Scope of Service

Rapid7 Managed Detection and Response (MDR Elite)

The Rapid7 MDR service is delivered as a collaboration between Rapid7 and the customer ("you", "your team", "your organization", "your environment"). The mission of Rapid7's MDR service is to leverage our experts to collaboratively advance your cybersecurity decision-making and maturity through our tailored guidance.

We pride ourselves on becoming a true extension of your security team. Our goal is to partner together to enhance your ability to detect and respond to threats with hands-on 24x7x365 monitoring, threat hunting, incident response, and customized security guidance to stop malicious activity and strengthen your security posture.

Any responsibilities and/or actions not explicitly defined in this Scope of Service are not considered part of the Rapid7 MDR Elite service. Additional details can be seen in the MDR Responsibilities Matrix.

All terminology not defined in this document can be viewed at: services.help.rapid7.com/docs/mdr-terms.

Joint Requirements For Ensuring Success

To ensure your organization realizes the full value of Rapid7 MDR, it is critical that both parties share in the responsibilities and requirements of the partnership for effective delivery of the MDR service:

Rapid7 Responsibilities and Requirements

	Responsibilities and Requirements
1	Monitor the customer's environment in accordance with the detection methodologies outlined in this Scope of Service and with the visibility provided by the Rapid7 MDR technology stack (InsightIDR & Insight Agent), and in conjunction with the event sources configured in InsightIDR from the customer environment.
2	Provide a knowledgeable information security professional with subject matter expertise in Rapid7 products to assist with the deployment of required and optional product features.
3	Provide a named security advisor ("Customer Advisor") as the point-of-contact for the MDR relationship and help to accelerate the customer's security maturity.
4	Perform Remote Incident Response engagement in addition to standard alert and incident investigations.
5	Provisioning and ongoing management of Rapid7 cloud services in the technology stack.
6	Delivery of all reports via the Rapid7 Services Portal in accordance with this Scope of Service.
7	Notify the customer to any Customer Advisor or service delivery changes to Rapid7 MDR.

Customer Responsibilities and Requirements

	Responsibilities and Requirements
1	Acknowledge, accept, and adhere to all requirements and actions outlined in this Scope of Service
2	License all endpoints within the in-scope environment(s). The in-scope environment(s) must be 'logically separated' from any other out-of-scope environments.
3	Designate a Point-of-Contact to work with Rapid7 for deployment and onboarding.
3a	Complete the Site Survey prior to starting the deployment.
3b	Complete the Service Request Form prior to being placed in service.
3c	Provide Rapid7 with relevant documentation (i.e., policies, procedures, diagrams, flow charts, etc.) so the MDR service can fit seamlessly into your process, including an escalation path for reporting incidents.

4	**Deploy Insight Agents** to all workstation, desktop, and server assets in the in-scope environment(s) and **connect all Insight Agents** to InsightIDR. ● Assets without the Insight Agent deployed will not be fully supported by the MDR service. ● Performing a Compromise Assessment and Monthly Threat Hunts require deployment of Insight Agents to at least 80% of the in-scope environment.
4a	Ensure the Rapid7 Insight Agent is up to date. The Rapid7 MDR service supports the current version of the Agent and up to two previous versions as signified by a change in either the ones (x), tenths (y) or hundredths (z) of a version (x.y.z).
5	Allocate and configure **Insight Collector(s)**. At least one Insight Collector must be provisioned. The collector is required in order to: ● Collect the event sources described in 6 and 7. ● Proxy connections from 'on premise' InsightAgents to the Insight Platform.
6	Connect **all available recommended data sources** to InsightIDR for each in-scope environment and ensure availability and connectivity to Rapid7 infrastructure for all MDR technology and Event Sources.
6a	For organizations that have Microsoft Windows domains, send the **Windows Security event logs** from each domain controller to InsightIDR using one of the supported methods. Without this event source, many of Rapid7 MDR's UBA detections will not be supported.
6b	For organizations that have Microsoft Windows domains, Microsoft Azure AD Domain Services, or Amazon AWS Domain Services, connect at least one **LDAP event source** for each domain. Without this event source, Rapid7 MDR will not have vital contextual information about users in your environment.
6c	Connect all supported **DHCP** log sources to InsightIDR. Without this event source, Rapid7 MDR's SOC may not be able to accurately attribute network traffic to the appropriate assets in your environment.
6d	Connect all supported network logs - **DNS, firewall, VPN,** and **Web Proxy** - to InsightIDR, particularly those network devices at your Internet egress points. These log sources are extensively used for real-time threat detection, hunting, and investigations.
6e	Connect all supported **Cloud Service** logs to InsightIDR. These log sources are extensively used for real-time threat detection, hunting, and investigations.
6f	Deploy **Insight Network Sensors** in your environment to analyze and log network traffic data.
6g	Set up **Honey Credentials, Honey Users, Honey files** and **Honeypots**.
7	Connect **any other security-relevant event sources** to InsightIDR *Note: All connected event sources may be leveraged for investigation and incident response purposes. Not all event sources will be used for real-time threat detection or hunting. For more detail, see 'Event Sources'.*
8	Notify Rapid7 to any personnel, technology, event source, or point-of-contact changes or modifications.
9	Configure the InsightIDR instance in accordance with the recommendations from Rapid7's deployment team and Customer Advisors.
10	Review alerts that are not in-scope for the MDR service. These alerts will **not** be reviewed by the MDR SOC. See 'Alert Review Limitations' appendix for details.
11	Respond to 'Requests for Information' (RFIs) from the MDR SOC regarding specific alerts. These requests may be sent via e-mail or as cases in the Support Portal. The MDR SOC requires additional feedback from the customer in order to accurately assess these alerts.

