14

The Cascade Vulnerability Problem for Open Distributed Systems: A Review

Stefanos Gritzalis
Department of Informatics
University of Athens
TYPA Buildings, Athens GR-15771, Greece
tel.: +30-1-7291885, fax: +30-1-7219561
Department of Informatics
Technological Educational Institute (T.E.I.) of Athens
Ag.Spiridonos St. Aegaleo GR-12210, Greece
tel.: +30-1- 5910974, fax.: +30-1-5910975
email: sgritz@teia.ariadne-t.gr

Sokratis K. Katsikas
Department of Mathematics
University of the Aegean
Samos GR-83200, Greece
tel.: +30-273-33919, fax: +30-273-35483
email: ska@aegean.gr

Diomidis Spinellis
Department of Mathematics
University of the Aegean
Samos GR-83200, Greece
tel.: +30-273-33919, fax: +30-273-35483
email: dspin@aegean.gr

Abstract

The Cascade Vulnerability Problem is a potential problem which must be faced when using the interconnected accredited system approach of the Trusted Network Interpretation. In this paper, we present the general Cascade vulnerability problem, describe the basic properties of the most important detection algorithms , and conduct a brief comparative analysis.

D. Gritzalis (ed.), *Reliability, Quality and Safety of Software-Intensive Systems*

Keywords
Cascade Vulnerability Problem, Network & Open Distributed Systems Security, Risk Analysis.

1 INTRODUCTION

The Cascade Vulnerability Problem was discussed in the Trusted Network Interpretation (NCSC, 1987) of the Trusted Computer System Evaluation Criteria. According to (NCSC, 1987), a Cascade Vulnerability Problem exists when a penetrator can take advantage of network connections to compromise information across a range of security levels that is greater than the accreditation range of any of the component systems one must defeat to do so.

In a distributed system with many nodes and interconnections, the existence of a Cascade Vulnerability Problem may not be obvious and can cause serious security problems (Katsikas, 1996). In this paper we present the most effective algorithms — published in the open literature — for the detection of the Cascade vulnerability Problem and the identification of the paths along which the problem exists. Then, we outline the basic semantics of an algorithm for the Restricted Cascade Correction Problem which proposes network modifications for reducing the risk of Cascade Vulnerability. Finally, we conduct a brief comparative analysis of the presented algorithms.

2 THE CASCADE VULNERABILITY PROBLEM

The Cascade Vulnerability Problem belongs to a subspace of the problem set that addresses the issue of whether the interconnection of secure systems via a secure channel results in a secure distributed system.

The term "*secure system*" is taken here to mean a system that has undergone not only a risk analysis evaluation with respect to the acceptable risk of operating the system, but a system security evaluation as well.

The assets of the system and the threats against each one of them are considered during the risk analysis review in order to identify the level of the security required. System security can be modelled as a function of many parameters, such as computer security, communications security, administrative security, personnel security, and physical security (Madron, 1990). For implementation purposes all these parameters must be categorised into classes of countermeasures that reduce the system risks. Therefore, a system security evaluation assesses the effectiveness of the countermeasures which were finally selected for a specific system, at a given time.

The Cascade Vulnerability Problem appears in the subset of networks that cannot be treated as a single system. There are different reasons why networks cannot be viewed as a single system. The main reasons can be:

- the large size of the network
- different administrative entities which may lead to different risk assessment methods.

In any case, it is necessary for the administrators of any two systems that are to be interconnected to mutually agree that both systems are secure as stand-alone systems; that is, both administrators need to accept the risk assessment and the security evaluation methods which are used for both systems.

In summary, one can argue that (Fitch, 1991) (Fitch, 1993) the Cascade Vulnerability Problem appears when independent mutually recognised secure systems are interconnected by secure channels to create a distributed system which is not as secure as its parts. In other words (Millen, 1988) the Cascade Vulnerability Problem appears when an adversary can take advantage of network connections to compromise information across a range of sensitivity levels that is greater than the accreditation range of any of the component systems s/he must defeat to do so.

