CITY OF O'FALLON, MISSOURI
PURCHASING DEPARTMENT

100 North Main Street
O'Fallon, MO 63366
T: (636) 379-5527

RFP NUMBER:	**26-009**	**RFP ISSUE DATE:**	**01/16/2026**
RFP DESCRIPTION:	**CLOUD-BASED BACKUP AND DISASTER RECOVERY SOLUTION**		
RFP OPENING DATE:	**02/19/2026**	**RFP OPENING TIME:**	**12:00 P.M. CST**

February 6, 2026

ADDENDUM #1

The purpose of this addendum is to address any questions received on or before January 30, 2026, at 12:00PM CST. The time period for questions is closed.

Questions:

1. We would appreciate clarification on the requirement stating 'no conversion formats whatsoever during fallback'. Could the City elaborate on the underlying objective of this constraint so we can align our proposed architecture accordingly? Specifically, we would like to understand which workflows or constraints make conversion unacceptable?

   No VM format conversion is permitted during migration, fallback, or recovery. Virtual machines must remain in their native format end-to-end across on-prem, cloud storage, and DR environments. The target architecture must natively support VMware VMDK on both on-prem and cloud file systems, with no format translation (e.g., VMDK to VHDX), and be fully managed.

2. Have you already paid or signed a contract for three or more years of the VMWare licensing?

   Yes, we are fully covered by VMware licensing for on-prem use.

A. Governance, Responsibility, & Vendor Managed Scope

1. Who is responsible for declaring a disaster, and is this process manual or automated?
   Process is manual and it is approved by upper management and DR board.

2. Who is responsible for approving failover outside of pre-scheduled testing (e.g. in the case of actual disaster recovery)?
   IT department with the approval of City Administrator and upper management.

3. Is City planning to involve staff in validation of failover (e.g. application-specific actions) after new SDDC capacity has been established?
   Yes. In addition, the validation failover needs to be scheduled and executed at least twice per calendar year.

4. Is it expected that the vendor that provides the cloud platform will operate the recovery process and conduct failover, validation, and failback activities or is the use of a separate managed service provider acceptable?
   The City requires the cloud platform vendor to directly provide recovery, failover, validation, and failback services. While the vendor may use a managed service provider, the City will work only with the awarded vendor for all support, billing, and services. All support requests must be handled directly by the vendor, and the City will maintain a single contract with no separate service providers or intermediaries.

B. Current Environment

1. What hardware manufacturer and model(s) and specifications are currently used on-premises for compute and storage?
   Due to security information, we will not reveal this information. We will provide awarding vendor with all necessary details including RV Tools exports.
   (ref. RFP p.8, vSphere environment requirements)

2. What version of vSphere is currently deployed (e.g. 7.0, 8.0, 9.0)?
   8.0

3. What is the current vSphere licensing edition (e.g., VVF, VCF) and total capacity (core/VSAN) licensed?
   VCF

4. Can the City provide an inventory of the 45 VMs to be protected, including per-VM details: vCPU count, memory allocation (GB), provisioned storage (GB), guest OS, and current network configuration?
   Awarding vendor will be provided with all inventory details along with full RV Tools information and environment insight.

5. What is the total current storage consumption (in TB) across all 45 VMs, excluding thin provisioning overhead? (*ref. RFP p.8 specifies 50TB total protected VM storage)*
   Approximately 35 TB with 10% yearly growth.

6. Does the City currently use a backup or replication solution (e.g., Veeam, Zerto)?
   Yes (Veeam) If yes:
   - What is the average daily change rate (GB/day or %) for the VMs in scope? Most VMs have minimal daily changes and stable storage footprints. A small number of VMs (fewer than four), such as file servers hosting video recordings, experience a higher daily change rate of approximately 1–2%.
   - Will this solution remain in place alongside the proposed cloud DR solution? (Yes, current backup solution will also remain in place)

7. What is the City's expected growth rate for VMs and storage over the 3-year contract term? 10% yearly storage growth

8. What is the current network bandwidth available for data replication to the cloud from on-premises (upload/egress)? Up to 1GBps could be available and allocated for upload/egress What network equipment is used to delivery connectivity to the Internet?RF We will not reveal this information due to security concerns, however, we could tell that our firewall and network equipment is capable of delivering 10GB connectivity.

C. Disaster Recovery Scope & Expectations

1. For failover testing where a minimum of seven test VMs are used, should bidders assume these seven test VMs represent the City's largest/most critical workloads, or an average workload profile? Six VMs will be of 60-250 GB workload profile. One of those VMs will be over 1TB (for testing failover time and capabilities).

2. During failover testing, is planned service downtime acceptable (even if it exceeds the 4-hour RTO), or must testing be performed in a non-disruptive manner? Testing must be performed in a non-disruptive manner. We cannot accept any downtime as we run critical services that require high availability.

3. Does the RTO apply to all 45 VMs or a subset of VMs? All 45 VMs

4. Is the RTO defined to be the time until VM power-on or other criteria? Yes

5. Are there preferred maintenance windows or blackout dates for the 60 hours of DR testing and 5 failover test days? We would prefer to run them over weekends or evenings, but we are also open for DR testing during weekdays (if this wouldn't produce any downtime).

6. Does the City expect to participate in validation testing (application-level verification) after SDDC capacity is launched during tests? Managed service provider will perform all validation tests, as well as all relevant documentation about SDDC capacity during launched tests.

7. What is the expected timeline for vendor response during a declared disaster (initial response and full SDDC activation)? Less than 30 minutes.

