limacharlie.io

Core Concepts

This document outlines Core concepts for those getting started with LimaCharlie: Security Infrastructure as a Service.

Any questions related to this document, or LimaCharlie in general, can be directed to: answers@limacharlie.io

Security Infrastructure as a Service

LimaCharlie is a cross-platform endpoint detection and response solution, with supporting infrastructure, delivered as a serverless utility. LimaCharlie provides all of the components required for modern information security operations, either through a web application or via the API/SDK/CLI. LimaCharlie provides a real-time connection to endpoints where activities are monitored against a set of rules to determine if an automated action or an investigation is warranted.

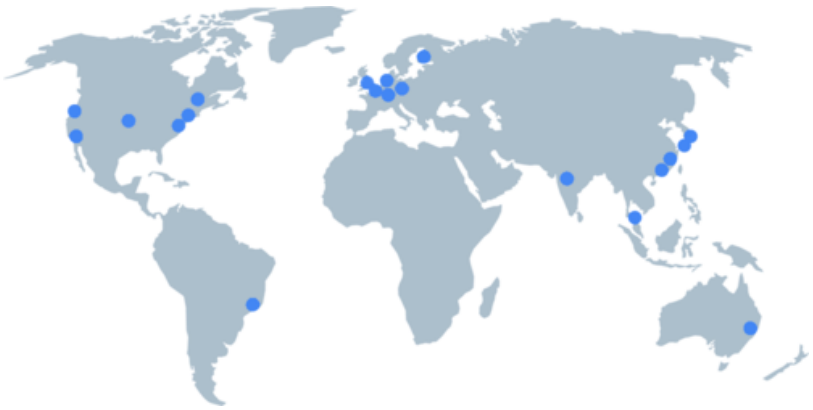

The LimaCharlie global infrastructure itself is built on the Google Cloud Platform (GCP) and currently has computing resources available in the USA, Canada, Europe, India and the United Kingdom. Choosing a geographical location ensures data will always be processed in this location and never moved outside. New data centres can be spun up anywhere GCP is available upon customer request.

Architecture

LimaCharlie employs a multi-tenant architecture. Every cluster is built across multiple availability zones to maintain high-availability. Clusters are available in multiple regions to deal with any legal requirements and all access to these clusters is transparent, allowing MSSPs to provision new customers wherever they see fit.

Connectivity between agents and the clusters is done point-to-point and is secured using pinned SSL certificates.

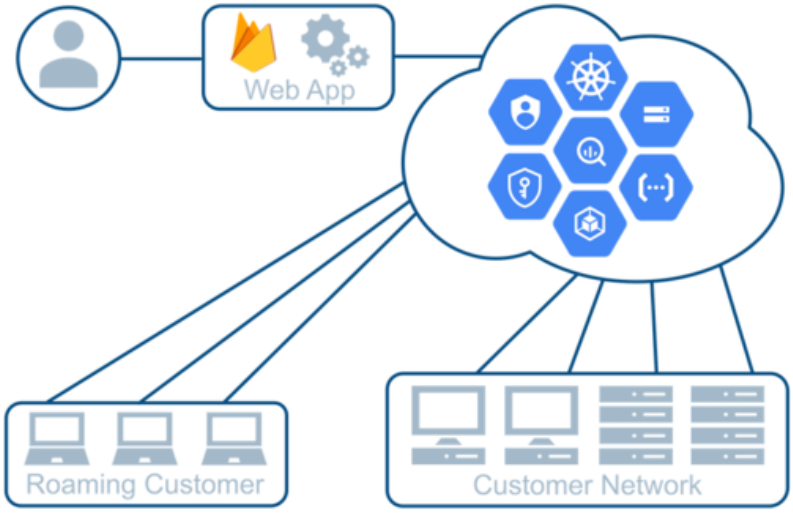

This enables extremely easy deployments and seamless transition of assets between offices and networks.

Each cluster has its own transport keys and each organization has its own command keys. All keys are encrypted in storage and runtime access to the envelope keys is limited to server containers. All sensitive output configurations are encrypted per organization.

Several components work together to provide custom solutions based on the needs of individual users. LimaCharlie is modelled after other cloud service providers, such as Amazon Web Services (AWS) or the Google Cloud Platform (GCP), but designed for the needs of information security operations.

Components (view from 30,000ft)

The components of the LimaCharlie stack, available in the cloud, are outlined at a high-level as follows.

