



Continent Enterprise Firewall Version 4

Basics

Administrator guide



© **SECURITY CODE LLC, 2024. All rights reserved.**

All rights to operation manuals are reserved.

This document is shipped along with the product kit. It is covered by all terms of license agreement. You may not copy this document in printed or electronic form, in whole or part, or deliver it to third parties on commercial purpose without a special written consent of Security Code LLC.

Security Code LLC reserves the right to change the information contained herein without special notice.

Mailing address:	115230, Russian Federation, Moscow, 1st Nagatinsky proezd 10/1
Phone:	+7 (495) 982-30-20
E-mail:	info@securitycode.ru
Web:	www.securitycode.ru

Table of contents

List of abbreviations	5
Introduction	6
Overview	7
Purpose	7
Continent product kit	7
Configuration Manager	7
Security Gateway	7
Identification Agent	8
General architecture of a protected corporate network	8
Continent management	11
Local management	11
Security Gateway centralized management	11
Continent administration	11
Licensing	12
User identification	12
Update	15
Integrity check	15
Security certificates	15
Backup and failover	17
Security Management Server database backup	17
Security Gateway failover	17
Security Management Server failover	19
IP packets processing	21
Receiving packets	21
Network Behavior Anomaly Detector	22
Connection monitoring	22
Destination NAT	22
Inbound traffic QoS	22
Firewall	22
Web filter	23
Application control	23
Managing connections	23
Related connection tracking	24
IP address geographic filtering	24
IPS	24
Routing	25
Static routing	25
Dynamic routing	25
Source NAT	25
Outbound traffic QoS	25
VPN principles	26
VPN tunnel	26
Encryption	27
Topology	27
Licenses	28
Intrusion detection and prevention	29
Remote user access	31
Access Server	31
Continent-RA	31
Remote user access to protected network resources	32
Access Server management	32

Networking functions	33
QoS support	33
Network device management using SNMP	34
Routing	34
Virtual routing	34
DNS	34
ARP	34
DHCP	34
Time synchronization	35
SSH	35
Bonds	35
Collecting data on neighboring network devices	36
Audit and monitoring	37
Documentation	39

List of abbreviations

AD	Active Directory
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
CA	Certification Authority
FTP	File Transfer Protocol
HFSC	Hierarchical Fair Service Curve
IP	Internet Protocol
IA	Identification Agent
LAN	Local Area Network
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NBAD	Network Behavior Anomaly Detector
SNMP	Simple Network Management Protocol
SPAN	Switched Port Analyzer
SSH	Secure Shell
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
URL	Uniform Resource Locator
USB	Universal Serial Bus
UTM	Unified Threat Management
VPN	Virtual Private Network

Introduction

This manual is designed for administrators of Continent Enterprise Firewall, Version 4 (hereinafter — Continent). It contains information about Continent core functions.

Website. Information about SECURITY CODE LLC products can be found on <https://www.securitycode.ru>.

Technical support. You can contact technical support by phone: +7 800 505 30 20 or by email: support@securitycode.ru.

Training. You can learn more about hardware and software products of SECURITY CODE LLC in authorized education centers. The list of the centers and information about the learning environment can be found on <https://www.securitycode.ru/company/education/training-courses/>.

You can contact a company's representative for more information about trainings by email: education@securitycode.ru.

Version 4.1.9 — Released on May 22nd, 2024.

Chapter 1

Overview

Purpose

Cutting-edge network technologies require a holistic approach to network asset protection.

UTM solutions are being rapidly developed in the network security field. UTM devices make it possible to implement multiple security measures such as firewall protection, IPS, application control, remote user access control, etc. within a single hardware or virtual platform.

Continent is a UTM solution that performs the following core functions:

- firewall protection (see p. [22](#));
- intrusion detection and prevention (see p. [29](#));
- remote user access to a corporate network (see p. [31](#));
- packet switching, routing, network address translation, VLAN creation, etc. (see p. [33](#));
- event logging, auditing and monitoring (see p. [37](#));
- centralized and local management of Continent components (see p. [11](#)).

Continent product kit

Continent product kit consists of:

- Identification Agent;
- Configuration Manager;
- Security Gateway.

Configuration Manager

The Configuration Manager is software that is installed on one or more computers (administrators' workstations) to manage the Continent security policy.

Security Gateway

The Security Gateway is a set of software tools that is installed on a hardware platform (hereinafter — Security Gateway). A Security Gateway can operate in three modes: UTM, IPS and NF2 (only for specific hardware platforms). According to the selected mode, you can enable the following software components on the Security Gateway:

Mode	Software component
UTM	Security Management Server Firewall QoS L2VPN L3VPN IPS Access Server User Identification Network Behavior Anomaly Detector
High-performance firewall	Firewall QoS
IPS	IPS

Security Management Server

The Security Management Server allows you to configure and manage Continent components, audit Security Gateways and monitor their status.

Firewall

Firewall is an essential component in Security Gateway UTM and High-performance firewall modes. In UTM mode, you can enable the **Advanced Protocol and Application Control** and **Malicious URL Blocking, SkyDNS URL Blocking, Antivirus** and **Geo Protection** modes.

L2VPN

L2VPN provides secure Ethernet frame transmission by using VPN tunneling over public networks.

L3VPN

L3VPN provides secure IP packets transmission by using VPN tunneling over public networks between network segments.

IPS

IPS is aimed at traffic analysis using heuristic and signature methods to detect and block malicious content and notify the administrator of such security policy violations. It is implemented as a separate device or as a part of Security Gateway software components in UTM mode.

IPS can be enabled in two modes: Inline and Monitor.

Access Server

The Access Server provides remote users with access to protected network resources.

User Identification

User identification provides identification of built-in users and users from Active Directory.

Network Behavior Anomaly Detector

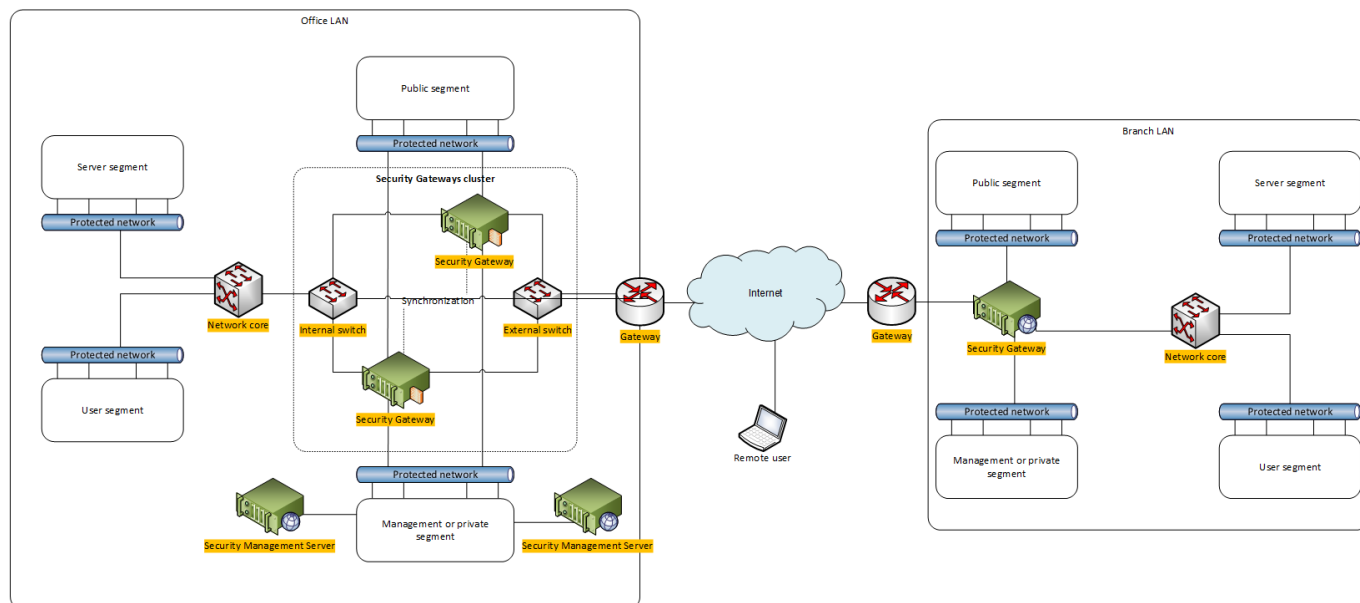
NBAD detects and blocks anomalous traffic and SYN-scan, FIN/RST-scan, SYN-flood, FIN/RST-flood attacks.

Identification Agent

Identification Agent is a program installed on a user's workstation within a protected network. When a user tries to get access to the Internet resources, the Identification Agent verifies user credentials to provide communication with the Security Server.

General architecture of a protected corporate network

The general architecture of a protected corporate network including several LANs is shown in the figure below.



LANs are connected over public networks. Each LAN is connected to a public network through a Security Gateway. It prevents LANs from the disclosure of their structure. Also, each IP address used in protected networks must be

unique for this corporate network. The Security Gateway can support several network interfaces, so it can be connected to several independent LANs. Several network interfaces can be bound together to enhance the reliability and increase the bandwidth of connection (for more details about bonds, see p. 35).

