DARKRA White Paper
FEBRUARY 2016

Amazon Cloud Services

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for DARKRA's products remains at the sole discretion of DARKRA.

Contents

Background on this White Paper

This document is provided for informational purposes only. It represents DARKRA's Amazon Web Services offerings and related success stories across clientele and practices as of the date of issue of this document, which are subject to change without notice.

DARKRA is the leading independent Oracle Consulting & Oracle Solution provider.

Our services are customized to the best possible level ensuring absolutely no compromise with regard to deployment of skilled personnel, following of standards and implementation of industry best practices. Our services encompass all the aspects of Oracle Applications / Database software that includes Implementation, Development, Upkeep, Maintenance and Administration.
We help clients to rapidly reengineer themselves and be flexible enough to change with the current volatile environment. We guide our customers in the transition phases and present them with diverse growth opportunities for expanding into unique and uncharted territories.
Having established a sizeable offshore development center, with clients across Globe. It is our intention to broaden our horizon serving our clients.
Our dedicated and committed team of qualified techno-functional consultants is well versed with emerging technologies and enable us to harness the latest technologies for boosting business capabilities. Our talented pool of project management personnel have helped us thus so far overcome every challenge as we strive to reach greater heights.

Highlighters

- **Offshore Development Center (ODC) in** India and a registered office in UAE
- Presence in Australia via local partners
- Oracle Gold Partner, Red Hat Partner and Amazon Cloud Consulting partner

Key Achievement

- 20+ certifications for Oracle Database and Oracle E-Business Suite
- 10+ Oracle Identity and Access Management Consultants
- 50+ Oracle consultants
- Carried out multiple Oracle Identity & access management projects in UAE & Australia
- 20+ Oracle customers
- A team of engineers fully certified on Oracle systems such as Database, E-Business Suite, Red Hat Linux, Identity Management.

OVERVIEW

In 2006, Amazon Web Services (AWS) began offering IT infrastructure services to businesses in the form of web services—now commonly known as cloud computing. One of the key benefits of cloud computing is the opportunity to replace up-front capital infrastructure expenses with low variable costs that scale with your business. With the cloud, businesses no longer need to plan for and procure servers and other IT infrastructure weeks or months in advance. Instead, they can instantly spin up hundreds or thousands of servers in minutes and deliver results faster. Today, Amazon Web Services provides a highly reliable, scalable, low-cost infrastructure platform in the cloud that powers hundreds of thousands of businesses in 190 countries around the world. This white paper is an introduction to the AWS cloud-computing platform. It discusses the advantages of cloud computing and the fundamentals of AWS. It provides an overview of the AWS services that comprise the platform.

At DARKRA, we understand how to create AWS solutions that deliver highly efficient performance and scale.
Our deep knowledge of cloud technologies and UNIX/Linux can help our clients migrate their platform to AWS as seamless and painlessly as possible.

We provide AWS consulting in architecture design, deployment and managed services for business critical and large-scale systems, with strong focus on performance and scalability for large Internet sites, mission critical systems and Big Data platforms

What Is Cloud Computing?

Cloud computing is the on-demand delivery of IT resources and applications via the Internet with pay-as-you-go pricing. Whether you run applications that share photos to millions of mobile users or you support the critical operations of your business, the cloud provides rapid access to flexible and low-cost IT resources. With cloud computing, you do not need to make large upfront investments in hardware and spend a lot of time managing that hardware. Instead, you can provision exactly the right type and size of computing resources you need to power your newest bright idea or operate your IT department. With cloud computing, you can access as many resources as you need, almost instantly, and only pay for what you use. Cloud computing provides a simple way to access servers, storage, databases, and a broad set of application services over the Internet. Cloud computing providers such as AWS own and maintain the network-connected hardware required for these application services, while you provision and use what you need using a web application.

Six Advantages of Cloud Computing

- Trade capital expense for variable expense

Instead of having to invest heavily in data centers and servers before you know how you are going to use them, you can pay only when you consume computing resources, and pay only for how much you consume.

- Benefit from massive economies of scale

By using cloud computing, you can achieve a lower variable cost than you can get on your own. Because usage from hundreds of thousands of customers is aggregated in the cloud, providers such as AWS can achieve higher economies of scale, which translates into lower pay as-you-go prices.