As a typical example of the Cascade Vulnerability Problem (NCSC, 1987), let us consider two systems, as shown in Figure 1. Host A is accredited for TS-Top Secret and S-Secret information and all users are cleared to at least the Secret level. Host B is accredited for S and C-Confidential and all users are cleared to at least the Confidential level; finally, there is a link at level S between the two systems.

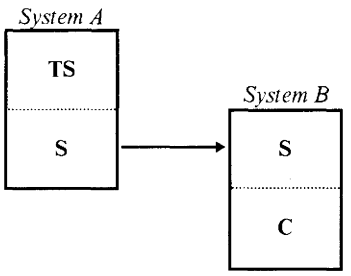

Figure 1. The generic Cascade Vulnerability Problem.

While the risk of compromise in each of these systems is small enough to justify their use with two levels of information the system as a whole has three levels of information. This increases the potential harm that an adversary could cause, since s/he could downgrade the TS-level information in system A to S-level, send it to system B, and further downgrade the information to C-level therein.

The adversary has to defeat the protection mechanisms of both systems A and B, but that is an easier job than defeating the protection mechanisms of a single system trusted to protect the whole range from TS-level to C-level. The network connection has, in essence, created a Trusted Computing Base (TCB) with users cleared to at least the C-level with data on it at the TS-level.

In this way (Millen, 1988) the network connection has invalidated the risk analysis that accredited the two systems, because such a networked system must have a more secure architecture, a TCB rating of B3, than either rating of the original individual sub-systems TCB (i.e. B1 or B2, Figure 2).

	Maximum Data Sensitivity							
Minimum Clearance or authorisation of System Users		U	N	C	S	TS	IC	MC
	U	C1	B1	B2	B3	*	*	*
	N	C1	C2	B2	B2	A1	*	*
	C	C1	C2	C2	B1	B3	A1	*
	S	C1	C2	C2	C2	B2	B3	A1
	TS(B1)	C1	C2	C2	C2	C2	B2	B3
	TS(SB1)	C1	C2	C2	C2	C2	B1	B2
	IC	C1	C2	C2	C2	C2	C2	B1
	MC	C1	C2	C2	C2	C2	C2	C2

Figure 2. Security Index Matrix for Open Environments (NCSC, 1985).

Let $R_j(t)$ be the probability that both TCBs can be penetrated if the joint combination of two TCBs is subject to a total threat of t units or less. Changing variables and taking into account that the probability of two independent events occurring together is the product of their separate probabilities (Freund, 1962), the value of R_j can be then computed as the convolution integral:

$$R_j(t) = \int_{-\infty}^{\infty} R_{B2}(x)\, R_{B2}(t-x)\, dx,$$

whose precise value in relation to the original $R_{B2}(x)$ is not intuitively obvious.

In Figure 3a and 3b the probability density functions for $R_{B2}(t)$ and of the Cascade $R_j(t)$ are shown.

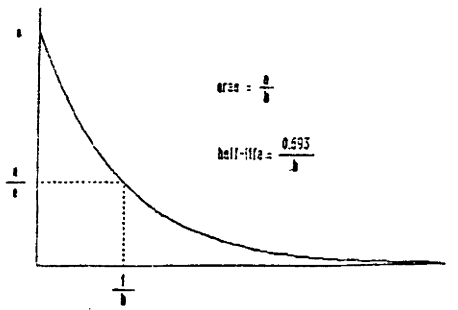

Figure 3a. $R_{B2}(t) = ae^{-bt}$

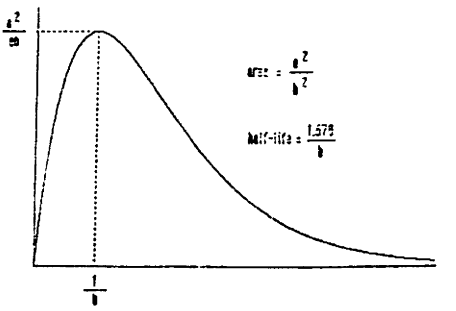

Figure 3b. $R_j(t) = a^2 t\, e^{-bt}$

It has been proven (Lee, 1989) that R_j is approximately equal to R^2_{B2} for the cascade of two B2 systems. This means that the resistance to threat of a cascade of two B2 systems is approximately the same as, or even better than, that of a B3 system.

