CERIAS Tech Report 2019-3
Moving to the Cloud? Points to Consider
by Eugene H. Spafford
Center for Education and Research
Information Assurance and Security
Purdue University, West Lafayette, IN 47907-2086

Moving to the Cloud?

Points to Consider

CERIAS TR 2019-3

Eugene H. Spafford - June 28, 2019 Purdue University CERIAS

Photo credit: Unknown. Found online with ambiguous origins. Version copied from https://karryon.com.au/industry-news/airline/buckle-up-or-is-the-future-really-supersonic/

Introduction

Over the last decade, there has been growing interest by organizations, big and small, in using cloud computing resources. Many organizations mix cloud services with local computing, and some organizations have even sought to make large-scale migrations of their local services to a cloud-based model. The use of cloud computing has a certain "buzz" associated with it, in large part because vendors with a vested interested in selling cloud services have sought to promote (directly and indirectly) its adoption, whether it is appropriate for a particular use case or not.

The cloud model of computing, in its various forms, may provide some advantages in deployment and management of IT in different organizations. What is often overlooked in deciding to move operations to a cloud environment are numerous risks and dangers posed by operations "in the cloud." These range from a lack of visibility into how the resources are managed, to flaws that can be exercised in a co-located environment, to potential unexpected financial burdens. Moving into the cloud will almost certainly increase the attack surface, and add another layer of complexity that may break or be misconfigured. It can also raise issues related to privacy and governance.

This document outlines some issues that should be considered when moving computing to a cloud environment.[2] Decisions about using cloud computing should include careful thought about a set of cost-benefit-risk evaluations, and the items in this paper are (at the least) some of the considerations of risk that should be included. The items discussed herein should be considered through the lens of particular organizational needs, as there may be additional, special considerations. For example, there may be legal constraints on the placement and use of particular resources.

It should especially be noted that cloud technologies are still evolving. There is considerable research being done on issues such as virtualization, containers, data protection in the cloud, encryption use in the cloud, and more.[3] Commercial providers will claim that their technology is mature, but the sheer number of active research projects coupled with unaddressed problems belies those claims. Significant security and operational concerns

[2] There are other documents available that also discuss security, privacy, and control issues with cloud computing. NIST Special Publication 800-144 https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-144.pdf offers a comprehensive view of security and privacy issues as of 2011, most of which are at least mentioned in this paper.

[3] E.g., the IEEE International Conference on Cloud Computing, ACM Symposium on Cloud Computing, etc.

have been identified[4] for cloud computing but these are often ignored by clients in favor of being "trendy" and "cost-effective."

What is the Cloud?

Many parties in different contexts have used the term "cloud." One common use is of remote, centralized online storage. However, "cloud computing" involves more than that. NIST provides[5] a comprehensive definition of cloud computing. Their definition involves three service models (Software as a Service (SaaS), Platform as a Service (PaaS), Infrastructure as a Service (IaaS)) in multiple deployment models: private, community, public, and hybrid clouds. One of the principal benefits claimed for cloud computing is it is based on elastic, on-demand service. Resources are made available on an as-needed basis, and returned to the cloud when no longer in use. Thus, resources are "leased" rather than having extra capacity pre-purchased and potentially idle: metered service replaces sunk costs.

Typically, a customer (client) will negotiate a long-term lease and service agreement with a cloud provider. This lease will specify issues of capacity, up-time, services, etc. This service level agreement (SLA) is a civil contract that also specifies fees and charges. The contract is interpreted according to some legal venue[6], which is also usually specified in the SLA.

In the remainder of this document, all three service models are addressed generically unless specifically noted; many of the risks of cloud computing are effectively the same across all three models. The discussion also focuses on public and community deployments and the public portion of a hybrid deployment, although many are also applicable to all deployments. Some differences between public and private clouds (such as physical security controls) are significant, and the reader is encouraged to consider these and consult related references, as appropriate, as they are not discussed in detail here.

Financial Considerations

There are some obvious potential financial benefits to migrating to use of a cloud computing platform. In particular, shifting costs from organizational computing support and licensing to

---

[4] See, for example, https://www.sdxcentral.com/cloud/definitions/11-critical-cloud-security-vulnerabilities/, https://www.mcafee.com/enterprise/en-us/security-awareness/cloud/security-issues-in-cloud-computing.html, https://www.csoonline.com/article/3043030/the-dirty-dozen-12-top-cloud-security-threats.html. Also see Symantec's June 2019 report, "Adapting to the New Reality of Evolving Cloud Threats."

[5] NIST SP 800-145, https://doi.org/10.6028/NIST.SP.800-145

[6] In the USA this is normally a "friendly" state, but may be in another country depending on the provider.

a more centralized service is often cited as having a financial advantage—not only for reducing the direct costs but reducing other costs, such as for personnel. It may also be the case that the routine services surrounding long-term maintenance (e.g., backups, security monitoring, patching) may more efficiently and consistently be applied by a cloud vendor than in a non-cloud installation. Security, in particular, is held out as an advantage because (it is claimed) that cloud operators can have full-time, trained staff and a more comprehensive selection of tools than most small businesses can afford.

