leidos

A Brief History of Separation Kernels

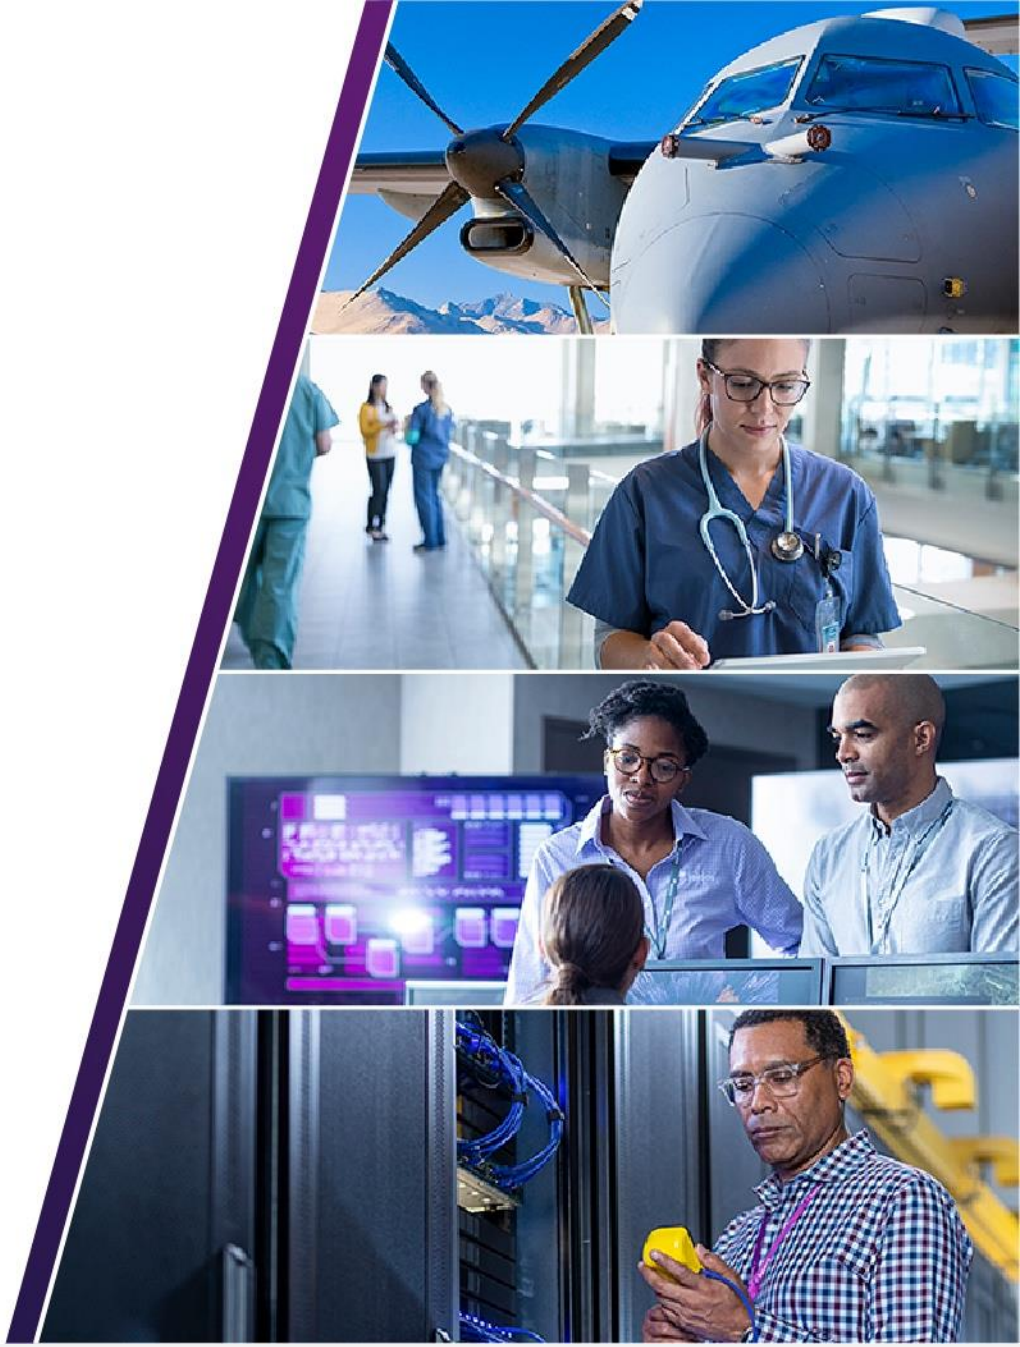

Dr. Raymond Richards

© 2021 Leidos. All rights reserved

In the Beginning…

- The concept of a security kernel that provides strong separation was introduced in 1981
  - John Rushby, "The Design and Verification of Secure Systems“, Eighth ACM Symposium on Operating System Principles, December 1981
- A separation kernel provides environments that are indistinguishable from a distributed (air-gapped) system
  - Only the explicitly authorized flow of information is allowed
  - Requires a proof of separability
- Provides a separation of concerns in the architecting of complex systems
  - Passes the Richards test

