Hosted PrismToken Service

Document number:	PR-D2-1133
Issue date:	2023/06/14
Prepared by:	TrevorD
Copyright:	© 2023 Prism Payment Technologies
Synopsis:	Important information about using our Hosted PrismToken service.

Dear customer,

Thank you for choosing our Hosted PrismToken service.

Getting started: actions you need to take

1. Read this document; it contains important information about using our service.
2. Share this document with your technical staff that develop, maintain, or operate systems that integrate with the service.
3. Send us your Public Static IP ranges for our firewall Allow List.
4. Configure your software with the correct hostname, port, and credentials supplied by us. Your software should respect the protocol and connection limits described in this document.
5. Test before going live! See the supplementary test case suggestions in this document.

Contents

Service description

The Hosted PrismToken service includes:

- Use of STSA Certified token issue software (STS531 “Type A” compliance), including:
    - A web-based User Interface to view the configuration of PrismToken, and to manage certain parts of the service.
    - A callable Application Programming Interface (API) to issue and verify STS tokens.
- Use of STSA Certified Security Modules (with latest STS firmware), which are accessed through the token issue software.
- Security Modules are loaded with Vending Keys per authorization from the SGC User(s). We also load the STS universal default SGC 991993 (base date 1993, terminating November 2024) and SGC 999014 (base date 2014).
- The token issue software and Security Modules are managed by Prism, and may be deployed in a redundant configuration (per contract) that provides availability while protecting against duplicate TokenIds.
- Reserved performance (tokens issued per second) and/or reserved number of vends, per contract.

Tel: +27 31 267 5500 | Email: info@prism.co.za

Security Notice

Cyber attacks are a threat to our business and to yours. Our systems log hundreds of intrusion attempts every day. We follow the best practice of Defense-in-Depth, which means multiple defensive layers to defend against these attacks. Several of these layers affect your interaction with the system:

- We maintain a firewall that only allows access to network ranges on our IP Allow List. You will need to send us your Public Static IP addresses (or address ranges) to add to this list, as described in the next section.
- Our services run on high-numbered non-default ports. You will need to take care to use the correct hostname *and port number* to access the service.
- All connections are protected by TLS v1.2 or higher with secure cipher suites, and our TLS certificates are signed by trusted Certificate Authorities. Your web browsers and vending software will need to support these modern security standards.
- All connections are authenticated. You will be issued with access credentials (typically a Username and Password), that you must present to gain access to the service.

  You must keep these credentials secret and safe, and prevent them from being misused.
- We log access to (and activity on) our systems, and reserve the right to log all traffic (and traffic metadata) on our network. Logs may be used for any legitimate purpose, and retained as long as necessary to serve the purpose. Logs may contain Personally Identifiable Information (PII) such as your username. We use IT security controls to protect these logs against unauthorized access.

  Examples of events we may log include: network packet, TCP connect, TCP disconnect, authentication attempt, HTTP request, HTTP response, API request, API response.

  Examples of what we may log include: date & time, source IP address, authenticated user, system or service accessed, page(s) or information record(s) accessed, nature of access (view/change), inputs & outputs of API calls.

  Examples of how we use logs include: security monitoring, incident analysis, troubleshooting, QOS monitoring, regulatory compliance, security compliance.
- We employ Intrusion Detection mechanisms, and may automatically block potentially-malicious network traffic, and may automatically block the source address(es) entirely, which may result in rejecting all connections from your systems, leading to service interruptions. We may also use Block Lists (DNSBL, IPBL, Bogons, etc.) for this purpose.

  Our Internet Service Providers (and their Carriers) reserve similar rights.

  Examples of malicious traffic include (but are not limited to): port scans, faked or malicious IP headers, vulnerability scans, Denial of Service (DoS) attacks, aggressive connection attempts, or repeated authentication failures.

  Removing such blocks may be time consuming, ranging from several hours for a block automatically added to our firewall, to days for a block automatically added by an ISP or Carrier, to months if your Public Static IP address is added to a public Block List. Removing a block does not prevent it from being added back automatically if further potentially-malicious traffic is detected.

  You must not scan our systems (or those of our Internet upstreams) or otherwise generate malicious traffic. You must take steps to secure your network, so as to prevent malicious traffic from originating from your IP addresses.

IP Allow List

To ensure the security and availability of our service, our firewall will only allow connections from specific IP address ranges.

Please provide us with one or more Public Static IP addresses (or address ranges) that you will use to connect to the service, bearing in mind:

- We need IPv4 addresses. Our service cannot be accessed over IPv6.
- We need the public IP address that our firewall will see, which is assigned to your router by your Internet Service Provider. The public IP address is likely to be different from the IP address of your PC or server. In particular the public IP address cannot be in the Private or "Bogon" ranges listed at https://ipgeolocation.io/resources/bogon.html .
- We need static IP addresses. Many Internet Service Providers will give you a dynamic IP address by default, which means that the public IP address of your router may change every time the router restarts or reconnects, which will cause our firewall to reject connections from your systems, leading to service interruptions.
- We strongly recommend that you check with your ISP to confirm the public static IP address assigned to you.
    - If you do not currently have an assigned static IP address then you should arrange with your ISP to obtain one. Alternatively you could use a commercial VPN service that can provide a static IP.
    - If you cannot use a static IP address at this time we will do our best to accommodate a dynamic IP address range on a temporary basis, using the address confirmation described below.
- Please confirm the public IP address of a PC or server by running the command line `curl https://ipinfo.io/what-is-my-ip` on that PC or server, and send us the entire output.

