



Your Video Infrastructure Checklist

Treating Video as a Strategic Asset

Video is increasingly seen as mission critical, whether it's being used for surveillance and security or for any number of other applications such as traffic analysis, facility utilization, energy management, production optimization, or telehealth. Unfortunately, many companies and organizations rely on outdated infrastructure that can't keep pace with the requirements of today's mission-critical video. It's not just about achieving faster throughput: it's about improving the ability to capture, protect and leverage video data as a strategic asset while reducing total cost of ownership (TCO), risks and liabilities. The checklist below can help you in assessing if the infrastructure in place for video surveillance and security operations is mature and capable of supporting mission-critical video.

Here are the hallmarks of a mature video infrastructure. Check all that apply:

RESILIENCE

- Infrastructure is optimized for write-intensive video workloads to eliminate the risk of dropped frames and image degradation.
- High availability and fault tolerance can be achieved without requiring redundant hardware and software or impacting storage efficiency by replicating data.
- Video and other data remains available and accessible even if multiple disks and/or an entire appliance are unavailable.
- The servers hosting your video recording and storage do not represent a single point of failure.
- Infrastructure includes intelligent monitoring and analytics for proactive system health monitoring and automated notification.

COST OF OWNERSHIP

- Large amounts of video can be stored in a compact and efficient footprint to reduce power, cooling, space, and management costs.
- Infrastructure resources can be shared or pooled to optimize performance and resource utilization and don't result in islands of stranded, underutilized resources.
- Multiple security applications (video management, access control, video analytics, intrusion detection, etc.) can be consolidated on common infrastructure to reduce costs and simplify management.

SCALABILITY/MANAGEABILITY

- Server and storage resources can be deployed on a pay as you grow basis to minimize unnecessary overprovisioning and up front expense.
- Storage, compute and bandwidth can be scaled independently or together as camera count, or retention times, or other requirements change.
- The infrastructure, even multiple systems or sites, can be managed and updated from a single pane of glass by non-technical administrators.
- Resources scale non-disruptively through automated and centralized management.



The Take-Away

To manage video data as a strategic asset, companies need a smart infrastructure that is purpose-built for mission-critical video – one that is designed for video-based workloads, reduces risk and liability, and is simple to manage and scale.



Purpose-Built: Recognize the Difference

Hyperconverged infrastructure (HCI) combines server, storage, and network resources in modular appliances to provide simplified deployment, scaling, and management. But is general-purpose HCI or virtualized server-SAN systems a good fit for mission-critical video systems. Check out the differences below:

Considerations	HCI Purpose-Built for Mission-Critical Video	Conventional HCI
Video Workload	Optimized to capture every frame at any frame rate	Optimized for read-heavy workloads (databases), not write-heavy workloads (video)
Resilience	Protects video data through a fault tolerance method known as Erasure Coding, a more storage-efficient method than replication	Relies on replication for data protection which consumes more available capacity and has substantial performance overhead, making it impractical for realtime processing of video data
Hardware Failure	Uninterrupted performance in the event of component failure	Cannot handle node failures without substantial performance degradation
Scalability	Easily scale up with cost-optimized storage-only nodes (a plus for video, when storage growth outpaces compute needs)	Requires you to buy more costly full-HCI nodes (which means deploying compute power that you don't need)

Are You Still Relying on NVRs?

The network video recorder (NVR) was the industry standard for video surveillance for many years. However, with the growing sophistication of video surveillance systems – and the growing importance of video data – NVRs have come into question.

Here are three reasons that NVRs can't meet the demands of mission-critical video:

- NVRs are based on Direct Attached Storage (DAS) technology, which are not designed to handle the highly variable bandwidth and intense throughput requirements of video. IT departments moved away from DAS over 15 years ago due to scalability, data availability, and resource sharing limitations, all key requirements in today's mission-critical video surveillance and security environments.
- The only way to “scale” NVRs is to add more NVRs, which requires complicated provisioning, constant management and manual load balancing
- When any component of an NVR fails, recording stops and both live and recorded video is inaccessible and possibly lost

To take care of your video data, work with a vendor that understands what it takes to manage video as a strategic asset for mission-critical use cases. Pivot3 offers a purpose-built infrastructure optimized for video capture, analytics, and visualization, helping simplify how you manage and extract maximum value of out of your video data. Specifically, with Pivot3 you can realize the following savings over conventional NVR and 3-tier approaches (separate server-storage-network systems)



- Cut TCO in half
- Reduce video storage requirement by 50%
- Reduce footprint and management by 85% vs NVRs
- Reduce power and cooling
- Reduce software management and maintenance costs

To learn more, visit pivot3.com/product/surveillance