not in the allowed list, then the attempt will be refused and logged by the Knowledge Platform Connector, disregarding the prohibited operation requested by the SKP.
- Syniti strongly encourages customers to configure their own network firewalls to only permit traffic originating from the Knowledge Platform Connector to expected internal and external hosts.

Detailed Description of Knowledge Platform Connector Communication

Installation Phase

During the installation phase, the Knowledge Platform Connector is installed by the customer on a locally-administered virtual or physical server.

A Wireguard key pair is generated on the Knowledge Platform Connector server when it starts for the first time, and the Knowledge Platform Connector begins polling the control plane by making a HTTPS request over TLS to onprem.syniti.com for the appropriate Wireguard configuration for its public key.

Because the public key has not yet been registered, a 404 Not Found response is sent, and the Knowledge Platform Connector simply continues polling periodically.

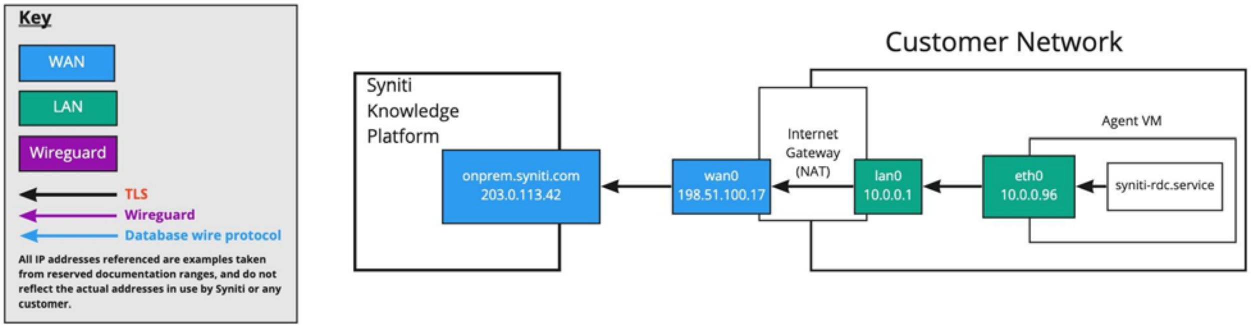

All IP addresses referenced are examples taken from reserved documentation ranges, and do not reflect the actual addresses in use by Syniti or any customer.

Registration

Customer links the Knowledge Platform Connector to their SKP tenant by registering its public key in the Admin area.

The Knowledge Platform Connector is assigned a private RFC 4193 Unique Local IPv6 Unicast Address by the control plane to use for its Wireguard interface. This ensures that the IP addresses used by Syniti for the secure tunnel will not conflict with any IP ranges in use by the customer. The IPv6 address is only used inside the tunnel; the Knowledge Platform Connector does not send IPv6 traffic over the customer's network.

The next time the Knowledge Platform Connector polls for a configuration, it will receive a Wireguard configuration. The configuration includes details on how the Knowledge Platform Connector can reach the SKP Wireguard Edge.

When the Knowledge Platform Connector receives this configuration, it will configure a Wireguard interface based on the received information.

