

The logo for NETSCOUT, with 'NETSCOUT' in white and 'O' in green. The background of the entire page is a black and white photograph of a city skyline at night, featuring a complex multi-level highway interchange in the foreground and numerous illuminated skyscrapers in the background.

NETSCOUT.

Getting Started with NETSCOUT nGeniusPULSE for Amazon Web Services

Virtual nGeniusPULSE

July 12, 2025

PN: 733-2143, Rev. B

Use of this product is subject to the End User License Agreement available at <http://www.NetScout.com/legal/terms-and-conditions> or which accompanies the product at the time of shipment or, if applicable, the legal agreement executed by and between NetScout Systems, Inc. or one of its wholly-owned subsidiaries ("NETSCOUT") and the purchaser of this product ("Agreement").

Government Use and Notice of Restricted Rights: In U.S. government ("Government") contracts or subcontracts, Customer will provide that the Products and Documentation, including any technical data (collectively "Materials"), sold or delivered pursuant to this Agreement for Government use are commercial as defined in Federal Acquisition Regulation ("FAR") 2.101 and any supplement and further are provided with RESTRICTED RIGHTS. All Materials were fully developed at private expense. Use, duplication, release, modification, transfer, or disclosure ("Use") of the Materials is restricted by the terms of this Agreement and further restricted in accordance with FAR 52.227-14 for civilian Government agency purposes and 252.227-7015 of the Defense Federal Acquisition Regulations Supplement ("DFARS") for military Government agency purposes, or the similar acquisition regulations of other applicable Government organizations, as applicable and amended. The Use of Materials is restricted by the terms of this Agreement, and, in accordance with DFARS Section 227.7202 and FAR Section 12.212, is further restricted in accordance with the terms of NETSCOUT'S commercial End User License Agreement. All other Use is prohibited, except as described herein.

This Product may contain third-party technology. NETSCOUT may license such third-party technology and documentation ("Third-Party Materials") for use with the Product only. In the event the Product contains Third-Party Materials, or in the event you have the option to use the Product in conjunction with Third-Party Materials (as identified by NETSCOUT in the Documentation provided with this Product), then such third-party materials are provided or accessible subject to the applicable third-party terms and conditions contained either in the "Read Me" or "About" file located in the Software or on an Application CD provided with this Product, or in an appendix located in the documentation provided with this Product. To the extent the Product includes Third-Party Materials licensed to NETSCOUT by third parties, those third parties are third-party beneficiaries of, and may enforce, the applicable provisions of such third-party terms and conditions.

Open-Source Software Acknowledgment: This product may incorporate open-source components that are governed by the GNU General Public License ("GPL") or licenses that are compatible with the GPL license ("GPL Compatible License"). In accordance with the terms of the GNU GPL, NETSCOUT will make available a complete, machine-readable copy of the source code components of this product covered by the GPL or applicable GPL Compatible License, if any, upon receipt of a written request. Please identify the product and send a request to:

NETSCOUT SYSTEMS, INC.
GNU GPL Source Code Request
310 Littleton Road
Westford, MA 01886
Attn: Legal Department

To the extent applicable, the following information is provided for FCC compliance of Class A devices:

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

Modifications to this product not authorized by NETSCOUT could void the FCC approval and terminate your authority to operate the product. Please also see NETSCOUT's Compliance and Safety Warnings for NetScout Hardware Products document, which can be found in the documents accompanying the equipment, or in the event such document is not included with the product, please see the compliance and safety warning section of the user guides and installation manuals.

No portion of this document may be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine form without prior consent in writing from NETSCOUT. The information in this document is subject to change without notice and does not represent a commitment on the part of NETSCOUT.

The products and specifications, configurations, and other technical information regarding the products described or referenced in this document are subject to change without notice and NETSCOUT reserves the right, at its sole discretion, to make changes at any time in its technical information, specifications, service, and support programs. All statements, technical information, and recommendations contained in this document are believed to be accurate and reliable but are presented "as is" without warranty of any kind, express or implied. You must take full responsibility for their application of any products specified in this document. NETSCOUT makes no implied warranties of merchantability or fitness for a purpose as a result of this document or the information described or referenced within, and all other warranties, express or implied, are excluded.

Except where otherwise indicated, the information contained in this document represents the planned capabilities and intended functionality offered by the product and version number identified on the front of this document. Screen images depicted in this document are representative and intended to serve as example images only.

Copyright © NETSCOUT 2009-7/12/25. All rights reserved.
PN: 733-2143

Contacting NETSCOUT SYSTEMS

Customer Support

The best way to contact Customer Support is to submit a Support Request:
<https://my.netscout.com/Pages/Overview.aspx>

Telephone: In the US, call **888-357-7667**; outside the US, call **001 978-614-4000**. Phone support hours are 8 a.m. to 8 p.m. Eastern Standard Time (EST).

When you contact Customer Support, the following information can be helpful in diagnosing and solving problems:

- Type of network platform
- Software and firmware versions
- Hardware model number
- License number and your organization's name
- The text of any error messages
- Supporting screen images, logs, and error files, as appropriate
- A detailed description of the problem

Sales

Call **800-357-7666** for the sales office nearest your location.

Education and Training

Education and training resources including course listings, product certification, webinars, and case studies are available at:
<http://www.netscout.com/education/overview/>

Solution Components	2
Detailed Deployment Architecture	3
System Requirements – Amazon Web Services	4
Skills and Specialized Knowledge Recommendations	5
About Pricing and Costs	5
About Service Quotas	5
Licensing Model – BYOL	6
About BYOL Licenses	6
Obtaining BYOL Licensing Information	6
Deploying nGeniusPulse from AWS Marketplace (25-30 Minutes)	7
Important – nGeniusPULSE Startup Takes Time!	12
Logging in to the Virtual nGeniusPULSE User Interface (5-10 Minutes)	13
Security Group Details	16
Instance Type Recommendations	17
Connecting to Instances (5-10 Minutes)	17
Operational Guidance	19
Snapshot and Backup Procedures	19
Backing Up nGeniusPULSE	19
Snapshot Examples by Target RPO	20
Routine Maintenance	20
Security Notes	20
Disaster Recovery	22
Disaster Recovery: Key Concepts	22
Sample Disaster Recovery Plans	22
Availability Zone Recovery	22
Region Recovery	23

Getting Started with NETSCOUT nGeniusPULSE

This document describes how to get started using nGeniusPULSE® with Amazon Web Services (AWS). See the following sections for details:

- "Solution Components" on page 2
- "System Requirements – Amazon Web Services" on page 4
- "Obtaining BYOL Licensing Information" on page 6
- "Deploying nGeniusPulse from AWS Marketplace (25-30 Minutes)" on page 7
 - "Logging in to the Virtual nGeniusPULSE User Interface (5-10 Minutes)" on page 13
 - "Security Group Details" on page 16
 - "Connecting to Instances (5-10 Minutes)" on page 17
- "Operational Guidance" on page 19
- "Security Notes" on page 20
- "Disaster Recovery" on page 22

Additional Resources

NETSCOUT® Systems strongly recommends that you read this document in its entirety, as well as the most recent versions of the following additional documentation available online at [My.NETSCOUT](#):

- *Virtual nGeniusPULSE Installation Guide*
- *nGeniusPULSE* documentation and Online Help

Note: For the most current and comprehensive information, visit the NETSCOUT Technical Support knowledge base at the following URL: <https://my.netscout.com/pages/mcplanding.aspx>. This site contains related documents, tips, FAQs, and suggested workarounds. You can also download updated copies of product documentation from this site.

Solution Components

The nGeniusPULSE solution consists of the **Virtual nGeniusPULSE** console to deliver an overarching view into the performance of all infrastructure and application components across geographically dispersed data centers and cloud (Figure 1).

