Quantifying Hardware Selection in an FTK 4.0 Environment

Introduction and Background

The purpose of this analysis is to evaluate the relative effectiveness of individual hardware component selection in the FTK 4.0 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.) Database
4.) Cache/TempDB
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 upgrades 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:

- Pre-Processing (factory defaults except as follows):
  - (Disable SHA-256)
  - (Enable Duplicate Files)
  - (Disable dtSearch)

Additional Processing Functions:

- dtSearch/Entropy
- Data and Meta Carving
- Known File Filter

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 disks to their baseline state. This would ensure that all residual data from the previous test would be eliminated including any file-system fragmentation, database fragmentation, or file relocation.

The test suite would be run utilizing both compressed (E01) and un-compressed (DD) evidence since both formats are supported by the application.

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 image (created with Tableau Imager) consisted of approximately 240 GB of uncompressed data.

The test disk was also imaged using the E01 compressed format the resulting dataset was approximately 60GB. The reduction in overall data output during the imaging process reduced acquisition times significantly.

Step 2 (I/O Channel Evaluation)

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

Separate Channel Configuration

Drive Letter	Contents
C:	Win7 x64 Operating System
F:	Evidence
G:	Database
H:	Case
L:	dbTemp

(A sixth MS-Dos partition, Drive E, was also present on the C: disk but was not employed during testing. Drive D was the CD ROM.)

The following charts were obtained from Windows Performance Monitor. An analysis of the disk activity indicated that both the O/S and Cache/dbTemp channels, as well as the Evidence and Case channels, could be combined with negligible performance impact.

Tests were then 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 2% 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 dbTemp files
F:	Evidence and Case
G:	Database

The following graphs compare the I/O usage of each of the processing steps with the “Separate Channel Configuration” on the left and the “Consolidated Channel Configuration” on the right.

Pre-Processing:
dtSearch/Entropy:

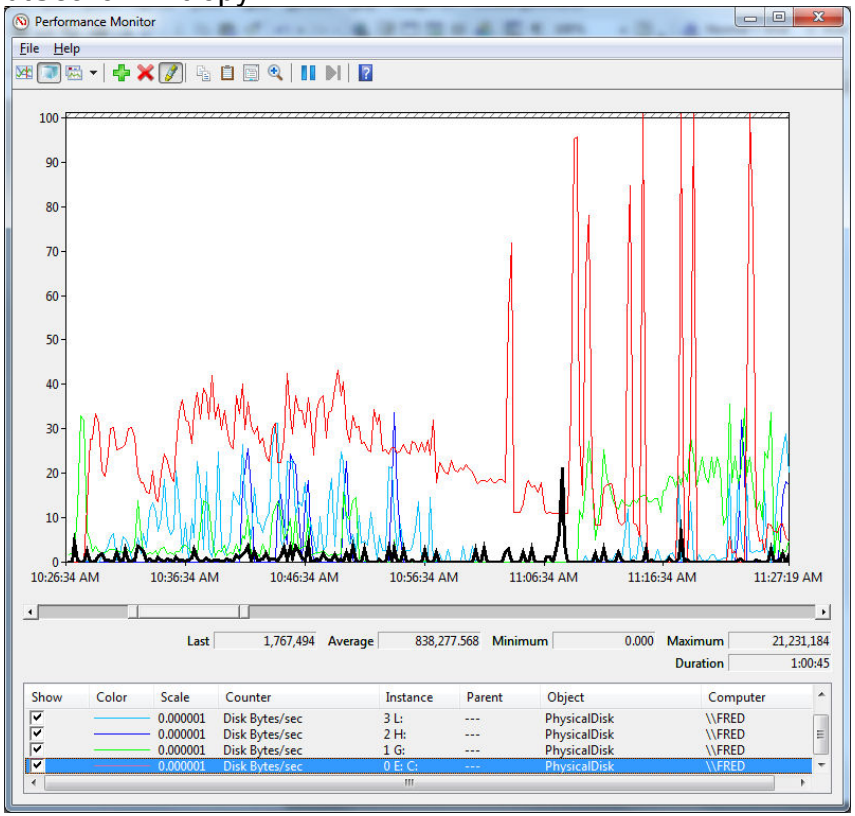

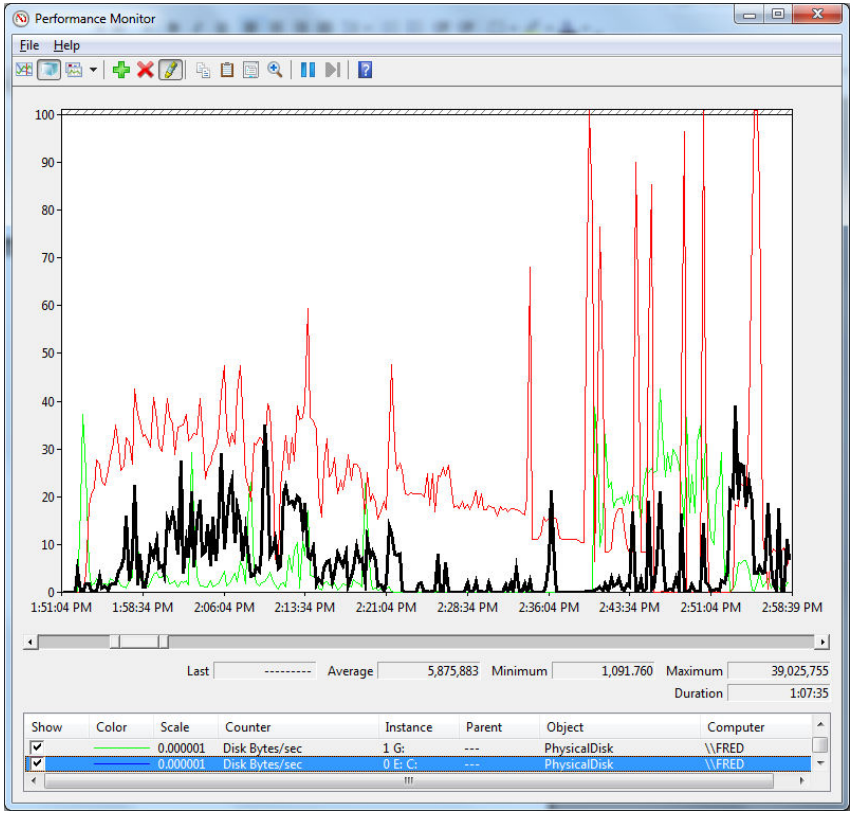

Legend:
Lt. Blue (L:) – dbTemp
Blue (H:) – Case
Green (G:) – Database
Red (F:) - Evidence
Black (C:) – O/S

