# Quantifying Hardware Selection in an EnCase v7 Environment

## Introduction and Background

The purpose of this analysis is to evaluate the relative effectiveness of individual hardware component selection in the EnCase v7 environment. While it is useful to document the individual hardware components which result in maximum performance, it is also important to identify those components which provide the best value. This effort is part of an ongoing commitment by Digital Intelligence to assist customers in making educated choices when selecting individual components for their forensic workstations.

## Approach

Four basic steps were used to evaluate the application's resource requirements.

Step 1 (Establish Test Environment): A suite of tests was developed for the application. These tests were intended to represent the demands of a typical forensic examination. These tests were then automated in order to provide accurate and repeatable recording of results.

Step 2 (I/O Channel Evaluation): The automated test suite was then used to determine the basic configuration of the I/O channels. As a starting point, the application manufacturer recommends up to 5 I/O channels:

1.) Operating System
2.) Casework
3.) Cache
4.) NSRL KFF Data
5.) Evidence

A demonstrated ability to combine two or more of these I/O channels could easily result in a less expensive and more manageable configuration. Evaluation of the I/O channel requirements would be essential in determining an optimal I/O configuration. A baseline system configuration can then be established using this information.

Step 3 (Resource Evaluation): Using the baseline configuration, individual components were identified for modification. These components consist of the general hardware options available for system configuration. By limiting baseline modifications to individual components, the relative importance of the associated resources can be evaluated.

Step 4 (Potential System Configurations): The final step was to identify and test several cumulative changes to the baseline configuration. The value of individual resource modifications, as identified in Step 3, would be essential in determining the hardware combinations to be tested. These hardware combinations would be good candidates for effective workstation configurations.

## Methodology

A test disk was created with the following attributes:

1.) Contains data which is generalized and varied.
2.) Contains data which is representative of what might be encountered in a typical examination.
3.) Contains data which is significant enough to result in a meaningful processing time.

A test suite was developed with the following attributes:

- File Verification (E01 format option only)

- Pre-Processing
  - Recover Folders
  - Protected File Analysis
  - Thumbnail Creation
  - Hash Analysis
  - Expand Compound Files
  - Find Email
  - Find Internet Artifacts

- Indexing
  - Index Text and Metadata

- File Carving
  - Modules – File Carver

A scripting tool was selected and implemented in order to automate the test suite. This tool not only allowed for the automation of testing but also ensured that the individual test times were accurately recorded. AutoIT™ was the tool selected to perform this task (http://www.autoitscript.com/autoit3/).

Before beginning each test, an imaging tool (Ghost™) was used to restore the O/S disk to its baseline state. The Cache Disk was formatted. This would ensure that all residual data from the previous test would be eliminated including any file-system fragmentation or file relocation.

The test suite would be run utilizing both compressed (E01) and un-compressed (DD) evidence images in order to evaluate the relative performance impact of each option.

## Step 1 (Establish Test Environment)

A Windows7(x64) workstation was installed and utilized to create the test disk. Files from the public domain Enron dataset were used to provide email and attachment content. A number of messaging programs were installed and used to simulate “chat” with other users. Additional emails were created and sent with

both browser-based and locally installed clients (Outlook). Web browsing was performed. All of these activities were intended to generate content similar to what might be encountered during a typical investigation. The resulting disk images (created with Tableau Imager) consisted of approximately 240 GB of uncompressed (DD) data and 60GB of compressed (E01) data.

## Step 2 (I/O Channel Evaluation)

The initial baseline test system for this analysis was a core i7 system with 16GB memory. Five identical 7200 RPM SATA drives were attached and a case was configured with 5 separate channels as follows:

**Separate Channel Configuration**

| Drive Letter | Contents |
|---|---|
| C: | Win7 x64 Operating System |
| D: | Evidence |
| J: | Cache |
| K: | Case |
| G: | NSRL KFF Database |