There are four related sets of potential costs related to cloud computing that are not routinely considered. These four may offset savings from moving from a local computing environment to a cloud environment, especially over time.

Auditing Cloud services are leased from cloud vendors. The performance, security, access, and other aspects of the service are stated in a service agreement. To ensure that the vendor is meeting the service agreement — especially at times of stress — the customer will need some form of auditing and logging of the services. This auditing may also require periodic on-site audits of the facilities to ensure that protection and equipment needs are met, if the cloud vendor permits. The cloud vendor may charge extra for any necessary auditing.

If the customer must meet certain audit conditions (e.g., GLB or FEDRAMP) then it may be more expensive to conduct the audits necessary to meet or maintain approval under those standards. This will be a function of the audit conditions, the cloud provider's services, the service contract, and the auditing entity, as well as the scope of the audit.

Legal The costs of monitoring contracts, licenses, regulations, and laws is a continuing need for the customer. Without the ability to directly interact with the IT resources, it is possible additional costs will be incurred to obtain data necessary to meet legal requirements. (Other legal issue are described in more detail, later.)

Future leases The cloud provider is a commercial entity. Once the customer's initial lease has reached its end, the vendor may seek to increase the pricing and alter terms of the SLA and the customer is in a (now) inferior position to negotiate. The long term costs of using the cloud service may eventually approach (or exceed) the costs of maintaining those services locally.

Migration/Lock-In Should the customer wish to migrate services and/or data away from the cloud environment, the vendor may charge for the additional bandwidth and services to make such a change. It is conceivable that a vendor might also somehow disadvantage the client for moving business away (e.g., by strictly controlling

bandwidth). If the customer needs to move to a local (traditional) system, there will be significant up-front costs in buying equipment, hiring personnel, installing software, and performing the transfer. There will also be a cost incurred for dual operating while testing that the transfer is complete and all services functioning normally. Thereafter, there will likely to a cost to audit the complete deletion and removal of customer content from the old cloud provider. These costs of a move may well discourage any such movement no matter how warranted, and thus help to "lock in" whatever pricing the cloud vendor wishes to charge.

Furthermore, as operation in a cloud environment of PaaS and SaaS proceeds, optimizations and creations will be be tailored to those environments, thus losing compatibility and portability. This will increase lock-in.[7]

When considering the costs of moving significant resources to a cloud environment, the future of computing should also be considered: secondary storage continue to drop in price while increasing in density[8], secondary component (e.g., racks, wires, routers, switches) prices are dropping or already quite low, and computing power continues to increase per unit price[9]. Any long-term cost projections should take into account both the decreasing cost of the IT infrastructure and the increasing sophistication of system management technologies.[10] Building-sized supercomputers of a decade ago can now be matched by a few racks of modern processors. This trend in reduced size with increased power for lower cost is likely to continue for some time to come and is likely to negate some of the cost advantages of current cloud environments.

Security Considerations

One of the most fundamental responsibilities of management of any organization is the protection of the operation and viability of that organization. Threats to the IT infrastructure can be existential threats and thus should be addressed accordingly. With the growing risk of attacks from criminal enterprises, ideologic groups, and nation-state actors, defense of both data and operational capabilities must be a primary concern of leadership. The number of

[7] See, for example, https://www.zdnet.com/article/game-of-clouds-lock-in-is-coming/. An extensive treatment of lock-in is available as https://www.intechopen.com/books/mobile-computing-technology-and-applications/taxonomy-of-cloud-lock-in-challenges

[8] Cf. https://aiimpacts.org/costs-of-information-storage/

[9] Cf. https://ourworldindata.org/uploads/2019/05/Transistor-Count-over-time-to-2018.png

[10] In particular, further development of big data handling and AI techniques are expected to make significant changes in how systems are managed.

(and trends in) vulnerabilities being present in commonly-used software[11] only makes the need for strong security more urgent.

Cloud computing has been held out as an attractive, partial solution to security problems. The arguments made by cloud proponents include:

- a cloud provider can centralize critical resources and apply uniform protection;
- a cloud provider can invest in better-trained personnel and tools than most organizations because their cost is spread across multiple customers;
- a cloud provider can do a better job with patches and updates when providing a centralized set of resources.

All of these arguments are true *in theory*. For them to be true in practice means that the cloud provider must be diligent in following best practices and invest wisely in resources. However, all cloud providers are commercial entities that are also seeking to minimize their expenses. Thus, a cloud provider may not achieve the best possible security, especially compared against larger organizations with mature security practices.[12]

Understanding the role of security in traditional versus cloud environments is complicated by the fact that there are no generally-accepted metrics of security…and not even a generally-accepted definition of what security entails. Privacy is also a concern that is closely related to security, and there are no accepted metrics for privacy, either.

