



Strategy | Technology | Results

200 Association Drive  
Charleston, WV 25311  
304 768-1645  
304 768-1671 (fax)



Mr. Frank Whittaker  
Department of Administration  
Purchasing Division  
2019 Washington Street, East  
Charleston, WV 25305

April 10, 2012

Dear Mr. Whittaker:

Please find attached the SIS response to Request for Proposal STO12007 due April 10, 2012. SIS agrees with the requirements of the RFP and is submitting a proposal offering the IBM BladeCenter H with IBM N-Series storage. This proposal meets or exceeds the mandatory requirements of the RFP and we look forward to the opportunity to discuss these solutions with the evaluation committee.

IBM has had the privilege to be the State Treasurer's IT vendor over these past many years and SIS and IBM are uniquely positioned to continue that high level of service.

Thank you for the opportunity to submit the attached responses and we look forward to further discussions

Sincerely,

Charles D. Arnett  
Senior Client Executive  
[carnett@thinksis.com](mailto:carnett@thinksis.com)

Enclosures

RECEIVED

2012 APR 10 PM 1:15

WV PURCHASING  
DIVISION



Strategy | Technology | Results

200 Association Drive  
Charleston, WV 25311  
(304) 768-1645  
(304) 768-1671 (fax)

West Virginia State Treasurer's Office  
Capitol Complex  
Charleston, WV 25305



RFQ:STO12007

10-Apr-12

RFP STO12007 Quotation

Part Number	IBM BladeCenter H for Production and DR	Qty
<b>Rack</b>		
93084PX	IBM 42U Enterprise Rack	2
69P9278	5 Year Onsite Repair 24x7 4 Hour Response	2
39Y8940	DPI 60amp/250V Front-end PDU with IEC 309 2P+Gnd connector	4
39Y8948	IBM DPI C19 Enterprise PDU w/o Line Cord	4
39Y8951	DPI Universal Rack PDU with Nema L5-20P and L6-20P (US line)	4
<b>Networking - 10GbE Switch(s)</b>		
7309G64	IBM System Networking RackSwitch G8264 (rear to front)	4
46C3447	IBM BNT 10GBASE-SR SFP+ Transceiver	8
81Y1618	IBM SFP RJ45 Transceiver	20
39Y7937	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	8
91Y6560	5 Year Onsite Repair 24x7 4 Hour Response	4
<b>BladeCenter 'H' Chassis</b>		
88524TU	IBM eServer BladeCenter(tm) H Chassis with 2x2980W PSU	2
46M0902	IBM UltraSlim Enhanced SATA Multi-Burner	2
2019A1X	IBM BladeCenter KVM/Advanced Management Module	2
40K5581	3m Blue Cat5e Cable	4
68Y6601	IBM BladeCenter H 2980W AC Power Modules w/Fan Pack	2
25R5783	4.3m Double 30A NEMA L6-30P (208V)	0
25R5784	4.3m, 230V, Dual 32A IEC 309 P+N+G/16AIEC320-C20	4
25R5785	2.8m, 200-240V, Triple 16A IEC 320-C20	0
2019B1X	IBM BladeCenter Open Fabric Manager Basic	2
4812S3X	BladeCenter Open Fabric Manager - Advanced Upgrade	2
69P9208	5 Year Onsite Repair 24x7 4 Hour Response	2
<b>1GbE Ethernet Switch Module</b>		
32R1860	BNT Layer 2/3 Copper Gb Ethernet Switch Module for IBM BladeCenter (VM Mgmt Only)	4
40K5581	3m Blue Cat5e Cable	8
<b>10GbE Ethernet Switch Module</b>		
46C7191	BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter	4
46C3447	IBM BNT 10GBASE-SR SFP+ Transceiver	8
90Y9430	IBM Passive DAC SFP+ cable, 3M (Switch Module to G8264)	16
<b>8Gb SAN Switch Module</b>		
42C1828	Brocade Enterprise 20-port 8 Gb SAN Switch Module for IBM BladeCenter	4
44X1962	Brocade 8 Gb SFP+ SW Optical Transceiver	24
39M5697	5m Fibre Optic Cable LC-LC	14
<b>HS23 Blades: 12 Cores / 192GB / Dual 1GbE / Dual 10GbE &amp; Dual 8Gb HBA's</b>		
7875B1U	HS23, Xeon 6C E5-2620 95W 2.0GHz/1333MHz/15MB, 4x4GB, O/Bay 2.5in SAS	14
81Y9295	Intel Xeon 6C Processor Model E5-2620 95W 2.0GHz/1333MHz/15MB	14
46C0599	16GB (1x16GB, 2Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP RDIMM	112
46C0588	8GB (1x8GB, 2Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP RDIMM	112
46M6140	Emulex 8Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter	14
00D4688	IBM SN Distributed Virtual Switch V5000 V1.x, 5 Yr S&S	28
95Y1115	IBM Systems Director Standard Ed for x86 V6-Srvr Lic w/3 Yr S&S	28
41Y8300	IBM USB Memory Key for VMware ESXi 5.0	14
00A4382	5 Year Onsite Repair 24x7 4 Hour Response	14





State of West Virginia  
Department of Administration  
Purchasing Division  
2019 Washington Street East  
Post Office Box 50130  
Charleston, WV 25305-0130

# Request for Quotation

RFQ NUMBER
STO12007

PAGE
1

ADDRESS CORRESPONDENCE TO ATTENTION OF:
FRANK WHITTAKER 304-558-2316

RFQ COPY  
TYPE NAME/ADDRESS HERE

Software Information System, LLC  
200 Association Drive, Suite 210  
Charleston, WV 25311

STATE TREASURER  
MAIN CAPITOL BUILDING  
SUITE E-145  
CHARLESTON, WV  
25305 304-343-4000

DATE PRINTED	TERMS OF SALE	SHIP VIA	F.O.B.	FREIGHT TERMS
03/06/2012				

BID OPENING DATE: 04/10/2012 BID OPENING TIME 01:30PM

LINE	QUANTITY	UOP	CAT. NO.	ITEM NUMBER	UNIT PRICE	AMOUNT
0001	1	LS		205-05		
COMPUTER NETWORK INFRASTRUCTURE						
REQUEST FOR PROPOSAL (RFP)						
THE WEST VIRGINIA PURCHASING DIVISION, FOR THE THE WEST VIRGINIA STATE TREASURER'S OFFICE, IS SOLICITING PROPOSALS FOR REPLACEMENT OF NETWORK INFRASTRUCTURE COMPONENTS PER THE ATTACHED SPECIFICATIONS.						
MANDATORY PRE-BID						
A MANDATORY PRE-BID WILL BE HELD ON 3/16/12 AT 10:30 A M RM EB96, BLDG. 1, 1900 KANAWHA BLVD. E. CHARLESTON. ALL INTERESTED PARTIES ARE REQUIRED TO ATTEND THIS MEETING. FAILURE TO ATTEND THE MANDATORY PRE-BID SHALL RESULT IN DISQUALIFICATION OF THE BID. NO ONE PERSON MAY REPRESENT MORE THAN ONE BIDDER.						
AN ATTENDANCE SHEET WILL BE MADE AVAILABLE FOR ALL POTENTIAL BIDDERS TO COMPLETE. THIS WILL SERVE AS THE OFFICIAL DOCUMENT VERIFYING ATTENDANCE AT THE MANDATORY PRE-BID. FAILURE TO PROVIDE YOUR COMPANY AND REPRESENTATIVE NAME ON THE ATTENDANCE SHEET WILL RESULT IN DISQUALIFICATION OF THE BID. THE STATE WILL NOT ACCEPT ANY OTHER DOCUMENTATION TO VERIFY ATTENDANCE. THE BIDDER IS RESPONSIBLE FOR ENSURING THEY HAVE COMPLETED THE INFORMATION REQUIRED ON THE ATTENDANCE SHEET. THE PURCHASING DIVISION AND THE STATE AGENCY WILL NOT ASSUME ANY RESPONSIBILITY FOR A BIDDER-S FAILURE TO COMPLETE THE PRE-BID ATTENDANCE SHEET. IN						

SEE REVERSE SIDE FOR TERMS AND CONDITIONS

SIGNATURE	Charles D. Arnett	TELEPHONE	304-768-1845	DATE	4-10-2012
TITLE	St. Clerk Executive	FAX	61-1371695	ADDRESS CHANGES TO BE NOTED ABOVE	

WHEN RESPONDING TO RFQ, INSERT NAME AND ADDRESS IN SPACE ABOVE LABELED 'VENDOR'



State of West Virginia  
Department of Administration  
Purchasing Division  
2019 Washington Street East  
Post Office Box 50130  
Charleston, WV 25305-0130

# Request for Quotation

RFQ NUMBER

STO12007

PAGE

5

ADDRESS CORRESPONDENCE TO ATTENTION OF:

FRANK WHITTAKER  
304-558-2316

RFQ COPY

TYPE NAME/ADDRESS HERE

V  
E  
N  
D  
O  
R

S  
H  
I  
P  
T  
O

STATE TREASURER  
MAIN CAPITOL BUILDING  
SUITE E-145

CHARLESTON, WV  
25305

304-343-4000

DATE PRINTED	TERMS OF SALE	SHIP VIA	F.O.B.	FREIGHT TERMS
03/06/2012				

BID OPENING DATE: 04/10/2012 BID OPENING TIME 01:30PM

LINE	QUANTITY	UOP	CAT. NO.	ITEM NUMBER	UNIT PRICE	AMOUNT
------	----------	-----	----------	-------------	------------	--------

PLEASE PROVIDE A FAX NUMBER IN CASE IT IS NECESSARY  
TO CONTACT YOU REGARDING YOUR BID:

304-768-1621

CONTACT PERSON (PLEASE PRINT CLEARLY):

Charles D. Arnett

\*\*\*\*\* THIS IS THE END OF RFQ STO12007 \*\*\*\*\* TOTAL:

SEE REVERSE SIDE FOR TERMS AND CONDITIONS

SIGNATURE TELEPHONE DATE

TITLE FEIN ADDRESS CHANGES TO BE NOTED ABOVE

WHEN RESPONDING TO RFQ INSERT NAME AND ADDRESS IN SPACE ABOVE LABELED "VENDOR"



# Attachment E: Purchasing Affidavit

RFQ No. STO12007

STATE OF WEST  
VIRGINIA Purchasing  
Division

## PURCHASING AFFIDAVIT

West Virginia Code §5A-3-10a states: No contract or renewal of any contract may be awarded by the state or any of its political subdivisions to any vendor or prospective vendor when the vendor or prospective vendor or a related party to the vendor or prospective vendor is a debtor and the debt owed is an amount greater than one thousand dollars in the aggregate.

### DEFINITIONS:

"Debt" means any assessment, premium, penalty, fine, tax or other amount of money owed to the state or any of its political subdivisions because of a judgment, fine, permit violation, license assessment, defaulted workers' compensation premium, penalty or other assessment presently delinquent or due and required to be paid to the state or any of its political subdivisions, including any interest or additional penalties accrued thereon.

"Debtor" means any individual, corporation, partnership, association, limited liability company or any other form or business association owing a debt to the state or any of its political subdivisions. "Political subdivision" means any county commission; municipality; county board of education; any instrumentality established by a county or municipality; any separate corporation or instrumentality established by one or more counties or municipalities, as permitted by law; or any public body charged by law with the performance of a government function or whose jurisdiction is coextensive with one or more counties or municipalities. "Related party" means a party, whether an individual, corporation, partnership, association, limited liability company or any other form or business association or other entity whatsoever, related to any vendor by blood, marriage, ownership or contract through which the party has a relationship of ownership or other interest with the vendor so that the party will actually or by effect receive or control a portion of the benefit, profit or other consideration from performance of a vendor contract with the party receiving an amount that meets or exceeds five percent of the total contract amount.

**EXCEPTION:** The prohibition of this section does not apply where a vendor has contested any tax administered pursuant to chapter eleven of this code, workers' compensation premium, permit fee or environmental fee or assessment and the matter has not become final or where the vendor has entered into a payment plan or agreement and the vendor is not in default of any of the provisions of such plan or agreement.

Under penalty of law for false swearing (*West Virginia Code §61-5-3*), it is hereby certified that the vendor affirms and acknowledges the information in this affidavit and is in compliance with the requirements as stated.

### WITNESS THE FOLLOWING SIGNATURE

Vendor's Name: Software Information Systems, LLC (SIS)

Authorized Signature: [Signature] Date: 3/27/2012

State of Kentucky

County of Fayette, to-wit:

Taken, subscribed, and sworn to before me this 27th day of March, 2012.

My Commission expires 10/26/2014, 2012.

AFFIX SEAL HERE



NOTARY

Darrie C. Cunningham  
State at Large



Strategy | Technology | Results

## Proposal: IBM BladeCenter H with IBM N-Series Storage



**State of West Virginia  
State Treasurer  
Main Capitol Building  
Suite E-145  
Charleston, WV 25305**

**RFP Number STO12007  
Computer Network Infrastructure  
April 10, 2012**

**SIS Contact Information:**  
Charles D. Arnett  
Senior Account Executive  
Software Information Systems, LLC  
200 Association Drive, Suite 210  
Charleston, WV 25311  
[carnett@thinksis.com](mailto:carnett@thinksis.com)  
Telephone: (304) 768-1645  
Fax: (304) 768-1671

**Prepared for:**  
State of West Virginia  
Department of Administration  
Purchasing Division  
Building 15  
2019 Washington Street, East  
Charleston, WV 25305-0130  
ATTN: Frank Whittaker  
Telephone: (304) 343-4000

CONSTRUCTION

CORPORATE

EDUCATION

ENTERTAINMENT

FINANCIAL

GOVERNMENT

HEALTHCARE



**1**

**RFP Document**

**2**

**Attachment A**

**3**

**Attachment B**

**4**

**IBM BladeCenter H**

**5**

**IBM N-Series Storage**



## PROJECT SPECIFICATIONS

- 2.1 **Location:** The Treasurer's Office is located at 1900 Kanawha Blvd. East, Bldg. 1 Suite E-145, Charleston, WV 25305

Delivery locations for equipment are as follows:

Primary site:

1900 Kanawha Boulevard East  
Capitol Complex, Building 6  
Charleston, WV 25305

Inside delivery required, exact location to be specified later.

Business Continuance Site:

89 Richard D. Minnich Drive  
Sutton, WV 26601

Inside delivery required, exact location to be specified later.

All freight and handling is to be included in the proposed cost and not billed as a separate item on the invoice(s); FOB Destination, Freight Prepaid are the preferred terms. Delivery should be made within 30 days of contract award.

### 2.2 **Background and Current Operating Environment:**

The West Virginia State Treasurer's Office (WVSTO) purchased an IBM Blade center in late 2003 and it was installed in early 2004. The WVSTO decided at that time to virtualize all physical servers into the new blade center. The WVSTO converted 36 servers at that time. During the purchasing process it was determined to have a storage allowance of 150% of current needs by the office. The WVSTO has strived to keep all systems virtualized since our original conversion, and due to the natural growth of the office and new programs and initiatives within the office, the total number of virtualized servers has increased to 80 and our current storage needs stands at 8 TB of usable space.

When the WVSTO purchased the equipment in 2003 there was also a need to replicate the data and functions of the office to an off-site location. Currently that location is in Flatwoods, WV which is located roughly 70 miles from the main data center located at the Capital complex. The WVSTO currently uses Veeam to perform all the replication and snap-shots to the DRIBC location, based on multiple schedules. The WVSTO had attempted to use a block level replication and found that a large majority of the servers were not functional or stable at the DR/BC location. The WVSTO also uses LTO 3 tape library and a secondary backup for all data on the blade center.

The WVSTO has been very satisfied with the performance, reliability, and support of the current system. However, as with all technology, at some point it must be replaced. With this in mind, the WVSTO is looking for solutions to our current problems, concerning storage and server performance. This RFP will be concerned with the performance needs and concerns of this office.

In an effort to help the vendor community and the WVSTO better assess the issues and problems with the current, dated solution, a solicitation was released in August of 2011 for services to be rendered in performing a health check of both the current SAN and server solution at the WVSTO. The solicitation



language is attached as Appendix B-1 of the document, and the results of the study are attached as an overall graph view as Appendix B-2 of this RFP.

This health check included an assessment of our current virtual infrastructure and storage array during which metrics were collected over the course of several weeks and a baseline was generated for the current performance profile of the WVSTO data center. This assessment found peak IOPS of over 4000 on our production array with a maximum throughput of about 185,000 KB/s. The storage array was noted in the assessment as being at capacity for IOPS performance and noted as a potential bottleneck that may be preventing the virtual environment from achieving higher IOPS and data throughput. This data along with other internal metrics, and the short and long term goals, of the WVSTO were the basis for the development of the technical requirements outlined later in this document.

The following information is provided to give prospective vendors and overview of the functionality the WVSTO is seeking out of the infrastructure acquired through this RFP process as well as a high level overview of how the WVSTO plans to utilize this infrastructure in our environment.

The WVSTO is a mostly virtualized environment running on VMware vSphere 4.0 with a total of about 90 virtual machines. The WVSTO plans to migrate or rebuild the existing virtual machines on the new infrastructure utilizing VMware vSphere 5.0. The WVSTP is also piloting VMware View 4.6 and plans to deploy approximately 60 virtual desktops on the new infrastructure utilizing VMware vSphere 5.0 and VMware View 5.0.

The unified storage array proposed as part of this RFP must be able to support peak loads of at least 20,000 IOPS for the virtualized server environment. The proposed storage array must provide redundant 4 Gbps fibre channel connections (8 Gbps connectivity preferred) to redundant 8 Gbps capable fibre channel switches. This will provide the vSphere hosts adequate bandwidth and I/O to meet the current and projected demands of the virtual server environment.

The unified storage array must provide additional capacity to support at least 60 virtual desktops with the potential to grow up to 100 virtual desktops. This is estimated to require an additional 8,000 IOPS.

The unified storage array must also provide additional capacity to be capable of hosting CIFS file shares consisting of approximately 3 TB of data with concurrent access by up to 120 users in the WVSTO without impact to either the virtual server or virtual desktop environments.

The WVSTO plans to build two clusters of computer resources running vSphere 5.0 to separate the virtual server and virtual desktop workloads. The WVSTO has a need for a minimum of 4 hosts for the virtualized server infrastructure and a minimum of 2 hosts for the virtualized desktop infrastructure. There is also a need to have at least 1 additional spare host for emergency hardware replacement or unexpected growth of the environment.

The WVSTO is seeking a mirrored solution to provide similar capacity and capabilities at both the production data center and the disaster recovery data center. This will allow the WVSTO to sustain business operations in the event of a long term outage at the primary data center.

The first exception to the mirrored requirement is that no solid state drives or flash drives will be required for the disaster recovery storage array. The disaster recovery storage array may utilize a mix of FC/SAS and ATA/NL-SAS to meet capacity requirements while still providing acceptable performance (approximately 60% of production IOPS) in the event of a fail-over to the DR site.



The second exception to the mirrored requirement is that the WVSTO requires at least 10 TB of additional capacity on the DR array. This additional capacity should be approximately an 80/20 split of lower speed (7.2K) and higher speed (10K or 15K) disks respectively. This extra storage should be in a RAID-S or RAID-6 configuration.

The WVSTO plans to utilize array-based, asynchronous replication with VMware Site Recovery Manager to facilitate the protection and fail-over/fail-back operations of the virtual machines critical to business operations.

The storage array included in the solution (as noted above) must integrate with VMware vSphere and VMware Site Recovery Manager to enable this plan to succeed. The storage array must also include all licensing required to enable replication and protection of data both locally within the array as well as to the remote array at the DR site.

The WVSTO will be moving to 10Gbps Ethernet for the network infrastructure in the vSphere environment. With that goal in mind, the WVSTO would like the NAS functionality of the storage array to utilize 10Gbps Ethernet although 1Gbps will be acceptable.

There are currently no plans to utilize the Ethernet network for storage traffic via either iSCSI or FCoE but support of these protocols within the Ethernet switches is desired to allow the WVSTO the flexibility to utilize these protocols in the future.

The WVSTO is currently not utilizing data-at-rest encryption on the storage array, but this capability has been discussed and may be implemented in the future. As such the proposed storage array should be able to support this feature.

The WVSTO is requesting at least 25 TB of useable capacity in the proposed unified storage array for the production site. This should be split across at least three physical or virtual storage pools configured to meet the specific performance and capacity requirements of their role as outlined above.

The first storage pool will be utilized for the vSphere cluster hosting the virtualized server environment. The second storage pool will be utilized for the vSphere cluster hosting the virtualized desktops for end users. The third storage pool will be utilized for the NAS functionality of the array to present file shares to end users.

The estimated minimum capacity requirements for each storage pool are approximately 17 TB for virtualized servers, 5 TB for file shares (NAS) and 3 TB for virtualized desktops.

The proposed storage array should have the capability to expand to at least double the initial capacity through the addition of drive trays and drives; no storage controller hardware upgrades or license upgrades should be required. Any additional costs to upgrade capacity such as license costs for array features such as data replication or snapshots should be detailed in the proposal.

A glossary of terms is provided as Appendix A just as a matter of reference.

## **2.3 Qualifications and Experience:**

Vendors shall provide in **Attachment A: Vendor Response Sheet**, information as follows:

2.3.1 An organization chart identifying the Vendor's overall business structure and locations, including an explanation of the various services offered by the company.



2.3.2 A minimum of three (3) references is requested. At least one (1) of these references should be from the public sector. All references should be from accounts of a similar scope and complexity as the project outlined in this RFP and include telephone number and email address.

2.3.3 Vendor should provide resumes of proposed project team members which provide adequate combined experience in completing similar projects; include copies of any staff certifications or degrees applicable to this project. .

2.3.4 Vendor response should provide a minimum of two (2) successful projects related to the project outlined in this RFP. The referenced projects should have a successfully completed delivery and implementation. Projects that are in process, but not completed, may be used as options. The Vendor should have had primary responsibility (not acting as a sub-contractor) for the various phases of the projects including: analysis, project/process design, pilot/test phases, and implementation. Vendor should clearly include the description of past projects completed entailing the location of the project, project manager name and contact information, type of project, and what the project goals and objectives were and how they were met.

2.3.5 Vendor should identify any and all subcontractors that will be involved in the delivery and ongoing support of this procurement. The primary vendor will be responsible for any and all work performed by the subcontractors.

## 2.4 Project and Goals:

The project should be a turnkey offering of all equipment, related software, applicable installation and training if/as needed. The WVSTO is aware that a single vendor may not be capable of meeting all goals and objectives. Vendors may elect to subcontract certain services. In such an event, the vendor will be solely responsible for all work performed under this contract, and will assume prime contractor responsibility for all services offered and products to be delivered under the terms of the contract. The State will consider the Vendor to be the sole point of contact with regard to all contractual matters. The Vendor may, with the prior written consent of the State, enter into written subcontracts for performance of work under the contract; however, the Vendor is responsible for payment of all subcontractors. Information/response must be included above, Section 2.3.5.

The WVSTO is aware that there may be multiple solutions that may be proposed as described in section 2.4.2, such as use of Rack Mount Servers or Blade Servers. The WVSTO desires the best solution to meet its current and future needs. If a Vendor plans to submit more than one solution, they may do so but it must be marked accordingly and prepared separately as each solution will need to be evaluated on its own capability and costs.

2.4.1 The following sets of questions are explanation based, concerning the Unified Storage Array that may be proposed. Each question should be responded to on Attachment A. Some questions will require a "yes or no" response while others will require a more detailed response on Attachment A.

