SIEMONSTER VERSION 3 HIGH LEVEL DESIGN

Document Version	1.4
Lead Designer	Chris Rock / James Bycroft
Authors	Chris Rock / James Bycroft
Last Change Date	Monday, 15 January 2018

Contact information

For more information on this document please contact:

Name	Chris Rock
Position	CEO
E-mail	info@siemonster.com

The following people can also be contacted in relation to this document:

Name	Position	Email
Chris Rock	Solution Lead	chris@siemonster.com
James Bycroft	Lead Architect	james@siemonster.com
Christopher Heuvel	Senior Engineer	christopher@siemonster.com

Glossary

The following terms and acronyms are used in this document:

Term	Definition
SOC	Security Operations Centre
SIEM	Security Information and Event Management
OS	Open Source
AD	Active Directory
TMG	Microsoft Threat Management Gateway
ASA	Cisco ASA IPS / Firewall
MSSP	Managed Security Service Provider
IIS	Internet Information Services
AWS	Amazon Web Services
ES	ElasticSearch
VM	Virtual Machine
VPN	Virtual Private Network
ELK	ElasticSearch Logstash and Kibana Open Source Data analytics
VPC	Virtual Private Cloud
RDS	Relational Database Service
EFS	Elastic File System, scalable file storage for use with Amazon EC2 instances
NFS	Network File System, It is a client/server system that allows users to access files across a network and treat them as if they resided in a local file directory
Cloud Formation	Service for rolling out AWS infrastructure
OS	Operating System
Rancher	Open source software platform that enables organizations to run containers in production
S3	Amazon S3 provides storage through web services interfaces
IAM	Identity and Access Management, Security of information and resources, controlling access to information and resources, managing levels of access for different users in system
Hydra	Log collection server on client site which securely sends data collected to the SIEMonster MSSP for processing and storage
Glacier	Amazon Glacier is a secure, durable, and extremely low-cost cloud storage service for data archiving and long-term backup
Prometheus	SIEMonster tool for system monitoring
HoneyNet	Honeypot platform, which is based on the well-established honeypots glastopf, kippo, honeytrap and dionaea, the network IDS/IPS suricata, elasticsearch-logstash-kibana, ewsposter and docker.
Suricata	Suricata is a free and open source, mature, fast and robust network threat detection engine. The Suricata engine is capable of real time intrusion detection (IDS), inline intrusion prevention (IPS), network security monitoring (NSM)
Cowrie	Cowrie is a medium interaction SSH and Telnet honeypot designed to log brute force attacks and the shell interaction performed by the attacker.

Glastopf	Glastopf is a Python web application honeypot founded by Lukas Rist.
Dionaea	Catches bugs, dionaea intention is to trap malware exploiting vulnerabilities exposed by services offered to a network, the ultimate goal is gaining a copy of the malware.
Elasticpot	Elasticpot elasticpot is a simple elastic search honeypot.
Honeytrap	Honeytrap is a low-interaction honeypot daemon for observing attacks against network services. In contrast to other honeypots, which often focus on malware collection, honeytrap aims for catching the initial exploit – It collects and further processes attack traces.
Kippo	Kippo is a medium interaction SSH honeypot designed to log brute force attacks and, most importantly, the entire shell interaction performed by the attacker.
EWSposter	EWSposter is a python application that collects information from multiple honeypot sources and posts it to central collection services like the DTAG early warning system and hpfeeds.
Grafana	Grafana is an open source, feature rich metrics dashboard and graph editor for Graphite, Elasticsearch, OpenTSDB, Prometheus and InfluxDB.
Prometheus	Prometheus, a Cloud Native Computing Foundation project, is a systems and service monitoring system. It collects metrics from configured targets at given intervals, evaluates rule expressions, displays the results, and can trigger alerts if some condition is observed to be true.
Netdata	Netdata is a system for distributed real-time performance and health monitoring.
Kopf/ES-Maintenance	kopf is a simple web administration tool for elasticsearch written in JavaScript + AngularJS + jQuery + Twitter bootstrap. It offers an easy way of performing common tasks on an elasticsearch cluster. Not every single API is covered by this plugin, but it does offer a REST client which allows you to explore the full potential of the ElasticSearch API.
LogTrail	Log viewer plugin for kabana
Dradis	Open Framework Vulnerability Assessment Import tool

Table of contents

1 AUTHORS PREFACE

In 2015, one of our corporate clients told us of their frustrations with the exorbitant licensing costs of commercial Security Information and Events Management (SIEM) products. The customer light heartedly asked whether we could build them an open source SIEM to get rid of these annual license fees. We thought that was a great idea and set out so to develop a SIEM product for Managed Security Service Providers (MSSP's) and Security Professionals. This product is called SIEMonster.

SIEMonster Version 1 was released in late April of 2016 and a commercial release in November 2016. The release has been an astounding success without over 100,000 downloads of the product. We have assisted individuals and companies integrate SIEMonster into small medium and extra-large companies all around the world. SIEMonster with the help of the community and a team of developers have been working hard since the Version1 release incorporating what the community wanted to see in a SIEM as well as things we wanted to see in the next release.

Along the way we have signed up MSSP's from around the world who have contributed to the rollout of SIEMonster and in return they have assisted us with rollout scripts, ideas and things we hadn't even considered.

We are now proud to release the latest Version 3.0 Beta, and finalized in February 2018 for Alpha Release. We have added the following features to this release

- ELK Stack updated to version 5.5
- Built in Searchguard open source RBAC & encrypted node to node transport
- Wazuh HIDS system with Kibana plugin and OpenSCAP options & simplified agent registration process
- Simplified installation process for both Rancher Docker orchestration & SIEMonster web application
- All new dashboard with options for 2fa, site administration with user role based access and faster load times
- Built in parsers for most proprietary devices
- Preloaded Minemeld threat intel feeds integrated with log ingest out of the box
- COREOS stable 1576.5.0 (Fix CPU disclosure of kernel memory to user process (CVE-2017-5754, Meltdown) with NFS support for configuration centralization.

We have also automated correlation with Palo Alto MineMeld Open Source Threat Intelligence and added two factor authentication and easier rollouts.

The transition has now been completed to a full containerize all aspects of the SIEMonster application pool using the popular Docker system. This allows us to run on any hardware, cloud or operating system. It also provides the architecture for docker containers to be moved to other servers during downtime without affecting the SIEM.

We welcome you to try out our fully functional SIEM product, and if you wish to upgrade to our Premium version with Advanced Correlation and Reporting please contact sales@siemonster.com.

2 INTRODUCTION

SIEMonster Version 3 is built on the best open source components and custom develop from a wish list from the SIEMonster community. This document will cover the architecture, the features and the open source components that make up SIEMonster, so that all security professionals can run a SIEM in their organisations with no budget. If you would like more information about the architecture please see our High-Level Design.

SIEMonster is built on CoreOS, Docker with Rancher, Kubernetes orchestration. The product comes in Vbox, VMware, Bare-metal or Cloud install on AWS/Azure. SIEMonster can scale horizontally and vertically to support any enterprise client.

Some of these features include.

- OSINT from PaloAlto Minemeld.
- OSSEC Wazuh fork. Full integration with OSSEC Wazuh fork for Host Intrusion Detection and PCIDSS ruleset incorporated into Elastic.
- 411 demonstrated at DEFCON. Instant Incident Alerting via email or SMS or Console view via a secure portal and integration with “Slack”/”PagerDuty”/”Jira” using 411 Streams.
- Open Source AuditIT by Opmantek.
- Open Source Incident Response. Alerts maybe escalated as tickets to other operators or a whiteboard to show night shift analysts current issues.
- Elastalert, alerting on anomalies, spikes, or other patterns within Elasticsearch.
- Prometheus metric exporters with Prometheus AlertManager for system monitoring.
- Data Correlation UI, community rulesets and dashboards, community and open source free plugins that make the SIEM.
- Incorporate your existing Vulnerability Scans into the Dashboard, (OpenVAS, Nexpose, Metasploit, Burp, Nessus etc.) using Dradis.
- We have also developed and built-in LDAP integration, advanced correlation and two factor authentication.

2.1 HIGH LEVEL COMPONENTS

This solution includes the following high-level components.

SIEM & SIM / SEM – Is built to provide 24x7 Security Event collection, correlation and Incident response. Risk Identification, visual alerting, analysis and secondary email/SMS/Slack alerts to the operator.

Software Components – All open source components that are included in SIEMonster.

Hardware Components – Virtual hardware requirements for Virtual environment

Configuration – SIEMonster configuration, rulesets, architecture, equipment requirements, backups and maintenance.

Dashboards – Visual representation of configured alerts and risks in the environment.

Rules and Search – How to configure rules and run searches using Elastic Search

Incident Response Alerting and Ticketing – Ticketing system to record Incident Response, and documentation for security analysts and 411 for alerting and UI

Host Based Intrusion Detection Alerting – Integration with OSSEC Wazuh fork to allow HIDS alerting and pre-built rulesets including PCIDSS.

OSINT – Integration with Open Source Threat Intelligence for real time threats in the wild incorporating Palo Alto Minemeld.

Vulnerability Scanning – Using your existing vulnerability scanning tool incorporated into a world view dashboard revealing hot spots in your network using Dradis IDE

Reporting – Reporting snapshots of your SIEM via scheduled emails in PDF and Excel format using Skedler.

Plugin Development – How to develop your own plugins for specific equipment like SCADA or custom equipment or how to get help from the community of Kustodian to monitor specific equipment like Blast Furnaces, paint guns, robotic arms or production lines.

.