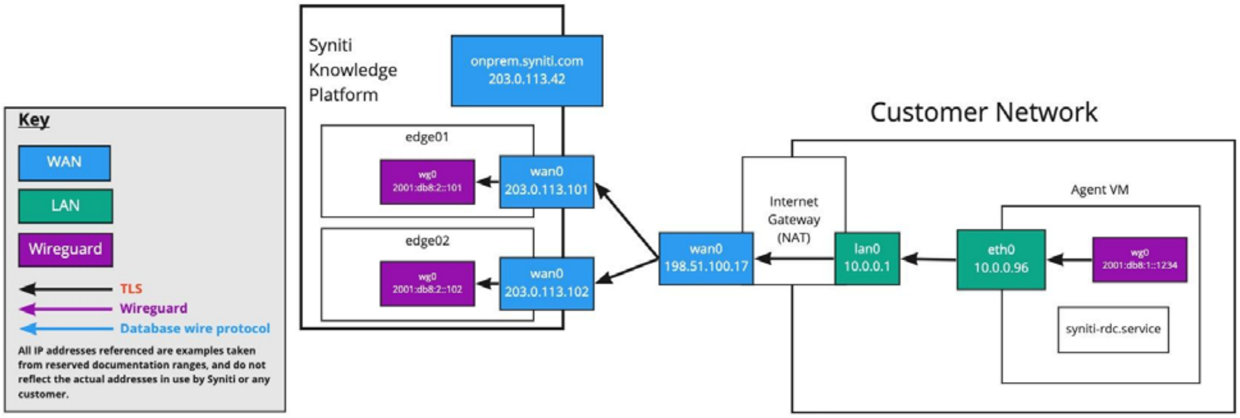

Usage

Once the Knowledge Platform Connector initiates the Wireguard connection to SKP, it is kept alive for bidirectional communication, without requiring the Knowledge Platform Connector VM on the customer's network to be exposed to the Internet directly.

A TCP proxy is started on the VM and bound only to the Wireguard interface, ensuring it cannot be misused by devices on the customer's network or on the public Internet.

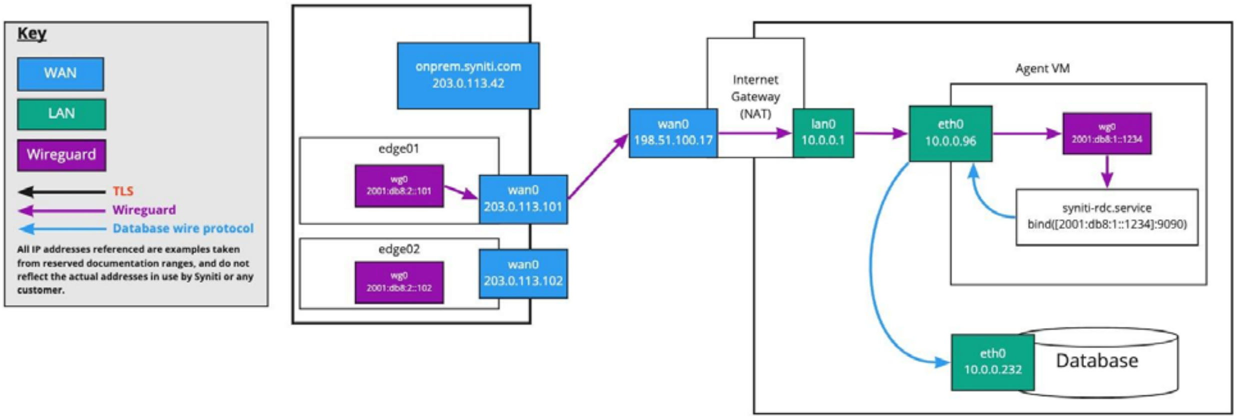

Protocols and Networking

The Knowledge Platform Connector is designed to operate without requiring a public IP address. This allows customers to more easily isolate it from the public internet.

The Knowledge Platform Connector requires inbound UDP Port 58120. All traffic is initiated from the Knowledge Platform Connector, and simply relies on the customer's NAT to permit response packets to be forwarded back to the Knowledge Platform Connector.

Customers are highly encouraged to configure their firewall to only permit egress to Syniti's published edge IP addresses, and the expected internal addresses/ports on their internal network needed to make use of the SKP product. The specifics will depend on the customer's internal data landscape but can be limited to the addresses and ports required to access data and metadata in scope for the SKP usage.

Furthermore, customers may install multiple Knowledge Platform Connectors if desired based on internal network segmentation. For example, if a customer has datacenters in multiple regions, they could install an Knowledge Platform Connector in each region rather than needing to up routing between their regions if desired.

Customers are encouraged to harden the machines used for the Knowledge Platform Connector following current industry best practices, including disabling unneeded services. If SSH is used by the customer for local administration, it is recommended to disable password-based login and configure sshd to listen only on the expected address to prevent it from being reached over the Wireguard interface.

Unless desired, customers should ensure that `net.ipv4.ip_forward` and `net.ipv6.conf.all.forwarding` are disabled in their kernel to prevent the VM from forwarding packets from one interface to another.

Egress Traffic

The Knowledge Platform Connector initiates communication by sending traffic out to the Syniti Cloud using the following ports and protocols:

- TCP port 443 outbound: the Knowledge Platform Connector requests configuration details from the Syniti Cloud using HTTP over TLS.
- UDP port 51820 outbound: the Knowledge Platform Connector initiates a Wireguard tunnel to the Syniti Cloud.