If users are connected to the Security Gateway, they can be authenticated by the Authentication Portal (see p. 12).

The Security Gateway routes IP traffic. If you need an additional router, you can place it before a Security Gateway (within a protected network segment) or after it (outside of a protected network segment). If a router is placed in a protected network, no additional security measures are required.

Besides traffic routing, the Security Gateway processes incoming and outgoing IP packets, filters and encrypts data transferred over public networks.

The Firewall filters and translates IP packets. You can configure this component by creating access control rules (Firewall and NAT rules).


The Firewall rules are used to control user access to protected internal and external networks by accepting or denying traffic through the Security Gateway.

Network address translation (NAT) rules modify address information of packets in transit.

The Firewall and NAT rules are grouped in lists in a strict order. The list defines a sequence of actions on packets processed by Security Gateways. These lists are empty by default.

The NAT mechanism is used to provide corporate users with access to public network resources.

To enable cryptographic protection of data transferred over public networks between remote subnetwork segments, you can use the L2VPN and L3VPN components. The remote segments interaction diagram is shown in the figure below.



Security Gateways are managed by the Security Management Server enabled on one of the Security Gateways. The Security Management Server operates similarly to other Security Gateways, that is, it can be used to receive, transfer, filter, route, encrypt IP packets.

You can manage Security Gateways locally using a keyboard and a monitor connected to the Security Management Server and configure a Security Gateway remotely through SSH.

You can manage Security Gateways using the Configuration Manager as well. The Configuration Manager must be installed on a computer within the protected corporate network segment (the administrator's workstation).

The Authentication Portal is designed to identify and authenticate users registered in the system and provide them with access to the Internet.

To establish connection from a remote workstation that is not included in the protected local network segment, enable the Access Server. The Access Server can be installed either on a standalone Security Gateway or on the Security Management Server. The administrator manages the Access Server via the Configuration Manager.

If control packets are routed through a firewall or other filtering device, create rules that allow passing service packets through these devices.

Continent Security Gateways do not conflict with devices that support NAT.

Chapter 2

Continent management

Local management

To configure a Security Gateway locally, connect a keyboard and a monitor to it. If necessary, instead of a keyboard, you can connect a computer or a laptop via the serial console.

The Security Management Server local menu allows you to:

- view system and license information;
- manage root, intermediate and control certificates;
- view system, network security and management logs;
- clear event logs;
- run different types of diagnostics;
- create and load a configuration backup copy of the Security Management Server;
- configure Security Gateways;
- perform initialization;
- configure Security Management Server connections;
- set system time;
- configure network settings, etc.

Security Gateway centralized management

You can manage the Security Gateway remotely using the Configuration Manager. The Configuration Manager is installed on one or multiple computers within the protected network segment (the administrator's workstation). The number of the administrator's workstations is not limited.

Typically, the administrator's workstation must be within the protected network that contains the Security Management Server.

The Configuration Manager connects to the Security Management Server and allows you to control all Security Gateways. The Configuration Manager connects to the Security Management Server only after the administrator identification and authentication using the password or the certificate.

You can also restrict access to the Security Management Server for the Configuration Manager by configuring the list of IP addresses from which the connection to the Security Management Server is allowed (see [4]).

User account security policy allows to configure password policy, automatic system blocking in case of entering an incorrect password and connection break in case of Configuration Manager inactivity time-out.

Continent administration

Administrator roles

Administrators manage Continent according to their roles.

During the Security Management Server initialization, a default administrator account is created. This account is granted the whole set of privileges to manage the Security Management Server and Security Gateways of the domain.

Continent provides two types of administrator roles:

- built-in role with a predefined set of privileges. This role cannot be edited or deleted;
- user role with a custom set of privileges. The administrator of the Configuration Manager can create and edit it.

The default roles are as follows:

- main administrator;
- security administrator;
- network administrator;

- audit administrator.

You can assign a role to an administrator in two ways:

- create or edit an administrator's account;
- create or edit a role (for example, to give more privileges to the administrator group).

An administrator can be assigned several default and user roles. They are granted privileges according to the roles.

Licensing

A Security Gateway license defines what functions of the Security Gateway are available.

To apply policies, you must have a license (the database of the Security Management Server is checked for licenses).

By default, a demo license with the null client ID is saved to the Security Gateway database during the Security Gateway initialization. This license allows you to use Continent without limitation for 14 days. After the license expiration, policies cannot be applied to the Security Gateway. When you save a configuration or log on to the system, a warning is displayed notifying you of the license expiration.

When you link a license for the first time, the client ID specified in the license is saved to the Security Management Server database. After that, only relevant licenses can be added to the repository.

An unlinked license remains in the repository. You can link it to another Security Gateway or delete it.

Use the Configuration Manager to add, link, unlink or delete a license.

User identification

Users working on workstations within a protected network are identified and authenticated by:

- Authentication Portal;
- Identification Agent on an end-user device;
- transparent authentication via Kerberos (Single Sign-On).

You can create user profiles in the local database of the Security Management Server or import them from AD.

Access is granted to the users/groups using firewall rules.


The data about registered users is stored in the Security Management Server database. The data about authenticated users is stored on the Security Gateway.

When connecting to the firewall, user authentication on the Security Gateway is performed using not cryptographic methods and based on user account data.

Data exchange between the Security Gateway and the connected computer is performed over HTTPS.

How Authentication Portal works

The Authentication Portal is a component of one of the Security Gateways of the network and authenticates users through the web interface.




The user sends HTTP (TCP/80) or HTTPS (TCP/443) requests (1). The Authentication Portal (2) start web page appears if there are rules to access the network for which redirection to the Authentication Portal is configured. When accessing the web pages through HTTPS, an intermediate certificate is required for a trusted connection to the Security Gateway.

The user enters their credentials on the Authentication Portal web page (3). The Security Gateway checks the local database of the Security Management Server for these credentials and if they are still valid (4). If the username looks like **username@domain**, the request is redirected to the AD server of the corresponding domain (for example, **usertst1@local.host**) (4a). If the match is found and the credentials are proved to be valid, then respective data is sent to the Security Gateway. After successful authentication on the Security Gateway, a respective temporary firewall rule containing the correspondence of the username and the IP address from which the authentication was performed is created and the user is granted access to the resource (5).

How Identification Agent works

The Identification Agent is software that is installed on users' workstations. It is used for connection with the Security Gateway in order to check user account data.



To get access to the Internet, a user runs the Identification Agent and enters their credentials (1). Then, the agent initiates identification in AD (2). The agent receives a confirmation for identification (3). When the user attempts to access the Internet (4) for the first time, the agent sends the confirmation to AD and receives a permission. The

user is granted access to the Internet based on the Security Gateway access control rules (5). When the user repeats their attempt to access the Internet (6), the agent checks in its cache if there is the permission. If the permission is expired, the Identification Agent requests it from AD again.

Access to the Internet and local network resources protected by the Security Gateway is granted to the user if the following conditions are met:

- user account data are confirmed and valid at a time when a request is made;
- there are access rules for the user.

Note.

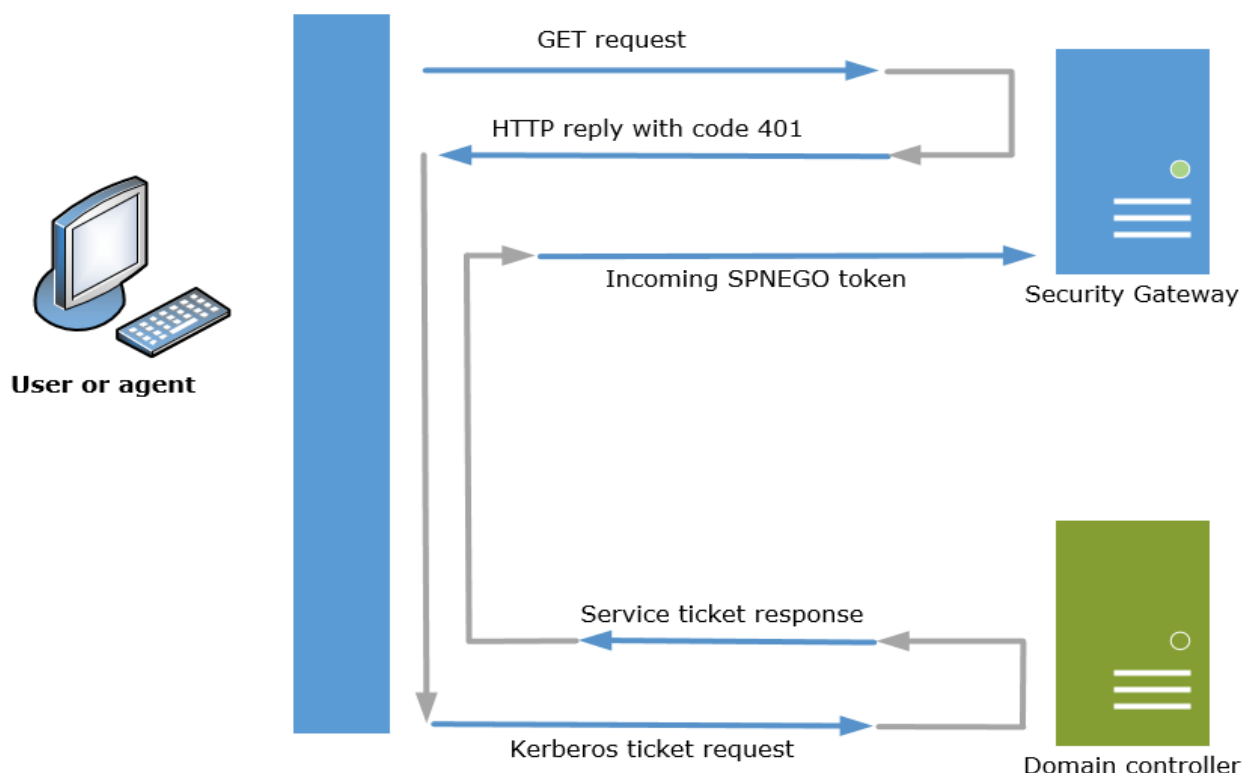
Identification Agent settings do not depend on Authentication Portal settings and use only a subordinate certificate specified in the Authentication Portal settings.

