Cisco Tetration Release Notes Release 3.3.2.2

This document describes the features, caveats, and limitations for the Cisco Tetration software, release 3.3.2.2.

The Cisco Tetration platform is designed to comprehensively address a number of data center operational and security challenges using rich traffic telemetry collected from servers, layer 4 through 7 service elements, and end-point devices (such as laptops, desktops, and smart phones). The platform performs advanced analytics using an algorithmic approach to offer a wholistic workload protection platform. This algorithmic approach includes unsupervised machine-learning techniques and behavioral analysis. The platform provides a ready-to-use solution supporting the following use cases:

- Provide behavior-based application insight to automate allowed-list policy generation
- Provide application segmentation to enable efficient and secure zero-trust implementation
- Provide consistent policy enforcement across on-premises data centers, and private and public clouds
- Identify process behavior deviations, and software vulnerabilities and exposure to reduce attack surface
- Identify application behavior changes and policy compliance deviations in near-real time
- Support comprehensive telemetry processing in a heterogeneous environment to provide actionable insight within minutes
- Comprehensive network performance metrics based on the telemetry collected from both switches and the servers
- Enable long-term data retention for deep forensics, analysis, and troubleshooting

To support the analysis and various use cases within the Cisco Tetration Analytics platform, consistent telemetry is required from across the data center infrastructure. Rich Cisco Tetration Analytics telemetry is collected using agents. There are different types of agents available to support both existing and new data center infrastructures. This release supports the following agent types:

- Software agents installed on virtual machine, bare-metal, or container hosts
- Embedded hardware agents in Cisco Nexus 9000 CloudScale series switches
- ERSPAN agents that can generate Cisco Tetration telemetry from copied packets
- NetFlow agents that can generate Cisco Tetration telemetry based on NetFlow v9 or IPFIX records
- Cisco AnyConnect and Cisco ISE integrations to collect telemetry from endpoints, such as laptops, desktops, and smartphones

Software agents also act as the policy enforcement point for the application segmentation. Using this approach, the Cisco Tetration platform enables consistent microsegmentation across public, private, and on-premises deployments. Agents enforce the policy using native operating system capabilities, thereby eliminating the need for the agent to be in the data path and providing a fail-safe option. Additional product documentation is listed in the “Related Documentation” section.

These Release Notes are sometimes updated with new information about restrictions and caveats. See the following website for the most recent version of this document:

https://www.cisco.com/c/en/us/support/security/tetration/products-release-notes-list.html

Table 1 shows the online change history for this document.

Table 1 Online History Change

Date	Description
August 26, 2019	Release 3.3.2.2 became available.

Contents

This document includes the following sections:

New and Changed Information

This section lists the new and changed features in this release and includes the following topics:

New Software Features

The following new software features are available in this release:

- This release supports Tetration hardware clusters built using Cisco UCS C220 M4 series as well as Cisco UCS C220 M5 series servers.
  - If you are running Tetration software on a Cisco UCS C220M4 series cluster, you can directly upgrade from 3.1.1.x to 3.3.2.2
  - If you are running Tetration software on a Cisco UCS C220M5 series cluster, you can directly upgrade from 3.2.1.x to 3.3.2.2
- Full visibility and policy enforcement support extended for the following operating system versions:
  - Red Hat Enterprise Linux Release 8
  - CentOS Release 8
  - Ubuntu 18.04 (except for forensics/CVE/process snapshot)
  - SUSE Linux 15 (except for forensics/CVE/process snapshot)
  - Windows Server 2019

- AIX deep visibility and enforcement are available for the following platforms as an ALPHA release, with the following caveats:
  - OS versions: 6.1, 7.1, 7.2 (PPC)
  - Agent deployment support is available only through installer script (no classic package download available)
  - AIX agent in this release does not support PID lookup for flows, TCP-related stats for flows, process forensics, software package inventory, software vulnerability and packet disposition
  - In order to use enforcement, ipfilter package is required on the workload
- Support is added to collect telemetry from end-point devices, such as laptops, desktops, smartphones, printers, HVAC systems, and other IoT devices through Cisco ISE integration. Cisco ISE provides the following benefits:
  - Augments end-point device information, device posture changes to provide visibility and stronger micro-segmentation policy based on this information.
  - Integration with LDAP server allows administrators to extend the micro-segmentation policy based on users, user groups, etc. (up to 6 LDAP attributes).
  - Integration with Cisco ISE uses pxGrid. See the following document for the details about this interface: https://www.cisco.com/c/en/us/products/security/pxgrid.html.
- Integration with Cisco ASA using NSEL (Network Secure Event Logging) and AVI load balancers for flow stitching and flow visibility.
- Deep visibility and enforcement software agent-related updates in this release for both new agent installations and agent upgrades:
  - For better manageability of agent SLAs, this release supports organizing agent configuration into three categories—Enforcement, Visibility and Forensics—and also introduces new configuration options to allow the user to define Memory Quota Limit, CPU Quota Limit, etc. for each of these functionalities.
  - All agent binaries are independently managed from the UI. If a given process detects that it is using more CPU/memory than configured for, it will restart itself.
  - Software agents will now report the status of the upgrade. If the upgrade has failed, the error will be shown in the UI.
  - Agents will now report the status of opening packet capture. If the interface is used for capturing, the status will be shown on UI.
  - New forensic event called “FollowProcess” is available in this release:
    - Users can define forensic rules in forensic configurations to follow processes based on certain forensic signals: ExecPath, CommandString, Username.
  - A new indicator for ‘DBR Ready’ (data backup and restore) is available. For this, DBR should be enabled.
  - The following additional information is available for agents:
    - A new indicator for ‘Interface Flow Extraction’ status
    - Agents can now be filtered based on tenants: the ‘Agent Filtered by Tenant’ feature
    - Agent status anomaly information is now available on the Agent Overview Page (for example, inactive, upgrade failed, enforcement policy out of sync).
    - Upgrade status “pending” and it is filterable on the Software Agents page.
  - In this release, there is a new dashboard for software vulnerabilities which let you focus your efforts on critical vulnerabilities and workloads that need the most attention.
    - The new page highlights the distribution of vulnerabilities in the chosen scope, as well as displaying vulnerabilities by different attributes; for example, complexity of exploits, can they be exploited over the

network, or attacker needs local access, etc. Furthermore, there are statistics to quickly filter out vulnerabilities that are remotely exploitable and have the lowest complexity to exploit.

- This new page is intended to help you identify workloads to focus on first and which packages to patch first.