- Stop guessing about capacity

Eliminate guessing on your infrastructure capacity needs. When you make a capacity decision prior to deploying an application, you often end up either sitting on expensive idle resources or dealing with limited capacity. With cloud computing, these problems go away. You can access as much or as little capacity as you need, and scale up and down as required with only a few minutes' notice.

- Increase speed and agility

In a cloud-computing environment, new IT resources are only a click away, which means that you reduce the time to make those resources available to your developers from weeks to just minutes. This result in a dramatic increase in agility for the organization, since the cost and time it takes to experiment and develop is significantly lower.

- Stop spending money running and maintaining data centers

Focus on projects that differentiate your business, not the infrastructure. Cloud computing lets you focus on your own customers, rather than on the heavy lifting of racking, stacking, and powering servers.

- Go global in minutes

Easily deploy your application in multiple regions around the world with just a few clicks. This means you can provide lower latency and a better experience for your customers at minimal cost.

DARKRA Expertise on Amazon Web Services Cloud Platform

AWS consists of many cloud services that you can use in combinations tailored to your business or organizational needs. This section introduces the major AWS services by category as per DARKRAs key offerings across clientele.

DARKRA Expertise Includes:

- **Architecture Design:** High availability engineering, infrastructure design, scalability engineering, infrastructure audit, best practices implementation, security audit and remediation.
- **Integration and Migration:** Application and infrastructure migration, performance optimization, load testing, capacity planning, big data, disaster recovery and backups.
- **Automation and Management:** Configuration management, cloud management, continuous integration and Chef.

Figure 1. Broad & Deep Core Cloud Infrastructure Services

Compute

Amazon EC2
Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances (called Amazon EC2 instances) to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers and system administrators the tools to build failure resilient applications and isolate themselves from common failure scenarios.

Benefits

Elastic Web-Scale Computing
Amazon EC2 enables you to increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds or even thousands of server instances simultaneously. Of course, because this is all controlled with web service APIs, your application can automatically scale itself up and down depending on its needs.

Completely Controlled
You have complete control of your Amazon EC2 instances. You have root access to each one, and you can interact with them as you would any machine. You can stop your Amazon EC2 instance while retaining the data on your boot partition, and then subsequently restart the same instance using web service APIs. Instances can be rebooted remotely using web service APIs. Instances can be rebooted remotely using web service APIs. You also have access to console output of your instances.

Flexible Cloud Hosting Services
You can choose among multiple instance types, operating systems, and software packages. Amazon EC2 allows you to select the memory configuration, CPU, instance storage, and boot partition size that are optimal for the operating system and applications that you choose. For example, your choice of operating systems includes numerous Linux distributions and Microsoft Windows Server10.

Designed for use with other Amazon Web Services
Amazon EC2 works in conjunction with Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), Amazon DynamoDB, and Amazon Simple Queue Service (Amazon SQS) to provide a complete solution for computing, query processing, and storage across a wide range of applications.

Reliable
Amazon EC2 offers a highly reliable environment where replacement instances can be rapidly and predictably commissioned. The service runs within Amazon's proven network infrastructure and data centers. The Amazon EC2 service level agreement (SLA) commitment is 99.95% availability for each Amazon EC2 region.

Secure
Amazon EC2 works in conjunction with Amazon Virtual Private Cloud (Amazon VPC) to provide security and robust networking functionality for your compute resources.

- Your compute instances are located in a VPC with an IP address range that you specify. You decide which instances are exposed to the Internet and which remain private.
- Security groups and network access control lists (ACLs) allow you to control inbound and outbound network access to and from your instances.
- You can connect your existing IT infrastructure to resources in your VPC using industry-standard encrypted IPsec virtual private network (VPN) connections.
- You can provision your Amazon EC2 resources as Dedicated Instances. Dedicated Instances are Amazon EC2 instances that run on hardware dedicated to a single customer for additional isolation.

Inexpensive
Amazon EC2 passes on to you the financial benefits of Amazon's scale. You pay a very low rate for the compute capacity you actually consume.

- **On-Demand Instances**—On-Demand instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into

much smaller variable costs. On-Demand Instances also remove the need to buy "safety net" capacity to handle periodic traffic spikes.

- **Reserved Instances**—Reserved Instances allow you to reserve Amazon EC2 computing capacity for 1 or 3 years, in exchange for a significantly discounted hourly rate (up to 75%), compared to On-Demand Instance pricing. You can use the Reserved Instance Marketplace to sell Reserved Instances if your needs change. (For example, you might want to move instances to a new AWS region, change to a new instance type, or sell capacity for projects that end before your Reserved Instance term expires).

