IMPERVA®

SECURESPHERE®

on Amazon AWS

On-Demand

Configuration Guide

Version 10.5

August 2015

Copyright Notice

© 2015 Imperva, Inc. All Rights Reserved.

Follow this link to see the SecureSphere copyright notices and certain open source license terms:

https://www.imperva.com/sign_in.asp?retURL=/articles/Reference/SecureSphere-License-and-Copyright-Information
This document is for informational purposes only. Imperva, Inc. makes no warranties, expressed or implied.

No part of this document may be used, disclosed, reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form or by any means without the written permission of Imperva, Inc. To obtain this permission, write to the attention of the Imperva Legal Department at: 3400 Bridge Parkway, Suite 200, Redwood Shores, CA 94065.

Information in this document is subject to change without notice and does not represent a commitment on the part of Imperva, Inc. The software described in this document is furnished under a license agreement. The software may be used only in accordance with the terms of this agreement.

This document contains proprietary and confidential information of Imperva, Inc. This document is solely for the use of authorized Imperva customers. The information furnished in this document is believed to be accurate and reliable. However, no responsibility is assumed by Imperva, Inc. for the use of this material.

TRADEMARK ATTRIBUTIONS

Imperva and SecureSphere are trademarks of Imperva, Inc.

All other brand and product names are trademarks or registered trademarks of their respective owners.

PATENT INFORMATION

The software described by this document is covered by one or more of the following patents:

US Patent Nos. 7,640,235, 7,743,420, 7,752,662, 8,024,804, 8,051,484, 8,056,141, 8,135,948, 8,181,246, 8,392,963, 8,448,233, 8,453,255, 8,713,682, 8,752,208, 8,869,279 and 8,904,558.

Imperva Inc.

3400 Bridge Parkway, Suite 200

Redwood Shores, CA 94065

United States

Tel: +1 (650) 345-9000
Fax: +1 (650) 345-9004

- **Website**: http://www.imperva.com
- **General Information**: info@imperva.com
- **Sales**: sales@imperva.com
- **Professional Services**: consulting@imperva.com

Technical Support: support@imperva.com

Imperva-SecureSphere-v10.5-Amazon-AWS-On-Demand-Configuration-Guide-Patch-14-v1

End User License and Services Agreement

BY CLICKING ON THE "ACCEPT" BUTTON, TAKING AN ACTION TO INDICATE ACCEPTANCE, OR USING THE PRODUCTS (AS DEFINED BELOW) END USER AGREES TO THE TERMS OF THIS END USER LICENSE AND SERVICES AGREEMENT ("AGREEMENT") WITH IMPERVA, INC. ("IMPERVA"). IN THE EVENT THE INDIVIDUAL IS ENTERING INTO THIS AGREEMENT ON BEHALF OF A CORPORATE OR OTHER PUBLIC OR PRIVATE ENTITY, END USER REFERS TO THAT ENTITY, AND SUCH INDIVIDUAL CERTIFIES THAT HE/SHE IS AN AUTHORIZED REPRESENTATIVE OF THE END USER. IF END USER DOES NOT AGREE TO ALL OF THE TERMS OF THIS AGREEMENT, CLICK THE "CANCEL" BUTTON, DISCONTINUE THE SET-UP AND INSTALLATION OR DISCONTINUE USE OF THE PRODUCT.   IF THE TERMS OF THE AGREEMENT ARE CONSIDERED AN OFFER, ACCEPTANCE IS EXPRESSLY LIMITED TO THESE TERMS.

1. **Definitions**.   The following capitalized terms shall have the meanings set forth below:

    a) "**Appliance**" means the Imperva branded computer hardware on which the Software operates.

    b) "**Delivery**" shall mean, (i) in the case of Software, when the Software is made available by Imperva for End User to electronically download or use on Amazon Web Services; (ii) in the case of Subscription Services, when the Subscription Service has been provisioned and made available to End User to access; and (iii) in the case of an Appliance, when the Appliance has been tendered by Imperva for shipment.

    c) "**Documentation**" means Imperva's technical specifications that accompany and describe the installation, use and operation of a Product.

    d) "**End User**" means the party that has purchased the Products for its own use, either directly from Imperva or through an authorized third party.

    e) "**Licensed Volume**" means the volume or other measurement of permitted use for the Products as agreed to by Imperva directly or through its resellers.

    f) "**Open Source Software**" means third party software that Imperva distributes with the Software pursuant to a license that requires, as a condition of use, modification and/or distribution of such software, that the software or other software combined and/or distributed with it be (i) disclosed or distributed in source code form; (ii) licensed for the purpose of making derivative works; or (iii) redistributable at no charge.

    g) "**Products**" mean Appliances, Software or Subscription Services, as the case may be.

    h) "**Services**" mean Professional Services or Support, as the case may be.

    i) "**Software**" means Imperva's or its licensors' software (in object code format) or content, any updates or upgrades thereto provided to End User by Imperva and any Documentation pertaining thereto.   Software may be delivered to End User on Appliances or on a standalone basis.   The term "Software" does not include Open Source Software.

    j) "**Subscription** Services" mean the subscription services, including content, updates and upgrades thereto, that may be made available to End User by Imperva directly or through its partners and suppliers.   Subscription Services include, without limitation, the ThreatRadar service and the Incapsula services.

    k) "**Support**" means the technical support and maintenance services for the Product and periodic bug fixes and updates to the Software that Imperva may make generally available at an annual subscription cost to end users.

    l) "**Professional Services**" mean the installation, configuration, and training services that Imperva may provide to an End User.

2. **Licenses and Restrictions**.

    a) **Software**.   Conditioned on the terms and conditions of this Agreement, Imperva grants End User a perpetual (subject to Section 11 or any term license restrictions), nonexclusive, nontransferable, nonsublicensable license to use the Software in accordance with its Documentation only for End User's internal business purposes on the Appliances or in the Licensed Volume licensed by End User.   If End User purchases Software on a standalone basis, the license granted herein shall include the right to copy the Software up to the Licensed Volume.

    b) **Subscription Services**.   Conditioned on the terms and conditions of this Agreement and any registration form that Imperva may require, Imperva grants End User a nonexclusive, nontransferable, nonsublicensable, revocable right to use and access the Subscription Services in accordance with its Documentation only for End User's internal business purposes.

c) **Restrictions**. End User may not (and may not permit any third party to): (i) modify, incorporate or use in any other works, translate, reverse engineer (except to the limited extent applicable statutory law expressly prohibits reverse engineering restrictions), decompile, disassemble, otherwise attempt to derive source code from or create derivative works based on the Products; (ii) make unauthorized copies of the Products; (iii) disclose, distribute, transfer or market the Products to third parties; (iv) remove or modify any proprietary notices, labels or marks on or in any copy of the Products; (v) distribute, sell, sublicense, rent, lease or use the Products (or any portion thereof) for time sharing, hosting, service provider or other computer services to third parties or otherwise make the functionality of the Software available to third parties; (vi) publicly disseminate performance information or analysis (including, without limitation, benchmarks and performance tests) from any source relating to the Products; (vii) access the database or any other third party product that is embedded in the Software with applications other than the Software; or (viii) use the Products other than as permitted herein.

d) **Appliance**. End User acknowledges that the Software included with the Appliance is licensed and not sold. Such Software is licensed solely in conjunction with such Appliance (and not separately or apart from such Appliance). If End User sells, leases, lends, rents, distributes or otherwise transfers any Appliance to any third party or if Imperva terminates this Agreement under Section 11.b or for a breach of Section 2.c, then End User will erase all Software from such Appliance.

e) **Open Source Software**. Open Source Software is copyrighted and licensed under the GPL/LGPL and other licenses. Copies of or references to those licenses are included with Software in the "Help" section. If delivery of source code is required by the applicable license, End User may obtain the corresponding Open Source Software source code for a period of three years after Imperva's last shipment of the Software, by sending a request to: Legal Department – Open Source Software Request, Imperva, Inc., 3400 Bridge Parkway, Suite 200, Redwood Shores, CA 94065, United States.

f) **US Government End User**. For purposes of this Agreement, "commercial computer software" is defined at FAR 2.101. If acquired by or on behalf of a civilian agency, the U.S. Government acquires this commercial computer software and/or commercial computer software documentation and other technical data subject to the terms of the Agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the Federal Acquisition Regulation ("FAR") and its successors. If acquired by or on behalf of any agency within the Department of Defense ("DOD"), the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of the Agreement as specified in 48 C.F.R. 227.7202-3 of the DOD FAR Supplement ("DFARS") and its successors. This U.S. Government End User Section 2(f) is in lieu of, and supersedes, any other FAR, DFARS, or other clause or provision that addresses government rights in computer software or technical data.

3. **Support and Subscription Services**. Provided End User has an active and fully paid contract for Support and has completed any required registration, Imperva will provide Support in accordance with its standard Support terms then in effect. Provided End User has an active and fully paid contract for Subscription Services, Imperva will provide such services in accordance with its standard Service Level Agreement (SLA) then in effect.

4. **Additional Terms for Subscription Services**.

a) **Accessing and Use of Subscription Services**. Except as explicitly set forth herein, End User is solely responsible for acquiring and maintaining all of the equipment, software, services and items necessary to access and make use of the Subscription Services, including without limitation paying all charges, taxes, and other costs and fees related to internet access. End User may access the Subscription Services only through the interfaces and protocols provided or authorized by Imperva and its partners, and agrees to set up, maintain and use the Subscription Services in strict compliance with Imperva's and its partners' instructions. End User is solely responsible for maintaining the confidentiality of any passwords and account information required to access Subscription Services, for all acts that occur in connection with End User's account and to immediately notify Imperva of any unauthorized use of End User's account. In the event of expiration or termination of any Subscription Services that require DNS routing, End User will be solely responsible for rerouting its DNS traffic and Imperva, its partners and suppliers shall have no liability for End User's failure to do so.

b) **Authorization**.   Certain Subscription Services are offered to cache, monitor and optimize web pages and web sites. In connection with such Subscription Services, End User hereby grants Imperva and its partners a nonexclusive, worldwide, fully paid-up, royalty-free license to use, transfer, display, minimize and compress the content and material on End User web sites ("**End User Content**"), in any media formats, solely as necessary for the performance of the Subscription Services.   Imperva and its partners do not provide backup services for End User Content, and if End User's use of the Subscription Services terminates for any reason,   Imperva and its partners may, without notice, delete or deny End User access to any of content or meta data that may remain in its/their possession or control.   In addition, End User agrees that if, at Imperva's and its partners' sole determination, End User is using the Subscription Services in a manner that violates laws, rules or regulations or creates an excessive burden or potential adverse impact on Imperva's, its partners' or its suppliers' systems, business or customers, Imperva, its partners or its suppliers may suspend or terminate End User's access to the Subscription Services without notice to or liability to End User.

5. **Professional Services**.   Professional Services, if any, to be provided by Imperva to an End User will be subject to a separate statement of work ("**SOW**") agreed to by Imperva and Imperva's standard Professional Services terms then in effect.

6. **Fees, Payment Terms, Shipment and Delivery**.

   a) For orders accepted directly by Imperva, End User shall pay Imperva the applicable fees designated by Imperva. Overage fees may apply if End User exceeds its allowance for any Subscription Services at Imperva's then-current overage rates. Any fees payable to Imperva are non-refundable and payable in US Dollars. End User shall also pay all sales, use, value-added and other taxes, tariffs and duties of any type assessed against End User, except for taxes based on Imperva's income. Should End User be required under any law or regulation of any governmental entity or authority outside of the United States, to withhold or deduct any portion of the payments due to Imperva, then End User shall increase the sum payable to Imperva by the amount necessary to yield to Imperva an amount equal to the sum it would have received had no withholdings or deductions been made.   Fees shall be invoiced as follows: (a) fees for all Subscription Services, term Software licenses and Support shall be invoiced at the time of the initial order and in advance of each renewal period; (b) fees for other Software licenses and Appliance purchases will be invoiced upon Delivery.   All payments from End User to Imperva are due net thirty (30) days after the date of invoice. If End User's account for Subscription Services or Support is thirty (30) days or more overdue, in addition to any of its other rights or remedies, Imperva reserves the right to suspend such services to End User, without liability to End User, until such amounts are paid in full. Imperva shall have the right to conduct and/or direct an independent accounting firm to conduct, during normal business hours, an audit of End User's facilities, computers and records to confirm End User's use of Products is in compliance with this Agreement.   End User shall provide reasonable cooperation with any such audit.

   b) Imperva will use commercially reasonable efforts to ship the Appliances and the Software license keys at the times requested in Orders (in partial or full shipments); provided, however, that Imperva shall in no event be liable for any delay in Delivery or for failure to give notice of delay. Without liability to any person and without prejudice to any other remedy, Imperva may withhold or delay shipment of any Order if End User is late in payment or is otherwise in default under this Agreement.   Title to Appliances and risk of loss shall pass to End User upon Delivery, and Products shall be deemed accepted by End User upon Delivery.   Appliances shall be delivered Ex Works (Incoterms 2000) Imperva's designated manufacturing facility. End User may specify shipping instructions with the Order. In the absence of specific shipping instructions from End User, Imperva will ship by the method it deems most advantageous.   End User shall pay and be exclusively liable for all costs associated with shipping and delivery including without limitation, freight, shipping, customs charges and expenses, cost of special packaging or handling and insurance premiums incurred by Imperva in connection with the shipment of Appliances to End User. In its discretion, Imperva may advance shipping charges on behalf of End User on Appliances shipped to End User, and End User agrees to reimburse Imperva for any such advanced charges and expenses.