How transparent authentication with Kerberos works

By using transparent authentication the domain user does not need to receive repeated authentication requests especially when they are granted access to network resources. The user enters the domain username and password only once when they log in to an operating system. When accessing resources next time, the authentication is performed automatically.

Kerberos authentication can be used for direct access as well as access via the Continent proxy server.

Continent implements the browser transparent authentication mechanism using the SPNEGO protocol. The authentication process is shown in the figure below.



1. The user logs in to a Windows domain from their workstation and tries to get access to the Internet using a web browser. The web browser sends an HTTP request which is intercepted by the Security Gateway.
2. Then, the Security Gateway sends back an HTTP response with the 401 code (Unauthorized) and the authorization header "WWW-Authenticate: Negotiate".
3. The web browser recognizes the Negotiate header. Next, the search for the Security Gateway name starts in DNS, using which the name of the participant (SPN service) is found.
4. Using SPN, the local system authenticity validation service requests a Kerberos ticket from the Key Distribution Center (KDC). The Kerberos authentication process starts. As a result, the client receives a service ticket (ST), based on which the Security Gateway will trust the client.

5. The web browser repeats the initial HTTP request. But, this time, user authentication data is contained in a Kerberos encrypted ticket encapsulated in a SPNEGO token that is sent in the HTTP authorization header.
6. The Security Gateway finds the SPNEGO token, retrieves information from the Kerberos ST, which contains all the data required for authentication.
7. Transparent authentication can be used alongside the explicit web browser authentication via the Authentication Portal web page or they can be used separately.

Update

Continent software is updated using the Security Management Server repository. Update files can be downloaded to the repository manually or automatically according to a set schedule.

IPS Protections, SkyDNS categories, Kaspersky hash databases, Kaspersky Feeds, Web/FTP filtering exclusions and GEO/IP updates are downloaded from the update server or from the local source and written to the Security Management Server database. Updates can be downloaded from the server both manually and automatically on schedule. Upon writing updates to the Security Management Server database, the version, date and time of the last update are displayed in the Configuration Manager.

You can install updates using the Configuration Manager. For more information about updating, see [10].

Integrity check

This mechanism checks the integrity of Continent software.

Integrity check can start:

- automatically during the Continent installation (Continent integrity check);
- automatically every time upon running Continent (checksum verification);
- manually.

Security certificates

You can view available certificates via any Continent component as well as import certificates, security keys and make requests to issue certificates.

The Security Management Server allows you to create and export certificates listed in the table below.

Certificate type	Maximum service life	Signature algorithm	Note
Root CAs			
Trusted issuer (Security Code)	11 years	GOST	Only for secure connection with the update server
Root certificate	5 years	GOST	See [4]
RSA root certificate	5 years	RSA	See[5], Configure the network connection
Intermediate CAs			
Authentication portal-redirect	1 year	RSA	See [6], Authentication Portal
SSL/TLS inspection	1 year	RSA	See [2], Initial configuration
Personal certificates			
Administrator	1 year	GOST	See [4]
User	1 year	GOST	See [6], Remote access
Authentication Portal	1 year	RSA	See [6], Authentication Portal
Access Server	1 year	GOST	See [7], Remote access
Security Gateway	1 year	GOST	Secure connection with the Security Management Server and other Security Gateways
Web-monitoring	1 year	RSA	Establish a secure connection between the Security Gateway and the Security Management Server to access the monitoring web console

After authorization, the main administrator and the security administrator are granted privileges to manage certificates. Network administrators and audit administrators are granted privileges to view certificates.

Chapter 3

Backup and failover

Security Management Server database backup

The Security Management Server backup and restoration make it possible to deal with a Security Management Server failure promptly.

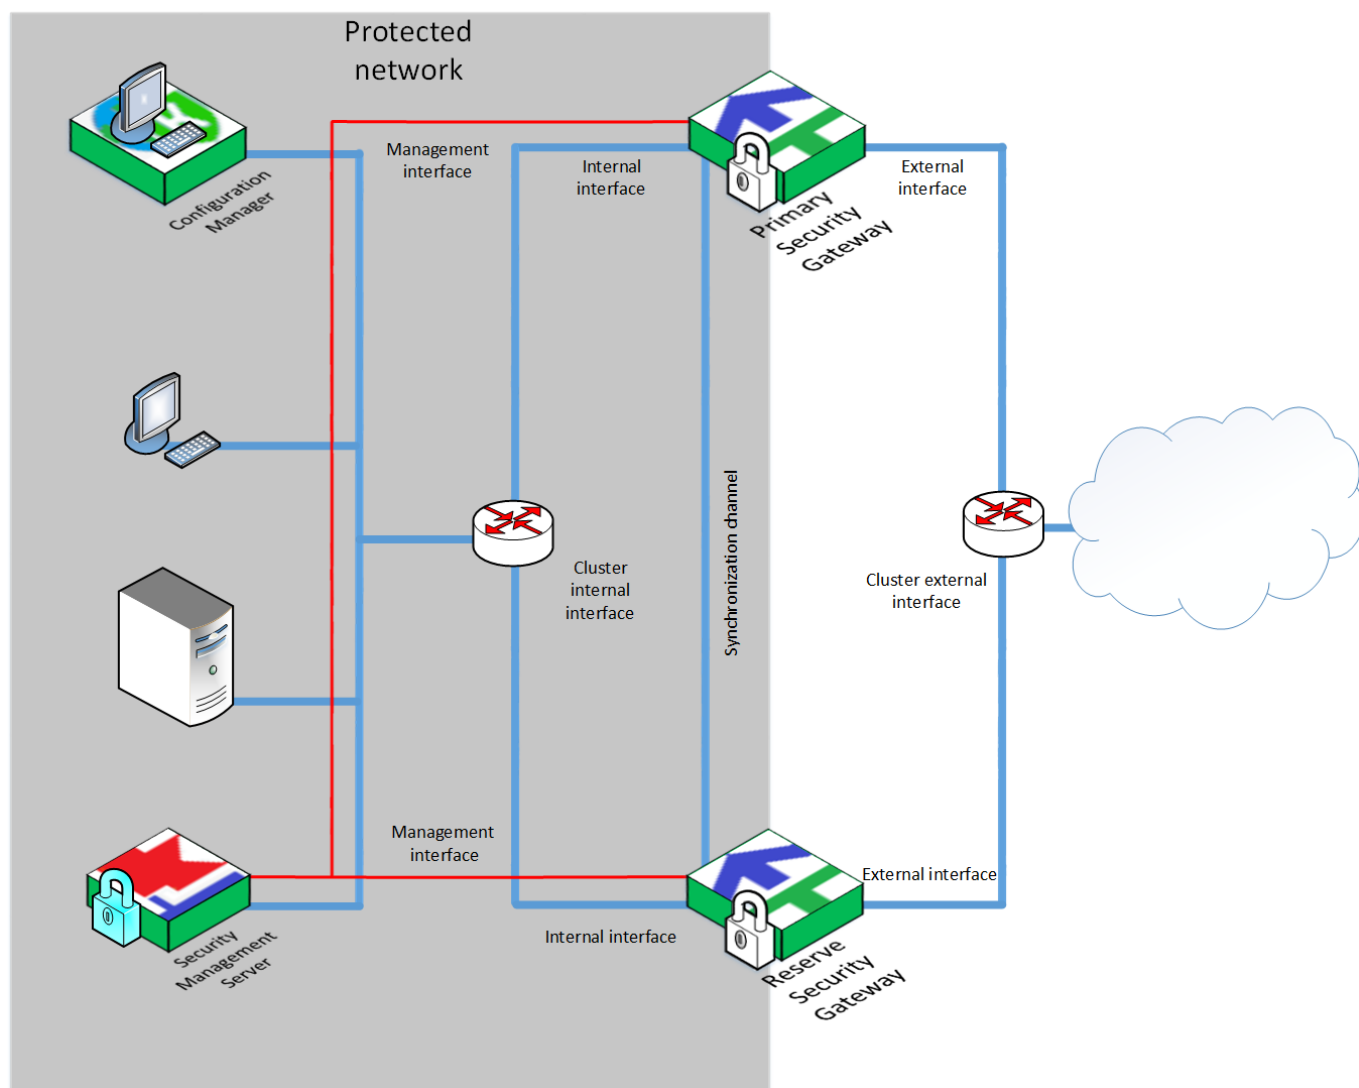
We recommend that you perform backup after changes in Continent settings.