- Spot Instances—Spot Instances allow you to bid on unused Amazon EC2 capacity and run those instances for as long as your bid exceeds the current Spot Price. The Spot Price changes periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. If you can be flexible about when your applications need to run, Spot Instances can significantly lower your Amazon EC2 costs.

Storage and Content Delivery

Amazon S3
Amazon Simple Storage Service (Amazon S3) provides developers and IT teams with safe, secure, and highly scalable object storage. Amazon S3 is easy to use. It has a simple web service interface for storage and retrieval of any amount of data from anywhere on the web. Amazon S3 offers a range of storage classes designed for different use cases, including Amazon S3 Standard for general-purpose storage of frequently accessed data, Amazon S3 Standard - Infrequent Access (Standard - IA) for long-lived, but less frequently accessed data and Amazon Glacier for long-term archiving. Amazon S3 also offers configurable lifecycle policies for managing your data throughout its lifecycle. Once a policy is set, your data will automatically migrate to the most appropriate storage class without any changes to your applications.

Amazon Glacier
Amazon Glacier is a secure, durable, and extremely low-cost storage service for data archiving and long-term backup. Customers can reliably store large or small amounts of data for as little as $0.007 per gigabyte per month, a significant savings compared to on-premises solutions. To keep costs low, Amazon Glacier is optimized for infrequently accessed data where a retrieval time of several hours is suitable.

Amazon Elastic Block Store
Amazon Elastic Block Store (Amazon EBS) provides persistent block-level storage volumes for use with Amazon EC2 instances in the AWS cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes—all while paying a low price for only what you provision.

Amazon Elastic File System
Amazon Elastic File System (Amazon EFS) is a shared file storage service for Amazon EC2 instances. Amazon EFS is easy to use and provides a simple interface that allows you to create and configure file systems quickly and easily. With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as you add and remove files, so your applications have the storage they need, when they need it.

AWS Storage Gateway
AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and the AWS storage infrastructure. The service allows you to securely store data in the

AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industrystandard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all of your data encrypted in Amazon S3 or Amazon Glacier.

Amazon CloudFront
Amazon CloudFront is a content delivery web service. It integrates with other services in AWS to give developers and businesses an easy way to distribute content to end users with low latency, high data transfer speeds, and no minimum usage commitments. Amazon CloudFront can be used to deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations. Requests for your content are automatically routed to the nearest edge location, so content is delivered with the best possible performance. Amazon CloudFront is optimized to work with other services in AWS, such as Amazon S3, Amazon EC2, Elastic Load Balancing, and Amazon Route. Amazon CloudFront also works seamlessly with any non-AWS origin server that stores the original, definitive versions of your files. Like other services in AWS, there are no long-term contracts or minimum monthly usage commitments for using Amazon CloudFront—you pay only for as much or as little content as you actually deliver.

AWS Import/Export Snowball
AWS Import/Export Snowball is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of AWS. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, secure, and can be as little as onefifth the cost of high-speed Internet.

Database

Amazon RDS
Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale relational databases in the cloud. Amazon RDS frees you up to focus on your application by managing time-consuming database administration tasks including backups, software patching, monitoring, scaling and replication. It is a fully managed cost-efficient service that allows you to launch secure, highly available, fault tolerant production ready databases in minutes. You can scale your database's compute and storage resources with only a few clicks or an API call, often with no downtime. Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora35, Oracle36, Microsoft SQL Server37, PostgreSQL38, MySQL39, and MariaDB40. For commercial database engines like Oracle and SQL Server, you can bring your existing licenses or pay for the license as part of the service with the “License Included” option.

Amazon Aurora
Amazon Aurora is a MySQL-compatible, relational database engine that combines the reliability, speed and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon Aurora provides up to five times better performance than MySQL at a price point that is one-tenth of a commercial database while delivering similar performance and availability. Amazon Aurora is available as a managed database service through Amazon RDS.

AWS Database Migration Service
AWS Database Migration Service helps you migrate databases to AWS easily and securely while the source database remains operational, minimizing downtime to applications that rely on the database. It can migrate your data to and from all widely used commercial and open-source databases and it supports homogenous migrations such as Oracle-to-Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora.