7. **Confidentiality**.

a) As used herein, "Confidential Information" means all confidential and proprietary information of a party ("**Disclosing Party**") disclosed to the other party ("**Receiving Party**"), whether orally or in writing, that is designated as confidential or that reasonably should be understood to be confidential given the nature of the information and the circumstances of disclosure. Confidential Information includes, without limitation, the Products, their performance (including any benchmarking information) and Imperva's pricing for direct sales of its Products and Services. Confidential Information shall not include any information that: (i) is or becomes generally known to the public without breach of any obligation owed to the Disclosing Party; (ii) was known to the Receiving Party prior to its disclosure by the Disclosing Party without breach of any obligation owed to the Disclosing Party; (iii) was independently developed by the Receiving Party without breach of any obligation owed to the Disclosing Party; or (iv) is received from a third party without breach of any obligation owed to the Disclosing Party.

b) The Receiving Party agrees that it will (i) use Confidential Information for the sole purpose of exercising its rights and performing its obligations under this Agreement, (ii) divulge Confidential Information only to those of its employees, contractors, directors, independent consultants and agents who have a need to know such Confidential Information and who are bound by professional duty or in writing (in advance) to confidentiality and non-use obligations at least as protective of such information as this Agreement, and (iii) not disclose any Confidential Information to any third party. The Receiving Party shall notify and cooperate with the Disclosing Party immediately upon discovery of any unauthorized use or disclosure of Confidential Information of the Disclosing Party. The Receiving Party may disclose Confidential Information to comply with an order from a court of competent jurisdiction or with a mandatory requirement of a governing regulatory body, provided such party, to the extent permitted by law and as soon as reasonably practicable under the circumstances, informs the Disclosing Party and allows the Disclosing Party the opportunity to object to the disclosure order or to take action to preserve the confidentiality of the information. The Receiving Party shall cooperate with the Disclosing Party in such party's reasonable efforts to limit the disclosure of the information. End User acknowledges, understands and agrees that Imperva may, as part of its provision of the Product and/or Services to End User, collect, store and use information obtained from End User, including, but not limited to, information about End User's users and customers ("**Information**") for the purposes of the provision of the Product, Services and other services to End User and may use such information for analysis and improvement of Imperva's products and services. End User represents and warrants that it has all rights and permissions necessary to transfer such Information and to grant Imperva access to such Information as contemplated herein.

c) Upon termination of this Agreement for any or no reason, the Receiving Party shall (i) immediately cease all use of the Disclosing Party's Confidential Information, (ii) upon request from the Disclosing Party, either promptly destroy all Confidential Information of the Disclosing Party or return all Confidential Information of the Disclosing Party. Notwithstanding the foregoing, both parties agree that the Receiving Party may retain copies of the Confidential Information that are necessary for (i) satisfying legal or regulatory requirements; and (ii) archiving consistent with good business practices, provided that the Receiving Party's obligations of confidentiality and restricted use under this Agreement shall continue with respect to such retained copies.

d) If the Receiving Party discloses or uses (or threatens to disclose or use) any Confidential Information of the Disclosing Party in breach of this Section 7, the Disclosing Party shall have the right, in addition to any other remedies available to it, to seek injunctive relief to enjoin such acts, it being specifically acknowledged by the parties that any other available remedies are inadequate.

8. **Proprietary Rights; Indemnity**.

a) All title and intellectual property rights in and to the Products are owned exclusively by Imperva and its partners and suppliers. Other than as expressly set forth in this Agreement, no license or other rights in or to the Products and intellectual property rights thereto are granted to End User, and all such licenses and rights are hereby expressly reserved. Any ideas, suggestions, modifications and the like made by End User with respect to a Product will be the property of Imperva regardless of whether Imperva chooses to exercise its rights to incorporate such ideas, suggestions or modifications into the Product.

b) Subject to the remainder of this Section 7.b, Imperva will indemnify End User from any Liability (as defined below) to a third party resulting from infringement of a U.S. patent or any copyright, or misappropriation of any third party trade secrets by the Product as delivered ("**Infringement Claim**"); provided that End User (1) promptly notifies Imperva of any and all threats, claims and proceedings of such Infringement Claim, (2) gives reasonable assistance in response to Imperva's request for assistance, and (3) grants Imperva sole control over defense and settlement thereof. For purposes of this section "Liability" means the resulting costs (including reasonable attorneys' fees) and damages awarded against End User to the third party making such Infringement Claim, by a court of competent jurisdiction or agreed in settlement.   The foregoing obligations do not apply with respect to Products or portions or components thereof, (i) that are modified after delivery by Imperva, (ii) combined with other products, processes or materials, where the alleged infringement relates to such combination, (iii) where End User continues allegedly infringing activity after being notified thereof and modifications that would have avoided the alleged infringement have been made available to End User, or (iv) where End User's use of such Product is not strictly in accordance with this Agreement.   In the event that Products are held to or believed by Imperva to infringe, Imperva at its discretion, will have the option to (A) modify the allegedly infringing Products to be non-infringing, (B) obtain for End User a license to continue using the Products, or (C) request the return of the Product and upon receipt thereof terminate this Agreement as to the infringing Product and refund to End User the unused portion of the fees paid under this Agreement for such Subscription Services or a pro rata portion of fees paid for the infringing Software or Appliance, depreciated on a straight-line basis over a three (3) year period.   End User will defend, indemnify and hold Imperva harmless against any claims, damages settlements and expenses (including attorneys' fees) excluded from Imperva's indemnity obligations in (i) – (iv) above.   THIS SECTION SETS FORTH IMPERVA'S SOLE OBLIGATION AND END USER'S SOLE AND EXCLUSIVE REMEDY IN THE EVENT OF VIOLATION OF THIRD PARTY RIGHTS.

9. **Warranty and Disclaimer**.
   a) Imperva warrants that during the sixty (60) day period commencing on the date of first Delivery, the Software and the Appliances will perform substantially in accordance with their Documentation.   In the event of a breach of the foregoing warranty with respect to the Software, as End User's sole and exclusive remedy, Imperva shall, at its sole expense, use reasonable efforts to modify the Software, so that it performs substantially in accordance with its Documentation.   In the event of a breach of the foregoing warranty with respect to an Appliance, as End User's sole and exclusive remedy, Imperva shall, at its sole expense and at its option, repair the Appliance or replace the Appliance with a new or reconditioned Appliance that performs materially in accordance with its Documentation.   The foregoing warranty extends only to the original purchaser and will not apply to errors or damage caused by misuse of the Products.   The rights and remedies granted End User under this Section state Imperva's entire liability, and End User's exclusive remedy, with respect to any breach of the warranty set forth in this Section 9(a).
   b) EXCEPT AS EXPRESSLY SET FORTH IN SECTION 9(a), THE PRODUCTS OR SERVICES ARE PROVIDED "AS-IS' AND IMPERVA MAKES NO WARRANTY OF ANY KIND, WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE. IMPERVA, ITS PARTNERS AND SUPPLIERS MAKE NO WARRANTY THAT USE OF THE PRODUCTS OR SERVICES WILL BE UNINTERRUPTED, ERROR-FREE OR DEFECT-FREE, OR AVAILABLE AT ALL TIMES. IMPERVA HEREBY SPECIFICALLY DISCLAIMS, ON BEHALF OF ITSELF AND ITS PARTNERS AND SUPPLIERS, ALL IMPLIED WARRANTIES, INCLUDING ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW.

10. **Limitations of Liability**.   IN NO EVENT WILL END USER'S OR IMPERVA'S (AND ITS PARTNERS' OR SUPPLIERS') LIABILITY FOR DIRECT DAMAGES HEREUNDER EXCEED THE TOTAL VALUE OF AMOUNTS TO BE PAID BY END USER FOR THE PRODUCT OR SERVICE AT ISSUE.   IN NO EVENT SHALL END USER OR IMPERVA (OR ITS PARTNERS' OR SUPPLIERS') HAVE ANY LIABILITY TO THE OTHER FOR ANY LOST PROFITS OR REVENUES, LOSS OF DATA OR USE, INTERRUPTION OF THE SERVICES, COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES HOWEVER CAUSED AND, WHETHER IN CONTRACT, TORT OR UNDER ANY OTHER THEORY OF LIABILITY, WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.   THE FOREGOING LIMITATIONS SHALL NOT APPLY TO END USER WITH RESPECT TO ANY CLAIMS ARISING UNDER SECTION 2 (LICENSES AND RESTRICTIONS), OR TO EITHER PARTY UNDER SECTION 7 (CONFIDENTIALITY).

11. **Term and Termination**.

a) The term of this Agreement will commence upon Delivery of the Products to End User and will continue in effect for such time as End User continues to have the right to access the Products. Support for Software and/or Appliances will automatically renew at the end of the applicable Support term unless either party gives the other at least thirty (30) days' notice of non-renewal prior to the end of the then current term.

b) Either party may terminate this Agreement due to a material breach of this Agreement by the other party if such material breach remains uncured for a period of thirty (30) days following receipt of written notice by the breaching party; provided that Imperva may terminate this Agreement and/or all licenses granted to End User hereunder immediately upon written notice to End User if End User breaches any provision of Section 2 (License & Restrictions), Section 4 (Additional Terms for Subscription Services) or Section 7 (Confidentiality).

c) Upon the earlier of expiration of End User's rights or termination of the Agreement, Imperva will cease providing Subscription Services, Support and Professional Services, and each party shall promptly return or destroy the other party's Confidential Information in accordance with Section 7 above. Termination shall not relieve End User of the obligation to pay any fees accrued or payable to Imperva prior to the effective date of expiration or termination. The following sections shall survive termination of this Agreement: Sections 2(c), 2(d), 2(f), 4, 6 -12, and 16.

12. **Compliance with Laws**; Export. End User acknowledges that the Software contains encryption technology that is subject to export restrictions by the U.S. government and import restrictions by certain other governments. End User will not and will not allow any third-party to remove or export from the U.S. or allow the export or re-export of any part of the Software or any direct product thereof: (i) into (or to a national or resident of) Cuba, Iran, North Korea, Sudan or Syria (to the extent the U.S. government or any agency thereof restricts export or re-export to such countries); (ii) to anyone on the U.S. Commerce Department's Table of Denial Orders or U.S. Treasury Department's list of Specially Designated Nationals; (iii) to any country to which such export or re-export is restricted or prohibited, or as to which the U.S. government or any agency thereof requires an export license or other governmental approval at the time of export or re-export without first obtaining such license or approval; or (iv) otherwise in violation of any export or import restrictions, laws or regulations of any U.S. or foreign agency or authority. End User agrees to the foregoing and warrants that it is not located in, under the control of, or a national or resident of any such prohibited country or on any such prohibited party list. The Software is restricted from being used for the design or development of nuclear, chemical, or biological weapons or missile technology without the prior permission of the U.S. government. End User agrees to indemnify and hold Imperva, its partners and suppliers harmless against any claims, losses or expenses arising out of End User's breach of this Section 12.

13. **End User Mention**. End User consents to Imperva using its name and logo to identify End User as a customer of Imperva, such as use on Imperva's web site. Any use shall be subject to Imperva complying with any guidelines that End User may deliver to Imperva from time-to-time regarding the use of its name and logo. This consent terminates upon termination of this Agreement.

14. **Force Majeure**. Neither party will be liable to the other for any delay or failure to perform any obligation under this Agreement (except for a failure to pay fees) if the delay or failure is due to unforeseen events, which occur after the signing of this Agreement and which are beyond the reasonable control of the parties, such as strikes, blockade, war, terrorism, riots, natural disasters, government order, refusal of license by the government or other governmental agencies, in so far as such an event prevents or delays the affected party from fulfilling its obligations and such party is not able to prevent or remove the force majeure at reasonable cost.

15. **Evaluation**.

a) **Evaluation Product**. If the order is for End User to evaluate Product and its related Documentation on a temporary basis for non-commercial use ("Evaluation Product") and Imperva agrees to such evaluation, conditioned on End User's compliance with the terms and conditions of this Agreement, Imperva grants to End User during the Evaluation Period (as defined below), a nonsublicensable, nontransferable, nonassignable and nonexclusive, revocable license to use the Evaluation Product, solely at the location identified in writing by End User and solely for End User's internal evaluation of the Evaluation Product. End User may only grant access to the Evaluation Product to employees, contractors, agents or consultants who are bound to confidentiality and non-use obligations no less protective of Imperva's proprietary rights than this Agreement. Notwithstanding anything to the contrary as stated in this Agreement, all worldwide right, title and interest to the Evaluation Product, and all intellectual property rights in and to such product, are and will remain the exclusive property of Imperva.

b) **Evaluation Period**.   Unless otherwise agreed to by the parties in writing or terminated earlier in accordance with this Agreement, an evaluation shall commence upon delivery of the Evaluation Product and continue for thirty (30) days thereafter ("Evaluation Period").   Upon the expiration or termination of the Evaluation Period, (i) all licenses granted under this Section 15 for such evaluation will cease, and (ii) End User will immediately return the Evaluation Product to Imperva and destroy or erase any intangible copies of the Evaluation Product, and certify in a writing signed by an officer of End User and delivered to Imperva that all such copies of have been returned, destroyed or erased.

c) **Additional Evaluation Terms**.   Notwithstanding anything to the contrary as contained in this Agreement, End User acknowledges and agrees that the Evaluation Product is provided for evaluation "AS-IS" and Imperva and its suppliers make no representations or warranties of any kind, express or implied, with respect to the Evaluation Product, including, without limitation, any implied warranties of merchantability, title, fitness for a particular purpose, informational content, system integration, enjoyment, non-infringement or any other warranties arising out of course of dealing, usage or trade, and no obligation under Section 8.b (Indemnity) shall arise with respect to an Evaluation Product.

16. **Miscellaneous Provisions.**   The parties are independent contractors under this Agreement and nothing in this Agreement authorizes a party to act as an agent of the other or bind the other to any transaction or agreement. This Agreement will bind and inure to the benefit of each party's permitted successors and assigns. Neither party may assign or transfer this Agreement in whole or in part by operation of law or otherwise, without the other party's prior written consent. Any attempt to transfer or assign this Agreement without such written consent will be null and void.   Notwithstanding the foregoing, however, Imperva may assign this Agreement without consent to the acquiring or surviving entity in a merger or acquisition in which Imperva is the acquired entity (whether by merger, reorganization, acquisition or sale of stock) or to the purchaser of all or substantially all of Imperva's assets. Imperva's licensors are intended third party beneficiaries of this Agreement.   In the event any provision of this Agreement shall be determined to be invalid or unenforceable under law, all other provisions of this Agreement shall continue in full force and effect.   Except as specifically provided in this Agreement, the exercise by either party of any rights and remedies under this Agreement will be without prejudice to its other remedies under this Agreement or otherwise.   This Agreement contains the entire agreement of the parties with respect to the subject matter of this Agreement and supersedes all previous communications, representations, understandings and agreements, either oral or written between the parties with respect to said subject matter. This Agreement may be modified or waived only in a written instrument signed by both parties. A waiver of any breach under this Agreement shall not constitute a waiver or any other breach or future breaches.   Notwithstanding the foregoing, if a separate, written and mutually signed agreement for the acquisition of the Products and/or Services exists between End User and Imperva, the terms of that written agreement (excluding any pre-printed terms of any purchase order, confirmation or similar document, all of which will have no effect and will not be considered agreed to by Imperva) shall take precedence over this Agreement. All notices, requests, demands and other communications hereunder shall be in writing to the address set forth below for Imperva and on the applicable order or registration form for End User and shall be deemed to have been duly given: (i) upon receipt if by personal delivery; (ii) upon receipt if sent by certified or registered mail (return receipt requested); or (iii) two (2) days after it is sent if by overnight delivery by a major commercial delivery service.   This Agreement will be interpreted and construed in accordance with the laws of the State of California and the United States of America, without regard to conflict of law principles. The parties hereby consent to the exclusive jurisdiction and venue of the state and federal courts located in Santa Clara County, California for resolution of any disputes arising out or relating to this Agreement. The provisions of the United Nations Convention on Contracts for the International Sale of Goods and the Uniform Computer Information Transactions Act will not apply to this Agreement in any manner whatsoever.