2.2 SCOPE

This document covers all the software and hardware infrastructure components for the Security Operations Centre SIEMonster product. Separate documents such as build guides, standard operating procedures, troubleshooting and maintenance are in other documents included in the document suite. Training videos, and how to use guides are on the SIEMonster website. http://www.siemonster.com

2.3 AUDIENCE

This document is intended for technical representatives of companies, SOC owners as well as security analysts and professionals. The audience of this document are expected to have a thorough level of knowledge of Security, Software and Server Architecture.

The relevant parts are included here for convenience, and may of course be subject to change. They will be updated when notification is received from the relevant owners.

2.4 SIEMONSTER AMAZON AWS BUILD OVERVIEW

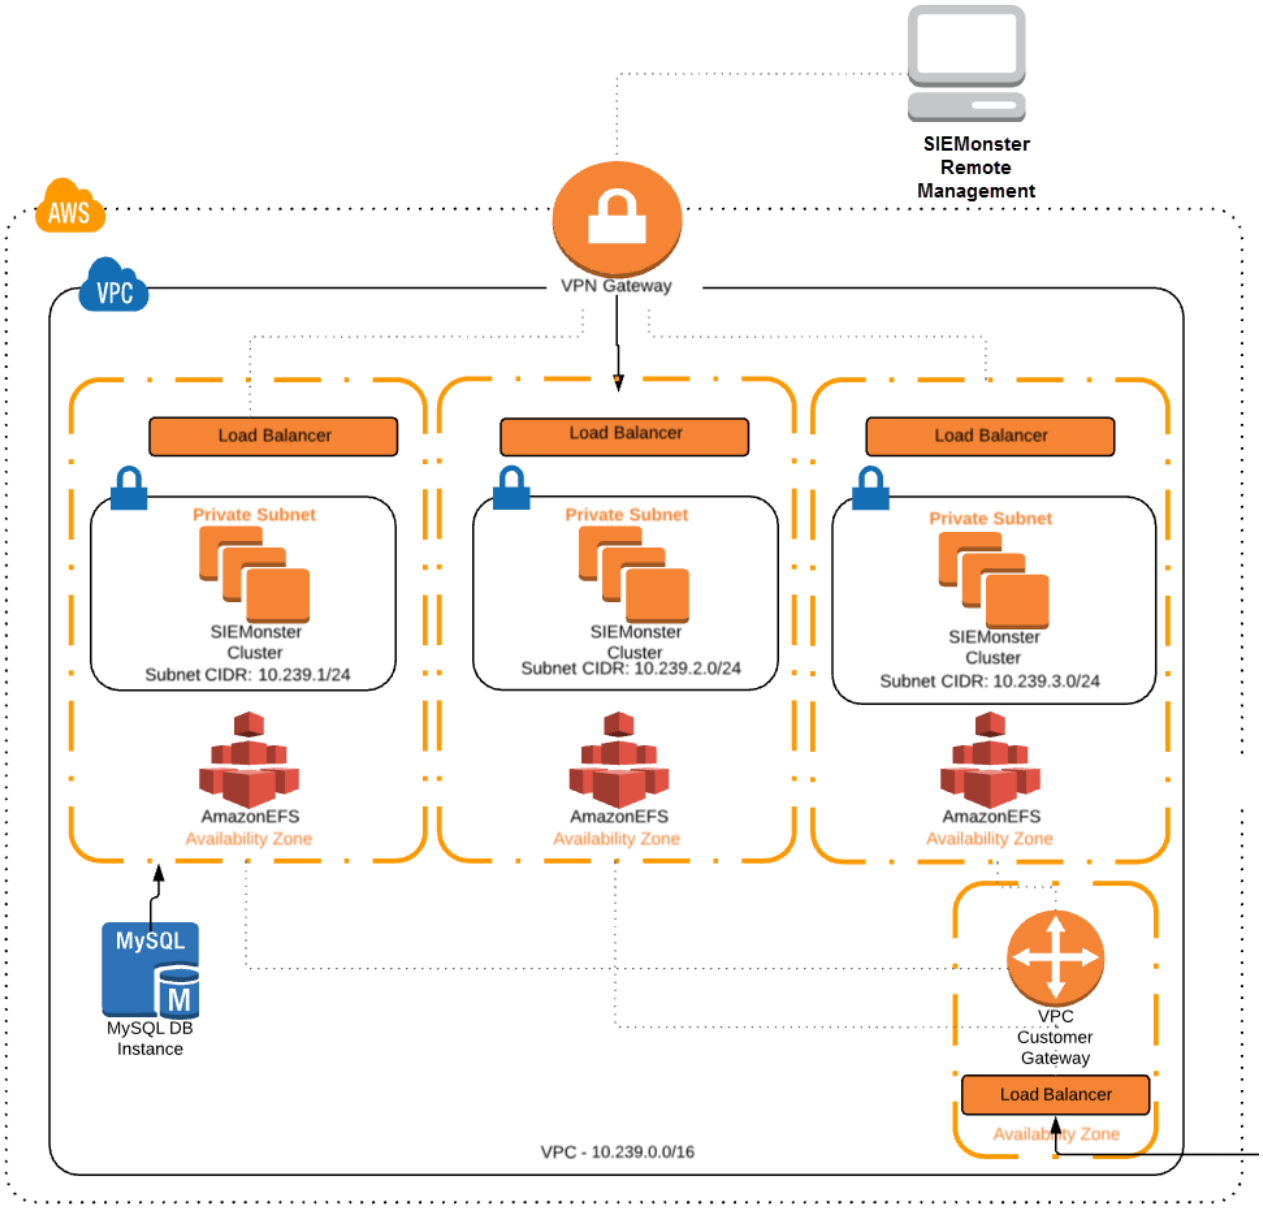

3 FUNCTIONAL OVERVIEW

3.1 HIGH LEVEL OVERVIEW

The SIEMonster Architecture is based on events being shipped, queued, and processed from endpoints within the organisation. Alerts based on rules stored and then Visualised using open source products only. Data Sources can be any device. Anything that produces an alert, a syslog or agent can be captured, correctly formatted and sent into the SIEM for analysis. This includes Network Appliances, Windows/Linux Servers, Applications, Security Appliances and SCADA. The events are filtered and stored. OSINT will alert the operator to known attack IP's and patterns from around the world, SIEMonster, 411 and OSSEC rules will identify any breaches or non-compliance and alerts.

3.2 HIGH LEVEL STACK IN AWS

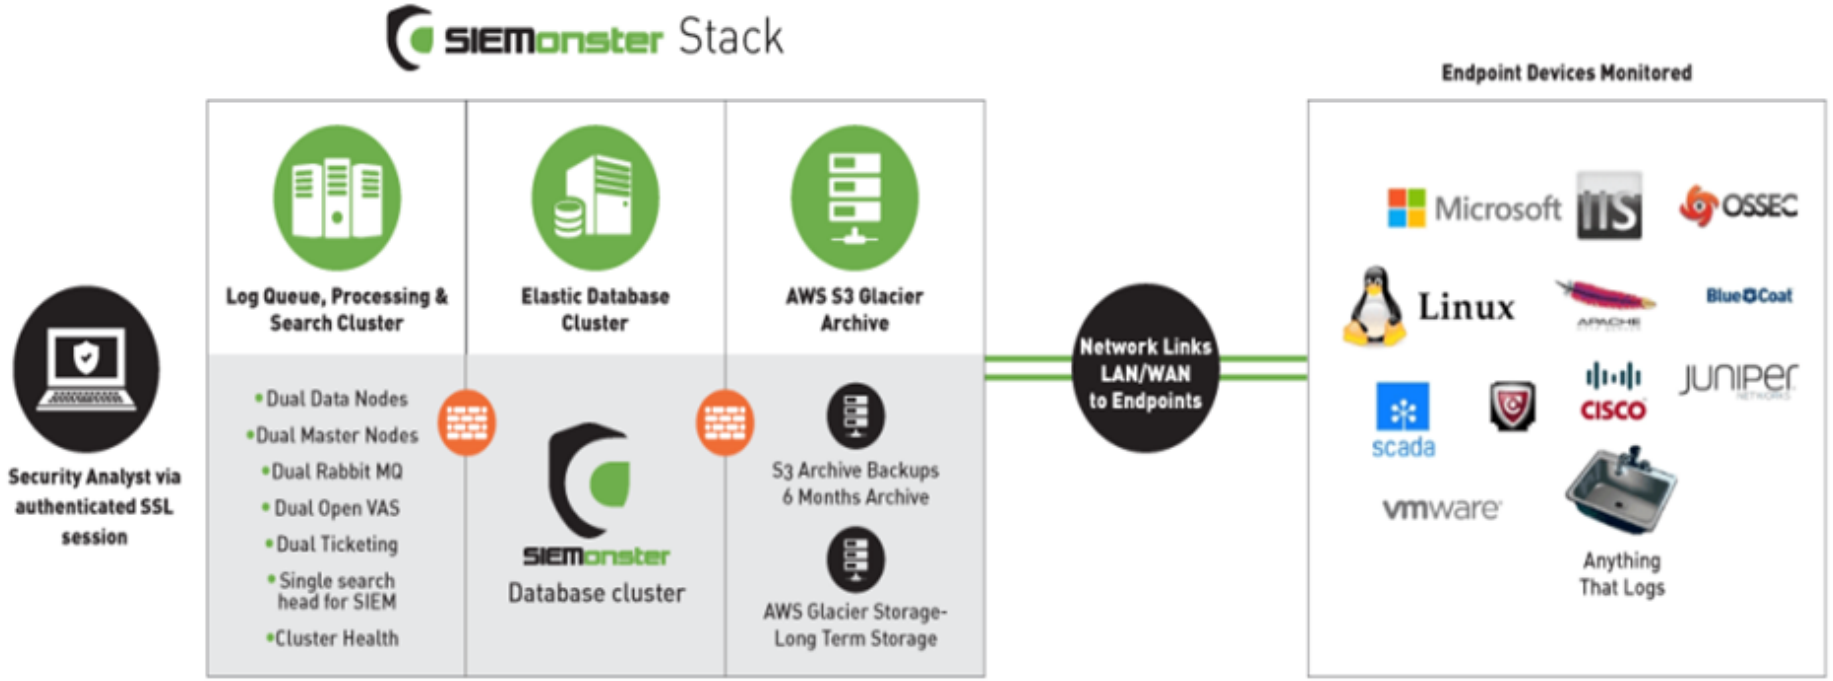

3.3 HIGH LEVEL OPEN SOURCE COMPONENTS

COLLECT	PROCESS	VISUALIZE	RISKS	TICKETING	OSINT
Collect events from your end point devices like Linux, Web Servers, Active Directory, Network Appliances	Analyse the events, process against rule sets and correlate	Show risks & alerts in the Dashboard, Web Interface, Email & or SMS the security analysts & provide reporting. User profile custom views	Vulnerability assessments against your endpoints, showing the results in the dashboard. Use OpenVAS or your commercial scanner.	The Security analyst can record incident & event for investigation & triage using included open source tool, or use your existing	Integration of Palo Alto Networks Mimemeld and Bro Intelligence Framework

