

IBM Security Solutions Architecture for Network, Server and Endpoint



Comprehensive discussion of the IBM Security Framework and IBM Security Blueprint

Detailed insight into the threat and vulnerability landscape

Extensive solution architecture and component introduction

Axel Buecker
Kent Browne
Louis Foss
Jaco Jacobs
Vladimir Jeremic
Carsten Lorenz
Craig Stabler
Joris Van Herzele



International Technical Support Organization

**IBM Security Solutions Architecture for Network,
Server and Endpoint**

February 2011

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

Second Edition (February 2011)

This edition applies to the following IBM Security Solutions products: IBM Security Network Intrusion Prevention System, IBM Security Network Controller, IBM Security SiteProtector, IBM Tivoli Endpoint Manager, IBM Security Server Protection, IBM Security Network IPS Virtual Appliance, IBM Security Virtual Server Protection for VMware, IBM RealSecure Server Sensor, IBM Tivoli Netcool Configuration Manager, IBM WebSphere DataPower XML Firewall, IBM Lotus Protector, and IBM Tivoli Application Dependency and Discovery Manager.

In addition, we examine many different IBM Security Service offerings that offer solutions for the Network, Server and Endpoint security domain.

© Copyright International Business Machines Corporation 2011. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	xi
Trademarks	xii
Preface	xiii
The team who wrote this book	xiv
Now you can become a published author, too!	xvii
Comments welcome	xvii
Stay connected to IBM Redbooks	xviii
Summary of changes	xix
Part 1. Business context and terminology	1
Chapter 1. Introducing the IBM Security Framework and IBM Security Blueprint	3
1.1 Business context for IT security	4
1.2 Drivers that influence security	4
1.2.1 Business drivers that influence security	5
1.2.2 IT drivers that influence security	7
1.3 Common industry approaches to IT security management	10
1.3.1 Control objectives for information and related technology	10
1.3.2 ISO/IEC 27002:2005	11
1.4 IBM Security Framework	12
1.4.1 Security Governance, Risk Management, and Compliance	14
1.4.2 People and Identity domain	15
1.4.3 Data and Information domain	16
1.4.4 Application and Process domain	18
1.4.5 Network, Server and Endpoint domain	19
1.4.6 Physical Infrastructure domain	20
1.5 IBM Security Blueprint	21
1.5.1 Foundational Security Management	24
1.5.2 Security Services and Infrastructure	25
1.5.3 Architectural principles	27
Chapter 2. The components of the IBM Security Blueprint	31
2.1 Foundational Security Management	32
2.2 Subcomponents	34
2.2.1 Command and Control Management	34
2.2.2 Security Policy Management	38

2.2.3 Risk and Compliance Assessment	44
2.2.4 Identity, Access, and Entitlement Management	54
2.2.5 Data and Information Protection Management	62
2.2.6 Software, System and Service Assurance	70
2.2.7 Threat and Vulnerability Management	75
2.2.8 IT Service Management	84
2.2.9 Physical Asset Management.	88
2.3 Conclusion.	92
Chapter 3. The Network, Server and Endpoint solution pattern	93
3.1 Deriving the solution patterns for the IBM Security Framework security domains	94
3.2 Examining the IBM Security Blueprint components for Network, Server and Endpoint	95
3.3 Using the solution pattern for Network, Server and Endpoint planning and design.	102
3.4 Conclusion.	104
Chapter 4. Common security architecture and network models	105
4.1 Security is omnipresent.	106
4.2 Enterprise Security Architecture model.	107
4.2.1 Security architecture delivery processes.	108
4.3 Common network components	111
4.3.1 Packet filter firewall	112
4.3.2 Circuit level firewall	112
4.3.3 Application layer firewall	113
4.3.4 Dynamic packet filter firewall.	113
4.3.5 Routers	114
4.3.6 Intrusion detection and prevention	114
4.4 Common network models and security domains.	116
4.4.1 Network zones	118
4.5 Practical designs	120
4.5.1 DMZ	120
4.5.2 Intranet	122
4.6 Additional components	123
4.7 Conclusion.	125
Chapter 5. Threat and vulnerability management	127
5.1 Security concepts and terminology	128
5.2 Malware.	129
5.3 Denial-of-service (DoS).	132
5.4 Advance Persistent Threat (APT)	133
5.4.1 Preventing Advance Persistent Threat attacks	136

5.5 Threat management	136
5.5.1 Threat mitigation architecture	138
5.6 Vulnerability management.	140
5.6.1 A need for vulnerability management	141
5.6.2 Comparing vulnerability assessment methods	142
5.7 Conclusion.	144
Part 2. IBM Security Solutions for Network, Server and Endpoint.	147
Chapter 6. Security intelligence, research, and technology	149
6.1 Security and cyber intelligence	150
6.1.1 Objective	150
6.1.2 IBM Security X-Force Research and Development Organization	151
6.2 Research	154
6.2.1 Research methods	155
6.3 Development	156
6.4 How can your business benefit	159
6.5 Protocol Analysis Module	165
6.5.1 Protocol Analysis Module internals	166
6.5.2 Protocol analysis module example	173
6.5.3 IBM Shellcode Heuristics	175
6.5.4 IBM Injection Logic Engine	176
6.5.5 JavaScript obfuscation detection	177
6.6 Content analysis research and technology	180
6.7 Spam protection	186
6.8 Security terms and definitions	190
6.8.1 X-Press update	190
6.8.2 Personal firewalls	190
6.8.3 Intrusion Detection Systems and Intrusion Prevention Systems.	191
6.8.4 Buffer Overflow Exploit Prevention	192
6.8.5 Application control.	192
6.8.6 Antivirus signatures.	193
6.8.7 SecurityFusion Module	193
6.8.8 White box testing.	194
6.8.9 Black box testing.	195
6.9 Conclusion.	197
Chapter 7. Centralized management	199
7.1 Benefits of centralized management.	200
7.1.1 Reducing cost	200
7.1.2 Demonstrating compliance and business value	200
7.2 Managing threats and vulnerabilities.	201
7.2.1 Asset and vulnerability prioritization	201
7.2.2 Modifying technical security controls.	204

7.2.3	Monitoring threats	204
7.3	IBM Security SiteProtector overview	205
7.3.1	SiteProtector components	205
7.3.2	Appliance and agent support	212
7.3.3	SiteProtector communication channels	212
7.3.4	Data redundancy	217
7.3.5	Authentication and encryption	219
7.3.6	Separation of duties and auditing	220
7.4	Managing operational security in SiteProtector	221
7.4.1	Managing assets	223
7.4.2	Managing policies	224
7.4.3	Performing event monitoring and analysis	229
7.4.4	Reporting functionality	233
7.4.5	Ticketing options	233
7.4.6	IBM Tivoli Security Information and Event Manager	235
7.4.7	IBM Tivoli Netcool/OMNIBus	239
7.5	Conclusion	242
Chapter 8. Network security solutions		243
8.1	IBM Security Network IPS	244
8.1.1	Zero-day protection	247
8.1.2	Next generation product enhancements	247
8.1.3	Next generation hardware	247
8.1.4	Next generation firmware	249
8.1.5	Next generation virtual appliances	258
8.2	Intrusion and intrusion prevention definitions	262
8.2.1	Intrusion prevention	262
8.3	Intrusion prevention policies	263
8.4	Intrusion prevention enforcement	266
8.4.1	General network requirements	267
8.4.2	Network IPS requirements	267
8.4.3	Performance	269
8.4.4	Security	269
8.4.5	Reliability	270
8.4.6	Deployment	271
8.4.7	Management	271
8.4.8	Confidence	271
8.5	Physical deployment model	272
8.5.1	Intrusion prevention architecture	272
8.5.2	IBM Security Network IPS: Modes of operation	273
8.5.3	External bypass	275
8.5.4	Copper and fibre connectivity	277
8.5.5	High availability	277

8.5.6	10 Gbps environments	282
8.5.7	SCADA environments	283
8.5.8	VoIP environments	283
8.5.9	FIPS 140-2 certification	284
8.6	IBM Tivoli Netcool Configuration Manager	284
8.6.1	Optimizing complex network environments	285
8.6.2	Conclusion	286
8.7	IBM WebSphere DataPower	287
8.7.1	XML and web services network security threats	288
8.7.2	IBM WebSphere DataPower and meeting SOA challenges	288
8.7.3	The IBM WebSphere DataPower SOA Appliance product line	289
8.8	IBM Lotus Protector for Mail Security	295
8.9	Conclusion	297
Chapter 9. Host security solutions		299
9.1	IBM Tivoli Endpoint Manager	300
9.1.1	Platform	302
9.1.2	Key components	305
9.1.3	Deployment architecture	307
9.1.4	Console	311
9.1.5	Computer groups	315
9.1.6	Fixlet sites	315
9.1.7	Relevance	316
9.1.8	Web Reports	317
9.1.9	Visualization Tool	317
9.2	Proventia Desktop Endpoint Security	318
9.2.1	Attack vectors covered	318
9.3	IBM Security Server Protection	322
9.3.1	Architecture overview	324
9.3.2	Windows and Linux server protection	326
9.3.3	UNIX server protection	331
9.4	Conclusion	335
Chapter 10. Virtual server security solutions		337
10.1	Virtualization defined	338
10.2	Virtualization threats	342
10.2.1	Virtual machine sprawl	343
10.2.2	Management console	344
10.2.3	Console operating system	344
10.3	IBM Virtual Server security solutions	344
10.3.1	Securing the virtual machine	345
10.3.2	Securing the management console	346
10.3.3	The IBM Security Network IPS Virtual Appliance	346

10.3.4 IBM Security Virtual Server Protection for VMware	348
10.4 IBM Security Virtual Server Protection for VMware component model .	352
10.4.1 Logical components	353
10.4.2 Physical components	358
10.4.3 Deployment architecture	359
10.5 Conclusion.	360
Chapter 11. Security services for Network, Server and Endpoint	363
11.1 Professional Security Services	366
11.1.1 Penetration Testing Service	367
11.1.2 Information Security Assessment	372
11.1.3 Deployment and Migration Services	378
11.1.4 Staff Augmentation Services.	382
11.1.5 Emergency Response Services	385
11.1.6 SCADA Assessment Service	389
11.2 Managed Security Services.	393
11.2.1 IBM MSS personnel qualifications	394
11.2.2 IBM MSS architecture	394
11.2.3 IBM Security Operations Centers	397
11.2.4 Managed Protection Services	403
11.2.5 Monitored and Managed Firewall Service.	405
11.2.6 Managed IDS and IPS Services for network and server.	408
11.2.7 Security Event and Log Management Services	411
11.2.8 Virtual-SOC Portal.	412
11.3 Cloud Security Services	417
11.3.1 Security challenges in the cloud	419
11.3.2 Advantages of cloud-based security	420
11.3.3 Using cloud-based security services.	420
11.3.4 Security for the cloud versus security from the cloud	421
11.3.5 IBM Security Services for the cloud	423
11.3.6 IBM Security Services from the cloud	425
11.4 Conclusion.	436
Part 3. Business scenarios	437
Chapter 12. A-B-C Government Agency	439
12.1 Company overview	440
12.1.1 Current IT infrastructure	441
12.1.2 Security issues within the current infrastructure	446
12.2 Business vision	447
12.3 Business requirements	448
12.3.1 IBM Security Framework mapping to business requirements.	448
12.4 Functional requirements	450
12.4.1 IBM Security Blueprint mapping to functional requirements	451

12.5 Design approach	454
12.6 Implementation approach	456
12.7 Conclusion	458
Chapter 13. X-Y-Z Cardio	459
13.1 Company overview	460
13.1.1 Current IT infrastructure	460
13.1.2 Security issues within the current infrastructure	467
13.2 Business vision	468
13.3 Business requirements	469
13.3.1 IBM Security Framework mapping to business requirements	470
13.4 Functional requirements	471
13.4.1 IBM Security Blueprint mapping to functional requirements	473
13.5 Design approach	475
13.6 Implementation approach	478
13.7 Conclusion	481
Related publications	483
IBM Redbooks	483
Other publications	483
Online resources	484
How to get Redbooks	484
Help from IBM	484

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	Lotus®	SecurityFusion™
AppScan®	Netcool®	Service Request Manager®
CICS®	Notes®	Smarter Planet™
Common Platform®	Proventia®	System z®
DataPower®	RACF®	Tivoli®
DB2®	Rational®	Virtual Patch®
Domino®	Real Secure®	WebSphere®
IBM®	Redbooks®	X-Force®
IMS™	Redguide™	z/OS®
InfoSphere™	Redpapers™	
Lotus Notes®	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Juniper, and Portable Document Format (PDF) are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Novell, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Threats come from a variety of sources. Insider threats, as well as malicious hackers, are not only difficult to detect and prevent, but many times the authors of these threats are using resources without anybody being aware that those threats are there.

Threats would not be harmful if there were no vulnerabilities that could be exploited. With IT environments becoming more complex every day, the challenges to keep an eye on all potential weaknesses are skyrocketing.

Smart methods to detect threats and vulnerabilities, as well as highly efficient approaches to analysis, mitigation, and remediation, become necessary to counter a growing number of attacks against networks, servers, and endpoints in every organization.

In this IBM® Redbooks® publication, we examine the aspects of the holistic Threat and Vulnerability Management component in the Network, Server and Endpoint domain of the IBM Security Framework. We explain the comprehensive solution approach, identify business drivers and issues, and derive corresponding functional and technical requirements, which enables us to choose and create matching security solutions.

We discuss IBM Security Solutions for Network, Server and Endpoint to effectively counter threats and attacks using a range of protection technologies and service offerings. Using two customer scenarios, we apply the solution design approach and show how to address the customer requirements by identifying the corresponding IBM service and software products.

This book is a valuable resource for senior security managers, security officers, and security architects who want to understand and implement enterprise security by following architectural guidelines.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.



Axel Buecker is a Certified Consulting Software IT Specialist at the International Technical Support Organization, Austin Center. He writes extensively and teaches IBM classes worldwide on areas of Software Security Architecture and Network Computing Technologies. He holds a degree in computer science from the University of Bremen, Germany. He has 24 years of experience in a variety of areas related to Workstation and Systems Management, Network Computing, and e-business Solutions. Before joining the ITSO in March 2000, Axel worked for IBM in Germany as a Senior IT Specialist in Software Security Architecture.



Kent Browne is a Worldwide Principal Security Architect for IBM Security Solutions. With more than two decades of experience in networking and security, Kent is one of those responsible for the technical vision and architecture of security solutions, and is regarded as an expert in technical subject matter for the IBM Internet Security Systems protection solution suite. He is also well versed in physical security infrastructure and strategy, and is asked to speak regularly on the convergence between physical and logical (network) security.

Kent has spoken at leading industry events, such as RSA, Interop, InfoSec, Bicsi, ISC, Infragard, ISACA, and The Conference Board. He has been a featured guest on television (BBC, RAI, ABC News, Silicon Spin, and so on), radio, and in magazines (Computer World, Business News, New York Times, Wired, and so on) to name a few.

Kent received a Bachelors Degree in Business Management/Marketing, Summa Cum Laude, from American Intercontinental University. He has been a penetration tester for the public and private sector, and uses this unique view of security to help others better understand the taxonomy of an attack and the strategies to prevent it.



Louis Foss is an Executive IT Architect supporting US federal customers that use IBM Software and IBM Security Solutions. Lou has 25 years of experience in the IT Industry and worked 18 years as an IT Specialist and Solutions Architect in a multitude of different roles, and projects, across the commercial and federal industries. Lou holds a Bachelor of Science degree in Computer Networking and a Master of Science degree in E-Commerce with a graduate certificate in Information Assurance from the University of Maryland University College. Lou is both a Certified Information Systems Security Professional (CISSP) and a VMware Certified Professional (VCP).



Jaco Jacobs is a Security and Privacy Consultant and Architect in the IBM Global Technology Services Middle East and Africa Technical Practice, where he delivers a wide variety of IBM Information Security services and solutions throughout the Middle East and Africa. He is a Certified Information Systems Security Professional and holds numerous Information Security Certifications. He has been a practicing Information Security Professional since 1998 with extensive experience in Vulnerability Management, Compliance Management, Incident and Event Management, and Intrusion Management across the Network, Server and Endpoint domain.



Vladimir Jeremic is a Security Enablement Instructor for the IBM Security Solutions portfolio. He primarily focuses on IBM Security Solutions for Network, Server and Endpoint (formerly known as the IBM Internet Security Systems (IBM ISS) portfolio). He has experience in designing, developing, and delivering learning materials. Vladimir also worked for many years as a Certified Security Managing Consultant with the IBM Global Services team, where he focused on architecture and implementation of the IBM Tivoli® Security portfolio. He has over ten years of experience in the IT field related to security, networking, and programming. He is Tivoli Certified Professional, IBM Certified Consultant, and he holds a Bachelor of Science degree in Electrical Engineering from the University of Novi Sad, in Serbia.



Carsten Lorenz is a certified Senior Managing Consultant at IBM UK and leads a team of security solution consultants and architects. He is responsible for managing, reviewing, and approving security solutions for large and complex IT infrastructure and business process outsourcing engagements for customers throughout Europe, the Middle East, and Africa. Carsten is involved in projects about IBM Security strategy and is currently engaged in the further advancement of the IBM Security Framework and the IBM Security Blueprint. With more than 10 years of experience in the security and compliance field, Carsten specializes in the areas of Security Governance, IT Risk Assessment, Compliance Management, Ethical Hacking, and Operational Risk Management. During his career, Carsten has engaged in consulting engagements with IBM customers in various industries, ranging from Fortune 500 companies to SMBs. Carsten is a CISSP, a CISA, and a CISM, and holds a Bachelors degree in European Studies from the University of Wolverhampton, UK, and a Masters degree in Business Science from the University of Trier, Germany.



Craig Stabler graduated from Heriot-Watt University in Edinburgh, Scotland in 1989 with a Master of Engineering (MEng) degree. Since graduating, Craig has built up his expertise in a number of European countries (UK, France, Benelux, and several East European countries) while working for various international data networking and security vendors, including Spider Systems, Shiva, Nortel, Internet Security Systems, and IBM. Craig has a keen interest in data security and completed his CISSP certification in 2009, which followed a CCNA certification that he completed in 2003. Craig was a member of the IBM Security X-Force® Research and Development Organization (X-Force) Roadshow team that presented in various European countries in 2007. Craig also presented the X-Force “24 hour vulnerability to malware life cycle” presentation at the ITWeb Security Summit in South Africa in May 2007.



Joris Van Herzele is an IT Specialist based in Brussels who provides research, design, and evaluation analysis for IBM Managed Security Services. He is a Certified Information Systems Security Professional with 10 years of experience in the network and security field and is a subject matter expert on threat management. Prior to his current role, Joris taught classes in the EMEA region through X-Force education services, and was a technical pre-sales engineer advocating the comprehensive product portfolio of IBM Security Solutions.

Thanks to the following people for their contributions to this project:

Greg Abelar, David Abercrombie, Jason Brewer, Leslie L. Burke, Tim Christensen, Tom Cross, Timothy Dodd, Jeffrey Douglass, Jonathan Fan, Brian Fitch, Robert Freeman, Bill Hines, Duncan Hoopes, James Innes, Naveed Makhani, David Ostrowski, Jeffrey Palatt, Peter Tosto, Lisa T. Washburn

IBM

Thanks to the authors of the previous editions of this book.

- Authors of the first edition, *Enterprise Security Architecture using IBM ISS Security Solutions*, SG24-7581, published in July 2008, were:

Per Andreas, Scott Paisley, Brian Reed, Rodrigo Antonio dos Reis, Prithvi Srihari

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>

Summary of changes

Because this is the second edition of the previously named IBM Redbooks publication *Enterprise Security Architecture using IBM ISS Security Solutions*, SG24-7581, you may expect an overview of all the details that were changed. Well, not in this case.

For this second edition, we have changed the title of the book to *IBM Security Solutions Architecture for Network, Server and Endpoint* to reflect its new positioning within the realigned IBM Security Solutions efforts.

The foundation to discuss a holistic approach to an enterprise wide security architecture is the IBM Redpapers™ publication *Introducing the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security*, REDP-4528. Because of the importance of this paper, we have included and extended the content of it in Part 1, “Business context and terminology” on page 1 of this book.

The IBM Security Framework maps out business related security issues into a total of six security domains:

- ▶ Governance, Risk and Compliance
- ▶ People and Identity
- ▶ Network, Server and Endpoint
- ▶ Data and Information
- ▶ Application and Process
- ▶ Physical Infrastructure

Based on the details provided in the IBM Security Framework, we have created this publication to focus on the IBM Security Solutions available to address the issues for the *Network, Server and Endpoint* security domain.

You will definitely still find some remnants of the IBM Internet Security Systems (IBM ISS) products you may have used in the past in this book. But we encourage you to walk through this completely redesigned IBM Redbooks publication to gain the maximum of technical understanding for addressing Network, Server and Endpoint related issues in your own environment.



Part 1

Business context and terminology

Security is a major consideration in the way that business and information technology systems are designed, built, operated, and managed. The need to be able to integrate security into discussions about business functions and operations exists more than ever.

In this part, we explore some of the concerns that characterize security requirements of, and threats to, business and information technology (IT) systems. We identify a number of the business drivers that illustrate these concerns, including managing risk and cost, and compliance to business policies and external regulations, showing how they can be translated into frameworks to enable enterprise security.

To help you with your security challenges, IBM has created a bridge to address the communication gap between the business and the technical perspectives of security to enable simplification of thought and process. The IBM Security Framework can help you translate the business view, and the IBM Security Blueprint describes the technology landscape view. In concert, they can help bring together the experiences we gained from working with many clients to build a comprehensive solution view.

The IBM Security Blueprint expands on this business oriented view of the IBM Security Framework by mapping the domains to a core set of security components representing capabilities and services. The IBM Security Blueprint aims to describe these security capabilities in vendor and product agnostic terms, using common, accepted industry definitions.

After we introduce the IBM Security Framework and IBM Security Blueprint, we take a closer look at architectures and network models. We also use the infrastructure context to introduce general concepts, such as IT perimeter security and preemptive security.



Introducing the IBM Security Framework and IBM Security Blueprint

To set the scene for the IBM Security Framework and IBM Security Blueprint, we start with a discussion of the typical business context when it comes to information technology (IT) security and how business leaders can use security, risk, and compliance related investments to competitively position their organization and satisfy complex regulatory guidelines. We describe two existing frameworks:

- ▶ CoBiT
- ▶ ISO27002

The remainder of this chapter is dedicated to introducing the IBM Security Framework and the IBM Security Blueprint.

1.1 Business context for IT security

Organizations rely on computing systems and automation more than ever to detect threats to intellectual property, reputation, and privacy. These organizations often adopt a piecemeal or technology-driven approach to security. Using this approach alone does not provide sufficient protection for business processes and assets against these business risks.

As the pace of globalization continues, traditional boundaries between organizations continue to disappear. The ideal response involves a level of planning and assessment to identify risks across key business areas, including people, processes, data, and technology throughout the entire business continuum. It is important to plan a holistic approach that can facilitate a business-driven security blueprint and strategy that can act as an effective shield of defense for the entire organization.

We believe that organizations have to build services that are *secure by design*, meaning that security is intrinsic to their business processes, their product development, and their daily operations. It is factored into the initial design, not bolted on afterwards. This allows an organization to securely and safely adopt new forms of technology, like cloud computing or virtualization, and business models like tele-working and outsourcing can be more safely leveraged for cost benefit, innovation, and shorter time to market.

With the security domains, capabilities, and services as a backdrop, this first section covers a detailed overview of the IBM Security Framework and IBM Security Blueprint. In the later sections we explain the IBM Security Blueprint in more detail by discussing its components and subcomponents. Later we take a closer look at the business context for areas such as identity management. We then look at the IBM Security Framework mapping and use the IBM Security Blueprint components and subcomponents and how they map to the needs of this scenario.

1.2 Drivers that influence security

Most of today's projects are driven by both business and IT drivers, although we can probably agree that business drivers are almost always the initiating factor. Let us take a closer look at these influencing factors:

- ▶ Business drivers: Business drivers measure value, risk, and economic costs that influence their approach to IT security. Value drivers determine the worth of assets of the system to the business and of the business itself.

Risk drivers involve compliance, corporate structure, corporate image, and the risk tolerance of the company. Economic drivers determine productivity impact, competitive advantage, and system cost.

- IT drivers: IT drivers represent operational constraints in the general IT environment. For example, the complexity of a system, including its environment, that is exposed to internal and external threats presents risks that the organization must address.

Business drivers also represent issues and consequences of significance to the stakeholders of the managed business system. This set of drivers might vary from industry to industry, from organization to organization in the same industry, and even from different business applications in an organization.

IT drivers represent technical considerations that affect the trustworthiness of the IT environment and likely the managed business systems as a whole. IT drivers are universal and must be considered within the context of the business drivers in all efforts. The combination of business and IT drivers represents the key initiatives for security management.

1.2.1 Business drivers that influence security

Business drivers represent a relationship between the IT organization and the rest of the business. They refer to business values that must be supported by the IT security infrastructure.

Correct and reliable operation

Correct and reliable operation is the degree to which the business must be accurate and consistent in its operation. Correct operation means that the operations perform the proper response or function with no errors. Reliable means that the same result occurs all the time. Any IT system must consistently provide stakeholders with the expected results.

Security events and incidents might impact the correct and reliable operation of these business processes. It might also affect the underlying IT infrastructure or upstream and downstream business processes. The consequences of a defective service (incorrect or varying results over time) might be significant to the consumer of the service, and therefore to the provider of the service.

Service-level agreements

This driver applies to circumstances where security threats and threat agents can impact an organization's ability to conduct business. Service-level agreements (SLAs) incorporate acceptable conditions of operation within an organization. SLAs might vary from business system to business system or application to application. Availability of systems, data, and processes is a condition commonly referenced within SLAs.

IT asset value

From a business perspective the IT asset value is directly related to the value of the business transactions that it supports. These might be tangible or intangible. For an e-retailer, these are tangible assets. For a financial services company, the asset might be the client information or other data used in transactions of the system.

Protection of the business asset value or brand image

This driver captures the firm's desire to protect its image. The loss of good will from a security incident or attack has a direct consequence to the business. Therefore, the security measures are likely to be proportional to the consequence. When the desire to avoid negative publicity increases, upon encountering a security breach, the stipulation for this driver becomes stronger.

Legal and regulatory compliance

Legal and regulatory compliance refers to the externally imposed conditions on the transactions in the business system and the company. This includes the rules and policies imposed by regulatory and government agencies. Civil, criminal liability, or regulatory penalty from a security incident or attack has a negative consequence on the business. Therefore, the amount of regulation and steps to ensure compliance should be factored in this driver. This includes privacy issues, the ability to prove the transaction initiator, and proving compliance.

Contractual obligation

Security measures for an IT system are likely to be proportional to the consequences encountered when the business encounters contractual liability from a security attack. Depending on the structure and terms of the contract, the consequence might lead to financial loss or liability. For example, when security incidents are encountered, the business might be unable to fulfill its contractual obligations of providing goods or services.

Financial loss and liability

Direct or indirect financial loss is a consequence to the business as a result of a security incident. Direct loss might include theft of asset, theft of service, or fraud.

Indirect loss might include loss based on civil or criminal judgment, loss of good will, or re-prioritized budget allocation. This driver identifies the fact that security measures for an IT system are likely to be in proportion to these consequences.

Critical infrastructure

This driver applies where security threats or threat agents can have a major impact on services or resources that are common to, or shared among, a community of businesses, the population at large, or both. Examples include telecommunications, electrical power, transportation systems, computing, and so on. The loss of critical infrastructure by its provider might have a ripple effect, causing secondary losses and driving security decisions for those affected. An important part of risk analysis is identifying critical infrastructure.

Safety and survival

This driver applies where security threats and threat agents can have a major impact on aspects of human life, government function, and socio-economic systems. Examples of processes to be considered for safety and survival impact include continuity of critical infrastructure, medical system, life support, or other high-impact or time-dependent process.

1.2.2 IT drivers that influence security

IT drivers comprise the second group of key security initiatives. These are considered universal drivers that must be considered in every modern IT solution in a manner commensurate with the risks and consequences of a related failure or incident.

Internal threats and threat agents

Security-related failures and incidents are caused by threats or threat agents found within the physical and logical boundaries of the organization or enterprise that operates and controls the IT system. These threats and threat agents might be associated with technology or people.

An example of an internal threat is a poorly designed system that does not have the appropriate controls. An example of a internal threat agent is a person who would use his ability to access the IT system or influence business or management processes to carry out a malicious activity.

External threats and threat agents

Security-related failures and incidents are caused by threats or threat agents found outside the physical and logical boundaries of the organization or enterprise that operates and controls the IT system. These threats and threat agents are also associated with technology or people. They seek to either penetrate the logical or physical boundary to become internal threats or threat agents, or to influence business or management processes from outside the logical or physical boundary.

Examples of external threats are single points of failure for one or more business or management processes that are outside the enterprise boundary, such as a power system grid or a network connection, or a computer virus or worm that penetrates the physical or logical network boundary. An example of an external threat agent is a hacker, or someone who has gained the ability to act as an insider, using personal electronic credentials or identifying information.

IT service management commitments

This driver identifies the fact that failure to manage the operation of the IT system might result in security exposures to the business. This driver can be divided into two categories, IT service delivery and IT service support.

- ▶ **Service delivery commitments**

The failure of the IT system to meet its metrics for managing itself can be viewed as a security exposure to both business or management processes.

An example of security exposure for service delivery is when IT operations processes cannot respond to critical events in a timely manner. Another is when IT resilience processes cannot recover from a denial of service attack in a timely manner, resulting in a loss of capacity or response time for business processes.

- ▶ **Service support commitments**

The failure of the business or IT management system to meet its service-level agreements can be viewed as a security exposure to business or management processes.

An example of security exposure for service support is a situation in which the customer relationship processes do not add, modify, or remove users from access control lists in a timely manner.

IT environment complexity

The complexity of the IT environment might contribute to the security or insecurity of the IT system. The IT environment reflects the infrastructure on which the business system will be placed.

For example, any IT environment that is connected to the intranet or extranet is exposed to internal or external threats or threat agents and requires specific security responses. A stand-alone facility for our system represents the lowest complexity. A hosting facility with other systems and other firms represents a more complex environment. An environment with a larger number of systems, varied network access paths, or a complex architecture, is a complex IT environment.

Business environment complexity

Because most businesses rely on IT, most business environments are an interconnected set of businesses, each with its own complex IT environment, business processes, and IT management processes. This complexity might contribute to the security or insecurity of the IT system.

Audit and traceability

This driver identifies the need for the IT system to support an audit of information contained within the system, whether it is associated with management data or business data.

IT vulnerabilities: Configuration

Configuration vulnerabilities are potentially present in every IT system, providing an opening to a potential attack based on the system and how it is designed and set up.

IT vulnerabilities: Flaws

Software flaws potentially exist in every IT system. These flaws represent vulnerabilities that were not detected and are not evident in the design documents. As such, they are an unexpected deviation from what was designed. An example is a defect in an operating system or application that is discovered after implementation.

IT vulnerabilities: Exploits

The basic design of software in any IT system might be exploited by threats or threat agents as a part of an attack on the IT system, the business, or the management processes. This might include the use of a function within a system in a way to compromise the system or underlying data. While certain people might define an exploit as both the flaw and the method, we treat them separately because an exploit might involve using normal functions as designed in an unusual manner to attack the system. The exploits can also be viewed as the openings or avenues that an attacker can use.

1.3 Common industry approaches to IT security management

IT security management is the term used for the set of management activities that are intended to address the business and technical issues described earlier, in accordance with the resilience and risk management objectives for the managed business system.

The business reasons depicted in 1.2.1, “Business drivers that influence security” on page 5 are leading to an evolving number of enterprises that adopt internationally accepted frameworks and best practices to help implement IT governance in their enterprise. Control Objectives for Information and related Technology¹ (CobiT), the International Organization for Standardization 27002:2005² (ISO/IEC 27002:2005), and the Information Technology Infrastructure Library³ (ITIL) have emerged worldwide as the most respected frameworks for IT governance and compliance. We take a closer look at CobiT and ISO/IEC 27002:2005 in the following sections because they have—in contrast to ITIL, which is more focussed on IT service management elements—a strong focus on IT security.

1.3.1 Control objectives for information and related technology

CobiT is a set of best practices (framework) for IT management created by the Information Systems, Audit and Control Association (ISACA) and the IT Governance Institute (ITGI) in 1996. It is an internationally accepted framework for IT governance and control. The current edition, 4.1, issued by the IT Governance Institute in 2007, includes the following sections:

- ▶ Executive summary (explains CobiT key concepts and principle)
- ▶ CobiT framework (explains the CobiT approach)
- ▶ Control objectives (defines a generic set of control requirements that need to be managed for each IT process to get effective control)
- ▶ Management guidelines (explains tools to measure, compare, and improve the performance of IT processes)

¹ For more information about CobiT, go to <http://www.isaca.org/Template.cfm?Section=COBIT6&Template=/TaggedPage/TaggedPageDisplay.cfm&TPLID=55&ContentID=7981>.

² To purchase a copy of ISO/IEC 27002:2005, go to http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=50297.

³ For more information about ITIL®, go to <http://www.itil-officialsite.com/home/home.asp>.

- ▶ Implementation guide (provides a tool set to implement CobiT)
- ▶ IT Assurance guide (explains methods to assess whether control objectives are achieved)

The underlying concept of CobiT is that it looks at *business information* that every enterprise needs to support its business decisions. Business information itself is again a result of IT-related resources, which CobiT defines as *applications, information, infrastructure, and people*. Finally, these IT-related resources are managed by IT processes to fulfill certain business information criteria (effectiveness, efficiency, confidentiality, integrity, availability, reliability, and compliance). CobiT defines 34 high-level processes that are grouped into the following four domains:

1. Plan and organize.

This domain focuses on IT strategy: How can IT contribute to business objectives?

2. Acquire and implement.

The topic of this domain is the identification, development, or acquisition and integration of IT solutions to realize IT strategy.

3. Deliver and support.

This domain is about delivering and supporting the entire range of IT services.

4. Monitor and evaluate.

This domain focuses on the continuous assessment of all IT process to ensure their quality and compliance.

These 34 processes are controlled by 210 control objectives. Therefore, choose a top-down approach when implementing CobiT, because business objectives must be clearly defined before the IT strategies can be aligned.

1.3.2 ISO/IEC 27002:2005

The British Standard 7799⁴ that preceded the International Organization for Standardization 27002:2005 (ISO/IEC 27002:2005) is the most widely recognized security standard in the world. The last major publication was in May 1999, an edition that included many enhancements and improvements over previous versions. When republished in December 2000, it evolved into the International Organization for Standardization 17799 (ISO/IEC 17799).

⁴ Information about RiskServer, Security Risk Analysis, ISO17799, Information Security Policies, and Audit and Business Continuity can be found at <http://www.riskserver.co.uk/>.

17799 was republished again in 2005 as ISO/IES 17799:2005(E) with more revisions. In 2007, the name of ISO17799 was, without further amendment, adapted to the new ISO/IEC numbering scheme for information security management standards and is now identified as ISO/IEC 27002:2005.

ISO/IEC 27002:2005 is comprehensive in its coverage of security issues. It contains a significant number of control requirements, some extremely complex. Compliance with ISO/IEC 27002:2005 is, consequently, a far from trivial task, even for the most security conscious of organizations.

A step-by-step manner of approaching ISO/IEC 27002:2005 is best. The best starting point is usually an assessment of the current position or situation, followed by an identification of the changes needed for ISO/IEC 27002:2005 compliance. From here, planning and implementing must be rigidly undertaken.

ISO/IEC 27002:2005 contains 11 categories that have to be considered when applying an overall enterprise security approach. The categories are:

- ▶ Security policy
- ▶ Organization of information security
- ▶ Asset management
- ▶ Human resources security
- ▶ Physical and environmental security
- ▶ Communications and operations management
- ▶ Access control
- ▶ Information systems acquisition, development, and maintenance
- ▶ Information security incident management
- ▶ Business continuity management
- ▶ Compliance

Now it is time for us to introduce the IBM Security Framework, which focuses on the *what*, not the *how*. It can help you translate your requirements into coarse-grained business solutions, not into specific IT components or IT services.

1.4 IBM Security Framework

Today's business leaders are expected to manage risk in their areas of responsibility in the same way that CFOs manage risks in their domains. Security risks and the potential impact on IT need to be communicated to executive peers in business terms. Additionally, they need to align IT security controls with their business processes, monitor and quantify IT risk in business terms, and dynamically drive business-level insight at the executive level.

Finally, business leaders need to manage risk and orchestrate security operations in a way that enforces compliance and optimizes business results.

As an organization secures its business processes, a business-driven approach needs to become the guiding influence for ensuring that all the different security domains work together in a holistic and synergistic manner, in alignment with the overarching business objectives. Otherwise, the organization's risk stance becomes vulnerable due to misalignment of priorities between IT and the business strategy. Using a standards-based approach to map business drivers to IT security domains is often difficult and is often an afterthought.

IBM created a comprehensive IT security framework (Figure 1-1) that can help ensure that every necessary IT security domain is properly addressed when using a holistic approach to business-driven security.

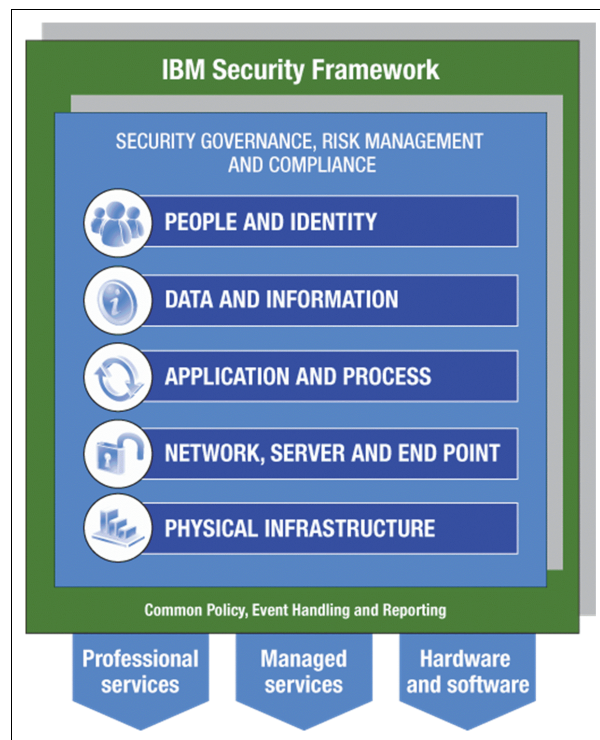


Figure 1-1 The IBM Security Framework

IBM provides the full breadth and depth of solutions and services that can enable organizations to take this business-driven, secure by design approach to security in alignment with the IBM Security Framework.

Comprehensive professional services, managed services, and hardware and software offerings are available from IBM to support your efforts in addressing the following security domains covered by the IBM Security Framework.

1.4.1 Security Governance, Risk Management, and Compliance

Every organization needs to define and communicate the principles and policies that guide the business strategy and business operation. In addition, every organization must evaluate its business and operational risks, and develop an enterprise security plan to serve as a benchmark for the execution and validation of the security management activities that are appropriate for their organization.

These principles and policies, the enterprise security plan, and the surrounding quality improvement process represent the enterprise Security Governance, Risk Management and Compliance model. Specifically, the requirements and the compliance criteria for the remaining security domains are:

- ▶ **People and Identity**

This domain covers aspects about how to ensure that the correct people have access to the correct assets at the correct time.

- ▶ **Data and Information**

This domain covers aspects about how to protect critical data in transit or at rest across the organization.

- ▶ **Application and Process**

This domain covers aspects about how to ensure application and business services security.

- ▶ **Network, Server and Endpoint (IT infrastructure)**

This domain covers aspects about how to stay ahead of emerging threats across IT system components.

- ▶ **Physical Infrastructure**

This domain covers aspects about how to use the capability for digital controls to secure events—on people or things—in the physical space.

Let us now take a closer look at these domains.

1.4.2 People and Identity domain

Organizations need to protect the assets and services that serve the business and support the business operation. One aspect of protection is provided by *access control*. The ability to provide effective access control services is based on the ability to manage people and identity as defined by the enterprise's security governance, risk, and compliance model.

The Security Governance, Risk Management, and Compliance model provides guidance about how identities are managed and how access control is to be conducted. Organizations register people and map them to identities. The relationships between people and organization are expressed in terms of role, rights, business policies, and rules. The ability to register people and describe their relationship with the enterprise is a key security enabler for the remaining security domains:

- ▶ Data and Information
- ▶ Applications and Process
- ▶ Network, Server and Endpoint (IT infrastructure)
- ▶ Physical Infrastructure

Operationally, people acting in authorized roles in an organization or as part of an extended relationship are granted access to infrastructure, data, information, and services. At the same time, people acting in unauthorized roles are denied access to infrastructure, data, information, and services if they are acting outside of the business policies and agreements.

Within an identity system, people can be issued a *credential*. A credential can take any of several forms, including a physical identity card or logical token or user identifier. The *trustworthiness* or *strength* of the credential is an important aspect of business policy and risk management. The ability to effectively manage the life cycle of identity, that is, the creation, removal, and role changes for dynamic populations of workforce, customer, or user communities, is extremely important. For example, the life cycle of identities and credentials can be influenced by business cycles, employment cycles, customer relationship, agreement, business, or calendar events, and so on.

Identity systems need to be integrated with appropriate sets of access controls. Identity systems need to manage user roles, rights, and privileges across the IT infrastructure that might contain multiple technology architectures, or multiple identity and access control systems will be required to ensure that users have access to the correct assets and services.

Compliance for identity and access is often externally motivated compliance. For example, legislated privacy and evidence recording is a significant driver for implementation of comprehensive user provisioning and identity-related record keeping.

Figure 1-2 shows a summary and additional aspects to be addressed within the People and Identity domain.

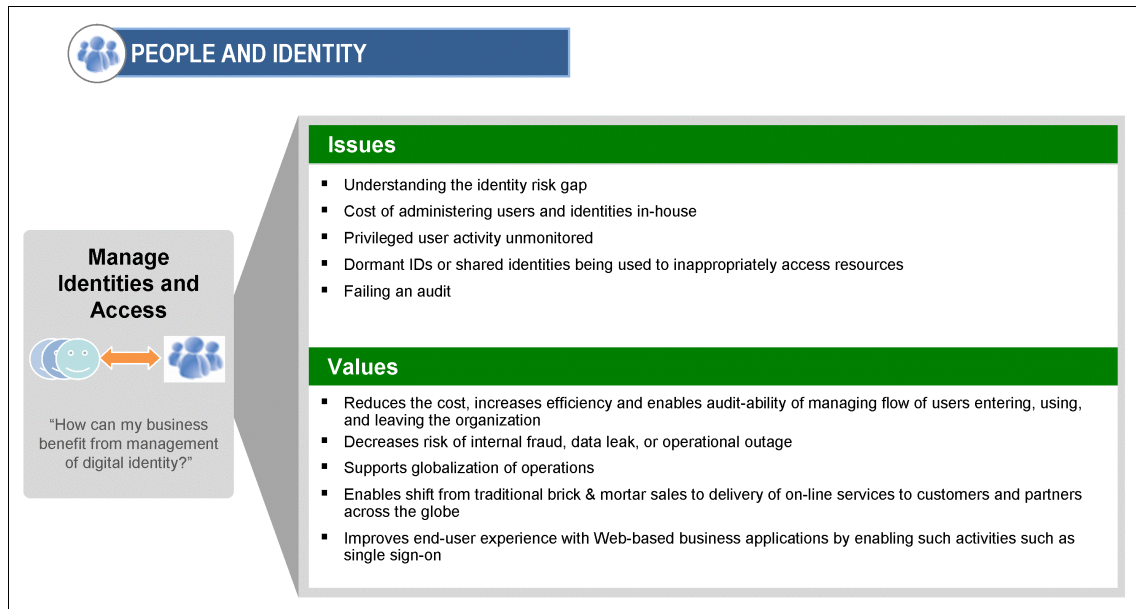


Figure 1-2 People and Identity domain

1.4.3 Data and Information domain

Organizations need to protect both the *raw data* and *contextualized information* that is within its span of control. The Security Governance, Risk Management, and Compliance model provides guidance about the value of data and information and how the risks to data and information must be managed.

An effective plan for data and information protection includes maintaining a catalog or inventory of these assets, along with attributes, policies, and enforcement mechanisms and services that govern the access, transformation, movement, and disposition of data and information.

This data and information protection plan can be applied to business processes, business transactions, or business and infrastructure support processes. The protection of data and information covers a full life cycle, from creation to destruction and across its various states, locations, and instantiations, and when it is stored or when it is being physically or electronically transported.

The term *data* can be applied to a wide range of electronically encoded assets. This includes software and firmware, which need to be protected against technical risks (to ensure that malicious code is not introduced) and business risks (to ensure that licensing terms have not been violated).

Protection of data and information is dependant on the definition and operation of all other operational security domains. Measuring and reporting on an organization's compliance with respect to protection of data and information is a tangible metric of the effectiveness of the enterprise security plan. A *data and information compliance report* reflects the strength or weakness of controls, services, and mechanisms in all domains.

Figure 1-3 shows a summary and additional aspects to be addressed within the Data and Information domain.

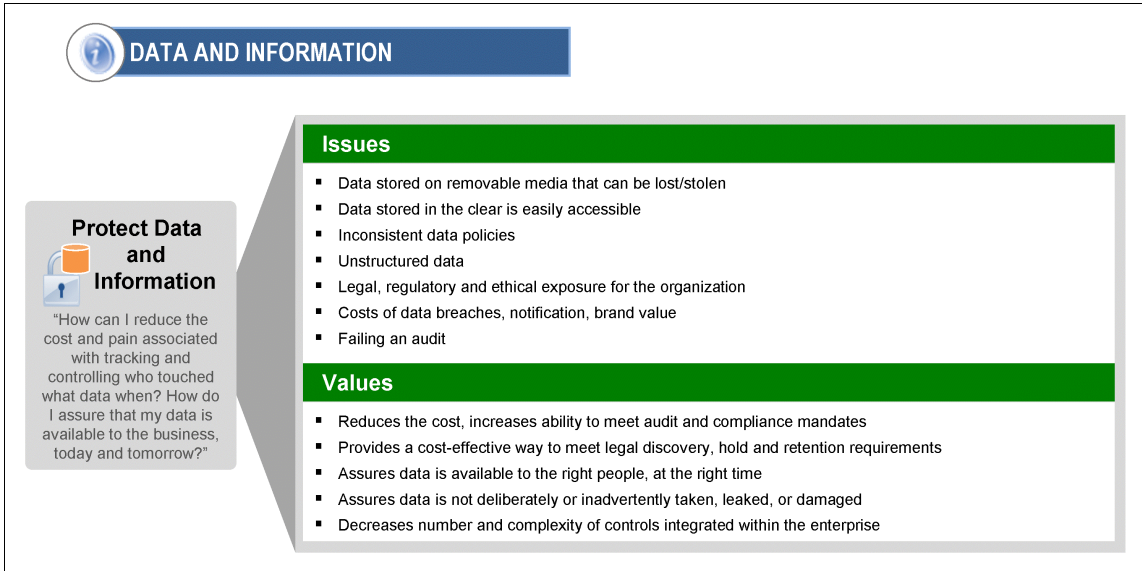


Figure 1-3 Data and Information domain

1.4.4 Application and Process domain

Organizations need to proactively protect their *business-critical applications* from external and internal threats throughout their entire life cycle, from design to implementation and production. Control throughout the application life cycle implies effective control and compliance in the remaining security domains. For example, whether an application is internally focused, such as a customer relationship management (CRM) system delivered through a service-oriented architecture (SOA), or is an externally facing application, such as a new customer portal, clearly defined security policies and processes are critical to ensure that the application is enabling the business rather than introducing additional risk.

Service management for all business and business support processes, including service management for processes within the security domain, is a critical part of ensuring that the business is operating within the appropriate risk management and compliance guidelines. Service management of security typically includes a combination of capabilities, such as centralized authentication, access and audit policy management, and web application vulnerability scanning and intrusion prevention.

Figure 1-4 shows a summary and additional aspects to be addressed within the Application and Process domain.

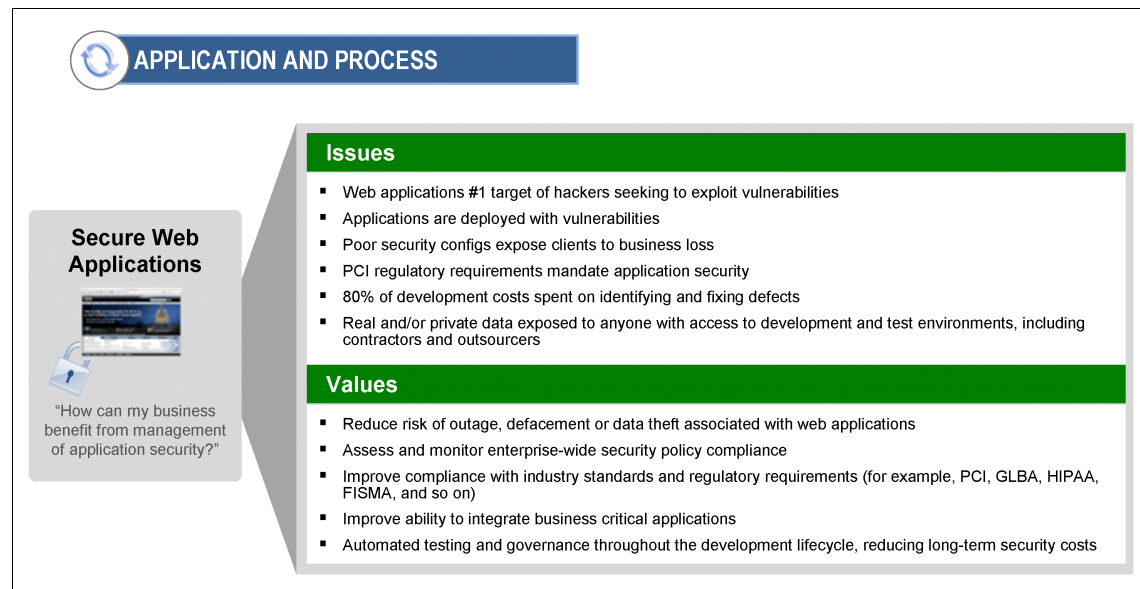


Figure 1-4 Application and Process domain

1.4.5 Network, Server and Endpoint domain

Organizations need to *preemptively* and *proactively monitor* the operation of the business and the IT infrastructure for *threats* and *vulnerabilities* to avoid or reduce breaches.

The Security Governance, Risk Management, and Compliance model can provide guidance on the business implications of technology-based risks. In practice, the definition, deployment, and management of technology-based threats, as well as the technical aspects of incident response, can be delegated to operational management and staff, or outsourced to a service provider.

Security monitoring and management of an organization's Network, Server and Endpoints is critical to staying ahead of emerging threats that can adversely affect system components and the people and business processes that they support. The need to identify and protect the infrastructure against emerging threats has dramatically increased with the rise in organized and financially motivated network infiltrations. While no technology is perfect, the focus and intensity of security, monitoring, and management can be affected by the type of Network, Server and Endpoints deployed in the IT infrastructure and how those components are built, integrated, tested, and maintained.

Organizations use *virtualization technology* to support their goals of delivering services in less time and with greater agility. By building a structure of security controls within this environment, organizations can reap the goals of virtualization—such as improved physical resource utilization, improved hardware efficiency, and reduction of power costs, while gaining peace of mind that the virtual systems are secured with the same rigor as the physical systems.

Figure 1-5 shows a summary and additional aspects to be addressed within the Network, Server and Endpoint domain.

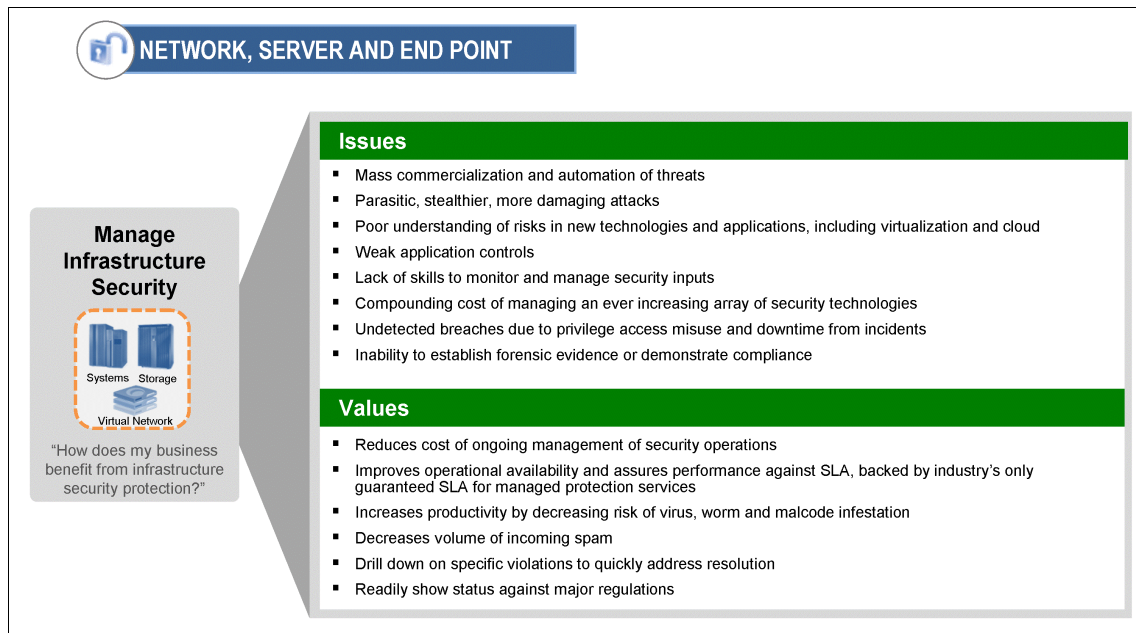


Figure 1-5 Network, Server and Endpoint domain

1.4.6 Physical Infrastructure domain

For an organization to effectively implement an enterprise security plan, the business and technical risks that are associated with the physical infrastructure must be understood and addressed. Security Governance, Risk Management, and Compliance provides guidance on the types of risks and the types of plans and responses for physical security.

Protecting an organization's infrastructure can mean taking precautions against a failure or loss of physical infrastructure that might impact business continuity. Protecting an organization's infrastructure can involve protection from indirect threats and vulnerabilities, such as the impact of loss of a utility service, a breach in physical access control, or loss of critical physical assets. Effective physical security requires a centralized management system that allows for correlation of inputs from various sources, including property, employees, customers, the general public, and local and regional weather.

For example, securing the perimeter of the data center with cameras and centralized monitoring devices is critical to ensure managed access to an organization's IT assets. Therefore, organizations concerned about theft and fraud, such as banks, retail stores, or public agencies, should define and implement an integrated physical security surveillance strategy that includes monitoring, analytics, and centralized control. This approach enables organizations to extract intelligent data from multiple sources and respond to threats sooner than manually monitored environments, resulting in reduced cost and risk of loss.

Figure 1-6 shows a summary and additional aspects to be addressed within the Physical Infrastructure domain.

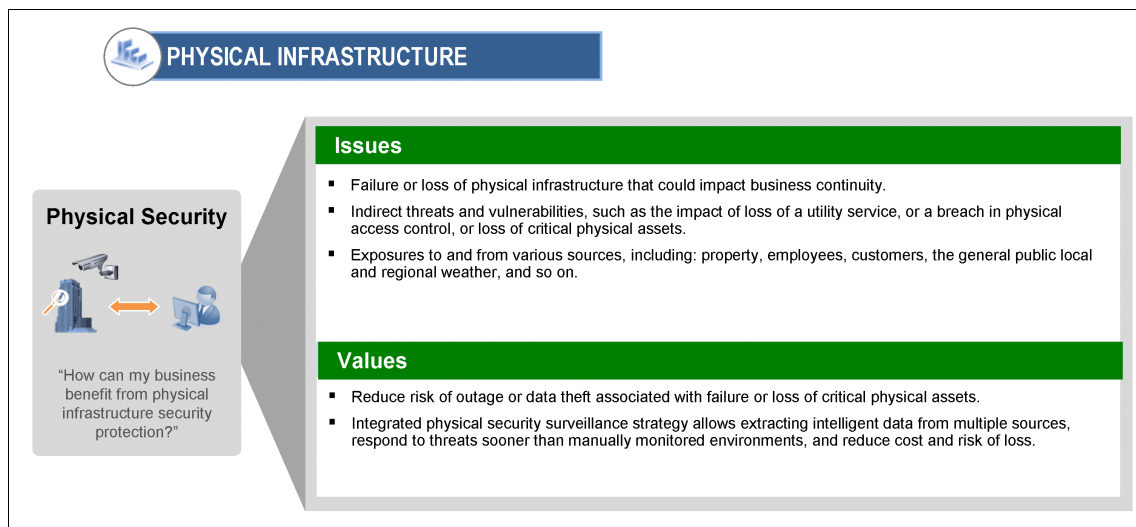


Figure 1-6 Physical Infrastructure domain

After having addressed and mapped the IT security domains into your business solutions, it is time to look at the component-oriented view of IT security in the IT Security Blueprint.

1.5 IBM Security Blueprint

The IBM Security Framework divides the area of business-oriented IT security into six domains. The next step is to break these down into further detail to work toward a common set of core security capabilities needed to help your organization securely achieve its business goals. These core security capabilities are called the *IBM Security Blueprint*.

The IBM Security Blueprint uses a product-agnostic and solution-agnostic approach to categorize and define security capabilities and services that are required to answer the business concerns in the IBM Security Framework.

The IBM Security Blueprint was created after researching many customer-related scenarios, focusing on how to build IT solutions. The intention of the blueprint is to support and assist in designing and deploying security solutions in your organization.

Building a specific solution requires a specific architecture, design, and implementation. The IBM Security Blueprint can help you evaluate these, but does not replace them. Using the IBM Security Blueprint in this way can provide a solid approach to considering the security capabilities in a particular architecture or solution.

IBM has chosen to use a high-level service-oriented perspective for the blueprint, based on the IBM service-oriented architecture approach. Services use and refine other services (for example, policy and access control components affect almost every other infrastructure component.)

To better position and understand the IBM Security Blueprint, see Figure 1-7.

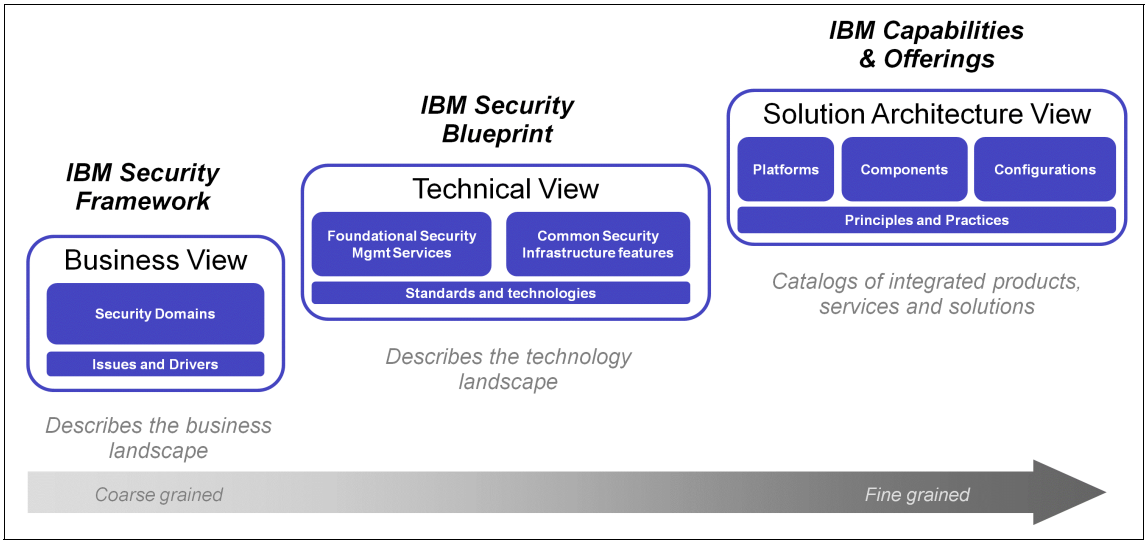


Figure 1-7 IBM Security Blueprint positioning

The left portion of Figure 1-7 represents the IBM Security Framework, which describes and defines the security domains from a business perspective. It was covered in 1.4, "IBM Security Framework" on page 12.

The middle portion in Figure 1-7 on page 22 represents the IBM Security Blueprint, which describes the IT security management and IT security infrastructure capabilities needed in an organization. As discussed earlier, the IBM Security Blueprint describes these capabilities in product and vendor-neutral terms.

The right portion of Figure 1-7 on page 22 represents the solution architecture views, which describe specific deployment guidance particular to a given IT environment. Solution architecture views provide details about specific products, solutions, and their interactions.

Figure 1-8⁵ shows the complete IBM Security Blueprint, and each layer and component are described in the following sections.

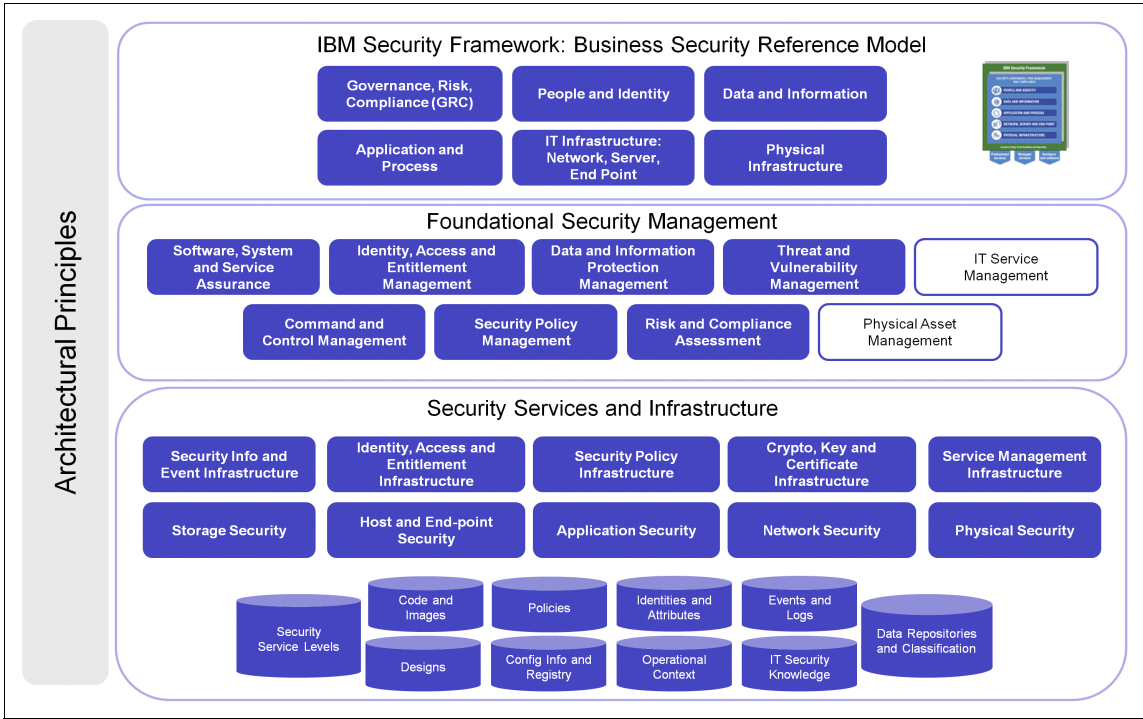


Figure 1-8 The IBM Security Blueprint

⁵ White boxes in Figure 1-8 on page 23 and other diagrams represent services or components that are not solely security related but might be connected with other IT service areas as well.

1.5.1 Foundational Security Management

The Foundational Security Management layer contains the top-level components used to direct and control IT security from a policy-based, risk management perspective. These components are described in more detail in Chapter 2, “The components of the IBM Security Blueprint” on page 31.

Let us take a closer look at each Foundational Security Management component:

- ▶ *Risk and Compliance Assessment* enables the IT organization to collect, analyze, and report security information and security events to identify, quantify, assess, and report on IT-related risks that can contribute to the organization's operational risk. This component covers *risk aggregation and reporting, IT security risk processes, business controls management, resiliency and continuity management, compliance reporting, and legal discovery services*.
- ▶ *Command and Control Management* provides the command center for *security management* and the *operational security capabilities* for non-IT assets and services to ensure protection, response, continuity, and recovery. It covers topics such as:
 - Ensuring that physical and operational security is maintained for locations, assets, humans, environment, and utilities
 - Providing surveillance and monitoring of locations, perimeters, and areas
 - Enforcing entry controls
 - Providing for positioning, tracking, and identification of humans and assets
 - Providing a focal point for continuity and recovery operations
- ▶ *Security Policy Management* provides all services and repositories to author, discover, analyze, transform, distribute, evaluate, and enforce security policies.
- ▶ *Identity, Access, and Entitlement Management* provides services related to roles and identities, access rights, and entitlements. The proper use of these services can ensure that access to resources has been given to the correct identities, at the correct time, and for the correct purpose. These services can also ensure that access to resources is monitored and audited for unauthorized or unacceptable use.
- ▶ *Data and Information Protection Management* provides services that protect unstructured and structured data from unauthorized access and data loss, according to the nature and business value of information. It also provides usage and access monitoring and audit services.

- ▶ *Software, System, and Service Assurance* addresses how software, systems, and services are designed, developed, tested, operated, and maintained throughout the software life cycle to create predictably secure software. This component covers:
 - Structured design
 - Threat modeling
 - Software risk assessment
 - Design reviews for security
 - Source code reviews and analysis
 - Dynamic application analysis
 - Source code control and access monitoring
 - Code/package signing and verification
 - Quality assurance testing
 - Supplier and third-party code validation
- ▶ *IT Service Management* provides the process automation and work flow foundation for security management. In particular, change and release management processes play a significant role in security management.
- ▶ *Threat and Vulnerability Management* provides services that identify vulnerabilities in deployed systems and receive reports of vulnerabilities from outside sources, determine the appropriate response, and make proactive changes to deployed systems to maintain the security of the deployed system.
- ▶ *Physical Asset Management* provides awareness of the location and status of physical assets and awareness of physical security controls and coordinates the security information for physical systems with the IT security controls.

1.5.2 Security Services and Infrastructure

The Security Services and Infrastructure layer contains components and sub-components that are being utilized by the Foundational Security Management components in their respective contexts:

- ▶ *Security Information and Event Infrastructure* provides the infrastructure to automate log aggregation, correlation, and analysis. It also enables an organization to recognize, investigate, and respond to incidents automatically, and streamline incident tracking and handling, with the goal of improving security operations and information risk management.
- ▶ *Identity, Access, and Entitlement Infrastructure* provides services to manage user provisioning, passwords, single sign-on, access control, and synchronization of user information across directories.

- ▶ *Security Policy Infrastructure* provides services to manage the development implementation of security policies in a consistent manner and automate the deployment of those policies to IT systems.
- ▶ *Cryptography, Key, and Certificate Infrastructure* provides services to perform cryptographic operations efficiently and provides operational processes and capabilities to manage cryptographic keys.
- ▶ *Network Security* consists of multi-layered network security to provide defense in-depth, deep inspection, and analysis of protocols, application level payloads, and user content to protect at all levels of the network stack. It extends to virtual networks for security in modern, heavily virtualized environments.
- ▶ *Storage Security* provides data-centric security capabilities for protecting data in use, in transit, and at rest through isolation and encryption capabilities. It also provides services to catalog and classify storage assets and associate control policies with them.
- ▶ *Host and End-point Security* provides protection for servers and user devices, such as mobile phones, desktop computers, and mobile computers using both host and network based technologies. This protection integrates into the virtualization infrastructure to provide security for virtual environments. It includes hardware-based attestation of host operating systems (OSs) and system resources to protect against malicious attacks.
- ▶ *Application Security* provides the infrastructure for testing, monitoring, and auditing deployed applications.
- ▶ *Service Management Infrastructure* consists of the infrastructure services to handle service management processes, such as incident, problem, change, and configuration management. Process automation is generic framework-based services to automate IT actions, including security-related activities.
- ▶ *Physical Security* is an IT infrastructure service to create awareness of physical security and coordinate it with IT security. This can include employee badges, RFID readers, surveillance systems, and associated technology or assets. Physical Security can include automation related to surveillance, motion detection, object and human identification and tracking, entry control, environmental system monitoring, perimeter control, and power and utility system monitoring.

1.5.3 Architectural principles

IBM security architects have defined the following *Architectural Principles* that accompany the service decomposition. These can be applied to all levels of the framework, blueprint, and solution designs, and are also guidelines for IBM products and solutions.

- Openness.

Openness is of primary importance in an enterprise environment. This includes support for all major platforms, run times, and languages, support for major industry standards, published interfaces and algorithms, no security by obscurity, documented trust and threat models and support for Common Criteria, and similar formal security validation programs.

- Security by default.

Security must not be an afterthought in IT solutions, but security policies must be secure out-of-the box. This is helped by a consistent definition and management of configurations, a consistent set of security roles and persona across products, and a consistent security management user interface.

- Design for accountability.

In today's environments, with many requirements in the compliance area, it is important that all security-relevant actions can be logged and audited, the audit infrastructure is scalable to handle these events, and audit information is immutable and non-reputable.

- Design for regulations.

Regulations drive many requirements in IT security projects, and regulations change over time. Handling this requires flexible support for the constraints set by government regulations and industry standards and traceability between regulations, standards, and business policies and the security policies used to implement them.

- Design for privacy.

In the current age of data sharing, privacy becomes increasingly important. Solutions must highlight the use of personally identifiable information and corresponding data protection mechanisms and enable the principles of notice, choice, and access.

- Design for extensibility.

Good solutions are component based and separate the management of mechanisms from the mechanisms themselves, to support a variety of mechanisms under the same framework. Already deployed systems must allow for the addition and extension of new mechanisms within the existing management framework.

- ▶ Design for sharing.

Multiple solutions can share a single IT environment, such as in a shared service center. To achieve this goal, security services and management must be able to span multiple domains, each domain potentially providing its own and independently set security policy, identity, models, and so on. Architectures must explicitly document the assumptions and limitations made in terms of span of control.

- ▶ Design for consumability.

All security services must be easily used by a variety of audiences. This includes programmers who develop and integrate applications with the security services, management systems that create, update, and manage security policies and other security artifacts, and people who manage security activities, audit security activities, and request access to protected resources.

- ▶ Multiple levels of protection.

Defense in depth is a general principle, which can be achieved by multiple levels of enforcement and detection. Resources must be designed to protect themselves as a first layer of defense. Intrusions can be contained through *isolation* and *zoning*. Multiple levels also minimize the attack surface to the outer-most accessible layer. *Least privilege* is a similar fundamental principle. Finally, the system should incorporate fail-safe principles.

- ▶ Separation of security management, enforcement, and accountability.

Security management services (identity, authorization, audit, and so on) are provided through a dedicated and shared security infrastructure, enabling consistent monitoring and enforcement. The enforcement itself (through cryptography, policy enforcement, or physical isolation) is typically distributed and kept close to the resources.

- ▶ Security-critical resources must be aware of their security context.

Resources and actors are kept aware of their environment (including physical location and logical co-location) and their security status and context.

- ▶ Security is model-driven.

Models are reflective of the operating environment, common models, and consistent formats for identity and trust, data, policy, applications, security information and events, and cryptographic keys. Models are consistently interpreted across the stack (for example, network identities are linked to application-level identities) and across units (for example, policies and trust are negotiated and understood within a federation). Models are consistently validated against reality (feedback from policy and model discovery).

- Consistency in approaches, mechanisms, and software components.

Two independent layers of protection for one resource might improve security. But using two different mechanisms for the same purpose for two resources increases the chances that at least one of them gets broken (plus, they increase management impact).

The IBM Security Blueprint lists the preferred standards and mechanisms.

This concludes the overview of the IBM Security Blueprint. In the next section, we discuss the components of the IBM Security Blueprint in more detail.



The components of the IBM Security Blueprint

In this chapter, we explain the IBM Security Blueprint in more detail by discussing the *components* and *subcomponents* of the IBM Security Blueprint.

The components in the IBM Security Blueprint describe the common security capabilities needed in any IT environment to manage IT security risks. Like the other elements of the IBM Security Blueprint, the components describe these security capabilities in vendor and product agnostic terms, using common, accepted industry definitions.

Each component is described in terms of the services that it provides, which can be combined with other components to create solution patterns. Key work products and artifacts for each component are also described, along with relevant industry standards.

The component descriptions often, but not always, correspond to market segments and product offerings. However, in many cases, a product offering might encompass multiple components. The intent of the components is not to describe a product or service taxonomy, but to provide a product and vendor agnostic way to describe IT security capabilities.

The components are organized into two layers. The *Foundational Security Management* components comprise the first layer. Each component is decomposed into a set of more detailed subcomponent descriptions. A set of key components in the *Security Services and Infrastructure* is identified on a second layer. While this section provides many details about the first layer, the second, more supportive layer is discussed more briefly because many terms should be familiar to the information technology (IT) security professional.

2.1 Foundational Security Management

In this section we explain the Foundational Security Management *components* of the IBM Blueprint and how they work together to govern the policies and deployed security capabilities in a way that supports the business objectives. Furthermore, we introduce their respective subcomponents.

The set of Foundational Security Management components form a closed loop management system. Figure 2-1 on page 33 depicts the continuous risk management cycle as it has to be practiced for comprehensive security management. *Command and Control Management* sets security directives and objectives, which are used by *Security Policy Management* to produce and set the policies and standards that have to be adhered to in the other functional areas of security, as they are represented on the right side of Figure 2-1 on page 33.

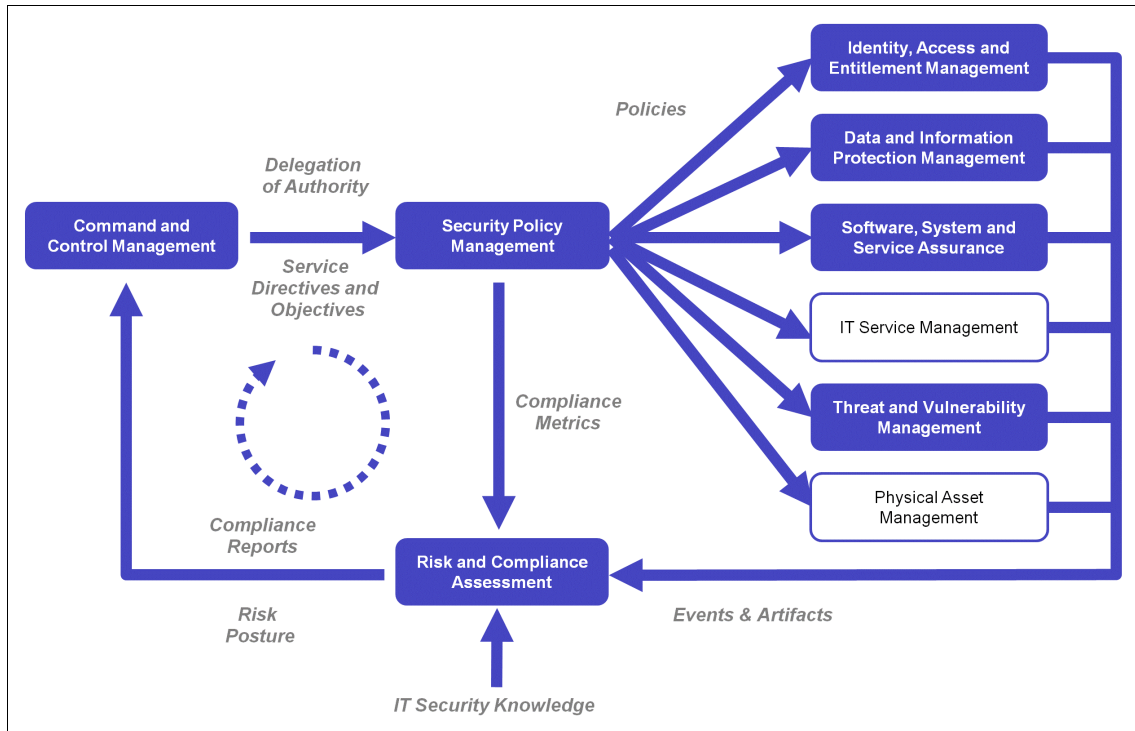


Figure 2-1 Foundational security components closed loop

Also, Security Policy Management delivers the compliance metrics as input to *Risk and Compliance Assessment*, which receives the security events and artifacts that are generated by the more IT delivery-centric security components. Next, Risk and Compliance Assessment combines these events and artifacts to match them with the compliance metrics to produce compliance reports and also to derive a related risk posture, both of which can serve as input into Command and Control Management, so the information can be used for further adjustments to directives and objectives.

Figure 2-1 also shows the security domains of the IBM Security Framework next to the respective matching Foundational Security Management services. The Command and Control Management, Security Policy Management, and Risk and Compliance Management components together reflect the Governance, Risk, and Compliance domain of the IBM Security Framework, the others have a one-to-one matching domain, with the exception of the Application and Process domain, which is matched by the Software, Systems, and Service Assurance component and the IT Service Management component.

In the next section, we provide further details about the Foundational Security Management components by deconstructing them into their subcomponents and listing the related common security infrastructure components.

2.2 Subcomponents

For each of the components on the Foundational Security Management layer, the IBM Security Blueprint provides a further deconstruction into *subcomponents*, as well as an alignment of key Security Services and Infrastructure components, which are essential to a given component in the Foundational Security Management layer. These components are presented in the following order:

- ▶ Command and Control Management
- ▶ Security Policy Management
- ▶ Risk and Compliance Assessment
- ▶ Identity, Access, and Entitlement Management
- ▶ Data and Information Protection Management
- ▶ Software, System and Service Assurance
- ▶ Threat and Vulnerability Management
- ▶ IT Service Management
- ▶ Physical Asset Management

2.2.1 Command and Control Management

The Command and Control Management component provides the command center for security management and the operational security capabilities for non-IT assets and services to ensure protection, response, continuity, and recovery. It covers many topics, such as:

- ▶ Approving authority for security
- ▶ Ensuring that physical and operational security is maintained for locations, assets, humans, environments, and utilities
- ▶ Providing surveillance and monitoring of locations, perimeters, and areas
- ▶ Enforcing entry controls
- ▶ Providing for positioning, tracking, and identification of humans and assets
- ▶ Providing a focal point for continuity and recovery operations

Command and Control Management encompasses situational awareness and reacting to urgent security issues. It also includes the ability to observe and react to long-term trends. In both cases, Command and Control Management includes the ability to trigger and initiate reactive and proactive changes in IT security.

Command and Control Management might utilize other Foundational Security Management services and can serve as the control point for them when knowledge, approval, situational analysis, risk mitigation, and delegation of authority decisions are needed. Figure 2-2 shows an overview of Command and Control Management components and the related components from the Security Services and Infrastructure layer.

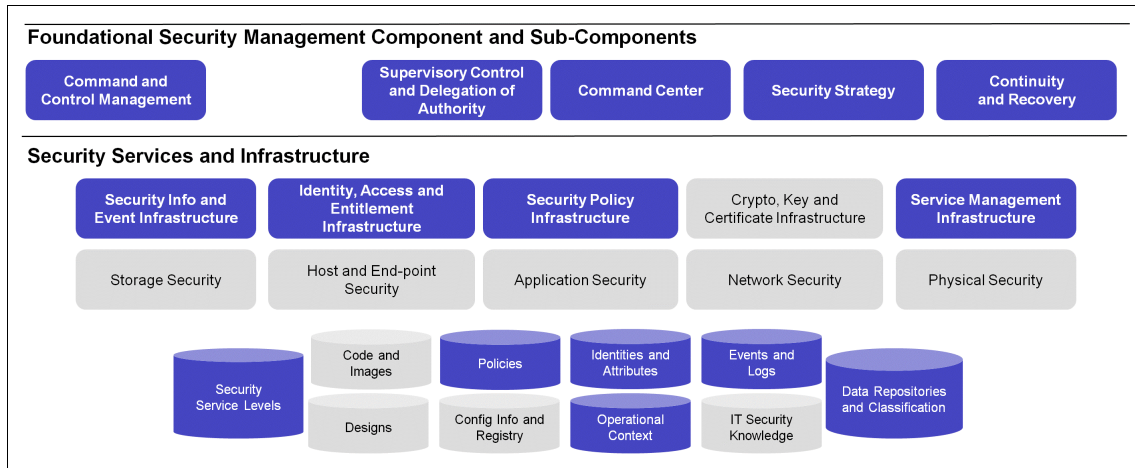


Figure 2-2 Command and Control Management subcomponents

Command and Control Management consists of the following subcomponents:

- ▶ Supervisory Control and Delegation of Authority
- ▶ Command Center
- ▶ Security Strategy
- ▶ Continuity and Recovery

These functional components are described separately to ensure that separation of duties can be achieved.

Supervisory Control and Delegation of Authority

Similar to the concept of Supervisory Control and Data Acquisition¹ (SCADA) systems in physical plants and industrial centers, this component represents the supervisory roles in information security management. This component includes the concepts of delegating authority for IT security to appropriate people and roles in the organization and remotely managing the IT security infrastructure.

¹ To learn more about SCADA, go to <http://www.cpni.gov.uk/ProtectingYourAssets/scada.aspx>.

As part of its supervisory duties, this component owns the responsibility for security as a whole and also for ensuring that policies, standards, and procedures comply with relevant elements of criminal, civil, administrative, and regulatory law to minimize adverse legal consequences.

This component is concerned with making sure that personnel and executives are safe and secure while on site or travelling for the company and knowing to whom authority should be delegated if a person becomes incapacitated or otherwise unavailable.

Command Center

The Command Center represents the service organization unit needed to respond to immediate physical or IT security threats, either through automated responses or scripted scenarios. It also encompasses the development and deployment of crisis management procedures.

The command center is also the focal point for managing communication to external organizations such as Emergency Management Services, fire, police, and other law enforcement agencies.

Security Strategy

The Security Strategy is closely aligned to the overall business strategy and, hence, Command and Control Management is the lead-in for business directives and thus owns responsibility for security strategy management.

Security strategy determines the overall direction of security and security-related compliance, it determines the level of security and protection targets that must be achieved, and it sets the overall boundaries for applicable controls to be deployed to meet the targets.

Continuity and Recovery

Continuity and Recovery represents a service applying a specialized set of skills, processes, and technology to recover from a major unexpected disruption or a disaster in service.

These services include emergency planning activities such as training of employees, escalation procedures, phone lists, procedures, and guidelines for all major types of emergencies, and classification of potential hazards. The services include the coordination of business continuity, that is, keeping the business running during and after a disaster with significant impact on key resources, as well as the coordination of disaster recovery (that is, re-establishing the key resources to a normal operations level).

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective Command and Control Management (depicted as blue-shaded objects in Figure 2-2 on page 35):

- ▶ Identity, Access, and Entitlement Infrastructure

The Identity, Access, and Entitlement Infrastructure is used by all services in 2.2.1, “Command and Control Management” on page 34, to delegate authority by authorizing appointed personnel to receive respective access rights.

- ▶ Security Policy Infrastructure

Security Policy Infrastructure is a key component for 2.2.1, “Command and Control Management” on page 34, as it provides access to the policy documents and also allows the *security policy owners* (the actors behind all four Command and Control Management services) to review and approve policies after confirming that their initially intended Security Service Levels are correctly reflected in the policies. The Security Policy Infrastructure is also used by the services to provide amendment requests to the policies if required.

- ▶ Security Information and Event Infrastructure

The Security Information and Event Infrastructure enables the services of Command and Control Management to retrieve security and event information. This can be valuable when the command center needs to confirm ad hoc occurrences of specific events during crisis management or in discussions with authorities.

- ▶ Service Management Infrastructure

The Service Management Infrastructure is fundamental to Command and Control Management because it relies on the Service Management Infrastructure to coordinate communication to other foundational security services. The personnel associated with the Command and Control Management services are also actors in the Service Management Infrastructure processes. For instance, a change with an impact on security might have to be approved by Supervisory Control and Delegation of Authority if the authority for approval of a specific level of changes (such as a major update to the security architecture of the network perimeter) has not been properly delegated and, hence, is above the clipping level of established delegations.

- Policies

Policies are important to the Command and Controls Management services as they, like all other foundational security services, adhere to policies irrelevant of the fact that the directions that are reflected in the policies had their origin in command and controls management itself. Specific examples for policies include policies around delegation of approval authorities and related clipping levels, in addition to escalation paths.

- Security Service Levels

Security Service Levels are the key output of Command and Control Management and, hence, are the most important data item for the Foundational Security Management services.

- Identities and Attributes

Identities and Attributes define the roles within the organization used in describing policies developed in Command and Control Management.

- Operational Context

Operational Context refers to the existing procedures and policies being followed in the IT organization so that Command and Control Management decisions can be made in a way that minimizes additional burden and disruption to the IT organization.

- Data Repositories and Classifications

Data Repositories and Classifications describe the information assets that are subject to the policies developed by Command and Control Management. Information assets have varying degrees of requirements for protecting confidentiality, availability, and integrity.

- Events and Logs

Events and Logs represent the evidence needed to assess the completeness and correctness of the security controls and to provide information that helps detect fraud and out of process changes to the environment.

2.2.2 Security Policy Management

Security Policy Management provides services and repositories to author, discover, analyze, transform, distribute, and evaluate IT security policies. This component represents a focal point for transforming security requirements needed to mitigate business risks into an IT perspective, which can then be consumed and enforced by the IT infrastructure.

Figure 2-3 shows an overview of Security Policy Management subcomponents and the related components from the Security Services and Infrastructure layer.

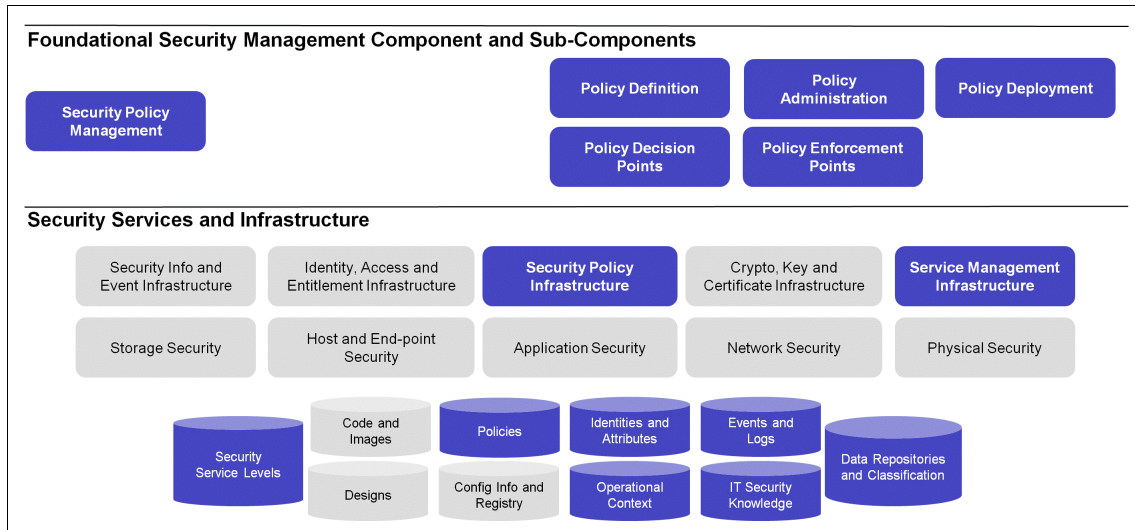


Figure 2-3 Security Policy Management subcomponents

Security Policy Management consists of the following subcomponents:

- ▶ Policy Definition
- ▶ Policy Administration
- ▶ Policy Deployment
- ▶ Policy Decision Points
- ▶ Policy Enforcement Points

These subcomponents are explained in more detail in the following sections.

Policy Definition

The Policy Definition subcomponent represents the ability to represent an IT security policy in human-readable terms, a machine-readable format, or both. It represents the translation of security directives and objectives (as derived by the Command and Control Management from the business security requirements) into actions that can be taken and enforced in the IT landscape. The policies are in scope on all levels and include the top-level security directive, underlying general security policies, deriving more technical policies and platform-specific security standards, in addition to related guidelines and procedures.

The Policy Definition is responsible for capturing the context and background of the IT security policy by tracking the *upstream* policy documents that influence it or rationalize and justify it.

Policy Administration

Policy Administration addresses the human-oriented workflow processes within the policy life cycle management, which includes the create, modify, and maintenance tasks for policies over time. It also addresses the need to manage multiple versions of a policy and transition from one to another over time.

Within the realm of Policy Administration, the policies are approved, announced, published, and commenced as part of Policy Deployment. Also, related activities for this subcomponent include policy education and security awareness.

Policy Deployment

As part of the policy life cycle management, business policies are refined to service-specific policies such as security, performance indicators and metrics, and trust policies. The security policies that result need to be translated and distributed to the technical enforcement and decision points.

For machine-readable policies, the policies are defined centrally and are distributed to the enforcement points in a canonical format (for example XACML, WS-Policy, or WS-SecurityPolicy). The binding information to enforce the policies is also distributed appropriately. These policies are often then transformed at the enforcement point to a local representation so that they can be enforced.

Policy Decision Points

Policy Decision Points (PDPs) represent the capability to evaluate a request and make a decision about whether the request is conformant to a policy. In certain cases, all the information needed to make the decision is contained within the request itself. In other cases, external context information is needed. Sometimes, the sources of context information are called Policy Information Points (PIPs).

There are important issues affecting the placement of PDPs in an IT environment. Centralization of the PDPs reduces administrative burdens and potential errors during deployment. Centralization of the PDPs also enables a PDP to serve multiple enforcement points. However, PDPs are often tightly bound to the Policy Enforcement Points for performance reasons.

Policy Enforcement Points

Policy Enforcement Points (PEPs) take action based on whether the request conforms to policy. The action might be an enforcement action, permitting or denying the request. The PEP might also monitor, log, and raise alerts without affecting the request.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective Security Policy Management (depicted as blue-shaded objects in Figure 2-3 on page 39):

- ▶ Security Policy Infrastructure

The Security Policy Infrastructure is the key component for Security Policy Management, as it provides the containers for the various policies, related standards, procedures, and guidelines. It can automate the workflow for the various administration activities and the deployment and the communication with Policy Decision Points and Policy Enforcement Points.

- ▶ Service Management Infrastructure

The Service Management Infrastructure provides the communication and coordination channels for Security Policy Management to reach all delivery units, which might not belong to security management, but perform some security delivery function and, hence, have to adhere to security policies. Also, this infrastructure component provides the capability to deploy and implement policy updates in line with standardized change and release structures.

- ▶ Security Service Levels

Security Service Levels represent the key input source for Security Policy Management as they set the overall targets that must be decomposed in more detail and then reflected in policy directions and related standards.

- ▶ Policies

Policies represent the key output of Security Policy Management and the most frequented data item for Policy Administration, Policy Decision Points, and Policy Enforcement Points.

- ▶ Operational Context

Operational Context is important for the policy definition service in Security Policy Management. The policies set for an environment should be achievable to a large extent, so it is important to establish the targeted controls documented in the policies with consideration of their achievability and their appropriateness. The Operational Context provides essential input for related evaluations of controls. This also helps to avoid the situation in which a policy would have to be accompanied with many policy exceptions to stay in control of the deviations of the deployed operational environment from the intended (and practically unachievable) state set out in the policies.

An unnecessary number of exceptions also requires a lot of avoidable administrative effort and leads to inefficient Security Policy Management, so evaluating the Operational Context thoroughly during the design of the policies helps to establish adequate policies and avoid situation in which policies take the form of a pure theoretical documentation.

The Operational Context is also important in Policy Administration and in Policy Deployment. Even when care is given to establish appropriate, practical, achievable control requirements in the policies, exceptions in an operational environment are unavoidable. Such exceptions can derive, for instance, from the lack of support of a given control by a particular system. While compliance can be achieved in such a case, an exception is documented to capture the particular deviation from the policy and the refined requirement of compensating controls for the particular deviation. To perform these actions, the Operational Context must be examined.

- Data Repositories and Classifications

Equally as important as the Operational Context, Security Policy Management services depend on reviewing and understanding the Data Repositories and Classifications. The Policy Definition service requires the structuring of the data repositories and identifying the confidentiality, integrity, and availability requirements of data repositories. Based on this, a sufficient yet manageable set of classifications of information assets has to be defined, and for each of the classifications the related security control requirements must be set in the designed policies. As explained in the Operational Context bullet above, the Data Repositories and Classifications are also examined as part of Policy Administration for the evaluation of exceptions from policy-mandated controls usually required for a given data repository in cases in which such controls cannot be maintained for technical or business reasons.

- Identities and Attributes

Besides the Operational Context, Data Repositories and Classifications, Identities and Attributes represent another data item that has to be fully understood to define appropriate policies for a given environment. While the Operational Context helps to evaluate the environment from a business and technical infrastructure perspective, and Data Repositories and Classifications help to understand it from a data and information perspective, Identities and Attributes provides the perspective onto an environment with a focus on users, administrators, and other actioners in the environment.

Like with the two aforementioned components, Identities and Attributes are taken as input to the activities of security control design in the Policy Definition service and also are used to set security requirements for these identities and the related attributes.

Also, the ongoing Policy Administration service will use Identities and Attributes and its evolution throughout operations to adapt policies with new or amended requirements to address identified operational security issues and to perform continued security improvement.

► Events and Logs

Events and Logs are important because they allow verification of completion of Security Policy Management activities, which have been performed with the help of the Security Policy Infrastructure. From this perspective, the Events and Logs serve as evidential records about activities (for instance, whether a specific control has received review and approval from stakeholders as part of Policy Administration activities before it is published in an updated security policy). But also from a perspective of Policy Definition, Events and Logs can be a helpful source of information. Alongside the traditional qualitative analysis of Operational Context, Data Repositories and Classifications, and Identities and Attributes when defining appropriate controls in the security policies, event and log data from the environment can be used to identify actions and behavior that happen in that environment. This allows for prioritization, especially in cases in which an environment is already in operation, but to a certain extent the security policy definition lags behind. When following security management approaches by the book, such situations should not exist (that is, no environment should go into operation without first defining adequate security policies, but in reality this is not always the case). Deriving the actions that caused specific events and log records (or combinations thereof) and examining the security requirements for these actions can be helpful, especially in situations in which information about the environment is not available or fully understood, deriving security-critical actions from Events and Logs help to find a start to fix situations in which an environment lacks security policies.

► IT Security Knowledge

IT Security Knowledge for Security Policy Management services includes general knowledge about how to create and maintain effective security policies, and also requires technical understanding of the security controls provided for various platforms, in case specific security standards for these technical platforms have to be established to provide clearer direction towards the implementation of respective policies. Also, general knowledge about well-established industry regulations and standards as well as about data privacy regulations in the various legal contexts in which a organization operates is required to better translate related directives and objectives coming from the business via Command and Control Management into the respective policies.

2.2.3 Risk and Compliance Assessment

Risk and Compliance Assessment enables the IT organization to collect, analyze, and report security information and security events to identify, quantify, assess, and report on IT-related risks that might contribute to an organization's operational risk. This component covers risk aggregation and reporting, IT security risk processes, business controls management, resiliency and continuity management, compliance reporting, and legal discovery services.

Figure 2-4 shows an overview of the Risk and Compliance Assessment subcomponents and the related components from the Security Services and Infrastructure layer.

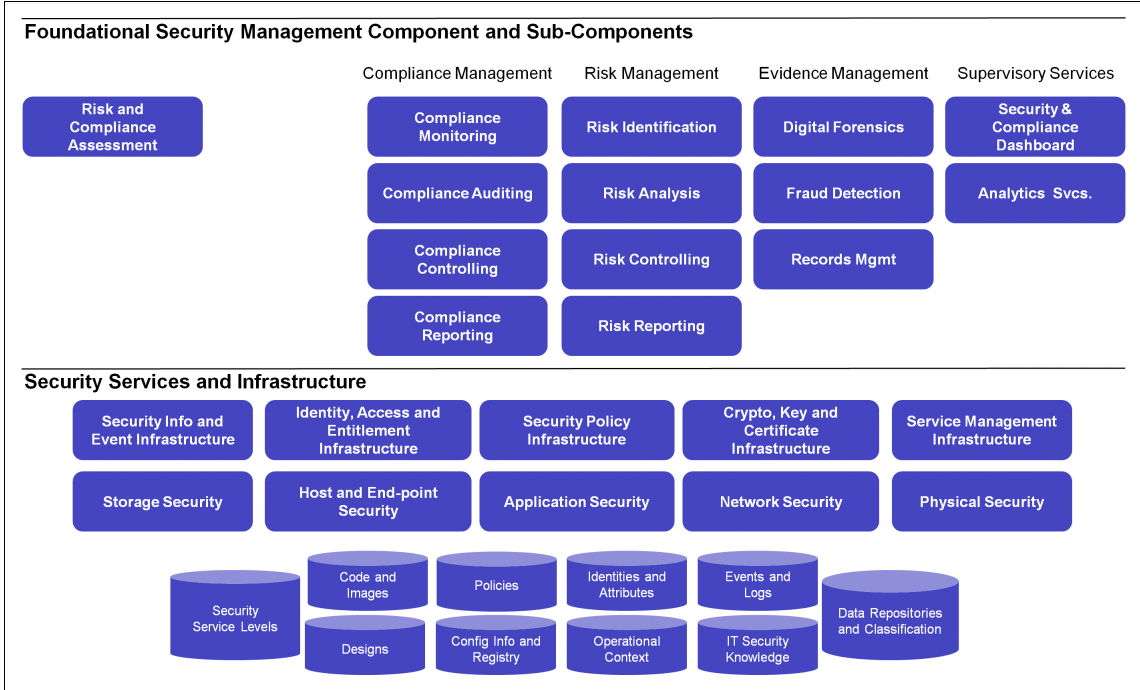


Figure 2-4 Risk and Compliance Assessment subcomponents

Risk and Compliance Assessment consists of the following subcomponents:

- Compliance Management
- Risk Management
- Evidence Management
- Supervisory Services

These four services are discussed in more detail in the following sections.

Compliance Management

Compliance Management covers all activities related to overlooking and driving the security compliance state of the IT environment.

Compliance Monitoring

Compliance Monitoring refers to observation of the environment to identify gaps between the actual operations, the internal policies and standards, and the requirements as they derive from external industry regulations, laws, and orders.

Compliance Auditing

Compliance Auditing refers to the ability to match event sources and their event streams to compliance reporting requirements for IT security and produce reports based on those event streams, either periodically or on demand as part of an audit. Managing the association between the event sources reports and the compliance reporting requirement is a key capability of this component. Also, compliance requirements often impose record retention requirements on audit data, which might be different than the retention requirements for the event streams in the IT environment in general. From an IT operations perspective, the event streams are more short lived, while data that supports compliance audits might have a life span of multiple years.

Compliance Controlling

Compliance Controlling stands for the continuous work that is contributed by IT security compliance experts throughout the various parts of an organization, focusing mostly on two key activities:

- ▶ Compliance support
- ▶ Compliance tracking

Compliance support refers to providing advice and guidance to those who are not necessarily compliance experts, but whose activities are subject to compliance. For example, compliance experts work with a business unit to help them prepare for an upcoming audit or to help during an audit. Similar to an attorney of law in court, a compliance expert can help an audited business unit with the preparation of paperwork requested by the auditors or in the preparation of audit interview partners for their meeting with the auditors.

The other aspect of Compliance Controlling is *compliance tracking*, which covers the structured documentation of follow-up activities after an audit and the progress of these activities until closure. The activities are either determined by the auditor directly or are derived by an analysis of audit results as those actions, which have to be implemented to mitigate identified compliance and security issues.

Compliance Controlling is a continuous process (*before, during, and after the audit*) and, hence, requires substantial ongoing efforts of a well-functioning compliance regime in an organization.

Compliance Reporting

Compliance Reporting refers to the ability to summarize analyzed event data and other security-relevant information for the specific use of demonstrating compliance. Most often, reporting is used to assess regulatory compliance or compliance with security service level agreements and overall compliance performance of the IT environment. From an internal security perspective, Compliance Reporting is most commonly used to demonstrate control over security policies and to identify trends in security compliance.

Risk Management

Risk Management covers all activities related to overlooking and driving the security risk posture of the IT environment.

Risk Identification

Risk Identification refers to the ability to discover, recognize, and verify the existence of specific risks. It also encompasses the structuring of risk by mapping it into clearly defined classification schemes that can be specific to the industry or even to the risk taxonomy of an individual organization.

Risk Analysis

Risk Analysis refers to activities related to the categorization, qualification, or quantification of the likelihood and impact of risks. It also covers the investigation of connections, dependencies, and correlations among various risks.

Risk Controlling

Risk Controlling covers the determination of activities that can be used to address given risks. The valid activities can range from *risk acceptance* over different approaches of *risk mitigation* to *risk transfer*. Risk Controlling also includes the determination of costs for such activities and the identification of potential risk and risk mitigation owners and actors. Another important part of Risk Controlling is tracking the status of identified and agreed risk mitigation activities until their closure.

Risk Reporting

Similar to Compliance Reporting, Risk Reporting refers to the ability to summarize analyzed risk data and other risk-relevant information and to provide different levels of detail about the security risk posture to different parts of the organization as input for further analysis and processing.

To a certain degree, Risk Reporting is also used as input into Compliance Reporting, because certain regulations might require that an organization provides information about key risk events to its stakeholders (for example, banks have to inform regulatory authorities about their operational risk, which also includes their security risk posture). From an internal security perspective, Risk Reporting is most commonly used to help make the correct decisions for investments in risk mitigation activities and to track the progress for these activities.

Evidence Management

Evidence Management covers services that are related to capturing and securing information in a form that can be used as legal evidence in court or that has to be preserved for other legal reasons.

Digital Forensics

Digital Forensics refers to the ability to retrieve and preserve the state of IT components that are subject to a legal investigation. In certain cases, forensics simply involves preserving the state of a system for future reference. In other cases, forensics requires the recreation of events that lead to the state of a particular component. For example, email is often subject to e-discovery requests in legal proceedings, and many organizations must be able to enforce *deletion holds* on email to prevent their destruction when subject to discovery proceedings.

As another example, a time line of configuration changes for a database might need to be recreated to identify why it failed and who authorized the changes that caused the failure.

Forensics investigations can be initiated internally or as part of a legal proceeding. When forensics investigations are initiated as part of legal proceedings, additional security issues can come into play, such as completeness and accuracy of the collected data, and chain of custody issues. The chain of custody issues cover situations in which data is transferred from one IT component to another or from one individual to another.

Fraud Detection

Fraud Detection covers the analysis of information and events within the IT environment relating to unsolicited business-level activity. Usually, Fraud Detection addresses the review of security information and events for a specific combination of occurrences, which not only indicate the abuse of user rights or bypassing of access controls in a pure policy context, but are targeted to perform fraudulent activities in a criminal and legal context.

Records Management

Records Management refers to the industry term that addresses the legal requirement to capture and keep specific records about business transactions and communications for potential submission as incriminating or discharging evidence.

Supervisory Services

Supervisory Services in Risk and Compliance Management provide monitoring, alerting, and analysis across all areas of compliance, risk, and evidence management.

Security and Compliance Dashboard

Like other business-related dashboards, the Security and Compliance Dashboard refers to a set of web interfaces to display the most current relevant reporting information for IT security events and the status and completeness of compliance efforts. Dashboards are based on event streams that have been collected over a period of time.

Analytics Services

Analytics Services help to find trends in correlated events and to make decisions based on the trends found. For example, an event analytics engine might match authorization events against human resources employee records to detect the use of orphaned accounts for people who have left the company, possibly indicating an attack from an ex-employee, or the use of a shared ID that is disallowed by corporate security policy.

In another example, a business activity monitoring control might require that each invoice be paired with an authorized purchase order. There might be multiple channels for purchase orders to come into the order system, each with a business event monitor sending purchase order events to the event correlator. Likewise, there might be two channels for invoices to be entered into the order system, each monitored by a business control that sends invoice events to the correlator. The event correlator might group these events into pairs based on the purchase order number, but emit a higher level invalid invoice event if it holds an invoice for more than 24 hours without receiving a corresponding purchase order event. The analytics engine can look for common patterns in the invalid invoice events and raise alerts to the appropriate departments or business control personnel.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key for an effective Risk and Compliance Assessment (depicted as blue-shaded objects in Figure 2-4 on page 44):

- Security Information and Event Infrastructure

The Security Information and Event Infrastructure is an important element for the Risk and Compliance Assessment, because it can help collect and provide information about events in a synthesized, consolidated, platform-independent, and less technical format. The aggregation of security logs and subsequent derivation of security information, which now is understandable by less technical people on the business level, is provided to the Risk and Compliance Assessment services to further analyze data in a risk context (that is, in terms of probability and business impact). In the context of Compliance Management, the Security Information and Event Infrastructure can help produce reports that are specifically designed for compliance to particular regulation and legal requirements. This infrastructure component also provides a substantial part of the evidence that has to be gathered and analyzed by the Evidence Management services.

Besides providing services *for* the Risk and Compliance Assessment, the Security Information and Event Infrastructure itself is also *subject to* Risk and Compliance Assessment.

- Identity, Access and Entitlement Infrastructure

The Identity, Access and Entitlement infrastructure is used by the Risk and Compliance Assessment services to analyze risk and compliance posture pertaining to insufficient separation of duty. Also, this infrastructure is used by the Risk and Compliance Assessment services to identify, verify, and further investigate activities of events resulting from malicious user behavior.

Besides providing services *for* the Risk and Compliance Assessment, the Identity, Access, and Entitlement Infrastructure itself is also *subject to* Risk and Compliance Assessment.

- Security Policy Infrastructure

The Security Policy Infrastructure provides structured access to the security policies and standards of an IT organization. Ideally, this infrastructure serves as the sole instance for compliance requirements and, thus, provides compliance-related information in an *end-to-end* fashion. End-to-end in this context implicates that the Security Policy Infrastructure must cover all possible applications and platforms and provide proper cross-referencing between the various compliance-related documents and, ideally, the individual requirements in these documents.

The Security Policy Infrastructure should follow the usual pyramid structure of a compliance documentation framework, with the top-level security policy and more detailed security policies and corresponding technical security standards underneath.

As policies and standards develop and change over time, the Security Policy Infrastructure is not only able to provide a snapshot of the policy framework at a given point in time, but it supports the evolution of policies and standards. It allows the recording of the state of approval for a given policy at a given point in time and provides convenient ways to examine differences between various compliance requirements. This can help identify and resolve potential contradictions between policies to prevent misunderstandings about the direction or the intent of a compliance requirement.

Finally, the Security Policy Infrastructure helps you to check whether the policy workflow for defining and establishing security policies and standards has been properly followed. From this perspective, the Security Policy Infrastructure itself must comply with requirements of the policies and standards that it holds and, thus, it is also subject to audits and reviews.

A policy infrastructure defined in this way can serve as a single consolidated reference of the intended state of compliance for any organization. It can be the key to an efficient security and compliance management implementation.

- **Cryptography, Key, and Certificate Infrastructure**

The Cryptography, Key, and Certificate Infrastructure provides the capability to perform cryptographic operations. As such, it is not directly used by the Risk and Compliance Assessment services. However, many organizations that utilize the Cryptography, Key, and Certificate Infrastructure have to abide by rigid laws and regulations on encryption key lengths and methods. That is why the Cryptography, Key, and Certificate Infrastructure is an area that needs to be thoroughly assessed by the Compliance Management pillar.

- **Service Management Infrastructure**

Risk and Compliance Assessment services operate under an agreed-upon Service Management Infrastructure and must utilize the services provided by that infrastructure. For example, accessing and transferring evidence from audited machines must be performed in line with the change management process (for instance, they must utilize proper change management ticketing and approval, and thus use the Service Management Infrastructure). Equally important, Evidence Management activities have to be performed in line with incident and problem management processes and utilize the related parts of the Service Management Infrastructure (for instance, mechanisms provided for incident and problem logs).

► Storage Security

Storage Security is tied to Risk and Compliance Assessment both from a direct and from an indirect perspective.

From a direct perspective, Storage Security is a target of many Risk and Compliance Assessment services. This means that Storage Security is assessed and examined by these services.

From an indirect perspective, Storage Security is heavily utilized by all five aforementioned management infrastructures (that is, Security Policy Management Infrastructure, Event and Log Management Infrastructure, Cryptographic Key Management Infrastructure, Identity and Access Management Infrastructure, and Service Management Infrastructure) required by the Risk and Compliance Assessment services. Risk Management, Compliance Management, and Evidence Management have high requirements for the integrity on stored data.

► Host and Endpoint Security

Host and Endpoint Security provides an indirect service to the Risk and Compliance Assessment component via the five security management infrastructures because all the infrastructure components run on actual hosts and use endpoints. Host and Endpoint Security is important for these management infrastructures to function. Several of those important services include agents and collectors that have to be distributed to the hosts and endpoints for the security management and aggregation layer infrastructure to be able to serve their purpose.

From a direct perspective, Host and Endpoint Security is a key examination point for Risk, Compliance, and Evidence Management services. Besides managing security aspects for physical systems, Host and Endpoint Security includes security configuration details for operating systems, middleware, software packages that provide a distinct security function, like antivirus software, personal firewalls, host intrusion detection and prevention systems, and hard disk, file or mail encryption software.

► Application Security

Application Security provides many events and logs that have to be analyzed for risk and compliance. Because it is the closest, most used interface to the business user, it is important that it be examined for Compliance Management and Evidence Management, especially for fraudulent activities.

► Network Security

Like many of the other technical platforms, network components and traffic provides a wide range of traces of events and general activities, which are considered important factors for all Risk and Compliance Assessment services.

- Physical Security

Physical Security, like the security of any of the technical platforms, can consist of a wide range of security controls that have to be functional to fulfill compliance requirements in mitigate risks and retain evidence.

Physical Security is essential because information that must be protected does not only exist in electronic forms, but also in traditional non-electronic forms. Good security practices require the management of the risks, compliance, and related evidence in the physical domain as well.

- Security Service Levels

Because the Security Service Levels provide the background for the policies, Risk and Compliance Assessment services can use them to understand and resolve potential different interpretations and ambiguities in the security controls as well as the security control objectives defined in the policies and, hence, in the measurement of compliance.

Also, Security Service Levels can be examined by Risk and Compliance Assessment services to evaluate whether they can cause risks by themselves and need to be adjusted.

- Code and Images

Code and Images are used to identify potential sources of risk and of non-compliance. Those risk and non-compliance issues might only surface on systems that are in production, but the issues have their origin in flaws in the source code and the base images. Comprehensive security policies typically define requirements and controls onto the source code and image composition themselves, so that Risk and Compliance Assessment services have to assess them before they are being put into production.

- Designs

Designs are important to Risk and Compliance Assessment services because they are used as (often graphical) representation systems, users, and processes and their relationships. Such representation must reflect the respective security policies, standards, and directives. Risk and Compliance Assessment services assess the designs and architectures for risks and for compliance within policy and regulatory requirements and also use them as reference for an intended state of something that has been implemented. In other words, Risk and Compliance Assessment services must verify whether the designs are in line with security and compliance requirements, and then again whether the implemented environment is in line with this verified design.

► Policies

One of the primary inputs into Risk and Compliance Assessment are the Policies. They define the compliance metrics that are used to identify non-compliance for many systems and services. Compliance Management assesses compliance of the IT environment by identifying and examining differences between actual and intended compliance values defined in the compliance metrics. Risk Management assesses the compliance metrics and the target values for the adequacy for mitigating related risks to the level set by Command and Control Management as acceptable.

► Configuration Information and Registry

The Configuration Information and Registry contains settings that have to be implemented to meet security controls defined in the policies. Compliance Management uses this information to verify that the security controls are properly implemented and, thus, compliance requirements are met. Risk Management assesses the configuration for the technical appropriateness of the settings to verify that risk mitigation targets are met. Evidence Management assesses the Configuration Information and Registry for any suspicious unauthorized changes that might allow fraudulent activities. Evidence Management also collects evidence about the security state of the IT environment to the extent as this is required from a legal perspective.

► Identities and Attributes

Directories contain important information about people's identities along with other key attributes, which is used to control access to data and other resources. Hence, major efforts within the Risk and Compliance Assessment services are focused on checking the compliance posture and the risks deriving from errors in Identities and Attributes. While the Identity, Access and Entitlement Management infrastructure is assessed from the perspective of procedural compliance, the Identities and Attributes are assessed from a perspective of factual or conclusive compliance. Both in combination can also reveal whether the Identity, Access and Entitlement Management services function as designed or whether they have been bypassed to make changes on Identities and Attributes. The Evidence Management services require access to Identities and Attributes to gather evidence about identities that have been used to perform potentially malicious behavior.

Besides assessing the technical compliance and related risks in the area of identity and access management, Risk and Compliance Assessment services assess identity and attributes information also from a more organizational perspective. In other words, the risk and compliance experts must not only check and verify whether technical settings for access administration are correct, but also whether the entitlements of a given user for a given resource are appropriate from a compliance and risk perspective.

For example, information access should not be granted to a user with specific *identity features* like nationality, security clearance, or location of that user. In another example, the information might be classified so that it cannot be changed by one user alone (four-eye-principle).

- ▶ Operational Context

The Operational Context can influence whether a given activity and, hence, related events are compliant or non-compliant. That is why the Operational Context has to be reviewed by Risk and Compliance Assessment services to come to correct conclusions towards compliance and evidential material and towards risk.

An example for such influence can be the execution of privileged activities with an unrestricted account. While an administrator might be granted—from a technical perspective—unrestricted access to a system, this administrator should only use a limited subset of commands to perform a given change. Hence, the execution of other privileged commands not related to this particular change, although still being perfectly OK for a different task, could be discovered by checking the Operational Context.

- ▶ IT Security Knowledge

Defining appropriate risk categories and applying security risk thinking requires specific experience and IT Security Knowledge. IT Security Knowledge for Risk and Compliance Assessment also includes detailed understanding of compliance and regulatory standards.

- ▶ Data Repositories and Classification

Data repositories are increasingly incorporating access control mechanisms to create an access control point as close to the data as possible. The classifications of the data in repositories must be compliant and set in a way that can possibly reduce risk to an acceptable level. Also, as evidential data has to be stored in data repositories (even if the repositories are taken offline), the access and the classification of this data is paramount to keeping them admissible for any legal activities.

2.2.4 Identity, Access, and Entitlement Management

Identity, Access, and Entitlement Management provides services related to roles and identities, access rights, and entitlements. The proper use of these services can ensure that access to resources has been given to the correct identities, at the correct time, and for the correct purpose. These services can also ensure that access to resources is monitored and audited for unauthorized or unacceptable use.

Figure 2-5 shows an overview of Risk and Compliance Assessment subcomponents and the related components from the Security Services and Infrastructure layer.

