



BeyondTrust

Remote Support Atlas Technology White Paper

Table of Contents

Introduction to Clustering with BeyondTrust Atlas Technology	3
Glossary	3
BeyondTrust Atlas Technology Prerequisites	4
BeyondTrust Atlas Technology Overview	5
Benefits	5
Technical Impact	5
Primary Node Configuration for Atlas Clusters	7
Traffic Node Configuration for Atlas Clusters	8
Connection Methods	8
Example: Using Time Zone Offset for Traffic Node Selection in an Atlas Cluster	10
Example: Using a DNS A Record for Traffic Node Selection in an Atlas Cluster	12
Summary	13
Appendix	14

Introduction to Clustering with BeyondTrust Atlas Technology

BeyondTrust Atlas Technology is designed for large scale geographical deployments of BeyondTrust. With Atlas, you use a single BeyondTrust site across multiple B Series Appliances. Since the administration is primarily performed on a primary B Series Appliance, Atlas has minimal administration impact.

This paper describes what comprises BeyondTrust Atlas Technology, how it works at a high level, and the different deployment options for you to consider.

Should you need any assistance, please log into the [Customer Portal](https://beyondtrustcorp.service-now.com/csm) at <https://beyondtrustcorp.service-now.com/csm> to chat with Support.



Glossary

Atlas	The BeyondTrust technology which enables B Series Appliances to be deployed in a cluster.
Cluster	The collective representation of all B Series Appliances that are participating in the same BeyondTrust environment.
Primary Node	The node where a majority of the configuration takes place, such as creating users, defining public sites, configuring support teams, defining traffic nodes, etc. Essentially, everything that you would typically do in a single B Series Appliance BeyondTrust installation /login interface will be done through the designated primary node for your clustered environment.
Backup Primary Node	This node is in a configured failover relationship with the primary node. In the event of a system failure on the primary node, the backup node can take over the role as primary node.
Traffic Node	This node normally handles the bulk of session traffic for the representative console and the customer client. Both the representative console and the customer client bind to a traffic node, as well as the primary node, during a session. The traffic node that is chosen will be determined by the various configuration options.
Inter-appliance Communication Pre-shared Key	This is a password that must be set on all of the B Series Appliances participating in a cluster. This key must match between B Series Appliances in order to replicate information between them and to allow them to participate in the cluster.


BeyondTrust Atlas Technology Prerequisites

In order to run a clustered B Series Appliance environment, the following is required:

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

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

- **One B300/B400/RS Virtual Appliance traffic node per geographic region in a minimum of two regions**

 *Traffic nodes can be a mix of B300, B400, and RS Virtual Appliances, as long as they are appropriately sized to handle the traffic. For recommended sizing, see [SRA Virtual Appliance Installation](#).*

- **Support site hostname**

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

- **Canonical node hostnames**


You must have a unique and unchanging hostname for each primary 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 : primary1.support.example.com
- Backup Primary : primary2.support.example.com
- Traffic Node 1: us-traffic1.support.example.com
- Traffic Node 2: us-traffic2.support.example.com
- Traffic Node 3: asia-traffic1.support.example.com

- **Valid SSL certificate for the BeyondTrust support 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.

You must send BeyondTrust Technical Support a copy of the SSL root certificate and/or B Series Appliance DNS address.

 **Note:** *If a self-signed certificate is used, the certificate serves as its own root certificate, and therefore, the self-signed certificate should be sent to BeyondTrust Technical Support. If a CA-signed certificate is used, contact the CA for a copy of their root certificate. If you have trouble contacting the CA, articles to assist with obtaining your root certificate can be found at beyondtrustcorp.service-now.com/csm. In either case, BeyondTrust Technical Support needs to know the DNS address of the B Series Appliance.*

- **TCP port 443 open bi-directionally on all B Series Appliances**

All B Series Appliances must be able to communicate over TCP port 443.

BeyondTrust Atlas Technology Overview

BeyondTrust Atlas Technology is intended for large enterprise customers performing more concurrent sessions than can be effectively or efficiently handled by a single existing B Series Appliance model. Atlas Technology allows a support organization to be effectively dispersed over different geographical locations and to support a global user base. Essentially, Atlas Technology enables large support organizations to scale horizontally across multiple B Series Appliances rather than vertically on a single B Series Appliance.

Creating a clustered Remote Support environment introduces new terminology: the primary and traffic node concept. The primary node serves as the main point of configuration for the site and also serves as the session initiation point of presence for the entire Remote Support deployment.

