Table 1: Privacy Design Guidance

ID / Category / Related Requirements	Guidance	Comments and Rationale
AC1 Architectural Concepts Enforce, Monitor, Authorize	Concept: Privacy Policy for real-time privacy decisions in an IT environment should be centrally managed and locally enforced.	• This concept is targeted at making real-time decisions consistent with privacy policy versus static manifestations of privacy policy such as user interface design. • Central and local are relative concepts here. Central might mean at the ministry or Program level for GoA and local might be just the applications and databases within the ministry. • Central privacy policy management provides consistency and the ability to implement cross-enterprise policy changes quickly and seamlessly. • Central privacy policy management also positions an enterprise for central reporting and audit (including the effectiveness of local enforcement). • Local enforcement improves response time for the user and balances work across IT resources. Network dependencies and bottlenecks are also avoided.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AC2 Architectural Concepts Enforce, Monitor, Authorize	Concept: Define "Privacy Boundaries" within the infrastructure that encompass all the places where PI is stored.	• This is really a definition/documentation activity that says objects within a privacy boundary, and transactions crossing a privacy boundary, will be treated a particular way. • For example, everywhere a data access request crosses the privacy boundary; a privacy conformance check should be done. • Very much like Security Zones of Control. An architecture would need to define what these boundaries are, and what controls would be implemented • Privacy boundaries are subsets of security zones. PI is placed within the security zone that affords it the required protection. Within this security zone is a privacy boundary (or boundaries) surrounding all PI. The only considerations for data crossing the privacy boundary are privacy ones because the security requirements have been taken care of. • Note, there is not much value in defining privacy boundaries until some sort of active control for privacy is actually implemented, as such, this concept for GoA might be just for future consideration
AC3 Architectural Concepts Acquisition	Concept: Privacy management should be handled by a specialized set of common components that provide services to existing elements of an IT infrastructure.	• Common components remove the burden for applications having to understand specific database technology or having to be privacy-aware. Aside from the efficiency and consistency of code re-use, applications will be independent of changes to database technology or privacy policy implementation. • Common components also provide a control-point through which all requests for PI can be funneled and audited • Clearly, these components need to be flexible enough to fit into the wide variety of application and data structures that exist in typical IT infrastructures. This would include everything from modules that might fit inside application and databases to components that intercept privacy-relevant requests between applications and databases.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AC4 Architectural Concepts Placement (Related to AC3)	Concept: Define an "Enterprise Application Layer" or an "Application Framework" to instantiate the common services and components for data access and privacy compliance.	• The Enterprise Application Layer approach physically puts the common services in a separate layer that applications call. (For example, "data access beans" in a Java environment or a "data abstraction API" in a native code environment). • The Application Framework approach provides common components that can be plugged into applications. (ex: servlet filter plug-ins or a STRUTS web application design). • The concepts are not mutually exclusive • The GAEA Application Architecture already has accepted the concept of an application layer and within that a data services layer so this is really just a point to support adoption of that concept.
AC5 Architectural Concepts Enforce, Monitor	Concept: For each element of PI, it should always be possible to determine the privacy policy that governs the collection, use, disclosure and retention of the PI. (The "Sticky Policy Paradigm")	• Privacy-related actions on PI should be made in the context of the Privacy Policy **currently** governing the PI - without this capability, privacy obligations cannot be assured of being met by definition • In many cases this will be the policy that was in force at the time of PI collection • However, if there has been a change in policy since collection, the policy that will apply depends on the organization's practices. Some organizations choose to have the new policy supercede the old one regardless of what the individual may wish. Others will just apply new policy to any subsequent PI collection leaving previously collected PI under the old policy. • Since GoA usually collects and uses PI under specific legal authority, the former case is likely to apply most of the time.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD1 Application Design Enforce, Monitor, Authorize (Related to AC3)	Concept: Applications should not implement logic to process real-time privacy decisions – these decisions will be processed by specialized common components (i.e. applications do not have to be "privacy aware")	• Real-time privacy decisions are those that need making dynamically against the privacy policy versus static manifestations of the privacy policy such as user interface design that probably will be implemented by the application. • Applications that imbed vs. call a specialized common component are consistent with this concept. • Applications do have an indirect role to play in real-time privacy decisions but they do not need to be "privacy aware" to do this – at a minimum they just need to supply context parameter values that can be mapped into privacy terms and must be able to tolerate a denial of access to PI if that is the outcome of the externally made privacy policy decision. • With this design point, application function does not have to be changed when privacy policy is changed. • There may need to be an interim state where applications do process some privacy logic. In this case they should mimic the relevant component function so that the eventual transition will be smoother.
AD2 Application Design Enforce, Monitor, Authorize (Related to AD1)	Concept: Applications should generate values for parameters that provide context for PI access requests	• Applications should be capable of supplying these values to accompany a request to access PI or of making them available when required (push or pull). • A possibility for the pull model is an Application Context Object, which has pointers to call-back[mts1] to an application for retrieving privacy relevant information on what the application is currently doing.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD3 Application Design Enforce, Monitor, Authorize (Related to AD2)	Consideration: Applications should be able to supply or make available the following values for parameters that provide context for PI access requests: Unique codes to identify both the application and the type of work that the application is doing at the time the access to PI is being requested. A key that can be used to determine the role of the person (or application) requesting access to the PI (ex: an authenticated credential) If relevant for the transaction, a key that determines the identity of the person who's PI is being requested (e.g. an IID)	• These are not abstract privacy notions but rather are concepts that the application knows and therefore can be expected to supply. • The application and type-of-work codes enable privacy enforcement systems to derive the business purpose of a PI request. • The key used to determine the identity of the person is ideally an MBUN (e.g. an IID) - a 'meaningless but unique number'.
AD4 Application Design Placement, Acquisition	Concept: Design applications such that unencrypted PI is never permanently stored on the client.	• Client workstations are generally less tightly controlled and more vulnerable than servers. Mobile devices such as laptops and PDA's are even less tightly controlled and potentially more vulnerable – the concept may be strengthened to say that no transient PI is stored on these devices either. • This concept is targeted at enterprise applications that manage a store of centrally collected PI (ex: an HR system) • For example, all temporary caches of PI used during a session should be erased when the session is ended • There may be good reasons to store some PI on the client (sometimes only on the client) – for example biometrics. But if it is, it should be encrypted. • If PI is stored on a smartcard it should be segregated into records corresponding to the most granular use of the PI and protected by a personal firewall (i.e. the cardholder controls release).

