

VMS Capacity Planning Worksheet

Size a surveillance system end to end: the eight-step method. Planning estimates - verify with a pilot.

The unit of demand is one camera

Capacity planning sizes FIVE resources - recording throughput, storage, network/power, viewing (client decode), and the management plane - and each hits its limit at a different camera count, so the resource you forget to size is the one that breaks. The camera is the unit of demand: list the cameras, multiply each camera's load by the count, sum, then add headroom. Worked example (40 cameras, 1080p, 2 Mbps, continuous, 30 days): ~26 TB video -> ~39 TB provisioned; 10 MB/s logical writes -> ~60 MB/s raw on RAID 6; 80 Mbps aggregate; ~400 W PoE.

The eight-step method (fill in the right-hand column for your system)

Step	Size this	Formula / rule	Your figure
1. Count cameras	Camera fleet	List cameras; fix resolution, bitrate, fps, recording mode	
2. Per-camera load	Load per resource	bitrate; GB/day = Mbps x 10.8; 1 switch port; decode share - then x count	
3. Recording servers	Ingest throughput	Sum bitrate as MB/s; size by SUSTAINED WRITE, not count (~3.1 Gbit/s ~ 700 cams/server)	
4a. Storage capacity	Retention TB	Mbps x 10.8 GB/day x cams x days x rec-factor, then x 1.4-1.6 provisioning	
4b. Storage throughput	Sustained write	logical MB/s x RAID penalty (RAID 6 ~ x6); size disks (CMR) for it	
5. Network + PoE	Switching + power	aggregate Mbps; 10G uplinks >16 cams; PoE budget +20-30%; own VLAN	
6. Viewing	Client decode	total decode Mbps; substreams for walls; hardware decode (~5x 1080p/GPU)	
7-8. Headroom + validate	Margin + redundancy	run at 70-80%; RAID6+spare, N+1, mgmt failover, UPS; pilot, then re-baseline	

Size by throughput, not camera count

The #1 error is sizing a recording server or a storage array by a round camera number. Recording servers are bound by sustained WRITE THROUGHPUT (MB/s), not a count - one server can record ~700 cameras at 1080p / 4.4 Mbps (~3.1 Gbit/s). Storage has TWO budgets: capacity (retention TB) AND throughput (write MB/s x the RAID penalty); a capacity-only array holds the video but drops frames under load. And the client/viewing tier is often the weakest link - a PC + GPU decodes only ~5 1080p streams (or one 4K), so use substreams + hardware decode.

Headroom & redundancy checklist

Run recording, storage, network, and clients at 70-80%, not 100%. Disk: RAID 6 (survives two failures) + a hot spare; size for the rebuild window. Recording: an N+1 failover recording server per pool. Management plane: management-server failover (clustered, e.g. WSFC). Power: UPS on cameras, switches, and servers. Network: redundant uplinks / dual NICs; cameras on their own VLAN. Validate with a pilot; re-baseline as cameras, bitrates, and retention change.

Engineering guidance and planning estimates - real installations differ (variable bitrate, scene activity). Vendor sizing tools (AXIS Site Designer, Milestone solution designer) automate steps 1-6 and are a useful cross-check; their output is an estimate. Designed to IEC 62676-4 (system-level application guidelines).