CDE DRDPtech Security

DRDPtech is operated and hosted by the University of California Berkeley, Berkeley Evaluation and Assessment Research Center for the California Department of Education. Extensive security measures have been put into place to meet State Department of Education regulations. These measures are described below.

Both Windows and Linux Servers

Ossec is installed on all our servers. Ossec is a host-based intrusion detection system (IDS). It is configured to perform the following actions:

Log monitoring and analysis: Collect, analyze and correlate application, device and network logs generated by the operating system. Some examples of log monitoring include, but are not limited to: new application installation on the system, firewall rule is changed, windows events, system reboot, IIS events.

Integrity checking: File integrity monitoring is performed on all critical systems to alert the system administrator about any tampering or modification of files. This can be as a result of an attack, misuse by an employee or even a typo by an admin, any file, directory or registry change will be alerted to you.

Windows registry monitoring: Detects any changes done in the registry. Every time an application is installed or uninstalled it changes the Windows registry.

Rootkit detection: Rootkits detection is performed to detect privileged and hidden access to the system. Alerts are also sent out to the system administrators if a rootkit is detected.

Active responses: Take immediate and automatic responses when something happens. Active responses include, but are not limited to:

- Alert server admin and Temporary block IP addressed that perform suspicious activity on the server.

- Alert server admin of failed login attempts to the server.

- Alert server admin of changes in system files. **Data Backup** Data from the server is regularly backed up to an external device in a secure way. **Strong passwords** Strong passwords are used for all accounts. All passwords are required to be at least 16 characters in

length and have a combination of letters in both uppercase and lowercase, contain numbers and special characters. **Password Expiration** User passwords are changed every 3 months. **Limited access** Only select key authorized personnel are given access to connect to the server. **Windows Servers Remote Desktop**

Remote Desktop is restricted to connections from the BEAR Center's Berkeley office only. Transport Layer Encryption has been enabled on the RDP protocol used by remote Desktop. The port to connect to the server is changed from the default.

Roles

Only required roles and services are installed on the server. All non-essential services and roles have been removed or disabled from servers.

IIS

DRDPtech resides in its own application pool.

Updates

The latest service packs and hotfixes from Microsoft are installed regularly.

Firewall

Windows firewall is configured to reject connections on all unused ports.

Linux Server

IPtables

IPtable is configured to close all unused ports and reject all packages that try to access a blocked port or an open port with a non-accepted protocol.

SELinux

SELinux is enabled to add extra layer of security. SELinux provides a mechanism for supporting access control security policies, including United States Department of Defense-style mandatory access controls, through the use of Linux Security Modules (LSM) in the Linux kernel.

Remote Access

FTP, telnet uninstalled. Only access using ssh and sftp trough ssh. shh is

a network protocol for secure data communication, remote shell services or command execution and other secure network services between two networked computers that it connects via a secure channel.

SSH Port

Port for ssh changed from default.

PAM

Linux Pluggable Authentication Modules are used to provide dynamic authorization for applications and services.

Security of AWS (Amazon Web Services)

(extracted form http://aws.amazon.com/security/ and http://d36cz9buwru1tt.cloudfront.net/pdf/AWS_Security_Whitepaper.pdf)

AWS builds services in accordance with security best practices, provides appropriate security features in those services, and documents how to use those features.

Control Environment Summary

AWS manages a comprehensive control environment that includes the necessary policies, processes and control activities for the delivery of each of the web service offerings. The collective control environment encompasses the people, processes, and technology necessary to maintain an environment that supports the effectiveness of specific controls and the control frameworks for which AWS is certified and/or compliant.

AWS is compliant with various certifications and third-party attestations. These include:

• SAS70 Type II. This report includes detailed controls AWS operates along with an independent auditor opinion about the effective operation of those controls. • PCI DSS Level 1. AWS has been independently validated to comply with the PCI Data Security Standard as a shared host service provider.

• ISO 27001. AWS has achieved ISO 27001 certification of the Information Security Management System (ISMS) covering infrastructure, data centers, and services. • FISMA. AWS enables government agency

customers to achieve and sustain compliance with the Federal Information Security Management Act (FISMA). AWS has been awarded an approval to operate at the FISMA-Low level. It has also completed the control implementation and successfully passed the independent security testing and evaluation required to operate at the FISMA-Moderate level. AWS is currently pursuing an approval to operate at the FISMA-Moderate level from government agencies.

Secure Design Principles

AWS' development process follows secure software development best practices, which include formal design reviews by the AWS Security Team, threat modeling, and completion of a risk assessment. Static code analysis tools are run as a part of the standard build process, and all deployed software undergoes recurring penetration testing performed by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through launch to ongoing operations.

