

Chapter 7

Topological Vulnerability Analysis

Sushil Jajodia and Steven Noel

Traditionally, network administrators rely on labor-intensive processes for tracking network configurations and vulnerabilities. This requires a great deal of expertise, and is error prone because of the complexity of networks and associated security data. The interdependencies of network vulnerabilities make traditional point-wise vulnerability analysis inadequate. We describe a Topological Vulnerability Analysis (TVA) approach that analyzes vulnerability dependencies and shows all possible attack paths into a network. From models of the network vulnerabilities and potential attacker exploits, we compute attack graphs that convey the impact of individual and combined vulnerabilities on overall security. TVA finds potential paths of vulnerability through a network, showing exactly how attackers may penetrate a network. From this, we identify key vulnerabilities and provide strategies for protection of critical network assets. TVA provides predictive context for network hardening, intrusion detection deployment and alarm correlation, and optimal attack response. Further, it employs efficient algorithms that scale well to larger networks.

7.1 Introduction

By their very nature, security concerns on networks are highly interdependent. Each machine's susceptibility to attack depends on the vulnerabilities of the other machines in the network. Attackers can combine vulnerabilities in unexpected ways, allowing them to incrementally penetrate a network and compromise critical systems. To protect our critical infrastructure networks, we must understand not only our individual system vulnerabilities, but also their interdependencies.

While we cannot predict the origin and timing of attacks, we can reduce their impact by knowing the possible attack paths through our networks. We need to transform raw security data into roadmaps that let us proactively prepare for attacks, manage vulnerability risks, and have real-time situational awareness. We cannot rely on manual processes and mental models. We need automated tools to analyze and visualize vulnerability dependencies and attack paths, so we can understand our overall security posture, providing context over the full security life cycle.

Our approach to such full-context security is called Topological Vulnerability Analysis (TVA) [1][2]. TVA monitors the state of network assets, maintains models of network vulnerabilities and residual risk, and combines these to produce models that convey the impact of individual and

Center for Secure Information Systems, George Mason University, Fairfax, VA 22030-4444, USA,
e-mail: {jajodia, snoel}@gmu.edu

combined vulnerabilities on overall security posture. The central product of this tool is an attack graph showing all possible ways an attacker can penetrate the network.

Our TVA approach provides a unique new capability, transforming raw security data into a roadmap that lets one proactively prepare for attacks, manage vulnerability risks, and have real-time situational awareness. It supports both offensive (e.g., penetration testing) and defensive (e.g., network hardening) applications. The mapping of attack paths through a network via TVA provides a concrete understanding of how individual and combined vulnerabilities impact overall network security. For example, we can

- Compare possible expenditures of resources to determine which will have the greatest impact on overall security
- Determine how much a new vulnerability will impact overall security
- Determine whether risk-mitigating efforts have a significant impact on overall security
- Analyze how changes to individual machines may increase overall risk to the enterprise

We have implemented our approach as a security tool that transforms raw security data into a model of all possible network attack paths. In developing this tool, we have met key technical challenges, including the design of appropriate models, efficient model population, effective visualizations and decision support tools, and the development of scalable mathematical representations and algorithms. The result is a working software tool that offers truly unique capabilities.

In the design of computers and networks, security is often not given adequate priority. This is compounded by the fact that each machine's exposure to attack depends on the vulnerabilities of the other machines in the network. Attackers can combine vulnerabilities in unexpected ways, allowing them to incrementally penetrate a network and compromise critical systems.

To protect critical networks, we must understand not only individual system vulnerabilities, but also their interdependencies. TVA places vulnerabilities and their protective measures within the context of overall network security by modeling their interdependencies via attack graphs. The analysis of attack graphs provides alternative sets of protective measures that guarantee safety of critical systems. Through this unique new capability, administrators are able to determine the best sets of protective measures that should be applied in their environment.

Still, we must understand that not all attacks can be averted in advance, and there must usually remain some residual vulnerability even after reasonable protective measures have been applied. We then rely on the detect phase to identify actual attack instances. But the detection process needs to be tied to residual vulnerabilities, especially ones that lie on paths to critical network resources as discovered by TVA.

Once attacks are detected, comprehensive capabilities are needed to react to them. TVA can reduce the impact of attacks by providing knowledge of the possible vulnerability paths through the network. TVA attack graphs can be used to correlate and aggregate network attack events, across platforms as well as across the network. These attack graphs also provide the necessary context for optimal response to attacks.

The next section shows the architecture of our TVA tool. Section 7.3 illustrates the TVA approach through a simple example. Section 7.4 describes the process for building models of the security environment for TVA, for generating multi-step attack graphs. Section 7.5 explains the visualization and analysis of TVA attack graphs. Section 7.6 examines scalability of TVA. Section 7.7 reviews related work, and Section 7.8 summarizes our approach.

7.2 System Architecture

Because of vulnerability interdependencies across networks, a topological attack graph approach is needed for defense against multi-step attacks. The traditional approach that treats network data and events in isolation without the context provided by attack graphs is clearly insufficient. TVA

combines vulnerabilities in ways that real attackers might do, discovering all attack paths through a network.

Figure 7.1 shows the architecture of our software tool for TVA attack graph analysis. *Network Capture* builds a model of the network, in terms of relevant security attributes. *Vulnerability Database* represents a comprehensive repository of reported vulnerabilities, with each vulnerability record listing the affected software (and hardware).