A Remote Support administrator accesses the primary node to create a cluster and define the structure of the traffic nodes and method of choosing a traffic node for a client connection. In addition, all configuration of the Remote Support site is handled on the primary node. So even though a cluster consists of multiple B Series Appliances, the /login administrative interface resides on the primary node and propagates most configuration settings to the traffic nodes automatically. The traffic nodes retain a /login interface on each respective B Series Appliance; however, the respective B Series Appliance has limited configuration settings available.

Licenses are designated for the site as a whole, and license utilization is not affected by the fact that there are multiple B Series Appliances involved.

All reporting is handled the primary. The session recordings reside on the respective traffic node where a customer client connects; however, when requesting to view any of the recordings, a dynamic link allows expected Remote Support reporting behavior just as if the recording resided on the primary itself.



For information on Atlas in the Cloud, please see [BeyondTrust Atlas in the Cloud](https://www.beyondtrust.com/docs/remote-support/getting-started/deployment/cloud/atlas-cloud.htm) at <https://www.beyondtrust.com/docs/remote-support/getting-started/deployment/cloud/atlas-cloud.htm>.

Benefits

A key benefit of clustering via BeyondTrust Atlas Technology is the ability to distribute a site geographically. This is important in situations where a support organization may span regions or have global reach. For instance, if a customer support request originates in Sydney and a traffic B Series Appliance residing in Australia handles the support session, then the support experience is more responsive and efficient. This is mainly due to the bulk of the session traffic staying local to the B Series Appliance in Australia versus the client using a traffic node in NYC, where it would have to traverse all traffic data back and forth to NYC, thereby increasing the transport latency of the session.

Technical Impact

In a clustered environment, all Remote Support traffic originates by first talking to the primary node. The representative console is downloaded from the primary node, and authentication into the representative console takes place against the primary node. Thus, any external authentication providers that need to be configured in your environment are done at the primary node level.

Initiating an attended support session is still done in the same method as a non-clustered environment. The public portal for your support site resides on the primary node B Series Appliance. From here, a customer can choose from the representative list, enter a session key, or use issue submission. Session initiation always occurs through the primary node and then bridges with the appropriate traffic node once the session is initiated.

Administrators control and define how a traffic node for a representative or customer client is chosen. The representative and customer independently bind to their own traffic nodes. Each may bind additionally to the other's traffic node depending on what is occurring within the session. If screen sharing is initiated, then the representative binds to the traffic node that the customer client is bound to, in order to receive the traffic stream that contains the actual screen sharing information.

Likewise, if the representative shares their screen within the session and chooses to send a file from their machine to the customer via the chat interface, then the customer client binds to the traffic node of the representative in order to receive the incoming file or to view the representative's screen. When a representative transfers a session or if a session is shared between representatives, then the incoming representative binds to the traffic node of the customer in order to view the customer's screen. All of this coordination between traffic nodes and clients is controlled by the primary node and happens automatically in the background.

When deploying Jump Clients in a clustered environment, the Jump Clients are initially deployed from, and communicate with, the primary node. Once deployed, they resolve a priority list of traffic nodes based on the site's currently set connection method. Jump Clients reconnect and use a traffic node to obtain future updates and to proxy communications to the primary node. This allows more Jump Clients to upgrade at once, and additionally allows the primary node to handle sessions and normal traffic with less customer impact during the process. If a Jump Client is unable to connect to its preferred traffic node, either due to capacity or an outage event, it falls back to another traffic node, or to the primary node if no traffic nodes are available.

As mentioned, all reporting is done from the primary node /login interface. The actual session recordings reside on the traffic node B Series Appliance that the customer client had bound to. If an aggregate, off-B Series Appliance session log store including session recordings is needed, a BeyondTrust Integration Client must be configured to talk to the primary node B Series Appliance and must be able to reach all traffic node B Series Appliances in the cluster.



Note: Representatives using a mobile console to provide support always bind to the primary node. Similarly, customers using a mobile customer client always bind to the primary node.

Primary Node Configuration for Atlas Clusters

The primary node configuration in itself is rather simple. First, you must choose which B Series Appliance will serve as the primary node. Unless the deployment is for a small number of representatives, the primary node will ideally be a B400 B Series Appliance. A second, matching backup will need to be used for the pairing of the primary role in a failover relationship.

Since the primary node will play a role in every support session, the network in which that primary node resides should be a central location in relation to your network as a whole.

