

WHITE PAPER

AV1 codec in video surveillance

Future-proof video compression in modern security systems

June 2026

Summary

The surveillance industry is transitioning from the long-standing H.264 video compression standard to AV1. This newer codec can better meet the increasing bandwidth and storage demands of modern high-resolution video. Unlike H.265, which was slowed by complex licensing, AV1 provides a royalty-free¹, highly efficient alternative supported by major technology leaders.

Axis has integrated AV1 support in cameras based on the system-on-chip ARTPEC-9. The implementation focuses on features that are relevant to surveillance use cases, rather than broadcast-centric functions. It provides compression efficiency comparable with H.265 without increased hardware costs or power consumption. AV1 is compatible with existing Axis technologies, such as Zipstream, while also enabling new features including togglable overlays and the AVIF format for still images.

AV1 builds on the same core principles of block-based compression, motion estimation, and transform coding used in H.264 and H.265, while introducing more flexible prediction mechanisms. It replaces rigid P- and B-frame structures with versatile inter-frames and more efficient key frames. Image partitioning is more precise, utilizing superblocks that can be split into various shapes and aspect ratios.

For surveillance specifically, global and warped motion tools allow the encoder to handle camera movements with minimal data. Switch frames enable adaptive bitrate streaming by allowing resolution changes without requiring a new independent key frame. Technical quality is maintained through higher internal bit depth, diverse transform types, and advanced in-loop filters. These improvements are supported by a non-binary arithmetic coder designed for efficient hardware implementation and faster processing.

¹ *Users should ensure they understand the licensing terms and conditions associated with AV1 video technology.*

Table of Contents

1	Introduction	4
2	The need for a new compression standard	4
3	Adopting AV1	4
4	Benefits of AV1 in video surveillance	5
5	Axis implementation of AV1	5
5.1	Zipstream and AV1: better together	5
5.2	New features with AV1	6
5.2.1	Togglable overlays	6
5.2.2	AVIF (AV1 Image Format)	6
6	Bitrate efficiencies of AV1, H.264, and H.265	6
6.1	Scene examples	6
6.2	Comparing codecs: forensic video quality versus bitrate	8
7	AV1 technical foundation	9
7.1	Refined frame types and prediction	9
7.1.1	The established I-P-B framework in H.264 and H.265	9
7.1.2	AV1's flexible approach to inter-prediction	9
7.2	Advanced coding toolsets	10
7.3	Operational efficiency features	10
7.4	Processing and filtering	10
8	Technical integration guide	11
8.1	Streaming and decoding AV1 video in Axis devices	11
8.2	Key considerations for developers	11
8.2.1	Retrieving AV1 streams	11
8.2.1.1	Standard protocols	11
8.2.1.2	WebRTC	11
8.2.2	Decoding and rendering AV1 streams	12
9	Trademark attributions	12

1 Introduction

For over two decades, H.264 (AVC) has been the dominant video compression standard in surveillance and similar applications, offering a reliable balance between bitrate and image quality. However, modern security systems are facing unprecedented data demands. Higher resolutions, larger camera counts, and the rise of cloud-based analytics are pushing network bandwidth and storage capacities to their limits. The industry has scaled, requiring compression technology to evolve with it.

While H.265 was once considered the natural successor, complex licensing has slowed its adoption. In contrast, the newer standard AV1 delivers superior compression efficiency combined with broad IT support and ecosystem maturity. It's an open, royalty-free codec backed by the world's leading technology companies and designed for the demands of modern video.

Many years ago, Axis led the industry into the H.264 era, by being the first security camera manufacturer to adopt H.264 in the products. Now, Axis is introducing AV1 to the industry.

This white paper examines AV1 in video surveillance, from compression efficiency to ecosystem support and practical integration. It includes bitrate comparisons, a technical deep dive into how AV1 works, and a developer's guide for integrating Axis devices.

2 The need for a new compression standard

Without video compression, streaming and storing video would be almost impossible. A standard 1080p video at 30 frames per second would consume almost 1 Gbit/s without any compression. With such a high bitrate, you would fill up a 1 TB disk in under three hours of recording, and a typical network couldn't transport the video from a single camera. With an 8K camera, which delivers extreme image quality, you would fill the same disk in about 30 minutes.

Video encoders compress video data dramatically, making it practical to stream, store, and transmit surveillance video at scale.