In case of a Security Management Server failure, the Security Gateway of the Security Management Server is switched to a properly functioning one. To do so, the administrator restores the Security Management Server database from a previously created backup copy. The administrator can create a backup copy or restore the database via the Configuration Manager or the local menu.

Security Gateway failover

A failover Security Gateway cluster ensures Continent continuous operation in the event of a Security Gateway failure. If the active security cluster member loses its operability, the system switches to a standby Security Gateway and it automatically takes on functions of the primary Security Gateway.

A security cluster circuit diagram is shown in the figure below.



Note.

The following Continent components cannot be included in a security cluster:

- Security Gateway with the Security Management Server;
- Security Gateway in IPS mode.

Security cluster members have the following features:

- Main IP addresses are assigned to internal interfaces of security cluster members. In the security cluster, a secondary IP address is assigned to the respective internal interface of the active Security Gateway to provide traffic transfer.
- Main IP addresses are assigned to external interfaces of security cluster members. In the security cluster, a secondary IP address is assigned to the respective external interface of the active Security Gateway to provide traffic transfer.
- External interfaces of security cluster members form a virtual external interface of the security cluster.
- Each cluster member has its own IP address.

Note.

Secondary IP addresses of cluster interfaces are assigned during the cluster creation (see [4]). IP addresses of a cluster interface and its members belong to the same subnet.

- Each cluster member has its own MAC address.
- Being a cluster member, they operate in one of the states: **Active** or **Busy**.
- In the ARP cache of neighboring hosts, the virtual IP address of the cluster corresponds to the MAC address of the cluster member in the **Active** state.

The proper operation of the security cluster requires the following:

- Security cluster members must have:
 - license for synchronization;
 - licenses that support the same components;
 - same platforms (hardware);
 - same software versions;
 - synchronized system time.
- The licenses must support the same number of components.
- The synchronization, external and internal interfaces of both security cluster members must not be connected to the same network. Use a separate network for synchronization interfaces.

If a security cluster operates properly, an administrator can switch the state of a security cluster member to **Active**. Whether to assign this state to an element is defined by:

1. Security cluster member operability.
2. An administrator.
3. The position in the security cluster member list. The operative Security Gateway on the top of the list is the primary one. The position is only considered if you selected the **Auto-switch at primary Security Gateway recovery** check box in the security cluster properties.

If all synchronization channels between cluster members fail, the most operative Security Gateway on the top of the security cluster member list will automatically become **Active** even if previously it was **Standby** (see [5]). After fixing a failure, only the administrator can make a standby Security Gateway active.

An active cluster member is marked with  in the Security Gateway list of the Configuration Manager.

The security cluster member can operate in the following states:

- **OK** (in case of software update of a security cluster component — **updated OK**);
- **Attention** (an active Security Gateway that can process traffic but there are limitations or events the administrator should be informed about);
- **OK, not ready** (a working Security Gateway that cannot process traffic because of incorrect settings or inaccessible interfaces);
- **Problem**;
- **Down** (a deactivated security cluster Security Gateway that does not pass traffic but is currently working and can be configured);
- **Unavailable** (shut-down or disconnected from the communication channel);

The security cluster can operate in the following states:

- **OK;**
- **Attention** (both Security Gateways can process traffic);
- **Critical** (only one Security Gateway can process traffic);
- **Problem** (both Security Gateways cannot process traffic);
- **Offline** (both Security Gateways are down or unavailable).

In the Configuration Manager, a cluster is displayed as a structure combining its elements. Common parameters are set via a cluster, other parameters are set via its elements.

In the Security Gateway state table of the Configuration Manager and in the Continent monitoring system, you can view information on a security cluster operation.

Security Management Server failover


Security Management Server failover is provided by creating backups of critical data stored in the Security Management Server database with the activated role of the standby Security Management Server.

Security Management Server failover prevents data loss in the Security Management Server database, provides backups of certificate authority files in the event of a Security Management Server failure, provides Security Management Server work continuity in case of planned maintenance works that require switching off the Security Management Server.

The failover mechanism supports the Security Management Server database synchronization in real time without rebooting the Security Gateway.

One or several standby Security Management Servers are added to provide the failover of the corporate network management system. One of the standby Security Management Servers can be set to active in the event of an emergency and can take on all of the functions of the Security Management Server. All standby Security Management Servers must have the same operating system version.

A standby Security Management Server can be located within the same network with the active Security Management Server as well as within other protected network segment.



As a standby Security Management Server is added, every Security Management Server must be defined as:

- active;
- standby.

All management operations (for example, installing and editing a policy, editing users and objects) must be performed on an active Security Management Server.

A standby Security Management Server can be enabled in management in one of the two ways:

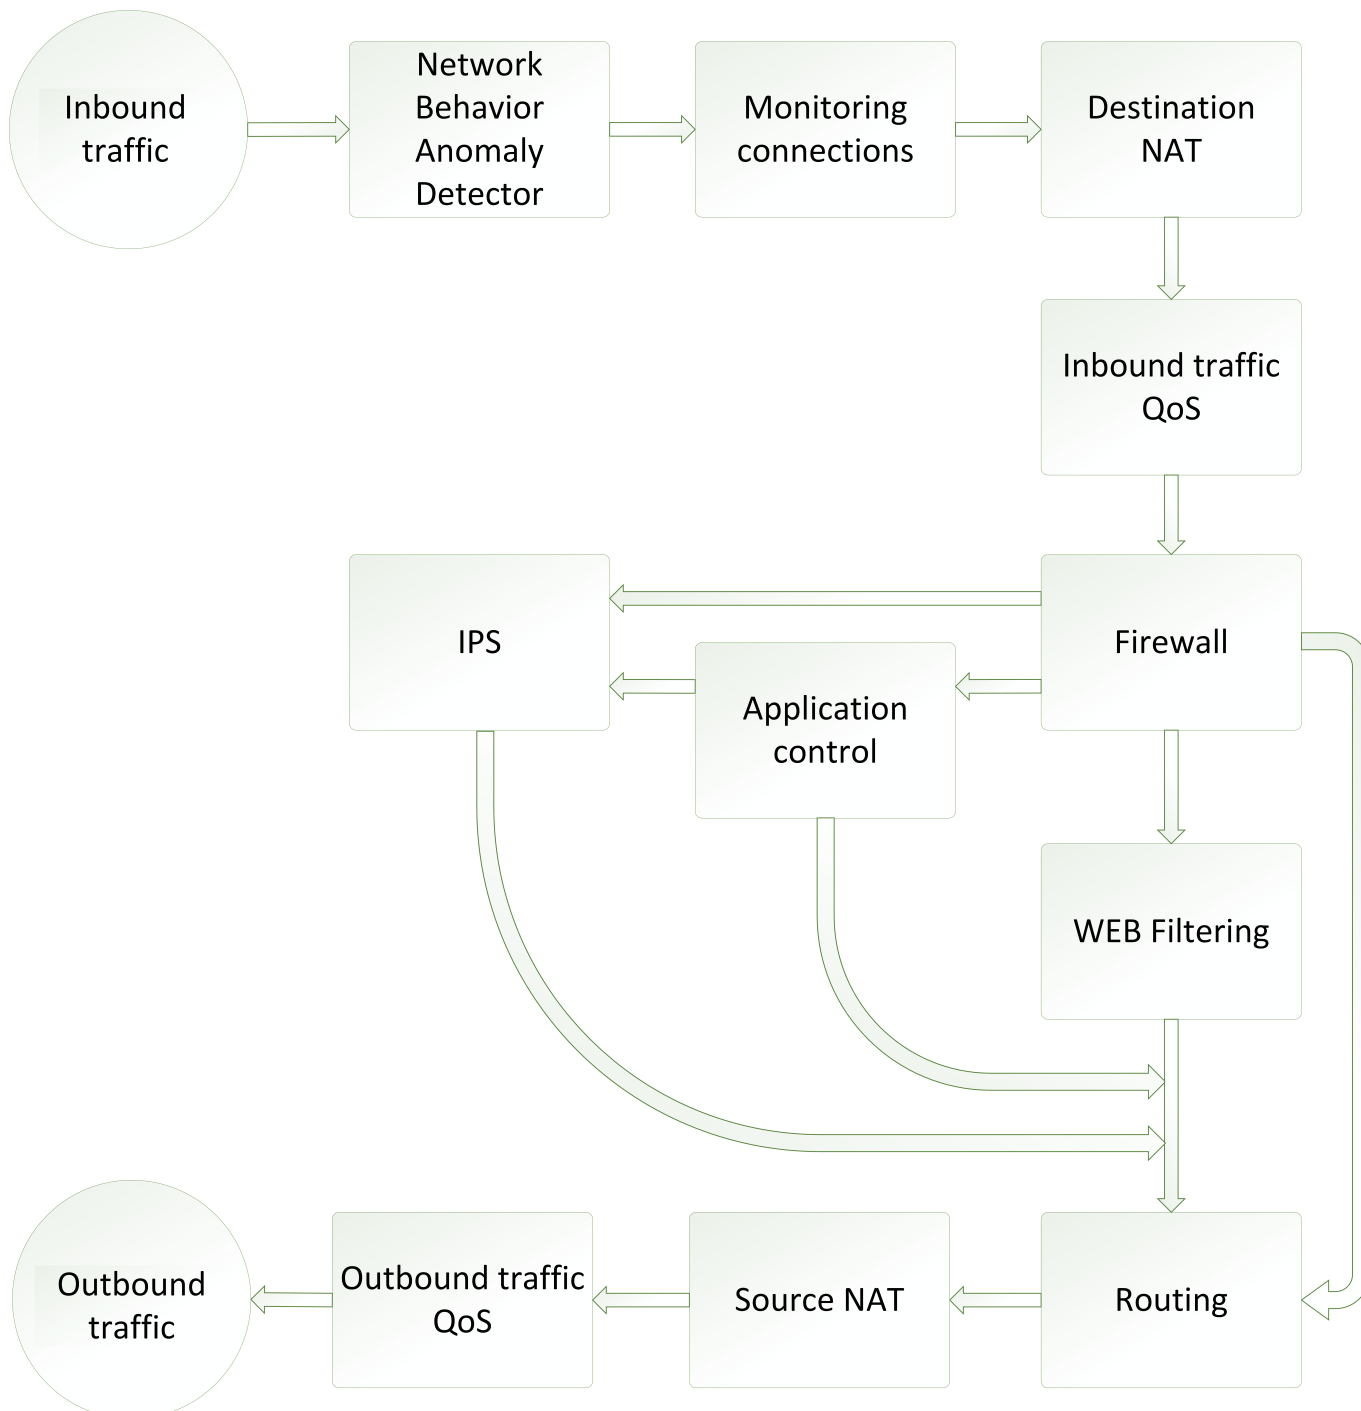
- a standby Security Management Server is always switched on, the Security Management Server database synchronization is performed by sending changes from the active Security Management Server to a standby one;
- a standby Security Management Server is always switched off, but it is switched on from time to time for the Security Management Server database synchronization and can be used as an active Security Management Server in the event of an active Security Management Server failure.