Management should balance risks against savings in any decision. Risks can be challenging to calculate, but many mature models are available for approximating long-term risks and costs

The following set of topics provide a summary of some potential risks that should be considered related to cloud security. The items are grouped by Donn Parker's Hexad[13] of security properties, which is an expansion of the more traditional C-I-A[14] model.

[11] See, for example, https://www.securitymagazine.com/articles/90175-open-source-vulnerabilities-increase-in-2018 and https://www.cvedetails.com/

[12] Most of the cloud security claims do apply to small organizations moving their operations into a cloud environment as those organizations have meager (or no) security infrastructure of their own.

[13] For an overview, see the Wikipedia entry https://en.wikipedia.org/wiki/Parkerian_Hexad

[14] Confidentiality-Integrity-Availability.

Confidentiality

In a cloud environment, customer data and processing are on computers that are potentially shared by other customers. Multiple paths of leakage have already been identified in the hypervisors controlling this sharing, with more likely to occur.[15] These include faults in the hypervisor software[16], faults in memory isolation[17], and faults in the processors[18] — among others. These flaws allow reading of memory or cache, and may even enable an unauthorized process to change values[19] or take control of a processor, thus exposing its contents.

In a cloud environment, all of the data and processing may be accessible to employees of the cloud provider — it may even be mandated in the SLA by the provider This access is necessary for them to troubleshoot problems and to monitor usage. Such access may also enable unauthorized access to memory and/or installation of hardware or software that may be used for eavesdropping[20]. Cloud providers usually indicate they strictly limit or prevent any form of unauthorized access to client assets, but the client must trust the provider's controls.

Encryption is often suggested as a solution for the confidentiality of data in a cloud environment. However, encrypted data must be decrypted to be processed. The threats mentioned above can then be executed against the unencrypted data or on the encryption keys as they are used. Furthermore, there must be a robust service to generate and distribute encryption keys for services and data sets, and to handle escrow and reset of the keys. These services add overhead locally. (It would be unwise to put the encryption key services in the cloud as that further reinforces a single point of failure/attack.) If the encryption algorithm is executed in the cloud, the client must trust the implementation to be correct and free of tampering — with little ability to verify that both are true.

---

[15] E.g., http://www.itnews.com.au/news/xen-bug-made-aws-rackspace-data-vulnerable-to-leaking-396393

[16] E.g., https://www.pcworld.com/article/3187782/critical-xen-hypervisor-flaw-endangers-virtualized-environments.html and https://www.csoonline.com/article/3193718/xen-hypervisor-faces-third-highly-critical-vm-escape-bug-in-10-months.html

[17] E.g., RAMBleed https://www.zdnet.com/article/rambleed-rowhammer-attack-can-now-steal-data-not-just-alter-it/

[18] E.g., Zombieload https://techcrunch.com/2019/05/14/zombieload-flaw-intel-processors/

[19] E,g,m https://arstechnica.com/information-technology/2019/06/researchers-use-rowhammer-bitflips-to-steal-2048-bit-crypto-key/

[20] Cf., https://ieeexplore.ieee.org/document/6296060

It is also the case that the users must now communicate over long-haul networks to access the data and communications. This communication provides another avenue of exposure of the data and results, both by access to the communications medium (e.g., cable) and to the switching and routing components.[21] Even if the communications channels operate without active eavesdropping, an adversary may reroute or duplicate traffic without it being noted by the customer.[22]

Availability

Cloud services are supposed to be more resistant to outages because of redundancy and careful administration. However, cloud providers are not immune to accidents or unexpected failures, and some can be significant.[23] Additionally, cloud providers tend to achieve their economy of scale by buying homogeneous systems in large quantities. Thus, common-mode failures may have a more substantial impact on cloud providers than on a system of local clients.

Because of the homogeneous nature of many cloud systems, a logical attack against those systems may have a wide-ranging denial of service effect on customers. For example, if an insider or externally-written malware (e.g., ransomware) were to attack remote management functions, it would be possible to disrupt — and perhaps wipe — a large percentage of systems within a cloud provider enclave. Proof-of-concept attacks have already been demonstrated, working against flaws in Intel AMT systems.[24]

Not all Cloud providers regularly practice addressing disruptions at scale, so a mistake or attack may have significant consequences.[25] If all the critical resources for a customer are in the cloud, the customer is effectively out of business until any such problems are resolved: the customer is at the mercy of the cloud provider's response time and effectiveness. In contrast, a diverse local computing environment can be designed to be resistant to such occurrences.