While H.265 offered improvements over H.264, allowing higher camera resolutions and reduced bitrate, the sheer number of deployed cameras continues to strain network bandwidth and storage infrastructure. As the industry moves beyond 4K resolution, incrementally better compression is no longer sufficient.

Furthermore, modern video encoders are built on compression methods developed by many different companies and researchers over the years. A significant challenge with both H.264 and H.265 has therefore been their fragmented and often unpredictable patent licensing models. H.265 faced multiple patent pools and individual patent holders, creating an opaque and costly ecosystem for manufacturers and users. This complexity has hindered adoption, limited compatibility, and added unforeseen costs during deployment. It has even led some major hardware manufacturers to begin removing H.265 hardware support from existing products.

The security industry requires a codec that not only delivers superior compression but also operates on a transparent, accessible, and future-proof foundation.

3 Adopting AV1

The AV1 video compression standard was developed by the Alliance for Open Media (AOMedia), a consortium founded with the goal of creating an open, royalty-free video format. Today, AV1 is the preferred codec for major streaming service providers like Netflix and YouTube. It's supported by all major desktop OS (Windows®, macOS® and Linux®) and mobile platforms (Android™ and iOS®).

As an open standard backed by companies such as Amazon, Apple, Cisco, Google, Intel, Meta, Microsoft, Mozilla, Netflix, Nvidia, and Samsung through AOMedia, AV1 has strong long-term prospects. Its open nature also means fewer licensing uncertainties and lower patent costs than other codecs.

VMS adoption is underway, with several industry-leading providers, including Milestone, Genetec, and Axis, already supporting it. The web client of many VMS solutions, including AXIS Camera Station Pro, benefits from

native AV1 support – something H.265 has struggled to achieve even after a decade on the market. Adding AV1 support is relatively straightforward thanks to mature open-source libraries and tools.

Axis is actively engaged with industry bodies like ONVIF® to drive interoperability around AV1. Axis views the current ONVIF video streaming method as already aligned with AV1, anticipating that broader industry adoption will pave the way for its formal addition to an ONVIF standard.

On the hardware side, modern CPUs now widely include integrated AV1 decoding. While some hardware manufactured between the years 2015 and 2020 supports H.265 but not AV1, the rapid expansion of AV1 hardware support in modern devices means this gap is closing quickly. For example, AXIS Camera Station S2216 Mk II Rack Appliance is a recording solution that offers hardware AV1 decode support for 16 channels without a dedicated GPU. But even typical laptop CPUs can handle 8K30 AV1 streams. Additionally, an increasing number of smartphones and tablets include AV1 hardware decoders, eliminating concerns about battery drain or latency during mobile playback.

4 Benefits of AV1 in video surveillance

While AV1 wasn't initially designed for surveillance, H.264 and H.265 weren't either. The underlying compression method of all the standards is, however, highly adaptable, and the AV1 implementation in Axis ARTPEC-9 system-on-chip is optimized for surveillance.

For a given amount of video data, the lower bitrates you can get with AV1 directly lead to reduced storage requirements, potentially extending the lifespan of existing storage infrastructure or allowing for longer retention periods. Furthermore, decreased network bandwidth consumption can defer or eliminate costly network upgrades. Overall this gives reduced total cost of ownership (TCO) for a surveillance system.

Also, cloud-based video surveillance is growing rapidly and this places significant demands on compression technology. Transmitting video from hundreds or thousands of cameras to cloud storage and analytics platforms requires a codec that can deliver high quality at low bitrates over internet connections. AV1 was designed for internet video transmission and its royalty-free licensing model means it's natively supported across browsers, platforms, and cloud services. This broad support simplifies cloud integration and removes the compatibility barriers that have historically made H.265 difficult to deploy at scale in cloud environments.

5 Axis implementation of AV1

Axis leads the security industry as the first manufacturer to integrate AV1 support in cameras. This reflects a commitment to both advancing the video surveillance industry and improve customer utility.

The AV1 implementation focuses on the codec's fundamental strengths for security applications, engineered for the specific demands of 24/7 continuous recording, AI analytics integration, and surveillance system interoperability. Broadcast-centric features that are not relevant to surveillance are deliberately excluded. Instead, the core compression capabilities are used, which are based on the codec's flexible block partitioning, advanced intra- and inter-prediction, and sophisticated filtering.

