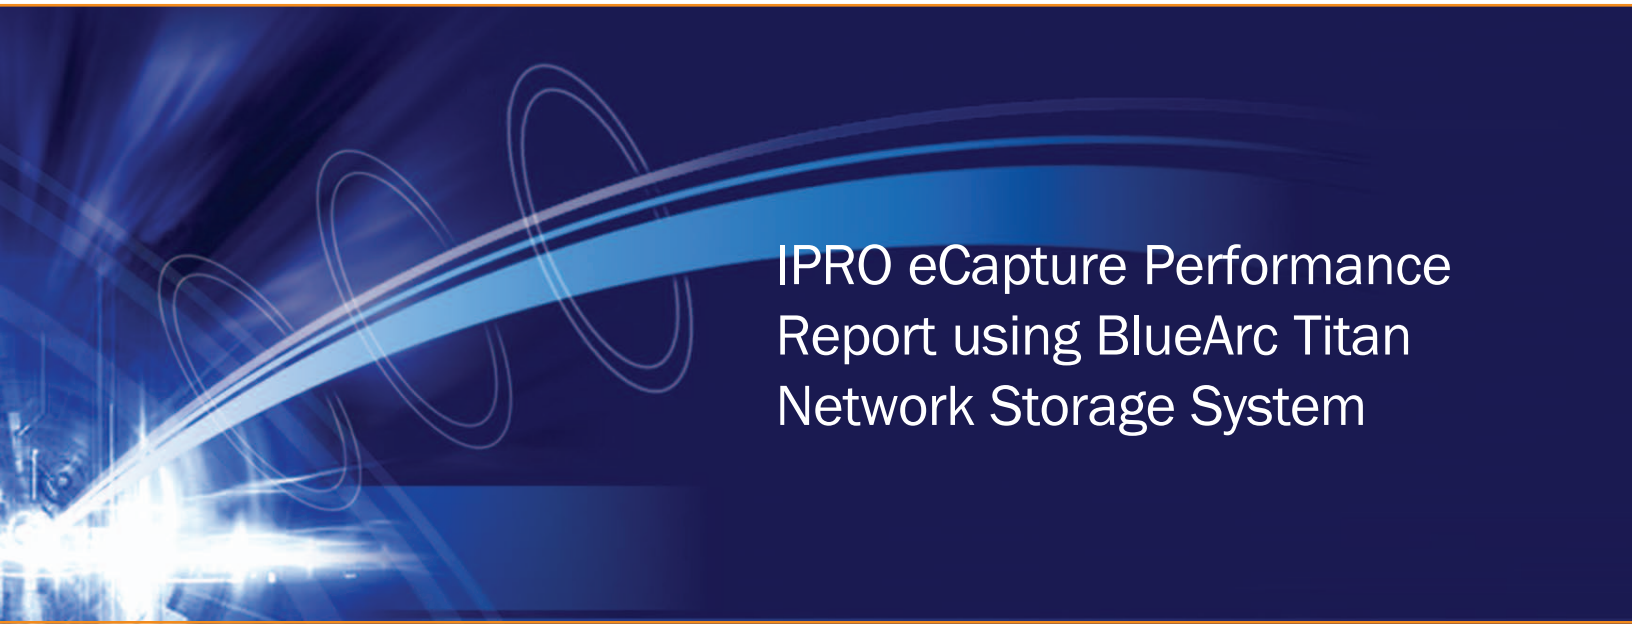

IPRO eCapture Performance Report using BlueArc Titan Network Storage System

Allen Yen, BlueArc Corp

Jesse Abrahams, IPRO Tech, Inc

IPRO Innovative Thinking ... Superior Software

BLUEARC®

Introduction

IPRO eCapture is an e-discovery application designed to handle large projects with speed and scalability. It is a SQL based, client-server application with multi-threading capability that can be distributed across any number of workstations. IPRO refers to these distributed workstations as "Workers." All processing is managed from a central console and all processing operations are performed on workers without operator intervention.

