Anexa nr. 22
la Documentația standard
aprobată prin Ordinul Ministrului Finanţelor
nr. 115 din 15.09.2021

Specificaţii tehnice

[Acest tabel va fi completat de către ofertant în coloanele 2, 3, 4, 6, 7, iar de către autoritatea contractantă – încoloanele 1, 5,]

Numărul procedurii de achiziție: *ocds-b3wdp1-MD-1740577309675*						
Obiectul achiziției: *Servere și sisteme de stocare (perioada 2024-2025)*						
Denumirea bunurilor/ serviciilor	**Denumirea modelului bunului/ serviciului**	**Ţara de origine**	**Produ-cătorul**	**Specificarea tehnică deplină solicitată de către autoritatea contractantă**	**Specificarea tehnică deplină propusă de către ofertant**	**Standarde de referinţă**
1	**2**	**3**	**4**	**5**	**6**	**7**
Bunuri/servicii						
Lotul nr. 3 Enterprise Storage (Sisteme de stocare) tip 1(SAS SSD**)**	IBM FlashSystem 9500	SUA	IBM	Echipament nou și nerecondiționat, produs minim trim. I anul 2024, corespunzător tipului de dispozitive de nivel Enterprise, produs de producători renumiți (Brand name internațional). Configuraţia echipamentului trebuie să fie compusă din componente reciproc compatibile şi să asigure funcţionarea optimă a sistemului în ansamblu. **Type:** Enterprise-grade Storage with SAS SSDs. **Form Factor:** min. 2U rack-mountable chassis, fully compatible with the EIA-310 standard for rack mounting. The solution must include all necessary components (e.g., rails, mounting brackets). **Availability requirements:** The equipment must be working in Symmetric Active-Active mode, which means that in the case of 100% utilization, ensures following: - The storage system architecture must ensure that, in the event of a controller failure, the write cache of the	**Conform anexei nr. 1 - Matrice conformitate storage (pct.21 din ofertă)**	

				surviving controller(s) remains fully operational and protected. The equipment must utilize mechanisms such as cache mirroring or equivalent protection to guarantee data integrity. Under no circumstances should the write cache be deactivated, operated without mirroring, or left without an alternative protection mechanism to prevent data loss or corruption. - The system must ensure a high availability rate of at least 99.9999%, minimizing downtime and guaranteeing continuous operation, - The system's efficiency must remain unaffected in the event of a failure of up to 50% of the controllers, maintaining consistent operational capability - alive with a single active controller, - The system must sustain its required performance levels without degradation in the event of a failure affecting half of the controllers, - The system must include robust, built-in mechanisms for non-disruptive software updates, ensuring no compromise in availability or loss of access to stored data during version upgrades. The storage system must ensure uninterrupted data availability and full operational continuity in the following failure scenarios: - failure of a single power supply line, ensuring redundancy in power management, - failure of any individual controller, with automatic failover mechanisms to maintain functionality - alive with a single active controller, - simultaneous failures of up to two user data storage drives, with no loss of data integrity or accessibility, - failures of any Fibre Channel (FC) or iSCSI port, with seamless rerouting of traffic to alternate pathways. The equipment must support hot-swappable replacement of critical components without interrupting access to data or degrading system performance. These components include, but are not limited to: controllers, power supplies, cooling fans, front-end and back-end ports, and storage drives. The hot replacement process must ensure seamless operation and maintain data availability throughout.		

				The system must be designed to withstand the simultaneous failure of at least two storage devices (e.g., drives, NVMe, or flash modules), regardless of the system's scale or configuration. In such scenarios, the equipment must ensure uninterrupted data access and maintain full data integrity. The system must include functionality to safely disable the storage drives without causing any loss or corruption of user data, ensuring seamless operational continuity during maintenance or decommissioning. **Type Drives:** **Enterprise-grade SAS SSD**s utilizing TLC (Triple-Level Cell) or eTLC (Enhanced Triple-Level Cell) technology, optimized for high-performance, high-reliability applications in enterprise environments. **Capacity:** The system must provide a marked usable storage capacity (before data reduction) of **minimum 200 TB**, ensuring sufficient space for high-demand enterprise applications. **Hot Spare Configuration(*optional*):** The solution must optionally support Hot Spare components, including spare controllers or disks, to enhance system redundancy. These spare components must remain inactive during regular operations but should automatically activate to maintain full system functionality in case of hardware failure. **RAID (*if the equipment involves the use of RAID*):** - The system must support advanced RAID levels, including minimum: RAID 6: Ensuring double parity protection, allowing the system to tolerate simultaneous failure of two drives without data loss. **Cache requirement(*if the equipment involves the use of memory cache for data*):** If the storage system includes a cache mechanism, the system must provide a **minimum of 512 GB** of		

