Is it Time to Switch to H.264?

The H.264 encoding standard promises better, clearer, and smoother video images over a wide range of applications including cell phones, video surveillance cameras, and HDTV broadcasts. In the video surveillance industry, the past year has seen a number of H.264 cameras hit the market, especially from Axis, Sony, and Arecont Vision. With increasing availability of these H.264 cameras, Milestone Systems was the first major Video Management Software to offer H.264 support in XProtect Enterprise version 6.5 (January, 20008).

The main advantage of moving to H.264 is that the cameras require much less network bandwidth to deliver the same or better quality than competing standards such as MJPEG and MPEG4. Less bandwidth means less disk storage, thereby decreasing the cost of a system and increasing archival times. If you have an opportunity to test one of these cameras, you will find they produce stunningly clear images and do indeed require much less bandwidth.

There's always a catch. In this case the H.264 encoder has to work much harder on encoding the video stream. Camera manufacturers have had to beef up their embedded processors to work with H.264. The concern is whether processing power needs to improve system-wide in order to support H.264. Will servers need more powerful CPU's? How will the client workstations be affected?

To try and understand the implications at a systems level, I conducted a series of benchmarks designed to compare and contrast system design using H.264, MJPEG, and MPEG4.

Camera Image Size Measurements

The first series of experiments involved capturing static images and calculating an average frame size. Testing was performed on four different multi-codec cameras from Axis, Sony, and Arecont Vision. Each camera supports both MJPEG and H.264; only the Sony cameras have MPEG4 as well. The Milestone Administrator program includes a Preview function where a sequence of images are retrieved and the average size per frame is calculated.

The data shows pretty much what was expected. H.264 has about a 7:1 advantage over MJPEG at the highest quality settings. At the lowest bandwidth settings, the curves do come together but H.264 continues to have smaller packets and the image quality is much better. MPEG4 also does quite well in producing small packets, but there are lots of video artifacts at the lower bandwidth settings.

How would this affect system design? If you have a collection of 16 cameras on a network, all set for 6 frames per second and mid-level quality*, the MJPEG cameras will produce 59.3 GB of data per day, while the H.264 cameras only need 3.7 GB. Stated another way, a system with 16 MJPEG cameras can store about 2 weeks of video on a 1 Terabyte drive. The H.264 cameras can record for 2 months. If the cameras are sharing a network with other business data, you will be much better off using H.264 cameras. They require only .38 Mbps of network bandwidth, compared to 1.5 Mbps for the each MJPEG camera.

* Based on 50% motion, 12 hours per day, image resolution 640x480, 6 fps, mid-level compression.

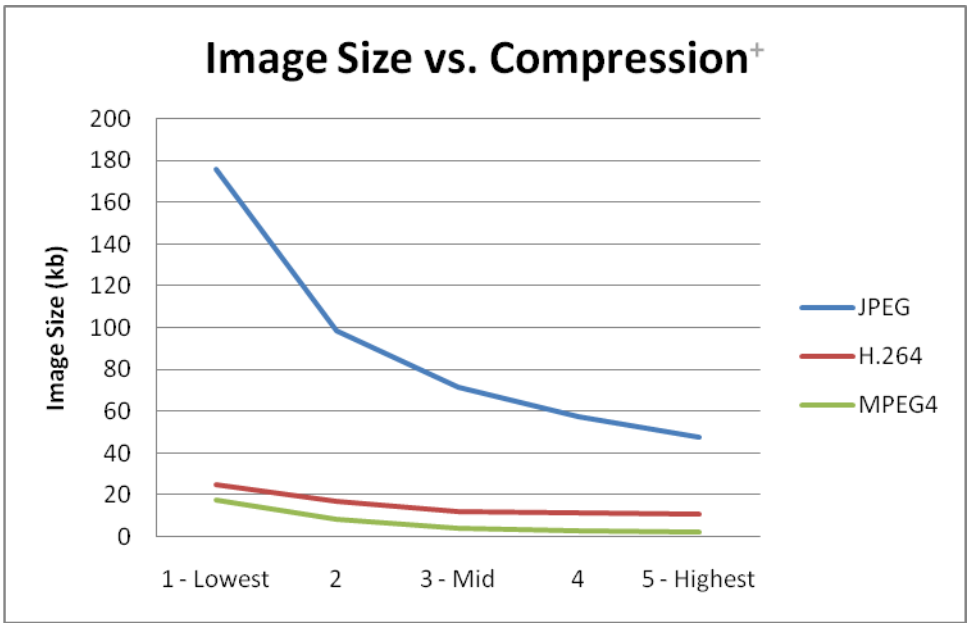

+ Image size data was averaged from the four cameras listed below. Only the Sony cameras support MPEG4. In the charts below, 'Dynamic Range' refers to the ratio of compression from maximum quality to minimum, and 'MJPEG:H.264' is a compression ratio is based on image sizes at highest and lowest quality.

Image Size Test Data		Dynamic Range (max:min)			Compression Ratio MJPEG:H.264		Frame Size at Max. Quality (KB)	
Camera	Resolution	MJPEG	H.264	MPEG4	Highest Quality	Lowest Quality	MJPEG	H.264
Sony CS-50	640x480	4.7	5.0	5.6	7.4	7.8	88.6	12
Sony RZ-550N	640x480	5.2	7.1	9.7	4.6	6.3	94.9	20.7
Axis Q1755	1440x1080	10.3	38.3	na	5.7	21.3	174.7	30.6
Arecont 2105N	1600x1200	2.5	1.0	na	9.3	3.7	346	37.2

Server Side Testing

The next series of tests were conducted on the server to measure the impact of codec selection on CPU performance, network and disk usage. The server is based on an Intel 2.0 GHz Quad Core Xeon processor with 2 GB of RAM and a 1 Terabyte SATA disk. The same four cameras were attached to the network, and performance was measured using the Microsoft Vista version of the 'PERFMON' program. The video management system is Milestone XProtect Enterprise, version 6.5f.

There are two questions to be answered:

- How does system performance vary based on the degree of compression for each codec?
- How does the H.264 codec compare to MJPEG under similar conditions?

Using the PERFMON traces, you will quickly see that MJPEG at high quality produces a heavy load on the system based on moving around massive amounts of data. The CPU is not working very hard because it does not need to decode the data – only store it. Nonetheless, this configuration will quickly become I/O bound if nothing is done to reduce bandwidth requirements.

By comparison, the H.264 test shows a lightly loaded system. It only has one fourth of the data load, so the server can handle many more cameras using H.264. Another conclusion we can reach from this data is that CPU performance requirements on the server side are slightly less for the H.264 codec. At maximum quality settings the MJPEG codec uses more CPU performance simply because there is more data to be moved around.

Server Side Performance Measurements

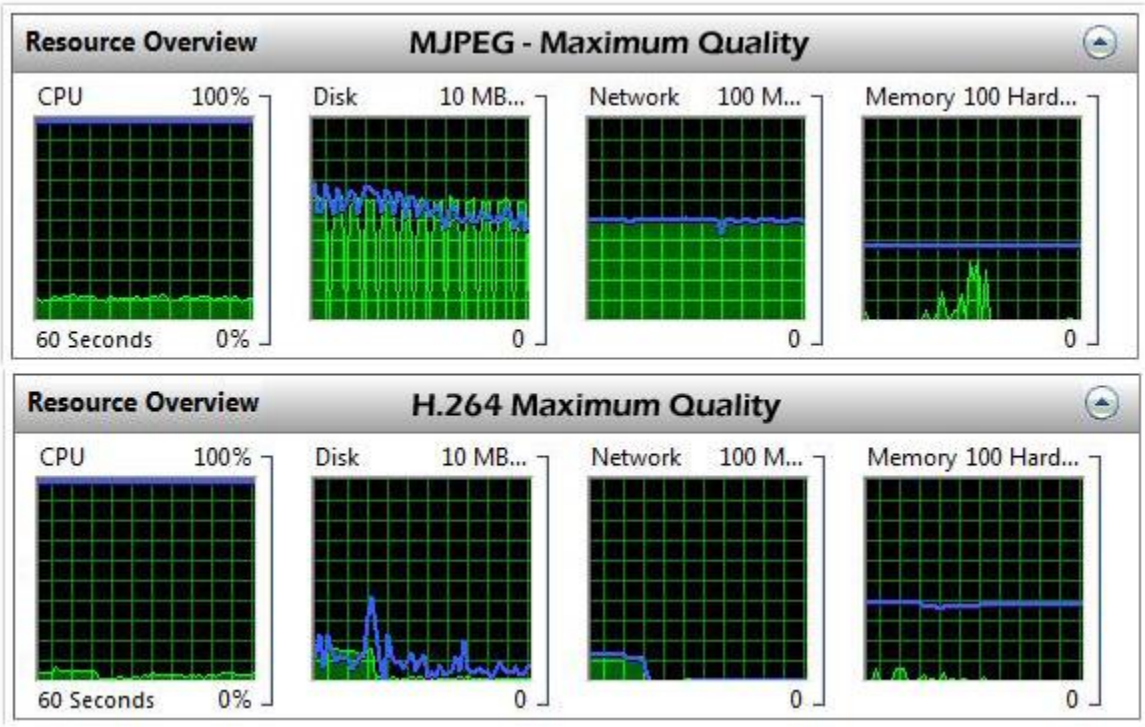

Client Side Testing

The expectation on the Client side was that the burden of decoding the H.324 packets would increase the CPU requirements for workstations. The data did not bear that out. The difference turns out to be minimal at high quality, and tips in favor of H.264 at lower quality.

The decoding work load is offset by having less data to handle. The ratio of network bandwidth for MJPEG vs. H.324 at highest quality is about 7:1. Disk performance is not a factor on the client side since the data is only viewed, not saved. All things considered, H.264 places no extraordinary burden on the client side when compared to MJPEG.

This finding means that any workstation or laptop which is currently capable of displaying MJPEG streams will do just as well (perhaps better) in handling H.264 video. In addition, the Milestone Smart Client application allows users to fine-tune the image quality and frame rate in the viewing window so that video can be viewed even from a low bandwidth wide area network connection.

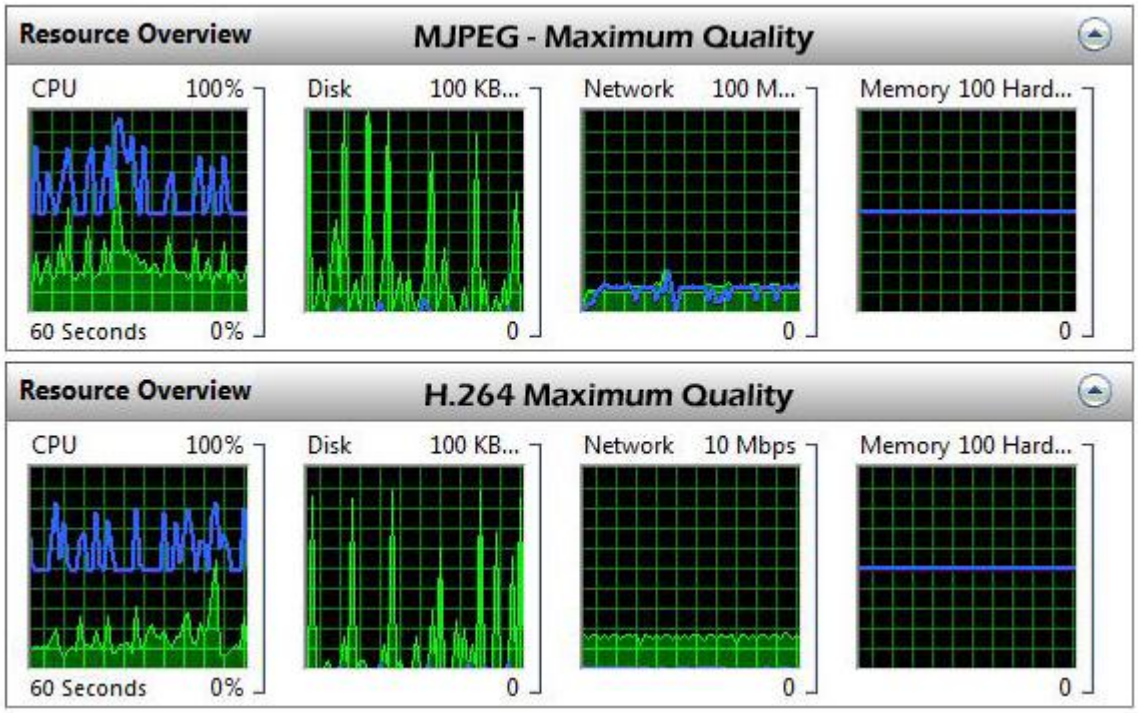

Note: The scale changes on 10:100 on the Network trace

Summary

Improved camera availability and software support from Milestone Systems means that there are no intrinsic obstacles to designing and deploying H.264 video surveillance systems. The benefits of using H.264 are:

- Lower network bandwidth requirements
- Increased archive retention due to lower storage requirements
- Higher camera counts per server
- Clearer video than MPEG4 and MJPEG at lower data rates

You don't have to dive into H.264 all at once. The Milestone XProtect software allows for multiple codecs to be used simultaneously so you can mix and match with other cameras. Existing Servers will run better because they are no longer I/O bound. You don't have to upgrade client workstations either – any reasonably current PC should be fine.

More H.264 cameras are on the way and it is clear that this codec will become mainstream in a very short time. This report finds no penalty to getting started early, and lots of reasons for why your next video surveillance design should include H.264 devices in it.