BlueArc's Titan Network Storage System represents an evolutionary step in file servers — by creating a hardware-based file system that can scale bandwidth, IOPS, and overall data capacity well beyond conventional software-based devices. With its ability to virtualize a massive storage pool of up to four usable petabytes of tiered storage, Titan can scale with growing data requirements, offering a competitive advantage for businesses, researchers, or other enterprises seeking to better manage data growth while still ensuring optimal performance.

The goals of this benchmark are to

- measure the workload characteristics and performance benchmarks of the Titan storage system in a processing architecture including IPRO eCapture running on a pool of Workers, and
- to relate Titan performance results with the workflow of IPRO eCapture.

This benchmark will also measure how IPRO eCapture and Titan respond in a variety of scalability scenarios.

1.0 Data Set & Test Equipment

1.1 Data Set

No standard data set for testing of electronic discovery processing solutions currently exists. There are industry heuristics that indicate what types of files are used in typical forensic data collections, typical data sets are based on how many pages per document per document type are typical, how many documents per gigabyte, how many gigabytes can be processed per day, etc. None of this is well defined. The lack of standard data sets that can be used to evaluate various systems and software architectures was a principal motivator for the development of this benchmark.

The distribution of file types in the dataset was jointly defined by Dutton, LLC and IPRO Tech and was then generated by IPRO. This data set is designed to be representative of collected ESI in a moderate-size electronic discovery project (45 Gigabytes) and was designed to yield approximately 1 million items and over 3 million pages. The data set was organized into 9 custodian directories of approximately 4-6 GB each. This benchmark uses this data set exclusively throughout.

1.2 Test Equipment

A summary of hardware and software used for this benchmarking follows:

Hardware		
IPRO eCapture Worker	One Dell R900 Server	• Four Quad-core Intel CPU at 2.33 GHz – 64 GB RAM
	One Dell R900 Server	• Four Quad-core Intel CPU at 2.33 GHz – 24 GB RAM
	Two Dell 2950	• Two Dual-Core Intel CPU at 2.0 GHz – 8 GB RAM
	One Supermicro 2U Server	• Two Dual-Core Intel CPU at 2.5 GHz – 16 GB RAM
SQL Server	One Dell 2900	• Dual Quad-Core Intel CPU at 3.16 GHz – 32 GB RAM
	One Dell 2950	• Dual Core Intel CPU at 2.0 GHz – 24 GB RAM
Storage	Titan 3201 Storage System	• 2ea RC16 arrays, each with 64 300GB 15K disks
Software		
ESX Server	3.5	
Application Worker	Window 2003 Server SP1 32 bit	• 2 VCPU • 2 GB RAM
SQL Server	SQL Server 2003 x64 SP2	• SQL RAM 28GB

2 Benchmark Methodology

The application used for this benchmark was the e-discovery software IPRO eCapture, version 3.3.0. The primary features of the software are divided into four job types: "Discovery", "Data Extraction", "Processing" and "Exporting". To understand workload characteristics for each job, all four jobs were tested and performance numbers collected.

Based on these tests, Discovery jobs were chosen as the most I/O-intensive jobs that would put the most stress on the storage environment, and were therefore used for the entire benchmark test.

Two approaches were used to benchmark the performance of IPRO eCapture and the Titan storage system: scale up and scale out. The scale-up approach utilized additional IPRO Workers to increase the workload and performance on a single SQL database. The scale-out approach utilized additional SQL servers and databases to distribute more data and tasks across all IPRO Workers while at the same time increasing the workload on the Titan.

3.1 Software configuration

IPRO eCapture is an e-discovery platform that uses a client-server model comprised of the following modules:

- Controller – Job creation, preparation and monitoring
- Queue Manager – Task distribution and coordination amongst all connected Workers
- Worker – Performs all required work as determined by task definitions. Examples of such work include: extracting metadata and text, converting electronic files to TIFF, JPG, or PDF, and creating word index files used for keyword searching.
- SQL Server – Primary communication method between all IPRO eCapture components, and manages SQL databases that store application settings, file information and file relationships.
- Storage – Used to access discovery directories and store application output (extracted files, indexes, images and text).