3 ALGORITHMS FOR CASCADE VULNERABILITY DETECTION

3.1 The Nesting and Cascade condition

The Nesting condition

The simplest approach for recognising a potential Cascade Vulnerability Problem is to test whether a network can or cannot face such a problem. This simple test is calling *the nesting condition* (NCSC, 1987).

The nesting condition is true if the accreditation ranges of each of the interconnected systems are either:

- nested - one range is included in the other
- disjoint - have no common level.

Fulfilment of the nesting condition implies that there can be no Cascade Vulnerability Problem in the network at hand. However, there are many cases in the literature (Millen, 1988) where the nesting condition is not fulfilled, yet there is actually no Cascade Vulnerability Problem.

A possible solution when the problem may exist is to eliminate certain network connections, either physically or by means of end-to-end encryption. The later solution allows hosts that need to communicate to do so, while eliminating additional unnecessary cascading risk on the path from one host to another.

The Cascading Condition

An attempt for a formal description of the *Cascading Condition*, which is more precise than the one described in (NCSC, 1987), is presented in (Millen, 1988). According to that, when we use a network, we know that it consists of some nodes h, and that every node has its accreditation range $A(h)$. This $A(h)$ consists of a set of sensitivity levels which, as a whole, form a lattice. Consequently, an accreditation range is a convex sublattice which is the formal notion corresponding to a range.

The protection regions are the pairs (h,s), for each sensitivity level $s \in A(h)$. A step is an ordered pair of protection regions (h_1, s_1), (h_2, s_2) such that either:

- $s_1=s_2$ and h_1 sends information to h_2 at level s_1 - a network link, or
- $h_1=h_2$ - an information flow within a component.

In the second case, if also $s_1 \leq s_2$, then the information transfer is permitted by the enforced security policy in this specific node. A downgrade is a flow such that $s_1 > s_2$.

A downgrade from s_1 to s_2 is always associated with a risk index $R(s_1,s_2)$. If $s_1 \leq s_2$, then $R(s_1,s_2)=0$, otherwise $R(s_1,s_2)>0$. The risk index of any convex sublattice can be defined as the least upper bound of all $R(s_i,s_j)$.

Two accreditation ranges - convex sublattices are:

- nested, if either $A \in B$ or $B \in A$,
- strictly ordered, if (for every $a \in A$ and $b \in B$) then $a<b$, *or*
- incomparable, if for every ($a \in A$ and $b \in B$) neither $a \leq b$ nor $b \leq a$.

For a certain path (h_1,s_1), (h_2,s_2),..., (h_n,s_n), its net downgrade is $R(s_1.s_n)$, and its difficulty is $max(R(A(h_i)))$, such that $s_i > s_{i+1}$.

According to the above formalism we can argue that a path is a Cascading path if its difficulty is at least as great as its net downgrade. Therefore, a network satisfies the Cascade condition if it has no Cascading paths at all.

In (Millen, 1988) one can find a program, written in Edinburgh Prolog, that can identify all cascading paths based on the previous formalism.

3.2 A heuristic procedure

The heuristic condition is a less conservative but much more complex heuristic that takes into account the connectivity of the network and the evaluation classes of the components. Given the goal of not allowing a risk greater than is recommended by the Environmental Guidelines, the *heuristic procedure* (NCSC, 1987) has been developed to examine systems and determine whether they fall within the bounds prescribed by these Guidelines.

In formal terms the heuristic procedure is an approximate test for the Cascade Condition, described in the previous section. It should be noted that this procedure is not intended to be prescriptive: it is merely one way of examining the problem.

It is obvious that the heuristic procedure — as every heuristic — has been derived through trial and error: it produces reasonable results and provides useful guidance to the prudence of interconnecting various systems.

In (NCSC, 1987) an algorithm is described for determining whether or not a given network, composed of evaluated components meets the risk categories of the Environmental Guidelines. The algorithm is based on the idea of dividing a network into groups. The risk presented by any given group can be compared to the maximum allowed risk as defined by the Yellow Book for a system at the given evaluation class to determine if any community presents an unacceptable risk.

