



# BeyondTrust

## **Privileged Remote Access Atlas Technology Deployment Guide**

## Table of Contents

---

<b>Site Configuration to Deploy BeyondTrust Atlas Cluster Technology .....</b>	<b>3</b>
<b>Prerequisites to Set Up Multiple Appliances in Atlas Clusters .....</b>	<b>4</b>
<b>Optional Considerations when Setting Up Multiple Appliances in Atlas Clusters .....</b>	<b>5</b>
<b>Configure the Inter-Appliance Communication Pre-Shared Key for Atlas Clusters .....</b>	<b>6</b>
<b>Configure the Primary Master Node in an Atlas Cluster .....</b>	<b>7</b>
<b>Configure the Traffic Nodes in an Atlas Cluster .....</b>	<b>8</b>
<b>Methods for Choosing Traffic Nodes in an Atlas Cluster .....</b>	<b>9</b>
<b>Perform a BeyondTrust Atlas Technology Cluster Data-Sync .....</b>	<b>11</b>
<b>Perform a BeyondTrust Atlas Technology Cluster Test .....</b>	<b>12</b>
<b>Appendix .....</b>	<b>14</b>

# Site Configuration to Deploy BeyondTrust Atlas Cluster Technology

BeyondTrust Atlas Technology is designed for large-scale geographical deployments of BeyondTrust. With Atlas, you use a single BeyondTrust site across multiple appliances, referred to as nodes in a cluster. Since the administration is primarily performed on a master appliance, BeyondTrust cluster configuration has minimal administration impact.

This guide describes the step-by-step setup details and the options you may wish to consider.



1. Review the [Prerequisites to Set Up Multiple Appliances in Atlas Clusters](#) and the [Optional Considerations when Setting Up Multiple Appliances in Atlas Clusters](#).
2. Configure each Privileged Remote Access appliance as if it were a stand-alone appliance, using the same software package on each appliance.



For more information about configuring your appliances, please see the [Secure Remote Access Appliance Hardware Installation Guide](#) and the [Appliance Administration Guide](#). For more information about BeyondTrust settings and sessions, please see the [Administrative User Guide](#) and [Access Console User Guide](#).

3. [Configure the Inter-Appliance Communication Pre-Shared Key for Atlas Clusters](#).
4. [Configure the Primary Master Node in an Atlas Cluster](#).
5. [Configure the Traffic Nodes in an Atlas Cluster](#), taking into consideration [Methods for Choosing Traffic Nodes in an Atlas Cluster](#).
6. [Perform a BeyondTrust Atlas Technology Cluster Data-Sync](#).
7. [Perform a BeyondTrust Atlas Technology Cluster Test](#).

Should you need any assistance, please contact BeyondTrust Technical Support at [beyondtrust.com/docs/index.htm#support](https://beyondtrust.com/docs/index.htm#support).

## Prerequisites to Set Up Multiple Appliances in Atlas Clusters

You must meet certain prerequisites before you can set up your BeyondTrust cluster.

- **Two B300, B400, or Virtual Appliances**

These appliances act as the master nodes. One will be designated the primary master node and the other will be a backup master node. Both master nodes must match same appliance type: B300 to B300, B400 to B400, or Virtual Appliance to Virtual Appliance. Your need for scalability, capacity, and redundancy will determine appliance needs.

- **Two B300/B400/Virtual Appliance traffic nodes per geographic region in a minimum of two regions**

Traffic nodes can be a mix of B300, B400, and Virtual appliances.

Note, however, that mixing appliance types will yield unbalanced capabilities and potential workflow conflicts. Therefore, it is recommended that all appliances be the same model or type.

You will also need the following hostnames, at a minimum:

- **Site hostname**

This is the hostname that customers will visit to initiate support. This hostname must route to the primary master node in the cluster.