Connecting to our services

Hostnames, IPs, Ports

We will supply you with hostnames, ports, and credentials to access our systems.

Take care to use the correct hostname and port. Your service is provisioned on a dedicated port that is linked to your credentials; using a default port will not work.

We strongly recommend that you connect to the hostname (DNS name), not to an IP address, for the following reasons:

- We change our DNS entries as necessary to manage the service, including redirecting access to secondary systems during network outages. If you use an IP address your access may be interrupted until you reconfigure your system.
- Some of our web-based services share IP addresses, and need the DNS name in the URL to determine which service is being accessed.

Tel: +27 31 267 5500 | Email: info@prism.co.za
Prism Payment Technologies (Pty) Ltd | Reg No. 1990/005062/07
Directors: H G Kotze, N Pillay, A M R Smith (British) | Company Secretary: C W van Straaten

- Our TLS certificates do not include IP addresses, so you will get security warnings if you connect by IP address. You should never ignore these security warnings!

API protocol & transport

The PrismToken API uses Thrift TBinaryProtocol over TFramedTransport over TLS.

How we manage connections

The Hosted PrismToken service manages connections slightly differently compared to an on-premises NSS:

- Your connection is routed through firewalls and security systems.
- The Thrift port is dedicated to you, and allows a maximum of 64 simultaneous connections. That's 64 total, not 64 per source IP address, or 64 per connecting system. Old connections may block new connection attempts, so take care to close your connections!
- The service will close connections that are idle for more than 600 seconds (10 minutes). If you attempt to send a request on a closed connection your software will report an exception (typically "Connection reset by peer" or "Connection aborted").
- Avoid establishing more than 16 new connections per second; connections made more rapidly than this may be delayed or rejected.

Recommendations

- Use a connection pool that grows on demand (up to a configurable maximum), and closes connections that have lived for longer than 8 minutes.
- Use a connection for only one request at a time. In other words: take a connection out of the pool (so this task has exclusive access; this could be an existing connection or a new one if the pool is empty), send one request, and wait (asynchronously or synchronously) for the response. If the response is a success or a PrismToken ApiException then the connection can be reused: return it to the pool. Otherwise the connection must be closed and discarded.
- Expect (and handle) unexpected disconnections. These can happen for a variety of reasons, including Internet outages between your ISP and ours, over which we have no control. If we have to switch to secondary systems due to a network outage, equipment fault, or regular maintenance, any existing connections will be terminated. We suggest that if you catch a network exception while sending a request, you close and discard the connection and retry with a new connection.
- PrismToken has high latency, so you must use multiple parallel connections to achieve adequate throughput. Plan for 1 connection per 15 TPS required performance, with a minimum of 2 connections. Remember that the service allows only 64 open connections in total.

Redundancy and availability

We operate redundant Security Modules and redundant instances of the token issue software. This section only applies if your hosting contract includes redundancy.

Tel: +27 31 267 5500 | Email: info@prism.co.za
Prism Payment Technologies (Pty) Ltd | Reg No. 1990/005062/07
Directors: H G Kotze, N Pillay, A M R Smith (British) | Company Secretary: C W van Straaten

In normal operation only our Primary systems are active; if any Primary system fails we switch over to a Secondary. We also switch over to Secondary systems during routine maintenance of the Primary systems. We keep the Secondary systems passive to avoid issuing duplicate TokenIds, which could cause STS tokens to be rejected by meters in rare circumstances.

When we switch from Primary to Secondary (or switch back), the service is interrupted briefly. Existing connections will be closed, and new connections will not be accepted during a 1-2 minute window. Thereafter normal API service will resume. When the Secondary is active only a limited Web UI is available.

Maintenance

We perform the following tasks to operate and maintain the hosted systems:

- Refresh Vending Keys on the Security Modules.
    - It is your responsibility to ensure that Vending Key authorization and expiry dates are managed between the SGC User and the Key Management Centre (KMC). We will indicate the specific Security Modules that require authorization.
- Upgrade token issue software and Security Module firmware.
- Check and update operating systems, system configurations, etc.
- Monitor availability.
- Respond to and resolve service interruptions.

Supplementary test case suggestions

In addition to testing all your integrations with PrismToken, we suggest that you consider these supplementary test cases:

- Automatic KCT: the API methods `issueCreditToken()` and `issueMseToken()` take a `meterConfig` parameter that includes flags `allowKrnUpdate` (default true) and `allowKenUpdate` (default true). If a new Vending Key is loaded on PrismToken, or the KeyExpiryNumber (KEN) attribute of a Vending Key is changed (via the KMC), these flags will cause a Key Change Token to be issued before the Credit or MSE token, in which case these methods will return 3 to 5 tokens instead of the usual 1.

Disclaimer

It is your responsibility to ensure that your system is STS compliant. PrismToken is STSA Certified, but incorrect use of the API and/or modification/truncation/deformation of tokens returned by the API may cause your overall system to be non-compliant.

It is your responsibility to ensure that Vending Key authorization and expiry dates are managed between the SGC User and the Key Management Centre (KMC). Prism is not responsible for service interruptions caused by Vending Key expiry.

This information supplement is not part of your contract. Prism makes no representations or warranties whether expressed or implied by or with respect to anything in this document, and shall not be liable for any implied warranties of merchantability or fitness for a particular purpose or for any indirect, special or consequential damages.

Tel: +27 31 267 5500 | Email: info@prism.co.za
Prism Payment Technologies (Pty) Ltd | Reg No. 1990/005062/07
Directors: H G Kotze, N Pillay, A M R Smith (British) | Company Secretary: C W van Straaten