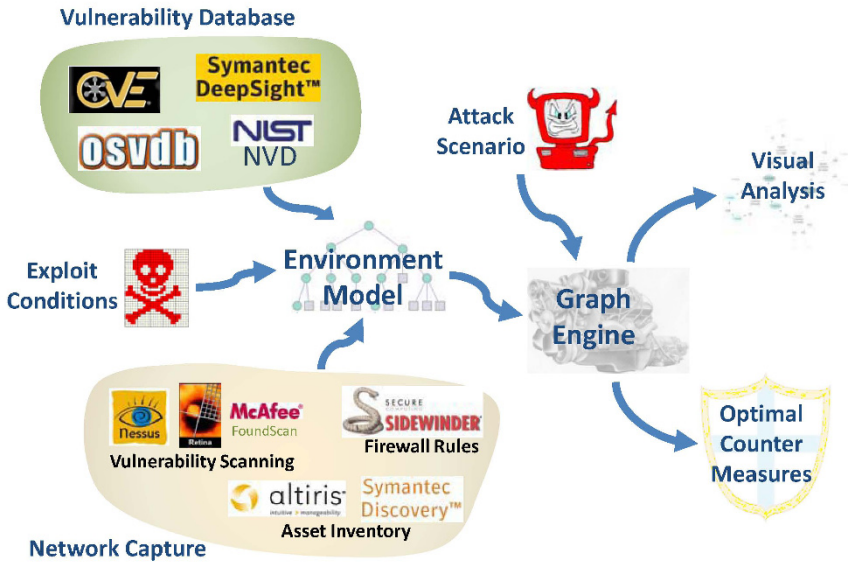


Fig. 7.1 Topological Vulnerability Analysis (TVA). Data from network scans and known vulnerabilities are combined into a model of the network security environment. Multi-step attack graph for this environment provides context for overall network security.

The *Exploit Specifications* encode how each vulnerability may be exploited (preconditions) and the result of its exploitation (postconditions). *Network Capture* represents data collection for a network to be defended, in terms of corresponding elements in *Vulnerability Reporting* and *Exploit Specifications*. Together, all these inputs are used to build an *Environment Model* for multi-step attack graph simulation.

In particular, the *Graph Engine* uses the *Environment Model* to simulate multi-step attacks through the network, for a given user-defined *Attack Scenario*. This engine analyzes vulnerability dependencies, matching exploit preconditions and postconditions, thus generating all possible paths through the network for a given attack scenario. The scenario may define particular starting and/or ending points for the attack, so that the graph is constrained to lie between them, or may be completely unconstrained (all possible starting and ending points).

The TVA tool provides sophisticated capabilities for interactive *Visual Analysis* of attack graphs [3]. It also computes *Optimal Counter Measures*, e.g., minimum number of network changes to thwart the attack scenario [4]. TVA attack graphs can also support proactive planning for optimal responses to attacks, based on known paths of residual vulnerability through the network. For

example, attack graphs can guide the placement of intrusion detection sensors to cover all attack paths, while minimizing sensors redundancy [5].

TVA attack graphs can filter false intrusion alarms, based on known paths of residual vulnerability. They also provide the context for correlating isolated alarms as part of a larger multi-step attack penetration [6]. The attack graph shows the next possible vulnerabilities that could be exploited by an attacker. This in turn supports optimal planning and response against attacks, while minimizing effects of false alarms and attacker misdirection.

As shown in Figure 7.1, our TVA tool integrates with vulnerability scanners Nessus [7], Retina [8], and FoundScan [9] for populating its network model. TVA processes data from the Sidewinder firewall [10] to capture network connectivity to vulnerable host services. TVA also integrates with host-based asset inventory technology, such as Centennial Discovery [11] and Symantec Altiris [12].

TVA matches the network model against a database of reported vulnerabilities. There are a number of such vulnerability databases, maintained by the government, commercial companies, and the security community. Examples include NIST's National Vulnerability Database (NVD) [13], the Bugtraq security database [14], Symantec DeepSight [15], the Open Source Vulnerability Database (OSVDB) [16], and the Common Vulnerabilities and Exposure (CVE) referencing standard [17]. We can thus leverage a storehouse of knowledge gathered by security researchers around the world, rather than being limited to vulnerabilities detected by a single tool like Nessus.

7.3 Illustrative Example

In Figure 7.2, a network is separated from the Internet by a firewall. The network is divided into 3 subnets, with one host in each subnet: a DMZ web server, an internal client, and an internal server. We wish to generate an attack graph showing whether an attacker can compromise the internal server from the Internet.

The DMZ web server is running Microsoft Windows Server, with Internet Information Services (IIS), Apache/MySQL/PHP, and Tomcat servlets. The client is running Microsoft Windows XP, client security software, an office productivity suite, and other utilities. The internal server is running Apache/MySQL/PHP, the Symantec Discovery asset inventory server, and Altiris Inventory Solution with associated software (e.g., IIS, Microsoft SQL Server). The Altiris Agent is deployed on each internal machine to collect asset inventory data.

The firewall blocks direct access to the internal server and client subnets from the Internet. Thus, from the outside, network vulnerability scanners such as Nessus are unable to detect any vulnerabilities on the internal server and client. In fact, if the firewall has network address translation (NAT) enabled, Nessus cannot even discover the existence of machines on these 2 subnets.

From the outside of the firewall, the only machine exposed is the DMZ server. In particular, the firewall blocks all traffic except HTTP to the DMZ server's TCP port 80. From behind the firewall, Nessus shows a variety of vulnerabilities on the DMZ server. But from the Internet, only the web server vulnerability is exposed. A Nessus scan from the DMZ to the internal subnets identifies any internal vulnerabilities permitted through the firewall. In our network, MySQL traffic is permitted between the DMZ web server and the internal server, and two exposed MySQL vulnerabilities allow an attacker to access the internal server (from the DMZ web server).

The question is whether an attacker can compromise the internal server from the Internet. Figure 7.3 shows the resulting attack graph for this network, using a TVA model populated by Nessus scans alone. This shows that an attacker starting on the Internet can first penetrate through the firewall and compromise the DMZ server, exploiting a vulnerability on its web server installation. Then, from the DMZ server, the attacker can access the internal server via exploitation of the two vulnerabilities.

