SecDATAVIEW: A Secure Big Data Workflow Management System for Heterogeneous Computing Environments

Saeid Mofrad, Ishtiaq Ahmed, Shiyong Lu, Ping Yang,

Heming Cui, **Fengwei Zhang***

{saeid.mofrad, ishtiaq, shiyong, fengwei}@wayne.edu

pyang@binghamton.edu

heming@cs.hku.hk

*The corresponding author, and he is currently affiliated with SUSTech.

Outline

- Introduction
- x86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

Outline

- Introduction
- x86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

Cloud Platform for Big Data Analytics

- Cloud platforms are common for big data analytics
- Isolation through software virtualization is used to achieve trusted execution environment (TEE) in cloud infrastructure
- **Downsides of virtualization [3]:**

1) Virtualization uses shared hardware, hypervisor and cloud system software thus increases the software and hardware TCB of the cloud platform
2) Hypervisor and cloud's system software contain thousands of lines of code and may have security flaws
3) Many hypervisor exploits have been reported in clouds [5,6]
4) Increased TCB size means less security

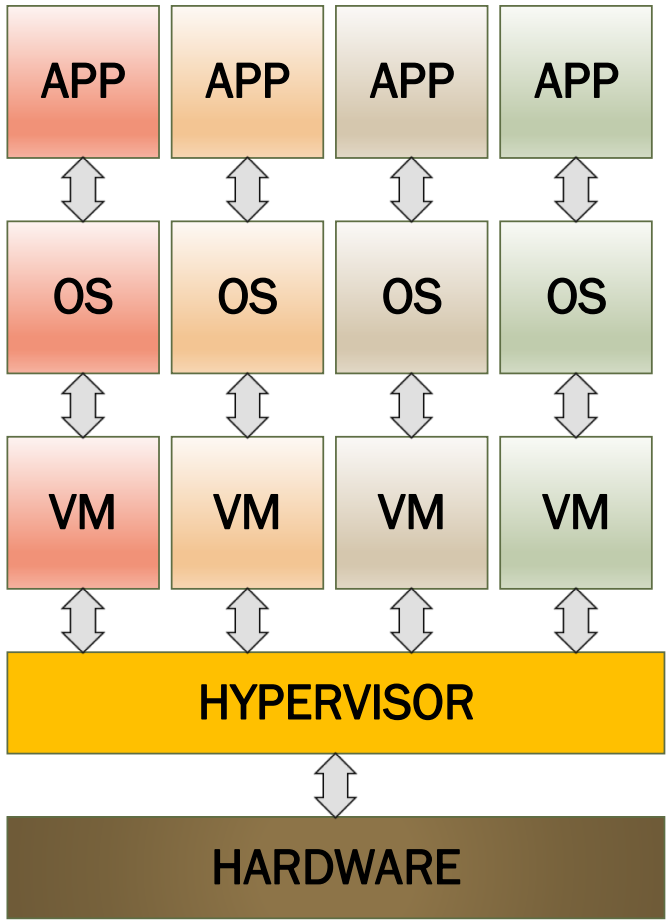

Outline

- Introduction
- x86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

Hardware-Assisted Trusted Execution Environment in x86 Architecture [3]

- Hardware-Assisted TEE couples hardware with TEE abstraction so mitigates the downsides of the software only TEEs
- Hardware-Assisted TEE may be faster since it uses dedicated hardware
- Hardware-Assisted TEE exposes small size of hardware TCB and smaller TCB means better security
- "Older" Hardware-Assisted TEE: Intel ME, AMD PSP, and x86 SMM [4]

- **Two general-purpose Hardware-Assisted TEE in x86 architecture:**

1. Intel Software Guard eXtensions (SGX) [HASP 2013], [1]
2. AMD Memory Encryption Technology [White Paper 2016], [2]

Background: Intel SGX and AMD SEV [3]

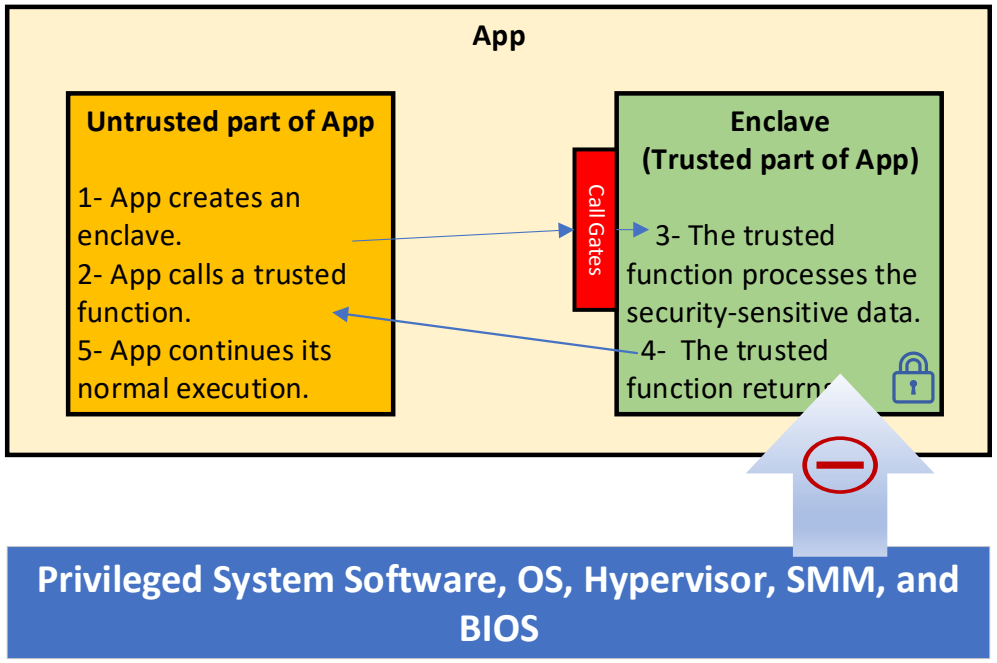

Intel SGX

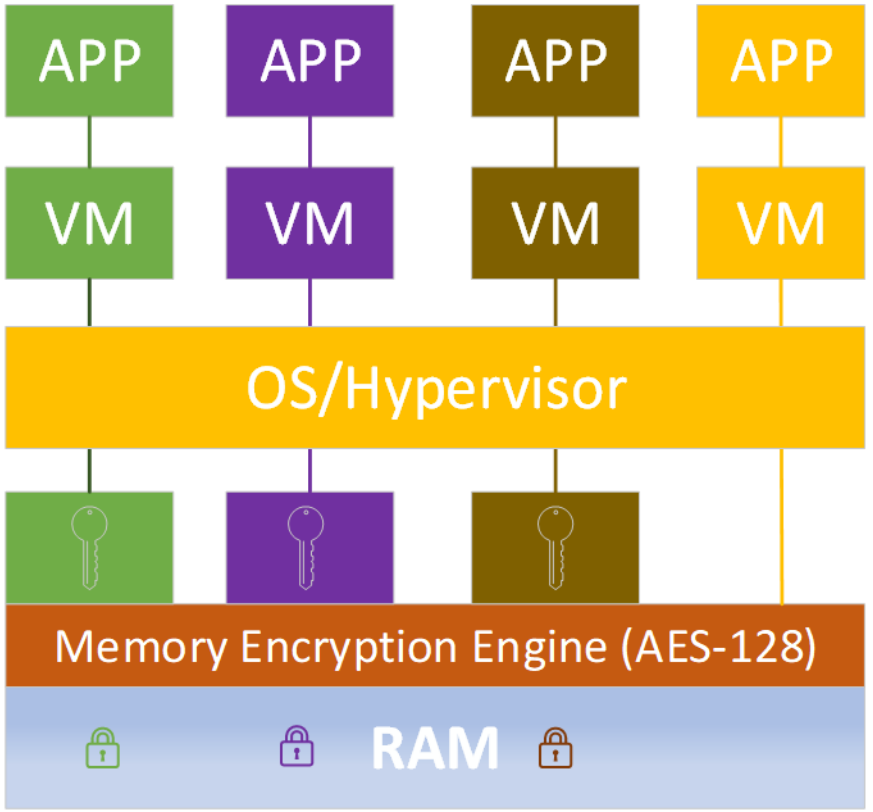

AMD SEV

Intel SGX VS AMD SEV [3]

TEE Technology	Runtime Access Privilege	Memory Size Limits	SDK	Software Change	Platform Attestation Mechanism	TEE Protection guarantee	TEE TCB SIZE	TEE performance
Intel SGX	Ring 3	Up to 128MB	Provided	Required	Attested through Intel remote attestation	Confidentiality and Integrity protection of the enclave's code and data at runtime	Smaller than SEV	Performs slower than SEV
AMD SEV	Ring 0	Up to available system memory	Not Required	Not required	Attested through AMD guest attestation	Confidentiality protection of the VM's memory image at runtime	Larger than SGX	Performs faster than SGX

Outline

- Introduction
- X86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

Previous Data Analytics Systems with TEE Support

- VC3: A trustworthy Hadoop based data analytics platform in the cloud that leverages SGX to protect unmodified Map-Reduce tasks written in C/C++ [S&P 2015], [7]
- A lightweight, Map-Reduce framework with Lua , a high-level language that interprets the Map-Reduce Lua scripts in Intel SGX [CCGRID 2017], [8]
- Opaque: An oblivious and encrypted distributed analytics platform that enhanced the security of the Spark SQL with SGX [NSDI 2017], [9]

Shortcoming with previous data analytics platforms with TEE support:

1. **Limited functionality:** They only support Map/Reduce or SQL query data types
2. **Lack of support for heterogeneous cloud infrastructure:** They only support Intel SGX platform

Outline

- Introduction
- X86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

SecDATAVIEW: A Secure Data Analytics System with Heterogeneous TEE Support

SecDATAVIEW main characteristics:

❑**Different data types:**

1. Supports scientific big data workflow [10] and considers each task as a black box
2. Supports many type of workflows (Map-Reduce, Query, Machine learning, Deep learning, Image-Video processing, etc..)

❑ **Heterogeneous TEE:**

1. Supports both Intel SGX and AMD SEV at the same time

❑**Strong security guarantee:**

1. Protects the confidentiality and integrity of **code** and **data** for workflows running on public untrusted clouds
2. Supports High-level and managed code programming language (Java) that protects memory leaks vulnerability (buffer overflow)
3. Provides minimal hardware and software TCB for general purpose cloud based big data analytics platform

❑**Flexible system settings (SGX mode, SEV mode, Hybrid mode) for enhanced security and performance requirements:**

1. Supports trade-off between enhanced security (SGX mode) and performance (SEV mode) for workflows with different user requirements.

SecDATAVIEW: Leverages Heterogenous Workers with TEE Support

❑ **SecDATAVIEW Intel SGX Worker:**

1) Uses SGX shield [19] programming model
2) SGX-LKL [20] is incorporated to provide the Java virtual machine in the SGX enclave
3) Encrypted SGX-LKL disk image is used to protect the confidentiality of user code and data at rest
4) Java reflection and class loader are incorporated to overcome lack of multi-process support in the SGX-LKL

❑ **SecDATAVIEW AMD Worker:**

1) Uses AMD Secure Encrypted Virtualization (SEV)
2) SEV-protected VM is used to protect the worker memory image at runtime
3) Java virtual machine is used in every SEV-protected VM

SecDATAVIEW: Adversary Model

- **SecDATAVIEW threat model targeted attacks that happen on untrusted cloud:**

1) Attacks that exploit flaws or vulnerabilities in the hypervisor, or cloud's system software layer trying to gain access to the user data or results stored on unprotected memory
2) Attacks that could happen by dishonest administrator to gain access to data or results stored on the user storage medium

- Attacks, including network traffic-analysis [11], denial-of-service, access pattern leakage [12], side-channels [13], and fault injections [14], are out of the scope

SecDATAVIEW System Architecture

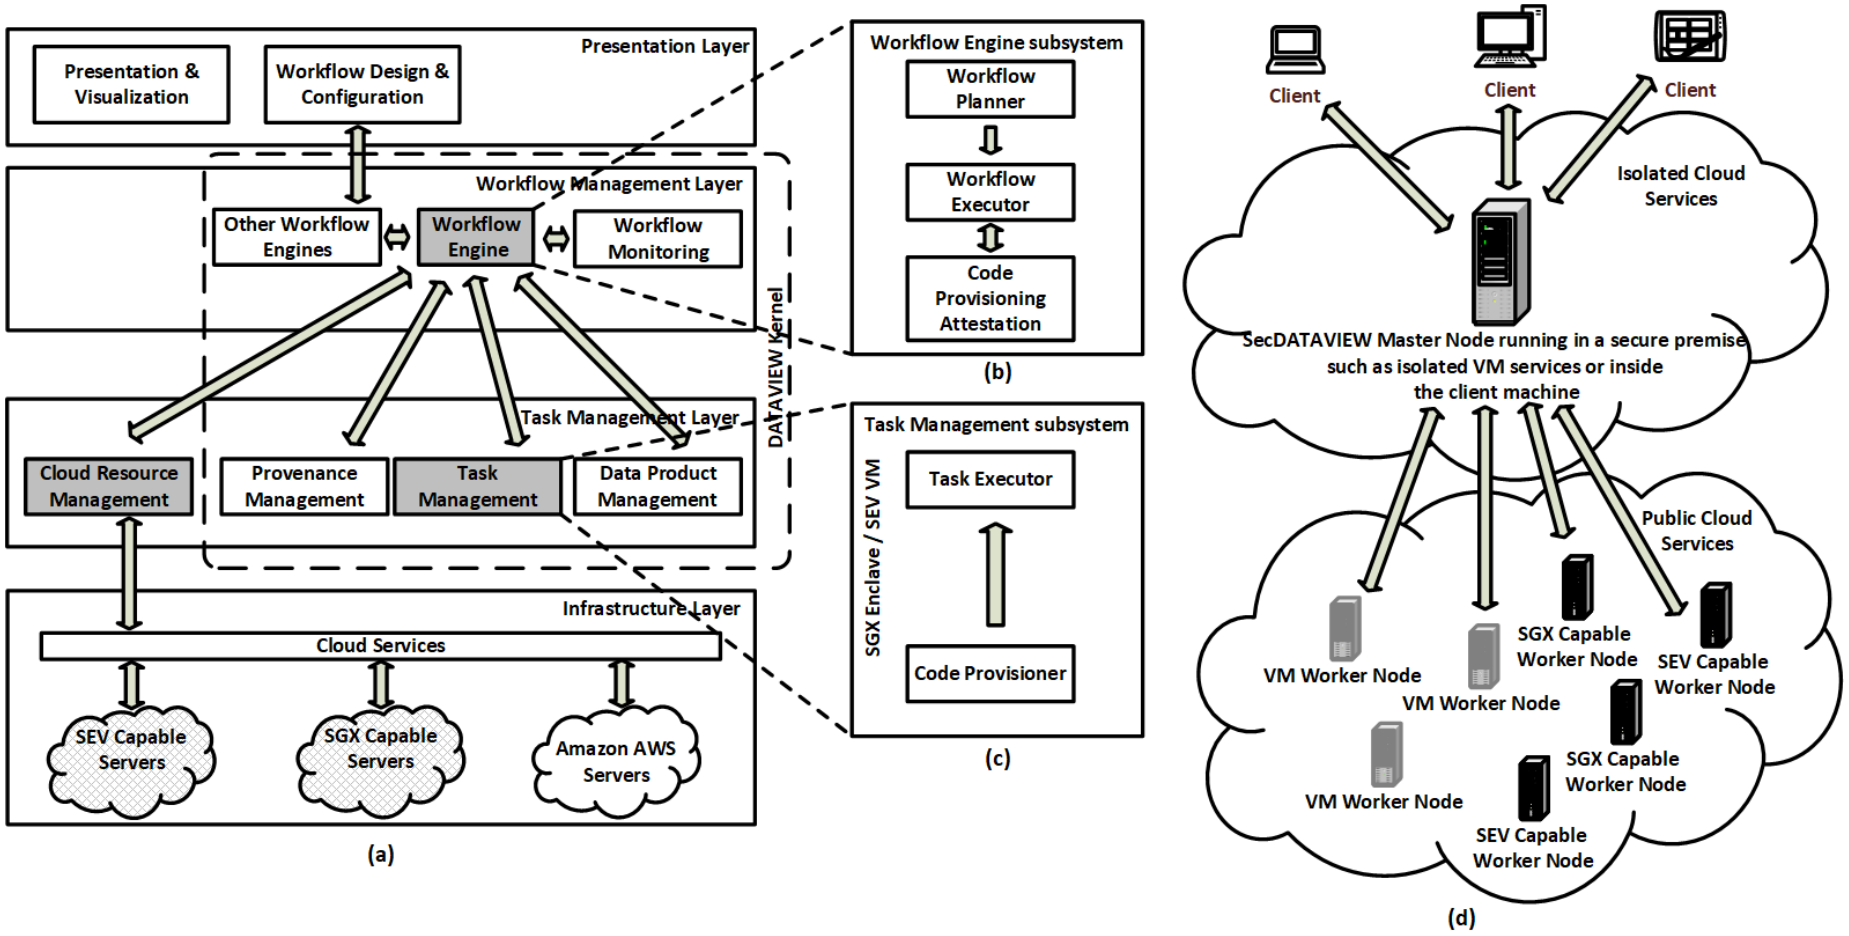

WCPAC: Workflow Code Provisioning and Communication Protocol

- **The WCPAC protocol's main functionality includes:**

1. to provision and attest secure worker nodes
2. to provision securely the code for the Task Executor and workflow tasks on each participating worker node
3. to establish the secure communication and file transfers between the master node and worker nodes
4. to ensure secure file transfers among worker nodes

WCPAC: Workflow Code Provisioning and Communication Protocol

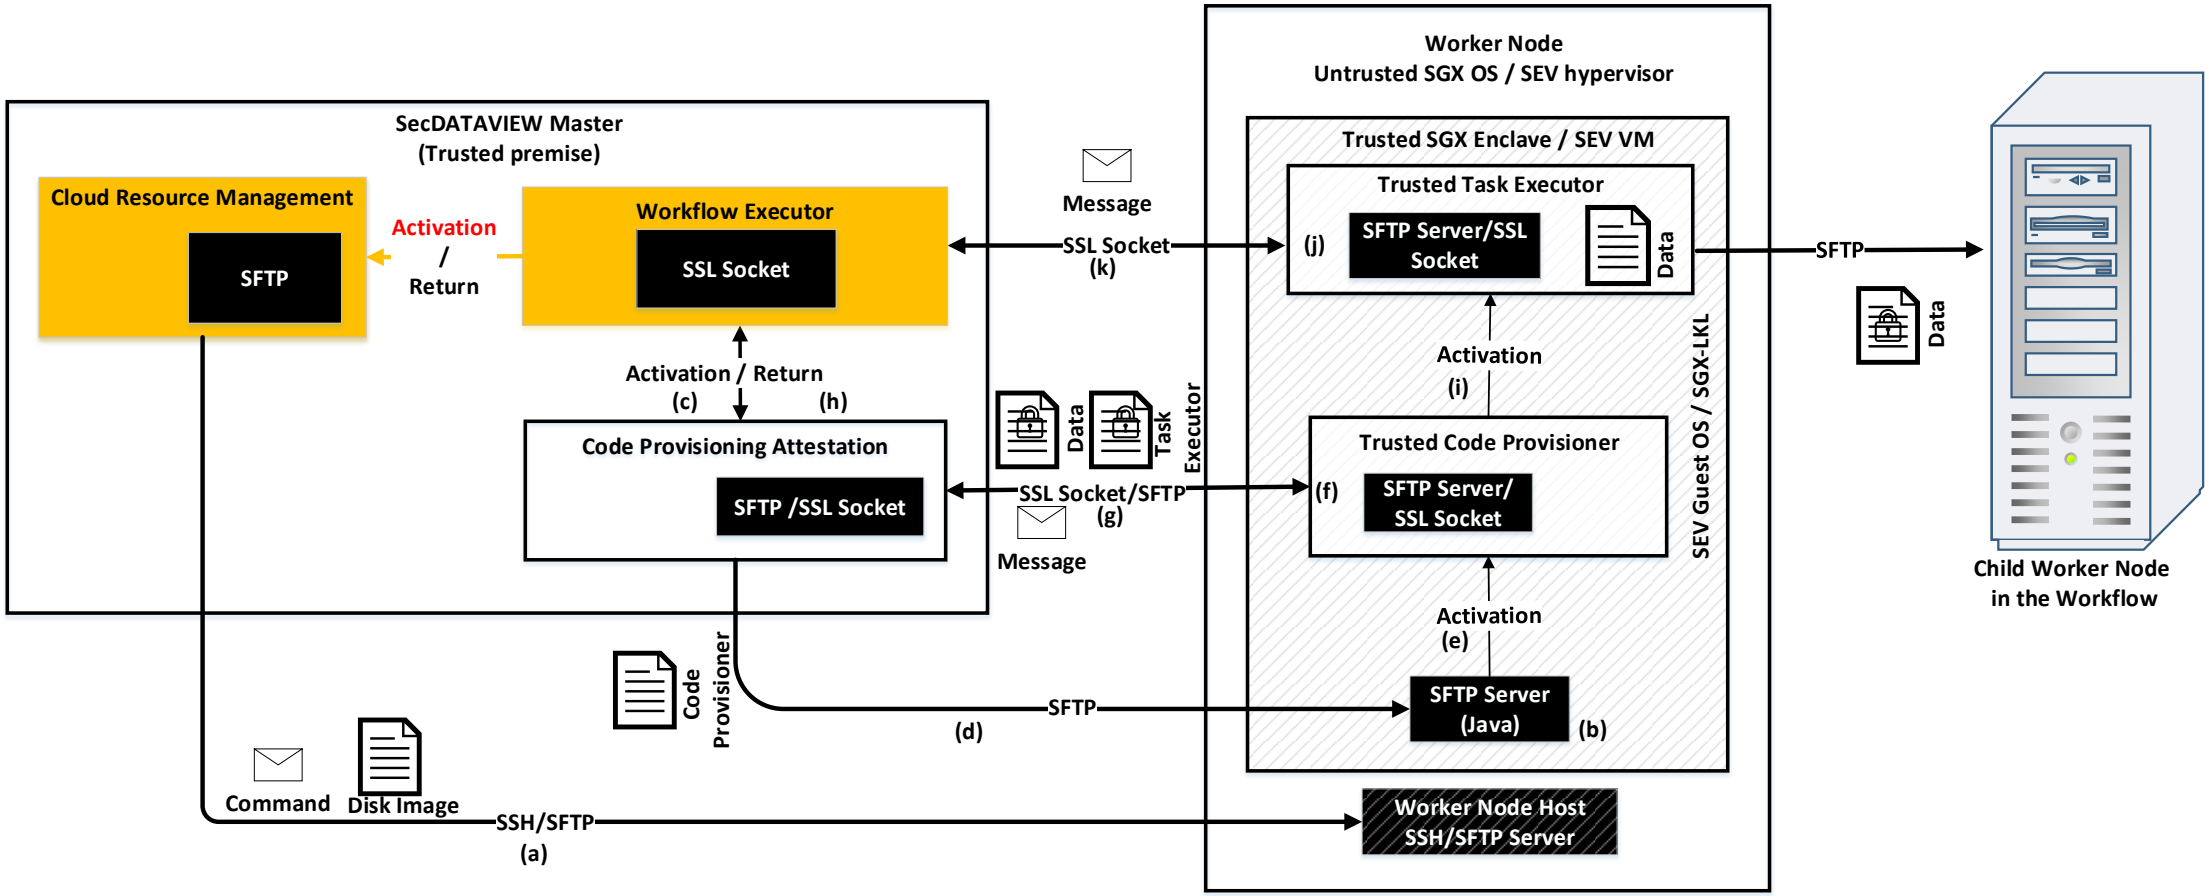

WCPAC: Workflow Code Provisioning and Communication Protocol

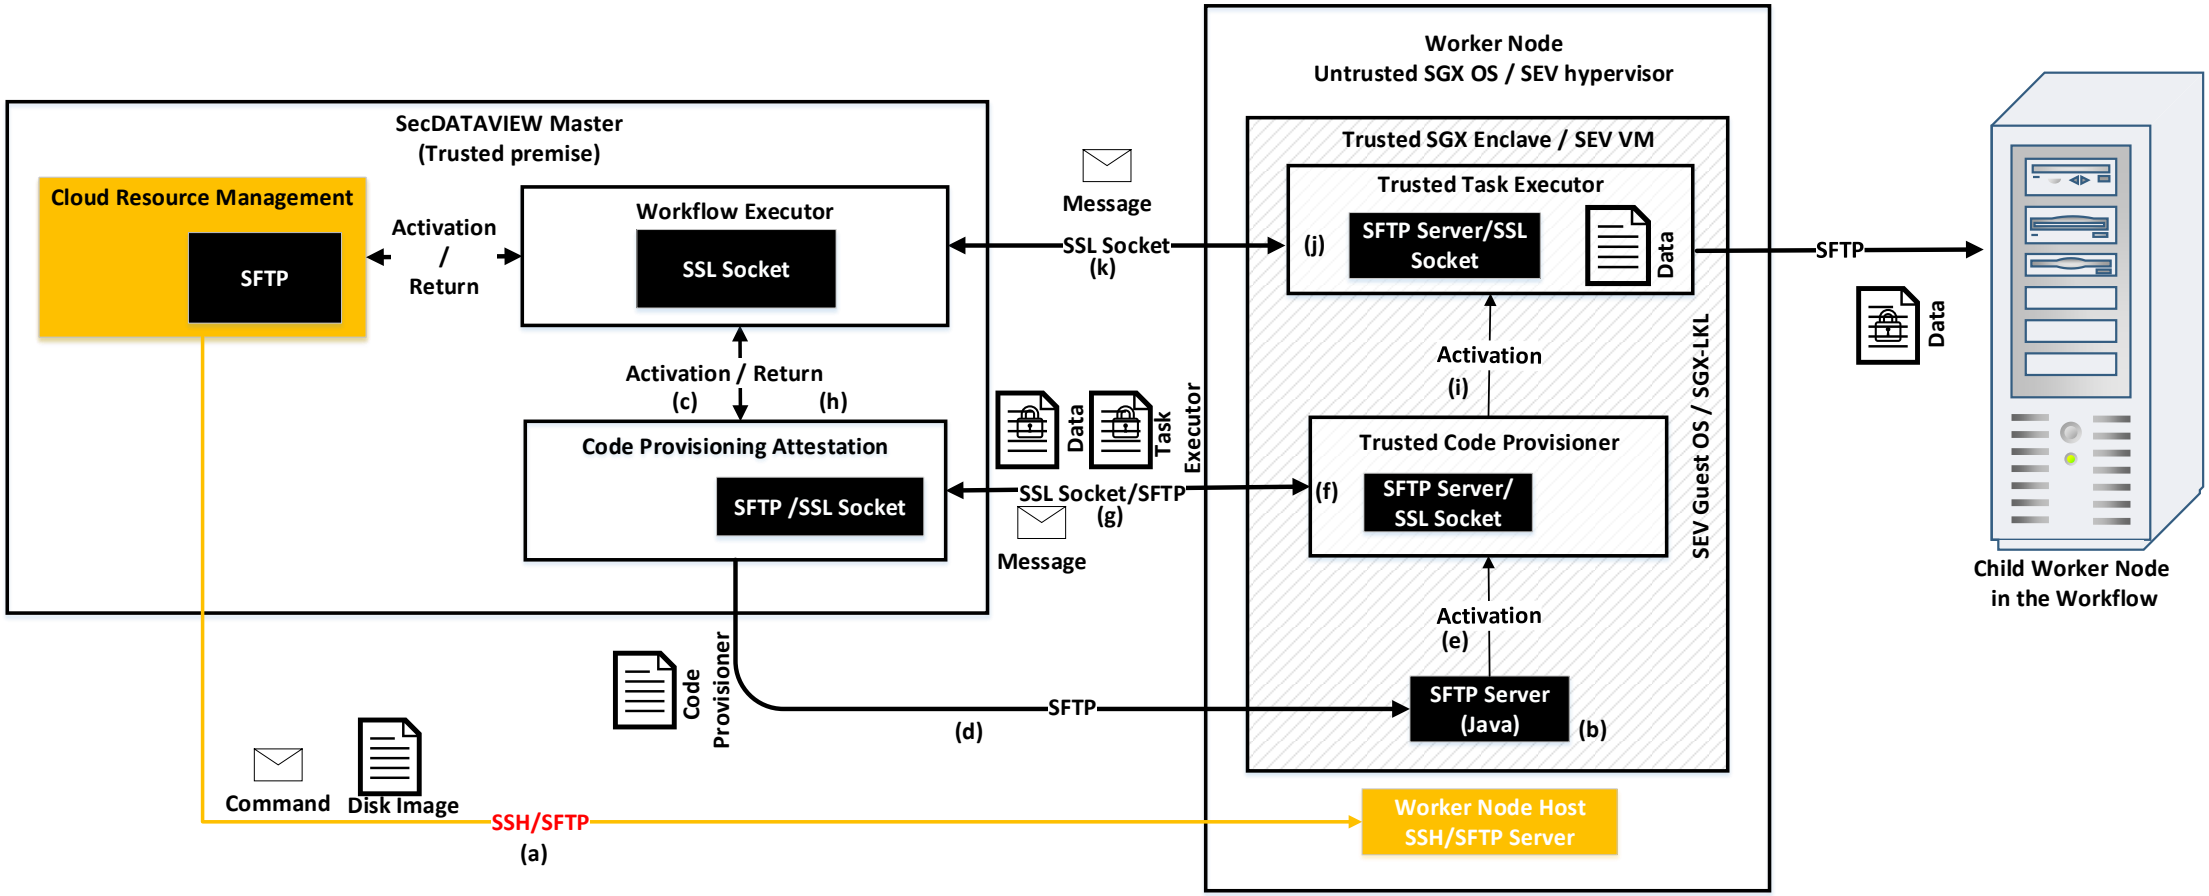

WCPAC: Workflow Code Provisioning and Communication Protocol

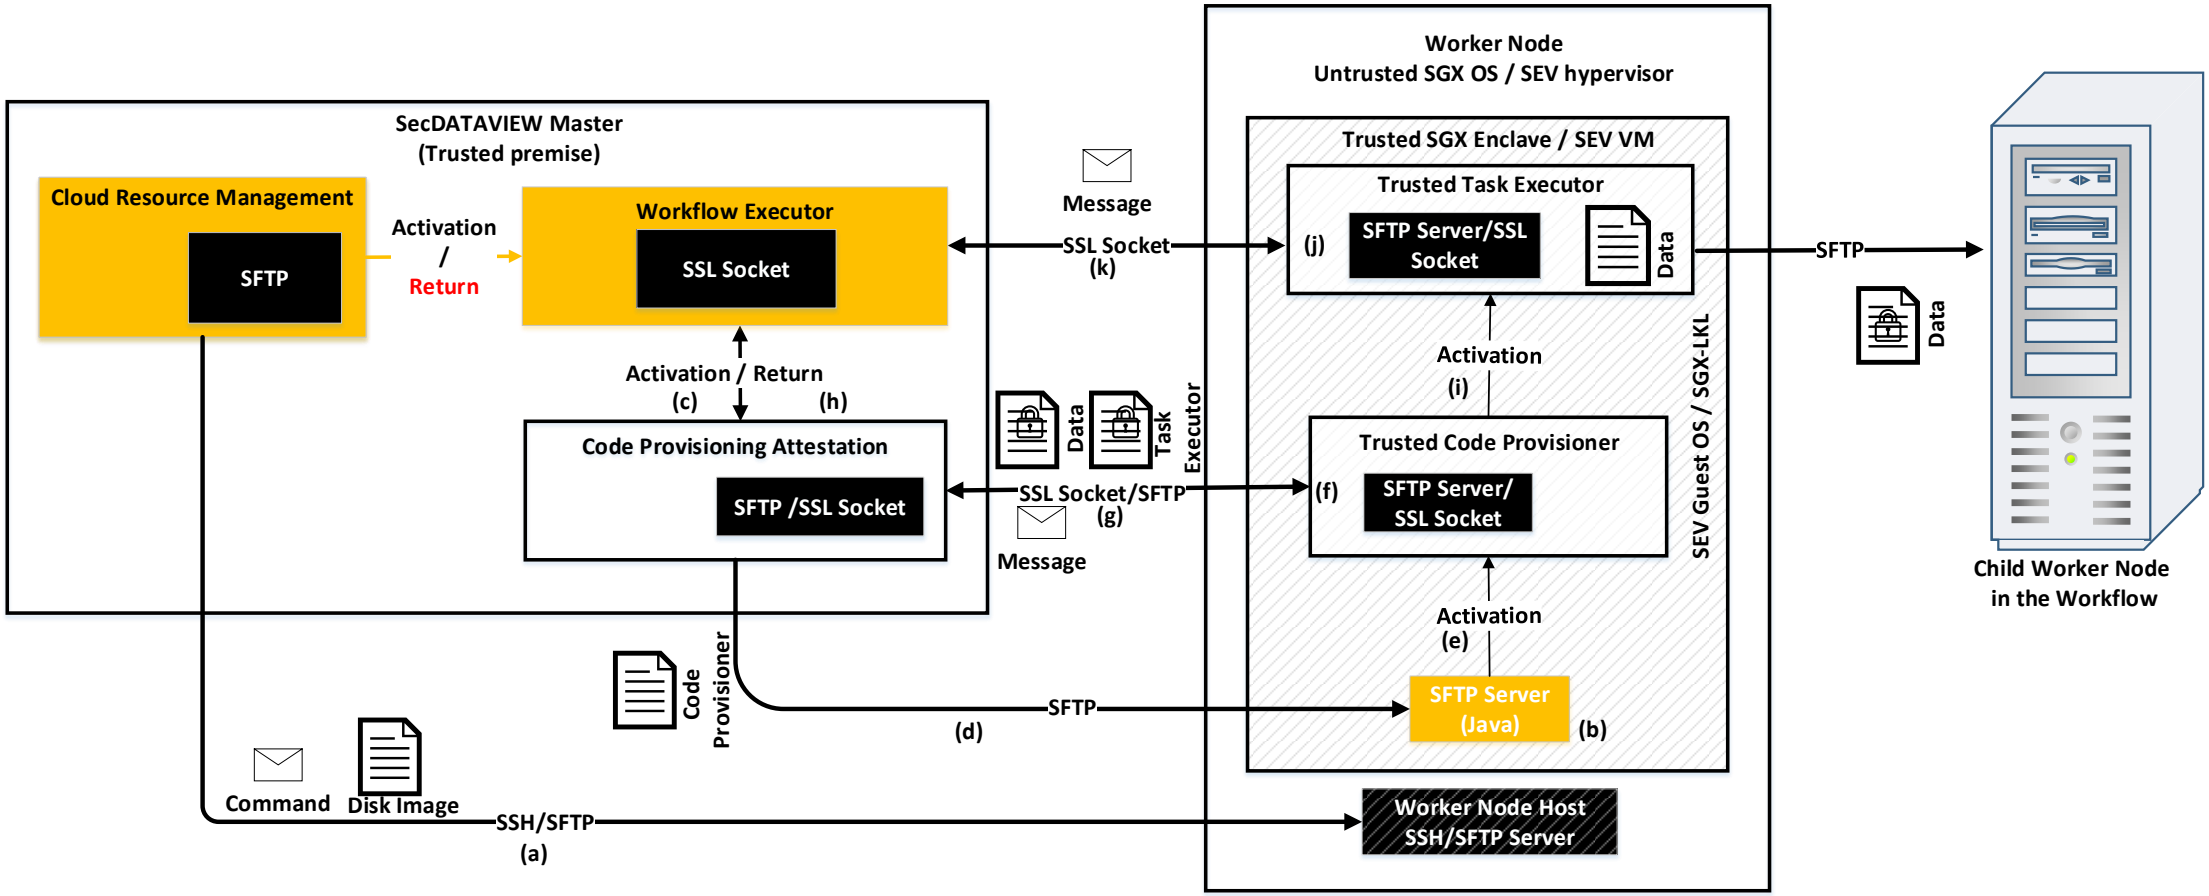

WCPAC: Workflow Code Provisioning and Communication Protocol

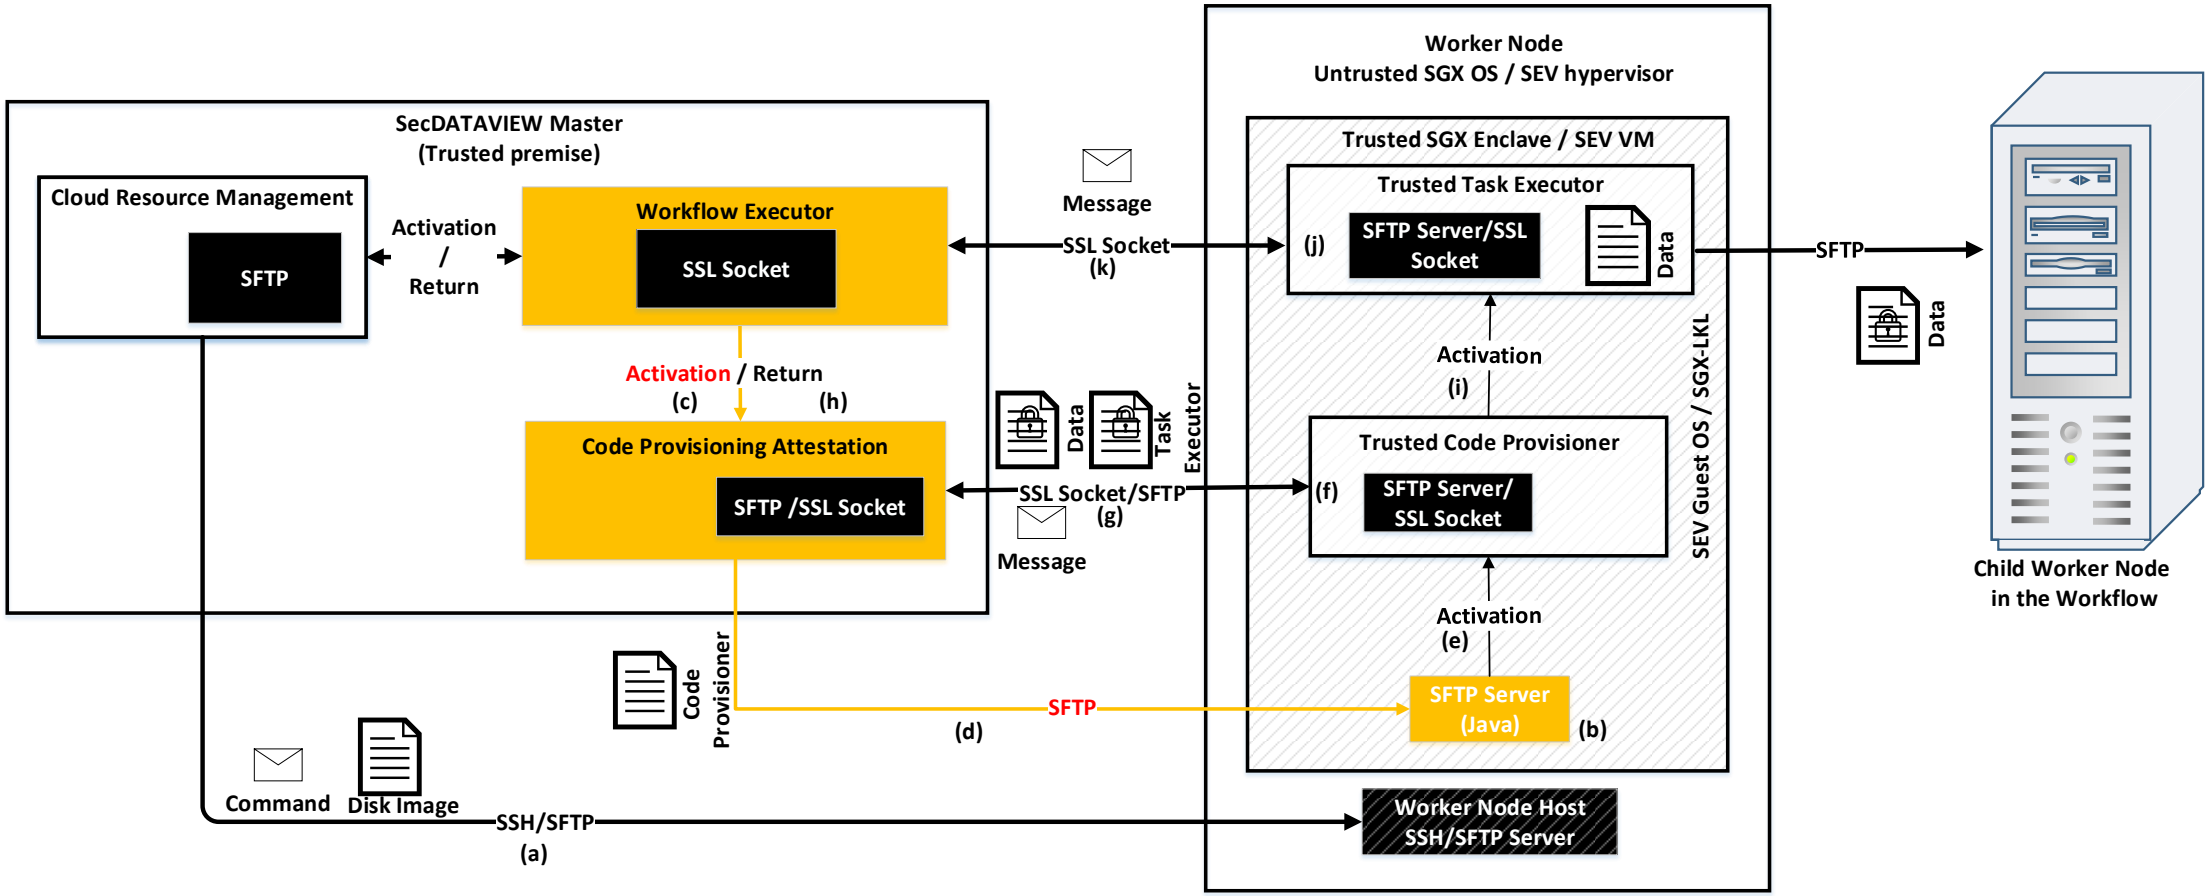

WCPAC: Workflow Code Provisioning and Communication Protocol

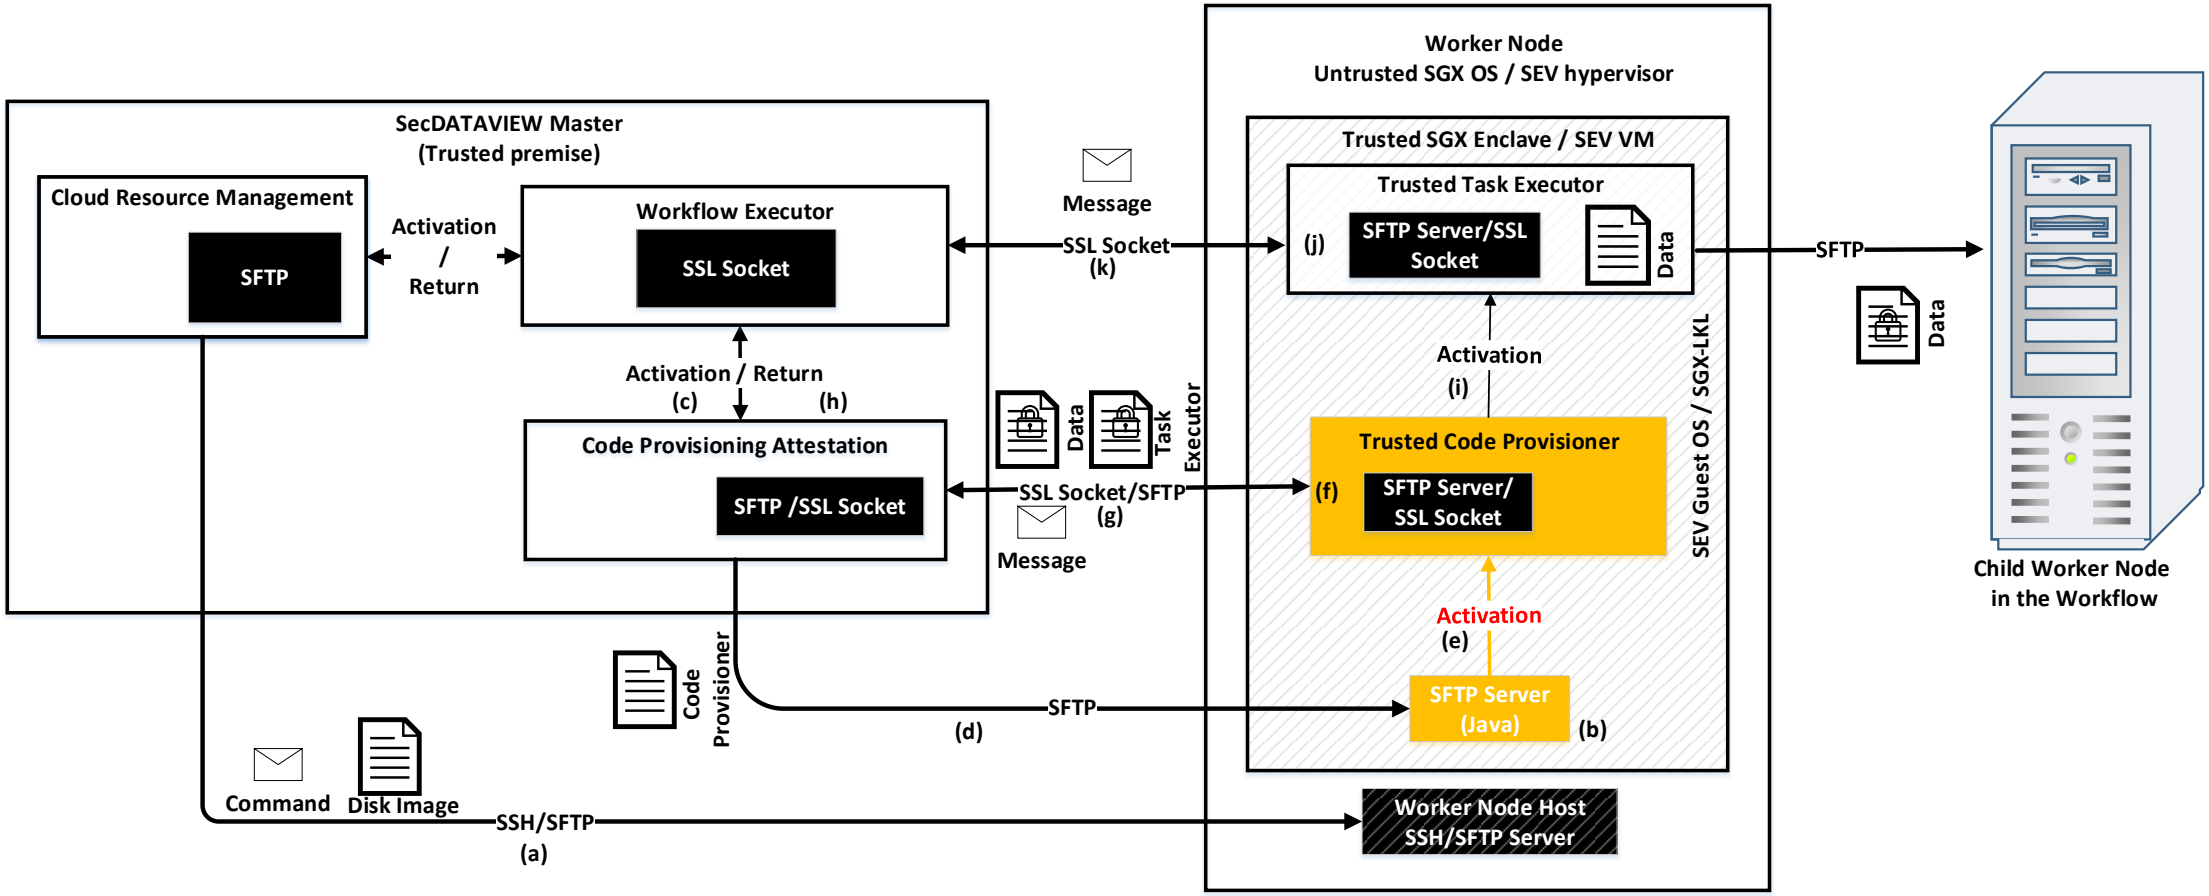

WCPAC: Workflow Code Provisioning and Communication Protocol

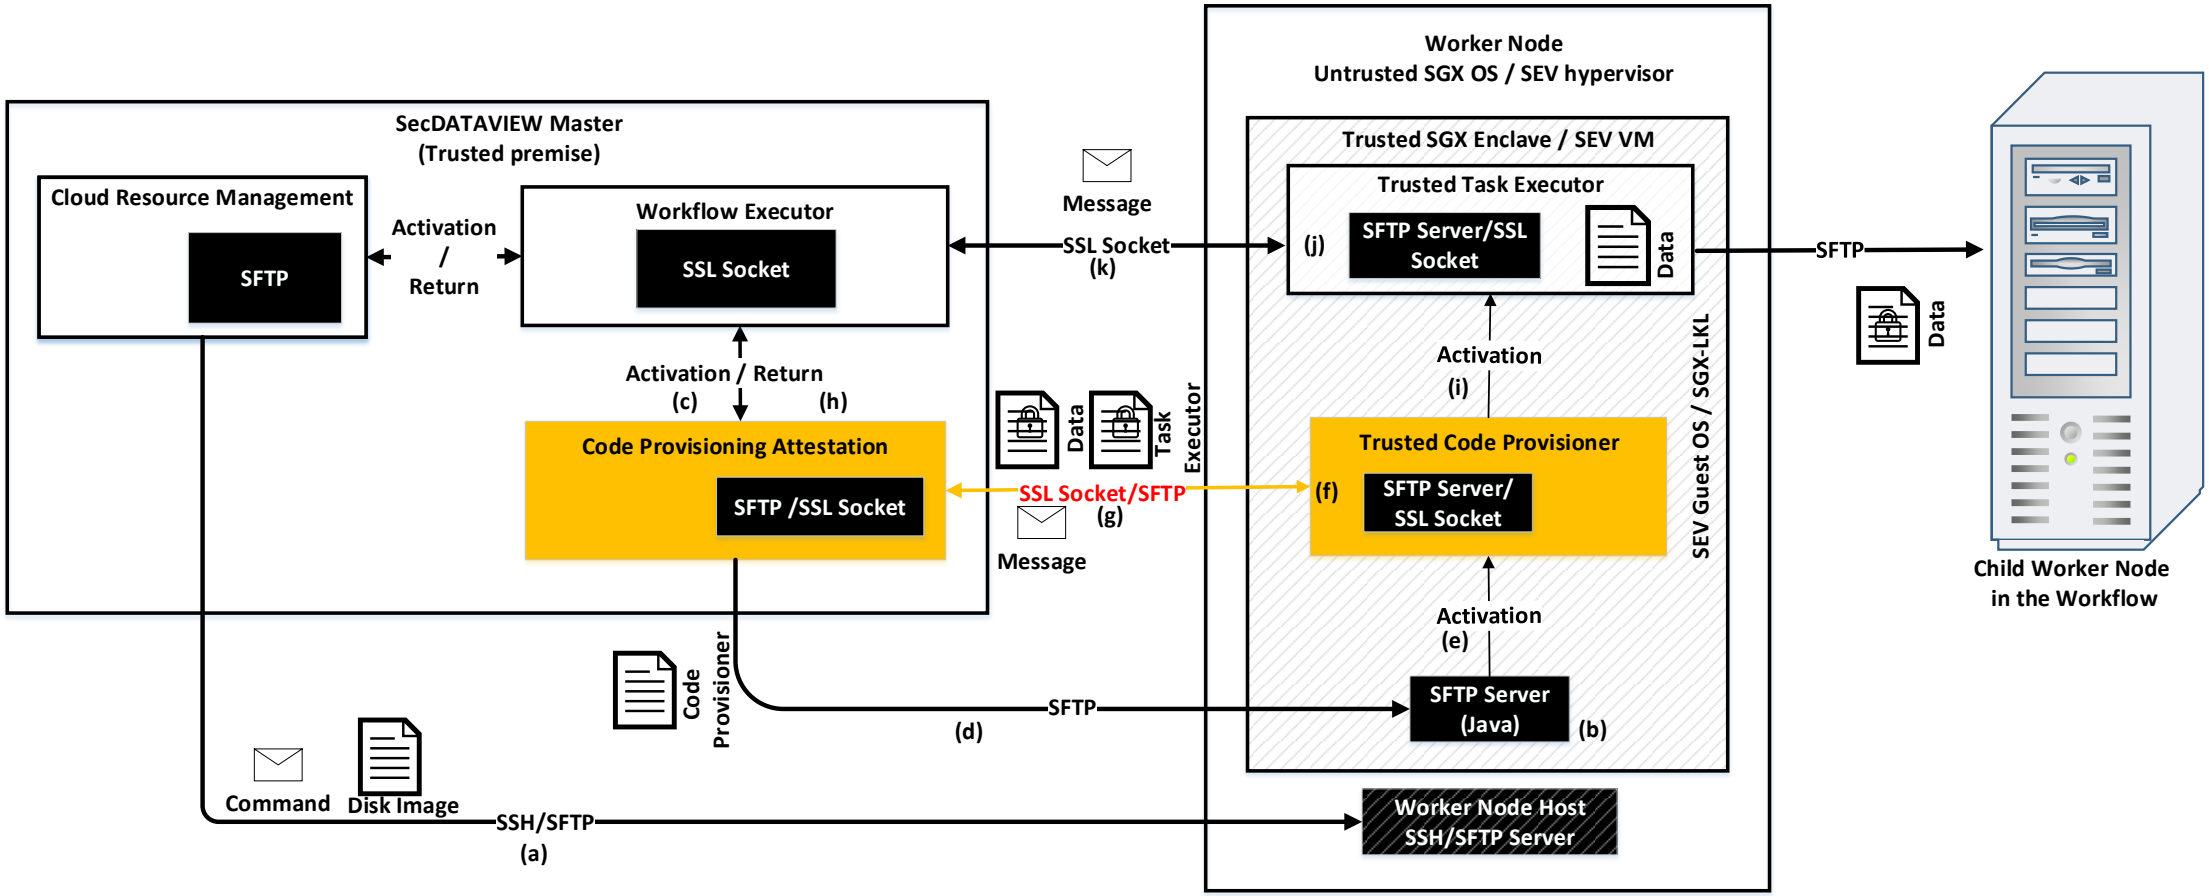

WCPAC: Workflow Code Provisioning and Communication Protocol

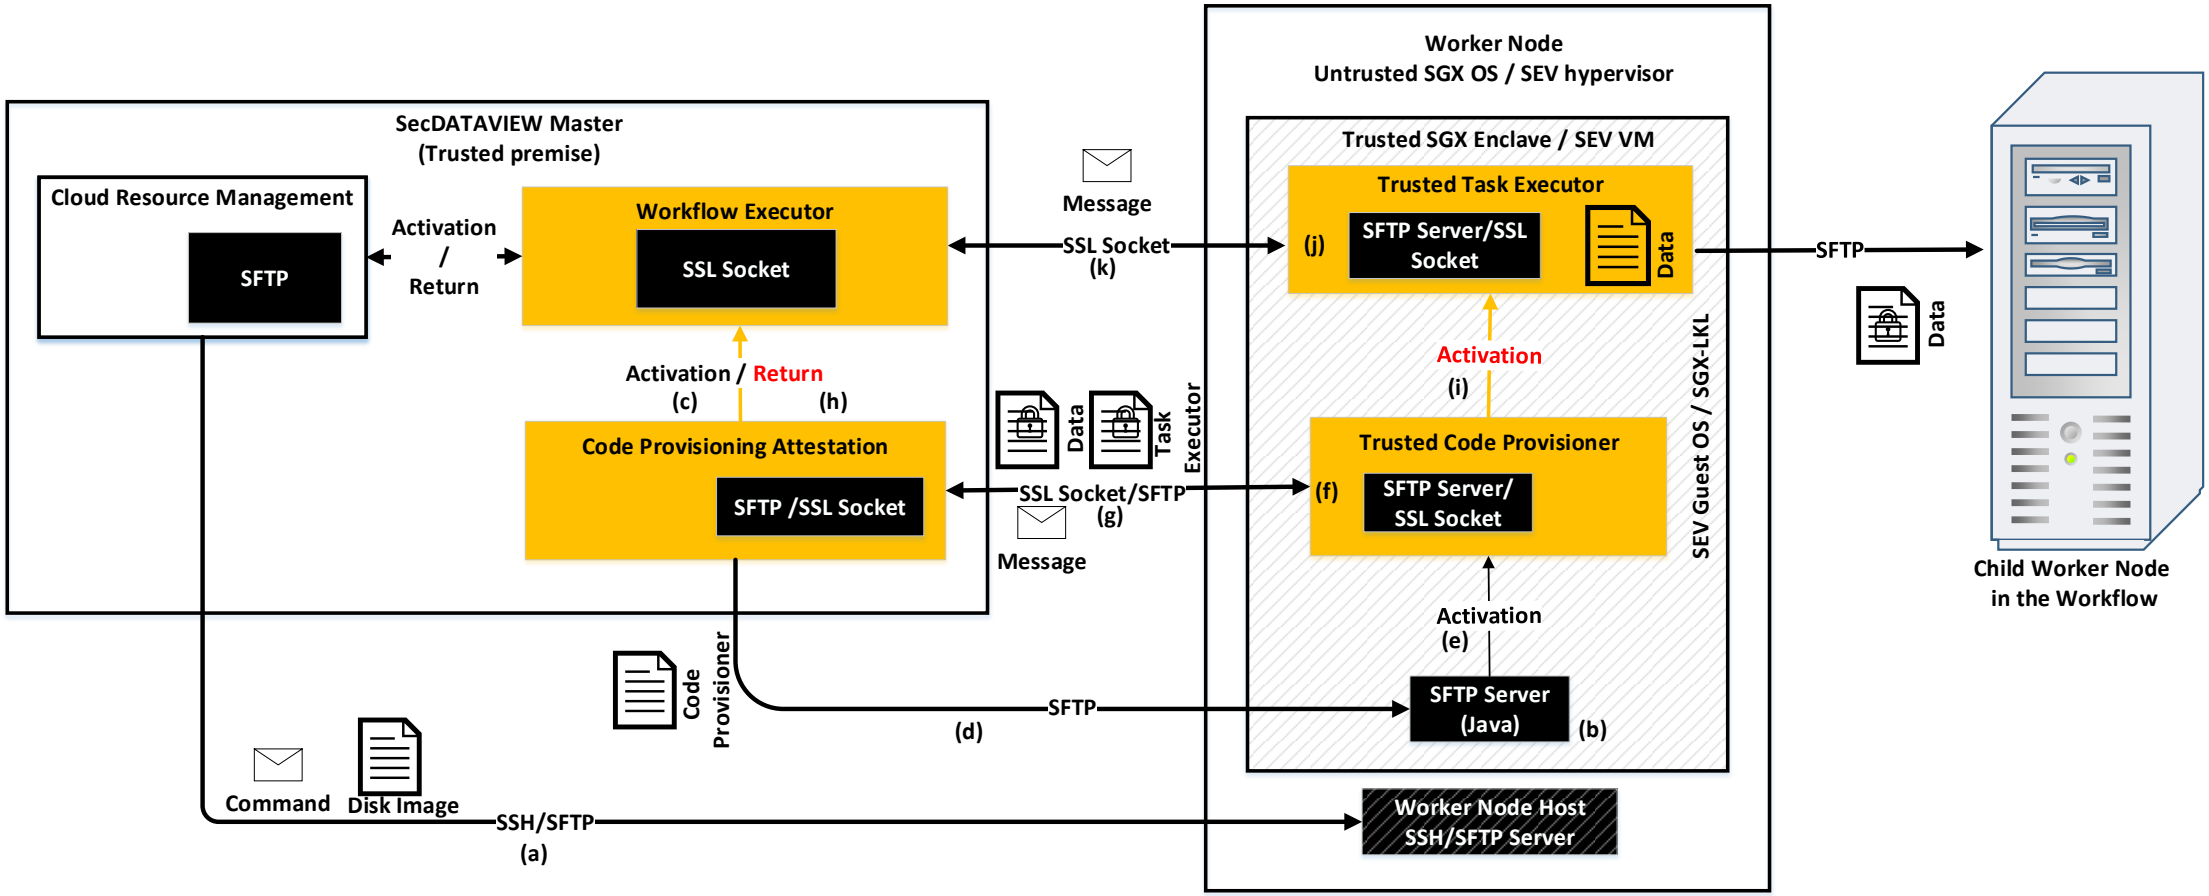

WCPAC: Workflow Code Provisioning and Communication Protocol

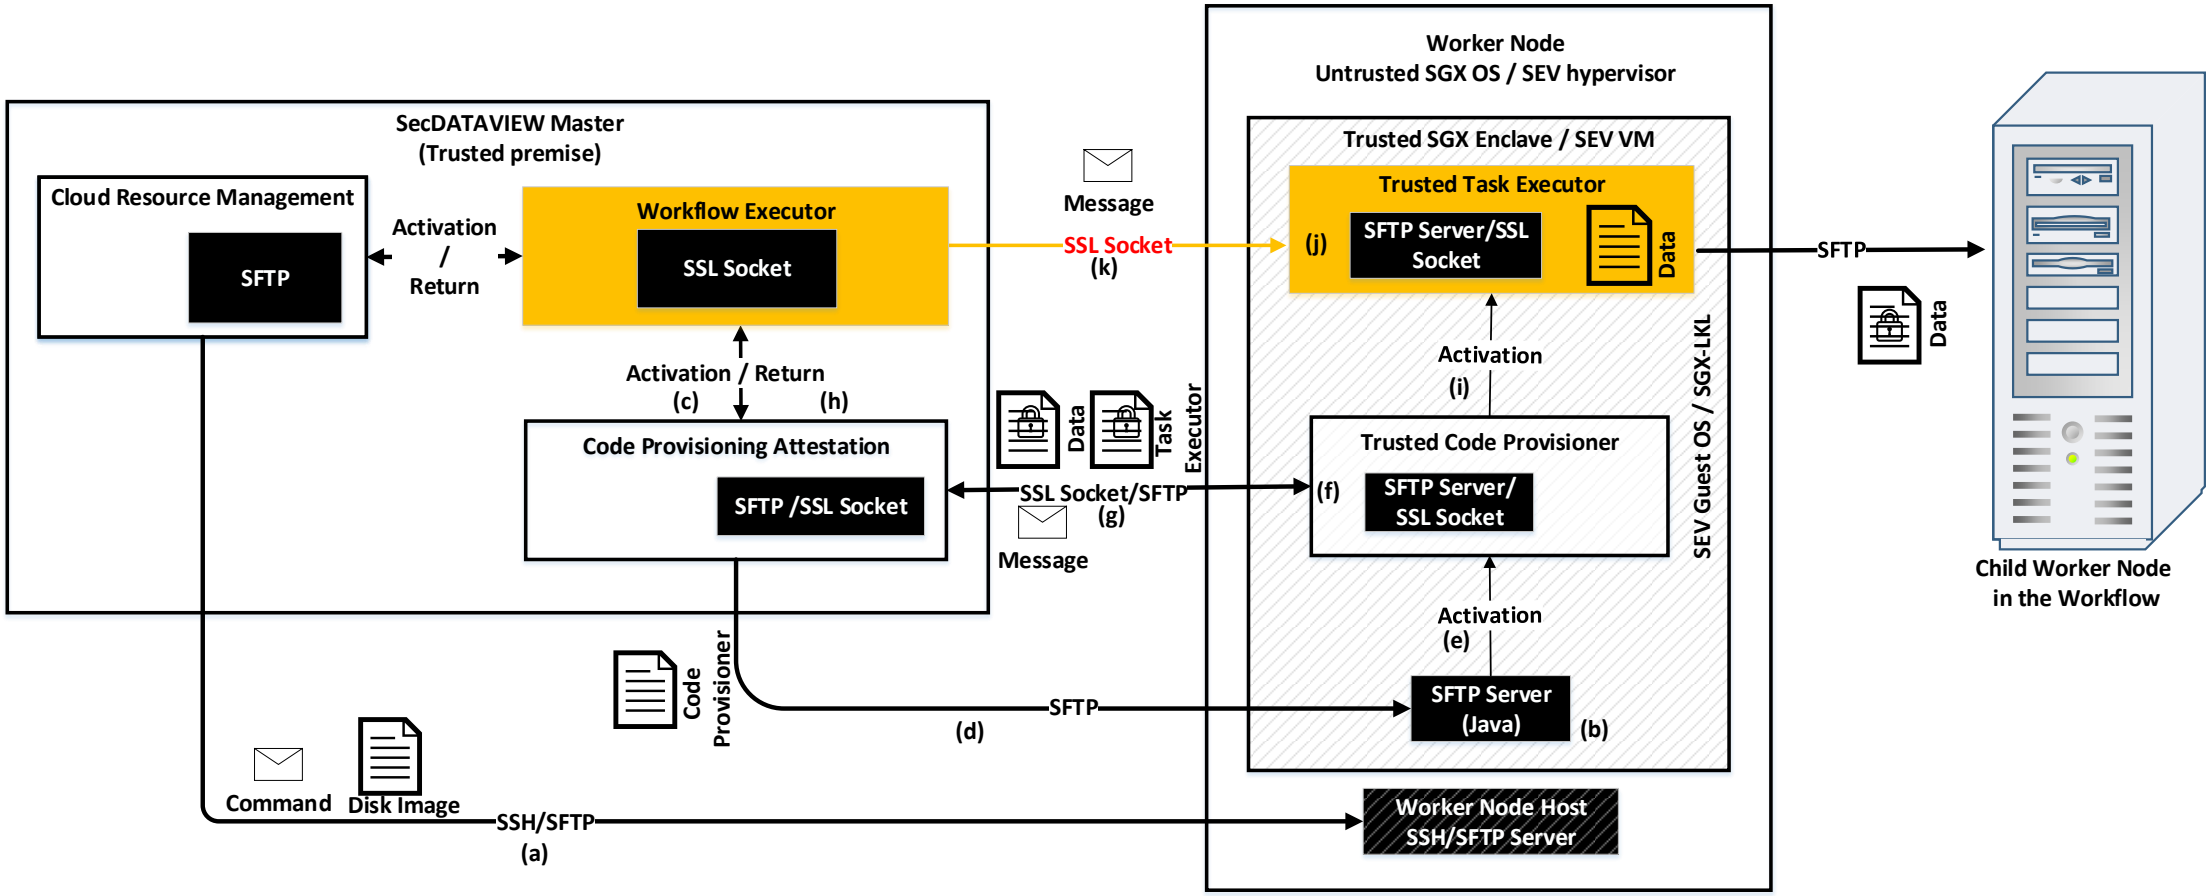

WCPAC: Workflow Code Provisioning and Communication Protocol

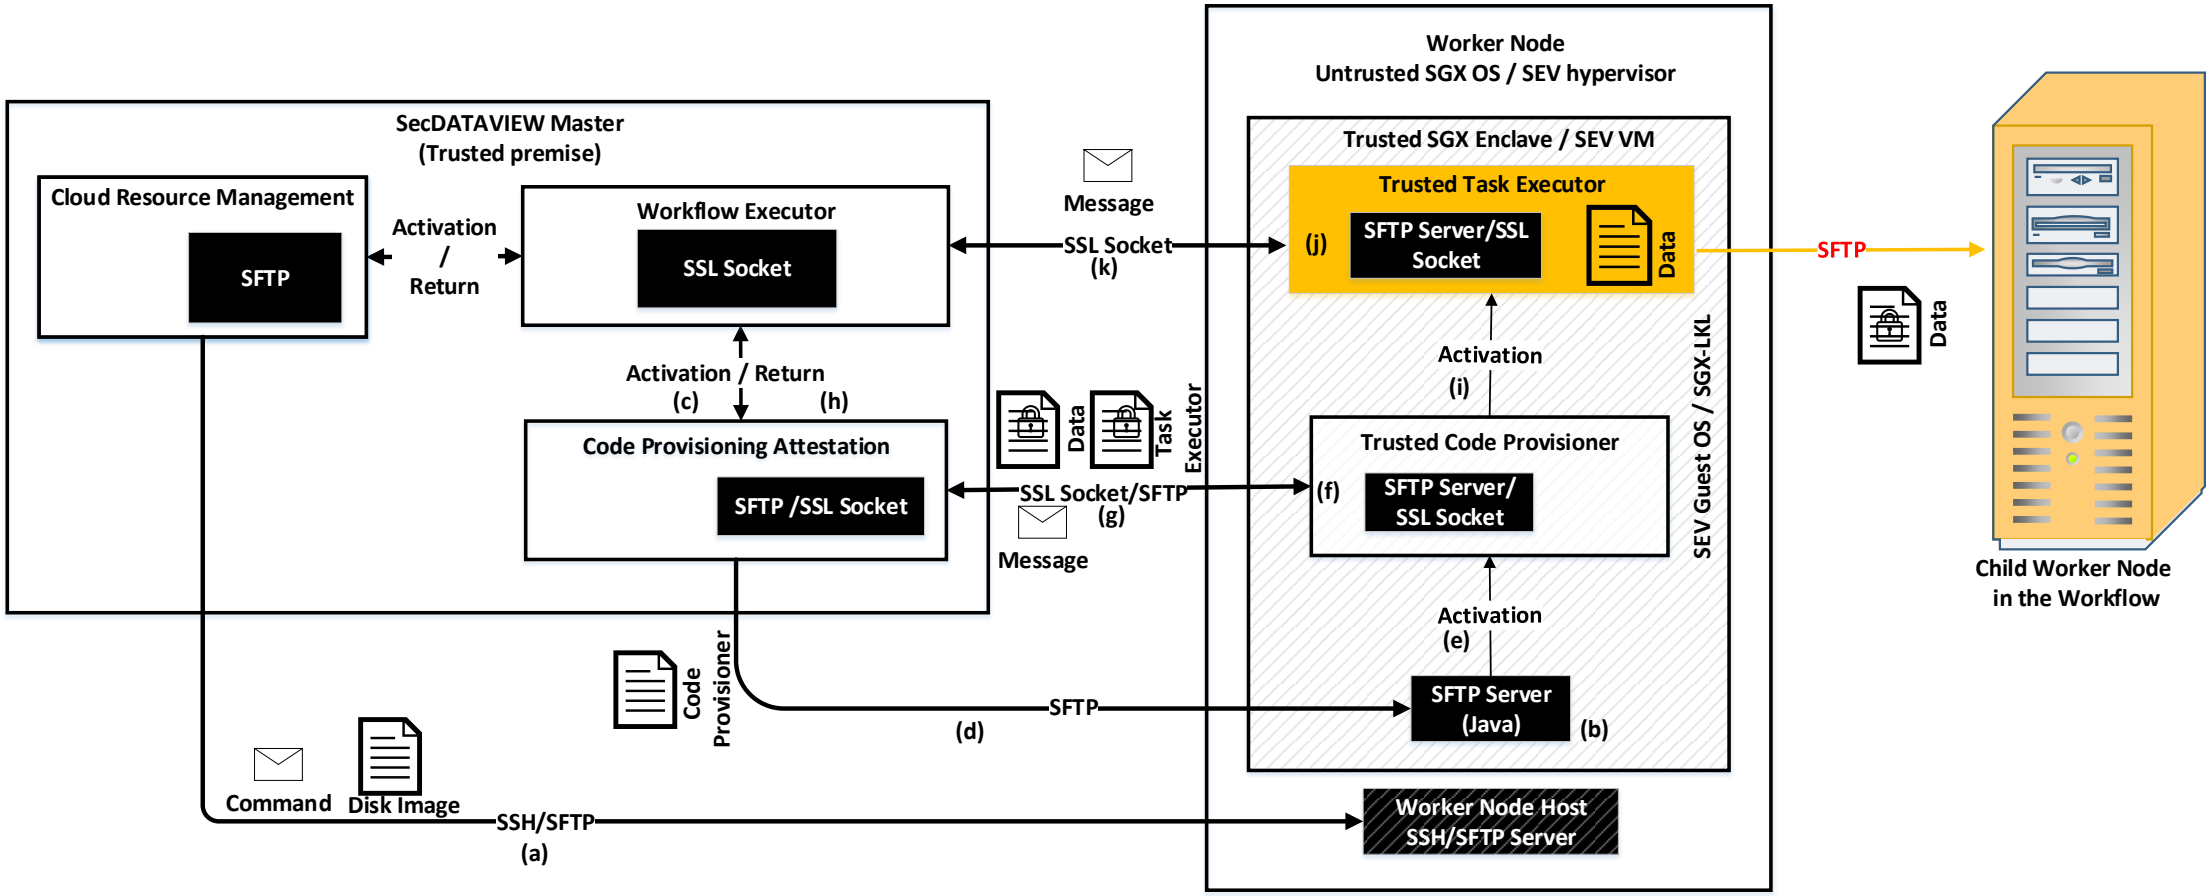

WCPAC: Workflow Code Provisioning and Communication Protocol

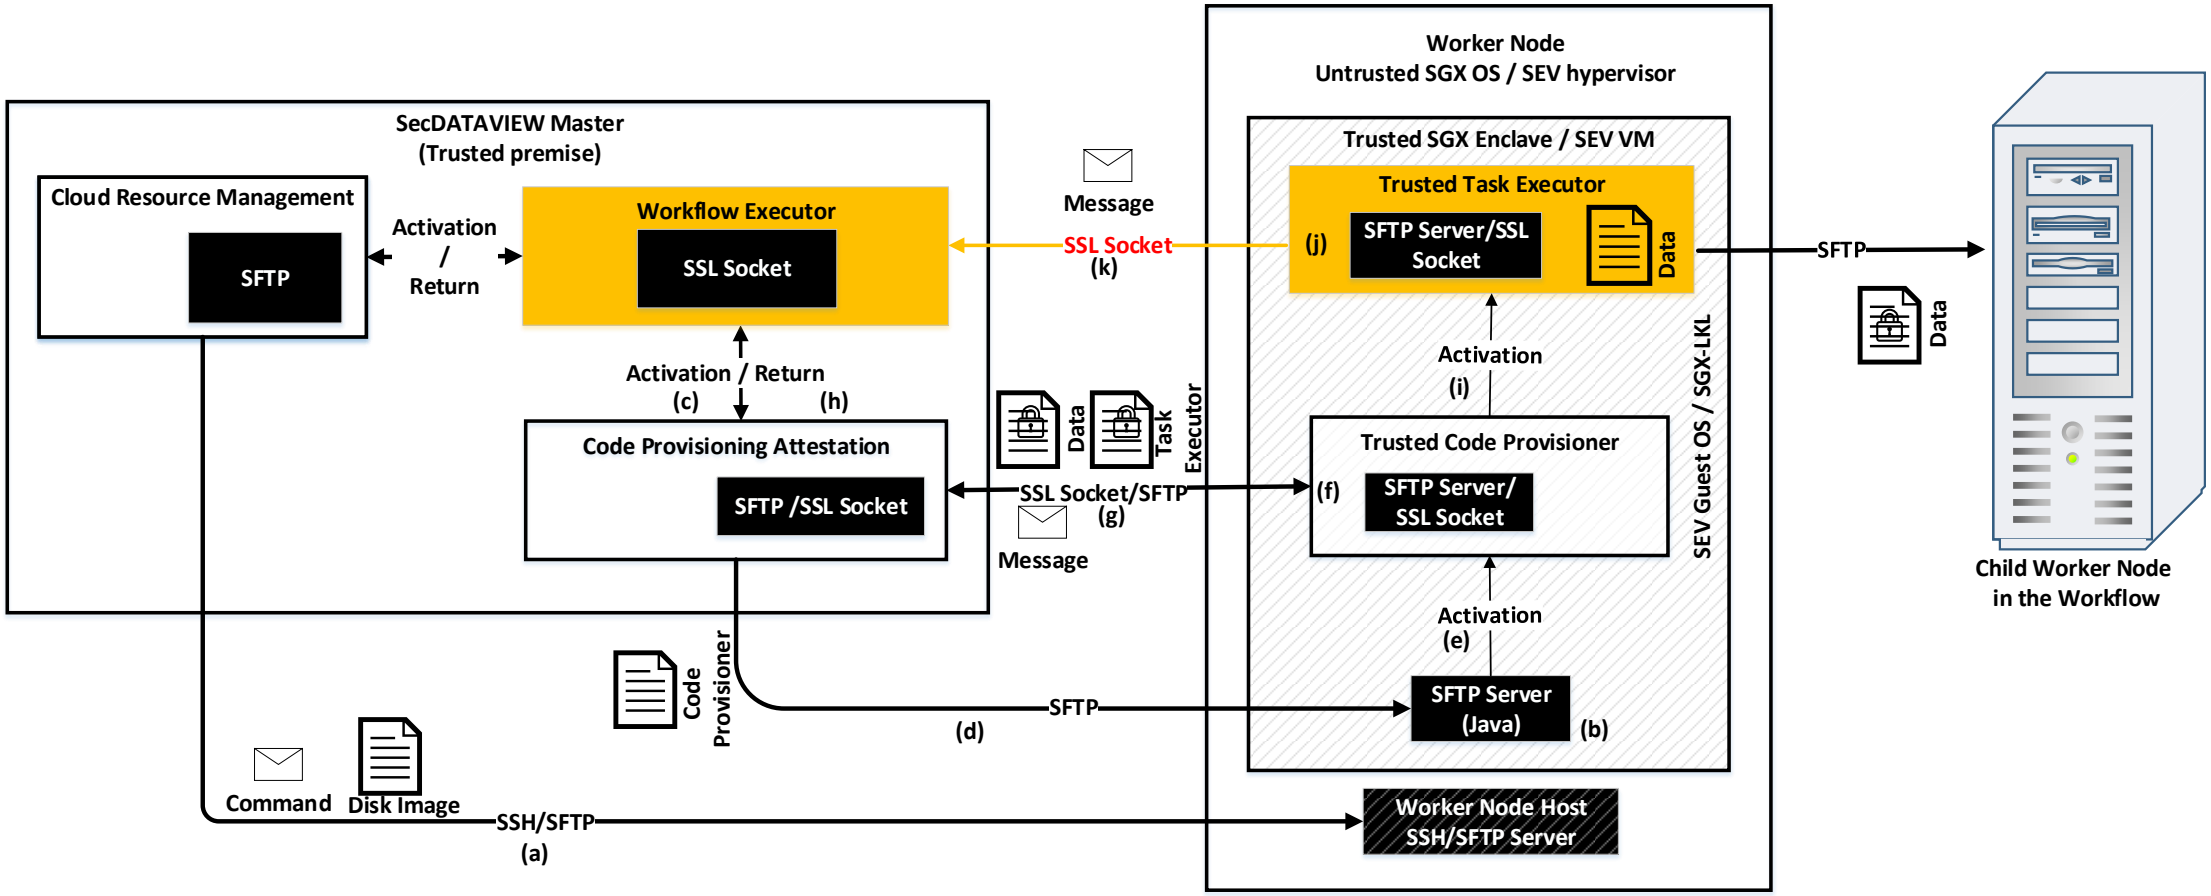

Outline

- Introduction
- X86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

Testbeds Configuration

Testbed Machine	SecDATAVIEW Master	Intel Worker	AMD Worker
CPU Model	Intel Core i7-6700T	Intel Xeon E3-1275 v5	EPYC-7251
Motherboard	Dell Inspiron 24-5459	Intel FOG	GIGABYTE MZ31-AR0
Memory	12GB DDR4 Non-ECC	32GB DDR4 No-ECC	32GB DDR4-ECC
Storage	Conventional HDD	NVME SSD	SATA SSD
Operating System	Linux 16.04 LTS	Linux 16.04 LTS	Ubuntu 18.04 LTS
OS, Hypervisor kernel	4.15.0-50-generic-x64	4.15.0-50-generic x64	4.20.0-sev-x64
TEE SDK Version	N/A	SGX SDK Ver 2.00	N/A
SGX-LKL	N/A	Hardware Mode	N/A
SGX-LKL Memory	N/A	2GB (Encrypted)	N/A
SGX-LKL Storage	N/A	2GB (Encrypted Disk Image)	N/A
SEV VM Kernel	N/A	N/A	4.18.20-generic-x64
SEV VM Memory	N/A	N/A	4GB (Encrypted)
SEV VM Storage	N/A	N/A	32GB (Disk Image)

The Diagnosis Recommendation Workflow [15]

- SGX mode overhead 2.62x
- SEV mode overhead 1.29x
- Hybrid mode overhead 1.20x
- Hybrid mode used two SGX and two SEV workers.

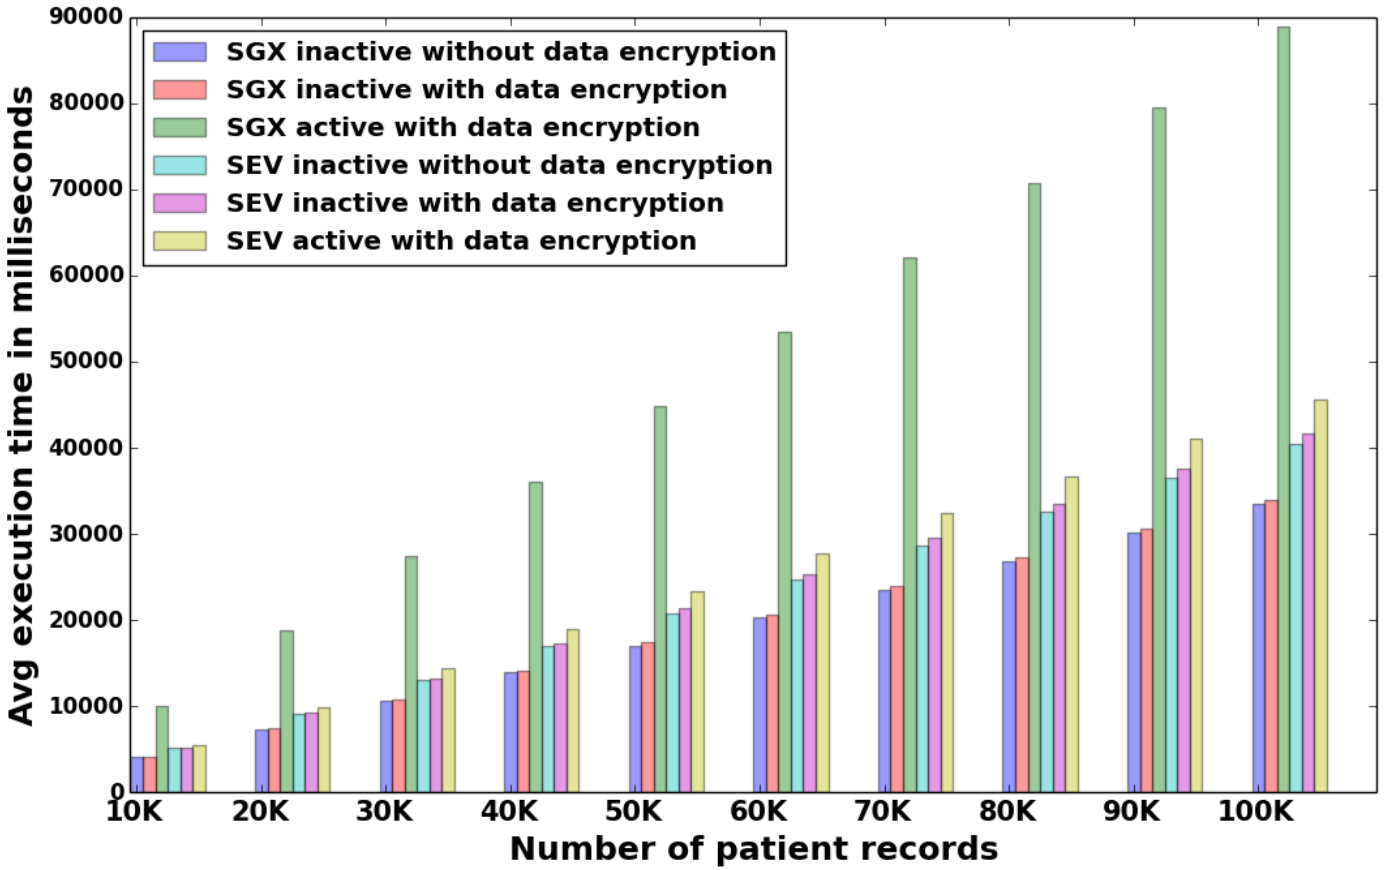

The Word Count (Map-Reduce) workflow [16]

- SGX mode overhead 1.89x
- SEV mode overhead 1.04x
- Hybrid mode overhead 1.33x
- Hybrid mode used two SGX and two SEV worker

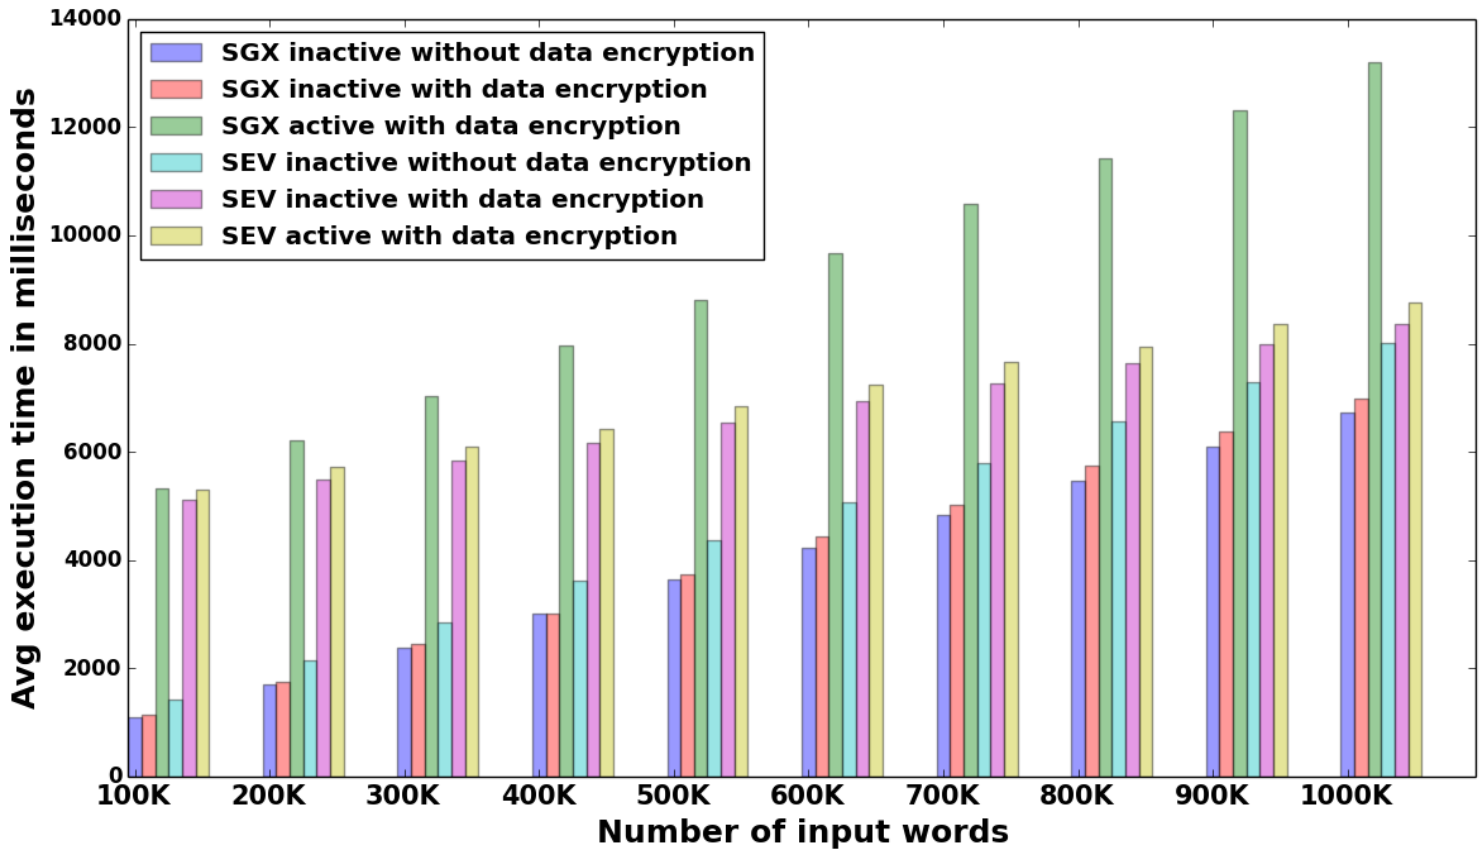

The Distributed K-means workflow [17]

- SGX mode overhead 1.69x
- SEV mode overhead 1.29x
- Hybrid mode overhead 1.43x
- Hybrid mode used two SGX and two SEV workers

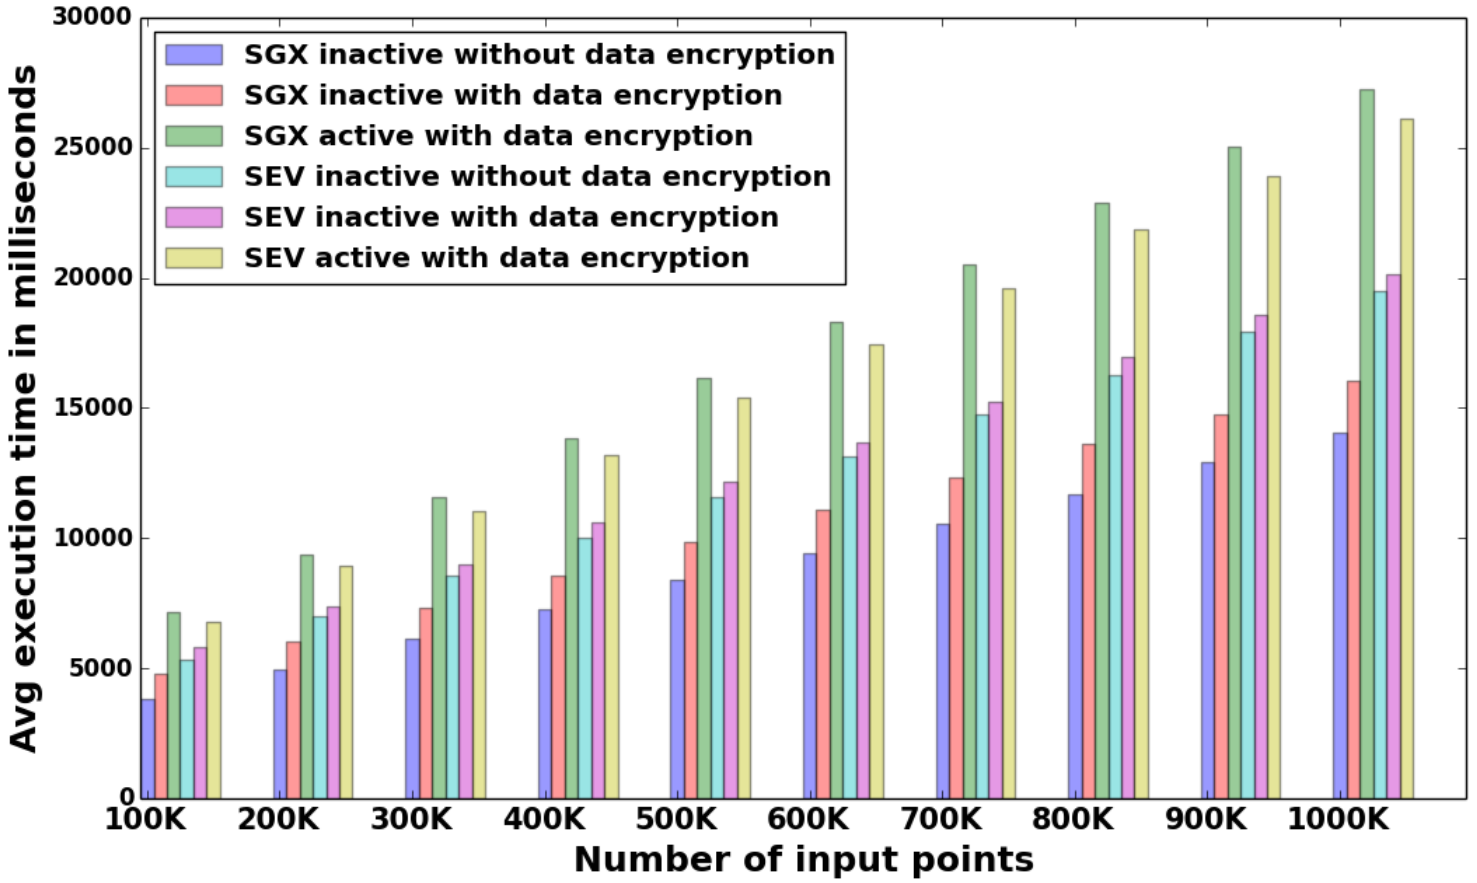

SecDATAVIEW: Security and TCB Analysis

SecDATAVIEW Intel SGX Worker:

- The software TCB is the LibOS, the JVM, the Code Provisioner, and the Task Executor
- The hardware TCB is the CPU package for the SGX workers

SecDATAVIEW AMD Worker:

- The software TCB is the guest OS, the JVM, the Code Provisioner, and the Task Executor
- The hardware TCB is AMD SoC and AMD secure processor for the SEV worker nodes

- SecDATAVIEW is protected against memory corruption vulnerabilities (Java)
- Workflow runtime is protected with hardware-assisted TEE
- Network traffic is protected with SSL protocol
- User data and results are protected with AES GCM-256 AEAD cryptography scheme

Functional Comparison with Existing Systems

Feature	SecDATAVIEW	VC3	Opaque	Lua Map/Reduce
Data confidentiality	AES-GCM-256	AES-GCM-128	AES-GCM-128	AES-CTR-128
Data integrity	Authenticated Encryption	Authenticated Encryption	Authenticated Encryption	No
Intel SGX	Yes	Yes	Yes	Yes
AMD SEV	Yes	No	No	No
Data structure compatibility	All types of workflow	Map-Reduce	SQL query	Map-Reduce
Job integrity verification	No	Yes	Yes	No
Access pattern leakage protection	No	No	Yes	No
Access pattern leakage overhead	N/A	N/A	1.6X-46X (oblivious mode)	N/A
Job performance overhead	1.2X-1.43X (hybrid mode)	1.04X-1.08X (base-encrypted mode)	0.52X-3.3X (encrypted mode)	1.3X-2X (encrypted mode)

Outline

- Introduction
- X86 TEE technology background
- Previous data analytics systems with TEE support
- SecDATAVIEW
- Performance results and security comparison
- Conclusions and future work

Conclusions and Future Work

- SecDATAVIEW, is an efficient and secure big data workflow management system that protects the confidentiality and integrity of Java-written tasks and data in the workflow with the help of SGX/SEV worker nodes.
- SecDATAVIEW significantly reduces the TCB size of the worker node and protects the Task Executor and individual workflow tasks by executing them inside the SGX enclave or the SEV-protected instance.
- Our experiments with different types of workflows show the usability of the system with a low-performance overhead while securing the confidential task execution at SGX enclave/SEV instance runtime.
- Future work: Investigate the security issues of collaborative scientific workflows [17]

References

[1] McKeen *et al.*, “Innovative instructions and software model for isolated execution.,” in HASP@ISCA, 2013, p. 10.

[2] Kaplan *et al.*, “AMD memory encryption,” White Paper. April 2016.

[3] Mofrad *et al.*, “A Comparison Study of Intel SGX and AMD Memory Encryption Technology,” in *Proceedings of the 7th International Workshop on Hardware and Architectural Support for Security and Privacy*, 2018, pp. 9:1–9:8.

[4] Zhang *et al.*, “SoK: A Study of Using Hardware-assisted Isolated Execution Environments for Security,” in *Proceedings of the Hardware and Architectural Support for Security and Privacy 2016*, 2016, p. 3.

[5] Ristenpart *et al.*, “Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds,” in *Proceedings of the 16th ACM conference on Computer and communications security*, 2009, pp. 199–212.

[6] Bugiel and others, “AmazonIA: when elasticity snaps back,” in *Proceedings of the 18th ACM conference on Computer and communications security*, 2011, pp. 389–400.

[7] F. Schuster *et al.*, “VC3: Trustworthy data analytics in the cloud using SGX,” *Proc. - IEEE S & P. Priv.*, vol. 2015-July, pp. 38–54, 2015.

[8] W. Zheng, A. Dave, J. G. Beekman, R. A. Popa, J. E. Gonzalez, and I. Stoica, “Opaque: An Oblivious and Encrypted Distributed Analytics Platform.,” in *NSDI*, 2017, pp. 283–298.

[9] Pires et al., “A lightweight MapReduce framework for secure processing with SGX,” in Cluster, Cloud and Grid Computing (CCGRID), 2017 17th IEEE/ACM International Symposium on, 2017, pp. 1100–1107.

References Continue

[10] A. Kashlev, S. Lu, and A. Mohan, “Big Data Workflows: a Reference Architecture and the {DATAVIEW} System,” *Serv. Trans. Big Data*, vol. 4, no. 1, pp. 1–19, 2017.

[11] J.-F. Raymond, “Traffic analysis: Protocols, attacks, design issues, and open problems,” in *Designing Privacy Enhancing Technologies*, 2001, pp. 10–29.

[12] Dinh et al., “M2R: Enabling Stronger Privacy in MapReduce Computation.,” in *USENIX Security Symposium*, 2015, pp. 447–462.

[13] Wang and others, “Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 2421–2434.

[14] Barenghi, L. Breveglieri, I. Koren, and D. Naccache, “Fault injection attacks on cryptographic devices: Theory, practice, and countermeasures,” *Proc. IEEE*, vol. 100, no. 11, pp. 3056–3076, 2012.

[15] Ahmed et al., “Diagnosis Recommendation using Machine Learning Scientific Workflows,” in *Big Data Congress, 2018 IEEE International Conference on*, 2018.

[16] Dean et al., “MapReduce: simplified data processing on large clusters,” *Commun. ACM*, vol. 51, no. 1, pp. 107–113, 2008.

[17] S. Lu and J. Zhang, “Collaborative scientific workflows,” in *2009 IEEE International Conference on Web Services*, 2009, pp. 527–534.

[18] https://www.flickr.com/photos/waynestateise/47529826741/

[19] Baumann and others, “Shielding applications from an untrusted cloud with haven,” *ACM Trans. Comput. Syst.*, vol. 33, no. 3, p. 8, 2015.

[20] Priebe, Christian, et al. "SGX-LKL: Securing the Host OS Interface for Trusted Execution." *arXiv preprint arXiv:1908.11143* (2019).

Thank You!

Email: saeid.mofrad@wayne.edu

The first release of SecDATAVIEW is available at

https://github.com/shiyonglu/SecDATAVIEW

Artifacts Evaluated – Functional