RED HAT® STORAGE

WHITEPAPER

NETWORK-ATTACHED STORAGE IN THE PUBLIC CLOUD

Red Hat Storage for Amazon Web Services

INTRODUCTION

Cloud computing represents a major transformation in the way enterprises deliver a wide array of IT services. Indeed, faced with the agility and cost control demands of a competitive global business environment, IT organizations are increasingly making the move to a shared infrastructure using public cloud services, such as the Amazon Web Services (AWS) cloud.

While cloud computing offers opportunities to deploy IT infrastructure and applications in more flexible and efficient ways, the true potential of the cloud can only be tapped when both computing and storage are working in lockstep to support applications.

The use of object-based storage services has grown in popularity with the advent of the cloud, due to the ubiquity of HTTP as an access protocol and the resulting simplicity of building these services into newly developed applications. However, services offering only object-oriented access fail to deliver the compatibility required by enterprises seeking to deploy existing applications to a public cloud.

To effectively migrate existing applications and data, enterprises require an approach that is at once compatible with these applications and at the same time exhibits the desired characteristics of cloud-based systems. For these users, Red Hat® offers Red Hat Storage Server for Public Cloud, a no-compromise Network-Attached Storage (NAS)-in-the-Cloud software solution that offers both traditional and object-based file access backed by a proven, scalable, and cloud-friendly distributed storage system.

In this paper, after reviewing the key requirements for an enterprise-grade cloud storage system, we introduce Red Hat Storage Server and review the architectural characteristics that make it the storage system of choice for enterprise users. We conclude with a case study exploring the real-word experiences of Intelitek, a manufacturer of high-tech educational products that successfully deployed Red Hat Storage Server for Public Cloud within AWS to meet the time-sensitive demands of its customer base.

ENTERPRISE REQUIREMENTS FOR PUBLIC CLOUD STORAGE

With the rise of cloud computing, object-based storage services, characterized by the use of HTTP-based access to stored objects, have gained popularity. These services can make perfect sense for new applications and services built in the age of cloud computing, because they allow storage transactions to take place purely via HTTP, the lingua franca of the Internet.

However, while these object-oriented storage services may provide time-to-market, scalability, and cost-of-ownership benefits for greenfield applications, enterprises seeking to migrate existing application portfolios to the cloud require a solution with greater compatibility.

Within a given application portfolio, traditional applications may require access to unstructured data via a variety of storage protocols and command sets, including CIFS, NFS, FTP, and others. In addition, these applications typically expect to see data organized via paths and directories, and with standard permissions applied. Presenting storage via these traditional access methods gives enterprises the ability to easily migrate existing unstructured data to public clouds, without having to rewrite their applications.

Solutions able to offer NAS-like storage in the public cloud, natively supporting traditional access mechanisms just like on-premise NAS, also allow enterprises to more easily integrate cloud environments into existing storage management frameworks and processes.

While a storage solution that integrates with existing applications and processes is necessary for enterprises to readily adopt cloud computing, it is far from the only requirement. In fact, when storage is running in the cloud, it is critically important than the storage solution is architected in a way that is compatible with cloud infrastructure, and exhibits:

- **Elasticity** is the ability to add and remove resources easily and with minimal friction. Elasticity ultimately drives many of the benefits of cloud computing for enterprises, and is a fundamental property of public cloud infrastructures like the AWS Elastic Compute Cloud (EC2). The cloud storage solution must be similarly elastic to take full advantage of cloud infrastructure.
- **Scalability** is the property of a cloud storage system that translates elasticity into tangible benefits for the user. More specifically, a scalable cloud storage system can be made to exhibit greater capacity and performance simply by making additional compute and storage resources available to it.
- Because cloud infrastructure is based on inexpensive commodity machines, the responsibility for ensuring application and data availability is relegated to software. To support enterprise service level requirements, a cloud storage system must have features that ensure high levels of data availability in the face of infrastructure failure.

Fortunately for enterprises embracing cloud computing, Red Hat offers Red Hat Storage Server for Public Cloud, which we explore in the following sections.

INTRODUCING RED HAT STORAGE SERVER FOR PUBLIC CLOUD

Red Hat Storage Server for Public Cloud is a software-based solution for scale-out network-attached storage (NAS), delivered as an Amazon Machine Image (AMI). Integrating GlusterFS, a proven, open-source, distributed file system, and Red Hat Enterprise Linux into an easy-to-consume package for cloud deployment, Red Hat Storage Server for Public Cloud lets enterprises easily deploy NAS in the AWS cloud.

