FIVE

INFORMATION TECHNOLOGY

ENGINEERING

CHALLENGES WITH

SCALING

ENTERPRISE

SECURITY

WORKFLOWS

AND HOW TO SOLVE THEM

AVERE

Data, data everywhere.

Ready or not, here comes the security log data—tens of gigabytes of it every day from end-user desktops, mobile devices, routers, switches, firewalls, VPN appliances, virtual machines, web servers, sensors, Internet of Things (IoT) devices, and more. It's a non-stop firehose of data from tens of thousands of users and assets at every layer of the IT infrastructure.

Keeping up with that intake is just one of the challenges organizations confront in scaling enterprise security log and analytics workflows. Regardless of the application, be it a log manager from SolarWinds, Sumo Logic analytics, or a Splunk Enterprise cloud-based service, large-scale log management presents a common set of engineering problems across the workflow, from data intake and accrual to normalization, archiving, pattern analysis, and remediation. Let's take a closer look at five of the biggest scale-related problems and see how implementing an Avere Systems high-performance caching layer can solve them.

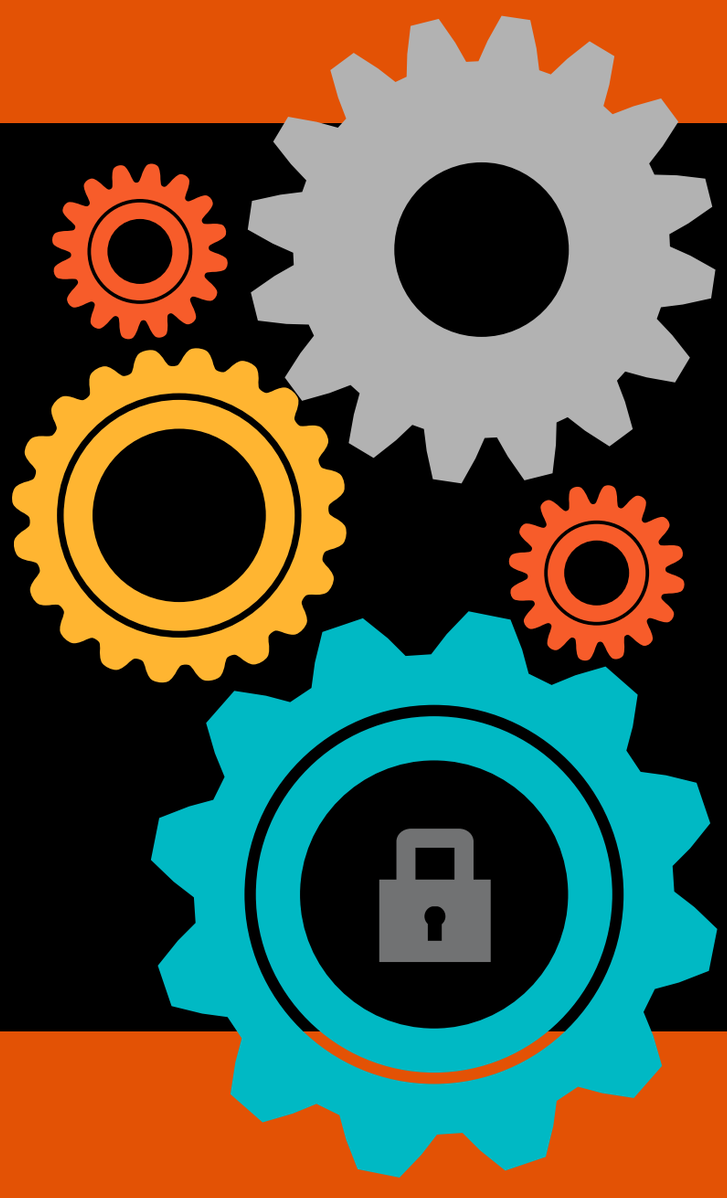

The Enterprise Security Workflow

- Acquire and aggregate inputs
- Normalize data
- Archive raw and/or normalized data
- Analyze for patterns
- Alert or remediate

