

AI-Powered Search User Guide



Table of contents

01	How to Get Started	3-6
-Setting up AI-powered search for your organization		
-Using AI-powered search		
-Using AI-powered search with reverse image analysis		
02	Optimal Query Structure & Best Practices for AI-Powered Search	8
03	Example Queries to Get Started	11
04	AI-Powered Alerts	12
-Generating an alert using quick alert creation		
05	Limitations	22
06	Query Moderation & Responsible Use	23
07	FAQ	24



Overview

Verkada's AI-powered search leverages state-of-the-art large language and vision models that let users search for relevant camera footage of people and vehicles by uploading an image or using everyday language, making investigations more intuitive and efficient. Users in Command are no longer restricted to a predefined list of search attributes of people or vehicles and can instead enter freeform text searches ("queries") or upload photos directly into the search bar to get more granular results. Users in Command can create searches like "person not wearing a safety vest" or "red 1970s muscle car outside on June 26th" to conduct highly-detailed investigations. In addition to using freeform text with AI-powered search, organizations can upload a photo of a person, vehicle, or object to Command to obtain relevant search results. Akin to Google's Reverse Image Search, this functionality pulls multiple aspects of an image, like shapes, pixels, patterns, and other granular details, then delivers the likeliest set of matches from one's camera fleet. Users can turn AI-powered search queries into actionable alerts for their organization, with a feature called AI-powered alerts (see page 12 for more detail on AI-powered alerts).

01

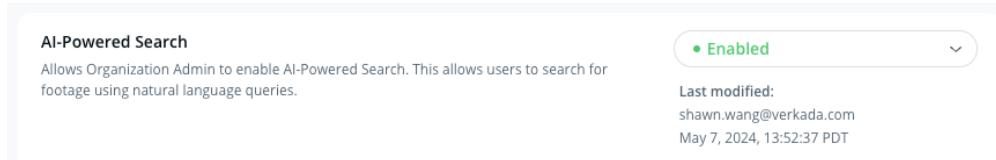
How to Get Started and Create Actionable Investigations with AI-Powered Search

Setting up AI-powered search for your organization

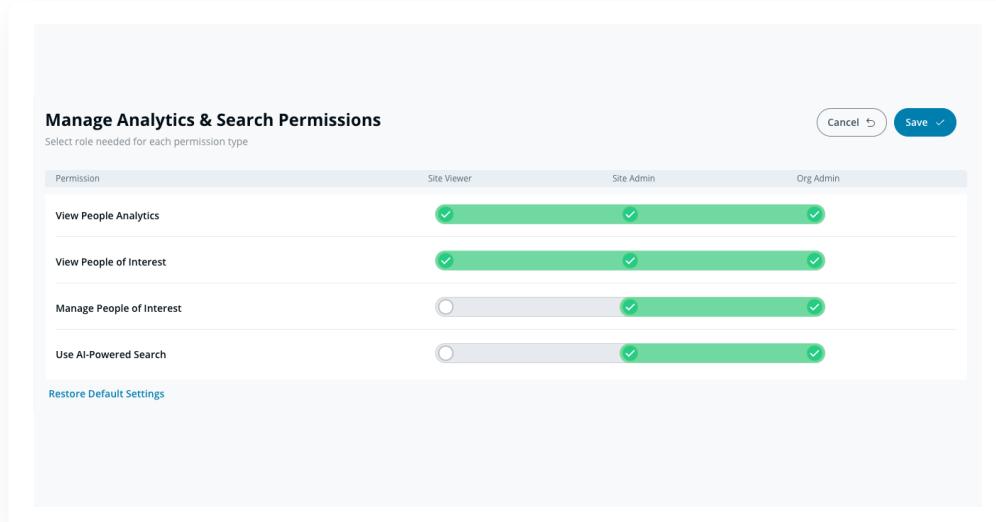
1. Organization administrators must first enable AI-powered search: i. have People and Vehicle Analytics enabled; ii. enable AI-powered search in the "Feature Manager" tab under "Privacy & Security" settings, pictured below under Step 1B and iii. offer the requisite permissions for "AI-powered search" to the desired group (i.e., "Site Viewer" or "Site Admin") within the "Roles & Permissions Customization" tab under "Org Settings," as pictured below under Step 1C.

Note: searches will not return results for Vehicle or People History if they are disabled. You can enable them by toggling them on in the "Feature Manager," under "Privacy."

Step 1B. for organization administrators enabling AI-powered search.



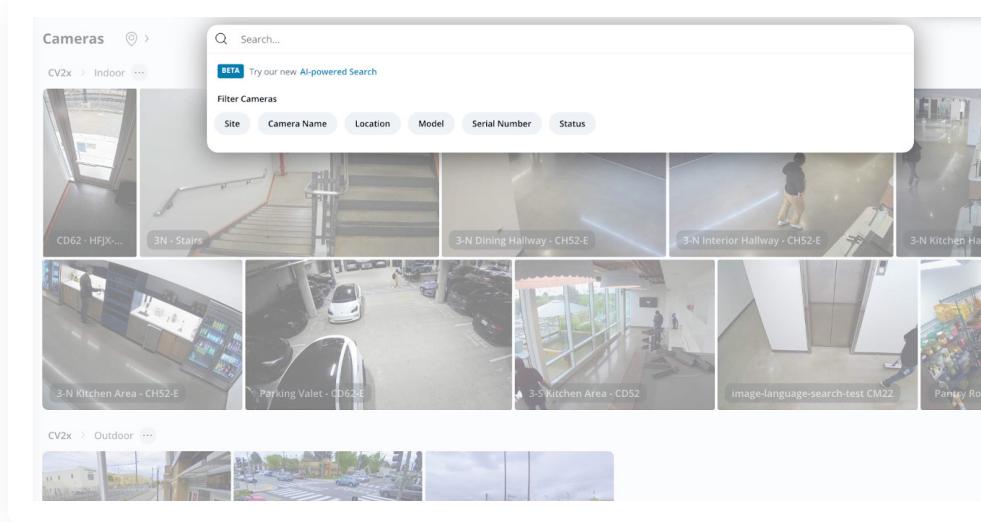
Step 1C. for organization administrators enabling AI-powered search.





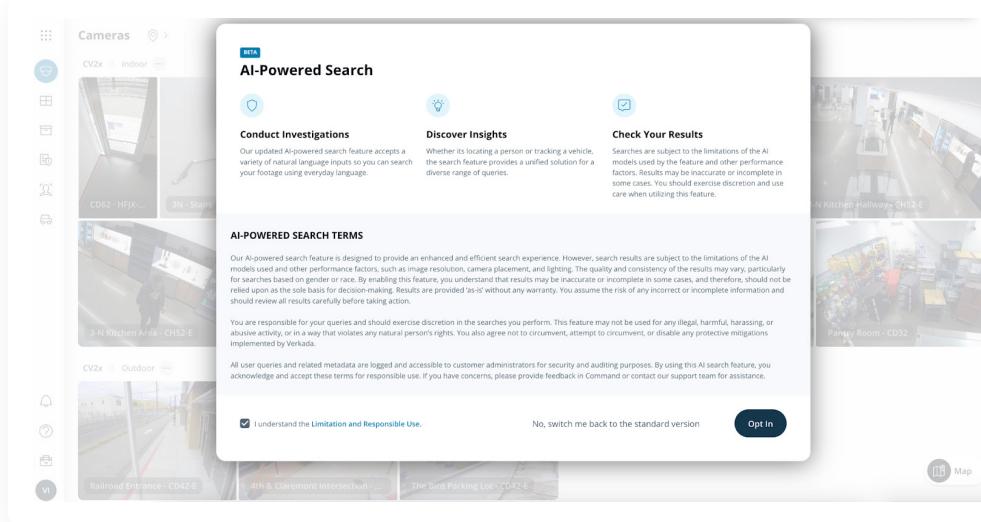
Using AI-powered search

2. Log in to Command and you'll see a prompt to "Try our new AI-powered search."



3. Once permissions are enabled and the user clicks "Try our new AI-powered search," the user will then be prompted to accept the AI-powered search Terms and Conditions (Terms). Please read through them carefully so you understand the capabilities and limitations of this feature.

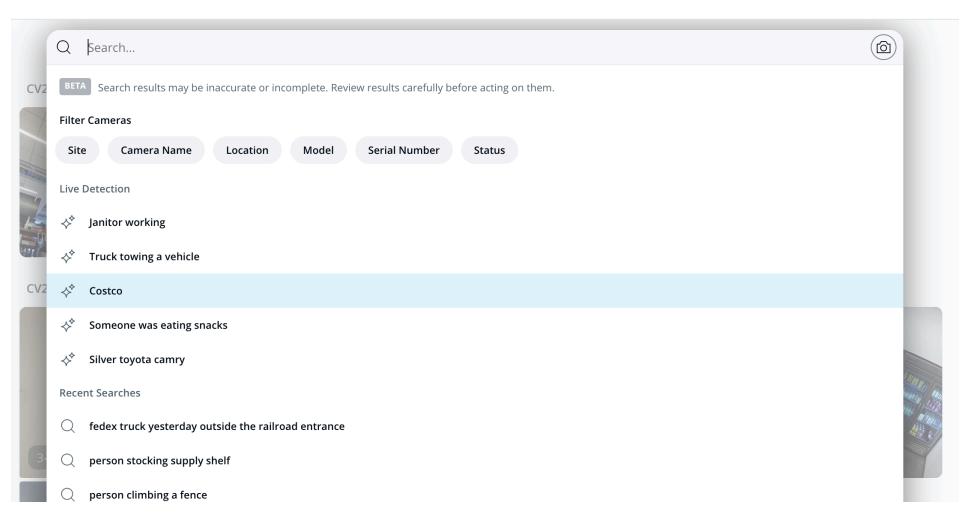
Note: the user will only need to opt in once.