Imperva, Inc.

3400 Bridge Parkway, Suite 200

Redwood Shores, CA 94065 USA

Revised October 2014

Table of Contents

CHAPTER 1

Introduction to SecureSphere on AWS

This publication is intended for administrators tasked with deploying an Imperva SecureSphere Web Application Firewall (WAF) configured as a Kernel Reverse Proxy (KRP) in an Amazon Web Services (AWS) environment. It assumes the reader has a working knowledge of AWS and details the configuration steps required to achieve a successful deployment.

Note: This guide assumes the user conducting deployment has good working knowledge of AWS.

Deployment Overview

This document describes deployment of SecureSphere on AWS in the order it should take place. It contains the following:

	Task/Subject	Description
1	Understanding SecureSphere Deployment in AWS on page 14	Provides an overview of Deploying SecureSphere in AWS, includes topology examples, and lists prerequisites.
2	System Prerequisites on page 18	Lists various items you should prepare before actually beginning to deploy SecureSphere in AWS.
3	Configuring AWS Infrastructure on page 19	Provides step-by-step instructions on how to prepare and configure the AWS infrastructure so that it is ready for the deployment of the SecureSphere Management Server and Gateway.
4	Deploying the SecureSphere Management Server on page 33	Provides instructions on how to deploy the SecureSphere Management Server once AWS infrastructure has been configured.
5	Licensing SecureSphere - BYOL	Once the SecureSphere Management Server has been deployed, you need to upload a license key to the Management Server before deploying gateways.
6	Deploying a SecureSphere Gateway on page 37	Once the license key has been uploaded, you need to deploy SecureSphere gateways, including creating a second interface.
7	Post Deployment Review on page 43	After having deployed both the SecureSphere Management Server, you should conduct a review to verify that you are ready to go online.

A Note for New SecureSphere Customers

Welcome to Imperva SecureSphere. If you are a new Imperva SecureSphere customer it is important to note that intial SecureSphere setup in Amazon Web Services is complex, requiring both the setting up of a number of native AWS components, as well as initial SecureSphere configuration.

Subsequently, it is important that you thoroughly review this guide and follow all required steps to get SecureSphere up and running.

If you encounter difficulties during setup, you can refer to the section Troubleshooting on page 48. Additionally, customers purchasing SecureSphere on AWS are entitled to Imperva's standard support package. After purchasing SecureSphere through AWS, you receive an email with instructions on how to contact support.

For more information on what this includes, please visit our website at www.imperva.com and under Support, choose Technical Support.

CHAPTER 2

Understanding SecureSphere Deployment in AWS

In AWS deployments, the SecureSphere Management Server, the SecureSphere Gateways, the protected web servers and the Elastic Load Balancers (ELBs) – everything is virtual. Moreover, the Gateways are scalable: in periods of peak demand, additional Gateways can be added to the Gateway Group and torn down when they are no longer needed. The web servers too can be scaled in the same way, in response to changes in the volume of traffic.

Notes:
- SecureSphere is available for AWS only in a WAF KRP configuration.
- SecureSphere AMIs are provided as Hardware Virtual Machines (HVM)

In AWS the Management Server holds the license for itself and the Gateways it manages. In order to enable auto scaling, your license should allow the number of desired gateways; otherwise a perpetual or time-based license must be acquired before being able to scale up.

Amazon Instance Type Mapping

The table below lists the recommended mapping of SecureSphere virtual appliance model numbers and the corresponding **minimum** Amazon Instance Types.

SecureSphere Virtual Appliance Model	SecureSphere Appliance Type	Minimum Amazon Instance Type
AV1000	Gateway	M3.large
AV2500	Gateway	M3.xlarge
AVM150	Management Server	M3.xlarge

Topology Overview

This section provides an overview of SecureSphere deployment in AWS and is meant to provide a model for deploying SecureSphere in AWS to match your requirements.

The configuration described here deploys SecureSphere in an existing AWS web server deployment, with two Availability Zones in the VPC.

> **Note:** Configuration in this document includes two availability zones to illustrate deployment which provides High Availability (HA). Users who don't require High Availability can deploy SecureSphere on AWS with one Availability Zone.

AWS Deployment Options

SecureSphere on AWS can be deployed in a number of configurations, as shown in the following table. Each row represents a different deployment and its options.

AWS Deployment Options

	# of External ELBs	Gateways	Internal ELB	Web Servers
1	Single	Scaling Group per Availability Zone	Spanning single or multiple Availability Zones	Scaling Group within Availability Zone or spanning multiple Availability Zones
2	Single	Single Scaling Group	Spanning single or multiple Availability Zones	Scaling Group within Availability Zone or spanning multiple Availability Zones
3	Multiple	Scaling Group per Availability Zone	Spanning single or multiple Availability Zones	Scaling Group within Availability Zone or spanning multiple Availability Zones
4	Multiple	Single Scaling Group	Spanning single or multiple Availability Zones	Scaling Group within Availability Zone or spanning multiple Availability Zones
5	**VPC Peering** across multiple VPCs, within a single region, across one or more AWS accounts. Note: DNS instances are required per VPC where the Imperva instances are deployed.			

Deployment Example without SecureSphere

The figure below shows a typical deployment before SecureSphere. It includes:

• Two availability zones

• An external Elastic Load Balancer
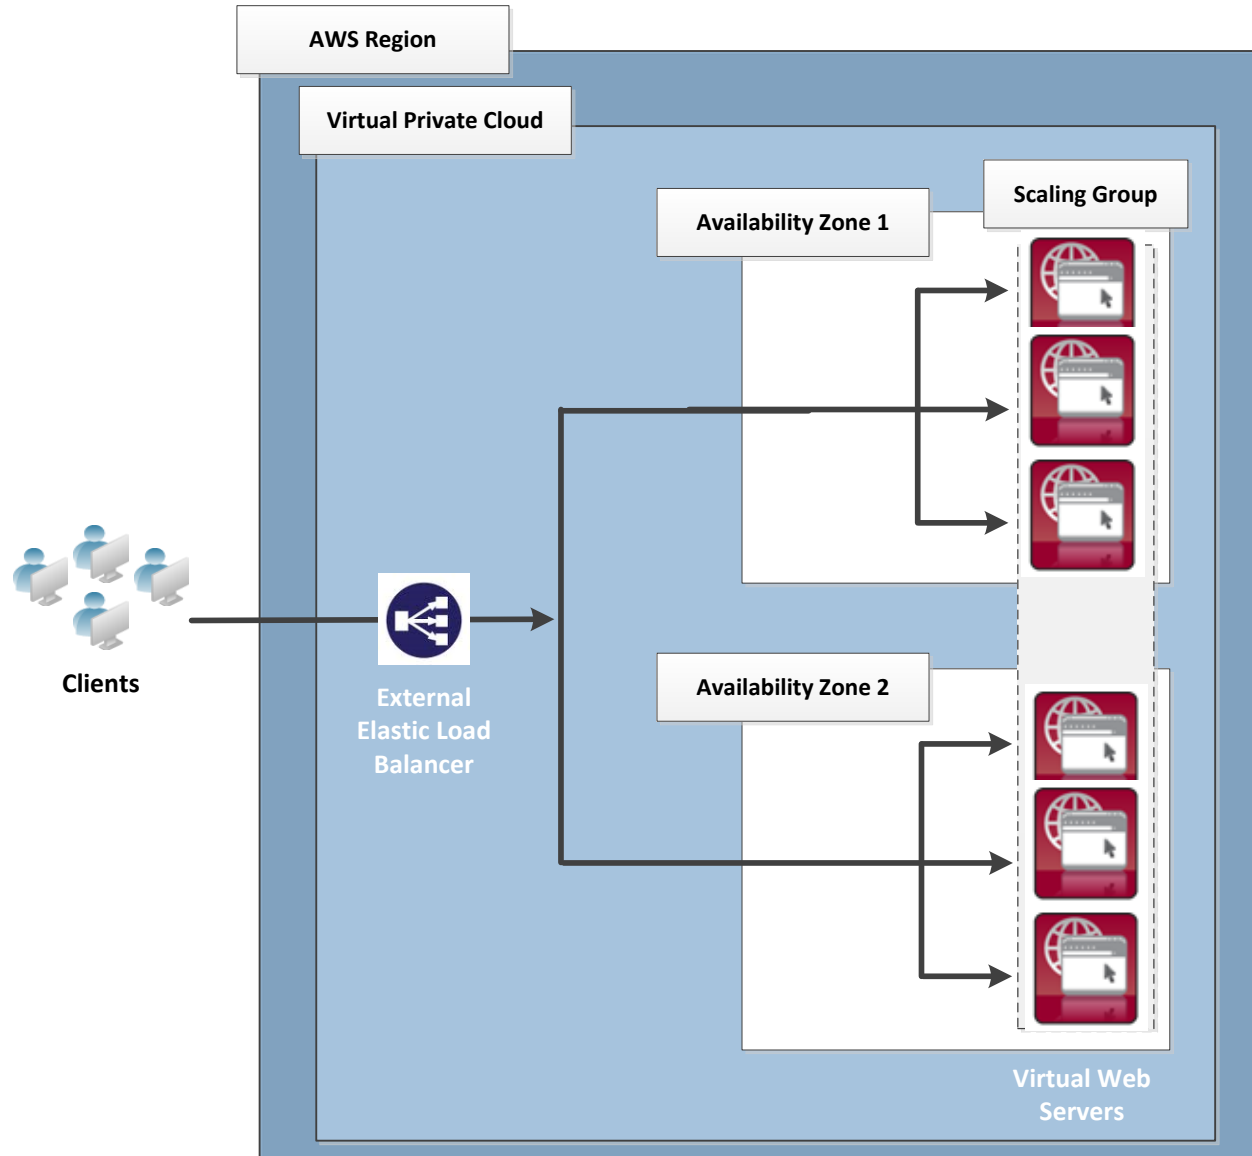
• A scaling group for each set of virtual web servers

Figure 1: BEFORE - Web servers deployed in AWS

Deployment Example with SecureSphere

The below image shows the deployment with SecureSphere. This example includes:

- Two availability zones
- Access for the SecureSphere Administrator, typically provided via a Windows Client installed from the AWS Marketplace or a VPN
- An external and two internal Elastic Load Balancers
- A scaling group for SecureSphere Gateways
- A scaling group for each set of virtual web servers
- A NAT Instance (or any other technology that provides external access)
- A Windows client to provide access to the browser based SecureSphere user interface
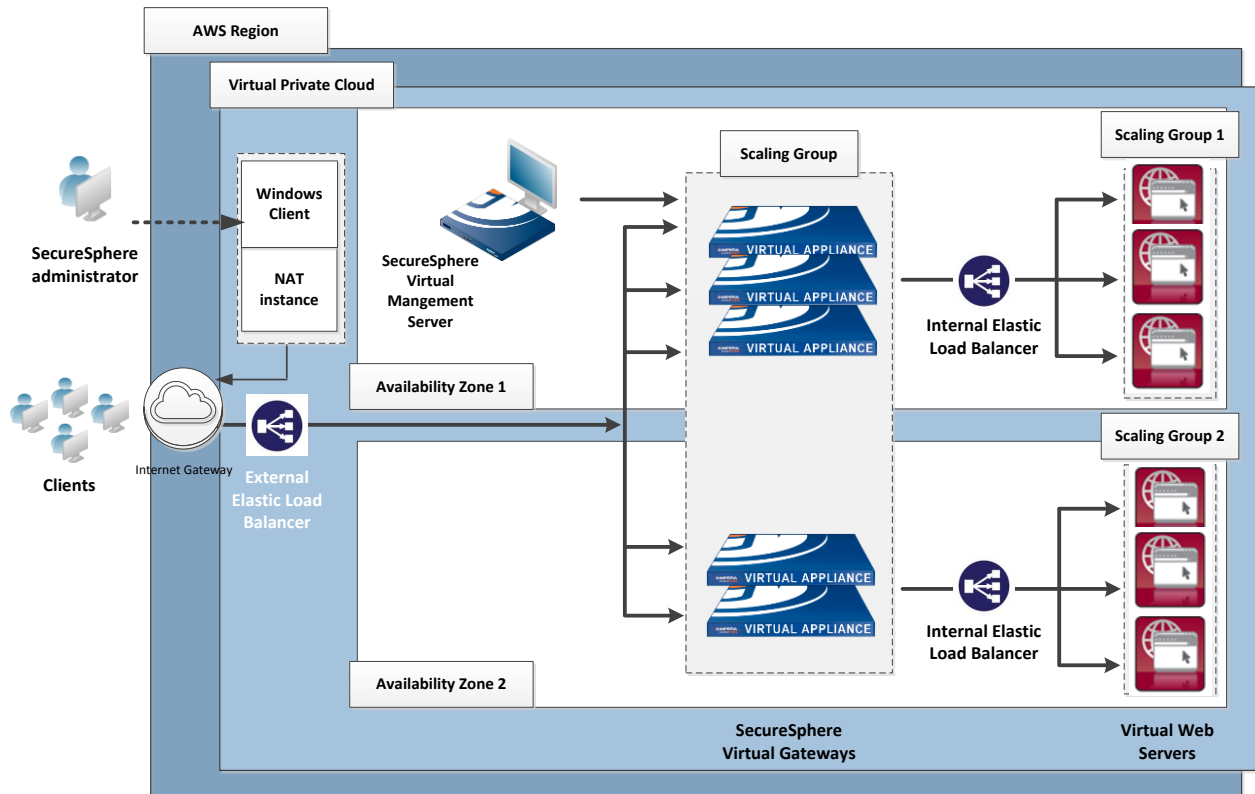
Figure 2: AFTER – SecureSphere WAF KRP added to the deployment

Please note that this is not the only deployment possibility. There are others: for example, a single internal ELB might be used for the web application servers (though this is not recommended for the reasons set out in Minimizing Traffic Costs "Minimizing Traffic Costs and Delays" on page 43) or several different applications might be deployed behind the SecureSphere WAF.

