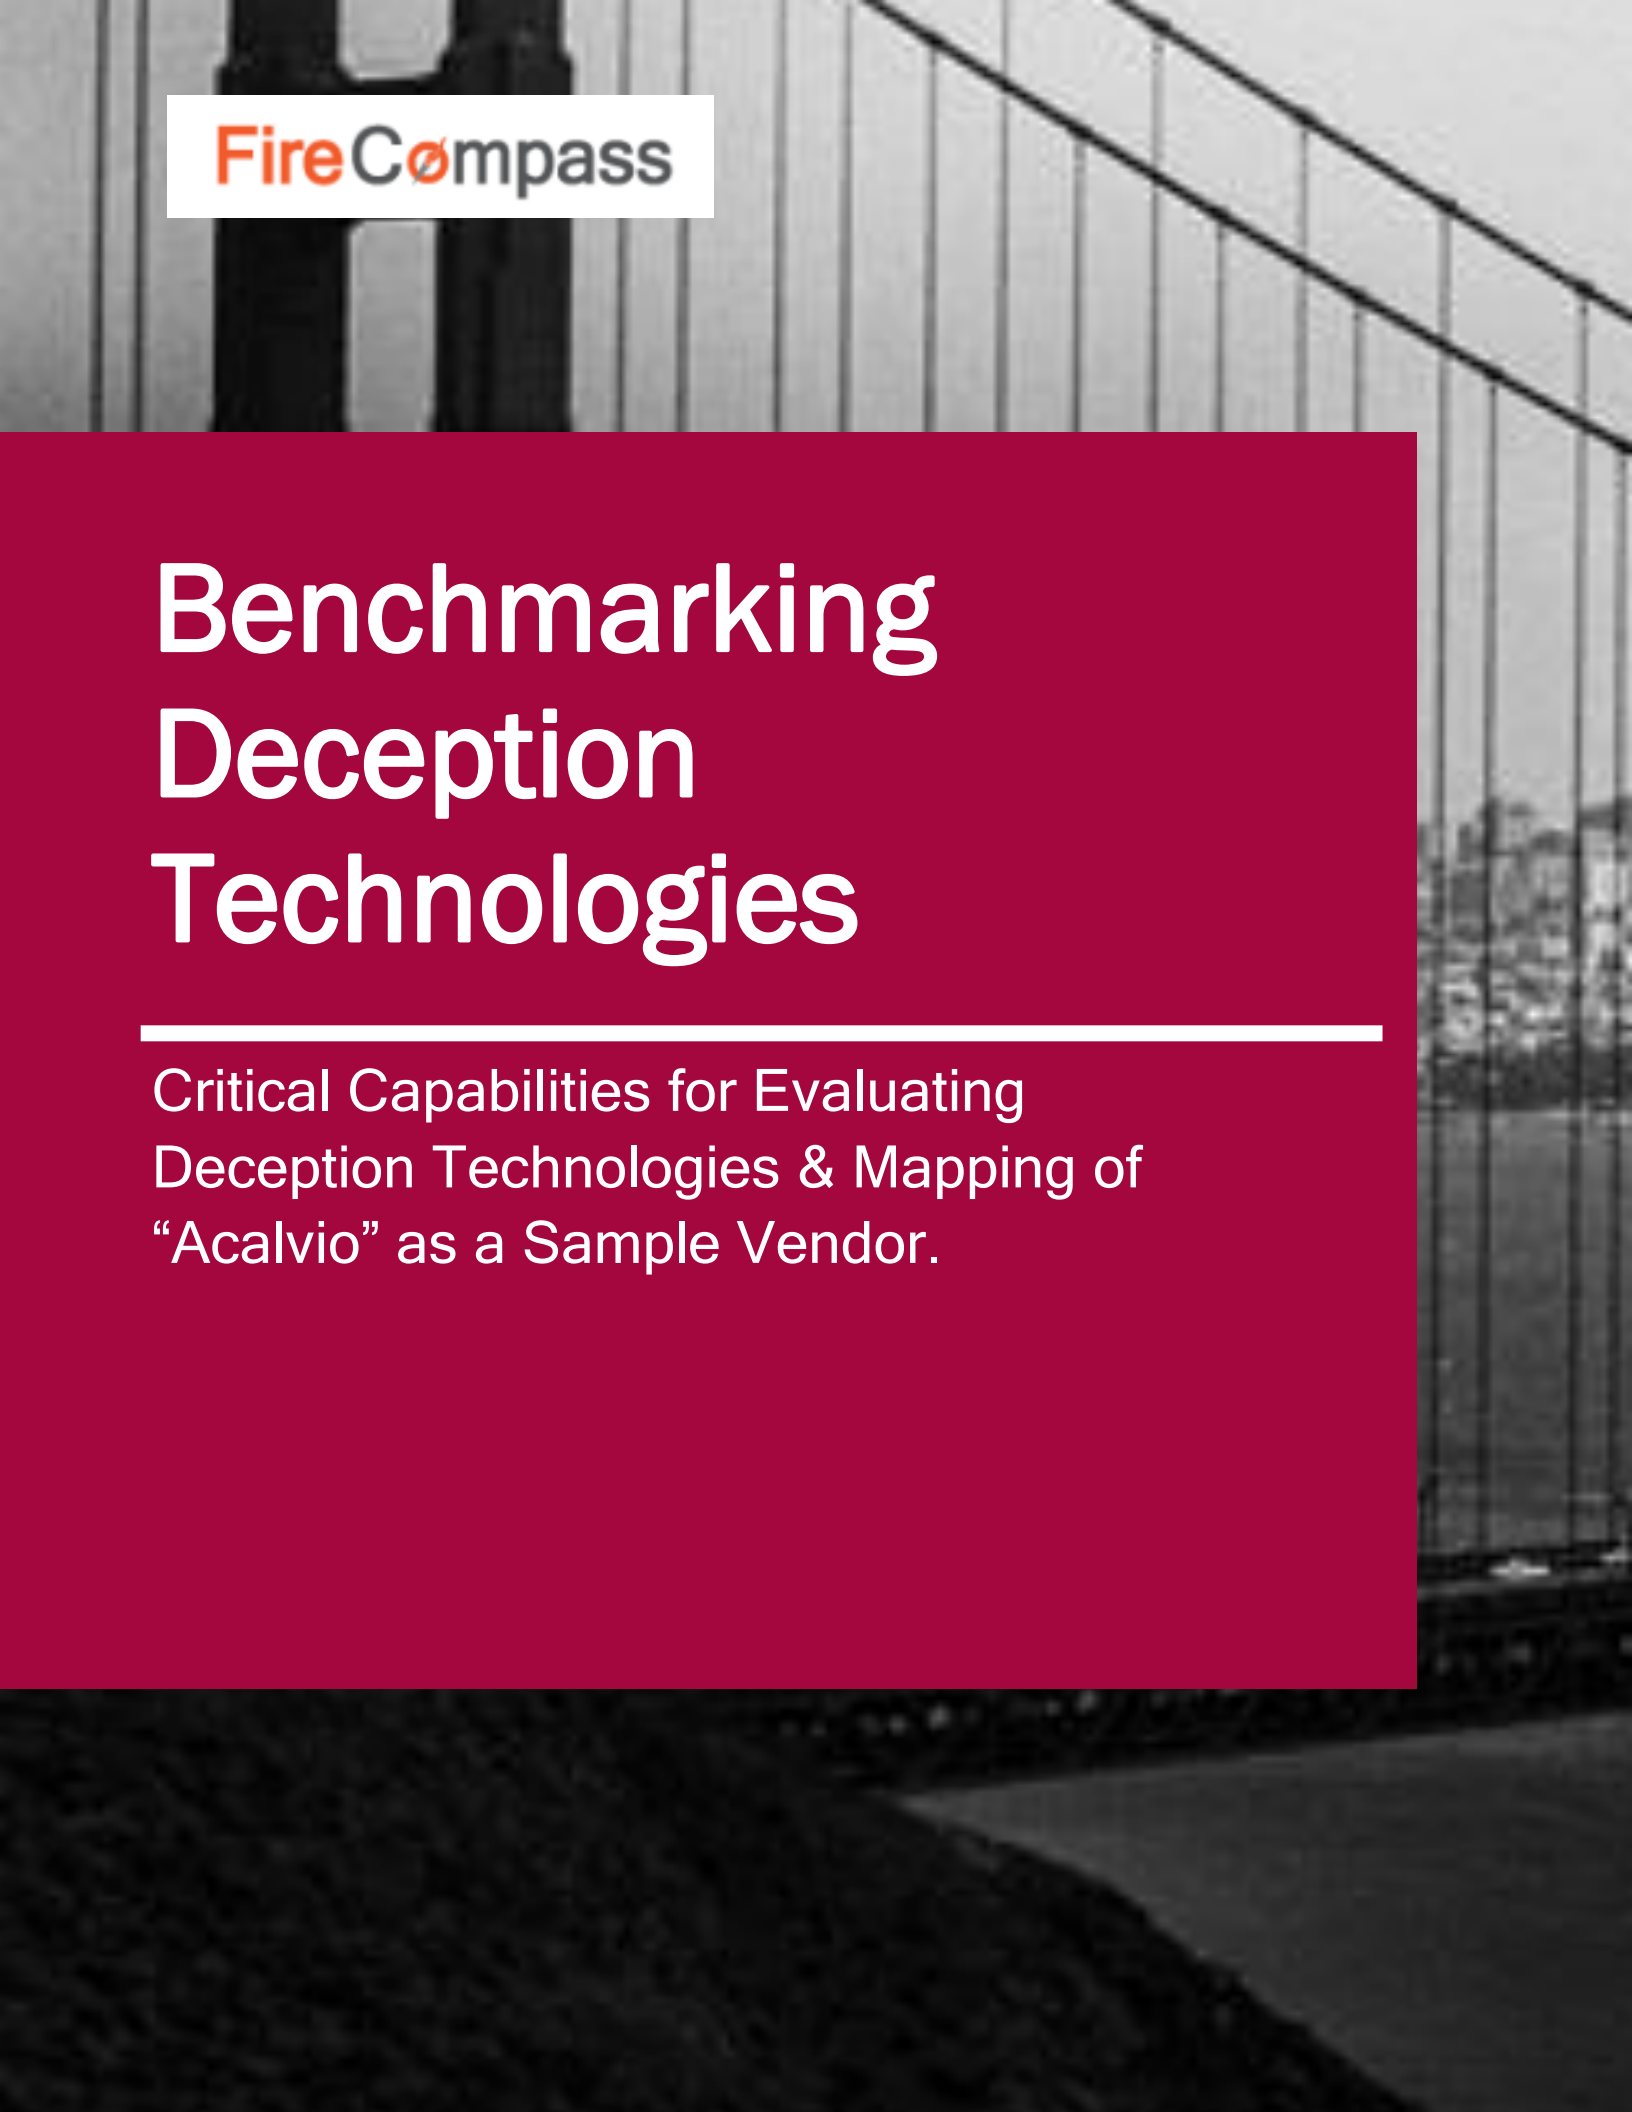

FireCompass

Benchmarking Deception Technologies

Critical Capabilities for Evaluating Deception Technologies & Mapping of “Acalvio” as a Sample Vendor.

Table of Contents

Introduction

Traditionally, organizations have relied on security controls such as host/network-based IDS, SIEM, and UTM as the primary mechanism for detecting threats inside the perimeter, which is no longer adequate. With the evolving threat landscape, breaches are becoming increasingly common. As per Ponemon Institute, in 2017, the average time required to detect a breach was 191 days and the average time required to contain a breach was 66 days. These time frames can have a significant impact on businesses relying on technology for their continued operations.

This has led to a renewed focus on threat detection and response—how to detect faster, analyze better, and respond quickly. While there are a number of technologies in this space (see below for the sample list), this report explores the rapidly evolving deception market, which is witnessing a speedy growth.