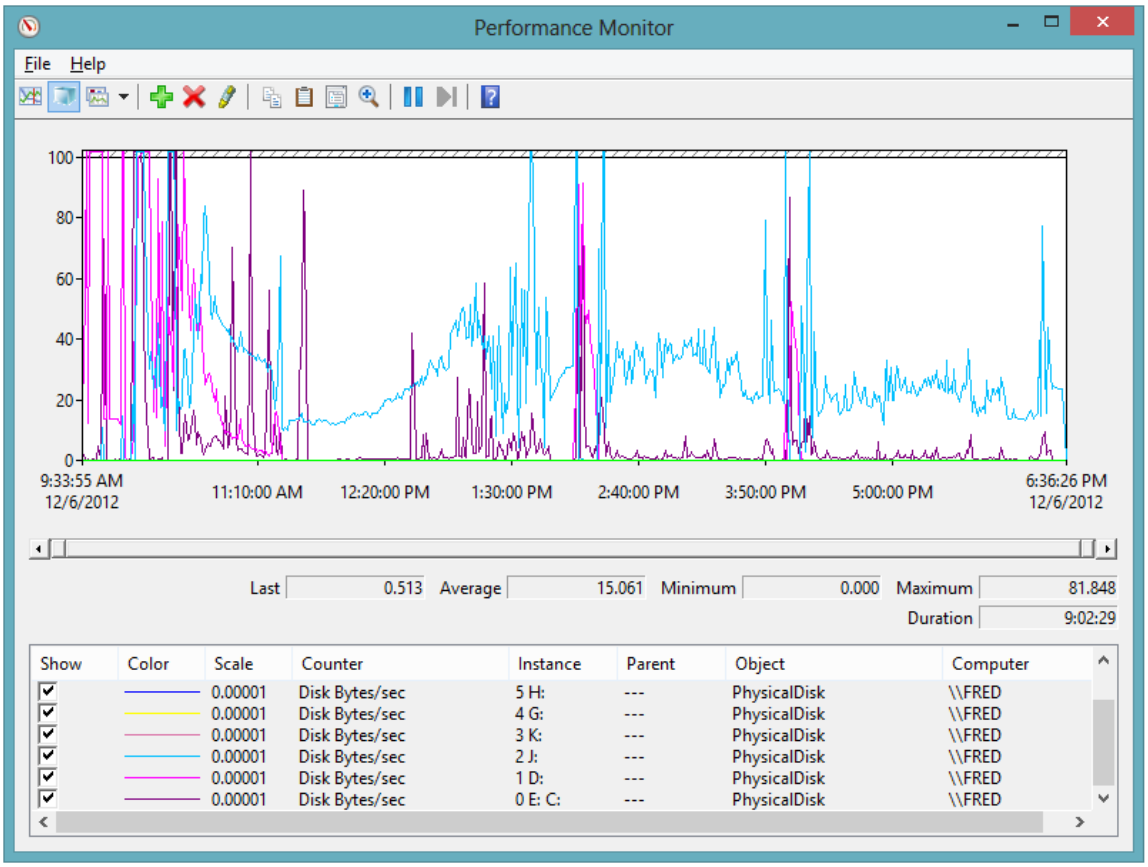


Using Windows Perfmon, an analysis of the disk activity indicated that both the O/S and NSRL KFF Database channels, as well as the Evidence and Case channels, could be combined with negligible performance impact.

Tests were subsequently run in the Consolidated Channel Configuration to validate this finding. The results showed that the five original channels could be reduced to three channels while only incurring a 3% loss in performance. Reducing the number of I/O channels results in a reduction in complexity, a reduction in cost, and the ability to combine case specific information on a single drive. The resulting configuration is shown below:

**Consolidated Channel Configuration**

| Drive Letter | Contents |
|---|---|
| C: | Win7 x64 Operating System and NSRL KFF Database files |
| D: | Evidence and Case |
| H: | Cache |

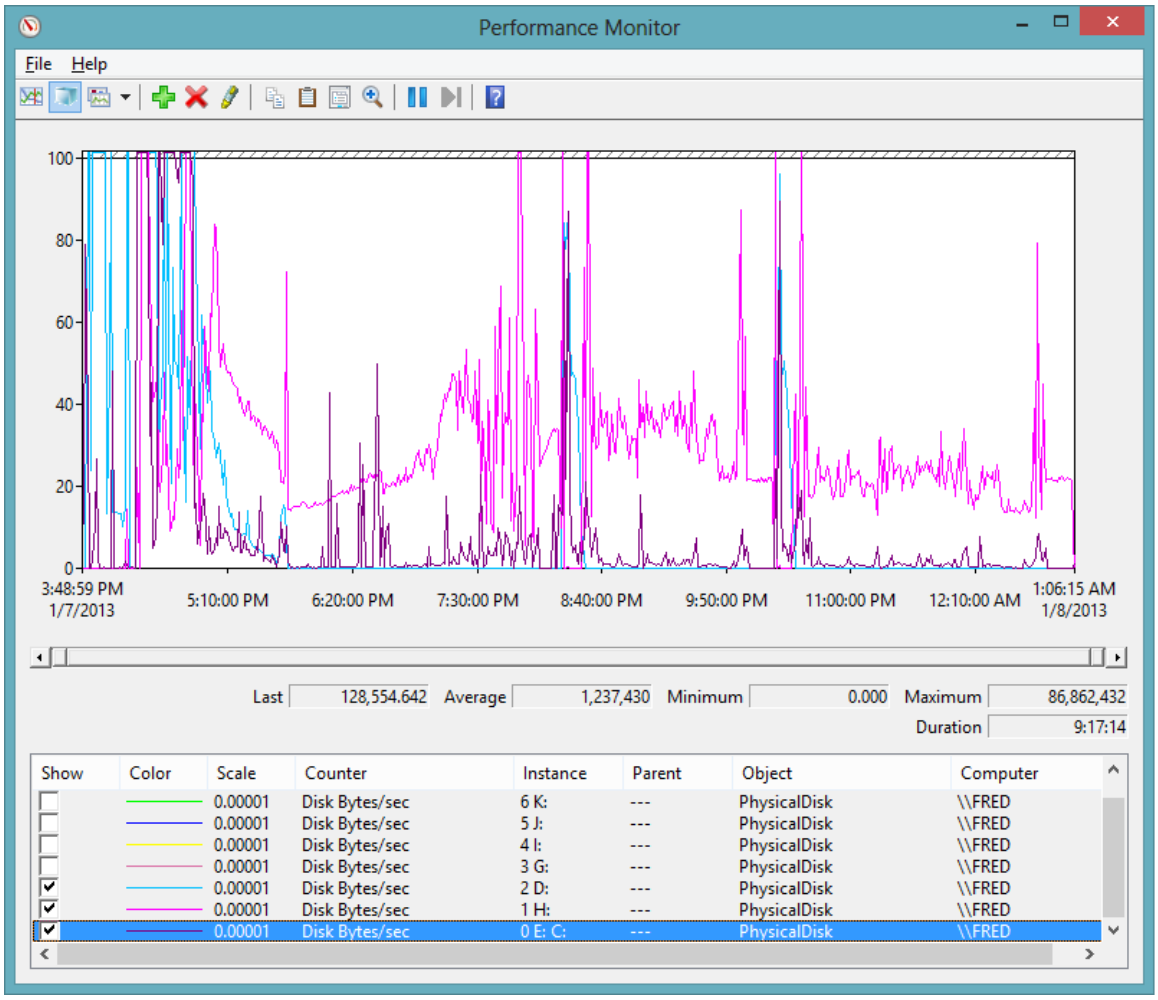

**Step 3 (Resource Evaluation)**

The results of Step 2 indicated that only three I/O channels would be needed. This helped define the testing matrix for resource evaluation. The following resources were to be evaluated:

- CPU/Processor
- Memory
- O/S & and NSRL KFF Database Drive
- Cache Drive
- Evidence and Case Drive

