

Messaging, Malware and Mobile Anti-Abuse Working Group

M³AAWG Compromised User ID Best Practices

Table of Contents

1. Executive Summary.....	1
2. Scope of this Document.....	2
3. Definitions.....	2
4. How Are Accounts Compromised?	2
5. Identify Compromised Accounts.....	3
5.1 Differentiating between compromised and malicious accounts	3
6. Mitigating Account Compromises	4
6.1 Remove inappropriate access and return control to the original account holder.....	4
6.2 Secure the account against re-compromise	6
6.3 Re-compromised accounts	6
7. Mechanisms that Need to Be Implemented	7
7.1 Detection of registration abuse.....	7
7.2 Notification of and interaction with account holders.....	7
8. Conclusion	9

1. Executive Summary

Spammers are constantly searching for new routes to deliver potentially dangerous and unwanted mail. For many years, spammers have worked with purveyors of viruses to surreptitiously infect users’ PCs with spam-sending software, creating networks of computers known as “botnets.”

Until recently, botnets typically sent spam directly from infected PCs on subscriber networks; however, vigilance by the anti-abuse community has made this direct form of spamming more difficult. As a result, spammers have begun to configure their botnets to send from a new source: compromised user email accounts.

This M³AAWG best practices document is focused on addressing problems associated with compromised user accounts. In order to address the problem, it is important to define what a compromised user account is and how user accounts become compromised. This document discusses mitigation techniques and ways of identifying compromised accounts. Rounding out the document is a set of recommendations to ensure the long-term security of accounts to prevent “re-compromise.”

This document is intended for operations staff involved in the creation and management of end user accounts, as well as for abuse-desk personnel who deal with the repercussions of compromised end user accounts.

2. Scope of this Document

This document provides advice on how to deal with compromised accounts: specifically, how to deny access to the unauthorized user of an account while ensuring continuity of access to the original account holder, if possible.

This document does not deal with the separate – and equally challenging – problem of accounts that were originally created for the purpose of abuse¹. We will touch on, but will not advise on, the topic of detecting and preventing malicious account registration.

This document will not address compromised employee accounts because these call for a different and a considerable remediation approach above and beyond that of a compromised end user account. We also will not deal with situations in which an account holder is tricked into a one-time action resulting in the propagation of click-jacking, trojans without credential-stealing behaviors, straightforward worms, etc.

3. Definitions

User Account: A set of services provided to a user for his or her exclusive use and secured through the use of a password or other authentication method.

Compromised User Account: A User Account that is fully or partially under the unauthorized control of someone other than the legitimate user.

4. How Are Accounts Compromised?

Accounts can become compromised in a wide variety of ways including, but not limited to, the following:

- **Credential compromise:** Unpatched network or software security vulnerabilities are exploited to leak users' credentials or an attacker compromises a legitimate site and steals the credential database for reuse.
- **Phishing:** An attacker pretends to be a legitimate site or sender and, as a result, users willingly provide their account credentials.
- **Keyloggers:** An attacker compromises the user's machine and steals all credentials entered on it.
- **Theft of temporary credentials (for example, cookies):** An attacker sniffs a network or performs cross-site scripting (XSS) attacks.
- **Password guessing (brute force attacks):** An attacker tries common passwords on many accounts or many passwords on individual accounts.

Other possible sources include social engineering of the account holder or the provider as well as compromise of provider employees.

While it is important for providers to prevent attacks when possible, there are always attacks that succeed by taking advantage of human susceptibility. Consequently, providers should assume there will always be compromised accounts.

¹In some cases, we will discuss this type of abuse, because it is often detected and handled via the same channels and must be considered when designing procedures to deal with compromised accounts.

5. Identify Compromised Accounts

Because the definition of an “account” is highly dependent on the particular factors local to a service provider organization, it is difficult to prescribe a single set of steps that will successfully identify compromised accounts in an organization. However, the recommendations below are provided as a general, high level approach and can be applied by any organization to make the identification of compromised accounts easier.

Attackers compromise accounts to satisfy a goal, such as sending spam. In working to satisfy their goals, attackers usually cause a compromised account to behave differently than it normally would if it were being used by the legitimate user. This varying behavior might be noticed by an organization’s own systems, by the legitimate user, by other users or by third parties.

