DARKTRACE

Darktrace Offering Product Specification

020 7930 1350

Maurice Wilkes Building, Cowley Rd, Milton, Cambridge CB4 0DS

darktrace.com

Preface

This document is the Darktrace Offering Product Specification for across the Darktrace Product range and should be read in conjunction with the Darktrace Master Service Agreement and the applicable Product Order Form. In the event of a lacuna for a given Product, the relevant technical documentation as hosted in the Customer Portal applies. Service Definitions for Darktrace Additional Services may be found at https://darktrace.com/legal/product-specifications-and-service-definitions.

Darktrace / NETWORK

Table of Contents

1. Introduction

Darktrace / NETWORK is a market leading cyber security solution that uses AI technology to build a dynamic understanding of customers' organizations. It models and evaluates network activity to gain visibility of threats in real-time and to respond autonomously and at machine speed.

Darktrace / NETWORK is comprised of two elements: Real-Time Detection and Autonomous Response. This Product Specification outlines the anticipated operation, architecture, implementation, and administration for each element.

This document it to be read in conjunction with the Darktrace Master Service Agreement which governs the usage of the Darktrace Product Offering. References to "Customer" throughout this document should be read to refer to the entity that is the owner of the software subscription and is ultimately responsible for its operation, whether as end-user or service provider.

Darktrace / NETWORK Subscriptions are sold according to the Usage Metrics bands set out in the relevant Product Order Form.

2. Real-Time Detection

Darktrace / NETWORK Real-Time Detection was previously named DETECT/Network and Enterprise Immune System. Darktrace / NETWORK that is implemented into cloud-hosted environments may be referred to in exceptional circumstances as Darktrace / CLOUD (Network). This specification applies for Customers who have purchased the product under any of the above naming conventions.

2.1. Product Features

2.1.1.High Level Summary

Darktrace / NETWORK Real-Time Detection evaluates network activity of devices for behavior outside of a “normal” state, created by an ongoing analysis of connections and other supplied data events across the enterprise network environment. The system profiles individual devices based upon their activity, and the activity of those it deems “peers” due to similar behavioral activity (“modelling”). Alerts are raised when specific criteria, or a minimum threshold of unusual behavior is met. The output is surfaced to operators in the Darktrace Threat Visualizer interface for investigation and resolution.

2.1.2.Data Retrieval and Ingestion

Darktrace / NETWORK Real-Time Detection retrieves and processes raw network traffic. Network traffic may be delivered directly to a physical Darktrace appliance ingestion port, to a virtualized Darktrace vSensor instance, or forwarded from Darktrace host-based osSensor agents to a Darktrace vSensor.

To ensure full visibility, all network traffic should be mirrored or ingested by Darktrace in some format. It is the responsibility of the Customer to ensure that Darktrace maintains visibility over network traffic in the event of network re-architecture outside the scope of the initial deployment design. The network traffic provided to Darktrace must be of suitable quality and not contain duplication, fragmentary data, or be incomplete in any fashion (for example, unidirectional).

Darktrace / NETWORK Real-Time Detection can also receive and parse log data in syslog from external tools (this requires additional configuration). Master, Unified View, physical Probe and vSensor instances (as defined at Section 2) support syslog ingestion.

Darktrace offers a suite of Threat Intelligence and Telemetry Integrations where data may be retrieved from other compatible security tools or OSINT sources. The retrieval method for each integration is detailed in the corresponding integration documentation. These alternative inputs (“Threat Intelligence Integrations” & “Telemetry Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page.

Darktrace also provides a REST API for automation and a subset of relevant data input.

2.1.3.Darktrace Analysis

Darktrace analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace platform, Threat Visualizer interface, or any underlying components.

Where source data is low quality, represents only a subset of all network activity, is delayed by external factors, or otherwise is incomplete, the quality of Darktrace analysis will be severely impacted. It is therefore imperative that data ingestion is configured comprehensively and to a high standard by the Customer. Analysis may also be impacted by overloading, such as that observed when traffic throughput significantly increases beyond scoped levels.

2.1.3.1. Darktrace Deep Packet Inspection

Darktrace Deep Packet Inspection is performed on raw ingested network traffic to produce metadata for analysis by other platform components. The list of supported L4 - L7 protocols for inspection is regularly expanded, and analysis is continuously refined.

The output of Deep Packet Inspection is passed to the Darktrace model engine, to the core Darktrace "classifier" engine, and made available for display in the user interface and any other platform components such as the AI Analyst.

2.1.3.2. Darktrace 'Pattern of Life' Analysis and Classification

Network events, user activity events, connection data and any other configured inputs are subject to Darktrace 'pattern of life' analysis. Darktrace will create individualized behavioral profiles for the network entities it observes and surface activity which is considered inconsistent with the expected norm. This baseline is derived from - but not limited to - an analysis of the behavior of the individual entity, analysis of one or more clusters created from similarly behaving entities, and many variable factors such as time or communication protocol. The 'pattern of life' data is continually updated in real time, and reflects the data that has been received, with a greater weighting to more recent data.

Analysis is performed in the "classifier" stage by a multitude of classifiers. This core analysis applies many approaches including unsupervised machine learning techniques such as Bayesian meta-classification, techniques derived from graph theory and network analysis such as node/graph centrality, approaches derived from statistical analysis such as spectral clustering and many other techniques. The previous examples of utilized techniques provided are for illustrative purposes and should not be considered exhaustive.

Darktrace does not offer any capability to access underlying behavioral models or classification output. Darktrace provides the ability to utilize the output of this analysis using the *models* framework. Darktrace will create alerts to indicate anomalous activity which will be inserted into the "event log" within the Darktrace Threat Visualizer interface of the corresponding device, user, or entity. These unusual activity "notices" are also consumed by Darktrace Real-Time Detection models.

2.1.3.3. Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of 'pattern of life' analysis is evaluated. Both analysis components described above submit data to the model engine for evaluation.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior. The models framework leverages both the underlying 'pattern of life' detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet any criteria; these are created and maintained in the Threat Visualizer "Model Editor" interface. Darktrace bears no liability for the operation or outcomes of custom models created by Customer.

Default Darktrace models are focused on 'pattern of life' anomaly detection, potentially malicious behavior and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically, clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

Models may also be used to trigger actions within the Darktrace Threat Visualizer platform; the output of the Darktrace Real-Time Detection model engine is described below.

2.1.3.4. AI Analyst

Darktrace AI Analyst performs a meta-analysis upon the previous layers of analysis described. Please refer to “AI Analyst” below.

2.1.3.5. Cross-Capability / Cross-Coverage Area Analysis

Where other Darktrace Real-Time Detection coverage areas are deployed, cross-platform analysis is performed by multiple components. Links are created for analysis purposes between network entities and entities modeled by other Real-Time Detection components such as, for example, user entities created by Darktrace / IDENTITY module monitoring known to be associated with a given network entity, or credential entries observed by Darktrace / NETWORK monitoring also observed as part of Darktrace / ENDPOINT coverage.

Darktrace AI Analyst will also link together entities across different coverage areas in the creation of AI Analyst Incidents and AI Analyst Incidents Events. Darktrace AI Analyst may also retrieve additional contextual data during investigation from other components if deployed, such as the retrieval of associated emails from Darktrace / EMAIL.

The output from Darktrace components may be fed to Darktrace Real-Time Detection as part of supplementary Threat Intelligence; Darktrace / Attack Surface Management can provide output of malicious asset identification into the “Watched Domains” list of Darktrace Real-Time Detection to create Model Breach alerts when a network entity accesses an Attack Surface Management-identified domain.

The above provides illustrative examples of collaborative scenarios between platform components for reference but is non-exhaustive.

2.1.3.6. Processed vs Modeled

Note, a distinction exists between activity which has been “processed” and that which has been “modeled”. Data may be processed and analyzed by Darktrace Deep Packet Inspection, but not further submitted to ‘pattern of life’ analysis or any other analysis components. In this scenario, no “device” entity is created as a result of the event ingestion, and activity performed by the entity is not evaluated for unusual activity or by Darktrace models, except in direct relation to modeled entities.

This distinction is only relevant where specific network ranges or devices have been intentionally excluded by an operator from further analysis. The responsibility for the exclusion of these devices from Darktrace analysis, and any reduction in analysis coverage as a result, is the responsibility of the Customer as operator.

2.1.4. Output

Activity processed is surfaced in the Darktrace Threat Visualizer interface, and is optionally available through the Darktrace Mobile App, compatible output formats, or the Darktrace Threat Visualizer REST API. Darktrace / NETWORK Real-Time Detection typically outputs connection metadata from Darktrace Deep Packet Inspection, results of ongoing ‘pattern of life’ analysis (including notification of unusual activity), model events, Darktrace AI Analyst alerts, contextual data retrieved from configured integrations (optional), and any other log-based event data created or processed by the Darktrace system.

The platform may also output alerts about overall system health and component health.

2.1.4.1. Darktrace Deep Packet Inspection Metadata

The metadata output of Darktrace Deep Packet Inspection is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Data retention for this output is on a rolling basis and is dependent upon a number of factors such as hardware capability, traffic makeup and other operational components which contribute entries to Advanced Search. Most deployments can expect around 30 days retention. Darktrace Deep Packet Inspection metadata can be exported to external tooling by supported methods for longer retention.

2.1.4.2. Alerts

Darktrace / NETWORK Real-Time Detection produces two primary alert types: Darktrace Model Breaches and Darktrace AI Analyst Incidents. Darktrace Model Breach alerts and Darktrace AI Analyst Incidents can be investigated in the Threat Visualizer interface, investigated in the Darktrace Mobile App, or exported to a compatible alert consumer. Alerts are categorized for priority and filterable by multiple factors, allowing for customization of alerts displayed in Darktrace interfaces and those exported to external tools.

A secondary alert type - Darktrace System Status alerts - may also be generated where the Darktrace platform is experiencing degraded service due to health issues, invalid input, or where it is necessary to highlight changes to system administrators. Alerts include details of the originating host, the severity of the event, and relevant links to investigate or resolve the issue. Notifications are sent when a system event becomes active and can optionally be sent on resolution.

Darktrace supports alert export in both industry-standard generic forms such as Syslog or Email and custom integrations with common tools such as Splunk and ServiceNow. Alert outputs (“Workflow Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page. The details included in each external output may vary due to third-party restrictions on content length or supported formats.

2.1.4.2.1. Model Breach Alerts

Darktrace Model Breach alerts are created as a result of Darktrace Real-Time Detection models; when conditions for a model are met, a model breach alert can be created in addition to other possible model actions. Darktrace Model Breach alerts contain details of the conditions that were satisfied in order to trigger the alert, the entity which met those conditions (for example, a device, or another Model Breach alert) and a description with recommended action points. Alerts will only be generated from ingested and inspected packets.

2.1.4.2.2. Darktrace AI Analyst Incident Alerts

Darktrace AI Analyst Incident alerts are created when Darktrace AI Analyst identifies activity considered significant enough to highlight to operators. Darktrace AI Analyst Incidents may contain a single “AI Analyst Incident Event”, or multiple linked findings aggregated together. External alerts are created when a new “AI Analyst Incident Event” is created, which may be associated with an existing “AI Analyst Incident”, or form an independent, new “AI Analyst Incident”. In the former case, mechanisms to identify the relationship between AI Analyst Incident Events are provided in the output.

In the Threat Visualizer and Darktrace Mobile App, the output of Darktrace AI Analyst analysis is aggregated on a per-AI Analyst Incident basis. New AI Analyst Incident Events added to existing Incidents will not produce independent alerts in these interfaces and are instead displayed as part of the existing Incident. Darktrace AI Analyst Incidents in these interfaces contains details of the activity identified, the Darktrace Model Breaches that triggered the initial AI Analyst investigation, the entities which performed the unusual activity, (for example, devices), the investigation steps AI Analyst performed, why activity was linked together (if multiple AI Analyst Incident Events), and a human-readable summary of the finding.

2.1.4.3. Reporting

Darktrace offers both manual and automatic PDF report generation, scheduled via the Darktrace Threat Visualizer “System Config” page. Details of reporting formats offered is provided on the Darktrace Customer Portal and relevant documentation.

Darktrace reserves the right to alter the content of reports offered to align with changing product and service offering.

2.1.5.Interface

The primary user interface of the Darktrace platform is the Threat Visualizer. The Threat Visualizer interface provides access to all major Darktrace user interfaces, consoles and product views - it contains both investigation interfaces and administration interfaces. A list of the user interfaces that comprise the Threat Visualizer - and detailed information on how to operate these interfaces - is provided in the relevant technical documentation.

Darktrace / NETWORK Real-Time Detection alerts are surfaced in the main Threat Visualizer dashboard which includes user customizable and filterable alert trays, a visualization of ongoing device activity over time, and access to the output of Darktrace analysis in both log and summary format. Alternative investigation views, access to review detailed metadata, report generation, and the ability to edit and create custom models (as described above) are also provided within the interface. Operators can also review and configure the devices and subnets observed by Darktrace / NETWORK Real-Time Detection, configure system settings, deploy integrations, review system health information, perform user management and other administrative tasks.

Physical Darktrace appliances and Darktrace vSensors maintain an administration console accessible only via SSH/22 (“Darktrace Console”). This console is used for advanced configuration tasks and is not available in Darktrace-hosted cloud environments.

2.1.6.AI Analyst

Darktrace models are used as a trigger to invoke AI Analyst. When the conditions for a model are met, a model breach is created; AI Analyst reviews and investigates all relevant model breaches that occur on the system as a starting point for its analysis process. The output from this analysis process is AI Analyst Incidents - a collection of one or more related events of anomalous activity. Incidents are formed through a meta-analysis of activity type, devices, and endpoints involved in each event. Each incident can encompass multiple stages of activity as it develops. The Darktrace AI Analyst operates a hypothesis-based analysis approach, where activity is evaluated against a number of possible, relevant hypotheses and a determination is taken of which (if any) hold based upon the evidence gathered and investigations performed. This investigation process involves numerous forms of data analysis including, but not limited to, AI and Machine Learning algorithmic approaches, statistical analysis techniques, and other forms of natural language and mathematical analysis.

The AI Analyst will combine activity across different Darktrace Real-Time Detection coverage areas where possible.

Although a model breach may be the trigger for an investigation, that does not mean the activity AI Analyst surfaces is directly related to the original model breach. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model breach but are worthy of investigation. Similarly, very few model breaches that are investigated will result in an incident - only activity the AI Analyst considers high priority. Whether AI Analyst has created a related incident is displayed in appropriate locations within the Darktrace Threat Visualizer and Darktrace Mobile App.

Please note, the Cyber AI Analyst will only investigate model breaches of default Darktrace models. Custom models are not currently supported. Users may manually trigger AI Analyst investigations into devices of interest, or trigger AI Analyst investigations through third-party telemetry inputs.

2.1.7.Reporting

The Threat Visualizer “Audit Log” records changes made by operators such as model alert acknowledgement; the audit log can be exported via Syslog for extended retention.

The metadata output of Darktrace Deep Packet Inspection is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Combined output from this metadata, from Darktrace analysis, and from any actions performed by the platform automatically (such as tagging as a result of a model) are combined into logs which are displayed for each device. Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, traffic makeup and other operational components which contribute data to the platform. Most deployments can expect around 30 days retention of Darktrace Deep Packet Inspection metadata and general device activity data. Automatic removal of older event log data does not affect the storage or training of the machine learning ‘pattern of life’ data.

Devices observed by Darktrace / NETWORK Real-Time Detection to be recently active in the network environment are detailed on the Threat Visualizer “Device Admin” page. Subnets observed by Darktrace / NETWORK Real-Time Detection in network connectivity are displayed on the Threat Visualizer “Subnet Admin” page. Metrics regarding data and event throughput are rendered on the Threat Visualizer “System Status” page.

Darktrace also provides a REST API for automated retrieval of a subset of data output.

2.2. Deployment Architecture

The appropriate combination of deployment components will vary depending on the customer network environment, for example, the balance of physical and virtualized infrastructure, or the physical location of covered networks. The appropriate deployment scenario may also change during the lifetime of the Darktrace deployment if significant changes are made to the network architecture. It is the responsibility of the client to ensure that Darktrace maintains visibility over network traffic in the event of network re-architecture outside the scope of the initial deployment design.

Each deployment requires a minimum of one “*Master*” instance to provide the capabilities outlined above. Deployments should be structured so that a single instance - Master or Unified View - sits at the top of the deployment topology to operate the Threat Visualizer interface and other relevant components.

Darktrace instances may operate in different roles within a distributed deployment; Darktrace vSensors can process raw network traffic ingested directly from virtualized networking equipment, or mirrored to it by connected osSensor agents, and forward the resultant metadata onward to a connected Darktrace instance. Darktrace instances may operate as *Probes* processing traffic and forwarding directly for analysis & storage on a *Master* or may store some processed traffic on-instance. Master instances may also directly ingest, process and analyze data.

The following outlines common Darktrace deployment architecture components. *Topology Components* are types of Darktrace hardware or software that may be deployed as part of a Darktrace deployment.

Topology Roles are particular purposes that instances may perform within the deployment that restrict their operation. There are five roles: “Master”, “Probe”, “Unified View”, “Subordinate Master” (or “SubMaster”), and “Agent”.

Some components are limited to specific roles only; others are able to operate in multiple roles.

2.2.1.“Physical Appliance” (Topology Component)

Darktrace offers a physical Darktrace appliance for installation in a datacenter or other compatible location. Darktrace appliances are highly tuned, high-performance pieces of hardware that host the Darktrace platform. There are multiple types of Darktrace appliance, with different throughput capacities and options for data ingestion. The appropriate model will depend upon the installation location, traffic volume, traffic composition and deployment role. Darktrace appliances can only run the Darktrace platform and no other software.

Customer has responsibility for managing software updates, backups, and other system health factors for Darktrace appliances. Master

2.2.2.“Cloud Instance” (Topology Component)

Darktrace provides virtualized Darktrace deployments by hosting a cloud-based instance within Darktrace cloud environments, hosted by AWS and Microsoft Azure. Virtualized deployments receive data from local Probes in the customer network. Cloud instances are managed and maintained by Darktrace; software updates, backups and instance scaling are managed by Darktrace operations.

Cloud instances can operate as “Master” and “Unified View” roles but cannot operate as Probes.

2.2.3.Master (Topology Role)

A Darktrace Master instance operates analysis, modeling and alerting capabilities and serves the Darktrace Threat Visualizer user interface. Darktrace Master instances can ingest and process network traffic directly, but this is not a necessity if this action performed by other instances in the topology. A Master instance may be a physical Darktrace appliance, or Darktrace Cloud Master, hosted by Darktrace in one of AWS or Azure in an appropriate availability zone. In the latter case (Cloud-hosted Master), Darktrace does not support direct traffic ingestion, and one or more Probe instance(s) will be necessary for network traffic retrieval.

2.2.4.Probe (Topology Role)

A Darktrace Probe receives data input such as network traffic (typically from SPAN, traffic mirroring or network TAPs) or syslog. Probes do not operate a user interface, analyze ‘pattern of life’ or have a model engine. A Probe may be a physical Darktrace appliance configured to operate in a Probe role, or a virtualized Darktrace vSensor Probe hosted by the client. Probes require a connected Master and will communicate in a method defined by the client during configuration; selection of an appropriate communication mode for the deployment scenario is the responsibility of the client.

2.2.4.1. Physical Probe (Topology Component Operating in Probe Role)

A physical Probe is a hardware appliance as described above, configured to operate as a Probe.

2.2.4.2. Virtual Probe (Topology Component with Exclusive Probe Role)

The Darktrace vSensor is a virtual Probe intended for deployment in cloud-based networks or environments where it is not feasible (such as virtualized networks), or not desired by the customer, to deploy a physical Probe. vSensor software is provided by Darktrace in various pre-packaged virtual machine formats or can be installed on virtual machine instances running a compatible Linux-based operating system. The vSensor supports autoscaling in compatible environments.

The vSensor receives data input such as network traffic or syslog but does not operate a model engine or serve a user interface other than a simple configuration console. vSensors require a connected Master to which they forward processed metadata and will communicate in a method defined by the client during configuration; selection of an appropriate communication mode for the deployment scenario, and definition of a secure access password (where relevant), is the responsibility of the client.

For usage in cloud or other environments where it is not possible to span a virtual switch, the vSensor also supports ingestion of traffic from multiple osSensors.

2.2.5.Darktrace Agents

Darktrace Agents are software components deployed on customer servers to collect traffic and enrichment data, which is forwarded to associated Probes or Masters; they do not perform any data processing themselves.

2.2.5.1. osSensor Agent (Topology Component with Exclusive Agent Role)

Darktrace osSensors are lightweight, host-based server agents. osSensors are provided for a range of compatible Windows and Linux-based operating systems. osSensors do not perform analysis or server a user interface. The osSensor forwards an unprocessed copy of network traffic to a connected vSensor for processing. Connection between the vSensor and osSensor is secured using a shared HMAC token. osSensors are not suitable for connectivity over untrusted networks.

The osSensor is also offered in a containerized format via Docker Hub for deployment in environments such as Kubernetes and AWS Fargate.