- This release includes the following workload protection features:
  - Security dashboard enhancements
    - Attack Surface score now takes into account CVE vulnerabilities associated with a process and an open port.
    - An open port can be added to the allowed list for a scope to avoid having that port affect the attack surface score calculation.
    - Additional attack surface details about CVEs and policy result counts are shown in the UI to aid in decision making.
  - Enhancements to detect malicious processes using hashes
    - Detecting known malicious hashes
    - Also adding known legitimate hashes to the allowed list can reduce false alarms
    - This feature requires Tetration Cloud Connection to be enabled; only information about running process hashes is sent to determine the verdict
  - Enhancements to network anomaly detection algorithms
    - Better seasonality detection
    - Per protocol (TCP and UDP) network anomaly detection
  - Enhancements to process forensics
    - Default forensics “MITRE ATT&CK Profile” is added. The profile contains 24 default rules that can detect a number of MITRE ATT&CK Techniques from executing, persistence, privilege escalation, and defense evasion categories (https://attack.mitre.org). The users can create additional rules to cover additional MITRE ATT&CK Techniques based for their needs.
  - The forensics rule can be configured to support the following new use cases:
    - Detect child process creation of a parent process based on process attributes
    - Detect child process creation of a parent process while adding selected child process subtree to the allowed list can reduce false positives
    - Detect child process creation of a parent process only when its parent process meets certain criteria
    - Detect child process creation only if its parent and its grandparent meet certain criteria
    - Certain regular expression in the rules; for example, Event Type = Follow Process with ancestor Process Info - Exec Path matches (.*)(winword\.exe|excel\.exe|powerpnt\.exe)
    - The users can copy an existing profile to a new forensics profile and modified the copied profile
    - Default forensics rule “Tetration - Raw Socket” is updated to filter out more potential false positives
  - In this release, enforcement engine supports a “NOT” filter without requiring frequent computation of address set, thereby optimizing agent CPU overhead
    - Example

    CMDB upload:

1.0.0.0/8 location=UNKNOWN,...
1.2.3.4/32 location=US

For inventory-filter = {location != UNKNOWN}
in the previous releases, Address-set for this filter would include all the IP address learned from flows resulting in frequent crunching of address-set with always increasing number of members.
In this release, the pipeline automatically does the negation and determine the membership,
in this example
address-set=
[
(0.0.0.0 - 0.255.255.255),
(1.2.3.4),
(2.0.0.0 - 255.255.255.255)
] without depending upon the flow learned inventories

Similarly, for inventory-filter = {subnet != 10.0.0.0/8} will translate into
address-set=
[
(0.0.0.0 - 9.255.255.255),
(11.0.0.0 - 255.255.255.255)
] without depending upon the flow learned inventories.

- In this release, there is a new dashboard for software vulnerabilities which let you focus your efforts on critical vulnerabilities and workloads that need the most attention.
  - The new page highlights the distribution of vulnerabilities in the chosen scope, as well as displaying vulnerabilities by different attributes; for example, complexity of exploits, can they be exploited over the network, or attacker needs local access, etc. Furthermore, there are statistics to quickly filter out vulnerabilities that are remotely exploitable and have the lowest complexity to exploit.
  - This new page is intended to help you identify workloads to focus on first and which packages to patch first.
- Policy Enforcement for F5 Load Balancer
  - Support for route domain is now added for F5 external orchestrator and load-balancer agent. When upgrading to this release, all existing F5 external orchestrators will have default route domain zero assigned. In case a different route domain is configured in F5, this needs to be manually changed in external orchestrator accordingly. The same needs to be done for load-balancer agents configuration.
  - With the support for route domain F5 external orchestrator and load-balancer agent will consider only virtual servers belonging to given route domain.
  - In this release, F5 load balancer agent will program policy rules on a per virtual server basis, as opposed to the F5 global policy list in the previous release. This means the load-balancer agent will filter the policies based on virtual servers VIP, protocol and ports, and place the policies rules including catch-all in the individual virtual servers policy.

- Data backup and restore
  - Data backup-and-restore copies cluster data from the Tetration cluster to an external storage device. In the event of a disaster, data can be restored from this external storage to any cluster with the same form-factor.
  - In this release, this feature requires an activation key. There is no separate license associated with this feature—contact Cisco Support to receive the activation key.

- Backup is triggered once a day at a scheduled time, based on your configuration. A successful backup is called a checkpoint. Checkpoint is a point-in-time snapshot of the cluster's primary datastores (Mongo, Druid, HDFS, Consul, and Vault).
- This release provides a configuration utility, data-backup configuration wizard, and a planner to enable this functionality.
- The planner can be used to test the access to the object store, determine the storage requirement, and the backup duration needed for each day.
- The configuration utility is used to configure and schedule back-up in the cluster. After the initial full backup, only incremental (newer changes) changes are backed up. This aids in keeping the bandwidth requirements lower. A full backup schedule can be configured along with the backup schedule, which will back up data that was backed up in the past.

- External orchestrator enhancements support integration with Infoblox.
  - This new feature imports Infoblox subnets and hosts automatically every minute, and lets you create inventory filters/queries using the generated annotations. The annotation key name consists of the prefix "orchestrator" and an Infoblox extensible attribute name separated by underscore character; for example, "orchestrator_Department".
  - Infoblox extensible attributes name and value(s) are imported as populated by Infoblox, and retrieved via Infoblox SDK API. Single- and multi-values are supported with this release.
  - Host names and references are imported into Tetration as "machine_name" and "machine_id" respectively.
  - Even though Infoblox subnets are imported, please note there is no direct use of them in inventory filters/queries, as Tetration inventory does not yet have support for subnets.

- External orchestrator enhancements now support integration with DNS servers.
  - This new feature automates ingestion of DNS name-to-IP mappings using zone-transfer protocol. When adding DNS servers as external orchestrators, you must specify the DNS zones for which the IP mapping information is to be ingested.
  - The annotation key name consists of the prefix *orchestrator_system/dns_name*.

- Secure Connector.
  - A new workflow for connecting to External Orchestrators through the Tetration Secure Connector has been introduced.
  - Instead of Tetration dialing out to connect to external orchestrators (VMWare vCenter, F5 BIG-IP, etc.), a new component, the Tetration Secure Connector client, can now be used. Connector dials in to the Tetration cluster and creates a cryptographically-secure reverse tunnel that can be used by Tetration to reach external orchestrators within the client network. This is especially useful for Tetration-as-a-Service customers, where the existing connectivity model would have required external orchestrators to be directly reachable from outside the client network.