				dedicated cache memory per node, ensuring high-speed data processing and optimal system performance. The cache must support advanced features such as: - Cache mirroring - to ensure data integrity and protection in the event of a node failure. - Dynamic allocation - enabling efficient use of cache resources based on real-time workload demands. - Non-volatile cache - to prevent data loss during power failures or unexpected shutdowns, ensuring all cached data is retained. The cache must be optimized for handling high IOPS workloads and ensuring low-latency operations, particularly for enterprise-grade applications. **Controllers requirement:** The storage system must include minimum one node equipped with a minimum of two fully redundant controllers configured in High Availability (HA) mode. The controllers must: - Operate in an Active-Active configuration, ensuring balanced workload distribution and seamless failover capabilities without performance degradation. - Support advanced fault-tolerant mechanisms to maintain uninterrupted access to data during hardware failures or maintenance. - Be hot-swappable, allowing replacement or upgrade without disrupting system operations or data availability. - Include built-in synchronization mechanisms to maintain consistency between controllers, including mirroring of critical operational data such as cache contents and configuration settings. The system must ensure that the failure of one controller does not impact the performance, availability, or operational integrity of the other controller. **Cluster and replication requirements:** 1. Synchronous replication capability: - The storage solution must support synchronous replication to enable the creation of an Active-Active cluster between two physically separated server rooms (located in separate buildings).		

				- The system must ensure zero Recovery Point Objective (RPO) by maintaining data consistency across the cluster in real time. 2. Comprehensive hardware inclusion: - The solution must include all necessary hardware components to fully implement synchronous replication functionality, utilizing Fibre Channel (FC) protocols for high-speed, low-latency data transmission. 3. Flexible volume replication: - The system must support synchronous replication for a minimum of one Logical Unit Number (LUN) and scale seamlessly to replicate multiple LUNs simultaneously. - Changes to the number of replicated volumes must not require modifications to the physical hardware configuration of the storage system. 4. Data consistency and synchronization: - The contents of all cluster volumes must remain identical across both systems in the cluster at all times, ensuring data consistency and integrity. - The system must include mechanisms to handle data synchronization efficiently during recovery scenarios, ensuring minimal impact on performance and availability. 5. Resiliency and high availability: - The cluster must provide continuous operation in the event of a hardware failure, network disruption, or planned maintenance at one site, without compromising data integrity or availability. - The system must be designed to support failover and failback between the two sites automatically and transparently. **Performance requirements:** 1. Minimum performance metrics: - the storage solution must deliver a combined performance of **minimum 300,000 Input/Output Operations Per Second (IOPS)** with inline data reduction (deduplication and compression). 2. Performance calculation parameters: IOPS performance must be evaluated based on the following metrics: - read/write ratio: 70% read / 30% write.		

				- block sizes: support for operations with block sizes of 16 KB, 32 KB, and 64 KB to accommodate varying workload requirements. - I/O patterns: include both sequential and random I/O workloads. - latency: ensure a maximum delay of 1 millisecond (0.001 s) under full load conditions. 3. Consistency of performance: - the system must maintain the required performance levels even under high concurrency and mixed workload conditions. - performance must remain unaffected during maintenance operations, including firmware updates, drive rebuilds, or component failures. 4. Performance verification: - vendors must provide detailed benchmark test results to validate the stated performance – **for operations with block sizes 16 KB(mandatory)**, 32 KB and 64 KB(optionall), using industry-standard tools such as IOmeter or FIO, under the specified conditions. - results must demonstrate compliance with all stated parameters, including latency and I/O patterns. 5. Monitoring and optimization: - the system must include tools to monitor and optimize performance dynamically, offering real-time insights into throughput, latency, and IOPS for proactive performance tuning. **Supported protocols:** - FC, - iSCSI, **Features:** Dedicated system management interfaces: 1. The system must include dedicated physical and/or virtual interfaces specifically for system management. 2. These interfaces should allow out-of-band management, ensuring that administrative tasks can be performed without impacting data traffic. 3. Management interfaces must support the following functionalities: - Web-based GUI for ease of access.		