Additionally, the Knowledge Platform Connector will send traffic out to whatever addresses and ports are configured for local data sources. For example, the Knowledge Platform Connector might open a TCP connection to a SQL Server instance on the customer’s LAN. The exact addresses and ports will vary based on the customer’s use of SKP. Some common ports are: 1433 for SQL server and 1521 for Oracle.

Ingress Traffic

The Knowledge Platform Connector VM does not accept any ingress traffic, other than in response to the egress connections it establishes and traffic which is encapsulated inside the Wireguard tunnel.

This also means that the Knowledge Platform Connector does not require a public IP address, it simply needs the ability to egress traffic to the Internet via a NAT gateway.

Encapsulated Traffic

Once the Wireguard tunnel is established, it is used to encapsulate traffic originating from the SKP. The Knowledge Platform Connector binds HTTP servers to the Wireguard interface specifically, enabling the SKT to make requests to the Knowledge Platform Connector on-demand, while protecting the endpoints from the public internet and other devices on the customer's network.

The Knowledge Platform Connector binds servers to the following ports on the Wireguard interface (these ports should not be opened to the Knowledge Platform Connector VM itself):

- TCP port 9090: TCP proxy
- TCP port 9100: This is used to send secure, web-service requests to the Knowledge Platform Connector in order to execute metadata scanning and profiling jobs via associated services running on the Knowledge Platform Connector machine

Outside of these Ports and Protocols **no other ports or protocols are used or should be allowed on the Knowledge Platform Connector machine.**

Example Firewall Rules

The below table serves as an example of recommended and typical firewall rules that can be used to secure the network connectivity of the Knowledge Platform Connector machine.

For the purposes of this example, assume the following network topology (note: these IP addresses are just examples):

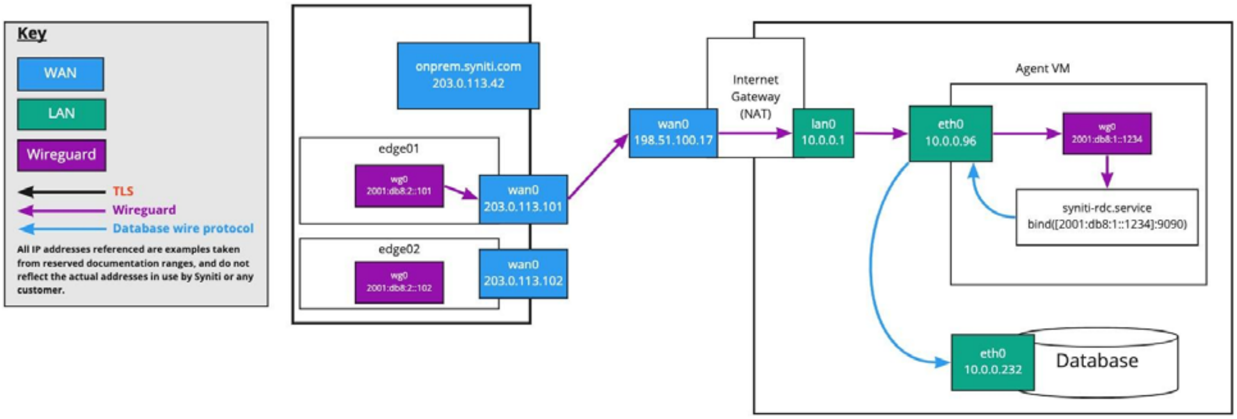

10.0.0.96: the Syniti Knowledge Platform Connector VM on the customer's local network

