Implementing Systems at Zero Incremental Cost

from

iOmniscient

Contents

1. Introduction

There is a general assumption that an organization will incur incremental costs as one adds new technology. This is normally the case in most domains. If one buys a car with advanced features it will normally cost more than a car with basic features.

iOmniscient decided to reverse this trend in the field of video analytics. It established a design principle that newer technologies must reduce infrastructure costs to the point where an intelligent iOmniscient system can be implemented at a lower cost than an equivalent system with no intelligence.

There are twelve areas where infrastructure costs can be reduced which are:

1. Type of Cameras
2. Number of Cameras
3. Storage
4. Network Bandwidth
5. Computing Power and avoidance of the Use of GPUs
6. Software
7. Reduction of Labour through building complex use cases
8. Command and Control Centers
9. Video Utilities
10. Installation
11. Maintenance
12. Response Time

Not every factor will be relevant to every user but generally the intent is that they can do more with less. And this is by design.

2. What are Intelligent Video Analytic Systems

To ensure we start from a common understanding of the technology let us define what we mean by a video analytic system. There are essentially four classes of intelligent video analytic system used for surveillance.

Behaviour Analytics Systems. The first class consists of Behaviour Analysis Systems. This involves a camera looking at a wide scene. The software in such a System essentially attempts to detect behaviours in the scene. The behaviour could be by humans (or animals) or it could be the behaviour of objects such as cars, trucks, boats or airplanes.

The system can determine what the person or object is doing (e.g. the direction and speed at which it is travelling). It can also understand its characteristics – for instance that it is person in a red shirt. Simple detection systems can detect basic behaviours such as detecting that someone has entered an area of interest. The more sophisticated systems can tell if someone has fallen down or if they have abandoned a bag in a crowded place and walked away.

The more advanced systems can also understand what is happening in a very complex scene such as the detection of smoke or fire in a crowd, determining if a person is attempting to commit suicide or the dynamics of how a queue is formed.

An IQ rating system can be used to rate the intelligence of Behaviour Detection systems which works just the way it does for humans. For humans, an IQ of 100 is the average IQ of the entire population. Hence 50% of the population has an IQ of less than 100 and of course 50% has more.

Similarly one can have systems with very basic IQ where they essentially work on Motion Detection and those that can be considered to be Genius systems which are able to detect things like theft and abandoned objects in very crowded and complex scenes and even see them when they are virtually invisible to the human eye.

Categorization Systems. At the next level these systems can categorize what has been detected. For instance if an abandoned object is detected it can determine if the object is a bag or trolley or even a sleeping child.

Recognition Systems. The third class of video analytic products focuses on recognition. Usually they require a slightly closer view of the persons or objects being observed and they can recognize who or what was involved in a particular incident. If it involves a person, a Facial Recognition system can be used to recognize who he is. If it involves a vehicle the system can recognize it by reading its number plate.

Even with recognition products there can be varied levels of sophistication. There are systems that can do face recognition in very controlled environments with very high resolution images. The system will only work when the person looks straight at the

camera and is in an environment with very constant lighting. A more advanced system can operate with much lower resolution cameras from a far distance where the person is often not even conscious that there is a camera watching him.

Convergent Systems. There is now a final class of products emerging. This is the class of convergent products that can do Behaviour Analysis, Categorization and Recognition at the same time on the same camera enabling the user to implement complex use cases relevant to the user.

3. Savings

3.1 Type of Cameras

iOmniscient systems are designed to use low resolution cameras. Behaviour Analytics can be performed at high accuracy using 1/4 CIF (QCIF) resolution (which is approximately 0.025 megapixels) and with a frame rate of 1 to 6 fps. If a higher resolution camera is used then this means that the behaviour can be detected at a much further distance.