---

[21] E.g., https://www.zdnet.com/article/cisco-warns-over-critical-router-flaw/

[22] E.g., BGP rerouting https://www.wired.com/story/google-internet-traffic-china-russia-rerouted/

[23] See, for example, https://www.crn.com/slide-shows/cloud/300097151/the-10-biggest-cloud-outages-of-2017.htm and https://www.crn.com/slide-shows/security/300107391/the-10-biggest-cloud-outages-of-2018-so-far.htm

[24] See, for example, https://thehackernews.com/2017/05/intel-amt-vulnerability.html

[25] This has happened many times. See, for example, https://appleinsider.com/articles/19/06/02/mass-google-outage-impacting-gmail-youtube-other-services

The services available from a cloud provider to a client depend on network access to the cloud provider's systems. There must be sufficient network bandwidth between the customer and the supplier to meet demand. Bandwidth is not necessarily under the control of either the customer or the cloud provider — at least one, and possibly many, network providers exist on the paths between the two. Maintaining high capacity for remote access comparable to access to local computing resources may involve extra expense. Furthermore, if access is critical, redundant methods of access may need to be provided. Even with high capacity, state-of-the-art limits, a concerted DDOS may limit or deny access to the cloud resources.[26] If network traffic for the customer is metered, a large-scale DDOS attack can both deny access and dramatically drive up costs. In contrast, a locally-sited computing infrastructure is not subject to such attacks, as they may be stopped at the perimeter.[27]

Rerouting or interference with the routing of network traffic is also a potential concern (see the discussion at the end of the Confidentiality section, above).

Physical disruption must also be considered. A major earthquake, power failure, or hurricane may result in a reduced or restricted capacity to and from the cloud provider even if the customer is not directly affected by such an event. Less dramatic threats also exist that can cause outages of power and telecommunications, including human error[28], wildlife[29], and simple wear-and-tear[30]. It is also the case that network connectivity, power, cooling, and other services necessary for the cloud provider could be deliberately targeted for sabotage, restricting or eliminating all computing access for a customer; in an adequately protected, diverse local computing environment any such a scenario can be addressed according to locally-calculated risk metrics.

---

[26] Recent examples included major stress on Akamai, a major DDOS defense, by parties attacking a journalist; see https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/; and ExploderBot, see https://www.researchgate.net/publication/328355903_ExploderBot_A_Slow_Drip_System.

Nation-states may have extreme systems that can be brought to bear, such as China's "Great Cannon." Cf. https://blog.thousandeyes.com/chinas-new-weapon-great-cannon/

[27] External connectivity may be affected by a DDOS attack, but the internal operations will still be available.

[28] E.g., https://www.wired.com/2006/01/the-backhoe-a-real-cyberthreat/

[29] E.g., https://www.washingtonpost.com/news/wonk/wp/2016/01/12/a-terrifying-and-hilarious-map-of-squirrel-attacks/

[30] See https://thebossmagazine.com/maintenance-statistics/, which indicates that 64% of unscheduled downtime for one class of systems is caused by poor or deferred maintenance. See also https://www.evolven.com/blog/downtime-outages-and-failures-understanding-their-true-costs.html

Special circumstances may also result in a forced network outage. For example, some security incidents[31] have led to organizations curtailing or shutting down their network access as a protective measure. If this were to occur at either the customer or provider sites then the customer would be unable to conduct business — including potential recovery operations. Thus, moving operations into the cloud effectively precludes the option to disconnect the customer site from the Internet for an arbitrary reason and length of time.

Integrity

Integrity is the property of protecting resources from changes, including corruption and deletion. The mechanisms for protecting data include proper access controls, protecting storage media, and keeping up-to-date archival copies. (Issues of access control and protecting storage are addressed throughout this document.)

Cloud providers may offer backup services (archival copying or mirroring) of critical data and programs. Any such offering must be carefully specified in any SLA, and periodically audited to ensure that the backups are viable and current. The backups must also be protected as they present another possible source of disclosure or attack.

The configurations of the applications and systems also need to be monitored. Processing may involve a set of interacting complex components provided by both the customer and the cloud provider. Changes to the configuration of these components may result in failed or incorrect results, as may updates of a subset of the components. This is a concern both because the cloud provider may make unannounced updates as part of regular upgrades to its infrastructure (e.g., OpenStack components), and because the configurations might be targeted by an adversary.

Of perhaps more concern is how the customer can know, without question, that the data and applications being used at the cloud site have not been altered in any unauthorized way from what the customer and vendors provided. The current state of data controls do not support this form of assurance for more than a small amount of data without extensive overhead.

Malware — particularly ransomware — can spread quickly in a connected environment. Removing boundaries and increasing homogeneity, as happens in a move to the cloud,

[31] E.g., external events such as the 9/11 attacks, and internal incidents such as a ransomware worm or discovering a significant penetration of sensitive systems.

provides a more susceptible environment for malware spread widely and with considerable speed. The resulting damage would therefore be more extensive.[32]

Authenticity

When operating in a controlled local environment, it is possible to determine with high confidence that the data, software, and hardware are authentic and what they are believed to be: the customer can be confident that the processing elements are the true and correct elements. When those items are located remotely, it is more complex to make this determination to any level of confidence. Issues of 2-way handshakes, encryption, and data digest signatures can all help establish authenticity, but these need to be layered on existing mechanisms and software.[33] Unfortunately, unless there is significant control over the hardware and boot stack (e.g., with a TPM) and associated software, it is not possible to conclude, with certainty, anything about the authenticity of a particular hardware/software entity. Furthermore, these can add additional processing time and reduce flexibility in most software solutions.