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD5 Application Design Acquisition	Concept: Destroy transient PI at the end of each session	• If an application creates a temporary store of PI it should be explicitly erased by the application at the end of the session (and not left to some other generic utility. This is especially important if the transient PI is created on a client device (AD4) • Examples of transient PI are local copies used to improve performance or interim PI created to arrive at some conclusion
AD6 Application Design Acquisition	Concept: Minimize the transmission of PI	• Transmission is a vulnerable point in data handling • This is a major design point for deciding which PI processing is done on the client vs. the server • Sending PI via email is especially vulnerable. Alternatives should be considered for sensitive PI such as sending an email notification that the PI is available and providing a link to a password-protected site where the PI can be viewed. • Alternatively, encryption and receipt acknowledgement features can be used • Some managed email systems provide a “do not forward” function which could be used if email is used to transmit sensitive PI • This concept also provides something of a control point on sharing of PI – once PI is transmitted to another system, it is usually going to a less controlled environment from the perspective of the sender and thus the sender has less control on subsequent sharing.
AD7 Application Design Acquisition	Concept: Process PI in the location that minimizes the transmission of PI	• This is basically an implication of AD6. If most of the PI is held at the host/server location, processing of PI should occur there and only the results transmitted to the client (often results do not have to contain PI, they can be just a confirmation). In a smartcard scenario this would imply a token-based implementation where the card just provides credentials or assertions, which allow the host application to process the PI it already possesses. • Classes of secure clients, such as smartcards, are starting to emerge where the reference PI can be stored securely on the client side. This is the converse case where as much processing as possible should be done on the card and only the results transmitted to the host/server.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD8 Application Design Acquisition	Concept: Use strong protection when transmitting PI	• If PI must be transmitted then it should be transmitted using strong protection. • In an IP environment this means SSL encryption • For web data collection forms passing PI this means use of "POST" vs. "GET" techniques • This concept can be tied into the security architecture, which may prescribe additional measures for transmission of very sensitive information.
AD9 Application Design Keys, Acquisition	Concept: Application transactions should be designed to use the maximum reasonable anonymity (or require minimum necessary identity)	• Self-serve functions for providing information etc. can often be anonymous. • In many cases, especially in a web environment, pseudonyms (e.g. NIDs) are sufficient. In other words, just being able to establish that the same individual is being interacted with over multiple sessions is sufficient and the actual identity of the individual is not required. • This concept could also be applied within a chain of processing and not just at an entry point. Some downstream functions may not require identity to be revealed. • The level of identity that is required may be determined by the need for non-repudiation regarding a given transaction. If non-repudiation is required then identity is required.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD10 Application Design Keys	Concept: Do not use PI to match identity. Use an internal key (i.e. an IID) to determine whether two individuals are the same person.	• The concept is for a specialized privacy-enabled component or service to link internal keys (e.g. an IID in GAEA) to all PI as it is submitted (DD1). Once the internal key (e.g. IID) has been linked, there is no need for downstream applications that run business processes to attempt to match identity. If the internal keys are the same then the PI relates to the same person. • This effectively enables de-identified sub-processes within the larger business-processes. • The registration process is the exception to this concept, as identity must be initially established for the internal key scheme to have integrity. However, the registration applications are specialized components and so other applications should be able to conform with this concept.
AD11 Application Design Authorization	Concept: Application functions accessing and manipulating PI for routine activities should be well-defined, structured (and not user-configurable) so there is a predicable output of PI consistent with “need to know”	• This helps prevent a person in a job role that does not need open-ended analytical capabilities, such as searches on data, from accessing PI outside the purpose defined by their role. • If exceptions are needed, separate functions with separate authorization (AD12) can be provided. Separation makes audit of these exceptional activities easier to accomplish. • This allows the parameters for evaluating the privacy significance of the function to be fixed and therefore quicker and simpler to process. • Avoid general-purpose query languages (such as SQL) on PI data for production applications. Analytic applications that require such general-purpose queries should be separate and carefully controlled – or at least the function should require higher authority (AD12) • General-purpose queries should be limited to outputting de-identified PI wherever possible.

