Sentinel

System Requirements

March 2020

NetIQ

Legal Notice

Contents

About this Book and the Library

The *System Requirements* document lists the hardware and software requirements for Sentinel and Sentinel Agent Manager.

Intended Audience

This guide is intended for Sentinel administrators and consultants.

Other Information in the Library

The library provides the following information resources:

Installation and Configuration Guide

Provides an introduction to Sentinel and explains how to install and configure Sentinel.

Administration Guide

Provides the administration information and tasks required to manage a Sentinel deployment.

User Guide

Provides conceptual information about Sentinel. This book also provides an overview of the user interfaces and step-by-step guidance for many tasks.

1 Product Requirements for Sentinel

Software Requirements

Operating Systems and Platforms for Sentinel Server

IMPORTANT:

The operating system for the Sentinel server must include at least the Base Server components of the SLES server or the RHEL server. Sentinel also requires certain RPMs to be available in the operating system. For more information about the required RPMs, see the Installation Checklist before you install Sentinel.

Software	Software
Sentinel Server, Collector Manager, or Correlation Engine	Sentinel runs on x86_64-based hardware and operating systems. It can run in Standard and FIPS 140-2 modes: ◆ SUSE Linux Enterprise Server (SLES) 15 SP1 64-bit ◆ SUSE Linux Enterprise Server 15 64-bit ◆ SUSE Linux Enterprise Server 12 SP4 64-bit (for both traditional and appliance installations) ◆ SUSE Linux Enterprise Server 12 SP3 64-bit (for both traditional and appliance installations) ◆ SUSE Linux Enterprise Server 12 SP2 64-bit ◆ Red Hat Enterprise Linux Server (RHEL) 8.1 64-bit ◆ Red Hat Enterprise Linux Server 8 64-bit ◆ Red Hat Enterprise Linux Server 7.7 64-bit **IMPORTANT:** Do not upgrade from RHEL version 7.7 to version 8.x because this is not supported by Red Hat. For more information, see the RedHat documentation. ◆ Red Hat Enterprise Linux Server 7.6 64-bit ◆ Red Hat Enterprise Linux Server 7.5 64-bit ◆ Red Hat Enterprise Linux Server 7.4 64-bit ◆ Red Hat Enterprise Linux Server 7.3 64-bit ◆ Red Hat Enterprise Linux Server 7.2 64-bit ◆ Red Hat Enterprise Linux Server 6.10 64-bit ◆ Red Hat Enterprise Linux Server 6.9 64-bit ◆ Red Hat Enterprise Linux Server 6.8 64-bit ◆ Red Hat Enterprise Linux Server 6.7 64-bit

Software	Software
Sentinel Server Software Appliance (includes SLES 12 SP3 operating system)	◆ ISO appliance **IMPORTANT:** For the ISO appliance to work properly, you must disable the EFI BIOS and use the Legacy BIOS. ◆ VMWare ESX 6.7 ◆ VMWare ESX 6.5 ◆ VMware ESX 6.0 ◆ VMware ESX 5.5 ◆ Hyper-V Server 2016 ◆ Hyper-V Server 2012 R2 (via DVD ISO) ◆ Hardware without a pre-installed operating system (via DVD ISO) ◆ OVF appliance ◆ VMWare ESX 6.7 ◆ VMWare ESX 6.5 ◆ VMware ESX 6.0 ◆ VMware ESX 5.5
Data indexing	◆ Elasticsearch 5.6.13 Download URL: https://www.elastic.co/downloads/past-releases/elasticsearch-5-6-13

Notes:

- Sentinel is certified on ext3 (SUSE), ext4 (Red Hat), and XFS file systems.
- Sentinel is not supported if the operating system is in FIPS mode.
- Sentinel is not certified on Open Enterprise Server installs of SLES.
- For SLES operating systems, use SLES 12 SP2 or later for CDH and Elasticsearch because the installation of some scalable storage services is more streamlined on these versions. For instance, the Elasticsearch RPM installer used on SLES 12 SP2 or later makes the installation easier.

Data Synchronization Platforms

Sentinel includes a feature to synchronize data subsets and summaries to a data warehouse.

Feature	Runs On
Data Synchronization	◆ Microsoft SQL Server 2017 ◆ Microsoft SQL Server 2016 ◆ Microsoft SQL Server 2014 ◆ Microsoft SQL Server 2012 ◆ Microsoft SQL Server 2008 R2 ◆ Oracle Database 12c ◆ Oracle Database 11g ◆ PostgreSQL ◆ IBM DB2 ◆ Sybase

Client Software

- **Java** Java 1.8 is required to launch Solution Designer and Sentinel Control Center.
- **Browsers** The Sentinel interface is optimized for viewing at 1280 x 1024 or higher resolution in the following supported browsers:
  - Microsoft Edge
  - Google Chrome
  - Mozilla Firefox
  - Microsoft Internet Explorer 11

  Although not officially certified, other modern browsers are known to work reasonably well with the Sentinel interface.