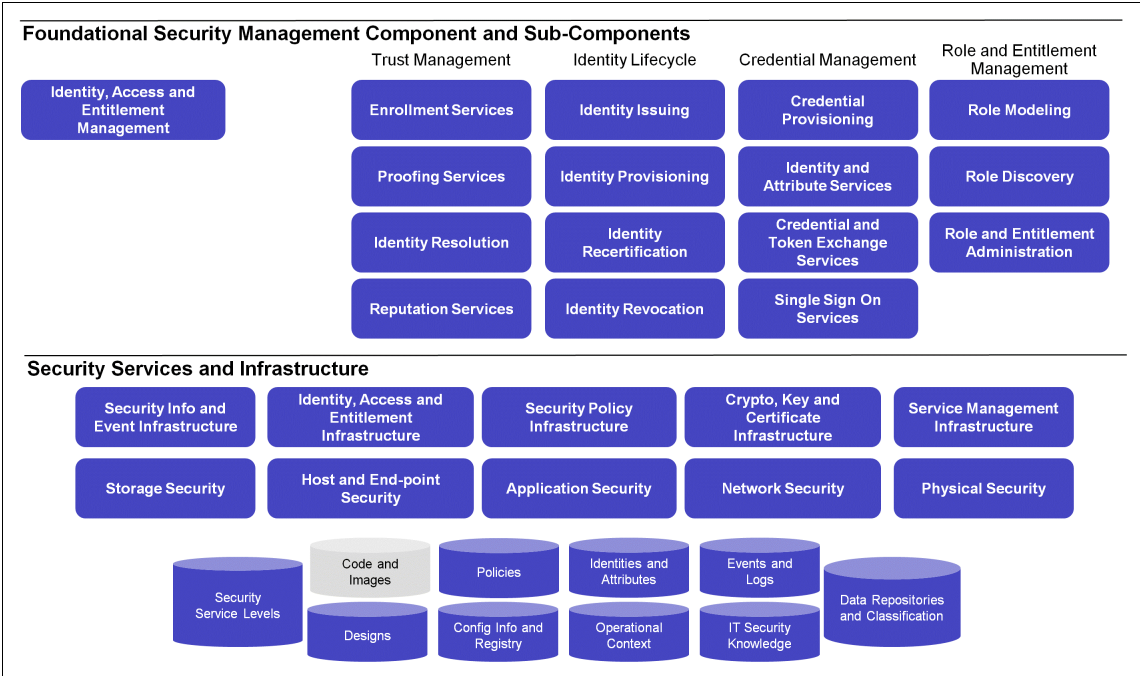


Figure 2-5 Identity, Access, and Entitlement Management subcomponents

Identity, Access, and Entitlement Management consists of the following subcomponents:

- ▶ Trust Management
- ▶ Identity Life cycle
- ▶ Credential Management
- ▶ Role and Entitlement Management

These services are explained in the next sections.

Trust Management

Trust Management refers to the activities needed to improve the reliability of identity management systems to ensure that credentials are issued to the correct people.

Enrollment Services

Enrollment Services cover the act of collecting initial documentation from the person who wants to be issued a credential, including things like birth certificates and other source documents. It might also involve collecting biometric and biographic information.

Proofing Services

Proofing Services are the processes and technology for verifying all the information collected from the individual with the enrollment services. In addition to verifying information against authoritative sources, it might also include using identity analytics to detect fraudulent applications.

Identity Resolution Services

Identity Resolution Services cover the processes and techniques to identify multiple records for the same person, whether by accident or fraud, and to resolve them into a single record for a single person.

Reputation Services

Reputation Services involves tracking an individual's actions over time, collecting data about the opinions others have of those actions either from other individuals or rating systems, and publishing an assessment of the opinions either publicly or to the subject individual as a feedback mechanism.

Identity Life cycle

The Identity Life cycle spans from the initial creation over specific events during the life of an identity through to the final deletion of an identity. The key elements of the Identity Life cycle are explained in the following sections.

Identity Provisioning

Identity Provisioning covers the processes and technology used to create the credential that will be used when issuing an identity token (for example, national ID card) and registering the credential to systems that need to authenticate the credential.

Identity Issuing

Identity Issuing covers the processes and technology used to create the physical components of the credential and securely deliver them to the owning individual.

Identity Recertification

Identity Recertification refers to the processes and technology used to re-validate a credential that has already been issued. In certain cases, this means updating the credential itself.

For example, a digital certificate has expired and another one has to be issued. In other cases, the recertification involves re-authorization and presenting proofing materials again.

Identity Revocation

Identity Revocation covers the processes and technologies used to de-certify a credential so that it can no longer be used as an identity token. This can happen through normal expiration processes or be initiated by an outside trigger event. For example, a revoked digital certificate might be published on a certificate revocation list, which is checked by the identity infrastructure.

Credential Management

Credential Management deals with the administration of credential information and related identity information. Besides the handling of credentials in electronic format, credential management also includes the administration of physical credentials, like tokens or badges.

Credential Provisioning

Credential Provisioning covers the activation of the issued credential so that it can be used to validate an individual's identity, in addition to services for updating, deleting, and managing trusted identity credentials through the entire life cycle.

Identity and Attribute Services

Identity and Attribute Services manage access to local user registries and databases that provide identity information. Typically, identity and attribute services are able to add and delete identity information in addition to reading it.

Identity and attribute services are used by authentication services when evaluating user-presented authentication credentials and to build privilege credentials used by session management services. The privileges are typically based on attributes of a user stored in the Identities and Attributes security service, such as group membership, roles, personal attributes, and so on.

Identities and Attribute services that also manage the attributes about a user are sometimes referred to as identity and attribute services (IdAS).

Credential and Token Exchange Services

Credential and Token Exchange Services combine token validation and issuance to convert one type of security token into another. Security tokens are validated in terms of signatures on the token, expected structure, and contents of the token. Token issuance involves creating a new, locally valid token based on the received, validated token.

When this new token is returned to the original requestor, the process is referred to as a token exchange. The requestor is in effect exchanging the token that it received on a request for a new token that is locally valid.

Single Sign-on Services

Single Sign-on Services implement a set of protocols designed to remove the burden of repeating actions placed on the requestor. Typically, an identity provider can act as a proxy on a requestor's behalf to provide evidence of authentication events to third parties requesting information about the requestor.

These identity providers (IPs) are trusted third parties and need to be trusted by both the person who originates the original request and the online service that allows the requestor to engage in sensitive or high-value transactions.

Role and Entitlement Management

Role and Entitlement Management embraces all functional services that relate to the grouping of identities and to the administration of access to information and resources at a group rather than an individual level.

Role Modeling

Role Modeling deals with the design of role structures to address requirements as they derive from the business and IT activities. The goal of role modeling is to reduce complexity of actors by grouping them, which can result in the capability to provide and to restrict access to information and other resources more efficiently and is less error prone when enforcing a separation of duties.

Role Discovery

Role Discovery refers to the identification of roles and their respective entitlement. The necessary information about roles and entitlements can be gathered either manually by observation and analysis of processes and interviews of the process actioner or process owners. Information can also be captured from systems supporting the processes. User activity, to a certain extent, can be automatically analyzed and structured to derive roles and entitlement patterns.

Role and Entitlement Administration

Role and Entitlement Administration deals with the activities around maintaining and updating the role and entitlement structures. It is similar to the management of identities.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective Identity, Access and Entitlement Management (depicted as blue-shaded objects in Figure 2-5 on page 55):

- ▶ Security Info and Event Infrastructure

Records of access attempts and whether they were granted is one of the most important records of activity for audit purposes, especially access records for privileged users. Policy enforcement points are responsible for generating appropriate audit records for these activities. These records are typically collected by a Security Information and Event Infrastructure for long-term tamper-proof storage, normalization, correlation with other events, and to provide appropriate evidence during audits.

- ▶ Identity, Access and Entitlement Infrastructure

The Identity, Access and Entitlement Infrastructure represents the Policy Decision Points and Policy Enforcement Points that make authorization decisions and enforce them during run time.

The Identity, Access and Entitlement Infrastructure includes access control points to prevent unauthorized access to data, applications, and other IT resources both from a business operations perspective and from an IT administration perspective. These control points are driven by policies and entitlements defined in the Identity, Access and Entitlement Management component in the Foundational Security Management layer.

The access control points rely on authentication mechanisms in the infrastructure and an identity management provisioning infrastructure that manages the accounts, passwords, public key certificates, and other materials needed for authentication.

The access control points are also the focal point for monitoring and enforcing segregation of duty policies as defined by the Identity, Access and Entitlement Management component in the Foundational Security Management layer.

- ▶ Security Policy Infrastructure

The Security Policy Infrastructure is responsible for taking a common access control policy defined in the Security Policy Management system, transforming it into a format that the Policy Decision Point can interpret, and securely delivering it to the Policy Decision Point.

- ▶ Cryptography, Key and Certificate Infrastructure

In many cases, the credentials used in an authentication request have been signed or encrypted so that a Policy Decision Point can properly validate the credentials. The Cryptography, Key and Certificate Infrastructure provides the ability to perform cryptographic operations and signature validation and creation as needed to process authentication requests.

- ▶ **Service Management Infrastructure**

Identity management processes in an organization help manage the entitlements to applications and data, typically using organizational roles as a basis for deciding who is entitled to which resources. The entitlements must be translated into specific credentials on target systems in the runtime environment so that the Policy Enforcement Points know which credentials to grant access and which to deny. The Service Management Infrastructure is responsible for interacting with the user repositories on target systems and creating and modifying the accounts on those systems so that the owner is granted the appropriate access based on his entitlements.

- ▶ **Storage Security**

The Storage Security infrastructure is responsible for protecting storage media from out-of-band attacks, such as theft of media, unauthorized duplication of media, or interception of traffic to and from the storage system. Storage Security relies on Identity, Access and Entitlement Management to define and manage the administrators and the runtime systems that have access to the storage system.

- ▶ **Host and End-point Security**

Host and End-point Security is tightly integrated with Identity, Access and Entitlement Management. End-point machines, by their nature, are often the initial point of contact with a user and are the first point that a user has the opportunity to authenticate to the IT environment. As a result, credentials established by the end-point often need to be propagated to back-end systems or translated into equivalent credentials used in back-end systems. Likewise, the end-points become a key component for single sign-on services.

- ▶ **Application Security**

As part of their design, applications typically use a set of application-specific roles. These roles define who can interact with an application, and in what way, to access the various services that the application provides. The application platform is typically responsible for defining associations between the application-specific roles and the organizational roles managed by the Identity, Access and Entitlement Management system. These associations are then translated into access control policies that the application platform uses to grant or deny access to the application at run time.

- ▶ **Network Security**

Granting and denying access to the network is a key component of Network Security. Network Security depends on the access, identity, and entitlement management system to manage who is granted access to which parts of the network and to generate the necessary credentials and access control policies for the Network Security infrastructure to use at run time.

- Physical Security

Physical Security increasingly relies on logical access security to protect physical access. The most common examples include access control systems on doors, such as password keypads, biometric scanners, or badge readers. In many cases, these access control systems require that access be granted on a per-person basis. In these cases, the Physical Security systems rely on the Identity, Access and Entitlement Management system to manage the identities and entitlements (who can access which parts of the physical facility) in an organization.

- Security Service Levels

The security service level agreements set objectives for managing access to key applications, data, networks, and physical facilities, in addition to the reporting and auditing requirements to demonstrate that the access controls are deployed and effective. Security service level agreements might also include provisions for various types of penetration tests of the access controls and performance metrics for the access control systems.

- Designs

Most IT-related designs in an organization define access control policies for the elements that they represent. These policies must be incorporated into the Security Policy Management system and represented as access control policies that the Identity, Access and Entitlement Infrastructure services must be able to enforce.

IT designs and other business-oriented domain designs are often considered to be high-value assets and are subject to access controls and auditing. So the document management systems used to store these designs rely on the Identity, Access and Entitlement Management system to define the policies about who can access which designs.

- Policies

One of the primary inputs into an Identity, Access and Entitlement Management system are policies, which define the organizational roles and their entitlements to applications, data, networks, and physical facilities. The Security Policy Management infrastructure is responsible for the authoring processes for these policies and the Identity, Access and Entitlement Management system is responsible for translating access control policies into machine-interpretable formats that can be understood by the Policy Decision Points and Policy Enforcement Points.

- Configuration Info and Registry

Because the configuration management databases and registries for IT resources represent valuable knowledge that can be used in an attack, access to those resources is typically tightly controlled and made available only to a few privileged users.

- Identities and Attributes

The directories that contains information about people in an organization and key attributes for them represent the primary data component for an Identity, Access and Entitlement Management system. These directories are typically tightly integrated with human resources systems or are synchronized with them so that they always reflect the current organizational structure for the enterprise. The Identity, Access and Entitlement Management system relies on the directories when mapping organizational roles to application roles and other sorts of entitlements.

- Operational Context

Access control policies increasingly depend on information that is not available at run time. For example, an access control policy can grant access to a resource only if the requester has been assigned to a unit of work that the resource is associated with. Or it might grant access to a resource only during certain times of day or from certain locations. The Identity, Access and Entitlement Management system must be aware of the Operational Context at run time to author policies that incorporate this runtime context.

- IT Security Knowledge

Defining appropriate access control policies to implement an organization policy requires a working knowledge of access control principles, such as granting of least privilege and how to combine entitlements when a person fulfills multiple roles in an organization. IT Security Knowledge is also important to choose appropriate Policy Decision Points and Policy Enforcement Points in an IT environment.

- Data Repositories and Classification

Data repositories are increasingly incorporating access control mechanisms to create an access control point as close to the data as possible. The classifications of the data in repositories must be available to the Identity, Access and Entitlement Management system to define access control policies that are appropriate for the data classification.

2.2.5 Data and Information Protection Management

Data and Information Protection Management provides services that protect unstructured and structured data from unauthorized access and data loss, according to the nature and business value of information. It also provides usage and access monitoring and audit services.

Figure 2-6 shows an overview of data and information protection management subcomponents and the related components from the Security Services and Infrastructure layer.

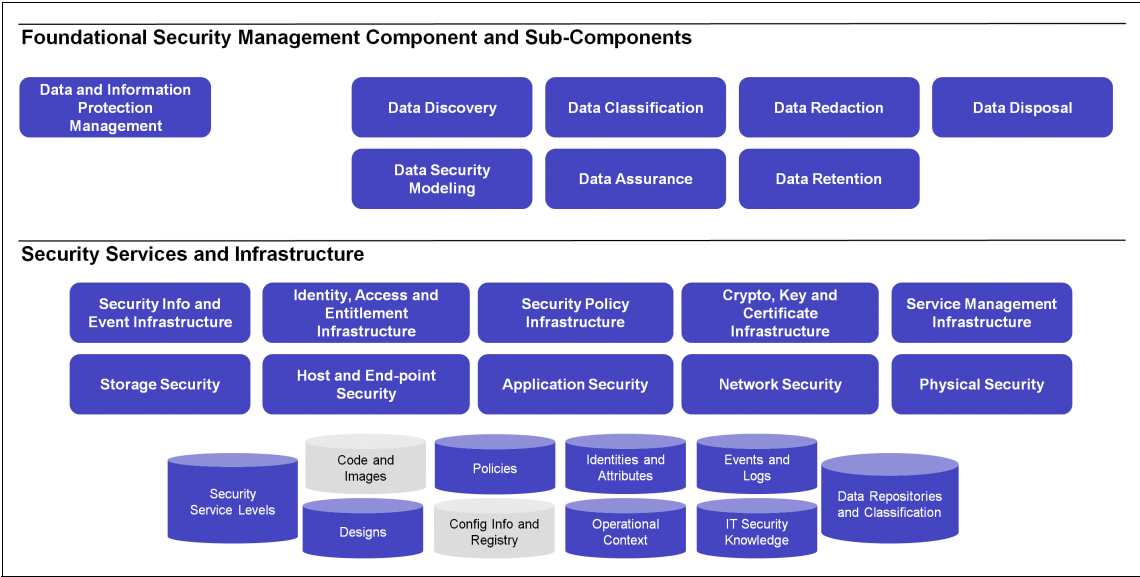


Figure 2-6 Data and Information Protection Management subcomponents

Data and Information Protection Management consists of the following subcomponents:

- ▶ Data Discovery
- ▶ Data Security Modeling
- ▶ Data Classification
- ▶ Data Assurance
- ▶ Data Redaction
- ▶ Data Retention
- ▶ Data Disposal

These services are explained in the following sections.

Data Discovery

The first step in securing data is having an accurate inventory of the organization's data repositories and understanding the security risks associated with them. Data Discovery is the process of identifying all the data repositories in an organization and analyzing the schema and data values and data patterns to identify relationships between the database elements.

Data Discovery looks at data relationships across repositories, understands how they relate to each other, and understands how the structured relationships are organized to represent business objects.

Data Discovery detects transformations and conditional logic that has been applied to data as it has been moved among repositories.

Building the business object view and understanding the transforms that the data has been subject to are key to planning master data management processes and business object archiving.

Data Security Modeling

Data Security Modeling refers to activities performed by a data architect to define domain-specific information models, logical data models, and physical data models.

Data Security Modeling captures the constraints on data types defined by an organization or an industry standard. They are business-oriented constraints that enable interoperability between systems and organizations. For example, bank routing codes represent a numeric string that follows certain rules in its format and interpretation for routing inter-bank transfers.

Logical data models are the semantic hub of an enterprise architecture. Logical data models are sometimes overlooked in the software development life cycle, but they have become increasingly important in the SOA context. A logical data model allows data architects to depict an overview of data entities in an application or an enterprise without having to look at overwhelming implementation details. Logical data models are often used as input into other enterprise architecture activities, such as defining message formats and service interfaces. The logical data model is also the starting point for transforming a domain model into a specific schema for a database instance.

Physical data models are database-specific models that represent relational data objects (for example, tables, columns, primary keys, and foreign keys) and their relationships.

Each of these layers of modeling can have security-related constraints attached to them that define requirements for confidentiality and encryption, access control, obfuscation, and redaction.

Data Classification

Data Classification refers to the tools and processes used to create a common set of semantic tags used by data modelers, data analysts, business analysts, governance stewards, and data architects.

Data Classification tools use rules and heuristics to examine logical data models from data repositories and associate business definitions with them.

In certain cases, business definitions relate directly to the format and constraints on the data (for example, the format of a telephone number). In other cases, there might be business-oriented definitions expressed in terms of the logical data model. For example, a *high value customer* might be defined as a customer who has bought products more than a specified number of times in a specified time period.

Data Classification manages both lower-level, logical data classification and business level classifications.

At the business level, a collection of business classifications, their definitions, and how they relate to the underlying logical data model creates a common business vocabulary, which can be used across the organization to ensure that every part of an organization agrees on the definition of the term. This helps to reduce confusion and miscommunication at the level of business discussions and also reduces interoperability errors across the IT organization.

A business vocabulary term might change over time and might have a significant impact on logical data models, database schema, and application logic. Data Classification tools can also enable data stewards to manage an orderly transition over time from one version of a business term to another.

Data Assurance

Data Assurance refers to tools and activities to make sure that data is cleansed and standardized to a defined model before it is used. Data Assurance also tracks the origin of the data when it is received through logging and auditing capabilities. Data Assurance processes also provide a governance checkpoint for aggregation, redaction, and obfuscation requirements to ensure confidentiality and privacy.

Data Redaction

Data Redaction refers to a set of tools and methods for eliminating sensitive or confidential data from a data set based on policy rules before it is given to a receiver.

Data Redaction techniques can be applied both to unstructured data, such as a collection of word processing documents, or structured data in databases.

A variety of techniques can be used in Data Redaction in addition to fully eliminating the data. For example, data can be partially obfuscated by masking out portions of the sensitive data. Data can be partially aggregated in ways that make it impossible to determine individual data records.

In certain techniques, errors can be deliberately introduced into data in ways that preserve confidentiality while preserving the ability to perform statistically valid operations on the data.

Data Redaction techniques enforce access control security policies while enabling the release of related and relevant data.

Data Retention

Data Retention capabilities cover both *backup* and *archive* tools and processes. Backup refers to the tools and activities needed to restore service to a well-known point in the event of system or media failure. Archiving refers to tools and processes to remove transactions from an active system that is no longer needed, but that might need to be preserved for legal requirements.

While backup techniques tend to apply to media or file-level activity, archiving often has to be aware of transactions. For example, a complete record of a business transaction might require preserving data from multiple tables in multiple databases and might even require preserving a variety of unstructured documents as well. Collectively, the set of structured and unstructured data that is needed to completely preserve a transaction is referred to as a *historical reference snapshot format*.

After archived, the snapshot files can remain on the local storage media or can be deleted. The organization controls how long an archive copy is to be retained, called the *retention period*. The retrieval process locates the copies within the archival storage and places them back into a designated system, which might be the active transactional system or a system specifically designed for displaying archived transactions.

Data Retention tools and techniques are an important component of a records management system that adds to these capability processes to manage decisions about what must be kept and in certain cases what must be deleted according to policy.

Data Disposal

Data Disposal refers to the tools and processes to delete data from a system that is no longer needed and required by law or policy to be retained. Disposing of data that is no longer needed reduces data management costs. In certain cases, regulations require that data be disposed of after certain time periods or when certain criteria are met.

Data Disposal processes can create a security risk if they inadvertently leave a way for the disposed data to be retrieved. Data Disposal tools and processes have to be designed to thwart likely threats to recovering the data, based on the value and sensitivity of the data and the techniques that an attacker might employ to retrieve the disposed data.

The techniques and processes for disposing of data are sometimes dictated by regulations and policy. For example, a regulation might require that data be overwritten a number of times with random information to reduce the possibility of retrieving it later.

Data Disposal tools and processes must also preserve sufficient records to show that the disposal processes have been followed.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective Data and Information Protection Management services (depicted as blue-shaded objects in Figure 2-6 on page 63):

- ▶ **Security Info and Event Infrastructure**

Interactions with databases and content repositories are one of the major sources of security events because they can create a log of data access attempts and logs of administrative activity by privileged users.

- ▶ **Identity, Access and Entitlement Infrastructure**

The Identity, Access and Entitlement Infrastructure translates the activity of privileged accounts to specific people who are responsible for data stewardship. In addition, database servers and content repositories are increasingly used to manage entitlements to access data. This can help tie database interaction to specific individuals, which is becoming more and more important due to increased compliance initiatives.

- ▶ **Security Policy Infrastructure**

Data access entitlements enforced by database servers and content repositories must be consistent with other access control policies in the organization. Integrating database servers and content repositories with the Security Policy Infrastructure helps to ensure this consistency. In addition, data retention policies, disposal policies, and other policies managed by data stewards should be authored, approved, and managed through a common Security Policy Infrastructure to ensure consistency with other security policies in the organization.

- **Cryptography, Key and Certificate Infrastructure**

Database servers, content repositories, and archive media capture *data at rest* and are subject to security requirements to encrypt data in case the storage media is subject to out-of-band attacks, such as media theft, making the data and information protection management systems dependent on the Cryptography, Key and Certificate Infrastructure.

- **Service Management Infrastructure**

Because data access management, data retention, and data disposal activities are often driven by regulatory requirements and are subject to audit, it is not sufficient to have the capability to perform the needed actions. It is necessary to show that the responsible people have configured the Data and Information Protection Management systems in accordance with agreed-upon policy and taken responsibility for the actions of the systems that they have configured. These activities require a robust Service Management Infrastructure to manage the work flow processes associated with these activities.

- **Storage Security**

In addition to cryptographic protection for data that is stored on storage media, additional Storage Security measures might be needed to protect the media and storage systems from tampering, theft, and copying.

- **Host and Endpoint Security**

Host and Endpoint Security is necessary for good Data and Information Protection Management to prevent access to the database servers and content repositories via the file system in the operating system.

- **Application Security**

Application Security is important for Data and Information Protection Management because compromised applications might be able to access the database servers and content repositories using the credentials of the application and issue unauthorized queries to them.

- **Network Security**

Network Security is important to Data and Information Protection Management to protect data while it is in transit. While message-level encryption and connection-level encryption can be used to protect data in transit, a secured network is important to prevent out-of-band attacks such as copying traffic for later decryption, man-in-the-middle attacks, DNS cache poisoning, and so on.

- **Physical Security**

Physical Security for the database servers and content repositories is important to prevent out-of-band attacks, primarily media theft.

- ▶ Security Service Levels

Security Service Levels can contain the agreed-upon data retention and disposal activities and the agreements regarding data access logging necessary for demonstrating policy compliance.

- ▶ Designs

Domain, logical, and physical data models for the organization are key designs that are used in a wide variety of enterprise IT architecture activities, including capacity planning for storage, application design, and message format design.

- ▶ Policies

Data retention policies and data disposal policies are key policies in every IT organization. Furthermore, the policies enforced by database servers and content repositories to manage access to the data must be consistent with access control policies at other layers of the application stack.

- ▶ Identities and Attributes

Credentials used by administrators and users to access data must be associated with individuals so that accountability for data access and usage can be managed. Often, attributes of individuals dictate the subset of data that they are authorized to see. For example, a sales manager might only be allowed to see the sales data for the region that he manages.

- ▶ Operational Context

Increasingly, database servers and content repositories are required to be aware of not only the credentials used to access them, but also the credentials that originated the request so that the data access logs can be associated with a responsible individual. The database servers and content repositories often need to log the transaction IDs or other unit-of-work identifiers for audit purposes.

- ▶ IT Security Knowledge

There are several areas of general IT Security Knowledge that are important to Data and Information Protection Management. For example, understanding the relative strength of encryption algorithms and key lengths is important when determining encryption protection for sensitive data. Understanding the most common ways that data media are stolen is important in determining media protection.

- ▶ Data Repositories and Classification

The first step in protecting data and information is keeping accurate inventories of where all the database servers and content repositories are and understanding their value and sensitivity.

2.2.6 Software, System and Service Assurance

Software, System and Service Assurance addresses how software, systems, and services are designed, developed, tested, operated, and maintained throughout the software life cycle to create predictably secure software. This component covers structured design, threat modeling, software risk assessment, design reviews for security, source code reviews and analysis, dynamic application analysis, source code control and access monitoring, code/package signing and verification, quality assurance testing, and supplier and third-party code validation.

Figure 2-7 shows an overview of the Software, System and Service Assurance subcomponents and the related components from the Security Services and Infrastructure layer.

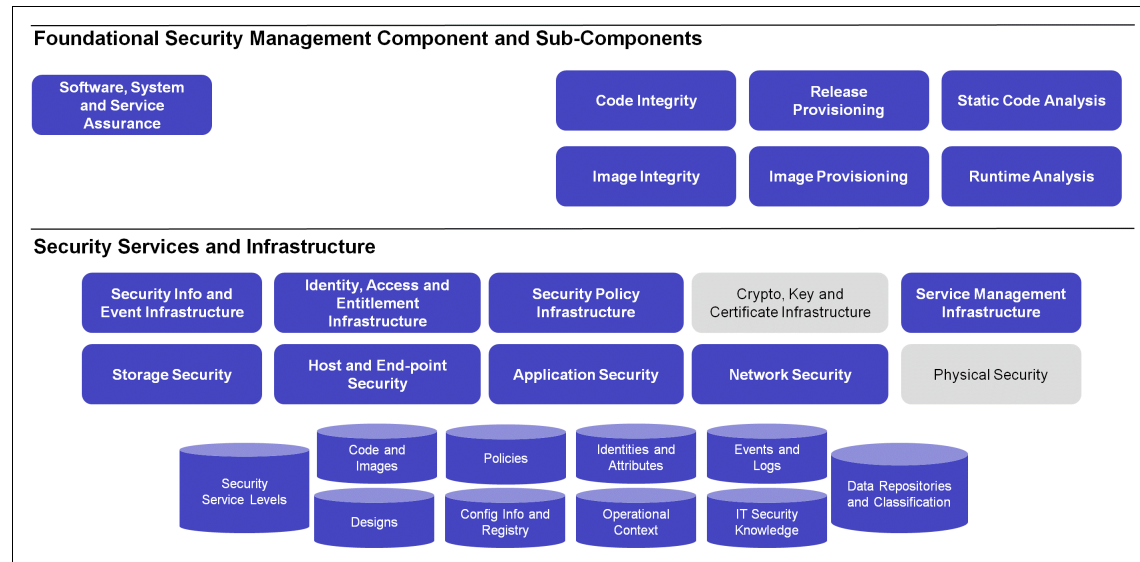


Figure 2-7 Software, System and Service Assurance subcomponents

Software, System and Service Assurance consists of the following subcomponents:

- ▶ Code Integrity
- ▶ Image Integrity
- ▶ Release Provisioning
- ▶ Image Provisioning
- ▶ Static Code Analysis
- ▶ Runtime Analysis

These services are explained in the following sections.

Code Integrity

Code Integrity refers to protecting assets used to build and run application object code to ensure that what is delivered to service management for deployment has not been tampered with or incorporated any unknown source code.

Code Integrity encompasses confidentiality of the source code from competitors and other unauthorized people. Code Integrity also ensures that all the proper licenses have been obtained for the running instance of the code and ensures compliance with any development team restrictions (for example, *clean room* rules might need to be followed to protect against charges of reverse engineering).

Image Integrity

Image Integrity covers the entire runtime stack, from operating system to middleware components and application platforms that are needed to run the application or service. Images might include definitions of runtime dependencies that are assembled during the deployment process, or an image might be an entire pre-built software stack packaged as a virtual machine image.

In the case of virtual machine images, image integrity refers to the tools and processes to track the provenance of all the software components that are included in the image. Image Integrity also ensures that the image has not been tampered with after it has been assembled.

Release Provisioning

Secure provisioning ensures that handing over code to release management for installation and configuration of dependent software infrastructure is done in accordance with security policy and, in certain cases, per contract with the customer. For example, release provisioning might include a mapping of organizational roles and individuals to application-defined entitlement roles to ensure that the correct people in an organization are granted access to the correct application roles. Release Provisioning might also dictate security requirements for the database middleware to define requirements for protecting data at rest.

Image Provisioning

Image Provisioning manages access to the image contents. For example, image administrators might not be authorized to see confidential data or code inside the image. Image provisioning also manages access to the image for deployment, defining who can access and deploy instances of the image in a production environment.

Finally, Image Provisioning might impose deployment restrictions, especially security-related restrictions, on the service deployment processes. For example, an image might have a requirement that it is not deployed in a DMZ, but only behind strictly controlled firewalls. Or an image might have a requirement to not be co-hosted with images from any other company.

Static Code Analysis

Static Code Analysis refers to the tools and processes that are usually instituted by a software development team or a build team to examine all the artifacts and components that are used to build an application. The analysis looks for security vulnerabilities and poor coding practices that can create security, performance, or other problems.

Static Code Analysis usually refers to automated tools that scan source code and report on potential problems. But Static Code Analysis can also include design model reviews and scanning, in addition to manual inspection of application artifacts.

Runtime Analysis

Runtime Analysis, or *software profiling*, refers to the ability to observe a running software system and analyze its behavior to detect vulnerabilities in the code.

While Runtime Analysis is often used to look for problems with memory usage, network usage, or other runtime resources, runtime analysis can also be used to identify potential security problems. For example, Runtime Analysis can highlight how an application fails to properly handle malformed messages resulting in failure to release allocated memory.

Runtime Analysis is an *internal view* of the running application, whereas *dynamic analysis* tests a running application by interacting with it from an *external perspective* in the same way that user or client software interacts with it.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective 2.2.6, “Software, System and Service Assurance” on page 70 (depicted as blue-shaded objects in Figure 2-7 on page 70):

- Security Info and Event Infrastructure

When planning the deployment of an application, the security-relevant events that it might generate need to be planned for in the Security Info and Event Infrastructure. IT operations must know how to enable the application-specific event logging and understand where the events are stored.

IT operations must also incorporate the application-specific logging into their Security Info and Event Infrastructure and understand how to recognize potential security incidents from the event stream.

- ▶ Identity, Access and Entitlement Infrastructure

Control of access to source code, images, and running applications must be tied to specific individuals by associating credentials with individuals and associating organizational roles with the management infrastructure roles and application-defined roles.

- ▶ Security Policy Infrastructure

Security policies regulating how applications are deployed into an environment and how machine images are deployed into a virtualization platform can be managed by a Security Policy Infrastructure to ensure consistency across the organization. Access control policies to applications and images should be coordinated with other access control policies.

- ▶ Service Management Infrastructure

Assurance activities define deployment and access requirements for applications and images that must be consumed and implemented by the release and deployment processes in the Service Management Infrastructure. Therefore, there needs to be coordination between development and operations to ensure that operations know how to implement the requirements defined by development.

- ▶ Storage Security

Applications define their dependency on storage infrastructure, and storage infrastructure components can be included in a virtual machine image. In both cases, the definitions might impose security requirements on the storage infrastructure, including requirements to encrypt storage media, locate it in a physically secure environment, and maintain data for a specified retention period.

- ▶ Host and End-point Security

Applications typically have limited awareness of the host environment and rely on security measures on the host to protect the application from out-of-band attacks. For example, it is the responsibility of Host and End-point Security to ensure that there are no processes on the host machine intercepting traffic between the application and its clients.

- ▶ Application Security

While secure coding practices, static analysis, and secure design practices can limit the vulnerabilities in an application, the applications typically rely on Application Security enforcement points to help detect and prevent attacks such as cross-site scripting, SQL injection, and so on.

- Network Security

Applications have dependencies in the Network Security infrastructure to make certain ports available for remote connections and to ensure the appropriate isolation of network traffic. Certain application-layer attacks can be detected and prevented via deep packet inspection and other types of network traffic analysis.

- Security Service Levels

Government agencies and companies are starting to require assurances that software code is free of viruses, malicious coding, vendor or programmer created backdoors or trapdoors (front and back), and other types of security vulnerabilities, which are considered a type of security service level for the software.

- Code and Images

Application Code and Images are the target resources protected by software, server, and security assurance.

- Designs

The application architecture as represented in the software designs is the first line of defense in software security. Equally important, the application designs represent the formal definition of what the software does and delivers and are part of the provenance of a software application. Good software provenance should be able to trace a chain of activity from the running application back to the design that it implements.

- Policies

There are a wide variety of Policies that affect software assurance and image integrity. Applications define sets of roles that dictate which features are accessible to which people. These roles need to be mapped to organizational roles by means of an access control policy specific to the application.

Likewise, security policies must define who has the authority to instantiate which virtual machine images under which circumstances. The security policies must also define where in the virtual environment these images can be present. For example, security policy might dictate when virtual machine images must be placed on a separate virtual network from other images.

- Identities and Attributes

Because applications are typically developed independently of any particular organization, they define access control mechanisms in terms of application-specific roles. These application-specific roles have to be mapped to the organizational roles, which requires an understanding of the directory information available about people and their credentials.

- ▶ **Operational Context**

Applications often rely on transactional context that comes from outside the application's environment. For example, an application might need to send SNMP events to a central management infrastructure, which must be defined to the application.

- ▶ **IT Security Knowledge**

An understanding of the current types of attacks that applications are typically subject to is important in planning application architecture and design and is crucial to static code analysis. Other types of industry knowledge, such as the most common programming errors that lead to security vulnerabilities, are also extremely important.

- ▶ **Data Repositories and Classification**

Virtually every application or web service relies on some sort of storage infrastructure for structured or unstructured data. These are typically defined at deployment time and must be registered with a data repository catalog and classified according to organizational policy to ensure they are adequately protected.

2.2.7 Threat and Vulnerability Management

Threat and Vulnerability Management provides services that can help determine security threats and identify vulnerabilities in deployed systems, collect security-related information from various internal and external sources, and determine the appropriate response.

Figure 2-8 shows an overview of the Threat and Vulnerability Management subcomponents and the related components from the Security Services and Infrastructure layer.

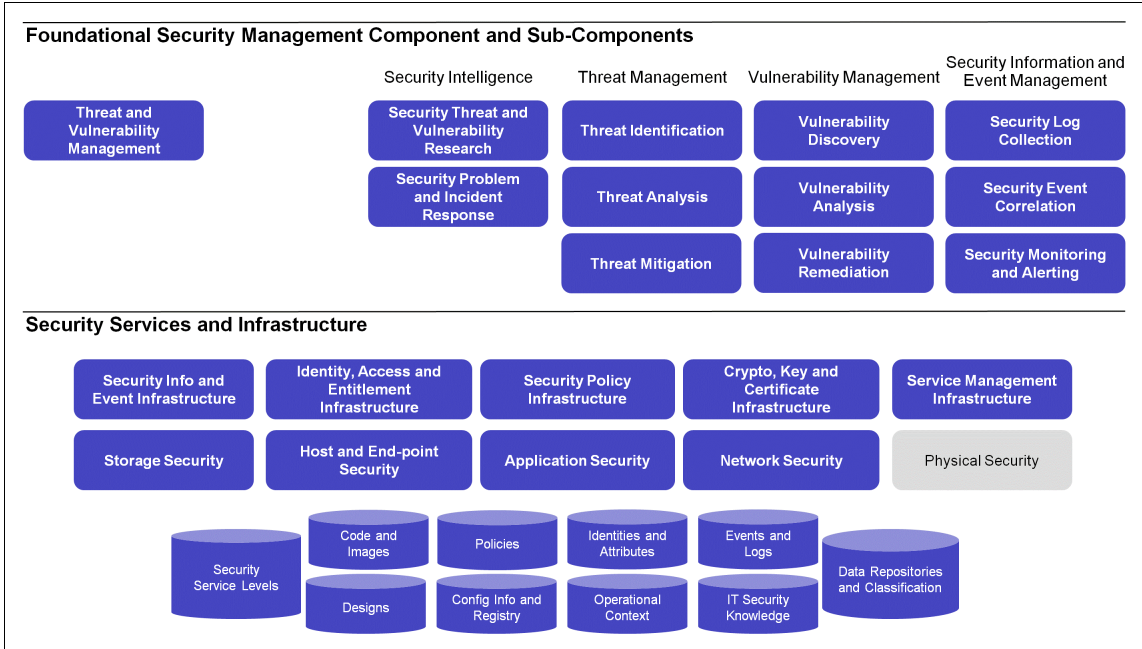


Figure 2-8 Threat and Vulnerability Management subcomponents

Threat and Vulnerability Management consists of the following subcomponents:

- ▶ Security Intelligence
- ▶ Threat Management
- ▶ Vulnerability Management
- ▶ Security Information and Event Management

These four services are explained in the following sections.

Security Intelligence

Security Intelligence provides security knowledge about threats and vulnerabilities.

Security Threat and Vulnerability Research

Security Threat and Vulnerability Research represents the ability to collect, analyze, and disseminate information as it pertains to computer security from reviewing and tracking a range of available information sources on potential

threats and potential vulnerabilities to determine the applicability to an organization's IT environment.

In a more sophisticated execution, Security Threat and Vulnerability Research also includes the detailed observation, manipulation, and analysis of the behavior of *threat agents* and of the composition of *vulnerability conditions* in attack scenarios to synthesize, create, and provide respective knowledge about potential future attacks from collected data. It takes into account external, situational awareness, identifies and examines possible new attack patterns, and monitors long-term trends that might lead to specific new threats against the security of information assets.

People and processes associated with this component are also responsible for gathering awareness of future potential threats and vulnerabilities to the facilities from law enforcement agencies, regulatory agencies, and industry trade groups.

Security Incident and Problem Response

Security Incident and Problem Response provides support to the related IT Service Management functions *Incident Management* and *Problem Management* by providing security expert knowledge on identified attacks and security-related anomalies and by recommending respective actions to manage security incidents and problems to closure. It embraces security incident containment, security incident recovery, root cause analysis, security problem analysis, and security problem resolution.

Threat Management

The Threat Management services deal with the identification, understanding, and counterfighting of specific threats that might exist for a given IT environment.

A *threat* is the intention of a *threat agent* to perform a *threat action* to exploit a specific vulnerability. Only the occurrence of both threat and vulnerability together define the likelihood of a risk. If either threat or vulnerability does not exist, that risk has a likelihood of zero. That is, there is no risk.

Threat Identification

Threat Identification embraces activities that help to discover actors and actions in the IT environment that might have a harmful effect on IT assets and the information stored and processed on them. Threat Identification can be performed purely manually, but today it can usually be based on the automated recognition of deviations from the usual operations in an IT environment. Any discovered anomalies can then be examined for their threat potential.

Threat Analysis

Threat Analysis is the continuous examination of available information related to *threat agents*, often called attackers, and their possible *threat actions*, the actual attack, to evaluate the severity of an identified threat, for instance, based on the potential occurrence of an attack due to the general awareness of the attack vector and on the presumed attractiveness of an organization as an attack target in the view of an attacker.

Threat Mitigation

Threat Mitigation embraces efforts taken to influence either the threat agents or to manipulate the threat actions to reduce the severity of a threat. Usual efforts taken include, for instance, declaring sanctions and disciplinary actions to threat agents upon discovery of their threat actions or even their planning of such threat actions as well as the deployment of measures that negate their actions or identify and deviate them away from a given vulnerability. Threat mitigation can consist of detective controls, such as the deployment of malware detection, intrusion detection, and honeypots.

Vulnerability Management

The Vulnerability Management services deal with the discovery, understanding, and reduction of specific vulnerabilities that might exist for a given IT environment.

A vulnerability is a weakness that can be exploited to compromise security.

Vulnerability Discovery

Vulnerability Discovery deals with the detection of vulnerabilities. Besides application of holistic security thinking, well-known methods for Vulnerability Discovery include:

- ▶ *Dynamic code analysis* to assess applications for vulnerabilities that might be exploited from an application user's perspective.
- ▶ *Network vulnerability scanning* to probe operating systems, databases, middleware, and firewalls, which protect all deployed IT services from vulnerabilities that are accessible from the internet. The difference from dynamic code analysis is that network vulnerability scanning focuses more on off-the-shelf software packages, whereas dynamic code analysis focuses mostly on custom-built applications.
- ▶ *Security healthchecking* to check systems with scripts or via a local agent from the inside and assess the configurations of local and network services of operating systems, databases, middleware packages, and applications for errors that could lead to potentially exploitable vulnerabilities.

- ▶ *Ethical hacking* to perform simulated attacks against a part of or the entire IT environment by applying human creativity and out-of-the-box-thinking, and by using a combination of automated discovery, probing and exploit tools, and manual or custom-scripted security tests. Such attacks can vary in scope, time, and resourcefulness as well as in the provision of inside knowledge and access rights to simulate different attack scenarios. Providing no inside knowledge is considered a blackbox test, whereas providing the testers with background information about the design and architecture is considered a whitebox test.

Vulnerability Analysis

Vulnerability Analysis covers the actual verification of vulnerabilities by eradication of false positives, and further covers the rating of such vulnerabilities in terms of criticality (for instance, based on their ease of discovery and the complexity of their exploitability by attackers, as well as on the level of resulting compromise of a tested system or environment).

Vulnerability Remediation

Vulnerability Remediation encompasses the combination of deterrent, preventive, detective, and corrective security controls to mitigate identified and verified vulnerabilities. The most commonly applied mitigation approaches to eliminate a vulnerability include the following measures:

- ▶ Fix the related code by patching.
- ▶ Change the configuration of the vulnerable service.
- ▶ Apply additional preventive security controls such as firewall and intrusion prevention systems with virtual patching capabilities.
- ▶ Employ additional corrective measures, such as increased frequency of system checks, data backups for quicker recovery, and enhanced emergency response procedures.

Security Information and Event Management

After the event data has been centrally collected, it can be consolidated and structured as well as combined and correlated to derive more meaningful and human-understandable security information.

Security Log Collection

Security Log Collection refers to the ability to collect security-related events from various collection points in the IT environment, usually in the form of system, network, and security log and alert data, and to store them in a structured way in order to have a redundant copy (alongside the logs on the originating systems) in order to retrieve and analyze them during security incidents and problems in case the logs on the originating systems have been compromised.

Security Event Correlation

Security Event Correlation builds upon Security Log Collection. After the log data has been centrally collected, it can be consolidated and structured, standardised, and combined and correlated into security events to derive more meaningful and human-understandable security information.

Security Monitoring and Alerting

Security Monitoring and Alerting refers to all activities related to the ongoing and frequent observation of the technical infrastructure for deviations from the standard operation, which confirm or at least indicate an impact on security.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective 2.2.7, “Threat and Vulnerability Management” on page 75 (depicted as blue-shaded objects in Figure 2-8 on page 76):

- ▶ **Security Information and Event Infrastructure**

The Security Information and Event Infrastructure collects security log data from various agents that are deployed throughout the IT environment. It has the ability to create events and incidents that can be combined with other events and incidents in a standardized format by consolidating, classifying, and correlating all collected information. The aggregation of security logs and subsequent derivation of security information is essential for all vulnerability-related services within the Threat and Vulnerability Management discipline. The large amount of data collected over time allows the Security Information and Event Infrastructure services to analyze trends of attack patterns as part of the security intelligence services and thus can also help to derive probabilities of threats.

- ▶ **Identity, Access and Entitlement Infrastructure**

The Identity, Access and Entitlement Infrastructure is used by the Threat and Vulnerability Management services to further analyze and tie back events to identities and entitlements to confirm whether events relate back to authorized activities or whether they occurred from unauthorized or even malicious activities.

- ▶ **Security Policy Infrastructure**

The Security Policy Infrastructure can help Threat and Vulnerability Management services to eliminate or reduce *false positive* events. A false positive is an event that, from a pure technical security perspective, is considered a threat. For example, a particular event has been granted a *policy exception* because a business application requires a specific network port to be used, although this port is known to be used for attacks.

By consolidating with the Security Policy Infrastructure, this particular event will no longer be flagged as a security event.

- **Cryptography, Key and Certificate Infrastructure**

Communication between distributed infrastructure components for Threat Management, Vulnerability Management, and between systems and the Security Information and Event Management infrastructure components is subject to encryption and secure authentication using certificates. Also, the log data might have to be encrypted or signed to protect against manipulation, so that Threat and Vulnerability Management services are dependant on the Cryptography, Key and Certificate Infrastructure.

- **Service Management Infrastructure**

Threat and Vulnerability Management services operate within agreed-upon service management infrastructures and must also utilize the services provided by that infrastructure. For example, performing vulnerability discovery activities and transferring evidence from the testing environment are typical operations that must be performed in line with change management and thus use the Service Management Infrastructure.

- **Storage Security**

Storage Security provides logging and alerting functionality that can be used and examined either directly by Threat and Vulnerability Management services or indirectly by the Security Information and Event Infrastructure. Storage Security can also employ dedicated monitoring agents that can provide a more comprehensive functionality than basic logging and alerting. Storage Security can also provide masking and filtering functionality that comes with most database products to allow improved vulnerability discovery and mitigation.

- **Host and End-point Security**

Like Storage Security, Host and End-point Security can provide security functionality that allows the Threat and Vulnerability Management services to identify and remediate vulnerabilities either proactively or reactively. Examples of such functionality includes malware scanning and remediation software, host intrusion detection and prevention systems, and security healthchecking software. Besides deploying additional software, many basic operating systems and middleware components provide configuration options to limit security vulnerabilities, or even the potential for future vulnerabilities by configuring stricter values and thus hardening systems against attacks.

- **Application Security**

Application Security provides options for security configurations and might include security defense mechanisms like input revalidation to close known and popular attack vectors.

- ▶ Network Security

Network Security provides filtering, monitoring, alerting, and discovery functionalities by using firewalls, routers, network device logging, network intrusion detection and prevention systems, and network protocol and application protocol vulnerability scanners.

- ▶ Security Service Levels

The Security Service Levels provide the operational background for the security policies. This information helps the Threat and Vulnerability Management services to better implement the required level of protection. Also, as Threat and Vulnerability Management services often operate using high privileges and access rights in the IT environment, it is important that these services follow the appropriate policies set for their activities.

- ▶ Code and Images

Code and Images are constantly examined by Threat and Vulnerability Management services for identified vulnerabilities within them.

- ▶ Designs

Designs are an important reference for Threat and Vulnerability Management services, as they allow you to derive potential attack and testing scenarios for vulnerability discovery and threat analysis services.

- ▶ Policies

Policies are required to be adhered to by Threat and Vulnerability Management services, especially as these services operate with high, sometimes ultimate, privileges in the IT environment. Because certain Threat and Vulnerability Management services emulate attacks, the approach and limits of such activities must be strictly regulated in policies before they can be executed.

- ▶ Configuration Information and Registry

The configuration management database and the registries of IT resources are used to store security settings and important asset information. This information needs to be available for a root cause analysis as part of a security threat investigation or a vulnerability examination as part of a security vulnerability assessment.

For instance, it is essential to check the configuration information to examine the reason for an identified dangerous configuration. It might have been introduced as part of an approved configuration change or it might have been introduced as part of a malicious system attack. By examining the recorded configuration information stored in the configuration management database, security administrators are able to determine either regular behavior or malicious intent and act accordingly.

Likewise, a vulnerability examination can greatly benefit from configuration and registry information because this information can be helpful to determine the number and location of systems that are exposed to a specific vulnerability.

- Identities and Attributes

Identities and Attributes are assessed by Threat and Vulnerability Management services as part of vulnerability discovery and incident and problem response tasks.

An example for the discovery of a security problem with an abuse of identity can be identified by cross-checking the user activity on systems with the attributes of the corresponding identity. In a case where the stored identity information for a particular user ID shows an attribute *revoked*, and there are still activities performed on systems in the context of this particular user ID, there is a high likelihood that this user ID is used in a malicious context.

- Operational Context

The Operational Context can help clarify whether an activity and its related events are harmless and intended or unplanned and potentially malicious. Thus, the Operational Context has to be reviewed by the Threat and Vulnerability Management services to come to a correct conclusion. For example, a discovered suspicious activity, like an internal network scan, might be related to authorized changes or problem determination activities. Because the Operational Context clarifies the legit intention, this event does not represent a potential attack.

- IT Security Knowledge

A deep and broad IT Security Knowledge is of key importance to Threat and Vulnerability Management. The type of knowledge required includes a deep technical understanding of platform-specific security functions and the ability to understand the performance of security attacks in a step-by-step manner. Besides the technical knowledge, it is also required that security experts working in Threat and Vulnerability Management are always up to date on new technologies so that they are able to identify potential new types of threats that might come with these innovations. Alongside of the IT Security Knowledge, it is also necessary to have skills in using the various security analysis and testing tools.

Finally, the provision of these services requires the ability to understand new security attack patterns and also the skills to efficiently keep up to date on newly discovered threats and vulnerabilities.

- Events and Logs

Event and logs are the most essential objects for the Threat and Vulnerability Management services, as they contain all the collected log and event information necessary to identify actual attacks.

► Data Repositories and Classification

Understanding of the Data Repositories and Classification is required by the Threat and Vulnerability Management services to allow thorough analysis of potential threats and targeted creative thinking about potential vulnerabilities and related attack patterns.

2.2.8 IT Service Management

IT Service Management provides the process automation and workflow foundation for all IT delivery activities, including security management. In particular, change management and incident management processes play a significant role in security management.

Restriction: This section is not intended to be a complete discussion of all IT Service Management domains. We focus on the key IT Service Management components that contribute to security.

Figure 2-9 shows an overview of the IT Service Management subcomponents and the related key components from the Security Services and Infrastructure layer.

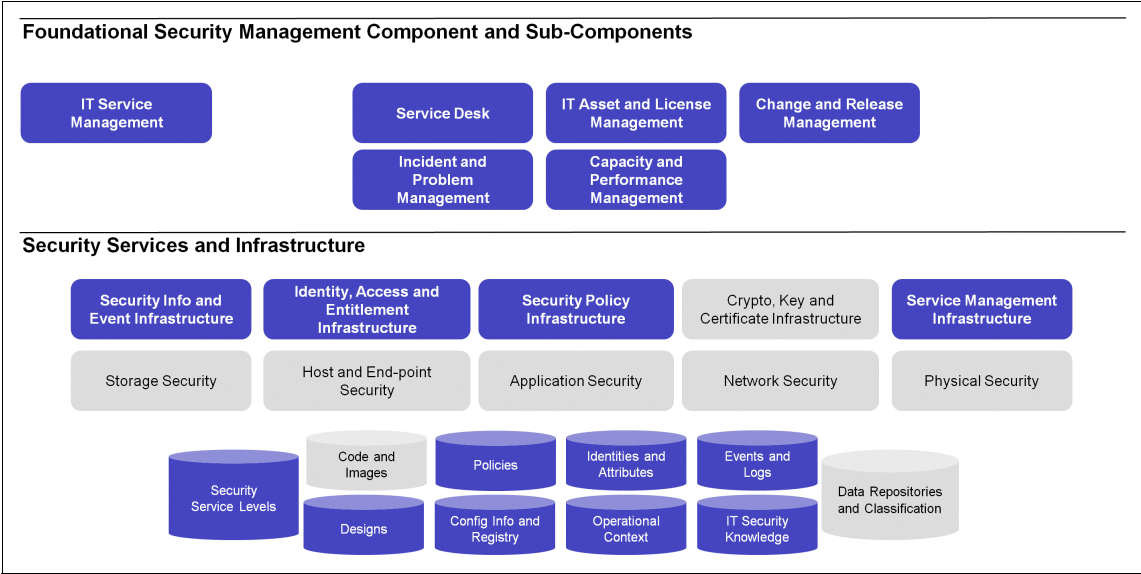


Figure 2-9 IT Service Management subcomponents

IT Service Management consists of the following subcomponents:

- ▶ Service Desk
- ▶ Asset and License Management
- ▶ Change and Release Management
- ▶ Incident and Problem Management
- ▶ Capacity and Performance Management

These services are explained in the following sections.

Service Desk

Service Desk refers to the *single point of contact* (SPOC) for all IT Service Management related matters where all service management functions are coordinated. In particular, the Service Desk provides a ticketing and tracking functionality for service delivery activities, including activities in the security management area.

Asset and License Management

Asset and License Management covers a set of capabilities to monitor deployed IT assets from a financial, compliance, and inventory perspective.

From a software perspective, Asset and License Management includes license management, certain aspects of configuration management (for inventory management purposes), and reporting for regulatory purposes. License management maintains an inventory of deployed software, measures usage activity, and manages entitlements to licensed software. It checks for adherence to license use requirements, summarizes software use for planning purposes, and assists in user charge-back activities.

Hardware asset management includes the physical characteristics of deployed hardware components in the IT environment, such as their make and model numbers, serial numbers, physical locations, and their role and placement in the network. Hardware asset management involves tracking regular maintenance of the hardware assets, tracking history of physical failures, and so on. Hardware asset management is often also involved in recording and tracking the financial view of the asset.

Change and Release Management

Change and Release Management covers the standardization of methods and work processes to manage changes to the configuration of deployed IT assets and to the upgrade of existing deployed software components and the deployment of new software components. The goals of this standardization are to minimize disruption of service and to ensure that software and hardware components are not deployed in ways that compromise any security or integrity aspects.

Incident and Problem Management

Incident and Problem Management handles the methods and processes used to restore service from any sort of disruption due to incidents and problems. An *incident* is considered a single event or a group of events that occur in parallel or in a short period in time and that trigger a negative impact on the level of service. A *problem* is considered a result of repetitive incidents of the same or a similar pattern, or as a result of an elevation of an incident due to its continued significant impact on the level of service or due to the increased efforts that are required to return to normal operations.

From a security perspective, incidents and problems can be classified as security incidents and security problems and will then require support from security incident and problem support or security emergency response teams.

Capacity and Performance Management

Capacity and Performance Management deals with the planning, provisioning, and optimization of IT resources that are required for the IT services. In a narrow sense this mostly refers to details like processing power and memory, system backup and archive storage, and network speed or bandwidth. In a wider context, Capacity and Performance Management can also include human resources and physical asset resources like floor space in a data center. This area is important to security, as the security services and their related infrastructure components deployed in an IT environment can consume a significant amount of resources, and thus can impact performance.

All too often such impact is ignored or not properly examined when security is not embedded in the planning of IT services from the start and also when the addition of new security controls are considered (for instance, as a result of a security incident remediation).

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective IT Service Management (depicted as blue-shaded objects in Figure 2-9 on page 84):

- Security Information and Event Infrastructure

The Security Information and Event Infrastructure is used by the IT Service Management services to monitor and observe changes in security-related assets that might result or relate to change and release execution or trigger incidents and problems, which, when confirmed, can become security incidents and security problems and be provided to the Threat and Vulnerability Management services for resolving.

► Identity, Access and Entitlement Infrastructure

The Identity, Access and Entitlement Infrastructure is used by the IT Service Management services to assign potential actioners for change and release, incident and problem, and capacity and performance management activities on systems.

The Identity, Access and Entitlement Infrastructure is also used by IT Service Management to review and authorize access to the components of the IT environment because IT Service Management is the owner of and thus overall is responsible for the IT services.

► Security Policy Infrastructure

The Security Policy Infrastructure is used by IT Service Management services to check and verify security requirements that must be adhered to (for example, under which conditions and in which timeframes) to avoid negative impact during changes and releases and during incident and problem handling activities.

► Service Management Infrastructure

The Service Management Infrastructure provides the overall ticketing and tracking, and progress and status reporting system for all IT Service Management services.

► Security Service Levels

The Security Service Levels are a subset of the overall IT service levels that IT Service Management must deliver and report on. The IT Service Management services (in particular the service desk) has to consider potential impact to the Security Service Levels by other service activities when planning and scheduling those.

► Designs

Designs are important to IT Service Management services to understand potential impacts to the services. For example, planned and accepted changes to one component can have possible effects on other components, which is of particular importance for the capacity and performance management services.

► Policies

Policies can help IT Service Management to identify and confirm security requirements that must be considered during any of the IT Service Management services activities.

- ▶ Configuration Information and Registry

The Configuration Information and Registry is most used and updated as a consequence of IT Service Management services and must be kept up-to-date in line with their activities to represent an accurate state of the deployed configurations.

- ▶ Identities and Attributes

Identities and Attributes feed directly into the Identity, Access and Entitlement Infrastructure, which is used by the IT Service Management services as described above.

- ▶ Operational Context

As with designs, IT Service Management services use and update the Operational Context for the IT environment in line with the change, release, and other IT Service Management activities.

- ▶ Events and Logs

Event and logs are created alongside the activities of IT Service Management services, and thus the event and log items are used to check and validate actual progress of initiated activities.

- ▶ IT Security Knowledge

The IT Security Knowledge required for IT Service Management activities consists mainly of the general understanding of security matters and of the security awareness required to prioritize and sufficiently consider security in general IT Service Management activities. For instance, incident and problem management must have sufficient security understanding to identify that an incident or problem might be related or have an impact onto the security posture.

2.2.9 Physical Asset Management

Physical Asset Management provides awareness of the location and status of physical assets and awareness of Physical Security controls and coordinates the security information for physical systems with the IT security controls.

Consideration: This section is not intended to be a complete discussion of all Physical Asset Management domains. We focus on the key Physical Asset Management components that contribute to security.

Figure 2-10 shows an overview of the Physical Asset Management subcomponents and the related components from the Security Services and Infrastructure layer.

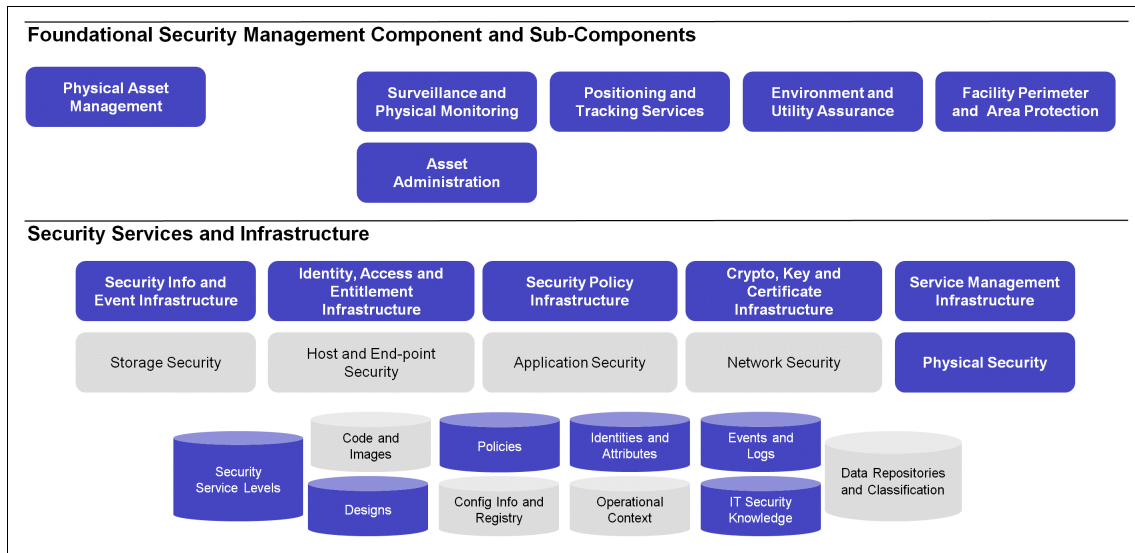


Figure 2-10 Physical Asset Management subcomponents

Because of the ongoing convergence of physical and IT security, Physical Asset Management is a major concern, although it builds its own discipline in IT management that has a much wider purpose.

Physical Asset Management consists of the following subcomponents:

- ▶ Surveillance and Physical Monitoring
- ▶ Environment and Utility Assurance
- ▶ Facility, Perimeter and Area Protection
- ▶ Positioning and Tracking Services
- ▶ Asset Administration

These services are explained in the following sections.

Surveillance and Physical Monitoring

Surveillance and Physical Monitoring covers all investigative physical security controls and is the equivalent of IT technical monitoring, including real-time observation of physical assets to detect physical attacks, theft, abuse, and other unusual and suspicious events. Such controls can include physical alarm systems triggered by opening doors and gates, breaking or opening windows and hatches, moving objects, or simple discovery of intruders due to motion

detection. Surveillance and Physical Monitoring can be performed using direct or indirect human supervision or automated systems that can analyze changes in normal and infrared light or sound patterns of the monitored area. Surveillance and Physical Monitoring can record evidence over a longer period of time to investigate security-related situations retrospectively.

Environment and Utility Assurance

Environment and Utility Assurance covers the provisioning of electricity and other power utility related supplies and climate controls. Environment and Utility Assurance is an integral part of facility management that can have a significant impact on the availability of IT services and hence on security.

Facility, Perimeter and Area Protection

Facilities, Perimeter and Area Protection covers the provisioning and management of preventive, deterrent, and reactive physical security and safety controls of a human or automatic nature. This service includes site-planning activities to address known risks from natural disaster, political events, and other external threats.

Positioning and Tracking Services

Positioning and Tracking Services are related to the identification of the location and movement of tangible physical assets, in this context, of those assets with valuable information that need to be protected. This can include short-range and long-range tracking, up to a worldwide scale.

Asset Administration

Asset Administration covers the coordination of activities related to the provisioning, building and procurement, maintenance and updating, movement, decommissioning, and destruction of primarily tangible but also non-tangible physical assets. These activities go beyond pure IT assets, but mostly focus on assets that have a direct or significant impact on information security. Examples of such assets include, but are not limited to, real estate buildings that provide office floor space or data centers, cable and utility channels, and data tape storage containers and their transportation vehicles.

Security Services and Infrastructure components

The following Security Services and Infrastructure components are key to effective Physical Asset Management (depicted as blue-shaded objects in Figure 2-10 on page 89):

- ▶ Security Info and Event Infrastructure

The information about physical environments recorded through surveillance and sensors is increasingly being indexed and converted to IT security events that can be correlated and combined with other IT events. For example, an authorization record regarding the access of an application can be correlated with an event representing a person using their badge to access a door. Likewise, these records can be correlated with segments of video surveillance footage with matching timestamps.

- ▶ Identity, Access and Entitlement Infrastructure

High-value assets in a physical environment are often protected by both physical controls (fences, guards, and so on) and logical access (badge readers, RFID detectors).

- ▶ Security Policy Infrastructure

The Security Policy Infrastructure that is used to manage organization roles and their entitlements to IT resources, such as applications, can also be used to manage the policies that govern activities in the physical environment. For example, the Security Policy Infrastructure can be used to author the policies that security personnel use to enforce access control if a person can or cannot pass a physical checkpoint on the premises.

- ▶ Cryptography, Key and Certificate Infrastructure

Many physical credentials, such as access badges, smart cards, or passports, are increasingly embedding logical credentials, such as public key certificates, which have to be managed by a Cryptography, Key and Certificate Infrastructure.

- ▶ Service Management Infrastructure

Service Management Infrastructure processes are often combined to manage both IT security and physical security incidents so that one service desk and one workflow infrastructure can manage both in one place.

- ▶ Physical Security

The Physical Security infrastructure, including barriers, fences, secure construction, and other types of inert security, can provide a base for providing an overall secure environment for an organization. The personnel, such as security guards and inspectors, add to the base security by enforcing operational processes on a day-to-day basis. The runtime aspects of Physical Security depend on the Physical Security infrastructure.

For example, the placement of surveillance equipment depends on the layout of the physical environment. If the physical environment is not designed with security in mind, it can be more difficult to place surveillance equipment effectively.

- ▶ Security Service Levels

The security service level agreements must, at least, delegate authority for Physical Security to an accountable person. Certain agreements even define fine-grained details, like specific physical controls (barriers, perimeter checkpoints, and so on).

- ▶ Designs

The designs of the physical layout of an organization's perimeter can have a large impact on the required surveillance and sensors that need to be in place. A good design includes Physical Asset Management requirements from the beginning.

- ▶ Policies

Policies related to the Physical Security of assets can depend on an organizational directory and organizational roles in the same way that access policies for securing IT resources do. Likewise, policies for securing physical assets are a necessary component to the overall IT security and should be included in the library of all security policies and be subject to the same review and change processes.

- ▶ Identities and Attributes

Physical asset security depends on directories of employees and their organizational roles to control access to physical assets and to manage who can use or maintain the high-value physical assets. For example, a Physical Security policy might require that only people who have completed a particular training program should be allowed to perform maintenance on a physical asset.

2.3 Conclusion

In this chapter we explained the IBM Security Blueprint in more detail by discussing the components and subcomponents of the IBM Security Blueprint. We described the subcomponents in detail and related them to the key infrastructure and security services components on which they depend. Next we look more closely at an example business.



The Network, Server and Endpoint solution pattern

In this chapter, we discuss how the IBM Security Blueprint can help you define a consistent *solution pattern* for the IBM Security Framework Network, Server and Endpoint security domain. This solution pattern can be used to isolate and better understand the functional capabilities for security solutions and to derive the elements that are required in the IT environment to establish those solutions. With this knowledge, IT security architects can feel confident when they have to design threat and vulnerability management solutions.

We will look at the following topics:

- ▶ Deriving the solution patterns for the IBM Security Framework security domains
- ▶ Examining the IBM Security Blueprint components for Network, Server and Endpoint
- ▶ Using the solution pattern for Network, Server and Endpoint planning and design

3.1 Deriving the solution patterns for the IBM Security Framework security domains

The IBM Security Blueprint can be used to derive a set of components for each security domain of the IBM Security Framework by outlining the respective security services and related security infrastructure components used by these services to address the issues and generate the values that are summarized under a given security domain. Such a set is called a *solution pattern* for a security domain.

Solution patterns can be useful for better understanding the requirements of the IT environment when designing a solution for a particular security domain. A solution pattern shows you the Foundational Security Management Components that you should consider along with the Security Services and Infrastructure on which they rely. A solution pattern can also help you better understand the internal relationships between the Security Services and Infrastructure components in the context given by a security domain.

Furthermore, a solution pattern can be used to scope projects in their early design stages due to providing a better understanding for external dependencies at a high level.

As shown in Figure 3-1, the security domains of the IBM Security Framework can be roughly mapped to the Foundational Security Management Components of the IBM Security Blueprint. It is important to understand that this mapping is not intended to be a perfect one-to-one match, but rather that the service or the services closest to a given domain provide the *main functionality* to address the issues and produce the value associated with that security domain. These *main services*, however, may require, and will at least in many cases, benefit from a combination with other services.

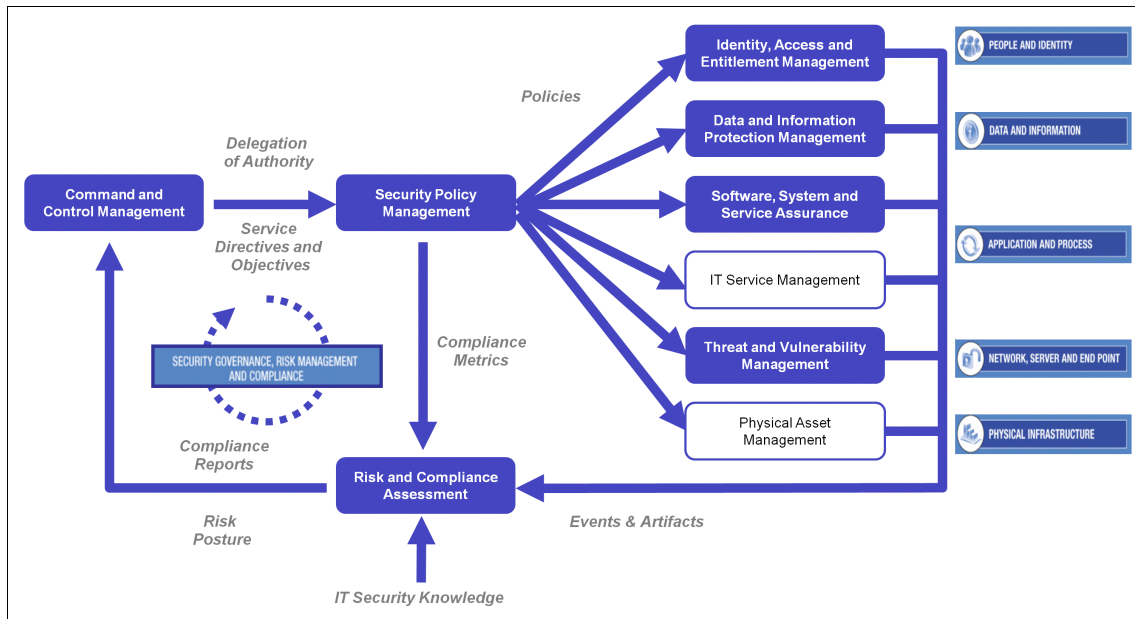


Figure 3-1 Mapping security domains to the Foundational Security Services of the IBM Security Blueprint

In the next sections, we focus on the solution pattern for the Network, Server and Endpoint security domain.

3.2 Examining the IBM Security Blueprint components for Network, Server and Endpoint

When dealing with any sort of security challenges for network, server and endpoint systems, business related issues are addressed by the Network, Server and Endpoint security domain of the IBM Security Framework, as discussed in 1.4.5, “Network, Server and Endpoint domain” on page 19.

From a functional security perspective, security issues for the Network, Server and Endpoint security domain are addressed (and potential values created) primarily by the Threat and Vulnerability Management components outlined in detail in 2.2.7, “Threat and Vulnerability Management” on page 75.

However, Threat and Vulnerability Management functionality cannot stand alone in battling Network, Server and Endpoint security concerns. When taking into account the dependencies outlined in the Foundational Security Management layer, these solutions also require Security Policy Management services to manage policies for the respective controls throughout their life cycle. This services can include:

- ▶ The analysis of potential security controls and possible corresponding technical settings on Network, Server and Endpoint devices, for example:
 - The frequency of patching servers and endpoints
 - The thresholds and triggers of security alerts
 - The administrative controls for centralized management and monitoring
- ▶ The publication, education, and awareness creation about security controls for Network, Server and Endpoint systems.
- ▶ The maintenance of review cycles and coordination of update efforts for Network, Server and Endpoint security policies.

Risk and Compliance Assessment is necessary to examine the risk posture and regulatory or internal compliance requirements for Network, Server and Endpoint security. This can include:

- ▶ The determination of business risks resulting from technical security issues (for example, based on how critical a server is for a given business process).
- ▶ The compliance posture of network controls.
- ▶ The investigation of security events to identify fraudulent or other legally relevant activities.

Command and Control Management is essential to structuring and orchestrating the various activities as a holistic entity to create the realization and ongoing maintenance of security controls for Network, Server and Endpoint operations. This can include:

- ▶ The day-to-day coordination of security activities for Network, Server and Endpoint across different departments and competencies in an IT organization.
- ▶ The management of continuity and recovery activities for Network, Servers and Endpoints during disasters and crisis situations.

Figure 3-2 highlights the components and subcomponents of the IBM Security Blueprint that have to be examined for every solution in the Network, Server and Endpoint security domain. Besides the Foundational Security Management services mentioned before, the IBM Security Blueprint enables you to determine the Security Services and Infrastructure components by reviewing the component catalogs for these Foundational Security Management services, as they can be found in 2.2, “Subcomponents” on page 34.

Each of these components can then be assessed by determining whether each particular infrastructure component is required to make a Foundational Security Management service functional so that it can address the issues or provide a prospected value associated with the particular business security domain, in this case, Network, Server and Endpoint.

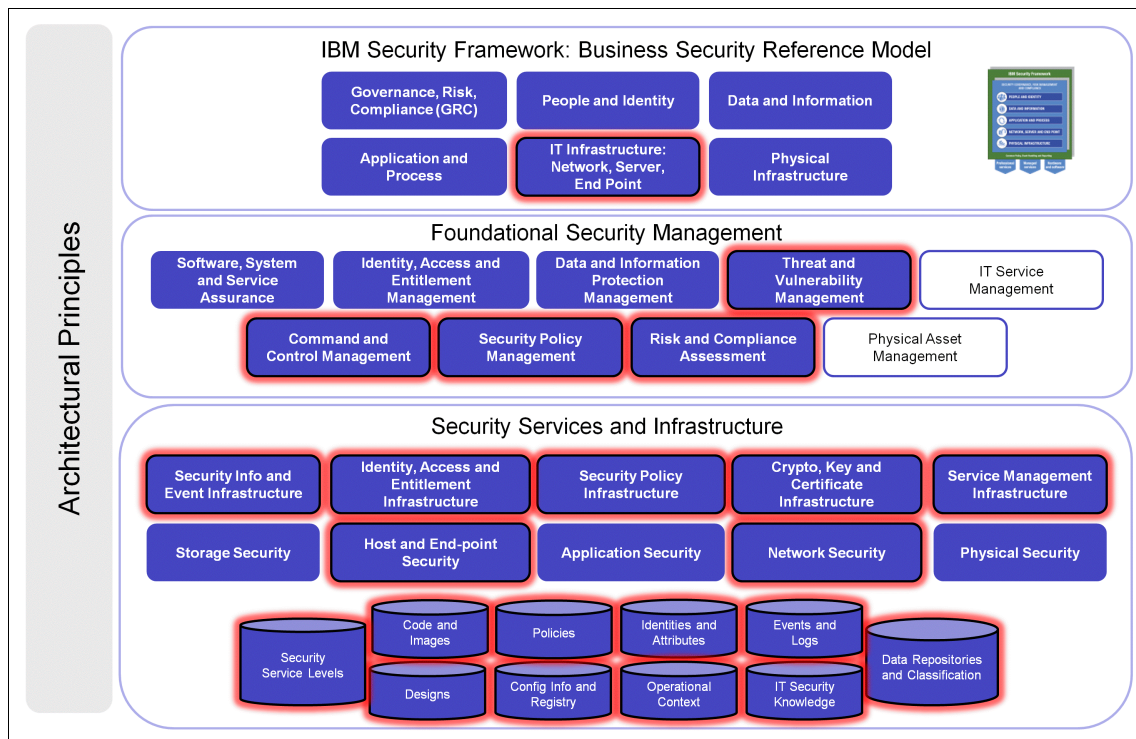


Figure 3-2 IBM Security Blueprint components for the Network, Server and End Point solution pattern

We can see in Figure 3-2 that almost all infrastructure components may be required for a Network, Server and Endpoint security solutions apart from Application Security, Storage Security, and Physical Security. The reason why those components are not included is that they are mostly covered by other domains of the IBM Information Security Framework.

Application Security is covered by the Application and Process domain, Storage Security is covered by the Data and Information domain, and Physical Security by the Physical Infrastructure domain.

However, there may be situations that require exceptions to this structure. For example, when you implement a particular project in your organization with a focus on networks, servers, or endpoints, you may have to consider the Physical Security subcomponent because your project delivery team also has to cover the physical protection for workstations. For this requirement, you have to include plans for cable locks and potentially a related master key solution. The physical separation of servers in a data center from other equipment and the restriction of physical access to a cabinet that contain the network devices that connect a floor of an office building can have similar requirements, which is why you want to consider the Physical Security subcomponent. These and other examples may also hold true for other subcomponents.

For larger organizations with a dedicated and holistic department for physical security, however, these functions may likely be provided by an established group responsible for facility and physical security management. Such a group can also use the IBM Security Blueprint to identify the required components for their needs and solutions around Physical Security, but these types of project are not covered in this book.

The component catalogs introduced in 2.2, “Subcomponents” on page 34 can be used to derive, in more detail, how the Foundational Security Management services address the issues and create the values associated with the Network, Server and Endpoint domain. This is accomplished by selecting those service subcomponents of the Foundational Security Management services that are found to be most relevant for Network, Server and Endpoint matters. We have done this task and the outcome is shown in Figure 3-3 on page 99.

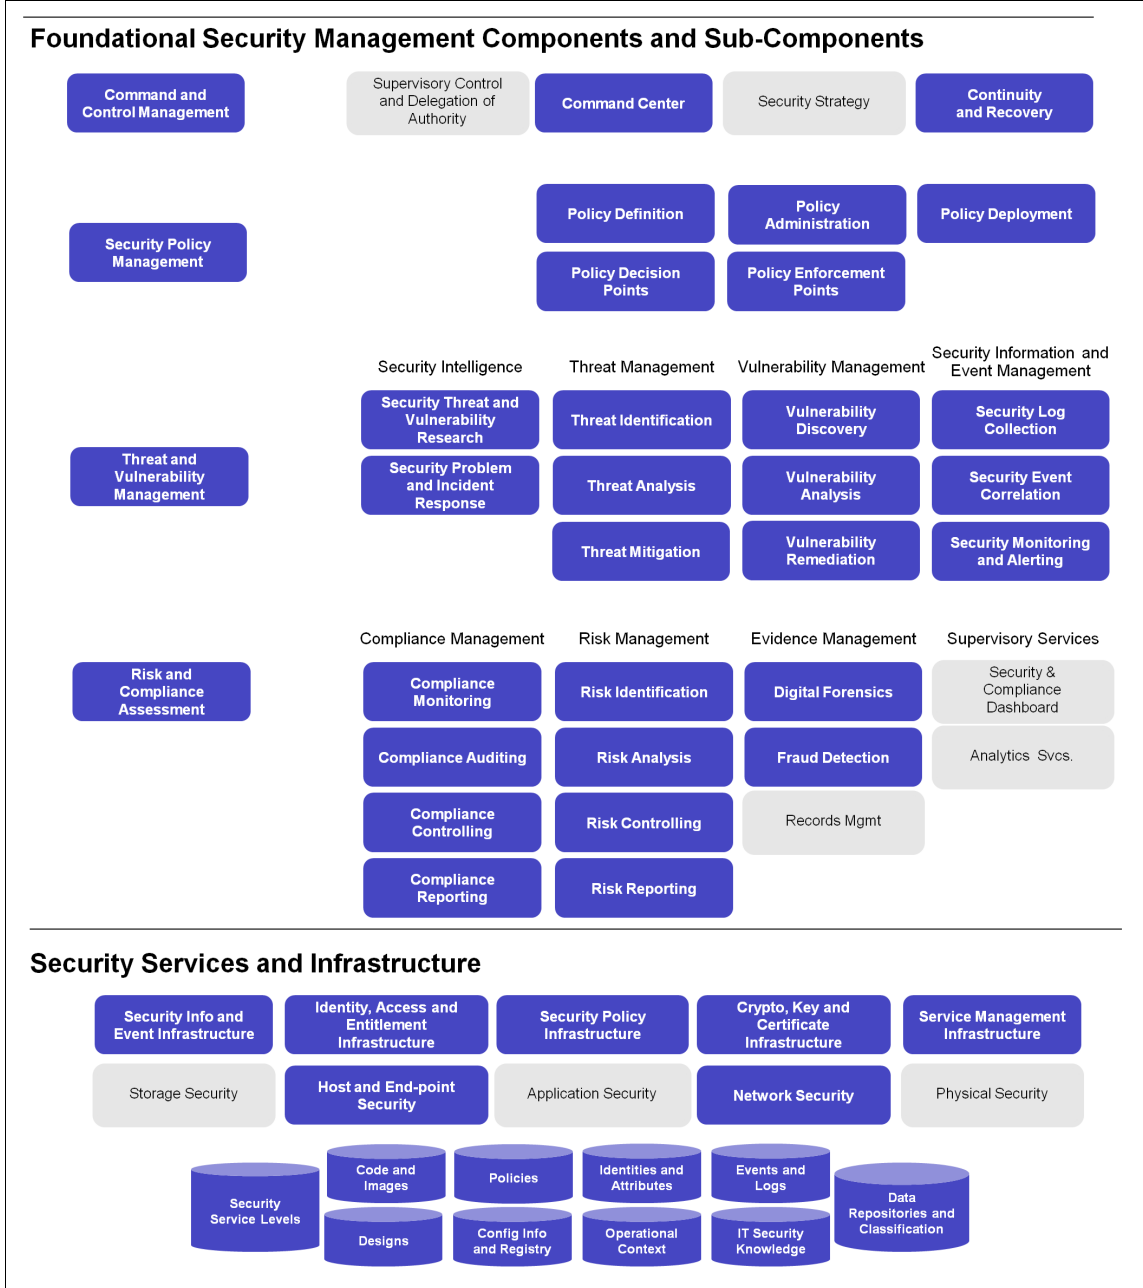


Figure 3-3 Foundational Security Management components and subcomponents

Figure 3-3 on page 99 shows the *Command Center* as the key service of Command and Control Management required for technical solutions for Network, Server and Endpoint systems and also lists (literally as a special case of the Command Center) *Continuity and Recovery*. The other two subcomponents *Security Strategy* and *Supervisory Control and Delegation of Authority* are not selected. Although they are important predecessors for all security services, the Security Strategy determines the security on a higher, less technical, level and Supervisory Control and Delegation of Authority deals, again from a business level perspective, with the establishment of authorities and delegation of decisive power and responsibility for the various facets of operations. Both these functions are more related to the business strategy than the technological level. Their essence is reflected in the policy framework of an organization, and these policies again are needed for an effective provisioning of Network, Server and Endpoint security.

From the Security Policy Management subcomponents, *Policy Definition*, *Policy Administration*, and *Policy Deployment* are required to manage the security controls and the respective technical security settings for the systems. The documentation of *Policy Decision Points* and *Policy Enforcement Points* is also necessary for the Network, Server and Endpoint domain.

At the core of the required functions for successful Network, Server and Endpoint security is Threat and Vulnerability Management with all its functional component groups of *Security Intelligence*, *Threat Management*, *Vulnerability Management*, and the technical *Security Log and Event Management*.

The subcomponents of Risk and Compliance Assessment, which are needed for Network, Server and Endpoint security, are all subcomponents of *Compliance Management*, as Network, Server and Endpoint security clearly requires *Compliance Monitoring* to track the level of adherence to security settings that match the security controls defined in the security policies for these platforms.

A holistic security management approach for Network, Server and Endpoint also includes *Compliance Auditing*. Compliance Auditing supplements Compliance Monitoring and addresses the manual processes and human resource factors of the security operations for Network, Server and Endpoint. Compliance Auditing also tracks the remediation of compliance deficiencies that have been identified. Finally, Compliance Reporting can consolidate and feed back findings of compliance monitoring, compliance auditing, and also compliance controlling activities to the different stakeholders in an organization or to external auditors.

The second key set of functional subcomponents in the Risk and Compliance Assessment area, which are required for a working security regime on Network, Server and Endpoint, are all subcomponents of *Risk Management*.

Risk Identification helps identify security events in the Network, Server and Endpoint domain, which are usually captured in the form of security event logs, and to determine the existence of security risks in the observed activities.

Risk Analysis further examines security risks of Network, Servers and Endpoints to determine the composition of these risks.

Risk Controlling is at the core of Risk Management, as it determines the actions to counter the risks. Usually these actions fall into the area of *risk mitigation*, *risk transfer*, or *risk acceptance* of a given risk.

Usually, risk mitigation, that is, the deployment of additional technical security controls, comes first to mind in the context of IT risk management, as it is the most used term when organizations act on risk. However, the other two are equally valid. A risk transfer can be achieved to a limited extent by buying insurance, or by outtasking or outsourcing scenarios. Often, for rare yet high-impact risks, risk acceptance is performed. Most of the times you find that all three options are being combined.

There is a fourth option, and it stands for the worst choice of all. Unfortunately, this choice, *risk ignorance*, is practiced much too often.

Let us put this into an analogy. Running a business is like constantly driving a car along a cliff road. Now let us assume it starts to rain.

- ▶ Risk mitigation is equivalent to the driver slowing down.
- ▶ Risk transfer is equivalent to either handing over the car to a more experienced driver or asking other vehicles to take your cargo load.
- ▶ Risk acceptance is equivalent to the driver just driving on like before.

So, what does risk ignorance mean in this case? Risk ignorance is equivalent to driving on by yourself while closing your eyes, so you do not have to see the rain.

The pity, and the limit of this analogy, is that it is harder to spot people who practice risk ignorance in your organization than it is to spot drivers with their eyes closed. But these people will wreck your business just like those drivers will wreck the car.

To conclude, *Risk Reporting* is the provision of results of aforementioned risk information processing to relevant stakeholders, like system owners or business process owners, whose processes are supported by the systems associated with the risk.

From the *Evidence Management* column, *Digital Forensics* is a component that is required. From a technical perspective, this covers required activities such as root cause analysis and securing of finely detailed evidence.

Fraud Detection is included as well, because technical analysis of threats and vulnerabilities can often lead to fraud investigations. The other Evidence Management subcomponent, *Records Management*, focuses more directly on the data, and thus is needed more when discussing the issues and values in the context of the Data and Information domain of the IBM Information Security Framework.

3.3 Using the solution pattern for Network, Server and Endpoint planning and design

The solution pattern for Network, Server and Endpoint can be used as a guiding structure throughout the development of a solution and can be useful in particular during the following stages or tasks:

- ▶ Better understand the solution requirements.
- ▶ Define and refine the solution scope.
- ▶ Review the capabilities and infrastructure of the current IT environment.
- ▶ Select software products for the new solution.
- ▶ Verify future capabilities and infrastructure of the new IT environment.

The general approach for all these activities is to take the solution pattern as a guide and carefully examine each of the elements. It is the solution pattern's goal to provide a comprehensive set of components that can help you address your requirements.

You may frequently come across an idea or concern about the need to improve your Network, Server and Endpoint security posture. At that time, however, the technical details for those new requirements are not yet understood and formulized. The solution pattern for Network, Server and Endpoint can be used to investigate these general concerns and ideas more closely. You can transform them into concrete requirements by going through each of the foundational subcomponents. You examine whether the improvements delivered by the new project, solution, or service should provide additional functionality in the area reflected by the respective functional components or whether the respective functional component may be already sufficiently covered by the existing environment and would not have to be addressed again.

The guiding question for this exercise needs to address each subcomponent listed in the solution pattern. You should ask yourself “*Does my Network, Server and Endpoint security require improvement in this functional area, and if so, what needs to be improved?*”. The result can be a list of requirements that are aligned to foundational security management components.

Best practice: At this time, you should take the opportunity to also document reasons why you might not want or need to address certain subcomponents. This can be important when you revisit your decisions at a later time, or when you have to consolidate with other departments in your organizations. This information can also be helpful in any legal disputes to prove your due diligence processes when you had to make these decisions and why you made them.

After the more granular requirements are identified, the results can optionally be further enhanced by assigning priorities to the identified requirements, which can help narrow down the solution scope. Throughout the solution development process, the resulting prioritization of requirements can also help deal with budget constraints and other scope-influencing events.

Besides structuring the requirements and scope of a future solution, the solution pattern can also be used as guidance for examining the current IT environment, in which a new solution may have to be embedded. For this, the solution pattern is used as a discussion guide, or checklist, to identify existing security-functional capabilities, including their potential limitations. In this discussion, you can also determine the existing infrastructure components and features to which the potential new solution will have to connect. The rather high level terminology of the solution pattern serves as an advantage here, because it allows you to run this exercise not only among security professionals, but to use the knowledge of experts for other technical disciplines as well as business level people.

Future elements of a new solution can also be examined with the solution pattern. Every new solution can be built with custom or pre-built software products, the underlying hardware (sometimes packaged together in form of an appliance), professional services, and the people and processes. For all these elements, the solution pattern can be used to identify their security-functional components and the provided or required infrastructure components. For example, a software product that will be used in the new solution can be examined on a high level towards the security-functional components and the infrastructure elements it requires and connects to or provides.

By using the solution pattern for this exercise, we can develop a structured mapping of the product functionality. If we have also used the solution pattern to develop the required functionality, it can be easy to determine the extent to which the software product covers the requirements and whether there are some gaps, which we then may have to close by either choosing a different product, or by accompanying the initial software with another software product and so on.

By using the solution pattern's high level of abstraction, we can quickly identify any significant gaps. But even if we want to dive deeper into the analysis of the capabilities of a product, the solution pattern helps us structure the analysis and focus on the required functional subcomponents, which we derive from the requirements and scope analysis, which also follow the solution pattern structure, and stay focused on a specific security functional subcomponent.

Using the solution pattern for such exercises can help keep the overall general picture in mind and stay focused on the intended functional coverage. Let us look at an example. During the implementation of a project, a software product, which was initially selected based on solution pattern functionality, is withdrawn and is no longer available. By examining the affected solution pattern components, you can see how the overall solution is impacted. This can help you find alternative solution offerings and get you back on track.

In the remainder of this book, we use the solution pattern for Network, Server and Endpoint to show the functional capabilities of the software and service solutions discussed in Part 2, "IBM Security Solutions for Network, Server and Endpoint". We also demonstrate that the approach outlined in this section can be used for the solution design of our practical customer scenarios in Part 3, "Business scenarios" on page 437.

3.4 Conclusion

In this chapter, we have shown how to use the IBM Security Blueprint to derive a solution pattern to outline the functional components for security management, as well as the technical components that need to be considered in Threat and Vulnerability Management solutions for the Network, Server and Endpoint security domain. We also explained how this pattern can be used as a guiding method during planning, assessing, designing, or deploying of respective security solutions.

Although we have approached the matter of Network, Server and Endpoint security from a business perspective in the first three chapters of this book, in the next chapter we explore the area of Network, Server and Endpoint security from a technical architecture perspective.



Common security architecture and network models

In this chapter, we start by briefly looking into *security omnipresence* in our daily lives and how it affects everything around us. Next, we bring this omnipresence into a business perspective by introducing an *enterprise security architecture model*. The more technical discussion begins by discussing *common network components* and a *common network model*. Finally, we map those components into some sample practical designs for a better understanding.

4.1 Security is omnipresent

As the planet is becoming *smarter and flatter* in a rhetorical way, it is also becoming more complex from the standpoint of communication and information exchange. Organizations, governments, and the general populace are all interconnected today and generating and exchanging zettabytes of data (a zettabyte is a 1 followed by 21 zeros). A large amount of that data is produced by individuals who do not have a core responsibility to secure that data, but more than 85 percent of that data ends up in organizations that have to carry that responsibility.

Across multiple industries and sectors of society, security has become a key component of our planet's vital systems. Those systems are shared and shaped by businesses, cities, government agencies, and communities. This increased complexity instills greater security risks for any organization looking to implement a smarter solution, regardless of industry. To realize the promise of future innovation, organizations have to take a more proactive approach to securing the infrastructure that powers new smart products and services.

As we have learned in the previous chapters, security in today's organizations cannot be solely regarded as a technological implementation. Every new project has to be business driven based on risk management, which has a global effect on the organization.

As shown at Figure 4-1, security is holistically present in all key segments (assets) within an organization:

- ▶ Business (process, management, governance)
- ▶ People
- ▶ Data (information)
- ▶ Assets and tools (including software tools and applications)
- ▶ IT infrastructure and communication (network)
- ▶ Customers




Figure 4-1 Holistic view of security

Every business related project that encompasses information assets (personal, organizational, or global) and the exchange of those assets must include proper risk assessment and management tasks to make information security one of the core responsibilities for everybody involved.

4.2 Enterprise Security Architecture model

With this responsibility in mind, it is obvious that every organization needs an *Enterprise Security Architecture* that follows a holistic approach to remediate critical business risks and to layout a proper security roadmap. These tasks can be accomplished by using IBM Security Framework and IBM Security Blueprint.

Because security envelops every business process and component shown in Figure 4-1 on page 107, it is imperative that security related tasks be addressed and well documented in all enterprise projects. Every IT related project is typically tied back to some business requirements. The IBM Security Framework can be used to develop a proper understanding of the business driven security issues. After those issues have been identified, a solution needs to be architected and then implemented. In the next section, we take you through a set of steps that can guide you in building an enterprise security architecture for your projects.

4.2.1 Security architecture delivery processes

When you are concerned with a standardized method for *solution design* and *solution delivery*, you need a common methodology that everyone who is involved in those tasks can follow. The following paragraphs provide you with some of the major tasks that you will encounter in your design phase:

- Project definition

The *project definition* is a document that describes the shape of the project and includes the objectives and scope, the stakeholders and proposed organization with responsibilities, and the major risks associated with the project. It is created by the project manager. An initial project definition is created after assessing the project goals and before refining the solution at a level that is appropriate for planning. It contains the reminders and the plan for the *defining* activities. The project definition is then developed during the defining activities by adding the planning framework and the plan for *planning* activities.

- Requirements matrix

The *requirements matrix* is usually delivered as a spreadsheet that captures all functional requirements that have to be delivered as part of the design and project delivery. It compiles input from multiple sources, such as a *statement of work*, business requirements based on project interviews, and the *project definitions* document. Besides covering the functional requirements as part of the project delivery, it usually also covers *non-functional requirements* (NFRs). Non-functional requirements should be organized by a number of common themes or subcategories. These include the areas of performance and capacity, availability, usability, certain security and privacy aspects (for example, disaster recovery and business continuity plans), maintainability, manageability, and flexibility.

► System context

The *system context*, usually shown as a diagram, defines and depicts the details on which the project is focused. It identifies the information and control flows that cross the system boundaries, and it assists in delineating the development team and client responsibilities. The system context highlights, at a high level, how the designed system interacts with external entities such as users, external systems, batch inputs and outputs, and external devices. The graphic (diagram) representation depicts:

- External events to which the system must respond.
- Events that the system generates that affect external entities.
- Data that the system receives from the outside world and that must be processed in some way.
- Data produced by the system and sent to the outside world.

► Architecture overview

The *architecture overview* document is one of the first design documents that provides an overview of the components that compose the architecture. The key aspects of these documents are:

- *Architectural goals* define how the system needs to respond to change over time.
- The *Architecture overview diagram* depicts the major elements of the architecture and their relationships. This diagram usually reflects network zones and how components fit into different zones.

► Architecture decisions

Architecture decisions document key decisions about the architecture and the rationale and reasons behind those decisions. It combines input from the system context, architecture overview, and requirements.

► Component model

The *component model* describes the structure of a system in terms of its software components with their individual responsibilities, interfaces, relationships, and the way the components collaborate to deliver the required functionality. The component model is the main artifact documenting the functional view of the architecture and serves as an abstraction of the design. The identified components may need to be decomposed into further component models before they meet the specifications required for the detailed design.

Component models are documented at two levels:

- The *logical level* focuses on specifying the components' responsibilities and the characteristics needed to deliver the requirements. These specifications are technology and product neutral.
- The *physical level* focuses on how to implement the components to meet the previously established specifications.

You may transform logical components into physical components via custom development, the purchase of products, or the reuse of assets.

► Operational model

The *operational model* describes the “operational” aspect of an IT system's architecture. It describes the required operational characteristics and capabilities of the IT system architecture and represents, at an architectural level, the network of computer systems and their associated peripherals, together with the systems software, middleware, and application software that they run to support the users of the system. It can be seen as a detailed *architecture overview diagram*.

► Use case model

The *use case model* presents an overview of the intended behavior of the system. It is the basis for agreement between stakeholders and the project team in regards to the intended functionality of the system that fulfills business, functional, and non-functional requirements. It also guides various tasks in the system development and test.

► Viability assessment

The *viability assessment* is an analysis of how viable the proposed system is and provides a measure of how likely it is that the system will satisfy critical requirements. It is usually a part of the *quality assurance* process.

Obviously, many security aspects have to be included in all those work products to reflect the current market trends and the concept *Secure By Design*¹. Integrating security into the initial design of any infrastructure and solution project can result in improved regulatory compliance posture, reduced risk and cost by assessing and identifying security vulnerabilities, and recommendations for prioritization and resolution.

¹ In the IBM Redguide™ *Security in Development: The IBM Secure Engineering Framework*, REDP-4641, you can learn more about secure engineering practices for software products, where we offer a description of the IBM end-to-end approach to product delivery, with security taken into account. To obtain this publication, go to <http://www.redbooks.ibm.com/abstracts/redp4641.html?open>.

As a part of security architecture related work products (such as architecture overview or component model), we talk about network segmentation. Network segmentation is the reflection of good security practices that can help isolate access to different types of data and assets by placing them into different zones.

Network segmentation establishes *security zones*, which include the segmentation of assets, users, and transactions based on both physical (network flows) and logical (access) infrastructure.

For the remainder of the chapter, we focus on the aspects of network components, network zones, and some practical designs.

4.3 Common network components

Networks are the mechanism for electronic communication between systems. The view of the network and security has changed over time. Network security has focused on hard boundaries, with limited access to and from the Internet. Now networks must provide a variety of communications inside and outside of an organization in a carefully controlled manner. There must be a balance between blocking malicious traffic and allowing traffic in a controlled manner.

Networks are the foundation for e-business transactions and Smarter Planet™ infrastructures providing communication pathways that help systems and organizations work and operate together. Networks must function in a secure manner aligned with the business context. This means the network structure must consider risks, and mitigate them through its design.

Most basic security architecture methods today are using common network models to better implement, control, and constrain the level of security for different network zones.

Network boundaries are used to isolate networking zones with different security policies. These boundaries are created to implement restrictions on the type of traffic that is allowed in a zone. For example, you can restrict access to only HTTP traffic on port 80 and HTTPS traffic on port 443 inbound from the outside to a zone of web servers. We use a firewall to allow this traffic and block all others. In its simplest case, a firewall is a device that implements a policy regarding network traffic. It creates boundaries between two or more segregated networks, and stands as a shield against unwanted penetrations into your environment. However, a firewall is not meant to be your only line of defense; it is a filter that can reduce the surface area for the threat to propagate.

One method of shielding information about the network that the firewall protects is through re-addressing the data packets that traverse the networks so that outbound traffic appears to have originated from an address associated with the firewall itself. This re-addressing is called *Network Address Translation (NAT)*, and its primary function is to hide the trusted network from untrusted networks.

Another method to protect network segments and their content is by using a network *Intrusion Prevention System (IPS)*. The IPS can complement a firewall by inspecting the content of traffic, which a firewall cannot. This method can prevent a threat from propagating from zone to zone, therefore significantly reducing the impact of the threat.

Firewalls and IPS devices can be bundled with other features, such as content filtering, Virtual Private Network (VPN) functionality, and even authentication. The next few sections describe several basic firewall and network intrusion prevention approaches.

4.3.1 Packet filter firewall

A *packet filter firewall* uses a rule set to decide what traffic is allowed and what traffic is blocked. It does this by analyzing individual network packets and matching them to a set of predefined rules. The packet filter allows or does not allow communication based on the information in the packet and the direction it is heading. Elements that are evaluated against the rules are:

- ▶ The physical network interface on which the packet arrives
- ▶ The IP address from which the data is coming
- ▶ The address to where the data is going
- ▶ The type of transport protocol being used (UDP, TCP, or ICMP)
- ▶ The source port
- ▶ The destination port

This type of firewall is simplistic. It does not look at the packets application layer data, and does not track the state of the connection. It allows access through the firewall with the least amount of inspection. Because it is simplistic, it is the fastest firewall technology available.

4.3.2 Circuit level firewall

Circuit level firewalls confirm that a packet is either a connection request or a data packet belonging to a connection. To validate the connection, the circuit level firewall examines each connection to ensure that it offers a legitimate handshake for the protocol being used. Data packets are not forwarded until the process is complete.

This type of firewall stores information as dynamic rules regarding that connection. These are in the form of a virtual state table about the session at the transport layer. All incoming packets are compared against rules on the transport layer. If the packet meets all conditions of the circuit table and rules, it is allowed.

4.3.3 Application layer firewall

Application layer firewalls examine the information in network packets, but operate at the application level. They view information as a data stream and not as a series of packets, so they are able to scan information being passed over them and ensure that the information is acceptable based on their set of rules and logic. This allows the firewall to make some intelligent decisions about what to do with packets that pass through it.

Application layer firewalls generally take the form of specialized software and proxy services, allowing no traffic directly between networks. They also have the added feature of performing logging and auditing of traffic passing through them. This enables them to communicate with an *Intrusion Detection System* (IDS), and log information regarding an attack.

4.3.4 Dynamic packet filter firewall

Also referred to as *stateful inspection*, dynamic packet filtering does not examine the contents of each packet. Instead, it compares certain key parts of the packet to a database of trusted information. Information traveling from inside the firewall to the outside is monitored for distinctive characteristics, then incoming information is compared to these characteristics. If the comparison yields a reasonable match, the information is allowed through. Otherwise, it is discarded.

The dynamic packet filter acts at the network layer, and tracks each connection negotiating all interfaces of the firewall to ensure that they are valid. It also monitors the state of the connection, and compiles the information in a state table. Because of this, filtering decisions are based not only on administrator-defined rules (as in static packet filtering), but also on context that is established by prior packets that pass through the firewall. It also has an added security measure that closes off ports until connection to the specific port is requested. This is an effective counter to port scanning.

4.3.5 Routers

A router is an interconnection device that links discrete networks and forwards packets between them. A router makes decisions on whether to forward a packet between networks based on a configuration table of routes, and addresses information in a packet. A router can be used to isolate the networks from one another, preventing the traffic on one from unnecessarily spilling over to the other. Why discuss routers within the context of firewalls? The two usually work in conjunction with each other. A solid firewall installation uses a combination of the technologies offered by routing and filtering.

Figure 4-2 outlines a basic firewall installation.




Figure 4-2 Basic Internet boundary network configuration

4.3.6 Intrusion detection and prevention

Network intrusion detection systems (NIDS) monitor network traffic for unwanted or improperly formatted traffic. This unwanted traffic is between systems in a network zone, or from the Internet into the network. Network sensors monitor the traffic in a passive mode, logging packet information based upon rules in the sensor. These rules can trigger alerts when suspicious or unwanted traffic occurs.

There are several methods of detecting intrusions, but most network IDS and IPS systems are based on signatures or heuristics. Signature based NIDS require individual rules to be constructed for types of traffic to either monitor or ignore. The rules tell the NIDS how to view traffic. Heuristic based NIDS use statistical or algorithmic techniques to determine what is normal traffic and what is suspicious. The advantage of the heuristic approach is that alerts are based on traffic patterns, and this allows for a more dynamic configuration of what is normal and what is not. Many claim this provides an advantage in day zero virus incidents, as unusual traffic activity is more likely to be detected by a heuristic based NIDS.

Network based intrusion detection and prevention systems are critical components when creating network segments. Most firewalls do not have a sophisticated ability to identify the threats in the data portion of the packet stream.

As you can see in Figure 4-3, the firewall inspects the header of the packet and compares the information in a firewall policy to determine if the packet is allowed to pass through the firewall.

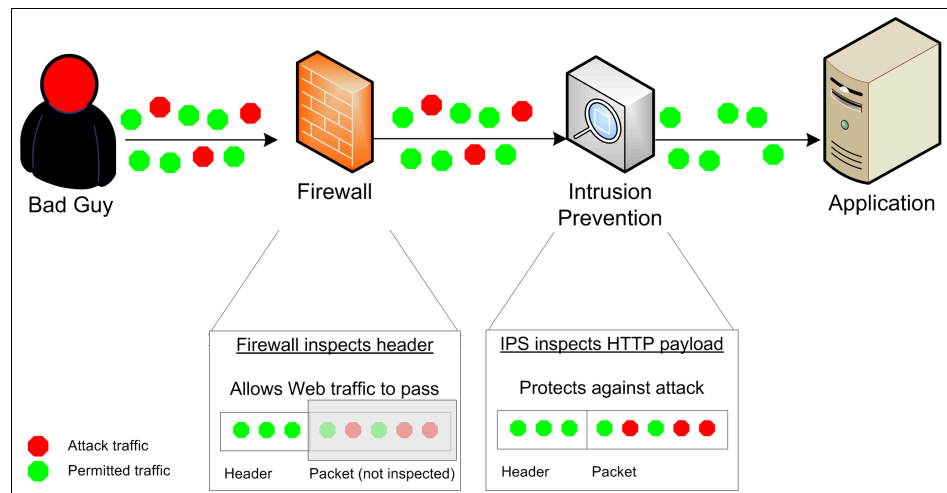


Figure 4-3 Firewalls typically inspect the header, while an IPS inspects the entire stream

The network IPS examines the stream completely, analyzing the data that is destined for the application host. Today, the majority of threats are being transported in the actual data-portion of the traffic, which is difficult for a firewall to inspect. This is why a network IPS is crucial to preventing a threat from getting into the network we need to segment.

We can regard these devices as necessary components of the information flow control and solution integrity subsystems. While intrusion detection and prevention are available for both network and hosts, for the purposes of network segmentation, we focus on the network IPS.

A growing interest is placed on extending the detection of suspicious traffic to a method to prevent it. This is the role of *network intrusion prevention systems*. These devices trigger an action to eliminate suspicious traffic when detected. Of concern is the impact if traffic viewed as suspicious is blocked, but the traffic is actually permissible.

4.4 Common network models and security domains

Every IT deployment architecture needs to be properly designed and documented. An important part of the documentation are visual representations of the logical networks and the components that are placed throughout those networks. In this section, we take a look at a network model for IT deployment architectures that introduces different *network zones* to allow the placement of IT components according to their risk and security classifications.

Figure 4-4 shows a typical network segmentation concept, which is based on many years of IT deployment and security practice across various industries and solutions.

Using a natural language, you can classify these areas as *uncontrolled*, *controlled*, *restricted*, *secured*, and *external controlled*. A client (for example, an application, a human being, or another intelligent system), uses the network to access information hosted within your network. This client can come from within your organization or from an external source (as shown in Figure 4-4). Using the concept of security domains, you can translate Figure 4-4 into something more targeted, as discussed in 4.5.1, “DMZ” on page 120.

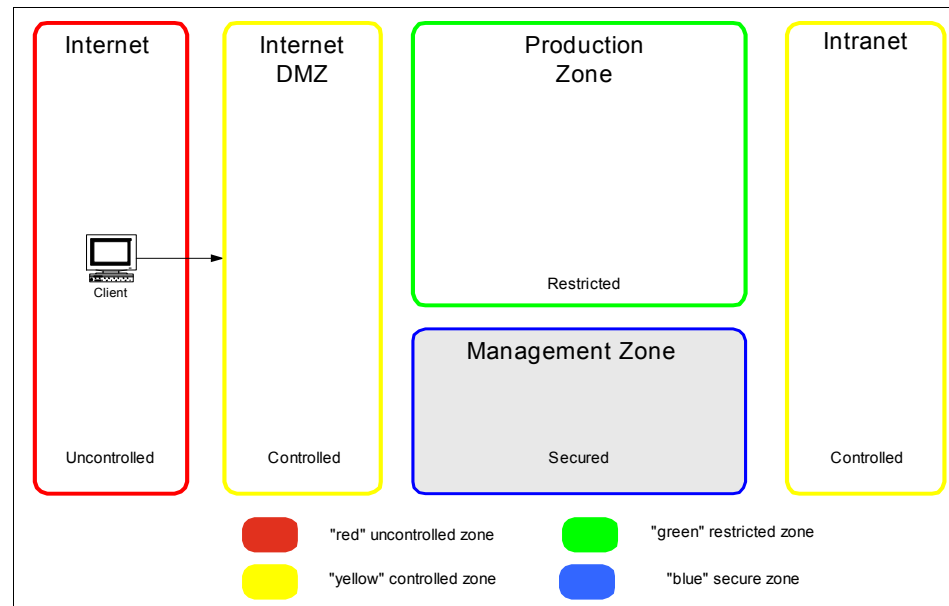


Figure 4-4 Security architecture: Network domain (zones) concepts

Traversing your networks: The breaks between each network zone indicate the use of either a firewall, network IPS, or both, that clearly delineates each perimeter from the next. Often you also find routers deployed within internal network structures.

Let us briefly explain what these domain categories stand for:

Uncontrolled	Refers to anything outside the control of an organization. Access from the uncontrolled environment to systems in the controlled zone can be through a wide number of channels.
Controlled	Restricts access between uncontrolled and restricted zones (for example, a traditional DMZ).
Restricted	Access is restricted and controlled. Only authorized individuals gain entrance, and there is no direct communication with external sources (Internet).
Secured	Access is available only to a small group of highly trusted users. Access to one secured area does not necessarily give access to another.
External controlled	An external zone in which data is stored by business partners external to the systems, where there is limited trust in the protection of data (for example, credit reporting agencies, banks, and government agencies).

Designing your IT environment in this manner enables internal users to see out, but prohibits external users to see anything inside your premises. The external users' access is restricted. The benefits of constructing security domains this way are:

- ▶ They are clear and efficient.
- ▶ They are easy to explain.
- ▶ They are easy to work with.
- ▶ They provide a complete design and implementation view, enabling you to avoid errors.
- ▶ Fewer errors mean a lower risk of exposure and loss.
- ▶ Each network deployment model can use any number of network zones.

A proper risk management strategy plays a big part in designing a secure solution, but so does security assurance. If we assess the risks for our systems, we must also ensure that we assign countermeasures for those risks providing assurance for the correctness and effectiveness of the security solution.

These network domain, or network zone, models are being used throughout this book. Figure 4-6 on page 121, and Figure 4-7 on page 122 show clearly marked firewalls to help you become familiar and comfortable with the placement of components and the overall domain concept.

4.4.1 Network zones

We have to consider four types of network zones and their transport classifications in our discussion:

1. Uncontrolled (the Internet)
2. Controlled (an Internet-facing DMZ and the intranet)
3. Restricted (a production network)
4. Secure (a management network)

Internet (uncontrolled zone)

The Internet is a global network (a network of networks) connecting millions of computers. It cannot be controlled, or have any components in it.

Internet DMZ (controlled zone)

The Internet DMZ is generally a controlled zone that contains components with which clients might directly communicate. It provides a buffer between the uncontrolled Internet and internal networks. Because this DMZ is typically bounded by two firewalls, there is an opportunity to control traffic at multiple levels:

- ▶ Incoming traffic from the Internet to hosts in the DMZ
- ▶ Outgoing traffic from hosts in the DMZ to the Internet
- ▶ Incoming traffic from internal networks to hosts in the DMZ
- ▶ Outgoing traffic from hosts in the DMZ to internal networks

The transport between a controlled and an uncontrolled zone is classified as *public*. The transport between a controlled and another controlled, or a restricted zone, is classified as *managed*.

Production zone (restricted zone)

One or more network zones might be designated as *restricted*, that is, they support functions to which access must be strictly controlled, and of course, direct access from an uncontrolled network is not permitted. As with an Internet DMZ, a restricted network is typically bounded by one or more firewalls and incoming/outgoing traffic might be filtered as appropriate.

The transport between a restricted and a controlled zone is classified as *managed*. The transport between a restricted and a secured zone is classified as *trusted*.

Intranet

Like the Internet DMZ, the corporate intranet is generally a *controlled zone* that contains components with which clients might directly communicate. It provides a *buffer* to the internal networks.

Management zone

One or more network zones might be designated as a *secured zone*. Access is only available to a small group of authorized staff. Access into one area does not necessarily give you access to another secured area.

The transport into a secured zone is classified as *trusted*.

Other networks

Keep in mind that the network examples we use do not necessarily include all possible situations. There are organizations that extensively segment functions into various networks. However, in general, the principles discussed here might be translated easily into appropriate architectures for such environments.

Placement of various data components within network zones is both a reflection of the security requirements in play, and a choice based on an existing, or planned network infrastructure, and levels of trust among the computing components within the organization. Requirement issues are often complex, especially with regard to the specific behavior of certain applications. With a bit of knowledge about the organizations network environment and its security policies, reasonable component placements are usually easy to identify.

Figure 4-5 summarizes general component-type relationships and the transport classifications to the network zones discussed above.

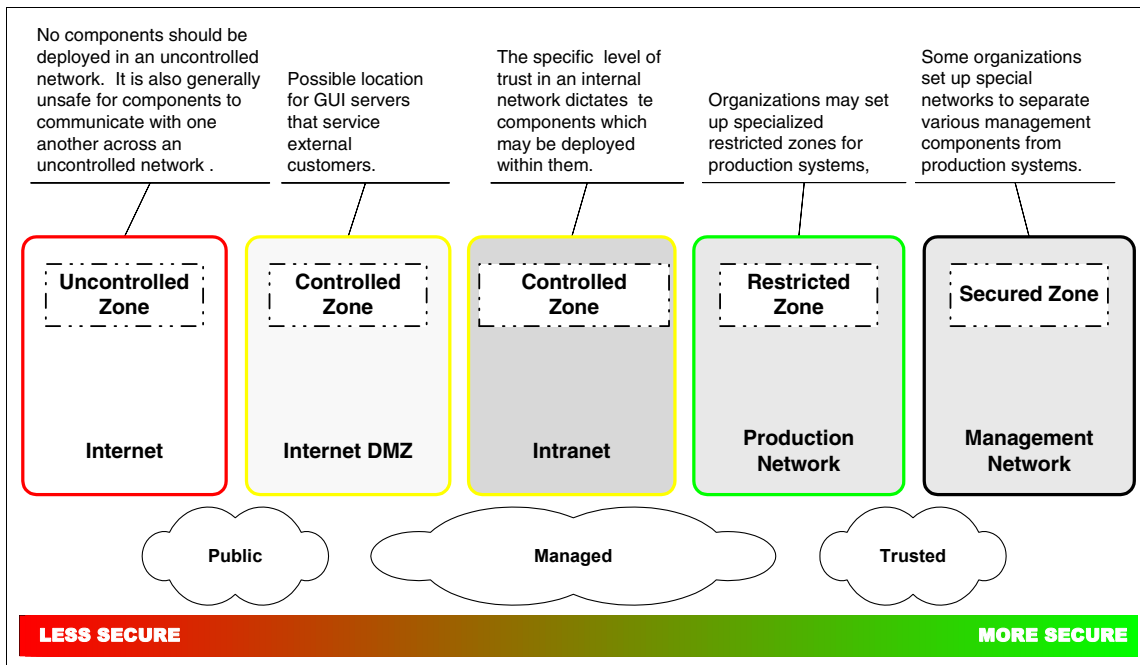


Figure 4-5 Graphic representation of network zones, transport classifications, and their level of trust

4.5 Practical designs

In this section, we describe several use cases for all zones using practical implementation scenarios. We start with a classical DMZ example and show how it delineates an organization's internal network from the Internet. Then we take a closer look at the typical intranet and production zones inside an organization and finish with the positioning of a management zone.

4.5.1 DMZ

The DMZ, or outermost perimeter network, is the separation point between your data or information (what you control) and the Internet (what you do not control). In physical terms, the DMZ connects to the Internet by means of a router, which is used to separate your network from your network provider, typically your Internet Service Provider (ISP). Here you exchange information with limited, calculated risk.

The classic approach in creating a DMZ involves adding firewalls for extra layers of security. Firewalls are often used in multi-machine systems to protect the resources that are placed within the private network, such as business applications and sensitive information. A wide variety of implementation topologies can be appropriate for a DMZ; however, the basic models usually look something like the layout in Figure 4-6.

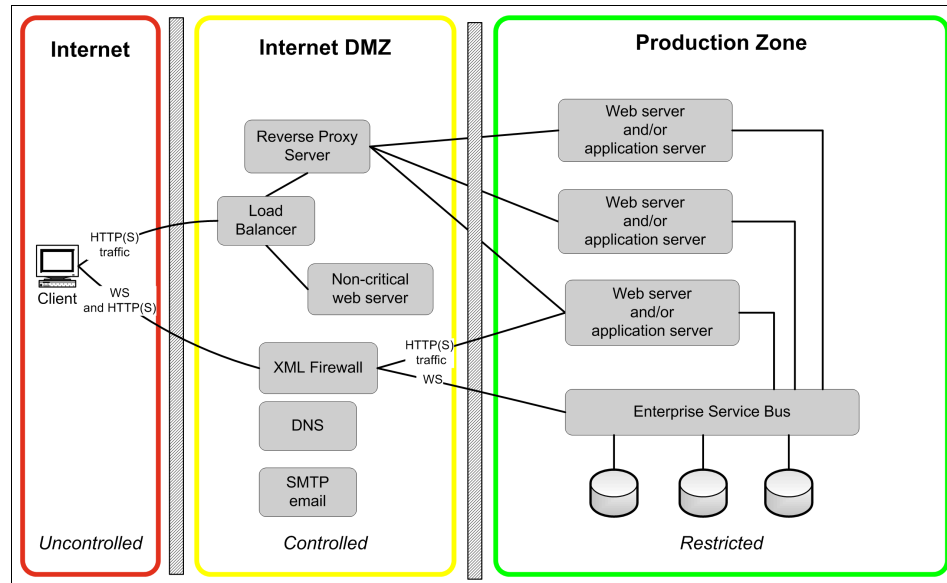


Figure 4-6 Basic DMZ design

This design approach allows for the separation of the presentation material on the non-critical web server and the business critical web application and web services. The publicly accessible read-only content can be located in the DMZ while the rest of the critical data and applications are hidden behind more strictly limited firewalls in the production zone of the network. The infrastructure allows secure transactions and processing in stages, reducing the demands on systems. The “bridge”, or proxy devices between production and client, are placed in the Internet DMZ. Even though the picture only shows the most utilized traffic schemes in the Internet DMZ, such as HTTP or HTTPS and web services (WS), some other services typically reside in this zone as well, such as DNS and email servers.

