INFORMATION ESCROW SERVICES TECHNOLOGIES

Long before the first computer was invented, financial banks learned to protect user's assets. Banks apply a large number of security technologies, policies, and procedures to ensure the integrity of your deposits. From security cameras to time locks on their vaults to insurance, banks continue to advance their security and protection schemes to respond to the ever-changing threat landscape. Only when banks can truly protect your assets, do you trust the bank with your deposits.

The information bank uses the same, time honored approach to protecting your digital assets. The Information bank orchestrates a large number of technologies to ensure that your information is protected and recoverable not matter what happens. The Information Escrow Service or IES was designed to provide this level of protection and yet make it hidden from users. Users do not have to do anything specific to have their information protected. Just like your financial bank, you can make a "deposit" and not have to do anything else to ensure that deposit is safe. However, unlike a financial bank, no humans are required to operate IES. IES is automatic, hidden, and always operational.

Information Escrow Service (IES)

Listed here are the core technologies required to implement IES. While each technology is described individually, they only become powerful once they are integrated together as a coherent service. Much like running a financial bank, configuring, coordinating, automating, and synchronizing this pool of technology is vital to providing this service that is hidden from the actual users of the bank. Some of these technologies currently exist and are simply incorporated into the product. Others had to be developed specifically for IES. Also note, just like a bank, there are a few technologies that are secret and are not disclosed in the list below.

The key technologies that make up the Information Escrow Service

	Technology	Description
1.	**Information Asset System**	One of the key, unique, and ground-breaking technologies built specifically for the bank is the Information Asset System. An Information Asset is a collection of data (files, metadata, logs, emails, contacts, etc.) that, collectively, are meaningful to a user. The fact is that users think, work, and communicate using Information Assets. See our website for additional details.
2.	**Process Controls**	One important part of managing valuable information is the proper use of procedures and controls around that information. Much like a financial bank, the Information Bank implements very powerful, yet easy to use procedures that maintain the integrity of your assets. These procedures ensure the proper steps are followed, all the steps are completed, and if something goes wrong, that everything is put back to where it started. The ability to automate certain functions is

	also an opportunity for the bank to continue to provide value far into the future as the needs and requirements expand and change.
3. File and Directory Security	If information is valuable to you, don't you think it needs to be locked down? You, or anyone else, should not be able to bump their mouse and have the file zoom to some faraway land. Criminals or ransomware should not be able to modify every file without a fight. Even though elaborate file and directory security has existed since the creation of computers, they are almost never setup correctly. They certainly can't be easily changed as the state of the information changes. Likewise, users shouldn't be able to store the funny cat video right in the middle of your tax records. The Information Bank completely automates the settings of every file and every directory under its control.
4. Security	The Information Bank utilizes and implements a large collection of features and technologies that fall under the category of "security". Note that while many technologies listed here have attributes that can be construed as "security" related, there are a large number of security specific tools, configurations, and procedures that enhance the protections of your assets. See our website for more details.
5. Metadata	Metadata is simply a fancy word for "data about your data". Think of it as tags you can place on your assets that help further describe the asset. Note that metadata support has been in file systems for ages. However, its use has been very limited. The key with IES is that the metadata is incorporated as part of the process automation. Only then is the metadata validity ensured. Metadata has also been described as "a love note to the future".
6. Asset account directed security	More than just the access methods can be controlled on the information. Based upon the needs and attributes of the information in the account, many of the security, process, and other requirements can be setup based upon the asset account (a collection of similar or related information).
7. Access method synchronization	The Information Bank supports a large number of ways users can access their information. It is important to ensure that if the different methods are being used on the same asset or account, the assets will remain consistent and correct for all the users.
8. Automation	Any collection of information will have some form of "Best Practices" that have been created and proven over time to manage that information most effectively. The bank will automate these processes and have them available at the press of a button. Automation insures everything is done correctly and completely every time.
9. Inventory	An inventory of every asset in every account is always maintained. This way the bank can tell what it is supposed to have, what state it is in, what every asset actually is, what they are related to, etc. Inventory – Know what you are supposed to have. This allows the bank to check, as part of its internal auditing system, l if something has