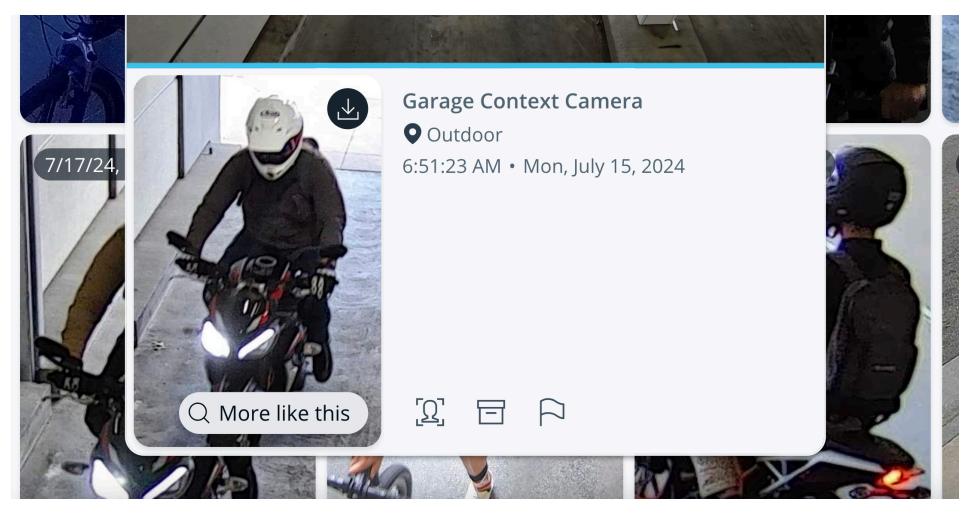


4. Upon accepting the Terms, the user will see a search bar that accepts virtually any text—a blank canvas for writing queries in everyday language, directly relevant to an organization's investigations. Users can also upload a photo to the AI-powered search bar using the camera icon on the right-hand side. This reverse image upload analyzes multiple aspects of an image, like shapes, pixels, patterns, and other granular details, then delivers the likeliest set of matches from one's camera fleet. See page 6 for more information on reverse image analysis with AI-powered search.

New queries search across all cameras by default, but users can also organize searches by camera sites, camera type, leverage searches of popular footage in one's organization (under the "Live Detection" banner) or revisit recent searches.



5. Hovering over a result gives the user the option to archive the clip (via the small box icon), flag the result to submit feedback, or view other similar results with the "More Like This" button.

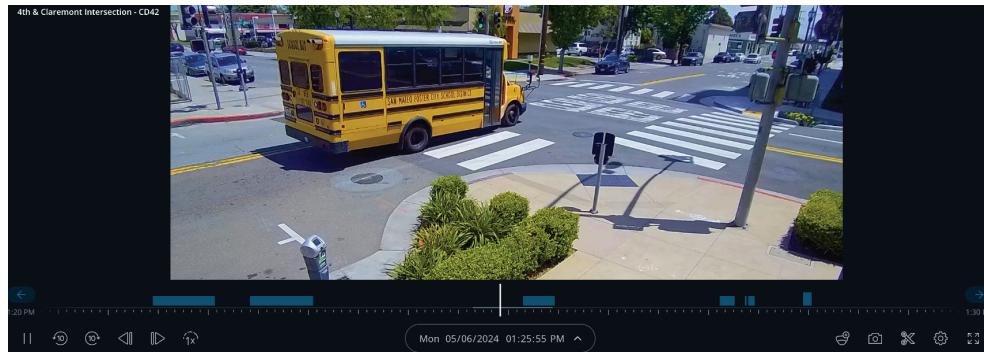


6. Clicking into the "More Like This" button searches through all footage to retrieve results most similar to the one the user selected.



The screenshot shows the Verkada search interface. At the top, there is a search bar with the query "person wearing a helmet but not protective padding", a "Clear" button, and a "Most Relevant" dropdown. Below the search bar, a message states: "This feature is in beta, search results may be inaccurate or incomplete in some cases. Review results carefully and act with discretion. Report [Issues](#) or [provide feedback](#)." A sidebar on the left contains various icons for navigation and settings. The main area displays a grid of 10 thumbnail images, each with a timestamp. The first thumbnail is labeled "Garage Entrance - CD41-E Outdoor 3:58:58 PM-Mon, July 8, 2024". A "Map" button is located in the bottom right corner of the thumbnail grid.

7. Clicking into a search result brings the user to the specific footage in the history player.



Clicking into a result for the query "San Mateo School Bus" brings the user to the history player at the time of the requested footage, as depicted above.

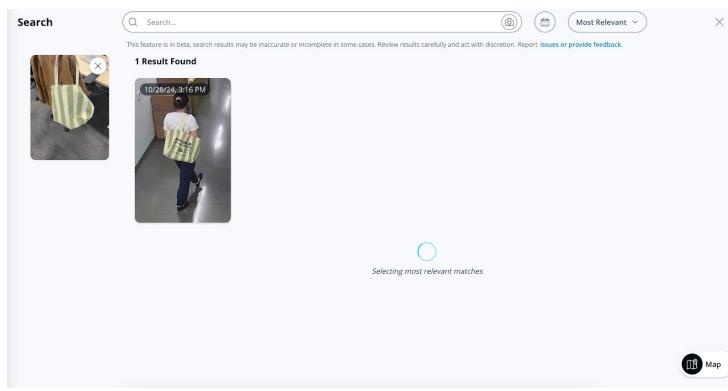
AI-powered search with reverse image analysis

8. Clicking the camera icon on the right-hand side allows a user to upload an image for further analysis. This reverse image search functionality pulls multiple aspects of an image, like shapes, pixels, patterns, and other granular details, then delivers the likeliest set of matches from one's camera fleet. While the results of this search will display thumbnails of people and vehicles,¹ organizations can use it to see if non-person and non-vehicle objects were carried, held, worn, or attached to a person or vehicle. Organizations can also use this functionality to help identify individuals based on other attributes of the uploaded image when an individual's face is partially visible or blurred.

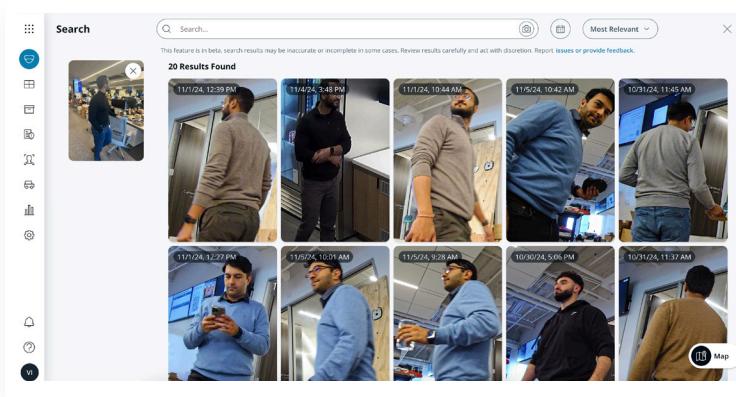
The screenshot shows the Verkada AI-powered search by image upload interface. At the top, there is a search bar with the placeholder "Q. AI-powered search by image upload..." and a camera icon. Below the search bar, a message states: "BETA Search results may be inaccurate or incomplete. Review results carefully before acting on them." A sub-instruction says "Search Verkada footage with a photo." A large dashed box with an upward arrow icon is provided for uploading a photo. Below the box, text reads: "Drop photo here or [choose a photo](#) to upload (JPG or PNG up to 15MB)".



As illustrated in the example below, organizations can upload a photo of a standalone object (a bag in this case) to retrieve results of any individual(s) who may have held (and left behind) the specific object – a use case with clear security benefits:

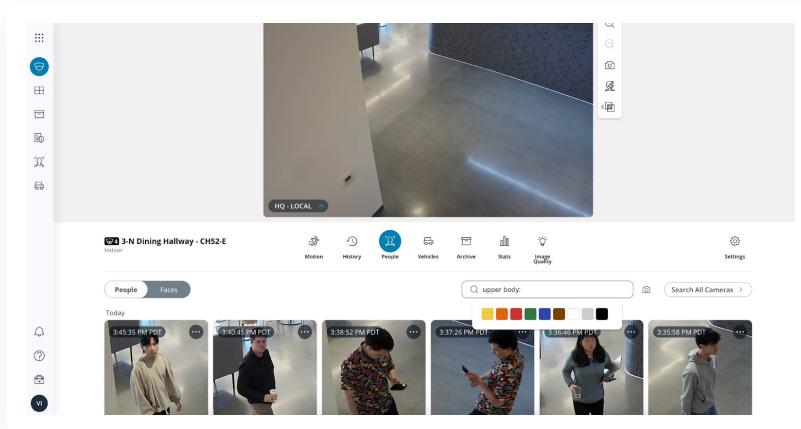


In another example below, we illustrate how this functionality can help identify individuals whose faces are partially visible—analyzing instead other granular aspects of the uploaded image:



Upon clicking into a specific camera on the home page, the user will see our classic attribute search with a predefined list of criteria (like selecting shirt color)—not AI-powered freeform search, as depicted below.