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD12 Application Design Authorization	Consideration: Authorization for analytical functions that allow query, search and output of multiple PI records should be more restrictive than those that allow access to individual PI records.	• Potential privacy abuse, liability and trust damage takes a quantum leap once it involves a file of multiple PI records vs. a single record. • There are many job functions that only require access to the PI of the person being served and so list functions should not be provided to people in these roles. Audit trails too are more specific with record-by-record access. • The counterbalancing consideration is that sometimes job roles are not split by routine vs. analytical tasks and so the overhead of separately managing these capabilities would not be justified. • If possible, analytical functions should be restricted to operations on attribute vs. identifier information. If the information is strongly anonymized then restriction may not be required at all.
AD13 Application Design Authorization	Consideration: Tie PI access to specific requests	• Many job roles that require access to PI only require access to a single record corresponding to the client making the current request. • In these situations, efforts should be made to tie access to the request by systematic means to prevent people prospecting through a database to look at records of people they know, celebrities etc. • For example, many Call Centre IVR applications are capable of using caller-id to bring up the customer record of the caller. If this is all that the Call Centre Agent needs to perform their role then they don't need a general access function (or if they do it would be easy to audit against PI access requests that were not driven by requests)
AD14 Application Design (Web) Placement	Concept: Do not use persistent "cookies" to store PI	• Supports the Privacy Principle of security • Persistent cookies are not secure repositories for PI • Many of the uses for cookies, such as web site optimization or personalization, can be done without associating the identity of the user so this concept should not be onerous to adhere to

ID / Category / Related Requirements	Guidance	Comments and Rationale
AD15 Application Design (Web) Placement, UI	Concept: Do not use hidden tracking mechanisms (e.g. web beacons) to collect PI and minimize the use of non-obvious mechanisms (e.g. cookies) to collect PI.	• Supports the Privacy Principles of openness, limited collection • Hidden tracking mechanisms would basically mean web beacons for the web environment – users do not have the tools to know when web beacons are being used • Non-obvious tracking mechanisms would basically mean (session) cookies for the web environment – users do have the tools to know when cookies are being used and can control them, however, notices to alert users to cookie use are often not effective
AD16 Application Design UI	Concept: Real-time collaboration applications should clearly communicate what information they relate to whom. By default, they should implement reciprocal visibility.	• For example: chatrooms, teamrooms, videoconferencing, instant messaging, calendaring etc. • This is to avoid the perception of surveillance because it says a user can always determine who sees them or sees information about them. • Reciprocal visibility says, “I can only obtain information on people that can obtain information about me". This is suggested as a default because it is what most users would assume is the case and is an intuitive way of understanding who can see you or your information. • An example would be an enterprise instant messaging system, which provides an option for the user to restrict who can see them if they are online. Such a function should have a default reciprocal design such that a user also cannot see anyone who they have restricted from seeing them. If the other party does not care then they could change the default and allow all to see them. • Similar situations occur in other collaboration tools like calendaring or teamrooms. The blind carbon copy function in many email systems may have valid business uses, but if used as part of a process to communicate to individuals it is questionable as to whether the notice and openness Privacy Principles have been observed.

ID / Category / Related Requirements	Guidance	Comments and Rationale
UI1 User Interface	Consideration: Collect PI in context – collect only the PI needed at the point-in-time that it is needed	• Incremental collection of PI whenever it is needed is perceived as less privacy invasive than collection of a larger set up front • This also supports the Privacy Principle of Notice since by matching up collection with granular purposes, the uses are much clearer to the Data Subject • Clearly this is most applicable when there are multiple optional application/process paths that a Data Subject may take – with correspondingly different PI requirements. • If all PI is ultimately going to be required, even if not immediately, and if collection closer to the time of use is not convenient for the individual or the program, then up-front collection is acceptable.
UI2 User Interface	Concept: Clearly distinguish PI collection fields that are optional from those that are required.	• Supports the Privacy Principles of openness and limited collection • Optional PI is not so common in a GoA context but may occur where service choices are provided (ex: “Please supply your email address if you wish to be notified by email”). • Needs to be accompanied by a corresponding description of the additional services or benefits that the individual will realize if they provide the additional PI.
UI3 User Interface	Concept: Employ validation checks when collecting PI that are commensurate with the consequences to the Data Subject of processing inaccurate data	• Supports the Privacy Principle of accuracy. • If consequences of inaccuracy are severe, then validation checking should be extensive. • Examples of validation techniques include format checking (ex: telephone numbers), confirmation (ex: “enter new password twice”) or checks against normal value ranges or existing data

ID / Category / Related Requirements	Guidance	Comments and Rationale
UI4 User Interface	Concept: Minimize the use of free-form input fields into which PI may be mistakenly entered and provide user guidance whenever they have to be used	• Free form fields for collecting comments etc. can be a source of liability if the responder mistakenly includes PI in their input. For example, a web survey responder could mistakenly enter their name and address into free form fields – then find this information published openly along with the other information in their survey response. • This includes direct input by the data subject and input about a data subject made by a data user • Ideally the PI collection should be broken down into structured fields of identifiers and attributes
UI5 User Interface	Concept: Provide a link to the organization's privacy statement on the home page and on every PI collection page	• Supports the Privacy Principle of openness • If it is easy, putting a link to the privacy statement on every page as part of a standard header/footer, is a good idea. • It is also suggested that a print function be provided for those wishing to retain a copy of the policy
UI6 User Interface	Consideration: Notify the user when they select a link that will take them to a site not covered by the current privacy policy	• Supports the Privacy Principles of openness • Would typically be used for links to 3rd party sites • Counterbalancing consideration is that it is more complex to set up and maintain and may be annoying for some users
UI7 User Interface Consent	Concept: On PI collection forms, for each purpose where consent is required: Set the initial consent value to "no consent" Provide information on the consequences of not providing consent Provide information on any options for subsequently withdrawing consent Do not allow the user to proceed without providing consent Store the result and context for the consent (individual's identity, time, date etc.)	• Consent for separate purposes should not be "bundled" into a single consent statement – each purpose should have its own consent. • Consent should not be derived from authentication activity either – establishing authentication and obtaining consent should be separate activities • Note that in a GoA context, consent is often not required as legislation provides the authority for collection. • These measures support implementation of obtaining express consent through an IT interface • If appropriate, the result and context for an individual who ends up not providing consent should also be stored • This is really intended for situations where the service or feature cannot be provided without consent.

ID / Category / Related Requirements	Guidance	Comments and Rationale
UI8 User Interface	Consideration: Include user guidance and warnings on PI handling for any functions that output PI to the user interface	• For extremely sensitive information (example: witness protection program information), the user should be advised to ensure they are in a secure and controlled environment before displaying the information (versus displaying on a laptop when traveling in a plane for example). • Wherever possible, and especially for non-routine situations, the user should be reminded as to the PI uses and disclosures that are allowed and those which are not. • The user should be cautioned about copying or printing the information unless the have authorization since this generates potentially uncontrolled copies of the PI.
UI9 User Interface	Consideration: Disable the “cut and paste”, print and “screen print” capabilities for application functions that output PI to a screen.	• This is a specific extension of UI8 in the situation where it has been determined that copying and printing of displayed PI is not appropriate. • This is not always technically feasible but should be used if it is. • This is more critical if the output comprises of multiple records versus just the record of one individual • This measure does not prevent copying, it just makes it very inconvenient • The counterbalancing consideration is that if other mechanisms are not provided that allow the user to transfer the information to another media for legitimate purposes then this just becomes an inhibitor to productivity.
UI10 User Interface Transformation	Concept: Provide the minimum PI and the maximum level of anonymity via the user interface that will still allow the purpose to be achieved	• This is clearly important for the different users of PI within the organization where transformation techniques should be considered to render the data provided via the user interface for each job role into the form that poses the lowest risk. • There is also an implication for data collection, especially where third parties may be providing consent – the user interface should not volunteer any stored PI that is not needed for the transaction. • Playing back stored PI to the individual for validation and correction must only be done if they have been strongly authenticated.

ID / Category / Related Requirements	Guidance	Comments and Rationale
DD1 Database Design Keys	Concept: Use an MBUN – a Meaningless But Unique Identifier to index PI	• The MBUN is meaningless when taken out of context. • This MBUN is not a secret like a password but its use is hidden as much as possible. It is used "under the covers" and is not something that gets printed out in reports • The association between the MBUN and public PI identifiers (e.g. Drivers License Number, Alberta Health Care number) constitutes a control point and should be closely guarded • The assignment of the MBUN (e.g. IID) should be performed by a specialized privacy-enabled component that uses appropriate identifying PI to match identity.
DD2 Database Design (Related toDD1) Keys, Placement, Trans, Authorization	Consideration: Separate PI into identifier information (e.g. PID) and attribute information such that access to each can be independently specified	• Access to identifiers should not automatically result in access to attributes and vice-versa. • This will often imply that identifiers and attributes need to be physically separated into separate tables or databases, but some databases provide field level access control, in which case, the identifier (E.g. PID) and attribute information can be stored in different columns in the same table. • This is a defense–in-depth security strategy approach that essentially de-identifies data and renders the component parts much less sensitive. An explicit join operation is required to re-identify the information and the differential access authorization allows this capability to be controlled. In addition, it may be the case that some of the separated PI is more sensitive than others and should be stored in a higher security zone. • Further degrees of separation (file, database, platform etc.) would mostly depend on the need to place different components of the PI into different security zones. • For very sensitive information, also separate identifiers from each other • Clearly, performance and complexity is the counter-balancing consideration

ID / Category / Related Requirements	Guidance	Comments and Rationale
DD3 Database Design Placement, Trans, Authorization	Concept: Store sensitive PI and PI linkages in encrypted form	• It is up to the organization to define what constitutes "sensitive", but at a minimum, biometric templates should be stored in encrypted form. • This also provides a "defense in depth" supplement to DD1 and DD2 where critical linkages and identifiers can be further protected
DD4 Database Design Consent	Consideration: Store consent, preference and choice values together with the PI they relate to.	• The determining factor for this should be the degree to which queries will be made against the database that could return multiple records • If a query would result in many matching records being returned subject to Data Subject consent, checking that consent in an external source for each matching record would be very inefficient vs. just doing a "join" in the same database. • This is also appropriate if the Data Subject is to be offered fine-grained choices (ex: selectively consent to use or disclosure of individual data items). In this case, recording consent in a central record would essentially require duplicating the data structure in that record. • The counterbalancing consideration is that recording, updating or summarizing consents back to the Data Subject is more complex.
DD5 Database Design Consent, Access	Consideration: Store consent, preference and choice history together with the PI they relate to.	• This is very similar to DD4 but recognizes that there are simple choice values like "preferred language" and then there are auditable consent values where it is important to be able to track the history (value and timestamp) of when a Data Subject gave consent or how they changed their values over time. • Counterbalancing considerations also very similar.

