intransa®

High Availability Concepts for Video Surveillance

White Paper

Table of Contents

Introduction

Video surveillance is often primarily a life safety application. In most installations video data must be *recorded* 24 hours a day, seven days a week, with no gaps, no frameloss, and without downtime.

Sometimes it is also important that video data also be available for immediate or near-immediate *playback*.

Recent growth in the video surveillance recording and retention systems market is exciting, but in any new marketplace, misuse of terms and exaggerated claims are a hazard. The market for video surveillance recording and retention systems has now evolved to the point where it is appropriate to define ever more frequently used terms. This is particularly true for those terms utilized to describe various forms of "High Availability" in video surveillance recording and retention solutions.

This paper reviews the concepts and terminology that evolved over several decades to describe High Availability in the computer systems industry. It then relates these concepts to product functionality in today's video surveillance recording and retention systems.

We will demonstrate that there were, and are, many different ways to maintain the availability of applications — *and the integrity and accessibility of data* — both historically in the computer systems business, and in today's video surveillance recording and retention environments.

High availability is an *optional* capability — *a characteristic* — that can be added to a computer or many video surveillance systems to mitigate or offset an identified risk of loss. Like many forms of insurance, its purchase is optional. Just as in other aspects of life and business, if you don't anticipate a particular event, need, or form of loss, you don't need that particular type of insurance. If you don't need to read the data in a short period of time, you may not need to pay for read data protection.

Learning about these different characteristics and techniques gives us a chance to understand our choices: agreement on concepts and terms is the foundation for the next generation of customer requirements, and the features — *as well as the functionality* — of future generations of video surveillance recording and retention products.

Defining High Availability

Availability[1] is the percentage of time that a system is capable of serving its intended function.

Availability is measured as a *level of service* applications on the system provide. *High availability* solutions are steps a business takes to ensure that their systems consistently provide the expected level of service, without significant interruption.

One aspect of High Availability is for the data *Writes* (recording). This is normally called *Failover*, whereas the application "fails over" to another resource, but continues to run.

The other aspect of High Availability is for the data *Reads* (playback). This is normally called *Read Protection*, where in spite of a loss of a storage element, the data is "available for the Read".

Each of these aspects of High Availability is driven by different customer requirements, and is mitigated with different potential solutions.

Because all businesses are different, their risk profiles are also different. A loss of service level for a few crucial moments may demolish one firm's reputation, create a significant liability for another, or cost a third organization a few hundred dollars, and another millions. Depending on a particular system's role — *the function in a company the system fulfills* — downtime has therefore has a specific and often overlooked cost. Having a High Availability solution therefore means different things in different environments, some critical and others perhaps not as much.

In the computer systems market, downtime comes in two forms: *Planned* and *Unplanned.*

Typically the IT (Information Management) department initiates *planned downtime*. During this period, management can schedule maintenance, patches to system software, and hardware upgrades for "off" hours.

Unplanned downtime is the result of events not within the direct control of IT administrators. Hardware failures, software problems, power failures, the Internet being down, administrator errors, and natural disasters cause unplanned downtime. Dramatic unplanned downtime events may never occur, but the cost

[1] A glossary containing detailed definitions of terms used to discuss High Availability is at the end of this document.

of such a catastrophe impacting some workloads is great enough that it can justify more expensive hedges against such threats.

The Cost of Downtime

When computer systems are running fine, it's easy to think nothing bad will ever happen. eBay's CIO probably felt that way on June 10th, 1999, the day before an administrative glitch brought the company's site down for more than seventeen hours, while investors on NASDAQ hammered the stock. eBay lost millions. A significant outage will cost companies in certain industries — *and different specific niches within those industries* — very different amounts on an hourly basis (Table 1).[2]

Industry	IT Service	Cost of Downtime per hour
Financial	Brokerage operations	$6,450,000
Financial	Credit card authorization	$2,600,000
Financial	ATM fees	$14,500
Media	Pay per view	$150,000
Media	Ticket sales	$69,000
Retail	Home shopping	$113,000
Retail	Catalog sales	$90,000
Transportation	Airline reservations	$89,000
Transportation	Package shipping	$28,000

Table 1 Cost of downtime per hour for IT services in the financial, media, retail, and transportation industries.