The hardware and software outlined above support the modules either directly (the software is installed on the hardware) or indirectly (the software uses the hardware to function normally).

The primary functions being performed during the Discovery portion of IPRO eCapture (a "Discovery Job") are:

- Discover all files in specified directories, including subdirectories.
- Extract embedded items (if necessary) from mail stores, email messages, archives, and other compound documents; put the files in a separate location on disk and record their document relations in the database.
- Index the contents and metadata of each item, for the purposes of filtering on responsive search terms; indexes are stored on disk.

The Discovery portion of IPRO eCapture performs several operations designed to facilitate and increase speed in other aspects of the processing workflow. Conversion to tiff and export are enabled for distributed processing due to the groundwork operations performed within the discovery phase.

3.2 Storage configuration

The following is the configuration of the Titan storage system:

- Each RAID group is RAID5, 4 data and one parity drive, with 64K block size.
- One hot spare per enclosure.
- Three storage spools were created:
  - App_pool: 12 RAID groups (Titan system drives). One 900 GB file system was created to serve application data, which was expanded as necessary.
  - SQLlog: 4 system drives. 2 iSCSI LUN were created for database logging. The size of each log was 150G and 200G each.
  - SQLdata: 8 system drives. 4 iSCSI LUN were created for base tables and indexing. The size of each LUN was 400G, 500G, 150G and 50GB each.
- Data Networking – Two aggregation port groups were created. Group One had 3ea Gigabit Ethernet ports configured as a VLAN to serve the application data. Group Two had another VLAN configured with 3ea Gigabit Ethernet ports to serve SQL databases and logs.

3 Benchmark Results

For each benchmark test, the Discovery feature of IPRO eCapture was performed on the same data set. To keep accurate comparisons, for every benchmark job fresh databases were created to ensure no pre-existing data in the databases caused additional overhead.

4.1 Scale-Up: Single Database with Variable Number of IPRO Workers

		Total GB (approx)		
Worker Count	Elapsed Time	Before	After	GByte/HR
10W	6:24:55	45	75	11.5
30W	4:45:36	45	75	15.8

Table 1: Each Worker has four threads running at the same time. Data set size increased from initial 45GB to 75GB at the completion of the Discovery job.

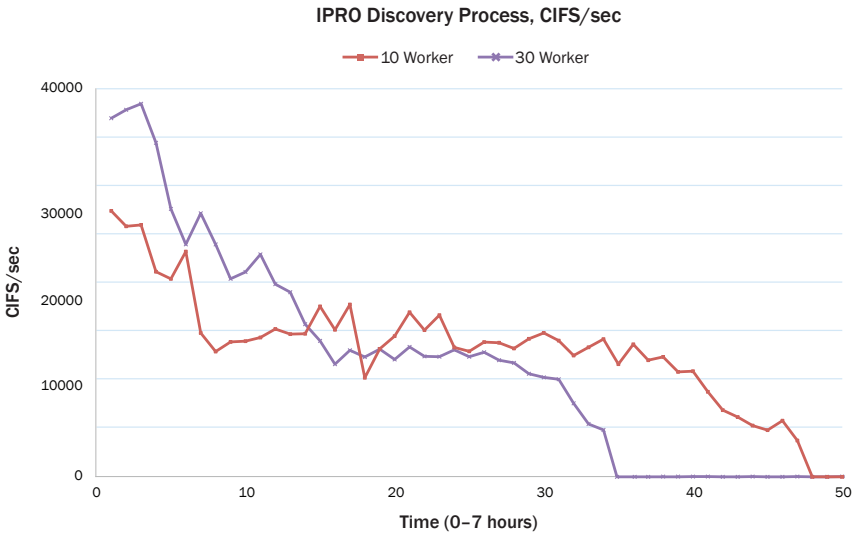

Figure 1 shows the results of scalability from 10 and 30 Workers.