8. What is the expected scope and format for the Disaster Recovery Plan document? The Disaster Recovery Plan must contain all information required for City staff to perform a complete recovery without any external assistance, including clear, step-by-step procedures for contacts, cloud access, SDDC setup, VM recovery, network configuration, and failbacks. The City will provide the selected vendor with current environment details and requirements. Should it include detailed technical runbooks for City staff to independently execute failover, **(Yes)** or focus on vendor-managed procedures and City approval workflows? **(Yes to Both)**. The managed service provider will perform failover after a disaster is declared; however, the City must also be able to initiate failover

independently if the provider cannot be reached in a timely manner. (For example, if communication is disrupted (such as phone outages), the City's support team must be able to access the DR management console and execute failover independently.)

9. Please confirm whether for the SDDC that some form of pre-provisioned or warm standby compute capacity is expected to be available or whether all recovery compute resources are expected to be provisioned only at the time of testing or disaster recovery declaration?
   The City does not require a pre-provisioned SDDC and expects recovery and compute resources to be provisioned only during testing, or a declared disaster. If the vendor includes pre-provisioned resources in its architecture, they must be included in the managed cloud solution at no additional cost. The SDDC must be available in a timely manner during an emergency to meet required RTO, RPO, and MDT.

10. For failover testing days, is it expected the SDDC will be created or destroyed per test or that the SDDC will remain persistent and temporarily activated for testing? The City is comfortable with the SDDC being created and decommissioned for each test.

D. Geographic Redundancy

1. The RFP requires a "minimum of two cloud hypervisor hosts in a redundant cloud geographic regions" (*ref. RFP p.8, General Requirements*). Does this mean that the City requires the ability to recover workloads into multiple independent cloud SDDCs in different regions (e.g. if a cloud datacenter in the US Midwest was down, the City also has the option of recovering to a datacenter on the US East Coast instead), or something else? **Yes.**
   If the cloud provider operates multiple data centers, the City requires workloads to run redundantly across regions. If one region (for example, US Midwest) becomes unavailable, workloads must be able to recover and run in another region (such as US East). During recovery, a minimum of two hypervisors must be available to support all workloads until on-premises systems are restored and failback is completed. The managed cloud provider must provide full redundancy, including redundant cloud regions, while all workloads are running in the cloud.
   - If backups are geographically isolated, is a single-region recovery SDDC acceptable? A single SDDC is acceptable, provided all workloads can be recovered and run from one cloud region and the recovery design eliminates the risk of SDDC backup failure due to a regional outage. The vendor must account for this requirement when designing the recovery architecture.

E. Ransomware Recovery

1. For ransomware recovery capabilities (*ref. RFP p.9, "40 VMs"*), what specific detection or analysis reporting does the City require to identify when ransomware payload was first detected? The City would allow and accept offers of various acceptable solutions, but our preference would lean toward real-time detection (for example, detecting unusually high network traffic, extensive hard disk writes etc.) The City would also like exact time when ransomware payload was planted

and executed on the system. (For example, Ransomware payload infected the system in January but didn't execute until June)

2. Should ransomware detection analysis be automated or on-demand? Automated

3. What is the expected response time for vendor-provided analysis after a suspected ransomware incident? Less than 15 minutes (The City would prefer real time analysis)

4. Are there specific immutability requirements for backup snapshots? Minimum of 10 days of immutable lock on the backup snapshots.

F. Networking & Security

1. What is the City's current on-premises firewall solution (manufacturer/model)? The City will not answer question related to equipment used (including exact firewall models) due to security concerns, however we can disclose that our firewall is a modern next generation application-level appliance capable of providing stateful packet inspection.

2. The RFP references WatchGuard FireBox deployment assistance (*ref. RFP p.9*). Does the City intend to:
   - Deploy WatchGuard as the virtual firewall in the cloud SDDC (The City is open to all options), OR
   - Continue using WatchGuard on-premises (Yes, preferred) and require VPN connectivity to the cloud SDDC? (Yes, in the case that on-prem VPN connectivity is no longer available due to disaster)

3. What is the current IP addressing scheme for the 45 VMs (subnet ranges, VLAN configuration)? IP addresses are Class 3 private subnets, allocated across several VLANs) (Multiple Class C subnet - 254 IPs per subnet)

4. For the 16 external public IPs specified (*ref. RFP p.8*), can the City provide the list of services requiring public exposure and expected concurrent usage? Yes, City will provide list of services requiring public exposure and concurrent usage to awarding vendor. Do the 16 external public IPs require static reservation or can they be dynamically allocated on demand during DR events? Yes, static reservation is required (as certain services rely on dedicated Static IP mappings).

5. Does the City require site-to-site VPN connectivity between on-premises and the cloud SDDC during failover scenarios? No, The City only requires that remote clients have ability to connect to available VPN resources (either to SDDC during disaster or when tests are executed, or to on-prem VPN resources (when on-prem firewall and domain controllers are operating properly).

6. Are there specific firewall rules or security policies that must be replicated in the cloud environment? Yes, the awarding vendor will be provided with all relevant information, including specific firewall rules or security policies that must be replicated in the cloud environment.

7. What DNS servers are currently used on-premises, and should these be replicated to the cloud SDDC? The City utilizes Microsoft Active Directory DNS for on-prem DNS servers. Requirement is that those need to be replicated to cloud SDDC environment, (however, as these DNS servers are part of protected VM sets, replication would happen automatically by scheduled snapshots). Exception to this would be if awarding vendor plans to offer additional built-in DNS redundancy with servers capable of Active Directory DNS zone replication) The City also uses extremal DNS services (for public facing websites and services), however those are administered by external managed providers and not required to be replicated.

