Network Working Group L. Amini
Internet-Draft IBM Research
Expires: March 31, 2003

A. Barbir
Nortel Networks

Oskar Batuner
Independent consultant

M. Day
Cisco Systems

O. Spatscheck
AT&T Labs

Kobus van der Merwe
AT&T Labs

Security Threat for Content Internetworking
draft-ietf-cdi-threat-00.txt

Status of this Memo

Security Threat for Content Internetworking

draft-ietf-cdi-threat-00.txt

Abstract

Content internetworking (also referred to as content distribution internetworking, or CDI) is the technology for interconnecting content networks. The CDI model allows for interconnecting various Content Networks. The internetworking task requires request routing and content distribution protocols. This document investigates the security risks and threats associated with the content internetworking. Proposed remedies are viewed not as design recommendations but more as illustrations of the nature of threats.

1. Introduction

Content internetworking (CDI) combines the resources of multiple content networks (CN) to increase their scale and reach. At the core of CDI are a request routing system and a distribution system. The request-routing system (RRS) directs client requests to surrogates and/or CNs that can best service the request. The internetworking of CNs is performed through Content Internetworking Gateway (CIG). The internetworking distribution system is responsible for moving content from one Distribution CN to another Distribution CN. Finally, the accounting infrastructure tracks and collects data on request-routing, distribution, and delivery functions within the CDN. The details of the CDI model can be found in [1].

The use of CDI - as any new mechanism - introduces new security risks and threats to the internetworked CNs. Some of these threats are specific to the CDI model, some are inherited from the CN systems. This document covers both new and inherited threats with distinctions made were appropriate.

The security risks within CDI can be classified along various dimensions including:

- the source of the threat ("insider" versus "outsider"),

- the level at which the attack occurs (network level attack versus application level attack),

- the type of harm that results from an attack (harm to content, harm to identity, harm to finances).

- the elements of the architecture attacked (e.g., the Distribution System, the Request Routing System, the Accounting System, the clients, or publishers)

All of these dimensions are considered in this document (some in greater detail) to develop a complete view of the threat model for content internetworking. However, this document focuses only on those threats specific to the content internetworking model. It does not consider, for example, the following issues:

- The security risks within an individual CN, such as denial of service attacks on individual surrogates, are beyond the scope of this document.

- Content security issues, such as the integrity of transformations or adaptations performed on content, are outside the scope of the current work.

- This document does not specify or recommend any particular solutions. In some cases however, potential threat mitigation steps are given to help illustrate a given threat.

The remainder of this document is organized as follows. We begin by describing the CDI Trust Model, and distinguish between "insider" and "outsider" attacks. Next, we broadly classify attacks as occurring at the network, content internetworking, or application level, and detail the resultant type of harm. We refine this list by detailing how the attacks might be perpetrated on specific components of the CDI architecture, and potential mitigation steps.

1.1 Conventions used in this document

Key terms in ALL CAPS, except those qualified with explicit citations, are defined in [1].

2. Content Internetworking Trust model

Relationships between CN's in the CDI model can be decomposed into relationships between individual pairs comprising a CONTENT SOURCE and a CONTENT DESTINATION. The ORIGIN refers to the point at which CONTENT enters the CDI model, and therefore is a specific type of CONTENT SOURCE. The trust model utilized within CDI is based on a transitive trust between a CONTENT SOURCE and a CONTENT DESTINATION. The transitive nature of the trust originates from the need of an ORIGIN to rely on one or more CONTENT SOURCE - CONTENT DESTINATION

pairs to deliver CONTENT to CLIENTs on the ORIGIN's behalf.

The trust model involves the following parties in trust relationships:

- CONTENT SOURCE and CONTENT DESTINATION
- CONTENT SOURCE and CLIENT
- CONTENT DESTINATION and CLIENT

We will use the term TRUSTED PARTY to refer to a party involved in a trust relationship.

We begin by classifying security risks into two main categories: threats from "insiders," and threats from "outsiders." Outsiders are those entities that have not established a trust relationship within the content internetworking system. Insiders are TRUSTED PARTIES that are participating in a trust relationship within the content internetworking system.

Threats from within the system may be intentional or unintentional. Intentional threats refer to the ability of a TRUSTED PARTY of a CDI relationship to mislead or harm, the party with which it has a trust relationship. For example, the TRUSTED PARTY, a CONTENT DESTINATION, might misrepresent quality or quantity of the service provided to the trusting party, a CONTENT SOURCE. This is distinct from the case when a TRUSTED PARTY's system is compromised by an outsider, which is covered as an "outsider" threat.

Unintentional threats refer to the ability of a TRUSTED PARTY, through improper implementation or configuration resulting in bad system behavior, to mislead or harm the party with which it has a trust relationship.

Content internetworking allows for relationships whose terms and conditions are partially or completely established outside the context of the content internetworking protocols, and refers to these relationships as NEGOTIATED RELATIONSHIPS. Just as trust relationships established completely within the context of content internetworking protocols, NEGOTIATED RELATIONSHIPS can result in intentional or unintentional threats.

Threats from outside the system, or outsiders, may also be intentional or unintentional. Since unintentional threats from outsiders do not rely on the trust model, and are not specific to the content internetworking model, this document will consider only outsider threats that are intentionally perpetrated.

In this document, we will focus on intentional and unintentional threats from within the system, and intentional threats from outside the system.

3. Threat classification by architectural level

In this section, we broadly classify threats according the architectural level -- network, content internetworking, or application -- at which the threat occurs. We refer to threats exploiting design or implementation weaknesses of internetworking and transport protocols (i.e., layer 3 and below of the TCP/IP protocol suite) as network level threats. We refer to threats exploiting weaknesses in content internetworking protocols as content internetworking level threats. We include in content internetworking level attacks, threats against CONTENT distributed using CDI specific protocols. Finally, we refer to threats to applications that utilize a content internetworking system as application level threats.

Where appropriate, the type of harm that can result from an attack is provided to show the complex interaction between different threats and/or attacks. For example, harm to content in the form of content degradation or content substitution might harm the finances of the content provider which might in turn harm the finances of the service provider. A denial of service attack or theft of identity might have a similar effect on parties involved with CDI.

3.1 Network Level Threats.

The content internetworking model comprises CONTENT NETWORKs, which in turn comprise CONTENT NETWORK ELEMENTS. A CONTENT NETWORK ELEMENT is a network device that performs at least some of its processing by examining CONTENT-related parts of network messages. Examples of CONTENT NETWORK ELEMENTS include CONTENT INTERNETWORKING GATEWAYs (CIG) and SURROGATES.

In IP-based networks, a CONTENT NETWORK ELEMENT is a device whose processing depends on examining some or all of an IP packet's body. As such, CONTENT NETWORK ELEMENTs are vulnerable to many types of network level attacks. Examples of TCP/IP attacks include IP spoofing and session stealing. The CERT Coordination Center [2] maintains an extensive repository of Internet Security vulnerabilities.

Harm specific to CONTENT NETWORK ELEMENTS, such as a CIG, achievable by hijacking a TCP/IP session includes the ability of outsiders to inject believable content distribution and request routing messages into the communication between CIG peers. This may lead to the injection of bogus content or bogus routing information that may lead to the breaking of the peer to peer connection. Any break in the peer to peer communication can have a ripple effect on the request routing system or the distribution system that could lead to disrupted services to end users.

CONTENT NETWORK ELEMENTS are also susceptible to a number of security threats commonly associated with network infrastructure. These threats include snooping, denial of service, sabotage, vandalism, industrial espionage, theft of service and inadequate system configuration that leaves unneeded ports and services open to the public.

3.2 . Content Internetworking Level Threats.

Content internetworking Level threats generally belong to one or more of the following categories:

- denial of service
- content distortion
- threats to identity
- threats to privacy
- content theft
- security threats
- threats to finances

In the following subsections we elaborate on these threats and potential resultant harm.

3.2.1 Denial of service threats.

At the Content Internetworking level, a denial of service (DoS) threat can be perpetrated on a number of levels. For example, an attack could be launched:

- specifically against a CONTENT SOURCE, thereby preventing any distribution from taking place
- against a content set, causing all CNs servicing this content set to be affected.
- against all SURROGATES of a specific CN.

A CONTENT SOURCE distributing streaming content, due to its high bandwidth nature and, in the case of live streaming, limited injection points, are likely to be especially vulnerable to DoS threats.

Misuse of a CN may make its facilities unavailable or available only at reduced functionality. Denial of service attacks can be targeted at a CN accounting system, distribution system, or request-routing system.

3.2.1.1. "Complexity threat": both CDN and CDI introduce many components and complex infrastructure. Malfunctioning of these components and infrastructure may result in DoS.

3.2.1.2. Misconfigured request routing (unintentional or malicious) may cause request loss or looping and result in DoS.

3.2.1.3. Conflicts between request routing and accounting mechanisms may create a DoS threat: a CN may refuse to deliver content because the authorization system treats a valid request as invalid (not coming from an authorized customer).

3.2.1.4. By redistributing the load between CNs CDI may cause DoS by unintentionally overloading one of CNs. Usually CNs have a specific (proprietary) adaptive mechanisms for load balancing. CDI load balancing mechanisms may be inadequate/malfunction or be incompatible with corresponding CDN load balancing.

3.2.1.5. A CN may cause problems in another CN by sending (unintentionally or with malicious intent) more content than advertised capacity permits.

