THE CLIPPER GROUP
Captain's Log™
Navigating Information Technology Horizons℠
Published Since 2001 Report #TCG2012010 April 24, 2012

Is IT Trustworthy? (or *Who Do You Trust?*)

Analyst: Jim Baker

This is a loaded question. Or is it two loaded questions, asked by business executives, staff, and users of all kinds, including those outside of your business? The first question, in the most general sense, can be read as ***Is it trustworthy?*** "It" might refer to a business system, database, cloud repository, webpage, or so many more things that make up information systems. The second question is more explicit: ***Is IT Trustworthy?*** This is more of an IT organizational question, but equally reasonable. Both readings rely on the concept of Trust, so this is where our discussion begins.

In our daily lives, most of us rely on a set of expectations, some simple and some more conceptual. We count on the power to be on, we count on our clock to awake us at the prescribed time, we count on the current weather report to be accurate, traffic to be as expected (not necessarily perfect, but within the realm of expectations), etc. You might say that "we trust" that these things will be true and to work as expected. If we could not put trust in the components of our daily lives, we would begin and end each day in chaos. **Trust is the backbone of a civilized world.** The same is true for commerce. When we buy or sell shares of a stock or an item on Amazon or eBay, we have a high expectation that the order will be faithfully executed and delivered without breach of confidential information. Without such trust, business would grind to a halt. **Trust is the essence of modern business.**

Bringing this closer to our IT responsibilities and duties, trust is a fundamental expectation to the users of information technology and services, whether businesses, government agencies, or consumers. While there are no signs that proclaim it, we all exist in a world where the motto is "In IT We Trust". Some of us are jaded (or perhaps burned) enough to remember when it was necessary to check to see if a file was actually copied when you moved it from one directory or server to another. Then the prudent method was "copy then verify". Now, we trust it to have been done unless we are informed otherwise. **The more trustworthy an information system is, the faster we can get done what needs to be done.**

This is true of underlying IT infrastructure, as well. If we can trust that the pieces will work well together without significant testing and oversight, the pace of its adoption and use can accelerate to near its theoretical potential. If we have to be cautious at every step of the way (over and over again) then business might never get done. When we make a deposit at a bank, we not only assume that the deposit will be entered correctly but we assume that the underlying systems and infrastructure are protected and safe.

This brings us to the topic for today, which is placing trust in information systems. In the old days of the hardened, glass-walled data center, access was limited, the pace could be measured in days, weeks, and months, intruders were few, and trust was high. Fast forward a few decades and we now have almost unlimited access, microsecond turnaround, vast numbers of unrelenting intruders and perpetrators of fraud, and increasing distrust of those holding our information and protecting it on our behalf. **Are we entering an era where trust can no longer be presumed?** ***What are the consequences? How does this affect our business and information system practices?*** Read on, to learn more, but first buckle your seatbelt, as this might be a scary ride.

Living in Fear

Pardon me for sounding somewhat cynical, but life in the 21st Century is grounded in *doubt* and its companion - *fear*. We worry about our identities being stolen, our credit cards falling into the wrong hands, our credit ratings becoming damaged, our data being misappropriated from our smartphones and then misused, our personal information being shared without our permission, and so much more. Recent news reports scare us that both our credit and our debit cards could have been compromised for millions of us unawares. I could go on, but I think that you get the idea. Whereas the past may have been grounded more in simplicity and naiveté, we now spend our lives attempting to minimize and/or mitigate a complicated web of fears.

We all yearn for more trustworthiness in all phases of our lives but, under a close inspection, we find much evidence that many things cannot be presumed to be trustworthy. Unfortunately, trust tends to be much more binary than shades of gray. **Either something is trustworthy, or it is not. Stated another way, either something is to be trusted or one must presume that it is not to be trusted.** This is true for people, institutions, and business entities, and also is true for IT infrastructure. For example, when you drive across a bridge, do you expect to arrive safely at the other side? What if there were a sign that said: "There is a 99.99% chance that you will cross the bridge without incident. Proceed at your own risk." (That's one failure per 10,000 crossings.) It might be considered a "little failure" or it might be catastrophic, depending on whether you are the unlucky statistic or not.