Once you have planned where your primary node will reside physically, the next step is to confirm the name of your support site. This hostname will serve as the central hub, essentially, where your customers will come seeking support (e.g., <http://support.example.com>). You will also need to have a canonical hostname registered in your DNS environment for each B Series Appliance in the cluster.

The primary node will also have the capacity to handle support sessions just as does a traffic node. If there are network or environmental conditions disrupting the availability of a traffic node (from a client's point-of-view) then a support session can fall back to the primary B Series Appliance. In this scenario, the primary B Series Appliance will handle all aspects of the session without utilizing a traffic node. An administrator can set how many concurrent sessions can fall back to the primary B Series Appliance at any given time.


Traffic Node Configuration for Atlas Clusters

When adding a traffic node to your clustered environment, you define the name of that node as well as its canonical hostname. Also, you may associate specific networks with the traffic node. This essentially predetermines a client's traffic node selection based on its network prefix mask. This type of configuration is more relevant for administrators who are deploying a clustered BeyondTrust site in a WAN environment.

One last configuration option available when initially defining a traffic node is the time zone offset of the B Series Appliance. This must be set if you plan to use the time zone offset method (which is discussed in more detail below) for clients deciding which traffic node to connect to.

After defining traffic nodes in your environment, you can decide on what process clients will use to determine which traffic node to connect to. BeyondTrust administrators have the following options to choose from:

Connection Methods

Time Zone Offset	<p>The time zone offset process involves detecting the time zone setting of the client machine and using that setting to match the nearest available traffic node. The time zone offset is derived from the client machine's time zone setting relative to Coordinated Universal Time (UTC). The time zone offset method is good for testing and can be used in production; however, a DNS-based solution would be a preferable method in a production environment.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  <p>Note: For environments where the time zone offset method is configured and a support session is initiated via the Session Generation API, the primary node redirects the customer client to the closest available traffic node.</p> </div>
IP Anycast	<p>The IP Any Cast method uses a shared IP address among all traffic nodes and relies on the network infrastructure to return the "closest" 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 Any Cast is a very robust solution but can be more complicated to implement and maintain. However, if you have this type of infrastructure in place, this will be your best method for customer and representative client traffic node selection.</p>
A Record	<p>The A record method instructs clients to attempt to connect 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 utilized within an environment where you have complete control of the DNS resources that all of your customers will be using. Also, there are third party DNS providers that can provide this service for you. With this method, you could have an A record defined for trafficnodepicker.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 trafficnode01.example.com resolves to 2.2.2.2.</p>

Random	<p>The random method randomly chooses which nodes a client will connect to. This method will most likely be used if you have taken the time to accurately define all the network prefixes for each respective traffic node. If a client's network doesn't match any of the predefined networks on any of the participating traffic nodes, then the client will be assigned a random traffic node at the discretion of the primary node. This method is simple and inexpensive, and it enables you to rely on the network prefix defined for each traffic node. However, if your clustered environment spans multiple regions or the globe and your network prefixes are left undefined, this method could yield less than desirable results.</p>
SRV Record	<p>The SRV record method is very similar to the A record process in that traffic node selection relies on the underlying DNS infrastructure to determine which node to connect to. The main difference between the two methods is that SRV records have the ability to assign a weight and 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.</p> <p>Note that this method requires that you have control over the DNS infrastructure that your clients will be using. 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.</p>

Example: Using Time Zone Offset for Traffic Node Selection in an Atlas Cluster

An example of a BeyondTrust clustered deployment that uses time zone offset traffic node selection is the Paxton Thomas support organization, <http://support.paxtonthomas.com>. Representatives are located in different geographic locations: Boston, Oakland, and London. Paxton Thomas has datacenters in Dallas, Oakland, Boston, and London. Paxton Thomas chooses the Dallas datacenter as the location for the primary node based on its available resources, such as rack space, adequate power and cooling, sufficient bandwidth, and its central location.

Primary Node Selection and Setup

Paxton Thomas chooses a B300 B Series Appliance to serve as their primary node, as they will have less than 300 concurrent logged in representatives at one given time. Paxton Thomas's failover strategy is to place a second B300 in its Boston datacenter to serve as the backup primary node for the cluster. Therefore, if there is a total outage in the Dallas datacenter, operations can fail over to the backup primary located in Boston. Since the primary and backup B Series Appliances reside on different network segments, Paxton will be required to use either the DNS or NAT swing approach as part of its failover process.