- **Canonical node hostnames**

You must have a unique and unchanging hostname for each master and traffic node. For geographic deployments, consider using the geographic region as part of the hostname. These hostnames should be registered in both the internal and external DNS. Here is an example:

- Primary Master: master1.access.example.com
- Backup Master: master2.access.example.com
- Traffic Node 1: us-traffic1.access.example.com
- Traffic Node 2: us-traffic2.access.example.com
- Traffic Node 3: asia-traffic1.access.example.com

- **Valid SSL certificate for the BeyondTrust site and for each traffic node**

It is recommended you use a valid third-party wildcard certificate that covers both your BeyondTrust support site name and each traffic node hostname. If a wildcard certificate is not used, adding additional traffic nodes that use different certificates may require a rebuild of the BeyondTrust software in order to provide support for mobile and Linux platforms.

- **TCP port 443 open bi-directionally on all appliances**

All appliances must be able to communicate over TCP port 443.

## Optional Considerations when Setting Up Multiple Appliances in Atlas Clusters


### Public vs Internal Hostnames:

You may optionally configure different hostnames for public traffic and private traffic per cluster node. This means you will need two different hostnames for each node in the cluster, one for the **Public Address** and one for the **Internal Address**. All client connections (e.g. access console, customer client, presentation client, etc.) will use the public address. All appliance-to-appliance communications (e.g. cluster data syncs) will use the internal address. This is useful if you want to keep traffic among appliances on a different network route than session traffic.

### Network Address Prefixes:

If you plan to use the **Network Address Prefixes** on the `/login > Management > Cluster > Cluster :: Traffic Nodes` page, then you must define the networks that this traffic node will serve. See "[Configure the Traffic Nodes in an Atlas Cluster](#)" on page 8 for more information.

# Configure the Inter-Appliance Communication Pre-Shared Key for Atlas Clusters

 **Note:** This step can be performed only after all appliances are configured. See the [Privileged Remote Access Appliance Hardware Installation Guide](#) and the [Appliance Administration Guide](#) for details.

Perform the following steps on **all nodes in the cluster**:

1. Go to the **/login > Management > Security** page.
2. Enter a secure password into the **Inter-appliance Communication Pre-shared Key** field. This password must match among all nodes.
3. Click the **Save** button.



## Configure the Primary Master Node in an Atlas Cluster

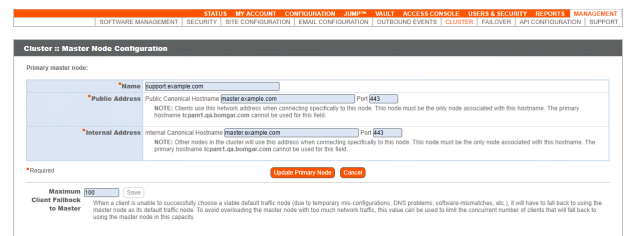
The primary master node is the master node in the BeyondTrust cluster that is configured as the primary site in failover. Information on configuring the Secure Remote Access Appliance for failover is available in the [Failover Configuration](#) document. The network in which the master lives should be a central location in relation to your network as a whole.



**Note:** It is a best practice to set up two master nodes in a failover relationship. However, it is possible though not recommended to have only one master node.

Perform the following steps on the **primary master node**:

1. Go to **/login > Management > Cluster**.
2. Look at the **Cluster :: Master Node Configuration** section.
3. Enter the following information in the **Primary master node** section:
  - **Name:** Enter a name that you will use to remember this node in the cluster. This name must be unique among all nodes in the cluster.
  - **Public Address:**
    - a. Enter the node hostname that you set up in DNS for this node. This should be the canonical hostname unique to the master node, not the primary hostname for the entire cluster (e.g., master.example.com, not access.example.com). (See ["Prerequisites to Set Up Multiple Appliances in Atlas Clusters"](#) on page 4.)
    - b. Enter the port over which clients will communicate with the node. This will usually be port 443.
  - **Internal Address:** This can be the same as the public address. Advanced configurations can optionally set this to a different hostname for inter-appliance communication. (See ["Optional Considerations when Setting Up Multiple Appliances in Atlas Clusters"](#) on page 5.)




**Note:** While IP addresses are recommended for failover setups, BeyondTrust does not currently support IP address usage for Atlas clusters. All master nodes (primary and failover) as well as traffic nodes should use their unique node hostnames. They should NOT use the primary public hostname of the entire cluster or their unique public and/or private IP addresses.