© 2021 Leidos. All rights reserved

Approved for Release: 21-leidos-1217-23862

Advances in trustworthy software architectures

- Research in Provably Dependable Software Architectures – 1990s
  - Victoria Stavridou, “Provably Dependable Software Architectures”, Proceedings of the Third International Workshop on Software architecture, November 1998
- Development of Integrated Modular Avionics -1990s
  - Integrated architecture with application software portable across an assembly of common hardware modules
- Multiple Independent Levels of Security (MILS) Architecture – 2000s
  - Integrated architecture with application software processing data at different levels of classification

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

A strong foundation

- ARINC 653 – 1990s
  - A standard for time and space partitioning of processes for safety critical systems
    - No formal specification of partitioning
  - Applications of mixed criticality hosted on a single processor
  - A fault in an application has no impact on any other applications
  - A variety of ARINC-563 compliant commercial RTOSes

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

An early separation kernel

- Mathematically Analyzed Separation Kernel (MASK) Early 2000s
  - A correct-by-construction kernel
  - A collaboration between Motorola, NSA, and Kestrel
    - Martin, W. & White, P. & Taylor, F.S. & Goldberg, A.. (2000). Formal construction of the Mathematically Analyzed Separation Kernel.
  - Employed SPECWare methodology

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

Microcode – Hardware of Software?

- The AAMP7 Microprocessor include intrinsic partitioning for mixed criticality avionics processing
  - Strong space and time partitioning implemented in microcode
  - ARINC 653 compliant scheduling
- A formal model developed of the partitioning system implemented in the AAMP7
  - A formal proof of separability has been developed
  - The NSA inspected the model and determined that it accurately represented the implementation
- Certified to be able to process unclassified through Top Secret Codeword information simultaneously
  - May 2005
- First Certified Separation Kernel

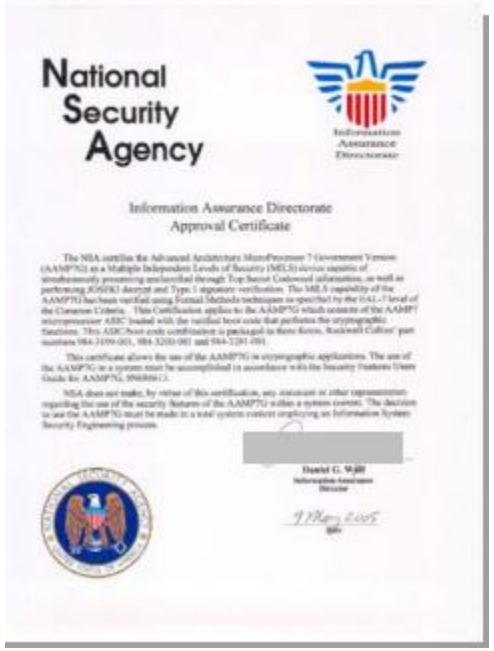
National Security Agency

Information Assurance Directorate

Information Assurance Directorate
Approval Certificate

© 2021 Leidos. All rights reserved

Approved for Release: 21-leidos-1217-23862

A Proof of Separability!

- The US Government wanted proof that the AAMP7 prevented unauthorized infiltration, exfiltration, and mediation
  - How are these expressed in a mathematical notation?
  - No formal specification of separation existed in the literature at this time

- A formal specification of separation was developed by Greve, Wilding, and Vanfleet
  - Came to be known as the GWV Theorem

```
(implies
   (and
      (equal (current st1) (current st2))
      (equal (selectlist srcsegs st1) (selectlist srcsegs st2))
      (equal (select seg st1) (select seg st2)))
   (equal (select seg (next st1)) (select seg (next st2)))))
```

- A later paper showed that the GWV theorem is only valid for kernels with a very restrictive scheduling model
  - Jim Alves-Foss and Carol Taylor, “An Analysis of the GWV Security Policy”, 2005 ACL2 Workshop