Traffic Node Considerations

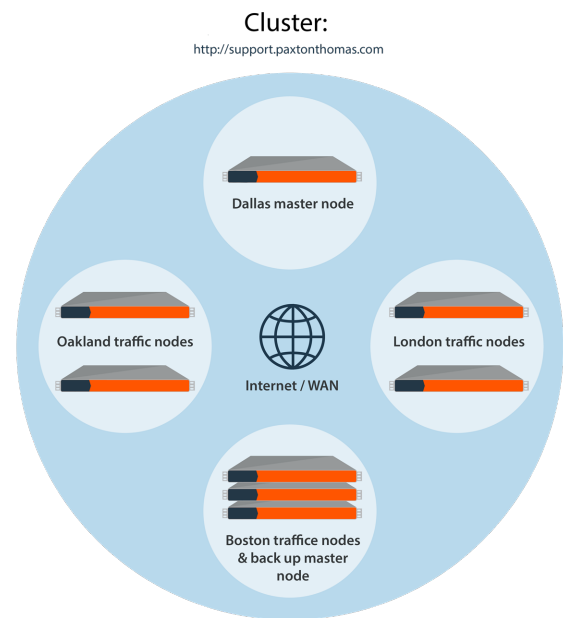
After deciding where the primary and backup nodes will reside, Paxton Thomas then examines where the majority of the customers being supported are located. After reviewing historical trends of closed support tickets, it is evident that a majority of the sessions are confined to either the east or west coast, with the remaining sessions located in either the central US or in Europe.

Based on this information, Paxton Thomas decides to place two traffic nodes in the Oakland data center, two traffic nodes in the Boston data center, and finally two nodes in the London data center. Each traffic node is assigned a unique hostname, and the time zone offset for each respective B Series Appliance is set according to its physical location. Once the traffic nodes are deployed, configured, and have joined the cluster, the setup is complete.

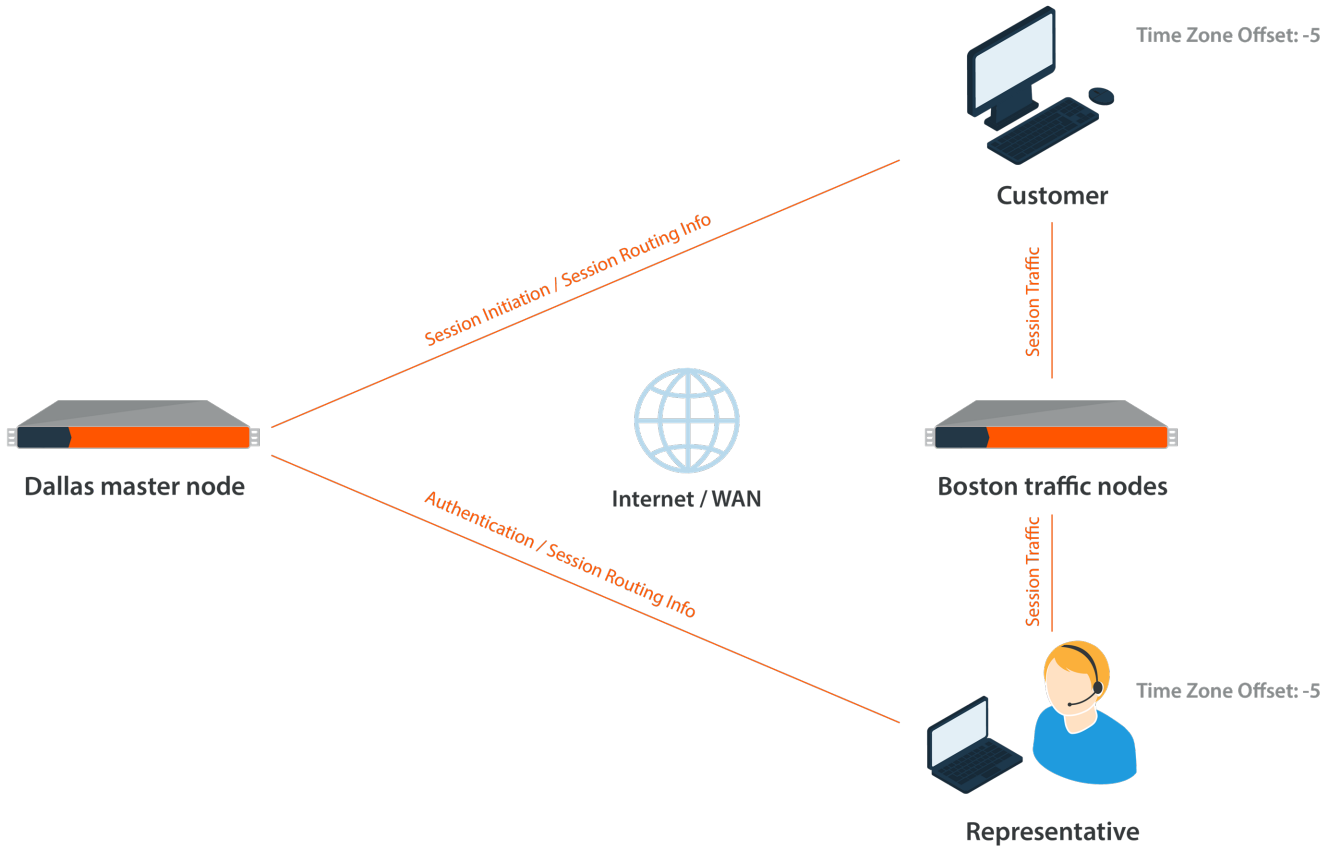
Node Selection in Action

With the clustered configuration completed, the Paxton Thomas support organization is now ready to take support sessions. A Paxton Thomas support representative logs into the representative console (authentication is taking place against the primary), and the console detects that the time zone offset on the representative's computer is -5, which means this representative is in the eastern time zone and, for this example, is based out of the Boston office. The primary tells the representative console to connect to one of the two traffic nodes in the Boston datacenter using the hostname for one of the two traffic nodes. Both of these nodes have the same time zone offset setting as the representative (-5), so the primary directs the representative to the less busy node.

The logged in representative now receives a call from a customer in Atlanta. The representative emails a session key to the customer, taking them to the <http://support.paxtonthomas.com> site, which resides on the primary node B Series Appliance in Dallas. Prior to the installation of the customer client, the time zone offset of the customer is determined (in this scenario, it is -5). The primary node determines that the closest traffic node for this specific customer is one of the two traffic nodes in the Boston datacenter, as they also have a -5 time zone offset. Therefore, the primary node chooses one of the two Boston traffic nodes from which to download the customer client, the client connects to that traffic node, and the session is initiated.



In this scenario, the representative has a connection to the primary node and the designated traffic node in Boston. The customer client has a connection to the primary B Series Appliance in Dallas as well as connection to a traffic node in Boston. Throughout the session, the primary node coordinates the traffic between the traffic nodes being used to conduct the session. The bulk of the session traffic takes place at the traffic node level, such as screen sharing and file transfer, while the connections from both the representative and the customer to the primary contain small pieces of information that coordinate the actions between traffic nodes and create the session log. The video of the session recording resides on the traffic node that the customer client is connected to. In this example, the video recording resides on the traffic node in Boston.