Note: AI-powered search only currently exists on the home camera screen page.



1. AI-powered reverse image search results leverage “hyperzooms,” which are highly-detailed segments of a video frame that contain images of a person or a vehicle only, typically the most important parts of an investigation. If an uploaded image does not include a person or vehicle (i.e., a standalone object), results will be directed to “hyperzooms” of people and vehicles that include the object (e.g., if uploading a standalone suitcase, the search will return results of a person holding a suitcase or a suitcase attached to a vehicle).

02

Optimal Query Structure and Suggested Queries for Useful Investigations

You can read more about the underlying models that support this feature in our [white paper](#) here to understand why we suggest structuring queries in the method we highlight below. We recommend the following “best practices” when using AI-powered search:

Best practices

1. Use AI-powered search for highly-detailed attribute searching for people or vehicles, not standalone objects

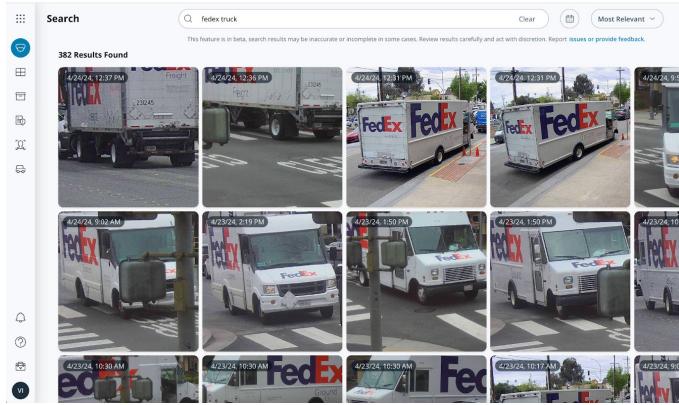
AI-powered search leverages “hyperzooms,” (HZs), which are highly-detailed segments of a video frame that contain images of a person or a vehicle only, typically the most important parts of an investigation. If your query does not include a person or vehicle, results will be directed to HZs of people and vehicles that include the object (e.g., if searching for “ski gear,” the search will return results of ski gear on a person or attached to a vehicle).

2. Use AI-powered search to filter further by time and space

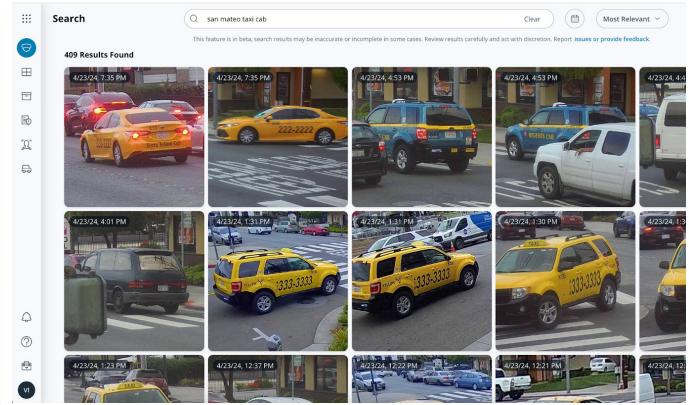
AI-powered search includes a concept known as “parsing,” or the ability to include date, time and location parameters in a query. To further narrow a query using parsing, a user can add to the search time specifications like a date range or specific day (e.g., “UPS truck between April 5th and April 6th,” or “UPS truck yesterday”), a time range (e.g., “person holding tote bag between 9am and 10am”), or a location tied to a camera (e.g., “people crowding inside the main lobby,” or “white van parked in the exterior parking lot”).

3. AI-powered search can be used to identify text, logos and brands

AI-powered search can assist you in identifying specific objects people hold or vehicles driven via recognition of **sufficiently visible** text, logos and brands, as pictured below. **Note:** this model was not trained using Optical Character Recognition (OCR) and therefore has limitations on identifying text and numbers. Please refer to the FAQ below for additional details.



Text recognition works best when the text, logo or brand is clearly visible as illustrated in the results above for “FedEx truck.”

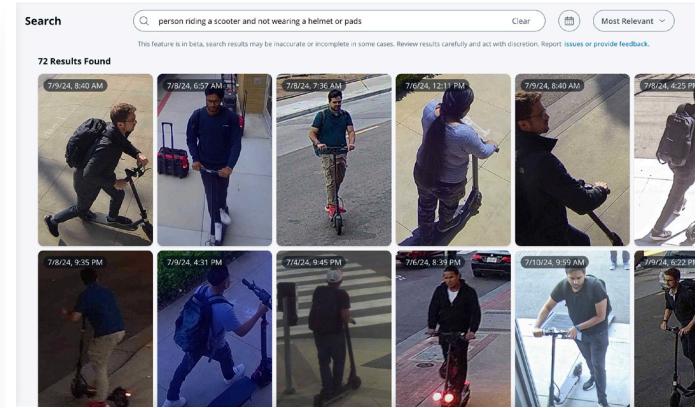
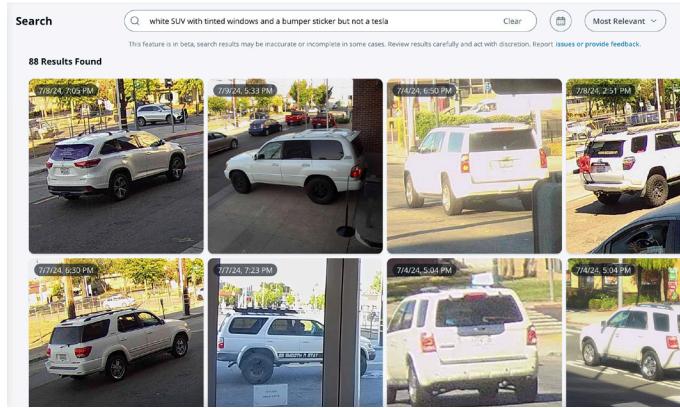


A search for “San Mateo taxi cab” returns mixed results. The model has detected the words “taxi” and “cab,” but “San Mateo” isn’t clearly visible in any of the results. Customers should note these limitations when attempting text-matching searches.



4. Use AI-powered search to search for multiple attributes of people and vehicles, or to search for attributes that are not present

AI-powered search supports queries that include the words "AND," "NOT," and "OR" so you can conduct searches with as much detail as you require. Use "AND" and "OR" to search for several attributes of people and vehicles and use "NOT" to exclude attributes or identify missing attributes.

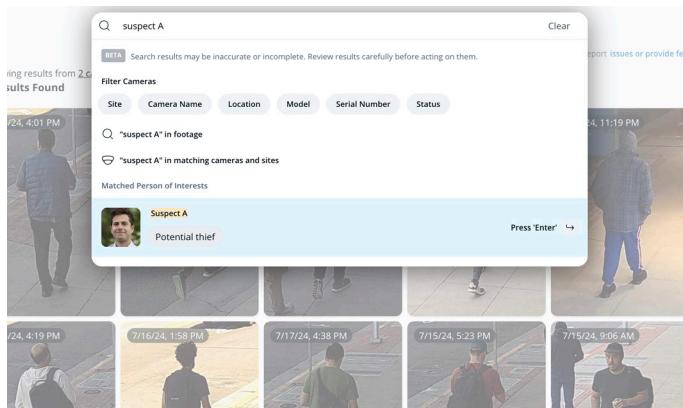


A highly-detailed vehicle query using "AND" and "NOT" for "White SUV with tinted windows and a bumper sticker but not a Tesla."

A highly-detailed person query using "AND," "NOT" and "OR" for "Person riding a scooter and not wearing a helmet or pads."

5. Include POIs and LPOIs in your search bar to quickly surface the time and location of known persons or license plates of interest

After storing a person or license plate of interest (LPOI), use the name of that saved individual or the specific alphanumeric digits of a saved license plate in the search bar to identify where that individual or vehicle appeared on camera footage and at what time.*



***Note:** we currently do not support POIs and LPOIs in structured queries as outlined on the following page. In other words, we do not currently support the ability to add the names of saved POIs and the alphanumeric digits of saved LPOIs as part of larger attribute searches. One can search, for example, "Suspect A in the garage," but not "Suspect A wearing a hoodie next to a van in the garage" (assuming "Suspect A" is the name of a saved POI).



With these best practices in mind, we suggest structuring queries as follows:

[person or vehicle] + [the specific action and/or descriptor of that person or vehicle] + [OPTIONAL “and” / “or” / “not” / “between” / “on” / “in” / “at” / “with”] + [OPTIONAL second or third descriptor] + AND/OR [OPTIONAL date / time / location]

Four example queries using the above structure:

“Person **wearing a green scarf **with** red hat **between** 8am and 9am **outside** the main entrance”**

“White Ford **pickup truck **with** red bumper sticker **yesterday**”**

“Person **unloading crates **and not** wearing a safety vest or safety goggles **at** the loading dock”**

“White van **with a red bumper sticker **and** a large scratch **outside** the parking garage **yesterday**”**

Important note on query inputs:

As the foundation model was trained on still images, it will perform less accurately on specific actions. Verbs such as “running” or “stealing” might produce less accurate results than “wearing” or “holding,” as “wearing” and “holding” can be more easily identified from still images than “running” or “stealing.”