It is important to note that accounts suspected of malicious or irresponsible activity are not necessarily compromised accounts. There may be accounts where the account holder is intentionally doing something bad or irresponsible with either a single account or multiple accounts. (See “[7.1 Detection of Registration Abuse](#)” below for information about mass account registrations.) If there is a pattern of previous valid use, however, the account may have been compromised. There is a good possibility in this case that there is a valid owner who has lost control of the account.

5.1 Differentiating between compromised and malicious accounts

One solution to identifying compromised accounts is to tap into the feedback from those individuals and systems that can detect and report that an account is behaving differently than it normally does.

Good sources of account behavior feedback include the following:

- Account owners themselves reporting that their account has been compromised, presumably after observing strange behavior in relation to their account; e.g., the emission of spam from their email address.
- Other users or third-party data providers reporting activities consistent with a compromised account.
- Behavior tracking and analysis – Accounts that have established behavior patterns may call for special actions when they show behavior outside those patterns. For example, common indications of compromise are accounts that log in from unexpected locations, accounts that suddenly send messages to all of their contacts, or accounts that delete all their sent mail.
- Detection of undesirable or unreasonable behavior — Some behaviors are consistently suspicious. For example, accounts with logins from widely separated geographical areas in short periods of time or that send messages containing hostile URLs may be detected as possibly compromised.
- Feedback loops provided by other service providers can help identify compromised accounts. An account may be in trouble if a feedback loop is suddenly generating a significant number of spam reports or complaints related to it.

Sometimes account compromise happens long before an attacker engages in detectable behavior. For example, an attacker might obtain a list of user names and passwords but not access the accounts for several weeks or months before finally logging in and engaging in a spamming campaign using the compromised accounts. By tapping into third-party reports and data sources, service providers can gain advance warning of such pending or probable account compromise.

The following is a list of good sources of information to help identify potentially compromised accounts:

- Third parties reporting found credentials; e.g., security researchers reporting they have discovered a cache of user names and passwords

- Attackers publishing lists of compromised accounts
- Notification of related compromise – Lists of compromised accounts at one site will often identify linked accounts that can be used to compromise accounts at another site. Attackers will follow these links and defenders need to do so, too.
- Detection of a pattern of account compromise – Once a compromised account is identified, it might be possible to use its characteristics to identify others.

Note that none of these methods will always reliably identify compromised accounts. Providers must verify that accounts are actually compromised — or at least, have a high probability of being compromised — before acting on them. Machine techniques may or may not be helpful in verifying whether an account has been compromised; consultation with the actual account owner may be required for accurate and full verification.

6. Mitigating Account Compromises

Mitigating compromised services varies significantly based on factors such as whether end users can create new accounts openly on the Web (versus requiring a formal interaction with customer support), whether user accounts are tied to other products or services, and the specific methods an organization uses to communicate with customers. Listed below are basic guidelines, which should apply to most scenarios.

Returning a compromised account to its rightful owner, if one exists, involves the following steps:

1. Identify compromised accounts.
2. Remove inappropriate access and return control to the original account holder.
3. Secure the account against re-compromise.

6.1 Remove inappropriate access and return control to the original account holder

The best way to remove the attacker’s control of the account depends on what was compromised and how it is being exploited. Therefore, the detection of compromised accounts should be categorized by compromise type, as each type requires different treatments.

Compromised temporary credential

Examples	Cookie theft via XSS or network sniffing; session hijacking
Distinguishing characteristics	Temporary credentials have a limited scope of action and these attacks can usually be distinguished from other types of compromise by the fact that they confine themselves to actions for which the temporary credentials can be used.
Mitigation	Invalidate the temporary credential and have the user re-authenticate with a permanent credential.
Requirements	Detection of the issue; methods for invalidating a temporary credential in use.
Recommendations	Temporary credentials should be issued with a limited lifespan. Providers should have a bulk mechanism for invalidating temporary credentials when compromise is detected. Users should also be able to invalidate temporary credentials they believe have been compromised.

