

NetScreen New Features

ScreenOS 3.0.1 Reference Guide

Version 3.0.1

P/N 093-0437-000

Rev. A

Copyright Notice

Copyright © 1998-2002 NetScreen Technologies, Inc. NetScreen Technologies, Inc., the NetScreen logo, NetScreen-5, NetScreen-5XP, NetScreen-10, NetScreen-25, NetScreen-50, NetScreen-100, NetScreen-204, NetScreen-208, NetScreen-500, NetScreen-1000, NetScreen-Global Manager, NetScreen-Global PRO, NetScreen-Global PRO Express, NetScreen-Remote, GigaScreen ASIC, and NetScreen ScreenOS are trademarks and NetScreen is a registered trademark of NetScreen Technologies, Inc. All other trademarks and registered trademarks are the property of their respective companies.

NetScreen Technologies, Inc.
350 Oakmead Parkway, Suite 500
Sunnyvale, CA 94085 U.S.A.
www.netscreen.com

FCC Statement

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 in a light commercial installation. This equipment generates, uses and can radiate radio frequency energy, and, if not installed and used in accordance with the instruction, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SPECIFICATIONS REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS. NO PART OF THIS DOCUMENT MAY BE REPRODUCED OR TRANSMITTED IN ANY FORM OR BY ANY MEANS, ELECTRONIC OR MECHANICAL, FOR ANY PURPOSE, WITHOUT RECEIVING WRITTEN PERMISSION FROM NETSCREEN TECHNOLOGIES INC.

Product License Agreement

PLEASE READ THIS LICENSE AGREEMENT (“AGREEMENTS”) CAREFULLY BEFORE USING THIS PRODUCT. BY INSTALLING AND OPERATING, YOU INDICATE YOUR ACCEPTANCE OF THE TERMS OF THIS LEGAL AND BINDING AGREEMENT AND ARE CONSENTING TO BE BOUND BY AND ARE BECOMING A PART TO THIS AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THIS AGREEMENT, DO NOT START THE INSTALLATION PROCESS.

1. **License Grant.** This is a license, not a sales agreement, between you, the end user, and NetScreen Technologies, Inc. (“NetScreen”). The term “Software” includes all NetScreen and third party Software provided to you with the NetScreen product, and includes any accompanying documentation, any updates and enhancements of the Software provided to you by NetScreen, at its option. NetScreen grants to you a non-transferable (except as provided in section 3 (“Transfer”) below), non-exclusive license to use the Software in accordance with the terms set forth in this License Agreement. The Software is “in use” on the product when it is loaded into temporary memory (i.e. RAM).

2. **Limitation on Use.** You may not attempt and if you are a corporation, you will use best efforts to prevent your employees and contractors from attempting to, (a) modify, translate, reverse engineer, decompile, disassemble, create, derivative works based

on, sublicense, or distribute the Software or the accompanying documentation; (b) rent or lease any rights in the Software or accompanying documentation in any form to any person; or (c) remove any proprietary notice, labels, or marks on the Software, documentation, and containers.

3. Transfer. You may transfer (not rent or lease) the Software to the end user on a permanent basis, provided that: (i) the end user receives a copy of this Agreement and agrees in writing to be bound by its terms and conditions, and (ii) you at all times comply with all applicable United States export control laws and regulations.

4. Proprietary Rights. All rights and title and interest in and to, and all intellectual property rights, including copyrights, to the software, and documentation, remain with NetScreen. You acknowledge that no title to the intellectual property in the Software is transferred to you and you will not acquire any rights to the Software except for the license as specifically set forth herein.

5. Term and Termination. The term of the license is for the duration of NetScreen's copyright in the Software. NetScreen may terminate this Agreement immediately without notice if you breach or fail to comply with any of the terms and conditions of this Agreement. You agree that, upon such termination, you will either destroy all copies of the documentation or return all materials to NetScreen. The provisions of this Agreement, other than the license granted in Section 1 ("License Grant") shall survive termination.

6. Limited Warranty. For a period of ninety (90) days after delivery to Customer, NetScreen will repair or replace any defective software product shipped to Customer, provided it is returned to NetScreen at Customer's expense within that period. NetScreen warrants to Customer that such product will substantially conform with NetScreen's published specifications for that product if properly used in accordance with the procedures described in documentation supplied by NetScreen. NetScreen's exclusive obligation with respect to non-conforming product shall be, at NetScreen's option, to replace the product or use commercially reasonable efforts to provide Customer with a correction of the defect, or to refund to customer the purchase price paid for the

unit. Defects in the product will be reported to NetScreen in a form and with supporting information reasonably requested by NetScreen to enable it to verify, diagnose, and correct the defect. For returned product, the customer shall notify NetScreen of any nonconforming product during the warranty period, obtain a return authorization for the nonconforming product, from NetScreen, and return the nonconforming product to NetScreen's factory of origin with a statement describing the nonconformance.

NOTWITHSTANDING ANYTHING HEREIN TO THE CONTRARY, THE FOREGOING IS CUSTOMER'S SOLE AND EXCLUSIVE REMEDY FOR BREACH OF WARRANTY BY NETSCREEN WITH RESPECT TO THE PRODUCT.

The warranties set forth above shall not apply to any Product or Hardware which has been modified, repaired or altered, except by NetScreen, or which has not been maintained in accordance with any handling or operating instructions supplied by NetScreen, or which has been subjected to unusual physical or electrical stress, misuse, abuse, negligence or accidents.

THE FOREGOING WARRANTIES ARE THE SOLE AND EXCLUSIVE WARRANTIES EXPRESS OR IMPLIED GIVEN BY NETSCREEN IN CONNECTION WITH THE PRODUCT AND HARDWARE, AND NETSCREEN DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. NETSCREEN DOES NOT PROMISE THAT THE PRODUCT IS ERROR-FREE OR WILL OPERATE WITHOUT INTERRUPTION.