3.2.1.6. Corruption (intentional or non intentional) of security related metadata (authentication data) might result in DoS: CN or CDI may refuse to perform a legitimate service.

3.2.1.7. False advertisement (unintentional or malicious) of nonexistent distribution/coverage capacity may result in failure of several CNs. Same problems may result when advertisement and usage policy do not reflect dynamic conditions.

3.2.1.8. Incompatible request routing systems may cause problems resulting in DoS.

3.2.1.9. Peering agreements may be vital for CDN functionality. This makes peering reliability a security issue. CIG (distribution CIG and request routing CIG) may introduce a single point of failure. Attack on (or malfunctioning of) a CIG may result in system disintegration and DoS for both CNs.

3.2.2 Content distortion threats.

3.2.2.1 An attacker may cause a CN to advertise bogus content, e.g. replacing proper content with bogus content either at the injection point of the system (CN or CDI) or inside elements of the system (e.g. surrogates inside the CN).

3.2.2.2. A CN may provide bogus information, e.g. a rogue "CN" inserting itself in the distribution path between two CNs to monitor and/or modify the content that they exchange.

3.2.2.3. A CN may advertise the availability of content which it doesn't have and can not distribute. This attacks can be the result of malicious CIG taking over the identity of a CIG to be able to inject bogus info into system, or a CIG that is compromised

3.2.3 Threats to user identity.

Identity/authentication threats may result from third party getting access to authentication data of end user or system component (surrogate, CIG) and this data permits unauthorized actions to be performed. Note that the last condition is essential: interception of session initiation packets of replay-resistant secure authentication protocol does not create such a threat.

Storage of security related data (user identities, passwords, etc.) creates an additional security threat.

3.2.4 Threats to privacy. Privacy threats may result in personal user information made available to third party without user's consent.

3.2.4.1. A CN may inadvertently or maliciously expose private information (passwords, buying patterns, page views, and credit card numbers) as it collects it and transits from surrogate to origin and/or publisher.

3.2.4.2. Accounting information transfer may jeopardize privacy.

3.2.4.3. Privacy threats may result from differences in privacy policy of Publisher, CDN and CDI.

3.2.4.4. Privacy and security threats from crossing jurisdiction boundaries: transfer and storage of sensitive privacy-related data (accounting, logs), transfer and storage of (secure) content and distribution of content from a different jurisdiction may create a security threat due to different level of legal protection.

3.2.5. Legal threats: by extending activities through jurisdiction boundaries CN and CDI may unintentionally violate local regulations (privacy and security policies).

3.2.6 Content theft.

Unauthorized access to non-public (secure or non-secure) content. For

secure content such unauthorized access clearly violates intention of security system and usually constitutes a content theft (paid content, proprietary data).

An example of unauthorized access to non-secure content is interception of form data in not-secure transmission or direct access to URL that is not supposed to be publicly available.

3.2.7 Security threats

3.2.7.1 Unauthorized access to metadata that is not supposed to be publicly available. This may include access to logs and accounting data containing private user's information, access to configuration data that may be used to facilitate future attacks and so on.

3.2.7.2 Exposure of Security Settings: There may be risks that expose client's security settings when content is served from surrogates as opposed to origin servers. Since the location of the surrogate is generally transparent to the client, the client may be aware that its protections are no longer enforced.

3.2.7.3 Improper enforcement of Security Policy

Policy information regarding security of the client may not be properly propagated when the requests are directed to surrogates in a CN that are different from the origin server. Client passwords and personal information may be less secure.

3.2.8. Improper Carriage of Security Policies

Surrogate may not employ the same security policies and procedures as the origin server. This may expose the client private information to access by unauthorized entities. The same threat may also result if the legal jurisdiction of the surrogate is different from that of the origin.

3.2.8.1. Different implementation of security at Publisher, CDN and CDI level may create security threats

3.2.8.2. Distribution of content from a different network location may create a security threat if client security policy depends on network location ("Internet Web Content Zone").

3.2.8.3. Transfer and storage of secure content create additional security threats.

3.2.8.4. The process of propagation of security policy and security related data (user identities, passwords, etc.) creates security

threats both at CDN and CDI level.

3.2.9 Threats to finances

Delivery of inaccurate accounting information or malicious distortion of this information may cause financial harm to all participating parties.

3.2.9.1 The client may be inappropriately charged for viewing content that was not successfully accessed or delivered according to some QoS criteria.

3.2.9.2 If a CN or Publisher is unable to collect or receive correct accounting information they may be unable to collect compensation for services.

3.3 Application level threats.

TBD (section should include attacks targeting applications that utilize the content internetworking system)

4. Threats against specific elements of the CDI architecture

In this section, we refine the list of threats by detailing how the attacks might be perpetrated on specific components of the CDI architecture. This section is intended to be used input to specify the security requirements for the content distribution and request routing protocols.

