THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2014004 March 10, 2014

A New Information Protection Paradigm is Needed — and AlephCloud has the Right Idea

Analyst: Mike Kahn

Management Summary

In days of yore, potentates built castles for protection, primarily their own, but also for the "citizenry" that served them. When the walls of the castle no longer seemed secure from the latest threats, they built wide moats around them and filled them with water. When that wasn't enough, popular lore suggests that hungry alligators were added. While this reptilian detail probably isn't true, there is an important point to this tale. Protecting what was inside the castle was an ongoing effort and new defenses needed to be added to fend off new risks. Further securing the castle made sense, as long as the potentate (or whatever was deemed to be valuable enough to protect) was within the castle. However, when that potentate or item of value needed to be outside the protected area, he or she or it was exposed and subject to much greater risk.

There is an information technology parallel to this story. Enterprises can do and, by and large, have done a very good job of protecting what it deems valuable (in this case – enterprise data) *by keeping it within its protected zone* (in this case, the data center) and *by rigidly restricting and controlling who had access to what parts of the hoard of enterprise data*. That made sense and was fine, as long as the information stayed within the data center and was accessed on enterprise-controlled devices over well-protected networks.

Well, you can see where this is going, can't you? Today, it's a mobile world, with laptops, tablets, and smartphones in the hands of the citizenry outside of the enterprise's safe zone of security. The enterprise's employees and partners and customers need access to that valuable enterprise data (on their many endpoint devices) and it even gets worse, because in spite of policies and procedures to the contrary, we now live in an age where convenience and instant gratification regularly trump enterprise policies. Hence, sensitive and important data tends to get emailed around, sent to convenient exchange sites in the public cloud, and stored in the public cloud and on endpoint devices, all of which are subject to theft, hacking, privacy invasions, and other malicious intent plus what may be the biggest threat of all –careless personal practices.

So enforcing the perimeter around the enterprise castle, widening the moats, and laying traps for the invaders all run head-on into the reality that the citizenry outside of the castle first needs protection from themselves and their sincere desires to do good by putting the enterprise's valuable information at risk. Clearly, telling them not to do it isn't good enough. And it gets worse, because not all of the citizenry have good and noble intentions. **So, what is an enterprise (or government agency) to do to protect its valuable information, whether made valuable by its nature (say, trade secrets or financial data) or by law and the penalties for failing to protect it (a.k.a. fiduciary responsibilities and compliance requirements)?**

The first thing to do is to recognize that castles, moats, and alligators are not the way to protect what is legitimately needed outside of the

Exhibit 1 —

Why Is Cloud Adoption So Hard for Enterprises?

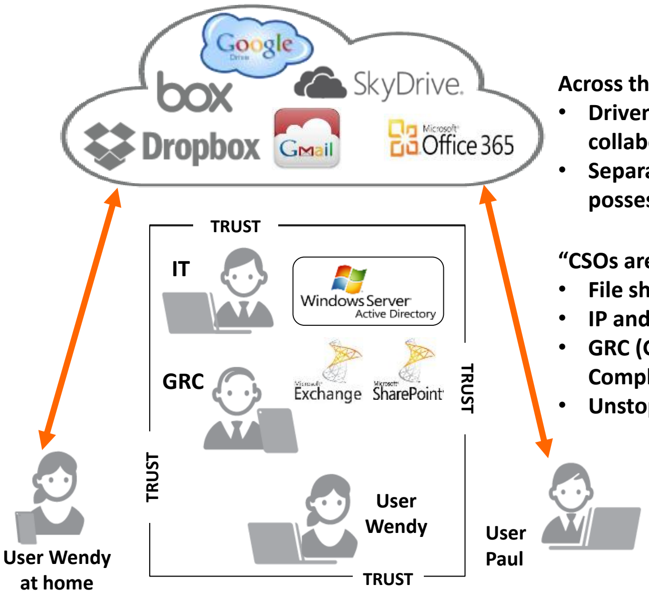

Across the trust boundary

- Driven by BYOD, cloud, and collaboration
- Separated ownership and possession

"CSOs are outflanked"

- File sharing outside IT process;
- IP and data leakage
- GRC (Governance, Risk and Compliance) problems
- Unstoppable yet dangerous

IT challenges

- Insufficient security
- Little or no visibility
- Lack of control & trust

Source: AlephCloud

castle's protective perimeter. What is needed is a new protection paradigm that works outside of the castle (as well as within); one that protects what the enterprise values and must protect. **In the simplest terms, the content needs to be protected, wherever it is located or being used, and its use must be restricted to authenticated users and also possibly for limited uses and/or for a limited amount of time. (See Exhibit 1, above, for a graphical summary of why this is hard to do and the problems that result.)**

With this kind of "transportable magical protection" surrounding the content, wherever the valuable information is being used (whether legitimately or not), the virtual equivalent of a castle with its guards and a moat filled with alligators surrounds the content and protects it, wherever it is. Plus, when the terms of the limited use or duration have expired, the content no longer is accessible. Where can you find this kind of transportable magical protection and privacy control? Read on to learn more about *AlephCloud*.

Understanding the Requirements

The Trust Requirements

Technical details aside (addressed below), **you need an information privacy and protection solution that you can trust to reduce the risks of information loss and/or misuse.** Finding a protection scheme that is trustworthy demands three important characteristics.

- **Its architecture must be inherently and demonstrably trustworthy** (by design). And it needs to be complete by design, i.e., doesn't rely on third parties (others) to make it whole.
- **It must be extensible**, so that it can work for new threats as they are invented and discovered, as well as those that are well known today, and also for new items of value that will need to be protected.
- **It must be scalable** to meet a range of needs from small projects to the largest enterprises.[1]

Of course, it needs to be reasonable (easy to deploy and use, seamless, and not very disruptive to the rapid exchange and use of information) and it needs to be affordable.

The Technical Requirements

As the castle analogy inferred, you need to protect what is valuable wherever is located and not just when it is within the perimeter of the castle. **With respect to enterprise data, you need**

[1] It is important to understand that the requirements of an enterprise data center and the requirements of a cloud provider essentially are the same. Each has one or more data centers in which different business segments or customers have stored their private data, which they want to protect.

to protect the content wherever it is and not focus on the container (device) where it is being held. This is a paradigm shift for most IT organizations, and while not a surprise, does represent a big challenge for most enterprises.

In addition, the solution needs to be self-contained and not requiring the involvement of trusted third parties to do the authentication and/or control. *Why not trust third parties?* Primarily because giving someone else the "keys to the castle" always injects increased risk into the equation and also lessens operational control. If you want to be truly comfortable, you need to do it yourself. Of course, this can't be burdensome to deploy and operate.

Finally, the solution must be complete. A fractional solution is just that and not adequate. The solution must include:

- Encryption and key management
- Key escrow and release
- Policy management and
- Access and compliance tracking

While the focus today clearly is on protecting files, this should not be the final focus. The long-term goal should be to apply the solution to many kinds of objects (different kinds of data at various levels of aggregation) and maybe even individual transactions.

The Scope Requirements

Many stakeholders will be involved with an information privacy and protection solution. Each has different requirements.

- **End users tend to be more concerned about their own productivity than security.** Productivity tends to be defined by the amount of time and effort it takes to get a desired task done or reach a result. In addition, end users tend to want any-to-any collaboration, which means that they want to select the data that is desired and the people with whom it should be shared, with all of the appropriate protections and limitations, of course.
- **Organizations (enterprises and government agencies) tend to be more concerned about enforcement of policies and rules.** They want content control (privacy, revocation, and recovery) and awareness (or visibility) of who is sharing and accessing what (and how and when) for purposes of auditing and compliance reporting.
- **Cloud storage and service providers (including in-house data centers) need scale and problem-free use by their business customers and users.** They are looking for a solution that gives them plausible deniability (i.e., something so clearly secure that it inherently provides protection through encryption and other means, thus ensuring that no data is usable except by its owners and those they authorize).

In a perfect world, one solution would meet the needs of all three. And that is what AlephCloud has done.

AlephCloud and Its Solution

AlephCloud is a Silicon Valley software company founded in 2011 by two technology executives[2] with deep knowledge and extensive experience in cryptography, policy languages, cloud storage, and data management. Its goal, now realized, was to bring to market a trustworthy cloud for seamless digital content access with full privacy and provenance protection.