03

Example Queries to Get Started

Retail (anti-theft)

- "Person with duffle bag between 9 and 10 am"
- "Person holding [insert retail item]"
- "Person wearing Patagonia vest with a scarf and blue jeans"
- "Person holding a Gucci handbag but not a Fendi or Prada handbag"
- "Person wearing a Louis Vuitton or Armani sweater"
- "Person carrying large shopping bag"

General security

- "Person climbing over fence in [insert specific camera or site] on April 5th"
- "Person scaling wall"
- "Person looking at camera"
- "Person wearing balaclava and holding a large bag"
- "Person with badge on chest"
- "Person walking near Sidewalk Closed sign"
- "Person near Amazon box"
- "People fighting or arguing with each other"

General people search enhancements

- "Person wearing green hat and a scarf and sunglasses"
- "Man with blue suitcase and brown jacket but not with sunglasses"
- "Person in spiderman costume"
- "Man wearing New York shirt"
- "Man holding an apple"
- "Person riding a bike and not wearing a helmet"
- "Kid riding a scooter and not wearing pads"
- "Person with brown hair looking at their phone"
- "Person playing with their phone"
- "Person making a phone call yesterday"
- "Person walking a dog and wearing headphones"
- "Person wearing a 49ers jersey"
- "Delivery person wearing an Amazon shirt and not a Fedex shirt"
- "Person wearing Canada Goose jacket"
- "Person holding a Redbull can and wearing sweatpants and a hoodie on April 7th"
- "Kids playing"
- "Person holding a cat and not a dog"

Education (student safety, rules compliance)

- "Person smoking"
- "People in a crowd in the parking lot"
- "People on football field between 12 am and 4 am"
- "Cars in parking lot between 8pm and 12am on April 6th"
- "[Insert school district name] school bus"

Manufacturing (workplace safety, compliance)

- "Person wearing a hardhat with safety vest"
- "Person not wearing a hardhat with a safety vest"
- "Person wearing a hardhat but not wearing a safety vest"
- "Person wearing a safety vest in [insert specific camera or site] on April 3rd"
- "Person carrying a box around 11am but not wearing PPE"
- "Person driving a forklift in [insert specific camera or site] on June 16th"
- "Person near wet floor"
- "Person carrying crates"
- "Person stocking supply shelf in kitchen or in pantry on April 15th"
- "Person using [insert tool, machine or instrument]"

General vehicle search enhancements

- "White Ford Explorer with tinted windows"
- "1970s red muscle car with a white roof"
- "Antique pickup truck with a roof rack and a bumper sticker"
- "FedEx truck at [insert specific camera or site]"
- "Gray BMW X7"
- "Black Tesla Model Y on April 4th"
- "White van with tinted windows and a bumper sticker and a scratch"
- "Blue Porsche Taycan"
- "Most expensive car"
- "McLaren"
- "Porsche 911"
- "Gray 1960s sports car"
- "1994 or 1995 Jeep"



04

AI-Powered Alerts

In addition to finding relevant footage using AI-powered search, customers who have enabled and opted into AI-powered search per the “How to Get Started” section on pages 3 - 4 above, can also receive alerts in real-time to the presence of the people or vehicles subject to their queries. Users can already use the Alerts section of Command to create dozens of alerts across Verkada products for notable events like a person or vehicle loitering in a specific area, a person tailgating in an office door, or a camera losing connectivity.

With AI-powered alerts, users can create alerts for any time their cameras detect any of the user’s queries. While AI-powered search proves useful for finding footage in the past, AI-powered alerts enable customers to receive notifications via text, email, or mobile app when the results of specific freeform queries appear in their footage. Users can use the queries in the “Example Queries” section on page 11, or create almost any query of relevance, to trigger real-time notifications of results. AI-powered alerts supports proactive, highly-customizable management of one’s premises.

Generating an alert using quick alert creation

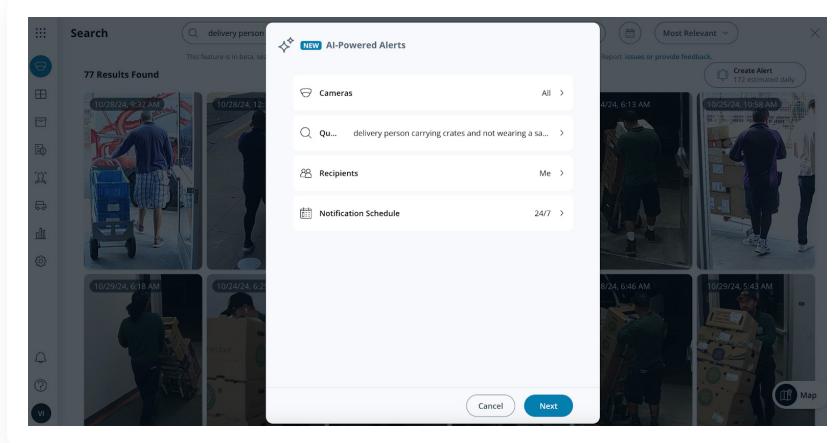
There are two paths for creating an AI-powered alert: quick alert creation from an existing query and creating an alert from scratch in the “Alerts” tab in Command. When a user writes a freeform query, the user will also have the option to create an alert from that specific query (via the “Create Alert” button on the right-hand side), as pictured below. The user will also see the number of estimated daily alerts (174 in this example) one will receive from the query.



A query for a workplace safety use case (“delivery person carrying crates and not wearing a safety vest”) not only shows corresponding results, but also gives the user the option (via the “Create Alert” button on the right) to be alerted to when results of this query appear on one’s footage.

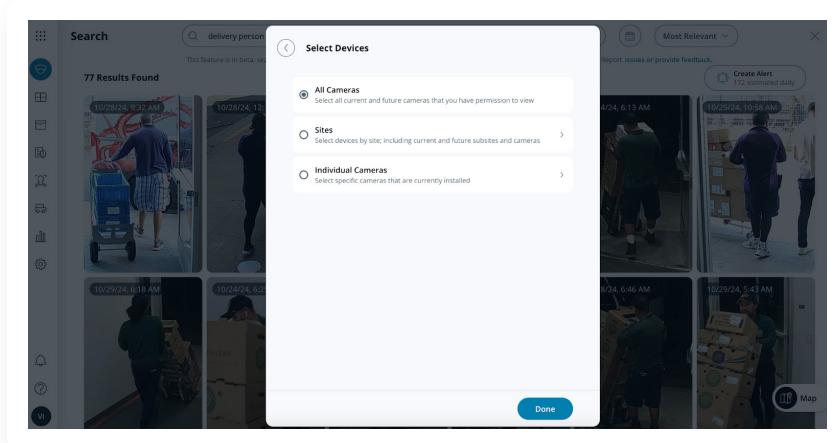


Clicking into the “Create Alert” button gives the user the option to further customize the alert’s prepopulated parameters: Cameras, Query, Recipients, and Notification Schedule. Below we expand on each parameter so users can understand how to optimize notification(s) for their organization’s needs:

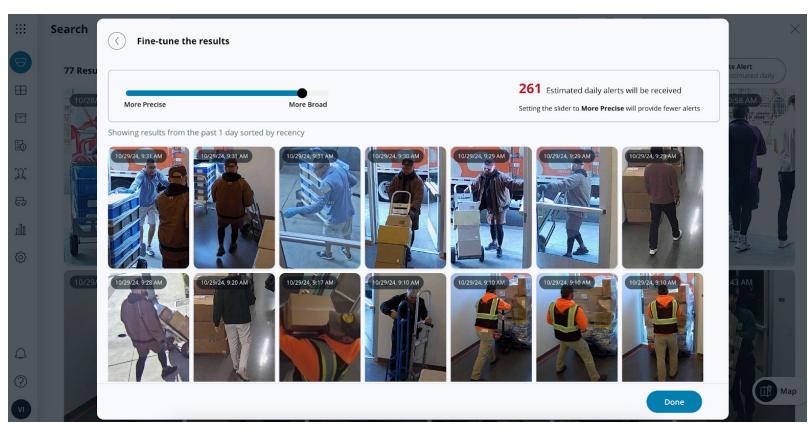


Quick alert creation offers a default list of alert criteria to expedite alert creation.

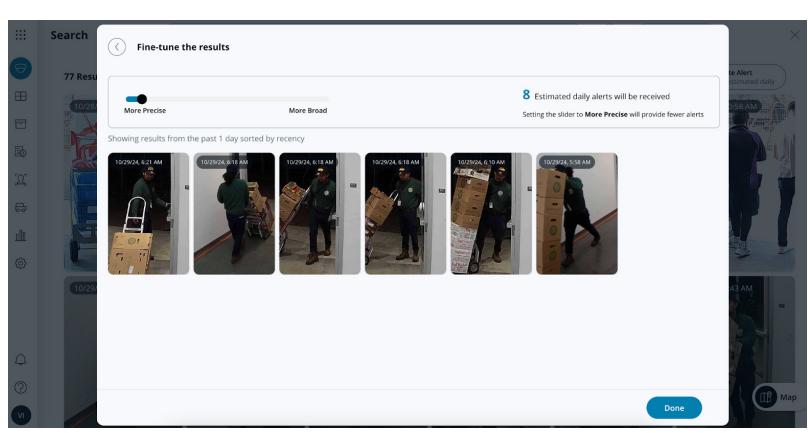
Cameras: as pictured below, users can choose the specific cameras to which the alert applies. Users can elect to have the alert apply to a single camera, a site, or all cameras in their organizations.