In Figure 1, the data load from the left side of the chart started at approximately 39,000 OP/sec with 30 Workers. It lasted for 30 minutes (10 Workers took longer to complete) and slowed down to 20,000 OP/sec in the middle of the job, then down to 10,000 OP/sec, eventually dropping to zero when the job was complete.

The Discovery function of IPRO eCapture is generally more I/O and CPU intensive on storage devices because existing files are continuously being accessed simultaneous to new files and folders being created and accessed. Additionally, after primary discovery functions complete the files are accessed again to create searchable indexes that are also stored on disk (see Section 3.1). The Data Extract, Processing, and Export functions of IPRO eCapture are more CPU intensive on the IPRO Worker because image and text files are created locally and then copied to a network location on the storage device. These functions are not as I/O intensive on the Titan because files are simply copied from the temporary local location to the network location.

Based on the results shown in Figure 1, it is clear the Titan is capable of handling more than 10 IPRO Workers for this specific data set, reaching only 27,000 OP/sec in performance with this number of multi-threaded Workers. For 30 IPRO Workers, the maximum OP/sec Titan achieved was around 39,000 OP/sec when CPU utilization was up to 100%. Additional IPRO Workers also raised the through-put rate from 11.5 to 15.8 gigabytes per hour, as shown in Table 1. Finally, additional Workers and increased Titan activity led to a faster completion time for the 30-Worker job.

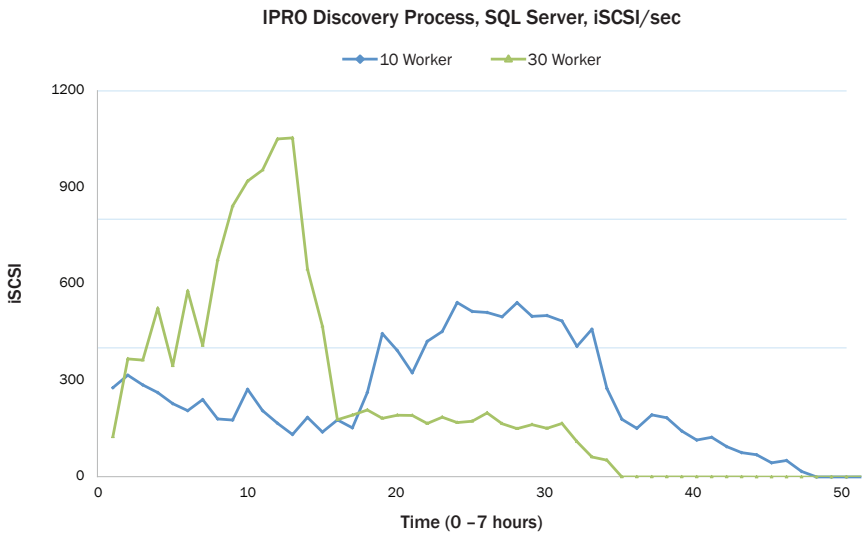

Figure 2: SQL Server database activity.

Figure 2 shows the overall SQL server activity including I/O for the database and log files. An iSCSI LUN was created and used by the SQL server to access the databases. It is important to know any potential bottlenecks in the database as that may degrade performance, especially for SQL logging which is a single thread sequential write workload. With 10 IPRO Workers (slowest completion time), the peak I/O activity started late and completed last. With 30 IPRO Workers, iSCSI OP/sec peaked at over 1000, dramatically decreasing as primary discovery tasks completed; SQL activity remained low for the remainder of the job as indexing tasks completed.

4.2 Scale-Out: Multiple Databases with Variable Number of IPRO Workers

			Total GB (approx)		
DB Instance	Worker	Elapse Time	Before	After	GByte/HR
1	30	4:45:35	45	75	15.8
2	30	6:22:14	90	150	23.6

Table 2: Elapsed time between one and two database instances. Each IPRO Worker has four threads running at the same time. For the two database instance, the same data set was run simultaneously on the two databases (two jobs at 45 GB each) for a total of 90 GB.