Each Security Gateway with the Security Management Server must be linked to the Security Management Server functions license for standby Security Management Server to work as a part of the management system.

Chapter 4

IP packets processing

Continent processes incoming IP packets according to the installed security and network service policies. The packet flow diagram is shown in the figure below.



Receiving packets

Ethernet frames are received by a Security Gateway network interface. Then, Ethernet header checksums are verified, the IPv4 is extracted and the header checksum is verified.

Network Behavior Anomaly Detector

The Network Behavior Anomaly Detector identifies and prevents scanning, protocol validation and DOS attacks. Network traffic analysis techniques taking into account traffic properties changes over time underpin the component's operation algorithm. The analysis is carried out using attack patterns. Attack patterns parameters are available for editing.

The Network Behavior Anomaly Detector analyzes both internal and external traffic as well as VPN tunnel traffic. To block traffic, the Network Behavior Anomaly Detector creates filtering rules.

If an attack is detected, the Network Behavior Anomaly Detector performs one of the following set of actions:

- registers the event in the network security log and displays the respective message in the Audit and Monitoring display area;
- registers the event in the network security log and temporarily blocks the attack source;
- collects statistical data.

Events related to the Network Behavior Anomaly Detector as a Security Gateway software component are registered in the system log. Events related to changes in the configuration of the Security Gateway or the Network Behavior Anomaly Detector are registered in the management log.

Connection monitoring

Connection monitoring is designed for identification of all packets that constitute traffic to process packets further. The collected data on connections are used during the Firewall and NAT operation.

Destination NAT

Request translation with destination address substitution.

NAT is used for transit IP address transformation. The features of IP packets for which address broadcasting is used are defined by NAT rules.

The tasks that are performed by this mechanism are described below.

Task	Rule
Hiding internal network structure behind one public address	Outgoing
Grant access to external public networks to users with non-unique intranet addresses	Outgoing
Provide access for third party resources to internal resources:	
• via set ports	Incoming
• via custom ports	Incoming
• via all ports (usually used to grant access to servers in the demilitarized zone)	1:1

Inbound traffic QoS

Inbound traffic prioritization for further processing on the Security Gateway.

Firewall

The Firewall protects network segments from unauthorized access.

You can configure the Firewall by creating lists of Firewall and NAT rules.

Firewall rules are used to control user access to external and protected internal networks by accepting or denying traffic to pass through the Security Gateway. NAT rules modify transit packet IP addresses.

Access control (Firewall and NAT) rules are grouped in lists in a strict order. A list defines a sequence of actions to perform with packets processed by Security Gateways. These lists are empty by default, and all traffic is denied by the Security Gateway except service packets.

Firewall rules are followed one by one without exceptions. If an IP packet matches rule parameters, it is processed with respect to the rule. The IP packet is not checked by other rules after that.

All IP packets passing through the Firewall are filtered. Packets are filtered twice: before and after passing through the cryptographic protection component.

IP packets are filtered according to rules that include source and destination IP addresses, a protocol name, UDP/TCP port numbers and network interface names. Time and the fact of authentication (for the protected segment) are checked as well. The packet contents are also analyzed during application protocol filtering. By default, all packets are denied if they are not explicitly allowed by filtering rules.

IP packets filtering rules are divided into two types:

- system — rules created by Continent automatically;
- custom — rules created by the administrator.

Rules are created automatically for the Security Gateway during the Security Management Server and the Security Gateway initialization. This type of rules is not available for management by the administrator.

The rules created by Continent automatically provide the following connections:

- between the Security Management Server and the Configuration Manager;
- between the Security Management Server and the registered Security Gateway;
- between the primary and the standby Security Gateway.

Filtering rules for other connections within a corporate network are created by the administrator.

If an IP packet does not meet the requirements, it is denied without notifying the sender.

Web filter

Continent provides the mechanism of additional filtration that enables you to analyze and process transit traffic at the level of some application protocols:

- HTTP;
- HTTPS;
- FTP.

Web/FTP filters are used in protection mechanisms against malicious web sites. Filters are grouped. Each filter is described by the following parameters:

- address — domain name of a destination server;
- scheme — web resources addresses scheme: FTP, HTTP or HTTPS;
- methods — methods of requesting data;
- content — list of MIME headers and file extensions;
- paths — names of web-pages, downloaded files or respective regular expressions;
- description — short description of a filter in any form.

Continent is supplied with pre-installed groups of filters produced by Kaspersky Lab and groups of SkyDNS vendor rules. They are forbidden to edit or delete.

The mechanism works as proxy, performs a man-in-the-middle attack by posing as the requested web resource for the request source and establishes a connection to the web resource on its own behalf. If HTTPS is in use for communication with a web resource, then this mechanism performs HTTPS inspection (traffic decryption) with certificate substitution. However, some web resources can use security mechanisms against HTTPS inspection, which makes it impossible to use the Web/FTP filtering mechanism as part of the Firewall. To keep access to such web resources, you can specify a set of exceptions for HTTPS inspection.

The Firewall comes with a built-in set of vendor exceptions. The administrator can create new exceptions and add them to this list.

Application control

When filtering, the Firewall provides step-by-step traffic processing. After checking an IP address, a protocol and a port, you can analyze traffic using application control mechanisms or protocol inspection. Application protocol inspection is performed only for TCP/UDP traffic. The destination port in server settings can be any, which allows determining a protocol on non-standard ports.

Managing connections

Managing connections means:

- Managing related connections tracking through certain protocols.
- Viewing and deleting connections created on a Security Gateway in the local menu.
- Configuring connection rematch after installing a policy in the Configuration Manager.

Related connection tracking

While Firewall is enabled, tracking of the related connections is performed for the following protocols:

- FTP;
- GRE;
- H.323;
- PPTP;
- SIP;
- TFTP.

Disabling related connections tracking is provided on the Security Gateways. That ensures the protocol inspection mechanisms. Disabling can be performed for the Firewall filtering rules and for the NAT rules as well.

IP address geographic filtering

Firewall provides inbound and outbound traffic blocking in accordance with source and destination countries. A country or a group of countries is specified as a source or a destination in Firewall rules. GeoProtection module specifies subnets that belong to the chosen countries for further filtering.

IPS

The Intrusion Prevention System analyzes network traffic to detect cyber attacks on the network level (L3 IPS).

Continent supports two operation modes for a Security Gateway with the enabled Intrusion Prevention System:

- **UTM** — packets are sent to the IPS after being processed by the Firewall. If you disable the Firewall on the Security Gateway with the enabled IPS, all traffic is automatically sent to the IPS. The IPS in **UTM** mode can be configured only in **Inline** mode (see below).
- **IPS** — the IPS does not modify packets. The Firewall is not enabled. The IPS appliance can be configured in **Monitor** and **Inline** modes.

The IPS can operate in the following modes:

- **Monitor**

In this mode, traffic is mirrored to the IPS from a SPAN port of a switch or a router.

If an attack is detected, the IPS appliance registers it and sends information about it to the Security Management Server.

- **Inline**

In this mode, the IPS appliance is placed between the Internet and a protected network. In case of a traffic analyzer software failure, the IPS appliance switches to the bypass mode for traffic to pass (when the respective option is enabled).