2.2.5.2. containerSensor Agent (Topology Component with Exclusive Agent Role)

The Darktrace containerSensor is a tracking sensor which enhances 'pattern of life' analysis in containerized environments. The containerSensor relays metadata about workloads to the connected Master, allowing traffic observed within the cluster to be assigned to consistent, recognizable entities.

2.2.5.3. Server Agents (Topology Component with Exclusive Agent Role)

Darktrace Server Agents collect enrichment data from customer servers which is not accessible from network traffic (typically remote access solutions), then relay this to the connected Master to improve user & device tracking. Further information can be found on the Customer Portal

2.2.6. “Unified View” (Topology Role) and “Subordinate Master” (Topology Role)

A Darktrace “Unified View” (UV) is a component of a distributed Darktrace deployment, where a single Master instance with or without Probes is not sufficient, or not suitable, for the deployment scenario. In this mode, an additional instance is deployed above, and receives data from, all Masters for display in a single user interface. Darktrace Master instances under a UV operate as subordinate Masters (“subMasters”), where a subset of capabilities is removed and controlled/operated centrally at the Unified View level. In this mode, each Master continues to operate a model engine and perform ‘pattern of life analysis’, but select components are synchronized from the Unified View (such as Darktrace models and system configuration).

2.3. Implementation

Darktrace / NETWORK Real-Time Detection is implemented through deployment of one or more physical or virtual components, including at least one Master instance. Full appliance specifications, implementation guides and administration guides are all found on the Customer Portal.

2.3.1.Physical Instances

Darktrace appliances are installed within the customer network to ingest network data into one or more packet capture ports. This information is processed by the platform and the output is displayed in the Darktrace Threat Visualizer of the associated Master.

2.3.1.1. Admin Interfaces

All Darktrace instances have a network port for the admin interface, which must be connected for: access to the User interface; connection between a Probe and its Master (dependent on the configured *role*); the Call-Home function; and additional network services such as third-party integrations, monitoring, inbound enrichment data, or outbound alerting.

Darktrace physical appliances also have an out-of-band (OOB) management interface. Darktrace strongly recommends connecting the OOB to the customer network for additional hardware-layer management, troubleshooting and maintenance capabilities, particularly when the admin interface is unavailable.

2.3.1.2. Analysis Interfaces

Darktrace instances feature different analysis interfaces for ingestion of network traffic into the Deep Packet Inspection engine. The available interfaces on each type of physical or virtual instance are detailed in appliance specifications available on the Customer Portal.

During implementation, customer must ensure a port mirroring, SPAN, TAP or other data forwarding method is set up through one or more interfaces on any Darktrace instance with the Probe role, or any Master also intended to ingest network traffic.

2.3.1.3. Data checks for new instances

To confirm instances are receiving traffic and the user interface is accessible, post installation, data checks should be carried out. Details of the necessary checks for each instance are provided in the Customer Portal documentation.

2.3.2.Darktrace Virtualized Instances

Darktrace provides virtualized Darktrace deployments by hosting a cloud-based Master instance within Darktrace cloud environments (hosted on AWS and Microsoft Azure). Within the scope of Darktrace / NETWORK, Cloud Masters receive data from local Probes in the customer network (physical or virtualized). Cloud Masters may also receive traffic as part of other Darktrace products (for example, client sensors as part of Darktrace / ENDPOINT)

Cloud Master instances can analyze and ingest the same data as physical Master instances.

Organizational virtual network traffic can be sent via a Darktrace vSensor in one of two encrypted communication modes. vSensors can be deployed as a standalone virtual machine in a traffic-mirroring scenario, or with up to 255 osSensor agents (per vSensor). Darktrace osSensors can be installed on devices running Windows, supported Linux distributions and any Linux environment running the Docker engine.

In addition to processing and transmitting network traffic, vSensors can ingest and forward syslog-format logs to the Darktrace Coud Master.

Individual, separate instances are provisioned for each customer within the relevant cloud-provider region. Data ingested from Probes in the customer network is encrypted in transit and will not leave the region. In addition, two-factor authentication is enforced on all user accounts. Darktrace is ISO27001 certified, ensuring we maintain a high standard of information security.

2.3.2.1. Configuration and Management of Cloud Masters

Cloud Masters are managed and maintained by Darktrace operations; the management and system administration console is not available. Darktrace manages instance scaling in line with traffic-load. If there is a hybrid Unified View deployment containing both physical and cloud submasters, Darktrace operations manages the cloud Master, whereas the physical Masters are managed by the customer. If you wish to modify a setting on a cloud master or submaster or configure a process that would normally require console access, please contact Darktrace support who will assist you with this process.

Threat Visualizer software on a Cloud Master is automatically updated when a new version becomes available; where possible, updates will be applied outside standard business hours. If this is not possible, the update process will cause minimal disruption for Threat Visualizer users. If automatic updates are turned off by customer request, it is customer responsibility to notify a Darktrace via the Customer Portal when they are ready to upgrade the deployment and the Darktrace team will update the deployment during the designated window provided by customer.

Multiple short-term snapshot backups of Cloud Masters are taken on a rolling basis to ensure continuity in a disaster recovery scenario.

2.3.2.2. Deployment process for Darktrace vSensors

In the same way as for physical instances, described in sections 3.1.1 and 3.1.2, vSensors need to be configured by the customer with both an admin interface and an analysis interface.

vSensors can ingest and process traffic from physical networks, in addition to virtualized traffic, with additional configuration. The vSensor supports VXLAN and ERSPAN traffics type I and type II, as well as GRE with transparent ethernet bridging.

For usage in cloud or other environments where it is not possible to capture data from a virtual switch, the vSensor also supports ingestion of traffic from native mirroring technologies, such as AWS VPC Traffic Mirroring and GCP Packet Mirroring, or multiple connected osSensors. The osSensor agent is installed on each customer device where visibility is desired and monitors all of the network traffic to/from configured interfaces of that device; the monitored traffic is then sent to the vSensor for analysis. osSensors utilize host resources to forward traffic, so should only be installed where it is not possible to retrieve traffic through other means.

2.3.2.2.1. vSensor Deployment Methods

Darktrace vSensors may be deployed via Standalone Image, Cloud QuickStart or Cloud CLI. recommended by Darktrace. The deployment guides for each method are available in the Customer Portal. It is the customer's responsibility to review the guides and decide which option works best for their organization.

2.3.2.2.2. vSensor health checks

Health checks must be performed by the customer to confirm that the vSensor is running and has connectivity with the associated master and osSensors (if applicable). Key areas to test to verify health include: (i) that the vSensor virtual machine is running; (ii) verification of incoming packets; and checking for vSensor overload.

2.3.2.3. Deployment process for Darktrace osSensors

The Darktrace vSensor coordinates with the osSensors associated with it, ensuring traffic is captured only once when osSensor devices communicate to each other. Each osSensor registers with a vSensor using a shared HMAC token which should be supplied to both ends. It is recommended that the associated vSensor is configured in advance of osSensor setup, in order to ensure the necessary HMAC token and IP Address for the vSensor have been collected.

The osSensor is deployed as a package installed on the customer server. Each osSensor registers with a vSensor using a shared HMAC token which should be supplied to both ends. The associated vSensor should be configured in advance of osSensor setup, in order to ensure the necessary HMAC token and IP Address for the vSensor have been assigned.

The Darktrace vSensor coordinates with the osSensors associated with it, ensuring traffic is captured only once when osSensor devices communicate to each other. The installation guides for each available formats of osSensor can be found on the Customer Portal

2.3.2.3.1. osSensor Health Check

It is recommended that the Customer checks: (i) that the osSensor service is running and that there is connectivity with the associated vSensor; (ii) how many osSensors are running; and (iii) for vSensor overloading.

Further information on the testing commands can be found in Customer Portal. It is a Customer responsibility to run periodical health checks to ensure the osSensor service is still running as optimally as initially setup.

2.3.3.Architecture Configuration

After each component supporting Darktrace / NETWORK is deployed, each needs to be configured into the wider Darktrace architecture. Darktrace can provide support and advise, but configuration is a Customer responsibility to implement. For more detailed information and examples of Darktrace architectures, see the Customer Portal documentation.

2.4. Administration

General administration of the Darktrace deployment is a Customer responsibility and is performed in the Threat Visualizer interface.

The majority of configuration, including administration of system settings such as proxies, authentication configuration such as LDAP and SAML SSO, deployment of alert and threat intelligence integrations, and other administrative tasks are performed in the “System Config” page. User and group management, including the assignment of data visibility and permissions, is performed on the “Permissions Admin” page. Device administration such as the setting of labels, application of tags, and altering of priority is performed on the “Device Admin” page. Management of subnets observed by the Darktrace / NETWORK Real-Time Detection instance are administered on the “Subnet Admin” page. System health information and system alert resolution is performed on the “System Status” page.

For physical Darktrace appliances of any role, additional administrative tasks such as interface configuration, setting of host variables, and Call-Home configuration may be performed in the appliance console, accessible over SSH. For cloud hosted instances, these administrative actions are managed by Darktrace operations directly and no access to the console is available.

Darktrace physical and cloud instances are each seeded with random passwords and two-factor authentication secrets at build time. These initial secrets are stored by Darktrace. Credentials granting access to the Darktrace Threat Visualizer interface and, for physical instances, the SSH administration console (“Darktrace Console”) are provided via the Customer Portal; these passwords can be optionally changed by the client at any time. Further access administration is a Customer responsibility.

Darktrace vSensors only provide access for administration tasks via the management console. osSensor, containerSensor and other Server Agents do not provide administration interfaces and are instead managed by the customer with configuration files and/or dialogues during installation, or through the customer server they are hosted on.

Detailed information about individual administrative tasks is outlined in the relevant Customer Portal documentation for the deployment of the specified component or administrative action intended to be performed.

3. Autonomous Response

Darktrace / NETWORK Autonomous Response was previously named RESPOND/Network and Antigena. Darktrace / NETWORK that is implemented into cloud-hosted environments may be referred to in exceptional circumstances as Darktrace / CLOUD (Network). This specification applies for Customers who have purchased the product under any of the above naming conventions.

3.1. Product Features

3.1.1.High Level Summary

Darktrace / NETWORK provides Autonomous Response capabilities against network entities. Actions are triggered by specific criteria in Darktrace / NETWORK real-time detection or can be taken manually by an operator.

The Darktrace / NETWORK component applies Autonomous Response capabilities to physical and cloud network devices by controlling connectivity. This control is enacted using spoofed TCP RST packets issued by a Darktrace master, a Darktrace probe, or a virtual sensor such as a vSensor or osSensor. The actions available to Darktrace / NETWORK Autonomous Response can also be extended through integrations with endpoint security providers and with a number of popular firewalls, allowing Darktrace to enact connectivity controls using the rule or policy capabilities available in the external platform.

Please note, Darktrace / NETWORK Autonomous Response was previously known as both “Antigena” and “Darktrace RESPOND NETWORK”. This terminology is preserved in many places for the purposes of ensuring continuity.

Darktrace / NETWORK Autonomous Response expands Cyber AI response to devices by severing network connections, restricting access and quarantining devices by limiting connectivity. Actions can be taken when a device exhibits significantly anomalous behavior, when it contravenes a compliance policy, when a device attempts to access a specific watched endpoint, manually taken by an operator, or as a result of any other custom criteria defined in a Darktrace / NETWORK real-time detection model.

3.1.2.Data Retrieval and Ingestion

Darktrace / NETWORK real-time detection retrieves and processes raw network traffic. Network traffic may be delivered directly to a physical Darktrace appliance ingestion port, to a virtualized Darktrace vSensor instance, or forwarded from Darktrace host-based osSensor agents to a Darktrace vSensor. Darktrace / NETWORK Autonomous Response may also process network traffic delivered to Darktrace / NETWORK real-time detection when necessary to create network-level actions.

Darktrace / NETWORK Autonomous Response uses metadata from individual network packets to create “spoofed” response packets which are then sent to participating devices. This metadata must be retrieved directly from network traffic in order to create accurate response. Darktrace / NETWORK Autonomous Response does not process all network traffic - only relevant traffic necessary for action creation. This processing and analysis are separate to the Darktrace Deep Packet Inspection performed by Darktrace / NETWORK real-time detection on raw ingested traffic. Processing is only performed for the purposes of action creation; Darktrace / NETWORK Autonomous Response does not create any metrics or surface any metadata from this analysis.

To ensure Darktrace can create action-level blocks, all network traffic should be mirrored or ingested by Darktrace in some format. It is the responsibility of the Customer to ensure that Darktrace maintains visibility over network traffic in the event of network re-architecture outside the scope of the initial deployment design. The network traffic provided to Darktrace must be of suitable quality and not contain duplication, fragmentary data, or be incomplete in any fashion (for example, unidirectional).

Darktrace / NETWORK Autonomous Response may also retrieve configuration data from third-party firewall components where required for the creation of corresponding firewall policies or rules, or to confirm that configuration was performed successfully. Darktrace is not liable for interconnectivity issues resulting from software changes or updates made to third-party components.

3.1.3.Darktrace Analysis

Darktrace / NETWORK Autonomous Response does not perform direct analysis on ingested or modeled data; actions are taken as a result of data processed and analyzed by the equivalent Darktrace / NETWORK real-time detection capability.

3.1.3.1. Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of Darktrace ‘pattern of life’ analysis are evaluated.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior.

The models framework leverages both the underlying ‘pattern of life’ detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace / IDENTITY, and Darktrace / CLOUD modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet different criteria; these are created and maintained in the Darktrace Threat Visualizer “Model Editor” interface. Care should be taken when defining custom models to ensure that the number of alerts produced is not excessive and does not impact the system’s ability to function, or otherwise degrade service. Similarly, modification of existing default models should not result in excessive activity or alter the logic to the extent that alerts criteria can no longer be met. The responsibility to ensure that models created or edited remain within these reasonable expected boundaries lies with the Customer.

Default Darktrace models are focused on ‘pattern of life’ anomaly detection, potentially malicious behavior and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically, Customers without automatic updates will receive all applicable model updates when Darktrace ActiveAI Security Platform software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

3.1.3.2. Autonomous Response Models

In default operation, Darktrace / NETWORK Autonomous Response responses are triggered by model alerts from a specific subset of Darktrace models, categorized as Autonomous Response (“Antigena”) models.

Darktrace Autonomous Response models may directly look for specific behavior, or for indicators identified by other models operating within the real-time Detection environment. Darktrace / NETWORK models fall typically into the second category - “meta-models” - which are triggered by an alert of another Darktrace Real-Time Detection model.

Any modification to an underlying real-time detection model which results in increased or reduced model alerts will subsequently impact Darktrace / NETWORK Autonomous Response models and actions. Any under-activity or over-activity of the corresponding Darktrace / NETWORK Autonomous Response model as a result of this type of modification is the responsibility of the Customer.

Darktrace provides the ability to limit actions created by these models to a minimum score threshold.

3.1.3.2.1. Darktrace / NETWORK Autonomous Response Models

When Darktrace / NETWORK Autonomous Response is enabled within a Darktrace environment, a subset of additional models become available. These models are categorized into four high-level categories, where each category contains Darktrace / NETWORK Autonomous Response models which target specific activity within that category: “Compliance”, “External Threat”, “Insider Threat” and “Significant Anomaly”.

Each model is intended to trigger on specific types of connection or activity and will perform different Autonomous Response actions depending on the activity identified.

Darktrace / NETWORK Autonomous Response models within these four pre-defined categories contain logical criteria that will only evaluate successfully if the device entity that has met the criteria is also tagged with the corresponding Autonomous Response tag for the given model category. This concept is referred to as “Tag-Based Eligibility” and is outlined below.

The contents of each model folder are subject to change as Darktrace personnel regularly revise and update categorized models. Darktrace may also provide models within these categories which are disabled by default to allow operators to perform customization and tuning, or where detection may only be relevant to specific compliance scenarios (optional). Models of this type can be enabled at any time by the Customer.

3.1.3.3. Darktrace / NETWORK Autonomous Response Success

Darktrace / NETWORK Autonomous Response may analyze network traffic ingested by Darktrace / NETWORK real-time detection to establish whether actions were sufficient to end the targeted connectivity.

This analysis identifies connections which have continued to transmit data after an RST packet was issued by one or more components of the Darktrace ActiveAI Security platform. The exact calculations and thresholds used to determine this success state remain subject to change.

3.1.4.Output

Darktrace / NETWORK Autonomous Response creates automatic, network-level responses to anomalous behavior detected by Darktrace / NETWORK real-time detection or when triggered manually by a Darktrace operator.

3.1.4.1.1. Action Methodology

There are two methods that Darktrace / NETWORK Autonomous Response may invoke to control connectivity when an action is created - spoofed TCP RST packets sent directly from a Darktrace instance, or policy/rule-based actions through third-party network components. The action methodology is selected at the time of action creation. Methods are not mutually exclusive; multiple methods may be invoked as the result of a single action.

3.1.4.1.1.1. TCP RST Actions

Standard Darktrace / NETWORK Autonomous Response actions utilize “spoofed” TCP RST packets to trigger the source and/or destination device to end the connection. These TCP RST packets may be issued by any compatible Darktrace instance including physical Darktrace appliances, virtualized probes such as vSensors, or on-host osSensors.

Darktrace / NETWORK Autonomous Response takes information from ingested network traffic to create packets imitating the ongoing connection, which are then sent to each end of the target connectivity imitating the other. Darktrace / NETWORK Autonomous Response will send a number of packets to each device with varying TCP and IP header properties to increase likelihood of successful reception by the endpoint (and therefore successful interruption of the TCP connection).

In internal-to-internal connection scenarios, spoofed TCP RST packets will be directed at both source and destination device - only one of the connected devices needs to receive an RST packet for the connection to be terminated. Where one device is external (internal-to-external, external-to-internal), packets will be targeted at the internal connection device.

In specific scenarios where the source or destination are unreachable, it may be necessary to target an intermediary device such as a proxy server.

The ability for Darktrace / NETWORK Autonomous Response RST packets to reach the target connection is referred to as “reachability”. Darktrace offers multiple implementation modes to ensure reachability across the network environment including dedicated “firing” interfaces connected to different areas of the network, configuration for custom packet traversal routes, and VLAN tag support. This concept is outlined in greater detail below under the administration section.

3.1.4.1.1.2. *Integration Actions*

Darktrace / NETWORK Autonomous Response extends control to non-TCP network traffic by integrating with a selection of compatible firewalls and networking components. These integrations can be configured and used independently of - although typically in conjunction with - Darktrace / NETWORK Autonomous Response's TCP RST capabilities.

Utilizing a Darktrace / NETWORK Autonomous Response firewall integration offers an alternative route to neutralize connections crossing the network boundary, where UDP traffic is frequently seen, or in scenarios where Darktrace / NETWORK Autonomous Response RST packets may not be able to traverse the network from the Darktrace instance to their intended destination.

Specific implementation varies between third-party product model and vendor - for example. Darktrace may programmatically access the networking component to create rules/policies corresponding to the desired connection-level blocks, or the network component may actively poll Darktrace for endpoints intended for blocking.

Actions created via integrations are mapped as closely as possible to the desired inhibitor (see below). If a firewall does not offer the highly specific, targeted action desired by Darktrace / NETWORK Autonomous Response, an alternative, less specific action may be taken.

3.1.4.2. Darktrace / NETWORK Autonomous Response Action States

Darktrace / NETWORK Autonomous Response actions can be taken when: Darktrace identifies significantly anomalous device behavior; Darktrace identifies activity which contravenes a compliance policy; a device attempts to access a specific watched endpoint; manually triggered by an operator; or as a result of any other custom criteria defined in a Darktrace model.

If the current system configuration permits autonomous actions - or the action was triggered manually - the action will be created in *active* state. This created action will appear on the Response Actions Page "Network" tab under "Active" actions, in the Darktrace Mobile App, and be accessible via the Darktrace API. The device will be visually distinguished (at present, with a green highlight) across the Darktrace platform for the duration of the action to indicate it is currently undergoing a Darktrace / NETWORK Autonomous Response action. If alerts for Autonomous Response actions ("*created*" state) are configured, an alert will also be sent to configured outputs.

An *active* Darktrace / NETWORK Autonomous Response action permits Darktrace / NETWORK to take the responses defined by the action configuration. For example, to send corresponding TCP RST packets automatically when matching activity is seen or make any necessary configuration changes in third-party environments such as firewalls to control connectivity. As manual Darktrace / NETWORK response actions are created directly by an operator, actions are always created in the *active* state.

If Darktrace / NETWORK Autonomous Response is compelled to request human confirmation due to the current system configuration, the action will be created in *pending* mode. This created action will instead appear on the Response Actions Page "Network" tab under "Pending" actions until a user confirms that Darktrace / NETWORK Autonomous Response can take action. A notification requesting that a human user approve the action is sent to the Darktrace Mobile App and any configured Darktrace Autonomous Response alert outputs ("*pending*" state). The device will also be visually distinguished (at present, with an orange highlight) to indicate human approval is pending for a Darktrace / NETWORK Autonomous Response action. Once confirmed by an operator, the action will move to "Active" as before.