Green (G:) – Database
Red(F:) – Evidence and Case
Black (C:) – O/S and dbTemp

Data and Meta Carving:

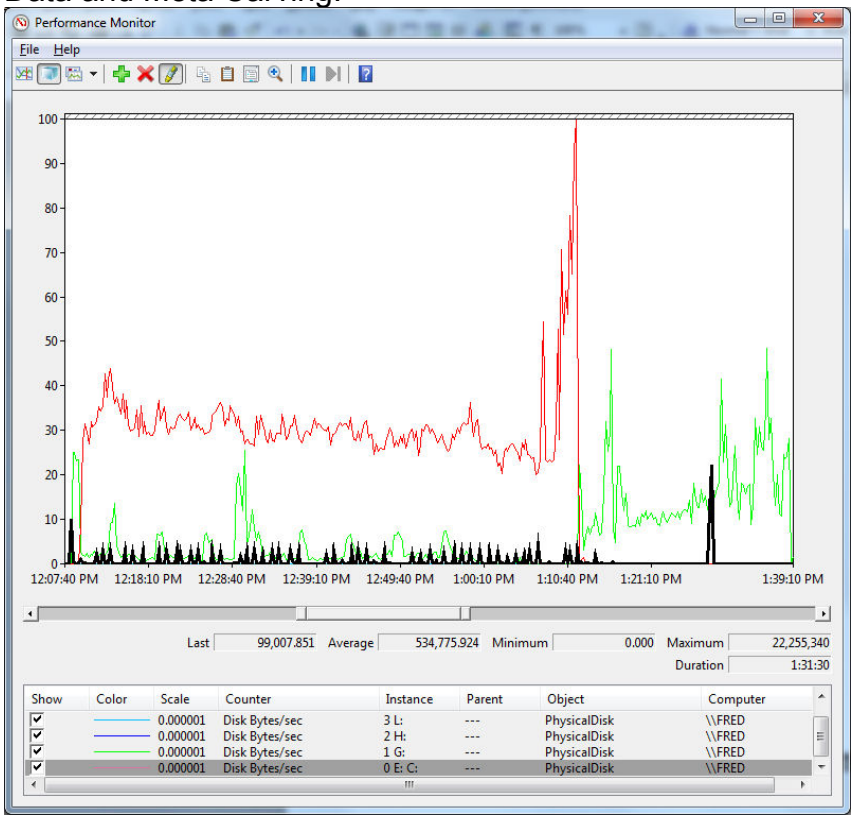

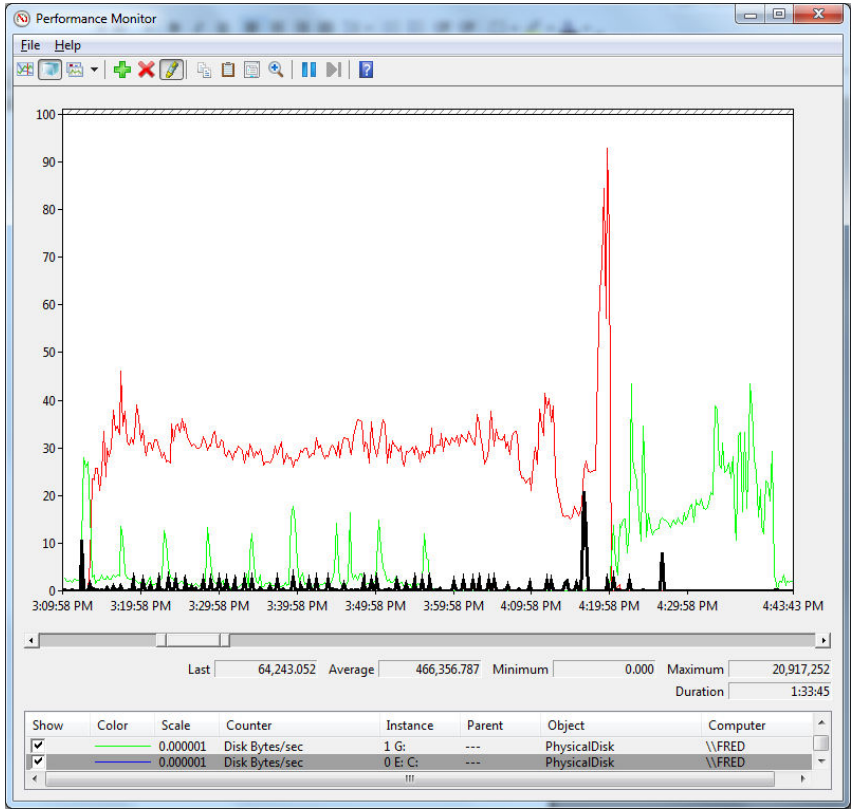

Legend:
Lt. Blue (L:) – dbTemp
Blue (H:) – Case
Green (G:) – Database
Red (F:) - Evidence
Black (C:) – O/S

Green (G:) – Database
Red(F:) – Evidence and Case
Black (C:) – O/S and dbTemp

Known File Filter:

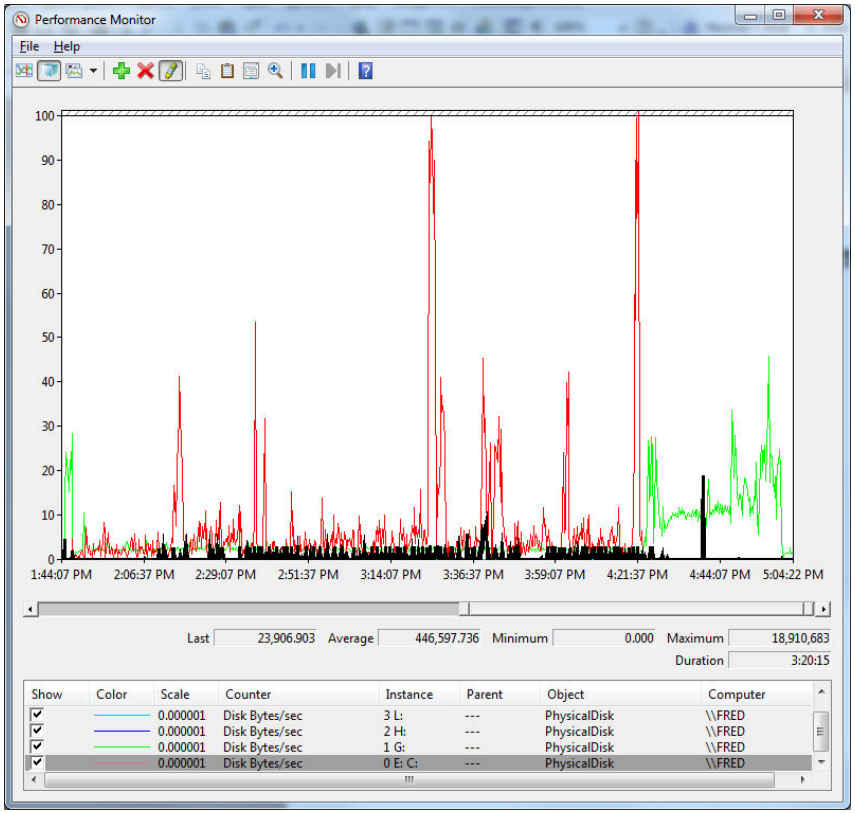

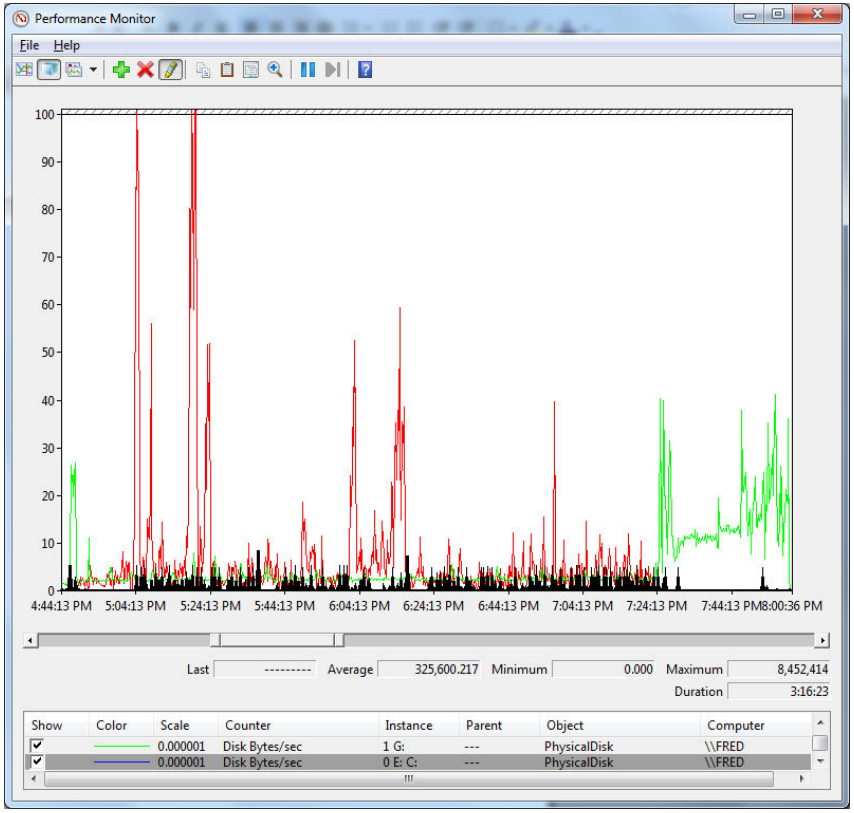

Legend:
Lt. Blue (L:) – dbTemp
Blue (H:) – Case
Green (G:) – Database
Red (F:) - Evidence
Black (C:) – O/S

Green (G:) – Database
Red(F:) – Evidence and Case
Black (C:) – O/S and dbTemp

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 & Temp Drive
- DB 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 3.6 Ghz Quad Core 10MB Cache	E5-2609 2.4 Ghz Quad Core (8 cores total) 10MB Cache
Chipset	X79	C602
Memory	16 GB	16 GB
O/S & DBTemp Drive	10K RPM SATA	10K RPM SATA
Database Drive	7200 RPM SATA	7200 RPM SATA
Evidence & Case Drive	7200 RPM SATA	7200 RPM SATA

- 16 GB Memory = DDR3-1600
- 10K RPM SATA = WD3000HLHX (Velociraptor) 300 GB
- 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 and un-compressed 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. FTK version 4 (4.0.2) and the Postgres SQL database was installed and configured per the manufacture's instructions.

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

Description	CPU	RAM	OSDrive	DatabaseDrive	EvidenceDrive	Pre-Processing	dtSearch	Carving	KFF	Total	%Change from baseline
Base System	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	338.6763	531.5186	484.0294	603.5341	1957.758	0%
SSD O/S Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	Vertex 4 SSD	7200 RPM SATA	7200 RPM SATA	296.2049	558.7994	492.8023	600.7014	1948.508	0%
Faster Processor	2-E5-2630@2.3Ghz - Hex - 15MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	259.8914	460.1538	501.9603	643.3581	1865.364	5%
SSD Data	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Vertex 4 SSD	309.7278	331.6214	462.9151	605.9666	1710.231	14%
RAID1 Data	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA RAID1+0	402.8598	304.7582	422.0155	580.0796	1709.713	15%
RAID5 Data	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID5	308.134	322.0277	444.2943	611.3797	1685.836	16%
RAID0 Data	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA RAID0	389.2505	304.6735	397.2652	559.7303	1650.92	19%
RAID5 Database	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA RAID5	7200 RPM SATA	290.446	495.6925	364.7414	378.4448	1529.325	28%
RAID1 Database	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA RAID1+0	7200 RPM SATA	339.9397	472.0692	379.779	293.4145	1485.202	32%
32 GB Memory	2-E5-2609@2.4Ghz - Quad - 10MB	32 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	364.7281	407.8952	272.1293	272.4531	1317.206	49%
SSD Database	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Vertex 4 SSD	7200 RPM SATA	297.0309	477.7824	360.6814	143.599	1279.094	53%
RAID0 Database	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA RAID0	7200 RPM SATA	272.8281	481.8646	273.9423	232.1	1260.735	55%