Note: The SecureSphere Virtual Management Server is in a private subnet, not a public subnet. See Subnets on page 24 for a more detailed diagram, and Windows Client - Connecting to the SecureSphere Management Server on page 34 for information on the connection from the SecureSphere GUI Client to the Management Server.

HTTP vs. HTTPS Support

Configuration described in this guide is based on a deployment where HTTPS communication is terminated at the External ELB, meaning that all traffic within the VPC is via HTTP. SecureSphere also supports full HTTPS traffic within the VPC.

The configuration below is relevant when working with HTTPS within the VPC.

To configure HTTPS support, you must:

* Configure a listener to HTTPS in the Load Balancer window.
* Upload an SSL certificate to the Load Balancer.
* Set the **Instance Protocol** to HTTPS.

Please note that using HTTPS will significantly impact performance, and traffic will be limited up to 100mbps. Therefore, if using HTTPS it is recommended to use the SecureSphere **AV1000**.

System Prerequisites

Before beginning the deployment, you need obtain the following:

Component	Definition
Amazon Account ID	You will get these when you open your account with Amazon Web Services.
Amazon Username and Password	The user who will perform the SecureSphere deployment should be defined as a **Power User** and granted the minimum required privileges for configuring the VPC and instantiating the SecureSphere Management Server, Gateway and the other components. See the Amazon IAM (Identity and Access Management) documentation for more detailed information about defining users and their privileges.
Amazon Access Key and Amazon Secret Key	These are the Amazon Security Credentials.
SecureSphere Template Files	These files are downloaded from the Imperva FTP site. There are two Cloud Formation template files (one for the Management Server and a second for Gateways). You need these files to create AWS stacks. These files are in JSON format. **Note**: Making changes to the original JSON files provided by Imperva beyond what is instructed in this document is not allowed. Any changes to these files without being approved by Imperva will not be supported.

In addition, you must ensure that your AWS Service Limits enable you to configure a sufficient number of instances and other resources for your specific deployment. Running out of instances in the course of the deployment will cause unnecessary delays. For more information, refer to the AWS documentation.

Note: This guide assumes that your application web servers are already deployed in AWS.

CHAPTER 3

Configuring AWS Infrastructure

The following sections – which assume a working knowledge of AWS on the part of the reader – describe the special considerations in configuring a SecureSphere deployment in AWS, and provide suggested configuration guidelines. However, since every deployment is unique, the reader should approach these guidelines with some flexibility and be prepared to depart from them if warranted by the particular characteristics, requirements and constraints of the specific environment.

The table below lists the required configuration steps to configure AWS Infrastructure.

	Configuration Step	
1	**Design the deployment.**	Determine which AWS components you will need and confirm that your AWS Service Limits enable you to configure them all.
2	**Create the VPC.**	In the deployment described here, the VPC already exists and the web servers are already deployed in it. If this is not the case in your deployment, create the VPC at this point. For more information, see VPC (Virtual Private Cloud) on page 22.
3	**Create the subnets.**	Create a subnet for each group of AWS components. For more information, see Subnets on page 24.
4	**Create a key pair.**	Key pairs are used to launch instances and to connect to them from outside the VPC. For more information, see Key Pair on page 26.
5	**Create a NAT instance.**	A NAT instance enables routing traffic from the Management Server and Gateways to the internet (for example, for to a syslog server), and to notify AWS of the success or failure of instance creation. For more information, see NAT Instance on page 26.

	Configuration Step	
6	**Update route tables.**	An AWS route table corresponds to the route table of a physical router. For more information, see Route Table on page 27.
7	**Create Elastic Load Balancers.**	See Elastic Load Balancer "Elastic Load Balancers" on page 28 .
8	**Elastic IP Addresses**	See Elastic IP Address on page 29.
9	**Create security groups.**	A security group corresponds to an internal firewall. For more information, see Security Group "Security Groups" on page 30.

AWS Console

When you log in to AWS, the Amazon Console is displayed.
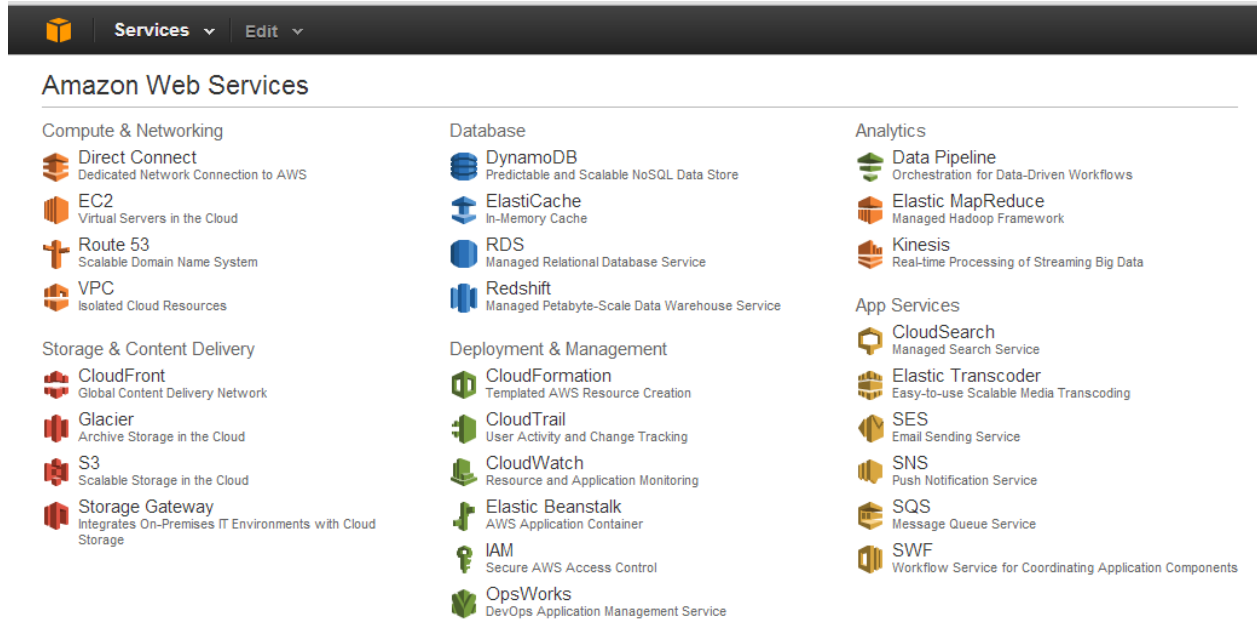
Figure 3: Amazon Console

The Amazon Web Services you will be using in this deployment are the following:

Amazon Web Services Component	will be used in order to ...
EC2	Create virtual servers in the Amazon Cloud. In this deployment, the web application servers have already been created and deployed (see the figure in Deployment Example without SecureSphere on page 16).
VPC	Create a Virtual Private Cloud in which the deployment will be located. In this deployment, the VPC already exists.
Cloud Formation	Create the SecureSphere Management Server and Gateways from templates provided by Imperva.
Cloud Watch	Monitor the success / failure of the deployment itself and afterwards, events in the VPC.

VPC (Virtual Private Cloud)

VPCs are configured by clicking Start VPC Wizard in the VPC Dashboard.
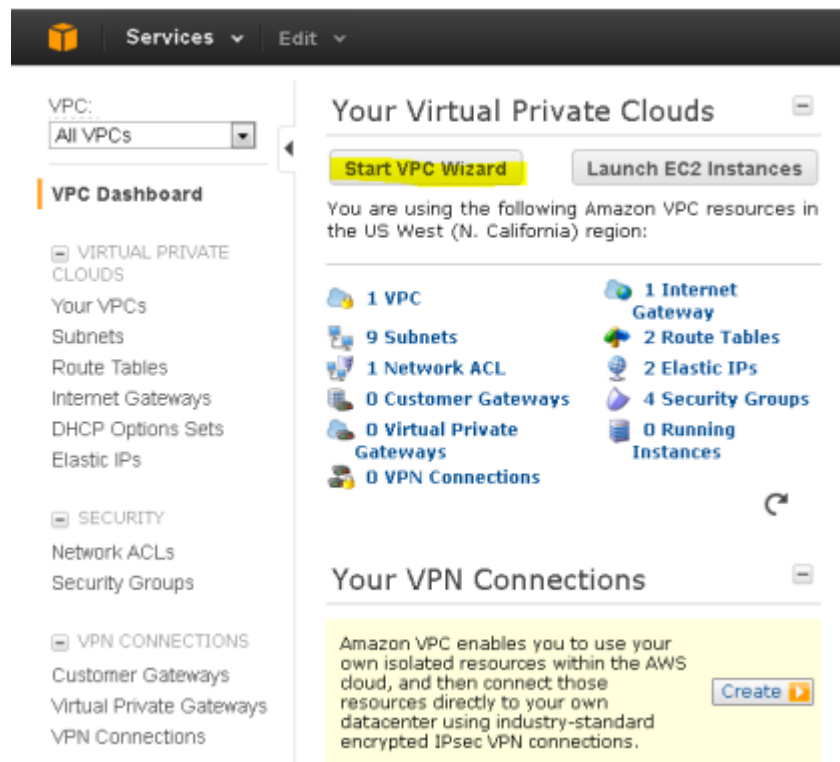
Figure 4: Creating a Virtual Private Cloud

Then the following is displayed once staring the VPC Wizard.
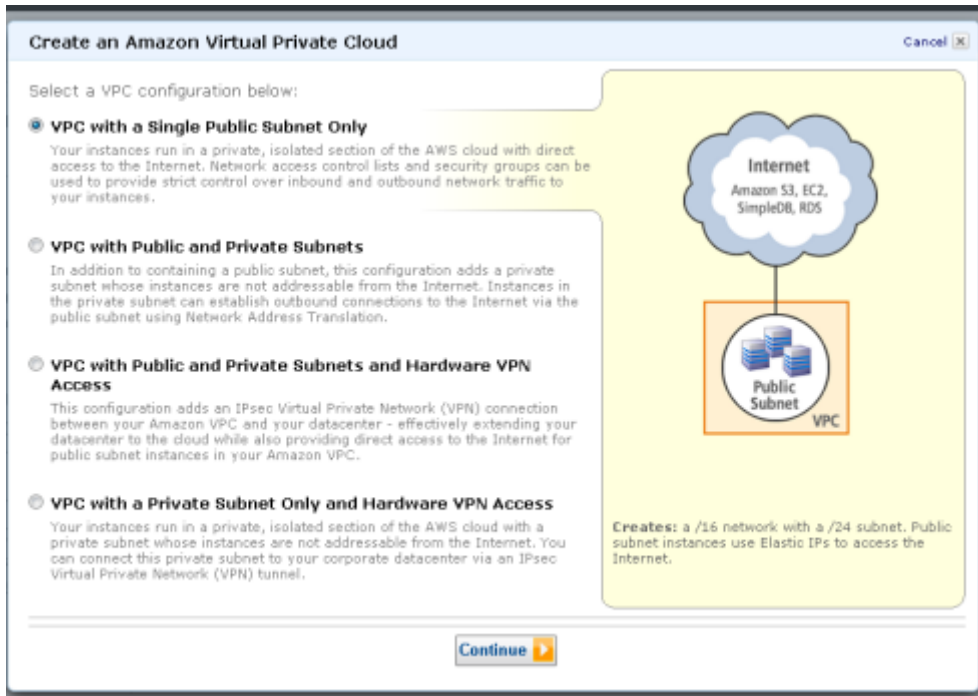
All the SecureSphere components should be instantiated in the same VPC as the web servers. Additionally, a SecureSphere Gateway scaling group should span all the Availability Zones in which the web servers are located. In this deployment, the VPC already exists, and consists of two Availability Zones (see Deployment Example with SecureSphere on page 17).

Subnets

Subnets are configured in VPC Dashboard > Subnets.

Figure 5: Creating a Subnet

The VPC should be configured with several subnets, for example as shown in the figure below.
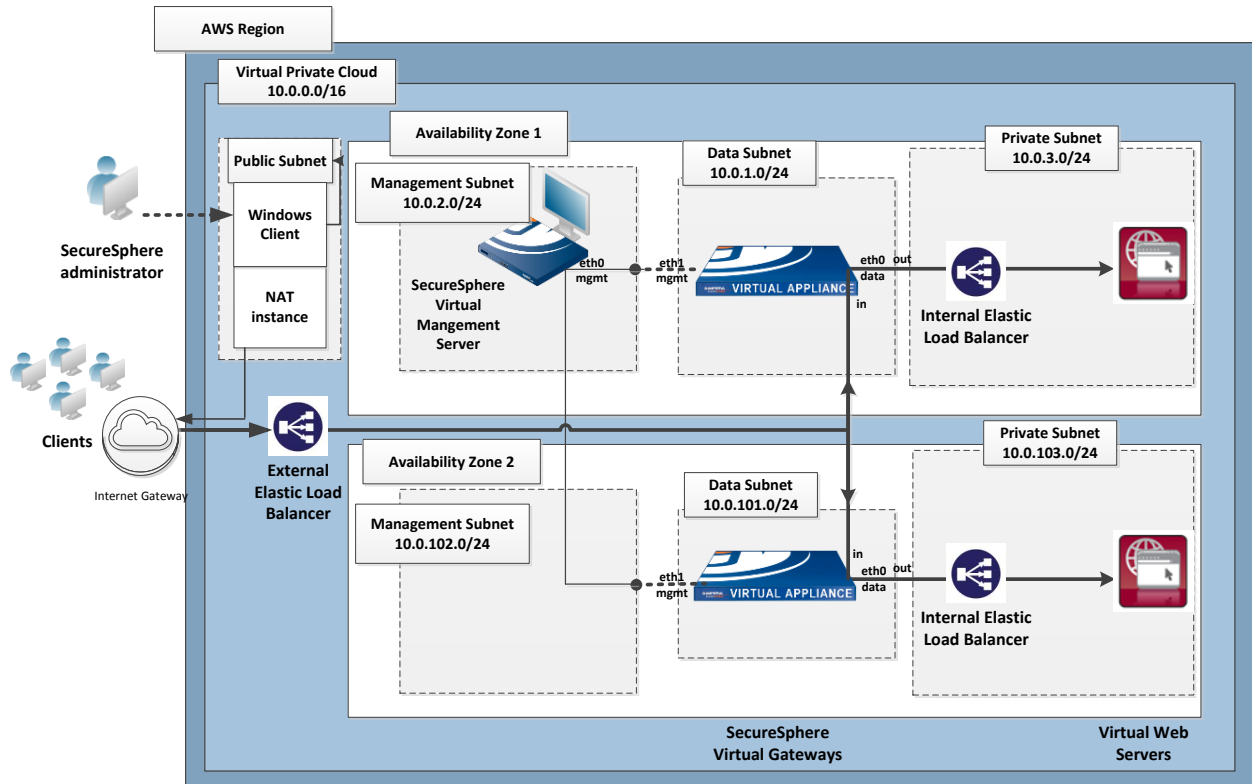
Figure 6: Subnets

The NAT instance and Windows client are located within a public subnet. While all other subnets are private, that is, they are not directly accessible from the internet.