System Requirements for Traditional Storage

This section provides sizing information based on the testing performed at NetIQ with the hardware available to us at the time of testing. Your results may vary based on details of the hardware available, the specific environment, the specific type of data processed, and other factors. It is likely that larger, more powerful hardware configurations exist that can handle a greater load, and for even greater scalability Sentinel is explicitly designed to support distributed processing across multiple systems. If your environment is at all complex, contact NetIQ Consulting Services or any of the Sentinel partners prior to finalizing your Sentinel architecture as they have additional spreadsheets and tools to calculate architectural constraints.

System Requirements for Sentinel

NOTE

- All-in-one configurations put all the varied processing loads (data collection, processing, analysis, user interface, search, etc) into one server rather than distributing it across multiple servers within the system. While an all-in-one configuration can work well for a smaller-scale environment that does not make heavy simultaneous use of all system features, the competing loads can potentially cause issues if the system is under stress (which is sometimes the case

exactly when you need it most). Sentinel will prioritize critical functions such as data collection and storage, but (for example) UI performance may suffer. For this reason, you should deploy remote Collector Managers and/or Correlation Engines in most environments.

- You can use Intel Hyper-Threading Technology (Intel HT Technology) with the Sentinel server to positively impact the load the system can handle. The following table specifies the scenarios in which Intel HT Technology was used in testing.

Similarly, you should enable multithreading on Collector Managers. You can configure a Collector instance to use multiple threads, which allows the Collector to process a higher number of events per second. To configure the number of threads, in the Edit Collector dialog box, click the Configure Collector tab. Set Number of Threads to the number of threads you want to use. With this feature, a single 8-core Collector Manager can process 10K EPS. However, the test results listed below do not include multithreading on Collector Manager.

NOTE: The CPU and memory resources for a Collector Manager are subject to change depending on the EPS and the number of Collectors. Therefore, you should use virtual machines for Collector Managers.

Hardware Requirements

- "System Requirements for Elasticsearch" on page 23
- "Elasticsearch Cluster Nodes" on page 24

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
Total System Capacity					
Retained EPS Capability: The events per second rate processed by real-time components and retained in storage by the system.	100 EPS	3000 EPS	2500 EPS	21000 EPS	21000+ EPS

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
Operational EPS Capability: The total events per second rate received by the system from event sources. This includes data dropped by the system's intelligent filtering capability before being stored and is the number used for the purposes of EPS-based license compliance.	100 EPS	3000+ EPS	2500+ EPS	21000+ EPS	25000+ EPS
Sentinel Server Hardware					

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
CPU	Intel(R) Xeon(R) CPU E5420@ 2.50GHz (4 CPU cores), without Intel HT Technology	Two Intel(R) Xeon(R) CPU ES-2650 O@ 2.00GHz (4 core) CPUs (8 cores total), without Intel HT Technology	Two Intel(R) Xeon(R) CPU ES-2680 O@ 2.70GHz (6 cores per CPU; 12 cores total)	Two Intel(R) Xeon(R) CPU ES-2695 v2@ 2.40GHz(12 core) CPUs (24 cores total), with Intel HT Technology	Contact Micro Focus Services.
Primary Storage: Primary indexed event data optimized for fast retrieval.	500 GB 7.2k RPM drive	10 x 300 GB SAS 15k RPM (Hardware RAID 10)	6 x 146 GB SAS 10K RPM (RAID 10, stripe size 128k)	12 TB, 20 x 600 GB SAS 15k RPM (Hardware RAID 10, stripe size 128k)	
Secondary Storage: Secondary indexed event data optimized for storage efficiency. Includes a copy of the data in local storage but is only searched if the data is not found in primary storage.	For information about configuring secondary storage, see Configuring Secondary Storage Locations in the Sentinel Administration Guide.				
Memory	4 GB 8 GB, when Sentinel Agent Manager, NetIQ Secure Configuration Manager, or NetIQ Change Guardian are connected	24 GB		128 GB	
Remote Collector Manager #1 Hardware					

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
CPU	Not Applicable (Local Embedded CM Only)	Intel(R) Xeon(R) CPU E5-2650 O@ 2.00GHz, 4 cores (virtual machine)	Two Intel(R) Xeon(R) CPU E5-2680 O@ 2.70GHz (4 cores per CPU; 8 cores total)	Two Intel(R) Xeon(R) CPU E5-2695 v2@ 2.40GHz(8 core) CPUs 16 cores total)	Contact Micro Focus Services.
Storage		100 GB		250 GB	
Memory		4 GB	8 GB	24 GB	
Remote Collector Manager #2 Hardware					
CPU	Not Applicable			Two Intel(R) Xeon(R) CPU E5-2695 v2@ 2.40GHz(8 core) CPUs 16 cores total)	Contact Micro Focus Services.
Storage				250 GB	
Memory				24 GB	
Agent Manager Hardware					
CPU	Not Applicable (Agent-less collection only)		Two Intel Xeon 5140 @2.33 GHz (2 cores per CPU; 4 cores total)	Not Applicable	Contact Micro Focus Services.
Storage			4 x 300 GB SAS 10K RPM (RAID 10, stripe size 128k)		
Memory			16 GB		
Remote Correlation Engine Hardware					
CPU	Not Applicable (Local Embedded CE Only)	Intel(R) Xeon(R) CPU E5-2650 O@ 2.00GHz, 4 cores (virtual machine)	Intel(R) Xeon(R) CPU E5-2680 O@ 2.70GHz, 4 cores (virtual machine)	Two Intel(R) Xeon(R) CPU E5-2695 v2@ 2.40GHz, 4 core per CPU (8 cores total)	Contact Micro Focus Services.
Storage		100 GB			
Memory		8 GB		16 GB	
Data Collection					

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
Collector Manager (CM) Distribution: The number of event sources and events per second load placed on each Collector Manager. The filtered percentage indicates how many normalized events were filtered out immediatel yafter collection, without being stored or passed to analytic engines. Note that the non-normalized raw log data that the normalized events are based off of is not affected by filtering and is always stored. The Local Embedded CM is located on the Sentinel Server machine.	Local Embedded CM ◆ Event Sources: 101 ◆ EPS: 103 ◆ Filtered: 0%	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 2500 ◆ EPS: 3000	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 3500 ◆ EPS: 2500 ◆ Filtered: 0%	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 200 ◆ EPS: 10200 ◆ Filtered: 1% ◆ Raw Data Enabled Remote CM #2 ◆ Event Sources: 200 ◆ EPS: 10200 ◆ Filtered: 1% ◆ Raw Data Enabled	Contact Micro Focus Services.

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
Collectors Used	Oracle Solaris 2011.1r2 ◆ Sources: 100 ◆ EPS: 100 Juniper Netscreen 2011.1r2 ◆ Sources: 1 ◆ EPS: 3	Each Collector had its own Syslog server. Oracle Solaris 2011.1r2 ◆ Sources: 1000 ◆ EPS: 1500 Microsoft AD and Windows version 2011.1r4 ◆ Sources: 1000 ◆ EPS: 1000 Sourcefire Snort 2011.1 r1 ◆ Sources: 450 ◆ EPS: 500 Juniper Netscreen 2011.1r2 ◆ Sources: 20 ◆ EPS: 10	Agent Manager event source server 1 ◆ Sources: 3500 ◆ EPS: 2500 IBM i series 2011.1r5 ◆ Sources: 1500 ◆ EPS: 1000 NetIQ Agent Manager 2011.1r4 ◆ Sources: 1150 ◆ EPS: 500 NetIQ Unix Agent 2011.1r4 ◆ Sources: 1150 ◆ EPS: 500 Juniper Netscreen 2011.1r2 ◆ Sources: 2 ◆ EPS: 1	Each of the following Collectors had its own Syslog server, parsing at the following EPS rates ◆ Fortinet FortiGate 2011.1r3 ◆ RCM #1: 1700 ◆ RCM #2: 1700 ◆ Palo Alto Networks Firewall 2011.1r2 ◆ RCM #1: 1700 ◆ RCM #2: 1700 ◆ Dumballa Failsafe2011.1r1 ◆ RCM #1: 1700 ◆ RCM #2: 1700 ◆ McAfee Firewall Enterprise 2011.1r4 ◆ RCM #1: 1700 ◆ RCM #2: 1700 ◆ Microsoft Active Directory and Windows 2011.1r7 ◆ RCM #1:	Contact Micro Focus Services.

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
Total	◆ Event Sources: 101 ◆ EPS: 103 ◆ Filtered: 0%	◆ Event Sources:2 500 ◆ EPS: 3010 ◆ Filtered: 0%	◆ Event Sources: 3500 ◆ EPS: 2501 ◆ Filtered: 0%	◆ Event Sources:4 00 ◆ EPS: 20411 ◆ Filtered: 1%	Contact Micro Focus Services.
Data Storage					

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
How far into the past will users search for data on a regular basis?. Amount of locally cached data for higher search performanc e	7 days				Contact Micro Focus Services.
What percentag e of searches will be over data older than the number of days above? Impacts the amount of input/ output operations per second (IOPS) for local or network storage.	10%				
How far into the past must data be retained? Impacts how much disk space is required to retain all the data. If secondary storage is enabled, this impacts the size of secondary storage. Otherwise,	14 days				

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
User Activity					

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
How many users will be active at the same time, on average? Impacts the amount of !OPS for primary and secondary storage and other items.	1				Contact Micro Focus Services.
How many searches will an active user be performing at the same time, on average? Impacts the amount of !OPS for primary and secondary storage.	1 100M events per search	1 300M events per search	Not tested with search or reporting load	1 2B events per search	
How many reports will an active user be running at the same time, on average? Impacts the amount of IOPS for primary and secondary storage.	1 200k events per report	1 500k events per report		1 600k events per report	
How many real-time alert views will be running at the same time, on average?	3 (whenever)				

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
Analytics					

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
What percentage of the event data is relevant to correlation rules? Amount of data the Correlation Engine will process.	100% (out of the box) (3 correlations per second)		100% (out of the box) (1 correlation per second)	100% (out of the box) (10 correlations per second)	Contact Micro Focus Services.
What percentage of the event data is relevant to Event Visualization? (Data indexed to Elasticsearch)	100% (out of the box)				
What percentage of the event data is relevant to IP Flows? (IP Flow events indexed to Elasticsearch)	3% (500 IP Flow events per second)		5% (100 IP Flow events per second)	10% (10 IP Flow events per second)	
How many source IPs or source host names are relevant to generic hostname resolution service? (Number of DNS lookups impacting the CPU utilization of the Collector Manager)	200			100	

Category	Demo All-in-One (Not intended for production)	Medium Distributed Agentless Data Collection	Medium Distributed Agent-based Data Collection	Large Distributed Agent-less Data Collection	Extra Large
How many events are relevant to threat intelligence feeds?	10 EPS				
High Availability	Not Used				
Notes: Notable functionality disabled or warnings of what happens when exceeding the system load described above.				Increasing Retained EPS will eventually cause instability in this system configuration.	

System Requirements for Elasticsearch

You must install and set up Elasticsearch nodes in a cluster mode if you want to use the Event Visualizations feature. For more information, see the “Configuring the Visualization Data Store” in the Sentinel Installation and Configuration Guide.

You must set up Elasticsearch as recommended in the following table:

Component	Recommendation
Indexing Node Data Storage	◆ Operating system and application binaries and configuration ◆ Fault tolerant RAID ◆ Data Storage ◆ Disks in JBOD (Just a Bunch Of Disks) configuration ◆ SSD or 15000 RPM SATA
CPU	Intel(R) Xeon(R) CPU ES-2695 v2@ 2.40GHz

Elasticsearch Cluster Nodes

	Elasticsearch Nodes	CPU per Node	Memory (GB) per Node	Disks per Node
100 EPS	1 data node + 1 master node (ES node in Sentinel)	4	4	2
3000 EPS	2 data nodes + 1 master node (ES node in Sentinel)	8	24	3
20000 EPS	4 data nodes + 1 master node (ES node in Sentinel)	8	32	4

System Requirements for Scalable Storage

This section provides sizing information based on the testing performed at Micro Focus with the hardware available to us at the time of testing. Your results may vary based on details of the hardware available, the specific environment, the specific type of data processed, and other factors. It is likely that larger, more powerful hardware configurations exist that can handle a greater load, and for even greater scalability Sentinel is explicitly designed to support distributed processing across multiple systems. If your environment is at all complex, consult with Micro Focus Consulting Services or any of the Sentinel partners prior to finalizing your Sentinel architecture as they have additional spreadsheets and tools to calculate architectural constraints.

The hardware requirements provided on this page are applicable only for Sentinel with the scalable storage option enabled. To perform functions that are not available with the scalable storage option, such as anomaly detection, you must install separate instances of Sentinel with traditional storage and route the specific event data to it by using Sentinel Link. In such a case, you must set up additional hardware for the traditional storage Sentinel servers based on the EPS you plan to filter and forward to the traditional storage Sentinel servers. For more information, see the hardware requirements for traditional storage.

- “Node Types” on page 24
- “System Sizing Information” on page 26

Node Types

- Grouping of services in the node types below is a suggested grouping aimed at achieving the following goals:
  - Minimize the number of nodes so that it is easier to manage.
  - Achieve good data reliability to avoid data loss even under typical system failure scenarios. Node redundancy is necessary to achieve this goal.
  - Isolate services whose performance profile would conflict with each other under load. For example, both Elasticsearch and Kafka make use of operating system file system caching. This would result in conflict with each other and with other memory intensive services running on the same operating system.
- Other arrangements of services can work very well if the goals of a scenario are different.

- If the appliance installer is selected for the SSDM server, it must be run on a separate node from the Master since installing CDH services on the appliance is not recommended for maintainability purposes.
- Each node can be a virtual machine or a bare-metal machine. For data reliability reasons, redundant nodes must be placed on separate bare-metal hardware. For example, if all nodes are virtual machines, then the minimum nodes for the Production System described below requires 3 bare-metal hosts with one virtual Worker node, one virtual Messaging node and one Indexing node on each bare-metal host.

Node Type	Services	Minimum Nodes for Production System
Worker	All Worker nodes include: ◆ HDFS DataNode ◆ HBase RegionServer ◆ YARN NodeManager Only 2 Worker nodes need: ZooKeeper Server	3
Messaging	All Messaging nodes include: Kafka Broker Only 1 Messaging node needs: ◆ HDFS SecondaryNameNode ◆ YARN Resource Manager (Standby)	3
Indexing	Elasticsearch	2
Master	◆ Cloudera Management Services ◆ HBase Master ◆ HDFS NameNode ◆ YARN ResourceManager ◆ YARN JobHistory Server ◆ Spark History Server ◆ ZooKeeper Server ◆ Sentinel Scalable Data Manager (SSDM) Server (traditional installer)	1

Hardware Recommendations

Component	Recommendation
CPU	Intel(R) Xeon(R) CPU ES-2695 v2@ 2.40GHz
Master Node Data Storage	Reliable storage (such as fault tolerant RAID)

Component	Recommendation
Worker and Messaging Node Data Storage	Operating system and application binaries and configuration: Fault tolerant RAID Data Storage ◆ Disks in JBOD (Just a Bunch Of Disks) configuration ◆ 7200 RPM SATA
Indexing Node Data Storage	Operating system and application binaries and configuration: Fault tolerant RAID Data Storage ◆ Disks in JBOD (Just a Bunch Of Disks) configuration ◆ SSD or 15000 RPM SATA
Network Technology	Bonded Gigabit Ethernet or 10 Gigabit Ethernet

System Sizing Information

8K EPS (Filtered data to Remote Correlation Engine, 4 Spark Jobs)

Layout of Services

Node Type	VM Nodes	vCPU per Node	vMemory (GB) per Node	Disks per Node
Worker	5	12	24	4
Messaging	4	4	32	3
Indexing	2	8	12	4
Master	1	16	32	1

Hardware Layout

Virtual Machines	Hardware Nodes	CPU per Node	Memory (GB) per Node	Disks per Node
1x Worker 1x Messaging 1x Indexing	1 (ESX1)	24	128, but only 68 will be used	12
1x Worker 1x Messaging 1x Indexing 1x Master	1 (ESX2)	24	128, but only 100 will be used	12
1x Worker 1x Messaging	1 (ESX3)	16	128, but only 58 will be used	9

CDH Setup Detail

		Components	RAM	CPU	HDD
Master (1)	Node1	◆ HBase Master ◆ HDFS NameNode ◆ Cloudera Management Service Alert Publisher ◆ Cloudera Management Service Event Server ◆ Cloudera Management Service Host Monitor ◆ Cloudera Management Service Reports Manager ◆ Cloudera Management Service Service Monitor ◆ Spark History Server ◆ YARN (MR2 Included) JobHistory Server ◆ YARN (MR2 Included) ResourceManager ◆ ZooKeeper Server ◆ SSDM Server	32	12	500 GB

		Components	RAM	CPU	HDD
Messaging (3)	Node2	Kafka Broker	32	4	//dev/sda4 408G /kafka1 /dev/sdb1 500G /kafka2 /dev/sdc1 500G /kafka3
	Node3		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
	Node4		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
Worker (3)	Node5	◆ HDFS DataNode ◆ HBase RegionServer ◆ YARN NodeManager	24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 6		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 7		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4

		Components	RAM	CPU	HDD
Indexing (2)	Node8	Elasticsearch	12	8	/dev/sda4 500G /es1 /dev/sdb1 500G /es2 /dev/sdd1 500G /es4 /dev/sdc1 500G /es3
	Node9		12	8	/dev/sda4 500G /es1 /dev/sdb1 500G /es2 /dev/sdd1 500G /es4 /dev/sdc1 500G /es3

Sentinel Components

	Number of instances	CPU	Memory	Disk Space
Collector Manager	2	8 cores	8 GB	100 GB free space (fixed storage)
Correlation Engine	1			

10K EPS (Filtered data to Remote Correlation Engine, 4 Spark Jobs)

Layout of Services

Node Type	VM Nodes	vCPU per Node	vMemory (GB) per Node	Disks per Node
Worker	5	12	24	4
Messaging	4	4	32	3
Indexing	2	8	12	4
Master	1	16	32	1

Hardware Layout

Virtual Machines	Hardware Nodes	CPU per Node	Memory (GB) per Node	Disks per Node
1x Worker 1x Messaging 1x Indexing	1 (ESX1)	24	128, but only 68 will be used	12
1x Worker 1x Messaging 1x Indexing 1x Master	1 (ESX2)	24	128, but only 100 will be used	12
1x Worker 1x Messaging	1 (ESX3)	16	128, but only 56 will be used	9
2x Worker 1x Messaging	1 (ESX4)	24	128, but only 24 will be used	12

CDH Setup Detail

		Components	RAM	CPU	HDD
Master (1)	Node1	◆ HBase Master ◆ HDFS NameNode ◆ Cloudera Management Service Alert Publisher ◆ Cloudera Management Service Event Server ◆ Cloudera Management Service Host Monitor ◆ Cloudera Management Service Reports Manager ◆ Cloudera Management Service Service Monitor ◆ Spark History Server ◆ YARN (MR2 Included) JobHistory Server ◆ YARN (MR2 Included) ResourceManager ◆ ZooKeeper Server ◆ SSDM Server	32	12	500 GB

		Components	RAM	CPU	HDD
Messaging (3)	Node2	Kafka Broker	32	4	/dev/sda4 408G /kafka1 /dev/sdb1 500G /kafka2 /dev/sdc1 500G /kafka3
	Node3		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
	Node4		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
Worker (4)	Node5	◆ HDFS DataNode ◆ HBase RegionServer ◆ YARN NodeManager	32	4	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 6		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 7		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node8		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4

		Components	RAM	CPU	HDD
Indexing (2)	Node9	Elasticsearch	12	8	/dev/sda4 500G /es1 /dev/sdb1 500G /es2 /dev/sdd1 500G /es4 /dev/sdc1 500G /es3
	Node10		12	8	/dev/sda4 500G /es1 /dev/sdb1 500G /es2 /dev/sdd1 500G /es4 /dev/sdc1 500G /es3

Sentinel Components

	Number of instances	CPU	Memory	Disk Space
Collector Manager	2	8 cores	8 GB	100 GB free space (fixed storage)
Correlation Engine	1			

11K EPS (Filtered data to Remote Correlation Engine, 4 Spark Jobs)

Layout of Services

Node Type	VM Nodes	vCPU per Node	vMemory (GB) per Node	Disks per Node
Worker	5	12	24	4
Messaging	4	4	32	3
Indexing	2	8	12	4
Master	1	16	32	1

Hardware Layout

Virtual Machines	Hardware Nodes	CPU per Node	Memory (GB) per Node	Disks per Node
1x Worker 1x Messaging 1x Indexing	1 (ESX1)	24	128, but only 68 will be used	12
1x Worker 1x Messaging 1x Indexing 1x Master	1 (ESX2)	24	128, but only 100 will be used	12
1x Worker 1x Messaging	1 (ESX3)	16	128, but only 56 will be used	9
2x Worker 1x Messaging	1 (ESX4)	24	128, but only 48 will be used	12

CDH Setup Detail

		Components	RAM	CPU	HDD
Master (1)	Node1	◆ HBase Master ◆ HDFS NameNode ◆ Cloudera Management Service Alert Publisher ◆ Cloudera Management Service Event Server ◆ Cloudera Management Service Host Monitor ◆ Cloudera Management Service Reports Manager ◆ Cloudera Management Service Service Monitor ◆ Spark History Server ◆ YARN (MR2 Included) JobHistory Server ◆ YARN (MR2 Included) ResourceManager ◆ ZooKeeper Server ◆ SSDM Server	32	12	500 GB

		Components	RAM	CPU	HDD
Messaging (3)	Node2	Kafka Broker	32	4	/dev/sda4 408G /kafka1 /dev/sdb1 500G /kafka2 /dev/sdc1 500G /kafka3
	Node3		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
	Node4		32	4	//dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3

		Components	RAM	CPU	HDD
Worker (5)	Node5	◆ HDFS DataNode ◆ HBase RegionServer ◆ YARN NodeManager	32	4	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 6		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 7		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node8		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node9		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4

Sentinel Components

	Number of instances	CPU	Memory	Disk Space
Collector Manager	3	8 cores	8 GB	100 GB free space (fixed storage)
Correlation Engine	1			

12K EPS (Filtered data to Remote Correlation Engine, 4 Spark Jobs)

Layout of Services

Node Type	VM Nodes	vCPU per Node	vMemory (GB) per Node	Disks per Node
Worker	5	12	24	4
Messaging	4	4	32	3
Indexing	2	8	12	4
Master	1	16	32	1

Hardware Layout

Virtual Machines	Hardware Nodes	CPU per Node	Memory (GB) per Node	Disks per Node
1x Worker 1x Messaging 1x Indexing	1 (ESX1)	24	128, but only 68 will be used	12
1x Worker 1x Messaging 1x Indexing 1x Master	1 (ESX2)	24	128, but only 100 will be used	12
1x Worker 1x Messaging	1 (ESX3)	16	128, but only 56 will be used	9
2x Worker 1x Messaging	1 (ESX4)	24	128, but only 80 will be used	12

CDH Setup Detail

		Components	RAM	CPU	HDD
Master/ Manager/ SSDM Server (1)	Node1	◆ HBase Master ◆ HDFS NameNode ◆ Cloudera Management Service Alert Publisher ◆ Cloudera Management Service Event Server ◆ Cloudera Management Service Host Monitor ◆ Cloudera Management Service Reports Manager ◆ Cloudera Management Service Service Monitor ◆ Spark History Server ◆ YARN (MR2 Included) JobHistory Server ◆ YARN (MR2 Included) ResourceManager ◆ ZooKeeper Server ◆ SSDM Server	32	12	500 GB

		Components	RAM	CPU	HDD
Messaging (4)	Node2	Kafka Broker	32	4	/dev/sda4 408G /kafka1 /dev/sdb1 500G /kafka2 /dev/sdc1 500G /kafka3
	Node3		32	4	dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
	Node4		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3
	Node5		32	4	/dev/sdb1 500G /kafka1 /dev/sdc1 500G /kafka2 /dev/sdd1 500G /kafka3

		Components	RAM	CPU	HDD
Worker (5)	Node 6	◆ HDFS DataNode ◆ HBase RegionServer ◆ YARN NodeManager	24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node 7		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node8		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node9		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G /hdfs4
	Node10		24	12	/dfs/dn 500G /dev/sdb1 500G /hdfs1 /dev/sdc1 500G /hdfs3 /dev/sdd1 500G

Sentinel Components

	Number of instances	CPU	Memory	Disk Space
Collector Manager	3	8 cores	8 GB	100 GB free space (fixed storage)
Correlation Engine	1			

SSDM Components	Number of instances	Disk Space
SSDM Server	1	100 GB free space (fixed storage)
Collector Managers		
Correlation Engines		

Storage Estimations (5K EPS)

CDH components	Total data nodes	Number of disks (per node)	Cluster Storage per day (all replicas) (in GB)	Cluster Storage per day (single replica) (in GB)	Notes
Messaging	3	3	900 (max 7 days retention by default) Replicas: 3X	300 (max 7 days retention by default)	◆ Event Analytics topic (4 GB for 10min retention) ◆ Events and raw data
Worker	3	4	750 Replicas: 3X	250	Events and raw data
Indexing	2	4	450 Replicas: 2X	225	Events (Default event fields are indexed)
Master	1	1	N/A	100 GB free space (fixed storage)	

Performance Test Details

Data Collection

	8000 EPS	10000 EPS	11000 EPS	12000 EPS
Collector Manager (CM) Distribution: The number of event sources and events per second load placed on each Collector Manager.	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 175 ◆ EPS: 5100 ◆ Filtered: 0% ◆ Raw Data Enabled Remote CM #2 ◆ Event Sources: 175 ◆ EPS: 3000 ◆ Filtered: 0% ◆ Raw Data Enabled	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 175 ◆ EPS: 5100 ◆ Filtered: 0% ◆ Raw Data Enabled Remote CM #2 ◆ Event Sources: 175 ◆ EPS: 5100 ◆ Filtered: 0% ◆ Raw Data Enabled	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 175 ◆ EPS: 5100 ◆ Filtered: 0% ◆ Raw Data Enabled Remote CM #2 ◆ Event Sources: 175 ◆ EPS: 5100 ◆ Filtered: 0% ◆ Raw Data Enabled Remote CM #3 ◆ Event Sources: 175 ◆ EPS: 1000 ◆ Filtered: 0% ◆ Raw Data Enabled	Local Embedded CM ◆ Not Used Remote CM #1 ◆ Event Sources: 175 ◆ EPS: 5100 ◆ Filtered: 0% ◆ Raw Data Enabled Remote CM #2 ◆ Event Sources: 175 ◆ EPS: 3000 ◆ Filtered: 0% ◆ Raw Data Enabled Remote CM #3 ◆ Event Sources: 175 ◆ EPS: 2000 ◆ Filtered: 0% ◆ Raw Data Enabled

	8000 EPS	10000 EPS	11000 EPS	12000 EPS
Collectors Used	Each of the following Collectors had their own Syslog server, parsing at the following EPS rates: ◆ Fortinet FortiGate 2011.1r3 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ Palo Alto Networks Firewall 2011.1r2 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ Dumballa Failsafe 2011.1r1 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ McAfee Firewall Enterprise 2011.1r4 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ Microsoft Active Directory and Windows 2011.1r7 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ Oracle Solaris 2011.1r2 ◆ RCM #1: 850 ◆ RCM #2: 500	Each of the following Collectors had their own Syslog server, parsing at the following EPS rates: ◆ Fortinet FortiGate 2011.1r3 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ Palo Alto Networks Firewall 2011.1r2 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ Dumballa Failsafe 2011.1r1 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ McAfee Firewall Enterprise 2011.1r4 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ Microsoft Active Directory and Windows 2011.1r7 ◆ RCM #1: 850 ◆ RCM #2: 85 ◆ Oracle Solaris 2011.1r2 ◆ RCM #1: 850 ◆ RCM #2: 850	Each of the following Collectors had their own Syslog server, parsing at the following EPS rates: ◆ Fortinet FortiGate 2011.1r3 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ RCM #3: 150 ◆ Palo Alto Networks Firewall 2011.1r2 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ RCM #3: 150 ◆ Dumballa Failsafe 2011.1r1 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ RCM #3: 150 ◆ McAfee Firewall Enterprise 2011.1r4 ◆ RCM #1: 850 ◆ RCM #2: 850 ◆ RCM #3: 150 ◆ Microsoft Active Directory and Windows 2011.1r7 ◆ RCM #1: 850 ◆ RCM #2: 850	Each of the following Collectors had their own Syslog server, parsing at the following EPS rates: ◆ Fortinet FortiGate 2011.1r3 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ RCM #3: 200 ◆ Palo Alto Networks Firewall 2011.1r2 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ RCM #3: 200 ◆ Dumballa Failsafe 2011.1r1 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ RCM #3: 200 ◆ McAfee Firewall Enterprise 2011.1r4 ◆ RCM #1: 850 ◆ RCM #2: 500 ◆ RCM #3: 200 ◆ Microsoft Active Directory and Windows 2011.1r7 ◆ RCM #1: 850 ◆ RCM #2: 500

	8000 EPS	10000 EPS	11000 EPS	12000 EPS
Remote Correlation Engine	Remote CE #1 ◆ EPS utilization: 80% ◆ CR fire rate: 1%	Remote CE #1 ◆ EPS utilization: 80% ◆ CR fire rate: 1%	Remote CE #1 ◆ EPS utilization: 80% ◆ CR fire rate: 1%	Remote CE #1 ◆ EPS utilization: 80% ◆ CR fire rate: 1%
Total	◆ Event Sources: 350 ◆ EPS: 8000 ◆ Filtered: 0%	◆ Event Sources: 525 ◆ EPS: 10000 ◆ Filtered: 0%	◆ Event Sources: 525 ◆ EPS: 11000 ◆ Filtered: 0%	◆ Event Sources: 525 ◆ EPS: 12000 ◆ Filtered: 0%

Data Storage

How far into the past will users search for data on a regular basis? Amount of locally cached data in Elasticsearch for higher search performance	7 days
What percentage of searches will be over data older than the number of days above? Impacts the amount of input/output operations per second (IOPS) for Elasticsearch	10%
How far into the past must data be retained? Impacts how much disk space is needed to retain all of the data. This also impacts the size of HBase and Elasticsearch.	90 days

User Activity

How many simultaneous users will be accessing the visualization dashboards at the same time, on average?	3
How many visualization dashboards will be running at the same time, on average?	3
How many widgets per visualization dashboard will be running at the same time, on average?	10
How many simultaneous visualization searches will be running at the same time, on average?	3
How many users will be accessing the Threat Response dashboard at the same time, on average?	3
How many Threat Response dashboards will be running at the same time, on average?	3

How many alert widgets (alert search queries) per dashboard will be running at the same time, on average?	2
How many real-time alert views will be running at the same time, on average?	3

Analytics

What percentage of the event data is relevant to correlation rules? Amount of data the Correlation Engine will process.	100% (out of the box) (10 correlations per second)
How many simple correlation rules (filter/trigger only) will be used? Impacts the CPU utilization of the Correlation Engine.	114 (out of the box)
How many complex correlation rules will be used? Impacts the CPU and memory utilization of the Correlation Engine.	1 (out of the box)
Correlation Engine (CE) Distribution	Local Embedded CE (75 rules) Remote CE (40 rules)
How many alerts will be created?	30 per minute

2 Product Requirements for Sentinel Agent Manager

Software Requirements for Sentinel Agent Manager

Software	Runs On
Sentinel Agent Manager Central Computer and Sentinel Agent Manager Console	◆ Microsoft Windows Server 2019 ◆ Microsoft Windows Server 2016 ◆ Microsoft Windows Server 2012 R2 ◆ Microsoft Windows Server 2012
Sentinel Agent Manager Database	◆ Microsoft SQL Server 2017 ◆ Microsoft SQL Server 2016 ◆ Microsoft SQL Server 2014 ◆ Microsoft SQL Server 2012 ◆ Microsoft SQL Server 2012 Express

System Requirements for Sentinel Agent Manager

NOTE: These are minimum recommendations.

	Requirements			
Sentinel Agent Manager Component	**Processor**	**Disk Space**	**Memory**	**Software**
Sentinel Agent Manager Central Computer	Dual processor dual-core AMD/Intel configuration	Depends on the event load estimated for your environment.	4 GB	◆ Microsoft Message Queuing (MSMQ) 3.0 ◆ Microsoft .NET Framework 2.0 Service Pack 1 or later ◆ Microsoft Visual C++ 2005 Service Pack 1 Redistributable Package ◆ Microsoft Core XML Services (MSXML) 6.0 or later
Sentinel Agent Manager Database	Dual processor dual-core AMD/Intel configuration Quad processors recommended in environments expecting more than one million total events per day.	100 GB	4 GB	See "Software Requirements for Sentinel Agent Manager" on page 45.
Sentinel Agent Manager Agent	500 MHz Intel Pentium or equivalent	100 MB	40 MB **NOTE:** The amount of memory usage varies and depends on the modules you have installed and the products you are monitoring	Microsoft Visual C++ 2005 Service Pack 1 Redistributable Package

3 Event Sources

Sentinel supports a wide variety of endpoint event sources that can deliver security and operational events to Sentinel for processing along with other types of contextual data using modular, pluggable components. Sentinel provides both agents and agent-less options. For more information about the specific endpoints monitored by these agents, follow the links below.

Module/Plug-in	Compatible Versions and Endpoints
Security Agent for UNIX	◆ Security Agent for UNIX 7.6.2 ◆ Security Agent for UNIX 7.6.1 ◆ Security Agent for UNIX 7.6
Windows Agent (available via Sentinel Agent Manager)	◆ Microsoft Windows Server 2016 ◆ Microsoft Windows Server 2012 R2 ◆ Microsoft Windows Server 2012 ◆ Microsoft Windows 10
Agentless data collection	Sentinel Collectors
ArcSight SmartConnectors	◆ AirMagnet Enterprise Syslog ◆ Amazon Web Services CloudTrail ◆ ArcSight CEF Cisco FireSIGHT Syslog ◆ ArcSight Common Event Format Hadoop ◆ Barracuda Email Security Gateway Syslog ◆ Box ◆ HPE Aruba Mobility Controller Syslog ◆ IP Flow (Netflow/J-Flow) ◆ IP Flow Information Export (IPFIX) ◆ Kaspersky DB ◆ Microsoft Office 365 ◆ sFlow ◆ Vormetric CoreGuard Syslog