Security Expertise

The MDR team is composed of a named **Customer Advisor,** the customer's **MDR SOC "Pod"**, and the **Rapid7 Threat Intelligence Team**.

Customer Advisor

Your **Customer Advisor ("CA")** is the main point-of-contact for the Rapid7 MDR service. This named resource works with your team as a strategic security partner—from initial technology deployment through incident remediation and ongoing security consultation—to shepherd your organization's security maturity. Your CA will be assigned to your organization at the time of your Service Delivery Kickoff call.

Throughout the service your CA will frequently communicate with your team to provide updates on service delivery, reporting, metrics, technology health, and ensure we are helping you address your security goals. Additionally, each CA is aligned to your assigned MDR SOC Pod ("Pod") to maintain constant communication and understand all relevant knowledge pertaining to investigations and incidents.

MDR SOC Pod

Your organization's environment will be assigned to one of our **SOC Pods** staffed by our world-class analysts to ensure continuous 24x7x365 monitoring coverage for real-time alert investigation, threat hunting, and incident response. Pods are comprised of three levels of Security Analysts and managed by a Pod Lead:

Analyst	Distribution	Description
Associate Analyst	3 per Pod	Responsible for alert triage and investigation and threat hunting.
Analyst	2 per Pod	Responsible for alert triage and investigation, threat hunting, alert tuning, and supporting Remote Incident Response engagements.
Senior Analyst	1 per Pod	Responsible for alert tuning, threat hunting, leading Remote Incident Response engagements, training other analysts, and handles escalated investigations.
Pod Lead	1 per Pod	Manages the SOC teams. Responsible (along with the named Customer Advisor) for the MDR service delivered to their team's assigned customers.

Threat Intelligence Team

Rapid7's **Threat Intelligence team** supports the MDR SOC and CAs with threat analysis and detections for new vulnerabilities, exploits, and attack campaigns found via their research. These new detections are added as detections for all MDR customers.

MDR Technology

The Rapid7 MDR service is powered by InsightIDR, Rapid7's own threat-focused Cloud SIEM, Endpoint Detection and Response (EDR), and User Behavior Analytics (UBA) solution, to provide comprehensive protection against intruders in your organizations internal network, devices, and cloud services. InsightIDR and the MDR SOC leverage the Insight Agent and other event sources from your existing security infrastructure to ensure visibility into threats across the environment.

A full list of all Rapid7 cloud technologies and customer-deployed software leveraged by the Rapid7 MDR service can be viewed at services.help.rapid7.com/docs/mdr-terms.

InsightIDR Instance Set-Up

By default, the MDR Elite service includes a single instance of InsightIDR for your entire organization. All log sources will be onboarded to this single instance. Additionally, all Insight Platform users (on both InsightIDR & Services Portal) will be assigned to, and will have access to, all data stored within this single instance.

In some cases, your business may want to deploy the MDR service to multiple 'organizations' as separate InsightIDR instances to provide separate visibility and reporting across these 'logically separated' organizations (such as business units). Details on the MDR service delivery for additional organizations can be provided in an additional Addendum titled "*Scope of Service Addendum - Additional Organizations*".

In-Scope Environment for MDR

Rapid7 MDR requires licensing and deployment of the Insight Agent across your organization's entire environment to have the best coverage and monitoring of malicious activity. In certain instances, your team may choose to license a partial portion of your environment for MDR ("in-scope") as long as this environment meets the qualifications of a 'logically separated' environment. For example, an Internet-facing production data center that is separate from your corporate IT end-user environment. Or your environment may include multiple subsidiaries with logically separate IT infrastructures.

