INTRUVENT
COMPREHENSIVE THREAT HUNTING GUIDE

Charming Kitten: Social Engineering, MFA Bypass, and Credential Harvesting Detection

Detecting phishing proxy attacks, MFA token interception, OAuth abuse, and custom backdoor deployment by APT35/APT42

Version 1.0
Published March 18, 2026
Threat Actor Charming Kitten / APT42 / Mint Sandstorm (IRGC)
Classification TLP:CLEAR
Total Hunt Time 5 Hours
Target Sectors Think Tanks, Academia, Media, Government, Defense
Intruvent Technologies
INT-THG-2026-APT35-SOCIAL-v1.0

Table of Contents

Overview

Executive Summary 3
MITRE ATT&CK Mapping 4

Hunt Modules

Module 1: Phishing Proxy Detection (Evilginx/Modlishka-Style Token Interception) 5
Module 2: MFA Token Replay and Session Hijacking Detection 6
Module 3: Suspicious OAuth Application Grants 7
Module 4: React Phishing Kit Domain Pattern Detection 8
Module 5: Email Impersonation and Persona-Based Spearphishing Detection 9
Module 6: PowerStar/CharmPower Backdoor Hunting 10
Module 7: BellaCiao Backdoor Detection (DNS-Based C2) 11
Module 8: Credential Harvesting Infrastructure Detection 12
Module 9: Compromised Account Lateral Movement 13
Module 10: Cloud Storage Exfiltration Detection 14

Appendices

Appendix A: Indicators of Compromise 15
Appendix B: Detection Rules 17
Appendix C: References 19

1 Executive Summary

AT A GLANCE

THREAT ACTOR Charming Kitten / APT42 / Mint Sandstorm (IRGC-IO, Iran)	PRIMARY OBJECTIVE Credential theft, surveillance, and intelligence collection via social engineering
ACTIVE SINCE 2014 (social engineering); 2025 (React phishing kit infrastructure)	TARGET SECTORS Think Tanks, Academia, Media, Government, Defense, Diaspora
KEY DIFFERENTIATOR Real-time MFA token interception via phishing proxy infrastructure	CUSTOM BACKDOORS PowerStar/CharmPower, NICECURL, TAMECAT, BellaCiao, BASICSTAR, AnvilEcho

Charming Kitten (also tracked as APT42 by Mandiant, Mint Sandstorm by Microsoft, TA453 by Proofpoint, MITRE G0059/G1044) is an Iranian state-sponsored threat group operating on behalf of the IRGC Intelligence Organization. The group specializes in highly targeted social engineering campaigns that build rapport with victims over weeks before delivering credential-harvesting links or malware. Since January 2025, the group has operated over 130 phishing domains using React-based phishing kits capable of real-time MFA token interception, defeating all MFA methods except FIDO2 security keys and passkeys.

The group's target set includes policy researchers, journalists covering the Middle East, nuclear security experts, political campaign staff, and Iranian diaspora activists. In August 2024, Google TAG publicly attributed APT42 to targeting both the Trump and Biden/Harris presidential campaigns. The group continues high-tempo operations against Israeli academics, cybersecurity professionals, and journalists through 2025 and 2026.

This guide provides **10 hunt modules** designed to help defenders: (1) detect phishing proxy infrastructure and MFA bypass attempts, (2) identify OAuth abuse and session hijacking, (3) hunt for custom Charming Kitten backdoors (PowerStar, BellaCiao, NICECURL, TAMECAT), and (4) detect credential harvesting, lateral movement, and cloud exfiltration patterns. Total estimated hunt time is **5 hours**.

Threat Overview

Charming Kitten's primary attack vector is social engineering. Operators create convincing personas impersonating journalists, academics, think tank researchers, or cybersecurity professionals. They initiate benign email or WhatsApp conversations, build trust over days or weeks, then deliver credential harvesting links disguised as conference invitations, shared documents, or Google Meet links.

The group's phishing infrastructure acts as a real-time proxy. When a victim enters credentials and an MFA code on the cloned login page, the infrastructure immediately relays those values to the legitimate service, captures the authenticated session cookie, and establishes attacker access before the TOTP token expires. This approach defeats SMS-based MFA, authenticator app codes, and push notifications. Only cryptographically bound methods (FIDO2/passkeys) are immune.

After gaining access, Charming Kitten targets cloud email (Gmail, Outlook), documents, and contacts. The group exfiltrates data using native cloud features and open-source tools. For higher-value targets, the group deploys custom backdoors including PowerStar (PowerShell), BellaCiao (.NET with DNS-based C2), NICECURL (VBScript), and TAMECAT (PowerShell).

KEY CHARACTERISTICS

- IRGC-IO state-sponsored (Iran)
- Multi-week relationship building
- Real-time MFA token interception
- 130+ phishing domains (NameCheap)
- React-based phishing kits
- WhatsApp for initial contact
- Fake Google Meet invites
- PowerStar/NICECURL/TAMECAT backdoors
- BellaCiao DNS-based C2
- Cloud email/document exfiltration

2 MITRE ATT&CK Mapping

The following techniques are sourced from MITRE ATT&CK (G0059, G1044), Mandiant, Google TAG, Volexity, Proofpoint, and Bitdefender:

`T1566.002` CORE
Spearphishing Link
Primary delivery method; links to React-based phishing pages mimicking Google, Outlook, Yahoo

`T1656` CORE
Impersonation
Operators impersonate journalists, academics, cybersecurity professionals to build trust

`T1111` CORE
Multi-Factor Authentication Interception
Real-time phishing proxy captures and replays MFA tokens before TOTP expiration

`T1539`
Steal Web Session Cookie
Session cookies captured via phishing proxy enable persistent account access

`T1598.003`
Spearphishing for Information
Multi-week email conversations to gather intelligence before payload delivery

`T1585.002`
Establish Email Accounts
Fake email accounts at plausible institutions to support impersonation personas

`T1583.001`
Acquire Domains
130+ phishing domains registered via NameCheap since January 2025

`T1059.001`
PowerShell
PowerStar, TAMECAT, AnvilEcho all use PowerShell for execution and C2

`T1059.005`
Visual Basic
NICECURL and BASICSTAR use VBScript for initial access and backdoor operations

`T1530`
Data from Cloud Storage
Exfiltration from OneDrive, Google Drive, and cloud mailboxes after credential compromise

`T1114.002`
Remote Email Collection
Bulk email download from compromised cloud mailboxes (Gmail, Outlook)

`T1098.002`
Additional Email Delegate Access
Mail forwarding rules and delegate access added for persistent email collection

3 Hunt Modules

1 Phishing Proxy Detection (Token Interception)

Duration: 30 min Priority: CRITICAL

Objective: Detect sign-ins that originated from phishing proxy infrastructure, where the user's credentials and MFA token were relayed through an attacker-controlled intermediary to the legitimate identity provider.