Amazon DynamoDB
Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed database and supports both document and key-value data models. Its flexible data model and reliable performance make it a great fit for mobile, web, gaming, ad-tech, Internet of Things, and many other applications.
The downloadable version of DynamoDB44 makes it easy to develop and test applications on your laptop or in an EC2 instance. Once your solution is ready, easily scale your application on the cloud with DynamoDB.

Amazon Redshift
Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small with no commitments or upfront costs and scale to petabytes for as little as $1000/terabyte/year. Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Amazon Redshift has custom JDBC and ODBC drivers that you can download from the AWS Management Console, allowing you to use a wide range of familiar SQL clients. You can also use standard PostgreSQL JDBC, and ODBC drivers. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis, or any SSH-enabled host.

Amazon ElastiCache
Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in memory caches, instead of relying entirely on slower disk-based databases. ElastiCache supports two open-source in-memory caching engines:

- Memcached, A widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service. Redis, a popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports master/slave replication and Multi-AZ, which can be used to achieve cross-AZ redundancy.

Networking

Amazon VPC
Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the AWS cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can easily customize the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your web servers that has access to the Internet, and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security (including security groups and network access control lists) to help control access to Amazon EC2 instances in each subnet. Additionally, you can create a hardware virtual private network (VPN) connection between your corporate data center and your Amazon VPC and leverage the AWS cloud as an extension of your corporate data center.

AWS Direct Connect
AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or co-location environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. AWS Direct Connect lets you establish a dedicated network connection between your

network and one of the AWS Direct Connect locations. Using industry standard 802.1Q virtual LANS (VLANs), this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources, such as objects stored in Amazon S3, using public IP address space, and private resources, such as Amazon EC2 instances running within an Amazon VPC, using private IP address space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

Amazon Route
Amazon Route is a highly available and scalable Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to Internet applications by translating human readable names, such as www.example.com, into the numeric IP addresses, such as 192.0.2.1, that computers use to connect to each other. Amazon Route effectively connects user requests to infrastructure running in AWS—such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets—and can also be used to route users to infrastructure outside of AWS. You can use Amazon Route to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of your application and its endpoints. Amazon Route makes it possible for you to manage traffic globally through a variety of routing types, including Latency Based Routing, Geo DNS, and Weighted Round Robin—all of which can be combined with DNS Failover in order to enable a variety of low-latency, fault tolerant architectures. Amazon Route also offers domain name registration—you can purchase and manage domain names, such as example.com, and Amazon Route will automatically configure DNS settings for your domains.

Security and Identity

AWS Identity and Access Management
AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. IAM allows you to do the following:

- Manage IAM users64 and their access—You can create users in IAM, assign them individual security credentials (access keys, passwords, and multi-factor authentication devices) or request temporary security credentials to provide users access to AWS services and resources. You can manage permissions in order to control which operations a user can perform.
- Manage IAM roles and their permissions—You can create roles in IAM, and manage permissions to control which operations can be performed by the entity, or AWS service, that assumes the role. You can also define which entity is allowed to assume the role.
- Manage federated users and their permissions—You can enable identity federation to allow existing identities (e.g., users) in your enterprise to access the AWS Management Console, to call AWS APIs, and to access resources, without the need to create an IAM user for each identity.

AWS Key Management Service
AWS Key Management Service (KMS) is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data, and uses Hardware Security Modules (HSMs) to protect the security of your keys. AWS Key Management Service is integrated with other AWS services including Amazon EBS, Amazon S3, and Amazon Redshift. AWS Key Management Service is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

AWS Directory Service
AWS Directory Service is a managed service that allows you to connect your AWS resources with an existing on-premises Microsoft Active Directory or to set up a new, standalone directory in the AWS cloud. Connecting to an on-premises directory is easy. Once this connection is established, all users can access AWS resources and applications with their existing corporate credentials. You can also launch managed, Samba-based directories in a matter of minutes, simplifying the deployment and management of Linux and Microsoft Windows workloads in the AWS cloud.

Amazon Inspector
Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed report with prioritized steps for remediation. To help you get started quickly, Amazon Inspector includes a knowledge base of hundreds of rules mapped to common security compliance standards (e.g., PCI DSS) and vulnerability definitions. Examples of built-in rules include checking for remote root login being enabled, or vulnerable software versions installed. These rules are regularly updated by AWS security researchers.

AWS WAF
AWS WAF is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow or block to your web application by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for your specific application. New rules can be deployed within minutes, letting you respond quickly to changing traffic patterns. Also, AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of web security rules.