3.4 SIEMONSTER HIGH LEVEL ARCHITECTURE ENDPOINT TO SIEMONSTER

.

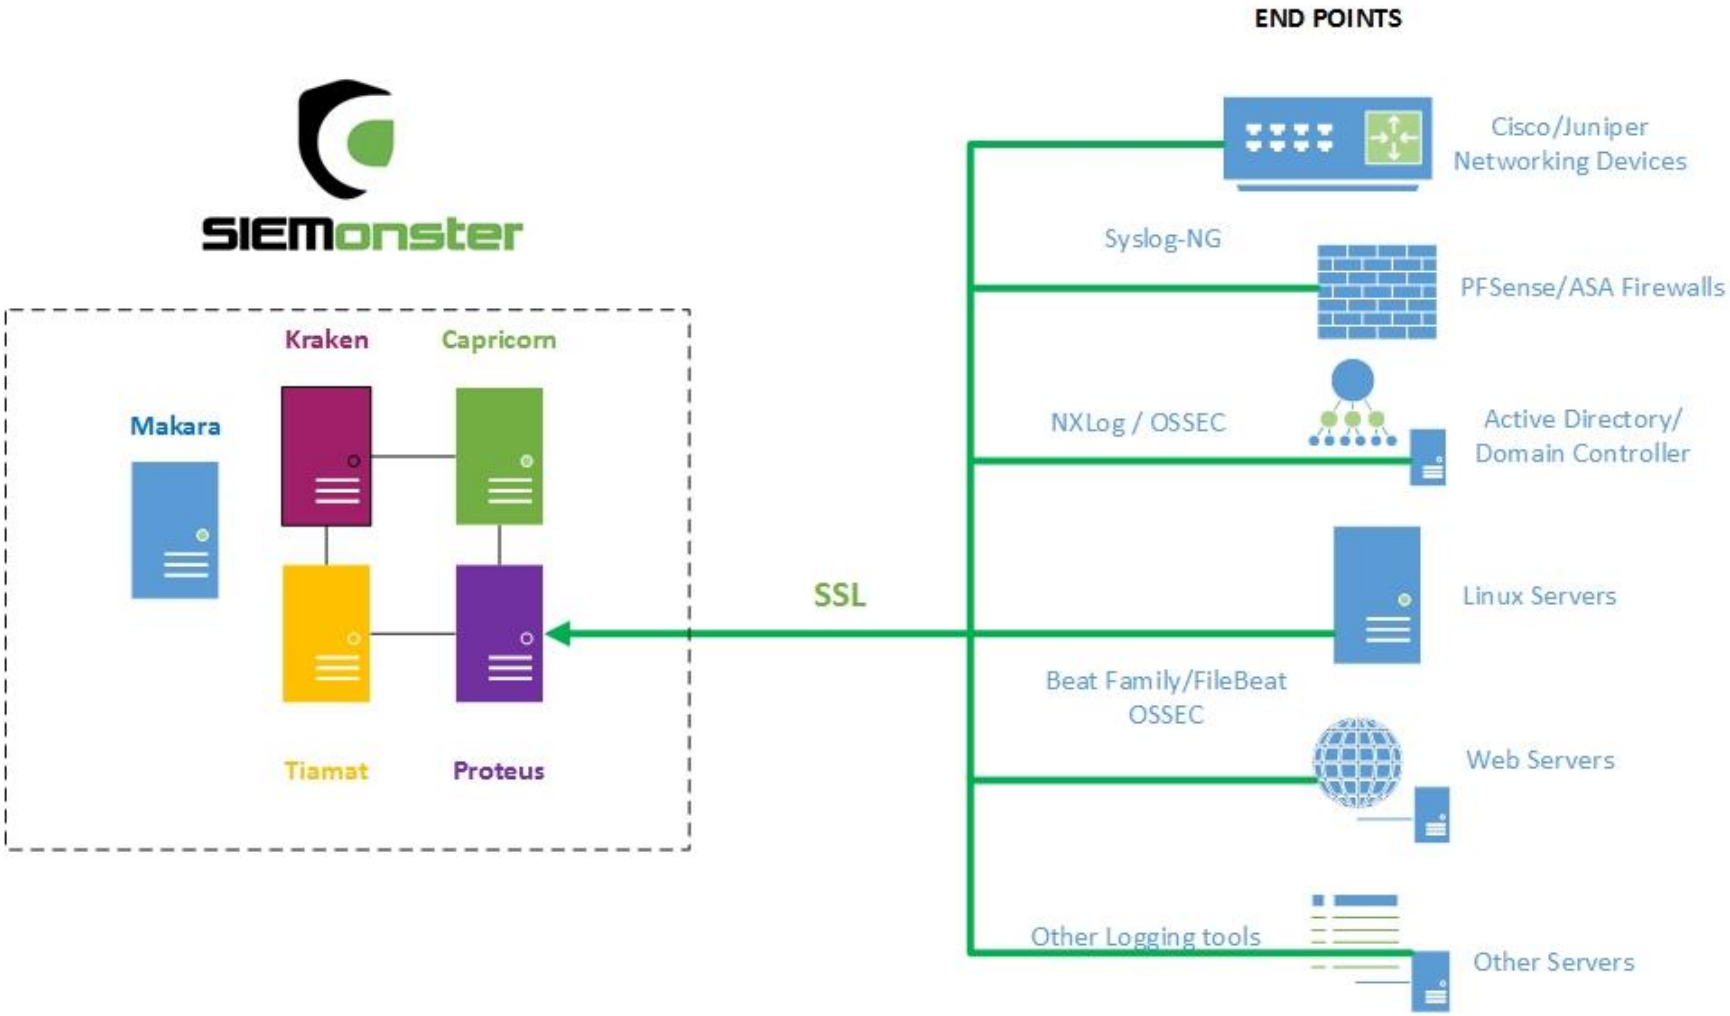

3.5 SIEMONSTER LOG FLOW

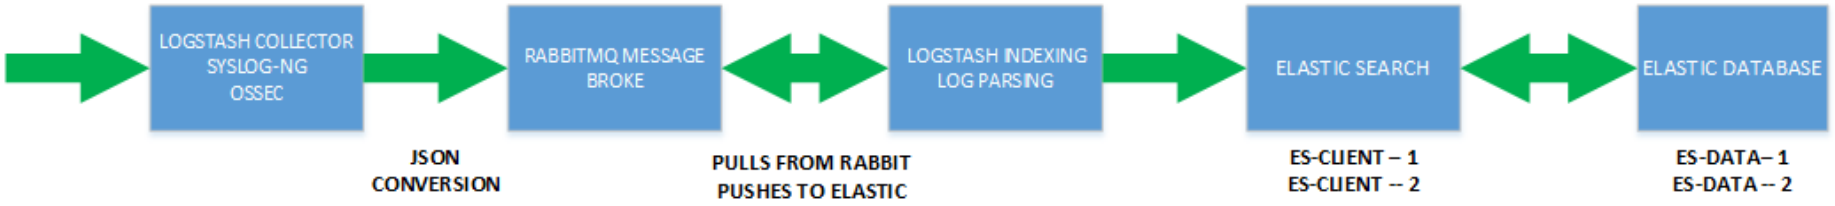

3.6 SIEMONSTER NFS MOUNTS

All of the containers configuration data is stored on shared NFS mounts across the cluster, storing the data on Makara. This is to make backups simpler going forward. Elastic Database data is still stored on Kraken and Tiamat. Within AWS these mounts are shifted to EFS.