- a. Does the array have 8Gbps Fibre Channel connections to the SAN switches?
- b. Does the array have the capability to support 10Gbps FCoE for storage presentation?
- c. Does the array have the capability to support 10Gbps iSCSI for storage presentation?
- d. Does the array support automatic, dynamic read/write memory (cache) allocation?
- e. Does the array support both 2.5" and 3.5" disk drives?
- f. Does the array support RAID 1/0 (striped and mirrored)?
- g. Does the array support RAID 5 (single parity)?



- h. Does the array support RAID 6 or RAID-DP (double parity)?
- i. Can the storage administrator choose which tier of disk in a storage pool is used when creating a new LUN?
- j. Can individual LUNs be expanded "on the fly" without down time on the system using the LUN?
- k. Can individual LUNs be converted from thick to thin provisioned and vice versa?
- l. Does the array support space reclamation on existing thin provisioned LUNs?
- m. Can individual LUNs be manually migrated between disk types in a storage pool without down time on the system using the LUN?
- n. Does the array support the exclusion of particular LUNs from automatic tiering?
- o. Are full copies (clones) of LUNs available for use immediately after initiating clone operation?
- p. Are full copies (clones) mountable by a different host?
- q. Does the replication technology in the array support both local and remote protection?
- r. Does the storage array utilize 10Gbps Ethernet for the NAS (CIFS/NFS) functionality?
- s. Does the storage array support NDMP for backup of raw file systems?
- t. Does the NDMP support allow for file and folder level restoration without the need to restore the entire NAS file system containing those items?
- u. Does the CIFS file server support Windows 2008 R2 native-mode Active Directory domains?
- v. Does the CIFS file server support Volume Shadow Copy to allow end-user or support staff recovery of files and folders using the "Previous Versions" features built into the Windows client operating systems when utilizing snapshot technology on the CIFS file shares?
- w. Does the CIFS file server support ABE (access-based enumeration)?
- x. Does the array support de-duplication of data presented via file protocols (CIFS/NFS)?
- y. Does the array support compression of data presented via file protocols (CIFS/NFS)?
- z. Does the array provide the ability to administer the system via a command line interface (CLI installed on a remote system or direct SSH/telnet interface)?
- aa. Does the array provide the ability to script administrative actions for bulk operations?

2.4.1.1 Does the array support the use of solid state drives (SSD) or enterprise flash drives (EFD) as an extension of read/write cache to enhance performance and alleviate hot spots from sudden, unexpected spikes in workload? If yes, response should provide details on how this works within the array and any limitations of this technology.

2.4.1.2 How many total active paths does a host have to an individual LUN?

2.4.1.3 What is the total bandwidth available for a host to an individual LUN?

2.4.1.4 What is the total number of drives and drive enclosures supported by the array (expandability)?

2.4.1.5 When implementing the data-at-rest encryption does the array provide internal key management system, utilize (or require) an external key management system or rely on drive-based encryption without the need for a key management system?

2.4.1.6 How many tiers of storage (drive types) may be placed in a single storage pool?

2.4.1.7 Does the array support automatic data tiering within a configured storage pool to allow migration of data to higher or lower speed disks based on an activity algorithm? If yes, response should explain how the automatic data tiering works in the proposed unified storage array.



2.4.1.8 Does the array support policies on automatic tiering to allow SAN administrators to designate particular LUNs that should only be migrated to higher (or lower) speed disks? If yes, response should explain how this functionality is implemented in the proposed array.

2.4.1.9 Does the array allow scheduled policy changes for the automatic tiering of individual LUNs based on regularly occurring events? (An example would be the ability to schedule a particular LUN to only be migrated to higher speed disks during a time period that is known to be very busy)

2.4.1.10 Does the array allow dynamic expansion of storage pools through the addition or more drives and/or RAID arrays into the storage pool? If yes, response should provide details on how this feature is implemented and any limitations imposed on this process.

2.4.1.11 Does the array support de-duplication of data contained on LUNs presented via block level protocol (FC/FCoE/iSCSI)? If yes, response should explain this functionality on the proposed array.

2.4.1.12 How many snapshots of a single LUN can be made? Response should include any details on performance degradation when utilizing multiple snapshots on a LUN.

2.4.1.13 Please outline the typical storage requirement for snapshots, both individual and multiple incremental snapshots of the same LUN. Also, response should provide a brief explanation of how snapshot technology is implemented on the array.

2.4.1.14 Does the replication technology in the array have the ability to take multiple snapshots of the LUNs to enable recovery or testing with copies of those LUNs at a user configurable interval? If yes, response should provide details on how this technology is implemented in the array.

2.4.1.15 What are the typical bandwidth requirements of the replication technology after initial seeding of the data to the remote site has been completed?

2.4.1.16 Is the data being replicated compressed or de-duplicated to reduce bandwidth requirements?

2.4.1.17 Is the data being replicated encrypted between the source and destination arrays?

2.4.1.18 Does the replication technology support RPO goals of 15 minutes or less using asynchronous replication to a remote site? Response should detail any bandwidth or latency requirements to meet this goal.

2.4.1.19 Does the array have the capability to serve as a CAS/WORM device to replace optical storage systems? If so, what level of compliance does the CAS functionality provide?

2.4.2 The following sets of questions are explanation based, concerning the Server Hardware that may be proposed. Each question should be responded to on Attachment A. Some questions will require a "yes or no" response while others will require a more detailed response on Attachment A.

- a. Do the proposed servers support 16GB DIMMs?
- b. Do the servers support 32GB DIMMs?
- c. Do the proposed servers contain more than the required minimum of 192GB RAM per server?

2.4.2.1 What is the total available processing power of the servers in the proposed solution? Response should provide a breakdown on core count, core speed and total processing power (GHz) for the proposed servers.

2.4.2.2 How many DIMMs can the servers in the proposed solution hold (without add-ons)?

2.4.2.3 Are add-ons (drawers, trays, add-on blades, etc.) available to increase the number of DIMMs that can be installed in a server? If so, what is the maximum number of DIMMs that can be installed in the servers with any available add-ons?

2.4.2.4 What is the maximum RAM supported by the servers without add-ons (drawers, trays, etc.)? With add-ons?

2.4.2.5 What size and speed DIMMs are being used in the proposed server configuration? Response should provide a detailed description of the RAM layout utilized on the servers.

2.4.2.6 If proposing Rack Mount Servers:

2.4.2.6.1 How many available PCI-Express slots do the servers in the configuration have?

2.4.2.6.2 What is the speed of the PCI-Express slots in the servers? Please provide a detailed listing of the available PCI-e expansion slots and their speeds and note which are already populated.

2.4.2.7 If proposing Blade Servers:

2.4.2.7.1 How many total slots are in the proposed chassis?

2.4.2.7.2 Are the blade chassis in this proposal equipped with all required power supplies, fans and I/O modules/switches to support fully populating the blade chassis without additional cost beyond the purchase of the blade servers?

2.4.2.7.3 How many slots are used by the servers included in this proposal?

2.4.3 The following sets of questions are explanation based, concerning the **Network Switches** that may be proposed. Each question should be responded to on Attachment A. Some questions will require a "yes or no" response while others will require a more detailed response on Attachment A.

a. Due to the core competency of the WVSTO staff as well as other WV state agencies, we would prefer to continue utilizing Cisco networking equipment within our data center for Ethernet connectivity. Does the proposed solution include Cisco network equipment?

b. Does the proposed solution include licenses for VMware distributed virtual switch modules to allow both the physical and virtual network infrastructure to be managed through a common interface (whether command line, browser-based GUI, etc.)?

2.4.3.1 Does the network equipment for server connectivity in the proposed solution have expansion capabilities (port modules, etc.), and, are those expansion slots available for future use or populated as part of the proposed solution? If yes, response should detail the expansion capabilities of the proposed network switches.

2.4.3.2 Does the propose network equipment include, or have the capability to support, other network protocols, specifically FCoE (fibre channel over Ethernet) and iSCSI? If yes, response should outline any additional modules or license costs to enable the support of these protocols on the proposed network switches.



2.4.4 The following sets of questions are explanation based concerning the **General Solution** being proposed. Each question should be responded to on Attachment A. Some questions will require a "yes or no" response while others will require a more detailed response on Attachment A.

- a. Does the proposed solution include a centralized, unified monitoring system that gives overall status information about the hardware included in the solution (switches, storage and servers)?
- b. Does the proposed solution include a single point of contact for all support issues (hardware and software) when utilized to run a vSphere environment?
- c. Does the propose solution include direct OEM support from the vendors of each component utilized in the solution to allow escalation of support issues to the OEM technicians by either the WVSTO or our single point of support for the propose solution?
- d. Does the proposed solution include regular (quarterly or bi-annually), pre-tested and validated firmware updates direct from a single source to allow the WVSTO to keep all hardware in the solution up-to-date without having to go through internal research, testing and validation of firmware as it is released by the OEMs?
- e. The proposed solution should take into consideration existing WVSTO licensing and should only include software licensing that is necessary to support the proposed solution that is not already owned by the WVSTO (see Appendix C for list of current VMware licensing). Have you taken existing WVSTO licensing into account and only included additional licenses, not already owned by the WVSTO in your proposed solution?

2.4.4.1 Does the proposed solution ship as a single unit (all hardware racked, all internal power, network, SAN and other cables connected) ready to connect to power and core networking equipment and begin deployment and configuration of storage, networking and the vSphere environment?

2.4.4.2 The WVSTO would like to keep the network traffic for the hosts, the network traffic for hardware management and the storage network traffic separated. This serves a few purposes, the first being segregation of traffic with dedicated resources for each type of traffic, to try and insure peak performance of the solution; the second being the ability to keep the management traffic on high performance (gigabit), but lower-cost switches that don't need the capabilities of the switches used to connect the VMware hosts to the network.

2.4.4.2.1 Does the propose solution include separate switch infrastructure for the hosts, the hardware management interfaces and storage (fibre channel) networks? If yes, response should provide some details on the internal network layout of the proposed solution and how it meets this goal.

2.4.4.3 Does the proposed solution include a centralized, unified management system that allows baseline configuration tasks to be performed? If it does, can the following tasks be performed through this management system? If so, response should outline the capabilities to perform the following functions:

- 2.4.4.3.1 Define VLANs available (trunked) into the network switches from the core network
- 2.4.4.3.2 Define storage available to the various vSphere clusters

2.4.4.3.3 Deployment of operating system (vSphere, Windows, etc.) to the physical servers included in the solution from user-provided ISO images

2.4.4.3.4 Creation of vCenter instances to manage vSphere hosts

2.4.4.3.5 Does the management system provide any additional capabilities not outlined above? If it does, response should detail any notable capabilities.

## **2.5 Mandatory Requirements**

The following mandatory requirements must be met by the Vendor as a part of the submitted proposal. Failure on the part of the Vendor to meet any of the mandatory specifications shall result in the disqualification of the proposal. The terms "must", "will", "shall", "minimum", "maximum", or "is/are required" identify a mandatory item or factor. Decisions regarding compliance with any mandatory requirements shall be at the sole discretion of the State.

### **2.5.1 Unified Storage Platform**

2.5.1.1 The unified storage systems must allow presentation of storage through block and file level protocols and meet the following requirements for useable capacity.

2.5.1.1.1 The storage array for the production data center must provide a minimum useable capacity of 17 TB for virtualized servers in a dedicated physical or virtual storage pool.

2.5.1.1.2 The storage array for the production data center must provide a minimum useable capacity of 5 TB for NAS file shares in a dedicated physical or virtual storage pool.

2.5.1.1.3 The storage array for the production data center must provide a minimum useable capacity of 3 TB for virtual desktops in a dedicated physical or virtual storage pool.

2.5.1.1.4 The storage array for the disaster recovery data center must provide a minimum useable capacity of 17 TB for replicated virtual servers.

2.5.1.1.5 The storage array for the disaster recovery data center must provide a minimum useable capacity of 5 TB for replicated NAS file shares.

2.5.1.1.6 The storage array for the disaster recovery data center must provide a minimum useable capacity of 3 TB for replicated virtual desktops.

2.5.1.1.7 The storage array for the disaster recovery data center must provide a minimum additional useable capacity of 10 TB.

2.5.1.2 The proposed storage array must be a unified storage array that allows presentation of storage via block (Fibre Channel) and file (CIFS, NFS) protocols.

2.5.1.3 The proposed storage array must have a minimum of 4Gbps fibre channel connectivity to the SAN switch infrastructure.

2.5.1.4 The proposed storage array must have two storage controllers for the block level protocol in an active/active configuration with at least two fibre channel connections to the SAN switch infrastructure providing a total of 4 paths to the storage array.



2.5.1.5 The proposed storage array must have two filers for the file level protocols in an active/passive or active/active configuration with a minimum of two (2) 1Gbps Ethernet connections per filer to the network infrastructure.

2.5.1.6 The proposed storage array for the production data center must provide a minimum of 20,000 IOPs dedicated to the virtualized server environment.

2.5.1.7 The proposed storage array for the production data center must provide a minimum of 8,000 IOPs dedicated to the virtual desktop environment.

2.5.1.8 The proposed storage array for the production data center must provide dedicated capacity to support NAS file shares for up to 120 users and 3 TB of data.

2.5.1.9 The proposed storage array for the disaster recovery site must provide a minimum of 60% of the total IOPs of the production storage array.

2.5.1.10 The proposed storage array must support Solid State Drives (SSD) or Enterprise Flash Drives (EFD) (Tier 0).

2.5.1.11 The proposed storage array must be able to use both 10K RPM and 15K RPM high speed Fiber Channel or Serial Attached SCSI (SAS) drives. The proposed storage array does not have to include both drive types in the configuration. Please use the drives appropriate to meet the capacity and IOP requirements as outlined in this RFP.

2.5.1.12 The proposed storage array must support 7.2K RPM near-line SAS or ATA drives (Tier 3).

2.5.1.13 The proposed storage array must support virtual (thin) provisioning for volumes presented via block level (FC) protocol.

2.5.1.14 The unified storage systems must support the ability to do snapshots and clones of volumes presented via block level protocols. It must also support the ability to do snapshots of the file systems presented via file level protocols.

2.5.1.15 The proposed storage array must include the ability to make clones of volumes presented via block-level (FC) protocol.

2.5.1.16 The proposed storage array must include the ability to take snapshots of volumes presented via block-level (FC) protocol.

2.5.1.17 The proposed storage array must include the ability to take snapshots of file systems presented via file-level protocols (CIFS, NFS).

2.5.1.18 The proposed storage array must include IP-based, asynchronous replication for the storage presented via block level (FC) protocol.

2.5.1.19 The proposed storage array must include IP-based, asynchronous replication for the file systems presented via file level (CIFS, NFS) protocols.

2.5.1.20 The proposed storage array must have the capability to support data-at-rest encryption.

2.5.1.21 The proposed storage array must have a single, unified management tool that allows the configuration and monitoring of all features and functionality of the array.

2.5.1.22 The proposed storage array must support all of the primitives defined in the VMware vSphere API for Array Integration (VAAI) specifications for vSphere 5.0 for storage presented via block level (FC) protocol.

2.5.1.23 The proposed storage array must include full, active-active, load balanced multi-path support for connected VMware vSphere 5.0 hosts (not the default most recently used or round robin provided by VMware).

2.5.1.24 The proposed storage array must include plug-ins for VMware vCenter to enable the creation and management of LUNs (from assigned storage pools) for the vSphere environment to ensure proper alignment and optimization of the LUNs.

2.5.1.25 The proposed array must include replication technology that integrates with VMware Site Recovery Manager (SRM) 5.0 to allow SRM to leverage the native replication technologies of the array to copy data to the disaster recovery site.

2.5.1.26 The proposed array must have the capability to enable call-home functionality for sending hardware alerts to the OEM when failures are detected on the array to enable rapid, pro-active response from technical support to replace or repair defective hardware.

2.5.1.27 The unified storage systems must have an expected product life of at least 5 years.

2.5.1.28 The unified storage systems must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 availability coverage.

## **2.5.2 Fibre Channel Switches**

2.5.2.1 The proposed solution shall include two independent fibre channel switches at each site.

2.5.2.2 The fibre channel switches must have autosensing 8 Gbps ports (support 8/4/2 Gbps).

2.5.2.3 The proposed fibre channel switches must have management capabilities via a command line interface (telnet/SSH).

2.5.2.4 The proposed fibre channel switches must have a browser-based management interface.

2.5.2.5 The proposed fibre channel switches must include some internal diagnostics.

2.5.2.6 The proposed fibre channel switches must include native alerting and reporting (without the need for a monitoring server).

2.5.2.7 The proposed fibre channel switches must include a native way to display performance metrics.

2.5.2.8 The proposed fibre channel switch configuration must support non-disruptive firmware upgrades.

2.5.2.9 The proposed fibre channel switches must have the capability to be either an NPV edge device or an NPIV core device.

2.5.2.10 The proposed fibre channel switches must have the capability to support multiple fabric environments in a single physical switch.



2.5.2.11 The proposed fibre channel switches must support aggregated ISL (inter-switch link) connectivity; i.e., several physical ISLs behaving as one virtual ISL.

2.5.2.12 The proposed fibre channel switches must support traffic engineering using FSPF.

2.5.2.13 The fibre channel switches must have at least 12 ports active each.

2.5.2.14 The fibre channel switches must have at least 24 ports total each.

2.5.2.15 The fibre channel switches must have redundant power supplies and fans.

2.5.2.16 The fibre channel switches must have an expected product life of at least 5 years.

2.5.2.17 The fibre channel switches include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 coverage.

### 2.5.3 Network Switches

2.5.3.1 The network switch(es) must support both 10Gbps and 1Gbps connectivity.

2.5.3.2 The network switch(es) must have a minimum of 16 ports available for connection of additional network devices not included in the proposed solution.

2.5.3.3 The network switch(es) must have redundant power supplies and fans.

2.5.3.4 The network switch(es) used for server connectivity must include layer 3 support (if a dedicated management network is present it does not need to support layer 3).

2.5.3.5 The network switch(es) must support Link Aggregation Control Protocol (LACP): IEEE 802.3ad.

2.5.3.6 The network switch(es) must support VLAN trunking.

2.5.3.7 The network switch(es) must support IEEE 802.1Q VLAN encapsulation.

2.5.3.8 The network switch(es) must support Jumbo Frames on all ports (up to 9216 bytes).

2.5.3.9 The network switch(es) must support CLI management (console, telnet and/or SSH).

2.5.3.10 The network switch(es) must support SNMP.

2.5.3.11 The network switch(es) must have an expected product life of at least 5 years.

2.5.3.12 The network switch(es) must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 coverage.

### 2.5.4 Server Hardware

2.5.4.1 There must be at least 7 identically configured servers per site (production and DR), 14 servers in total.

2.5.4.2 The proposed servers must be dual CPU socket servers.

2.5.4.3 The proposed servers must use 6-core Intel 5600 series or 10-core Intel E7 series processors or superior.

2.5.4.4 Each server must have at least 192GB of RAM.

2.5.4.5 Each server must include a minimum of two (2) 10Gbps network connections.

2.5.4.6 Each server must include a minimum of two (2) 8Gbps fibre channel (SAN) connections.

2.5.4.7 The servers must include remote management capabilities (DRAC, iLO or equivalent).

2.5.4.8 The servers must have fully redundant internal components (power supplies, fans, etc.).

2.5.4.9 The servers must have an expected product life of at least 5 years.

2.5.4.10 The servers must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 coverage.

#### 2.5.5 Rack Mount Servers (If this solution is proposed)

2.5.5.1 All of the PCI-Express slots in the servers must run at a minimum of 4x speed.

2.5.5.2 The servers must have at least two available PCI-Express slots for expansion capabilities.

#### 2.5.6 Blade Servers (If this solution is proposed)

2.5.6.1 In the proposed blade solution the individual blade servers at each site must be split as evenly as possible across two blade chassis (elimination of single point of failure and provide extra expansion capabilities through number of available slots for blades).

2.5.6.2 Each blade chassis must include fully redundant I/O and management modules.

#### 2.5.7 **General Requirements for Proposed Solution**

2.5.7.1 The proposed solution must include the cabinet(s) for the equipment. The cabinet(s) must have locks and provide adequate air flow to keep installed equipment cool.

2.5.7.2 The proposed solution must include all required power distribution units (PDUs) inside the cabinet(s) to provide power to all installed equipment.

2.5.7.3 All equipment must be mounted in the cabinet(s).



2.5.7.4 All wiring inside the cabinet must be complete. This includes power cables from equipment to PDUs and network and fiber interconnects between servers, storage and switches.

2.5.7.5 The proposed solution must include maintenance agreements on the hardware and software included in the proposal for five years. These maintenance agreements must, at a minimum, provide break fix support, software upgrades and general technical support (knowledge base, trouble tickets) on the included components via phone and/or online support sites.

2.5.7.6 The proposed solution must include only the professional services required to get the equipment into our data centers, installed in the cabinet(s) as outlined above and brought online and ready for end-user configuration. The vendor must also perform any configuration on the equipment (for example, the storage pools on the storage array and other array features such as cache settings) required to ensure the equipment meets the capacity and performance requirements outlined in this RFP. Any additional professional services such as configuration of replication features on the arrays, configuration of VMware Site Recovery Manager (SRM) to work with array replication or any assistance with data migration or other tasks will be acquired through a separate procurement(s).

2.6 Oral Presentations (Agency Option): State agencies have the option of requiring oral presentations of all Vendors participating in the RFP process. If this option is exercised, it would be listed in the Schedule of Events (Section 1.3) of this RFP. During oral presentations, Vendors may not alter or add to their submitted proposal, but only clarify information. A description of the materials and information to be presented is provided below:

2.6.1 Materials and Information Required at Oral Presentation:

- a. Vendor should present information about any technologies being offered, as part of the response(s) to the RFP that include technologies that the Agency may not be currently aware of, or being introduced into the market place.
- b. Vendor will have the opportunity to provide insights into technological or other aspects of their proposed solution(s) that may be unique to their solution(s).
- c. Vendor will have the opportunity to explain the differences between multiple proposed solutions (such as described in sections 2.4.2.6 and 2.4.2.7 as well as sections 2.5.5 and 2.5.6) when/if a vendor proposes more than one solution, in response to the RFP.

## Attachment A: Vendor Response Sheet

---

### ***Qualifications and Experience contained in Section 2.3:***

Provide staff qualifications and experience in completing similar projects, as a part of this section the Vendor should provide:

2.3.1 An organization chart identifying the Vendor's overall business structure and locations, including an explanation of the various services offered by the company.

Vendor Response: *SIS has included the requested information as Attachment A-1, attached hereto.*

2.3.2 A minimum of three (3) references is requested. At least one (1) of these references should be from the public sector. All references should be from accounts of a similar scope and complexity as the project outlined in this RFP and include telephone number and email address.

Vendor Response: *SIS has included the requested information as Attachment A-2, attached hereto.*

2.3.3 Vendor should provide resumes of proposed project team members which provide adequate combined experience in completing similar projects; include copies of any staff certifications or degrees applicable to this project.

Vendor Response: *SIS has included the requested information as Attachment A-3, attached hereto.*

2.3.4 Vendor response should provide a minimum of two (2) successful projects related to the project outlined in this RFP. The referenced projects should have a successfully completed delivery and implementation. Projects that are in process, but not completed, may be used as options. The Vendor should have had primary responsibility (not acting as a sub-contractor) for the various phases of the projects including: analysis, project/process design, pilot/test phases, and implementation. Vendor should clearly include the description of past projects completed entailing the location of the project, project manager name and contact information, type of project, and what the project goals and objectives were and how they were met.

Vendor Response: *SIS has included the requested information as Attachment A-4, attached hereto.*

2.3.5 Vendor should identify any and all subcontractors that will be involved in the delivery and ongoing support of this procurement. The primary vendor will be responsible for any and all work performed by the subcontractors.