G. Active Directory & Authentication

1. What is the current Active Directory forest/domain structure? Single Forest, single Domain, multiple Domain Controllers. DC site-to-site replication.

2. How many domain controllers are included in the 45 VMs to be protected? 4 Domain Controllers

3. For SAML/MFA requirements (*ref. RFP p.9*), what identity provider does the City currently use (e.g., Azure AD, Okta, on-premises AD FS)? Azure AD, third-party provider (Duo)

4. Does the City expect domain controllers to be operational in the cloud SDDC during failover (Yes), or will authentication rely on connectivity to on-premises DCs? (Depending on scenario. They will authenticate to on-prem DC (if those are available), however, if on-prem connectivity is lost, authentication will rely solely on Domain Controllers that are already replicated & running in the cloud SDDC environment.

H. Built-in Platform Services

1. The RFP requests built-in DNS, DHCP, and NTP services to support the case where internal DNS servers become unavailable.
   - Are these services (DNS, DHCP, and NTP) currently running on VMware VMs on-premises, on bare metal servers managed by the city, or somewhere else? (Yes, DNS and DHCP are running on the VMs, while NTP runs as a stand-alone server appliance). Solution needs to be capable of providing a minimum basic redundancy to DNS, DHCP and NTP services (should these services become unavailable after SDDC environment is deployed).
   - Is it expected the virtual services within the SDDC will supply the functions (Yes) or will cloud-native managed services external to the SDDC be required? Not necessarily, however managed service provider could also offer it as an optional add-on, included into the final solution.

I. Compliance & Identity

1. For CJIS compliance, is this required from the cloud provider only, or from all entities involved in operating and managing the DR environment? (All entities involved in operating and managing the DR environment)

2. For SAML/MFA access, should bidders assume integration with the City's existing identity provider (Yes, and this is mandatory for VM protection), or instead with a vendor-provided platform? (Vendor provided platform will also be mandatory for all Cloud Management Console access (but it will only be used as redundancy should existing SAML/MFA identity provider become unavailable. (For example, if existing identity provider is down when trying to access cloud management console, offered solution would switch to vendor-provided identity provider).

3. Can the City provide the specific CJIS compliance requirements or documentation that the cloud solution must meet? (*ref. RFP p.8, "Cloud needs to be CJIS compliant")* *These compliance requirements are standard (Key 13 Security Policy Areas) and could be retrieved from CJIS website, as well as NIST 800-53 regulation document. The City is primary concerned with key elements (regarding solution being FIPS 140-3 compliant for encryption of data in transit and at rest), multi-factor authentication control (MFA), and strict access control for personnel operating cloud site.*

4. Does the City's cyber insurance policy specify required content, format, or attestations for disaster recovery test reports? *(ref. RFP p.2, mentions "Cyber Insurance Compliance Policy")* *No specific requirements*

5. What is the City's data retention policy for deleted or modified data beyond the snapshot retention schedule specified in the RFP *(ref. RFP p.9)*? Both deleted and modified data should be consistent with snapshot retention schedule. In addition, deleted and modified data should be retained for a minimum of 12 weeks.

J. Billing & Commercial Terms

1. The RFP specifies "Annual billing over the course of the agreement" (*ref. RFP p.8*). Does the City prefer:
   - Three separate annual invoices (one per year), OR
   - A single invoice for the full 3-year term payable annually? (We would want to sign a three-year agreement but pay annually)

2. For the on-demand SDDC environment *(ref. RFP p.8*), please clarify which costs are:
   - **Fixed/Flat**: Included in the base annual subscription price (SDDC offered by managed provider as part of solution, which include cloud storage used for snapshot retention & backups, operational support or recovery services, cloud management services, licenses or services that are already bunded and not passed to the customer)
   - **Variable/Usage-based**: Billed based on actual consumption

1. Some cloud providers require purchase of pilot-lights, with a commitment of prior purchasing cloud resources regardless of their utilization and consumption (e.g., Hypervisor Hosts & Live Cloud Storage File System which are being operational and running only when SDDC is created) The City would like to be billed for actual cloud consumption. (e.g., object storage utilization for stored backup, vs maintaining persistent cloud infrastructure for hypervisors that aren't even running or created)

2. Another example:
   If offered SDDC pilot-lights bundled package services are exceeded beyond agreed SLA (and those include running full operational workloads on cloud SDDC during recovery), the City would be responsible for actual cloud over-consumption. (For example, 30 days free of charge are exceeded and the City still continues consuming public cloud hosts & storage beyond agreed failback timeline)

3. What is the billing mechanism if the City exceeds 45 VMs or 50TB during the contract term? This will be negotiated & agreed during contract signing (awarding) phase.

4. For the 5 failover test days (*ref. RFP p.9*), does 'SDDC operation cost will be included in the price' mean that SDDC consumption during these specific 5 test days is included in the base subscription price? Yes
   - Separately, are the 60 hours of DR testing also included in the base price, or billed as on-demand SDDC usage? Included in the base price.

5. Are data egress and ingress charges expected to be capped or uncapped? They should be uncapped and truly unlimited. City doesn't expect to pay additional charges for data egress/ingress during daily snapshot backups, failback operations, or when running on public SDDC cloud (during a disaster).

6. Is there a not-to-exceed hour cap expected for professional services (deployment, setup, training), or should vendors propose a fixed project price? Vendors should offer fixed project price (everything should be included in the final price).