In these situations, Rapid7 recommends that your organization deploy our MDR service to all of your environments for the following reasons:

- Attackers often move laterally within an organization from one environment to another, and without a full deployment to all environments we may be unable to detect or respond to the full scope of an attack.
- If traffic/activity from 'out of scope' environments is logged by 'in scope environments', this causes additional detection and response work for the MDR SOC that your business is not licensed for.

However, Rapid7 MDR will still support only a subset of their logically separated environments if required by your team. When determining if one or more environments are 'logically separated' and therefore can be considered in-scope, consider the following criteria:

- Do the environment(s) have their own authentication and access control infrastructure? Specifically, do they have their own Windows domain?
- Are the environment(s) on a network that are logically segmented from the in-scope environment(s)?
- Does the environment(s) serve a distinctly different purpose than other environments? For example, a production data center (versus a corporate IT end user network).
- Do the environment(s) have their own Internet egress points?

If the in-scope environment(s) meets the criteria above **and** your team accepts the risk of a limited deployment of the Insight Agent in their in-scope environment, we can consider this a partial deployment and exclude other 'out of scope' environment(s) from MDR licensing and deployment.

Detection Methodologies

Below are the detection methodologies employed by the Rapid7 MDR to detect anomalous and malicious activity:

Detection Method	Description
User Behavior Analytics (UBA)	InsightIDR creates a baseline of normal user activity within your environment and generates alerts when there is a deviation.
Attacker Behavior Analytics (ABA)	InsightIDR applies behavioral analytics to generate alerts, built from our experience and understanding of attacker tools, tactics, procedures, and methodologies.
Network Traffic Analysis (NTA)	InsightIDR detects intrusions or other potential security events on your network through traffic analysis.
Threat Intelligence Detections (Intel)	Proprietary threat intelligence indicators derived from research, previous investigations, MDR monitoring findings, and third-party sources.
Threat Hunting	The MDR team performs monthly forensic analysis based on Insight Agent data and other log sources to identify unknown threats in your environment based on emerging trends in the threat landscape. Threat hunting requires deployment of the Insight Agent to at least 80% of your logically separated environment.

Event Sources

InsightIDR supports a wide range of security-relevant event sources. These event sources are leveraged by the MDR SOC as described in the table below. Specifically:

- **Real-time detection**: These sources are processed by our threat detection engine and may generate real-time alerts that are reviewed by our 24x7x365 SOC.
- **Threat Hunting**: Data from these sources are aggregated and leveraged by analysts when performing monthly threat hunts.
- **Investigation**: Data from these sources may be leveraged to accurately attribute other activity to an asset or user, and to provide other useful context data in the course of investigating alerts or performing Remote Incident Response.

Source	Real-time Detection				Threat Hunting	Investigation	
	UBA	ABA	NTA	Intel		Asset/User Attribution	Log Search
Insight Agents	✓	✓		✓	✓	✓	✓
Active Directory	✓	✓		✓	✓		✓
VPN Logs	✓	✓		✓	✓		✓
Cloud Services Logs	✓	✓		✓	✓		✓
Deception Technology	✓						✓
DNS Logs		✓		✓	✓		✓
Firewall Logs		✓		✓	✓		✓
Web Proxy Logs		✓		✓	✓		✓
Insight Network Sensor			✓	✓			✓
DHCP						✓	✓
LDAP						✓	
All Other Log Types							✓

See the full list of event sources supported by InsightIDR: insightidr.help.rapid7.com/docs/insightidr-event-sources.

Data Retention Policy

MDR offers your organization unlimited data ingestion into InsightIDR with access to 12 months hot storage and one month cold storage. Your team can add or remove event sources with no incremental data charges, with exception to Insight Agents which must be licensed and deployed to as much of your entire environment as possible. Exporting data is possible, but must be set up on a separate S3 bucket instance managed by your team and limited to a go-forward basis.

Deployment and Configuration Tasks

We encourage your team to begin the deployment as soon as possible by self-deploying InsightIDR in your environment. If deployment assistance is requested, Rapid7 will assign a Product Consultant to work with your team to ensure a successful and timely deployment. To expedite this deployment process, we ask that your team also elect a Project Manager to expedite the process.

Deployment Process

The Rapid7 InsightIDR product can be deployed with or without the assistance of a Rapid7 Product Consultant. After purchasing Rapid7 MDR, you will be sent a welcome email that explains your options for both self-deployment and scheduling time with a Rapid7 Product Consultant ("Deployment Sessions").

While you have the option to complete deployment on your own and go directly into service, an enablement session with a Product Consultant is still highly encouraged.

If you opt to work with a Product Consultant, Rapid7 recommends that you begin deployment on your own and then work with a Product Consultant to complete any outstanding items, answer questions, and get enablement on the product.

Deployment Sessions

Upon completion of the pre-deployment tasks, your team can leverage remote deployment assistance with a Rapid7 Product Services Consultant. This assistance should not to exceed 3 sessions, but can be reviewed on a case-by-case basis for complex environments.

During the Deployment Session(s), your assigned Rapid7 Product Consultant will assist your team with any remaining InsightIDR deployment related tasks, including the configuration of Collectors, event sources and product settings. Rapid7 will not assist the customer with agent roll-outs. Custom integrations, additional deployment time, training, and other services are not included in the Deployment Sessions and must be purchased separately.

Activating your MDR Service

In order to activate your MDR service, you must fill out a Service Request Form to provide Rapid7 with information regarding your environment. You will be provided with a link to this form in your welcome email. Once you have completed this form and had a kickoff meeting with your Customer Advisor, your service will begin.

InsightIDR Access

Rapid7 MDR has the right to access your InsightIDR instance as necessary to deliver service/support. A list of users who have access to Rapid7 InsightIDR can be seen inside the product. Your team is required to add users for your organization to InsightIDR; Rapid7 will not add users for your organization to InsightIDR once the Deployment Phase is completed. In the case that all of your organization's existing Insight Platform Admins are no longer with the organization, someone from your organization must provide Rapid7 a written request for access.

Compromise Assessment

Once your team has deployed the Insight Agent to 80% or more of endpoints in your in-scope environment, a Compromise Assessment will be performed.

Deliverable	Frequency	Description
Compromise Assessment	One-time	After deployment, Rapid7 MDR will evaluate whether there is malicious activity in your network or evidence of previous compromise(s). This report contains any detected active or historic compromises, potential avenues for future breaches, and prioritized remediation and mitigation recommendations.

If the Compromise Assessment finds that there is currently a compromise or detected malicious activity, Rapid7 will suggest that you utilize one of your Remote Incident Response engagements.

Validation of Alerts

When a threat is detected, your assigned Pod of analysts will act as an extension of your team, manually validating each detection by gathering context from your endpoints and logs to assess the severity. Validation is defined as

the Rapid7 MDR SOC performing initial triage and investigation to determine with a high degree of confidence that the event is non-benign and requires a communication to your security team.

Threat Findings and Reporting

Rapid7 MDR service reports are delivered via the Rapid7 secure file transfer system located in the Rapid7 Services Portal. These include:

Deliverable	Frequency	Description
Findings Report	Ad-hoc, after validated attacker activity	Provides written analysis ("attack storyboard"), criticality, raw details, remediation and mitigation recommendations, and suggested containment actions at the conclusion of each validated incident investigation. Rapid7 will notify the customer of any malicious activity ("incident") discovered via your preferred method within the timeframes outlined in the 'Response Times'.
Hunt Reports	Monthly	Once your compromise assessment is complete, you will begin receiving monthly hunt reports. These reports provide metrics and findings related to endpoint forensic analysis activities performed by the MDR analysts. Our analysts leverage the Rapid7 Insight Agent to collect metadata from multiple locations on the customer's endpoints to identify persistent malware, historical application execution, unusual processes and network communications, and per-system anomalies.
Monthly Service Reports	Monthly	Provides metrics and context surrounding analysis activities, technology health, and findings summaries for an at-a-glance overview of MDR activities.
Quarterly Service Reports	Quarterly	Provides metrics and context surrounding analysis activities, technology health, and findings summaries for an at-a-glance overview of MDR activities for the customer.
Threat Intelligence Reports	Ad-hoc, MDR finds new attack patterns	Provides a highly targeted analysis of new and emerging threats to inform the customer of the findings based on Rapid7's Threat Intelligence infrastructure or third-party threat intelligence partners.

Response Actions

You have the option to enable the Active Response service capability. Active Response gives your security program immediate response capabilities—initiated by our MDR experts—to stop attacks and contain confirmed threats in your environment. Details on the 'Active Response' service capability for MDR Elite customers can be provided in an additional Addendum titled "Scope of Service Addendum - Active Response".

Customer Advisor Engagement

During the course of your MDR service, your team will engage with your assigned CA. This resource is available to answer any questions about the MDR service and offer security advisorship to advance your security maturity.

Customer Advisors are available during normal business hours by phone and email. During non-business hours, a member of the CA team is on-call via the CA Hotline if malicious activity is detected in your environment.

Outlined below are frequent interaction touchpoints that your team will have with your CA:

Communication	Frequency	Method	Description
Monthly Meeting	Scheduled Monthly	Online, Phone, or Screen Share	• Review monthly hunt reports • Address questions about alerts • Walk through threat intel reports • Build custom alerts or other use cases • Answer questions related to the current program

Request For Information (RFI)	Ad-Hoc, when SOC analysts need additional details	Email from CA to Customer	When an incident is identified and the MDR SOC needs additional context (input from you), your CA will reach out via email to ask for more information to assist the investigation. (Ex. user running an abnormal process that we must confirm is related to malicious activity or intentional).
Findings Report Communication	Ad-Hoc, with Findings Report	Phone and/or Customer Portal	CA notifies you of a validated incident and presents remediation recommendations.
Quarterly Service Review	Quarterly	Online, Phone, or Screen Share	CA will walk through a summary of the service for the quarter and present customer recommendations for how to further advance your security maturity.
Customer Requested Meeting	Ad-Hoc, requested through Online Support Portal	Online, Phone, or Screen Share	You request a meeting to address concerns or questions regarding the service, technology, or alerts with your CA -- both for Rapid7 MDR or outside of Rapid7 advice.
Customer Questions	Unlimited, Ad-Hoc	Online Support Portal	You can leverage the online Support Portal to request help or voice concerns/questions.

Remote Incident Response

Remote Incident Response ("Remote IR") engagement is a technical response process handled remotely by the Managed Services SOC team. The customer is allotted two (2) Remote IR engagements per contract year.

Customers will be able to invoke a Remote IR for any incident discovered -- by the customer or the Managed Services SOC -- in the in-scope environment at any time after contract execution. An 'incident' is defined as 'a confirmed or reasonably suspected compromise of customer systems or data'. An 'in-scope environment' is one in which the customer has either deployed or is licensed to deploy the MDR service. This would include all systems for which an Insight Agent license has been purchased, and all common cloud services used by customer users.

In the event of a validated security incident, the customer will have the option to exercise a Remote IR engagement per the service level objectives outlined below:

Activity	Definition
Remote Technical Analysis & Incident Scoping	Analysis of any data source including data generated by the Insight Agent, and other analysis techniques to include full disk forensics.
Communications & Updates	Daily verbal debrief of the day's investigation results and progress. Substantive findings (such as increase in incident scope or impact) will be communicated regularly as discovered. Written weekly Summary Status Reports will be produced if engagement exceeds a week.
Remote Incident Response Report	This report will provide an overview of the Incident and a retrospective to include an executive summary, findings details, analysis, root cause, and recommendations, within 10 business days from completion of investigation

Additionally, you may use a Remote Incident Response engagement to test your MDR security control during a Penetration Testing engagement (through Rapid7 or otherwise). Details are available in the Appendix section titled, "MDR Reporting During Penetration Testing".

Service Level Objectives

Response Times

The following response times are included as a part of the Rapid7 Managed Detection and Response service:

Investigation Validation Response Time

Based on the level of severity of an incident, the MDR team will alert you per the response times outlined in the table below. It should be noted, criticality of an event is determined by the Rapid7 MDR SOC during the course of an investigation into an identified event. It is not possible to assign criticality before the scope of the event is determined and the incident is validated.

Severity	Example behaviors	Target time to notification	Time to Findings Report
Critical	An incident created via non-commodity malware deployed via spearphishing, social engineering, zero-day exploitation, or strategic web compromise, specifically targeted towards a target or organization.	Within one (1) hour of validation; Ongoing communications as they become available, but at a minimum, every 4 hours. Significant findings will be communicated as they are identified.	Findings Report will be posted in the Services Portal within 24 hours upon completion of investigation
High	An incident created using targeted off-the-shelf software backdoor deployed via spearphishing, social engineering, or strategic web compromises. Planned and targeted, but using common malware.	Within one (1) hour of validation; Ongoing communications as they become available, but at a minimum, every 4 hours. Significant findings will be communicated as they are identified.	Findings Report will be posted in the Services Portal within 24 hours upon completion of investigation
Medium	An incident created using common threat malware, typically non specifically targeted, but rather opportunistic and automatic.	Within eight (8) hours of validation; Ongoing communications as they become available. Significant findings will be communicated as they are identified.	Findings Report will be posted in the Services Portal within 24 hours upon completion of investigation
Low	An low-risk threat, not capable of remote code execution, credential harvesting, or data theft. (ex: Spam email delivering adware).	Within eight (8) hours of validation; Ongoing communications as they become available. Significant findings will occur as they are identified.	Findings Report will be posted in the Services Portal within 24 hours upon completion of investigation

Customer Advisor Response Times

The Customer Advisor team is held to the following response times for notifications of validated threats:

Trigger	Time to Action	Method	Action
Critical Severity Threat	Up to 1 hour after posting the RFI	Phone	Proactively reach out to the customer for validated critical severity threats by phone to provide relevant details while the SOC generates a Findings report.
High Severity Threat	Up to 1 hour after posting the RFI	Phone	Proactively reach out to the customer for validated high severity threats by phone to provide relevant details while the SOC generates a Findings report.
Medium Severity Threat	Up to 1 hour after posting the RFI	Email	Proactively reach out to the customer for validated medium severity threats by email to provide relevant details while the SOC generates a Findings report.
Low Severity Threat	Up to 1 hour after posting the RFI	Email	Proactively reach out to the customer for validated low severity threats by email to provide relevant details while the SOC generates a Findings report.

The following are response times to inquiries from your team:

Response Trigger	Time to Action	Action
Urgent Request	2 business hours	Reactively respond to an urgent request from the customer's team. Urgency is based on the discretion of the Customer Advisor team.
Non-urgent Request	24 business hours	Reactively respond to a non-urgent request from the customer's team. Urgency is based on the discretion of the Customer Advisor team.

Remote Incident Response

Each Remote IR is governed to the response times outlined below:

Action	Time to Action
Remote IR Trigger	Remote IR will begin as requested by the customer, or once any investigation performed by the MDR SOC exceeds 8 hours.
Time to begin Remote IR	1 hour from customer request/approval to initiate Remote IR.

Technology Uptime

Rapid7 InsightIDR and the Insight Platform, which powers the MDR service, follows the same uptime availability reflected by Rapid7's overall Insight Platform Service Level Agreement.

Additional Terms

This Scope of Services contract is governed by Rapid7's standard Master Services Agreement available at https://www.rapid7.com/legal/terms/ unless the parties have a fully executed Master Services Agreement which supersedes such standard terms. Any changes in materials or scope of work as defined in this document must be agreed upon in writing by the customer and Rapid7. Customer deployed software and related services are governed by the Rapid7 Terms of Service available at https://www.rapid7.com/legal/terms.

Signature

The undersigned has read, understands, and agrees to the Rapid7 MDR Scope of Service (including Appendix).

Company _____ Title _____

Signature _____ Name _____

Date _____

APPENDIX

MDR Responsibilities Matrix

	Rapid7	Main POC	Security & IT	C-Suite
Initiation Phase				
Complete Site Survey		✓	✓	
Define internal remediation escalation path(s) for MDR reporting		✓		
Set up customer in InsightIDR	✓			
Enable customer in customer's services portal	✓			
Deployment Phase				
Deployment Intro call	✓	✓		
Download and install Collectors		✓		
Deploy Insight Agent to all servers and workstations		✓	✓	
Deploy Insight Network Sensor(s)		✓	✓	
Install Orchestrator and workflows		✓	✓	
Configure event sources				
Configure all required event sources	✓	✓		
Optional - Configure recommended event sources (if available); Ex. Firewall, VPN, Web Proxy	✓	✓		
Firewall rules complete		✓		
Deploy Deception Technologies				
Deploy honeypots		✓	✓	
Deploy honeycredentials		✓	✓	
Deploy honeyfiles		✓	✓	
InsightIDR walkthrough	✓	✓		

	Rapid7	Main POC	Security & IT	C-Suite
Service Delivery Phase				
Service Delivery Kickoff call	✓	✓		
Compromise Assessment *(If in Full Service)*	✓			
Customer Advisor communication process				
Monthly Meeting	✓			
Quarterly Meeting	✓			
Availability for Board and Executive calls (optional)	✓	✓		✓
Ad-Hoc Calls	✓	✓	✓	