7. Limitation of Liability. IN NO EVENT SHALL NETSCREEN OR ITS LICENSORS BE LIABLE UNDER ANY THEORY FOR ANY INDIRECT, INCIDENTAL, COLLATERAL, EXEMPLARY, CONSEQUENTIAL OR SPECIAL DAMAGES OR LOSSES SUFFERED BY YOU OR ANY THIRD PARTY, INCLUDING WITHOUT LIMITATION LOSS OF USE, PROFITS, GOODWILL, SAVINGS, LOSS OF DATA, DATA FILES OR PROGRAMS THAT MAY HAVE BEEN STORED BY ANY USER OF THE SOFTWARE. IN NO EVENT WILL NETSCREEN'S OR ITS LICENSORS' AGGREGATE LIABILITY CLAIM BY YOU, OR ANYONE CLAIMING THROUGH OR ON BEHALF OF YOU, EXCEED THE ACTUAL AMOUNT PAID BY YOU TO NETSCREEN FOR SOFTWARE. Some jurisdictions do not allow

the exclusions and limitations of incidental, consequential or special damages, so the above exclusions and limitations may not apply to you.

8. Export Law Assurance. You understand that the Software is subject to export control laws and regulations. **YOU MAY NOT DOWNLOAD OR OTHERWISE EXPORT OR RE-EXPORT THE SOFTWARE OR ANY UNDERLYING INFORMATION OR TECHNOLOGY EXCEPT IN FULL COMPLIANCE WITH ALL UNITED STATES AND OTHER APPLICABLE LAWS AND REGULATIONS.**

9. U.S. Government Restricted Rights. If this Product is being acquired by the U.S. Government, the Product and related documentation is commercial computer Product and documentation developed exclusively at private expense, and (a) if acquired by or on behalf of civilian agency, shall be subject to the terms of this computer Software, and (b) if acquired by or on behalf of units of the Department of Defense (“DoD”) shall be subject to terms of this commercial computer Software license Supplement and its successors.

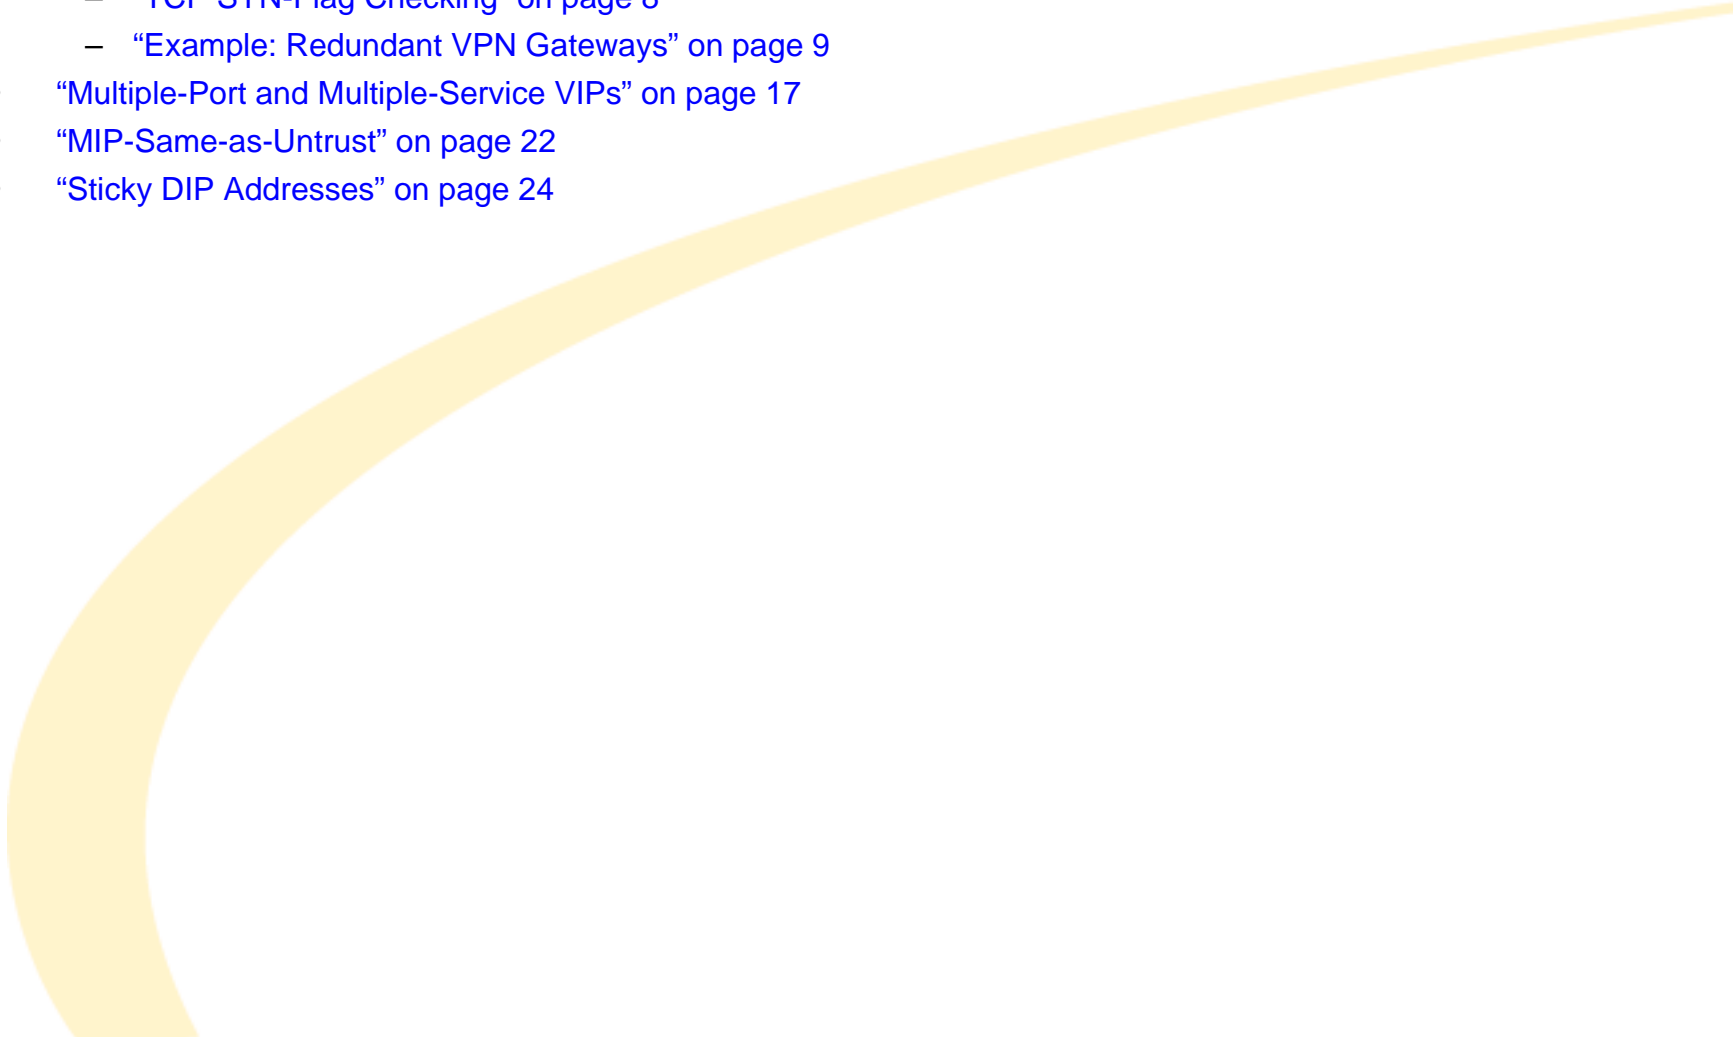
10. Tax Liability. You agree to be responsible for the payment of any sales or use taxes imposed at any time whatsoever on this transaction.