Monitoring

AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Documentation is maintained to aid and inform operations personnel in handling incidents or issues. If the resolution of an issue requires collaboration, a conferencing system is used which supports communication and logging capabilities. Trained call leaders facilitate communication and progress during the handling of operational issues that require collaboration. Post-mortems are convened after any significant operational issue, regardless of external impact, and Cause of Error (COE) documents are drafted so the root cause is captured and preventative actions are taken in the future. Implementation of the preventative measures is tracked during weekly operations meetings.

Information and Communication

AWS has implemented various methods of internal communication at a

global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employees; regular management meetings for updates on business performance and other matters; and electronics means such as video conferencing, electronic mail messages, and the posting of information via the Amazon intranet.

AWS has also implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "Service Health Dashboard" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact. A “Security and Compliance Center” is also available to provide customers with a single location to obtain security and compliance details about AWS. Customers may subscribe to Premium Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

Employee Lifecycle

AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS requires that staff with potential access to customer data undergo an extensive background check (as permitted by law) commensurate with their position and level of data access. The policies also identify functional responsibilities for the administration of logical access and security.

Account Provisioning The responsibility for provisioning employee and contractor access is shared across Human Resources (HR), Corporate Operations and Service Owners. A standard employee or contractor account with minimum privileges is provisioned in a disabled state when a hiring manager submits his or her approval. The account is automatically enabled when the employee’s record is activated in Amazon’s HR system. Access to other resources including Services, Hosts, Network devices, Windows and UNIX groups must be explicitly approved in Amazon’s proprietary permission management system by the appropriate owner or manager. All changes affected in the permissions management tool are captured in an audit. When changes in an employee’s job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked.

Account Review Every access grant is reviewed every 90 days; explicit re-approval is required or access to the resource is automatically revoked.

Access Removal Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. Windows and UNIX accounts are disabled and Amazon's permission management system removes the user from all systems.

Password Policy Access and administration of logical security for Amazon relies on user IDs, passwords and Kerberos to authenticate users to services, resources and devices as well as to authorize the appropriate level of access for the user. AWS Security has established a password policy with required configurations and expiration intervals.

Physical Security

Amazon has many years of experience in designing, constructing, and operating large-scale datacenters. This experience has been applied to the AWS platform and infrastructure. AWS datacenters are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access datacenter floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

AWS only provides datacenter access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to datacenters by AWS employees is logged and audited routinely.

Environmental Safeguards

Amazon's data centers are state of the art, utilizing innovative architectural and engineering approaches. Fire Detection and Suppression Automatic fire detection and suppression equipment has been installed to reduce risk. The fire detection system utilizes smoke detection sensors in all data center environments, mechanical and electrical infrastructure spaces, chiller rooms and generator equipment rooms. These areas are protected by either wet-pipe, double-interlocked pre-action, or gaseous sprinkler systems.

Power The data center electrical power systems are designed to be fully redundant and maintainable without impact to operations, 24 hours a day, and seven days a week. Uninterruptible Power Supply (UPS) units provide back-up power in the event of an electrical failure for critical and essential loads in the facility. Data centers use generators to provide back-up power for the entire facility. Climate and Temperature Climate control is required to maintain a constant operating temperature for servers and other hardware, which prevents overheating and reduces the possibility of service outages. Data centers are conditioned to maintain atmospheric conditions at optimal levels. Personnel and systems monitor and control temperature and humidity at appropriate levels. Management AWS monitors electrical, mechanical and life support systems and equipment so that any issues are immediately identified. Preventative maintenance is performed to maintain the continued operability of equipment.

Configuration Management

Emergency, non-routine, and other configuration changes to existing AWS infrastructure are authorized, logged, tested, approved, and documented in accordance with industry norms for similar systems. Updates to AWS' infrastructure are done to minimize any impact on the customer and their use of the services. AWS will communicate with customers, either via email, or through the AWS Service Health Dashboard (http://status.aws.amazon.com/) when service use is likely to be adversely affected.

Software AWS applies a systematic approach to managing change so that changes to customer impacting services are thoroughly reviewed, tested, approved and well communicated. AWS' change management process is designed avoid unintended service disruptions and to maintain the integrity of service to the customer. Changes deployed into production environments are: • Reviewed: Peer reviews of the technical aspects of a change • Tested: being applied will behave as expected and not adversely impact performance • Approved: to provide appropriate oversight and understanding of business impact

Changes are typically pushed into production in a phased deployment starting with lowest impact areas. Deployments are tested on a single system and closely monitored so impact can be evaluated. Service owners have a number of configurable metrics that measure the health of the service's upstream dependencies. These metrics are closely monitored with thresholds and alarming in place. Rollback procedures are documented in the Change Management (CM) ticket. When possible, changes are scheduled during regular change windows. Emergency

changes to production systems that require deviations from standard change management procedures are associated with an incident and are logged and approved as appropriate.

Periodically, AWS performs self-audits of changes to key services to monitor quality, maintain high standards and to facilitate continuous improvement of the change management process. Any exceptions are analyzed to determine the root cause and appropriate actions are taken to bring the change into compliance or roll back the change if necessary. Actions are then taken to address and remediate the process or people issue.

Infrastructure Amazon's Corporate Applications team develops and manages software to automate IT processes for UNIX/Linux hosts in the areas of third-party software delivery, internally developed software and configuration management. The Infrastructure team maintains and operates a UNIX/Linux configuration management framework to address hardware scalability, availability, auditing, and security management. By centrally managing hosts through the use of automated processes that manage change, the Company is able to achieve its goals of high availability, repeatability, scalability, robust security and disaster recovery. Systems and Network Engineers monitor the status of these automated tools on a daily basis, reviewing reports to respond to

hosts that fail to obtain or update their configuration and software. Internally developed configuration management software is installed when new hardware is provisioned. These tools are run on all UNIX hosts to validate that they are configured and that software is installed in compliance with standards determined by the role assigned to the host. This configuration management software also helps to regularly update packages that are already installed on the host. Only approved personnel enabled through the permissions service may log in to the central configuration management servers.

Business Continuity Management

Amazon's infrastructure has a high level of availability and provides customers the features to deploy a resilient IT architecture. AWS has designed its systems to tolerate system or hardware failures with minimal customer impact. Data center Business Continuity Management at AWS is under the direction of the Amazon Infrastructure Group.

Availability Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In

case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load- balanced to the remaining sites. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. This means that Availability Zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by Region). In addition to discrete uninterruptable power supply (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability Zones are all redundantly connected to multiple tier-1 transit providers. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Distributing applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes including natural disasters or system failures. Incident Response The Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide 24 x 7 x 365 coverage to detect incidents and to manage the impact and resolution. Company-Wide Executive Review Amazon's Internal Audit group has recently reviewed the AWS services resiliency plans, which are also periodically reviewed by members of the Senior Executive management team and the Audit Committee of the Board of Directors. Note that on April 21, 2011, EC2 suffered a customer-impacting service disruption in the US East Region. Details about the service disruption are described in "Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region" (http://aws.amazon.com/message/65648/).

Backups

Data stored in Amazon S3, Amazon SimpleDB, or Amazon Elastic Block Store (EBS) is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. Amazon S3 and Amazon SimpleDB provide object durability by storing objects multiple times across multiple Availability Zones on the initial write and then actively doing further replication in the event of device unavailability or detected bit-rot. Amazon EBS replication is stored within the same Availability Zone, not across multiple zones and therefore it is highly recommended that customers conduct regular snapshots to Amazon S3 for long-term data durability. For customers that have architected complex transactional databases using EBS, it is recommended that backups to

Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed. AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices.

Fault Separation

AWS provides customers the flexibility to place instances and store data within multiple geographic Regions. Each Region is an independent collection of AWS resources in a defined geography. AWS currently supports five Regions: US East (Northern Virginia), US West (Northern California), EU (Ireland), Asia Pacific (Singapore) and Asia Pacific (Tokyo). The Amazon S3 US Standard Region includes the US East facilities in Northern Virginia and facilities in Western Washington State. The selection of a Region within an acceptable geographic jurisdiction to the customer provides a solid foundation to meeting location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between Regions unless proactively done so by the customer, thus allowing customers with these types of data placement and privacy requirements the ability to establish compliant environments. It should be noted that all communications between Regions is across public Internet infrastructure. Appropriate encryption methods should be used to protect sensitive data. Within a given Region, Amazon EC2, Amazon EBS and Amazon Relational Database Service (RDS) allow customers to place instances and store data across multiple Availability Zones. See the "Business Continuity Management" section for more information on availability. Amazon S3, Amazon SimpleDB, Amazon Simple Notification Service (SNS), and Amazon Simple Queue Service (SQS) do not expose the concept of Availability Zones to customers. With these services, data is automatically stored on multiple devices across multiple facilities within a Region.

Network Security

The AWS network provides significant protection against traditional network security issues and the customer can implement further protection. The following are a few examples:

Distributed Denial Of Service (DDoS) Attacks AWS Application Programming Interface (API) endpoints are hosted on large, Internet-scale, world-class infrastructure that benefits from the same engineering expertise that has built Amazon into the world's largest online retailer. Proprietary DDoS mitigation techniques are used. Additionally, AWS's networks are multi-homed across a number of providers to achieve Internet access diversity.

Man In the Middle (MITM) Attacks All of the AWS APIs are available via SSL-protected endpoints which provide server authentication. Amazon EC2 AMIs automatically generate new SSH host certificates on first boot and log them to the instance's console. Customers can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. Customers are encouraged to use SSL for all of their interactions with AWS.

IP Spoofing Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.

Port Scanning Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on our website at: http://aws.amazon.com/contact-us/report-abuse/ When unauthorized port scanning is detected it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by the customer. The customer's strict management of security groups can further mitigate the threat of port scans. If the customer configures the security group to allow traffic from any source to a specific port, then that specific port will be vulnerable to a port scan. In these cases, the customer must use appropriate security measures to protect listening services that may be essential to their application from being discovered by an unauthorized port scan. For example, a web server must clearly have port 80 (HTTP) open to the world, and the administrator of this server is responsible for the security of

the HTTP server software, such as Apache. Customers may request permission to conduct vulnerability scans as required to meet their specific compliance requirements. These scans must be limited to the customer's own instances and must not violate the AWS Acceptable Use Policy. Advanced approval for these types of scans can be initiated by submitting a request via the website at: https://aws-portal.amazon.com/gp/aws/html-forms- controller/contactus/AWSSecurityPenTestRequest Packet sniffing by other tenants It is not possible for a virtual instance running in promiscuous mode to receive or "sniff" traffic that is intended for a different virtual instance. While customers can place their interfaces into promiscuous mode, the hypervisor will not deliver any traffic to them that is not addressed to them. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic. Attacks such as ARP cache poisoning do not work within Amazon EC2 and Amazon VPC. While Amazon EC2 does provide ample protection against one customer inadvertently or maliciously attempting to view another's data, as a standard practice customers should encrypt sensitive traffic.

Amazon Elastic Compute Cloud (Amazon EC2) Security

Security within Amazon EC2 is provided on multiple levels: the operating system (OS) of the host system, the virtual instance operating system or guest OS, a firewall, and signed API calls. Each of these items builds on the capabilities of the others. The goal is to protect against data contained within Amazon EC2 from being intercepted by unauthorized systems or users and to provide Amazon EC2 instances themselves that are as secure as possible without sacrificing the flexibility in configuration that customers demand.

Multiple Levels of Security

Host Operating System: Administrators with a business need to access the management plane are required to use multi-factor authentication to gain access to purpose-built administration hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane of the cloud. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems are revoked.

Guest Operating System: Virtual instances are completely controlled by the customer. Customers have full root access or administrative control over accounts, services, and applications. AWS does not have any access rights to customer instances and cannot log into the guest OS. AWS

recommends a base set of security best practices to include disabling password-only access to their hosts, and utilizing some form of multi-factor authentication to gain access to their instances (or at a minimum certificate-based SSH Version 2 access). Additionally, customers should employ a privilege escalation mechanism with logging on a per-user basis. For example, if the guest OS is Linux, after hardening their instance, they should utilize certificate-based SSHv2 to access the virtual instance, disable remote root login, use command-line logging, and use ‘sudo’ for privilege escalation. Customers should generate their own key pairs in order to guarantee that

_____ _____ _____

they are unique, and not shared with other customers or with AWS.

Firewall: Amazon EC2 provides a complete firewall solution; this mandatory inbound firewall is configured in a default deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic. The traffic may be restricted by protocol, by service port, as well as by source IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

The firewall can be configured in groups permitting different classes of instances to have different rules. Consider, for example, the case of a traditional three-tiered web application. The group for the web servers would have port 80 (HTTP) and/or port 443 (HTTPS) open to the Internet. The group for the application servers would have port 8000 (application specific) accessible only to the web server group. The group for the database servers would have port 3306 (MySQL) open only to the application server group. All three groups would permit administrative access on port 22 (SSH), but only from the customer’s corporate network. Highly secure applications can be deployed using this expressive mechanism. See diagram below:

The firewall isn’t controlled through the Guest OS; rather it requires the customer's X.509 certificate and key to authorize changes, thus adding an extra layer of security. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling the customer to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports are opened by the customer, and for what duration and purpose. The default state is to deny all incoming traffic, and customers should plan carefully what they will open when building and securing their applications. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages customers to apply additional per-instance filters with host-based firewalls

such as IPtables or the Windows Firewall and VPNs. This can restrict both inbound and outbound traffic on each instance. API calls to launch and terminate instances, change firewall parameters, and perform other functions are all signed by the customer's Amazon Secret Access Key, which could be either the AWS Accounts Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to the customer's Secret Access Key, Amazon EC2 API calls cannot be made on his/her behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints. AWS IAM also enables a customer to further control what APIs a user created with AWS IAM has permissions to call.

The Hypervisor

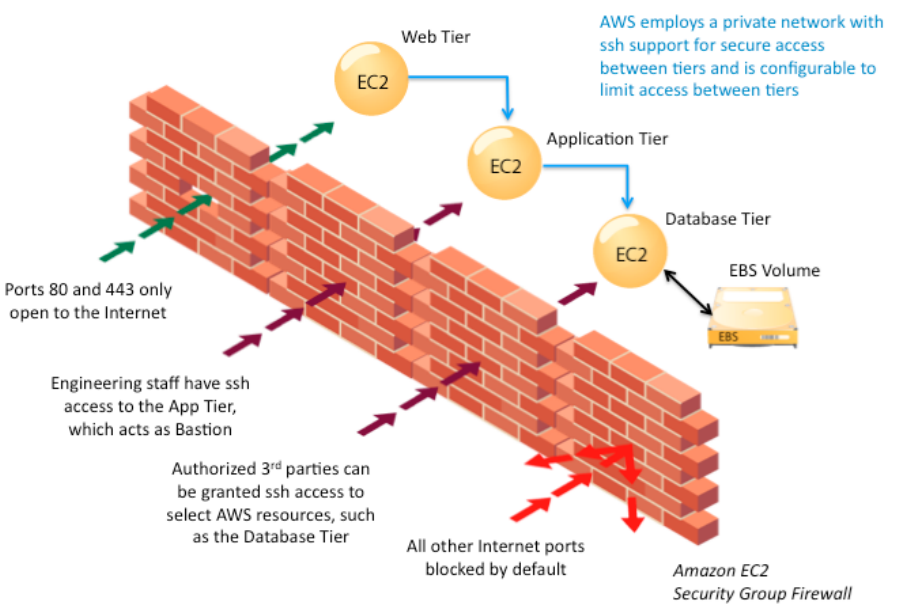

Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor, taking advantage of paravirtualization (in the case of Linux guests). Because paravirtualized guests rely on the hypervisor to provide support for operations that normally require privileged access, the guest OS has no elevated access to the CPU. The CPU provides four separate privilege modes: 0-3, called rings. Ring 0 is the most privileged and 3 the least. The host OS executes in Ring 0. However, rather than executing in Ring 0 as most operating systems do, the guest OS runs in a lesser-privileged Ring 1 and applications in the least privileged Ring 3. This explicit virtualization of the physical resources leads to a clear separation between guest and hypervisor, resulting in additional security separation between the two.

Instance Isolation

Different instances running on the same physical machine are isolated from each other via the Xen hypervisor. Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance’s neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.

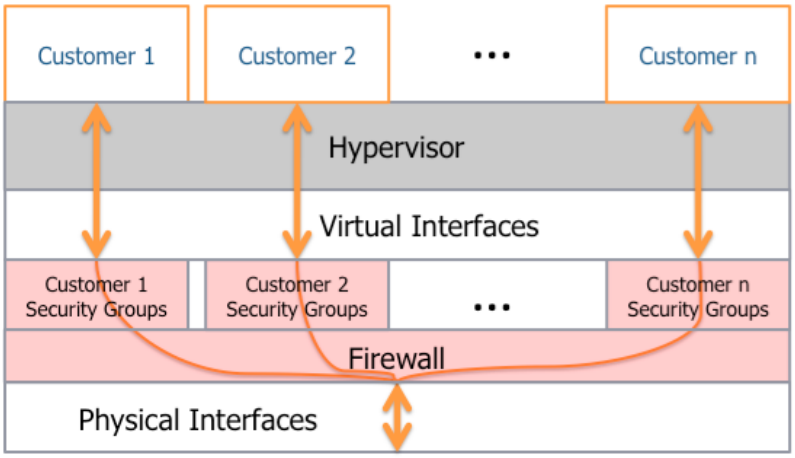

Customer instances have no access to raw disk devices, but instead are presented with virtualized disks. The AWS proprietary disk virtualization layer automatically resets every block of storage used by the customer, so that one customer’s data are never unintentionally exposed to another. AWS recommends customers further protect their data using appropriate means. One common solution is to run an encrypted file system on top of the virtualized disk device.