AV1 support isn't a cost driver in Axis cameras since the camera-integrated ARTPEC architecture allows for efficient implementation and integration of new codecs. The customers get the full efficiency benefits of AV1 without additional expense. ARTPEC-9 delivers full real-time AV1 encoding performance with minimal overhead, without additional CPU or power load on the camera.

All Axis features available with H.264 and H.265 are supported from launch, including Zipstream, signed video, bitrate control algorithms, and Zipstream profiles. AV1 is also fully supported in AXIS Camera Station and AXIS Site Designer.

5.1 Zipstream and AV1: better together

AV1, H.264, and H.265 were all developed for broad use across many industries. Because the standards focus on the decoder rather than the encoder, incremental improvements to the encoder can be developed independently. This is what allowed Axis to develop Zipstream, a surveillance-specific optimization designed to work together

with the underlying video codec. It delivers at least 50% savings on bandwidth and storage while protecting the forensic value of the video.

Zipstream analyzes video scenes in real time, dynamically identifying areas of interest such as human faces and vehicle license plates, and preserving their detail while compressing less critical areas more aggressively. Combined with AV1's inherently superior compression of complex moving objects, high-quality forensic details are maintained where they matter most, while overall bitrate is significantly reduced.

5.2 New features with AV1

Already in its first ARTPEC implementation, AV1 consistently demonstrates a compression efficiency comparable to the third-generation implementation of H.265. Also, AV1 offers a more advanced toolset that unlocks completely new capabilities. Features that soon will be available in Axis cameras include toggable overlays and AVIF format for still images.

5.2.1 Toggable overlays

Toggable overlays is a feature that lets users switch graphical elements in the video on or off during both live and recorded playback. These elements include bounding boxes, text annotations, and MQTT data visualizations. The overlays are permanently embedded in a dedicated layer within the video, so they are always there even when not visible. This means that operators can choose to view clean video when they don't need the overlay information, without losing any of the embedded data. Toggable overlays demonstrate how AV1's more advanced architecture enables entirely new capabilities that go beyond what H.264 and H.265 can support.

5.2.2 AVIF (AV1 Image Format)

AVIF is a modern still image format based on the same compression technology as AV1 video. Just as AV1 improves on H.264 and H.265 for video, AVIF offers significantly better compression efficiency than older image formats such as JPEG, while maintaining high image quality. For video surveillance, this means that still images captured from cameras, such as snapshots used in alerts or evidence, can be stored and transmitted more efficiently. AVIF also supports wide color gamut and high dynamic range, preserving more detail in challenging lighting conditions that are common in surveillance.

6 Bitrate efficiencies of AV1, H.264, and H.265

The built in video encoder in ARTPEC-9 can compress video using MJPEG, H.264, H.265, or AV1.

While MJPEG is extremely bitrate inefficient, it's a format that is easy to use and, after 30 years, it still has some applications. Since the JPEG algorithm produces very high bitrates, the camera's embedded software puts a limit on the total MJPEG throughput.

H.264, H.265, and AV1 have much lower base bitrates than MJPEG, and ARTPEC-9 can deliver streams with all these three encoders simultaneously. You can mix resolution, frame rate, and encoding settings as your application needs, up to the device's performance limit. In its highest-performance memory configuration, ARTPEC-9 delivers a 4K stream at 60 fps, reaching a total throughput of 540 MP/s. Read more in the white paper *Streaming performance for ARTPEC-9 products*

The bitrate from a new scene is hard to predict because video codecs are complex, and bitrate efficiency varies significantly from scene to scene. More advanced codecs tend to show greater variation. For help with estimating bitrate and storage needs for complete installations with H.264, H.265, and AV1, you can use the design tool *AXIS Site Designer*, axis.com/support/tools/axis-site-designer

6.1 Scene examples

Axis tests with simultaneous streams in the same camera, filming selected surveillance scenes, show that bitrate with AV1 is typically similar to H.265 bitrate, while H.264 bitrate is significantly higher.