Background

Charming Kitten operates React-based phishing kits that clone Google, Microsoft, and Yahoo login pages. These kits act as real-time proxies: when a victim enters credentials and an MFA code, the kit immediately forwards them to the real identity provider, captures the resulting session token, and passes the legitimate page back to the victim. The attack completes in seconds. Detecting this requires examining sign-in logs for anomalies such as unfamiliar IP addresses, non-compliant devices, or impossible travel patterns associated with successful MFA completions. (Source: Mandiant, May 2024; Google TAG, August 2024)

Hunt Procedure: Successful Sign-ins from Suspicious Locations After MFA

DETECTION ONLY

KQL (MICROSOFT ENTRA ID / AZURE AD)

```
// Detect successful sign-ins from atypical locations that completed MFA
// Phishing proxy attacks complete MFA successfully from attacker infrastructure
// Source: Mandiant APT42 reporting (May 2024)
SigninLogs
| where ResultType == 0  // Successful sign-in
| where AuthenticationRequirement == "multiFactorAuthentication"
| where RiskLevelDuringSignIn in ("medium", "high")
    or RiskState == "atRisk"
| project TimeGenerated, UserPrincipalName, IPAddress, Location,
    DeviceDetail, AppDisplayName, AuthenticationMethodsUsed,
    RiskLevelDuringSignIn, ConditionalAccessStatus
| sort by TimeGenerated desc
```

Hunt Procedure: Impossible Travel with MFA Completion

DETECTION ONLY

KQL (MICROSOFT ENTRA ID / AZURE AD)

```
// Detect impossible travel: user signs in from two geographically distant
// locations within a short timeframe, both completing MFA
let timeWindow = 2h;
SigninLogs
| where ResultType == 0
| where AuthenticationRequirement == "multiFactorAuthentication"
| summarize Locations=make_set(Location),
    IPs=make_set(IPAddress),
    LocationCount=dcount(Location)
    by UserPrincipalName, bin(TimeGenerated, timeWindow)
| where LocationCount > 1
| sort by TimeGenerated desc
```

Hunt Procedure: Non-Compliant Device Sign-in After MFA (Splunk)

DETECTION ONLY

SPL (SPLUNK)

```
Detect MFA-completed sign-ins from devices that are not managed/compliant.
Phishing proxy sessions originate from attacker infra, not corporate devices.
index=azure_ad sourcetype="azure:signinlogs" ResultType=0
    AuthenticationRequirement="multiFactorAuthentication"
| where DeviceDetail.isCompliant="false" OR
    DeviceDetail.isManaged="false"
| stats count by UserPrincipalName, IPAddress, Location,
    DeviceDetail.operatingSystem, AppDisplayName
| where count > 3
| sort -count
```

- ☐ Review all successful MFA sign-ins from IP addresses not previously seen for each user (past 90 days)
- ☐ Investigate sign-ins flagged with "atRisk" or "medium/high" risk that still completed successfully
- ☐ Check for impossible travel patterns where both sessions completed MFA
- ☐ Examine the user-agent string for anomalies (e.g., browser mismatch with user's known devices)

MFA Completion Does Not Guarantee Legitimacy
Charming Kitten's phishing proxies complete MFA successfully because they relay the real token in real time. Do not dismiss sign-in alerts simply because MFA was completed. The combination of unfamiliar IP + successful MFA + sensitive account is the primary indicator of phishing proxy compromise.

2 MFA Token Replay and Session Hijacking Detection

Duration: 30 min Priority: CRITICAL

Objective: Detect session token reuse and replay attacks where an attacker uses a stolen session cookie to access cloud services without re-authenticating.

Background

After capturing an authenticated session via phishing proxy, Charming Kitten operators use the stolen session cookie from their own infrastructure. This creates a pattern where the same session token is used from two different IP addresses or geographic locations. Microsoft Entra ID's Continuous Access Evaluation (CAE) and token protection features can detect and mitigate some of these attacks, but many organizations have not enabled these controls. (Source: Mandiant, May 2024; Google TAG)

Hunt Procedure: Session Token Used from Multiple IPs

DETECTION ONLY

KQL (MICROSOFT ENTRA ID)

```
// Detect the same session ID used from multiple IP addresses
// Indicates token theft and replay from attacker infrastructure
SigninLogs
| where ResultType == 0
| where isnotempty(CorrelationId)
| summarize IPCount=dcount(IPAddress),
    IPs=make_set(IPAddress),
    Locations=make_set(Location),
    Apps=make_set(AppDisplayName)
    by UserPrincipalName, CorrelationId
| where IPCount > 1
| sort by IPCount desc
```

Hunt Procedure: Token Replay Detection via AADNonInteractiveUserSignInLogs

DETECTION ONLY

KQL (MICROSOFT ENTRA ID)

```
// Non-interactive sign-ins using refresh tokens from new IP addresses
// After phishing proxy capture, attacker uses refresh token silently
let KnownUserIPs = SigninLogs
    | where TimeGenerated > ago(30d)
    | where ResultType == 0
    | summarize KnownIPs = make_set(IPAddress) by UserPrincipalName;
AADNonInteractiveUserSignInLogs
| where TimeGenerated > ago(7d)
| where ResultType == 0
| join kind=leftouter KnownUserIPs on UserPrincipalName
| where not(IPAddress in (KnownIPs))
| project TimeGenerated, UserPrincipalName, IPAddress, AppDisplayName, ResourceDisplayName, DeviceDetail
| sort by TimeGenerated desc
```

- ☐ Search for session tokens (CorrelationId) used from more than one IP address
- ☐ Review non-interactive sign-in logs for refresh token usage from previously unseen IPs
- ☐ Verify Continuous Access Evaluation (CAE) is enabled for critical applications
- ☐ Check if Token Protection (token binding) is configured for high-value accounts

Enable Continuous Access Evaluation (CAE)
CAE forces real-time policy evaluation for token usage, reducing the window for stolen token exploitation from hours to near real-time revocation. Enable CAE in Conditional Access policies for all critical applications. This is the most effective server-side defense against session token replay.