1. INGEST LATENCY & THROUGHPUT

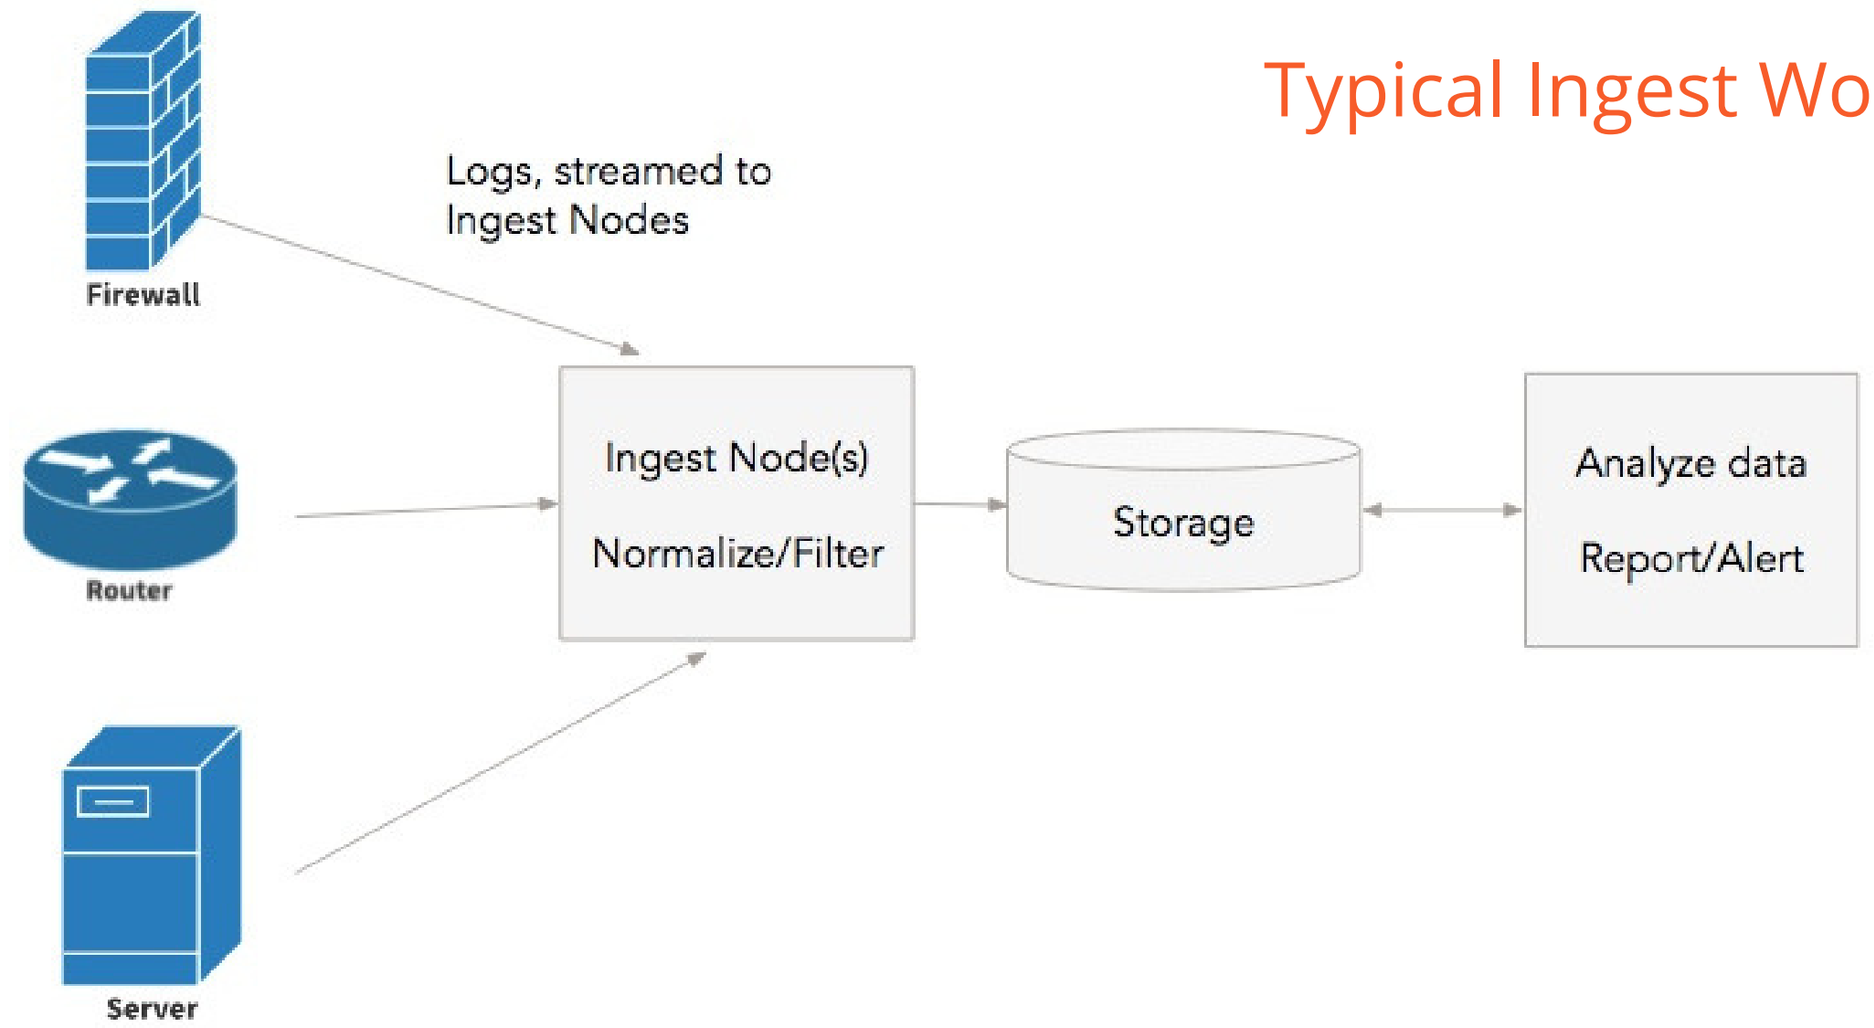

In a typical ingest workflow, intake nodes receive machine-generated log data directly or from system logging agents. Logs are queued up on the ingest node(s) for filtering and normalization. Once those processes are complete, the log data is sent to storage. Backend analytics and reporting servers will then access this stored log data to determine patterns and possible issues. The challenge is to avoid bottlenecks to and from storage as the infrastructure scales to ingest more and larger data sets. Even with high-speed network links, the immense amount of incoming log data can slow down storage and begin to back up the entire workflow. Back pressure can force I/O ingest throttling to avoid data loss.

On the other side, heavy reads from analysis nodes can further bog down storage. If the analytics servers are not co-colocated, latency will also impede analytics.