10.0.0.232: an internal data source to be used with SKP

203.0.113.42: onprem.syniti.com

203.0.113.101: Syniti Wireguard Edge 1

203.0.113.102: Syniti Wireguard Edge 2

Protocol	Source	Destination	Action	Purpose
TCP	10.0.0.96:*	203.0.113.42:443	Allow	Allow Knowledge Platform Connector to poll control plane for configuration
UDP	10.0.0.96:*	203.0.113.101:51820	Allow	Allow Knowledge Platform Connector to initiate Wireguard tunnel to SKP
UDP	10.0.0.96:*	203.0.113.102:51820	Allow	Allow Knowledge Platform Connector to initiate Wireguard tunnel to SKP
TCP	10.0.0.96:*	10.0.0.232:?	Allow	Allow Knowledge Platform Connector to query internal data source
*	10.0.0.96:*	*	Deny	Deny all other outbound traffic from the Knowledge Platform Connector
*	*	10.0.0.96:*	Deny	Deny all inbound traffic to the Knowledge Platform Connector. *Note: customers may choose to enable certain inbound traffic based on their internal management needs.*

Associated Services

While the Syniti Knowledge Platform Connector provides automated and secure connectivity, there are a series of associated services Syniti recommends be installed on the same physical or virtual server as the Knowledge Platform Connector itself. These include:

- Data Profiling - The engine that is used to provide automated, technical profiling of data sources
- Advanced Metadata Scanners - Scanners that connect to and read metadata from data sources to be brought into the Data Catalog
- Data Matching (coming soon) - Execution of de-duplication and matching jobs

Over time, additional services may be added to the collection of functionality that Syniti recommends be installed near or on the same environment as the Knowledge Platform Connector.

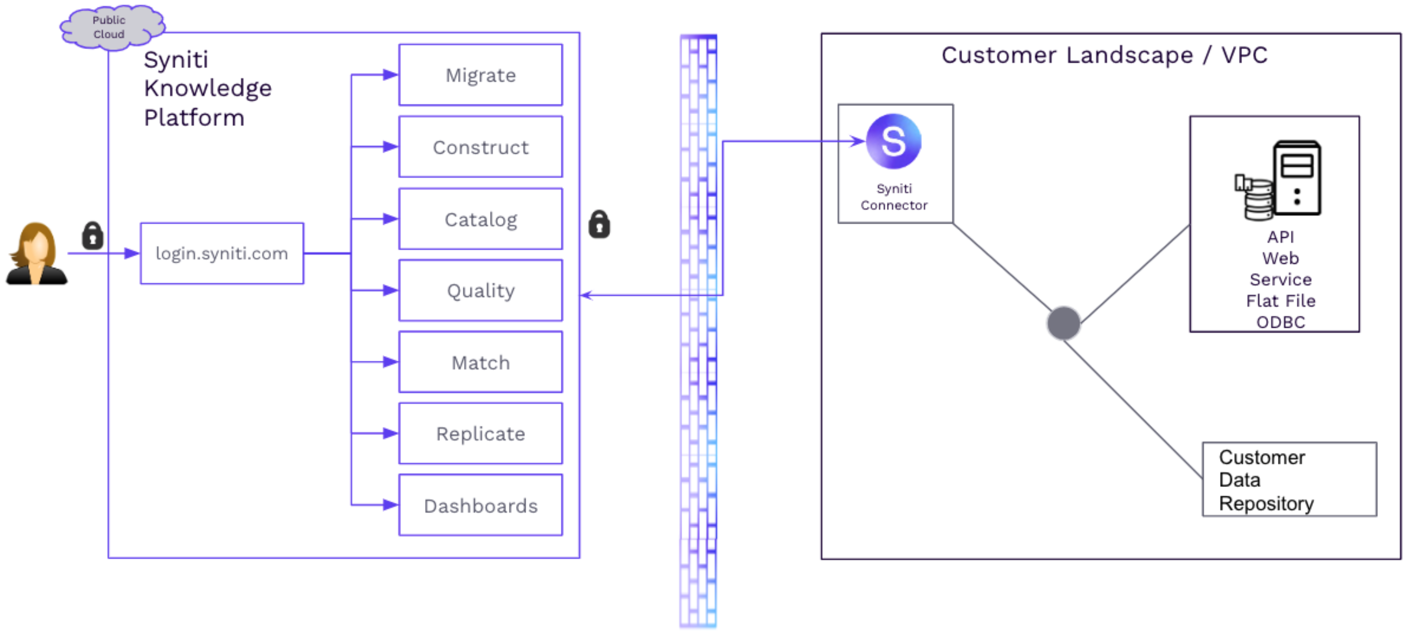

Figure 3: Knowledge Platform Connector and Associated Services