				- Command-line interface (CLI) for advanced configuration. - Support for industry-standard protocols such as SSH, SNMP, and REST API for integration with monitoring and orchestration tools. - Role-based access control (RBAC) to ensure secure system administration. 4. Redundancy for management interfaces: - to ensure availability, the management interfaces must support redundancy, allowing continuous system management even in the event of a single interface failure. 5. Protocol optimization: The system must include protocol-specific optimizations such as: - Multipath I/O (MPIO) for FC and iSCSI to ensure high availability and load balancing. - Support for jumbo frames in iSCSI for improved performance in high-throughput environments. 6. Compliance and Interoperability: The system must be compliant with industry standards for both FC and iSCSI protocols. It must ensure interoperability with third-party devices, including servers, switches, and network adapters. **Deduplication and compression requirements:** 1. Functional capabilities: The storage system must provide deduplication functionality for data stored at the block level (iSCSI/FC LUN) and file level, with the following specifics: - Deduplication must operate both at the volume level and globally across the system, ensuring optimal storage efficiency. The system must also include compression functionality for: - Block-level volumes (iSCSI/FC LUN). 2. Interoperability and unrestricted functionality: Deduplication and compression features must operate seamlessly without introducing limitations or restrictions on simultaneous use of other critical functionalities, including but not limited to: - Data replication. - Thin provisioning.		

				- Backups. - Volume cloning. 3. Inline deduplication and compression: - Both deduplication and compression mechanisms must function in in-line mode, ensuring real-time data optimization without requiring post-processing. - Deduplication must remain continuously active and cannot be disabled or bypassed by system administrators or any other means, ensuring consistent storage efficiency and data integrity. **- Storage solutions that rely on scheduled or job-based data reduction processes are not acceptable.** 4. Licensing and support: All features related to deduplication and compression must be: - Fully licensed (if required by vendor provisions) and included in the offer, eliminating additional licensing costs for essential functionality. - Supported by the storage system in its maximum configuration, ensuring scalability and compatibility across all deployment scenarios. 5. Performance and reliability considerations: - The deduplication and compression mechanisms must not introduce significant latency or impact the system's performance metrics, such as IOPS or throughput. - Mechanisms should include built-in error detection and correction to maintain data integrity during deduplication and compression processes. 6. Management and monitoring: The system must provide a dedicated interface or tools for monitoring deduplication and compression efficiency, including: - Space savings metrics. - Real-time and historical performance impacts. - Detailed logs of deduplication and compression activities. **Snapshot requirements:** 1. General functionality: - The system must support snapshot functionality at a minimum for block-level volumes (LUNs), ensuring operational flexibility.		

				- The snapshot functionality must be applicable to both LUNs and other supported volumes without imposing restrictions on the simultaneous use of other critical system functions, including replication, backups, and cloning. 2. Snapshot quantity and retention: - The system must provide the ability to create and manage a minimum of 365 snapshots per shared volume, supporting long-term operational and recovery needs. - Snapshots must be configurable with retention policies to optimize storage space and align with data governance requirements. 3. Performance efficiency: - The implementation of snapshots must not degrade overall system performance, regardless of the number of active snapshots or system workload. - The system must include optimization mechanisms, such as metadata indexing and intelligent snapshot scheduling, to minimize latency and maintain high performance. 4. Space efficiency: - Snapshot functionality must employ a cost-effective approach by storing only the delta (changes) from the original data. This ensures minimal storage consumption while preserving full data access and recovery capabilities. 5. Integration with storage QoS: - The system must support performance monitoring and prioritization mechanisms for snapshots, enabling administrators to enforce Storage QoS (Quality of Service) policies at both the volume and LUN levels. - These QoS policies should dynamically allocate resources to prioritize performance-critical snapshots, ensuring minimal impact on other operations. 6. Advanced features: Snapshots must support: - Application-consistent snapshots, ensuring data integrity for workloads such as databases and virtualized environments. - Writable snapshots, allowing clones to be created for development, testing, or analytics without affecting the production environment.		

				Snapshots must be compatible with data replication workflows, ensuring consistent replication of both primary data and snapshot states across systems. 7. Monitoring and reporting: - The system must include a dedicated interface or tools for managing, monitoring, and reporting on snapshot performance, space utilization, and recovery operations. - Real-time alerts and historical logs must be available for visibility into snapshot performance and potential bottlenecks. **Encryption requirements:** 1. Encryption standard: - The solution must support encryption of all stored data using a minimum of the AES-256 algorithm or a stronger industry-standard encryption algorithm, ensuring compliance with modern security and regulatory standards. 2. Scope of encryption: - Encryption must be applied to all drives, NVMe, and flash storage within the device, covering the entire data storage ecosystem. - Encryption must extend to data at rest across all volumes, snapshots, backups, and metadata associated with the system. 3. Performance integrity: - Encryption functionality must operate with no measurable impact on system performance, ensuring IOPS, throughput, and latency metrics remain consistent with non-encrypted operations. - The system must leverage hardware-accelerated encryption or equivalent technologies to maintain optimal performance during data encryption and decryption processes. 4. Key management: - The solution must generate encryption keys using a secure hardware-based random number generator, ensuring keys are robust and resistant to attacks. - Encryption keys must be securely stored on the equipment, leveraging a dedicated hardware security module (HSM) or equivalent secure enclave to isolate keys from unauthorized access.		