Pending actions are those which Darktrace / NETWORK Autonomous Response wishes to take but must require approval before it can act automatically; pending actions may therefore be referred to as "human confirmation mode" actions. Pending actions do not create firewall rules (or equivalent) or cause the sending of TCP RST packets until approval is granted in some format, at which the point the action becomes "active".

In the event that an Autonomous Response action is delayed due to the system configuration compelling the action to be created in pending mode, Darktrace bears no liability for any resulting connections, actions or damage that occurs as a result of the delay.

Darktrace provides mechanisms to alert users on the creation of pending actions requiring human approval to proceed. These notifications are surfaced in the Threat Visualizer interface, in the Darktrace Mobile App, and in compatible alert outputs.

3.1.4.2.1. Additional Action States

Darktrace Autonomous Response actions proceed through several states during their lifecycle. For example, an action is automatically created that *requires confirmation* (“pending”), it is then *confirmed* by a user. Another user *extended* the action, then *cleared* it. Finally, it was *reactivated* by another user, before *expiring*.

There are four key states (“pending”, “active”, “cleared”, “expired”) and two additional states (“extended”, “reactivated”):

- The first two states - pending and active are outlined above.
- Cleared actions are those which have been manually ended by an operator. Clearing informs Darktrace to cease the action, and to suppress the combination of Darktrace / NETWORK Autonomous Response action and model breach conditions for the remainder of the action’s set duration.
- Expired actions are those which are historic, regardless of their state (pending, active) before the action period passed.
- Reactivated actions are those which were made active again by an operator.
- Extended actions are those which are currently active and have had their duration manually extended by an operator.

Actions can be extended, reactivated or cleared from the Response Actions page, the Darktrace Mobile App, and the Darktrace Threat Visualizer API.

3.1.4.3. Darktrace / NETWORK Autonomous Response Actions in Default Operating Mode

Default operating mode (*default operation*) is here defined as Darktrace / NETWORK Autonomous Response operating in a state where autonomous (*active*) actions can be taken automatically against eligible devices when triggered by a default Darktrace / NETWORK Autonomous Response model.

Default operation presumes that Darktrace / NETWORK Autonomous Response is licensed and enabled globally on the Darktrace ActiveAI Security Platform instance, and Darktrace / NETWORK Autonomous Response has been configured according to the requirements of the network architecture and tested to ensure actions can be routed to the desired destination.

The conditions for the creation of action in any state in default operation are:

1. A Darktrace / NETWORK Autonomous Response model which has its logical criteria sufficiently met.
2. The prior criteria are met by a device or entity with the appropriate eligibility for the given model category.

3. A Darktrace / NETWORK Autonomous Response action (inhibitor) configured within the triggered model to define the action taken.
4. A system operating mode (schedule) which does not prevent the creation of a Darktrace Autonomous Response action at the given time.

The conditions for an *active* action in default operation are:

1. A Darktrace / NETWORK Autonomous Response model which has its logical criteria sufficiently met.
2. The prior criteria being met by a device or entity with the appropriate eligibility for the given model category.
3. An appropriate action (inhibitor) applied to the triggered model to define the action taken.
4. The model in question is configured to *permit* or *force* autonomous actions.
5. The system operating mode permits the creation of an active Autonomous Response action at the given time for the given device.
6. The model criteria are met by a device which does not possess configuration preventing autonomous actions.
7. The system operating mode permits the creation of an active Autonomous Response action at the given time for the given device.

Manual actions triggered by an operator are not subject to the conditions outlined above; these actions are always active on creation.

3.1.4.4. Manual Darktrace / NETWORK Autonomous Response Actions

Darktrace provides mechanisms to create Darktrace / NETWORK response actions manually. These actions can be created by a user with appropriate permissions from the Darktrace ActiveAI Security Platform Threat Visualizer interface, from the Darktrace Mobile App, or via the Threat Visualizer API.

Darktrace provides the option to enforce a free-text justification which must submitted by the end-user when the action is created.

Manual response actions require Darktrace / NETWORK Autonomous Response to be enabled globally. Actions are automatically created in an “active” state and do not require human activation. All devices are considered eligible for manual Darktrace response actions - the responsibility to ensure that the targeted device is appropriate lies with the Customer.

The type of action taken (*inhibitor*) is selected during creation. The responsibility to select an appropriate inhibitor is also the responsibility of the manual action creator.

3.1.4.5. Model-Triggered Darktrace / NETWORK Autonomous Response Actions

In default operating mode, Darktrace / NETWORK Autonomous Response actions are triggered by models from the four categories outlined above.

This triggered action is referred to as a “model action” - a system action taken in response to a specific model criterion being met. The creation of a Darktrace Autonomous Response action is a model action which exists by default on applicable Darktrace-created models contained within the “**Antigena**” category.

When the criteria for these models are met, Darktrace will invoke all model actions applied to the model, triggering a Darktrace / NETWORK Autonomous Response action as a result.

Operators may also add the Darktrace Autonomous Response model action to modified or newly created custom models. The responsibility for selecting an appropriate action inhibitor (see below), and for any actions created as a result of a custom model configuration, lies with the Customer.

3.1.4.6. Darktrace / NETWORK Autonomous Response Inhibitors

The actions Darktrace / NETWORK takes in response to a trigger will vary based upon the action configuration and the activity detected. The types of response actions taken are henceforth referred to as “inhibitors”, as they *inhibit* a specific form of behavior or connectivity.

The type of action Darktrace / NETWORK Autonomous Response takes is referred to as an inhibitor. Inhibitors can be “automatic”, where Darktrace / NETWORK selects the most appropriate action at the time, or pre-defined from a list of options.

In default operation, when the criteria for a Darktrace / NETWORK Autonomous Response model are met, Darktrace / NETWORK will attempt to create an action corresponding to the selected inhibitor for the trigger device or entity. The inhibitor is defined in the model directly.

Manual actions require the inhibitor to be selected during creation. This is applicable to actions created through the Darktrace Threat Visualizer, the Darktrace Mobile App, or the Darktrace Threat Visualizer API.

For actions created as a result of a user-created or user-modified model, Darktrace / NETWORK Autonomous Response will attempt to create an action corresponding to the inhibitor as defined in the model for the trigger device or entity.

At present, there are six pre-defined Darktrace / NETWORK inhibitors of increasing severity which may invoked through Darktrace / NETWORK Autonomous Response models or manual response actions: “Automatic”, “Block Matching connections”, “Enforce pattern of life”, “Block all outgoing traffic”, “Block all incoming traffic” and “Quarantine device”.

Typically, Darktrace / NETWORK Autonomous Response will take the inhibitor selected in the model action in response to the Darktrace / NETWORK Autonomous Response model criteria being met. However, there are necessary exceptions to this behavior:

- Actions are proportional to threat and may be escalated if granular blocks are not sufficient.

  For example, an action that targets *matching connections* is no longer sufficient where the number of connections deemed to be anomalous increases rapidly - Darktrace / NETWORK Autonomous Response will therefore escalate the severity of the action taken.

- Where it is necessary to invoke integrations with firewall vendors, the actions possible in the vendor environment may not be sufficiently granular to achieve the targeted block initially selected.

  In this case, Darktrace / NETWORK Autonomous Response will modify the action to match the capabilities of the integration. The limitations of specific integrations are outlined in the appropriate integration documentation.

- Where the device or entity has a protected characteristic which excludes it from the selected inhibitor.

3.1.4.7. "Flag for Antigena" / iagn

Darktrace / NETWORK includes a "Watched Domains" component - a list of hostnames, IPs and endpoints which, if observed in network traffic or operation, will trigger an alert of some form. Entries can be added manually, through the API (*/intelfeed*), through the Darktrace / Attack Surface Management integration, or via configured STIX/TAXII intelligence feeds. No default list exists for watched domains.

When endpoints are populated by operators - programmatically or through manual entry - the option is offered to take automatic Darktrace / NETWORK Autonomous Response actions against network connectivity to/from these endpoints. If this option is selected by an operator and an endpoint is subsequently seen in network connectivity, any Autonomous Response actions triggered are not surfaced in the user interface and are taken regardless of eligibility, device type, or how autonomously Darktrace / NETWORK Autonomous Response is permitted to act at the given time.

3.1.4.8. Autonomous Response Actions Initiated by Cyber AI Analyst

Cyber AI Analyst may initiate action creation - this functionality is optional and disabled by default. Darktrace Cyber AI Analyst can create Autonomous Response actions for detected activity which does not have corresponding active Darktrace Autonomous Response actions at the time of creation. This capability creates actions for "edges" between Darktrace Cyber AI Analyst incident events.

Actions can be created as a result of Cyber AI Analyst Incidents. AI Analyst links related devices and unusual activity together to create an overall incident structure. If AI Analyst detects that a link between two entities does not have a corresponding Autonomous Response action - and an inhibitor is available which would suitably target the activity exists - it may create an action to target this behavior.

3.1.4.9. Tag-Based Eligibility for Darktrace / NETWORK Autonomous Response

Device and credential entities can be "tagged" - marked with a specific, predefined identifier - within the Darktrace system. This may be performed manually within the Darktrace Threat Visualizer, programmatically via the Darktrace Threat Visualizer API, or by utilizing a model with a "tagging" action. The tags applied to a device entity, or a credential present on that device entity, are available to the Darktrace Model Engine during model evaluation, in addition to surfacing on the user interface.

A subset of tags necessary for system operation are supplied by default. Users may also define custom tags through the Darktrace Threat Visualizer User Interface or via the Darktrace Threat Visualizer API. Darktrace / NETWORK Autonomous Response adds five additional tags necessary for operation: "Antigena Compliance", "Antigena External Threat", "Antigena Insider Threat", "Antigena Significant Anomaly" and "Antigena All". Four of these tags correspond to a model category outlined above, with the exception of "Antigena All" which is equivalent to the previous four.

Darktrace / NETWORK Autonomous Response models within the four pre-defined categories contain logical criteria which will only evaluate successfully if the device entity which has met the prior criteria *is also tagged* with the corresponding Darktrace / NETWORK Autonomous Response tag for the given model category. For example, when a device is given the **Antigena Compliance** tag, it becomes eligible to receive Darktrace / NETWORK Autonomous Response responses triggered by models in the **Compliance** category.

This is referred to as *tag-based eligibility*.

The possession of a corresponding tag is therefore presumed as "eligibility" for a category of Darktrace / NETWORK Autonomous Response in default operating mode. Without eligibility, Autonomous Response inhibitors will not be effective.

Any modifications made by Customer operators to remove tag requirements from the default models outlined above are not considered within the scope of default operation. Darktrace does not recommend making modifications of this kind.

Ensuring effective tag-based eligibility across an environment is the responsibility of the Customer.

3.1.4.10. Darktrace / NETWORK Response Autonomy

Darktrace / NETWORK Autonomous Response can create actions in an "active" or "pending" state, were pending actions require human approval before any direct action is actually taken. The state actions are created in - referred to here as "autonomy" - is defined by the following factors and conditions.

3.1.4.10.1. Manually-Triggered Darktrace / NETWORK Response Actions

Manual actions are always automatically created in an "active" state.

3.1.4.10.2. Model-Triggered Darktrace / NETWORK Autonomous Response Actions

The status of actions created as a result of a Darktrace model alerts (both default Darktrace / NETWORK Autonomous Response models, and those user-created or user-modified) is defined by a combination of granular eligibility controls.

A high-level summary of the current autonomy level is provided to operators in the main Darktrace Threat Visualizer homepage. A detailed summary is also accessible from the Darktrace / NETWORK Response Quick Setup Process.

Darktrace / NETWORK Autonomous Response autonomy can be configured on a per-model basis to meet the varied requirements of each organization. This autonomy is controlled in the model configuration: models may "force human confirmation", "force autonomous action", or "permit autonomous action". Actions created by models with a "force" state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action created as a result of the model with "permit autonomous action" will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

Templates provided as part of the Darktrace / NETWORK Response Quick Setup Process include a subset of high-severity models recommended by Darktrace to be placed in "force autonomous action" state. A configuration interface to modify these states on a per-model basis is also provided as part of the process.

Darktrace also provides a seven-day, hourly timetable ("Response Schedule") which allows blocks of autonomous action or enforced human oversight to be scheduled. The schedule is applied across Darktrace / NETWORK, Darktrace / ENDPOINT and Darktrace / IDENTITY.

The schedule is applied to actions created by models with "permit autonomous action". Actions created manually, or by models in a "force" state ("force human confirmation", "force autonomous action"), are not impacted by the schedule state. Darktrace provides recommended timetables for the schedule as part of Darktrace Response Quick Setup Process templates. Access to configure the schedule directly is also provided within the Darktrace Response Quick Setup Process and from the Threat Visualizer interface.

For global organizations, the schedule can also be localized at the subnet-level. In this mode, the hours defined in the schedule are adjusted to apply in the time zone matched to the latitude and longitude information defined for the subnet in the Darktrace Threat Visualizer. Individual schedules can also be defined on a per-submaster basis in a distributed Darktrace deployment.

3.1.4.10.3. Per-Device Human Confirmation Mode

Darktrace also provides an advanced option to override autonomous actions on a per-device basis using a pre-defined tag.

3.1.4.10.4. Cyber AI Analyst-Triggered Actions

Actions created as a result of a Darktrace incident can "force human confirmation", "force autonomous action", or "permit autonomous action". This setting is set on the Darktrace System Config page and applies to all actions created by Darktrace Cyber AI Analyst.

Actions created when this setting is in a "force" state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action when set "permit autonomous action" will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

3.1.5.Interface – Darktrace Threat Visualizer

Read-only access to view Darktrace / NETWORK Autonomous Response actions is available to all users with access to the Darktrace Threat Visualizer interface. Modification of actions within these interfaces - including the creation of new actions - is permission restricted.

3.1.5.1. Overview

The main Threat Visualizer homepage displays a high-level overview of how autonomously Darktrace Autonomous Response can act at the current moment.

This overview summarizes across all possible factors, configurations, and settings to give a true representation of how autonomously Darktrace can act at each point of the day. When collapsed, the element will indicate the overall autonomy level and how long this state is applicable for.

A high-level, exportable summary of Darktrace Autonomous Response actions over the last seven days is also available from the Threat Visualizer Response Actions Summary view. This overview provides high level statistics regarding Darktrace / NETWORK Autonomous Response activity including the action types created, human approval granted, and the meantime to activation for pending actions.

Darktrace / NETWORK Autonomous Response actions are also included in reporting outputs (Executive Threat Report, Operations Report, Cyber AI Insights Report).

3.1.5.2. Actions

Darktrace / NETWORK Autonomous Response actions are recorded in the Response Actions window within the Threat Visualizer. Actions are sorted by type, then by state. A per-device filtered view is also offered. Users with appropriate permissions may alter the state of an action from this view (clear, activate, extend, or reactivate). Darktrace / NETWORK Autonomous Response actions are also surfaced in, and can be modified from, the Darktrace Mobile App.

Pending Darktrace / NETWORK Autonomous Response actions produce a notification above the threat tray in the Threat Visualizer. The count of pending actions is also displayed in the homepage summary. Darktrace offers alerting when Darktrace / NETWORK response actions are created or change state; notifications are surfaced in the Threat Visualizer interface, in the Darktrace Mobile App, and sent via compatible alert outputs.

Devices currently undergoing actions are visually distinguished (at present, with a green highlight) across the Darktrace Threat Visualizer interface for the duration of the action. Devices with pending actions are also separately distinguished (at present, with an orange highlight) to indicate human approval is pending for a Darktrace / NETWORK response action.

Connections prevented by Darktrace / NETWORK Autonomous Response are indicated in device event logs throughout the interface. Where network traffic ingested by Darktrace is found to include Darktrace / NETWORK Autonomous Response TCP RST actions, Darktrace Deep Packet Inspection will distinguish these packets from standard TCP RSP packets with a special code (“G/g”) in the Darktrace Advanced Search interface entry.

3.1.5.3. Autonomous Response Models

Darktrace / NETWORK Autonomous Response models in default operation mode will create standard model alerts in the Darktrace Threat Visualizer interface. Where an action is created as a result of a model, a pivot point from the model alert is also provided.

3.1.6.Reporting

The Darktrace Threat Visualizer “Audit Log” records changes made by operators to Darktrace Autonomous Response actions; the audit log can be exported via Syslog for extended retention.

User interactions with Darktrace / NETWORK Autonomous Response actions which alter the state (e.g. “activate”, “extend”, “clear”) are also recorded in the action history, accessible from the Darktrace Response Actions window. Users can also be compelled to provide a free-text justification when an action state is altered, which is also displayed where applicable.

Individual action-level information is surfaced primarily in the Darktrace Threat Visualizer Response Actions window and Darktrace Mobile App.

Darktrace provides a REST API for automated retrieval of a subset of data output.

3.2. Deployment Architecture

The appropriate combination of deployment components will vary depending on the network environment, for example, the balance of physical and virtualized infrastructure, or the physical location of covered components. The appropriate deployment scenario may also change during the lifetime of the Darktrace deployment if significant changes are made the network architecture. It is the responsibility of the Customer to ensure that Darktrace maintains both continued visibility and continued *reachability* over network traffic to facilitate operation of Autonomous Response.

Darktrace strongly recommends that deployments are architected in line with best operating practice. In the event of re-architecture of the environment, Darktrace recommends Customer reaches out to the Darktrace support team to determine the impact on the operation of Autonomous Response.

3.2.1.Topology Components

Darktrace / NETWORK Autonomous Response TCP RST actions may be issued from any compatible Darktrace ActiveAI Security Platform instance. All network architecture components will attempt to send spoofed Darktrace / NETWORK Autonomous Response TCP RST packets in response to targeted connectivity.

Compatible topology components include physical Darktrace probes, physical masters, physical subordinate masters, physical Unified View instances, vSensors with appropriate configuration and Darktrace osSensors. Darktrace-hosted virtual cloud masters are located within the Darktrace cloud environment and therefore cannot issue TCP RST packets to targeted endpoints. To take action in this scenario, a firewall integration or other Darktrace topology component (e.g. as a physical or virtual probe located within reach of the target entity) is required. For more information on the Topology Components, please refer to the Real-Time Detection section of the Product Specification.

3.2.2.Packet Traversal & Stateful Boundaries

In the majority of operating scenarios, Darktrace / NETWORK Autonomous Response TCP RST packets will be unable to cross stateful boundaries or may be limited in traversal by network routing components. This may be due to a number of factors including configuration restrictions in place within the networking component, anti-spoofing measures present in the component, or incomplete information in the network traffic sent to Darktrace which prevents it from appropriately routing the spoofed packet back to the target entity. The deployment must therefore be architected to ensure Darktrace / NETWORK Autonomous Response packets are able to traverse to all locations, including those behind stateful boundaries.

Darktrace provides mechanisms to ensure that packets can be routed to all areas of the network including dedicated “firing” network interfaces and custom-defined routing rules for spoofed packets.

Best practice recommendations are available for the implementation of Darktrace / NETWORK Autonomous Response which provide guidance on these considerations.

3.2.3.Firewall Integration

Darktrace / NETWORK Autonomous Response integrates with a selection of firewall vendors to allow Darktrace / NETWORK Autonomous Response’s targeted network blocking actions to be performed by third-party firewalls. These integrations can be configured and used independently of Darktrace / NETWORK Autonomous Response’s TCP Reset capabilities and are included in the Darktrace / NETWORK product.

Utilizing a Darktrace / NETWORK Autonomous Response firewall integration offers an alternative route to neutralize connections crossing the network boundary, where UDP traffic is frequently seen, or in scenarios where Darktrace / NETWORK Autonomous Response RST packets may not be able to traverse the network from the Darktrace instance to their intended destination. As such, Darktrace recommends integrating with all compatible firewalls to supplement the native TCP RST blocking to provide redundancy in case of unforeseen network changes.

3.3. Implementation

3.3.1.Licensing

Darktrace / NETWORK Autonomous Response will only take autonomous actions when a valid License Key is configured on the Threat Visualizer's System Config page. In Unified View environments, the License Key should propagate from the UV master instance to the subordinate Master instances unless specifically configured otherwise. Adding the License Key will also activate Darktrace / NETWORK Autonomous Response capabilities on all connected virtual sensors (vSensors). Instructions on adding a License Key can be found in the relevant documents in the Customer Portal.

3.3.2.Tagging & Scope

There are multiple approaches that can be taken to optimizing Darktrace / NETWORK Autonomous Responses for a given network environment. The exact methodology for deploying Darktrace / NETWORK Autonomous Response will depend both on the composition of the Customer's network and the Customer's own preference.

Darktrace / NETWORK Autonomous Response models indicate how, when and which response actions can be taken against devices. For a device to be in scope of Darktrace / NETWORK Autonomous Response it must have an "Antigena" tag applied to it. There are 5 default tags, as described above, that can be applied to devices to include them in the scope. Custom tags can be defined by modifying existing models or creating new models. Some tags may be included in models under other folders/categories, so it not advised to modify legacy models. If the Customer decides to modify the default tags or create custom ones, Darktrace bears no liability for the consequential impact on the operation of the Product.

Additionally, Customer can also add devices under scope by using the Darktrace / NETWORK Response Quick Setup. The Darktrace / NETWORK Autonomous Response Quick Setup Process offers one click setup option to select which devices are eligible, define any activity-based overrides, and to set the timed schedule.

