Video Surveillance System Anatomy

The seven-stage pipeline — the typical technology and the typical failure mode at every hop.

Stage	Typical technology	Typical failure mode
1 · Camera **capture + compress**	IP camera; H.264 / H.265 encoder; sensor + lens chosen by task (DORI)	Wrong camera for the task; bitrate set needlessly high
2 · Network **carry stream + power**	PoE switch (IEEE 802.3af/at/bt); structured cabling; dedicated VLAN	PoE budget or uplink under-provisioned; cameras brown out
3 · Ingest **find + pull stream**	ONVIF discovery; RTSP control (RFC 7826); RTP transport (RFC 3550)	'ONVIF just works' — only the baseline streams; rest needs SDK
4 · Storage **retain footage**	Recording server; surveillance-grade drives; RAID; recording mode	Sized for the demo, not the retention policy; fills early
5 · Analytics **video into events**	Edge / edge-server / cloud computer vision; ONVIF Profile M metadata	Tiers blurred into 'AI'; false-alarm flood; unbudgeted cloud bill
6 · Client **view + respond**	VMS desktop / browser / mobile; role-based access control (RBAC)	Too many feeds per operator; flat access fails a privacy audit
7 · Integration **events into action**	APIs and SDKs; access control; alarms; notifications	Alert fatigue; island deployment reviewed only after the fact

Wrapping layer — Privacy & retention: GDPR / EDPB Guidelines 3/2019 / Illinois BIPA; encryption in transit and at rest; retention limits. Biometrics (face, license plate) are a legal gate before they are a feature.

Storage math: storage = bitrate × cameras × hours × days. Worked example: 40 cameras × 4 Mbps × 24 h × 30 days ≈ 51.6 TB in H.264; ≈ 26 TB in H.265; less again on motion or event recording.

Engineering guidance, not legal advice. Confirm specifics with qualified counsel.