				- The system must ensure that data stored on drives/NVMe/flash cannot be accessed if the storage media is removed from the device or if the device itself is compromised. 5. Key backup and recovery: - The system must include mechanisms for secure backup and recovery of encryption keys, supporting integration with external key management systems (KMS) compliant with KMIP (Key Management Interoperability Protocol) standards. - Key rotation and lifecycle management should be automated and configurable to align with organizational policies and compliance requirements. 6. Encryption for replication and snapshots: - The encryption functionality must extend to replicated data and snapshots, ensuring consistency in encryption across all replicated sites or volumes. - Encryption must not disrupt or degrade replication workflows, including synchronous and asynchronous modes. **Monitoring requirements:** 1. Analytical platform or portal: - The system must include a robust analytical platform or virtual machine (VM) accessible via a web browser-based portal. - The platform must provide an intuitive, user-friendly interface with interactive dashboards for data visualization and management. 2. Log collection and reporting: The platform must automatically collect and analyze logs from the device and present them as customizable graphs, reports, and alerts, covering the following: 2.1. Storage utilization: - Real-time and historical monitoring of used space. - Display of the data reduction indicator, accounting for deduplication and compression (excluding thin provisioning, if applicable). - Granular visibility at both the global device level and the local LUN level. 2.2. Space growth prediction:		

				- Advanced forecasting tools for predicting space growth, factoring in deduplication, compression, and provisioning trends. - Tools for future expansion analysis, including recommendations for scaling. 3. Component monitoring: The system must include an application or hardware-based monitoring solution to oversee and report detailed events for the following physical and logical components: - Physical components: controllers, drives, ports, power supplies, and network interfaces. - Logical components: volumes, LUNs, replication processes, deduplication, and compression algorithms. 4. Performance monitoring: The portal must provide minimum: - Real-time and historical performance metrics for individual resources. - Key parameters to monitor: Latency, Read and Write IOPS, Bandwidth. Performance data must be available at both the global system level and the LUN level. 5. Storage QoS and prioritization: - The system must include a performance monitoring and prioritization mechanism for Storage QoS, configurable at both the volume and LUN levels. - QoS metrics should be adjustable in real-time to meet dynamic workload demands. 6. Reporting and alerting: The portal must provide comprehensive reporting capabilities, including at least: - Capacity reports: current usage, available space, and forecasted capacity needs. - Performance reports: historical trends and real-time analytics of system performance. - Future space predictions: automated simulations for capacity increases based on application type and workload. - Event logs: authorization attempts, executed commands, and system alerts for security and operational events. - Technical support logs: level of support received, resolution times, and incident history.		

				7. Operational monitoring: - Snapshot and replication status: display the real-time status of operations such as snapshots, synchronous/asynchronous replication, and recovery tasks. - Threat alerts: warnings related to system integrity, user activity, or misconfigurations. - Optimization insights: recommendations for system performance improvement, resource reallocation, or energy efficiency. 8. Configuration verification and upgrades: - The platform must include an algorithm for verifying configuration correctness and compatibility with potential device or cluster upgrades. 9. Simulation and optimization: - The platform must enable capacity simulation tools to project storage needs based on application types and expected workloads. - Display real-time system consumption metrics with actionable optimization guidelines for improving performance and efficiency. **NICs included per controller:** Min. 1 x 1GE for management; Min 2 x 32G FC SFP28(850nm SFP+ SR MM module included) for data transfer; Min. 2 x 32G FC dedicated for replication (metro cluster). **Supported operating environments:** Microsoft Windows Server; Red Hat Enterprise Linux; VMware (VMware ESXi); **Power supplies included:** The system must include a minimum of two (2) hot-swappable (hot-plug) Power Supply Units (PSUs). The PSUs must support at least 1+1 redundancy, ensuring continuous operation in case of failure of one PSU. Power cables included must meet the following specifications: - Type: IEC C13 to C14. - Minimum length: 0.6 meters (24 inches).		