[2] These figures are used in a variety of places on the Internet with attributions to Gartner, the Meta Group, and others, but apparently date to a Contingency Planning Research study in 1996. (http://www.eaglerockalliance.com). Current costs of downtime per hour for many IT services are potentially much higher.

Reducing the Probability of Downtime

Availability is expressed as the percentage of time that a system is available for use.

For example, a requirement of 99.9% availability over a one-year period allows only 8.8 hours of downtime. The relationship between some common systems availability targets and the amount of downtime that they permit annually is shown in Table 2.

Level of Availability (%)	Annual Downtime
100%	0 minutes
99.999%	5 minutes
99.99%	53 minutes
99.9%	8.8 hours
99 to 99.5%	87.6 to 43.8 hours

Table 2 Decreasing levels of availability translated to minutes and hours of annual downtime.

Whether the task is to avoid complete loss of a service or to maintain a service in a degraded form, High Availability solutions rely on the simple principle of supplying redundancy (extra, costly hardware and software) to possible points of failure to maintain service availability.

Striving for inappropriate High Availability in the computer systems market is not cost effective. Resources devoted to addressing various unplanned downtime scenarios should reflect both the cost of a major outage as well as its relative probability. Spending $100,000 on a High Availability solution that shifts a workload's uptime from 99.99% to 99.999% is not justifiable if the loss of that particular workload only costs your business $1,000 per minute.

Customers in different IT service environments face different risks of loss from system downtime, and so make very different decisions about what types of redundancy they need to purchase.

Achieving High Availability

Marcus and Stern describe four stages of availability*: Basic, Data, System,* and *Organization* (Table 3).[3]

Availability Index	Stage	Example
1	Basic	No specific extra measures
2	Data	Data backup / archiving
3	System	Redundancy, fail-over
4	Organization	Alternate sites, systems, services

Table 3 Marcus and Stern's four stages of availability.

There are three pillars to configuring a highly available system

- *Technology*
- *People*
- *Process*

Technology gets the most attention, although People and Process have a significant impact on how much additional availability gets delivered.

Most estimates are that achieving "High Availability" in the computer systems market is 20% technology, 40% vendor support, and 40% IT process.

Technology

The traditional approach to deliver High Availability is through hardware and software redundancy. All computer systems are comprised of several critical components, each of which may fail in different ways.

There is at least one, and often several, approaches to mitigating each type of failure for each component. Most commonly computer systems fail at one of the following levels:

- *Applications*
- *Operating System*
- *Virtual Server*
- *Physical Server*
- *Networking*
- *Storage system*

[3] Marcus, Evan and Hal Stern. *Blueprints for High Availability* (Wiley, 2003)

- *Power*

Component Layer	Sources of Downtime	Solutions for High Availability
Applications	Software bugs, software interactions, and user error	Some applications handle their own failover or recovery (e.g., Oracle Database or Microsoft SharePoint.)
Operating System	Software bugs, software interactions, administrative error	Add Clustering software (active/passive or active/active)[4]
Server Virtualization Software (optional)	Software bugs, software interactions, administrative error	Add Hypervisor Failover/Restart (active/passive or active/active)
Physical Server	Memory fault, CPU or circuit board problem, power supply or fan failure	Add ECC or parity memory Use Fault Tolerance: dual independent (active/active) systems running the critical application(s) in lockstep.
Network	Stale network directory (DNS), switches failure	Add redundant networking cards, redundant network cabling, bonded networks, MPIO
Storage System	Enclosure or individual disk failure	Add mirrored disk drives or RAID parity. Use Tape backup and offsite archiving
Power	Power supply failure, loss of power from the electrical grid	Add redundant power supplies, Uninterruptable Power Supply, and external backup generator

[4] Server clusters can take two forms: active/passive clusters and active/active clusters. In active/passive clustering, the cluster includes active nodes and passive nodes. The passive nodes are only used if an active node fails. In active/active clusters, all nodes are active. In the event of a failover, the remaining active node takes on the additional processing operations, which causes a reduction in the overall performance of the cluster. (*Server Cluster Overview*, http://technet.microsoft.com/en-us/library/cc759183%28WS.10%29.aspx) Active/Active clusters are more cost-effective, Active/Passive cluster configurations are better if degraded performance after a failover is not acceptable.

Table 4 Availability concerns and solutions in the functional layers of a computer system. Failures in any one part of the stack can bring down the entire system.

- **Applications**
  Applications can fail by being in an unexpected state, due to coding problems, a fault in the underlying Operating System, memory parity errors, or some other hardware problem. Some applications, including databases like Microsoft SQL Server, make discrete writes to disk, transaction by transaction, and can easily be restarted — *without loss of data* — on another server in a cluster. These applications can also be restarted in a virtualized environment spanning multiple, similar servers. Applications can continue to run through a failure, without a restart or other impact on users, if they are running on multiple CPUs in lockstep on a Fault-Tolerant system.

- **Operating System**
  An Operating System (OS) fails for many of the same reasons an application fails. In a Clustered environment, an Operating System can continue to provide services even when all the systems in the cluster, except one, have failed. Clustering technology can both be extended so applications and services fail over to a remote site, continuing to provide services, at the cost of another server, disks, software, and a high bandwidth connection to write data — *in real time, every day, all day* — from the first cluster to the secondary site.

- **Server Virtualization Software (Hypervisor)**
  Server virtualization software allows a single physical server to run multiple Operating Systems. This is a component of *server consolidation,* achieved by using a Hypervisor. Some types of server virtualization software can be run across multiple physical servers – applications running on top of such a virtualized environment have constant access to physical server resources even if one server fails. This is functionally equivalent to the application running on a high-availability Cluster.

- **Physical Server**
  Many things can bring down a server, including memory errors, a bad CPU chip or logic board, bad drivers, or overheating. Solutions include using parity-checking memory, paying attention to console error messages, conducting regular system maintenance updates, cleaning dust out of the system regularly, and using redundant fans.

- **Network**
  Networks are comprised of switches, routers, and physical cabling. Switches can fail, routing tables can be corrupted, and Network Interface Cards (NIC) can fail, like any other electronic component. Physical cabling can be

damaged.

- **Storage Systems**
  Storage virtualization, delivered either as block-level Storage Area Network (SAN) technology, or file-level Network Attached Storage (NAS) technology, can fail. There are multiple strategies for providing High Availability in SAN/NAS environments including Storage Clustering and Redundant Data Paths.

  Disks drives are ultimately mechanical devices, operating within very fine tolerances and hence relatively likely to fail. To protect vital business data, disk drives can be mirrored, where the operating system writes data identically on two drives at once. Disk drives can also take advantage of RAID — *redundant array of independent disks* — a technology that provides increased storage reliability through redundancy, combining multiple low-cost, less-reliable disk drives into a logical unit where all drives in the array are interdependent. In the event of a failure, data written with RAID parity striping can be recovered.

  RAID techniques can be extended across multiple drive enclosures or nodes in a Storage Cluster, defending against the failure of a single array.

- **Power**
  Servers can fail because of a failed power supply device or a power failure at the local electric utility. Servers can have redundant power supplies (and cords) or an Uninterruptable Power Supply (UPS) and an Auxiliary Power Generator between servers that need to stay available.

People

New versions of applications and updates from one or more vendors can conflict, or not have been tested on a firm's particular hardware or software configuration. At mission-critical sites, IT typically tests software updates and new applications on a test-bed system, verifying them before running them on the production server. This is called Change Management and Configuration Control.

Process

A trusted vendor or other service provider can use remote monitoring to track service events and manage the service process, speeding the initiation of a service action when a problem is detected.

A mission critical site can purchase a premium service contract to get fast track call resolution and 24x7 access to top-level software specialists. With a premium contract for hardware, the systems vendor may store critical replacement parts on site, while assigning a full time service technician to the customer's facility.

High Availability Concepts for Video Surveillance Systems

Customers in the video surveillance recording and retention industry include casinos, banks, airports, prisons, schools, retailers, museums, and many others.

The number of surveillance cameras can range from a handful to several thousand. Basic requirements for video surveillance include:

- *"Always On" — no downtime, continuous streaming video*
- *Quality metric of minimal or* no *frameloss*
- *Affordability / cost-effectiveness*
- *Scalability of storage from 2 to 2000 Terabytes*

There are two basic, fundamental aspects to High Availability in video surveillance: ***Write Failover*** and ***Read Data Protection***.

Write Failover (recording)

Write Failover means no matter what, keep the camera *recording.*

Can you keep the cameras recording even when there is a system failure?

- *Application failover built into the Video Management Software (VMS)*
- *Operating System failover (active/active) with Windows Clustering with shared memory*
- *Virtual Server failover (active/passive or active/active) from Hypervisor*
- *Physical Server failover (active/passive or active/active) with "hot standby" server or shared memory*
- *Storage Virtualization (SAN) with failover and automatic load balancing*

Read Protection (playback)

Read Data Protection means the camera video is *available for playback.*

How long can the customer wait to have access to recorded video after a system failure? 2 nanoseconds? 2 minutes? 2 hours? 2 days?

Most customers do not need 2 nanosecond, 2 second, or even 2 minute access to video data. Why should a customer pay a premium for availability they don't need? Understand the real business and technical requirement before buying redundant components.

- *Access to recorded video in SAN by multiple servers*
- *High Availability virtualized storage (clustered storage, auto load balancing)*

- *RAID0 (striping), RAID1 (mirror), RAID5(n+1), RAID6(n+2)*
  - Disks within an Enclosure
  - Across Enclosures
  - Across Nodes (cluster)
- *Hot-swap drives*

For <u>both</u> Write and Read High Availability

- *Network Interfaces Failover*
  - Multiple networks
  - Bonded NICs
  - MPIO (in Operating System)
- *Power Failover*
  - Dual power supplies
  - Uninterruptable Power Supply (UPS)

Summary

As we have shown, High Availability is a *characteristic* and not a unique product.

This characteristic (High Availability) can be implemented for Write Failover, *or* Read Data Protection, or *both*.

It is important to understand the customer's cost of downtime *and* the customer's sensitivity to the time it takes to be able to read the video after a system failure.

Not all customer situations require High Availability. To impose the complexity and cost of High Availability on a customer who doesn't really need it is a disservice. Choose the appropriate Availability Tier for the customer situation.

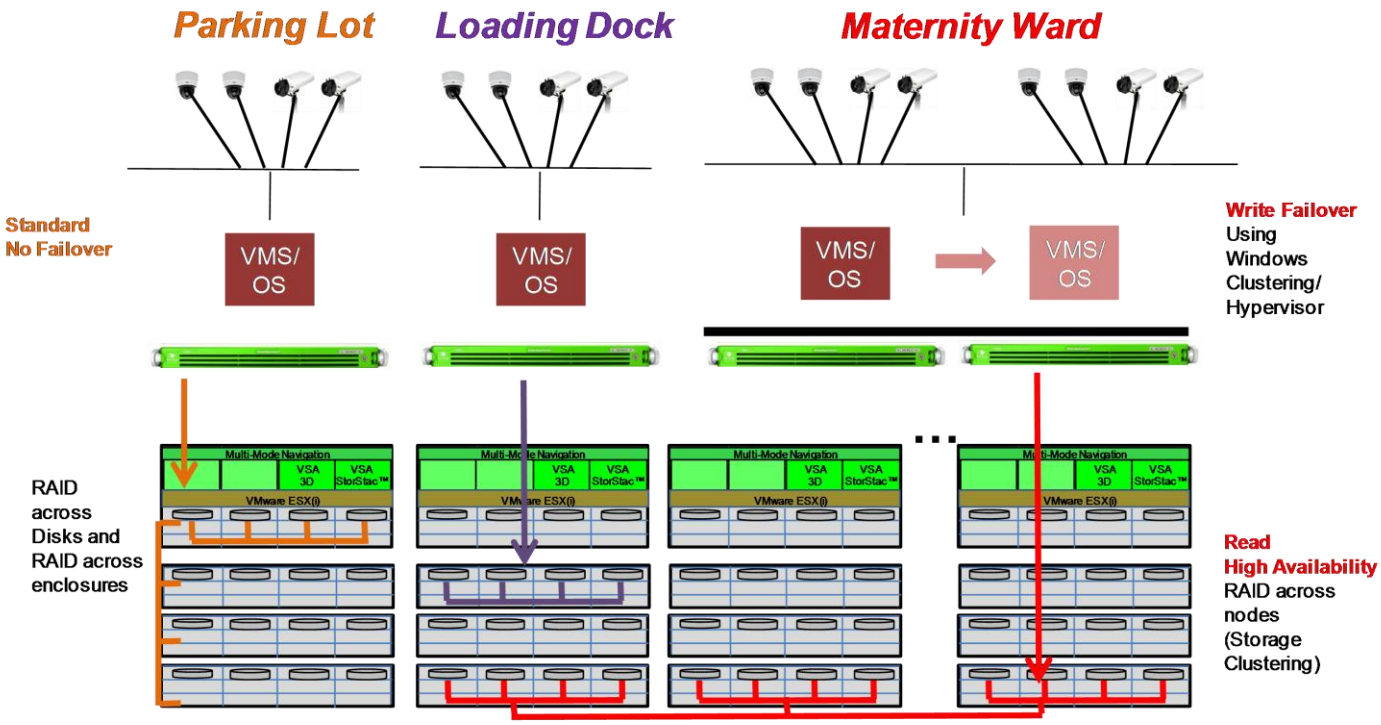

Table 5 In a single environment, such as this example of a modern hospital, all areas require no loss of video, and the Intransa VideoAppliance™ system is designed to support that need. But only the Maternity Ward requires always-ready, instant access to live and recorded video, for legal and life safety needs and thus requires full High Availability. The Parking Lot and Loading dock can be supported with less expensive resources delivering Fault Tolerant protection, and is all possible on either single or multiple appliances with Intransa VDMR's VSA Tiered Availability feature.

Glossary

Availability	The degree to which a system or subsystem is in an operable and committable state at the start of a mission, when the mission is called for at an unknown time. Simply put, availability is the proportion of time a system is in a functioning condition. This is often described as a **mission-capable rate**. Mathematically, this is expressed as A=1 - unavailability or A= uptime/(uptime + downtime).
Fault tolerance	**Fault-tolerance** or **graceful degradation** is the property that enables a system (often computer-based) to continue operating properly in the event of the failure of (or one or more faults within) some of its components. If its operating quality decreases at all, the decrease is proportional to the severity of the failure, as compared to a naïvely-designed system in which even a small failure can cause total breakdown. Fault-tolerance is particularly sought-after in High Availability or life-critical systems. Recovery from errors in fault-tolerant systems can be characterized as either roll-forward or roll-back. When the system detects that it has made an error, roll-forward recovery takes the system state at that time and corrects it, to be able to move forward. Roll-back recovery reverts the system state back to some earlier, correct version, for example using checkpointing, and moves forward from there. Some systems make use of both roll-forward and roll-back recovery for different errors or different parts of one error.
High availability	High availability is a system design approach and associated service implementation that ensures a prearranged level of operational performance will be met during a contractual measurement period.
Mission-capable rate	See **Availability**.

Parity protected memory	The purpose of parity RAM is to detect when a memory error (corrupt bit or bits) has occurred. Undetected memory errors can have varying results; from simply annoying to catastrophic. In the case of the home PC where data integrity is often perceived to be of little importance, non-parity memory is an affordable option. However, if any sort of data integrity is required, parity memory would be the minimum level of protection. RAM with ECC or Error Correction Code can detect and correct errors. Additional information needs to be stored and more processing needs to be done, making ECC RAM more expensive and a little slower than non-parity and parity RAM. This type of ECC memory is especially useful for any application where uptime is a concern. Failing bits in a memory word are detected and corrected on the fly with no impact to the application.
Planned downtime	Downtime is used to refer to periods when a system is unavailable. Scheduled downtime is a result of maintenance that is disruptive to system operation and usually cannot be avoided with a currently installed system design. Scheduled downtime events might include patches to system software that require a reboot or system configuration changes that take effect upon a reboot. In general, scheduled downtime is usually the result of some logical, management-initiated event.
RAID	RAID, an acronym for redundant array of inexpensive disks or redundant array of independent disks, is a technology that provides increased storage reliability through redundancy, combining multiple low-cost, less-reliable disk drives into a logical unit where all drives in the array are interdependent. RAID is often used as an umbrella term for computer data storage schemes that can divide and replicate data among multiple disk drives. The schemes or architectures are named by the word RAID followed by a number (e.g., RAID 0, RAID 1). The various designs of RAID systems – whether software or hardware -- involve two key goals: increase data reliability and increase input/output performance. When multiple physical disks are set up to use RAID technology, they are said to be in a RAID array. This array distributes data across multiple disks, but the array is addressed by the operating system as one single disk.

Read data protection	The data is available for playback when the customer needs it.
Redundancy	In engineering, redundancy is the duplication of critical components of a system with the intention of increasing reliability of the system, usually in the case of a backup or fail-safe. In many safety-critical systems, such as fly-by-wire and hydraulic systems in aircraft, some parts of the control system may be triplicated. An error in one component may then be out-voted by the other two. In a triple redundant system, the system has three sub components, all three of which must fail before the system fails. Since each one rarely fails, and the sub components are expected to fail independently, the probability of all three failing is calculated to be extremely small. Redundancy may also be known by the terms "majority voting systems" or "voting logic".
Reliability	Reliability is the ability of a person or system to perform and maintain its functions under stated conditions for a specified period of time.
Remote monitoring	Provides event monitoring and service management for critical production systems via a modem or the Internet.
Service Level Agreement	A service level agreement (SLA) is a negotiated agreement between two parties where one is the customer and the other is the service provider. The SLA records a common understanding about services, priorities, responsibilities, guarantees, and warranties. Each area of service scope should have the "level of service" defined. The SLA may specify the levels of availability, serviceability, performance, operation, or other attributes of the service, such as billing. The "level of service" can also be specified as "target" and "minimum," which allows customers to be informed what to expect (the minimum), whilst providing a measurable (average) target value that shows the level of organization performance.

Uninterruptable Power Supply	An uninterruptible power supply (UPS) is an electrical apparatus that provides emergency power to a load when the input power source, typically the utility mains, fails. A UPS differs from an auxiliary or emergency power system or standby generator in that it will provide instantaneous or near-instantaneous protection from input power interruptions by means of one or more attached batteries and associated electronic circuitry for low power users, and or by means of diesel generators and flywheels for high power users. The on-battery runtime of most uninterruptible power sources is relatively short—5–15 minutes being typical for smaller units—but sufficient to allow time to bring an auxiliary power source on line, or to properly shut down the protected equipment.
Unplanned downtime	Unscheduled downtime events typically arise from some physical event, such as a hardware or software failure or environmental anomaly. Examples of unscheduled downtime events include power outages, failed processor or memory components (or possibly other failed hardware components), an over-temperature related shutdown, logically or physically severed network connections, catastrophic security breaches, or various application, middleware, and operating system failures.
Write failover	No matter what happens, data is being written to disk.

Source: http://en.wikipedia.org

For a more extensive physical security and Information Technology Glossary of Terms please visit *Intransa.com*, at http://www.intransa.com/resources/glossary.php.

intransa®

Intransa, Inc. / www.intransa.com / www.videoappliance.com

10710 N. Tantau Avenue, Cupertino, CA 95014 USA

Toll free 866.446.8726 **International** +1.408.678.8600 **Email** sales@intransa.com