Even for Recognition, an iOmniscient system can accurately recognize a person in a crowd at 20 to 25 meters on a 2 megapixel camera. This is because the system is designed to require only 22 pixels between the eyes for good recognition (and even down to 12 pixels with a slightly lower level of accuracy). This compares with the 60 to 100 pixels required by Face Recognition systems that are generally available on the market. In fact NIST does not even have a data set to test below a resolution of 45 pixels reflecting the types of systems generally available in the market today.

Overall therefore, iOmniscient systems can achieve the same results on low resolution cameras that others require high resolution for.

3.2 Number of Cameras

Because an iOmniscient Face Recognition system requires only 22 pixels between the eyes compared to the 60 to 100 pixels required by other systems, it can cover a far wider area. In a typical situation if the requirement is to recognize people walking down a 10 meter wide corridor iOmniscient would require a single camera. Others would require between 5 to 10 cameras forcing them to operate only at narrow choke points.

Using fewer cameras is not reflected just in a lower cost for cameras. It translates into fewer software licenses, less computing and storage, reduced network bandwidth requirements and less installation and wiring.

In a recent large airport the consultant had suggested that they required 600 cameras. On redesigning it using an iOmniscient system the airport needed 120 cameras to achieve the desired results.

The same is true for behaviour analytic. An iOmniscient system can detect behaviours at very low resolution and low frame rate. For instance it can detect an abandoned object when it is only 4x4 pixels in size in a video stream running at 1 frame per second or less. Others require the object to be 20x20 pixels before it can be accurately detected. Again this means that a given camera can cover a much larger area. Hence one requires fewer cameras overall.

3.3 Reduced Storage and Network Bandwidth

New Cameras have increasingly higher resolution. Today it is impractical to transmit video from a large number of high resolution video streams to the Cloud or other central location with the internet bandwidth that is available today even in many advanced countries.

One then has a choice.

One can transmit and record the video in the resolution of the camera - which requires an enormous amount of bandwidth and storage (and of course one needs fibre or other ultra-high bandwidth network)

Or

One can transmit and record at a lower resolution which means that one would lose the detail that the cameras can provide. That is a waste – why deploy high resolution cameras if one cannot take advantage of that resolution.

The patented Smart Compression from iOmniscient solves this dilemma.

It records every face, license number and significant event that is seen in the camera in high resolution while the rest of the video is recorded at low resolution. In every single frame of the video one never misses anything that is important. Equally one does not waste bandwidth and storage on anything that is unimportant.

To enable Cloud Computing with high resolution cameras one needs enormous bandwidth to stream the video. With iOmniscient's Smart Compression very low resolution video can be streamed even while some parts of the video (e.g. important details such as the face or license plates or key events) are retained in high resolution. In other words within the video itself, faces can be streamed at high resolution and the rest of the person and the background would be at low resolution. This allows you to stream at low bandwidth while still transferring all the important information.

This system has the advantage of saving over 90% of the network bandwidth and storage needed. If the previous system could store video for 6 months, an iOmniscient’s Smart Compression one could store that video for 60 months (5 years) using the same amount of storage. Or looking at it another way you need only 10% of the storage you previously required.

Consider a user who installs a 2 mega-pixel camera. He can record the video at 2 mexa-pixels or he could reduce the image size and record at 1xCIF (equivalent to 0.1 megapixel).

If he records at 2 megapixels (mp), it takes an enormous amount of storage and requires significant bandwidth for transmission. However, much of this resource is wasted as the network will be transmitting and the recording system will be recording a lot of useless information continuously because the user does not wish to miss the small amount of important information that may be there.

If he records at 1xCIF or at anything less than the full resolution of the camera, he loses the detail which he might need for investigation in the future and hence has got little benefit from using a 2 mp camera.

Smart Compression from iOmniscient

iOmniscient's Smart Compression is an internationally patented technology that enables the user to save enormous amounts of network bandwidth and storage by transmitting and recording only the important details at high resolution and the rest at a lower resolution (even down to 1xCIF).

What is important?