				<u>Cerințe obligatorii pentru prestarea serviciilor de punere în funcțiune, a garanției și a serviciilor de suport (deservire și mentenanță) a bunurilor -</u> <u>conform Anexei la Anunțul de participare</u>**.** Toate licențele necesare (dacă se aplică conform termenilor și condițiilor producătorului) pentru caracteristicile platformei/portalului de monitorizare (analitică) și software-ului/firmware-ului specific sistemului de stocare, inclusiv actualizările/patch-urile periodice, trebuie să fie incluse în ofertă și furnizate pe o bază perpetuă - valabile obligatoriu pentru durata integrală de viață a sistemului de stocare. **Termeni și condiții:** Toate cerințele sunt minime și obligatorii; O cerință nu trebuie să limiteze o altă cerință; Toate componentele trebuie să fie actuale și să nu fie promovate ca EOS (sfârșitul vânzării/suportului) / EOL (sfârșitul duratei de viață); Extinderea memoriei (ram) și a capacității de stocare nu trebuie să includă limitări hardware sau software.		
Lotul nr. 4 Enterprise Storage (Sisteme de stocare) tip 2 (Full flash)	IBM FlashSystem 9500	SUA	IBM	Echipament nou și nerecondiționat, produs minim trim. I anul 2024, corespunzător tipului de dispozitive de nivel Enterprise, produs de producători renumiți (Brand name internațional). Configurația echipamentului trebuie să fie compusă din componente reciproc compatibile și să asigure funcționarea optimă a sistemului în ansamblu. **Type:** Enterprise-grade Storage with Full Flash **Form Factor:** min. 2U rack-mountable chassis, fully compatible with the EIA-310 standard for rack mounting. The solution must include all necessary components (e.g., rails, mounting brackets). **Availability requirements:** The equipment must be working in Symmetric Active-Active mode, which means that in the case of 100% utilization, ensures following: - The storage system architecture must ensure that, in the event of a controller failure, the write cache of the	**Conform anexei nr. 1 - Matrice conformitate storage (pct.21 din ofertă)**	

				surviving controller(s) remains fully operational and protected. The equipment must utilize mechanisms such as cache mirroring or equivalent protection to guarantee data integrity. Under no circumstances should the write cache be deactivated, operated without mirroring, or left without an alternative protection mechanism to prevent data loss or corruption. - The system must ensure a high availability rate of at least 99.9999%, minimizing downtime and guaranteeing continuous operation; - The system's efficiency must remain unaffected in the event of a failure of up to 50% of the controllers, maintaining consistent operational capability - alive with a single active controller; - The system must sustain its required performance levels without degradation in the event of a failure affecting half of the controllers; - The system must include robust, built-in mechanisms for non-disruptive software updates, ensuring no compromise in availability or loss of access to stored data during version upgrades. The storage system must ensure uninterrupted data availability and full operational continuity in the following failure scenarios: - failure of a single power supply line, ensuring redundancy in power management, - failure of any individual controller, with automatic failover mechanisms to maintain functionality - alive with a single active controller, - failures simultaneous failures of up to two user data storage drives, with no loss of data integrity or accessibility, - failures of any Fibre Channel (FC) or iSCSI port, with seamless rerouting of traffic to alternate pathways. The equipment must support hot-swappable replacement of critical components without interrupting access to data or degrading system performance. These components include, but are not limited to: controllers, power supplies, cooling fans, front-end and back-end ports, and storage drives. The hot replacement process		

				must ensure seamless operation and maintain data availability throughout. The system must be designed to withstand the simultaneous failure of at least two storage devices (e.g., drives, NVMe, or flash modules), regardless of the system's scale or configuration. In such scenarios, the equipment must ensure uninterrupted data access and maintain full data integrity. The system must include functionality to safely disable the storage drives without causing any loss or corruption of user data, during maintenance or relocation of the device.. **Type Drives**: **Enterprise-grade NVMe/Flash** utilizing TLC (Triple-Level Cell) or eTLC (Enhanced Triple-Level Cell) technology, optimized for high-performance, high-reliability applications in enterprise environments. **Capacity:** The system must provide a marked usable storage capacity (before data reduction) of **minimum 600 TB**, ensuring sufficient space and maximum performance for high-demand enterprise applications. **Hot Spare Configuration(*optional*):** The solution must optionally support Hot Spare components, including spare controllers or disks, to enhance system redundancy. These spare components must remain inactive during regular operations but should automatically activate to maintain full system functionality in case of hardware failure. **RAID (*if the equipment involves the use of RAID*):** - The system must support advanced RAID levels, including minimum: RAID 6: Ensuring double parity protection, allowing the system to tolerate simultaneous failure of two drives without data loss. **Cache requirement(*if the equipment involves the use of memory cache for data*):**		

				The storage must provide a **minimum of 512 GB** of dedicated cache memory per node, ensuring high-speed data processing and optimal system performance. The cache must support advanced features such as: - Cache mirroring - to ensure data integrity and protection in the event of a node failure. - Dynamic allocation - enabling efficient use of cache resources based on real-time workload demands. - Non-volatile cache - to prevent data loss during power failures or unexpected shutdowns, ensuring all cached data is preserved and immediately available after hardware recovery from power failures or unexpected shutdowns. The cache must be optimized for handling high IOPS workloads and ensuring low-latency operations, particularly for enterprise-grade applications. **Controllers requirements:** The storage system must include minimum one node equipped with a minimum of two fully redundant controllers configured in High Availability (HA) mode. The controllers must: - Operate in an Active-Active configuration, ensuring balanced workload distribution and seamless failover capabilities without performance degradation and data loss. - Support advanced fault-tolerant mechanisms to maintain uninterrupted access to data during hardware failures or maintenance (until the technical interventions are provided). - Be hot-swappable, allowing replacement or upgrade without disrupting system operations, performance or data availability. - Include built-in synchronization mechanisms to maintain consistency between controllers, including mirroring of critical operational data such as cache contents and configuration settings. The system must ensure that the failure of one controller does not impact the performance, availability, or operational integrity of the other controller. **Cluster and replication requirements:** 1. Synchronous replication capability:		

				- The storage solution must support synchronous replication to enable the creation of an Active-Active cluster between two physically separated server rooms (located in separate buildings). - The system must ensure zero Recovery Point Objective (RPO=0) by maintaining data consistency across the cluster in real time. 2. Comprehensive hardware inclusion: - The solution must include all necessary hardware components to fully implement synchronous replication functionality, utilizing Fibre Channel (FC) protocols for high-speed, low-latency data transmission. 3. Flexible volume replication: - The system must support synchronous replication for a minimum of one Logical Unit Number (LUN) and scale seamlessly to replicate multiple LUNs simultaneously. - Changes to the number of replicated volumes must not require modifications to the physical hardware configuration of the storage system. 4. Data consistency and synchronization: - The contents of all cluster volumes must remain identical across both systems in the cluster at all times, ensuring data consistency and integrity. - The system must include mechanisms to handle data synchronization efficiently during recovery scenarios, ensuring minimal impact on performance, availability and corrupted/degraded data. 5. Resiliency and high availability: - The cluster must provide continuous operation in the event of a hardware failure, network disruption, or planned maintenance at one site, without compromising data integrity or availability. - The system must be designed to support failover and failback between the two sites automatically and transparently. **Performance requirements:** 1. Minimum performance metrics: - the storage solution must deliver a combined performance of **minimum 500,000 Input/Output Operations Per Second (IOPS)** with inline data reduction (deduplication and compression). 2. Performance calculation parameters:		

				IOPS performance must be evaluated based on the following metrics: - read/write ratio: 70% read / 30% write. - block sizes: support for operations with block sizes of 16 KB, 32 KB, and 64 KB to accommodate varying workload requirements. - I/O patterns: include both sequential and random I/O workloads. - latency: ensure a maximum delay of 1 millisecond (0.001 s) under full load conditions. 3. Consistency of performance: - the system must maintain the required performance levels even under high concurrency and mixed workload conditions. - performance must remain unaffected during maintenance operations, including firmware updates, drive rebuilds, or component failures. 4. Performance verification: - vendors must provide detailed benchmark test results to validate the stated performance – **for operations with block sizes 16 KB(mandatory)**, 32 KB and 64 KB(optionall), using industry-standard tools such as IOmeter or FIO, under the specified conditions. - results must demonstrate compliance with all stated parameters, including latency and I/O patterns. 5. Monitoring and optimization: - the system must include tools to monitor and optimize performance dynamically, offering real-time insights into throughput, latency, and IOPS for proactive performance tuning. **Supported protocols:** - FC, - iSCSI; **Features:** Dedicated system management interfaces: 1. The system must include dedicated physical and/or virtual interfaces specifically for system management. 2. These interfaces should allow out-of-band management, ensuring that administrative tasks can be performed without impacting data traffic.		

				3. Management interfaces must support the following functionalities: - Web-based GUI for ease of access. - Command-line interface (CLI) for advanced configuration. - Support for industry-standard protocols such as SNMP and REST API for integration with monitoring and orchestration tools. - Role-based access control (RBAC) to ensure secure system administration. 4. Redundancy for management interfaces: - to ensure availability, the management interfaces must support redundancy, allowing continuous system management even in the event of a single interface failure. 5. Protocol optimization: The system must include protocol-specific optimizations such as: - Multipath I/O (MPIO) for FC and iSCSI to ensure high availability and load balancing. - Support for jumbo frames in iSCSI for improved performance in high-throughput environments. 6. Compliance and Interoperability: The system must be compliant with industry standards for both FC and iSCSI protocols. It must ensure interoperability with third-party devices, including servers, switches, and network adapters. **Deduplication and compression requirements:** 1. Functional capabilities: The storage system must provide deduplication functionality for data stored at the block level (iSCSI/FC LUN) and file level, with the following specifics: - Deduplication must operate both at the volume level and globally across the system, ensuring optimal storage efficiency. The system must also include compression functionality for: - Block-level volumes (iSCSI/FC LUN). 2. Interoperability and unrestricted functionality: Deduplication and compression features must operate seamlessly without introducing limitations or		