A Windows client needs to be separately added from the Amazon Marketplace. It is then used so the SecureSphere Administrator can communicate with the SecureSphere Management Server.

Notes:

- A Gateway has two NICs, each one of which is connected to a different subnet in the same Availability Zone.
- The Gateway's data interface cannot be on the same subnet as the Management Server.

This configuration forces all traffic to the web servers to pass through the SecureSphere Gateways. When the deployment is complete, it is essential that you confirm that there are no alternate routes to the web servers that bypass the SecureSphere Gateways.

Key Pair

Amazon EC2 uses public–key cryptography to encrypt and decrypt login information. Public–key cryptography uses a public key to encrypt a piece of data, such as a password, then the recipient uses the private key to decrypt the data. The public and private keys are known as a key pair.[1]

To log in to your instance, you must create a key pair in AWS, then specify the name of the key pair when you launch an instance, and provide the private key when you connect to the instance.

Linux/Unix instances have no password, and you use a key pair to log in using SSH. With Windows instances, you use a key pair to obtain the administrator password and then log in using RDP.

Key Pairs are configured in **Services > EC2 Dashboard > Key Pairs**.

You will need to define at least one key pair that will be used, for example, when you:

- Launch an instance, for example, a SecureSphere Gateway, Management Server or NAT instance
- Connect to an instance, using SSH for example

If lost, the key file cannot be restored, so you must store the Key Pair file in a safe and secure manner.

NAT Instance

NAT Instances are used in this deployment in order to provide access to the SecureSphere Management Server and Gateway from the public domain. However you can alternatively decide to use other means, such as an HTTP Proxy.

NAT instances are configured in **Services > EC2 > Instances**.

Note: For detailed information about NAT instances, refer to the Amazon documentation.

Figure 7: Configuring a NAT Instance

A NAT instance enables SecureSphere Gateways and Management Servers in the private subnets to access the internet, for example, for syslogs, Imperva FPS services, and to notify AWS of the success or failure of instance creation.

Note: For information about configuring a Security Group for a NAT instance, see NAT Instance Security Groups on page 32.

A NAT instance is located in a public subnet (see the figure in Subnets on page 24).

Additionally, in order to provide public access you need to configure the associated **Security Group** protocol with a Source IP address of 0.0.0.0.

You should disable source/destination checking for the NAT instance, because it must be able to send and receive traffic when the source or destination is not itself. For more information, search for **Disabling Source/Destination Checks** in the Amazon AWS documentation.

Route Table

Route tables are configured in Services > VPC > Route Tables.
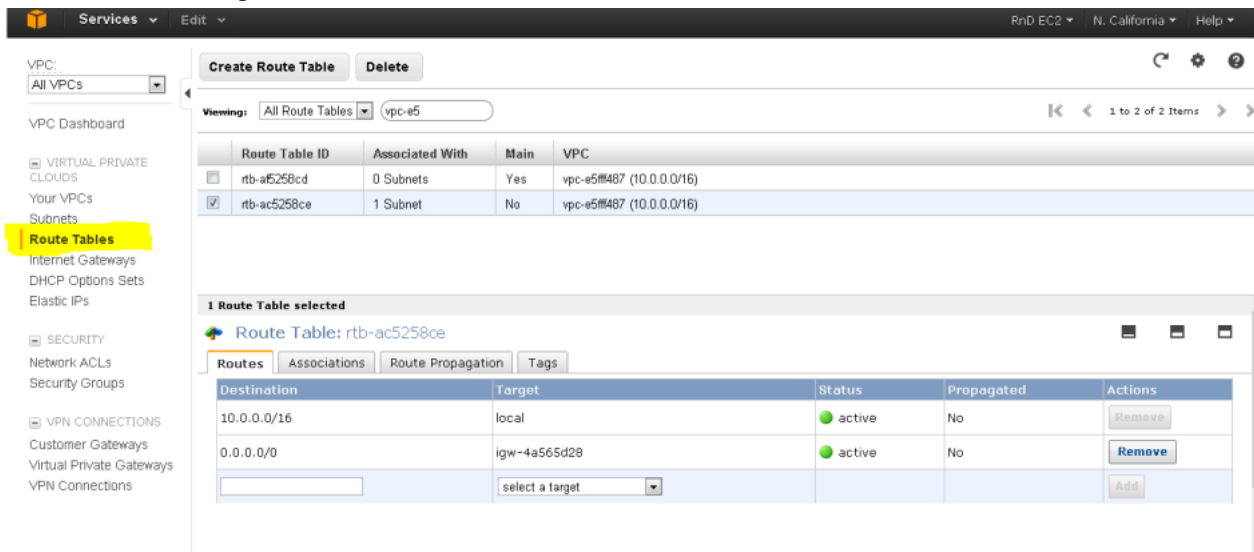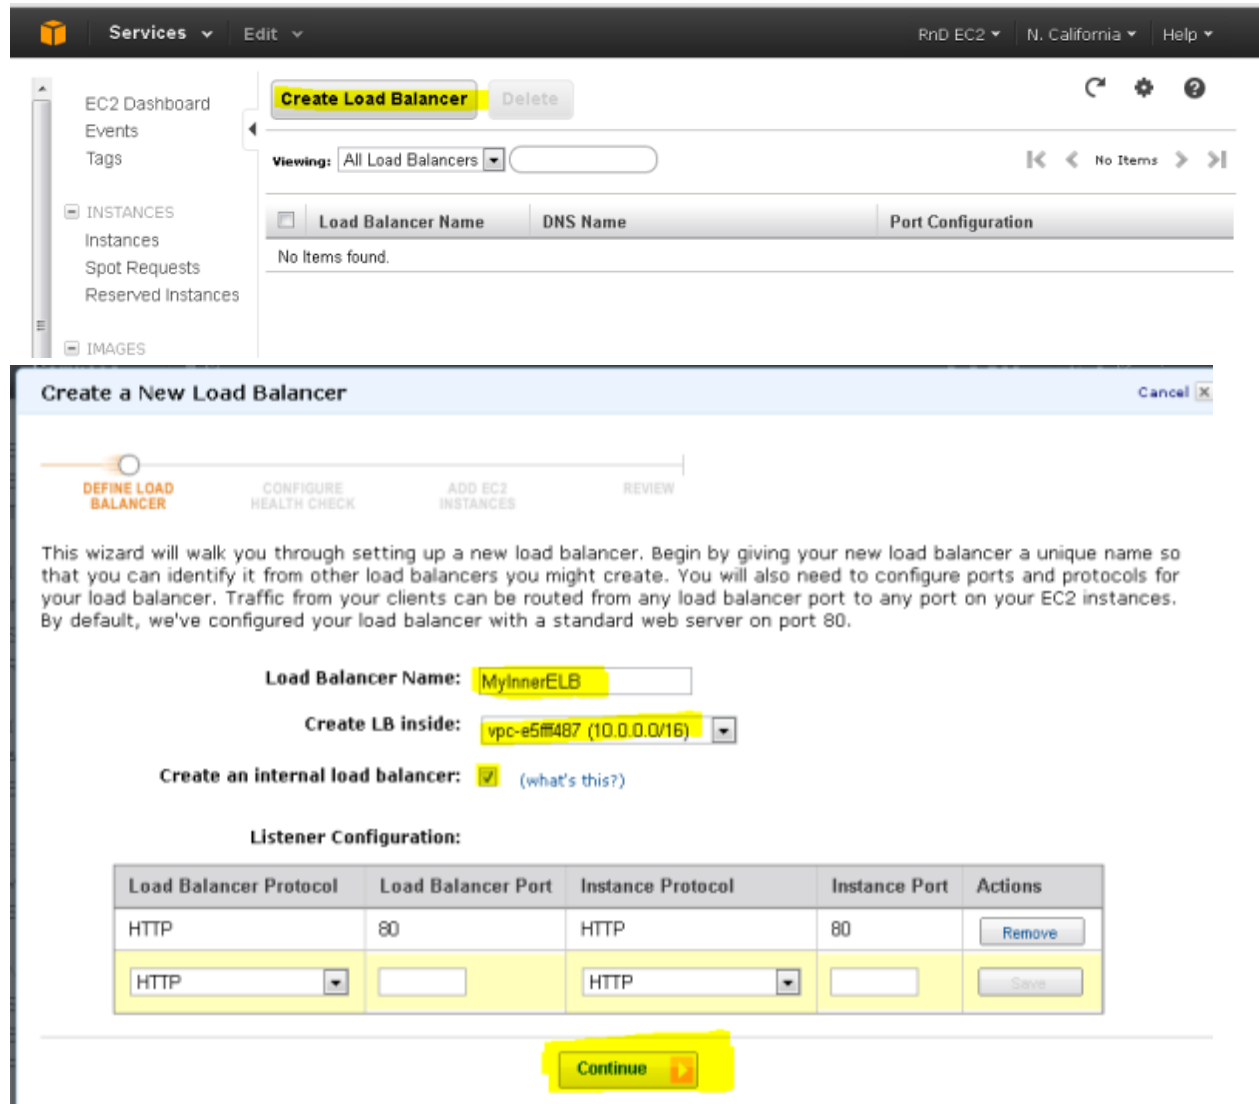

Figure 8: Create Route Table Window

A route table associated with the public subnet should route permitted outbound traffic through the NAT instance to the internet gateway that was automatically created as part of the VPC.

Elastic Load Balancers

Elastic Load Balancers are configured by clicking Create Load Balancer in Services > EC2 > Load Balancers.

Figure 9: Create an Elastic Load Balancer

In the deployment shown in Deployment Example without SecureSphere on page 16 , there is only an external Load Balancer fronting the web servers.

In the deployment shown in Deployment Example with SecureSphere on page 17, you will need three ELBs: an external ELB to balance traffic between the SecureSphere Gateways in the two Availability Zones, and an internal ELB in each Availability Zone to balance traffic among the web servers in that Availability Zone.

External ELB

This topic addresses details regarding the following items:

DNS

After completing the deployment you must associate the DNS entry for the web application with the external ELB.

Health Check

When configuring the external ELB's Health Check, you must set the **Ping Protocol** to TCP and the **Ping Port** to 80.

XFF

Because the external ELB is in effect a Reverse Proxy, XFF should be enabled in SecureSphere. For more information, see **Supporting Web Load Balancers** in the **SecureSphere Web Security User Guide**.

SSL

It is recommended, for performance reasons, that SSL be offloaded to the external ELB. For more information, see **Adding SSL Keys** in the **SecureSphere Web Security User Guide**.

Session Stickiness (Affinity)

If the web application does not maintain its own session cookies, then an ELB can be configured – by defining a stickiness policy in the AWS Console – to insert its own cookie to bind subsequent requests from the same user to the same web application instance.

Elastic IP Address

Elastic IP addresses are configured in **Services > EC2 > Elastic IPs**.

An Elastic IP address is a publicly available IP address, enabling access from the internet.

- If you are connecting to the Management Server from a Windows client using RDP, you should define an Elastic IP address for the Management Server.

- If you are connecting to the SecureSphere Management Server over a VPN, you may not need to define any Elastic IP addresses.

Security Groups

A Security Group acts as a firewall for instances, for example, for the SecureSphere Gateway instance. CloudFormation automatically creates two Security Groups: one for the Management Server and another for the Gateways. You may wish to modify the automatically-created Security Groups, for example, to protect ports other than the default port 80.
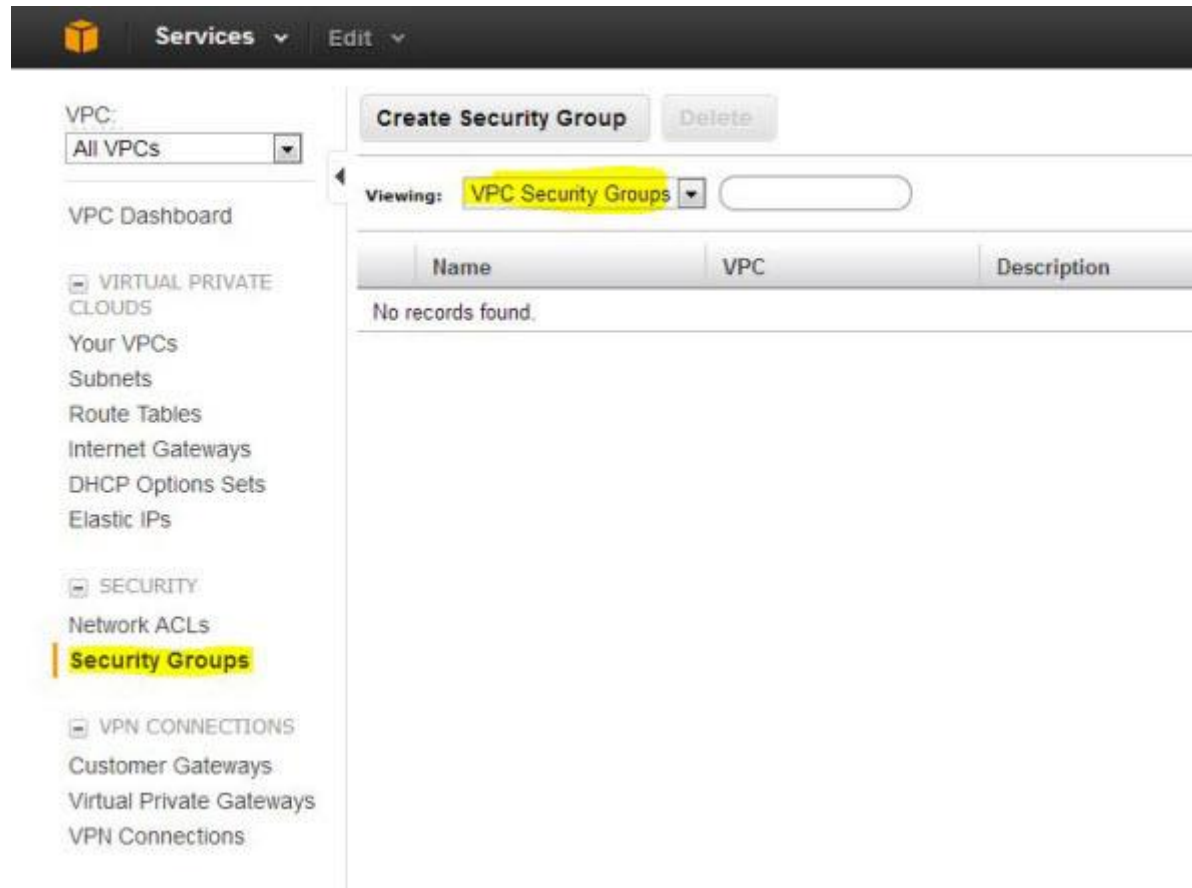