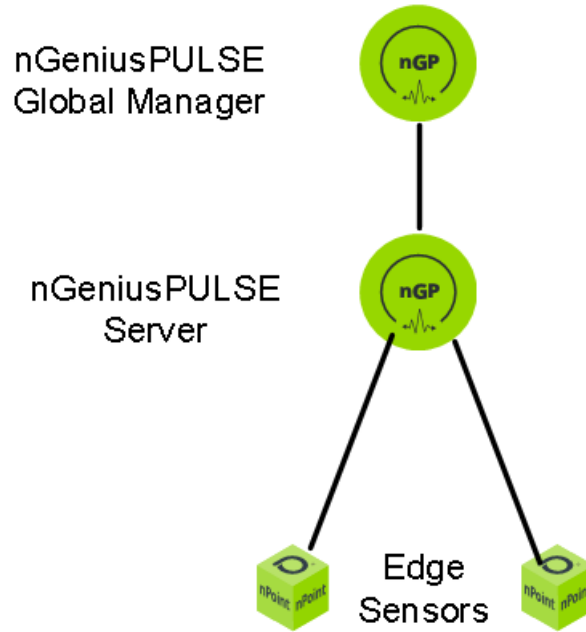





Figure 1 Detailed View of NETSCOUT Synthetic Transaction Testing Components

The table below summarizes the role of each of these components:

<p>nGeniusPULSE Server</p> <ul style="list-style-type: none"> • Collects and stores Synthetic test results from hardware or virtual nPoints or Edge Sensors • Configure/Manage nPoints or Edge Sensors • Configure/Manage Synthetic Tests for nPoints/Edge Sensors • Integration with nGeniusONE for Synthetic Test packet data 	
<p>nGeniusPULSE Global Manager</p> <ul style="list-style-type: none"> • Provides a single-pane of glass for deployments where there are multiple nGeniusPULSE Servers • Supports connecting up to 10 nGeniusPULSE Servers • Provides consolidated Global views of nGeniusPULSE synthetic test data as well as direct access to the connected nGeniusPULSE Servers • Can be integrated with nGeniusONE Dedicated Global Manager 	
<p>nGeniusPULSE Standby Server</p> <ul style="list-style-type: none"> • nGeniusPULSE Server that is connected to a Primary nGeniusPULSE Server that is operating in Standby mode (services stopped) and synchronizes data between the two servers • This is meant as a “Warm” Standby setup. Meaning if the Primary server goes offline, the user needs to manually start the nGeniusPULSE Services • In an Active/Standby Deployment, a proxy or load balancer is required between the nPoints and the nGeniusPULSE servers. 	

nPoint/Edge Sensor

- nPoints or Edge Sensors are instrumentation that are connected to the nGeniusPULSE server and run synthetic tests
- Synthetic tests are configured and managed on the nGeniusPULSE Server and pushed down to the individual nPoints/Edge Sensors
- This instrumentation can be hardware (Edge Sensor 490, rISNG 690/695 or 3000H), or Virtual (Edge Sensor 290 or 3000V).
- nPoints/Edge Sensors are meant to be deployed in remote locations such as Offices, Data Centers, Retail locations, Remote Edge locations, Work-from-home, and so on.



Detailed Deployment Architecture

Figure 2 illustrates a sample of a multi-VPC, load-balanced deployment, including an auto-scaling application with multi-AZ databases. Note the following:

- nGeniusPULSE resides in a separate VPC from the monitored application deployment. Although this example shows both VPCs in the same AWS Region, they can also be in separate regions.
- NETSCOUT's CloudFormation templates in the AWS Marketplace are used to perform the deployment of nGeniusPULSE instances.

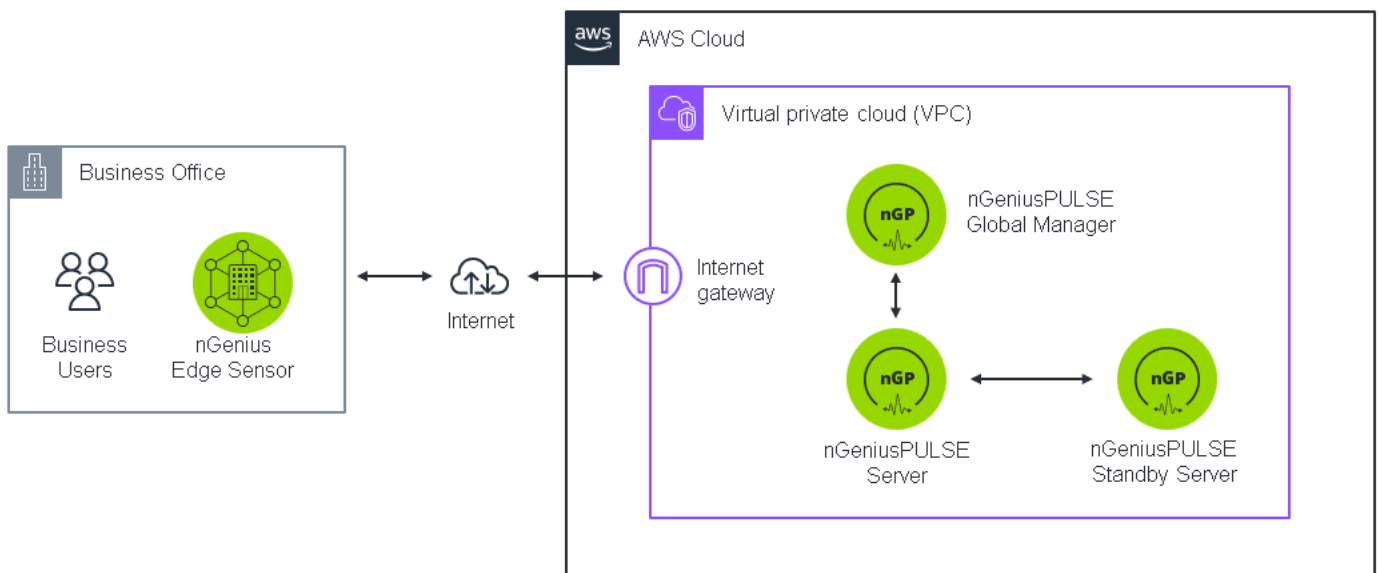


Figure 2 Detailed Deployment Diagram

System Requirements – Amazon Web Services

Table 1 summarizes the necessary requirements to deploy the NETSCOUT Smart Data solution:

Table 1 Deployment Requirements

Component	Description
Amazon Web Services Account	You must have an active Amazon Web Services account with access to the EC2 Management Console to deploy in an AWS environment.
Amazon Web Services Permissions	<p>The Amazon Web Services account used to deploy NETSCOUT Smart Data solutions must have appropriate permissions granted. The simplest way to do this is to grant the AdministratorAccess policy. However, if granting administrator access is not acceptable in your environment, assign the following policies to the account used to deploy NETSCOUT components:</p> <ul style="list-style-type: none">• Assign the built-in AmazonEC2FullAccess policy.• Create a custom policy with a permission for Full access to the CloudFormation service and assign it. <p>It's easiest to grant these permissions in the AWS Organizations visual editor. Note that granting these permissions complies with the "principle of least privilege" – these are the minimum permissions required to deploy the solution.</p> <p>Refer to "Security Notes" on page 20 for more information on best practices for the security of NETSCOUT Smart Data solutions.</p>
Static Private IP Address & License Information	<p>Bring Your Own License (BYOL) Deployments</p> <p>If you are deploying NETSCOUT Smart Data solutions using the BYOL model, you will need a static private IP address for Virtual nGeniusPULSE. You use this IP address to complete the product registration procedure and obtain the Serial Number and Password to be entered in the CloudFormation templates and deploy the BYOL AMIs from the AWS Marketplace. Refer to "Obtaining BYOL Licensing Information" on page 6 for details.</p> <p>Note: A static IP address is only needed for BYOL deployments.</p>
Existing AWS VPC	An existing AWS VPC with subnets for both Management and Monitoring.
Route Tables/Security Groups	Appropriate Route Tables and Security Groups for communication to/from nGeniusPULSE.
NTP Server Access	Access to an NTP Server for accurate timestamps in NETSCOUT analysis. NETSCOUT recommends using Amazon Time Sync Services . Note that NTP is enabled by default.
Access to Marketplace Images	You must have access to the NETSCOUT nGeniusPULSE AMI images in the AWS Marketplace in the AWS region you are using.
SSH Key Pair	<p>You must have a key pair for SSH access to deployed AMIs. You can create or import the key pair in AWS using these instructions.</p> <p>Note: SSH key pairs are created in AWS:</p> <ul style="list-style-type: none">• Public keys are stored in AWS, are not confidential, and are protected at the account level.• Private keys are stored by the user and are their responsibility to protect.