A Nessus scan from the Internet would reveal no vulnerabilities on the internal server or client, so that there is no direct attack from the Internet to the internal machines. Because the firewall

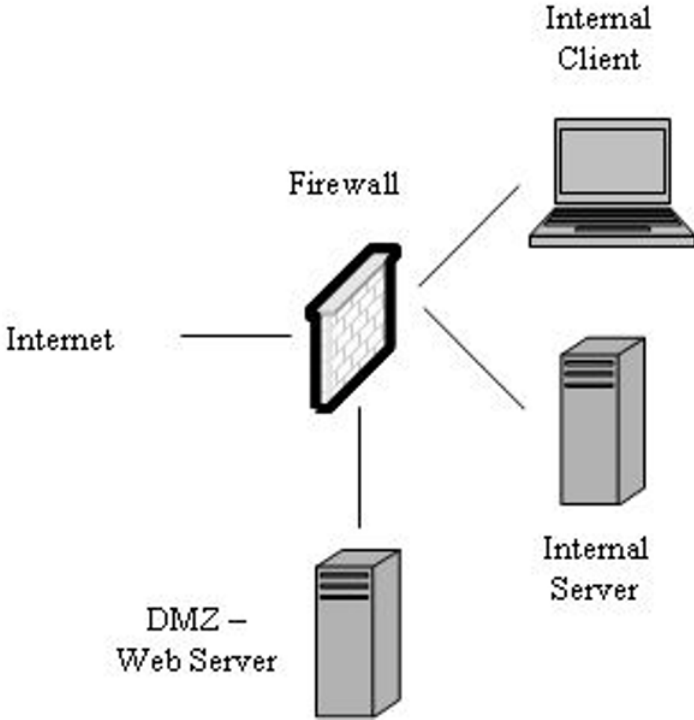


Fig. 7.2 Example network. Firewall is intended to protect network from Internet attackers.

blocks traffic originating from the server subnet to the client subnet, and blocks all traffic from the DMZ to the client subnet, there is no attack path to the client at all. That is, a Nessus scan from the DMZ to the client reveals no vulnerabilities. In this case, a Nessus scan of the client within its own subnet (no intervening firewalls) detects no vulnerabilities (or even open ports), because of the client’s personal firewall (part of the security suite).

We then augment the Nessus-based model, using software inventory data from Altiris. In particular, we compare products and versions from Altiris against a vulnerability database to determine the vulnerabilities associated with each application. For Microsoft products, this correlation process involves an extra step. We first determine how many vulnerabilities are associated with the Microsoft product and version installed, and then compare the Microsoft patches and hotfixes installed (collected by Altiris). This determines which Microsoft vulnerabilities are unpatched.

So, for example, while there are almost 200 vulnerabilities associated with Microsoft Windows XP SP2, many of the older vulnerabilities are patched on the client machine. On the other hand, no Microsoft hotfixes are applied on the client for Microsoft Office components after Office SP2. Thus all of the dozens of vulnerabilities associated with that version of Microsoft Office are relevant to the client.

These client-side vulnerabilities are associated with software that has no network service running. Still, these vulnerabilities represent vectors by which an attacker might obtain access to the client machine. Significant numbers of vulnerabilities are associated with web browsers and plug-ins. The typical scenario for these increasingly widespread vulnerabilities is that a user running a vulnerable web browser or plug-in visits a web site with malicious content that exploits the

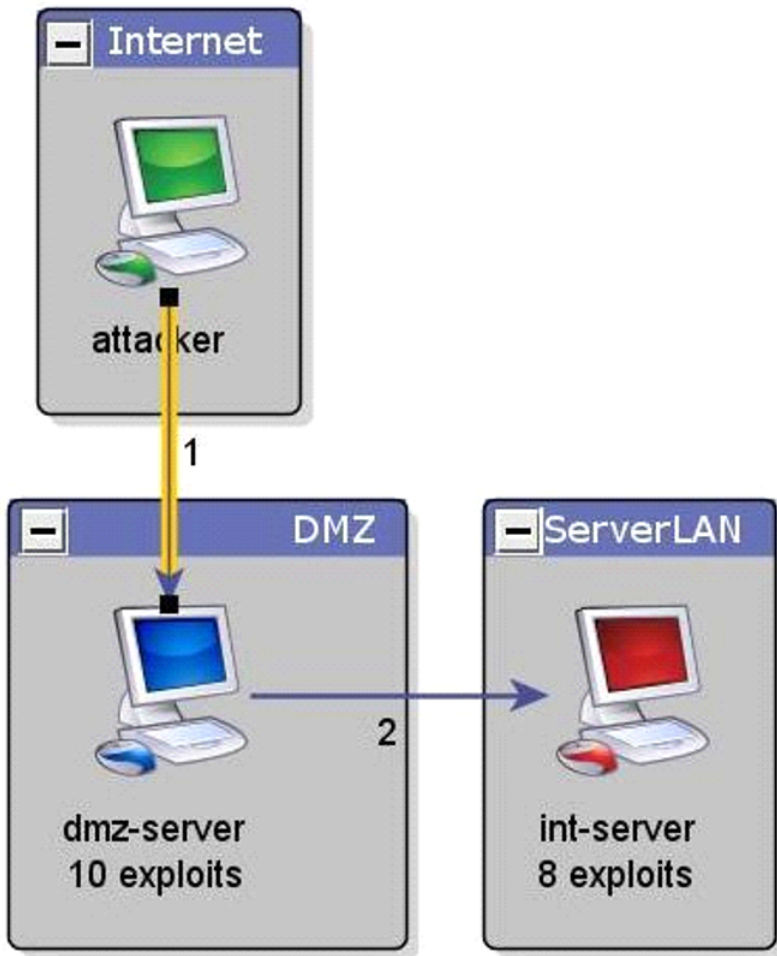


Fig. 7.3 Attack graph from internet attacker to internal server. The network model is populated from Nessus vulnerability scans.

client-side vulnerability. Another important class of vulnerabilities is associated with document processing applications, e.g., infected documents via e-mail.

Figure 7.4 shows the attack graph augmented with Altiris data. This shows that it is actually possible to attack the client directly from the Internet, via 12 different client-side vulnerabilities. Strictly speaking, the client needs to make an outbound connection (e.g., web site visit) to a compromised server. But the firewall allows this, so it is correct to model the server as the attacker and the client as the victim. The attacker on the client can then compromise the internal server, through a firewall hole allowing access from client to server.

This example illustrates the importance of accounting for such client-side vulnerabilities. Services on an internal server would typically be exposed to internal clients, and may be vulnerable. In this case, our augmented model has uncovered attack paths that would have otherwise been undetected by a vulnerability scanner.