Compromised permanent credential, mass exploit

Examples	Compromised accounts that are sending spam to their contacts or operating as part of a botnet, or accounts whose credentials have been compromised as part of a known exploit.
Distinguishing characteristics	Involves large numbers of accounts and cookie-cutter exploitation with little or no customization. It is distinguished from abusive registration by accounts talking to contacts or accessing information about contacts after a pattern of previous good behavior.
Mitigation	Invalidate the compromised credentials and force the account owner to set a new one. Mechanisms for preventing re-compromise are discussed in Section 6.2 , and specific details for resetting the user's credentials are discussed in Section 7.2 .
Requirements	Prompt the user to set a new password upon login. In addition, a mechanism for locking out the compromised credential while still allowing the account-holder to reset it should be put into place.
Recommendations	Providers should have a bulk mechanism for forcing users to change compromised credentials. However, we recommend making the credential change process difficult to automate (e.g., by incorporating CAPTCHAs or requiring answers to password-recovery questions) to increase the cost to the attacker.

Compromised permanent credential, customized exploit

If mitigation attempts inconvenience mass exploiters, they will change their tactics to the point where human intervention is required to mitigate the abuse.

Examples	Stranded traveler; accounts sending mail to the users' contacts claiming that the account owner has been mugged while on an unexpected business trip and asking for money to be wired.
Distinguishing characteristics	The compromise is hand-crafted for each victim, but there are numerous victims. The attacker is at a social and physical distance from the victim.
Mitigation	These attacks will always require human intervention by someone who is familiar with both customer service and fraud, and who has a technical background. In some cases, the only available mitigation will be to cancel the account because it is impossible to determine the original owner. In other cases, it may be necessary to set up a new account and transfer some contents of the old account to it.

Compromised permanent credential, human exploit

Examples	Stalking attacks on celebrities, ex-partners and family members.
Distinguishing characteristics	Attack is on a more personal scale; the attacker is physically or socially close to the victim. The attacker may be willing to invest substantial amounts of effort into compromising the account.
Mitigation	Mitigation and requirements are similar to those defined in the previous section, “Compromised permanent credential, customized exploit.”
Requirements	Mechanisms to identify and freeze accounts in contention; trained personnel needed to resolve issues; mechanisms for the provider personnel to find relevant account data.
Recommendations	Providers should have an escalation process allowing accounts in contention to be referred to staff for a more thorough evaluation and investigation, as appropriate. Providers should have the ability to freeze accounts pending staff intervention by the provider. Providers should have the ability to transfer account contents to a new account, if necessary.

6.2 Secure the account against re-compromise

Several steps can help to prevent re-compromise after an account is returned to the legitimate user. These steps include:

- Assist the user in identifying and removing any back doors left by the exploit, i.e., changes to contact mechanisms, email forwarding, reply-to addresses.
- Prevent password reuse.
- Institute stronger password requirements.
- Request the user to implement two-factor authentication, if available.
- Educate the user about the means of compromise, if known.
- Notify other providers if the compromise involves their accounts.
- Encourage users to only access account resources via secure mechanisms, if available (e.g., IMAPS instead of IMAP)

6.3 Re-compromised accounts

Accounts that are compromised multiple times call for special handling rather than putting them through the same recovery process over and over again. Deploying measures to look for re-compromise should be put into place to limit recovery cycles, effort and further abuse.

Accounts may be detected as compromised multiple times due to:

- Flaws in the detection process causing false positives.
- Compromise of something other than the account credentials, such as a compromise of the account holder’s computer, a network the account holder frequently uses, or the provider itself.
- Account holders who do not have adequate skills to protect their account credentials.

These cases require separate handling. In the last case, someone must be authorized to make a business decision on the next steps: Do you retain the customer, possibly with special requirements, or send their business elsewhere?

7. Mechanisms that Need to Be Implemented

Much of the real work in mitigating the effects of compromised accounts requires building mechanisms to detect and mitigate the compromise. These mechanisms range from provider-specific work making automated detection and user notification systems effective to research on methods for self-service credential change.

7.1 Detection of registration abuse

Abusive registration is outside the scope of this document and is of interest here only because providers should be able to distinguish abusively registered accounts from accounts with a valid account holder.

For free and easily created accounts, registration abuse detection usually takes the form of detecting registration patterns. For accounts that cost money, registration abuse often involves stolen or disposable credit cards. For accounts that require significant human intervention to set up — advertiser accounts, for instance — registration abuse typically involves account holders who are suspiciously lax in their requirements. These account holders might make surprising claims about who they are, what they do, and where they are, e.g., an unknown ad agency with a top-name client list; an interior design firm in a distant country with a burning desire to have local representative; or a large financial services firm whose address is a mail drop. In their zeal to create the account, they often ignore delays or rules that other users frequently question or they are willing to quickly change their business practices to fit the account requirements.