It is the Customer's responsibility to determine the scope of Darktrace / NETWORK Autonomous Response by the application of tags.

Darktrace recommends all devices are in scope of Darktrace / NETWORK Autonomous Response as any untagged devices will not be protected during a compromise and any activities from those devices cannot be contained.

3.3.3.Testing & Reachability

Darktrace / NETWORK Autonomous Response's ability to access different areas of the network environment ("reachability") must be tested. If reset packets cannot traverse the network from the Darktrace instance to their intended destination, customer network will not be fully covered by Autonomous Response; Darktrace / NETWORK Autonomous Response will not be able to quarantine a device or perform the full range of autonomous actions.

Testing of actions and reachability should be done for every part of the customer's network, ensuring that each separated segment (whether logically/physically etc.) is reachable by the spoofed TCP RSTs. The Customer is responsible for running these tests and implementing any remediations. Given the complexities of networks, their propensity to change over time, and the reliance on other 3rd party software/hardware/firmware, the Customer should confirm reachability across all network segments every 6 months or if there is a significant network change.

Testing can be done manually or automatically using the Darktrace / NETWORK Autonomous Response Quick Setup Process "Spot Testing" capability (recommended). Instructions on how to perform testing can be found in the relevant document in the Customer Portal.

3.3.4.Dedicated Firing Interfaces (DFIs)

Darktrace / NETWORK Autonomous Response fires reset packets from the administrative interface of Darktrace instances by default. If the administrative interface is placed in a restricted portion of the network, the reset packets may not reach their destination. To ensure that packets are able to reach destination in separated parts of the network, it is possible to configure an additional or alternative firing interface from the Console called Dedicated Firing Interface (DFI).

For every appliance, at least one DFI is recommended to be installed. Customers using non-flat networks should configure a Custom Route to send RST packets through this DFI on each appliance, using the PACKET method initially. Customers using VLANs should additionally configure the Custom Route to use VLAN tags and ensure that all expected VLAN IDs are included in the mirror feed into the relevant Darktrace appliances/sensors.

In highly segregated networks it may be necessary to use multiple DFIs in order to achieve 100% reachability. The Customer may have to physically/logically plug a DFI from multiple Darktrace appliances across the network to bypass certain switches/firewalls or other network limitations. For vSensors, it is not usually needed to install a physical dedicated firing interface as these are mostly virtual. Additional virtual interfaces can be assigned to the vSensor VM and custom routes sent accordingly through this interface. However, it is also possible to attach a physical interface, providing the VM host machine has a spare NIC interface. This interface can then be used as a normal physical DFI plugged into a Darktrace appliance. Further details on how to configure DFIs can be found in the relevant documentation on the Customer Portal.

3.3.5.Schedule/Going Autonomous

Whether Darktrace / NETWORK can take action autonomously, or must wait for human approval, is defined by a schedule or model setting. The seven-day, hourly timetable allows blocks of autonomous action or enforced human oversight to be scheduled. For some high severity activity types, it is recommended to ensure Darktrace / NETWORK Autonomous Response can always take an autonomous action regardless of the scheduled state; in this case, an override can be configured in the model setting.

When using the Darktrace / NETWORK Autonomous Response Quick Setup, One Click Setup options are graded from least to most autonomous. In the most autonomous mode, Darktrace / NETWORK Autonomous Response will take action autonomously in all cases, requiring no human oversight. In partially autonomous operating mode, Darktrace / NETWORK Autonomous Response is restricted to taking action autonomously only outside of business hours, when there may be no one to approve pending actions. Finally, human confirmation mode compels Darktrace / NETWORK Autonomous Response to request human confirmation before acting in all but the most severe cases.

But for models that block trigger on privileged user activities, Darktrace recommends that Darktrace / NETWORK Autonomous Response models to be set to autonomous every hour of every day (24/7). Human Confirmation mode should be considered for initial setup as it allows customizations to be made based on the behavior Darktrace is seeing in different networks but is not the designed end-state for Darktrace / NETWORK Autonomous Response. Please see above for the possible methods of implementation that allow gradual implementation, via a hybrid deployment, to reach a fully autonomous deployment. A hybrid deployment is considered any that is neither fully autonomous nor fully human confirmation mode.

The Customer is responsible for configuring Darktrace / NETWORK Autonomous Response to run fully autonomously. Fully autonomous mode is the end goal of all Darktrace / NETWORK Autonomous Response deployments; reaching a fully autonomous state where the AI can take action whenever unusual or concerning behavior is detected, without the need for human oversight. This mode lends itself to a minimal-interaction workflow, where Customer may infrequently modify actions through the Darktrace Threat Visualizer interface, API, or Darktrace Mobile App, but on the whole leave Darktrace / NETWORK Autonomous Response to operate with little intervention.

3.3.6.External Alerting & Mobile App

Darktrace supports a number of Workflow Integrations directly or through generic formats such as syslog; all integrations can be reviewed on the Modules page of the Darktrace System Config page. Darktrace recommends external alerting from the UI to be configured for Darktrace / NETWORK Autonomous Response actions.

Darktrace / NETWORK Autonomous Response action alerts are a new alert format for Darktrace 6.1, providing even greater visibility over Darktrace Autonomous Response activity and notifying users when the AI is waiting for Customer approval. Each alert includes a link back to the relevant action to quickly grant approval or alter the action in some way. By default, alerts are created when an action is pending and Darktrace needs human authorization to step in but can be configured for any change of state. For Darktrace environments running in a fully autonomous mode, customers who wish to keep abreast of Darktrace activity can receive alerts when autonomous actions are created automatically, or when they are cleared or expired.

Darktrace also recommends using the mobile app and configure it for Autonomous Response notifications. The Darktrace mobile app allows users to easily access Darktrace alerts when they are on the move. This allows alerting similar to that mentioned above but also allows interaction with response actions from the app itself. This improves ease of interaction with Darktrace / NETWORK Autonomous Response actions for the Customer. Find information to configure the mobile app on the Customer Portal.

3.3.7.Active Integrations

Darktrace provides a wide range of components and integration methods that extend Darktrace / NETWORK Autonomous Response capabilities across various vendors - and ensure that Darktrace fits seamlessly into any existing security stack. Darktrace's Autonomous Response component offers a selection of direct integrations with enterprise firewall models. In scenarios where RST packets may not be able to traverse the network from the Darktrace appliance to their intended destination, or UDP traffic is frequently seen, firewall integration can ensure that connections are neutralized at the boundary. This is particularly useful in network environments with multiple internal boundaries, or where firewall operators desire more control of Darktrace / NETWORK blocking capabilities. Details on what integrations Darktrace offer can be found in the Customer Portal.

Darktrace recommends customers integrate with all available EDR and firewall third parties to supplement the native spoofed TCP RST blocking mechanisms. This allows blocking of UPD traffic but also adds redundancy for complex networks where 100 % reachability is still being addressed and to ensure no single point of failure.

Firewall integrations vary significantly between vendors and ensuring operational integration is a Customer Responsibility. In most cases, lists are required to specify the type of block the firewall will undertake on behalf of Darktrace. For such firewalls, wherever possible, Darktrace recommends configuring 3 lists in the Darktrace module and in the firewall itself. These should correspond to:

- block all outgoing connections from a source IP,
- block all incoming connections to a destination IP, and
- block matching connections (all connections from a source IP to a destination IP).

Configuring these three lists allows Darktrace / NETWORK Autonomous Response to leverage the firewall actions without significantly reducing the precision of the native spoofed TCP RSTs.

3.3.8.Model Editing

As stated, edits made to Darktrace / NETWORK Autonomous Response models are made at the Customer’s own risk. Darktrace recommends that the best practice is to edit the underlying real-time detection models, if necessary, rather that the Autonomous Response models. In the event automated scripts and systems unintentionally trigger models, it is recommended that any defeating of the underlying detection is conducted as precisely as possible.

Autonomous Response is configurable to the Customer’s security-availability requirements; however, the highest level of security is achieved when no model tuning is conducted. Darktrace only recommends defeating detection triggers for response models or actions that the Customer deems an impact to business continuity and at the Customer’s discretion. Note that there will always be some actions against benign activity, but precise actions should not cause business impact. If business processes would be impacted by the action taken or suggested, then Darktrace recommends addressing this.

In most deployments, the data collect from one month’s worth of Autonomous Response model alerts provides sufficient insight into the impact of actions and their frequency, to allow for informed decisions to be made regarding tuning. For details on the possible edits and how to implement them, see the relevant documentation on the Customer Portal.

3.4. Administration

Darktrace strongly recommends that deployments are architected in line with best operating practice. This includes administration of physical implementations or networking implementations outside the Darktrace Threat Visualizer environment, and the administration of device eligibility, action autonomy, and other components within the Threat Visualizer platform.

Darktrace provides mechanisms to assist with the administration of Darktrace / NETWORK Autonomous Response within the Darktrace Threat Visualizer. These mechanisms are outlined below.

Deployments should be tested frequently to ensure that Darktrace / NETWORK Autonomous Response retains access to control all connectivity as desired. This responsibility for regular testing and resolution of issues identified during tests lies with the Customer. Please refer to “Reachability” below.

3.4.1.Darktrace / NETWORK Response Quick Setup

Darktrace provides a “Quick Setup Process” (hereafter *QSP*) for Darktrace / NETWORK Autonomous Response. This process guides users through enabling Darktrace / NETWORK Autonomous Response, the selection of eligible devices, the types of activity Darktrace / NETWORK Autonomous Response is permitted to target, and the level of autonomy available to Darktrace / NETWORK response in each case. This process offers both templated operating scenarios and granular configuration steps for each deployment aspect.

The QSP is intended for standard Darktrace / NETWORK Autonomous Response deployment and operation. The QSP does not include initial roll-out processes for Darktrace / NETWORK Autonomous Response routing or integrations - these elements must be completed before attempting the process.

3.4.2.Device Eligibility

The responsibility to ensure that devices are suitably tagged - and remain tagged - lies with Customer. Devices which are not tagged for Darktrace / NETWORK Autonomous Responses will not have actions taken against them in default operating mode.

Darktrace provides a mechanism for automatic device tagging via the Darktrace / NETWORK Response Quick Setup (method A). All templated deployment scenarios provided within the Response Quick Setup include a minimum recommended tag eligibility for devices. Eligibility is divided between client and server devices. An advanced option to define eligibility rules on a *per device type* or *per subnet basis* is also offered within the Response Quick Setup.

Separate to the Response Quick Setup, Darktrace also provides a default subset of customizable models (method B) which can be modified to opt devices into Darktrace / NETWORK Autonomous Response coverage.

These mechanisms (A & B) for defining tag eligibility are rules-based and will automatically apply tags to newly observed devices which meet the criteria.

Tags can also be added manually to individual devices using any mechanism outlined above (in "Tag-Based Eligibility"). Manual tagging is not recommended as the primary eligibility method; the static nature of this eligibility is unlikely to keep step with asset or network changes over time.

Eligibility should be regularly reviewed to ensure it has remained in-step with your network environment. Darktrace provides mechanisms to review the breadth of tag eligibility coverage within the Darktrace Threat Visualizer.

3.4.3. "Reachability"

Reachability refers to the ability for Darktrace / NETWORK Autonomous Response to reach all connectivity targeted by actions.

The responsibility to ensure reachability is tested regularly and any factors impacting reachability are resolved lies with the Customer operator. Reachability testing can be performed manually, or via the mechanisms provided by Darktrace.

Darktrace / NETWORK Response Quick Setup is accompanied by an automated testing system in technical preview. This reachability *spot tester* will take quarantine actions against nominated devices in each subnet over the duration of test; results are then reported in the Threat Visualizer, and Darktrace / NETWORK Autonomous Response will also suggest possible reasons for unsuccessful tests.

Testing Darktrace / NETWORK Autonomous Response is an iterative process - every subnet should be tested to ensure that Darktrace / NETWORK Autonomous Responses can reach all eligible devices if the need arises. Network configurations change frequently; tests should also be repeated regularly to ensure routes have not been severed by new stateful boundaries.

4. Requirements

4.1. Physical Install Requirements

4.1.1.Network Access Ports

It may be necessary to configure your network security devices to permit access to and from the Darktrace physical appliances for administration and other services. The required connections below must be enabled before attempting initial setup:

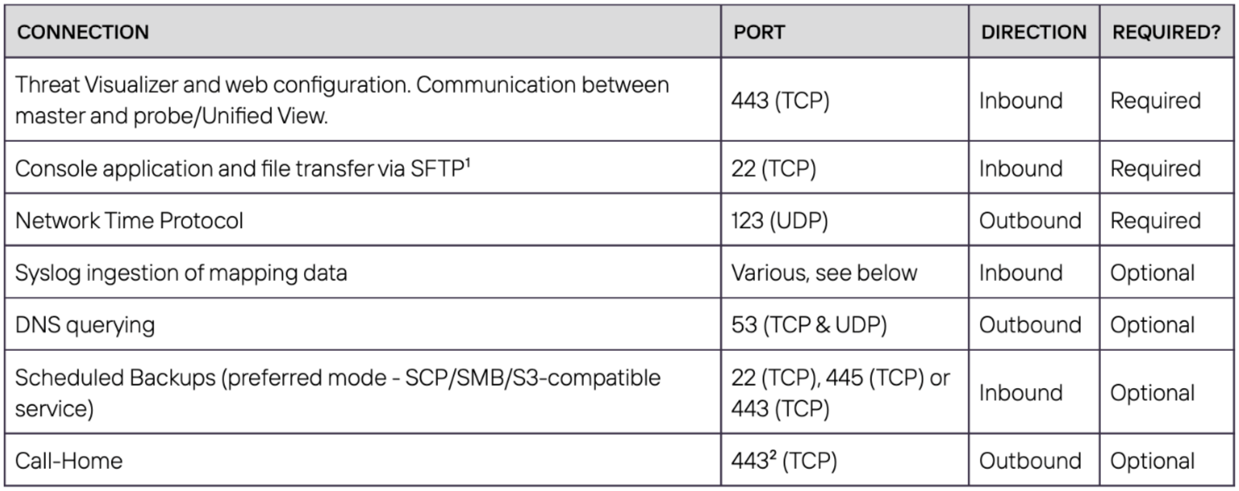

CONNECTION	PORT	DIRECTION	REQUIRED?
Threat Visualizer and web configuration. Communication between master and probe/Unified View.	443 (TCP)	Inbound	Required
Console application and file transfer via SFTP[1]	22 (TCP)	Inbound	Required
Network Time Protocol	123 (UDP)	Outbound	Required
Syslog ingestion of mapping data	Various, see below	Inbound	Optional
DNS querying	53 (TCP & UDP)	Outbound	Optional
Scheduled Backups (preferred mode - SCP/SMB/S3-compatible service)	22 (TCP), 445 (TCP) or 443 (TCP)	Inbound	Optional
Call-Home	443[2] (TCP)	Outbound	Optional

Connectivity between a Darktrace Unified View instance and any subordinate Masters requires a minimum bandwidth of 5Mbps.

4.1.2.Call Home

Outbound firewall exceptions may be necessary for the appliance to connect to the Call-Home service. Each instance is provided with a unique Call-Home endpoint, which can be retrieved from the console of the instance after boot.

4.1.3.System Backups

Automatic software backups can be configured and exported to an external location over SCP, SMB or to a S3-compatible location. Please ensure the selected backup protocol/mode is allowed between the backup storage location and the Darktrace appliance. It is a Customer responsibility to configure backups.

4.1.4.Syslog Input

Inbound syslog ingestion is supported on several port/protocol combinations. The most suitable is dependent upon your network environment and the Darktrace components deployed, such as the supported export formats of the syslog sender or whether syslog must traverse untrusted networks.

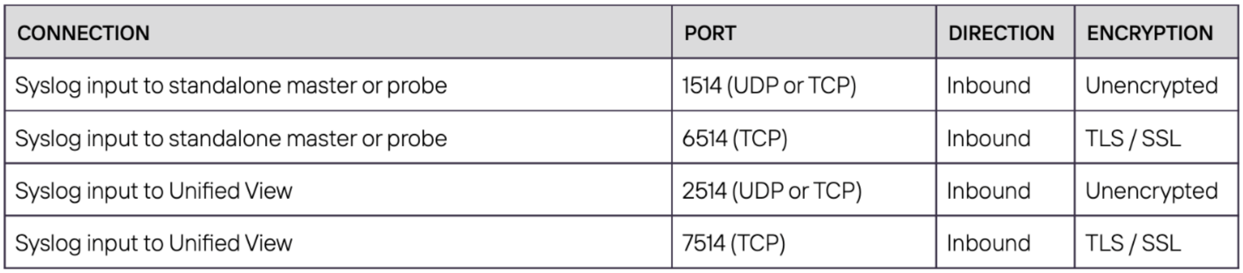

CONNECTION	PORT	DIRECTION	ENCRYPTION
Syslog input to standalone master or probe	1514 (UDP or TCP)	Inbound	Unencrypted
Syslog input to standalone master or probe	6514 (TCP)	Inbound	TLS / SSL
Syslog input to Unified View	2514 (UDP or TCP)	Inbound	Unencrypted
Syslog input to Unified View	7514 (TCP)	Inbound	TLS / SSL

The default ingestion port can be altered in the host variables of the appliance after configuration. Unused syslog ports can also be disabled by the customer if desired for compliance purposes.

4.1.5.Powering the Appliance

Appliances shipped with Darktrace Threat Visualizer 4.0.7 and above will auto-boot on power. It is strongly recommended that redundant and uninterruptible power supplies are used to maintain service in the event of primary power failure.

For appliance models that contain dual-power supply units for redundancy, both supplies must be connected at all times. The units will charge when the power supplies are connected.

Some models may contain hardware RAID battery backup units (BBU’s). In common with other RAID battery backup units, and to ensure data integrity, it is recommended that the units have at least 20% charge before full disk access can be assured.

4.2. Darktrace Cloud Region Availability & Security

If the customer organization is subject to regional or geographic restrictions on data flow that would prevent use any of the following cloud regions, the customer must disclose this to Darktrace. The customer must enable network access from probe components to the relevant cloud region.

For security reasons, the cloud-hosted master will not accept or ingest unencrypted data. Network traffic, log data, and other data types must be sent over a secure channel, such as a vSensor operating in an approved mode.

4.2.1.AWS Regions

Darktrace offers AWS cloud-based deployments hosted in Europe (AWS region "eu-west-1" or "eu-west-2"), the United States (AWS regions "us-west-1" and "us-west-2"), Canada (AWS region "ca-central-1"), Singapore (AWS region "ap-southeast-1") or Australia (AWS region "ap-southeast-2").

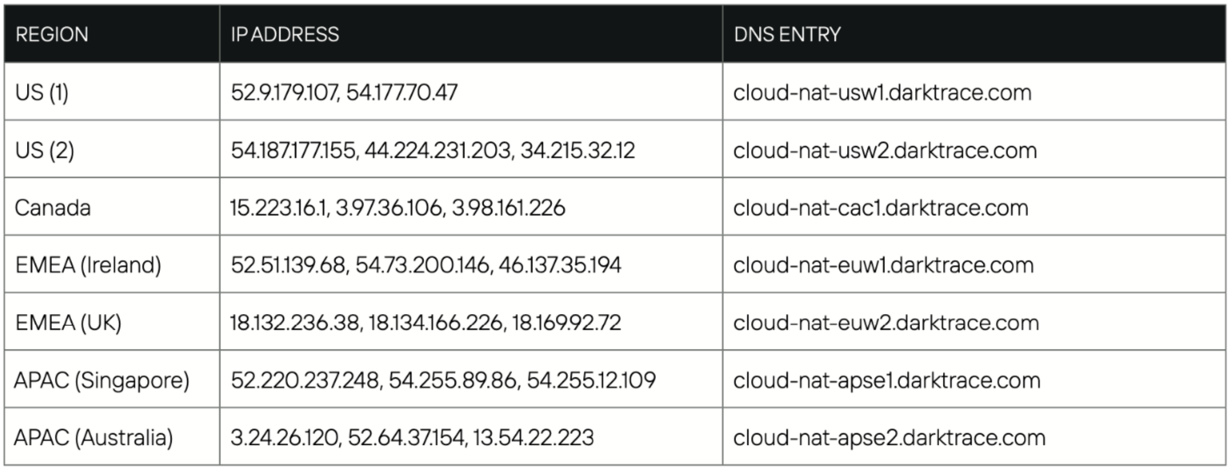

REGION	IP ADDRESS	DNS ENTRY
US (1)	52.9.179.107, 54.177.70.47	cloud-nat-usw1.darktrace.com
US (2)	54.187.177.155, 44.224.231.203, 34.215.32.12	cloud-nat-usw2.darktrace.com
Canada	15.223.16.1, 3.97.36.106, 3.98.161.226	cloud-nat-cac1.darktrace.com
EMEA (Ireland)	52.51.139.68, 54.73.200.146, 46.137.35.194	cloud-nat-euw1.darktrace.com
EMEA (UK)	18.132.236.38, 18.134.166.226, 18.169.92.72	cloud-nat-euw2.darktrace.com
APAC (Singapore)	52.220.237.248, 54.255.89.86, 54.255.12.109	cloud-nat-apse1.darktrace.com
APAC (Australia)	3.24.26.120, 52.64.37.154, 13.54.22.223	cloud-nat-apse2.darktrace.com

4.2.2.Microsoft Azure Regions

Darktrace can offer Azure cloud-based deployments hosted in EMEA (Azure regions "UK South", "West Europe", "UAE North" or "South Africa North"), the United States (Azure region "East US"), Canada (Azure region "Canada Central"), South East Asia (Azure region "Southeast Asia") or Australia (Azure region "Australia East").

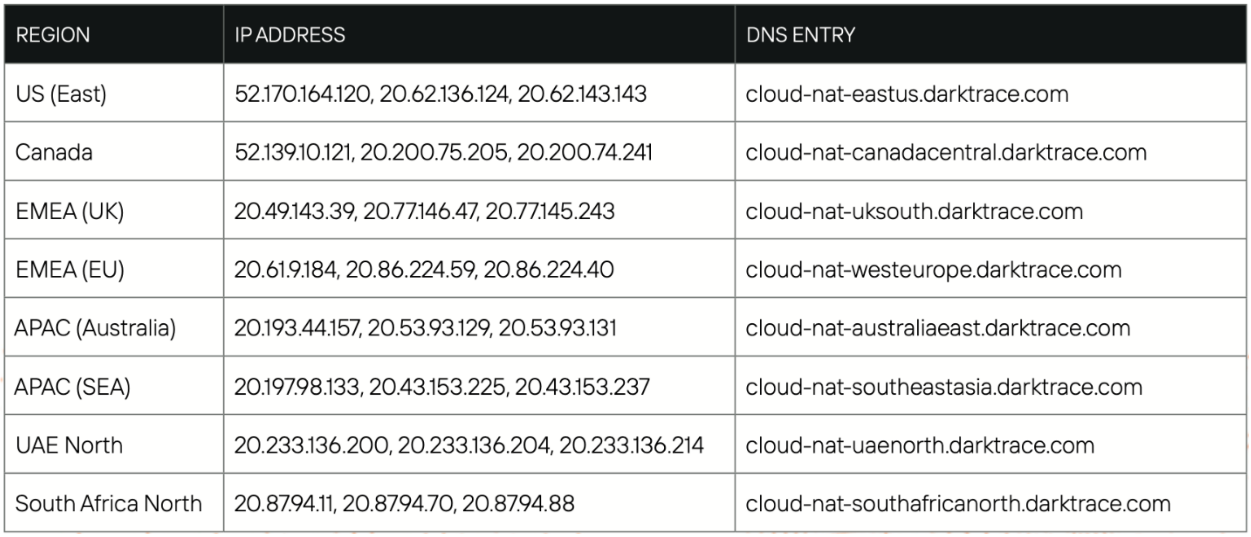

REGION	IP ADDRESS	DNS ENTRY
US (East)	52.170.164.120, 20.62.136.124, 20.62.143.143	cloud-nat-eastus.darktrace.com
Canada	52.139.10.121, 20.200.75.205, 20.200.74.241	cloud-nat-canadacentral.darktrace.com
EMEA (UK)	20.49.143.39, 20.77.146.47, 20.77.145.243	cloud-nat-uksouth.darktrace.com
EMEA (EU)	20.61.9.184, 20.86.224.59, 20.86.224.40	cloud-nat-westeurope.darktrace.com
APAC (Australia)	20.193.44.157, 20.53.93.129, 20.53.93.131	cloud-nat-australiaeast.darktrace.com
APAC (SEA)	20.197.98.133, 20.43.153.225, 20.43.153.237	cloud-nat-southeastasia.darktrace.com
UAE North	20.233.136.200, 20.233.136.204, 20.233.136.214	cloud-nat-uaenorth.darktrace.com
South Africa North	20.87.94.11, 20.87.94.70, 20.87.94.88	cloud-nat-southafricanorth.darktrace.com

4.2.3.Darktrace Probe Requirements

4.2.3.1. vSensor requirements

The vSensor ingests and processes network traffic before sending it to the Darktrace Master instance. The size of data transferred across the network is approximately 1-4% of the traffic ingested by the vSensor. The table below shows the required network ports for communication.

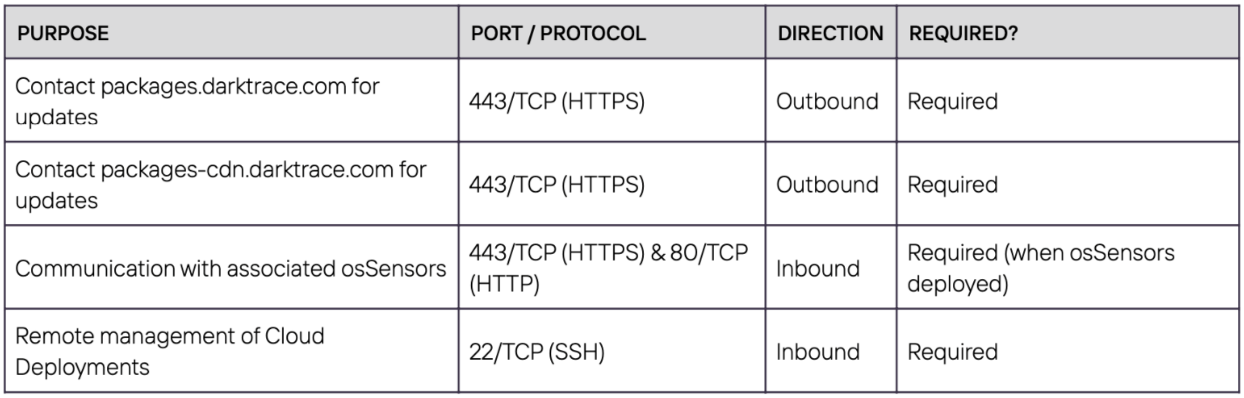

PURPOSE	PORT / PROTOCOL	DIRECTION	REQUIRED?
Contact packages.darktrace.com for updates	443/TCP (HTTPS)	Outbound	Required
Contact packages-cdn.darktrace.com for updates	443/TCP (HTTPS)	Outbound	Required
Communication with associated osSensors	443/TCP (HTTPS) & 80/TCP (HTTP)	Inbound	Required (when osSensors deployed)
Remote management of Cloud Deployments	22/TCP (SSH)	Inbound	Required

Ubuntu 20.04 Focal Server is required for Darktrace vSensor 5.1 and above. This operating system is supported by Linux. Darktrace will align to update the vSensor to newer operating system before end-of-life support by Linux.

vSensors require at least two CPU cores but works best with three or more. 2GB RAM is required, however 3.75-4 GB RAM is a highly recommended minimum for substantial performance.

Deep Packet Inspection workers scale automatically, it is recommended to keep at least a 2GB per core ratio, with 4GB per core being optimal.

PCAPs are stored on the vSensor disk. A minimum disk of 20GB is recommended, retention is dependent on available storage and traffic ingestion, so better retention can be achieved with more disk space provisioned.

4.2.3.2. osSensor Requirements

To perform the installation process for a Darktrace osSensor, there must be a Darktrace Master appliance running the most recent version of the Darktrace Threat Visualizer, and the IP of a configured vSensor must be able to communicate with the intended osSensor location.

4.2.4.Darktrace Server Agent Requirements

The Server Agent must be installed on each instance of the Server hosting centralized applications. The TSA customer server on which the Agent is deployed location(s) must be able to contact the Darktrace Master instance in order to send connection information.

5. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer's responsibilities are to ensure that:

- Darktrace / NETWORK Real-Time Detection has been deployed in line with Darktrace best practice recommendations and, if changes are made to the network configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.
- Darktrace maintains visibility over all network traffic to be monitored.
- The network traffic provided to Darktrace is of suitable quality and does not contain duplication, fragmentary data, or is incomplete in any fashion.
- The network traffic provided to Darktrace components is within the relevant specification of those components and that deployment components are not overloaded.
- System health issues are monitored using the built-in features of the Darktrace / NETWORK product, and if alerts arise, are addressed and rectified in a timely manner.
- The subnets detected by Darktrace / NETWORK Real-Time Detection are configured to be tracked by the optimum available identifier.
- End-user access is managed securely and within recommended best practice.
- Changes made to components or overall system configuration by users of the Darktrace platform do not impact the system's ability to function, or otherwise degrade service.
- Darktrace / NETWORK Autonomous Response has been deployed in line with Darktrace best practice recommendations and, if changes are made to the network configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.
- Darktrace / NETWORK Autonomous Response has - and maintains - reachability to all devices intended for Darktrace / NETWORK Autonomous Response actions. Where reachability is impeded, it falls within the customer's responsibility to identify the impedance and ensure that reachability is restored.
- Where Darktrace / NETWORK Autonomous Response TCP RST packets are unable to cross stateful boundaries or are limited in traversal by network routing components, to ensure that a dedicated firing interface or alternative configuration is in place to route packets to all network locations.

- There is knowledge and understanding of customer network topology, systems, etc. such that appropriate best practices and necessary modifications to Autonomous Response can be implemented.
- Darktrace / NETWORK Autonomous Response is enabled in the appliance console (where relevant), in the deployment settings, and in any other applicable locations to ensure actions can be taken.
- Darktrace / NETWORK Autonomous Response is granted necessary autonomy to enact actions, or, where human confirmation mode is partially or fully in place, ensure that actions are activated by human operators in a timely manner.
- Suitable workflows and processes are in place to remediate any cyber-attacks/compromises that are contained by Darktrace / NETWORK Autonomous Response
- Devices are opted appropriately into Darktrace / NETWORK Autonomous Response using tags.
- Alterations are not made to Darktrace / NETWORK Autonomous Response models, or other Darktrace models utilized by Darktrace / NETWORK Autonomous Response, which would significantly impede the ability to create targeted actions.
- Darktrace / NETWORK Autonomous Response actions are not created excessively by custom models or existing models modified to contain actions.
- All desired integrations with third party firewalls and EDRs are configured to enable action by Darktrace on all desired devices.
- The outputs of the Darktrace / NETWORK Autonomous Response are monitored, and its use complies with all applicable laws and regulations.

Examples of configuration changes which may result in degraded service include the exclusion of traffic or devices from deployment scope, the incorrect modification of subnet tracking configuration, misidentification of device types, the creation of excessively overactive models, and the modification of existing models to the extent that alerting criteria can no longer be met.

The RACI matrix below details the accountability and responsibility between Darktrace and Customer where no Additional Services have been purchased. Customer should be familiar with their existing security & network infrastructure and must be able to perform the actions outlined.

Role	**Activity/Responsibility**
R – Responsible	The party is responsible for implementation of the activity; owns the problem or project
A - Accountable	Right to make decisions, signs, and improves work
C – Consulted	Has information, resources, and/or capacities necessary to assist the job
I – Informed	Must be informed of the results, but does not need to be consulted

Activity/Responsibility	Darktrace	Customer
Implementation, Administration & Maintenance		
Product Implementation, Administration and Maintenance	CI	RA

Provision of guides to support Implementation, Administration & Maintenance	RA	CI
Product Licensing	RA	CI
Maintain coverage, including visibility, tagging and eligibility within product license	CI	RA
Monitor and resolve system status alerting	CI	RA
Administrate access to the Threat Visualizer	CI	RA
Administrate access to the Darktrace Customer Portal	CI	RA
Performing regular testing to ensure continued reachability	CI	RA
Managing active integrations and model editing	CI	RA
Ongoing maintenance of Autonomous Response configurations	CI	RA
Physical components		
Provide shipping and appliance configuration information	CI	RA
Provide appliances, subject to customer-provided information	RA	CI
Configure network settings for appliance	CI	RA
Maintain & modify traffic feeds as required, within appliance specifications	CI	RA
Modify Firewall rules to allow communication with Darktrace infrastructure (Call Home)	CI	RA
Modify firewall rules to allow communication between Darktrace Master and Probe components	CI	RA
Virtual components		
Make software available as detailed in customer contract terms	RA	CI
Download software from Darktrace infrastructure	CI	RA
Ability to access and deploy new virtual machines to support sensor components	CI	RA
Modify network access rules to allow communication between Darktrace sensors and Master instance	CI	RA

Modify network access rules on endpoint to allow communication & updates from Darktrace infrastructure	CI	RA
Maintain & modify traffic feeds as required	CI	RA

6. Considerations

Where source data is low quality, represents only a subset of all network activity, is delayed by external factors, or otherwise is incomplete, the quality of Darktrace analysis will be impacted. It is therefore imperative that data ingestion is configured comprehensively and to a high standard by the client. Analysis may also be impacted by overloading, such as that observed when traffic throughput significantly increases beyond scoped levels.

Each deployment has a specified capacity, determined by reference to the Usage Metrics band that has been purchased. Exact performance may vary depending on several factors that may be unique to the network being analyzed. Staying within Usage Metrics band allows for optimal Darktrace performance.

In the event that Customer has provided incorrect or incomplete sizing information or goes allows their usage to exceed the Usage Metrics band, Darktrace bears no liability should the actual traffic in Customer environment exceed the sizing information and degrade service.

In the event that actual traffic in Customer environment exceeds the Fencing Metrics as calculated for the relevant Usage Metrics band, Customer acknowledges this may result in, without limitation: overloading, inconsistent service, delayed Autonomous Response actions, a processing queue for new traffic, packet drops, and unsupervised learning being turned off.

Environments which have not been adequately configured to track devices by the most appropriate, consistent factor, or where tracking has not been updated in line with network changes will be unable to develop long-term ‘pattern of life’ behavioral profiles for unusual activity detection.

Where a Darktrace environment undergoes significant system load - typically due exceeding recommended operational limits - “High performance mode” may be activated. This mode disables a subset of high load ‘pattern of life’ classifiers to reduce CPU/memory usage within acceptable limits.

Darktrace / NETWORK Autonomous Response can only take automatic actions when granted appropriate autonomy. If actions are left in a pending state, Darktrace / NETWORK Autonomous Response cannot perform the required actions to mitigate connectivity.

Darktrace / NETWORK Autonomous Response TCP RST packets are typically unable to cross stateful boundaries without configuration specific to customer environment and may also be limited in traversal by network routing components. Darktrace provides best practice guidance to mitigate these limitations. The customer must ensure that these guidelines are followed to provide routing of response packets to all locations.

Where the source and destination device are located in close physical proximity, the round-trip time for the ingestion of the mirrored traffic into the Darktrace instance and subsequent RST packet creation may prevent successful connection interruption.

Darktrace / NETWORK Autonomous Response will be unable to target short-lived connectivity which has completed before traffic is received by the Darktrace instance. Where Darktrace / NETWORK Autonomous Response creates an active action, but no subsequent connectivity which meets the criteria is observed during the duration, no RST packets will be issued despite the action creation.

Darktrace / NETWORK Autonomous Response TCP RST actions are not compatible with UDP traffic.

In the event that a bug is identified, the Customer must notify Darktrace in writing. Darktrace agrees to make commercially reasonable efforts to fix or provide a workaround for any critical bugs but makes no guarantees regarding the time required to resolve non-critical issues.

7. Roles

Customer Role	Responsibility
Darktrace Product Owner / Project Manager	Coordinates Customer resources as necessary. Serves as the primary point of contact between Customer and Darktrace. Drives communication from the Customer side. Serves as the point of escalation for issue resolution and service-related issues.
Users of Darktrace Platform	Leverages Product functionality; monitors & actions output. Provides feedback to other Customer & Darktrace roles
Network/System Administrator	Ensures Darktrace / NETWORK coverage is maintained. Performs regular checks of data inputs. Responds accordingly to System Status alerts to ensure optimal traffic quality and delivery. Maintains call-home connectivity with Darktrace (where enabled).
Customer Portal Primary User	Manages customer portal access and contact information for all other customer roles. Ensures that Service Contacts are verified.

Darktrace Role	Responsibility
Customer Success Manager	Oversee customer’s experience using Darktrace. Act as customers’ sponsor & escalation path internally within Darktrace.
Darktrace Customer Support	Provides Support Services as per Master Customer Agreement
Account Executive	Facilitate commercial arrangements between Customer and Darktrace
Solutions Engineer	Provides technical expertise to support Account Executive and Customer Success Manager, regarding expansions or changes of coverage.

Darktrace / IDENTITY

Table of Contents

1. Introduction

Darktrace / IDENTITY is a market leading cyber security solution that integrates with enterprise software and cloud platform solutions to bring visibility and threat analysis to critical systems. Extending Darktrace's ActiveAI Platform beyond the physical enterprise network, each module brings the insight of the Cyber AI Analyst and Darktrace's unique 'pattern of life' anomaly detection to enterprise software and cloud-based environments, monitoring user activity whether it originates inside the network or from remote locations.

This document it to be read in conjunction with the Darktrace Master Service Agreement which governs the usage of the Darktrace Product Offering. References to "Customer" throughout this document should be read to refer to the entity that is the owner of the software subscription and is ultimately responsible for its operation, whether as end-user or service provider.

Darktrace / IDENTITY Subscriptions are sold according to the Usage Metrics bands set out in the relevant Product Order Form.

Darktrace / IDENTITY was previously named Darktrace DETECT/Apps, Darktrace RESPOND/Apps Darktrace DETECT/Zero Trust, and Darktrace RESPOND/Zero Trust (excluding Zero Trust/Netskope and Zero Trust/Zscaler ZIA/ZPA). Darktrace / IDENTITY that is implemented into cloud-hosted environments may be referred to in exceptional circumstances as Darktrace / CLOUD (Identities). This specification applies for Customers who have purchased the Product under any of the above naming conventions.

2. Product Features

2.1. High Level Summary

Darktrace / IDENTITY modules retrieve activity data from third-party SaaS and enterprise software platforms.

Activity data is processed, parsed, and analyzed to identify users or other entities which are active within the third-party platform. Each user or entity identified is then evaluated for behavior outside of a "normal" state, created by an ongoing, real-time analysis of audited activity and other supplied data events across the enterprise network environment. The system profiles individual users or entities based upon their activity, and the activity of those it deems "peers" due to similar behavioral activity ("modelling").

Alerts are raised when specific criteria, or a minimum threshold of unusual behavior is met. The output is surfaced to operators in the Darktrace / IDENTITY Console and Darktrace Threat Visualizer interface for investigation and resolution. Some Darktrace / IDENTITY modules can also respond to unusual activity ("Autonomous Response") by modifying the users, entities, or relevant configuration settings within the third-party platform.

2.1.1.Data Retrieval and Ingestion

Darktrace / IDENTITY modules retrieve log-based activity data from third-party platforms; this is typically achieved via APIs or export formats provided by the third-party platform. Due to the variation in APIs, authentication methods, and SDKs, provided by each separate platform, the methodology for ingestion will differ between Darktrace / IDENTITY modules. The method of authentication and retrieval is outlined in the relevant documentation provided for each module.

Delays may be incurred where the external platform does not make events available to the Darktrace module for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the associated third-party platform vendor. Latency between event occurrence and when it was made available to the module are indicated in the event metadata within the Threat Visualizer.

A subset of modules will also retrieve contextual enrichment data on users, files, and other relevant entities. The data available and the method by which it is retrieved are outlined in the relevant documentation provided for each module.

Darktrace offers a suite of Threat Intelligence and Telemetry integrations where data may be retrieved from other compatible security tools or OSINT sources. The retrieval method for each integration is detailed in the corresponding integration documentation. These alternative inputs (“Threat Intelligence Integrations” & “Telemetry Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page.

Darktrace also provides a REST API for automation and a subset of relevant data input.

2.1.1.1. “Open” Ingestion

Data ingested by Darktrace / IDENTITY modules is retrieved directly from the audited activity logs of the relevant third-party platforms - returned information is therefore limited to the events that each vendor chooses to audit, and the data recorded as part of those audit log entries. Variation in both the events available to each Darktrace / IDENTITY module, and the level of detail, are expected.

Darktrace / IDENTITY modules operate in a passive, “open” ingestion mode - modules do not actively request specific events or place restrictions on the data retrieved from the third-party platform. The range of events available to the module will therefore vary both between each monitored platform and between each client.

Some third-party platforms may also restrict the events and APIs available behind license requirements. Darktrace will outline where possible the minimum license requirements required for basic real-time monitoring operation.

Exceptions exist where the resource types to be monitored are pre-defined during configuration (e.g. Salesforce, Google Workspace), or where individual event types must be removed due to excessive volume.

2.2. Darktrace Analysis

Darktrace analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace platform, Threat Visualizer interface, or any underlying components.

Where source data is low quality, represents only a subset of all activity, is delayed by external factors, or otherwise is incomplete, the quality of Darktrace analysis may be severely impacted. Analysis may also be impacted by overloading, such as that observed when event throughput significantly increases beyond scoped levels.

2.2.1.Darktrace / IDENTITY Event Inspection

Darktrace / IDENTITY modules will retrieve activity log information made available by third-party APIs (or other comparable methods) for analysis. Event and activity data is then parsed into a series of standard metadata fields and analyzed to identify the type of activity, any entities involved (e.g., users, files, or other resources), and to extract key information.

The output of this analysis is passed to the Darktrace model engine, to the core Darktrace "classifier" engine, and made available for display in the user interface and any other platform components such as the AI Analyst.

2.2.2.Darktrace 'Pattern of Life' Analysis and Classification

Network events, user activity events, connection data and any other configured inputs are subject to Darktrace 'pattern of life' analysis. Darktrace will create individualized behavioral profiles for the network entities it observes and surface activity which is considered inconsistent with the expected norm. This baseline is derived from - but not limited to - an analysis of the behavior of the individual entity, analysis of one or more clusters created from similarly behaving entities, and many variable factors such as time or communication protocol. The 'pattern of life' data is continually updated in real time, and reflects the data that has been received, with a greater weighting to more recent data.

Analysis is performed in the "classifier" stage by a multitude of classifiers. This core analysis applies many approaches including unsupervised Machine Learning techniques such as Bayesian Meta-classification, techniques derived from graph theory and network analysis such as Node/Graph Centrality, approaches derived from statistical analysis such as Spectral Clustering and many other techniques. The previous examples of utilized techniques provided are for illustrative purposes and should not be considered exhaustive.

Darktrace does not offer any capability to access underlying behavioral models or classification output. Darktrace provides the ability to utilize the output of this analysis using the *models* framework. Darktrace will create alerts to indicate anomalous activity which will be inserted into the "event log" within the Darktrace Threat Visualizer interface of the corresponding device, user, or entity. These unusual activity "notices" are also consumed by Darktrace Real-Time Detection models.

2.2.3.Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of 'pattern of life' analysis is evaluated. Both analysis components described above submit data to the model engine for evaluation.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior. The models framework leverages both the underlying 'pattern of life' detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace / IDENTITY and Darktrace / CLOUD modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet any criteria; these are created and maintained in the Threat Visualizer "Model Editor" interface. Care should be taken when defining custom models to ensure that the number of alerts produced is not excessive and does not impact the system's ability to function, or otherwise degrade service. Similarly, modification of existing default models should not result in excessive activity or alter the logic to the extent that alerts criteria can no longer be met. The responsibility to ensure that models created or edited remain within these reasonable expected boundaries lies with the Customer.

Default Darktrace models are focused on ‘pattern of life’ anomaly detection, potentially malicious behavior and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically, clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

Models may also be used to trigger actions within the Darktrace Threat Visualizer platform; the output of the Darktrace Real-Time Detection model engine is described below.

2.2.3.1. Autonomous Response Models

In default operation, Darktrace / IDENTITY Autonomous Response responses are triggered by model alerts from a specific subset of Darktrace models, categorized as Autonomous Response (“Antigena”) models.

Darktrace Autonomous Response models may directly look for specific behavior, or for indicators identified by other models operating within the real-time Detection environment. Darktrace / IDENTITY models fall typically into the second category - “meta-models” - which are triggered by an alert of another Darktrace Real-Time Detection model.

Any modification to an underlying Real-Time Detection model which results in increased or reduced model alerts will subsequently impact Darktrace / IDENTITY Autonomous Response models and actions. Any under-activity or over-activity of the corresponding Darktrace / IDENTITY Autonomous Response model as a result of this type of modification is the responsibility of the Customer.

Darktrace also provides the ability to limit actions created by these models to a minimum score threshold.

2.2.4.AI Analyst

Darktrace AI Analyst performs a meta-analysis upon the previous layers of analysis described. Please refer to “AI Analyst” below.

2.2.5.Cross-Capability / Cross-Coverage Area Analysis

Where other Darktrace ActiveAI Security Platform coverage areas are deployed, cross-platform analysis is performed by multiple components. Links are created for analysis purposes between network entities and entities modeled by other Darktrace Real-Time Detection components such as, for example, user entities created by Darktrace / IDENTITY module monitoring known to be associated with a given network entity, or credential entries observed by Darktrace / NETWORK monitoring also observed as part of Darktrace / ENDPOINT.

Darktrace AI Analyst will also link together entities across different coverage areas in the creation of AI Analyst Incidents and AI Analyst Incidents Events. Darktrace AI Analyst may also retrieve additional contextual data during investigation from other components if deployed, such as the retrieval of associated emails from Darktrace/Email. Please refer to “AI Analyst” below.

The above provides illustrative examples of collaborative scenarios between platform components for reference but is non-exhaustive.

2.2.6. “User” Device Entities

As part of the analysis process, Darktrace will model individualized behavioral profiles in the form of “user” device entities.

Each “user” device represents a distinct, unique actor that has performed a *meaningful* action in the third-party platform. This is defined as any successful login, or any “outbound” event, where outbound indicates the user or entity was the actor who performed the action. This is distinct from “inbound” actions, where the user or entity was observed within the context of another user modifying them in some way but was never seen performing any actions themselves.

Examples of outbound actions may include a user accessing files, creating API keys, or modifying a resource in some way. In this definition, a “resource” is any modifiable element, including but not limited to a group, a file, a virtualized computing resource, a password policy, or another user.

Failed login events alone will not create user entities but will be added to user entities if created already by a meaningful action.

Where a user or entity is only observed performing an event which does not meet the criteria for the creation of a “user” device - and no prior user device exists for them - the event will be aggregated into a generic log for the Darktrace / IDENTITY module. Behavioral analysis and Darktrace Real-Time Detection will be severely limited for this aggregate entity.

Many third-party platforms permit external entities to influence resources. Similarly, many third-party platforms and systems operate “service” accounts or system-owned entities that may alter internal resources. Therefore, Darktrace creates no expectation that the count of user device entities created from activity monitoring will directly align with license or seat counts negotiated with the third-party platform.

Darktrace provides mechanisms to limit the scope of entities modeled in this way.

2.2.6.1. Restricting “User” Device Entities

Organizations may wish to restrict visibility over the event data retrieved and modelled by Darktrace / IDENTITY modules due to regional monitoring restrictions, multi-tenant environments, or other compliance policies. Darktrace offers both user-level and ingestion-level restrictions to limit the processing of this data. Typically, the filtering of data is performed *after it is retrieved* from the third-party platform due to the open ingestion model outlined above.

The “Activity Filter” and “Group Activity Filter” (*limited modules only*) settings restrict the creation and modelling of users and groups respectively to those that meet or do not meet a specified criterion. Users which take actions that impact explicitly monitored users, or are directly actioned by monitored users, will still be modeled if they do not match the criteria to ensure that the two-way interaction is captured. This is the recommended operating mode to ensure full visibility of relevant information for operators.

If it is necessary to ensure that information about these unmonitored users is not preserved, Darktrace provides mechanisms to collapse all users outside the scope of monitoring to a single entity. Behavioral analysis and Darktrace Real-Time Detection will be severely limited for this aggregate entity.

2.3. Output

Darktrace / IDENTITY modules can possess both Darktrace Real-Time Detection and, where applicable and licensed, Darktrace Autonomous Response capabilities. The output of these two capabilities are outlined individually.

2.3.1.Real-Time Detection

Darktrace / IDENTITY Real-Time Detection outputs activity-based metadata for users and entities, formatted into consistent categories and with the most relevant information pre-extracted. As noted above (refer to *1.2.1.1 "Open" Ingestion*), each module will retrieve all possible relevant event activity from the third-party platform - the events retrieved and surfaced will therefore differ between monitored platforms and between client deployments.

Darktrace / IDENTITY Real-Time Detection will typically surface activity data, results of ongoing 'pattern of life' analysis (including notification of unusual activity), model events, Darktrace AI Analyst alerts, contextual data retrieved from configured integrations (optional), and any other log-based event data created or processed by the Darktrace system. The platform may also output alerts about overall system health and component health.

Deploying a Darktrace / IDENTITY module will provide access to the Darktrace / IDENTITY Console, a specialized interface for investigating SaaS and Cloud activity. The Console provides access to global maps of activity, detailed logs of user activity and visualization of anomalous chains of behavior. Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, activity processed is also surfaced in the main Darktrace Threat Visualizer interface. Data is optionally available through the Darktrace Mobile App, compatible output formats, or the Darktrace Threat Visualizer REST API.

2.3.1.1. Processed Event Metadata

The output of Darktrace / IDENTITY Real-Time Detection is displayed in the Darktrace Threat Visualizer "Advanced Search" interface. Data retention for this output is on a rolling basis and is dependent upon a number of factors such as hardware capability, event volume and other operational components which contribute entries to Advanced Search. Most deployments can expect around 30 days retention.

2.3.1.2. Alerts

Darktrace / IDENTITY Real-Time Detection produces two primary alert types: Darktrace Model Alert (*formerly known as "Model Breach" alerts*) and Darktrace AI Analyst Incidents. Darktrace Model Alerts and Darktrace AI Analyst Incidents can be investigated in the Darktrace / IDENTITY Console, the main Threat Visualizer interface, investigated in the Darktrace Mobile App, or exported to a compatible alert consumer. Alerts are categorized for priority and filterable by multiple factors, allowing for customization of alerts displayed in Darktrace interfaces and those exported to external tools.

A secondary alert type - Darktrace System Status alerts - may also be generated where the Darktrace platform is experiencing degraded service due to health issues, invalid input, or where it is necessary to highlight changes to system administrators. Alerts include details of the originating host, the severity of the event, and relevant links to investigate or resolve the issue. Notifications are sent when a system event becomes active and can optionally be sent on resolution.

Darktrace supports alert export in both industry-standard generic forms such as Syslog or Email and custom integrations with common tools such as Splunk and ServiceNow. Alert outputs ("Workflow Integrations") are configured on the Darktrace Threat Visualizer "System Config" page. The details included in each external output may vary due to third-party restrictions on content length or supported formats.

2.3.1.2.1. Model Alerts (formerly “Model Breach” alerts)

Darktrace model alerts are created as a result of Darktrace / IDENTITY Real-Time Detection models; when conditions for a model are met, a model alert can be created in addition to other possible model actions. Darktrace model alerts contain details of the conditions that were satisfied in order to trigger the alert, the entity which met those conditions (for example, a device, or another model alert) and a description with recommended action points.

Model alerts will only be created for users who are included in monitoring.

2.3.1.2.2. Darktrace AI Analyst Incident Alerts

Darktrace AI Analyst Incident alerts are created when Darktrace AI Analyst identifies activity considered significant enough to highlight to operators. Darktrace AI Analyst Incidents may contain a single “AI Analyst Incident Event”, or multiple linked findings aggregated together. External alerts are created when a new “AI Analyst Incident Event” is created, which may be associated with an existing “AI Analyst Incident”, or form an independent, new “AI Analyst Incident”. In the former case, mechanisms to identify the relationship between AI Analyst Incident Events are provided in the output.

In the Darktrace / IDENTITY Console, Darktrace Threat Visualizer and Darktrace Mobile App, the output of Darktrace AI Analyst analysis is aggregated on a per-AI Analyst Incident basis. New AI Analyst Incident Events added to existing Incidents will not produce independent alerts in these interfaces and are instead displayed as part of the existing Incident. Darktrace AI Analyst Incidents in these interfaces contains details of the activity identified, the Darktrace model alerts that triggered the initial AI Analyst investigation, the entities which performed the unusual activity, (for example, devices), the investigation steps AI Analyst performed, why activity was linked together (if multiple AI Analyst Incident Events), and a human-readable summary of the finding.

2.3.2. Autonomous Response

Darktrace / IDENTITY Autonomous Response actions can be taken when Darktrace / IDENTITY Real-Time Detection identifies significantly anomalous user behavior, when Darktrace identifies activity which contravenes a compliance policy, when manually triggered by an operator, or as a result of any other custom criteria defined in a Darktrace model. Darktrace / IDENTITY Autonomous Response is only available for a specific subset of supported Darktrace / IDENTITY modules.

If the current system configuration permits autonomous actions - or the action was triggered manually - the action will be created in *active* state. This created action will appear on the Darktrace / IDENTITY Console Response Actions Page under “Active” actions, in the Darktrace Threat Visualizer Response Actions Page “Platform” tab, and in the Darktrace Mobile App. If alerts for Autonomous Response actions (“*created*” state) are configured, an alert will also be sent to configured outputs.

If Darktrace / IDENTITY Autonomous Response is compelled to request human confirmation due to the current system configuration, the action will be created in *pending* mode. This created action will appear on the Darktrace / IDENTITY Console Response Actions Page (and locations outlined above) under “Pending” action until a user confirms that Darktrace / IDENTITY can take action. A notification requesting that a human user approve the action is sent to the Darktrace Mobile App and any configured Darktrace Autonomous Response alert outputs (“*pending*” state). Once confirmed by an operator, the action will move to “Active” as before.

Read-only access to view Darktrace / IDENTITY Autonomous Response actions is not restricted for users with appropriate permissions to view Darktrace / IDENTITY Real-Time Detection data.

Darktrace / IDENTITY operates an “opt-out” approach to Autonomous Response eligibility. Please refer to “1.2.8 Autonomous Response Eligibility” below for further information.

2.3.2.1. Pending vs Active Actions

Darktrace / IDENTITY Autonomous Response actions created as a result of Darktrace Real-Time Detections are initialized in two possible states: *active* and *pending*.

An *active* Darktrace / IDENTITY Autonomous Response action permits Darktrace / IDENTITY to take the responses defined by the action configuration. For example, to log a user out or block a specific IP via the third-party platform API or supported integration method. As manual Darktrace / IDENTITY response actions are created directly by an operator, actions are always created in the *active* state.

Pending actions are those which Darktrace / IDENTITY Autonomous Response wishes to take but must require approval before it can act automatically; pending actions may therefore be referred to as “human confirmation mode” actions. Pending actions do not cause modifications in the third-party platform until approval is granted in some format, at which the point the action becomes “active”.

Darktrace provides mechanisms to alert users on the creation of pending actions requiring human approval to proceed. These notifications are surfaced in the main Threat Visualizer interface (*excludes Darktrace / IDENTITY Console*), in the Darktrace Mobile App, and in compatible alert outputs.

Please refer to “1.2.9 Darktrace / IDENTITY Response Autonomy” for more information about these system conditions

The alternative terminology “Human Confirmation” or “Human Confirmation Mode” may also be used, respectively, when referring to pending actions or a system state where pending actions are created.

2.3.2.2. Action States

Darktrace Autonomous Response actions proceed through several states during their lifecycle. For example, an action is automatically created that *requires confirmation* (“pending”), it is then *confirmed* by a user. Another user *extended* the action, then *cleared* it. Finally, it was *reactivated* by another user, before *expiring*.

There are four key states (“pending”, “active”, “cleared”, “expired”) and two additional states (“extended”, “reactivated”):

- The first two states - pending and active are outlined above
- Cleared actions are those which have been manually ended by an operator. Clearing informs Darktrace to cease the action, and to suppress the combination of Darktrace / IDENTITY Autonomous Response action and model breach conditions for the remainder of the action’s set duration.
- Expired actions are those which are historic, regardless of their state (pending, active) before the action period passed.
- Reactivated actions are those which were made active again by an operator.
- Extended actions are those which are currently active and have had their duration manually extended by an operator.

Actions can be extended, reactivated or cleared from the Response Actions page and the Darktrace Mobile App.

2.3.2.3. Actions in Default Operating Mode

Default operating mode (*default operation*) is here defined as Darktrace / IDENTITY Autonomous Response operating in a state where autonomous (*active*) actions can be taken automatically against eligible users on supported platforms, when triggered by a default Darktrace / IDENTITY Autonomous Response model. Autonomous actions are actions active without human intervention.

Default operation presumes that Darktrace / IDENTITY Autonomous Response is enabled globally on the Darktrace ActiveAI Security Platform instance, licensed for the appropriate module, and the module has been granted the appropriate permissions in the third-party platform to enact actions. In some cases, it may be necessary to re-authenticate the module to grant the necessary permissions. Please refer to the corresponding documentation for each module for further guidance.

The conditions for the creation of action in any state in default operation are:

1. A Darktrace / IDENTITY Autonomous Response model which has its logical criteria sufficiently met.
2. The prior criteria met by a user or IP Address which has not been explicitly excluded from actions.
3. A Darktrace / IDENTITY Autonomous Response action (inhibitor) configured within the triggered model to define the action taken, where:
   - The prior criteria are not met by an entity which is “immune” to the inhibitor in the module configuration.
4. A system operating mode (schedule) which does not prevent the creation of a Darktrace Autonomous Response action at the given time.

The conditions for an *active* action in default operation are:

5. A Darktrace / IDENTITY Autonomous Response model which has its logical criteria sufficiently met.
6. The prior criteria met by a user or IP Address which has not been explicitly excluded from actions.
7. A Darktrace / IDENTITY Autonomous Response action (inhibitor) configured within the triggered model to define the action taken, where:
   - The prior criteria are not met by an entity which is “immune” to the inhibitor in the module configuration.
8. The model in question is configured to *permit* or *force* autonomous actions.
9. The system operating mode permits the creation of an active Autonomous Response action at the given time for the given device.

Manual actions triggered by an operator are not subject to the conditions outlined above; these actions are always active on creation.

2.3.3.Darktrace / IDENTITY Actions

Darktrace / IDENTITY response actions can be triggered by Darktrace / IDENTITY Autonomous Response models, by custom models possessing the required configuration, or manually by an operator.

The actions Darktrace / IDENTITY takes in response to a trigger will vary based upon the action configuration and the third-party platform the trigger user has been observed upon. The types of response actions taken are henceforth referred to as “inhibitors”, as they *inhibit* a specific form of behavior or activity.

The range of inhibitors available differ between Darktrace / IDENTITY modules.

2.3.3.1. Manual Darktrace / IDENTITY Response Actions

Darktrace provides mechanisms to create Darktrace / IDENTITY response actions manually. These actions can be created by a user with appropriate permissions from the Darktrace / IDENTITY Console, the main Threat Visualizer interface, or from the Darktrace Mobile App.

Darktrace provides the option to enforce a free-text justification which must submitted by the end-user when the action is created.

Manual actions are automatically created in an "active" state and do not require human activation. The type of action taken (*inhibitor*) is selected during creation. The responsibility to select an appropriate inhibitor lies with the Customer.

Where a user or targeted IP Address is excluded from the specific action inhibitor, or from all actions, triggered manually by an operator, this will be indicated in the action status. In this case, an active action will be created but the inhibitor will be prevented from taking effect.

2.3.3.2. Model-Triggered Darktrace / IDENTITY Autonomous Response Actions

In default operating mode, Darktrace / IDENTITY Autonomous Response actions are triggered by models. This trigger is referred to as a "model action" - a system action taken in response to a specific model criterion being met.

The creation of a Darktrace Autonomous Response action is a model action which exists by default on applicable Darktrace-created models created for the purposes of Autonomous Response. When the criteria for these models are met, Darktrace will invoke all model actions applied to the model, triggering a Darktrace / IDENTITY Autonomous Response action as a result. In this context, "trigger all model actions" refers to all possible model actions applied to the model (for example, "tag device" or "create alert"). It does not refer to triggering all *inhibitors* applied to the model.

Operators may also add the Darktrace Autonomous Response model action to modified or newly created custom models. The responsibility for selecting an appropriate action inhibitor (see below), and for any actions created as a result of a custom model configuration, lies with the Customer.

2.3.3.3. Darktrace / IDENTITY Autonomous Response Inhibitors

The type of action Darktrace / IDENTITY Autonomous Response takes is referred to as an inhibitor. The suite of available inhibitors differs between each platform monitored by Darktrace / IDENTITY - models therefore accept multiple Darktrace / IDENTITY inhibitors to ensure an appropriate action exists across all relevant platforms. The supported inhibitors for each platform can be found in the relevant Darktrace documentation.

In default operation, when the criteria for a Darktrace / IDENTITY Autonomous Response model are met, Darktrace / IDENTITY will attempt to create an action corresponding to the selected inhibitor for the combination of trigger user and platform. The inhibitor is defined in the model directly (see "1.2.3.4.2 Model-Triggered Darktrace / IDENTITY Autonomous Response Actions").

Manual actions require the inhibitor to be selected during creation. This is applicable to actions created through the Darktrace / IDENTITY Console, main Threat Visualizer, and the Darktrace Mobile App.

For actions created as a result of a user-created or user-modified model, Darktrace / IDENTITY Autonomous Response will attempt to create an action corresponding to the inhibitor as defined in the model (see "1.2.3.4.2 Model-Triggered Darktrace / IDENTITY Autonomous Response Actions") for the trigger user and platform.

Darktrace / IDENTITY allows individual inhibitors to be disabled (*specific modules only*). Users and IP addresses can also be made "immune" on a per-inhibitor basis.

2.3.3.3.1. Modification and Restriction of Inhibitors

Typically, Darktrace / IDENTITY Autonomous Response will take the action ("inhibitor") selected in the model action in response to the Darktrace / IDENTITY Autonomous Response model criteria being met. However, there are exceptions to this behavior:

- Where the user who met the model criteria was observed on a platform which does not support Autonomous Response. In this case, no actions are created.
- Where Darktrace / IDENTITY possesses some of the necessary permissions to take actions but is unable to target all user types, the action may be downgraded to a supported capability (*relevant modules only*).

