On Trusting Third-party Satellite Data

Sean Crosby and Kurt Brenning
Sandia National Laboratories
P.O. Box 5800, Albuquerque, NM 87185-0980
smcrosb@sandia.gov

ABSTRACT

Increased access to space has opened the door to many satellite vendors. These vendors are collecting data using a variety of sensors, including electro-optical, radio frequency, and synthetic aperture radar. Customers want true-sourced, authentic data. However, as with any lower barrier to entry, the risk of counterfeit, tampered, or low-quality products increases. In this work, we describe the key requirements for trusting imagery and present a hardware design and a system of controls that meet those requirements of trust for Earth imaging satellites. Our trusted hardware provides assurance of capture time, location, and preserves the content and origin by capturing and digitally signing the information end users need to make trust decisions about the data. Our hardware functions as an independent witness that oversees and signs off on satellite collection activities. Anti-tamper, inspection, and verification measures protect and verify the secure operation of our hardware. Satellite operators that use this approach in their satellites and operations will offer their end users greater assurance in the authenticity of the produced satellite imagery products.

INTRODUCTION

Satellite imagery is more available now than ever before. Increased access to space and improved constellation management are revolutionizing the way we see the Earth. Service costs are decreasing and ground coverage and revisit frequencies are increasing. Output image products are used for many objectives, including agricultural management, environmental monitoring, and defense. These eyes in the sky are feeding decision makers on the ground.

Customers need true-sourced, authentic data. However, as with any lower barrier to entry, the risk of counterfeit, tampered, or low-quality products increases. As satellite production and operation moves out of house, end users have less knowledge and control of the satellite systems and their ground-based components. This separation decreases the end users' ability to verify the integrity of the systems that capture, process, and disseminate the data.

Good data can be poisoned through interception by an adversary or through processing on a tampered or hacked device. Devices can be impacted through direct access or through a supply-chain attack. Bad data can also originate from a simulator, a generative adversarial network (GAN), a replay attack, or some other means. Even satellite imagery is susceptible to image compromising attacks and new machine learning-based deepfake geography generators[1–3] are making it harder for humans and machines to detect fake data. Without knowledge of the original data, it may be very difficult to identify bad data. End users need a means to verify the origin of satellite imagery and the correctness of its content and metadata.

Methods for providing assurance are under development, but none of them offer end-to-end guarantees of what was seen by a third-party sensor. Cryptography is being used to protect data in flight and preserve actions performed by trusted and/or semi-trusted entities, but by itself cannot validate the output of a third-party remote sensor. Forensic analysis algorithms can identify some types of image tampering, but this is a cat and mouse game as detectors are not created until threats are identified. Background checks of company personnel and accreditation programs for information technology can help, but the inherent complexity of these systems makes it hard to know whether a system is truly trustworthy or not.

While all of these methods have merit, they must be coupled with an onboard verification mechanism and end-to-end cryptographic protections, to deliver the highest assurance to end users. In this work, we identify the attributes of the image collection process that must be verified and propose a trusted hardware architecture for doing so. This trusted hardware acts as a digital notary public that signs off on the collection activities of the third-party satellite. Securely deploying trusted hardware on an untrusted

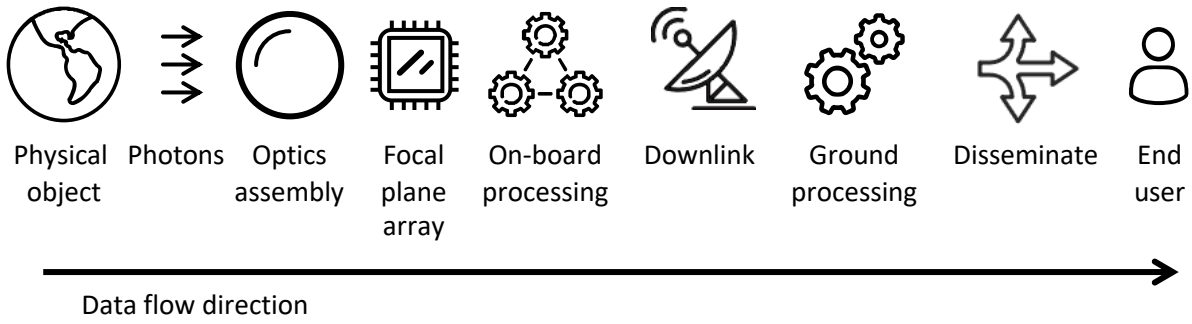

Figure 1: The lineage of a satellite image

platform poses many challenges. Anti-tamper technology, inspections, and remote verification are all required. Our intent is for this technology to be implemented, keyed, deployed, and monitored by an independent trusted party.

Satellite operators that utilize these principles will offer true end-to-end assurance of their image products to their end users, regardless of company or government affiliation. We are preparing a hardware prototype to fly and remotely verify on the International Space Station early next year.