Traffic is captured and sent using physical interfaces. You can use several pairs of interfaces.

If an attack is detected, the IPS appliance registers it and drops malicious traffic if it is prescribed in an IPS policy. The Security Management Server receives information about the attack.

The IPS analyzes data using a signature method based on IPS protections. You should upload IPS protections to the Security Management Server, then include the required ones in the IPS profile. To apply the IPS protections included in the IPS profile to the IPS appliance, an administrator must create an IPS policy rule that includes the required IPS profile, then install it on the required IPS appliance.

The IPS profile contains a custom heuristic analyzer to control application traffic.

The Security Code IPS protections set is divided into groups by default. You cannot modify a single vendor IPS protection or the whole set. The IPS administrator can create and modify custom IPS protections and groups. You may use a vendor IPS protection as a template for a custom one.

Each IPS protection defines a counteraction way (alert, drop or pass) to an attack signature for each IPS profile separately. The IPS administrator can modify the attack counteraction way of the IPS protection or the IPS profile according to the IPS appliance operation mode. In **Monitor** mode, the IPS appliance can only notify the administrator about a detected attack. In **Inline** mode, the IPS appliance can counteract the attack in any existing way.

Routing

Static routing

A list of paths is displayed as a table where each line corresponds to one route.

If you want to modify network settings of an edge device, it is enough to change them in the routing table in advance.

Dynamic routing

The following dynamic routing protocol versions are supported:

- OSPF version 2;
- BGP version 4.

Protocol support is performed based on BIRD (Internet Routing Daemon).

To perform dynamic routing, create a BIRD configuration file. You can use any available text editor or the built-in configuration editing window in Security Gateway properties.

Source NAT

IP packet broadcasting with source address substitution.

Outbound traffic QoS

Outbound traffic prioritization.

Chapter 5

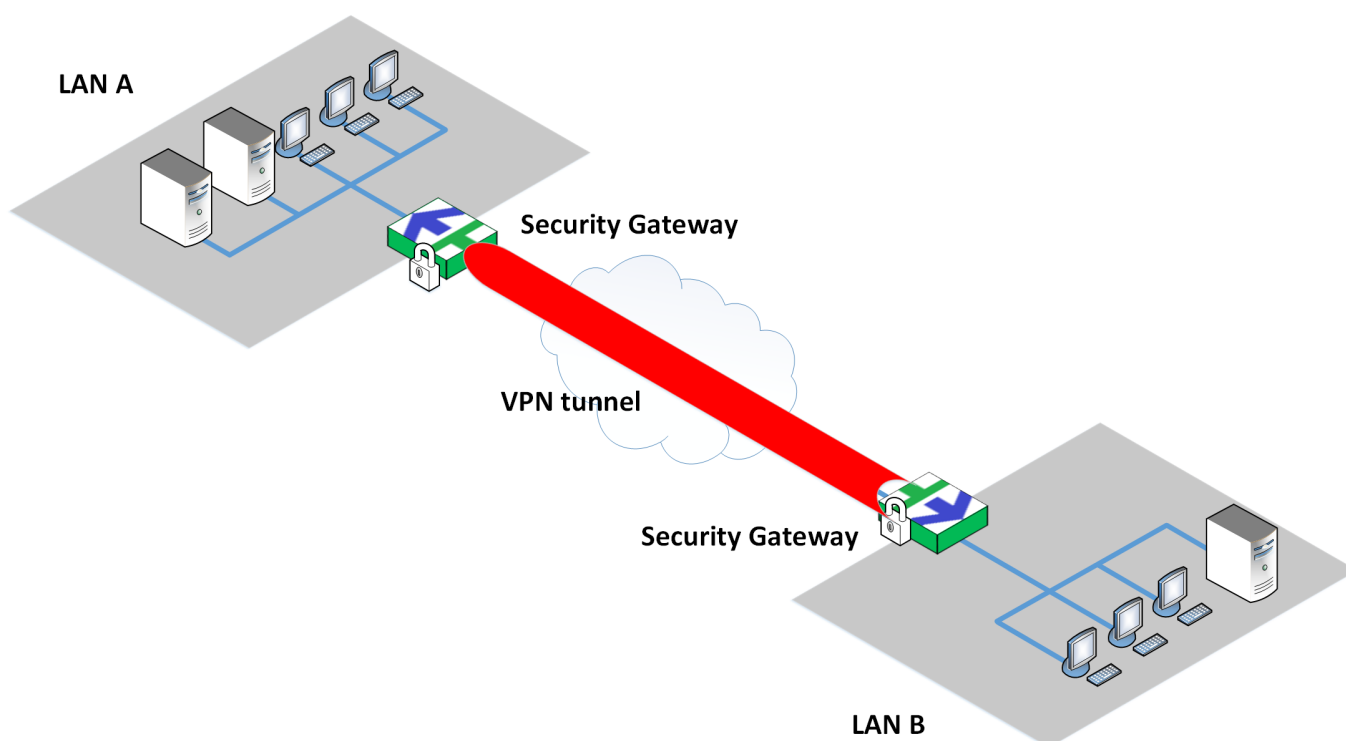
VPN principles

VPN technology makes it possible to combine local area networks, network segments and workstations into a secure virtual network over public TCP/IP networks. This technology is replacing dedicated channel communication, which allows companies to reduce operational costs. However, the use of public networks imposes additional requirements for information resource protection.

VPN tunnel

Using Continent VPN, you can create secure channels (tunnels) between two Security Gateways. Each Security Gateway that acts as the Firewall protects a LAN or a network segment. When traffic is transferred from one protected network to another one, it is encrypted before getting into the tunnel and decrypted after getting out of the tunnel.

In the figure below, there are two LANs. Each of them has its own Continent Security Gateway. They are connected using public data networks (for example, via the Internet).



Hosts of the **A** and **B** LANs exchange data using a tunnel between Security Gateways:

1. The Security Gateways are connected to each other and create a tunnel.
2. The **A** host sends packets to the **B** host.
3. The **A** Security Gateway encrypts the packets and sends them through the tunnel.
4. The **B** Security Gateway decrypts the packets and sends the packets to the destination host.
5. The **B** host receives the decrypted data.

Packets are transferred backwards in the same way — when the **B** host starts to exchange data with the **A** host. The packet encryption prevents unauthorized access when a third party intercepts it.

A VPN tunnel does not impose any additional requirements on the **A** and **B** users. Applications on their computers keep generating packets with respective source and destination addresses as usual. Security Gateways perform all operations for encrypting, encapsulating and transferring packets.

Encrypted data is transferred only within the tunnel between two Security Gateways. A host and a Security Gateway exchange unencrypted data.

One Security Gateway can maintain more than one tunnel at the same time and each tunnel can support more than one connection.

Encryption

In Continent, VPN uses symmetric-key cryptography. A connection between Security Gateways is based on the shared secret key mechanism. Each IP packet is encrypted using one packet encryption key based on a shared secret key.


Data encryption meets GOST R 34.12-2018 (Magma) requirements in Cipher Feedback Mode. Message authentication meets GOST R 34.12-2018 (Magma) requirements using a message authentication code.

Topology


You can create VPN tunnels between Continent Security Gateways and build a single VPN with the **star** or **full-mesh** topology.

In **full-mesh**, tunnels are created for each Security Gateway pair.

In **star**, there is a central Security Gateway. Other Security Gateways can be connected only to the central one.




Full-mesh



Star

You can see the mixed use of base topologies in the figure below.



The **1** and **2** Security Gateways form a full-mesh VPN segment, but the both are the central Security Gateways in the star segments.

Licenses

VPN is available only if all the Security Gateways (of the Continent domain) that transfer encrypted traffic are licensed including Security Gateways with the Security Management Server and Security Gateways with the Access Server. L3VPN must be marked in the licenses linked to all the Security Gateways including Security Gateways with the Security Management Server. Access Server must be marked in the licenses linked to Security Gateways with the Access Server.

For more information about license types and how to use them, see [4].

Chapter 6

Intrusion detection and prevention

All traffic passing through the Security Gateway in Continent can be analyzed for unauthorized access (network attacks). To do so, enable the IPS component.

The IPS analyzes the following traffic data:

- network address;
- port in use;
- packet field values;
- protocol identifiers;
- packet field sizes;
- traffic intensity.

The Intrusion Prevention System analyzes network traffic to detect cyber attacks on the network layer (L3 IPS).

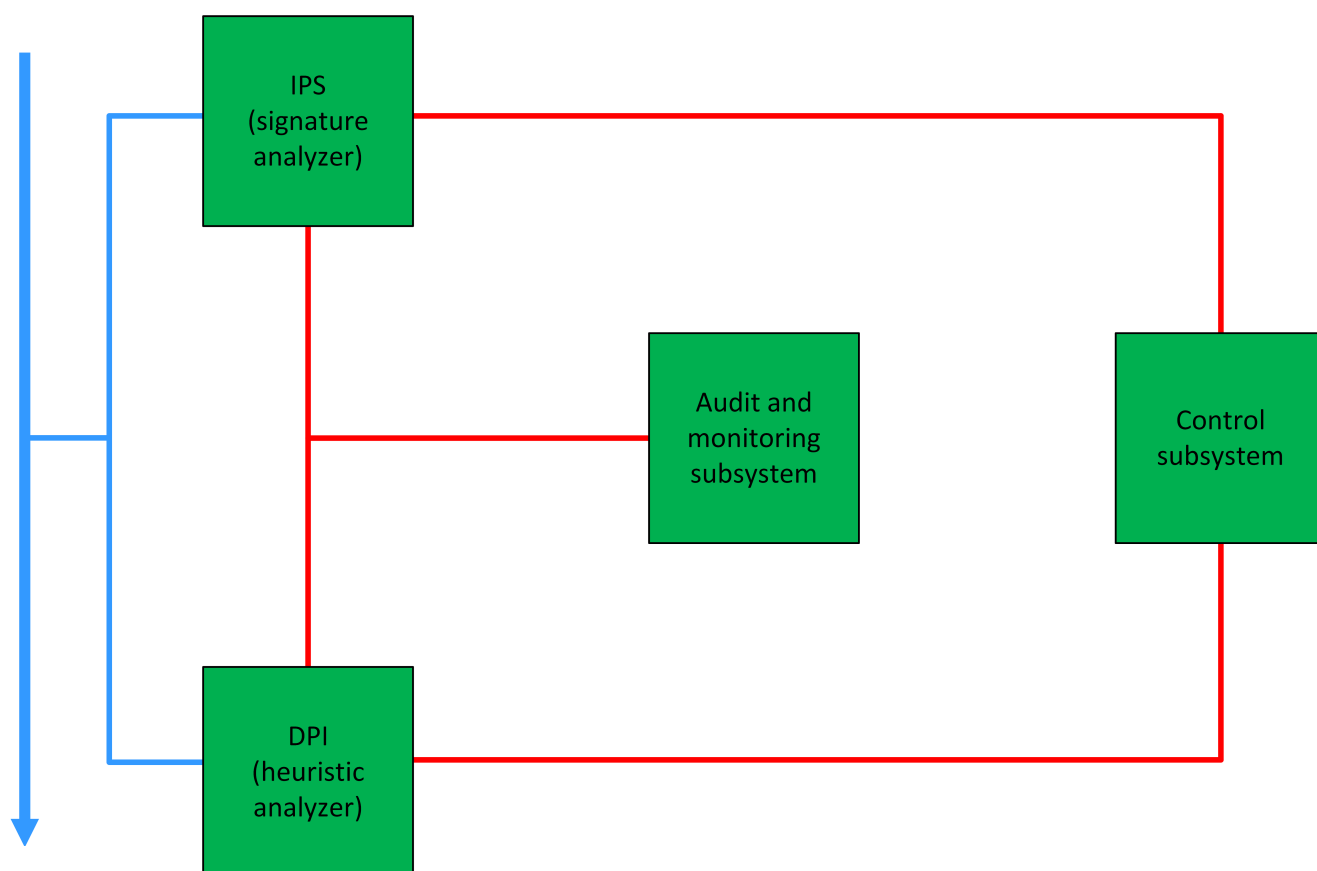
Continent supports two operation modes for a Security Gateway with the enabled Intrusion Prevention System:

- **UTM** — packets are sent to the IPS after being processed by the Firewall. If you disable the Firewall on the Security Gateway with the enabled IPS, the traffic is automatically sent to the IPS. The IPS in **UTM** mode can be configured only in **Inline** mode (see p. 30).
- **IPS** — the IPS does not modify packets. The Firewall is not enabled. The IPS appliance can be configured in **Monitor** and **Inline** modes.

The IPS can operate in the following modes:


- **Monitor**

In this mode, traffic is mirrored to the IPS from a SPAN port of a switch or a router. If an attack is detected, the IPS appliance registers it and sends information about it to the Security Management Server.



- **Inline**

In this mode, the IPS appliance is placed between the Internet and a protected network. In case of a traffic analyzer software failure, the IPS appliance switches to the bypass mode for traffic to pass (when the respective option is enabled). Traffic is captured and sent using physical interfaces. You can use several pairs of interfaces. If an attack is detected, the IPS appliance registers it and drops malicious traffic if it is prescribed in an IPS policy. The Security Management Server receives information about the attack.



The available options for using the IPS appliance in various Security Gateway operation modes are shown in the table below.

	Inline	Monitor
UTM	+	-
IPS	+	+

The IPS profile contains a custom heuristic analyzer to control application traffic.

The Security Code IPS protections set is divided into groups by default. You cannot modify a single vendor IPS protection or the whole set. The IPS administrator can create and modify custom IPS protections and groups. You may use a vendor IPS protection as a template for a custom one.

Chapter 7

Remote user access

Access Server

The Access Server supports Continent-RA secure connections initiated by a user.

The Access Server component interacts with the User Identification and the Firewall components.

User Identification is required because exchange parties must be authenticated.

The Firewall is a basic Security Gateway component that controls access of exchange parties to resources (both internal and external).

You can manage and configure the Access Server via the Configuration Manager.

TLS and VPN provide secure connections between Continent-RA and the Access Server. VPN tunnels are created based on cryptographic data gained during a TLS connection between Continent-RA and the Access Server.

Client authentication means user authentication. Users with or without domain credentials can access protected network resources remotely.

Domain user credentials are stored on a server and controlled by the AD. These credentials are read-only. For user authentication, a domain user name and password pair is used.

You can manage other credentials locally. In this case, when establishing a TLS connection, both server and client are authenticated (mutual authentication). The client is authenticated using a Continent-RA certificate. It is not necessary to use a user name and a password.

A user can be authenticated by using:

- domain credentials (one-way authentication when establishing a TLS connection, domain authentication with a user name and a password);
- local credentials (one-way or mutual authentication when establishing a TLS connection, domain authentication is not performed).

Upon successful authentication, a user is granted access to a resource according to remote access rules. The resource can be within a protected office network controlled by the Security Gateway with the Access Server as well as within a remote office network. If the resource is within a remote office network, access to it is provided by VPN Routing.

Transmitted data encryption and received data integrity control are performed using GOST cryptographic algorithms.

The Access Server is licensed according to the number of supported connections with Continent-RA. However, only a pre-established TLS connection is taken into account when using the tunneling UDP connection.

Continent-RA

Continent-RA is software that is installed on remote user workstations to provide them with access to protected network resources.

Continent-RA includes Security Code CSP, a management program for keys and a Continent-RA management program.

Continent-RA provides:

- secure connection with the Access Server;
- mutual authentication between the Access Server and Continent-RA when establishing a secure connection;
- secure connection with the Access Server over TCP;
- access to protected network resources when connected to the Access Server;
- tools to work with public key certificates;
- Access Server connections logging;
- connection events display and management.

Remote user access to protected network resources

The Access Server is enabled on the Security Gateway to provide remote users with access to corporate resources. Remote users in turn need to have Continent-RA installed on their workstations. The Access Server provides its functions only for protected networks of the Security Gateway on which it is installed.

A remote user registered on several Access Servers can connect to any of them using the same Continent-RA. Remote users communicate with internal resources over public networks. The Continent-RA cannot connect to several Access Servers simultaneously.

A remote user can initiate a connection between Continent-RA and the Access Server. Both a remote user and the administrator can break the connection. In some cases, the connection can be broken by the Access Server automatically.

After the connection is established, the Access Server uploads user's IP packet filtering rules to the Security Gateway IP packet filter. Besides, the list of protected subnetworks that are available for the user is sent to Continent-RA. After that, Continent-RA and protected network resources communicate via the Access Server. Besides, all traffic transmitted between Continent-RA and the protected network is encrypted using the GOST 28147-89 algorithm.

Continent-RAs can be connected to each other if they are assigned static IP addresses. Respective filtering rules must be created additionally to allow such connections.

All insecure connections with third parties (for example, with web-sites or FTP servers) can be prevented during communication with protected network resources.

Access Server management

You can manage the Access Server via the Configuration Manager that is installed on one or several workstations (administrator's workstations) within a protected network.

You can control the Access Server state and manage its database.

Chapter 8

Networking functions

QoS support

Continent supports the following QoS management mechanisms:

- traffic prioritization;
- real-time traffic jitter minimization;
- IP packet marking;
- VPN tunnel traffic prioritization management.

Traffic prioritization

Traffic prioritization is configured in the Security Gateway settings. There are 8 types of traffic prioritization:

- Real time;
- Top;
- High;
- Medium High;
- Medium;
- Medium Low;
- Low;
- Lowest.

For more information about traffic prioritization, see [8].

Real-time traffic jitter minimization

Note.

Jitter is a deviation from the true periodicity of a presumably periodic signal. Jitter is a significant and undesirable factor that must be taken into account during the design of communication channels.

Jitter minimization is given priority when streaming traffic and is usually achieved by assigning the traffic the top priority when it passes through network equipment.

Jitter required for IP telephony to function is achieved by assigning the real-time priority to traffic.

IP packet marking

An IP packet mark is defined using the DSCP tag in the IP packet header. You can automatically process the mark values in the following ways:

- save an existing value;
- change the DSCP tag value;
- reset the DSCP value (traffic is prioritized as **Lowest**).

Note.

The DSCP tag is compatible with IPP.