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

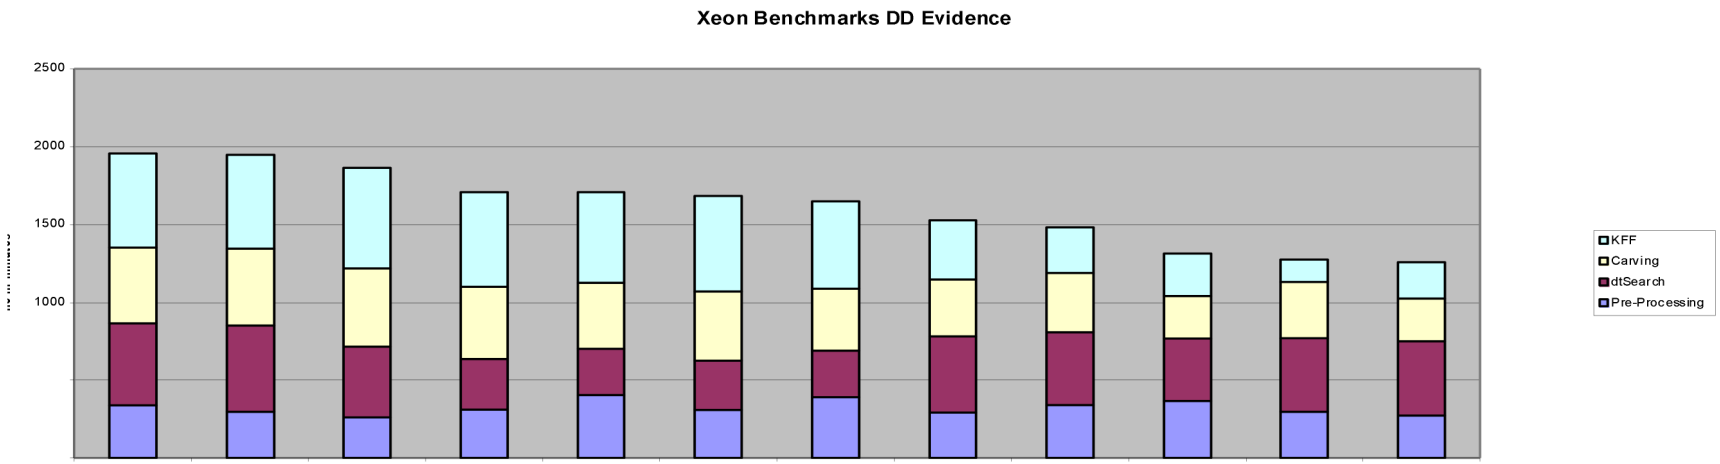

Description	CPU	RAM	OSDrive	DatabaseDrive	EvidenceDrive	Pre-Processing	dtSearch	Carving	KFF	Total	%Change from baseline
PCIe Evidence & Case Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	PCIe SSD	281.1762	430.5167	421.9005	500.6462	1634.2396	-15%
SSD O/S Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	Vertex 4 SSD	7200 RPM SATA	7200 RPM SATA	295.6153	311.5974	376.6228	637.65	1621.4855	-14%
Base System	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	359.7857	247.7694	304.8353	474.9433	1387.3337	0%
RAID1 Evidence & Case Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID1+0	306.2967	248.853	318.5605	470.5847	1344.2949	3%
RAID5 Evidence & Case Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID5	311.8057	249.6804	308.1161	457.7409	1327.3431	5%
Faster Processor	2-E5-2630@2.3Ghz - Hex - 15MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	229.023	251.6775	336.6487	495.968	1313.3172	6%
RAID0 Evidence & Case Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID0	293.4432	245.9345	305.0217	459.2524	1303.6518	6%
PCIe Evidence & Case Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	PCIe SSD	298.9644	245.2616	298.2354	458.4955	1300.9569	7%
64 GB Memory	2-E5-2609@2.4Ghz - Quad - 10MB	64 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	436.7485	205.2038	267.0517	361.5297	1270.5337	9%
32 GB Memory	2-E5-2609@2.4Ghz - Quad - 10MB	32 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	368.4801	194.5161	259.0134	355.3428	1177.3524	18%
SSD Evidence & Case Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Vertex 4 SSD	311.1356	268.1167	198.993	317.2063	1095.4516	27%
RAID5 Database Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA - RAID5	7200 RPM SATA	291.3167	183.3894	219.8323	291.7533	986.2917	41%
RAID1 Database Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA - RAID0	7200 RPM SATA	345.0293	113.5549	121.3996	171.7315	751.7153	85%
RAID0 Database Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA - RAID1+0	7200 RPM SATA	260.2833	125.7064	125.1628	178.9993	690.1518	101%
SSD Database Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	Vertex 4 SSD	7200 RPM SATA	357.5504	100.9533	103.8784	100.3895	662.7716	109%
PCIe Database Drive	2-E5-2609@2.4Ghz - Quad - 10MB	16 GB	10k Raptor	PCIe SSD	7200 RPM SATA	293.5348	111.5606	103.462	102.9498	611.5072	127%