Query: using quick alert creation, the user’s query (“delivery person carrying crates and not wearing a safety vest”) prepopulates. The user can, however, click into “Query” and fine-tune the accuracy of the results with a slider as desired to receive fewer notifications (“More Precise”) or more notifications (“More Broad”).

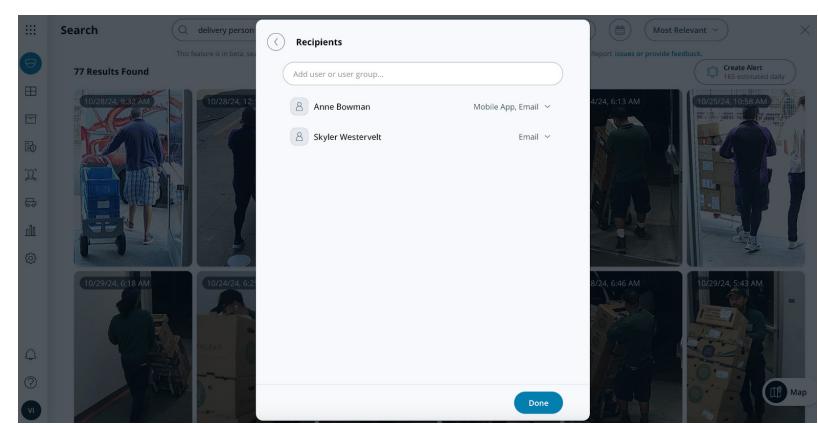


Moving the slider to the right towards “More Broad” increases the number of estimated daily alerts associated with the query (“delivery person carrying crates and not wearing a safety vest”) to 261.



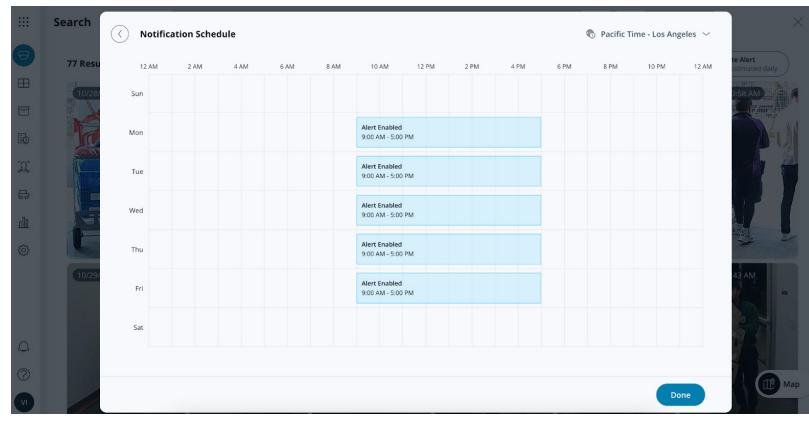
Moving the slider to the left towards “More Precise” decreases the number of estimated daily alerts associated with the query (“delivery person carrying crates and not wearing a safety vest”) to 8.

Recipients: the user then has the option to decide who and how individuals in the organization receive the alerts. Users can be notified via Command Mobile App, email, or text.

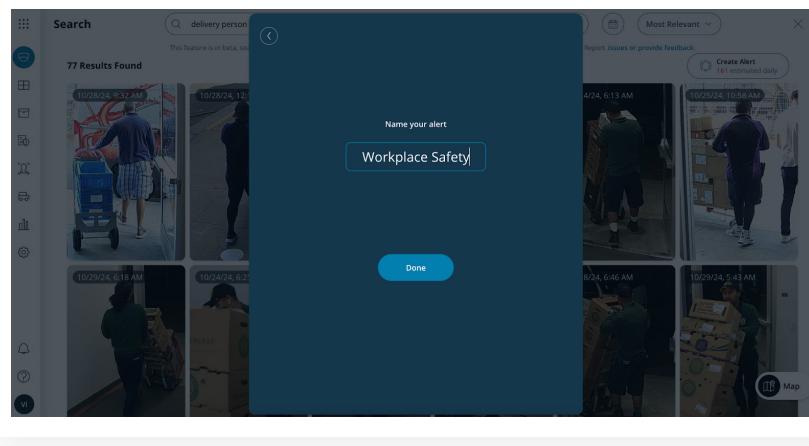




Notification Schedule: users can customize when alerts are received. Perhaps, as illustrated below, the user only requires notifications during workdays between 9am and 5pm.



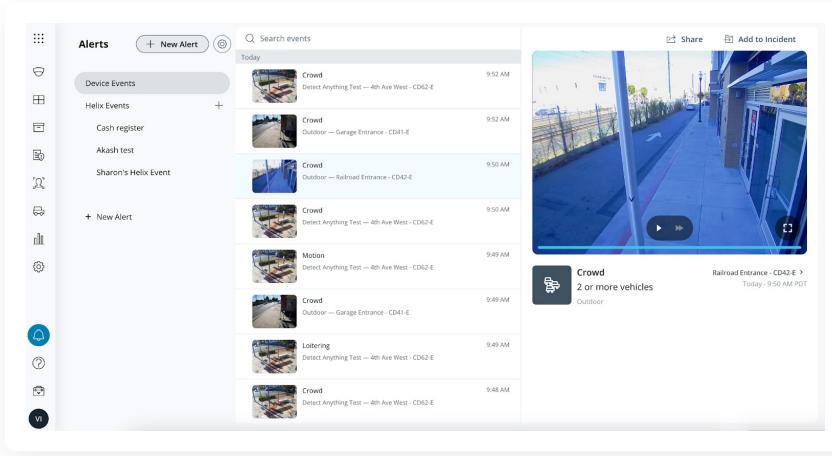
Finally, the user then names the alert ("Workplace Safety") in this example.



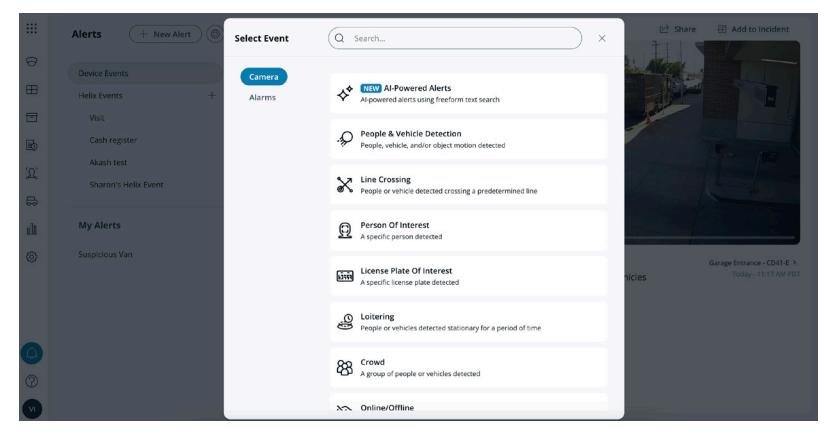


Generating an alert from the Alert creation page in Command

Users can also create an AI-powered alert directly from the Alerts section in Command and follow a similar setup flow. First, navigate to the Alerts section of Command (the bell symbol on the left-hand side) and select “New Alert” in the top-left, as depicted below:

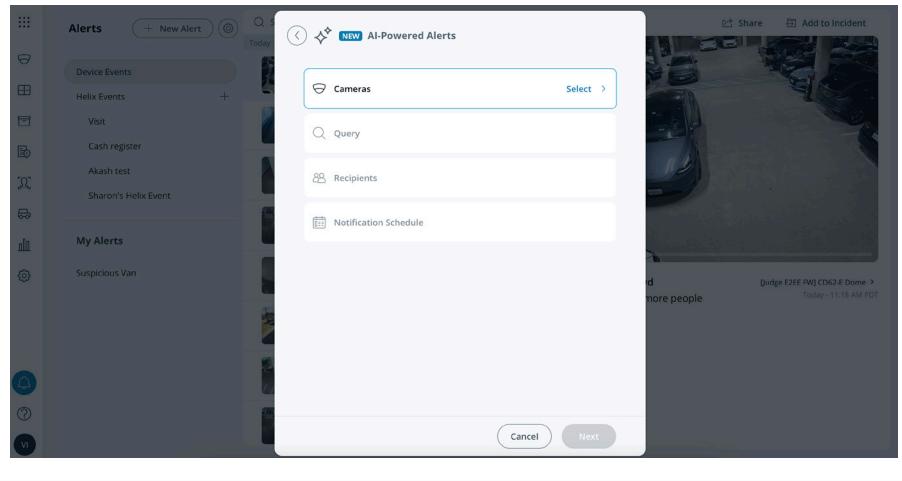


Clicking into “New Alert” will return the typical alerting page, and users will see “AI-Powered Alerts” directly at the top:

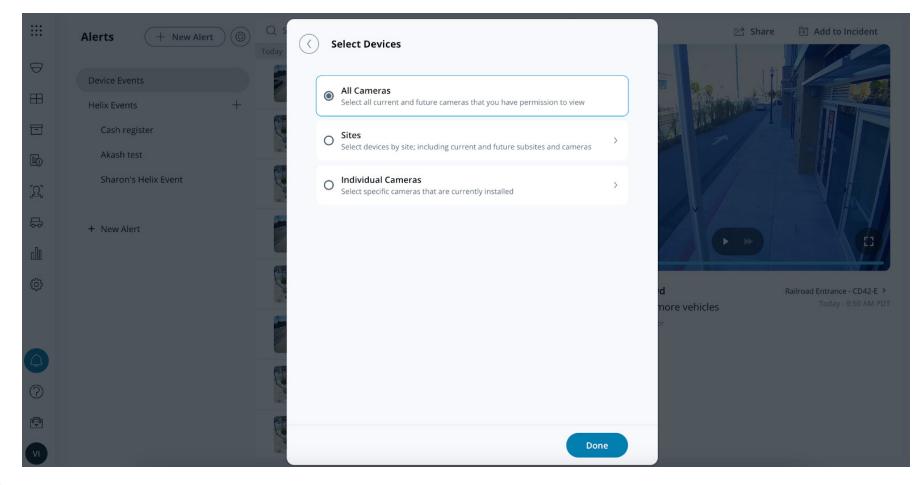




Once selected, the user will follow a similar setup flow as outlined in the quick alert creation section above. The user will see the same list of criteria (pictured below), but without any prepopulated information.

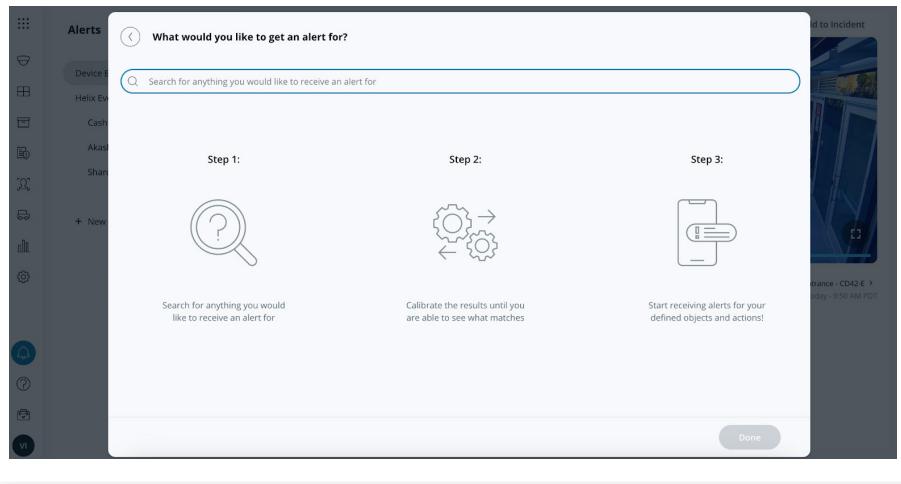


In the same manner as the quick alert creation example above, the user first selects the cameras in their organization to which the alert applies. In this example, we'll select "All Cameras".

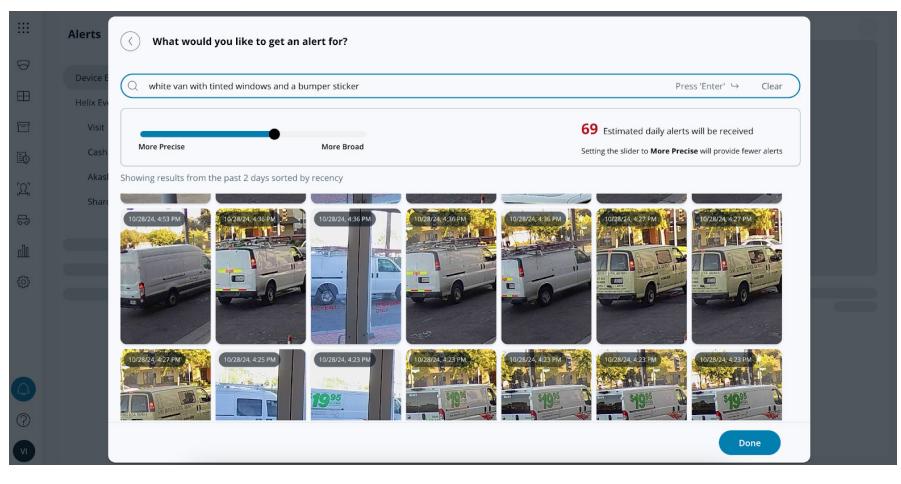




Clicking then into "Query," the user will be prompted to set up their alert using AI-powered freeform text, as illustrated below:

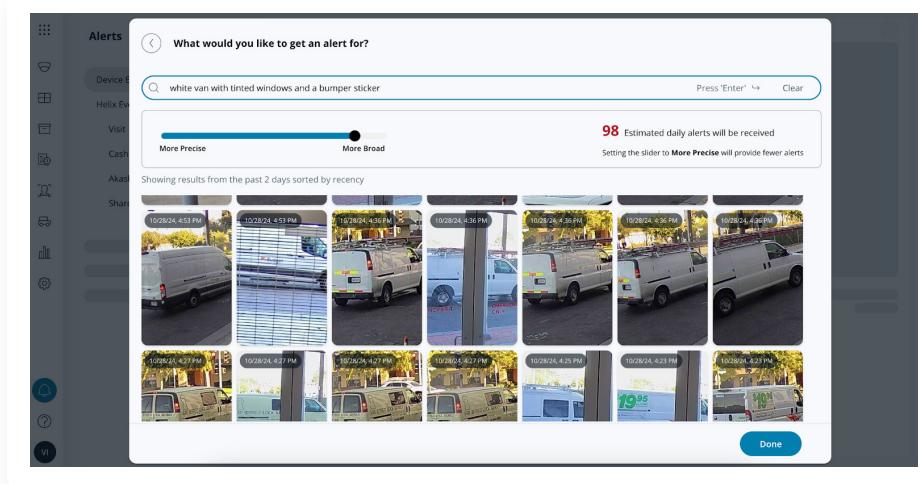


Focusing on a security use case, we can create an alert to see when a suspicious vehicle with specific attributes appears on any camera in an organization's fleet. In this example, we'll create an alert for "white van with tinted windows and a bumper sticker."

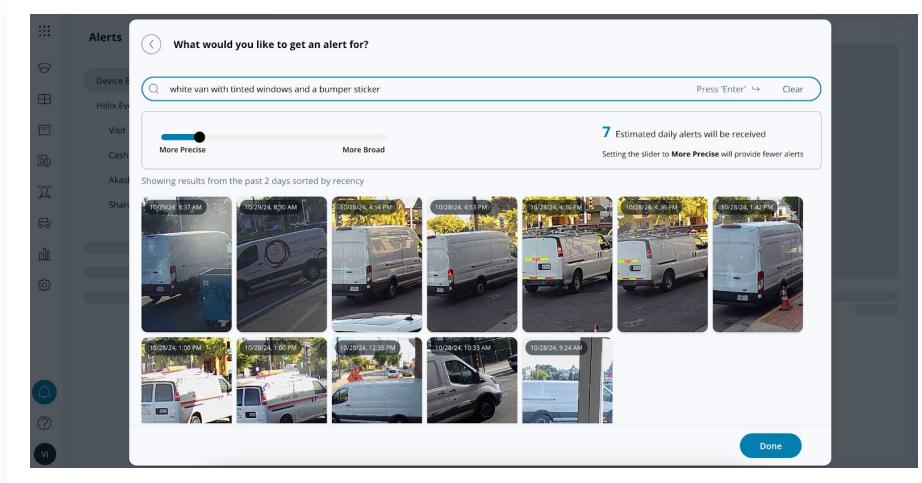




Using the same fine-tuning method as we used in the quick alert creation example above, we can opt for greater accuracy of results (fewer notifications) or less accuracy (more notifications) using the slider on the left-hand side.



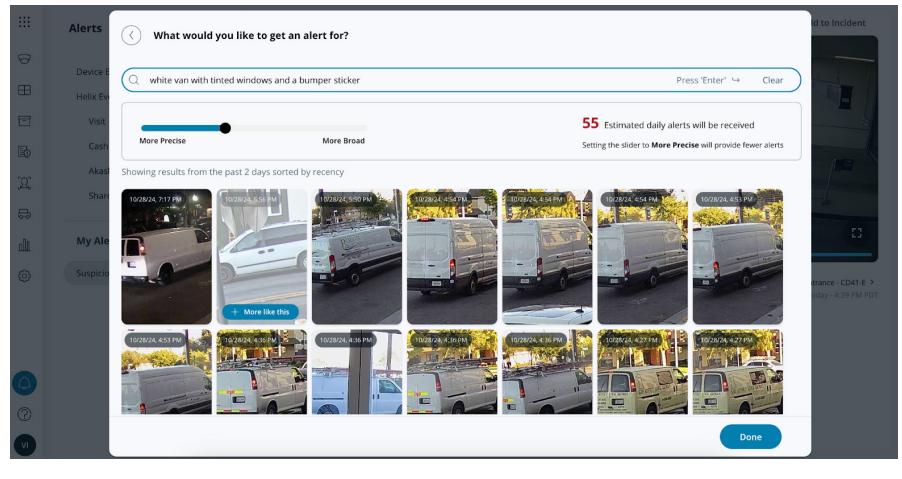
Moving the slider to the right ("More Broad") generates more estimated daily results for this specific query (98 in this example).