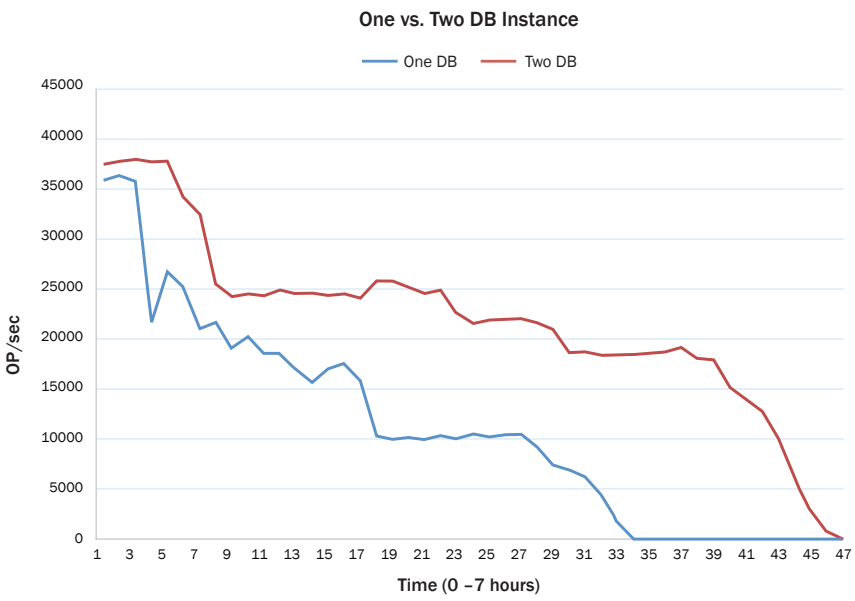

Figure 3. 30 Workers for one database, 30 Workers for two databases.

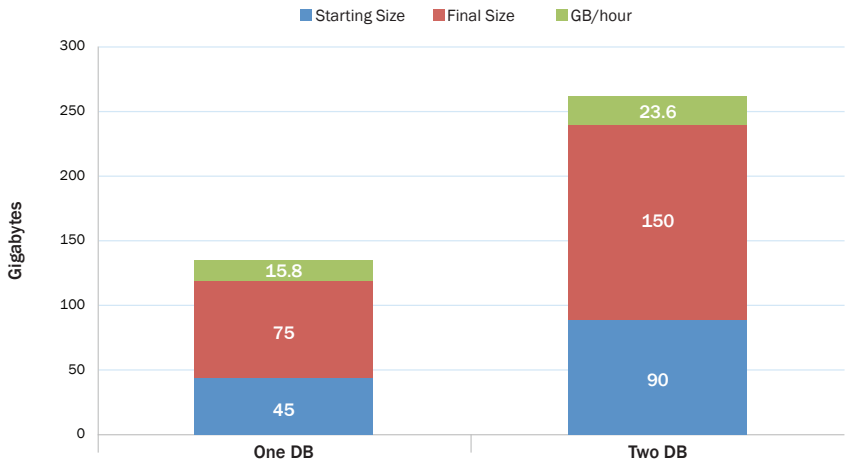

Figure 4. Starting (compressed) and Final (uncompressed) data set sizes for one- and two-database tests, and the calculated GByte/hour from Table 2.

Table 2 shows the elapsed time to completion for Discovery Jobs run on two SQL servers (one database on each SQL server). Both jobs were run simultaneously, with the longest-running job finishing in less than 6.5 hours. A total of 90 GB of data, twice the original amount of data, was discovered with the same number of Workers. Figure 4 shows the original size and the final size of the data sets for each job. These results show that in only 50% more time over the single-database job (less than 2 additional hours), twice the amount of data was discovered. A linear increase would have yielded over 8 elapsed hours instead of 6.5 (actual) elapsed hours. The results also show that data was discovered at a higher rate per hour during the two-database job (see Figure 4). In addition, Figure 3 shows that Titan I/O peaked during the initial discovery tasks and remained high 1.5 times longer than the I/O produced on a single database (30 min for the single database, 45 min for two databases). This is due to duplicate data sets and database activity producing twice the workload on the Titan.