ID / Category / Related Requirements	Guidance	Comments and Rationale
DD6 Database Design Enforce, Monitor, Access	Consideration: Provide for storage of "Privacy unit of work" codes against each PI record or element.	• A "Privacy unit of work" code allows all the elements touched by a particular Privacy transaction to be tied together and linked to a central record describing the Privacy transaction. The concept is useful for situations where the Privacy transaction or the elements touched by it may vary for each execution. • Any privacy transaction that is non-routine and/or which impacts PI spread over many storage locations, is a candidate for using the unit of work concept • An obvious use is in support of the Stick Policy Paradigm (AC5) where the Policy is the central record and the corresponding unit of work identifier is used to tag all PI collected under that policy • This concept can also support privacy transactions such as requests for access or non-routine disclosures where a central annotated record may be required to describe the transaction along with links to all PI affected by the transaction. • This technique can also be used to record consent, or change of consent, where the choice is made once centrally and then propagated to the storage locations of all the affected PI
DD7 Database Design (Queries)	Consideration: Implement query rules such that queries on de-identified data don't return results if the result set is very small	• For further guidance, refer to the "K-Anonymity" guidance in the Data Transformation section. • This helps to prevent identification through triangulation • This is really "second order" privacy protection • The Statistics Canada guideline is for a minimum cell size of 5
DD8 Database Design (Queries) Authorization	Consideration: If a multi-level security scheme has been implemented, leverage it to classify both PI and queries by level. Then mandate that a query cannot return PI classified at a higher level than the query level	• Not applicable for many organizations, but this may be applicable to some applications within the GoA..

ID / Category / Related Requirements	Guidance	Comments and Rationale
DD9 Database Design Tax, Placement	Consideration: Create a Privacy Metadata Schema for each database containing PI	• The Privacy Metadata Schema uses the Privacy Taxonomy (TX1) to describe the PI contained in the database and the policies that apply to it in a consistent way. • The initial use for this is to provide consistent privacy-relevant documentation for the database. The ultimate use is to provided a rule-set for future database resident policy engines. These engines conceptually provide the most practical way to ensure that queries returning multiple PI entries are privacy compliant (DD4). • This is not in conflict with AC1, the policy can still be centrally managed. This is just an illustration of local enforcement to provide optimum performance • The counterbalancing consideration is that in the absence of the policy engine, will documentation alone be sufficient benefit to justify the effort.
DD10 Database Design (Related to AD6) Placement	Concept: Do not allow central storage and transmission of the biometric templates	• The biometric template here is defined as the one-way hash or representation of the original biometric – not the biometric itself. • This guidance is intended for situations where a biometric template is used for authentication (vs. matching) purposes. • The way to achieve this is to have a secure, trusted client device (like a smartcard); which contains the biometric template. Authentication happens within this device and all that is transmitted to the entity requesting authentication is the result of the authentication test, not the biometric template itself. • Clearly, biometrics are extremely sensitive since once they are compromised they cannot be re-issued. • Also, not building a central database of biometric templates avoids the possibility of inappropriate matching activity • The counterbalancing consideration is that it may not be practical or affordable for the client population to be equipped with the trusted devices

ID / Category / Related Requirements	Guidance	Comments and Rationale
DD11 Database Design	Concept: Capture and store the authority and currency status of the PI within the database	• The authority just means identifying whether the PI in the database is a "source of truth" or whether it is just a copy of another database. This information can be included in the metadata described in DD9 • Currency status means being able to identify when the PI was last updated or validated. This will likely require a separate column as this information often varies by record (i.e. by individual).
LR1 Logging, Retention and Audit (Related to AD2) Monitor	Concept: Applications should write values for parameters that provide context for PI access requests into audit logs of their activities	• Even if the common services or components that perform privacy conformance checks maintain an audit log of their activities, having applications track activities in a form that can be manipulated to reveal privacy significance, is a good integrity check. • It may also be useful in an environment where common services and components have not been implemented and a manual audit of privacy compliance is required (or for forensic investigations). • This will enable future "privacy violation detection sensors" that report violations to an IDS (Intrusion Detection Service).
LR2 Logging, Retention and Audit Monitor (Related to DD1)	Concept: Audit records of transactions involving PI should be indexed using an internal key (e.g. IID)	• This in effect means that audit logs are not PI in themselves which in turn means audit activities are not privacy invasive • For enhanced protection, this can be a special logging sub-ID that can only be linked to the internal key (DD1) by a well-protected identity protection component.
LR3 Logging, Retention and Audit Monitor	Consideration: Develop usage profiles for job roles and audit against normal behaviour patterns	• This is really more of an opportunity to provide additional privacy audit measures in environments containing sensitive PI that is facilitated by some of the other concepts already described (AD2 and LR1)

