An Oracle White Paper
April 2011

Information Security:

A Conceptual Architecture Approach

ORACLE®

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 Oracle's products remains at the sole discretion of Oracle.

Introduction

Information is the lifeblood of every organization. If this Information is compromised there can be a wide range of consequences ranging from damage to a company's reputation through to financial penalties such as regulatory fines and cost of remediation. Therefore, why is it that many organizations take a tactical approach to addressing Information Security? All too often, individual projects implement point solutions to address their specific project's requirements without considering the wider implications of security

Information Security is a strategic approach that should be based on a solid, holistic framework encompassing all of an organization's Information Security requirements, not just those of individual projects. The framework should be based on a reference architecture that takes into account key security principles such as 'Defence in Depth' and 'Least Privileges'. By taking this approach to Information Security, organizations can ensure that the components of their Information security architecture address all business critical Information and are driven by the requirements of the business.

Oracle has invested in providing an authorised library of guidelines and reference architectures that will help organisations better plan, execute and manage their enterprise architecture and IT initiatives. The Oracle Reference Architecture (ORA) defines a technology architecture blueprint that includes an Enterprise Security domain. In order to complement the Enterprise Security domain, this document has taken a different viewpoint and looked at Enterprise Security from an Information Security perspective. It focuses on how to protect the most important part of any organisation; the Information. To support this, a top-level, technology agnostic, conceptual architecture that is based on industry security domains and that provides a common architecture and security language has been produced by the document author and contributors. It will show how the architectural concepts associated with Information Security can be represented within a reference architecture.

Information Security Architecture Requirements

All too often within organisations, IT projects are designed to address the core business requirements without thinking about the security of the solution. As a result, security is often only considered as an afterthought and only included during last phases of project design. This leads to security silos within tactical solutions being adopted to address individual security requirements. In addition, as organisations evolve their IT architectures, they must ensure that the security within their IT remains up-to-date and capable of addressing constantly changing security threats, including external threats such as Cyber-Security.

A reference architecture for Information Security helps to address these issue by providing a modular set of reusable security services that can be deployed strategically and reused across all IT projects. As a modular, loosely-coupled architecture, components can be upgraded or replaced as required to meet changing threats without having to re-architect the entire strategic infrastructure.

The domain of Information Security is a mature area with a number of industry recognised best practices and architectural principles already clearly understood, for example, through standards such as ISO27001. Therefore, it is important when producing an Information Security conceptual model that these principles are taken into account and can be tested against the model to ensure it is fit for purpose.

Below are some of the core principles that were taken into account when building the Information Security architectural blueprint. A more complete list of security principles can be found in the ORA Security document referenced on the previous page.

Core Principles of Confidentiality, Integrity and Availability

The three common goals of Confidentiality, Integrity, and Availability are key requirements within Information Security. The definitions of each as provided by Wikipedia[1] are:

> CONFIDENTIALITY - Confidentiality is the term used to prevent disclosure of information to unauthorised individuals or systems.
>
> INTEGRITY - In Information Security, integrity means that data cannot be modified undetectably.
>
> AVAILABILITY - For any information system to serve its purpose, the information must be available when it is needed.

Further details on the three core principles above can be found in the ORA Security document (see Further Reading section at the end of this document.)

Defence in Depth

The principle of Defence in Depth is used to describe the concept of implementing multiple layers of security so that if one layer is breached, the asset being protected, in our case Information, is not compromised. By implementing enough layers of protection the likelihood of compromise is drastically reduced.

One of the key facets of this principle is the heterogeneity of each layer, ensuring that each layer is implemented using distinctly different patterns to another layer. In Information Security terms this could be through using different technologies, encryption standards, protocols, deployment methods, etc.

Comprehensive Domain Coverage

Many industry standards organisations divide the area of Information Security into a number of distinct domains of knowledge. These reflect the different areas of interest within Information Security. The table below highlights some of the key industry standards bodies and the different groupings they apply to Information Security realms.

[1] http://en.wikipedia.org/wiki/Information_security

TABLE 1. SAMPLE KEY REQUIREMENTS FROM INDUSTRY STANDARDS ORGANISATIONS

INDUSTRY BODY	AREA CATEGORISATION	EXAMPLE DOMAINS
International Information Systems Security Certification Consortium (ISC2)	Common Body of Knowledge (CBK) Domains	• Access Control • App Development Security • Cryptography • Operations Security • IS Governance & Risk Mgmt
National Institute of Standards and Technology (NIST)	Topic Clusters	• Authentication • Biometrics • Cryptography • Risk Assessment • Digital Signatures
British Standards Institute (BSI)/International Standards Organisation (ISO) 27001:2005	Control Objectives	• Information Security Policy • Physical and Environmental Security • Backup • Network Security Management • Electronic Commerce Services
Cloud Security Alliance (CSA)	Critical Areas of Focus	• Compliance and Audit • Legal & Electronic Discovery • Traditional Security, Business Continuity and Disaster Recovery • Identity & Access Management • Virtualisation
European Network & Information Security Agency (ENISA)	*Provide Positioning Papers and Reports on a number of key security areas*	• Web 2.0 Security & Privacy • Reputation-based Systems: a security analysis • Recommendations for Online Social Networks

An Information Security architecture must ensure that it has comprehensive coverage of all of the different domains (sometimes called Control Areas). Many of the domains defined by the industry standards bodies are the same or similar and can be grouped into a number of key security requirement areas:

- Confidentiality
- Integrity
- Availability
- User Management
- Network Security
- Key Management
- Security Management
- Governance
- Risk
- Regulation
- Audit
- Access Control
- Standards for Interoperability

Information Security Conceptual Architecture

After defining the requirements that will feed into the conceptual architecture, as detailed in the previous section, the next step is to determine the purpose of the conceptual architecture. At the highest level, it can be represented by Figure 1.

