Watson Developer Cloud Security Overview

Introduction

This document provides a high-level overview of the measures and safeguards that IBM® implements to protect and separate data between customers for implementations of IBM Watson™ Services. While data that is subject to regulation must not be used with Watson ™ Services, IBM recognizes that other types of customer data must be appropriately protected and segregated. Additionally, users might submit Personally Identifiable Information (PII) and regulated data through a Watson query. Procedures and governance are used to handle these unintended submissions in addition to maintaining the security of the Watson environment.

Deployment and Security Overview

Watson services may be deployed in a variety of ways such as the IBM Watson public cloud, premium plans which provides additional data isolation within the public cloud, or a dedicated cloud environment for when clients need infrastructure that only supports them. In each case the service remains the same and IBM ensures that the security architecture remains consistent.

Data: Watson services manage data in a variety of ways. Many of the Watson services are designed to be stateless in nature, meaning that while they may process data, they do not store it; the data is only used to complete the transaction and when that call to the service is completed the data is not retained. Some of the services allow customization which enables customers to help bring specific context to the data being submitted for processing. This configuration information when stored is isolated and encrypted. By default, transaction logs for each service are stored but users can opt out if so desired.

Authentication & Authorization: The services are instantiated within Bluemix. Once a service has been requested for use credentials are generated and managed through Bluemix. When a call to a Watson service is made by an application, the credentials are transmitted via HTTPS for authentication and authorization. This allows only authorized users access to their content. Once this step is completed; a temporary token is generated that is good for 60 minutes.

Encryption: Watson services only accept and send client data over the Internet using HTTPS via TLS connections with support for TLS version 1.2. Any client data stored is encrypted per IBM policy.

Base Security: IBM Watson security policy requires that all services include network and storage encryption, circuit and application level firewalls, security information and event management, intrusion detection, application source code scanning, 3rd party penetration testing, and regular vulnerability scanning. Figure 1 shows how these standards are used together.

Backup & Redundancy: Watson services leverage replication and snapshots to support these requirements. Implementation may vary depending on the service. Generally, data is replicated

securely across multiple instances or data store locations where daily snapshots are taken and stored using encrypted storage. Note: IBM may use outside vendors to assist with backup requirements.

Figure 1. Watson Services Environment Security

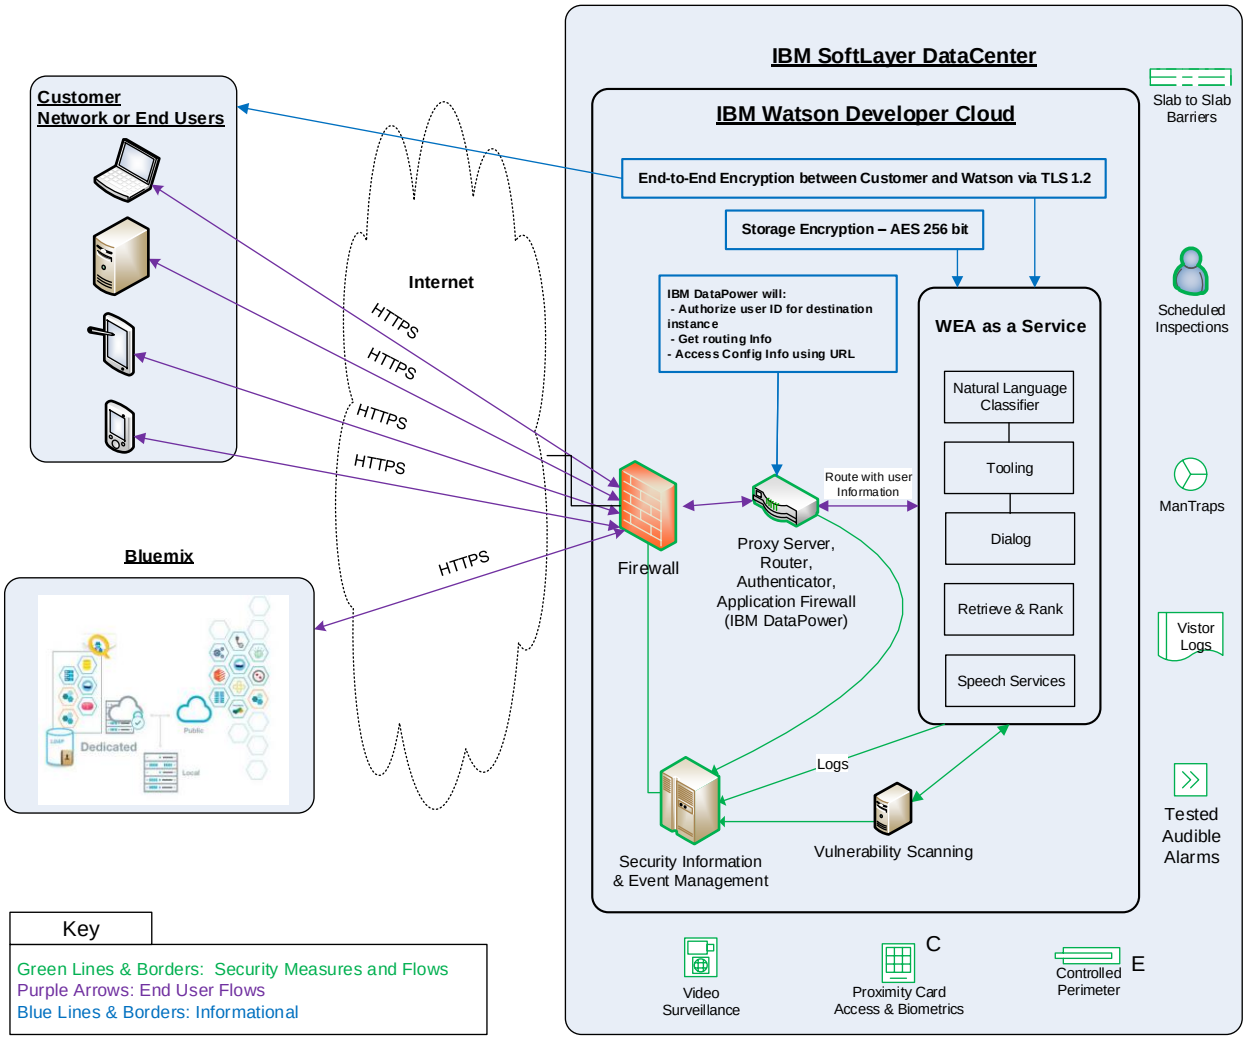

Watson Deployment Options

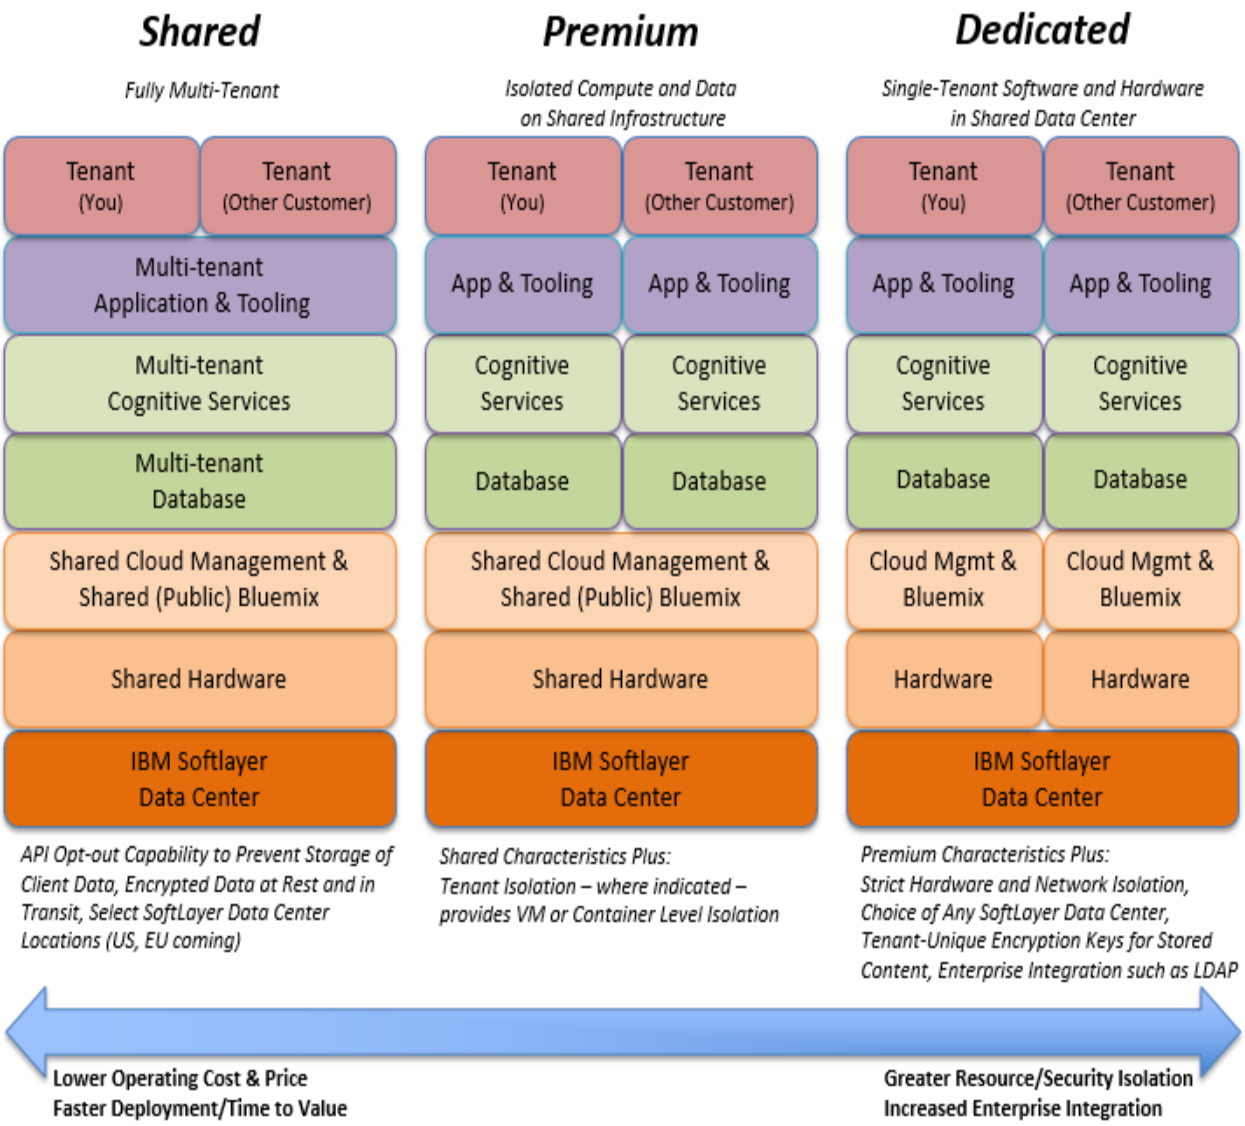

Watson works to satisfy a broad set of enterprise security and compliance requirements. Three cloud deployment models are available to meet various data requirements and business needs. All of our deployments reside within hardened enterprise class IBM SoftLayer data centers that are ISO27001 and SOC2 certified*.

1. Public - The Public Cloud is the most cost effective and provides a shared tenancy model which allows users to embrace the power of Watson services while sharing the infrastructure cost needed to run Watson. Each service provides unique credentials, API Opt-out capability (should users not want to share their data with Watson for service improvement), encryption of data in motion and at rest, and all of the enterprise security controls you expect from IBM. Public plans are a great option for companies not looking to include regulated or personally identifiable data (PII) into their Watson Services.
2. Premium Plans – Provide all of the features above with the added benefits of data isolation and service SLA's. Enterprise plans provide customers a unique instance of a Watson service that is dedicated for their use leveraging containers and dedicated

database instances to isolate client data. This option still leverages the advantage of shared hardware within the Watson Cloud environment. Enterprise plans are suggested for customers looking to use Watson services with non-regulated PII data or that may have other data isolation requirements.

3. Dedicated Deployments – Allow customers full data isolation by implementing a dedicated Watson Cloud for each customer. Customers get a dedicated instance of Bluemix and Watson services, which allows for integration with most enterprise single sign on solutions, tenant unique encryption keys, and added logging and monitoring capabilities, including detailed access logs. This not only allows customers to see who accessed their environment and when it was accessed, but it also provides the added benefit of knowing the complete solution is running on hardware dedicated for them in the geography of their choice. Dedicated deployments are appealing to enterprises with workloads that include sensitive data and have a need for additional transparency into where and how their data is managed.

Conclusions and Recommendations

Being aware of the features and benefits of Public, Premium, and Dedicated Cloud is key to designing the right transformation to the Cloud. The added transparency provided by extended security controls in Dedicated Cloud deployments can boost Cloud initiatives. It is important to uncover security and compliance requirements up front and to tackle the hard questions early by including all internal stakeholders into Cloud initiatives. Making sure to include the security and compliance and risk teams early in the process will help ensure that the Cloud adoption program remains on track.

- Additional information on SoftLayer security certifications can be found here: http://www.softlayer.com/compliance
- Additional information on Bluemix security certifications can be found here: http://www.ibm.com/cloud-computing/bluemix/trust/
- Additional information on Watson security certifications can be found here: http://www.ibm.com/watson/watson-security.html

Governance, Risk, and Compliance

The Watson data compliance strategy is built upon widely accepted Governance, Risk, and Compliance (GRC) principles as shown in Figure 2.

Figure 2. Watson Governance, Risk & Compliance

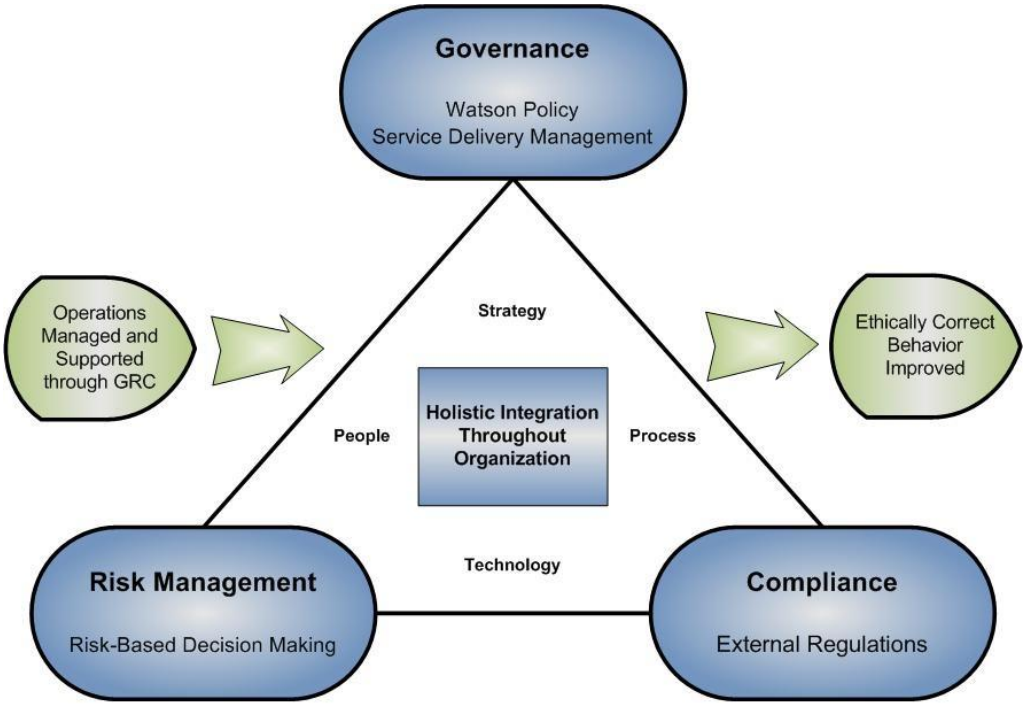

Security Policy

The Watson Cloud Security Policy is established by the IBM Corporate Directives that are defined at the highest level of IBM.

The Watson Security Policy maps to the ISO27002 structure. Watson security controls are designed to meet industry standard controls and are intended to assist with compliance to external regulations in the healthcare and financial sectors such as HIPAA, HITECH, and FFIEC when and where applicable.

Audits and Self-Assessment

IBM assesses and audits compliance with HIPAA, FFIEC, and IBM internal security policies.

Assessments can include:

- Self-assessment of security controls.
- Independent internal audits that are performed by using the security principle of separation of duties.
- 3rd party auditors (SSAE16, ISO27001, and government regulatory agencies)

External Audit

Watson Cloud Technology and Support has a team of professionals that are prepared to respond to external audits that are required under applicable law or regulation.

Risk Assessment

IBM recognizes risk assessment to be an important factor in security and has established a periodic risk assessment process that is applicable to the systems that host Watson as a Service. Assessments are entered into the IBM Governance, Risk, and Compliance program to determine & manage the current risk posture.

Physical security

Physical security of IBM property is defined at the global level and includes a layered approach that includes site, building, data center, and data center partitions. Employees have limited physical access based on their job requirements to systems that host "Watson as a Service" offerings.

Physical building security is maintained at various levels that are based on a categorization of security requirements for any physically partitioned area. The security includes but is not limited to gates, badge locks, cipher locks, key locks and biometrics, video monitoring and access logs. Data centers do not have first floor windows. Data center emergency doors are alarmed.

Logical security

Logical security consists primarily of technical means as specified by the IBM CIO and other security authorities within IBM. Watson as a Service logical security uses the following safeguards:

- Activity logging that includes suspicious activity monitoring of protected logs.
- End-to-end encryption of data.
- Isolation of customer data.
- Procedures for an emergency shutdown to prevent data leakage.
- Technical specifications that detail allowable configurations for devices.
- Application of security patches.
- Network configuration that includes zoned security layering that is enforced by mandatory firewall and router rule sets.
- Security for user devices.
- Antivirus and anti-malware protection with automated workstation compliance tools.
- Vulnerability scanning and intrusion detection.
- Change management process and information systems maintenance.
- DDoS protection of inbound circuits to data centers.
- Ongoing internal & external penetration testing/ethical hacking program.
- Regular application source code reviews, threat modeling, and application scanning.

Human Resource security

IBM Human Resource policies determine the required background checks and monitoring for employees. These policies are based on applicable local laws. Employees with elevated system privileges are subject to more stringent requirements.

All IBM employees are required take annual security education and to read and certify annually that they comply with established IBM Business Conduct Guidelines (BCGs). For more information about the BCGs, see http://www.ibm.com/investor/governance/business-conduct-guidelines.wss.

Secure Engineering

The IBM Watson Development teams institute the IBM Secure Engineering Framework which reflects best practices from across the company and directs our development teams to give proper attention to security during the development lifecycle. These practices are intended to help enhance product security, protect IBM and customer intellectual property and support the terms of warranty of IBM products.

Access control

Watson as a Service uses a provisioning system with robust security attributes that is used to manage access for IBM administrators and to retain audit trails of access control workflow. Secondary controls are used to enforce periodic revalidation of users that are based on continued business need and employment verification.

Access to systems that host the Watson as a Service offering is granted by management and is based on role requirements. Access is decided by using the principle of least privilege as a guide.

Cryptography

IBM employs the latest cryptographic technologies when available and technically feasible to protect customer data while at rest and in motion, examples include TLS/SSL, IPSEC, Third Party CAs, Encrypted File Systems, Encrypted Storage Systems, Key Management Systems, etc.

Deviations

On occasion, deviations from the written security practice might be discovered through an audit or other means. When conditions warrant, systems might be taken offline for a deviation until remedial actions are taken. Deviations must be applied for by using a defined process, tracked to closure, and remediated with approved interim measures until a final remediation is completed.

IBM Vendor Partners

Different companies can have different security practices while still conforming to prudent security principles. Watson vendor partners are carefully vetted and required to provide equally robust security practices in the area for which they provide their services.

Security incident management

A global management process for security incidents is employed and is applicable to the systems that host the Watson as a Service offering. This process is communicated to IBM employees and management, and is monitored 24x7x365 by trained IBM employees.

Notes Section

1. Natural Language Classifier: The database used in this solution is backed up at regular intervals and is protected while in transit and at rest. Backup copies are encrypted in transit via SSL and then written on encrypted disk using a unique key. The current backup location utilizes Amazon Simple Storage Service.

NLC Tooling stores the following information:
1. Questions used to train instances of Natural Language Classifiers (NLC)
2. Corresponding Classes / Intents
3. Training data for untrained classifiers
4. Test data from executed tests