11. General. If any provisions of this Agreement are held invalid, the remainder shall continue in full force and effect. The laws of the State of California, excluding the application of its conflicts of law rules shall govern this License Agreement. This Agreement will not be governed by the United Nations Convention on the Contracts for the International Sale of Goods. This Agreement is the entire agreement between the parties as to the subject matter hereof and supersedes any other Technologies, advertisements, or understandings with respect to the Software and documentation. This Agreement may not be modified or altered, except by written amendment, which expressly refers to this Agreement and which, is duly executed by both parties.

You acknowledge that you have read this Agreement, understand it, and agree to be bound by its terms and conditions.

New Features for ScreenOS 3.0.1


This document describes the new features added to ScreenOS 3.0.1. The specific topics covered are as follows:

- [“Redundant VPN Gateways” on page 2](#)
 - [“VPN Groups” on page 3](#)
 - [“Monitoring Mechanisms” on page 4](#)
 - [“TCP SYN-Flag Checking” on page 8](#)
 - [“Example: Redundant VPN Gateways” on page 9](#)
 - [“Multiple-Port and Multiple-Service VIPs” on page 17](#)
 - [“MIP-Same-as-Untrust” on page 22](#)
 - [“Sticky DIP Addresses” on page 24](#)
- 

REDUNDANT VPN GATEWAYS

The NetScreen redundant gateway feature provides a solution for continuous VPN connectivity during and after a site-to-site failover. You can create a VPN group to provide a set of up to four redundant gateways to which LAN-to-LAN or LAN-to-LAN Dynamic Peer AutoKey IKE IPsec¹ VPN tunnels can connect. When the NetScreen device first receives traffic matching a policy referencing a VPN group, it performs Phase 1 and Phase 2 IKE negotiations with all members in that group. The NetScreen device sends data through the VPN tunnel to the gateway with the highest priority, or weight, in the group. For all other gateways in the group, the NetScreen device maintains the Phase 2 SAs and keeps the tunnels active by sending IKE keepalive packets through them. If the active VPN tunnel fails, the tunnel can fail over to the tunnel and gateway with the second highest priority in the group.

Note: This scheme assumes that the sites behind the redundant gateways are connected so that data is mirrored among hosts at all sites. Furthermore, each site—being dedicated to high availability (HA)—would have a redundant cluster of NetScreen devices operating in HA mode. Therefore, the VPN failover threshold must be set higher than the device failover threshold or VPN failovers might occur unnecessarily.

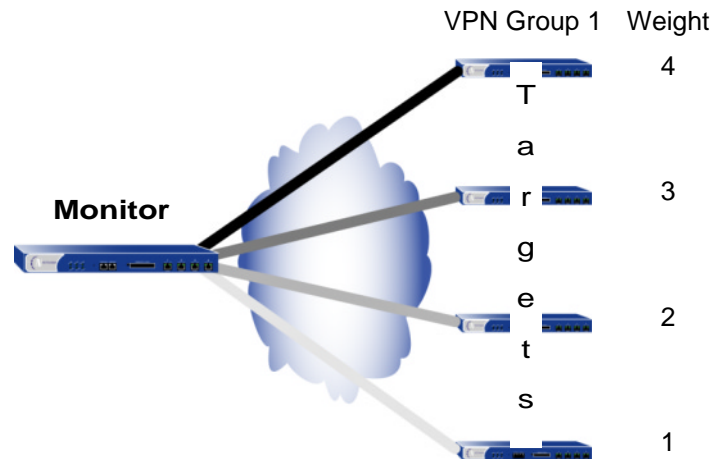


- VPN groups do not support L2TP, L2TP-over-IPsec, dialup-to-LAN, or Manual Key VPN tunnel types. In a LAN-to-LAN Dynamic Peer arrangement, the NetScreen device monitoring the VPN group must be the one whose untrust IP address is dynamically assigned, while the untrust IP addresses of the VPN group members must be static.