In order to effectively compare two different architectures, a baseline was established for both an Intel i7 and an Intel Dual-Xeon system as follows:

| Component | I7 Baseline | Dual-Xeon Baseline |
|---|---|---|
| Processor | I7-3820<br>3.6 Ghz<br>Quad Core<br>10MB Cache | E5-2609<br>2.4 Ghz<br>Quad Core (8 cores total)<br>10MB Cache |
| Chipset | X79 | C602 |
| Memory | 16 GB | 16 GB |
| O/S & KFF Drive | 7200 RPM SATA | 7200 RPM SATA |
| Cache Drive | 7200 RPM SATA | 7200 RPM SATA |
| Evidence & Case Drive | 7200 RPM SATA | 7200 RPM SATA |

- 16 GB Memory = DDR3-1600
- 7200 RPM SATA = WD2002FAEX 64MB Cache 2TB

Two systems were built utilizing the baseline configurations in the table above. The performance test was run on each and the results recorded. The entire suite of tests would then be run, modifying a single component, in order to quantify the impact of the associated resource on overall system performance. Both compressed (E01) and un-compressed (DD) evidence was processed.

Each system was installed with Microsoft Windows 7 Ultimate (64 bit version) and all patches applied. The Windows Firewall, Search Service, Scheduled Defragmentation, and Windows Update were turned off or disabled. The Auto-IT (scripting environment) was installed and configured. EnCase Version 7 (7.05.02) was installed configured per the manufacturer's instructions.

The following tables identify the associated hardware permutations and the resulting impact on system performance:

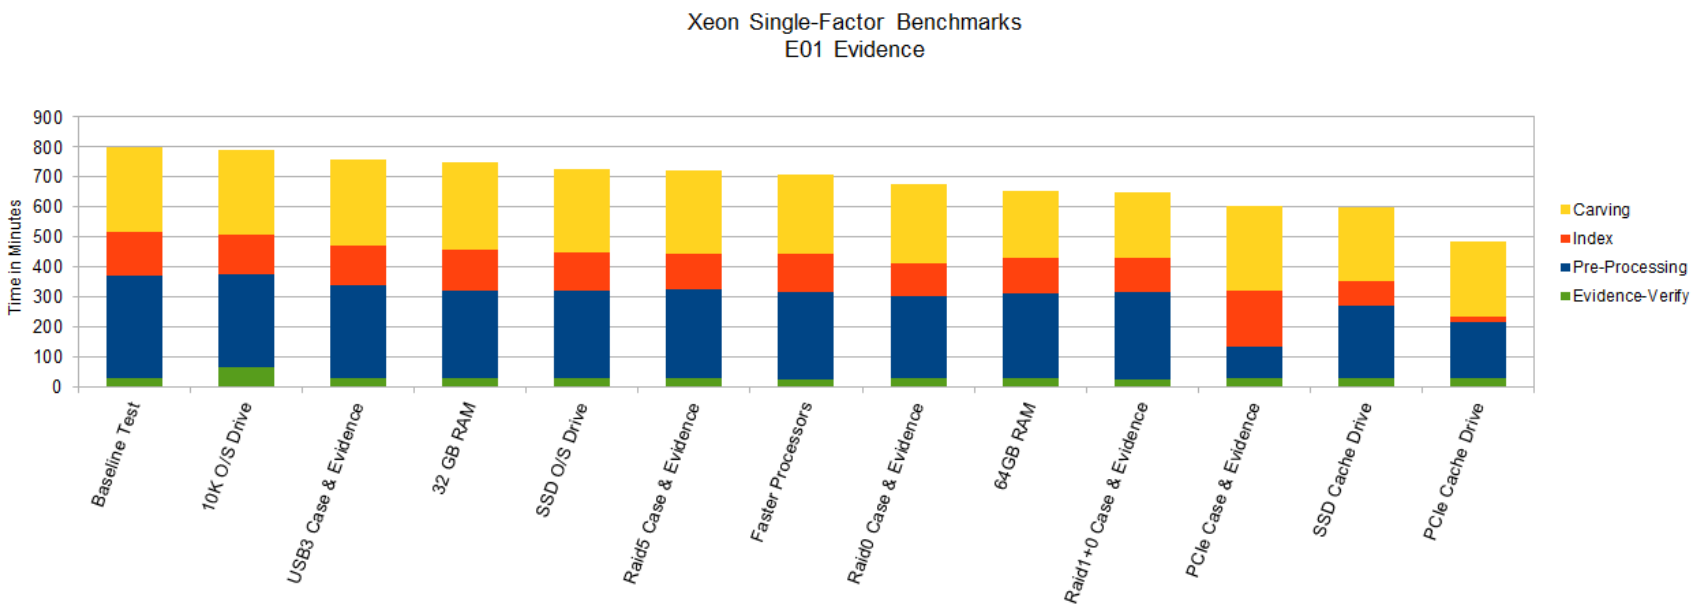


| Description | CPU | RAM | OSDrive & KFF Drive | Case and Evidence Drive | Primary Cache Drive | Evidence-Verify | Pre-Processing | Index | Carving | Total | % Change |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Baseline Test | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 31.0957 | 342.2671 | 143.0776 | 281.4888 | 797.9292 | 0% |
| 10K O/S Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 10k Raptor | 7200 RPM SATA | 7200 RPM SATA | 68.4855 | 305.7112 | 135.428 | 279.3033 | 788.928 | 1% |
| USB3 Case & Evidence | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | USB3 7200 RPM SATA | 7200 RPM SATA | 31.1344 | 306.8155 | 134.2953 | 288.1202 | 760.3654 | 5% |
| 32 GB RAM | 2-E5-2609@2.4Ghz - Quad - 10MB | 32 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 30.0080 | 291.4356 | 135.4269 | 290.2875 | 747.158 | 6% |
| SSD O/S Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | Vertex 4 SSD | 7200 RPM SATA | 7200 RPM SATA | 30.0081 | 290.3419 | 128.8049 | 277.1095 | 726.2644 | 9% |
| Raid5 Case & Evidence | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Areca 5X7200 RPM SATA - RAID5 | 7200 RPM SATA | 28.9093 | 296.9258 | 116.7185 | 279.3024 | 721.856 | 10% |
| Faster Processors | 2-E5-2630@2.3Ghz - Hex - 15MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 27.8369 | 288.1394 | 128.7988 | 263.919 | 708.6941 | 11% |
| Raid0 Case & Evidence | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Areca 5X7200 RPM SATA – RAID0 | 7200 RPM SATA | 30.0000 | 274.9585 | 106.8319 | 262.8553 | 674.6457 | 15% |
| 64GB RAM | 2-E5-2609@2.4Ghz - Quad - 10MB | 64 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 30.0362 | 282.6511 | 118.9176 | 223.2783 | 654.8832 | 18% |
| Raid1+0 Case & Evidence | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Areca 5X7200 RPM SATA - RAID1+0 | 7200 RPM SATA | 27.8427 | 290.3443 | 112.3254 | 218.8845 | 649.3969 | 19% |
| PCIe Case & Evidence | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | PCIe SSD | 7200 RPM SATA | 30.0363 | 103.5876 | 188.1594 | 281.5313 | 603.3146 | 24% |
| SSD Cache Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | Vertex 4 SSD | 31.1059 | 242.0051 | 79.3987 | 246.3409 | 598.8506 | 25% |
| PCIe Cache Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | PCIe SSD | 31.1331 | 183.7824 | 22.2439 | 248.5445 | 485.7039 | 39% |

Other Components were:

- Vertex 4 SSD = OCZ-VERT Ex4 1.4 SATA
- RAID configurations = RAID0, 1, and 5 using an ARECA ARC-1882ix-12 Raid controller and 5-WD2002FAEX 7200 RPM SATA 64MB Cache 2TB drives
- PCIe SSD = OCZ-REVO3X2 PCIe

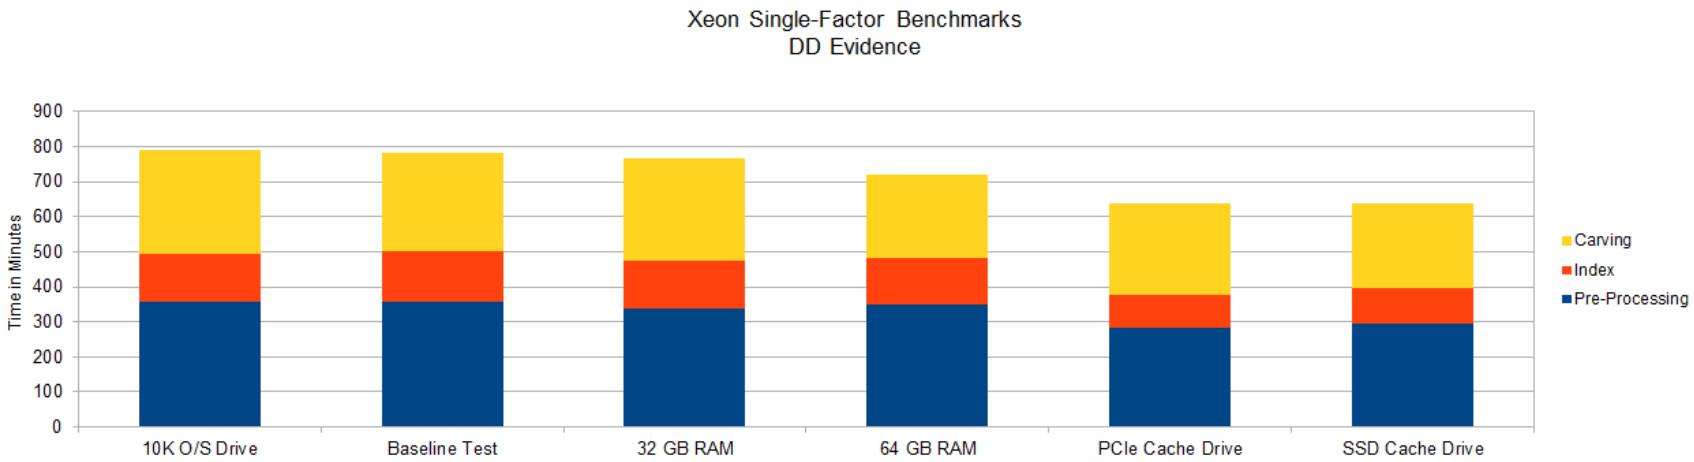


| Description | CPU | RAM | OSDrive & KFF Drive | Case and Evidence Drive | Primary Cache Drive | Pre-Processing | Index | Carving | Total | % Change |
|---|---|---|---|---|---|---|---|---|---|---|
| Baseline Test | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 357.3099 | 143.0779 | 282.5932 | 782.981 | 0% |
| 10K O/S Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 10k Raptor | 7200 RPM SATA | 7200 RPM SATA | 356.2084 | 138.6804 | 294.6946 | 789.5834 | -1% |
| 32 GB RAM | 2-E5-2609@2.4Ghz - Quad - 10MB | 32 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 338.6551 | 135.3854 | 291.3783 | 765.4188 | 2% |
| 64 GB RAM | 2-E5-2609@2.4Ghz - Quad - 10MB | 64 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 349.6551 | 133.197 | 236.4563 | 719.3084 | 8% |
| PCIe Cache Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | PCIe SSD | 282.6276 | 95.8387 | 260.6178 | 639.0841 | 18% |
| SSD Cache Drive | 2-E5-2609@2.4Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | Vertex 4 SSD | 295.7944 | 100.2338 | 243.0459 | 639.0741 | 18% |

Other components were:

- Vertex 4 SSD = OCZ-VERT Ex4 1.4 SATA
- RAID configurations = RAID0, 1, and 5 using an ARECA ARC-1882ix-12 Raid controller and 5-WD2002FAEX 7200 RPM SATA 64MB Cache 2TB drives
- PCIe SSD = OCZ-REVO3X2 PCIe