  For example, if Darktrace / IDENTITY does not have the appropriate user role to target Microsoft 365 "Global Admin" users with a "disable user" inhibitor, it will attempt to downgrade to a "force logout" action. This will be noted in the action status on the Response Actions page.

  Downgrading can be disabled in the relevant module configuration.
- Where the targeted user is configured to be "immune" from the specific inhibitor, or globally from actions. In this case, an active action will be created but the inhibitor will be prevented from taking effect. This will be noted in the action status on the Response Actions page.
- Where the targeted IP Address is configured to be "immune" from the specific inhibitor, or globally from actions. In this case, an active action will be created but the inhibitor will be prevented from taking effect. This will be noted in the action status on the Response Actions page.
- Where the targeted entity is of an unsupported type (as outlined in the relevant documentation). Darktrace / IDENTITY may attempt to take action - an active action will be created - but be prevented from succeeding. This will be noted in the action status on the Response Actions page.
- Where the triggered inhibitor is "Block IP", and the model name meets the criteria of the "Block IP without Force Logout Models" field in the module configuration. In this case, "Block IP" will be applied but only in a partial format.

2.3.4.Autonomous Response Actions Initiated by Cyber AI Analyst

Cyber AI Analyst can now initiate action creation - this functionality is optional and disabled by default.

Actions are created as a result of Cyber AI Analyst Incidents. AI Analyst links related devices and unusual activity together to create an overall incident structure. If AI Analyst detects that a link between two entities does not have a corresponding Autonomous Response action - and an inhibitor is available which would suitably target the activity exists - it can now create an action to target this behavior.

2.3.5.Autonomous Response Eligibility

Darktrace / IDENTITY operates an "opt-out" approach to Autonomous Response eligibility. All user entities detected by Darktrace / IDENTITY modules with Autonomous Response capabilities are considered eligible for Autonomous Response actions.

Eligibility is highly configurable - both users and IP addresses (relevant to the “Block IP” inhibitor) can be made *immune* from actions on a global-level, module-level or inhibitor-level.

Where the Darktrace / IDENTITY module is active on a Unified View instance (*specific modules only, advanced operating mode*), immunity must be configured on the corresponding subordinate master where the user entity is modelled.

2.3.6.Darktrace / IDENTITY Response Autonomy

Darktrace / IDENTITY Autonomous Response can create actions in an “active” or “pending” state, were pending actions require human approval before any direct action is actually taken. The state actions are created in - referred to here as “autonomy” - is defined by the following factors and conditions.

2.3.6.1. Manually Triggered Darktrace / IDENTITY Response Actions

Manual actions are always automatically created in an “active” state.

2.3.6.2. Model-Triggered Darktrace / IDENTITY Autonomous Response Actions

The status of actions created as a result of a Darktrace model alerts (both default Darktrace / IDENTITY Autonomous Response models, and those user-created or user-modified) is defined by a combination of granular eligibility controls.

2.3.6.2.1. Model-Based Autonomy

Darktrace / IDENTITY Autonomous Response autonomy can be configured on a per-model basis to meet the varied requirements of each organization. This autonomy is controlled in the model configuration: models may “force human confirmation”, “force autonomous action”, or “permit autonomous action”. Actions created by models with a “force” state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action created as a result of the model with “permit autonomous action” will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

2.3.6.2.2. Time-Based Autonomy

Darktrace also provides a seven-day, hourly timetable (“Response Schedule”) which allows blocks of autonomous action or enforced human oversight to be scheduled. The schedule is applied across Darktrace / NETWORK, Darktrace / ENDPOINT and Darktrace / IDENTITY.

The schedule is applied to actions created by models with “permit autonomous action”. Actions created manually, or by models in a “force” state (“force human confirmation”, “force autonomous action”), are not impacted by the schedule state.

Localization is not available for Darktrace / IDENTITY module responses. If “local subnet time” is active, the schedule will apply Darktrace / IDENTITY actions in UTC.

2.3.6.2.3. *Cyber AI Analyst-Triggered Actions*

Actions created as a result of a Darktrace incident can “force human confirmation”, “force autonomous action”, or “permit autonomous action”. This setting is set on the Darktrace System Config page and applies to all actions created by Darktrace Cyber AI Analyst.

Actions created when this setting is in a “force” state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action when set “permit autonomous action” will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

2.4. Interface

The primary user interface of the Darktrace platform is the Threat Visualizer. The Threat Visualizer interface provides access to all major Darktrace user interfaces, consoles and product views - it contains both investigation interfaces and administration interfaces. A list of the user interfaces that comprise the Threat Visualizer - and detailed information on how to operate these interfaces - is provided in the relevant technical documentation.

The Darktrace / IDENTITY Console is a purpose-built investigation interface for Darktrace / IDENTITY that forms part of the Threat Visualizer.

- If Darktrace / IDENTITY is the only Darktrace coverage area, only the Darktrace / IDENTITY Console and a subset of relevant Threat Visualizer pages will be accessible.
- Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, the Darktrace / IDENTITY Console will available alongside the main Threat Visualizer interface. Darktrace / IDENTITY data will be surfaced in both views.

Both the Darktrace / IDENTITY Console and main Threat Visualizer interface include user customizable and filterable alert trays, a visualization of ongoing device activity over time, and access to the output of Darktrace analysis in both log and summary format. Alternative investigation views, access to review detailed metadata, report generation, and the ability to edit and create custom models (as described above) are also provided within these interfaces. Operators can also review the detected users observed by Darktrace / IDENTITY Real-Time Detection (including any relevant contextual data), configure system settings, deploy integrations, review system health information, perform user management and other administrative tasks.

Darktrace / IDENTITY Autonomous Response actions are recorded in the Response Actions window within the Darktrace / IDENTITY Console. This view restricts all actions to only those taken by Darktrace / IDENTITY modules. A per-user filtered view is also offered from each user profile.

Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, response actions for Darktrace / IDENTITY are also surfaced in the main Darktrace Threat Visualizer Response Actions window. Here, Darktrace / IDENTITY actions are located under “Platform”.

Users with appropriate permissions may alter the state of an action from this view (clear, activate, extend, or reactivate). Darktrace / IDENTITY Autonomous Response actions are also surfaced in, and can be modified from, the Darktrace Mobile App.

Darktrace offers alerting when Darktrace / IDENTITY response actions are created or change state; notifications are surfaced in the Threat Visualizer interface, in the Darktrace Mobile App, and sent via compatible alert outputs.

Darktrace / IDENTITY modules display a current “state” indicating whether actions could be successfully applied, whether the user was immune in some format, or whether the action has subsequently expired.

2.5. Cyber AI Analyst

Darktrace models are used as a trigger to invoke Cyber AI Analyst. When the conditions for a model are met, a model alert is created; Cyber AI Analyst reviews and investigates all relevant model alerts that occur on the system as a starting point for its analysis process. The output from this analysis process is Cyber AI Analyst Incidents - a collection of one or more related events of anomalous activity. Incidents are formed through a meta-analysis of activity type, entity type (such as devices, identities), and endpoints involved in each event. Each incident can encompass multiple stages of activity as it develops.

The Darktrace Cyber AI Analyst operates a hypothesis-based analysis approach, where activity is evaluated against a number of possible, relevant hypotheses and a determination is taken of which (if any) hold based upon the evidence gathered and investigations performed. This investigation process involves numerous forms of data analysis including, but not limited to, AI and Machine Learning algorithmic approaches, statistical analysis techniques, and other forms of natural language and mathematical analysis.

The Cyber AI Analyst will combine activity across different Darktrace Real-Time Detection coverage areas where possible.

Although a model alert may be the trigger for an investigation, that does not mean the activity Cyber AI Analyst surfaces is directly related to the original model alert. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model alert but are worthy of investigation. Similarly, very few model alerts that are investigated will result in an incident - only activity the Cyber AI Analyst considers high priority. Whether Cyber AI Analyst has created a related incident is displayed in appropriate locations within the Darktrace Threat Visualizer and Darktrace Mobile App.

Users may manually trigger Cyber AI Analyst investigations into devices of interest or trigger Cyber AI Analyst investigations through third-party telemetry inputs.

2.5.1.Cyber AI Analyst Hypotheses

As of Darktrace 6.2, users may alter the hypotheses investigated by Darktrace Cyber AI Analyst as a result of a model alert or add hypotheses for investigation to custom models.

Darktrace strongly recommends that operators do not alter the hypotheses associated with Darktrace-maintained models; any activity which is not investigated due to a modification of this type is the responsibility of the Customer.

Darktrace also strongly recommends limiting the number of custom models Cyber AI Analyst is invoked by; a significant number of additional investigations will create additional load on the Cyber AI Analyst engine, potentially resulting in failure to investigate key activity.

2.6. Reporting

Darktrace offers both manual and automatic PDF report generation, scheduled via the Darktrace Threat Visualizer “System Config” page. Details of reporting formats offered is provided on the Darktrace Customer Portal and relevant documentation. Darktrace reserves the right to alter the content of reports offered to align with changing Product and Service Offering.

The Threat Visualizer “Audit Log” records changes made by operators such as model alert acknowledgement; the audit log can be exported via Syslog for extended retention. User interactions with

Darktrace / IDENTITY Autonomous Response actions which alter the state (e.g. “activate”, “extend”, “clear”) are also recorded in the action history, accessible from the Darktrace Response Actions window. Users can also be compelled to provide a free-text justification when an action state is altered, which is also displayed where applicable.

The output of Darktrace / IDENTITY data retrieval and analysis is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Combined output from this processing, from Darktrace analysis, and from any actions performed by the platform automatically (such as tagging as a result of a model) are combined into logs which are displayed for each user.

Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, event volume, and other operational components which contribute data to the platform. Most deployments can expect around 30 days retention of Darktrace / IDENTITY metadata and general device/user activity data. Automatic removal of older event log data does not affect the storage or training of the machine learning ‘pattern of life’ data.

User entities observed by Darktrace / IDENTITY Real-Time Detection to be recently active in monitored third-party platforms are detailed on the Darktrace / IDENTITY Console “Profiles” page. Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, user entities will also be reported on the “Device Admin” page. Metrics regarding data and event throughput are rendered on the Threat Visualizer “System Status” page.

Darktrace also provides a REST API for automated retrieval of a subset of data output.

3. Deployment Architecture

Darktrace / IDENTITY modules can operate on Darktrace physical or cloud-hosted master instances, or on Unified View instances in select cases. Internet connectivity to the relevant endpoints - both those associated with the third-party platform, and with relevant Darktrace services - is necessary for data retrieval. Darktrace / IDENTITY modules only operate on physical probe instances, vSensors, osSensors, or other topology components referenced in the Darktrace / NETWORK specification.

The relevant combination of Darktrace / IDENITITY modules will vary depending on the customer technical stack and what they currently have in their environment. The list of currently available Darktrace / IDENTITY modules are as follows:

- Asana
- Box
- CloudFare
- DropBox
- Duo
- Egnyte
- Google Workspace
- HubSpot
- JumpCloud
- Microsoft 365
- Okta
- Salesforce
- Slack
- Zoom

The applicable Darktrace / IDENTITY modules may vary during the lifetime of a Darktrace deployment should new tools be introduced to the Customer’s architecture. It is the responsibility of the client to ensure that Darktrace maintains visibility over all traffic in the event of additions that will fall outside the scope of the initial deployment design.

Each deployment requires a minimum of one “*Master*” instance to provide the capabilities outlined above. Deployments should be structured so that a single instance - Master or Unified View - sits at the top of the deployment topology to operate the Threat Visualizer interface and other relevant components.

For further details on the relevant Topology Components and Topology Roles, please refer to the / NETWORK Product Specifcation.

4. Implementation

The Customer Portal contains technical guides for implementing the various Darktrace / IDENTITY modules. Customers should follow the best practices set out in these guides when implementing applicable modules.

4.1. License Keys

Darktrace / IDENTITY will only be available when a valid License Key is configured on the Threat Visualizer's System Config page. In Unified View environments, the License Key should propagate from the UV master instance to the subordinate Master instances unless specifically configured otherwise. Once the license is entered, Customer will get access to all available Darktrace / IDENTITY modules. Adding the License Key will also activate Autonomous Response for the modules that have this feature available. Instructions on adding a License Key can be found in the relevant documents in the Customer Portal.

4.2. Deploying Modules

Deploying one or more Darktrace/ IDENTITY modules will provide access to the SaaS Console, a specialized interface for investigating SaaS and Cloud activity. The console is powered by the Cyber AI Analyst and Darktrace's unique 'pattern of life' anomaly detection; each element is purpose built for monitoring and analysis in these environments whilst maintaining existing workflows for operators already familiar with the Darktrace Threat Visualizer. The SaaS console provides access to global maps of activity, detailed logs of user activity and visualization of anomalous chains of behavior.

Each module comes with its own configuration considerations and the implementation process required to ensure operationality will vary accordingly. Processes that vary across modules include, but are not limited to: means of data retrieval; permissions requirements; limitation considerations; and license requirements. The specifics configuration requirements of each Darktrace / IDENTITY module and detailed instructions for each setup can be found in the relevant technical documentation in the Customer Portal. Customers should consult this documentation ahead of attempting deployment.

4.2.1.Considerations

Third-party platforms may limit the events in the applicable modules that Darktrace is able to access. Each module has its unique set of considerations that should be reviewed before deploying the module. Delays may be incurred where the external platform does not make events available to the Darktrace module for processing and analysis within the expected timeframe. It is a Customer responsibility to identify and minimize latency between the relevant Darktrace / IDENTITY module and a given third-party platform, Darktrace does not bear responsibility for latency resulting from third-party platforms. Latency between event occurrence and when it was made available to the module are indicated in the event metadata within the Threat Visualizer.

4.2.2.Permissions

Each Darktrace / IDENTITY Module has its unique set of permissions that are required to deploy the module. The permissions vary by third-party platform and details about the permissions needed can be found in the relevant technical documentation in Customer Portal. Different permissions apply in relation to data retrieval than apply in relation to taking a response action.

4.2.3.Deploying Autonomous Mode

Darktrace / IDENTITY modules must be reauthorized to enable Darktrace Autonomous Response functionality after a Darktrace Autonomous Response license key is added, regardless of whether this permission was present at initial authorization. Additional permissions are required in order for Darktrace to take actions within a given third-party platform. The permissions vary per vendor as outlined above, and the information needed on the Darktrace end can be found in the relevant technical documentation in Customer Portal.

Whether Darktrace / IDENTITY can take action autonomously, or must wait for human approval, is defined by a schedule or model setting. The seven-day, hourly timetable allows blocks of autonomous action or enforced human oversight to be scheduled. For some high severity activity types, it is recommended to ensure Darktrace / IDENTITY Autonomous Response can always take an autonomous action regardless of the scheduled state; in this case, an override can be configured in the model setting.

The Customer is responsible for configuring Darktrace / IDENTITY Autonomous Response to run fully autonomously. Fully autonomous mode is the end goal of all Darktrace / IDENTITY Autonomous Response deployments; reaching a fully autonomous state where the AI can take action whenever unusual or concerning behavior is detected, without the need for human oversight. This mode lends itself to a minimal-interaction workflow, where the Customer may infrequently need to modify actions through the Darktrace Threat Visualizer interface, API, or Darktrace Mobile App, but on the whole can leave Darktrace / IDENTITY Autonomous Response to operate with little intervention.

4.2.4.License Requirements

Some third-party platforms require special licenses to allow events to be shared with Darktrace. For best results the license requirements should be reviewed before implementing the Darktrace / IDENTITY module. The requirements are listed in the Darktrace / IDENTITY section in Customer Portal.

4.3. External Alerting, Mobile App and Model Editing

Details on External Alerting, use of the Darktrace Mobile app and managing edits made to Autonomous Response models can be found in the applicable sections of the Darktrace / NETWORK product specification.

5. Administration

Detailed information about individual administrative tasks is outlined in the relevant documentation for the deployment of the specified component, module, or administrative action intended to be performed.

General administration of the Darktrace deployment is performed in the Threat Visualizer interface:

- The majority of configuration including administration of system settings such as proxies, authentication configuration such as LDAP and SAML SSO, deployment of alert and threat intelligence integrations, and other administrative tasks are performed in the “System Config” page.
- User and group management, including the assignment of data visibility and permissions, is performed on the “Permissions Admin” page.
- The application of tags to user entities is performed on the “Tag Manager” or “Profiles” pages.
- System health information and system alert resolution is performed on the “System Status” page.

Darktrace physical and cloud instances are each seeded with random passwords and two-factor authentication secrets at build time. These initial secrets are stored by Darktrace. Darktrace will provide the client with username/password combinations granting access to the Darktrace Threat Visualizer interface and, for physical instances, the SSH administration console (“Darktrace Console”); these passwords can be optionally changed by the client at any time. Access to the underlying backend systems of the Darktrace instance is reserved for Darktrace only.

5.1. Administration of Darktrace / IDENTITY Modules

Darktrace / IDENTITY modules are configured from the Darktrace “System Config” page. Each module has an individual entry with module-specific configuration settings. Users can modify Autonomous Response eligibility, user activity filters, and authenticate - or reauthenticate - the module from this location. It is a Customer responsibility to ensure that each module is configured correctly to ensure correct operation and interactivity with the associated third-party platform.

5.2. De-Commissioning

The process for removing a Darktrace / IDENTITY module will differ between third-party platforms.

Modules can be “de-authenticated” from the Darktrace System Config page - this will remove any authentication information used by the Darktrace ActiveAI Security Platform to contact the third-party platform. After de-authentication, any components created during the authentication process such as OAuth app registrations, service principals, API keys, or other authentication methods can be removed or deleted by the client from the relevant third-party platform. Darktrace strongly recommends that the client remove any remaining components when Darktrace / IDENTITY module service is ended.

6. Requirements

For specific technical requirements for a given module, consult the Customer Portal for the relevant technical documentation.

7. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer’s responsibilities are to ensure that:

- The Darktrace / IDENTITY module has been configured and authorized according to Darktrace instructions.

This includes ensuring that all permissions requested by the module are granted, whether during an authentication flow such as OAuth, granted separately as permissions or roles in the third-party platform, or granted/ provided in another comparable deployment process.

- Authentication is performed by - or granted via - a user with the appropriate permissions as outlined in the relevant Darktrace / IDENTITY module documentation.

  Recommendations or requirements placed on authorizing users are due to technical necessities. Darktrace cannot guarantee that OAuth grants, API keys, or other authentication methods performed by users without the appropriate permissions will result in functional module operation.

- The Darktrace / IDENTITY module retains authentication and, if lost, is re-authenticated in a timely manner. Unauthorized modules are unable to continue with monitoring or Darktrace Autonomous Response.

- The Darktrace / IDENTITY module is re-authenticated when requested by Darktrace. The module may require re-authentication to add additional capabilities (for example, after Autonomous Response is licensed).

- The Darktrace / IDENTITY module operates within the maximum threshold for event and user coverage, and that deployment components are not overloaded.

- Modifications are not made to the third-party platform configuration or licensing that would prevent operation.

  Examples include the removal of auditing within the platform, removal of licenses required to contact specific APIs, removal of permissions from associated API keys, or any other configuration alteration that would interrupt service operation.

- System health issues are monitored using the provided tools, then if alerts arise, addressed and rectified in a timely manner.

- End-user access is managed appropriately and within recommended best practice.

- Changes made to components or overall system configuration by users of the Darktrace platform do not impact the system’s ability to function, or otherwise degrade service.

  Examples of user error which may result in degraded service include the exclusion of key users using activity filters, the exclusion of all users from Darktrace Autonomous Response, the creation of excessively overactive models, and the modification of existing models to the extent that alerts criteria can no longer be met.

- Darktrace / IDENTITY Autonomous Response has been deployed in line with Darktrace best practice recommendations and, if changes are made to the third-party platform configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.

- Darktrace / IDENTITY Autonomous Response is enabled globally, in the deployment settings, and in any other applicable locations to ensure actions can be taken.

  Deployment configuration must also not preclude the taking of reasonable actions, for example, where all users are placed in an “immune” state.

- Darktrace / IDENTITY Autonomous Response is granted necessary autonomy to enact actions, or, where human confirmation mode is partially or fully in place, ensure that actions are activated by human operators in a timely manner.

- Alterations are not made to IDENTITY Autonomous Response models or models utilized by IDENTITY Autonomous Response that would significantly impede the ability to create targeted actions.
- IDENTITY Autonomous Response actions are not created excessively by custom models or existing models modified to contain actions.

8. Considerations

Due to the variation in third-party platforms, considerations will differ between individual Darktrace / IDENTITY modules. Please refer to the individual module documentation for further information.

Delays may be incurred where the external platform does not make events available to the Darktrace module for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the associated third-party platform vendor. Latency between event occurrence and when it was made available to the module are indicated in the event metadata within the Threat Visualizer.

Where delays are incurred receiving events from the third-party platform, Darktrace / IDENTITY modules may be unable to take Autonomous Response actions in a timely manner due to the latency between event creation and its ingestion and processing by Darktrace / IDENTITY.

If latency exists between Darktrace / IDENTITY and the third-party platform (for example, due to limited outbound bandwidth or delays in the third-party platform processing of requests), Darktrace / IDENTITY Autonomous Response actions may incur delays.