- Connectors and External Appliances: An entirely new workflow and deployment model is introduced for managing Tetration integrations and external appliances. This new workflow removes many manual steps for deploying appliance agents and a TAN appliance. These appliance agents and connectors are enabled and managed (including configuration management) directly through the Tetration UI. The number of supported appliances (and unique OVAs) are consolidated to three:
  - Tetration Ingest Appliance: Support for high-volume endpoint and flow-data ingestion through standard protocols such as NetFlow and IPFIX. Supported connectors include:
    - F5 BIG-IP
    - Citrix Netscaler

        - AVI (new)
        - ASA (new)
        - NetFlow
        - AWS
        - Meraki (new)
        - AnyConnect
    - Tetration Edge Appliance: Support for alert notifications or other low-volume data ingestion such as inventory enrichment. Supported connectors include:
        - Syslog
        - Email
        - Slack
        - PagerDuty
        - Kinesis
        - ISE (new)

- Configuration Management of Connectors and External Appliances: Configurations for virtual appliances and connectors can be created, updated, and removed from Tetration directly. Configurations can be applied in one of two modes.
    - Test and Apply: Test for the validity of the configuration and apply/commit the configuration. Examples of this configuration include: NTP, AWS, Syslog, Email, Slack, PagerDuty, Kinesis.
    - Discover: Test for the validity of the configuration, discover additional properties of the configuration, enhance the configuration using these properties, and apply/commit the configuration. LDAP is an example of a configuration that supports discovery mode. The basic configuration of LDAP is first tested for validity (for example, connectivity to the server). Once the basic configuration is validated, the list of common single-valued attributes are discovered and presented to the user. Subsequently, the user selects an attribute that corresponds to username and a list of up to 6 attributes that should be fetched/annotated for each inventory item. The final complete configuration is then applied to the connector.
- Data Exporter is a new capability designed to enable exporting aggregated flows and host-inventory data from a Tetration cluster.
    - Use Explore commands to set up data export for flows or inventory
    - Data is exported through managed data taps (MDT)
    - Exported data can be consumed outside of the Tetration cluster
    - Tetration Export appliance can be deployed to consume data which uses ELG stack (Elasticsearch, Logstash and Grafana), and can be used for further analysis and visualizations
    - Limit on export is 1.5 million (flows + inventory records) per minute across all the Tenants
    - Needs licensing to be enabled for “Data Export” Feature flag for Data Export to be used
- Compliance alerts can be configured on a Live Analysis policy.
    - The alert trigger condition and generated alert text will indicate whether the alert is for the enforced policy or live policy for the workspace.
- Geo information has been moved from Visit History tab to two separate tabs (Geo Inbound and Geo Outbound), and now has a map view, in addition to the tabular view.

- This release introduces the following enhancements to ADM:
  - An applications landing page offers an overview of not just application workspaces, but also all policies (analyzed or enforced). The page also provides buttons for various functionality such as adding a policy or creating a new filter. As before, clicking the Applications menu toggles between the most-recently viewed workspace and the overview page.
  - Published versions are limited to 100 total per workspace. Once this limit is reached, you will need to delete old versions using the UI or API.
  - New options to generate only policies and skip clustering upon an ADM run (Advanced Configurations). This feature is useful for those who understand their application component grouping and just want to generate the policy edges between the application component groupings. This feature is also useful for policy generation between coarse collections of inventory.
  - When generating policies, ADM uses filters in the External Dependency list to map IP addresses to filters (scopes or user inventory filters). If an IP address does not match any filter, in previous releases it would get assigned to the last filter/scope in the External Dependency list. Starting with this release, that address will map to root scope.
  - Ungrouped Policy table view: This view is differentiated by port (port-range) in addition to consumer/provider/action. Thus, you can search or filter the rows easily based on ports. In particular, you can view policy confidence (or confidence on the server port classification).
- Enhancements to neighborhood graphs
  - Two additional tabs in the neighborhood application to show inbound and outbound Geo information.
  - Geo tabs include map view with aggregated information, tabular view and detail
  - Alerts can be configured on this Geo information

  Note: This feature relies on the external Geo information dataset shipped with Tetration images. To keep the Geo dataset up to date, enable open-threat telemetry updates between the cluster and the latest data from the Tetration Cisco cloud.
- This release includes the following platform-related enhancements:
  - Yarn HA has been introduced with this release
  - Hadoop has been upgraded to 2.7.3 from 2.4.0 (apache version)
  - Beginning with this release, it is mandatory to register all on-premise Tetration appliances. When you deploy or upgrade Tetration software to this release on any on-premise appliance, you will enter a 90-day evaluation period and you have to register the cluster with Cisco within this period. Otherwise, your appliance will be considered to be out-of-compliant. Detailed instructions on how to obtain a license and use it in the appliance will be available to site administration users. Once the license is applied on the cluster, the appliance could be in a state of either in-compliance or over-use. Please note that Tetration features are not blocked due to out-of-compliance or over-use.
  - A new option to add additional debug log messages for external authorization to help debug connection issues, sign-in issues, and so on. Additional log messages are written into `external_auth_debug.log` if this option is on.
  - A new option in LDAP external authorization mode to 'Auto Create Users' if they are successfully authenticated with LDAP but do not exist in our database. If this option is turned off, the site administrator will have to pre-provision the user before the user tries to sign in.
  - A new option in LDAP to enable/disable authorization with LDAP. If this option is enabled, the site administrator will have to set up group-to-role mappings; that is, Active Directory group names to Tetration role mappings. And these mappings will be applied to users when they authenticate with LDAP. If this option is disabled, users are assigned roles based on the Tetration roles assigned to them when they were provisioned.

- A new option to enable/disable outbound HTTP connections.
- In the HTTP proxy portion, proxy port numbers other than 80 can be used.
- Sanitization of x.509 certificates before accepting for import.
- A new workflow to import certificate and key by creating a Certificate Signing Request.
- An external switch is no longer needed for initial bare-metal imaging. The bare-metal CIMCs can be connected to the spine (39RU) or leaf1 (8RU) and left there after the initial bare-metal imaging is complete.

- Data platform enhancements
  - Additional options for reading and writing JSON blobs using IO.read/IO.write APIs.
  - Improvements to ExternalApi:
    - Error messages on ExternalApi api calls are more clear; OpenApi error codes will be returned to the user application
    - Added the API call `ExternalApi.delete()`
  - New use-case example notebooks: these show how to aggregate hourly data into daily and weekly presentations, and show how to get long-term average scope-to-scope traffic, and how to calculate the segmentation policy effectiveness score

Changes in Behavior

These are changes in behavior for this release:

- Enforcement functionality is not available for the following Windows OS versions:
  - Windows Server 2008
  - Windows 7
- Forensics, software package, CVE and Process Snapshot (or file hashes) are not available on the following OS versions:
  - Ubuntu 18.04
  - SUSE 15
  - AIX 6.1, 7.1, 7.2 (PPC)
- The following OS versions are obsolete and are no longer available in this release. If they were running previous agent versions they won't be able to upgrade.
  - Ubuntu 12.04
  - Ubuntu 14.10
  - RHEL 5.0 to 5.6
  - CentOS 5.0 to 5.6
- The forensics process, "tet-worker", has been replaced by "tet-main", as the host forensics binary now operates independent of data gathering and manages its own back-end connections. However, tet-main will only connect to one collector at a time, instead of all available collectors like tet-sensor.
- When running on a Kubernetes node, the Tetration Enforcement agent no longer delays the starting of pods until their policies arrive. In previous releases, a CNI plugin was added to Kubernetes to pause the pod initialization for 15 seconds or until the policy arrives, whichever is earlier. This behavior has been discontinued in this release. In this

release, the Tetration Enforcement agent does not interfere with the pod initialization process; policies are applied to the pod as soon as they are received.

- The explore powerdown command/open API-based behavior is deprecated in favor of the UI workflow.
- The expected CIMC and individual bare-metal component firmware versions are dynamically loaded from the UCS Firmware RPM resulting in more accurate firmware version comparisons.
- The network performance monitoring feature set is now deprecated. Starting with this release, the performance monitoring and fabric features are disabled by default. They will be removed in subsequent releases.
  - If you were using this feature prior to this release or bought Tetration for its NPMD (network performance monitoring & diagnostic) capability, please contact your Cisco account team to learn how you can unlock this feature for a limited time.
- The following appliance agents are now supported as connectors:
  - NetFlow, Citrix NetScaler, F5 BIG-IP, AWS VPC Flow Logs, and AnyConnect Proxy. These connectors can only be enabled using the connector workflows on a Tetration Ingest appliance. When enabled, these connectors register as agents (similar to earlier releases).
- Appliance agents in 3.1, namely, NetFlow, Citrix NetScaler, AWS VPC Flow Logs, F5 BIG-IP, and AnyConnect Proxy, are supported in 3.3 only as connectors. These agents will not auto-upgrade from 3.1 to 3.3. The administrator has to redeploy these agents using the new connector workflows for managing Tetration integrations.
  - Cisco Tetration NetFlow Virtual Appliance: In this release, the equivalent of this appliance is to deploy a NetFlow connector on a Tetration Ingest appliance.
  - Citrix NetScaler AppFlow Appliance: In this release, the equivalent of this appliance is to deploy a Citrix NetScaler connector on a Tetration Ingest appliance.
  - AWS VPC Flow Logs Collector Appliance: In this release, the equivalent of this appliance is to deploy an AWS VPC Flow Logs connector on a Tetration Ingest appliance.
  - F5 BIG-IP IPFIX Collector Appliance: In this release, the equivalent of this appliance is to deploy an F5 BIG-IP connector on a Tetration Ingest appliance.
  - Cisco Tetration AnyConnect Proxy Appliance: In this release, the equivalent of this appliance is to deploy a AnyConnect connector on a Tetration Ingest appliance.
- Cisco Tetration Alert Notifier (TAN): Similar to appliance agents, alert notifieres are supported in 3.3 only as connectors. TAN Appliance (where these notifiers are instantiated) will not auto-upgrade from 3.1 to 3.3. The administrator has to redeploy these notifiers using the new connector workflows. In this release, the equivalent of TAN appliance is to deploy Syslog, Email, Slack, Pagerduty, and/or Kinesis connectors on a Tetration Edge appliance. This means Alert Notifier configuration can be done only from the connectors page, once the Tetration Edge appliance is deployed.
- Spark is upgraded from 1.6 to 2.3.
  - This may require some modifications to existing user applications.
  - Known potential changes (such as getting the sqlContext from SparkSession) are highlighted in the user guide.
- Data Lake Machine and Inventory data are removed (deprecated last release).
- Data Lake Shallowflows is deprecated. This data will have a substantially reduced storage window.
- TA_bogon_ipv4 and TA_zeus tags that were pushed by Lookout Annotation are not pushed to User Annotations; they are pushed to a separate annotation space so the user cannot accidentally delete them while modifying User Annotations. The switch happens after the next data pack update.

- Tetration will discard tags with duplicate keys for a workload. If a workload consists of such tags, one of the tags is randomly accepted and the rest are discarded.
- Non UTF-8 characters are not accepted in tags. If such character(s) is present, the tag's key/value is trimmed by discarding the invalid character(s). For example:

`abc\xc6de\xc8s` will produce `abcdes` as the new key/value
`\xc6` will be discarded since the key/value has no valid characters

Also, the maximum length of key/value of a tag must be less than 512 characters.

- Manual switchover of Yarn is not required in case of a node failure.
- User Defined policies have been migrated to Approved Policies. Policies were marked as user defined unless they were created via an ADM run. Policies created via the UI or a JSON import were marked as user defined. You can switch policies to and from approved.
- Beginning with this release, policy backdated experiment is conducted on conversations instead of flow data. This has been done to speed up backdated experiments.
- Static-mode application workspaces are deprecated. All new workspaces will be in dynamic mode. The main differentiator of dynamic-mode workspaces is the ability of clusters to process dynamic queries and not be limited to a static set of IP addresses. All workspaces will be upgraded to dynamic mode in the next release.
- For the FlowSearch API (POST /openapi/v1/flowsearch) , `scopeName` was an optional parameter prior to this release. Beginning with this release, `scopeName` is a mandatory parameter in the parameters passed to the FlowSearch API.
- The scope selector input box has enhancements to indicate that it is clickable, and the box displays "Select a scope" when clicked. It also displays a list of scopes for auto-complete with recently visited scopes in the suggestions list.
- The user management UI now includes a wizard workflow to create and edit users. This matches the wizard workflow for roles management which was introduced in the previous release.
- Site administrators can create snapshots for diagnostics.
- New user creation results in an email for password reset if the appliance is in local database authentication mode.
- Agents have been reorganized into three groups (Workloads, Endpoints and Flow Ingest). All other appliance agents are moved to Tetration Connectors Page.

Caveats

This section contains lists of open and resolved caveats, as well as known behaviors.

- During upgrade when a new RPM is uploaded, adhocKafka is gracefully shutdown. This is done to avoid Kafka index corruption. Kafka comes up after the upgrade. If upgrade is aborted after uploading the RPM, adhocKafka should be restarted using explore command.

Open Caveats

The following table lists the open caveats in this release. Click a bug ID to access Cisco's Bug Search Tool to see additional information about that bug.