NETSCOUT nGeniusPULSE Solution is Supported in All AWS Regions

The NETSCOUT nGeniusPULSE solution is supported in all Amazon Web Services regions.

Skills and Specialized Knowledge Recommendations

Table 2 summarizes recommended skills and specialized knowledge for deployment of the NETSCOUT Smart Data solution:

Table 2 Skills and Specialized Knowledge Recommendations

AWS Component	Description
AWS Core Services	<ul style="list-style-type: none">• Understanding of EC2 Core services, including Marketplace.• Understanding of EC2 backup, snapshot, and restore processes.• High level understanding of AWS networking services, including VPCs, Subnets, Route Tables, Elastic/Public IP addresses, and Security Group.
AWS CloudFormation	<ul style="list-style-type: none">• Able to launch a Stack from a predefined CloudFormation Template.• Optional – Understanding of YAML.
AWS IAM	<ul style="list-style-type: none">• Able to attach AWS Managed IAM Policies to an IAM User running the deployment, either directly or via a Group.
Tools for AWS	<ul style="list-style-type: none">• Able to write scripts for regular maintenance of the EC2. There are multiple tools for scripting available, including AWS Command Line Tools and AWS SDKs. You can see a list of all supported Tools for Amazon Web Services here.

About Pricing and Costs

The NETSCOUT site on the AWS Marketplace provides helpful tools that let you estimate the costs of using NETSCOUT Smart Data solutions with different configuration choices. After navigating to one of the NETSCOUT solutions available on the AWS Marketplace, click on the **Pricing** tab and fill out the fields to estimate your costs. Keep in mind that your usage and costs may vary from the estimate depending on actual usage. In addition, Support is included as part of the pricing on the page referenced above.

About Service Quotas

AWS accounts typically have a variety of service quotas assigned that limit the quantity of resources a given account is allowed to consume. Keep in mind that the NETSCOUT nGeniusPULSE solution consumes resources from the following AWS service quotas:

- Amazon EC2
- Amazon EBS
- Amazon VPC
- Network interfaces, security groups, elastic IP addresses, subnets, routing tables, and traffic mirroring.

If your NETSCOUT deployment exceeds your AWS service limits, you can use [this procedure in the AWS documentation to increase your service quotas](#).

Licensing Model – BYOL

NETSCOUT nGeniusPULSE solutions are available in the AWS Marketplace in **BYOL** (Bring Your Own License) deployments for both **Commercial** and **GovCloud** environments:

Table 3 summarizes the available CFT templates:

Table 3 Available CFT Templates

Deployment Type	Description	Available CFT Templates
BYOL	<ul style="list-style-type: none">• Purchase nGeniusPULSE Server license from NETSCOUT based on static IP add Install using license utility	<ul style="list-style-type: none">• Virtual nGeniusPULSE

About BYOL Licenses

This section describes licensing for deployments for the BYOL CFT template.

License Type: Permanent Licenses

Description: Each nGeniusPULSE server requires a separate BYOL license.

Obtaining BYOL Licensing Information

Use the following procedure to obtain the **Serial Numbers** and **Passwords** from the NETSCOUT registration site to enter in the BYOL CloudFormation Templates as part of the deployment for both Virtual nGeniusPULSE.


- 1 When you purchase Virtual nGeniusPULSE from NETSCOUT, you receive a registration form that includes a registration key. Locate this form.
- 2 Open a web browser and navigate to <https://my.netscout.com/mcp/Pages/default.aspx>.
- 3 Navigate to **Licensing & Downloads** and follow the instructions there to enter your registration key. You will also enter the static, private IP address to be used for Virtual nGeniusPULSE in the AWS public cloud.
- 4 When you complete the registration procedure, you receive both a serial number and a password (license key). Print the screen that contains this information. You will enter these values in the CloudFormation templates when you deploy the Virtual nGeniusPULSE.

Deploying nGeniusPulse from AWS Marketplace

(25-30 Minutes)

This section describes how to deploy the nGeniusPulse server from the NETSCOUT site in the AWS Marketplace:

Note: Do not install or configure nGeniusPULSE as the **root** user.

- 1 Search the Amazon Marketplace for NETSCOUT.
The Amazon Marketplace shows the available NETSCOUT product offerings.
- 2 Click the entry for nGeniusPulse server.
- 3 Click the  button.
- 4 Accept the Terms and Conditions.
- 5 The **Fulfillment option** dropdown is already set to the product you selected in Step 2.
- 6 Use the **Software Version** dropdown to select the version of the selected template to deploy.
- 7 Use the **Region** dropdown to specify the Availability Zone where the software should be deployed.
- 8 Click **Continue to Launch** to continue.
- 9 Review the configuration details in the Launch page. Set the **Choose Action** dropdown to **Launch CloudFormation** when you are ready to launch.
- 10 Click **Launch** when ready to continue.

The **Create Stack** wizard appears with the listed options preconfigured for the selected template:

- **Prerequisite - Prepare template** is set to **Choose an existing template**.
- **Specify template** is set to **Amazon S3 URL** with the URL field pointing to the NETSCOUT CFT for your selected option.

Figure 3 shows the options for an nGeniusPulse server template:

Create stack

Prerequisite - Prepare template
You can also create a template by scanning your existing resources in the [IaC generator](#).

Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

Choose an existing template
Upload or choose an existing template.

Build from Infrastructure Composer
Create a template using a visual builder.

Specify template Info
This [GitHub repository](#) contains sample CloudFormation templates that can help you get started on new infrastructure projects. [Learn more](#)

Template source
Selecting a template generates an Amazon S3 URL where it will be stored. A template is a JSON or YAML file that describes your stack's resources and properties.

Amazon S3 URL
Provide an Amazon S3 URL to your template.

Upload a template file
Upload your template directly to the console.

Sync from Git
Sync a template from your Git repository.

Amazon S3 URL

Amazon S3 template URL

S3 URL: `https://netscout-cloudformation.s3.us-east-1.amazonaws.com/test/BYOL-nGP.cf.yaml` [View in Infrastructure Co](#)

[Cancel](#)

Figure 3 Sample Create Stack Options for nGeniusPULSE Server

11 Click **Next** to continue.

The **Specify Stack Details** screen appears, as illustrated in [Figure 4](#).

Specify stack details

Step 1 Create stack
Step 2 **Specify stack details**
Step 3 Configure stack options
Step 4 Review and create

Provide a stack name

Stack name

Enter a stack name

Stack name must contain only letters (a-z, A-Z), numbers (0-9), and hyphens (-) and start with a letter. Max 128 characters. Character count: 0/128.

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

General Configuration

Availability Zone

Availability Zone

us-east-1b

Key Name

Name of an existing EC2 KeyPair to enable SSH access to the instances

Select AWS::EC2::KeyPair::KeyName

Virtual nGeniusPULSE Configuration

Virtual nGeniusPULSE Name

AWS resource name for this Virtual nGeniusPULSE instance (2 to 240 characters)

Virtual nGeniusPULSE

Virtual nGeniusPULSE Host Name

Linux hostname for the virtual machine (must be between 3 and 57 characters and contain lowercase letters, numbers, and internal hyphens only)

ngp

Virtual nGeniusPULSE Instance Type

Virtual nGeniusPULSE Instance Type. Refer to <https://aws.amazon.com/ec2/instance-types/> for details on instance types.

m6i.4xlarge

Virtual nGeniusPULSE IP Address

Private IP Address of the virtual nGeniusPULSE within the management subnet (e.g., 192.168.100.100)

Enter String

Virtual nGeniusPULSE Enable Public IP Address

Enable Public IP Address of the virtual nGeniusPULSE

true

Virtual nGeniusPULSE Volume Size

Size of the Virtual nGeniusPULSE dbPULSE storage volume in GB

500

Virtual nGeniusPULSE Volume Encryption

Encrypt the Virtual nGeniusPULSE dbPULSE storage volume

false

Virtual nGeniusPULSE Additional Volume Size

Size of the Virtual nGeniusPULSE dbPULSE additional storage volume in GB

10

Virtual nGeniusPULSE Additional Volume Encryption

Encrypt the Virtual nGeniusPULSE dbPULSE additional storage volume

false

Network

VPC Identifier

VpcId of your existing Virtual Private Cloud (VPC)

Select AWS::EC2::VPC::Id

Management Subnet

Subnet ID for the Management Interface

Select AWS::EC2::Subnet::Id

Access Location

IPv4 CIDR address range that can SSH access the Virtual nGeniusPULSE Console (e.g., 10.0.1.0/24)

Enter String

Network Security Groups

Virtual nGeniusPULSE Management Security Group ID

Virtual nGeniusPULSE management security group ID

CREATE

AWS Marketplace Parameters

Virtual nGeniusPULSE disk image ID

Virtual nGeniusPULSE 3.14.0-0.306.0 fPC build 1

ami-0619dcd2bfc3954c

Figure 4 Stack Details for nGeniusPULSE Server

12 Use the information in Table 4 to supply the stack details for nGeniusPulse server:

Table 4 Stack Details for nGeniusPulse Server

Parameter	Description
Stack name	Provide a unique name for this stack.
General Configuration	
Availability Zone	Select an AWS Availability Zone to be used for the deployment from the dropdown list. The list includes the Availability Zones accessible from your account
Key Name	Select an existing keypair from the dropdown to be used for access to the instance once deployed. You can review your existing keypairs in Network & Security > Key Pairs from the EC2 Dashboard.
Virtual nGeniusPULSE Configuration	
Virtual nGeniusPulse Name	Supply a resource name for the nGeniusPulse instance. This is the name by which the nGeniusPulse instance will be known in AWS displays.
Virtual nGeniusPulse Host Name	Supply a hostname for the Linux virtual machine on which the Virtual nGeniusPulse software is installed.
Virtual nGeniusPulse Instance Type	Choose an AWS Instance Type for the Virtual nGeniusPulse deployment from the dropdown list. Virtual nGeniusPulse supports only the m6i.4xlarge instance type. NOTE: Refer to https://aws.amazon.com/ec2/instance-types for details on instance types.
Virtual nGeniusPulse IP Address	Supply a static, private IP address to be used as the private IP address of nGeniusPULSE's management port. The address you supply must be in the Management Subnet specified further down in the stack details. Note: The CloudFormation template only supports IPv4 addresses in this release. Contact NETSCOUT for assistance if you require IPV6 support.
Virtual nGeniusPULSE Enable Public IP Address	Specify whether AWS should provision a public, Internet-accessible public IP address for Virtual nGeniusPULSE.
Virtual nGeniusPulse Volume Size	Specify the size of the Virtual nGeniusPulse database (dbPULSE) in GB. The default value is 500 GB.
Virtual nGeniusPulse Volume Encryption	Specify whether the nGeniusPulse database (dbPULSE) should be encrypted. By default, it is not.
Virtual nGeniusPULSE Additional Volume Size	Specify the size of the additional nGeniusPULSE storage volume.
Virtual nGeniusPulse Additional Volume Encryption	Specify whether the additional nGeniusPULSE storage volume should be encrypted. By default, it is not.
Network	
VPC Identifier	Use the dropdown to select an existing VPC for the deployment. If you have many VPCs associated with your account, you can type an entry in the field to narrow the results to matching IDs or name tag values.

Table 4 Stack Details for nGeniusPulse Server

Parameter	Description
Management Subnet	Use the dropdown list to select an existing subnet for Virtual nGeniusPULSE's management traffic. The dropdown lists the subnets already provisioned for your account. If you have many subnets associated with your account, you can type an entry in the field to narrow the results to matching IDs or name tag values.
Access Location	Use this field to limit the range of IP addresses from which the deployed instance will accept SSH connections. This field is mandatory. However, if you want to allow SSH connections from any location, you can enter a value of 0.0.0.0/0 . You can edit the Security Group settings later on to change the IP addresses for which access is allowed. Refer to Working with Security Groups in the AWS documentation for details.
Network Security Groups	
Virtual nGeniusPULSE Management Security Group ID	Use this field to assign the Virtual nGeniusPULSE Management interface (eth0) to a Security Group: <ul style="list-style-type: none"> • If you leave this option set to CREATE (the default), the template automatically assigns the Management interface to a Security Group with the necessary permissions and open ports to allow communications with other NETSCOUT solutions. Ports are opened in accordance with the principle of <i>least privilege</i> – only the ports required for successful communications are opened. • The CREATE keyword should be used the first time a NETSCOUT product is deployed in a VPC. If additional NETSCOUT products are deployed that need to communicate with each other, the CREATE keyword should be replaced with the security group ID from the original deployment. Security group IDs can be found in the Outputs tab of the original CloudFormation stack deployment. • You can also supply the name of an existing Security Group. If you use an existing Security Group, you must open the necessary ports manually using the information in "Security Group Details" on page 16.
AWS Marketplace Parameters	
Virtual nGeniusPULSE disk image ID	This field reports the disk image ID of the AMI file to be used for the deployment.

13 When you have finished configuring the stack details, click **Next** to continue.

14 The **Configure stack options** page appears, allowing you to configure the standard CloudFormation Stack settings listed below. These are all optional; none are required. Use the links below to learn more about these AWS options.

- [Tags \(key-value pairs\)](#)
- [Permissions](#)
- [Stack failure options](#)
- [Additional settings](#)

When you have finished setting stack options, click **Next** to continue.

15 The **Review and create** page displays a summary of the settings for the new stack. Review the settings and use the **Previous** button to correct if necessary. When you are satisfied with your settings, click **Submit** to launch the new instance(s).

The Stack Wizard begins to create the requested resources (Figure 5).

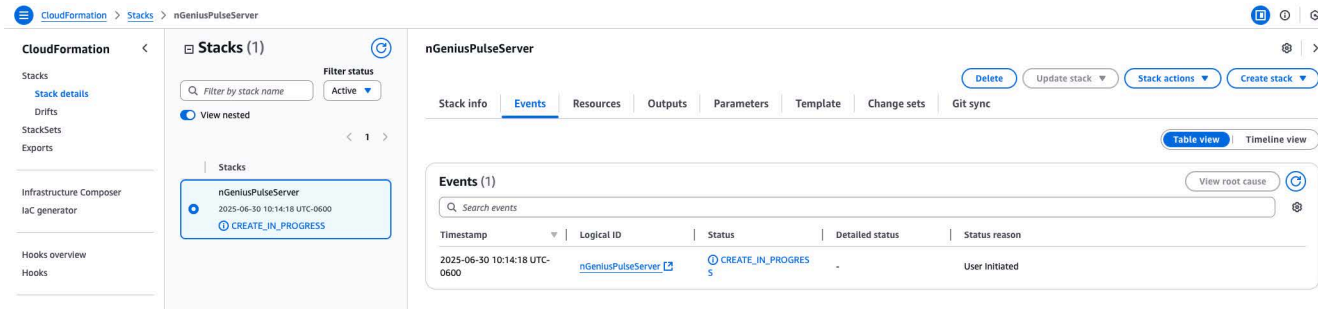


Figure 5 Stack Creation in Progress

When stack creation is complete, the nGeniusPULSE instance appears in the EC2 Management Console's **Instances** list (Figure 6).