	Rapid7	Main POC	Security & IT	C-Suite
Service Delivery Phase (Continued)				
24x7 environment monitoring				
Investigate MDR ABA, UBA, & NTA alerts	✓			
Validate investigated alerts to remove false positives (including RFI confirmation from the customer when needed)	✓			
Validate alerts and remove false positives from Customer's custom alerts and 3rd party alerts		✓	✓	
Detail malware and/or malicious activity analysis	✓			
Attack storyboarding	✓			
Write and assemble Findings Reports detailing any verified suspicious or malicious activity	✓			
Outreach to Customer ensure Findings Report and all recommendations are understood	✓	✓		
Initial Containment (one or the other)				
Without Active Response set up		✓		
With Active Response set up	✓			
Additional remediation & mitigation actions		✓	✓	
Monthly Threat Hunting *(If in Full Service)*				
Analyze historical data	✓			
Threat hunting	✓			
Threat reporting	✓			
Remediation & mitigation actions performed		✓	✓	
Threat Intelligence				
Monitor global attacks and vulnerabilities	✓			
Share research and findings in Threat Intel reports	✓			
Add Threat Intel findings to monitoring detections	✓			
Remediation & mitigation actions performed		✓	✓	
Remote Incident Response (Remote IR)				
Suggestion for Remote IR if a breach is confirmed	✓	✓	✓	
Acceptance of Remote IR		✓		
Scoping of breach	✓			
Reporting of findings	✓			
Presentation of findings	✓			
Remediation & mitigation actions performed		✓	✓	

Technology-Dependent Service Limitations

Some aspects of the Rapid7 MDR service may be degraded as described below if technology deployment or coverage requirements are not fully met.

Service Limitations of a Partially Deployed Environment

Rapid7 MDR recommends full deployment of Insight Agents to all in-scope assets. However, for a partial deployment of the Insight Agent to the environment your organization understands, agrees, and accepts the limitations and risk of service degradation. Specifically, the following aspects of the MDR service are unavailable to assets without the Insight Agent installed:

Detection Aspect	Limitation
Attacker Behavior Analytics	A significant portion of MDR's threat detection power lies in the ability to detect specific events (file system changes, network connections, process start/stop) on each of the assets. This data can only be provided by the Insight Agent.
Manual Human Threat Hunting	The MDR monthly threat hunts rely on the endpoint agent to collect the data in scope for threat hunts. Assets without the Insight Agent will be excluded from threat hunts. Threat hunting requires deployment of the Insight Agent to at least 80% of the in-scope environment.
Threat Intelligence matching	All executable processes run on any asset with the Insight Agent are matched against known threat intelligence. Assets without the Insight Agent will not have running processes matched against threat intelligence.
Alert validation and Remote IR investigations	MDR's incident investigations rely on the Insight Agent to collect data for analysis. Assets without the Insight Agent will be out of scope for both the typical validation process conducted by the SOC team for an alert as well as any Remote IR investigation.
Local authentications and group membership changes	The Insight Agent is required to identify authentications using local accounts, such as a local administrator account, and is required to identify local group membership changes (ex: user added to local administrators group). Assets without the Insight Agent will be excluded from local authentication and UBA, where UBA is the act of tracking per-user and per-system actions to build statistical models of user activity and identify anomalies.
Attacker ingress detection	The most common methods of compromise are via Phishing (malicious emails) and malicious web sites, both of which require end-user interaction to succeed. As the majority of internet browsing and email activity occurs on end-user workstations, Rapid7 is unable to identify initial methods of compromise and lateral movement from those systems to servers and other critical assets without the Insight Agent.

Alert Review Limitations

Rapid7 MDR reviews real-time alerts based on event sources described above. The following limitations are part of the MDR service related to event sources configured in InsightIDR:

Alert Source	Limitation
Third party alerts	These alerts (listed at https://docs.rapid7.com/insightidr/third-party-alerts) are not reviewed by the MDR SOC and do not generate alerts for the MDR SOC team to investigate. These log sources are used only for informational purposes to add fidelity and evidence during an investigation performed by the MDR SOC analysts.
Custom InsightIDR alerts	Alerts you add are not reviewed by the MDR SOC. It is up to you to review these alerts.

Linux & MacOS Limitations

Rapid7's MDR service on endpoints (including workstations and servers) is delivered via the Insight Agent on devices running the Windows operating system. While the Insight Agent can be installed on Linux and Mac endpoints, these agents do not currently have full parity with the Windows Insight Agent. Therefore, your team understands and accepts that the limited capabilities of the Insight Agent on Linux and Mac assets present a risk of coverage for those devices, including the following:

Service	Limitation
Compromise Assessment	A Compromise Assessment is performed to determine whether there is malicious activity in your network or evidence of previous compromise(s). The data leveraged for the Compromise Assessment is not collected from Mac and Linux endpoints, and as a result these endpoints are not included in the Compromise Assessment. *NOTE: If all systems in a customer environment are Mac/Linux systems, a Compromise Assessment report will **not** be provided.*
Threat Hunting	The MDR team performs monthly forensic analysis from Insight Agent data to identify unknown threats in the customer's environment based on emerging trends in the threat landscape. This data is not collected from Mac and Linux endpoints, and as a result these endpoints are not included in monthly threat hunts. *NOTE: If all systems in a customer environment are Mac/Linux systems, Monthly Threat Hunt reports will **not** be provided.*

MDR Reporting During Penetration Testing

Many MDR customers integrate routine penetration testing into their holistic security strategy to perform end-to-end evaluations of all cyber security measures, including prevention, detection, and response. As an integral aspect of your cyber security controls, MDR works on your behalf to identify malicious activity, provide initial containment actions (if Active Response is enabled), and guide your team with detailed remediation and mitigation steps.