Table 2 Open Caveats

Bug ID	Description
CSCvq72540	Broadcast addresses may be appear in the wrong scope in ADM
CSCvq20740	ADM AppView Reversing Consumer and Provider Labels
CSCvq48913	PowerShell Script needs more info when failing
CSCvq91327	An extra unknown node may be shown in cluster status
CSCvq82858	The site linter and site checker passed for invalid site_ssh_key
CSCvq26107	Enforcement ruleset breaks Linux UDP-Based Traceroute
CSCvo26666	After a node recommission a key file to process CIMC commands is missing from bare metal node
CSCvp10656	CIMC Internal Network is not tested after cluster deployment, upgrade, or reboot (edited)
CSCvp10580	Unable to change CIMC Internal Network outside of upgrade or reboot
CSCvq96155	Manual upgrade required for load balancer (AVI, F5, Citrix) enforcement agent
CSCvo42565	Cannot use # in ldap password on Anyconnect Proxy VM without double quotes around password
CSCvo19895	/local/tetration/log/tet-ldap-loader log requires timestamps in anyconnect VM
CSCvo17238	Update iRules to better handle possible iRule errors and add logging
CSCvq85892	Allow manual upgrade of netflow/span/f5/netscaler/aws/anyconnect via software agent upgrade page
CSCvr03130	Unable to change CIMC Internal Network Gateway through upgrade/reboot in 3.3.2.2

Resolved Caveats

The following table lists the resolved caveats in this release. Click a bug ID to access Cisco's Bug Search Tool to see additional information about that bug.

Table 3 Resolved Caveats

Bug ID	Description
CSCvn78222	enable to update the batch of lookout annotation after installing rpms by manual in offline state
CSCvo89242	ability not delete Default, Unknown and Tetration vrf name
CSCvp33648	Tetration V reboot fails with error pyVmomi.VmomiSupport.InvalidPowerState

CSCvf78109	Upgrade to Collectd ver. 5.72 / New version offers enhancements and better memory management.
CSCvm85308	Validate Qualys scan results: HTTP Security Header Not Detected
CSCvn86706	Lookout Annotation:Upon new rootscope addition zeus tags are not added unless UAS service is enabled
CSCvo59068	Outbound HTTP Connection fails on port 8080
CSCvo78365	Missing checks in Linux installer scripts
CSCvp18606	Server power off from UI fails with error Shutdown validation step failed due to AttributeError
CSCvq17715	Internal haproxy certificates are expired
CSCvq29036	Flow Search is displayed as "unknown" when annotation columns includes /
CSCvq52863	Vault tokens no longer valid one year after major upgrade
CSCvq54100	incorrect cabling image for M5 8RU in deployment guide
CSCvm68801	LACP on external ports is not configured which causes vPC for traffic into the cluster to fail.
CSCvp89096	Cable checks raises CABLECHECK_UNBOUND_10G_INTF on G2 clusters
CSCvq77108	Install npcap in noncompatible mode
CSCvq78946	[3.1.1.67] Anyconnect flows show incorrect LDAP user annotation
CSCvq21346	npcap 0.995 unstable - do not use with Tetration Deep Visibility or Enforcement Agent
CSCvp98092	Tetration Alerts in Platform don't work after enabling Email in Platform
CSCvn50243	redhat 5.11 legacy deep visibility and deep visibility sensor failed send data to tetration
CSCvo87526	Netflow Sensor stops sending flows after ~30 to 60 minutes
CSCvo98967	Linux Enforcement Agent installer fails due to SafeModeException when Namenode is not in safe mode
CSCvp06054	Upgrade and Reboot Fails in Pre Upgrade Checks at switch_config.yml
CSCvp24340	Inbound Rules for Windows Enforcement Agent are not allowing inbound Multicast or Broadcast traffic
CSCvp31371	Flow stitching fails on Citrix NetScaler AppFlow Appliance
CSCvp40579	Policies regenerated by ADM if an existing policy has ANY keyword
CSCvp71462	Unable to choose default timeframe for forensic event data
CSCvn90943	Hourly Fabric Alert Summary is not reported
CSCvn52935	Software Sensor captures packets only on one interface if its friendly name is non-English
CSCvn78767	Notifier docker container failed as missing tet-alert-notifier.tar.gz from the iso

CSCvo02165	Agents do not upgrade after changing Agent intent in the Agent profile to Auto-upgrade enabled.
CSCvn49926	Command get_cimc_techsupport fails to produce a Tech Support file.
CSCvn49906	External orchestrator lost from GUI after upgrade from 2.3.1.52 to 3.1.1.53
CSCvn50142	AIX universal sensor installation selftest failure due to idle/defunc process
CSCvn20704	Failed to enable policy analysis
CSCvn28898	MsiInstaller incorrectly installs WindowsAgentEngine enforcement service for deep visibility agents
CSCvn30664	UUID case change from dmidecode causing duplicate sensors in Tetration
CSCvn34366	Period (".") is no longer acceptable in CMDB attribute names
CSCvn37738	Make msi installer for windows software sensor not force a system reboot on sensor (auto)upgrade.
CSCvn64220	Sensor not reporting agent stats after upgrade to 3.1.1.54
CSCvh97957	Cisco Tetration Analytics Cross-Site Request Forgery Vulnerability
CSCvm84884	Tetration UI Possibly Impacted by CVE-2014-8730 (Poodle Attack - TLS)
CSCvk51665	Adding k8s external orchestrator using default api port (6443) does not import metadata from k8s
CSCvm88166	Define/tune retention for data under /app-logs
CSCvn12781	Druid Services fail causing "internal server error" returned in UI
CSCvn18783	Unable to access to CIMC with incorrect default gateway
CSCvn20511	Flows in flow search not updating while services are healthy and pipelines are running
CSCvi19170	Number of ECC memory errors reported by the service status page does not match Show Tech
CSCvi20538	Bosun alert: Correctable ECC errors should be for individual DIMMs not a sum of errors for the node.
CSCvj86257	Stop reporting correctable ECC errors in the Service Status page.
CSCvk34853	Clear text admin passwords written to the orchestrator.log during reimaging.
CSCvm35195	CitrixParser in tetration may crash after parsing a SLB config file
CSCvm57680	keepalived does not failover VIPs on appServer when an interface for Public Network is down.
CSCvm63714	Tetr-V // Graceful cluster power down fails
CSCvm85033	Qualys scan - AutoComplete Attribute Not Disabled for Password in Form Based Authentication
CSCvn46417	Tetration impacted with HDFS-6870
CSCvm90092	Tet-sensor may cause an overwrite of `/etc/audit/audit.rules` in RHEL 6