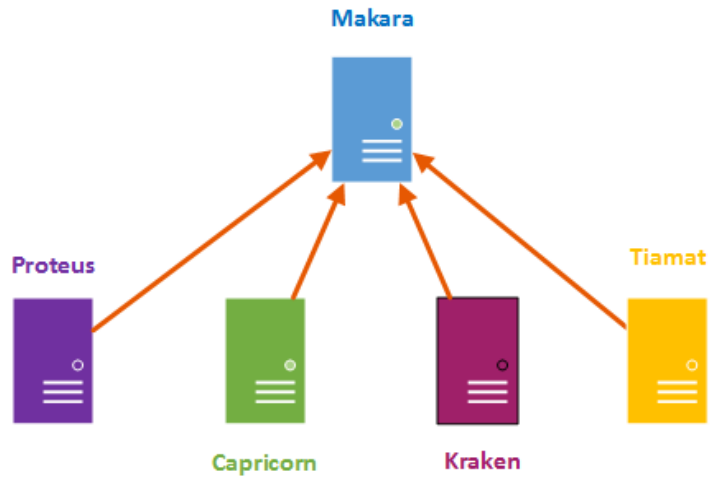

3.7 SIEMONSTER MSSP COMMERICAL EDITION OVERVIEW

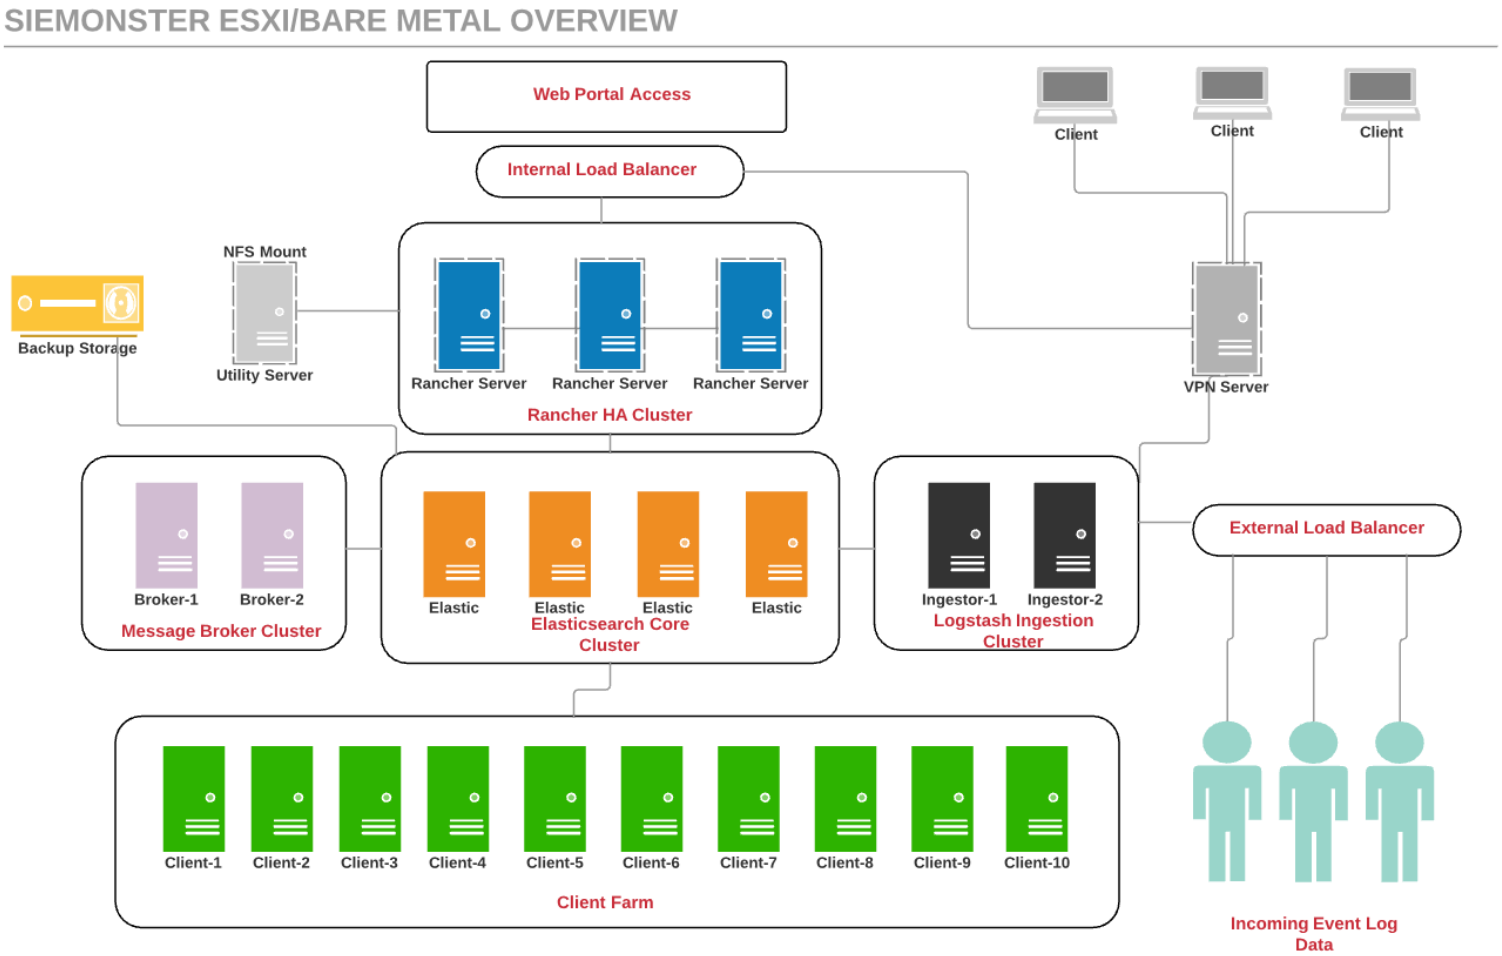

3.8 SIEMONSTER AS A SERVICE MSSP EDITION USING HYDRA EDITION

A full overview of the Log collection to visualisation process. **Client 1** represents the example client premises, the **MSSP** representing the AWS VPC for the SIEMonster MSSP. The diagram represents the entire process of log collection for client end devices all the way to storage and subsequent visualisation of data stored in the ES cluster.