1. On-Premises Environment & Workload Scope

Can you confirm the current number of vSphere clusters, hosts, and vCenters supporting the 45 protected VMs? 2 Clusters, 2 hosts per cluster (Total of 4 hosts), single vCenter

Are all 45 VMs considered Tier-1 / mission-critical, or should workloads be prioritized by criticality? Workloads should be prioritized by criticality. (e.g., Domain Controllers, DNS, DHCP, Database Servers - approx. 7-10 VMs), Tier-2 all other VMs

Do any workloads have special boot order, dependency, or application-consistency requirements during recovery? Domain Controllers, DNS and DHCP should boot first, followed by Database Servers, and rest of the VMs.

Are there any non-VMware workloads that must be considered in the DR plan? Only VMware workloads are part of this DR. Protection for non-VMware workloads (e.g. Tier-2 Physical servers) would be treated as a benefit, but it is not requirement for the scope of this project.

What is the average and maximum VM size, and are there specific high-IO or latency-sensitive workloads? 50-250 GB, however, there are several VMs going well over 1TB. No specific high-IP or latency-sensitive workloads.

2. Recovery Objectives & Expectations

Is the stated 4-hour RTO intended to apply to all workloads collectively or only to critical systems first? Collectively

Do you have defined RPO targets, or should the solution propose best-practice RPOs? Solution should propose best-practice RPOs, however, RPOs cannot go over certain thresholds for certain workloads (for example, one day maximum would be considered acceptable data loss for certain Servers that are hosting files & data capable of being recovered by other means (e.g. nightly snapshots stored on third-party cloud objects such as Wasabi or Amazon S3).

Are there specific applications or services that must be available first during a disaster declaration? (Domain Services, Authentication, VPN capabilities)

How will a disaster event be formally declared, and who has authority to initiate failover? IT Department management in coordination with managed cloud provider.

3. Disaster Recovery Testing & Validation

Do you have preferred timing windows for DR test days and failover tests? No preferred timing windows, both could be run either during week and/or over the weekend.

Should DR testing be IT-only or include full business-impact simulations? Full business-impact simulations is required

Are there specific compliance or reporting formats required for DR test reports? No, there aren't any.

Should external/public-facing services be validated during DR testing? Yes

4. Cloud DR SDDC Design & Operations

Are there any cloud provider preferences or restrictions? No preferences nor restrictions, however, cloud provider needs to provide references on similar projects and demonstrate ability to recover quickly and effectively during a disastrous event. The City will look closely into provider's regional and national geographical datacenter presence (including geo-zone network redundancy) during contract award selection phase.

Should the on-demand SDDC be sized for minimum capacity or full production equivalency? Either way is fine. This is left to awarding vendor to provide best service and industry recommendations.

Are there licensing constraints during DR operation? There aren't any.

How long should the City expect to operate in the cloud during a prolonged outage? Cannot answer this question as it depends on actual outage circumstances. (For example, replacement parts aren't readily available, or network connectivity cannot be restored due to

widespread disruption of internet services in the affected area). Standard 30-days is assumed as a minimum in the event of disaster.

5. Networking, IP Addressing & Access

Is IP address consistency required for all recovered VMs or only selected systems? IP address consistency is required on all recovered VMs

Are there existing VPN or SD-WAN solutions to integrate with DR? Yes and No. We have an existing firewall but in the event of disaster it's safe to assume that this firewall would no longer be available (if our datacenters are lost due to disaster). Therefore, by no longer having existing VPN or SD-WAN solutions anymore, there wouldn't be any third-party vendors or solutions to integrate with DR.

Who manages DNS changes and public IP cutovers during DR events? This is handled by the City Communications Department with direct IT Department input. The awarding vendor will be coordinating directly with the IT Department for any potential public IP cutover during DR events.

Are external partners required to access systems during DR? Yes

6. Security, Compliance & Identity

Are there specific CJIS audit or documentation requirements? No. However, awarding vendor will need to prove that they are CJIS compliant. (Clarified in another answer)

Should MFA/SAML integrate with an existing identity provider? Yes

Are additional security monitoring or logging tools required during DR? Preferrable

Are there defined retention or chain-of-custody requirements for DR logs? Currently, there aren't any, but this is subject to change as our DR strategy evolves. As a scope of this project, the City will require all logs being retained for potential cyber-forensics, however, log retention and administration would be responsibility of managed service provider (where IT Department will have ability to review them when needed)

7**. Ransomware Recovery & Cyber Events**

How is a ransomware event detected and declared? The City maintains a relationship with cyber-security managed service provider, which monitors our environment in real-time for potential security threats and breaches. Based on severity of a threat, an emergency event is declared, and authorization for ransomware protection response is granted. Contracting vendor will be included in the ransomware protection chain and part of disaster recovery process.

Is forensic validation required before recovery? Yes

Are there workloads requiring longer retention or immutable recovery points? No, everything is already defined in RFP.

Should ransomware recovery exercises be included in DR testing? Yes. (But, those could be exercised with a set of test VMs and not necessarily with production workloads)

8. Disaster Recovery Plan & Operational Ownership

Who are the primary and secondary DR contacts? Those would be defined in SLA (with IT Department acting as primary DR contact).

Should the DR plan be operational, executive-level, or both? Both

Will City staff participate in recovery operations or be fully vendor-executed? City staff (IT Dept) will participate in recovery operations providing guidelines and monitoring recovery process; however, recovery operations will be fully vendor-executed. In addition, (and as part of requirement), the City will also have ability to execute recovery operations independently (if needed and required)

Are there approval or change-management processes during failover? Yes, however, this is done in coordination with City's IT Department (who manages and approves change management)

9. Commercial & Engagement Clarifications

Are there SLA expectations for recovery execution? Yes, and those will be negotiated in the contract award phase.

Should professional services be fixed-scope or time-and-materials? Fixed

Are there any budgetary constraints or cost caps vendors should be aware of? No. (However everything is subject to City council final vote and approval).

