
INTEL REQUEST FOR PROPOSALS (RFP)

SUBJECT

Intel's University Research & Collaboration Office requests proposals for academic research on *Trustworthy Data Center of the Future (TDCoF)*.

KEY DATES

Proposal Submission Deadline (PIs): July 18th, 2022

Proposal Responses from Intel: September 2022

Planned Research Start: Late Q4 2022 or early Q1 2023

OVERVIEW

Intel's University Research & Collaboration Office (URC) invites proposals to establish a new research center which will focus on the Trustworthy Data Center of the Future (TDCoF).

Increasingly, private data and computations are migrating from on-premises to the public cloud and edge. While many individuals, companies, and organizations embrace the resulting efficiencies, for others, trust poses a barrier to adoption and is reflected in their reluctance to migrate confidential workloads to what they consider untrusted environments. Removing this barrier to adoption by securing data and compute has become a business imperative.

At the same time a new class of legal requirements for privacy and data protection is emerging as regulations such as GDPR, Schrems II and HIPAA take effect across the globe. These legal requirements mandate heightened confidentiality guarantees for personal data at all stages: during storage, during transmission and during processing of this data. Therefore, companies that are subject to these regulations have additional incentives to provide elevated security of their clients' personal data.

To secure data at all stages of processing, Confidential Computing has emerged as a new form of computing that protects data-in-use by performing computation in a hardware-based Trusted Execution Environment (TEE). Recognizing these trends, major Cloud Service Providers (CSPs) have adopted Confidential Computing solutions to stay out of the tenant trust boundary such that they have no visibility into tenants' data. As a result, we are witnessing a trend where

security-sensitive services are migrating to the public cloud. At the edge, similar trends are driving an increased adoption of Confidential Computing.

In tandem, motivated by developer productivity, microservices and Functions-as-a-Service (FaaS) have become de facto development practices in the cloud. This shift is resulting in software architectures where CSPs provide the infrastructure software to run and connect microservices and functions as Software-as-a-Service offerings. Developers favor microservices and FaaS for their reusability and logical decomposition, however this comes at a cost with CSPs reporting a large infrastructure tax of more than 25%¹. To overcome these limitations CSPs are responding with a shift to Infrastructure Processing Units (IPUs), offloading the burden of infrastructure services to more cost-effective hardware.

Compute intensive workloads, for example, AI/ML and data analytics are driving accompanying trends in HW accelerator (FPGA, GPU, TPU, ASIC) deployments due to their power and performance efficiencies. To maximize use of these high value assets, there is a growing trend towards disaggregated computing where distributed compute resources are dynamically assembled into virtual platforms to meet workload requirements. Emerging high-performance network technologies and protocols for efficient memory copy across the network have been instrumental in making disaggregated computing practical.

The TDCoF collaborative research center will conduct innovative research on the security and privacy implications of the above shifts in datacenter architecture and software abstraction innovations. While we observe from industry and research that the trend toward disaggregated architectures has already begun, e.g., with foundational infrastructure workloads already running on IPUs, it is unclear what the optimal balance of disaggregation is. We expect that identifying the right performance, flexibility, pooling and security tradeoffs will be challenging in this dynamic and evolving space.

At the intersection of these trends, we expect an evolution in Confidential Computing to deliver the Trusted Datacenter of the Future, while maintaining the promises of infrastructure tax reduction and increased performance and efficiency.

PROGRAM SCOPE AND FUNDING

Intel intends to fund a new collaborative, multi-university research center addressing the four Research Vectors (RVs) described in detail below. We suggest that each submitting organization focus their proposal on one or two research vectors (in line with their primary expertise) and identify key contributions that are expected.

¹ Google. Profiling a warehouse-scale computer. <https://dl.acm.org/doi/10.1145/2749469.2750392>

RESEARCH VECTORS IN DETAIL

RV1: SECURE HW ARCHITECTURES

This research vector will focus on the security implications of evolving hardware architectures in the Trusted Data Center of the Future.