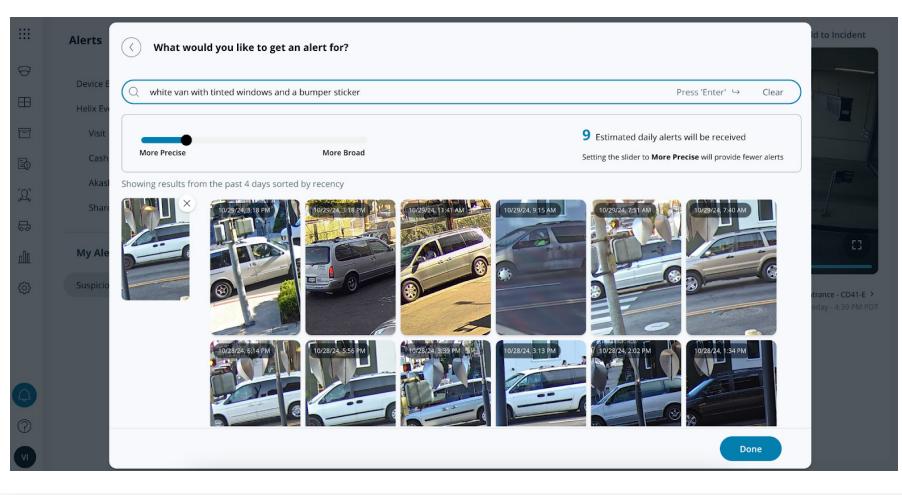
Moving the slider to the left ("More Precise") generates fewer estimated daily results for this specific query (7 in this example).



If users see a specific result in the corresponding thumbnails that more closely aligns with the desired result, users can click “More like this” to create alerts more similar to the selected image.

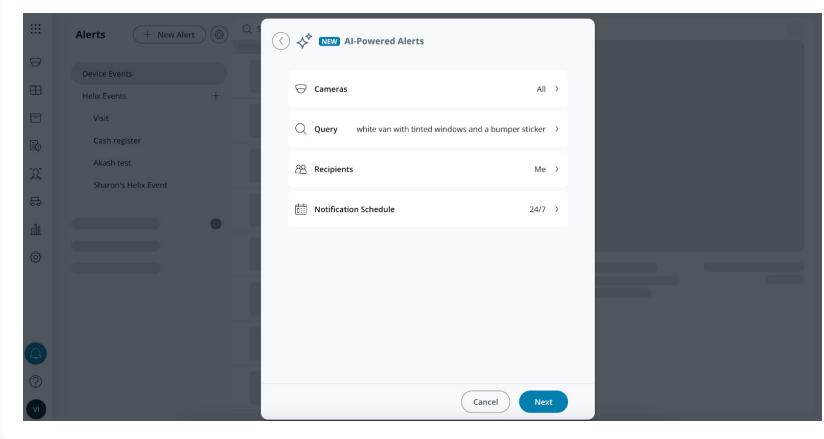


Clicking “More like this” will recalibrate the results accordingly to be more consistent with various aspects (i.e., pixels, shapes, colors) of the selected image.

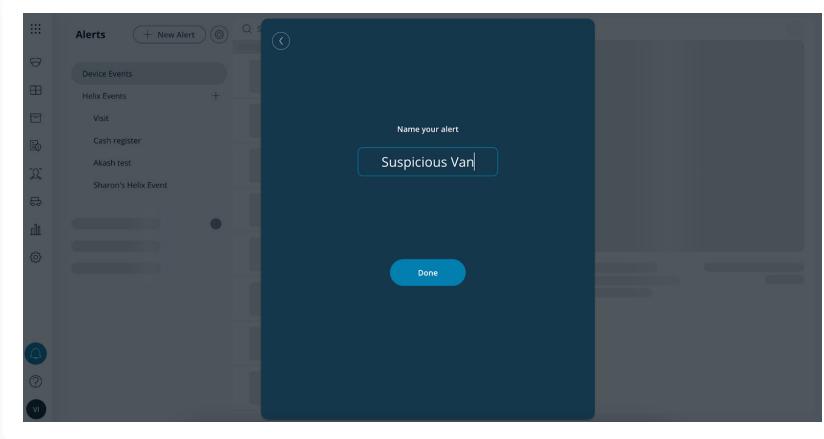




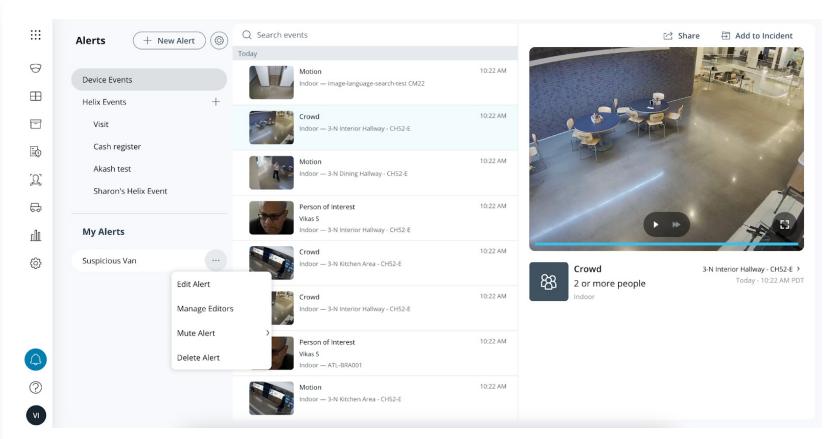
Following the same method as quick alert creation above, the user can then customize which organizational users receive the corresponding alert (via Command Mobile App, email, or text) and can further specify when the alert should be received. In this example, we'll stick with the default 24/7 notification schedule as we'd like to be notified anytime the suspicious van appears on any cameras.



The user can then name the alert accordingly ("Suspicious Van" in this example).



Finally, the user can then edit, mute, and otherwise manage the alert on an ongoing basis from the Alerts page in Command.

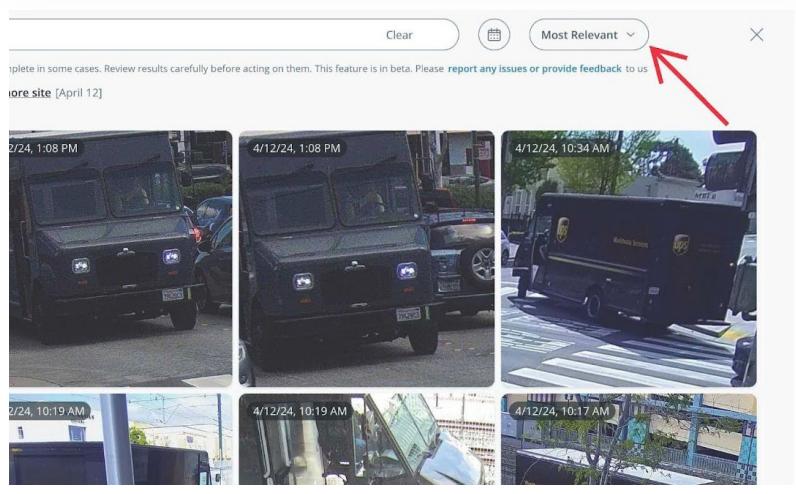
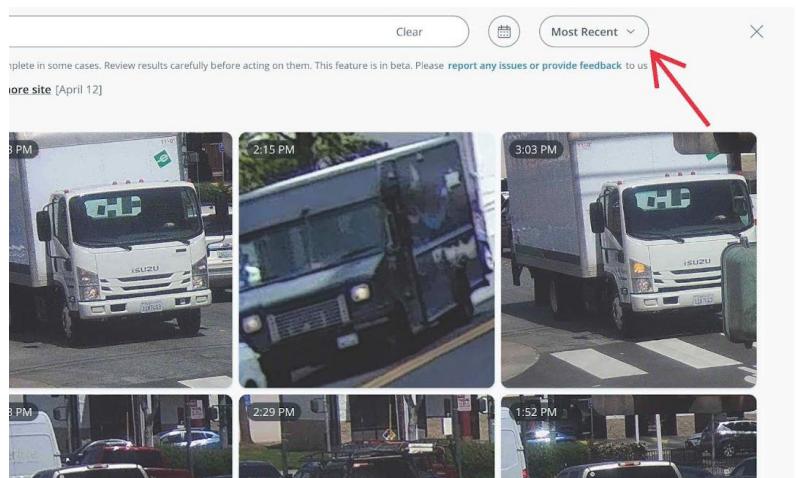




05

Limitations

Customers using AI-powered search may experience incorrect or incomplete results, which we are working to improve, both for recall and precision.² We've opted for a model that has higher recall than precision.³ As a result, users will see all likely results of their investigation (inclusive of true positives and some false positives), and can determine the accuracy of their search results themselves. If we opted for a model with higher precision, some results identified by the system as false negatives would not be displayed to the user, which could be more costly than a false positive for an investigation. As designed, we put the customer in control to filter out false positive results for themselves. Users can also boost precision by selecting the “**Most Relevant**” button instead of the “**Most Recent**” button, as depicted below:



Two different sets of results from the same query “UPS truck yesterday.” In the first set of results with “Most Recent” toggled on, the model displays the most recent results from all cameras that could reasonably resemble a UPS truck (i.e., higher recall than precision). In the second set of results when the user selects “Most Relevant,” the model discounts more recent results for results with higher precision and produces fewer false positives.