Penetration Test Communications to Rapid7 MDR

Communication Prior to Penetration Test

Rapid7 encourages your team to communicate upcoming or ongoing penetration testing to your CA. If you do choose to communicate details of upcoming penetration tests to your CA, you agree to provide:

- Beginning and end dates of penetration test
- Systems, networks, or subnets in scope
- Whether the test is a 'blackbox or no knowledge', 'limited knowledge', or 'full knowledge' test

You may also opt to not inform your CA as part of your testing regimen.

Communication During a Penetration Test

During the course of your Penetration Test, your CA will communicate Findings Reports or Requests for Information generated by the MDR SOC with a request to clarify whether MDR findings are related to penetration test activities.

Upon confirmation of penetration tester activity, MDR analysts will continue to monitor for related testing activities and use additional context provided by your team (or determined through analysis), to differentiate between penetration tester activity and potential attacker actions. As such, it is required that your team acknowledge if there is pen testing activity happening in the in-scope environment if MDR detects the test.

Once confirmed, you will be asked which engagement model you prefer ('no further notifications', 'rollup' or 'purple team exercise, see below). The default for Rapid7 MDR reporting for Penetration Testing is 'no further notifications'

which means your team will not be alerted to further security testing activity.

It is required that your team notifies your CA at the conclusion of the penetration test; following the completion of scheduled penetration testing activities and review of the MDR deliverables generated by the test, your CA will work with your team to identify detection gaps, if any, and additional measures to assist in adding context to findings.

Service Delivery Considerations

Should the scope of attacker activity increase beyond MDR's initial findings (or you communicate to your CA that a Pen Test is taking place), we will ask you to choose one of the following options:

Option 1: Roll Up Reporting (Preferred)

Once Rapid7 identifies a Pen Tester in your environment, Rapid7 will only provide an initial notification on events related to security testing. Following initial notification and confirmation that findings are associated with ongoing penetration testing activities, the MDR team will continue to collect related activities and deliver our findings as an aggregate 'roll-up' report that demonstrates which alerts the SOC would investigate had this been a real attack. You are required to specify "Roll Up" reporting as the default is 'no further notification' to further security testing activity.

Roll Up reporting contains the following deliverables:

- An aggregate 'roll-up' report of all alerts generated by the penetration test activity.
- Optional: A penetration test 'post mortem' report upon request (see below)

Option 2: "Purple Team Exercise"

You may want our SOC to treat the penetration test as an actual attack with all the resulting deliverables and updates. This type of engagement is considered a 'purple team exercise'. Rapid7 can support this 'purple team exercise' upon request; however, the 'purple team exercise' is subject to the following requirements listed below.

Note: Rapid7 MDR reserves the right to refuse the 'Purple Team Exercise' in rare instances. It is best practice to notify your CA in advance if your team wishes to engage in a Purple Team Exercise in order to accommodate requests.

Customers are required to:

1. Timeline for the Penetration Test must be **limited to under seven (7) days** from the confirmation of the Pen Testing activity to the end of the engagement. Rapid7 holds the right to not continue providing real-time reporting relating to the Pen Test after these 7 days.
2. Your organization agrees to **use one (1) of their allotted Remote IR engagements** due to the amount of SOC resources required to respond to penetration test activity.
3. Your team agrees to **take all recommended containment and/or remediation actions in a timely manner** to simulate an active response to these attacks. If Active Response is enabled, Rapid7 will perform actions to contain the threat as it relates to the Active Response service. Performing response actions in real-time will result in a more accurate simulation of an attack and subsequent response. Additionally, this prevents the MDR SOC against 'unbounded' engagements that impact Rapid7's service delivery to all customers.

Purple Team Exercise reporting contains the following deliverables:

- All Remote Incident Response deliverables.
- Optional: A penetration test 'post mortem' report upon request (see below).

Penetration Testing Post-Mortem

Upon request, Rapid7 can also generate a penetration test 'After Action Report' to address any specific detection and response gaps (if applicable) and planned/recommended remediation steps. In order for Rapid7 to generate this report, your team must provide Rapid7 with sufficient detail about the activity performed during the penetration test (dates, systems, usernames), ideally in the form of the final penetration test report.