The steps for the heuristic procedure are (NCSC, 1987):

1. Create a Network Table listing all components within the network. This table should include the following information for every component:

    1.1. Component ID,
    1.2. Evaluation Class,
    1.3. range of security classifications at which the component sends data to the network,

1.4. list of security classifications at which the component receives data from the network,
1.5. maximum of (highest level of data received from network, highest level of data processed by component), and
1.6. minimum of (clearance of the user with the lowest clearance of the users with direct access to the component, lowest level of data sent to the network from the component).

2. Produce three tables: a Network Table Evaluation Class, a Network Table Maximum and a Network Table Minimum. The Network Table Evaluation Class will be the highest evaluation class of any component listed in the table. The Network Table Maximum will be the maximum of the maxima associated with all the components listed in the table which send data to the network. The Network Table Minimum will be the minimum of the minima associated with all the components listed in the table which receive data from the network.

   If the Network Table Evaluation Class is greater than B1 then tables for each evaluation class lower than the class of the Network Table, must be produced down to tables for the C1 class. These tables will be produced for each evaluation class by first listing any one component whose evaluation class is less than or equal to the evaluation class for the table. Then add to the table all components that meet all of the following conditions:

   2.1. they have an evaluation class less than or equal to the class of the table,
   2.2. they receive data from the network at a level that is being sent by a component which is already in the table, and
   2.3. they send data to the network at a level that is equal to or less than any node already in the table.

3. After all the tables have been constructed, the Network Table Evaluation Class of each table is compared to the maximum and minimum for the table with regard to the rules specified by the Environmental Guidelines.

4. If all tables satisfy the assurance requirements for the Environmental Guidelines then the network passes the assurance requirements. If any of the tables provide a greater risk index than is permitted by the Environmental Guidelines then the network provides a high level of risk and should not be connected as currently designed.

The reader can find an analytical application of the heuristic procedure in an example in (NCSC, 1987).

3.3 Shortest path network security model

The formulation of the Cascade Vulnerability Problem as a Resource-Constrained Shortest Path Problem (Fitch, 1991) (Fitch, 1993) leads to an efficient algorithm for determining whether a network presents a Cascade Vulnerability Problem.

The resource-constrained shortest path is based on three phases: Preprocessing, Shortest Path Calculation, and Postprocessing.

1. The Preprocessing step consists of the following actions:

   - defining the Cascade Vulnerability Problem as a graph by identifying nodes, edges, and weights,
   - viewing the problem from the penetrator's perspective by allocating the penetrator a set of resources, and
   - defining the resource consumption function that determines how the network consumes the penetrator's resources.

2. The Shortest Path Calculation step determines the paths through the graph that minimise the cost to the penetrator under the consumption function. The Shortest Path Calculation step may require careful selection of the algorithm or algorithms used so that the path calculation is computationally efficient. The appropriate selection may be based on the size of the network problem, the user-defined consumption function, and whether the penetrator's resources are scalar (e.g. money), or vector (e.g. (money, time))

3. The Postprocessing step analyses the shortest path results to determine the network security metric. For some applications determining that the network is not secure may require rearchitecting the network connectivity and reiterating the model steps.

   In this approach the algorithm is flexible in two ways:

   - There is a wide choice of what is meant by *cost*. In the standard Cascade Vulnerability Problem the *cost of a path* is defined by the maximum TCB rating of a machine on the path at which the security level of the data is compromised by illegally downgrading it to a lower security level. This makes the conservative assumption that once one machine of a given TCB rating is penetrated, then all others of equal or lower rating can be penetrated easily. However, the cost could also be defined as the sum of the costs of defeating the security protection mechanisms of all computers on the path independently.

   - There is a flexibility for the choice of which shortest path algorithm to use. In (Fitch, 1991) (Fitch, 1993) it is apparent that the Floyd-Warshall all-pairs algorithm (Aho, 1974) was intended to be used. This algorithm has a very good worst-case complexity to solve the all-pairs shortest path problem. However the Dijkstra single-source algorithm (Aho, 1974) could also be used. This algorithm is generally faster for sparse graphs like a computer network; therefore it is possible to implement the current algorithm using a somewhat faster algorithm than which is suggested in (Fitch, 1991) (Fitch, 1993).