For the Cloud-based SDDC, is the expectation of the customer that the DR instance will be created from start to finish and then managed by the solution provider once completed? Yes Or will the customer take over management responsibilities once it is created and handed over to the customer? No, the City is primarily looking for a managed service provider (who will take ownership and management of Disaster Recovery capabilities)

-Can the amount of data within each of these segments be estimated: Not exactly (as workloads run within Virtual environment which is currently allocating approximately 35 TB of physical storage)

Virtual Machines, (approx. 35 TB for all virtual workloads)

Databases, (2-4 TB part of above mentioned 35 TB)

Media data(photos, videos, GIS), (approx. 8TB with 4TB being part of storage consumed and referenced in the above-mentioned 35 TB storage consumption)

File server / NAS 4 TB part of above mentioned 35 TB

-Are the costs for the separate cloud environment to be included in the total solution's cost?
Yes, all costs should be included in the total solution cost.

1 Cloud based initial 3-year subscription to host 45 VMs (Fully vSphere Compatible especially when falling back – no conversion formats whatsoever) "License procurement and endpoint:
Are you open to running your DR site on Azure VMware Solution or Google Cloud VMware Engine with BYOL VCF, or VMC on AWS via Broadcom purchase, provided all paths keep workloads 100 percent vSphere-native with no conversion during failback? " "vSphere components that must be preserved: (No we are not open to purchase any additional VMware components and services)

Beyond ESXi and vCenter, do you require NSX-backed networking constructs to be identical in the cloud for seamless failback and IP preservation, or is basic VM-level recovery sufficient? "
Yes. Everything on the cloud must match on-prem IP scheme.

2 Annual billing over the course of the agreement is required "Scope of annual billing:
Does the annual billing requirement apply only to the cloud DR hosting platform, or must all services, including recovery execution and managed operations, be billed annually as a single line item?" All services

"Fiscal alignment:
Does the city require true annual invoicing aligned to fiscal years, or is an annualized contract with installment payments acceptable as long as costs are fixed and predictable?" True annual invoicing

3 Total Protected VM storage 50TB "Definition of “50 TB protected”:
Is the 50 TB intended to represent logical source VM data size, or the actual consumed backup storage after retention, snapshots, compression, and deduplication?" Actual consumed backup (logical size is much larger) "Retention expectations (ties directly to 50 TB):
How long must protected VM data be retained in the cloud (for example 30 days, 90 days, 1 year), and does retention differ between backup and DR recovery copies?" 12 weeks minimum. No difference between backup and DR recovery copies.

4 Cloud should be built-out as an On Demand Environment without Pilot-Lights. (Billed by actual usage once SDDC is constructed/operated until SDDC is destroyed) "Hyperscaler flexibility:
If we can achieve no pilot-light and near-zero steady-state cloud cost on Azure or Google but not on AWS due to the end of on-demand elasticity for VMC, is the city open to Azure VMware Solution or Google Cloud VMware Engine as the DR endpoint to honor the “on demand” intent?" Yes as long as it retains native VM data format and not subject to any VM data conversions (e.g. vmdk instead of vhdx) "Contracting intent:
Does “billed by actual usage” mean only the cloud compute should be usage-metered, Yes, while software subscriptions and management services can follow annual terms, or must all components strictly follow a per-use model? However, The City would sign a three-year contract locking in the prices on all services but would need to pay for each year annually. This determines how we structure the Broadcom VCF BYOL subscription and IBM managed services." We prefer to sign a contract with a single vendor which will negotiate its own BYOL subscriptions or managed services.

5 "Billing for Cloud Storage Consumption (for Backup Snapshots) would be separated from

On-Demand Environment Billing" "Invoicing and reporting:
Does the city require separate invoices for cloud storage consumption and on-demand SDDC runtime, or is separate line-item reporting within a single invoice sufficient for audit and budgeting purposes?" Separate line-item reporting within a single invoice "Retention sensitivity:
Should storage costs be optimized using multiple storage tiers based on snapshot age (for example standard vs archive), or is the expectation that all backup snapshots remain immediately recoverable regardless of age?" All backup snapshots remain immediately recoverable, regardless of age.

6 Minimum of two cloud hypervisor hosts in a redundant cloud geographic regions "Intended distribution of the “two hosts”:
Do you expect two hosts per region to achieve in-region HA (with an additional region for geographic redundancy), Two hosts per region or do you literally mean two hosts total split across two regions? This affects both supportability and RTO expectations. " "Region pairing policy:
Does the city require specific region pairings for compliance or latency, or can IBM recommend the closest compliant region pair on your chosen hyperscaler to meet the 4-hour RTO with an on-demand SDDC spin-up? GCVE’s documented fast provisioning helps, and AVS is fully managed operationally, but we must match region proximity and pre-approved quotas to your recovery objectives." Cloud provider can recommend the closest compliant region pair.

7 Datacenters must be based within the United States of America "Geographic scope within the U.S.:
Is redundancy required across different U.S. regions (for example Midwest to East Coast), or must regions also be in separate U.S. compliance zones to mitigate regional disasters?" Redundancy is required. For example, if the Midwest region becomes unavailable, data must be available in the East region. All City data and backups must be stored only in U.S. data centers (in separate U.S. compliance zones) and must not be stored outside the United States.
"Data mobility restrictions:
Should backup data, metadata, logs, and DR test artifacts be restricted from ever leaving the U.S., including during vendor support or troubleshooting activities?" Yes. Everything should be restricted from ever leaving the U.S, including during vendor support or troubleshooting activities.