Red Hat Storage Server running on AWS aggregates Amazon Elastic Block Storage (EBS) and Amazon EC2 instances, creating a highly available, virtualized storage pool that can scale on demand to gigabits of throughput and petabytes of capacity.

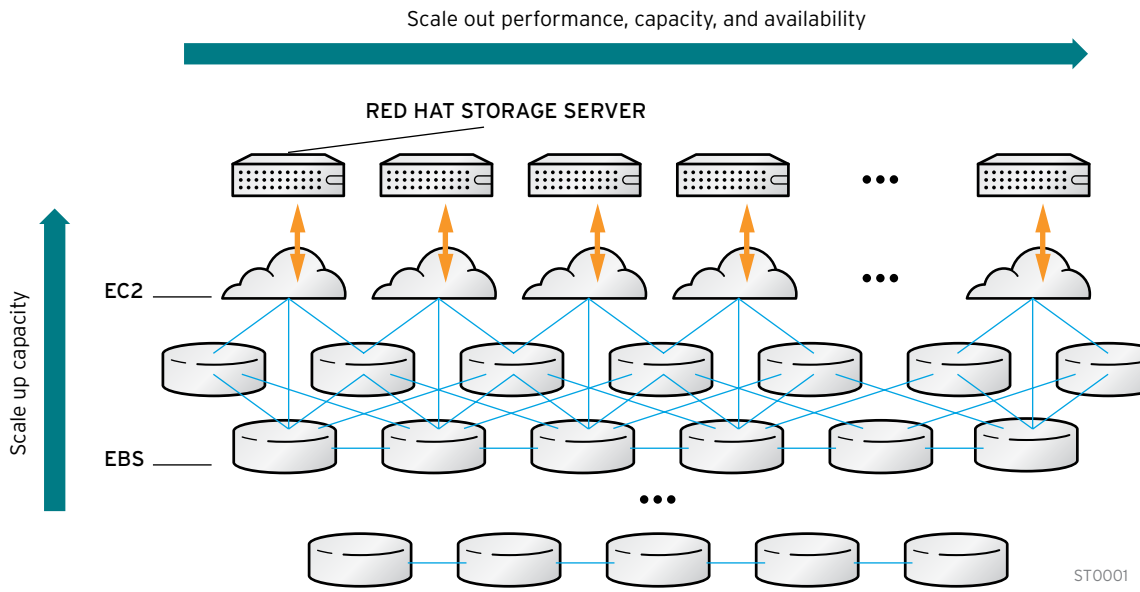

Red Hat Storage Server delivers a POSIX-compatible file system providing full support for industry-standard NAS protocols such as NFS and CIFS. As a result, enterprises can deploy existing unstructured data applications unmodified in AWS, accelerating their cloud adoption efforts and improving the overall return-on-investment (ROI) for their cloud computing projects.

In addition, Red Hat's Unified File and Object (UFO) technology allows users and applications to access the same data both as a file (via NFS or CIFS) and as an object (via HTTP), thus ensuring compatibility across applications, future-proofing the storage system, and simplifying management.

By aggregating large numbers of AWS server and storage resources into a unified storage pool, Red Hat Storage Server delivers elasticity and superior storage scalability in both the capacity and performance dimensions. Additional capacity for the storage pool can be easily provisioned by allocating additional EBS storage capacity to each Red Hat Storage Server instance. Similarly, as I/O performance requirements increase, additional Red Hat Storage Server instances running on discrete EC2 virtual machines can be non-disruptively added to the storage pool for increased throughput.

Red Hat Storage Server can ensure high levels of data availability with AWS, protecting against localized failures as well as the loss of entire AWS availability zones and regions. For complete and dynamic fault tolerance in face of local hardware or software failures, Red Hat Storage Server supports n-way synchronous replication, in which each node's file data is automatically replicated to two or more alternate nodes across availability zones. To protect against site failures, Red Hat Storage Server supports WAN-optimized long-distance asynchronous replication, in which server nodes replicate data across AWS regions for non-stop application availability.

By eliminating the performance, capacity, and availability limitations of the Amazon cloud's storage offerings and enhancing accessibility and manageability, Red Hat Storage Server for Public Cloud delivers a superior storage experience for AWS users while preserving the benefits of on-demand utility pricing.

RED HAT STORAGE SERVER: ARCHITECTED FOR CLOUD

The unique architecture of Red Hat Storage Server and its underlying GlusterFS file system enable it to deliver NAS storage that linearly scales-out in both performance and capacity, while meeting the resiliency requirements of the wide range of applications deployed in the cloud.

Built from the ground up for commodity computing environments like cloud, Red Hat Storage Server does not create, store, or use a separate metadata repository in any way. All nodes in the Red Hat Storage Server cluster have the intelligence to locate any piece of data without looking it up in an index or querying another server.

As a result, Red Hat Storage Server is fully distributed and decentralized, with no single point of failure. The system's advanced file management algorithms allow it to efficiently support multi-petabyte repositories. In addition, because compute nodes can access storage nodes directly, hot spots, choke points, and other I/O bottlenecks are eliminated and contention for data is reduced.

Red Hat Storage Server supports a global namespace that aggregates EBS disk and EC2 memory resources into a single pool with a common mount point, simplifying management of the cloud storage environment and eliminating data silos. Namespaces may be grown and shrunk dynamically, with no interruption to client access.

Offering storage features required for success in the cloud, and backed by a proven technology architecture, Red Hat Storage Server delivers outstanding value for enterprises like Intelitek, whose experience we briefly explore in the next section.

CASE STUDY: INTELITEK

Intelitek is a US-based, world-leading developer, producer and supplier of comprehensive learning solutions for training in engineering, automated manufacturing, mechatronics, industrial maintenance, and advanced manufacturing technologies. The company's broad product line includes CAD, CAM, CNC, robotics, machine vision, hydraulics and pneumatics, PLCs, sensors, process control and data acquisition. Intelitek also designs and produces automated manufacturing workcells for training, ranging from small-scale flexible manufacturing systems (FMS) to complete computer integrated manufacturing (CIM) systems.

A large number of educational institutions heavily rely on Intelitek to provide leading-edge technology to meet the rapidly changing need for learning solutions in various industries including engineering, automated manufacturing, mechatronics, industrial maintenance, and advanced manufacturing technologies. For one of its largest customers, Intelitek knew that utilizing the cloud was the best choice to ensure the customer's success, as it would easily deliver the flexibility and scalability required.

Intelitek found Amazon Web Services (AWS) to be an ideal cloud environment for its project, however in order to use the Amazon S3 object storage service, Intelitek would need to undertake a costly and time-consuming rewrite of parts of its existing POSIX-compatible applications.

As Intelitek learned about Red Hat Storage Server for Public Cloud and its ability to deliver NAS in the AWS cloud, it quickly became clear that pairing the two was the right solution for cloud enabling its current applications. Intelitek is now using Red Hat Storage Server on EC2 instances running in the AWS cloud. Red Hat Storage eliminated the need for application changes due to its POSIX compatibility and easy deployment in the Amazon cloud.

"We needed a storage solution that allowed us to keep our current API, while providing cloud services and capabilities to our customers," said Alon Kadury, vice president of research and development at Intelitek. "I immediately recognized the innovation of Red Hat Storage,

especially its ability to move POSIX applications seamlessly, and was amazed by the solution's scalability and high performance capabilities. Instead of months and months to get up and running, we were able to deliver the product to our customer quickly."

CONCLUSION

Red Hat Storage Server provides extremely scalable and highly available storage, with superior performance and seamless access to data, via both traditional networked protocols and HTTP-based, object-oriented methods. Packaged specifically for the Amazon Web Services cloud, Red Hat Storage Server for Public Cloud allows customers to accelerate their use of the cloud and easily migrate existing enterprise applications and data to the cloud environment.

Easy to deploy and manage across on-premise datacenters and public and private clouds, Red Hat Storage Server offers a single, unified storage experience for a wide array of enterprise applications. In delivering the flexibility and compatibility of NAS in the public cloud, Red Hat Storage Server helps enterprises unlock the promise of cloud computing, yielding enhanced agility and high ROI for new and existing applications.

SALES AND INQUIRIES

NORTH AMERICA
1-888-REDHAT1
www.redhat.com

EUROPE, MIDDLE EAST AND AFRICA
00800 7334 2835
www.europe.redhat.com
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
www.apac.redhat.com
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
latammktg@redhat.com

redhat