For those of you old enough to remember, we used to build the equivalent of moats around our data centers. Access was limited intentionally. That was considered a good and prudent behavior, because what was inside was presumed to be safe and, thus, the data center was considered trustworthy. Have we come a long way! In many cases, the moats have been filled in and paved over by the Internet. Anyone on the outside can search the castle's walls continually to see if there is a vulnerable point where it can be breached. We have scores of skilled craftspeople and exotic applications on the lookout for scoundrels and they always are kept busy plugging the newly found cracks and holes. *Would anyone consider this level of protection to be worthy of our trust?* Maybe, but just by a little, at best. ***Does "good enough most of the time" still seem adequate?***

Focusing on Just Part of the Problem

Wanting to keep this Captain's Log to a reasonable length, I will focus on just one part of the "Trust Problem" and save the remainder for later discussions. **Trust by IT's users is an obvious big issue, one that is staring into the face of each provider of IT infrastructure and keeper of business information.**

Let's continue with the moat analogy. If everything of value to a modern business is kept in the castle behind the moat, *what are today's most potent weapons for breaching the castle? What do the warriors and attackers outside the castle have in vast numbers, that they might use to infiltrate the numerous layers of protection, whether by malicious intention or not?* Time's up! **Many have keys to the castle.**

These are the laptops/desktops, tablets and smartphones that users, partners, and employees have been provided or have acquired on their own. In a business world, these devices are connected in one way or another to the treasures of the castle, in this case, data and business processes. The problem is intensified by the increasing trend of BYOD (Bring Your Own Device) strategies, where the instruments of connectivity are no longer provided by the lord of the castle and therefore are out of his control.

Question: ***How big is this device problem anyway?*** **Answer:** ***It is BIG!*** The problem is that everyone in the village (which is now the whole planet) is working against the protection system, whether intentionally or not. The villagers all now have sophisticated tools to make protection many orders of magnitude harder. Many businesses aggravate the problem by allowing the villagers to keep business data and applications on these out-of-castle devices mixed with personal data, contacts, and notoriously insecure social apps. **Building more, deeper, and or wider moats around the castle will do little good, if the treasures can be or already are being transported beyond the moats by unknowing villagers. Storing the treasures elsewhere just brings up the same issues of protection and trustworthiness.** This last point is very important, since many are contemplating using

"The Cloud" as an operational approach to delivering IT or as an outsourcing of responsibilities. This leads us to the heart of today's discussion: ***Is "The Cloud" to be trusted or does it exacerbate the problem?***

Another way to look at this is through the perspective of failure. With all of these devices in the hands of villagers, what's going to happen when their company-provided or BYOD endpoint is lost, stolen, or falls into the wrong hands? *Is this going to happen?* Of course! *Will it be a catastrophe?* That depends on how "fail-safe" the device is. The fail-safe concept has many meanings but one definition fits our discussion: *something designed to work or function automatically to prevent breakdown of a mechanism, system, or the like*.[1] **We need protection that works even when the endpoint device is compromised. Again, this is true regardless of where the treasures are stored.**

Is "Good Enough" Really Good Enough?

This is more than a philosophical discussion; it gets to the heart of the trust issue. *Can you place trust in systems that are invulnerable 99.999% of the time?* This is not about up time (how long your system will run without becoming unavailable), as that is another issue. This is about the number of "nines" of trust that you can place on the fortress holding your data and running your processes without uncertainty.

Many industries build "losses" into their business models. Banks assume that a certain percentage of borrowers will not be repaying their debts. Credit card processors assume that a certain percentage of the transactions that they process will be fraudulent. You might say that they are presuming a level of breaching, whether it is .001% or much more. Someone pays for these losses, and typically it is the good customers that do pay their bills.

Here's an important question. *Could any parts of your business sustain a .001% breach of your data or processes?* Let's assume that you said "yes"; make a mental note of what percentage of the total treasure chest that this might be. Now, let's focus on those remaining parts where you say, "No, I cannot allow this to ever happen for these more valuable assets." *What are your alternatives?*

Let's dismiss all of those situations that are probabilistic in nature. These are where you are willing to take a non-fatal hit – somewhat occasionally – and make that assumption part of the cost of doing business in your business plan. These really are not "good-enough" solutions in a technical sense, because they presume that there will be regular losses that have been predetermined to be acceptable.

This leaves two classes of solutions to consider further: (1) those that are very unlikely to fail and (2) those that cannot fail. OK, I hear you thinking, *can anything be absolutely safe?* Humor me, for a short moment, and this will become more evident.

Let's assume that you have a solution that is said to be very unlikely to fail. It is not fail-safe, but it seems like the small risk is tolerable probabilistically. Now, let's set all of these treasures aside. *Do you have any treasures still left that are worth protecting from <u>any</u> loss or failure?* Maybe, you have some really valuable treasures remaining. In an era of mandated business responsibilities (which varies by nation, applicable laws, and governance policies), of course you do. *Do you want to tell your customers that their data that you were holding has been breached by evil-doers or just made public?* **All of a sudden, "absolutely safe" is beginning to look like it is worthy of investigation and consideration.** Just being "good enough" in terms of your protection strategy clearly doesn't make a very good excuse, whether in a court of law or in the public's eyes.