VPN Groups

A VPN group is a set of VPN tunnel configurations for up to four targeted remote gateways. The Phase 1 and Phase 2 security association (SA) parameters for each tunnel in a group can be different or identical (except for the IP address of the remote gateway). The VPN group has a unique ID number, and each member in the group is assigned a unique weight to indicate its place in rank of preference to be the active tunnel. A value of 1 indicates the lowest, or least preferred, ranking.

Note: In this illustration, the shading symbolizes the weight of each tunnel. The darker the tunnel is shaded, the higher its priority.



The NetScreen device communicating with VPN group members and the members themselves have a monitor-to-target relationship. The monitoring device continually monitors the connectivity and wellbeing of each targeted device. The tools with which the monitor does this are as follows:

- IKE heartbeats
- IKE recovery attempts

These two tools are presented in the next section, [“Monitoring Mechanisms” on page 4](#).

Note: The monitor-to-target relationship need not be one way. The monitoring device might also be a member of a VPN group and thus be the target of another monitoring device.

Monitoring Mechanisms

NetScreen uses two mechanisms to monitor members of a VPN group to determine their ability to terminate VPN traffic:

- IKE heartbeats
- IKE recovery attempts

Using these two tools, plus the TCP application failover option (see [“TCP SYN-Flag Checking” on page 8](#)), NetScreen devices can detect when a VPN failover is required and shift traffic to the new tunnel without disrupting VPN service.

IKE Heartbeats

IKE heartbeats are hello messages that IKE peers send to each other through the VPN tunnel to confirm the connectivity and wellbeing of the other. If, for example, device_m (the “monitor”) does not receive a specified number of heartbeats (the default is 5) from device_t (the “target”), device_m concludes that device_t is down. Device_m clears the corresponding Phase 1 and Phase 2 security associations (SAs) from its SA cache and begins the IKE recovery procedure. (See [“IKE Recovery Procedure” on page 5](#).) Device_t also clears its SAs.

Note: The IKE heartbeats feature must be enabled on the devices at both ends of a VPN tunnel in a VPN group. If it is enabled on device_m but not on device_t, device_m suppresses IKE heartbeat transmission and generates the following message in the event log: “Heartbeats have been disabled because the peer is not sending them.”



IKE Heartbeats must flow both ways through the VPN tunnel.

To define the IKE heartbeat interval and threshold for a specified VPN tunnel (the default is 5), do the following:

WebUI

VPN >> Gateway (P1) >> Edit (for the gateway whose IKE heartbeat threshold you want to modify): Enter the new values in the Heartbeat Hello and Heartbeat Threshold fields, and then click **OK**.

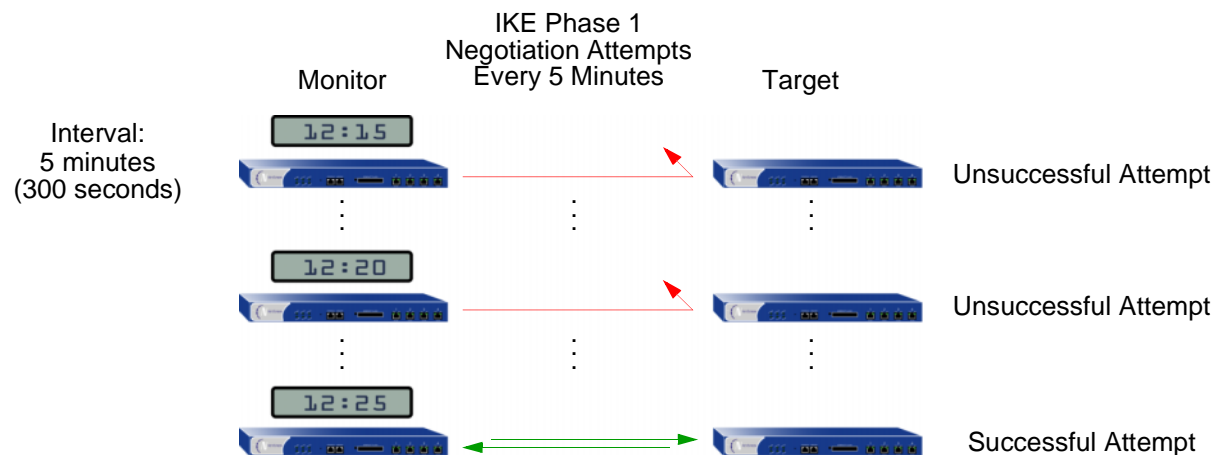
CLI

```
set ike gateway <name> heartbeat hello <seconds>
```

```
set ike gateway <name> heartbeat threshold <number>
```

IKE Recovery Procedure

After the monitoring NetScreen device determines that a targeted device is down, the monitor stops sending IKE heartbeats and clears the SAs for that peer from its SA cache. After a defined interval, the monitor attempts to initiate Phase 1 negotiations with the failed peer. If the first attempt is unsuccessful, the monitor continues to attempt Phase 1 negotiations at regular intervals until negotiations are successful.