PROBLEM STATEMENT

Satellite imagery collection and dissemination processes involve several systems and components as depicted in Figure 1. The start is when a physical object emits photons, which propagate up and through the satellite lens onto the focal plane array (FPA). The FPA transforms the signal into a digital frame and delivers it to an onboard processor. The onboard processor may perform some enhancement or correction and then it packages the data for downlink. When the data is received at the ground station it may undergo processing. Only after that processing is completed is the imagery then delivered to the end user.

In this process, different components of this data path may be owned/managed by different entities or organizations. Current methods require each stage of the data pipeline to blindly trust that the previous or following stages did not maliciously modify the data, utilizing only error correction codes to ensure data is not corrupted accidentally. Additional protections should be implemented to more effectively secure image data, and give the end user higher assurance that the data received has not been tampered with. Adding cryptographic signatures to data at the time of creation will increase confidence in the accuracy and authenticity of satellite imagery, and will reduce the risk of making critical decisions with incorrect or out-of-date information.

In addition further security mechanisms must also be included in a reliable system to ensure the cryptographic integrity of the system is preserved, and to protect from malicious attackers.

Threat Summary

Each stage of the data processing pipeline is subject to a variety of threats that may impact the accuracy and usability of the final product. Without mitigation, there is an increasing risk that counterfeit or modified data will appear. It is our observation that it is harder for intruders to gain access to the systems and components further upstream than it is for them to access those systems on the ground and closer to the user. Here is a summary of some potential attacks.

There are multiple opportunities to replace or modify the imagery data after it is downlinked from the satellite. At this stage, during ground processing or dissemination, there is no way for the end user to know whether or not the image has been modified or replaced. To insert false data at this stage, an attack could fabricate new satellite imagery or modify existing imagery to create new data, and deliver that to the end user. Powerful image editing tools and deepfake models can produce an image that is nearly indistinguishable from an authentic image. This replacement or modification can be accomplished by compromising the downlink, the ground processing, or the dissemination mechanism.

To compromise the image earlier in the data pipeline, an adversary would need to somehow compromise the space vehicle itself. This would require physical access to the system, or a modification of

the hardware/software being designed and deployed. While less likely, as many space vehicles are joint efforts, it still represents a threat to the design. An attacker could attempt to compromise the security of the focal plane array or on-board processing by replacing or modifying the device during system integration.

Obviously, an adversary can also protect a physical object from being imaged by simply obscuring it from view so it can't be seen from orbit. This is only possible with small and/or portable objects, and isn't particularly feasible as the adversary rarely knows precisely when an image will be collected.

In addition, once cryptographic signatures have been added to the pipeline, they also present an attack surface to a potential adversary. It is essential to protect the integrity of the device by ensuring that the secret key is sufficiently protected, and verifying that the device implementing these security features has not been compromised. If these protections fail, an adversary could use the device or the secret key to sign invalid or modified data.

Risks

There is a fundamental integrity risk in using imagery collected by third-party remote sensors as it is difficult to determine whether the integrity of an image has been compromised. The components of third-party remote sensing platforms and ground stations are untrusted by the end user, and it is very difficult to verify the authenticity of the imagery, as third party companies will commonly use proprietary algorithms to process their data.

The use of satellite imagery as input data for decision makers and algorithms has become ubiquitous in the 21st century. The cost of using data that is incorrect is difficult to quantify. In some cases the risk can be small, such as inaccuracies in the in-game map for Microsoft Flight Simulator. In many cases however, inaccurate data could result in loss of life. Militaries rely on this data for training and deployment. Autonomous car companies rely on this data for navigation. This data is used to track forest conditions to estimate risk of wildfires. In these cases and more, increased trust in an images authenticity reduces risk.

RELATED WORK

Researchers have proposed many methods for ensuring or checking image authenticity, including digital signatures, digital watermarks, forensic analysis, and so forth.[4] Forensic analysis is a continual cat-and-mouse game between attackers and defenders[5] does not provide the level of assurance we desire. For this reason, we chose to focus on cryptographic methods for establishing image authenticity.

Authenticated cameras produce imagery where the content and source can be verified. These cameras were conceived many years ago.[6,7] Researchers have developed cameras that use Field Programmable Gate Arrays (FPGAs) with embedded Physical Unclonable Functions (PUFs), to digitally sign imagery with asymmetric keys as it is captured.[8] Our focus is in adapting authenticated camera technology to a space environment.

Several organizations are working to establish standards and means for enabling authentication of imagery[9–12] for combating disinformation. The Coalition for Content Provenance and Authenticity (C2PA)[9] is standardizing approaches for securely establishing the pedigree of multimedia. The basic premise is to compute a digital signature when the image is created, and then when the multimedia is updated, the changes are added to the history and the multimedia is re-signed by the editor. Downstream users can then verify that the data was created and edited by the said parties and that no changes were made elsewhere. The C2PA specification does not try to establish that the provenance data is 'true', but enables verification of its association with the asset, correct formation, and tamper-free status.[13] We take this approach one step further by sourcing the original data from a trusted source.