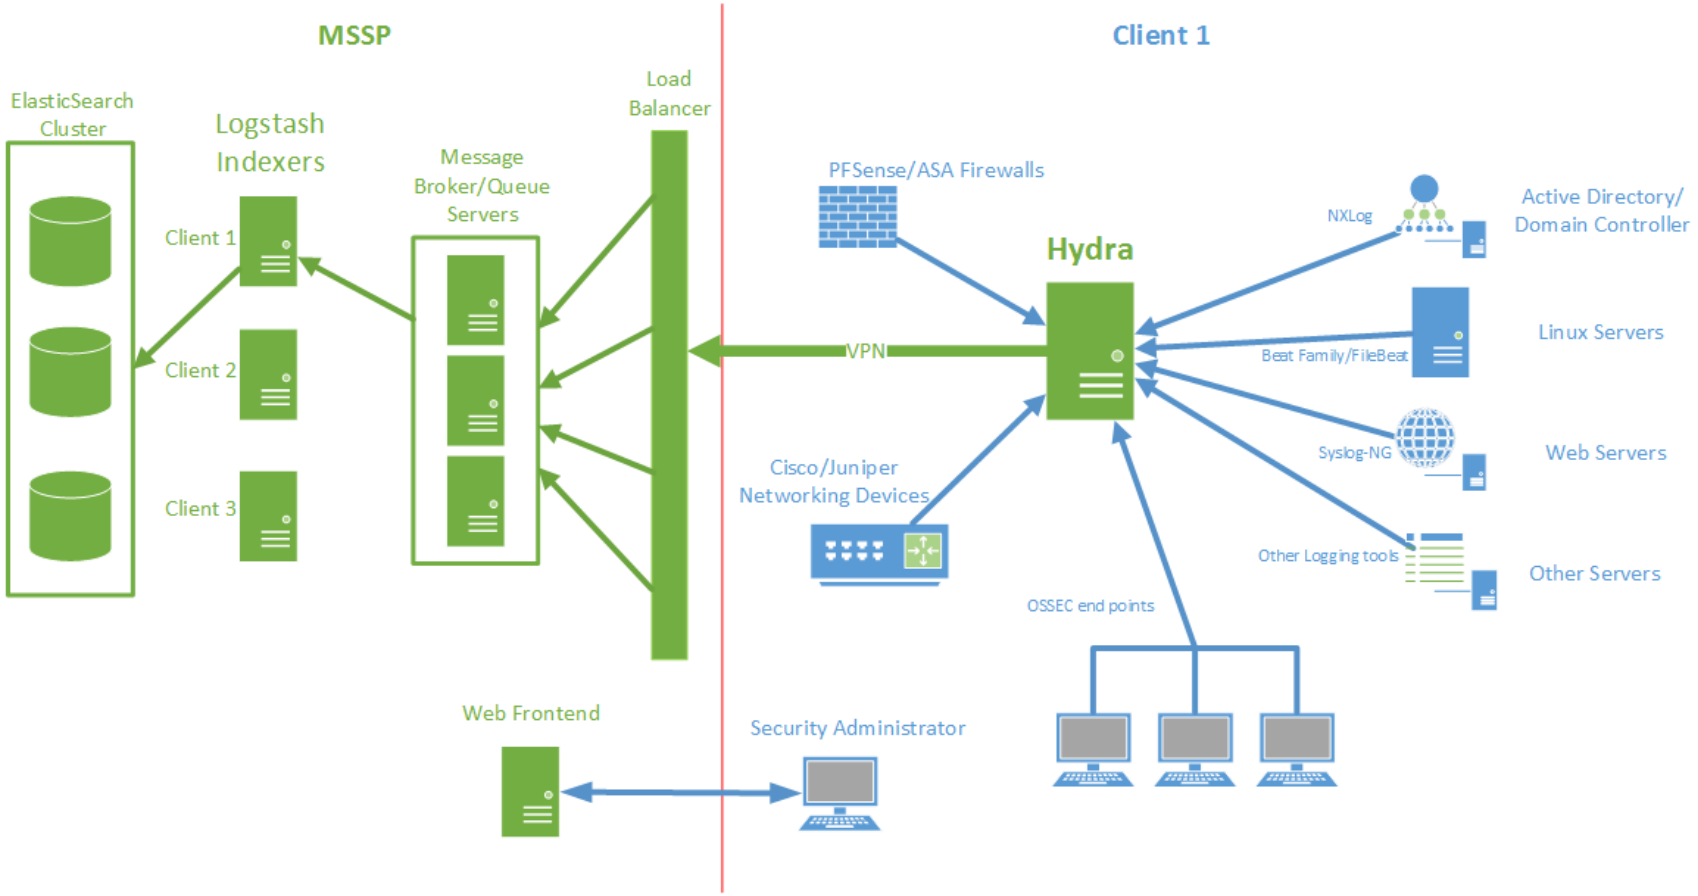

3.9 SIEMONSTER HIGH LEVEL ARCHITECTURE DOCKER

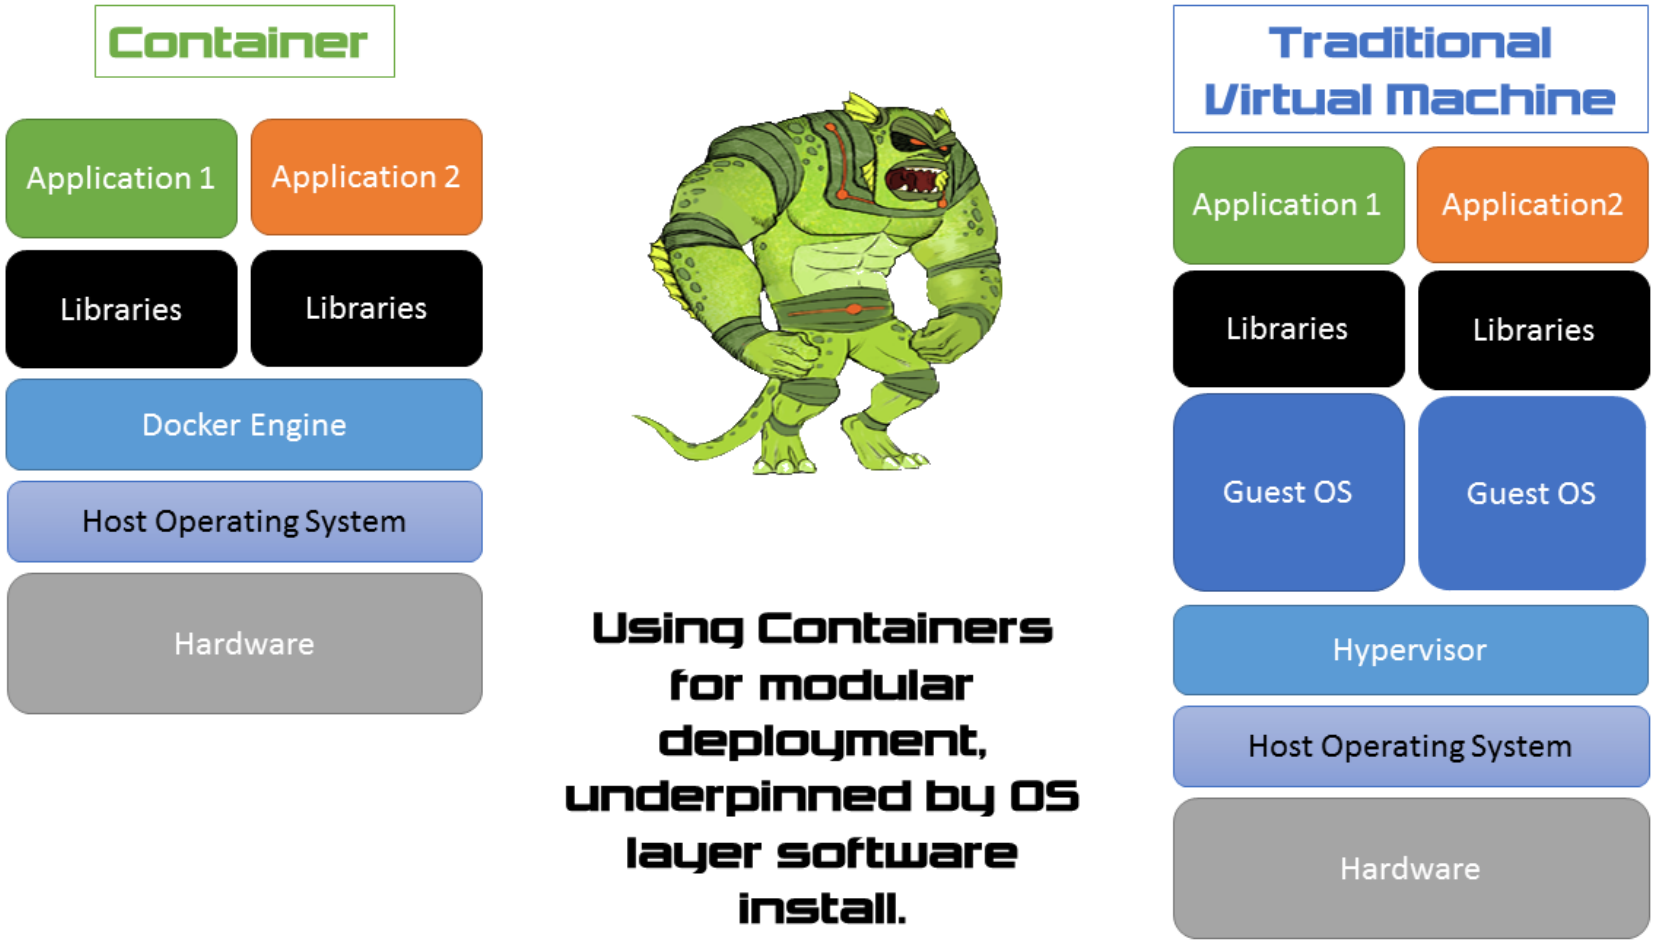

3.10 COREOS

SIEMonster has chosen CoreOS as the default operating system for the SIEMonster build. Of course, you can choose your own operating system. SIEMonster have been testing operating systems and container systems over the last 18 months to find the right OS, in terms of performance and stability. CoreOS has fit these criteria.

CoreOS is a lightweight Linux operating system designed for clustered deployments providing automation, security, and scalability for your most critical applications. CoreOS Container Linux is the leading container operating system, designed to be managed and run at massive scale, with minimal operational overhead. Applications with Container Linux run in containers, providing developer-friendly tools for deploying software. Container Linux runs on nearly any platform whether physical, virtual, or private/public cloud. This is the reason we chose this platform it was ideal for our VM/Vbox, bare metal, or multi cloud providers.

The current version in use for deployments is COREOS stable 1576.5.0 (Fix CPU disclosure of kernel memory to user process (CVE-2017-5754, Meltdown) with NFS support

3.11 CONTAINERS

Containers are key to the modern datacenter. For developers, it has never been easier to ship new application versions. Containers easily plug into your CI/CD pipeline for automated build, test and deployment with an audit trail.

Hybrid infrastructure requires containers – containers are the consistent, portable object that can be safely transported between environments. The tools to accomplish this are open source and standards driven, giving you a truly vendor-neutral solution.

The container engines Docker and are configured out of the box, ready to run your applications. Through the continuous stream of updates, Docker are automatically and continuously updated with the operating system.

3.12 SIEMONSTER HIGH LEVEL POOLING

The items marked with (Container Pool) are container services that will be deployed automatically within the cluster pending available resources. This is designed so that if any server is underloaded it can move services around to distribute the workload. Kraken and Tiamat do not allow containers to move around into the pool. Each server has a series of default installation services in Green below. These are essential for the build including NFS Mount points and Rancher Agents.The items marked with (Pinned) are services designated for fixed deployment.

3.13 FUNCTIONAL ARCHITECTURE OVERVIEW – SOFTWARE STACK

The solution had to be completely scalable, open source and completely free without exception. Using SIEMonster you can use it for free and as many nodes/clusters as you need without data restrictions.

Purpose	Description
Log Retrieval	Filebeat, Syslog, NXlog, Ossec
Process	RabbitMQ, Logstash, Kustodian
Intrusion Detection	OSSEC with Wazuh fork
Rules, Storage, Alerting	Logstash, Elasticsearch, 411, Kustodian. OSSEC Wazuh fork
Security	SearchGuard, Firewalls & Lockdowns
Audit and Discovery	Open AudIT by Opmantek
Threat Intelligence	Minemeld, open source feed aggregation
Visualize	Kibana, Slack, Dradis IDE
Backup scripts, Maintenance	Kustodian scripted AWS Glacier backup, archiving and restoring
Vulnerability Scanning	Use OpenVAS or use your own commercial scanner, Includes Dradis IDE
Ticketing	411/FIR Incident Response

3.14 RULESETS

Pre-configured alerting rule sets have been created for your use and well as integrated rules from the community are included. For example, if you want to know when a user is added to the Domain admin group, an alert will be issued or a PCIDSS environment has detected open telnet in a Card Holder Data's environment.

Sample Rule-sets to get the idea.

Rule Name	Description
Logon	A user fails to login 10 times or more within a 60-minute period
Trojan	A Trojan virus outbreak spikes at more than x events per hour
DDOS	A webserver experiences a denial of service attack when the rate of events spikes beyond a set limit
User Lockout	The rate of user lockouts spikes beyond a set limit, indicating a brute force attack
User privilege uplift	A user adds another user to a Domain Admin or administrators group.
Weird network activity	Anomalous network activity detection outside of expected patterns
Email box compromised	A user logs on to OWA from multiple IP's in different countries
Redundant Staff	A redundant staff member group for catching IP leaving your Enterprise "4 Weeks' Notice period"
Custom Alerts	Anything can be made into an alert to an email, a Dashboard alert or an SMS
Email	User using the word Confidential, Top Secret, private to a sender to an outside organisation or their own personal Hotmail/Gmail account against company policy

3.15 SIEMONSTER FUNCTIONAL TECH / METRICS OVERVIEW

5 Node - Dual Node Cluster (Makara-Proteus-Capricorn-Kraken-Tiamat)

The following metrics is the typical Enterprise size of 10,000-20,000 seats with 1000 monitored endpoints as a baseline of what you archive monitoring an array of equipment from around the organisation can. These cluster nodes are flexible, and can be scalable to a Quad node cluster at any time to suit any company size. If your organisation has WAN LINK of 256k-2MG and you are concerned about WAN overload, the event collectors can be configured to a 32-64k feed to not flood your links.