Ingest Latency Scales with Growth

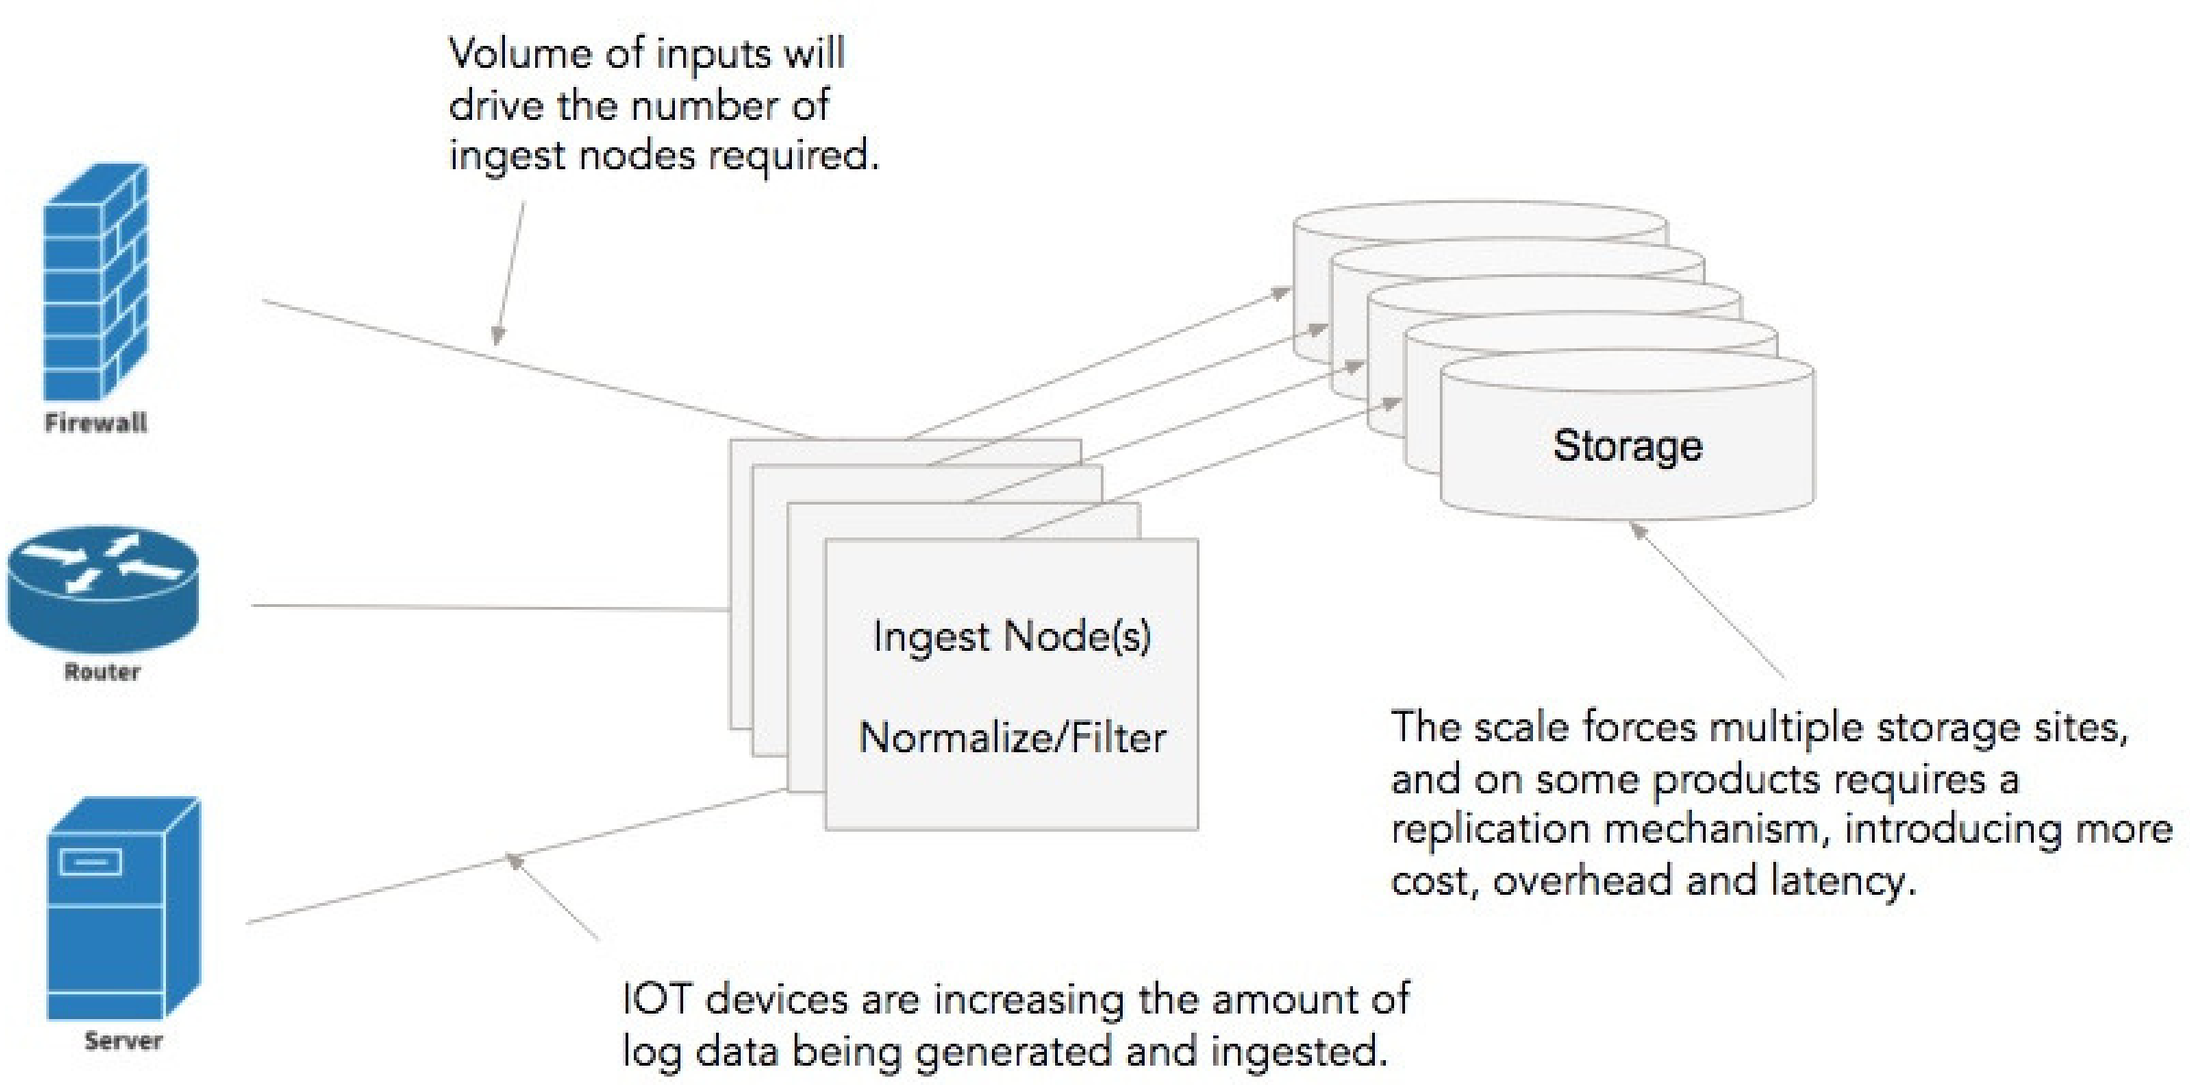

Who is Avere Systems?

Avere Systems was created by file systems experts determined to reinvent storage by changing the way enterprises thought about and bought storage resources. With decades of experience behind the company's founders, Avere got its start in 2008 with a mission to use fast, flash-based storage in the most efficient, effective manner possible.

What the team discovered was a technology that optimized storage resources and reduced dependencies on sprawling storage installations. Launched as the Avere OS, this advanced file system not only boosted performance within standard, on-premises, network-attached storage systems but also within diverse, hybrid architectures that include cloud provider services and other types of object storage.