Figure 10: Creating a Security Group

The following diagram shows the Security Groups that should be defined for the deployment depicted in Deployment Example with SecureSphere on page 17.
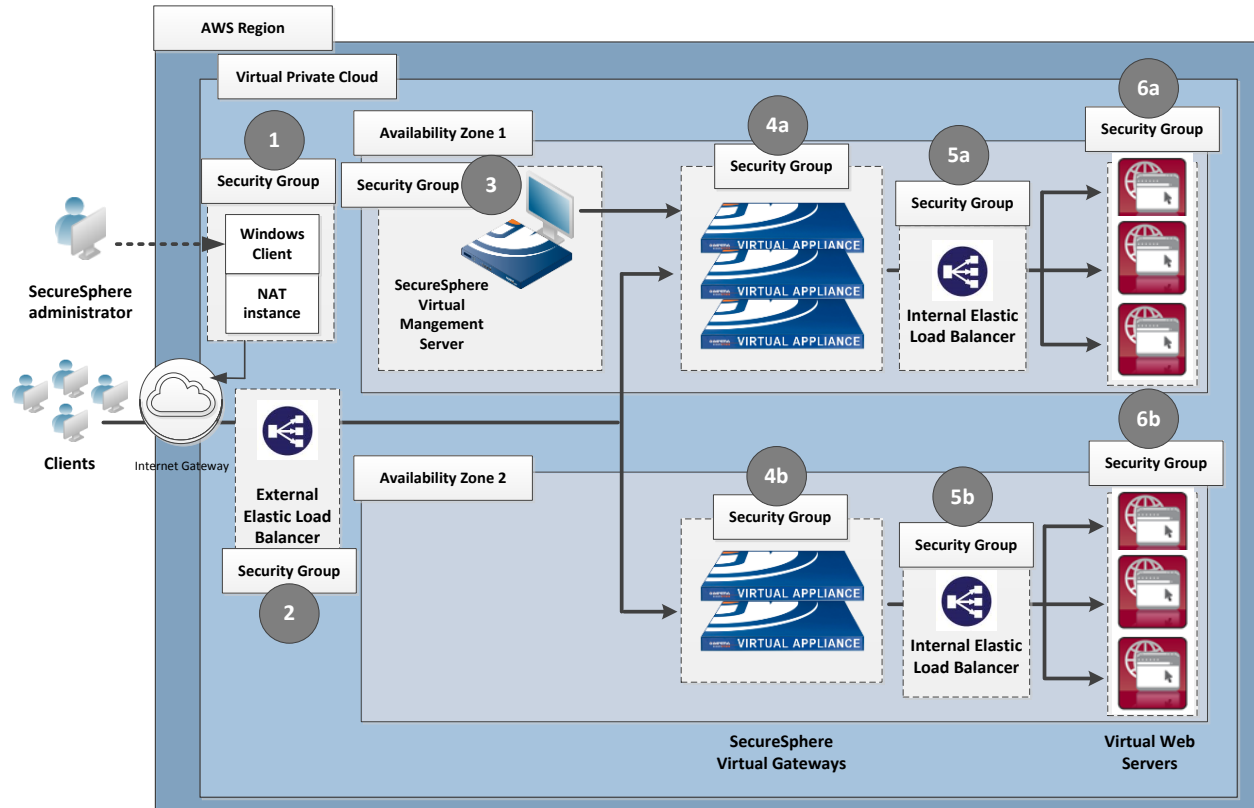

Figure 11: Security Groups

	Security Group	Rules
1	**Windows Client/NAT Instance**	Allows permitted outbound traffic to access the internet and responses to that traffic to enter the VPC (see NAT Instance Security Groups on page 32 below).
2	**External ELB**	Accepts traffic from the internet and sends traffic only to the SecureSphere Gateways.
3	**Management Server**	Sends and accepts SSH and inbound HTTP traffic from the SecureSphere administrator, either over a VPN or a remote client (see Windows Client - Connecting to the SecureSphere Management Server on page 34). Sends and accepts traffic from the SecureSphere Gateways it manages, as well as the enabled external systems, for example, imperva.com, the Imperva Cloud, syslog, *etc*.
4a 4b	**Gateways**	Sends and accepts SSH from the SecureSphere administrator, traffic from the Management Server, as well as the enabled external systems, for example, imperva.com, the Imperva Cloud, syslog, *etc*. Sends and accepts traffic to the internal ELBs.
5a 5b	**Internal ELBs**	Accepts traffic only from the SecureSphere Gateways and sends traffic only to the web servers.

	Security Group	Rules
6a 6b	**Application web servers**	Typically, accepts and sends traffic only from/to the internal ELBs, but other rules are possible as well as long as the only publicly-available path to the application web servers is through the SecureSphere Gateways.

You may find it necessary to adjust these rules for your specific deployment. In addition, you must take care that the Security Group rules are synchronized with changes to the configuration, for example, when new services are added.

NAT Instance Security Groups

Configure the NAT instance Security Group as described in the table below in order to allow the NAT instance to:

receive internet-bound traffic from the private subnet

receive SSH traffic from your network

send traffic to the internet, which enables the instances in the private subnet to get software updates

Inbound			
Source	**Protocol**	**Port**	**Comments**
the private subnet	TCP	80	Allows inbound HTTP traffic from servers in the private subnet.
the private subnet	TCP	443	Allows inbound HTTPS traffic from servers in the private subnet.
the public IP address range of your network	TCP	22	Allows inbound SSH access to the NAT instance from your network (over the Internet gateway)
Outbound			
Destination	**Protocol**	**Port**	**Comments**
0.0.0.0/0	TCP	80	Allows outbound HTTP access to the internet.
0.0.0.0/0	443	443	Allows outbound HTTPS access to the internet.

Note: For detailed information about NAT instances and Security Groups, refer to the Amazon documentation.

VPC Peering

In a situation where there is more than one VPC within the same region, there are Gateways on all the VPCs, but only one Management Server on one of the VPCs, you use VPC Peering to enable communication between these VPCs to enable the Management Server to connect with all the Gateways.

Note: A VPC with no NAT instance has no internet access, and VPC Peering alone does not solve this. In order for SecureSphere to work across multiple VPCs, you must create a NAT instance for each VPC.

CHAPTER 4

Deploying the SecureSphere Management Server

To deploy a SecureSphere Management Server in AWS, use **Services > CloudFormation** to create a stack, and upload the Management Server template file available on the Imperva FTP site.

During the course of creating the Management Server instance, you will be asked to enter a number of parameters. Some of these are the standard First Time Login parameters while others are specific to the AWS environment.

Management Server instance parameters	Description
AvailabilityZone	Specify the Availability Zone in which the Management Server is to be instantiated. For example: us-west-1a.
DNSDomain	The default is to use Amazon's DNS server domain, but you can specify another domain.
DNSServers	The default is to use Amazon's DNS servers, but you can specify the IP addresses of other servers.
ImpervaLicenseKey	This license key is provided in the Imperva license email you receive after activation.
InstanceType	Specify the AWS instance type, for example **m3.large**.
KeyPairName	Specify the name of the key pair to be used when accessing the Management Server with SSH.
NewUserName	The name of a SecureSphere administrative user who will be accessing the Management Server using CLI commands. **Note**: The user will authenticate using the key pair (see the **KeyPairName** parameter above).

Management Server instance parameters	Description
NTPServers	The default is to use Amazon's NTP servers, but you can specify the IP addresses of other servers.
PrivateIPAddress	The IP address to which SecureSphere administrators will connect in order to configure SecureSphere. If you use the default, the IP address will be assigned by AWS DHCP.
SecurePassword	The password used for Gateway-Management Server communication. The same password must be entered when configuring the SecureSphere Gateways.
SecurePasswordRetype	
SubnetId	Specify the ID of the subnet in which the Management Server is to be instantiated.
SystemPassword	The password used by the Management Server to access the SecureSphere database.
SystemPasswordRetype	
Timezone	The timezone of the Management Server in POSIX format. A list of valid time zone formats is available at http://il1.php.net/manual/en/timezones.php http://il1.php.net/manual/en/timezones.php - http://il1.php.net/manual/en/timezones.php. Click on the continent and enter one of the time zones exactly as displayed there, for example, **America/Halifax** or **Africa/Nairobi**.
VpcId	Specify the ID of the VPC in which to instantiate the Management Server. For example: vpc-a0f832c5.

Windows Client - Connecting to the SecureSphere Management Server

There are two possibilities for connecting a client to the SecureSphere Management Server and configuring SecureSphere using the GUI:

Over a VPN – This method takes advantage of an existing internal enterprise VPN to provide security. For more information, search for OpenVPN in the Amazon documentation.

Remotely running a Windows client within the VPC – This method uses Microsoft RDP. You will have to create a Windows client instance and assign a key pair for this option.

Note: Windows clients are available in the Amazon marketplace.

Sealed CLI

The following commands are available in the Command Line interface (CLI) in On-Demand deployments.

Command	Function
date	Display the system date and time.
df	Display file system disk space usage.
hades-show	Display hades parameters.
hades-watch	Monitor changes to hades parameters.
help	List available commands and their parameters.
history	Display command history.
hostname	Display the system hostname.
ifconfig	Display network interface settings.
impctl	The Imperva SecureSphere command line control utility.
netstat	Display network statistics.
ping	Send ICMP ECHO_REQUEST to network hosts.
quit	Exit SecureSphere shell.
reboot	Reboot the machine.
shutdown	Shutdown the machine.
traceroute	Print the route packets trace to network host.
patch-remote	Download and install SecureSphere patch from URL.
export-local	Export database to local file.
export-remote	Export database and upload it to URL.
import-local	Import database from local file.
import-remote	Import database from URL.
gti-gw-remote	Download the Gateway's get-tech-info from the Management Server.
gti-remote	Upload get-tech-info to URL.
version	Display version information.
wget	Test URL connectivity.

To obtain help for any of these commands (except exit), enter either of the following commands:

help <command name>

or

<command-name> -h

or

? <command-name>

Terminating a Management Server

Termination Protection is configured by default for Management Servers. To terminate a Management Server instance, you must first turn off Termination Protection (**Services > EC2**) and then terminate the Management Server.

Licensing SecureSphere - On-Demand

Your copy of SecureSphere On-Demand does not require a license. You can begin using SecureSphere On-Demand immediately after you install and configure it.

After configuring SecureSphere, it is recommended that you notify Imperva of your contact details. To do this, go to **Admin** > **Licenses** and click in the **Amazon On-Demand Registration Form** section.

> **Note:** SecureSphere supports both BYOL and On-Demand Licenses at the same time. If your implementation uses both of these license types, they will both appear in the Management Server licensing window.

After completing the registration process, you will be have access to the following:

1. Imperva Support and Professional Services.
2. The full range of ThreatRadar feature included with SecureSphere On-Demand.

CHAPTER 5

Deploying a SecureSphere Gateway

Note: Before deploying the SecureSphere Gateway, install the Imperva license on the Management Server and start the Management Server. For more information, see Licensing.

To deploy a SecureSphere Gateway in AWS, use **Services > CloudFormation** to create a stack, and upload the Gateway Server template file available on the Imperva FTP site.

During the course of creating the Gateway instance, you will be asked to enter a number of parameters. Some of these are the standard First Time Login parameters while others are specific to the AWS environment.

A Gateway should have exactly two NICs on different segments (subnets) in the same Availability Zone:

One NIC (eth0) is for monitoring traffic and should be facing the internal and external ELBs.
This interface cannot be on the same subnet as the Management Server.

The other NIC (eth1) is for management (including health checks) and communication with the SecureSphere Management Server.

Note: In contrast to a physical deployment, in the AWS environment the Gateway is configured so that management traffic is on eth1 and monitored traffic is on eth0.

A gateway, like any other EC2 instance is created by default with one NIC. You need to create an IAM user with limited permissions to create the second NIC. Once you create this user, you are provided an AWSAccessKey and AWSSecretKey which is described in the table below, which is then used to create the second NIC. Configure the limited user in the IAM with the following permissions:

- ec2:DescribeSubnets
- ec2:AttachNetworkInterface
- ec2:CreateNetworkInterface
- ec2:DescribeNetworkInterfaces
- ec2:ModifyNetworkInterfaceAttribute
- cloudformation:DescribeStackResources
- iam:GetRole
- iam:PassRole
- autoscaling:DescribeLifecycleHooks
- autoscaling:PutLifecycleHook
- autoscaling:DescribeAutoScalingInstances

Gateway instance parameters	Description
AvailabilityZones	Specify a comma-separated list of Availability Zones which the Gateway scaling group is to span. For example: "us-west-1a, us-west-2a".
AWSAccessKey and AWSSecretKey	These keys, which you receive when creating a limited user in the IAM (Identity and Access Management) as described in the previous section, are required to instantiate the Gateways.
DataSubnets	Specify a comma-separated list of subnet IDs, one for each Availability Zone which the Gateway scaling group is to span, in the same order as the list of Availability Zones specified under **AvailabilityZones**.
DNSDomain	The default is to use Amazon's DNS server domain, but you can specify another domain.
DNSServers	The default is to use Amazon's DNS servers, but you can specify the IP addresses of other servers.
ELBNames	Specify the names of the external ELBs. The Gateways will be assigned these ELBs, and will be scaled up and down based on the ELBs health check.
GatewayAssetTag	Specify the Gateway's asset tag (model number), for example **AV1000**.
GatewayGroup	Enter the name of the Gateway group to which the Gateway belongs. If the **UseSingleGatewayGroup** parameter (see below) is set to **False**, the Availability Zone will be appended to the name. This will be updated on the Management Server.
ImpervaLicenseKey	This license key is provided in the Imperva license email you receive after activation.
ImpervaPassword	The password for the **Imperva** user.
ImpervaPasswordRetype	
InstanceType	Specify the AWS instance type, for example **m3.large**.
KeyPairName	Specify the name of the key pair to be used when accessing the Gateways with SSH.
ManagementServerIP	The IP address of the Management Server which manages the Gateway.
ManagementSubnets	The IDs of the subnets in which the Management Server is located.