What is a Trustworthy Cloud?

AlephCloud takes seriously its trademarked moniker "Trustworthy Cloud". It lists the following characteristics as key.

- **Object-level security (encryption)** to protect the content, wherever it is and however its use and duration are limited.
- **Federation and mediation technology** to allow any-to-any sharing across trust boundaries, without reliance on third parties.
- **Relevance to mitigate insider attacks**, such as taking insider documents and being able to access them outside of the control of the enterprise or government agency.
- **Technology platform to support data management strategies of existing enterprise and solution provider storage offerings.** AlephCloud isn't looking to establish a protected storage service or device; it wants to protect data where it sits by wrapping existing content with secure envelope that protects the content and restricts its use,

[2] Jieming Zhu, CEO (who previously worked at HP, Xsigo, LogLogic, and Brocade) and Roy D'Souza, CTO (who previously worked at Microsoft, Mimosa, Brocade, and Intel).

Exhibit 2 —

AlephCloud Content Canopy Solution Overview

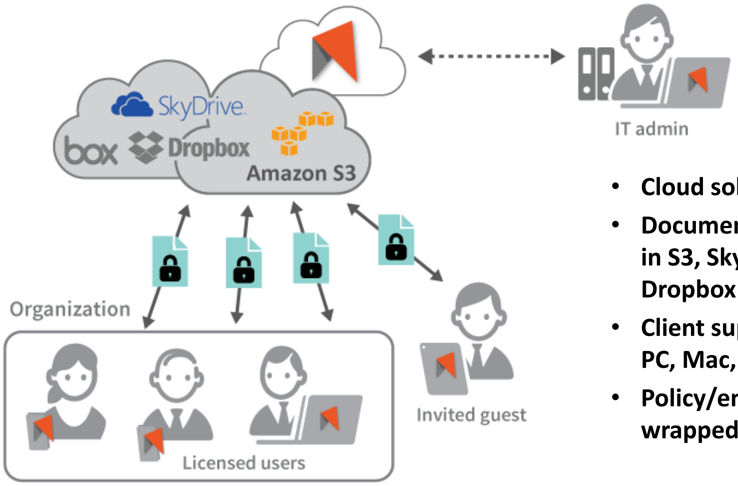

- Cloud solution
- Documents stored in S3, SkyDrive, Box, Dropbox
- Client support for PC, Mac, iPhone/Pad
- Policy/encryption is wrapped around file

Cloud-based with free guest users for low-cost, secure share and sync

Source: AlephCloud

even when away from the security of the data center.

How Does AlephCloud Work?

First, AlephCloud uses the concept of a *canopy* (think "umbrella") to cover the organization's information assets. If the goal is to protect everything that is under the canopy, then this is a good concept on which to focus. (See Exhibit 2, above, for a good visualization.) **Whatever information is held within the cloud (whether externally located or internal) and needs to be protected (because it is a source of risk), i.e., must be brought under the canopy. Thus, internal and external users all will be using protected data, regardless of where they happen to be and via whatever devices and networks they want to use.**

As Exhibit 2 describes, with AlephCloud the policies and encryption are wrapped around each file. What this means is that the data is encrypted and protected at rest (in the data center or in the cloud), in transit, and on whatever device on which it resides (like a laptop or smartphone). Without the right key to unwrap it for use, the file is protected by NSA-sanctioned encryption[3] and is unreadable.

There are three components of the AlephCloud solution.

1. **The Cloud Platform** is fully cloud-based. The platform is the management and administration engine for all Content Canopy users. It is what wraps each file and protects it. Although SaaS-based, the Cloud Platform does not store any user data or passwords, guaranteeing confidentiality and protection from deliberate or accidental breaches. AlephCloud's unique key federation, mediation, and adjudication technology also runs on this high-availability platform.
2. **The Connector** is a software component installed on all devices that need to access protected data. Connectors are available for *Windows*, *Mac OS*, *iOS* for *iPad* and *iPhone*, with other platforms (including *Android*) coming soon, according to AlephCloud. The key function of the connector is to encrypt data before it leaves the device, ensuring that any data transmitted or stored in the cloud can be read only by authorized users.
3. **Third-party cloud storage** is where the wrapped files are kept for the convenience of those needing access to it and for sharing with others. The files encrypted by Content Canopy are stored in a cloud-based file sharing

[3] Described in more detail on the next page.

service, such as *Box*, *SkyDrive* or *Dropbox*. Authorized users can use their personal or business account to store the files they contribute to a shared virtual directory. AlephCloud uniquely allows each member of a single security group to use their own account for the documents they post, while still allowing the sharing of documents between authorized group members. This flexibility makes it much easier to share documents among companies, customers, and partners, because there is no need to mandate where the data must be hosted.

Why is AlephCloud Trustworthy?

AlephCloud's model is trustworthy, because:

- **It uses a cryptography-based mechanism that eschews the need for conventional trust**, that is, a "trusted third party." This "zero knowledge" model makes sure only authorized parties are privy to the information exchanged and stored, while nobody else has visibility, not even AlephCloud. Objects put into the cloud repository are always encrypted before leaving the customer's trusted domain (or trusted device), ensuring strong security and privacy.
- **It enables all parties to securely communicate across trust boundaries through the central intermediary** (the technology platform provided by the Content Canopy). Federated data privacy and policy controls attach directly to the data to protect it, while also ensuring ease of use and the flexibility for businesses to utilize a wide range of cloud service providers.
- **It does so without the service, that is, Content Canopy, ever being able to access the data**, the secret keys, or the data access policies.

All of this needs to be performed at cloud scale in order to leverage the cost benefits of clouds. All end-points in such a workflow can count on the same degree of trustworthiness as point-to-point secure communications supported by protocols such as SSL/TSL and IPSec.

AlephCloud's Content Canopy provides a trustworthy workflow technology that pro vides object-level data protection and access control across organizational trust boundaries. It is trustworthy for the following reasons.

- **It is neutral** with regards to the cloud repository.
- **It supports multiple geographically or technologically disparate cloud repositories**.
- **It supports a broad range of cloud storage access protocols** to enable tracking of objects that are viewed, without the need to control object possession.
- **It is able to instantaneously revoke access when needed**, because the Content Canopy is checked before the file can be unwrapped.

AlephCloud's Content Canopy uses accepted standards-based key federation and key adjudication to support the use case described above. It is based entirely on *NSA Suite B* cryptography.[4] The net benefit of using standards-based key federation is the ability to extend trustworthiness across trust boundaries through the cloud, all while preserving existing business processes and organizational identify management solutions.

Conclusion

In the beginning of this paper, I talked about wrapping the benefit of having a well-fortified moat around a castle and why you would want to do that. I also explained why this moat failed to be effective once the object being protected left the castle. Now, you should understand that you need to protect your data wherever it resides or is being used or shared.

AlephCloud's Content Canopy solution provides an imaginative way of doing this, all why maintaining the trustworthiness of the solution. Content Canopy may just be what you need to provide information utility and sharing, along with a very-necessary dose of protection. Check it out!

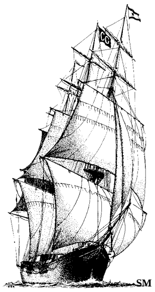

[4] Suite B cryptography has been selected from cryptography that has been approved by NIST for use by the U.S. Government and specified in NIST standards or recommendations. Suite B Cryptography is formalized in CNSSP-15, National Information Assurance Policy on the Use of Public Standards for the Secure Sharing of Information Among National Security Systems, dated March 2010. See http://www.nsa.gov/ia/programs/suiteb_cryptography/index.shtml.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-second year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Mike Kahn **is Managing Director and a co-founder of The Clipper Group.** Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 21 years at Clipper. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation – Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; at Power Factor Corporation, a Boston-based electronics start-up – President; at Honeywell Bull – Director of International Marketing and Support; at Honeywell Information Systems – Director of Marketing and Director of Strategy, Technology and Research; at Arthur D. Little, Inc. – a consultant specializing in database management systems and information resource management; and at Intel Corporation – Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan, where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- ***Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com or via phone at (781) 235-0085 Ext. 121. (Please dial "121" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, *Clipper Notes*, *The Clipper Group Calculator*, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.