4.5.2 Intranet

Most firewalls and security schemes are built to segregate the Internet from the internal networks. However, in some situations, you might want to protect parts of the internal network from other areas of your internal network. It makes sense that not everyone needs access to the same services, information, or security protection. Figure 4-7 shows the segregation of the intranet client from the production environment. Some parts of your organization need to be more secure than others, such as demonstration or test networks (where there are often people from outside of the organization present), Human Resources data, development projects, financial data, and so on.

Use of internal facing network IPS: Even though the prime location of a network IPS is behind your Internet DMZ firewall, a network IPS can also address the need for internal network segmentation. For example, if you use a network IPS between the intranet and the production zone, a threat can be contained within the zone it initially affected. This means that if a worm infects systems in the intranet, the production zone is not affected because the worm cannot pass through the network IPS. This also provides the added benefit of requiring far fewer policy rules to manage.

Adding the additional security of another reverse proxy and XML firewalls to the network gives you central manageability of the internal access as well. In most cases, it is sufficient to separate the internal user network from the production zone(s) by using firewalls, reverse proxies, and ESB/XML firewalls.

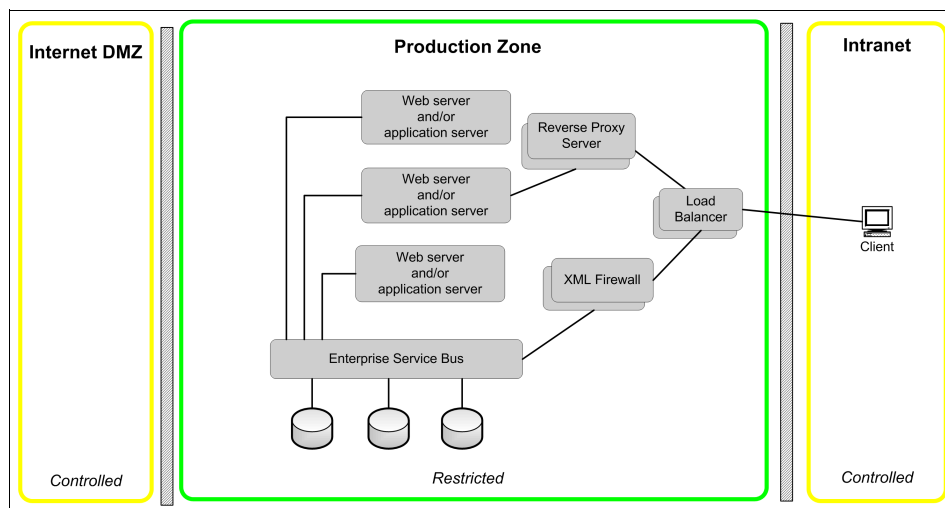


Figure 4-7 Segregating the intranet client

Let us take this concept one step further. In Figure 4-8, we add an additional zone of protection. By moving the central security management components into their own network zone, which is physically and virtually secure, you can create a zone where security administration is performed, and all of the necessary data is contained within that zone.

You can establish this type of network segregation for various reasons. You may want to create another protected area for Human Resources, where all applications and data pertaining your employees are contained inside that specific network, with access granted only as needed. Be careful when you apply this type of design; separate the things that absolutely must be protected. Keep your solution straightforward and easily scalable for future growth.

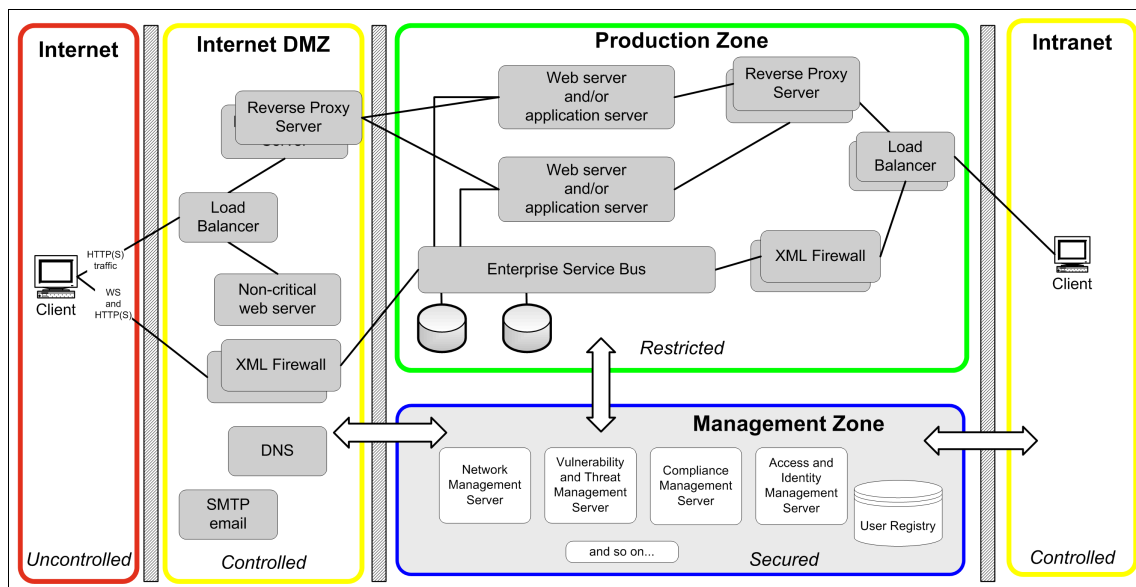


Figure 4-8 Management zone, high availability, and load balancing

4.6 Additional components

The previous discussion in 4.5, “Practical designs” on page 120 has introduced system components into our discussion. Most of these components are going to be used throughout the remainder of this book, so let us spend a few pages explaining what those components are.

Web server

A *web server* is a server that processes HTTP and HTTPS requests. Those requests can be intended for content stored on, or developed on the server itself. The web server can also act as an intermediary, where it interacts with incoming client web browser requests to pass on those requests into the web application server infrastructure for further processing. When results are returned, the web server displays them back to the client web browser.

Web application servers

A *web application server*, or application server, is a server that is executing an application. This application can be coupled with a web server on the same system, or it can receive requests from a separate web server. The application server and web server, when separate, might exist in separate network zones.

Web Proxy and Reverse Web Proxy

This component is usually deployed as a device (for example, in form of an appliance) or a software solution that usually controls the HTTP(s) traffic between the uncontrolled zone (Internet) and the organization's network. It can also be deployed on the internal networks to control access to production systems for internal users. Besides controlling the traffic, it is most often used to apply authentication and authorization rules to enforce security policies and controls.

Portal

A *portal* represents a way to provide access to a variety of applications from a single web location. The portal represents a single interface to the user, making the transition to various locations, or applications, seamless and transparent.

Back-end systems

The *back-end systems* refer to the part of the overall IT system that actually processes the requests and provides the information. This often includes database and mainframe systems that are located in the deeply protected production zone.

Database

A *database* is a collection of data stored for use by applications. The data might or might not be related. Database servers are typically part of the *back-end systems*.

Messaging services

The *messaging services* and messages can take many forms. At the basic level, a message consists of data sent between two devices. This can take several forms, such as web services, email, text messages in a wireless environment, and so on. Messaging services deal with the transport and delivery of these messages. For our purposes, the messaging service enables communications between devices.

Service-oriented architecture

A *service-oriented architecture* (SOA) reflects distributed services that communicate with each other to meet service requirements. These services are *orchestrated* to process data and data requests. Each service operates independently with its own state and context. Each service has a clearly defined method to use. Most service-oriented architectures involve web services using Simple Object Access Protocol (SOAP) and Web Services Description Language (WSDL). SOAP is a method for exchanging XML messages over the HTTP protocol. WSDL is an XML method for describing available web services.

XML firewall

The XML firewall protects against a new class of XML-based threats introduced with the deployment of web services and SOA architectures. It is usually designed to process generic XML requests and responses transmitted over HTTP or HTTPS. Although the design of the XML firewall is to process XML documents of all types, including SOAP-formatted messages, it can also accept unprocessed (text/binary) documents. Through the processing policy, the XML firewall can apply all of the various processing actions to the request and response message, regardless of format. Processing can include transformations, schema validation, logging, and cryptographic operations.

Additional reading: To learn more about XML threats, go to:

http://www.ibm.com/developerworks/websphere/techjournal/0603_col_hines/0603_col_hines.html

4.7 Conclusion

In the previous four chapters, we have seen that security has evolved into a major consideration about the way that business and information technology systems are designed, built, operated, and managed. The need to be able to integrate security into discussions with business functions and operations is more relevant than ever.

To help you with your security challenges, IBM has created a bridge to address the communication gap between the business and the technical perspectives of security to enable simplification of thought and process. The IBM Security Framework can help you translate the business view, and the IBM Security Blueprint describes the technology landscape view. In concert with a methodical design approach for your network zones, these services can help bring together the experiences that we gained from working with many clients to build a comprehensive IT solution architecture view that includes security concepts from the beginning, and not just as an afterthought.

But building a secure IT system is not enough; keeping it functional, and constantly testing, improving, and reviewing it with your team of professional security, network, and development professionals is mandatory.

For every security related measure or process that should be put in place, you need to ask the question: “Is it more cost effective to protect against or mitigate a business risks, rather than react to the result of a compromise of your organization?”

In the following chapter, we take a closer look at the foundation to provide comprehensive threat and vulnerability management.



Threat and vulnerability management

In this chapter, we focus on threat and vulnerability management for networks, servers and endpoints by exploring the current landscape of malware and the important concept of the advanced persistent threat. We then take a closer look at how to deal with threat and vulnerability management in an organized way.

Some of the major challenges that we are facing in today's world of IT security in the Network, Server and Endpoint domain are:

- ▶ Increasing number and sophistication of threats.
Organizations face more than just viruses and worms now. You have to be able to defend against and stay ahead of a great variety of threats rather than just respond to intrusions.
- ▶ IT security resources are stretched thin.
Today, threats seem to evolve faster than IT budgets and resources can keep up with. Every organization needs an efficient, integrated approach to threat and vulnerability management.
- ▶ Intrusions and malicious disruptions impact your bottom line in both customer confidence and business productivity.
Security breaches can destroy your brand image and impact your critical business processes, both of which can cost you big dollars.

At the bottom line, effective threat and vulnerability management processes need to be proactive rather than reactive, preventing problems rather than responding to them. To be efficient and effective, organizations need to address prevention, detection, and compliance in an integrated way.

5.1 Security concepts and terminology

The term *security* has many definitions and is being used in many situations. In this section, we focus on defining the major terms that we use throughout the book.

Security represents almost an industry in itself that is focused on confidentiality, integrity, and availability of information. *Information security* stands for a broader spectrum than pure IT security, because it can include aspects that do not have a direct relationship with technology, such as physical security, or an organization's security policies at the business level.

Information security has to deal with *vulnerabilities* and *exploits* to properly address threats to the organization and its information, assets, and networks.

A *threat* can be defined as events, people, or forces (*threat agents*) that can pose a risk to our assets by exploiting a vulnerability.

A *vulnerability* represents a weakness in the systems. Vulnerabilities come from deficiencies in legitimate code that is running on internal computer systems, or a system misconfiguration that can lead to an unexpected outcome. For example, SQL injection vulnerabilities are well known for being easily exploited to gain knowledge of internal database structure and contents.

A well known vulnerability category is a *software bug*, a cozy name for an application that malfunctions due to a programming mistake or error. Other common vulnerabilities in relation to applications is the misconfiguration or lack of properly implemented access control.

An *exploit* is the result of a vulnerability. It can be a piece of software or a command that takes advantage of the vulnerability. Based on an exploit, we experience different forms of attacks, for example, buffer overflows (BOFs), Denial of Services (DoS), worms, and so on. An exploit can be revealed by using a *signature*, and it can be prevented from propagating by blocking it. You can also apply a proper software patch at the target system to remediate the vulnerability.

Thus, a *signature* is a piece of information that describes a specific attack pattern. Signatures are used in IT security devices, such as a network intrusion prevention or detection systems, to detect and hopefully block the attack.

Exploits that do not yet have defined signatures or cannot be remediated by a software patch are called *zero-day exploits*. They typically attack undisclosed or unknown vulnerabilities.

Fast attack patterns: It is generally known that most antivirus vendors require between 7 to 30 hours after revealing an exploit to reverse-engineer it, create a signature to stop the attack, and send out the update to their customers. However, exploits such as Slammer^a are known to propagate worldwide in just 15 minutes. The damage is already done before an antivirus signature can be constructed. This is a typical example of a reactive approach to security that does not help with new and undisclosed vulnerabilities and the zero-day effect.

- a. To find out more about the Slammer worm, go to:
<http://www.wired.com/wired/archive/11.07/slammer.html>

5.2 Malware

Many of today's IT related attacks are implemented by developing and using malicious software, which is also called *malware*.

Malware stems from programs, scripts, or macros that can execute on almost any computer, and are malicious in nature. This category of threat is often subdivided into viruses, worms, and trojan horses.

A *virus* is code attached to or contained within a legitimate program or document. Self-propagating code is often designated as a *worm*.

Trojan horses (also called *trojans*) are old threats now returning to the forefront of IT. A trojan is a piece of code that uses *trickery* to get people to run it for a visibly legitimate purpose, but in reality the code hides its intended malicious behavior, which is unknown to the user. A trojan might perform key logging or password stealing activities. Because the motives for hacking have shifted from fame and satisfaction to financial or political profit, trojans are becoming a more and more significant threat vector. As of 2006, trojans represent the vast majority of malicious code (75 percent). Stealthy trojans might not even replicate; they are intended to steal data, or gain access to systems for future exploitation. Examples include keyloggers and password stealers that can enable financial profit through inappropriate access to accounts.

Malware can contain many components, and its categorization is subdivided according to the component's purpose (password stealers, keyboard loggers, botnets, droppers, and so on). A variety of stealth technology can be deployed to keep malware installed without detection (for example, rootkits).

Some of those common and destructive types of malware include:

► *Designer malware*

Designer malware is a piece of malcode written to infect or compromise either one or a small number of organizations with similar profiles. For example, designer malware can be a trojan horse written specifically for a single bank.

Threats using designer malware are targeted and specific. Targeted attacks and designer malware take a laser-focused approach on which organization to infect, and at the most simplistic level might target a single company or user population. In the past, antivirus vendors always prioritize threats by the *total number of infected systems*. As a more targeted attack mode, designer malware takes advantage of the old view of risk and stays under the radar of antivirus systems.

It is possible to develop antivirus signatures for designer attacks. However, attackers have come to understand the traditional responses to virus outbreaks and have crafted attacks that carefully avoid the trigger points that start the typical response. When the attack does not propagate beyond a small user community, it greatly decreases its chance of being detected at all.

Although most modern hackers eschew headlines in favor of profits, designer malware is responsible for several notable attacks. In Israel, a trojan horse attack conducted industrial espionage, and remained undetected for 18 months. This attack directly mirrors the trend of new attacks to fly under the radar of existing protection, and steal data for as long as possible before being found. In the Israeli incident, intellectual property was stolen during 18 months of infection.

The other recent example is *Stuxnet*, a trojan discovered in June 2010 by the antivirus company VirusBlokAda. It is one of the most sophisticated malware ever written. It targets industrial control systems and can modify code on programmable logic controllers (reprogram PLCs) that drive industrial processes. It is also the first malware that included a PLC rootkit.

This again demonstrated that traditional antivirus software can be totally ineffective against malware that exploits undisclosed vulnerabilities until a sample is discovered and a signature can be developed and distributed.

With millions of dollars invested in proprietary research, the biotech industry is another target of designer malware. Imagine the value of stealing the recipe for the next wonder drug. Two biotech firms have been infected with designer malware specifically targeted to steal research secrets for new projects. Designer malware has the potential to steal research findings and trade secrets, undetected, and in a relatively short period of time.

► *Ransomware*

Ransomware is malware that executes on an infiltrated computer system and packs important files into an encrypted archive and deletes the original files, thereby making access to the source information impossible unless a ransom is paid. More advanced ransomware scenarios now use multiple forms of user manipulation and extortion.

Ransomware is a growing and significant trend dealing with data, file, and user manipulation. With ransomware, attackers encrypt a user's documents and force the user to pay a ransom to regain access to the files. After paying the ransom, the user is given the password to unlock the files. Typically, users pay the ransom by visiting a website devised by the hacker and making a *purchase* of some high-priced product. Ransomware attacks also employ fear and embarrassment by telling victims the ransomware is caused by visiting inappropriate websites, or from storing pornography on their computers. Whether these accusations are true or false, such ransomware tactics can prevent users from working with security teams to cure the problem. New threats like ransomware employ technology, as well as engaging the user, which escalates damage beyond traditional Internet worm outbreaks. Certain ransomware uses stealth tactics that can cause code to self-destruct after encrypting a user's files. This makes unlocking the files even more difficult without dealing with the attacker.

► *Rootkits*

With the ability to make malware completely invisible to operating system (OS) and antivirus signature scans, rootkits can be combined with multiple types of malicious code to enter enterprise systems undetected and launch multifaceted attacks.

The rootkit is one of the most significant threats in practice today due to its stealthy nature and its ability to work with other malware. Rootkits help make malware invisible to signature antivirus scans. A rootkit is simply a *shielding technology* that can be used by any type of malware. By insinuating itself into the operating system of the compromised system, it can effectively prevent detection of whatever elements of the attack it wants to hide. Basic requests, like asking for a list of all files in a directory might be unreliable because the rootkit might hide files in the directory.

Dealing with rootkits can be a little like a game of hide-and-seek. If you watch the person hide, you have a much better chance of finding him. If the person hides and you did not see where, you might never find the person. Using behavior-based protection technology can help to identify rootkits before they can establish themselves. After the rootkit hides, it might be too late and damage can be irreversible.

Many firms attempt to clean up the rootkit after infection. However, best practices suggest that reimaging is preferable to restoring the system. Even if the time is taken to restore the system, the real damage to the enterprise is already done. If the rootkit enabled the theft of strategic corporate data or intellectual property, the enterprise cannot retrieve information that becomes public or is revealed to their competition.

Besides malware, computer crime is focused on other types of attacks, such as denial of service, social engineering, phishing and spear phishing, and so on. Let us look a bit closer at one of the most prominent of these types of attacks.

5.3 Denial-of-service (DoS)

Denial-of-service originates from external users or systems attacking a systems infrastructure with the general idea to disrupt the operation of the system. There are various forms of denial-of-service attacks. One is the vulnerability denial-of-service. There are vulnerabilities that might not be able to exploit remote code execution, but can crash the system. An attacker can crash a computer by sending a single packet to the vulnerable host.

More common are denial-of-service disruptions that come from generating a volume of traffic that overwhelms a network or host computer in the network. DNS servers are particularly vulnerable when dealing with malformed DNS requests. If an attacker can find a packet that causes a lot of cycles to be spent by the host computer, then a flood of these packets to the host can cause a denial-of-service. Bandwidth denial-of-service attacks seek to exhaust the network capacity by flooding the network with traffic. Often these attacks are mounted from thousands of different host computers (distributed denial-of-service), and usually the computers that are attacking are compromised with bot-net malware installed on the machines

5.4 Advance Persistent Threat (APT)

The term *Advanced Persistent Threat* (APT) originated in U.S. Government circles. It refers to a variety of different groups from different nation states that attack computer networks to steal intelligence information, as opposed to groups with a more direct financial motivation, such as those who target caches of credit card numbers.

The word *advanced* is used because APT groups are using exploits for unreported vulnerabilities (zero-day). The tools are advanced, custom malware that is not detected by antivirus products, and they coordinated attacks using a variety of vectors.

The word *persistent* is used to characterize the capacity that APT groups have for maintaining access to and control of computer networks even when the network operators are aware of their presence and are taking active steps to combat them. Also, APT groups are patient, as they slowly develop access to the information they want while staying below an activity threshold that would attract attention. So the attack can potentially last for months or years.

And it is a *threat*, as APT groups are dedicated to the target. The attacks are not random; they are “out to get you”, targeting at specific individuals and groups within an organization, and aimed at compromising confidential information.

The concepts of APT are:

- Reconnaissance

Includes identification of a target and method of compromise. They use a lot of investigation and information collection about the target before they execute the attack.

- Social engineering

This most commonly comes in the form of spear-phishing (email or instant messaging that appears to come from a known trusted source). The message typically contains a malicious payload or a link to a web page that has malicious code.

- Use of zero-day tools

Attacks involve exploitation of never-before-seen vulnerabilities discovered by the attackers. Not all malware in APT cases is undetectable, but the majority of malware used during the initial compromise is custom.

- Covert

The attackers remain patient and attempt to conceal their activity by masquerading as normal users. Attackers attempt to cover their actions by using legitimate accounts and protocols when possible.

- Privilege escalation and lateralization

Most often the attackers attempt to utilize a current account and obtain any information they can with those privileges. Some APT cases have involved the creation of new accounts with administrative privilege.

- Adaptive

The attackers observe remedial actions and adjust accordingly. They use their least sophisticated attacks first.

- Persistence

Attackers are patient and watch targets for long periods of time. Attackers install multiple backdoors to ensure continued access to the target network.

The level of sophistication of attack techniques seen in APT cases is often directly proportional to the level of sophistication of the capabilities of the people defending a particular network. What all sophisticated, targeted attacks have in common is that the first step for the attackers is *reconnaissance*. Often, an initial target is not always the true target. Although this may include the traditional network probing and scanning activities that we associate with computer intrusions, sophisticated attackers think outside of that box.

There is a wealth of information available on the Internet regarding many people working in the business world. We publish profiles on personal and professional social networking sites, we send out status updates that indicate where we are traveling, we engage in online forums relevant to our jobs, we talk at public conferences, we write articles and papers, we take news media interviews, and in doing all of these things we leave a large number of bread crumbs that malicious persons can use to reconstruct not just a picture of our own personal lives, but of the organizations that we work for and how we fit into them.

Sophisticated attackers use this public information to develop a complete picture of a targeted organization; who works there, what they do, and who they report to within the organization. This picture enables them to identify the particular individuals who may have access to the kind of information that they seek. Those individuals are targeted with various kinds of *social engineering attacks* intended to trick them into running a malicious exploit. We can say that social engineering exploits the “bugs” in the human brain and behavior that will help the attacker gain control of the victim’s workstation. From that point, all of the victim’s work and communications become an open book. These attacks often involve

malformed documents or web pages that target zero-day vulnerabilities with obfuscated exploits.

Spear phishing is a combination of phishing and social engineering that targets a single person or a single group of people. Spear phishing is hyper-focused to lend added credibility to the attack. Spear phishing combines the standard phishing attack with additional social engineering techniques to build super targeted attacks. Spear phishing is used heavily in state sponsored attacks, and attacks against financial institutions. The attacker takes advantage of personal or public information about individuals to customize an email that appears to come from a legitimate source and tricks people into responding with personal information such as user names and passwords.

In the following sample spear phishing scenario, John Smith's name and professional contact information is published in an industry magazine based on his recent promotion. A spear phishing attacker uses Smith's information to send a spoofed, but official looking, email to Smith posing as a professional service and requesting that he activate his new complimentary account. In responding, Smith inadvertently allows the attacker to install a trojan horse or backdoor on his computer.

That is one of the reasons why Adobe® PDF and Microsoft® office product exploits and attacks are rising. Focus is always on the weakest link in the security chain, that is, the users. By exploiting the vulnerabilities of those vastly used tools in combination with phishing, the attackers are able to insert different malware on the employee's workstations and use them for further attacks.

The attack might come as an email, addressed from a business partner or colleague, with a malicious attachment that sounds directly relevant to the victim's job function. It might be a link to a juicy document that is hosted on a competitor's website, or perhaps a USB token handed to the victim at a trade show with an interesting presentation. The custom malware that is installed by the exploit uses covert channels to communicate over the network without being noticed. After the attackers have their malware running on one victim's machine, they often try to spread their control to other systems in the targeted network. They also try to exploit business relationships to use their control over one company's network to break into others. For network security professionals in the private sector, the line between intelligence-related APT activity and financially motivated attacks is blurry at best.

Power plants have been attacked by state-sponsored cyber warriors as well as criminal groups who are simply interested in blackmail. The same sort of sophisticated spear phishing attacks that have been used to target government strategists have also been directed at executives in financial institutions who have access to funds transfer systems.

Social engineering samples: Email samples of APT social engineering attacks can be found at <http://contagiodump.blogspot.com/>.

5.4.1 Preventing Advance Persistent Threat attacks

Deploying proper sophisticated security protection mechanisms (such as state of the art identity and access management (IAM) systems, email antispam filters, or the latest intrusion prevention system (IPS)) is just half of the story. Education of employees is the other important factor.

One of the most effective countermeasures that you can employ to combat these threats is to enlist your people. If you can identify the people who are most at risk for this kind of attack in your organization, and you sit down with them and explain the nature of the threat and how it works, they can become your first line of defense. They can report suspicious emails to you. After you have received a sample of an exploit being used by these attackers, you have got a foothold on the problem. You may be able to identify other targeted victims, identify malware command and control patterns, and begin to unravel the infestation.

5.5 Threat management

Threat management is the process of identifying, understanding, and fighting threats to network infrastructure (including wireless networks), hosts, and end points. With increased market focus on cloud computing and virtualization, those security issues are prevalent in virtual environments as well. You can read more about security in virtualized environments in Chapter 10, “Virtual server security solutions” on page 337.

In addition, we talk about Cloud computing related threats and vulnerabilities in 11.3, “Cloud Security Services” on page 417.

As we introduced in “Threat Management” on page 77, the threat management discipline consists of the following categories:

- ▶ Threat identification
- ▶ Threat analysis
- ▶ Threat mitigation

Threat mitigation describes the ability to reduce risk by identifying and preventing malicious attacks from being successful in your network, on your hosts, and compromising your entire IT environment.

In many cases, threat mitigation is often overlooked as part of the necessary security infrastructure. One reason it is overlooked is because many people assume that the security, which is included in applications, operating systems, and traditional network infrastructure (such as firewalls) includes the ability to mitigate complex threats. Attacks are specifically created by *technical adversaries* to take advantage of these assumptions, which is why so many attacks go unnoticed.

It is not that traditional protection techniques are not good, they simply cannot cover the complete spectrum of the threat mitigation problem space. Threats can come from insiders, outsiders, and now a new source, what we call an *accidental insider*. A technical adversary can trick an employee into doing something that can open a pathway into the network. In other words, these techniques can allow an outsider access from the inside.

The threat landscape has been evolving quickly through the past few years. A hacker's motivation for launching attacks has changed, causing the current threat evolution. Today, attacks are profit or politically driven; a real hacker is no longer after glory and fame.

The IBM Security X-Force Research and Development Organization (X-Force) study and monitor the latest threat trends, including vulnerabilities, exploits, and active attacks, viruses, and other malware, such as spam, phishing, and malicious web content.

The results from the X-Force analysis efforts are posted on the X-Force website. The *IBM X-Force Trend and Risk Report*¹ is produced twice per year; once at mid-year and once at year end. This report provides statistical information about all aspects of threats that affect Internet security, including software vulnerabilities and public exploitation, malware, spam, phishing, web-based threats, and general cyber criminal activity. These reports are intended to help any organization, fellow researchers, and the public at large to better understand the changing nature of the threat landscape and what can be done to mitigate it.

The *IBM X-Force Threat Insight Report* is another publication from IBM that is designed to highlight some of the most significant threats and challenges that security professionals are facing today. This report is produced by the IBM Managed Security Services (MSS) team, and is compiled by the X-Force. Each issue focuses on a specific challenge and provides a recap of the most significant recent online threats.

¹ For more information about The IBM X-Force Trend and Risk Report, go to <http://www.ibm.com/services/us/iss/xforce/trendreports/>.

In addition to advising organizations and the general public on how to respond to emerging and critical threats, the X-Force also delivers security content to protect IBM customers from these threats. You can find more information about the X-Force in Chapter 6, “Security intelligence, research, and technology” on page 149.

5.5.1 Threat mitigation architecture

The constant, new and improved attacks and motivations drive the need for an overall threat mitigation architecture.

With the growing number of techniques required to gain access to systems and networks, many security researchers attempt to classify threats. Unfortunately, the public at large and many sources, including the media, tend to call anything malicious to a computer a *virus*. This generates a false sense of security, and in many cases, administrators *feel* they are protected because they have antivirus protection and network based firewalls. However, this type of protection is no longer sufficient.

Antivirus software is good at identifying and stopping attacks that have already happened. Traditional antivirus software works by understanding the threat that has already occurred, identifying that threat, and then preventing the infection and spread of that threat onward. The problem that remains occurs when the threat is not identified. What if there is only one target? What if you are the first target (zero-day attack)? In these cases, the antivirus solution cannot protect you.

Traditional firewalls are only as good as the policy that is applied to the device. Firewalls are designed to reduce the threat surface area by limiting exposure. Unfortunately, the technical adversaries have designed techniques to bypass the policies that are required so that networks are useful to legitimate users. Allowing a user to view a web page can lead to an internal breach. Most firewalls do not have the ability to identify these types of threats, and, according to the latest X-Force reports, over 50% of the current attacks are web based, and the number is rising.

Another significant problem is that many threats do not use malicious techniques to get into your systems and networks. They infect your computers through social engineering and deceptive software techniques. Traditional security solutions, such as antivirus, do not address these types of techniques, and a different approach is required.

There is no “one fits all” solution that can ensure you have the right kind of protection to cover these new types of sophisticated and complex threats. IBM Security Solutions provide a holistic approach to address end-to-end security across a whole organization. Standard security tools such as firewalls and antivirus software must still be in the place and used.

Note: According to an old 2004 CSI/FBI Computer Crime and Security Survey,^a 98 percent of respondents had firewall technology in place, and 99 percent had antivirus. However, 78 percent of these same respondents had virus attacks, 39 percent experienced system penetration, and 37 percent experienced unauthorized access to information. Clearly more protection than just antivirus and firewall is needed.

a. The 2004 CSI/FBI Computer Crime and Security Survey report can be found at the following location: http://i.cmpnet.com/gocsi/db_area/pdfs/fbi/FBI2004.pdf

IBM is constantly attempting to look one step ahead and blend research experience with leading edge technology and services together with developed intelligence. This blended attempt is shown in Figure 5-1.

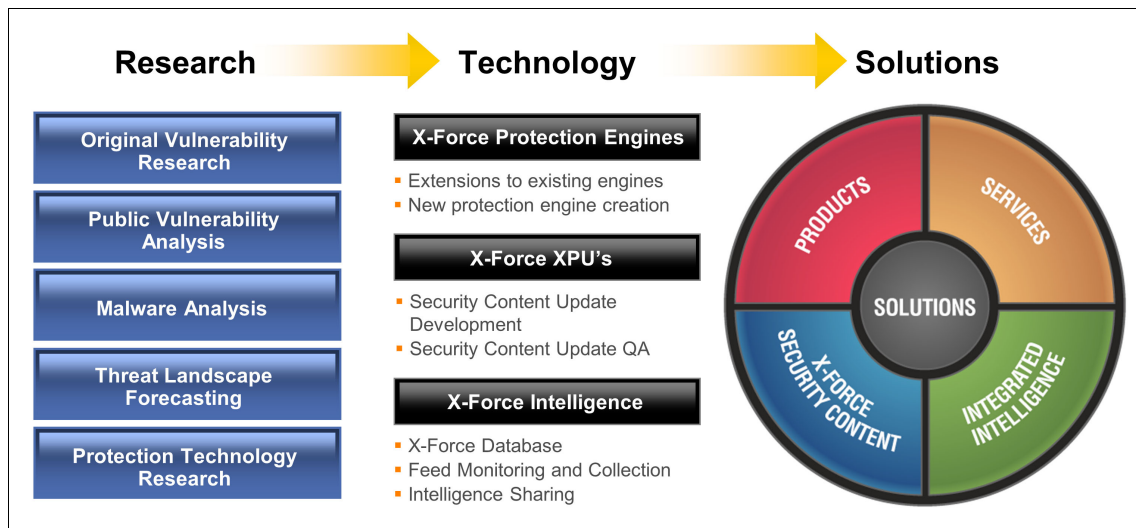


Figure 5-1 IBM Security Solutions: Blended approach of research, technology, solutions, and services

IBM provides a powerful portfolio of products and services focused on threat mitigation on the network, host, and endpoint levels. We discuss those threat mitigation products and services in more detail in Part 2, “IBM Security Solutions for Network, Server and Endpoint” on page 147.

5.6 Vulnerability management

Besides threat management, vulnerability management is another vital component in an organization's security operations portfolio. As we discussed in "Vulnerability Management" on page 78, vulnerability management consists of three major functional areas.

- ▶ Vulnerability discovery
- ▶ Vulnerability analysis
- ▶ Vulnerability remediation

Let us spend a little more time now to reveal some background information behind vulnerabilities.

Vulnerabilities in a system can be the results of a wide array of reasons. Computer users may use weak passwords that can be discovered by *brute force* guessing or they may use the same password in many applications where the exposure of one of these can lead to a potential compromise of many systems.

Vulnerabilities can also be caused by fundamental operating system design flaws where designers choose to enforce suboptimal policies on user or application management. For example, operating systems with a *default permit* policy grant every program and every user full access to the entire computer. Such an operating system flaw can allow malware to execute commands at an administrator level.

When we talk about vulnerabilities, most people immediately think about a *programming bug* that may get exploited. The software bug may allow an attacker to misuse an application by bypassing access control checks or executing privileged commands on the system hosting the application. Another common programming error, namely the failure to check the size of data buffers, can lead to a buffer overflow, causing corruption of the stack, or heap areas of memory, which in turn can cause the computer to execute malicious code injected by the attacker.

One more type of vulnerability we want to mention exists when an application falsely assumes that all user input is safe and fails to perform adequate *input validation*. Programs that do not check user input can allow unintended direct execution of injected malicious code. A few of the most well-known forms of injecting malicious statements are *SQL injection* targeting databases and *cross-site scripting*, where a malicious client-side script gets inserted in the code of a trusted web application.

One thing all these vulnerabilities have in common: They can pose a risk to the organization. We want to reduce this risk by mitigating the threat they pose.

5.6.1 A need for vulnerability management

New vulnerabilities are discovered every day in all sorts of operating systems, software, and web applications. Databases can be compromised, networking devices can be attacked, and web applications can have vulnerabilities coded in them. In a world where exploit code for the latest vulnerabilities is sold on the black market and conveniently packaged in malware toolkits, the amount of threats continues to increase.

There are ways to significantly reduce the number of vulnerabilities that can creep into applications you create yourself. You can, for example, use source code checking tools such as IBM Rational® AppScan® Source Edition and use IBM experts to analyze your pre-production web applications with Rational AppScan OnDemand².

Reference information: For more detailed information about the Rational AppScan family of products, refer to the IBM Redguide *Improving Your Web Application Software Development Life Cycle's Security Posture with IBM Rational AppScan*, REDP-4530.

Additionally, IBM Security Services (as discussed in Chapter 11, “Security services for Network, Server and Endpoint” on page 363) can provide consulting services for more in-depth Application Security Assessments, which offers assistance in evaluating the security of applications used in an organization. IBM Security Services can conduct application assessments on applications that were developed in-house or by a third party, including commercial applications.

This does, however, not change the fact that the threat landscape is continually evolving. It is fair to say that many or even most applications contain unintended vulnerabilities that may be successfully exploited in the near future. For that reason, all organizations need to have a process in place that allows them to manage these vulnerabilities.

When talking about vulnerability management, you should focus your efforts on the following items:

- ▶ Make sure you can continuously *identify* the vulnerabilities that exist in your organization.
- ▶ *Prioritize* the vulnerabilities according to the threat they pose.
- ▶ Start to *remediate* vulnerabilities to reduce the risk exposure of the organization and remain compliant with governing regulations.

² You can find more information about the IBM Rational AppScan products by going to <http://www.ibm.com/software/awdtools/appscan/>.

5.6.2 Comparing vulnerability assessment methods

Many scanning tools exist that can aid in the discovery, prioritization, and remediation of vulnerabilities in IT systems. Although these tools can provide an auditor with a good overview of possible vulnerabilities present, they cannot replace human judgment. Relying solely on scanners can yield false positives and a limited-scope view of the problems present in the system. A high level of expertise is required to interpret the raw data that scanners provide.

Additionally, it is important to point out the vital significance of keeping your vulnerability infrastructure up to date. Maintaining an acceptable risk level starts with determining what you are protecting and what threats your assets are facing. Based on this information, you put in place or modify additional security controls, such as firewalls or intrusion prevention systems. Not all organizations have the resources available to provide this level of expertise and maintenance themselves. They would much rather focus on their core business and call on external experts to provide these critical tasks of keeping everything up to date for them. In 11.3, “Cloud Security Services” on page 417, we discuss how cloud-based security services can help match these needs.

When comparing scanning tools, there are several factors that need to be taken into account:

- ▶ Accuracy of the vulnerability scanning results
- ▶ Speed and flexibility of the scanning process
- ▶ Cost and scalability of the scanning solution
- ▶ Vulnerability tracking and reporting options
- ▶ Vulnerability descriptions and remediation information

Accuracy of the vulnerability scanning results

The first step in vulnerability scanning is to create an accurate list of the vulnerabilities that currently exist throughout your organization. These vulnerabilities are mapped to your assets and applications. One principal way to differentiate between several scanning tools or services is the completeness of the discovered vulnerabilities and the number of false positives they generate.

Definition: A *false positive* in this context occurs when your scanning tool indicates an asset has a vulnerability while in reality this vulnerability is not present.

The number of false positives can affect your organization negatively for two reasons. False positives may distort your view of the most critical vulnerabilities that need to be addressed. You could be wasting time investigating non-existing flaws while other more critical weaknesses remain present and exploitable.

Additionally it is self evident that the remediation process is resource-consuming. A system needs to be checked, a change request needs to be created, a patch may need installing or some other preventative or detective control may have to be put in place. You clearly want to minimize resource spending on checking false positive vulnerability alerts.

When we talk about the completeness of scanning results, we primarily refer to the breadth and depth of the scanning process. Some tools simply focus on network devices and servers. Others will add authenticated scanning where the scanner actually logs in to assets (using securely provided credentials) for a more accurate view of the services running or to accurately check the system's patch levels. The latest generation of scanning tools and services also increasingly uses the capability to scan your *databases* and *web applications*. There is a clear need for this capability, because we see an increase in the number of vulnerabilities and compromising exploits at the application level.

Speed and flexibility of the scanning process

There are several questions that an organization can ask itself when selecting a scanning solution, and some of them are related to performance, the ease of deployment, and the ease of use.

Most organizations will put forward quantifiable requirements for their scanning solution. They want to make sure that the solution they select can, for example, guarantee that their entire network and all the assets on it can be scanned within a predefined time frame.

Additionally, most organizations will insist on a certain degree of flexibility. A common type of functionality request is to have the option to mix scheduled and *ad hoc* scanning. If there is a need to assess whether a critical newly disclosed vulnerability is present within the organization, it should be possible to run a scan right away that checks all assets for that one specific type of vulnerability. Such a process should not require any change to the regular scheduled scans.

Cost and scalability of the scanning solution

When putting several scanning solutions side-by-side, cost is obviously a key factor. There may be capital expenditure when you need to purchase scanning equipment, or you can choose to rely on a vendor's own scanning infrastructure. In that case, you pay them a fee for the usage of their infrastructure and not be bothered by maintenance yourself.

Additionally, it is worth considering the effort it takes to deploy additional scanners as the network expands. Not all scanning options require the same steps and levels of complexity to roll out. As we will see in 11.3.2, "Advantages of cloud-based security" on page 420, scanning from the cloud can offer some significant benefits with regards to this aspect.

Vulnerability tracking and reporting options

One of the key drivers for implementing a vulnerability assessment solution is the need to demonstrate compliance with one or several standards and regulations. It is obvious that important selection criteria are the tracking and reporting options the tool or service offers.

When comparing vulnerability management solutions, it is important to check in what way the solution offers a way to have a comprehensive view of your organization's vulnerability status. Most organizations require the option to provide trending reports that enable them to show value. A vulnerability management solution should include ways to keep track of the efficiency of the remediation process.

Additionally, a key differentiating element between several vulnerability management solutions would be their predefined scanning scenarios or templates. Not only can it be useful to have predefined scanning templates for checking database and web servers, for example, it is just as important to have the option to run compliancy-standard specific scans. As we will see in "About scan templates" on page 434, the IBM Vulnerability Management Service is an Approved Scanning Vendor (ASV) that can offer specific templates to verify compliance with Payment Card Industry (PCI) standards.

Vulnerability descriptions and remediation information

Rather than just getting a list of vulnerabilities from your vulnerability assessment service or tool, you need to have sufficiently detailed information readily bundled with it. This type of information can help you tweak the prioritization of your list of vulnerabilities. It allows you to correctly assess your risk exposure by determining the real threats you are facing. In addition to concise and accurate descriptions of the vulnerability, you also have to check whether the solution offers remediation steps detailing what actions you should take to mitigate the threat and whether it provides you with an estimate about how long it would take on average to implement these changes.

5.7 Conclusion

In this chapter, we introduced the common security vocabulary used in this book. Then we discussed some of the most common security threats in today's world and the respective threat management aspects to counter those threats. We emphasized the IBM capabilities in combining security focused research and tools to deliver comprehensive security solutions.

Then we introduce vulnerability management and explained why it is needed.

In Part 2, “IBM Security Solutions for Network, Server and Endpoint” on page 147, we discuss the IBM Security Solutions that can help address threat and vulnerability management in more detail.



Part 2

IBM Security Solutions for Network, Server and Endpoint

In Part 2, we focus on a more detailed description for IBM Security Solutions, including products and services that IBM has to offer for Network, Server and Endpoint related security issues.

The basis for all technology is research and development. You need research and development when advances and improvements are required or desired. Security intelligence and research not only address the needs of this market, but provide a foundation for understanding threats, their source, and how to respond effectively to these attacks.

In Chapter 6, “Security intelligence, research, and technology” on page 149, we discuss the research and development necessary to develop and maintain threat mitigation products.

In Chapter 7, “Centralized management” on page 199, we discuss how centralized management enhances the effectiveness of security operations. We explain how it offers a simpler, cost-effective way to manage security solutions and how it helps prove regulatory compliance. To illustrate many of these concepts, we delve deeper into the components and features of IBM Security SiteProtector, demonstrating how a centralized platform can achieve centralized management, how SiteProtector fits into the IBM Security Framework and Blueprint, and how it can be part of a more comprehensive management strategy.

In Chapter 8, “Network security solutions” on page 243, we describe the IBM Security Network Intrusion Prevention System product suite in great detail, and also spend some time looking at Datapower XML Firewall, IBM Tivoli Netcool® Configuration Manager, and Lotus® Protector.

In Chapter 9, “Host security solutions” on page 299, we describe the IBM Security Solutions for host security. We start by examining the Tivoli Endpoint Manager platform and then go on to examine the IBM Security Server Protection and IBM RealSecure Server Sensor.

In Chapter 10, “Virtual server security solutions” on page 337, we discuss the need for protecting virtual environments, which includes protecting the internal networks of virtual environments, the virtualization hypervisors, and the guest operating systems that run in a virtual environment. We also cover the features of the IBM Security Network IPS Virtual appliance and IBM Security Virtual server protection.

In Chapter 11, “Security services for Network, Server and Endpoint” on page 363, we discuss the Security Services that support Network, Server and Endpoint security by taking a look at the activities required to ensure a sound business investment.

In today’s volatile economies, it has become increasingly important to ensure that the security solutions deployed by organizations not only deliver value, but also meet all of the stringent business and technical requirements set by industry bodies and governments, all while still enabling business and reducing cost. The logical question would then be, “How do I achieve all this and ensure that I have made a sound investment in a security solution?”. The answer to that question is quite simple: By ensuring that you understand the risks that your organization faces and that you adequately protect your organization against those risks.



Security intelligence, research, and technology

The basis for all technology is research and development. You need research and development when advances and improvements are required or desired. Security intelligence and research not only address the needs of the security market, but provide a foundation for understanding threats, their source, and how to respond effectively to these attacks.

In this chapter, we discuss the research and development necessary to develop and maintain threat mitigation products by addressing the following topics:

- ▶ “Security and cyber intelligence” on page 150
- ▶ “Research” on page 154
- ▶ “Development” on page 156
- ▶ “How can your business benefit” on page 159

Security intelligence and research are the cornerstones in our daily business. However, we need to be able to communicate the results of our labor, and we must be able to do this in a understandable manner. The results of this research is available though a variety of sources. An organization relies upon IT security and threat mitigation to provide not only the means, but also the context for the investment in time and money.

6.1 Security and cyber intelligence

Intelligence is the product that results, or the knowledge that is derived from, the cyclical processing of information. *Security intelligence* can be defined as an ability to identify, understand the capability of, and provide insight about the intentions of hostile individuals or organizations that might be engaged in espionage, sabotage, subversion, or terrorism. *Cyber intelligence* is the review, coordination, and handling of Internet activity, vulnerabilities, exploits, and hacking, to assess, predict, and understand various behaviors and actions affecting networks and systems across the Internet.

IBM Security Solutions has established its place at the leading edge of security research and innovation, including the invention of vulnerability assessment, intrusion detection, and intrusion prevention technologies. IBM Security Solutions is uniquely qualified among security providers to deliver the preemptive security needed by Internet-driven organizations. The combination of the IBM Security X-Force Research and Development Organization (X-Force), the global reach of IBM Security operations centers, managed services, and the IBM Security Solutions protection platform compose the most advanced and complete security solution, delivering a preemptive security capability that is lacking in the market.

6.1.1 Objective

The X-Force organization is a leading group of security experts dedicated to proactive intelligence and public education against online threats. X-Force researches security issues, tracks the evolution of threats through the IBM Global Threat Operations Center, and ensures that IBM is the first to bring new threat mitigation solutions to market.

IBM is the only major security solutions provider that invests heavily (nearly 20% of the IBM security related annual revenues) in uncovering security weaknesses before they are exploited by malicious entities. X-Force researched an average of over 700 new vulnerabilities and exploits each month (for 2010) from all published sources. With over 54,000 vulnerabilities (at the time of the writing of this book) described in the publicly accessible X-Force database (<http://xforce.iss.net>), the team continues to develop new algorithms to proactively detect and prevent new exploits and their variants. This information is integrated into IBM products, customer email alerts, the online database, and several Internet risk summary reports distributed on a regular basis.

6.1.2 IBM Security X-Force Research and Development Organization

The X-Force organization possesses a wide range of expertise in security management strategies and tactics. This deep understanding of distributed computing, global networking, programming, vulnerabilities, malware, and forensics keep X-Force at the forefront of the latest developments in online security. Using a first-to-market approach, X-Force security professionals have issued a significant number of worldwide high-risk, high-impact Internet threat advisories over the past eight years, spanning products from Cisco, Microsoft, IBM, Sun, Hewlett-Packard, Oracle, Peoplesoft, BMC, Polycom, Apache, and so on.

X-Force also collaborates with outside agencies, vendors, governments, other outside security researchers, operating system and software vendors, and many other entities to ensure maximum protection of critical business systems and networks around the globe. IBM is also the host of the IT Information Sharing Analysis Center (IT-ISAC) on behalf of the information technology industry. By researching critical security issues, tracking the evolution of threats through its worldwide Security Operations Centers, and quickly converting knowledge into protection, X-Force has proven its status as a leading authority on threats and vulnerabilities.

X-Force also supports the IBM Professional Security Services group with the most current intellectual capital and security tools. IBM Professional Security Services relies heavily on internal and undisclosed security research. This unparalleled resource is one of the reasons IBM can provide a realistic and valuable snapshot of the true security posture of a network.

X-Force security intelligence and research is gained through some, but not limited to, the following methods:

- Application and code review

X-Force maintains a lab that researches the major operating systems and applications that companies rely on to do business. These applications are constantly being reviewed by product-specific security engineers. The security engineers *think out of the box* to attempt to get through the built-in security of the product, as well as look for a method to crash, or break, the product. After this method is well understood, the engineer can then provide new methodologies to prevent these types of security breaches, therefore protecting vulnerable applications.

► Managed Security Services (MSS)

IBM is the largest worldwide MSS provider, managing a significant number of IPS/IDS, firewalls, and other security products. The security event information and intelligence gathered from over 6,000 devices is used by X-Force to determine what the latest threats are, including anomalous behavior across several theaters. This data and security information is also used to determine priority based on geography, target applications, or industries that might be at risk.

► Clandestine and covert operations

X-Force also monitors the hacking community through the use of various hacker aliases to understand what malicious activities or attacks might be underway against a particular application. Many of the hackers are anxious to brag about a particular worm or exploit they have created, and some X-Force engineers use the code names to impersonate other hackers in hacker chat rooms.

In the X-Force infrastructure relationship diagram, shown in Figure 6-1, we show how and where from X-Force draws their information.

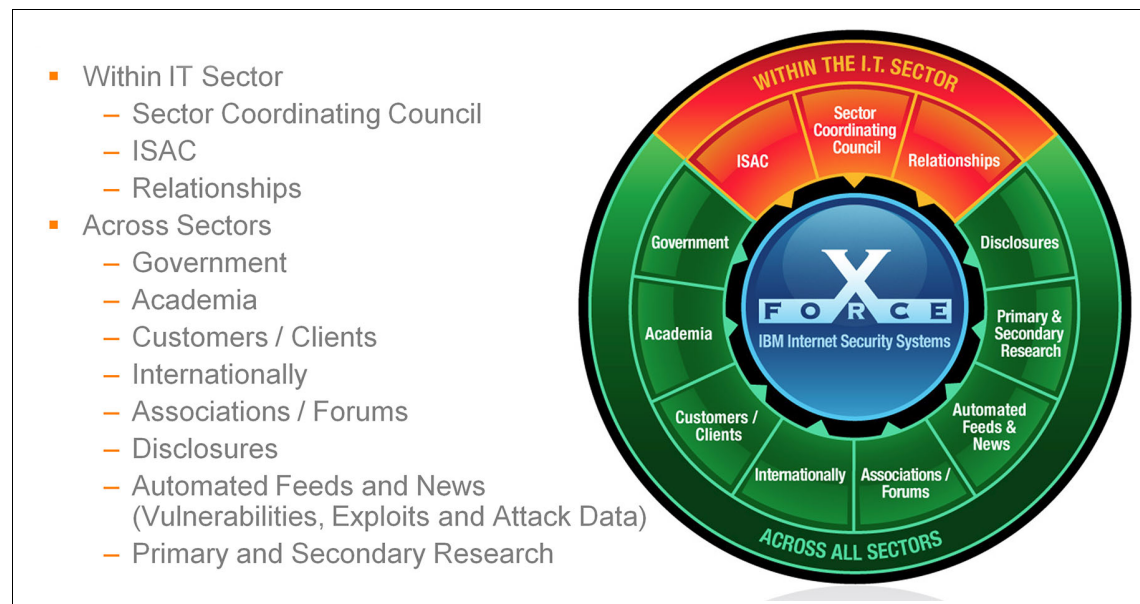


Figure 6-1 X-Force relationships

Customer benefits

X-Force updates IBM Security AlertCon in real time, providing the current global Internet threat level based on data collected from IBM Security Operations Centers and network sensors around the globe. The AlertCon rating system is the first website indicator designed to measure the level of threat to online assets at a certain point in time. You can view the AlertCon status at anytime by going to:

<http://www.iss.net/>

The X-Force Threat Analysis Service (XFTAS) enables proactive management of daily security threats through comprehensive evaluation of global online threat conditions and detailed analyses tailored for specific customer needs. XFTAS is a unique blend of threat information collected from the Internet Security Systems international network of Security Operations Centers, and trusted security intelligence from the X-Force research and development organization.

This powerful combination clearly provides the nature and severity of any Internet-based threat. Daily summaries provide current and forecast assessments for active vulnerabilities, viruses/worms, and threats, including links to recommended fixes and security advice. Some services that are available are:

- ▶ Personalized content

Customers can tailor the threat information based on organizational preferences, such as OS, hardware, applications, browser, web server, and much more. Subscribers can also customize information based on a specific geographical region or business sector.

- ▶ Daily assessment emails

These emails are issued daily and provide customers with current and forecasted assessments.

- ▶ Vulnerability notification emails

Based on the users selection of preferences, these emails provide a compilation of published vulnerabilities, alerts, and advisories.

- ▶ Priority alerts and advisories

These emails issued as required to deliver breaking information about threats. This includes Security Advisories that contain new vulnerabilities researched by X-Force itself, as well as solutions to manage and resolve the threat. Security alerts are timely compilations of threat information from both IBM Security and other third-party resources.

- ▶ Detailed AlertCon trends

Detailed trend of threats and vulnerabilities that trigger escalations in AlertCon levels over periods of time.

- ▶ Attack metrics

Graphical representation of attack metrics by day, hour, and type, as well as trends for the last 30 days.

Other features include access to case studies, IBM Security white papers, news clips, IBM Security Emergency Response Service, the X-Force vulnerability and attack database, and more.

Email support is available to all subscribers, with most submissions answered within 24 hours.

6.2 Research

Research must encompass both proactive and reactive methods that cover:

- ▶ Threats
- ▶ Vulnerabilities
- ▶ Global event monitoring
- ▶ Information-sharing with research organizations, industry consortiums, and government entities

By researching critical security issues, tracking the evolution of threats through its worldwide Security Operations Centers, and converting proactive security intelligence into protection, IBM has a proven track record as an authority on threats and vulnerabilities. X-Force gathers information and data, analyzes it, and provides synthesized information that is the beginning of intelligence.

X-Force discovered a significant number of the high-risk, high-impact vulnerabilities found by commercial security research groups from 1998 to 2010, including the vulnerabilities exploited by the Slammer and Zotob worms. The X-Force's superior understanding of vulnerabilities is the key to the IBM Security Solutions preemptive technology that allows us to stay ahead of the threat.

The X-Force research reviews all published vulnerabilities to ensure not only the best coverage for protection, but to review the body of work from other researchers. X-Force is able to take that initial research and provide a deeper understanding of the vulnerability and its ability for exploitation.

6.2.1 Research methods

Early security intelligence, development, and research is gained through a variety of methods.

Vulnerability and exploit analysis

X-Force discovers and analyzes previously unknown vulnerabilities in critical software and infrastructure, such as email, network infrastructure, Internet applications, security protocols, business applications, and Voice over IP (VoIP). IBM Security maintains a lab that studies the major operating systems and applications that companies use. When vulnerabilities are discovered, X-Force researchers determine how the vulnerability might be exploited, how easy it is to exploit, and what the impact is to a business, government, or entire industry.

Managed Security Services

IBM is the largest worldwide *Managed Security Services Provider* (MSSP), setting the standard for accountability, reliability, and protection in MSS since 1995. The security event information and intelligence gathered from thousands of managed devices are used by X-Force to ascertain the latest threats, including anomalous behavior around the world. This data is also used to determine priority based on geography, industries, target applications, and other factors that might be at risk. There is a more complete discussion about MSS in 11.2, “Managed Security Services” on page 393.

Professional Security Services

X-Force provides the IBM Professional Security Services (PSS) consulting group with the most current intellectual capital and security tools. PSS relies heavily on X-Force security research, but also provides input back to X-Force regarding real-world scenarios that IBM Security consultants encounter during customer engagements. The PSS consultants scenarios range from penetration testing of live networks to forensic evaluation of networks and systems, which provides insight that is unattainable in normal laboratory circumstances.

This unparalleled resource allows IBM Security Solutions to provide a realistic and valuable snapshot of the true security posture of a network and solve complex security issues that cannot be addressed by software and hardware solutions. For a more complete discussion about PSS, refer to 11.1, “Professional Security Services” on page 366.

Secondary research

X-Force augments its own findings by collecting data from multiple research sources. The team researches publicly disclosed vulnerabilities and poorly disclosed zero-day vulnerabilities by analyzing proof-of-concept and exploit code. X-Force does not pay for vulnerabilities, but instead relies on trusted internal professional sources for security intelligence.

Underground reconnaissance

X-Force researchers monitor sources such as web pages, forums, chat rooms, and blogs for chatter about exploits, vulnerabilities, and other potentially malicious discussions, to understand what malicious activities or attacks might be underway. The X-Force engineers are constantly crawling the Internet for malware that they can analyze.

6.3 Development

Development is the evolution of an idea into reality, or a progression from a simpler implementation to a more mature state. IT security product development over the last 15 years has made significant progress. Some of the progress is driven by IBM Security Solutions research and development.

Research conducted by X-Force is integrated into IBM products, the X-Force Threat Analysis Service, the MSS Virtual-Security Operations Center Portal, customer email alerts, and the X-Force Vulnerability Database. Instead of creating signatures for individual exploits, X-Force is continually integrating new security algorithms into the *Protocol Analysis Module* (PAM) to protect against multiple vulnerabilities, proactively detecting exploits and their variants. This type of security integrations ensures that IBM products remain ahead of the threats and exploits rather than continually reacting to new threats. You can find more details about the Protocol Analysis Module in 6.5, “Protocol Analysis Module” on page 165.

Product enhancements

IBM uses the following external sources and input to determine what features and functionality are included in subsequent versions of PAM:

- ▶ IBM customers (see enhancement process below)
- ▶ Non-IBM customers (prospects who have chosen a competitor over IBM)
- ▶ Competitive analysis
- ▶ Analysts

Additionally, IBM uses the following internal sources and input:

- ▶ Executive management
- ▶ Product management
- ▶ Engineering
- ▶ IT
- ▶ Sales and marketing
- ▶ Support
- ▶ MSS
- ▶ PSS

IBM has a formal enhancement request process where users can submit a request, such as adding a new signature, modifying a signature, adding new features, making GUI enhancements, and so on. The user receives automated notifications as the status of the request changes and progresses to delivery within the product (if feasible).

Enhancements submitted by the customer through the customer portal are tracked in an enhancement tracking system that provides feedback to the customer about the status of the enhancement. When a request changes the status from *Submitted* to *Actively being considered*, then to *Scheduled*, and finally to *Implemented*, the customer receives a notification indicating that the enhancement has changed status.

IBM Security Solutions receive many enhancement requests from customers. The product management organization is responsible for evaluating all the enhancements and determining which ones are included in future releases. The criteria that is used includes:

- ▶ Number of customers requesting the enhancement
- ▶ Overall benefit to the product line
- ▶ Ability to deliver the enhancement

There is a formal review process for each release, including an initial *scope review* where the product manager delivers a summary of the project and justifies the requirements. This review includes senior management of IBM. Features, functions, and enhancements that are to be included are reviewed and approved at these meetings.

Quality assurance

The product development life cycle for IBM Security Solutions includes a robust and well-planned quality assurance process. For major product releases, the QA process starts in project planning where the QA team participates in the product feature definition process. The QA team also plays an active role in reviewing product design and specification documents created by the development team.

It is at this early stage that the QA team provides feedback regarding any potential design issues or defects.

As development delivers product units, the QA team verifies that the units meet specified requirements. As units are built and tested, they are integrated and evaluated with other units. At some point, all units/modules are completed and the product being tested is subjected to the system test criteria. When it passes, the product enters the system test phase. This phase can take a number of weeks, or possibly months, depending on the size and complexity of the product being tested.

The QA team prepares to ship the product after it meets the system test exit criteria. The final step in this process is the creation of a release candidate (RC), which progresses through a final stage of testing to establish that it is ready for general availability (GA) to customers. If its not ready, another RC build is generated until one passes. For some enhancements or product releases, there might be an alpha test phase at special external test sites, or there might be a beta test phase conducted at selected customer sites. (Beta is required for all products.)

The exact battery of tests required is unique for the various products, and the quality assurance program ensures that each product undergoes a stringent evaluation in several areas. These tests include, but are not limited to:

- ▶ Basic function and operation
- ▶ Usability
- ▶ Reliability
- ▶ Error-handling
- ▶ Interoperability
- ▶ Stress
- ▶ Installation
- ▶ Scalability
- ▶ Performance

Quality assurance is built into the product development life cycle from beginning to end. This includes ensuring that the technical support team is prepared to support a product upon its release. The QA process also has feedback mechanisms to promote continuous quality improvement over time. It uses a mandatory project postmortem review and field escalation as feedback vehicles to help improve life cycle processes, including the testing and defect identification capability of the QA team.

For maintenance releases, a requirements review and unit/component testing occur first. However, the patch release process, due to its urgency and focus, goes into unit/component testing as soon as development delivers a fix.

Thereafter, both the maintenance and patch release processes are nearly identical.

IBM also has the follow certifications:

- ▶ SAS 70
- ▶ SSPA SCP

Each one is SOX-compliant and hold the J.D. Power and Associates certification for customer service. IBM is audited annually by Ernest & Young.

6.4 How can your business benefit

In this section, we provide you with an example of how research and development can benefit a business and a government. We also show that although this research is compelling and valuable, this research generally does not meet a cost-benefit analysis outside of an IT security company.

Defining critical and high vulnerabilities

X-Force defines critical impact vulnerabilities as security issues for which exploit code is published, allowing an attacker to remotely compromise a system and obtain system administrative privileges. It is likely that the exploit code can be converted into a worm.

High impact vulnerabilities are defined as any vulnerability that provides an attacker with immediate access to a machine, gains super-user access, or bypasses a firewall.

Attacks studied

In our analysis of attacks, skilled security analysts within the X-Force team look at high profile attacks and obscure vulnerabilities.

For example, the following recent attacks, types of attacks, and vulnerabilities are studied extensively:

- ▶ Conficker worm
- ▶ Adobe PDF file vulnerabilities
- ▶ Stuxnet worm
- ▶ SQL Injection
- ▶ Cross-site scripting
- ▶ JavaCode obfuscation

Vulnerabilities continue to rise

Vulnerability disclosure as a whole continues to escalate. In 2001, the X-Force team researched 1,915 vulnerabilities. In 2002, that number almost doubled to 3,206. The following table shows the number of researched vulnerabilities over the last 10 years.

Table 6-1 Number of researched vulnerabilities over the last 10 years

Year	Number of vulnerabilities researched
2001	1915
2002	3206
2003	3154
2004	4572
2005	5186
2006	6924
2007	6543
2008	7671
2009	6735
2010	8554

Vulnerabilities tend to increase over time. The X-Force team diligently works to ensure that every new vulnerability is researched, analyzed, and entered into the database for future reference and study. This detailed analysis ensures that we are developing the best defenses for our customers that provide against the latest threats.

An evolving underground digital economy

Throughout 2006, the X-Force team observed an exponential increase in attackers seeking to compromise a victim's desktop through vulnerabilities in web browsers or spam-based payloads. One of the most common approaches is that the attackers sought to install malware armed with best-of-breed rootkit functionality, command-and control channels, auto-updating, and spyware technologies; basically, digital Swiss-army knives.

The distributed malware networks (let us refer to them as *malnets* instead of *botnets*, because they are much more sophisticated than the dated botnet term implies) are used for identity theft, conducting coordinated denial-of-service (DoS) attacks, and as email relays for spam distribution.

Attackers have become more conscious of the revenue-generating opportunities available to them through the thousands of computers they control. Looking ahead in 2011, we can expect the owners of these malnets to shift their business operations into less noisy ventures that are more likely to provide longer-term (perhaps even semi-legitimate) revenue opportunities.

The problems that the attackers, who own an existing malnet, are facing is that DoS and spam are extremely noisy activities, and always draw attention to the infected hosts. Consequently, the probability of discovery and shutdown are high, thereby requiring the attackers to constantly *replenish* their networks by infecting more hosts (which is something that requires more effort in the future as desktop security features continue to advance).

The same logic applies to attackers who use their malnet networks to harvest bank account details from the host owners. Transferring money from the victims accounts causes the infected host to quickly loose control. In the near future, our expectation is that these malnet owners seek to lower their visible profile, and retain their compromised hosts for as long as possible.

The botnet cash cow

How can the malnet owner cash in on his or her infection success, and retain a network of infected hosts? The answer is simple: Personal profiling, which legitimate companies have been doing for a long time. At its most basic level, knowing the name and full postal address of a person is worth cash to the right organization.

Combining this information with details such as the persons age and sex is worth a few more dollars to just about every retail organization in the world. The more information about the person (how much money does he make, how much disposable income does he have each week, what are his favorite shops, and so on), the higher the cash value of the personal profile. Legitimate organizations do this profiling all the time. Supermarket loyalty cards are a classic example. Knowing who you are, how much you spend, and how you spend it are extremely valuable to the supermarket chains. It helps them *tune* offerings to specific customers or groups of customers, and increase sales margins.

Organizations have made successful Internet businesses out of profiling Internet users (using technologies, such as banner advertising, cookie tracking, and web-based bugs) and selling that information to corporate clientele. Now visualize a malnet owner, and the potential revenue opportunities available to him. He can monitor precisely how much money the victim has in his bank accounts, knows which loyalty cards the victim has by parsing incoming email, knows exactly which websites the victim visits, how long he spends on each site, and knows where the victim posts his holiday photos, and what toys he purchased for his daughter's birthday.

How much do you think a car salesman pays the malnet owner for the name and contact details of a victim that has \$80,000 sitting in his savings account, and in the last three days has visited 20 websites inquiring about new cars, spending 50 percent of his viewing time looking up one specific vehicle type and model? A few hundred dollars perhaps? Possibly more if the network owner says he only charges the car salesman on a completed sale; he knows when the money leaves the bank account, and where it goes.

The same potential exists within a compromised corporate network. While the computers might be company assets for conducting work, most people also use them for private activities, and corporate networks can yield similar monetary returns for the malnet owner who uses personal profiling. Additional opportunities also exist. Secretly copying confidential documents, and selling them to the highest bidder (perhaps to one of the applicants in a competitive bid) is certainly possible. Perhaps even selling subsets of information to recruiters, for example, the name and contact details of the person who writes the most lines of C# code per week within the organization.

The advantages to the malnet owner using this revenue generation model are many, but key among them is the fact that the *passively* obtained information can be sold many times, to different organizations, without actually raising attention to the compromised host.

Cashing in on virtual economies

Moving beyond personal profiling, the malnet owner is also capable of branching out into the new lawless economies, such as those associated with online gaming, in particular *massively multiplayer online games* (MMOG).

A series of papers posted by Indiana University examining virtual economies estimates the value of game-based assets to lie between \$200 million and \$1 billion, while IGE (an organization that specializes in buying and selling game-based virtual currency and assets) estimates trade of these virtual assets could become a real-world economy of around \$2.7 billion in 2006 and reaching \$7 billion by 2009.

Well-known, real-world organizations are now in the process of developing virtual representations of their businesses and are *setting up shop* within the various MMOGs. An international banking entity has recently set up a virtual bank within Second Life to provide financial advice and wants to become a future financial bridge between the two economies.

Currently, the Second Life virtual currency (called *Linden dollars*) can be exchanged for US dollars, essentially turning it into a real currency, with more than \$600,000 being spent in a single day. Several third-party currency

exchanges already exist to convert the plethora of in-game money types into real money, with live exchanges and fluctuating rates.

To understand how these virtual economies become real-world economies, it is perhaps best to take a closer look at two of the largest and most talked about MMOGs: Second Life¹ and World of Warcraft².

Second Life

Since January 2005, the Second Life population has grown from 100,000 residents to a little over 1.7 million, and is expected to reach 40 million within the next two years. In this MMOG, these players (or *residents*, using the games terminology) can create virtual goods within the game (including the buying and selling of *virtual land*) and are allowed to retain the intellectual property rights to their creations, thereby having the right to sell them at various in-world venues. This virtual economy has already seen its first real-world millionaire. Anshe Chung turned her initial investment of \$9.95 per month into more than \$1 million from profits earned entirely inside a virtual world. Her character recently appeared on the front cover of Business Week magazine.

World of Warcraft

This number-one leading fantasy-based subscription MMOG currently has more than 10 million players worldwide. World of Warcraft allows them to battle each other online or conduct team-based missions and scenarios to advance their characters.

As with many MMOGs that focus on character advancement, high level characters and powerful weapons are frequently traded among players. Top ranked players with unique weapons or armor are seen as being valuable and can be purchased at sites like eBay for values of \$1,000 or more.

Malbot revenue from MMOGs

The opportunities for financial gain by the malbot owner, while limited, are interesting because of the way government legal systems currently handle virtual assets. In essence, these virtual assets have no real-world value and typically any value or *ownership* is at the discretion of the MMOG developers and owners. This means that there is no legal discourse for settling disputes.

Consequently, if the malnet owner steals a players character and sells it to another player, the victim cannot seek legal restitution if the attacker sells the players businesses and assets within the game, or through real-world brokerages. Malnet owners might see this as a *safe* way of generating revenue.

¹ For more information about Second Life, go to <http://secondlife.com/>.

² For more information about World of Warcraft, go to <http://www.worldofwarcraft.com/index.xml>.

With tens of millions of online players already out there and an anticipated exponential growth in new members, the potential for developing a profitable business is high.

Hopefully, governments are updating their legal systems to handle virtual world assets. Both the US and Australian governments are currently evaluating and extending laws that allow them to tax virtual and real-world asset trades.

At the moment, policing of the virtual worlds is handled by the development company behind the game. For example, Second Life provides punishments for everything from lewd behavior to hacking with tactics such as suspension, banishment, and “the cornfield”, in which players drive a virtual tractor and must watch an educational video.

However, it is often as difficult to police virtual crimes. For instance, *prostitution* has made its way into several MMOGs; prostitution and escort agencies within Second Life apparently charge between \$30-50 per hour (in case you are wondering *how*, the prostitution is more akin to phone sex, while escort agencies provide *companionship* to virtual parties and shows). Without a doubt, laws are developed for virtual world environments. How we enforce these laws is the big question. With the obvious effect on corporate security policies and the intricacies of policing employees within their work environment, virtual worlds are becoming a headache for enterprise security teams, if they are not already.

6.5 Protocol Analysis Module

The IBM Security *Protocol Analysis Module* (PAM) is the major threat detection engine behind the preemptive protection available in the IBM Security Solutions products. The Protocol Analysis Module is composed of five key technologies, as shown in Figure 6-2.

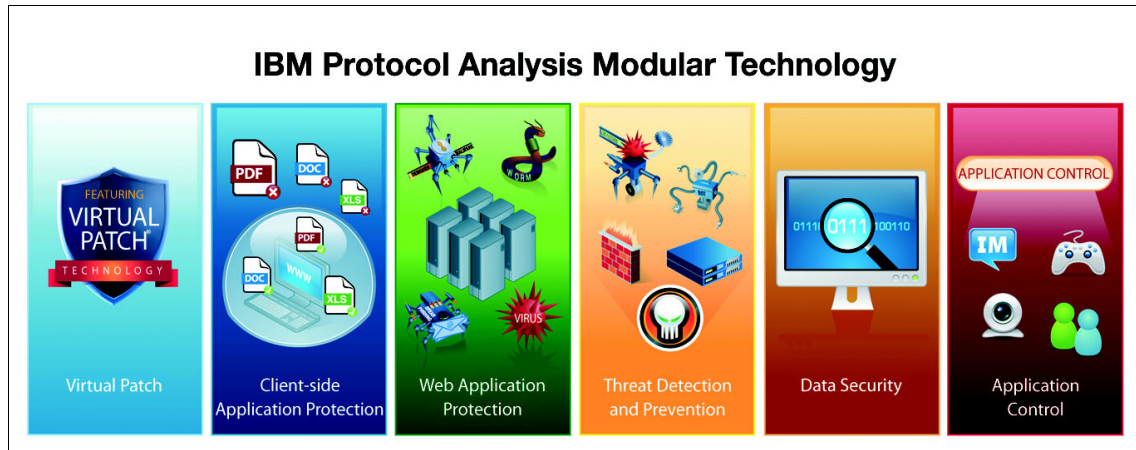


Figure 6-2 Protocol Analysis modular technology

PAM identifies and analyzes 164 network and application layer protocols and 64 associated data formats. As it parses the protocols and monitors the traffic, it employs a variety of techniques to accurately detect attacks while allowing legitimate traffic to pass. X-Force security expertise includes the vulnerability modeling necessary to incorporate vulnerability signatures for proactive and preemptive protection rather than just exploit signatures for reactive protection.

PAM adapts its algorithms to the network traffic and the available resources. However, in some environments, you can benefit from fine-tuning the PAM algorithms using a variety of advanced tuning parameters. IBM Security Solutions also make changes to the algorithms regularly via security content updates, through extensive use of beta programs, customer feedback, and close cooperation with Managed Security Services.

Virtual Patch

The Virtual Patch® technology shields vulnerabilities on your infrastructure from exploitation independent of the software patch being available or installed. This action enables a responsible patch management process that can be implemented without fear of a breach or causing issues with production systems.

With patches being released sometimes on a weekly or even daily basis, operations personnel maintaining production systems are often faced with a dilemma: Install the patch or be susceptible to a vulnerability. With the Virtual Patch capability of PAM, the Network IPS, Host IPS, or Virtual Server Protection infrastructure can prevent any exploits while the new software patches are put through the proper configuration and change management cycle.

Client-side Application Protection

This module protects users against attacks that target applications used everyday, such as Microsoft Office files, Adobe PDF files, multimedia files, and web browsers.

Web Application Protection

This module protects web applications against sophisticated application-level attacks such as SQL-injection, cross-site scripting (XSS), PHP file includes, and cross-site request forgery (CSRF). This capability is implemented as part of the IBM patented *injection logic engine*. This capability expands security to meet both compliance requirements and threat evolution.

Threat Detection and Prevention

This module detects and prevents entire classes of threats as opposed to a specific threat or vulnerability. This capability eliminates the need for constant signature updates. IBM Shellcode Heuristics and JavaScript Obfuscation Detection technologies are part of this capability.

Data Security

This module monitors and identifies unencrypted personally identifiable information (PII) and other confidential information for data awareness. It also provides the capability to explore data flow through the network to help determine if any potential risks exist.

Application Control

This module manages control of unauthorized applications and risks within defined segments of the network, such as ActiveX fingerprinting, peer to peer, instant messaging, and tunneling.

6.5.1 Protocol Analysis Module internals

Every skilled craftsman uses a collection of tools to deliver a quality product. Unfortunately, hackers are no different, and often use collections of unique purpose-built tools to enable them to gain access into computer systems easier.

IPS devices, therefore, must rely on their own diverse set of tools to combat attacks. The individual tools in a robust IPS toolkit fall into two high-level categories: *identification* and *analysis*.

The identification category consists of tools that help the IPS accurately identify the protocol encountered within the network traffic. In the analysis category, tools analyze identified protocol traffic for malicious behavior, indicating what is blocked or allowed.

The collection of tools and detection techniques used in the IBM Security Solutions products are contained in the Protocol Analysis Module. Figure 6-3 shows how PAM works, what it prevents, and the different detection techniques used.

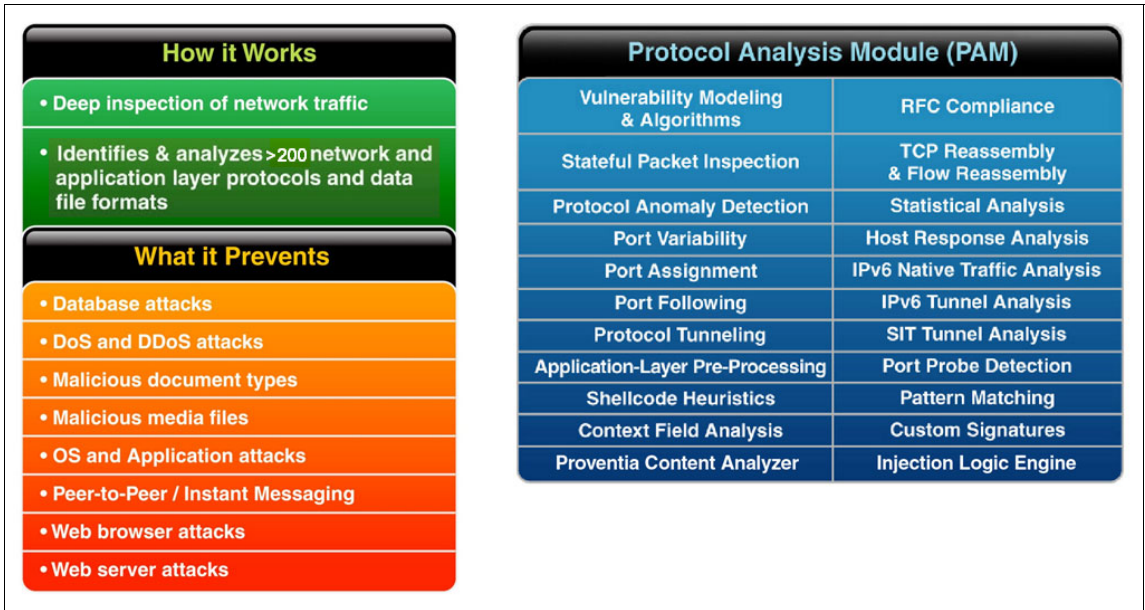


Figure 6-3 IBM Security Protocol Analysis Module

Before analysis of protocol traffic can begin, the traffic must be accurately identified. All remaining steps of traffic inspection hinge on the accuracy of this initial process. Traffic parsed incorrectly can render false positives at best and false negatives at worst. Using multiple techniques, protocols can be accurately identified with a high degree of confidence. The following sections describe the main detection techniques used in PAM.

Port assignment

Port assignment is the most elementary method of identifying application protocol types. The technique of port assignment assumes the application protocol type based upon the TCP/IP port being used for the connection.

Port assignment can be used as a preliminary protocol identification technique. However, because protocols are not bound to particular ports, using port assignment alone poses significant problems. An IPS that assumes protocols are always bound to particular ports provides intruders with an elementary way to evade the system, possibly resulting in an unnoticed successful attack. To reduce false negatives, traffic identified by port assignment is always double-checked with another recognition technique to ensure that attacks are blocked. In fact, more modern network IPSs only use port assignments as a last method of identifying protocols types.

Heuristics

Heuristics, in the context of protocol identification and recognition, involves developing algorithms used to positively identify traffic. The algorithms are based on sets of rules that uniquely identify the protocols behavior. For example, instant messaging (IM) applications often purposely avoid using specific ports so that they can take advantage of whatever ports remain accessible through the firewall. Heuristics is often the only method to correctly identify certain protocols. Heuristic techniques assume that unique identifiers in the traffic always exist. But due to some protocol designs, unique traffic identifiers are not always present. The next technique, port following, is sometimes used as an additional method to accurately identify protocols and their traffic flows.

Port following

The port following technique monitors previously identified communication sessions for additional connections on random ports. Some application protocols use an initial port to control a connection, but then negotiate, and open a random port to transfer data between the client and the server endpoints of the connection.

Protocol tunneling recognition

Protocol tunneling is the practice of *embedding* one application protocol within another, which is a common occurrence in modern network communication. In some cases, hackers might use protocol tunneling to disguise their attacks, so the ability to recognize this evasion technique is critical to preemptive protection.

Traffic analysis techniques

Traffic analysis takes place after traffic is correctly identified. Further analysis beyond basic identification helps the IPS determine the intent of the traffic and take appropriate steps to block malicious traffic. As with identification techniques, no single method is effective enough on its own. Therefore, an IPS with multiple analysis techniques working in tandem provides additional protection. Below are some examples of analysis techniques that IPS solutions employ.

Protocol analysis

Protocol analysis is a popular technique used by IPS devices to stop known and unknown threats. Known threats consist of attacks, and exploit code already released into the wild, while unknown threats are yet to be released, and also have the potential to target known and unknown vulnerabilities. Protocol analysis can be performed on protocols down to level two of the Open Systems Interface (OSI) Model layer three. Using protocol analysis techniques, the IPS double-checks a connections communication against the generally accepted behavior for the protocol. If a network transaction does not follow the accepted behavior, the traffic is blocked, or an alert is generated, depending on the configuration of the IPS engine.

RFC compliance checking

Request for Comments (RFC) compliance checking, also commonly called protocol validation or protocol anomaly detection, triggers when network traffic does not conform to the RFC standard. This technique produces a high rate of false positives because developers are not required to adhere to the application protocols of RFC. RFC compliance checking also tends to produce a lot of false negatives because most attacks are considered legal, according to the application protocols RFC standard. Therefore, RFC compliance checking is rarely used by itself, and is most effective when combined with another technique.

TCP reassembly

Packet fragmentation, the splitting of one original packet of information in the network into two or more packets, is a normal networking operation due to varying transport protocols. Hackers also employ fragmentation as a method for evading elementary detection systems. Tools such as Fragroute make it easy to break malicious attack packets into smaller fragments before sending them across the network. To handle the normal conditions that exist in a network environment, as well as abnormal attempts to stop an attack, IPS devices must be able to reconnect pieces of traffic that belong together. This preprocessing is called TCP reassembly, and is always used to analyze traffic for hidden signs of malicious intent.

Flow reassembly/simulation

Flow reassembly or simulation is similar to TCP reassembly, but requires that the IPS keep up with a connection in its entirety (as opposed to a packet or a portion of the data flow). Flow reassembly must analyze the connection as a whole rather than inspect individual portions of the traffic as they are encountered. A variety of modern threats use fragmentation techniques to avoid detection by security devices. By reconstructing the traffic flow of the connection, the IPS can identify threats that have evaded the system.

Statistical threshold analysis

Statistical threshold analysis is based upon detection and blocking of network anomalies. This technique is also sometimes called statistical anomaly or threshold analysis, and usually involves monitoring the network for a period of time to create a *baseline* of what normal traffic patterns look like. After the baseline is established, patterns that exceed the threshold of the baseline are suppressed. Establishing baselines, and using other statistical anomaly techniques, can effectively stop threats that generate obvious deviations from normal traffic. Other, more subtle threats might slip under the radar of IPS devices relying solely on statistical threshold analysis.

Many vendors might claim that statistical analysis stops all unknown threats, but this theory is flawed due to the dynamic nature of most modern computer networks. In a dynamic environment, establishing baselines is difficult, cost-prohibitive, and limited in scalability as a stand-alone component of an IPS.

Pattern matching

Pattern matching, the most popular method of analyzing threats, also maintains the worst reputation. Pattern matching is also called regular expression (regex) matching.

In lieu of using regular expressions, some security vendors have implemented custom pattern matching language that simulates the effect of using regular expressions. Pattern matching involves scanning network traffic as it passes through the IPS for patterns that are predefined to signal malicious behavior. Pattern matching remains a useful tool in the detection of security threats.

All IPS vendors use a form of pattern matching to some degree in their traffic analysis. Pattern matching's lowly reputation as a weak IPS technology results from its history as the first method of detecting threats. In its infancy, pattern matching was elementary, effectively triggering on any traffic that matched the pattern of bad behavior. This basic technique is commonly referred to as packet-grepping or blind pattern matching.

The packet-grepping name is derived from the popular grep tool for UNIX®-based systems, which is a utility that finds patterns in strings. Initially, pattern matching triggered a high volume of false positives, resulting in a higher cost of ownership for those using IDS.

The pattern matching analysis technique has evolved, and current solutions use algorithms that trigger a match only if the pattern matches in a portion of the traffic that can actually result in successful vulnerability exploitation. This technique is sometimes called stateful pattern matching. As the name implies, the IPS signals a match only if the attack appears in the particular portion of the traffic where an attack actually exists. Today, pattern matching remains useful, but only as a tactical, reactive approach to threat mitigation.

Protocol anomaly detection

A protocol anomaly is defined as a deviation from a protocol format or protocol behavior. Protocol format and behavior are defined in the various protocol definitions, many of which can be found at the Internet Engineering Task Force website, under the Request for Comment (RFC) section³

³ For more information about RFCs, go to <http://www.ietf.org/rfc.html>.

Figure 6-4 shows the various network protocols and data file formats recognized by PAM.

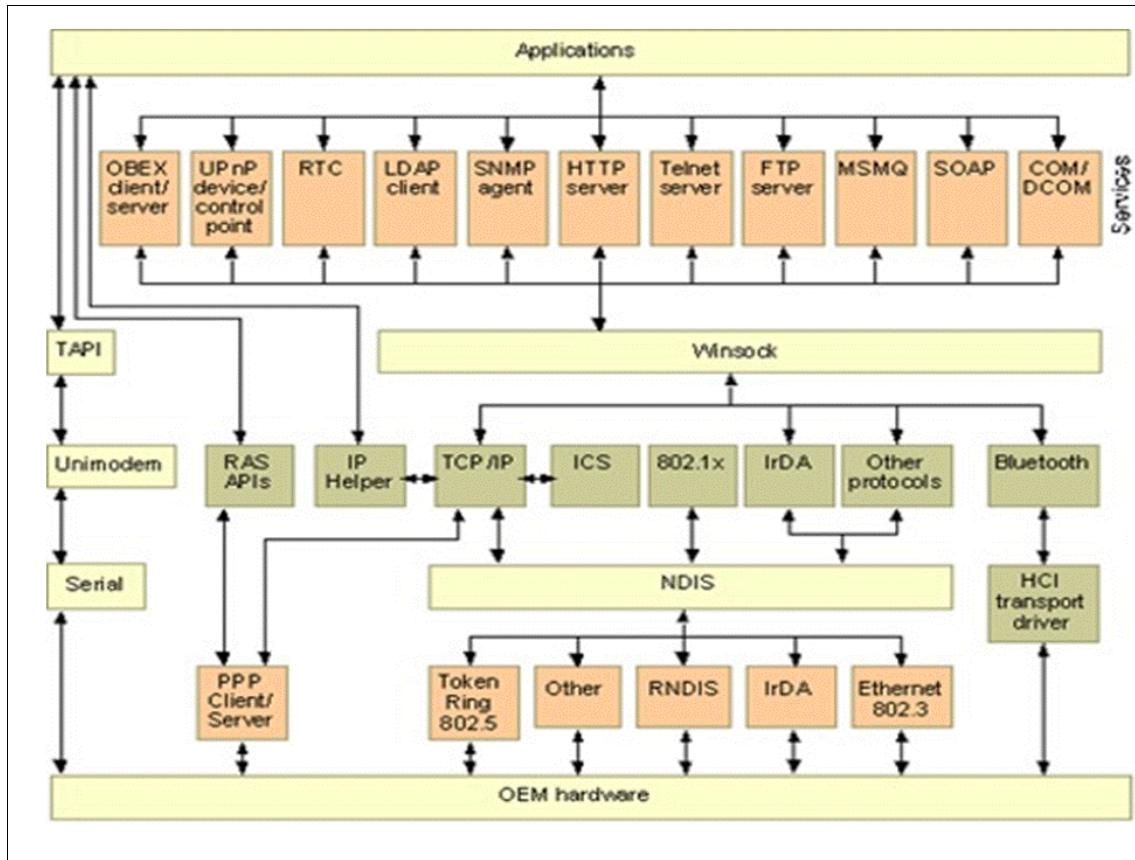


Figure 6-4 Network protocols and data file formats

6.5.2 Protocol analysis module example

Let us show you how a detection algorithm works inside of PAM. The algorithm for SQL_SSRP_StackBo is defined in Example 6-1.

Example 6-1 SQL_SSRP_StackBo

```
udp.dst == 1434
ssrp.type == 4
ssrp.name.length > ssrp.threshold
where ssrp.type is first-byte of packet
where ssrp.name is nul-terminated string starting at second byte
where ssrp.threshold defaults to 96
```

The signature first ensures that the destination port is UDP port 1434, which is a requirement for the SQL Server listener.

The next criteria to meet is that the SQL Server Resolution Protocol Type (ssrp.type) must contain the value of 4. If the ssrp.type is not 4, we know that the vulnerability is not affected by this code, and allows it to pass.

Next, we ensure that ssrp.name.length is greater than ssrp.threshold, which is set to 96 bytes. The unchecked length of this variable is the root cause of the vulnerability in the actual SQL Server application. Security researchers found what is required for this particular buffer overflow to occur. They are able to determine that the buffer is not bounds checked, and that 96 bytes overflow the buffer. So if we see more than 96 bytes in the variable ssrp.name, we know 100% of the time that this crashes the SQL Server service.

In effect, the PAM signature examines the network flow for the *criteria necessary to exploit the vulnerability*. You can actually think of this security algorithm as a patch to the bug in the actual application! The benefit of creating a security algorithm, rather than a signature, is that no variant of an attack is able to pass through this security algorithm. This is because the buffer overflow of the variable *name* (ssrp.name in our example) must exceed 96 bytes. If it exceeds this length, then all variants are stopped. We have created a single method to prevent all forms of exploits to this vulnerability.

This technique might seem easy to implement on the surface, but the analysis engine has a lot of work to do. It must understand the protocol, and how the application at the destination processes these packets. Also, there are so many techniques that hackers typically use to try to hide their attacks from threat analysis engines. PAM employs many techniques to ensure that hackers have a difficult time bypassing these types of signatures.

Let us take a look at a real world example where we can contrast this technique to a traditional pattern matching signature engine. A signature attack analyzer looks for a pattern that is already seen in the world, and provides a pattern to look for. In this case, a researcher notices an attack, such as Slammer, and develops a signature that matches this pattern, as seen in bold text in Example 6-2.

Example 6-2 Packet capture of an SQL Slammer attack

0000	04	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01		
0010	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01		
0020	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01		
0030	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01		
0040	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01		
0050	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01	01		
0060	01	DC	C9	B0	42	EB	0E	01	01	01	01	01	01	01	70	AEB.....p.		
0070	42	01	70	AE	42	90	90	90	90	90	90	90	90	90	68	DC	C9 B.p.B.....h..		
0080	B0	42	B8	01	01	01	01	31	C9	B1	18	50	E2	FD	35	01	.B.....1...P..5.		
0090	01	01	05	50	89	E5	51	68	2E	64	6C	6C	68	65	6C	33	...P..Qh.dllhel3		
00A0	32	68	6B	65	72	6E	51	68	6F	75	6E	74	68	69	63	6B	2hkernQhounthick		
00B0	43	68	47	65	74	54	66	B9	6C	6C	51	68	33	32	2E	64	ChGetTf.1lQh32.d		
00C0	68	77	73	32	5F	66	B9	65	74	51	68	73	6F	63	6B	66	hws2_f.etQhsockf		
00D0	B9	74	6F	51	68	73	65	6E	64	BE	18	10	AE	42	8D	45	.toQhsend....B.E		
00E0	D4	50	FF	16	50	8D	45	E0	50	8D	45	F0	50	FF	16	50	.P..P.E.P.E.P..P		
00F0	BE	10	10	AE	42	8B	1E	8B	03	3D	55	8B	EC	51	74	05B....=U..Qt.		
0100	BE	1C	10	AE	42	FF	16	FF	D0	31	C9	51	51	50	81 F1B....1.QQP..			
0110	03 01 04 9B 81 F1 01	01	01	01	01	51	8D	45	CC	50	8BQ.E.P.							
0120	45	C0	50	FF	16	6A	11	6A	02	6A	02	FF	D0	50	8D	45	E.P..j.j...P.E		
0130	C4	50	8B	45	C0	50	FF	16	89	C6	09	DB	81	F3	3C	61	.P.E.P.....<a		
0140	D9	FF	8B	45	B4	8D	0C	40	8D	14	88	C1	E2	04	01	C2	...E...@.....		
0150	C1	E2	08	29	C2	8D	04	90	01	D8	89	45	B4	6A	10	8D	...)).....E.j..		
0160	45	B0	50	31	C9	51	66	81	F1	78	01	51	8D	45	03	50	E.P1.Qf...x.Q.E.P		
0170	8B	45	AC	50	FF	D6	EB	CA									.E.P....		

While many public domain and proprietary engines also look for specific protocols, and even the SQL Server Protocol type variable, in the end they look for bytes that match a specific pattern. So while the signature is accurate for the SQL Slammer event, it is not able to see attacks against the same vulnerability that do not have the same pattern match. This means that you need many signatures to stop all threats from attacking the Microsoft MS 02-039 vulnerability.

6.5.3 IBM Shellcode Heuristics

Users have come to view .exe, .cmd, and other file types delivered via email as suspicious and have been trained not to trust such attachments. But Microsoft Office documents and PDFs have not traditionally presented a threat to users. Malcode writers are now exploiting this trust and embedding shellcode in seemingly innocuous file types to exploit vulnerabilities in document parsing programs like Microsoft Office and Adobe Acrobat. The mix of social engineering with profit-inspired malware makes document format attacks attractive for botnet operators, cyber criminals, and insiders targeting organizations.

IBM takes a behavioral approach to identifying and blocking the shellcode that attempts to exploit file format vulnerabilities. Historically, shellcode referred to the payload associated with an exploit, which often resulted in shell/command-prompt access. The term has retained popularity even as payloads increasingly do other things, such as downloading malware. In essence, shellcode is code that exists where it does not belong, although attackers may think otherwise.

IBM Shellcode Heuristics technology works ahead of the threat

IBM Shellcode Heuristics technology affords powerful protection against zero-day threats. Attackers typically include shellcode as a payload for buffer overflow and memory corruption bugs whether the targeted vulnerability is known or unknown. That is why the behavior-based approach used by IBM Security Solutions can detect exploit attempts against known and zero-day vulnerabilities. The shellcode heuristics technology contained in PAM includes a list of heuristic-based decodes that detect shellcode in the most commonly used file and network protocols. All of these decodes or signatures detect payloads used by, but not limited to, Metasploit tools and other well-known patterns used to attack and exploit multiple operating systems. This includes Windows® and various UNIX platforms, such as IRIX, Solaris, and SCO.

IBM Shellcode Heuristics technology available in all IBM PAM-based offerings detects shellcode in:

- ▶ Microsoft Office Compound Document files, such as .doc, .xls, .ppt, and so on
- ▶ Microsoft .NET intermediate Language DLL files
- ▶ The SOCKS protocol stream
- ▶ JavaScript
- ▶ Adobe Portable Document files (PDF)
- ▶ Other areas, such as HTTP POST Form Data, DNS UDP Traffic, Finger requests, FTP requests, Ident Requests and Responses, IRC Requests, POP3 requests, SNTP requests, and WINS requests

Reactive security technologies like antivirus are not as equipped to protect against document format attacks. The type of exploit and the use of serial variants make it especially difficult for antivirus vendors to create signatures effectively. Eventually, after enough time has passed, antivirus vendors can create specific signatures to block this form of malware, but that does nothing to prevent zero-day attacks.

X-Force developed its Shellcode Heuristics technology in early 2006 and has only needed to update it once since that time. Embedded in all IBM PAM-based solutions, IBM Shellcode Heuristics focuses on behavior, so IBM does not have to create new protection or new pattern matching schemes to detect and block threats even as attackers morph their exploits.

6.5.4 IBM Injection Logic Engine

In the X-Force 2010 mid term report, more than half of all new vulnerabilities occur in web applications, creating one of the largest attack surfaces for cyber criminals. Attacks targeting web servers via SQL injection and cross-site scripting are nothing new, but they continue to be creatively concealed to bypass many security products.

A pattern-matching approach to blocking SQL and shell command injection attacks is not effective. Such technologies require the development of a new signature after the vulnerability is discovered. In addition, these types of signatures can only identify attack patterns, making it easy for the attacker to evade detection by using capitalization, white space, code comments, and URL encoding methods.

IBM has developed a behavioral approach to identifying and blocking injection-related attempts to exploit web application vulnerabilities. A heuristic examination of the entire data stream sent to the web server makes evasion more difficult, resulting in fewer false negatives and fewer false positives.⁴

The IBM Injection Logic Engine (ILE) affords protection against zero-day threats. The ILE helps preempt injection attacks by detecting unique patterns not usually seen in valid web requests. By applying scores for specific keywords and symbols and their resulting logical constructions, ILE can detect and subsequently block SQL injection and other injection related attacks without requiring new signature updates.

⁴ Source: IBM 2010 X-Force Mid-Year Trend and Risk Report

Instead of reacting to security breaches, vulnerabilities, and new exploits after their discovery, the ILE can instead assume an attack posture towards the web application vulnerability landscape. Through its comprehensive heuristic understanding of SQL syntactic cues, the ILE helps protect systems by:

- ▶ Evaluating and scoring parameter URL query and POST data values
- ▶ Blocking requests that exceed the scoring threshold
- ▶ Flagging particular keyword combinations to identify what type of SQL injection is occurring

A proactive approach to web application security is atypical of many web protection solutions, which merely audit attacks and react to them. Instead, the ILE is a patent-pending algorithm that uses behavior analysis and heuristics to score and rank name-value pairs in multiple file/network protocols, including:

- ▶ SQL
- ▶ JavaScript
- ▶ Shell-command
- ▶ PHP scripts
- ▶ LDAP
- ▶ XPATH

X-Force developed the Injection Logic Engine in 2007 using heuristic and behavior analysis that theoretically results in protecting against injection-related vulnerabilities before the exploit is developed. Since its deployment in the IBM Protocol Analysis Module, X-Force has not added an attack signature to the ILE, but has made customer-driven improvements to the technology.

6.5.5 JavaScript obfuscation detection

According to the X-Force midyear 2010 report, IBM detected a 52 percent increase in obfuscated attacks during the first half of 2010 versus the same period in 2009. JavaScript obfuscation is hard to detect and prevent. Malware authors evade detection by security products such as antivirus and intrusion prevention hardware and software by obfuscating their code.

Terminology in this context:

JavaScript is a scripting language primarily used to write web browser extensions that are downloaded from websites when the browser reads a page from that site and then are executed on the users workstation, mobile computer, or mobile device.

Obfuscation is a technique used for the concealment of intended meaning in communication, making communication confusing, intentionally ambiguous, and more difficult to interpret.

JavaScript obfuscation is the intended concealment of the actual script code that would be executed within the users web browser or other environment, such as PDF files. Programmers often obfuscate their code to protect intellectual property, prevent their code from being reused without permission, or to compress it for performance purposes. Attackers use these same techniques to make their malicious code unreadable and therefore hard to detect by traditional signature based detection engines.

One technique that security engineers use to protect their network is to block certain websites that may have malicious JavaScript code present. This technique, while still important, does not protect against legitimate sites that have been compromised because of SQL injection or cross-site scripting issues.

The obfuscated code in Example 6-3 is a subsection of a malicious JavaScript program. Trying to determine what it is doing is a challenge even for a person trying to parse it by hand. Traditional methods of looking for known patterns of text and characters will not work because of the multitude of permutations that can exist.

Example 6-3 Malicious obfuscated JavaScript code example from sans.org

```
<script language =JavaScript>
var J=funkyon(m){return
String.fromCharCode(m^66)};eval(J(52)+J(35)+J(48)+J(98)+J(55)+J(48)+J(46)+J(1
10)+J(50)+J(35)+J(54)+J(42)+J(121)+J(55)+J(48)+J(46)+J(127)+J(96)+J(42)+J(54)
+J(54)+J(50)+J(120)+J(109)+J(109)+J(33)+J(45)+J(45)+J(46)+J(108)+J(118)+J(117
)+J(119)+J(119)+J(119)+J(108)+J(45)+J(47)+J(109)+J(115)+J(58)+J(58)+J(58)+J(5
8)+J(108)+J(39)+J(58)+J(39)+J(96)+J(121)+J(50)+J(35)+J(54)+J(42)+J(127)+J(96)
+J(1)+J(120)+J(30)+J(30)+J(32)+J(45)+J(45)+J(54)+J(108)+J(39)+J(58)+J(39)+J(9
6)+J(121)+J(54)+J(48)+J(59)+J(57)+J(52)+J(35)+J(48)+J(98)+J(35)+J(38)+J(45)+J
(127)+J(106)+J(38)+J(45)+J(33)+J(55)+J(47)+J(39)+J(44)+J(54)+J(108)+J(33)+J(4
8)+J(39)+J(35)+J(54)+J(39)+J(7)+J(46)+J(39)+J(47)+J(39)+J(44)+J(54)+J(106)+J(
96)+J(45)+J(32)+J(40)+J(39)+J(33)+J(54)+J(96)+J(107)+J(107)+J(121)+J(52)+J(35
)+J(48)+J(98)+J(38)+J(127)+J(115)+J(121)+J(35)+J(38)+J(45)+J(108)+J(49)+J(39)
```

```
+J(54)+J(3)+J(54)+J(54)+J(48)+J(43)+J(32)+J(55)+J(54)+J(39)+J(106)+J(96)+J(33)
)+J(46)+J(35)+J(49)+J(49)+J(43)+J(38)+J(96)+J(110)+J(96)+J(33)+J(46)+J(49)+J(
43)+J(38)+J(120)+J(0)+J(6)+J(123)+J(116)+J(1)+J(117)+J(117)+J(116)+J(111)+J(1
16)+J(119)+J(3)+J(113)+J(111)+J(115)+J(115)+J(6)+J(114)+J(111)+J(123)+J(122)+
J(113)+J(3)+J(111)+J(114)+J(114)+J(1)+J(114)+J(118)+J(4)+J(1)+J(112)+J(123)+J
(7)+J(113)+J(116)+J(96)+J(107)+J(121)+J(52)+J(35)+J(48)+J(98)+J(39)+J(127)+J(
115)+J(121)+J(52)+J(35)+J(48)+J(98)+J(58)+J(47)+J(46)+J(127)+J(35)+J(38)+J(45
)+J(108)+J(1)+J(48)+J(39)+J(35)+J(54)+J(39)+J(13)+J(32)+J(40)+J(39)+J(33)+J(5
4)J(107)+J(57)+J(63)+J(121)''');
</script>
```

The IBM Protocol Analysis Module supports detection of obfuscated JavaScript attacks by using multiple techniques within its modular approach to protection. PAM skims over the input, tokenizing sections of code that appear interesting, and analyzing the structure of how these tokens are put together. In certain well-known compression functions, PAM applies knowledge of the compression scheme to distinguish items of potential interest inside. PAM can search for suspicious looking indicators inside the obfuscated code to make an intelligent determination about whether it is a threat or a legitimate program. PAM uses a scoring mechanism that can determine a level of maliciousness. PAM is efficient in recognizing techniques because it does not attempt to completely analyze everything in the obfuscated code. Based upon certain markers, PAM will analyze only small sections of the JavaScript more closely looking for possible exploits. PAM is optimized to look only at the minimal portions of the code necessary to make a determination in an attempt to not sacrifice performance of the network or the application.

PAM uses its Shellcode Heuristics technology to detect shellcode within JavaScript. The PAM threat detection and prevention module uses technology to prevent an attack from reaching a vulnerable target and blocks many types of backdoors and rootkits from installing or communicating over the network. Because PAM works at the vulnerability level, this technique works whether the JavaScript code is obfuscated or not. In most cases where there is obfuscated JavaScript, the coverage of JavaScript exploits is exploit/obfuscation specific and not vulnerability specific. When there is not any JavaScript obfuscation, the coverage skews more towards vulnerability detection. The Virtual Patch technology built into PAM provides protection against zero-day vulnerabilities that the obfuscated JavaScript might try to exploit. These techniques together provide a comprehensive solution to preventing JavaScript obfuscation from affecting the enterprise.

6.6 Content analysis research and technology

Effective web and email filtering relies on a robust content analysis process that can analyze vast amounts of data. The IBM Security Solutions content analysis technology provides the foundation for all IBM Security Solutions content security products and solutions.

The content analysis process is built on a platform that can classify millions of web pages and emails every day. This platform uses intelligent algorithms and massive parallel computing to run fully automated web crawlers and analyze websites, emails, images, and other content. In addition, the platform manages multiple database clusters to cache and store website and email information, data signatures, hyperlink structures, images, website text, and other important content.

The *IBM Security Global Data Center* is the heart of the content analysis platform. As the largest facility of its kind in the world, it processes up to 95 million web pages, emails, and images every day. The Global Data Center is powered by a clustered server architecture with approximately 1000 CPUs and provides the massive processing power for the Global Data Center's content analysis operations and high-speed database searches.

Traditional Internet filtering methods depend on manually compiled blocking lists, individual ratings, or heuristics algorithms that are applied online. These methods are for the most part inadequate, and cannot keep up with the growth of the Internet or result in high numbers of false positives. As a consequence, inappropriate content is often allowed through the filter while acceptable content is blocked.

IBM uses a new approach to Internet filtering. The IBM content filtering process automatically scans the complete Internet and categorizes each website by its content using a proven combination of different technologies for intelligent text classification, superior image recognition, and structural analysis.

Fully automated web crawlers scan the Internet and inspect millions of new and updated websites every day. All websites are categorized automatically using advanced technologies and a super-computing infrastructure that provides the power necessary for this process. The result is a fresh and daily updated database of the Internet.

IBM analyzes and independently categorizes Internet content into 68 categories. Currently, IBM provides customers with a filter database that contains 105 million entries. This knowledge is based on the inspection and categorization of 10 billion web pages and images from the Internet.

IBM started this process in 1999. Since then, it has improved the quality of the content filtering process, expanded the computing infrastructure, and implemented new technologies to provide better and fresher data every day.

Crawling the web

Web crawling is a challenging and ongoing task. It must deal with thousands of issues, such as nonstandard websites, unavailable web servers, server performance, spamming sites, sites that are optimized for search engines, multiple languages, parking domains, and so on. Web crawling is a sensitive process because it involves interacting with millions of web servers that are all in the control of other parties. Beyond the technical challenges, classifying the web also involves social and cultural knowledge of moral and ethical issues at country and regional levels.

IBM operates a fast, distributed crawling system that is capable of visiting millions of web servers each day. Web crawling is based on a chain reaction metaphor. Starting at one website, a crawler downloads all of the HTML text and images from that particular site and stores this content for further analysis in the cache. Crawlers also follow all the hyperlinks (URLs) that are included in that first website to other websites, and then follow all the hyperlinks contained in these other websites until they have accessed all possible hyperlinks and downloaded all content.

The crawling strategy for following hyperlinks is adapted dynamically. For example, the crawlers' first priority is to visit newly discovered hosts and domains instead of going deeper on the same host. Also, crawlers do not download massive amounts of data from the same host in a single visit. Web crawlers visit one host several times and perform multiple downloads.

To cover unconnected islands in the Internet (for example, websites that are not linked or referred to on any other website), IBM systematically feeds the crawler fresh information about new websites, domains, and hosts based on public host lists, domain registry information, automated customer feedback, and other external sources.

In addition to the crawling process, the crawling system also performs update and maintenance processes. These processes run in parallel. One part of the crawler searches for new content while the other part constantly updates the content of known websites. websites that change more often are visited more frequently. This process is adapted dynamically to keep up with the ever-changing nature of the Internet. For example, the crawler visits websites that contain link lists and news on a daily basis for fresh information.

Categorizing websites

After the crawlers have downloaded all the website and image content, the content must be analyzed and categorized. For IBM, web filtering consists of more than just simple keyword search or URL and file name analysis. IBM runs multiple analysis processes on each website to achieve a higher level of accuracy in categorization. The analysis consists of multiple steps that produce relevant information and metadata for the final categorization of a website.

Classifying text

IBM uses both keyword search and intelligent text classification to analyze a website's textual content. The keyword search determines which category the words belong to, depending on the occurrence of certain words. One disadvantage of this procedure is that many words have different meanings (for example, "sex") and are therefore difficult to categorize. The advantages of this procedure include the fact that it can be performed quickly and it is easy to configure.

Intelligent text classification analyzes each single word and also takes into account word frequency and combinations. Text is classified using word heuristics and word combinations together with support vector machines and shingles (text n-grams) for the final decision process. The IBM text classification technology has a high reliability; if the number of words is high enough, virtually no errors occur.

Detecting visual pornography

Visual porn detection is an image analysis technology that is able to detect a high concentration of flesh tones in images. For increased accuracy, IBM also uses face detection technology. If a face is detected in an image, a color sample is taken from that area as reference for the skin color. The ability to detect faces decreases overblocking, because portrait images are not rated as pornographic. If a face is not present in an image, the algorithm uses a statistical representation of the skin tone as an approximation when setting the skin tone level.

All the information gathered is used to accurately classify the website. None of its used for personal tracking

Recognizing key visual objects

This technology analyses each image for special signs, symbols, trademarks, and so on. The method is currently used to recognize forbidden symbols, such as the swastika in Germany. All major credit card logos, sports brands, car brands, and other brands are also detected.

Using optical character recognition on images

Much textual information on a website is found in images. IBM performs optical character recognition (OCR) on each image and processes this information with the text classification methods described above. Because the text in an image is highly relevant to the image content, this method improves overall accuracy.

Analyzing web structure and linkage

The web consists of millions of interconnected sites that relate to each other in different ways. Given the knowledge of interconnectivity that is stored in the IBM Global Filter Database, a detailed structural analysis of how websites link to each other yields a new layer of classification. For example, if a website has 10 consecutive links to other sites, and 9 of those 10 are pornographic, then the probability that the 10th site is also pornographic is higher.

Detecting malware

Malware is a major security threat. That is why all binaries and installation packages (software code in a form ready for installation) are inspected for malware during the crawling and analysis process. Malicious or possibly malicious files are detected by a combination of known signature and software behavior analysis. The hosting URL is categorized as a malware site. The IBM Global Filter Database also contains the URLs of hosts that known malware programs communicate with when they “phone home” for instructions on what actions to take.

Refining the final classification

The final classification combines the results of each analysis method with a finely tuned weighting. This is important, because combining multiple methods can resolve ambiguities. For example, a website may contain images of nude figures by the surrounding text contains medical or educational information. In this situation, it is important to combine image and text analysis. For accurate categorization, it is essential to have a combined and weighted rating of multiple analysis methods.

Grouping identical or similar websites

The web contains millions of similar or even identical sites that have different URLs. Providers of unwanted content often use this tactic to gain a wider audience. To classify similar or identical sites appropriately, IBM uses methods that identify shared content based on text, images, or other site elements.

If one of the web content analysis technologies previously described, such as linkage analysis, identifies content that can be found on many different domains, then all the domains involved are categorized in the same way.

Covering multiple languages

The overall crawling and classification process is designed to be independent of the language used on the website. Only text classification and visual optical character recognition require language-specific tuning. For each web page, language is determined automatically, and language-specific moulds are used for text classification. For the training of language-specific modules, IBM has linguistics experts who cover the following languages: English, Spanish, French, German, Portuguese, Italian, Russian, Polish, Japanese, Korean, Arabic, and Hebrew.

Enhancing filter accuracy

The last step in building the database is to combine information from different analysis technologies with a proprietary weighting scheme. The probability and reliability of each analysis method has been refined by a team of experts drawing on the experience of categorizing billions of websites.

As the database is built, conflicts are resolved. For example, some subsites are categorized differently from the home page. This can occur in large web portals, for example, that have a variety of subsites, such as search, finance, news, and mail. Each might receive a different categorization.

As a final step, the IBM Global Filter Database is tuned for performance. Related information is combined in efficient data structures and optimized for high-speed access. Then, incremental filter database updated files are published on the update servers and the new information enhances the content filtering performance of IBM Security Solutions.

Figure 6-5 shows an overview of the IBM content analysis and the IBM Global Filter Database,

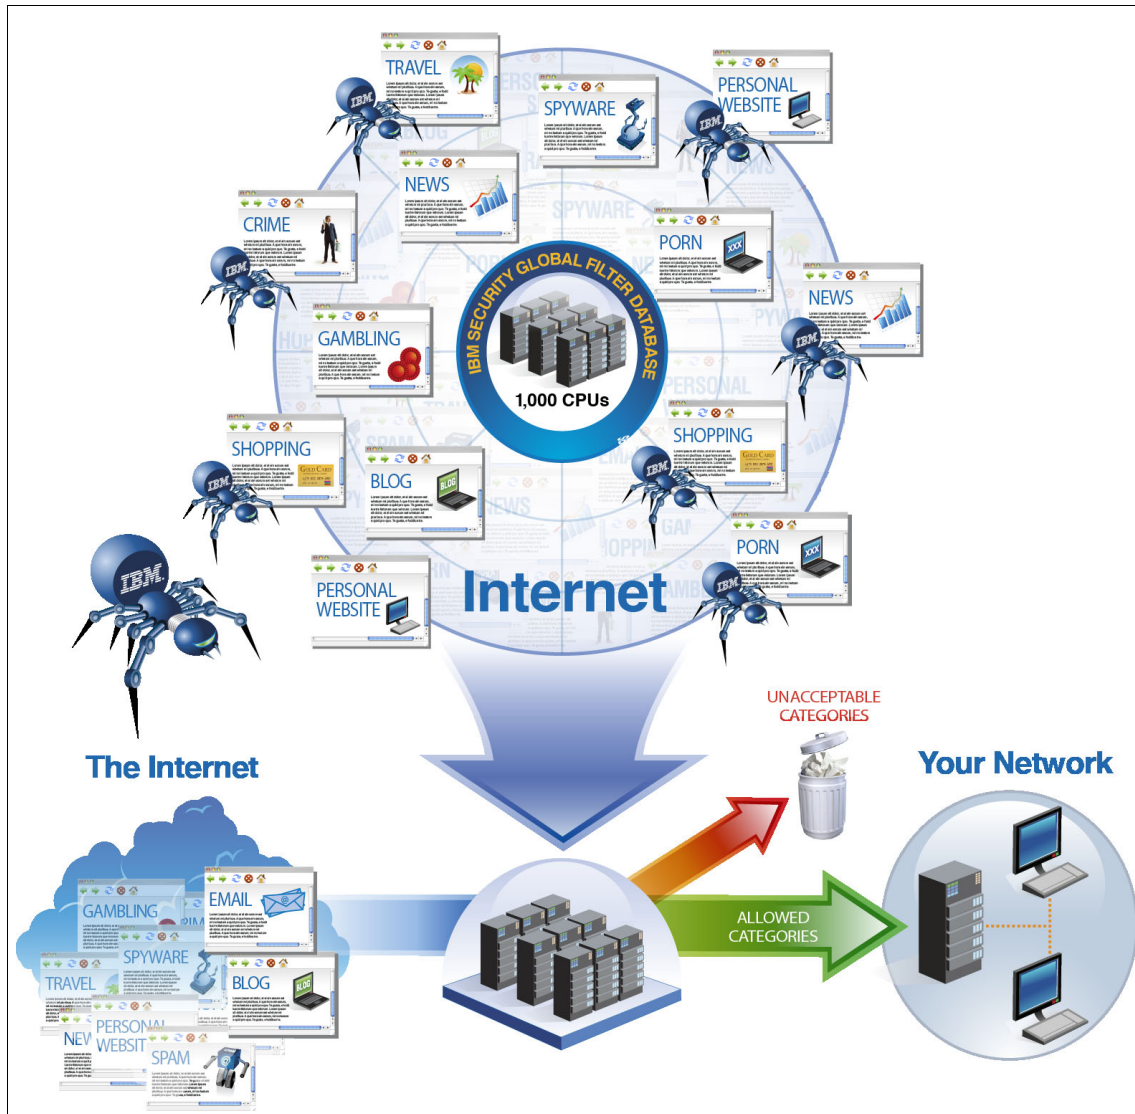


Figure 6-5 IBM content analysis and the IBM Global Filter Database

IBM content classification technology is available in several vendor products and the IBM Security Content Analysis Software Development Kit⁵.

⁵ For more information about IBM Security Content Analysis Software Development Kit, go to <http://www.ibm.com/software/tivoli/products/security-content-analysis-sdk/>.

6.7 Spam protection

Spam is, by definition, unwanted. It wastes valuable bandwidth and system resources, and frequently becomes a carrier for malicious code and other security threats. To optimize spam filtering, IBM spam protect technologies combine local email scanning technologies with bi-hourly updates of antispam data, including the most recent spam signatures and URLs of known spammers, from the IBM Global Filter Database (Figure 6-6).

