



(U) Update on the Knowledge System Prototype -- It's in the BAG!

FROM: [REDACTED]
Chief, Knowledge Systems, Advanced Analysis Lab (S211K)
Run Date: 04/23/2004

FROM: [REDACTED]
Chief, Knowledge Systems, Advanced Analysis Lab (S211K)

(U) We ran an [article back in September 2003](#) on the development and status of the Knowledge System Prototype (KSP), aka "the BAG." A lot has happened in the program since then, so we thought it would be a good idea to give everybody an update.

(S//SI) **Background** : The Knowledge System Prototype (KSP) was initially envisioned as a four-year advanced technology demonstration program. The initial KSP prototype was delivered in March 2002, and a second, improved prototype was delivered in March 2003. The current version of the KSP recognizes 300+ different concepts and relationships, has been linked to dozens of disparate open source and SIGINT data sets, and can rapidly correlate and display billions of instances from these data sets.

(S//SI) **"Challenge Problems" - Moving towards operations** : The second KSP prototype provided us a basic technical capability. While we are continuing technical development work, the delivery of the second prototype allowed us to start looking much harder at just how the KSP capability would be used in operations. To explore this question, we have initiated a series of "challenge problem" with Analysis & Production (S2) product lines. Specifically, we have asked A&P analysts to attempt to use the KSP to work on real analytic problems using real target data. We use what we learn from these problems to drive further development of the system. Our first "challenge problem" demonstrations were completed back in November 2003 with analysts for the Proliferation and Arms Control Office (S2G) showing how they could use the KSP to identify middleman "cut-outs" between suspected arms traffickers. This month, analysts from the Counterintelligence and HUMINT Support Office (S2D) are demonstrating their use of the KSP on a variety of analytic tasks. We have about twenty additional "challenge problems" lined up that we hope to get to over the next eighteen months.

(S//SI) **"Demonstration Deployments" - Moving further towards operations** : We are also looking at a few deeper operational applications for the KSP via what we call "demonstration deployment." The biggest of these will be an end-to-end deployment of the KSP into the National Security Operations Center (NSOC). The idea is to integrate all data sources coming into NSOC and all applications used by NSOC into the KSP to gain knowledge about how this improves NSOC's ability to monitor and control the SIGINT system. Work on this task started in January. Another demonstration deployment we are planning is with the new SIGDEV Data/Network Operations Center (DNOC). The goal is integrate data sources and applications via the KSP and then learn how this impacts DNOC's ability to understand and exploit the global network. Two additional "demonstration deployments" are underway or planned with other program. The first is with Project MAINWAY to provide a rapid telephone chaining query capability to a small group of CT analysts. The second is with the Advanced Analysis Lab's FOSSWAY Project where we are ingesting open-source data sources provided by FOSSWAY to "enrich" SIGINT. FOSSWAY will then use this enriched SIGINT to develop advanced traffic categorization and characterization tools to attack the "unknown unknowns" problem.

(S) **"Technology Thrusts"** : While we are doing a lot of interesting work with operational elements, the KSP is still a prototype and still needs to develop technically. In addition to what we learn from the challenge problems and demonstration deployments, we are also sponsoring a series of "technology thrusts" to identify, evaluate, compare, and demonstrate advanced technologies of potential value to KSP program. Current or upcoming efforts include work on entity extraction tools, entity naming, belief maintenance systems, metadata registries, graph query languages, database evaluations, conceptual graph generators, and terminology management systems.

(S//SI) **Preparing for the future** : Finally, we are planning for success by taking the steps necessary to make the KSP a robust system available to support all of operations. Specifically, we have brought in additional support to develop and implement configuration management, help desk, training, website, and system documentation. We have signed long-term maintenance agreements for key KSP hardware and software components. We have taken action to ensure all commercial hardware and software used in the KSP is on the NES-Baseline. We have arranged space in the [REDACTED] [mission assurance facility](#) . We are moving one of existing KSP systems (there are two) to the SERVERWORLD facility where it will have 24/7 ITIS support as well as an UPS connection. In conjunction with TRAILBLAZER, we just purchased our next-generation KSP hardware platform, a Silicon Graphics ALTIX 3700. The Altix is the largest memory system ever built by Silicon Graphics and, to our knowledge, is the first high-performance, large-memory, scalable system specifically tailored to support Analysis & Production missions. Not stopping there, we are already talking with industry leaders to plan moving us up to even bigger and more capable hardware platforms.

(U//FOUO) **Learn more** : For those of you interested in learning more about the KSP, we hold open briefings and demonstrations in the Advanced Analysis Lab Conference Room (2N039-1) every other Friday from 1100-1200. The next scheduled briefing is 14 May. If you are interested in attending, please contact or email me ([REDACTED]@nsa) to reserve a seat.

"(U//FOUO) SIDtoday articles may not be republished or reposted outside NSANet without the consent of S0121 ([DL sid comms](#))."

DYNAMIC PAGE -- HIGHEST POSSIBLE CLASSIFICATION IS
TOP SECRET // SI / TK // REL TO USA AUS CAN GBR NZL
DERIVED FROM: NSA/CSSM 1-52, DATED 08 JAN 2007 DECLASSIFY ON: 20320108