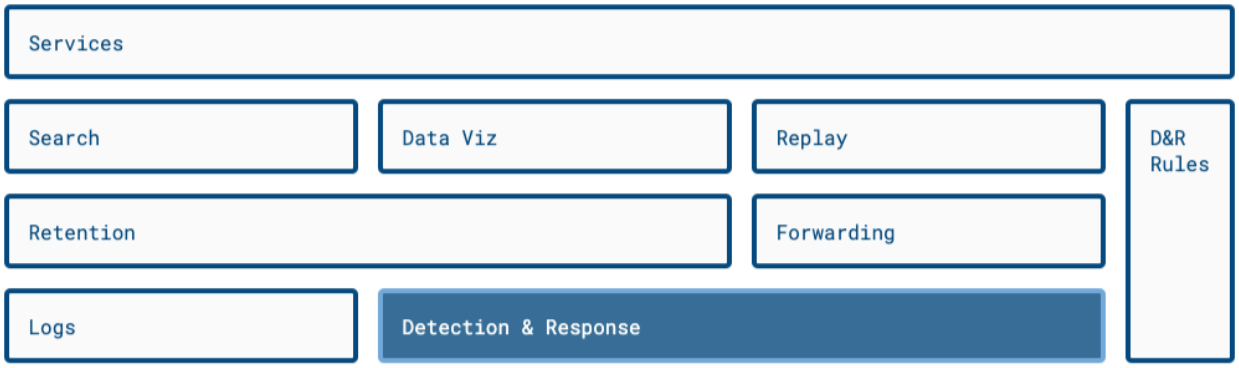

Endpoint Detection & Response

At its core LimaCharlie is an Endpoint Detection & Response engine that operates at wire speed.

The well documented Detection and Response engine consumes instructions using the YAML format. Examples of how this engine works along with documentation can be found here.

LimaCharlie also provides a mechanism for running payloads through the agent. Payloads can be any executable and this feature is only available with specific permissions enabled. Details on this feature can be found here.

Detect & Response Rules (D&R)

With LimaCharlie you can create custom Detection & Response rules. The Detection component is a scriptable rule system that can be chained with multiple conditions that will match certain events. When the Detection component matches the Response component is actioned. Responses are actioned in real-time by the sensor via a secure semi-persistent TLS connection. The Detection & Response engine automates your ability to investigate, mitigate or apply tags and more using

serverless functions. Roundtrip time from detection until the response action is typically under 100ms.

Documentation on custom D&R rules can be found here.

Data Forwarding (outputs)

The LimaCharlie endpoint agent produces telemetry in well documented JSON that can be output wherever you want. LimaCharlie has modules supporting Slack, Google Cloud Storage, S3, SFTP, Syslog, SMTP and SCP. Out of the box we provide you with a quickstart script to get a Splunk instance setup to receive data in minutes.

You can have as many Output modules active as you want and adjust the verbosity of the telemetry (event, detect and audit levels). This means you can stream to two different syslog destinations using the Syslog Output module and then send the same data to cold storage over an Scp Output module.

Documentation on setting up output connections can be found here.

Data Retention (telemetry storage)

The LimaCharlie agent produces a lot of telemetry (currently 54 unique even types). Long-term telemetry storage with search capability is available as an optional add-on. LimaCharlie can still operates as middleware but with this feature enabled we offer one year of **full** telemetry storage and search capability. Telemetry storage allows MSSPs and SOCs who do not already have their own EDR infrastructure to gain a completely functional information security centre upon signing up.

Once enabled, LimaCharlie Insight will automatically send all telemetry data to secure storage on the Google Cloud Platform.

A full list of events can be found here.

External Logs

LimaCharlie can automatically collect and store logs with no additional configuration or without installing another agent. Users

select how long the logs are to be stored and can even send logs to LimaCharlie through the API or CLI manually.
LimaCharlie can consume logs from any OS. Logs can be unstructured (no parsers necessary) and we even support pcap and Windows logs. Logs are indexed across common indicators of compromise and can be searched. D&R, YARA or Sigma rules may also be run against log files at ingestion or historically.

Advanced Searching

Telemetry and logs stored with LimaCharlie are indexed across common indicators of compromise (IOCs) and can be searched in a number of ways. Using the web application or API users can search an organization for an IP, file path, hash or user name which will bring back stats around the prevalence of the given datapoint: where when and how often it has been observed. Users are also able to drill down to the raw telemetry for events of interest.

Using the command line interface (CLI) for LimaCharlie users can search for IOCs across multiple organizations. The new CLI command supports multiple arguments and the output is written human-readable to stdout or to a file as YAML. The following man page outlines all available options and provides an example.

Replay (Historical Search)

LimaCharlie is able to perform retroactive hunting on up to a year's worth of log and telemetry data. This ability to retroactively apply Detection & Response rules to historical data is extremely powerful and through the use of the API can be used to build Continuous Delivery (CD) / Continuous Integration (CI) into detection systems.