The User can decide what information is important to him. He may want to capture the detail of every person or vehicle that enters the scene or he may want to be selective and only capture the details of persons or vehicles that exhibit certain behaviours (eg only people who fall down or only blue vehicles that have an accident).

If one is trying to recognize people, the Smart Compression unit only needs to maintain the face of the person in High Resolution. If the rest of the image is transmitted/ stored in a lower resolution it will not negatively impact the recognition of the person.

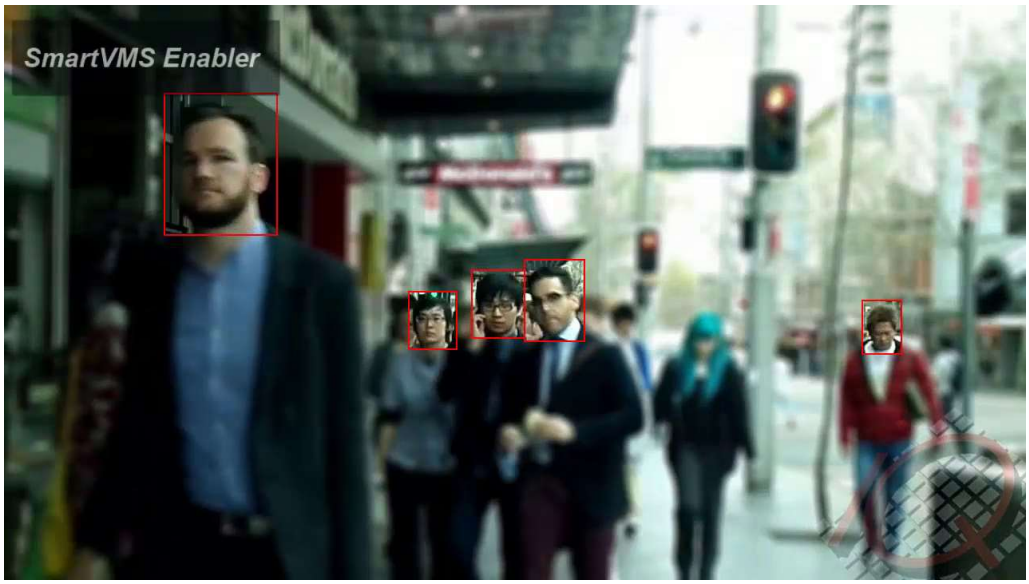

How smart can the Smart Compression be? As smart as the user wants it to be. As the provider of the world's most comprehensive multimedia analytics capability, iOmniscient can help to detect and recognize events with varying levels of intelligence including with its patented technology for crowded and complex scenes. For details visit iOmniscient's website.

Benefits

The main benefit of the iOmniscient's Smart Compression is the enormous saving in Storage and Network Bandwidth. For a 2 megapixel camera the saving can be enormous. To make this tangible, if the normal cost of storage required to record the video is $20,000 it could be reduced to $2,000.

Example:

An airport has 3,000 cameras, each with a resolution of 2 megapixels. It will require 8,500 terabytes of storage for recording 30 days of video using H.264. If the video is not stored at a resolution of 2 megapixels they will lose all the detail.

The iOmniscient Smart Compression will allow the airport to store the video at whatever resolution they like (e.g. at 1xCIF) without losing the detail because the detail will still be recorded at a high resolution.

The average cost of storage today is US$300 per terabyte. With Smart Compression the airport will save around 90% of the storage – a saving of US$2.3m – which can be used to purchase intelligence that makes the system more useful and the user more productive.

In environments where only low bandwidth networks (e.g. 3G wireless systems) are available, the Smart Compression can still transmit providing the user with the important information at high resolution.

In this case the analytics and compression must be performed at the Edge using a Super Edge device of the types shown below. These devices are placed near the camera.

Today our software is implanted on these devices and they can help to make any existing camera smart. This strategy has been employed to ensure that the software is camera agnostic and widely available. There is no technical constraint to embedding the software in the camera itself (though this has not been done for strategic reasons).

Wired

Wireless

Ruggedized for outdoor use

Edge devices embedded with iOmniscient's Smart Compression.

The table below shows the savings in network bandwidth that can result from transmitting only the important information at high resolution and the rest at a lower resolution.

Bandwidth Savings:

Camera Resolution	Frame Rate	Compression Format	Bandwidth Calculation
2 megapixel	15FPS	H.264	**5.51 MBit/s**
With Smart Compression	15FPS	H.264	**0.55 MBit/s**

Note: that only unimportant information is transmitted at low resolution. Important information would still be transmitted at the maximum resolution of the camera.

How does it Work?

The Smart Compression determines the parts of the video input that are important based on the comprehensive video analytics provided by iOmniscient. This could be several seconds of video or even areas of interest within a single image (e.g. just the face of a person or a License Plate) which is maintained by the software in higher resolution.

The rest of the video is compressed to a lower resolution to whatever level is required (eg 1 xCIF).

An inbuilt Transcoder ensures that input video in any format, resolution or frame rate is transformed into the format, resolution and frame rate required for recording in the VMS.

The entire system only works reliably because of the very robust and sophisticated Video Analytics from iOmniscient.

The video can then be stored on ANY VMS. A unique feature of the system is that the information recorded in the VMS can be viewed continuously with the faces, plates etc at high resolution with the rest in low resolution.

Alternatively, important information on key events can be displayed independently at the desired higher resolution. The results can be portrayed on dual monitors with the low resolution video being displayed separately from the components of the image in high resolution.

Architecture:

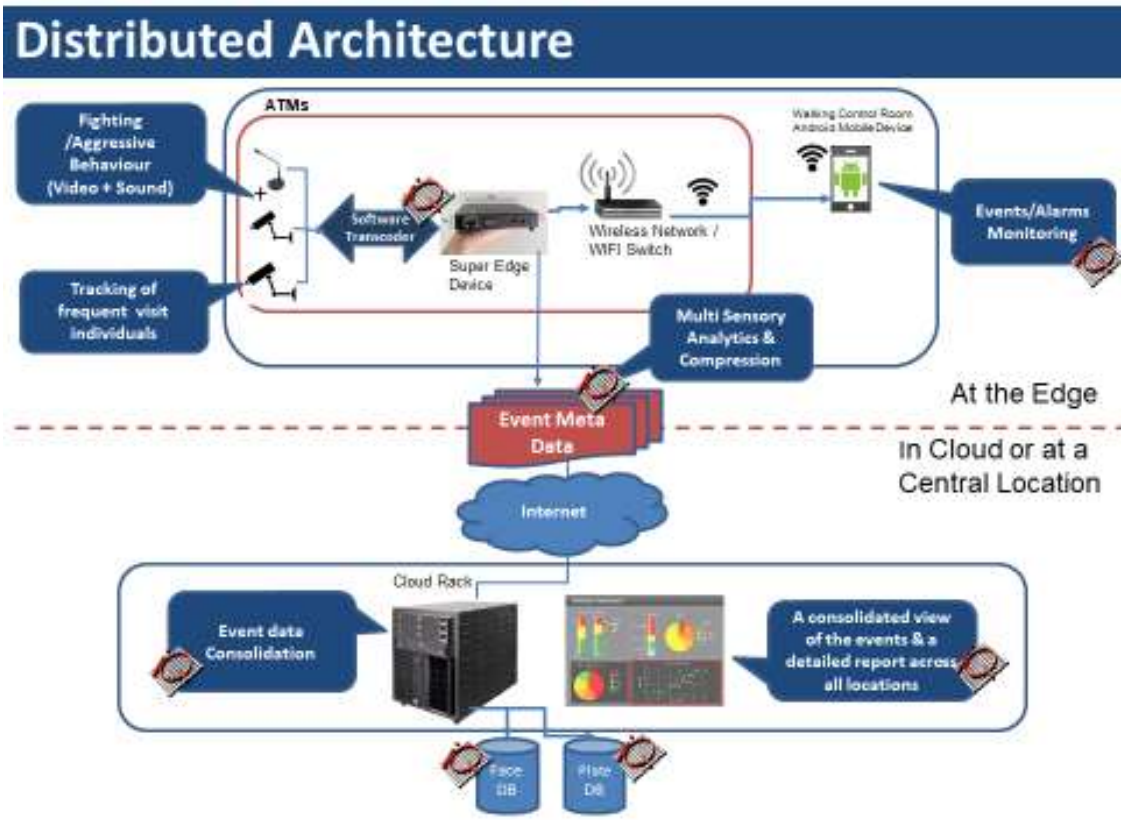

In a distributed design, the Smart Compression can be deployed on a Super Edge device close to the camera resulting in a saving of both storage and network bandwidth.

Hybrid designs are of course possible.

Can the video be used for evidentiary purposes?

Yes. Today video compressed using compression techniques such as MJPEG, MPEG4, H.264 and H.265 are already accepted as being of evidentiary quality even though the video is being compressed. They use techniques such as having I frames and P frames where some frames are transmitted whole while in others one sends changes since the last whole frame.

In iOmniscient's Smart Compression one is merely compressing the same video more intelligently without in any way modifying the original video.

Don't other suppliers have similar Smart Compression?

No. Some supplier have attempted to emulate iOmniscient's smart compression by compressing certain areas in the image at different resolutions. So they can draw an area of interest around a door and the whole area of the door can be recorded in high resolution and the rest in low resolution. The recording in high resolution will continue even if no one is walking through the door.

Others have attempted to use motion to determine what to compress. So if there is motion in the image the moving part can be left in high resolution and the rest in low resolution.

But motion detection is a crude tool for determining what is important. Every moving cloud and waving tree would be recorded. And in a busy airport where there is continuous motion everywhere there would be no benefit.

To recognize a person one does not need to see their shoe laces in high resolution - only their face. Only iOmniscient can be selective down to that level.

Why is this offering unique to iOmniscient?

To be effective the Smart Compression concept requires 2 key capabilities.

1. It requires the Smart Compression to separate out the important parts of the video from the less important ones – a concept patented by iOmniscient.

2. It requires a robust ability to determine what is important. For this, iOmniscient has the most advanced analytics systems in the industry including the patented ability to recognize people IN A CROWD and detect behavior IN A CROWD.

How much does iOmniscient's Smart Compression cost?

Almost nothing. The savings in storage is so huge that, today a user can implement a system with advanced video analytics and the savings in storage alone would pay for the intelligent software – hence the value proposition is that with this technology a user can implement a smart system at a lower price than he could implement a standard recording system.

Can any VMS system be made Smart?

Yes. While iOmniscient offers its own VMS, it is also integrated with most of the major VMS systems available in the market.

The Smart Compression and Intelligence can be added to any existing VMS system to make it smart.

3.4 Reducing the Cost of Computing

The prevalence of deep learning based technologies has resulted in the extensive use of GPUs. These add to the cost of the computing.

But humans do not have a single way for thinking. They use deductive and inductive logic, they can be creative. They may use their memory and their experience. Deep Learning assumes that there is only one way to perform analytics. iOmniscient has developed patented techniques to use the many different types of intelligence that humans use simultaneously. Deep learning is just one of these techniques. These methods greatly reduce the computing requirement for the analytics making the use of GPUs unnecessary.

In a typical example the user wanted to detect little nails and screws that fell on the steps of an escalator which could damage the mechanism. They did not want to detect pieces of paper or other benign objects that might fall on the step. If deep learning was used the system would be trained on thousands of different nails and screws. Then every image frame would be analysed to determine if there was a screw there. This involves a lot of heavy computing would require a GPU.

In an alternative method a simple computing light, heuristic system can determine if “something” is on the step of the escalator. Such detections may occur only once a day or even once a week. Once an object is detected only then is that single image sent to the categorization system to determine what the object is. As this categorization is done occasionally on a single image rather than on a video stream at 25 fps, no GPU is required.

Computing hardware has become a commodity pushing down the cost of computers. Hardware suppliers are attempting to lock their customers in to sustain higher prices by creating proprietary languages to access their GPUs. This threatens the concept of Open Systems that users have pushed for over time.

By having systems that do not require GPUs the user can not only reduce his costs but he also gains independence as his systems can be open and not locked in with a specific supplier or hardware.

3.5 The cost of the Software

How can one compare the cost of software between suppliers. There is often a very large difference between the cost of the software from different suppliers. There is also a big difference in what these different systems can achieve.

All iOmniscient systems are armed with an Artificial Intelligence based module called NAMS (Nuisance Alarm Minimization System). It helps to reduce the number of false alarms from a variety of factors such as variations in the light or other environmental conditions.

In a recent trial with 40 cameras one supplier who purportedly was offering a system at a quarter of iOmniscient's price, was generating 200 false alarms per camera each night. The iOmniscient system generated an average of 1 per night across the whole network.

What is the cost of running a system that generates 200 false alarms per camera. How many extra people would it take to just look at each false alarm? Is a system with that performance useable?

All iOmniscient systems are designed to offer visibly superior performance to anything else that is available on the market and the price is always competitive for a comparable system.

Free VMS from iOmniscient

Most users believe they need a Video Management System (VMS). A VMS is essentially a recording (and occasionally a display) system. Generally there is little to differentiate one from another and hence VMS suppliers attempt to lock in customers by creating proprietary interfaces.

But the VMS system is like the foundation of the house. iOmniscient is offering the whole house and not just the foundation. It is therefore agnostic with respect to the foundation on which the house is built. If the user already has a VMS system iOmniscient can integrate with it. If the user does not already have a VMS system, the iOmniscient VMS system is provided FREE for all cameras which use its analytics.

This too can significantly reduce the overall cost of the software required.

3.6 Reduction of Labour through Complex Use Cases.

Earlier it was mentioned that the latest new class of capabilities involves convergent technologies. Let us consider how one of these works.

Let us assume that there is a vehicle accident and the car is on fire.

A video analytics system could detect the accident. The alarm would be sent to a control room where a human was analyse the results and determine the next course of action. If the vehicles are on fire he might call the fire brigade

The operator can then zoom in on the number plates of the vehicles involved, capture these and put them through the License Plate Recognition system to check on whether they were stolen vehicles.

On viewing the video the human operator may realize that one of the drivers is running away from the scene. He may then initiate the Face Recognition system to attempt to recognize the driver.

All of this involves the operator in reviewing results from one analytics system and initiating action in another one. In an iOmniscient system most of these actions can be automated.

When there is an accident the system will itself determine if the vehicles are on fire. It can automatically initiate the reading of the number plates of vehicles involved and check if the vehicles are stolen. It can check for the faces of people who might be seen leaving from the vehicle and see if they are on any criminal list.

By addressing the entire use case the system can relieve the operator's workload and enable him to be used in more productive activities.

3.7 And why do you need a Command and Control Center?

Large organizations revel in having impressive Command and Control Centers. It gives them a sense of power and the illusion of control.

But independent studies have shown that an operator watching just 2 cameras misses half the action after 10 minutes and 95% of the action after 22 minutes. Does it make sense to have a huge control room (one organization we worked with had a staff of 1500 manning a control center which showed 12000 cameras on huge video walls. It cost the tax payer a huge amount and was totally ineffective in pre-emptively detecting events).

With a good analytics system the organization can operate with a blank screen. Only events detected by the analytics system are brought to the attention of the operator.

One argument advanced against such a scheme is: What if the analytics system misses the event. Does it not require a human to ensure that the event is detected? But we have already seen that the chance of a human detecting an event is miniscule. They only get involved in the post event investigation.

With the proliferation of iOmniscient's IQ-Mobile system all events can be seen on a Smart Phone. Everything that can be done in the Control Room can be done on the phone. The phone becomes a walking Control Room. It can be more useful to have a person in the field making their presence felt to help the public than to be sitting in a control room. And if there is an event he would receive the notification in any case and, being in the field, he is in a better position to respond to an incident.

But don't you require someone in the control center to receive an alarm and decide which field operator to send it to? The patented iOmniscient Automated Response System can achieve this with minimal human intervention. Once the system understands the situation it can find the nearest appropriate responder and tell them where to go and what to do.

Organizations may still require a Command and Control Center for those times when human intervention is required. However their size and scale can be greatly reduced as much of the work that is supposed to be done there (and rarely is) can be taken over by an Automated Response System.

3.8 Video Utility – reducing the cost of duplication

Large organizations often have multiple stakeholders with different objectives. Invariably each group implements its own infrastructure (on the assumption that their data is confidential and they can only maintain this confidentiality through a separate, isolated system) multiplying the costs for the organization as a whole.

In very large organizations different departments can put in their own surveillance system to meet their specific narrow objectives. In a shopping mall the marketing department may want to count the people traffic through the store. The security department may worry about recognizing known shoplifters. The Health and Safety department may have a need to detect accidents such as people falling down.

In a City Surveillance system, the traffic department may be interested in accidents and traffic jams, the local city council may be interested in ensuring that dustbins do not overflow while the police may be interested in apprehending known criminals.

In such organizations we often find that each group will be protective of their information and would set up their own independent CCTV network. As an example, in the City Surveillance system, the police may be doing Face Recognition and they may be concerned that by making their video footage available to others the privacy of citizens could be breached. They would therefore not want to share their video with other groups. The department concerned about detecting overflowing dustbins would be forced to install their own infrastructure. In one situation we found 7 government departments (each concerned about the confidentiality of their information) had installed cameras next to each other watching the same scene.

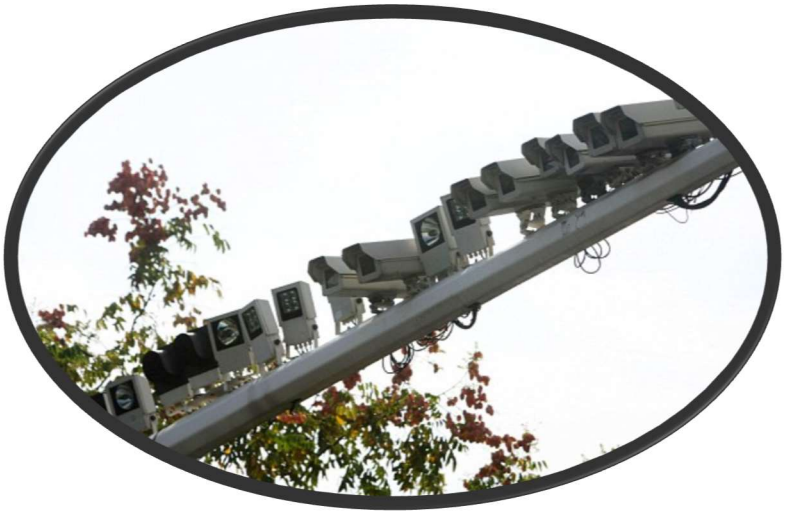

Consider an analogy. In times past when an organization wanted electricity they would build their own generator. Over time electricity utilities were created and they provided electricity on a user pays basis in a very cost effective way to multiple users.

Similarly it is possible to set up a private Video Utility where multiple users in an organization can all access the same video but in different ways. Each one can only receive the information that they are authorized to see and each one is charged based on their usage. This concept is patented by iOmniscient and is already available today.

The Video Utility concept can reduce the costs for a large organization by an order of magnitude. Of course it would require leadership at the top of the organization to force different stakeholders to co-operate but the benefits can be enormous.

3.9 Implementation and Maintenance.

Software implementation and maintenance can be an expensive exercise especially if an engineer has to be sent on site to a remote location every time some activity is required on the software.

For this reason all iOmniscient systems are designed to be remotely implemented, configured, diagnosed and maintained. Application based encryption and inbuilt cyber security protections ensure that the system cannot be hacked during such sessions.

Hardware implementation and maintenance costs are reduced because fewer cameras need to be implemented

Camera maintenance costs can also be reduced using the iOmniscient HealthCheck System. At any time in a large network of cameras about 20% of the cameras are not operational for some reason. This does not mean that the camera is necessary disconnected which is easy to detect. It may not be operational because it has moved from its required position or it has lost focus due to vibrations in the environment or even that it is working perfectly but nothing is visible due to fog or rain or snow. The user usually does not know which cameras are affected until the operator observes that something is wrong. The IQ-HealthCheck system can detect such anomalies and enable the user to take counter measures proactively.

The cost of an in-operational system can be huge as many users have experienced. For instance, paintings by five masters were stolen from the Museum of Modern Art in Paris. This was undetected because while the CCTV system was working the camera lens had been coated with dust during a construction event and they were effectively not able to see what they were supposed to see – something that is easily detected by the IQ-HealthCheck system.

3.10 Response Time

Long response time carry consequential costs to users. In a typical city surveillance system long response times for accident type events result in traffic congestion and the time it takes for rescue workers to arrive can determine whether a victim lives or dies.

iOmniscient's patented Automated Response system has been shown to reduce response times by 80% in City Surveillance applications. Where it used to take 25 minutes to respond to a street incident the response time was reduced to under 5 minutes.

- **Automated Response Systems ----------** ***http://youtu.be/n_2cGKtLqmE***

The cost of time is often not considered in Cost Benefit calculations as these are often performed by the departments incurring the costs (e.g. the IT department) whereas the benefit (such as a faster response time) accrues to the end user or even to the organization's customer.

A comprehensive cost benefit analysis would take these benefits into account as they are important to the organization as a whole.

4. Key to a Successful Implementation

Very few intelligent surveillance systems are deemed to have been successful. iOmniscient has developed a strong reputation for successful projects by understanding the 5 key steps to implementing successful projects. These are:

1. Understand the objectives of the system
2. Determine the appropriate software for the requirement
3. Select the cameras
4. Design the network and the systems needed
5. Implement

In unsuccessful projects we find that users often start at Step 3. They then complete steps 4 and 5 and then go back to step 2 (sometimes without ever doing step 1). If the correct process is followed it maximizes the chance that the organization can recognize the full benefit of implementing an intelligent system.

5. Tools to calculate Cost Savings

Users need to be able to quantify the savings they can make. iOmniscient offers tools that can enable users to calculate the savings that they can achieve through using iOmniscient systems.

6. Summary

With good planning, an organization can implement advanced Intelligent Video Surveillance Systems without increasing the overall cost of the system. Often the final cost of such a smart system will be lower than the cost of using traditional, passive, unintelligent technologies.

The summary below shows how the various costs of the system can be minimized

	Cost	How the Cost is Minimized
1	Type of Cameras	Cameras with lower resolution
32	Number of Cameras	Fewer cameras through doing analytics at lower resolution
3	Storage	Through Smart Compression
4	Network	Through Smart Compression
5	Computers and GPUs	Use of combinational AI, processing at lower resolution and frame rate and mo need for GPUs
6	Software Analytics	NAMS
7	Labour	Through building complex Use Cases
8	Control Room	IQ Mobile – control room on a phone and Automated Response Systems
9	Cost of Duplication	Video Utility
10	Installation	Fewer cameras means less installation
11	Maintenance	Fewer cameras means less maintenance, remote maintenance and real time Healthchecks
12	Response Time	Automated Response Systems

NOTE: Virtually all of the methods described for reducing costs in this paper are internationally patented by iOmniscient.