CSCvi71219	MSServer2016Standard Software Agents Cannot be Manually Upgraded
CSCvj23172	Commands get_cimc_techsupport, clear_ecc, clear_sel don't work on Tetration 2.3.1.41
CSCvj46846	Tech Supports are not gathered / incorporated into snapshot files on Tetration 2.3.1.41 clusters
CSCvk38762	Teteration agent update may fail on SUSE endpoints because tet-sensor process locks RPM database
CSCvm49542	Agent tet-worker process has high CPU utilization due to large system log file
CSCvf80588	Cisco Tetration Analytics Authentication Bypass Vulnerability
CSCvf80617	Cisco Tetration Analytics Remote Command Execution Vulnerability
CSCvf80602	Cisco Tetration Analytics Reflected XSS Vulnerability
CSCvf71955	Cisco Tetration Analytics Hardcoded SSH Authorized Key Vulnerability
CSCvh21899	Cisco Tetration Analytics Certificate Validation Vulnerability
CSCvh21844	Cisco Tetration Analytics ACL Bypass Vulnerability
CSCvj45311	HbaseRegion server periodicFlusher requesting flush and stops receiving requests
CSCvc64131	Not all services are displayed on the service nor cluster status pages.
CSCvc65452	Windows PATH variable getting modified
CSCve15116	Hadoop service switchover for HDFS and Yarn needs Tetration engineering involvment
CSCvf22308	Some clusters cannot run restartservices
CSCvf22828	Hardware Swith Agent reporting to default Scope while re-configured to use another VRF/Scope
CSCvf93113	Windows policy restricts sensor installation
CSCvg69762	IPV6 addresses show up in inventory when all IPv6 traffic is excluded in collection rules
CSCvg69774	With subnet in collection rules, ip addresses are displayed in Flows but not in Inventory
CSCvg72893	Changes in Apps-Policy views behaviors are not documented in TA user guide
CSCvh06306	Enforcement Policies not being pushed to the endpoints in any workspace
CSCvh08287	Enforcement Policies Not Enforced Due to inconsistent enforcement order
CSCvh47800	Running adm with fine granularity on external dependencies does not create policies requests.
CSCvh48928	Patch CentOS VMs on Tetration clusters for Spectre/Meltdown vulnerabilities
CSCvh48941	Python 2.6.6 used for several apps on CollectorDatamover VMs has security vulnerabilities.
CSCvh87245	After cluster upgrade some hardware sensors do not automatically upgrade to the new version

CSCvh88191	Tetration opens multiple TCP connections into VCenter causing other connections to drop
CSCvi19883	EX Switches Hardware Agents Reversing UDP and TCP Traffic, Misleading information displayed in ADM
CSCvi20041	Different threshold for notification of correctable ECC errors for Service Status vs Bosun alerts.
CSCvi59083	External Orchestrator configuration requires correct plugin name
CSCvc69960	Linux Tetration Agent fails install without "which"
CSCvd80405	Tetration cluster making GET requests to Google Analytics.
CSCve52628	Flow Search returning Error 504. Druid query timeout.
CSCve53091	tetpyclient module not compatible with
CSCve53686	TSDB Not Reporting on Itself in 2.0
CSCvg74804	Release Notes should document all behavior changes
CSCvh89813	Single node reimage fails due to non responsive bmmgr service on bare metal(s).
CSCvh89652	Tetration upgrade guide on Cisco website does not have the complete upgrade path.
CSCvi21617	ReadOnly user should not have privilege to create API KEY beyond their capability
CSCvi23470	SW Windows Sensor Universal Visibility 2.2.1.34-lw only sending ARP_REQUEST and UDP data
CSCvi60993	Sensor installation should not rely on tet-sensor's ability to use PAM and 'su'
CSCvi61862	SLES11 zypper install wil not pick the right RPM from repo
CSCvi63860	Not Possible to Install Tet Agent in Ubuntu when the OS is not using systemd
CSCvd86311	diskIsOff Alert Misreporting
CSCve62618	error - {df_instance=run-user-1000} for sys.diskUsage : error calling Eval: no results returned
CSCve95757	Tetration 2.0 Misdropped criteria has been changed
CSCve98414	Data Tap not selectable when activating Compliance App in Data Platform
CSCvf67422	Certs expire causing collectors to deny access to agents, which in turn causes flows to stop.
CSCvf68866	Shallow Flows are missing in Grafana and ADM Flows out of date
CSCvg44736	Source/Target Cluster number indicates Empty until clicked on
CSCvg44965	Some "row locks" in the hbase regions are not being released which cause inconsistencies in the UI
CSCvn70337	TAN needs a way to export logs out of the TAN appliance VM
CSCvn04971	ASA NSEL netflow flows not decoded by Neftlow appliance

CSCvk62307	Confirm minimum version of f5 supported for iRule
CSCvj89293	Option to collect logs from Virtual Appliances like f5 and Citrix
CSCvp89285	Tetration Netflow Virtual Appliance needs to use port 4739 for IPFIX protocol

Known Behaviors

- The enforcement engine does not enforce back-end policies for external Kubernetes services
  - An external Kubernetes service is one whose endpoints are manually defined; they are not automatically associated to the service through the use of a selector. An example of an external service is the default "kubernetes" service created by the Kubernetes system to connect with the api-server pods. This service's endpoints are created manually by Kubernetes on initialization. When such services are used as providers in a Tetration policy, Tetration does not write rules on the back-end endpoints.
  - Tetration will only write policies allowing the CoreDNS pods to access the Kubernetes service ClusterIP. No rules will be automatically written on the nodes/pods providing the Kubernetes service. Such rules will have to be defined manually.
- Container enforcement does not support Kubernetes clusters with kubeproxy running in IPVS mode
  - Starting with Kubernetes 1.11, kubeproxy supports handling services using IPVS instead of IPTABLES. This configuration is not currently supported by Tetration for enforcement.
- Policy Enforcement for F5 Ingress controller
  - Tetration software only supports F5 Ingress controller in the current release.
  - In the allowed-list policy model [CATCH ALL rule is DROP], you have to create a policy to allow traffic between F5 ingress controller pods and the Kubernetes API server, along with another rule to allow traffic between an F5 ingress controller pod and the F5 load balancer.
  - Tetration software only supports port 80 and port 443 for the F5 ingress controller.
- Load Balancer Agents for Policy Enforcement
  - Auto-upgrade to this release is not supported because of changes to the new deployment mechanics described here. This means you will need to reinstall this appliance agent.
    - There will be only one OVA image for all supported F5, Avi and Citrix load balancers. The agent will be installed and run directly in the created VM, and requires only one IP address with access to the Tetration cluster and the connected load-balancer appliance. That means no docker container(s) are configured as in previous releases.
    - The OVA provided with this release forbids root log-in via a console. You are recommended to set a new password for the built-in user "tetuser" via a console immediately after boot-up of the created VM. Without setting a new password, it will not be possible to enter the VM console.
    - If an "authorized_keys" file is given in the configuration iso during first boot-up of the VM, the start-up script will enable SSH service with no root login, and public key authentication only. The found "authorized_keys" file is then set for the built-in "tetuser," allowing you to log into the VM with given public keys.
    - In addition to the OVA deployment, you can download the agent's RPM (`tet-lbenforcer-f5/avi/citrix-3.3.1-el7_x86_64.rpm`) from the Tetration UI page "Software Agents Download" and install it on any Linux CentOS-7 compatible platform, whether it is a docker container, VM, or bare metal machine.