				restrictions on simultaneous use of other critical functionalities, including but not limited to: - Data replication. - Thin provisioning. - Backups. - Volume cloning. 3. Inline deduplication and compression: - Both deduplication and compression mechanisms must function in in-line mode, ensuring real-time data optimization without requiring post-processing. - Deduplication must remain continuously active and cannot be disabled or bypassed by system administrators or any other means, ensuring consistent storage efficiency and data integrity. **- Storage solutions that rely on scheduled or job-based data reduction processes are not acceptable.** 4. Licensing and support: All features related to deduplication and compression must be: - Fully licensed (if required by vendor provisions) and included in the offer, eliminating additional licensing costs for essential functionality. - Supported by the storage system in its maximum configuration, ensuring scalability and compatibility across all deployment scenarios. 5. Performance and reliability considerations: - The deduplication and compression mechanisms must not introduce significant latency or impact the system's performance metrics, such as IOPS or throughput. - Mechanisms should include built-in error detection and correction to maintain data integrity during deduplication and compression processes. 6. Management and monitoring: The system must provide a dedicated interface or tools for monitoring deduplication and compression efficiency, including: - Space savings metrics. - Real-time and historical performance impacts. - Detailed logs of deduplication and compression activities. **Snapshot requirements:** 1. General functionality:		

				- The system must support snapshot functionality at a minimum for block-level volumes (LUNs), ensuring operational flexibility. - The snapshot functionality must be applicable to both LUNs and other supported volumes without imposing restrictions on the simultaneous use of other critical system functions, including replication, backups, and cloning. 2. Snapshot quantity and retention: - The system must provide the ability to create and manage a minimum of 365 snapshots per shared volume, supporting long-term operational and recovery needs. - Snapshots must be configurable with retention policies to optimize storage space and align with data governance requirements. 3. Performance efficiency: - The implementation of snapshots must not degrade overall system performance, regardless of the number of active snapshots or system workload. - The system must include optimization mechanisms, such as metadata indexing and intelligent snapshot scheduling, to minimize latency and maintain high performance. 4. Space efficiency: - Snapshot functionality must employ a cost-effective approach by storing only the delta (changes) from the original data. This ensures minimal storage consumption while preserving full data access and recovery capabilities. 5. Integration with storage QoS: - The system must support performance monitoring and prioritization mechanisms for snapshots, enabling administrators to enforce Storage QoS (Quality of Service) policies at both the volume and LUN levels. - These QoS policies should dynamically allocate resources to prioritize performance-critical snapshots, ensuring minimal impact on other operations. 6. Advanced features: Snapshots must support: - Application-consistent snapshots, ensuring data integrity for workloads such as databases and virtualized environments.		

				- Writable snapshots, allowing clones to be created for development, testing, or analytics without affecting the production environment. Snapshots must be compatible with data replication workflows, ensuring consistent replication of both primary data and snapshot states across systems. **Encryption requirements:** 1. Encryption standard: - The solution must support encryption of all stored data using a minimum of the AES-256 algorithm or a stronger industry-standard encryption algorithm, ensuring compliance with modern security and regulatory standards. 2. Scope of encryption: - Encryption must be applied to all drives, NVMe, and flash storage within the device, covering the entire data storage ecosystem. - Encryption must extend to data at rest across all volumes, snapshots, backups, and metadata associated with the system. 3. Performance integrity: - Encryption functionality must operate with no measurable impact on system performance, ensuring IOPS, throughput, and latency metrics remain consistent with non-encrypted operations. - The system must leverage hardware-accelerated encryption or equivalent technologies to maintain optimal performance during data encryption and decryption processes. 4. Key management: - The solution must generate encryption keys using a secure hardware-based random number generator, ensuring keys are robust and resistant to attacks. - Encryption keys must be securely stored on the equipment, leveraging a dedicated hardware security module (HSM) or equivalent secure enclave to isolate keys from unauthorized access. - The system must ensure that data stored on drives/NVMe/flash cannot be accessed if the storage media is removed from the device or if the device itself is compromised. 5. Key backup and recovery:		

				- The system must include mechanisms for secure backup and recovery of encryption keys, supporting integration with external key management systems (KMS) compliant with KMIP (Key Management Interoperability Protocol) standards. - Key rotation and lifecycle management should be automated and configurable to align with organizational policies and compliance requirements. 6. Encryption for replication and snapshots: - The encryption functionality must extend to replicated data and snapshots, ensuring consistency in encryption across all replicated sites or volumes. - Encryption must not disrupt or degrade replication workflows, including synchronous and asynchronous modes. 7. Audit and compliance: - The system must provide audit logs and reports detailing encryption operations, key management activities, and access attempts, ensuring transparency and regulatory compliance. - Logs should be exportable and compatible with industry-standard security information and event management (SIEM) systems. **Monitoring requirements:** 1. Analytical platform or portal: - The system must include a robust analytical platform or virtual machine (VM) accessible via a web browser-based portal. - The platform must provide an intuitive, user-friendly interface with interactive dashboards for data visualization and management. 2. Log collection and reporting: The platform must automatically collect and analyze logs from the device and present them as customizable graphs, reports, and alerts, covering the following: 2.1. Storage utilization: - Real-time and historical monitoring of used space. - Display of the data reduction indicator, accounting for deduplication and compression (excluding thin provisioning, if applicable). - Granular visibility at both the global device level and the local LUN level.		