Vendor Response: *SIS has included the requested information as Attachment A-5, attached hereto.*

### ***Project Goals and Objectives contained in Section 2.4:***

The project should be a turnkey offering of all equipment, related software, applicable installation and training if/as needed. The WVSTO is aware that a single vendor may not be capable of meeting all goals and objectives. Vendors may elect to subcontract certain services. In such an event, the vendor will be solely responsible for all work performed under this contract, and will assume prime contractor responsibility for all services offered and products to be delivered under the terms of the contract. The State will consider the Vendor to be the sole point of contact with regard to all contractual matters. The Vendor may, with the prior written consent of the State, enter into written subcontracts for performance of under the contract; however, the Vendor is responsible for payment of all subcontractors. Information/response must be included above, Section 2.3.5.

The WVSTO is aware that there may be multiple solutions that may be proposed as described in section 2.4.2, such as use of Rack Mount Servers or Blade Servers. The WVSTO desires the best solution to meet its current and future needs. If a Vendor plans to submit more than one solution, they may do so but it must be marked accordingly and prepared separately as each solution will need to be evaluated on its own capability and costs.



## Attachment A-1: Organizational Charts

---

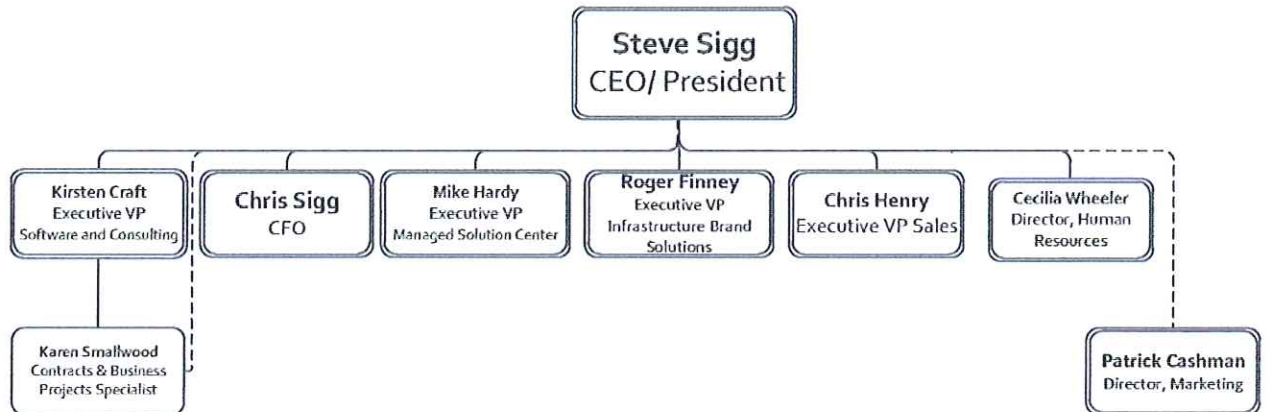
SIS is headquartered in Lexington, KY with regional offices in Louisville and London, KY; Cincinnati, OH; Indianapolis, IN; Charleston, WV; Troy, MI and Alexandria, VA. Our teams deliver a full range of technology infrastructure and managed IT services to help clients achieve measurable business results:

- [Analytics Optimization](#)
- [Business Continuity](#)
- [Data Center & Cloud Services](#)
- [Data and Backup Storage Management](#)
- [Security](#)
- [Software & Applications](#)
- [Virtualization](#)
- [Professional Services](#)

Members of the Executive Team are located in Lexington, KY; Indianapolis, IN; Alexandria, VA; and Troy, MI.



## Executive Team



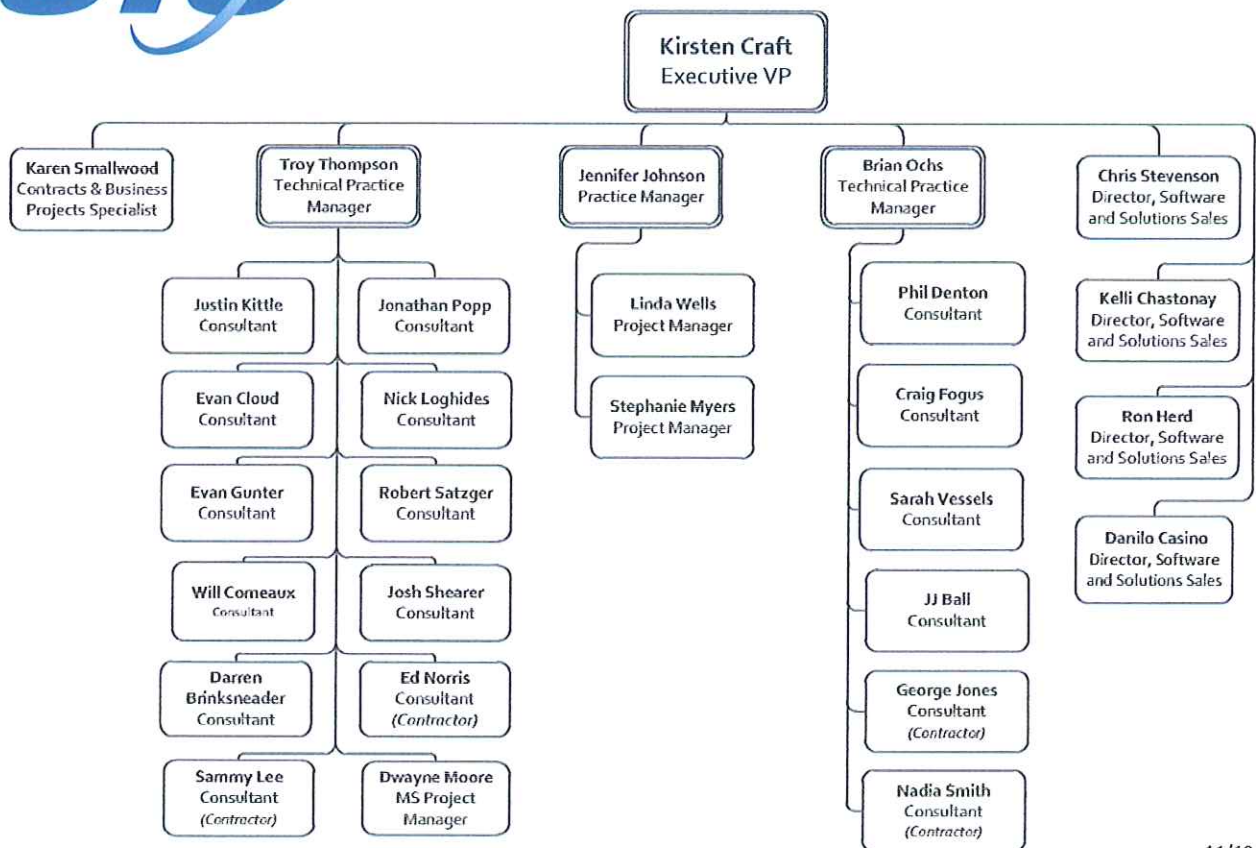
11/10/2011



Members of the Software Services & Consulting team are located in Lexington, Louisville, and London, KY; Cincinnati, OH; and Indianapolis, IN.



## Software Services & Consulting



11/10/2011

Members of the Marketing team are located in Lexington, KY.



## Marketing



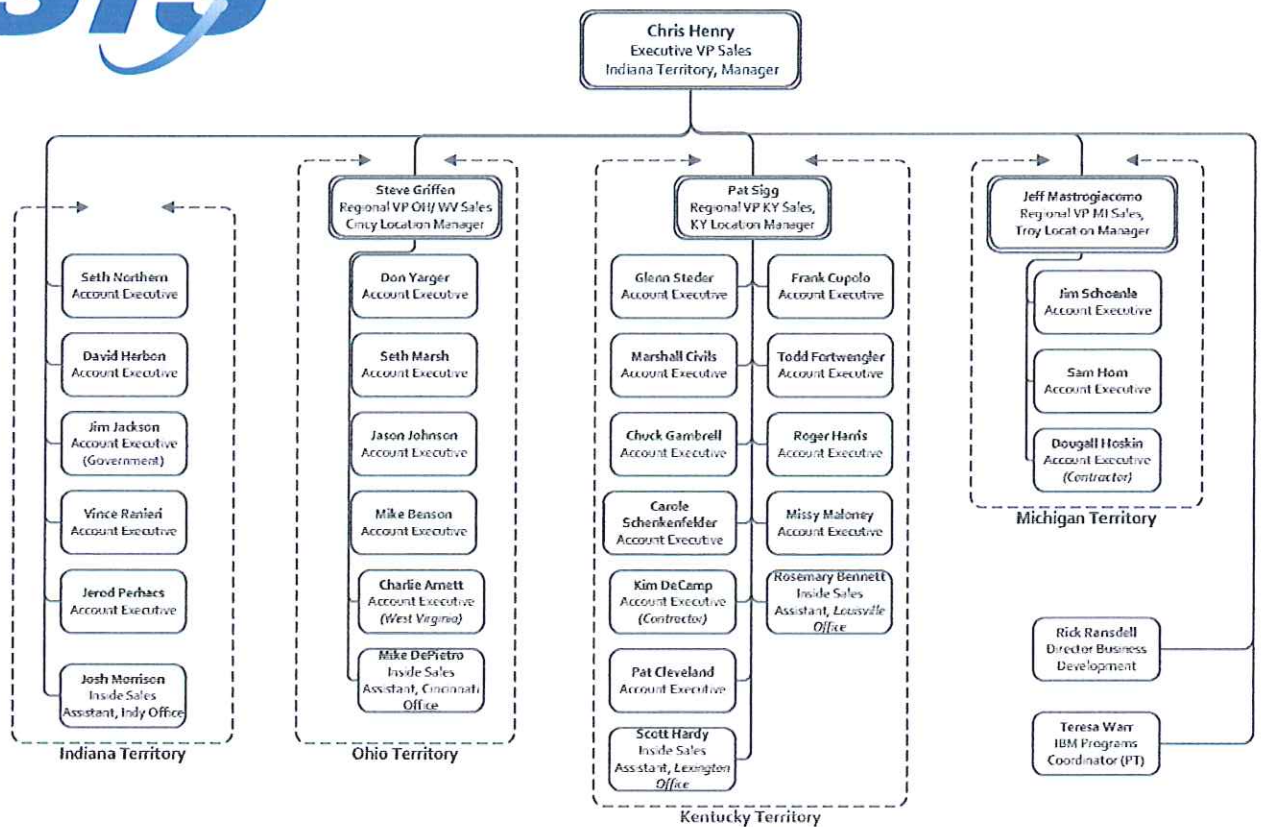
11/10/2011



Members of the Sales team are located in Lexington, Louisville, and London, KY; Cincinnati, OH; Troy, MI; and Indianapolis, IN.



## Sales Location Structure

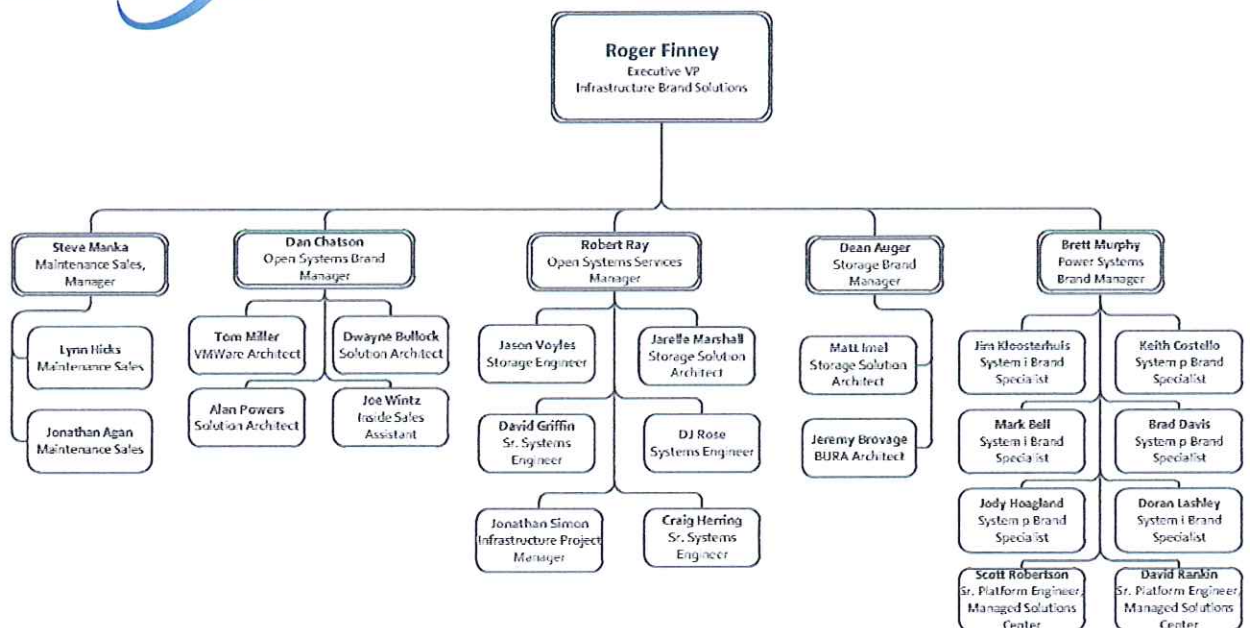


11/10/2011

Members of the Infrastructure Brand Solutions team are located in Lexington and Louisville, KY; Cincinnati, OH; Troy, MI; and Indianapolis, IN.



## Infrastructure Brand Solutions

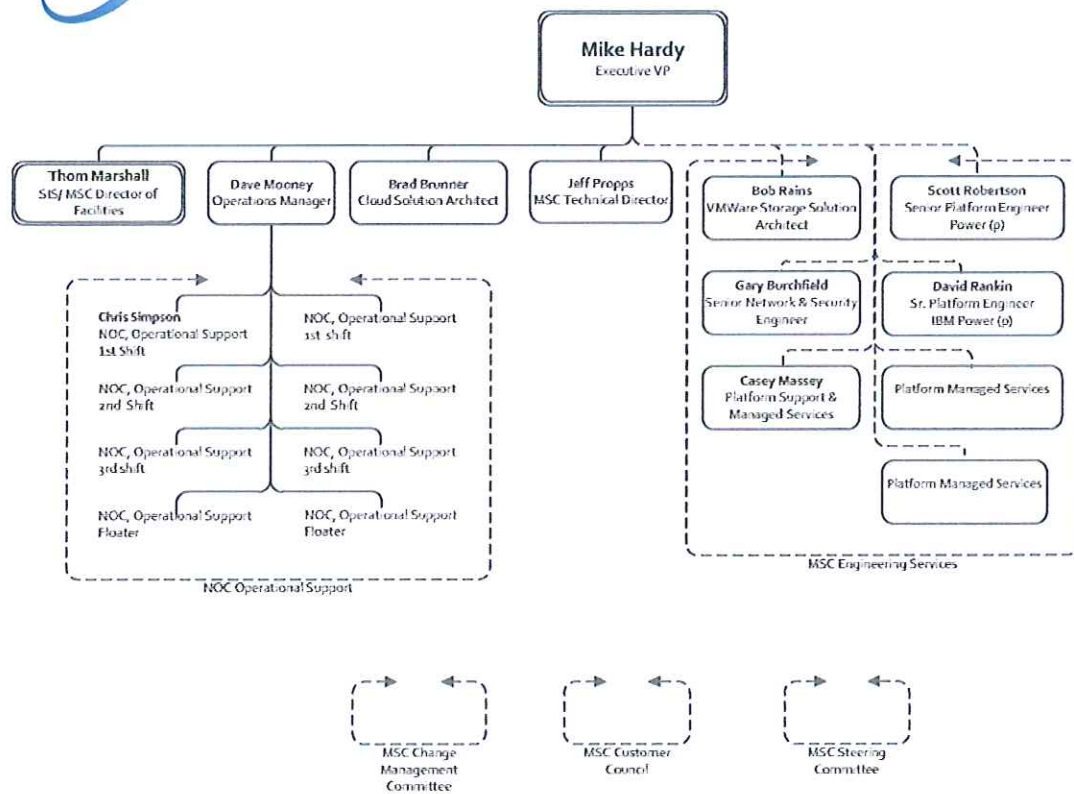




Members of the Managed Solution Center ("MSC") team are located in Lexington, KY. The MSC is located in Lexington, KY.



## Managed Solution Center

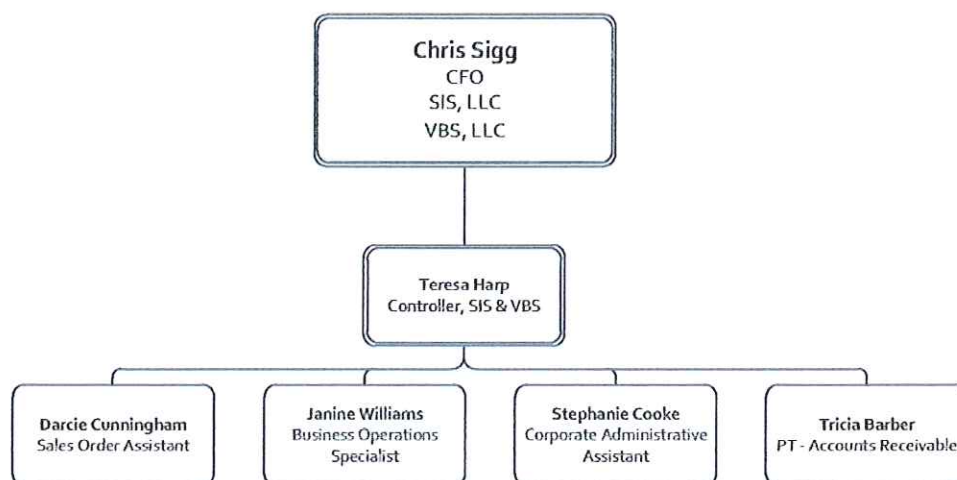


11/10/2011

Members of the Financial & Administrative team are located in Lexington, KY and Alexandria, VA.



## Financial & Administrative



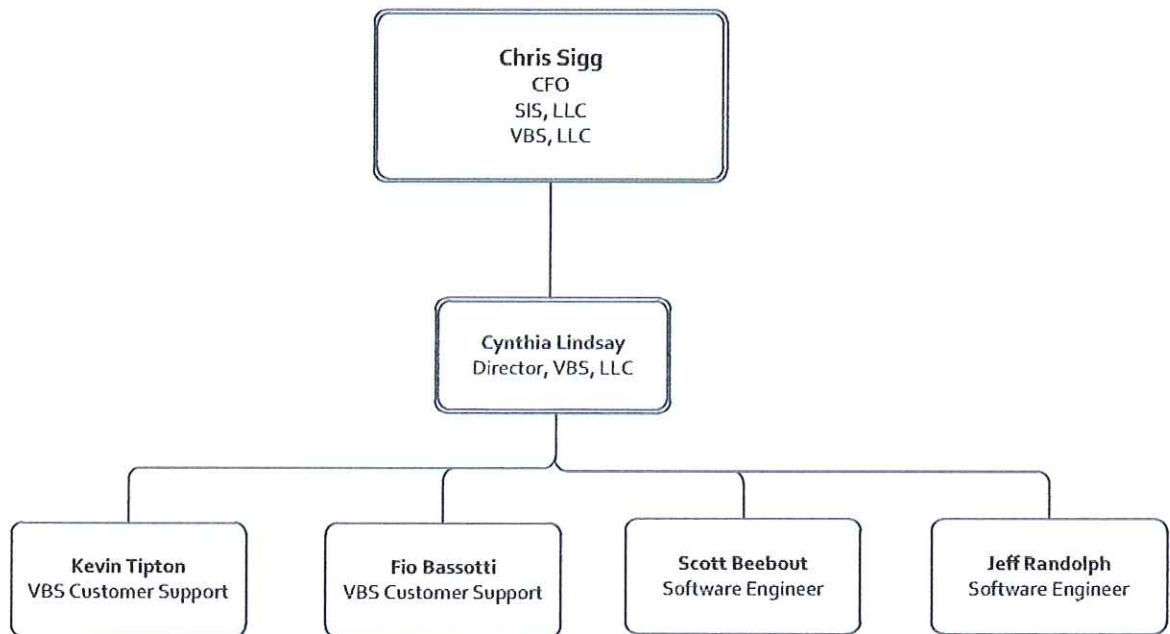
11/10/2011



Members of the VBS team are located in Lexington, KY and Alexandria, VA.



VBS, LLC



11/10/2011

## Attachment A-2: References

---

A minimum of 3 references requested. At least one reference should be from the public sector. All references should be from accounts of a similar scope and complexity as the project outlined in this RFP and include telephone number and email address.

Ohio Dept of Jobs and Family Services

Contact : Mark Thrasher

Phone: (614) 387-8391

Email: [Mark.thrasher@jfs.ohio.gov](mailto:Mark.thrasher@jfs.ohio.gov)

Peoples Bank of Marietta

Contact : Steve Nulter

Phone: 740 376-7106

Email: [snulter@pebo.com](mailto:snulter@pebo.com)

First Community Bank, Bluefield VA

Contact: Brian Broyles

Phone: 304 323-6400

Email: [bbroyles@fcbinc.com](mailto:bbroyles@fcbinc.com)

WVU

Contact: Mark Six

Phone: 304 293-9083

Email: [mark.six@mail.wvu.edu](mailto:mark.six@mail.wvu.edu)

WVNET

Contact: Richard Lynch

Phone: 304 293-5192 Ext 243

Email: [rich@mail.wvnet.edu](mailto:rich@mail.wvnet.edu)

Office of the State Comptroller

Contact: Brad Lawyer

Phone: 518 486-6753

Email: [blawyer@osc.state.ny.us](mailto:blawyer@osc.state.ny.us)



### **Attachment A-3: Team Resumes and Certifications**

---

Below are SIS resumes of proposed project team members which provide adequate combined experience in completing similar projects.



Open Systems Implementation Engineer

---

#### Certifications and Related Skills

---

- VMware Certified Professional
- VMware Sales Professional
- HP AIS Networking
- Compellent Storage
- HP P4000 (Lefthand)
- Novell Certified Linux Professional

---

#### Project History

---

**Industry: Education**

**Project: Infrastructure Virtualization**

**Technologies Used: VMware ESX, VMware Converter, HP Lefthand**

Provided virtualization design and implementation using two different sites to assist with a disaster recovery strategy, server refresh cycle and consolidation project.

**Industry: Corporate**

**Project: VMware installation and storage configuration**

**Technologies Used: VMware ESX, VMware Converter, Compellent, HP Lefthand**

Implemented many standard VMware ESX installs and storage configurations.

Assisted clients with physical-to-virtual process training and technical assistance.

Also provided client with documentation and training to manage new system.

**Industry: Education**

**Project: Infrastructure Upgrade**

**Technologies Used: HP Procurve and Cisco**

Designed and managed client resources with the implementation and upgrade of a new switch infrastructure. Provided the replacement of 71 switches with HP Procurve and the reconfiguration of 30 Cisco switches totaling approximately 3100 ports of connectivity. Also reconfigured the network from a flat network and no VLANs to a



segmented network to limit broadcast domains. The network consisted of 5 sites and different interconnect mediums. Time given for project – 1 week.

**Industry: Various**

**Project: Linux Training**

**Technologies Used: Linux**

Provided basic Linux training for clients interested in enhancing management of systems. The training provided foundational concepts of a Linux system concerning structure and management. Distributions used were RedHat/CentOS, SLES/OpenSUSE, Ubuntu.

**Industry: Various**

**Project: Linux Implementations**

**Technologies Used: Linux**

Provided various services using Linux based technology. Services implemented were Linux Terminal Services, File/Print, Web and application services. Distributions used were RedHat/CentOS, SLES/OpenSUSE, Ubuntu.

**Industry: Various**

**Project: Linux – Windows Integration**

**Technologies Used: Linux, Windows Server**

Provided Linux and Windows integration from basic file, print and web services to full user integration . Distributions used were RedHat/CentOS, SLES/OpenSUSE, Ubuntu.