8 Cloud needs to be CJIS compliant "CJIS program scope:
Will the city require a provider signed CJIS Security Addendum and state-specific CJIS Information Agreement Yes, the City would at least require a notarized statement of CJIS compliance, or is a controls-based approach under CJIS v6.0 sufficient if encryption and data boundary controls eliminate any unescorted cloud provider access to unencrypted CJI? ""Audit posture and baseline:
Does the city’s CJIS Systems Agency plan to audit against CJIS v5.9.5 through March 2026 or will it audit directly against the v6.0 requirements and priorities outlined in the December 27, 2024 policy and companion document? This determines exact control wording and evidence packages. " We are not aware of any current audit plans; however, we prefer staying in compliance with all the latest audit requirements and priorities.

9 "Cloud provider will perform all necessary, virtual networking migration and replicate on-prem environment settings (so that virtual machines cloud IP allocations match the IP assignments on on-prem environment" "Scope of network fidelity:
Does the city require full NSX network parity Full NSX network parity including segments, firewall rules or routing policy in the case of SDDC Cloud Operation. (For example, client machines will be logging into SDDC environment, adhering to all routing policies and firewall rules. (segments, firewall rules, routing policies), or is the primary requirement VM IP

preservation and subnet consistency during recovery?" IP preservation and subnet consistency is required in both cases, but this only applies during failback recovery. Full NSX parity is required when operating strictly from the SDDC cloud data centers. "Connectivity expectations during DR:
During a disaster, should recovered workloads be accessible via existing on-prem IPs through VPN/SD-WAN, Yes, everything should be accessible via existing on-prem IPs through VPN/SD-WAN or is access limited to internal users and systems until failback occurs?"

10 "Public IP addresses allocated for both firewall and Cloud entry points would be included and part of the solution (City will not be responsible for acquiring external public IP addresses for the managed Cloud). Minimum of 16 external public IPs will be reserved and consumed on demand, needed to support various external City services (e.g. GIS servers on public IPs and similar)" "Public IP usage model:
Should the 16 public IPs be permanently reserved for exclusive City Yes, they should, (as those IPs would be used for certain workloads that require additional configuration to ensure minimum downtime during a disaster. (for example: static IP mappings baked into SSL certificates or application aware GIS workloads which strictly operate in static IP settings) use throughout the contract term, or can they be reserved at the provider level and instantiated on demand during DR events?" "Service exposure expectations:
Are externally facing services expected to be reachable via direct IP access (Yes, direct IP access), or is it acceptable (and preferred) to front them with firewall-based NAT or application gateways that cleanly separate public exposure from internal VM networking?" (This is also acceptable if application-level firewall exists and separating public exposure from internal VM segments. Those workloads could be configured and mapped via NAT/PAT)

11 Managed Cloud should also support some sort of build-in DNS & DHCP services (for privateinternal IP resolution) should internal DNS server become unavailable "Scope of DNS fallback:
Is the expectation that built-in DNS only support infrastructure and application startup during DR, This, only applies for a simple DNS redundancy (e.g., DNS for resolving Internet Queries and does not have to replicate AD. or must it fully replicate Active Directory–integrated DNS behavior for the duration of the outage?" Not required but certainly beneficial (if those services are offered as part of built-in service bundle).
"DHCP behavior preference:
Should DHCP only be used for infrastructure continuity if primary services fail Yes, or must it support dynamic re-allocation and scaling of recovered VMs during extended DR operations?" Dynamic re-allocation not required

12 "Managed Cloud Provider will also offer built-in, highly available NTP services for instances,
typically accessed via simple hostnames to keep VMs and services synchronized, reducing dependency on external internet servers while ensuring accuracy and compliance within their ecosystems" "Time hierarchy expectations:
During disaster recovery, should the cloud environment be authoritative for time Yes (as on-prem NTP time services wouldn't be available), or must it resynchronize back to on-prem sources once they become available during failback? (It must re-synchronize back to on-prem only after on-prem NTP appliance becomes available. If time appliance is still unavailable after failback, vendor can temporary substitute with public time servers until NTP appliance gets redeployed).

" "Compliance validation:

Does the city require documented evidence (for audits) that recovered workloads are synchronized to provider-managed, internal NTP sources, rather than public internet NTP pools?" Yes

13 "Virtual Firewall needs to have ability to support VPN Clients and Cloud Provider will assist with VPN Configurations for Client Machines connecting to Cloud Environment. Virtual firewall needs to support Ike2 and MFA authentication." "VPN user scope:
Is remote VPN access intended primarily for IT administrators and support staff, or must it also support broader user or inter-agency access during DR?" It must support all City personnel, not only IT administrators and support staff. Broader user support is required "MFA integration preference:
Does the city have an existing MFA provider that must be integrated with the VPN (for example a centralized identity platform), Yes, we have the existing MFA provider. or is the managed cloud provider expected to supply and operate MFA as part of the solution?" This would be beneficial to have but not necessarily part of project requirement. If offered, this could be used as part of built-in redundancy service bundle.

14 Cloud provider will also help with third party firewall deployments (e.g. Watchguard FireBox) "Firewall continuity expectation:
Is the intent to use third-party firewalls (such as WatchGuard Firebox) as the primary ingress firewall in the DR environment, or as a supported alternative alongside native cloud firewalls?" Supported alternative. We would prefer the cloud native firewall with Ikev2 and MFA capability. "Operational responsibility:
During a disaster event, should the managed cloud provider be responsible for operating and modifying third-party firewall policies Yes, this will be the responsibility of the managed cloud provider, or is support limited to deployment and baseline configuration?"