To define the IKE recovery interval for a specified VPN tunnel (the minimum setting is 60 seconds), do either of the following:

WebUI

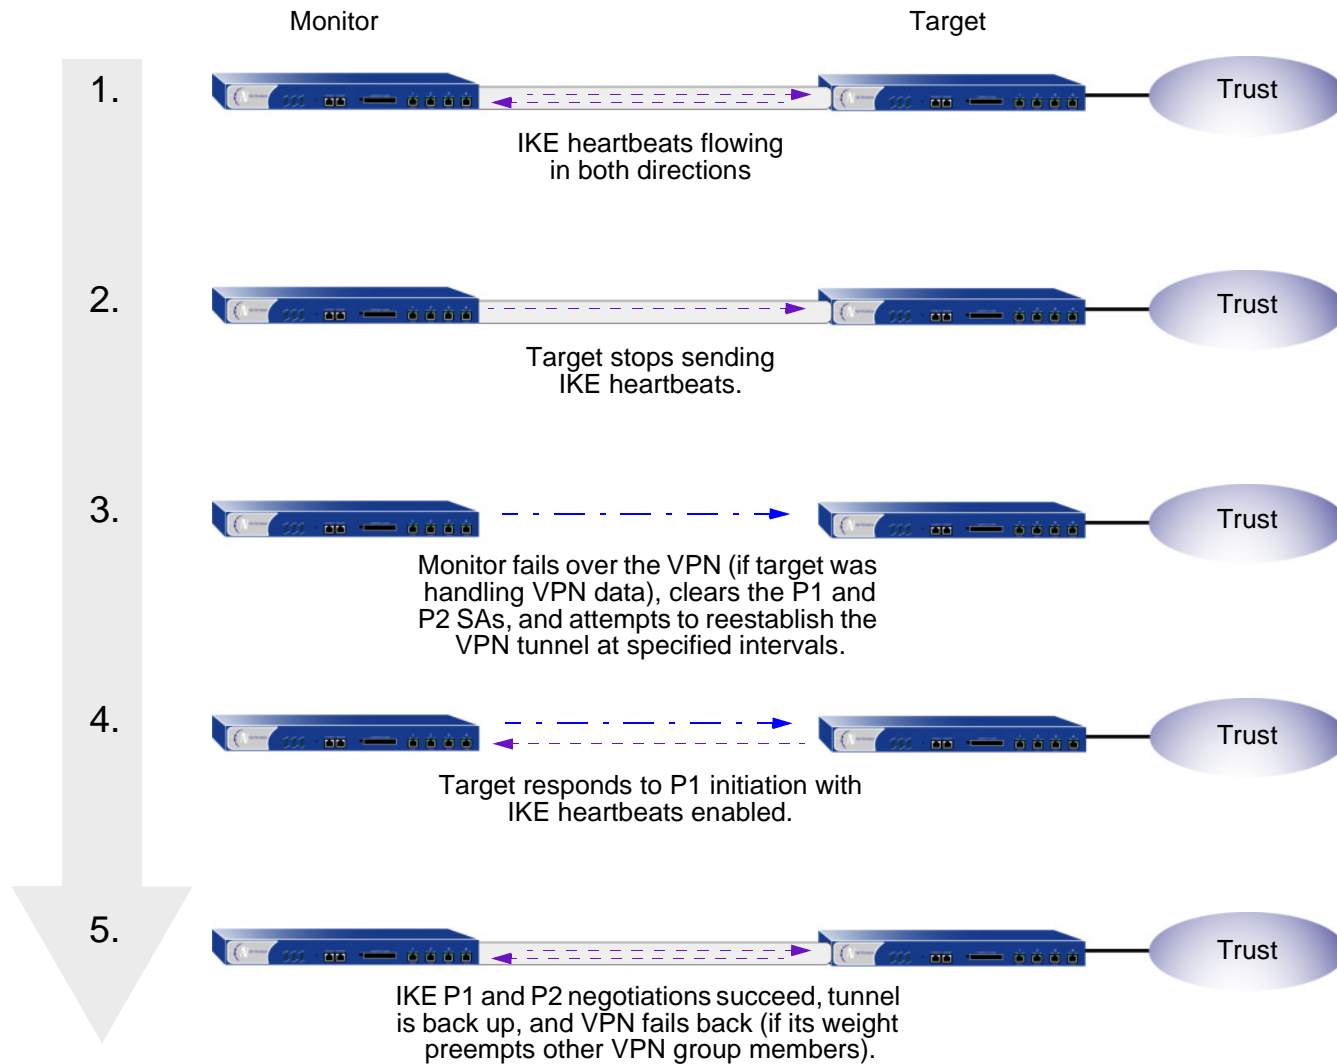
VPN >> Gateway (P1) >> Edit (for the gateway whose IKE heartbeat interval you want to modify): Enter the value in seconds in the Heartbeat: Reconnect field.

CLI

```
set ike gateway <name> heartbeat reconnect <seconds>
```

When a VPN group member with the highest weight fails over the tunnel to another group member and then reconnects with the monitoring device, the tunnel automatically fails back to the first member. The weighting system always causes the best ranking gateway in the group to handle the VPN data whenever it can do so.

The process that a member of a VPN group undergoes when the missing heartbeats from a targeted gateway surpass the failure threshold, is presented in the following illustration.



TCP SYN-Flag Checking

For a seamless VPN failover to occur, the handling of TCP sessions must be addressed. If, after a failover, the new active gateway receives a packet in an existing TCP session, the new gateway would treat it as the first packet in a new TCP session and check if the SYN flag is set in the packet header. Because this packet is really part of an existing session, it does not have the SYN flag set. Consequently, the new gateway would reject the packet. With TCP SYN-flag checking enabled, all TCP applications would have to reconnect after the failover occurs.

To resolve this, you can disable SYN-flag checking for TCP sessions in VPN tunnels, as follows:

WebUI

You cannot disable SYN-flag checking via the WebUI.

CLI

```
unset flow tcp-syn-check-in-tunnel
```

Note: By default, SYN-flag checking is enabled.

Example: Redundant VPN Gateways


In this example, a corporate site has one VPN tunnel to a data center and a second tunnel to a backup data center. All the data is mirrored via a leased line connection between the two data center sites. The data centers are physically separate to provide continuous service even in the event of a catastrophic failure such as an all-day power outage or a natural disaster.