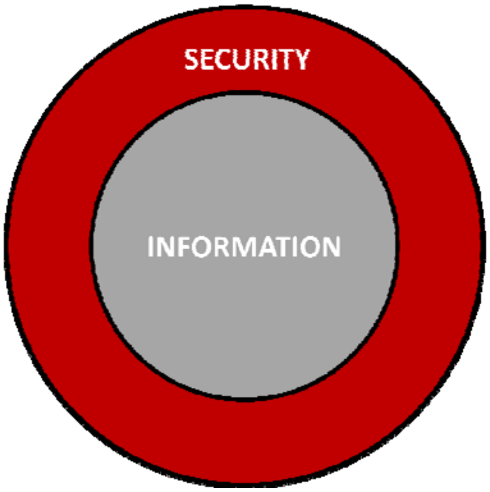

Figure 1 - The Essence of Information Security

As the name suggests Information Security is about providing security to protect an organisation's most sensitive asset, its Information. All of the key requirements defined by the various industry standards bodies in the previous section are all different mechanisms to achieve the security boundary surrounding the Information.

However, whilst this model provides the Information-centric security that is required, how that level of protection is actually achieved is where the next level of detail appears. It is the layers of the Information Security architecture and the relationship between those layers that gives us the ability to implement and manage the security boundary protecting our Information.

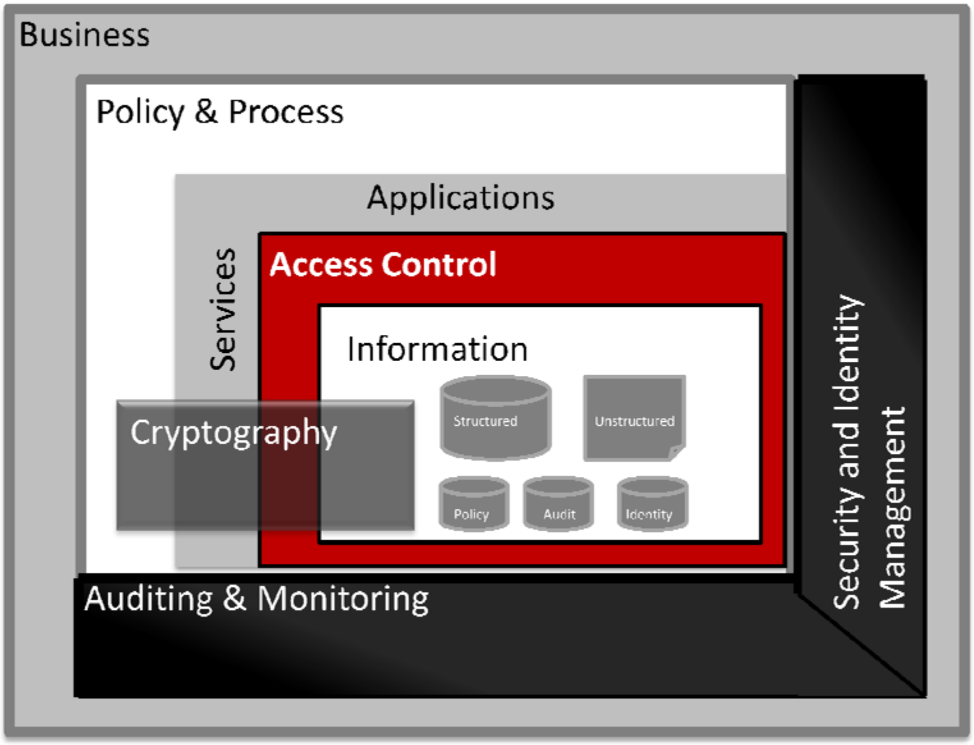

Figure 2 - Information Security Conceptual Architecture

The diagram in Figure 2 represents a conceptual model for the way that we view Information Security. As you can see, **Information** is at the centre of the diagram. As has been discussed previously, Information is the lifeblood of any organisation. This data can be structured or unstructured, has value and therefore is defined as information. In addition to Business Information, the scope of Information also encompasses Policy & Process, Audit, and Identity information, which also needs to be secured.

The next layer surrounding the Information on all sides is **Access Control**. This is a way of showing that all access to the Information is pursuant to controls and secured if necessary, regardless of the method or channel used or where the Information resides

Beyond the Information Access Control layer, it is appropriate to start to look at ways of interacting with the Information. This could be **Applications** or **Services** which are the information processors. This includes not just the traditional business applications (e.g. Siebel etc) or SOA-based services such as (secure) web services, but also includes applications required to manage the Information Access Control, such as a provisioning application.

As has been seen above, the Information Access Control layer controls access to the information for applications and services. The rules specifying that access needs to be defined, subject to the **Policy and Process** layer, that, provides that level of design and the appropriate stipulation of control.