Equipment:

5 Server instances Redundant ES Cluster comprising of

- C5.4Xlarge Rancher / Orchestration / NFS Mount – Server Name **Makara**
- C5.4Xlarge AWS spec Non-Data cluster health / OSSEC/Logstash/Syslog/RabbitMQ, – Server name **Proteus**
- C5.4Xlarge AWS spec Alerting incorporating 411 & Alerta OSINT, Incident Response & Reporting – Server name **Capricorn**
- 2 X C5.4Xlarge AWS spec End DATA Elastic Cluster (redundant pair) Server name **Kraken and Tiamat**

Enterprise Sample metrics:

Event Rate: ~ 47 million events per hour scalable upwards
30-day time frame for open indices
6 months' data maintained onsite / Weekly snapshots
Data archive to AWS Glacier/S3

13 Node Quad Cluster (Makara-Proteus-Capricorn-Kraken(2)-Tiamat(2))

The following metrics is the typical Enterprise size of 10,000-20,000 seats with 3000+ monitored endpoints as a baseline.

Equipment is based on an event rate of 50 million events per hour scalable upwards with a 30-day time frame for open indices. Data storage volume based on 6 months' data maintained onsite / Weekly snapshots and data archive to AWS Glacier/S3.

Equipment:

13 Server instances Redundant ES Cluster comprising of

- 3 x Orchestration Server - AWS type M4.large, 2 core CPU. (Makara)
- 1 x ES Non-Data Master Eligible, OSSEC server - AWS type C5.2Xlarge, 8 core CPU. (Proteus)
- 1 x Alerting/Kibana/Reporting/OSINT Node AWS type C5.2Xlarge, 8 core CPU. (Capricorn)
- 4 x ES Data Nodes - AWS type C5.4Xlarge,16 core CPU, 32GB Mem. (Kraken x2, Tiamat x 2).
- 4 x Logstash/RabbitMQ/SearchHead servers - AWS type C5.2Xlarge, 8 core CPU, 8GB Mem.

Monitoring aspects Sample:

- Multiple Domain Controllers Security Event Logs
- External Websites IIS & Apache
- Exchange OWA and Message Tracking
- Multiple Cisco Devices
- IPS devices
- VPN Concentrators
- Internal Asset Vulnerability Analysis Data
- Bluecoat Proxy
- Ironport Firewalls
- McAfee ePO Orchestrator
- OSSEC HIDS Data
- Any device that's produce a log, syslog SNMP or agent installed.

Alerting – Events and Metrics Sample:

- Administrator Actions
- Logon Failures
- Anomalous Activity – Spikes/Flatlines
- Brute force attacks
- Multiple Logon Source IPs
- Email phishing and virus attacks
- Denial of Service Attacks
- Web Application Hacking Attempts
- Honeypot activity
- HIDS
- Virus Outbreaks
- Heartbeat

Visualisation: Dashboards for event visualisation Sample

- SOC Dashboard with breakdowns of relevant DC security Events
- 2007, 2010-2013 Microsoft Exchange Dashboards for Tracking Logs
- Exchange OWA Activity
- External Website Dashboards for IIS and Apache
- Cisco
- Threat Intelligence OSINT
- IPS
- Antivirus
- OSSEC HIDS
- Bluecoat Proxy
- Syslog
- Vulnerability Data
- Anything that logs, you can visualize

4 VIRTUAL HARDWARE

4.1 SIEMONSTER AWS VIRTUAL SERVERS - CORPORATE

Total Servers: 5

Server 1: Proteus (Logstash-SyslogNG-RabbitMQ-OSSEC/Wazuh)

SIEMonster Recommended Server Specifications – Front End	
Build	Cloud AWS C5.4XLarge
CPU	16 Cores
Memory	32 GB RAM
Storage	SSD
Operating System	CoreOS

Server 2: Capricorn (411-Kibana-Reporting-OSINT-Health Monitor-Event Monitor-FIR)

SIEMonster Recommended Server Specifications – Front End	
Build	Cloud AWS C5.4XLarge
CPU	16 Cores
Memory	32 GB RAM
Storage	SSD
Operating System	CoreOS

Server 3 & 4: Kraken Data Cluster Node 1 & Tiamat Data Cluster Node 2

SIEMonster Recommended Server Specifications Data Cluster 1 and 2	
Build	Cloud AWS C5.4Xlarge
CPU	16 Cores
Memory	32 GB RAM
Storage	SSD
Operating System	CoreOS

Server 5: Makara

SIEMonster Recommended Server Specifications – Orchestration / Rancher	
Build	Cloud AWS C5.4XLarge
CPU	16 Cores
Memory	32 GB RAM
Storage	SSD
Operating System	CoreOS

4.2 STORAGE RECOMMENDATIONS

Using Amazon AWS storage or local storage the costs are negligible. Data volumes of course will vary from client to client depending on how many agents your data is being transmitted but as a guide, here is a good way of managing your data requirements. Based on premise that the SIEM will store 100 GB per month and increase by 100 GB per month that's approximately 1.2 TB per year at AWS S3 storage costs of less than $500 per year. You will notice that at any stage at 6 months, 1 year, 2 years you can move your data from fast S3 storage to Amazon Glacier storage at no cost. This means you can archive your data whenever you like. You can still access the data (for a fee) but you remove the load from the Elastic Database. You might grow at 1 TB per month, even so the costs are very small for backed up storage. Of course, if your data is 10-100 times this in size, just multiply it out to give yourself a guide.

	S3 Storage	S3 Cost	Glacier Storage	Glacier Cost	**Total month**
Month 1	100	3.30	0.00	0	3.30
Month 2	200	6.60	0.00	0	6.60
Month 3	300	9.90	0.00	0	9.90
Month 4	400	13.20	0.00	0	13.20
Month 5	500	16.50	0.00	0	16.50
Month 6	600	19.80	0.00	0	19.80
Month 7	700	23.10	0.00	0	23.10
Month 8	800	26.40	0.00	0	26.40
Month 9	900	29.70	0.00	0	29.70
Month 10	1000	33.00	0.00	0	33.00
Month 11	1100	36.30	0.00	0	36.30
Month 12	1200	39.60	0.00	0	39.60

4.3 END USER AGENTS

Windows Endpoint Host: On a Microsoft host, NX-Log and OSSEC Agent is to be installed. This will collect event logs and OSSEC HIDS events and send them to Proteus encrypted via TCP/UDP for Logstash analysis.

Linux Endpoint Host: On a Linux host, Filebeat is to be installed. This will collect event logs and send them to Proteus via SSL for Logstash analysis.

Agentless Endpoint Hosts (SYSLOG): On a network appliance/printer/SCADA device that does not have an agent point the hosts SYSLOG configuration to Proteus. Proteus is running Syslog-ng and will capture all syslog's and insert them into the Logstash stream procedure. The data will then be queried for Alerting/Analysis and relevant data will be fed back into Kraken/Tiamat for long term storage.

5 FUNCTIONAL SERVERS

SIEMonster is built on 5 servers, some of the Docker containers are pinned to certain hosts and the others are free floating between the servers to distribute load. The tables in each Monster starting in 6.2 contain the pinned containers that never move from that host. Below however are the pooled containers can reside on any host and can move to free up resources. Only Makara, Capricorn and Proteus are hosts that can have containers move around on. Kraken/Tiamat containers are all pinned.

For example, the SIEMonster Website must stay on Makara for DNS resolution but containers like 411 or Grafana can float between Makara, Capricorn and Proteus.

Pooled Containers	Function
Dradis	Vulnerability Management
411 Alerts	Event Management Stream/Alerting Real Time alerts
Incidence Response FIR	Ticketing system for incidents and investigations
Prometheus + Alert Manager	Alert Management
Grafana	Alert Visualization
Reporting	Skedler Reporting Tool
RabbitMQ	Messaging queuing
MongoDB	Database for Webserver
Kibana	Visualization tool through a Browser (Dashboards)
Gmail Relay	Gmail Relay option for email alerts
MineMeld	Open Source Threat Intelligence
ElastAlert	Alerting system for integration into SMS/Email/Slack etc.
Open AudIT	Open Source Auditing tool
MYSQL	MYSQL database for 411