This ability enables you to look for specific indicators of compromise and run complete D&R rules, including threat feeds, APIs or operators against historical telemetry.

Data Visualization

Through the web application, LimaCharlie provides an interface where users can explore prevalence and timing for domains, IP addresses, files, file paths, hashes and user names across their entire fleet for up to a year's worth of data. This console also allows users to search

for a given sensor ID and deep dive all of the aforementioned data points for a single endpoint over the given time period.

When searched, an indicator of compromise will show up in the graph along with information relating to when it was observed, how often and across how many endpoints. Users can then click on any given graph node to get more information and start to drill down.

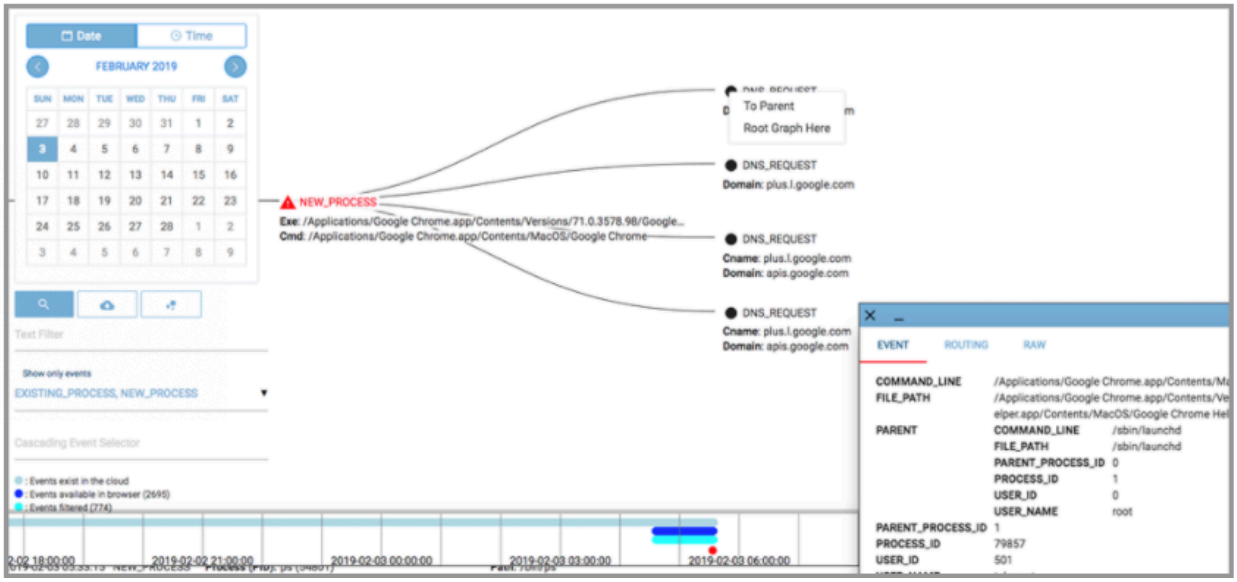

Services

The LimaCharlie platform is designed to leverage external resources and provides a framework for easily building your own which you can monetize or keep private.

Out-of-the-box LimaCharlie provides integrations with the following (and the list is always growing).

- Backstory
- YARA
- SIGMA
- VirusTotal
- Humio
- Zeek
- SocSoter
- AlienVault OTX
- MISP

The Agent

A software sensor - or agent - installed on the endpoint maintains a persistent TLS connection to the cloud where events are monitored and evaluated. Round-trip times for a detection and response to take place are generally under 100ms.

The total footprint of the agent on disk combined with what is in memory is approximately 2MB. The agent typically runs under 1% CPU with small spikes around certain events and usually uses less than 50MB of RAM.

The amount of data that is produced by the agent is dependent on the activity that is taking place on the endpoint. That being said, the average data produced per endpoint across thousands of deployments is approximately 1MB per day.

The LimaCharlie agent has zero impact on endpoint performance, even when analyzing, searching and investigating incidents. The agent deploys in seconds with no reboot required and will not be picked up by antivirus software (it does not hook into the kernel and abides by all APIs).

The agent is supported on Windows XP and up, MacOS, Linux, Docker and Chrome OS, with distributions across x86, ARM and MIPS architectures (32 & 64 bit). Custom builds of the agent for OpenBSD and Solaris can be made upon request.

Onboarding Checklist

LimaCharlie is a large, evolving set of terms and concepts, all of which are document here: doc.limacharlie.io

This checklist is meant as a discussion on topics that are relevant prior to onboarding your first customer.