AVERE

2. VENDOR LOCK-IN

Demand Locks Vendors Into Your Operations

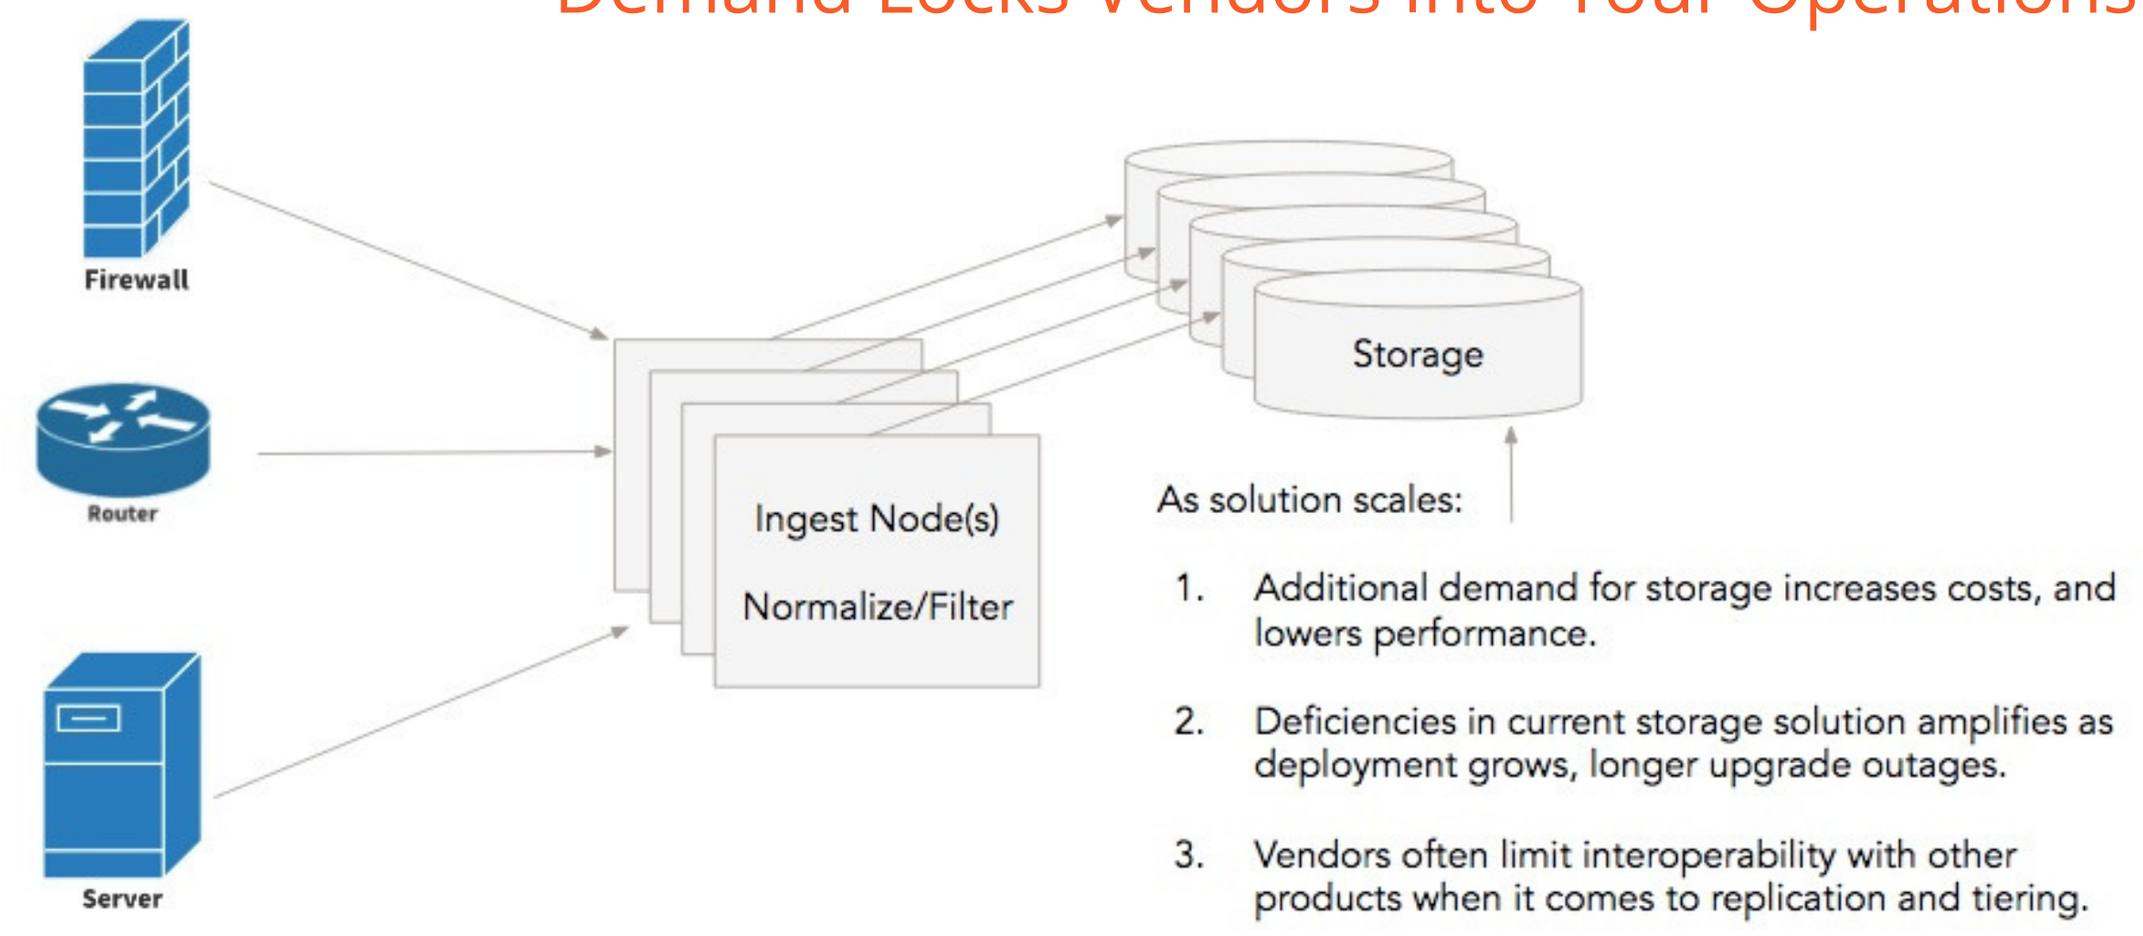

As the security infrastructure scales to take in and analyze more log data, scaling storage capacity and performance can become a costly, disruptive proposition. Few IT budgets can indefinitely absorb the costs of scaling or forklift-upgrading high-end NAS systems. Furthermore, taking an outage to make a change can dangerously disrupt the 24x7 security log infrastructure, impacting both protection levels and business continuity.

At some point, performance requirements may also push the upper bounds of available NAS infrastructure, obliging the enterprise to consider even more costly block-storage options. The prospect of changing architectures or vendors for any reason is daunting from multiple perspectives, including cost, the physics of data migration, and interoperability between old and new storage infrastructure.