    - You will need to create policies to allow health-status checks between load balancers and backend servers in allowed-list mode.

- Identifying process hash anomalies
    - Frequency analysis (hence the output score) is done at rootscope level only.
    - Analysis is run once per hour.
    - File Hashes tab on Workload Profile page only shows process hash details analyzed in the last hour.
- Detecting network anomalies
    - Previously detected “Data Leak” events continue to be shown as “Data Leak” events
- FollowProcess forensic events based on ancestor lineage are up to four levels
- The Tetration Cluster does not throttle network bandwidth to object store
- If the Performance and Fabric Monitoring page is set as landing page

  If you have already set your landing page preference to the performance/fabric page, upon upgrade to this release, you may still automatically navigate to the performance/fabric page but you will encounter an authorization error message. This is expected behavior as these pages are disabled by default starting with this release. You can change your default landing page to a different page using Setting → Preferences → Select (the landing page).
- Lookout Annotation tag source may become disabled (in a few cases) during or shortly after an upgrade to this release as the switch happens. You should verify Lookout Annotation is still enabled for the sources you want after the TA_* tags are moved, and manually re-enable them if they have become disabled.
- On a NetFlow connector, NetFlow v9 or IPFIX records with custom enterprise information elements may not get exported to Tetration.
- When deploying a virtual appliance using the new connectors workflow, for optional fields, if the user enters a field, he/she will have to explicitly clear the field to avoid a warning being shown during appliance VM setup. The workaround is to click `cancel` button and restart from beginning
- For connectors on Tetration Ingest appliance and ISE connector on Tetration Edge appliance, upgrade is managed through Agent upgrade workflow. Agent config intent with Auto-Upgrade marked True should be applied to all these connectors. Unless this is set, these connectors would not be upgraded when Tetration gets upgraded. For other connectors (esp., alert notifier connectors), upgrade happens automatically.
- When deploying a virtual appliance using the new connectors workflow, for optional fields, if the user enters a field, he/she will have to explicitly clear the field to avoid a warning being shown during appliance VM setup. The workaround is to click `cancel` button and restart from beginning

Compatibility Information

The software agents in the 3.3.2.2 release support the following operating systems (virtual machines and bare-metal servers) for deep visibility:

- Linux:
    - CentOS-5.x: 5.7 to 5.11
    - CentOS-6.x: 6.1 to 6.10

- CentOS-7.x: 7.0, 7.1, 7.2, 7.3, 7.4 and 7.5
- Redhat Enterprise Linux-5.x: 5.7 to 5.11
- Redhat Enterprise Linux-6.x: 6.1 to 6.10
- Redhat Enterprise Linux-7.x: 7.0, 7.1, 7.2, 7.3, 7.4 and 7.5
- Redhat Enterprise Linux-8.0
- Oracle Linux Server-6.x: 6.0 to 6.10
- Oracle Linux Server-7x: 7.0, 7.1, 7.2, 7.3, 7.4 and 7.5
- SUSE Linux-11.x: 11.2, 11.3, and 11.4
- SUSE Linux-12.x: 12.0, 12.1, 12.2,12.3 and 12.4
- SUSE Linux-15.x: 15.0, 15.1
- Unbuntu-14.04
- Ubuntu-16.04
- Ubuntu-18.04

- Windows Server (64-bit):
  - Windows Server 2008R2 Datacenter
  - Windows Server 2008R2 Enterprise
  - Windows Server 2008R2 Essentials
  - Windows Server 2008R2 Standard
  - Windows Server 2012 Datacenter
  - Windows Server 2012 Enterprise
  - Windows Server 2012 Essentials
  - Windows Server 2012 Standard
  - Windows Server 2012R2 Datacenter
  - Windows Server 2012R2 Enterprise
  - Windows Server 2012R2 Essentials
  - Windows Server 2012R2 Standard
  - Windows Server 2016 Standard
  - Windows Server 2016 Essentials
  - Windows Server 2016 Datacenter
  - Windows Server 2019 Standard
  - Windows Server 2019 Essentials
  - Windows Server 2019 Datacenter
- Windows VDI desktop Client:
  - Microsoft Windows 8
  - Microsoft Windows 8 Pro
  - Microsoft Windows 8 Home
  - Microsoft Windows 8 Enterprise
  - Microsoft Windows 8.1
  - Microsoft Windows 8.1 Pro
  - Microsoft Windows 8.1 Home
  - Microsoft Windows 8.1 Enterprise

- Microsoft Windows 10
- Microsoft Windows 10 Pro
- Microsoft Windows 10 Home
- Microsoft Windows 10 Enterprise

- Container host OS version for full visibility:
  - Red Hat Enterprise Linux Release 7.1, 7.2, 7.3, 7.4
  - CentOS Release 7.1, 7.2, 7.3, 7.4
  - Ubuntu Release 16.04

The 3.3.2.2 release supports the following operating systems for the policy enforcement add-on capability:

- Linux:
  - CentOS-6.x: 6.1 to 6.10
  - CentOS-7.x: 7.0, 7.1, 7.2, 7.3, 7.4 and 7.5
  - Redhat Enterprise Linux-6.x: 6.1 to 6.10
  - Redhat Enterprise Linux-7.x: 7.0, 7.1, 7.2, 7.3, 7.4 and 7.5
  - Redhat Enterprise Linux-8,0
  - SUSE Linux-11.x: 11.2, 11.3, and 11.4
  - SUSE Linux-12.x: 12.0, 12.1, 12.2, 12.3 and 12.4
  - SUSE Linux-15.x: 15.0, 15.1
  - Oracle Linux Server-6.x: 6.0 to 6.10
  - Oracle Linux Server-7.x: 7.0, 7.1, 7.2, 7.3, 7.4 and 7.5
  - Ubuntu-16.04
  - Ubuntu-18.04

- Windows Server (64-bit):
  - Windows Server 2008 Datacenter
  - Windows Server 2008 Enterprise
  - Windows Server 2008 Essentials
  - Windows Server 2008 Standard
  - Windows Server 2008R2 Datacenter
  - Windows Server 2008R2 Enterprise
  - Windows Server 2008R2 Essentials
  - Windows Server 2008R2 Standard
  - Windows Server 2012 Datacenter
  - Windows Server 2012 Enterprise
  - Windows Server 2012 Essentials
  - Windows Server 2012 Standard
  - Windows Server 2012R2 Datacenter
  - Windows Server 2012R2 Enterprise
  - Windows Server 2012R2 Essentials
  - Windows Server 2012R2 Standard

- Windows Server 2016 Standard
- Windows Server 2016 Essentials
- Windows Server 2016 Datacenter
- Windows Server 2019 Standard
- Windows Server 2019 Essentials
- Windows Server 2019 Datacenter

- Windows VDI desktop Client:
  - Microsoft Windows 7
  - Microsoft Windows 7 Pro
  - Microsoft Windows 7 Home
  - Microsoft Windows 7 Enterprise
  - Microsoft Windows 8
  - Microsoft Windows 8 Pro
  - Microsoft Windows 8 Home
  - Microsoft Windows 8 Enterprise
  - Microsoft Windows 8.1
  - Microsoft Windows 8.1 Pro
  - Microsoft Windows 8.1 Home
  - Microsoft Windows 8.1 Enterprise
  - Microsoft Windows 10
  - Microsoft Windows 10 Pro
  - Microsoft Windows 10 Home
  - Microsoft Windows 10 Enterprise
- Container host OS version for policy enforcement:
  - Red Hat Enterprise Linux Release 7.1, 7.2, 7.3, 7.4
  - CentOS Release 7.1, 7.2, 7.3, 7.4
  - Ubuntu Release 16.04

The 3.3.2.2 release supports the following operating systems for the universal visibility agent :

- Linux 32-bit and 64-bit (CentOS 4.x, RHEL 4.x, CentOS 5.x, RHEL 5.x, and so on)
- Windows Server (32-bit and 64-bit)
- Solaris 11 on x86 (64-bit)
- AIX 5.3 (PPC)

The 3.3.2.2 release supports the following Cisco Nexus 9000 series switches in NX-OS and Cisco Application Centric Infrastructure (ACI) mode:

Table 4 Supported Cisco Nexus 9000 Series Switches in NX-OS and ACI Mode

Product line	Platform	Minimum Software release
Cisco Nexus 9300 platform switches (NX-OS mode)	Cisco Nexus 93180YC-EX, 93108TC-EX, and 93180LC-EX	Cisco NX-OS Release 9.2.1 and later
	Cisco Nexus 93180YC-FX, 93108TC-FX, and 9348GC-FXP	Cisco NX-OS Release 9.2.1 and later
	Cisco Nexus 9336C-FX2	Cisco NX-OS Release 9.2.1 and later
Cisco Nexus 9300 platform switches (ACI mode)	Cisco Nexus 93180YC-EX, 93108TC-EX, and 93180LC-EX	Cisco ACI Release 3.1(1i) and later
	Cisco Nexus 93180YC-FX, 93108TC-FX**	Cisco ACI Release 3.1(1i) and later
	Cisco Nexus 9348GC-FXP	Cisco ACI Release 3.1(1i) and later
	Cisco Nexus 9336C-FX2	Cisco ACI Release 3.2 and later
	Cisco Nexus 9500 series switches with N9K-X9736C-FX linecards only	Cisco ACI Release 3.1(1i) and later

**Network performance features using hardware agents is supported only in Cisco ACI mode with release 3.1 or later.

Usage Guidelines

This section lists usage guidelines for the Cisco Tetration Analytics software.

- You must use the Google Chrome browser version 40.0.0 or later to access the web-based user interface.
- After setting up your DNS, browse to the URL of your Cisco Tetration Analytics cluster: https://<cluster.domain>

Verified Scalability Limits

The following tables provide the scalability limits for Cisco Tetration (39-RU), Cisco Tetration-M (8-RU), and Cisco Tetration Cloud:

Table 5 Scalability Limits for Cisco Tetration (39-RU)

Configurable Option	Scale
Number of workloads	Up to 25,000 (VM or Baremetal)
Flow features per second	Up to 2 Million
Number of hardware agent enabled Cisco Nexus 9000 series switches	Up to 100

Note: Supported scale will always be based on which ever parameter reaches the limit first

Table 6 Scalability Limits for Cisco Tetration-M (8-RU)

Configurable Option	Scale
Number of workloads	Up to 5,000 (VM or Baremetal)
Flow features per second	Up to 500,000

Number of hardware agent enabled Cisco Nexus 9000 series switches	Up to 100

Note: Supported scale will always be based on which ever parameter reaches the limit first

Table 7 Scalability Limits for Cisco Tetration Virtual (VMWare ESXi)

Configurable Option	Scale
Number of workloads	Up to 1,000 (VM or Baremetal)
Flow features per second	Up to 70,000
Number of hardware agent enabled Cisco Nexus 9000 series switches	Not supported

Note: Supported scale will always be based on which ever parameter reaches the limit first.

Related Documentation

The Cisco Tetration Analytics documentation can be accessed from the following websites:

Tetration Datasheets: https://www.cisco.com/c/en/us/products/security/tetration/datasheet-listing.html

General Documentation: https://www.cisco.com/c/en/us/support/security/tetration/tsd-products-support-series-home.html

The documentation includes installation information and release notes.

Table 8 Installation Documentation

Document	Description
Cisco Tetration Analytics Cluster Deployment Guide	Describes the physical configuration, site preparation, and cabling of a single- and dual-rack installation for M4 based Cisco Tetration (39-RU) platform and Cisco Tetration-M (8-RU). Document Link: https://www.cisco.com/c/en/us/td/docs/security/workload_security/tetration-analytics/hw/installation_guide/Cisco-Tetration-Analytics-Cluster-Hardware-Deployment-Guide.html Describes the physical configuration, site preparation, and cabling of a single- and dual-rack installation for M5 based Cisco Tetration (39-RU) platform and Cisco Tetration-M (8-RU). https://www.cisco.com/c/en/us/td/docs/security/workload_security/tetration-analytics/hw/installation_guide/Cisco-Tetration-M5-Cluster-Hardware-Deployment-Guide.html
Cisco Tetration Virtual Deployment Guide	Describes the deployment of Tetration virtual appliance. Document Link: https://www.cisco.com/c/en/us/td/docs/security/workload_security/tetration-analytics/sw/install/b_Tetration_Analytics_Virtual_Appliance_Deployment

	_Guide.html
Cisco Tetration Cluster Upgrade Guide	Document Link: https://www.cisco.com/c/en/us/td/docs/security/workload_security/tetration-analytics/sw/install/b_Tetration_Analytics_Upgrade_Guide.html
Latest Threat Data Sources	https://updates.tetrationcloud.com/