LEVEL 8 SOLUTIONS

i-Vertix

Technical Specifications

Date: February 5, 2024

# i-Vertix Log & Data Management


**MSP Edition**

A Detailed Insight into Architecture, Capabilities & Technical Specifications

**PGUM S.r.l**

NOI Techpark
Via Ipazia 2
39100 Bolzano - Italia

Mail: info@i-vertix.com
PEC: pgumgmbh@pec.it
Web: i-vertix.com

VAT: 02923310219
REA: BZ - 216954

# Table of contents

# THE IMPORTANCE OF A LOG MANAGEMENT SOLUTION

In order to properly manage and secure an IT environment, administrators need to know, in detail and in real-time, what's happening on all of its assets. Such information is spread across large volumes of devices, OSs and applications logs.

Securely collecting, archiving and analyzing log data of each asset in the IT environment is key to keep it secure and compliant with internal policies and Data Protection regulations, detect early indicators of attacks, reveal issues, provide forensic evidence, demonstrate compliance.

Common challenges IT managers have to face are the huge volumes of heterogeneous logs that are to be managed, and the prohibitive costs of SIEMs.

i-Vertix Log & Data Management provides efficient, scalable and reliable centralization, normalization and analysis of logs across both centralized and geographically distributed environments. It helps making IT operations secure and compliant and can also drastically reduce the cost associated with a SIEM by prefiltering and preprocessing logs and forwarding to the SIEM only the ones that really matter and require further processing.

## Real-time data analysis

i-Vertix Log & Data Management collects, archives, signs, parses, filters, classifies, enriches, encrypts and purges any kinds of logs coming from different sources and customers. Logs are processed in real-time by a high-performance log management engine.

The architecture is highly scalable, so the solution fits any environments, from small centralized deployments, to large distributed multitenant ones with huge amounts of data.

The system comes with several preconfigured dashboards, but users are free to create new custom ones. All collected data can be enriched with additional information.

## Authentication analysis

Authentication data provide visibility on users that access (or try to access) systems. i-Vertix Log & Data Management collects all the authentication events (successful or failed logins and logoffs) across different device types and platforms (Windows, Office365, Linux SSH, IBM iSeries, Network devices, storages, etc.) and normalizes them, providing a consolidated, consistent overview of authentications that take place throughout the whole IT environment.

## Optimized Log collection strategy

While i-Vertix Log & Data Management meets the requirements of most customers, it can also be extremely useful in complex IT environments where a SIEM is used.

Large customers can use their SIEM to manage only critical logs/events and use i-Vertix Log & Data Management to:

- perform effective data aggregation, parsing and normalization across all different data sources,
- discard background noise/lower value events,
- optimize log retention, reducing storage requirements and maintenance costs,
- classify/enrich your data according to policies and Data protection regulations,
- send only critical, filtered and normalized data to your SIEM.

## Network traffic analysis

Using i-Vertix it is also possible to visualize and analyze network traffic. In fact, i-Vertix Log Collector can collect, aggregate and analyze traffic flow records (NetFlow, sFlow, IPFIX, ...). In addition to real-time traffic analysis, it can also perform malicious sources detection and alert on network traffic related events, in combination with the i-Vertix IT Monitoring module.

## Data Privacy Regulations Compliance

i-Vertix helps meet Data Protection and Data Privacy regulations compliance requirements. This is done by providing automation, advanced security controls, log analysis, log classification and incident response, and is applicable to Windows EventLog and Syslog messages.

The default i-Vertix ruleset provides support for GDPR, PCI DSS, HIPAA, NIST 800-53 and TSC frameworks and standards.

Leveraging integrations with constantly updated public blocklists of malicious IP addresses, domains and URLs and public knowledge bases of adversary tactics and techniques based on real-world observations, i-Vertix rules can be used to detect early indicators of attacks and reveal potential issues.

# ARCHITECTURE

i-Vertix Log Management architecture is purpose-bult for MSPs and Enterprises. It is distributed, multi-tenant, highly scalable and supports granular Role Based Access Control (RBAC).

It has made up of 2 components: i-Vertix **Log Collector** (iLC) and i-Vertix **Data Node** (iDN).

The **Log Collector** is responsible for log collection, parsing, normalization, ingestion, archival, compression and encryption. It is typically installed close to the log sources, on premises at customers' sites / main remote sites.

It has a local storage were logfiles are saved and signed.

i-Vertix **Data Node** is the central collector/aggregator and analyzer of the logs coming from the Log Collectors. It can be installed on premises at the MSP SOC/NOC/Datacenter or in private or public cloud. It can be deployed in High availability & load balancing mode.

