guardtime

Virtualization and Attribution

Matthew C. Johnson,
CTO of Guardtime

www.guardtime.com

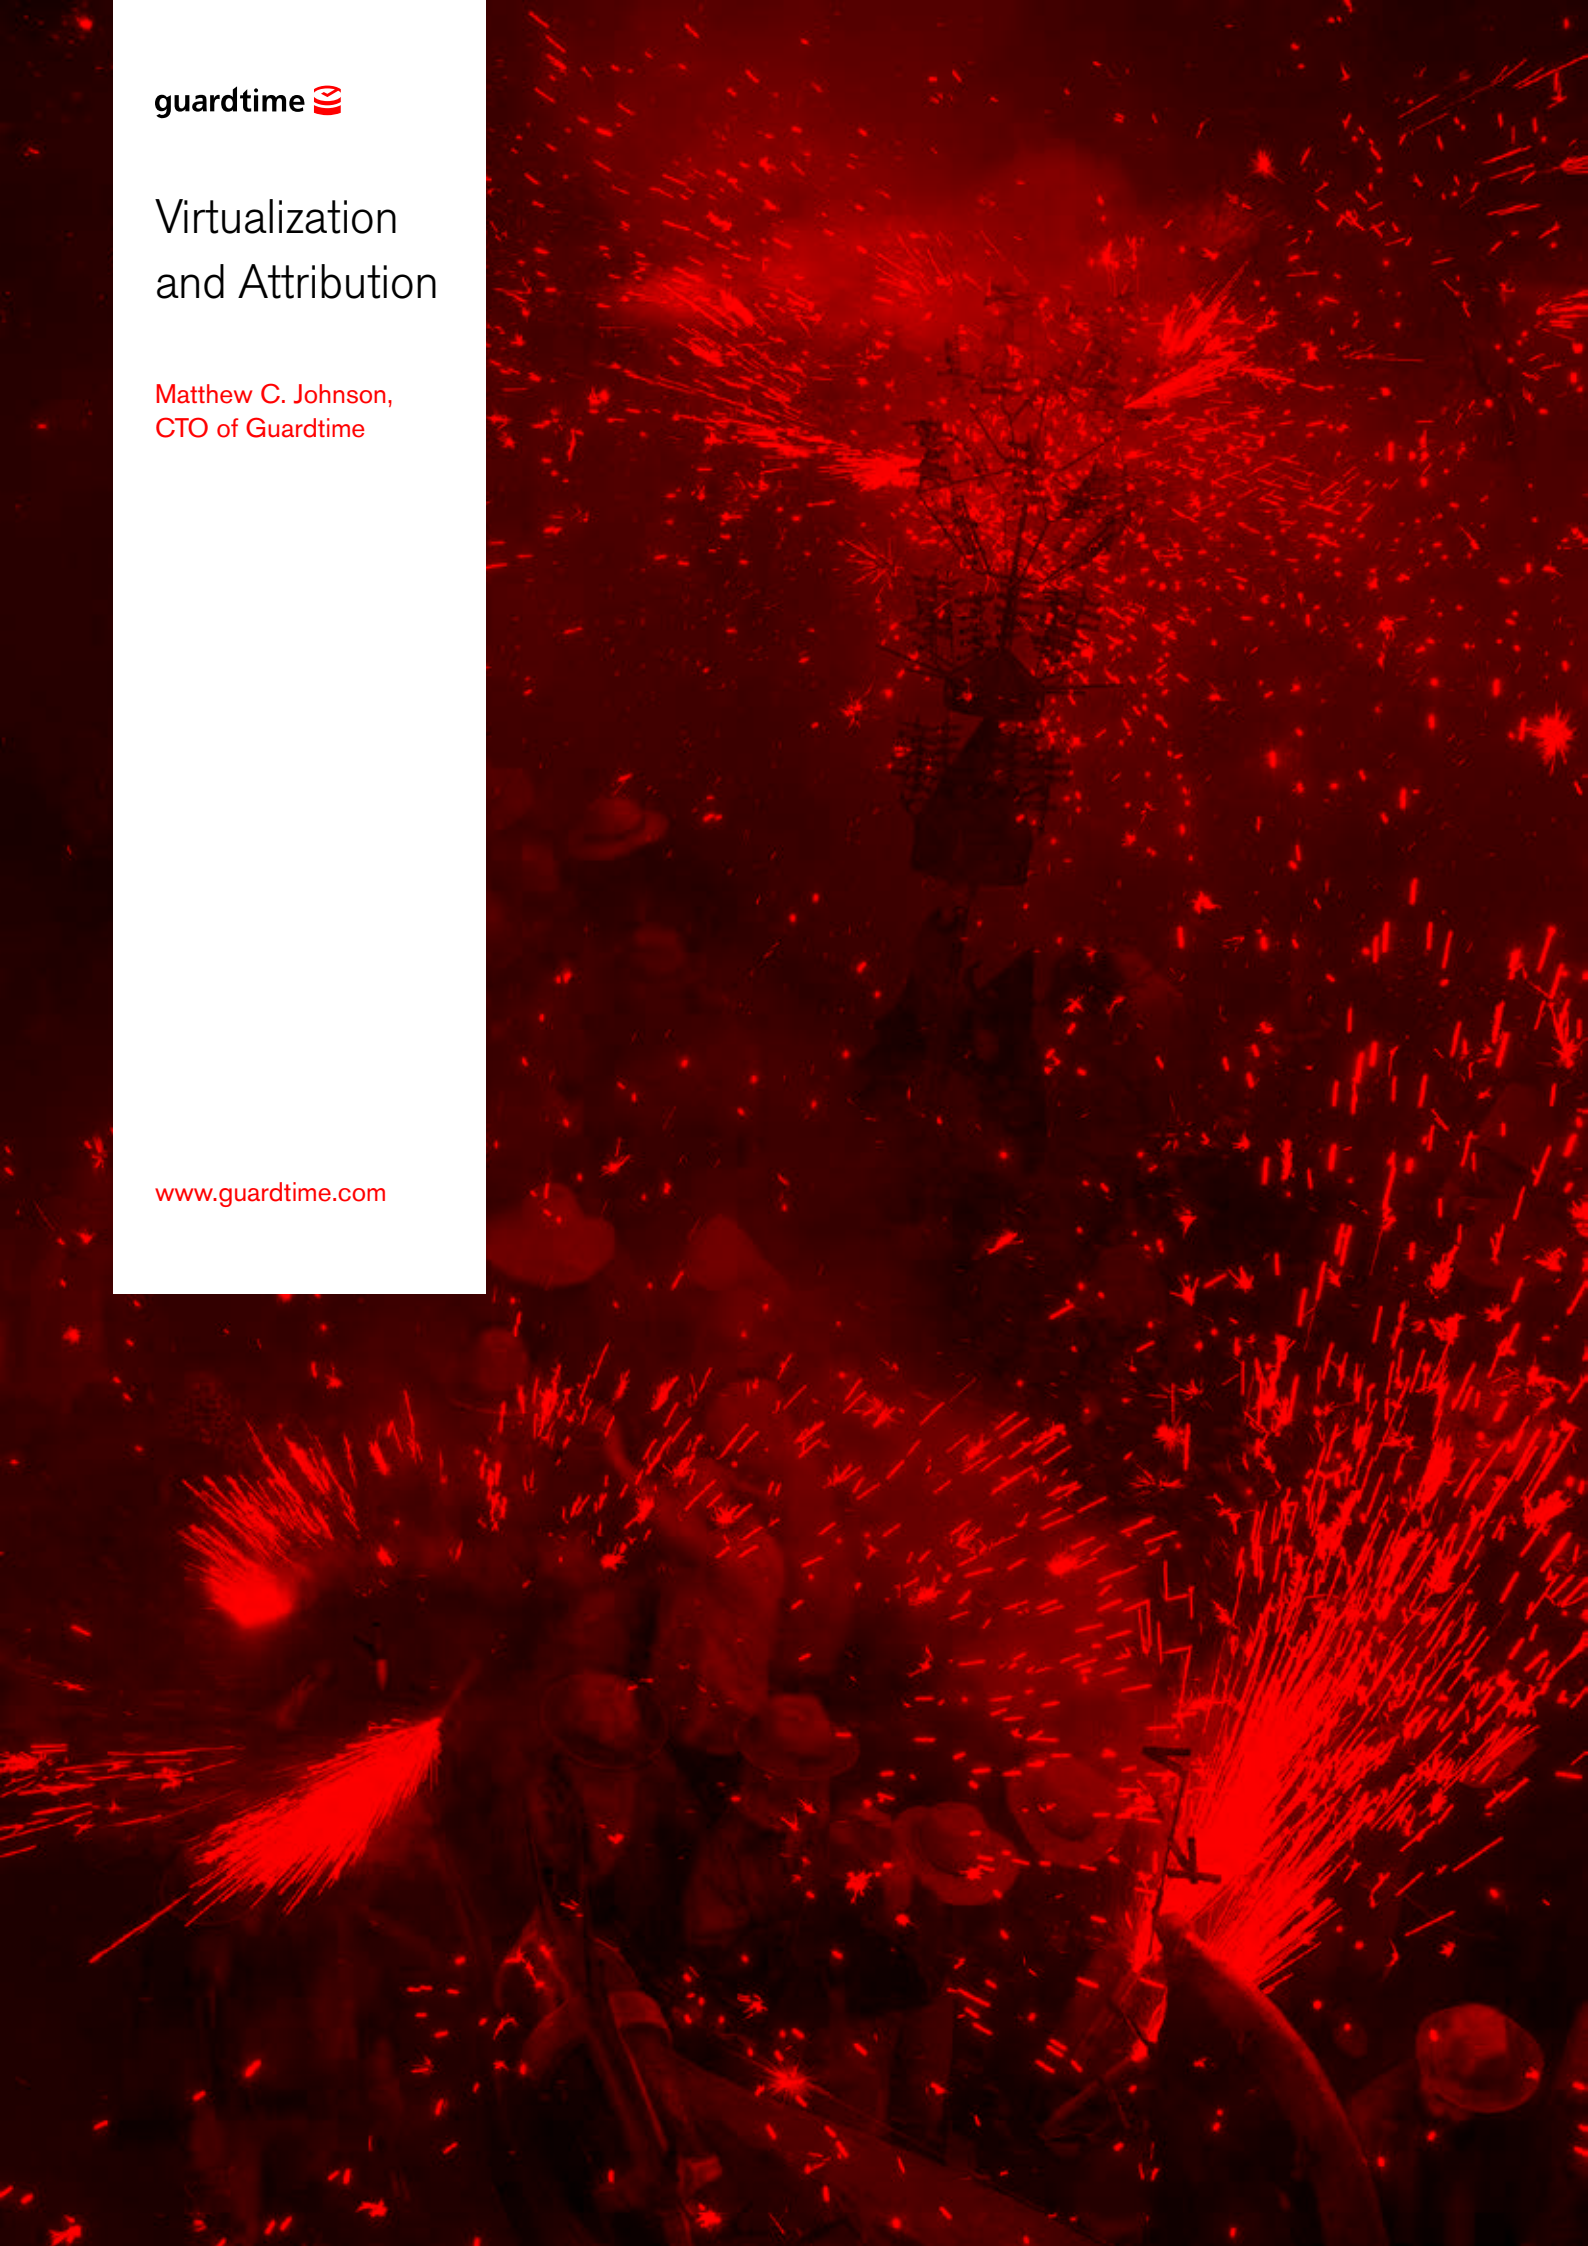

The Evolution of Systems Security

Of all the claims being made about the evolution of technology, the one that everything is getting more complex is surely the most likely to be true.

Applications are dividing into orchestrated components, networks are serving "things" instead of people, and servers/network devices are becoming virtual instead of real. Subsequently, control over these applications and systems are being increasingly abstracted to machine-based rules and decisions.

The thing that makes all these shifts complicated is that the complexity of a system is proportional not to the number of elements but to the number of relationships, and that's growing exponentially.

You can argue that the seminal trend in technology is virtualization. It's the foundation of a re-architecting of data centers, the basis for cloud computing, a key paradigm in framing the network services businesses buy, and the defining element in the most significant carrier network evolution of our time — network functions virtualization. Virtualization is the key to agility in resource use, in application design, in service creation.