Absolutely Safe

What is the best way to be absolutely safe? Philosophically, this can only be done through abstinence. This seems extreme, of course, but if no one has access to the treasures, they might be absolutely safe. Obviously, that is not going to be the solution for controlling and protecting access to informational treasures and processes. What might be the next best thing, thinking in 21st-Century terms? *How about a virtual solution, say where possession is real but less tangible or more fleeting?* Here's how that might work.

---

[1] "fail-safe." Dictionary.com Unabridged. Random House, Inc. 02 Apr. 2012. <Dictionary.com http://dictionary.reference.com/browse/fail-safe>.

First, none of the most-valuable treasures ever can leave the castle. You cannot have the treasures sitting on endpoint devices, whether "only a copy" or the only instance. That seems like an extreme requirement, but absolute safety requires a fail-safe solution. Bear with me, please.

Second, access to the treasures must be virtual. When you see a balance online for your checking account, you are seeing a representation not a picture of actual currency and change being held in your name. When you own shares of stock in this post-paper-certificate era, you do not have a physical item to possess or even to see. Virtual access might be all that you need to provide to your users.

Third, access must be fail-safe. If an endpoint device can no longer be guaranteed to be proper and secure, whether because a virtual session was terminated, a device was turned on by an unauthorized user, or even if the session seems compromised (say, not coming from an expected network path), then access can be denied and the user instructed to follow proper access or recovery protocols.

This all has led to a virtual conclusion. **By turning the endpoint device into a secure, fail-safe virtual access point, we have found a way to be absolutely safe, except when a gun is being held to the head of one of the villagers or one of the villagers has decided to "go renegade".** *Is this better than paying the piper (or the lawyers) to get you out of a mess?* Sure, it is a lot better.

Three Important Questions

OK, I have led you to the concept of a Virtual Desktop or something very similar, where each endpoint device serves as a "smart terminal" to the remote castle and its treasures. There are many VDIs (Virtual Desktop Interfaces) on the market today. Two questions are paramount. First, *how good are they?* Second, *are they really affordable at large scale?* And then there is a third, ancillary question: *what does this have to do with The Cloud?*

- **The goodness question brings us quickly to the issue of "good enough".**
  - ***Does it allow the needed access to the treasures?***
  - ***Is the protection really good enough for your most important treasures? Where's the proof?*** Often the proof can be found in certification levels published by third party organizations or the government. **You need to ask!**
- **The large-scale affordability question leads to questions about its business value (in several dimensions).**
  - ***How valuable are the IT treasures and how great is the cost if they are compromised?***
  - ***What is the cost for a solution that will support your target population?*** The situation is a little more complex, due to the issue of scale. What may seem affordable for 100 or 1000 endpoint users may become cost prohibitive when the number of users is measured in tens or hundreds of thousands or millions. Unfortunately, many of the castles with treasures that most need protection are used by many villagers.
  - ***Does the cost per user make sense, given the value of accessing and using the treasures?***
- **The Cloud question is one of perspective and vocabulary. Is the keeper of the treasures a cloud provider?**[2] This might be broken into three sub-parts.
  - ***Does the solution act like a cloud?*** If it is totally virtual and the user has no knowledge of physicality, it probably is cloud-like, even if not labeled as a cloud service.
  - ***Is it expandable like a cloud?*** Clouds are presumed to have near infinite capacity. Does the capacity seem unlimited (i.e., is it able to scale to meet extreme demands)?

[2] This may be seen as a "trick question", as the service provider may or may not call the solution a cloud service.

- ***Can it be managed like a cloud?*** Clouds are presumed to be easy (or easier) to administer and manage, from an IT operations perspective. This becomes very important when large scaling is needed.

Often, it is said that if it looks like a duck, swims like a duck, and quacks like a duck, then it probably is a duck. Similarly, if it acts like a cloud, scales like a cloud, and can be managed like a cloud, it probably is a cloud![3] Call it what you want, but considering it to be a cloud service may help to explain it. Whether you see your business needing to deliver a trusted solution, trusted data, or a trusted cloud, they all have much in common, starting with the need for trustworthiness.

What Constitutes a Good and Trustworthy Solution?

What follows is a list of general questions for you to consider that serves as just a good starting point for your own consideration. It is not intended to be perfect or universally applicable. You need to define your own set of business-specific requirements and frame your questions accordingly.

1. ***Does it offer Operational Fidelity?*** *Does it work as designed, especially under heavy use?*
2. ***Is it Operationally Protective?*** *Does it protect what needs protecting and only allow access to those allowed?*
3. ***Is it Externally Protective?*** *Does it protect from outside elements? Is it safe from unauthorized intrusion, phishing, and other malicious threats?*
4. ***Does it deliver the necessary Operational Availability?*** *Is it reliably available when needed (i.e., short of a locational disaster)? Restated: Is it sufficiently reliable, so that access to the treasures is always there, when needed?*[4]
5. ***Is it Locationally Protective?*** *Is it able to protect against locational failure (outage), say from a region-wide natural disaster? Are data and procedures available at another site farther away? At what cost?*
6. ***Does it enforce Safe Sharing?*** *Does it provide protected multi-residency? Restated: Are classes of users only allowed to access the specific treasures and processes for which they are approved? Is there adequate fencing of resources and users?*
7. ***Does it offer Historical Protection?*** *Does it provide backup and archival services? Restated: Does it treat backup separately from archiving.*[5]
8. ***Does it provide Physical Protection for data and processes?*** *Can it encrypt data on the fly? Can it encrypt data at rest? How well (how fast) does it do encryption and, possibly more importantly, decryption?*
9. ***Is it Topically Relevant?*** *Can it handle Big Data and business analytics? More specifically, can it deliver the most current data on the fly, when needed?*
10. ***Is it Generationally Progressive?*** *Is the underlying infrastructure's architecture contemporary and forward focused? Restated: Has the underlying architecture and components (servers, storage, middleware, key applications, etc.) demonstrated a history of technological evolution? How costly is that evolution? Can evolutions be accomplished in a reasonably fluid manner?*
11. ***Is it Operationally Efficient?*** *Is it straightforward to use, does it require minimal*

---

[3] Ownership and physical location are important characteristics of a cloud service. While the natural inclination for many is to assume a public cloud when talking about clouds in general, that would be a severe limitation or even a mistake. If cloud delivery is a service philosophy then it is applicable to public, private and hybrid clouds. For an in-depth discussion on private clouds, see the January 25, 2012, issue of Clipper Notes entitled *Moving to a Private Cloud? Infrastructure Really Matters!*, available at http://www.clipper.com/research/TCG2012001.pdf.

[4] For more on high availability ,, see the September 11, 2011, issue of Clipper Notes entitled *Disasters Remind Us Why We Need a High-Availability Data Center - How Prepared are You?*, available at http://www.clipper.com/research/TCG2011032.pdf

[5] For more on archiving and it differences from backup, see the February 29, 2012, issue of Clipper Notes entitled *Regaining Control Over Hoards of Enterprise Data - Why You Need an Archive*, available at http://www.clipper.com/research/TCG2012005.pdf.

(hands-on) administration, and is it driven by policy? Restated: As the volume of treasures and users grow, does the administrative workload grow more slowly (i.e., can more work be done with proportionally fewer administrators?) Does is optimize itself autonomically?

12. ***Does it Scale Manageably?*** *Can systems resources be scaled (both upward and downward) to meeting changing demands without compromising in-process activities?*
13. ***Is it Application Transparent?*** *Is it able to run a diversity of applications with few or no changes?*
14. ***Is it Cloud-like?*** *Can a multiplicity of services be delivered to a multiplicity of users from a common management point? How easily can users provision by themselves (i.e., without special IT assistance)?*

The Bottom Line

If you have read this far, thank you and congratulations. You have gone from fear, to the need for trustworthy protection and access to treasures, to the need to do something about the big problem presented by insufficiently secure endpoint devices, to the need for virtual access to your treasures with many questions that you need to consider. Along the way, we have repeatedly explained how IT infrastructure really matters, a continuing theme for The Clipper Group. What a ride! Back to the question in the title: *Is IT Trustworthy?* The answer sounds simple but is not – it is trustworthy only if you make it so. Trustworthiness should be your goal and you need to plot your course carefully. Keep that in mind when you begin (or continue) your quest for this new-age holy grail.

As often is the case with a *Captain's Log*, we end with more questions than when we began. That's OK, as our goal is to stimulate your thinking. Thus, this paper is only a beginning. Look for future bulletins that will build upon what was introduced herein.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twentieth 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 25 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

Jim Baker **is a Senior Contributing Analyst for The Clipper Group.** Mr. Baker specializes in storage hardware and software, the relationship between technology and its application in business, and emerging technological trends affecting users and vendors in the Information Technology community. He recently joined The Clipper Group after three decades in hardware and software product marketing and product management at companies as large as EMC and Honeywell Information Systems and as small as IPL Systems and Datapoint Corporation. He most recently served as Research Manager for Storage Software at IDC in Framingham, MA. Mr. Baker earned a Bachelor of Science in Marketing degree from the University of Illinois at Urbana-Champaign and an MBA, also from UIUC.

- ***Reach Jim Baker via e-mail at jim.baker@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" 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**, 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.