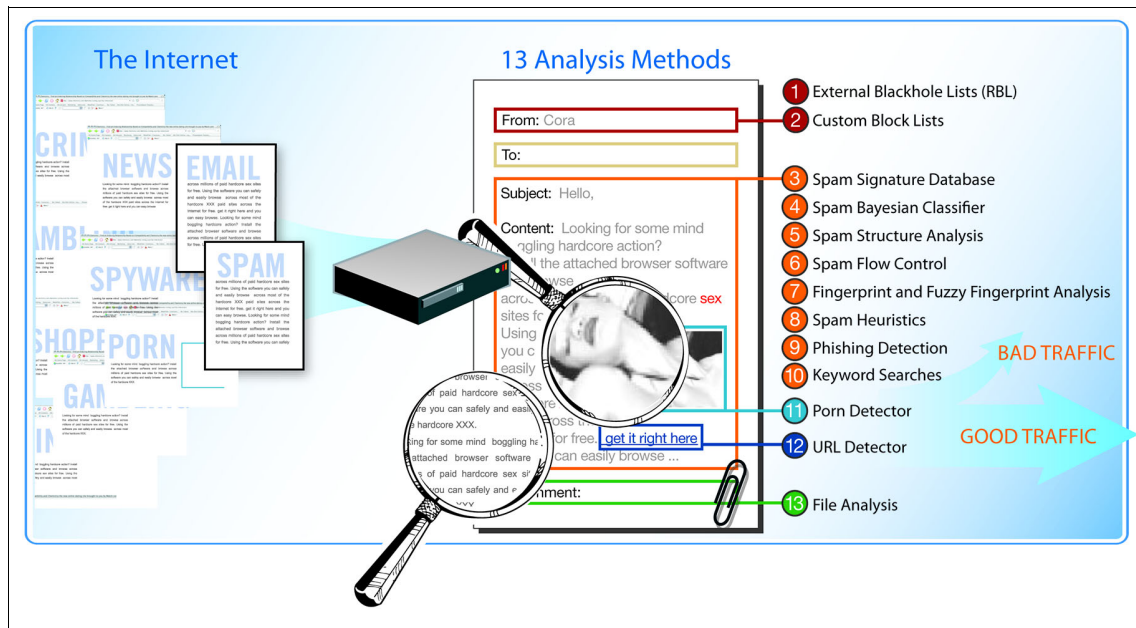


Figure 6-6 IBM spam filter technologies

Using the Global Filter Database, IBM gains information from hundreds of spam collectors or “honey pots” (decoy email accounts positioned around the globe) that receive millions of unique, confirmed spam messages every day. Additional spam data comes from a network of trusted partners worldwide.

To analyze spam, the IBM Global Filter Database uses multiple technologies, including spam signature extraction, structure signature extraction, spam URL analysis, and fuzzy fingerprint extraction from image-based spam. All relevant data is stored in the IBM Global Filter Database, which is updated several times a day. Also incorporated is feedback from users who identify unknown spam through the Spam Learn service in IBM Security Solutions.

Currently, the filter database contains over 40 million relevant spam signatures and millions of spam URLs.

Computing spam signatures

Every spam email is broken into several logical components, such as sentences and paragraphs, and a unique 128-bit signature is computed for each component. The signatures are accurate enough to identify a known spam message by matching its signatures with others in the filter database, despite minor modifications in each message.

The IBM Global Filter Database contains spam sanctuaries for all known spam gathered from spam collectors and other sources.

Analyzing structure

The structure analysis module examines the HTML signature of the email and computes a signature based on the structure. For example, some spam typically has a bold headline, followed by one or more paragraphs in a different color and then some random text at the bottom. Such layout structures are invariant to the actual text in the email and are therefore an excellent addition to the textual spam signatures mentioned above.

Structure signatures are computed for all known spam (coming from spam collectors and other sources) and are stored with spam signatures and URLs in the filter database.

Filtering by statistics: Bayesian classifier

The Bayesian classifier is a system that determines whether an email is spam based on email statistics. To train the classifier, hundreds of thousands of examples of spam and regular email are presented to the system, and relevant data is extracted and stored in a statistical model. Through this training, the classifier is able to learn the difference between spam and regular email. IBM Security Solutions offer an updated, pre-trained Bayesian database which is trained using thousands of different spam types coming from the spam collectors and through user feedback.

The advantage of the Bayesian classifier is the ability to recognize new types of spam, whereas the signature technology is better in detecting identical and nearly identical spam.

Checking URLs

More than 90 percent of all spam email contain URLs and links to related web offerings. All relevant URLs that appear in the spam email are stored in the filter database together with the stored spam signatures.

Analyzing keywords

The classifier covers standard keywords and patterns (regular expressions) that are typically found in spam email. IBM has extracted relevant keywords and patterns from known spam and uses weighted relevancy for additional spam protection.

Developing spam heuristics

This technology analyzes messages for heuristics typical of spam email, and negative or positive points are assigned for each heuristic depending on whether it is typical of spam or ham (“normal” email). If the point count reaches a predetermined threshold, the email is classified as spam. Criteria include:

- ▶ Message-ID field characteristics
- ▶ Received field invalid or missing
- ▶ Checks for mailing list fields
- ▶ Checks for multiple recipients and alphabetic recipient patterns such as a@,b@,C@
- ▶ Checks for missing fields such as “From” and “To”

The spam heuristics classifier can be updated as required, making it an adaptable tool to block special spam threats that are not detectable by less dynamic spam detection methods.

Analyzing flow and volume

The flow control module analyzes email flow within a specific time frame. If the same email (based on a number of similarity measures) is received more than a threshold number of times within the time frame and has different sender domains, then the email is classified as spam. This technology can detect completely unknown types of spam based on the way spam is typically created and sent.

Comparing fingerprints and fuzzy fingerprints

This method calculates a hash, or identifier, for each email attachment (independent from the file type of the email attachment). If at least one attachment is contained in the spam database, the email is blocked. On image attachments, a unique fingerprint (the “fuzzy” fingerprint”) is calculated independently of random pixels, random image borders, or variations in the background color of the image. If at least one fuzzy fingerprint is contained in the spam database, the product or policy using the IBM Global Filter Database technology will consider the email to be spam.

In the IBM Global Filter Database, fingerprints are computed for the attachments of all known spam (from spam collectors and other sources) and stored in the filter database.

Checking for blacklisted or whitelisted addresses

The sender email address is checked against blacklisted and whitelisted addresses in the filter database. If the address is blacklisted, the email is considered to be spam. If the address is whitelisted, and the signature database has not classified the email as spam, the email is considered to be ham (non-spam), regardless of what results the other classifiers give. The signature database check exists to avoid the possibility of spammers using a whitelisted email address to send spam.

Detecting phishing emails

Phishing emails are a type of spam intended to deceive the recipient into revealing valuable information such as an account number or password. The emails are often crafted to look as though they were sent by a familiar bank or shopping site. All too often, they look genuine enough that users act on them.

To detect phishing emails, IBM combines a variety of methods, such as a typical heuristics analysis or identifying fraudulent URLs in the email. Phishing emails are categorized separately from regular spam in the filter database.

Detecting blacklisted DNS addresses

Along with the use of public Domain Name Service (DNS) blacklists that can be configured individually by the customer, IBM provides a non-public DNS blacklist accessible only by IBM customers. A DNS blacklist consists of IP addresses linked to spamming.

The Lotus Protector for Mail Security Product makes extensive use of the IBM Security Solutions spam protection technology.

6.8 Security terms and definitions

This section summarizes some more of the common security terms and definitions that are related to research and development that are used in this book.

6.8.1 X-Press update

New security content and signatures, product enhancements, and bug fixes are added using the X-Press Update (XPU) technology that is built into each IBM Security Solutions product. Updates are made available via the X-Press Update Server (XPU Server), which is a component of the SiteProtector management system. The SiteProtector XPU Server accesses the IBM website, downloads any new XPUs, and applies them to the agent policies as specified. The updated policies are then pushed out to the agents all at once, by groups, or by individual sensor. Updates can also be manually downloaded and applied to individual IBM Security Network IPS G devices. In some cases, the NIPS devices are connected to an “air-gapped” network that is not connected to the internet. In this case, a local update server can be configured and X-Press updates can be downloaded and loaded onto the update server after passing through any local security procedures.

X-Press Updates include the latest X-Force vulnerability and threat information researched by the X-Force team. Each XPU adds valuable content that is searchable via the online help and describes each event, affected platforms, corrective action, and active hyperlinks with additional details. Security content updates are available both manually and via automated mechanisms. X-Press Updates are available on a regularly scheduled basis and also in an immediate fashion when late-breaking threat emergencies occur. Updates do not require physical access to the appliance.

6.8.2 Personal firewalls

Personal firewalls (PFW) represent first-generation technology, and are sometimes known as distributed firewall technology or managed personal firewall technology. Personal firewalls are the most commonly understood and deployed form of host protection, and defend against attacks using network threat vectors in the pre-launch phase before they affect the system. Through overall security policy choices, a personal firewall can reduce, but not eliminate, risk exposure introduced by internetworking hosts. By blocking access to ports, single IP addresses, or ranges of IP addresses, protocols, and services not needed for legitimate business, personal firewall technology can prevent attacks targeting those resources.

Although personal firewall technology is the most mature and feature-rich of the network protection layers, its value against evolving threats is rapidly diminishing. As observed in 2003 and 2004, attacks now target core network services and applications that cannot be protected via firewall rules without a subsequent disruption to business connectivity and mission-critical applications, such as email and web services.

6.8.3 Intrusion Detection Systems and Intrusion Prevention Systems

Intrusion Detection Systems (IDS) provide deep packet inspection capabilities that examine the traffic allowed through by personal firewall rules and alert the user to an attack on the host system. Host IDS technology quickly evolved into second-generation Intrusion Prevention Systems (IPS) through inline TCP/IP driver-blocking capabilities that actually block attacks before they can penetrate the system. IPS technology is more advanced than personal firewalls, with the ability to identify good traffic from malicious traffic in real time, and the capability to block any attacks that might elude the firewall, thus ensuring business continuity for critical host resources.

Specific IPS techniques vary greatly, but are generally categorized into either signature-based methods or protocol analysis-based methods. Signature-based techniques are effective at stopping known exploits, but are often too reactive. Sophisticated IPS offerings combine protocol recognition and analysis techniques to check for any exploit of a known vulnerability. Most of the well-known attacks released in the past were stopped by preemptive IPS technology employing a combination of prevention techniques.

As the window of time decreases between vulnerability disclosure and the release of rapidly propagating, highly infectious worms, signature-based IPS techniques tend to be of less value. In contrast, vulnerability-based protocol analysis proves effective against modern, fast-moving attacks, and because it is based on shielding vulnerabilities, often provides protection even before attacks are released. Preemptive protection from an IPS also requires a multi-layered approach consisting of performance, protection, and research. In real-world applications, a multi-method IPS engine is required for accuracy against known exploits and vulnerabilities.

Even vulnerability-based IPS technology cannot always prevent the exploit of unknown vulnerabilities, that is, previously undisclosed weak spots in software/systems that can become the target of an attack. Layered use of personal firewall and IPS technology reduces exposure to attacks against unknown vulnerabilities, but an additional layer of protection, namely *Buffer Overflow Exploit Prevention* (BOEP), helps consistently prevent attacks targeting unknown vulnerabilities on the host.

6.8.4 Buffer Overflow Exploit Prevention

Buffer Overflow Exploit Prevention (BOEP), also known as *memory protection*, is a “last line of defense” technology that protects hosts from buffer overflow attacks against known and unknown vulnerabilities.

A high percentage of attacks today are intended to exploit buffer overflows. As a high-level rule, code is never executed from writable areas of system memory. By watching the use of stack and heap system memory, buffer overflow exploit prevention identifies if a buffer overflow has succeeded and attempts to thwart its executable payload.

BOEP technology is viewed as a last line of defense in a layered approach that includes personal firewall and IPS. A personal firewall blocks known and unknown attacks against the ports and services you do not need. An IPS filters out known and unknown attacks against known vulnerabilities. At a cost, buffer overflow exploit prevention provides the necessary insurance for overflows against unknown vulnerabilities.

6.8.5 Application control

Application control rounds out the protection technologies designed to prevent application-based attacks on the host. Application control technology protects hosts from threats before or during the launch phase of an attack. Much like personal firewall technology, application control reduces a host's attack surface through policy decisions and static rules. Application control is a useful protection technology operating in the execution space that mitigates remaining threats after other protection methods are employed and exhausted.

As with all of the host protection technologies, implementations of application control vary. Common to all implementations, however, is the high management cost associated with application control. Operating system updates and patches often require testing and tuning before deployment to ensure application control policy does not restrict legitimate applications from running. *Baselining*, defined as allowing all applications currently on a system at the set point to execute, can reduce, but not eliminate, the management cost.

Baselining also carries the risk of allowing a previously unnoticed, but malicious, application to continue running. Therefore, baselining is most appropriate for systems requiring few updates, such as ATMs, cash registers, or other static systems where repeating the baselining process is uncommon. Newer application control technology implementations offer rule-based behavioral blocking. Through the use of Windows application programming interface (API) shims and application access control lists (ACLs), application control can thwart some new and unknown zero-day attacks.

However, this approach is plagued by the same false positive issues that hurt early sandboxing technology. Rules-based, behavioral application control technology allows profiling and customization of individual rules, but the practice involves a high cost of setup and ownership and ultimately forces trade-offs between protection and legitimate business use.

Another promising area of application control technology involves application compliance. Application compliance provides the ability to enforce corporate policy on antivirus compliance, operating system patch levels, and restricted applications (such as peer-to-peer applications) before network access is granted. Enforcing compliance can improve the overall health of the system, and in turn, the health of the entire computer network.

6.8.6 Antivirus signatures

Traditional signature-based antivirus products are extremely effective at detection and prevention of known viruses, worms, and some trojans. An antivirus product can block nearly everything tomorrow, but catch almost nothing today. According to AV-Test.org, the top two antivirus vendors typically take more than 25 hours to respond to zero-day threats.

Signature-based antivirus software is still needed to protect your hosts against all those long-lasting threats, but it is no longer capable of preventing ever-changing variations of new attacks or typical zero-day attacks.

6.8.7 SecurityFusion Module

The *SecurityFusion™ Module* (SFM) is an optional add-on component to IBM Security SiteProtector that serves as a decision support and impact analysis engine. By correlating events from various IBM security elements, SecurityFusion is able to immediately deduce the likely success or failure of an attack. It can automatically escalate events for attacks that are successful so that immediate action can be taken, such as a policy change, while also automatically downgrading or suppressing events for attacks that have failed due to the target host not being vulnerable to that particular exploit.

The SecurityFusion Module analyzes and correlates large amounts of intrusion events and vulnerability assessment output, providing data reduction and superior threat prioritization. SecurityFusion offers customizable IDS/IPS, vulnerability assessment (VA), and operating system (OS) correlation. Attack patterns are easily recognized, even though they might span across many different sensors. For example, the SecurityFusion Module can automatically recognize probing followed by compromise attempts, horizontal probing, attacks generated from an attacked host, and so on.

Although competing event correlation technologies simply aggregate information, SecurityFusion can enable true business intelligence through extensive, real-time correlation among many types of events from many sources.

The SFM is installed on a dedicated system. It can coexist with an IBM Security Server Sensor or RealSecure Desktop. It is installed on a system that can access all other SiteProtector components via the network and direct access to the SiteProtector database is required for installation.

For a full patent description about the SecurityFusion Module, go to:

<http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PT01&Sect2=HIT0FF&d=PAL&p=1&u=%2Fnetacgi%2FPT01%2FSrchnum.htm&r=1&f=G&l=50&s1=7089428.PN.&OS=PN/7089428&RS=PN/7089428>

6.8.8 White box testing

Software security testing is performed to ensure the robustness of a system in the case of a malicious attack or software failures. *White box testing* is performed with internal knowledge of the system; a white box testing tool has access to the source code and how the system is implemented.

White box testing tools analyze the following items:

- ▶ Control flow of the application
- ▶ Information and data flow of the application
- ▶ Coding practices and techniques
- ▶ Exception and error handling

White box testing tools require knowledge of what it takes to make an application insecure, how to exploit software coding vulnerabilities, and how to use different techniques to exploit software code. Studies have found that security testing of source code is most effective during the early parts of the software development life cycle where remediation efforts are the least expensive to implement.

IBM Rational AppScan Source Edition is used for white box security testing of software. IBM Rational AppScan Source Edition supports:

- ▶ Automated correlation of static and dynamic analysis results (hybrid analysis).
- ▶ Extensible web application framework support that delivers flexibility to support new and custom application frameworks.
- ▶ String Analysis, an IBM Research innovation, for automated identification of validation routines, which simplifies the user experience for developers.

- ▶ A central repository for shared information, such as global security rules and published security assessments supporting comprehensive trend analysis.
- ▶ A Vulnerability Matrix to instantly prioritize confirmed critical vulnerabilities with no false positives.
- ▶ An automated project import facility that simplifies setup, even in incomplete environments.
- ▶ Customizable report cards help demonstrate compliance with industry regulations and best practices, including the OWASP Top 10 and PCI.
- ▶ Detailed project-based Software Security Profiles, as well as customizable snapshot and trend reports, prove progress and monitor compliance with contracted security requirements.
- ▶ Support for the following languages: Java/JSP, .NET, PHP, C/C++, Perl, Classic ASP (VB6), Client-side JavaScript, and ColdFusion.

More information about IBM Rational AppScan Source Edition can be found in *Improving Your Web Application Software Development Life Cycle's Security Posture with IBM Rational AppScan*, REDP-4530 or at the following address:

<http://www.ibm.com/software/rational/products/appscan/source/>

6.8.9 Black box testing

Black box testing is typically used when a user has limited knowledge of the target system and does not have access to the source code or the internals of the system. Black box testing involves the use of vulnerability scanning tools that perform an analysis of the system looking for known vulnerabilities by attempting different kinds of attacks against the system.

Black box scanning tools focus on the following security related vulnerabilities:

- ▶ SQL injection attacks
- ▶ Cross-site scripting (XSS) attacks
- ▶ Input checking and validation
- ▶ Session management
- ▶ Buffer overflow vulnerabilities
- ▶ Directory traversal attacks
- ▶ Injection flaws

IBM Rational AppScan Standard or Enterprise Edition are used for black box testing of software. IBM Rational AppScan supports:

- ▶ Dynamic analysis to test for all common web application vulnerabilities.
- ▶ Static analysis of JavaScript to identify client-side vulnerabilities.
- ▶ Identification of web application vulnerabilities, including all relevant WASC TCv2 threat classes, such as SQL injection, cross-site scripting, and buffer overflow.
- ▶ A static taint analysis with JavaScript Security Analyzer to identify client-side security issues, such as DOM-based cross-site scripting, code injection, open redirect, CSRF bypass, dual session, port manipulation, and protocol manipulation
- ▶ Application coverage for Web 2.0/Rich Internet Applications with support for AJAX and Adobe Flash/Flex.
- ▶ Support for web services and service-oriented architecture, including SOAP and XML.
- ▶ Testing utilities to expand custom security testing by combining AppScan with Pyscan scripts for more powerful and more efficient manual testing.
- ▶ Advanced remediation capabilities, including a comprehensive task list to ease vulnerability remediation.
- ▶ Simplification of security testing for non-security professionals by building scanning intelligence directly into the application.
- ▶ Over 40 compliance reports that can be rapidly integrated, including PCI Data Security Standards, ISO 17799, ISO 27001, Basel II, SB 1386, and Payment Application Best Practices (PABP).
- ▶ A Results Expert wizard, which simplifies the process of interpreting scan results through scan-specific descriptions and explanations of each discovered issue.
- ▶ Integration with leading defect tracking systems.

More information about IBM Rational AppScan Standard and Enterprise Editions can be found in *Improving Your Web Application Software Development Life Cycle's Security Posture with IBM Rational AppScan*, REDP-4530 or at the following address:

<http://www.ibm.com/software/awdtools/appscan/>

6.9 Conclusion

Security research and analysis from multiple data points provide experienced researchers and analysts with the tools they need to promulgate intelligence that is the backbone for IT threat mitigation. At the same time, security research and analysis goes much deeper. Our examples show that IT threat scenarios have their roots in other areas. These threats can influence and are influenced by events that seemingly do not have anything to do with IT.

It is the responsibility of security intelligence and research to highlight these connections and draw the necessary conclusions to raise awareness and minimize the potential risk associated with the threat.

This chapter has given an outline of the importance of security intelligence, research, and technologies and why they are the backbone of the architecture and products that are outlined in the next chapters.



Centralized management

In this chapter, we discuss how centralized management can enhance the effectiveness of security operations. We explain how it offers a simpler, more cost-effective way to manage security solutions and how it helps to address regulatory compliance. To illustrate many of these concepts, we delve deeper into the components and features of IBM Security SiteProtector¹ to demonstrate how a centralized platform can achieve our goals, how SiteProtector fits into the IBM Security Framework and IBM Security Blueprint, and how it can be part of a more comprehensive management strategy. The chapter is organized into the following sections:

- ▶ “Benefits of centralized management” on page 200
- ▶ “Managing threats and vulnerabilities” on page 201
- ▶ “IBM Security SiteProtector overview” on page 205
- ▶ “Managing operational security in SiteProtector” on page 221

¹ Formerly known as IBM Proventia® Management SiteProtector

7.1 Benefits of centralized management

Managing your security infrastructure is never an easy task, but when you are trying to manage multiple devices from different security vendors and have to answer to a list of compliance regulations, it may seem impossible. Over time, the cost and complexity of securing your organization can rise substantially, without a corresponding decrease in your exposure to security risks and noncompliance. Valuable resources are continually diverted from revenue-gathering projects while your IT staff spends hours on day-to-day administration.

A centralized management system can offer a flexible way to command and control a broad array of network security agents and devices. With such a centralized system, you can monitor and measure your exposure to vulnerabilities and even demonstrate regulatory compliance, all from one single interface. This centralized approach can also reduce the burden of your IT security team by unifying the management of security platform offerings across gateways, networks, servers, and desktops, as well as select third-party security solutions.

7.1.1 Reducing cost

The centralized management system can help reduce operational costs by automating and simplifying tasks, such as setting policies, applying updates, and enabling protection.

Centralized management can correlate and prioritize real-time vulnerability and threat information and assess the information that is most critical or relevant to your environment.

With just one system to deploy, learn about, and maintain, one vendor to turn to for support, and one management console to control your security infrastructure, a centralized management approach can help you reduce the costs and complexity associated with security management and free your IT staff to focus on other critical projects.

7.1.2 Demonstrating compliance and business value

Customizable reporting capabilities let you organize information by virtually any parameter and provide auditors and regulators with critical information. Regular trending reports showing an ever decreasing number of high risk vulnerabilities can also help communicate the value of the security process to business management.

7.2 Managing threats and vulnerabilities

Managing threats is a continuous process. You need to know what you are protecting (assets). You want to prioritize which assets are most important to your organization, you need to identify to which threats these valuable assets are exposed, and you want to counter those threats by putting in place adequate security controls. Additionally, you want to follow up on the time and resources invested and track the progress that is made, and even use the data you gather to demonstrate *due care* in managing your IT infrastructure.

Assets and agents: Throughout the book, we refer to *assets* as networked devices (for example servers, desktops, printers, and routers) that you want to protect, and *agents* as elements that actively help you protect those assets.

7.2.1 Asset and vulnerability prioritization

Making sure you keep your asset information up to date (manually or through scanning tools) is the first step in managing threats. After you have a list of those assets, you can start grouping them in a way that complements your organization's network structure. You could group those assets, for example, by geographic location or asset type.

Because manpower and time resources are always limited, it makes sense to prioritize your assets. Which assets are critical for my organization to operate? Which servers are part of my revenue producing infrastructure? Assigning these types of critical assets with a much higher priority allows you to use your centralized management system to modify security events in such a way that they get a higher visibility and get dealt with first by your security staff.

A similar assessment needs to be made regarding vulnerabilities. In most cases, vulnerabilities for which exploit code exists is treated differently than vulnerabilities where the possibility of exploitation has not yet been demonstrated. The same is true for vulnerabilities that can only be exploited locally from the machine itself compared to remotely exploitable weaknesses.

Having a prioritized list of your assets and a prioritized list of the vulnerabilities they face helps you better assign the resources available to start countering the threats.

IBM Tivoli Application Discovery and Dependency Manager

One option for discovering assets in your network is to use a discovery tool for finding assets. The optimal solution is a discovery tool that does not require an agent or any software installed on the target device. IBM Tivoli Application Dependency and Discovery Manager² provides agentless discovery of assets and also builds dependency maps of servers, networks, and applications and how they are interconnected. IBM Tivoli Application Dependency and Discovery Manager has three levels of discovery:

1. Discovery using NMAP and IBM technology (known as the stack scan sensor)

This level of discovery can fingerprint all the systems it finds on a network either by IP address, by range, or by indicating a subnet to discover. This data is then loaded into the IBM Tivoli Application Dependency and Discovery Manager relational database and can be exported in CSV format for use in IBM Security SiteProtector.

2. Discovery using operating credentials

This second level of discovery in IBM Tivoli Application Dependency and Discovery Manager allows the software to log into the servers and network devices as though they were a user using SSH login, Telnet, or using SSH keys; **sudo** is also supported. By running basic operating system commands such as **ls**, **netstat**, and **rpm**, IBM Tivoli Application Dependency and Discovery Manager is able to take a complete inventory of the server or network device and store that information in the database. In addition, IBM Tivoli Application Dependency and Discovery Manager is able to discover dependencies between applications on the same server, and across the network by resolving the IP addresses and the applications that are running on those servers.

3. Discovery using application credentials

This third level of discovery in IBM Tivoli Application Dependency and Discovery Manager allows the deep discovery of application configuration information such as IBM WebSphere® Application Server configuration, Oracle WebLogic Configuration, IBM DB2®, Oracle, and Sybase database configuration data, and database layout. To perform this level of discovery, read/only application server or database credentials are supplied to the IBM Tivoli Application Dependency and Discovery Manager server.

² To discover more about Tivoli Application Dependency and Discovery Manager, go to <http://www.ibm.com/software/tivoli/products/taddm/>. Make sure you also consult *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1*, SG24-7616.

Although the IBM Tivoli Application Dependency and Discovery Manager first level of discovery is sufficient for IBM Security SiteProtector, the ability of IBM Tivoli Application Dependency and Discovery Manager to perform deep discovery, dependency mapping, and versioning of all the configuration data that it finds is additional functionality that can greatly enhance the security architect's ability to understand the network and its vulnerabilities. Using this information, security architects can understand changes as they occur in the network and see how outages in the network can affect critical infrastructure.

IBM Rational AppScan

IBM Rational AppScan³ enables comprehensive and automated testing of web applications for vulnerabilities. It can also scan your websites for embedded malware or links to malicious sites. IBM Rational AppScan Version 7.8 or later can fully integrate with IBM Security SiteProtector Version 2.0 SP8 and later using a no cost extension. The Security SiteProtector Publisher extension for IBM Rational AppScan allows you to publish Rational AppScan results to Security SiteProtector.

With this broader view, you can get a more complete understanding of your security posture, so that you can better understand the risks that your organization faces and set priorities accordingly. SiteProtector can, for example, use this type of vulnerability data to help you focus on attacks targeting the vulnerabilities it discovered while lowering the priority of attacks that target non-vulnerable assets. We discuss this topic in the “SecurityFusion Module” on page 211.

Most events in SiteProtector are reported for assets, which are usually servers. However, web application security issues are detected on more granular elements, such as a specific URL, and, possibly, even a specific parameter or cookie. To support greater granularity, SiteProtector has been enhanced in Version 2.0 SP8 and later to ensure that the new events are marked as distinct events.

If you prefer to continue looking at the web application security issues in the context of the application being scanned, SiteProtector Version 2.0 SP8 and later also provides new *analysis views* that display the events in a view similar to the AppScan user interface.

³ You can find more information about the IBM Rational AppScan products at <http://www.ibm.com/software/awdtools/appscan/>. Make sure you also consult *Improving Your Web Application Software Development Life Cycle's Security Posture with IBM Rational AppScan*, REDP-4530.

7.2.2 Modifying technical security controls

Even after making sure you abide by your own security standards, configuration requirements, and patch management guidelines, you still have to add or modify additional security controls. Not every vendor can provide a patch at the time of disclosure or even discovery of a vulnerability in their software. There is a need to add additional layers of security. You most likely want to add additional protection by putting in place preventive and detective controls in the form of Network Intrusion Prevention System appliances (for more details, see Chapter 8, “Network security solutions” on page 243) and run Host Intrusion Prevention System agents (see Chapter 9, “Host security solutions” on page 299) on your physical and virtual servers that can inspect all network traffic directed towards them.

A centralized management system allows you to manage all necessary changes in network, host, and endpoint security policies from one central point.

7.2.3 Monitoring threats

A final critical task your centralized management system allows you to perform is monitoring attacks. It can filter out *false positives* created by valid custom applications you have running, which may violate RFCs and other Internet standards, and as a result trigger an alert. You can detect external and internal attackers and respond by blocking them from accessing your systems. And you can detect compromised hosts in your own organization so they can be separated from the rest of the network to be checked, cleaned, or forensically analyzed.

False positive and false negative: A *false positive* in this context is an event or alarm generated by a protective agent (such as a Network Intrusion Prevention System) when no real attack has taken place. The protective agent has incorrectly categorized benign traffic as malicious.

A *false negative*, however, occurs when an attack has taken place but the protective agent has not identified it as such and as a result failed to generate an alert.

All this is possible because a centralized management system bundles asset and vulnerability information, allows you to modify security policies, monitor security alerts generated by your protective agents, and respond to them by staffing incident escalations and reports.

7.3 IBM Security SiteProtector overview

IBM Security SiteProtector is a centralized management system that consists of multiple components, as shown in Figure 7-1 on page 207.

When you install SiteProtector, you need to decide what installation option you want:

- ▶ *Express* installation (all components installed on one host)
- ▶ *Recommended* installation (components installed on a minimum of two hosts)

In a default *recommended* installation, one host is used for the SiteProtector Database and the Event Collector. All other essential SiteProtector components are then installed on a second host.

Stay flexible: Most organizations should choose a *recommended* installation, because it offers better performance and more flexibility to expand later on when scalability and redundancy may become a factor.

7.3.1 SiteProtector components

To provide some insight into how the SiteProtector platform works and to illustrate the scalability of deployments, let us discuss the different components that make up SiteProtector:

- ▶ Event Collector
- ▶ Database
- ▶ Application Server
- ▶ Sensor Controller
- ▶ Console
- ▶ Agent Manager
- ▶ X-Press Update server

The components in this list are present in all SiteProtector deployments and are used in daily operations. They are also schematically depicted in Figure 7-1 on page 207.

Some other, often optional, components can be part of SiteProtector as well. We make sure to include them in the sections below. Some of these components are only present depending on the type of SiteProtector installation you prefer.

SiteProtector is available in a hardware appliance form or you can opt to go with a software installation on your own server hardware. The main benefit of the hardware SiteProtector appliance is the reduced time and effort required to deploy and maintain a SiteProtector installation. The software deployments of SiteProtector have the advantage of being more scalable for larger networks.

Deployment Manager

To roll out a software version of SiteProtector, the first step is to install the Deployment Manager. The Deployment Manager (not depicted in Figure 7-1 on page 207) is a local web server that lets you download all the software packages needed to install all the other SiteProtector components and various IBM Security agents. The Deployment Manager is not used during daily operations afterwards.

How to connect to the Deployment Manager: The Deployment Manager uses Apache HTTP Web Server Version 2.0, which is installed as part of the Deployment Manager installation. Once installed, the Deployment Manager is reachable at:

`https://deployment-manager-ip:3994/deploymentmanager/`

In a typical SiteProtector deployment, shown in Figure 7-1, the Sensor Controller is at the heart of the command and control center while the Event Collector plays the central role in the Event Channel, and the X-Press Update server provides updates to all other components.

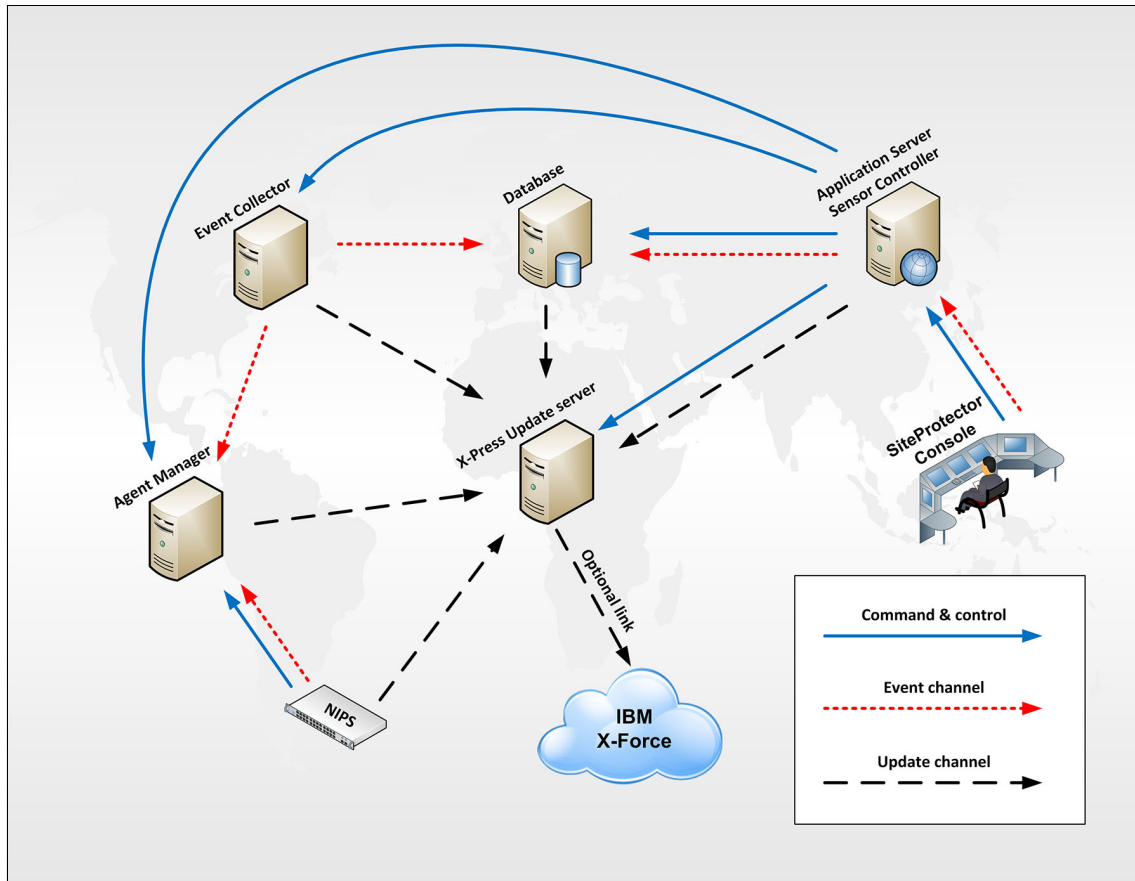


Figure 7-1 SiteProtector component deployment

Event Collector

The Event Collector gathers the data detected by agents in *real time* and directs it to the SiteProtector Database for storage. The agents send event data either directly to the Event Collector, or through an Agent Manager to the Event Collector. All event data is sent in real time.

If communication is lost between the Event Collector or Agent Manager and an agent, a cache of events is maintained on the agent until communication is restored.

Maximum numbers: You can install up to 255 Event Collectors on one site.

SiteProtector Database

The SiteProtector Database stores the agent data gathered by the Event Collector. The SiteProtector Database uses a Microsoft SQL Server database.

Tip: The SiteProtector Database can also be installed on an SQL Server Cluster to achieve high availability with MS SQL Server and support operations 24x7.

Application Server

The Application Server facilitates communication between the SiteProtector Console and other SiteProtector components and agents. The Application Server makes it possible for multiple SiteProtector Consoles to communicate with the SiteProtector Database, and to monitor and manage the same set of Event Collectors and agents.

Sensor Controller

The Sensor Controller is at the heart of any SiteProtector deployment. The Sensor Controller manages the *command and control* activities of agents and the other SiteProtector components, such as the command to start or stop collecting events.

Another term: The Application Server combined with the Sensor Controller is also referred to as the *SiteProtector Core*.

Agent Manager

The Agent Manager facilitates *command and control* between the SiteProtector Console and various agents. Agents *heartbeat* into an Agent Manager on a configurable schedule to obtain policy and configuration changes. The Agent Manager also mediates the transmission of event data between various agents and the Event Collector.

Workload and availability: In most deployments, multiple Agent Managers are installed to spread the workload and provide high availability.

X-Press Update server

An X-Press Update server facilitates the downloading and delivery of component and agent updates.

Update servers can download updates automatically from the IBM website or you can choose to disable this feature and isolate them from the outside world. You would then simply obtain the update packages on a different host and, after they are scanned and approved, copy these updates over to your local X-Press Update server.

After the update packages are available locally, the X-Press Update servers then work in conjunction with the Application Server to make the updates available to the SiteProtector components and agents in your network.

SiteProtector Firmware

SiteProtector Firmware facilitates updates to the hardened operating system, SQL Server software, BIOS, and drivers installed on the IBM Security SiteProtector appliance.

Appliance only: This component applies only to the IBM Security SiteProtector appliance shown in Figure 7-2. It is not relevant if you opt for a server-based deployment of the SiteProtector software.




Figure 7-2 IBM Security SiteProtector appliance SP2001 model

Event Archiver

The Event Archiver is a stand-alone SiteProtector component that archives event data in flat text files on a system that is separate from the SiteProtector Database. Flat text files are considered to be live data. By default, a new file is created each hour. The file name contains the time and date the events were archived.

By off-loading events to one or more Event Archivers, you can purge the SiteProtector Database more frequently, which reduces the number of events stored in the database and improves performance. You can also define filters to distinguish different types of event data and send them to separate Event Archivers.

Most importantly, Event Archiver can help meet your compliance regulations without overloading the SiteProtector Database. Event Archiver can, for example, help comply with Sarbanes-Oxley requirements, because it offers a simple way to maintain live data for years.

Be aware: The Event Archiver is optional and is not installed by default.

SiteProtector Console

The SiteProtector Console can manage multiple SiteProtector installations and all the IBM Security appliances, agents, and scanners deployed across the network. You also can use the Console to monitor and analyze event data (as shown in Figure 7-3) and perform maintenance tasks on the SiteProtector Database. Some of the other functions you can perform from a Console include:

- ▶ Monitor the status of components, appliances, agents, and scanners.
- ▶ Create, save, and print analysis views.
- ▶ Monitor and filter event alerts.
- ▶ Run vulnerability scans.
- ▶ Manage appliance, agent, and scanner policies.
- ▶ Generate and schedule reports.
- ▶ Configure SiteProtector Database maintenance options.
- ▶ Create SiteProtector tickets for assets and events.

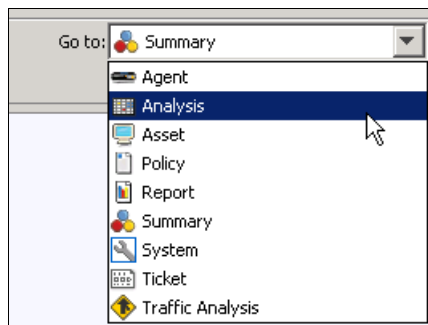


Figure 7-3 Different functional views available in a SiteProtector Console

SiteProtector web portal

Technical note: A SiteProtector console currently (SP 8.1) requires a 32-bit or 64-bit Windows operating system and requires a Java™ Runtime Environment to run.

The SiteProtector web portal is a web-based, *read-only* console that provides remote access to your SiteProtector data. The web portal gives you access to the same data views and filters that are available in the SiteProtector Console, allowing you to:

- ▶ Monitor SiteProtector assets and network assets.
- ▶ View SiteProtector groups.
- ▶ View security events, incidents, and exceptions, including event details.
- ▶ Generate and view reports.

SecurityFusion Module

The SiteProtector SecurityFusion Module correlates event data to identify relationships between intrusions and vulnerabilities. The *impact analysis* module allows you to filter out alerts where the target asset is not vulnerable to a certain attack. The SecurityFusion Module can:

- ▶ Reduce the number of security events displayed in the Console.
- ▶ Lower or raise the severity of security events based on vulnerability information.

For example, if a scan using IBM Rational AppScan (discussed in “IBM Rational AppScan” on page 203) reveals that your web application running on a specific host contains a vulnerability, the SecurityFusion Module can raise the severity level of a security event generated by a Network Intrusion Prevention System appliance when it sees an attack was targeting this vulnerability. Similarly, hosts scanned and proven to not be vulnerable can have the severity of specific security event lowered when they are targeted. All this is configurable in the SecurityFusion Module.

Be aware: A SP2001 SiteProtector appliance has SecurityFusion pre-installed. If you opt for a *recommended* software installation of SiteProtector, you must install SecurityFusion separately using the Deployment Manager.

7.3.2 Appliance and agent support

SiteProtector supports, among others, the following appliances and agents:

- ▶ IBM Security Network Intrusion Prevention System appliances
An *inline* intrusion prevention platform that monitors and blocks malicious traffic to protect your network from attacks. It is discussed in Chapter 8, “Network security solutions” on page 243.
- ▶ IBM Security Virtual Server Protection for VMware
Protects your virtualized infrastructure at the hypervisor level. It is discussed in Chapter 10, “Virtual server security solutions” on page 337.
- ▶ IBM RealSecure Server agent
Monitors network traffic to the host, as well as host-based operating system logs, for suspicious activity. It is discussed in Chapter 9, “Host security solutions” on page 299.
- ▶ IBM Security Server and Desktop Endpoint Security agents
Uses firewall, intrusion prevention, and application control technologies to block improper activity at the host level. It is discussed in Chapter 9, “Host security solutions” on page 299.

7.3.3 SiteProtector communication channels

We distinguish three types of communication in SiteProtector:

- ▶ A command and control channel that allows SiteProtector to control other protective agents, such as Network and Host Intrusion Prevention Systems.
- ▶ An event channel where SiteProtector gathers all security event data from Network and Host Intrusion Prevention Systems and stores it centrally.
- ▶ An update channel in which SiteProtector and the protective agents make sure they are loaded with the latest features and content updates.

Command and control channel

At the heart of the command and control channel is the Sensor Controller. Authenticated users contact the SiteProtector Core (Application Server + Sensor Controller) through their Console. They can issue commands to change policies, update an agent, adapt device interface settings, and so on. The Sensor Controller then either contacts the affected protective agents directly or uses an intermediary Agent Manager to make sure the agents get the instructions issued by the authorized user. This process is shown in Figure 7-4.

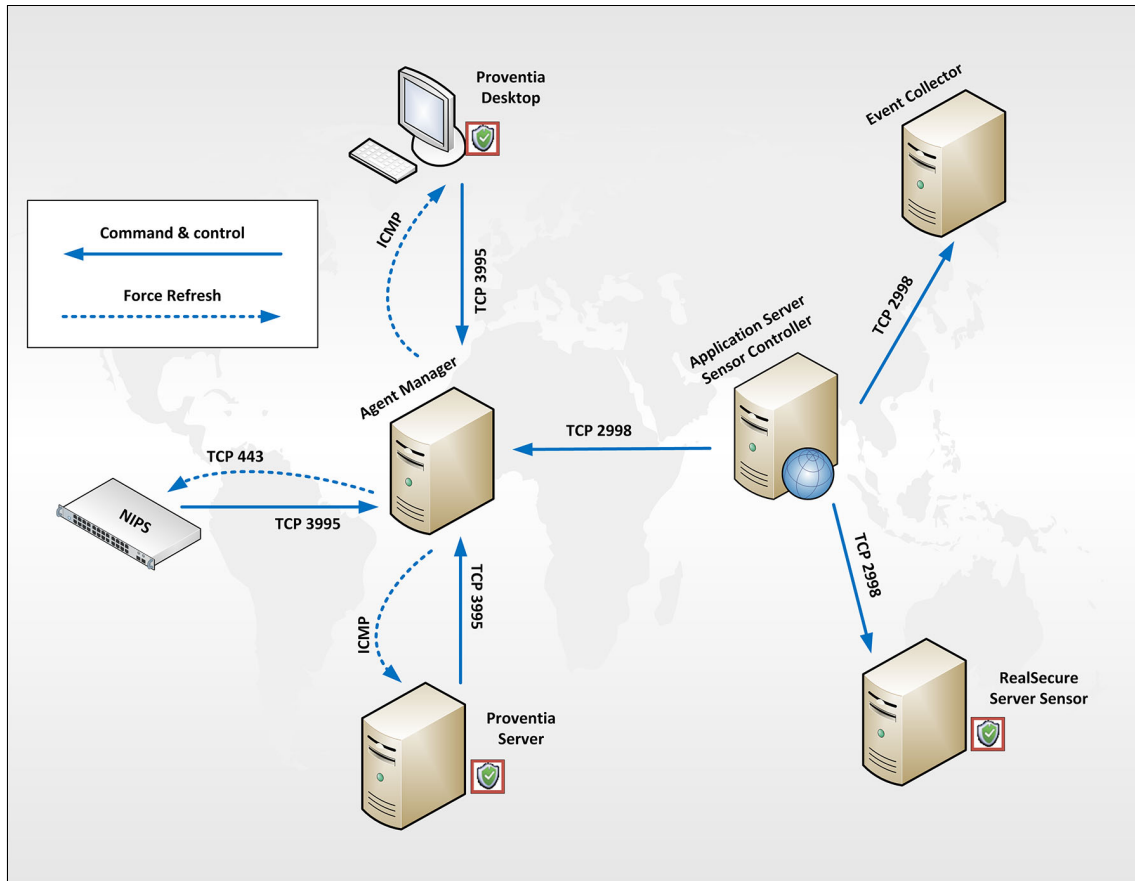


Figure 7-4 Command and control channel communications

Remember: All IBM Security agents (software and appliances) only communicate with SiteProtector through intermediary Agent Managers. Each agent can have multiple Agent Managers listed; if one of them is unavailable, then the agent moves to another Agent Manager. Existing RealSecure agents do not use Agent Managers.

Event channel

Agents (Network Intrusion Prevention System appliances and Host Intrusion Prevention System software agents) generate security events. These security events are collected by one or more Event Collectors and stored in the central database, as shown in Figure 7-5. There they can be queried and processed further. The Event Collector either contacts the agent directly or contacts an intermediary Agent Manager.

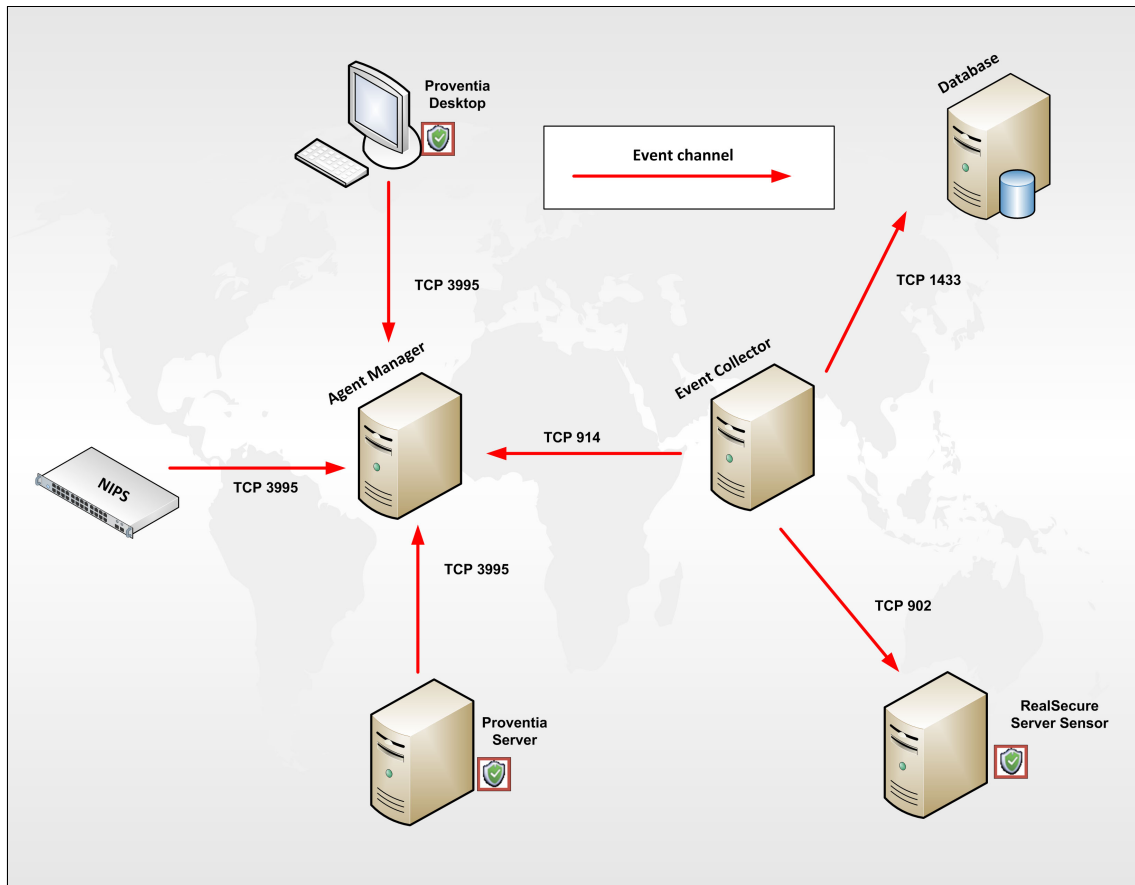


Figure 7-5 Event Collector communication

Update channel

Every SiteProtector installation comes with at least one X-Press Update server. Agents (software and hardware appliances) get their updates from an X-Press Update server. The ability to run your own X-Press Update servers allows you to proxy the required update packages locally and save on bandwidth consumption in case you have many agents requiring the same updates.

As already mentioned in “X-Press Update server” on page 208, you can choose to allow your update servers to directly download the update packages from IBM servers, but this is not a requirement. You can always choose to maintain an “air gap” between your SiteProtector and the outside world. Organizations can opt to control how update packages can be brought into their network and only install them on the update servers after following their own internal scanning and auditing requirements.

In Figure 7-6, we show a scenario where two X-Press Update servers are used for load-distribution and to make sure the default X-Press Update server, which is installed on the same host as the Application Server, does not initiate a direct connection to the Internet.

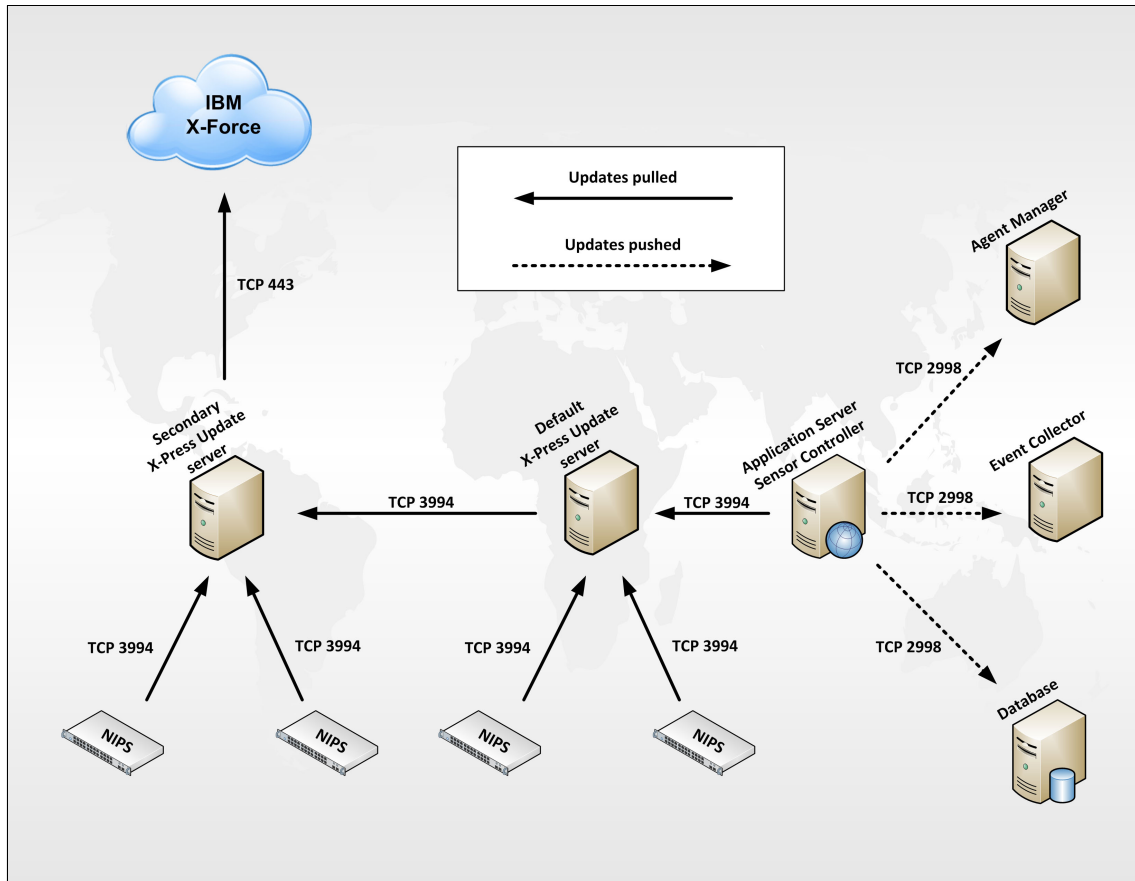


Figure 7-6 Deployment example with multiple X-Press Update servers

It is your choice: IBM offers a Manual Upgrader Tool for SiteProtector, which enables you to select and download all update packages available from the IBM owned X-Press Update servers to a separate system in your organization where you can first scan and audit these packages before bringing them into a more secure zone where your SiteProtector would be located.

7.3.4 Data redundancy

Several components in SiteProtector offer a level of built-in redundancy.

As already discussed in “SiteProtector Database” on page 208, you can install the SiteProtector Database on an SQL Cluster, and in 7.3.3, “SiteProtector communication channels” on page 212, we mention that agents can use multiple Agent Managers. However, three other SiteProtector redundancy mechanisms warrant some mentioning as well:

- ▶ Event Collector stacking
- ▶ Event Collector partnering
- ▶ SecureSync

Event Collector stacking

Event Collector *stacking* (shown in Figure 7-7) allows you to configure an Event Collector to send the events it receives to a second Event Collector. This feature is useful if you have an additional SiteProtector installed in a separate business unit or remote office and direct all events from that entity to the primary SiteProtector site via a single, consolidated event stream.

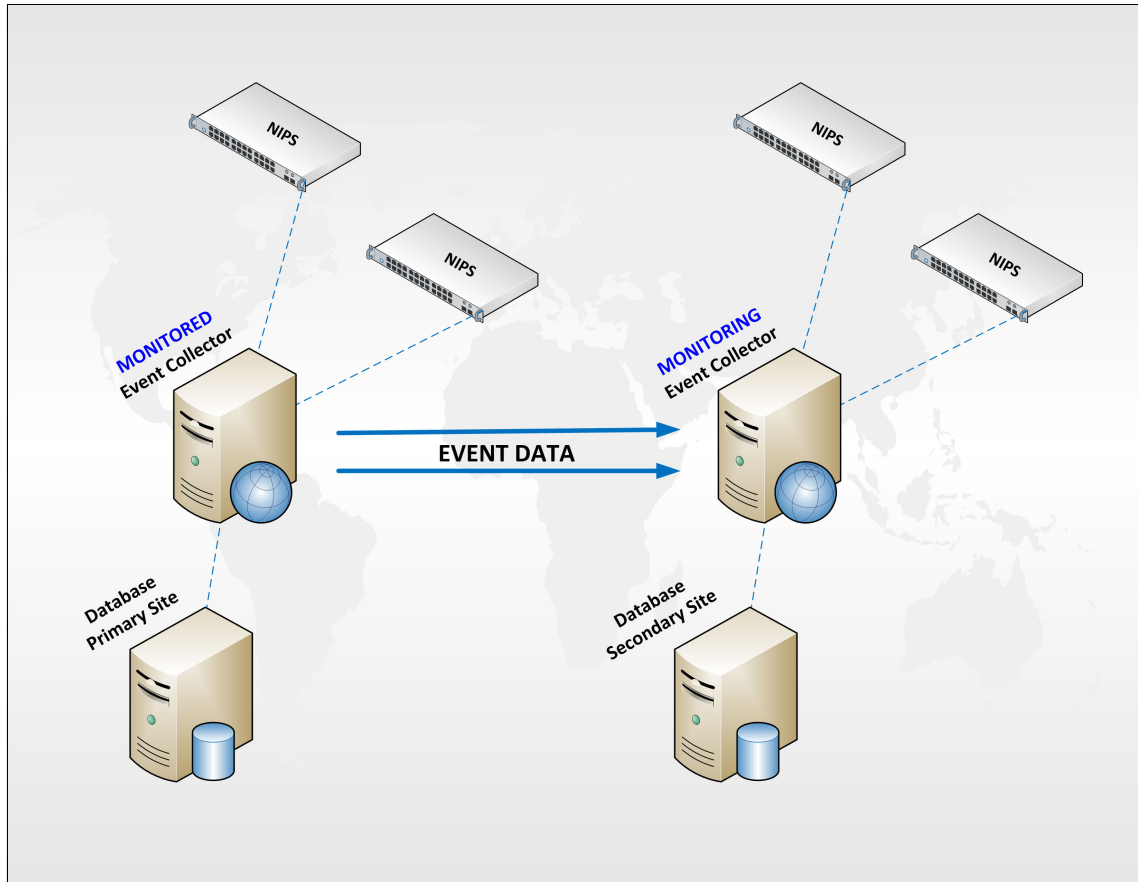


Figure 7-7 Event Collector stacking

When configuring Event Collector stacking, you need to distinguish between the two Event Collectors you are stacking:

- ▶ *Monitored* Event Collector: The monitored Event Collector sends events to the monitoring Event Collector.
- ▶ *Monitoring* Event Collector: The monitoring Event Collector receives events from the monitored Event Collector.

Event Collector partnering

To ensure that you do not lose data if an Event Collector goes offline, you may want to use an Event Collector *partnering* strategy. With this *failover* mechanism activated, you have a way to automatically switch between primary and secondary Event Collectors. If the primary Event Collector shows a status of *stopped*, *error*, *unknown*, or *not responding*, then data is redirected to the secondary Event Collector in the partnership.

SecureSync

The SiteProtector SecureSync process allows you to transfer management and data collection responsibilities to another SiteProtector instance in the event of catastrophic failure, network outage, or disaster affecting a site. This failover process ensures that you do not lose important data if a SiteProtector Core server becomes unavailable.

You can also transfer operations back to a primary SiteProtector Core server when it is once again operational. This process is referred to as failback.

When you fail over or fail back, SecureSync redirects the following items:

- ▶ Event Collector(s), including agents that report to the Event Collectors
- ▶ Agent Manager(s), including agents that report to the Agent Managers
- ▶ SecurityFusion Module

Be aware: SecureSync functionality requires a separate license.

For more detailed information about how to configure Event Collector stacking, partnering, or SecureSync, refer to the respective configuration guides at:

<http://www.ibm.com/support/>

7.3.5 Authentication and encryption

Users can authenticate using their SiteProtector console in several ways. The SiteProtector two-factor authentication feature provides a plug-in interface that supports any authentication software you use and already provides specific plug-in interfaces for RADIUS and LDAP certificate authentication.

When one SiteProtector component communicates with another component, authentication is also used, and in the case of IBM Security products, it relies on one of the following mechanisms:

- ▶ SSL certificates (used, for example, between an IBM Security Network Intrusion Prevention System and an Agent Manager).
- ▶ A Public/Private key pair created by the cryptographic provider selected when you install the component (used, for example, between Sensor Controller and Event Collector).

Technical note: The IBM Security SiteProtector Cryptographic Module has received Federal Information Processing Standard (FIPS) 140-2 certification.

More information about this standard can be found at:

<http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140crt/140crt1402.pdf>

7.3.6 Separation of duties and auditing

The SiteProtector system lets you log almost all actions that are performed in the SiteProtector system and lets you track this activity for auditing purposes. Pre-formatted reporting is available as well.

A record appears in the Audit Detail report for each action that is logged by the SiteProtector system if the specified action was performed. Audit records typically contain the following information:

- ▶ The type of action
- ▶ The time and date an action occurred
- ▶ The user or SiteProtector system component that performed the action
- ▶ Location where the action was performed

The permissions associated with each user or group enable or restrict individual users in the tasks they can perform while logged on.

By default, SiteProtector provides the following user groups:

- ▶ Administrator
- ▶ Analyst
- ▶ Operator
- ▶ Assessment Manager
- ▶ Desktop Manager
- ▶ Network Manager
- ▶ Server Manager

You can also add new user groups to accommodate your specific needs. You can, for example, define certain groups of users who can only perform specific tasks on a limited set of specific devices. Some users may only be allowed to run reports, others would be allowed to modify policies for a specific group of Network Intrusion Prevention System appliances, and yet another team would be allowed to apply the prepared changes.

7.4 Managing operational security in SiteProtector

In this section, we discuss several security capabilities that the IBM Security SiteProtector product offers.

To understand how the security capabilities of SiteProtector can be mapped to the IBM Security Blueprint⁴, see Figure 7-8 on page 222. This diagram shows the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using SiteProtector. This functional highlighting is applicable for the infrastructure service components as well.

Besides the fully highlighted elements, Figure 7-8 on page 222 also shows some medium highlighted elements. Although SiteProtector can be used to address such a component to some degree, the respective area of coverage is not considered a core function of the product and thus is considered to be limited.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 7-8 on page 222 can be used as a quick reference of the functional security management aspects of SiteProtector. This reference can help us determine which functions of a solution can be covered by selecting this product.

⁴ For a detailed discussion of the elements, see Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

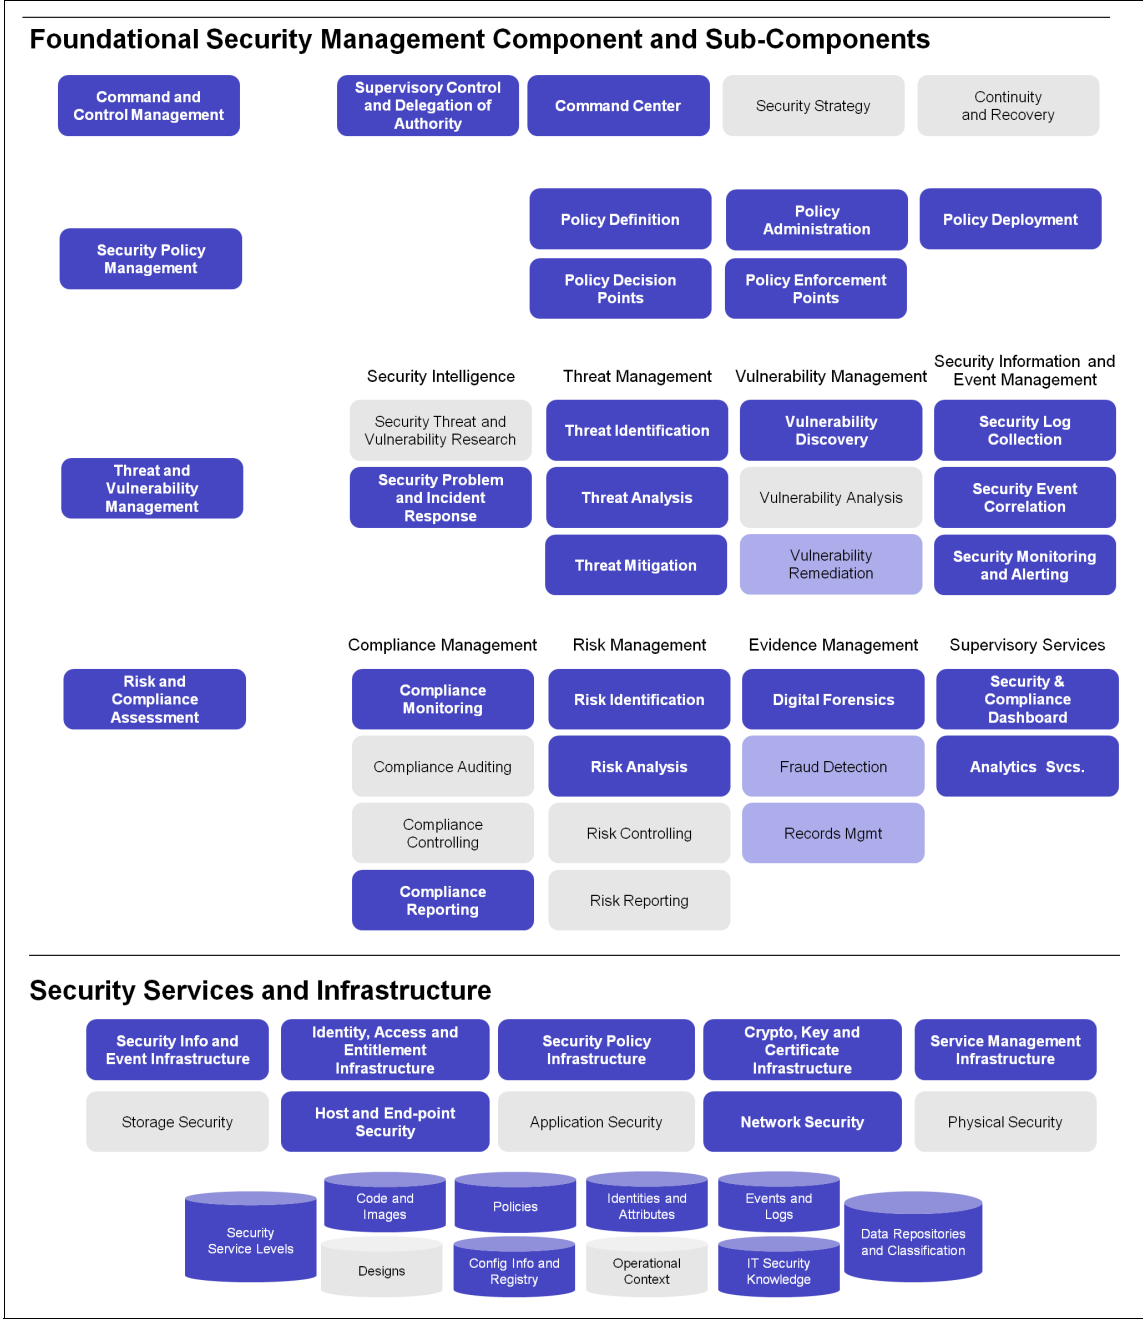


Figure 7-8 Mapping of IBM Security SiteProtector to the IBM Security Blueprint

7.4.1 Managing assets

The importance of having an accurate asset list in your centralized management system was discussed in 7.2.1, “Asset and vulnerability prioritization” on page 201.

There are several methods you can use to populate SiteProtector with asset data:

- ▶ Gather host information using agents.
- ▶ Add host information manually.
- ▶ Use a scanning agent.
- ▶ Import host information from an Active Directory container.

SiteProtector allows you to organize your network assets and your SiteProtector components and agents into asset groups to make it easier to manage and monitor the security of your network. For example, you can create groups to:

- ▶ Apply agent policies.
- ▶ Apply X-Press Updates.
- ▶ Analyze event data.

Within each group, you also can create subgroups to further organize your SiteProtector assets and network hosts.

Using effective grouping strategies can help you align SiteProtector functionality with your company’s organization and security processes. The following categories of activity serve as the primary strategic focus for grouping SiteProtector assets and network assets:

- ▶ *Command and control*: These are tasks that relate to controlling SiteProtector functionality, as well as the operation of agents.
- ▶ *Data analysis*: These are tasks that relate to analyzing event information.

Organizational tip: The more granular your groups, the more command and control you gain over those groups.

You may organize groups and subgroups using any criteria your organization requires. Some common organizational strategies include grouping SiteProtector assets and other network assets by:

- ▶ Geography
- ▶ Topology
- ▶ Services
- ▶ Business function
- ▶ Scope of responsibility

An example of a grouping strategy is shown in Figure 7-9, where groups organized by geography are at the top of the hierarchy, with subgroups organized by services at lower levels of the hierarchy.

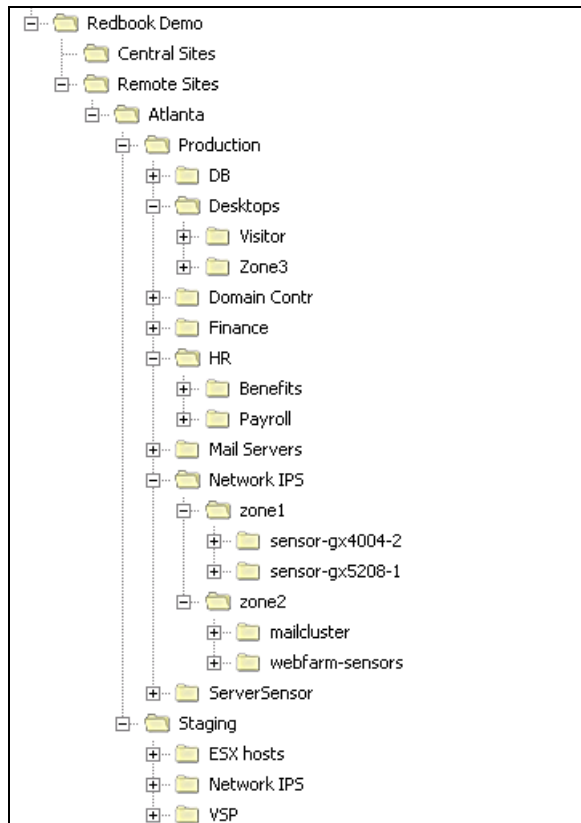


Figure 7-9 View of an asset and agent grouping structure in the SiteProtector Console

7.4.2 Managing policies

In this section, we discuss the powerful hierarchical model that SiteProtector uses to assign policies and policy elements. We explain how SiteProtector keeps track of policy changes, how you can compare policies, and how you can manage, deploy, and roll back several versions of a given policy.

Hierarchical policy model

SiteProtector uses a *hierarchical* inheritance model to manage policies across all of the asset groups in a site. In this environment, you can apply a single, distributed policy element to multiple agents and groups.

In Figure 7-10, you can see how policies in SiteProtector use the asset tree structure. The group zone1 has a Security Events policy directly assigned to it (indicated by the red arrow). The Security Events policy assigned to the group zone1 is the second version of a policy labeled Redbook Demo. As a result, all agents in this group zone1 use this policy. The same is true for its subgroups, sensor-gx-4004-2 and sensor-gx5208-1; they too use this same policy unless they get their own different policies assigned directly to them.

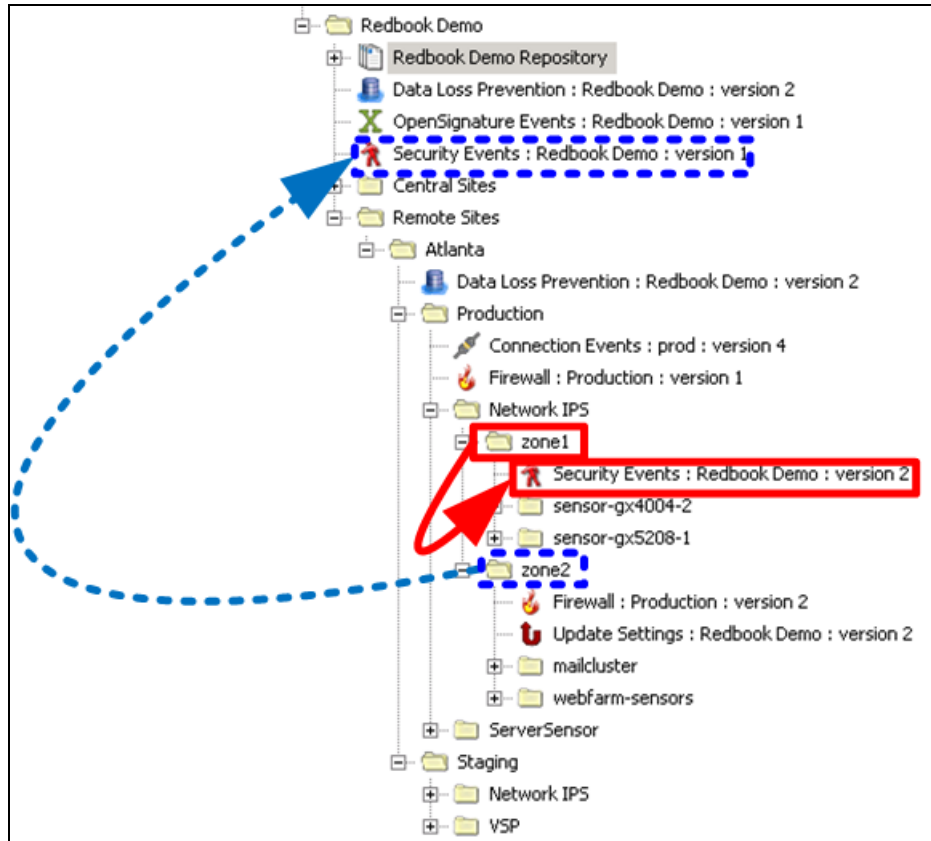


Figure 7-10 View of a policy hierarchy in the SiteProtector Console

Remember: A group always uses a policy directly assigned to it. Only when it does not have a policy assigned to it will it inherit a policy from another group higher up the tree structure.

Policy inheritance is also shown in Figure 7-10 on page 225, as indicated by the blue arrow. The group zone2 has a Firewall and an Update Settings policy assigned directly to it, but it does not have a Security Events policy assigned to it. To determine which Security Events policy is used by group zone2, we have to go up the hierarchy. When we go up one level to the group Network IPS, we see that this group also does not have a Security Events policy assigned to it. We have to look even higher up. The same is true for the groups Atlanta and Remote Sites. Only when we reach the group Redbook Demo do we find that it has a Security Events policy assigned to it. Because of the inheritance model, group zone2 (and all agents belonging to it) inherits this policy and uses the first version of the Security Events policy labeled Redbook Demo.

The main advantage of having the ability to use *policy inheritance* is the fact that you can share policies among several devices while maintaining sufficient flexibility to tailor individual devices or software agents.

Often you find that many configuration elements are the same throughout your organization. You can use the same type of update settings for all your IBM Network Intrusion Prevention System appliances, but give them a different type of response to attacks depending on their location in the network.

Remember: You can also assign policies to specific agents and not just to groups. Agent-specific policies are not shared among devices.

Policy repositories

A policy repository, shown in Figure 7-11, is a work space that you can use to store, modify, and deploy the hierarchical agent policies used in your organization. To maximize your control over hierarchical policies, SiteProtector allows you to create a separate *repository* for any group.

Default Repository						
Right-click on a column header to filter.						
Policy Type	Name	Latest Version	Agent Version(s)	Agent Mode(s)	Last Modified	Last Modified By
Alerts	Default	2	4.1	All	2010-11-18 21:59:30 GMT	jvanherzele
Authentication Servers	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Connection Events	Default	2	4.1	All	2010-11-18 22:02:00 GMT	jvanherzele
Data Loss Prevention	Default	2	4.1	All	2010-11-18 21:59:01 GMT	jvanherzele
Firewall	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Group Settings	Default	5	4.1,3.3,3.2,3.1,3.0,2.5,...	All	2010-11-18 22:01:11 GMT	jvanherzele
OpenSignature Events	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Response Filters	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Rolling Packet Captur...	Default	4	4.1	All	2010-11-18 21:59:19 GMT	jvanherzele
Security Events	Default	1	4.1	All	2010-06-11 13:47:18 BST	
SNMP	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Tuning Parameters	Default	2	4.1	All	2010-07-19 16:16:18 BST	Internal System
Update Settings	Default	1	4.1	All	2010-06-11 13:47:18 BST	
User Defined Events	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Web Application Prot...	Default	1	4.1	All	2010-06-11 13:47:18 BST	
Version History for 'Connection Events : Default'						
Version	Status	Created On	Created By	Comment		
2	Archived	2010-11-18 22:02:00 GMT	jvanherzele	enabled FTP detection		
1	Archived	2010-06-11 13:47:18 BST				

Figure 7-11 View of a Default Repository showing two versions of a policy

There are two types of repositories:

- ▶ A Default Repository

This repository, which is associated with the top-level group, contains the default policies installed by SiteProtector. The Default Repository allows you to push policies from the top-level group down through all the other groups.

- ▶ Group-level repositories

Group-level repositories represent a policy work space that you can create for a group in the tree structure. This type of repository takes the name of the group with which it is associated, and contains any new policies you create for the group. Group-level repositories allow you to push hierarchical policies from the local group down through any child groups.

Repositories facilitate policy management in several ways. For example, you can use repositories to create and safely modify new versions of active production policies, without having to directly configure your active policies. Then after thoroughly reviewing the revised policies, you can *deploy* them to your production environment.

Policy repositories also help you:

- ▶ Track the deployment of your policies.
- ▶ Archive previously deployed policies.
- ▶ Reapply previously deployed policies, if necessary.

Policy versioning

You can open and edit a policy at any time, even if the policy has been deployed. Whenever you edit and save a policy, SiteProtector always saves a new version of the policy. The original policy version is unaffected by your changes. If you want to apply your changes to an agent, you must deploy the new version of the policy.

After you modify a policy in a repository, you may want to deploy the policy. SiteProtector gives you the option to automatically deploy a policy when you save the policy, or you can deploy the policy at a later (scheduled) time.

Remember: When deploying policies from a repository, it is important to realize that you can deploy the policies only to the group associated with the repository and to its child groups.

Comparing policies

Since SP 8.0, SiteProtector also offers the option to run a *diff*, which is a line-by-line comparison of two policies. This feature and the policy reporting capability offers a way to quickly map the differences between policies that may contain thousands of Security Event signatures, each with a dozen parameters.

An example of a result of such a policy comparison is shown in Figure 7-12.

	Enabled	Protection...	Event Name	Severity	Protocol	Ignore Ev...	Display	Block	Log Evidence...
Attack/Audit: Attack (1 items)									
User Overridden: false (2597 items)									
	<input type="checkbox"/>	Global	DNS_Tunnel	Low	dns	<input type="checkbox"/>	Without Raw	<input type="checkbox"/>	None
	<input type="checkbox"/>	Global	LDAP_Respo	Low	ldap	<input type="checkbox"/>	Without Raw	<input type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	HTTP_IIS_A3	Medium	http	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	LDAP_LSASS	High	ldap	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	AVI_Cinepak	High	riff	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	HTTP_Lnk_F	Medium	url	<input type="checkbox"/>	Without Raw	<input type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	Content_TNE	High	tnef	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	HTML_MS_H	High	html	<input checked="" type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	HTTP_OpenV	High	url	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input type="checkbox"/>	Global	PE_WinVerif	High	pecoff	<input type="checkbox"/>	Without Raw	<input type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	Script_Java	High	js	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None
	<input type="checkbox"/>	Global	DNS_TCP_M	Low	dns	<input type="checkbox"/>	Without Raw	<input type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	Oracle_Too	Medium	tns	<input type="checkbox"/>	Without Raw	<input type="checkbox"/>	None
	<input checked="" type="checkbox"/>	Global	DataProtect	High	data_protect	<input type="checkbox"/>	Without Raw	<input checked="" type="checkbox"/>	None

Figure 7-12 A comparison between two policy versions (changes are marked by triangles)

7.4.3 Performing event monitoring and analysis

Because organizations must address a large number of security events each day, it is important that security administrators use the right tools and strategies to analyze these events. SiteProtector provides event analysis functions that are designed to support a straightforward methodology to help you perform security analysis more efficiently.

IBM recommends that you use the following analysis strategy when using SiteProtector:

1. View summary event information.

SiteProtector includes predefined Analysis views with summary counts of important data points. As shown in Figure 7-13, the “Event Analysis - Event Name” view displays all events, along with aggregated counts of sources, targets, and target objects (usually a port number). This gives you a broad view of the events affecting your network.

Event Analysis - Event Name (Agent) [Load View](#) [Save View](#) Data last

Data Filters (3 applied)

Time Filter

Today

Tag Name Filter

Source IP Filter

Target I

Right click on the column header to group by that column.

Tag Name	Severity	Event Count	Source Count	Target Count	Status
HTTP_Tivoli_Prov_Mgr_Malformed_Post	High	596	39	25	Detected
SQL_Injection	High	304	41	30	Detected
MSRPC_Svcctl_Remote_Control	High	29	3	3	Detected
IMAP_Tag_Overflow	High	16	1	8	Detected
Swf_RealPlayer_Frame_Overflow	High	15	10	10	Detected
HTML_DOMINO_Web_Access_Overflow	High	8	1	2	Detected
ASP_IIS_File_Change_Notification	High	6	1	1	Detected
HTTP_DotDot	High	4	2	1	Detected
HTTP_Cisco_Catalyst_Exec	High	2	1	1	Detected
Content_Compound_File_Bad_Extension	High	1	1	1	Detected
Smurf_Attack	Medium	8,477	1	1	Detected
Ping_Sweep	Medium	2,266	5	4	Detected
HTTP_URL_Name_Very_Long	Medium	2,240	200	142	Detected
TCP_Short_Header	Medium	1,721	4	3	Detected
XPATH_Injection	Medium	1,060	6	15	Detected
Stream_DoS	Medium	645	17	1	Detected
HTTP_IIS_Hex_Evasion	Medium	492	115	104	Detected
HTTP_Html_In_Ref	Medium	368	39	31	Detected

Figure 7-13 Event analysis overview in the SiteProtector Console

2. View high-level event details to determine importance.

To facilitate analysis, SiteProtector provides predefined Analysis views with detailed information, for example, the “Event Analysis - Details” view, as shown in Figure 7-14. Such views allow you to look more selectively at detailed event information, including sources, targets, target objects, other events directed at a target, vulnerabilities on a target, and the agents that detected an event.

Time *	Tag Name *	Event Co...	Status *	Severity *	Source IP	Event-type	Field	Value	Object
2010-11-11 14:3...	SQL_Injection	1	Detect...	High	10.74.1.61	Attack	unknown	exec says ...	Target
2010-11-11 19:0...	SQL_Injection	1	Detect...	High	10.0.108.23	Attack	query	SELECT na...	Target
2010-11-11 18:5...	SQL_Injection	1	Detect...	High	10.0.8.87	Attack	query	SELECT firs...	Target
2010-11-11 18:5...	SQL_Injection	1	Detect...	High	10.0.8.87	Attack	query	SELECT firs...	Target
2010-11-11 18:4...	SQL_Injection	1	Detect...	High	10.0.8.87	Attack	query	SELECT firs...	Target
2010-11-11 18:4...	SQL_Injection	1	Detect...	High	10.0.8.87	Attack	query	SELECT firs...	Target
2010-11-11 18:4...	SQL_Injection	1	Detect...	High	10.57.0.56	Attack	query	SELECT na...	Target
2010-11-11 18:4...	SQL_Injection	1	Detect...	High	10.36.1.83	Attack	query	SELECT firs...	Target
2010-11-11 18:4...	SQL_Injection	1	Detect...	High	10.36.1.83	Attack	query	SELECT firs...	Target
2010-11-11 18:2...	SQL_Injection	1	Detect...	High	10.42.25.53	Attack	query	select mess...	Target
2010-11-11 18:1...	SQL_Injection	1	Detect...	High	10.42.25.53	Attack	query	select mess...	Target
2010-11-11 18:1...	SQL_Injection	1	Detect...	High	10.26.0.5	Attack	query	select url,s...	Target
2010-11-11 18:1...	SQL_Injection	1	Detect...	High	10.42.25.53	Attack	query	select mess...	Target

Figure 7-14 Detailed view of a single type of security event

3. Create a prioritized subset of the event data.

Sort and filter the event data to focus on critical and high value assets first.

4. Perform additional filtering on the events you have prioritized.

Use *guided* analysis and other filtering options to focus on the most critical events.

5. Manually correlate events of undetermined importance.

If you cannot make a final determination as to the importance of an event, leave the event in your list so that you can manually correlate it with future events.

Guided analysis

The console provides context-sensitive event data that can help you access the detailed information required in your event analysis strategy. This event data is accessible through an event’s right-click menu, which includes a view-specific set of options presented in the form of *guided* questions.

When you access detail information through the contxt menu, as shown in Figure 7-15, the information does not appear in a separate window. Instead, it appears as a new data view within the current Analysis view. After you review the detail information, you can return to the previous parent view by clicking the **Back** toolbar button.

HTTP_repeated_		74
HTTP_Cross_Siti	What are the Virtual Infrastructure details?	50
ICMP_Unreacha		40
HTTP_URL_repe	What are the event details?	34
LanMan_Share_	What are the target objects of this event?	25
XML_EntityRef_	Which agents detected this event?	17
TCP_Port_Scan	What are the sources of this event?	17
SNMP_Default_E	What are the targets of this event?	12
HTTP_Fields_Wil	What are the ADS Event Details?	10
HTML_NullChar		10

Figure 7-15 Guided analysis options in the SiteProtector Console

Incidents and exceptions

After you have determined whether a specific event is important by using the Console’s predefined Analysis views and drill-down options, you can filter the event as appropriate. SiteProtector allows you to filter event data globally at the site level. These “global” filters affect all Analysis views, regardless of any other filters defined at the view level.

SiteProtector allows you to create two types of global filters:

- *Exception*: Used to filter information of questionable value, such as false positives.
- *Incident*: Used to combine and filter important information that is being duplicated unnecessarily.

Analysis views

The Console includes several predefined Analysis views that allow you to examine data from various perspectives, and at different levels of detail. These views can help you perform the first two steps of the event analysis strategy to review summary and detail information. You can access predefined views from the drop-down list in the Analysis view.

7.4.4 Reporting functionality

SiteProtector offers built-in reporting capabilities that allow you to either choose from predefined reporting templates, as shown in Figure 7-16, or you can define your own report templates with custom data filters and formatting options that meet your specific needs. If you want to create your own templates, you can install the Business Intelligence and Reporting Tools (BIRT).⁵

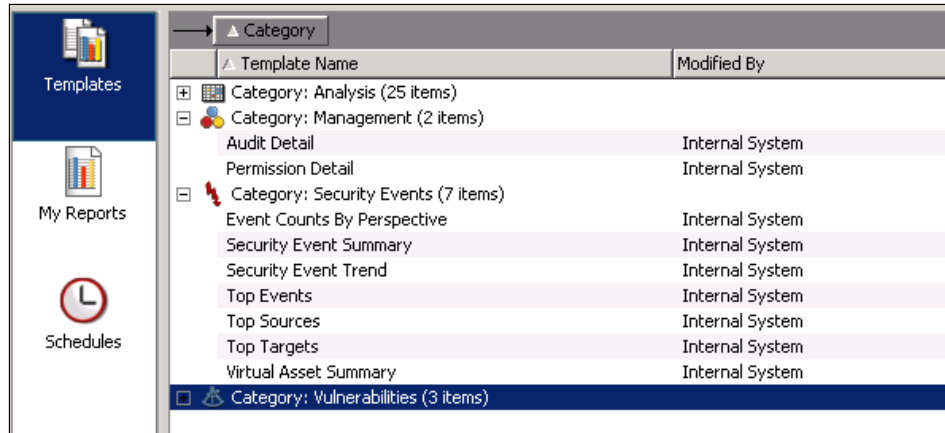


Figure 7-16 View of the default reports available in the SiteProtector Console

Be aware: SiteProtector requires an optional Reporting Module license to use the reporting options discussed.

7.4.5 Ticketing options

SiteProtector has built-in ticketing functionality that can help organizations to systematically track and manage tasks, issues, and requests submitted by SiteProtector users. It helps an organization ensure that:

- ▶ Security incidents are tracked and addressed.
- ▶ SiteProtector assets are properly managed and updated.
- ▶ Issues and tasks are directed to the appropriate teams.

⁵ For more information about BIRT, go to <http://www.eclipse.org/birt/phoenix/>.

You can create new tickets from context menus on the Analysis, Agent, and Assets views available in the SiteProtector Console, as shown in Figure 7-17.

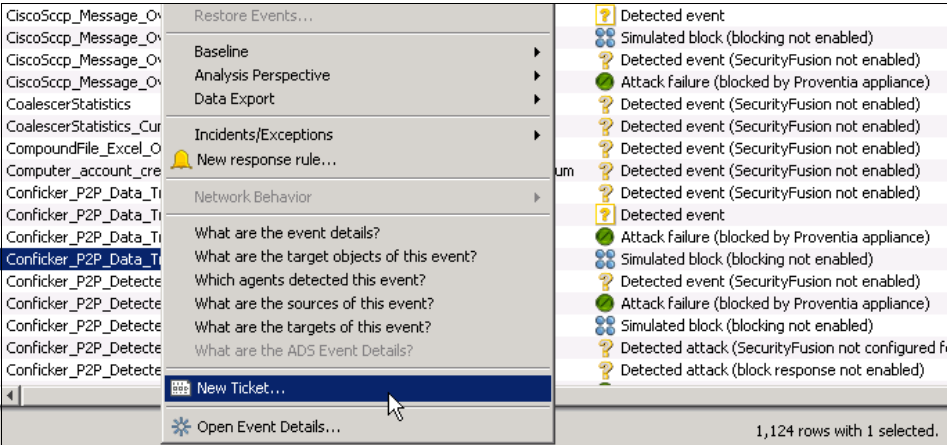


Figure 7-17 The SiteProtector console allows ticket generation through a context menu

SiteProtector also allows you to integrate SiteProtector ticketing with the Remedy Action Request System⁶.

SiteProtector and Remedy integrate at the server level, using a Remedy/Java application program interface (API). When you create SiteProtector tickets in this integrated environment, the tickets are saved to both the SiteProtector Database and the Remedy server.

Be aware: In the integrated SiteProtector/Remedy environment, you can create tickets in SiteProtector, but you can edit them only in Remedy. A copy of each ticket created in SiteProtector is saved to the SiteProtector Database.

⁶ For more information about Remedy, go to <http://www.remedy.com/>.

7.4.6 IBM Tivoli Security Information and Event Manager

Often, organizations need to expand their security event management system to include the ability to perform log monitoring, auditing, and reporting. IBM Tivoli Security Information and Event Manager (TSIEM) is the IBM SIEM solution, and is shown in Figure 7-18. Tivoli Security Information and Event Manager can manage and monitor logs from a variety of technology platforms and provide analytics that help the security manager to protect intellectual property, privacy, and support compliance mandates.

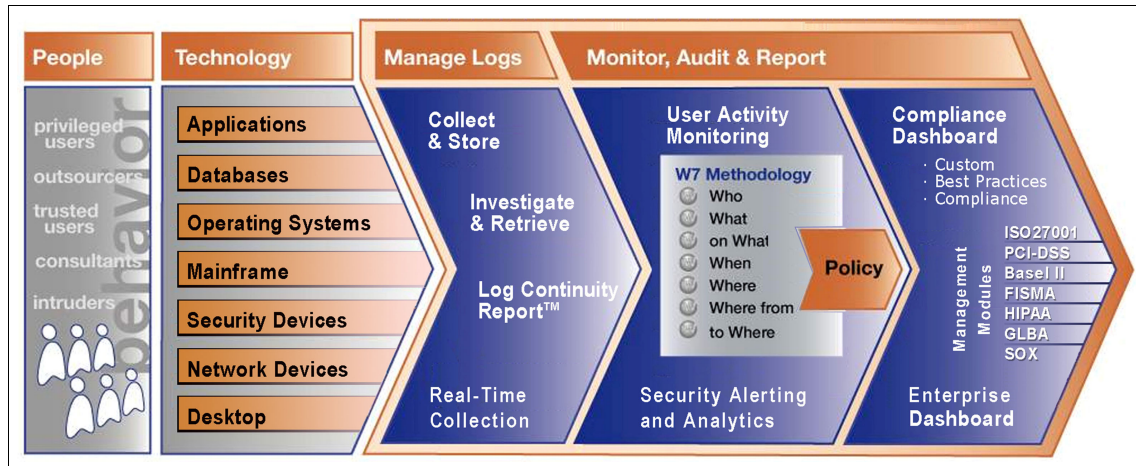


Figure 7-18 IBM Tivoli Security Information and Event Manager

The combination of the Tivoli Security Information and Event Manager and IBM Security SiteProtector products can provide organizations with an integrated SIEM solution combined with the event and policy management of their threat protection devices.

In addition, the Tivoli Security Information and Event Manager solution offers other capabilities that are often key decision factors for organizations that are looking for these types of technologies. Among these additional factors are:

- ▶ A rapid and scalable deployment and support process

The Tivoli Security Information and Event Manager solution is modular and flexible and includes the capability to start with a minimal implementation to address key aspects of SIEM requirements and then grow to encompass the complete SIEM requirements of an organization at a later time.

The Tivoli Security Information and Event Manager implementation can start to address the log management functionality first and address the other SIEM functionality over time, in phases. Sometimes, organizations that deploy SIEM solutions place an emphasis on their vendor providing an appliance type solution.

IBM provides the flexibility of a rapidly deployable software solution that can be installed on the most appropriate hardware configuration to meet the log data collection, archiving, and reporting requirements, and custom IT hardware requirements, such as brand, type, and configuration. The installation of a Tivoli Security Information and Event Manager solution can take less than an hour on the various hardware and operating system platforms that it supports; therefore, an organization can focus on the configuration activities that are more important during project implementation.

Typically, appliance-based solutions do not offer much flexibility, for example, for growth, you only have the option of buying more appliances. With the Tivoli Security Information and Event Manager solution, you have options to increase the amount of hardware or the capacity of the hardware that you have and to change the operating system platform to a more scalable approach.

- ▶ Reliable and secure log collection and archiving

Tivoli Security Information and Event Manager offers the possibility to begin an SIEM implementation with log data collection and archiving only. Tivoli Security Information and Event Manager can process up to 30,000 syslog messages or SNMP traps per second and archive them using a FIPS-certified communication protocol.

- ▶ Capability to collect and archive any type of log data

IT departments in many organizations must support business processes with customized software or with in-house developed software. Because these highly customized systems typically support essential business processes, their log data is subject to audit and analysis, to reduce the business risks or, at least, prove that these sensitive business processes are monitored continually. Custom built applications can be integrated with Tivoli Security Information and Event Manager to collect any type of log data and archive it using the FIPS certified communication layer.

- ▶ Integration with identity and access management solutions

Tivoli Security Information and Event Manager supports integration with a large amount of user directory types. The user and group/role/profile information that is maintained in these user directories can be applied to the security policy compliancy reports to show user behavior that does not follow the proper user's profile/role. These reports can be used for role life cycle

management and to automate remediation of user directory configuration errors.

► Built-in best practice reporting and analysis of log data

Deployment of an SIEM tool can sometimes be delayed because there is no clear idea what reports can help lower operational risks. Tivoli Security Information and Event Manager provides many report templates that can help monitor business processes following best practices recommendation for the IT audit field. These reports can help lower the operational risk that is related to the use of privileged user accounts. These reports are available when Tivoli Security Information and Event Manager is deployed as an SIEM solution. In the case where Tivoli Security Information and Event Manager is deployed as a Log Management only solution, similar reports can be generated to search the log data for suspicious events.

► Predefined audit and regulatory compliance reports cover the following standards and regulations:

- SOX
- FISMA
- HIPAA
- PCI DSS
- BASEL II
- GLBA
- ISO 17991
- ISO 27001
- COBIT
- NERC

The reports that are needed to support the regulatory compliancy process are, in some cases, predefined, and in other cases they rely on the implemented security controls in an organization. Tivoli Security Information and Event Manager provides built-in reports for regulations with specific requirements. For other, more flexible requirements, Tivoli Security Information and Event Manager uses the *ISO Code of Practice for Information Security Management* as the framework for reporting. When an organization decides they must comply with an IT-related regulation, they may not have defined the necessary security controls that are required for the monitoring and audit process yet. Tivoli Security Information and Event Manager can help the organization by providing a best practice starter set of reports. Gradually, the organization can create and tune its own set of regulatory compliance reports that it can use to support the compliancy claim.

- ▶ Capability to normalize any type of log data for audit and regulatory compliancy processing

Collecting and archiving log data from software that is developed in-house is possible with Tivoli Security Information and Event Manager. But, most of the time, collecting and archiving this log data is not good enough, because this log data can contain information that is crucial for effectively monitoring the sensitive business processes. Therefore, being able to compare this log data with business process rules such as IT security policies is essential.

Tivoli Security Information and Event Manager provides a normalization library that allows organizations to create scripts that can normalize any type of log data into a database model. This data can then be used for reports to facilitate a comparison of the log data against the IT security policy and business process rules. Tivoli Security Information and Event Manager has built-in normalization scripts for over 300 various types of log data that are generated by well-known operating systems, databases, applications, data management systems, and network devices.

- ▶ IBM mainframe integration

Among the well-known supported operating systems is IBM z/OS® for System z® servers. Tivoli Security Information and Event Manager covers a wide range of System Management Facility (SMF) types and subtypes, including SMF records that are generated by DB2 and CICS®, in addition to RACF® events. When a System z system is used in an IT environment, this system probably manages business critical data, and therefore, these systems are most likely audited and monitored.

- ▶ High performance syslog and SNMP collector

Many IT systems that support an overall IT infrastructure, such as routers and firewalls, historically do not have a sophisticated subsystem for auditing and monitoring. Operational logs that are generated by such systems are mostly stored on syslog-enabled host machines. The syslog protocol was originally designed for system maintenance messages and the protocol itself does not support a guaranteed message delivery mechanism. Environments where large amounts of syslog messages must be analyzed therefore require a syslog-enabled host that can process high volumes of syslog messages. Tivoli Security Information and Event Manager uses a high performance syslog collector that can process 30,000 events per second. This system is not restricted to syslog messages, but can also process SNMP traps at the same rate. The Tivoli Security Information and Event Manager syslog collector is software based and does not require dedicated hardware, which makes it scalable.

- ▶ Real time event correlation and incident response functionality
Tivoli Security Information and Event Manager can process the collected and archived log data in near real time and perform basic event correlation to identify a possible security breach. Security incidents can trigger Tivoli Security Information and Event Manager alerts using SMTP, SNMP, or executables to process the security incident's information.

More information about the IBM SIEM Solution can be found in *IT Security Compliance Management Design Guide with IBM Tivoli Security Information and Event Manager*, SG24-7530.

7.4.7 IBM Tivoli Netcool/OMNIBus

IBM Tivoli Netcool/OMNIBus is the IBM high-performance event Manager of Managers (MOM) for consolidating complex IT and network operation management tasks. Often, the security and network operations domains are implemented in separate service operation centers. When attacks occur in the network, information from the network domain and security domain can provide crucial information to both the network operators and the security engineers. Sharing this information between the domains, or providing a single *pane of glass* to display network and security events, is often critical to getting these threats under control.

IBM Tivoli Netcool/OMNIBus provides the following capabilities:

- ▶ Delivers a central point of real-time service management for business applications, network devices, Internet protocols, and security devices.
- ▶ Enables you to identify and resolve the most critical problems with automated event correlation, isolation, and resolution capabilities.
- ▶ Consolidates data in operational silos into real-time web dashboard views with customizable displays of events, service views, and operational indicators.
- ▶ Supports current and evolving standards and uses approved cryptographic providers to help ease security audits.
- ▶ Uses customizable lightweight agents to collect business and technology events from more than 1000 sources in real time.
- ▶ Has a clustering capability that allows the distribution of Tivoli Netcool/OMNIBus servers across an enterprise and the consolidation to a centralized server.

- ▶ Provides tight integration with many vendor monitoring solutions, such as IBM Tivoli Monitoring, HP Openview, CA Unicenter, and BMC Patrol through the use of probes designed specifically for these solutions. Over 200 generic and vendor specific probes have been developed.
- ▶ Provides generic probes for ODBC, SNMP, HTTP, Syslog, and other organization specific integration with many other platforms.
- ▶ More than 30 vendor alliances with major network and security vendors, including Cisco, Juniper, Brocade, and Checkpoint.
- ▶ Gateway integrations with IBM Tivoli Service Request Manager®, BMC Remedy, HP Service Desk, and others provide integration with existing trouble ticketing systems.

IBM Tivoli Netcool/OMNIbus is used by many of the major telecommunication companies for its performance and flexible integration. Service providers report to IBM that Tivoli Netcool/OMNIbus has been able to support rates of 1 billion raw events per week and event storms of 1 million events per minute over several hours, which is critical to maintaining control of the network when major outages or sustained security events occur.

The components of IBM Tivoli Netcool/OMNIbus are:

- ▶ Probes
- ▶ Gateways
- ▶ The Object Server
- ▶ Tivoli Integrated Portal for Event Visualization and configuration
- ▶ Desktop Tools
- ▶ Administration Tools

Probes are software tools that provide integration with other solutions and forward their events to an object server. They are visually represented in Figure 7-19. IBM offers a probe for IBM Security SiteProtector that reads the database tables in SiteProtector directly and then sends those events to the Tivoli Netcool/OMNIBus Object server. The Object server is a memory-resident database that can be configured in a high-availability failover configuration that is responsible for the correlation and deduplication of events as they arrive. Gateways are used to pass events to ticketing and other systems, or used to enrich events with additional data, such as location, point of contact, or service level agreement information. A browser or desktop interface can be used to interact with the system, display events, configure policies and actions, and manage the system.

A complete listing of the IBM Tivoli Netcool/OMNIBus probes can be found in the Omnibus Information Center at:

http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.nam.doc/welcome_ptsm.htm

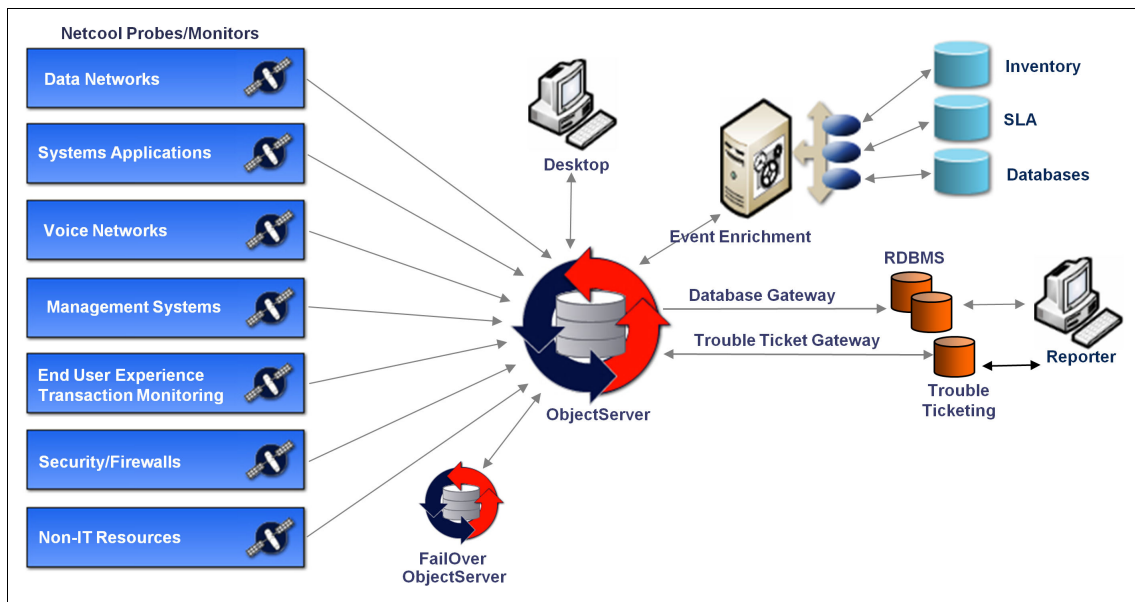


Figure 7-19 IBM Tivoli Netcool/OMNIBus architecture

Additional information about IBM Tivoli Netcool/OMNIBus and other complementary products that are integrated to create a robust systems management infrastructure can be found in *Integration Guide for IBM Tivoli Netcool/OMNIBus*, *IBM Tivoli Network Manager*, and *IBM Tivoli Netcool Configuration Manager*, SG24-7893.

7.5 Conclusion

We started this chapter by explaining how a centralized management platform offers a way to reduce costs and demonstrate compliance and business value. We mentioned how IBM Tivoli Application Dependency and Discovery Manager can be used to provide you with an accurate asset overview of your organization and how the integration with IBM Rational AppScan can provide a more detailed view of your security status.

We explained how SiteProtector controls the various agents discussed in Chapter 8, “Network security solutions” on page 243 and Chapter 9, “Host security solutions” on page 299. We illustrated how SiteProtector manages the various types of security policies and how it can be used to monitor threats. We provided an overview of the different components that make up a SiteProtector deployment and elaborated on several architectural elements, such as the communication channels, and ways to provide redundancy and enforce a strict separation of duties. We concluded this chapter by referring to the fact that SiteProtector integrates with several other tools, such as IBM Tivoli Security Information and Event Manager and how it, as such, can fit in a larger, more comprehensive management strategy.



Network security solutions

In this chapter's context of network security solutions, we focus on the IBM Security Network Intrusion Prevention System product suite. We also provide a solution overview of the IBM WebSphere DataPower® SOA Appliances, the IBM Tivoli Netcool Configuration Manager, and the IBM Lotus Protector product.

In addition, we revisit the definition of an intrusion and how to prevent it, explain the concepts of the IBM Security protection engine, review the different security policy options, and describe the available deployment options. This content is mapped out in the following sections:

- ▶ “IBM Security Network IPS” on page 244
- ▶ “Intrusion and intrusion prevention definitions” on page 262
- ▶ “Intrusion prevention policies” on page 263
- ▶ “Intrusion prevention enforcement” on page 266
- ▶ “Physical deployment model” on page 272
- ▶ “IBM Tivoli Netcool Configuration Manager” on page 284
- ▶ “IBM WebSphere DataPower” on page 287
- ▶ “IBM Lotus Protector for Mail Security” on page 295

8.1 IBM Security Network IPS

The IBM Security Network IPS delivers preemptive network protection through its combination of line-speed performance, security intelligence, and a modular protection engine that delivers security convergence.

In this section, we discuss several security capabilities that the IBM Security Network IPS product offers.

To understand how the security capabilities of the IBM Security Network IPS can be mapped to the IBM Security Blueprint¹, see Figure 8-1. This diagram shows the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate the functional components that can be fulfilled, or implemented, using IBM Security Network IPS. This functional highlighting is applicable for the infrastructure service components as well.

Besides the fully highlighted elements, Figure 8-1 also shows some medium highlighted elements. Although the IBM Security Network IPS can be used to address such components to some degree, the respective area of coverage is not considered a core function of the product and thus is considered to be limited.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 8-1 can be used as a quick reference of the functional security management aspects of the IBM Security Network IPS. This reference can help us determine which functions of a solution can be covered by selecting this product.

¹ For a detailed discussion about the elements, see Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

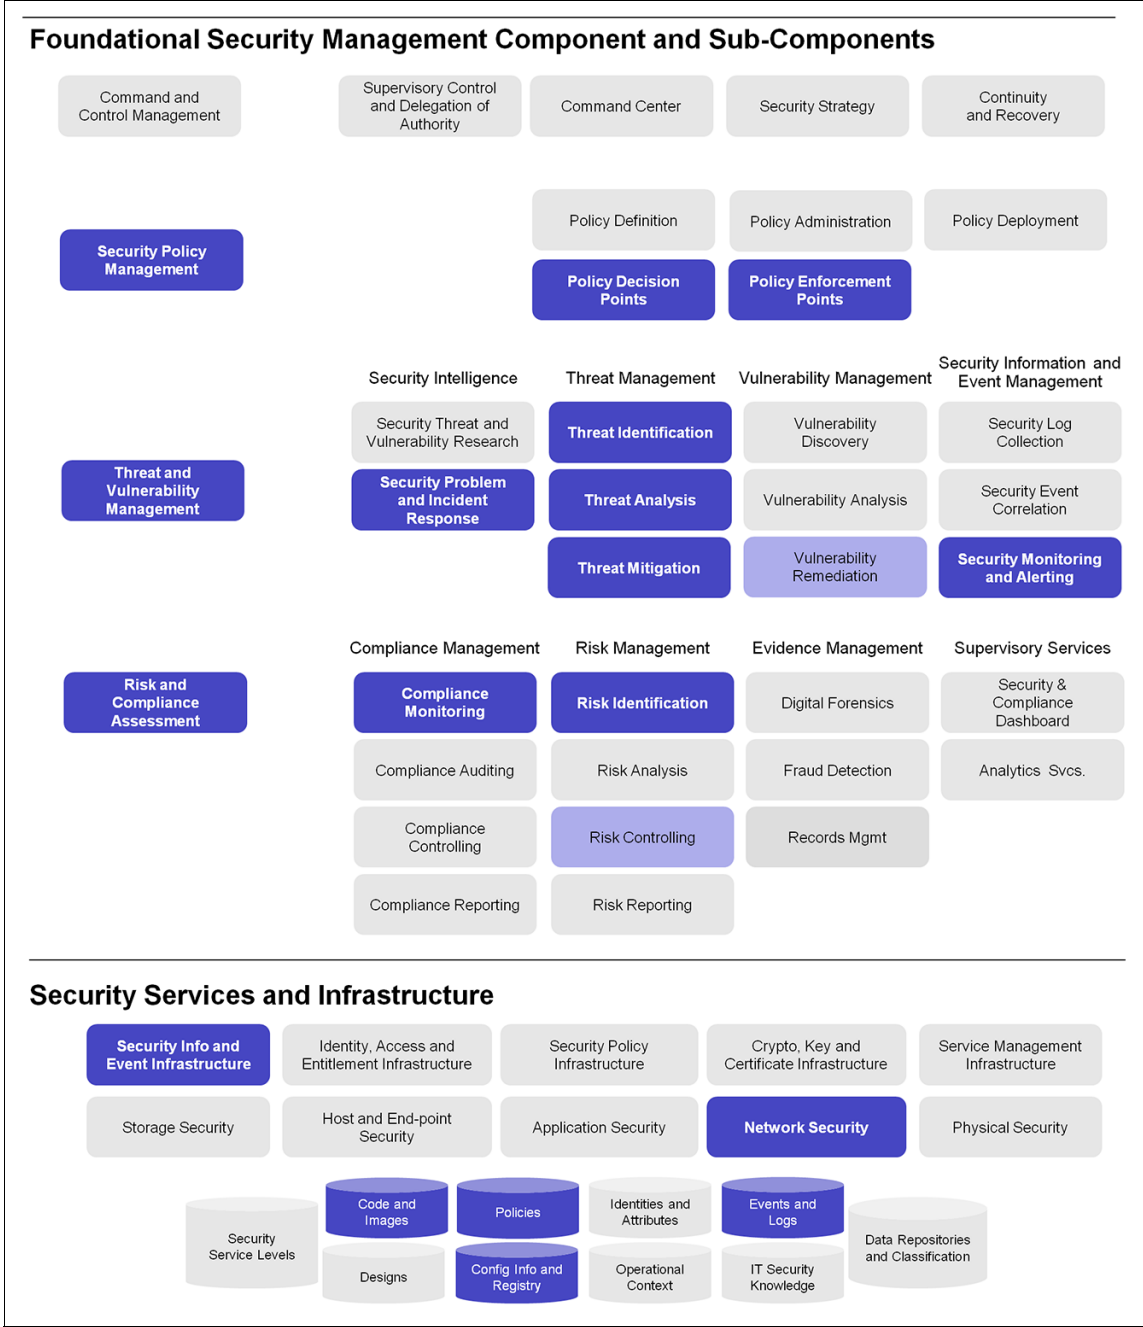


Figure 8-1 Mapping of the IBM Security Network IPS to the IBM Security Blueprint

The IBM Security Network IPS delivers network protection that is designed to:

- ▶ Stop threats before they impact network assets without sacrificing high-speed network performance.
- ▶ Provide a platform for security convergence that eliminates the costs of deploying and managing point solutions for web application and data security.
- ▶ Protect networks, servers, desktops and revenue-generating applications from malicious threats.
- ▶ Conserve network bandwidth and prevent network misuse/abuse from instant messaging and peer-to-peer file sharing.
- ▶ Prevent data loss and aids compliance efforts.

The IBM Security Network IPS can stop Internet threats before they impact your organization. It delivers protection to all three layers of the network: core, perimeter, and remote segments.

The IBM Security X-Force Research and Development Organization (X-Force) enables *ahead of the threat* protection for an IT infrastructure before vulnerabilities are made public and before exploits against those vulnerabilities become available.

By consolidating network security demands for data loss prevention and protection for web applications, IBM Security Network IPS serves as the security platform that helps reduce the costs of deploying and managing point solutions.

When evaluating intrusion prevention technology, organizations often struggle to balance and optimize the following six areas:

- ▶ Performance
- ▶ Security
- ▶ Reliability
- ▶ Deployment
- ▶ Management
- ▶ Confidence

The IBM Security Network IPS delivers on all six areas, with performance, preemptive protection, high availability, simple deployment and management, and excellent customer support. Organizations can manage the IBM Security products themselves or can decide that they want to transfer the risk of protecting their network to a trusted security partner. The IBM Security Services division can manage the IBM Security Network IPS product family of network, server, and endpoint protection. Working with IBM organizations provides benefits from a range of complementary consulting services for assessment, design, deployment, management, and education.

8.1.1 Zero-day protection

IBM Security Network IPS delivers zero-day protection through the X-Force Virtual Patch technology.

By providing the Virtual Patch, the IBM Security Network IPS allows organizations to avoid emergency patching because Virtual Patch blocks attacks against vulnerabilities at the network level before they can reach their targeted system, application, or network resource.

X-Force research and the IPS engine identify and block attacks based on the vulnerability, so IBM Security Network IPS does *not* require a new signature for every exploit. The X-Force-powered protection engine employs multiple intrusion prevention technologies in tandem to monitor, detect or block these classes of network threats, such as cross-site scripting, drive-by downloads, and web browser attacks.

8.1.2 Next generation product enhancements

In 2010, IBM Security Solutions released their next generation versions of the IBM Security Network IPS products:

- ▶ Version 2.0 hardware
- ▶ Version 4.1 firmware
- ▶ New Virtual IPS platform

In the following sections, we explain the features and benefits of these new product versions.

8.1.3 Next generation hardware

Version 2.0 of the IBM Security Network IPS hardware delivers a doubling in performance compared to Version 1.0 hardware.

Performance was optimized by deploying the following items:

- ▶ Multi-core 64-bit CPU
- ▶ Increased memory
- ▶ Improved motherboard for faster bus speeds

Figure 8-2 shows the 2010 product range and the associated performance figures.

IBM Security Network IPS Throughput Metrics						
	Remote Segments	Perimeter			Core	
Model	GX4004-V2-200	GX4004-V2	GX5008-V2	GX5108-V2	GX5208-V2	GX6116
Inspected Throughput	200 Mbps	800 Mbps	1.5 Gbps	2.5 Gbps	4 Gbps	8 Gbps
Protected Segments	2	2	4	4	4	8

Figure 8-2 Next generation IBM Security Network IPS: Models and performance figures

10 Gbps (10G) support is described fully in 8.5.6, “10 Gbps environments” on page 282.

Each appliance provides two management interfaces, a *TCP_Reset* interface and either four, eight, or 16 monitoring interfaces. The number of monitoring interfaces corresponds to the last two digits of the model number. For example, a GX4004 has four monitoring interfaces and a GX6116 has 16 monitoring interfaces.

In Figure 8-3, you can see the front of a GX5008 appliance. You can clearly see the two management interfaces, eight monitoring interfaces, two USB interfaces, and the TCP_Reset interface. The monitoring interfaces are labeled as follows:

- ▶ 1A and 1B
- ▶ 2C and 2D
- ▶ 3E and 3F
- ▶ 4G and 4H




Figure 8-3 IBM Security Network IPS GX5008 appliance

The number corresponds to the *network segment* being protected (1, 2, 3, or 4) and the letter corresponds to the physical interface (A - H).

Version 2.0 Network IPS appliances can be upgraded using a USB memory stick. The procedure to do this is detailed in the *IBM Security Network Intrusion Prevention System Installation Guide, Firmware Version 4.1*².

Remote segment NIPS

The GX4004C-V2-200 allows IBM to serve the low-end IPS market at a competitive price. This GX4004C-V2-200 is a 4-port appliance capable of protecting two network segments. It is licensed for up to 200 Mbps.

The GX4004C-V2-200 uses the same hardware as the GX4004C-V2, but the client is limited by the license agreement to only deploy this product on networks with a maximum of 200 Mbps of traffic.