The Log Collectors securely uploads logs to the Data Node(s), using HTTPS TLS encryption.

## Deployments examples

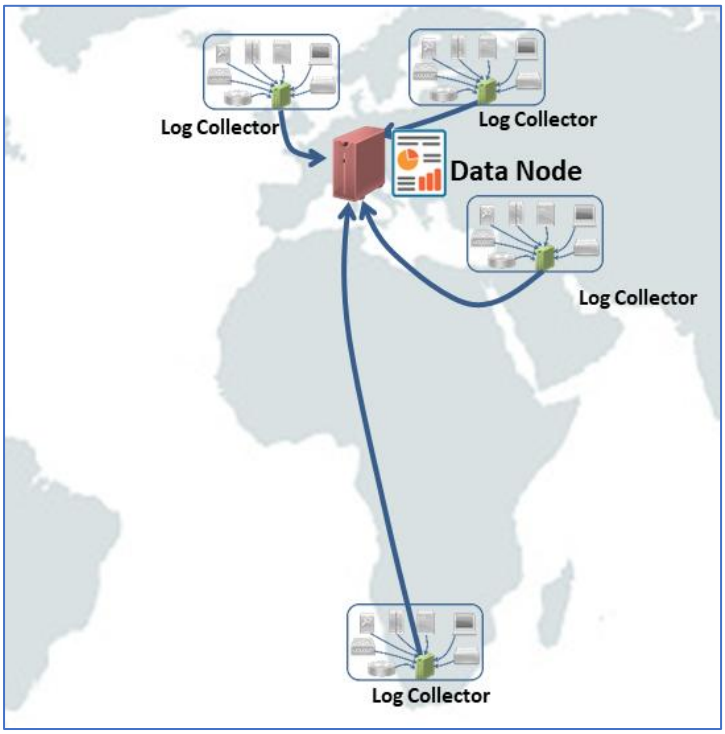


Figure 1 - Single Data Node on prem deployment

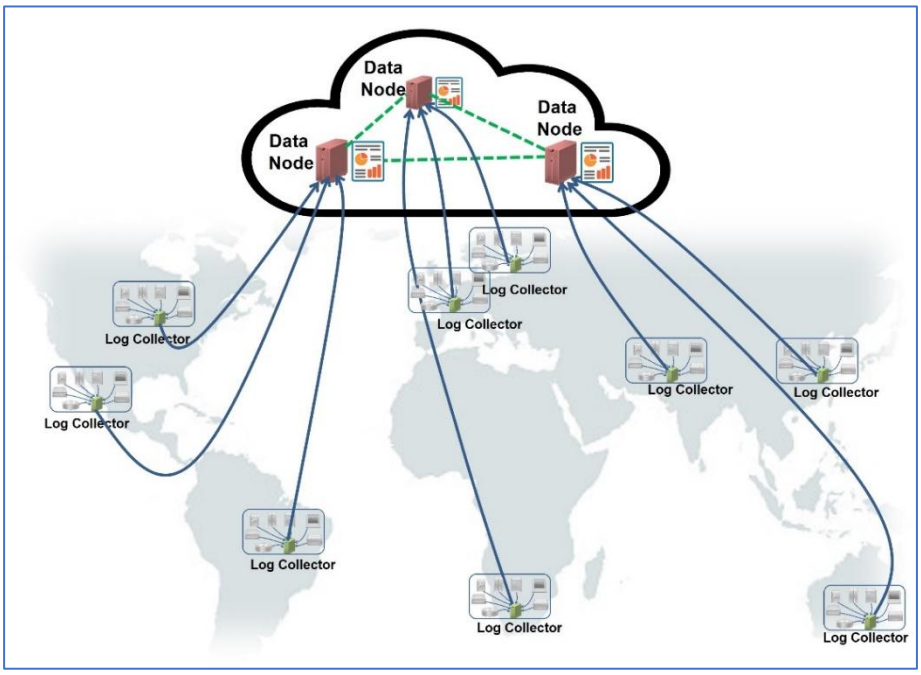


Figure 2 - 3-Data Node HA cluster in cloud deployment

## Clustering Data Nodes

Data Nodes can be clustered (software cluster). To create a cluster at least 3 Data Nodes are to be used. In clustered Data Nodes scenarios, Log Collectors upload logs to all Data Nodes that are part of the cluster. Each Data Node has its own dedicated storage.

One of the Data Nodes is the **Master** one and manages the cluster.