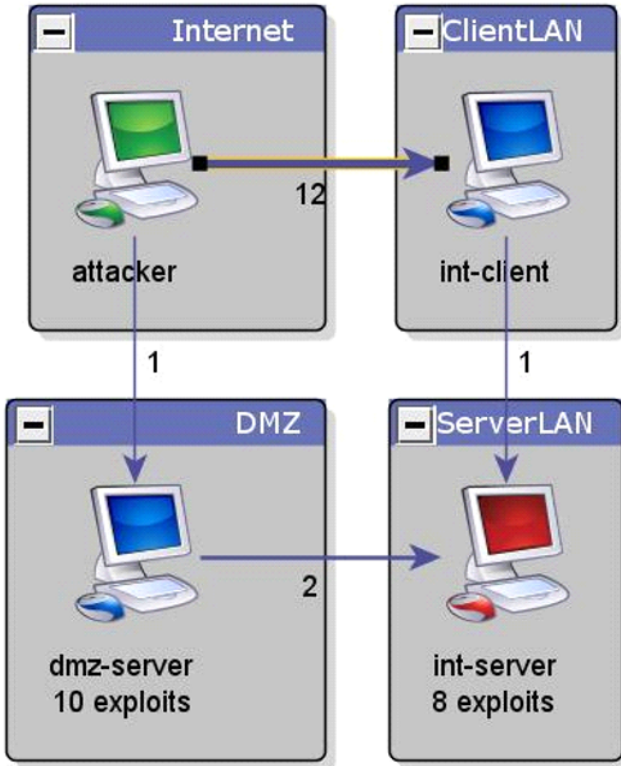


Fig. 7.4 Attack graph augmented with Altiris data. The host-based asset inventory uncovers additional client-side vulnerabilities.

Our attack graphs show how hosts on a network can be exploited through multiple steps, even when the attacker cannot access them directly. It is not directly possible to compromise the internal server from the outside because of the policy enforced by the firewall. But TVA shows that the attack goal can be reached indirectly, through two different attack paths. While it may be possible for an experienced analyst to find such paths manually for a small network, for enterprise networks an automated tool is needed.

7.4 Network Attack Modeling

In TVA, the environment model includes aspects of the network configuration relevant to attack penetration, as well as a set of potential attacker exploits that match attributes of the configuration. The TVA approach can apply to many different types of attack models (even non-cyber models) as long as a common schema is employed across the model.

Typically, we model a network as machines and collections of machines into protection domains. Protection domains capture the idea that the set of machines in a domain have (implicitly) unrestricted access to one another's vulnerable services. A machine may include attributes for modeling network attack penetration, such as operating system, application programs, and connections to vulnerable services.

A connection describes how a machine connects to potentially vulnerable services across the network, to ports on other machines or to its own ports. This mirrors the Transmission Control Protocol/Internet Protocol (TCP/IP) reference model, in which a layered connectivity structure represents the various network architectures and protocols [18].

To keep pace with emerging threats, we continually monitor sources of reported vulnerabilities, and add those to our database of modeled TVA exploits. We model an attacker exploit in terms of preconditions and postconditions, for generic attacker and victim machines, which are subsequently mapped to the target network.

The detection of a host vulnerability (e.g., via Nessus) represents a precondition for a corresponding TVA modeled exploit. In reality, the vulnerability exists because of a particular combination of software components. A detected vulnerability is thus a convenient higher level abstraction, in comparison to the combination of components that causes it. There are usually many possible software components and configurations that can cause a particular vulnerability.

We can thus model vulnerable software components themselves in TVA. In particular, host-based asset management solutions like Centennial Discovery and Symantec Altiris collect specific information such as operating system, service packs, and installed software through on-host agents. Many exploits work only with specific versions of operating systems and/or applications, and vulnerabilities can be removed by applying the appropriate patches.

To keep TVA exploits continually updated as new vulnerabilities are reported, we may leverage the component-to-vulnerability mapping information provided by Symantec DeepSight. Also, while TVA incorporates exploits against known vulnerabilities into its analysis, it is a general modeling approach that can also accommodate unknown vulnerabilities. Anything that can be described in terms of preconditions and postconditions can be included in the model. Also, tools such as Nessus generate many alerts that are merely informational, i.e., irrelevant to network penetration. We carefully exclude these from our database of modeled exploits.

One important class of vulnerabilities detected by our approach is client applications. Client-side vulnerabilities have a major impact on enterprise network security posture, a trend that will no doubt continue. For example, web applications (an important source of client-side vulnerabilities) represent about 60% of vulnerabilities documented in 2007 [19].

In our approach, the correctness of attack graphs depends on the quality of the input datasets. The network scanning tools must have a complete and current database of detected software and/or vulnerabilities. Standardization efforts such as Security Content Automation Protocol (SCAP) [20] help facilitate the management of mapping data between asset inventory and vulnerabilities. Otherwise, each tool may use its own naming conventions. Within SCAP, the Common Platform Enumeration (CPE) [21] provides standard naming of operating systems and applications, and CVE provides mappings to CPE for vulnerable software. The OVAL Language [22] standardizes the vulnerability assessment process and results. Adoption of such standards among tool vendors will greatly facilitate the building of attack models for TVA.

The purpose of modeling network configuration in TVA is to support preconditions of modeled attacker exploits. Machines are explicitly included in the model only if they offer services over the network that can be exploited. The effects of other network devices such as routers and firewalls is captured implicitly, in the way that connectivity is modeled at the various network layers, i.e., in the way that they provide/restrict connectivity to vulnerable services on other machines.

For remote (versus host-based) scanners such as Nessus, Retina, and FoundScan, we can capture the implicit effects of devices such as firewalls by scanning from different network vantage points, targeting hosts through the firewall. We can then combine multiple scans from various network locations, building a complete map of connectivity to vulnerable services throughout the network. Alternatively, we can analyze firewall rules directly, adding the resulting vulnerable connections to the model, eliminating the need for scanning through firewalls.

7.5 Analysis and Visualization

Based on a given attack scenario, the attack graph may be constrained by specific starting and ending points. The scenario may also be less constrained, such as finding all possible attack starts leading to one or more goals, or finding all possible paths from particular starting points. For example, one may wish to know how a particular critical system can be compromised from all possible starting points. Or, one may want to know all systems that could be compromised from a particular starting point, or even from all possible starting points. Our TVA tool supports each combination of specified/unspecified attack start/goal.