Another approach for ensuring third-party correctness is to investigate the third-party's infrastructure and establishment. For example, Kinser et al. proposed a methodology for establishing cyber trust scores for space enterprises.[14] While these accreditation processes are good, there can be a gap between perceived state and actual implementation.

APPROACH

In this section, we discuss the requirements for trusting an image and propose an architecture[1] and complimentary routines for meeting those requirements in remote sensing applications.

Requirements for Trusting Imagery

An end user can only trust an image if the user trusts the image origin, has assurance of the time and loca-

[1] Our approach to creating authenticated imagery for remote sensing applications is patent pending

tion of the capture, and has assurance that the image and its corresponding metadata were not modified by an unauthorized party (See Figure 2). Here we build upon the early work by Kelsey, Schneier, and Hall[6] to define six requirements for trusting an image.

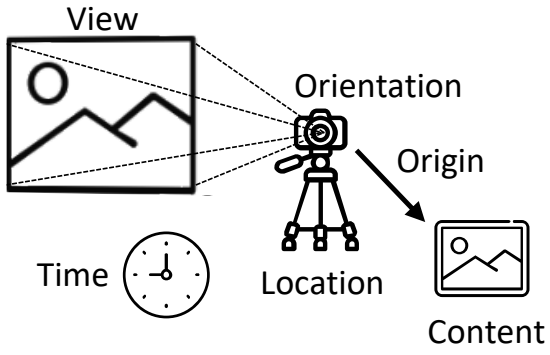

Figure 2: Six attributes of image capture that must be trustworthy

Trust Requirement 1 (Time) *The provided capture time must be trustworthy.* Without a well founded timestamp, the image could be incorrectly associated with the wrong time, as is the case with a replay attack.

Trust Requirement 2 (Location) *The provided capture location must be trustworthy.* This includes the location of the imaging sensor and geospatial or other location information associated with the scene. Without verified location information, the image could be incorrectly associated with another location.

Trust Requirement 3 (Orientation) *The provided orientation of the imaging sensor must be trustworthy.* The orientation of the imaging sensor provides important perspective information, particularly in the case when only sensor location information is available. An image without orientation information could be incorrectly associated with the wrong scene.

Trust Requirement 4 (View) *The view of the imaging sensor must be real and unobstructed.* Blocking an image sensor or feeding in a fabricated scene can circumvent all other security controls. Images should be free from staging or other in front of the lens manipulations. Unauthorized staging is a risk whenever the operator of the imaging sensor does not control every aspect of the scene, which is always the case in remote sensing.

Trust Requirement 5 (Origin) *The true origin of the resulting image must be known.* Associating an image with a specific piece of hardware connects the image with the organization that is operating the hardware. This requirement is not met if the imaging hardware can be cloned or if another device can produce an image that appears to have come from the listed source.

Trust Requirement 6 (Content) *The content and any associated metadata of the resulting image must be authentic and free from any unauthorized modifications.* An undetected change can fully compromise an image.

When all of these requirements are met, end users are able to have greater trust in the authenticity of an image. If one or more of the requirements are not met, then additional measures should be taken, such as using forensic analysis tools for identifying tampering. While later measures are helpful, they may not offer the same assurance as meeting the trust requirements.

Authenticated Imagery for Remote Sensing

While meeting these requirements in a controlled environment on the Earth is possible using existing authenticated camera technology, meeting these requirements for an imaging satellite in space brings additional challenges. When the satellite is built, owned, and operated by an untrusted or semi-trusted third party, additional measures are required.

In the remainder of this section, we lay out our approach to meeting the trust requirements for imaging satellites. While our approach is suitable for third-party satellites, it offers benefits to satellites built and operated in-house.

Our approach to meeting these requirements is to employ trusted hardware, inspections, and verification routines. A trusted witness is needed to sign off on the collection of an image. This is analogous to a notary public that would observe a client taking a photo and then notarize the resulting image. We describe this *Trusted Signing Hardware* below. As hardware is vulnerable to attack, we propose inspection routines to keep the hardware safe prior to launch into orbit. This is analogous to a notary public going through an exam and a background check prior to receiving a commission. Once deployed, routine verification tests of the Trusted Signing Hardware is required. Continuing the analogy of the notary public, these verification routines are the random reviews of a notary's journal records.