5.1 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
Dradis	A reporting and collaboration tool that maps and leverages all the knowledge you generate in real time, in terms of vulnerability management.
411 Alerts	411 is an open source logging and event management engine. Due to its' extensive Stream and Pipeline features and intuitive GUI, alerts can be easily configured by end users. The provided Slack & Pagerduty plugins provide notification to private channels easily viewable on a smart phone or tablet.
Incidence Response FIR	FIR is tracking system used for bug tracking, help desk ticketing, customer service, workflow processes, change management, network operations. FIR has been included into SIEMonster to record, report, and escalate Incident Responses to other security analysts
Prometheus + Alert Manager	Prometheus is an open-source system monitoring and alerting toolkit. Prometheus's provides a multi-dimensional data model with time series data identified by metric name and dashboard graphic.
Grafana	Visualization tool Grafana allows you to query, visualize, alert on and understand your metrics no matter where they are stored.
Reporting	The Skedler tool provides PDF/XLS Style reporting in the Premium Build
RabbitMQ	RabbitMQ is used as buffer funnel header that allows data flowing in from 1000's of endpoint sources quickly and orderly and holds, stores and flows into the SIEM in an orderly rate.
MongoDB	The MongoDB database provides the SIEMonster user groups
Kibana	Kibana is Elasticsearch's data visualization engine, allowing you to natively interact with all your data in Elasticsearch via custom dashboards. Kibana's dynamic dashboard panels are saveable, shareable and exportable, displaying changes to queries into Elasticsearch in real-time. You can perform data analysis in Kibana's user interface using pre-designed dashboards or update these dashboards in real-time for on-the-fly data analysis.
Gmail Relay	This option allows emails to be relayed outside of the organisation.
MineMeld	SIEMonster provides OSINT (Open-Source Intelligence) threat intelligence gathering from PA Minemeld feed aggregation and support for BRO NIDS and SNORT. OSINT data is sent to the SIEM and is used by security analysts for event context attack prediction, prevention and detective controls with real time visualization and alerting
ElastAlert	A secondary alert system, ElastAlert works by querying Elasticsearch with configured rules. It periodically queries Elasticsearch depending on the rule type, which determines when a match is found in Elasticsearch. When a match occurs, it is given to one or more alerts, which take action based on the match. This is configured by a set of rules, each of which defines a query, a rule type, and a set of alerts.
Open AudIT	An open Source tool for auditing, asset management to prioritize assets by Opmantek.
MYSQL	MySQL is provided for the Alert tool 411 configuration and alerts.

5.2 MAKARA

Makara's primary function is the Rancher Orchestration, web application front end, NFS Server and event log parsing and processing from the RabbitMQ message queue before passing on the ES Client 1 Elastic node.

5.3 SEA MONSTER NAME ORIGINS

Makara is a sea-creature in Hindu culture. It is generally depicted as half terrestrial animal in the frontal part (stag, deer, crocodile, or elephant) and half aquatic animal in the hind part (usually a fish or seal tail, though sometimes a peacock or even a floral tail is depicted.) Makara take many different forms throughout Asia. In Hindu astrology, Makara is equivalent to the sign of Capricorn, tenth of the twelve symbols of the Zodiac.

5.4 SOFTWARE OVERVIEW PINNED CONTAINERS

Makara	Function
Rancher Server	Installation and Orchestration
Rancher Agent	Rancher connection
Logstash Indexer	Log retrieval processing
Web Server Front End	SIEMonster Front End Application
Cadvisor + Nodevisor	Docker Monitor and Host Metrics
Container Volume Mounts	SIEMonster data store for all server configs

5.5 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
Rancher Server	Rancher provides Installation mechanism and Kubernetes Orchestration for unpinned container to move around servers to minimise load on certain servers.
Logstash Indexer	Logstash helps to take logs and other time-based event data from any system and stores it in a single place for additional transformation and processing. Logstash will scrub the logs and parse all data sources into an easy to read JSON format.
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
Container Volume Mounts	NFS Storage location for all SIEMonster data configuration for simple backup.
Rancher Agent	Rancher Orchestration Agent controlled by Makara.

5.6 PROTEUS SERVER

Proteus function in SIEMonster is to ingest and process incoming endpoint data and forward on to the RabbitMQ message broker and provide hosting for any of the pooled apps in Section 6.1

5.7 NAME ORIGINS

In Greek mythology, Proteus is an early sea-god or god of rivers and oceanic bodies of water. Some describe him a specific domain call him the god of "elusive sea change", which suggests the constantly changing nature of the sea or the liquid quality of water in general. He can foretell the future, but, in a mytheme familiar to several cultures, will change his shape to avoid having to; he will answer only to someone who can capture the beast. From this feature of Proteus comes the adjective protean, with the general meaning of "versatile", "mutable", "capable of assuming many forms". "Protean" has positive connotations of flexibility, versatility and adaptability.

5.8 SOFTWARE OVERVIEW PINNED CONTAINERS

Proteus	Function
ES Client Node 1	Open source, distributed, real-time search and analytics engine
Logstash Collector	Log retrieval, processing
SYSLOG - NG	SYSLOG engine for incoming logs
OSSEC - WAZUH	HIDS, Rulesets, PCIDSS, CIS benchmarks, Forensic analysis
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount to store container data remotely
Rancher Agent	Rancher connection

5.9 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
ES- Client 1	Elasticsearch is running on Proteus Client Node 1
Logstash Collector	Logstash helps to take logs and other time-based event data from any system and stores it in a single place for additional transformation and processing. Logstash will scrub the logs and parse all data sources into an easy to read JSON format.
SYSLOG - NG	SYSLOG-NG provides data ingestion from Syslog devices such as routers and firewalls.
OSSEC - WAZUH	Host Intrusion Detection Server Ingestion from clients for SIEM correlation and alerting.
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount storing configuration and server data on Makara
Rancher Agent	Rancher Orchestration Agent controlled by Makara.

5.10 PROTEUS AND END USER AGENTS

Proteus receives logs from all Windows, Linux, Application and hardware providing syslog's. Agents provide TLS/SSL encryption using purchased certificates or in-house self-signed or propriety certs included in the OVA image. By using this encryption and other methods there is no need for the Elastic Shield product which means support is free.

Preconfigured Nxlog agents with SSL certificates is used for Windows hosts is used for log collecting and sending to Proteus TCP Ports 3520-3529

Preconfigured Filebeat agents with SSL certificates is used on Linux hosts for log collecting and sending to Proteus TCP Ports 3520-3529

Hosts that don't support an agent such as Network appliances can be configured to send all alerts SYSLOGS (0,1,2,3,4+) UDP Ports 514/1516 and TCP 1516,55000 to Proteus which has Syslog-ng installed to collect these logs.

5.11 CAPRICORN SERVER

Capricorns function in SIEMonster is an Elastic Client Node 2 and provide hosting for any of the pooled apps in Section 6.1
.

5.12 NAME ORIGINS

The Babylonians connected the Zodiac sign and the constellation with the mythological animal, sort of a mermaid goat. They called it the Goat-Fish. Here is a feature of the goat that deserves special attention in our understanding of this sign. The name 'Capricorn' draws our attention to it and comes from the Latin caper ('goat') and cornu ('horn') - literally 'the Goat's horn'. In the ancient world horns were symbols of royalty, strength and power, as well as fertility and abundance. Cornucopia in mythology was the goat Amalthea who nourished the infant Jupiter with her milk, though the term remains in use today as the 'copious horn' or 'horn of plenty' which symbolises prosperity and growth. The goat is one of the three horned creatures in the zodiac; these were also the creatures celebrated in ancient religious festivals and used in sacrifice to draw power from the gods. The use of the goat as a 'scapegoat' in the biblical ritual of Atonement (seethe star lore of Capricorn) has led to goat deities accumulating a reputation as icons of evil occult powers rather than the neutral symbols of earthly fertility and focused power that was implicit in the older customs.

5.13 SOFTWARE OVERVIEW PINNED CONTAINERS

Capricorn	Function
ES Client 2	Open source, distributed, real-time search and analytics engine
ES Master	Open source, distributed, real-time search and analytics engine
NGINX	NGINX provides HTTPS services as well as reverse Proxy
SearchGuard	Elastic Security
SearchGuard Admin	Admin for Searchguard
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount to store container data remotely
Rancher Agent	Rancher connection

5.14 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
ES- Client 2	Elasticsearch is running on Capricorn Client Node 2
ES- Master	Elasticsearch Master node is running on Capricorn
SearchGuard	SearchGuard is an Elasticsearch plugin that offers encryption, authentication and authorisation. It builds on Search Guard SSL and provides pluggable auth/auth modules in addition. Search Guard is an alternative to ES Shield, and offers all basic security features for free. If you need enterprise features, we offer a very flexible licensing model and support. Tailored to your needs if none of our packages fit.
SearchGuard Admin	Docker container called ES_master which provides Searchguard Admin
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount storing configuration and server data on Makara
Rancher Agent	Rancher Orchestration Agent controlled by Makara.

5.15 KRAKEN

Kraken's primary function is Cluster Node 1 Elastic storing all your long term SIEM data in the database. When a user performs a Kibana search on. "All users who used the word confidential in an email sending to an external email domain" Elasticsearch database will locate the entries and present the lookup to the user in Kibana. Cluster Node 2 called Tiamat is identical and provides redundancy for Kraken. The health and controlling of the cluster is done by Proteus. In the event of hardware failure, a cluster node can be bought offline and another replaced.

5.16 SEA MONSTER NAME ORIGINS

The Kraken is a legendary sea monster of giant size that is said to dwell off the coasts of Norway and Greenland. Several authors over the years have postulated that the legend originated from sightings of giant squids that may grow to 12–15 meters (40–50 feet) in length, even though the creature in the original tales was not described as having tentacles and more closely resembled a whale or crab. The sheer size and fearsome appearance attributed to the kraken have made it a common ocean-dwelling monster in various fictional works.

5.17 SOFTWARE OVERVIEW PINNED CONTAINERS

Kraken	Function
Elastic Search – ES Data Node 1	Open source, distributed, real-time search and analytics engine Version 5.5.2
SearchGuard	Elastic Security
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount to store container data remotely
Rancher Agent	Rancher connection

5.18 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
Elastic Search	Elasticsearch is a flexible and powerful open source, distributed, real-time search and analytics engine. Architected from the ground up for use in distributed environments where reliability and scalability are must haves, Elasticsearch gives you the ability to move easily beyond simple full-text search.
SearchGuard	SearchGuard is an Elasticsearch plugin that offers encryption, authentication and authorisation. It builds on Search Guard SSL and provides pluggable auth/auth modules in addition. Search Guard is an alternative to ES Shield, and offers all basic security features for free. If you need enterprise features, we offer a very flexible licensing model and support. Tailored to your needs if none of our packages fit.
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount storing configuration and server data on Makara
Rancher Agent	Rancher Orchestration Agent controlled by Makara.

