ROANOKE COUNTY

Purchasing Division
5204 Bernard Drive, Suite 300-F, P.O. Box 29800
Roanoke, Virginia 24018-0798
TEL: (540) 772-2061 FAX: (540) 772-2074

April 29, 2022
ADDENDUM NO. 1 TO ALL BIDDERS/OFFERRORS:

Reference – RFP #2022-082

Description: Answers to Vendor Submitted Questions

Issue Date: April 29, 2022 Proposal Due: May 10, 2022

The above Project is hereby changed as addressed below:

Section 2 – Background:

1. **What are the current deduplication/compression ratios for each identified workload?**
   a. **SQL Virtual Servers**
   b. **Non-SQL Virtual Servers**
   c. **Physical Servers**

Response: The current solution does not provide that information per that break down. The current solution is called a Core server. On the Core server there may several repositories for the backup data. These Core servers will contain SQL, non SQL and physical server data. It is not separated out in the manner listed. A compression percentage is shown for each repository. A list of the core servers and compression percentage is shown below.

PSFCCORE3	Repository 1	27% compression	
PSFCCORE4	Repository 1	47% compression	
	Repository 2	8% compression	
	Repository 3	6% compression	
RVRVCORE	Repository 1	12% compression	
SOCOCORE	Repository 1	17% compression	
VTDRCORE3	Repository 1	33% compression	These last two Core servers are
	Repository 2	29% compression	replication targets only
VTDRCORE4	Repository 1	18% compression	

The current solution in place does not provide any deduplication.

2. **In the Network Connections sub-section, item 2, a "third site" … "no data is on site to backup" is described.**
   a. **Is the 45 mile geographic diversity considered sufficient for the County's future BCDR purposes, and if so, could it serve as a location for replication targets for the proposed solution?**

Response: The purpose of this location is to be a replication target only. It would never be considered for a true DR recovery site. It's desired to be a third location for our data with some geographic diversity.

Section 6 – Services and/or Items Required:

1. **Desired System has three options. Do you want all three options presented or are you preferential to one? Option 2 and 3 will require in depth conversations with the hosting service provider for quoting**

   Response: Currently the backup solution is all on premise. Our biggest desire in changing systems is to get an immutable backup for the recovery from a ransomware attack. We understand different vendors may approach this with a different solution. That is why we are open to on premise, hybrid or all cloud if that solution is cost effective and provides us this immutability. If on premise is the most cost effective way to present your solution, then that is fine.

2. **Operational documentation is in the help files, is this a request for a high level overview of steps or is there a need for table top exercise to identify required steps for business continuity with respect to how systems need to come back in order for operation? That could be a bit more involved.**

   Response: This was meant as a high level overview of how the system works and how to recover data from backups. It was not intended to include documentation for our own business continuity or recovery from a disaster. In other words, we are just looking for good system documentation to be contained within the help files in the system.

3. **Required Features and Functionality item 2.18: "Provide continued functionality in the event of a power loss from outside the network." Clarification, please?**

   Response: The thought process here was if the solution had a centralized management of the system on premise. If we lost power at the main site and lost access to the centralized management, then we wanted to have some sort of access via the remote or replication site to get to the data. In other words, we don't have all the eggs in one basket at the main site. The solution could be managed from any of the backup locations. Wording could probably have been better

4. **Required Features and Functionality item 2.23: "Geodiversity of replication hardware" Clarification, please?**

   Repsonse: We are referring to the current replication site located about 45 miles from our main site. We are looking for something to be located outside of the Roanoke Valley. Although our alternate data center is over 10 miles away, we are accustomed to this third site that's more than a half hour away from any of our sites. That is assuming an on premise solution. A hybrid or cloud solution would by default have geodiversity.

5. **If the County provides these ports via existing LAN switches, then Professional Services item 4.5 "Set up all required infrastructure for successful backup and full recovery processes." will involve the Offeror configuring production equipment ostensibly during production hours. Will this be acceptable to the County?**

   Response: This could have been more clearly stated. We are looking for the professional services to set up the solution on premise or the on premise portion of a hybrid solution. Roanoke County will provide the underlying infrastructure such as LAN ports and configure them per system requirements. We do not expect or want vendors configuring anything but the proposed backup solution.