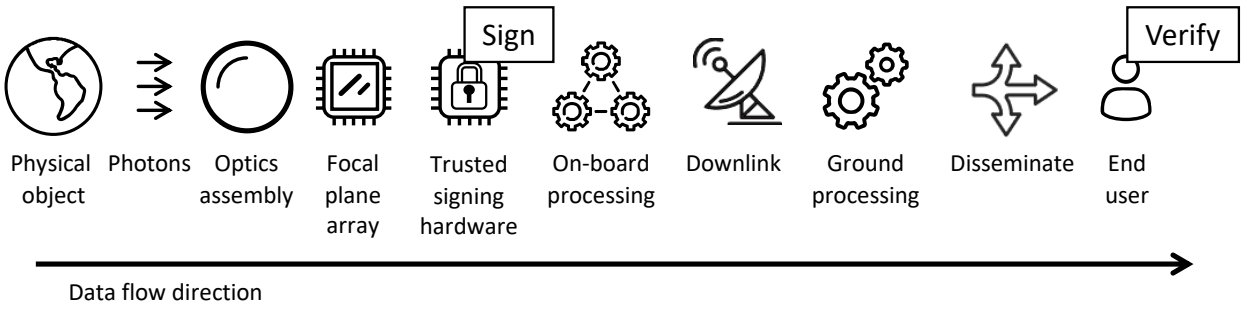

Figure 3: Independent trusted hardware signs the data immediately after capture. Any end user with the public key can verify the data.

The organization that builds the trusted hardware and performs inspections and verification should be trusted by both the satellite operator and the end user. Having an organization that is independent from the satellite operator can lower the chances of insider threat. One option is to utilize a third-party standards laboratory.

Trusted Signing Hardware

A typical authenticated camera produces imagery where the origin and content can be verified. This is performed by capturing the image and related metadata (including time and location) and then encrypting or signing the data in some fashion using encryption keys. If using asymmetric encryption, the public keys can be widely distributed and anybody with the public key can verify the authenticity of the content.

In a remote sensing application, it is best to sign the data as early as possible. The data path (as depicted in Figure 1) makes up a broad attack service. However, a digital signature can protect the image from unauthorized modification and secures portion of the data path and attack surface that falls between the generation of the signature and verification by the end user. For this reason, we sign the data as close to the analog to digital conversion as possible (see Figure 3).

Our approach is for the Trusted Signing Hardware to subsume the responsibility for reading the raw frame off of the focal plane array (FPA). It also gathers time, location, orientation, and other metadata and generates a digital signature for both the image content and the metadata. Two of the hard problems associated with authenticated cameras are binding time and location to an image.[6] The hardware package must provide an independent time source and an independent sensor for providing location and orientation information. In our design we use an Inertial Measurement Unit (IMU) to capture the payload orientation and to show that the payload is in a space environment. After the digital signature is computed, the image and the signature (either separately or together). At that point, the image and signature can be downlinked by the untrusted satellite and distributed by untrusted distribution networks. End users with the appropriate encryption certificate can verify the origin and content of the image by checking the digital signature. This design is depicted in Figure 4.

As this hardware can be deployed on an untrusted platform, it must be protected from modification. Anti-tamper methods, such as tamper evident packaging must be utilized to protect the Trusted Signing Hardware from modification. Industry best-practices must be followed for key distribution and protection. Theft of a private or secret key would compromise the ability to know the origin of a piece of data. As such, the trusted hardware needs to employ measures to prevent theft. Some modern Field Programmable Gate Arrays (FPGAs) support secure boot and so configuration and keys are stored unencrypted can can only be decrypted with a Physical Uncloneable Function (PUF) derived encryption key.[15] We recommend these sorts of measures for protecting encryption keys.

Performing the image signing in an isolated FPGA reduces the risk of compromise. This is similar to an approach taken by researchers working in participatory sensing,[16] who have identified solutions for enabling trust of crowdsourced data from mobile phones and other Internet of Things (IoT) devices.[17,18] Many of the proposed techniques depend on each device having a root-of-trust established in trusted hardware and software, such as Arm's TrustZone[19] or other Trusted Execution Environ-

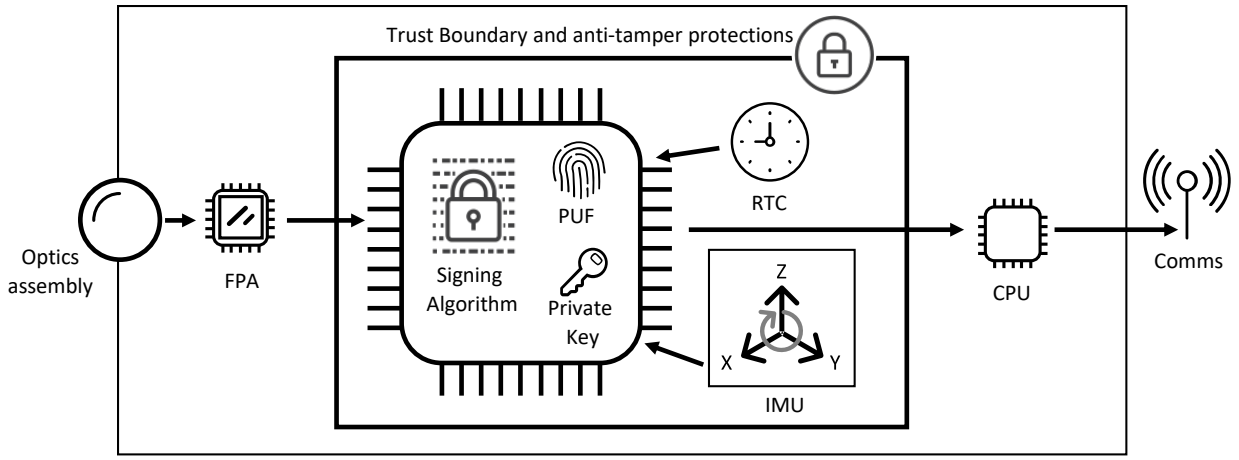

Figure 4: The architecture of the Trusted Signing Hardware

ment (TEE) technology. Modern computers perform some of their most critical operations in TEEs. A TEE protects operations from application-level threats living on the same physical machine. The use of such hardware can enable the secure generation of data on third-party platforms, given that the data is digitally signed or otherwise encrypted before leaving the secure space.

As size, weight, and power (SWaP) is an ever present concern for satellite missions, the form factor and specific anti-tamper techniques utilized will need to be crafted for each satellite. As FPAs are generally implemented as an application-specific integrated circuit (ASIC), there is an opportunity to integrate the signing logic into the FPA's ASIC. This could save space and power, though the independent time source and location sensor and its connection to the ASIC would still need to be protected.

Once an image is digitally signed, it cannot be modified without detection. A digital signature is created by hashing the data with a hash function and then encrypting the hash with symmetric or asymmetric encryption. Our current prototype uses SHA-3-384 as the hash function and RSA-4096 as the asymmetric encryption algorithm. However, raw satellite imagery typically needs to undergo correction (such as radiometric correction and orthorectification) and enhancement (such as geocorrelation) prior to being useful to non-expert end users. These actions could be performed onboard, within the Trusted Signing Hardware, prior to generating the digital signature. Otherwise, the modification would need to be performed downstream by the end user or by a system trusted by the end user.

Theoretically, a hash collision could be found which could result in a duplicate digital signature. Currently, the key sizes selected meet or exceed the minimum recommended sizes in the CNSA standard.[20] As an additional protection and to ensure no images are omitted when receiving a continuous sequence of collects, we suggest including the hash of the previous image in the metadata for the following image to establish a chain of hashes. The availability of this linked list will enable end users to detect issues by locating gaps or forks in the linked list of hashes, and identify potentially missing data.

Other challenges will need to be addressed on a platform-by-platform basis, including the power usage required by encryption algorithms,[21] identification of image modifications required by commercial imagery providers, and design of interfaces and protocols between the Trusted Signing Hardware and the onboard processor.[22] The trusted hardware must interface with the FPA and the untrusted onboard processor. Great care must be taken to limit interfaces and protocols to the required set of operations.

Inspection Plan

Securely deploying trusted hardware on an untrusted payload poses many challenges, especially when the satellite engineers and integrators are themselves untrusted. The anti-tamper controls mentioned in the last section protect the hardware and keys from certain types of threats. However, given sufficient access and enough time, they can be defeated. As such, additional process is required to protect these mechanisms. For pre-launch protection for production units, we recommend the following actions:

1. Provide representative test units for satellite developers to use for design and test.
2. Manufacture the Trusted Signing Hardware in a trusted facility.
3. Limit the amount of time the third-party has unsupervised access to the Trusted Signing Hardware.
4. Review designs and visually inspect the satellite to ensure that nothing sits between the FPA and the Trusted Signing Hardware. Ensure that the communication lines can not be reconfigured post launch.
5. Inspect the Trusted Signing Hardware for tampering at all key points during integration, especially before it is placed into a compartment that reduces visibility. Use seals when appropriate.
6. Burn-in the encryption keys into non-volatile memory as close to launch as possible.
7. Retrieve and inspect all test and otherwise unlaunched units for tampering.

Once the satellite is launched, access to the payload becomes much harder and ongoing physical inspection is not required. Tamper events must still be reported to operators and end users.

Verification Routines

Additional actions are required to verify the proper deployment of the Trusted Signing Hardware to orbit and the proper function of the hardware. All three of the following routines are required:

In addition to these two tests, ongoing inspection of data products is required. While it is theoretically possible to find a hash collision with SHA-384 or to attack the encryption of a single record, it is challenging and computationally infeasible to do this on an ongoing basis. For this reason, we provide the hash from the previous image in the metadata for each image. This creates a chain of trust that can be checked, even by a standard user. Other ongoing tests include ensuring that the timestamp is monotonic in that it never is earlier than the timestamp of the previous data record and checking IMU readings against the expected look angle for the reported imaging position in orbit.

Verification Routine: Image Product Verification
Frequency: Upon receipt of each image
Requirements checked: Time, orientation, origin, content
Description: The trust attributes of image product are to be checked upon receipt. Consistently checking each product provides a consistent view of clock and IMU readings which can reveal a compromise of the clock or the IMU sensor within the Trusted Signing Hardware. The digital signature should be checked for every image using the public key of the Trusted Signing Hardware that signed the image.

The clock reading should be monotonic. If the timestamp ever moves backward, then this indicates tampering or an update to the time. The IMU sensor reading should be evaluated to ensure that it matches the off-nadir angle reported elsewhere in the image.

Finally, each image will include the hash of the previous image in its metadata. This chaining forms a linked list of hashes that should be checked. The chain can be broken by lost images or by receiving an image that was signed by a leaked private key. This check is important in detecting leaked private keys.

Verification Routine: Challenge/Response Test
Frequency: Post-launch then at regular intervals during the life of the payload
Requirements checked: Time, location, view, and origin
Description: Using a trusted radio frequency antenna, communicate with the satellite directly to have a message encrypted. The ground operator generates a nonce (a one time message, such as a random number) and transmits it to the satellite over commanding channels with a request to have it timestamped and encrypted by the Trusted Signing Hardware. The satellite receives the request and forwards it on to the Trusted Signing Hardware. The Trusted Signing Hardware collects the received nonce, the current time, the IMU reading, and the presence of any tamper events into a package and encrypts it with the private key. The encrypted content is then returned to the satellite processor and then downlinked to the ground operator.

When the response is received, it is decrypted using the Trusted Signing Hardware's public key. The ground operator ensures that the returned nonce matches the one that was sent, that the returned timestamp is correct, and that the IMU reading reports a space environment. The presence of any tamper events is noted. This test is important for checking the timestamp because the operator can bound

the expected round-trip time using the distance to the satellite and the expected time for processing the request. If a response is unduly delayed, then the test fails. This test ensures that the private key is located at a specific place in the sky and that the clock has not been tampered with. The time bounds must be tight enough so that the satellite would not have time to defer the request to a neighboring satellite or a ground station where the Trusted Signing Hardware is actually located. An untrusted ground station should not be used for this test, as it could redirect the request to a different antenna or to the Trusted Signing Hardware that was retained on the ground. This test is depicted in Figure 5.

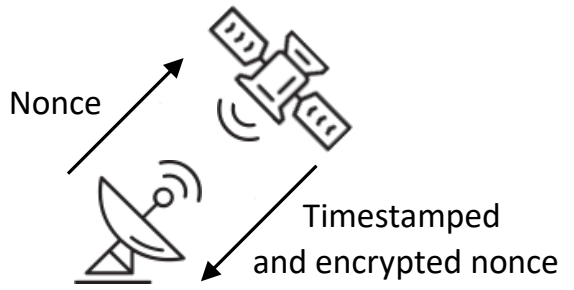

Figure 5: The messaging between a trusted ground station and the Trusted Signing Hardware for the challenge response test

---

Verification Routine: Emitter Test
Frequency: Post-launch then at regular intervals during the life of the payload
Requirements checked: Check location, orientation, view, origin, and content
Description: A ground-based optical signal is directed at the satellite hosting the Trusted Signing Hardware. This signal ultimately represent a nonce and that can be done by varying the signal's color, intensity, and/or location. Multiple optical emitters can also be used within the satellites field of view to increase the amount of information transmitted. The signals color, intensity, location, and emission time are logged. The diameter of the signal should be limited to reduce the chances of interception by unintended viewing satellites. For testing with satellites in non-geosynchronous, the signal emitter should be on a motorized mount that can track the satellite. This test must be performed at a time that the satellite will be imaging. This can be done by tasking the satellite to collect at a particular place and time or by calculating when and where the satellite will be collecting.

When imagery from the time period of the test becomes available, it will be checked using the standard steps for image product verification. The test operator then attempts to detect and extract the signal from the message. This can be done by the geospatial coordinates of the emitter at the time of the test or by searching for the signal in the image. The location, color, and intensity of the signal in the image are noted. The ground operator then looks in the emitter log for the signal attributes at the time the image was taken. If the signal matches, then the test passes. If the signal is missing or if the attributes do not match, then the test fails. Failures could be the result of unsynchronized clocks, poor geolocation, or poor pointing of the emitter. This test is depicted in Figure 6.

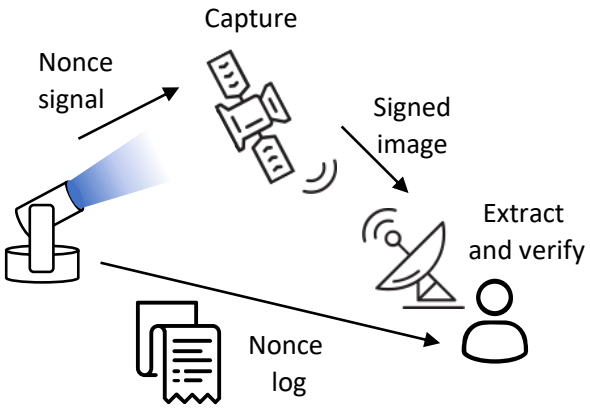

Figure 6: Component interactions in the emitter test

This test is important because it verifies that the imaging sensor had a real, unobstructed view of the ground. This test can be performed passively if the collection schedule is known or if tasking is submitted through an unattributed account. This test also confirms the location and orientation of the satellite by the sheer fact that the satellite was able to see the directed signal.

This test has similarity with Satellite Laser Ranging (SLR)[23] in that a signal is transmitted from the ground to the satellite. The difference is that instead of using a reflector on the surface of the satellite to return the signal to Earth, the return signal returns in the data stream. Since the return signal is digitally signed in the data stream, we sometimes call this test "Verifiable Satellite Ranging" or VSR. The signal does not have to be a light, but could be an object visible from space. Historically, objects like these used for communication are called optical telegraphs.[24]

The risk of the integrity box being compromised or

Table 1: Summary of how the trust requirements are covered by the four classes of protections

Trust Requirement	Potential Threats	Architecture Properties	Anti-tamper Measures	Inspection Processes	Verification Routines
Time	Replay attack, time falsification	Independent time source	Anti-tamper packaging	Access restriction, package check	Challenge/ response test, image product verification
Location	Not deployed/operated elsewhere	Independent IMU - space environment	Anti-tamper packaging	Access restriction, package check	Challenge/ response test, emitter test
Orientation	Looking in wrong direction	Independent IMU - attitude	Anti-tamper packaging	Access restriction, package check	Emitter test, image product verification
View	Fake inputs, Blocked view	-	Restricted interfaces and connections	Check connection to FPA	Emitter test
Origin	Private key leak, falsification of origin	Private signing key	Protect private key with secure boot	Access restriction, package check	Challenge/ response test, emitter test, image product verification
Content	Private key leak, falsification of content	Private signing key, chain of hashes	Protect private key with secure boot	Access restriction, package check	Challenge/ response test, emitter test, image product verification

tampered with drops significantly once it is fielded in a space environment. The majority of the protections we have architected are focused on confirming that the device was fielded to the correct orbit, and that it was not modified or compromised before launch. This necessitates aggressive anti-tamper protections as well as inspections to verify no tamper-events occurred.

As with any authenticated camera, ours must be checked for proper function. These tests must take place once the satellite is delivered on orbit and can be performed on a routine basis after that.

Limitations

While our approach is the most advanced we know of, it does have some potential weaknesses. First, physical access to a launched satellite or the ability to remotely reconfigure the connection between the FPA and the Trusted Signing Hardware, can violate trust requirement 4 (view). Second, digitally signed imagery cannot be modified for any reason, including standard image reconstruction and enhancement pipelines processing, without invalidating the previously computed digital signature. Solutions include running image processing routines inside the Trusted Signing Hardware or deferring them to the end user. Finally, if the host satellite doesn't communicate challenge/response test requests and responses in a timely manner, we will be unable to accurately validate the clock. In this case, checking that timestamps are monotonically increasing in image product verification will buy down risk of clock tampering.

Solution Summary

Our approach addresses the six trust requirements for imagery by utilizing trusted hardware to capture the image and the collection time, location, and orientation. The hardware is protected by anti-tamper mechanisms and pre-launch controls and inspections. Post launch, the hardware is routinely checked at its expected location for proper function. These safety measures enable end users to verify the image products they receive. Table 1 provides a summary of these protections.

PATH FORWARD

Our design is complete and testing of our verification routines is underway. We are preparing to deploy a Trusted Signing Hardware prototype to the International Space Station (ISS) early next year. That flight will enable us to check our remote verification routines and location sensors together.

We have started investigating approaches for verifying the outputs of image processing algorithms executed at untrusted ground stations.[25] Our next steps include looking into secure multi-party computation to enable processing imagery without impacting the validity of a digital signature.

The techniques in this paper may also apply to other remote sensing modalities, including Radio Frequency (RF) and Synthetic Aperture Radar (SAR). We reserve those as future work.

CONCLUSION

Commercial remote sensing is changing the way we see the world. The availability and coverage of imagery is increasing while the costs are decreasing. The architecture and methods outlined in this paper provides a new framework for ensuring trust in satellite imagery, addressing an increasing accessible range of attack vectors.

Our system can create trust in:

1. The time an image was captured through an independent time source and remote verification of that clock.
2. The location of the imaging sensor through direct communication with the satellite.
3. The orientation of the imaging sensor through an independent IMU reading.
4. The imaging sensor having an unobstructed view of the scene through routine spot checking.
5. The origin of the image through proper key distribution, key protection, and image signing and verification.
6. The authenticity of the image content by signing and verifying the content.

This approach of providing independent verification of satellite collection activities offers trust to end users. We are not aware of any other solution for trusting third-party satellite data in a cryptographically secure fashion.

ACKNOWLEDGEMENTS

This work was supported by the U.S. Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2022-7590 C

Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. DOE or the United States Government.

References

[1] Abraham Noah Wu and Filip Biljecki. Ganmapper: geographical content filling. *arXiv preprint arXiv:2108.04232*, 2021.

[2] Bo Zhao, Shaozeng Zhang, Chunxue Xu, Yifan Sun, and Chengbin Deng. Deep fake geography? when geospatial data encounter artificial intelligence. *Cartography and Geographic Information Science*, 48(4):338–352, 2021.

[3] Christopher X Ren, Amanda Ziemann, James Theiler, and Juston Moore. Deepfaking it: experiments in generative, adversarial multispectral remote sensing. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11727, page 117270M, 2021.

[4] Paweł Korus. Digital image integrity–a survey of protection and verification techniques. *Digital Signal Processing*, 71:1–26, 2017.

[5] ShiYue Lai and Rainer Böhme. Countering counter-forensics: The case of jpeg compression. In *International Workshop on Information Hiding*, pages 285–298. Springer, 2011.

[6] John Kelsey, Bruce Schneier, and Chris Hall. An authenticated camera. In *Proceedings 12th Annual Computer Security Applications Conference*, pages 24–30. IEEE, 1996.

[7] David Skogmo. Trace-tamper resistant authenticated camera enclosure. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 1989.

[8] Ihtesham Haider and Bernhard Rinner. Private space monitoring with soc-based smart cameras. In *2017 IEEE 14th International Con-*

ference on Mobile Ad Hoc and Sensor Systems (MASS), pages 19–27. IEEE, 2017.

[9] The coalition for content provenance and authenticity (c2pa). `https://c2pa.org`. Accessed: 2022-02-11.

[10] Three challenges of our digital age. `https://www.starlinglab.org/challenges/`. Accessed: 2022-02-16.

[11] Project origin - protecting trusted media. `https://www.originproject.info/`. Accessed: 2022-02-16.

[12] Content authenticity initiative. `https://contentauthenticity.org`. Accessed: 2022-02-16.

[13] C2pa explainer. `https://c2pa.org/specifications/specifications/1.0/explainer/Explainer.html`. Accessed: 2022-02-14.

[14] Sean Kinser, Pete de Graaf, Matthew Stein, Frank Hughey, Rob Roller, David Voss, and Amanda Salmoiraghi. Scoring trust across hybrid-space: A quantitative framework designed to calculate cybersecurity ratings, measures, and metrics to inform a trust score. Technical report, THE MITRE CORPORATION HANSCOM AFB United States, 2020.

[15] Nathan E Price and Alan T Sherman. How to generate repeatable keys using physical unclonable functions. *Department of Computer Science and Electrical Engineering*, pages 1–9.

[16] Salil S Kanhere. Participatory sensing: Crowdsourcing data from mobile smartphones in urban spaces. In *International Conference on Distributed Computing and Internet Technology*, pages 19–26. Springer, 2013.

[17] Heejin Park, Shuang Zhai, Long Lu, and Felix Xiaozhu Lin. Streambox-tz: secure stream analytics at the edge with trustzone. In *2019 {USENIX} Annual Technical Conference ({USENIX}{ATC} 19)*, pages 537–554, 2019.

[18] Kwstantinos Papadamou, Riginos Samaras, and Michael Sirivianos. Ensuring the authenticity and fidelity of captured photos using trusted execution and mobile application licensing capabilities. In *2016 11th International Conference on Availability, Reliability and Security (ARES)*, pages 706–714. IEEE, 2016.

[19] Sandro Pinto and Nuno Santos. Demystifying arm trustzone: A comprehensive survey. *ACM Computing Surveys (CSUR)*, 51(6):1–36, 2019.

[20] Commercial national security algorithm suite (cnsa). `https://apps.nsa.gov/iaarchive/programs/iad-initiatives/cnsa-suite.cfm`. Accessed: 2022-05-30.

[21] Swapnil Sayan Saha, Shafizur Rahman, Mosabber Uddin Ahmed, and Subrata Kumar Aditya. Ensuring cybersecure telemetry and telecommand in small satellites: Recent trends and empirical propositions. *IEEE Aerospace and Electronic Systems Magazine*, 34(8):34–49, 2019.

[22] Daria Lane, Enrique Leon, Dexter Solio, Daniel Cunningham, Dmitrie Obukhov, and Francisco C Tacliad. High-assurance cyber space systems for small satellite mission integrity. 2017.

[23] John J Degnan and Erricos C Pavlis. Laser ranging to gps satellites with centimeter accuracy. 1994.

[24] Amanda Dominguez. Optical telegraph (semaphore system). Accessed on 25 August 2021.

[25] Jack Toomey and Sean Crosby. A step toward working with untrusted ground stations. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2022.