	been added or missing, and can correct it before it becomes a problem.
10. Directional Mirroring	Traditional data mirroring will automatically make a copy of any change. However, if that change happens to be bad, such as a corruption or unauthorized deletion, normal mirroring will simply propagate the disaster. Directional Mirroring, on the other hand, can tell if the change as supposed to happen or if a problem is afoot. The mirror can then automatically and transparently know which direction to make the copy, thus the directional name. Should the change be authorized, a mirrored copy will be made. However, if the change was not desired, the information will be recovered from the mirror. It seemed to us that wiping out a good copy of something with a bad copy is not what users expect.
11. Cloud-based object storage	The advent of the "Cloud" has allowed us to easily make copies of important information offsite such that should some catastrophic problem occur, the odds of having the cloud copy also lost is very low. However, raw object storage is not the easiest to use directly. The information bank takes advantage of this technology but hides all the complexity from the users.
12. Inventory Mirroring	Just like all the data being mirrored, the inventory itself is also mirrored to the cloud, for the same reason. Both the data AND the inventory have to be kept offsite and synchronized to insure that a recovery can be made should something happen. See the 3-2-1 entries below for additional details.
13. Cloud Failure Recovery	Cloud vendors are, by their very definition, a service that users really don't have much control over. If they foul up, well, not much you can do about it but complain on social media. No service is perfect. Cloud vendors, unfortunately, have a whole lot of new ways they can make your life miserable. We have suffered just about every one of them personally! Our view of cloud services is that because we have no control over them, we should not and will not trust them. We treat cloud services just like any other technology that can store information. It has its advantages, but it also has its disadvantages. It is our job to use these services in such a way that no matter how they fail, foul up, or "break", your information is safe and recoverable.
14. Separate cloud services	The bank separates the different aspects of the offsite protection system across at least three different service vendors so that the compromise of any one, and even two, and yes, even all three services does not put any user information at risk. Isn't it about time we make things more difficult for cyber-criminals?
15. Cloud vendor switching	If a cloud vendor does something to make them inappropriate for use with the bank, the bank has the capability to switch to a different cloud service transparently. These services often rely on the fact that migrating your data to a different vendor is so difficult, they are free to make changes without losing customers. Not so with the bank. We can monitor changes and change services transparently to all users.
16. Asset Watcher	One of the key technologies developed specifically for the Information

	Bank is the Asset Watcher. This technology acts like a security guard monitoring the activity on all your assets. The technology is used for many different things including detecting if something was supposed to happen or not and triggering the directional mirroring to perform the correct response. This service is in operation 24-7 to ensure nothing escapes the eye of the watcher.
17. File System encryption	Files are stored in a file system. However, it is vital that, should the Information Bank be lost or stolen, none of your information would be at risk. All of your Information Assets stored in the bank are stored in an encrypted file system.
18. Protocol Encryp0tion	When your computer is talking to the bank, you don't want anyone to be able to eavesdrop on your conversations. The bank uses powerful encryption for all bank communications from your computer.
19. Inventory Encryption	The Information Asset system keeps information about your files within its inventory system. Just like your file's content, you don't want to have any of the inventory content to be compromised should the bank be lost or stolen, so it is encrypted as well.
20. Cloud Encryption	A copy of every asset in your bank is sent to the cloud just in case the bank fails and needs to be recovered. However, despite what cloud vendors say, they cannot prevent every single attack from having your copies stolen. Therefore, we encrypt the asset in the bank before the data goes to the cloud. Therefore, if the cloud is ever breached, and all of them eventually will be, your information is still safe.
21. Key Encryption	Encryption technology uses a set of keys to perform the encryption, and more importantly the decryption of your information. It is important that whenever those keys are stored, the keys themselves are encrypted. With the Information Bank, all encryption keys are then encrypted and escrowed to the central bank. (See central bank and key escrow below)
22. Recovery Package Encryption	Every bank is unique for each customer. While the inventory and asset contents are stored in the cloud, the rest of the data such as configurations, keys, etc., need to be saved. We call this data that is unique to the bank the recovery package. This recovery package is encrypted and sent to the central bank for safe keeping.
23. Central Bank Encryption	The central bank is a special Information Bank that your bank will talk to in order to ensure your bank and all your assets can be recovered. The central bank itself uses specialized encryption techniques such that even if the central bank is compromised, none of your information is at risk.
24. Client and server side certificates	Whenever banks talk to each other, each bank needs to ensure that they are talking to the correct bank and not some impersonator. The way this is done is via what are called client and server side certificates. These are created when a consumer first sets themselves as the owner of the bank. From that point forward, each bank challenges the other to ensure they are the same bank as when the communication first occurred.
25. Bank ID,	Anonymization is a fancy name for not being able to make any sense

Metadata, Filename, and Directory name anonymization	out of the names. Should the bank be stolen or the cloud service get breached, the thief (or some over zealous government agency!) will not be able to tell anything about the information stored, even though they would not be able to read the asset contents. All names are replaced with what appear to be random strings of letters and numbers. They won't even be able to tell that any particular stored object even belongs to you. Only your bank can tell what is going on.
26. Encryption Key Management	One of the weakest links with encryption is the management of the keys. As anyone knows who has gone up to their friends house only to find a key under the mat knows, it is important to protect all the keys. Protecting encryption keys is even more difficult since they are just a small digital file with a bunch of numbers in it. Forcing every user to keep track of their own keys would only add to the complexity of the system, not make it easier to use. The Information Bank manages all the keys in such a way that users don't even need to know what a key is, much less be responsible for not losing one. As you can see from the list of the different encryption functions, there are a lot of keys in use. The synchronization of keys with the central bank escrow service ensures the keys are not lost, stolen, or compromised, without any user interaction.
27. Central Bank	Each Information Bank, referenced as a "Consumer Bank", communicates with another bank called the Central Bank. The central bank keeps all the information necessary to recover from a failed or lost bank. As mentioned above, another level of encryption is added so the compromise of a central bank does not compromise any user information. In case you need to recover your bank after theft or failure, the critical information necessary to recover your bank can be recovered from the central bank. (See Disaster Recovery section below.)
28. Recovery Package Escrow	Each bank uses the Central Bank escrow service to hold all the rest of the data needed to completely recover the operation of the bank after a disaster or replacement. In this case, the word "escrow" means to give someone something for safe keeping so you can get it back later. Many different pieces of data are escrowed to the central bank in order to be able to recover your bank should something bad happen to it.
29. Disaster Recovery	Disaster Recovery is a process that will completely restore all operations and content of your bank should it be lost, stolen, or break. A new "empty" bank replaces your lost hardware and the disaster recovery process is initiated. That process is not much more than a button press by you. Once started, the recovery package is first recovered from the central bank and the configuration reestablished. Then the inventory is recovered from the cloud, which is a reasonably quick process. All the assets are then marked as being "migrated" to the cloud. (See ILM below). Then a background process is started called the "dribble syncher" (I love that name!) that will slowly restore each asset back into your bank. Should an asset be needed before it is