Authenticity also applies to issues of the supply chain. Local processing can employ mechanisms and procedures to gain greater confidence in the authenticity of the hardware, software, and data in use. In a cloud environment, the customer must trust the expertise and good intentions of the cloud provider. When one considers that the cloud provider is a commercial entity intent on keeping its costs low, it should raise concerns about the surety of authenticity and control over the supply chain that may be exerted by the cloud provider.

Control/Ownership

One of the claimed advantages of moving to a cloud environment is that the updating, maintenance, and security of customer resources will be under the control of highly-trained personnel with no other responsibilities. This advantage may not be the case, as cloud providers are also seeking to minimize expenses and maximize their own profits. As is noted in comments throughout this document, visibility into cloud operations, personnel, and infrastructure is different from the local case. Thus, one fundamental supporting element of control — visibility — is degraded. The client will have reduced visibility into important aspects of care for the client's resources, and this may have negative consequences.

---

[32] Cf., https://krebsonsecurity.com/2016/01/ransomware-a-threat-to-cloud-services-too/

[33] Also, as noted earlier, this presents a bootstrapping problem — how does the client verify that these remote mechanisms are functionally correct and not maliciously coded?

More to the point, if an organization's intellectual data and resources are present on a different party's infrastructure, the organization no longer has control, although it may maintain ownership in a legal sense. If the operator of the cloud service decides, with or without notice, to wipe all storage and shut down its systems, it can do so as it controls those resources.[34] The only real control held by the customer is the ability to enforce the SLA using civil law. That control is neither absolute nor timely, and may well be insufficient.

This issue is one of the most serious security concerns about moving to the cloud — the surrender of control. If the data so transferred is highly valuable, or if the ability to conduct transactions is critical (or existential) to the organization, movement to the cloud results in reduced control and additional risk.

Utility

Most cloud deployments are made with the expectation that the cloud environment will provide equal or greater utility compared to what is currently present. Certainly, the expectation is that the cloud at least will provide elasticity (capacity) in response to needs. It is also the expectation that the cloud provider will make available necessary updates and resources in a manner that will at least match what can be performed locally. However, unless the SLA is written carefully and dutifully enforced, these expectations may not be met. The cloud provider may overload its resources, run outmoded or unpatched software, fail to provide sufficient communication bandwidth, or fail to respond to requests in a timely fashion.

With respect to concerns such as the ones outlined above, a customer with a local infrastructure instead of the cloud can choose to devote resources to alleviate the problems. In a cloud-based infrastructure, the customer can only seek to enforce terms of the SLA — if the agreement covers the situation.

The fact that the cloud is non-local means, at a minimum, that there is additional communications delay in common operations. Thus, if the resources provided by the cloud operator are equivalent to localized resources, the overall performance and utility will be less. (See also the "availability" considerations.) The impact of this difference may or may not be significant.

[34] This has already happened, as detailed in https://www.networkworld.com/article/2173255/cloud-computing-cloud-s-worst-case-scenario-what-to-do-if-your-provider-goes-belly-up.html

Other Considerations

Audit

When computing is local, the owners/operators can perform an audit of accesses and operations at will — even continuously, as might be used as part of an IPS.[35] If the data and services are located remotely and operated by another party, any such audits will necessarily be more limited in scope and frequency. It may be the case that the fidelity and scope of any such audit are also limited. Cloud providers may be willing to provide higher-fidelity access to audit and audit logs, but it should be anticipated that enhanced access will be at an increased cost.

As a matter of security, audit data supplied by a cloud provider cannot easily be verified as complete and correct in the same manner as locally-generated data. This lack of assurance not only presents a problem for the customer seeking to monitor and protect operations, but it may present sufficient uncertainty as to interfere with legal investigation and prosecution of cases of misuse.

Legal[36]

An overriding point about cloud service providers: they are for-profit businesses. This means that those businesses seek to maximize their profits and minimize their expenditures and risks. They do not operate their clouds as a public service! To ensure that certain levels of service are delivered, customers must negotiate an SLA (contract) that fully specifies those services and the recourse if they are not met. Thus, it is important to keep potential risks and limitations in mind when negotiating those contracts — the cloud service providers are incentivized to omit or limit their services and thus minimize their expenses.

Of further note is the consequence of failure to deliver stated services or service levels. As the only specification of services is the service contract, it may require a lawsuit to obtain recompense or adjustment. That lawsuit may occur in a venue that requires years to come to trial. Furthermore, the venue of any such suits might be more favorably disposed to the service provider than the customer so there is always some doubt about a successful conclusion. This is may especially be the case if the SLA specifies arbitration rather than trial.

---

[35] Intrusion Prevention System.

[36] This statements in this section are based, in part, on informal discussions with several attorneys. It does not represent a qualified legal opinion, nor does it cite to relevant statute or case law. The reader is advised to consult with well-qualified legal counsel to understand any particular set of circumstances.

In environments involving materials legally protected from certain forms of exposure or transfer, putting those materials into a public cloud environment may be a violation of the law. In particular, putting into an external cloud any material that is classified, export controlled, or regulated under laws such as the Atomic Energy Act of 1954 (as amended) and ITAR might be chargeable as one or more felonies. Other materials, such as health records (protected by HIPAA) and personal information about minors (COPPA) are examples of additional protected material that should be stored off-site with great caution.

It should be noted that liability for any form of disclosure, loss, or misuse of legally protected material cannot be completely transferred to another party. Thus, if a cloud provider discloses, loses, or misuses protected information, the persons who directed placement of those materials in the cloud are potentially liable, both financially and criminally.

The cloud provider may insist on being shielded from any liability for any loss or exposure of customer data. The provider may even go so far as to include indemnification in the SLA, thus shifting all responsibility for damages and costs of defending any suits by third parties onto the customer.

Material stored in off-site cloud systems may be subject to access under subpoena or warrant of the local jurisdiction, and possibly without notice to the customer. This lack of notice is even in the United States: the process of seeking material from a cloud provider under Patriot Act provisions[37] may prohibit the cloud provider from informing customers that their data is being sought. When the data is maintained locally, the customer is the one notified of the order and has an opportunity to address it.

Storage that is located outside the U.S.A. — even in part — is governed by the laws of that locale, and may also allow access to the contents without notifying the customer or allowing an opportunity to contest the grounds of access. Note that some locales also allow the surreptitious insertion of tracking devices or recording of transactions based on their local regulations.

Other local laws may apply to the cloud systems that will, in turn, impact the customer. For instance, if some data is transmitted through or stored in systems subject to EU laws, the GDPR may apply to all information stored and processed in that cloud system. Local taxes may also be assessed in circumstances such as these.

An order executed against a cloud provider to cease or alter operations, including under terms of a bankruptcy case, contempt order, or another legal issue, may have direct effects on

[37] Cf. https://www.aclu.org/other/surveillance-under-usapatriot-act

the customers of the cloud provider. Those customers may not have legal standing to contest such orders even if they are provided timely notice.

Any investigation of criminal actions — by insiders, outside entities, or cloud provider personnel — is more difficult in a cloud environment.[38] Issues of evidence collection, analysis, and chain of custody are all made more difficult by the split nature of location/control and ownership.[39] As such, attempts at either criminal or civil prosecution may be hindered by exclusion of potential evidence.

Supply Chain

Customers have little or no say in where a cloud provider gets its hardware, software, or 3rd-party services. Few cloud providers have the resources (or will) to carefully examine and consider all items brought into the ecosystem it operates. Thus, the provider may have "backdoors" that allow access to customer data and services that are unknown to both the provider and the customer. Once a customer of interest moves to an external cloud provider, 3rd parties — particularly nation-state actors and well-resourced criminal enterprises — will have increased motivation to gain such access.[40]

There is also the issue of examining and securing the software and hardware used by 3rd parties that cloud providers themselves employ. The majority of companies rely on (misplaced) trust alone — with no verification — to secure their resources.[41] The current concerns[42] about use of Huawei 5G equipment[43] and Kaspersky software[44] are two current manifestation of this problem. For example, the cloud provider or some of the communication providers between the customer and cloud provider may well be using some

---

[38] See https://onlinelibrary.wiley.com/doi/full/10.1002/sec.1688 and https://www.reuters.com/investigates/special-report/china-cyber-cloudhopper/

[39] Cf. https://jolt.richmond.edu/jolt-archive/v19i1/article1.pdf and https://josiahdykstra.com/wp-content/uploads/2016/06/Dykstra-acquiring_forensic_evidence_from_infrastructure-as-a-service_cloud_computing.pdf

[40] See, especially, https://www.reuters.com/investigates/special-report/china-cyber-cloudhopper/ and https://krebsonsecurity.com/2019/06/breach-at-cloud-solution-provider-pcm-inc/

[41] Cf. https://www.infosecurity-magazine.com/news/most-firms-rely-trust-alone-supply-1/

[42] This paper takes no position on the actual or potential risk of use of Huawei equipment or Kaspersky software..

[43] Cf. https://www.channele2e.com/business/enterprise/huawei-banned-in-which-countries/

[44] Cf. https://www.nextgov.com/cybersecurity/2018/07/governments-kaspersky-ban-takes-effect/149758/

Huawei equipment for 5G support without the customer knowing it. This use may not be disclosed by or even known to the cloud provider.

