Keyfort Cloud Services (KCS)

Data Location, Security and Privacy

1. Executive Summary

   The purposes of this document is to provide a common understanding of the data location, security, privacy, resiliency and disaster recovery aspects of KCS.

2. Corporate Security Policy

   Keyfort's security policy defines and facilitates the secure operation of the KCS systems, data therein and its supporting organization.

   For systems, the security policy addresses constraints on functions and flow among them, constraints on access by external systems and adversaries including programs and access to data by people.

   For Keyfort, it addresses the constraints on behavior of its members as well as constraints imposed on adversaries by mechanisms such as encryption, firewalls, intrusion detection and physical security.

3. Organizational Security

   Keyfort:

   - Reviews security plans for Keyfort's networks, systems, and services using multi layered processes.
   - Conducts security design and implementation level reviews.
   - Provides ongoing consultation on security risks associated with a given project.

- Monitors for suspicious activity on Keyfort's networks, systems and applications, and follows formal incident response processes to recognize, analyse, and remedy information security threats.

- Drives compliance with established policies through security evaluations and internal audits.

- Develops and delivers training for employees on complying with Keyfort security policy, including in the areas of data security and secure development.

- Runs a vulnerability management program to help discover problem areas on Keyfort's networks, and participates in remedying known issues within expected time-lines. Keyfort also works with the security community outside of Keyfort.

- Works with physical security teams dedicated to the physical security of Keyfort's data centres. These security officers are qualified with training to protect high security enterprises with mission-critical infrastructures.

4. Data Asset Management

Data assets comprise customer and Keyfort data.

Operation of KCS includes the use of virtual machines running across multiple physical machines. This provides for processing, memory, data storage and network bandwidth on demand in an automatic, dynamic and scalable environment. As a result in the event of physical hardware failure the virtual machines are automatically redistributed to other hardware without service interruption. Likewise additional hardware can be added and taken out for maintenance without loss of service.

The virtual machines, each with its encapsulated customer services, are automatically backed up complete with operating system, application software, data and configuration enabling rapid restoration of service in the event of disruption.

The data is stored across multiple data storage devices complete with RAID10 mirroring and auto backup. Thereby avoiding disruption in the event of hardware failure and also past data versions can be restored as required.

Dual system monitoring complete with alerts and logging is implemented throughout KCS.

Automatic data replication from customer site(s) to KCS is available for disaster recovery purposes.

Automatic data replication from KCS and any linked cloud services such as Google Apps for Work or Microsoft Office 365 to the customer site using a Keyfort data storage device, provided for the purpose and located on the customer's site, is available. This ensures that the customer always has a copy of all their data which can be used for disaster recovery and independence of Cloud purposes.

All data is held in country of origin unless requested otherwise by the customer. Keeping data in country simplifies legal jurisdiction and reduces security risk.

Keyfort can supply, as appropriate, a Written Statement of Assertion in accordance with the International Standard on Assurance Engagements ISAE 3402 standard if so required by the customer.

5. Access Control

Keyfort employs a number of authentication and authorization controls that are designed to protect against unauthorized access.

Authentication Controls

Keyfort requires the use of a unique user ID for each employee. This account is used to identify each person's activity on Keyfort's network, including any access to employee or customer data. This unique account is used for every system at Keyfort. When an employee commences work with Keyfort they are assigned a user ID and granted a default set of privileges. At the end of a person's employment, their account's access to Keyfort's network is disabled and passwords reset.

Keyfort's password policies include as appropriate password expiration, restrictions on password reuse, sufficient password strength, two-factor authentication, certificates and one-time password generators.

Authorization Controls

Access rights and levels are based on an employee's job function and role, using the concepts of least-privilege and need-to-know to match access privileges to defined

responsibilities. Keyfort employees are only granted a limited set of default permissions to access company resources, such as their email, and Keyfort's internal portal. Employees are granted access to certain additional resources based on their specific job function. Requests for additional access follow a formal process that involves a request and an approval from a data or system owner, manager, or other executives, as dictated by Keyfort's security policies. An employee's authorization settings are used to control access to all resources, including data and systems for Keyfort's cloud technologies and products.

Accounting

Keyfort's policy is to log administrative access to every Keyfort production system and all data. These logs are reviewable by Keyfort on an as-needed basis.

6. Personnel Security

Keyfort employees are required to conduct themselves in a manner consistent with the company's guidelines regarding confidentiality, business ethics, appropriate usage, and professional standards. Upon employment, Keyfort will verify an individual's education and previous employment, and perform internal and external reference checks. Where appropriate and statutory regulations permit, Keyfort may also conduct criminal, credit, immigration, and security checks. The extent of background checks is dependent on the desired position. Upon acceptance of employment at Keyfort, all employees are required to execute a confidentiality agreement as part of their employment contract.

7. Physical and Environmental Security

KCS includes power feeds (A&B), uninterruptible power supply (N+1), generator backup (N+1), air conditioning (N+1) and fire suppression systems to ensure its server environment is resilient.

KCS features.

- Direct connection to services via encrypted data links from your designated office(s) and roving staff. Ensuring that your data is available in part or full only as you designate for staff and selected partners if any.

- Secure resilient 100Mbps high performance Internet connection. Greater bandwidth speed is available on request.

- 24x7 security and support.
- In the event of a power failure your services are automatically protected through an uninterruptible power supply (large mains batteries) with further protection via generators with many hours of fuel.
- All servers are above flood level.
- Security personnel ensure that only authorised members of staff are permitted access.

8. Infrastructure Security

KCS includes a series of threat prevention and management services.

- Automatically updated malware prevention
- Automatically updated intrusion prevention.
- Automatically updated anti-virus.
- Firewall protocol, application, access and bandwidth control.[1]
- Encrypted network links.
- White lists.
- Black lists.
- Dual monitoring with auto alerts.
- Traffic and activity logging.
- Change management to provide a centralized mechanism for registering, approving, and tracking changes.
- Operating systems and application services are 'patched' and upgraded as an when required on a proactive basis.

[1] See Appendix A, *Firewall Management Policy*, version 1.4.

9. Systems and Software Development and Maintenance

Security is a key component of our design and development process. Keyfort's engineering organization does not require development teams to follow a specific software development process; rather, teams choose and implement processes that fit the project's needs. As such, a variety of software development processes are in use at Keyfort, from Agile Software Development methodologies to more traditional, phased processes. Keyfort's security review processes are adapted to work within the chosen framework. Engineering management has defined requirements for project development processes:

- Peer reviewed design documentation.
- Adherence to coding style guidelines.
- Peer code review.
- Multi-layered security testing.

The above mandates embody Keyfort's software engineering culture, where key objectives include software quality, robustness, and maintainability. While the primary goal of these mandates is to foster the creation of software artifacts that excel in all aspects of software quality, Keyfort's experience also suggests that they can reduce the incidence of security flaws and defects in software design and implementation.

- The existence of adequately detailed design documentation is a prerequisite of the security design review process, since in early project stages it is generally the only available artifact on which to base security evaluations.
- Many classes of implementation-level security vulnerabilities are fundamentally no different from low-risk, common functional defects. Many implementation-level vulnerabilities are caused by fairly straightforward oversights on the developer's part.
- Given developers and code reviewers who are educated with respect to applicable vulnerability patterns and their avoidance, a peer review-based development culture that emphasizes the creation of high-quality code supports a secure code base.

10. Disaster Recovery and Business Continuity

The objective is to minimise service interruption due to hardware failure, natural disaster, or other catastrophe. Keyfort implements a disaster recovery program at all of its data centres. This program includes multiple components to minimize the risk of any single point of failure, including the following:

- Data replication and backup. Application data is replicated to multiple systems within a data centre and in some cases also replicated to multiple sites.

- Keyfort operates a geographically distributed set of data centres/customer sites that are designed to maintain service continuity in the event of a disaster or other incident in a single region. High-speed connections between these centres help to support swift failover. Management of these centres is also distributed to provide location-independent, around-the-clock coverage, and system administration. Keyfort conducts regular testing of its disaster recovery plans. For example a disaster in a centre is simulated by taking IT systems and business and operational processes in that location off-line, and allowing such systems and processes to transfer to fail-over services designated by the disaster recovery (DR) plan. During the course of the test, it is verified that business and operations functions can operate during fail-over and hidden/unknown dependencies on the off-line location are identified and logged for later remediation.

11. Summary

KCS provides a resilient, secure and cost effective environment for customer services and data. From initial porting of the data and services to the cloud to its subsequent use, data privacy is respected as the right of the customer and limited in accordance with the customer's requirements. All data is transmitted over encrypted links and is stored in country of origin unless otherwise requested by the customer. Keyfort also provides an option for keeping a synchronized updated copy of the data held within KCS on the customer premises with a data storage unit supplied for the purposes thereof.

Appendix A
Firewall Management Policy

Firewall Management Policy

Table of Contents

1. Introduction

Firewall management is resource intensive and requires a high level of expertise to prevent unauthorised access and costly breaches. Devices must be provisioned, deployed, upgraded, licenced where appropriate, and patched to be kept up to date to meet the latest threats.

Security policies are subject to change in accordance with an organization's evolving requirements and these changes need to be reflected in the firewall policies for the associated firewalls, thereby necessitating firewall configuration updates to ensure appropriate access controls are in place.

Firewall availability and network traffic is subject to 24x7 secure monitoring by Keyfort's *KCS* cloud computing dual monitoring systems.[2]

KCS Managed Firewalls. The key benefits are the:

- Protection of systems and data with current corporate strength firewall services complete with 24x7 monitoring for alert purposes.

- Removal of the management and monitoring burden overhead for customers through the use of purpose provided expert systems engineers and services.

- Support of compliance initiatives including annual vulnerability scans.

[2] Keyfort cloud computing services known collectively as, *'KCS'.*

2. Firewall Management Policy

This policy is vital to the pursuit of external connectivity, data protection and commerce in a secure manner. This policy governs everything from acceptable use to response scenarios in the event a security incident occurs. It also serves as a customer briefing re Keyfort's *Firewall Management Policy*. This firewall management policy is distinct from the Information Security Policy, in as much as it is simply a description of how the Information Security Policy will be implemented by the firewall and associated security mechanisms. The purpose of this firewall policy is to provide a constructive framework and guidance for systems administrators and organizations to promote a secure environment for systems and data. Firewalls can be complex to manage and security incidents can occur daily. Without a policy to guide firewall implementation and administration, the firewall itself may become a security problem.

3. Firewall Policy Creation

A firewall policy dictates how the firewall should handle applications traffic such as web, email, or telnet. The policy describes how the firewall is managed and updated. Before a firewall policy can be created, some form of risk analysis must be performed on the applications that are necessary for accomplishment of the organization's mission. The results of the analysis will include a list of the applications and how these applications will be secured. The process to create this list will require knowledge of the vulnerabilities associated with each application and the cost-benefits associated with the methods used for securing the applications. Risk analysis of the organization's information technology infrastructure includes an evaluation of threats, vulnerabilities and counter measures required to mitigate vulnerabilities and the impact if sensitive data is compromised. The goal is to understand and evaluate these elements prior to establishing the firewall policy. The result of the risk analysis dictates the manner in which the firewall system handles network applications traffic.

The steps involved in creating the firewall policy are as follows:

3.1. Identification of network applications deemed necessary.

3.2. Identification of vulnerabilities associated with applications.

3.3. Cost-benefits analysis of methods for securing the applications.

3.4. Creation of applications traffic matrix showing protection method.

4. Firewall Policy Implementation

The firewall policy implementation is performed by creating and activating rulesets as the mechanism for implementing security controls. The contents of the rulesets determine the actual functionality of the firewall. Depending on the firewall platform architecture, firewall rulesets can contain various pieces of information. Nearly all rulesets, however, will contain the following fields, as a minimum:

4.1. The source address of the packet i.e. the Layer 3 address of the computer system or device the network packet originated from (an IP address such as 192.168.1.1).

4.2. The destination address of the packet, in other words, the Layer 3 address of the computer system or device the network packet is trying to reach (e.g., 192.168.1.2).

4.3. The type of traffic, in other words, the specific network protocol being used to communicate between the source and destination systems or devices, usually Ethernet at Layer 2 and IP at Layer 3.

4.4. Possibly some characteristics of the Layer 4 communications sessions i.e. the protocol TCP and the source and destination ports of the sessions e.g. TCP:80 for the destination port belonging to a web server, TCP:1320 for the source port belonging to a personal computer accessing the server.

4.5. Sometimes, information pertaining to which interface of the router the packet came from and which interface of the router the packet is destined for. This is useful for routers with three or more network interfaces.

4.6. An action, such as Deny or Permit the packet, or Drop the packet which does not return a response to the packet's source.

Customers should be aware that firewall rulesets tend to become increasingly complicated with age as new rules are added and old ones are modified or deleted in accordance with the customer organization's evolving requirements. New users or organizational requirements typically drive these changes but they can also reflect political forces within an organization or compliance requirements.

The firewall ruleset is assembled after completing the applications traffic matrix and are built to be as specific as possible with regards to the network traffic they control. Rulesets are not kept as simple as possible, as this makes them vulnerable to inadvertently leaving access routes which may be exploited. Thereby permitting unauthorized or unwanted traffic to traverse the firewall. The default policy for the firewall for handling inbound traffic is to block all packets and connections unless the traffic type and connections have been specifically permitted. This approach is more secure than permitting all connections and traffic by default and then blocking specific traffic and connections.

The firewall ruleset should always block the following types of traffic:

4.7. Inbound traffic from a non-authenticated source system with a destination address of the firewall system itself. This type of packet normally represents some type of probe or attack against the firewall. One common exception to this rule would be in the event the firewall system accepts delivery of inbound email (SMTP on port 25). In this event, the firewall must allow inbound connections to itself, but only on port 25.

4.8. Inbound traffic with a source address indicating that the packet originated on a network behind the firewall. This type of packet likely represents some type of spoofing attempt.

4.9. Inbound traffic containing ICMP (Internet Control Message Protocol) traffic. Since ICMP can be used to map the networks behind certain types of firewalls, ICMP should not be passed in from the Internet or from any untrusted external network.

4.10. Inbound or outbound traffic from a system using a source address that falls within the address ranges set aside in *RFC 1918* as being reserved for private networks. For reference purposes, *RFC 1918* reserves the following address ranges for private networks: 10.0.0.0 to 10.255.255.255 (Class A) 172.16.0.0 to 172.31.255.255 (Class B) 192.168.0.0 to 192.168.255.255 (Class C). Inbound traffic with these source addresses typically indicates the beginning of a denial of-service (DoS) attack involving the TCP SYN flag. Some firewalls include internal functionality to combat these attacks, but this

particular type of network traffic should still be blocked with ruleset entries. This rule is now frequently automatically implemented by firewalls.

4.11. Inbound traffic from a non-authenticated source system containing SNMP (Simple Network Management Protocol) traffic. These packets can be an indicator that an intruder is probing a network, but there are few reasons an organization or agency might want to allow inbound SNMP traffic and it should be blocked in the vast majority of circumstances.

4.12. Inbound traffic containing IP source routing information. Source routing is a mechanism that allows a system to specify the routes a piece of network traffic will employ while traveling from the source system to the destination system. From a security standpoint, source routing has the potential to permit an attacker to construct a network packet that bypasses firewall controls. In modern networks, IP source routing is rarely used, and valid applications are even less common on the Internet.

4.13. Inbound or outbound network traffic containing a source or destination address of 127.0.0.1 (localhost). Such traffic is usually some type of attack against the firewall system itself.

4.14. Inbound or outbound network traffic containing a source or destination address of 0.0.0.0. Some operating systems interpret this address as either localhost or as a broadcast address, and these packets can be used for attack purposes.

4.15. Inbound or outbound traffic containing directed broadcast addresses. A directed broadcast is often used to initiate a broadcast propagation attack such as a Smurf attack.[3]

Some types of firewalls are also capable of integrating user authentication into ruleset enforcement. For example, many firewalls have the capability of blocking access to certain systems until a user authenticates to the firewall. This authentication can either be internal to the firewall or external to the firewall. Firewalls that implement application proxies can also integrate with advanced enterprise authentication schemes e.g. Active Directory. Most firewalls also support multiple options for logging. These options range

[3] A Smurf attack is a distributed denial of service (DDoS) attack in which large numbers of Internet Control Message Protocol (ICMP) packets with the intended victim's spoofed source IP are broadcast to a computer network using an IP broadcast address.

anywhere from the creation of simple log entries, up to options for alerting users that a certain event has occurred. Depending on the alert implementation, this action can include a range of options, from sending email notification, to paging appropriate personnel.

When the firewall rule set has been created:

- It is then activated.
- Its logging parameters set.
- Its configuration is (auto) backed-up.
- The firewall is placed on *KCS* 24x7 dual monitoring for availability status and alerts.

The firewall administrator, working via a secure connection, then has the option to print a copy of the ruleset that has been implemented on the firewall by printing out or down-loading a copy of the firewall's technical support report.

5. Firewall Testing

Firewall policies are implemented every day but many organizations rarely check and verify their firewall policies. With the increase in data protection rules such as the implementation of European Union (EU) general data protect regulations (GDPR) 25th May 2018 and the growing threat of cyber-crime the case for verifying the efficacy of firewall rules on a regular basis has now been made.[4] For nearly all organizations, firewall and security policies should be audited and verified annually at the very least. In many cases, firewall policy can be verified using one of two methodologies. The first methodology is to obtain hardcopies of the firewall configurations and compare these hardcopies against the expected configuration based on defined policy. The second methodology involves in situ configuration testing generally known as vulnerability scanning. In this methodology, Keyfort utilizes semi-automated commercial tools that assess

[4]Information Commissioner's Office, *Guide to the General Data Protection Regulation (GDPR)*, https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/ .

the configuration of a device by attempting to perform operations that should be prohibited. This provides a vulnerability report with findings categorised and prioritised. The results of which should be actioned as appropriate to address any vulnerability issues. Although these reviews can be completed with public domain tools, many organizations, especially those subject to regulatory requirements, will choose to employ commercial tools. While the second methodology is more rigorous, both methodologies may be employed. The goal is to make sure that the firewalls, as well as any other security-related devices, are configured exactly as they should be, as per the written policy. It is also important that the firewall system itself be tested using security assessment tools. These tools should be used to examine the underlying firewall operating system, as well as the firewall software and implementation.

6. Firewall Failover Strategies

Many options exist for providing redundancy and failover services for firewalls. These options range from using specially designed network switches to using firewall heartbeat mechanisms to assess and coordinate the availability of the primary firewall so that a backup can take over in the event of a failure. Network switches and firewalls can provide load balancing and failover capabilities. In a failover configuration, the switches monitor the responsiveness of the production firewall and shift all traffic over to a backup firewall in the event that there is a failure on the production system. The primary advantage to this type of solution is that the switch masquerades both firewalls behind the same MAC (Media Access Control) address. This functionality allows seamless failover; in many cases, established sessions through the firewall are not impacted by a production system failure. The firewall heartbeat solutions involve a firewall pair, with a single virtual MAC address, with a custom network interface to notify the backup firewall to assume primary firewall duties in the event that the primary firewall fails to respond with a heartbeat. These systems rely on established, reliable technology to handle failover. The primary drawback to this approach is that established sessions traversing the production firewall may be lost in the transition from the primary to the backup firewall. A significant advantage of the heartbeat approach is that one of the firewall pair may be taken out of service for updates or replacement and then added back without interrupting service. The decision on which failover method to implement is often reduced to cost; the network switch-based failover solution is generally more expensive than a heartbeat-based system.

7. Firewall Maintenance, Management, Backup and Logs

Generally firewall appliances, including virtual firewall appliances, are more secure than those implemented on top of commercial operating systems as firewall appliances do not suffer from the security vulnerabilities associated with an underlying commercial operating system.

Firewall appliances generally employ application specific integrated circuit (ASIC) technology, with the actual firewall software being present as firmware driving the ASICs. These firewalls also tend to be faster than firewalls implemented on top of commercial operating systems. The advantage of implementing firewall appliances as virtual systems is scalability. If an environment requires improved performance, organizations can buy additional computing capacity in the form of processing power (cpu), random access memory (RAM), network interface connections (NICs) and hard storage for logs etc.. Most appliances do not offer this level of flexibility or scalability.

Firewall appliances employ one of two mechanisms for configuration and ongoing maintenance. The first mechanism is command line interface (CLI) configuration, which enables an administrator to configure the firewall by typing commands into a command prompt. This technique is prone to error due to typing mistakes. The primary advantage to command line configuration is that a skilled and experienced administrator can configure the firewall and react to emergency situations more quickly than with a graphic interface. The second and more common mechanism for firewall configuration is through a graphic user interface. Graphic interfaces are simpler and enable an administrator to configure advanced systems in a reasonable amount of time. The major issue with graphic interfaces is configuration granularity. In many modern firewall platforms, there are options available in the firewall that cannot be configured using the graphic interface. In these circumstances, a command line interface must be used. For either option, great care is taken to ensure that all network traffic dealing with the firewall system management is secured. Thus, Keyfort uses encrypted links with network and user lockdown, monitoring and logs when accessing firewall administration interfaces, whether they be CLI or graphic interface. It is a matter of policy that all firewall management functions take place over secured links using strong authentication and encryption. This requires firewall firmware upgrades or failing that firewall replacement to ensure that the secured links meet the evolving current standards. Firewall firmware upgrades require firewall down time which is

subject to customer agreement. Where customer agreement for the upgrade is not received the customer will be advised of the security deficiency and the onus thereof will be on the customer.

Firewall rulesets are subject to change to meet organizational requirements hence Keyfort conducts scheduled monthly backup of configurations for firewalls whether they are in the cloud or on site. All backups are conducted exclusively over encrypted links with restricted network and user access.

Firewall logs are automatically sent via encrypted links to the *KCS* management system. Extended reporting options are available. In the event of a security breach or warning thereof by the monitoring systems Keyfort system engineers will liaise with the customer and close down the breach as appropriate. It should be noted that security is best effected by a multi-layered approach including LAN anti-virus, anti-spam, anti-phishing, email filtering, in and out-bound filtering and appropriate access controls. Firewalls, though significant, are only part of an organization's information security plan. The firewall logs may be of significant service in analysing any security breach.

In the event of a security breach restoring data and systems from backup is a key element in effectively responding to a breach. Therefore Keyfort recommends automatic off site daily backup of an organization's key data to Keyfort's *KCS* cloud service.

If the security has impacted personal data held by the organization, as Data Controller, or on behalf of the organization by another party, as Data Processor, then the Information Commissioner's Office must be informed, by the Data Controller, in accordance with the current data protection/GDPR procedures. Keyfort if acting as a Data Processor for the organization will provide appropriate supporting information.

8. Physical Security of the Firewall

The physical security of the firewall including its environment should not be overlooked. If the devices are located in a non secure area, they are susceptible to damage from intruders and at a higher risk to accidental damage. Therefore, firewall devices should be secured behind locked doors.

Another factor in physical security is the quality of the electrical and network connections and environment control. The firewall facility should have backup power supplies and possibly redundant connections to external networks. Some form of air-conditioning and air filtration is also typically a requirement.

Lastly, the firewall facility should be protected, as is reasonable, from natural disasters such as fire and flood. Fire suppressant systems are usually standard equipment in computing facilities.

Keyfort locates its *KCS* cloud service firewalls in secured computing facilities, complete with N+1 power and computing resilience, restricted access and other physical security alarms.

9. Periodic Review of Information Security Policies

As with any type of policy, information security policies must undergo periodic review in order to ensure accuracy and timeliness.[5] Best practice dictates that information security policies should be reviewed and updated at least annually. Best practice further dictates that several events can trigger a review of information security policies. These triggers include events such as the implementation of major enterprise computing environment modifications and any occurrence of a major information security incident. A review may be requested by either the customer or Keyfort and actioned subject to commercial agreement.

[5] An Information Security Policy is implemented, in part, via firewall policies. Hence a change to the Information Security is likely to require changes to the associated firewall policies.