Mike Neal

#### Project History

---

**Industry: Government**

**Project: HIPAA Compliance**

**Technologies: HIPAA Security Rule, NIST SP800-30/33/53**

Project: HIPAA Compliance Gap Analysis and HIPAA Compliance Risk Assessment for statewide Health Information Exchange (HIE). Served as Project Manager over internal and subcontractor resources providing on-site project management, leading meetings, providing weekly status reports, reporting findings from Gap Analysis, and contributing consulting-level expertise to guide client through the process.

**Industry: Healthcare**

**Project: Risk Assessment**

**Technologies: HIPAA Security Rule, NIST SP800-30/33/53**

Project: Perform Risk Assessment for healthcare entity who was attesting to Meaningful Use for federal funding. Served as both Consultant and Project Manager for the engagement - as Consultant, conducted System Characterization, Threat Identification, and worked with internal team to perform Risk Analysis and as Project Manager, lead meetings and provided weekly status reports keeping the client in tune with progress against budget. Project was completed 25% under budget and one-week ahead of schedule.

**Industry: Financial**

**Project: Assessment, Consulting, IT Steering Committee**

Project: Chaired IT Steering Committee for a client, acting as interim IT Director after termination of VP of Technology which was a result of IT Assessment. Conducted introductory meetings to determine alignment of business goals to technology, SWOT analysis, and strategic IT planning. Set company direction on timeline for implementation of various technologies that supported new customer services, facilitated company growth,



and met regulatory requirements. Ongoing remote and onsite meetings during year-long transition process of IT, which culminated in infrastructure overhauls and IT Managed Services.

**Industry: Financial**

**Project: Infrastructure Overhaul**

**Technologies:** HP/EMC hardware, Cisco networking, VMWare, Veeam, Managed Services

**Project:** Complete overhaul of IT infrastructure in conjunction with new data center build-out and company growth. Managed project for virtualization of Windows server infrastructure, network overhaul for full redundancy across branch locations, implementation of upgraded security solutions, and centralized management of all infrastructure/solutions through IT Managed Services.

**Industry: Legal**

**Project: Cloud-based Infrastructure**

**Technologies:** HP/EMC hardware, Check Point, VMWare, Veeam, Citrix, Microsoft, Symantec, Managed Services

**Project:** New startup law firm required entirely outsourced IT operations to support all file and print services, legal software applications, email, etc. Architected and designed entire solution for connectivity, applications delivery, email communications, etc. on top of VMWare-backed cloud infrastructure on HP/EMC hardware. Developed model for cloud services and managed IT services coupled to provide a fully outsourced IT operation.

**Industry: Financial / Healthcare / Government / Other**

**Project: Information Security Assessment and Audit**

**Project:** Conducted and project managed 40-50 engagements over 2 years for various regulated industries with evaluation of security posture and controls against regulatory guidelines and information security best practices. Projects included both technical and non-technical assessments of security controls, social engineering, application security, and management process review.



*"Thank you for everything you do. It is a pleasure working with you."*

#### Certifications and Related Skills

---

- Microsoft Certified Systems Engineer (Windows 2003)
- VMware Certified Professional
- DataCore Certified Install Engineer
- NetApp Certified Design Architect
- NetApp Certified Implementation Engineer

#### Project History

---

**Industry: Manufacturing**

**Project: VMware and Storage Upgrade**

**Technologies: VMware vSphere and DataCore SANmelody**

This project upgraded two high availability DataCore SANmelody 2.0.2 servers to version 2.0.4 with no service outage. VMware Virtual Center 2.5 was upgraded to VMware vCenter 4.0 with no service outage. Three VMware ESX 3.5 servers were upgraded to VMware ESX 4.0 with no service outage to running virtual servers. Classroom instruction and lab time was provided for the customer covering the version changes for the upgrades.

**Industry: Financial**

**Project: Infrastructure Refresh**

**Technologies: Microsoft Windows 2003 Server, Cisco ASA Firewall, Symantec Backup Exec and Endpoint Protection, IBM Intel Hardware and workstation deployment**

This project upgraded a Windows 2000 environment to Windows 2003 R2 (Windows 2008 is not supported from an application vendor). The upgrade included server replacements at three locations, sixty workstation replacements, Active Directory upgrade and application replacements including Symantec Backup Exec, Endpoint Protection and third party banking applications. The three locations also had a Cisco PIX to Cisco ASA firewall migration.



**Industry: Broadcasting**  
**Project: Microsoft Exchange 2007 deployment**  
**Technologies: Microsoft Exchange 2007**

This project replaced a POP email server with Microsoft Exchange 2007 for a local television station. There were 120 mailboxes migrated to Exchange 2007 with no downtime or message loss.

**Industry: Sports Office**  
**Project: High Availability Environment**  
**Technologies: VMware vSphere, DataCore SANmelody, Microsoft Windows 2008 Server**

This project provided a high available storage and VMware ESX environment. A Windows 2003 Small Business Server was migrated to Windows 2008 Servers and Active Directory. Classroom knowledge transfer was provided to the customer.

**Industry: Manufacturing**  
**Project: High Availability Environment and Disaster Recovery location**  
**Technologies: VMware ESX and DataCore SANmelody**

This project provided high available storage, VMware ESX environment and storage replication to a disaster recovery site. DataCore SANmelody was utilized for local storage high availability and remote asynchronous data transfer to the disaster recovery site. Classroom knowledge transfer was provided to the customer.



Open Systems Implementation Engineer

Certifications and Related Skills

- VMware Certified Professional – VCP5
- VMware Technical Sales Professional
- VMware Sales Professional
- VMware Certified Associate - Desktop Technologies VCA4-DT
- Veeam Technical Sales Professional
- Veeam Sales Professional
- Microsoft Certified Professional

Project History

**Industry: Manufacturing**

**Project: Infrastructure Virtualization and Relocation**

**Technologies Used: VMware ESX, VMware Converter, Vizioncore vConverter**

I was part of a 6-week project that involved the virtualization of this company's entire Intel-based infrastructure. There were approximately 50 physical and virtual machines that had to be converted to virtual machines or upgraded from their existing virtualization. Phase 1 was the virtualization, phase 2 involved relocating their datacenter approximately 110 miles away with less than expected issues.

**Industry: Various**

**Project: VMware installation and jumpstart presentation**

**Technologies Used: VMware ESX, VMware Converter, Veeam Backup and Replication**

Provided jumpstart training on VMware ESX 3.5 and 4.0. Jumpstarts involve classroom sessions as well as hands on learning leaving the customer with a fully functioning production environment that they created with our direction.

**Industry: Manufacturing**

**Project: VMware ESX Deployment**

**Technologies Used: VMware ESX, Ultimate Deployment Appliance**

Using open-source applications I was able to deploy 56 ESX hosts with very little time involved. The deployment was done in phases of 28 and upon completion were 75% ready. The deployment of 28 hosts took less than an hour.


**Industry: Various**

**Project: Desktop Virtualization**

**Technologies Used: VMware ESX, VMware View, Wyse Thin Clients**

Provided VMware View jumpstarts to get a customer on the road to virtualizing not only their server environment, but also their desktops. Jumpstarts include the installation and configuration of software, as well as the specific needs of the customer's desktop requirements.





*"Thank you for everything you do. It is a pleasure working with you."*

### Certifications and Related Skills

#### Certifications

- Microsoft MCSE 2000
- VMware
  - VCP 3.5 and 4.0 (VMware Certified Professional)
  - VTSP (VMware Technical Sales Professional)
  - VSP (VMware Sales Professional)
  - VAC Accreditations
    - Infrastructure Virtualization
    - Desktop Virtualization
    - Business Continuity
    - System Management
- HP Blade ASE
  - HP Certified Professional Implementing HP BladeSystem
  - Integrating and Managing HP BladeSystem in the Enterprise
- LPI
  - Linux Professional Institute
- Hitachi Data Systems
  - HDS Modular Storage Certification
- CompTia
  - Security+

#### Skills Assessment

- Microsoft Windows XP, Vista and Windows 7
- Microsoft Windows 2003 Server & Enterprise Edition
- Microsoft Windows 2000 and 2003 with Clustering and Network Load Balancing
- Microsoft Windows 2008
- Microsoft Active Directory
- Microsoft Exchange 2003 & 2007
- Macintosh Hardware & Operating Systems
- LAN/WAN Administration
- EMC Clarion and Symmetrix
- Hitachi AMS Sans
- Redhat Linux 7.x – RES 5 & Fedora Core 1 – 11
- Ubuntu 7.x -10.x

- o VMware ESX 1.x – 4.X, P2V, Converter, VirtualCenter, VMware Server 3.x and Workstation 3.0 6.X, vSphere infrastructure
- o VMware Capacity Planner Assessments
- o HTML with ASP and Java Scripting
- o Functional skills Microsoft SQL and Microsoft Access
- o Extensive scripting knowledge with Linux Bash Shell, Batch, Windows Shell, VBScripting, Perl, Powershell
- o IBM Director 4.x-5.x and Remote Deployment Manager
- o Server Security Administration on Windows, Linux and Active Directory including GPO and associated tasks.
- o Assessing, Designing and Implementation of General Network Security, Hardening and Best Practices
- o Assessing, Designing and Implementation of Microsoft and Linux Server Security, Hardening and Best Practices.
- o Assessing, Designing and Implementation of Internet Connectivity Security, Hardening and Best Practices
- o General education of Management and Users of Social Engineering techniques and Best Practices.

## Project History

---

### **Industry: Banking**

**Project: Virtualization, Storage and Microsoft Enterprise Solutions**

**Technologies: VMware, Intel Class Servers, Pillar SAN, Brocade Fibre Switches, Microsoft Group Policies, Disaster Recovery, Business Continuity, Site to Site Replication, Exchange 2003, VizionCore Replication and Backup**

This Project was to Architect, Design and Implement a complete site to site failover utilizing industry standard best practices for Business Continuity. The project had a RPO and TRO of less than one hour, the client had Data Centers in Cincinnati and Cleveland that required site to site replication and security compliance.

### **Industry: State/Health Care**

**Project: Infrastructure Refresh**

**Technologies: VMware, Hitachi SAN, Intel Servers, Brocade fibre switches, HP Data Protector Backup Software, HP Tape Array, Disaster Recovery, Business Continuity**

This Project was to Architect, Design and Implement a complete server and application consolidation solution for a Poison Control call center for the state of Illinois. This design had very high levels of availability and performance to handle the call load that the center would receive. This Project also involved me creating a face to face interview process to acquire the true business DR rating of the applications and viability of the application becoming highly available.

**Industry: Automobile Manufacturing**

**Project: Infrastructure Refresh**

**Technologies: VMware, EMC SAN, Intel Servers, Brocade fibre switches, Disaster Recovery, Business Continuity**

This Project was to Architect, Design and Implement a complete server consolidation solution for a high volume automobile manufacturing plant. This design had very high levels of availability and performance to handle the 24/7 manufacturing of the production lines in the plant. This company had an estimate that for every 1 minute of downtime meant millions of dollars in lost revenue. So this solution had to be extremely solid and versatile





"I have known Eric Lovelace at the highest level, and with his assistance, I have been able to successfully complete a number of projects. Eric is a professional who is always willing to go the extra mile to ensure that a project is completed on time and to the satisfaction of the client. He is a true professional and a pleasure to work with. I highly recommend Eric Lovelace for any project that requires a professional with a strong background in IT and a commitment to excellence." **Eric Lovelace, Itelligence**

## Jarelle Marshall

## Profile

SNIA Certified Storage Professional  
ITIL V3 Foundations Certification  
IBM Certified Specialist - System Storage SAN Volume Controller  
IBM Certified Specialist - Systems Storage Implementation Entry/Midrange for DS series,  
IBM Certified Specialist – Storwiz V7000 Technical Solutions V2  
Netapp Certified Implementation Engineer  
Netapp Certified System Administrator  
Microsoft Certified Professional (MCP)  
Comptia A Plus Certification (A+)  
Comptia Network Plus Certification (Net)  
Dell Parts Direct Certification, Printer, and Servers, Notebooks (DWD)  
HP Printers, Desktop, and Laptop Certification (ASE)

## Project History

### Industry: Aerospace

### Project: Datacenter Consolidation

### Technologies: IBM XIV, SVC, Softek TDMF, NexSAN SATABEAST

Project to consolidate 13 datacenter spread across the globe into 3 datacenters in the United States. / My role was to design and implement a solution to migrate data quickly and efficiently given time constraints, and bandwidth restrictions. / Migrated over 500TB of data using a mix of SVC tools, Softek TDMF, and Lift and Ship utilizing Nexsan Satabeast with zero unscheduled downtime, and zero data loss or corruption over the entire project.

**Industry: Healthcare**

**Project: Storage Refresh**

**Technologies: IBM SVC, XIV, DS4800, and DS8100, Softek TDMF**

Project to remove old hitachi, and hp storage, and virtualize there entire storage landscape and unilize Boot from SAN. / Assessed storage demands and performance metrics to create a tiered data storage environment. After virtualizing various IBM storage subsystems, behind SVC, we had the ability to move data from the old san to the new SAN with as little and a reboot, and we were able to phase our moves based on hosts, and not on a storage subsystem basis. This also provided an very efficient way for our customer to stay current of firmware patches with NO downtime or impact to there hosts. / main customer benefits



## Software Information Systems, LLC

---



Jeremy Brovage

Backup Architect

---

### Certifications and Related Skills

IBM Certified Tivoli Storage Manager Deployment Professional for all TSM from v4.1 to v6.1

IBM Certified Systems Expert - Enterprise Technical Support for AIX and Linux  
Virtual Storage and Virtual Tape Solution Engineer

IBM Mid-level and Enterprise Storage Administration and Implementation

VMWare Professional with virtual storage and virtual server backup/recovery expertise

---

### Project History

**Industry:** healthcare

**Project:** TSM Backup solution

**Technologies:** Backup and DR, Oracle RMAN, VTL, HSM, Shared Storage,

Chief TSM Architect for a large hospital network's initial TSM implementation, with continued role as primary Architect over a ten year period. Environment included TSM SAN agents, as well as TDPs for multiple applications (Oracle, SQL, MS Exchange, etc). Solution design included a SAN-based backup solution for a critical 10TB Oracle database application. Designed and implemented TSM/HSM dual-active AIX clusters for additional remote hospital locations. Service delivery work included upgrades and cleanup over a nine month period. Additional service delivery work involved conducting performance tests and design recommendations for backup of the large Oracle database.

**Industry:** Manufacturing

**Project:** Automated TSM rollout

**Technologies:** automated software distribution, TSM architecture, VTL



Architected, and Lead Engineer role on an automated TSM server installation system that used CYGWIN shell and Windows Scripting to provide a way to roll out fully customized TSM servers in remote offices with no skilled TSM personnel. This process was sold and used in an implementation across 30+ sites in Europe for a global corporation based in Toledo.

**Industry:** Insurance

**Project:** TSM Backup solution

**Technologies:** Backup and DR, VTL

Architected and implemented a 300-node TSM solution for a client on MS Windows in 2001-2002. Initial architecture requirements presented unique challenges and part of large Windows TSM solutions. The client continues to use TSM on Windows, with environment expanded to 3 TSM servers, VTL and various other more advanced TSM technologies. Continuing role as provider of TSM advice, and strategy planning, as well as hands-on upgrade and troubleshooting work.

**Industry:** Advertising

**Project:** Server Consolidation

**Technologies:** virtualized servers, virtualized storage, data protection.

Lead the implementation of a physical to virtual migration. Implemented a VMWare multi-host cluster utilizing IBM BladeCenter and a FalconStor IPStor virtual storage cluster. The IPStor cluster virtualized about 6TB of IBM DS4000-series disk for the implementation. The goal was to provide increased uptime due to the extensive use of in-house web-hosting provided by the customer to a major automotive manufacturer. The final implementation allowed for hardware failure, storage failure, Firmware upgrades, and snapshot-based backups with near-zero downtime.



Jon Simon

*Project Manager*

#### Certifications

ITIL V3 Foundations Certification

#### Project History

**Industry:** Automotive Leasing

**Project:** Storage Upgrade and Virtualization, NetBackup, Exagrid

**Technologies:** IBM, p-series, x-series, VMware, NetApp, Cisco

[Project: Upgrade existing network design to accommodate both virtual workload traffic and backup traffic while not impacting the effectiveness of the core network stack; upgrade the existing NetApp storage to replace end of life equipment and to allow for virtualization of over 30 servers, including primary production servers such as Microsoft Exchange and both Oracle and SQL database servers supporting over 500 employees. This project included a complex coordination of VMware, AIX, NetApp and Cisco resources to meet all customer requirements and a continuation of effort to ensure the success of the project and the satisfaction of the customer.]

**Industry:** Financial

**Project:** Backup and Disaster Recovery

**Technologies:** VMware, Veeam

[Project: Institute a successful backup and disaster recovery process for a local bank with multiple locations. My role as Project Manager was to oversee the discovery, design and implementation of the company's Veeam Backup and Replication of over 20 virtual machines running on VMware servers.]

**Industry:** Financial

**Project:** Storage Upgrade and Virtualization Upgrade

**Technologies:** VMware, IBM, Veeam

[Project: Upgrade the existing VMware virtual environment hardware, migrate the virtual machines to the new hardware, upgrade the existing back-end storage from the existing DS4700 to a new IBM V7000 and relocate the old equipment to the company's disaster recovery location to allow for off-site replication of the production environment.]



Tom Miller

Tom Miller is a former VMware employee who has been with SIS for over 10 years. He is responsible for our VMware practice and has been instrumental in the success of our VMware engagements. He is also responsible for mentoring other SIS engineers to deliver high quality consulting engagements. Also VMware program development and offerings are Tom's responsibility. This includes package development, sales and engineer training. Tom has been a key component of several large data center consolidation and relocation projects.

#### Certifications and Related Skills

- VMware – VCP 2,3,4,5 VSP VTSP
- Microsoft – Microsoft Certified Professional
- Novell – CNE 3,4,5,6
- ITIL Certified

#### Server Virtualization

At SIS Tom is our VMware Practice lead. Tom developed our VMware "Jumpstart" methodology where the training comes to you. As a former VMware employee Tom has transitioned that experience where Tom is responsible for our quality of service as it pertains to VMware engagements. Tom is also responsible for mentoring other SIS engineers to deliver high quality consulting engagements. Also VMware program development and offerings are Tom's responsibility. This includes package development, sales and engineer training. Tom has been a key component of several large data center consolidation and relocation projects.

**Project:** Data Center Consolidation and Relocation

**Technologies:** VMware, Dell, EMC, Emulex, Cisco

Project coordinated, designed, and implemented a 4 month project to relocate and upgrade 30 ESX host and 230 Virtual Machines. Project consisted of upgrading client from VMware 3.5 to 4.0. Also 85 physical machines consisting of Linux and Windows were converted to virtual machines verses a physical relocation.



**Project:** Disaster Recovery with Site Recovery Manager

**Technologies:** VMware, IBM, EMC

Project consisted of training customer on installing, managing and utilizing Site Recovery Manager for a Disaster Recovery initiative. The infrastructure consisted of IBM Bladeservers utilizing EMC SAN technology with replication. Design, implementation, and Disaster Recovery execution were required deliverables.

## Software Information Systems, LLC



Jeff Barratt-McCartney

Storage/Power Architect

---

### Certifications and Related Skills

IBM Certified Advanced Technical Expert pSeries and AIX 2003-2009  
IBM Certified Specialist - Enterprise Storage Technical Support V2  
IBM Certified Specialist - Midrange Storage Technical Support V2  
IBM Subject Matter Expert XIV migration/replication developed test 000-967  
IBM Subject Matter Expert developed tests 000-100, 000-102 and others

---

### IT Experience

15+ years industry experience. Initial experience as a Unix administrator (AIX) and Lead Storage Administrator in a rapidly growing insurance/financial services organization (6 years). 9 years as a technical and delivery consultant for IBM business partners, my focus is twofold. I concentrate my efforts on storage migrations: cross platform, virtualization, replication, disaster recovery. I also focus on IBM storage and platform (AIX) performance optimization in Oracle environments.

### Project History

**Industry:** Healthcare

**Project:** XIV remote migration/replication and SVC virtualization

**Technologies:** IBM XIV, IBM SVC, AIX, VMware, MScluster.

Primary architect and delivery consultant for remote migration of legacy datacenter to multiple remote datacenters. Used multiple platform (AIX/VMware/Microsoft) and storage migration techniques to migrate 100TB of legacy data from disparate storage subsystems to XIV storage, then replicated the host data to remote site and facilitated the physical movement of all hosts to multiple independent data centers. Following physical migrations; all hosts were migrated to a virtualized SVC stretch cluster to provide redundancy and recoverability for all critical infrastructure.

**Industry:** Service

**Project:** Storage migration

**Technologies:** IBM XIV, AIX, HP-UX, MS, linux, VMware, Cisco SAN

Migrated multiple platforms from underperforming legacy storage subsystems to IBM XIV and Cisco SAN. Provided repeatable migration strategy to customer for additional migrations.

**Industry:** Manufacturing

**Project:** Storage migration

**Technologies:** IBM SVC, EMC, AIX, Solaris, MS, IBM DS4xxx, IBM DS8100, IBM I5, IBM 800, Z, Brocade, Connectrix

Provided technical design, architecture and delivery for replacement of antiquated multi-tier EMC infrastructure. Replacement architecture included SVC virtualized IBM DS4xxx, and IBM DS8100, and non-virtualized ESS 800 for zSeries 3390 drives. OS platforms included antiquated/unsupported Zseries, E10K, AIX 4.3.3, i5, MS. Customer was faced with less than a month to replace antiquated EMC storage and SAN infrastructure, or face significant maintenance and software renewal charges. Delivery required around the clock migrations with minimal input from the customer. Delivered a comprehensive multi-tiered virtualized solution for 100+ vintage hosts, on time.

**Product expertise and experience:**

IBM Power Systems, AIX, IBM SVC, IBM XIV, IBM DS4xxx, IBM 5xxx, IBM DS8xxx, IBM F20, IBM 800, EMC Clariion, NetApp filers and gateways, brocade SAN, Oracle, Oracle Financials, DB2, etc



Below is a list of certifications applicable to this project:

First	Last	Certification Description	Test Nu
Dean	Auger	IBM Certified Specialist - Enterprise Storage Technical Support V2	119
Jeff	Barratt-McCartney	IBM Certified Specialist - Enterprise Storage Technical Support V2	119
Alan	Powers	IBM Certified Specialist - High Volume Storage Fundamentals V1	203
Robert	Ray	IBM Certified Specialist - High Volume Storage Fundamentals V1	203
Dan	Chatson	IBM Certified Specialist - High Volume Storage Fundamentals V2	955
Dean	Auger	IBM Certified Specialist - Midrange Storage Technical Support V2	118
Jeff	Barratt-McCartney	IBM Certified Specialist - Midrange Storage Technical Support V2	118
Carrie	Hughes	IBM Certified Specialist - Midrange Storage Technical Support V2	118
Jarelle	Marshall	IBM Certified Specialist - Midrange Storage Technical Support V2	118
Jeff	Barratt-McCartney	IBM Certified Systems Expert - XIV Replication and Migration Services V1	967
Roger	Finney	IBM System Storage Sales Professional V2	S32
Robert	Ray	IBM Systems Storage Sales Professional v1	M53C
David	Griffin	NetApp Certified Design Architect	
Jarelle	Marshall	NetApp Certified Design Architect	
David	Griffin	NetApp NCIE SAN	
Jarelle	Marshall	NetApp NCIE SAN	

## **Attachment A-4: Similar Project Descriptions**

---

Vendor response should provide a minimum of two (2) successful projects related to the project outlined in this RFP. The referenced projects should have a successfully completed delivery and implementation. Projects that are in process, but not completed, may be used as options. The Vendor should have had primary responsibility (not acting as a sub-contractor) for the various phases of the projects including: analysis, project/process design, pilot/test phases, and implementation. Vendor should clearly include the description of past projects completed entailing the location of the project, project manager name and contact information, type of project, and what the project goals and objectives were and how they were met.

