B. Project Summary

Insider attacks present an extremely serious and pervasive security problem. In the national defense realm such attacks indicate espionage activities or at the very least lead to breaches that compromise national security. In the finance and banking sector, a single insider attack might cost billions of dollars in damages. Accurate insider threat detection has proved to be a very challenging problem, as many studies have shown that it is difficult to even cleanly define the notion of insider threats.

This proposal takes a step toward addressing the above challenge by formulating and devising machine learning-based solutions to the insider attack problem on *relational database management systems* (RDBMS), which are ubiquitous and are highly susceptible to insider attacks. We propose concrete plans to validate and evaluate our solutions by intimately collaborating with a large financial institution to build a prototype insider threat detection engine operating on their presently operational RDBMS. Limiting the scope of the insider threat detection problem to RDBMS allows us to make measurable progress on various technical fronts. We have a novel approach for insider threat detection in RDBMS whose potential has been demonstrated in a preliminary work. The main idea, also our conviction, is that the best way to distinguish normal vs. abnormal access patterns is to statistically model the *semantics* of users' queries. The key piece of information capturing a query semantics is the *data region* that the user is trying to access. This *query-semantics-based* approach to user profiling is the foundation upon which our proposed research is built.

The main thrusts of the proposed research program include: (1) expressive and accurate nested hierarchical and mixture learning models and effective detection algorithms based on the query semantics idea, developed along side a realistic threat model and evaluated on large and real-world RDBMS, (2) incorporate expert domain knowledge into the statistical model to improve false positive and false negative rates, (3) develop a set of statistical and mathematical tools for dealing with dynamic database updates, stealthy and collaborative attacks, and (4) address database performance and software tool development issues arising when evaluate and deploy the system in practice.

Intellectual Merit. First, the insider threat problem is a notoriously hard security problem to pin down. A technically feasible definition of "insider threats" not only helps design a solution, but also paves the way for feasible independent benchmarking and validation. Second, coming up with learning models and algorithms for capturing users behavior, taking into account data updates and temporally correlated query results, is an important problem in itself. Third, our problem domain provides an important and challenging ground for applications and further developments of a number of areas of active interest in both statistics and machine learning communities, including hierarchical and non-parametric modeling and structured prediction. Fourth, a good threat model and assessment methodology for insider attack mitigation in RDBMS is an important problem to the security research community at large. Fifth, our partnership with a large financial institution directly bridges the gap between theory and practice. The evaluation is done via red-teaming exercises performed by banking security experts, whose results feedback immediately into our models. Finally, by integrating ideas from statistics, machine learning, algorithms, databases and security, and by working directly with an industrial partner, this project will introduce new issues and develop new techniques that have the potential of making a tangible and profound impact on these fields.

Broader Impact. Our integrated research, evaluation, and deployment plan gets as close as possible to a practical tool for insider threat mitigation in RDBMS in the the real world. This research uniquely connects database systems, security and machine learning, opening up new research directions interconnecting those areas. The research will be conducted under the aegis of the SUNY Buffalo's NSA-certified center of excellence (CAE) in Information Assurance whose mission includes involvement of minority and women students in the research. We shall continue the center's tradition of mentoring undergraduate minority students and doctoral women students, and of conducting cyber security awareness workshops in high schools and middle schools. The results of this research will be transformed into publications and several course modules planned at SUNY Buffalo and University of Michigan. Dissemination of research will take place through a new project webpage at the CAE and through conferences and journals and the center's future workshops.