	Detect	Respond
Devices	Deception, EDR	Deception, EDR
Applications	DDoS	DDoS
Networks	Deception, NTA	Deception
Data	Web monitoring	
Users	UEBA	

Table 1: Sample List of Cyber Security Technology for Detection and Response

Deception as a strategy has always been a critical element for detecting threats and adversary behavior. Honeypots have been used as a deception mechanism for over a decade, both by organizations and security vendors, to research attack tactics, techniques, and procedures and provide threat intelligence.

Organizations can deploy a deception solution:

- **Inside the network** – with the primary aim of detecting and responding to attacks that have bypassed the perimeter controls, that is, the attacker has breached the perimeter and is inside the network. Various types of decoys can be set at different places inside the network (e.g., devices, applications, network, and data) to try and detect the attacker. This is the most common approach and will be the focus of this report.

- **Outside the perimeter/in DMZ** – with the primary aim of gathering high-quality threat intelligence. If not tuned carefully, the deception solution can generate many alerts for common activities (e.g., automated web security scanners), which will defeat the purpose.

Deception can help disrupt the following phases of the kill chain (assuming the attacker is already inside the network):

- **Recon** – With a number of high-quality decoys (details on decoys provided later in the text) placed in the environment, the solution makes it extremely difficult for attackers to create an accurate picture of the infrastructure. If an attacker interacts with any of the decoys, they have a high chance of being detected.

- **Lateral movement** – When attackers try to move laterally, there is a high chance that they might move toward a decoy and get caught. Lures and breadcrumbs are spread throughout the infrastructure, making it even more likely that the attacker will move toward a decoy.

- **Exfiltration** – Once the attacker reaches a decoy, organizations can closely observe TTPs of attackers, identify factors such as the C&C destination and type of data being targeted, and use this information to block exfiltration from the other parts of the organization.

Types of Threats

Deception solutions are primarily used to address the following types of threats:

- **Malware (including ransomware)**: Deception solutions can help prevent the spread of malware/ransomware by luring the malware to decoys (e.g., network file shares). Devices can be set up with lures and breadcrumbs in such a manner that if a machine is infected, the malware is directed toward decoys where they can be investigated and contained.

- **Advanced persistent threats**: Because deception does NOT rely on signatures, but on interactions with decoys, the solution can be used to detect targeted attacks (including use of zero days), understand the adversary behavior, and provide an appropriate response.

- **Insider threats and stolen credentials**: Deception can help flag suspicious insider behavior when insiders access decoys. Accidental access can be filtered out automatically by the solution, but suspicious activities or repeated visit to decoys will trigger alerts, which can be investigated further. External attackers using stolen credentials can be caught in a similar manner, and if the organization already knows about compromised credentials, then they can straightaway divert attackers towards decoy networks.

Deception Technology Use Cases

Here are some of the key benefits of using deception technology:

1. Enhanced Threat Detection with High Accuracy

With sophisticated attacks becoming increasingly common, organizations have increasingly started focusing on detection and response capability when (and not if) they get breached.

Deception solutions rely on obfuscating the real assets/infrastructure with decoys, making it highly likely that a single interaction with a decoy can lead to discovery. Vendors have multiple approaches to determine the quantity and quality of obfuscation, but the intent is the same—to provoke the attacker to make some sort of interaction with a decoy. As soon as the interaction happens, alerts/automated responses get triggered.

Accuracy- Because the alerts are only triggered when interactions happen with decoys, and most deception solutions have some sort of basic filtering in place for accidental interactions, the alerts are highly accurate and need deeper investigation/response.

2. Understand Attacker TTPs and Motivation, Engagement

Very few technology solutions have the capability of engaging the attackers to understand the TTPs and attacker motivation. Deception solutions allow organizations to continually monitor the attack behavior and conduct an in-depth analysis.

After detection, deception solutions can help as follows:

- View the trail – Trace the attacker movement during the preceding steps and investigate the TTPs used at various stages in the past.