In their raw form, attack graphs can be much too complex for easy understanding. Figure 7.5 is an example of such a raw attack graph, for a network of only 20 hosts in four subnets. Computational complexity for this graph is quadratic in the number of hosts, and we compute it in a fraction of a second. Still, it is difficult to for an analyst to understand its complexity. To help manage this, we aggregate attack graphs, providing better situational awareness at higher levels of abstraction. We define multiple levels of abstraction by recursively aggregating elements from one level to the next, forming a hierarchy of abstraction levels.

Figure 7.6 illustrates our hierarchy of attack graph abstractions. From the bottom, the preconditions or postconditions for an exploit are aggregated into a condition set. This provides the higher-level viewpoint that exploits depend on one another, without showing the particular conditions that cause the dependency.

Sets of conditions for a machine are aggregated into a machine abstraction, so that one can consider a machine without knowing its particular conditions. Exploits, with their preconditions and postconditions, are aggregated if they involve the same pair of attacker/victim machines. This shows that there are exploits from one machine to another, while hiding the details about them.

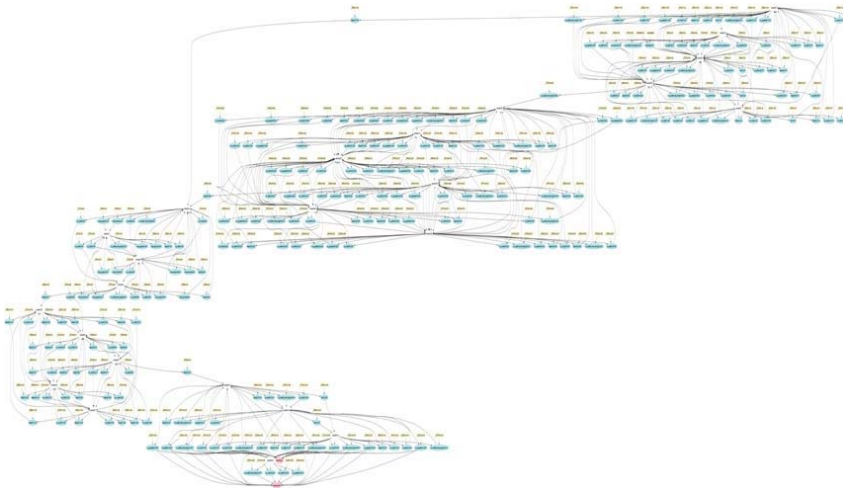


Fig. 7.5 Raw attack graph. This graph is computed quickly, but cannot be quickly understood.

An important high-level abstraction in TVA is the protection domain, which represents a set of machines that have full access to one another’s vulnerabilities. In general, connectivity can be

restricted across a network through a variety of mechanisms, such as firewalls, trust relationships, etc. But within a protection domain, implicitly there are no connectivity limitations among machines. Thus all of a machine’s vulnerabilities are potentially exploitable from anywhere within the protection domain.

In a raw non-aggregated form, the attack graph is fully connected within a protection domain. Through this implicit full connectivity, we employ a much more efficient representation. The machines in a domain are listed, along with exploits against each of their vulnerabilities. Then implicitly, once an attacker takes control of a machine within a protection domain, he can exploit all vulnerabilities on machines within it. In this way, we avoid explicitly listing every n^2 (fully-connected) exploit dependency within the protection domain. Thus, within each domain, complexity is linear in the number of hosts.

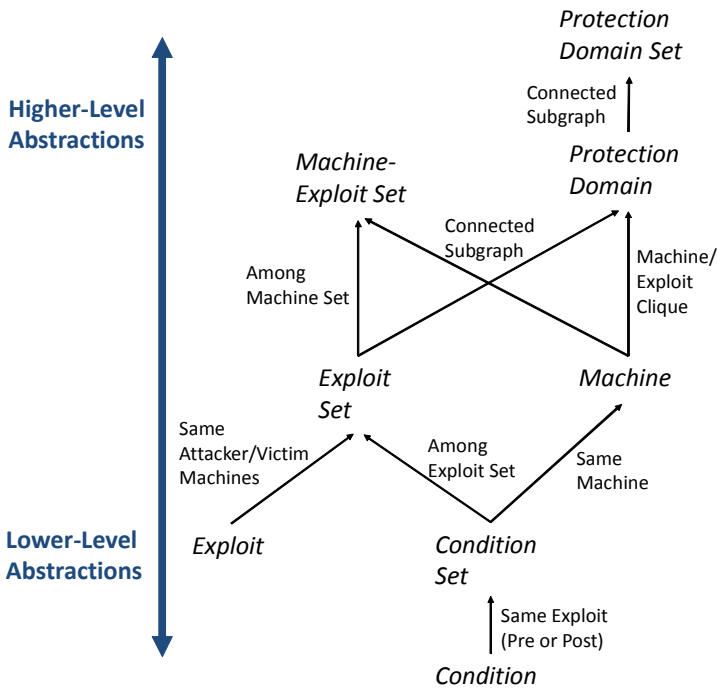


Fig. 7.6 Levels of attack graph aggregation. Lower-level details are aggregated into progressively summarized sets.

Figure 7.7 demonstrates the effectiveness of our graph aggregation approach. This is an attack graph for a network of six subnets, with 150 hosts. In its raw form, this attack graph is actually much more complex than Figure 7.5. But through our graph aggregation, the flow of the attack paths is clear. Individual machines appear within protection domains, and are visible when a domain is expanded. Graph edges from one machine to another represent the set of exploits from attacker to victim machine. Graph edges between domains are aggregated exploit sets between domain pairs.

Our TVA software tool includes a full-featured attack graph visualization interface. In this interface, a high-level view clearly displays attack relationships among protection domains, which

can be opened individually or in groups for deeper views of attack properties and relationships. In this process, no graph information is lost; one has merely to expand a folder to acquire information at a lower level.

A complete listing of exploits and associated details for any selected component is available at all times. This supports in-depth analysis of exploit details, while overall topology and network relationships are kept simple and understandable within the main graph view.