The responsibility to ensure that Darktrace / IDENTITY has an appropriate level of outbound access and bandwidth lies with the Customer. The responsibility for any delays incurred between instruction by Darktrace / IDENTITY and action occurrence in the third-party platform lies with the associated third-party platform vendor.

Individual vendors may place rate-limiting restrictions on the APIs and methods utilized by Darktrace / IDENTITY modules - please refer to the individual module documentation for further information.

Darktrace / IDENTITY modules monitor and retrieve data from third-party platforms using the APIs and integration methods presently offered by those platforms. If this coverage or available capabilities are modified or revoked by the third-party, Darktrace will endeavor to maintain continuity but is ultimately limited in capacity to restore service.

Darktrace Autonomous Response is subject to the restrictions placed by third-party platforms on API capabilities. This can result in specific user types being ineligible for actions, specific circumstances impacting the efficacy of actions, and the requirement for additional configuration (in both Darktrace and/or the third-party platform). Individual considerations are outlined in the individual module documentation.

Changes made to third-party platform capabilities or APIs are outside the scope of Darktrace control and may be introduced at any time. Changes to the known considerations due to alterations by vendors of these platforms may arise at any time.

In many third-party platforms, Darktrace / IDENTITY modules are reliant upon specific vendor licenses to retrieve the events necessary for operation or to take Autonomous Responses.

If the required licenses are not present, or are revoked, degradation or interruption of service will occur.

Product Agnostic Details

Table of Contents

1. Documentation

Documentation for the installation, operation and administration of Darktrace environments and Products is provided on the Darktrace Customer Portal.

Customer has the right to make a commercially reasonable number of copies of the Documentation, provided however, that Customer must reproduce and include all of Darktrace's and its suppliers' copyright notices and proprietary legends on each such copy.

2. Operational Change Management and Version Updates

2.1. Operational Change Management

Darktrace follows a set of standardized methodologies and procedures for efficient and prompt handling of changes to infrastructure and application, which enables beneficial changes to be made with minimal disruption to the service.

All Darktrace software is developed under secure development policies and practices. This is applicable to software run on Darktrace supplied hardware, managed by Darktrace operations in private cloud environments, or operated on third-party hardware (for example, Darktrace / ENDPOINT cSensor agents).

2.2. Version Updates

Version updates to Darktrace / NETWORK and Darktrace / IDENTITY are supplied through Threat Visualizer software bundles.

When “Call-Home” is enabled, all Darktrace Master appliances will automatically be upgraded by Darktrace to the latest Threat Visualizer release, unless an approval requirement has been configured and registered with Darktrace operations. Where possible, updates will be applied outside standard business hours. If this is not possible, the update process will cause minimal disruption for Threat Visualizer users.

If approval is required, an upgrade can be manually initiated from the management console or by Darktrace operations.

Where “Call-Home” is not enabled, software bundles are available from the Darktrace Customer Portal and can be unpacked and installed within the management console. Configuration of both automatic and manual updates in the management console is described in the relevant product documentation.

Darktrace vSensors will automatically update if granted access to the Darktrace update packages infrastructure. The update schedule for the Darktrace osSensor and other Agents offered in container format is defined by the client during configuration.

For cloud-hosted environments, software updates are managed by Darktrace operations. Threat Visualizer software is automatically updated when a new version becomes available; where possible, updates will be applied outside standard business hours. If this is not possible, the update process will cause minimal disruption for Threat Visualizer users.

Software releases are subject to the Darktrace Product Support and End Of Life (EOL) Policy, which can be found in the Customer Portal.

2.2.1.Model Updates

Darktrace will periodically update the standard supplied Darktrace Real-Time Detection models - customers with Call-Home or cloud-hosted instances will receive updates automatically, clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

3. Data Backups and Retention

For cloud-hosted instances managed by Darktrace operations, multiple short-term snapshot backups are taken on a rolling basis to ensure continuity in a disaster recovery scenario.

For physical appliances, Darktrace provides mechanisms to create backups of physical appliance configuration which can be exported to an external server. Backups should be treated as sensitive data and access should be protected. Backups can be created either manually, or automatically each day. By default, only the three most recent backups will be retained in the Darktrace appliance’s local storage. If a new backup file is created, the oldest backup in local storage will be discarded. A backup will include all Darktrace machine learning, models, and model breaches, as well as subnet information, device information, and configuration settings on the Threat Visualizer. A backup will *not* include transactional data such as connections in the Event Log, Advanced Search entries, and PCAP files, nor will it include configuration settings on the console menu.

For Darktrace deployments with Probe and Master appliances, only the Master appliance needs to be backed up. For Unified View deployments, or deployments with more than one Master appliance, all Master appliances must be backed up individually.

It may be necessary to re-authenticate a Darktrace / IDENTITY module after a restoration from backup, depending upon the specific module and deployment configuration.

Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, traffic makeup and other operational components which contribute data to the platform. Transactional data has a typical data retention of around 30 days.

4. Appliances

The technical specifications of each physical appliance size are detailed below:

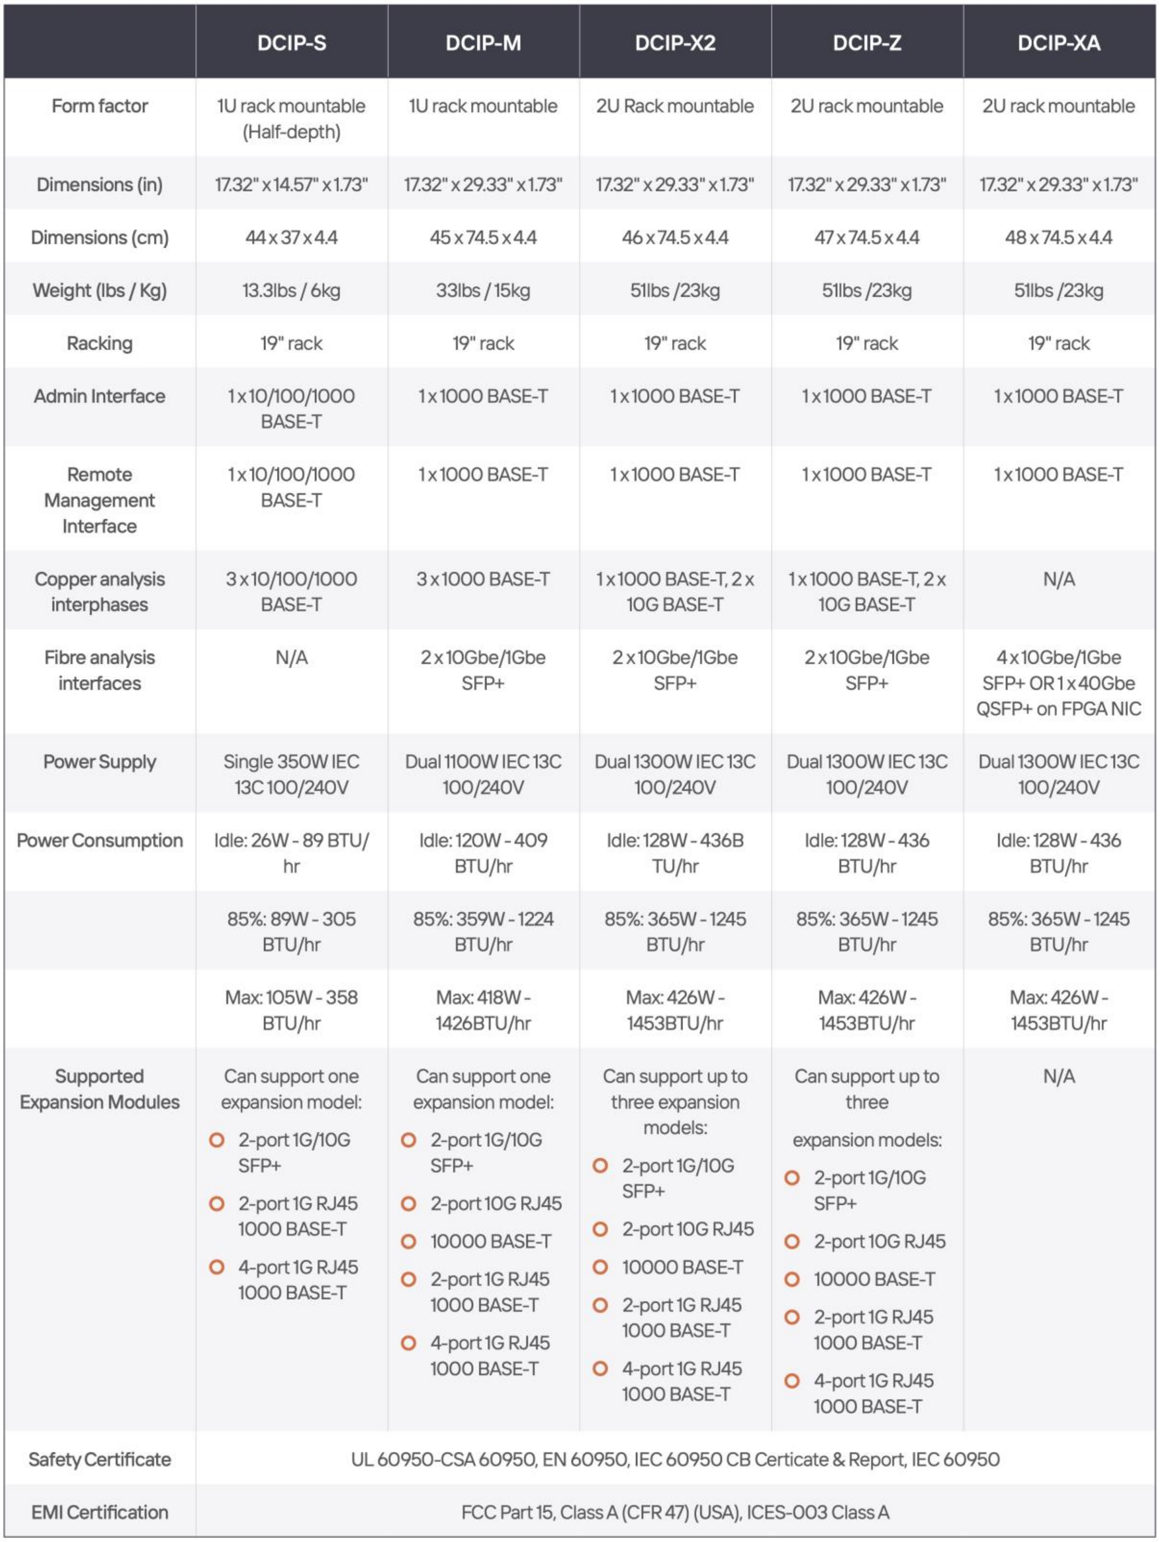

	DCIP-S	DCIP-M	DCIP-X2	DCIP-Z	DCIP-XA
Form factor	1U rack mountable (Half-depth)	1U rack mountable	2U Rack mountable	2U rack mountable	2U rack mountable
Dimensions (in)	17.32" x 14.57" x 1.73"	17.32" x 29.33" x 1.73"	17.32" x 29.33" x 1.73"	17.32" x 29.33" x 1.73"	17.32" x 29.33" x 1.73"
Dimensions (cm)	44 x 37 x 4.4	45 x 74.5 x 4.4	46 x 74.5 x 4.4	47 x 74.5 x 4.4	48 x 74.5 x 4.4
Weight (lbs / Kg)	13.3lbs / 6kg	33lbs / 15kg	51lbs /23kg	51lbs /23kg	51lbs /23kg
Racking	19" rack	19" rack	19" rack	19" rack	19" rack
Admin Interface	1 x 10/100/1000 BASE-T	1 x 1000 BASE-T	1 x 1000 BASE-T	1 x 1000 BASE-T	1 x 1000 BASE-T
Remote Management Interface	1 x 10/100/1000 BASE-T	1 x 1000 BASE-T	1 x 1000 BASE-T	1 x 1000 BASE-T	1 x 1000 BASE-T
Copper analysis interphases	3 x 10/100/1000 BASE-T	3 x 1000 BASE-T	1 x 1000 BASE-T, 2 x 10G BASE-T	1 x 1000 BASE-T, 2 x 10G BASE-T	N/A
Fibre analysis interfaces	N/A	2 x 10Gbe/1Gbe SFP+	2 x 10Gbe/1Gbe SFP+	2 x 10Gbe/1Gbe SFP+	4 x 10Gbe/1Gbe SFP+ OR 1 x 40Gbe QSFP+ on FPGA NIC
Power Supply	Single 350W IEC 13C 100/240V	Dual 1100W IEC 13C 100/240V	Dual 1300W IEC 13C 100/240V	Dual 1300W IEC 13C 100/240V	Dual 1300W IEC 13C 100/240V
Power Consumption	Idle: 26W - 89 BTU/hr	Idle: 120W - 409 BTU/hr	Idle: 128W - 436B TU/hr	Idle: 128W - 436 BTU/hr	Idle: 128W - 436 BTU/hr
	85%: 89W - 305 BTU/hr	85%: 359W - 1224 BTU/hr	85%: 365W - 1245 BTU/hr	85%: 365W - 1245 BTU/hr	85%: 365W - 1245 BTU/hr
	Max: 105W - 358 BTU/hr	Max: 418W - 1426BTU/hr	Max: 426W - 1453BTU/hr	Max: 426W - 1453BTU/hr	Max: 426W - 1453BTU/hr
Supported Expansion Modules	Can support one expansion model: ○ 2-port 1G/10G SFP+ ○ 2-port 1G RJ45 1000 BASE-T ○ 4-port 1G RJ45 1000 BASE-T	Can support one expansion model: ○ 2-port 1G/10G SFP+ ○ 2-port 10G RJ45 ○ 10000 BASE-T ○ 2-port 1G RJ45 1000 BASE-T ○ 4-port 1G RJ45 1000 BASE-T	Can support up to three expansion models: ○ 2-port 1G/10G SFP+ ○ 2-port 10G RJ45 ○ 10000 BASE-T ○ 2-port 1G RJ45 1000 BASE-T ○ 4-port 1G RJ45 1000 BASE-T	Can support up to three expansion models: ○ 2-port 1G/10G SFP+ ○ 2-port 10G RJ45 ○ 10000 BASE-T ○ 2-port 1G RJ45 1000 BASE-T ○ 4-port 1G RJ45 1000 BASE-T	N/A
Safety Certificate	UL 60950-CSA 60950, EN 60950, IEC 60950 CB Certicate & Report, IEC 60950				
EMI Certification	FCC Part 15, Class A (CFR 47) (USA), ICES-003 Class A				

4.1. Safety Information for Darktrace physical instances

Darktrace appliances must operate within thermal limits, such that the ambient inlet temperature must never exceed 35°C (95°F). Continued operation close to this limit may impair the long-term reliability of the system. Appliances are intended to operate in environments that meet ASHRAE Class A2 data center guidelines. A specialized rugged hardware Probe is available for industrial environments unsuitable for standard hardware.

All drive bays must be occupied. Empty drive bays should be occupied with a blank drive supplied by Darktrace.

Darktrace appliances are provided with a physical security seal on the chassis and a front bezel to protect the drive bays. The cover should not be removed by anyone other than a Darktrace engineer or at a minimum, under the remote supervision of a Darktrace engineer. When a security seal must be removed, new seals will be provided and should be applied immediately after work requiring the seal to be removed is complete.

4.2. Ownership and Return

Unless otherwise agreed to in writing, title to all Appliances (and all components thereof) provided by Darktrace to Customer will always remain with Darktrace. Customer's use of any Appliance is subject to this Product Specification.

Upon termination or expiration of the Evaluation Period or Subscription Period (as applicable), Customer shall:

(a) promptly return all Appliances to Darktrace or the applicable Partner, in accordance with Darktrace's (or the applicable Partner's) instructions; and

(b) ensure all Customer Data is removed from the Appliance. Darktrace shall not be responsible for maintaining or protecting any configuration setting or data found on the returned Appliance or component part of the Appliance and it is Customer's sole responsibility to delete any such information prior to return.

4.3. Delivery

Darktrace uses commercially reasonable efforts to ship Appliances on the delivery dates agreed in writing by Darktrace and the Customer (or Partner, if applicable); provided, however, that Customer's sole and exclusive remedy for any delay in delivery or for failure to give notice of delay shall be for Darktrace to make such delivery as soon as practicable thereafter.

Darktrace is not liable for the acts or omissions of any third-party courier or shipping provider. Darktrace may withhold or delay shipment of any order if Customer, or the applicable Partner, is late in payment or is otherwise in default under the Agreement or the Partner Arrangement.

The Appliance is provided solely as the medium for delivery and operation of the Software and must not be used for any other purpose. Whilst the Appliance is in Customer's possession, Customer must: (i) store and use the Appliance in a proper manner in conditions which adequately protect and preserve the Appliance; (ii) not sell, charge, pledge, mortgage or otherwise dispose of the Appliance or any part of it; (iii) not permit any lien to arise over the Appliance (or part thereof); and (iv) keep the Appliance free from distress, execution and other legal process.

Customer shall be responsible for preparing the delivery location for the delivery of the Appliance(s) and for the provision of all necessary access and facilities reasonably required to deliver and install the Appliance(s). If Darktrace is prevented from carrying out delivery or installation of any Appliance because no such preparation has been carried out, Darktrace may levy additional charges to recover its loss arising from such event.

In the event that there are additional fees attributed to delivering to the Customer’s designated delivery site, Darktrace shall deliver the Appliance(s) FCA (Incoterms 2010) to the agreed sites and unless otherwise set out in the applicable Product Order Form, Customer shall pay and be exclusively liable for all costs associated with shipping and delivery including without limitation, freight, shipping, customs charges and expenses, cost of special packaging or handling and insurance premiums incurred by Darktrace in connection with the shipment of the Appliance(s). Unless otherwise set out in the applicable Product Order Form or agreed in writing by Darktrace and the Customer (or Partner, if applicable), Darktrace will ship by the method of its choice.

Darktrace shall identify itself in all documents related to the shipment of the Appliance(s) as the exporter of record from the applicable jurisdiction of export, and Customer (or its agent, as applicable) as the importer of record into the country of delivery. Darktrace can provide Customer with reasonable assistance and support relating to the import of the Appliances.

5. Assumptions

The Software may contain or be accompanied by certain Third- Party Products including Open-Source Software. Any Open-Source Software provided to Customer as part of the Offering is copyrighted and is made available to Customer under the GPL/LGPL and other Open-Source Software licenses. Copies of, or references to, those licenses may be set out in a Product Order Form, the Third-Party Product packaging or in a text file, installation file or folder accompanying the Software. If delivery of Open Source Software source code is required by the applicable license, Customer may obtain the complete corresponding Open Source Software source code for a period of three years after Darktrace’s last shipment of the Software by sending a request to: Attn: Legal Department - Open Source Software Request, Darktrace Holdings Limited, Maurice Wilkes Building, Cowley Road, Cambridge CB4 0DS, United Kingdom. All other implied licenses are disclaimed, and all rights not expressly granted herein are reserved to Darktrace.

Unless expressly agreed between the Parties in writing, the Offering does not include the monitoring, interpretation or corrective action with respect to any Alerts. No advice, report, or information, whether oral or written, obtained by Customer from Darktrace or through or from the use of the Offering shall create any warranty not expressly stated in this Agreement. Customer understands that: (a) any outcome of the use of the Offering involving security assessment is limited to a point-in-time examination of Customer’s security status; and (b) the Offering does not constitute any form of representation, warranty or guarantee that Customer’s systems are secure from every form of attack, even if fully implemented. Customer understands and acknowledges that not all anomalies / intrusions may be reported or prevented.

The Customer shall perform or procure the performance of the responsibilities set out in this Service Definition in a manner which ensures Darktrace is not delayed from performing its obligations in accordance with this Agreement. To the extent that the Customer's delay or failure to comply with a responsibility does or may cause Darktrace to miss any timeframe for the performance of an obligation, Darktrace shall be entitled to an extension of time equivalent to the delay caused by the failure of the Customer.

Customer will own all rights, titles and interests in and to the Customer Data and the contents of any Alerts. In respect of any Customer Data stored on the Appliance, Customer grants to Darktrace a limited and non-exclusive license to access and use the Customer Data only to the extent necessary for Darktrace to perform the Services. Customer agrees Darktrace may utilize the details of any Alerts evaluations occurring in Customer’s network and any connected data source on an anonymized basis and excluding any Customer Confidential Information and / or Personal Data, to develop and improve the Darktrace technology. Customer is solely responsible for its use of the Offering, the activities of its users and for the accuracy, integrity, legality, reliability and appropriateness of all Customer Data. Customer expressly recognizes that Darktrace does not create nor endorse any Customer Data processed by or used in conjunction with the Offering. Customer further acknowledges that Darktrace and its Affiliates do not provide or undertake backup or maintenance services for Customer Data and Customer undertakes that it shall be solely responsible for backup of all Customer Data.

Upon expiry of the Subscription, Darktrace shall maintain Customer Data and grant Customer access to the cloud Services, solely to download and delete any Customer Data. Thereafter, Darktrace will delete or destroy all copies of Customer Data without liability or additional notice, unless legally prohibited from doing so. Customer Data cannot be recovered once deleted or destroyed.