2. In machine learning modeling, “precision” is the percentage of true positive results out of the entire batch of “positive” (inclusive of true and false positives) returned results. It’s expressed formulaically as $[Precision = \text{True Positive Results} / (\text{True Positive Results} + \text{False Positive Results})]$. It’s a measure of the model’s accuracy in identifying positive results. “Recall,” on the other hand, measures the percentage of true positives that are correctly identified. It is expressed formulaically as $[Recall = \text{True Positive Results} / (\text{True Positive Results} + \text{False Negative Results})]$. Whereas precision measures the overall rate of correct predictions, recall is helpful to identify the rate of false negatives. See [here](#) for a depiction of precision and recall.

3. Note that with Large Language Models (LLMs), recall and precision statistics will differ depending on the measured sample sizes. Our model’s average recall is greater than its average precision.



06

Query Moderation & Responsible Use

Unlike our filter-based attribute search, which has a predefined set of descriptors for narrowing search results, our AI-powered search lets customers search their footage for a wide range of attributes using their own words. Because our foundation model has been trained on publicly available data from the Internet, results may still be inaccurate, inappropriate or offensive. We have built query moderation into our platform to help reduce the risk that our AI-powered search is used maliciously or in ways that may be harmful. Moderation is a critical safeguard for this powerful feature.

At the same time, we implemented moderation in a way that isn't overly restrictive to the point that it compromises the feature's usability. We also leverage industry recognized practices, including open source data and OpenAI's moderation APIs, to make moderation more effective. Striking the right balance between respectful and useful searches is paramount for us.

Our approach to moderation

Just as we've chosen to use a publicly available model for our AI-powered search feature, we developed our query moderation using publicly available practices for effective moderation. OpenAI, for instance, has published guides and blogs for developing suitable moderation techniques.⁴ We further augmented these capabilities with additional proprietary protections, which include our own list of prohibited search categories:

- Information about the race, ethnicity or nationality of a person
- Information about the educational qualifications or religious beliefs of a person
- Subjective descriptions of people (e.g., attractive, ugly, wealthy)
- Inferred content from an image (e.g., "coolest person in the world")
- Sexual content or innuendo
- Names of specific people (i.e., public figures)
- Inappropriate or offensive descriptions of people or objects (e.g., "dumb person")
- Slurs equating people with animals (e.g., monkeys)

We also tapped into [open source data](#) to generate a list of banned text strings across all supported languages.

When words or phrases in a query appear on the list, we reject the search and mask any results—a process known as string-based matching. We also complemented this technique with [OpenAI's moderation API](#) to cross-check queries that create harmful or biased results.

We also hold frequent engineering “bug bashes,” analyze thousands of our own engineering test queries (i.e., classify thousands of queries as “appropriate” or “inappropriate”) to modify our moderation rules and optimize them for usability and to avoid bias. We can update our “blocklist” in minutes and thereby continuously improve our moderation techniques.

4. See blog: <https://openai.com/blog/using-gpt-4-for-content-moderation>. For more detailed, technical approaches to moderation that we leveraged, see: https://cdn.openai.com/papers/DALL_E_3_System_Card.pdf and https://cdn.openai.com/papers/GPTV_System_Card.pdf.



07

FAQ

1. Why did we develop AI-powered search?

Beyond common characteristics such as the presence of a backpack or the color of a car, many of our customers need the ability to search for people and vehicles across a much wider range of descriptive features. This capability is important for security and physical operations professionals, as it enables them to search through their footage quickly with a much more expansive set of descriptors than before. It also addresses operational issues across a variety of industries—from retailers looking to identify shoplifters or manufacturing customers looking to support workplace safety (e.g., "person not wearing a safety vest").

2. How did we develop the AI-powered search capabilities?

When developing this feature, we employed a two-pronged approach. First, we leveraged a publicly-available model with large language and large vision functionalities. This model enabled direct comparisons between text and images by training a multi-billion parameter neural network to bring related images and texts closer together while pushing unrelated ones apart. The foundation model also provided the basis for image classification and retrieval, allowing users to search for images using natural language.

Our approach further builds upon the foundation model in several important ways. We have, for instance, built a scalable, logically separated vector storage and retrieval system that processes customer video data before the search query is performed (in advance in a cache) to enable faster queries in real time, (e.g., no need to run the customer's video footage through the foundation model each time a query is run). This allows our AI-powered search to index and retrieve relevant footage for our customers quickly and at scale. See [here](#) for more information about our foundation model and how we've improved upon it.

3. What data do we send to OpenAI and what do we get back in return?

We send user search queries to OpenAI and receive back in return a "pass" or "fail" response. When using it for query parsing, we send the query to OpenAI and get it back parsed with time and camera/site information extracted from the original query which helps provides more accurate results.

4. Does Verkada use my queries to develop or improve its products or services?

We use queries only to provide the AI-powered search service, and will not use queries to develop or improve our model or other offerings without the customer's consent. If a customer submits feedback on specific results from AI-powered search, however, we may use that feedback to improve our products or services or to enhance the customer experience.

5. Can I use AI-powered search for all Verkada products? Is AI-powered search compatible with all Verkada cameras?

No. This feature applies only to video security cameras with people and vehicle analytics enabled. Please visit <https://www.verkada.com/docs> and look up the datasheet corresponding to your camera to find out if it's compatible with AI-powered search.

6. Can I use AI-powered search to find anything?

No. As described above, there are limitations as this search is tied to people and vehicles (and not standalone non-person, non-vehicle objects). Also, we're employing query moderation to reduce the risk that our AI-powered search is used maliciously or to promote harm.



7. Can AI-powered search be used to detect lethal weapons (e.g., guns, knives)?

No. While AI-powered search can identify certain objects that people hold or carry, our current models are not designed to identify lethal weapons. Verkada continues to explore integrations with third-party tools designed to detect weapons, but cannot currently support this functionality. Users should not rely on AI-powered search to detect guns, knives, or other lethal weapons.

8. Can AI-powered search be used to identify text, numbers, logos or brands on specific objects?

In order for AI-powered search to identify specific text, numbers, logos or brands on objects, two criteria **must** be met: i. the object needs to be tied to a person (e.g., a person is holding a bag with a clearly visible brand) or a vehicle (e.g., text on the side of a truck or bus) and ii. the text, number, logo or brand that you're trying to identify must be **sufficiently visible** for the camera to detect. Ideally, the text or logo should be large, clearly visible and the camera should have a high resolution. If these criteria are not met, then the results will be less accurate. **Note: this feature was not trained using Optical Character Recognition (OCR), so it cannot be used reliably to read license plates, barcodes, SKU names or other alphanumeric codes.**

9. Is there an additional charge to access AI-powered search?

No. All camera customers with Command licenses receive access to AI-powered search without additional charge.

10. What languages does AI-powered search currently support?

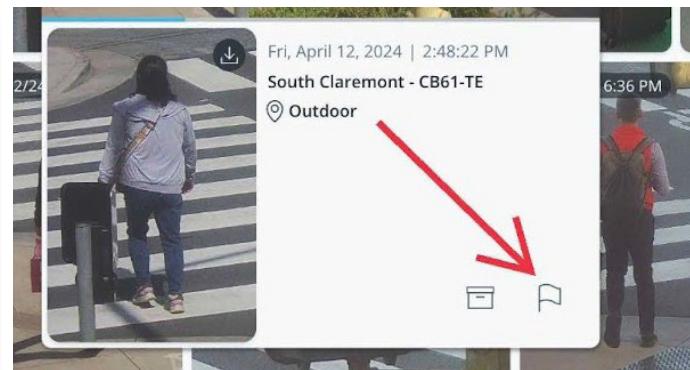
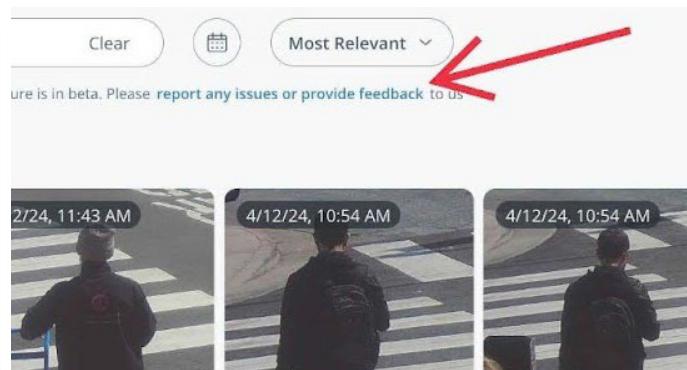
AI-powered search currently supports searches in English, Spanish, French, German, Japanese, and Korean.

11. Does AI-powered search work on the Command Mobile app?

No. In its current state, AI-powered search works on the Command web browser. Customers can access AI-powered search from their computers or mobile phone using a web browser.

12. What if I see inappropriate or inaccurate results?

Please report your issue(s) to us using either option depicted below. We appreciate your feedback in helping us to continually improve our model and offering:



The user can submit feedback on overall results and experience using AI-powered search via the button in the left-side image, or submit feedback on specific results using the flag button by hovering over a specific result.