Example: Using a DNS A Record for Traffic Node Selection in an Atlas Cluster

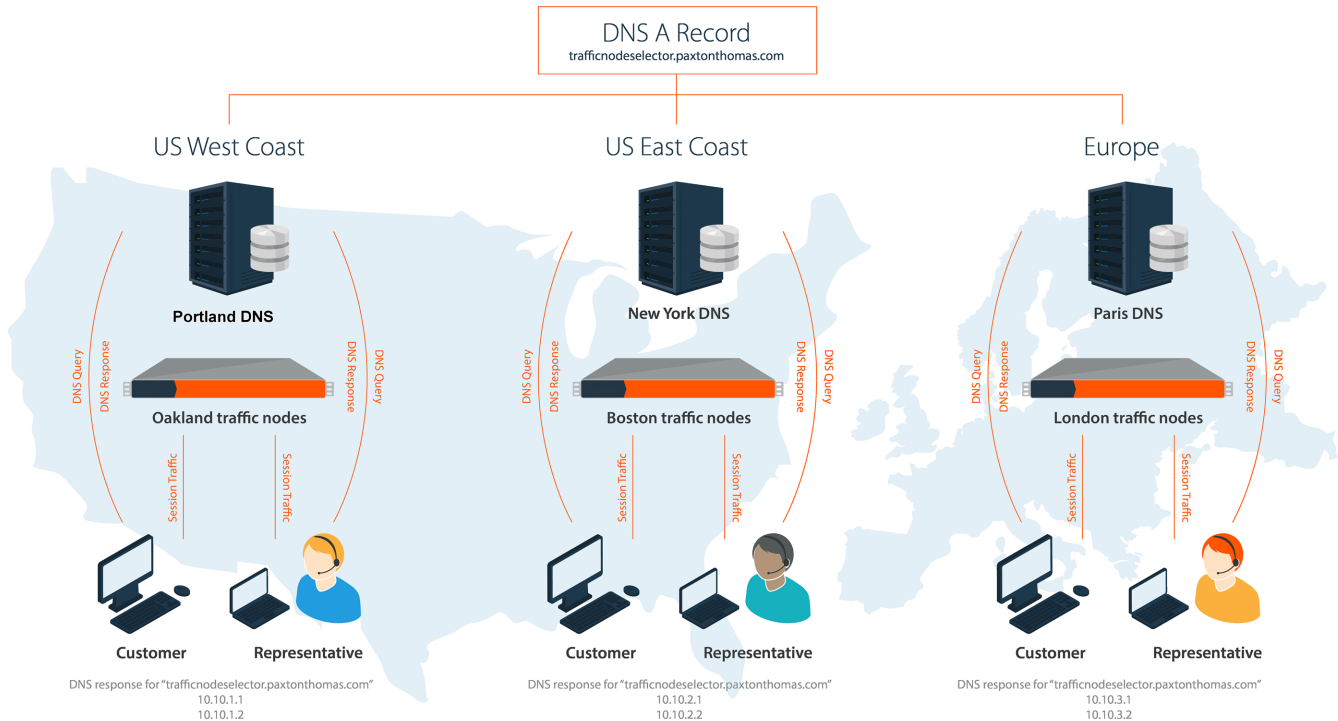
In this example clustered configuration, we use a different traffic node selection method. The Paxton Thomas support organization has subscribed to a third-party DNS solution provider. The DNS provider provides hosting for the paxtonthomas.com domain and also has a special offering that allows Paxton Thomas to create a single A record that utilizes their traffic management functionality, which essentially determines where a DNS request should be routed.

Node Selection in Action

The third-party DNS provider has DNS servers strategically placed in the different geographic locations throughout the US, Europe, and Asia. BeyondTrust creates an A record for the name of trafficnodepicker.paxtonthomas.com and within the DNS management interface specifies the IP address of each traffic node that is in the specific region for each DNS server.

For example, on the DNS servers that are responsible for the "US East Coast" region and physically reside in New York, the server resolves the DNS name trafficnodepicker.paxtonthomas.com to one of the two IP addresses for the B Series Appliances located in the Boston datacenter. Likewise, the DNS server responsible for regions in Europe and physically residing in Paris resolves the DNS name trafficnodepicker.paxtonthomas.com to one of the two IP addresses for the B Series Appliances located in the London datacenter.

This traffic node selection process works well. However, it does require more administrative overhead to maintain the DNS infrastructure. Also, there exists a potential for additional cost when utilizing a third party to host DNS. So, when choosing which method to pursue for your traffic node selection process, it is important to consider factors external to your BeyondTrust environment, which may increase cost and complexity to maintain your clustered environment.



Summary

BeyondTrust Atlas Technology enables the clustering of multiple B Series Appliances in your support environment. Atlas Technology is a highly configurable and robust solution to ensure the best possible support experience for your customers. The ability to efficiently scale geographically is crucial where support organizations may span regionally or globally. BeyondTrust is the industry leader in remote support and privileged access applications relied upon by thousands of customers worldwide. Our Atlas Technology is an exciting component of our proven ability to deliver rock solid secure access software to customers around the globe.

Appendix

Peer-to-Peer Functionality

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

i To learn more about peer-to-peer functionality, please see [Options: Manage Session Queuing Options, Record Sessions, Set Up Text Messaging](https://www.beyondtrust.com/docs/remote-support/getting-started/admin/options.htm) at <https://www.beyondtrust.com/docs/remote-support/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/remote-support/getting-started/deployment/web/security-appliance-administration.htm) at <https://www.beyondtrust.com/docs/remote-support/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 Remote Support can be configured to use either the BeyondTrust public STUN server, or the BeyondTrust Appliance B Series (primary node) can act as a STUN server for connections.

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

If the BeyondTrust Appliance B Series (primary node) is used as the STUN server, the clients reach out to the primary 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 BeyondTrust Appliance B Series 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 primary node acts as the STUN server when the BeyondTrust Appliance B Series is configured for this role.