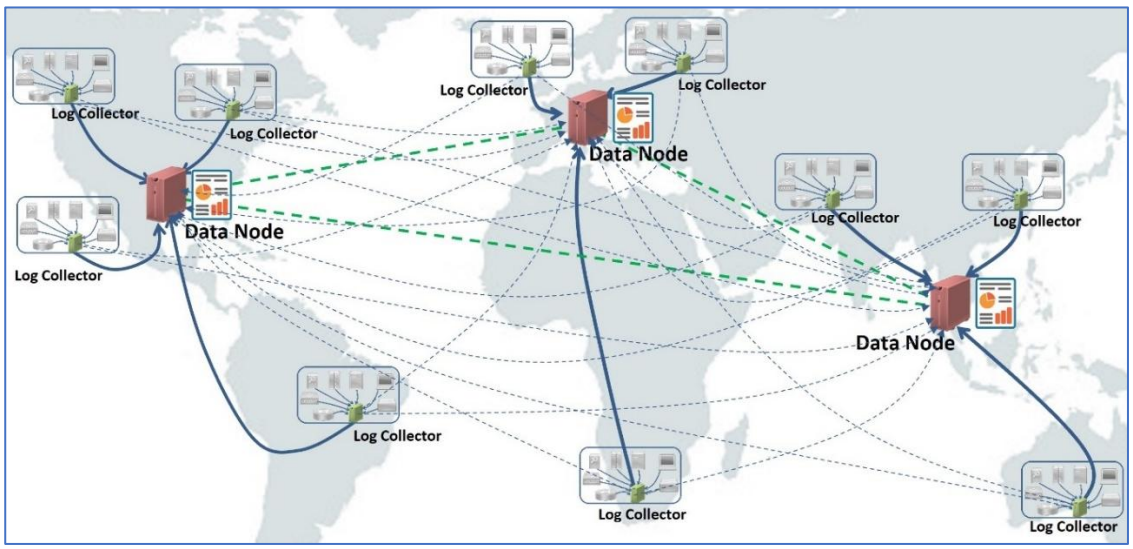


Figure 3 - 3-Data Node Cluster with Log Collectors sending logs to all members

Logs generated by log sources are saved into **Indexes**, made up of **documents** (indexes can be thought of as being like database tables while documents are like database records).

Each document is saved in an Index.

An Index doesn't have a fixed structure. For each type of log source, a specific Index type is created.

Indexes are created **per log source per day**. For example, *rsyslog-2022.09.10* is the index created by the product to save logs generated by source "rsyslog" on September 10th 2022.

Indexes can be of two types: "**primary**" or "**replica**". Each "primary" index can have N "replicas". By default, N = 1 (like in a RAID 5 disks array configuration).

In a 3-Data Node cluster, if N is set to 2 replicas, then each log/document is saved on each Data Node.

Should a Data Node go DOWN, Master Data Node triggers replication of faulty Data Node primary indexes to another Data Node and then triggers a synchronization process. This way, indexes are always available at least on one Data Node.

Data synchronization across Data Nodes is asynchronous. There is no specific minimum network performance required to run a healthy Elasticsearch cluster. In theory, a cluster will work correctly even if the round-trip latency between nodes is several hundred milliseconds. In practice, if your network is that slow then the cluster performance will be very poor.

# TECHNICAL SPECIFICATIONS

## Software Requirements

### Operating System - Log Collector and Data Node

i-Vertix Log Management is released as a software Image that includes CentOS 7.9 (latest version) and all the required packages.

### Elastic stack (Data Node)

The following table lists the software dependencies:

| Software | Version |
|---|---|
| Elasticsearch | 7.17.9 |
| Logstash | 7.17.9 |
| Kibana | 7.17.9 |
| Beats | 7.17.9 |

### Web browser

i-Vertix Log Management web GUI is provided by Data Nodes on a dedicated host.

It is compatible with the following web browsers:

- Google Chrome (latest version)
- Mozilla Firefox (latest version)
- Apple Safari (latest version)

Screen resolution must be 1280 x 768, 1920 x 1080 or higher.

# Hardware Requirements

## i-Vertix Log&Data Collector (iLC)

### Minimum requirements:

- Hypervisor:
  - VMware
  - XenServer
  - AWS Elastic Compute Cloud
  - Oracle VM
  - VirtualBox
  - Proxmox VE
  - Hyper-V by using Microsoft conversion tools.
- Processor: 4 vCPU
- Memory: 4GB RAM
- Disk: 250GB SSD

### Recommended requirements

- Hypervisor:
  - VMware
  - XenServer
  - AWS Elastic Compute Cloud
  - Oracle VM
  - VirtualBox
  - Proxmox VE
  - Hyper-V by using Microsoft conversion tools.
- CPU: 8 vCPU
- RAM: 12GB
- HDD: 500GB

System specification may have to be increased based on events/sec rate and amount of logs that are to be saved into the local storage.

Also in heavily loaded scenarios, a physical server is recommended.

# i-Vertix ELK Data Node (iDN)

The i-Vertix ELK Data Node is the central collector and aggregator of the logs sent by Log Collectors

## Minimum requirements:

- Hypervisor: VMware
- CPU: 2 vCPU
- RAM: 8GB
- HDD: 250GB (SSD)

## Recommended requirements

- Hypervisor: VMware or physical host
- CPU: 8-12 vCPU
- RAM: 64GB
- HDD: 2TB (SSD or NVME)

System specification may have to be increased based on events/sec rate, amount of logs that are to be saved into the local storage and queries done.

Also in heavily loaded scenarios, a physical server is recommended.

## Network flows table

| From | To | Protocol | Port | Application |
|---|---|---|---|---|
| iLC | iDN | HTTP/HTTPS | TCP 9200 | Send data to iDN |
| iLC | NTP server | NTP | UDP 123 | Synchronization of the system clock |
| iLC | DNS server | DNS | UDP 53 | Domain name resolution |
| iLC | Internet | HTTP/HTTPS | TCP 80<br>TCP 443 | System updates<br>Blacklist updates |
| iDN | NTP server | NTP | UDP 123 | Synchronization of the system clock |
| iDN | DNS server | DNS | UDP 53 | Domain name resolution |
| iDN | Internet | HTTP/HTTPS | TCP 80<br>TCP 443 | System updates |

| From | To | Protocol | Port | Application |
|---|---|---|---|---|
| iDN<br>Beats | iDN | HTTP/HTTPS | TCP 9200 | Access to API<br>Send monitoring data |
| iDN | iDN | HTTP/HTTPS | TCP 9300 | Cluster synchronization (on standalone not needed) |
| Client | GUI | HTTP/HTTPS | TCP: 5601 | Access to GUI |
| Source | iLC | SYSLOG | UDP 514<br>TCP 514 | Syslog receiver |
| Source | iLC | Windows Event Logs | TCP 1514 | Windows Event Log receiver |
| Source | iLC | Windows Event Logs | TCP 1515 | Windows Event Log receiver with TLS encryption |
| Source | iLC | Beats | TCP 5044 | Beats receiver |
| Source | iLC | Netflow | UDP 9901 | Netflow receiver |
| Source | iLC | CEF | UPD 1515 | CEF receiver |
| Source | iLC | Netscaler | UDP 1516<br>TCP 1516 | Citrix Netscaler Syslog receiver |
| Source | iLC | sFlow | UDP 6343 | sFlow receiver |

## Standard datasources table

| Source | Description | Agent | Info |
|---|---|---|---|
| Syslog | Syslog RFC3164, RFC5424 | Agentless | Configure Syslog forwarder to i-Vertix Log Collector (TCP/UDP 514)<br>If a tenantcode is needed define some rules on Log Collector |
| Windows Event Log | Windows Event Logs | Nxlog | Configure nxlog to forward to i-Vertix Log Collector (TCP 1514). Optionally you can enable TLS encryption.<br>You can define a tenantcode |
| Office 365 | Office 365 Audit logs | Filebeat | Configure filebeat to forward to i-Vertix Log Collector (TCP 5044).<br>O365 tenant name should be the tenantcode |

| Source | Description | Agent | Info |
|---|---|---|---|
| CEF | Common Event Format | Agentless | Configure CEF forwarder to i-Vertix Log Collector (UDP 1515) |
| Netlfow | Netflow v5, v9, v10 | Agentless | Configure Netflow forwarder to i-Vertix Log Collector (UDP 9901) |
| sFlow | sFlow | Agentless | Configure sFlow forwarder to i-Vertix Log Collector (UDP 6343) |
| IIS | Internet Information Server | Filebeat | Configure filebeat to forward to i-Vertix Log Collector (TCP 5044) using tag “iis”. |
| Apache | Apache Webserver | Filebeat | Configure filebeat to forward to i-Vertix Log Collector (TCP 5044) using tag “apache”. |
| Windows DNS | Windows DNS Server Logs | Filebeat | Configure filebeat to forward to i-Vertix Log Collector (TCP 5044) using tag “dns”. |
| Exchange | Exchange 2013, 2016, 2019 Transport logs | Filebeat | Configure filebeat to forward to i-Vertix Log Collector (TCP 5044) using tag “exchange”. |
| Citrix Netscaler | American format, local time | Agentless | Configure log forwarder to i-Vertix Log Collector (TCP/UDP 1516)<br><br>If a tenantcode is needed define some rules on Log Collector |

# LICENSING - MSP PAY PER USE MODEL

**MSP pay-per-use model** is available for MSPs/MSSPs that provides Log Management as a managed service to their customers.

It provides MSPs with the following **benefits**:

- **No license tiers**, so no unused/wasted licenses,
- **No upfront costs**: monthly fee based on **number of Data Nodes and number of Log Collectors**.
- Fast & easy new customers enrollment & predictable budgeting
- Maximum flexibility: MSPs can dynamically adjust the extension of their services
- It includes:
    - i-Vertix technical support and product updates
    - It also **includes the ordinary maintenance of the whole solution: periodic installation of product updates & patches performed by i-Vertix engineers**

This licensing, paired with **large product scalability**, that sustains growing log rates and volumes, while meeting systems specs. and storage budget constraints, is key to MSPs to offer a reasonably-priced managed service to their customers.

# SOLUTION SIZING

As the licensing is per-Data Node and per-Log Collector, in order to properly size the solution, the following information should be collected:

- **Number and type of log sources** (source hosts, PCs/Servers/VMs/DB&App servers, etc.)
- **Total number of users** (across all systems) for whom logs are generated
  **Note**: this is not the number of users who access the i-Vertix solution
- **Approximate average logs/second rate, per source**
- **IT environment architecture**: is it centralized (single site) or distributed across multiple sites/Datacenters? How are log sources distributed across sites?
  This is useful to determine the **number of Log Collectors and Data Nodes**, and which **type of**

**deployment** is required according to Customer's datacenter and networking specification and constraints.

- **Log retention period** (days)
- Is **High Availability** required?
- In addition to Log Management, will traffic analysis (based on Netflow family traffic flows technologies) be performed?

  If so:

  - **Number of traffic flows sources** (devices exporting traffic flows records).
  - **Approximate average traffic flows/second rate, per source**
  - **Traffic flows retention period** (days)

As many parameters are to be taken into consideration here, and it can happen that not all of the aforementioned information is available (in particular determining the log rates could be difficult). If necessary i-Vertix can make Log Collectors available to help the MSP understand how many logs are being generated and their rates.

The most common deployment is a cluster of 3 Data Nodes installed on premises at MSP's NOC/SOC and 1 or multiple Log Collectors installed on premises at customers' sites.

## Log archival strategy

Provided that the amount of logs that are to be archived can vary a lot, depending on number and type of log sources, log rates and log retention periods, the best practice is:

- Retain in Data Nodes Elastic database a certain amount (number of days) of logs (typically some hundred GB or a few TB)
- Archive older logs on a central archive/storage, for example an S3 (WORM) which is cheaper.

Logs that have been archived on the "long term archival storage" have to be reviewed, for example because of an Audit, can simply be imported on the Data Node Elastic database for an **a-posteriori analysis**.

# SIZING EXAMPLES

In this section we provide a list of real cases / deployment scenarios, to help MSP partners size not only the systems on which i-Vertix solutions are running, but also related storage and network bandwidth.

| Customer vertical | Flow rate: average fps | Log & Netflow retention period | Log & Netflow rollup period | N. Log Collector | N. Data Node | Server and storage specifications |
|---|---|---|---|---|---|---|
| Food | 500 eps | 2 weeks – 6 months | No rollup | 1 | 1 | 400m documents, 300Gb Data<br>8 core, 12GB RAM |
| GOV | 550 eps | 2 weeks – 6 months | No rollup | 1 | 1 | 2.000m documents<br>900GB Data<br>8vCPU, 32GB RAM |
| Manufacturing | 500 eps | 4 weeks – 6 months | No rollup | 1 | 1 | 8 Core, 600m documents (DNS, Win, Syslog, Web Server)<br>600GB |
| MSSP | 2.500 eps | 4 weeks (syslog, win), 6 months garante privacy | No rollup | 133 | 5 | 8.000m documents, > 4TB data |
| System Integrator | 9.000 eps | 2 weeks | 6 months | 1 | 1 | 3 TB data (80 Palo Alto FW Cef, > 300 Win) |
| GOV | 5.000 eps netflow<br>3.000 eps syslog, win<br>Users: ~5.500 | 2-4 weeks<br>6 months garante privacy | No rollup | 2 | 3 | 7.500m Docs, 3,3TB<br>Data Node: 2 physical hosts |