Gateway instance parameters	Description
NewUserName	The name of a SecureSphere administrative user who will be accessing the Gateway using CLI commands. **Note**: The user will authenticate using the key pair (see the **KeyPairName** parameter above).
NTPServers	The default is to use Amazon's NTP servers, but you can specify the IP addresses of other servers.
ScalingMaxSize	Specify the maximum number of Gateways in the Gateway scaling group. You must ensure that you have a license for the number of Gateways specified; otherwise the scale up will fail.
ScalingMinSize	Specify the minimum number of Gateways in the Gateway scaling group. The minimum supported is one gateway. However the default value for this field is two. It is recommended that you use two gateways to provide full gateway high availability.
SecurePassword **SecurePasswordRetype**	The password used for Gateway-Management Server communication. The same password must be entered when configuring the SecureSphere Management Server.
Timezone	The timezone of the Management Server in POSIX format. A list of valid time zone formats is available at http://il1.php.net/manual/en/timezones.php http://il1.php.net/manual/en/timezones.php - http://il1.php.net/manual/en/timezones.php. Click on the continent and enter one of the time zones exactly as displayed there, for example, **America/Halifax** or **Africa/Nairobi**.
UseSingleGatewayGroup	Select one of: • True – There will be only one Gateway group (specified by the **GatewayGroup** parameter - see above). • False – The Availability Zone name will be appended to the Gateway group names created by SecureSphere. So if there are multiple Availability Zones, there will be multiple Gateway groups.
VpcId	Specify the ID of the VPC in which to instantiate the Gateway scaling group. For example: vpc-a0f832c5.

The SecureSphere Gateways are in effect reverse proxies for the internal ELBs, and the details of the web application servers can be hidden from the Gateways. See the figure 4 in VPC "VPC (Virtual Private Cloud)" on page 22 for an example of this configuration.

HTTP Proxy

If your deployment includes an HTTPS proxy, you must edit the Gateway and MX JSON template files before deploying it and specify values for the following parameters as appropriate to your configuration.

Parameter	Definition

https.proxyHost	The IP address or resolvable hostname of the HTTPS proxy.
https.proxyPort	The port number on the HTTPS proxy to which to send HTTPS traffic.
http.proxyUser	The user name for the connection to the HTTP proxy.
http.proxyPass	The password for the connection to the HTTP proxy.

Note: HTTPS support is required in a proxy-only environment for AWS deployment.

Configuring a SecureSphere Gateway

Once you have configured AWS infrastructure, deployed the SecureSphere management server, licensed SecureSphere, and deployed SecureSphere gateways, you need then access SecureSphere via the SecureSphere GUI and configure a number of items to get up and running. These items include:

	Configuration in SecureSphere	For more information...
1	Create a server group.	Creating a SecureSphere Server Group "Creating a SecureSphere Server Group and HTTP Service" on page 40
2	Define one or more KRP rules for the SecureSphere Gateway, so that traffic is correctly directed to the web servers.	Configuring KRP Rules on page 41
3	Configure the Server Group's Operation Mode.	Configuring Operation Mode on page 41

For more information on configuring a SecureSphere gateway, see the **SecureSphere Web Security User Guide**. The following sections describe configuration issues specific or particularly important in the AWS environment.

Creating a SecureSphere Server Group and HTTP Service

Once the gateway has been deployed, access SecureSphere configuration via a web browser using the following path: https://<Your Management IP address>:8083/ and log on. Then:

- In the SecureSphere **Main** workspace under **Setup > Sites**, right click the **Sites** tree and click **Create Server Group**.
- Right click the new **Server Group** and click **Create > HTTP Service**.

See the SecureSphere Web Security User Guide for assistance with these tasks.

Configuring KRP Rules

When defining a KRP rule in SecureSphere, the traffic should be directed to the internal ELB fronting the web application servers.

To access KRP rule configuration, in the Server Group you created, create an **HTTP service** and select it in the **Sites** tree. Click the **Reverse Proxy** tab. Then under **Gateway IP alias**, click **New** and configure the relevant details.

In an AWS deployment, in the field **Internal IP / Hostname**, enter the **AWS DNS Name** of the **internal ELB** rather than an IP address, because the internal ELB's IP address may change from time to time. You can see the DNS Name in **Load Balancers** under **EC2 Dashboard > Network & Security**.

For more information about defining SecureSphere KRP rules, see the **SecureSphere Web Security User Guide**.

Configuring Operation Mode

There are three operation modes:

- **Active**: SecureSphere monitors traffic and apply policies. This means that alerts are generated and traffic is blocked when required by a policy.
- **Simulation** (default): SecureSphere simulates monitoring, that is, it monitors traffic and generates alerts and violations, but does not block traffic. It is recommended to use simulation mode when SecureSphere is learning traffic.
- **Disabled**: SecureSphere neither monitors nor blocks traffic, which flows uninterrupted through the SecureSphere gateway to the web servers. This mode should be used only in exceptional circumstances, for example, for debugging.

Initially, after SecureSphere is first installed, it is configured in simulation mode. While in this mode, you can identify suspicious traffic by examining whatever alerts are generated. At some point, you will want to switch to active mode in order to block this traffic, some of which will be blocked by SecureSphere's default security policies. There will be a period of time during which you may wish to fine tune the policies, possibly defining new ones, in order to adapt your configuration to your specific requirements, for example, to eliminate false positives.

To configure the operation mode:

1. In the **Main** workspace, select **Setup > Sites**.
2. In the **Sites** window, click the server group whose operating mode you want to modify.
3. Click the **Definitions** tab.
4. Select the desired **Operation Mode**.
5. Click **Save**.

For information on fine tuning your configuration, see the **SecureSphere Web Security User Guide**.

Important Notes

- **Gateway Group Names**: SecureSphere Gateway group names should not be changed after they are initially defined. The reason is that scaled-up Gateways will continue to be assigned to the old group name.

- **Auto Scaling**: You can change the default auto scaling rules in Amazon CloudWatch if you find them unsuitable for your specific environment. Auto scaling cannot be disabled, but you can configure the rules in such a way that no auto scaling ever takes place.

- **Gateway Configuration Changes**: Changing the configuration of Gateways is complicated by the need to ensure that scaled-up Gateways receive the changed configuration rather than the previous one. To do this, update the stack (**Services > CloudFormation**) with the changed values, then manually scale the Gateways down and then up.

CHAPTER 6

Post Deployment Review

This section lists some of the issues you should review after completing the deployment.

Secure Access

Ensure that the only publicly-available path to your application web servers is through the SecureSphere Gateways.

Minimizing Traffic Costs and Delays

You should configure your deployment in a way that reduces to a minimum the traffic flowing among multiple Availability Zones. This has two benefits: you will avoid the extra costs as well as the delays associated with cross-Availability Zone traffic. For example, if you have two Availability Zones, configure an internal ELB for each Availability Zone (as in Deployment Example with SecureSphere on page 17) rather than a single internal ELB for both.

Scaling Rules / License

You should take care to configure the scaling rules to match your Imperva license, that is, to ensure that AWS will not attempt to scale up SecureSphere Gateways beyond the number permitted by the license and to scale down Gateways when their time-based licenses expire.

Cross-Region Load Balancing

If you require the capability to perform load balancing across Amazon regions, you should employ third-party tools for this purpose.

IP Address Change

It may happen that the IP address of the SecureSphere Management Server changes, for example, after a crash or after Amazon maintenance. If this occurs, the Gateways will be unable to communicate with the Management Server until you manually restore the previous IP address.

You should enable detailed AWS monitoring so that you can respond in a timely manner to IP address changes and other important events.

CHAPTER 7

Patching AWS

The procedure for patching an AWS SecureSphere deployment is different for Management Servers and for Gateways.

For information on patching an AWS Gateway, see Patching an AWS Gateway on page 46.

For information on patching an AWS Management Server, see Patching an AWS Management Server on page 47.

Patching an AWS Gateway

Gateways are deployed using Cloud Formation. Because a Gateway instance does not store persistent data, it can be created and destroyed by the AWS scaling mechanism. For this reason, directly patching a running Gateway instance is not effective: new Gateways will not include the patch.

The patch procedure therefore consists of updating the scaling group (Cloud Formation stack).

To patch an AWS Gateway:

1. Back up your current CFN template as follows:

    • Select the Cloud Formation stack.

    • Click the **Template** tab.

    • Save the template.

    • Click the **Parameters** tab.

    • Save the parameters.

    This procedure is described in the AWS documentation.

> **Note:** If you are already running the maximum number of Gateways allowed by the Management Server license, the new Gateway will not be able to register to the Management Server. You will have to take down one of the Gateways before continuing with this procedure. If you have a High Availability deployment and do not want to lose High Availability functionality during the patch update, create an additional On-Demand Gateway stack to handle the traffic during the update.

2. In Cloud Formation, select the Gateway stack.

3. Click **Update Stack**.

4. Click **Upload a template to Amazon S3**.

5. Upload the JSON file of the patch AMI.

6. Review the JSON parameters and verify them. The parameter values are carried over from the previous version, and you can change them if required.

7. Password parameters are empty and you must select **Use existing value** to copy them to the new template.

8. In the **Options** screen, leave the default settings and click **Next**.

9. Click **Update** to start the patch installation.

10. The Gateway stack will now perform a "rolling update," that is, it will take down one Gateway after another and bring up a new instance (which includes the patch) in its place, by default at ten minute intervals until all the Gateways are updated.

If for whatever reason you want to restore the previous version, update the Cloud Formation stack to use the previous template and parameters and the stack will be updated, that is, restored to the previous unpatched version.

For information on patching an AWS Management Server, see .

Patching an AWS Management Server

This procedure describes how to patch an AWS Management Server.

To patch an AWS Management Server:

1. Stop the Management Server (using the impctl teardown CLI command).

2. Create a snapshot of the Management Server.

3. Run the patch file using the patch-remote command, as follows:

   ```
   patch-remote [-h] [--insecure] [--user USER:PASSWORD] [--proxy HOST:PORT]
   [--proxy-user USER:PASSWORD] PATCH_URL.
   ```

   For details regarding optional download flags, see the table **Optional Download Flags** below.

4. When the installation ends, reboot the Management Server instance.

If for whatever reason you want to restore the previous version, you can do this from the snapshot.

For information on patching an AWS Gateway, see

Optional Download Flags

Name	Description
--insecure	Allows connections to SSL sites without certificates
--user USER:PASSWORD	Sets the server's usename and password
--proxy HOST:PORT	Tells SecureSphere to use an HTTP proxy on the listed port
--proxy-user USER:PASSWORD	Defines the proxy's username and password

APPENDIX A

Troubleshooting

This section includes some guidelines for troubleshooting deployment problems. The following information for troubleshooting is available:

Troubleshooting Checklist

The following is a list of items that you can check when trying to troubleshoot installation or operation of SecureSphere on AWS.

Troubleshooting Checklist

Item to Verify	Description
AWS System Log	The AWS System Log lists errors encountered in AWS infrastructure and can give an indication as to problems that may be occurring. When experiencing issues during setup, it is recommended you examine this log. For more information see Get AWS System Log on page 50
MX and Gateway have outside access	The NAT instance must be configured to enable outside access for a variety of purposes. The following provides an overview of what needs to be configured for the NAT instance: • The MX and gateway must have a route table associated with the subnets directing to the NAT Instance - e.g : "0.0.0.0 -> Nat instance id" • The NAT instance must have a public route table connected to the internet gateway - e.g. "0.0.0.0 - > igXXXX instance ID" Note that the gateway will probably be stuck (in "WaitCondition" state) as part of the first time login if access is not be available You must disable the source/destination check in the NAT instance. For more information, see Amazon documentation on **Disabling Source/Destination Checks**. For information on configuring a NAT instance, see NAT Instance on page 26.
Connectivity between Devices	For SecureSphere to properly operate, there needs to be connectivity between its various devices. When encountering issues you should check to verify the below: • **Security Groups**: Needed for relevant ports between the MX, gateway and ELBs. usually ports 8083, 22, 443, 80. For more information on security groups, see Security Groups on page 30. • **Subnets and routing**: Configured accordingly to allow access. For information on configuring subnets, see Subnets on page 24. • **Elastic's IPs**: Must be configured to reach within the VPC. Alternatively a VPN can be configured to allow for direct access from within your IP addresses. For information on configuring Elastic IPs, see Elastic IP Address on page 29.
ELB is configured to listen to and Forward HTTP Protocol	The Elastic Load Balancer (ELB) must be configured to listen for the HTTP protocol and forward HTTP traffic on the ports used by the SecureSphere gateway and web servers. To verify this is taking place: • Make sure that the health check is configured per the recommended guidelines. For information on the ELB Health Check, see Health Check on page 29. • Check that ELBs show instances as in service (Health check are getting response). For more information on ELBs, see External ELB on page 29

Item to Verify	Description
Scaling Group is configured as required and that scaling policies applied there have not been changed	Scaling policies should be based on CPU and Throughput , for 5 minutes polling interval.
KRP Rules are Configured	KRP rules must be defined so that the KRP will listen to the relevant port. The server IP address should be the IP address associated with the relevant Inner ELB name. For more information on KRP Rules, see Configuring KRP Rules on page 41.
ec2_auto_ftl.log	A log is available at **/var/log/ec2_auto_ftl.log** which contains an overview of initial setup of SecureSphere listing the status of tasks including first tome login, creating a new users, setting passwords, asset tag, timezone and more. On-Demand customers can user the command ec2log to display the log file.

Troubleshooting Errors

The following is a list of errors that may be encountered and suggested resolutions to these errors

Problem	Possible Cause
Instance creation fails (ROLLBACK_COMPLETE)	The reason for the failure is available in the stack's **Services > CloudFormation > Events** tab. The event shows the reason for the failure. If the event shows a WaitHandle Timeout message, the NAT instance through which communication with AWS takes place may be misconfigured, or the proxy is not configured properly.
Deletion of Management Servers fails (DELETE_FAILED).	Termination Protection is configured by default for Management Servers. To delete the Management Server stack you must manually disable Termination Protection in **Services > EC2**.
Failed connecting to CloudFormation, validate NAT Instance or Proxy Configuration	Displayed in the AWS System Log, this error indicates there is a problem with the NAT configuration. It is recommended that you check NAT settings as described in NAT Instance on page 26.