6. **Professional Services item 4-10 "Perform failover testing." Please clarify the nature and scope of this requirement.**

Response: The intention here was to have the vendor demonstrate recovery of data. For example, if a physical domain controller fails, we would want to see how to perform a bare metal recovery to new hardware. Keep in mind our reference point is today's solution. So in today's environment, if we lost a Core backup server at the primary site, I would want to have the backups "failover" to the alternate site where the backup solution could back up vm's and servers still running in the primary site until that core backup server could be replaced.

7. **Service and Support item 5.1 "The vendor must provide 24x7 on-site support with a 4 hour maximum response time" - Please characterize 'response time'?**
    a. **Example: some support arrangements consider acknowledging a request for support via phone or email as responding within four hours**
    b. **Example: some support arrangements consider having a technician on site within four hours as responding within four hours**
    c. **Example: some support arrangements assign response urgency by the nature of the reported issue. A failed drive within a redundant system may receive a next-day response while an inoperable system that exceeds configured redundancies receives immediate dispatching.**

    Response: This requirement was made with hardware failures in mind. Failed components where the system is down, a four hour response time is desired to minimize the outage. Similarly with software support in the event system is inoperable. Would expect the vendor to provide some level of support within a four hour time frame. Understand resolution may not occur always within four hours but at least it's being worked on. Example A – If the issue is below a Priority 2 (business function somewhat impacted), this sort of response could be adequate. Example B – If the issue is a Priority 1 (business function down or severely impacted), we would expect a four hour response time on parts or software support. It would be critical to get the system back up and running normally. Example C – The example given is what would be expected. If the system is still functional and the part or service can be provided next business day, that is fine.

8. **Final Testing Item 8.1 "The proposed Backup hardware and software will be tested before any commitment to purchase" – Please clarify?**

    Response: Ideally, we would be able to run a proof of concept in a limited capacity before making the final purchase decision. Akin to test driving a car. However, we understand in today's environment with supply chain issues, hardware may not be available to perform this type of testing. As mentioned, a site visit with an existing customer to review and discuss the solution may be an alternative. In a cloud based solution we would expect to be able to test within a specified period of time (i.e. 14 to 30 days).

9. **In Section 2 for the Primary Data Center, we get several different numbers:**
    a. **6. Backup Data size: approximately 65 TB**
    b. **6.1 SQL Data size = 2.2 TB (only SQL backups are backed up to disk, not databases)**
    c. **6.2 Non-SQL size = 46.5 TB**
    d. **6.3 Physical servers size = 575 GB**

    Are there any other workloads protected at this site? Do the data sizes noted here for SQL, non-sql, and physical servers stand for the total FETBs? Is their total size, which is ~50 TB considered the FETB or the 65 TB?

    Response: According to the current solution, the front end amount of data being backed up is 63+ TB. We rounded up to 65. The SQL Data was based on the amount of space occupied on the SQL servers where the SQL backups are stored. The physical server size was determined from the physical servers drive space currently in use.

10. **Secondary or Alternate Data Center:**
    a. **Backup Data size approximately 35 TB**
    b. **6.1 SQL Data Size = 1.1 TB (only SQL backups are backed up to disk, not databases)**
    c. **6.2 Non-SQL size = 19 TB (server data, file system, application data)**
    d. **6.3 No Physical servers**

Are only ~20 FETBs being protected at this site? (Non-sql and sql size)

Response: The secondary or Alternate site is primarily for test application, test SQL, and a few Library related production servers. The 35 TB total FE is according to our current backup solution. The SQL size was determined from the SQL server drive size used where the SQL backups are placed.