- Engage attackers – Transition the attackers to a sandboxed environment where the attacker TTPs and behavior can be observed closely to determine the likely motivation.

3. Slowing Down the Attackers

Deception can help significantly slow down the attackers as they try to move laterally within the network. Identifying the right target among a large number of decoys will be tough and require attackers to do a significant level of analysis before they pick their next target within the network.

4. Accelerated Response to Attacks

Ability to help substantially automate and orchestrate response for minimizing the impact is one of the drivers for deception technology. Response capabilities vary widely based on vendors, with some being as simple as triggering an incident for SIEM/SOAR solutions to some having the ability to define pre-defined workflows for investigation and response.

5. Automate Threat Hunting

Deception solutions can help automate a part of the threat-hunting process. Alerts/notable events from SIEM can be passed to the deception solution, which can then dynamically set up relevant decoys around the hosts, which are a source of alerts. These decoys can then help refine the alerts for further investigation (e.g., based on interactions with decoys).

6. Secure Specialized Environments

Deception solutions can help protect specialized environments, some of which are difficult to secure using traditional enterprise security controls. These include the following environments:

- **IoT, industrial IoT, and SCADA** – These are usually highly specialized environments, where minimal security controls (if any) exist. Traditional approaches such as using AV or EDR will not work. Most environments will also have a large number of unpatched vulnerabilities due to various reasons. Deception solutions provide a unique ability of protecting these environments by creating a large number of decoys, which makes it difficult for attackers to identify the right target.

- **Healthcare** – EHR terminals, medical diagnostic terminals, etc. can be protected in a similar manner as IoT environments. Deception solutions can help protect against ransomware attacks (the healthcare sector has been a frequent target), as well as in regulatory compliance (e.g., HIPAA).

- **Banking infrastructure** – Deception solutions can be used to create different decoys in banking infrastructure such as ATM and SWIFT servers, and any interaction with decoys can trigger response playbooks.

7. Accurate Asset Inventory

Some deception solutions (based on how they are deployed) can more accurately provide a picture of the current state of IT asset inventory in various network segments (e.g., type of assets, basic configuration, and logon/logoff times) than the other existing solutions can.

Technology Overview

Deception technology comprises a number of key components, which are briefly discussed below:

	Decoys	Breadcrumbs	Lures
Purpose	Fake systems and services mimicking real IT assets in the infrastructure	Redirecting attackers from real assets to decoys	Makes decoys attractive to attackers
Example	Devices, servers, and applications	Mapped network drive and credentials	Vulnerability with known exploit
IT impact	Requires virtualization, cloud, and/or hardware. Requires IPs to be allocated	Minimal, uses the existing infrastructure to deploy baits	NA, uses decoys
Impact on actual users	Nil/minimal	Some, as users might accidentally interact	Nil

Table 2: Key Components of Deception Technology

- **Decoys:** These are perhaps the single most important factors that determine the success of a deception solution. Decoys or traps are fake systems, applications, or services that are added to the network and look similar to regular assets in an organization's IT infrastructure (via either virtualization or emulations).

- **Breadcrumbs (aka Baits):** These are planted in IT assets to redirect/divert attackers from real assets toward decoys, for example, hidden mapped network drives that lead malwares to decoys.

- **Lures:** These are used to make decoys attractive to attackers and malwares. These are deliberately placed to make a decoy more enticing for attackers when they are looking for targets during reconnaissance, for example, a decoy server with apparent misconfiguration or unpatched vulnerability like Heartbleed.

- **Management/administration:** This is a key component as it connects everything together (decoys, lures, breadcrumbs, and alerts). Management of decoys includes the following:
    - Where to deploy what type of decoy, breadcrumb, and lure
    - Manage and keep the decoys, breadcrumbs, and lures updated
    - Personalization and customization of decoys
    - Manage deployments at scale
    - Monitor alerts
    - Manage integrations with other systems (e.g., active director, SIEM, and EDR) and playbooks

Key Evaluation Parameters

- **Variety of deception**- Different vendors support a variety of deceptions. Check what type of deceptions the solution supports: decoys, lures, and breadcrumbs.