Listed below are two (2) successful SIS projects related to the project outlined in this RFP:

### **Project 1: Industry: Public Sector (Higher Ed)**

This was a two-phased project that involved a full rollout of a new server, storage and virtual infrastructure plus a complete DR replication to an offsite location 40 miles away. Phase one and phase two were spread out over multiple years and separate purchases. Phase one was the rollout of the production location with new IBM servers and along with NetApp 2040 Storage (IBM nSeries 3400). This phase consisted of product rollout along with onsite education/training/implementation woven into a single deliverable designed to have the customer in production according to best practices by phase one completion.

Phase two of the project was the upgrade of the production storage to a new NetApp 2240 and the movement of the previous storage into a separate DR location 40 miles away. This phase consisted of a mix of storage based replication as well as MS SQL replication to the DR facility as necessary to meet the customer RPO and RTO objectives. This second phase also included the addition of new servers in the DR facility.

SIS provided all the project management as well as all the education, training, and implementation work. In addition to the infrastructure project management work performed SIS also provided a SQL expert to perform the following:

1. Providing Offsite HA/DR via log shipping against 9 databases and 2TB data
2. Installed and configured SQL Diagnostic Manager - this will allow CRD to capture and report on 1000s of SQL Server and Server metrics for performance diagnostics, forensics etc...
3. Used SQLdm to do a health assessment on a very, very busy OLTP server with a significant amount of data and simultaneous users
4. Used SQLdm to monitor on-site POC for the log shipping POC watching network bandwidth used, transactions per second, IO, memory usage and overall server stability
5. Configure HA/DR alerts, baselined SQL Server performance and configured alerts based on significant deviations from baseline.



## **Project 2: Industry: Automotive Leasing**

This was a three phased project that involved a complete overhaul of the customer's current virtual workload, back-end storage, backup strategy and network layout. In phase one of this project the first goal was to upgrade the existing NetBackup server and institute an Exagrid appliance to support a long-term goal of eliminating tape backups and using off-site replication for restore and recovery purposes. Also in this phase we laid the groundwork of a network re-design which included use of both iSCSI and fiber to allow for separation of production workload and backup traffic and to eliminate backup traffic from traversing the core network stack.

The primary focus and largest phase of this project came in phase two when SIS upgraded the customer's storage to a new NetApp FAS 3240(IBM nSeries 6240). This consisted of an extensive detailed design, implementation and customer education of the new NetApp FAS 3240 which would support both VMware virtual workload as well as AIX servers. To accomplish phase two, SIS performed a NetApp base installation of the new FAS 3240 (basic rack/cabling, initial system set-up, verification/testing, enable AutoSupport and initial customer orientation). Following the base install, SIS proceeded with the implementation to get the system to a production-ready state. SIS worked closely with the customer to assemble a complete design layout and configuration of the NetApp storage prior to implementation which reduced the need for configuration changes during the set-up of the storage.

Wrapping up the project in phase three, SIS built a new VMware virtual environment and attached it to the NetApp SAN that was set up in phase two. This consisted of SIS performing a VMware Jumpstart at the customer location, during which SIS worked side by side with the customer to discuss, design and build out the virtual environment as well as educate the customer and put them in a position to be able to manage their environment moving forward. Once the new virtual environment was set up, SIS assisted the customer to begin the conversion process to migrate their current physical servers to the new virtual workspace.

**Attachment A-5: Subcontractor Information: None being used.**



#### 2.4.1 The following sets of questions are explanation based, concerning the Unified Storage Array that may be proposed.

- a. Does the array have SGBps Fibre Channel connections to the SAN switches?  
Vendor response: **yes**
- b. Does the array have the capability to support 10Gbps FCoE for storage presentation?  
Vendor response: **yes**
- c. Does the array have the capability to support 10Gbps iSCSI for storage presentation?  
Vendor response, **yes**
- d. Does the array support automatic, dynamic read/write memory (cache) allocation?  
Vendor response: **yes**
- e. Does the array support both 2.5" and 3.5" disk drives?  
Vendor response: **yes**
- f. Does the array support RAID 1/0 (striped and mirrored)?  
Vendor response: **yes**
- g. Does the array support RAID 5 (single parity)?  
Vendor response: **yes**
- h. Does the array support RAID 6 or RAID-DP (double parity)?  
Vendor response: **yes**
- i. Can the storage administrator choose which tier of disk in a storage pool is used when creating a new LUN? "  
Vendor response: **yes**
- j. Can individual LUNs be expanded "on the fly" without down time on the system using the LUN?  
Vendor response: **yes – LUNs can also be reduced without down time.**
- k. Can individual LUNs be converted from thick to thin provisioned and vice versa?  
Vendor response: **yes**
- l. Does the array support space reclamation on existing thin provisioned LUNs?  
Vendor response: **yes**
- m. Can individual LUNs be manually migrated between disk types in a storage pool without down time on the system using the LUN?  
Vendor response: **Yes Datamotion for SAN LUNs will move volumes with LUNs between aggrs on a controller**
- n. Does the array support the exclusion of particular LUNs from automatic tiering?  
Vendor response: **N/A**
- o. Are full copies (clones) of LUNs available for use immediately after initiating clone operation?  
Vendor response: **yes – without increasing storage at time of clone.**
- p. Are full copies (clones) mountable by a different host?  
Vendor response: **yes**
- q. Does the replication technology in the array support both local and remote protection?  
Vendor response: **yes – via IP based replication**
- r. Does the storage array utilize 10Gbps Ethernet for the NAS (CIFS/NFS) functionality?  
Vendor response: **yes**
- s. Does the storage array support NDMP for backup of raw file systems?  
Vendor response: **yes**

t. Does the NDMP support allow for file and folder level restoration without the need to restore the entire NAS file system containing those items?

Vendor response: **yes**

u. Does the CIFS file server support Windows 2008 R2 native-mode Active Directory domains?

Vendor response: **yes**

v. Does the CIFS file server support Volume Shadow Copy to allow end-user or support staff recovery of files and folders using the "Previous Versions" features built into the Windows client operating systems when utilizing snapshot technology on the CIFS file shares?

Vendor response: **yes**

w. Does the CIFS file server support ABE (access-based enumeration)?

Vendor response: **yes**

x. Does the array support de-duplication of data presented via file protocols (CIFS/NFS)?

Vendor response: **yes – in production and secondary tiers across all applications**

y. Does the array support compression of data presented via file protocols (CIFS/NFS)?

Vendor response: **yes**

z. Does the array provide the ability to administer the system via a command line interface (CLI installed on a remote system or direct SSH/telnet interface)?

Vendor response: **yes**

aa. Does the array provide the ability to script administrative actions for bulk operations?

Vendor response: **yes**

2.4.1.1 Does the array support the use of solid state drives (SSD) or enterprise flash drives (EFD) as an extension of read/write cache to enhance performance and alleviate hot spots from sudden, unexpected spikes in workload? If yes, response should provide details on how this works within the array and any limitations of this technology.

Vendor Response: Today No, available in future version of OnTap. Happy to discuss under NDA

2.4.1.2 How many total active paths does a host have to an individual LUN?

Vendor Response: 2

2.4.1.3 What is the total bandwidth available for a host to an individual LUN?

Vendor Response: Capable of 8Gb

2.4.1.4 What is the total number of drives and drive enclosures supported by the array (expandability)?

Vendor Response: 600 Drives, 1.8PB

2.4.1.5 When implementing the data-at-rest encryption does the array provide internal key management system, utilize (or require) an external key management system or rely on drive-based encryption without the need for a key management system?

Vendor Response: Both Internal Key Management or Drive-Based encryption are available

2.4.1.6 How many tiers of storage (drive types) may be placed in a single storage pool?

Vendor Response: One

2.4.1.7 Does the array support automatic data tiering within a configured storage pool to allow migration of data to higher or lower speed disks based on an activity algorithm? If yes, response should explain how the automatic data tiering works in the proposed unified storage array.

Vendor Response: **No – Please consider the below explanation as to why.**



## Tier Less, Cache More

### IBM N-Series Caching Trumps Traditional Tier Thinking

If you were delivering IT services and you had complete freedom to do what you wanted, how would you architect a solution to best support your customer and business requirements?

#### SATA + Flash Cache: Much Less Expensive, Just as Fast, More Flexible

We benchmarked<sup>[1]</sup> Flash Cache with SATA disk drives (aka intelligent caching) under an industry-standard file services workload to demonstrate that you can achieve performance comparable to that provided by a configuration with Fibre Channel disk drives while reducing the purchase price and operating cost of the storage system.

The cost of storage capacity on the configuration with Flash Cache and SATA drives was 56% lower on a per-terabyte basis than that of the baseline configuration. The SATA + Flash Cache configuration also uses 2/3 less power and 59% less rack space.

<sup>[1]</sup>These test results are [posted](#) on the Web site of the Standard Performance Evaluation Corporation (SPEC).

You could build multiple infrastructures for each customer or user. Each built with different architectures and management. Each designed to satisfy the needs of an individual client or workload. We typically call this practice of putting different workloads on different architectures “tiering.” Of course, as those needs change over time, and data becomes more or less critical, you would need to move the workloads and data from one tier to another. That often requires migration, new training, and process changes.

Or, you could build an intelligent infrastructure that can meet the needs of diverse clients and workloads, with automated management tools and processes. This infrastructure would allow you to offer multiple classes of service based on customer requirements, much like yesterday’s dedicated tiers. However, it would be flexible, so you could easily alter it over time to meet changes in business requirements and workloads. Data could be stored and managed without a trade-off in performance, cost, flexibility, or efficiency. Better yet, you don’t need to build yet another new infrastructure or forklift the data from one to another.

We think the answer is pretty clear.

Yesterday’s traditionally tiered infrastructures simply cannot affordably deliver the speed of response and diverse service levels required by business. To meet the demands of business today, IT must create a shared IT infrastructure that allows for rapid, affordable provisioning of services. This infrastructure flexes to meet diverse client and workload needs, enabling classes of service. Equally important is making sure you can dynamically manage data across storage resources, based on business need, with minimal human intervention. It is a far more flexible and efficient way to run things. It takes an intelligent, automated approach.

Flexible and efficient IT services are better enabled by [intelligent caching](#), which combines high performance and efficient cost technologies to enable people to better manage data with improved price/performance.

#### Minimize Costly Tiering

Intelligent caching lets you create a large pool of shared storage divided into virtual classes of service rather than create many hard-to-manage tiers. The cache inside the large pool of storage is dynamic, so the most important data remains in cache automatically. You can set your cache as big or as small as you want.

With the traditional physical tiering approach, once you’ve designated an infrastructure as a certain tier, it is expensive to change. Tiers are more difficult to manage because you not only need to assign data to each tier, but you also need to continually copy and delete data as you move among tiers. You need to provision new storage as

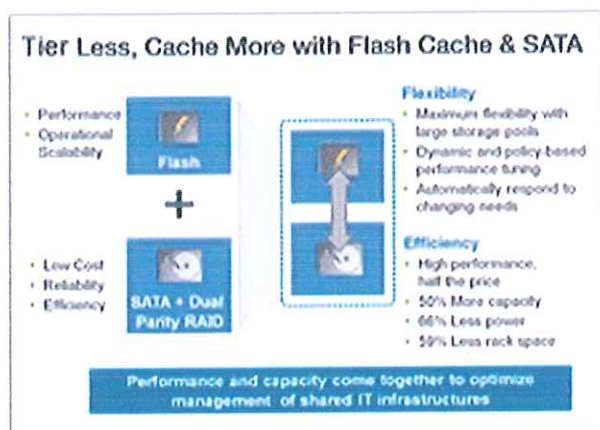


tiers grow and shrink, a job that is particularly challenging in an environment with different platforms at each tier. Worst of all, the tiered strategy ties workloads to particular services. If tier three doesn't have disaster recovery services, but business requirements change, IT has a project on its hands! In a tierless environment, meeting changing requirements doesn't require copy and delete operations, it delivers [Flexible IT](#). A simple policy change, and the challenge is met!

With intelligent caching, you use flash memory intelligently to optimize performance as needed. For example, if you couple it with SATA to manage capacity requirements, you can use far fewer [high-performance disk drives](#). Together, SATA and Flash Cache can significantly reduce costs, thereby, helping you save a boatload of money without affecting performance for many workloads. More importantly, you manage different classes of service virtually, which dramatically increases your flexibility.

Best of all, by taking an intelligent caching approach, your data center staff no longer needs to waste cycles managing tiers. We've simplified that. Intelligent caching automatically puts a copy of hot data where access will be fast and erases it when the data becomes cold. The persistent copy can stay in low cost storage. You can make the cache even smarter. Our FlexShare software lets you set up classes of service. For example, you can give a high priority to your OLTP database and a low priority to SharePoint files. Your OLTP data will be kept in the cache. When push comes to shove and SharePoint files will get kicked out.

#### Flash + SATA: Same Speed, Lower Cost, Greater Business Acceleration



Flash Cache plus cost-effective, high-capacity SATA drives changes the economics of storage and data management. Now, you have very dynamic data management at an extremely low cost. Perhaps most important, you have the flexibility to automatically adjust IT resources as needed to support business priorities and keep pace with business demands.

Most of you know that historically, the labor cost of managing storage has been bigger than the cost of acquiring the storage. Now, work with me here.

[Click to enlarge](#)

*If...*

A single large pool of storage is easier to manage than multiple smaller pools  
(IPB is more efficient to operate than three smaller pools of 200TB, 300TB, and 500TB.)

*And...*

The cost of managing storage is equal to or greater than the cost of buying storage  
(If you pay \$100 for your storage, it will cost you at least \$100 to manage that storage over its useful life.)

“We need solutions that can keep up with the pace and scale of exploration and give us a competitive advantage, because increasing our time to discovery is vital to the company’s continued success. Ultimately, we need to help users spend more time developing prospects and reservoir models and spend less time loading data and managing data between applications.”

- Bradley Lauritsen,  
[Apache Corporation](#)

Why does it make sense to break your storage into smaller tiered pools of storage? It seems silly. If you can fulfill all of your customers’ business application and workload needs with a single infrastructure, why would you want to build and manage multiple tiers?

Many companies have already taken advantage of the powerful combination of Flash Cache and SATA to drive business advantage. [Apache Corporation](#), an independent energy company, easily accommodated 700% data growth with no additional administrative storage costs, and its data is delivered in seconds rather than minutes because of IBM N-Series® Flash Cache.

There are definitely some workloads that require the raw performance that SSDs can provide. We have SSD shelves available to help. We simply do not believe in tiering *all* of your data.

Don’t waste precious IT resources moving lots of data, or start with your data on SSDs, then move them to lower cost architectures. The same objective of delivering classes of service and lowering costs can be met much more effectively with the data residing on one tier of disk drives combined with intelligent caching and software that sets service priorities.

Take advantage of intelligent caching to consolidate or even eliminate physical tiers, while still delivering varied service levels and diverse workloads with unprecedented efficiency, ease of management, and cost savings from a single shared storage infrastructure. With it, you can satisfy a broad range of business needs in real time and maximize efficiency and lower cost.

If you want to keep pace with business ... if you want to drive competitive advantage ... if you want your IT resources to be more responsive to strategic business projects, you need to intelligently manage your IT environment to maximize flexibility and efficiency.

2.4.1.8 Does the array support policies on automatic tiering to allow SAN administrators to designate particular LUNs that should only be migrated to higher (or lower) speed disks? If yes, response should explain how this functionality is implemented in the proposed array.

Vendor Response: Since Question is dependent on a “Yes” response to question 2.4.1.7, this question is N/A

2.4.1.9 Does the array allow scheduled policy changes for the automatic tiering of individual LUNs based on regularly occurring events? (An example would be the ability to schedule a particular LUN to only be migrated to higher speed disks during a time period that is known to be very busy)

Vendor Response: Since Question is dependent on a “Yes” response to question 2.4.1.7, this question is N/A

2.4.1.10 Does the array allow dynamic expansion of storage pools through the addition of more drives and/or RAID arrays into the storage pool? If yes, response should provide details on how this feature is implemented and any limitations imposed on this process.

Vendor Response: Yes – Simply add any number of shelves and assign the number of drives you desire to the target aggregate, up to the maximum per TB per aggregate limit.




2.4.1.11 Does the array support de-duplication of data contained on LUNs presented via block level protocol (FC/FCoE/iSCSI)? If yes, response should explain this functionality on the proposed array.

Vendor Response: Yes

### Reduce Capacity Requirements with Deduplication

With IBM N-Series deduplication, WVSTO can store just one copy of each unique data object, substantially reducing capacity requirements. Deduplication automatically removes duplicate data blocks on a 4KB level across an entire volume, reclaiming wasted storage to achieve significant space savings.




*IBM N-Series deduplication removes duplicate data blocks.*

IBM N-Series deduplication can be implemented across a wide variety of applications and file types, including:

- Data backup
- Data archiving
- Primary data volumes

Depending on the application and file type, IBM N-Series deduplication can help WVSTO reclaim up to 95% of storage space.





*Typical savings with IBM N-Series deduplication technology*

2.4.1.12 How many snapshots of a single LUN can be made? Response should include any details on performance degradation when utilizing multiple snapshots on a LUN.

**Vendor Response:** 255 – No Performance impact

2.4.1.13 Please outline the typical storage requirement for snapshots, both individual and multiple incremental snapshots of the same LUN? Also, response should provide a brief explanation of how snapshot technology is implemented on the array.

**Vendor Response:** We do not rely on “Copy on Write” Snapshots so no specific Volume or LUN of any size required.

### **Speed Backup and Recovery with IBM N-Series Snapshot Technology**

Snapshot technology is available from a variety of data storage vendors, but not all snapshots are created equal. IBM N-Series Snapshot™ technology enables WVSTO's IT administrators to create point-in-time copies of virtual machines or entire data stores. Then, using SnapRestore®, you can restore from these backup copies at any level of granularity—single files, directories, or entire volumes—simply and quickly when required. Many copies can be made at any time increment in less than one second, with no performance impact, no matter how many Snapshot copies are taken. These are not full copies of data; they are only tracking changes and are very efficient in terms of overall storage capacity. Restores can be done rapidly from any of the copies, providing WVSTO with an exceptional recovery time objective (RTO).

Up to 255 Snapshot copies can be created automatically or manually on each volume. WVSTO can use IBM N-Series Snapshot technology to perform backups as often as needed—daily, hourly, etc. In the event of a recovery, more frequent backups will reduce data lost since the last backup was taken. This greatly enhances WVSTO recovery point objective (RPO). Each Snapshot copy is RAID protected for reliable backup.

*IBM N-Series Snapshot technology. – Protect more of your data with minimal performance impact and consumption of storage space. You can take a Snapshot copy in less than a second while applications are running, regardless of volume size or level of activity on your IBM N-Series system.*

*Anadarko Petroleum Corporation uses IBM N-Series Snapshot technology to reduce backup times by 90% and recover 18 times faster. Backups that used to take 28 to 34 hours are now completed in no more than 2.5 hours. Recovery time has been reduced from 72 hours to just 4 hours.*

2.4.1.14 Does the replication technology in the array have the ability to take multiple snapshots of the LUNs to enable recovery or testing with copies of those LUNs at a user configurable interval? If yes, response should provide details on how this technology is implemented in the array.

Vendor Response: Snapshots are a function of our OnTap OS and not our Replication software. However, our replication software is very closely tied into our SnapShot technology and is dedupe aware to replicate a volume and it's various snapshots so that you can have multiple, on disk recovery points both locally, on production storage or tier 2, as well as at the DR site. We have multiple Application Specific software to help WVSTO manage this process. Please read more below about SnapMirror Thin Replication and our SnapManager suites for Oracle, SQL, Exchange, VMWare Virtual Environments.

## **IBM N-Series SnapMirror: Technical Overview**

When designing the disaster recovery component of WVSTO's business continuity plan, it is critical that the solution you deploy can meet requirements around:

- Recovery time objectives (RTOs)
- Recovery point objectives (RPOs)
- Distance between sites
- Cost (acquisition, operational, infrastructure)

It's also helpful if the solution delivers additional capabilities that improve IT operations, such as disaster recovery testing and data migration, as well as business initiatives, such as business intelligence or development and testing. These capabilities can reduce risk and show business leaders how disaster recovery investments can be leveraged to help drive the business.

SnapMirror® enables WVSTO to address all these requirements with a single solution. SnapMirror is a multipurpose replication product available on all N-SERIES platforms and IBM N-Series® V-Series systems. It helps you satisfy strict RPO and RTO requirements while controlling costs and improving operational processes.



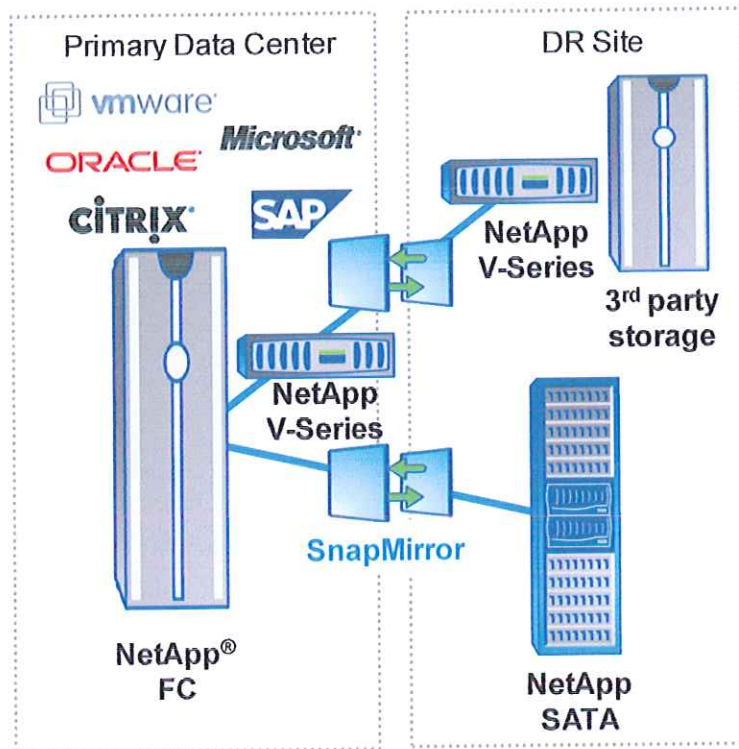


Figure 1: IBM N-Series SnapMirror: An easy-to-use, all-purpose data replication solution.

With SnapMirror, WVSTO can:

- Minimize downtime
- Reduce storage requirements and bandwidth utilization
- Actively use your DR site

### Minimize Downtime

Depending on the application and your distance requirements, the amount of downtime you can tolerate will vary. For applications with high change rates and a short distance requirement (200km), WVSTO might want to implement an RPO of zero. If longer distances are required, an RPO of seconds, minutes, or hours might be appropriate.