4. Click **Create New Cluster**.




**Note:** If the primary master appliance is in a failover synchronization with a backup master, then the backup master is added to the cluster automatically. If the primary master node is put in failover with a new backup appliance, the new backup appliance is added to the cluster automatically, and the cluster settings are synced to the new backup during the process of establishing failover.

# Configure the Traffic Nodes in an Atlas Cluster

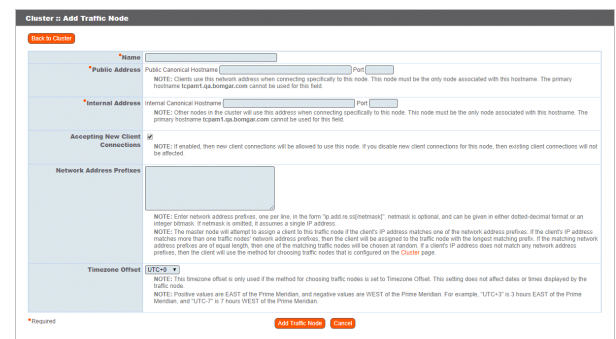
For each traffic node that you want to have in your BeyondTrust cluster, perform the following steps on the **primary master node**:

1. Go to **/login > Management > Cluster**.
2. Look at the **Cluster :: Traffic Nodes** section.
3. Click **Add New Traffic Node**.



4. Enter the following information on the **Cluster :: Add Traffic Node** page:

- **Name:** Enter a name that you will use to remember this node in the cluster. This name must be unique among all nodes in the cluster.
- **Public Address:**
  - a. Enter the node hostname that you set up in DNS for this node. (See ["Prerequisites to Set Up Multiple Appliances in Atlas Clusters"](#) on page 4.)
  - b. Enter the port over which clients will communicate with the node.
- **Internal Address:** This can be the same as the public address. Advanced configurations can optionally set this to a different hostname for inter-appliance communication. (See ["Optional Considerations when Setting Up Multiple Appliances in Atlas Clusters"](#) on page 5.)
- **Accepting New Client Connections:** Be sure this is checked; otherwise, clients will not use the traffic node. As soon as this option is checked, the new setting takes effect automatically via Ajax scripting, and all new BeyondTrust customer client connections are routed to the affected traffic node per the Atlas cluster configuration.
- **Network Address Prefixes:** When this field is populated, the master node attempts to assign a client to this traffic node if the client's IP address matches one of the network address prefixes. If the client's IP address matches more than one traffic node's network address prefixes, the client is assigned to the traffic node with the longest matching prefix. If the matching prefixes are of equal length, one of the matching traffic nodes is chosen at random. If a client's IP address does not match any network address prefixes, the client is assigned using the method configured on the main **Cluster** page (see ["Methods for Choosing Traffic Nodes in an Atlas Cluster"](#) on page 9).



Enter network address prefixes, one per line, in the form of **ip.add.re.ss[/netmask]**. Netmask is optional and can be given in either dotted-decimal format or as an integer bitmask. If netmask is omitted, a single IP address is assumed.

You may leave this field blank.

- **Timezone Offset:** This is used only if the method for choosing traffic nodes is set to **Timezone Offset**. This process involves detecting the time zone setting of the machine hosting the client and using that setting to match the appropriate traffic node which has the closest time zone offset. The time zone offset is derived from the customer time zone setting relative to Coordinated Universal Time (UTC). For countries or zones that use Daylight Saving Time (DST), enter the currently active time zone for the node in question. The cluster offsets for DST automatically.

5. Click **Add Traffic Node**.



## Methods for Choosing Traffic Nodes in an Atlas Cluster

After defining traffic nodes in your BeyondTrust Atlas Technology environment, you can decide on the process which clients use to connect to them.

Session initiation always occurs through the master node and then bridges with the appropriate traffic node. Administrators control and define how a traffic node for a access console or customer client is chosen using the **Method for Choosing Traffic Nodes** dropdown in the master node. Once a node is chosen, the master node usually provides the unique DNS address of the respective traffic node to the client software. The only traffic node setting where an IP address would be specifically provided would be when using the **IP Anycast** selection method.