- **Decoys** – Genuineness of decoys is perhaps the single most important evaluation criterion. The success of deception deployment depends on the ability to attract attackers to decoys and engage them. Below are some parameters that can be used to evaluate decoys:

    - **Types**: The ability to emulate a variety of IT assets present in the organization. These may include the following:
        - Devices (including various environments such as Linux, Mac, and Windows)
        - Applications
        - Network
        - Servers (including DB server)
        - Industry-specific environments (e.g., SCADA, IoT, Healthcare, and SWIFT)

    - **Authenticity**: Decoys have to appear authentic; otherwise, attackers will ignore them. Most deception solutions support integration with AD/LDAP servers to make the decoys look more authentic. Some decoys can more easily be made to appear authentic (e.g., using gold images for device decoys) than others (e.g., applications and data).

    - **Level of interactivity/credibility of decoys**: To make decoys credible, they should have some level of interactivity. Low interactivity decoys can be easily detected, whereas high interactivity decoys can be used to engage the attackers.

    - **Scalability/density**: One way to increase the chances of catching attackers is to increase the quantity of decoys spread across the network. The solution should be able to scale thousands of decoys, without degrading the performance of the environment, while providing the ability to engage and respond to attacks.

    - **Relevance**: When deploying decoys, they should be relevant to specific environments. Some deception solutions can provide recommendations around this (by placing sensors that can analyze the environment), while others need to be set up manually.

- **Adaptability/blending/updation**: Staleness is an enemy of deception. As the network and threat environments evolve, deception must adapt. The solution should have the ability to blue print environment profiling capabilities to dynamically adapt decoys to the environment around them, be it applications, OS, naming conventions, services, patch levels, file contents, etc.

- **Fingerprinting**: Decoys should be difficult to fingerprint and not follow easy-to-guess patterns. Some examples include MAC addresses that are not sequential/follow a pattern or do not align with the other devices and using an OS flavor that is not used in the environment (e.g., using Ubuntu when the organization is on RedHat).

- **Breadcrumbs**: They play a crucial role in getting attackers from real assets to decoys. Support for a wide variety of breadcrumbs is important. However, keep in mind that breadcrumbs should not get easily triggered by legitimate users of the system; otherwise, this can lead to many false alerts. Some examples of breadcrumbs to lookout for are as follows:
  - Mapped network drives
  - Credentials
  - Browser content (saved passwords, history, cookies, and bookmarks)
  - Local files

- **Lures**: These play an important role in attracting attackers to decoys. A variety of lures supported by a deception solution can be an important factor during evaluation. Some examples are as follows:
  - Running seemingly important services on decoys, which get discovered during network discovery (e.g., DB or FTP services)
  - Vulnerable Services on decoys (with known exploits)

- **Engage capabilities**: Although this is not for every organization, for some organizations, this can help shortlist deception solutions to a considerable degree. Check the type of engagement capabilities that the solution supports, some of which are as follows:
  - Ability to seamlessly transition attackers to a different environment for observing TTPs and intentions
  - Beaconing capabilities – Check the decoy file movement after it leaves the source

- **Response capabilities**: These vary considerably based on vendors. Although most support integration with SIEM and incident management tools, some also offer the following advanced capabilities:
    - Tracing attacker trajectory before reaching the decoy
    - Automated IOC generation
    - Inbuilt workflows and orchestration support

- **Automatic and intelligent**: An enterprise-scale deception solution needs to lay out a multitude of deceptions and manage them dynamically. Automation of every step is a requirement for practical deception at scale. Machine intelligence is imperative for automation.

- **Data driven**: A deception solution must be driven by the vulnerabilities in the network and the current threat landscape. Integration with the SIEM and cyber threat feeds is essential for effective deception.

- **Fitment with existing security stack**: Integrations with other security solutions in your environment, ranging from SIEM to EDR, will reduce operational overheads and improve effectiveness.

Closely watch for vendor capabilities in this space, as they vary considerably.