The "*Internet of Things*" is a network of virtual users or virtual processes. Virtualization is our tool for managing that complexity explosion, in fact. It's also the thing that breaks every notion of security, compliance, and even business promises that we've ever had. We've fulfilled these requirements in the past by managing not our assets but our relationships, and those are exploding with complexity, disguised by abstractions, or both.

CIOs have the parallel mission of exploiting technology to serve the needs of their business, and protecting their businesses from the risks that improper applications of technology can bring. There can be no compromise in either of these areas, and in fact advances that impact one will nearly always impact the other.

You'd think that virtualization and the Internet of Things would have driven business-benefit exploitation and governance and security processes in parallel, but that hasn't happened.

We hear about issues with sensor, network and cloud security because of that disconnect, and we hear about specific remedies to get things back on track. A lot of CIOs have looked at, or adopted, these remedies thinking they'd bring their two missions into harmony, and most find out quickly that they were wrong.

You can't glue security and compliance onto the end of process changes, technology changes. You can't assure services or guarantee performance without an authenticated chain of responsibility that all parties involved can verify.

You have to create all these capabilities at the source, at the architecture level, and that's what Guardtime has done. Every CIO knows that architecture is the only defense against multiple, incompatible, inefficient, solutions. Guardtime provides that defense, provides that architecture, and most importantly shows the real relationship between trust and truth.

A Different Path for a Different Future

Everything in networking and information technology is based on resources that are digital assets. Whether they're real like servers, switches, and routers or virtual software instances representing a function or device, they're still the things that support our business processes. Logically everything we do to secure, govern, and monitor our processes should be focused on these assets, but if you look at our practices today we find that's not the case.

When we talk about encryption today, we're talking most often about encrypting traffic, which is managing the relationship between things, not the things themselves. We apply governance principles to application access—again a relationship-based approach. Service-level agreements are also applied to relationships, not to what provides the service or if it's working properly. This approach is illogical on the face because our assets are what we want to validate, but it also creates some critical problems in an age of virtualization. Virtualization creates abstract assets that relate to real ones through a chain of commitments, where the commitments are expressed through relationships. The problem is that these relationships are hidden in the abstraction process so there's no "chain of custody evidence" to independently verify their integrity. This problem impacts everything about virtual machines, clouds, virtual network services, and even APIs and software development.

Guardtime changes all of this, by applying the same basic principle of "layering" that's common today in networking. The approach is called the Attribution Envelope, and when you wrap an asset in a Attribution Envelope you create an Attributed Object. The concept is illustrated in Figure 1. At the core of an Attributed Object is the crucial concept of Identity, which means a unique mapping between an attributed "object" and the digital asset it represents. This mapping conveys authenticity to the asset; we know what it is in absolute terms.

The next layer in the Attribution Envelope is the Assertions made about the asset, which means the parameters, properties, features, interfaces it exposes.

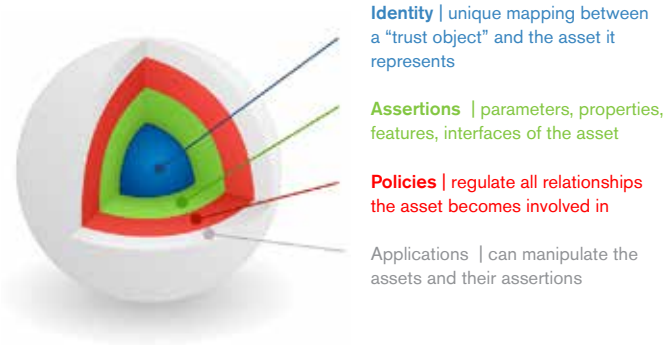

Figure 1 | The "Attribution Envelope" Concept

These two inner layers of the architecture create, for each digital asset, an authentic representation of the asset — a kind of "black box" that has properties we can now depend on. Guardtime's technology is the first approach to recognize the difference between "trust" in the sense of believing something is what it says it is, and truth, which is knowing that not only is a parameter or value or device what it says it is, but it's also exhibiting properties that are themselves authentic. These two layers of Attributed Objects are the core of Guardtime's difference.

The next two layers add in the relationships we've been accustomed to believing were the core of our security, governance, and monitoring. Guardtime's approach provides enormous flexibility in managing relationships without any loss of security, transparency, or accountability because we build relationships on top of an authentic core, which can be independently verified and attested to. With Guardtime technology we can now reliably verify the truth of the digital asset and assure its coherence independent of the service provider maintaining it on your behalf.