5.19 TIAMAT

Tiamat's primary function is Cluster Node 2 Elastic storing all your long term SIEM data in the database. When a user performs a Kibana search on. "All users who used the word confidential in an email sending to an external email domain" Elasticsearch database will locate the entries and present the lookup to the user in Kibana. Cluster Node 1 called Kraken is identical and provides redundancy for Tiamat. The health and controlling of the cluster is done by Makara. In the event of hardware failure, a cluster node can be bought offline and another replaced.

5.20 SEA MONSTER NAME ORIGINS

Tiamat is a primordial goddess of the ocean, mating with Abzû (the god of fresh water) to produce younger gods. She is the symbol of the chaos of primordial creation, depicted as a woman, she represents the beauty of the feminine, depicted as the glistening one It is suggested that there are two parts to the Tiamat mythos, the first in which Tiamat is a creator goddess, through a "Sacred marriage" between salt and fresh water, peacefully creating the cosmos through successive generations. In the second "Chaoskampf" Tiamat is considered the monstrous embodiment of primordial chaos. Some sources identify her with images of a sea serpent or dragon.

5.21 SOFTWARE OVERVIEW PINNED CONTAINERS

TIAMAT	Function
Elastic Search – ES Data Node 2	Open source, distributed, real-time search and analytics engine. Version 5.5.2
SearchGuard	Elastic Security
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount to store container data remotely
Rancher Agent	Rancher connection

5.22 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
Elastic Search	Elasticsearch is a flexible and powerful open source, distributed, real-time search and analytics engine. Architected from the ground up for use in distributed environments where reliability and scalability are must haves, Elasticsearch gives you the ability to move easily beyond simple full-text search.
SearchGuard	SearchGuard is an Elasticsearch plugin that offers encryption, authentication and authorisation. It builds on Search Guard SSL and provides pluggable auth/auth modules in addition. Search Guard is an alternative to ES Shield, and offers all basic security features for free. If you need enterprise features, we offer a very flexible licensing model and support. Tailored to your needs if none of our packages fit.
Cadvisor + Nodevisor	Cadvisor monitors Docker container usage and performance. Nodevisor monitors host metrics eg Disk Space, CPU.
NFS Mount	NFS Mount storing configuration and server data on Makara
Rancher Agent	Rancher Orchestration Agent controlled by Makara.

5.23 IKUTURSO

Ikuturso role is a network sensor placed away from SIEM sitting in a DMZ or network edge, running BRO, LOGSTASH, RabbitMQ, Suricata with the ability to block known traffic using Threat Intelligence. It also provides Forensic capabilities of known attacks with deep application and network packet inspection.

5.24 SEA MONSTER NAME ORIGINS

Iku Turso is a famous monster in Finnish mythology. Iku turso was described as an evil sea-monster and is dated back to the 16th century. Iku Turso was described as many different things but mainly symbolizes death and evil. In Finnish, the word for octopus, (merituras) is named after Iku Turso and in WW2 the Finns named one of their submarines Iku Turso.

Iku Turso's appearance is described in many ways. Some say he is a thousand headed, or a thousand horned, or the one that lives on the edge. In Finnish mythology, he is known as the ox of death, the god of war, and the demon of diseases. So, he is respected in some ways for being a warrior but overall he symbolizes evil. He is said to be from the far north land of Pohjola, which (per Finnish mythology) is forever cold and the heart of all evil.

5.25 SOFTWARE OVERVIEW FUNCTION TABLE

Ikuturso	Function
SURICATA	Network Packet Analysis
BRO IDS	Packet inspection Network Tool
RabbitMQ	Messaging Queueing Service
Logstash Collector	Log retrieval, processing

5.26 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
SURICATA	Network Packet Analysis
BRO	Bro is a powerful network analysis framework that provides site-specific monitoring policies that does not rely on specific network signatures. It provides Forensics capabilities with full network logging. .
RabbitMQ	Messaging Queueing Service
Logstash Collector	Log retrieval, processing

5.27 HYDRA (AWS COMMERCIAL OPTION)

Hydra is used by SIEMonster as a server that collectors logs at a customer's site who requires SIEM as a Service. Instead of all of the customer's endpoints sending logs directly into the AMAZON VPC tunnel, Hydra collects all the logs ensures correct queuing and in the event of a Cloud outage stores the SIEM logs until it comes back online. Hydra then passes the SIEM event into Amazon AWS to Proteus/Capricorn and Kraken/Tiamat.

5.28 SEA MONSTER NAME ORIGINS

Hydra or Hydra of Lerna more often known simply as the Hydra, was a serpentine water monster in Greek and Roman mythology. Its lair was the lake of Lerna in the Argolid, which was also the site of the myth of the Danaids. Lerna was reputed to be an entrance to the Underworld and archaeology has established it as a sacred site older than Mycenaean Argos. In the canonical Hydra myth, the monster is killed by Heracles, using sword and fire, as the second of his Twelve Labors. According to Hesiod, the Hydra was the offspring of Typhon and Echidna. It possessed many heads, the exact number of which varies per the source. Later versions of the Hydra story add a regeneration feature to the monster: for every head chopped off, the Hydra would regrow one or multiple heads. The Hydra had poisonous breath and blood so virulent that even its scent was deadly.

5.29 SOFTWARE OVERVIEW FUNCTION TABLE

Hydra	Function
Logstash	Log retrieval, processing
RabbitMQ	Messaging queuing
OpenVPN	OpenVPN client software
OSSEC WAZUH	OSSEC WAZUH HIDS Software
SYSLOG-NG	Collecting syslog's from proprietary devices.

5.30 SOFTWARE DETAIL FUNCTION TABLE

Software	Function
Logstash	Logstash helps to take logs and other time based event data from any system and stores it in a single place for additional transformation and processing. Logstash will scrub the logs and parse all data sources into an easy to read JSON format.
RabbitMQ	RabbitMQ is used as buffer funnel header that allows data flowing in from 1000's of endpoint sources quickly and orderly and holds, stores and flows into the SIEM in an orderly rate.
OpenVPN	OpenVPN client software to allow encrypted traffic between datacentres or AWS/AZURE infrastructure Direct Connect is not in place.
OSSEC WAZUH	OSSEC WAZUH HIDS Software
SYSLOG-NG	Collecting syslog's from proprietary devices.

6 SECURITY

The SIEM servers must reside in a secure subnet protected by a Firewall blocking all but the required ports. The SIEM contains the most sensitive logs and must be protected from the other network equipment. Administration access should be restricted via ACLS and SIEMonster is configured for user/password access. Some of the technologies are listed below.

NOTE: In other words, don't stick it on the Internet with SSH facing outwards for anyone to access. Use the firewall and ACLS settings on the net page as a guide.

- Flexible REST layer access control (User/Role based; on aliases, indices and types)
- Flexible transport layer access control (User/Role based; on aliases, indices and types)
- Document level security (DLS): Retrieve only documents matching criteria (Premium Customers)
- Field level security (FLS): Filter out fields/source parts from a search response (Premium Customers)
- HTTP authentication (SPNEGO/Kerberos, Mutual SSL/CLIENT-CERT)
- HTTP session support through cookies
- Authentication backends (LDAP(s)/Active Directory)
- Node-to-node encryption through SSL/TLS (Transport layer)
- Secure REST layer through HTTPS (SSL/TLS)
- X-Forwarded-For (XFF) support
- Audit logging (Optional Extra)
- LUKS Disk encryption at rest (bare metal installs only recommended) (Premium Customers

6.1 FIREWALL SETTINGS FOR ISOLATED SIEM MULTI NODE CLUSTER

It is recommended that SIEMonster sits behind an internal firewall to remove unnecessary surface areas. The following tables will help setup these firewall rulesets.

6.2 UNTRUSTED TO TRUSTED (LAN TO SIEMONSTER)

Description: Incoming Ports required to be open from untrusted network (LAN) to SIEMonster environment

Source LAN	Destination	Port & Transport
NxLog Agent / OSSEC / Filebeat	Proteus	3520-3529 TCP
Syslog (Network Appliances)	Proteus	514,1514 UDP 1516, 55000 TCP
Security Professional PC for SSH	SIEMonster Servers	22 TCP
Security Professional PC SIEMonster Interface	Capricorn	443 TCP
Security Professional PC for Rancher	Makara	8080 TCP

6.3 TRUSTED TO TRUSTED (SIEMONSTER TO SIEMONSTER TRAFFIC)

Description: Ports required between all the SIEMonster servers

Source SIEMonster	Destination	Port & Transport
Internal Traffic UDP	SIEMonster	500 UDP 4500 UDP
Internet Traffic TCP	SIEMonster	2049 TCP 2379 TCP 2380 TCP 4001 TCP

6.4 AGENT PORT QUICK PORT CHECK

Source (End Points)	Destination	Port & Transport
NxLog Agent / OSSEC / Filebeat	Proteus	3520-3529 TCP
Syslog (Network Appliances)	Proteus	514,1514 UDP 1516, 55000 TCP

6.5 SECURITY ANALYST QUICK CHECK

Source (Admin Desktop)	Destination	Port
Security Analyst	All Servers	22 TCP (SSH)
Security Analyst	Capricorn	443 TCP
Security Analyst	Makara	8080 TCP

Note: Internet access for SIEMonster for daily rule updates via secure Internet gateway. These are just recommendations to protect your SIEM