Table 6.1 *Bitrates and bitrate comparison between the codecs.*


Scene	H.264 (Mbit/s)	H.265 (Mbit/s)	AV1 (Mbit/s)	AV1 vs H.264	AV1 vs H.265
Busy city surveillance	23.558	16.842	15.862	-33%	-5.8%
Traffic monitoring	20.525	15.773	15.660	-24%	-0.72%
Daytime city surveillance	16.211	13.249	13.244	-18%	-0.04%
Nighttime city surveillance	4.7725	3.5333	3.3637	-30%	-4.8%




*Busy city surveillance scene, from a standard test sequence video.
Bitrate with AV1: -33% compared with H.264, -5.8% compared with H.265.*



*Traffic monitoring scene, filmed with AXIS Q1728 Block Camera.
Bitrate with AV1: -24% compared with H.264, -0.72% compared with H.265.*



*Daytime city surveillance scene, filmed with AXIS Q1728 Block Camera.
Bitrate with AV1: -18% compared with H.264, -0.04% compared with H.265.*



*Nighttime city surveillance scene, filmed with AXIS Q1728 Block Camera.
Bitrate with AV1: -30% compared with H.264, -4.8% compared with H.265.*

6.2 Comparing codecs: forensic video quality versus bitrate

Comparing bitrates is meaningless if the video quality differs. Unfortunately, there's no reliable way to measure image quality automatically, and manual evaluation is very time consuming.

If you use two cameras you would have to either:

- tune both cameras to deliver the same bitrate and then compare video quality by stepping through frames and looking for artifacts
- or tune both cameras to deliver the same image quality (time consuming and difficult to verify), then compare the bitrate.

However, two cameras side by side will only have approximately the same field of view. Bitrate uncertainty can be up to 20% in such test setups, which is far too much for a meaningful comparison. You would need to film a wide variety of scenes to compensate for the fact that you're filming only almost-identical content.

To fairly compare two codecs, you need to encode the exact same video with both. The video must be frame-by-frame identical, which means all camera settings, field of view, sharpness, and scene content must match. The only way to achieve this from a live view is to encode both streams simultaneously in the same camera. Recent ARTPEC chips have enough processing power to handle this. They can encode two streams in parallel and save both to an SD card or stream them to a VMS. The recommended approach is to use one camera with

preconfigured stream profiles and triggers, and this is how the bitrate efficiency comparison presented here was conducted.

7 AV1 technical foundation

While the transition to AV1 introduces new efficiencies, it is best understood as a refinement of the existing standards. AV1 builds on the same core principles of block-based compression, motion estimation, and transform coding used in H.264 and H.265. AV1 also introduces a suite of advanced tools and prediction mechanisms.

7.1 Refined frame types and prediction

The main differences between the codecs lie in the degree of flexibility the encoder has when predicting pixel movement.

7.1.1 The established I-P-B framework in H.264 and H.265

H.264 and H.265 rely on a standard hierarchy of three primary types of compressed frames to achieve interframe compression. A sequence of these frames forms a group of pictures (GOP), typically starting with an I-frame, followed by a mix of P- and B-frames. The encoder can freely select frame type after the I-frame.

- **I-frames**, or intra-coded frames, are independent anchor frames that store a complete image without referencing other frames. They typically provide random access points for decoding, seeking, and error recovery. Using only spatial prediction from neighboring pixels within the same frame, I-frames are the largest frame type. H.265 enhanced intra-prediction by introducing more directional modes and larger prediction block sizes than the H.264 standard.
- **P-frames** are predictive frames, encoded by predicting differences from one or two previous I- or P-frame(s). They are significantly smaller than I-frames because they only store motion vectors (indicating how blocks of pixels have moved) and residual data (the differences between the prediction and the actual image). While P-frames efficiently represent movement between frames, their reliance on a past reference can limit prediction accuracy during complex motion or substantial scene changes.
- **B-frames** are bi-predictive frames that reference both past and future data to maximize coding efficiency. They achieve the highest compression ratios by leveraging information from two directions and utilizing a hierarchy that prevents artifacts from spreading to other frames. While ideal for smooth motion, B-frames introduce decoding delay due to their dependency on future frames. H.265 enhanced B-frame performance through flexible reference lists and adaptive motion vector prediction.

7.1.2 AV1's flexible approach to inter-prediction

AV1 maintains the concept of referencing other frames to reduce data, while replacing rigid P- or B-frame labels with a versatile inter-frame structure.