WVSTO can use SnapMirror® to satisfy all of these RPO requirements using any of the three common levels of data replication:

- **Synchronous** replication (up to 200km) makes sure of zero data loss
- **Semisynchronous** replication supports RPOs in seconds with minimal impact on the host application
- **Asynchronous** replication, the most flexible and cost-effective solution, can meet RPOs ranging from one minute to one day without any limits on distance



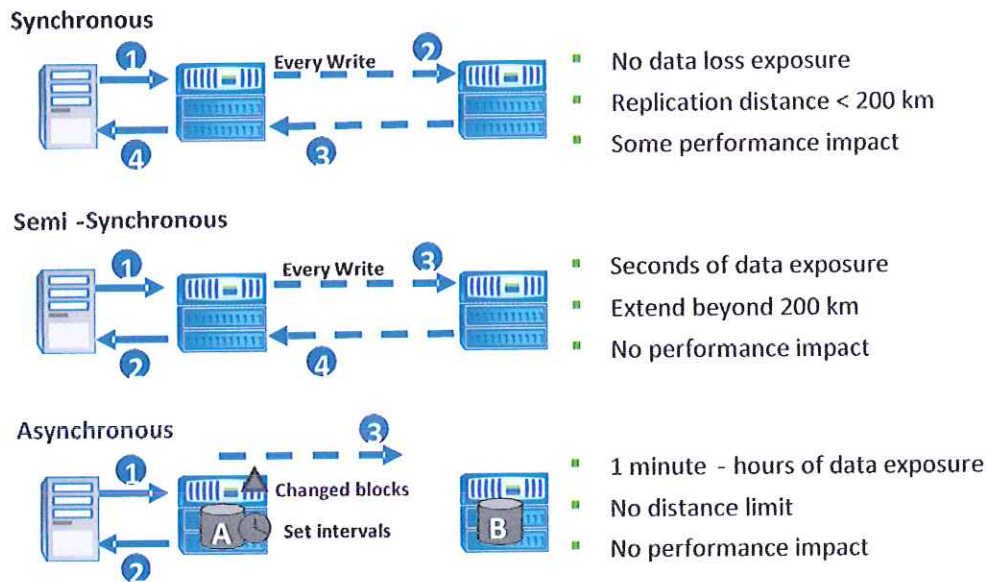


Figure 2: SnapMirror offers three levels of synchronization to address RPO requirements and minimize downtime.

For simplicity of deployment, SnapMirror works between all N-SERIES platforms and can be implemented in many-to-one, one-to-many, and cascading configurations. Integration with SnapManager® enables application-consistent protection and reduces RTO using advance restore capabilities such as automated VM failover and Exchange single mailbox restore.

*"With the common operating system across all controller heads, our storage system is easier to manage than other systems of the same size. With IBM N-Series®, we deliver 100% uptime."*

— Mark Tuttle, Senior Manager of Engineering, AutoTrader.com

## Reduce Storage Requirements and Bandwidth Utilization

Disaster recovery solutions hinge on storage availability and network speed. To implement a successful DR solution and better meet RPOs, WVSTO must have sufficient storage and bandwidth available to house and transport DR data. Using asynchronous SnapMirror® together with the core technologies of Data ONTAP®, including IBM N-Series® Snapshot™ and deduplication, you can extend IBM N-Series space saving technologies to your DR storage environment. These efficiencies can help WVSTO:

- Reduce secondary storage requirements
- Minimize bandwidth utilization
- Accelerate data transfers

WVSTO can use Snapshot and deduplication to reduce the amount of data that is actually transmitted over the network. Replication based on Snapshot sends only changed blocks, and deduplication enables only unique blocks to be sent. WVSTO can further cut bandwidth needs and costs by using built-in SnapMirror network compression tools to compress data before sending it over the wire. With less data to send and lower bandwidth utilization, you can complete transfers more quickly and reduce both RTO and RPO.

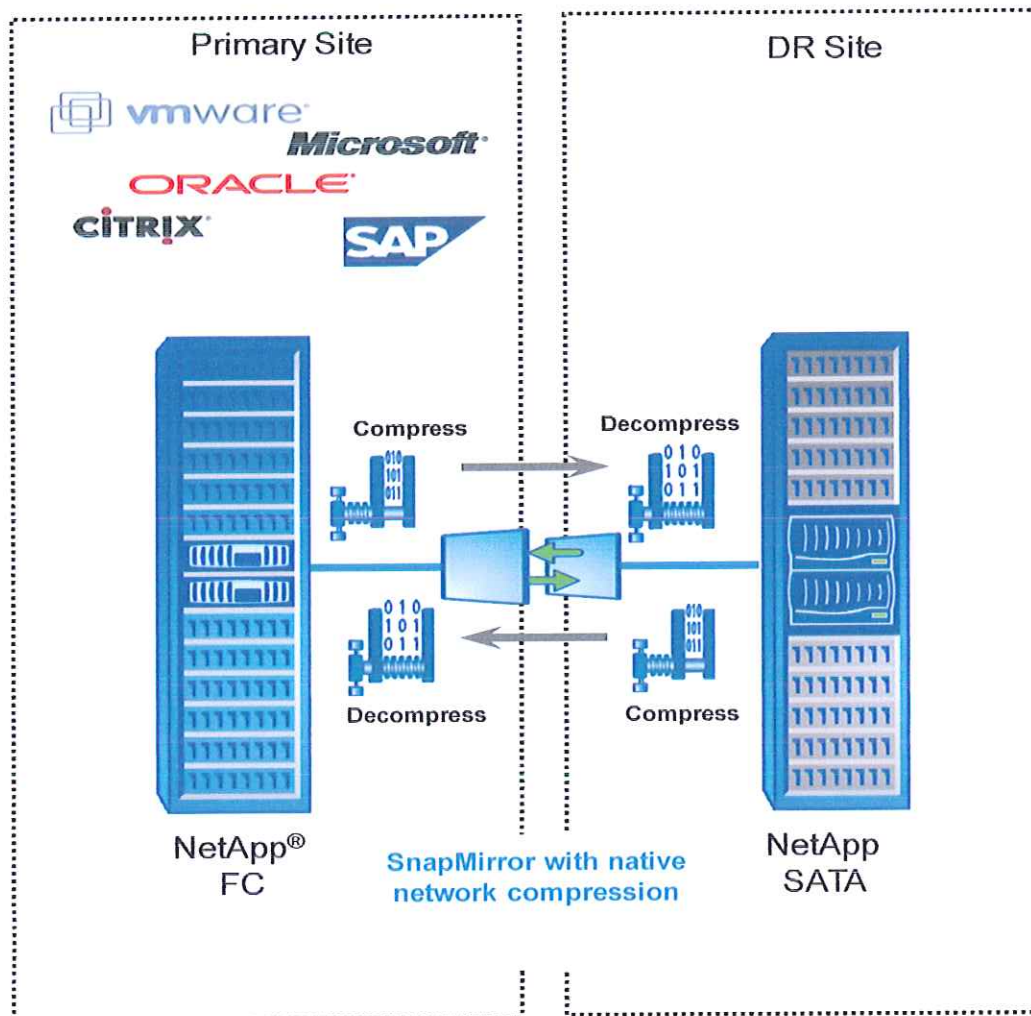


Figure 3: SnapMirror offers native network compression. Speed data replication, minimize bandwidth utilization, and reduce the performance impact on your network.

To further reduce your network bandwidth utilization, SnapMirror automatically takes checkpoints during data transfers. If your system goes down, the transfer restarts from the most recent checkpoint.

WVSTO can also use SnapMirror to perform intelligent resynchronization and eliminate the need for full transfers when recovering from a broken mirror or loss of synchronization. If data on the mirrored copy was modified during application testing, it can be quickly resynchronized with the production data by copying the new and changed data blocks from the production system to the mirrored copy.

### Actively Use Your DR Site

Being able to use your disaster recovery facility for development, testing, and business intelligence can help WVSTO turn your DR solution into a business accelerator. With IBM N-Series® FlexClone®, you can create near-instantaneous, space-efficient copies of your most current data. The clones are fully functional copies of the original information, perfectly suited for development and testing, DR testing, or other business intelligence functions.

Colocating your DR and development and testing environments can:



- Significantly improve utilization
- Speed test and development
- Improve the reliability of your DR solution by enabling consistent testing

*"We can schedule a DR test without interrupting production or IT administrators' regular activities. This capability gives us added security and confidence, knowing we can recover because we perform the tests more frequently."*

— Rick Stuart, Manager, UNIX and Storage System, Anadarko Petroleum Corp.

2.4.1.15 What are the typical bandwidth requirements of the replication technology after initial seeding of the data to the remote site has been completed?

Vendor Response: SnapMirror is IP Based Thin Replication that utilizes the fact it is Dedupe aware, Data OnTap's Compression features to extremely minimize bandwidth it requires.

### **Reduce Storage Requirements and Bandwidth Utilization**

Disaster recovery solutions hinge on storage availability and network speed. To implement a successful DR solution and better meet RPOs, WVSTO must have sufficient storage and bandwidth available to house and transport DR data. Using asynchronous SnapMirror® together with the core technologies of Data ONTAP®, including IBM N-Series® Snapshot™ and deduplication, you can extend IBM N-Series space saving technologies to your DR storage environment. These efficiencies can help WVSTO:

- Reduce secondary storage requirements
- Minimize bandwidth utilization
- Accelerate data transfers

WVSTO can use Snapshot and deduplication to reduce the amount of data that is actually transmitted over the network. Replication based on Snapshot sends only changed blocks, and deduplication enables only unique blocks to be sent. WVSTO can further cut bandwidth needs and costs by using built-in SnapMirror network compression tools to compress data before sending it over the wire. With less data to send and lower bandwidth utilization, you can complete transfers more quickly and reduce both RTO and RPO.

*Figure 4: SnapMirror offers native network compression. Speed data replication, minimize bandwidth utilization, and reduce the performance impact on your network.*

To further reduce your network bandwidth utilization, SnapMirror automatically takes checkpoints during data transfers. If your system goes down, the transfer restarts from the most recent checkpoint.

WVSTO can also use SnapMirror to perform intelligent resynchronization and eliminate the need for full transfers when recovering from a broken mirror or loss of synchronization. If data on the mirrored copy was modified during application testing, it can be quickly resynchronized with the production data by copying the new and changed data blocks from the production system to the mirrored copy.

2.4.1.16 Is the data being replicated compressed or de-duplicated to reduce bandwidth requirements?



Vendor Response: Both if you like. In addition, SnapMirror will look at the data that is to be replicated and if the same blocks are already found at the DR, that data isn't sent and an addition pointer is created on the DR site to the block. This further reduces bandwidth requirement and is core to N-Series Dedupe functionality of both removing duplicate 4k blocks while also NOT creating duplicate data to begin with.

2.4.1.17 Is the data being replicated encrypted between the source and destination arrays?

Vendor Response: Yes

2.4.1.18 Does the replication technology support RPO goals of 15 minutes or less using asynchronous replication to a remote site? Response should detail any bandwidth or latency requirements to meet this goal.

Vendor Response: YES – We recommend Asynchronous replication in most cases, the most flexible and cost-effective solution, can meet RPOs ranging from one minute to 15 minutes without any limits on distance

2.4.1.19 Does the array have the capability to serve as a CAS/WORM device to replace optical storage systems? If so, what level of compliance does the CAS functionality provide?

Vendor Response: YES

IBM N-Series SnapLock® enables WVSTO to meet the strictest data retention regulations or to comply with internal IT governance policies by committing files to WORM storage for compliance. IBM N-Series SnapLock is a flexible, write once, read many (WORM) solution that allows WVSTO to create nonrewritable, nonerasable volumes to prevent files from being altered or deleted until a set retention date. Any file that WVSTO needs to archive—e-mails, documents, audit information, or other data—SnapLock can put in an unalterable state for as long as required. SnapLock files carry a retention date that cannot be shortened but can be lengthened. Even if those files are replicated or backed up to a different storage system or different type of storage system, WORM compliance remains intact. This makes SnapLock ideal for distributed businesses. The level of compliance is US SEC Rule 17a-4 that require the use of nonerasable storage. Only an act of willful destruction, such as physically removing disks from a NetApp® system, can result in record deletion prior to the specified expiration date,

## 2.4.2 The following sets of questions are explanation based, concerning the Server Hardware that may be proposed.

- a. Do the proposed servers support 16GB DIMMs?

Vendor response: yes or no

Ans. YES – The HS23 Blade has up to 16 DDR3 DIMM sockets (eight DIMMs per processor) using very low profile (VLP) DIMMs. Four memory channels per processor (two DIMMs per channel).

- b. Do the servers support 32GB DIMMs?

Vendor response: yes or no

Ans. No – Not necessary to meet the 192GB requirement. If more than 192GB is required than 8GB DIMM's can be replaced with additional 16GB DIMM's allowing for 256GB. Note: If more than 192GB is required per server (96GB per Processor Socket) than additional VMware licenses will be required because of vRAM limits.

- c. Do the proposed servers contain more than the required minimum of 192GB RAM per server?

Vendor response: yes or no

Ans. No – Note: If more than 192GB is required per server (96GB per Processor Socket) than additional VMware licenses will be required because of vRAM limits.

2.4.2.1 What is the total available processing power of the servers in the proposed solution? Please provide a breakdown on core count, core speed and total processing power (GHz) for the proposed servers.

Vendor Response:

Ans. HS23 will have two Xeon Six Core E5-2620 95W 2.0GHz/1333MHz/15MB, Two QPI links up to 8.0 GT/s each.

2.4.2.2 How many DIMMs can the servers in the proposed solution hold (without add-ons)?

Vendor Response:

Ans. HS23 has up to 16 DDR3 DIMM sockets (eight DIMMs per processor) using very low profile (VLP) DIMMs. Four memory channels per processor (two DIMMs per channel).

2.4.2.3 Are add-ons (drawers, trays, add-on blades, etc.) available to increase the number of DIMMs that can be installed in a server? If so, what is the maximum number of DIMMs that can be installed in the servers with any available add-ons?

Vendor Response:

Ans. No – Note: If more than 192GB is required per server (96GB per Processor Socket) than additional VMware licenses will be required because of vRAM limits.

2.4.2.4 What is the maximum RAM supported by the servers without add-ons (drawers, trays, etc.)? With add-ons?

Vendor Response:

Ans. HS23 Blade can support internally up to 256 GB with 16 x 16GB RDIMMs and two processors.

2.4.2.5 What size and speed DIMMs are being used in the proposed server configuration? Response should provide a detailed description of the RAM layout utilized on the servers.

Vendor Response:

Ans. In order to maximize performance and satisfy the 192GB requirement the HS23 has 8 x 8GB (1x8GB, 2Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP RDIMM and 8 x 16GB (1x16GB, 2Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP RDIMM's. DIMM's will be populated two per channel across four channels per processor.

2.4.2.6 If proposing Rack Mount Servers:

2.4.2.6.1 How many available PCI-Express slots do the servers in the configuration have?

Vendor Response:

Ans. N/A – Recommending Blade Servers

2.4.2.6.2 What is the speed of the PCI-Express slots in the servers? Response should provide a detailed listing of the available PCI-e expansion slots and their speeds and note which are already populated.



Ans. N/A – Recommending Blade Servers

#### 2.4.2.7 If proposing Blade Servers.

##### 2.4.2.7.1 How many total slots are in the proposed chassis?

Vendor Response:

Ans. IBM BladeCenter 'H' Chassis being proposed has fourteen blade slots available. With the current requirement we are quoting two chassis (one at Primary DC & one at DR Site). Each chassis will have only Seven Blade slots consumed. Leaving WVSTO with seven slots of additional Blade growth.

##### 2.4.2.7.2 Are the blade chassis in this proposal equipped with all required power supplies, fans and I/O modules/switches to support fully populating the blade chassis without additional cost beyond the purchase of the blade servers?

Vendor Response:

Ans. YES – All components of the IBM BladeCenter 'H' Chassis and HS23 Blades being quoted have redundant Power, Fans, I/O Modules/Switches and Management Modules. Including PFA (Predictive Failure Alerts) for CPU, Memory, Power Supplies, Fans/Blowers, Hard Drives, & Voltage Regulator Modules. All of which can trigger VM migration via VMotion utilizing IBM Systems Director and it's additional extensions/plugins.

##### 2.4.2.7.3 How many slots are used by the servers included in this proposal?

Vendor Response:

Ans. IBM BladeCenter 'H' Chassis being proposed has fourteen blade slots available. With the current requirement we are quoting two chassis (one at Primary DC & one at DR Site). Each chassis will have only Seven Blade slots consumed. Leaving WVSTO with seven slots for additional Blade growth.

#### 2.4.3 The following sets of questions are explanation based, concerning the Network Switches that may be proposed.

a. Due to the core competency of the WVSTO staff as well as other VW state agencies we would prefer to continue utilizing Cisco networking equipment within our data center for Ethernet connectivity. Does the proposed solution include Cisco network equipment?

Vendor response:

Ans: No – IBM Networking products proposed allow seamless, standards-based integration into existing Cisco and other networks reducing down time and learning curve. For network virtualization, VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement and reconfigures VM network policies across VLANs to keep the network up and running without impacting traffic or performance.

b. Does the proposed solution include licenses for VMware distributed virtual switch modules to allow both the physical and virtual network infrastructure to be managed through a common interface (whether command line, browser-based GUI, etc.)?

Vendor response:

Ans: YES – HS23 Blades have the IBM SN Distributed Virtual Switch V5000 which provides a managed, distributed virtual switch with advanced networking capabilities for VMware vSphere 5.0 Enterprise Plus. It replaces the native distributed virtual switch, making it visible to the administrators just like any other IBM physical switch. Advanced networking capabilities such as Ethernet Virtual Bridge (IEEE 802.1Qbg) and load balancing, and advanced troubleshooting capabilities such as SPAN and ERSPAN make it ideally suited for large-scale Virtual Machine deployments such as WVSTO in planning for.

2.4.3.1 Does the network equipment for server connectivity in the proposed solution have expansion capabilities (port modules, etc.), and, are those expansion slots available for future use or populated as part of the proposed solution? If yes, response should detail the expansion capabilities of the proposed network switches.

Vendor Response:

Ans: YES – The IBM System Networking G8264 being proposed comes with Forty-eight ports 10 GbE, Four 40 GbE, One 10/100/1000 Ethernet RJ45 port for out-of-band management. G8264



*delivers exceptional performance, being both lossless and low latency. In addition it has a feature-rich design standard with key virtualization features such as CEE/FCoE, high availability, and enterprise class Layer 2 and Layer 3 functionalities. This switch also provides the full range of Layer 3 protocols from static routes for technologies such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) for enterprise customers.*

2.4.3.2 Does the proposed network equipment include, or have the capability to support, other network protocols, specifically FCoE (fibre channel over Ethernet) and iSCSI? If yes, response should outline any additional modules or license costs to enable the support of these protocols on the proposed network switches.

Vendor Response:

*Ans: YES - The G8264 supports the newest protocols including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) for support of Fibre Channel over Ethernet (FCoE) in addition to iSCSI and NAS. It is designed to support CEE and connectivity to FCoE gateways. CEE will help enable WVSTO to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE will help enable highly efficient block storage over Ethernet for consolidating server network connectivity. As a result, clients can deploy a single server interface for multiple data types, which can simplify both deployment and management of server network connectivity, while maintaining the high availability and robustness required for storage transactions.*

2.4.4 The following sets of questions are explanation based concerning the General Solution being proposed.

a. Does the proposed solution include a centralized, unified monitoring system that gives overall status information about the hardware included in the solution (switches, storage and servers)?

Vendor Response: yes or no

*Ans: YES – IBM Systems Director enables integration with Tivoli, and third party management platforms, providing the building block for integrated services management. It automates data center operations by implementing cloud-ready virtual infrastructures. Unifies the management of physical and virtual resources for IBM servers, storage and networks and simplifies the management of optimized systems. Achieve a single view of the actual energy usage throughout your data center.*

b. Does the proposed solution include a single point of contact for all support issues (hardware and software) when utilized to run a vSphere environment?

Vendor Response: yes or no

*Ans: YES - Benefits to WVSTO of having ALL the components from one vendor (Servers, LAN Infrastructure & SAN Infrastructure) hardware support is just a single phone number away. Even though WVSTO has not requested software (ref. Sever Addendum to questions) for either Hardware, OS or Application(s) IBM can if needed provide WVSTO with a single phone number for all your support needs.*

c. Does the proposed solution include direct OEM support from the vendors of each component utilized in the solution to allow escalation of support issues to the OEM technicians by either the WVSTO or our single point of support for the proposed solution?

Vendor Response: yes or no

*Ans: YES – Benefits to WVSTO of having ALL the components from one vendor (Servers, LAN Infrastructure & SAN Infrastructure) hardware support is just a single phone number away. Even though WVSTO has not requested software (ref. Sever Addendum to questions) for either Hardware, OS or Applications IBM can if needed provide WVSTO with a single phone number for all your support needs.*

d. Does the proposed solution include regular (quarterly or bi-annually), pre-tested and validated firmware updates direct from a single source to allow the WVSTO to keep all hardware in the solution up-to-date without having to go through internal research, testing and validation of firmware as it is released by the OEMs?

Vendor Response: yes or no

*Ans: YES – IBM Systems Director update manager will acquire, install, and manage updates to Systems Director, its components, and the systems it manages. Updates can include new functions as well as fixes to certain firmware, operating systems, and platform managers.*

e. The proposed solution should take into consideration existing WVSTO licensing and should only include software licensing that is necessary to support the proposed solution that is not already owned by the



WVSTO (see appendix for list of current VMware licensing). Have you taken existing WVSTO licensing into account and only included additional licenses, not already owned by the WVSTO in your proposed solution?

Vendor Response: yes or no

Ans: YES – *The proposed solution will allow WVSTO to grow beyond current Server & Desktop requirements.*

2.4.4.1 Does the proposed solution ship as a single unit (all hardware racked, all internal power, network, SAN and other cables connected) ready to connect to power and core networking equipment and begin deployment and configuration of storage, networking and the vSphere environment?

Vendor Response:

Ans: YES – *With the racks being proposed this solution can be built and shipped fully integrated or while on site.*

2.4.4.2 The WVSTO would like to keep the network traffic for the hosts, the network traffic for hardware management and the storage network traffic separated. This serves a few purposes, the first being segregation of traffic with dedicated resources for each type of traffic, to try and insure peak performance of the solution; the second being the ability to keep the management traffic on high performance (gigabit), but lower-cost switches that don't need the capabilities of the switches used to connect the VMware hosts to the network.

Vendor Response:

Ans: YES – *Three separate infrastructures have been proposed (mgmt LAN, Production LAN & SAN).*

2.4.4.2.1 Does the proposed solution include separate switch infrastructure for the hosts, the hardware management interfaces and storage (fibre channel) networks? If yes, response should provide some details on the internal network layout of the proposed solution and how it meets this goal.

Vendor Response:

Ans: YES – *Switches for three separate infrastructures have been quoted (1GbE for mgmt LAN, 10GbE for Production LAN & 8GB Fibre for Storage).*

2.4.4.3 Does the proposed solution include a centralized, unified management system that allows baseline configuration tasks to be performed? If it does, can the following tasks be performed through this management system? If so, response should outline the following capabilities to perform that function:

Vendor Response:

Ans: YES – *IBM Systems Director enables integration with optional third party management platforms, providing the building block for integrated services management. It automates data center operations by implementing cloud-ready virtual infrastructures. Unifies the management of physical and virtual resources for IBM servers, storage and networks and simplifies the management of optimized systems. Achieve a single view of the actual energy usage throughout your data center.*

2.4.4.3.1 Define VLANs available (trunked) into the network switches from the core network

Vendor Response:

Ans: YES – *The G8264 includes up to 40 Gb uplinks. Trunking via LACP, Static Trunks (EtherChannel), Configurable Trunk Hash algorithm. The chassis BNT Virtual Fabric 10 Gb Switch Module for IBM BladeCenter supports up to ten 10 Gb SFP+ ports per switch, clients can easily support 4 - 6 chassis per Rack Switch, and if they are not concerned about oversubscription, they can support up to 24 or more per Rack Switch G8264. With 1.2 Tbps throughput, the G8264 provides massive scalability and low latency ideal for latency-sensitive applications such as high-performance computing clusters and financial applications.*

2.4.4.3.2 Define storage available to the various vSphere clusters

Vendor Response:

Ans: YES – *The two clusters are based on the pre-defined storage specification.*

2.4.4.3.3 Deployment of operating system (vSphere, Windows, etc.) to the physical servers included in the solution from user-provided ISO images

Vendor Response:

Ans: YES – *Tivoli Provisioning Manager for OS Deployment - IBM Systems Director Edition provides better management and remote deployment of operating systems on IBM Systems Director. Enables rapid provisioning, through accelerated and automated OS and software management on Systems Director. Leverages industry standards and provides a unmatched*

flexibility to automate the installation of operating systems and selected software on tens, or even hundreds, of computers simultaneously. Includes a comprehensive Web-based console, allowing configuration of the provisioning server. Offering a full set of operating system provisioning features designed to address a wide array of WVSTO specific needs. Provides cloning and unattended setup for Windows and Linux, servers enabling ready to use installation of operating systems and selected software on hundred of computers simultaneously. Saves travel and labor cost. Implements a PXE network boot application. PXE is a type of DHCP extension. It is supported by the vast majority of hardware vendors, including IBM, Lenovo, HP, and DELL.

#### 2.4.4.3.4 Creation of vCenter instances to manage vSphere hosts

Vendor Response:

Ans: YES – IBM Systems Director included in this WVSTO proposed solution provides a foundation for systems management. It provides the integrated tools needed to efficiently visualize and communicate the relationships of physical and virtual systems that are discovered, monitor their health, define and receive threshold alerts, and update system firmware and operating environments. Built-in automation capabilities enable IT administrators to schedule updates and configuration changes to proactively avoid problems, and reduce the administrative burden of routine maintenance. Advanced OS deployment, monitoring and control features under the same console. WVSTO can create and deploy virtual images that encapsulate the Microsoft Windows or Linux operating environment as well as any middleware and applications, while automatically configuring the storage attached network (SAN). If WVSTO is facing energy cost or capacity concerns you will have the benefit of extensive electrical power and thermal monitoring capabilities across the data center, as well as energy control of selected Power servers and networking systems.

#### 2.4.4.3.5 Does the management system provide any additional capabilities not outlined above? If it does, response should detail any notable capabilities.

Vendor Response:

Ans: YES – IBM Systems Director: •Unifies the essential management of IBM servers, storage and network devices delivering a consistent look and feel for common management tasks that reduces operational complexity •Integrates IBM's best-of-breed virtualization capabilities to provide new ways to simplify the management of physical and virtual platform resources •Reduces energy costs and usage by monitoring and managing the energy and cooling needs of servers and storage •Easy integration with enterprise service management tools from Tivoli® as well as other third-party providers •IBM® Systems Director Editions provide options to select the level of leading management solutions based on your needs.



## Attachment B: Mandatory Specification Checklist

---

The mandatory requirements in Section 2.5 will be deliverables upon award of this RFP.

An affirmation to each requirement is required; if you cannot mark "yes or V" to each requirement, you will be considered non-responsive and disqualified.

### 2.5 Mandatory Requirements

The following mandatory requirements must be met by the Vendor as a part of the submitted proposal. Failure on the part of the Vendor to meet any of the mandatory specifications shall result in the disqualification of the proposal. The terms "must", "will", "shall", "minimum", "maximum", or "is/are required" identify a mandatory item or factor. Decisions regarding compliance with any mandatory requirements shall be at the sole discretion of the State.

2.5.1.1 The unified storage systems must allow presentation of storage through block and file level protocols and meet the following requirements for usable capacity.

Affirm: **Yes**

2.5.1.1.1 The storage array for the production center must provide a minimum usable capacity of at least 17TB for virtualized servers in a dedicated physical or virtual storage pool.

Affirm: **Yes**

2.5.1.1.2 The storage array for the production data center must provide a minimum useable capacity of 5 TB for NAS file shares in a dedicated physical or virtual storage pool.

Affirm: **Yes**

2.5.1.1.3 The storage array for the production data center must provide a minimum useable capacity of 3 TB for virtual desktops in a dedicated physical or virtual storage pool.

Affirm: **Yes**

2.5.1.1.4 The storage array for the disaster recovery data center must provide a minimum useable capacity of 17 TB for replicated virtual servers.

Affirm: **Yes**

2.5.1.1.5 The storage array for the disaster recovery data center must provide a minimum useable capacity of 5 TB for replicated NAS file shares.

Affirm: **Yes**

2.5.1.1.6 The storage array for the disaster recovery data center must provide a minimum useable capacity of 3 TB for replicated virtual desktops.

Affirm: **Yes**

2.5.1.1.7 The storage array for the disaster recovery data center must provide a minimum additional useable capacity of 10 TB.

Affirm: **Yes**

2.5.1.2 The proposed storage array must be a unified storage array that allows presentation of storage via block (Fibre Channel) and file (CIFS, NFS) protocols.

Affirm: Yes

2.5.1.3 The proposed storage array must have a minimum of 4Gbps fibre channel connectivity to the SAN switch infrastructure.

Affirm: Yes

2.5.1.4 The proposed storage array must have two storage controllers for the block level protocol in an active/active configuration with at least two fibre channel connections to the SAN switch infrastructure providing a total of 4 paths to the storage array.

Affirm: Yes

2.5.1.5 The proposed storage array must have two filers for the file level protocols in an active/passive or active/active configuration with at least two (2) 1Gbps or two (2) 10Gbps Ethernet connections per filer to the network infrastructure.

Affirm: Yes

2.5.1.6 The proposed storage array for the production data center must provide a minimum of 20,000 IOPs dedicated to the virtualized server environment.

Affirm: Yes

2.5.1.7 The proposed storage array for the production data center must provide a minimum of 8,000 IOPs dedicated to the virtual desktop environment.

Affirm: Yes

2.5.1.8 The proposed storage array for the production data center must provide dedicated capacity to support NAS file shares for up to 120 users and 3 TB of data.

Affirm: Yes

2.5.1.9 The proposed storage array for the disaster recovery site must provide a minimum of 60% of the total IOPs of the production storage array.

Affirm: Yes

2.5.1.10 The proposed storage array must support Solid State Drives (SSD) or Enterprise Flash Drives (EFD) (Tier 0).

Affirm: Yes

2.5.1.11 The proposed storage array must support high speed (10K and 15K RPM) Fibre Channel (FC) or Serial Attached SCSI (SAS) drives (Tier 1 and Tier 2).

Affirm: Yes

2.5.1.12 The proposed storage array must support 7.2K RPM near-line SAS or ATA drives (Tier 3).

Affirm: Yes

2.5.1.13 The proposed storage array must support virtual (thin) provisioning for volumes presented via block level (FC) protocol.

Affirm: Yes

2.5.1.14 The unified storage systems must support the ability to do snapshots and clones of volumes presented via block level protocols. It must also support the ability to do snapshots of the file systems presented via file level protocols.

Affirm: Yes

2.5.1.15 The proposed storage array must include the ability to make clones of volumes presented via block-level (FC) protocol.

Affirm: Yes

2.5.1.16 The proposed storage array must include the ability to take snapshots of volumes presented via block-level (FC) protocol.

Affirm: Yes

2.5.1.17 The proposed storage array must include the ability to take snapshots of file systems presented via file-level protocols (CIFS, NFS).

Affirm: Yes

2.5.1.18 The proposed storage array must include IP-based, asynchronous replication for the storage presented via block level (FC) protocol.

Affirm: Yes

2.5.1.19 The proposed storage array must include IP-based, asynchronous replication for the file systems presented via file level (CIFS, NFS) protocols.

Affirm: Yes

2.5.1.20 The proposed storage array must have the capability to support data-at-rest encryption.

Affirm: Yes

2.5.1.21 The proposed storage array must have a single, unified management tool that allows the configuration and monitoring of all features and functionality of the array.

Affirm: Yes

2.5A.22 The proposed storage array must support all of the primitives defined in the VMware vSphere API for Array Integration (VAAI) specifications for vSphere 5.0 for storage presented via block level (FC) protocol.

Affirm: Yes

2.5.1.23 The proposed storage array must include full, active-active, load balanced multi-path support for connected VMware vSphere 5.0 hosts (not the default most recently used or round robin provided by VMware).

Affirm: Yes

2.5.1.24 The proposed storage array must include plug-ins for VMware vCenter to enable the creation and management of LUNs (from assigned storage pools) for the vSphere environment to ensure proper alignment and optimization of the LUNs.

Affirm: Yes

2.5.1.25 The proposed array must include replication technology that integrates with VMware Site Recovery Manager (SRM) 5.0 to allow SRM to leverage the native replication technologies of the array to copy data to the disaster recovery site.

Affirm: Yes

2.5.1.26 The proposed array must have the capability to enable call-home functionality for sending hardware alerts to the OEM when failures are detected on the array to enable rapid, pro-active response from technical support to replace or repair defective hardware.

Affirm: Yes

2.5.1.27 The unified storage systems must have an expected product life of at least 5 years.

Affirm: Yes

2.5.1.28 The unified storage systems must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 availability coverage.

Affirm: Yes

## 2.5.2 Fibre Channel Switches

2.5.2.1 The proposed solution shall include two independent fibre channel switches at each site.

Affirm: Yes

2.5.2.2 The fibre channel switches must have autosensing 8 Gbps ports (support 8/4/2 Gbps).



Affirm: Yes

2.5.2.3 The proposed fibre channel switches must have management capabilities via a command line interface (telnet/SSH).  
Affirm: Yes

2.5.2.4 The proposed fibre channel switches must have a browser-based management interface.  
Affirm: Yes

2.5.2.5 The proposed fibre channel switches must include some internal diagnostics.  
Affirm: Yes

2.5.2.6 The proposed fibre channel switches must include native alerting and reporting (without the need for a monitoring server).  
Affirm: Yes

2.5.2.7 The proposed fibre channel switches must include a native way to display performance metrics.  
Affirm: Yes

2.5.2.8 The proposed fibre channel switch configuration must support non-disruptive firmware upgrades.  
Affirm: Yes

2.5.2.9 The proposed fibre channel switches must have the capability to be either an NPV edge device or an NPIV core device.  
Affirm: Yes

2.5.2.10 The proposed fibre channel switches must have the capability to support multiple fabric environments in a single physical switch.  
Affirm: Yes

2.5.2.11 The proposed fibre channel switches must support aggregated ISL (inter-switch link) connectivity; i.e., several physical ISLs behaving as one virtual ISL.  
Affirm: Yes

2.5.2.12 The proposed fibre channel switches must support traffic engineering using FSPF.  
Affirm: Yes

2.5.2.13 The fibre channel switches must have at least 12 ports active each.  
Affirm: Yes

2.5.2.14 The fibre channel switches must have at least 24 ports total each.  
Affirm: Yes

2.5.2.15 The fibre channel switches must have redundant power supplies and fans.  
Affirm: Yes

2.5.2.16 The fibre channel switches must have an expected product life of at least 5 years.  
Affirm: Yes

2.5.2.17 The fibre channel switches must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 coverage.  
Affirm: Yes

### 2.5.3 Network Switches

2.5.3.1 The network switch(es) must support both 10Gbps and 1Gbps connectivity.  
Affirm: Yes

2.5.3.2 The network switch(es) must have a minimum of 16 ports available for connection of additional network devices not included in the proposed solution.  
Affirm: Yes

2.5.3.3 The network switch(es) must have redundant power supplies and fans.  
Affirm: Yes

2.5.3.4 The network switch(es) used for server connectivity must include layer 3 support (if a dedicated management network is present it does not need to support layer 3).

Affirm: Yes

2.5.3.5 The network switch(es) must support Link Aggregation Control Protocol (LACP): IEEE 802.3ad.

Affirm: Yes

2.5.3.6 The network switch(es) must support VLAN trunking.

Affirm: Yes

2.5.3.7 The network switch(es) must support IEEE 802.1Q VLAN encapsulation.

Affirm: Yes

2.5.3.8 The network switch(es) must support Jumbo Frames on all ports (up to 9216 bytes).

Affirm: Yes

2.5.3.9 The network switch(es) must support CLI management (console, telnet and/or SSH).

Affirm: Yes

2.5.3.10 The network switch(es) must support SNMP.

Affirm: Yes

2.5.3.11 The network switches must have an expected product life of at least 5 years.

Affirm: Yes

2.5.3.12 The network switches must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 coverage.

Affirm: Yes

## **2.5.4 Server Hardware**

2.5.4.1 There must be at least 7 identically configured servers per site (production and DR), 14 servers in total.

Affirm: Yes

2.5.4.2 The proposed servers must be dual CPU socket servers.

Affirm: Yes

2.5.4.3 The proposed servers must use 6-core Intel 5600 series or 10-core Intel E7 series processors or superior.

Affirm: Yes

2.5.4.4 Each server must have at least 192GB of RAM.

Affirm: Yes

2.5.4.5 Each server must include a minimum of two (2) 10Gbps network connections.

Affirm: Yes

2.5.4.6 Each server must include a minimum of two (2) 8Gbps fibre channel (SAN) connections.

Affirm: Yes

2.5.4.7 The servers must include remote management capabilities (DRAC, iLO or equivalent).

Affirm: Yes

2.5.4.8 The servers must have fully redundant internal components (power supplies, fans, etc.).

Affirm: Yes

2.5.4.9 The servers must have an expected product life of at least 5 years.

Affirm: Yes

2.5.4.10 The servers must include 5 years of support with a guaranteed response time of 4 hours and 24x7x365 coverage.

Affirm: Yes

## **2.5.5 Rack Mount Servers (If this solution is proposed)**

2.5.5.1 All of the PCI-Express slots in the servers must run at a minimum of 4x speed.

Affirm: N/A

2.5.5.2 The servers must have at least two available PCI-Express slots for expansion capabilities.

Affirm: N/A

## **2.5.6 Blade Servers (If this solution is proposed)**

2.5.6.1 In the proposed blade solution the individual blade servers at each site must be split as evenly as possible across two blade chassis (elimination of single point of failure and provide extra expansion capabilities through number of available slots for blades).

Affirm: Yes

2.5.6.2 Each blade chassis must include fully redundant I/O and management modules.

Affirm: Yes

## **2.5.7 General Requirements for Proposed Solution**

2.5.7.1 The proposed solution must include the cabinet(s) for the equipment. The cabinet(s) must have locks and provide adequate air flow to keep installed equipment cool.

Affirm: Yes

2.5.7.2 The proposed solution must include all required power distribution units (PDUs) inside the cabinet(s) to provide power to all installed equipment.

Affirm: Yes

2.5.7.3 All equipment must be mounted in the cabinet(s).

Affirm: Yes

2.5.7.4 All wiring inside the cabinet must be complete. This includes power cables from equipment to PDUs and network and fiber interconnects between servers, storage and switches.

Affirm: Yes

2.5.7.5 The proposed solution must include maintenance agreements on the hardware and software included in the proposal for five years. These maintenance agreements must, at a minimum, provide break fix support, software upgrades and general technical support (knowledge base, trouble tickets) on the included components via phone and/or online support sites.

Affirm: Yes

2.5.7.6 The proposed solution must include only the professional services required to get the equipment into our data centers, installed in the cabinet(s) as outlined above and brought online and ready for end-user configuration. The vendor must also perform any configuration on the equipment (for example, the storage pools on the storage array and other array features such as cache settings) required to ensure the equipment meets the capacity and performance requirements outlined in this RFP. Any additional professional services such as configuration of replication features on the arrays, configuration of VMware Site Recovery Manager (SRM) to work with array replication or any assistance with data migration or other tasks will be acquired through a separate procurement(s).

Affirm: Yes

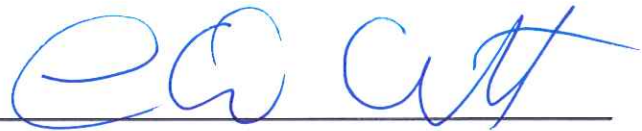


---

I certify that the proposal submitted meets or exceeds all the mandatory specifications of this Request for Proposal. Additionally, I agree to provide any additional documentation deemed necessary by the State of West Virginia to demonstrate compliance with said mandatory specifications.

**Software Information Systems, LLC**  
(Company)

**Charles D. Arnett, Senior Client Executive**  
(Representative Name, Title)



**304 768-1645/304 768-1671**  
(Contact Phone/Fax Number)

**4-10-2012**  
(Date)



---

### Highlights

- Delivers high performance, energy efficiency and versatility to run the most demanding applications in larger data centers
  - Reduces energy bill with leadership high-efficiency power supplies that are 94 percent efficient and 80 PLUS® Platinum Certified<sup>1</sup>
  - Provides more than 12 percent greater performance per watt compared to the competition<sup>2</sup>
  - Protects your investment by being compatible with the entire IBM BladeCenter® family
- 

## IBM BladeCenter H

*Extending the value of blades to higher performance in an enterprise environment*

In today's high-demand enterprise environment, organizations need a reliable infrastructure to run compute-intensive applications with minimal maintenance and downtime. IBM BladeCenter H is a powerful platform built with the enterprise customer in mind, providing industry-leading performance, innovative architecture and a solid foundation for virtualization.

### Powerful performance

This next-generation chassis is a workhorse—it runs applications quickly and efficiently, enabling organizations to maximize their investment and succeed in a highly competitive market. The powerful BladeCenter H delivers the performance customers need to make informed decisions.

With IBM Open Fabric Manager, BladeCenter H can help organizations run applications even more efficiently by delivering a flexible, open and connected infrastructure.

### Driving innovation

IBM continues to build on its reputation for helping customers find innovative IT solutions with BladeCenter H, which delivers even more capabilities to the BladeCenter family. The 9U chassis provides the



standard BladeCenter functionalities with the added performance of high-speed I/O and the fastest blades. Other highlights include:

- High-speed fabric with eight data channels to every blade, four of which can be 10 Gb
- Leadership high-efficiency power supplies
- Flexible, scalable design so the system can grow as your business grows
- Sustainable platform life for enhanced investment protection
- Multi-Switch Interconnect Module for additional connectivity to blade servers
- Light path diagnostics for fast, simple, accurate troubleshooting to help guard against system downtime

BladeCenter H also drives innovation through its open architecture, supporting industry-standard I/O switches from Cisco, Brocade and others, so organizations have more choice and flexibility when it comes to their IT environment. Likewise, IBM's affiliation with blade.org helps open the architecture even more by supporting an online, collaborative environment to promote advances in technology solutions.

### Managing growth, minimizing risk

With its unsurpassed architecture, BladeCenter H efficiently integrates servers, storage, networking, I/O and applications, enabling organizations to build flexible IT infrastructures using common building blocks. This helps customers control costs by protecting their investments—and with a powerful set of tools that are open and equally integrated, organizations can quickly deploy and easily manage their systems.

As a rack-dense, space saving chassis, BladeCenter H provides a superior foundation for virtualization. And in a virtualized environment, organizations can run a more flexible, better utilized and potentially more cost-effective IT infrastructure with a common management construct.




### Easy and energy efficient

Because BladeCenter H is more than 12 percent more energy efficient than the IIP C-class chassis,<sup>1</sup> organizations know they're making the right choice when it comes to using less power. Innovative tools such as IBM Systems Director Active Energy Manager™ track power usage so organizations can better manage energy and help prevent outages. And not only does BladeCenter H help organizations go green and save, it also makes it easy to deploy, integrate and manage applications—so customers can focus on their business and not on their IT.

### IBM: The right choice

As an industry leader in innovative IT systems, IBM delivers yet another important solution for the IT environment. BladeCenter H is the ideal chassis for running the most demanding applications and simulations at high speeds, giving enterprise organizations a distinct advantage over competitors.





**Run Cool  
Run Reliably  
Run With 80 PLUS®**

[www.80PLUS.org](http://www.80PLUS.org)

The BladeCenter family of products includes five chassis—BladeCenter E, BladeCenter H, BladeCenter HT, BladeCenter S and BladeCenter T—that meet a wide range of customer needs.

BladeCenter H delivers high performance by supporting 4X InfiniBand and 10 Gb Ethernet switches. It also supports the largest power supplies and the entire IBM family of blade servers.

BladeCenter H helps reduce energy costs with leadership high-efficiency power supplies that are 94 percent efficient and 80 PLUS Platinum Certified.

#### IBM BladeCenter H at a glance

Form factor/height	Rack-mount chassis/9U
Blade bays	Up to 14
Switch modules	Up to four traditional, up to four high-speed, and up to four bridge module bays
Power supply module	Up to four hot-swap and redundant 2980 W ac high-efficiency power supplies with load-balancing and failover capabilities. Operating at 200-240 V
Cooling modules	Two hot-swap and redundant blowers standard, additional fan packs on power supplies
Systems management hardware	Advanced Management Module standard; add an optional Advanced Management Module for redundancy
I/O ports	USB-based keyboard, video, mouse (KVM), Ethernet, USB
Media	Two USB connections and an optional DVD multiburner
Systems management software	Systems management and trial deployment tools
Predictive Failure Analysis	Internal storage, processors, blowers, memory
Light path diagnostics	Blade server, processor, memory, power supplies, blowers, switch module, management module, internal storage and expansion card
Limited warranty	Three-year customer replaceable unit and onsite limited warranty
External storage	Support for IBM System Storage® solutions (including DS and NAS family of products) and many widely adopted non-IBM storage offerings

## For more information

### World Wide Web

U.S. [ibm.com/systems/bladecenter](http://ibm.com/systems/bladecenter)

Canada [ibm.com/systems/ca/en/bladecenter](http://ibm.com/systems/ca/en/bladecenter)



---

© Copyright IBM Corporation 2010

IBM Systems and Technology Group  
Route 100  
Somers, New York 10589

Produced in the United States  
July 2010  
All Rights Reserved

IBM, the IBM logo, [ibm.com](http://ibm.com) and BladeCenter are trademarks or registered trademarks of IBM Corporation in the United States, other countries or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Other company, product and service names may be trademarks or service marks of others.

<sup>1</sup> 80 PLUS Organization, 2010, visit [www.80plus.org](http://www.80plus.org) for more information.

<sup>2</sup> Edison Group, Inc., 2010, visit [www.TheEdison.com](http://www.TheEdison.com) for more information.




Please Recycle



---

### Highlights

- Provides outstanding performance with the latest Intel Xeon processor E5-2600 product family and 1600 MHz memory
- Integrated Virtual Fabric 10 Gigabit Ethernet (GbE) to gain flexibility and easy scalability



## IBM BladeCenter HS23

*Next-generation two-socket system optimized for performance and scalable networking*

Do more in the data center you own with IBM® BladeCenter® HS23. This efficient server helps you manage big data and is ideal for a wide variety of workloads including virtualization and cloud infrastructure solutions. New integrated 10 GbE Virtual Fabric provides high networking bandwidth. Sixteen DIMM slots, supporting up to 256 GB of DDR3 memory, deliver more and larger virtual machines per blade compared to previous generations. Simplify infrastructure deployment and complexity for faster time-to-value with IBM FastSetup.

### Outstanding performance

The HS23 offers outstanding performance, with support for the new 8-core Intel Xeon processor E5-2600 product family that processes up to 16 simultaneous threads. It has fast memory throughput and integrated RAID-0, -1.


### Flexible, scalable networking

The HS23 provides integrated 10 GbE Virtual Fabric, delivering increased high-speed performance and I/O flexibility, running multiple protocols (FCoE/iSCSI), and freeing your CFFh slot to do more. Connect to any networking by choosing from a variety of adapters to suit your data center needs.

### Investment protection

Protect your investment by integrating the HS23 into your existing infrastructure. The HS23 is backward compatible with your existing BladeCenter chassis. It supports multiple technologies enabling zero rip-and-replace, and allows you to do more now or later as your business grows.

Select configurations of the HS23 are part of the IBM Express Portfolio™ designed to meet the needs of small and mid-sized businesses. Easy to manage, Express models vary by country.