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

Proof of Separability on an ARINC 653 RTOS

- The F-22 and F-35 programs jointly funded the certification of a commercial RTOS against the *U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness* (A.K.A. SKPP)
  - Common Criteria EAL 6+
- Green Hills Integrity-178 awarded certification in 2011
  - Certified in DO-178C DAL A environments
- Formally modeled at the source code level
- GWVr2, a generalization of GWV
  - Greve, Wilding, Vanfleet and ray richards
  - Correspondence proofs only required to be “semi-formal”
  - Only a proof of separability, not of functional correctness

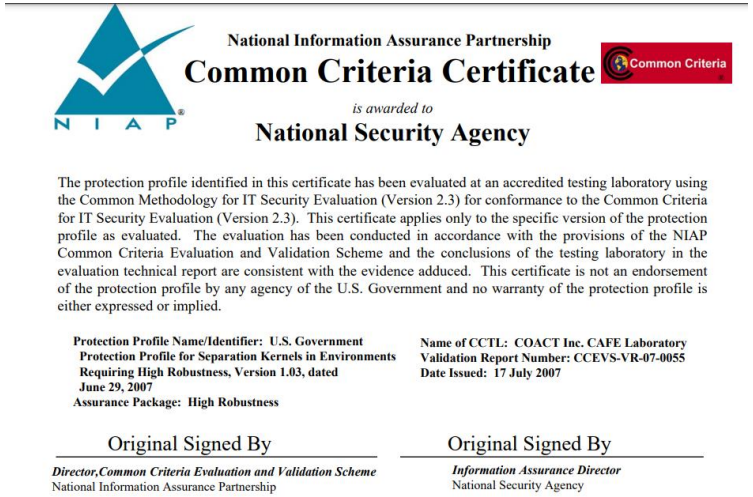

National Information Assurance Partnership

Common Criteria Certificate

is awarded to

National Security Agency

The protection profile identified in this certificate has been evaluated at an accredited testing laboratory using the Common Methodology for IT Security Evaluation (Version 2.3) for conformance to the Common Criteria for IT Security Evaluation (Version 2.3). This certificate applies only to the specific version of the protection profile as evaluated. The evaluation has been conducted in accordance with the provisions of the NIAP Common Criteria Evaluation and Validation Scheme and the conclusions of the testing laboratory in the evaluation technical report are consistent with the evidence adduced. This certificate is not an endorsement of the protection profile by any agency of the U.S. Government and no warranty of the protection profile is either expressed or implied.

Protection Profile Name/Identifier: U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness, Version 1.03, dated June 29, 2007
Assurance Package: High Robustness

Name of CCTL: COACT Inc. CAFE Laboratory
Validation Report Number: CCEVS-VR-07-0055
Date Issued: 17 July 2007

Original Signed By
Director,Common Criteria Evaluation and Validation Scheme
National Information Assurance Partnership

Original Signed By
Information Assurance Director
National Security Agency

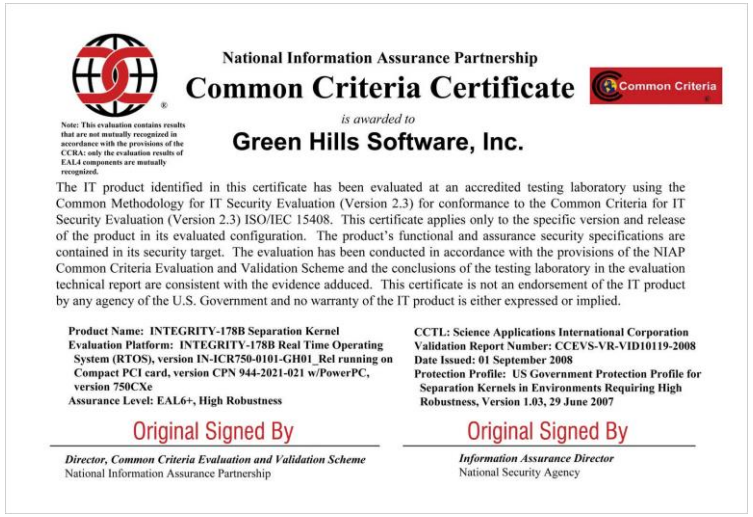

National Information Assurance Partnership

Common Criteria Certificate

is awarded to

Green Hills Software, Inc.

Note: This evaluation contains results that are not mutually recognized in accordance with the provisions of the CCRA: only the evaluation results of EAL4 components are mutually recognized.

The IT product identified in this certificate has been evaluated at an accredited testing laboratory using the Common Methodology for IT Security Evaluation (Version 2.3) for conformance to the Common Criteria for IT Security Evaluation (Version 2.3) ISO/IEC 15408. This certificate applies only to the specific version and release of the product in its evaluated configuration. The product's functional and assurance security specifications are contained in its security target. The evaluation has been conducted in accordance with the provisions of the NIAP Common Criteria Evaluation and Validation Scheme and the conclusions of the testing laboratory in the evaluation technical report are consistent with the evidence adduced. This certificate is not an endorsement of the IT product by any agency of the U.S. Government and no warranty of the IT product is either expressed or implied.

Product Name: INTEGRITY-178B Separation Kernel
Evaluation Platform: INTEGRITY-178B Real Time Operating System (RTOS), version IN-ICR750-0101-GH01_Rel running on Compact PCI card, version CPN 944-2021-021 w/PowerPC, version 750CXe
Assurance Level: EAL6+, High Robustness

CCTL: Science Applications International Corporation
Validation Report Number: CCEVS-VR-VID10119-2008
Date Issued: 01 September 2008
Protection Profile: US Government Protection Profile for Separation Kernels in Environments Requiring High Robustness, Version 1.03, 29 June 2007

Original Signed By
Director, Common Criteria Evaluation and Validation Scheme
National Information Assurance Partnership

Original Signed By
Information Assurance Director
National Security Agency

© 2021 Leidos. All rights reserved

Approved for Release: 21-leidos-1217-23862

Winter has come

- The US Government was interested in certifying a second commercial RTOS
  - The US National Information Assurance Partnership (NIAP) listed Wind River's Multilevel Secure Platform as "in evaluation for EAL 6+" in 2009
    - https://www.windriver.com/news/press/news-6841
  - Certification is never awarded
- NIAP suspends evaluations at EAL 5 and above
  - The US Common Criteria Evaluation and Validation Scheme (CCEVS) and NIAP needed to focus their resources on a backlog of EAL 2, 3, and 4 evaluations critical to national security
  - NIAP was not getting the necessary bang for their buck on high assurance (EAL 5, 6,and 7) evaluations
    - From my conversations with Margaret Salter, NSA circa 2013
- Winter has fallen on stand-alone certifications of separation kernels
  - SKPP sun set date September 1, 2011
  - No principled linkage between kernel-level assurance and system-level assurance
  - No one ever asked the proof developers how to leverage the proof of separability in their system architecture assurance story

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

Kernels with formal roots

- seL4 – seL4 Foundation
  - Extensive effort to develop formal specification and implementation of a microkernel with formal proofs that high level security properties are provided by the specification, and thet the implementation down to the binary is faithful to the specification

- CertiKOS – Yale University
  - A compositional approach of formally specified abstraction layers. Note that references to having a certified kernel are using the word “certified” in an academic sense. They have completed formal proofs, it has not withstood evaluation and been certified by a government agency.

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

Other secure kernels

- Hypervisors are the backbone of cloud computing
  - By and large, hypervisors lack proofs of separability
  - Virtual machine escapes is the term used to describe a breach of separability in a hypervisor
    - A hot topic of research and discussion
  - An ARINC-653 scheduler has been developed for the Xen hypervisor (open source)
    - Temporal partitioning
    - Steven Vanderleest, “ARINC 653 Hypervisor”, 29th Digital Avionics Systems Conference, 2010

- ARES Secure Kernel
  - Developed on the AFRL ARES and HADES efforts
  - The ARES platform is a manifestation of the MILS architecture
  - ARES Kernel developers started with seL4 source code, but quickly diverged into a unique kernel

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

Is that all?

- No, plenty of separation kernel technologies are not included in this Brief History
- For a more in depth treatment of this topic
  - Yongwang Zhao, David Sanán, Fuyuan Zhang, Yang Liu, “High-Assurance Separation Kernels: A Survey on Formal Methods”, 2017

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862

Conclusions

- Security is a race to the bottom
  - Separation kernels are the bottom of the software stack
    - Assuming secure boot
  - Formal verification provides assurance that nothing gets below the kernel
- There are still gaps in the assurance story
  - Hardware-Software boundary
    - This is where covert channels are found
  - Kernel-Software stack boundary
    - Principled use of the proof of separation in the composition of a secure system
- It is easy to build an insecure system on a proof of separatbility
  - Be careful out there!

© 2021 Leidos. All rights reserved
Approved for Release: 21-leidos-1217-23862