- **Key frames** serve the same fundamental role as I-frames by providing self-contained pictures that can be decoded independently. They remain essential for playback initiation, seeking, and error resilience. Enhanced intra-prediction ensures that even the "full" key frames in AV1 are encoded more efficiently than I-frames in older codecs, reducing their size while maintaining quality. Specific intra-prediction tools include more directional modes, Paeth and palette predictors, and Chroma from Luma.
 - More directional modes means a wider range of 56 angles for predicting pixel values from neighbors, compared to 35 in H.265 and nine in H.264.
 - The Paeth predictor selects directions based on gradients while the palette predictor efficiently manages content with limited color sets such as computer graphics.
 - Chroma from Luma (CfL) is an intra-prediction technique that predicts color information (chroma) from brightness data (luma).
- **Inter-frames** represent motion and changes between key frames. They are fundamentally similar to P-frames but exceed the capabilities of both P- and B-frames by utilizing a multi-directional reference structure. Unlike P-frames with one reference and B-frames with up to two, AV1 inter-frames can reference

up to six previously decoded frames for any given block. This allows the encoder to find the best match across multiple temporal distances for more accurate prediction and smaller residuals.

7.2 Advanced coding toolsets

AV1 introduces several sophisticated tools to refine how the blocks are processed.

Advanced block partitioning. While AV1 supports 128x128 superblocks, the main strength of the codec lies in the flexibility of its block partitioning. Superblocks of 128x128 or 64x64 pixels can be divided into highly flexible patterns including T-shaped splits and 4:1 or 1:4 aspect ratios, down to 4x4 pixels. This allows the encoder to partition images more precisely around fine details than the rigid structures in H.264 or H.265.

Compound prediction. This feature enables the blending of two different predictions for a single block. Inter-inter compound combines different reference frames or motion vectors, while inter-intra compound blends inter and intra-predictions to improve object boundaries. Tools like wedge partitioned prediction creates smooth oblique transition lines between blended predictions to avoid blocky artifacts and accurately separate moving objects.

7.3 Operational efficiency features

Some AV1 tools are designed specifically for camera movement and streaming performance.

Global motion and warped motion. AV1 accounts for camera movements like pan, tilt, or zoom by applying a single set of warping parameters across large regions or entire frames. This eliminates the need to encode individual motion vectors for every block during camera movement and is highly efficient for surveillance footage.

Switch frames (S-frames). These are specialized inter-frames designed for adaptive bitrate streaming. They can be predicted from a higher-resolution reference frame. This allows a decoder to switch to a lower resolution stream without needing a full key frame, optimizing for varying network conditions.

7.4 Processing and filtering

AV1 introduces several mathematical and post-processing improvements that enhance final image quality and encoding efficiency.

Higher internal precision. AV1 performs internal processing at 10 or 12 bits per sample to reduce rounding errors and preserve detail. This is especially beneficial in high dynamic range scenarios.

Transforms. AV1 supports a significantly broader range of transforms than H.265. It utilizes various types to map residual patterns into the frequency domain, including rectangular DCT (discrete cosine transform) and asymmetric DST (discrete sine transform). The codec can combine different 1D transforms horizontally and vertically to better adapt to the characteristics of residual data.

In-loop filtering. AV1 utilizes the constrained directional enhancement filter (CDEF) to smooth ringing artifacts along dominant edges. It also employs a Wiener-based loop restoration filter to reduce blur artifacts throughout the frame.

Film grain synthesis. Rather than inefficiently encoding random noise, AV1 transmits parameters to synthesize grain at the decoder. This preserves perceived texture while saving significant bitrate. This feature is designed for film and broadcast use rather than surveillance.

Non-binary arithmetic coding. AV1 utilizes a non-binary arithmetic coder that can be more efficient and faster to implement in hardware than the binary CABAC systems used in H.264 or H.265.

8 Technical integration guide

8.1 Streaming and decoding AV1 video in Axis devices

Axis supports the AV1 video encoding standard on devices equipped with the ARTPEC-9 system-on-chip (SoC) or later. The implementation is fully compliant with the AV1 specifications, supporting the Main Profile with 8-bit color depth for compatibility across a wide range of platforms.

There is support for streaming AV1 video over RTSP/RTP in the Axis streaming library. This implementation is written in C# and uses VAPIX to communicate with Axis devices.

To ensure high performance playback, specific decoding strategies are used based on the viewing platform. For desktop applications, FFmpeg is used with the dav1d decoder library enabled. Web clients rely on the built-in AV1 video support in modern browsers.