AWS CloudHSM
The AWS CloudHSM service helps you meet corporate, contractual, and regulatory compliance requirements for data security by using dedicated Hardware Security Module (HSM) appliances within the AWS cloud. The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance. AWS CloudHSMs are provisioned inside your Amazon VPC with an IP address that you specify, providing simple and private network connectivity to your Amazon EC2 instances. AWS provides dedicated and exclusive access to AWS CloudHSMs, isolated from other AWS customers. Available in multiple regions and Availability Zones, AWS CloudHSM allows you to add secure and durable key storage to your Amazon EC2 applications.

Success Story

Changing consumer lifestyles and the explosion of web-connected devices such as tablets, mobiles, consoles, and Internet-enabled televisions led our client to develop digital and video-on-demand (VOD) services so its viewers can consume content in new and diverse ways.

The Challenge

The business success of these services created new challenges for Channel MBC: the volume of largely unstructured web data flowing from the VOD platforms grew faster than its existing on-premises environment could manage.
With its business and brand missions to connect more personally with viewers and more closely with advertisers, Channel MBC realized it needed smarter, more agile decision-making in the commissioning, scheduling and monetizing of content.

Why Amazon Web Services

To realize its business objectives, our client needed a proven high-performance data-analytics solution—a flexible, integrated service capable of capturing, storing, indexing, searching, mapping, analyzing and matching high volumes of viewer and platform data on demand.
When our client investigated the feasibility of using its existing data-analytics software to perform key tasks, the company found that it would take eight months to deliver the relevant base data for analysts to access. In contrast, by using Amazon Web Services (AWS) and Amazon Elastic MapReduce (Amazon EMR), our client could begin analyzing and modeling the data in just two and one-half days.
Amazon EMR runs across a cluster of Amazon Elastic Compute Cloud (Amazon EC2) instances that feature Intel Xeon E5 family processors. The combination of AWS and Intel together provided our client with the availability, flexibility, and virtually unlimited capacity of a cloud-based solution, with the proven power to process huge volumes of data.
Our Client also developed a Big Data control panel (BDCP), a web-based interface to allow analysts to spin up and spin down Amazon EMR clusters, submit queries through Apache Hive and Pig, monitor job status, and extract sample or actual data to run in modeling applications. BDCP puts convenience, control, power, and speed right in the hands of the people who need it, while providing a simple solution to a complex business problem.

KEY RESULTS Achieved by DARKRA

- Using the AWS Cloud allowed our client to avoid the need to invest in massive server infrastructure, disks, and CPUs and bring its registered VOD viewers closer to both content and advertisers, with benefits for both parties.

- By analyzing repeat viewers, the broadcast company can better understand their preferences and behavior, schedule more relevant content, commission enticing new content, and connect more closely with viewers.

- ✓ Its modeling systems can even predict the socioeconomic class, age, and gender of anonymous or unregistered viewers from content viewed, giving advertiser's confidence in optimizing media placements.

- Our Client analysts can instantly provision the capacity they need from Amazon EMR to perform complex predictive modeling and other business-intelligence tasks with maximum speed, cost-efficiency, and return on investment.

- ✓ Our Client can have a direct relationship with its nine million registered viewers for the first time, even modeling in-session data to deliver relevant ads before a program ends—an enviable advantage for any broadcaster. "In 60 minutes, our client can analyze and model in-session data to deliver highly targeted ads to viewers-before a program ends.

- Using AWS has helped our client build ad revenues, enhance its marketing of content, and use detailed, dynamic data to open up exciting new business opportunities.

CONCLUSION

- AWS provides building blocks that you can assemble quickly to support virtually any workload.
- With AWS, you will find a complete set of highly available services that are designed to work together to build sophisticated scalable applications.
- You have access to highly durable storage, low-cost compute, high-performance databases, and management tools. All this is available without up-front cost, and you pay for only what you use.
  .


ORACLE
HUTCHISON GLOBAL SERVICES
Hexaware TECHNOLOGIES
Accionlabs
HCL
XTglobal
exemplifying excellence
CSC
REDTAG
mbc
Galana
our energy at your service
AIRBUS
KHALIFA UNIVERSITY
OASIS investment
عبد اللطيف جميل
Abdul Latif Jameel
TADMUR
HOLDING
AL HARAMAIN
THE ABRAAJ GROUP
beIN SPORTS
NIPPON MARKETING
مستشفى الزهراء دبي
AL ZAHRA HOSPITAL DUBAI