8.1.4 Next generation firmware

In firmware release 4.1, a major redesign of the local management interface was made. This new release provides a significantly improved user experience when managing individual Network IPS appliances from a secure web browser session and provides the following main benefits:

- ▶ Improved, easier to use navigation
- ▶ Simpler transfer of control of the appliance between IBM Security SiteProtector and Proventia Manager

The menu items shown in Figure 8-4 make it easy to navigate through the five most common user tasks:

- ▶ Home - Appliance Dashboard
- ▶ Monitor - Health and Statistics
- ▶ Secure - Protection Settings
- ▶ Manage - System Settings
- ▶ Review - Analysis and Diagnostics




Figure 8-4 Firmware 4.1: Easy to use navigation

² The IBM Security Network IPS product documentation can be found at http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp?topic=/com.ibm.ipsec/IBMSecNetIPS_landing_page.html.

Figure 8-5 shows the Firmware 4.1 Appliance Dashboard, which is the first window that opens after connecting to the appliance.

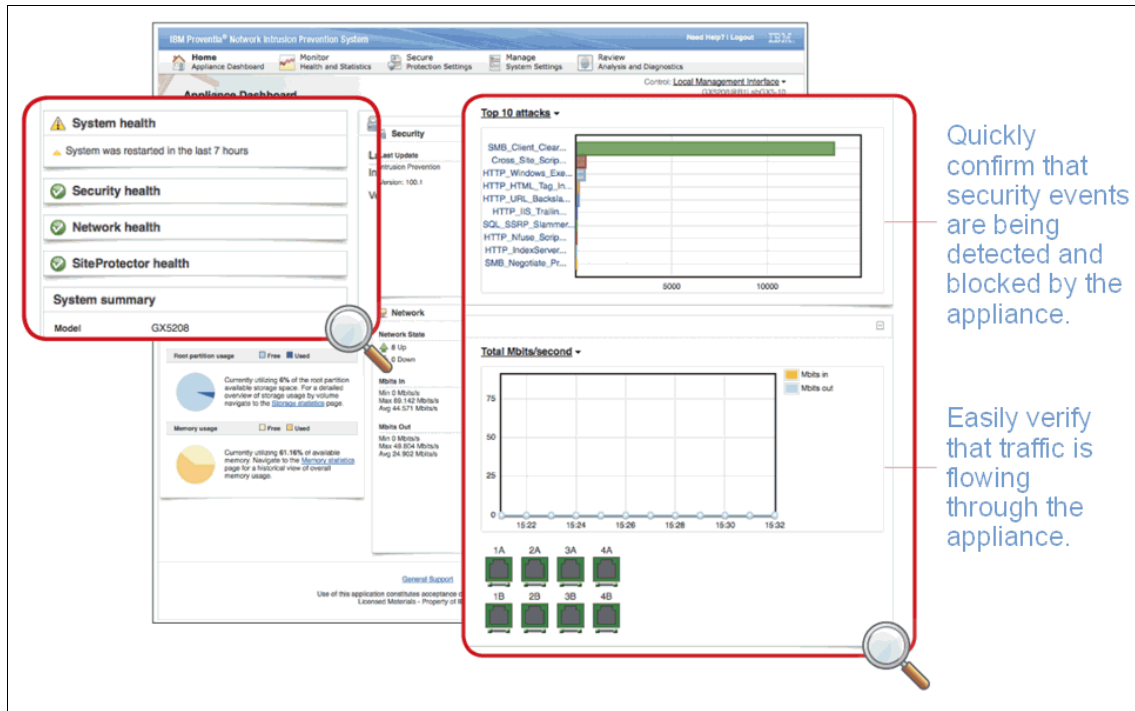


Figure 8-5 Firmware 4.1: Appliance Dashboard

The dashboard provides an at-a-glance view of the health of the key components of the solution:

- ▶ Network health
- ▶ Security health
- ▶ System health
- ▶ SiteProtector health

In addition, the top 10 attacks, memory and disk usage, and network interface utilization are also shown.

Users can drill down to more detailed data by clicking through the dashboard summary information.

Security Modules menu

To simplify the configuration of the most important protection aspects of the IBM Security Network IPS product suite, a separate menu option was introduced in firmware 4.1, that is, Security Modules, as shown in Figure 8-6.

This menu option simplifies the configuration of the following items:

- ▶ Data Loss Prevention
- ▶ Web Application Protection
- ▶ X-Force Virtual Patch

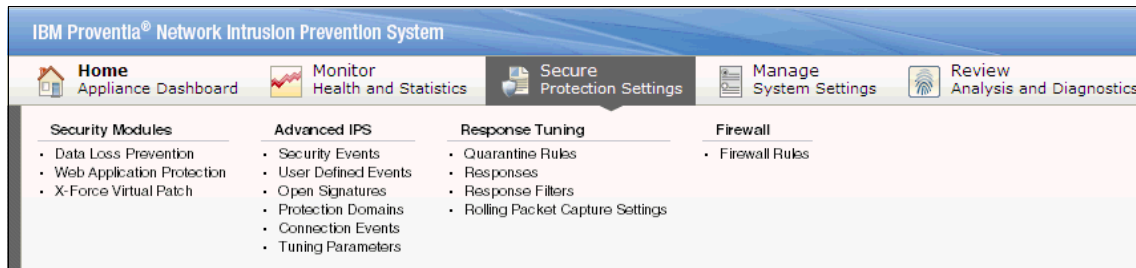


Figure 8-6 Firmware 4.1: Protection Settings menu item showing the Security Modules

In the following sections, we explain these new security modules.

Data Loss Prevention (DLP)

Data loss can be costly for any organization, and it can occur in many different ways. Data can leak purposely from an insider who intentionally takes information out of the network, or it can happen accidentally. For the scope of this book, we address how threat mitigation techniques can complement data loss prevention techniques.

There are many ways that information can be digitally stored. It can be structured, unstructured, images, video, voice, or many other types of formats. To make matters even more complex, data can be stored on devices of many types, such as cellphones, laptops, USB drives, iPods, PDAs, or even in briefcases.

Of course, it is the access to, and changes to the data, that introduces the risk. If someone accesses the data, that person has many opportunities to move the data across the digital network, over either voice, audio channels (or both), cut and paste the data into an email, print or fax the data, and so on.

The process of outlining the risk of data leakage is further discussed in the following list:

► Assess

With regards to the access of data at rest, you might want to ask, “Do I have intellectual property, confidential records, or personally-identifiable information that potentially violates policy, government regulations, or is on the verge of being compromised (or both)?” You can only adequately protect your data when you know how it is classified.

► Protect

When you identify the need to protect the data usage at the endpoint, you might want to think about methods to categorize the data, standardize policies, manage data protection issues at the point of use, and keep those policies manageable.

► Defend

With regards to data in transit, guarding the data while it is at rest might not be an issue. However, we need to manage policies that can help protect the data while sharing it, and take note of the issue surrounding the internal threats to the organization.

► Monitor

The number of integrated solutions keeps growing, and the ability to report and track information about multiple consoles can be a daunting task. It is as though we have so much control, yet we are out of control.

► Control / Respond

The amount of data we have to respond to can create additional problems.

With false positives, and the amount of data to sort through, the real violations can be lost. Data leakage strategies typically do not begin with protection technologies.

However, threat mitigation technologies can assist with identifying data leakage problems in the environment, and can also provide some insight as to where the data leakage problem might exist, and therefore provide a nice complement to address the problems listed above.

Although data leakage is typically addressed at the point where the data resides, there are two key technologies that help complement a data leakage network architecture: network intrusion prevention systems, and mail security systems.

By understanding the communication between hosts, we can get a visual representation of the data movement in our network. Many times, this representation can lead us to detect areas of concern that previously went unnoticed.

The IBM Security Network IPS can also help us gain some insight into improving the protection of data movement. In this case, the data must be unencrypted and available for the IPS to see, which allows the IPS to locate misuse and abuse from an insider.

Modern IPS systems have the ability to examine files and content within the data stream that passes through the IPS. By looking for keywords and information associated with important business aspects, you can use an existing IPS infrastructure to watch for data leakage by looking for data signatures, inspection of the use of common and uncommon protocols, and the content that can be carried within the protocol.

- ▶ Content types

Typically, the data content that is being searched for includes Personally Identifiable Information (PII), where someone is trying to gain access to steal another person's identity. Examples include credit card numbers, Social Security numbers, and email or postal addresses.

- ▶ Protocols

Protocols that are inspected include the following: HTTP, FTP, SMB, IMAP, POP3, SMTP, IRC, and peer-to-peer protocols, such as Yahoo! Messenger, Microsoft Messenger, and AOL IM.

- ▶ File types

The types of files being inspected include Microsoft Office documents, PDF files, HTML files, and compressed files.

Mail security gateways can also complement a data leakage protection strategy. Most modern mail security gateways can filter the content that is not only sent into the organization, but also filter the content that leaves the organization. When looking for information that flows outside, a mail security gateway can ensure that policies within the corporation are not violated, as files and message content can be inspected for appropriate use. See 8.8, "IBM Lotus Protector for Mail Security" on page 295 for more information.

The DLP policy editor shown in Figure 8-7 shows that a DLP policy can be made up of a combination of predefined events and user-defined events.

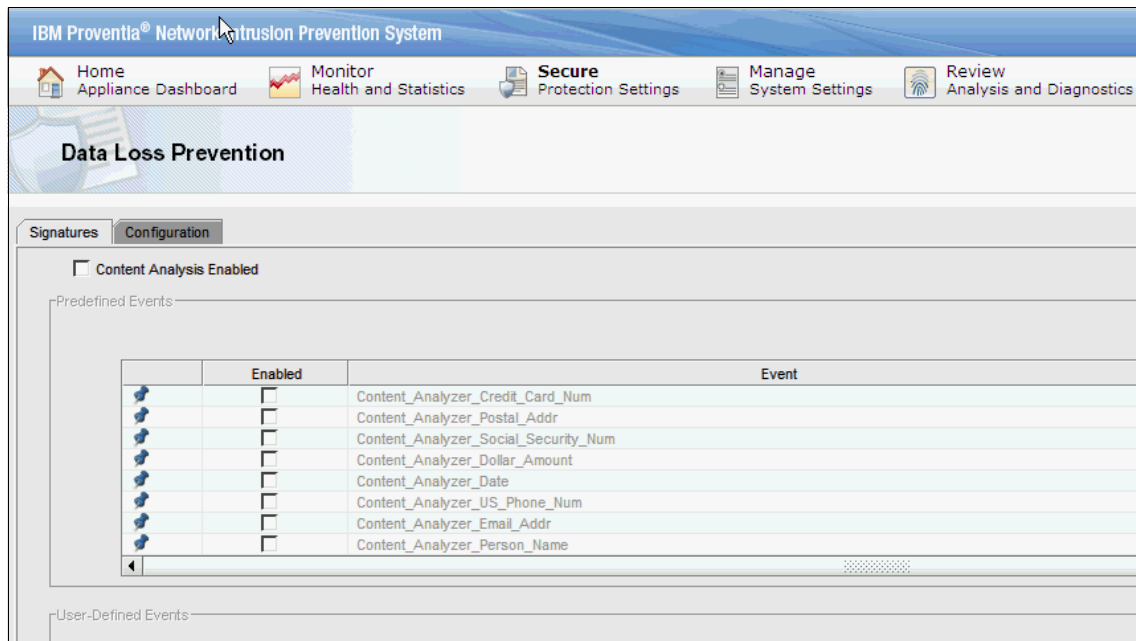


Figure 8-7 Data Loss Prevention policy editor

Web Application Protection

The security communities behind the Web Application Security Consortium (WASC)³ and the Open Web Application Security Project (OWASP)⁴ continue to develop and refine a common testing and evaluation standard for web applications. In line with this refined testing methodology, IBM introduced *Web Application Protection* as a new threat category.

The goals behind this introduction are:

- ▶ To group together the web application attack signatures into a single category.
- ▶ To simplify the policy management of these signatures.

³ For more information about WASC, go to <http://www.webappsec.org/>.

⁴ For more information about OWASP, go to <http://www.owasp.org/>.

This new threat category includes the following attack types:

- ▶ Client-side attacks
- ▶ Injection attacks
- ▶ Malicious file execution
- ▶ Cross-site request forgery (CSRF)
- ▶ Path traversal
- ▶ Buffer overflow
- ▶ Directory indexing

The new policy editor that was introduced in firmware 4.1, as shown in Figure 8-8.

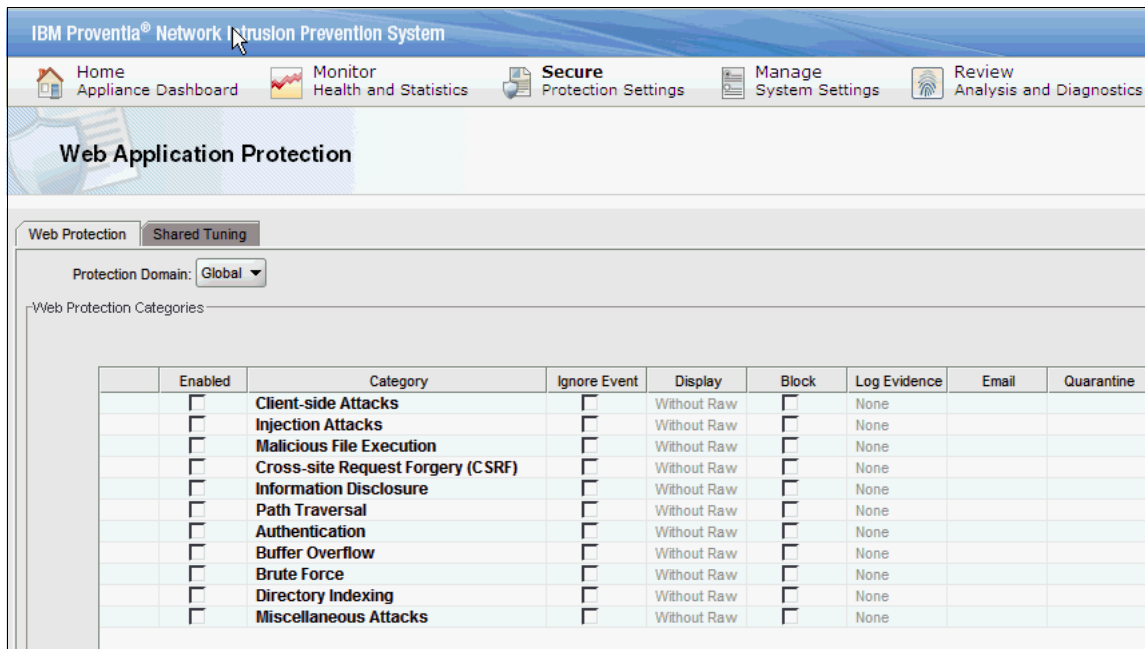


Figure 8-8 Web Application Protection policy editor

When you configure the IBM Security Network IPS and want to check the individual signatures associated with each attack category, you need to double-click the attack category and then click the **Show Security Events** button.

Figure 8-9 shows the signatures associated with client-side attacks.

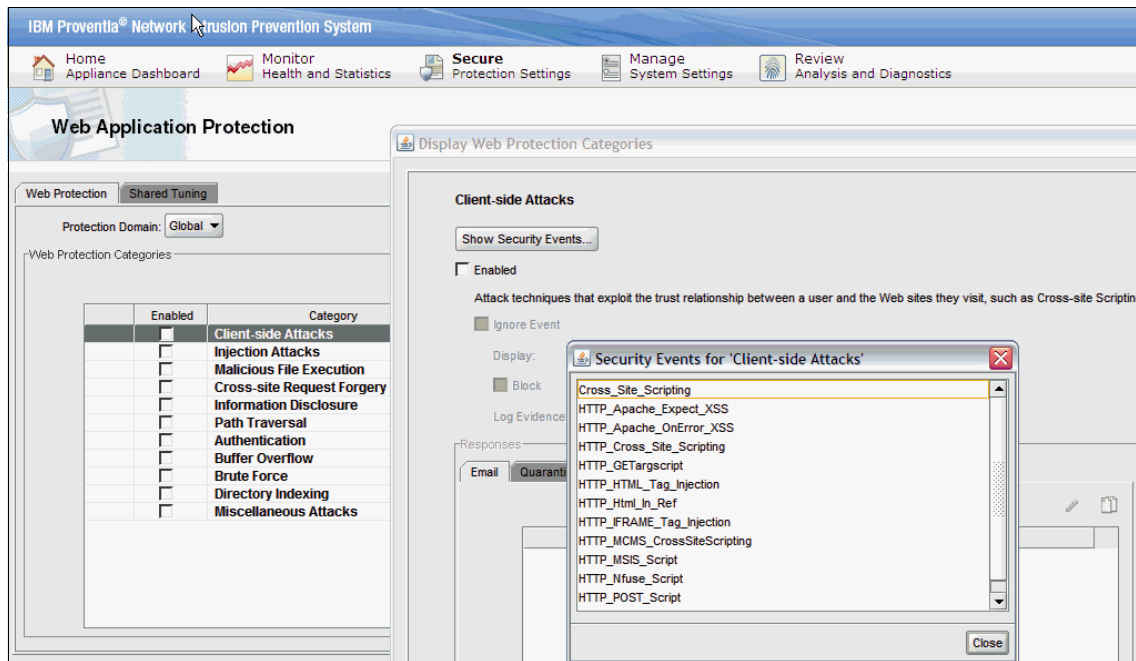


Figure 8-9 Client-side attack signatures

To see how closely the Web Application Protection categories match OWASP's Top Ten Web Application Security Risks or to obtain more information about this component of Web Application Protection, you may want to refer to the OWASP Top Ten Project⁵ or their testing guide⁶.

Where to look: The attack signatures contained in the Web Application Protection policy editor are not present in the general *Security Events* policy editor.

X-Force Virtual Patch

The third security module is the X-Force Virtual Patch module. Within this menu option, the user has the ability to define whether the default *block responses* defined by X-Force are turned on or off.

⁵ For more information about the OWASP Top Ten Project, go to http://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project.

⁶ For more information about the OWASP Top Ten Project testing guide, go to http://www.owasp.org/images/5/56/OWASP_Testing_Guide_v3.pdf.

In order for a specific signature to not only detect an attack, but also to block it, it must have the associated *block response* enabled. To simplify the deployment of the products in real-life environments, the X-Force team always specifies the default block responses in each X-Press Update (XPU). Administrators can then decide if they want to trust the X-Force default block responses or not.

The X-Force Virtual Patch menu option allows an administrator to change this setting, as shown in Figure 8-10.

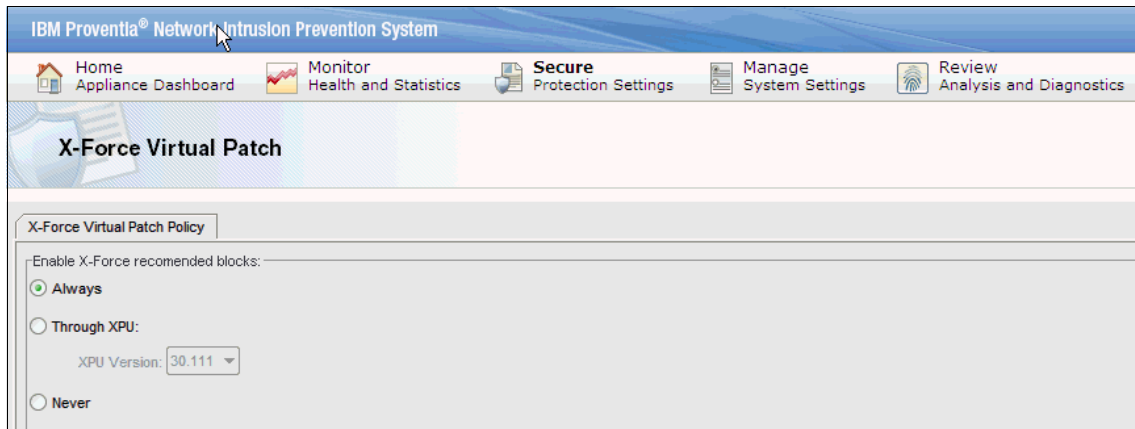


Figure 8-10 X-Force Virtual Patch policy editor

Fine grained control: Individual block responses for specific signatures can still be modified using the policy editor in the *Security Events* menu option.

Open Signatures

In addition to the protection signatures provided by X-Force in the monthly X-Press Updates (XPUs), customers also have the ability to define their own customized signatures, known as Open Signatures⁷ or Open Signature rules.

Let us take a look at an example of an Open Signature rule. The general syntax options for Open Signature rules are as follows:

```
<action>: alert
<protocol>: tcp, udp, icmp, ip
<IP and netmask>: single IP address (a.b.c.d), range of IP addresses
(a.b.c.d - w.x.y.z), network address using CIDR notation (a.b.c.0/24)
```

⁷ For more information, refer to the Open Signature User Guideline document at <http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/topic/com.ibm.opnsg.doc/pdfs/0penSignatureUserGuide.pdf>.

An example of a rule that detects the word “yahoo” in the data being inspected by the IBM Security Network IPS would be:

```
alert tcp any any -> any any
(msg:"Yahoo accessed";content:"yahoo";nocase; sid:5000;)
```

The `sid:5000` parameter above indicates the unique *rule ID*, which is always a unique 1 to 4 digit identifier.

The Open Signature rules engine adds 6,000,000 to this identifier to create a *custom issue ID*.

Thus, the custom issue ID that appears in the Alerts overview for the rule ID of *5000* is *6005000*.

Performance considerations: If good quality Open Signatures are written, the performance impact on the IBM Security Network IPS is negligible. The important detail to remember is not to rely too heavily on “wild cards” in each signature, as this will have a detrimental effect on the performance of the IBM Security Network IPS.

8.1.5 Next generation virtual appliances

As part of the overall IBM Security Network IPS portfolio, there are two virtual appliances available, that is, the GV200 and GV1000. They are provided as preconfigured Virtual Machine (VM) packages.

Deploying an IBM Security Network IPS Virtual Appliance provides the following benefits:

- ▶ X-Force powered protection in a virtual environment
- ▶ Lowered complexity with centralized operations
- ▶ An upgrade path from previous IBM Internet Security Systems (IBM ISS) software-based IDS solutions

The following virtual interfaces are automatically created during the installation process:

- ▶ TCP Reset port: Used for resetting TCP connections.
- ▶ Management port: Used for connection to SiteProtector and the Local Management Interface (LMI).
- ▶ Two sensor ports.

To understand the system requirements of the server onto which you are deploying these products, refer to the online *System Requirements for Virtual IPS Appliances* guide⁸.

Figure 8-11 shows a deployment in which both sensor ports are connected to physical NICs on the server, resulting in a two port *physical* appliance.

Monitoring limits: Unlike the physical appliances, the virtual appliances only have *two* monitoring (sensor) ports.



Figure 8-11 Virtual appliance on a physical server protecting a remote server farm

⁸ http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/topic/com.ibm.ips.doc/ProventiaIPS_Virtual_Security_Platform_Sys_Req.pdf

Figure 8-12 shows a deployment in which the virtual appliance provides protection to the cluster of virtual machines.




Figure 8-12 Single virtual appliance on a physical server protecting a virtual server farm

Figure 8-13 shows a deployment in which two virtual sensors are deployed on one physical server, protecting the virtual servers of company “A” and “B”.




Figure 8-13 Multiple virtual appliances on a physical server

8.2 Intrusion and intrusion prevention definitions

An intrusion is defined by any or all of the following definitions:

- ▶ An unauthorized act of bypassing the security mechanisms of a computer system to gain access to it, or cause a denial of service condition.
- ▶ Attacks that are attempted from outside the network security perimeter in attempts to access a secured computer system.
- ▶ An uninvited and unwelcome entry into a computer system by an unauthorized source.
- ▶ An entrance by force or without permission.
- ▶ An attempt to compromise the integrity, confidentiality, or availability of a system.

8.2.1 Intrusion prevention

Intrusion prevention implies the ability to prevent or deny an attempt to access an unauthorized portion of data, computer system, or network service. The goal is to prevent the alleged intrusion or at least report on it.

Another way to think about intrusion prevention is that it is a preemptive approach to network security, used to identify potential threats, and respond to them swiftly.

However, because an exploit might be carried out quickly after the attacker gains the knowledge or ability to bypass traditional security precautions, intrusion prevention systems also have the ability to take immediate action, based on a set of rules established by the administrator of the IPS. For example, an IPS might drop a packet that it determines to be malicious, and block all further traffic from that IP address or port. Legitimate traffic is forwarded to the recipient with no apparent disruption or delay of service.

An effective intrusion prevention system must be able to perform more complex monitoring and analysis, such as tracking and identifying protocols based on content and behavior (instead of port numbers), and watching and responding to traffic patterns and individual packets.

8.3 Intrusion prevention policies

The need to enforce a policy with the goal of preventing intrusions is readily understood and accepted.

A granular intrusion prevention policy must be defined, enforced, audited, revised, and re-enforced. Due to the impact of managing a granular policy, many system owners opt to outsource these types of burdens to a trusted security vendor. However, you should still take responsibility for evaluating and updating the policy at specific intervals.

Organizations expect their security technologies to protect them against each and every new threat with the same efficiency and level of performance as the day the solution is first purchased and installed.

New threats are not the only unauthorized activity that protection technology must combat. Older threats still plague the Internet, and must continually be prevented, whether they are old login bypass vulnerabilities, web browser exploits, worm infestations, or new vectors for previously misclassified vulnerabilities. Though many of these threats are in the eradicated phase of the vulnerability life cycle for many years, they continue to find avenues of success if the diligence to continue to prevent them is not present. Therefore, older threats must continue to be monitored and prevented.

To provide organizations with granular policy management capabilities, IBM Security Solutions use the concept of *Protection Domains*. By default, each IBM Security Network IPS has its own *Global Protection Domain*. Organizations can specify their own specific Protection Domains, for example, for specific mail servers, and apply a specific subset of signatures to that Protection Domain. Protection Domains can be defined according to the following parameters (including combinations thereof):

- ▶ Network IPS physical interface (port)
- ▶ VLAN tag/range
- ▶ IP address/range

After one or more Protection Domains have been configured, you can then allocate specific signatures to each one. To simplify this process, the policy editor allows signatures to be grouped, for example, by XPU version number or by protocol (Figure 8-14).

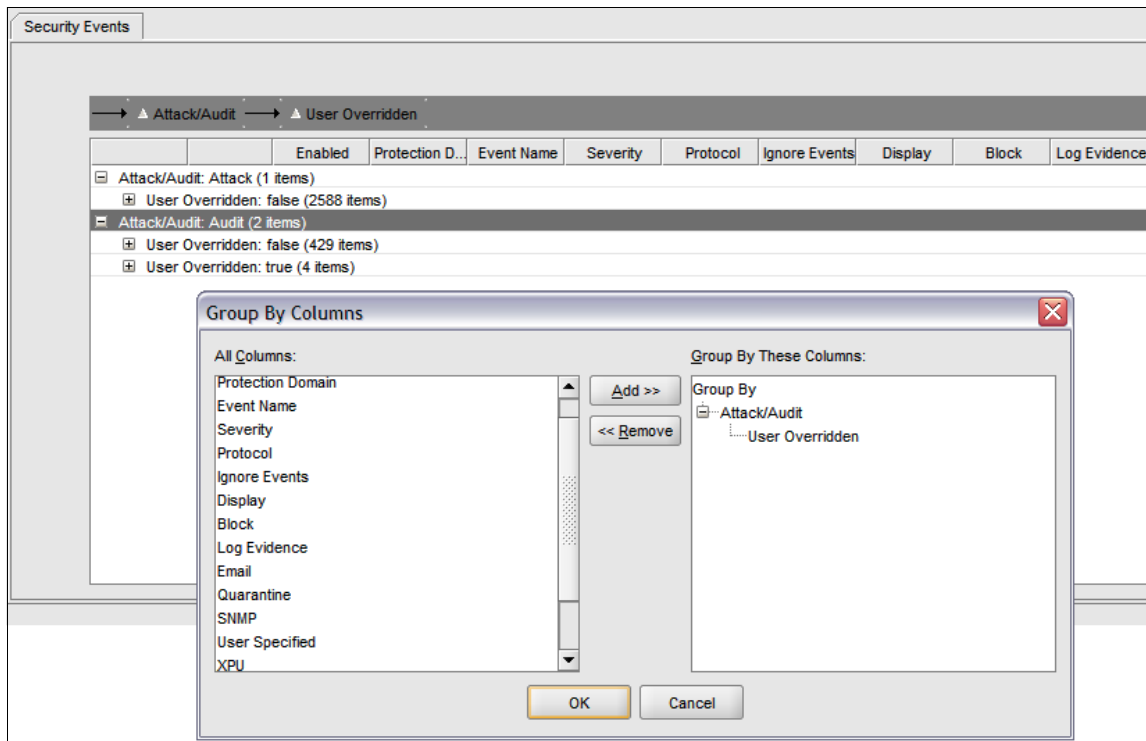


Figure 8-14 Security Events: Signature grouping

In addition, the following intrusion responses (and combinations thereof) are configurable within the policy editor:

- ▶ Block
- ▶ Ignore
- ▶ Log and log evidence
- ▶ Email
- ▶ Quarantine
- ▶ SNMP
- ▶ User-defined

When it comes to logging data that may be necessary for auditing purposes, for example, the following log options are available:

- ▶ attack packet logging
- ▶ pcap file

Log files can be accessed from the Local Management Interface via the Logs and Packet Capture Menu item (Figure 8-15).




Figure 8-15 Logs and Packet Captures Menu Item

In addition, in firmware 4.1 you have the option for the administrator to interact with the system, firewall, and security event logs directly from the local management interface.

Figure 8-16 shows an example where the administrator has the option to directly block an intruder by *right-clicking* the Source IP address field.

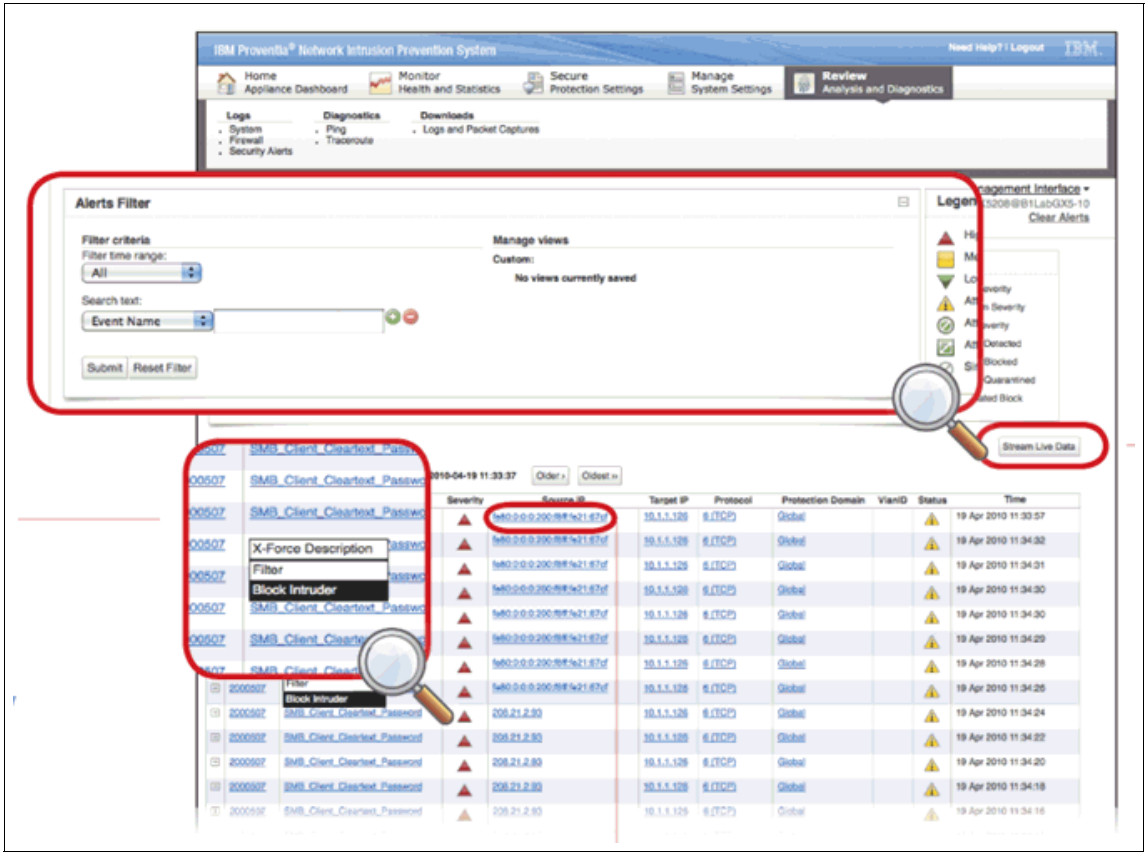


Figure 8-16 Direct administrator interaction with Security Event logs

8.4 Intrusion prevention enforcement

In many cases, the security policy is viewed as a burden in the network and the infrastructure, when in fact it can increase throughput of the networks and provide the means to discover misuse and abuse within the infrastructure, which is actually slowing transactions down.

8.4.1 General network requirements

Let us start outlining our security requirements and applying them to our network infrastructure. The network infrastructure is one of the access methods an attacker can use to carry out his threat against a vulnerable or accessible server, or even an unsuspecting user. Our conceptual architecture is shown in Figure 8-17.

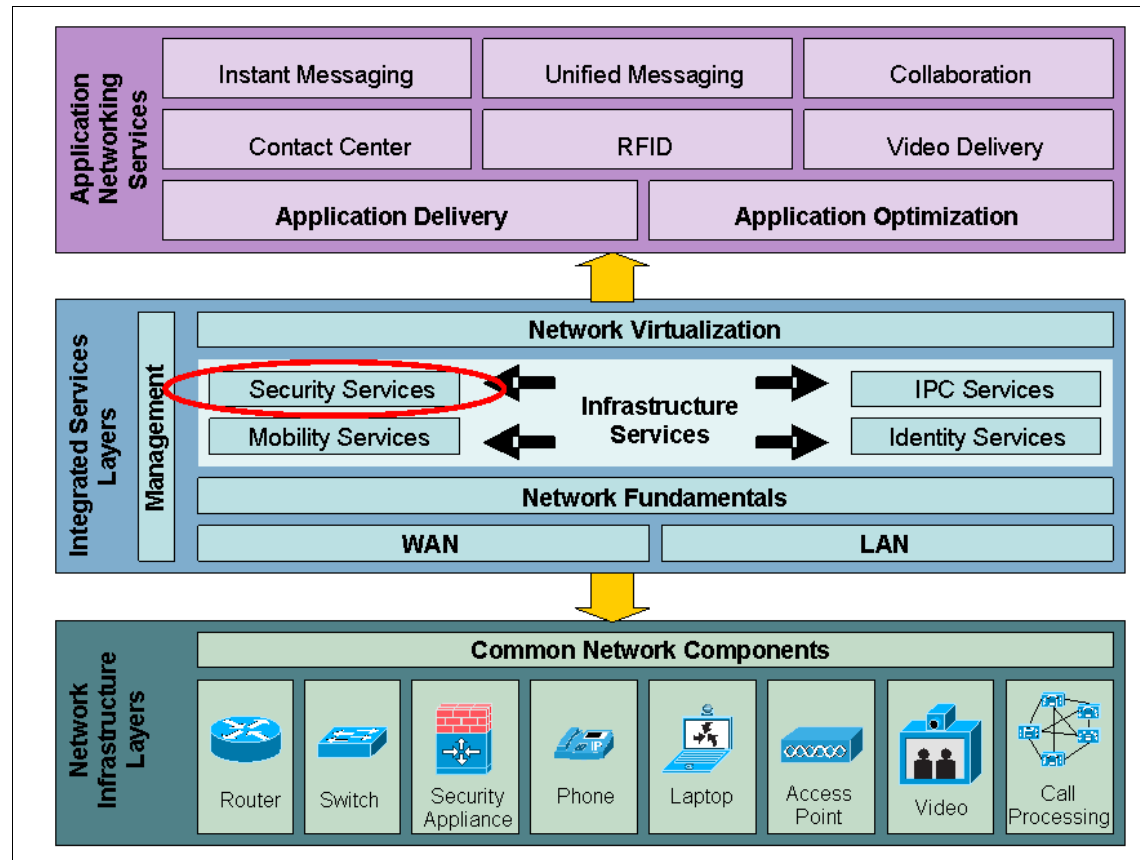


Figure 8-17 Conceptual architecture

8.4.2 Network IPS requirements

Network IPS products must not only provide the highest level of protection to block attacks, but must offer optimum performance and rely on a solid foundation of research to understand and block known and unknown vulnerabilities and exploits.

Figure 8-18 illustrates the six key attributes provided by the IBM Security Network IPS solutions that help achieve superior performance, along with the ability to prevent known and unknown threats.

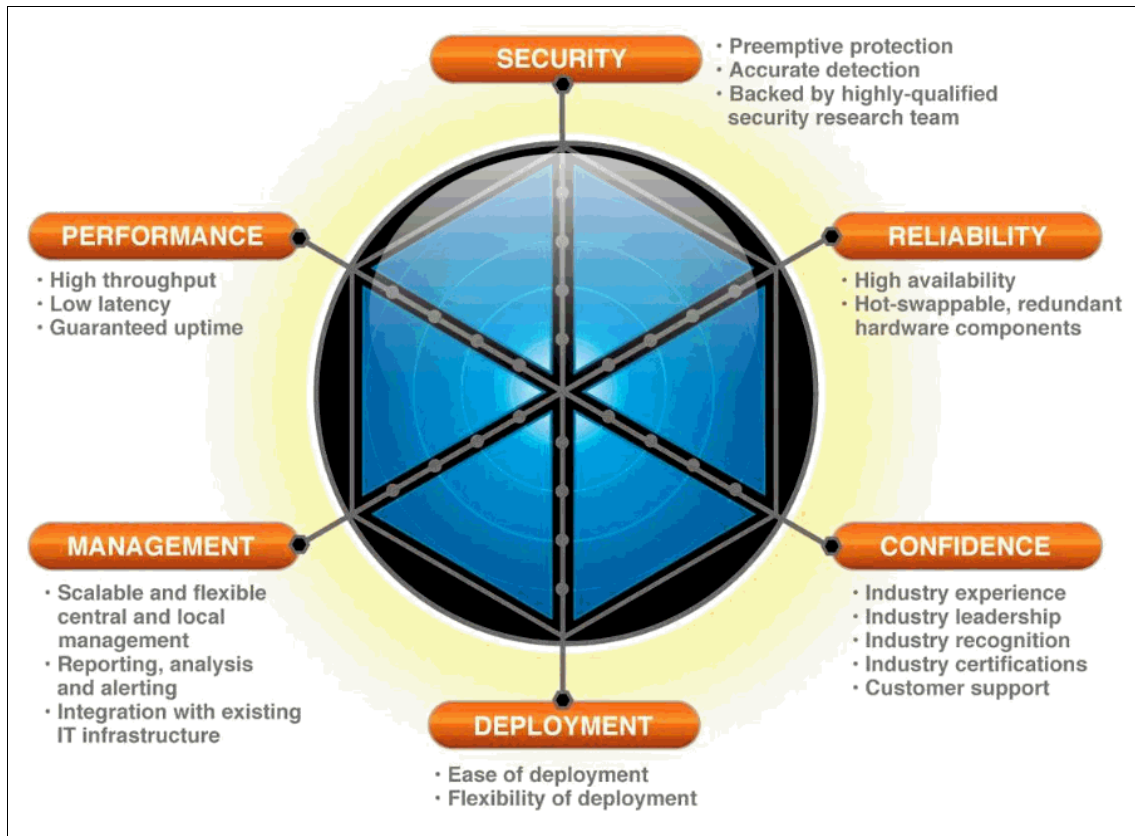


Figure 8-18 IBM Security NIPS: Six key attributes

In the following sections, we explain how the IBM Security Network IPS delivers on these six key areas:

- ▶ Performance
- ▶ Security
- ▶ Reliability
- ▶ Deployment
- ▶ Management
- ▶ Confidence

8.4.3 Performance

The first rule of preemptive protection of an intrusion prevention system is performance. IPS performance is ideally matched to the environment being protected. Several subcategories contribute to the overall performance of an IPS.

Inline operation

An effective IPS must operate transparently inline in the network. Transparent inline operation results in minimal impact to the IT infrastructure.

Low latency

Network-based IPS devices must introduce a minimal amount of latency to network traffic. Low latency is often the most critical performance factor for network intrusion prevention.

High performance

A network-based IPS must exhibit many of the performance characteristics of switching and routing equipment, while simultaneously blocking threats to the network and the devices connected to it.

8.4.4 Security

Up-to-the-second security research must be incorporated into the IPS as rich security content, often in the form of logic or algorithms. Not all IPS vendors conduct the same caliber of research, resulting in security content that varies in effectiveness. As with the performance and protection components of an IPS, no single research methodology is adequate. For preemptive IPS solutions, research must encompass both proactive and reactive methods, covering both threats and vulnerabilities, global event monitoring, and information-sharing with other research organizations, industry consortium, and government entities.

Reactive research

The fact remains that hackers still manage to bypass many security devices, primarily because of the security industry's overdependence on reactive threat coverage.

Many security vendors who do not possess an internal security research team are forced to rely solely on security intelligence available in the public domain, including attack exploit code posted on public security websites, hacker websites, and vulnerability announcements released by software vendors. These IPS vendors are caught in a holding pattern, waiting for attack code to be publicized before they can update the protection in their IPS.

Collecting research on exploit tools commonly used by hackers, and studying different vulnerabilities after they are announced does provide insight into the nature of attacks, but it is not the only form of security research supporting an IPS. Hackers generally do not post their exploit code until they have already used it to break into a system, and by then it might be too late for the reactive security vendor and its customers. Plus, no security vendor relies entirely on the hacker community for education about stopping attacks. Preemptive protection against Internet attacks requires proactive research above and beyond what is made public after an exploit or vulnerability is released.

Proactive research

Sophisticated modern attacks move across the network at a rapid pace. If an IPS merely reacts to new threats after they appear, organization systems are likely to suffer negative impacts, such as corruption and downtime. Proactive security research is a pivotal requirement for preemptive protection. To conduct proactive security research, vendors must maintain a highly trusted and qualified team of security professionals who conduct primary research on the nature of vulnerabilities and attacks. Proactive research also requires extensive capital resources to acquire thousands of different types of hardware and software that are studied for vulnerabilities, and tested when new attacks appear. An IPS updated with proactive vulnerability-based research focuses on the weak spots that are targets of attack, rather than the actual attack payload.

Vulnerability-based protection is preemptive because it blocks any attack targeting the known weakness in the system, whether that attack has been seen before, or represents a new variant or unknown threat. IPS solutions powered by proactive research can also provide a viable alternative to the current patching crisis by offering a *virtual patch*. As software vulnerabilities continue to increase, so do the number of patches. Today, many organizations remain in an ongoing *triage mode* by trying to determine which critical patches to apply first. If an IPS has the benefit of proactive security content updates focused on vulnerabilities, the system can protect those vulnerabilities during the window of exposure between vulnerability announcement and patch application (hence the term virtual patch).

8.4.5 Reliability

Intrusion prevention is usually applied at critical network infrastructure points. Therefore, IPS failures have the potential to cause system outages. With crucial information and systems on the line, IPS solutions must be highly reliable with a long *mean time between failure* (MTBF).

Availability

At a minimum, a network IPS must not interfere with traffic malfunctions, or go into an offline state. To avoid this outcome, network IPS devices fail open, regardless of the network media.

8.4.6 Deployment

At the network level, IPS devices must scale to a large number of user sessions and transactions without disrupting organization continuity. IPS performance requirements and characteristics differ slightly depending on whether intrusion prevention is deployed in the network or within host-based systems, such as servers and desktops. In either case, performance remains a key purchase consideration to ensure that the IPS causes no disruption to applications residing on servers, desktops, or to application communication within the network.

8.4.7 Management

IBM Security Network IPS solutions can be centrally managed by using IBM Security SiteProtector (see Chapter 7, “Centralized management” on page 199) or locally managed from the Local Management Interface (LMI). All communication is encrypted using TLS/SSL. Extensive reporting, analysis, and alerting options are available in both management options.

8.4.8 Confidence

The X-Force conducts original, primary research on vulnerabilities and threats, which is applied to the IBM Security Network IPS in the form of security content updates known as X-Press Updates (XPU). The X-Force team is credited with discovering and mitigating more major software vulnerabilities since 1998 than all other commercial security research organizations combined.

8.5 Physical deployment model

In the following sections, we explain the different modes of operation, the high availability options, how to deploy the IBM Security Network IPS in a 10G environment, and the support for SCADA and VoIP network environments.

8.5.1 Intrusion prevention architecture

When designing an IPS architecture for your network, it is important to answer the following questions:

- ▶ How and where does my network traffic flow?
- ▶ How much traffic does my network transport?

Although the IBM Security Network IPS solutions do contain signatures capable of detecting SSL traffic, they are not able to inspect the encrypted packet's payload. In order to provide *Defense in Depth*, a combination of network, server and endpoint IPS protection is often proposed.

In this way, the deployment of a host-based IPS (for example, IBM Security Server for Linux®) on your web server allows inspection of SSL encrypted data before it is passed to the associated web application.

See Figure 8-19 for a typical deployment architecture.

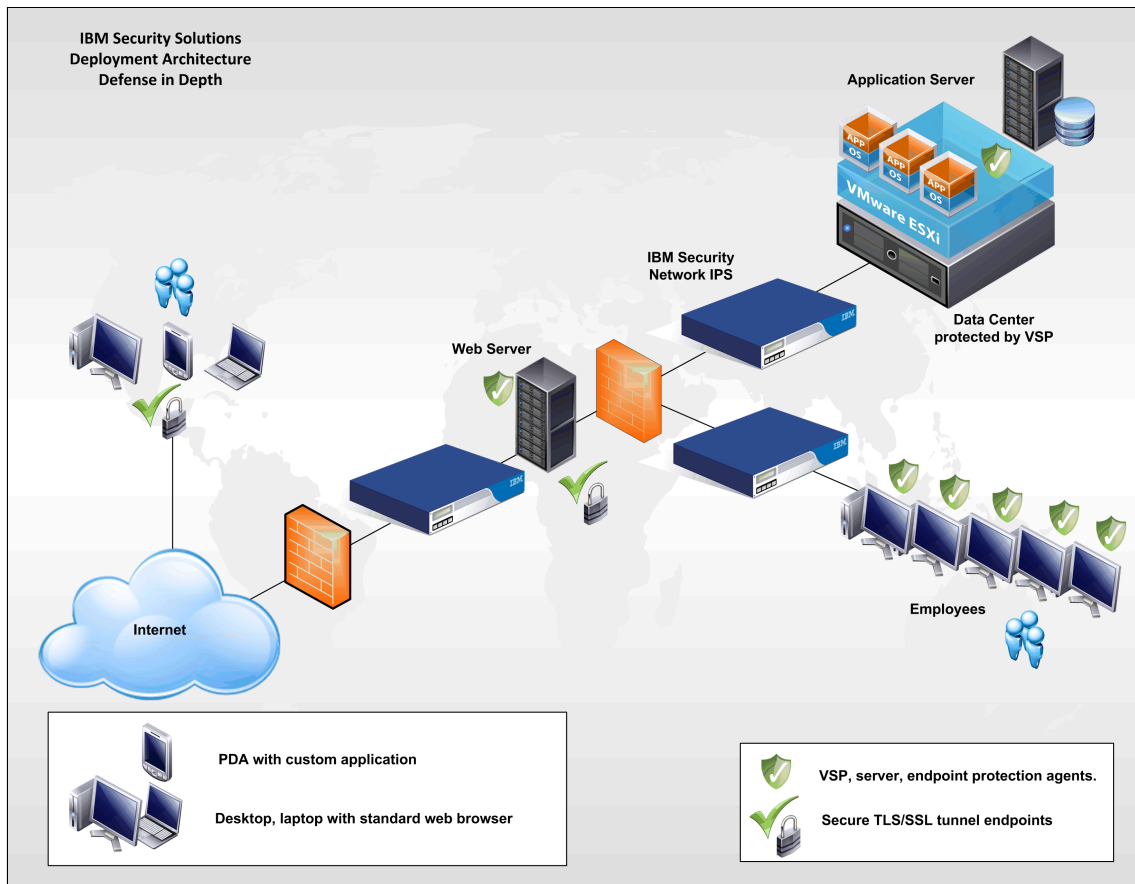


Figure 8-19 Typical IBM Security Network IPS deployment architecture

8.5.2 IBM Security Network IPS: Modes of operation

The IBM Security Network IPS supports three modes of operation:

- ▶ Passive monitoring
- ▶ Inline simulation
- ▶ Inline prevention mode

Passive monitoring mode acts like a traditional IDS, analyzing traffic using a promiscuous interface. In this mode, there is a manual way to respond to TCP attacks by sending TCP reset packets to both source and destination hosts to prevent certain attacks.

Inline simulation mode acts as a learning mode, alerting you about traffic that would otherwise have been blocked (in inline prevention mode).

Inline prevention mode actively blocks malicious and unwanted traffic according to the security policy that has been applied, without user intervention.

Figure 8-20 shows these three modes.

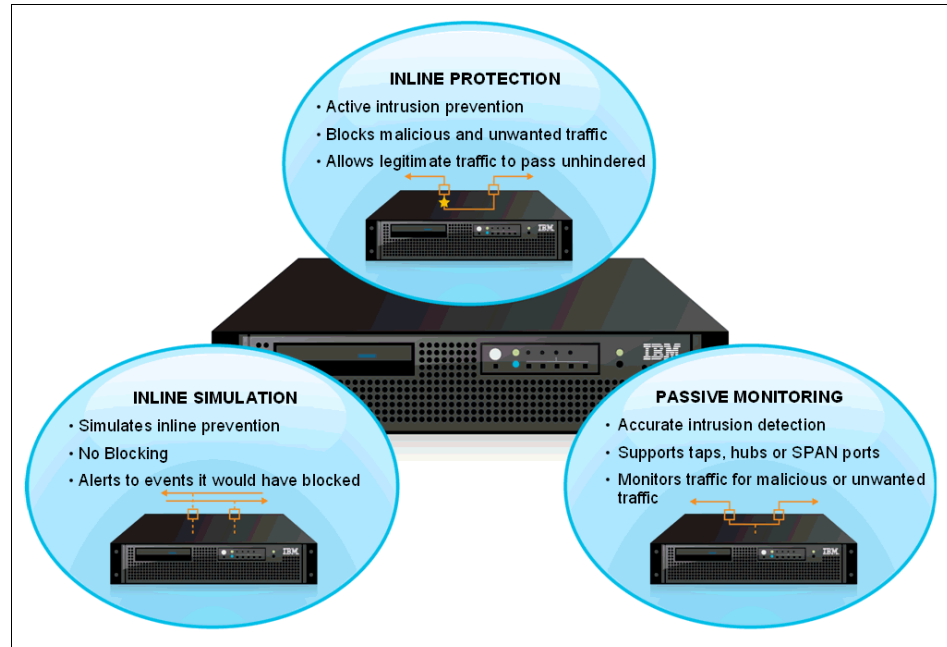


Figure 8-20 IBM Security Network IPS: Three modes of operation

Passive monitoring mode can be used in conjunction with the *TCP_Reset* port on the Network IPS appliances to block certain types of attacks that use TCP as the transport layer protocol and when it is not physically possible to place a Network IPS appliance between two devices. An example would be a TCP session between two hosts on the same Layer 2 switch.

By using Passive Monitoring mode and connecting one of the available monitoring ports on the Network IPS to a SPAN port on that switch and the *TCP_Reset* port to another available port on the same switch, TCP sessions between the two hosts that are deemed to be malicious can be reset.

It is important to note that the IBM Security Network IPS is a pre-configured appliance and operates effectively using an easily integrated configuration, whether deployed in passive monitoring, inline simulation, or inline prevention mode.

When installed as an inline prevention device, IBM Security Network IPS immediately protects the environment with minimal tuning. The default blocking policy protects against a large amount of threats, including hybrids such as MS Blaster, Welchia, Nachi, LoveSan, Code Red, Nimda, WebDAV, SQL Slammer, LASSER, Conficker, Stuxnet, and future worm propagations, without requiring advanced knowledge of the network infrastructure.

Layer two operation

IBM Security Network IPS operates at layer two (Data Link) of the OSI model, so it does not require an IP address on its monitoring interfaces. It acts as a bridge, and is invisible in the network. No network reconfiguration is required.

8.5.3 External bypass

The IBM Security Network Active Bypass⁹ intelligently incorporates two bypass modes:

- ▶ Active bypass capabilities provide maximum flexibility and deliver an uninterrupted communications session.
- ▶ Passive bypass capabilities deliver traditional static bypass.

⁹ For more information about the IBM Security Network Active Bypass, go to <http://www.ibm.com/software/tivoli/products/network-active-bypass/>.

See Figure 8-21 for a typical network configuration that shows traffic flows with and without bypass.




Figure 8-21 IBM Security Network Active Bypass traffic flows

The active bypass units can be configured to go from *inline* or active mode to bypass mode if a number (1 - 10) of heartbeats get lost. The fastest way to switch to bypass mode is to configure heartbeat=1. In addition, you can configure the maximum time allowed between heartbeat acceptance from 100 to 25500 ms.

In order for the bypass unit to switch back to *inline* or active mode, you can configure the same threshold as a number (1 - 10) of heartbeats. When you set heartbeat=1, the bypass unit switches back to active mode after only one heartbeat is accepted from the IPS appliance.

If the IPS appliance fails for any reason, the bypass is designed to ensure that the network remains functional and users have unimpeded access to important applications.

8.5.4 Copper and fibre connectivity

Each IBM Security Network IPS appliance can be fitted with either copper or fibre interfaces. Fibre connectivity is provided by either LD or SD connectors.

Connection details: The interfaces in the IBM Security Network IPS products are modular and use the *Small Form-Factor Pluggable (SFP)* format. These are also known as *mini-GBIC* connectors.

8.5.5 High availability

To ensure connectivity between two distinct networks, organizations often deploy redundant network architectures. These architectures make use of network protocols such as the *Spanning Tree Algorithm* or *High Speed Routing Protocol* (HSRP) to ensure that data loops do not occur between the two distinct networks.

More connection details: IBM Security Network IPS products do not actively participate in the protocol negotiations; they simply forward the packets.

However, in the event of a failure of the primary Network IPS, the redundant Network IPS still needs to be able to protect against attacks. To understand the various high availability options available with the IBM Security Network IPS products, it is imperative to firstly understand the two appliance failure modes: *fail_open* and *fail_closed*.

fail_open Traffic will pass through the failed appliance uninspected.

fail_closed Traffic will no longer pass through the failed appliance.

In the event of a failure of the IBM Security Network IPS GX4004 product, it fails in *fail_open* mode.

In the event of a failure of the IBM Security Network IPS GX5008, 5108, 5208, or GX6116 products, they fail in *fail_closed* mode.

In *fail_closed* mode, there are two options available to ensure that traffic continues to pass through your network:

- ▶ Use bypass units, as described in 8.5.3, “External bypass” on page 275.
- ▶ Configure a *pair* of appliances in a high availability (HA) cluster configuration.

Supported network configurations

High availability networks are typically configured in one of two ways:

- ▶ Active/passive configuration
Network traffic flows on the primary (active) network segment and the devices in that segment handle all of the network traffic. If one of these devices fail, the network traffic fails over to the secondary (passive) network segment, and the secondary devices take over.
- ▶ Active/active configuration
Both devices in the HA pair are active and handle traffic all of the time.

When deploying the IBM Security Network IPS appliances, there are two types of high availability embedded in firmware 4.1.

- ▶ Legacy high availability (legacy HA): Active/passive or active/active configuration
- ▶ Geographical HA (geo HA): Active/passive configuration

We explain these types in more detail in the following sections.

Legacy HA: Active/passive

Legacy HA has been available for a number of years in the IBM Security Network IPS product range and it has the following characteristics:

- ▶ Does not support asymmetric traffic.
- ▶ Determines the primary network path by network protocols (for example, HSRP and STA).
- ▶ Only works with GX5x08 and GX6116 (fail_closed devices).
- ▶ Does not share Dynamic Blocking Table (quarantine) or state information between IPS devices.
- ▶ Does not drop network flows if devices fail.
- ▶ Might drop some TCP packets while network protocols converge.
- ▶ Might drop some UDP packets while network protocols converge.

VoIP information: In a VoIP environment, there may be some loss of voice quality while the network protocols converge and when UDP packets are dropped.

To successfully deploy the legacy HA option, *mirroring traffic* links are used to pass the traffic from one Network IPS to its partner unit in the HA pair.

Port specifics: On a GX5x08 Series Network IPS, if ports A and B are used as the inline monitoring ports, ports C and D are used as the mirroring ports. Likewise, if ports E and F are used as the inline monitoring ports, ports G and H are used as the mirroring ports.

Both appliances in the HA pair process packets inline and block attack traffic that arrives on their inline monitoring ports, not on their interconnection/mirror ports. Both appliances also report events received on their inline monitoring ports to the SiteProtector management console.

Figure 8-22 shows a legacy HA configuration with two IBM Security Network IPS appliances configured as an HA pair.

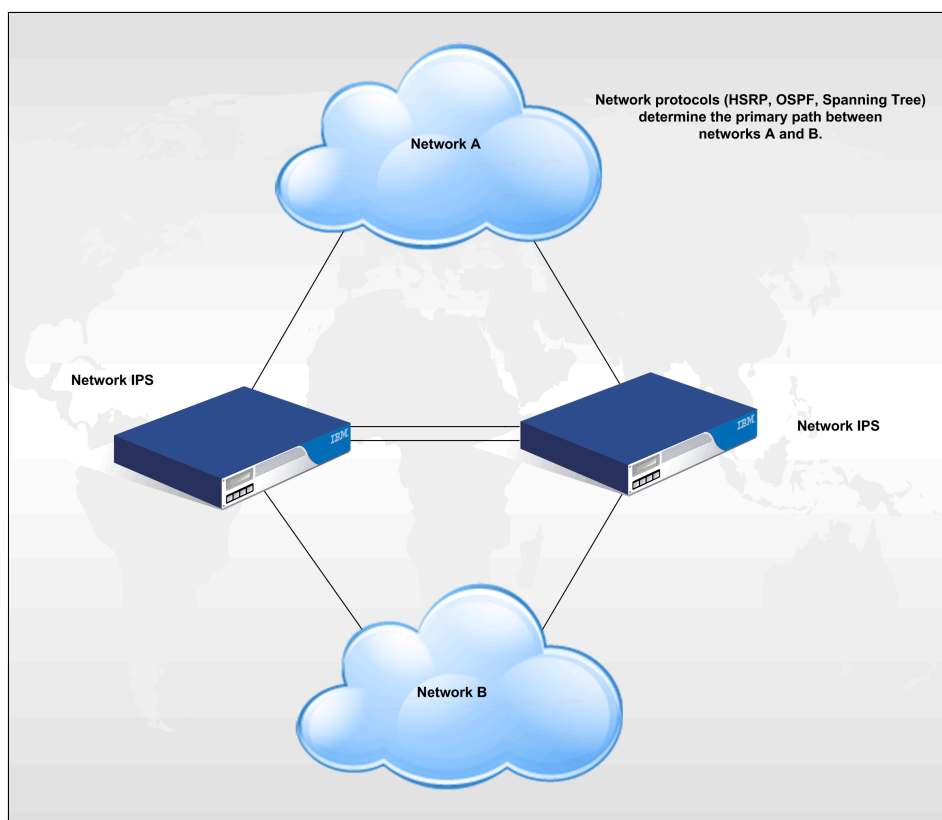


Figure 8-22 Legacy HA: Active/passive configuration

Configuration note: If you run the Setup Utility when the HA feature is enabled, you cannot modify the network settings.

In the event of a failure of the primary (active) network IPS, traffic is unable to flow from network A to network B via the primary network path. The network protocol being used renegotiates the (new) active path, as shown in Figure 8-23.

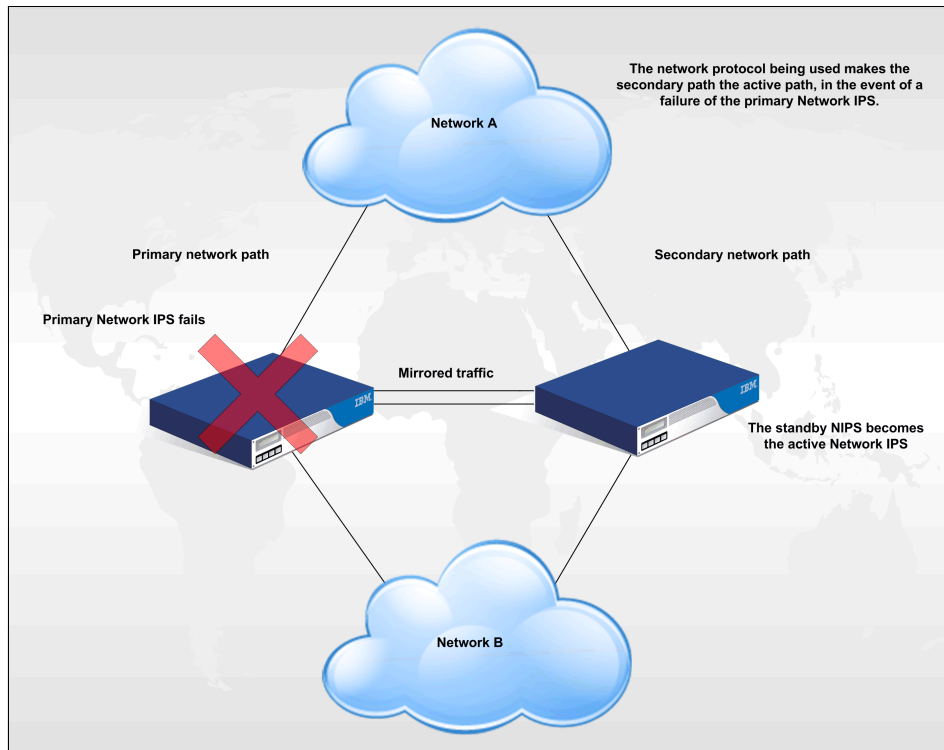


Figure 8-23 Failure of the primary network IPS

Legacy HA: Active/active

In the above example, the primary Network IPS was active and the secondary Network IPS was passive.

In active/active mode, both Network IPS appliances are inspecting live traffic at the same time. As a result, asymmetrically routed traffic is supported in this configuration, as long as the traffic passing through the monitoring interfaces is *mirrored* to the second Network IPS appliance in the HA cluster, as in the legacy HA example.

Geographical HA (geo HA)

Whereas legacy HA required the use of dedicated mirroring ports on each Network IPS, geographical HA makes use of the Network IPS management interfaces to exchange the information needed for the HA cluster to operate correctly. This configuration has the following advantage over legacy HA:

- ▶ Data exchange can take place over a routed network.
- ▶ Monitoring ports are not used up as mirroring ports.

Geographical HA has the following characteristics:

- ▶ Does not support asymmetric traffic.
- ▶ Exchanges Dynamic Blocking Tables via Network IPS management interfaces.
- ▶ Exchanges quarantine information between the active and standby devices.
- ▶ Does not support NAT between the geographically separated devices.

Both appliances process packets received from all redundant segments, but only need to block attack traffic that arrives on their inline ports when appropriate. Both appliances report events to the management console at all times. However, responses are processed only for events that are generated by packets that arrive on inline ports. Appliances process but do not block or report events that are generated by traffic that arrive on mirroring ports. As both appliances see all the traffic at all times, the failover time for response processing is eliminated. Both appliances maintain a current state, so if one HA network segment fails, the other appliance receives all the packets on its inline ports, resulting in events being generated as soon as the network fails over.

In an HA configuration, the appliance can only operate in either inline simulation or inline protection modes. Passive monitoring mode does not need to be supported because a passive configuration is not an inline deployment. Adapter level operation modes supported in normal mode are not supported in an HA configuration. If HA simulation mode is selected, all monitoring adapters are put in inline simulation mode automatically. If HA protection mode is selected, all monitoring adapters are put in inline protection mode automatically. In normal mode, each adapter can be configured to run in a different operational mode.

HA addresses the operation of inline appliances in a high availability environment. Again, note that all traffic is copied to each device through the mirrored traffic connections, because the Protocol Analysis Module (PAM) needs to see all the data to make a decision about whether traffic is malicious or not. In the event of a failover scenario, the IBM Security Network IPS that remains functioning in the network fabric contains all the data necessary to prevent attacks.

8.5.6 10 Gbps environments

In order to use IBM Security Network IPS solutions in a 10 Gbps (10G) environment, the IBM Security Network Controller product is required.

The IBM Security Network Controller product is a four-port 10 Gb aggregator switch with 24 Gb ports and active bypass switch capabilities, as shown in Figure 8-24.

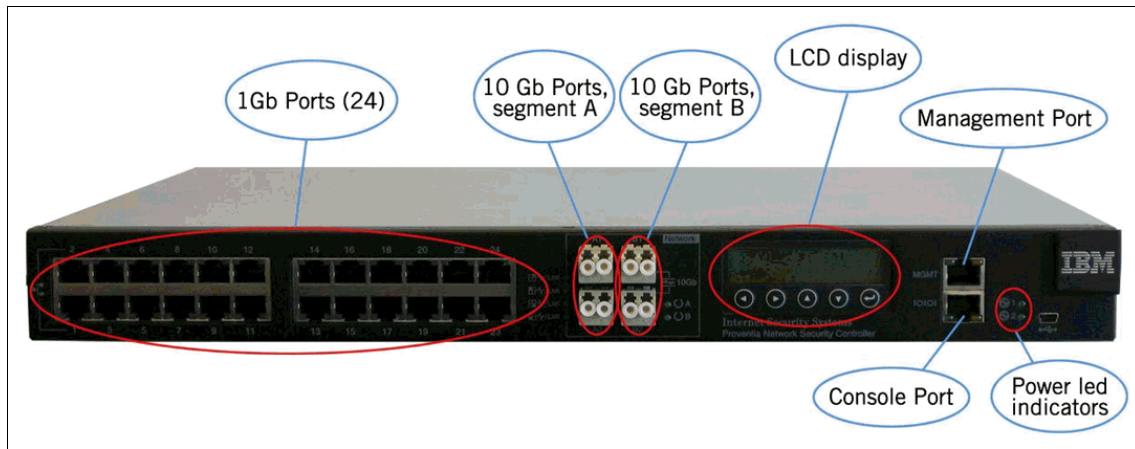


Figure 8-24 IBM Security Network Controller

The combination of 1 and 10 Gb interfaces provides a seamless connection between IBM Security Network IPS appliances and 10G networks.

See Figure 8-25 for a typical deployment architecture.

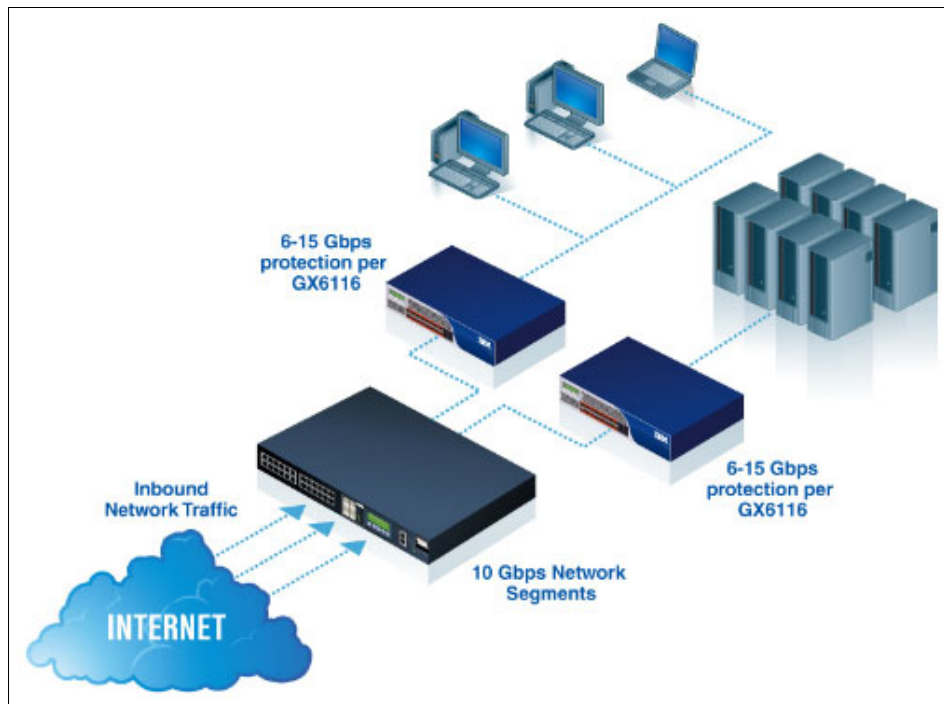


Figure 8-25 Typical 10G deployment of the IBM Security Network Controller

8.5.7 SCADA environments

IBM Security Network IPS solutions are designed to work in Supervisory Control and Data Acquisition¹⁰ (SCADA) environments. Examples of SCADA protocols supported within PAM include MODBUS and DNP3.

8.5.8 VoIP environments

One of the attack categories supported by PAM is the Voice over IP (VoIP) set of protocols. SIP and SCCP are supported, and there are currently 39 different attack signatures for these two protocols alone.

¹⁰ For more information about Supervisory Control and Data Acquisition Systems, go to http://www.ncs.gov/library/tech_bulletins/2004/tib_04-1.pdf.

Real world test: Proof-of-concept testing in 2010 by the world's leading VoIP vendor showed that the IBM Security Network IPS solution provided real-time inspection of VoIP traffic with no jitter or VoIP packet delay.

8.5.9 FIPS 140-2 certification

All the IBM Security Network IPS appliances (GX4004, GX5008, GX5108, GX5208, and GX6116) and the associated software and encryption development have passed the FIPS 140-2 certification.

8.6 IBM Tivoli Netcool Configuration Manager

IBM Tivoli Netcool Configuration Manager¹¹ automates network configuration management tasks, controls network device access, and ensures network policy compliance. To discover more about Tivoli Netcool Configuration Manager integration, refer to *Integration Guide for IBM Tivoli Netcool/OMNIBus, IBM Tivoli Network Manager, and IBM Tivoli Netcool Configuration Manager*, SG24-7893.

Tivoli Netcool Configuration Manager allows network operators and administrators to automate routine network configuration management tasks, enhance network security by controlling access by users, devices and commands, and maintain the real-time state of the network. The following list details the main feature set provided by IBM Tivoli Netcool Configuration Manager:

- ▶ Allows network device configuration through multiple modes, supporting users of various skill levels, from expert network engineers to the novice operator.
- ▶ Manages the complete device configuration life cycle: Baseline, search, configure, test, approve, track, and rollback of network device configurations.
- ▶ Provides an intuitive user interface so that all relevant device information (for example, hardware, configuration history, changes, and activities) and frequently performed functions are no more than a click away.
- ▶ Limits user access to part of a device configuration or selected commands, which is an essential feature for virtualized or multi-service environments, or where there is a need for greater device control.

¹¹ For more information about Tivoli Netcool Configuration Manager, go to <http://www.ibm.com/software/tivoli/products/netcool-configuration-manager/>.

- ▶ Provides a reusable policy framework for configuration policy validations of regulatory mandates, security, and operational policies (for example, engineering standards).
- ▶ Enables rapid integration with third-party applications, for example, IT automation, Configuration Management Database (CMDB), network management, help desk, and enterprise workflow. Every configuration function is exposed through open APIs.
- ▶ Operates on a variety of server platforms, including Red Hat Enterprise Linux and Sun Solaris, as well as in a virtual VMWare environment. It also requires an Oracle relational database to be installed for persistent data storage.

8.6.1 Optimizing complex network environments

There is no question that in today's highly instrumented, interconnected, and intelligent world, the technology infrastructures that support and connect the enterprise grow more complex every day. How do we handle that complexity? How do we discover the components of the enterprise network? How do we understand the network topology? And how do we optimize the network so the business can build a competitive advantage from its investments?

Enterprise network management may begin with the pieces, that is, the routers, switches, hubs, and other devices that tie the network together. But network administrators need to see the big picture. They need ways to meet complex challenges posed by the introduction of new technologies, mergers and acquisitions, compliance requirements, and rising customer expectations. They need sophisticated tools and meaningful insights that allow the enterprise to better use its network to grow services, enhance the customer experience, become more agile and flexible, boost operational efficiency, and reduce costs.

Tivoli Netcool Configuration Manager provides automation capabilities that enable businesses to better control, manage, and grow their networks. Scalable to tens of thousands of devices that a large, multisite, and multivendor enterprise network can have, these solutions provide a critical view into the real-time state of devices. They support configuration and change management across the complete network life cycle. They help ensure compliance, security, and effective resource provisioning as the network evolves to support today's dynamic business.

A functional diagram of the solution is shown in Figure 8-26.

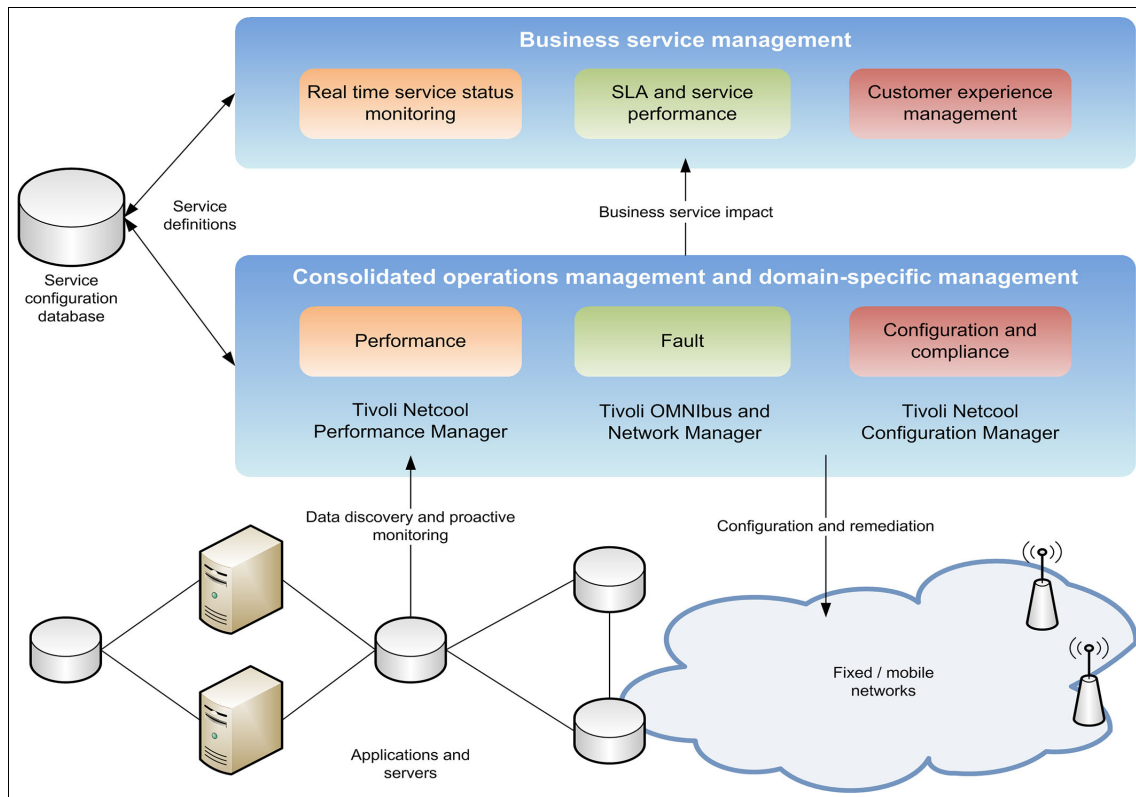


Figure 8-26 IBM Tivoli Netcool Configuration Manager: Managing the impact of change

8.6.2 Conclusion

IBM Tivoli Netcool Configuration Manager works together with the IBM Tivoli Integrated Service Management portfolio to deliver comprehensive solutions for automating and simplifying configuration and change management, demonstrating regulatory compliance, speeding deployments, and reducing operational costs.

Real-time visibility and monitoring help ensure network availability and quality across the network management life cycle. The IT department can both increase productivity for routine functions and take proactive steps to building a strategic network that will help the business move forward.

The ability to confidently respond to changing business needs can give today's enterprises a competitive edge, helping to ensure that the network infrastructure is configured to support business goals.

8.7 IBM WebSphere DataPower

Network Intrusion Prevention devices provide a significant level of protection against threats to the network and server infrastructure. In addition, application level protection mechanisms are emerging that protect specific application level threats. With the advent of service-oriented architectures (SOAs), new avenues of opportunities are available to attackers and these areas need to be protected.

IBM WebSphere DataPower SOA Appliances represent an important element in the holistic approach of IBM Security to service-oriented architecture (SOA) environments. These appliances are purpose-built, easy-to-deploy network devices that can simplify, secure, and accelerate your XML and web services deployments while providing additional SOA functionality. These appliances offer an innovative and pragmatic approach to harness the power of SOA. Through their use, you are able to enhance the value of your existing application, security, and network investments. The emergence and proliferation of XML and web services has seen an increase in the middleware infrastructure required to support them. An important component in this middleware architecture is the enterprise service bus (ESB), a collection of runtime components that provide services such as routing, transformation and bridging, management, security, and other control functions. Although XML and SOAP enable a rich application-aware communication, their emergence has resulted in challenges in terms of consumability, performance, and especially security.

WebSphere DataPower SOA Appliances provide the ability to understand and act upon application data as it traverses the network. Although this application awareness is not, in itself, a new networking concept, XML has accelerated its appeal and difficulty. That is, application awareness comes with many security, complexity, and performance challenges. As a result, a new genre of hardened software, hardware, and XML-centric appliances have arisen to bridge this gap.

These appliances focus on providing consumability, performance, and hardened security. They can extend the ESB into the network and also provide an SOA gateway for business-to-business integration.

8.7.1 XML and web services network security threats

The advent of SOA has created a common communication framework to understand and operate on application data unlike anything that has been seen before. With self-describing XML, intermediaries are able to extract portions of the data stream and effect application-aware policies. Unfortunately, this has also enabled a new opportunity for malicious attacks. As XML regularly flows from client to enterprise through IP firewalls without much impediment, the obvious place to attack is in the application data stream itself, that is, XML.

Although we are just beginning to understand the repercussions of these types of attacks, they are emerging. XML denial-of-service (XDoS) attacks seek to inject malformed or malicious XML into middleware servers with the goal of causing the server to churn away valuable cycles processing the malicious XML. Enterprise-ready application servers are susceptible to many of these types of attacks, leaving open a security hole that must be closed.

8.7.2 IBM WebSphere DataPower and meeting SOA challenges

The IBM WebSphere DataPower SOA Appliances family contains rack-mountable network devices that overcome many of the challenges facing SOA and XML today. At a high level, the IBM WebSphere DataPower SOA Appliances offer:

- ▶ 1U (1.75-inch thick) rack-mountable, purpose-built network appliances, with a tamper-proof enclosure.
- ▶ XML/SOAP firewalls, field-level XML security, data validation, XML web services access control, and service virtualization.
- ▶ Lightweight and protocol-independent message brokering, integrated message-level security and fine-grained access control, and the ability to bridge mission-critical transaction networks to SOAs and ESBs.
- ▶ High performance, multi-step, wirespeed message processing, including XML, XSLT, XPath, and XSD.
- ▶ Centralized web services policy and service-level management.
- ▶ WS-* standard support, such as WS-Security, SAML 1.0/1.1/2.0, portions of Liberty Alliance protocol, WS-I Basic Security Profile, WS-Federation, WS-Trust, XKMS, Radius, XML Digital Signature, XML-Encryption, WSDM, WS-SecureConversation, WS-Policy, WS-SecurityPolicy, WS-Policy Framework, WS-ReliableMessaging, SOAP 1.1 and 1.2, WSDL, UDDI, and others.
- ▶ Transport layer flexibility, which supports HTTP/HTTPS, MQ, SSL, FTP, and others.

- ▶ Scalable, wirespeed, any-to-any message transformation, such as arbitrary binary, flat text, and XML messages, which include COBOL Copybook, CORBA, CICS, ISO 8583, ASN.1, EDI, and others.
- ▶ Consumable simplicity: An easy-to-install and easy-to-maintain network appliance that can satisfy both application and network operational groups, supporting current and emerging standards, as well as XML web services standards in an easily integrated fashion.
- ▶ Enhanced security: Key support includes, but is not limited to, XML/SOAP firewall and threat protection, field-level XML security, data validation, XML web services access control, service virtualization, and SSL acceleration.
- ▶ Acceleration: A drop-in solution that can streamline XML and web service deployments, helping lower the total cost of ownership and accelerating return on your assets as you continue to move to SOA. WebSphere DataPower SOA Appliances are purpose-built hardware devices capable of offloading loads off of overtaxed servers by processing XML, web services, and other message formats at wirespeed.

8.7.3 The IBM WebSphere DataPower SOA Appliance product line

The product line consists of three appliances. The higher model numbers are a functional super set of lower-numbered appliances. When examining the model numbers, keep in mind that A stands for acceleration, S stands for security, and I stands for integration.

All appliances share a basic common engine as well as management/monitoring interfaces.

IBM WebSphere DataPower XML Accelerator XA35

The most basic appliance, the XA35, delivers a number of routing and transformation features. It supports routing based on IP, TCP, HTTP, XML, and SOAP criteria. Routing tables can be defined (or queried) to determine the appropriate routing action. Load balancing algorithms, such as least-used and round-robin, are also available. Additionally, the XA35 can throttle or adjust its request rate based on routing criteria to employ traffic shaping.

From a transformation perspective, the XA35 provides basic XSLT processing. Built on the underlying premise of all appliances, special-purposed hardware converts one XML schema to another at wirespeed. XSLT 1.0 and XPath 1.0 support (with some 2.0 support) is available. The appliance is able to use and apply internal as well as external schema.

This product can help speed up common types of XML processing by offloading this processing from servers and networks. It can perform XML parsing, XML Schema validation, XPath routing, Extensible Stylesheet Language Transformations (XSLT), XML compression, and other essential XML processing with wirespeed XML performance.

The XA35 has the following features:

- ▶ Unmatched performance: DataPower's purpose-built message processing engine can deliver wirespeed performance for both XML-to-XML and XML-to-HTML transformations with increased throughput and decreased latency.
- ▶ Ease of use: The XA35 provides drop-in acceleration with virtually no changes to the network or application software. No proprietary schemas, coding, or APIs are required to install or manage the device. In addition, it supports popular XML Integrate Development Environments (IDEs) to help reduce the number of hours spent in the development and debugging of XML applications.
- ▶ Helps reduce infrastructure costs: Unlike simple content switches that only redirect business documents, the DataPower XML Accelerator XA35 fully parses, processes and transforms XML with wirespeed performance and scalability to help reduce the need for stacks of servers. The XA35 also supports accelerated SSL processing to help further lessen the load on server software.
- ▶ Helps cut development costs: The XA35 can enable multiple applications to use a single, uniformed XML processing layer for all XML processing needs. By standardizing high-performance hardware appliances, enterprises can deploy sophisticated applications while helping eliminate unnecessary hours of application debugging and tuning for marginal performance gains.
- ▶ Intelligent XML processing: In addition to wirespeed processing, appliances support XML routing, XML pipeline processing, XML compression, XML/XSL caching, as well as other intelligent processing capabilities to help manage XML traffic.
- ▶ Advanced management: The XML Accelerator model XA35 provides real-time visibility into critical XML statistics, such as throughput, transaction counts, errors, and other processing statistics. Data network level analysis is provided, and includes server health information and traffic statistics, as well as management and configuration data.

IBM WebSphere DataPower XML Security Gateway XS40

The XS40 offers security capabilities, such as XML Encryption and XML digital signature processing, in addition to the accelerator capabilities of the XA35. Supporting WSS 1.0, WS-Trust, and WS-SecureConversation, the XS40 is the foundation upon which Web Services Security can be deployed in an enterprise. Support exists for authorization (access) checking to offboard entities such as the IBM Tivoli Access Manager, IBM Tivoli Federated Identity Manager, RSA, Netegrity, Oblix, and more. There is partial SAML 2.0 support as well as the ability to extract and use security credentials from LDAP, RADIUS, and XKMS. But perhaps the most significant security feature of the XS40 is its firewall capabilities. In addition to filtering input traffic (as is done in the XA35 routing capabilities), the XS40 firewall enables XML Denial of Service protection. Used in conjunction with schema validation and well-formedness checking, the XS40 is an integral part of an enterprise's network security. This appliance provides a security-enforcement point for XML and web services transactions, including encryption, firewall filtering, digital signatures, schema validation, WS-Security, XML access control, XPath, and detailed logging.

The XS40 has the following features:

- ▶ **An XML/SOAP firewall:** The DataPower XML Security Gateway XS40 filters traffic at wirespeed, based on information from layers 2 through 7 of the protocol stack, from field level message content and SOAP envelopes to IP address, port/host name, payload size, or other metadata. Filters can be predefined with an easy point-and-click XPath filtering GUI, and automatically uploaded to change security policies based on the time of day or other triggers.
- ▶ **XML/SOAP data validation:** With its unique ability to perform XML Schema validation as well as message validation at wirespeed, the XS40 ensures that incoming and outgoing XML documents are legitimate and properly structured. This protects against threats such as XML Denial of Service (XDoS) attacks, buffer overflows, or vulnerabilities created by deliberately or inadvertently malformed XML documents.
- ▶ **Field-level message security:** The XS40 selectively shares information through encryption/decryption and signing/verification of entire messages or of individual XML fields. These granular and conditional security policies can be based on nearly any variable, including content, IP address, host name, or other user-defined filters.
- ▶ **XML web services access control:** The XS40 supports a variety of access control mechanisms, including WS-Security, WS-Trust, X.509, SAML, SSL, LDAP, RADIUS, and simple client/URL maps. The XS40 can control access rights by rejecting unsigned messages and verifying signatures within SAML assertions.

- ▶ Service virtualization: XML web services require companies to link partners to resources without leaking information about their location or configuration. With the combined power of URL rewriting, high-performance XSL transforms, and XML/SOAP routing, the XS40 can transparently map a rich set of services to protected back-end resources with high performance.
- ▶ Centralized policy management: The XS40's wirespeed performance enables enterprises to centralize security functions in a single drop-in device that can enhance security and help reduce ongoing maintenance costs. Simple firewall functionality can be configured via a GUI and be running in minutes, and using the power of XSLT, the XS40 can also create sophisticated security and routing rules. Because the XS40 works with leading Policy Managers, it is an ideal policy execution engine for securing next-generation applications. Manageable locally or remotely, the XS40 supports SNMP, script-based configuration, and remote logging to integrate seamlessly with leading management software. Its emerging support for WS-Policy and WS-SecurityPolicy further augments these capabilities.
- ▶ Web services management/service level management: With support for Web Services Distributed Management (WSDM), Universal Description, Discovery, and Integration (UDDI), Web Services Description Language (WSDL), Dynamic Discovery, and broad support for service level management configurations, the XS40 natively offers a robust web services management framework for the efficient management of distributed web service endpoints and proxies in heterogeneous SOA environments. Service level management (SLM) alerts and logging, as well as pull and enforce policies, help enable broad integration support for third-party management systems and unified dashboards, in addition to robust support and enforcement for governance frameworks and policies.

IBM WebSphere DataPower Integration Appliance XI50

The XI50 adds integration capabilities to its accelerator and security abilities. It enables a key concept of any-to-any transformation, where data can be received in any format over any protocol and be converted to any other format over any other protocol using DataPower core high-performance transformation technology. In DataPower everything is a transformation.

Supported protocols include HTTP, HTTPS, and MQ. Formats include SOAP/XML, EDI, CICS COBOL Copybook, Corba, ISO 8583, CSV, ASN.1, ebXML, and more. This appliance provides transport-independent transformations between binary, flat text files, and XML message formats. Visual tools are used to describe data formats, create mappings between different formats, and define message choreography. This appliance can transform binary, flat text, and other non-XML messages to help offer an innovative solution for security-rich XML enablement, ESBs, and mainframe connectivity.

The XI50 has the following features:

- ▶ Any-to-any transformation engine: XI50 can parse and transform arbitrary binary, flat text, and XML messages, including EDI, COBOL Copybook, ISO 8583, CSV, ASN.1, and ebXML. Unlike approaches based on custom programming, DataPower's patented DataGlue technology uses a fully declarative, metadata-based approach.
- ▶ Transport Bridging: With support for a wide array of transport protocols, the XI50 is capable of bridging request/response flows to and from protocols such as HTTP, HTTP MQ, SSL, IMS™ Connect, FTP, and more
- ▶ Integrated message-level security: XI50 includes mature message-level security and access control functionality. Messages can be filtered, validated, encrypted, and signed, helping to provide more secure enablement of high-value applications. Supported technologies include WS-Security, WS-Trust, SAML, and LDAP.
- ▶ Lightweight message brokering: Sophisticated multi-step message routing, filtering, and processing.
- ▶ Support for JavaScript Object Notation (JSON) and REpresentational State Transfer (REST) applications.
- ▶ Multiple synchronous and asynchronous transport protocols.
- ▶ Detailed logging and audit trail, including non-repudiation support.

More information about the IBM WebSphere DataPower Appliances can be found in *DataPower Architectural Design Patterns: Integrating and Securing Services Across Domains*, SG24-7620.

To see how DataPower and the IBM Security Network IPS appliances can be deployed together, refer to Figure 8-27. As shown, a firewall protects the network from traditional threats. The XML Firewall (implemented as the DataPower XS40) provides application level security functions directly on the web services, or XML payload. Additionally, in this implementation, the XS40 can provide the SSL termination point for inbound transactions. An IBM Security Network IPS device behind the DataPower appliance then can inspect for additional threats. This architecture provides a holistic, defense in-depth approach across both the network and the application layers.

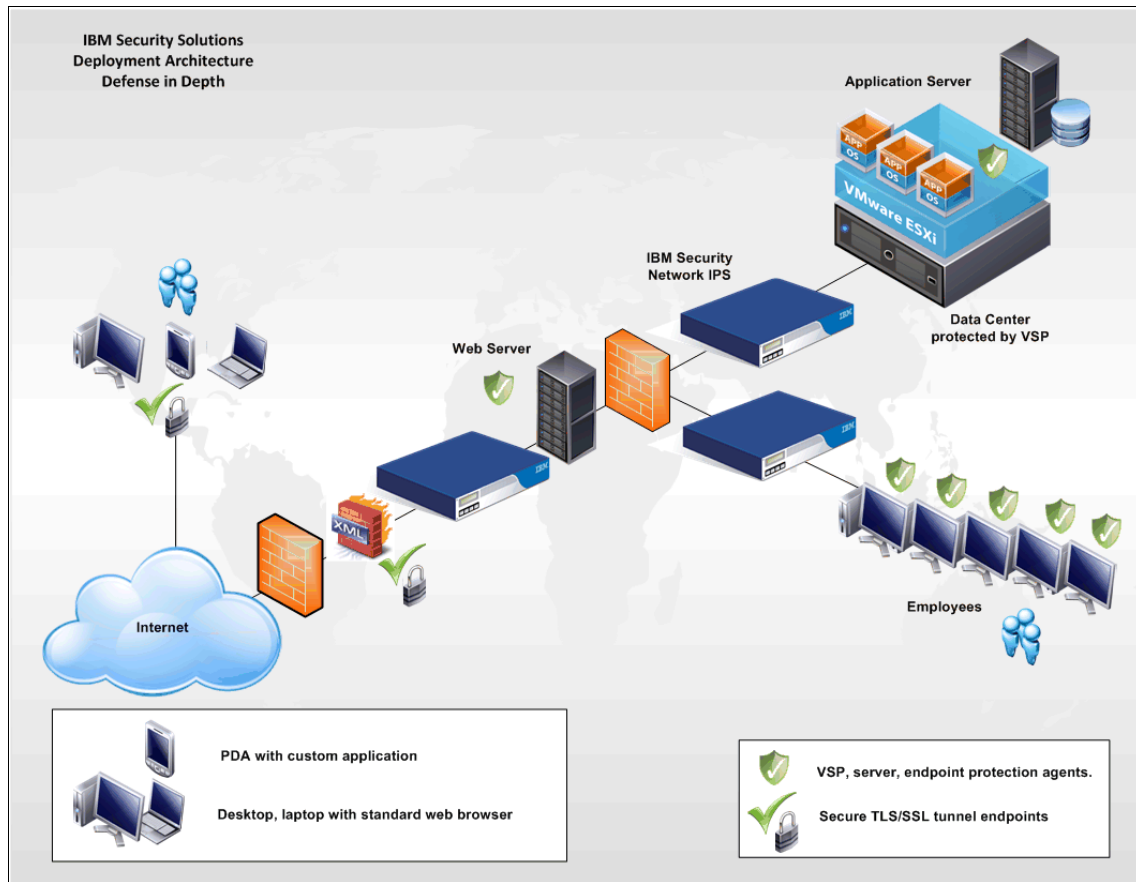


Figure 8-27 DataPower and Network IPS: Typical deployment architecture

8.8 IBM Lotus Protector for Mail Security

Email is one of the primary communication methods for an enterprise and spam continues to find its way into inboxes daily, sometimes outnumbering legitimate email. The fight against spam can negatively impact productivity and strain network and server capacities, affecting your users and your system administrators. Additionally, malware is often distributed through inbound email, infecting user workstations and their productivity as it passes from mailbox to mailbox.

Optimized for the IBM Lotus Notes® and IBM Domino® platform, the IBM Lotus Protector for Mail Security solution lightens the spam burden in several ways. First, it can be quickly configured to block spam using either default or custom content filtering policies. Content filtering innovation is provided by the X-Force team. The X-Force team routinely monitors new spam techniques and distribution methods. By default, the Lotus Protector for Mail Security solution checks each hour for updates from IBM that include new spam signatures and potentially dangerous URLs. As a result, Lotus Protector for Mail Security technology helps you keep ahead of the latest spam trends, including social penny stock schemes and image-based spam.

In addition, the solution includes dynamic host reputation filtering technology, which uses sophisticated IBM research on where spam is likely to originate, to help stop spam before it even reaches your system. By analyzing the source IP address on each incoming email, it can make a mathematical judgment about whether or not the source of the email is reliable. When an email is deemed to be coming from an unreliable source, the connection is dropped before the email is delivered. As a result, you can help reduce the system load associated with managing spam by preventing spam from reaching your filter in the first place.

At the IBM Security Solutions Global Data Center, IBM maintains a security database containing more than 95 million web pages and relevant spam signatures to date. IBM operates spam collectors worldwide using email accounts known as “honey pots”, which receive hundreds of thousands of confirmed spam emails every day. Data gathered from these messages is fed into the Global Data Center. In addition, IBM has established a network of trusted IBM Business Partners and IBM corporate users that contribute spam data to the database. IBM spam protection technologies use data gathered from all of these sources to increase the efficiency of their spam filtering. Plus, with the optional Spam Learn feature, IBM offers anonymous, automated reporting of new spam back to the Global Data Center to be included in the database.

Lotus Protector for Mail Security technology combines multiple analysis modules for greater customization, enabling you to define policies or tailor modules to help meet legal and regulatory compliance requirements for data. Messages can also be scanned for offensive words, customizable keywords, and attachment types, and specialized analysis capabilities help prevent sensitive information such as Social Security and credit card numbers from leaving your network. In addition, the phishing module provides a separate analysis technique to protect your employees against email messages that target their personal information.

Granular policy control includes simple rules-based policy creation (enabling you to take action based on factors such as whom, what, and when) and more than 10 different customizable action types, such as modifications and notifications. Policies can be applied globally, by user group or by individual user. Plus, the Lotus Protector for Mail Security solution supports Lightweight Directory Access Protocol (LDAP), including Lotus Domino and Microsoft Active Directory technologies. Users can control their own allow and block lists, giving them personalized control over their own spam preferences. They can also view and control their quarantined messages if granted permission by an administrator.

Beyond spam control, the Lotus Protector for Mail Security solution is equipped with advanced protection technologies to provide security features that are ahead of the threat. With the IBM Security Network IPS engine and IBM Virtual Patch technology, the application supports the vital security necessary in today's IT environments.

Support for the Transport Layer Security (TLS) protocol enables you to automatically encrypt all emails between your company and trusted partners and suppliers. By establishing mutual public certificates on your server, you can make sure that communication between your company and these organizations is protected. The message transport agent at the edge of your network automatically encrypts all emails to and from such organizations, providing a seamless user experience.

The solution's recipient verification technology and queuing mechanism helps protect your mail server from zero-day attacks, including denial-of-service and directory harvest attacks. Many spammers direct spam at a particular domain simply by guessing at user names or naming conventions. Recipient verification technology helps minimize the effects of this practice by confirming that the specific user name to which each email is addressed actually exists before accepting the message. Any message that is addressed to an unknown recipient is rejected before the connection is accepted, helping to save valuable bandwidth.

The queuing mechanism is designed to provide multiple levels of protection against spam-based, denial-of-service attacks. The application has two predefined thresholds for its unchecked queue, which begins to grow during a denial-of-service attack. When the total number of messages in the unchecked queue reaches the first threshold, the application begins throttling new Simple Mail Transfer Protocol (SMTP) connections based on a predefined period of time. When the number of messages in the unchecked queue reaches the second threshold, all new SMTP connections are answered with a “temporarily not available” message and a request to try again later, based on SMTP standards. Typical spam bots cannot handle this type of rejection and will fail at this point, whereas valid SMTP servers will try again after a predefined period of time.

The Lotus Protector for Mail Security solution includes remote malware detection, which is automatically distributed to your application via signature updates to the filter database. In addition, behavioral genotype and signature antivirus technologies take action against suspicious code for known and unknown viruses. This technology analyzes both incoming and outgoing email in parallel with the application’s antispam features.

8.9 Conclusion

In this chapter, we explained the next generation IBM Security Network IPS solution and how it delivers *ahead of the threat* preemptive network protection in a wide variety of deployment scenarios. We also introduced three complementary network security products: IBM WebSphere DataPower SOA Appliances, IBM Tivoli Network Configuration Manager, and IBM Lotus Protector for Mail Security.



Host security solutions

In this chapter, we describe the IBM host security offerings. First we examine the Tivoli Endpoint Manager platform, and then we take a closer look at the host security solutions.

Host security solutions are composed of agents for servers, desktops, and mobile computers. In this chapter, we explain which operating systems are supported and which feature sets are available for the different operating systems in the following sections:

- ▶ “IBM Tivoli Endpoint Manager” on page 300
- ▶ “Proventia Desktop Endpoint Security” on page 318
- ▶ “IBM Security Server Protection” on page 322

9.1 IBM Tivoli Endpoint Manager

Organizations are more widely distributed than ever before, something that can make systems management tasks, such as distributing software and patches, extremely challenging.

The IBM Tivoli Endpoint Manager platform is architected for today's highly diverse, distributed, and complex IT environments. It provides real-time visibility and control through a single infrastructure, a single agent, and as single console for systems life cycle management, endpoint protection, and security configuration and vulnerability management. This platform enables organizations to securely manage their global IT infrastructures faster and more accurately, resulting in improved governance, control, visibility, and business agility. Plus, it gives organizations the ability to handle tomorrow's unforeseen challenges.

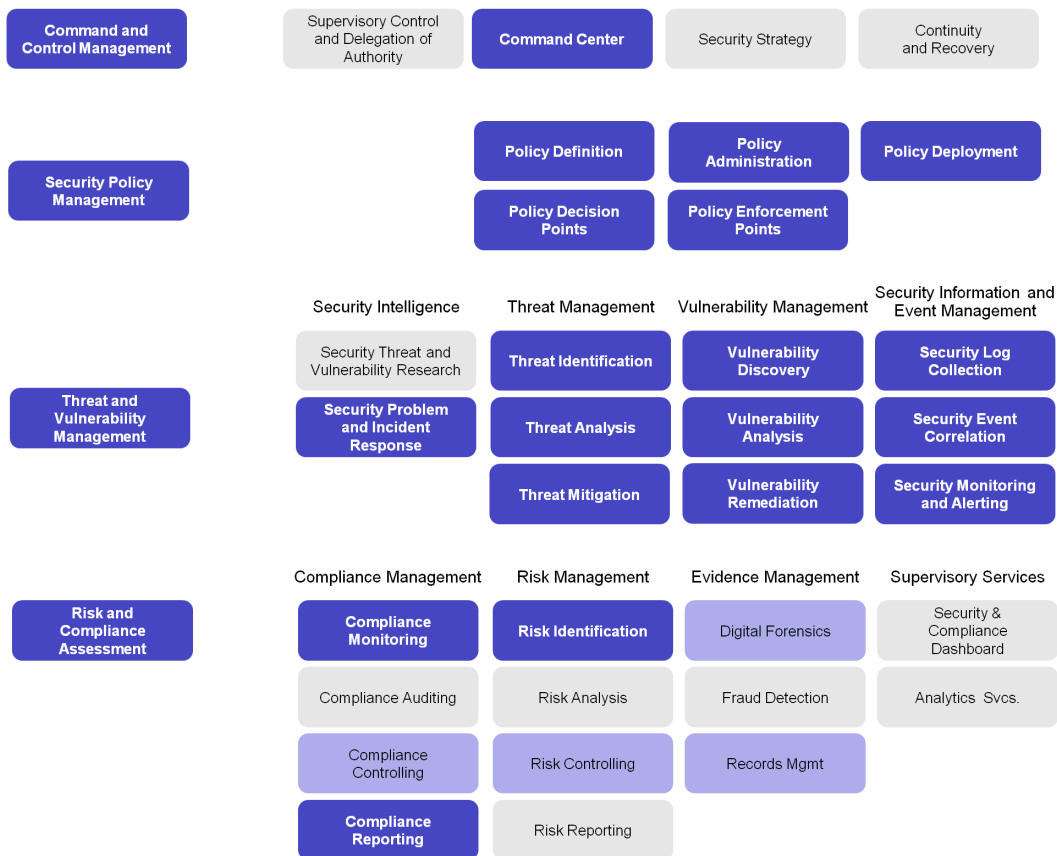
To understand how the security capabilities of Tivoli Endpoint Manager can be mapped to the IBM Security Blueprint¹, see Figure 9-1 on page 301. This diagram shows the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using Tivoli Endpoint Manager. This functional highlighting is applicable for the infrastructure service components as well.

Besides the fully highlighted elements, Figure 9-1 on page 301 also shows some medium highlighted elements. Although Tivoli Endpoint Manager can be used to address such components to some degree, the respective area of coverage is not considered a core function of the product and thus is considered to be limited.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 9-1 on page 301 can be used as a quick reference of the functional security management aspects of Tivoli Endpoint Manager. This reference can help determine which functions of a solution can be covered by selecting this product.

¹ For a detailed discussion of the elements, refer to Chapter 2, "The components of the IBM Security Blueprint" on page 31 and Chapter 3, "The Network, Server and Endpoint solution pattern" on page 93.

Foundational Security Management Component and Sub-Components



Security Services and Infrastructure

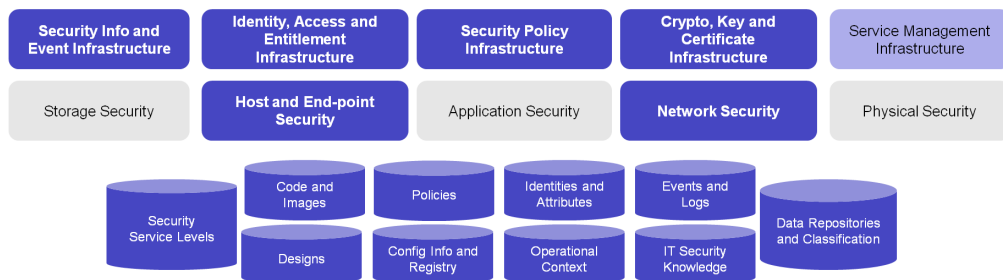


Figure 9-1 Mapping of Tivoli Endpoint Manager to the IBM Security Blueprint

Tivoli Endpoint Manager is a multilayered technology platform that acts as the central nervous system of your global IT infrastructure. As a dynamic, content-driven messaging and management system, the technology distributes the work of managing IT infrastructures out to the managed devices themselves. As a result, the Tivoli Endpoint Manager platform is able to operate in near real time, delivering the scalability and performance that large organizations demand.

Acquisition information: As of February 1, 2011, BigFix has been officially integrated into IBM. All BigFix products have been integrated into the Tivoli Software portfolio and rebranded as IBM Tivoli Endpoint Manager.

IBM Tivoli Endpoint Manager is derived from the *BigFix*^a Unified Management Platform. Because IBM has only recently acquired this technology, this chapter sometimes refers to the platform as both Tivoli Endpoint Manager and BigFix, for example, in representative screen captures.

a. <http://www.bigfix.com/>

9.1.1 Platform

The Tivoli Endpoint Manager solution has four packages available as part of the offering:

- ▶ Tivoli Endpoint Manager for Lifecycle Management
- ▶ Tivoli Endpoint Manager for Security and Compliance
- ▶ Tivoli Endpoint Manager for Patch Management
- ▶ Tivoli Endpoint Manager for Power Management

These packages are shown in Figure 9-2.

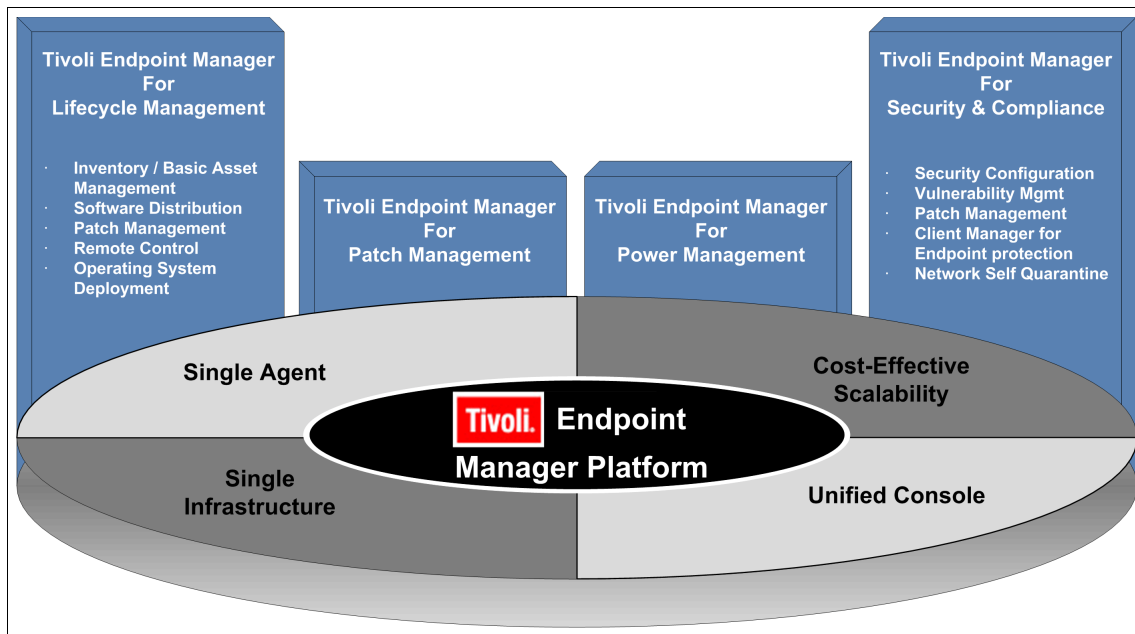


Figure 9-2 Tivoli Endpoint Manager offerings

More detailed descriptions of these packages are given in the following sections.

Tivoli Endpoint Manager for Lifecycle Management

Tivoli Endpoint Manager for Lifecycle Management addresses the management challenges of distributed environments with real-time visibility and advanced functionality that help users see and update systems, and remediate issues with continuous management, enabling IT to maintain service levels and focus on critical issues. This offering:

- ▶ Manages hundreds of thousands of endpoints regardless of location, connection type, or status.
- ▶ Manages heterogeneous platforms, such as Microsoft Windows, UNIX, Linux, and Mac operating systems, running on physical or virtual machines.
- ▶ Employs an agent-based approach that delivers up-to-date visibility and automatically remediates issues.
- ▶ Reduces management complexity and cost, increases accuracy, and boosts productivity.

- ▶ Simplifies and streamlines help desk calls and problem resolution with remote desktop control.
- ▶ Shortens update cycles, improves the success rates for provisioning, and reduces IT labor requirements.

Tivoli Endpoint Manager for Security and Compliance

IBM Tivoli Endpoint Manager for Security and Compliance addresses the security challenges of distributed environments with endpoint management and security in a single solution that supports security among distributed endpoints, reducing costs and complexity of management while increasing business agility, speed to remediation, and accuracy. This offering:

- ▶ Automates time-consuming device configuration and change management tasks.
- ▶ Effectively manages the compliance life cycle with an ongoing, closed-loop process.
- ▶ Gains greater visibility into network resources in dynamic and complex environments.
- ▶ Provides accurate, precise, and up-to-the minute visibility into and continuous enforcement of security configurations and patches.
- ▶ Centralizes the management of functions that provide advanced antivirus and firewall protection.
- ▶ Employs a unified management infrastructure to coordinate among IT, security, desktop, and server operations.
- ▶ Reaches endpoints regardless of location, connection type, or status with comprehensive management for all major operating systems, third-party applications, and policy-based patches.

Tivoli Endpoint Manager for Patch Management

Tivoli Endpoint Manager for Patch Management provides unified, real-time visibility and enforcement to deploy and manage patches to all endpoints from a single console that supports comprehensive patch management capabilities among distributed endpoints, reducing business risk, costs, complexity, and time while enhancing security. This offering:

- ▶ Automatically manages patches for multiple operating systems and applications across hundreds of thousands of endpoints regardless of location, connection type, or status.
- ▶ Reduces security risks by slashing remediation cycles from weeks to days or hours.

- ▶ Provides greater visibility into patch compliance with flexible, real-time monitoring and reporting.
- ▶ Provides up-to-date visibility and control from a single management console.
- ▶ Efficiently deploys patches, even over low-bandwidth or globally distributed networks.
- ▶ Automatically remediates problems related to previously applied patches.
- ▶ Patches endpoints on or off the network, including roaming devices using Internet connections.

Tivoli Endpoint Manager for Power Management

Tivoli Endpoint Manager for Power Management addresses the power management challenges of distributed environments with a policy-driven power management solution that supports comprehensive power control among distributed endpoints, reducing energy usage and costs while avoiding disruptions in systems management. This offering:

- ▶ Controls energy costs with a centralized, scalable, policy-driven power management system for all endpoints running Microsoft Windows and Mac operating systems.
- ▶ Manages power settings for hundreds of thousands of endpoints regardless of location, connection type, or status, from a single console.
- ▶ Empowers users with an opt-in approach that allows them to select the appropriate power profile.
- ▶ Applies easily integrated capabilities to deal with common power management issues, such as placing a always running PC into hibernation during off-hours to save energy.
- ▶ Creates “what if” energy usage scenarios to encourage conservation initiatives.
- ▶ Displays energy consumption as measures of power or carbon dioxide to help encourage organizational efforts to “go green”.
- ▶ Provides real-time visibility into current power usage and costs.

9.1.2 Key components

The Tivoli Endpoint Manager platform consists of the following three key components, which work in concert to enable real-time visibility from a single, central point of control:

- ▶ Single management agent
- ▶ Single management console and server
- ▶ Single policy-based model

Single management agent

The lightweight, intelligent agent can be deployed on every desktop, mobile computer, mobile device, and server that you want to be managed. It has the following characteristics:

- ▶ Has a multipurpose agent that offers the ability to consolidate and replace existing point-product software.
- ▶ Requires only 2 - 4 MB of endpoint system memory.
- ▶ Has real-time and continuous policy processing, remediation, validation, and reporting.
- ▶ Its policies remain enforced even when remote devices roam from the enterprise network.
- ▶ Has support for on-the-fly queries and management actions.
- ▶ Has policy-based and dynamic bandwidth throttling to work over Very Small Aperture Terminal (VSAT), Multi-Protocol Label Switching (MPLS), and other bandwidth-constrained networks.
- ▶ Has broad platform support, including virtualized operating systems, such as VMware ESX Server 3 and Microsoft's Hyper-V.

Single management console and server

The console and server work together to orchestrate a high level of visibility and control and have the following characteristics:

- ▶ The basic server model can manage up to 250,000 devices.
- ▶ Has built-in reporting and analysis tools.
- ▶ Has support for automatic multiserver synchronization and nonstop services, even during a disruptive event.
- ▶ The integrated security infrastructure controls agent actions and ensures administrator accountability.
- ▶ Has the ability to set configuration standards and baselines from defined groups of managed clients.
- ▶ Has standard SQL and SOAP interfaces for integration with other database applications and systems.

Single policy-based model

The policy language, referred to as the Fixlet Relevance language (see 9.1.7, “Relevance” on page 316), is a published command language that enables customers, business partners, and developers to create custom policies and services for managed assets.

As a single paradigm for interrogating and managing endpoints irrespective of platform or domain, the policy language can be used to solve common problems experienced by most organizations, such as deployment of patches, configuration management, antivirus management, or on-the-fly queries and remediation to manage the unforeseen and unstructured problems encountered by almost every enterprise. Without Tivoli Endpoint Manager, these problems either cannot be solved or must be solved manually, which takes a long time to complete and requires excessive amounts of resources. The main characteristics are as follows:

- ▶ Cloud-based service delivery of policy content for on demand functionality.
- ▶ New solutions are provisioned without additional hardware, infrastructure, or network changes.
- ▶ Open architecture for easy policy customization and development.

9.1.3 Deployment architecture

Figure 9-3 shows a typical Tivoli Endpoint Manager deployment architecture with a central server deployed at the Customer HQ and remotely deployed agents.




Figure 9-3 Tivoli Endpoint Manager: Typical deployment architecture

The four main components of the Tivoli Endpoint Manager platform are:

- ▶ Tivoli Endpoint Manager Client (the software agent required for each endpoint to be managed)
- ▶ Tivoli Endpoint Manager Server
- ▶ Tivoli Endpoint Manager Console
- ▶ Tivoli Endpoint Manager Relay

These four components are described in the following sections.

Tivoli Endpoint Manager Client

The Tivoli Endpoint Manager Client has to be deployed on every desktop, mobile computer, mobile device, and server that you want to manage using the Tivoli Endpoint Manager platform.

The client accesses a collection of Fixlet messages (see “Fixlet messages” on page 312) that seek out security holes, vulnerabilities, and deviations from the desired operating environment. If a vulnerability is found, the client can then implement corrective actions received from the console. In most cases, the client operates silently, without any direct intervention from the user. However, should you need to solicit user response, the solution also allows you to provide prompts.

The client is capable of assessing the state of the endpoint against policy and bringing the endpoint back into compliance with policy, without any instruction from the management server. This is true of the whole platform’s security and systems management applications, from security configuration management to software distribution to power management; a single agent is all that is necessary.

Clients can be deployed using the Client Deployment Wizard, as shown in Figure 9-4.

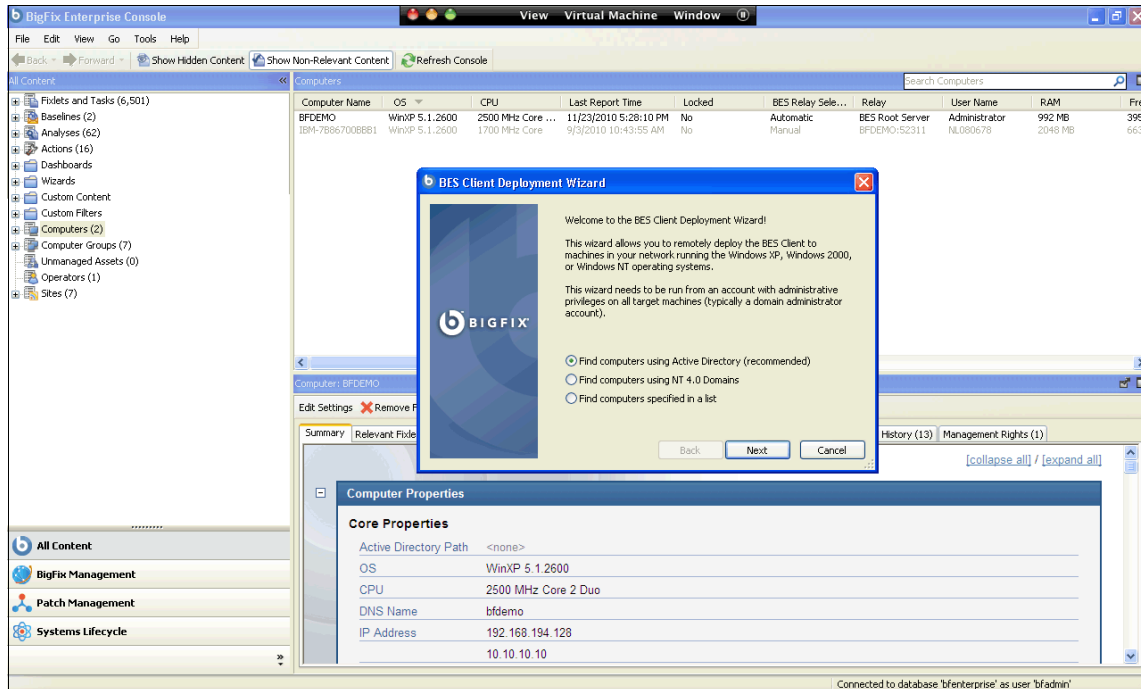


Figure 9-4 TEM Client Deployment Wizard

Broad support: The Tivoli Endpoint Manager Client can be deployed on a wide range of endpoints, including Mac, Windows, Windows Mobile, VMWare ESX Server, Linux, and UNIX operating systems. For a more detailed list of supported operating systems, go to:

<http://support.bigfix.com/bes/misc/supportpolicy.html>

Tivoli Endpoint Manager Server

In most deployments, there is usually only one server (aside from high-availability scenarios), and this single server can typically manage up to 250,000 endpoints.