Given the evolving environment outlined in the overview, this RV should advance research that enables trusted and scalable workload isolation and elasticity while comprehending the control point shift from general purpose CPUs to IPU and heterogeneous accelerators. Novel security models to segment, slice and isolate heterogeneous resources, enable secure multitenancy and enable a security plane at the system-level are required. A key goal is to understand how Confidential Computing can evolve to comprehend the trends toward disaggregation and heterogeneity.

In this RV we will explore the solution space for end-to-end secure microservices and functions with accelerator disaggregation, scalable fabrics, infrastructure offload, multiple roots-of-trust and composable attestation as key focus areas.

Example research directions, including, but not limited to:

- Secure isolation architectures with a focus on efficiency and high-performance communication
- IPU isolation architectures to deliver security assurances for offloaded infrastructure and services, including tenant services
- Architectures to provide security agility, enabling device vendors, CSPs, and tenants to make security, performance, power and manageability tradeoffs
- Trusted Execution Environments for heterogeneous compute elements
- Attestation capabilities that comprehend composability, dynamically changing pools of heterogeneous resources and multiple Roots-of-Trust
- The implications for security of migrating higher layers of the tenant stack to IPU e.g., RPC, TLS, Service Meshes, Message Buses.

RV2: SECURE SYSTEM ARCHITECTURES

This RV aims to develop innovative and secure software architectures for TDCoF.

In a disaggregated data center running microservices, there is a need for scalable and secure communication protocols achieving low tail latencies with a focus on zero copy, efficient data encryption and key management schemes for cloud native environments. With offloading of infrastructure services, this RV should explore the security implications of such approaches as well

as research novel alternatives. As a complimentary approach CSPs are developing custom software architectures tailored to workloads achieving 100x higher density and elasticity².

At the same time many of these workloads demand higher security guarantees given a shift to non-traditional data center deployments such as Edge Cloud. The RV will also focus on the evolution and role of confidential computing in this evolving data center that is distributed, disaggregated, and decentralized.

Example research directions, including, but not limited to:

- Research and develop techniques and methods for attaching trust and confidence levels to data in data centric computing
- Explore distributed data containers and associated storage engines for cloud native data management systems that self-organize into optimal solutions. One key aspect will be to research solutions that lead to data-containers that are “inherently” secure and privacy-preserving
- Research and enhance existing authorization and authentication mechanisms with code-based identities via attestation integration
- Understand the implications of confidential computing in a disaggregated data center environment spanning from cloud to edge
- Research the evolution of orchestration methods to tackle the challenges emerging from multi-tenant, heterogeneous and disaggregated compute across heterogeneous networks in a secure manner
- Secure software isolation architectures with a focus on high performance communication
- Develop novel HW/SW architecture (e.g Memory-Safe Architecture) to not only eliminate memory safety bugs but to also eliminate tail latencies due to cloud native microservices/FaaS while delivering on Confidential Computing
- Build next generation serverless runtimes using in-process isolation to improve elasticity and communication overheads
- Explore novel two-way sandboxes for high assurance solutions running in a hostile environment like untrusted edge nodes

RV3: APPLICATIONS AND WORKLOADS

The goal of this RV is to understand key characteristics of the top workloads driving the definition of the Trustworthy Data Center of the Future. The intersection of the TDCoF with Edge computing creates opportunities for new use cases, for example, the metaverse, autonomous vehicles, smart home/city/factory/warehouse, and the retail, medical and industrial verticals.

Example research directions, including, but not limited to:

² <https://blog.cloudflare.com/mitigating-spectre-and-other-security-threats-the-cloudflare-workers-security-model/>

- What are the key characteristics of the top workloads driving the definition of trustworthy warehouse-scale platforms?
- How to deliver the required security properties for disaggregated, massively distributed workloads in diverse and heterogeneous environments?
- Security requirements for future apps (functions) to execute anywhere (edge to cloud) while maintaining the security properties of defined policies
- Develop threat models and benchmarks

RV4: WILDCARD RESEARCH VECTOR

GOAL: IMPACTFUL RESEARCH TOWARDS OUR GOAL NOT COVERED IN RV1-3.

It is possible that we overlooked an essential research question for reaching the goals of this program. If you have a strong belief (and ideally also some evidence) that a major research question has been overlooked, you can submit a research question as “wild card”. In this case, the burden is higher, and you need to argue that the research that you plan to pursue (a) is a promising approach to reach the goals stated above and (b) that your research does not fit into the research vectors documented above.

PROPOSAL FORMAT

Please note that Intel is unable to receive proposals under an obligation of confidentiality. All proposals submitted should therefore include only public information.

Proposals should be 4-8 pages, not including citations or cost volume. We slightly prefer proposals that define a project for one Principal Investigator (PI) for up to three years. Collaborative proposals between two Principal Investigators with complementary domain expertise are also encouraged; for example, one PI with expertise in silicon circuit design may collaborate with another PI on hardware architecture and performance optimization to fully evaluate a resilient approach. Researchers can be part of only one proposal. Each proposal should comprise the following sections:

- **Proposal cover page (max 1 page)**
 - **Organization**
 - **List of PIs and the main contact person**
 - **List one or at most two targeted research vectors**
 - **Executive summary** including intended outcomes. Summarize the key elements of the proposal.
- **High-level motivation, preliminary results, approach, and proposed goals for the research questions (<= 3 pages).** Briefly describe the motivation for the proposed project, preliminary results, techniques (especially novel ones) that underpin the approach, and

the plan of tackling the proposed research questions. Summarize what will have been accomplished after 3 years if all goes according to plan. Be sure to detail the current state-of-the-art for the proposed technology (or nearest related technologies). This section must also include an explicit statement of the Intellectual Property (IP) status for all background IP related to this technology (i.e., are the property rights to this technology protected, and if so, who owns those rights).

- **Statement of work, schedule, milestones, success criteria and deliverables (<=3/4 page).** For each of the goals addressed, outline the 3-year scope of the effort including tasks to be performed, schedule, milestones, deliverables, and success criteria. It is understood that aspects of this research effort may be exploratory in nature and schedules/deliverables reflect intentions rather than a firm commitment.
- **Personnel plan and expertise statement (max. 1/4 page per Researcher).** Include a list of key personnel (at most 6) plus a statement on each person's role and time commitment. For each person, please add a brief bio or web page link and list their 6 most relevant prior publications (within the last 8 years) for the selected research questions.
- **Student plan (<1 page).** Please provide information about the PhD students and postdocs you envision to assign to this project (if known). Outline the approach and plan whereby PhD students will be recruited and incorporated into the team, and any plans for encouraging/supporting those students in collaborations with Intel (e.g., availability for Internships should a mutually interesting opportunity arise). If the PIs have a pre-existing relationship and history of student hiring by Intel please discuss issues/plans/ideas to continue or strengthen that connection.
- **Diversity and Inclusion (<1 page).** In light of Intel's strong commitment to diversity and creating an inclusive environment, please address: (a) your organization's commitment to diversity and inclusion with respect to race, national origin, gender, veterans, individuals with diverse abilities and LGBTQ, and (b) a summary of your performance in this area and any initiatives you are pursuing.
- **Prior Intel Collaborations (max 1/3 page per project).** If you collaborated with Intel in the past, please list the project/institute, the year, and the main contact(s) at Intel. Furthermore, add a short abstract outlining the scope.
- **Past Successful Technology Transfers (<1 page).** Evidence of past successful industry collaborations and technology transfers. Examples include startups, products, and other evidence of tangible business impact of the involved academics.
- **Budget and Financials (1/3 page).** Typical grants are USD\$70-140K per year for three years. We plan to work under an Open Intellectual Property model (results are published, code is open sourced). Our goal is to maximize the available research ideas for our fixed amount of total funding. Universities may propose how to achieve this. Please also indicate how many researchers (FTE) can contribute their research for the proposed funding.
- **IP-compatible funds amplification (no limit).** If the team can obtain funding for related work from other sources (including the University) and the sponsor commits to follow a public dedication approach for that project or provide Intel with non-exclusive, royalty-free **research and commercial** licenses to any IP, the team may list funding that would be considered to amplify the proposed project.
- **Citations {unlimited}.**

- **Cost volume {unlimited}**. Cost proposal in Excel or another format as appropriate.

EVALUATION CRITERIA

In order of importance, the evaluation criteria for this solicitation are as follows:

1. **Potential contribution and relevance to Intel and the broader industry:** The proposed research should directly support a technology solution that addresses the RVs outlined above, leading to technological advances with the potential for ongoing technology transfer in collaboration with Intel and the broader industry.
2. **Technical innovation:** Proposed solutions of interest should clearly push the boundaries of technical innovation and advancement. Research that is not of interest in this program include incremental advancements to state-of-the-art and current design practices. Feasibility of new algorithms/techniques should be demonstrated through SW/HW implementations.
3. **Clarity of overall objectives, intermediate milestones and success criteria:** The proposed Research Plan should clearly convey that the PIs have the knowledge and capability to achieve the stated research goals. It is understood that any research program will have uncertainties and unanswered questions at the proposal stage, but a clear path forward in key challenge areas must be identified and justified. Teams are expected to demonstrate progress toward project goals at quarterly milestones and monthly project status updates. As detailed in “Program Scope and Funding” section, the proposal should explicitly point out which RV is being addressed, the synergy among them if more than one RV, the plan and milestones towards building research prototypes, plan for ongoing technology transfers, and the anticipated proof of concept outcome. Strength of project management will also be considered.
4. **Qualification of participating researchers:** The extent to which expertise and prior experience bear on the problem at hand. Please elaborate on track records of building research prototypes (e.g., open-source research code/collaterals on GitHub) and resulting publications from past relevant projects.
5. **Cost effectiveness and cost realism:** The extent to which the proposed work is both feasible and impactful within the proposed resource levels will be examined.
6. **Potential for co-funding:** Opportunity for closely synergistic matching grants and co-funding with other funding entities, such as SRC, NSF, DARPA, NSERC, etc. will be given significant consideration.
7. **Potential for broader impact:** As an industry leader, Intel pushes the boundaries of technology to make amazing experiences possible for every person on earth. From powering the latest devices and the cloud you depend on to driving policy, diversity, sustainability, and education, we create value for our stockholders, customers, and

society. Intel expects the academic community to be strong partners in making Intel successful through support of Intel's goals and commitments to diversity, sustainability, and education. Intel supports the advancement of computing education and diverse participation in STEM. Significant consideration will be given to proposals in which the outcome of the research can influence the development of new curriculum initiatives impacting undergraduate or graduate education at the respective universities (e.g., exposure to latest industry technologies/tools in classroom setting). Proposals are encouraged to elaborate on how the proposed work is anticipated to impact student education on campus and/or the broader academic community.

PI MEETINGS AND COLLABORATION STRUCTURE

Intel will be deeply engaged with the center and will assign partner technologists/collaborators across RVs to interact with the academic community to produce a stream of innovation proof-points, publications, demonstrations, and technology transfers into Intel and the broader industry throughout the duration of the program. We aim for the interaction to be bi-directional where Intel collaborators are part of the research team. Not only will they provide research feedback, but they will also actively contribute and co-develop the research to amplify the center outcome and enable continuous technology transfers into Intel and the broader industry.

It is expected the PI and student researchers will collaborate on a daily or weekly basis. Monthly PI, student and Intel collaborator meetings will be used to review research results, present significant updates, and provide feedback.

Semi-annual face-to-face or virtual meetings will be held to facilitate center-wide information exchange, review, and discussion of research. Researchers should anticipate one annual face-to-face meeting to be held at an Intel site in the US or Europe and one annual face-to-face meeting to be held at a university associated with this center. Associated travel costs for two annual meetings should be considered and included in the proposed budget. In the event unexpected travel restrictions prohibit a face-to-face meeting, a virtual meeting will be held.

To aid in collaboration across projects within the center and communication of research findings to the public, it is anticipated that a center website will be established, hosted, and maintained and Intel requests the right to host the associated website link on their respective university program websites.

Intel will offer free access to Intel's Academic Compute Environment, a resource for academia researchers in the center to exercise their workloads on Intel's latest hardware.

For those researchers who are already funded and seek collaboration opportunities with Intel and other researchers in the area of this RFP, please let us know. One option is to participate in center activities (e.g., seminars, workshops, and hardware access) without Intel funding.

ELIGIBILITY

This RFP is open only to academic researchers and institutions that have been specifically invited to participate in the proposal process. However, invitees may freely select additional academic collaborators. Any questions regarding eligibility should be directed to Richard Chow and Frank McKeen (contact info below).

INTELLECTUAL PROPERTY

This solicitation affords proposers the choice of submitting complete program proposals for the award of a grant, a Sponsored Research Agreement, or other agreement as appropriate. Intel reserves the right to negotiate the final choice of agreement. Intel prefers that university research in the program be placed in the public domain (patentable inventions dedicated to the public and source code distributed under an open-source license similar to the Apache, BSD or MIT license). The final award terms are expected to follow a public dedication model. This means that either (1) Intel and the university will jointly agree that IP developed under an award will be placed in the public domain, including offering software under an open source license (Intel's preference as referenced above), or (2) if IP is not placed in the public domain, then all parties (the university, Intel and all third parties) will be afforded equivalent non-exclusive no-fee royalty-free rights to the research results for any commercial or non-commercial purposes, preferably with the right to sublicense third parties under such rights.

INTEL TEAM CONTACT INFO

The following individuals from Intel Labs are actively involved with the creation of this center.

Richard Chow, Program Director (richard.chow@intel.com)

Mona Vij, Co-Principal Investigator (mona.vij@intel.com)

Patrick Koeberl, Co-Principal Investigator (patrick.koeberl@intel.com)

Anand Rajan, Managing Sponsor (anand.rajan@intel.com)

Please submit proposals using the website:

<https://academic-rfp.intel-research.net/cur-rfp/TDCoF>.

Please send any inquiries to Mona Vij and Patrick Koeberl and copy Richard Chow. Please include "Submission for TDCoF" in the Subject of your email.

FAQ

WHAT IS THE TYPICAL GRANT AND PROPOSAL TEAM SIZE?

Proposals generally request grants in the range of \$70-140K per year. This would typically support 1 or 2 graduate students advised by 1 or 2 PIs.

WHAT IS THE ENVISAGED PROJECT DURATION?

Three years (there is a renewal process each year, but proposals should outline all 3 years with more details on year 1).

DO YOU CONSIDER PROPOSALS PRIMARILY CONCENTRATING ON THEORY/ALGORITHMS?

Proposals without a strong implementation/validation component are of interest, although it is strongly encouraged to provide evidence that the theory/algorithms will have use in practical systems. However, significant theoretical advances that will lead to practical solutions in the future are also welcome.

CAN YOU SPECIFY WHICH RESEARCHERS HAVE BEEN INVITED TO THIS RFP?

We don't release the names of invited researchers. Keep in mind that if you are seeking to partner with a specific academic PI, your partners do not have to be invited; you can choose to partner with any PI and share the RFP with them.

ARE WE ENCOURAGED TO SEEK CO-FUNDING OPPORTUNITIES?

While co-funding is not required, a proposal with co-funding or matching funding would be a strong plus.