1. Multiple IPRO eCapture jobs running simultaneously on one or many SQL servers and/or databases is a common realistic scenario; see Section 6 for details.

More processing resources on the Titan became available as the Discovery jobs moved toward completion since OP/sec decreases over the duration of the job. Due to hardware availability constraints additional IPRO Workers and SQL servers could not be added for these benchmark tests. Based on the above results, however, the Titan is capable of handling additional IPRO Workers and additional SQL servers and database instances for long periods of sustained activity without seeing any performance degradation. IPRO eCapture performance is not affected by this increased activity as well; according to the results, a set number of Workers processed more data using two database instances, faster than expected[1].

4.3 Titan Utilization and Disk Read/Write Latency

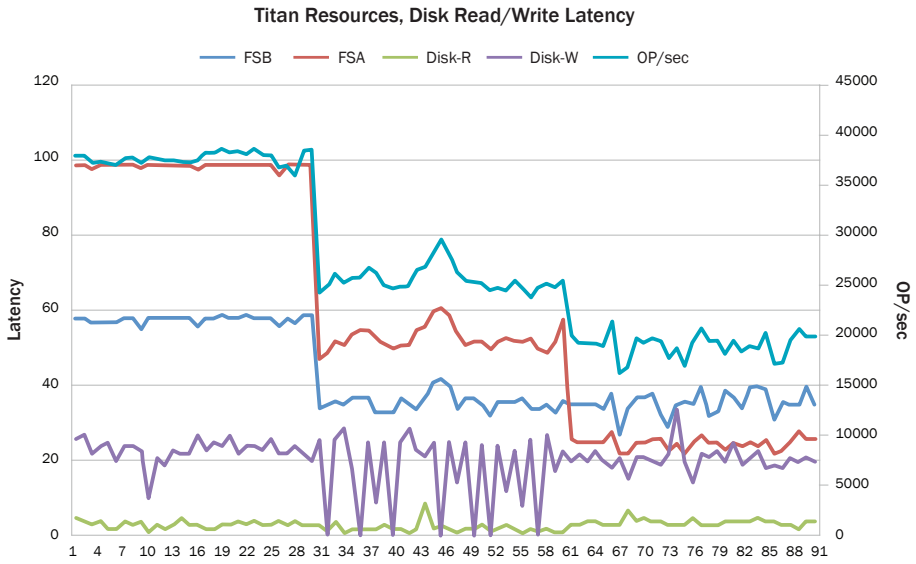

Figure 5: Titan OP/sec, CPU utilization and Disk latency performance on one database with 30 IPRO Workers.

Figure 5 above consists of snapshots from different stages of the 30 Worker test run on a single SQL server and database. The snapshots were collected at stages of the job when OP/sec was at 39K, 26K and 18K OP/sec; Figure 4 shows the results of each variable and how they relate to each other. When OP/sec was in the 39K range, the Titan FSA and FSB were 99% and 60% utilized. When OP/sec was reduced to 26K, both the FSA and FSB were down to the mid 50% and the mid 40% range. At the end of benchmark run, FSB had higher utilization then FSA because of indexing tasks which require more resources of FSB.

In addition to resource utilization, disk I/O latency is also important to measure and probably is the most critical in the entire data path. When a read operation was "cache-missed" in the Titan or storage array controller, it will be accessed from disk which involves a mechanical operation and will degrade performance. When I/O is write-missed, it has to wait for all the "dirty" data de-staging to the disk before operation can be performed.

When the FSA and FSB resources were up to 100%, it is possible that the resources were busy waiting for I/O to complete. However, when disk latency was in the optimal range: read latency was less than 5 milliseconds and write latency was no more than 25 milliseconds throughout the entire benchmark measurement, regardless of CPU utilization. These numbers indicate that IPRO eCapture was not disk bound, and that the Titan is capable of receiving more I/O by adding additional front end load, such as more IPRO Workers and additional SQL servers.