Tivoli Endpoint Manager Console

The console ties all these components together to provide a system-wide view of your networked computers, along with their vulnerabilities and suggested remedies. As an authorized user, the Tivoli Endpoint Manager Console allows you to quickly and simply distribute a fix to exactly those computers that need it, with zero impact on the rest of the network.

The Tivoli Endpoint Manager Console is a thick-client Windows application that can be run on any Windows computer that has network access to the Tivoli Endpoint Manager Server. Because the Tivoli Endpoint Manager Console is the management center of operations, we provide more detailed information about in 9.1.4, “Console” on page 311.

Tivoli Endpoint Manager Relay

Relays are optional network components that can significantly improve the performance of your platform deployment. A relay simultaneously mitigates two bottlenecks:

- ▶ It can relieve the load on the server(s). The server has many duties, among them the taxing job of distributing patches and other files. A relay can be set up to ease this burden so that the server does not need to distribute the same files to every client. Instead, the file is sent once to the relay, which in turn distributes it to the other clients. The impact on the server is reduced, on average, by the ratio of relays to clients.
- ▶ It can reduce congestion on low-bandwidth connections. If you have a server communicating with a dozen computers in a remote office over a slow VPN, designate one of those computers as a relay. Then, instead of sending patches over the VPN to every client independently, the server only sends a single copy to the relay. That relay, in turn, distributes the file to the other computers in the remote office over the LAN. This effectively removes the VPN bottleneck for remote groups on your network.

A key benefit of deploying relays is that they can be deployed on shared hardware, such as file, print, or domain servers, or other computers, such as kiosks that are operational all the time. This way, organizations can scale with minimal hardware requirements. Any Windows, Solaris 10, or RHEL 4 or 5 computer with an agent installed can be dynamically designated to be a relay.²

² For more information about operating systems support for relays, go to <http://support.bigfix.com/bes/misc/supportpolicy.html>.

9.1.4 Console

The Tivoli Endpoint Manager Console allows you to connect to the server and manage the endpoints on which you have deployed the client.

The console consists of four main parts, as shown in Figure 9-5:

- Content: Specific to the *domain* selected
- Computers: Endpoints on which the client is already deployed
- Computer specifics: Details about a specific endpoint
- Domain: For example, Endpoint Protection or Patch Management

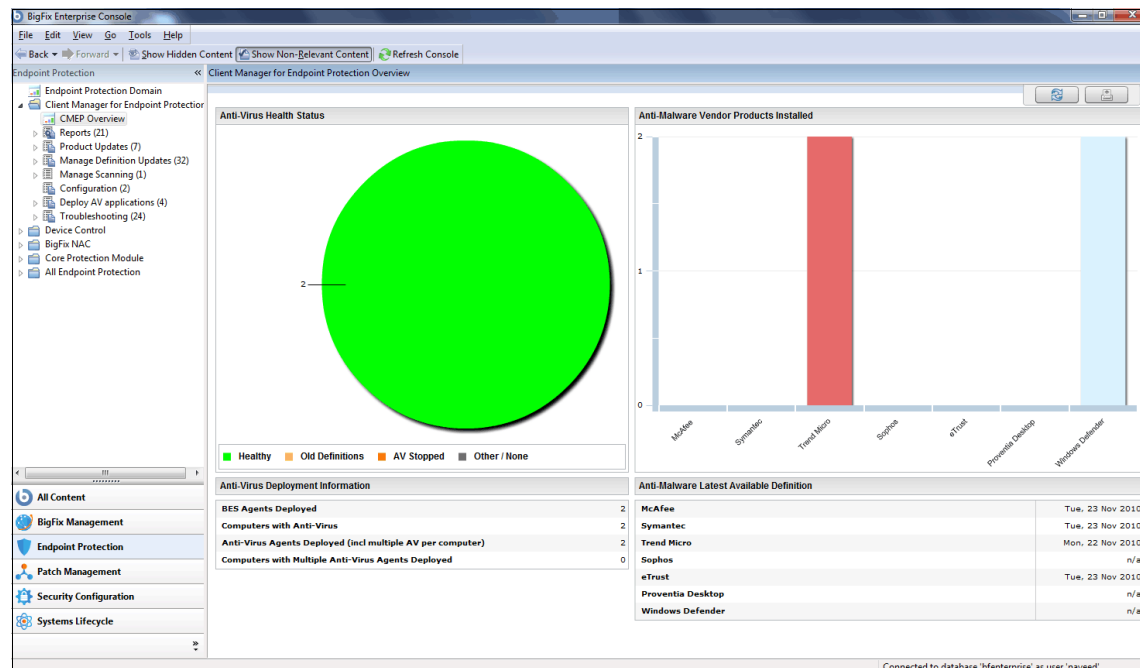


Figure 9-5 Tivoli Enterprise Manager: Console showing the Endpoint Protection domain

Domain portlet

At the bottom left of the console window, you see the available *domains*. The organization by domain makes it easier for functional administrators (for example, antivirus or patch management) to see only those content items that are relevant to them and remove those not relevant from the content portlet (top left). Clicking the **Endpoint Protection** domain displays only Endpoint Protection related dashboards and wizards.

Content portlet

The first four menu items in the content portlet are as follows:

- ▶ Fixlets and tasks
- ▶ Baselines
- ▶ Analyses
- ▶ Actions

Fixlets and tasks

Because fixlets and tasks perform similar functions, they are grouped together into a single menu item in the console.

- ▶ Fixlets

Definition: A fixlet is a piece of code within the Tivoli Endpoint Manager solution that first identifies a “problem situation”, such as a missing operating system patch.

The fixlet instructs the agent on the endpoint to query the server for this missing patch, and download it to the endpoint.

Fixlets are grouped into “sites” that denote a collection of fixlets that apply to a certain category of issues to be managed, such as patch management.

At this point, we begin to see messages in the console alerting the operator to the current process, and showing progress towards completion.

These are called *fixlet messages*.

- ▶ Fixlet messages

Fixlet messages lay at the core of the Tivoli Endpoint Manager functionality. Using *Relevance statements*, they can target specific computers, remediating just those client computers with an issue and never bothering the computers that do not have an issue. Fixlets come with an *action script* that can resolve the issue with a single mouse click. Typically, when the action has completed, the fixlet detects that the issue is no longer applicable to that computer. As fixlet actions propagate through your network, you can track their progress with the console, Web Reports, and the Visualization Tool. After every computer in your network is remediated, the fixlet message disappears from the list. Should the issue reappear, the fixlet once again shows up in the list, ready to address the issue once again.

Fixlet messages contain a text description of the issue and may offer several different actions, including links to more information. Often a fixlet message will have a default action, allowing you to simply click from the fixlet list to deploy it.

Fixlet messages can be grouped into *baselines*, allowing even higher levels of automation. If you create a baseline of fixlet messages that contain default actions, you can turn the tedious chore of maintaining a common operating environment into a single-click operation.

At any time, you can open a fixlet message to inspect the Relevance Expressions that target the clients and the action script that will remediate the issue. This gives you a high degree of confidence in the applicability and efficacy of the remedial action. You can also see exactly which computers on your network are affected by each fixlet message and view a history of the actions taken on a client-by-client basis.

► Tasks

Tasks are similar to fixlet messages, but are designed for ongoing tasks, and as a consequence, they are more persistent. Tasks come with one or more action scripts that can help you to adjust settings or run maintenance tasks.

Tasks contain a description of the issue and may have a default action, allowing you to simply click from the Task list to deploy it. Tasks and fixlet messages can be grouped into baselines, allowing even higher levels of automation.

At any time, you can open a task to inspect the Relevance Expressions that qualify the clients and the action script that will address the task.

Baselines

Baselines are collections of fixlet messages and tasks. They provide a powerful way to deploy a group of actions across an entire network with a single command.

Baselines provide a way to maintain a common operating environment, making sure that all users in any given domain have the same software, patches, and drivers. Baselines are easy to set up, simply by selecting the fixlet messages, tasks and other baselines that you want to be a part of the group. To limit the scope of a baseline, a Relevance Expression can be used to target any subset of your network, using IP addresses, computer names, operating systems, and many other qualifiers.

For example, you might make a baseline named “All critical hotfixes”, and populate it with all the current critical hotfixes available in the fixlet list. Or you might create one named “Finance department baseline”, designed to keep that particular group of computers updated with the latest financial programs, financial tables, updates, and patches.

Analyses

Analyses allow an operator to view and summarize various properties of client computers across a network. There are several pre-made analyses supplied by IBM that examine various aspects of your networked computers, including their hardware, applications, and server, relay, and client relationships.

Studying these pre-made analyses can be instructive when you want to make your own or customize existing ones. Custom analyses can help you monitor aspects of your network that are interesting or vital to your organization's operation.

The Retrieved Properties that underlie each analyses are created with Relevance Expressions. For example, to make sure you have fully deployed the most recent client software, you might use an expression such as `version of client`. This simple expression is evaluated on every computer where the analysis is targeted, allowing you to see explicitly which version of the client is running on each computer, or to view a summary of how many machines are running each version.

Analyses are targeted with another Relevance statement, which may be as simple as `TRUE`, to include all connected clients. Generally, you want to narrow the scope with a Relevance statement, such as the name of the operating system as lowercase starts with "win", which would limit the analysis to Windows computers only.

Actions

Actions represent the core functionality of the system. Fixlet messages, tasks, and baselines depend on actions to execute their remediation mission.

Actions are typically scripts that can customize a specific solution for each client, using the power of Relevance Expressions. Although the Relevance language itself cannot alter a client, it can be used to direct *actions* in a way that parallels the original trigger. For example, a fixlet might use the Relevance language to inspect a file in the system folder. Using a similar Relevance clause, the action can then target that same file without knowing explicitly where that folder resides. This allows the action author (and issuer) to concentrate on the issue at hand without worrying about the vagaries of each individual computer system.

You can inspect an action script before you execute it by looking at the Details tab of Fixlet messages and Tasks. You can also write your own custom action scripts.

9.1.5 Computer groups

The Tivoli Endpoint Manager Console allows you to group your computers so that you can target them appropriately. You might want to group your development computers, for example, to make sure you do not interfere with certain earlier software projects. There are several ways to group computers, but the two most common techniques are *manual* and *automatic grouping*.

Manual groups are static, but automatic groups can change dynamically, depending on the current values of the inclusion properties. See Figure 9-6 for an example of how to manually add a computer to a group.

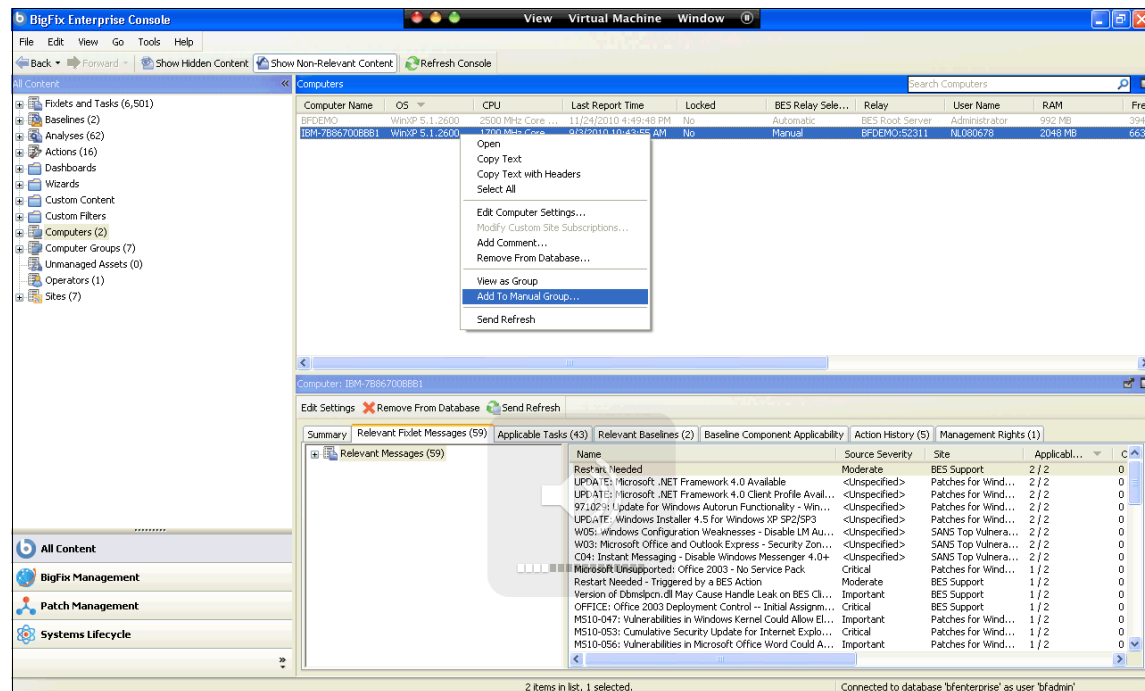


Figure 9-6 Tivoli Endpoint Manager: Adding a computer manually to a group

9.1.6 Fixlet sites

Upon installation, the platform automatically subscribes itself to the fixlet sites that you select during installation. Each fixlet site contains a collection of fixlet messages that perform certain tasks.

See Figure 9-7 for examples of fixlet sites.

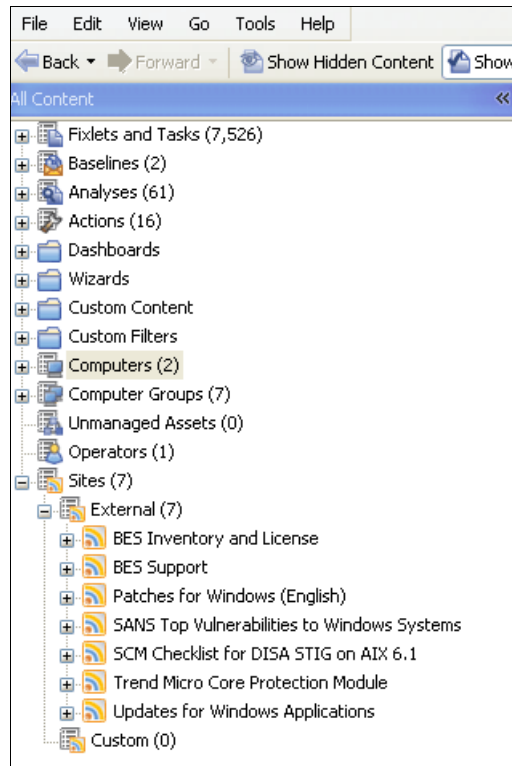


Figure 9-7 Tivoli Endpoint Manager console showing fixlet sites

9.1.7 Relevance

To quickly inspect various aspects of a computer, the Relevance language described in “Single policy-based model” on page 306 was developed. This human-readable language is at the heart of Tivoli Endpoint Manager and allows fixlet authors to target actions at just those computers that need the fix, and no others. Thus, you can be confident that only broken machines are being fixed and that the rest are never bothered.

The Relevance language can query an exhaustive set of computer properties, and do it quickly. Most console operators rely on third parties to write fixlet messages, and so their exposure to the Relevance language is not critical to operating the console. However, you can customize the console with short lines of code from the Relevance language (called Relevance Expressions) that grant you an unprecedented amount of control over the client computers in the network.

A typical Relevance Expression might be:

```
vendor name of processor
```

This expression returns the name of the manufacturer of the CPU (Intel® or AMD, for example), which can then be used to determine relevance.

You can use Relevance Expressions to create retrieved properties, which you can then use to organize and filter the clients in the network. You can experiment and debug your Relevance Expressions using the Relevance Debugger, which is automatically installed with the console. The program can also format your expression for easier reading. There are literally thousands of useful Retrieved Properties, and there are far too many to list here.

9.1.8 Web Reports

The Web Reports program can monitor, print, or archive the status of the local database. It also has the ability to read the databases of other servers and aggregate the data, which offers you a top-level view of the overall organization with multiple database servers. Aggregation servers allow you to view information from multiple networks that have hundreds of thousands of computers.

The Web Reports program enables you to get an overview of your relevant fixlet messages and your remediation efforts. You find charts summarizing the number of administered computers in your network and your vulnerability status. In addition, you find overall statistics and a list of the most common issues detected. You can click these commonly relevant fixlet messages to see them in greater detail. In addition, you can see at a glance how remediation and policy enforcement efforts are progressing.

9.1.9 Visualization Tool

The Visualization Tool allows you to view and manipulate data from your entire managed network. It lets you visualize various hierarchical relationships in your network, using a three-dimensional sphere to map the data.

The Visualization Tool makes it possible to view a real-time graphical network map showing fixlet relevance, action status, and Retrieved Properties.

As an example, you could view all computers that are currently unpatched for the MS04-011 vulnerability across an enterprise network, displayed as an Active Directory hierarchy. Then you could watch the clients change from red to green as they get patched in real time across your network.

As another example, you can view all clients that are currently using Microsoft Office, colored according to version. Or you can create your own hierarchy. You can assign settings on all your machines named “city”, “building” and “floor”. You could then create a dynamic setting called “location” that concatenates these properties.

9.2 Proventia Desktop Endpoint Security

IBM has been in the endpoint security market since 2007, when they acquired Internet Security Systems. Internet Security Systems themselves acquired Network ICE in 2001.

The endpoint security agent protects Windows endpoints and contains two layers of security: a personal firewall and host intrusion prevention system.

The Proventia Desktop Endpoint Security agent is part of the IBM Security offerings and offers preemptive security to Windows endpoints using the Protocol Analysis Module described in 6.5, “Protocol Analysis Module” on page 165.

9.2.1 Attack vectors covered

The following attack vectors are covered by the Proventia Desktop Endpoint Security solution.

Network vector

Network-based attacks use malicious network traffic to remotely compromise their target systems. Unlike other threats, network based attacks can penetrate, launch, and propagate *without human intervention*. Network-based attacks on the host predominantly exploit vulnerabilities in protocols and network-aware processes. These vulnerabilities are typically the result of programming errors that provide opportunities for a buffer overflow. Exploit types include, but are not limited to, direct hacking and theft, network-based worms, denial-of-service, attacks, and the installation of remote access backdoors, and robot (bot) footholds for future use by a hacker.

The Sasser³ worm, being a good example of a network vector attack, infects hosts from the network by exploiting vulnerabilities in the Microsoft Local Security Authority Subsystem Service (LSASS). After penetration of the host through exploitation of the LSASS vulnerability, the worm’s payload is transferred to the compromised system and launched remotely, without user involvement.

³ For more information about the Sasser worm, go to <http://www.viruslist.com/en/viruses/encyclopedia?virusid=50204>.

After a system is infected, the malicious executable associated with the Sasser worm intends to propagate by scanning other network hosts for the vulnerable LSASS service.

Three technologies work to defend host systems against network-based attacks, including *personal firewalls*, *intrusion detection and prevention systems*, and *Buffer Overflow Exploit Prevention*. A subset of network-based attacks can utilize file executables to further propagate from the host. In such cases, application-based prevention technologies might provide detection post-launch and prevent attack propagation.

Personal firewalls

The personal firewall is the method used by the first generation of firewalls. It was the first security type created and used, and is sometimes known as *distributed firewall technology* or *managed personal firewall technology*.

Some common features of a personal firewall include:

- ▶ Monitors all incoming and outgoing packets.
- ▶ Controls the acceptance or denial of applications that attempt to access the Internet.
- ▶ Logs and gives alerts about all connection attempts.

Intrusion detection and prevention system

As IP traffic enters your system, the intrusion prevention system (IPS) analyzes it for malicious data. The IPS drops offending packets, and allows the original traffic to continue travelling up the network stack unhindered. If the IPS component detects an event serious enough to pose an immediate threat to the system, the IPS component blocks, drops, or resets the connection. It drops the current packet and instructs the firewall to block all traffic to or from the intruder, depending on user input. This action stops any further communication with the intruder.

Application vector

Application-based attacks use program files to attack and compromise target systems. Unlike network-based attacks, program files typically require some form of user involvement to launch an attack. Email, web downloads, removable media, and peer-to-peer applications are all common sources of application-based attacks. Attack types include, but are not limited to, viruses, email worms, trojan horse applications, and spyware. More information about typical application-based attacks can be found in 5.2, “Malware” on page 129.

A good example of an application vector attack is the MyDoom⁴ worm. Variants of this worm are sent to users through email, and contain a semi-randomly named executable attachment. Upon user execution of the attachment, the worm launches, installs itself onto the system, and begins propagating through email to other email addresses found on the compromised host. Some variants of the MyDoom worm can also launch DDoS attacks against specific targets.

Antivirus

Antivirus software is the computer's first line of defense against an application-based attack. Signature-based antivirus and anti-spyware uses a library of known virus signatures to check files, email messages, and attachments for viruses. Trend Micro's Office Scan engine is a signature-based antivirus product and is effective in detecting and preventing known viruses, worms, spyware, and some trojans.

With Tivoli Endpoint Manager you can control which antivirus software version is required, and set actions to be taken if the computer is not running the required software. Using this approach, the agent scans for patterns that match signatures of known worms and viruses, and handles the viruses according to the options you select. You must enable and configure the signature antivirus and antispysware feature before you can use it to protect your organization's endpoints. The agent can enforce antivirus compliance to give you added protection. With antivirus compliance enabled, the agent can enforce the presence of antivirus software on a user's computer.

⁴ For more information about the MyDoom worm, go to <http://www.viruslist.com/viruses/encyclopedia?virusid=57410>.

To see an example of the type of report available from Tivoli Endpoint Manager that shows antivirus versions deployed on the managed endpoints, see Figure 9-8.

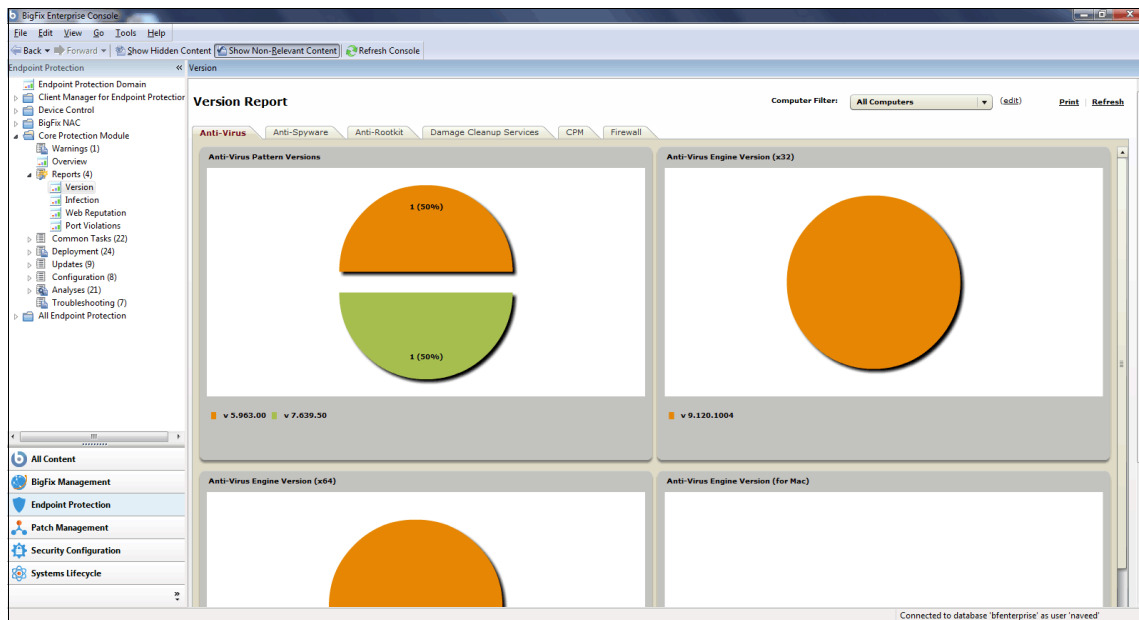


Figure 9-8 Tivoli Endpoint Manager: Antivirus version report

This concludes the overview of the IBM Security endpoint protection capabilities. We now continue with the IBM Security Server Protection agents.

9.3 IBM Security Server Protection

IBM Security Server Protection agents combine powerful protection technologies into a single multi-layered agent to guard business critical systems and data from attack.

To understand how the security capabilities of the IBM Security Server Protection can be mapped to the IBM Security Blueprint⁵, see Figure 9-9 on page 323. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using IBM Security Server Protection. This functional highlighting is applicable for the infrastructure service components as well.

Besides the fully highlighted elements, Figure 9-9 on page 323 also shows some medium highlighted elements. Although the IBM Security Server Protection can be used to address such a component to some degree, the respective area of coverage is not considered a core function of the product and thus is considered to be limited.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 9-9 on page 323 can be used as a quick reference of the functional security management aspects of the IBM Security Server Protection. This reference can help determine which functions of a solution can be covered by selecting this product.

⁵ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

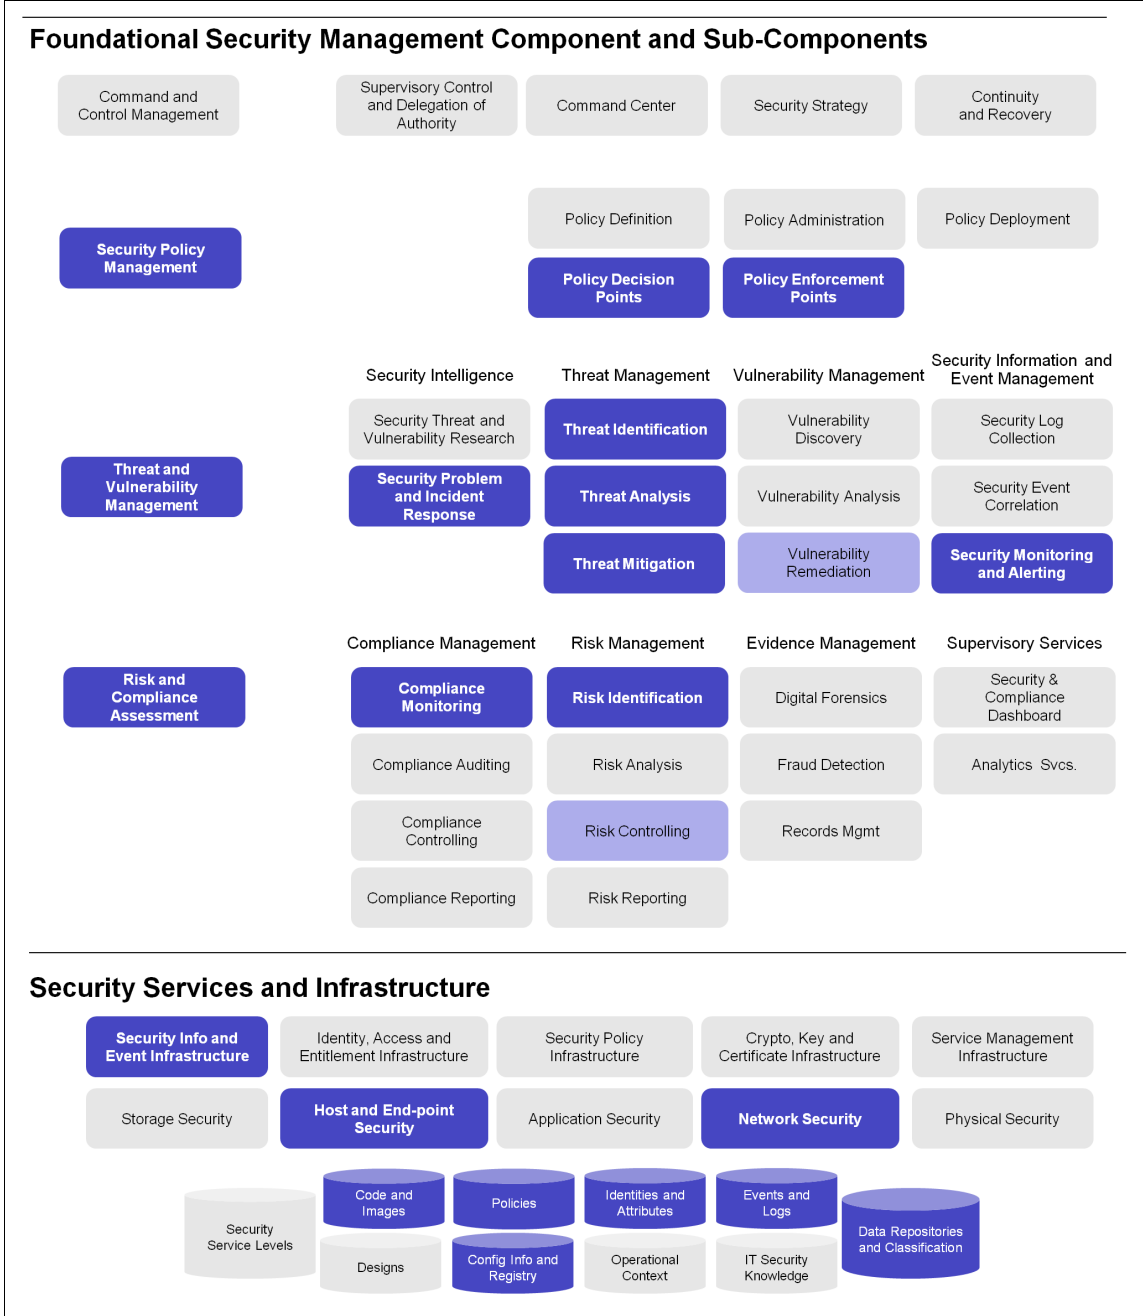


Figure 9-9 Mapping of the IBM Security Server Protection suite to the IBM Security Blueprint

9.3.1 Architecture overview

Figure 9-10 shows the progression of data as it passes through the multiple protection layers of the IBM Security Server Protection agents. The features described in this section are not universally available in each operating system.

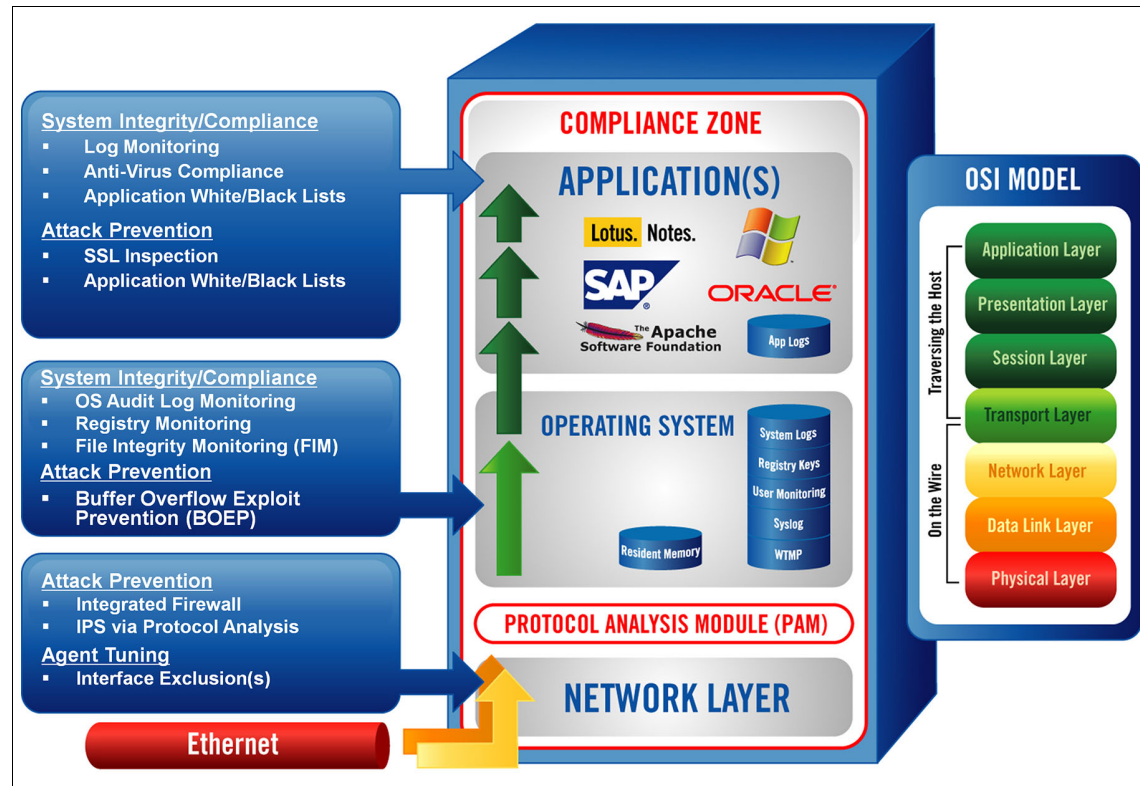


Figure 9-10 Multilayer server protection

Passing through the physical, data link, network, and transport layers of the OSI model, packets traverse the network interface card (NIC) and onto the TCP/IP stack, where data is queued for delivery to the execution space. During this initial phase, a mechanism called the *Packet Capture Driver* (PCD) is used to capture packets for inspection. The PCD passes the data to the *Protocol Analysis Module* (PAM) for inspection before allowing the delivery of potentially malicious information to vulnerable areas of the operating system and applications.

The following actions occur chronologically after being queued by the PCD:

- ▶ Interface exclusions (performance enhancer)
- ▶ Inspection bypass (performance enhancer)
- ▶ Firewall (intrusion prevention)
- ▶ Intrusion prevention (intrusion prevention)

Although the design strategy to *hook* or *shim* the kernel differs in each OS, it is clear that capturing traffic at the network level is required to prevent potentially malicious data from reaching vulnerable areas of the machine. This approach is unequivocally the most secure method of halting and inspecting traffic. When assessing a host intrusion and prevention system product, it is critical to understand the design of each competitor and the security benefits or trade-offs for where they capture and inspect traffic.

To be clear, interface exclusions, inspection bypass, and firewall rules can be assessed at low layers of the OSI model, but the intelligence required to perform deep packet inspection lies at the application layer (layer 7 or the OSI model). Queued packets are sent to the PAM in the application layer, where a signal is sent back to inform the PCD whether to release or block the data in question.

One caveat to mention is SSL encrypted traffic. Although encrypted traffic cannot be inspected through the initial PAM screening, our host agents can integrate with Microsoft IIS and Apache web servers, allowing the decryption process to occur before capturing this data and sending it for inspection by a light-weight PAM parser.

Following inspection, data is forwarded to the execution spaces of the operating system and applications. At this juncture, two technologies intercept the execution of this data:

- ▶ Buffer Overflow Exploit Prevention (BOEP)
- ▶ Application control (white and black lists)

Buffer Overflow Exploit Prevention

BOEP identifies attempts to execute code (system calls) on writable memory regions. This is an important distinction to make. BOEP does not prevent all buffer overflows, but only those that overrun their bounds and attempt to execute in writable regions of memory. More information about BOEP can be found in 6.8.4, “Buffer Overflow Exploit Prevention” on page 192.

Application control

Application control, a technology that enables the specification of trusted versus untrusted applications, is also available. If used correctly, application white and black lists are an effective method of policy enforcement (through either the file name, MD5 checksum, or both).

The remaining features of the IBM Security Server Protection products are classified as compliance or intrusion detection features. File, registry, and system integrity monitoring (as well as third-party log monitoring) provide the ability to correlate network activity with user and file behavior at the operating and file system levels. The ability to correlate certain events make IBM Security Server Protection and the RealSecure Server Sensor a powerful security and compliance tool. Not only can we provide four of the five Ws of a potential attack (who, what, when, and where), we can also observe and correlate these events to distinguish internal versus external attacks, as well as malicious versus unintentional actions. The IBM Security Server Protection and the RealSecure Server Sensor can provide certain historical and forensic data required to pass even the most restrictive compliance regulations.

9.3.2 Windows and Linux server protection

IBM Security Server Protection agents offer multi-layered protection for Windows and Linux servers, designed to keep your data and applications reliable, available, and confidential.

IBM Security Server Protection can assist organizations who need to comply to the Payment Card Industry Data Security Standard (PCI DSS) by providing functionality such as *Registry Integrity Monitoring* (RIM) and *File Integrity Monitoring* (FIM).

In addition, the Application Control feature can lock down user access to programs that might be used to take control of a server and extract customer data or confidential intellectual property. Also, IBM Security Server Protection uses the IBM Security SiteProtector as a single management console for multilayered protection. You can obtain more information about SiteProtector in Chapter 7, “Centralized management” on page 199.

IBM Security Server Protection can integrate seamlessly with existing IT infrastructures to preserve legitimate traffic flows without interruption and the unification of compliance and Defense-in-Depth lowers your total cost of ownership (TCO) while strengthening your return on investment (ROI).

Components

Multiple technologies deliver server protection that prevents attackers from exploiting vulnerabilities in operating systems and applications. Some of them use the same technologies performed by Proventia Desktop, but other features are incorporated exclusively to IBM Security Server Protection. Let us have a closer look at those components.

Console

The local console provides the user interface for IBM Security Server Protection. It reports information about the intrusions IBM Security Server Protection has detected, the intruders performing attacks, and the responses of IBM Security Server Protection. It also displays recent network traffic and attacks graphically so that you can observe patterns or spikes in network activity. The local console also provides tools to manually create firewall entries and control how tightly IBM Security Server Protection guards your computer.

Intrusion prevention

As IP traffic enters your system, the intrusion prevention system (IPS) in IBM Security Server Protection analyzes it for malicious data. The IPS drops offending packets, and allows the original traffic to continue travelling up the network stack unhindered. If the IPS component detects an event serious enough to pose an immediate threat to the system, the IPS component blocks, drops, or resets the connection. It drops the current packet and instructs the firewall to block all traffic to or from the intruder, depending on user input. This action stops any further communication with the intruder.

The core of IBM Security Server Protection is its seven-layer decoding engine, which analyzes incoming and outgoing network traffic in real time. If IBM Security Server Protection detects an intrusion, it displays an event on the local console, reports the event to SiteProtector (if configured), and, if necessary, commands the firewall to automatically block the intruder's communications.

Most modern intrusion prevention systems compare the information in each incoming transmission to a huge database of intrusion patterns. This *pattern matching* approach is slow and prone to errors. The IBM Security Server Protection engine instead analyzes the protocols that carry the information. This enables IBM Security Server Protection to work even on fully loaded, high-speed networks without slowing traffic. This engine can run *invisibly* on your system to ensure that an intruder does not accidentally or purposely disable IBM Security Server Protection.

Firewall capabilities

IBM Security Server Protection provides powerful firewall capabilities that inspect all inbound and outbound traffic on your computer for unauthorized activity. IBM Security Server Protection can control outbound communication based on port, IP address, and protocol. IBM Security Server Protection blocks unauthorized activity without affecting normal traffic.

The IBM Security Server Protection firewall collects valuable information about intrusion activities. This evidence can be used to analyze an intruder's attacks or provide proof that certain events took place. When IBM Security Server Protection detects an intrusion, it immediately traces the origin of the attack, which helps IBM Security Server Protection determine the intruder's IP address, computer name, and hardware address. You can use the local console or SiteProtector to specify how extensively IBM Security Server Protection traces each intrusion.

Application control

IBM Security Server Protection can control whether specified applications can run on your computer and whether applications on your computer can connect to a network. You can control communications for unknown and known applications.

Buffer Overflow Exploit Prevention

IBM Security Server Protection can prevent exploits based on buffer overflows. Some intruders attempt to send more data to the buffer than it can handle. This can enable intruders to effectively take control of the computer.

Virtual Patch protection

Virtual Patch protection allows you to preemptively protect your servers without the pain and expense associated with emergency patch rollouts.

Web application protection

Web application protection protects web applications by inspecting traffic for malicious activity and is capable of adding an additional level of protection to those applications running on Apache 2.0.x by inspecting Secure Sockets Layer (SSL) encrypted traffic.

Antivirus enforcement

IBM Security Server Protection defends all the computer systems on your corporate network from intrusions. IBM Security Server Protection agents are designed to run on all the workstations and servers on your enterprise network, providing local intrusion prevention and protection for each and every machine.

Antivirus compliance lets you specify antivirus software requirements that must be met before a computer accesses your network, and what actions the IBM Security Server Protection Agent Manager and Proventia Desktop Agent take if these requirements are not met. Antivirus compliance operates the same way regardless of which VPN software is installed on the computer, and does not affect the computer's ability to access the Internet. The actions the agent takes if these requirements are not met include:

- ▶ Displaying warning messages
- ▶ Blocking the user from most or all of the corporate network
- ▶ Blocking the user from using the VPN, in effect blocking network access

When you implement antivirus compliance, you include the antivirus compliance configuration as part of a policy, and then apply the policy to a group of agents. You must ensure that each group using the policy has the required antivirus software installed. If an agent in the group does not have the required antivirus software installed, the agent might block access to the network. Use the Agent Manager to instruct the agent to check for the following antivirus software requirements:

- ▶ Whether the required antivirus software is installed on the computer.
- ▶ Whether the antivirus software is running on the computer.
- ▶ Whether the antivirus software has up-to-date virus definition files.
- ▶ How often the agent checks for antivirus compliance.
- ▶ What message the user receives when antivirus software requirements are not met.

Whether or not the agent blocks the computer's VPN ability to access the network if antivirus software requirements are not met each time a computer in the group starts, the IBM Security Server Protection Agent Manager starts the Proventia Desktop Agent to run a series of checks to ensure that the computer meets the antivirus software requirements you specify. If the computer does not meet these requirements, the agent notifies the IBM Security Server Protection Agent Manager and displays the text message you specify. You can also instruct the agent to block the computer's access to the IP addresses or address ranges on the corporate and VPN IP address lists.

Spyware and riskware protection

Spyware and riskware, software that gathers user information without the knowledge or informed consent of the user, is most often installed on host computer by another application from the Internet. IBM Security Server Protection provides protection against the most commonly known spyware and riskware applications.

Policies

A policy consists of *policy elements*. A policy contains all the configuration, installation, and security settings applied to agents within a group. Policies are required for all IBM Security Server Protection deployments.

Have a carefully planned corporate security policy in place before you attempt to configure IBM Security Server Protection policies. A corporate security policy allows access to required resources, while also providing protection against the risks that open resources can represent. IBM Security Server Protection policies define the behavior of the agent. You must correctly configure and apply policies to obtain the protection you desire.

IBM Security Server Protection policies are configurable through policy elements and an actual policy; these elements have related components as well. The following list shows the definition of the policy elements:

Administration	This element configures administrative settings, such as heartbeat interval, agent protection properties, and failover configuration.
Agent build settings	This element sets the agent version. IBM Security Server Protection for Windows offers the ability to create silent installation packages (Agent Builds) with predefined behavior for easy deployment.
Group settings	This element defines communication properties, such as the Agent Manager with which the IBM Security Server Protection agent communicates.
Install and update settings	This element is used to update agents as new versions or security content becomes available.

The IBM Security Server Protection Administrators Guide⁶ provides good descriptions for the other policy elements:

- ▶ Application compliance
- ▶ Buffer Overflow Exploit Prevention (BOEP)
- ▶ Bypass filters
- ▶ File Integrity Monitoring (FIM) and Registry Integrity Monitoring (RIM)
- ▶ Firewall
- ▶ Security events
- ▶ System integrity monitoring

⁶ You need to download the appropriate document for your operating system platform from <http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp>. At the time of the writing of this book, you still need to look up the product name under Proventia Server IPS.

Operating system support

IBM Security Server Protection supports Microsoft Windows Server and several Linux server operating systems, as well as operating systems running within certain virtual environments, such as Hyper-V and VMware. For an overview of the exact operating system requirements for the current versions of the IBM Security Server Protection products, refer to the IBM Security product Information Center at:

<http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp>

With IBM Security Server Protection for Windows V2.1, IBM introduced support for Windows Server 2008 32-bit and 64-bit operating systems.

This concludes our overview of the IBM Security Server Protection for Windows and Linux. In the remainder of this chapter, we focus on the UNIX protection.

9.3.3 UNIX server protection

The IBM Security Server Protection Agent for UNIX protects business critical servers from both internal and external threats. Its real-time intrusion detection and prevention can reduce network security costs while protecting enterprise server environments reducing downtime. The IBM Security Server Protection agent analyzes events, host logs, and inbound and outbound network activity on critical enterprise servers to block malicious activity from damaging critical assets. The solution applies built-in signatures and sophisticated protocol analysis to prevent known and unknown attacks.

Legacy product connection: At the time of the writing of this book, the UNIX server protection agents are still named *IBM Real Secure® Server Sensor*. To navigate to the appropriate product documentation, visit the IBM Security product Information Center at

<http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp> and then navigate to **Legacy products** → **Real Secure Server Sensor**.

The IBM Security Server Protection Agent preemptively combats threats and addresses vulnerabilities at the network and application levels while performing security compliance auditing. Here is an overview of its benefits:

- **Server protection**

Designed to protect the underlying operating system by helping to prevent attackers from exploiting operating system and application vulnerabilities.

- ▶ Web application protection

Provides Secure Sockets Layer (SSL) encrypted application layer intrusion monitoring, analysis, and response capability for both Apache and IIS web servers.
- ▶ Advanced intrusion prevention/blocking

Monitors all traffic to and from the server or network to detect and prevent inbound attacks as well as block new and unknown outbound attacks, such as buffer overflows, trojans, brute force attacks, unauthorized access, and network worms.
- ▶ Console and network-based intrusion protection

Provides the flexibility to detect and prevent both console and network-based attacks through log monitoring capabilities that detect malicious activity before it causes any damage.
- ▶ Broad platform coverage

Provides you with the flexibility to grow your server protection strategy regardless of the environment (Windows, Solaris, HP-UX, AIX® and Linux).
- ▶ Audit policy management

The centralized management of the operating system audit policy helps ensure that all critical servers have consistent and effective audit policy and allows for the management of true kernel-level auditing.
- ▶ Global technical support

Provides customers with a wide array of support offerings, specifically designed to meet the cost and service demands of diverse networking environments.

The IBM Security Server Protection Agent for UNIX provides automated, real-time intrusion protection and detection by analyzing events, host logs, and inbound and outbound network activity on critical enterprise servers to block malicious activity from damaging critical assets.

The IBM Security Server Protection Agent for UNIX combines built-in signatures and sophisticated protocol analysis with behavioral pattern sets and automated event correlation to help prevent known and unknown attacks.

Analyzing and blocking network-based threats

Network-based intrusion detection is good at providing an early warning of attacks. By monitoring the traffic stream in real time, a network sensor can see a threat and often neutralize it before it has a chance to do any damage.

The IBM Security Server Protection Agent for UNIX protects against network vector attacks, including worms, bot worms, trojans, and denial-of-service (DoS) attacks through a local firewall and inline vulnerability-centric intrusion prevention. IBM Security Server Protection provides the following items:

- ▶ Firewall
- ▶ Advanced Intrusion Prevention/Blocking
- ▶ Protocol Analysis Module (PAM)
- ▶ Centralized Management

Analyzing and blocking host-based threats

Because network IPS products cannot tell you whether an attack was successful or not, host-based intrusion detection systems complement their network counterparts nicely. Host-based sensors can provide confirmation of an attack's success or failure and can yield system-specific event data, such as the user name and file name during an unauthorized access attempt. Host-based sensors are important for another reason: Local users can attack a system without being detected by the network sensor. Let us look at an example.

Somebody with access to a server console can try passwords all day without a network sensor detecting it. A valid user running the hacker utilities *getadmin* or *sechole* to add himself to the administrators group, or someone trying to open a file without permission or insert a trojan into a system file, can do so without being detected by the network sensor. The IBM Security Server Protection Agent for UNIX can detect all of these examples. The agent's layered intrusion prevention inspects and blocks application traffic with malicious code activity, including applications running on both Apache and Internet Information Services (IIS) web servers. The IBM Security Server Protection Agent for UNIX provides:

- ▶ Preemptive web application protection with SSL inspection
- ▶ HTTP/application protection
- ▶ Buffer Overflow Exploit Prevention

Deployment

The IBM Security Server Protection Agent for UNIX works in network environments where it is either impractical or too costly to deploy network IPS products. As networks become faster, it becomes more difficult and costly to monitor all inbound and outbound traffic.

By deploying both network sensors and host-based sensors, you can have the best of both worlds: ultra-fast detection and response at the network level with rich, system-specific confirmation of events at the host level. In addition, the combination of network sensors and host-based sensors is the most effective way to provide threat coverage to a switched network.

The actions taken upon detection of an attack or unauthorized activity are determined by the administrator and fall into three categories:

- Notification** A server sensor notification can display alerts on the console, send an email (SMTP), or send an SNMP v3 trap, for every attack found.
- Log** All environment attack results and packet payloads are logged to the database.
- Activity** To block an unauthorized activity or attacks, the IBM Security Server Protection Agent for UNIX can disable a user account, block any network attack, or run a user-specified program.

Architecture

The IBM Security Server Protection Agent for UNIX uses a distributed architecture. The sensors perform the threat detection and response functions on critical network segments and servers. The event collector(s) collect events from the sensors for storage in the enterprise database and the IBM Security Server Protection console displays alarms, consolidates engine data, provides report generation capabilities, and acts as a centralized engine management point.

The relationship between sensors and managers is *many-to-many*. Several IBM Security Server Protection Agents for UNIX can report to a single event collector. Up to five event collectors can send data to a single enterprise database. This is all independent of the number of consoles used for reporting and command and configuration. This flexibility is useful for environments where there are geographical or organizational management boundaries.

Regarding placement of IBM Security Server Protection agents, the best rule is to install an agent on all servers containing business critical information.

In summary, IBM Security Server Protection Agents for UNIX can be installed on a variety of important systems: DMZ web security servers (such as a reverse proxy), internal web servers, critical file and application servers, and so on. These internal servers might be hosting such critical data as financial plans, sales databases, or engineering source code archives.

In addition to deploying server protection agents, IBM Security Network IPS appliances can also be deployed at strategic points in the network infrastructure, as explained in Chapter 8, “Network security solutions” on page 243.

Security compliance auditing

IBM Security Server Protection agents for UNIX help achieve regulatory compliance and provide centralized management of OS audit policies to protect against vulnerabilities that can arise from application design, development, or deployment flaws. This also ensures that all critical servers have a consistent audit policy to protect data confidentiality, and accessibility by monitoring logins, privilege escalations, and other system-level activity. IBM Security Server Protection Agents for UNIX integrate with the existing infrastructure and enable:

- ▶ Policy management and enforcement auditing
- ▶ Log monitoring
- ▶ Registry integrity monitoring
- ▶ OS auditing
- ▶ File integrity monitoring

Operating system and server sensors can automatically clear and set audit flags when you make changes to the policy so the user does not have to know what flags to set. They can also enforce auditing to be sure auditing flags are not accidentally changed by users or programs.

Operating systems

Windows, Solaris, HP-UX, and AIX operating systems are supported. See the online system requirements document in the IBM Security product Information Center for more specific details at:

<http://www.iss.net/>

9.4 Conclusion

We started the chapter by examining the Tivoli Endpoint Manager platform and explaining how it can be deployed to provide organizations with visibility and control of all their distributed endpoints. We then described the key features of the server and endpoint protection agents available today.



Virtual server security solutions

In this chapter, we discuss the need for protecting virtual environments. This discussion includes protecting the internal networks of virtual environments, the virtualization hypervisors, and the guest operating systems that run in a virtual environment. The functionality and components of the IBM Security Network IPS Virtual Appliance and IBM Security Virtual Server Protection for VMware are explained in more detail.

10.1 Virtualization defined

Although *virtualization* of hardware server resources was invented in the 1960s by IBM, recently with the advancement of x86 virtualization technologies such as VMware, Xen, and KVM, virtualization has become mainstream and has allowed individuals, corporations, and governments to get more effective use out of their existing hardware resources. There are many business drivers for virtualizing a server hardware environment, the most important of which are power, space, and cooling, which are the key factors for the widespread adoption of this technology. As with any computing resource, careful consideration must be taken with regard to security and protecting these new virtual environments.

Virtualization continues to expand; according to a recent IDC press release¹, 18.2 percent of all new servers shipped in the fourth quarter of 2009 were virtualized, representing a 20 percent increase over the 15.2 percent shipped in the fourth quarter of 2008. The size of the virtualization market in 2009 was US\$15.2 billion. Growing interest in cloud computing will fuel further demand for virtualization solutions as private cloud infrastructures are built within businesses and government. Often, these cloud infrastructures are built using virtualization technology, such as IBM Cloudburst or IBM Service Delivery Manager (ISDM).

Currently, the IBM vulnerability and threat protection technologies for virtualization have been focused on the VMware environment, specifically VMware ESX and VMware ESXi. VMware ESX and VMware ESXi are bare-metal hypervisors, meaning that they execute directly on the server hardware resources. Other technologies exist, such as VMware Server and VMware Workstation, which first require an operating system such as Linux or Windows to run the virtualization technology. Technologies such as VMware Server and VMware Workstation do not focus on data center virtualization capabilities, and therefore the IBM virtualization protection technologies focus on the bare-metal hypervisors VMware ESX and VMware ESXi (to be referred collectively as VMware ESX in the remaining sections).

¹ <http://www.idc.com/about/viewpressrelease.jsp?containerId=prUS2231661>

As shown in Figure 10-1, before virtualization, the operating system controlled all hardware resources and a typical model was to run a single application on a single physical server. With virtualization, the server runs a hypervisor, such as VMware ESX, and multiple operating systems and multiple applications can then share all the hardware resources, including any network interface cards, processors, memory, and storage. Each virtual machine (VM) has control of a set of resources provided to it by the hypervisor. In the case of VMware ESX, configured amounts of storage, memory, network connectivity, and CPU resources are defined per VM. Each VM operates independently of the others while the hypervisor is responsible for ensuring the adequate sharing of physical resources. With this shared mode, business applications executing on a virtual infrastructure need to be protected from threats between virtual machines, the internal virtual network, from the hypervisor, and from other threats on the physical network.

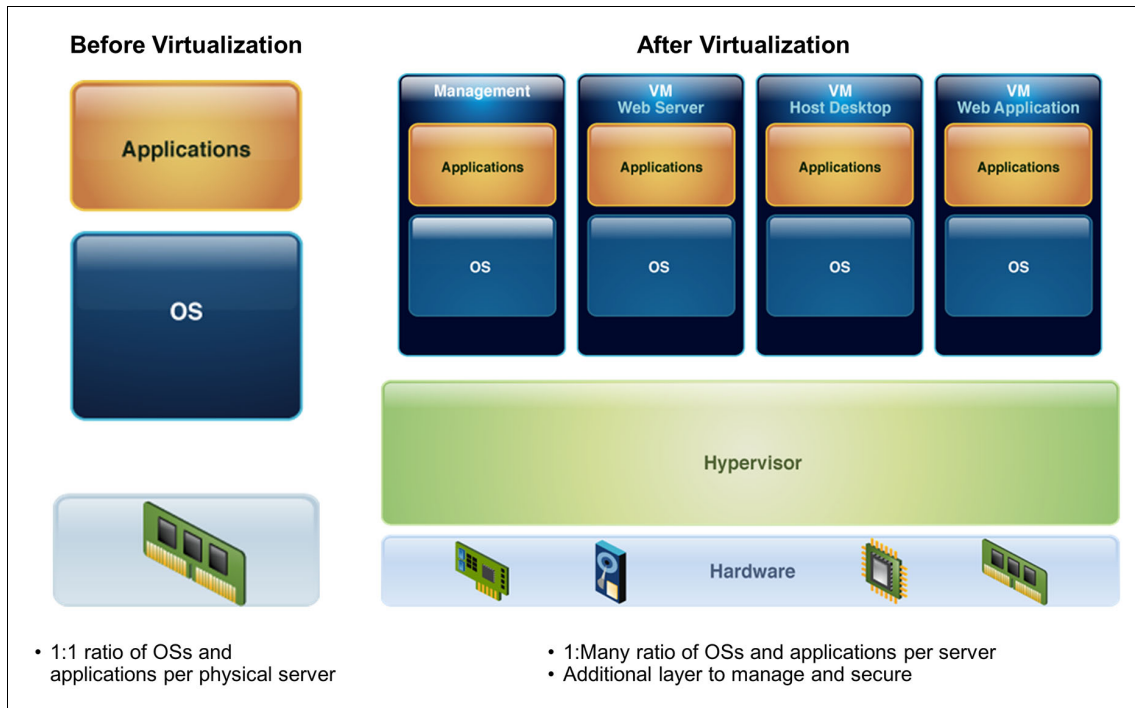


Figure 10-1 Physical server implementation versus virtual server implementation

Management server

VMware ESX and VMware ESXi can operate in a stand-alone mode, meaning that a single server running ESX or ESXi can host multiple different virtual machines. However, to gain more capabilities from the virtual environment, VMware has introduced the concept of a management server known as VMware Virtual Center or VMware vCenter, as shown in Figure 10-2.




Figure 10-2 Management Server, two VMware servers, with networking

This management console runs natively on its own physical hardware server on the Windows operating system. With VMware Virtual Center and VMware vCenter, new capabilities are available to VMware hypervisor and the virtual machines:

- **VMotion**

This capability allows a virtual machine to move transparently from one ESX host to another based on workload capability or through manual movement through the vSphere console while maintaining the application execution state.

- Distributed Resource Scheduling (DRS)

This capability monitors the usage of the ESX host and automatically rebalances the workload across the hosts by using VMotion to move the workload.

- High availability

This capability allows for the restarting of a virtual machine on another ESX host if the physical machine where the Virtual Machine was executing has failed.

- Fault tolerance

Provides synchronization and failover between primary and secondary virtual machines.

- VMSafe

A set of security APIs that allow for the creation of security products to protect the virtual server infrastructure.

In order for these capabilities to work, several new concepts have been introduced to the virtual environment, including shared storage, where the actual virtual machine files are available and seen across the multiple VMware ESX hosts, and shared networking, which is used to manage the virtual machines, maintain heartbeat, and synchronize memory between the physical hosts. Recently, vNetwork Distributed Switches were introduced, which allow a shared network switch to be available across multiple physical hosts, simplifying network configurations for the administrator. These capabilities are used considerably in virtual infrastructures and data centers. However, the capability for workload to move from one machine to another, the new connections between hosts, and the shared storage model all increase the security exposure for those environments.

With the ability to run many workloads within a virtual environment, additional capabilities have been implemented in the form of virtual network switches, where guest virtual machines can implement complex networking scenarios and support advanced features, such as VLANs and trunking. Often this virtual switch infrastructure, as shown in Figure 10-2 on page 340, is not visible to network administrators through their traditional network management and provisioning tools. Additionally, often the virtual server administrator has the rights to make modifications to the virtual network switches instead of the traditional network engineer. Recently, vendors such as Cisco systems have introduced the Nexus 1000V virtual switch, which runs its traditional switch operating system and has many of the features and management capabilities of traditional physical switches.

10.2 Virtualization threats

Organizations are under continuous pressure to deliver more functionality to their customers, provide more services to drive revenue, and attract customers while trying to do more with less. While power, space, and cooling limitations are a considerable barrier overcome by virtualization, often overlooked are the security implications regarding a virtual server deployment. The IBM Security X-Force Research and Development Organization (X-Force) continues to spend considerable resources researching virtualization vulnerabilities and threats, and the output of their work is reflected in the IBM Security virtualization products described in this book.

According to the IBM X-Force 2010 Mid-Year Trend and Risk Report², as shown in Figure 10-3, over half of the virtualization vulnerabilities are classified as medium or high severity. Because high severity vulnerabilities tend to be the easiest to exploit and provide full control over the attacked system, virtualization vulnerabilities represent a significant security threat to the organization.




Figure 10-3 Virtualization vulnerability severity by year reported

² The IBM X-Force Trend Reports are available at <http://www.ibm.com/services/us/iss/xforce/trendreports/>.

The X-Force categorizes virtualization vulnerabilities for bare-metal hypervisor environments into six categories that should be of concern to the security architect:

- ▶ Guest
Vulnerabilities that affect a guest virtual machine without affecting the bare-metal hypervisor.
- ▶ Web application
Vulnerabilities in web applications (typically the virtualization management applications) that affect the system on which the client browser is running.
- ▶ Virtualization system
Vulnerabilities that affect the virtualization system itself, that is, the entire virtualized environment, but do not arise from guest virtual machines.
- ▶ Escape to hypervisor
Vulnerabilities that allow an attacker to “escape” from a guest virtual machine to affect other virtual machines, or the hypervisor itself.
- ▶ Console
Vulnerabilities that affect custom management consoles or the management server.
- ▶ Web server
Vulnerabilities that affect a web server that implements a web application used by the virtualization system.

As identified in the IBM X-Force 2010 Mid-Year report, of the 373 virtualization vulnerabilities reported since 1999, 51 (14%) have known exploits.

10.2.1 Virtual machine sprawl

Although organizations implement virtualization to reduce or prevent physical server sprawl, reducing their power, space, and cooling needs, a new phenomenon has occurred that needs to be addressed by virtualization management and security products. *Virtual machine sprawl* occurs when the IT organization does not have the proper controls over the deployment of virtual machines within an environment, which exacerbates the configuration and change management issues that an organization often experiences. Virtual machines that have been turned off for months, or new virtual machines that are provisioned from an original standard image, often do not have the level of patching or security configuration changes necessary. When these servers are provisioned or turned on, old threats that previously have been addressed suddenly need to be handled by the security response team again.

10.2.2 Management console

The virtualization management console is responsible for many important functions. In a VMware ESX or VMware vSphere environment, the management console has access and control over the configuration of the hypervisors, the virtual machines, the network, and many additional details. The management console and its physical hardware, operating system, network connectivity, and application stack must be protected from threats just like any other server. Administrators often access the management console through a web browser or through a Windows-based client application, so the management infrastructure must be protected from threats originating from the administrator's workstation.

10.2.3 Console operating system

VMware ESX has a console operating system that allows an administrator to log on locally to the VMware ESX System using a command-line interface. This console operating system is based on Red Hat Enterprise Linux and provides most of the functionality of Red Hat Enterprise Linux. Some capabilities, such as an X Windows System interface, are not available, and the console operating system is executing inside a special virtual machine that has access to the VMware file system and special utilities that are available for manipulating virtual machines. The console operating system is a security exposure because vulnerabilities that exist in Red Hat Enterprise Linux can also exist in the console operating system.

The VMware ESXi system has a more restrictive console interface that does not provide all the functionality of an operating system, such as Red Hat Linux, that the VMware ESX console provides, and was developed by VMware to be more secure.

10.3 IBM Virtual Server security solutions

IBM offers several solutions for protecting virtualization environments. These solutions can be segmented into four areas.

1. Securing the individual virtual machine
2. Securing the management console
3. Securing the network within the virtualized infrastructure
4. Securing the entire virtualized environment

This concept is shown in Figure 10-4. In the following sections, we provide more details about each of the areas.

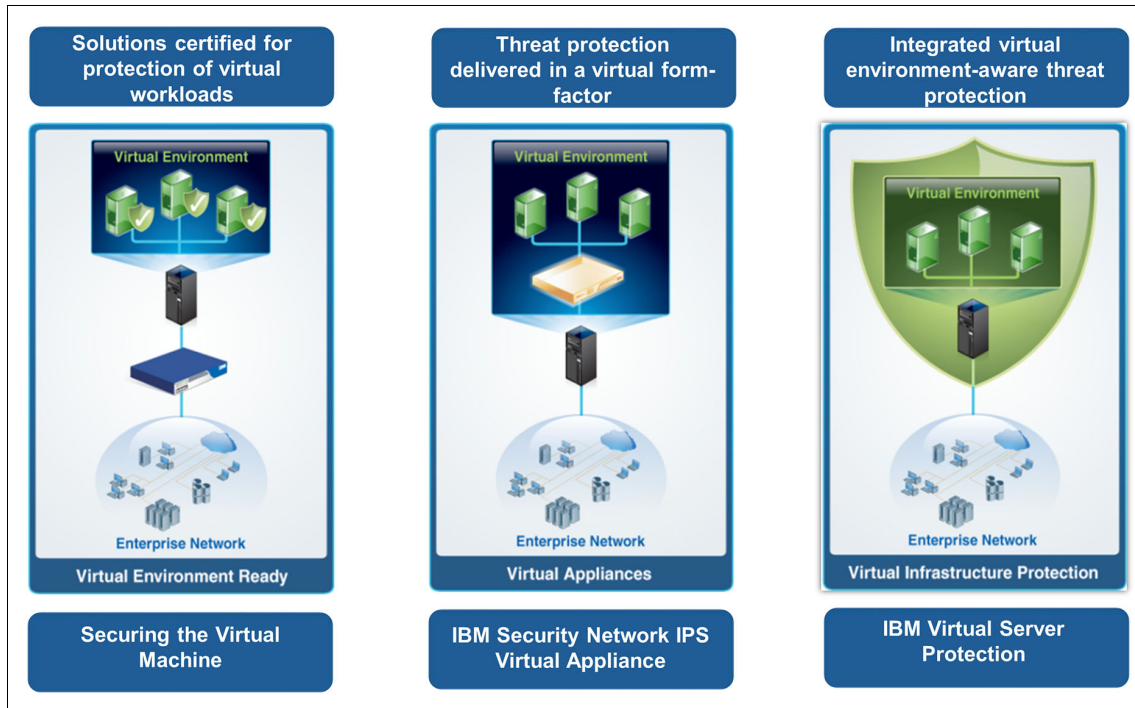


Figure 10-4 IBM Security Solutions for virtualized environments

10.3.1 Securing the virtual machine

IBM Security Server Protection, which we described in 9.2, “Proventia Desktop Endpoint Security” on page 318, can be installed directly onto running virtual machines to protect them by using a traditional host-based intrusion prevention system approach (HIPS). By providing this minimum level of protection, the security engineer ensures that active virtual machines are protected against the same types of threats to which physical servers would be exposed.

This solution provides no capability to protect the virtual switch infrastructure or to protect the hypervisor from threats.

10.3.2 Securing the management console

The virtual management console for VMware ESX is known as VMware Virtual Center, and for VMware vSphere, it is known as VMware vCenter. The VMware management console application has direct control over all of the hypervisors and virtual machines under its control, and it can be a target of threats against the virtual infrastructure.

The VMware management console is a Windows-based application that uses a relational database to store configuration and event data from the virtual infrastructure. VMware Virtual Center or VMware vCenter can be configured to use a Microsoft SQL Express database that ships with the VMware product, the full Microsoft SQL Server product, or an Oracle database. These database engines can be installed directly on the VMware management console machine, they can be configured to run on a separate physical machine, or even on a virtual machine residing in the infrastructure. VMware fully supports running both the VMware management console and its database in a virtual machine or installed directly on a physical machine, but most implementations install this software on its own physical infrastructure. The reason is that most implementations do not want to put the management application on the same infrastructure that it is managing.

If the VMware management console is installed on a physical server and the database to support this infrastructure is installed on its own physical server, this infrastructure must be protected on its own using a Host Intrusion Prevention System (HIPS) and Network Intrusion Prevention System (NIPS) separate from what is implemented for the actual virtual infrastructure. The HIPS and NIPS can be implemented by installing IBM Security Server Protection directly on the VMware management console and database server hosts and by ensuring that an IBM Security Network IPS is in place on the management network between the users and the physical server running the VMware management console software.

10.3.3 The IBM Security Network IPS Virtual Appliance

The IBM Security Network IPS Virtual Appliance provides intrusion detection and prevention for three different use cases:

1. Intrusion detection when the IBM Security Network IPS Virtual Appliance, running on a VMware server, is attached through a physical NIC to the SPAN port of a network switch.

2. Intrusion prevention, as shown in Figure 10-5, where the IBM Security Network IPS Virtual Appliance examines all traffic between virtual switches within a VMware ESX host.
3. Intrusion prevention where the IBM Security Network IPS Virtual Appliance is used instead of a physical Network IPS device and it is connected between two physical switches in the network.

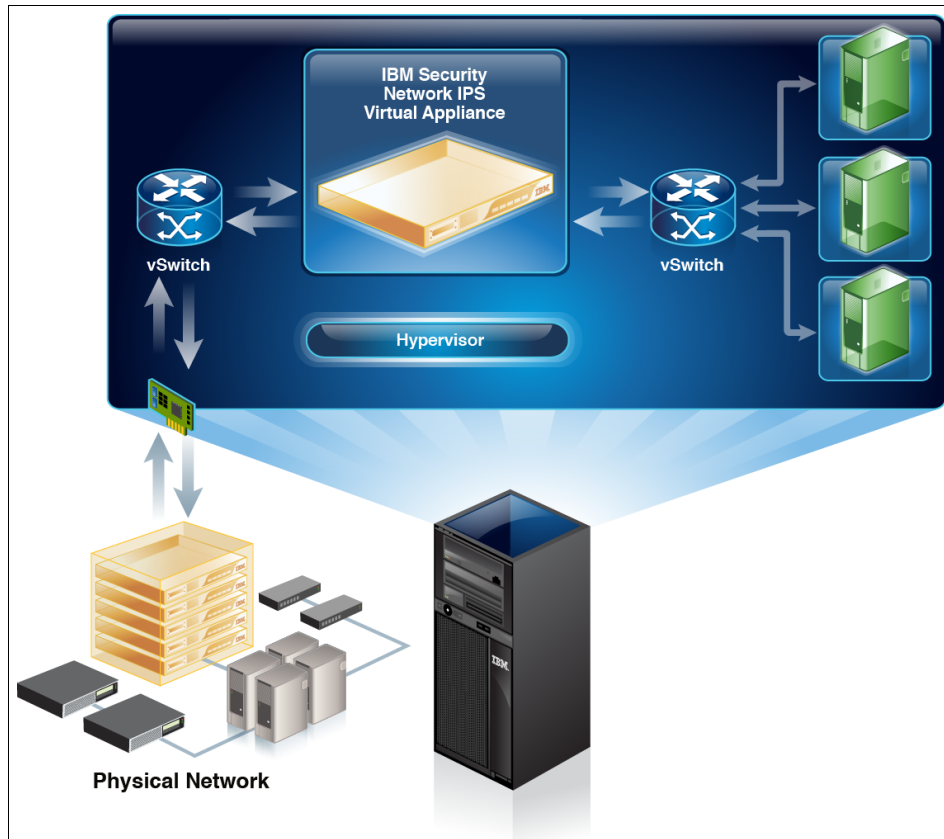


Figure 10-5 IBM Security Network IPS Virtual Appliance

The IBM Security Network IPS Virtual Appliance provides all the functionality of the IBM Security Network Intrusion Prevention System physical devices.

For more information about the IBM Security Network IPS Virtual Appliance, refer to 8.1.5, "Next generation virtual appliances" on page 258.

10.3.4 IBM Security Virtual Server Protection for VMware

The IBM Virtual Server Protection for VMware (VSP) offers integrated threat protection for VMware ESX and VMware ESXi that provides protection for multiple layers of the virtual infrastructure, including network, virtual machine (VM), and traffic between VMs. The transparent intrusion prevention and firewall in VSP provides multilayered IPS and firewall technology to protect the virtual environment at the core of the infrastructure. A high-level view of IBM Security Virtual Server Protection for VMware is shown in Figure 10-6.

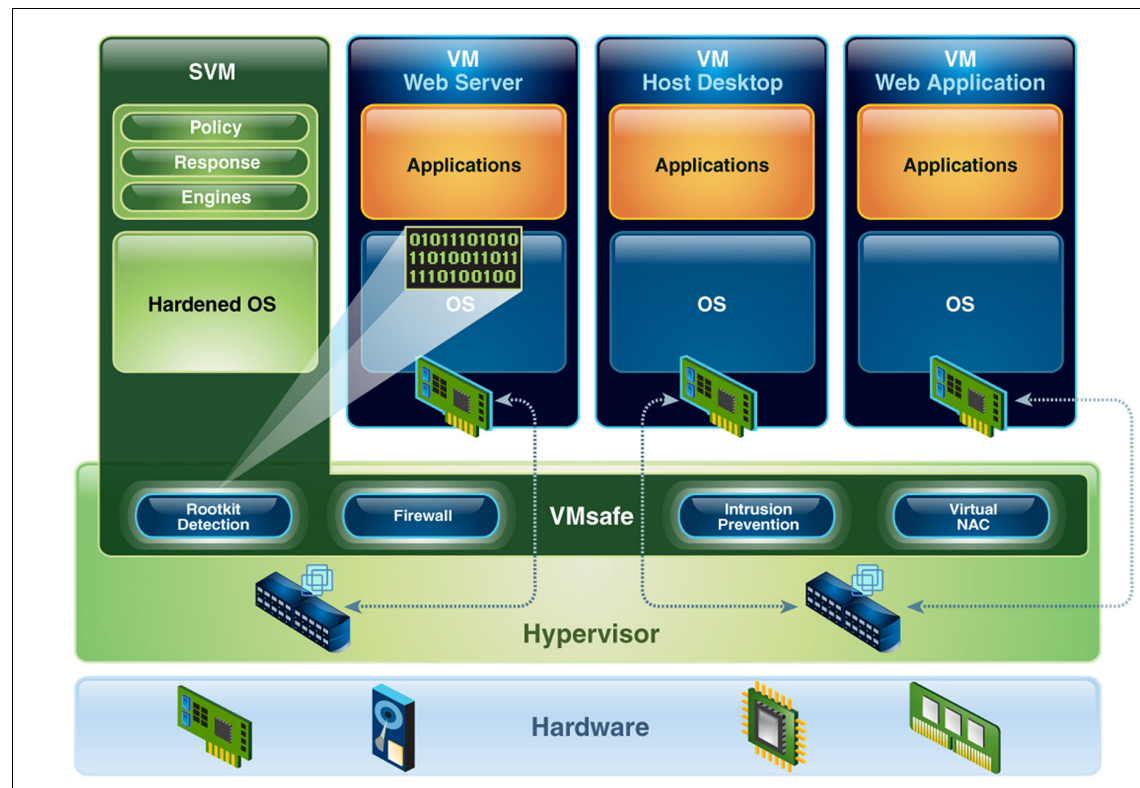


Figure 10-6 IBM Security Virtual Server Protection for VMware

To understand how the security capabilities of IBM Security Virtual Server Protection for VMware can be mapped to the IBM Security Blueprint³, see Figure 10-7 on page 350. This diagram shows the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using IBM Security Virtual Server Protection for VMware. This functional highlighting is applicable for the infrastructure service components as well.

Besides the fully highlighted elements, Figure 10-7 on page 350 also shows some medium highlighted elements. Although IBM Security Virtual Server Protection for VMware can be used to address such components to some degree, the respective area of coverage is not considered a core function of the product and thus is considered to be limited.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 10-7 on page 350 can be used as a quick reference of the functional security management aspects of IBM Security Virtual Server Protection for VMware. This reference can help us determine which functions of a solution can be covered by selecting this product.

³ For a detailed discussion of these elements, see Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

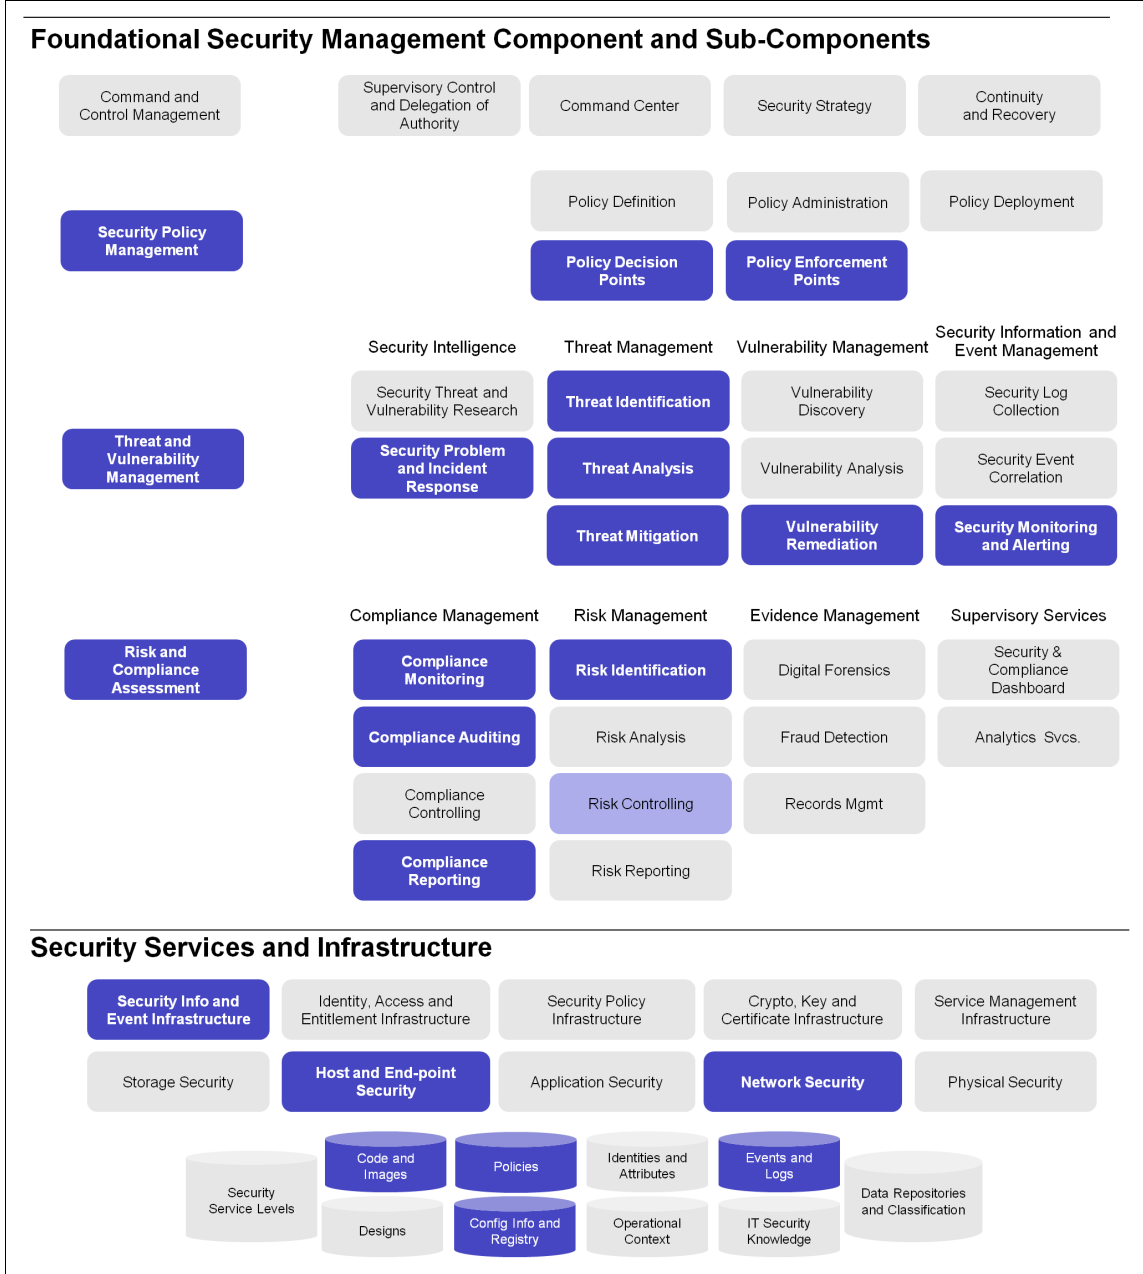


Figure 10-7 Mapping of IBM Security Virtual Server Protection for VMware to the IBM Security Blueprint

IBM Security Virtual Server Protection for VMware provides the capabilities shown in the following sections.

Virtual network access control

IBM Security Virtual Server Protection for VMware performs virtual network access control to quarantine or limit network access from a virtual server until the VM security posture has been confirmed. Through IBM Security SiteProtector, the administrator can define “trusted” or “allowed” lists and decide when a guest virtual machine is “approved” for network access. The administrator can then add that VM to the “trusted” list so that it is no longer subject to quarantine.

VM rootkit detection

Rootkits are particularly worrying because they can cause the kernel of the operating system to become infected with malware, and can conceal themselves from traditional security tools. Traditional rootkit detection is done by agents resident on the infected host and those agents are therefore subject to disruption or deception by the rootkit. IBM Security Virtual Server Protection for VMware inspects the guest virtual machine from outside the guest virtual machine and is therefore not subject to the rootkit disrupting the technology used to detect the rootkit.

Intra-ESX traffic analysis

IBM Security Virtual Server Protection for VMware monitors traffic between virtual servers on the ESX host to stop threats before they impact the host. Because VMWare ESX hosts are usually deployed in clusters, other tools such as an IBM Security Network IPS device should be used to inspect traffic between ESX hosts.

Inter-VM intrusion prevention

The IBM Security Network IPS Virtual Appliance is used to protect VM to VM communication and inbound communication from the physical network interface card(s).

Virtual infrastructure auditing

IBM Security Virtual Server Protection for VMware reports on privileged user activity, such as VMotion events, VM state changes (start, stop, and pause), and login activity to VMware.

Security Virtual Machine

A hardened version of Linux containing all the IBM Security Virtual Server Protection for VMware components is delivered as a virtual appliance for the VMware infrastructure, eliminating the need for administrators and software engineers to build the environment and install the software.

Automatic virtual machine discovery

Virtual servers operate in a dynamic state that can render traditional security technology ineffective. With IBM Security Virtual Server Protection for VMware, the security virtual machine can perform automatic discovery of all VMs, which helps increase security awareness and visibility across the virtual environment.

Integration with IBM Security SiteProtector

IBM Security Virtual Server Protection for VMware ships with a limited use version of IBM Security SiteProtector for managing the events and policies of the VSP infrastructure. Additionally, IBM Security Virtual Server Protection for VMware can integrate with an existing IBM Security SiteProtector infrastructure for those customers with Network Intrusion Prevention, Host Intrusion Prevention, and other SiteProtector-managed solutions from IBM.

Although the IBM Security Virtual Server Protection for VMware provides protection for the VMware virtual infrastructure, a defense-in-depth approach to enterprise security should be implemented as part of any implementation. The IBM Security Virtual Server Protection for VMware solution provides defense-in-depth for the virtual infrastructure, but it is only one layer of a larger enterprise security strategy. Network Intrusion Prevention and Host Intrusion Prevention (for physical hosts) should also be implemented to complement the overall VSP solution.

10.4 IBM Security Virtual Server Protection for VMware component model

The IBM Security Virtual Server Protection for VMWare component model is shown in Figure 10-6 on page 348. The component model of VSP provides Security and System Architects an understanding of how the VSP product works and how robust and secure architectures can be built using this technology.

IBM Security Virtual Server Protection for VMware is built using the VMware VMSafe API, which allows vendors to develop security products that integrate tightly with the VMware infrastructure. IBM has chosen to use specific parts of the VMware VMSafe API to enable robust security protection for the guest virtual machine environment while using and integrating with other IBM Security Solutions components, such as IBM Security SiteProtector and the technology within the IBM Security Network IPS product line.

10.4.1 Logical components

The logical component model shown in Figure 10-8 reveals the components within the IBM Security Virtual Server Protection for VMware product. Some optional components, such as the accelerator module, are shown, but can be configured separately, as described in this section. Additionally, each guest virtual machine is shown as having two connections to both the production and the management network. This configuration can be implementation specific and can be configured differently depending on the environment in your organization.

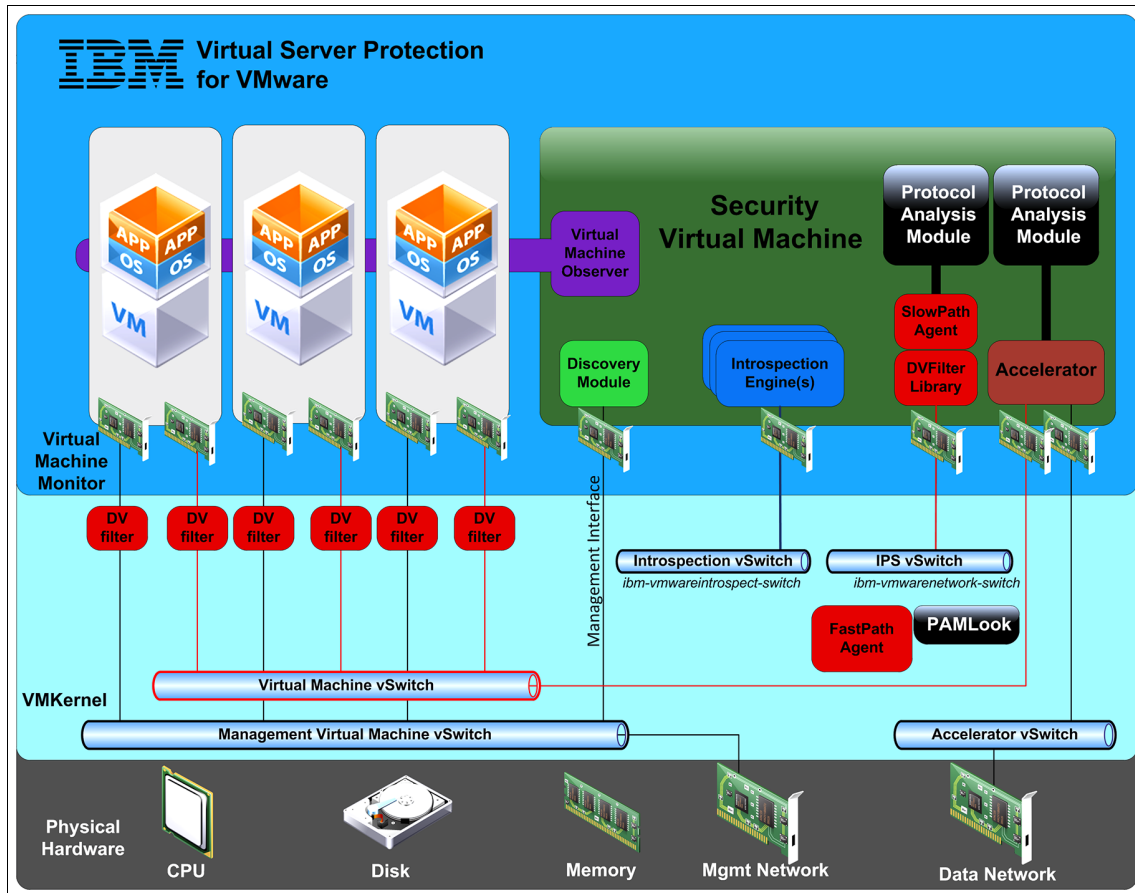


Figure 10-8 IBM Security Virtual Server Protection for VMware component model

vNIC

vNIC stands for Virtual Network Interface Card. Each virtual machine in a VMware environment has a vNIC, and it appears to the operating system as though it is a real network card. The normal parameters that can be changed in the operating system, such as link speed, duplex, and other parameters, are also available. VMware recommends that virtual machines are installed with the latest “VMware Tools”. These tools include a recommended virtual network and graphics driver specific for the guest operating system.

DVFilter Library

The DVFilter Library is a software interface provided by VMware that allows a solution to insert itself into the virtual packet stream between the vNIC on the virtual machine and the virtual switch. The FastPath Agent and the SlowPath Agents use the DVFilter Library to exchange data through a configured virtual switch in the VMKernel. The DVFilter Library also provides a mechanism for detecting and handling events, such as the attachment of a vNIC to a virtual switch.

FastPath Agent

The FastPath Agent executes in the VMware Hypervisor kernel (VMKernel) and processes captured packets from the vNICs and their associated DVFilter. There is one capture point (DVFilter) per monitored vNIC. The association between the FastPath Agent and the vNIC is defined in the .vmx configuration file for the monitored virtual machine.

The FastPath Agent can:

- ▶ Pass a packet without intervention
- ▶ Drop a packet from the packet stream
- ▶ Modify a packet
- ▶ Forward a packet to the SlowPath Agent (removing the packet from the packet stream)
- ▶ Inject a packet into the packet stream

For each captured packet, the FastPath Agent does some preliminary processing to decide whether or not the packet merits further analysis. If deep packet inspection of a packet is warranted, the FastPath Agent uses the DVFilter Library to send the packet through a configured virtual switch to the SlowPath Agent running in the Security Virtual Machine (SVM).

SlowPath Agent

The SlowPath Agent analyzes captured packets and decides whether or not they are safe for transmission. Safe packets are returned to the FastPath Agent via the DVfilter Library for reinsertion into the packet stream. The SlowPath Agent is where deep packet inspection is performed by the Protocol Analysis Module (PAM).

PAMLook

PAMLook is a precursor to the Protocol Analysis Module and runs in the VMKernel with the FastPath Agent. PAMLook examines packets at the FastPath Agent and determines if those packets should be sent on to the SlowPath Agent for deeper inspection or returned to the packet stream. PAMLook maintains a flowtable containing information about the current state of monitored connections.

Protocol Analysis Module

The Protocol Analysis Module is an integral component of the IBM Security Solutions product line and provides deep packet inspection functionality for the Security Virtual Machine. Details about the functionality of PAM and its features can be found in 6.5, “Protocol Analysis Module” on page 165.

IPS vSwitch (ibm-vmwarenetwork-switch)

The IPS vSwitch is used to forward packets (if selected by the FastPath Agent for further processing) from the FastPath Agent to the DVFilter Library and the SlowPath Agent running the Security Virtual Machine for deep packet inspection by the Protocol Analysis Module.

Accelerator vSwitch

The administrator can optionally configure an Accelerator vSwitch to connect the physical NIC used for data traffic into and out of the VMware ESX or vSphere environment to the vNIC on the Security Virtual Machine, which is connected to the optional Accelerator component. All traffic coming into and out of the VMware ESX or vSphere host would then pass through the Accelerator vSwitch and the Accelerator component.

Accelerator

The optional Accelerator component is used for the analysis of the physical network traffic coming into and out of the VMware ESX or vSphere host. Processing is sent directly through the Accelerator vSwitch to a vNICS running in the SVM for deep packet inspection through PAM. Packets are then forwarded through a VMware vswitch to the running virtual machine.

Use of the Accelerator and the Accelerator vSwitch allows improved processing of inbound and outbound packets of the VMware ESX or vSphere host itself, thereby improving the performance of the packet inspection process.

Implementation suggestion: The Accelerator and the Accelerator vSwitch are optional components and are disabled by default. For improved performance, use these components.

Introspection Engine(s)

The purpose of the Introspection Engine is to use the VMSafe CPU/Mem API functionality to provide rootkit detection capability. The rootkit detection capability in VSP uses a list of known rootkit attacks and inspects executing code within the guest virtual machines to ensure that a rootkit attack is not occurring. The VMKernel forwards information about the memory and CPU calls and states through the Introspection vSwitch to the Introspection Engine for the guest virtual machine being monitored running in the Security Virtual Machine. There can be multiple introspection engines in the one Security Virtual Machine for each supported guest virtual machine.

IBM customers have the ability to add additional rootkit attack signatures to the system in addition to the ones provided by IBM by contacting IBM Support. IBM Support can provide instructions about where to find information about an exploitation in the VSP logs, and how to use that information to construct a new rootkit signature.

Example 10-1 shows an example rootkit signature that was taken from the log files of the Security Virtual Machine.

Example 10-1 Example rootkit signature from the log files on the Security Virtual Machine

```
Anti-Rootkit:      [ProcessID:15290] engine:1 :  
ibm.research.antirootkit.monitors.interrupt_descriptor_table.idtcheckcallback  
<vm=WinXP SP3> SIG: Unauthorized change to idt@46:  
SHA256=189d43a4353d3c7c38595b8cc68069908fffeb6be0e66358f5f9cb119d954c7b;  
addr=0xf9f0d488; driver=\\??\\C:\\Documents and  
Settings\\Administrator\\Desktop\\ARK\\EvilRootKit.sys
```

Introspection vSwitch (ibm-vmwareintrospect-switch)

The Introspection vSwitch (virtual switch) is used by VMSafe to forward packets to the VMSafe CPU Mem Module and Introspection Manager running in the Security Virtual Machine.

Virtual Machine Observer

The Virtual Machine Observer (VMO) component is responsible for tracking the state of the virtual machines. The VMO monitors the guest virtual machines for events, such as starting, stopping, creating, and removing them. The observer function is facilitated by the VMware vSphere API. The Virtual Machine Observer is also responsible for associating DV filters with guest virtual machine vNICs in the guest .vmx files when new virtual machines come online.

VMware vSphere API

The VMware vSphere API is provided by VMware and provides access to the VMware infrastructure management components that are the managed objects that can be used to manage, monitor, and control life cycle operations of virtual machines and other VMware infrastructure components (data centers, data stores, networks, and so on). The Virtual Machine Observer uses the VMware vSphere API to determine events relevant to virtual machine status.

Discovery Module

The Discovery Module needs a routable path to reach each of the guest virtual machines, so it is connected through the management vNIC to a virtual switch, which has access to all existing and new virtual machines in the environment. The Discovery Module performs network scans of virtual machines and detects operating systems and open ports. It is used to determine if a virtual machine should exist within the infrastructure. The Discovery Module monitors for the threat of unauthorized creation or movement of virtual machines. The Discovery Module can be set to periodically scan the environment and maintains asset information (operating information and open ports) on the virtual machines that it finds.

Implementation note: VMware recommends separation of the management and production networks. If this recommendation is followed, guest virtual machines will not have any connection to the management network and therefore will not be scanned by the Discovery Module. For discovery scans to succeed, there must be a route from the management to the production network. This can be accomplished through additional vNICs on the guest virtual machines, as shown in Figure 10-8 on page 353, or by providing TCP/IP routing between the management network and the production network.

10.4.2 Physical components

When constructing a secure virtual environment, you must consider both the virtualized environment and the components that surround the virtual environment, which includes servers that host the management software components, such as VMware vCenter and IBM Security SiteProtector. Additionally, a Network IPS may be required to protect the virtual infrastructure and any other physical servers from network threats. Figure 10-9 identifies many of these components.

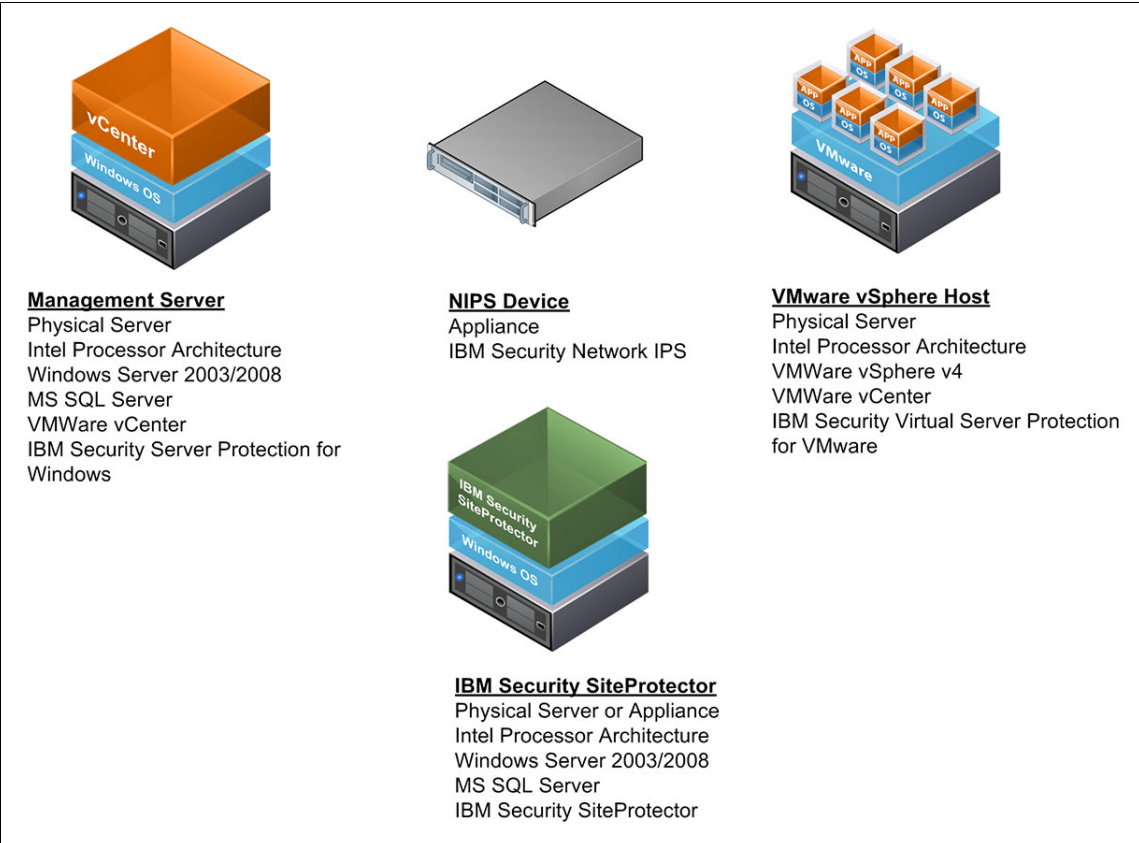


Figure 10-9 Physical components that enable a secure virtual infrastructure

10.4.3 Deployment architecture

An example deployment architecture for a web-based application is shown in Figure 10-10. This generic web server and application server architecture can be built on many different technologies, such as .NET or Java. In this architecture, traditional firewalls are deployed within an Internet DMZ, which also includes a Network IPS. The Network IPS protects any inbound or outbound traffic within the Internet DMZ from threats. An additional Network IPS device is deployed behind the DMZ firewall and in front of the virtualized server running the application and database server. Additionally, a Host IPS is installed on the web server to protect it and its operating system from threats.

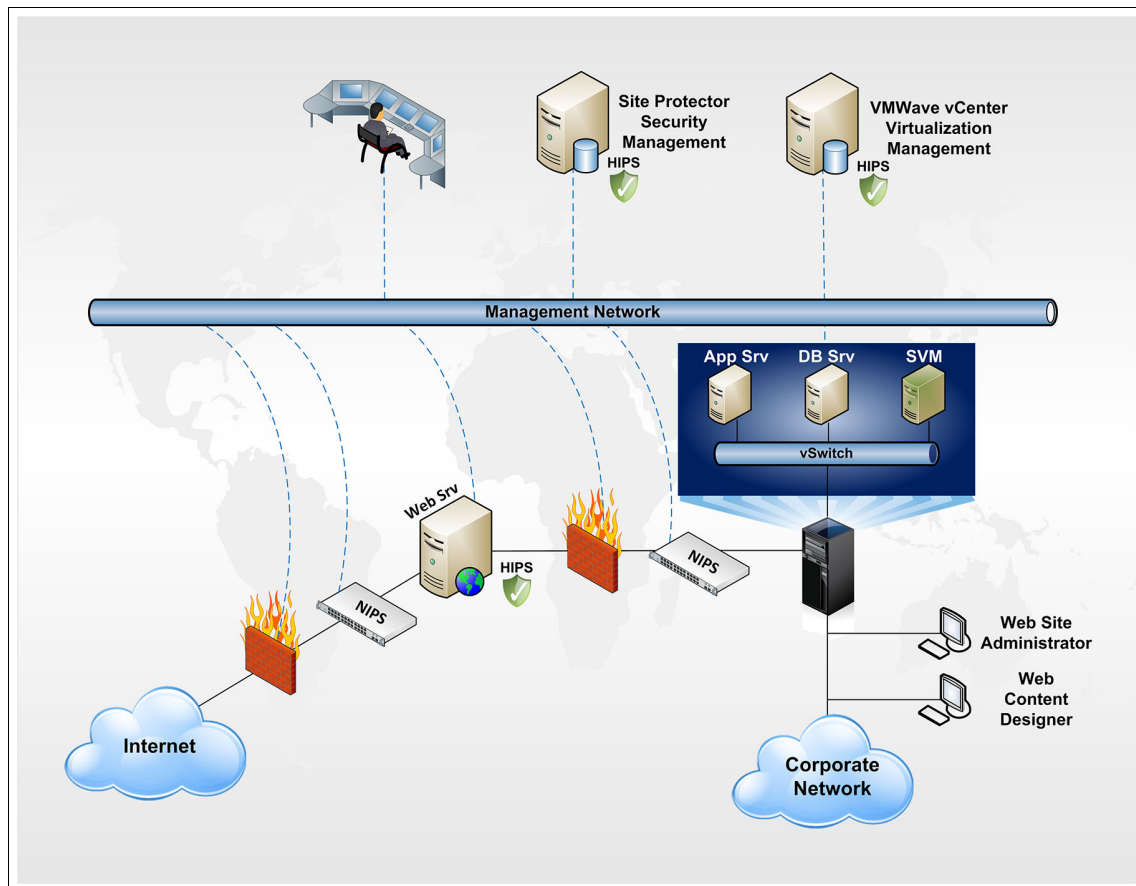


Figure 10-10 Example deployment architecture

In this example, a virtualized system contains three virtual machines:

- ▶ The application server, which contains the business logic of the application
- ▶ A database server for storing persistent data within the application
- ▶ The Security Virtual Machine (SVM), which is an instance of the IBM Security Virtual Server Protection for VMware software (VSP).

As explained earlier in this chapter, the VSP software provides robust capabilities for protecting the virtualized infrastructure.

Often overlooked in the deployment of a virtualized infrastructure, a *management network* is defined in this architecture as where the Network IPS, Host IPS, and VSP elements are connected, as well as the IBM Security SiteProtector server, the VMWare vCenter application, and the Security Operations Center (SOC). Best practices suggest that this management network be configured separately from the production data network, and that it provides the ability for all the security agents and devices to report their status to the SiteProtector server. The SiteProtector server and the VMWare vCenter servers are also protected with Host IPS by the installation of the IBM Security Server protection for Windows software.

Virtual desktop protection


Many organizations are now deploying *thin client architectures* using VMware vSphere and virtual desktops for their workstation environments. This is a cost-effective way of reducing power, space, and cooling as well as centralizing the management of desktops. Because desktops are susceptible to web browser vulnerabilities, including botnet harvesters, data theft, and malicious email-based attacks, a virtualized desktop deployment is an optimal use-case for IBM Security Virtual Server Protection for VMware. IBM Security Virtual Server Protection for VMware can prevent these types of attacks from spreading within the virtual desktop environment.

10.5 Conclusion

IBM Security Virtual Server Protection for VMware provides a comprehensive security platform for protecting your VMware vSphere environments. It includes a runtime version of IBM Security SiteProtector or integrates with your existing IBM Security SiteProtector infrastructure.

IBM Security Network IPS Virtual Appliance can provide intrusion detection and prevention when the IBM Security Network IPS Virtual Appliance is attached through a physical NIC to the SPAN port of a network switch. It can also provide intrusion prevention where the IBM Security Network IPS Virtual Appliance

examines all traffic between virtual switches within a VMware ESX host. Finally, you can use the IBM Security Network IPS Virtual Appliance instead of a physical Network IPS device, connected between two regular physical switches in the network.



Security services for Network, Server and Endpoint

In today's volatile economies, it has become increasingly important to ensure that the security solutions deployed by organizations not only deliver value, but also meet all of stringent business and technical requirements set by industry bodies and governments, all while still enabling business and reducing cost.

The logical question would then be, "How can I achieve all this and ensure that I have made a sound investment in a security solution?" The answer to that question is straight forward: By ensuring that you understand the risks that your organization faces and that you adequately protect your organization against those risks.

In this chapter, we discuss the IBM Security Services that address Network, Server and Endpoint security by taking a look at the activities required to ensure a sound business investment.

We distinguish between three service types delivered by IBM Security Services, namely *Professional Services*, *Managed Services*, and *Cloud Security Services*, and explore the value that each service component and its associated services can add towards maturing the security posture of organizations and how each of these service components fit into the IBM Security Framework and the IBM Security Services Integrated Lifecycle Methodology. An overview of all three IBM Security Services is shown in Figure 11-1.



Figure 11-1 IBM Security Services for Network, Server and Endpoint

Where:

► Professional Security Services

Professional Security Services from IBM Security Services deliver comprehensive, enterprise-wide security assessment, design, and deployment services to help you build effective information security solutions. Expert security consultants can show you how to implement network security best practices that can reduce online threats to your critical business assets.

► Managed Security Services

IBM Managed Security Services provide 24x7 monitoring and management of security technologies that you have deployed in your own IT environment. IBM provides a single management console and view of your entire security infrastructure, allowing you to mix and match by device type, vendor, and service level to meet your individual business needs while reducing your security costs, simplifying security management and accelerating your speed to protection.

► Cloud Security Services

IBM Cloud Security Services offerings harness the power of the IBM Virtual Security Operations Center platform to deliver high-value services that require little or no security device investment or maintenance, making the total cost of ownership much lower than what you would incur performing these security services in-house.

The IBM Security Services Integrated Lifecycle Methodology (Figure 11-2) enables organizations to see tangible results from services delivered within their environments based on their uniquely defined requirements.

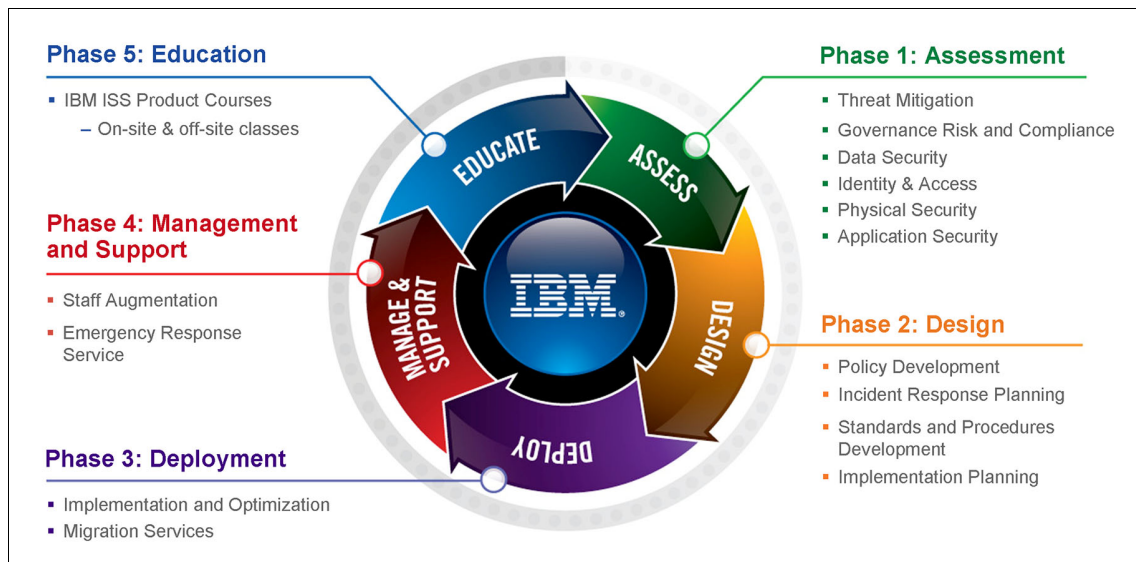


Figure 11-2 IBM Security Services Integrated Lifecycle Methodology

The IBM Integrated Lifecycle Methodology can, at a high level, be described as delivering the following items:

► Phase 1: Assessment

Action: Assess the current level of security effectiveness and strengthen network and security posture by identifying vulnerabilities and weakness against best practices and business requirements.

Result: Provide gap analysis and resolution recommendations between the current state and requirements.

- ▶ Phase 2: Design
Action: Design and documentation of policies, procedures, and architecture/solutions to ensure protection and extension of business capabilities.
Result: Creation of a gap closure plan for short and long-term resolution to ensure optimization of security infrastructure.
- ▶ Phase 3: Deployment
Action: Expert deployment, implementation, tuning, and change support.
Result: Helps organizations execute the gap closure plan, improve performance, and save costs.
- ▶ Phase 4: Management and Support
Action: Management of security infrastructure/program to meet defined business objectives.
Result: Ensures gaps remain closed and new gaps are not opened by providing improved protection, lowering total cost of ownership (TCO) and demonstrating compliance.
- ▶ Phase 5: Education
Action: Education and knowledge transfer of security best practices
Result: Improved employee understanding and skills related to security

In the remainder of this chapter, we focus on the following three IBM Security Service types:

- ▶ “Professional Security Services” on page 366
- ▶ “Managed Security Services” on page 393
- ▶ “Cloud Security Services” on page 417

11.1 Professional Security Services

IT security is critical to most organizations today. Industry bodies and governments exert a significant amount of pressure on organizations to comply with IT security regulations, standards, and legislation. Security, however, is rarely ever the core business of these organizations who, more often than not, lose focus of their primary business drivers while trying to address all of their IT security concerns.

Professional Security Services (PSS) for Network, Server and Endpoint refers to services used by organizations to elevate their security posture by making use of external entities that possess the necessary skills to address their IT security requirements, enabling them to remain focused on their primary business objectives.

IBM Security Services PSS focuses solely on security, with industry-leading security research, methods and tools that help organizations achieve their risk-management goals by reducing risk, achieving regulatory compliance, and maintaining business continuity. These goals are accomplished by using both strategic and technical consulting capabilities and using proven consulting methods, based on the globally accepted International Organization for Standardization (ISO) 27000 series best security practices. All of this, coupled with the use of proprietary tool sets, the latest threat intelligence, and advanced countermeasures, helps organizations build effective security programs that protect and enhance business operations through the value delivery chain.

This section describes Professional Security Services, as delivered by IBM Security Services, and is intended to facilitate the architecting of security solutions for Network, Server and Endpoint.

11.1.1 Penetration Testing Service

Let us examine this service offering in more details.

What is a penetration test

A penetration test offers an invaluable and compelling way to establish a baseline assessment of security as seen from outside the boundaries of an organization. Properly executed penetration tests provide evidence that vulnerabilities do exist and that network penetrations are possible. More importantly, they provide threat information that used to further enhance security through remediation and adoption into the overarching security program.

A penetration test simulates covert and hostile network activities to identify specific exploitable vulnerabilities and to expose potential entryways to vital or sensitive data that could pose increased risk and liability to the organization.

Qualified security professionals who perform penetration tests attempt to gain access to online assets and company resources through the network, servers, and endpoints from both the internal and external perspectives, simulating both internal and external threats. The results of a penetration test clearly articulate security issues and recommendations while creating a compelling business case for management to support a security program.

Why is penetration testing effective

Penetration testing provides a sound baseline for enterprise security risk. Most organizations underestimate their security exposure and overestimate their capacity and resources for their internal IT and IT security staff to address risk. Conducting a penetration test can be considered as the first step in understanding the organization's current security posture, and a vital building block in the organization's risk mitigation, reduction, and remediation strategy. Penetration tests help identify security gaps and outline where to implement solutions and services to enhance security.

Penetration testing helps create a compelling business case. Penetration tests help administrative users and security managers justify budget increases where needed and also deliver a message of urgency and criticality at the executive level. Well documented penetration test results that expose the susceptibility of customer and financial information create a compelling business case to enable executives concerned with company finances, liability, and reputation to timeously mitigate the discovered risk.

Penetration testing performs due diligence and independent audits. Most organizations are required to assess their security posture on a regular basis to ensure compliance, evaluate the progress of mitigation plans, and to account for new threats. An unbiased analysis and penetration test can help focus the appropriate resources where they are needed most. An independent security audit also provides evidence of due diligence in the legal context of protecting online assets by limiting and minimizing risk to C-level management and shareholders.

Penetration testing can aid in achieving regulatory compliance requirements. Regulatory and legislative requirements often identify penetration testing as a business necessity specifically for the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), OCC Built-in 2014, and Graham Leach Bliley as a component of demonstrated due diligence.

Penetration testing can establish risk baselines for new applications and systems. By conducting penetration tests against new applications, especially web applications, and newly deployed systems, organizations are able to further pinpoint potential risk areas, as penetration test results illustrate how vulnerabilities are exploited to gain access to systems and information.

Penetration testing helps protect against inter-connected business partner risk. Online commerce initiatives require organizations to grant access to business partners, business-to-business exchanges, customers, and several other trusted entities. It is therefore in the best interest of organizations to conduct regular penetration tests on all connected entities to ensure that they maintain an appropriate security baseline.

Penetration testing offers validation. Penetration tests provide critical validation feedback of newly established security practices and controls.

Requirements for a successful penetration test

Key parameters must be established for a penetration test to deliver useful, timely, and accurate results. Penetration tests should cover the full range of the threat spectrum, from the presence of antivirus software to the presence of malicious code and vulnerabilities that might enable denial of service and other sophisticated attacks. Penetration testing should also deliver clear and unambiguous results that address both technical and business objectives. The consultants who perform penetration tests should follow clearly defined and robust testing methodologies, use the most up to date vulnerability research available, and possess creative instincts to manipulate tools and systems in both typical and unconventional ways. Experienced technicians who administer penetration tests must know how to gain the maximum results with minimal disruption to normal business operations.

The penetration testing process

Organizations must be able to select their preferred approach to penetration tests. Working with consultants, organizations should indicate a particular system or part of a network that they want tested. Alternatively, organizations can choose to follow the open or “blind” approach, which gives the consultants performing the penetration test the freedom to assess the organization’s entire infrastructure to identify exploitable vulnerabilities. Penetration tests are also often used to determine the effectiveness of external monitoring services. After a set of basic parameters have been established, penetration tests are usually conducted following a three-phase process.

Phase one: Reconnaissance

In the reconnaissance phase, penetration testing consultants gather as much information as they possibly can. This normally includes finding network, system, and endpoint information that consists of network ranges used by the organization, virtual private networks (VPNs), firewalls, network intrusion prevention systems (NIPS), web applications, and other security protection technologies that might be used by the organization.

The intention of the reconnaissance phase is to mimic the actions that an external attacker would take to enumerate resources that might be exploitable from the Internet. Effective penetration tests follow industry accepted principles.

Phase two: Vulnerability analysis

During phase two, penetration testing consultants try to detect exploitable vulnerabilities in the systems discovered in the reconnaissance phase. In addition to using a variety of automated vulnerability scanning tools, penetration testing should also make use of manual probing to determine if there are any exploitable components in systems or their configurations.

Phase three: Penetration

In phase three of the testing, consultants exploit the systems and associated vulnerabilities that were discovered during the previous phases. This is done to confirm that valuable information can be obtained from the systems being tested. This phase of the process requires experienced consultants knowledgeable in application development, scripting, attack techniques, and various security technologies. After the conclusion of the testing, consultants should deliver a comprehensive report detailing all the vulnerabilities found during the penetration test. This report usually includes remediation actions accompanied by a remediation plan to enable the organization to quickly and effectively remediate all or most of the discovered vulnerabilities. The report is also accompanied by evidence gathered during the penetration test.

IBM Security Blueprint solution pattern

To understand how the security capabilities of the Penetration Testing Service can be mapped to the IBM Security Blueprint¹, see Figure 11-3 on page 371. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the Penetration Testing Service. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-3 on page 371 can be used as a quick reference of the functional security management aspects of the Penetration Testing Service. This reference can help us determine which functions of a solution can be covered by selecting this product.