The location, name, and IP addresses for the trust and untrust interfaces for each NetScreen device are as follows:

Device Location	Device Name	IP Address (Trust Interface)	IP Address (Untrust Interface)	VPN Group ID and Weight
Corporate	Monitor1	10.10.1.1/24	1.1.1.1/24	--
Data Center (Primary)	Target1	10.1.1.1/16	2.2.2.1/24	ID = 1, Weight = 2
Data Center (Backup)	Target2	10.1.1.1/16	3.3.3.1/24	ID = 1, Weight = 1

Note: *The internal address space at both data center sites must be identical.*

Both Phase 1 and Phase 2 SAs of the LAN-to-LAN AutoKey IKE tunnels employ 3DES encryption, SHA-1 authentication, and Diffie-Hellman Group 2 (with Perfect Forward Secrecy enabled in Phase 2). Phase 1 is in Main Mode. Preshared keys authenticate the participants, and the Encapsulating Security Payload (ESP) protocol provides authentication and encryption.



WebUI (Monitor1)

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:
 IP Address: 10.10.1.1
 Netmask: 255.255.255.0
2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save and Reset**:
 IP Address: 1.1.1.1
 Netmask: 255.255.255.0

3. Address >> Trusted >> New Address: Enter the following, and then click **OK**:
 - Address Name: in_trust
 - IP Address/Domain Name: 10.10.1.0
 - Netmask: 255.255.255.0
4. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:
 - Address Name: data_ctr
 - IP Address/Domain Name: 10.1.0.0
 - Netmask: 255.255.0.0
5. Configure >> Route Table >> New Entry: Enter the following, and then click **OK**:
 - Network Address: 0.0.0.0
 - Netmask: 0.0.0.0
 - Gateway IP Address: 1.1.1.2
 - Interface: untrust
 - Metric: 1
6. VPN >> VPN Group >> New VPN Group Entry: Enter **1** in the VPN Group ID field, and then click **OK**.
7. VPN >> Gateway (P1) >> New Remote Tunnel Gateway: Enter the following, and then click **OK**:
 - Gateway Name: target1
 - Remote Gateway: Static IP Address: 2.2.2.1
 - Mode (Initiator): Main (ID Protection)
 - Phase 1 Proposal: pre-g2-3des-sha
 - Preshared Key: SLi1yoo129
 - Heartbeat: Hello: 3 Seconds
 - Reconnect: 60 seconds
 - Threshold: 5

8. VPN >> AutoKey (P2) >> New AutoKey IKE Entry: Enter the following, and then click **OK**:
 - Name: to_target1
 - Remote Gateway Tunnel: target1
 - Phase 2 Proposal: pre-g2-3des-sha
 - VPN Group: VPN Group-1
 - Weight: 2
9. VPN >> Gateway (P1) >> New Remote Tunnel Gateway: Enter the following, and then click **OK**:
 - Gateway Name: target2
 - Remote Gateway: Static IP Address: 3.3.3.1
 - Mode (Initiator): Main (ID Protection)
 - Phase 1 Proposal: pre-g2-3des-sha
 - Preshared Key: CMFwb7oN23
 - Heartbeat: Hello: 3 Seconds
 - Reconnect: 60 seconds
 - Threshold: 5
10. VPN >> AutoKey (P2) >> New AutoKey IKE Entry: Enter the following, and then click **OK**:
 - Name: to_target2
 - Remote Gateway Tunnel: target2
 - Phase 2 Proposal: pre-g2-3des-sha
 - VPN Group: VPN Group-1
 - Weight: 1

11. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

Source Address: in_trust

Destination Address: data_ctr

Service: ANY

Action: Tunnel

VPN: VPN Group-1

Create matching incoming VPN policy: (select)

WebUI (Target1)

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:

IP Address: 10.1.1.1

Netmask: 255.255.0.0

2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save and Reset**:

IP Address: 2.2.2.1

Netmask: 255.255.255.0

3. Address >> Trusted >> New Address: Enter the following, and then click **OK**:

Address Name: in_trust

IP Address/Domain Name: 10.1.0.0

Netmask: 255.255.0.0

4. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name: 10.10.1.0

Netmask: 255.255.255.0

5. Configure >> Route Table >> New Entry: Enter the following, and then click **OK**:
 - Network Address: 0.0.0.0
 - Netmask: 0.0.0.0
 - Gateway IP Address: 2.2.2.2
 - Interface: untrust
 - Metric: 1
6. VPN >> Gateway (P1) >> New Remote Tunnel Gateway: Enter the following, and then click **OK**:
 - Gateway Name: monitor1
 - Remote Gateway: Static IP Address: 1.1.1.1
 - Mode (Initiator): Main (ID Protection)
 - Phase 1 Proposal: pre-g2-3des-sha
 - Preshared Key: SLi1yoo129
 - Heartbeat: Hello: 3 Seconds
 - Reconnect: 0 seconds
 - Threshold: 5
7. VPN >> AutoKey (P2) >> New AutoKey IKE Entry: Enter the following, and then click **OK**:
 - Name: to_monitor1
 - Remote Gateway Tunnel: monitor1
 - Phase 2 Proposal: pre-g2-3des-sha
8. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
 - Source Address: in_trust
 - Destination Address: corp
 - Service: ANY

Action: Tunnel

VPN: monitor1

Create matching incoming VPN policy: (select)

Note: Follow the same steps for configuring Target1 to configure Target2, but define the untrust interface IP address as 3.3.3.1/24, the default gateway IP address as 3.3.3.2, and use CMFwb7oN23 to generate the preshared key.