The top layer is the **Business** which provides the feed into the policies (and business processes), determining what the business drivers are that lead to specific information security decisions.

Complementing the layers are conceptual **Security & Identity Management** and **Auditing & Monitoring** services. These provide the capabilities to support the management and governance of the security within the Access Control layers as well as the monitoring and remediation of those enforcement policies.

Finally, there needs to be an underlying capability to support the security mechanics of the architecture. This is delivered through the **Cryptography** domain which is pervasive due to its necessity to be included at multiple layer, from definition within the policy(s), to implementation within the applications and services and the Access Control layers.

The following sections examine each of those components in more detail.

Business

All Information Security decisions must start from the top, which means they must to start with the business. It is the responsibility of the business to decide what controls are required, what is driving this to happen and what the business needs from Information Security. These decisions are governed by a number of different factors and influences. Some of these will be external whilst others will be driven from an internal perspective.

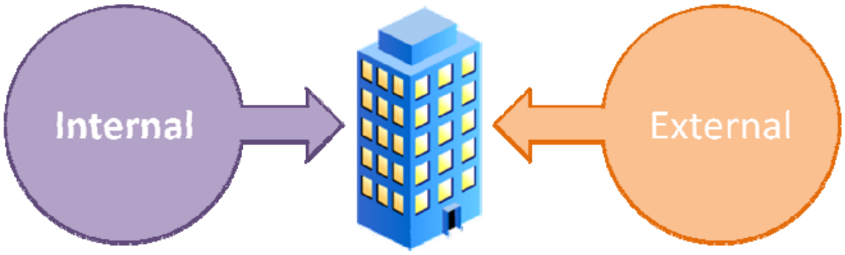

Figure 3 - Factors Affecting Information Security Decisions

The most common **External** factors are laws and regulation, i.e. the legal and regulatory requirements that organisations must conform to. The exact list of requirements applicable to any organisation varies depending on a number of factors such as an organisation's industry and country of operation. However, a few of the common requirements include:

- Data Protection Act
- Sarbanes Oxley (SOX)
- Health Insurance Portability & Accountability Act (HIPAA)
- Payment Card Industry Data Security Standards (PCI/DSS)
- European Data Protection Directive

These legal and regulatory requirements must not only be complied with, but regular audits undertaken to prove compliance against these requirements. However, these are not the only external factors.

Market Conditions also have a strong influence on the Information Security needs of an organisation. This is a broad influencing factor that can take many forms. For example, anti-animal testing campaigners may constantly target life sciences organisations, or anti-war campaigners may target defence contractors. However, the influence of market conditions also changes depending on nations' social and economic stature. For example, in poor economic times (e.g. during a recession), fraud may be more prevalent. This can affect many industries including online retailers and insurance companies. Similarly, during a recession, an increase may be seen in industrial espionage.

Another key external factor is external commercial relationships. Organisations need to form partnerships and in doing so share information to enable effective joint working. This impacts on some of the legal requirements mentioned above but also impacts on security policy. These partnerships drive a need for interoperability standards for ensuring information access can be granted in a controlled manner whilst the information exchange can be controlled and secured. There are many such standards already in existence that cover areas such as:

- Common signing and encryption formats (XML-Signature, XML-Encryption, AES, 3DES)
- Common cross-domain access mechanisms (SAML, WS-Trust, WS-Federation)
- Common cross-domain user management standards (LDAP, SPML)

The need for Information Security to be addressed within these organisational partnerships extends not only within a B2B context but also within Cloud architectures where Information needs securing whilst being made accessible to the appropriate people once the correct agreements are in place.

It is not just external factors that affect an organisation's approach to Information Security, **Internal** factors also have a big influence. A prime example of this is an organisation's route to market. An organisation that sells merchandise only through high-street stores will have differing Information Security attack vectors and challenges than an online retailer. Both will have a common set of core requirements but they will also have requirements specific to their chosen routes to market. Another internal factor is the diversity of staff employed within the organisation. Factors such as education level of staff, use of permanent, temporary or contract workers, on-shore or off-shore workers will all affect the approach to Information Security and education and training programmes that need to be implemented.

All of these factors affect the way an organisation must apply Information Security.

Policy and Process

Taking all of the relevant external and internal factors into account, it is then the organisation's responsibility to specify what Information Security rules must be applied. This is done through a combination of policies and processes (or procedures). It is these policies and processes that provide the rules that govern the 'use' of the Information. As part of the process of determining what threats to protect against within the policies, an organization must first undertake an exercise to assess the value of the Information (via classification). It can then define the security risks to the valuable Information, the consequences of the defined security breach and finally the likelihood of the event happening. All of this Information is fed into the policy ensuring that the policy provides protection for the correct risks.

Any organisation will have a wide range of policies and processes. However, it is the ones that are directly related to Information Security that are taken into account within this document.

The role of Policy is to specify what controls must be put in place. A control is a rule or set of rules designed to counter a particular threat or vulnerability. These controls are typically high-level statements and should not specify the technical details of how a particular control should be implemented. For example, a control may specify:

All data on backup tapes must be encrypted

The policy does not typically stipulate what software or hardware should be used to encrypt the data. Additional controls may specify the strength of the encryption algorithm used or this may be left to the technical implementation. This decision will be governed by the input from the business in the previous layer (for example, depending on legal or regulatory requirements).

Whilst the policy documents specify what controls must be in place, the process documents will typically define the approved way of carrying out specific business practices. For example, this type of document may define the process for a creating a new user and giving them access to all of the appropriate systems and applications that they will need. This may be a manual or automated process but should still be documented.

Policy and Process documents cover a wide range of Information Security topics, some of which may not always be directly associated with Information Security. Table 2 shows some of the documents that you would expect to find within an organisation and its relevance to Information Security.

TABLE 2. EXAMPLE TYPES OF POLICIES AND PROCESSES WITHIN ORGANISATIONS

POLICY & PROCESS TYPE	RELEVANCE TO INFORMATION SECURITY
Security Policy	This will contain all of the organisation's requirements and documented controls for its security, including, for example, access control, data handling, data encryption, data retention, physical security and network security,
Governance, Risk & Compliance	This will typically not be a single document but will be a set of documents containing details of information such as risks, e.g. risk register, risk mitigation etc and compliance requirements. It is also commonly where the requirements for governance boards will be defined, such as an Information Security Governance Board to oversee Information Security within an organisation
Audit Policy	Sometimes contained with the Security Policy, the audit policy will specify the level of audit data that needs to be maintained, for example, in order to meet regulatory requirements. It will also typically specify controls around storage, access, use and destruction of the audit data.
Organisation Structure	This structure of the organisation can be important to Information Security as it can be used to determine actions such as approvals for roles and/or Information access.
Business Continuity Planning/Disaster Recovery	Linking back to our core principles of Information Security defined on page 4, an organisation's BCP and DR strategies are important to Information Security in order to maintain Information integrity and availability in the event of a business critical event.
Development and Deployment Practices	Ensuring that any hardware and software is designed, deployed and managed in an appropriately secure manner is important to maintain the security of the Information that it is storing, processing and presenting to authorised users. This is the same for both internally developed applications as well as COTS[2] based applications
Data Classification	This defines the different levels of sensitivity for an organisation's Information. Each level will be related to sets of policies that define the amount of protection necessary for that classification of data. Examples of categories may include "Internal Only" and "Board Members Only" (see ORA Sec chapter 1.2 for more classification schemes.)

[2] Commercial Off The Shelf (COTS) is the term used to describe software written by 3rd parties software development company and sold to the market.

In addition to the above policies and processes, additional documents may be produced for specific purposes. For example, UK public sector organisations are required to produce Information Assurance documentation in-line with HMG Security Policy Framework[3] and organisations complying with ISO27001:2005 are required to produce an Information Security Management System (ISMS) Policy document.

By defining the appropriate policies and processes, organisations stipulate the rules that govern access to and use of Information. Once these have been established, it is then the responsibility of the organisation to implement the prescribed controls and for the governance function to maintain adherence to said controls.

Applications and Services

Information can only provide business value if it can be accessed when it is needed and by authorised individuals. Information Security should not only be seen as a defence mechanism to protect the data, but, in addition, as a business enabler, ensuring that the Information is available to the right people at the right time. So, how do users gain access to this information?

All access to Information is via an application. Whether it is a business application (e.g. Siebel, Oracle e-Business Suite etc), a middleware component such as an Enterprise Service Bus or a command-line utility, it is the interface that provides the link between the end-user and the Information they are trying to access. Even if a disk containing information is stolen, access to the information on the disk will still need some form of 'application'. In this case, this would probably be a low-level disk reading utility.

Furthermore, applications are not just those components that interface to the 'business' information such as CRM data or HR data, it is also those components that are used to manage the system information. For example, Database Management tools or Identity Management tools. All of these components require access to Information that is sensitive to the organisation and therefore needs security.

Applications and Services can be divided into three types (see ORA Sec chapter 1 for finer-grained definition); Information Consumers, Information Producers and Information Providers. Information consumers may access a wide range of Information types but are typically consuming the information. Information Providers are accessing Information to pass onto the user's requesting it. For example, this could be a customer accessing their order history via an online retailer's website. Information Producers are the applications and services that are actively creating, updating and deleting the information. Again, they can require access to many types of information. Applications and services may function as Information consumers, providers and/or Information producers.

When considering an organisation's Information Security requirements, it is important to understand the different applications and services that require access to the information and furthermore, the type of interaction required, i.e. consumer, provider or both.

[3] http://www.cabinetoffice.gov.uk/content/government-security/

Access Control

Access Control is at the heart of Information Security. It is the red ring of security surrounding an organisation's Information. All access to that information MUST flow through the Access Control layer irrespective of whether the request is internal or external, business application or command-line utility, employee or customer. No information can be accessed without first passing the access control layer.

The implementation of the access control layer can (and will) exist within multiple places within an organisation - in accordance with the principle of "Defence in Depth". The security may be built into the applications and services or externalised within the access control layer itself. However, it is important to understand that it is not just talking about logical (IT) security such as firewalls, access management and database security systems at this layer. Physical security, such as building entry, CCTV cameras and security guards is also being referenced at this layer.

When thinking about Access Control, it is also important to understand where the access is coming from. Here, we are not just talking about local users sat in an organisation's office accessing internal Information. We are also talking about more modern architectures such as Cloud architectures where the Information could be sat within a 3rd party organisation and may be accessed by internal users and/or external users. Access Control must still play a key role in controlling access to this Information. As described in the Business section, there are many standards, such as SAML, that can be adopted to help implement and enforce this level of access control.