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

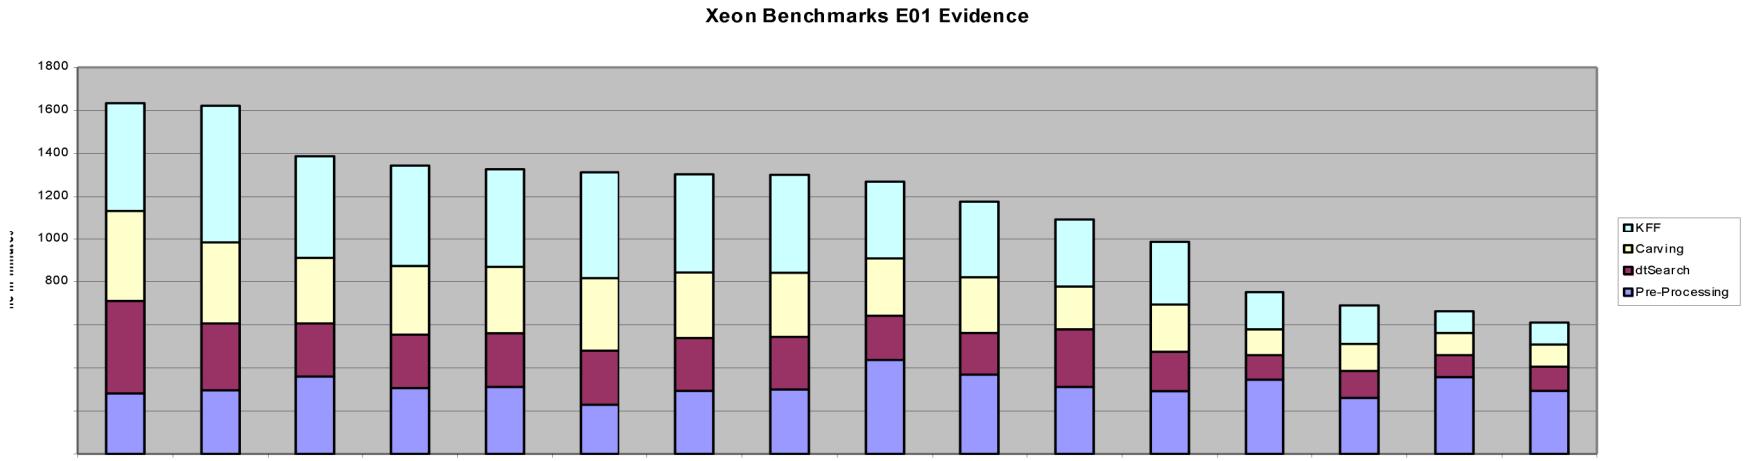

Description	CPU	RAM	OSDrive	DatabaseDrive	EvidenceDrive	Pre-Processing	dtSearch	Carving	KFF	Total	%Change from baseline
SSD O/S Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	Vertex 4 SSD	7200 RPM SATA	7200 RPM SATA	400.7672	519.0058	489.7701	600.8633	2010.406	-3%
Base System	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	444.6675	505.6481	474.1929	529.3455	1953.854	0%
Faster Processor	core i7-3960X@3.3Ghz - Hex - 15MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	316.7296	577.1724	435.1224	553.1401	1882.165	4%
RAID1 Data	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	ARECA 5X7200 RPM SATA RAID1+0	400.0267	353.2456	455.0443	594.8655	1803.182	8%
SSD Data	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Vertex 4 SSD	343.9371	333.1365	433.1946	577.594	1687.862	16%
RAID5 Data	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	ARECA 5X7200 RPM SATA RAID5	310.6354	345.6578	433.5258	583.6931	1673.512	17%
RAID5 Database	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	ARECA 5X7200 RPM SATA RAID5	7200 RPM SATA	327.5033	522.671	407.6679	409.7482	1667.59	17%
RAID0 Data	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	ARECA 5X7200 RPM SATA RAID0	285.433	335.2983	433.6147	590.3512	1644.697	19%
RAID1 Database	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	ARECA 5X7200 RPM SATA RAID1+0	7200 RPM SATA	323.5007	496.465	415.2411	296.0835	1531.29	28%
RAID0 Database	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	ARECA 5X7200 RPM SATA RAID0	7200 RPM SATA	346.1913	473.7424	339.2267	221.0384	1380.199	42%
32 GB Memory	core i7-3820@3.6Ghz - Quad - 10MB	32 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	344.5412	428.1997	231.7255	224.3171	1228.784	59%
SSD Database	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	Vertex 4 SSD	7200 RPM SATA	270.4248	524.9491	291.1309	128.7667	1215.272	61%