Requests for information not contained in the original RFP

1. **Workload sizing details – for each workload type (see list below for all types) that will be included in the backup sets:**
    a. **SQL**
    b. **Non-SQL (i.e., unstructured data)**
    c. **Virtual servers**
    d. **Physical servers**

Response: What information desired is unclear but will attempt to answer. Current front end data is around 65 TB of space according to the current backup solution. The overall storage in use including full and incremental backups is 271 TB.
SQL – According to our DBA, 8.26 TB would be a full backup of all SQL each night. The backups are compressed by 40% so about 5.87 TB. About 25% of this is in the test environment.
Non-SQL – If SQL is about 8 TB of our front end workload, then the remaining amount would be approximately 58 TB of data.
Virtual Servers - 63 virtual servers in the alternate site. 58 of them run a Windows operating system. 4 are linux and 1 is vmware photon os. Typical Windows server has a C: drive of 100 to 125 GB C: drive for the operating system. Some will have a D: drive between 50 and 100 GB. All of them have a W: drive for the swap file and it's the same size as the memory (i.e. 8 Gb, 16 Gb, etc).
161 virtual machines in the primary data center. 3 are Windows 10 machines, 3 are linux machine and the remainder are Windows Server from 2012 to 2019 versions.
Physical Servers – Six (6) physical servers. Right at 2 TB total space being used between all six servers.

2. **Please provide the following for each workload type:**
    a. **# of Files per TB (in millions)**
    b. **Data Size (TB)**
    c. **Annual Growth Rate (%)**
    d. **Daily Change Rate (%)**
    e. **Backup Frequency**
    f. **Backup Window (Start – Stop – Duration)**
    g. **Replication Requirements**
    h. **Retention/Archiving Requirements if different than previously given**
    i. **Indexing Requirements (does this workload require file-level granularity of restores)**
    j. **Is the workload type considered to be non-deduplicable/non-compressable?**

Response:

# of Files per TB (in millions) - – Unknown. Current system does not provide that sort of statistical information and is somewhat hard to determine.
Data Size (TB) – As previously stated, current system indicated approximately 65 TB front end data. Current backups which include a full and number of incrementals and retention is 272 TB of space.
Annual Growth Rate (%) – Unknown. Current system does not provide that sort of statistical information and is somewhat hard to determine.
Daily Change Rate (%) – Unknown. Current system does not provide that sort of statistical information and is somewhat hard to determine. Best guess is minimal change rate per day. Probably less than 100 GB.
Backup Frequency – Current Retention policy is:
Keep all recovery points for 3 days
And then keep one recovery point per hour for 2 days
And then keep on recovery point per day for 7 days
And then keep one recovery point per week for 4 weeks
And then keep one recovery point per month for 12 months

We would like to have a retention policy to keep
1 day of hourly backups
7 days of daily backups
5 weekly backups
13 monthly backups
This would give us slightly over a year worth of data to recover from. If we could define what days to roll up, that would be a bonus. Current system does not allow that.

Backup Window (Start – Stop – Duration) – Currently we are backing up SQL files between 6:00 AM and 11:00 PM. Typical file server backups start between 6:00 and 7:00 AM and go to 7:00 PM or some will go to 11:00 PM. Typically the backup runs every two hours within those time frames. Test server are only backed up once a day. The time vary but most occur after working hours. The overnight hours were reserved for the system to perform roll ups and deletion of old backup files outside the retention period. We are open to recommendations on backup schedules.

Replication Requirements – Currently replication runs between 6:00 AM and Midnight. However, when replication falls behind, it will run 24 hours a day to try and catch up. One pain point we have with current replication is the internet connection to our off site location 45 miles away. Once replication falls behind, it may take days to replicate one server. Some of this will be alleviated with an upgrade to the internet connection within the next 60 to 90 days. The internet pipe will triple in size.

Retention/Archiving Requirements – Prefer to retain backup information for a least one year. Retention requirements are defined by the Library of Virginia and it's up to each department to notify IT if a file needs to be kept for a longer period. As of now, no department has notified IT that a Word, Excel, Powerpoint, PDF or other document is required to be kept longer than the year time frame.