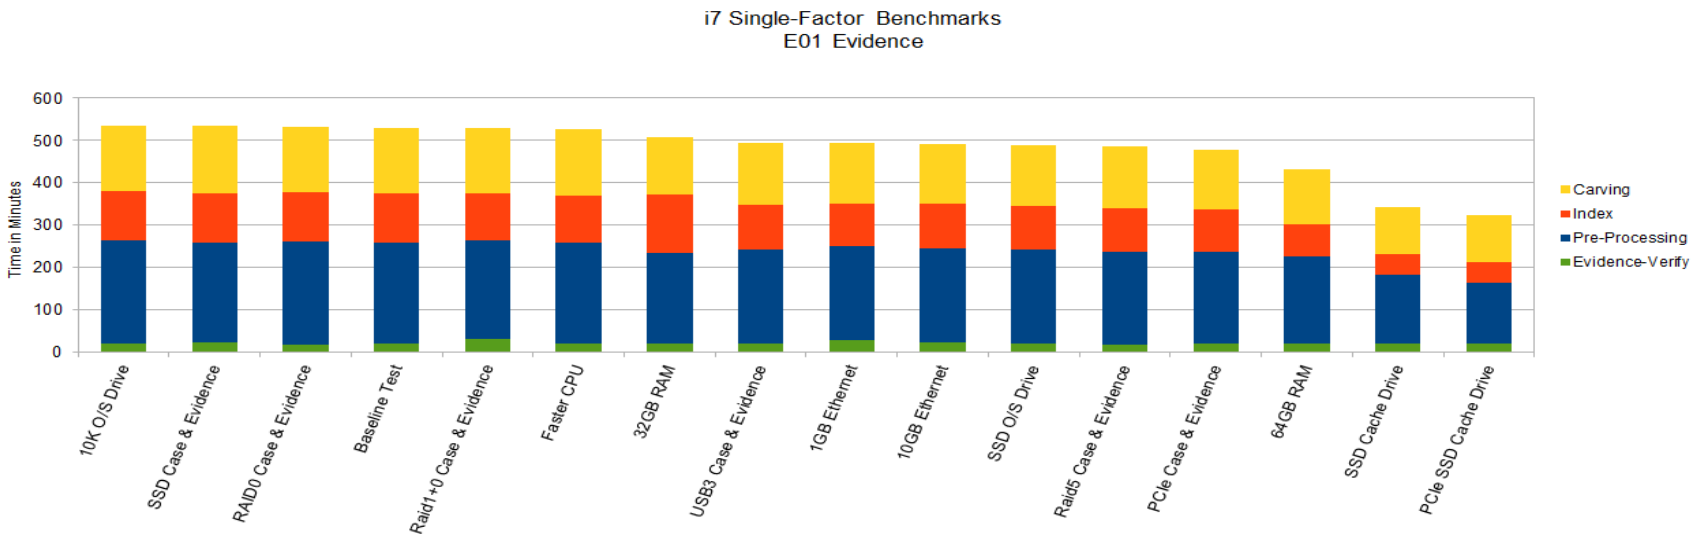


| Description | CPU | RAM | OSDrive & KFF Drive | Case and Evidence Drive | Primary Cache Drive | Evidence-Verify | Pre-Processing | Index | Carving | Total | % Change |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 10K O/S Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 10k Raptor | 7200 RPM SATA | 7200 RPM SATA | 20.1107 | 245.2397 | 114.5139 | 156.2439 | 536.1082 | -1% |
| SSD Case & Evidence | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Vertex 4 SSD | 7200 RPM SATA | 22.3456 | 236.4963 | 115.6186 | 159.5466 | 534.0071 | -1% |
| RAID0 Case & Evidence | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Areca 5X7200 RPM SATA – RAID0 | 7200 RPM SATA | 17.9236 | 243.0905 | 117.8417 | 154.0536 | 532.9094 | -0% |
| Baseline Test | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 21.2091 | 237.5512 | 115.6054 | 156.2426 | 530.6083 | 0% |
| Raid1+0 Case & Evidence | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Areca 5X7200 RPM SATA - RAID1+0 | 7200 RPM SATA | 31.1363 | 234.3069 | 109.053 | 155.1538 | 529.65 | 0% |
| Faster CPU | core i7-3960X@3.3Ghz - HEX - 15MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 21.2516 | 236.4965 | 112.3417 | 156.2824 | 526.3722 | 1% |
| 32GB RAM | core i7-3820@3.6Ghz - Quad - 10MB | 32 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 20.1106 | 213.3803 | 139.5101 | 134.2722 | 507.2732 | 4% |
| USB3 Case & Evidence | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | USB3 7200 RPM SATA | 7200 RPM SATA | 21.2444 | 221.1171 | 106.8186 | 146.3647 | 495.5448 | 7% |
| 1GB Ethernet | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 1GB Ethernet | 7200 RPM SATA | 27.8103 | 223.3027 | 99.1339 | 145.2655 | 495.5124 | 7% |
| 10GB Ethernet | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 10GB Ethernet | 7200 RPM SATA | 23.4462 | 222.2065 | 104.6545 | 140.9058 | 491.213 | 7% |
| SSD O/S Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | Vertex 4 SSD | 7200 RPM SATA | 7200 RPM SATA | 20.5 | 222.625 | 102.4156 | 144.1572 | 489.6978 | 8% |
| Raid5 Case & Evidence | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | Areca 5X7200 RPM SATA - RAID5 | 7200 RPM SATA | 17.947 | 220.0191 | 102.4242 | 145.2645 | 485.6548 | 8% |
| PCIe Case & Evidence | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | PCIe SSD | 7200 RPM SATA | 20.1513 | 217.8199 | 100.2285 | 139.8058 | 478.0055 | 10% |
| 64GB RAM | core i7-3820@3.6Ghz - Quad - 10MB | 64 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 20.1486 | 206.8371 | 73.8959 | 132.0837 | 432.9653 | 18% |
| SSD Cache Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | Vertex 4 SSD | 20.1107 | 161.7532 | 50.7902 | 109.006 | 341.6601 | 36% |
| PCIe SSD Cache Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | PCIe SSD | 21.2084 | 141.9843 | 49.6931 | 110.1044 | 322.9902 | 39% |