**Note:** If the network prefixes are defined, the **Method for Choosing Traffic Nodes** setting will be overridden. See ["Configure the Traffic Nodes in an Atlas Cluster"](#) on page 8.

**Cluster :: Traffic Nodes**

Method for Choosing Traffic Nodes: Random (dropdown menu open showing: Random, SRV Record Lookup, A Record Lookup, IP Anycast, Timezone Offset)

NOTE: Clients will use this method to choose a traffic node only if the client's IP address does not match any of the traffic nodes' network address prefixes.

[Save Changes](#)

---

[Add New Traffic Node](#)

Name	Appliance GUID	Public Address	Internal Address	Accepting New Client Connections		
los-angeles	584433b4734b1c3ffc9c117837c4d950	la.support.example.com:443	la.support.example.com:443	<input checked="" type="checkbox"/>	<a href="#">Edit Node</a>	<a href="#">Remove Node</a>
atlanta	3e00068123f0716cffa9f36415bfc0f5	atl.support.example.com:443	atl.support.example.com:443	<input checked="" type="checkbox"/>	<a href="#">Edit Node</a>	<a href="#">Remove Node</a>

The available methods for defining the connection are:

- **Random:** Randomly chooses the node to which a client will connect.

This method will most likely be used if you have taken the time to accurately define all network address prefixes for each traffic node. If a client's network does not match any of the predefined networks on any of the participating traffic nodes, then the client will be assigned a random traffic node. Each traffic node's network address prefixes should be well-defined so that client network matching will be automatic.

This method is simple and inexpensive and enables you to rely on the network prefix defined for each traffic node. However, if your clustered environment spans multiple regions and your network prefixes are left undefined, this method could yield less than optimal results.

- **SRV Record Lookup:** Similar to **A Record Lookup**, SRV traffic node selection will rely on the underlying DNS infrastructure to determine the node to which to connect. The main difference between the two methods is that SRV records have the ability to assign a weight and a priority to a specific host entry. The advantage that this gives you is a method for providing load balancing and backup service at the network level.

Note that this method requires that you have control over the DNS infrastructure used by your clients. If you are deploying in a WAN environment, the use of SRV records is probably already a common practice which you can leverage to provide an extra layer of redundancy and load balancing to your clustered BeyondTrust environment.

- **A Record Lookup:** Instructs clients to attempt connection to a specified (shared) hostname and rely on the DNS configuration to return the appropriate IP address of the traffic node for connection.

This method can be used within an environment where you have complete control of the DNS resources which all of your customers will be using. For instance, you could have an A record defined for traffic1.access.example.com. For your customers in the US who use DNSserver01, the A record points to IP address 1.1.1.1. For your customers in Europe who use DNSserver02, the A record for traffic1.access.example.com resolves to 2.2.2.2.

- **IP Anycast:** Uses a shared IP address among all traffic nodes and relies on the network infrastructure to return the nearest traffic node to the client.

If you are part of an organization that already has a global content delivery network in place, this may be a preferable option for you. IP Anycast is a robust solution but can be complicated to implement and maintain. However, if you already have this type of infrastructure in place, this will be your best method for customer and representative client traffic node selection.

- **Timezone Offset:** A simple and inexpensive method for configuring a BeyondTrust cluster.

The time zone offset process involves detecting the time zone setting of the machine hosting the client and using that setting to match the appropriate traffic node which has the closest time zone offset. The time zone offset is derived from the customer time zone setting relative to Coordinated Universal Time (UTC). The time zone offset method is good for testing and can be used in production. Specifically, in cases where multiple traffic nodes are located in the same time zone, this method may not be the most effective solution. A DNS-based solution would be the preferable method in a production environment.

# Perform a BeyondTrust Atlas Technology Cluster Data-Sync

After you have configured the other settings on the site, you must perform a cluster data sync to make sure the traffic nodes have the same configuration.

When a cluster sync is performed, all configuration data in /login is synced with the traffic nodes, and the failover master syncs via the failover configuration. The latter is controlled from the failover configuration settings. After the initial sync, subsequent synchronization should be performed when changes are made to the **/login > Management > Cluster** page and/or the /login sync warning message appears, prompting you to perform the cluster sync.

