Keystone Aerial Surveys, Inc.

Background

This case study details the procedures, tools and necessity of a thorough examination of a firm's entire processing infrastructure. When Keystone Aerial Surveys, Inc. needed to upgrade its processing hardware Keystone first set out to determine how the software it used interacted with its enterprise and workstation hardware. As a primary data acquisition company using Microsoft\Vexcel UltraCams, Keystone uses the UltraMap suite to process the hundreds of thousands of raw images it collects every year.

Many primary acquisition vendors creating downstream products can benefit from the implementation of the test structure, tools and techniques demonstrated.

Software

UltraMap is a software package produced by Microsoft\Vexcel of Austria to support the UltraCam line of large format imagery sensors. The software has many modules, including Imagery processing, Aerial Triangulation (AT), Radiometry, Digital Surface Model (DSM) and Orthophoto generation. For this study, only the Raw Data Center (RDC), AT and Radiometry modules where considered. The UltraMap software is based on a distributed processing model where one controller computer doles out tasks to worker nodes.

Hardware

Keystone's existing hardware infrastructure consisted of a NetApp Network Attached Storage device using Serial Advanced Technology Attachment (SATA) hard disks with 10 GB Ethernet and Fiber Channel connectivity. The existing servers were a combination of Intel and AMD processing cores running on various ages of Dell servers with at least the minimum of recommended RAM (though at various speeds).

Workflow

UltraCam sensors use multiple calibrated cameras to generate its final, full frame images. The raw data recorded by the sensor consists of multiple pieces of imagery and metadata that must be assembled for use in geospatial products. This state of the imagery is called Level0 (L0). The transfer of the L0 data from the sensor to network storage is handled by the RDC module within UltraMap. The initial phase of post processing primarily involves the seaming of the panchromatic images into on single virtual image. RDC is again used to do this in a distributed fashion. The controller server receives the job and instructs the processing nodes to retrieve the L0 imagery from the storage location, process it to the next level and store it at a shared location. The imagery and associated metadata are called Level02 (L2) at this stage in the process.

The L2 imagery is then run through an initial automatic point extraction using the AT module. Finally, color balance across the project is performed using the Radiometry module and a delivery data, or Level03 (L3), job is run. This is distributed among the processing nodes in the same manner as the L0 to L2 processing.

Test Results For Processing Server

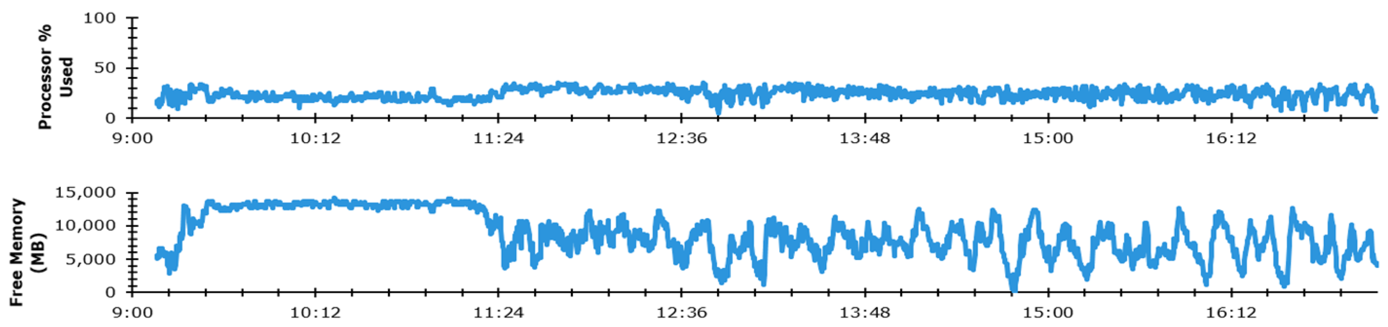

Test Results For New Processing Server

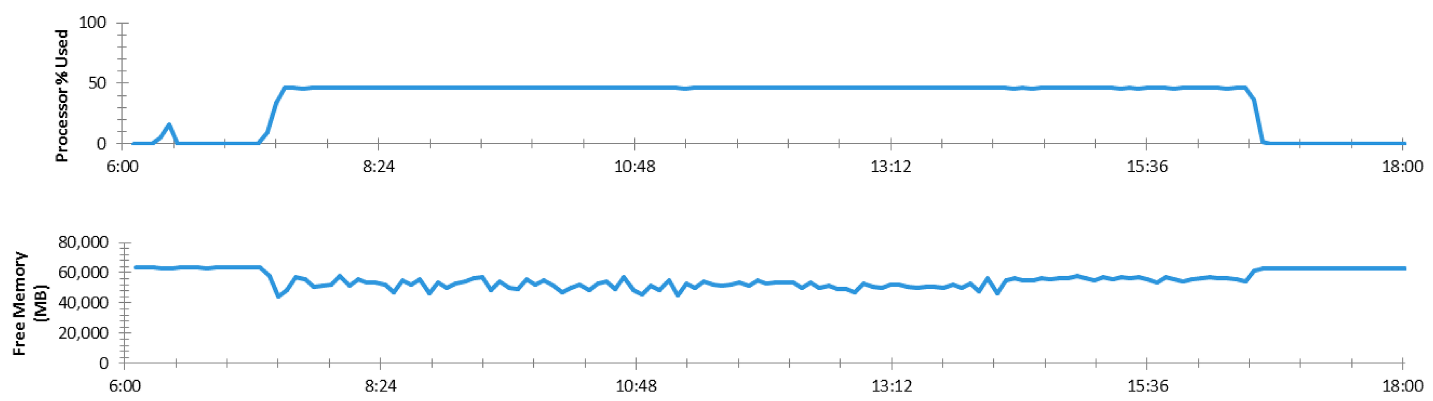

Latency Test Results for NAS

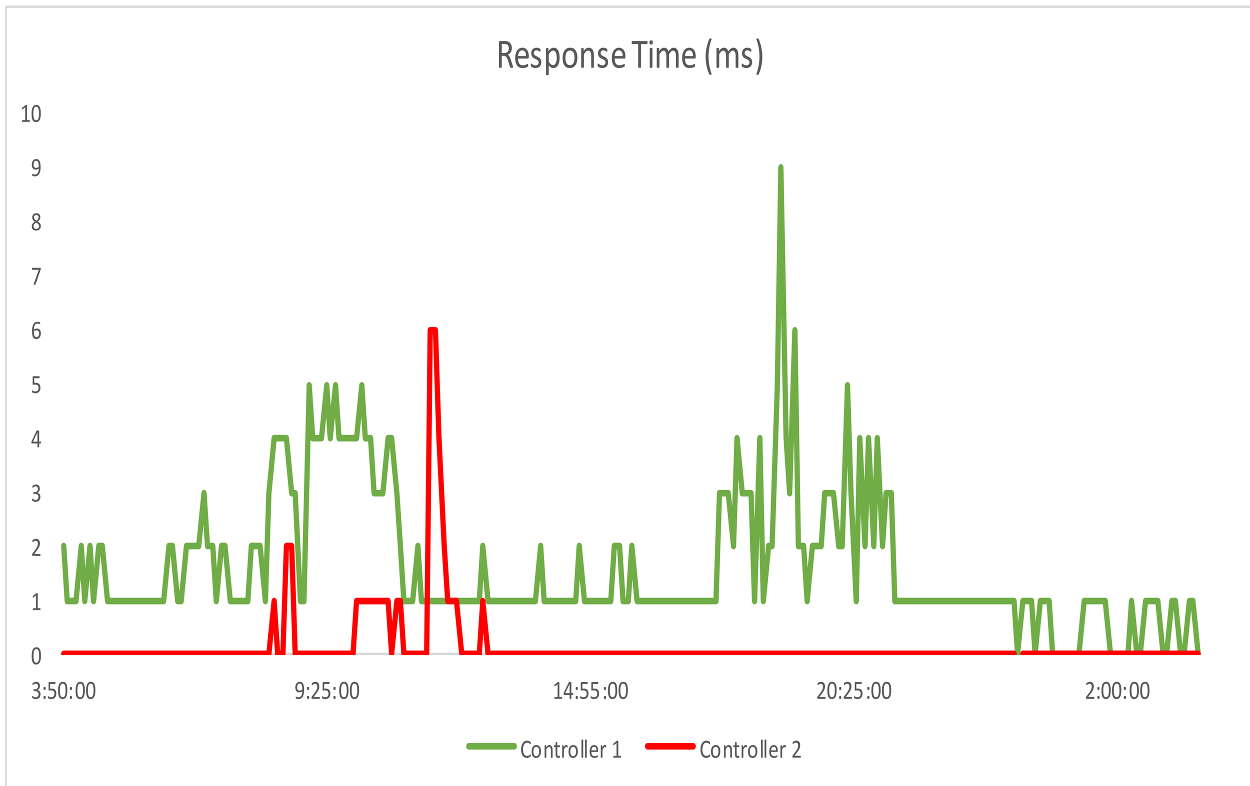

Latency Test Results for New NAS

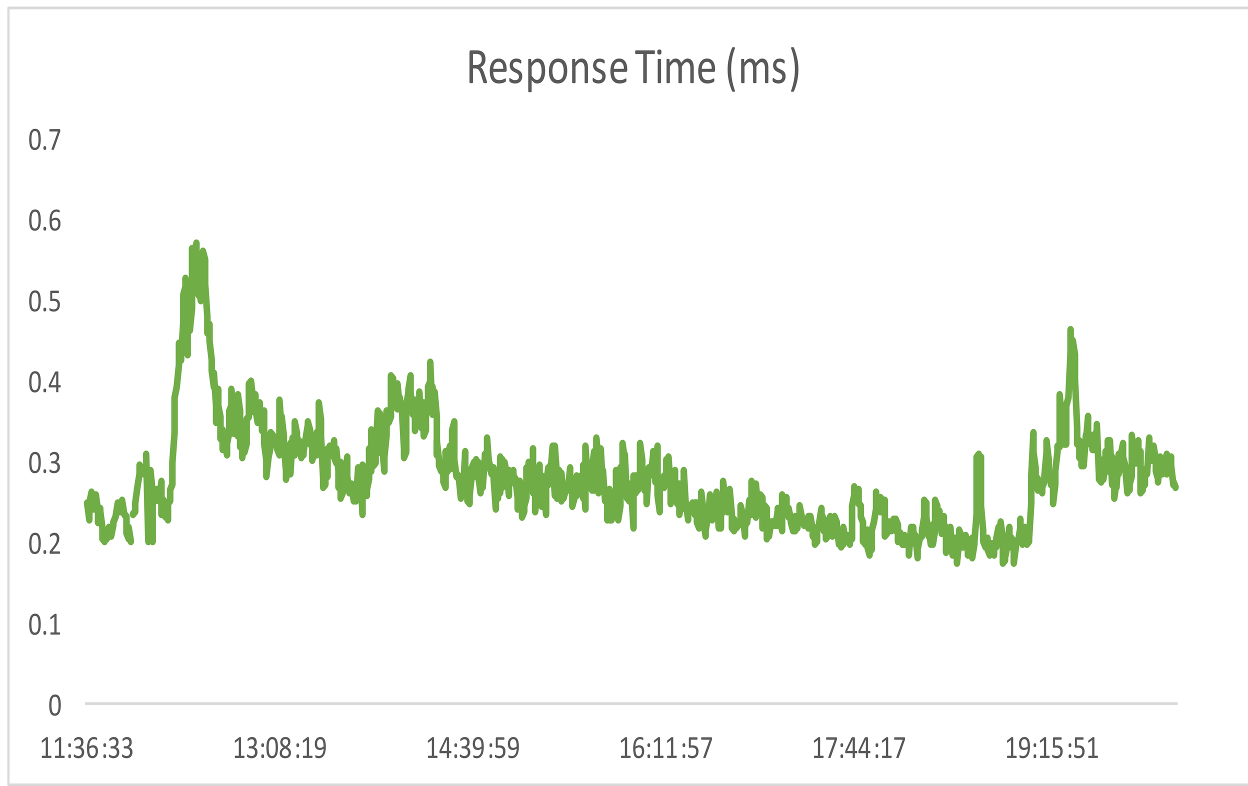

Performance Test Results for NAS

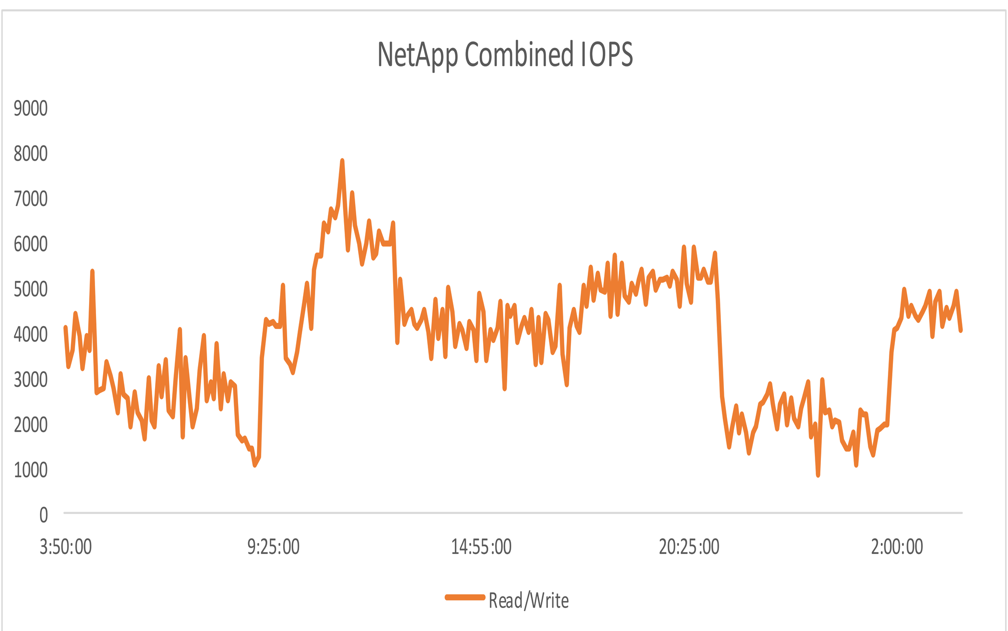

Performance Test Results for New NAS

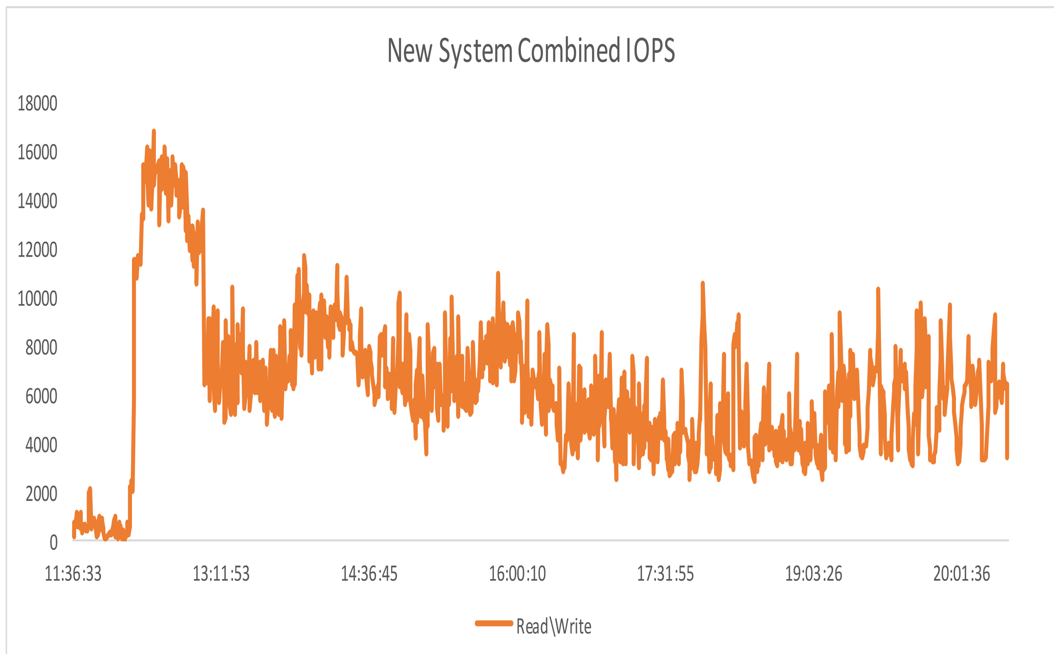

Methods

The system test was designed to mimic a real workflow event where the entire processing architecture and storage architecture are stressed. This included massive movement of data during all stages of the usage of UltraMap software (RDC, AT and Radiometry). A small footprint software package from EMC Corporation was installed on each processing server that monitored the system Input/Output (I/O), RAM, processor usage, network traffic, etc. for an entire 24 hour period. Using software designed for the storage system, it was also monitored for a complete 24 hour period. The results for throughput, latency, etc. on the existing system were established and used for bench marks and insight into a possible replacement. The test was then simulated on a prospective system and, finally, the full system test was repeated on the installed solution after the purchase.

Conclusions

- The initial assumption that increased CPU core count would generate increased throughput proved false. CPU processor speed proved a more vital factor in processing speed.
- Unlike other business workflow scenarios SSD usage on large storage systems does not produce good results for geospatial imaging workflows.
- The use of solid state disk in the processing servers produced significant speed increases. The results show the reduction in disk read/write speeds of 90% and 95% respectively. Particularly interesting are the reduced average queue length by 80% to 0.35.
- The storage area network also showed remarkable improvements over the previous system with a peak IOPS of nearly 17,000 while maintaining response times of half a millisecond or less compared to the previous system's maximum IOPS of 5000 and response times consistently above 1 millisecond and peaking at 8.
- Processing gains individual image processing range from 11% to 64%. On a 60 core distributed system, 1000 images will be processed between 60 and 100 minutes faster than with the previous system while allowing other much needed processes to run unaffected.

Image Processing Time Comparison (L2)

Level 02 Processing	Old Storage System (In Minutes				New Storage System (In Minutes)				
Sensor	"A" Group	"C" Group	"D" Group	Weighted Ave	New Group	"C" Group	"D" Group	Weighted Ave	Gain %
Eagle	16.6	15.8	10.8	15.8	9.3	14.4	10.7	11.0	30%
Falcon	11.5	11.4	10.3	11.3	6.6	8.2	7.9	7.2	36%
UCX	11.1	10.3	7.8	10.5	4.5	5.2	4.8	4.7	55%
UCO	21.6	19.1	27.1	21.3	16.4	21.4	26.5	18.9	11%

Image Processing Time Comparison (L3)

Level 03 Processing	Old Storage System (In Minutes)				New Storage System (In Minutes)				
Sensor	"A" Group	"C" Group	"D" Group	Weighted Ave	New Group	"C" Group	"D" Group	Weighted Ave	Gain %
Eagle	12.0	7.1	6.3	9.9	5.7	6.5	6.3	6.0	39%
Falcon	8.3	4.8	4.5	6.8	4.1	4.6	4.6	4.3	37%
UCX	11.5	6.9	8.5	9.7	3.3	4.0	3.5	3.5	64%

Latency Comparison of Server Disk

System	Average Read Response Time (ms)	95th Percentile Read Response Time (ms)	Average Write Response Time (ms)	95th Percentile Write Response Time (ms)	Average Queue Length	95th Percentile Queue Length
Original System	2.40	7.51	36.32	98.92	1.81	8
New System	0.35	1.10	3.15	5.10	0.35	3

Follow the QR Code to view the entire pdf document

Contact:
www.KASurveys.com
info@kasurveys.com
215-677-3119

KEYSTONE AERIAL SURVEYS