¹ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

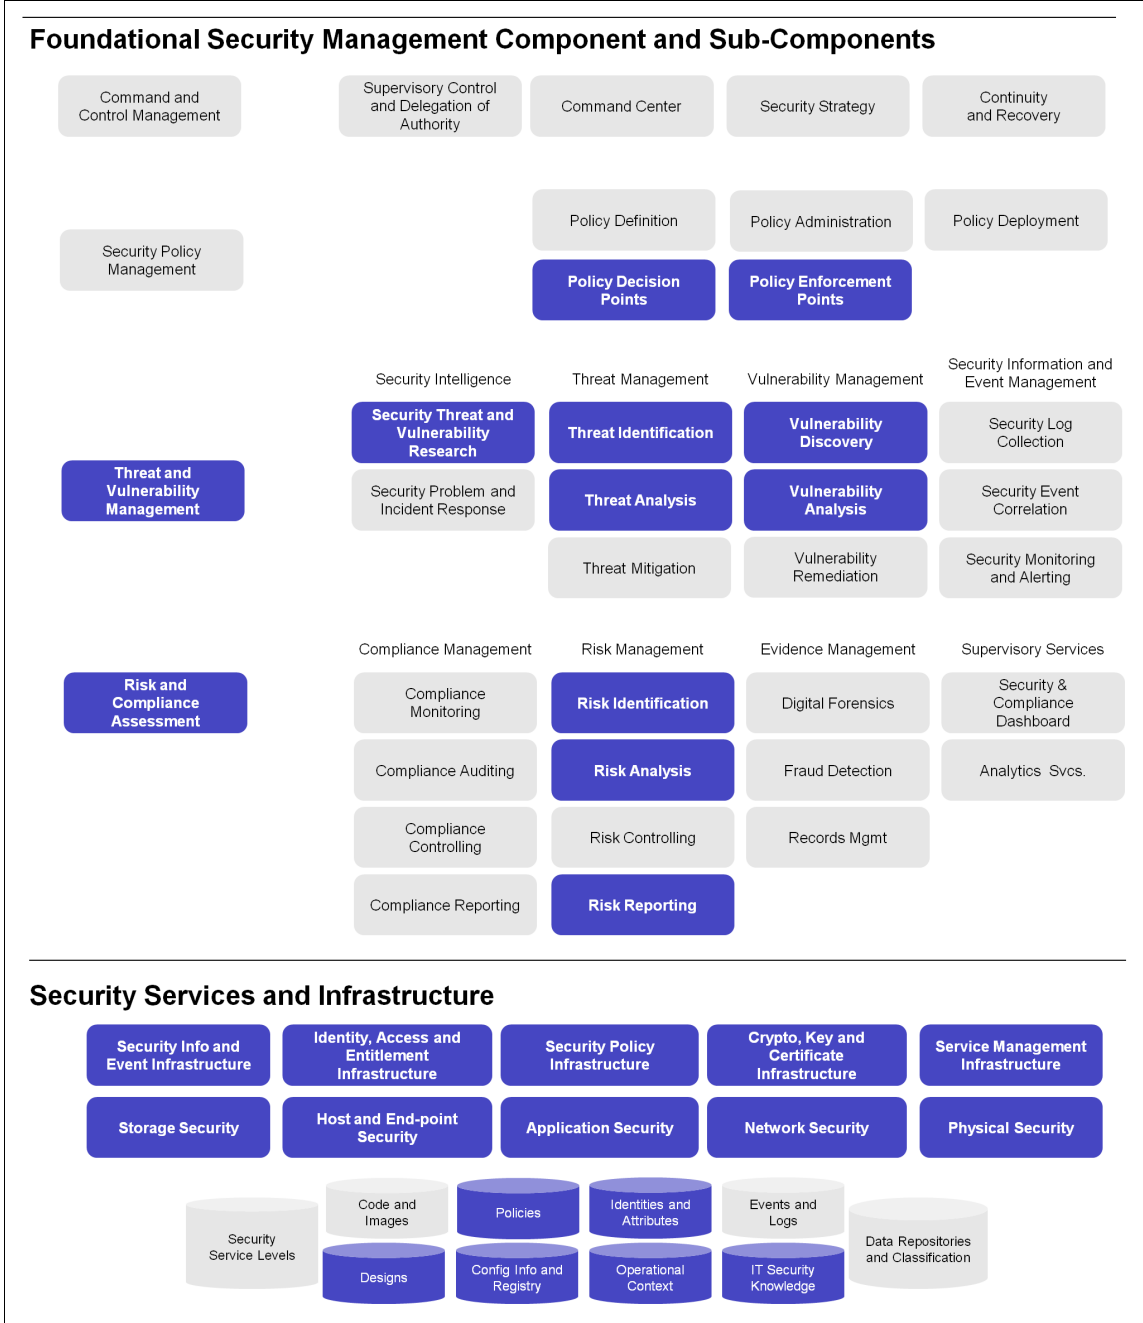


Figure 11-3 IBM Security Blueprint solution pattern for the Penetration Testing Service

11.1.2 Information Security Assessment

Let us examine this service offering in more detail.

What is an Information Security Assessment

An Information Security Assessment (ISA) affords organizations the ability to comprehensively evaluate their security landscape in relation to industry best practices and regulatory requirements. Information regarding current controls is gathered and their effectiveness evaluated to identify risks and to provide detailed and actionable mitigation and protection information.

An Information Security Assessment provides organizations with a broad view of the security landscape that spans not only technical controls, but also looks at components that are often overlooked in traditional technical assessment services, such as Information Security Policies and regulatory and standards requirements.

Why are Information Security Assessments effective

An Information Security Assessment provides organizations with a clear understanding of their current information security risks and the potential impact that the exploitation of vulnerabilities associated with those risks could have on affected network, server, and endpoint infrastructures. This action leads to increased internal awareness of information security risks which in turn leads to more informed decision making based on the gaps in security controls, policies, and processes.

It also provides specific, actionable plans to improve the overall security posture of organizations based on business needs, which encourages and enables a more proactive approach towards threat protection and enhances efforts to meet compliance requirements.

Requirements for a successful Information Security Assessment

Key parameters must be established for an Information Security Assessment to deliver useful, timely, and accurate results. Information Security Assessments should offer comprehensive features that ensure that the evaluation of information security architectures, policies, and processes are based on accepted standards, such as ISO 27002, and industry best practices derived from experience in delivering services, such as Information Security Assessments. Information Security Assessments should also follow strict, comprehensive methodologies that yield clear, unambiguous results that address not only technical but also business objectives.

Consultants who conduct Information Security Assessments should have a broad enough knowledge base to assess vulnerabilities throughout the computing environment, analyze security controls for both internal and perimeter infrastructure, and conduct meaningful interviews with technical and business-oriented staff members while ensuring that the final deliverable speaks to technical and business audiences and provides specific recommendations and action items for addressing identified issues.

Information Security Assessment delivery process

Because an Information Security Assessment is such a comprehensive service, it is essential that organizations understand what it is that is being delivered and which parts of their security environments will be assessed. A strong, proven methodology must be followed to identify and ultimately reduce risk.

Phase one: Initiation

Phase one of the Information Security Assessment delivery is made up of all the activities that the security consultants assigned to the Information Security Assessment perform to ensure that the scope of the Information Security Assessment is defined and agreed upon and that all relevant materials for assessment has been received and approved for the assessment activities.

The consultants typically review the scope of the project, discuss potential scheduling options with the customer, discuss potential project kickoff timing options, and verify all customer contact information.

After the scope discussion is complete, the team generates the appropriate pre-assessment questionnaires that align with the scope of the project and submit the pre-assessment questionnaires to the customer for completion. They also at this time verify that the project schedule is appropriate and finalize it with a customer's point of contact. Consultants are also provided with the Work IDs and any other necessary information to perform the assessment activities.

When the scope has been finally agreed upon by all parties, the team verifies receipt of all completed questionnaires and schedules a kick-off meeting with the customer. During the kick-off meeting, all of the team members are introduced to each other and roles and responsibilities are finalized. All questionnaires are reviewed for accuracy and the consultants verify that all information contained in the questionnaires is accurate and agreed upon by both the delivery team and the customer. They then go on to discuss the preliminary schedule in detail with all participants to verify that the schedule is still accurate and ensure that the customer has contact information for all of the project resources.

Phase two: Remote activities

In phase 2, the following activities are usually conducted remotely:

- ▶ External Vulnerability Assessment
The purpose of this activity is to identify and validate potential vulnerabilities that exist on a customer's Internet facing infrastructure.
- ▶ War-Dial Activities
The purpose of this activity is to identify active modems within a pool of in-scope phone numbers and identify potential weaknesses associated with modem access.
- ▶ Application Assessment Vulnerability Scan
The purpose of this activity is to identify and validate vulnerabilities that exist in the in-scope applications.

Phase three: Onsite activities

In phase 3, the following activities are usually conducted onsite:

- ▶ Security Architecture Review
The purpose of this activity is to identify gaps within the customer's environment as it relates to security best practices associated with logical placement of devices, secure device configurations, and remote access configurations.
- ▶ Security Policy and Procedure Review
The purpose of this activity is to determine gaps within the customer environment as it relates to security policies and procedures based on the industry acceptable ISO 27002 framework. This activity typically includes the Information Security Policy with the following subcomponents:
 - Organization of Information Security
 - Asset Management
 - Human Resources Security
 - Physical and Environmental Security
 - Communications and Operations Management
 - Access Control
 - Information Systems Acquisition, Development, and Maintenance
 - Information Security Incident Management
 - Business Continuity Management
 - Compliance

This is not meant to be an in-depth review of the organization based on ISO 27002, however, the review of policies and procedures should highlight the major gaps found within them.

► Technical Security Controls and Mechanisms Review

The purpose of this activity is to identify gaps associated with the security controls and mechanisms that are currently in place within the customer environment as they relate to industry best practices. This activity typically includes interviewing key staff members responsible for, and the policies and configurations, of the following systems:

- Antivirus configurations
- Standard Operating System Configuration/Policy documentation
- Network/Host-based IDS/IPS, File Integrity configurations
- Vulnerability Management process documentation
- Encryption Standards
- Content Filtering configurations

► Internal Vulnerability Assessment

The purpose of this activity is to identify and validate potential vulnerabilities that exist on the customer's internal networks.

► System Security Assessment

The purpose of this activity is to assess the security of systems and to identify gaps associated with the operating system configurations.

► Database Vulnerability Testing

The purpose of this activity is to identify the vulnerabilities associated with databases in the customer's environment.

► Social Engineering

The purpose of this activity is to test the physical security controls and security awareness of employees within the customer's environment. There are typically two phases to this activity: remote social engineering and onsite social engineering.

– Remote Social Engineering Activities

The purpose of this activity is to remotely test the security awareness, processes and procedures of the customer's employee base by calling either predefined phone numbers or through passive blind reconnaissance efforts.

– Onsite Social Engineering Activities

The purpose of this activity is to test the security awareness, processes, and procedures of the customer's employee base while onsite.

► Physical Assessment Review

The purpose of this activity is to test the physical security of the customer's locations.

- ▶ **Wireless Penetration Testing**
The purpose of this activity is to attempt to penetrate into a customer's internal network via wireless network vulnerabilities.
- ▶ **Wireless Assessment**
The purpose of this activity is to assess the effectiveness of security controls that have been implemented as part of the wireless network design.
- ▶ **Mainframe Assessment**
The purpose of this activity is to review the current configurations and settings of the mainframe (through technical and interview based verification) for security best practices associated with a mainframe.

Phase four: Documentation, review, and project closure

When all onsite and remote activities have been completed, final documentation is created. This typically includes a summary of strengths, findings, recommendations, overall risk rankings, detailed findings and recommendations, and any other relevant supporting documentation deemed necessary for the final report. When the final documentation has been submitted, the customer is allowed an appropriate amount of time to review the final documentation. This gives the customer the ability to ask questions or raise concerns about the documentation.

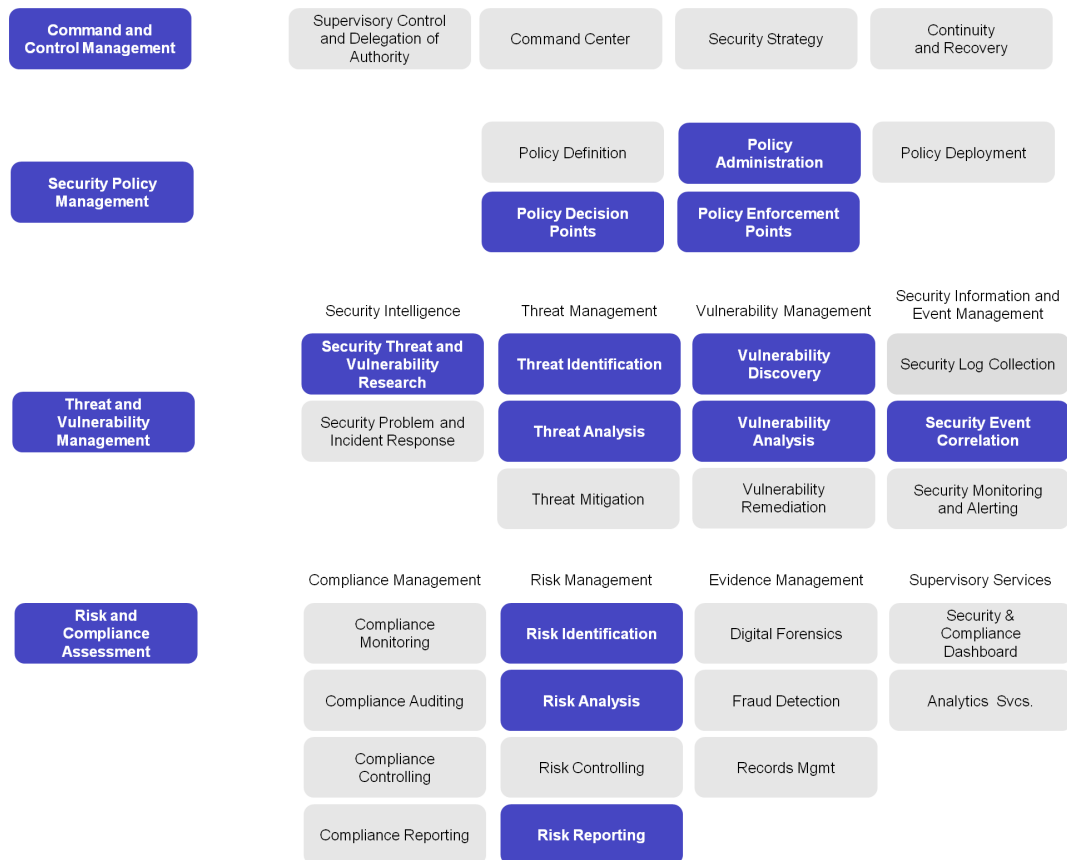
IBM Security Blueprint solution pattern

To understand how the security capabilities of the Information Security Assessment can be mapped to the IBM Security Blueprint², see Figure 11-4 on page 377. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the Information Security Assessment. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-4 on page 377 can be used as a quick reference of the functional security management aspects of the Information Security Assessment. This reference can help us determine which functions of a solution can be covered by selecting this product.

² For a detailed discussion of the elements, refer to Chapter 2, "The components of the IBM Security Blueprint" on page 31 and Chapter 3, "The Network, Server and Endpoint solution pattern" on page 93.

Foundational Security Management Component and Sub-Components



Security Services and Infrastructure

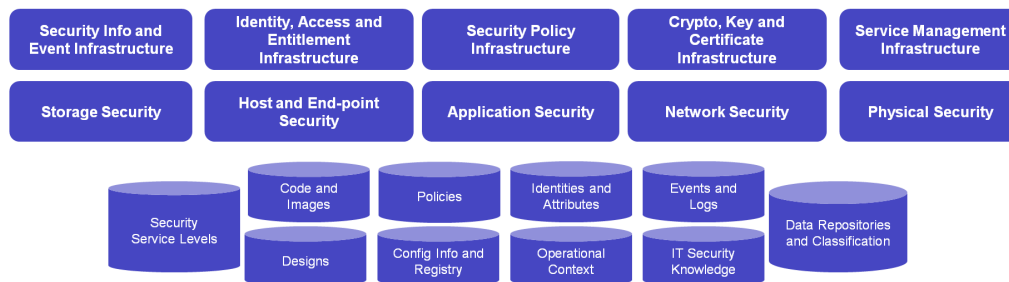


Figure 11-4 IBM Security Blueprint solution pattern for the Information Security Assessment

11.1.3 Deployment and Migration Services

Let us examine this service offering in more detail.

What are Deployment and Migration Services

Deployment and Migration Services are designed to assist organizations with the deployment, configuration, tuning, installation, and integration of Network, Server and Endpoint solutions. The appropriate deployment of these solutions into existing environments often means the difference between trying to put out fires and providing proactive protection to the organization.

Deployment efforts create the perfect opportunity to integrate information security technologies for maximum effectiveness while ensuring that the deployment of new solutions do not hinder employee productivity, overtax IT resources, or elevate deployment costs.

Why are Deployment and Migration Services effective

As most organizations do not consider information security their core business, they do not hire security experts capable of deploying complex solutions within their environments, but rather invest in building internal skills to enable the day-to-day operational running of information security solutions.

This effectively leaves a gap when it comes to the acquisition of new solutions that must be deployed into and interact with existing infrastructures, so organizations look towards external entities that are capable of providing these types of services.

Deployment and Migration Services are therefore designed to facilitate collaborative deployment that ensure seamless integration to provide better protection faster and more effectively while causing the least amount of disruption to the daily operations of the company.

Experienced consultants define strategy, scope, objectives timelines, and milestones for the successful deployment of solutions, giving organizations the opportunity to minimize the impact on their internal IT resources and shortening implementation cycles.

Requirements for successful Deployment and Migration Services delivery

To be fully effective, Deployment and Migration services must provide at least the following items to the organization:

- ▶ Assist in the completion of crucial planning for deployment projects in an efficient time frame.

- ▶ Provide information security expertise capable of developing a sound approach and strategy for the implementation and management of security solutions.
- ▶ Identify and agree upon ownership of potential project risks.
- ▶ Enhance communication among the various teams responsible for the implementation, integration, and management of the new solution.
- ▶ Promote security awareness within the organization.
- ▶ Avoid making costly mistakes or delays by using design and deployment expertise.

To be fully effective, information security solutions for Network, Server and Endpoint must be carefully deployed, configured, tuned and integrated within an organization's unique environment,

Deployment and Migration Services delivery process

The methodologies and processes used to effectively deliver Deployment and Migration Services are vital for these projects to be successful. As with most services, a phased approach is used to ensure precision and effectiveness of deployment and migration projects.

Phase one: Architecture review and planning

Phase one is used to identify network infrastructure and communications requirements by reviewing specific business requirements, which typically include reporting and management processes. It also reviews existing network infrastructure to determine whether network requirements for the selected solution have been met, and identifies specific network and communications issues that have to be addressed and resolved before deployment or migration commences. All findings and recommendations are carefully documented for use in a final report.

Phase two: Architecture and design

Phase two focuses on the architecture and design of the selected solution(s) for deployment or migration into the existing infrastructure. This phase ensures that the new solution(s) seamlessly integrate into the existing environment and conforms to the organizations existing security architecture and design standards and practices.

Phase three: Deployment or migration

Phase three is the actual deployment or migration of the chosen security solution(s) into the organization's existing environment. This phase includes the installation of the security technology on applicable systems, integration within the network infrastructure, optimization of the setup and configuration,

implementation of appropriate management and communication components, and establishment of the appropriate alerting and reporting capabilities.

Phase four: Skills transfer and hand over

The fourth and final phase focuses on skills transfer to the organizations IT resources to ensure that the new solution(s) is efficiently and effectively managed and finally handing over the operational solution to the organization.

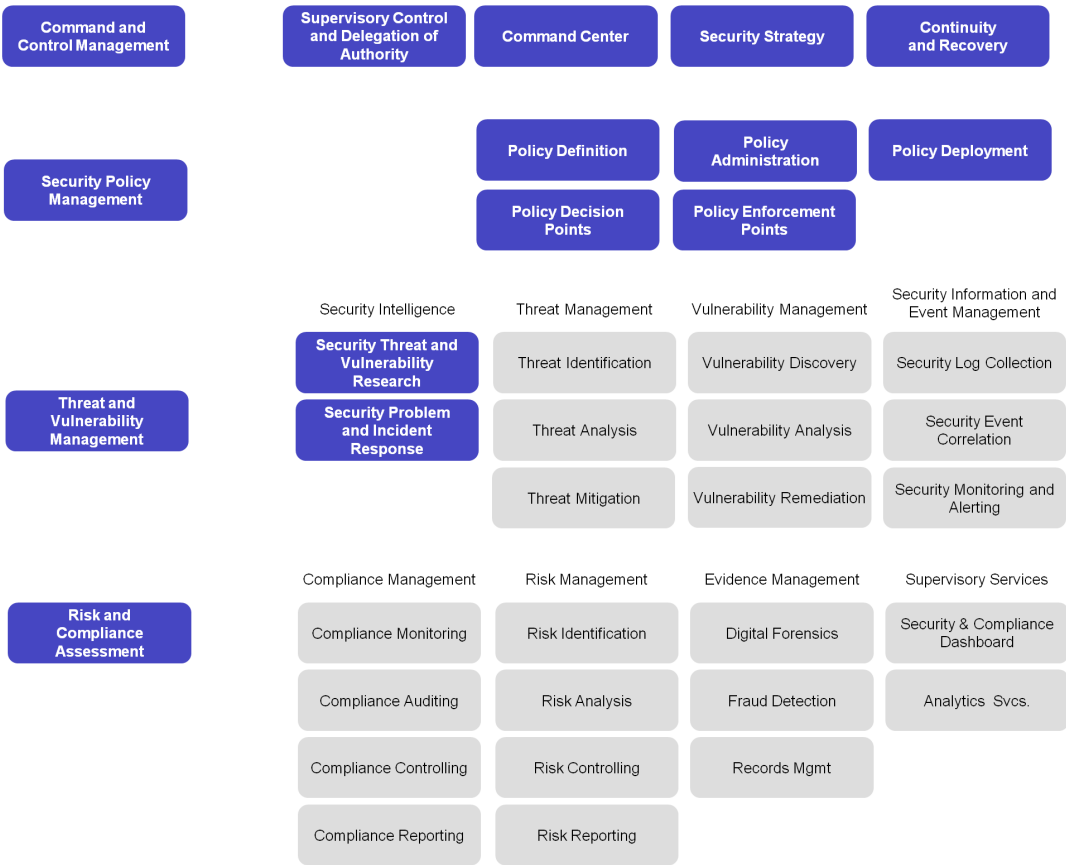
IBM Security Blueprint solution pattern

To understand how the security capabilities of the Deployment and Migration Services can be mapped to the IBM Security Blueprint³, see Figure 11-5 on page 381. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the Deployment and Migration Services. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-5 on page 381 can be used as a quick reference of the functional security management aspects of the Deployment and Migration Services. This reference can help us determine which functions of a solution can be covered by selecting this product.

³ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

Foundational Security Management Component and Sub-Components



Security Services and Infrastructure

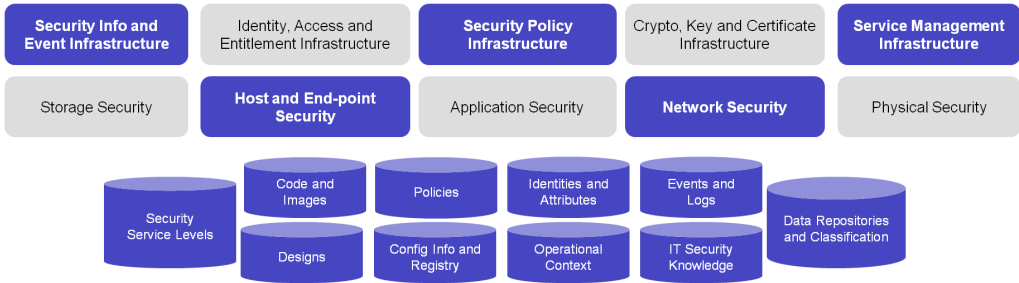


Figure 11-5 IBM Security Blueprint solution pattern for the Deployment and Migration Services

11.1.4 Staff Augmentation Services

Let us examine this service offering in more detail.

What are Staff Augmentation Services

Staff Augmentation Services are designed to help organizations focus the efforts of their internal resources on maintaining their normal business operations while external specialist consultants handle designated security tasks on their behalf. Augmentation services are highly customized to fit the unique needs of the organization they are intended to serve.

Organization's that make use of Staff Augmentation Services usually do so because they are faced with situations that include the following:

- ▶ IT initiatives for security that are planned or underway, but the organization does not have the required capacity or skills to support the initiatives.
- ▶ Security staffing needs are not clearly defined, but the organization could benefit from the addition of experienced security resources to carry out critical security management activities in accordance with industry best practices, standards, and regulations.
- ▶ The organization is in the process of hiring security resources and needs transitional support.
- ▶ The organization does not focus enough on information security to warrant the need for permanently employed security resources, but has need of *ad hoc* security services.

Why are Staff Augmentation Services effective

Staff Augmentation Services provide organizations the ability to use security experts' skills to assist with a variety of security staffing needs. By using proven methodologies, these experts foster a deep understanding of the organization's business environments while establishing seamless relationships with internal IT staff members.

Consultants provide on demand security expertise that helps reduce risk and demands on internal staff, thus ensuring optimal resource use in providing cost-effective and efficient security operations support. This ultimately leads to an improved security posture and reduced complexity by applying expert knowledge to maximize the value of security technology.

Requirements for a successful Staff Augmentation Services delivery

The delivery of successful Staff Augmentation Services relies on how well the organization acquiring the services understands their information security requirements and operations.

There are a multitude of ways in which the requirements can be determined, such as Security Risk Assessments, Information Security Assessments, and Strategy and Architecture consulting engagements. These activities typically look at both business and technical requirements with associated risk ratings, gaps, and mitigation plans that all help with determining the need for additional staff requirements.

After an organization understands all of its security requirements, it is able to focus the required experts' skills in the appropriate areas, whether business or technical focused. This in turn ensures the most effective and efficient use of augmentation services.

Staff Augmentation Services delivery process

As Staff Augmentation Services are generally *ad hoc* services, there is no set delivery process. There is, however, still the need to timeously identify appropriate resources to address specific requirements, so it is therefore vital to ensure that the required resources are available to deliver the needed services in the required time frames.

IBM Security Blueprint solution pattern

To understand how the security capabilities of the Staff Augmentation Services can be mapped to the IBM Security Blueprint⁴, see Figure 11-6 on page 384. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the Staff Augmentation Services. This functional highlighting is applicable for the infrastructure service components as well.

Because the Staff Augmentation Services depend on each individual customer situation we have highlighted every functional component in this solution pattern. The intention here is to convey the message that this service is highly customizable and can be used to address any challenge an organization faces.

⁴ For a detailed discussion of the elements, refer to Chapter 2, "The components of the IBM Security Blueprint" on page 31 and Chapter 3, "The Network, Server and Endpoint solution pattern" on page 93.

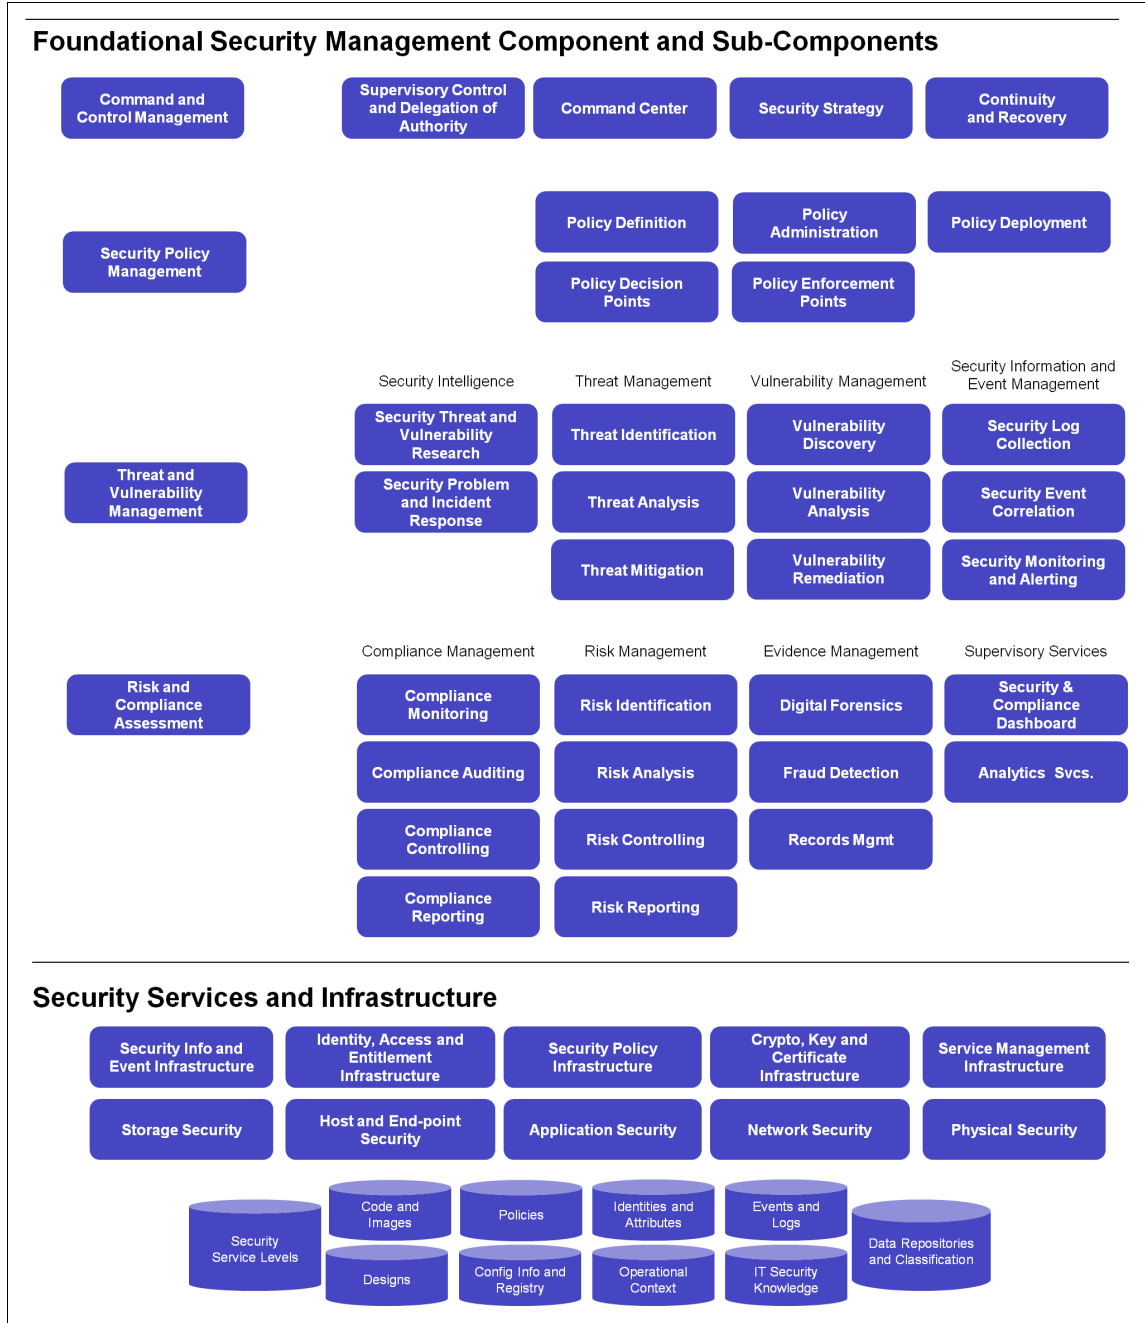


Figure 11-6 IBM Security Blueprint solution pattern for the Staff Augmentation Services

11.1.5 Emergency Response Services

Let us examine this service offering in more detail.

What are Emergency Response Services

Unforeseen security breaches can have devastating consequences, such as data corruption, identity theft, brand damage, and loss of business for organizations. Emergency Response Services are designed to help organizations prepare for, manage, and respond to incidents quickly and effectively.

Emergency Response Services give organizations the ability to retain expert emergency response security consultants prior to incident occurrence, thereby ensuring incident response and management, data preservation, and in-depth data analysis through 24x7 availability. This can help reduce damages and ensure systems and services are restored quickly while maintaining data integrity.

Why are Emergency Response Services effective

Emergency Response Services provide organizations with experienced emergency response security consultants capable of quickly reacting to breaches, thereby reducing the complexity of remediation and ensuring speedy recovery to reduce damages and restore systems and services. Downtime for critical systems such as e-commerce solutions are significantly reduced while protecting internal and external communications and maintaining data integrity. Immediate services with a response time of less than 24 hours ensures that attacks in progress are stopped, thereby reducing their impact and enabling sufficiently detailed data analysis to be performed.

By effectively and efficiently executing Emergency Response Services, organizations are assured of significant cost savings by reducing the impact of breaches and prolonged business disruptions while providing a unique way to gain a wider knowledge base and experience in different environments.

Requirements for successful Emergency Response Services delivery

Due to the *ad hoc* nature of these services, the organization has to ensure that all contractual obligations towards the service provider are maintained at all times. This ensures that the security consultants responsible for delivering the Emergency Response Services stay abreast of developments within the organization's environment and are able to rapidly respond to all reported incidents. In addition, it is vital that the Emergency Response Services consist of comprehensive incident response components that can be rapidly executed.

The service must be designed to cater for the unique requirements of the organization and ensure that it is able to maintain the continuous availability of business systems and reduce risks typically associated with breaches.

Emergency Response Services delivery process

Emergency response services are, as mentioned before, *ad hoc* services that are executed when a breach occurs. That being said, the security team responsible for the delivery of the services will make use of predetermined processes and methodologies to perform all emergency response related tasks.

These tasks typically includes the following phases with associated tasks:

Phase one: Identification

- ▶ Assign an incident owner.
- ▶ Verify that an incident has occurred.
- ▶ Ensure that a provable chain of custody is maintained.
- ▶ Ensure the coordination of infrastructure and operations services.
- ▶ Execute the communications plan.

Phase two: Containment

- ▶ Deploy an onsite team.
- ▶ Ensure that discretionary action is taken.
- ▶ Ensure the containment of potentially breached systems.
- ▶ Perform the appropriate backups.
- ▶ Determine the risk of containment.
- ▶ Execute the containment technical processes and procedures

Phase three: Eradication

- ▶ Determine the cause of the incident.
- ▶ Harden the infrastructure components as and where needed.
- ▶ Perform a vulnerability analysis.

Phase four: Recovery

- ▶ Restore the affected systems.
- ▶ Validate the integrity of the restored systems.
- ▶ Restore operations.
- ▶ Execute monitoring protocols.

Phase five: Lessons learned and incident analysis

- ▶ Develop the lessons learned report.
- ▶ Develop the incident analysis report.

IBM Security Blueprint solution pattern

To understand how the security capabilities of the Emergency Response Services can be mapped to the IBM Security Blueprint⁵, see Figure 11-7 on page 388. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the Emergency Response Services. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-7 on page 388 can be used as a quick reference of the functional security management aspects of the Emergency Response Services. This reference can help us determine which functions of a solution can be covered by selecting this product.

⁵ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

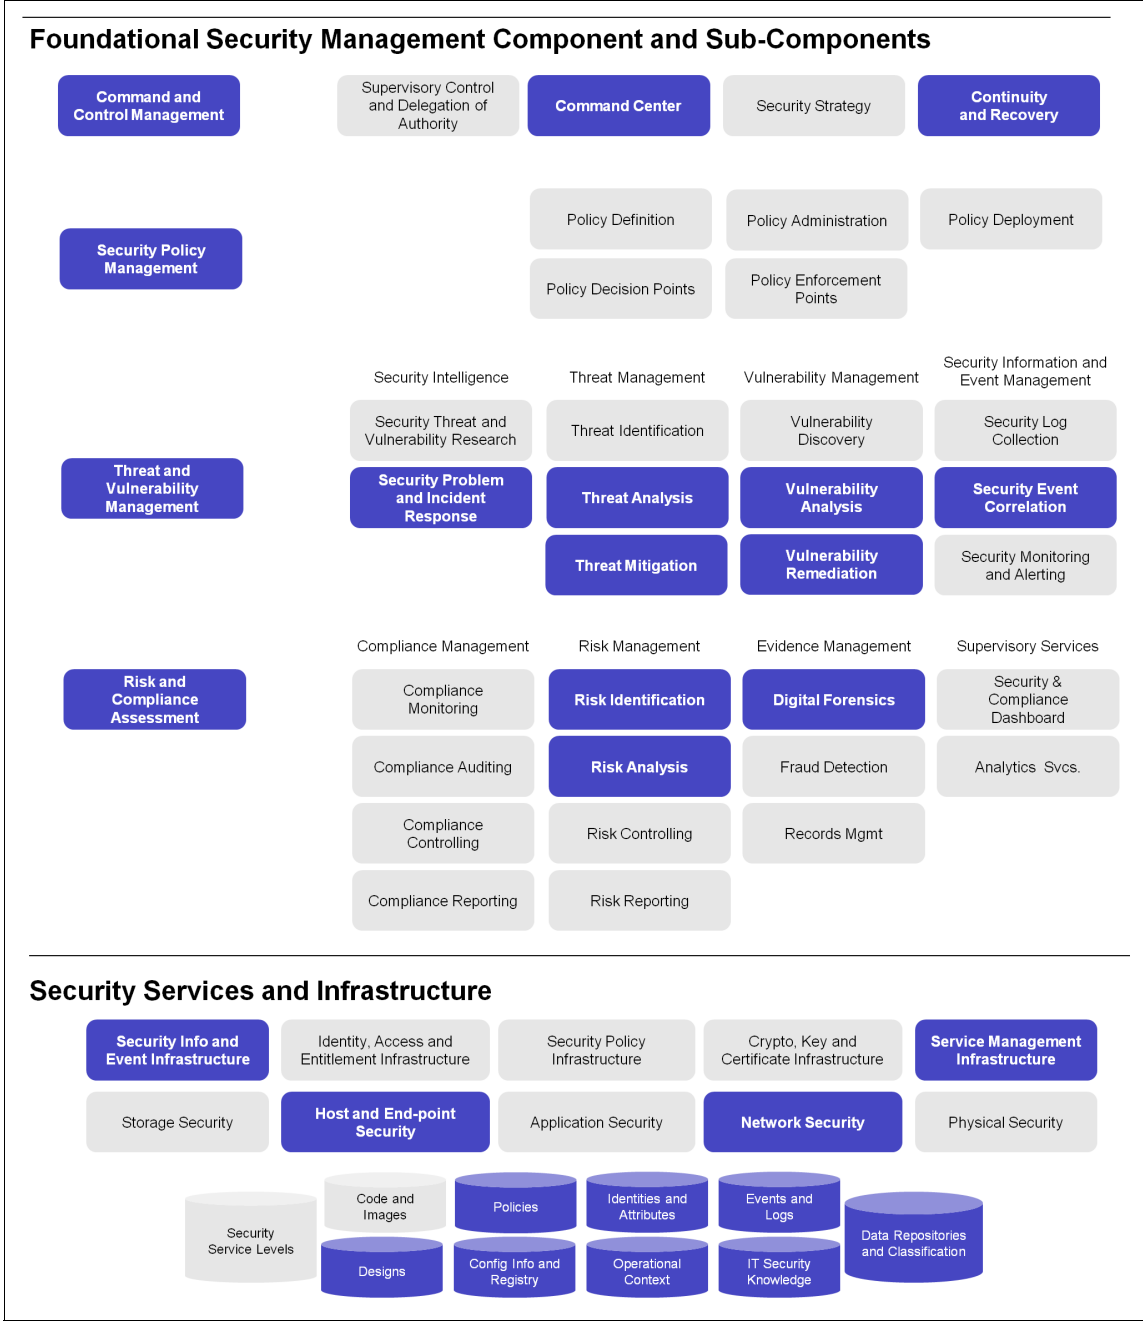


Figure 11-7 IBM Security Blueprint solution pattern for the Emergency Response Services

11.1.6 SCADA Assessment Service

Let us examine this service offering in more detail.

What are SCADA Assessment Services

Process Control Systems (PCS) refer to the overall set of systems that remotely monitor and measure remote sensors from a centralized location. These sensors also typically possess some type of automated response capability when certain criteria are met. A subset of PCS systems that manage systems over large geographic areas are typically referred to as Supervisory Control and Data Acquisition (SCADA) systems. SCADA systems make up the critical infrastructure associated with electric utilities, water and sewage treatment plants, and large-scale transportation systems.

Most SCADA systems used by companies today were developed years ago, long before public and private networks or desktop computing were a common part of business operations. As a result, the need to incorporate security measures in these systems was not anticipated. At the time, good security for SCADA systems meant limiting and securing the physical access to the network and the consoles that controlled the systems. Engineers rationalized that if the systems were suitably isolated from any physical entryways, and if access was limited to authorized personnel only, the systems were fully secure and unlikely to be compromised. This is no longer the case.

SCADA Assessment Services provide customers with a team of highly qualified security consultants to ensure a comprehensive approach to analyzing the vulnerabilities of PCS and SCADA systems. The consultants include Certified SCADA Security Architects who have extensive experience conducting these assessments in the field. A consulting team will coordinate the SCADA and Process Control Systems Assessment closely with customers to ensure an efficient use of resources and minimal impact on personnel.

Why are SCADA Assessment Services effective

A good approach to protect these critical systems is to use existing security best practices with SCADA priorities in mind. This action requires the use of both administrative and technical controls and a Defense In-Depth strategy that encompasses vulnerability assessment, threat prevention, and policy. A Defense in-Depth approach recognizes that the information security problem is not just about IT, but also involves people, process, and technology.

The basic process for securing SCADA environments is made up of four steps:

- ▶ Establishing security
- ▶ Educating employees
- ▶ Enforcing policy
- ▶ Evaluating results

When selecting a security provider to secure a SCADA environment, several factors must be considered; primary among them is SCADA system expertise. Security consultants should be Certified SCADA Security Architects with real-world experience in the field working with these critical, real-time processing systems. In addition, it is helpful to work with a security provider that uses the SCADA Data Dictionary, which can provide organizations with normalized data on threats targeting SCADA systems.

SCADA Assessment Services delivery process

SCADA Assessments Services typically include the following tasks:

- ▶ Network Security Architecture Review
The purpose of this activity is to identify gaps within a customer's network environment as they relate to security best practices associated with logical placement of devices, secure device configurations, remote access configurations, and interconnectivity with other networks.
- ▶ Security Policy and Procedure Review
The purpose of this activity is to determine gaps within a customer's environment as it relates to security policies and procedures based on the industry accepted SCADA policies and the ISO 27002 framework.
- ▶ SCADA Vulnerability Analysis
The purpose of this activity is to identify and validate potential vulnerabilities that exist on a customer's core systems within the SCADA control center.
- ▶ SCADA Operations Network Analysis
The purpose of this activity is to assess the security of the SCADA operations network to identify security issues associated with the SCADA network. This activity may seem similar to Network Security Architecture Review, but is different in that it involves review and technical analysis of the SCADA specific portion of the network.
- ▶ Wireless Assessment
The purpose of this activity is to assess the effectiveness of security controls that have been implemented as part of a customer's wireless network design.

► Physical Assessment Review

The purpose of this activity is to test the physical security of a customer's locations.

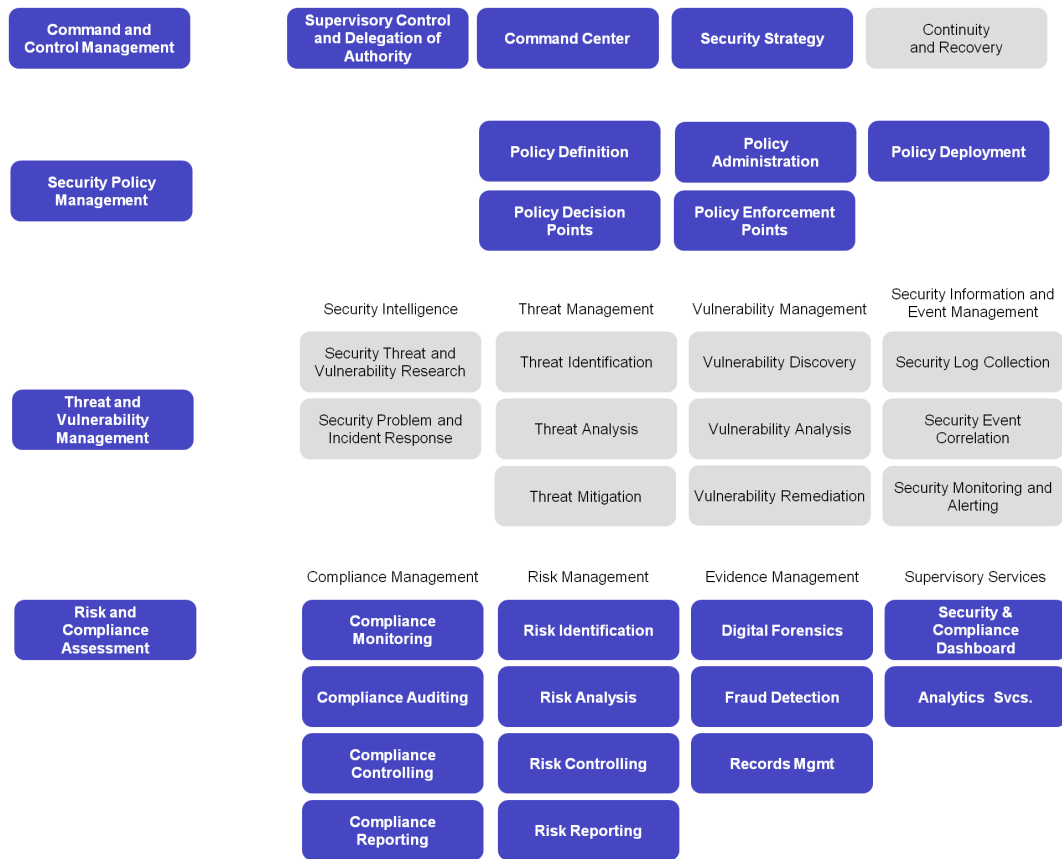
IBM Security Blueprint solution pattern

To understand how the security capabilities of the SCADA Assessment Services can be mapped to the IBM Security Blueprint⁶, see Figure 11-8 on page 392. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the SCADA Assessment Services. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-8 on page 392 can be used as a quick reference of the functional security management aspects of the SCADA Assessment Services. This reference can help us determine which functions of a solution can be covered by selecting this product.

⁶ For a detailed discussion of the elements, refer to Chapter 2, "The components of the IBM Security Blueprint" on page 31 and Chapter 3, "The Network, Server and Endpoint solution pattern" on page 93.

Foundational Security Management Component and Sub-Components



Security Services and Infrastructure

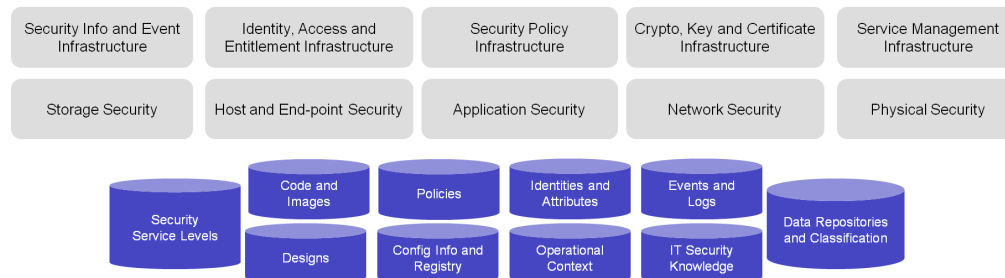


Figure 11-8 IBM Security Blueprint solution pattern for the SCADA Assessment Services

11.2 Managed Security Services

With rapidly changing economic landscapes, it has become increasingly important for organizations to revisit their strategic considerations towards business and IT. More and more pressure is being put on organizations to show value in their spending for security programs, as the proverbial purse-strings are being tightened by finance departments. It is not surprising that, during this strategic reshuffling that most organizations have gone through, they have realized that information security does not lie at the core of their business objectives and drivers, but that it is merely an enabler, albeit a vitally necessary one. This realization, together with the rapidly evolving threat landscape and its associated complexities, has given many organizations reason for exploring ways to reduce cost, while at the same time (at the least) maintaining, or possibly increase, their information security capability.

For most organizations, security has traditionally been an area upon which money is grudgingly spent. This perception has changed with the commoditization of many traditional security solutions and services, such as antivirus solutions and penetration testing services, over the last few years, as most of these commoditized solutions and services now form part of the day-to-day operational functions of most organizations. The one consistent factor that has remained is that it is still difficult to show the value of security solutions and services coupled with the cost associated with hiring, developing, and retaining the required skills to ensure the effective and efficient management of all of the necessary security solutions that we find in most corporate environments today. The introduction of new legislation, regulations, standards, and technologies on an increasingly more regular basis also means that most organizations end up spending a small fortune on training in a futile attempt to keep their security personnel abreast of new developments in the security landscape and battling with operational issues while the security personnel are away from the business.

Managed security services are designed to alleviate all the typical pains associated with keeping up to date with the threat landscape and enable organizations to focus on their core business by managing their information security solutions for them.

Through state-of-the-art Security Operations Centers (SOCs), Managed Security Service Providers (MSSPs), with their compliment of experienced security engineers, consultants, analysts and architects, not only manage security infrastructure components, such as firewalls and IPS's, but also monitor security infrastructures to protect data and communications from misuse and attacks.

MSSPs who are considered to be leaders in their field typically have clearly defined roles and responsibilities within their structures to ensure that the most appropriate skills are assigned to the relevant areas in which they are required.

As an industry leading MSSP, IBM Security Services is able to demonstrate the capabilities shown in the following sections.

11.2.1 IBM MSS personnel qualifications

The IBM MSS organization is made up of approximately 225 people responsible for service delivery and support for IBM MSS customers. All IBM MSS staff members are highly skilled, experienced, and certified on all of the technologies that are managed and monitored as part of the IBM MSS offerings. Many of the staff members hold vendor and vendor-independent certifications, such as CCSA, CCSE, CCNA, CCNE, CCIE, CISSP, CISM, and CISA, and have an average of ten years of information security experience. This expertise is invaluable, as it forms the foundation of the daily operation of the Security Operation Centers worldwide, and as such, IBM continues to invest in the education of our MSS staff to retain high quality security professionals while continuing to provide high quality services to our customer base.

11.2.2 IBM MSS architecture

The IBM Security Operations Centers run from a single, custom-built, and back-end infrastructure capable of delivering all MSS services effectively and efficiently.

In this section, we discuss the IBM MSS architecture components (Figure 11-9) to assist organizations in their decision making process where MSS is concerned.

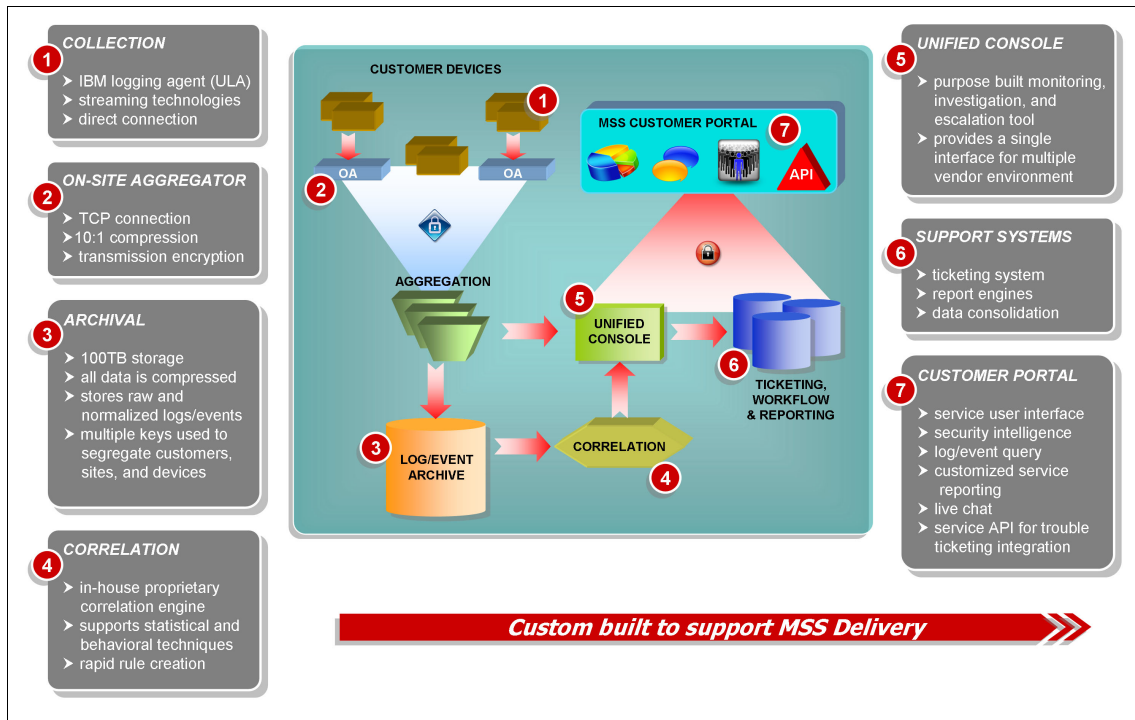


Figure 11-9 IBM MSS architecture overview

IBM MSS architecture overview

The IBM MSS architecture is designed to ensure that potential threats are identified in the quickest time possible by facilitating rapid uptake of security log and event information from a myriad of sources. This, in turn, ensures rapid response to potential threats before they impact customer environments negatively.

The architecture consists of the following items:

1. Collection

The IBM logging agent (ULA) collects flat-file and system information. This built-in watchdog ensures proper functionality while being automatically updated by the IBM SOC. It makes use of unidirectional transmission buffering to prevent data loss message throttling.

2. Onsite aggregator

The onsite aggregator and infrastructure aggregation collects events from the IBM ULA and other streaming technologies. The built-in watchdog ensures proper functionality while being automatically updated by the IBM SOC. In addition, it offers unidirectional transmission, buffering to prevent data loss, and message throttling and transmission windows.

3. Archival

The SOC archive is a proprietary IBM developed storage infrastructure. It is a hierarchical data storage solution using three primary-keys for access control and authentication. It stores both raw and normalized logs and events and supports the IBM MSS customer portal for log and event queries and event monitoring, offering full text indexing and fault tolerance, and is DR mirrored. Customer log and event storage is provided for up to seven years.

4. Correlation

The correlation engine is in-house proprietary technology that supports statistical and behavioral techniques and rapid rule creation. The controller provides management of event streams, converts streams into discrete correlation tasks, and ensures sequential processing, rebalancing, and completion. Correlation drones perform correlation tasks, and use external data stores and a custom correlation database to maintain the state of event stream processing.

5. Unified Console

The Unified Console is a purpose-built monitoring, investigation, and escalation tool that provides a single interface for a complex multivendor environment.

6. Support systems

The SOC support systems consist of the ticketing system, report engines, and data consolidation mechanisms. Automated escalations (XPS Alerts) are made available in the customer portal while additional, optional email escalations can be configured. SOC escalation priorities range from actions and events documentation in the portal to email and phone calls. Analyst investigations are available in the customer portal.

7. Customer portal

The customer portal is the user interface used by IBM MSS customers to view and manage their IBM MSS services and all associated information. It offers a unique blend of security intelligence services, log and event queries, customized service reporting, live chat, and an API for trouble ticketing integration. Using the tools provided in the services layer, the portal user interface provides customers with real-time views of service data and methods of communication to the SOC.

The portal services layer provides a single interface between the data management layer and both internal and external programs. Services supported here are secured through access control and encryption. The data management layer is composed of the same infrastructure components that are used for service delivery and it provides instant views and access to the current status and configuration of the customer's security landscape.

11.2.3 IBM Security Operations Centers

IBM operates nine Security Operations Centers in seven countries and nine Security Research Centers in six countries around the world. The SOC's manage and monitor more than 30,000 devices for 3,700 customers in 133 countries. This ensures 24x7 coverage and access to international expertise for IBM MSS customers. Figure 11-10 shows the IBM SOC worldwide distribution scheme.



Figure 11-10 IBM Security Operations and Research Centers

SOC infrastructure

The SOC infrastructure must be in place and maintained constantly, especially the security surrounding the locations.

Physical security

The physical security is as important for an SOC as it is for any business environment today. Security guards do SOC visitor identity screening. Proper identification must be presented, for example, drivers license, passport, and so on, for a visitor badge to be obtained.

CCTV monitoring is part of the security, as are biometric (handprint) access devices to the SOC itself. Procedures are in place for intruder response. Employee background check is an integrated part of the new employee hiring procedure.

Facilities and infrastructure protection

The facilities and the infrastructure are protected and secured to allow for multiple failures at the same time, while still having the ability to operate. Facilities are below ground. Redundant network circuits, HVAC, and power protection are a given. To ensure an uninterrupted power supply, diesel backup power generation is present at each SOC, while still having contingency mobile power supply and fuel.

Fire protection, a pre-action sprinkler system, and FM-200 gas is part of the overall protection precautions, as is dual-zone smoke/heat detection, all supported by a 24x7 onsite fire response team.

Policies, procedures, and programs

The SOC is governed by a comprehensive set of plans and policies. These plans and programs are developed over the past years and are revised on an annual basis.

The policies, procedures, and programs include:

- ▶ Security and privacy policies
- ▶ Disaster recovery plan
- ▶ Data backup program
- ▶ Security incident response team (SIRT)
- ▶ Change control program
- ▶ Penetration testing program
- ▶ Independent compliance audits

Network and computer security

The security infrastructure is a *must-have* in an SOC. IBM ensures a secure SOC environment to house all customer information. The SOC consists of the following security infrastructure:

- ▶ Network
 - Internet perimeter firewalls
 - Intrusion detection/prevention systems
 - 24x7 intrusion monitoring
 - URL blocking
 - SecureID authentication
 - Virus protection
- ▶ Computer
 - Password protection
 - Virus protection
 - OS and database scanning
 - Secure access placement
 - Audit trail

SOC disaster recovery

The IBM SOC's are located in low-risk threat zones and operations are placed within disaster resistant or hardened facilities. The design of the facilities and structural hardening is commensurate with the proximal risk profile where the facilities are located. Accepting the fact that no geographical location is risk free, facilities are located in areas where the Annualized Loss Expectancy (ALE) is acceptably low. Known risks can subsequently be managed or mitigated through the application of generally accepted risk management and disaster recovery strategies.

We have also adopted the basic precept of risk management; geographically distributing assets to reduce the loss potential. In other words, *never leave all your eggs in one basket*. To that end, our operations are physically and geographically separated among multiple locations. These facilities can back up one another, which reduces the hardening requirements and cost of any single facility.

For power redundancy, we ensure that the facilities infrastructure design is commensurate with the level of protection of the facility itself. The SOC infrastructure includes the concept of N+1 (needed plus one) from the Uptime Institute.⁷ The N+1 strategy ensures that the failure of any single infrastructure system does not adversely affect the whole of the operation.

⁷ For more information about the Uptime Institute, go to <http://www.uptimeinstitute.org/>.

Each SOC has multiple power feeds, duplicate switch gear, UPS, and backup power generation. Additionally, the SOC facilities can accept remote power connectivity in the event of either catastrophic power equipment failures, extended periods of backup power generation (or both), when primary power is lost.

Network transport in our SOC's uses multiple Internet connections. A circuit, typically a DS3 running at 10 Mbps is fully redundant and load balanced with the primary circuit. In each case, the circuits are provisioned from alternate carriers. Multiple Internet connections support the SOC's in a number of ways: First, a second Internet connection ensures that the SOC's stay connected to the Internet. In a bandwidth critical role such as MSS, it is vital that connectivity to the Internet is never lost. This strategy overcomes any number of problems at either the ISP or the telco. With a second connection, our SOC's are still able to service clients in the event of a single ISP outage.

Second, having multiple connections enables us to perform an expanded range of preventative maintenance without affecting production. It also allows clients to have more confidence in the ability of IBM to deliver on our commitment of 24x7 uptime and helps demonstrate our commitment to redundancy in all crucial areas.

Multiple Internet connections also enable us to have two paths to choose from when communicating with clients, as well as offering the client two paths to reach an SOC. This ensures that the quickest path is chosen, reducing latency and improving response time. Our two service providers, AT&T and UUNet, have large and diverse backbones, maximizing the likelihood that our clients are close to either carrier and minimizing their latency over the Internet. In the event that an SOC does lose one of the two Internet connections, the failover is instantaneous, transparent, and automatic.

Security incident escalation procedures

As shown in Figure 11-11, all data from customer devices is fed through two redundant data analysis paths to ensure both real-time monitoring by our security experts, as well as a *second option* through the use of the *X-Force Protection System* (XPS), a proprietary tool used exclusively for our customers. Information from customer devices is monitored through the IBM Security SiteProtector (management console) instances in real time by security analysts, who can escalate and open tickets based on the triage and analysis process. In parallel, the data is fed through the XPS, which can issue alerts based on predefined rules (event thresholds, deviations from baseline, and so on), and this system independently creates tickets for security analysis and follow-up. Based upon the outcome of analyst investigations, the XPS system *learns* and can modify the alerting for future events of the same type for the same customer or system. All alerting and ticketing is completely transparent and available in the customer portal, including raw and normalized logs and all ticket notes.

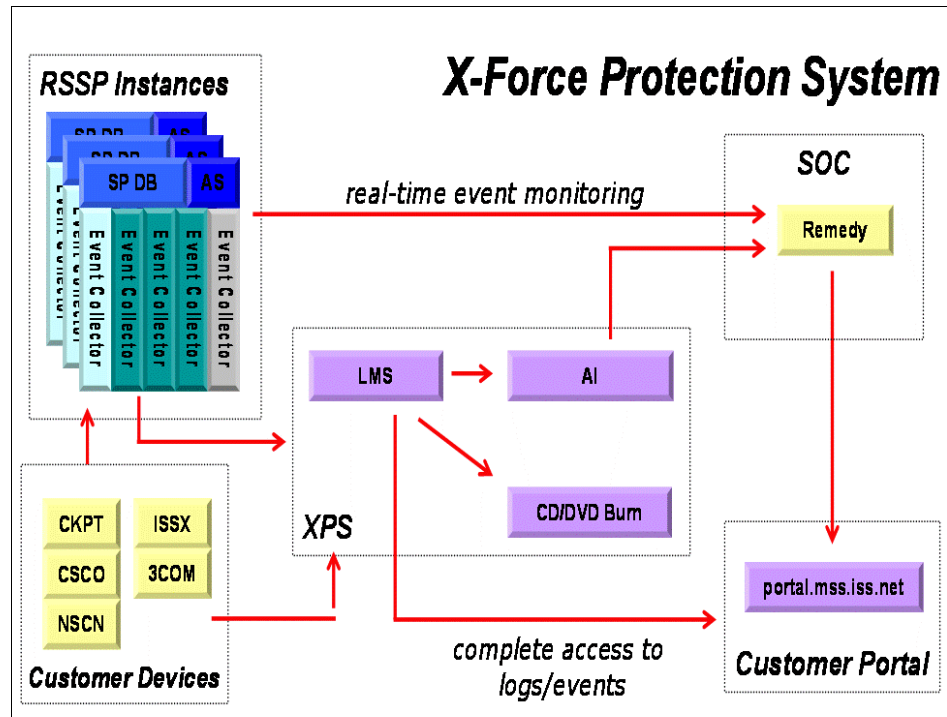


Figure 11-11 X-Force Protection System data flow

IBM is dedicated to providing customers with the highest level of protection services to guard against Internet-based threats and vulnerabilities. As part of those services, highly trained security experts are constantly monitoring and evaluating real-time intrusion event data, and systematically categorizing and classifying each threat.

As shown in Figure 11-12, SOC analysts inspect all incoming events in real time. Leveraging X-Force security knowledge capital, in-depth documentation of customer environments, and state of the art technologies, these events can quickly be evaluated and prioritized. Incident classification is not solely based on the vendor's predetermined event priority, but rather on the SOC analysts' correlation of X-Force security intelligence, global threat awareness, and customer security posture.

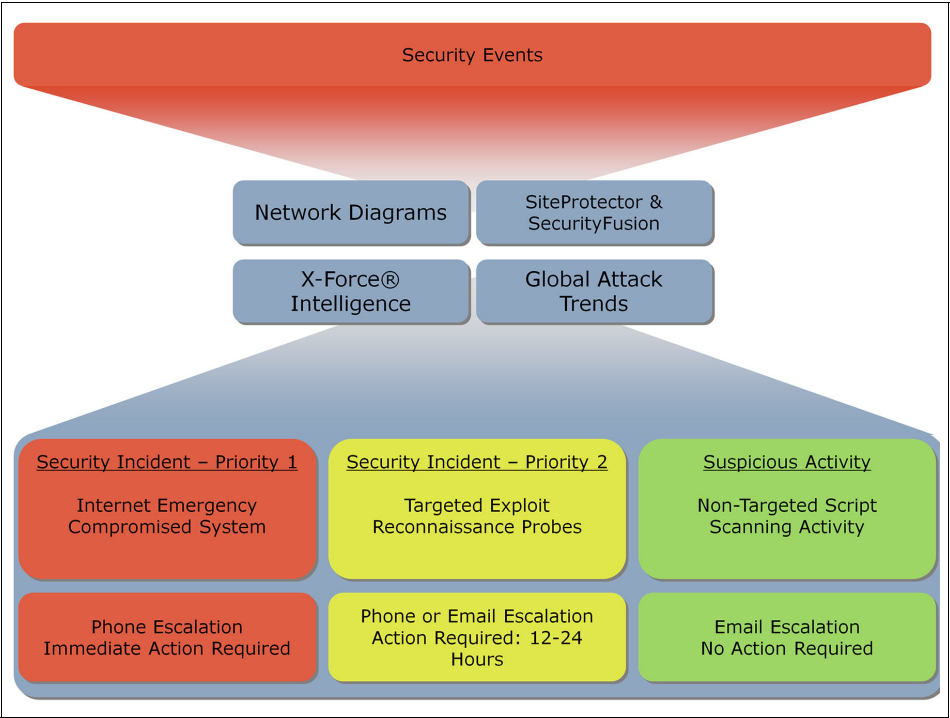


Figure 11-12 SOC Analyst Security Event Investigation

11.2.4 Managed Protection Services

IBM Managed Protection Services offers the industry's only guaranteed protection solution for real-time, 24x7 expert monitoring, management, and escalation across a variety of platforms and operating systems for networks, servers, desktops, and wireless applications. IBM Managed Protection Services customers are able to access real-time service-level data, proactive updates of security content, customized reports, and intelligence through the IBM Virtual-SOC Portal. Figure 11-13 shows the service deployment and integration process.

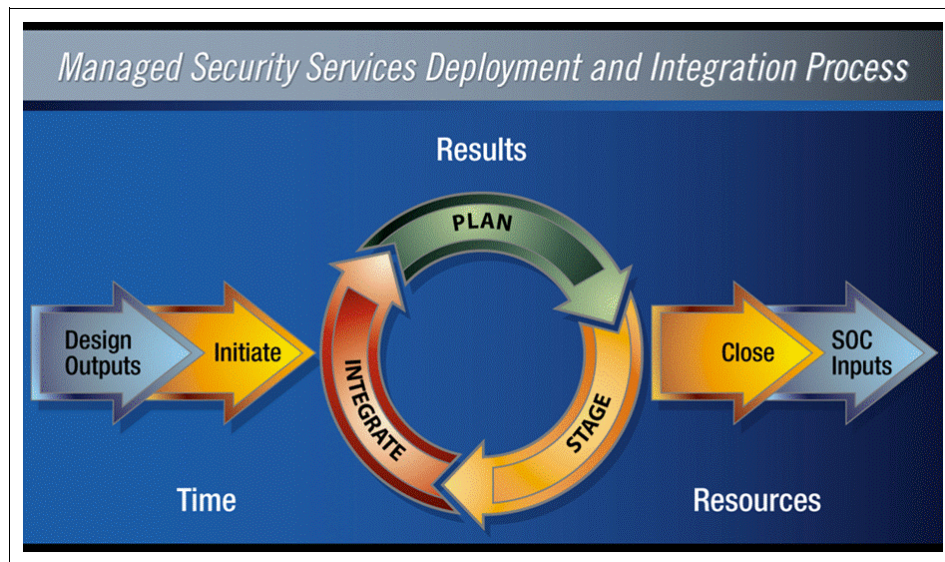


Figure 11-13 Managed Protection Services deployment and integration process

IBM Managed Protection Services are customized to protect your most critical networks, servers, desktops, and wireless applications. If you have less-critical networks, servers, or desktops that require management and reporting, the Managed Protection Services are also available without real-time, 24x7 live event monitoring and escalation, but still include access to an IBM 24x7 team of trained security analysts to assist in resolving problems or answering any service related questions as well as on-demand reporting.

The Managed Protection Services are:

- ▶ Managed Protection Services for Networks

This offering provides around-the-clock protection and expert management and monitoring of the firewall, intrusion prevention, antivirus, antispam, content security, and VPN capabilities of the IBM Security suite of protection appliances.

- ▶ Managed Protection Services for Servers

This offering provides real-time, 24x7 protection and live expert management, monitoring, and escalation for critical server devices across a variety of platforms and operating systems, exactly where you need it, and on the resources you value most.

- ▶ Managed Protection Services for Desktops

This offering provides real-time protection and expert management for desktop environments incorporating IBM market-leading firewall, intrusion prevention, antivirus compliance, spyware prevention, and Buffer Overflow Exploit Protection.

The Managed Protection Services are the only Managed Security Services offering in the industry that offer the following items:

- ▶ Custom policy development and regulatory compliance strategy development.
- ▶ Scheduled penetration testing with quarterly vulnerability scans and assessment security meetings designed to continuously improve security posture.
- ▶ 24x7 expert monitoring, management, incident response and support for networks, servers, and desktops through state-of-the-art, certified, and secure redundant Security Operations Centers (SOC).
- ▶ Real-time response and escalation of unauthorized activities and security events.
- ▶ Virtual Patch protection provides identification and automated patching of vulnerable systems.
- ▶ Around-the-clock access to professionally trained and certified security experts.
- ▶ The state-of-the-art X-Force Protection System and security professionals provide accelerated aggregation, advanced correlation, and event prioritization.
- ▶ Onsite deployment and emergency response workshops.
- ▶ Bundled value-added information security assessments, emergency response, and strategic and technology planning services (available for Premium service level only).

- ▶ Virtual-SOC Portal provides secure, real-time access for all client/Security Operations Center communications, trouble ticket entry, event handling, incident response, data presentation, report generation and trends analysis for all devices under management.
- ▶ Comprehensive reporting through Virtual-SOC Portal provides executive, technical and compliance reporting.
- ▶ IBM X-Force Threat Analysis Service subscription.

11.2.5 Monitored and Managed Firewall Service

Today's complex IT environment presents significant challenges for enterprises that require a high degree of secure network connectivity to efficiently conduct business.

The IBM Monitored and Managed Firewall Service provides real-time security monitoring and management that delivers customized protection at a fraction of the cost of traditional solutions.

This service offers a vendor-neutral approach to maximize your existing security investments while delivering around-the-clock monitoring, management, and analysis of firewall logs. We provide companies of all sizes with the ability to stay ahead of threats while reducing risk and improving regulatory compliance.

The Managed and Monitoring Firewall Services are ideal for organizations that do not have the resources, expertise, or monitoring capabilities to perform required tasks in-house, or who need additional resources to supplement their existing security team.

The IBM Managed Firewall Services can protect company assets with 24x7 expert daily management and maintenance of a variety of certified firewall platforms, ensuring expert monitoring, maintenance, and configuration at a fraction of the cost required in-house.

The IBM Monitored Firewall Service is an optional extension of the IBM Managed Firewall Service, providing 24x7 monitoring of a variety of certified firewall platforms, enabling enhanced detection and cross-correlation of detected attacks features:

- ▶ Access to professionally trained and certified security experts.
- ▶ The state-of-the-art X-Force Protection System provides accelerated aggregation, advanced correlation, and event prioritization.
- ▶ Provides 24x7 expert monitoring and best practices firewall device management through a state-of-the-art, certified, and secure-redundant SOC.

- ▶ Virtual-SOC Portal provides secure, real-time access for all client/Security Operations Center communications, trouble ticket entry, event handling, incident response, data presentation, report generation, and trends analysis for all devices under management.
- ▶ Customized planning, design, and configuration of contracted firewall devices.
- ▶ Ongoing maintenance, including system patches, upgrades, and security content updates.
- ▶ 24x7 proactive monitoring for health and performance of the device.
- ▶ Support for Virtual Private Networks (VPNs).

IBM Security Blueprint solution pattern

To understand how the security capabilities of the IBM Monitored and Managed Firewall Service can be mapped to the IBM Security Blueprint⁸, see Figure 11-14 on page 407. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the IBM Monitored and Managed Firewall Service. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-14 on page 407 can be used as a quick reference of the functional security management aspects of the IBM Monitored and Managed Firewall Service. This reference can help us determine which functions of a solution can be covered by selecting this product.

⁸ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

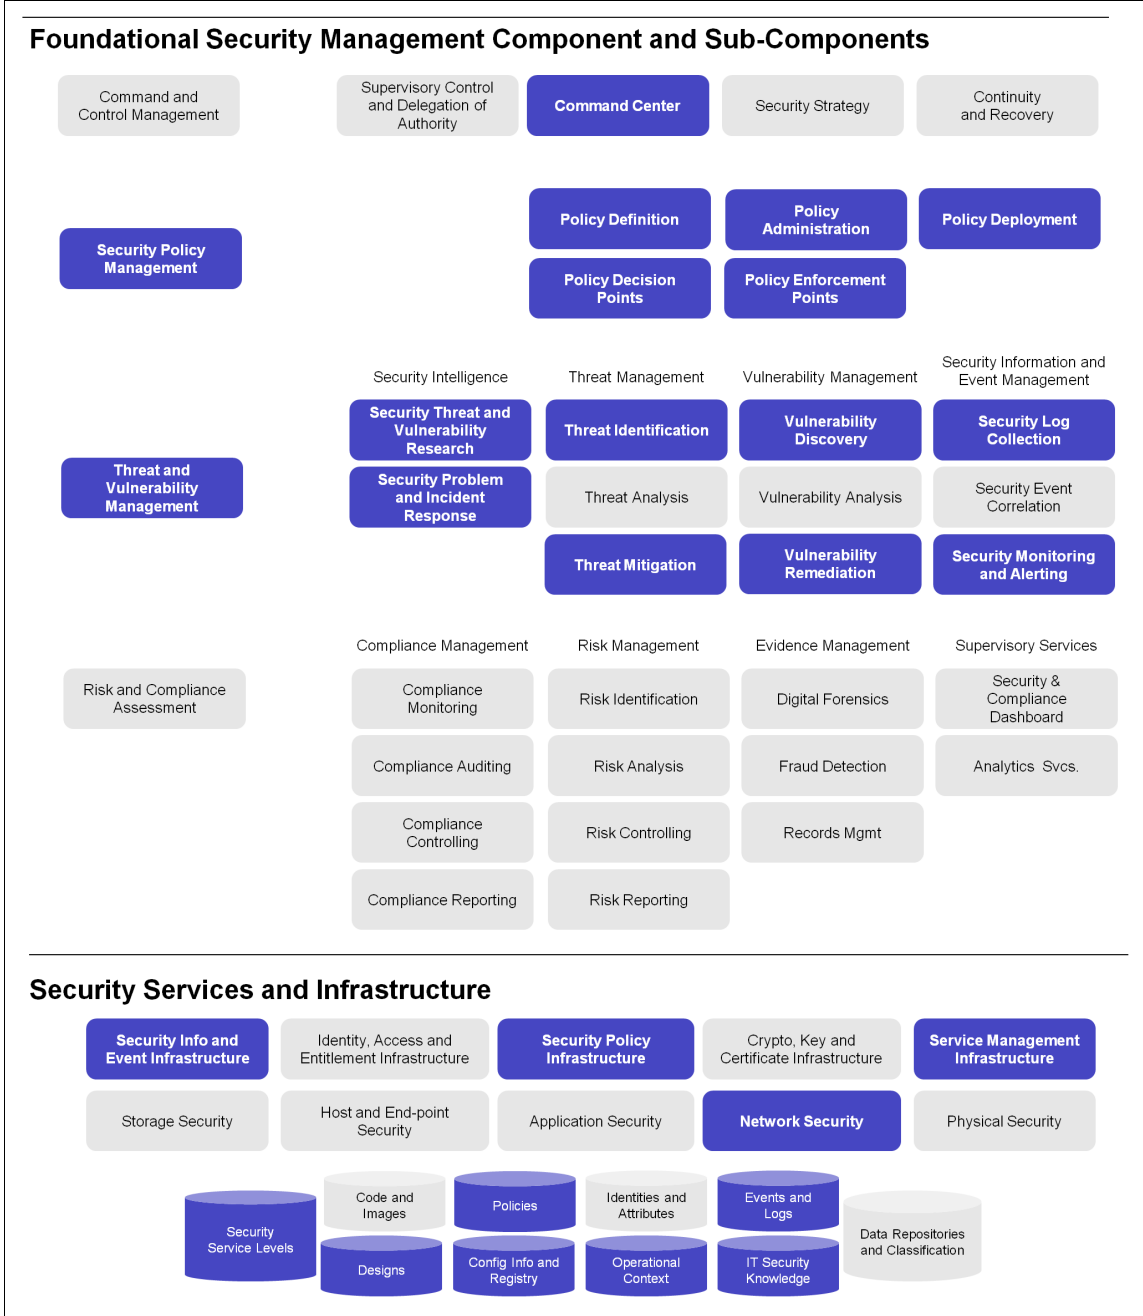


Figure 11-14 IBM Security Blueprint solution pattern for the IBM Monitored and Managed Firewall Service

11.2.6 Managed IDS and IPS Services for network and server

Internet security is now a critical factor in IT performance, impacting everything from business continuity to cost management. Because catastrophic Internet attacks can disrupt your business operations, top security expertise is more valuable and essential than ever before. The IBM Managed Intrusion Detection and Prevention Service (IDS/IPS) provides comprehensive protection for networks and servers against threats and unauthorized intrusions from both internal and external sources. This offering can improve your security posture at a fraction of the cost of traditional solutions.

This service delivers 24x7 expertly monitored and managed intrusion detection and prevention services that keep networks and critical applications protected, enabling companies to save money and reduce the overall security risk. Because these services are vendor-neutral, they are available in a wide variety of options to maximize the existing security investments while reducing risk and improving regulatory compliance.

The IBM comprehensive suite of 24x7 proactive monitoring, intrusion detection, and incident response services protect against inside and outside threats and unauthorized intrusions. The IBM Managed Intrusion Detection and Prevention Service are ideal for organizations that do not have the resources, expertise, or monitoring capabilities to perform required services in-house, or who need additional resources to supplement their existing security team. The IBM Managed Intrusion Detection and Prevention Service are available for protecting networks and servers, and are recommended in conjunction with the IBM Managed Firewall Services and Vulnerability Management Service for an optimal security program. A list of the various services is given below:

- ▶ The Managed IDS/IPS Service for Networks provides 24x7 monitoring, intrusion detection, and incident response services of network segments against unauthorized network intrusions.
- ▶ The Managed IDS/IPS Service for Servers provides 24x7 monitoring, intrusion detection, and incident response services of critical email, web, database, and other types of servers against hacker threats and intrusions, and 24x7 expert monitoring, best practice IDS/IPS device management, incident response, and support for both networks and servers.
- ▶ The services are provided through state-of-the-art, certified, and secure-redundant SOC.
- ▶ Integrated vulnerability intelligence improves attack identification and reduces false positives.
- ▶ Customized configuration and tuning, performance, and fault management of IDS/IPS sensors.

- ▶ Customized prevention baselines when implementing IPS blocking policies.
- ▶ Real-time response and escalation of unauthorized activities and security events.
- ▶ Virtual Patch protection provides identification and automated patching of vulnerable systems.
- ▶ Real-time correlation of vulnerability and attack data prevents the spread of attacks.
- ▶ Around-the-clock access to professionally trained and certified security experts.
- ▶ Uses the state-of-the-art X-Force Protection System to provide accelerated aggregation, advanced correlation, and event prioritization from a broad range of heterogeneous security devices.
- ▶ Virtual-SOC Portal provides secure, real-time access for all client and security operations.
- ▶ Center communications, trouble ticket entry, event handling, incident response, data presentation, report generation, and trends analysis for all devices under management.

IBM Security Blueprint solution pattern

To understand how the security capabilities of the Penetration Testing Service can be mapped to the IBM Security Blueprint⁹, see Figure 11-15 on page 410. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the IBM Managed Intrusion Detection and Prevention Service. This functional highlighting is applicable for the infrastructure service components as well.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-15 on page 410 can be used as a quick reference of the functional security management aspects of the IBM Managed Intrusion Detection and Prevention Service. This reference can help us determine which functions of a solution can be covered by selecting this product.

⁹ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

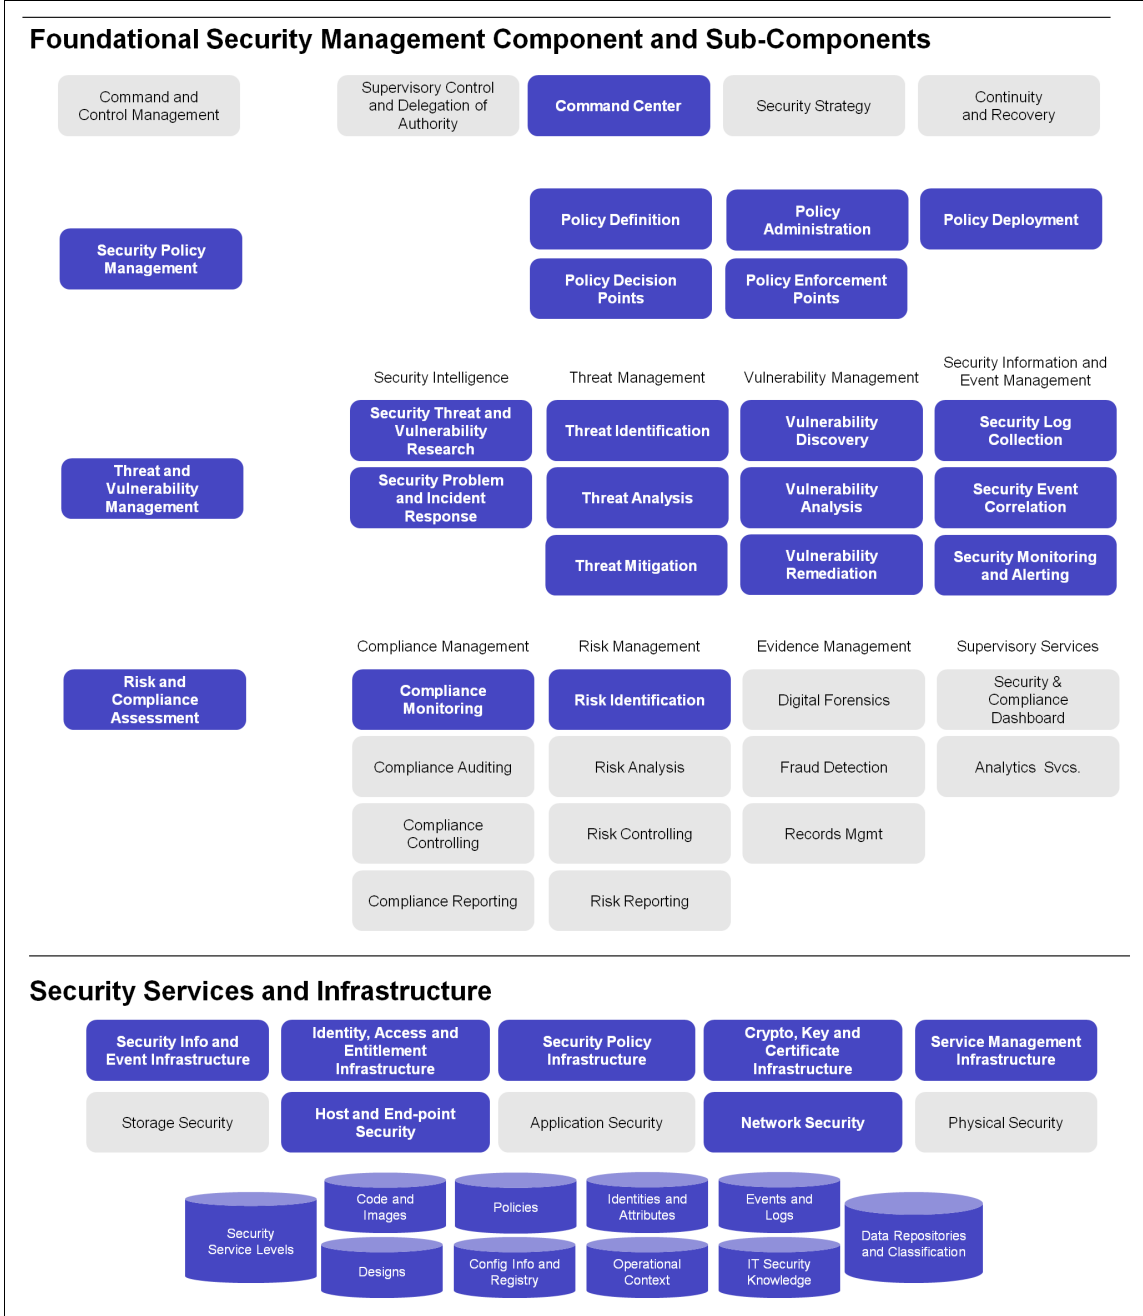


Figure 11-15 IBM Security Blueprint solution pattern for the IBM Managed IDS and IPS Services

11.2.7 Security Event and Log Management Services

The IBM Security Event and Log Management Services (SELM) enable the compilation of the event and log files from network applications, operating systems, and security technologies into one seamless platform. The IBM Security Event and Log Management Services offering allows for automated analysis of IPS data as well as robust query and research capabilities against a variety of disparate log types.

► Key features

- Two tiers of service

The IBM Security Event and Log Management Service is available in *Standard* and *Select* service levels, allowing for varying degrees of analysis and analytics to be applied to varying data types.

- Integrated workflow and analysis capabilities

With the IBM Security Event and Log Management Service integrated workflow and analysis capabilities, security issues can be investigated, escalated, and recorded using IBM web-based tools.

- Seamless blending of MSS and non-MSS data

The IBM Security Event and Log Management Service allows for data of managed as well as unmanaged devices to be stored in the same systems and seamlessly interacted with as though all data is part of a common data set.

- Log collection using an onsite aggregator

Logs are aggregated, encrypted, and compressed prior to transmission to IBM to ensure optimal efficiency and security of sensitive data.

- Forensically sound storage and archival

The IBM Security Event and Log Management Service employs best practice processes for data in motion and at rest as suggested by the IBM Emergency Response Services team.

► Core capabilities

- Streamlined installation and registration

There is automatic registration of IBM Security Event and Log Management Service components within the VSOC portal, eliminating manual registration.

- Enhanced monitoring and troubleshooting

There is proper indication of the health information and installation state of new devices.

- Custom user driven rules
Allows customers to use a simple interface for creating and applying multiple user-defined rules.
- Reporting and compliance
The creation of audit readiness reports are used to help demonstrate compliance based on security information analyzed by the service.

11.2.8 Virtual-SOC Portal

Managing enterprise security can be quite a challenge. With the combination of hackers, regulatory compliance issues, and the difficult task of finding internal resources with the time and expertise to handle enterprise-wide security, it comes as no surprise that many organizations are still being hampered by Internet attacks.

A critical component of partnering with a managed security services provider is its ability to enable customers to remain in control of their security operations and interactively communicate with security experts, as well as show the value of security investments and provide on demand security intelligence.

The IBM Virtual-SOC Portal integrates what an organization needs to know to stay ahead of the threat in an easy-to-use interactive format and offers the tools needed to quantify security investments, refine and deploy security policies, control the security program, and meet compliance requirements.

Let us take a look at the benefits of the IBM Virtual-SOC Portal:

- ▶ Provides access to early warning threat detection to ensure organizations are protected in case of an attack.
- ▶ Complies with federal and state regulations.
- ▶ Refines and deploys security policies to ensure organizations are protected.
- ▶ Quantifies the value of security investments.
- ▶ Prioritizes initiatives within the security program.
- ▶ Reduces risk to the organization.
- ▶ Improves internal security awareness.

The Virtual-SOC Portal includes:

- ▶ Early warning X-Force security intelligence and threat analysis, which details global online threat conditions tailored for specific needs.

- ▶ 24x7 event monitoring, event handling, and security analysis, which synchronizes the management of deployed security devices, networks, and applications.
- ▶ Real-time integration of functionality and research, which provides interactive content, daily assessment and proactive vulnerability notifications, alerts and advisories, detailed trend analysis, and attack metric reporting.
- ▶ A mobile interface (WAP) option, which enables remote access from portable electronic devices, phones, and PDAs.
- ▶ Easy-to-read, business focused reports with configurable views ranging from the enterprise level, work groups, or an individual device level.
- ▶ Supports on demand services, which includes interactive, real-time customer and security operations center communications, security incidents, and ticket data, including:
 - Event handling
 - Help desk requests
 - Trouble ticketing
 - Incident response tracking
 - Event history consolidation
 - Review logs
 - Policy changes
 - Service level agreement (SLA) performance metric reporting
 - Trend analysis
- ▶ Intelligent event correlation and analysis, which matches all security data sets against predetermined alert and response criteria in addition to having the capability to initiate actions to thwart malicious activity.
- ▶ Full-site search capabilities, which quickly assembles security tickets generated for worms and viruses, vulnerabilities, intrusion detection logs, and security news.
- ▶ Periodic scanning and penetration testing with access to tools to scan security devices, networks, and applications to determine vulnerabilities.

The IBM Virtual-SOC is a framework for integrating security tools, services, and intelligence into a consolidated view. It can manage and monitor security operations from a single point. The Virtual-SOC architecture, shown in Figure 11-16, provides a comprehensive view of your security posture combined with actionable security information to help you stay ahead of threats. Using advanced artificial intelligence systems that aggregate security events and network logs, the Virtual-SOC correlates this information against security vulnerabilities. It then produces prioritized actions based on your individual organization's security posture.

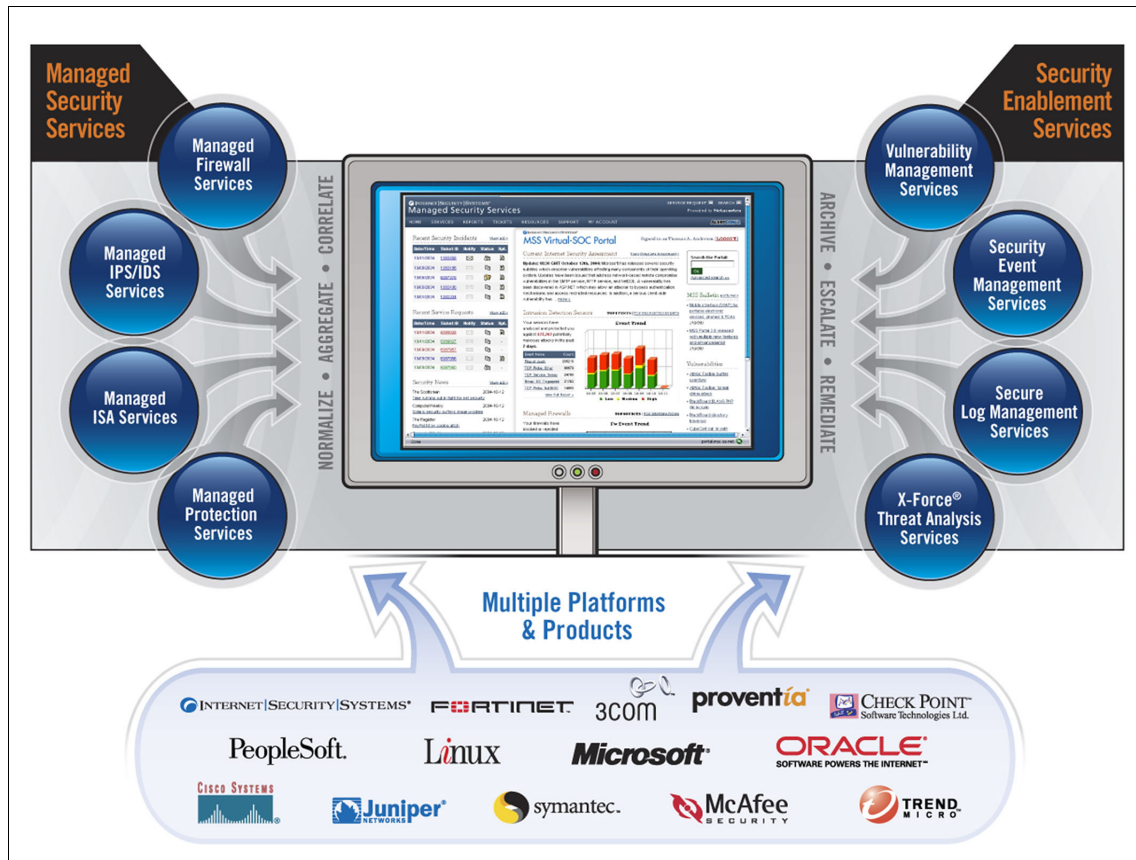


Figure 11-16 Virtual-SOC

Using the IBM Virtual-SOC, you are able to manage your entire security operation, that is, managed and unmanaged devices (both IBM and third-party), security intelligence, reporting, archiving, remediation, escalation, and collaboration with the IBM Managed Security Services analysts.

The Virtual-SOC is accessible anytime and anywhere through the powerful, easy-to-use Virtual-SOC Portal that allows for real-time decision-support.

The Virtual-SOC solution can be customized from an array of managed and monitored security services along with the suite of Security Enablement Services, which provides you with the flexibility to outsource the management and monitoring of certain devices to IBM, while using the Virtual-SOC to monitor other security with in-house resources. The Virtual-SOC further simplifies security management, monitoring and reporting by consolidating and normalizing logs and events across your deployment of multivendor security technologies, whether those technologies are managed by us or by you in-house.

These are packaged solutions that can be easily integrated into your security program to provide the tools and intelligence required to proactively secure the network and maintain regulatory compliance. The Security Enablement Services include the IBM Vulnerability Management Service, Security Event and Log Management Services, and X-Force Threat Analysis Service. IBM is the trusted security partner for many commercial organizations and governments worldwide. The Virtual-SOC architecture is developed using a decade of experience designing, managing, and monitoring thousands of security solutions.

The Virtual-SOC architecture is an extensive network of intelligent systems and processes that enables seamless integration between the Managed Security Services and Security Enablement Services delivered through a secure, web-based portal. This integration can provide IT organizations the intelligence, tools, and capabilities necessary to make real-time decisions when immediate action is required. Some of the most popular features of the Virtual-SOC Portal and architecture are discussed in the following list.

- Open vendor architecture

The Virtual-SOC accommodates a wide variety of best-of-breed IDS, IPS, and firewall technologies from multivendor systems, including products from IBM, Cisco, Check Point, Juniper, 3Com, McAfee, Sun, Microsoft, and other vendors.

- Consolidated security views

An IT organization can monitor and control all Virtual-SOC services through a centralized command center using the Virtual-SOC Portal. Subscribing to any mix of IBM Security Enablement Services or traditional managed services can alleviate your IT staff's struggle to monitor a mixture of multivendor security devices, such as firewalls, IDS, or IPS. The IT organization can view all security events and logs from a single location through the Virtual-SOC Portal. In addition, the IT organization can monitor security events or logs for one device, all devices, or anything in between with easy-to-use filters.

- ▶ Powerful query and reporting options

The Virtual-SOC normalizes all events and logs published to the Virtual-SOC Portal, which enables the IT security organization to run queries and generate reports on any or all security devices, security events, service level agreement (SLA) activity, and many other parameters through a robust query and report engine. This capability greatly reduces the time needed to conduct investigations and identify abuse trends across the enterprise. You can use one of the IBM recommended Virtual-SOC report templates or create your own. You might add a logo or other personalized branding to the reports and tailor them for your organization. All reports can be exported to commonly supported formats, such as CSV, PDF, DOC, and others.

- ▶ Automated event and log analysis

The Virtual-SOC services include automated analysis of security events and logs through our network of intelligent systems within the Virtual-SOC architecture. These services, events, and logs received by IBM are analyzed by expert systems to uncover trends, anomalies, activity spikes, and subtle, under-the-wire attacks. When these systems identify an event or log trend that is indicative of abnormal activity, they generate an alert or ticket that is posted within the Virtual-SOC Portal.

- ▶ Unlimited event and log archive

Many Virtual-SOC services include one year of unlimited online event and log storage accessible through the Virtual-SOC Portal, and seven years of unlimited offline archiving in the forensically sound IBM archival system. We maintain the integrity of all events and logs by storing the original logs in their raw native formats; copies of these original logs are used for normalization, monitoring, or reporting purposes.

- ▶ Granular permissions system

Access to and within the Virtual-SOC is driven by a granular permission system. This can enable an IT organization to determine who can access the portal, what users can see when they are logged in, what they have the right to change, and which users are authorized to contact the IBM SOC's. User permissions can be granted as read or read/write at the device level, site level, division or department level, by service type, and by technology type. These granular permission capabilities enable an IT organization to use the Virtual-SOC Portal as a collaboration tool.

- ▶ Integrated trouble ticketing and workflow

The Virtual-SOC Portal includes a trouble ticketing workflow system by which an IT organization can create, assign, and track ticket status for collaboration within an enterprise. These trouble tickets are not accessible or viewable by IBM analysts. Authorized Virtual-SOC Portal users are able to view their own trouble tickets side by side with the trouble tickets being shared with IBM.

Using this capability, an IT organization can streamline remediation and change control management efforts through their own private tickets. IBM can also provide an application program interface (API) to integrate with common trouble ticketing systems.

11.3 Cloud Security Services

Cloud computing provides flexible, cost-effective delivery of business or consumer IT services over the Internet. Cloud resources can be rapidly deployed and easily scaled, with all processes, applications, and services provisioned on demand, regardless of the user location or device. As a result, cloud computing helps organizations improve service delivery, streamline IT management, and better align IT services with dynamic business requirements. Cloud computing can also simultaneously support core business functions and provide capacity for new and innovative services.

Both public and private cloud models, or a hybrid approach using both models, are now in use. Available to anyone with Internet access, public clouds are acquired as a service and paid for on a per-usage basis or by subscription. Private clouds are owned and used by a single organization. They offer many of the same benefits as public clouds, but give the owner greater flexibility and control.

Although the benefits of cloud computing are clear, so is the need to develop proper security for cloud implementations, whether public or private. Embracing cloud computing without adequate security controls can place the entire IT infrastructure at risk. Cloud computing introduces another level of risk because essential services are often outsourced to a third party, making it harder to maintain data integrity and privacy, support data and service availability, and demonstrate compliance. Even if IT workloads are transitioned to the cloud, users are still responsible for compliance and data security. As a result, subscribers must establish trust relationships with their cloud providers and understand the risk posed by public and private cloud computing environments.

However, cloud computing changes some of the basic expectations and relationships that influence how we assess security and perceive risk. Although the intent of security remains the same, that is, to ensure the confidentiality, integrity, and availability of information, cloud computing shifts control over data and operations in the following ways:

- ▶ **Delivery**

In addition to the usual challenges of developing secure IT systems, cloud computing presents an added level of risk because essential services are often outsourced to a third party. The externalized aspect of outsourcing makes it harder to maintain data integrity and privacy, support data and service availability, and demonstrate compliance.

The cloud shifts much of the control over data and operations from the client organization to their cloud providers, much in the same way organizations entrust part of their IT operations to outsourcing companies. Even basic tasks, such as applying patches and configuring firewalls, can become the responsibility of the cloud service provider, not the user. This means that clients must establish trust relationships with their providers and understand the risk in terms of how these providers implement, deploy, and manage security on their behalf.

- ▶ **Multi-tenancy**

In addition, the massive sharing of infrastructure with cloud computing creates a significant difference between cloud security and security in more traditional IT environments. Users spanning different corporations and trust levels often interact with the same set of computing resources. At the same time, workload balancing, changing service level agreements, and other aspects of today's dynamic IT environments create even more opportunities for misconfiguration, data compromise, and malicious conduct.

- ▶ **Self-service**

Gone are the days when the IT and security staff completely dictate how computing is consumed. IT is becoming more and more self-serve; but who is helping to make sure this is done securely?

- ▶ **Slow-provisioning**

Gone, too, are the days when it took days and weeks, possibly even months, to order hardware, set up resources, and provision applications. With services a few clicks and a credit card away, how can we be sure that control processes, security policies, and proper decisions are made? Will the speed of the cloud put us at risk?

11.3.1 Security challenges in the cloud

One of the most significant differences between cloud security and traditional IT security stems from the sharing of infrastructure on a massive scale. Users spanning different corporations and trust levels often interact with the same set of computing resources. Public cloud services are increasingly being offered by a chain of providers, all storing and processing data externally in multiple unspecified locations.

Inside the cloud, it is difficult to physically locate where data is stored. Security processes that were once visible are now hidden behind layers of abstraction. This lack of visibility can cause concerns about data exposure and compromise, service reliability, ability to demonstrate compliance and meet SLAs, and overall security management.

Visibility can be especially critical when it comes to compliance. The Sarbanes-Oxley Act, the Health Insurance Portability and Accountability Act (HIPAA), European privacy laws, and many other regulations require comprehensive auditing capabilities. Many public clouds may indeed be a black box to the subscriber, so clients may not be able to demonstrate compliance. A private or hybrid cloud, conversely, can be configured to meet those requirements.

In addition, providers are sometimes required to support third-party audits, or support e-discovery initiatives and forensic investigations. This adds even more importance to maintaining proper visibility into the cloud. Legal discovery of a co-tenant's data may affect the confidentiality of other tenants' data if the data is not properly segmented. This may mean that some sensitive data may not be appropriate for certain cloud environments.

Organizations considering cloud-based services must understand the associated risks and ensure appropriate visibility. IBM guidelines for securing cloud implementations focus on the following areas:

- ▶ Building a security program
- ▶ Confidential data protection
- ▶ Implementing strong access control and identity life cycle management
- ▶ Application provisioning and de-provisioning
- ▶ Governance audit management
- ▶ Vulnerability management
- ▶ Testing and validation

Because cloud computing is available in several service models (and hybrids of these models), each of these models presents different levels of responsibility for security management. Trusted third parties can help companies apply cloud security best practices to their specific business needs.

11.3.2 Advantages of cloud-based security

Cloud-based security services can offer the following advantages over traditional security deployments.

- ▶ No expensive, on-premise security hardware to purchase, install, and maintain.
- ▶ No stand-alone software to constantly update and patch.
- ▶ Rapid deployment and self-service through a web-based portal.
- ▶ Ability to scale and expand security coverage quickly, without investing in additional infrastructure.
- ▶ Flexible, service-oriented pricing and service level agreements.

Traditional approaches to security require the purchase of multiple security technologies and management systems, along with manpower for integration, configuration, and patching. Cloud-based services can reduce those expenses and provide an ideal delivery method for many external security functions, including:

- ▶ Vulnerability scanning
- ▶ Web/URL filtering
- ▶ Security event management
- ▶ Security log management
- ▶ Email security

In addition to security functionality, cloud-based security services also provide up-to-the-minute security intelligence and analytics to keep all security technologies patched and up-to-date.

11.3.3 Using cloud-based security services

Cloud-based security services benefit the largest enterprises and small and medium-size businesses. Relative to on-premise security software, cloud security requires little to no upfront capital investment and deployment costs, and the cloud delivery model lowers ongoing operational management costs. It can also satisfy risk management needs across industries, such as government, retail, and manufacturing. In addition to reducing costs and addressing risk management requirements, cloud security services can enable an organization's limited operational resources to direct their efforts towards more strategic initiatives that drive real business value.

Compared to traditional, on premise security implementations, organizations using the cloud spend less time managing systems, troubleshooting technical problems, and responding to the most recent security threat.

11.3.4 Security for the cloud versus security from the cloud

Security services *for the cloud* provide organizations with assistance they need to assess, design, and implement a sound security strategy for cloud computing.

IBM offers security services across the multiple domains of the IBM Security Framework that take into account these unique cloud computing challenges, which includes the IBM Cloud Security Assessment Services, which assess the architecture, practices, and policies of a specific cloud environment, and the IBM Cloud Security Strategy Roadmap, which can help organizations to develop a high level strategy for addressing risks associated with new cloud initiatives. An overview of the IBM Security Services offered for the cloud are shown in Figure 11-17.

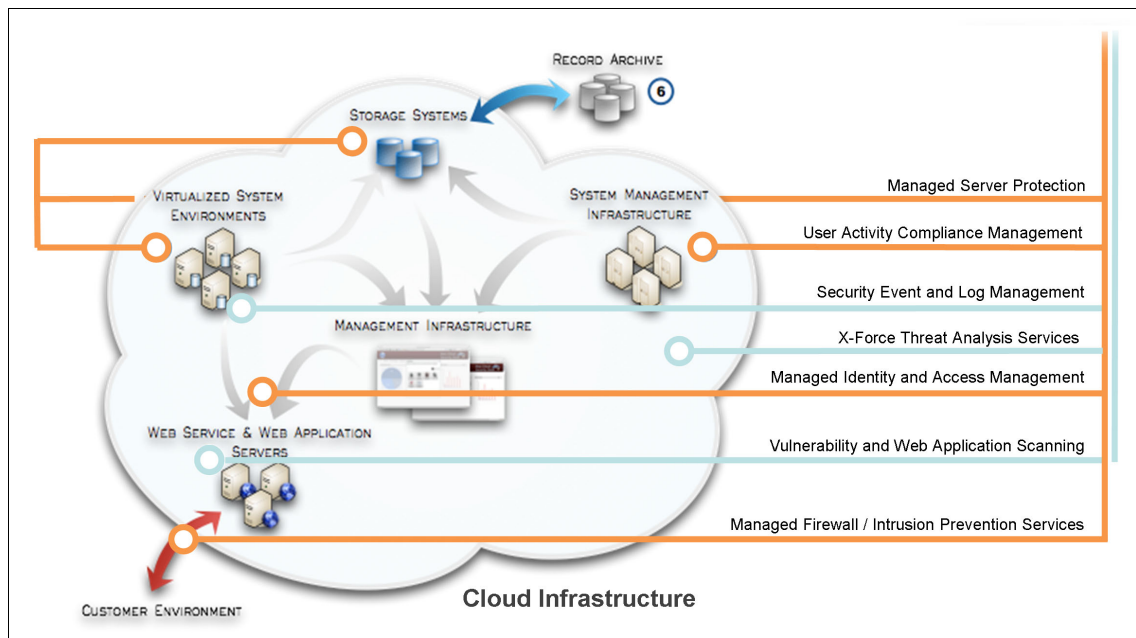


Figure 11-17 IBM Security Services for the cloud

IBM also offers cloud-based security services *from the cloud*. IBM has built a cloud security services portfolio based on security activities organizations are already performing or need to perform for compliance and general security best practices.

The IBM cloud security services *from the cloud* include the Hosted Security Event and Log Management Service (SELM), which enables the compilation of event and log files from network applications, operating systems, and security technologies into one seamless platform. This offering allows for the automated analysis of IPS data and robust query and research capabilities against a variety of disparate log types. It also includes the Hosted Vulnerability Management, which provides 24x7 cloud-delivered vulnerability management scanning services providing vulnerability discovery, prioritization, remediation, dynamic protection, verification, and customizable reporting. This service delivers comprehensive visibility into each area of potential exposure within a distributed network environment.

These cloud-based services allow organizations to use sophisticated security technology without having to purchase, deploy, and maintain expensive software and hardware solutions. These services help reduce the cost and complexity of security management and allow them to use security best practices all from a common web-based portal.

An overview of the IBM Security Services offered *from the cloud* are shown in Figure 11-18.

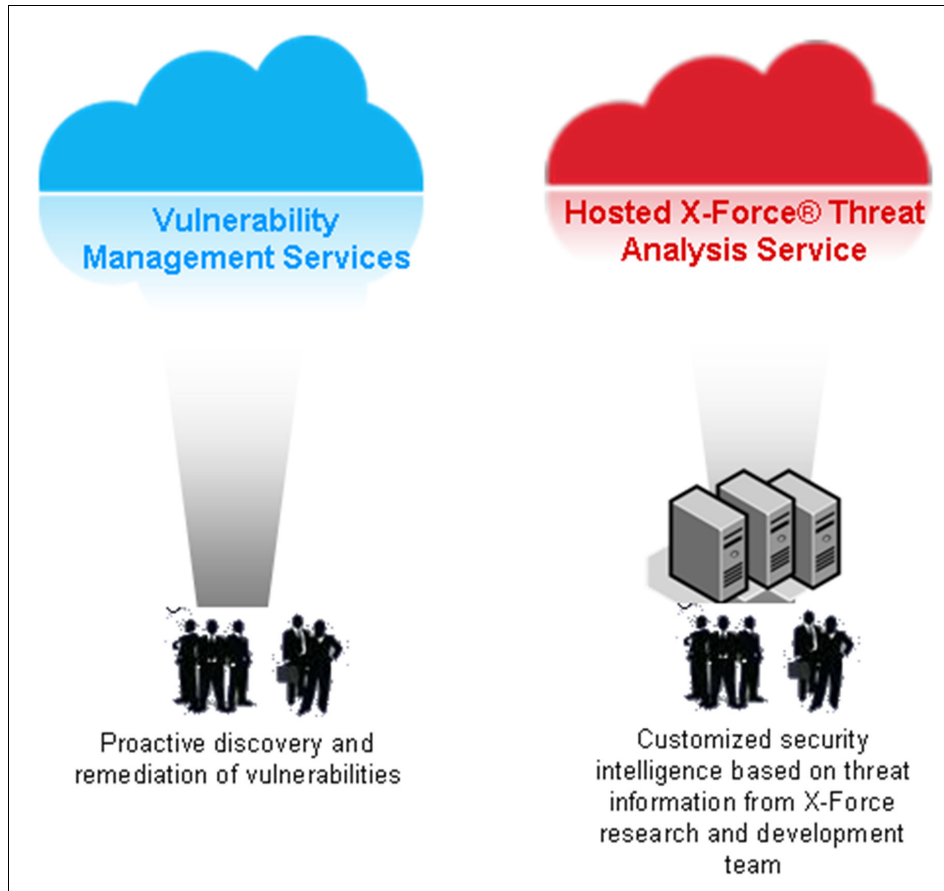


Figure 11-18 IBM Security Services from the cloud

11.3.5 IBM Security Services for the cloud

The IBM Security Services *for the cloud* are designed to help organizations in varying stages of their cloud adoption. The services can be applicable to cloud providers, both public and private, and cloud subscribers, and ultimately help organizations in their pursuit of the many benefits of cloud computing.

Cloud Security Assessment Service

The IBM Cloud Security Assessment Service is designed to address the following tasks:

- ▶ Assess the maturity of cloud solution security controls and mechanisms by comparing against best practices.
- ▶ Develop a ranking of existing security postures in consideration of cloud security goals and identified gaps.
- ▶ Provide specific recommendations for addressing identified issues, and use IBM expertise to provide an actionable plan to close security gaps within cloud initiatives.

IBM Security Consultants conduct a thorough security analysis of current or planned cloud solutions for both private and public cloud environments. The review assesses the effectiveness of cloud security as it relates to identity and access management, data protection, application security, infrastructure protection, physical security and overall security governance and compliance. The solution's security architecture, practices, and policies are reviewed in comparison to industry best practices. The assessment helps identify the strengths and vulnerabilities within existing cloud infrastructure security programs and recommends steps to improve security posture.

After the assessment phase is completed, cloud security professionals compare current state assessment findings against industry best practices. This gap analysis enables consultants to evaluate the overall maturity of existing security programs and identify areas where security objectives are not being met. The analysis provides recommendations to improve the overall security posture of your cloud environment.

IBM Cloud Security Consultants guide customers through the unique challenges associated with securing the cloud environment while using a proven best practices methodology to enable them to quickly identify security gaps and outline remediation recommendations for improvement of security posture.

Cloud Security Strategy Roadmap

IBM security experts guide the customer through an onsite working session to help define the cloud computing initiative and goals, identify associated security and privacy risks, determine appropriate risk mitigation strategies, and develop a high-level security strategy roadmap for achieving cloud security objectives.

The experts help identify and prioritize cloud computing scenarios for specific security requirements and business needs while using IBM expertise to help assess risks and develop a high-level roadmap for risk mitigation.

11.3.6 IBM Security Services from the cloud

Cloud-based security services are delivered remotely to provide security functionality and intelligence from an offsite service provider. These providers can enable organizations to perform routine security activities more efficiently and cost-effectively by using technology and infrastructure provided by a trusted third party. They help organizations contain costs with flexible pricing based on usage. The cloud security service provider takes responsibility for application functionality, deployment, performance, and maintenance, liberating an organization from these burdensome activities.

IBM X-Force Hosted Threat Analysis Service

The IBM X-Force Threat Analysis Service (XFTAS) is a cloud-based service that delivers a unique blend of threat information collected from globally networked security operations centers and trusted security intelligence from the IBM Security X-Force Research and Development Organization (X-Force). The service provides daily summaries that clearly denote the nature and severity of threats and vulnerabilities, and include links to recommended fixes and security advice.

The IBM X-Force Threat Analysis Service delivers customized information about a wide array of threats that can affect network security. The service can help organizations to proactively protect networks with detailed and customized analyses of global online threat conditions. The IBM X-Force Threat Analysis Service combines high-quality, real-time threat information from the IBM international network of Security Operations Centers with security intelligence from the renowned X-Force team.

Many organizations lack the time or resources to constantly monitor, research, and analyze security information. Administrators are often forced to react to threats instead of actively preparing for them. The difficulty of sorting through the myriad security issues of the day often causes companies to be blindsided by truly important issues. Without timely and relevant security information, organizations spend resources reacting to threats (and often times disrupting the business) instead of preparing for them.

The IBM X-Force Threat Analysis Service analyzes global security issues to deliver tailored, timely, and relevant information to you so that you can take decisive, proactive measures to increase your organization's security.

The service provides the following benefits:

- ▶ A single source for trusted, expert analysis of global security threats.
- ▶ Actionable data and recommendations that help you maintain your network security.

- ▶ Customized information to provide insight into the issues that pertain to an organization.
- ▶ Timely notification to keep you aware of critical issues.
- ▶ Threat forecasts to help you prepare for increased due diligence.

The X-Force is the foundation of the preemptive approach to Internet security pioneered by Internet Security Systems. Started in 1997 by ISS founder Chris Klaus, the X-Force is the oldest, best-known commercial security research group in the world. This leading group of security experts researches and evaluates vulnerabilities and security issues, develops assessment and countermeasure technology for IBM products and educates the public about emerging Internet threats. The X-Force provides the ability to stop more threats because of its knowledge base of information, understanding the tools and techniques used to create attacks, and collaborating with government agencies, industry consortia, and software developers. X-Force security intelligence, combined with 24x7 threat tracking and analysis through the IBM Global Threat Operations Center, ensures that IBM stays ahead of threats.

The X-Force discovered 51% percent of the high-risk, high-impact vulnerabilities found by commercial security research groups from 1998 to 2005, including the vulnerabilities exploited by the Slammer and Zotob worms. The X-Force's superior understanding of vulnerabilities is the key to IBM preemptive technology that allows you to stay ahead of threats as well.

Through the Global Threat Operations Center (GTOC) at the Atlanta, Georgia (US) headquarters, the X-Force monitors and analyzes global Internet threats 24x7. The X-Force updates the AlertCon¹⁰ in real time, providing the current global Internet threat level based on data collected from IBM Security Operations Centers and network sensors around the globe. The AlertCon rating system is the first web site indicator designed to measure the level of threat to online assets at a certain point in time.

IBM Vulnerability Management Service

The IBM Vulnerability Management Service is a unified vulnerability solution that scans networks to identify the devices running on them and to probe these devices for vulnerabilities. It analyzes the scan data and processes it for reports. You can use these reports to help you assess your network security at various levels of detail and remediate any vulnerabilities quickly. The IBM Vulnerability Management Service provides clear remediation steps for the vulnerabilities it discovers and even provides an estimate of how long these steps will take to implement.

¹⁰ To monitor the current Internet Threat Level, go to <https://webapp.iss.net/gtoc/index.html>.

The vulnerability checks in the IBM Vulnerability Management Service identify security weaknesses in all layers of a network computing environment, including operating systems, databases, web applications, and files. The IBM Vulnerability Management Service can detect malicious programs and worms, identify areas in your infrastructure that may be at risk for an attack, and verify patch updates and security compliance measures.