In reality, regardless of how many different layers access control is represented as, it is about providing two key functions: Authentication and Authorisation.

Authentication is the process of validating an asserted claim to an identity. Before a decision can be made to see if a user should be allowed to access Information it is a pre-requisite to need to know who that user is. Whilst a user can provide their identification, it is of no use until it has been suitably assured. This is done through the process of authentication. The level of assurance will be dependent on the sensitivity of the Information the user is trying to access. The table below shows two different types of Information and the levels of authentication that may be required at both the physical and logical access control layers.

TABLE 3. EXAMPLE AUTHENTICATION TYPES FOR DIFFERENT LEVELS OF INFORMATION

TYPE OF INFORMATION	PHYSICAL SECURITY	LOGICAL (IT) SECURITY
Corporate white pages directory in a company office	Verified Employee ID Badge	Username and Password
Corporate product designs in a data centre	ID Badge + Biometric verification	Biometric or suitable strong authentication
Corporate Marketing Material on a public website	None	Unauthenticated (anonymous) access

Also note that an identity may also be assured through trust between organisations. For example, in physical security terms, a contractor may be able to gain access to a data centre to carry out some maintenance if they show their ID card, asserting their identity from the maintenance company. In logical security terms, their identity may be assured through mechanisms such as Federation where their Identity Provider asserts their identity to the organisation providing a service to them. Once a suitably assured identity has been obtained, Authorisation can be checked.

Authorisation is the process of deciding whether the assured identity is allowed to access the information that they have requested and are allowed to perform the requested action on that information. There are many ways that this can be implemented. For physical security this could be through allocation of certain building zones so that the contractor's ID card can only be used to open certain doors. In logical security, this could be through an application checking its roles and responsibilities to ensure they have been assigned the correct permissions or asking the same question to a centralised policy decision point. However the decision is made, it is the access control layer that enforced the authorisation decision, ultimately allowing or denying them access to the Information that they have requested.

Of course, the reference architecture is not trying to represent access control as only protecting information within the perimeter. As a conceptual model it is showing that access control surrounds the Information irrespective of where it sits. What about when the Information is accessed by an authorised user and then taken out of an organisation's control (e.g. downloaded to a spreadsheet or emailed to an external party etc)? Information Security is not constrained to protecting access to Information whilst it is within an organisations control. It is also important to provide the policies, processes and controls for the Information once it has left our control.

Information

Within this architectural model, Information is at the heart of the architecture. It is the business sensitive information that an organisation needs to protect. As expected, this information includes the types of data traditionally associated with business sensitivity. For example, CRM databases containing customer data, HR databases containing employee data and finance databases containing accounting data. All of this data falls into the category of structured data. This refers to data held in a structured form such as a database. However, there are large amounts of information within organisations that contain business value and therefore need protecting but aren't contained in a structured way.

For example, a organisation's 'Go To Market' strategy may be held in a word processing document, sensitive product designs and diagrams may be held in images or bespoke applications formats. In addition, structured data that has been exported from the database for local reporting (e.g. into a spreadsheet) is also in need of protection.

Beyond pure 'business' data, there are other sets of information that are crucially important to any organisation and therefore need consideration in respect of an Information Security architecture. These include the following types of information:

- Policy
- Audit
- Identity

Policy information contains the rules that govern Information Security access. Typically these will be rules implemented within a variety of policy decision points and enforced through policy enforcement points. Failure to protect this information sufficiently can mean changes that would allow access to information by unauthorised users.

Audit information is often a critical component of any organisation's compliance team. It provides the visibility of key security operations such as authentication attempts, authorisation attempts, policy changes and identity changes. Failure to protect this information can result in users altering the data to hide activity such as unauthorised access to data or changes to identities and their roles.

Identity information is key to maintaining a strong Information Security architecture. It provides the knowledge on who a user is and what they are entitled to do. It also provides the visibility of where users have different identities and accounts throughout an organisation. It must be protected from unauthorised changes in order to ensure correct business processes are followed when managing an Identity and its access.

One of the key challenges for Information Security is control for all of the above data regardless of where it resides. Traditional security frameworks provide a security boundary around an organisation and implement tight controls to ensure that users entering that boundary are controlled. However, as highlighted in the previous section, with increased de-perimeterisation and increased partnerships and information sharing, it is critical that Information can be controlled beyond the traditional organisational IT boundaries.

At the opposite end of the scale is protection of the information from inside threats from both standard end-users as well as power-users (e.g. DBAs). The following output from IOUG's 2010 Data Security Report shows some of the challenges faced by organisations when trying to protect their data from internal threats.