15 Total vCPU 178 GHz "Source of the 178 GHz figure:
Was the 178 GHz value derived from on-prem host specifications, VMware reports, or application performance baselines? No, it is not. Actual on-prem host specifications are much larger. This specification is based on metrics measured for designing our current DR solution. SDDC minimum capacity was calculated through reports, and baselines so total vCPU represents minimum host specifications needed to run everything from an SDDC data center. This helps us normalize capacity accurately in the cloud SDDC." "Performance expectation during DR:
Is the expectation that all workloads run at full production performance during DR (Yes, full production performance), or is reduced but functional performance acceptable for the duration of an outage?"

16 Total vRAM 1.5 TB "Memory reservation expectations:
Does the 1.5 TB value represent total allocated VM memory, Yes, total minimum allocated VM memory expectations or does it include memory reservations There aren't any memory reservations or peak usage that must also be honored during DR?" "Duration of DR operation:
If the DR environment must operate for an extended period, should the 1.5 TB sizing support steady-state operations Yes, it should support steady-state operations, or is temporary performance degradation acceptable during prolonged outages? No performance degradation is acceptable during prolonged outages."

17 RTO of 4 hours "Scope of the 4-hour RTO:
Does the 4-hour RTO apply to all 45 VMs Yes, or to a defined subset of priority systems required for continuity of operations?" "Start point definition:

Should the 4-hour clock begin at formal disaster declaration Yes, formal disaster declaration, or at loss of primary datacenter services? This affects how the SLA is measured and audited."

18 Ransomware Recovery for 40 VMs "Definition of the 40 VMs:
Are the 40 VMs a fixed set of known critical systems (No, as some of the workloads are Linux based appliances or not critical), or should the solution support dynamic selection of which VMs require ransomware recovery based on the incident?" (Yes, and it applies only on 85% of VMs – mostly VMs running Windows) "Containment expectations:
During a ransomware event, does the city expect the managed cloud provider to assist with containment and isolation Yes (for example, disabling connectivity until clean recovery is verified), or only with VM restoration?" Yes, for both. (Both containment and isolation, including VM restoration. Not just VM restoration)

19 60 Hours of Disaster Recovery test days should be included for City to test the ability to recover in the event of disaster. In addition, two Recovery Tests Yearly need to be performed by managed provider, who will also provide bi-yearly reports on Disaster Recovery Capabilities
"Use of the 60 test hours:
Should the 60 hours be consumed only during formal DR tests, or may they also be used for ad-hoc validation, ransomware simulations, or City-initiated testing?" "Test scope expectations: The allotted hours may be used for other activities as long as they do not reduce the time reserved for formal DR testing. For example, if two annual DR tests require 40 total hours, the remaining 20 hours should be available for City-initiated or ad-hoc testing and should not be fully consumed by the provider.
During the two annual provider-executed DR tests, is the expectation to recover all protected workloads Yes, all protected workloads during two annual provider-executed DR tests, or a defined priority subset aligned to the 4-hour RTO?" Only, for ad-hoc validation, ransomware simulations or City-initiated testing.

20 "5 Failover test days per year (Meaning that Provider will also demonstrate ability to failover to SDDC and back to on-prem environment). SDDC operation cost will be included in the
price during these failover test days (where City will not be billed for SDDC consumption during times these failover tests are performed). These tests will be performed by Cloud
Managed Provider with minimum of 7 (seven) test VMs (where provider will measure and report on times it took to perform failover on a test VM set, with an average of 250GB in VM size)."
"Failover scope per test:
During each of the five failover test days, should the provider recover only the defined 7-VM test set, Yes, only defined 7 VM test sets (which are actual Virtual Machine replicas) or is periodic testing of additional workloads expected over the course of the year?" Not required, but certainly a benefit. "Measurement definition:
Should reported failover timing be measured:
- From test initiation to VM power-on, Yes, (from test initiation to VM power-on) or
- To application or service availability within the VM?" (Not required – optional)

21 Ability to use both Active Directory Native, as well as other forms of replications for both Virtual Environment and Windows Active Directory (e.g. Zerto, Veeam B&R or similar technologies for VM replications) For Active Directory, should the recovery design use a dedicated recovery forest for ransomware-assured restores when needed Yes, as this will restore active directory in a clean isolated environment so that no ransomware or malware is re-introduced into production network during restoration process, or does the City require in-place recovery of existing domains using native replication only? For VM replication tooling, does the City have an existing operational preference or licenses for Zerto or Veeam that you want the provider to adopt, or should the provider recommend the primary tool after discovery and

workload characterization? Provider should recommend and include primary tool. (For example: If provider recommends Zerto, those licenses should be part of managed services solution bundle)

22 200 IOPS/TiB "Scope of the IOPS requirement:
Should the 200 IOPS/TiB requirement apply only to active VMware datastores during DR, or does the City expect this performance level for backup or recovery staging storage as well?" This applies to SDDC for both VMware datastores during DR as well as recovery staging storage as well. "Latency expectations:
Is there a target or maximum acceptable storage latency during DR operations (for example, under normal load vs failover scenarios), or is IOPS the primary performance metric of concern?" This hasn't been defined and it's up to vendor to recommend and offer best industry standard latency as a part of managed solution.

23 Steady state CPU headroom should equal 15% "Scope of the 15 percent headroom:
Should the 15 percent reserve be measured cluster-wide across all DR hosts (preferred and more resilient), or per-host (stricter and may require more capacity)?" "Measurement window:
Should compliance be demonstrated over a defined steady-state window during tests (for example, a continuous 60–120 minutes after all VMs reach green state) to avoid counting the initial boot storm and background catch-up tasks?" Not necessarily required but certainly a benefit if provider could demonstrate those as well.

