



Software Group

# Some Patterns of Knowledge Management in Secure Environments

Alan D. Marwick  
IBM Software Group

Fred Maymir-Ducharme  
IBM Federal Systems

# Agenda

- Patterns
- Ad-hoc Collaboration
- Text Search
- Conclusions

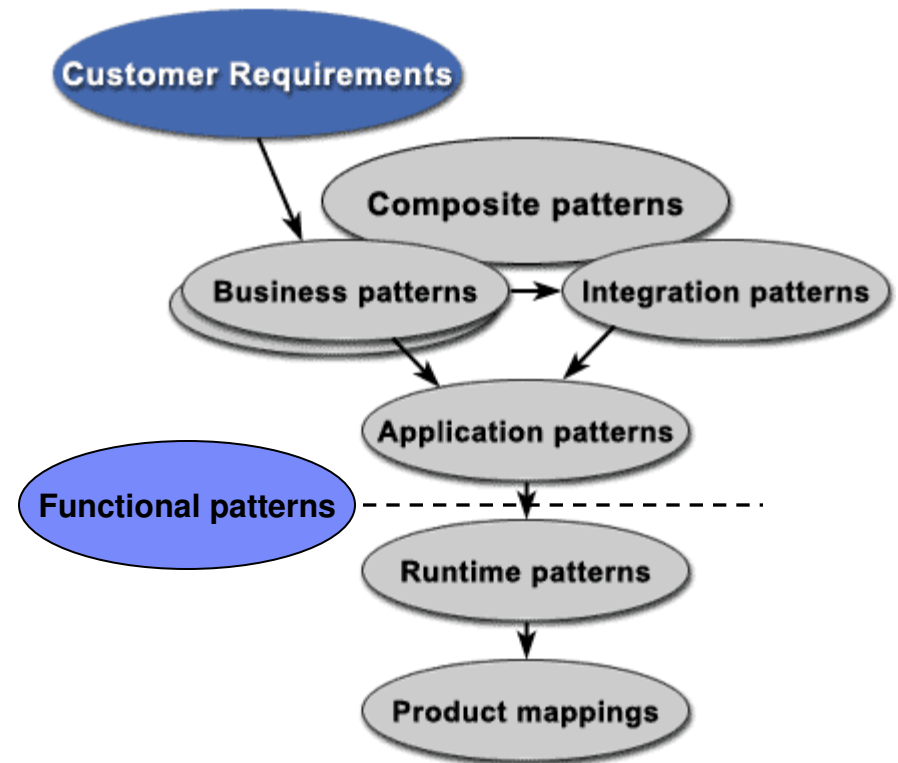
## Patterns describe elements that recur in many Knowledge Management implementations

- “A solution cannot be a pattern unless it has been found over and over again”. (L. Rising, “The Patterns Handbook”)
- Patterns provide a high-level description of solution elements, abstracted from their implementation – a language\*.
- In Knowledge Management, people are in the loop. KM Patterns describe solution elements that:
  - ▶ Promote effective collaboration,
  - ▶ Help people to use large amounts of information for analysis of problems and synthesis of understanding
- Examples are found in products, customer solutions built by IBM teams, and in research projects.

\* Christopher Alexander et al “A Pattern Language” (Oxford, 1977)

## Our Knowledge Management patterns are functional patterns, at a level of abstraction between Application patterns and Runtime patterns

- Patterns are increasingly used to describe elements of solutions:
  - ▶ Business processes
  - ▶ High level architectures
  - ▶ Application architectures
  - ▶ Software design
- KM patterns describe elements of functionality that are important to users
- Functional patterns are an addition to the pattern taxonomy proposed by Adams et al.



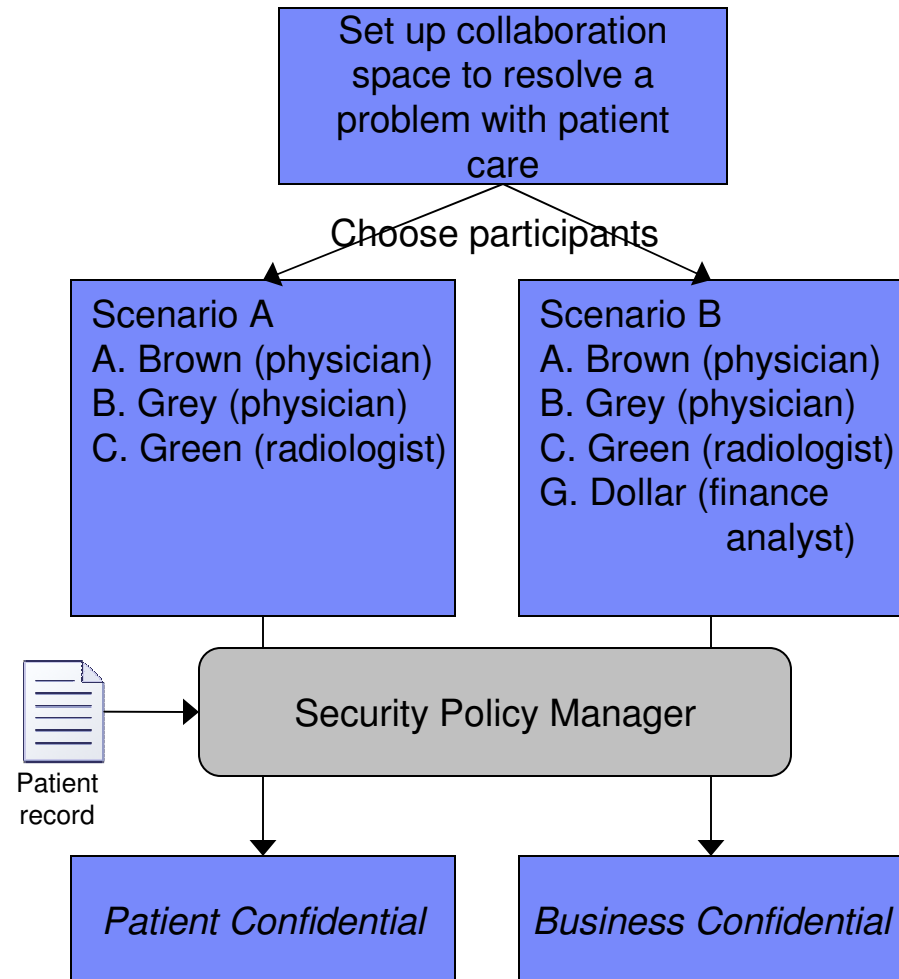
Source: "Patterns for e-business, a Strategy for Reuse", J. Adams, S. Koushik, G. Vasudeva and G. Calambos (IBM Press, 2001).

## The *Ad-hoc collaboration* pattern requires rapid setup and lightweight administrative processes

- Assemble a distributed team to rapidly solve an unanticipated problem
  - ▶ Gives access to the tacit knowledge of the team by working with the people who have it
  - ▶ Quickly provision the team with an on-line “place” for discussions, meetings and documents
- Examples:
  - ▶ Knowledge Management: need people with relevant expertise & understanding to discuss an issue and advise
  - ▶ Collaborative e-commerce: people in roles appropriate to resolve e.g. supply chain problems between companies
- Must be easy to set up – minutes, not days
  - ▶ Security admin must be easy

## Setup must be easy. Flexible policy based security allows access rights to be inferred automatically

- Goal: resolve an problem with the care of hospital patient Jane Doe
- Issue: the appropriate level of classification of the shared collaboration space depends on who is participating:
  - ▶ Patient Confidential: Allows details of patient's illness to be discussed. Only the patient's physicians can access (as determined from the patient record)
  - ▶ Business Confidential: No medical information, but billing and financial info. Physicians and admin staff can access
- Policy-based access control avoids the need to explicitly assign roles to people in other departments or organizations
- Can be implemented with standard products (Goodwin 2002)



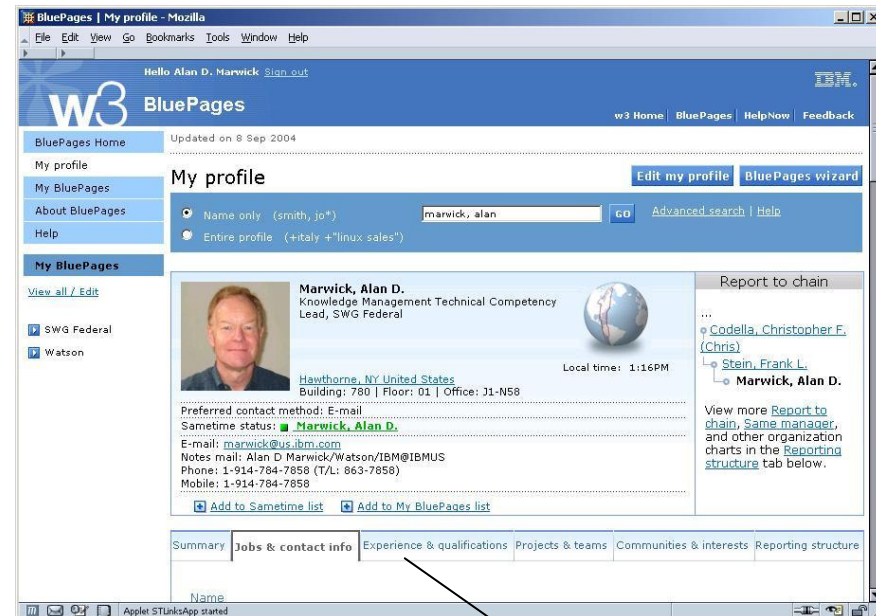
## The *See Participant Details* pattern helps people in an ad-hoc collaboration to adhere to security policies

- Where adherence to security policies cannot be completely automated, participants need knowledge of the other people involved
  - ▶ In large distributed or virtual organizations, may not know the other people involved
  - ▶ In a face-to-face, easier to understand people's roles, affiliations
  - ▶ This pattern also facilitates building of vital inter-personal trust
- Need on-line access to authoritative information about participants in a collaboration



## An augmented directory provides authoritative source of information about participants in a collaboration

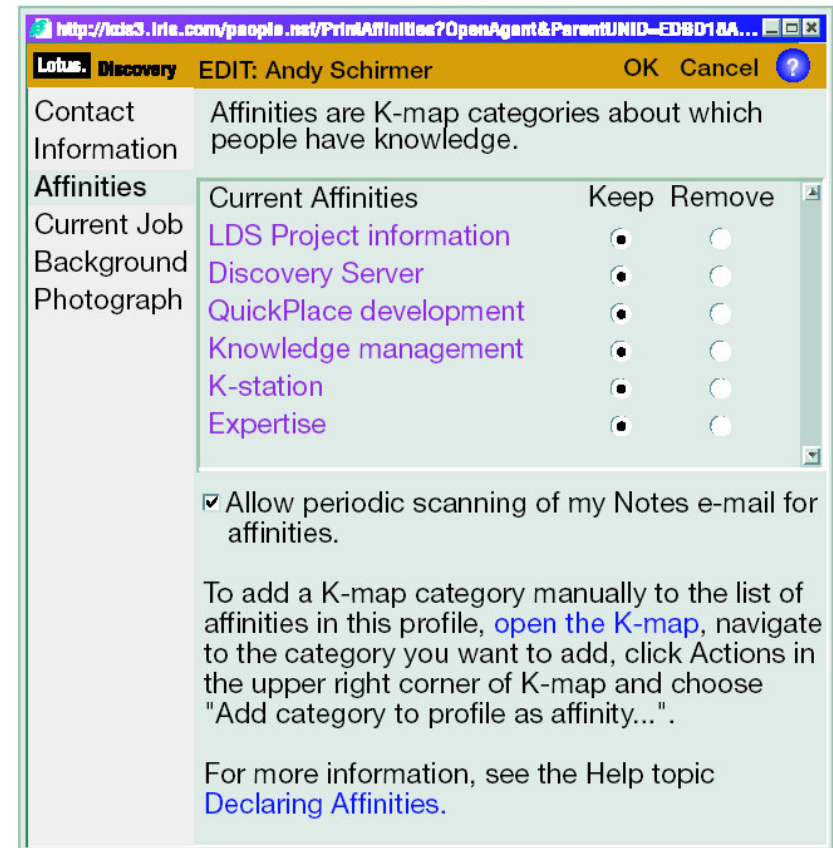
- Includes
  - ▶ Phone book info
  - ▶ Position, role
  - ▶ Regular/supplemental employee
  - ▶ Manager status
  - ▶ Potentially, clearances etc.
- Create from trusted data, and/or validate updates through a business process
- LDAP access available to applications
  - ▶ But the extra information is for *people* to use





## Expertise location with automated profile gathering requires the *Approve Expertise Profile* pattern

- To find participants with relevant knowledge and experience, expertise location features can be used
- Search an index created from either
  - ▶ Unstructured fields in employee directory
  - ▶ A profile automatically created using analysis of documents associated with person
- To ensure that the automated profile does not breach privacy or security, applications implement the *Approve Expertise Profile* pattern



Source: A.L. Schirmer "Privacy and Knowledge management: Challenges in the design of the Lotus Discovery Server" IBM Sys. J. 42 (2003) 519

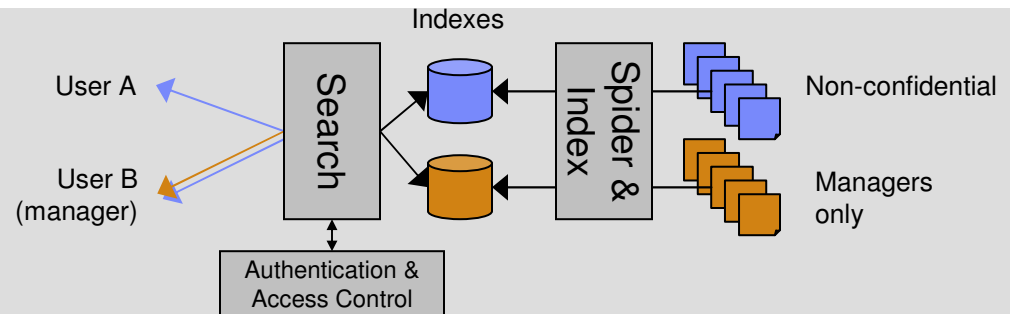
## Knowledge Management in distributed environments requires the functionality of *One Query Searches Everything*

- Basic function of KM systems; predominant paradigm for access to unstructured information.
- Search with one query gives users easy virtualized access to all available information
  - ▶ Overcomes stovepiping of information within organizations or systems
- Enterprise search is more difficult than Internet search (R. Mukherjee and J. Mao. Enterprise search: Tough stuff. *ACM Queue*, 2(2):37, 2004)
- How is security handled in search?

## Implementations of *One Query Searches Everything* can balance security, flexibility and efficiency


- Two indexes (or one per stovepipe)

- Simple, fast
- User identity determines index access
- OK if access policies are simple



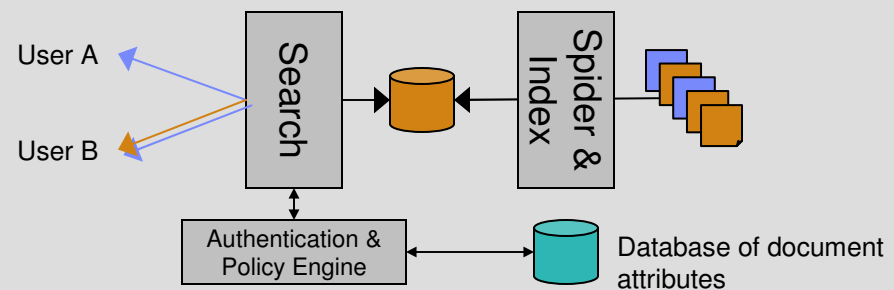
- Access info is in index (e.g. roles that can access each document)

- OK if roles seldom change
- Efficient search, as permitted roles are easily added to query
- Can filter in search engine's inner loop



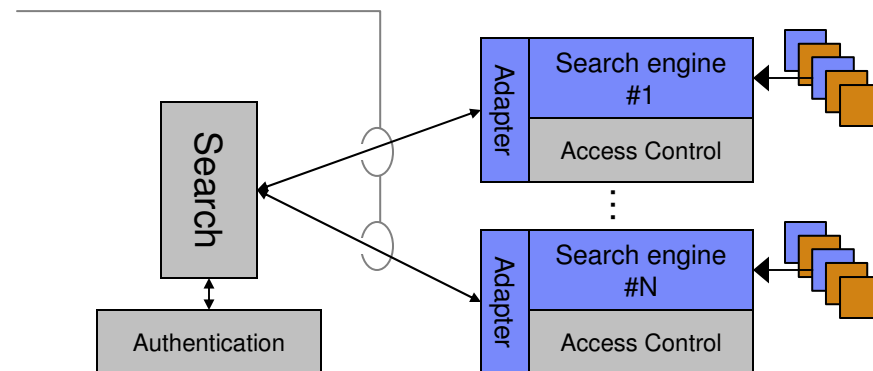
- Only document attributes are in index; access is computed at search time through application of policies

- Must check each document that satisfies query
  - Requires fast policy engine
  - Pref. integrated with search engine
- All security models can be supported



## A Distributed search implementation allows the content providers to implement their own security policies

- Search is delegated to search engines attached to remote content
  - ▶ Use a proprietary protocol, or ISO 23950 / Z39.50
- Each repository can implement its own policies
  - ▶ Even do authentication if necessary
- Disadvantages
  - ▶ Interleaving ranked results lists
  - ▶ Tends to least-common-denominator function

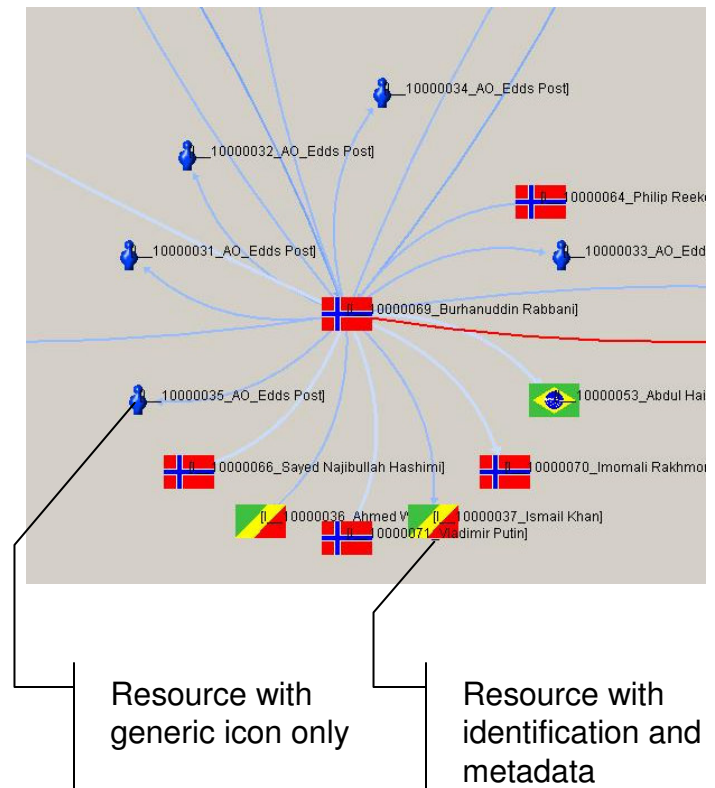


## *Inaccessible Documents are Invisible* is a common pattern applied to search, but potentially reduces effectiveness of KM

- User never sees documents s/he is not allowed to access, even if they satisfy the query
  - ▶ Easy to implement e.g. with two index scheme
  - ▶ Security by concealment
- However, user gets incomplete picture of available information
  - ▶ Could seek access, or
  - ▶ Could ask cleared team member to review the document and find if it is relevant
- Following pattern is an alternative, if security policy permits

## The pattern *See It Exists* allows a user to demonstrate need-to-know

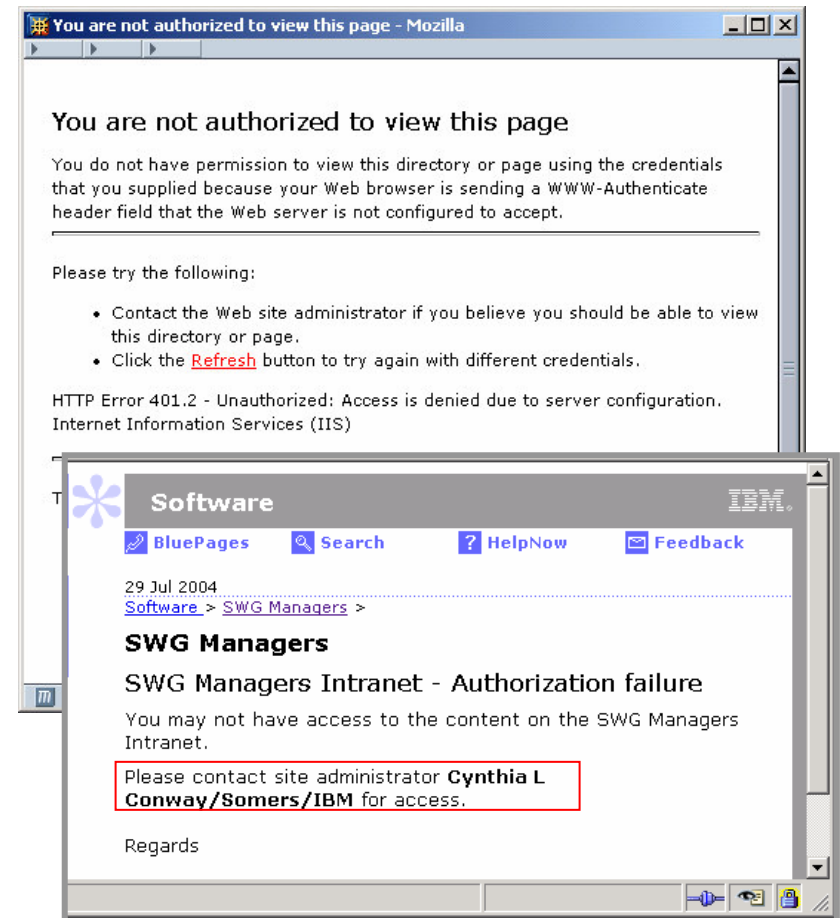
- Documents or resources that the user is not allowed to see are represented with reduced detail
  - ▶ Representation is within the users permissions
  - ▶ Could be unclassified summary
  - ▶ Uses metadata when access to the full resources is not possible
  - ▶ Where people are represented, issues are like privacy
- Allows user to seek access
- Appropriate for “need to know”
  - ▶ User gets opportunity to demonstrate k2k
- Implement with a “Discover” access right
  - ▶ As in InfoWorkSpace application (Brindley, 2000)



Source: Graham Bent. An On Demand Data and Text Mining Application based on DB2 and WebSphere. To be published.

## It is very helpful if the pattern *System Advises How to Get Access* is implemented

- It may not be obvious what a user has to do
  - ▶ Especially in a inter-organization collaboration
  - ▶ Simple approach: nominate an access officer
- By implementing this pattern, the system helps the user
  - ▶ improves efficiency and reduces frustration
- Already implemented in e-commerce systems\*
  - ▶ Can get some answers to questions:
    - Who can do action X to object Y?
    - Systems lists constraints “A manager in SWG”, or actual people
- Future systems perhaps can inference over security policies, and produce a plan that will least inconvenience the user



\* R. Goodwin, A. Raina, A Rajasekharan, W. Philip, J. Thomas, J. Nuzzo, and R. Balakumaran. Advances in Policy Based Authorization in WebSphere Commerce Business Edition. In Proceedings 5th International Conference on Electronic Commerce Research (ICECR-5), 2002. Also, R. Goodwin, private communication (2004).

## Conclusions

- For text search that is both secure and effective, close integration of the search engine with the security infrastructure is needed.
- Future systems may advise users how to get access to resources by using inferencing about security policies, and planning.
- Flexible policy-based security models, already used in e-commerce, can be applied to ad-hoc collaboration
- Still many challenges to fully support these patterns



Thank you

# Backups/drafts

## We focus on aspects of KM patterns that help to resolve the tension between the goals of knowledge sharing and security

- Tension:
  - ▶ Knowledge sharing: make potentially relevant information available for decision making and to allow people to build their tacit knowledge
  - ▶ Security: limit access to authorized people with a need to know
- Resolution:
  - ▶ Within the standard definition of Information Systems Security (Confidentiality, Integrity, Availability),
  - ▶ Availability for knowledge users subsumes the knowledge sharing goal above
  - ▶ Must be balanced against the Confidentiality goal to meet overall organizational objectives
- KM patterns provide a framework within which to discuss how the goals can be balanced when a system is designed and implemented

While policy-based access control can allow automation of some security tasks, users must appropriately label information

The screenshot shows a web browser window displaying the 'Kbase Submission Form'. The page has a blue header with the 'Kbase' logo and navigation links for 'Admin', 'Search', and 'Submit'. A left sidebar lists various communities such as 'C4ISR', 'COMBAT SYSTEMS', 'ENGINEERING', 'EDGE PROTECTION', 'JAG', 'JN J, IN', 'MAINTENANCE', 'MEDICAL', 'METOC AND NAVIGATION', 'ORONANCE', 'PERFORMEL AND SUSTA', 'PLANNING AND READINESS', 'SECURITY', and 'SUPPLY AND LOGISTICS'. The main form area contains several fields: 'Title' (Air Intercept Controller Debriefing), 'Submission Date' (August 20, 2002), 'Submission Type' (Tools), 'Classification' (SECRET), and 'Releasability' (US, AUS, CAN, UK). There are also dropdown menus for 'Community of Interest' (Supply and Logistics) and 'Approval Status' (Unapproved). A 'User Ratings' section shows 'Not yet rated'. Below this is a 'POC Information' section with fields for 'Last Name' (Alec), 'First Name' (James), 'Rank' (CIV), 'Location' (DAB), and 'E-mail'. A 'Brief Description' section contains the text 'ATC Debriefing Database in CAS architecture'. The browser's status bar at the bottom shows 'Done' and 'Internet'.

# Parking