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

Core-I7 Benchmarks DD Evidence

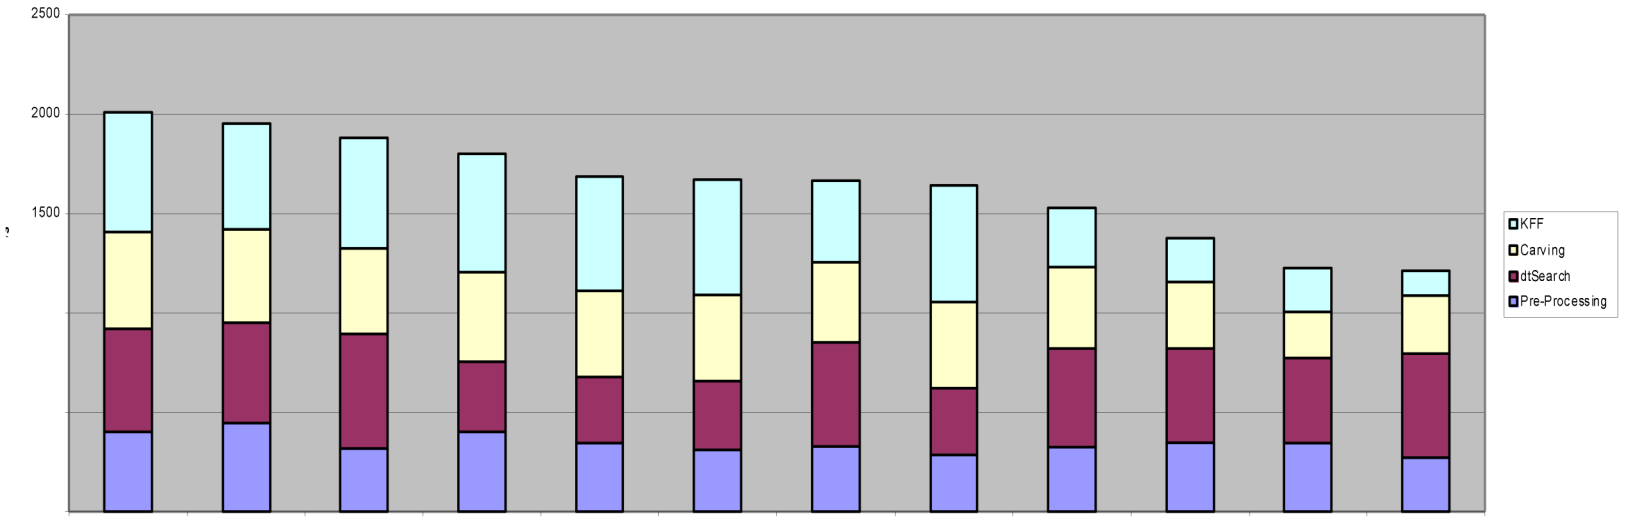

Description	CPU	RAM	OSDrive	DatabaseDrive	EvidenceDrive	Pre-Processing	dtSearch	Carving	KFF	Total	%Change from baseline
RAID0 Evidence & Case Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID0	410.8711	269.5431	332.3162	491.0501	1503.7805	-2%
Base System	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	370.3913	273.7311	332.2209	502.4039	1478.7472	0%
SSD O/S Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	Vertex 4 SSD	7200 RPM SATA	7200 RPM SATA	373.1444	281.9813	315.1991	506.4	1476.7248	0%
Faster Processor	core i7-3960X@3.3Ghz - HEX - 15MB	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	314.3941	285.9842	350.0042	521.6879	1472.0704	0%
RAID1 Evidence & Case Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID1+0	287.2785	275.2232	334.1563	509.0858	1405.7438	5%
Hot Swap USB3 Evidence & Case Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	USB3 7200 RPM SATA	289.8641	270.7911	343.1807	499.9704	1403.8063	5%
PCIe Evidence & Case Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	PCIe SSD	334.6706	251.7804	320.1254	485.373	1391.9494	6%
RAID5 Evidence & Case Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Areca 5X7200 RPM SATA - RAID5	274.7469	269.9841	325.4505	500.2475	1370.429	8%
Hot Swap USB3 Database Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	USB3 7200 RPM SATA	7200 RPM SATA	328.4973	256.7856	279.4774	432.8635	1297.6238	14%
64 GB Memory	core i7-3820@3.6Ghz - Quad - 10MB	64 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	366.578	196.5164	266.1331	367.6144	1196.8419	24%
32 GB Memory	core i7-3820@3.6Ghz - Quad - 10MB	32 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	371.0833	197.8015	261.9968	359.6004	1190.482	24%
RAID5 Database Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA - RAID5	7200 RPM SATA	319.8286	196.181	235.0261	326.3785	1077.4142	37%
SSD Evidence & Case Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	7200 RPM SATA	Vertex 4 SSD	266.6667	171.4291	222.3944	345.248	1005.7382	47%
RAID1 Database Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA - RAID1+0	7200 RPM SATA	340.3436	130.1883	141.6009	223.2231	835.3559	77%
RAID0 Database Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	Areca 5X7200 RPM SATA - RAID0	7200 RPM SATA	337.9365	117.8202	112.1362	169.2321	737.125	101%
PCIe Database Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	PCIe SSD	7200 RPM SATA	326.5893	102.0325	86.3614	88.8569	603.8401	145%
SSD Database Drive	core i7-3820@3.6Ghz - Quad - 10MB	16 GB	10k Raptor	Vertex 4 SSD	7200 RPM SATA	263.2338	105.0475	88.4543	88.2807	545.0163	171%

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

Core I7 Benchmarks E01 Evidence

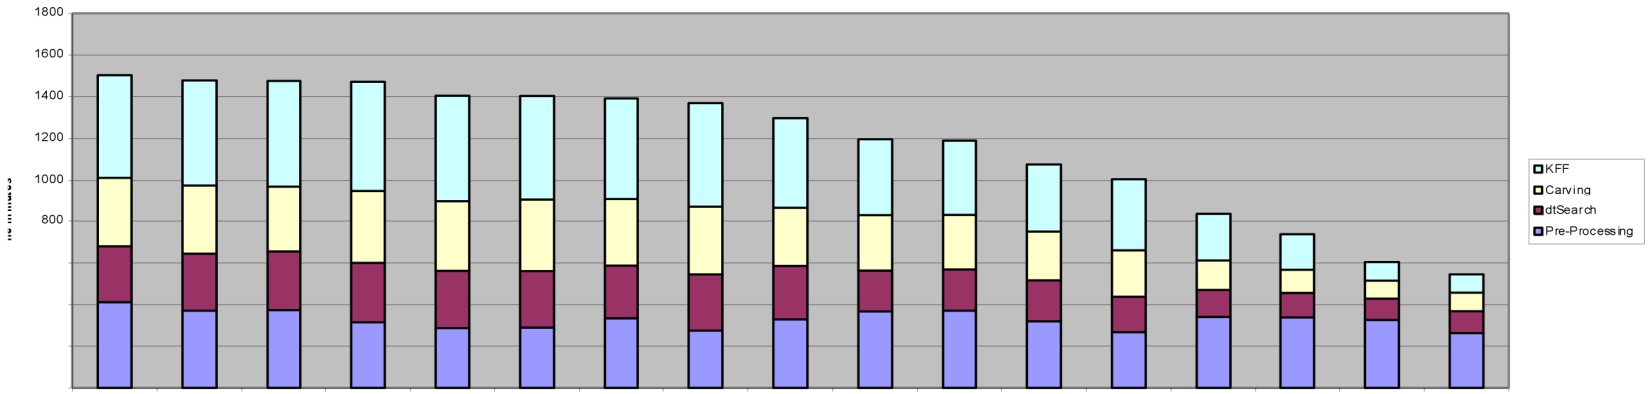

Resource Utilization Analysis

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

- Increasing throughput for the Database I/O channel significantly improves performance
- An increase from 16GB to 32GB of memory improves performance
- Increasing throughput for the Case and Evidence I/O channel improves performance
- CPU selection or architecture does not significantly affect performance
- Increasing the throughput for the O/S and dbTemp I/O channel does not significantly improve performance

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

Increasing memory from 16 to 32 GB resulted in a large 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 other resources had become the limiting factor.

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

Increasing throughput for the Operating System I/O channel did not yield a performance improvement. This suggests that the I/O requirements of this channel can be met by storage device of modest architecture.

Step 4 (Potential System Configurations)

With a better understanding of application resource utilization, it becomes possible to develop several relevant system configurations. As little benefit could be demonstrated with the upgrade of CPU speed, count, or architecture, further testing would be performed using only an i7 processor system.

The single most significant improvement in performance resulted from the increase in throughput of the database 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, additional storage devices to be tested included:

- USB 3.0 Hot-Swap Drive
- Network-based storage
- RAID-5 Array
- PCIe-based Solid State Disk (SSD)

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), when considered for use in other storage positions where long term data preservation is also critical.