Other components were:

- Vertex 4 SSD = OCZ-VERT Ex4 1.4 SATA
- RAID configurations = RAID0, 1, and 5 using an ARECA ARC-1882ix-12 Raid controller and 5-WD2002FAEX 7200 RPM SATA 64MB Cache 2TB drives
- PCIe SSD = OCZ-REVO3X2 PCIe

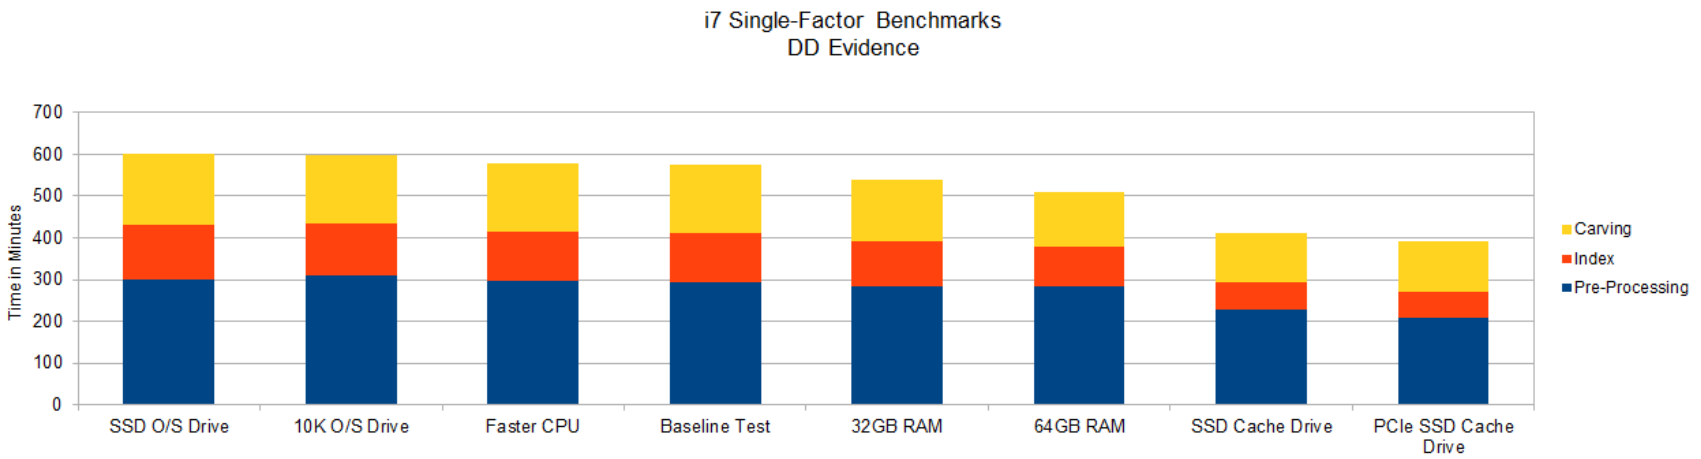


| Description | CPU | RAM | OSDrive & KFF Drive | Case and Evidence Drive | Primary Cache Drive | Pre-Processing | Index | Carving | Total | % Change |
|---|---|---|---|---|---|---|---|---|---|---|
| SSD O/S Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | Vertex 4 SSD | 7200 RPM SATA | 7200 RPM SATA | 301.299 | 129.9177 | 169.4299 | 600.6466 | -0% |
| 10K O/S Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 10k Raptor | 7200 RPM SATA | 7200 RPM SATA | 311.1839 | 124.3953 | 162.8414 | 598.4206 | 0% |
| Faster CPU | core i7-3960X@3.3Ghz - HEX - 15MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 295.8104 | 120.0025 | 163.94 | 579.7529 | 3% |
| Baseline Test | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 293.5692 | 117.7965 | 163.9311 | 575.2968 | 4% |
| 32GB RAM | core i7-3820@3.6Ghz - Quad - 10MB | 32 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 284.7812 | 105.7055 | 149.6501 | 540.1368 | 10% |
| 64GB RAM | core i7-3820@3.6Ghz - Quad - 10MB | 64 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 282.6238 | 96.9308 | 129.8848 | 509.4394 | 15% |
| SSD Cache Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | Vertex 4 SSD | 227.6589 | 65.0692 | 118.8928 | 411.6209 | 31% |
| PCIe SSD Cache Drive | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | PCIe SSD | 208.9841 | 60.674 | 122.1876 | 391.8457 | 35% |

Other Components were:

- Vertex 4 SSD = OCZ-VERT Ex4 1.4 SATA
- RAID configurations = RAID0, 1, and 5 using an ARECA ARC-1882ix-12 Raid controller and 5-WD2002FAEX 7200 RPM SATA 64MB Cache 2TB drives
- PCIe SSD = OCZ-REVO3X2 PCIe

## Resource Utilization Analysis

With the Step 3 tests completed, the quantitative impact of hardware selection becomes more obvious. In order of effectiveness:

- Increasing throughput for the Cache I/O channel significantly improves performance
- An increase in memory improves performance
- Increasing throughput for the Case and Evidence I/O channel improves performance
- Increasing the throughput for the O/S and KFF Database channel improves performance
- Optimized Network Storage improves performance
- CPU selection or architecture does not significantly improve performance

Most notably, increasing throughput to the Cache showed significant improvements in performance. Additionally, further incremental improvements in throughput appeared to scale accordingly.

Increasing memory resulted in a performance improvement. This suggests that the application can take advantage of additional memory resources.