### IBM BladeCenter HS23 at a glance

Form factor/height	Single-width (30 mm)
Processor (max)	Intel Xeon processor E5-2600 product family, with up to 8 cores, processing up to 16 simultaneous threads
Number of processors (std/max)	1/2
Cache (max)	20 MB
Memory (max)	256 GB at 1333 MHz (assuming 16 GB DIMMs), 16 DIMM sockets - VLP Registered ECC DDR3, at 1333 or 1600 MHz, memory mirroring and sparing supported
Expansion slots	1 CIOv slot (standard PCIe daughter cards) and 1 CFFh slot (high-speed PCIe daughter card)
Disk bays (total/hot-swap)	2 hot-swappable HDDs (SATA/SAS) or SSDs
Maximum internal storage	2 TB
Network interface	Integrated Virtual Fabric for BladeCenter, Emulex BE3 Dual 10 Gb/1 Gb
RAID support	RAID-0, -1
System management	Unified Extensible Firmware Interface (UEFI), IBM Integrated Management Module, Predictive Failure Analysis, optional embedded hypervisor for virtualization, IBM Systems Director Active Energy Manager™, light path diagnostics, IBM Systems Director, IBM Server Guide, UpdateXpress, Dynamic System Analysis with Integrated RT Diagnostics, BladeCenter Open Fabric Manager
Operating systems supported	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise, VMware
Limited warranty	3-year customer replaceable unit and onsite limited warranty, next business day 9x5, service upgrades available

### For more information

To learn more about the HS23 blade visit:

[ibm.com/systems/bladecenter/hardware/servers/hs23/index.html](http://ibm.com/systems/bladecenter/hardware/servers/hs23/index.html)

or contact your IBM representative or IBM Business Partner.



© Copyright IBM Corporation 2012

IBM Systems and Technology Group  
Route 100  
Somers, New York 10589

Produced in the United States of America  
March 2012

IBM, the IBM logo, ibm.com, and BladeCenter are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows Server are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product, company or service names may be trademarks or service marks of others.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.



Please Recycle

## IBM System Storage N6000 series, Filer and Gateway



*Easily respond to current and future storage demands  
with advanced capabilities in a midrange storage system*

---

### Highlights

- Increase NAS storage flexibility and expansion capabilities by consolidating block and file data sets onto a single multiprotocol storage platform
  - Get performance when your applications need it most with high bandwidth, 64-bit architecture and the latest I/O technologies
  - Maximize storage efficiency and growth and preserve investments in staff expertise and capital equipment with data-in-place upgrades to more powerful IBM System Storage N series
  - Improve your business efficiency by taking advantage of the N6000 series capabilities, also available with a Gateway feature, to reduce data management complexity in heterogeneous storage environments for data protection and retention
- 


Today's business environment demands innovation and flawless execution. You have to manage and protect valuable data to support business growth and success. Your IT operations have to evolve with the business while meeting budget, staffing and infrastructure limits. Virtualized computing requires networked storage systems supporting diverse data sets to unlock the full potential of virtualized servers.

With IBM® System Storage® N6000 series systems, you can meet your Network Attached Storage (NAS) needs and provide high levels of application availability for everything from critical business operations to technical applications. You can also address NAS and Storage Area Network (SAN) as primary and secondary storage requirements. In addition, you get outstanding value—our flexible systems offer excellent performance and impressive expandability at a low total cost of ownership.

N series systems enable easy provisioning, managing and upgrading so you can quickly adapt your storage infrastructure to meet your changing business and technical needs. To help you maximize staff productivity, all N series systems use the Data ONTAP operating system and the same suite of application-aware management software. Also, OnCommand enables the consolidation and simplification of shared IT storage management.

### Versatility for your diverse business needs

The N6000 series systems offer a versatile storage platform for handling the large amounts of diverse data moving through your business. With an N6000 series system, you can consolidate varied data sets simultaneously—whether block- or file-based—onto a single storage platform.





With N6000 series, you can unlock the full potential of your growing virtualized server environment by enabling virtual machine mobility and offloading the work of data protection. The N6000 systems enable you to connect your heterogeneous server environment (including Microsoft Windows, UNIX, and Linux servers) and clients to one storage system by using standard storage protocols and interfaces.


### Increase data and application availability

N6000 series systems can help you spend less time on backup and recovery, so you can focus on growing your business. A full range of enterprise-class, high-availability and disaster-recovery products provides affordable support for data protection to safeguard your business-critical application data. N series Snapshot technology reduces backup times to minutes. SnapRestore software also enables recovery of point-in-time data in minutes.

N series SnapManager software quickly returns applications to the same point in time as recovered data. All of this capability is built on the solid foundation of our low-overhead, dual-parity RAID-DP—the N series implementation of high-performance RAID 6 for data protection and capacity use.

### Performance when your applications and users need it

The N6000 series offers extraordinary performance to help you meet demanding service levels and get your products and services to market faster. The high bandwidth, 64-bit controller architecture with large memory cache and the latest I/O technologies provides data at the rates you need to keep your demanding business and technical applications running smoothly. Your critical applications can take priority under peak load conditions with FlexShare quality of service software.



The Performance Acceleration Module (Flash Cache), an intelligent read cache, improves throughput and reduces latency to optimize the performance of your storage system.

### Respond to your data growth challenges

In today's business environment, it seems the data your systems collect grows relentlessly, regardless of your company's size. With versatile N6000 series systems, you can combine high-performance Fibre Channel and large-capacity Serial Advanced Technology Attachment (SATA) disk drives in storage tiers to optimize performance and cost. In addition, you can seamlessly consolidate block and file storage on the same system. N series makes this possible by providing native support of the Network File System (NFS) and Common Internet File System (CIFS), Fibre Channel over Ethernet (FCoE), Fibre Channel Protocol (FCP), Internet Small Computer System Interface (iSCSI) storage protocols through both Fibre Channel and Ethernet interfaces.



Innovative thin provisioning capability eliminates stranded storage by expanding or contracting logical unit numbers (LUNs) and volumes through a common pool of spare capacity, without IT staff intervention. When more performance or expandability is required, you can preserve your investment by installing a more powerful N series controller that enables you to keep your data in place and use the same management tools.

### Maximize your resources

N6000 series systems can help you reduce costs in many aspects of your storage environment by simplifying data management and maximizing storage use to conserve raw storage, power, cooling and data center space. N6000 systems can help you spend less time waiting and more time innovating, due to high system performance, fast backup and recovery, and rapid cloning of data sets.

### Improve your business efficiency

You can take advantage of the N6000 series capabilities, ordered with a Gateway feature, in heterogeneous storage environments to help improve business efficiency and reduce data management complexity. N6000 series systems ordered with a Gateway feature can support attachment to IBM Enterprise Storage Server® (ESS) series, IBM XIV® Storage System, IBM System Storage DS8000® and DS5000 series and a broad range of IBM, EMC, Hitachi, Fujitsu and HP storage subsystems.

### Automation, consolidation and simplification

N6000 systems are empowered with OnCommand, a multiprotocol manager that delivers a single code experience to manage physical and virtual storage environments using integrated workflows and policy-driven automation. From a single interface, OnCommand enables the consolidation and simplification of shared IT storage management, delivering greater flexibility and efficiency.

Also, the capabilities of expandability, connectivity, data protection and retention, and copy recover availability integrate with leading IT management frameworks. The software provides common management services, integration, security and role-based access controls. OnCommand offers data and storage management tools that increase productivity, storage efficiency and agility for organizations of all sizes. It supports and integrates with higher-level IT orchestration and management frameworks, which helps you manage data from a business perspective, enabling administrators to manage data across applications, databases, servers and storage.

---

#### Software

Operating system	Data ONTAP
Operating systems supported	Windows 2000, Windows Server 2003, Windows Server 2008, Windows XP, Linux, Oracle Sun Solaris, IBM AIX®, HP-UX, Apple Macintosh OS, VMware ESX
Software features	See <a href="http://ibm.com/systems/storage/network/software">ibm.com/systems/storage/network/software</a> for a full list of software features

## Specifications

	N6210	N6210	N6240	N6240	N6240	N6270	N6270	N6270
Machine type model	2858-C10	2658-C20	2858-C21	2858-E11	2858-E21	2858-C22	2858-E12	2858-E22
Gateway Machine type model		2658-C20 (with feature code 9551)	2858-C21 (with feature code 9551)		2858-E21 (with feature code 9551)	2858-C22 (with feature code 9551)		2858-E22 (with feature code 9551)
Controller configuration	Single (C)	Dual (active/ active) (CC)	Dual (active/ active) (CC)	Single + IO Exp (CI)	Dual + IO Exp (active/ active) (CI-HA)	Dual (active/ active) (CC)	Single (CI)	Dual (active/ active) (CI-HA)
Processors speed and type	2.3 GHz Intel (Dual Core)		2.3 GHz Intel (Quad Core)			3.00 GHz Intel (Quad Core)		
Number of processors (cores)	2	4	8	4	8	8	4	8
Random access memory	4 GB	8 GB	16 GB	8 GB	16 GB	32 GB	16 GB	32 GB
Nonvolatile memory	512 MB	1 GB	2 GB	1 GB	2 GB	4 GB	2 GB	4 GB
<b>Integrated I/O Ports</b>								
Fibre Channel ports (speed)	2 (4 Gbps)	4 (4 Gbps)	4 (4 Gbps)	2 (4 Gbps)	4 (4 Gbps)	4 (4 Gbps)	2 (4 Gbps)	4 (4 Gbps)
Ethernet ports (speed)	2 (1 Gbps)	4 (1 Gbps)	4 (1 Gbps)	2 (1 Gbps)	4 (1 Gbps)	4 (1 Gbps)	2 (1 Gbps)	4 (1 Gbps)
SAS ports (speed)	2 (6 Gbps)	4 (6 Gbps)	4 (6 Gbps)	2 (6 Gbps)	4 (6 Gbps)	4 (6 Gbps)	2 (6 Gbps)	4 (6 Gbps)
<b>Storage scalability</b>								
Maximum number of Fibre Channel loops	10	10	8	13	26	8	13	26
Maximum raw capacity	720 TB	720 TB	1800 TB	1800 TB	1800 TB	2880 TB	2880 TB	2880 TB
Maximum number of disk drives	240	240	600	600	600	960	960	960
Maximum volume size	16 TB (32-bit) 75 TB (64-bit)	16 TB (32-bit) 75 TB (64-bit)	16 TB (32-bit) 90 TB (64-bit)	16 TB (32-bit) 90 TB (64-bit)	16 TB (32-bit) 90 TB (64-bit)	16 TB (32-bit) 105 TB (64-bit)	16 TB (32-bit) 105 TB (64-bit)	16 TB (32-bit) 105 TB (64-bit)

## Specifications

	N6210	N6210	N6240	N6240	N6240	N6270	N6270	N6270
Maximum number of volumes / LUNs	4096	4096	4096	4096	4096	4096	4096	4096
Maximum number of storage enclosures	20	20	42	42	42	48	68	68
Maximum number of Fibre Channel or iSCSI SAN connected servers	Up to 512 hosts per HA pair Up to 24 directly connected servers per HA pair							
Disk expansion units supported	EXN1000 - SATA Disk Storage Expansion							
	EXN3000 - SAS/SATA Disk Storage Expansion							
	EXN3500 - SAS Disk Storage Expansion (SFF)							
	EXN4000 - 4 Gbps Fibre Channel Disk Storage Expansion Unit							

## I/O Scalability

PCI-e expansion slots	2	4	4	6	12	4	6	12
Maximum number Fibre Channel ports	10	20	20	26	52	20	26	52
Maximum number of Ethernet ports	10	20	20	22	44	20	22	44
Maximum number of SAS ports	20	20	20	26	52	20	26	52
Maximum number of optional adapters	2	4	4	6	12	4	6	12



## For more information

To learn more about the IBM System Storage N6000 series systems, please contact your IBM representative or IBM Business Partner, or visit: [ibm.com/systems/storage/network](http://ibm.com/systems/storage/network)

For N6000 series modular disk storage system technical specifications and optional adapter cards, visit: [ibm.com/systems/storage/network/n6000/appliance](http://ibm.com/systems/storage/network/n6000/appliance)

For N6000 series interoperability and tape drive support, visit: [ibm.com/systems/storage/network/interophome.html](http://ibm.com/systems/storage/network/interophome.html)

Additionally, IBM Global Financing can tailor financing solutions to your specific IT needs. For more information about great rates, flexible payment plans and loans, and asset buyback and disposal, visit: [ibm.com/financing](http://ibm.com/financing)



© Copyright IBM Corporation 2012

IBM Systems and Technology Group  
Route 10  
Somers, NY 10589

Produced in the United States of America  
February 2012  
All Rights Reserved

IBM, the IBM logo, ibm.com, System Storage, Enterprise Storage Server, XIV, DS8000, and ALX are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product and service names may be trademarks or service marks of others.

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. IBM shall have no responsibility to update such information.



Please Recycle

## IBM System Storage N3000 Express

*Advanced system designed to enable outstanding,  
cost-effective deployment versatility*



---

### Highlights

- Takes advantage of proven features, including a high-performing and flexible operating system, data management software and redundancy features for continuous operations
  - Supports disk-based backup, with file or application-level recovery using SnapMirror, Snapshot and SnapRestore software features through a simplified replication, backup and recovery system for better data protection and retention
  - Enables on-the-fly provisioning with self-diagnosing systems to simplify management
  - Supports concurrent file and I/O block serving over Ethernet and Fibre Channel storage area network infrastructures using a single, integrated architecture for improved versatility
- 

The promise of expanding a data center with small, low-cost servers has led to an unintended consequence—"stranded storage" from internal disks or directly attached storage (DAS) solutions. IT professionals today are overwhelmed by the amount of data they have to manage. They are challenged by the need to keep pace with their companies' growing business, to improve backup and restore effectiveness, and to implement disaster recovery solutions without overwhelming IT staff—often on a shoestring budget. This leads us to develop solutions that improve storage efficiency and data protection considering assets that companies currently have.

The IBM® System Storage® N3000 Express systems are designed to provide primary and secondary storage for midsize enterprises, consolidating all of their fragmented application-based storage and unstructured data into one single-code system. Easily managed and expandable, this platform can help IT generalists increase their effectiveness. N3000 Express systems offer integrated data access, intelligent management software and data protection capabilities—such as those found in higher-end IBM System Storage N series systems—all in a cost-effective package. N3000 Express series innovations include internal controller support for serial-attached SCSI (SAS) or serial advanced technology attachment (SATA) drives, expandable I/O connectivity and onboard remote management.

The N3000 Express series is compatible with the entire family of N series storage systems, which feature a comprehensive line-up of hardware and software designed to address a variety of possible deployment environments.






The new N3220 series squeezes 7.2 TB of internal raw capacity into a 2U enclosure. Optional external expansion can increase total system raw capacity to 374 TB. The new N3240 series squeezes 36 TB of internal raw capacity into a 4U enclosure. Optional external expansion can increase total system raw capacity to 432 TB. The N3400 series can expand up to 24 TB of internal raw capacity and increase total raw capacity to 368 TB. Whether used for primary or secondary storage, the N3000 Express systems are intended to provide outstanding deployment versatility and connectivity to help satisfy your data protection and recovery needs at an affordable cost, improving storage efficiency.

### Easy to use

N3000 Express systems offer versatility through a single code file storage—Common Internet File System (CIFS) and Network File System (NFS)—that support block storage. The Internet Small Computer System Interface (iSCSI) and Fibre Channel protocol can be used as primary or secondary storage. These systems are designed to address storage consolidation challenges as well as application server virtualization projects in an environment of continuous operations. With NetApp Data ONTAP Snapshot technology, the N3000 Express systems offer the ability to use storage efficiently by helping increase utilization through thin provisioning (FlexVol and FlexClone) and reduce storage space requirements.

### More business uptime

The N3000 Express systems support dual-controller configuration with automated active-active failover for continuous operations. Using the N series SnapManager set of features, multipath high availability for business continuity, and intelligent data protection and disaster recovery software, the N3000 Express systems are intended to help keep your business running smoothly.



### Designed to help keep costs low

The N3000 Express systems are designed as the entry point to the entire N series family. The systems provide multiple I/O connectivity options, a small footprint to hold high-density SAS or SATA drives, and external expansion using low-cost SATA drives and Fibre Channel disks for production applications that leads to an improvement in storage efficiency. They also utilize Data ONTAP Snapshot technology. The systems are truly versatile products that can be deployed to address some of the most demanding application environments. For further systems administration time and cost advantages, the N3000 Express systems come standard with remote onboard management capabilities to help simplify remote system monitoring, cycle power, execute firmware upgrades, enter console commands and run diagnostics to help maintain the reliability of the system and your business-critical data.

### Highly flexible storage solution


The N3000 Express series is designed for a broad range of deployment scenarios. The flexibility of the N3000 Express allows it to address the storage needs of a wide range of organizations, including distributed enterprises and data centers for midrange enterprises. The N3000 Express also supports sites with compute- and data-intensive enterprise applications, such



as database, data warehousing, workgroup collaboration and messaging. The N3000 Express system functions as an “integration engine,” which is designed to enable you to simultaneously serve both file- and block-level data across a single or multiple networks—demanding procedures that for some solutions require multiple, separately managed systems. The N3000 Express supports Ethernet and Fibre Channel environments, enabling economical NAS, Fibre Channel and iSCSI deployments.

### Affordable data protection for distributed enterprises

N3000 Express storage systems can offer significant advantages for distributed enterprises with remote and branch office sites. These organizations and others can use the SnapVault and SnapMirror software functions to implement a cost-effective data protection strategy by mirroring data back to a corporate data center. N3000 Express systems can help improve data availability and simplify backup and restore operations by implementing centralized backup through a single methodology. This helps reduce tape management requirements and the need for remote systems administration. Recovering data backed up on N3000 Express systems can be faster than recovering from tape.



### Support for low TCO and long-term investment protection

N3000 Express systems support a low total cost of ownership (TCO) with an affordable price point, easy installation, configuration and ongoing maintenance. Standardization on the N series storage architecture can help your organization reduce complexity and take advantage of staff IT skills. The innovative design of the N3000 Express results in a small form-factor appliance that conserves scarce and valuable space in data centers or remote office locations. In addition, the ability to support a single-code storage networks enables you to make the most of your current network investment while deploying a long-term, highly expandable and easily upgradeable storage solution.

### Automation, consolidation and simplification

N3000 systems are empowered with OnCommand, a multiprotocol manager that delivers a single code experience to manage physical and virtual storage environments using integrated workflows and policy-driven automation. From a single interface, OnCommand enables the consolidation and simplification of shared IT storage management, delivering greater flexibility and efficiency.

Also, the capabilities of expandability, connectivity, data protection and retention, and copy recover availability, integrates with leading IT management frameworks. The software provides common management services, integration, security, and role-based access controls. OnCommand offers data and storage management tools that increase productivity Storage efficiency, and agility for organizations of all sizes. It supports and integrates with higher-level IT orchestration and management frameworks, which helps you manage data from a business perspective, enabling administrators to manage data across applications, databases, servers and storage.

## N3000 Express at a glance

Specifications						
	N3220	N3220	N3240	N3240	N3400	N3400
Machine type model	2857-A12	2857-A22	2857-A14	2857-A24	2859-A11	2859-A21
Controller configuration	Single	Dual (active/active)	Single	Dual (active/active)	Single	Dual (active/active)
Random access memory	6 GB	12 GB	6 GB	12 GB	4 GB	8 GB
Fibre Channel ports (speed)	2 (8 Gbps)*	4 (8 Gbps)*	2 (8 Gbps)*	4 (8 Gbps)*	2 (4 Gbps)	4 (4 Gbps)
Ethernet ports (speed)	4 (1 Gbps) onboard 2 (10 Gbps)*	8 (1 Gbps) onboard 4 (10 Gbps)*	4 (1 Gbps) onboard 2 (10 Gbps)*	8 (1 Gbps) onboard 4 (10 Gbps)*	4 (1 Gbps)	8 (1 Gbps)
Maximum raw capacity	374 TB		432 TB		368 TB	
Maximum number of disk drives	144 24 internal		144 24 internal		136 12 internal	
Disk drives supported in controller (type, size, speed)	SAS: 450 GB, 600 GB; (10k RPM)		SATA: 1 TB, 2 TB, 3 TB; (7.2k RPM)		SAS: 300 GB, 450 GB, 600 GB, 600 GB FDE (15k RPM) SATA: 1 TB, 2 TB, (7.2k RPM)	
Disk expansion units supported	EXN3000 - SAS Disk Storage Expansion Unit: SAS: 300 GB, 450 GB, 600 GB (15k RPM) SATA: 1 TB, 2 TB, 3 TB (7.2k RPM),  EXN3500 - SAS Disk Storage (SFF) Expansion Unit: SAS: 450 GB, 600 GB (10k RPM)				EXN3000 - SAS Disk Storage Expansion Unit: SAS: 300 GB, 450 GB, 600 GB, 600 GB FDE (15k RPM) SATA: 1 TB, 2 TB, 3 TB (7.2k RPM)  EXN3500 - SAS Disk Storage (SFF) Expansion Unit: SAS: 450 GB, 600 GB (10k RPM)  EXN1000 - SATA Disk Storage Expansion Unit: SATA: 1 TB, 2 TB (7.2k RPM)  EXN4000 - Fibre Channel Disk Storage Expansion Unit: 4-Gbps Fibre Channel: 300 GB, 450 GB, 600 GB (15k RPM)	
Warranty	3 years for hardware and licensed software, CRU and on-site service, next business day 9x5, service upgrades available					

Software		
Operating system	Data ONTAP	
Operating systems supported	Microsoft Windows 2000, Windows Server 2003, Windows XP, Linux, Oracle Solaris, IBM AIX®, HP-UX, Apple Macintosh OS, VMware ESX	
Software features	<b>Standard</b> Integrated RAID manager, including RAID-DP Snapshot Fast Boot NIS DNS FlexVol FlexShare Network Data Management Protocol (NDMP) Data Compression** Storage Encryption	<b>Licensed</b> CIFS NFS HTTP iSCSI FCP FTP FlexCache FlexClone MultiStore Clustered Failover SnapLock SnapMirror SyncMirror SnapRestore Single Mailbox Recovery SnapVault SnapMover NearStore Deduplication SnapValidator SnapManager for Microsoft Exchange SnapManager for Microsoft SQL Server SnapManager for Microsoft Office SharePoint (SMSP) SnapManager for Oracle SnapManager for SAP SnapManager for Virtual Infrastructure SnapManager for Hyper-V SnapDrive OnCommand



## For more information

To learn more about the IBM System Storage N3000 Express systems, please contact your IBM representative or IBM Business Partner, or visit: [ibm.com/systems/storage/network](http://ibm.com/systems/storage/network)

For N3000 Express series modular disk storage system technical specifications and optional adapters available, please visit: [ibm.com/systems/storage/network/n3000/appliance](http://ibm.com/systems/storage/network/n3000/appliance)

For N3000 Express series interoperability, visit: [ibm.com/systems/storage/network/interphome.html](http://ibm.com/systems/storage/network/interphome.html)

For an overview of the N3000 Express series software features, functions and benefits, please visit: [ibm.com/systems/storage/network/n3000/appliance/features.html](http://ibm.com/systems/storage/network/n3000/appliance/features.html)



© Copyright IBM Corporation 2012

Systems and Technology Group  
Route 100  
Somers, NY 10589

Produced in the United States of America  
February 2012

IBM, the IBM logo, ibm.com, System Storage, and AIX are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. IBM shall have no responsibility to update such information.

IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

\* Optional ports available via add-on mezzanine cards

\*\* Available only in Data ONTAP 8.1 or later and internally in N3400.



Please Recycle