Perform the following steps on the primary master node to do a cluster data sync:

1. Go to **/login > Management > Cluster**.
2. Click **Sync Now**.
3. Wait for the sync to finish.
4. Check the **Status History** table to see if the sync succeeded.

After the initial sync, subsequent syncs are also manual. These should be performed as part of regular maintenance.

**Cluster :: Status**

Current Status: This site instance is the primary master node in the cluster. The cluster configuration for the primary master node, backup master nodes, and traffic nodes can be configured below.

**Sync Now** NOTE: This command will sync all configuration data on the master node with the traffic nodes.

**Disband Cluster** NOTE: This command will remove all nodes, including this one, from the cluster. This appliance will attempt to notify the traffic nodes that they are no longer in a cluster. If this appliance is unable to notify a traffic node, then you will have to log in to the traffic node and remove the traffic node from the cluster manually. The cluster status history below will show the result of the notification attempts.

Time	Status	Message
2019-04-23 13:25:45 CDT	0	Data Sync: Finished data sync with traffic nodes.
2019-04-23 13:25:45 CDT	0	Data Sync: Finished configuration data sync with 10.10.10.10.
2019-04-23 13:25:27 CDT	0	Data Sync: Started configuration data sync with 10.10.10.10.
2019-04-23 13:25:20 CDT	0	Data Sync: Started configuration data sync with traffic nodes.

# Perform a BeyondTrust Atlas Technology Cluster Test



**Note:** This step is optional for your deployment.

Given the extreme flexibility of BeyondTrust's Atlas technology, it is impossible to give a detailed and rigorous set of testing steps which will apply in all cases, but a general process with guidelines and expected behaviors should allow administrators to develop more detailed test procedures specific to their environments.

## Review the Planning Process

An Atlas deployment revolves around the master appliance routing client traffic to various nodes. Therefore, testing an Atlas cluster involves three basic steps:

1. Identify which node should be expected to handle any given client connection.
2. Run one or more test connections using BeyondTrust software from the Atlas cluster.
3. Check which traffic node a given test client connects with.

The following sections explain how to plan and implement a testing methodology based on these steps.

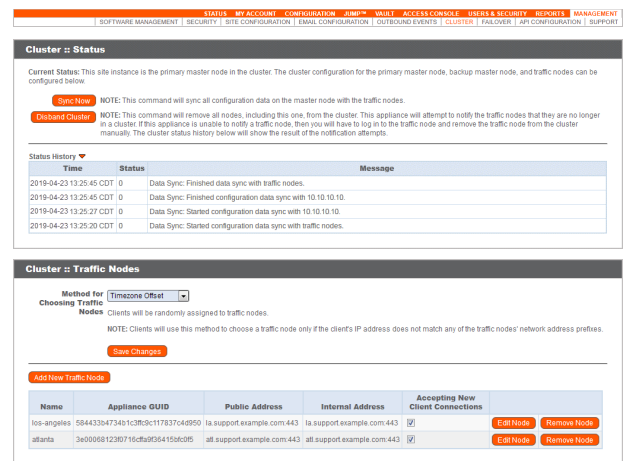
## Identify the Expected Node

The traffic node chosen to handle any given client connection is based on the **Method for Choosing Traffic Nodes** setting. See "[Methods for Choosing Traffic Nodes in an Atlas Cluster](#)" on page 9. The current setting can be checked from the **/login > Management > Cluster** page of the BeyondTrust interface. The first steps of any test, therefore, are to verify the current settings and status of the cluster. To do this, perform the following steps on all traffic nodes in your Secure Remote Access Appliance cluster.

1. Log into the **/login** interface as an administrator.
2. Go to **/login > Management > Cluster**.
3. Verify the configuration details and review the status history.

Depending on the settings, it is possible to artificially route new connections to different traffic nodes by modifying the settings of the client's local host. For example, if **Method for Choosing Traffic Nodes** is set to **Timezone Offset** and the local host's timezone setting is modified such that it matches the timezone offset of the desired traffic node, new BeyondTrust client connections made from the modified host will go to the desired traffic node.