3. LIFE CYCLE MANAGEMENT

Efficient Data Management Throughout the Lifecycle

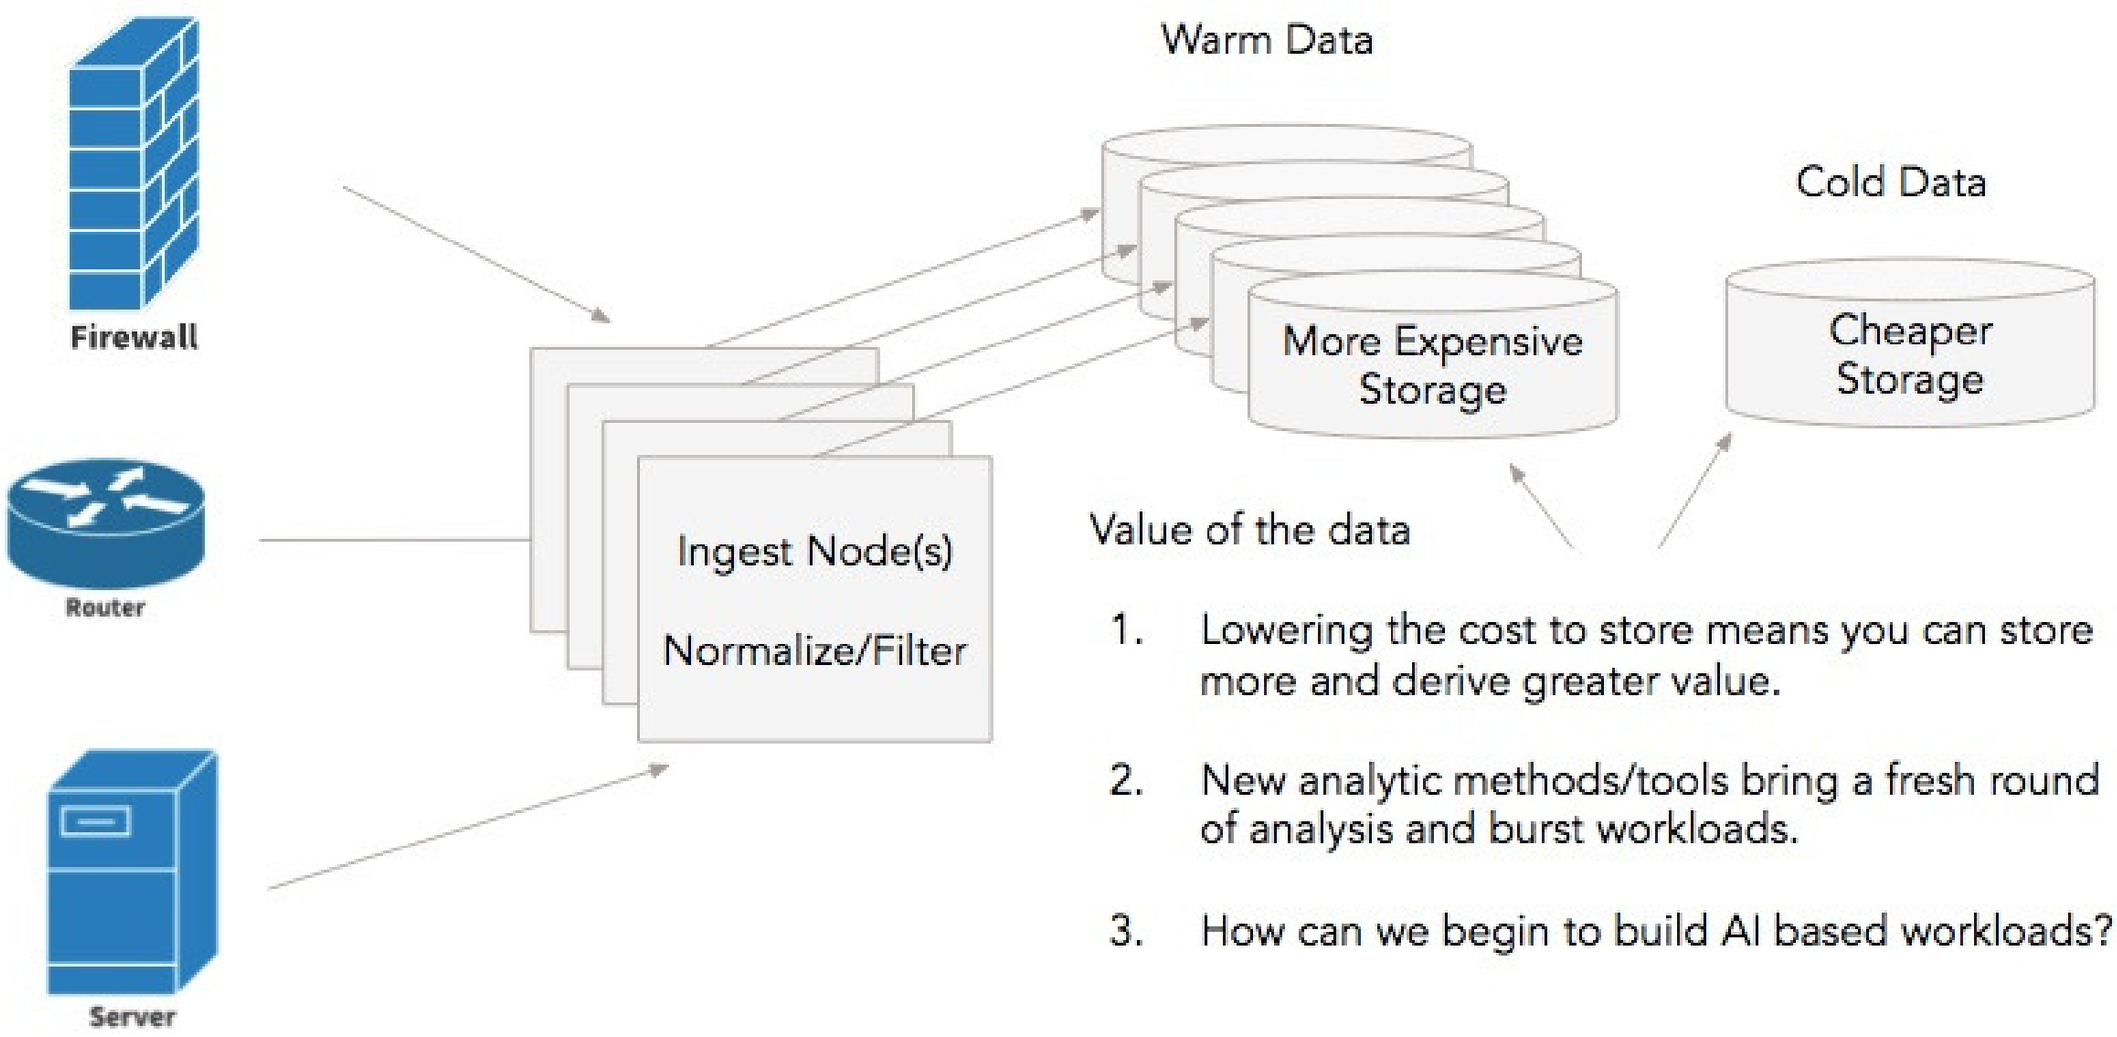

More data for longer periods of time—that's the objective of most enterprise security data-retention policies. But cost and inevitable data center changes such as platform upgrades, new applications, and physical moves can make it difficult to store, secure, and protect access to log data over time. Information security officers need to reduce the costs of petabyte-scale archive capacity and at the same time, despite underlying infrastructure changes, ensure accessibility for long-term forensic use of security log data.

4. Data Availability & Redundancy

Security officers also have a fiduciary responsibility to maintain the availability of security data, ensuring redundancy and protecting against any single point of failure in the security infrastructure. But as the infrastructure scales, so does the complexity of backup and recovery.

Traditional methods like mirroring can create new back pressure against ingest processes and rob IOPS from reporting and processing nodes. Multi-site mirrors add to the complexity, and latency issues can make replication of a golden version to a remote site unreliable. While many security-application vendors offer replication solutions, their toolsets typically exacerbate performance issues and detract from core application functions.

Data Availability and Redundance

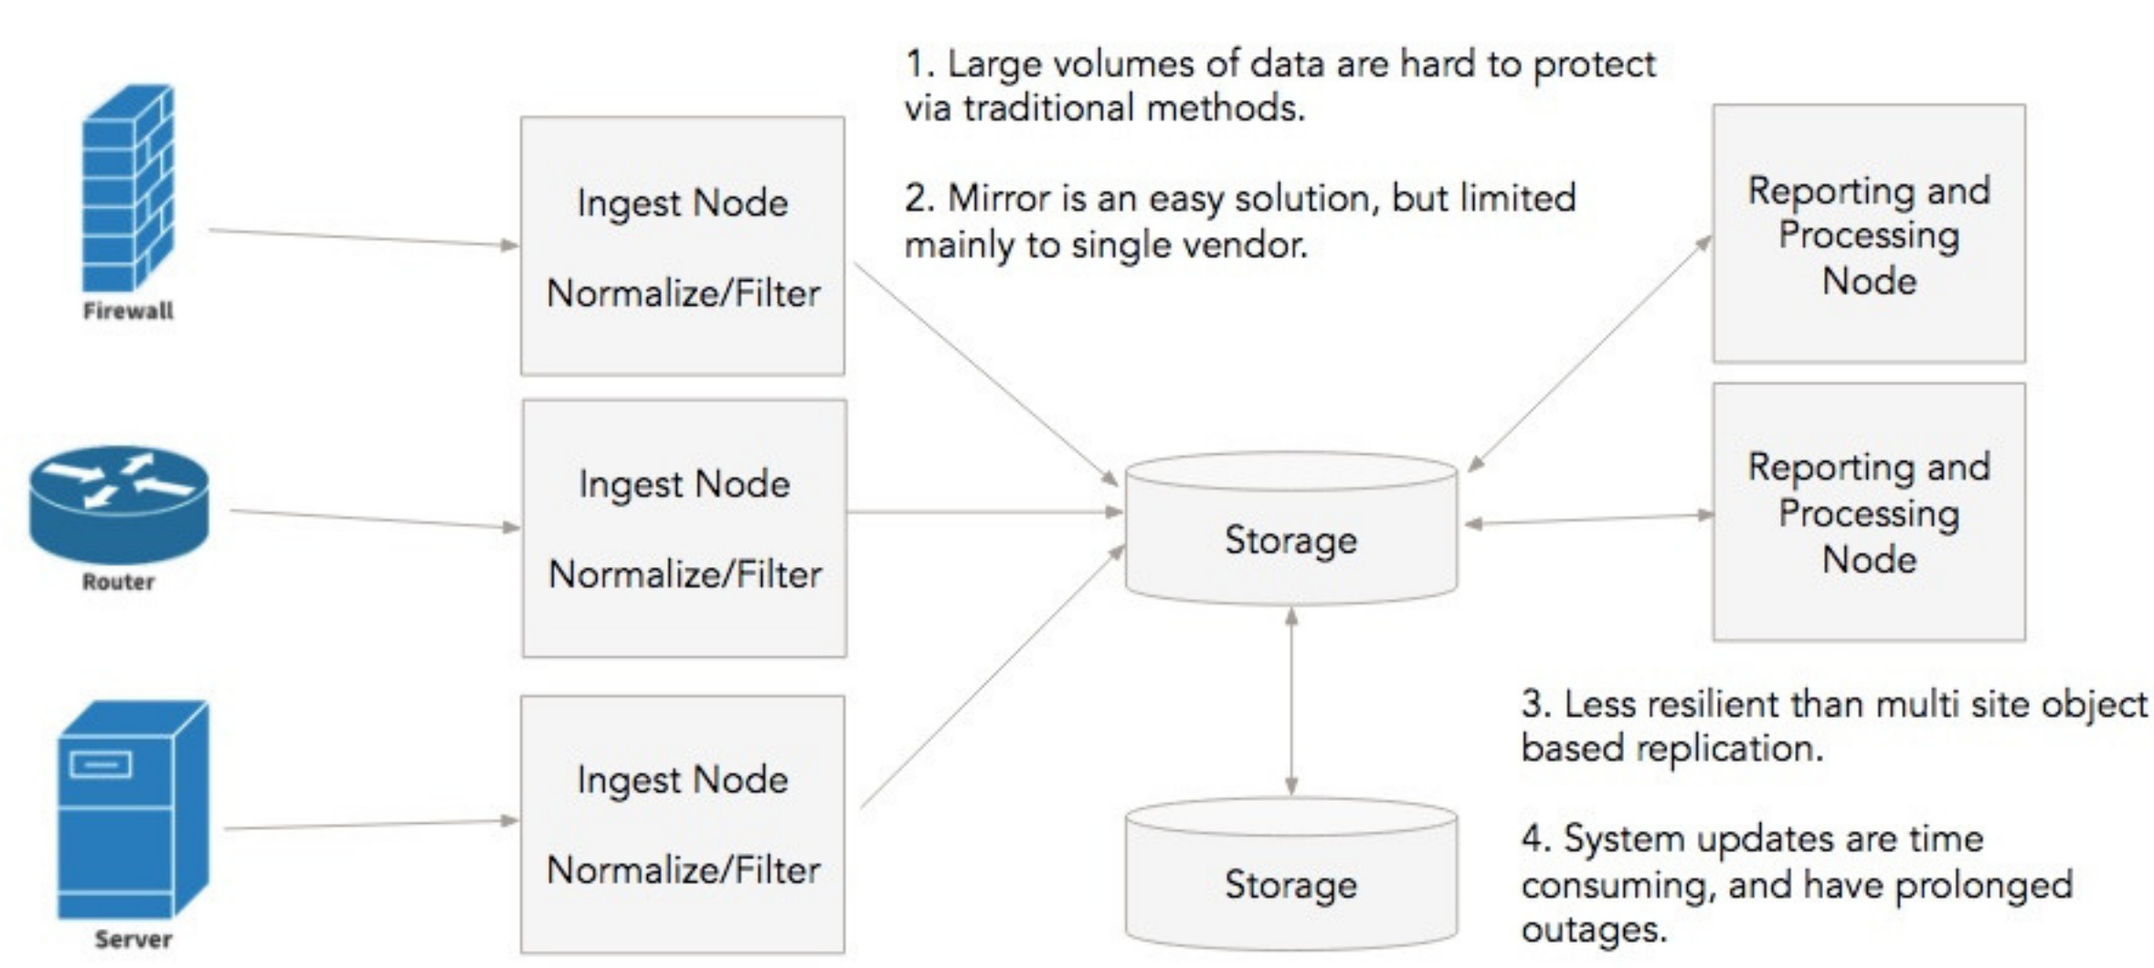

5. Cloud Integration

Cloud-based resources offer potential solutions to the problems of storage cost and redundancy. But the challenge comes in the technical adoption of public cloud, including its use as long-term data storage. Accessing cloud storage, for example, requires an object API, and sufficiently performant data access requires overcoming the inherently high latency between on-premises or cloud compute infrastructure and cloud storage.

Adding a Cloud Storage Tier

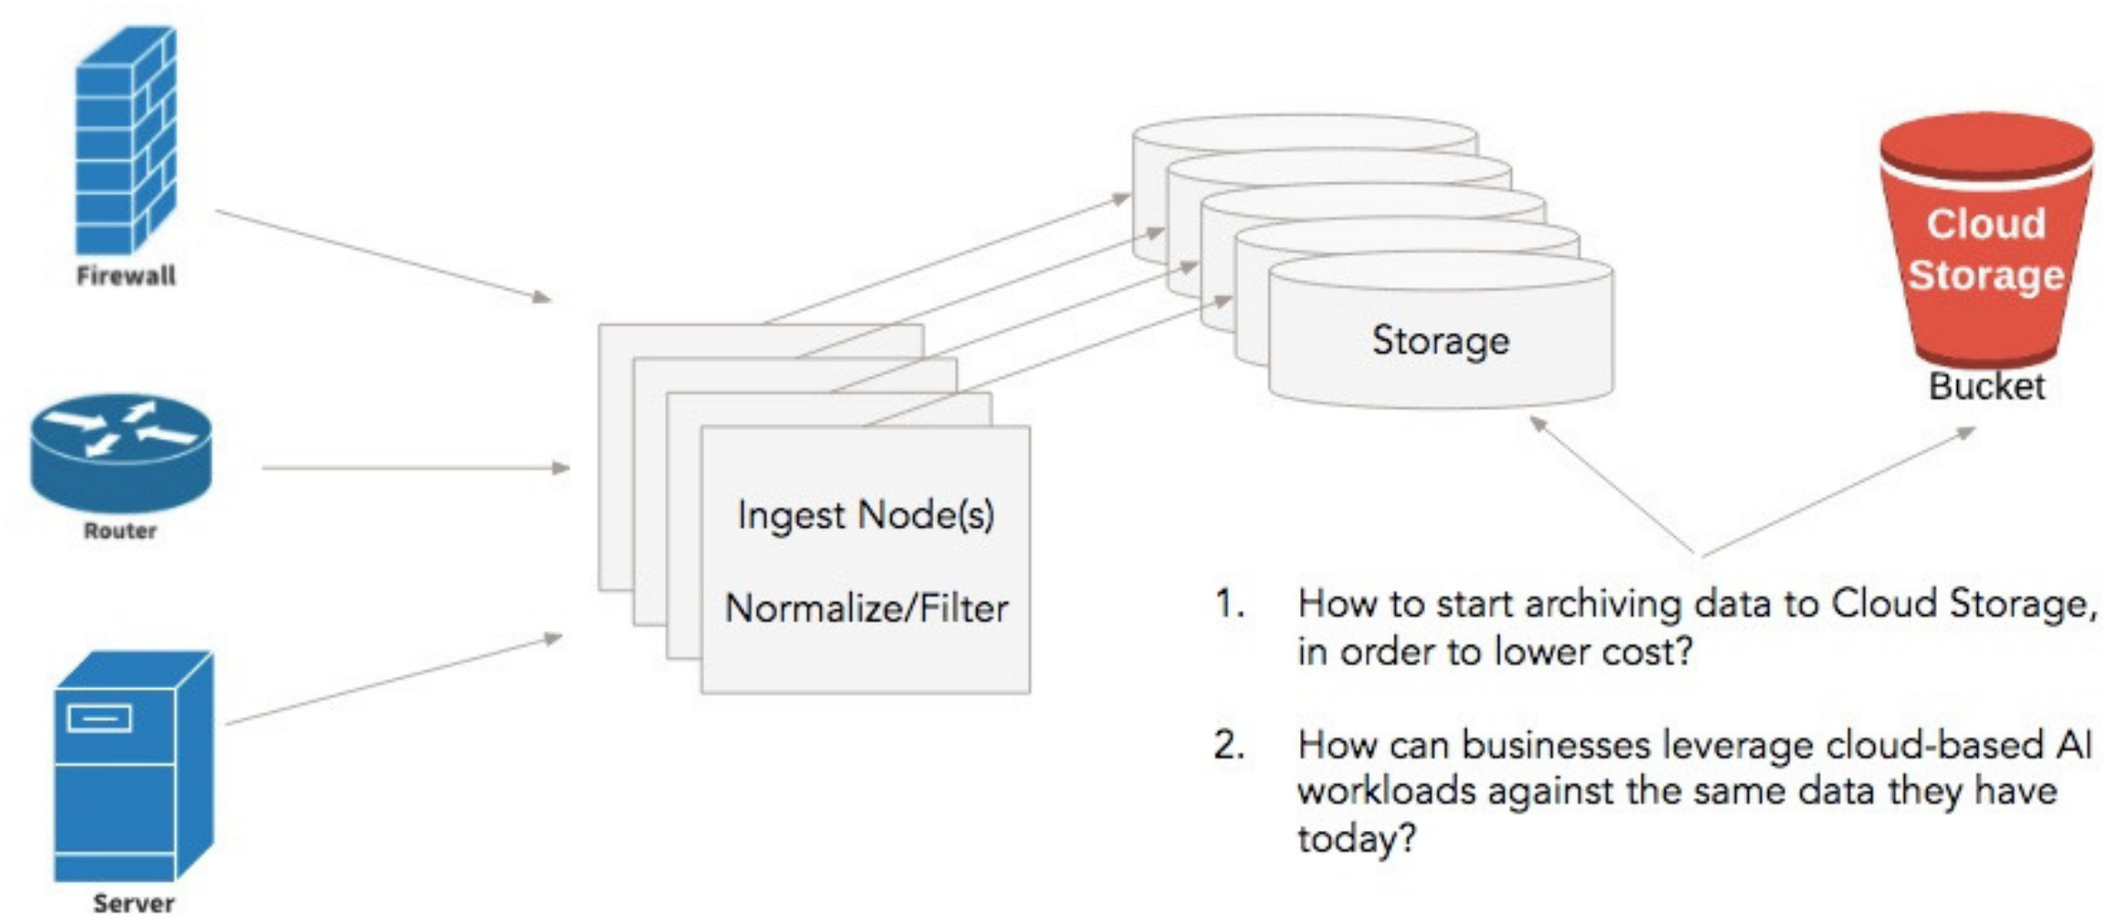

1. How to start archiving data to Cloud Storage, in order to lower cost?
2. How can businesses leverage cloud-based AI workloads against the same data they have today?

CACHE IT IF YOU CAN

Avere Systems solutions ideally suit both the scale and the workflow characteristics of log-management applications, effectively mitigating the most common problems of scaling enterprise security workflows. Avere technology enables drop-in, high-performance file system access for security log workflows. More affordable, less complex, and faster to implement than NAS upgrades or SAN build out, high-performance Avere FXT Edge Filers can better equip the enterprise infrastructure for exponential log-data growth.

Basic Caching Architecture

Avere solutions help protect against gaps in security data, effectively eliminating the problems of back pressure. A key feature of Avere's FXT Edge filer is write-back caching that allows caching of write operations to an Avere cluster that in turn smooths and gradually writes log data to the NAS back end, alleviating ingest latency, congestion, and bottlenecks. Avere FXT Edge filers also deliver more affordable throughput scaling.

Avere FXT Edge filers also deliver more affordable throughput scaling. Avere filers feature dynamic tiering that ensures active data is stored appropriately on the high-performance storage media (DRAM, SSD, and SAS) contained within the FXT Edge filer and that read, write, and metadata operations are handled with the lowest latency possible. Clustered FXT solutions enable performance scaling to millions of input/output operations per second (IOPS) and throughput in excess of hundreds of gigabytes/second. In addition, clustering scales the capacity of high-performance media contained within the FXT cluster to a capacity of total addressable storage beyond an exabyte.

Avere also prevents vendor lock-in and facilitates data lifecycle management by enabling consolidated user access through a single mount point that can include unified access to a wide variety of supported vendor options. The global namespace allows management of all log data—across public and private object storage and heterogeneous NAS—as a single logical pool.