	copied back, the system will retrieve that file immediately. The goal for this is to be able to get your bank running as quickly as possible so you can get on with your life and put the loss or failure behind you.
30. Traceability	Just like every financial bank, the Information Bank keeps track of every transaction that has been made to every asset. Without it, you would wonder who did what and when they did it. The logs are automatic, hidden, and record who did it, what they did, when they did it, and what computer they used to do it.
31. Audits	Much like a financial bank, the Information Bank performs periodic checks as to the integrity of the contents of the bank. These checks, called audits, verify the integrity of the information you have and will recover from any problems detected. Trust but verify.
32. Template Files	When an asset is created by a process within the bank, the asset is first created from a protected template file. These files, which are themselves stored in the bank, can be verified as to their integrity. Thinks like avoiding unauthorized macros in Word documents or unsupported fonts can be avoided. The process used to create a new asset can also specify a specific template that has to be used. For example, if a contract is created, the proper contract template is used to create the initial asset eliminating a lot of the "cut-and-paste" problems that can so easily occur. Note that the bank is controlling your information BEFORE it is even created!
33. Data Deduplication	Traditional mirroring (see directional mirroring above) makes a copy of your information each time it changes. Beyond the problem of propagating corruption, mirroring cuts your available capacity in half effectively doubling your storage costs. The Information Bank uses directional mirroring but with a twist. The second copy created is done using deduplication meaning that two independent copies of the files appear to exist but only use the capacity needed to store one copy. Should one of the copies be deleted, the other file still exists. Such a scheme allows the bank to make mirrored copies of every asset without an impact in available capacity or storage costs. (Note that even with a second copy, the storage device can still fail or be stolen. This is why we also implement cloud mirroring!)
34. Asset level versioning	Storing previous versions of assets makes identifying changes and even recovering from certain edits much easier. However, previous versioning techniques were based upon the file itself, not the asset. Without going into specifics, the Asset Versioning stores the version based upon the process that governs the asset, not just whenever the file system detects a change. Excessive versioning, incomplete versioning, and trimming old versions are a real challenge with traditional versioning techniques. The Information Bank uses Asset versioning to get this done automatically and with just the right amount, frequency, and at just the right time, making the versions much more useful and therefore, valuable.
35. Execution Prevention	No files imported or created within the bank can be executed by the bank itself. This eliminates a common "attack vector" used by hackers

	to compromise normal computer systems. The bank simply will not execute anything put into it.
36. Anti-virus Protection	When an asset is imported into the bank, the source of that asset cannot be determined with any certainty. It could include a computer virus that will spread to another computer if it is accessed from the bank. The bank uses antivirus checking technology to detect these nasty things before they spread.
37. Platform for future security protections.	Computer security is an arms race. The attackers keep coming up with new ways to compromise systems with the defenders responding with new ways to fend off these attacks. The only long-term strategy to survive in such a digital warzone is to have a defensive position that can be continuously fortified as new defensive weapons are developed. The Information Bank is a platform for future security and information integrity protection advancements. The automated system updates feature will ensure that when these new protections become available, they will become part of your bank's arsenal.
38. 3-2-1 strategy for the asset contents.	The 3-2-1 strategy has been recommended to protect digital information for a long time. However, actually implementing the scheme turns out to be lot more difficult than simply recommending it! What the numbers mean is that for everything valuable, keep three copies of it, two on different storage technologies, and one offsite. The Information Bank automatically does this for you so you don't even need to figure this out, much less try to implement it!
39. Triple redundant databases	The Asset Inventory mentioned above is implemented in a database. Databases are very reliable but are not 100%. The information bank uses three different types of databases to ensure that a problem in one type is not devastating to the integrity of your inventory.
40. 3-2-1 strategy for databases	Similar to the 3-2-1 strategy on your asset contents, the Inventory also used a similar scheme. Three copies of the database entries exist for every asset, two on different technologies (we actually have all three different), and one is maintained offsite.
41. Modern File System	A computer file system implements the ability to store random length files, supports file names and properties, and can store them in a hierarchical set of directories. The Information Bank uses a modern Copy On Write (COW) file system that checksums each data block in order to detect read/write errors and report them. This prevents low level storage problems from showing up as corrupted files.
42. File Transfer Functionality (Data mover)	One would think that the ability to copy of file across a network or via the internet would long since be perfected. However, the truth is that in practice, all sorts of bad things can happen when copying information. The Information Bank uses a battle-hardened file transfer capability whenever it is moving information around. Our scheme ensures that information copied between computers, either locally or across the internet, arrive complete and intact at the other end. The requirement for this, by the way, was that a dude in a pickup truck out in the middle of the Texas prairie with just a portable generator, a laptop, and a satellite dish, had to be able to upload a large number of