ID / Category / Related Requirements	Guidance	Comments and Rationale
LR4 Logging, Retention and Audit Monitor	Consideration: Log all queries against databases containing PI	• It is impractical to log each record that is touched by a particular query (the log would be comparable in size to the database). However, if the query is logged, events can be reconstructed if need be to determine if a particular record was touched. • Clearly, if the content of the database is very dynamic, re-running the query may not return the same result as it originally did. However, just the parameters of the query alone may be sufficient for audit purposes.
TX1 Taxonomy Tax	Concept: Use a Privacy Taxonomy to distinguish PI from other types of information	• This is a minimal requirement to be able to apply any technology to support privacy compliance (AC2)
TX2 Taxonomy Tax, Trans	Concept: Data classification schemes should distinguish PI that is used as an index key or identifier from PI that is just an attribute	• Index keys include unique or relatively unique information such as name, account number, email address, employee number etc. • Attribute information such as age, salary, evaluation etc. is not unique or personal without an associated index key or identifier • Clearly, there is no black and white rule as to whether a particular type of PI constitutes an identifier but it should be possible to make a reasonable choice based on the organization's use of the data.
TX3 Taxonomy Tax	Consideration: Classification schemes for data types, uses, disclosures and retention should be compatible with P3P	• The Data Taxonomy section of the GAEA PA incorporates P3P • P3P is the de facto standard and compatibility with P3P will position an organization to implement available and planned privacy technology. • Since P3P is primarily Business-to-Consumer, most organizations will have to extend the P3P scheme to meet their needs

ID / Category / Related Requirements	Guidance	Comments and Rationale
AA1 Authentication, Authorization and Identity protection	Concept: Authentication, Authorization and Identity protection should be performed by specialized enterprise-wide common components.	• This facilitates the development of common services and components to mediate data access and privacy compliance (AC3) – these will require authenticated credentials as input and so consistency of these credentials makes these services and components simpler to write • This also avoids the need for each application to build in it's own authentication and authorization function and changes to security policy are less likely to require application updates. • This guidance essentially supports the Authentication, Authorization, and Identity Nodes and components proposed in the GAEA Security Architecture
AA2 Authentication, Authorization and Identity protection	Concept: Authorization should be roles-based. Access to information or capability to execute a function should be defined in terms of roles.	• Privacy policy is expressed in terms of roles versus individuals, which makes articulation of policy much more compact and static. • This is especially important when attempting to automate access decisions against privacy policy since it removes the need for a mapping layer. • This guidance supports the role-based security proposed in the GAEA Security Architecture
AA3 Authentication, Authorization and Identity protection	Concept: Ensure the original biometric cannot be reverse-engineered from the biometric template	• In other words, using a strong, one-way hash • Once compromised, biometric features cannot be re-issued!!
AA4 Authentication, Authorization and Identity protection (Related to AD9)	Concept: Authenticate only the minimum required identity.	• Do not require identification if the service can be done anonymously or under a pseudonym. • For instance, a smartcard can validate that the holder is the true owner (via PIN, biometric etc.) and then send a message to the requesting system that the holder is valid. The system can then proceed based on the identifier associated with the smartcard (once it has established that the card itself is still valid). • Especially do not use a biometric for identification if this is not crucial. • If needed, re-authenticate with a stronger identity at the point it is needed (step-up authentication).

ID / Category / Related Requirements	Guidance	Comments and Rationale
AA5 Authentication, Authorization and Identity protection (Related to AA4)	Consideration: Do not use a biometric for identification	• Using a biometric template for authentication (a one to one match) is not privacy invasive if done correctly, however, using a biometric template for identification implies searching a database of biometric templates until a match is found (a one to many match), and the existence of such a database is open to potential privacy abuse (DD10) • Use of a one to many matching process should be very carefully considered and controlled (usually used for law enforcement and legitimate investigations).