24 CPU utilization 30% "Validation intent:
Should the 30% CPU utilization requirement be treated as a design target validated during DR tests and reported bi-yearly Yes, design target validated during Full DR tests and reported bi-yearly, rather than a hard operational limit at all times?" "Relationship to headroom:
Is the 15% steady-state headroom the mandatory threshold, with 30% representing a desired operating profile under normal DR conditions?" Yes

25 Memory utilization 100% "Intent of "100 percent":
Should we treat "100 percent memory utilization" as a burst tolerance requirement (acceptable during boot-storm) rather than a steady-state operating target?" Yes "Steady-state window:
Over what period should "steady state" be measured during tests, for example, the first 60–120 minutes after all recovered VMs reach green state, to exclude initial boot-storm effects?" Yes 60-120 minutes after all recovered VMs reach green state.

26 CPU overcommit factor 4 "Scope of enforcement:
Should the 4:1 overcommit be enforced cluster-wide (preferred and most practical) Yes if all objectives are met. or per-host (stricter and may require additional capacity to maintain)?" "Exception policy:
May the provider declare exceptions for specific latency-sensitive or licensed workloads (kept below 4:1) as long as the overall cluster average remains ≤ 4:1 and all RTO/performance objectives are met?" Yes, provider could do that, if and only if, all RTO/performance objectives are met.

27 Memory overcommit factor 1.25 "Enforcement scope:
Should the 1.25 memory overcommit factor be enforced cluster-wide (standard and recommended), Yes, just standard and recommended or per-host (stricter and may require additional capacity)?" No, not needed. "Steady-state target:
After stabilization during DR tests, is there an expected steady-state memory utilization band (for example ≤ 85–90%) even though the platform must tolerate temporary peaks up to full

utilization?" Not required during DR tests, only if (and when) switching to run full production workloads (completely from SDDC cloud datacenter in event of a disaster).

28 Average boot-up time of a VM after power-on begins should be no greater than 8 minutes "Definition of “boot-up complete”:
Should boot time be measured until the guest OS is responsive and reachable, or until all applications inside the VM are fully started?" Until guest OS is responsive and reachable "Test measurement scope:
Should the 8-minute average apply to the 7-VM failover test set, or a broader set of recovered VMs during full DR exercises?" It could be combination of both however, 8-minutes average should be maintained during full DR exercises.

29 De-dupe and compression ratio should meet or exceed industry standard "Scope of “industry standard”:
Should the ratio be evaluated primarily on backup/immutable storage (Yes, evaluated on backup / immutable storage) (where dedupe/compression are most impactful), or do you also want a target for replication/journal storage used for short-RPO recovery?" (Yes, also on target for replication/journal storage used for short-RPO recovery) "Evidence expectation:
Are periodic KPI reports showing logical vs physical footprint and effective reduction (by tier and by dataset) sufficient to demonstrate compliance (Yes this would be sufficient), or do you want the provider to contractually commit to a minimum reduction figure (not recommended without analyzing data composition)?"

30 "Snapshots and retention should follow:
o Every 4 hours and retained for 2 days
o Daily at 12AM and retained for 7 days
o Weekly on Sundays and retained for 3 weeks
o Monthly on the 1st and retained for 6 months" "Snapshot technology preference:
Should these snapshots be implemented exclusively using backup tool like managed snapshots (No, they should not) (preferred for ransomware safety and retention enforcement), or is the City expecting hypervisor-native snapshots to also be retained across this entire schedule?" Hypervisor-native snapshots should also be retained across entire snapshot service beside backup managed snapshots. "Immutability expectation:
For the weekly and monthly snapshots, does the City require immutability/WORM protection for ransomware assurance, or is standard retention enforcement sufficient?" Standard retention enforcement is sufficient.

31 List cost for Egress per GB and Ingress per GB if any. N/A

32 List cost for professional services needed for deployment, initial setup and training. This should be part of the vendor’s bid.

33 Creation of a Disaster Recovery Plan document that includes, but is not limited to a background to the City’s DR capabilities, and a step by step direction of how to utilize SDDC and full restoration of data onto on-premise equipment "Format preference:
Does the City require the Disaster Recovery Plan in a specific format (for example Word or PDF), and should it align to any existing City or CJIS documentation standards?" "Audience scope: Word and PDF formats are acceptable. Documentation should not align to any existing standards.
Should the plan include both technical runbooks and a high-level executive summary, or will those be delivered as separate documents?" Two separate documents.

34 SAML/MFA for access to system is required "IdP & factors:

Which Identity Provider and MFA methods must we integrate with (for example, Entra ID with FIDO2 and push, or Okta with WebAuthn and OTP), and are there any City-mandated Conditional Access rules we must inherit?" Entra ID with FIDO2 and Push. Duo Security MFA

"Break-glass expectation:

Does the City require a dual-custodian break-glass model with time-boxed access (No, it's not required) and mandatory post-event review (Yes), and should we schedule an annual drill to prove the process?" (Yes)

Reminder:

The due date and time remain the same for this project. This is to remind all bidders that sealed proposals for **RFP #26-009 Cloud-Based Backup and Disaster Recovery Solution** must be submitted by 12:00 P.M. CST, February 19, 2026, to:

Christine Grabin, Purchasing Agent
City of O'Fallon, Missouri
100 North Main Street
O'Fallon, Missouri 63366

Addendum information will be available over the Internet at www.ofallonmo.gov Adobe Acrobat® Reader may be required to view this document. We strongly suggest that you check for any addenda a minimum forty-eight hours (48) in advance of the bid deadline. Due to revisions, the bidders must acknowledge the Addendum(s) on the bid form.

END OF ADDENDUM #1