	files with perfection. As long as you are not out in the middle of nowhere in your pickup, you should be good to go.
43. Asset Search	Being able to locate information when needed can often be a frustrating process. If you can't find something, it is not that different from not even existing! The Information Bank uses what we call the Asset Search method to help you locate information. It is designed to help you locate things based upon what it is, not what you happen to call it and where you happen to put it last. Yes, it will also search the content of your collections as well. The key thing is that it will help you locate the correct copy of something, not the 50 other things that traditionally pollute your set of files. How do you know which one is right? Are they all a little different? No. The Information assures that if it locates something, it is the correct, most recent ONE copy of something.
44. Intelligent Indexing	In order to support the asset search capability (above), the files have to be "indexed" which means each file is read, the words or values extracted, and then put into a special database to facilitate the search process. Reading every asset can be a time consuming task. Just checking to see if the file has really been changed can tie up storage devices for lengthy periods of time. The Information Bank uses its knowledge of the asset and the processes that govern the asset to intelligently know the appropriate time to perform this index operation. It does not, for example, have to look at the directory entries for every file to see if any of them have changed. The index is performed at the right time and only once per file change. .
45. Intelligent remote caching	When the Information Bank is accessed across the internet (we all that internet banking), your computer can't access the files directly. If an asset is accessed, the bank will first transfer the file to the remote computer, and then open the file using the appropriate application. This means that the copy made is maintained, repaired, and removed when not needed. If computer is stolen, it will limit the exposure of the information to the thieves. Another advantage of this is that your remote computer does not have to hold a complete copy of the contents of the bank, like many of the file synch products require.
46. Asset account-level policies.	Certain policies, for example the retention time on an asset, are set at the account level based upon the lifecycle needs of those assets. Such an approach allows for different setups and configurations for different classes of data. Also, it makes it much easier to maintain these properties by setting them at such a high level. They don't have to be done on each individual file. If users had to do this, they would simple skip it, kind of like they are doing now!
47. ILM finally done right	Information Lifecycle Management, or ILM, has had many names over the past half century but all the previous attempts have not been very successful. The idea is that files that are "not important" are migrated to cheaper, slower storage. If they are accessed, they will be migrated back into faster storage first, and then made available. The idea is valid but every implementation (until now) has had limitations that,

	tragically, lead to the demise of many implementation and a large amount of venture capital. The reasons for these problems are vast and too extensive to cover here. However, the Information Bank's implementation of ILM solves all the previously unsolvable problems. The Information Bank's implementation of ILM is used to migrate assets back quickly from the cloud and in the background yet still allowing users to access information that may not be back onto their bank yet. We could do things like being able to only store asset contents in the cloud but not sure we want to bother with that. That brings up an interesting point about capacity exhaustion. Like a vault that is not big enough to hold all the cash. What do you do? Leave it outside the vault on the floor?
48. Asset crypto signatures	Knowing for sure that the copy of the asset that is being stored is really correct can be a challenge. The Information Bank Independently maintains crypto signatures of all assets. This is a way to create a big number that, should the contents of the asset vary even slightly, will be different when recomputed. These allow the bank to ensure the correctness of each copy of any asset content.
49. Asset Account Priorities	Each bank can hold any number of asset accounts (collections of similar or related assets). Each asset account can have a priority set by the owner. This will allow higher priority assets to be sent to the cloud or recovered first during the disaster recovery process.
50. Asset Transactions	Many of the processes implemented within the bank result in changes to both the inventory (database entries) and the file content. A transaction is wrapped around these changes so that should something happen between these changes, the asset is not left in an intermediate state. Either both are changed or neither is changed.
51. Obsolescence Process	Removing an asset from the system is called a "withdrawal". Instead of letting users (or viruses!) remove the file, a process is executed to ensure the user has the capability to do this, the retention schedule is not violated, etc. The user will need to enter a reason for the removal and copies of the entire asset is saved in a temporary directory for a period of time.
52. Automatic serial numbering	Each asset in an account can be configured to have its own unique serial number. Much like numbers on your checks, serial number can be used to ensure uniqueness, track time-relationships among the assets (which were created in what order), and tests for completeness. If there is a gap in the serial numbers, you can immediately tell that something is missing.