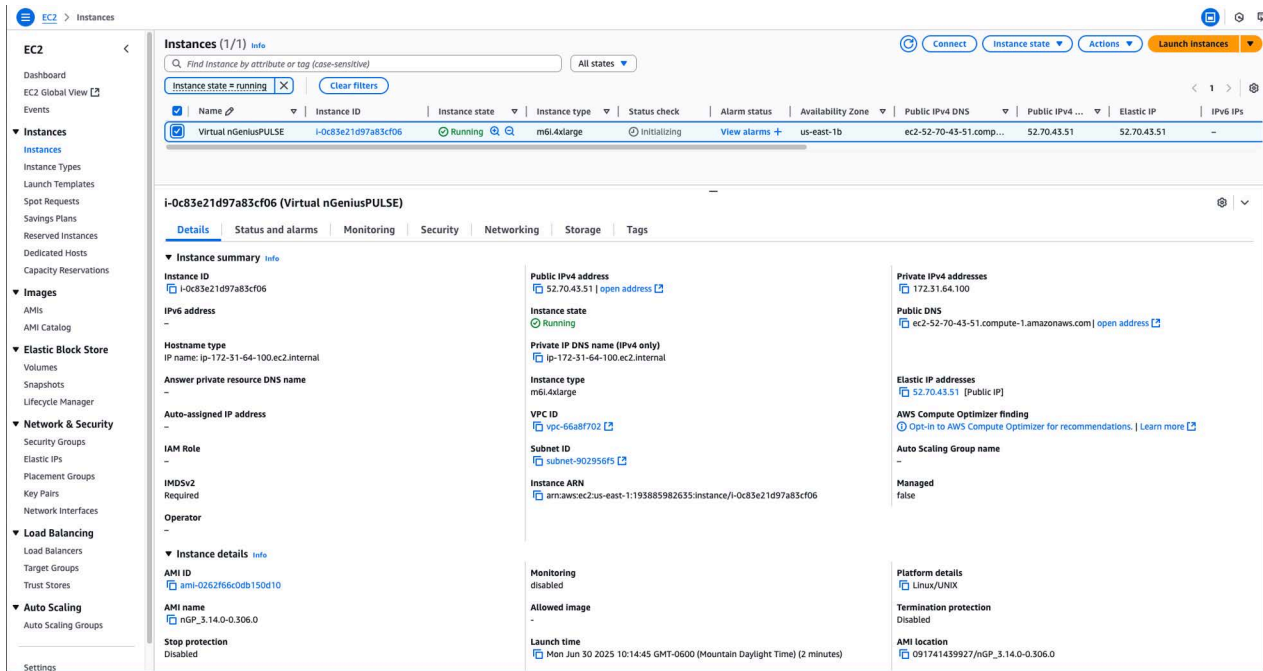


Figure 6 Newly Created Instance

Important – nGeniusPULSE Startup Takes Time!

Once the nGeniusPULSE stack is successfully created, it takes approximately 12 minutes for all processes to start. Make sure you wait at least 12 minutes before attempting to connect to the nGeniusPULSE user interface in the next section.

Logging in to the Virtual nGeniusPULSE User Interface

(5-10 Minutes)

If you opted to assign a public IP address to nGeniusPULSE, you can connect to the instance from the Internet using the instructions in this section:

Important: Make sure you wait at least 12 minutes after the nGeniusPULSE stack is created before connecting to the user interface. This allows sufficient time for all nGeniusPULSE services to start.

- 1 Navigate to the **Instances** list in the EC2 Console and copy the **Public IPv4 address** assigned to nGeniusPULSE (Figure 7):

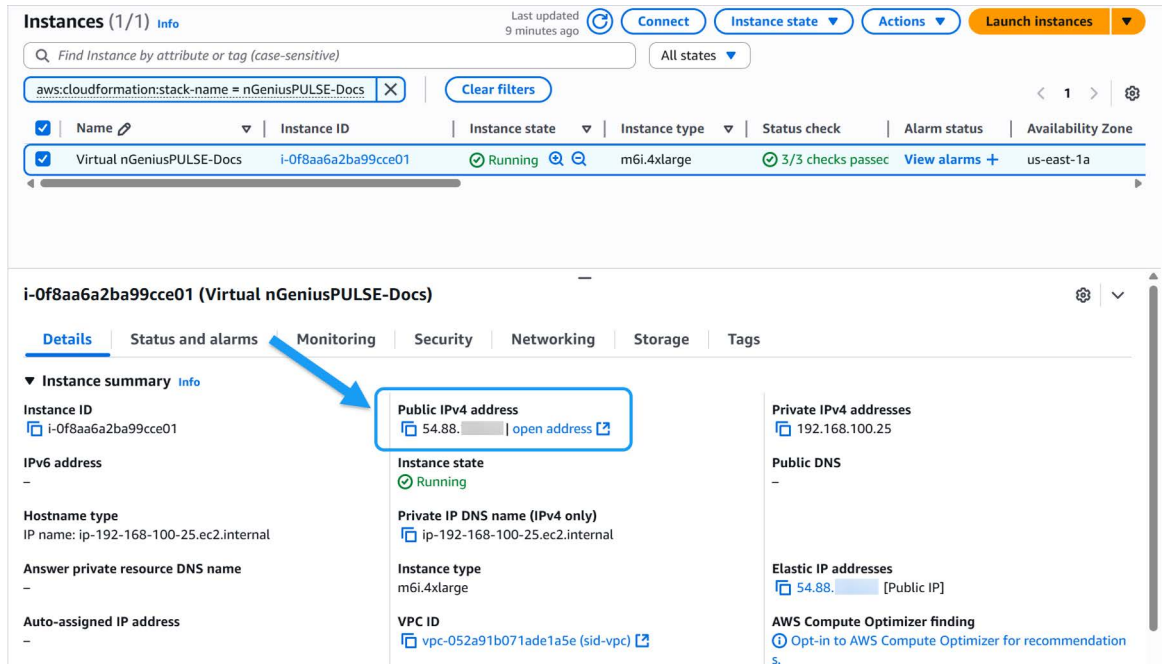



Figure 7 Copying the Public IP Address of the Deployed Instance

- 2 Open a web browser and type in the following URL, substituting the IP address you copied from the Instances list in the EC2 Console:

`https://<IP Address from Instances List>/console/`

For example:



- 3 Click through any security warnings that appear and proceed to the nGeniusPULSE server's login page.
- 4 Log in to nGeniusPULSE. The default login credentials for Virtual nGeniusPULSE are as follows:

- Username: **administrator**
- Password: **Instance ID**

The password is the Instance ID of the deployed Virtual nGeniusPULSE virtual machine. You can copy this from the Instances list in the EC2 Console as illustrated in Figure 8.

It's best to copy and paste the password to prevent the possibility of a typo. Using the dedicated **Copy** button in the **Details** tab is an easy way to get it right, as illustrated in [Figure 8](#)

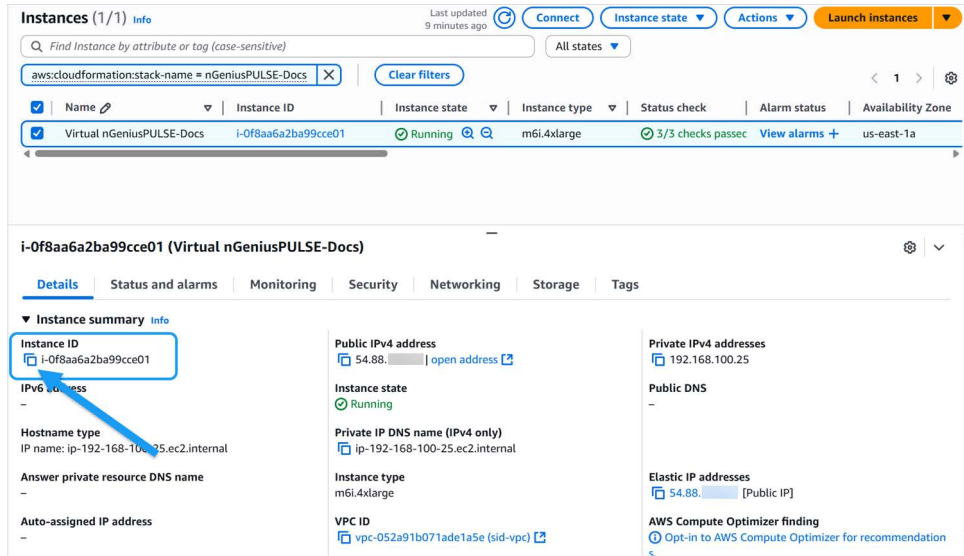


Figure 8 Copying the Instance ID for the Default Password

5 The first time you log in, nGeniusPULSE requires you to change your password ([Figure 9](#)):

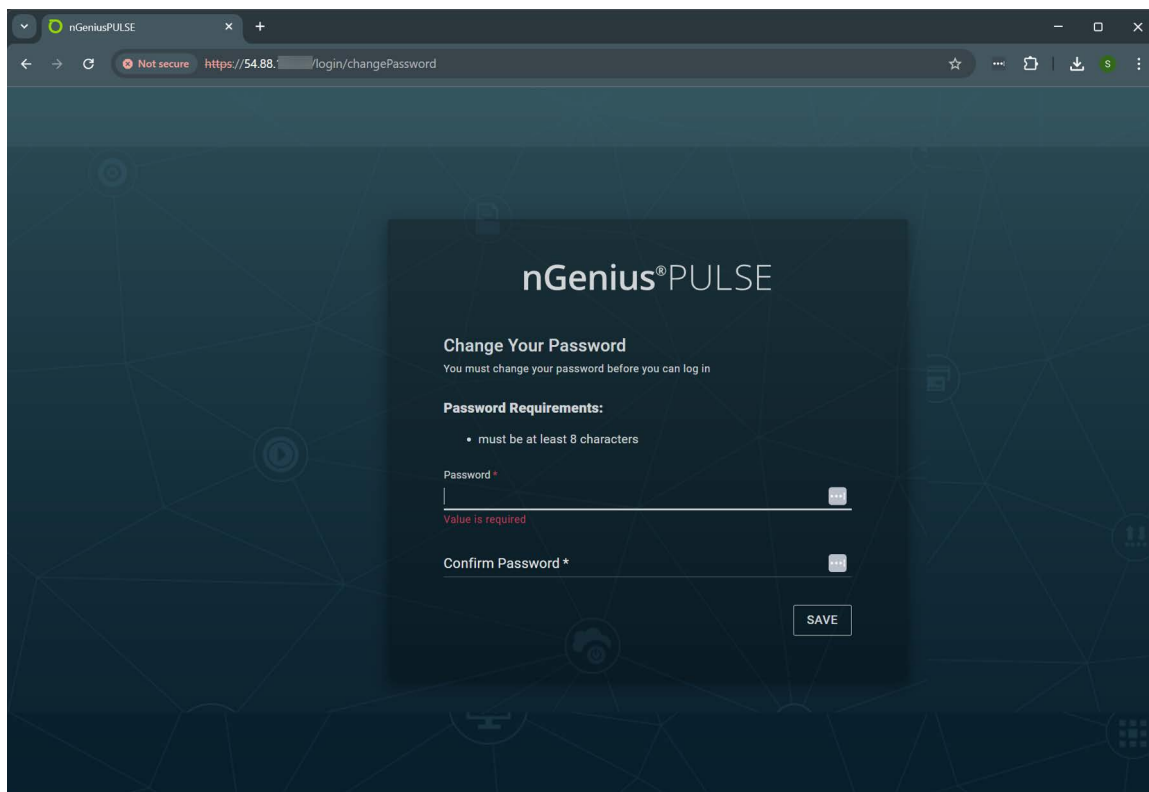


Figure 9 First Login to nGeniusPULSE

6 Supply a new password meeting the requirements and click **Save**.

After a successful login, nGeniusPULSE appears, as shown in [Figure 10](#). From here, you can use the comprehensive online help available in the console to get started with nGeniusPULSE.

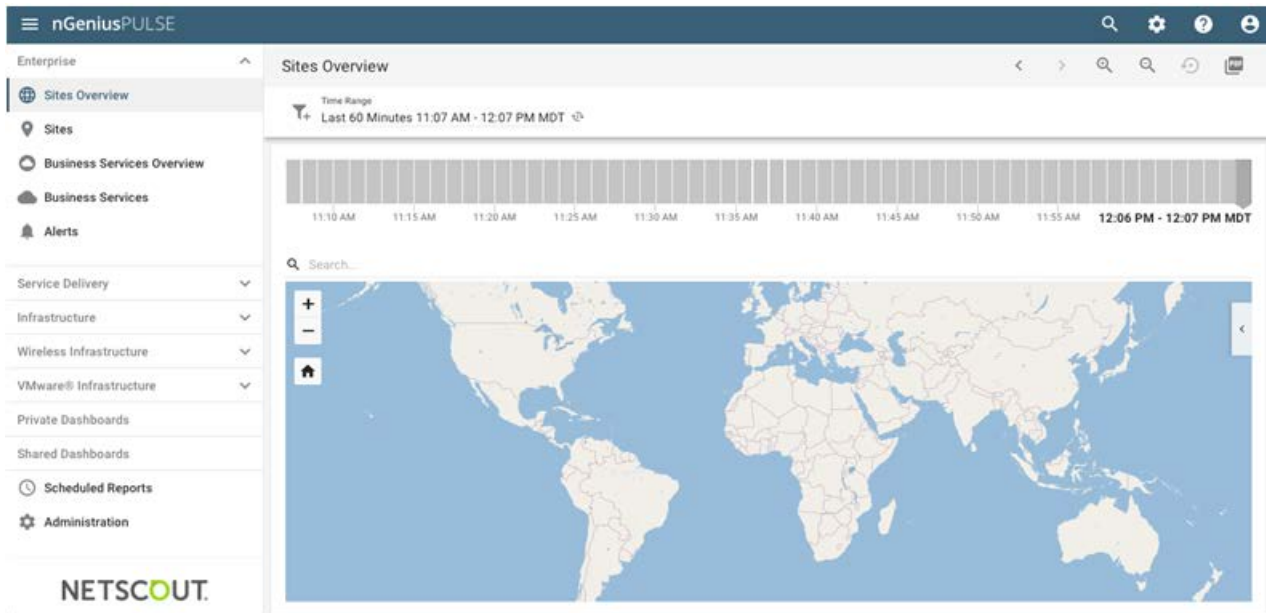


Figure 10 Virtual nGeniusPULSE Server Deployed in AWS

Note: Refer to "[Connecting to Instances \(5-10 Minutes\)](#)" on [page 17](#) for information on opening an SSH connection to the operating system of the new instances.

Security Group Details

The nGeniusPULSE template provides the option of creating the AWS Security Group summarized in [Table 5](#) for the Virtual nGeniusPULSE management interface. This section describes the ports opened by this Security Group.

If you did not create a Security Group as part of the template, you can also use the information in this section to open the necessary ports for nGeniusPULSE management communications in your own Security Group:

- ["About the Virtual nGeniusPULSE Mgmt Security Group" on page 16](#)
- ["Instance Type Recommendations" on page 17](#)

[Table 5](#) lists the default Security Groups created by the NETSCOUT CFT templates. Following the table, [Figure 6](#) illustrates sample creation of these groups.

Table 5 nGeniusPULSE Management Security Group

Name	Group Name	Instance	Interface
sg-vnGP-mgmt	Virtual nGeniusPULSE Management Security Group ID	Virtual nGeniusPULSE eth0	eth0




Figure 11 NETSCOUT Security Groups

About the Virtual nGeniusPULSE Mgmt Security Group

The Virtual nGeniusPULSE Security Group allows packets and selected ports from interfaces in the groups as summarized in [Table 6](#).

Table 6 Traffic Allowed by Virtual nGeniusPULSE Mgmt Security Group

Description	Protocol	Port Range
HTTPS from interfaces in Security Groups.	TCP	443
SSH, as configured by Access Location parameter in stack details.	SSH	22

Instance Type Recommendations

The CloudFormation templates for the NETSCOUT nGeniusPULSE solution requires the m6i.4xlarge Instance Type for the Virtual nGeniusPULSE virtual machine.

Note: Refer to <https://aws.amazon.com/ec2/instance-types> for details on instance types.

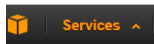
Table 7 describes the instance type for Virtual nGeniusPULSE:

Table 7 Instance Type Recommendations

Instance Type	vCPUs	Memory	Supported nPoints	Type
m6i.4xlarge Recommended for high usage environments.	16	64 GB	2000	nGeniusPULSE Server

Connecting to Instances (5-10 Minutes)

Connect to the operating system of NETSCOUT instances using the key pair you selected as part of the CloudFormation template as follows:

- 1 Click the **Services** dropdown  in the AWS Management Console and select **Compute > EC2**.
- 2 Click the **Instances** entry in the left column.
- 3 Make sure the desired instance is selected.
- 4 Click the **Connect** button (Figure 12).




Figure 12 Connecting to the Instance

- 5 The **Connect To Your Instance** window provides guidance on using SSH to connect to the instance remotely, either using the Linux **ssh** command or a Windows client, such as PuTTY. Keep in mind the following:
 - You will need access to your private key file. The **Connect To Your Instance** window reminds you of the name of the private key file you associated with the instance.
 - Your private key file must not be publicly viewable for SSH to work. You can use **chmod 400 netscout-keys** to make your private key file not publicly viewable.

- The **Connect To Your instance** window shows you the IP address you should use to connect to your instance along with the correct SSH syntax. For example, in [Figure 13](#), we can use the following SSH command to log in to the default **cloud-user** account provided with NETSCOUT AMIs:

```
$ ssh -i "netscout-keys.pem" cloud-user@34.203.23.249
```

Connect Info

Connect to an Instance using the browser-based client.

EC2 Instance Connect | Session Manager | **SSH client** | EC2 serial console

Instance ID
I-0c83e21d97a83cf06 (Virtual nGeniusPULSE)

1. Open an SSH client.
2. Locate your private key file. The key used to launch this instance is netscout-keys.pem
3. Run this command, if necessary, to ensure your key is not publicly viewable.

```
chmod 400 "netscout-keys.pem"
```
4. Connect to your instance using its Public DNS:

```
ec2-52-70-43-51.compute-1.amazonaws.com
```

Example:

```
ssh -i "netscout-keys.pem" root@ec2-52-70-43-51.compute-1.amazonaws.com
```

Note: In most cases, the guessed username is correct. However, read your AMI usage instructions to check if the AMI owner has changed the default AMI username.

Figure 13 The Connect To Your Instance Window

- 6 Click **Close** on the Connect To Your Instance window.
- 7 Open a terminal window and use the **ssh** command to connect to the NETSCOUT instance:

```
$ ssh -i "<keyfile.pem>" cloud-user@<NETSCOUT_IP>
```

Operational Guidance

This section provides information on assessing and monitoring the health of the NETSCOUT nGeniusPULSE for AWS solution. If you experience operational or performance issues not covered by this section, contact your NETSCOUT Support representative using the information in ["Contacting NETSCOUT SYSTEMS, INC." on page iv](#).

- ["Snapshot and Backup Procedures" on page 19](#)
- ["Routine Maintenance" on page 20](#)

Snapshot and Backup Procedures

This section describes how to perform routine snapshot and backup procedures of Virtual nGeniusPULSE using EBS Snapshots and EC2 Images as part of a standard Disaster Recovery Plan. NETSCOUT recommends that you use an Automation process based on one of the following tools for creating snapshots and images:

- [CreatelImage API](#)
- [CreateSnapshot API](#)
- [AWS Data Life Cycle Manager](#)

In general, NETSCOUT recommends that you back up custom AMIs instead of creating snapshots for easier orchestration of a disaster recovery. However, snapshots are also acceptable.

Backing Up nGeniusPULSE

Use the AWS **CreatelImage** API to back up an AMI image and the AWS **CreateSnapshot** API to take a snapshot. Refer to the following AWS documentation for details on using these tools:

- **CreatelImage** – [Creating an Amazon EBS-Backed Linux AMI](#)
- **CreateSnapshot** – [Creating an Amazon EBS Snapshot](#)

In order to preserve file system integrity, you should avoid running **CreatelImage** with the **NoReboot** flag or **CreateSnapshot** on a live EC2 or EBS instance. NETSCOUT recommends the following procedure to preserve file system integrity:

- 1** Shut down the Virtual nGeniusPULSE EC2s targeted for backup.
- 2** Execute the necessary **CreatelImage/CreateSnapshot** API calls.
- 3** Power the EC2s back on after **CreatelImage/CreateSnapshot** has started.

Although these steps will cause a small service disruption, they do ensure file system integrity is maintained. Keep in mind that while **CreatelImage** or **CreateSnapshot** processes are running, the performance of Virtual nGeniusPULSE performance may be degraded.

If it is not possible to shut down the target EC2 instances, NETSCOUT recommends at least stopping Virtual nGeniusPULSE processes before running **CreatelImage/CreateSnapshot**. In addition, NETSCOUT recommends running **CreatelImage/CreateSnapshot** during a time of day with the least network traffic.

Snapshot Examples by Target RPO

Backup/snapshot options for nGeniusPULSE are based on the target Recovery Point Objectives (RPO) in [Table 8](#):

Table 8 Backup/Snapshot Techniques by Target RPO

RPO Target	Backup/Snapshot Technique
12-24 Hours	Life Cycle Manager
12 Hours or Less	CreateImage or CreateSnapshot APIs
1 Hour or Less	Direct backup of nGeniusPULSE database using AWS S3 Sync in addition to CreateImage/CreateSnapshot APIs.

For successful Disaster Recovery after an AWS Region failure, NETSCOUT recommends the following regular backups/snapshots:

- 1 Multi-Region backup of nGeniusPULSE database by [Cross-Region Replication](#).
- 2 Multi-Region copies of AMIs by [CopyImage API](#).
- 3 Multi-Region copies of Snapshots by [CopySnapshot API](#).

Because each of these steps can be executed programmatically using [AWS Tools](#), you can extend them into your existing Backup Automation strategy using the tool of your choice.

Routine Maintenance

NETSCOUT recommends that you follow industry standard best practices for security with the NETSCOUT nGeniusPULSE solution, including key rotation and certificate maintenance.

[Refer to this article for instructions on changing access keys](#) on a regular schedule.

In addition, it is important that you apply patches and upgrades for NETSCOUT components as they become available. The Virtual nGeniusPULSE and physical nGeniusPULSE upgrade procedures are identical. Refer to the *nGeniusPULSE Administrator Guide* for details on the procedure

Security Notes

[Table 9](#) lists and describes some best practices for the security of NETSCOUT Smart Data solutions:

Note: Refer to the [AWS Identify and Access Management User Guide](#) for detailed information on security best practices in AWS.

Table 9 Security Notes

Topic	Notes
Permissions and Roles	The minimum permissions for the account used to deploy NETSCOUT components are as follows: <ul style="list-style-type: none">• Assign the built-in AmazonEC2FullAccess policy.• Create a custom policy with a permission for Full access to the CloudFormation service and assign it.
Key Rotation	NETSCOUT recommends that you follow industry standard best practices for the rotation of SSH keys used to access NETSCOUT components.

Table 9 Security Notes

Topic	Notes
Security Groups	NETSCOUT recommends that you use AWS Security Groups and VPC access control lists to limit access to the networks where NETSCOUT components are deployed. Avoid using “open” security groups and consider limiting access to certain IP addresses or ranges.
Data Encryption	You can optionally encrypt the storage volume by enabling the VolumeEncrypt option in the CloudFormation template during deployment. Doing so enhances the security of stored network data.
CloudTrail	NETSCOUT recommends that you enable and use the AWS CloudTrail feature for enhanced logging capabilities, including flow logs and access logs. In addition, you can use nGeniusPULSE logging features to monitor usage of the solution. Refer to “Working with Activity Logs” in the nGeniusPULSE online help for details.
Resource Tagging	Resources deployed as part of NETSCOUT Smart Data solutions are typically tagged with the NETSCOUT name to allow easy monitoring of the usage of its components.

Disaster Recovery

This section discusses Disaster Recovery procedures for NETSCOUT solutions. Separate sections discuss Disaster Recovery for both AWS Availability Zone and Region failures:

- [Disaster Recovery: Key Concepts on page 22](#)
- [Sample Disaster Recovery Plans on page 22](#)
 - [Availability Zone Recovery on page 22](#)
 - [Region Recovery on page 23](#)

Disaster Recovery: Key Concepts

Successful recovery of the NETSCOUT solutions from a failed hosting environment depends on preservation of the following:

- Private IP addresses of nGeniusPULSE.

It's crucial to prepare a Disaster Recovery Plan in such a way that these items are preserved. Note that the procedures described in [Snapshot and Backup Procedures on page 19](#) all ensure that the private IP addresses are retained for Disaster Recovery.

Additional Recommendations

- NETSCOUT recommends operating nGeniusPULSE in its own VPC. Because private IP addresses are unique to a [subnet in a VPC](#) and a subnet cannot span multiple Availability Zones, this results in an improved Recovery Time Objective (RTO).
- Disaster Recovery also benefits from a carefully designed [VPC Peering](#) and [Transit Gateway](#) implementation in a multi-VPC architecture. Keep in mind that successful VPC Peering requires that VPC CIDRs do not overlap and that the associated [Route Tables](#) are configured properly.

Sample Disaster Recovery Plans

The sample Disaster Recovery plans in this section can be used as a reference when creating your own plans for recovery from both AWS Availability Zone and Region failures. Because the steps in these sample plans can all be executed programmatically using [AWS Tools](#), you can extend them into your existing Disaster Recovery Automation strategy using the tool of your choice.

Recovery Plan Assumptions

The examples in the sections below assume the following scenario:

- The CIDR for the VPC where nGeniusPULSE are installed is **10.10.10.0/23 (VPC1 and VPC3)**
- **VPC1** is where the assumed disaster takes place (either an Availability Zone or Region failure).
- **VPC2** hosts the Application Stack and resides in healthy Availability Zones.
- **VPC3** will host recovered nGeniusPULSE in a healthy Availability Zone.

Availability Zone Recovery

An availability zone recovery is summarized in [Table 10](#).

- 1 Create a new **VPC3** with Management and Monitoring Subnets in a healthy AZ within the same AWS Region while maintaining the IP Addresses from VPC1, as summarized in [Table 10](#).

Table 10 Pre- and Post-Disaster Network Configuration

Network Element	Failed (VPC1)	Recovered (VPC3)
VPC CIDR	10.10.10.0/23	10.10.10.0/23
Management Subnet CIDR	10.10.10.0/24	10.10.10.0/24
nGeniusPULSE Management Network Interface IP	10.10.10.10	10.10.10.10

- 2 Remove the VPC Peering between **VPC1** and **VPC2**.
- 3 Remove the Peering Route Table entry for **VPC1**'s CIDR from **VPC2**'s **Route Table 2**.
- 4 [Create VPC Peering](#) between **VPC2** and **VPC3** (for example, **pcx-123abc456def**).
- 5 Create [Route Table](#) entries for the following CIDR blocks from [Table 10](#):
 - a **VPC2**'s CIDR of **10.0.0.0/16** in **VPC3**'s **Route Table 3** via the **pcx-123abc456def** VPC Peering.
 - b **VPC3**'s CIDR of **10.10.10.0/23** in **VPC2**'s **Route Table 2** via the **pcx-123abc456def** VPC Peering.
- 6 [Disassociate the Elastic IP addresses](#) from the Virtual nGeniusPULSE Management interfaces
- 7 Replicate the Security Groups from **VPC1** to **VPC3**. We will use these for Management interfaces later.
- 8 Create an nGeniusPULSE Management Interface in the Management Subnet (**10.10.10.0/24**) in **VPC3**.
 - a Assign the Private IP address of **10.10.10.10** to the nGeniusPULSE Management interface, identical to what it was in **VPC1**.
 - b [Associate](#) the existing nGeniusPULSE Elastic IP address (**x1.x2.x3.x4**) to the new interface in **VPC3**.
- 9 Assign the matching Security Groups from **VPC1** (created in [Step 7](#)) to the new network interfaces you just created in VPC3.
- 10 Launch a new nGeniusPULSE Instance in the new Management Subnet of **VPC3** using the Custom nGeniusPULSE AMI created in "[Snapshot and Backup Procedures](#)" on [page 19](#).
 - a Add the Management Network Interface created in [Step 8](#) to the new nGeniusPULSE instance in VPC3.
- 11 If you created an S3 Sync Backup of the nGeniusPULSE Database as described in "[Backing Up nGeniusPULSE](#)" on [page 19](#), use the following steps to restore the database once nGeniusPULSE is up and running:
 - a Open an SSH connection to the nGeniusPULSE instance.
 - b Stop nGeniusPULSE processes by issuing the **ngp-stop** command:
 - c Restore the nGeniusPULSE database from the S3 bucket to EC2 using S3 Sync.
 - d Start nGeniusPULSE processes by issuing the **ngp-start** command.

Region Recovery

A region recovery is summarized in [Table 11](#).

- 1 Create a new **VPC3** with Management and Monitoring Subnets in an AZ within a separate, healthy AWS Region while maintaining the IP Addresses from the previous configuration, as summarized in [Table 11](#).

Table 11 Pre and Post-Disaster Network Configuration

Network Element	Failed (VPC1)	Recovered (VPC3)
VPC CIDR	10.10.10.0/23	10.10.10.0/23
Management Subnet CIDR	10.10.10.0/24	10.10.10.0/24
nGeniusPULSE Management Network Interface IP	10.10.10.10	10.10.10.10
Monitoring Subnet CIDR	10.10.11.0/24	10.10.11.0/24

- 2 Remove the VPC Peering between **VPC1** and **VPC2**.
- 3 Remove the Peering Route Table entry for **VPC1**'s CIDR from **VPC2**'s **Route Table 2**.
- 4 [Create VPC Peering](#) between **VPC2** and **VPC3** (for example, **pcx-123xyz456abc**).
- 5 Create [Route Table](#) entries for the following CIDR blocks from [Table 11](#):
 - a **VPC2**'s CIDR of **10.0.0.0/16** in **VPC3**'s **Route Table 3** via the **pcx-123xyz456abc** VPC Peering.
 - b **VPC3**'s CIDR of **10.10.10.0/23** in **VPC2**'s **Route Table 2** via the **pcx-123xyz456abc** VPC Peering.
- 6 [Allocate Elastic IP addresses](#) for the nGeniusPULSE (**x5.x6.x7.x8**) Management interfaces in the same healthy Region of **VPC3**.
- 7 Replicate the Security Groups from **VPC1** to **VPC3**. We will use these for Management and Monitoring interfaces later.
- 8 Create an nGeniusPULSE Management Interface in the Management Subnet (**10.10.10.0/24**) in **VPC3**.
 - a Assign the Private IP address of **10.10.10.10** to the nGeniusPULSE Management interface, identical to what it was in **VPC1**.
 - b [Associate](#) the nGeniusPULSE Elastic IP address (**x5.x6.x7.x8**) to the new interface in **VPC3**.
- 9 Assign the matching Security Groups from **VPC1** (created in [Step 7](#)) to the new network interfaces you just created in VPC3.
- 10 Launch a new nGeniusPULSE Instance in the new Management Subnet of **VPC3** using the Custom nGeniusPULSE AMI created in "[Snapshot and Backup Procedures](#)" on [page 19](#).
 - a Add the Management Network Interface created in [Step 8](#) to the new nGeniusPULSE instance in VPC3.
- 11 [Update the Route 53 DNS entries](#) for the nGeniusPULSE Management interfaces with new Elastic IP addresses from [Step 8](#) and [Step 9](#) respectively.
- 12 If you created an S3 Sync Backup of the nGeniusPULSE Database as described in "[Backing Up nGeniusPULSE](#)" on [page 19](#), use the following steps to restore the database once nGeniusPULSE is up and running:
 - a Open an SSH connection to the nGeniusPULSE instance.
 - b Stop nGeniusPULSE processes by issuing the **ngp-stop** command.
 - c Restore the nGeniusPULSE database from the S3 bucket to EC2 using S3 Sync.
 - d Start nGeniusPULSE processes by issuing the **ngp-start** command.

NETSCOUT®

NETSCOUT SYSTEMS, INC.
310 Littleton Road
Westford, MA 01886-4105

Tel. 978-614-4000
888-999-5946

Fax 978-614-4004

E-mail info@netscout.com

Web www.netscout.com

© 2025 NETSCOUT SYSTEMS, INC.
All rights reserved.
733-2143 Rev. B