The time-complexity of the algorithm is at most $O(n'^3)=O(a^3n^3)$, where n' is the number of nodes in the expanded graph. The space-complexity of the algorithm is $O(n'^2)=O(a^2n^2)$.

3.4 The Horton algorithm

An efficient algorithm for the detection of cascading vulnerability paths is presented in (Horton, 1993). In this algorithm the interconnection network of trusted computer systems is modelled as a directed graph with *n* nodes and *m* edges. The nodes have a

TCB rating associated with each other and represent the trusted subsystems. The edges represent the interconnection on which data can flow. The information needed to be associated with a node is the lowest user security level for which some user on the node is cleared, as well as the highest security level of labelled data at the node.

A new data structure is needed to represent the paths that data can follow. Each data path is represented by a directed edge or arc from the starting to the ending node of the path. With each arc a pair (d,u) is associated, where d is the security level of the data at the beginning of the path, and u the security level of the user at the end of the path. In a cascading vulnerability path the proposition $u<d$ holds true.

The algorithm (Horton, 1993) begins by creating arcs for each edge in the graph in which u and d are equal. The remaining part of the algorithm consists of two phases. In the first phase all possible legal paths through the network are found. A legal path is one for which $d \leq u$.

In the second phase the new arcs represent illegal data paths as well as legal data paths. The new step in this phase, is the answer of the question whether the TCB rating of node i allows the accreditation range of the $arc(j, k)$. If it does, then the path corresponding to the new arc is not a cascading vulnerability path. If it does not, then the path from which this arc is constructed is a cascading vulnerability path. For determining the above the author proposes the use of the Floyd-Warshall shortest path algorithm (Aho, 1974).

A cascade vulnerability correction algorithm should include suggestions for these network modifications which eliminate or reduce the risk of cascade vulnerability. (Horton, 1993) supports the idea that an algorithm that solves the cascade vulnerability correction problem would also solve the vertex cover problem for planar graphs (which is known to be intractable). Based on the above, the cascade vulnerability correction problem appears to be NP-hard. The detection algorithm shows that the correction problem is in NP and therefore the problem is NP-complete. Thus, an efficient algorithm that would give the optimal solution is unlikely to be found.

It is only possible that by using generalised techniques (e.g. ILP-Integer Linear Programming) a reasonable initial network could be defined, with all known constraints incorporated; this network could then be modified as unfulfilled requirements are identified.

Although solving integer linear programs problems is also NP-complete, there are efficient techniques which give acceptable solutions. The multiple-path problem can be stated as an ILP as follows:

p is a Cascading path in the network
n is a node in the network
s is the TCB rating for a system
$c_{n,s}$ is the cost of upgrading node n to rating s
$n,s \in p$ means that upgrading node n to TCB rating s would correct path p
$x_{n,s}$ is a variable that is either 1 (node n is to be upgraded to TCB rating s) or 0 (node n is not to be upgraded to TCB rating s).

Then

Minimise $\sum_{n,s} c_{n,s} x_{n,s}$, subject to $\sum_{n,s \in p} x_{n,s} \geq 1, \forall p$

The basic disadvantage of Horton's algorithm is that not all pairs of nodes connected by cascading vulnerability paths are found. However, if all the pairs of nodes connected by cascading vulnerability paths that are found are corrected, then all the cascading vulnerability paths are corrected. Any unreported cascading vulnerability paths will contain all of some reported cascading path.

3.5 Algorithm comparative effectiveness analysis

As stated above, the Horton algorithm has reduced time-complexity $O(an^3)$ as compared to $O(a^3n^3)$ for the (Fitch, 1991) algorithm. The space-complexity for the (Horton, 1993) algorithm is $O(an^2)$ as compared to $O(a^2n^2)$ for the (Fitch, 1991) (Fitch, 1993) algorithm. In (Millen, 1990) the resistance of all paths in the network is calculated by a matrix computation which requires $O(N^3 \log_2 N)$ steps.

A problem common to (Fitch, 1991) (Fitch, 1993) (Horton, 1993) is the following: if there are multiple cascading vulnerability paths between a pair of nodes, then only one of the paths will be detected using the straightforward version of the corresponding algorithm.