- Some vendors only support breadcrumbs or lures to attract the attacker to the decoy server and do NOT support a variety or large volume of decoys

- Some vendors only support low interaction decoys for scalability

- Check for level of automation offered by the deception solution (e.g., by using AI or ML) in terms of planning and executing a deception strategy. For instance, can the solution automatically do the following:

    - Recommend the type of deception (decoys and breadcrumbs) and the location for placing them

- Deploy realistic/genuine-looking decoys and breadcrumbs, with minimal human intervention, which cannot be easily fingerprinted. Other parameters to check may include naming conventions followed, quality of data/files, placement in the network, and blending with the environment

- Keep the decoys fresh/updated. Keeping pace with the changing IT environment

- Refine and prioritize the alerts (either from decoys or fed from SIEM). Can the solution automatically deploy based on alerts from SIEM/SOC to help filter or prioritize alerts?

Evaluation Checklist (Sample Vendor: Acalvio)

A checklist can be used to evaluate and compare different deception solutions as outlined below. Organizations can customize this based on their specific requirements.

	Acalvio	Comments
1. Use cases		
a. What is required		
i. Detect, engage, and respond	Yes	All three capabilities are natively supported
b. Detection of attack chain/kill chain		
i. Recon	Yes	
ii. Lateral movement	Yes	
iii. Exfiltration	Yes	
c. Attacks types		
i. Malware (e.g., ransomware)	Yes	
ii. Advanced persistent threats	Yes	
iii. MITM	Yes	
d. Internal vs external threats	Both	
e. Industry-specific capability	Yes	
2. Deception capabilities		
a. Types of deceptions		
i. Decoys	Yes	
ii. Breadcrumbs/baits	Yes	
iii. Lures/vulnerabilities	Yes	
b. Decoys		
i. Types of Decoys—enterprise endpoint decoy, database decoy, network shares decoy, cloud decoy, social decoys	All	Doesn't currently supports SCADA decoys
ii. Platforms		
1. OS—Windows, Linux (RHEL, CentOS,	All	

Debian, etc.), Mac		
2. Types—Workstations, Servers	All	
iii. Interaction levels—low, medium, high	All	
iv. Support for virtualization and real system	Both	
v. Industry-specific decoys (e.g., IoT and SCADA)	Yes	IoT devices such as cameras and printers are supported but doesn't currently supports SCADA decoys
c. Breadcrumbs		
i. Types		
1. Files/folders	Yes	
2. Browser profiles	Yes	
3. Registry	Yes	
4. In-memory	Yes	
5. User profiles	Yes	
6. Network connections/listeners	No	
ii. Platforms (specify versions)		
1. Windows	Yes	
2. Linux	Yes	
3. Mac	Yes	
d. Use custom images (e.g., gold image) as decoys	Yes	
e. Deception freshness/update frequency for decoys, breadcrumbs and lures	Real-time	
f. Containment of high-interaction decoy—access to internet, neighborhood, access enterprise resources (AD, etc.), activity against another machine	Yes	
g. Visualization of deception deployment	Yes	

h. Blending of decoys with the IT infrastructure	Yes	
i. Memory forensics and sandbox for malware detonation	No	No automatic memory forensics after an attack on the decoy
3. Management of deceptions		
a. Deployment automation	Yes	
b. Auto refresh	Yes	
c. Personalization	Yes	
d. Customization	Yes	
4. Reporting and logging		
a. Customize reports	Yes	
b. Export report	Yes	
c. Pre-built reports	Yes	
d. Dashboard	Yes	
e. Schedule/email reports	Yes	
5. Scalability		
a. Decoys across geo-locations	Yes	
b. Decoys across subnets	Yes	
c. Maximum number of decoys	-	
6. Form factors (deployment options)		
a. Cloud to on-premise	Yes	
b. On-prem to on-prem (hardware appliance—single centralized appliance)	Yes	
c. On-prem to on-prem (software appliance—single centralized appliance)	Yes	
d. Cloud to cloud	Yes	
e. MSSP/MDR support	Yes	
7. Product security		
a. Password policy	Yes	
b. DOS protection	Yes	
c. Hardening	Yes	
d. Pen test/red teaming	Yes	
e. Deception fingerprinting	Yes	
f. Custom OS used for the	No	