For mobile applications built for iOS and Android, Axis AV1 implementation balances flexibility and performance. FFmpeg is used to retrieve the video frames from the stream. The mobile platform itself performs the actual decoding. Native video players handle the rendering of frames to ensure smooth playback and efficient power usage.

8.2 Key considerations for developers

VAPIX® APIs provide the necessary tools and documentation for customized integration with AXIS OS. For guidance on how to integrate AV1 encoded video from our devices into your integrations and solutions, see Axis developer documentation, especially: <https://developer.axis.com/video-streaming-and-recording/av1/how-to-guides/integration-guide-av1/>

Before integrating with AV1-enabled Axis devices, make sure that your system's hardware and software can support AV1 video decoding. AV1 video can require more CPU resources than H.265 if hardware support is missing or the decoder library isn't optimized. Consider enhancing your application's performance by employing multi-threading, GPU acceleration, or other techniques.

There are licensing terms and conditions associated with using AV1 video technology. It is your responsibility to obtain any necessary licenses from third parties for any patents or other intellectual property that may be required for the use of AV1 video technology.

8.2.1 Retrieving AV1 streams

8.2.1.1 Standard protocols

Developers can request AV1 streams using standard protocols such as RTSP/RTP or HTTP, by accessing a specific URL that accepts streaming parameters, such as codec, resolution, frame rate, and bitrate.

8.2.1.2 WebRTC

You can also use WebRTC to retrieve AV1 streams. For this, you must use a WebRTC client application to view the video stream, at least one signaling server to set up the initial connection, and a STUN or TURN server for NAT traversal.

For the WebRTC client application, the most common and flexible option is to use a modern web browser and develop a web page using JavaScript to interact with the WebRTC APIs (such as `RTCPeerConnection` and `getUserMedia`). For desktop or mobile, you could instead develop a custom application using WebRTC libraries (such as `LibWebRTC`) to integrate the video stream.

If you're not using an Axis Cloud Connect solution, you'll need to host your own signaling server. It will act as a coordinator, allowing the Axis camera (as one WebRTC endpoint) and your client (the other endpoint) to exchange necessary information to establish a direct connection. This includes:

- SDP (Session Description Protocol) offers/answers. These describe the media capabilities and configurations of each peer.
- ICE (Interactive Connectivity Establishment) candidates. These are potential network addresses (IP and port combinations) through which each peer might be reachable.

You can build a signaling server using various backend technologies (such as Node.js with WebSockets, Python, Java, Go) to handle these exchanges.

Most devices, including cameras and client computers, are behind NAT (Network Address Translator) routers and firewalls. This makes direct peer-to-peer connections difficult because their internal IP addresses aren't publicly routable. A STUN (Session Traversal Utilities for NAT) server helps peers discover their public IP addresses and port mappings. It's relatively lightweight. You can use public STUN servers or host your own.

If STUN isn't enough to establish a direct connection (for example due to restrictive symmetric NATs or corporate firewalls), a TURN (Traversal Using Relays around NAT) server is needed. TURN acts as a relay, forwarding media traffic between peers. This consumes more bandwidth and introduces a slight delay, but it guarantees connectivity. You would need to host your own TURN server.

8.2.2 Decoding and rendering AV1 streams

To decode AV1 video streams in your application you can use FFmpeg, a popular open-source multimedia framework. To render AV1 video streams you can use most modern browsers, media players, video hardware, or public frameworks.

9 Trademark attributions

© 2026 Axis Communications AB. AXIS COMMUNICATIONS, AXIS, ARTPEC and VAPIX are registered trademarks of Axis AB in various jurisdictions. Microsoft and Windows are registered trademarks of the Microsoft group of companies. MacOS and Apple are trademarks of Apple Inc., registered in the U.S. and other countries. IOS is a trademark or registered trademark of Cisco in the U.S. and other countries and is used under license. Android and Google are trademarks of Google LLC. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. ONVIF is a trademark of ONVIF, Inc. All other trademarks are the property of their respective owners.

About Axis Communications

Axis enables a smarter and safer world by improving security, safety, operational efficiency, and business intelligence. As a network technology company and industry leader, Axis offers video surveillance, access control, intercoms, and audio solutions. These are enhanced by intelligent analytics applications and supported by high-quality training.

Axis has around 5,000 dedicated employees in over 50 countries and collaborates with technology and system integration partners worldwide to deliver customer solutions. Axis was founded in 1984, and the headquarters are in Lund, Sweden.