Apart from modifying host settings per the **Method for Choosing Traffic Nodes** setting, it is also possible to simply hard code the network prefixes of the appropriate client hosts into the configuration of the respective traffic node. Once done, clients on the given networks will always route to the traffic nodes assigned to those networks regardless of which method is being used for choosing traffic nodes. This configuration is done from the **Edit Node** option in the cluster configuration page of the primary master node. Simply enter the network prefixes in the **Network Address Prefixes** field of the traffic node to override the extant method for choosing traffic nodes.



The screenshot shows the 'Cluster :: Status' page with a navigation menu at the top. Below the status information, there is a 'Status History' table with columns for Time, Status, and Message. The table shows four entries related to data sync operations. Below the status page, the 'Cluster :: Traffic Nodes' configuration page is visible, showing the 'Method for Choosing Traffic Nodes' set to 'Timezone Offset'. A table at the bottom lists traffic nodes with columns for Name, Appliance GUID, Public Address, Internal Address, and Accepting New Client Connections. Two nodes are listed: 'los-angeles' and 'atlanta', both with checkboxes for accepting new connections and 'Edit Node' and 'Remove Node' buttons.

## Run Test Connections

In general, all BeyondTrust Clients are always connected to the master node while they are online. Once a session is started, the client makes an additional connection to the appropriate traffic node (its home traffic node) based on the cluster configuration logic. In addition, the access console involved in the session will make a third connection, which is to the home traffic node of the remote client involved in the session. Finally, if the representative uses Show My Screen during the session, the remote client makes a connection to the representative's home traffic node.

For example, if a representative in the US remotely connects to a customer in EMEA, the customer client in EMEA connects to the master and the EMEA traffic node (its home traffic node). The access console connects to the master and the US traffic node (its home traffic node). Once the representative starts screen sharing, the access console also connects to the customer's traffic node in EMEA in order to receive the incoming stream of the customer's screen. Thus, the access console is connected to the master, its own home traffic node, and the customer's home traffic node.

To take the scenario one step further, if the representative starts Show My Screen with the customer, then the customer client in EMEA connects to the representative's home traffic node in the US to receive the stream from the representative.

## Appendix

### Peer-to-Peer Functionality

BeyondTrust Privileged Remote Access's peer-to-peer technology is compatible with Atlas deployments.

**i** For more information about peer-to-peer functionality, please see [Options: Manage Session Queuing Options, Record Sessions, Set Up Text Messaging](https://www.beyondtrust.com/docs/privileged-remote-access/getting-started/admin/options.htm) at <https://www.beyondtrust.com/docs/privileged-remote-access/getting-started/admin/options.htm> and [Appliance Administration: Restrict Accounts, Networks, and Ports, Set Up Syslog, Enable Login Agreement, Reset Admin Account](https://www.beyondtrust.com/docs/privileged-remote-access/getting-started/deployment/web/security-appliance-administration.htm) at <https://www.beyondtrust.com/docs/privileged-remote-access/getting-started/deployment/web/security-appliance-administration.htm>.

However, there are a few considerations when attempting to use peer-to-peer with an Atlas architecture:

#### How can BeyondTrust's peer-to-peer functionality be used in an Atlas-configured environment?

For Atlas deployments, BeyondTrust Privileged Remote Access can be configured to use either the BeyondTrust public STUN server, or the Secure Remote Access Appliance (master node) can act as a STUN server for connections.

#### What impact will the availability of the STUN server have on the deployment?

If the Secure Remote Access Appliance (master node) is used as the STUN server, the clients reach out to the master node for session initiation. If the public BeyondTrust STUN server is used, the clients reach out to the public BeyondTrust STUN server for session initiation. Peer-to-peer connections are attempted like any non-Atlas deployment; however, the main difference is the connection falls back to a selected traffic node at session start if the connection attempt to the STUN server is unsuccessful.

#### Are there any special considerations for using the Secure Remote Access Appliance as a STUN Server in an Atlas environment?

The same firewall considerations apply for peer-to-peer in an Atlas deployment as in a non-Atlas deployment. The clients need to reach out to a STUN server, and in this case, the master node acts as the STUN server when the Secure Remote Access Appliance is configured for this role.