KETRA N4 FAQ

April 8, 2019

Q: My network has a firewall. How do I need to configure it for my N4 devices?

A: Your firewall should be configured to allow outbound traffic on TCP ports 443 and 8883 and UDP port 123. If your N4s have a firmware version earlier than 1.12.3, you must also allow ICMP (ping) traffic to & from addresses 8.8.8.8, 8.8.4.4, 208.67.222.222 and 208.67.220.220. This is not a requirement with N4 firmware 1.12.3 or later but is still recommended

If your N4s are on a different subnet from your computer running Design Studio or the Ketra mobile app, and you have a firewall between the subnets, you should configure the firewall as follows:

Subnet "A" (PC running Design Studio or mobile device running Ketra app)

Subnet "B" (N4s)

Subnet "B" must allow the following traffic from Subnet "A":

- TCP ports 443, 8124, and 8125
- UDP port 4934
- ICMP (ping) traffic
- For discovery of unprovisioned N4s one or more of the following must be the case:
  - SSDP multicast traffic (UDP multicast group "239.255.255.250" port 1900) is routed from Subnet A to B

    Or

  - Subnet A and Subnet B have the same public IP address

Subnet "A" may block all traffic originating from Subnet "B"

If your firewall only allows outbound access to whitelisted servers, you should whitelist my.goketra.com, mqtt.ketra.comand and time.windows.com.

Q: Why do N4s require internet access for a secure installation?

A: This is no longer a requirement as of Design Studio 2.0. An installation can be created without Internet access, but authentication of clients (e.g. Design Studio or the Ketra Mobile Apple) will be disabled. It is recommended to enable authentication if the N4s are connected to the internet at a later time. This can be done in the Device Settings tab of Design Studio. If an installation is created with internet access, authentication will be enabled by default.

continued on next page

N4s use internet access in order to verify that a connecting client (e.g. Design Studio or Ketra Mobile App) has permissions to access the installation. We use OAuth2 for this purpose so that an installation can be shared with other Design Studio user accounts (and permissions later revoked) without the need to inform the N4 of the permissions change.

Here is a brief overview of the process:

- The client makes a request to the Ketra cloud server, providing the Ketra account credentials, and receives an OAuth token. The OAuth token is used as basic authentication credentials for any request to the N4.
- The N4 will contact the Ketra server to verify the validity of the token and allow or deny the request.
- If the request is allowed, the N4 will cache that token so that future requests using that same token don't require validation. Periodically the N4 will contact the Ketra server to verify the validity of each of the OAuth tokens in its cache (to check for permission revocation).
- If the Ketra server can't be contacted at that time, the OAuth tokens will still remain cached so client access is maintained during periods of internet outage.

The N4s also use the internet to update their real-time clock using the NTP protocol, and to also assist in network discovery. See the FAQ entry on network discovery for more details on this topic.

Q: Can I assign a static IP address to my N4s?

A: Static IP configuration is available with Design Studio / Ketra Tech Tool version 1.7 or greater. For version 1.7, this is done in the **Advanced Features** tab of Ketra Tech Tool. In Design Studio version 1.8 this is done in Design Studio's "Device Settings" section. Note that the N4s must be first assigned a DHCP address in order to configure them with a static IP address.

Q: Can I set a custom DNS address on my N4s?

A: If your hubs have firmware version 1.14, you can configure custom DNS addresses with the "advanced features" tab of Ketra Tech Tool version 1.

Q: What protocols are used for N4-to-N4 communication? How much broadcast traffic is generated?

A: The N4s communicate with each other using a compact binary protocol over a TLS 1.2 socket connection with the server on port 13107. In normal use there is no broadcast traffic generated, except in network discovery. See the FAQ entry on network discovery for more details.

Q: Do the N4s have unique SSL certificates?

A: Yes, each N4 has its own SSL certificate, issued by a common self-signed CA. The public certificate of the CA is available here: https://s3.amazonaws.com/ketra-software/KetraMobileAPI/ketra-ca.pem.

Q: What mechanisms are used for discovering N4s on a LAN?

A: The N4s register their local IP address with the Ketra cloud server if they have internet access. Any computer on a network with the same public IP address as the LAN of the N4s can see the registration data of the N4s by viewing https://my.goketra.com/api/n4/v1/query in a web browser.

The N4s also support the UPnP / SSDP discovery protocol, using multicast group 239.255.255.250, port 1900. See section 1.3 of http://www.upnp.org/specs/arch/UPnP-arch-DeviceArchitecture-v1.1.pdf for more details.

If your client is on the same LAN subnet as the N4 hubs, you can send a UDP broadcast packet containing the character '*' (hex 0x2A) and the hubs will respond and indicate their serial number and other information. Even if you are not on the same subnet you can send the same data to a unicast address and an N4 hub at that address will respond as well.

In a multi-pan installation, N4s will use a combination of queries to https://my.goketra.com/api/n4/v1/query and SSDP multicasts to discover their peer N4s. Once all hubs have discovered their peers they no longer need to send any broadcast / multicast messages.

Q: Why do all N4s have to be on the same subnet?

A: This is no longer a requirement, as long as the N4s have firmware version 1.12 or greater and have internet access.

Q: Does my computer running Design Studio need to be on the same subnet as the N4s?

A: When adding new hubs to your installation, your PC must be either on the same subnet, or on a subnet with the same public IP address as the LAN containing the new hubs. Once all your hubs are added, this is no longer a requirement; your PC can be on a different network as long as you can ping the hubs.

Q: Does the N4 support 802.1Q-tagged ethernet frames?

A: No; any tagged frames will be dropped by the N4 but will be forwarded across the N4's L2 switch as appropriate.

Q: What are the pros and cons of daisy-chaining N4s together via the N4's two-port ethernet switch?

A: The benefit of daisy-chaining N4s is that it typically simplifies your wiring and only consumes a single port of your main network switch. The downside is that any power outage to devices in the chain will cut off network access to all downstream devices.

Q: What is the "remote access" feature and how does it work?

A: When "remote access" is enabled for an N4-based installation, the N4 hubs establish a secure connection to a cloud-based message queueing service on port 8883. The message queueing service can route messages between authorized applications on external networks and the N4 hubs, just as if the applications were on the same network as the hub. This allows the application on the external network to perform remote configuration, diagnostics, and lighting control operations. It does not provide a direct bridge to the subnet that the hubs are on; only Ketra-specific protocols are supported. Note that remote access can only be enabled or disabled by Design Studio when it is on the same network as the hubs. Also note that in order to connect remotely to an installation from Design Studio, the network you connect from must not block TCP ports 8883 or 443.