Increasing throughput of the Case & Evidence I/O channel also showed a performance improvement. However, this performance improvement appeared to be somewhat limited regardless of the storage device employed. This could be an indication that demands of this channel are modest.

When testing with E01 (Compressed) data, improving the O/S and KFF Database channel improved performance. This is due to the demands of the decompression process.

The application performs comparably regardless of the architecture, speed, or number of processors. This strongly suggests that the application is in not processor bound. (Tests with other software products have also indicated that the forensic process, in general, is not processor bound).

Utilizing optimized network storage for the Case and Evidence I/O channel showed improvement comparing favorably to local RAID5 storage. This suggests that network storage is a viable alternative to local storage for Case and Evidence.

E01 (Compressed) data was processed faster than DD (Uncompressed) data when tested in the same hardware permutations. Limited DD tests validated similar trends found with the E01 tests; namely, that the Cache I/O channel was the most sensitive to performance improvements.

## Step 4 (Potential System Configurations)

With a better understanding of application resource utilization, it becomes possible to develop several relevant system configurations. Since the Dual-Xeon based architectures actually demonstrated lower performance, further testing would be performed using only i7 processor architectures.

The single most significant improvement in performance resulted from the increase in throughput of the Cache I/O channel; specifically with SSD architecture. To further explore this we evaluated:

- Standalone SATA Solid State Disk (SSD)
- PCIe-based Solid State Disk (SSD)

For the Case and Evidence I/O channel, storage devices to be tested included:

- USB 3.0 Hot-Swap Drive
- RAID-5 Array

Ultimate performance should not overshadow reliability - especially for the Case and Evidence channel. It should be noted, while unprotected storage environments (like RAID-0) might deliver marginally better performance, a RAID-5 volume is proven to provide critical data protection with only a very small decrease in performance. The same can be said for individual hard drives (including SSD – SATA or PCIe based), when considered for use in other storage positions where long term data preservation is also critical.

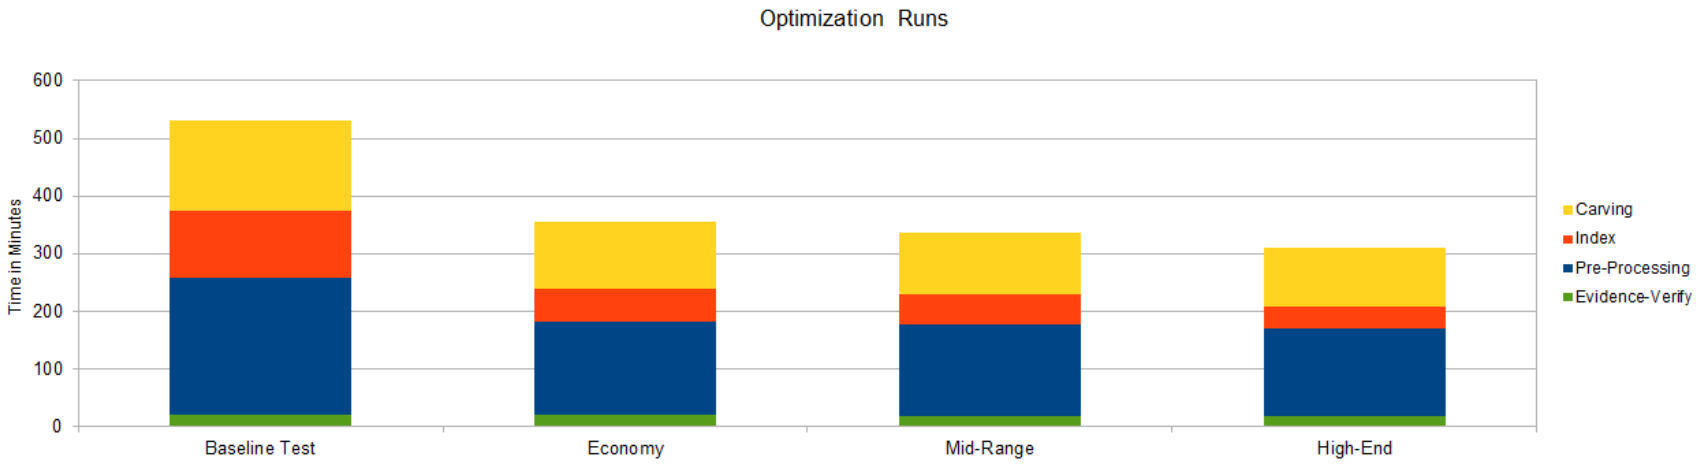


| Description | CPU | RAM | OSDrive & KFF Drive | Case and Evidence Drive | PrimaryCacheDrive | Evidence-Verify | Pre-Processing | Index | Carving | Total | % Change |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Baseline Test | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 7200 RPM SATA | 7200 RPM SATA | 7200 RPM SATA | 21.2091 | 237.5512 | 115.6054 | 156.2426 | 530.6083 | 0% |
| Economy | core i7-3820@3.6Ghz - Quad - 10MB | 16 GB | 10k Raptor | USB3 7200 RPM SATA | Vertex 4 SSD | 21.2425 | 161.7889 | 56.3276 | 115.6075 | 354.9665 | 33% |
| Mid-Range | core i7-3820@3.6Ghz - Quad - 10MB | 32GB | Vertex 4 SSD | USB3 7200 RPM SATA | Vertex 4 SSD | 21.2184 | 156.2904 | 55.22 | 111.2114 | 343.9402 | 35% |
| High-End | core i7-3960X@3.3Ghz - HEX - 15MB | 64GB | Vertex 4 SSD | Areca 5XWD2002FAEX 7200 RPM SATA 64MB Cache - RAID5 | PCIe SSD | 17.923 | 153.0016 | 37.6122 | 102.4226 | 310.9594 | 41% |

Other components were:

- Vertex 4 SSD = OCZ-VERT Ex4 1.4 SATA
- RAID configurations = RAID0, 1, and 5 using an ARECA ARC-1882ix-12 Raid controller and 5-WD2002FAEX 7200 RPM SATA 64MB Cache 2TB drives
- PCIe SSD = OCZ-REVO3X2 PCIe

## Analysis of Combined Components

An analysis of the results of Step 4 testing confirmed the following resources continued to benefit from further enhancement:

- Increased throughput on the Cache I/O channel: PCIe SSD showed improvement over SATA SSD for the Cache I/O channel
- Additional memory: increasing memory from 32 to 64 GB showed improvement
- Increased throughput on the Case and Evidence I/O channel: RAID-5 showed a small improvement over SATA and USB3 SATA Hot Swap for the Case and Evidence I/O channel

## Final Results

EnCase 7.05.02 benefits from improving the I/O frequency (IOPS) to the Cache channel, increasing memory, and, to a lesser extent, improving the I/O throughput to the case and evidence channel. There is no significant return on investment in utilizing more than three I/O channels for an EnCase 7 system. Additional processor speed, number of cores, or processor cache only improves performance at the very high end of I/O channel improvements. It should also be noted that the Dual-Xeon architecture did not distinguish itself in these tests.

Processing compressed evidence (E01) files was significantly faster (even with the additional "File Verification" step) than processing un-compressed evidence (DD) files. This was likely due to the reduced I/O requirements on the Case and Evidence I/O channel. With significant reductions in acquisition time, processing time, and reduced storage space requirements, E01 files should be considered a preferred format for evidence.

The following three configurations represent a range of component choices for the EnCase 7 environment when processing compressed (E01) Evidence:

| | Economy | Mid-Range | High-End |
|---|---|---|---|
| CPU | core i7-3820@3.6Ghz - Quad - 10MB | core i7-3820@3.6Ghz - Quad - 10MB | core i7-3960X@3.3Ghz - Hex - 15MB |
| Memory | 16 GB | 32 GB | 64 GB |
| O/S and dbTemp I/O Channel | 10K SATA | SATA SSD | SATA SSD |
| Cache I/O Channel | SATA SSD | SATA SSD | PCIe SSD |

| | | | |
|---|---|---|---|
| Case and Evidence I/O Channel | USB3 Attached SATA | USB3 Attached SATA | RAID-5 |
| Time (in minutes) | 355 | 344 | 310 |

## Other considerations

It should also be noted that raw PC performance is not the only factor to be considered when working to minimize case processing times. Functional convenience can also play a large part in minimizing overall case processing requirements. Little value can be demonstrated if a relatively expensive hardware selection generates a small performance gain but also brings with it significant administrative overhead. Although it has been demonstrated that higher cost fixed disk systems can provide measurable performance benefit, these fixed resources must be re-imaged or recreated each time the contents are to be replaced or updated. Depending on the amount of data involved, this re-imaging, recreation, or copying can take a significant amount of time. The resulting managerial overhead might easily be displaced through the use of removable media. As a result, any time advantage seen in the relatively high-cost, high-end solution might quickly be overcome through thoughtful management of casework data. This could easily include the use of paired sets of removable database and case/evidence drives as benchmarked. Similarly, although the location of the case/evidence on high speed network storage resulted in slightly lower performance, the administrative benefit is substantial.

## Observations and Summary

With the completion of over 45 iterations of EnCase 7.05.02 benchmarks, a number of interesting observations have been recorded. While many of our observations might be as expected, some were more interesting than others. The following observations appear to be the most relevant when selecting hardware components for processing in the EnCase 7 environment:

- **I/O Channel Configuration**: The analysis of bandwidth utilization for the 5 identified areas of I/O (O/S, Cache, Case, Evidence, and KFF Database drives) supported a reduction in I/O channels to a consolidation of 3 (O/S and KFF Database, Case and Evidence, and Cache). Testing demonstrated relatively insignificant change in case processing times while resulting in a much less complicated and expensive solution. Additionally, this configuration also lends very well to simplified case management as it maintains both Casework and Evidence on the same storage device.

- **I/O Component (Drive) Selection**:

    - The Cache Drive: Careful selection of the Cache Drive is proven to be the most important drive choice with respect to performance. Selecting a Cache Drive capable of supporting a very high level of I/O Operations per Second (IOPS) results in significant performance gains. Solid State Disks are most beneficial in this position with further gains delivered by the PCIe based implementations.

    - The Case and Evidence Drive: The demands of Case and Evidence drive are served very well by a local USB3 connected (Hot-swappable) SATA drive or Network Storage. RAID arrays appear to work very well in this position as they provide improved throughput as well as increased local redundant storage capacity. The demands of the test suite do not require high I/O throughput from this channel.

    - The O/S and KFF Database Drive: The choice of the O/S and KFF Database drive appeared to have a minimal effect on system performance. An improvement is to be found in selecting an SSD in this position when working with E01 (Compressed) evidence.

- **System Memory:** Increasing the system memory reduced case processing times. The improvement did scale suggesting that maximizing memory is beneficial.

- **CPU:** CPU clock rate, number of cores, or multiple CPU architectures did not have a significant impact on processing times until the I/O subsystems were fully optimized. This is due to the I/O bias of case processing tasks. Ultimately, Dual-Xeon systems do not justify added expense over the i7 processor based systems. This is likely a result of newer i7 systems having much more capable I/O subsystems when compared to the more "mature" implementations typically found on Xeon based platforms.

- **Evidence File Format**: The difference in performance between processing Compressed (E01) and Uncompressed (DD) file formats was quite significant. As we have seen that the ultimate limitation on processing performance is often the I/O throughput capacity of the system, lessening I/O requirements can be of obvious benefit. By using a compressed data source, we are able to trade some CPU activity in lieu of I/O demands. Additionally, testing has demonstrated that Image decompression is one of the few processing activities which place any significant demands on the CPU. Using a compressed image format quite simply helps offload a portion of the very busy I/O demands onto a much less used CPU resource.