4.4 Command queuing and queue depth on Titan

Command tag queuing is also an important variable in benchmark activity. When commands are queued on the Titan, it has to wait for storage resource availability before it can process the request from the client. In the entire benchmark testing, there was not a high number of commands queuing from the Titan to the storage array. This shows that disk latency is not a bottleneck in this test.

5. Conclusion

The results in this benchmark study show that the Titan storage system is capable of handling extended periods of high activity and workload resulting from running the IPRO eCapture application on an array of Workers. The disk read/write latency is not affected when Titan resource utilization reached 100%, and the maximum performance was up to 40000 OP/sec for the 30-45 minutes of peak time. Results also show that there would be potential for further expansion by either adding more Workers or adding additional SQL servers and/or database instances, all without any performance degradation on the Titan.

6. Notes

In real-world scenarios, multiple Discovery Jobs are run simultaneously on one system, discovering data sets of a size the same or greater than what was discovered for these benchmark tests (45 GB and 90 GB compressed). This, in addition to Process Jobs, Data Extract Jobs and Export Jobs (see Section 3) also running simultaneously, would produce consistent activity and workload on the Titan storage system, as well as increased data through-put per hour through the IPRO eCapture system. Based on the conclusions above, such load would also not negatively affect performance on the Titan. It is also common to have a higher number of IPRO Workers per system, with the largest IPRO eCapture customers implementing 50 to 100 IPRO Workers across a single SQL server or multiple SQL servers. Additional IPRO Workers, as shown in Section 4.1, increases activity on the Titan as well. Improvements to the IPRO eCapture task distribution model, available in the new release of IPRO eCapture (not used for these benchmark tests), will allow the user to make adjustments to task distribution settings to reflect the number of available IPRO Workers connected to the system; this in turn will improve completion times for all job types.

7. Recommendations

A good deal of time was spent configuring a bare-metal Worker environment, then a virtualized Worker environment. It is recommended that the customer properly configure the SQL server and storage system to the potential workload that they intend to run. With these important aspects of the environment set, time can be spent configuring IPRO eCapture to run IPRO Workers on bare-metal machines (a “traditional” configuration) or in a virtualized environment. If virtualization is necessary, it is recommended that separate benchmark tests be performed to determine what impact (if any) a shift from bare-metal to virtualization may have on performance.

There are many potential solutions for scaling IPRO eCapture. Using each individual solution effectively can be complicated. Real benefits are gained when different technologies are combined together to produce a high performance solution for the customer.

About BlueArc

BlueArc is a leading provider of high performance unified network storage systems to enterprise markets, as well as data intensive markets, such as electronic discovery, entertainment, federal government, higher education, Internet services, oil and gas and life sciences. Our products support both network attached storage, or NAS, and storage area network, or SAN, services on a converged network storage platform.

We enable companies to expand the ways they explore, discover, research, create, process and innovate in data-intensive environments. Our products replace complex and performance-limited products with high performance, scalable and easy to use systems capable of handling the most data intensive applications and environments. Further, we believe that our energy efficient design and our products' ability to consolidate legacy storage infrastructures, dramatically increases storage utilization rates and reduces our customers' total cost of ownership.

About IPRO Tech

IPRO offers a diverse suite of superior litigation software tools for e-Discovery, database management, review and production and a complete line of products for both workgroups and high-volume production environments. IPRO has thousands of systems of various sizes in law firms, corporate and government legal departments, and litigation service organizations. Millions of documents are processed using IPRO software annually. The American Lawyer Magazine recently completed a survey of the top 100 law firms in the United States and found that IPRO software is by far the most widely used litigation document software. To learn more about IPRO software tools, visit www.iprotech.com or call 888-477-6463.

BLUEARC®

BlueArc Corporation
Corporate Headquarters
50 Rio Robles Drive
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202