First, Policies are defined to be applied by the "owner" of the asset. These policies can regulate any and all relationships the asset becomes involved in and their linkage with the digital assets themselves is similarly authenticated. Policies are, in a sense, a mixture of Assertions and an outer layer, Applications. They let asset owners set boundaries on things and provide a mechanism to relate or "chain" Attributed Objects, which can then can be asserted to act as guides for how other things can relate to them.

Second, Applications are defined, which are simply things that, by operating through Policies, can manipulate the assets and their assertions. Guardtime provides certain Applications themselves, but also APIs that let the users

or third parties develop their own applications or build a Attributed Object model into existing applications.

One powerful property of the Attribution Envelope model is that a digital asset can be a collection of other digital assets, as shown in the Figure 2 below:

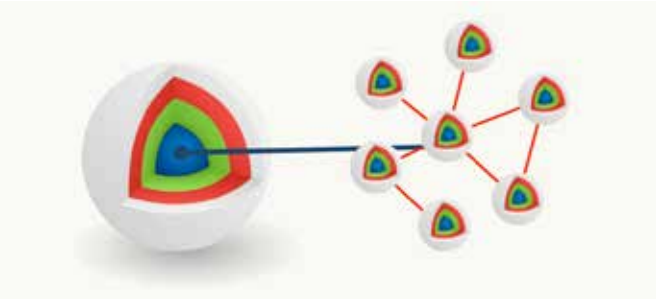

Figure 2 | Collection of Assets in the "Attribution Envelope"

In the Figure 2, Attributed Objects are linked by policies to show relationships where trust is inherited or derived from other elements. That means that in the Guardtime architecture, a complex system of assets like a cloud service can be represented by an Attributed Object and the policies in that high-level object can be used to relate it to the real assets that make it up—the servers, software, network connections, etc. This means that every complex system can be attributed to the combined behavior of its nuclear resources. With this attribution, trust, governance, and service accountability are never abstracted away. Remember the "attributed" concept, because we're going to build a whole system of asset security, governance, and management on it.

The most significant point about an Attributed Object is that it is also a "Truth Object" because we extend basic identity assurance to assertions or properties. Most security/compliance processes rely on "trust" in a very explicit sense, meaning that while they may secure the bond between things, the path or link, the user is expected to trust that an authentic link leads to an authenticated partner. However, we know this not to be true in this age of system exploitation and hacking. With Guardtime, any Attributed Object is explicitly attributable to the entity that it represents itself to be, with all the properties it represents itself as having, as precisely as DNA would identify an individual. That's what is meant by starting with the asset, and that's why the approach is so powerful.

How it Works

The critical technical barrier to making the notion of a Attributed Object work is the current infatuation with public-key encryption and key management. Assets are dynamic, hierarchical, and have different rules of association among pairs of them or within object communities. Traditional encryption from a confidentiality perspec-tive is hard to sustain at scale due to the complexities of key management, particularly in a dynamic environment, and it can't be applied in a policy-driven hierarchy unless you build a structure to do that completely outside the structure that provides encryption and certificates.

What's needed here is clearly something that's asset-based, that lets assets generate assertions that are as authentic as the assets themselves, sign exchanges among assets, verify authenticity a as needed, and above all create policy-based hierarchies where trust flows to a point of connection between two authentic domains. That's what Guardtime does with something called "Keyless Signature Infrastructure" or KSI.

KSI is based on proven principles of hashing, signing, distributed consensus and widely witnessed evidence, as Figure 3 shows. The basic concept starts with signing an asset, and that is done by first creating a hash from the asset's value, then aggregating all hash-values from all objects to be signed into a binary tree and making the root hash of that tree widely witnessed (by publishing it in a distributed data structure known as a hash calendar). The necessary hash-values to allow the original asset hash-value to reconnect to the root are then inserted into a "keyless signature" and distributed back to the asset.

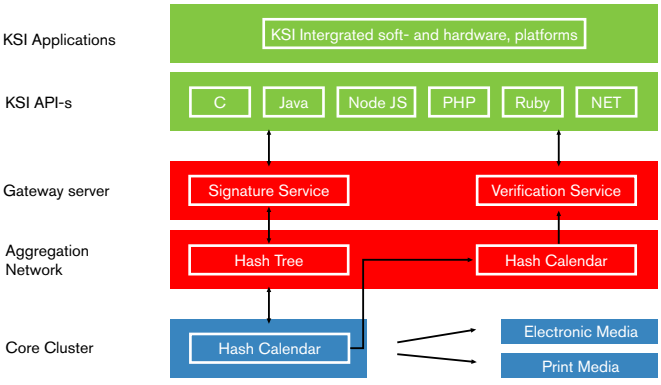

Figure 3 | KSI Architecture "Attribution Envelope"

To verify an asset, a receiver or user of that asset runs the verification by extracting the hash-values from the signature, rebuilding the tree and verifying that the root-hash value matches that in the calendar, which determines whether the asset is correctly signed and if the value of the asset is unchanged from the time of signature. If either of those conditions is not met then the asset has been contaminated.

The Hash Calendar is refreshed every second and redistributed, but gateway points cache a copy and can continue to authenticate objects offline if connection to the core is lost. A key innovation is that the Hash Calendar measures the passage of time without reliance on a trusted time-source and thus the signatures automatically include a time component effectively replicating PKI digital signatures and digital timestamps using only hash-function cryptography.

The services of KSI are distributed through a set of APIs supporting virtually all the popular programming languages and through a port to the gateway device. Applications can use these services to sign assets or validate signatures as needed, and users and partners can construct any software structure they like around the KSI framework or integrate the APIs into any existing structure. However, the optimum model for KSI application is the Attributed Object, and Guardtime builds their own applications of KSI around the Attributed Object model.

Logically speaking, digital assets are known by their properties, called Assertions. The contents of a log file, the state of network or application variables, the interfaces to devices or services—all these things are Assertions made by digital assets. Assertions of a Attributed Object can be attributed to that object through the signature and verification process. The Assertions can be native to a real "atomic" asset or derived from Assertions made by lower-level assets in a complex structure like a cloud service. Because Assertions are attributable to Attributed Objects they create a mathematically verifiable chain in such a complex structure, and that means that they can form the basis for claims of security or compliance, including SLA compliance.

Because Assertions are the properties of Attributed Objects, they are also the logical place to define and apply Policies that either show how the Assertions can be used or changed by others, or how a "service-level" Assertion is derived from the Assertions of the digital assets that fulfill the service. A service, whether it's a network service, an application service, a cloud service, or whatever, is considered to be Attributed when it is made up of Attributed Objects and its own service-level Assertions are policy-derived from those objects. There are specific examples of implementing this model in Attributed Networks and Attributed APIs models, described later.

Any application can tag/sign assets or verify signatures, and any asset can be signed. That means that attributability can be added to any interface or set of assets, and that asset collections can be made attributable. Tagged assets carry a history of changes with them as they move as workflow elements, or as workflows operate on them. Tagging can be carried over a network, applied to control and management exchanges or even to workflows. This means that Attributed Objects can have attributed exchanges with each other, and that all parties involved in any transaction can be authenticated at any point by anyone who receives a tagged asset.

It's important to note here that while the KSI model is the most comprehensive security and compliance tool available, it has the ability to authenticate any and all assertions made by an Attributed Object, so any essential property of a business digital asset can be guaranteed through attribution.

In the Attributed Object structure, the outer layer is the Application layer, where attributed services are consumed by users to solve business problems. Guardtime's own applications for managing the KSI framework and creating security and compliance solutions from Attributed Objects. These applications include:

- **GuardVision** – Attributed Networking Integrity Management and Advanced Analytics
- **GuardView & Videri** – Advanced Integrity Analytics and Big Data Security Platform. Digital Asset Protection, Data Loss Prevention, Security Management Rules and Governance based on Integrity Monitoring and Evidence

- **GuardSweep** – Integrity Baseline of Network and Object Store Resources, Associated Applications, Configurations, and Access, Authorization, and Authentication Policies

- **GuardLog** – Integrity Log and Audit Monitor for Applications Responsible for Audit and Log Monitoring and SIEM Escrow. Benefits Risk, Compliance, Retention & Audit Solutions

- **GuardView TPV (Third Party Verification) Managed Security Services** – Integrity as a Service and Integrity Escrow Services to Perform Real-Time Incident Response, SLA Assurance for Users/Operators to Verify the Integrity of Contract SLAs and Technical Security Controls Described in Service Provider Contracts.

In operation, KSI would be applied to sign the digital assets of a network, an application, a data center, or an entire business. These signatures allow KSI applications to validate the state of any asset, and thus the entire asset set, based on a 'clean' reference state. When any asset deviates from the baseline the asset can no longer be trusted, and the state of the assets in any asset system can be escrowed and exchanged for validation without providing access to the assets themselves. That means that anyone can be given proof of the authenticity of an asset, even a complex one like a cloud or data center, and can validate that proof independently. KSI-specific applications can receive real-time notifications of trust issues to provide asset owners with alerts to problems with security, governance, and even SLAs.

The SLA application of KSI may be the most compelling. SLAs are expressed by a provider and relied on by service consumers, and yet it's normally almost impossible to determine if the promises they contain have been met without providing monitoring of service resources that would itself breach security and confidentiality for the providers and even other users. By making the SLAs Assertions of the appropriate service interfaces, and by linkingthem back to the assets that fulfill the service requests, the SLAs themselves become Attributed and a promise of compliance is enforced just as a promise of security would be enforced.

An Assertion is a promise—of security, compliance, or conformance — as much as it is a parameter or data value. By creating the link from a provider's promise to the behavior of the infrastructure that fulfills it, KSI Assertions convey attributable commitment without exposing the details of how the commitment was fulfilled, and even a change in the mechanisms for fulfillment can be made into Assertions and published so that those who receive a promise know not only that it's being kept, but that the mechanism for doing so was changed. In every area where digital assets represent commitments to something, KSI lets both the provider and the consumer of the assets believe what they're told.

Critical Applications: Attributable Networks and APIs

It's time now to return to that earlier point about having the properties of complex systems attributed. Even today, services like VPNs or cloud computing are consumed through APIs that are themselves abstractions of what's likely complex infrastructure behind the scenes. The key term here is "behind the scenes" because while consumers of services are expected to trust not only their security but their governance and even their SLA, there's no way to know whether the service is in fact trustworthy. Attributability is a key element in our trust/truth relationship because in order for something to be attributable, it has to be both associated with an authentic source and second be an authenticated assertion of that source.

The challenge with attributability is that we have typically focused on individual interfaces and connections to secure and trust, and yet we consume complex systems of devices and software that have complex relationships and properties. If a network or cloud user has to establish trust and truth a device or software component at a time, there will never be a useful model of either.

Fortunately, a complex system of any sort can be Attributed by tagging its critical assets and their assertions/properties for verification. The attributes can include the physical devices, software images, log files, data files, management state and parameters, or any other information. The current state of an "Attributed Complex" can be measured against the Clean State for the system to validate the system is what it claims to be, operating as it should be. There are two examples of this process that are highly important in today's technology market; one is the network and the other is the generic notion of the "service API" that is prevalent in componentized applications, virtualization, and the cloud.

The Attributed Network model for KSI is an example of chaining across multiple assets. The process starts by making all the network devices into Attributed Objects that contain a chain of Attributed Objects representing the versions of hardware, software, and firmware, the operating parameters, and critical state variables. These Attributed Objects can then be chained upward to define the network, either as a single step for homogeneous networks or through one or more layers of administrative or physical segments. A Clean State is established for the network complex, based on a baseline operating state that is known to be valid.

This process in itself will validate the elements and configuration of an entire network. It's also possible to extend this basic model by adding tagging to critical management and control exchanges and even to data exchanges. Tagged messages such as this will allow elements of the Attributed Network to validate each other by authenticating key traffic that could change network behavior.

The interfaces and services of the Attributed Network are delivered through interfaces and/or APIs and these can also be made Attributed Objects, which means that security, compliance, and performance can be validated to network users without exposing detailed management and configuration data or exposing network devices to hacking. This process is similar to that used to create Attributable APIs, which is our next topic.

In a software-driven age, application program interfaces (APIs) of all types are used to deliver application, network, storage, and compute services to users and to connect elements of complex services, applications, and experiences. In virtualization and cloud computing, we use APIs to represent resource pools that can be assigned to applications ad hoc. Increasingly, we're orchestrating application components using workflow engines, and with the advent of public cloud services some of these components will have to be cloud-hosted. That creates a major risk of security problems, compliance/governance issues, and even SLA disputes.

To Guardtime and the KSI architecture, the property of "attributability" can be gained by having a Attribution Chain established to authenticate the constituent elements of a complex system. That's how any complex system with one or more APIs can be attributed. The application of Attributed APIs to security and governance are clear, we think, but the fact that they can also apply to service level agreements is revolutionary, and demonstrates the flexibility of our approach.

Any "service" has to guarantee something to the user, whether it's simple availability or more stringent perfor-

mance boundaries. It's common to have SLAs and just as common to have disputes over whether they are being met. It's hard to resolve SLA issues because of visibility; no service provider will offer users management visibility into their infrastructure, and without common monitoring there's no basis for agreement on performance.

What KSI can do here is to make SLA conformance of a service into an Assertion that is linked to an Attribution Chain that leads back into the provider's network. If the API's SLA Assertion(s) are related to deeper assertions of performance, customers can be allowed to follow the chain to see that indeed network/service elements are meeting an SLA without seeing the objects themselves. SLAs are then authentic because they are attributable to the performance of the network and IT elements within.

Any attribute of an API or complex system, even the security and authenticity of a sensor network in M2M, can be authenticated this way. Anything that is guaranteed, then, can be made trusted and thus verified independently from those operating the system.

All of this is possible because of the dynamic nature of KSI's sign/validate approach. Because Attributed Objects can be generated with Assertions in any volume, there is no system too complex to be made Attributable. In a world where "real" resources and services are a thing of the past, that's a critical benefit.

Facing the Future Means Facing the Right Way

Despite the fact that we've had sensor networks for decades, the Internet of Things would change everything simply because of the scale of what is available and the ways that dissemination of data and access rights then needs to be controlled. Despite the fact that nearly every enterprise uses virtualization today, the extensions of virtual networking and cloud computing will explode the number of services and resources that depend on virtualization as we move to the second half of this decade. Mobility, experience management, and customization will further componentized software and demand software composition processes be even more agile. And these trends will reinforce each other, as we already know.

What we "know" in a truth-and-trust sense is critical in this future filled with elastic and agile concepts and relationships. Knowledge of this kind can't be conveyed by a simple exchange of a certificate, can't be protected by encrypting known paths of information exchange. We must fall back to protecting the digital assets themselves, the things that demand protection. We must build our secure, compliant, conformant service relationships from assets that we can validate, not only in terms of identity but in terms of properties. A hacked secure system is not only not secure in itself, it's a conduit of contamination to every other IT element that trusts it, yet traditional security and compliance processes would not spot the problem.

Any system that's more comprehensive, more agile, more flexible is also more complex. Practices that manage complex systems but don't authenticate their elements or properties are not guaranteeing proper operation at all. We are building toward a future where our current practices will fail us at the most fundamental level, the level of those attributes of trust and truth that we've talked about so many times here.

KSI can authenticate not only the identity of elements but their properties, make trust-and-truth a combination and not a contradiction. By building on the inner layer of Identity and Assertions in Attributed Objects, Policies and Applications can build a knowledge of the trust and truth of IT and network elements that has never been available before, and use that to set service levels, apply governance, secure assets, and build applications that are able to do what they promise because they use only what is secure and authentic.

Guardtime is offering a whole new way of looking at digital assets, services, networks, and applications. It's a way that reflects the critical importance of all aspects of information technology to modern businesses and the need to secure the technical elements of our infrastructure and applications that support workers, partners, and customers.

About Guardtime

Founded in 2007, Guardtime invented Keyless Signature Infrastructure (KSI) - a technology that allows any type of electronic activity to be independently verified using only formal mathematical methods, without the need for trusted administrators or cryptographic keys.

Deployed by world governments, KSI provides real-time validation and an independent audit trail for everything that happens in digital society, limiting liability and making it impossible for insiders or sophisticated cyber attackers to manipulate data and cover their tracks.

Read more at http://www.guardtime.com.


guardtime
© 2014 Guardtime