7.2 Notification of and interaction with account holders

Detecting compromised accounts does not prevent current or future abuse—the compromise must be resolved and ownership restored to the valid owner or the account completely taken offline. For any reasonably-scaled compromise scenario, this means implementing an automated mechanism that forces compromised accounts to re-authenticate.

Self-service credential change

This is the single largest area of innovation in dealing with compromised accounts. It is desirable, from both the provider's and the account holder's perspective, to allow the account holder to re-authenticate using an automated process whenever possible. This usually requires the existence of an independent authentication channel allowing the account holder to prove they are not only human but also are the correct person to control the account.

The authentication channel must rely on information or access that, at a minimum, does not become available to an attacker automatically upon compromise of a primary credential. Some mechanisms include the following:

Secret questions

A question or set of questions set up, usually at account creation time, with answers theoretically known only to the account holder that can be used to authenticate the user after a compromise.

Pros	Cons
<ul style="list-style-type: none">• Familiar to the user	<ul style="list-style-type: none">• Hard to automatically verify because answers are relatively long free-form text susceptible to all sorts of variation that are unimportant to humans but important to machines.• Most users pick questions where the answer is either easily guessable, e.g., “What color is the sky,” or researchable, e.g., “What school did you go to,” or both, e.g., “What was your first pet’s name?”• Phishing and social engineering attacks routinely ask for the answers to secret questions so spammers may have already obtained the information and can provide the correct answer.• Users frequently forget the answer they gave to a secret question, a problem that becomes worse if the answer changes over time, e.g., “What is your favorite movie?”

Alternate contact methods

Require a phone number or independent email address associated with the account where the authentication data can be sent.

Pros	Cons
<ul style="list-style-type: none">• Particularly easy to automate• Understandable to the account holder	<ul style="list-style-type: none">• Account holders may be unwilling to provide cell phone numbers for privacy or expense reasons and some may not have cell phones.• Account holders may not wish to either link the account to an existing email address or set up a second email address.• Account holders that provide the information often make errors when entering the answers. This can be mitigated through a verification processes but legitimate account holders often fail to complete the verification processes. Holders of inaccurately entered email addresses who receive verification notices in error rarely respond to deny the linkage.• Account holders often set up two email accounts as reciprocal alternate contacts, meaning that as soon as one is compromised, the other is also compromised.• Account holders often use the same password on multiple accounts and spammers know to try the same credentials on various accounts held by the same person. As a result, even if the primary account that is compromised is not an email account, the spammer may gain access to the owner’s email account by trying the previously obtained user name and password.• Account holders often fail to update the information when they lose control of the alternate channel, turning a safety measure into another route for compromise.

Providers can mitigate many of these issues with close control of users’ accounts by requiring alternate information, requiring positive verification of it, and repeating the process of positive verification on a regular basis.

Other information about the account

A service provider may know the zip code, birth date, registration geolocation, previous passwords, or secret questions and use that information as a form of verification. For providers with a business relationship with the account holder, it could be information about business interactions; e.g., “What was the last order you placed?” or “What’s the serial number on your modem?” This could be information generated by the account, information about messages sent or received, links with friends, and content related to the account.

Pros	Cons
<ul style="list-style-type: none">• Does not require special set up for account recovery	<ul style="list-style-type: none">• Providers of free accounts, for business reasons, generally minimize the information collected at account creation and they do not or often cannot verify the answers. As a result, the providers do not have this type of data available.

While this approach can be useful, it is often difficult to find information that the attacker cannot change or discover if the attacker has access to the account. Automating processes like this often exposes possible verification information to an attacker who has access but has not yet compromised the account. For example, if you display a set of pictures for the user to choose from, a potential attacker who has access to the account knows that at least one of the presented options is known to the legitimate holder.

8. Conclusion

Account compromise is a common occurrence and will continue to occur in the future. Account providers must be prepared to detect when an account has been compromised and provide steps to restore it back to the original user.

From small- to wide-scale attacks, spammers and botnet operators are leveraging compromised user accounts to send massive amounts of spam without authorization from – and often, without knowledge of – the actual account owners. The issue requires carefully defined approaches for remediation from mailbox providers.

The processes in this document are intended to serve as guidelines for account providers to mitigate the compromise, restore the account, prevent future re-compromise, and provide the user a self-serve mechanism. By implementing some or a combination of these, account providers can better serve their users and the Internet community as a whole.