3 Suspicious OAuth Application Grants

Duration: 25 min Priority: HIGH

Objective: Detect malicious or suspicious OAuth application consent grants that Charming Kitten may use to maintain persistent access to compromised cloud accounts.

Background

After initial credential compromise, threat actors sometimes register OAuth applications or trick users into granting consent to attacker-controlled applications. This provides persistent access to cloud resources (email, files, calendar) even if the user changes their password, because OAuth tokens operate independently of the user's credentials. Mandiant has documented APT42 using cloud service features for data exfiltration. (Source: Mandiant, May 2024)

Hunt Procedure: New OAuth App Consent Grants

DETECTION ONLY

KQL (MICROSOFT ENTRA ID)

```
// Detect OAuth application consent grants, especially with Mail.Read,
// Files.Read, or Contacts.Read permissions
AuditLogs
| where OperationName == "Consent to application"
| extend AppName = tostring(TargetResources[0].displayName)
| extend ConsentedBy = tostring(InitiatedBy.user.userPrincipalName)
| extend Permissions = tostring(AdditionalDetails)
| project TimeGenerated, ConsentedBy, AppName, Permissions,
    OperationName
| sort by TimeGenerated desc
```

Hunt Procedure: Google Workspace Third-Party App Access (Splunk)

DETECTION ONLY

SPL (SPLUNK)

```
Detect new third-party app authorizations in Google Workspace.
APT42 targets Gmail accounts; OAuth persistence is a common follow-up.
index=gws sourcetype="gws:reports:token"
    event.name="authorize"
| stats count by actor.email, app_name, scope, timestamp
| where like(scope, "%mail%") OR like(scope, "%drive%")
    OR like(scope, "%contacts%")
| sort -timestamp
```

- ☐ Review all OAuth application consent grants from the past 90 days
- ☐ Flag applications requesting Mail.Read, Files.ReadWrite, or Contacts.Read that are not on the organization's approved list
- ☐ Check for applications registered within the last 30 days with high-privilege permissions
- ☐ Verify Google Workspace Connected Apps for unknown third-party authorizations

OAuth Tokens Survive Password Changes
If an attacker has established OAuth application access, changing the compromised user's password alone will not revoke access. The OAuth tokens must be explicitly revoked, and the malicious application must be removed from the user's authorized applications list.

4 React Phishing Kit Domain Pattern Detection

Duration: 30 min Priority: HIGH

Objective: Detect network connections to domains matching patterns associated with Charming Kitten's phishing infrastructure, including recently registered domains mimicking legitimate login services.

Background

Since January 2025, Charming Kitten has registered over 130 unique domains and numerous subdomains through NameCheap, resolving to approximately a dozen IP addresses. These domains mimic Google, Microsoft, and Yahoo login pages using React-based single-page applications (SPAs). The phishing kits include real-time keyloggers that capture credentials as they are typed. Domains often use patterns like "accounts-google-[keyword].com", "login-outlook-[keyword].com", or similar variations. (Source: SecurityAffairs 2025, Google TAG)

Hunt Procedure: Connections to Recently Registered Domains Mimicking Login Services

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect connections to domains containing login-service keywords
// that were recently registered (common Charming Kitten infrastructure pattern)
DeviceNetworkEvents
| where RemoteUrl has_any (
    "accounts-google", "login-google", "signin-google",
    "accounts-outlook", "login-outlook", "signin-microsoft",
    "login-yahoo", "mail-verification", "account-verify",
    "security-check", "2fa-verify", "mfa-verify")
| where RemoteUrl !endswith "google.com"
    and RemoteUrl !endswith "microsoft.com"
    and RemoteUrl !endswith "yahoo.com"
    and RemoteUrl !endswith "live.com"
| project Timestamp, DeviceName, RemoteUrl, RemoteIP,
    InitiatingProcessFileName, InitiatingProcessCommandLine
| sort by Timestamp desc
```

Hunt Procedure: DNS Queries to NameCheap-Hosted Suspicious Domains (Splunk)

INVESTIGATION AID

SPL (SPLUNK)

```
Hunt for DNS queries to recently registered domains hosted on NameCheap
infrastructure that mimic authentication services.
NOTE: NameCheap hosts millions of legitimate domains; use as enrichment only.
index=dns sourcetype="dns"
| where query_type="A"
| regex query="(?i)(account|login|signin|verify|auth|security).*(google|outlook|microsoft|yahoo)"
| where NOT match(query, "(?i)\.(google|microsoft|yahoo|live|outlook)\.(com|net|org)$")
| stats count by query, answer, src_ip
| sort -count
```

Hunt Procedure: Proxy/Firewall Logs for Phishing Page Indicators

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect connections to phishing-pattern domains from non-browser processes
DeviceNetworkEvents
| where Timestamp > ago(24h)
| where ActionType == "ConnectionSuccess"
| where RemoteUrl matches regex @"(login|signin|verify|account|secure|auth)[.-]"
| where RemoteUrl !endswith "microsoft.com"
    and RemoteUrl !endswith "google.com"
    and RemoteUrl !endswith "apple.com"
    and RemoteUrl !endswith "okta.com"
    and RemoteUrl !endswith "auth0.com"
| where InitiatingProcessFileName !in~ ("chrome.exe", "msedge.exe", "firefox.exe", "brave.exe")
| project Timestamp, DeviceName, RemoteUrl, RemoteIP,
    InitiatingProcessFileName, InitiatingProcessCommandLine
| sort by Timestamp desc
```

- ☐ Query DNS logs for domains containing authentication-related keywords that do not resolve to legitimate service provider infrastructure
- ☐ Review proxy logs for HTTPS connections to recently registered domains (under 30 days old) mimicking login services
- ☐ Cross-reference any identified domains with threat intelligence feeds for known Charming Kitten infrastructure
- ☐ Check for URL shortener usage (bit.ly, tinyurl) in email or messaging logs that redirect to suspicious login pages

Domain Lists Require Continuous Updates
Charming Kitten registers new domains frequently. Static blocklists become stale within days. Complement domain-based detection with behavioral indicators (MFA anomalies, impossible travel) and category-based blocking (newly registered domains) at the proxy/DNS layer.

5 Email Impersonation and Persona-Based Spearphishing Detection

Duration: 30 min Priority: HIGH

Objective: Detect inbound emails that impersonate known contacts, journalists, or academic researchers, consistent with Charming Kitten's multi-stage social engineering approach.

Background

Charming Kitten operators invest significant time building convincing personas. They register email accounts using names of real journalists, professors, or think tank researchers at lookalike domains or free email providers. The initial email is typically benign (a question about research, an interview request, or a conference invitation) and contains no malicious content. The malicious link or attachment arrives only after multiple rounds of correspondence. Proofpoint's TA453 reporting documents operators impersonating NPR journalists, university professors, and research institute staff. (Source: Proofpoint 2024, Volexity February 2024, Google TAG)

Hunt Procedure: Emails from Lookalike Domains of Known Institutions

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR OFFICE 365)

```
// Detect inbound emails from domains that closely resemble known research
// institutions, media outlets, or universities
// Customize the domain patterns for your organization's contacts
EmailEvents
| where EmailDirection == "Inbound"
| where SenderFromDomain has_any (
    "harvard", "mit", "stanford", "brookings", "rand",
    "carnegie", "chatham", "reuters", "bbc", "nytimes",
    "washingtonpost", "rasanah", "middleeast")
| where SenderFromDomain !endswith ".edu"
    and SenderFromDomain !endswith ".org"
    and SenderFromDomain !in (
        "harvard.edu", "mit.edu", "stanford.edu",
        "brookings.edu", "rand.org", "reuters.com",
        "bbc.co.uk", "nytimes.com")
| project Timestamp, SenderFromAddress, SenderFromDomain,
    RecipientEmailAddress, Subject, DeliveryAction
| sort by Timestamp desc
```

Hunt Procedure: Emails with Conference/Webinar/Podcast Invitation Lures

INVESTIGATION AID

KQL (MICROSOFT DEFENDER FOR OFFICE 365)

```
// Detect emails with subject lines matching Charming Kitten lure themes
// Sources: Volexity (fake webinar), Proofpoint (fake podcast)
EmailEvents
| where EmailDirection == "Inbound"
| where Subject has_any (
    "webinar invitation", "podcast invitation",
    "interview request", "research collaboration",
    "conference invitation", "panel discussion",
    "keynote invitation", "feedback form")
| where SenderFromDomain !in (
    // Add your organization's known partner domains here
    "zoom.us", "webex.com", "teams.microsoft.com")
| project Timestamp, SenderFromAddress, RecipientEmailAddress,
    Subject, DeliveryAction, AuthenticationDetails
| sort by Timestamp desc
```

- ☐ Review inbound emails from domains that closely resemble known research institutions or media outlets
- ☐ Search for emails with subjects containing "interview", "webinar", "podcast", "conference", or "research collaboration" from external senders
- ☐ Check email authentication results (SPF, DKIM, DMARC) for suspected impersonation emails
- ☐ Interview high-risk users (researchers, executives, journalists) about unexpected contact from unknown individuals claiming shared professional interests

Initial Emails May Be Entirely Benign
Charming Kitten's first email in a social engineering chain typically contains no malicious links, attachments, or indicators. Email security tools will not flag it. The malicious payload arrives only after trust has been established through multiple exchanges. Train users to verify the identity of new contacts through independent channels (e.g., calling the institution's published phone number) before engaging with unsolicited correspondence.

6 PowerStar/CharmPower Backdoor Hunting

Duration: 30 min Priority: HIGH

Objective: Detect PowerStar (also known as CharmPower, MITRE S0674) and related PowerShell-based backdoors (TAMECAT, AnvilEcho) used by Charming Kitten for persistent access and intelligence collection.

Background

PowerStar/CharmPower is a modular PowerShell backdoor with capabilities for remote command execution, persistence management, screenshot capture, process enumeration, file search, and a cleanup module. TAMECAT (S1193) is a PowerShell toehold that communicates via HTTP with Base64-encoded data. AnvilEcho is a PowerShell intelligence collection trojan with data exfiltration via FTP and Dropbox. These tools are typically delivered through macro-enabled documents, LNK files, or RAR archives after the social engineering phase. (Source: MITRE ATT&CK S0674, Mandiant May 2024, Proofpoint August 2024)

Hunt Procedure: Suspicious PowerShell Module Loading Patterns

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect PowerShell execution patterns consistent with modular backdoors
// PowerStar downloads modules from C2; TAMECAT receives Base64 commands
DeviceProcessEvents
| where FileName in~ ("powershell.exe", "pwsh.exe")
| where ProcessCommandLine has_any (
    "IEX", "Invoke-Expression",
    "DownloadString", "DownloadFile",
    "Net.WebClient", "Invoke-WebRequest",
    "[Convert]::FromBase64String",
    "New-Object System.Net.Sockets",
    "Start-BitsTransfer")
| where ProcessCommandLine !has "WindowsUpdate"
    and ProcessCommandLine !has "PSWindowsUpdate"
| project Timestamp, DeviceName, ProcessCommandLine,
    InitiatingProcessFileName, AccountName
| sort by Timestamp desc
```

Hunt Procedure: NICECURL VBScript Execution Chain

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// NICECURL (S1192) is delivered via LNK files that execute VBScript
// The VBS downloads a PDF decoy and the NICECURL backdoor
// Source: Mandiant (May 2024)
DeviceProcessEvents
| where FileName in~ ("wscript.exe", "cscript.exe")
| where InitiatingProcessFileName has_any (
    "cmd.exe", "explorer.exe", "powershell.exe")
| where ProcessCommandLine has ".vbs"
| project Timestamp, DeviceName, FileName, ProcessCommandLine,
    InitiatingProcessFileName, InitiatingProcessCommandLine
| sort by Timestamp desc
```

Hunt Procedure: LNK File Delivering Backdoor (NICECURL Pattern)

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect LNK files that spawn scripting engines (VBS/PowerShell)
// NICECURL uses LNK -> VBS -> download chain
DeviceProcessEvents
| where InitiatingProcessCommandLine has ".lnk"
| where FileName in~ ("wscript.exe", "cscript.exe",
    "powershell.exe", "cmd.exe", "mshta.exe")
| project Timestamp, DeviceName, FileName, ProcessCommandLine,
    InitiatingProcessFileName, InitiatingProcessCommandLine
| sort by Timestamp desc
```

- ☐ Search for PowerShell processes using IEX, DownloadString, or Base64 decoding in command lines
- ☐ Hunt for VBScript execution spawned by LNK files (NICECURL delivery chain)
- ☐ Check for macro-enabled documents (.docm, .xlsm) that spawned PowerShell (TAMECAT delivery)
- ☐ Review scheduled tasks for persistence entries executing PowerShell scripts from unusual paths

7 BellaCiao Backdoor Detection (DNS-Based C2)

Duration: 30 min Priority: HIGH

Objective: Detect BellaCiao implants that use DNS resolution as a covert command-and-control channel, and BellaCPP variants that establish SSH tunnels for data exfiltration.

Background

BellaCiao is a .NET dropper that performs a DNS query every 24 hours to resolve a victim-specific subdomain. The DNS server (attacker-controlled) returns an IP address that encodes instructions. The implant compares this IP against hardcoded values to determine its next action. This produces no active C2 traffic visible to network monitoring; only periodic DNS queries are generated. Each BellaCiao sample is customized per victim, containing hardcoded company names and subdomains. BellaCPP (December 2024) is a C++ variant that replaces the web shell component with SSH tunneling. (Source: Bitdefender April 2023, Kaspersky December 2024)

Hunt Procedure: Periodic DNS Queries to Unusual Subdomains

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect processes making periodic DNS queries to subdomains with
// unusual naming patterns (BellaCiao uses victim-specific subdomains)
DeviceNetworkEvents
| where ActionType == "DnsQueryResponse"
| where RemoteUrl matches regex @"^[a-z0-9]{8,32}\.[a-z0-9-]+\.[a-z]{2,6}$"
| where RemoteUrl !endswith ".microsoft.com"
    and RemoteUrl !endswith ".windows.net"
    and RemoteUrl !endswith ".azure.com"
    and RemoteUrl !endswith ".google.com"
    and RemoteUrl !endswith ".googleapis.com"
| summarize QueryCount=count(),
    UniqueHosts=dcount(DeviceName)
    by RemoteUrl, bin(Timestamp, 1d)
| where QueryCount >= 1 and QueryCount <= 3  // BellaCiao queries once per 24h
| sort by QueryCount desc
```

Hunt Procedure: .NET Process Making Periodic DNS Queries

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// BellaCiao is a .NET executable; hunt for .NET processes with
// periodic DNS resolution behavior
DeviceNetworkEvents
| where ActionType == "DnsQueryResponse"
| where InitiatingProcessFileName !in~ (
    "chrome.exe", "firefox.exe", "msedge.exe",
    "outlook.exe", "teams.exe", "svchost.exe")
| summarize DailyQueries=count(),
    Domains=make_set(RemoteUrl)
    by DeviceName, InitiatingProcessFileName, bin(Timestamp, 1d)
| where DailyQueries between (1 .. 5)
| where array_length(Domains) <= 3
| sort by DailyQueries asc
```

Hunt Procedure: SSH Tunnel Establishment by Non-Standard Process (BellaCPP)

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// BellaCPP establishes SSH tunnels for covert data exfiltration
// Hunt for non-standard processes connecting on port 22
DeviceNetworkEvents
| where RemotePort == 22
| where InitiatingProcessFileName !in~ (
    "ssh.exe", "putty.exe", "plink.exe",
    "scp.exe", "sftp.exe", "winscp.exe",
    "git.exe", "code.exe")
| project Timestamp, DeviceName, InitiatingProcessFileName,
    RemoteIP, RemotePort, LocalPort
| sort by Timestamp desc
```

- ☐ Review DNS query logs for single daily queries to unusual subdomains (BellaCiao beacon pattern)
- ☐ Search for .NET executables in non-standard locations that perform periodic DNS resolution
- ☐ Hunt for SSH connections (port 22) initiated by non-standard processes
- ☐ Check for scheduled services or startup entries running unknown .NET assemblies

BellaCiao Is Victim-Specific
Each BellaCiao sample contains hardcoded victim information (company name, subdomains, IP addresses). If a BellaCiao implant is detected, the DNS subdomain and resolved IP addresses are forensic evidence specific to this intrusion. Preserve all DNS query logs for the affected host immediately.

8 Credential Harvesting Infrastructure Detection

Duration: 30 min Priority: HIGH

Objective: Detect indicators that credentials have been harvested from the organization, including password spray patterns, LSASS access, browser credential extraction, and keylogging behavior.

Background

Beyond phishing proxy operations, Charming Kitten also conducts on-host credential harvesting. MITRE documents the group using Mimikatz (S0002) for LSASS memory dumping (T1003.001), keylogging (T1056.001), and browser credential extraction (T1555.003 via APT42). Recent reporting indicates a custom Chromium-based credential stealer targeting Chrome, Opera, Brave, and Edge that extracts encrypted keys from Local State files. (Source: MITRE ATT&CK G0059/G1044, Dark Atlas 2025)

Hunt Procedure: Browser Credential Store Access

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect processes accessing browser credential stores
// APT42 uses custom Chromium-based credential stealers (T1555.003)
DeviceFileEvents
| where FileName in~ ("Login Data", "Local State", "Cookies")
| where FolderPath has_any (
    "Google\\Chrome\\User Data",
    "Microsoft\\Edge\\User Data",
    "BraveSoftware\\Brave-Browser\\User Data",
    "Opera Software\\Opera Stable")
| where ActionType in ("FileRead", "FileModified")
| where InitiatingProcessFileName !in~ (
    "chrome.exe", "msedge.exe", "brave.exe", "opera.exe")
| project Timestamp, DeviceName, FileName, FolderPath,
    InitiatingProcessFileName, InitiatingProcessCommandLine
| sort by Timestamp desc
```

Hunt Procedure: LSASS Memory Access (Mimikatz Pattern)

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// Detect LSASS credential dumping (Mimikatz, S0002)
// Magic Hound / G0059 uses Mimikatz per MITRE ATT&CK
DeviceProcessEvents
| where FileName in~ ("mimikatz.exe", "procdump.exe",
    "procdump64.exe", "MiniDump.exe")
    or ProcessCommandLine has_any (
    "sekurlsa::logonpasswords", "sekurlsa::wdigest",
    "lsass.dmp", "lsass.exe")
| project Timestamp, DeviceName, FileName, ProcessCommandLine,
    InitiatingProcessFileName, AccountName
| sort by Timestamp desc
```

Hunt Procedure: Keylogger Detection

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR ENDPOINT)

```
// APT42 deploys keyloggers (T1056.001) for credential capture
// Detect processes using keyboard hook APIs
DeviceProcessEvents
| where ProcessCommandLine has_any (
    "SetWindowsHookEx", "GetAsyncKeyState",
    "GetKeyState", "keylog", "KeyLogger")
| where InitiatingProcessFileName !in~ (
    "explorer.exe", "svchost.exe", "dwm.exe")
| project Timestamp, DeviceName, FileName, ProcessCommandLine,
    InitiatingProcessFileName, AccountName
| sort by Timestamp desc
```

- ☐ Search for non-browser processes accessing Chrome, Edge, Brave, or Opera credential stores
- ☐ Review all LSASS access events and correlate with authorized security tooling
- ☐ Hunt for processes using keyboard hook APIs that are not part of known accessibility tools
- ☐ Check for browser processes being terminated by external processes (credential stealer pattern)

9 Compromised Account Lateral Movement

Duration: 25 min Priority: MEDIUM

Objective: Detect signs that a compromised cloud account is being used for lateral movement, including email forwarding rule creation, delegate access, and contact list harvesting for further targeting.

Background

After compromising a cloud account via phishing proxy, Charming Kitten operators establish persistence and expand access. Common post-compromise actions include creating inbox forwarding rules to an external address (T1098.002), adding delegate access to allow another account to read the mailbox, harvesting the contact list for further spearphishing, and using the compromised account to send phishing emails to the victim's contacts (leveraging the trust relationship). (Source: Mandiant May 2024, Proofpoint TA453 reports)

Hunt Procedure: New Email Forwarding Rules

DETECTION ONLY

KQL (MICROSOFT 365 / EXCHANGE ONLINE)

```
// Detect new inbox forwarding rules (persistence indicator)
// Charming Kitten may forward email to attacker-controlled accounts
OfficeActivity
| where Operation in ("New-InboxRule", "Set-InboxRule",
    "Set-Mailbox", "New-TransportRule")
| where Parameters has_any ("ForwardTo", "ForwardAsAttachmentTo",
    "RedirectTo", "DeliverToMailboxAndForward")
| project TimeGenerated, UserId, Operation,
    Parameters, ClientIP
| sort by TimeGenerated desc
```

Hunt Procedure: Delegate Access or Permission Changes

DETECTION ONLY

KQL (MICROSOFT 365)

```
// Detect mailbox delegate access additions
// Attackers add delegate access for persistent email reading
OfficeActivity
| where Operation in ("Add-MailboxPermission",
    "AddFolderPermissions", "ModifyFolderPermissions")
| project TimeGenerated, UserId, Operation,
    Parameters, ClientIP
| sort by TimeGenerated desc
```

Hunt Procedure: Compromised Account Sending Phishing to Contacts

DETECTION ONLY

KQL (MICROSOFT DEFENDER FOR OFFICE 365)

```
// Detect outbound emails from organizational accounts containing
// suspicious URLs or matching phishing lure patterns
EmailEvents
| where EmailDirection == "Outbound"
| where Subject has_any (
    "webinar invitation", "conference invitation",
    "shared document", "interview request",
    "research collaboration", "Google Meet")
| join kind=inner (
    EmailUrlInfo
    | where Url has_any ("accounts-google", "login-outlook",
        "signin-google", "verify-account")
) on NetworkMessageId
| project Timestamp, SenderFromAddress, RecipientEmailAddress,
    Subject, Url
| sort by Timestamp desc
```

- ☐ Review all new email forwarding rules created in the past 30 days
- ☐ Audit delegate access permissions for all high-value mailboxes
- ☐ Check for bulk contact list exports from compromised accounts
- ☐ Monitor for outbound emails from organizational accounts containing phishing URLs

Compromised Accounts Amplify the Campaign
Charming Kitten leverages compromised accounts to send phishing emails to the victim's contacts. The emails come from a trusted sender, significantly increasing click-through rates. If an account compromise is confirmed, immediately review all outbound email from that account for the past 30 days and notify recipients of any suspicious messages.

10 Cloud Storage Exfiltration Detection

Duration: 25 min Priority: MEDIUM

Objective: Detect data exfiltration from cloud storage services (OneDrive, Google Drive, SharePoint) after credential compromise, and identify abuse of cloud services for outbound data transfer.

Background

Mandiant documents APT42 exfiltrating data using native features and open-source tools. After gaining access to cloud accounts, operators download email archives, documents, and contact lists. AnvilEcho (Proofpoint, August 2024) exfiltrates data via FTP and Dropbox. The group also uses cloud storage sharing features to share sensitive documents with attacker-controlled external accounts. (Source: Mandiant May 2024, Proofpoint August 2024, MITRE T1530)

Hunt Procedure: Bulk File Downloads from Cloud Storage

DETECTION ONLY

KQL (MICROSOFT 365)

```
// Detect bulk file downloads from OneDrive/SharePoint
// APT42 exfiltrates data using native cloud features (T1530)
OfficeActivity
| where TimeGenerated > ago(24h)
| where Operation in ("FileDownloaded", "FileSyncDownloadedFull")
| where OfficeWorkload in ("OneDrive", "SharePoint")
| summarize DownloadCount = count(),
    Files = make_set(OfficeObjectId, 50)
    by UserId, ClientIP, bin(TimeGenerated, 1h)
| where DownloadCount > 50
| sort by DownloadCount desc
```

Hunt Procedure: External Sharing of Sensitive Documents

DETECTION ONLY

KQL (MICROSOFT 365)

```
// Detect documents shared externally from OneDrive/SharePoint
// Attackers may share files with external accounts for exfiltration
OfficeActivity
| where Operation in ("SharingSet", "AddedToSecureLink",
    "AnonymousLinkCreated")
| where OfficeObjectId has_any (".pdf", ".docx", ".xlsx",
    ".pptx", ".zip", ".rar")
| project TimeGenerated, UserId, Operation, OfficeObjectId,
    ClientIP
| sort by TimeGenerated desc
```

Hunt Procedure: Google Workspace Drive Exfiltration (Splunk)

DETECTION ONLY

SPL (SPLUNK)

```
Detect bulk file downloads or sharing from Google Drive.
APT42 targets Google Workspace accounts.
index=gws sourcetype="gws:reports:drive"
    (event.name="download" OR event.name="change_acl_editors"
    OR event.name="change_document_access_scope")
| stats count by actor.email, target_user, doc_title, event.name
| where count > 20
| sort -count
```

- ☐ Review bulk file download events (>50 files per hour) from cloud storage for each user
- ☐ Audit external sharing links created in the past 30 days, especially for sensitive document types
- ☐ Check for Google Takeout or similar bulk export requests
- ☐ Monitor for FTP or Dropbox uploads from endpoints where AnvilEcho may be present

A Appendix A: Indicators of Compromise

Blocking Guidance: 3-Tier IOC Classification

Tier 1: Block (High Confidence)

All malware file hashes (PowerStar, NICECURL, TAMECAT, BellaCiao, AnvilEcho, BlackSmith): Confirmed Charming Kitten backdoors with no legitimate use case. Safe for hash-based blocking and EDR prevention policies.

Confirmed phishing domains from vendor-published IOC lists: Safe for DNS/proxy blocking.

Tier 2: Alert / Monitor (Shared Infrastructure)

Cloud service URLs: Charming Kitten abuses Google Sites, Dropbox, and other legitimate cloud platforms for phishing page hosting and data exfiltration. Blocking these services will disrupt legitimate operations. **Recommend contextual monitoring only.**

VPN/VPS provider IPs: The group uses commercial VPN and VPS infrastructure. Blocking provider IP ranges will affect legitimate users. **Recommend monitoring and analyst review before any action.**

Tier 3: Investigate (Unconfirmed)

NameCheap-registered domains: Over 130 phishing domains were registered through NameCheap (source: SecurityAffairs, Google TAG). However, NameCheap hosts millions of legitimate domains. NameCheap registration alone is not an IOC. Use as contextual enrichment only.

Recently registered domains mimicking login services: Not all such domains are Charming Kitten. Verify against published vendor IOC lists before blocking.

Malware Families

Name	MITRE ID	Type	First Seen	Source
CharmPower / PowerStar	S0674	PowerShell Modular Backdoor	2022	MITRE ATT&CK, Volexity
NICECURL	S1192	VBScript Backdoor	January 2024	Mandiant (May 2024)
TAMECAT	S1193	PowerShell Toehold	2024	Mandiant (May 2024)
BellaCiao	N/A	.NET Dropper (DNS C2)	April 2023	Bitdefender (April 2023)
BellaCPP	N/A	C++ Backdoor (SSH tunneling)	December 2024	Kaspersky (December 2024)
BASICSTAR	N/A	VB Backdoor	2023	Volexity (February 2024)
PowerLess	S1012	PowerShell/.NET Backdoor	2022	Cybereason (February 2022)
AnvilEcho	N/A	PowerShell Intel Trojan	August 2024	Proofpoint (August 2024)
BlackSmith	N/A	Malware Loader	August 2024	Proofpoint (August 2024)

IOCs Requiring Verification
The following vendor reports contain specific file hashes, domains, and IPs that should be obtained directly from the original publications for deployment in detection rules:

- **Mandiant**: SHA256 hashes for NICECURL and TAMECAT; phishing domain list; C2 infrastructure
- **Volexity**: YARA rules for BASICSTAR and POWERSTAR (GitHub)
- **Proofpoint**: SHA256 hashes for BlackSmith loader and AnvilEcho trojan
- **Bitdefender**: SHA256 hashes for BellaCiao samples; victim-specific DNS subdomains
- **Google TAG**: 130+ phishing domain list from 2025 credential harvesting campaign

B Appendix B: Detection Rules

KQL Detection Rules

Phishing Proxy MFA Bypass Detection

KQL

DETECTION ONLY

MICROSOFT ENTRA ID

```
// RULE: Detect successful MFA sign-ins from high-risk or atypical locations
SigninLogs
| where ResultType == 0
| where AuthenticationRequirement == "multiFactorAuthentication"
| where RiskLevelDuringSignIn in ("medium", "high")
    or RiskState == "atRisk"
| extend AlertSeverity = "High"
| project TimeGenerated, UserPrincipalName, IPAddress, Location,
    AppDisplayName, RiskLevelDuringSignIn, AlertSeverity
```

Session Token Replay Detection

KQL

DETECTION ONLY

MICROSOFT ENTRA ID

```
// RULE: Detect same session used from multiple IP addresses (token theft)
SigninLogs
| where ResultType == 0
| summarize IPCount=dcount(IPAddress),
    IPs=make_set(IPAddress)
    by UserPrincipalName, CorrelationId
| where IPCount > 1
| extend AlertSeverity = "Critical"
```

Email Forwarding Rule Creation

KQL

DETECTION ONLY

MICROSOFT 365

```
// RULE: Detect new email forwarding rules (post-compromise persistence)
OfficeActivity
| where Operation in ("New-InboxRule", "Set-InboxRule", "Set-Mailbox")
| where Parameters has_any ("ForwardTo", "ForwardAsAttachmentTo",
    "RedirectTo", "DeliverToMailboxAndForward")
| extend AlertSeverity = "High"
| project TimeGenerated, UserId, Operation, Parameters,
    ClientIP, AlertSeverity
```

Phishing Domain Pattern Detection

KQL

DETECTION ONLY

MICROSOFT DEFENDER FOR ENDPOINT

```
// RULE: Detect connections to domains mimicking login services
DeviceNetworkEvents
| where RemoteUrl has_any (
    "accounts-google", "login-google", "signin-google",
    "accounts-outlook", "login-outlook", "signin-microsoft")
| where RemoteUrl !endswith "google.com"
    and RemoteUrl !endswith "microsoft.com"
| extend AlertSeverity = "High"
| project Timestamp, DeviceName, RemoteUrl, RemoteIP,
    InitiatingProcessFileName, AlertSeverity
```

PowerShell Backdoor Execution Pattern

KQL

DETECTION ONLY

MICROSOFT DEFENDER FOR ENDPOINT

```
// RULE: Detect PowerShell backdoor patterns (PowerStar, TAMECAT, AnvilEcho)
DeviceProcessEvents
| where FileName in~ ("powershell.exe", "pwsh.exe")
| where ProcessCommandLine has_any (
    "IEX", "Invoke-Expression", "DownloadString",
    "[Convert]::FromBase64String", "Net.WebClient")
| where ProcessCommandLine !has "WindowsUpdate"
| extend AlertSeverity = "High"
| project Timestamp, DeviceName, ProcessCommandLine,
    InitiatingProcessFileName, AlertSeverity
```

Sigma Rule: Phishing Proxy MFA Bypass

Sigma Rule Disclaimer
This Sigma rule requires conversion to your SIEM's native query language using sigmac or pySigma. Test thoroughly in a non-production environment before deployment.

DETECTION ONLY

SIGMA (YAML)

```
title: Successful MFA Sign-in from Risky Location (Charming Kitten Phishing Proxy)
id: c3d4e5f6-a7b8-9012-cdef-b02216000001
status: experimental
description: |
    Detects successful sign-ins that completed MFA from locations flagged
    as risky by the identity provider. Charming Kitten's phishing proxy
    infrastructure relays MFA tokens in real time, causing sign-ins to
    appear as successful MFA completions from attacker infrastructure.
references:
    - https://cloud.google.com/blog/topics/threat-intelligence/untangling-iran-apt42-operations
    - https://attack.mitre.org/groups/G1044/
author: Intruvent Technologies
date: 2026-03-18
tags:
    - attack.credential_access
    - attack.t1111
    - attack.t1539
logsource:
    product: azure
    service: signinlogs
detection:
    selection:
        ResultType: 0
        AuthenticationRequirement: 'multiFactorAuthentication'
    risk_condition:
        RiskLevelDuringSignIn:
            - 'medium'
            - 'high'
    condition: selection and risk_condition
falsepositives:
    - Legitimate users traveling to new locations
    - VPN exit nodes triggering location-based risk
    - New device enrollment with MFA
level: high
```

C Appendix C: References

MITRE ATT&CK

MITRE ATT&CK: Magic Hound (G0059) - Primary MITRE group profile for APT35/Charming Kitten

MITRE ATT&CK: APT42 (G1044) - MITRE group profile for APT42

MITRE ATT&CK: CharmPower (S0674)

MITRE ATT&CK: NICECURL (S1192)

MITRE ATT&CK: TAMECAT (S1193)

Mandiant / Google Threat Intelligence

Mandiant (May 2024): "Uncharmed: Untangling Iran's APT42 Operations"

Mandiant (2022): "APT42: Crooked Charms, Cons, and Compromises"

Google TAG (August 2024): APT42 targeting U.S. presidential campaigns

Volexity

Volexity (February 2024): "CharmingCypress: Innovating Persistence"

Volexity YARA Rules for CharmingCypress (GitHub)

Proofpoint

Proofpoint (August 2024): "TA453 Targets Religious Figure with Fake Podcast Invite Delivering New BlackSmith Malware Toolset"

Bitdefender / Kaspersky

Bitdefender (April 2023): "Unpacking BellaCiao: A Closer Look at Iran's Latest Malware"

CyberSRC (December 2024): "Iran's Charming Kitten Deploys BellaCPP"

Other Vendor Research

SecurityAffairs (2025): "APT42 impersonates cyber professionals to phish Israeli academics and journalists"

The Hacker News (May 2024): "APT42 Hackers Pose as Journalists to Harvest Credentials"

Dark Atlas: "Threat Profile: APT35 (Charming Kitten)"

Intruvent Golden CTI Database (Codex): Actor profiles for Magic Hound (G0059) and APT42 (G1044)

NEED HELP IMPLEMENTING THESE DETECTIONS?

Our threat hunting team can deploy, tune, and operationalize these detection rules in your environment.

✓ Detection rule deployment and tuning
✓ Managed threat hunting operations
✓ Incident response retainer

Contact Our Team

or

Automate continuous detection with BRACE, our sector-specific threat intelligence platform. See BRACE plans →