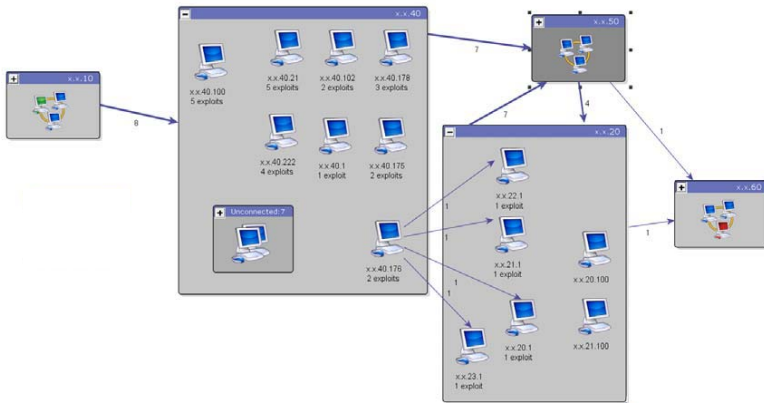


Fig. 7.7 Aggregated attack graph. This is much easier to understand overall attack flow.

Our TVA tool also emulates the hardening of machines and exploitable vulnerabilities to study the effects of remediation and what-if scenarios. Exploring the attack graph, the analyst is often faced with multiple options for remediation. This involves choosing a machine or set of machines to protect (harden), or identifying specific exploits to protect against.

We display the attack graph effects that occur when a specific machine or protection domain is hardened or when a specific exploit is neutralized. Hardened elements are maintained in a log, e.g., for reporting. The TVA tool also generates recommendations automatically, i.e., first layer (from start), last layer (from goal), and minimum set that that separates start from goal.

To aid user navigation, the TVA tool maintains a global overview of the entire attack graph at all times, which can be used to pan the main graph view. The tool also has a graphical (tree) attack dictionary of all graph elements. The various graph views are linked, so that selecting an element in one view cause it to be selected in all view. A variety of toolbars are available for commonly used tools. This includes a suite of interactive layout tools, with manual repositioning as well as full-scale layout algorithms, continuously available to restructure the display.

7.6 Scalability

For computing attack graphs for larger networks, we need scalable mathematical representations and algorithms. We assume the attacker’s control over the network increases monotonically over time. This corresponds to the conservative assumption that once an attacker gains control of a network resource, there is no need to relinquish it to further advance the attack.

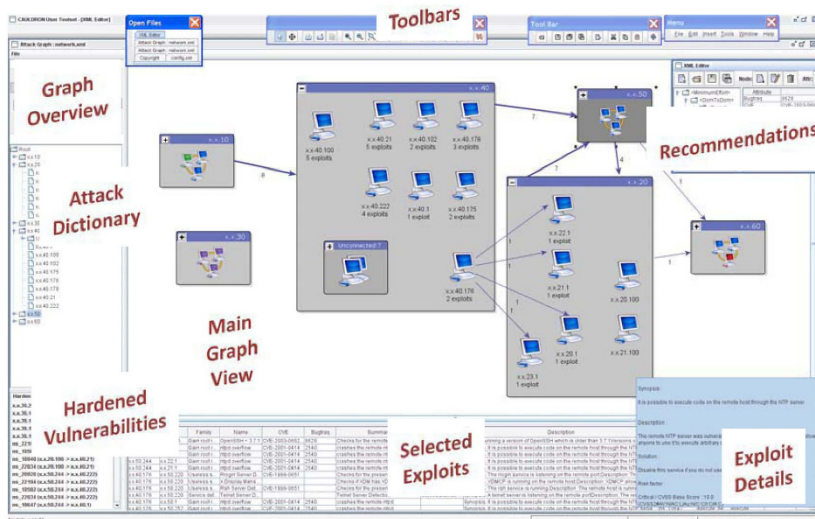


Fig. 7.8 Attack graph visualization. Supports graph navigation with high-level overviews and detail drilldowns.

Under this monotonicity, it is sufficient to represent the dependencies among exploits, rather than explicitly enumerating every sequence of exploits [23]. The resulting exploit-dependency attack graphs grow quadratically rather than exponentially [24]. In particular, worst-case complexity for n network hosts is $O(n^2)$. By grouping hosts into protection domains, complexity is reduced to $O(n)$ within each domain. In terms of the database of potential attacker exploits, complexity is $O(e)$, for e exploits.

Figure 7.9 shows attack graph computation times for networks of various sizes. In each case, a subnet contains 200 hosts, and each host has 5 vulnerabilities. Each subnet has incoming vulnerable connections from 2 other subnets, and symmetrically, outgoing vulnerable connections to 2 other subnets. This is a ring topology, in which the number of network connections grows linearly with the number of subnets.

From one subnet to another, there are 500 connections to vulnerabilities in the victim subnet. Thus there are $2 * 500=1,000$ incoming and $2 * 500=1,000$ outgoing vulnerable connections (a grand total of 2,000) for each subnet. Computation times (total run time in seconds) are based on increasing numbers of subnets, from 20 subnets (4,000 hosts) to 200 subnets (40,000 hosts). Run times are for a quad-core Intel Xeon CPU at 1.86 GHz, with 4 GB RAM.

In this experiment, overall network size (number of vulnerable connections) grows linearly with the number of subnets (and hosts). This clearly shows how graph generation time depends proportionally on the size of the input network. This excludes any time for generating the input model (network and exploits) itself, although this has the same worst-case complexity, and these can be pre-computed.

Graph visual layout performance is a separate issue, and is not included in the execution times of Figure 7.9. For example, Figure 7.10 shows a layout of the 20-subnet case (4,000 hosts) of Figure 7.9. Computing this layout only a few seconds, but computing layout for the 100-subnets case (20,000 hosts) takes 14 minutes. But in all cases, once the initial layout is computed, performance of user interaction (repositioning, drilldown, etc.) is immediate.