However, the (Horton, 1993) algorithm does not handle partially-ordered security levels as the (Fitch, 1991) (Fitch, 1993) (Millen, 1990) algorithms do.

4 CONCLUSIONS

A network of computers is exposed to the Cascade Vulnerability problem when data of a security level *d* can be passed to a user with a lower security clearance *u* elsewhere on the network, without having to defeat any single component of the system that has an accreditation range great enough to allow users of level *u* and data of level *d* on a single system.

Many efficient algorithms have been proposed in the literature to deal with the cascade vulnerability detection. In the previous sections we reviewed the basic properties of the most important algorithms and conducted a brief comparative analysis of them and explained why the cascade vulnerability correction problem is NP-complete. Possible future work could focus on finding reasonable approximation heuristic procedures for the cascade vulnerability correction problem.

5 REFERENCES

Aho. A., Hopcroft J., Ullman J., (1974) *The Design and Analysis of Computer Algorithms*, Addison-Wesley, Reading, MA.

Department of Defence (1985) *Trusted Computer System Evaluation Criteria*, DOD 5200.28-STD.

Fitch J.A. III, Hoffman L.J., (1991) The Cascade problem: Graph Theory can help, *Proceedings of the 14th National Computer Security Conference*, pp. 88-100.

Fitch J.A. III, Hoffman L.J., (1993) A shortest path network security model, *Computers and Security*, Vol. 12, pp. 169-189.

Freund J. E., (1962) *Mathematical Statistics*, Prentice Hall, Englewood Cliffs, N.J.

Horton J.D., et al., (1993) The Cascade Vulnerability Problem, *Journal of Computer Security*, Vol. 2, No. 4, pp. 279-290.

Katsikas S., Spyrou T., Gritzalis D., Darzentas J., (1996) Model for network behaviour under viral attack, *Computer Communications*, Vol. 13, No. 2, pp. 124-132.

Lee T. M. P., (1989) Statistical models of trust: TCBs vs. People, *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 10-19.

Madron T., (1990) *Network Security in the 90's*, J. Wiley & Sons, Inc.

Millen J. K., Schwartz M.W., (1988) The Cascading problem for interconnected networks, *Proceedings of the 4th Aerospace Computer Security Conference*, pp. 269-273.

Millen J. K., (1990) Algorithm for the Cascading problem, in J.P.Anderson ed., *Internet IEEE Cipher News Group*, June 25 IEEE Cipher Forum on *dockmaster.ncsc.mil*.

National Computer Security Center, (1985) *Technical Rationale Behind CSC-STD-003-85: Computer Security Requirements*, National Computer Security Center, USA.

National Computer Security Center, (1987) *Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria*, Red Book NCSC-TG-005, Library No. S228, 526, Version 1, National Computer Security Center, USA.

6 BIOGRAPHIES

Stefanos Gritzalis holds a BSc (Physics) and an MSc (Electronic Automation) degree, both from the University of Athens, Greece. He is also pursuing a PhD degree on Distributed Systems, with the Department of Informatics of the University of Athens, Greece. Currently, he is an Assistant Professor with the Department of Informatics of the Athens Technological Educational Institute (TEI). His research interests include Distributed Systems, Operating Systems and Information Security.

Sokratis Katsikas holds a Diploma (Electrical Engineering) from the University of Patras, Greece an MSc (Electrical and Computer Engineering) from the University of Massachusetts at Amherst, USA, and a PhD (Computer Engineering) from the University of Patras, Greece. He is an Associate Professor of Informatics with the Department of Mathematics of the University of the Aegean, Greece. Prof. Katsikas is the Vice-president of the Greek Computer Society and the representative of Greece to CEPIS SIG on Security and Legal Issues.

Diomidis Spinellis holds an MEng in Software Engineering and a PhD in Computer Science both from Imperial College (University of London). He has provided consulting services to a number of Greek and international Information Technology companies and is a four times winner of the International Obfuscated C Code Contest. Currently he is lecturing at the University of the Aegean, Greece. His research interests include Software Engineering, Programming Languages, and Information Security.