Description	RAM	OSDrive	DatabaseDrive	EvidenceDrive	Pre-Processing	dtSearch	Carving	KFF	Total	% Improvement from baseline
Base System	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	444.6675	505.6481	474.1929	529.3455	1953.854	0%
Base System + 32GB	32 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	344.5412	428.1997	231.7255	224.3171	1228.784	59%
Base System + 64GB	64 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	249.2	301.4815	316.1907	283.5444	1150.417	70%
SSD Database + SATA Data	64 GB	10k Raptor	Vertex 4 SSD	7200 RPM SATA	297.3833	331.0562	242.4126	78.7481	949.6002	106%
SSD Database + Network Data	64 GB	10k Raptor	Vertex 4 SSD	FREDC EXTRAID	327.7425	134.8374	172.3926	99.1098	734.0823	166%
SSD Database + USB3 Data	64 GB	10k Raptor	Vertex 4 SSD	USB3 7200 RPM SATA	229.8507	184.4829	189.4117	78.0851	681.8304	187%
SSD Database + RAID5 Data	64 GB	10k Raptor	Vertex 4 SSD	ARECA 5X7200 RPM SATA RAID5	236.65	103.4449	112.4685	73.7307	526.2941	271%
PCIe Database + PCIe Data	64 GB	10k Raptor	PCIe SSD	PCIe SSD	301.5401	82.091	68.0702	58.8064	510.5077	283%
PCIe Database + RAID5 Data	64 GB	10k Raptor	PCIe SSD	ARECA 5X7200 RPM SATA RAID5	201.4778	101.1269	107.2116	72.1571	481.9734	305%
Faster Processor + PCIe Database + RAID5 Data	64 GB	10k Raptor	PCIe SSD	ARECA 5X7200 RPM SATA RAID5	175.361	88.343	100.9371	73.5869	438.228	346%

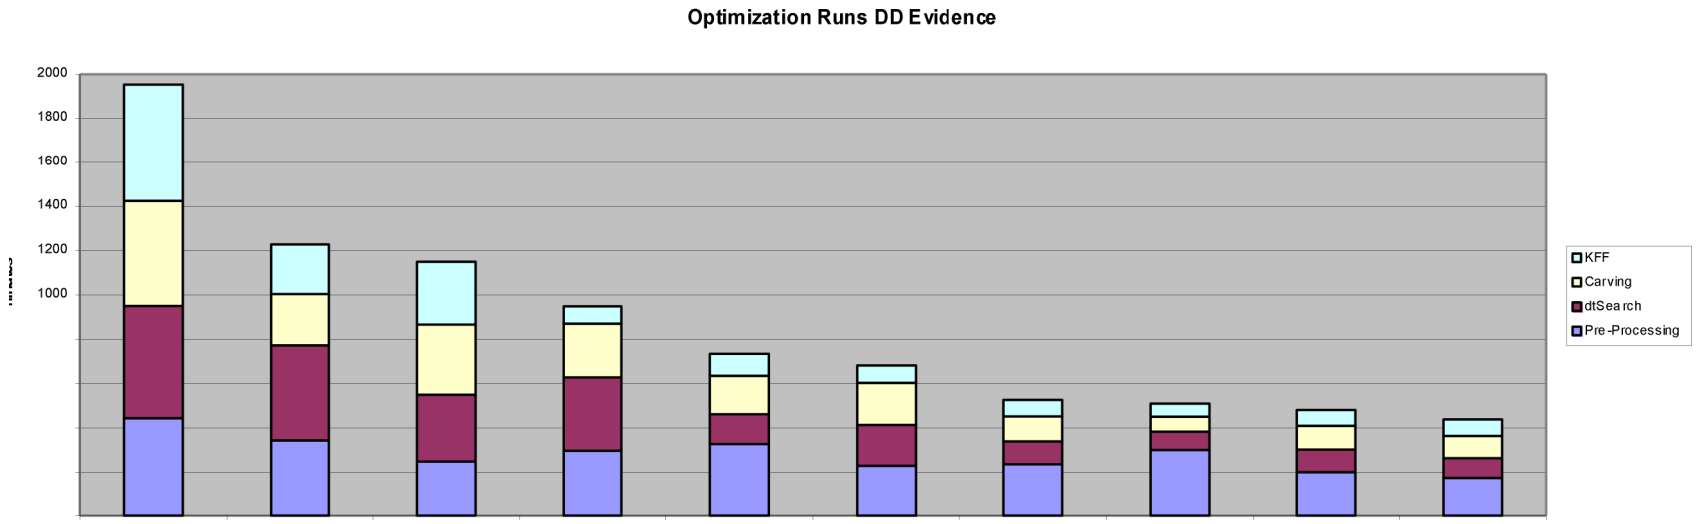

Description	RAM	OSDrive	DatabaseDrive	EvidenceDrive	Pre-Processing	dtSearch	Carving	KFF	Total	% Improvement from baseline
Base System	16 GB	10k Raptor	7200 RPM SATA	7200 RPM SATA	370.3913	273.7311	332.2209	502.4039	1478.747	0%
32GB + SSD Database + USB3 SATA DATA	32 GB	10k Raptor	Vertex 4 SSD	USB3 7200 RPM SATA	318.2403	80.085	69.5041	60.4075	528.2369	180%
64GB + PCIe Database + Raid5 DATA	64 GB	10k Raptor	PCIe SSD	Areca 5X7200 RPM SATA - RAID5	263.8131	79.2694	66.244	61.4847	470.8112	214%
32GB + SSD Database + Raid5 DATA	32 GB	10k Raptor	Vertex 4 SSD	Areca 5X7200 RPM SATA - RAID5	244.6276	87.4292	70.0803	61.5761	463.7132	219%
32GB + PCIe Database + Raid5 DATA	32 GB	10k Raptor	PCIe SSD	Areca 5X7200 RPM SATA - RAID5	201.3041	82.7745	64.8292	57.2178	406.1256	264%
Faster Processor + 32GB + PCIe Database + Raid5 DATA	32 GB	10k Raptor	PCIe SSD	Areca 5X7200 RPM SATA - RAID5	150.0648	66.4882	57.0659	49.8788	323.4977	357%