Get AWS System Log

The AWS System Log can contain information that will assist you in diagnosing problems you're encountering.

When encountering deployment issues, its always recommended that you download the AWS system log to see if there are any messages.

To access the AWS System Log:

1.   In AWS, under **Instances**, right click the Instance.

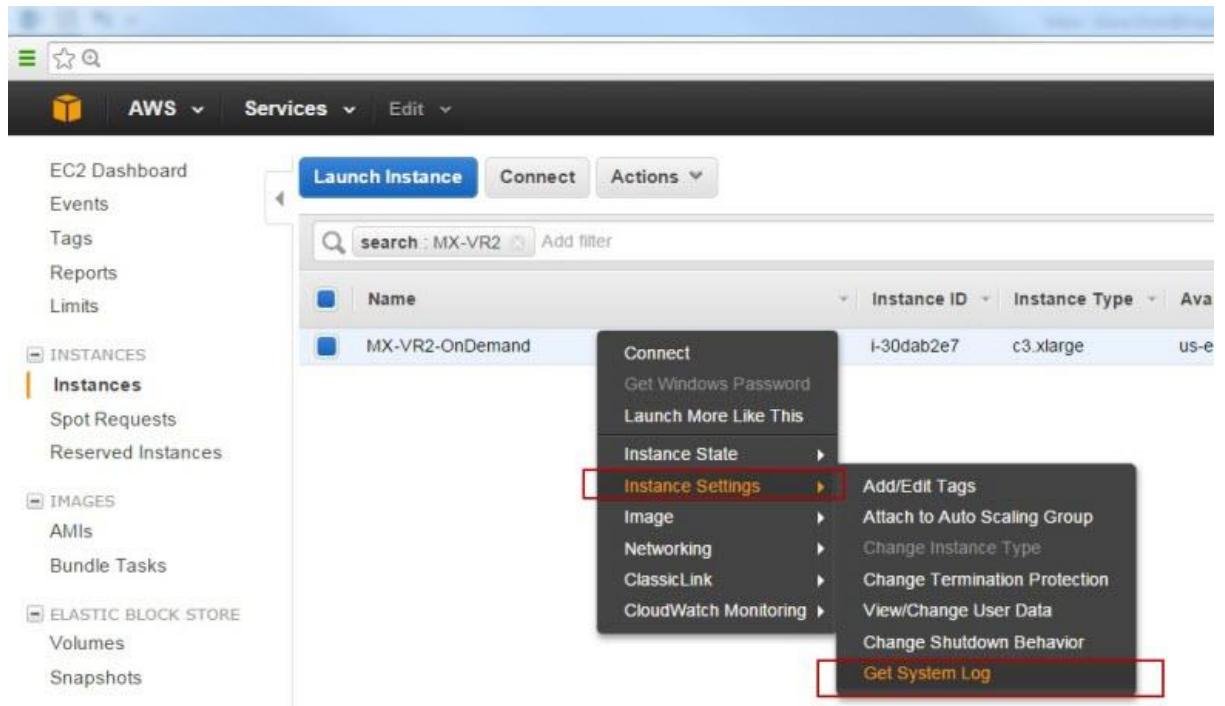2. Under **Instance Settings**, Select **Get System Log**.

Figure 12: Get System Log

The System Log opens in a separate window as shown below. in this example, the message states "Failed connecting to CloudFormation, validate NAT Instance or Proxy Configuration." So to troubleshoot, you could check NAT instance configuration as described in
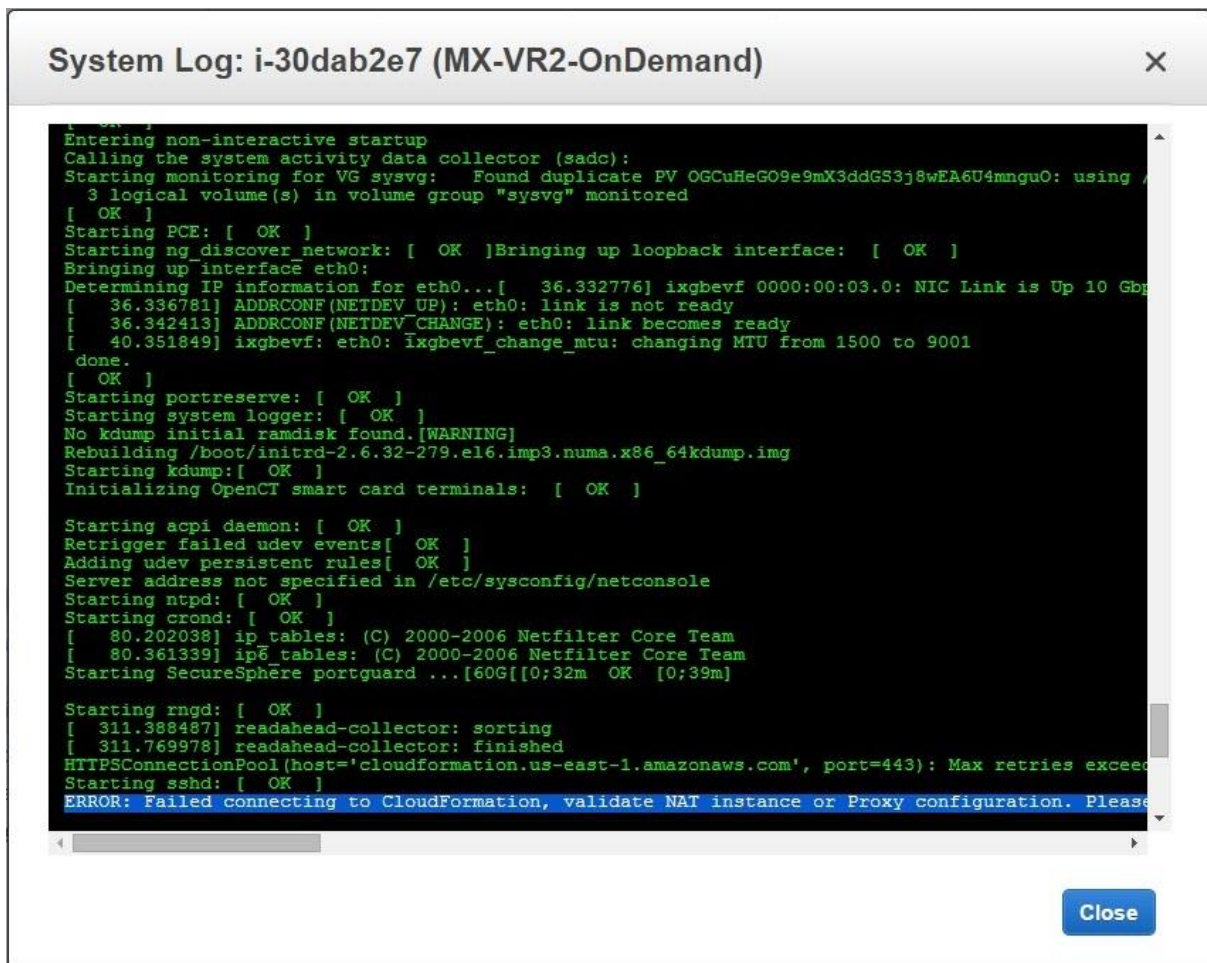

Figure 13: AWS System Log Example

HTTP Health Check

If the ELB Health Check is configured to use HTTP, the Gateway can be configured to periodically log the health check and its "health status" in a cyclical log. The health check confirms that the connection through the External ELB > Gateway > Internal ELB > Web Server path and back is up, using HTTP packets.

The log file is in the `/proc/hades/http_health_check` directory, and it is limited to 3,000 entries, after which the entries are overwritten cyclically.

If you are using this feature, you should configure it for all Gateway stacks and Gateway Groups managed by the Management Server.

To configure the HTTP health check:

1. In the **Main** workspace, select **Setup** > **Gateways**.
2. Select the Gateway group.
3. In the **Details** tab, open the **Advanced Configuration** section.
4. Enter the following text and the click **Save**.

```
<http-health-check-monitoring>
    <health-check-url url="<URL>" interval-sec="<n>" />
</http-health-check-monitoring>
```

For example:

```
<http-health-check-monitoring>
    <health-check-url url="/healthCheck/health.html" interval-sec="300" />
</http-health-check-monitoring>
```

Name	Description
url	The URL of the host whose health is to be checked. You can define multiple URLs, each one on a separate line. Note: It is recommended you created a custom page for the Health Check URL and not use a default page. Otherwise the log may fill up and it may be difficult to identify relevant issues.
interval-sec	The interval in seconds between health checks.

Click **Save**.

The log entries contain the following information:

- timestamp
- details of the IP addresses and ports in the path
- URL
- status

To delete the HTTP health check:

1. In the **Main** workspace, select **Setup** > **Gateways**.
2. Select the Gateway group.
3. In the **Details** tab, open the **Advanced Configuration** section.
4. Delete all the text between the opening and closing `http-health-check-monitoring` tags, so that all that remains is the following, and click **Save**.

```
<http-health-check-monitoring>
</http-health-check-monitoring>
```

Debugging a Failed Gateway

The Management Server monitors the AWS SQS for requests to terminate a Gateway, which can occur for one of 3 reasons:

- The Gateway instance failed an ELB health check.

- The Gateway instance was scaled down.

- The Gateway instance was manually terminated by the user.

In all of these cases, AWS removes the Gateway from the stack and scales up another Gateway in its place, and the Management Server unregisters the Gateway, so that it is no longer counted against the license.

If the reason for the Gateway instance termination is that the Gateway failed an ELB health check, the following events occur:

- AWS delays terminating the failed Gateway instance for a pre-defined period (up to 30 minutes).

- The Management Server runs the get-tech-info command on the failed Gateway (rebooting the Gateway if necessary) and stores the output file locally (on the Management Server).

- The Management Server requests that AWS terminate the failed Gateway instance.

The SecureSphere administrator can then review the get-tech-info file to determine the reason for the Gateway's failure.

If the reason for the Gateway instance termination is that the Gateway instance was scaled down or terminated by the user, the Management Server does not generate a get-tech-info file but allows AWS to immediately terminate the Gateway instance.

If more than one Gateway instance fails, the Management Server deals with them successively. There can be complications in unusual circumstances, for example, if many Gateways fail at the same time and AWS terminates a failed Gateway instance before the Management Server is able to run get-tech-info on it.

During the time a failed Gateway instance remains up pending the retrieval of its get-tech-info file and before AWS terminates its instance, the string "Under_Log_Retrieval" is appended to its instance name.

APPENDIX B

Backup and Restore

This section review the process of backup and restore, as follows:

Backup

Backing up the SecureSphere Management Server EBS (the Management Server data) is performed by running a Full System Export and then creating an AWS snapshot. These two actions should be coordinated as follows:

1. Create an Action Set of type Archiving.

2. To the Action Set, attach an action interface of type **AWS Snapshot > Create AWS Snapshot**.

3. Provide the Amazon credentials (Access Key ID and Secret Access Key) for creating the AWS snapshot.

4. Configure the Action Set as an Archiving Action in Full System Export (**Admin > Maintenance > Export Settings**).

5. Schedule the Full System Export.

A system event is issued upon the completion of the Archiving Action.

These actions should be scheduled at regular and frequent intervals and run by a user whose permissions are restricted to taking snapshots. The user must have URL access to the relevant Amazon Region in order to perform the snapshot. See the Amazon documentation for detailed information about defining users and their privileges.

The snapshot is created in the AWS environment. Only the last two snapshots are saved.

Restore

This section reviews the various procedures involved with restoring a system and includes:

Management Server

If the SecureSphere Management Server fails, you can restore it as follows:

To restore the SecureSphere Management Server:

1. Create a new SecureSphere Management from the AMI as described in Deploying the SecureSphere Management Server on page 33.
2. Confirm that Management Server still has the same IP address; otherwise the Gateways will be unable to communicate with it.
3. Restore the EBS from the snapshot, as described below.

Management Server EBS

If the SecureSphere Management Server EBS fails, it can be restored from the snapshot.

To restore the SecureSphere Management Server EBS:

1. In **AWS EC2 > Snapshots**, from the **Snapshots** page, right-click the snapshot from which to restore the SecureSphere Management Server EBS.
2. Select **Create volume**.
3. In **Type**, select the appropriate volume type, for example, Standard volume.
4. In the **Instances** page, stop the Management Server.
5. Wait for the Management Server to shut down completely.
6. Locate the **volume** used by the Management Server in the **Volumes** page.
7. Right-click the **volume** and select **Detach volume**.
8. Locate the **volume** created from the snapshot and select **Attach volume**.
9. Select the Management Server instance.
10. Select the Device **/dev/sda1**.
11. Click **Attach**.
12. Start the Management Server instance.
13. Connect to the Management Server using SSH.
14. Stop the Management Server with the `impctl server stop` command.
15. Import the export file that was created with the system export function. To import the file use the full_expimp.sh utility.
16. The export file location in the snapshot is `/var/tmp/server-backup`.
17. Start the Management Server with the `impctl server start` command.

For more information on exporting and importing the SecureSphere configuration and on starting and stopping the Management Server, see **SecureSphere Administration Guide**.

APPENDIX C

Migrating an AWS On Demand Deployment to a BYOL Deployment

Note: Before beginning this procedure, contact Imperva and obtain a BYOL license suitable to your requirements.

The individual steps in this procedure are documented in detail elsewhere in this Deployment Guide and in the SecureSphere User and Administration Guides.

To migrate an AWS deployment from On-Demand to BYOL:

1. Login to the On-Demand Management Server using `ssh`.
2. Run the CLI "`export-remote`" command and save the exported file to a remote location.

    You will need the exported file later in this procedure.
3. Using Cloud Formation, create a new BYOL Management Server.
4. Login to the BYOL Management Server using `ssh`.
5. Stop the BYOL Management Server.
6. Run the CLI "`import-remote`" command and import the file you previously exported.
7. After the import successfully completes, start the BYOL Management Server.
8. Log into the BYOL Management Server GUI.
9. Apply the BYOL license you obtained before you began this procedure (see the note above) to the BYOL Management Server.

10. Remove the individual On-Demand Gateways from the BYOL Gateway screen, but **do not remove the Gateway Group**.
11. Create a new BYOL Gateway stack pointing to the BYOL Management Server.
12. Associate the Gateway stack with the Gateway Group.
13. After determining that the BYOL deployment is functioning correctly, you can stop and terminate the On-Demand deployment.

Index