Organization

Adding new organizations is simple. Organizations are free and begin with a two agent free tier. If you hit the maximum number of organizations just get in touch with us and we'll increase your limit.

Scaling up and down as well as granting access to an organization is easy. This means the first step to onboarding a new customer, even if only for a trial, should be to create an organization.

Users

LimaCharlie supports role-based access control (RBAC) for user accounts. Through the web interface you can create additional users while controlling what it is they are able to see and do. Use pre-built access control templates designed for account owners, administrators, operators, view-only or create your own.

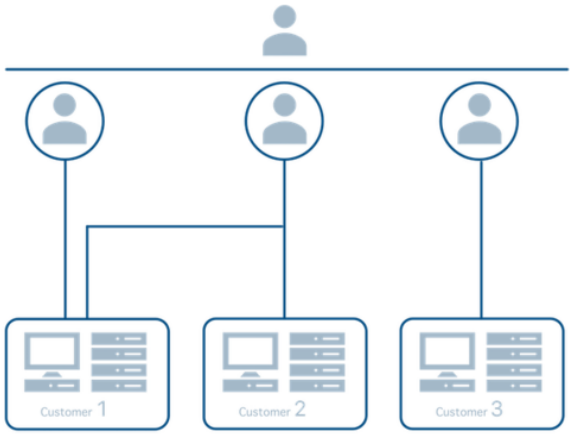

Installation Keys

The Installation Keys (IK) are used to enroll new agents. One IK can be reused as many times as you want and it remains valid until it is deleted. Agents that are already enrolled will keep their access but new agents will not be able to enroll with the deleted key. This model makes it easy to revoke keys and maintain tight control over installer agents.

The main reason to have multiple IKs is to associate tags with agents. Each IK has a list of tags that are attached to agents enrolled using it. For an MSSP this provides an easy first pass at categorizing agents.

For example, you can create a key with the tag “server”, and another with “desktop”. Then it is just a matter of asking your customer to use the relevant IK for the various assets.

Tags can also be added/removed either programmatically or interactively after enrolment.

Outputs

Outputs allow you to forward data to your infrastructure. If storage is enabled all telemetry will automatically be stored and accessible through the LimaCharlie infrastructure. Storage does not have to be enabled for Outputs to work. Outputs allow the user to direct an adjustable amount of telemetry to any location. For example, users can direct the full telemetry to cold storage on S3 and send alert level data to Slack. The data is composed of three main streams.

1. Event
2. Detect
3. Audit

The audit stream allows you to choose where you want to store a copy of the audit messages. It is not technically required but we recommend sending it to a write-only location like an AWS S3 bucket.

Next is the detection stream. This is high priority data and you will want it to go to a live system rather than an archiving one. It is common to send it to the SIEM or a product like Splunk.

If you find some of the detections to be overly verbose you can also customize the outputs so that only certain detection types go to a cold storage and others to a live system.

Finally, the event stream is by far the most verbose. It contains the raw events coming from the agents. Most cost-conscious MSSPs will, as a general rule, send all the events to cold storage. A popular method is to use an AWS S3 bucket with a retention period set (like 3 months), and to upload data to it with compression enabled.

The event data is very verbose so that it can be easily used operationally, but it also compresses extremely well (~90%). So for cold storage, enabling compression is a good idea.

Another common setup is to tag sensitive agents as “vip” and configure an Output to send the event stream from these agents to a live system like Splunk so that it is always readily available.

Insight [telemetry storage]

Long-term telemetry storage with search capability is available as an optional add-on. LimaCharlie still operates as middleware but with this feature enabled we offer one year of storage and search capability for a small additional cost. This move allows MSSPs and SOCs who do not already have their own EDR infrastructure to gain a completely functional information security centre upon signing up.

Once enabled, LimaCharlie Insight will automatically send all telemetry data to secure storage on the Google Cloud Platform. The Insight user interface allows you to select a date and time from which to start and then perform complex investigations.

Automation

The first few organizations you onboard will likely be done manually. But you will quickly want to clone your configuration and store it for use setting up future organizations. LimaCharlie provides tools to export your configuration and store your “infrastructure as code”.

Using the sync tool part of the Python API, you can backup an organization’s configuration as YAML files and organize them so that you re-use base configurations across all your customers.

The tool will allow you to export to YAML and apply the YAML back up to the cloud. This results in a more efficient onboarding process and a better repeatable process.

For a more complete discussion, see the related blog post: Infrastructure as Code

Additional Resources

Full documentation can be found here: doc.limacharlie.io

REST API documentation via Swagger here: Swagger Doc

Free online course here: edu.limacharlie.io