CLI (Monitor1)

1. set interface trust ip 10.10.1.1 255.255.255.0
2. set interface untrust ip 1.1.1.1 255.255.255.0
3. set address trust in_trust 10.10.1.0 255.255.255.0
4. set address untrust data_ctr 10.1.0.0 255.255.0.0
5. set route 0.0.0.0 0.0.0.0 gateway 1.1.1.2
6. set ike gateway target1 ip 2.2.2.1 main preshare SLi1yoo129 proposal pre-g2-3des-sha
7. set ike gateway target1 heartbeat hello 3
8. set ike gateway target1 heartbeat reconnect 60
9. set ike gateway target1 heartbeat threshold 5
10. set vpn to_target1 gateway target1 tunnel proposal g2-esp-3des-sha
11. set ike gateway target2 ip 3.3.3.1 main preshare CMFwb7oN23 proposal pre-g2-3des-sha
12. set ike gateway target2 heartbeat hello 3
13. set ike gateway target2 heartbeat reconnect 60
14. set ike gateway target2 heartbeat threshold 5
15. set vpn to_target2 gateway target2 tunnel proposal g2-esp-3des-sha
16. set vpn-group id 1 vpn to_target1 weight 2
17. set vpn-group id 1 vpn to_target2 weight 1

18. unset flow tcp-syn-check-in-tunnel
19. set policy outgoing in_trust data_ctr any tunnel vpn-group 1
20. set policy incoming data_ctr in_trust any tunnel vpn-group 1
21. save

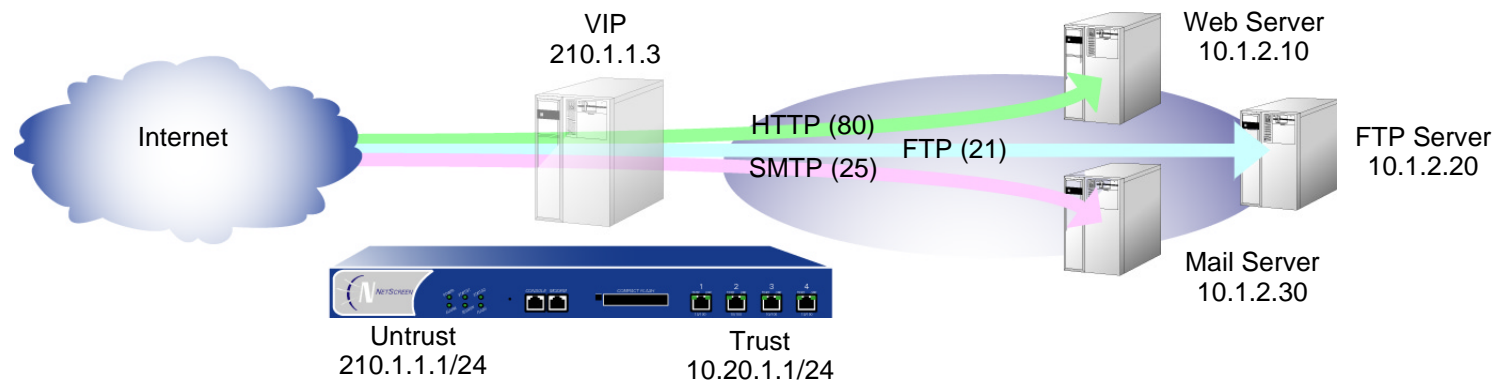
CLI (Target1)

1. set interface trust ip 10.1.1.1 255.255.0.0
2. set interface untrust ip 2.2.2.1 255.255.255.0
3. set address trust in_trust 10.1.0.0 255.255.0.0
4. set address untrust corp 10.10.1.0 255.255.255.0
5. set route 0.0.0.0 0.0.0.0 gateway 2.2.2.2
6. set ike gateway monitor1 ip 1.1.1.1 main preshare SLi1yoo129 proposal pre-g2-3des-sha
7. set vpn to_monitor1 gateway monitor1 tunnel proposal g2-esp-3des-sha
8. set ike gateway target1 heartbeat hello 3
9. set ike gateway target1 heartbeat threshold 5
10. set policy outgoing in_trust corp any tunnel vpn to_monitor
11. set policy incoming corp in_trust any tunnel vpn to_monitor
12. save

Note: Follow the same steps for configuring Target1 to configure Target2, but define the untrust interface IP address as 3.3.3.1/24, the default gateway IP address as 3.3.3.2, and use CMFwb7oN23 to generate the preshared key.

MULTIPLE-PORT AND MULTIPLE-SERVICE VIPs

A virtual IP (VIP) address maps traffic received at one IP address to another address based on the destination port number in the packet header. For example, an HTTP packet destined for 210.1.1.3:80 (that is, IP address 210.1.1.3 and port 80) might get mapped to a web server at 10.1.2.10. An FTP packet destined for 210.1.1.3:21 might get mapped to an FTP server at 10.1.2.20. An SMTP packet destined for 210.1.1.3:25 might get mapped to a mail server at 10.1.2.30. The destination port numbers determine to which host the same VIP sends them.



New VIP features and capabilities in ScreenOS 3.0.1:

- You can map up to 64 services from a single VIP to one or more servers².
- You can map more predefined services³ and user-defined services.
- Custom services with the same source and destination port numbers but different transports are distinguishable and can be individually handled by the same VIP.
- A single VIP can support custom services with multiple port entries by creating multiple service entries under that VIP—one service entry in the VIP for each port entry in the service.


2. On NetScreen devices that support load balancing, you can map up to eight services per VIP.

3. By default you can use the single-port services. To be able to use multi-port services in a VIP, issue the CLI command **set vip multi-port**.

Example: VIP with Custom and Multiple-Port Services

In the following example, you configure a VIP at 210.1.1.3 to route the following services to the following internal addresses:

Service	Transport	Virtual Port Number	Actual Port Number	Host IP Address
DNS	TCP, UDP	53	53	10.20.1.3
HTTP	TCP	80	80	10.20.1.4
PCAnywhere	TCP, UDP	5631, 5632	5631, 5632	10.20.1.4
LDAP	TCP, UDP	5983	389	10.20.1.5



Note: Because a NetScreen device with load balancing does not support more than eight ports per VIP, you must configure extra VIPs to complete the configuration in this example.

The VIP routes DNS lookups to the DNS server at 10.20.1.3, HTTP traffic to the web server at 10.20.1.4, and authentication checks to the database on the LDAP server at 10.20.1.5. For HTTP, DNS, and PCAnywhere, the virtual port numbers remain the same as the actual port numbers. For LDAP, a virtual port number (5983) is used to add an extra level of security to the LDAP authentication traffic.