Avere clusters ensure high availability of data. Writes are mirrored onto fast SSDs for data redundancy, and NVRAM cards help protect writes in the event of a power outage. Avere also makes it easier to move and mirror data without impacting data collection or analytics. Avere FlashMove software enables migration of live data, and Avere FlashMirror software replicates data for protection and recovery. The Avere solution lets companies implement needed storage and data center changes—consolidations, renovations, upgrades—without slowing down or stopping critical security applications and workflows.

The Power of High-Performance Caching

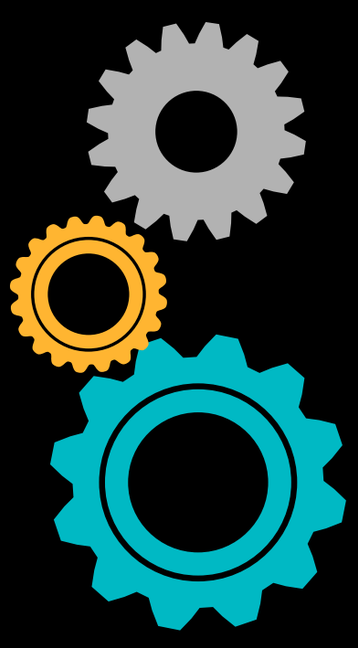

Speed Ingest with Write-back Caching

Gather writes (ack'ing clients immediately) and flushing in parallel

Hardware: NVRAM for write protection and caching

Clustered caching solution distributes writes across multiple nodes

Accelerate Read Performance with Distributed, Read-Ahead Caching

Read-ahead a request (read a bit more than what was requested)

Cache requests for other readers (typical in analytics)

Writes cached as written, speeding analysis workloads

Another benefit of the Avere technology is high-speed access to economical object-based archive storage on premises or in the public cloud. Avere hides latency to remote storage and translates NAS protocols to object APIs and back to NAS to facilitate use of services like Google Cloud Storage or Amazon Simple Storage Service (Amazon S3) from Amazon Web Services (AWS) for an active archive. Avere delivers the benefit of performant access to petabytes of cold log data and at the same time helps prevent lock-in to any particular archive tier or vendor.

Avere solutions can help enterprises reliably collect security log data, more quickly search and analyze data for improved intelligence and threat detection, and ultimately protect at-scale functionality to support expanding enterprise security workflows.

Avere solutions are helping Fortune 500 companies successfully tackle the challenges of scaling enterprise security workflows. Avere technology is used for ingest and analytics performance, vendor choice protection, and migration and mirroring of log data across data centers.

To learn more about how Avere technologies address common engineering challenges, watch the Scaling Security Workflows for Enterprise webinar at: https://www.averesystems.com/news-and-events/webinars

Or download the whitepaper: http://lp.averesystems.com/security-analytics-workflows-storage-performance

AVERE