Visual layout computation is needed for a cyber view of network attack graphs. Such layout induces spatial coordinates onto an abstract information graph. But in some situations, understanding the actual physical location of attacks may be important. We can embed the attack graph in a geo-spatial visualization, as illustrated in Figure 7.11. When spatial coordinates are given, no

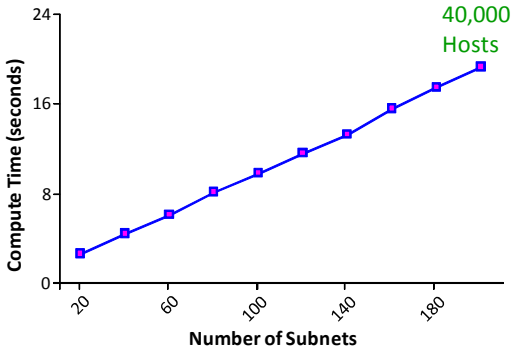


Fig. 7.9 Computation time for attack graph generation. Network complexity grows linearly with the number of subnets. Attack graph generation time is proportional to the size of the network.

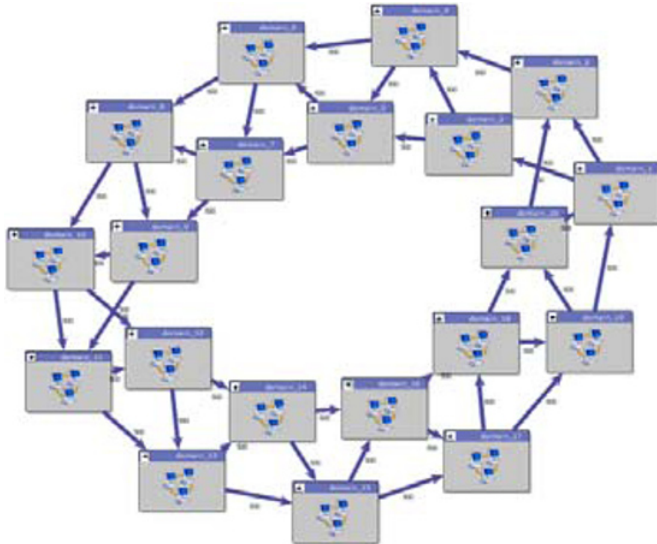


Fig. 7.10 Attack graph layout solution. For this 20-subnet case (4,000 hosts), graph visual layout is computed in only a few seconds.

additional graph layout computation is needed. In such cases, visualizing complex attack graphs is much faster than for abstract cyber views.

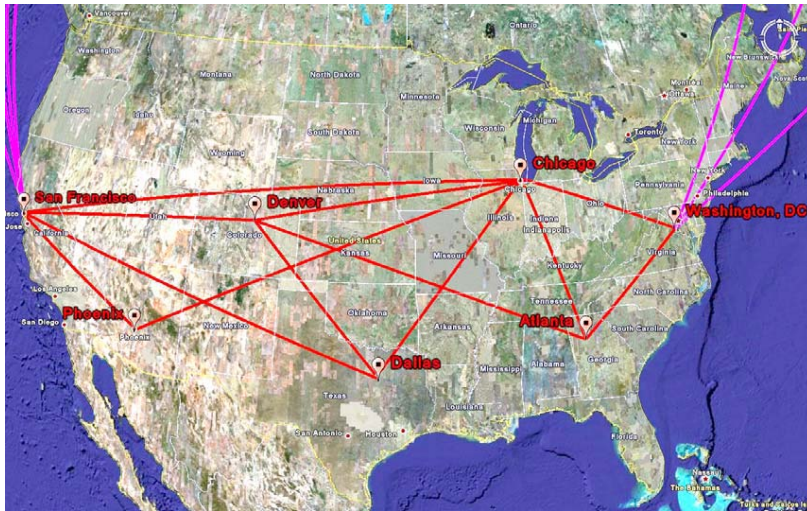


Fig. 7.11 Geo-spatial attack graph visualization. When spatial coordinates are given, computational requirements are reduced.

7.7 Related Work

In early attack graph formalisms, algorithmic complexity is exponential [25][26][27] [28] because paths are explicitly enumerated. Under reasonable assumptions attack graph analysis can be formulated as monotonic logic, making it unnecessary to explicitly enumerate states. This leads to polynomial rather than exponential complexity [29]. Our protection domain abstraction reduces complexity further, to linear within each domain [30]. Complexity can be further reduced based on host configuration regularities [31].

Attack graph research has largely focused on scalability, with relatively little work on aspects of model population. Notable exceptions include [18][32][33], although these are more theoretical frameworks than practical model population. Commercial capabilities for attack graph analysis remain limited, especially in the area of visualization for large-scale graphs [34][35]. A review of attack graph research circa 2005 is given in [36].

7.8 Summary

TVA shows paths of vulnerability allowing attackers to penetrate through a network. It identifies critical vulnerabilities and provides strategies for protection of critical network assets. This allows

us to harden networks before attacks occur, and handle intrusion detection more effectively, and respond appropriately to attacks.

TVA captures the network configuration, including software, their vulnerabilities, and connectivity to vulnerable services. It then matches the network configuration against a database of modeled attacker exploits for simulating multi-step attack penetration.

During simulation, the attack graph can be constrained according to user-defined attack scenarios. From the resulting attack graphs, TVA computes recommendations for optimal network hardening. It also provides sophisticated visualization capabilities for interactive attack graph exploration and what-if analysis. TVA employs algorithms with worst-case quadratic complexity in the number of network hosts.

Further, TVA attack graphs then provide the necessary context for correlating and prioritizing intrusion alerts, based on known paths of vulnerability through the network. Overall, TVA offers powerful capabilities for proactive network defense, transforming raw security data into actionable intelligence.

Acknowledgements

This material is based upon work supported by Homeland Security Advanced Research Projects Agency under the contract FA8750-05-C-0212 administered by the Air Force Research Laboratory/Rome; by Air Force Research Laboratory/Rome under the contract FA8750-06-C-0246; by Federal Aviation Administration under the contract DTFWA-08-F-GMU18; by Air Force Office of Scientific Research under grant FA9550-07-1-0527 and FA9550-08-1-0157; and by the National Science Foundation under grants CT-0716567, CT-0716323, and CT-0627493. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsoring organizations.