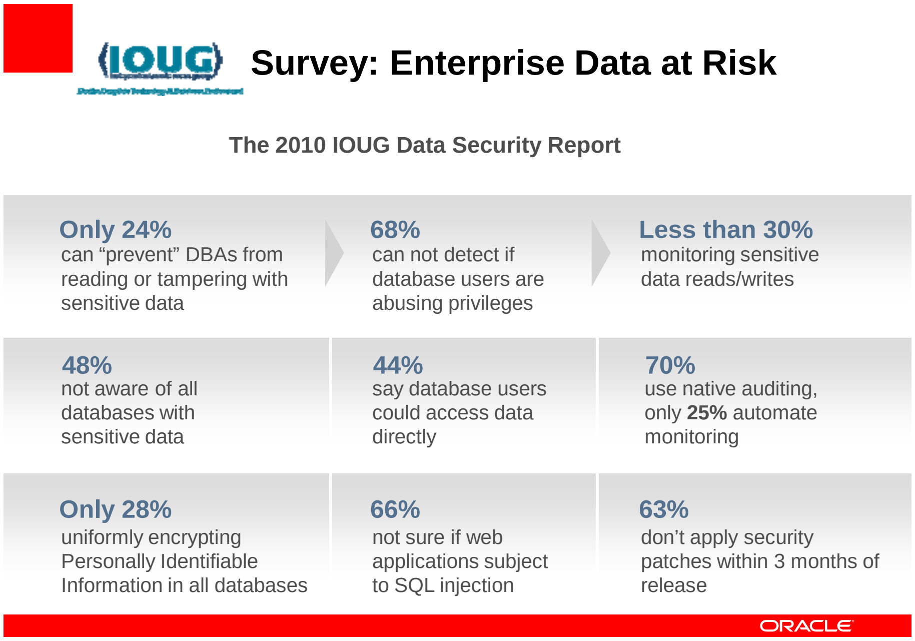

IOUG Survey: Enterprise Data at Risk

The 2010 IOUG Data Security Report

Only 24% can “prevent” DBAs from reading or tampering with sensitive data

68% can not detect if database users are abusing privileges

Less than 30% monitoring sensitive data reads/writes

48% not aware of all databases with sensitive data

44% say database users could access data directly

70% use native auditing, only **25%** automate monitoring

Only 28% uniformly encrypting Personally Identifiable Information in all databases

66% not sure if web applications subject to SQL injection

63% don’t apply security patches within 3 months of release

ORACLE

Figure 4 - IOUG Data Security Report

Security and Identity Management

So far within the architecture all of the components to secure the Information have been delivered. However, no consideration has yet been made as to how security controls protecting and enabling access to the Information will be implemented and managed. For example, the Access Control layer may include one or more Policy Enforcement Points (PEP), but where are the Policy Administration Point (PAP) and Policy Decision Points (PDP) that determine the policies to enforce and respond to requests from the PEPs for security decisions?

Security and Identity Management is the container for these components. It is responsible for providing the ability to:

1. Implement the access control policies that are enforced by the Access Control layer
2. Manage identities
3. Manage the access control
4. Convert the security policy into usable access control

As well as access control tools, this component in the architecture is also where user management tools such as Identity Management capabilities will reside. Whilst not strictly mandatory for pure access control, Identity Management provides a mechanism to improve the efficiency of the processes around identity creation, modification and removal, commonly referred to as the "Joiner, Mover, Leaver" process. Identity Management can also be used to automate the assignment of permissions to users to enable them access to the Information through the Access Control component. Open standards can play a key part in Security and User Management. For example, standards such as SPML can be used to exchange Identity information and entitlements in a common format. This is especially applicable to Cloud architectures where direct access to applications and their underlying Identity repositories may not be available or where a standard method of user management is required.

When defining security policy and managing identities, the Security and User Management tools will require access to the Policy and Identity repositories within the Information layer. Of course, this access is not permitted directly. As with ALL access to Information, it must pass through the Access Control layer. Therefore, at this point the security and user management tools are no different to the applications and services that reside in the Applications and Services Layer.

Audit and Monitoring

The Auditing and Monitoring layer is responsible for recording all of required security-related operations to maintain a non-repudiated and tamper-evident trail of evidence. The level of information that needs auditing will be determined within the Policy and Processes. This may include:

- Information Access
  - e.g. Successful and/or Failed Authentication/Authorisation attempts
  - e.g. Failed access attempts to sensitive tables within databases by privileged user
- Policy Administration
  - e.g. Changes to access control policies
- User Administration
  - e.g. Addition/Removal of users' roles and privileges
- Information Change
  - e.g. Changes to sensitive Information

As with Security and Identity Management, once a policy has determined that a particular operation needs auditing, then each operation will be recorded within the audit repository in the Information layer. As before, this repository must be accessed through the Access Control layer.

Another important characteristic of this layer is monitoring. Ensuring availability of the Information is a key security principle of the Confidentiality, Integrity and Availability triad. Therefore, monitoring is used to ensure that all of the necessary security components are running in-line with the Service Level Agreements (as defined by Policy and Process).

Cryptography

At the heart of many IT-related security controls is cryptography, the process of using complex maths and algorithms involving large numbers, to protect information through enciphering and deciphering messages. This can be used for a number of different purposes including, maintaining the confidentiality of Information through encryption and the integrity of Information through digital signing.

Within the reference architecture Cryptography has been shown spanning multiple layers from Policy and Process all the way through to Information. This is because cryptography is pervasive across all of these layers. Table 4 shows the relevancy of cryptography to each layer of the model.

TABLE 4 - RELEVANCE OF CRYPTOGRAPHY IN EACH LAYER

LAYER	RELEVANCE	EXAMPLES
Policy and Process	Determines what level of cryptography is required	Key lengths, algorithms
Access Control	Implements cryptography for authentication	X.509 Authentication
Applications and Services	Implements cryptography for data in transit and data at rest	SSL, TLS
Information	Implements cryptography to secure data at rest	AES, RSA
Security & Identity Management[4]	Manages cryptographic keys and certificates	X.509 user certs, SSL certs, PKI keys

Validation of the Architecture

When defining a conceptual Information Security architecture it is important that it is validated against a number of key Information Security principles to ensure that it does not omit any capabilities that are required. Within this section the model is tested against three key facets of Information Security.

Validation against Confidentiality, Integrity & Availability

As defined earlier in the document, the CIA triad one of the core principles of Information Security. Therefore, how can it be ensured that the architecture addresses these three key factors?

It is the responsibility of ***Policy and Process*** to define the requirements for CIA within the organisation. . For example, the different levels of data classification, the metrics for SLAs, the data encryption policies etc. This is led from the ***Business,*** determining the level that the Information needs to be covered by CIA. The policies are then implemented through the ***Security and Identity Management*** and enforced by a combination of ***Access Control*** and ***Cryptography***. In addition, ***Monitoring*** is used to ensure that the availability of the ***Information*** is maintained in accordance with the SLAs.

[4] Not shown in reference architecture model

Validation against Defence in Depth

The next principle to validate our architecture against is Defence in Depth. This is ensuring an 'onion skin' style approach to security by not relying on a single component or layer to protect an organisation's Information, but instead providing multiple different levels of protection. Conceptually, this is shown within the architecture model as a single ***Access Control*** layer. However, the key to this layer is that it will be implemented at many layers irrespective of the method of access to the Information or where the Information is located.

To validate this principle it is important to understand the different locations that an organisation's Information may reside and therefore the different levels that the Access Control layer(s) need to work at. The diagram in figure 5 below shows these layers.

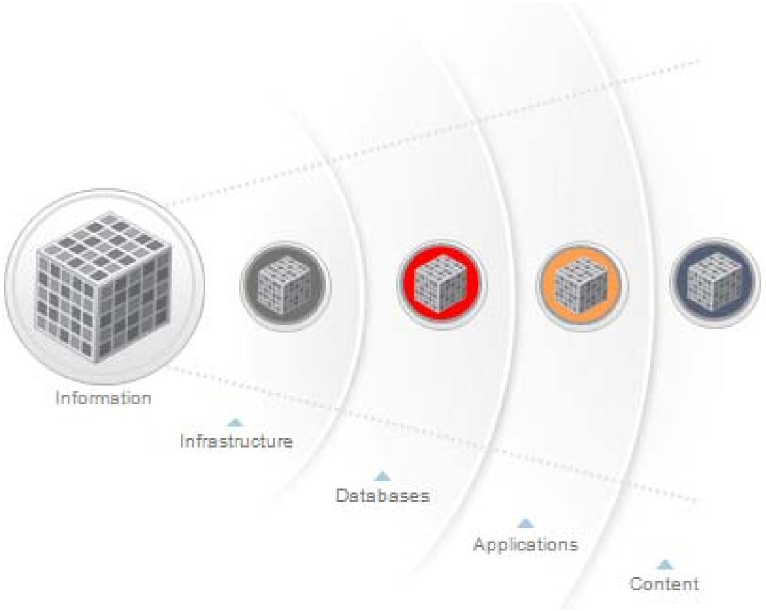

Figure 5 - Layers of Information Access

Table 5 below explains each of the layers from figure 5 and the type of defence in depth controls that can be implemented within ***Access Control***.

TABLE 5. LOCATIONS OF INFORMATION AND DEFENCE IN DEPTH ACCESS CONTROLS

LAYER	DESCRIPTION	EXAMPLE ACCESS CONTROLS
Infrastructure	This is the data sat on physical media (e.g. tapes and disks) within data centres as well as the networking components	• Security Guards • Door Locks, Rack Locks • Door Entry Systems • Disk Encryption • Network Encryption
Database	This covers an organisations sensitive Information when it is residing in a database, e.g. CRM, ERP, HR Information	• Database Access Controls • Database Encryption • Data Classification • Segregation of Duty • Auditing
Applications	Applications covers the Information whilst it is being processed through Applications and/or Services	• Role-Based Access Control • Authentication • Auditing
Content	This covers the Information when it is beyond the application. For example, exported data in a spreadsheet or an organisation's product design documents. This layer of access of relevant regardless of whether the Information is internal or external	• Information Rights Management • Endpoint Access Control

As a proof point, all of these access controls are included within the reference architecture model.

Validation against Governance, Risk & Compliance

Governance, Risk and Compliance (GRC) has many of the same requirements as Confidentiality, Integrity and Availability. The essence of both principles require the implementation of Policies and Processes. Specifically for GRC, based on the country or sector an organisation operates in, the controls required are based upon ***Business*** factors and context, and defined by ***Policy & Process***. GRC is then implemented using ***Access Control*** and ***Audit and Monitoring.***

Conclusion

The purpose of this document is to expand on the common industry Information Security domains and provide an in-depth view on how the architectural concepts associated with Information Security can be represented within a re-usable reference architecture.

The document is positioned as a discussion document outlining the building blocks required when considering the topic of Information Security, but is by no means an exhaustive list of the controls, frameworks or challenges related to Information Security.

Our aim is to discuss the importance of providing an end-to-end, defence in depth enterprise-wide Information Security architecture with practical proof points, to meet both business and IT requirements for control as well as enabling the organisation to meet their desired goals.

Further Reading - Oracle Reference Architecture

Oracle has invested in an authorised library[5] of guidelines and reference architectures that will help organisations better plan, execute and manage their enterprise architecture and IT initiatives.

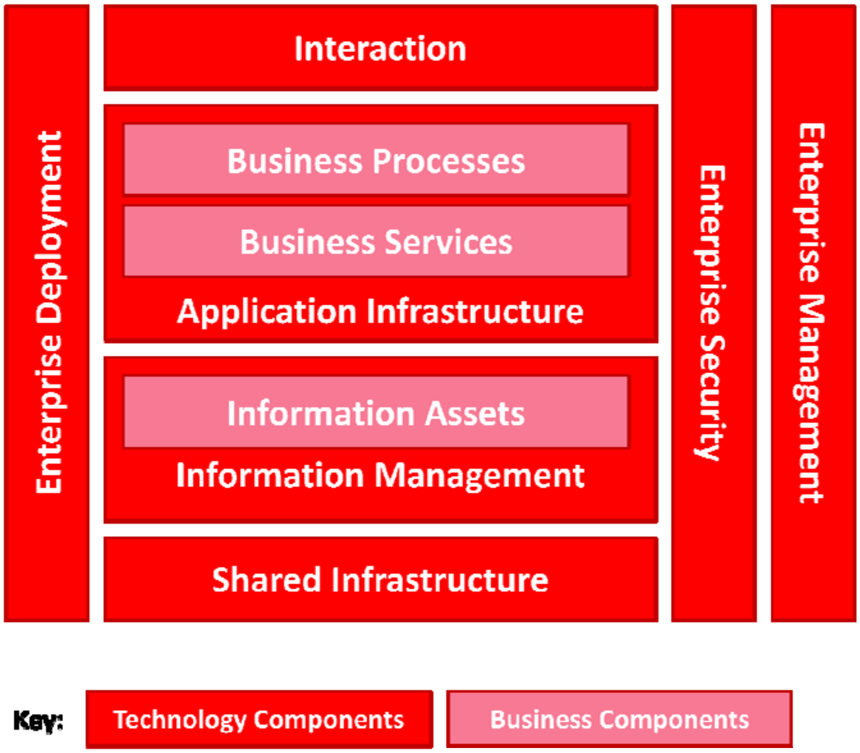

Figure 6 - Oracle Reference Architecture

Within ORA, one of the key domains is Enterprise Security. As can be seen in Figure 6, Enterprise Security is a pervasive domain, spanning all of the layers of ORA. The ORA Security document provides a detailed description of common security concerns, descriptions of relevant industry security standards, as well as a reference architecture for designing an enterprise security framework.

In order to complement the Enterprise Security domain, this document has taken a different viewpoint and looked at Enterprise Security from an Information Security perspective. It has focused on how to protect the most important part of any organisation; the Information. To support this, a top-level, technology agnostic, conceptual architecture that is based on industry security domains and that provides a common

[5] http://www.oracle.com/goto/itstrategies

architecture and security language has been produced by the document author and contributors. It shows how the architectural concepts associated with Information Security can be represented within a reference architecture.

Alignment to ORA Security Architecture Principles

The ORA Security Architecture document referenced within the security architecture requirements section of this document provides a comprehensive list of security principles that should be applied to a security architecture. The security architecture requirements section of this document addresses some of these as inputs into the architecture design.

Below is an example of some of the Information Security architecture principles that Oracle software is based upon. This list is derived from Oracle's own expertise as well as industry best practice and documented to enable organisations to use them as a basis for their own Information Security architectures.

- **Defence in Depth** - Failure of a single component of the security architecture must not compromise the entire IT environment.
- **Least Privilege** - Users and other consumers of resources must operate using the least set of privileges necessary to complete the job.
- **Security as a Service** - Business solutions must be designed to consume common security services where possible as opposed to implementing custom security logic and replicating copies of security data.
- **Identity Federation** - Security infrastructure must provide identity mapping, credential mapping, and identity propagation necessary to support federation across security domains.
- **Secure Web Services** - End-to-end system security and adherence to other security principles must not be compromised by the use of Web Services.
- **Secure Management of Security Information** - Security information such as users, credentials, groups, roles, attributes, must be managed centrally (holistically) across the entire organization in a secure and auditable manner.
- **Active Threat Detection & Analysis -** Security infrastructure must be capable of detecting abnormal behavior and adapting accordingly to protect vulnerable resources.
- **Secure, Complete Audit Trail** - The security system must be able to identify when changes have been made to data within the organization, what changes have been made, and by whom.
- **Data Security** - The confidentiality, integrity, and availability of data must be ensured at all times.
- **System Availability** - Systems must be adequately protected to ensure their intended degree of availability, but not overly constrained by security measures to unnecessarily impede normal operations.

Information Security - Architectural Approach
Apr 2011
Author: Paul Toal
Contributors: Angus Herron, Jason Rees,
Patrick McLaughlin and Dale Young
Reviewer: Dave Chappelle

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Oracle is committed to developing practices and products that help protect the environment