Traffic prioritization management

You can use the HFSC method for priority processing. It has the following additional parameters:

- proportion of total bandwidth allocated for each priority. If necessary, the specified value can be exceeded if the maximum bandwidth value is defined;
- maximum bandwidth set for each priority. The parameter value must be greater than or equal to the value specified for the total bandwidth proportion parameter. By default, it matches the guaranteed bandwidth.

Network device management using SNMP

Continent provides tools for network device management over SNMP, so that you can control, for example, the following parameters:

- network device operation time;
- the number of sent/received packets;
- interface status (**Up/Down**), etc.

Read requests to network devices are supported as well as sending service messages (**traps**). Traps are sent in the event of:

- cold start (**coldStart**);
- physical interface communication failure (**linkDown**);
- interface communication restoration (**linkUp**).

For more information about the module, see [9].

Routing

If there are other network devices between a protected network or the Security Management Server and a Security Gateway, specify routes for them.

Continent supports dynamic and static routing.

Static routing

Static routing makes routing decisions based on the routing table that is configured by the administrator for each Security Gateway. You can set routes both for IP addresses and for particular network objects.

The **metric** parameter defines the preference of a route. The lower the value of this parameter, the more preferable the route is.

Dynamic routing

Employing the BIRD daemon, Continent supports the following dynamic routing protocols:

- OSPFv2;
- BGP-4.

To configure dynamic routing, edit the BIRD configuration file and then load it to the Security Gateway.

Virtual routing

Continent supports VRF (Virtual Routing and Forwarding). You can create and use VRF zones as well as routing tables associated with them (kernel tables).

VRF zones are assigned interfaces of an individual Security Gateway. Therefore, isolated VRF zones are created within a Security Gateway, which makes it possible to process traffic including the one with overlapping address ranges.

DNS

In Continent, you can edit the DNS server list via the Configuration Manager as well as via the local menu.

ARP

ARP settings make it possible to view, change, update, delete record in the ARP table (in which IP addresses match MAC addresses), as well as configure ARP proxy. Depending on the record type (static or dynamic) operations are performed in the Configuration Manager and/or by the monitoring tools via the web console.

DHCP

You can use the Security Gateway as a DHCP server or DHCP relay. This allows assigning IP addresses and other network configuration parameters automatically to workstations within a protected network.

The DHCP service is available only for workstations of a network segment protected by the Security Gateway. In this case, a Security Gateway internal interface can hide only one domain.

The DHCP service can be disabled or operate in one of the following modes:

- **Server;**
- **Relay.**

Continent provides specifying pre-installed and configured DHCP server options.

If a workstation cannot connect to the DHCP server directly, then the **Relay** mode is used. The DHCP relay processes the broadcast DHCP request and redirects it to the DHCP server as a unicast packet. Then, it processes the DHCP server response and redirects it to the DHCP client.

The DHCP service is disabled on the Security Gateway by default.

You can manage the DHCP service remotely via the Configuration Manager.

Time synchronization

To set the synchronization mode with the NTP server, specify its name or the IP address. The Security Management Server can be used as the NTP server. Time synchronization is performed once every hour. You can also create a list of external accurate time servers. In this case, the most accurate one is chosen for synchronization.

You can set up system time synchronization locally as well as remotely.

By default, Continent components synchronize with each other over NTP (Security Gateways synchronize with the Security Management Server of the respective domain automatically).

SSH

The Security Gateway local menu can be accessed remotely over SSH. You can configure access for administrators via the Security Gateway local menu.

Bonds

Bonds are combinations of several physical interfaces into a logical one (a bond) to increase the bandwidth and enhance the connection reliability.

The administrator can create a bond in one of the following ways:

- by creating a new bond;
- by creating a VLAN interface using a bond;
- by adding a physical interface to a new bond.

A bond can include 1–8 physical interfaces.

Physical interfaces can be added to a bond. Physical interface settings do not influence the bond.

The following bond parameters can be set up:

- name;
- type;
- IP address;
- subnet mask;
- interfaces included;
- MTU;
- operation mode.

A bond can operate in the following modes;

- **active-backup.** Only one network bond is active. Other interfaces can become active only in case of an active interface failure. With this policy, you can see the bond MAC address only from one network port, so you can avoid problems with a switch.
- **802.3ad.** Bond network card groups are created with the same speed and duplex mode. With this bond, transmission uses all channels in the active bond according to the IEEE 802.3ad standard. The choice of interface for sending packets is defined by the default policy.

If you delete a bond, its settings are assigned to the physical interface with the highest priority. Other physical interfaces are set by default.

Collecting data on neighboring network devices

Continent enables network devices to receive information about the presence and characteristics from other network devices located in the same network and in turn send the same information about themselves. The LLDP protocol is used for data exchange.

Access to the received data on the neighboring devices is granted under the SMTP protocol. The collected data are displayed in the Continent monitoring subsystem.

The detection mechanism of the neighboring network objects is configured in the Configuration Manager, in the Security Gateway properties (**Configuration Manager | Structure | Security Gateway | Properties**), in **LLDP**.

To configure the neighboring network objects detection mode:


1. In the Configuration Manager, select the required Security Gateway and click **Properties** on the toolbar.
2. On the left, select **LLDP**.

The **Network Device Discovery** section appears.

If the LLDP component was not configured earlier or was disabled, the **Network Device Discovery** parameters will be unavailable.

3. Turn on the toggle.

The **Network Device Discovery** parameters become available for editing.

4. Specify the general detection parameters:
 - **Transmit interval (seconds)** — time period of sending data on Security Gateways to neighboring devices.
 - **Hold multiplier** — parameter that defines the lifetime together with the transmit interval (TTL). TTL is the product of multiplication of the transmit interval and the hold multiplier.
5. If event registration of network device detection under the LLDP protocol is required, select **Event logging**.
6. In **Optional TLVS**, select the required options if additional data on Security Gateways are to be sent to the neighboring network objects.
7. In **Interfaces**, add the interfaces that detect neighboring network objects by clicking .
8. For each interface, specify the operation mode: **Receive/Transmit/Receive and transmit**.
9. Click **Apply** and **OK** consistently after configuring the required parameters.

Chapter 9

Audit and monitoring

Continent audit and monitoring system allows you to monitor Security Gateways parameters. The system provides the following:

- logging and audit of security, management and system events;
- centralized monitoring of Security Gateways.

Data about Security Gateways events is logged and stored on Security Gateways and sent to the Security Management Server. Continent logs system, network security and management events. In each log, you can filter the records to find the required ones. The system log contains data about subsystem events. The network security log contains data about IPS, Firewall and unauthorized access events. The management log contains data about administrator and user actions.

Managing SMTP Server configuration, a list of users and email alerts about policy installation on Security Gateway is performed in the Configuration Manager.

The audit administrator performs an audit. Audit tasks are as follows:

- view logs regularly;
- configure log storage;
- manage log records.

How audit and monitoring works

The monitoring objects are:

- Security clusters;
- Security Gateways;
- Security Gateway groups.

A user with the right to access **Group management** can create new groups and move unsorted Security Gateways into them as well as move Security Gateways from group to group.

The audit and monitoring system supports the following information types:

- events;
- data;
- status.

The information type and source are the parameters that are used for information display about the objects of monitoring.

The sources for each information type are shown in the table below.

Information type	Source
Events	Management system Monitoring and audit system Integrity control Access control Application control Firewall Application filtration Secure communications Intrusion Prevention System Remote access Computing platform VPN
Data	Network interfaces Monitoring and audit system IPS signature triggering
Status	Monitoring and audit system

To display data about an object, you need to create an object monitoring rule.

Continent supports four types of monitoring rules:

- cluster rule — for Security Gateways within a cluster;
- Security Gateway rule — for the respective Security Gateway;
- group rule — for all Security Gateways in the group and subgroup;
- common rule — for all the Security Gateways and Security Gateway groups.

A template is a rule or a set of rules that is applied to Security Gateways, their groups, and determines event triggers.

A trigger priority depends on the rule type. A cluster rule has the highest trigger priority and precedes Security Gateway and group rules. A common rule has the lowest priority.

An object is displayed with its status. The status can be one of the following values:

Status	Description
Critical	This status denotes that a critical level event occurred. To change the status, choose the Closed status, i.e. change the parameter caused the event in accordance with the security policy
Warning	This status denotes that an event of the respective level occurred. The object remains in the status until you choose the Closed status or until the status does not change to critical
Info	This status denotes that an information level event occurred. The status remains until the parameter is changed

Documentation

1. Continent Enterprise Firewall. Version 4. Administrator guide. Deployment.
2. Continent Enterprise Firewall. Version 4. Administrator guide. Firewall.
3. Continent Enterprise Firewall. Version 4. Administrator guide. Intrusion Prevention System.
4. Continent Enterprise Firewall. Version 4. Administrator guide. Management.
5. Continent Enterprise Firewall. Version 4. Administrator guide. Monitoring and Audit.
6. Continent Enterprise Firewall. Version 4. Administrator guide. User Authentication.
7. Continent Enterprise Firewall. Version 4. Administrator guide. VPN.
8. Continent Enterprise Firewall. Version 4. Administrator guide. Networking functions.
9. Continent Enterprise Firewall. Version 4. Administrator guide. SNMP.
10. Continent Enterprise Firewall. Version 4. Administrator guide. Installation and Update.