For managing the HTTP server remotely, you define a custom service and name it PCAnywhere. PCAnywhere is a multiple-port service that sends and listens for data on TCP port 5631 and status checks on UDP port 5632.

WebUI

Custom Service

1. Service >> Custom >> New Service: Enter the following, and then click **OK**:

Service Name: PCAnywhere

No 1:

Source Port Low: 0

Source Port High: 65535

Destination Port Low: 5631

Destination Port High: 5631

Transport: TCP

No 2:

Source Port Low: 0

Source Port High: 65535

Destination Port Low: 5632

Destination Port High: 5632

Transport: UDP

VIP Address and Services⁴

2. Virtual IP >> Virtual IP 1 >> Virtual Server IP: click here to configure: Type **210.1.1.3** in the Virtual IP Address field, and then click **OK**.
3. Virtual IP >> Virtual IP 1 >> New Service: Enter the following, and then click **OK**:
 - Virtual Port: 53
 - Map to Service: DNS
 - Map to IP: 10.20.1.3⁵
4. Virtual IP >> Virtual IP 1 >> New Service: Enter the following, and then click **OK**:
 - Virtual Port: 80
 - Map to Service: HTTP
 - Map to IP: 10.20.1.4
5. Virtual IP >> Virtual IP 1 >> New Service: Enter the following, and then click **OK**:
 - Virtual Port: 5631⁶
 - Map to Service: PCAnywhere
 - Map to IP: 10.20.1.4
6. Virtual IP >> Virtual IP 1 >> New Service: Enter the following, and then click **OK**:
 - Virtual Port: 5983
 - Map to Service: LDAP
 - Map to IP: 10.20.1.5

-
4. To enable the VIP to support multiple-port services, you must use enter the CLI command **set vip multi-port**, save the configuration, and then reboot the device.
 5. For NetScreen devices with load balancing capabilities, select **None** from the Load Balance drop-down list, and type **1** in the Server Weight field.
 6. For multiple-port services, enter the lowest port number of the service.

CLI

Custom Service

1. set service pcanynwhere protocol udp src-port 0-65535 dst-port 5631-5631
2. set service pcanynwhere protocol tcp src-port 0-65535 dst-port 5632-5632

VIP Address and Services

3. set vip multi-port
4. save
5. reset

When prompted if you are sure you want to reset the device, press the Y key.

6. set vip 210.1.1.3 + 53 dns none 10.20.1.3/1
7. set vip 210.1.1.3 + 80 http none 10.20.1.4/1
8. set vip 210.1.1.3 + 5631 pcanynwhere none 10.20.1.4/1
9. set vip 210.1.1.3 + 5983 ldap none 10.20.1.5/1
10. save

MIP-SAME-AS-UNTRUST

As IPv4 addresses become increasingly scarce, ISPs are becoming increasingly reluctant to give their customers more than one or two IP addresses. If you only have one IP address for the untrust interface—the trust zone is in Network Address Translation (NAT) mode—you can use the untrust interface IP address as a mapped IP (MIP) to provide inbound access to an internal server or host. A MIP maps traffic arriving at the one address to another address, so by using the untrust interface IP address as a MIP, all inbound traffic arriving at the untrust interface will be mapped to a specified internal address.

If you create an access policy in which the destination address is a MIP using the untrust interface IP address and the service specified in the policy is HTTP, you lose web management of the NetScreen device via the untrust interface (because all inbound HTTP traffic to that address is mapped to an internal server or host). You can still manage the device via the untrust interface using the WebUI by changing the port number for web management. To change the web management port number, do the following:

1. Admin >> Web: Enter a registered port number (from 1024 to 65,535) in the HTTP Port field. Then click **Apply**.
2. When you next connect to the untrust interface to manage the device, append the port number to the IP address. (For example, `http://209.157.66.170:5000`)

Example: MIP on the Untrust Interface

In this example, you select the IP address of the untrust interface (210.1.1.1/24) as the MIP for a web server whose actual IP address is 10.10.1.5. Because you want to retain web management access to the untrust interface, you change the web management port number to 8080. You then create an incoming policy permitting HTTP service from the untrust zone to the MIP.

WebUI

1. Admin >> Web: Type **8080** in the HTTP Port field, and then click **Apply**.

The HTTP connection is lost.

2. Reconnect to the NetScreen device, appending 8080 to the IP address in the URL address field in your web browser. (If you are currently managing the device via the untrust interface, type **http://210.1.1.1:8080**.)
3. Interface >> Untrusted >> Mapped IP >> New Entry: Enter the following, and then click **OK**:

Mapped IP Address: 210.1.1.1

Netmask: 255.255.255.255⁷

Host IP Address: 10.10.1.5

4. Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:

Source Address: Outside Any

Destination Address: MIP(210.1.1.1)

Service: HTTP

Action: Permit

7. The netmask for a MIP using the untrust interface IP address must be 32 bits.

STICKY DIP ADDRESSES

When a host initiates several sessions that match an access policy with network address translation (NAT) enabled and is assigned an address from a dynamic IP (DIP) pool, the NetScreen device assigns a different source IP address for each session. Such random address assignment can be problematic for services that create multiple sessions that require the same source IP address for each session.

For example, it is important to have the same IP address for multiple sessions when using the AOL Instant Messaging (AIM) client. You create one session when you log in, and another for each chat. For the AIM server to verify that a new chat belongs to an authenticated user, it must match the source IP address of the login session with that of the chat session. If they are different—possibly because they were randomly assigned from a DIP pool during the NAT process—the AIM server rejects the chat session.

To ensure that the NetScreen device assigns the same IP address from a DIP pool to a host for multiple concurrent sessions, you can enable the “sticky” DIP address feature. To enable this feature, enter the CLI command **set dip sticky**.