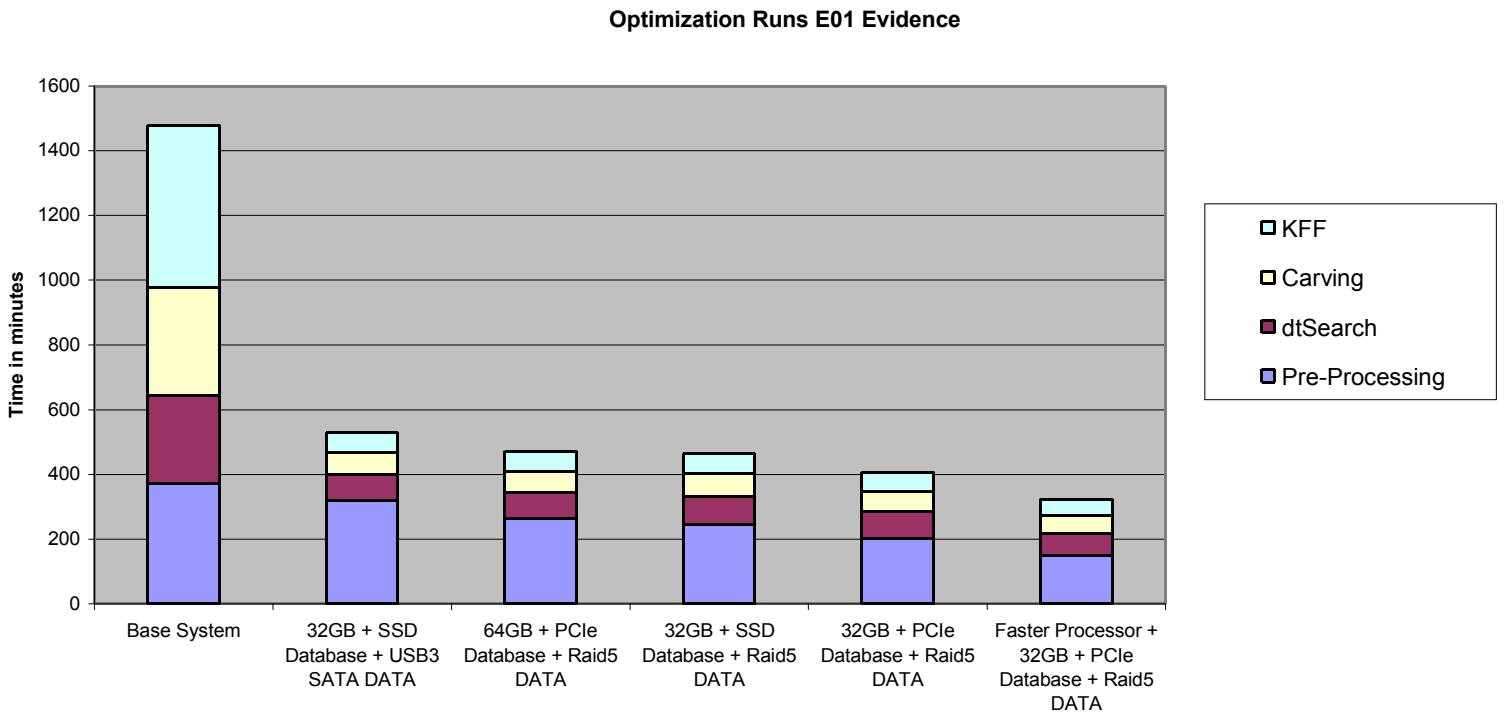

Analysis of Combined Components

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

- Increased throughput on the Database I/O channel: PCIe SSD showed improvement over SATA SSD for the Database I/O channel
- Additional memory: increasing memory from 32 to 64 GB showed improvement with DD Evidence but a slight degradation in performance with E01 Evidence
- Increased throughput on the Case and Evidence I/O channel: RAID-5 showed improvement over SATA and USB3 SATA Hot Swap for the Case and Evidence I/O channel

Accessing the Case & Evidence via a network file server was tested to assess the relative performance of storing the casework and evidence in a centralized location. There are many benefits such as backup management, redundancy, and shared storage, which could mitigate a small reduction in overall performance.

Final Results

With over 70 different configurations tested (between the 2 architectures and 2 evidence formats), the relative value of hardware selection becomes more obvious. Significant improvements can be obtained with an optimized choice of components.

FTK 4.0.2 benefits from increasing memory, improving the I/O frequency (IOPS) to the database channel, and 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 FTK 4 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 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. Systems which will be processing un-compressed evidence will benefit from maximizing memory while systems processing compressed evidence will find memory above 32 GB to be slightly detrimental. In all other aspects, processing compressed or un-compressed evidence had similar resource requirements. With significant reductions in acquisition time, processing time, and reduced storage space requirements, E01 files might be considered a preferred format for evidence.

The following three configurations represent a range of component choices for the FTK 4 environment when processing un-compressed (DD) 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	32 GB	64 GB	64 GB
O/S and dbTemp I/O Channel	10K SATA	10K SATA	10K SATA
Database I/O Channel	SATA SSD	SATA SSD	PCIe SSD
Case and Evidence I/O Channel	SATA	USB3 SATA	RAID-5
Time (in minutes)	1078	734	438

The following three configurations represent a range of component choices for the FTK 4 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	32 GB
O/S and dbTemp I/O Channel	10K SATA	10K SATA	10K SATA
Database I/O Channel	SATA SSD	SATA SSD	PCIe SSD
Case and Evidence I/O Channel	SATA	RAID-5	RAID-5
Time (in minutes)	545	464	323

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 even higher.

Observations and Summary

With the completion of over 70 iterations of FTK 4.0.2 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 FTK 4.02 environment:

- **I/O Channel Configuration**: The analysis of bandwidth utilization for the 5 identified areas of I/O (O/S, dbTemp, Case, Evidence, and Database drives) supported a reduction in I/O channels to a consolidation of 3 (O/S and dbTemp, Case and Evidence, and Database). 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 Database Drive: Careful selection of the Database drive is proven to be the most important drive choice with respect to performance. Selecting a Database 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 second most important drive selection was proven to be the Case and Evidence Drive. The Case and Evidence drive is served very well by a storage device of very high I/O Bandwidth (MB/Sec). RAID arrays appear to work very well in this position as they provide very high throughput as well as increased storage capacity. Specifically, RAID-5 volumes delivered performance almost as high as RAID-0 volumes, but have the significant benefit of data protection in the event of drive failure.

- The O/S and Database Temp Drive: The choice of the O/S and dbTemp drive appeared to have the least effect on system performance. A relatively inexpensive 10K RPM SATA drive performed well in this position.

- **System Memory:** Increasing the system memory significantly reduced case processing times. However, it should also be noted, that when working with compressed evidence files (E01), an increase beyond 32 GB was actually detrimental to performance. This is particularly noteworthy, as there are significant performance gains which can be had by using compressed evidence drives (as discussed below).

- **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 places 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.