platform		
8. Use of AI/machine learning		
a. For a deception campaign strategy	Yes	
b. Deception campaign execution	Yes	
c. Threat analysis and threat hunting	Yes	
9. REST API support	Yes	
10. Incidents (alerts)		
a. Attacker activity monitoring—PCAP, TTP (IOC), commands typed	Yes	
b. Incident management—alert prioritization	Yes	
c. Whitelisting	Yes	
d. Incident policies—what to send to SIEM/notification	Yes	
11. Out-of-the-box integrations (ecosystem fitment)		
a. SIEM	Yes	
b. EDR	Yes	
c. Sandbox/ATP (Cuckoo, FireEye)	Yes	Some
d. Orchestration engine (Phantom, Demisto)	Yes	Some
e. Active directory	Yes	
f. Firewall and IDS	Yes	
g. TI consumption and creation	Both	
h. Network integrations	Yes	
i. IP address—static, dynamic	Yes	
ii. Switch integration—trunk, access	Yes	
iii. DHCP	Yes	
iv. DNS	Yes	
12. Product administration		
a. Support SSO	Yes	
b. Built-in RBAC	Yes	
c. Create users/roles	Yes	

13. Support		
a. Remote troubleshooting	Yes	
b. Incident management	Yes	
c. UI-based upgrade/patch management	Yes	
d. 24*7*365 support available	Yes	
e. Quality of technical support	Yes	
f. Support presence in a local city, country, etc.	Yes	Multi-country presence
g. Direct support or partner	Yes	
h. SLA, TAT, escalation matrix, etc.	Yes	
14. Total cost of ownership		
a. Initial cost	-	
b. Setup and implementation cost	-	
c. Recurring subscription costs	-	
d. Patch update and upgrade cost	-	
e. Software licensing cost for decoys (e.g., windows license)	-	
f. Integration cost	-	
g. Hardware/appliance cost	-	
h. Any other cost	-	
15. Vendor background		
a. Market share, turnover, profitability	<$10M/year	
b. Number of patents filed (attach links)	30+	
c. Customer base	<20	

Table 3: Evaluation Checklist for Deception Technology

Besides the parameters outlined above, organizations can also consider the following:

- Attaching the vendor response to evaluation criteria (such as one outlined above) or RFP, which specifies the features that are out of the box, configuration changes, or customizations, to the final contract.

- The level of support a vendor can provide for defining the deception strategy/program. Deception is not just about technology; organizations should consider the full spectrum of the deception strategy/program and how a vendor can support them.

- Doing a red teaming exercise during PoC (if budget permits), where the red team knows that the deception solution is already in place, and evaluating the impact.

- Researching the vendor's R&D team on platforms such as LinkedIn. Examples include the following:
    - Size and qualifications of the R&D team
    - Background of senior R&D personnel—CTO, VP Engineering, Head of Research, etc.
    - If a vendor is claiming AI/machine learning capabilities, does the R&D team have people with credible background in the field

Disclaimers & Disclosures

The report has been written by the cyber security product analyst team from FireCompass – a SaaS platform for cyber security preparedness assessment and technology evaluation. This report should not be reproduced or distributed without written consent from FireCompass Inc.

Acalvio has procured the rights for distribution of this particular report from FireCompass. FireCompass team has conducted multiple interviews with Acalvio technical team and went through product demos to validate the said capabilities. We advise organization to use this report only as an example approach document for an evaluation process. For thorough and accurate evaluation we suggest organizations to conduct internal assessment, employ third party or engage with FireCompass professional team. Any opinion expressed in the report by FireCompass should not be considered as statement of facts and any such opinion is subject to change without any prior notice. FireCompass Inc disclaims all warranties including (by way of example and not by limitation) that of completeness, accuracy or adequacy of any information contained in the report and shall have no liability of any kind including (by way of example and not by limitation) that for errors & omissions.