				2.2. Space growth prediction: - Advanced forecasting tools for predicting space growth, factoring in deduplication, compression, and provisioning trends. - Tools for future expansion analysis, including recommendations for scaling. 3. Component monitoring: The system must include an application or hardware-based monitoring solution to oversee and report detailed events for the following physical and logical components: - Physical components: controllers, drives, ports, power supplies, and network interfaces. - Logical components: volumes, LUNs, replication processes, deduplication, and compression algorithms. 4. Performance monitoring: The portal must provide minimum: - Real-time and historical performance metrics for individual resources. - Key parameters to monitor: Latency, Read and Write IOPS, Bandwidth. Performance data must be available at both the global system level and the LUN level. 5. Storage QoS and prioritization: - The system must include a performance monitoring and prioritization mechanism for Storage QoS, configurable at both the volume and LUN levels. - QoS metrics should be adjustable in real-time to meet dynamic workload demands. 6. Reporting and alerting: The portal must provide comprehensive reporting capabilities, including at least: - Capacity reports: current usage, available space, and forecasted capacity needs. - Performance reports: historical trends and real-time analytics of system performance. - Future space predictions: automated simulations for capacity increases based on application type and workload. - Event logs: authorization attempts, executed commands, and system alerts for security and operational events.		

				- Technical support logs: level of support received, resolution times, and incident history. 7. Operational monitoring: - Snapshot and replication status: display the real-time status of operations such as snapshots, synchronous/asynchronous replication, and recovery tasks. - Threat alerts: warnings related to system integrity, user activity, or misconfigurations. - Optimization insights: recommendations for system performance improvement, resource reallocation, or energy efficiency. 8. Configuration verification and upgrades: - The platform must include an algorithm for verifying configuration correctness and compatibility with potential device or cluster upgrades. 9. Simulation and optimization: - The platform must enable capacity simulation tools to project storage needs based on application types and expected workloads. - Display real-time system consumption metrics with actionable optimization guidelines for improving performance and efficiency. **NICs included per controller:** Min. x 1GE for management; Min. 2 x 32G FC SFP28(850nm SFP+ SR MM module included) for data transfer; Min. 2 x 32G FC dedicated for replication (metro cluster). **Supported operating environments:** Microsoft Windows Server; Red Hat Enterprise Linux; VMware (VMware ESXi); **Power supplies included:** The system must include a minimum of two (2) hot-swappable (hot-plug) power supply units (PSUs). The PSUs must support at least 1+1 redundancy, ensuring continuous operation in case of failure of one PSU.		

				Power cables included must meet the following specifications: - Type: IEC C13 - C14. - Minimum length: 0.6 meters (24 inches). **<u>Cerințe obligatorii pentru prestarea serviciilor de punere în funcțiune, a garanției și a serviciilor de suport (deservire și mentenanță) a bunurilor -</u>** <u>conform Anexei la Anunțul de participare**.**</u> Toate licențele necesare (dacă se aplică conform termenilor și condițiilor producătorului) pentru caracteristicile platformei/portalului de monitorizare (analitică) și software-ului/firmware-ului specific sistemului de stocare, inclusiv actualizările/patch-urile periodice, trebuie să fie incluse în ofertă și furnizate pe o bază perpetuă - valabile obligatoriu pentru durata integrală de viață a sistemului de stocare. **Termeni și condiții:** Toate cerințele sunt minime și obligatorii; O cerință nu trebuie să limiteze o altă cerință; Toate componentele trebuie să fie actuale și să nu fie promovate ca EOS (sfârșitul vânzării/suportului) / EOL (sfârșitul duratei de viață); Extinderea memoriei (ram) și a capacității de stocare nu trebuie să includă limitări hardware sau software.		

Numele, Prenumele: Bîrsan Vitalie, În calitate de: Administrator

Ofertantul: Î.C.S. RELIABLE SOLUTIONS DISTRIBUTOR S.R.L., Adresa: mun. Chișinău, Str. Alexandru cel Bun, 85

Digitally signed by Bîrsan Vitalie
Date: 2025.03.18 23:45:37 EET
Reason: MoldSign Signature
Location: Moldova
MOLDOVA EUROPEANĂ