Along the dimension of threats against specific elements of the architecture, threats against the accounting system should also be noted. A detailed analysis of the threats against the accounting system can however only be done within the framework of a specific accounting system and is considered outside the scope of this document.

4.1 Threats to the Content Internetworking Gateway The CIG is the connecting point for the CNs that are participating in the CDI model. CIGs from various CNs establish peer to peer relationships in order to exchange content distribution and request routing information. Threats on the CIG can be perpetrated at all levels, the network, content internetworking, and application level.

A CIG must be accessible at the network level from many other CIGs. The CIG is vulnerable to any of the network level attacks specified in Section 3.1. The CIG is susceptible to network level

attacks from outsiders, which may or may not be posing as the CIG of a TRUSTED PARTY, and from CIGs of TRUSTED PARTIES.

4.2 Threats to Distribution System

Threats to distribution system from insiders can be intentional or the result of bad implementation. Outsiders can pose the same threats if they acquire access to the distribution system. The threats include:

4.2.1 Advertising of unavailable content.
4.2.3 Advertising of bad metrics that are associated with a given content.
4.2.4 Delivery of bad content to surrogates in the connected CN
4.2.5 Using badly formed messages for advertisements

4.3 Threats to Request Routing System

Threats to the request routing system from insiders or outsiders include:

4.3.1 Advertising of wrong metrics to force unfair or inaccurate redirection to a given CN.
4.3.2 Redirection to a CN that does not have the content.
4.3.3 The introduction of loops in the requesting routing system.
4.3.4 Redirection to an inappropriate surrogate.
4.3.5 Forwarding request when no forwarding is appropriate.
4.3.6 Failing to forward requests when forwarding is appropriate.
4.3.7 Using badly formed messages for advertisements

h) TBD

5. CDI Security Threat Mitigation

The main security issues for the CDI model are focused on the Trust model. Insiders are TRUSTED PARTIES, while outsiders are not.

Threats from outsiders are primarily at the network level. There are well known solutions to network level threats that are practiced in the industry. In this work, it is recommended that the security of the CONTENT NETWORK ELEMENTs at the network level be enhanced using standard techniques and methods that minimize the risks of IP spoofing, snooping, denial of service and session stealing.

Threats at the content internetworking and application levels can be mitigated by using strong authentication and encryption techniques. Therefore, there may be the need to make strong authentication and encryption a requirement for the CDI model. IPSec

and TLS are solutions for this requirement. Regardless of the choice of the protocol, the solution must scale to accommodate large number of interconnected CNs. Furthermore, it is recommended not to send passwords in the clear.

To mitigate threats from insiders CDI must implement appropriate monitoring, signaling, logging, dynamic authorization and verification mechanisms. The following sections provide more detailed guidelines for development of request routing and distribution protocols for content internetworking.

5.1 Treatment of malformed messages

Malformed message can be the result of bad implementation or a consequence of an outside attack on a given CN whereby, the attacker gains access of the peering system. A Malformed messages is a message that does not comply with the message format for the distribution (or request routing) protocol. A malformed message may be a message that has wrong content attributes in it or wrong IP footprint. A malformed IP or IPSec packet is not considered a malformed message.

In the event that a CN detect malformed messages terminating the session appears to be the only safe way to handle it. Terminating a session does not mean terminating the peering relationship. The session can be restarted after termination. If the problem of malformed messages persists, the interconnected CNs must verify the cause of the problem and proceed with a solution.

The treatment of malformed messages is different than the case where a peer intentionally or unintentionally sends incorrect advertisements which might lead to incorrect selections. For example, a CN might incorrectly advertise low load, low cost and good coverage and therefore attract a large proportion of traffic. This problem can be somewhat mitigated through filtering of advertisements and local policies but ultimately comes down to a trust relationship between peers.

5.2 General Distribution and Request Routing Protocol Requirements

Based on the security threats that are faced by other peer-to-peer based protocols such as BGP, this section provide some guidelines that should be used during the design of the request routing and content distribution protocols.

5.2.1 There should be a mechanism that provides strong protection of the integrity, freshness and source authenticity of the messages in

the protocol. Techniques such as digital signature may be used.

5.2.2 There should be a mechanism to validate the authenticity of a CN_Path value.

5.2.3 There should be a mechanism to use IP level protection that can be used to provide connectionless integrity, data origin authentication, and secure authentication.

5.2.4 There should be a mechanism to protect the peer-to-peer connection by applying cryptographic protection at the TCP level to provide connectionless integrity and data origin authentication.

References

[1] Day, M., Cain, B. and G. Tomlinson, "A Model for Content Distribution Internetworking", January 2001.

[2] CERT Coordination Center (CERT/CC). http://www.cert.org/nav/index_main.html