Indexing Requirements – Our interpretation of this question may not be what is intended. Obviously, we want the ability to restore a single file such as a Word or Excel document contained within a folder. If the question refers to a file within the operating system or application, probably not as much.

Is the workload type considered to be non-deduplicable/no-compressible? – We would prefer the backup data be deduped and compressed as much as possible to save on space. We understand SQL backups are not compressible and there may be other files that also can not be compressed or deduped due to other circumstances.

3. **Please provide this for all source data at all sites, broken out by workload data types:**
    a. **VM's**
    b. **Oracle**
    c. **SQL**
    d. **SAP Hana**
    e. **Exchange On-Prem (note: customer is using GroupWise on premises)**
    f. **Microsoft 365 Exchange Online**
    g. **Microsoft 365 Sharepoint Online**
    h. **Microsoft 365 OneDrive Online**
    i. **Other Databases**
    j. **SharePoint On-Prem**
    k. **Unstructured Data**
    l. **Non-Dedupable Data**
    m. **Archive/Transaction Logs**

    Response:
    VM's – Alternate data center has 63 VM's. Four (4) of those are Linux servers and the remainder are Windows. Current backups utilize 36 TB of space.
    Primary data center has 161 VM's. Three of those are Windows 10 machines, 3 are Linux machines and the remainder are Windows Server versions 2012 through 2019. Current backups utilize 222 TB of space.
    Oracle – None
    SQL –14 VM Servers in Prodcution – 20 SQL instances
    10 VM Servers in Test – 13 SQL instances
    SAP Hana – None
    Exchange on Prem – None
    Microsoft 365 Exchange Online – None at this time but will have in the future
    Microsoft 365 Sharepoint Online – None at this time but will have in the future
    Microsoft 365 OneDrive Online – None at this time but will have in the future
    Other databases – Unknown. Some applications utilize SQL Express but those are not supported direction by our DBA
    SharePoint On Prem – included in VM information above
    Unstructured Data – Primarily this would be our CIF servers. We have two CIFS servers. One contains User (personal drive) and Departmental Shared drive data. The personal drive data is 3.8 TB. The shared departmental data is 6.9 TB. Most of this data is static and grows at a slow rate. The second CIFS server contains application data such as Laserfiche files and Pictometry data for GIS. There is 1.5 TB of Fiche data; 2.4 TB of General application data, and 3.5 TB of Pictometry data. Other than Laserfiche, most of this data is static and changes very little.
    Non-Dedupable Data – Other than SQL backups, unknown.
    Archive/Transaction Logs – N/A

4. **What is the function of each physical server?**

    Response – Three (3) of the servers are domain controllers. Two (2) of the servers run the software for the scale operations at the local trash transfer station. And the one (1) server is a file server used by the local access television station to store older video files.

5. **Are you looking for any cloud archive component in the solution?**

    Response: As previously stated, we currently all on premise. We are open to staying all on premise, hybrid (on premise and some cloud) or all cloud. This is a solution we expect to be in place for a number of years so we can not ignore the cloud as an option.

6. **Are you interested in defining a "low RPO" critical workload segment, rather than the entire workload?**

Response: Unfortunately, we don't have the knowledge that we need concerning our data to completely identify the most critical components. We think we know but in actuality we probably don't at this point. That process is under way now. An overall RPO of 12 to 24 hours to restore data is considered reasonable. Our most critical data would probably be in the less than 4 hour range but not all of that data has been identified.

7. **Is an electronic only submission suitable to lower environmental impact?**

Response: As outline in the RFP, we do not accept electronic proposal submissions.

Note: A signed acknowledgment of this addendum must be received at the location indicated on the original solicitation either prior to the proposal due date or attached to your proposal. Signature on this addendum does not substitute for your signature on the original proposal/bid document. The original proposal/bid document must be signed.

Thank you,

Neil Huss
Phone: (540) 283-8151
nhuss@roanokecountyva.gov

Sign Name:	Print Name:
Name of Firm:	Date: