PALLADYNE AI

From Remote Control to Collaborative Swarm Intelligence

Defining the True Levels of Swarm Autonomy for Drones

Denis Garagić, CTO and co-founder, Palladyne AI
Ben Wolff, CEO and co-founder, Palladyne AI

Executive Summary

In recent months, drone warfare has increasingly dominated defense headlines and congressional hearings. Reports from conflicts around the world regularly describe the growing use of unmanned aerial systems operating in large numbers, often accompanied by references to “autonomous swarms” or “AI-enabled drone swarms.” Senior officials within the U.S. Department of War have discussed the strategic importance of swarming technologies, while members of Congress have raised both interest and concern about the rapid emergence of autonomous capabilities in aerial systems. Media coverage frequently echoes these discussions, describing everything from coordinated drone attacks to experimental AI-enabled systems as examples of “autonomous swarms.”

At the same time, the terminology has spread rapidly across the defense technology industry. Drone manufacturers, defense startups, and major aerospace contractors increasingly describe their platforms using phrases such as “AI-enabled autonomy,” “collaborative swarming,” or “autonomous swarm capability.” In fact, more than two dozen drone manufacturers and defense technology companies have recently described their systems using terms such as “AI-enabled autonomy,” “swarm capability,” or “autonomous swarms” in press releases, investor briefings, and product demonstrations. In many cases, these claims refer to capabilities that range from simple formation flight, waypoint coordination or collision avoidance.
For policymakers, operators, and technology developers alike, this ambiguity creates a fundamental problem: different companies are often using the same words to describe very different capabilities. Systems that rely on pre-programmed automation, tightly scripted coordination, and fully autonomous collaborative intelligence are frequently described using identical terminology despite representing dramatically different levels of machine authority and decision-making.

The automotive industry confronted an almost identical problem more than a decade ago. As driver-assistance technologies improved, manufacturers began marketing vehicles as “self-driving.” Consumers encountered terms such as lane assist, adaptive cruise control, highway autopilot, and full autonomy with little clarity about what those distinctions meant in practice. Regulators, engineers, and customers needed a shared framework.
The Society of Automotive Engineers responded by formalizing the now widely recognized Levels 0 through 5 system. This framework did not slow innovation. Instead, it clarified responsibility. At its core was a simple organizing principle: as levels increase, responsibility for the driving task transitions from human to machine.

Understanding those levels in detail provides an essential foundation for understanding drone autonomy.

I. The SAE Automotive Autonomy Framework: Levels 0–5 Explained

Level 0 – No Automation
At Level 0, the human driver performs all aspects of the driving task. The vehicle may issue warnings, such as blind-spot alerts or forward-collision warnings, but it does not actively control steering, acceleration, or braking. Responsibility rests entirely with the driver.

Level 1 – Driver Assistance
At Level 1, the vehicle can assist with either steering or acceleration/braking, but not both simultaneously. For example, adaptive cruise control may adjust speed to maintain following distance, or lane-keeping assistance may help keep the vehicle centered. The human remains fully responsible and must continuously supervise.

Level 2 – Partial Automation
At Level 2, the vehicle can control both steering and acceleration/braking under certain conditions. Highway autopilot systems often fall into this category. However, the human must continuously monitor the environment and be prepared to intervene immediately. Responsibility still lies with the driver.

Level 3 – Conditional Automation
At Level 3, the vehicle can manage most aspects of driving under defined conditions. The system monitors the environment and handles routine driving tasks. However, when the system encounters conditions outside its operational domain, it requests that the human take over. Responsibility shifts temporarily to the machine, but fallback remains human-dependent.

Level 4 – High Automation
At Level 4, the vehicle can perform all driving tasks within a defined operational domain, such as a geofenced urban area or highway corridor, without requiring human intervention. If the system reaches the edge of its domain, it can bring itself to a safe state without relying on the driver.

Level 5 – Full Automation
At Level 5, the vehicle can operate under all conditions that a human driver could manage. No steering wheel or pedals are required. The human is no longer part of the control loop. Responsibility resides entirely with the system.

The power of the SAE framework lies in its clarity. Autonomy is not defined by how impressive a feature appears. It is defined by who is responsible for perception, decision-making, and execution.

This organizing principle translates directly to drones.

II. Remote Control: The Human as the Cognitive Core

At the most fundamental level, a drone is a remotely piloted aircraft. The human operator controls direction, speed, altitude, and heading. Every maneuver is deliberate. Every adjustment reflects a human decision.

In operational settings, the operator typically manages more than flight alone. Most drones carry gimbal-mounted cameras. The operator must therefore independently control pan, tilt, and zoom while flying the aircraft. This requires dividing attention between navigation and sensor management. Maintaining stable flight, avoiding obstacles, adjusting altitude relative to terrain, monitoring telemetry and battery reserves, interpreting live video feeds, and communicating with other personnel frequently occur at the same time.

The cognitive load can be substantial. Piloting demands spatial awareness and precision. Managing a camera requires visual focus and interpretive judgment. When both tasks are performed concurrently, particularly in complex environments, the human becomes the central processing unit of the system.

In the automotive analogy, this corresponds to Level 0. The machine executes but the human decides.

III. Automation in Drones: The Equivalent of Driver Assistance

As drone systems evolved, autopilot features were introduced to reduce workload and increase reliability. Stabilization compensates for wind. Altitude hold maintains consistent elevation. Return-to-home sequences provide safety if communications degrade. Waypoint navigation allows a drone to follow a predefined route without continuous joystick input.

These capabilities parallel automotive Levels 1 and 2. They provide assistance and partial automation, but the human remains responsible for oversight and intervention.

When a drone executes a waypoint mission, a human defines the route, altitude, speed, and often sensor orientation in advance. The drone follows the plan as programmed. If unexpected conditions arise, it can only respond within predefined rules. It does not reinterpret mission intent.

Object tracking operates in a similar way. The drone maintains relative distance to a selected subject and adjusts dynamically as that subject moves. The responsiveness is real, but the objective remains externally defined. The drone does not assume mission-level responsibility or independently determine whether the tracked subject remains the correct focus.

These features improve performance and reduce workload. However, they do not transfer responsibility from human to machine in the way higher autonomy levels require.

IV. Higher Levels of Drone Autonomy

Drone Level 0 – Manual Remote Operation
Equivalent to automotive Level 0. The human controls all aspects of flight and payload. This is 100% remote control with human-in-the-loop.

Drone Level 1 – Assisted Flight
Stabilization and altitude hold reduce workload. Responsibility remains with the operator, but cognitive load is somewhat reduced for the human operator.

Drone Level 2 – Automated Behaviors
Waypoint missions, return-to-home, and object tracking execute predefined routines. The drone reacts within programmed boundaries but does not interpret mission objectives.

Drone Level 3 – Conditional Autonomy
The drone can detect obstacles, adjust routes dynamically, and manage defined contingencies. If it encounters conditions outside its operational parameters, human supervision is required.

Drone Level 4 – High Autonomy
The drone executes missions independently within a defined operational envelope. It adapts to environmental changes and maintains safety without continuous human input. Humans are on-the-loop and are able to intervene or make ultimate go/no-go decisions as necessary or appropriate.

Drone Level 5 – Full Mission Autonomy Within Defined Domain
The system assumes responsibility for mission-level decisions within its domain. It adapts to unforeseen conditions and continues operating without requiring human fallback during execution. However, as with Level 4, humans remain on-the-loop and are able to intervene or make ultimate go/no-go decisions as necessary or appropriate.

As with automobiles, the defining factor is responsibility.

V. From Individual Autonomy to Collective Intelligence

When multiple autonomous drones operate together under decentralized principles, a second dimension of capability emerges. Swarming introduces structural advantages that no single drone can achieve alone.

Distributed sensing enables simultaneous observation from multiple vantage points. Dynamic role assignment allows responsibilities to shift fluidly as conditions evolve. Collective adaptation enables coordinated response to environmental or adversarial changes. Resilience under attrition allows missions to continue despite individual losses. Coordinated mission optimization reallocates sensing and positioning resources continuously to maximize effectiveness under uncertainty.

These properties mirror the decentralized, predictive principles described in Decentralized Embodied Collaborative Autonomy (DECA) . Intelligence resides at the edge. Perception is filtered locally. Decisions are made locally and collaboration emerges through decentralized interaction.

The behaviors described above begin to resemble a concept often used in military discussions of advanced swarming: the idea of a "wolf pack." The analogy comes from cooperative hunting behavior observed in nature. Wolves do not hunt as isolated individuals executing rigid roles. Instead, they operate as a coordinated unit in which sensing, pursuit, and containment are distributed across the pack. Some wolves push prey in a particular direction. Others flank. Others anticipate escape routes. Roles shift dynamically depending on terrain, prey behavior, and the moment-to-moment evolution of the hunt.

The importance of the wolf pack analogy lies in how clearly it illustrates the difference between synchronized automation and collaborative autonomy. A formation of drones flying predetermined routes may appear coordinated, but each aircraft is still executing a script. A wolf pack swarm behaves differently. Individual drones perceive locally, share information with neighboring agents, and adjust behavior continuously in response to the actions of other members of the swarm and the movement of targets within the environment.

In operational terms, wolf pack swarming represents a level of capability beyond simple multi-drone coordination. Drones distribute sensing responsibilities across multiple vantage points while maintaining collective awareness of the environment. If a target moves, some drones maintain persistent observation while others reposition to monitor escape routes or secondary threats. If one drone is lost, nearby units shift roles automatically to preserve coverage.

Wolf pack swarming already represents a major step beyond most current implementations of multi-drone operations. However, wolf pack behavior remains largely reactive. The swarm adapts quickly and intelligently to observed changes in the environment, but its behavior is primarily driven by immediate conditions.

A further level of capability emerges when predictive reasoning is incorporated into swarm decision-making. At this level, the swarm does not merely respond to what it sees, but it begins to anticipate what is likely to happen next. This combination of decentralized collaboration and predictive reasoning represents what we refer to as Oracle-Class Wolf Pack Swarming.

With Oracle-Class Wolf Pack Swarming, the collective intelligence of the system incorporates predictive models that estimate the likely future behavior of both friendly and adversarial actors. Instead of simply reacting to a target's movement, the swarm evaluates probable trajectories, potential objectives, and likely escape paths. Sensing resources and flight paths can then be allocated in anticipation of those possibilities.

VI. Operational Illustration

To understand how these concepts translate into operational capability, it is useful to examine how different levels of drone autonomy would perform in the same mission scenario.

Consider a reconnaissance mission supporting a ground convoy moving through a contested urban corridor. The convoy consists of several armored vehicles transporting personnel and supplies between two forward operating positions. The route passes through a dense section of city blocks where tall buildings create blind corners, narrow streets restrict maneuvering space, and multiple intersecting alleys provide potential ambush points.

A longer-endurance drone operating at higher altitude maintains wide-area surveillance above the district, monitoring overall traffic patterns and providing communications relay if needed. Lower-altitude drones operate closer to the convoy, positioned to observe intersections, rooftops, and adjacent streets.

At Drone Autonomy Levels 0–2, the mission resembles traditional remotely piloted operations. A drone operator must simultaneously fly the aircraft, manage the camera payload, and monitor the convoy's route. Even with assisted flight or automated behaviors such as waypoint navigation or object tracking, the system remains fundamentally human-directed.

In practice, this means each drone requires a dedicated operator, and sometimes additional personnel for payload management or mission coordination. If suspicious activity appears near an intersection—perhaps a parked vehicle suddenly pulling away from the curb or several individuals moving quickly between buildings—the operator must redirect the drone to investigate.

During that moment, the convoy temporarily loses visibility elsewhere along the route. The aircraft can only observe one location at a time, forcing the operator to constantly decide which potential threat deserves attention. If the drone turns to inspect a suspicious vehicle, it may lose visibility on a nearby rooftop or alleyway.

Scaling coverage under these conditions requires scaling personnel. Expanding aerial awareness from one drone to five drones often means assigning five operators and supporting crew, increasing both manpower demands and cognitive workload. If a drone is jammed or lost, aerial awareness may disappear entirely until another human-operated asset can be deployed.

At Drone Autonomy Level 3, the aircraft begins to demonstrate conditional autonomy. It can avoid obstacles, adjust routes dynamically, and respond to predefined contingencies. For example, if the convoy slows unexpectedly or traffic blocks an intersection, the drone may automatically reposition to maintain line-of-sight observation.

However, the drone still relies on human supervision when conditions exceed programmed parameters.

If several suspicious events occur simultaneously—such as a vehicle accelerating toward the convoy while individuals move across nearby rooftops—the operator must still prioritize which situation to investigate first. While workload may be reduced, human oversight remains necessary for each aircraft, limiting the degree to which personnel requirements can be reduced.

The dynamics change significantly at Drone Autonomy Level 4, where high autonomy allows multiple drones to coordinate with minimal direct control. In a wolf pack swarm configuration, drones distribute themselves across the environment rather than relying on a human operator to manually position each platform.

As the convoy moves through the corridor, one drone maintains forward reconnaissance above the lead vehicle. Another monitors rooftops along the route where observers or attack teams might position themselves. Additional drones spread out along adjacent streets to watch likely approach paths.
The higher-altitude drone continues to maintain a wide-area perspective, tracking vehicle movement across several blocks and alerting the swarm if unusual traffic patterns appear — such as multiple vehicles converging toward the convoy's route.

If suspicious movement appears near an intersection, one drone moves closer to confirm the observation while others automatically reposition to maintain coverage across the corridor. If buildings begin interfering with signal strength, a drone positioned behind the convoy may rise slightly to act as a communications relay between the convoy and command elements.

If one drone experiences mechanical failure or is lost to hostile action, neighboring drones expand their coverage zones to compensate. Coverage gaps close automatically without requiring an operator to manually reposition aircraft.

At this level, a single operator can supervise multiple aircraft simultaneously, shifting the role of the human from active pilot to mission supervisor. Instead of flying each drone individually, the operator manages mission objectives while the swarm handles navigation, positioning, and coordination.

Now consider the same scenario operating at Drone Autonomy Level 5, representing full mission autonomy within a defined operational domain — what we describe as an Oracle-Class wolf pack swarm.

Instead of merely reacting to suspicious activity, the swarm continuously evaluates patterns of movement across the environment. The higher-altitude drone observes vehicle traffic across the surrounding district while lower drones monitor street-level activity. Together they analyze pedestrian movement, vehicle flow, terrain features, and known choke points to estimate where an ambush is most likely to occur.

Before the convoy even approaches a potentially dangerous intersection, several drones have already repositioned. One observes rooftops overlooking the intersection. Another watches an alley that could conceal a vehicle or attack team. A third monitors a narrow side street that provides the most direct approach toward the convoy’s path.

Another drone quietly shifts position to watch a corridor that adversaries would most likely use to escape after an attack.

When a hostile vehicle suddenly accelerates from a side street toward the convoy, the swarm does not simply react. Two drones already positioned along the predicted approach corridor immediately confirm the threat and track the vehicle’s movement. The higher-altitude drone widens its observation area to detect whether additional vehicles are moving to support the attack.

At the same time, a drone that had been monitoring a potential escape route begins tracking a second vehicle attempting to leave the area. Another drone shifts to maintain persistent observation over the intersection as the convoy moves through.

In this configuration, the swarm is not merely observing events—it is continuously shaping the sensing network in anticipation of them. Sensors are positioned where activity is most likely to occur before events unfold.

The operational implications are significant. What previously required multiple operators controlling individual aircraft can now be accomplished by a small supervisory team overseeing an entire swarm. Human personnel remain on the loop, retaining authority to intervene or make ultimate go/no-go decisions, but the swarm assumes responsibility for positioning sensors, maintaining coverage, and adapting to changing conditions.

What Oracle-Class capability adds is foresight. Instead of repositioning drones only after a target changes course, the swarm can allocate coverage across likely future locations in advance. Instead of reacting to adversarial maneuvering, it can shape the operational environment by anticipating how those maneuvers will unfold. By combining predictive reasoning with mission-level autonomy, they enable persistent situational awareness while dramatically reducing the number of personnel required to sustain aerial reconnaissance in complex environments.

The difference between reacting to an adversary’s movement and predicting it may represent the most important qualitative leap in the evolution of autonomous swarming.

VII. From Theory to Real World Demonstration

The transition from remote control to automation to autonomy, and from individual autonomy to decentralized swarm intelligence, has progressed beyond theory. Cross-platform demonstrations by Palladyne AI have shown heterogeneous drones operating together with Oracle Class Wolf Pack Swarming capabilities under decentralized architectures consistent with DECA principles. These systems perceive, decide and coordinate locally without centralized real-time command. Palladyne AI's patented SwarmOS enables individual drones to operate with high levels of autonomy while simultaneously participating in decentralized swarm collaboration, allowing the swarm to demonstrate wolf pack behaviors — distributed sensing, dynamic role assignment, collective adaptation, and resilience under attrition — while also incorporating predictive reasoning that allows the swarm to anticipate how the operational environment will evolve. In other words, SwarmOS is not merely enabling drones to coordinate; it is enabling drones to coordinate and predict as a cohesive unit consistent with Oracle-Class Wolf Pack Swarming capabilities.

Conclusion

The SAE automotive framework clarified autonomy by focusing on responsibility. Applying the same principle to drones reveals clear distinctions between remote control, automation, conditional autonomy, and full mission autonomy. Extending autonomy into decentralized collaborative swarming introduces distributed sensing, dynamic role assignment, collective adaptation, resilience under attrition, and coordinated optimization.

Adding predictive capabilities to deliver Oracle-Class Wolf Pack Swarming represents the next evolutionary step in autonomous aerial capability. Oracle-Class Wolf Pack Swarming is not a distant research objective, it is an operational capability that is beginning to redefine how autonomous aerial systems collaborate, sense, and make decisions in complex environments.

This isn't science fiction. We don't have to wait until 2032 to deliver this capability. It exists today.