Deskilling

If the customer has need of skills involving knowledge of the entire Internet service stack, network protocols, cybersecurity, incident response, or other aspects of cyber/IT operations, moving to the cloud may result in "deskilling" of the customer's staff.[45] Having systems locally that must be maintained, updated, and monitored provides on-going practice and experience for local subject matter experts (SMEs). Moving those activities to the cloud means less (or no) practice for the local SMEs. This will likely result in reduced skills, and may even lead to departure of talented staff who are interested in being more involved in these activities.

Cloud Provider Personnel

It would seem apparent that cloud providers would seek to hire qualified personnel. However, customers have little leverage to ensure that this is the case. Furthermore, customers of cloud providers may be unable to ensure that the employees of the provider have desirable industry certifications, clearances, no criminal records, or are even citizens of countries judged to be trusted by the client. Many providers have a global presence, so the development and administration of their services is conducted worldwide. The customer has legal rights to enforce that its employees meet specific requirements, but the only leverage over an external provider is clauses in the SLA. Violations of these conditions are only remedied in an after-the-fact civil lawsuit, assuming the violations are discovered at all. Lawsuits can drag on for years, and outcomes are never certain. As noted above, even if the client wins a lawsuit in a dispute, the recompense is usually only financial, and that does nothing to recover lost data or opportunity.

The customer is also unlikely to have at-will audit rights to determine if cloud provider personnel meet required standards. Even if there are rights to audit, the form of in-depth background investigation that might be the norm for the client organization (especially if it has some legal authority, such as handling classified information) is unlikely to apply to the personnel employed by the cloud provider.

Cloud providers may also provide access to their physical plant and/or logical infrastructure to cleaning personnel, maintenance personnel, vendor representatives, service personnel, prospective customers, and more. The customer engaging the cloud provider is likely to

[45] See, for example, https://shawnharry.co.uk/2019/01/28/deskilling-in-it-due-to-the-cloud/

have no say in the screening of those people, and may never be notified when (or if) they use that access. The customer is also highly unlikely to have any visibility or control of other customers of the cloud provider, who may use that access in an untoward manner. This is in contrast to a locally-maintained computing facility where access can be determined *and enforced* by the owner of the facility.

Inference & OPSEC

From the standpoint of both data confidentiality (secrecy) and operational security, moving to a cloud structure provides greater exposure. Data and operations must transit the organizational perimeter to reach the cloud provider and are thus exposed to both interception and traffic analysis. An adversary can gain considerable knowledge by observing patterns of usage, sources/sinks of traffic, activation of certain services, etc. When this traffic occurs solely within protected organizational boundaries, any such observations can be contained or prevented; when the traffic must reach a cloud provider, there is little or no such protection.

Furthermore, if a large enterprise moves to a cloud environment, the traffic that would normally have been contained in local subnets will now all flow to the cloud infrastructure. This provides a greater opportunity for an adversary to aggregate disparate information into a larger whole.

Large-scale data mining techniques can be used on records of network traffic to draw conclusions about usage patterns, actors, and items of interest. Even partial information derived from the cloud provider via compromise or insider can significantly improve such inferences, and these activities will not be visible to the security personnel of either the cloud provider or the customer. This information can then be used for more specific targeting for collection or attack.

Ripple Effects

Some salutary effects are common in an environment that maintains a "critical mass" of trained personnel. Among these effects are issues of training and awareness, creativity in solving problems, and cross-training. In environments where a significant transfer of duties off-site occurs, there will be a corresponding reduction in both staff and other resources locally that could lead to these effects. This will likely then lead to a decrease in readiness and effectiveness of retained staff corresponding with a reduction in headcount. (See also the section on deskilling.)

Off-loading to a cloud environment will also lead to an eventual decrease in available hardware/software that may be quickly reconfigured and deployed in cases of critical need, as well as a reduction in personnel who could make such a response. Alternatively, a cloud provider should theoretically have the capacity to meet such a need seamlessly with its resources (especially if IaaS is involved) but if the service is not geared to meet similar needs for all of its customers, the response is uncertain. To date, we have not seen a situation of sufficient stress to challenge a major cloud provider in this manner — but that does not mean one is not possible, especially if a nation-state actor or major geophysical event is involved.

Real-time and Direct Connections

Systems requiring a real-time response or that have connections to sensors or actuators with time-sensitivity are currently not feasible to move to a cloud environment. The speed of connections and the de/encoding of signals mean real-time is not currently supportable with a remote cloud system. Thus, any systems requiring (near) real-time input, output, or calculation will need to be retained locally; the expense for operation, maintenance, and protection of these systems will still need to be maintained.

Homogeneity

As has been noted elsewhere in this paper, homogeneity introduces both opportunities and risks. The opportunities may lead to better consistency and lower cost. The risks, however, are that common failures and common attacks are likely to be more devastating if they occur. Furthermore, nuance and special features may be filtered out of reporting and audits to support the common case.

Heterogeneity provides a certain amount of isolation and "friction" that can prevent common mode problems. Cloud providers, however, base their offerings on the benefits of homogeneity. It is difficult to clearly calculate the trade-offs in this space, but for organizations the are high-value targets or that must maintain continuity of operations in the face of any failure, homogeneity of environment is likely to be a net negative.

Flaws in Migration

Care must be taken to ensure that no flaws are introduced in the migration to a cloud-based system, especially if performed by a third party that is unfamiliar with the customer's special needs.[46] Even when correct, high-security options are available they must be properly

[46] See, for example, https://www.us-cert.gov/ncas/analysis-reports/AR19-133A

configured and enabled. Furthermore, the transfer of customer resources must be protected and carefully checked to ensure that the resources are fully and correctly transferred.

Any backups or hot spares of customer resources should be protected and kept operational for a reasonable length of time. This retention will allow recovery or rollback if the transition encounters difficulties or failures. At the least, local backups should be kept until appropriate backups are present at the cloud provider, and those backups have been properly tested.

Legacy and Custom Code

Using cloud services for SaaS or PaaS may be infeasible when legacy and custom applications are involved. Legacy code designed to run in specialized or unusual environments may not have a matching virtual environment to run on a cloud platform. Thus, legacy systems that are important for enterprise operations including those that have been tuned for custom or deprecated hardware and software platforms will need to be retained locally or re-engineered; the expense for operation, maintenance, and protection of these systems will still exist.

Concluding Remarks

For many small enterprises, moving operations into a cloud-based environment usually has clear advantages over self-maintained systems. Such a move includes the potential for better reliability and security compared to what a small organization can afford, and reduced cost to upgrade to new services. Even so, a careful examination should be made of the risk posture.

The case for a large, established organization is far less certain. Considering the points covered in this paper, an organization with sensitive data—especially with processing regulated by law or with existential value—should be considerably more cautious about any such move into anything other than a privately-operated cloud. *At the least*, serious doubts should be raised about any wholesale move to a public, community, or hybrid cloud from a well-protected and functional local infrastructure; to move critical data and functions "outside" would likely result in an overall increase in the risk of exposure and/or loss that may not be justified by any cost savings. A partial move to a cloud service — with careful planning and long-term consideration — *might* result in some short-term savings, but long-term savings are *not* assured. Furthermore, splitting the budget between local computing and a remote cloud service could result in overall reduced resources both for computing and for security.

If there is some overriding reason to begin (or increase) usage of cloud services that provides greater benefit than all the risks attendant to use of those services, it should be approached in

a cautious, piecemeal manner so as to minimize the chance of unexpected damage. In particular, one or more providers should be identified as presenting the lowest risk. A subset of non-critical functions that are already public-facing should be chosen to migrate to the cloud, while a ready fallback to the original service is kept viable. The subset should be operated for a non-trivial amount of time, and assessed for exposures and problems, as well as convenience and functionality. This process may be repeated, slowly, for more critical materials, with a careful and impartial evaluation made as to cost-benefit-risk.

Building a private cloud and slowly migrating some services to it is another approach that may satisfy economic concerns but also ameliorate some of the risks outlined in this document. A private cloud can also serve as a testbed to determine if currently operational elements will even work in a cloud environment. It is notable that some organizations that were in the public cloud have decided to move to their own private service, for reasons including better reliability and control.[47]

Decision-makers should recognize that there are no good metrics for security exposure and risk when applied to enterprise computing. Making decisions to embrace particular technologies—especially technologies that have achieved "fad" status (e.g., cloud technologies, blockchain, AI) —based primarily on financial models is extremely ill-advised. This caution is especially true for moving significant portions of IT operations into a cloud environment.

It is conceivable that movement of a functioning organization into the cloud could result in the demise of the organization. Clayton Christensen's books, especially "The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail" and work at Harvard Business School have illustrated that "sound" business decisions can ultimately destroy an industry or company.[48] This is because of unintended side-effects, such as some of the ones illustrated in this document. It is therefore critical that a comprehensive analysis of potential problems is conducted. Candid input from personnel associated with the acquisition, operation, and protection of the current infrastructure should be a principal component of any such process as those professional are likely to have the best understanding of current practices and challenges.

Cloud technologies are similar to other technologies — they are not inherently good or bad. Technology provides tools to use. It is up to those wielding the tools to understand the consequences and make good decisions. Thus, it is incumbent on decision-makers to gather

[47] See, for example https://www.networkworld.com/article/3045570/why-dropbox-dropped-amazons-cloud.html

[48] Cf. also: http://sc10.supercomputing.org/files/SC10-ChristensenPhone.m4v

and evaluate complete information so as to make wise decisions. Cloud computing is no different than any other technology, at least in this respect.