References

- [1] S. Jajodia, S. Noel, and B. O’Berry, “Topological Analysis of Network Attack Vulnerability,” in *Managing Cyber Threats: Issues, Approaches and Challenges*, V. Kumar, J. Srivastava, A. Lazarevic (eds.), Kluwer Academic Publisher, 2005, pages 248–266.
- [2] S. Jajodia, S. Noel, “Topological Vulnerability Analysis: A Powerful New Approach for Network Attack Prevention, Detection, and Response,” in *Algorithms, Architectures and Information Systems Security (Indian Statistical Institute Platinum Jubilee Series)*, B. B. Bhattacharya, S. Sur-Kolay, S. C. Nandy, A. Bagchi, eds., World Scientific, New Jersey, 2009, pages 285–305.
- [3] S. Noel, M. Jacobs, P. Kalapa. S. Jajodia, “Multiple Coordinated Views for Network Attack Graphs,” in *IEEE Workshop on Visualization for Computer Security (VizSEC2005)*, Minneapolis, MN, October, 2005, pages 99–106.
- [4] L. Wang, S. Noel, S. Jajodia, “Minimum-Cost Network Hardening Using Attack Graphs,” *Computer Communications*, 29(18), 2006, pages 3812–3824.
- [5] S. Noel, S. Jajodia, “Optimal IDS Sensor Placement and Alert Prioritization Using Attack Graphs,” *Journal of Network and Systems Management*, 16(3), 2008, pages 259–275.
- [6] S. Noel, E. Robertson, S. Jajodia, “Correlating Intrusion Events and Building Attack Scenarios through Attack Graph Distances,” in *Proceedings of the 20th Annual Computer Security Applications Conference (ACSAC)*, 2004, pages 350–359.
- [7] R. Deraison, *Nessus*, <http://www.nessus.org>.
- [8] eEye Digital Security, *Retina Network Security Scanner*, <http://www.eeye.com/html/Products/Retina/index.html>.

- [9] Foundstone, *FoundScan Frequently Asked Questions*, <http://www.foundstone.com/us/index.asp>.
- [10] Secure Computing, Sidewinder Firewall Device, <http://www.securecomputing.com/>.
- [11] Centennial Software, Discovery Asset Management, <http://www.centennial-software.com/>.
- [12] Symantec, Altiris, <http://www.altiris.com/>.
- [13] NIST, National Vulnerability Database (NVD), <http://nvd.nist.gov/>.
- [14] Security Focus, *Bugtraq Vulnerabilities*, <http://www.securityfocus.com/vulnerabilities>.
- [15] Symantec Corporation, *Symantec DeepSight Threat Management System*, <https://tms.symantec.com/Default.aspx>.
- [16] Open Source Vulnerability Database, <http://osvdb.org/>.
- [17] MITRE Corporation, *CVE - Common Vulnerabilities and Exposures*, <http://cve.mitre.org/>.
- [18] R. Ritchey, B. O'Berry, S. Noel, "Representing TCP/IP Connectivity for Topological Analysis of Network Security," in *Proceedings of the 18th Annual Computer Security Applications Conference (ACSAC)*, 2002, pages 156–165.
- [19] D. Turner, M. Fossi, E. Johnson, T. Mack, J. Blackbird, S. Entwisle, M. K. Low, D. McKinney, C. Wueest, *Symantec Global Internet Security Threat Report Trends*, 2008.
- [20] NIST, Security Content Automation Protocol (SCAP), <http://nvd.nist.gov/scap.cfm>.
- [21] MITRE, Common Platform Enumeration (CPE), <http://cpe.mitre.org/>.
- [22] MITRE, Oval Language, <http://oval.mitre.org/>.
- [23] P. Ammann, D. Wijesekera, S. Kaushik, "Scalable, Graph-Based Network Vulnerability Analysis," in *Proceedings of the 9th ACM Conference on Computer and Communications Security*, Washington, DC, pages 217–224.
- [24] S. Noel, J. Jajodia, "Understanding Complex Network Attack Graphs through Clustered Adjacency Matrices," in *Proceedings of the 21st Annual Computer Security Applications Conference (ACSAC)*, 2005, pages 160–169.
- [25] D. Zerkle, K. Levitt, "Netkuang: A Multi-Host Configuration Vulnerability Checker," in Proceedings of the 6th USENIX Unix Security Symposium, 1996.
- [26] R. Ritchey, P. Ammann, "Using Model Checking to Analyze Network Vulnerabilities," in Proceedings of the IEEE Symposium on Security and Privacy, 2000.
- [27] L. Swiler, C. Phillips, D. Ellis, S. Chakerian, "Computer-Attack Graph Generation Tool," in Proceedings of the DARPA Information Survivability Conference & Exposition II, 2001.
- [28] O. Sheyner, J. Haines, S. Jha, R. Lippmann, J. Wing, "Automated Generation and Analysis of Attack Graphs," in *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA.
- [29] R. Lippmann, K. Ingols, C. Scott, K. Piwowski, K. Kratkiewicz, M. Artz, R. Cunningham, "Validating and Restoring Defense in Depth Using Attack Graphs," in Proceedings of the MILCOM Military Communications Conference, 2006.
- [30] S. Noel, S. Jajodia, "Managing Attack Graph Complexity through Visual Hierarchical Aggregation," in *Proceedings of the ACM CCS Workshop on Visualization and Data Mining for Computer Security* Fairfax, Virginia.
- [31] W. Li, An Approach to Graph-Based Modeling of Network Exploitations, PhD dissertation, Department of Computer Science, Mississippi State University, 2005.
- [32] F. Cuppens, R. Ortalo, "LAMBDA: A Language to Model a Database for Detection of Attacks," in 3rd International Workshop on Recent Advances in Intrusion Detection, 2000.
- [33] S. Templeton, K. Levitt, "A Requires/Provides Model for Computer Attacks," in New Security Paradigms Workshop, 2000.
- [34] Skybox Security, <http://www.skyboxsecurity.com/>.
- [35] RedSeal Systems, <http://www.redseal.net/>.
- [36] R. Lippmann, K. Ingols, An Annotated Review of Past Papers on Attack Graphs, Lincoln Laboratory, Technical Report ESC-TR-2005-054, 2005.