As we discuss in “Hosted scan engines (global scan engine pool)” on page 430, the IBM Vulnerability Management Service provides a way to scan your public facing infrastructure without any onsite equipment deployed or installed. For internal scanning needs, the same web site can be used to control and report on optional local scanners, as we see in “Distributed scan engines” on page 431.

As mentioned in Chapter 6, “Security intelligence, research, and technology” on page 149, the X-Force research team has a privileged view of the evolving threat landscape. One of the unique benefits of relying on IBM to provide the service to manage your vulnerabilities is the added value they can provide through their continuous vulnerability research.

To understand how the security capabilities of IBM Vulnerability Management Service can be mapped to the IBM Security Blueprint¹¹, see Figure 11-19 on page 428. This diagram depicts the functional components of the Threat and Vulnerability Management solution pattern, and the highlighted elements indicate those functional components that can be fulfilled, or implemented, using the IBM Vulnerability Management Service. This functional highlighting is applicable for the infrastructure service components as well.

Besides the fully highlighted elements, Figure 11-19 on page 428 also shows some medium highlighted elements. Although the IBM Vulnerability Management Service can be used to address such a component to some degree, the respective area of coverage is not considered a core function of the product and thus is considered to be limited.

If we determine the desired functionality of a solution using the Threat and Vulnerability Management solution pattern, the mapping shown in Figure 11-19 on page 428 can be used as a quick reference of the functional security management aspects of the IBM Vulnerability Management Service. This reference can help us determine which functions of a solution can be covered by selecting this product.

¹¹ For a detailed discussion of the elements, refer to Chapter 2, “The components of the IBM Security Blueprint” on page 31 and Chapter 3, “The Network, Server and Endpoint solution pattern” on page 93.

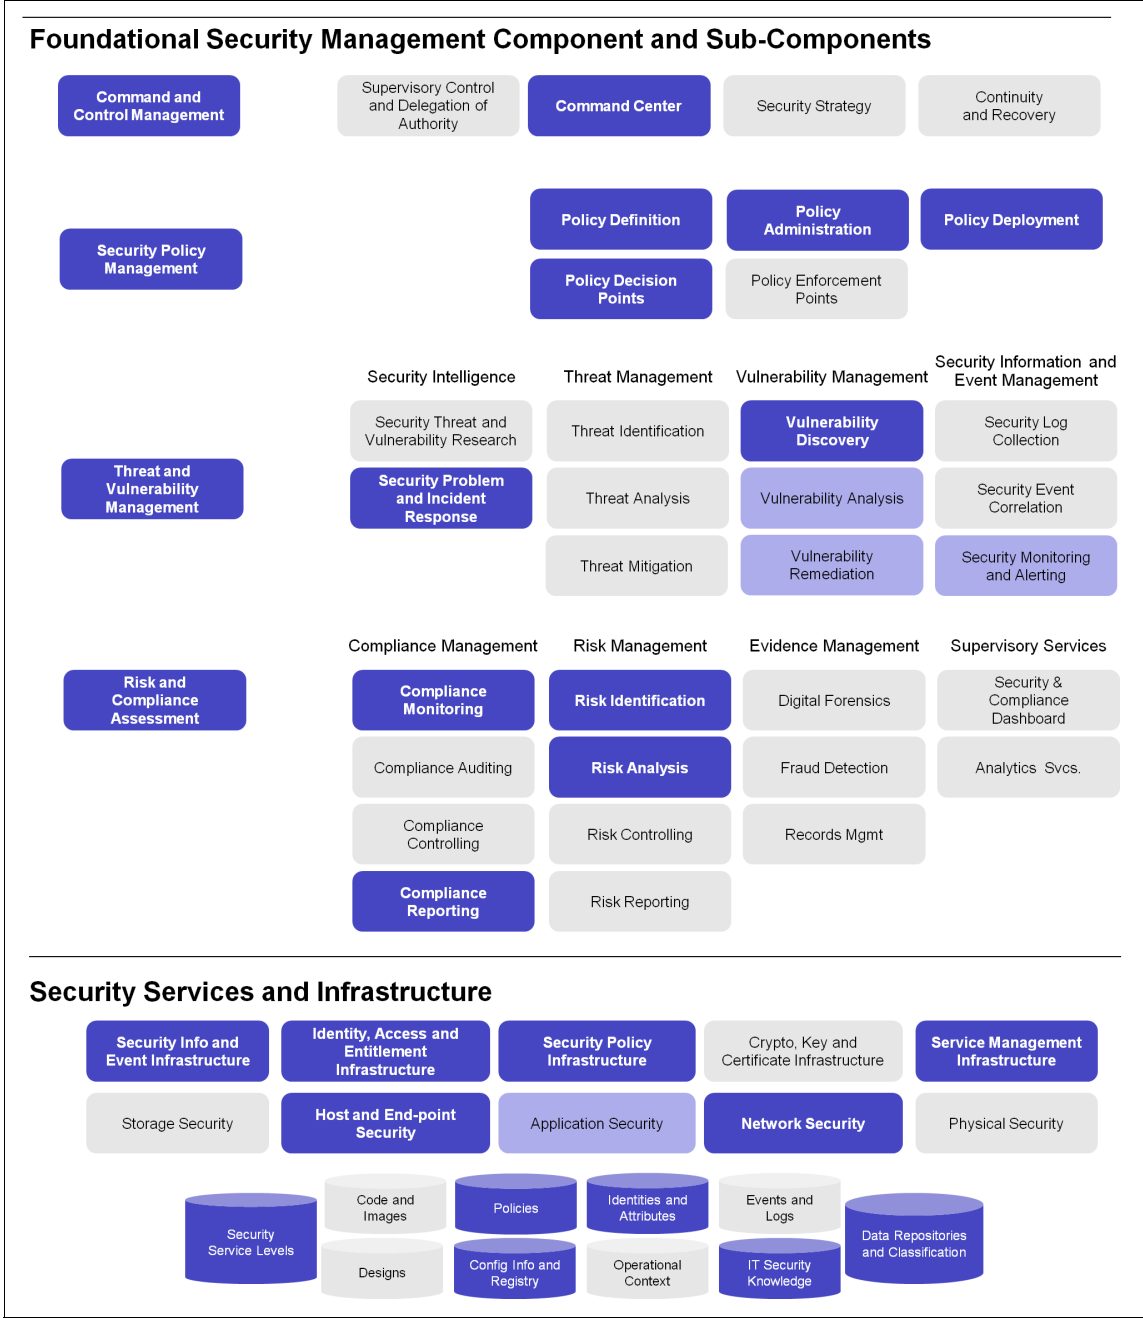


Figure 11-19 Mapping of the IBM Vulnerability Management Service to the IBM Security Blueprint

Components of the IBM Vulnerability Management Service

The IBM Vulnerability Management Service consists of three main components:

- The Virtual Security Operations Center portal website
- An authorized user within your organization can log on to the Virtual Security Operations Center securely, using HTTPS, to perform any vulnerability-related task that his or her role permits. (See “Understanding user roles and permissions” on page 434.)

The Virtual Security Operations Center is the front end of the system that communicates with local scan engines (see “Distributed scan engines” on page 431) or the IBM global scan engine pool (see “Hosted scan engines (global scan engine pool)” on page 430) to start scans and retrieve scan information. The welcome window for the Vulnerability Management Service section in the Virtual Security Operations Center is shown in Figure 11-20.

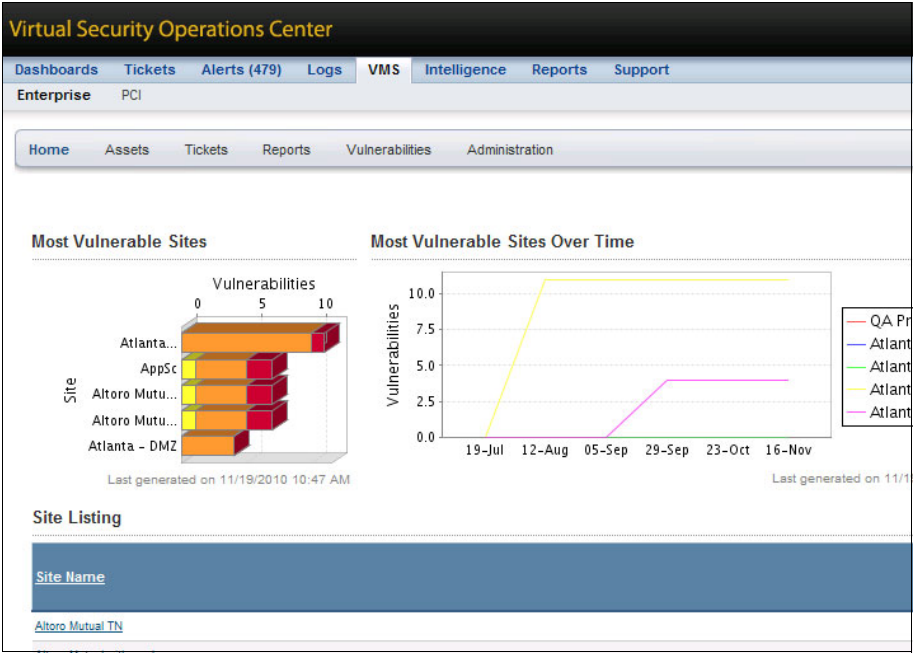


Figure 11-20 The Vulnerability Management Service web page

Other important portal functions include obtaining vulnerability information, assigning remediation tickets to people within your organization, and configuring, running, and distributing reports.

► Hosted scan engines (global scan engine pool)

Authorized users can schedule and start scans from the Virtual Security Operations Center portal website. When they want to scan their *public facing* infrastructure, the IBM Vulnerability Management Service back end will assign the scanning tasks to scanners in the IBM Global Scan Engine Pool.

These *hosted scan engines* allow you to see your network in the way an external attacker with no access permissions would see it. They scan everything on the periphery of your network, outside the firewall. These are assets that, by necessity, provide unconditional public access, such as websites and email servers. The process of using the global scan engine pool is shown in Figure 11-21.

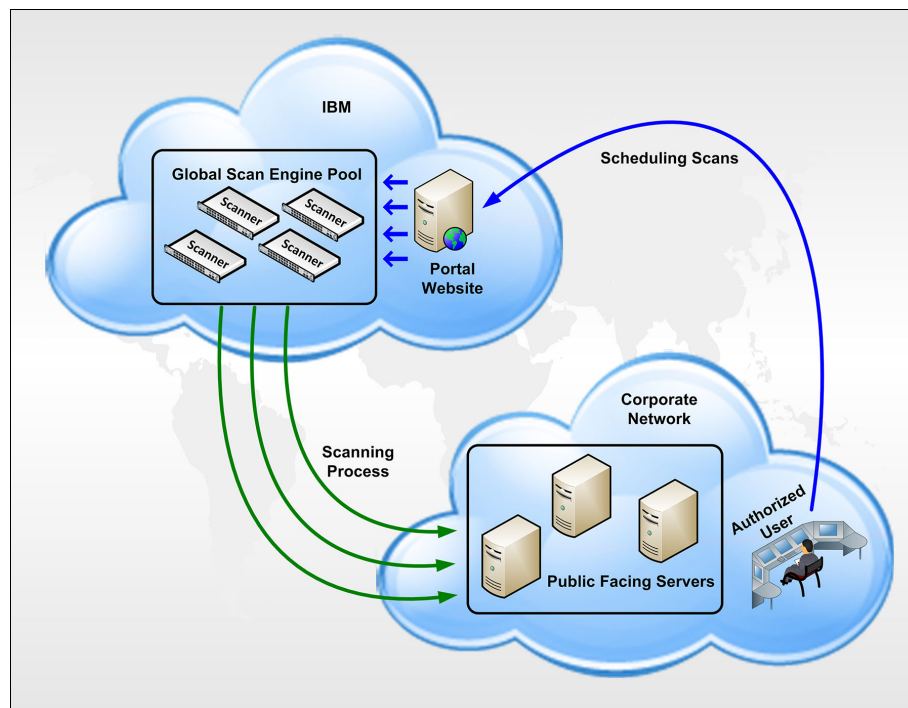


Figure 11-21 Public facing servers get checked by a pool of scanners

IBM hosts and maintains these hosted scan engines, which entails several benefits. You do not have to have to install or manage them, and the engines reside in continuously monitored data centers, ensuring high standards for availability and security.

With these advantages, it may be tempting to deploy hosted scan engines exclusively. However, hosted engines have limitations in certain use cases that warrant deploying distributed scan engines.

► Distributed scan engines

The IBM Vulnerability Management Service also offers a fully managed way to scan your organization from within. Unlike *hosted engines*, distributed scan engines allow you to inspect your network from the inside. They are ideal for core servers and workstations. You can deploy distributed scan engines anywhere on your network to obtain multiple views. This flexibility is especially valuable when it comes to scanning a network with multiple subnetworks, firewalls, and other forms of segmentation. The process of using distributed scan engines is shown in Figure 11-22.

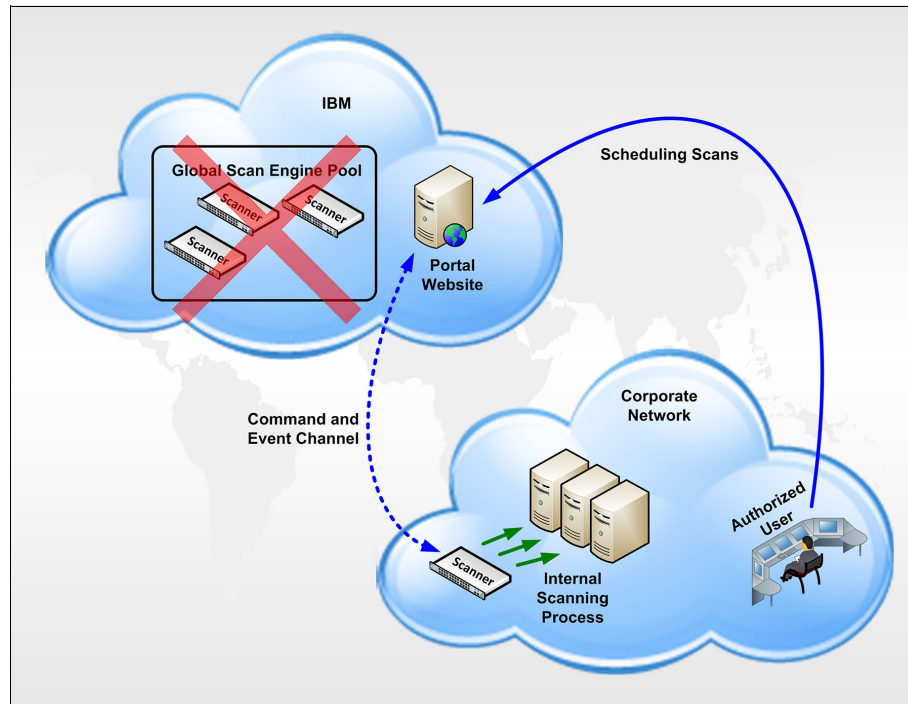


Figure 11-22 Internal scanning engines scan assets from an insider's perspective

To run or schedule a local internal scan, the process is completely parallel to external scans: An authorized user connects to the IBM Vulnerability Management Service Portal website and specifies parameters, such as how to scan and which machines to scan, just as they would with a remote scan, but now they change the scan engine to be used from the default Global Scan Engine Pool at IBM to a *distributed scan engine* locally within their organization.

Common communication channels: The IBM Vulnerability Management Service performs all of its scanning operations over the network, using common Windows and UNIX protocols to gain access to target assets. This architecture makes it unnecessary for you to install and manage software agents on your target assets.

In determining where to put scan engines, it is helpful to look at your network topology. What are the areas of separation? Where are the connecting points? If you can answer these questions, you have a pretty good idea of where to put scan engines.

The division of a network into subnetworks is often a matter of security. Communication between subnetworks may be severely restricted, resulting in slower scans. Scanning across subnetworks can be frustrating, because they are often separated by firewalls or have access control lists (ACLs) that limit which entities can contact internal assets. For both security and performance reasons, assigning a scan engine to each subnetwork is a best practice.

Planning for capacity: Empirical lab data indicates that one engine can completely scan 400 to 500 targets IP addresses in an hour. So, for example, if you have 30,000 live IP addresses and an 8 hour scan window, you need eight scan engines.

Features of the IBM Vulnerability Management Service

When you are ready to use the IBM Vulnerability Management Service, it is a good idea to follow a general sequence. Certain tasks are dependent on others being completed.

1. Install IBM scan engines, and pair them with the IBM hosted solution.
2. Create one or more *sites*.
3. Assign a scan engine and *scan template* to each site.
4. Schedule scans.
5. Create user accounts and assign site-related *roles and permissions* to them.
6. Run scans.
7. Configure and run *reports*.
8. Create asset groups to view reports and asset data.
9. Assign *remediation tickets* to users.
10. Re-run scans to verify remediation.

Understanding sites and asset groups

The IBM Vulnerability Management Service enables you to plan scans effectively by organizing your network assets into *sites* and *asset groups*.

A *site* is a physical group of assets assembled for a scan by a specific, dedicated scan engine. The grouping principle may be something meaningful to you, such as a common geographic location or a range of IP addresses, or, you may organize a site for a specific type of scan.

When you create a site, you identify the assets to be scanned, and then define scan parameters, such as scheduling and frequency. You assign a scan engine to that site, whether it is a scan engine installed locally or a hosted scan engine that is run remotely by IBM. You can assign only one scan engine to a given site. However, you can assign many sites to one scan engine.

Watch your IP addresses: If you are using RFC1918 addressing (192.168.x.x or 10.0.x.x addresses), different assets may have the same IP address. You can use site organization to enable separate scan engines located in different parts of the network to access assets with the same IP address.

You also define the type of scan you want to run for that site. Each site is associated with a specific scan. The IBM Vulnerability Management Service supplies a variety of scan templates, which can expose different vulnerabilities at all network levels. Template examples include Penetration Test, Microsoft Hotfix, Denial of Service Test, and Full Audit. You can also create custom scan templates.

Another level of asset organization is an *asset group*. Like the site, this is a logical grouping of assets, but it is not defined for scanning. An asset group typically is assigned to a non-administrative user, who views scan reports about that group to perform any necessary remediation. An asset must be included within a site before you can add it to an asset group.

An *asset group* is a logical collection of assets to which specified users have access to view data about these assets. These users are typically in charge of monitoring these assets and reporting or remediating any vulnerabilities that the IBM Vulnerability Management Service discovers on them.

Asset groups can include assets listed in *multiple sites*. They may include assets assigned to multiple scan engines, whereas sites can only include assets assigned to the same scan engine. Therefore, if you want to generate reports about assets scanned with multiple scan engines, use the asset group arrangement. You also can configure reports for combination of sites, asset groups, and assets.

Only designated IBM Vulnerability Management Service administrators are authorized to create sites and asset groups. For more details about access permissions, see “Understanding user roles and permissions” on page 434.

About scan templates

As with all other deployment options, scan templates map directly to your security goals and priorities. If you need to become HIPAA compliant, use the HIPAA Compliance template. If you need to protect your perimeter, use the Internet DMZ audit or Web Audit template.

PCI DSS scans: The IBM Vulnerability Management Service includes three default report templates mandated for PCI scans, namely Attestation of Compliance, PCI Executive Summary, and Vulnerability Details,

Alternating templates is a good idea, as you may want to look at your assets from different perspectives. The first time you ever scan a site, you may just do a discovery scan to discover exactly what is running on your network. Then, you could run a vulnerability scan using the Full Audit template, which includes a broad and comprehensive range of checks.

If you have assets that are about to go into production, it may be a good time to scan them with a Denial-of-Service template. Exposing them to unsafe checks is a good way to test their stability without affecting workflow in your business environment.

Tuning your scans by customizing a template is, of course, an option. However, keep in mind that the preset templates are themselves, best practices. The design of these templates is intended to balance three critical performance factors: time, accuracy, and resources. If you customize a template to scan more quickly by adding threads, for example, you may pay a price in bandwidth.

PCI DSS approved: IBM is an Approved Scanning Vendor (ASV), qualified by the Payment Card Industry Security Standards Council (PCI SSC), and the IBM Vulnerability Management Service is a Payment Card Industry (PCI)-sanctioned tool for conducting compliance audits.

Understanding user roles and permissions

User access to the IBM Vulnerability Management Service functions is based on roles. After you give a role to a user, you restrict access in the IBM Vulnerability Management Service to only those functions that are necessary for the user to perform that role.

Six roles exist in the IBM Vulnerability Management Service:

- Enterprise Admin: A silo owner responsible for configuring sites and assigning users to sites.

- ▶ Enterprise Regular: A security or compliance person who can run scans but cannot create sites or scan targets.
- ▶ Enterprise Restricted: An asset owner who can see vulnerabilities that were detected and manage remediation tickets.
- ▶ PCI Admin: A security or compliance person who can define users and run scans but cannot create sites or scan targets.
- ▶ PCI Regular: A security or compliance person who can run scans but cannot create sites or scan targets.
- ▶ PCI Restricted: An asset owner who can see vulnerabilities that were detected and manage remediation tickets.

Assigning tasks through tickets

You can use the IBM Vulnerability Management Service ticketing system to manage the remediation work flow and delegate remediation tasks. Each ticket is associated with an asset and contains information about one or more vulnerabilities discovered during the scanning process.

You can update other teams about the status of a remediation project, or note impediments, questions, or other issues, by annotating the ticket history. As IBM Vulnerability Management Service users and administrators add comments related to the work flow, you can track the remediation progress.

Working with reports

Reports allow you to distribute critical security data to stakeholders in your organization who do not have access to the IBM Vulnerability Management Service interface. Different export formats also make it possible to integrate IBM Vulnerability Management Service with external systems and databases.

Report configuration entails selecting a report template, assets to report on, and distribution options. You may schedule automatic reports for generation and distribution after scans or on a fixed calendar timetable, or you may run reports manually.

Several formats¹² make report data easy for security team members to distribute, open, and read immediately:

- ▶ PDF can be opened and viewed in Adobe Reader.
- ▶ HTML can be opened and viewed in a web browser.
- ▶ RTF can be opened and viewed in Microsoft Word.

¹² If you are using one of the three report templates mandated for PCI scans as of September 1, 2010, we recommend you use the RTF format. These three templates require ASVs to fill in certain sections manually.

- ▶ Text can be opened and viewed in any text editing program.
- ▶ Comma separated value (CSV) text can be opened in Microsoft Excel, and the data can easily be manipulated with macros.
- ▶ A Database Export can be output to Oracle, SQL/Server, and external databases through ODBC drivers.
- ▶ An XML Export, also known as raw XML, contains all possible data from a scan with minimal structure. Its contents must be parsed so that other systems can use its information.
- ▶ NeXpose Simple XML is also a raw XML format. It is ideal for integration of scan data with the Metasploit vulnerability exploit framework.
- ▶ SCAP Compatible XML is also a raw XML format that includes Common Platform® Enumeration (CPE) names for fingerprinted platforms. This format supports compliance with Security Content Automation Protocol (SCAP) criteria for an Unauthenticated Scanner product.
- ▶ XML arranges data in clearly organized, human-readable XML and is ideal for exporting to other document formats.
- ▶ Qualys XML Export is intended for integration with the Qualys reporting framework.

11.4 Conclusion

In this chapter, we discussed how IBM Security Services for Network, Server and Endpoint can help organizations add value to their existing and new information security environments by understanding what they have, how it is integrated, and offsetting services that do not form part of the organizations core business, to external entities in the form of managed services.

For more information related to the IBM Threat Mitigation services, go to:

<http://www.ibm.com/software/tivoli/solutions/threat-mitigation>

For more information related to the IBM Professional Security Services, go to:

<http://www.ibm.com/services/us/index.wss/offerfamily/iss/a1026711>



Part 3

Business scenarios

In the final part of this book, we illustrate two scenarios about a government agency who operates in the Public Sector Financial industry, and a health care provider that focuses on providing specialized cardiovascular related health care services. Based on their current IT infrastructure and existing security issues within that infrastructure, we help them articulate business and functional requirements. We then develop a design and implementation approach to address those security issues by using the IBM Security Blueprint solution pattern for Network, Server and Endpoint.



A-B-C Government Agency

In this chapter, we discuss a typical business scenario for A-B-C Government Agency, who operates in the Public Sector Financial industry, illustrating how a fictional organization would apply the methods and solutions discussed in this book to develop a security solutions architecture for Network, Server and Endpoint security. In this chapter, we introduce our fictional organization, A-B-C Government Agency, including its organization profile, its current solution landscape, and its medium to long-term business requirements, vision, and objectives with regard to Network, Server and Endpoint security by looking at:

- ▶ “Company overview” on page 440
- ▶ “Business vision” on page 447
- ▶ “Business requirements” on page 448
- ▶ “Functional requirements” on page 450
- ▶ “Design approach” on page 454
- ▶ “Implementation approach” on page 456

Warning: All names and references for company and other business institutions used in this chapter are fictional. Any match with a real company or institution is coincidental.

12.1 Company overview

A-B-C is a government agency that is responsible for the collection of revenue tax and income tax from citizens and businesses operating within the borders of the agency's country. A-B-C has been mandated to comply with all regulations and laws relating to privacy and financial information that have been instituted by its government.

Because information is continuously created, processed, transmitted, and stored by A-B-C business processes, applications, employees, and customers, it is the primary security business objective of A-B-C to protect this information. A-B-C must also ensure uninterrupted service delivery to internal and external customers, and that all information processing activities adhere to the legislation and regulations that influence the information and focus, in particular, for the following information:

- ▶ Personal information of customers and employees
- ▶ Financial information of customers
- ▶ Corporate and government financial information

In the past year, A-B-C has strategically moved away from paper based revenue collection and administration systems in an effort to “go green” and also to reduce labor costs traditionally associated with labor intensive, paper-based collection methods and organizational administration. This has caused significant changes to A-B-C's IT landscape due to the acquisition of virtualized solutions that have enabled A-B-C to consolidate their earlier document management systems, human resources systems, administration systems, and the distributed infrastructure used to run the systems into multiple data centers across the country.

A-B-C has also become increasingly aware of the potential dangers posed by insecure endpoint devices after concluding extensive auditing and assessment exercises. Based on the recommendations resulting from the audits and assessments, A-B-C has made a strategic decision to exert stricter control over their endpoint devices by acquiring a single solution that will allow central management, control, and reporting of security on endpoints.

A-B-C is also considered as the leading IT innovator and adopter within the public sector and has forged a close relationship with IBM over a course of several years to establish itself as such.

Let us now provide you with an overview of the IT infrastructure that supports this operation.

Staying focused: The following sections describe company information that is relevant to Network, Server and Endpoint security solutions. It is not intended to provide a complete description of the company, and the subsequent sections do not cover all the necessary activities related to information security in detail.

12.1.1 Current IT infrastructure

Being a national agency, A-B-C has its head office in the financial capital of the country, and is supported by several hundred branch offices dispersed throughout the country. The head office serves as the primary switching center, and all the branch offices connect to secondary switching centers that serve as primary access points for each of the seven major provinces. A-B-C has an IT environment with elements commonly found in financial services institutions and public sector organizations, as shown in Figure 12-1.

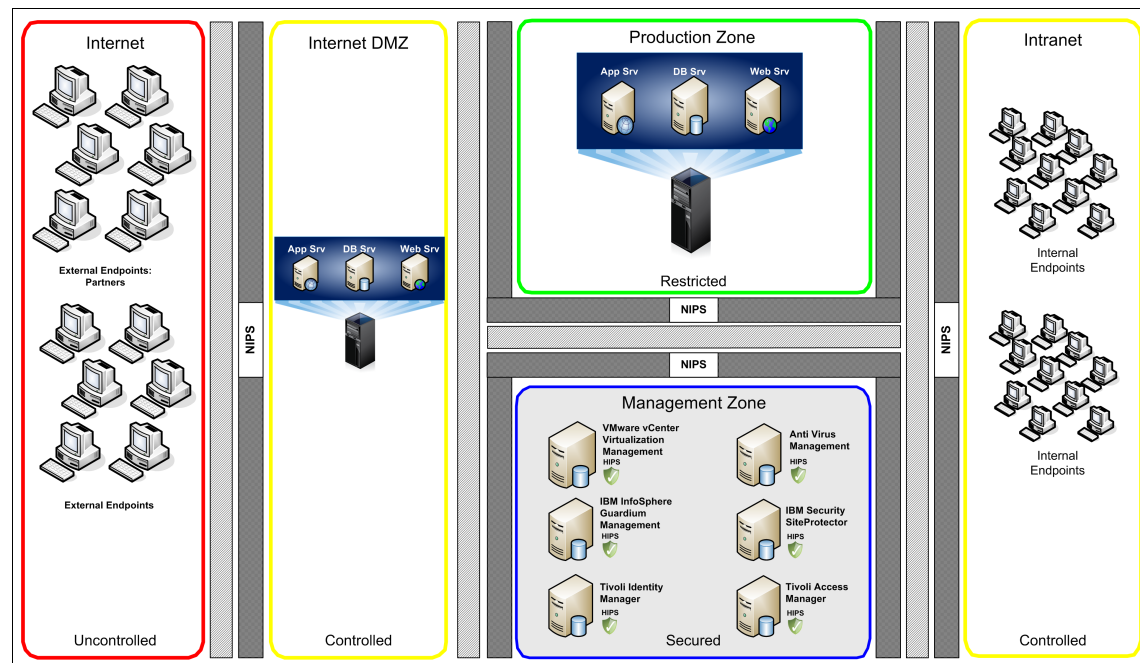


Figure 12-1 A-B-C current IT infrastructure: Network zones

A-B-C has made a strategic decision to consolidate most of its server infrastructure to virtualized environments for Microsoft Windows images. A-B-C has an internal user base of approximately 30,000 users in dispersed geographic regions who make use of desktops and mobile computers running Microsoft Windows XP, which are part of A-B-C's Microsoft Active Directory structure. Utility servers used for printing and file services all run on MS Windows platforms within VMware ESXi environments.

In addition, A-B-C supports approximately 5 million users (that is, citizens and businesses) who make use of the A-B-C web portal and conduct all tax related transactions, including the submission of tax returns. The web portal is considered the most critical business application used by both A-B-C internal and external users and is a custom developed web application used by customers of A-B-C to submit yearly income tax returns and also accessed by internal administrative users who perform day-to-day administrative and maintenance functions within the application. This application is called Online Tax Return Application for Citizens (OTRAC).

A-B-C has invested significantly in securing their infrastructure on a network and application level and has a comprehensive solution suite that includes firewalls, network intrusion prevention systems, proxy servers, endpoint firewalls and intrusion prevention systems, antivirus for endpoints and servers, which includes malware protection and email security solutions, to name but a few. A-B-C has made a significant investment in IBM Security Solutions, having deployed IBM Security Network Intrusion Prevention Systems, IBM Security Server Protection, IBM Security SiteProtector, IBM Tivoli Identity Manager, and IBM Tivoli Access Manager solutions within their current environment. In addition, A-B-C scans all of their web applications on a regular basis, including OTRAC, for vulnerabilities using IBM Rational AppScan.

Figure 12-1 on page 441 shows an overview of the current IT infrastructure of A-B-C using the network zone representation introduced in 4.4, "Common network models and security domains" on page 116.

► Internet

The Internet zone represents all of the external entities who connect to the OTRAC application and other external facing web services. The external users primarily consist of citizens who connect to the OTRAC application to submit online tax returns, but there is also a group of employees who connect to the OTRAC system during and also after normal business hours to provide application support if and when required. In addition, there are business partners who connect to web services to register custom transactions.

► Internet DMZ

The Internet DMZ, as shown in the diagram, hosts services essential to the day-to-day business of A-B-C, in particular the OTRAC application. The services are hosted on infrastructure that includes Microsoft ISS web servers, MS SQL, IBM WebSphere Application Servers, and other, earlier application servers running as images on VMWare ESXi. The servers are all protected with antivirus software and IBM Security Server host intrusion prevention systems. The Internet DMZ is protected by a firewall and an IBM Security Network IPS. Some of the components not depicted in the diagram are:

- IBM InfoSphere™ Guardium, which monitors the database activities.
- IBM Tivoli Access Manager for e-business Web Security Servers, which control access to the OTRAC application and other systems.
- DNS and email servers.

► Production Zone

The Production Zone currently hosts all the back-end servers and services that support the daily operations of A-B-C. These servers and services include web, application, and database servers and services running on virtual Microsoft Windows 2003 and 2008 servers deployed on VMWare ESXi servers. The servers are all protected with antivirus software and IBM Security Server host intrusion prevention systems. The Production Zone is protected by a separate firewall and a separate IBM Security Network IPS. Some of the components not depicted in the diagram are:

- IBM InfoSphere Guardium, which monitors the database activities.
- IBM Tivoli Access Manager for e-business Web Security Servers, which control employee access to systems and applications.

► Management Zone

The Management Zone currently hosts all the management systems for the security solutions deployed throughout the enterprise.

- IBM Security SiteProtector management system.
- IBM InfoSphere Guardium management backend.
- IBM Tivoli Identity Manager.
- IBM Tivoli Access Manager.
- Anti-virus management system.
- VMWare vCenter management backend.

The Management Zone is currently protected by a firewall, IBM Security Network IPS, and IBM Security Server Protection.

► Intranet

The Intranet Zone hosts all of the internal users that make use of A-B-C's internal and external business systems. The intranet hosts a variety of devices, including desktop computers, mobile computers, and mobile devices, such as PDAs and cellphones. The desktop and mobile computers connected to the corporate network mostly run Microsoft Windows XP, as per A-B-C's standard. These are protected by antivirus, antimalware, and desktop firewall applications. Mobile devices are only approved to connect to the corporate network if they have the same types of protection mechanisms installed, and this is controlled by MAC-based Network Access Control on the wireless access points deployed throughout the environment.

Note: Due to the limited scope of this exercise, we are unable to apply the exercise to the A-B-C environment in its entirety and therefore assume that the architecture as shown in this chapter applies to the switching centers and all associated sites in the environment.

After having discussed the logical network zone layout, a more operational architecture diagram of the current IT Infrastructure of A-B-C with specific focus on the network security devices is usually needed. Those type of diagrams can vary in their designs. It is the intent of those diagrams to provide a more specific angle on functional design within the architectural documents.

The A-B-C operational architecture diagram is provided in Figure 12-2.

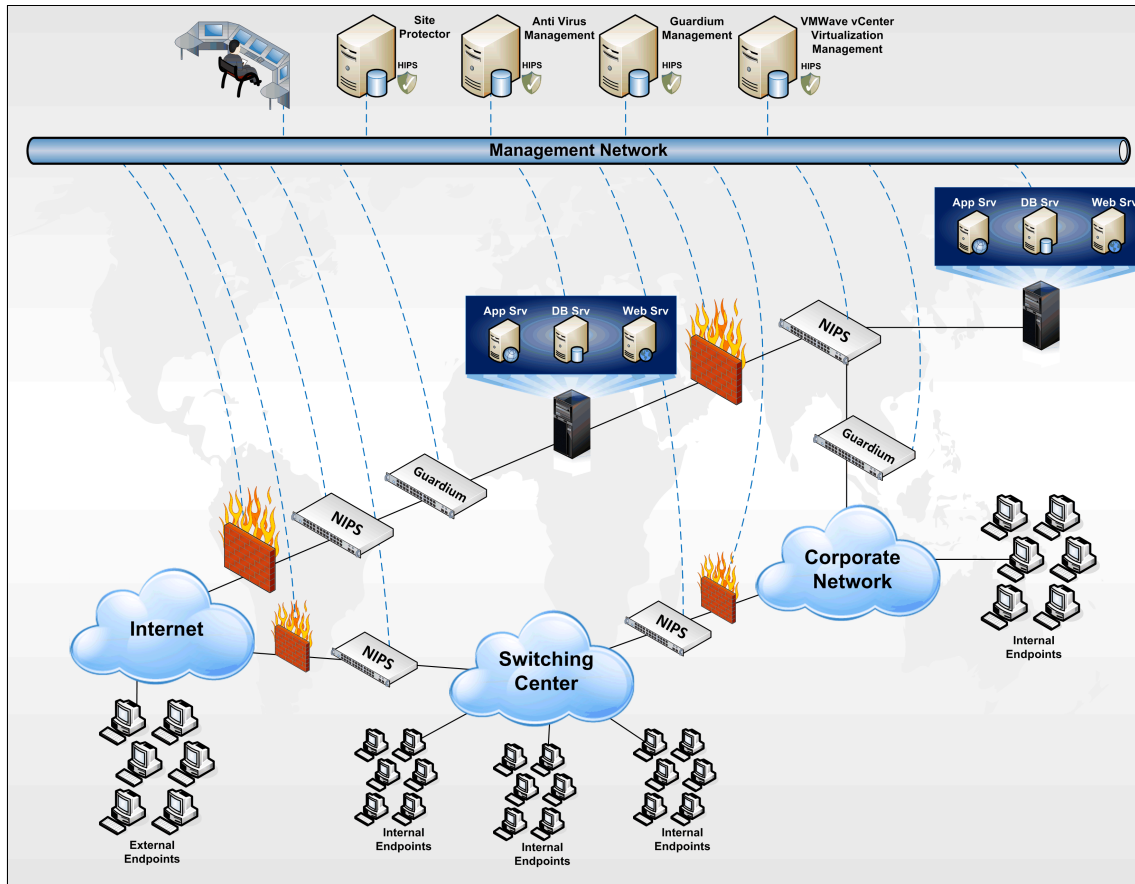


Figure 12-2 Current architecture

While A-B-C has significant protection measures in place for their IT environment, the current structure of the security controls were designed with the traditional environment in mind before A-B-C performed a major transformation from traditional platform systems into a more virtualized environment. Thus, A-B-C now needs to investigate whether this shift to the virtualized environment caused a change to the security posture and whether the security architecture has to be amended to counter potentially new security threats. A-B-C would like to review whether the security architecture could be streamlined to achieve similar efficiency gains derived for overall IT operations on the back of virtualization.

12.1.2 Security issues within the current infrastructure

After having IBM conduct a thorough Security Risk Assessment and Information Security Assessment, A-B-C has discovered significant room for improvement with regards to security for their virtual environments and endpoint devices deployed throughout the organization. With these results in mind, A-B-C has, in line with its business strategy, vision, and goals, approached IBM to enhance their security posture for these environments.

IBM conducted both a Penetration Test (discussed in 11.1.1, “Penetration Testing Service” on page 367) and an Information Security Assessment (discussed in 11.1.2, “Information Security Assessment” on page 372,) which exposed several threats and vulnerabilities within A-B-C’s operating environment. Most of the findings in the assessments indicate that the major exposures in the environment lay in the newly deployed virtual systems and the lack of consolidated management of the endpoint environment. Some of the most prominent findings in the reports indicate that the following areas are of particular concern:

- ▶ Insufficient overall endpoint security

Lack of consolidation in the endpoint computing environment brings additional security risks to A-B-C, as there is no view of the overall security level of said endpoint devices throughout the organization. This includes security patch levels, antivirus signature updates, and overall compliance settings as prescribed by A-B-C’s internal information security policy. Additionally, the lack of the ability to rapidly deploy information security countermeasures to the endpoint environment adds to the risk level, as most known attacks target endpoint devices. Consolidated management of endpoint security increases visibility and agility around endpoint security while decreasing management and personnel cost.

- ▶ Lack of overall virtual system security

Even though the overall security of virtual systems are adequate at this time, there is the potential for exposure within the virtual system environment due to lack of security within the underlying hypervisor and its operating system. This can lead to an attack vector for potential rootkit attacks with severe exposure to further privilege escalation onto the virtual operating systems running on the hypervisor, as the current environment lacks security controls limiting exposure associated to virtual machine sprawl, and also does not provide controls to allow detailed monitoring of traffic on the virtual network.

- ▶ Suboptimal usage of system resources for security services

Currently, all operating systems in the images running on the hypervisor do run a number of distinct security services, such as antivirus. If such a system becomes compromised in an attack, the associated security services are considered to be compromised as well. By introducing the appropriate security measures on the hypervisor layer, A-B-C can provide additional security controls that would not automatically fail in case of a virtualized image system compromise, and even further, could use these hypervisor layer security controls to limit the security footprint per guest operating system, eliminating redundant resource consumption and reducing overall security management complexity.

12.2 Business vision

This section reflects the core values, purpose, and goals of IT and IT security at A-B-C. These values constitute a statement of the business vision of IT security at A-B-C, which intends to achieve the following goals:

- ▶ Protect A-B-C information and IT resources at a business acceptable level by aligning with A-B-C's business strategy and business objectives.
- ▶ Provide innovative, cost-effective information security solutions that transparently secure the enterprise.
- ▶ Be responsive to changing business needs and technology directions.
- ▶ Protect A-B-C information assets while balancing security investment with risks to the business.
- ▶ Establish information and IT security as an enabler for information sharing, and enforce information security policy compliance.
- ▶ Consolidate the IT infrastructure to conserve power and utility resources in an effort to reduce the carbon footprint of the business while enhancing systems management capabilities within the business and reducing year-to-year IT spend.
- ▶ Protect and extend A-B-C's status as the leading IT innovator and adopter within the public sector.

12.3 Business requirements

To meet its statutory and legal obligations and to fulfil its government mandate, following due care and due diligence, A-B-C wants to:

- ▶ Implement a comprehensively architected and expertly implemented enterprise information security architecture encompassing, among others, the following items:
 - Governance, Risk and Compliance Management as it relates to the development, maintenance, and enforcement of governance, risk, and compliance components that are both business oriented and technical in nature.
 - Application and Process Management as it relates to secure application and process development and its associated life cycles.
 - People and Identity Management as it relates to employee life cycle management, role-based access controls, and geographic disbursement.
 - Data and Information Management as it relates to the confidentiality, integrity, and availability of all data and transactions within the business.
 - Network, Server and Endpoint Management as it relates to the mitigation of threats and management of geographically disbursed systems.
 - Physical Security as it relates to the access controls and monitoring of physical locations.
- ▶ Comply with legal and regulatory requirements.
- ▶ Improve efficiency and productivity while reducing costs.
- ▶ Cope with a growing dependency on complex IT infrastructure, especially the virtualization efforts.
- ▶ Address properly an increasing frequency of business changes.

12.3.1 IBM Security Framework mapping to business requirements

Using the IBM Security Framework definitions for business-driven security and our knowledge of the business requirements discussed in 12.3, “Business requirements” on page 448 and the current organizational infrastructure discussed in 12.2, “Business vision” on page 447, we can engage in a discussion with A-B-C to better articulate their needs. This discussion helps us to derive the proper functional requirements using the underlying IBM Security Blueprint.

► Governance, Risk and Compliance

A-B-C is obliged to protect citizen information and financial tax information. A-B-C uses IBM InfoSphere Guardium Database Monitoring and Protection, yet they might want to consider enhancing their capabilities in security log management in the future. However, this is currently not seen as the most pressing concern, because public sector regulation currently is undergoing review and any investment decision will not be made before new regulation standards are concluded. A-B-C is forced to use the country-wide security policy framework and has implemented respective controls for the policies and the related processes for security governance and incident response. With regard to virtualization, however, A-B-C is on the edge of public sector transformation, and hence is required to obey the existent policy regime, while using its flexibility to follow new paths.

► People and Identity

A-B-C has strict identity and access management processes and tools that help lower the costs related to this domain. Many processes are automated, such as password reset, process on-boarding, and terminating users. Although the assessment notes that there is room for improvement towards automation and integration of identity and access management, the current operations fulfill the compliance needs. An improvement of identity and access management is currently not seen as a premier target of A-B-C's IT budget.

► Data and Information

Access to the database servers is strictly real-time monitored and enforced, including privileged users, without the performance impact and separation of duties issues of native database logging by using IBM InfoSphere Guardium Database Monitoring and Protection. A-B-C currently uses the IBM InfoSphere Guardium Database Monitoring and Protection reporting capabilities to meet the security policies requirements for non-repudiation and access violation reporting.

► Application and Processes

Application development is mainly managed by other authorities and business partners. For live applications, A-B-C uses IBM Rational AppScan software to discover new vulnerabilities and test for new application exploits on a regular basis.

► Physical Security

Physical security enforcement is also present in all locations, but not discussed in this book.

The key area of concern with A-B-C at this time is the Threat and Vulnerability Management capabilities for the recently virtualized environment and their endpoints. Therefore, we can conclude that the new solution has to address the Network, Server and Endpoints domain of the IBM Security Framework.

Let us take the next step in understanding the functional requirements and mapping them to the IBM Security Blueprint, followed by high level discussion of the implementation approach.

12.4 Functional requirements

Based on A-B-C's business requirements and vision coupled with the findings of the Penetration Test and Information Security Assessment Service engagements, IBM and A-B-C have developed a set of functional requirements needed for the successful implementation of security solutions within the endpoint and virtual environments required for mitigating the risks exposed by the assessments. Here are the functional requirements:

- ▶ Establish a consolidated back-end view of server and endpoint security within A-B-C's environment. This requirement ultimately enables A-B-C to view the security state of every server and endpoint in their environment from a single console by a single individual.
- ▶ Establish consolidated policy enforcement for endpoint protection and virtual security throughout the enterprise to ensure that all systems make use of the appropriate policies. These policies have to be rapidly deployed from a single console that is also capable of providing a view of the overall coverage throughout the environment.
- ▶ Provide sufficient oversight of security artifacts on endpoint protection and virtual security systems during problem handling and incident response. This requirement is an increasingly vital one, as it enables A-B-C to effectively respond to security incidents that occur in any part of its geographically distributed environment quickly and efficiently.
- ▶ Establish a threat identification and mitigation capability for A-B-C's virtual security environment. Because no such view currently exists, A-B-C needs to ensure that the virtual environments are provided with the same level of security as other existing environments capable of interfacing with A-B-C's technology of choice and seamlessly integrating with current security solutions deployed by A-B-C.

- ▶ Establish effective patch deployment mechanisms for endpoints throughout the organization. By ensuring that the level of security is consistent throughout the environment, A-B-C will be able to more effectively identify threats and the impact that these threats might have on the endpoint environment.
- ▶ Establish timely security alerting capability in the virtual security environment. As the virtual cloud hosts most of the critical systems in A-B-C's environment, it is vital that A-B-C is able to identify threats within that environment and rapidly deploy countermeasures to mitigate these threats.
- ▶ Establish an enterprise view of the compliance status of all endpoints throughout the organization. By gaining a view of the compliance state of each endpoint, A-B-C is able to determine the effectiveness of compliance within their environment and also identify weak areas where attention is required.

12.4.1 IBM Security Blueprint mapping to functional requirements

Although we now understand the functional requirements for the additional security measures that A-B-C needs to implement, we still have to determine which specific solutions can potentially fulfill the functional requirements. By using the IBM Security Blueprint, we can map the functional requirements into specific blueprint areas, thereby identifying the appropriate solutions to implement within A-B-C's IT environment.

Figure 12-3 shows the mapping of the functional requirements to the IBM Security Blueprint.

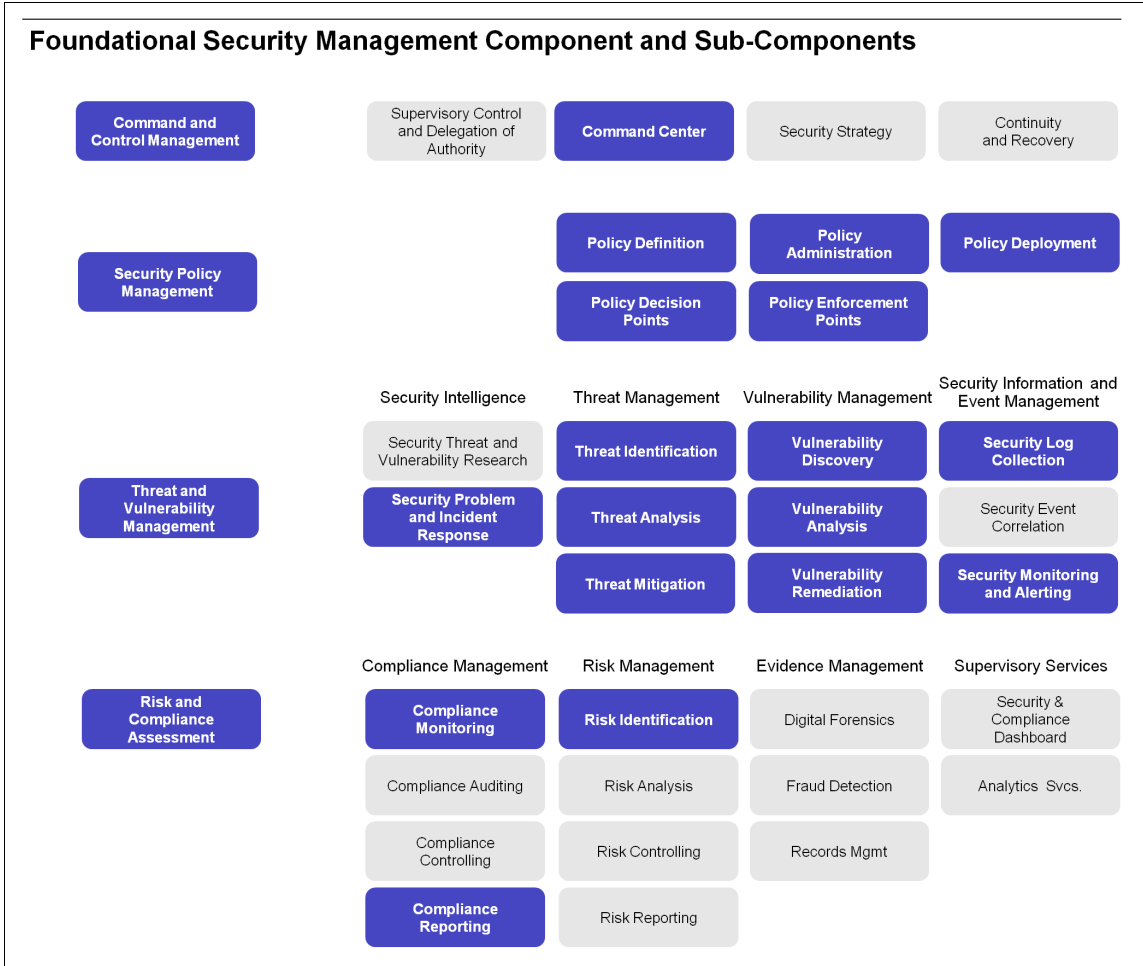


Figure 12-3 IBM Security Blueprint: Foundational components for functional requirements

Let us take a closer look at each of the functional requirements derived from the IBM Security Blueprint workshop and map them to each of the required Foundational Security Management Components and Subcomponents:

- Command and Control Management
 - Command Center: Establish a consolidated back-end view of virtual server and endpoint security.

- ▶ Security Policy Management
 - Policy Deployment: Establish mechanisms to rapidly deploy security policies throughout the enterprise.
 - Policy Enforcement: Establish mechanisms to enforce policies throughout the enterprise.
- ▶ Threat and Vulnerability Management
 - Security Problem and Incident Response: Provide sufficient oversight of security information artifacts on endpoint and virtual systems during problem handling and incident response.
 - Threat Identification: Deploy threat identification mechanisms in virtual environments.
 - Threat Mitigation: Deploy threat mitigation mechanisms in virtual environments.
 - Vulnerability Remediation: Deploy sufficient patch deployment mechanisms to endpoint systems throughout the enterprise.
 - Security Monitoring and Alerting: Introduce sufficient monitoring and alerting mechanisms within virtual environments.
- ▶ Risk and Compliance Assessment
 - Compliance Monitoring: Establish mechanisms capable of measuring compliance on endpoint devices.
 - Compliance Reporting: Establish mechanisms capable of reporting on the state of compliance of endpoint devices.

Although the business and functional requirements are both critical when identifying solutions, it is also important to consider the non-functional requirements that can influence the decision making process. After considerable analysis, the consulting team identified the following items as being critical non-functional requirements:

- ▶ A-B-C must at all times ensure the confidentiality of citizen's private information that reside on their networks and systems.
- ▶ A-B-C must show agility and resilience in enforcing the security guidelines mandated by the government, thereby ensuring the safety of confidential government information.

All of these requirements, functional, non-functional, and business, indicate the need for best of breed security solutions to mitigate the risks associated with virtual and endpoint security. These solutions should primarily provide consolidated management for endpoints, enabling the rapid deployment of countermeasures and security updates, and provide security to virtual systems by protecting the hypervisor level within the virtual systems.

The following sections show how to further use the IBM Security Framework and IBM Security Blueprint in both the design and implementation of new security solutions.

12.5 Design approach

Now that we have determined which areas of the IBM Security Blueprint our new solutions must fulfill to adequately address all of our requirements, we can map our technical requirements into the Security Services and Infrastructure components of the IBM Security Blueprint. The purpose of this exercise is to help us determine which security solutions would ultimately best satisfy all of our requirements, business, functional, non-functional and technical.

Figure 12-4 shows how the mapping was done for A-B-C using the functional requirements and existing architecture.

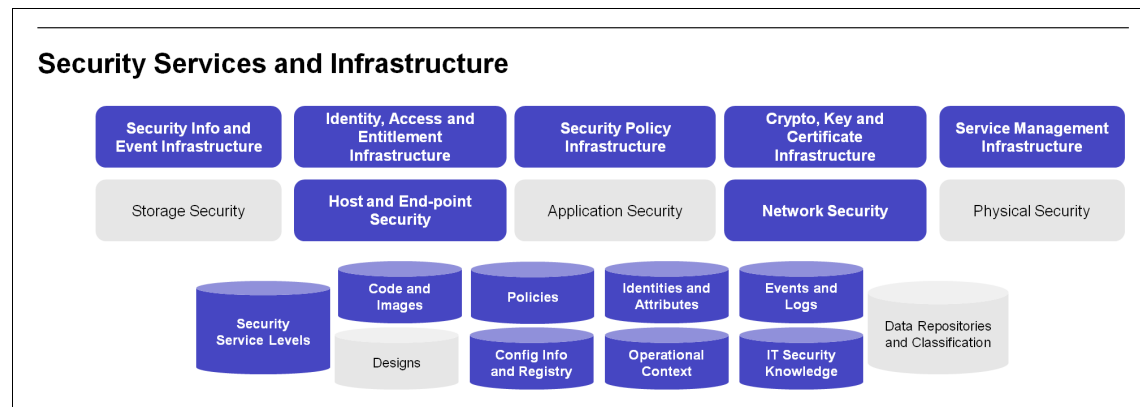


Figure 12-4 IBM Security Blueprint: Technical components for design

Based on the mapping, we now know that the solutions must provide the following Security Services and Infrastructure components:

- Infrastructure
 - Security Info and Event Infrastructure: Shows that the chosen solution must be able to integrate with the current security management solution used by A-B-C, or at least be able to serve as an effective integration point into the current management system.

- Identity, Access and Entitlement Infrastructure: Indicates that the chosen solution must ensure that appropriate access control mechanisms exist to reach agents in a controlled manner. This can be provided by encrypted communications channels between the management infrastructure and agents, and appropriate authentication and access control mechanisms within the solution.
 - Security Policy Infrastructure: Shows that the chosen solution should provide policy deployment and enforcement mechanisms that can ensure that the policies are not only delivered, but also enforced, on the network, server, or endpoint device. In addition, the solution should be able to integrate with the existing policy management repository and enforcement point.
 - Crypto, Key and Certificate Infrastructure: Indicates that the chosen solution must be able to provide secure communication paths to current and new agents in the existing environment.
 - Service Management Infrastructure: Shows that the chosen solution should provide delivery processes for the secure paths outlined above. The new security solution should integrate with the existing service management infrastructure for ticketing of security incidents, problems, and other security service events.
 - Host and End-point Security Infrastructure: Indicates that any new solution must address the virtualized environment and endpoints and provide agents for security patching, policy enforcement, compliance reporting, and vulnerability remediation.
 - Network Security Infrastructure: Indicates that the chosen solution must provide additional functionality for and additional features to the network security solutions currently being used by A-B-C.
- Services
- Security Service Levels: Are used by other infrastructure components to consistently derive frequencies and thoroughness of security settings throughout the environment.
 - Code and Images: Are used by the back-end systems to communicate with agents and to supply effective patch and deployment mechanisms to the required devices.
 - Policies: Are used to store technical settings mapping to security requirements.
 - Identities and Attributes: Holds security related technical requirements for each user by associating the business role of the user to the access requirements for each of the roles.

- Events and Logs: Holds credentials and access control information to allow communication between components and adds a layer of auditability to transactions in the environment.
- Config Info and Registry: Holds artifacts for compliance monitoring and security monitoring and alerting, as well as completion activities for all other components.
- Operational Context: Used to determine specific features and states of endpoint components.
- IT Security Knowledge: Used by process actors to efficiently use the functionality provided by the other infrastructure components.

After having completed the mapping, we are able to use the outputs to determine which solutions we need and ultimately produce an implementation plan for the selected solutions. We accomplish this goal by mapping all of our stated requirements into product features using the IBM Security Blueprint diagrams.

12.6 Implementation approach

Now that we understand all of our requirements and how they map into the IBM Security Framework and IBM Security Blueprint, we can apply the knowledge that we have gained to select the appropriate solutions to satisfy all of our requirements. Based on our design approach discussed in 12.5, “Design approach” on page 454 and by scrutinizing the chapters in this book that discuss IBM Security Solutions for Network, Server and Endpoint, we can conclude that the solutions that will best satisfy our requirements are:

- ▶ In 10.3.4, “IBM Security Virtual Server Protection for VMware” on page 348, we describe how to provide security for all of the virtualized servers.
- ▶ In 9.1, “IBM Tivoli Endpoint Manager” on page 300, we describe how to provide management and control for the endpoint environment.

Figure 12-5 shows how we used the solutions identified above to develop a new deployment design.

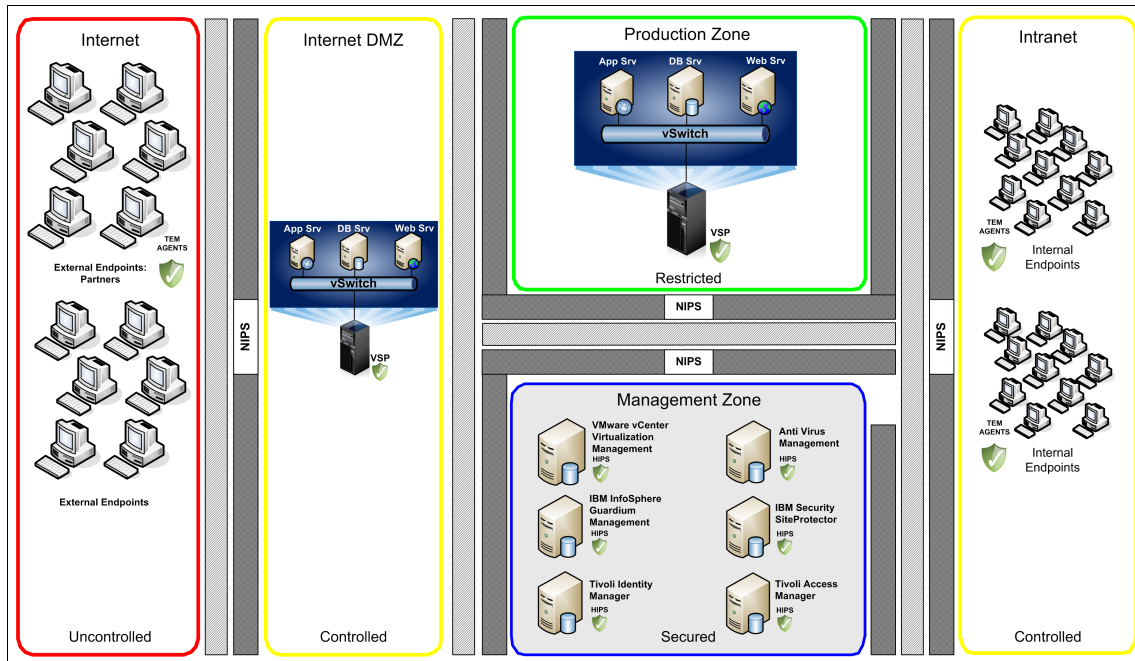


Figure 12-5 Deployment design: Network zones

Some of the key features provided by the selected solutions that address our requirements are:

- ▶ IBM Virtual Server Protection for VMware (VSP) offers integrated threat protection for VMware ESX and VMware ESXi, which provides protection for multiple layers of the virtual infrastructure, including network, virtual machine (VM), and traffic between VMs. The transparent intrusion prevention and firewall in VSP for VMware provides multilayered IPS and firewall technology to protect the virtual data center in a solution that is purpose-built to protect the virtual environment at the core of the infrastructure.
- ▶ Tivoli Endpoint Manager provides real-time visibility and control through a single infrastructure, single agent, and single console for endpoint life cycle management, endpoint protection, security configuration, and vulnerability management.

Figure 12-6 shows the details of what the final implementation architecture looks like for A-B-C.

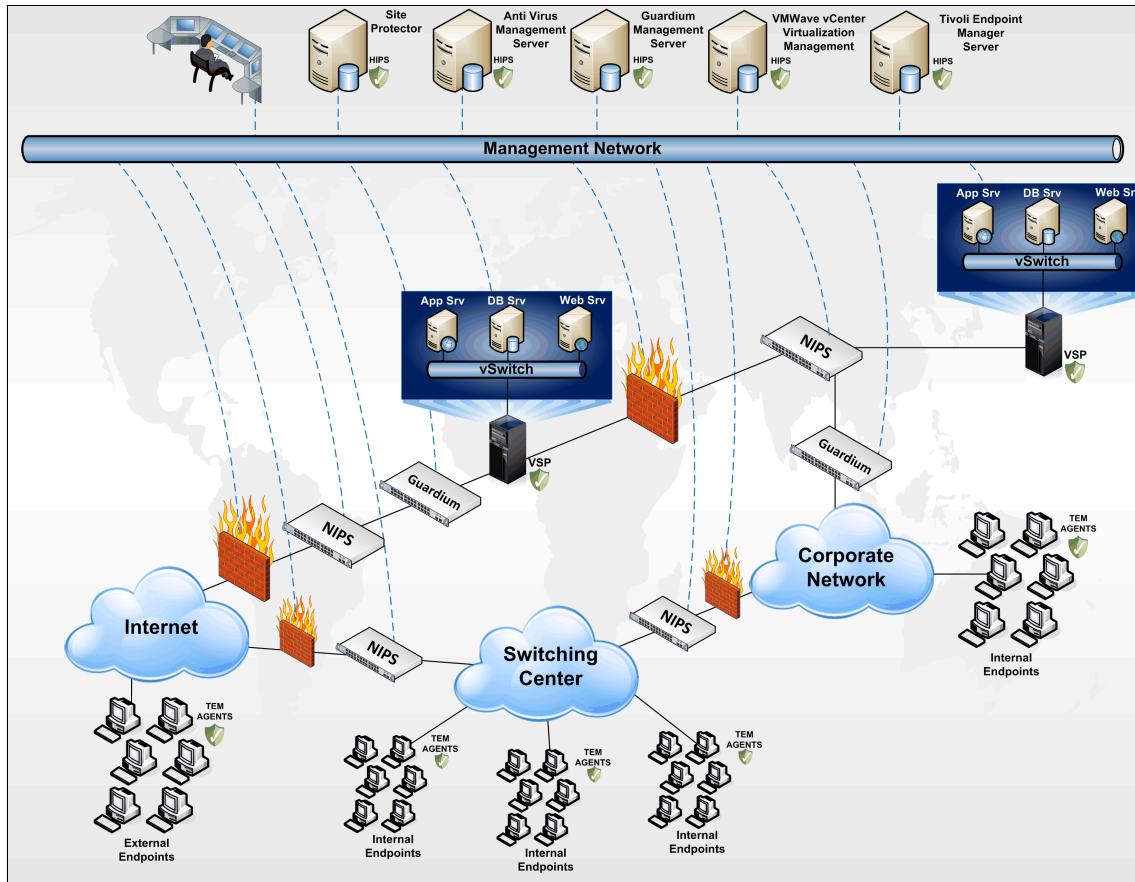


Figure 12-6 Implementation architecture

12.7 Conclusion

In this chapter, we used the IBM Security Framework and IBM Security Blueprint to develop a solution architecture for the Network, Server and Endpoint domain based on the business requirements of a fictional organization named A-B-C.

We discussed the process of developing the business vision and requirements of A-B-C into functional and non-functional requirements, and how to apply those requirements to a set of pre-determined areas within the IBM Security Blueprint, thereby identifying solutions that satisfy the stated requirements of A-B-C.



X-Y-Z Cardio

In this chapter, we discuss a typical business scenario of a health care company, X-Y-Z Cardio, and how it can use the IBM Security Framework and IBM Security Blueprint to help protect their servers and network from various security threats. We cover the following aspects:

- ▶ “Company overview” on page 460
- ▶ “Business vision” on page 468
- ▶ “Business requirements” on page 469
- ▶ “Functional requirements” on page 471
- ▶ “Design approach” on page 475
- ▶ “Implementation approach” on page 478

Warning: All names and references for company and other business institutions used in this chapter are fictional. Any match with a real company or institution is coincidental.

13.1 Company overview

X-Y-Z Cardio is a health care provider that focuses on providing specialized cardiovascular related health care services in the United States. The company was founded in California and then expanded across the country. They operate stand-alone clinics in several states. Each of these clinics occupy their own building. In those locations, X-Y-Z Cardio provides preventive care and outpatient services. For surgery and other inpatient services, X-Y-Z Cardio uses operating environments in partner hospitals. X-Y-Z Cardio also participates in research programs.

X-Y-Z Cardio maintains financial and confidential health information about their customers (patients, research partners, and affiliated hospitals). All records are kept in electronic form. One of their key applications is the *Patient Web Portal*, where patients can access their personal health records, payment information, and so on, by using a personal portal page. In addition, there is email communication between patients and service providers.

Because X-Y-Z Cardio works closely with a few pharmaceutical companies on the latest drugs for heart disease, the exchange of confidential research related information is quite extensive. Research information is also kept in electronic form and shared over the network.

X-Y-Z Cardio has built a strong and long-term reputation and financial stability over the past 15 years in the US. The company's plan is to expand their operations within the US and to open health care centers in international markets.

Let us now provide you with an overview of the IT infrastructure that supports this business.

Staying focused: The following sections describe company information that is relevant to Network, Server and Endpoint security solutions. It is not intended to provide a complete description of the company, and the subsequent sections do not cover all the necessary activities related to information security in detail.

13.1.1 Current IT infrastructure

X-Y-Z Cardio relies on two data centers: a *primary site* (located in Phoenix, AZ) and a *backup site* (located in Raleigh, NC). All production related operations are performed in the primary data center. In terms of production, the backup data center is used for disaster recovery only.

However, the backup data center is also used for development and quality assurance (QA) tests on the applications and the infrastructure. The majority of the business applications are web based. All clinics are considered to have isolated internal networks that communicate with the production servers at the primary site. The endpoint systems in the intranet networks are primarily workstations running Microsoft Windows. In addition, the majority of the clinics' modern health care appliances (such as ECG, nuclear diagnostic imaging systems, and so on) are also connected to the clinic's network and generate patient related data, which is considered part of a patient's data record.

Figure 13-1 shows the geographical distribution of the provider.

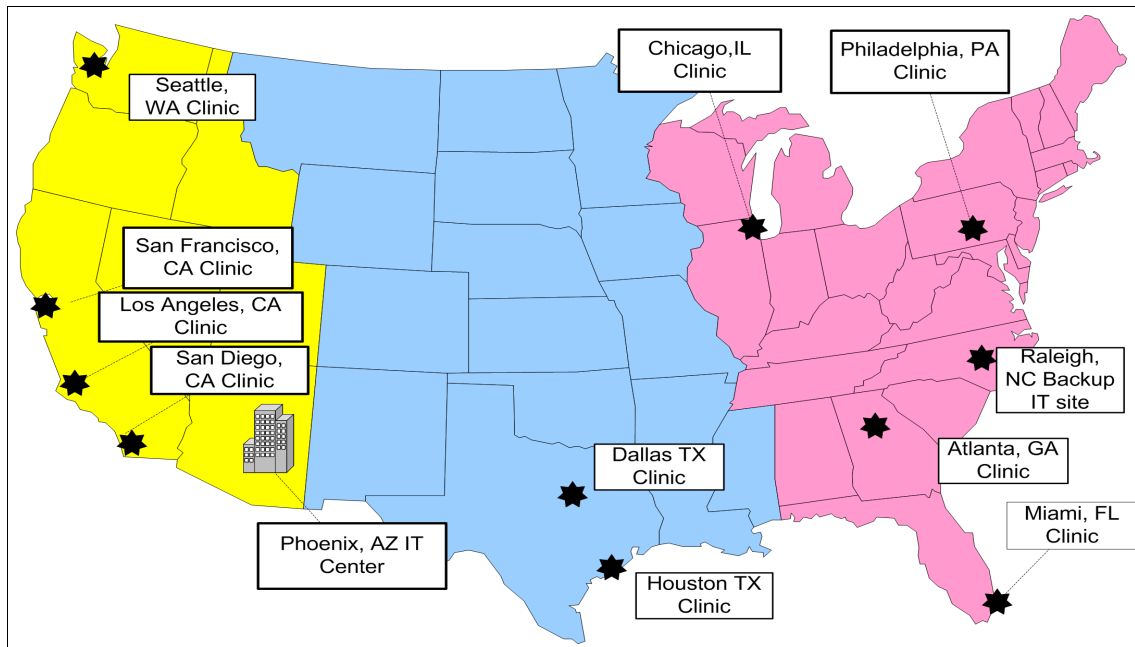


Figure 13-1 Geographical distribution of X-Y-Z Cardio

Clinics

X-Y-Z Cardio runs clinics in multiple US states. Each clinic operates its own network with multiple zones and communicates with the primary data center.

Primary data center

All customer related information is stored on separate database entities that are clustered to fulfill high availability requirements. The majority of business critical web applications are deployed in a highly available configuration using IBM WebSphere Application Server Network Deployment.

Web Security Servers (built on IBM Tivoli Access Manager technology) are located in the Internet DMZ to manage access to the applications from the Internet. Those Web Security Servers help consolidate access management for the external users accessing web applications. The Web Security Servers perform centralized authentication and authorization before allowing access to the web applications. Public web content is isolated on separate web servers and is not protected with SSL. All existing network infrastructure components (such as firewalls, switches, and routers) are designed and implemented in a high availability (redundancy) configuration.

Application servers and database servers are located in separate network zones and are isolated from each other using firewalls.

X-Y-Z Cardio's IT standards require that all servers use a UNIX or Linux based operating system, where:

- ▶ Application and database servers operate on AIX.
- ▶ Tivoli Access Manager Web Security Servers operate on Linux.
- ▶ The secure FTP (SFTP) server operates on Linux.
- ▶ DNS and email servers operate on Linux.
- ▶ The server components that are deployed in the management zone operate on AIX.

Some of the other security components that are shown in Figure 13-2 on page 463 are:

- ▶ A centralized identity management solution is based on IBM Tivoli Identity Manager. That system manages the full identity life cycle for internal as well as external users. Tivoli Identity Manager workflows are used to implement business processes for onboarding of new users, in case of role changes, and at the time when termination of access is required. The company also implements a Role Based Access Model (RBAC) that ties into the user management processes with Tivoli Identity Manager. Self-service password reset functionality provided by Tivoli Identity Manager helps lower IT help desk costs.
- ▶ Centralized access control management is implemented based on IBM Tivoli Access Manager software. All web-based access is controlled by using Web Security Servers. In addition, operating system level access is enforced on critical servers using Tivoli Access Manager for Operating Systems agents. Identified critical servers are:
 - The secure FTP server (Linux) containing confidential research reports.
 - Application and database servers in the production zone (AIX).

- ▶ Centralized log collection, analysis, and reporting on compliance for HIPAA, PCI DSS, SOX, and ISO/IEC 27002:2005 is enforced by using Security Information and Event Management (SIEM) technology. This technology is implemented by IBM Tivoli Security Information and Event Manager, which integrates with the Tivoli Identity Manager and Tivoli Access Manager infrastructure and also offers operating system level actuators (agents) that can collect logs from critical AIX and Linux servers.
- ▶ A distributed database real-time monitoring system using IBM InfoSphere Guardium Database Monitoring and Protection, which monitors all database activities in real time, including privileged user access, without the performance impact and separation of duties issues of native database logging. This solution also provides capabilities such as blocking, workflow management, and vulnerability assessments for the databases. The solution can be integrated with the Tivoli Security Information and Event Manager enterprise dashboard. The resulting compliance reports include IBM InfoSphere Guardium Database Monitoring and Protection events.

All of the systems described above are shown in Figure 13-2. The components are deployed in different network zones, which are separated by firewalls.

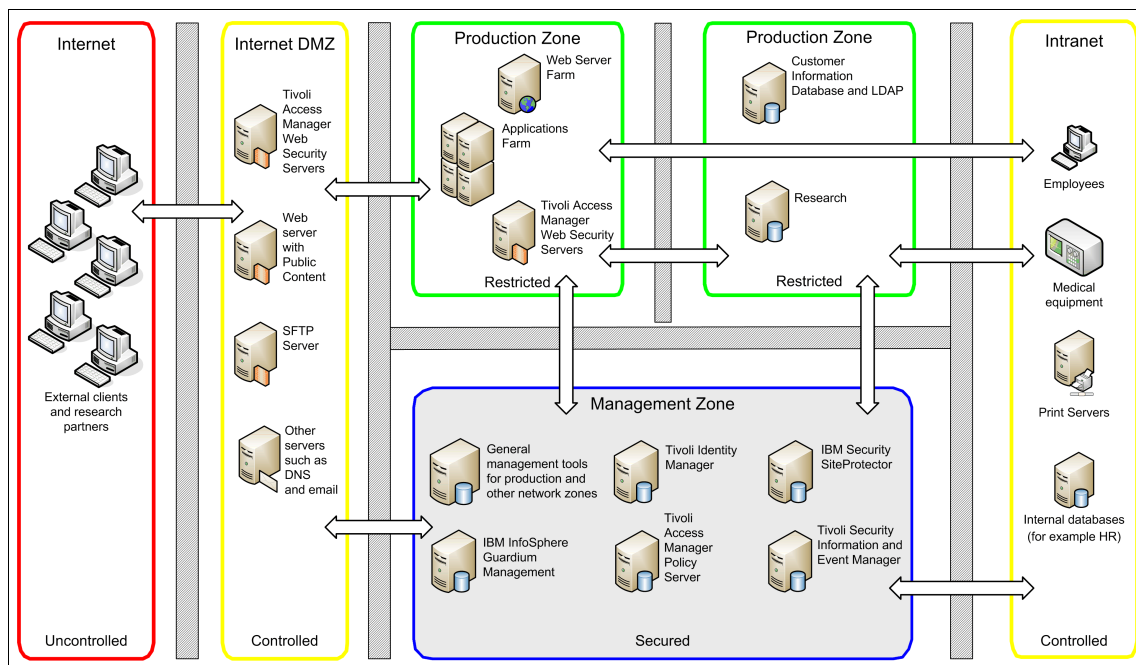


Figure 13-2 X-Y-Z Cardio current IT architecture: Network zones¹

¹ To minimize complexity, the diagram does not contain Tivoli Identity Manager, Tivoli Security Information and Event Manager, Tivoli Access Manager for Operating System, and InfoSphere Guardium agents deployed on various servers.

Figure 13-2 on page 463 shows an overview of the current IT infrastructure of X-Y-Z Cardio using the network zone representation introduced in 4.4, “Common network models and security domains” on page 116.

► Internet DMZ

The Internet DMZ hosts Tivoli Access Manager Web Security Servers that enforce access control policies for clients and research partners accessing the applications and servers from the Internet. The architecture diagram also shows standard application servers located in the Internet DMZ, such as an email and DNS server and a web server with public content. The SFTP server, which is used to exchange research data and reports with the business partners in the pharmaceutical industry, is also located in the DMZ. It operates over a secure connection.

Design consideration: You often find that storing production data on any server in the Internet DMZ is not acceptable due to compliance and policies issues in many customer environments. In that case, an SFTP server should be replaced with a Secure FTP proxy server, and data should be placed on a server hosted in the Production Zone. In our case, we used a Linux server with IBM Tivoli Access Manager for Operating Systems to enforce tight access control on the SFTP server.

► Intranet

In the diagram, we represent the intranet, which depicts the IT resource location for a single clinic, as a single network zone. Each clinic operates its own network with local print servers, HR servers for that clinic, various workstations, and clinic devices.

Design consideration: For the overall scenario, the figure would have to have multiple intranet zones, each representing a separate clinic, and each clinic having multiple zones for their own network. For the purpose of simplification in this scenario, we omit this complexity.

► Production Zone

The Production Zone is split into two segments. One segment is used to host the production web and application servers, while the second segment shows that database servers are securely isolated by another firewall. This setup helps enforce strong control on data access, often implemented in the health care industry. In addition, the production zone hosts Tivoli Access Manager Web Security Servers that enforce access control policies for employees accessing the servers from the internal network.

► Management Zone

The Management Zone shows that X-Y-Z Cardio has implemented an IBM Tivoli Security Information and Event Manager (TSIEM) solution for centralized log and event management. Many Tivoli Security Information and Event Manager agents are deployed throughout the IT infrastructure, but for the sake of simplicity, we are not specifically showing those agents.

Identity and access management solutions are also in place. The centralized Tivoli Access Manager and Tivoli Identity Manager components are securely deployed in the Management Zone to restrict access to security administrative personnel only.

Database access controls and logging is strictly enforced using a IBM InfoSphere Guardium Database Monitoring and Protection solution and collector component, which offloads the tasks from the database server (continuous analysis, reporting, and storage of audit data) to avoid any negative impact on database performance.

In addition, this zone hosts other centralized management components, such as an incident and problem management solution based on IBM Tivoli Service Request Manager.

The core applications used by external users (patients and business partners, such as pharmaceutical and research partners) are hosted on the Applications Farm:

- The Patient Web Portal, which provides the single web interface where patients can look up their health record, check and provide payments, and appointment management.
- The Secure FTP server, which is used for information exchange with pharmaceutical companies.

Internal users have access to more applications besides the ones mentioned above. Some of those are:

- HR database and applications that are locally deployed in each clinic.
- Research information shared with pharmaceutical business partners, and so on.

After discussing the logical network zone layout, a more operational architecture diagram of the current IT Infrastructure of X-Y-Z Cardio, with specific focus on the network security devices, is needed. This type of diagram can vary in design. It is the intent of this diagram to provide a more specific angle on functional design within the architectural documents.

The X-Y-Z Cardio operational architecture diagram, which involves firewalls, network switches, and major communication lines between the separate zones, is provided in Figure 13-3.

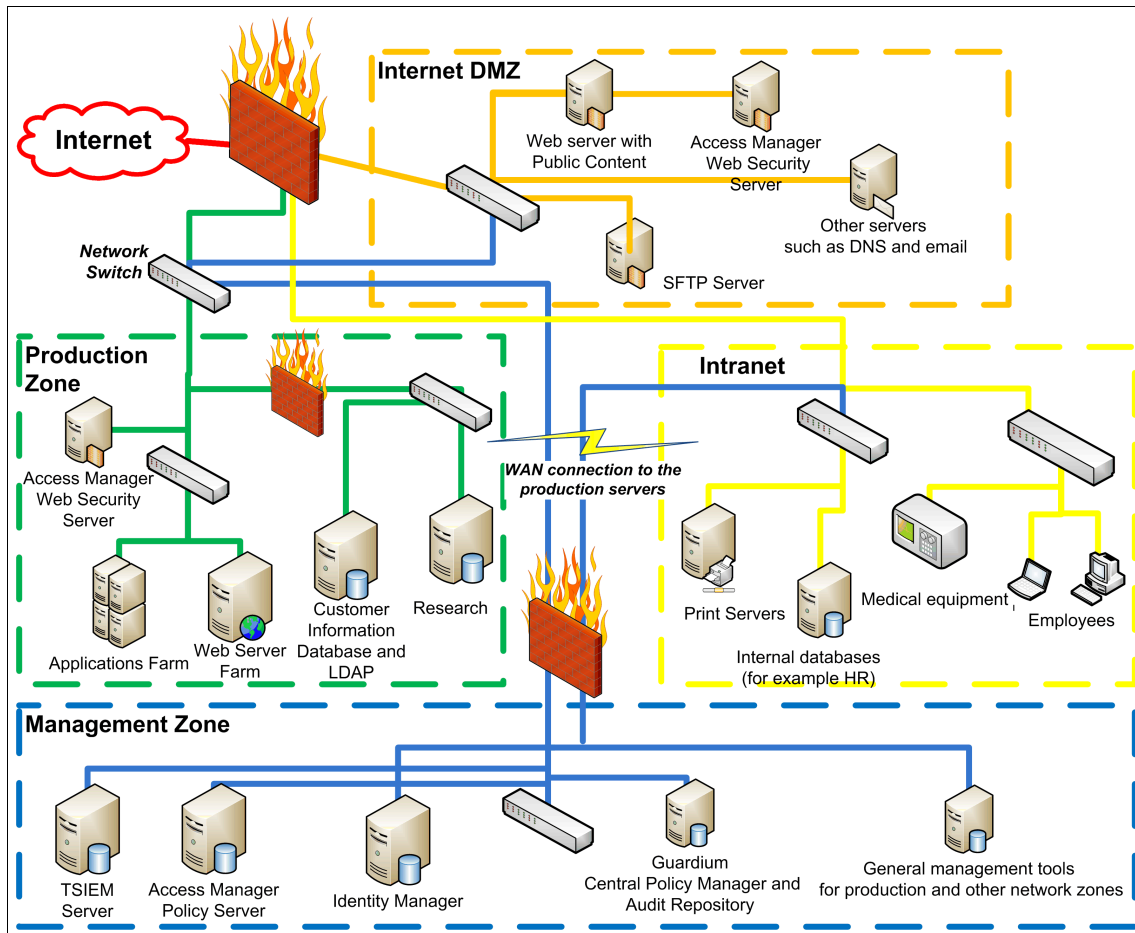


Figure 13-3 Current X-Y-Z Cardio architecture overview

High availability capability: To simplify the diagram shown in Figure 13-3, we did not include any high availability aspects of the current architecture, and we will not discuss those aspects in the proposed solution diagrams. However, all IBM security products in this scenario can be implemented in a high availability infrastructure design.

Backup site

The backup site is not designed with high availability capabilities. The disaster recovery plan target is to recover the primary data center within one day. All system snapshots from the primary site are taken nightly and transferred to the backup site. In addition to this main purpose, the backup site is used as a test and integration environment for various application development and other IT projects. As a part of application development cycles, IBM Rational AppScan is used to perform web application scanning and testing to various security vulnerabilities.

Note: For the purpose of the book, we are going to focus only on the security of the primary site. However, all security aspects and diagrams can be applied to the backup site, and the local networks at clinic locations, as well.

13.1.2 Security issues within the current infrastructure

Let us take a closer look at some of the other security related issues in the current IT environment. The majority of web applications are running on different versions of IBM WebSphere Application Servers and web servers that require regular patching and maintenance. However, change management processes are usually slow and can take some time to approve specific patch implementation in the production environment, which can potentially expose systems to security risks.

In addition, IT management is aware that the current threat management approach is passive because it only involves antivirus software, intrusion detection devices, and firewalls. This type of threat management does not address zero-day attacks. X-Y-Z Cardio is looking to enhance their security threat management with a more active approach, such as intrusion prevention devices.

A constantly increasing number of medical devices are connected to the internal IT network. Those devices operate on different versions of embedded operating systems that are located within the firmware of those devices, and the firmware is usually hard to patch. Because these devices run the same or similar code that they regular operating systems (and, to some extent, providing the same system service) do, those devices are exposed to the same threats.

Due to increased electronic communication (email) between doctors and patients, there is a higher risk of worms and other malicious code being introduced into this environment. Although X-Y-Z Cardio has antivirus software in place, it does not provide increased protection against zero-day attacks, or more focused and advanced attacks that can target a patient's information.

Besides the web applications, there are a certain number of other systems (including some health care devices) that rely on JDBC and ODBC connectivity (non-encrypted) to the patient record database. This requirement poses a constant threat for possible information leakage over those types of connections.

13.2 Business vision

This section discusses the future direction that X-Y-Z Cardio is looking to move their business development towards in the next five years:

- Expand business to European Union by opening a clinic in Munich, Germany.

By collaborating with some pharmaceutical companies on research in the EU, X-Y-Z Cardio wants to expand their business related to heart diseases by opening a clinic in Europe. The project is scheduled to begin in two years and the opening is planned in the next four years.

- Reduce costs by reusing the solutions and using the lessons learned from the current IT infrastructure.

X-Y-Z Cardio is looking to reuse their architectural and implementation approach wherever possible. While copying the general infrastructure design, they will try to fix problems found during operations and remediate them at an early phase.

- Respond to changing business needs and technology directions that can help improve customer experience by using new technologies.

Business goals are always reassessed and they are constantly changing depending of the customers needs. The Internet is becoming more and more a part of everyone's life. An increasing number of patients are using email as a communication tool with their doctors. In addition, patients are using the web applications for reviewing their own medical records and to manage appointments online.

Long term, X-Y-Z Cardio is looking to interconnect with other health care providers through the Smart Health Care system, which involves other aspects of health business, such as pharmaceutical, insurance, and so on.

This myriad of constant changes require a greater flexibility in IT technology. However, new technologies and means of communication open the possibility for new threats and vulnerabilities.

- Manage the budget by avoiding penalties due to non-compliance with major regulations, such as HIPAA, SOX, and DSS PCI.

In the health care industry, as in many others, non-compliance with regulations and standards can lead to significant financial fines and other types of penalties.

X-Y-Z Cardio is successful in managing compliance with major regulations and is looking to maintain this good practice while expanding the business.

- ▶ Protect the company image and reputation by avoiding patient information leakage, preventing security attacks, and practicing security due care and due diligence.

Any security intrusions, or leakage of any type of patient information (health, financial, or personal type), can lead to a loss of trust and damage to the X-Y-Z Cardio's reputation. Bottom line, besides losing money on penalties, it leads to loss of customers and missed revenue opportunities.

13.3 Business requirements

Based on the visionary aspects described in 13.2, "Business vision" on page 468, and by reviewing 13.1.2, "Security issues within the current infrastructure" on page 467, X-Y-Z Cardio wants to achieve the following short-term business goals:

- ▶ Improve the quality and availability of patient care and satisfaction by delivering an excellent, individualized health care experience.
- ▶ Increase the protection of all patient related information and address the diverse security risks that are driven by eHealth initiatives, emerging technologies, data explosion, and so on.
- ▶ Facilitate the management of the overall compliance posture with data privacy laws and industry regulations, such as HIPAA and PCI-DSS.

Overall, X-Y-Z Cardio wants mature security solutions that can prevent information leaks, as well as stay ahead of constantly evolving threats.

By addressing these pressing business requirements, X-Y-Z Cardio is trying to achieve the following goals:

- ▶ Continue to manage an acceptable balance between preventing security risks and adversely impacting the business.
- ▶ Constantly look for new, innovative solutions in all areas of the business, and always take security aspects into account.
- ▶ Be more proactive in security measures.
- ▶ Raise security awareness throughout the company by practicing corporate security education. We do not explicitly cover any of the solutions for this goal in this book.

13.3.1 IBM Security Framework mapping to business requirements

Using the IBM Security Framework definitions for business-driven security and our knowledge of the business requirements discussed in 13.3, “Business requirements” on page 469 and the current organizational infrastructure discussed in 13.2, “Business vision” on page 468, we can engage in a discussion with X-Y-Z Cardio to better articulate their needs. This discussion helps us to derive proper functional requirements using the underlying IBM Security Blueprint.

- ▶ People and Identity

X-Y-Z Cardio uses mature identity and access management processes and tools that help lower the costs related to this domain. Many processes are automated, such as password reset and onboarding and terminating users. The implementation uses IBM Tivoli Identity Manager and IBM Tivoli Access Manager software.

- ▶ Data and Information

X-Y-Z Cardio uses a granular information asset classification scheme paired with a least privilege principle. Access to the database servers are strictly real-time monitored and enforced, including privileged users, without the performance impact and separation of duties issues of native database logging by using IBM InfoSphere Guardium Database Monitoring and Protection. The solution is integrated with the Tivoli Security Information and Event Manager enterprise dashboard.

- ▶ Application and Processes

Application development focuses on the secure by design principle. X-Y-Z Cardio follow a rigorous release management process with a granular promotion-to-production path that specifies security testing criteria. X-Y-Z Cardio uses IBM Rational AppScan software for testing during early development stages through to applications running in the production environment. This approach helps with practicing security during the application development phase, and also helps discover any application vulnerabilities.

The processes of X-Y-Z Cardio have achieved a high level of automation and embrace security controls, such as separation of duties and creation of auditable records.

- ▶ Physical Security

Physical security controls are also embraced in X-Y-Z Cardio’s security program, and respective controls for physical access controls to facilities and systems are also present in all locations.

► Governance Risk and Compliance

X-Y-Z Cardio practices strong compliance enforcement by managing a security controls framework and strict audit and security awareness program. From a security monitoring perspective, running a Security Information and Event Management solution with compliance reporting modules for HIPPA and with other specially developed custom reports address important regulations for X-Y-Z Cardio operations in the health care industry.

XYZ Cardio designed and implemented a security policy framework and supporting processes for security governance.

X-Y-Z Cardio proactively works on identifying and eliminating security threats that enable attacks against systems, applications, and devices by using various products mentioned above (such as IBM Rational AppScan and IBM InfoSphere Guardium Database Monitoring and Protection) and integrating these products with IBM Tivoli Security Information and Event Manager for compliance reporting purpose.

Based on the business requirements discussed in 13.3, “Business requirements” on page 469 and various security issues highlighted in 13.1.2, “Security issues within the current infrastructure” on page 467, we realize that the main requirements indicate a solution in the Network, Server and Endpoint domain of IBM Security Framework.

Let us take the next step in understanding the functional requirements and mapping them to the IBM Security Blueprint, followed by high level discussions of the implementation approach.

13.4 Functional requirements

As mentioned in 13.1.1, “Current IT infrastructure” on page 460, X-Y-Z Cardio has a mature security infrastructure in place to address compliance needs by using Tivoli Security Information and Event Manager. X-Y-Z Cardio has also deployed a strong identity and access control management solution using Tivoli Identity Manager and Tivoli Access Manager. Governance, change management, and separation of duties are strictly enforced across the organization.

However, to properly address the new business requirements, X-Y-Z Cardio must enhance their security solution infrastructure. X-Y-Z Cardio defines the following high level functional requirements:

- ▶ To better manage their compliance posture with data privacy laws and industry regulations, X-Y-Z Cardio needs to employ a cost-effective centralized management solution for security configuration policies and audit data. They also have to integrate the proposed new security solution to the existing incident and problem management solution.
- ▶ To better protect all their patient related information, and to address the diverse security risks driven by eHealth initiatives, emerging technologies, data explosion, and so on, X-Y-Z Cardio needs to protect against information leakage due to intrusions and zero-day attacks. They must also protect their critical servers with additional layers of intrusion prevention.
- ▶ To improve the quality and availability of patient care and satisfaction by delivering an excellent, individualized health care experience, and to increase caregiver productivity and reduce administrative costs, X-Y-Z Cardio must address sometimes unavoidable delays in the IT change management processes to improve the security posture of their servers and of all nonstandard (embedded) operating systems of medical appliances that are connected to the network.

In addition to these three distinct functional requirements that are in line with the business requirements, X-Y-Z Cardio has some additional, more generally valid functional requirements that we need to examine:

- ▶ More real-time responsiveness to intrusion detection and prevention (blocking) events.
- ▶ Detect and, if possible, automatically counteract detected attacks.
- ▶ Provide a solution that is more proactive to security threats.

X-Y-Z Cardio already uses some solutions to identify and eliminate security threats that enable attacks against systems, applications, and devices, yet the level of automation and the speed of these activities, as well as the information at hand for Threat and Vulnerability Management, can be improved.

13.4.1 IBM Security Blueprint mapping to functional requirements

Although we now understand the functional requirements for the additional security measures that X-Y-Z Cardio needs to implement, we still have to determine which specific solutions can potentially fulfill the functional requirements. By using the IBM Security Blueprint, we can map the functional requirements into specific blueprint areas, thereby identifying the appropriate solutions to implement within X-Y-Z Cardio’s IT environment.

Figure 13-4 shows the mapping of the functional requirements to the IBM Security Blueprint.

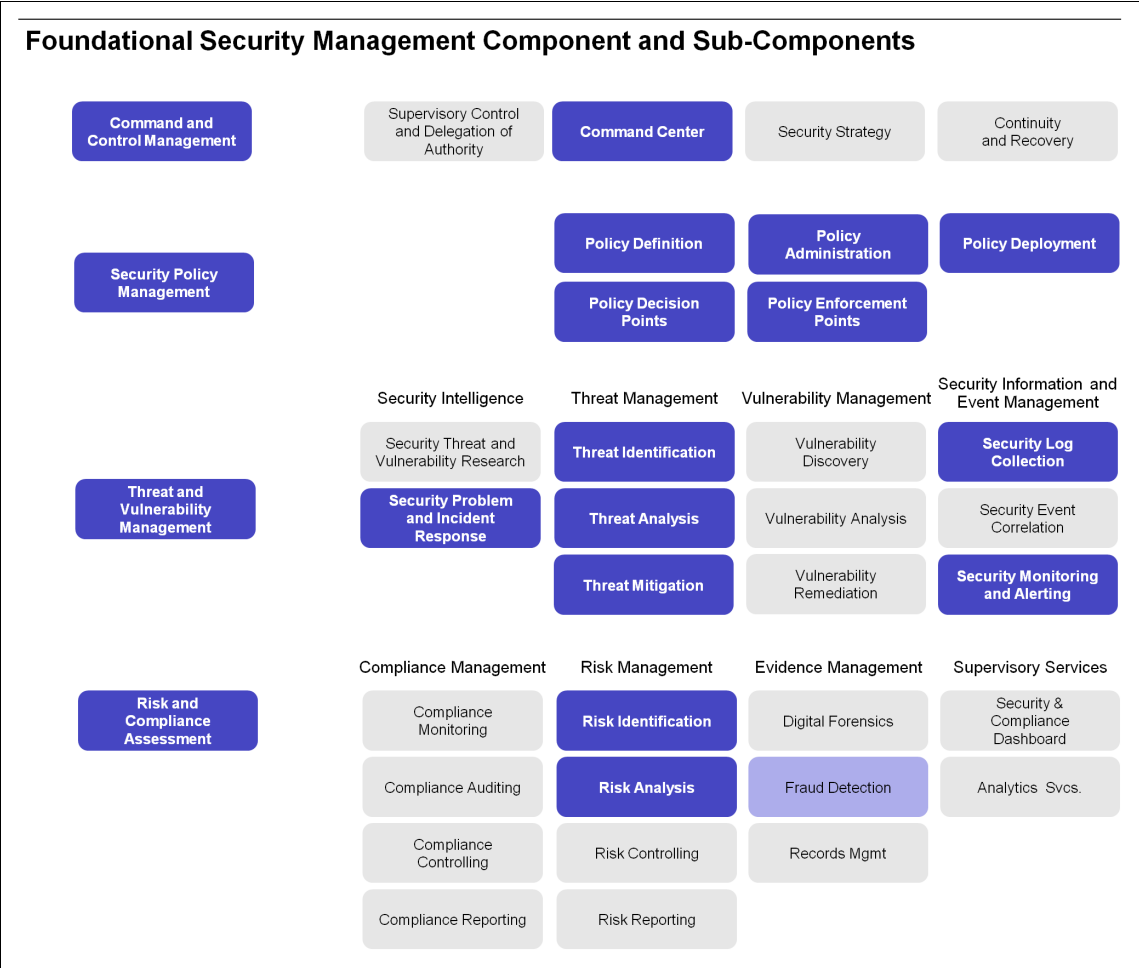


Figure 13-4 IBM Security Blueprint: Foundational components for functional requirements

Let us take a closer look at each of the functional requirements derived from the IBM Security Blueprint workshop and map them to each of the required Foundational Security Management Components and Subcomponents.

Note: When you compare this section with the one in 12.4.1, “IBM Security Blueprint mapping to functional requirements” on page 451, you realize that we are using two different ways of articulating our findings. Depending on your personal preference, you can use a bulleted list, a table, or freeform text (like in this section) to describe each of the components for your particular customer environment.

- ▶ Continue good management practices by adding a centralized management platform (*command center*) for IPS components.
- ▶ Centralize life cycle management of security configuration policies (*definitions, administration, deployment, and archiving*) for the Network IPS infrastructure.
- ▶ Better understand the actual threat posture’s (*identification*) ability to perform attack (*threat*) analysis and threat *mitigation*.
- ▶ Continue to comply with HIPAA regulations, which is currently managed with Tivoli Security Information and Event Manager infrastructure. In other words, have the ability to integrate with current Security Information and Event Management infrastructure by *collecting security logs* from different security tools deployed with the new solution, and have the ability to perform more frequent and efficient *monitoring and alerting*.
- ▶ *Report* all security related tickets and *incidents* using a standard company enterprise ticketing system (IBM Tivoli Service Request Manager).
- ▶ Use *Policy Enforcement Points (PEP)* and *Policy Decision Points (PDP)* for intrusion prevention on all network segments and strategic (production) servers. Identified strategic servers are:
 - Secure FTP server containing confidential research reports
 - Application and database servers in the production zone
- ▶ Capture and block information leakage of customer sensitive information and research data from the production databases.
- ▶ Protect financial and health records information on all systems and prevent any leakage of that information to the public web pages.
- ▶ Continue practicing strict governance, change management, and separation of duties rules.
- ▶ Rely on multiple levels of security protection by using network zoning.

Although business and functional requirements are the main parts of the security design objectives, we also must consider other non-functional requirements and constraints. These non-functional requirements and constraints might include objectives that are necessary to meet general business requirements or practical constraints about constructing security subsystems. The architectural team performed analyses of non-functional requirements, and the key non-functional requirements and constraints are:

- ▶ Provide 99.99% availability (access to the system) that translates into (“four nines”) 52.56 minutes downtime per year.
- ▶ Provide a scalable solution that can be replicated to any other data center.

Note: Because we focus on the network, server and endpoint security architecture in this book, we do not look in detail at these non-functional requirements. However, all IBM products mapped into this IBM Security Framework component are able to satisfy the non-functional requirements and constraints mentioned above.

The following sections show how to further use the IBM Security Framework and IBM Security Blueprint in both the design and implementation of new security solutions.

13.5 Design approach

Now that we have determined which areas of the IBM Security Blueprint our new solutions must fulfill to adequately address all of our requirements, we can map our technical requirements into the Security Services and Infrastructure components of the IBM Security Blueprint. The purpose of this exercise is to help us determine which security solutions would ultimately best satisfy all of our requirements, whether they are business, functional, non-functional, or technical.

Figure 13-5 shows how the mapping was done for X-Y-Z Cardio using the functional requirements and existing architecture.

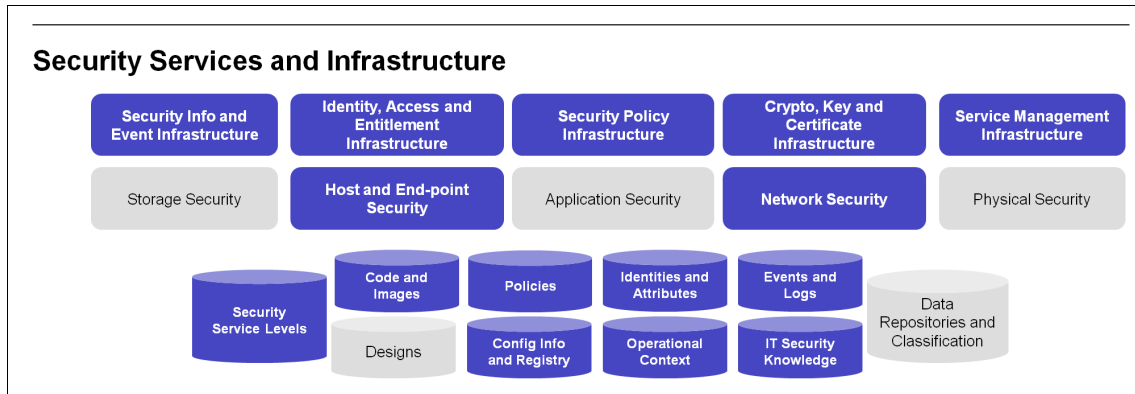


Figure 13-5 IBM Security Blueprint: Technical components for design

As part of the design, we can produce an implementation plan for our deployment that involves following steps:

1. Prioritize the requirements.
2. Map the requirements to IBM product features.
3. Define the tasks involved in using those features to satisfy the requirements and estimate the effort required for each task.
4. Divide the tasks into phases.

However, the details about those more project management related steps extend beyond the scope of this book.

Therefore, let us now focus on the technical components of the IBM Security Blueprint and how they can be mapped into technical and operational requirements.

Based on the mapping, we now know that the solutions must provide the following Security Services and Infrastructure components:

- ▶ Integrate log and intrusion events with existing *Security Information and Event Infrastructure* based on IBM Tivoli Security Information and Event Manager product.
- ▶ Define network *access control infrastructure* and mechanisms (network traffic security access policies) to allow proper communication between the components.

- ▶ Centralize management of technical level *security policies* with a single location to design, update, change, and roll back the policies.
- ▶ Secure communication between key components using *certificate infrastructure*, or some type of public/private key infrastructure.
- ▶ Integrate event, incident, and problem handling with existing *services management procedures*.
- ▶ Integrate the ticketing mechanism of the existing service management infrastructure (Tivoli Service Request Manager).
- ▶ Support the existing operating systems (*hosts*), identified as key components, for UNIX (AIX) and Linux.
- ▶ Support the exiting network infrastructure, and facilitate the deployment of *network security* mechanisms (tools) with a minimal disruption of the current network design and address schema. At the same time, the components need to be deployed into key network communications paths.
- ▶ Implement proactive protection against security threats to help stop attacks against known vulnerabilities at network zone borders.
- ▶ Mitigate the security risk from delays in change management for operating systems and middleware by stopping the threats at the network level before they reach vulnerable targets.

Note: Physical security and storage security are also an important part of the overall solution. However, they are treated as a separate project that is not in scope of this book, and they are analyzed in more detail in the other parts of the IBM Security Framework.

- ▶ Meet *security service level* agreements (SLAs) that are in line with the company's policies standards and procedures (for example, the 99.99% availability target mentioned above).
- ▶ Use a central repository for *code updates* and new release of agents.
- ▶ Provide *policy life cycle* management. The solution has to store, maintain, and provide versioning of security policies.
- ▶ Store the credential (*identities and attributes*) for communication between components; for interactive login, the credentials need to be stored in accordance with company security policies.
- ▶ Create and map *events and logs* by all agents, and integrate them with the overall event management solution (Tivoli Security Information and Event Manager) and allow for real-time alerting.
- ▶ Provide an efficient mechanism for keeping a record of the actual *configuration*.

- ▶ Consider the criticality, the layout, and the structure of communication patterns (*operational context*) to determine the threat landscape of observed vulnerabilities.
- ▶ Establish educated resources that possess the appropriate *security knowledge* and analytic skills, and also have support from the vendor.

After completing the mapping, we are able to use the outputs to determine which solutions we need and ultimately produce an implementation plan for the selected solutions. This is done by mapping all of our stated requirements into product features using the IBM Security Blueprint diagrams.

13.6 Implementation approach

Now that we understand all of our requirements and how they map into the IBM Security Framework and IBM Security Blueprint, we can apply the knowledge that we have gained to select the appropriate solutions to satisfy all of our requirements. Based on our design approach discussed in 13.5, “Design approach” on page 475 and by scrutinizing the chapters in this book that discuss IBM Security Solutions for Network, Server and Endpoint, we can conclude that the solutions that will best satisfy our requirements are:

- ▶ IBM Security Network Intrusion Prevention System (NIPS) as the robust proactive intrusion prevention system for the networks, which is discussed in more detail in Chapter 8, “Network security solutions” on page 243.
- ▶ IBM Security Real Secure Server Sensor as the host based intrusion prevention system for the AIX platform, which is discussed in more detail in Chapter 9, “Host security solutions” on page 299.
- ▶ IBM Security Server Protection for Linux, as the host based intrusion prevention system for Linux, which is discussed in more detail in Chapter 9, “Host security solutions” on page 299.
- ▶ IBM Security SiteProtector, as the centralized management platform for intrusion prevention systems (network and hosts), which is discussed in more detail in Chapter 7, “Centralized management” on page 199.

Using the components mentioned above, we can position them in the new solution using a network zone diagram, as shown at Figure 13-6.

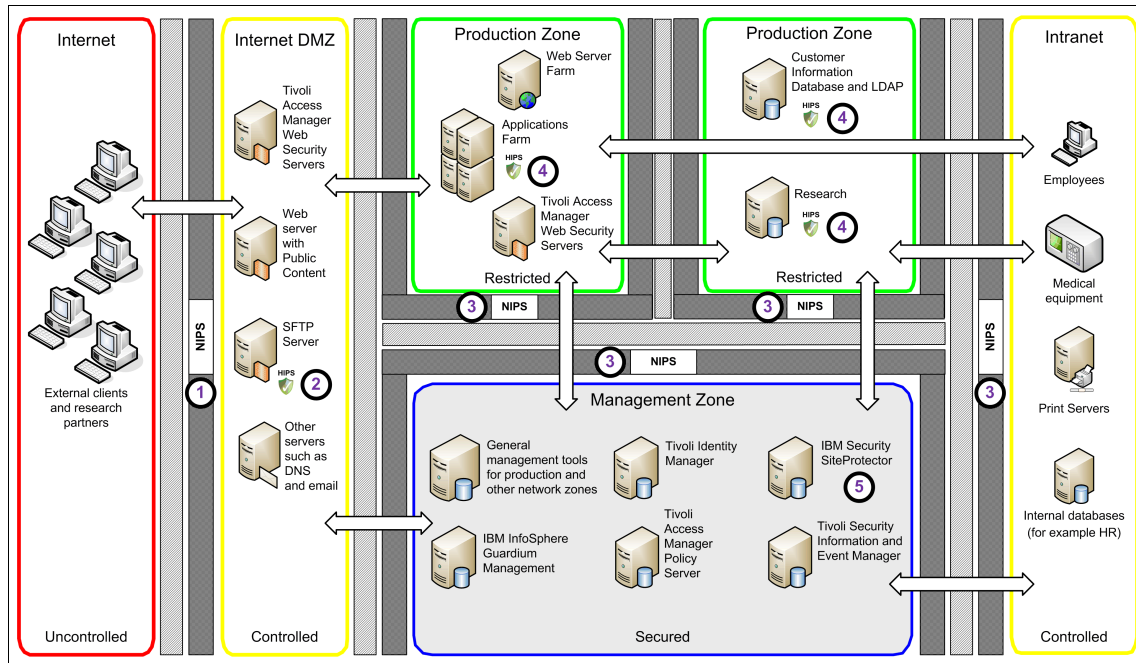


Figure 13-6 New proposed solution using a network zone

Some of the key features provided by the selected solutions that address our requirements are:

1. An IBM Security Network IPS device deployed in the area marked with ❶ in Figure 13-6 applies a “Virtual Patch” to the all systems behind the Internet DMZ firewall. We can apply policies that will stop sensitive data transfers (such as email and email attachments) that contain, for example, private patient information (such as credit card data, Social Security numbers, and so on). This device can help protect against zero-day attacks and the propagation of known and unknown viruses, worms, and other types of malware.
2. IBM Security Server Protection for Linux implemented on the SFTP server adds an additional layer of protection to a server that stores reports and data on drugs and other research related information that are exchanged with pharmaceutical companies. Because the Network IPS device deployed in ❶ is not able to block encrypted FTP traffic, adding host based IPS protection adds significant value.

3. Network IPS devices in other zones help satisfy the layered protection requirement. These devices can contain attacks within one network zone. In addition, the devices in the database zone can be configured with Information and Data Leakage Protection policies to prevent uncontrolled use of patient sensitive information across ODBC and JDBC connections.
4. Host IPS agents deployed on servers in the production zone can also provide additional security layer and allows you to apply policies specific to hosts, such as a Buffer Overflow Exploit Prevention (BOEP) policy.

All devices above provide the Virtual Patch technology to all critical and non-critical assets.

5. IBM Security SiteProtector helps deliver centralized management of intrusion prevention agents (host and network). It can analyze intrusions, handle real-time discovery, and monitor subsequent attacks. Integration with the Tivoli Security Information and Event Manager can help maintain further compliance to critical regulations. In addition, SiteProtector has an API that can integrate the company ticketing solution and incident management system based on Tivoli Service Request Manager. Going a step beyond, we can integrate SiteProtector with IBM Rational AppScan and discover any new vulnerabilities, which can help create new signatures (policies) in SiteProtector. SiteProtector can then again push such policies to all deployed Network IPS systems.

The final architecture overview diagram in Figure 13-7 shows more details about how we place the Network Intrusion Prevention Systems with inline protection mode, showing the position of firewalls and network switches in the architecture.

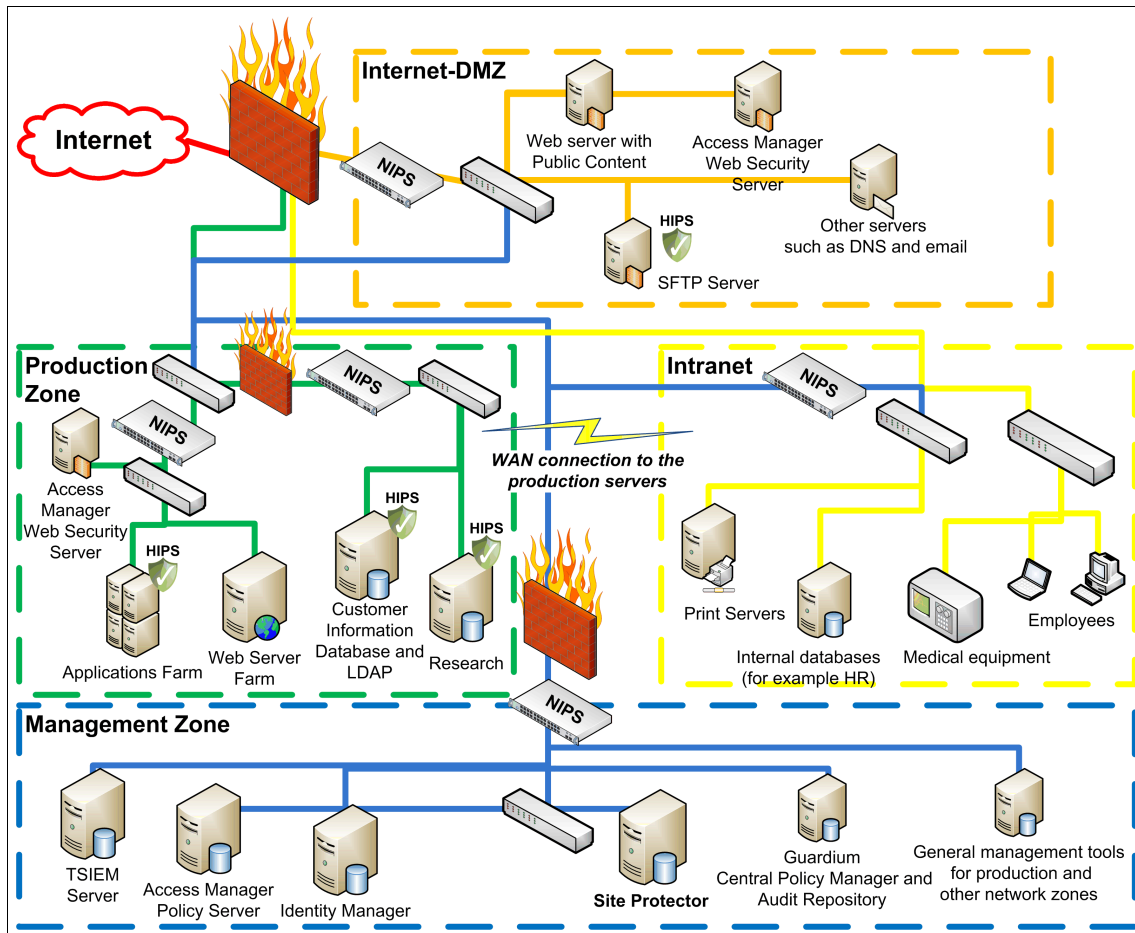


Figure 13-7 Architecture overview of the new proposed solution

13.7 Conclusion

In this final chapter, we combined several IBM Security Solutions products to help an organization fulfill the requirements for Threat and Vulnerability Management. We also showed how the IBM Security Framework and the IBM Security Blueprint can provide a structure to derive the IT functional and technical requirements from the business vision, goals, and requirements.

First, we introduced the X-Y-Z Cardio health care corporation. We discussed the company profile, its current IT infrastructure, and its issues with networks, servers, and endpoints.

Next, we discussed the business requirements and the associated functional requirements. After refining these requirements to a more detailed technical level, we described the design approach that X-Y-Z Cardio took for their solution following the IBM Security Framework and the Threat and Vulnerability Management Solution Pattern of the IBM Security Blueprint.

When applied to their unique IT environment, this process of analysis and design helped X-Y-Z Cardio define an implementation plan.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 484. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *DataPower Architectural Design Patterns: Integrating and Securing Services Across Domains*, SG24-7620
- ▶ *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1*, SG24-7616
- ▶ *Enterprise Security Architecture Using IBM Tivoli Security Solutions*, SG24-6014
- ▶ *Identity Management Design Guide with IBM Tivoli Identity Manager*, SG24-6996
- ▶ *Improving Your Web Application Software Development Life Cycle's Security Posture with IBM Rational AppScan*, REDP-4530
- ▶ *Integration Guide for IBM Tivoli Netcool/OMNIBus, IBM Tivoli Network Manager, and IBM Tivoli Netcool Configuration Manager*, SG24-7893
- ▶ *Introducing the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security*, REDP-4528
- ▶ *IT Security Compliance Management Design Guide with IBM Tivoli Security Information and Event Manager*, SG24-7530

Other publications

These publications are also relevant as further information sources:

- ▶ Saltzer, J.H., and Schroeder, M.D., *The Protection of Information in Computer Systems*, The Proceedings of the IEEE, 1975, ISBN 00189219

Online resources

These websites are also relevant as further information sources:

- ▶ An Introduction to ISO 27001, ISO 27002 through ISO 27008:
<http://www.27000.org/>
- ▶ The IBM product documentation for IBM Security Network IPS, IBM Security SiteProtector, and IBM Security Virtual Server Protection:
http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp?topic=/com.ibm.ips.doc/IBMSecNetIPS_landing_page.html

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Redbooks

IBM Security Solutions Architecture for Network, Server and Endpoint

(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages



IBM Security Solutions Architecture for Network, Server and Endpoint

**Comprehensive
discussion of the
IBM Security
Framework and IBM
Security Blueprint**

**Detailed insight into
the threat and
vulnerability
landscape**

**Extensive solution
architecture and
component
introduction**

Threats come from a variety of sources. Insider threats, as well as malicious hackers, are not only difficult to detect and prevent, but many times the authors of these threats are using resources without anybody being aware that those threats are there. Threats would not be harmful if there were no vulnerabilities that could be exploited. With IT environments becoming more complex every day, the challenges to keep an eye on all potential weaknesses are skyrocketing.

Smart methods to detect threats and vulnerabilities, as well as highly efficient approaches to analysis, mitigation, and remediation, become necessary to counter a growing number of attacks against networks, servers, and endpoints in every organization.

In this IBM Redbooks publication, we examine the aspects of the holistic Threat and Vulnerability Management component in the Network, Server and Endpoint domain of the IBM Security Framework. We explain the comprehensive solution approach, identify business drivers and issues, and derive corresponding functional and technical requirements, which enables us to choose and create matching security solutions.

We discuss IBM Security Solutions for Network, Server and Endpoint to effectively counter threats and attacks using a range of protection technologies and service offerings. Using two customer scenarios, we apply the solution design approach and show how to address the customer requirements by identifying the corresponding IBM service and software products.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks