

Real-Time, Mission-Critical Business Intelligence: Lessons from the Military and Intelligence Community

By

Alan R. Simon

Introduction

For most organizations, the quest for *business intelligence* (BI) has been closely tied with early 1990s-style *data warehousing*: batch-oriented data flows from source systems into a single data warehouse, from which individuals run reports and analyses of interest to their respective job functions.

Since the late 1990s, some “early adopter” organizations have dabbled in real-time business intelligence for parts of their respective enterprises, but the results have been decidedly mixed. On the one hand, end users now have access to critical pieces of data much more rapidly than before as a result of the real-time data flows, and – in theory at least – are now in a much better position to make critical business decisions in a more timely manner.

The reality, however, is that many of these real-time BI environments do little else other than speed up the flow of data from the point of entry into the enterprise until available for reports and analysis...but with little or no actual benefit to decision making, business process monitoring and management, organizational productivity, and all the other laudable goals which drove their BI initiatives in the first place.

Real-time, mission-critical business intelligence requires much more than simply real-time enabling data flows between source systems and the target data warehouses and analytical environments. Indeed, an entirely new way of thinking about what one wants to achieve is required to help avoid the missteps and disappointments that so many others have encountered in their real-time BI initiatives.

As it turns out, one needs only to look towards the military and intelligence communities and the ways in which their command, control, communications, and intelligence (C³I) or command, control, communications, computers, and intelligence (C⁴I) systems are constructed. This chapter looks at best practices in creating a hybrid architecture in which the advantages of “classical” business intelligence environments are blended with those of C³I/ C⁴I systems to provide a solid foundation for mission-critical BI. It also looks beyond technology and architecture fusion at some key operations principles from military and intelligence systems that have rarely been part of commercial or civilian government BI environments but, upon closer examination, should be key objectives of any mission-critical BI initiative.

Definitions

No doubt the reader has come across numerous definitions for the term “business intelligence.” The BI definition I have always preferred and used goes as follows:

The pursuit of timely, actionable, high-value insight.

Each of three descriptors in the above phrase – timely, actionable, and high-value – is equally important regardless of what one’s specific objectives might be. If the sought-after insight isn’t timely, it is of marginal use at best...perhaps even useless. Similarly, very timely *low-value* insight isn’t likely to be of much use nor worth the cost of building a complex system, either.

Then we have the word in the middle – actionable – that provides a nice segue into the body of the material in this chapter. Anyone who has watched cable or network news in the past few years and followed the search for and capture of a particular deposed Iraqi dictator or the search for high-profile terrorists, or most anything else dealing with military matters and combat operations, has heard the phrase *actionable intelligence* used. The premise of actionable intelligence is exactly as the words themselves read: gathering intelligence not for the sake of the intelligence information itself, but specifically to be able to take precise action based on that information and insight.

Therefore, as we proceed with looking at how one can apply principles from the C³I/ C⁴I realm to mission-critical business intelligence, it’s important to keep the three phrases in my preferred definition in mind and, in particular, to remember that unless one plans to actually do something with the insight gained – that is, take action – and architects the environment to achieve and enable that goal, the result will likely be just one more exercise in futility.

A Tale of Two Architectures

Figure 1 below shows a high-level conceptual architecture diagram of the type that the reader has most likely seen dozens, perhaps even hundreds, of times: a “classical” data warehousing/business intelligence environment built on a foundation of batch-based integration of information from multiple sources, with subsequent “spin-off” of portions of the newly consolidated and integrated data into more tightly focused *data marts* – geographically-based, in this particular example.

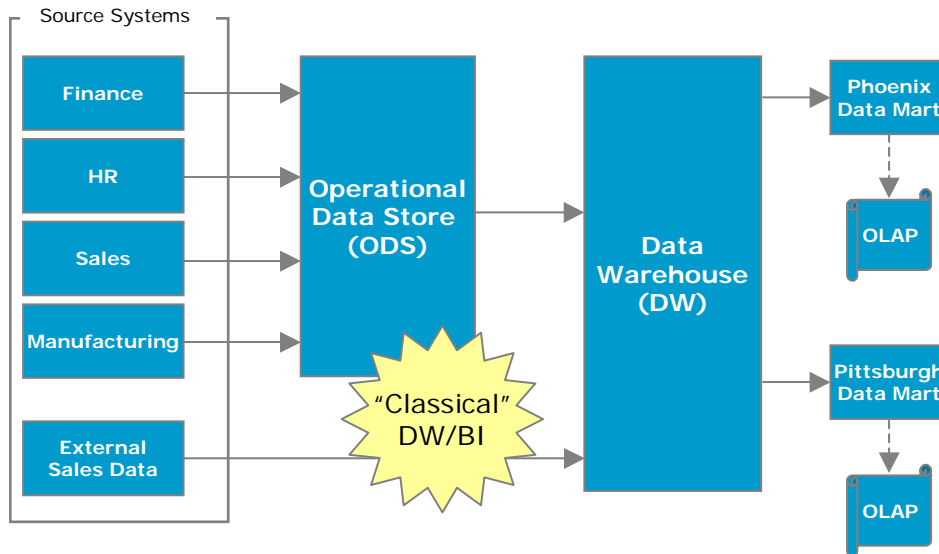


Figure 1 – “Classical” Data Warehousing and Business Intelligence High-Level Architecture

Even though one might draw the conclusion from my dismissive remarks at the outset of this chapter that I see little or no value in business intelligence when implemented as shown in Figure 1, that conclusion would be an overstatement. In fact, this early 1990s-approach to BI does have its merits, such as:

- The ability to integrate, consolidate, and organize very large volumes of data and to make that data readily available to a large number of users for analysis
- The ability to summarize the data at multiple levels and for users, during their analysis, to “drill down” and “drill up” as desired to various levels of granularity as dictated by the needs of the moment
- Equipping users with graphically oriented tools that are relatively easy to use...once users have been properly trained, of course

At the same time, a look inside the classical approach to business intelligence shows some glaring weaknesses that have tripped up almost every BI practitioner at one time or another:

- A lack of confidence in the results one receives from a particular report or analytical function; “Can I base a critical, bet-my-job decision on the answer the system is giving me?”
- A reliance on structured data – numbers, dates, and character strings – with inadequate linkage (or, more commonly, no linkage at all) to unstructured information found in compound documents and multimedia
- A dominance of individual analysis with little or no collaborative...and when collaborative work processes do exist they are almost always “kludged” together in a highly manual manner, rather than on any kind of workflow or collaborative computing infrastructure

- An operating model that essentially leaves users to their own devices and assumptions for linking the results of reports and analyses with their job functions, rather than actively guiding their usage of the environment

The approach many architects and planners have taken to date in the pursuit of real-time, mission-critical BI has been to simply transition most or all of the batch data feeds to real-time replacements. But regardless of the specific technology used for the real-time data flows – message queuing; an enterprise application integration (EAI) tool; the real-time capabilities of a data warehousing extraction, transformation, and loading (ETL) tool; or even “brute force” database stored procedures and triggers for database-to-database transfers – what most of those architects and planners wind up with is faster transfer of data to and loading the data into the target database...but little or no help with the critical shortcomings listed above.

To address these procedural and operating model shortcomings, one needs to take a look at the way the military and intelligence community addresses their own mission-critical needs of air defense, missile attack warning, space defense, and intelligence-based combat operations. Figure 2 illustrates the high-level conceptual architecture common to many military and intelligence systems that are charged with a mission of warning and defense.

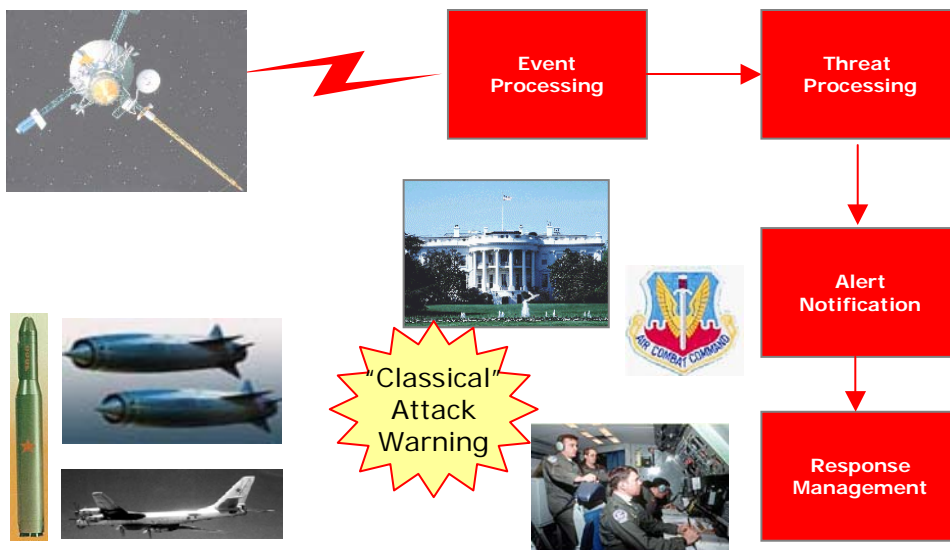


Figure 2 - C³I/ C⁴I Conceptual Architecture for Military Warning and Defense Missions

The architecture illustrated above is not only based on a foundation of “ultra-real-time” data transmission from sensors (e.g., satellites, radars, etc.) and similar sources into “the environment” but also:

- Determining if the just-received information denotes a threat or possible threat that must immediately be addressed

- Putting certain high-ranking authorities “on notice” that a threat determination is forthcoming and their response may or may not be required, depending on the outcome of that threat determination
- “Forcing” all those involved in either monitoring automated functions or interjecting human decision-making into the situation to follow their work processes to a conclusive determination and, if applicable, decisive action

Still, many C³I/ C⁴I systems have some shortcomings as compared with traditional business intelligence environments we looked at a moment ago:

- Data volumes tend to be “modest” rather than the multi-terabytes of information increasingly found in data warehouses that support BI functionality
- Many C³I/ C⁴I are built around “fixed-screen” user interfaces with little or no variability dynamically adjustable by users to specific needs of the moment
- Data tends to be “flat” rather than dimensionally structured along multiple hierarchies with various levels of summarization, as would be found in the typical data warehouse
- C³I/ C⁴I environments that aren’t hindered by the above shortcomings – typically those for the intelligence community rather than military attack warning and response systems – have their own issues such as information overload

In many ways, the strengths of data warehousing-based BI are weaknesses in C³I/ C⁴I systems, and vice versa. The question must then be asked: is it possible to create a hybrid architecture that combines the best aspects of data warehousing-based BI with those of C³I/ C⁴I warning and defense systems? The answer: absolutely! Figure 3 shows the result.

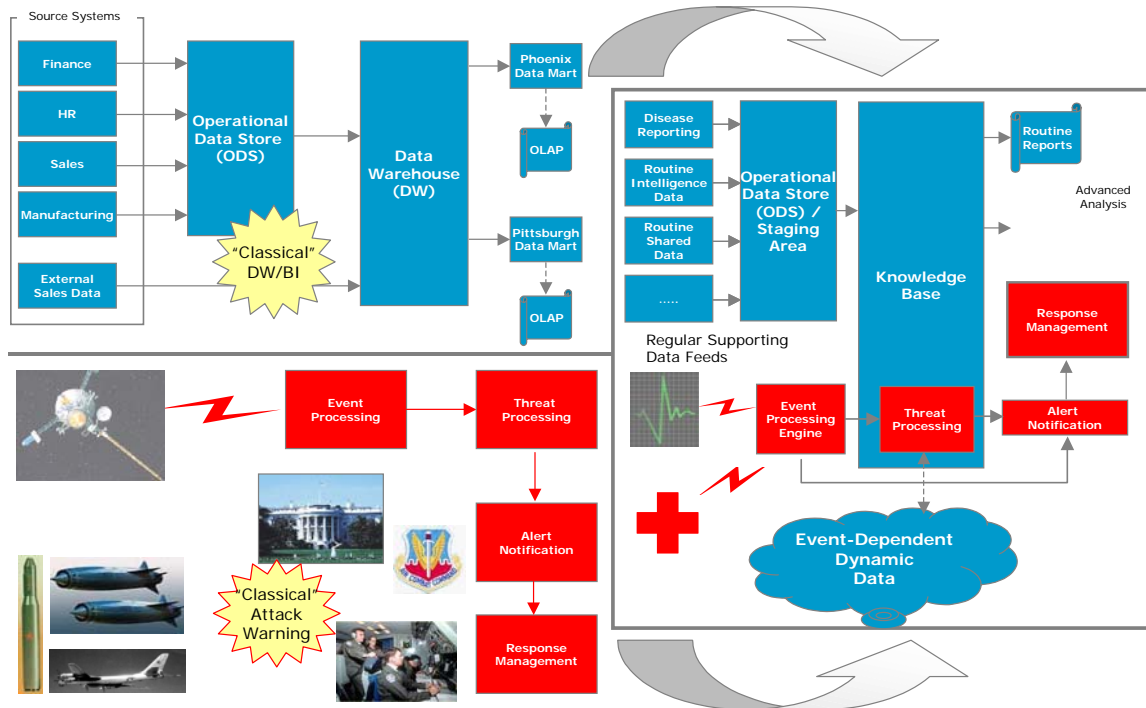


Figure 3 – Hybrid Architecture Suitable for Real-Time Mission-Critical BI

The architecture illustrated in Figure 3 was specifically created for military and homeland security “BI”-style functionality such as bioterrorism, law enforcement, and similar missions. It is, however, just as applicable to commercial and civilian government business intelligence that in those respective environments are absolutely mission-critical. Essentially, almost any problem to which BI planners and strategists have applied real-time data flows in the pursuit of real-time business intelligence would likely be better served by the hybrid architecture shown above than by simply replacing batch data flows with real-time counterparts.

Beyond Architecture: Key Principles

Achieving real-time, mission-critical BI is more than a matter of fusing two different architectural approaches into a new one similar to that illustrated in Figure 3. When architecting the *operating model* of the environment, one must keep a number of key principles “front and center” (in the spirit of the military side of this discussion!), and infuse most or all of them into the resulting system. These key principles are:

- Events and event management
- Hypothesis formulation
- Correlation
- Interdiction
- Definitive Conclusions

- Collaboration
- Outbound information flows
- Fail-safe functionality replication
- Well-defined roles and responsibilities
- Planning for personnel succession

Events and Event Management

Many mission-critical military or intelligence systems are built around the concept of *events*: an occurrence of some type that must be immediately addressed, with potentially (though not definitely) dire consequences for failing to do so. Examples of events include:

- A missile launch picked up by a “first detection” sensor
- The sudden appearance of several unidentified aircraft on the perimeter of sovereign airspace
- A person under surveillance taking specific actions categorized as “significant”

The premise behind events is that they must be managed, i.e., not “lost in the shuffle” or set aside to be picked up at some later point at which time it might be too late to take appropriate action. Further, the system components (communications links, protocols, screening and filtering software, etc.) dedicated to receiving and doing the initial pre-processing of events are architected and engineered to get the applicable information to some sort of threat processing (or equivalent) algorithms as quickly as possible.

Applying this concept to commercial or civilian government mission-critical BI is very much a situational matter, with the types of events, how they’re handled, and other properties being very much dictated by the specific mission(s) that are driving the development of the environment in the first place. The lesson that does pervade every single mission-critical BI development effort, though, is that simply having a collection of real-time source-to-target data interfaces is not enough; the system must be architected and engineered to take full advantage of those real-time data flows in an almost fanatic pursuit of (referring back to my preferred definition of business intelligence) timely, actionable, high-value insight.

Hypothesis Formulation

Traditional BI environments typically deal in cold, hard facts: for example, last quarter’s revenue produced by each sales person in each of the company’s regions with various levels of aggregation and “slicing” of information available upon demand. Or, for a public sector organization such as a city-level public welfare department, a complete report of all disbursements organized and categorized by neighborhood cross-referenced with information drawn from state-level job training programs.

C³I/ C⁴I systems must also deal in cold, hard facts but before those facts are determined, many of these systems formulate *hypotheses* based on the first fragments of information.

Consider a military system responsible for detecting inter-continental ballistic missile (ICBM) launches and determining whether or not that launch is a threat to the United States or its allies. As discussed next, making an “official” threat/no threat call may require the consolidation and correlation of many different data points from multiple sensors. However, upon the first detection of a launch, the system may make a “best guess” – that is, formulate and report a hypothesis – of whether or not the launch is likely to be a threat. The system may factor in data points such as the missile’s initial trajectory; if it appears headed towards the other side’s regular testing target zone, in the opposite direction from the U.S. mainland or the territory of an ally, then the initial hypothesis may be that the launch is likely to *not* be a threat.

Note, however, that the word “hypothesis” was selected for a reason. Just like in geometry, where an initial hypothesis is essentially little more than an educated guess until proven correct, a warning system’s initial hypothesis is reported to the users *but processing and tracking continues until the “proof” occurs*.

Applying this concept to real-time, mission-critical BI, one might architect an environment in which a system not only receives real-time data inputs but passes the newly received data through a series of filters and algorithms looking for an early indication of a potential problem – not a definitive indication of a problem but rather a first hint that something *might* be amiss...and absolutely must be monitored closely until it can be determined with absolute certainty whether or not a problem actually does exist.

Correlation

Most BI environments provide a wealth of information to users, but within the portfolio of available reports and analyses, it is rare to find any two that work cooperatively with one another. Many reports may be available for supply chain performance and even more analyses available for customer satisfaction, but it is inevitably left up to the user to make the linkage between the two functionality areas and determine, for example, whether or not supply chain problems materially affect customer satisfaction and subsequent buying patterns.

In the military and intelligence world, however, *correlation* is often an integral part of a system’s functionality and the underlying architecture. Consider the missile warning example above. Suppose that several sensors are solely responsible for detecting launches, while other sensors are primarily responsible for detecting a missile’s trajectory based on factors such as altitude, direction, speed, and other data points. Essentially, a single missile launch may be detected by different sensors at different points in time.

A missile warning system must therefore correlate all of its data inputs using algorithms that are “aware” of the characteristics and missions of the sensors providing information into the system. The reason? It makes a big difference to folks such as four-star generals and the President of the United States whether a barrage of inputs may actually indicate only one or two missiles that in fact aren’t headed in our direction as contrasted with an

erroneous (and potentially catastrophic) indication that dozens of missiles are flying, and nobody is quite sure in what direction they're headed.

BI strategists and planners should look for opportunities to correlate different portions of the overall base of information in the quest for high-value, actionable insight.

Interdiction

Most BI systems built around batch processing do provide insight to their users...after-the-fact, “tell me what *did* happen” insight. Too often, this insight is neither timely nor actionable, and the users can do little more but plan workarounds, recovery plans, and other reactive steps.

In contrast, think about a warning system responsible for air defense. Detecting incoming enemy aircraft penetrating North American airspace is one thing; doing something about that penetration – sending intercept aircraft or preparing to launch air defense missiles is typically called for in these situations. Detecting a potentially problematic situation isn't enough; what's called for is *interdicting* the situation by taking appropriate action in time to positively influence the outcome (for our side, at least).

How might interdiction be accomplished in a mission-critical BI environment? Perhaps by triggering an automated feed of critical data into an enterprise's transactional systems for immediate action, or by persistently alerting and updating appropriate senior-level executives of the nature of the problem along with recommended solutions based on the system's knowledge base and embedded business rules.

Definitive Conclusions

In the typical BI environment, information is delivered to users in the forms of query results, reports, or through some other vehicle...and then it's often up to those information recipients to draw their own conclusions from the information they receive. Certainly, those users typically receive hard facts, as discussed earlier: actual sales dollar volume; performance by all the organization's sales people; changes in the market value of real estate property “owned” by a particular governmental or quasi-governmental organization; upward or downward trends in the number of problem properties on the roles of a government agency responsible for subsidized housing; and so on.

But despite the millions of discrete facts delivered every single day to millions of users of BI systems, almost every one of those users would be hard pressed to answer, with absolute confidence, the following question: “So what does it all mean?”

Consider, in contrast, a military or intelligence system that relates to decision makers definitive conclusions such as:

- “Yes, a missile was launched from a particular site in Russia, but it was a test launch and is not a threat to the U.S. or its allies”

- “Yes, an unidentified aircraft suddenly appeared on the fringes of North American airspace, but based on further identifying data it was determined that this aircraft was an international commercial flight that had gone off-course due to poor visibility and possible instrumentation failure...but regardless of the cause, it’s not a threat.”
- “(Person X) who has been under surveillance by the intelligence community for the past 18 months has been gathering documents of the type that it is a virtual certainty he/she will try to enter the United States under a false identity, and will try to do so within the next 15 days.”

Mission-critical BI environments often stop short of getting to the point where the facts they produce can be drawn together and, based on how those facts all fit in the context of one another, point to a definitive conclusion. Absent this final step, users often find themselves in “analysis paralysis” where they do have all the information they need, but they are unable to perform the proper correlation and synthesis to get to the point where they are absolutely certain of what the situation is and what actions to take.

By equipping these environments with the procedural logic and other facilities to get to this final stage, the value of BI environments will be far, far greater than we typically see today.

Collaboration

BI environments can be lonely affairs, despite being architected for hundreds or even thousands of users. Why? Because for the most part, a single user issues queries or information requests against this vast wealth of information, and the response comes back to that user...and nobody else.

Certainly, a user could e-mail report results to a co-worker with a message asking for that person’s thoughts about the results, or even call that co-worker over to his or her PC screen to discuss what the report is showing. But doing so is very much a reactive action, done only on the initiative of the first user. Should that user decide to “go it alone” there is no system-managed facility that would otherwise require that person to collaborate with anyone else, even if doing so would bring tremendous value to the analysis being undertaken.

In the C³I/ C⁴I world, however, collaborative processing is usually built into the support systems as well as the work processes associated with that particular mission. Actually, “built into” is an understatement: collaboration is an integral, immutable aspect of these systems and their mission. Operators in the missile warning center in Cheyenne Mountain, Colorado work in concert with other operators elsewhere in “The Mountain,” who also work in concert with operators at other places: Offutt Air Force Base in Nebraska, the Pentagon, and so on. And somewhere in the chain are codified communications processes with generals and civilian command authorities, all of whom have a well-defined, well-rehearsed role in the mission.

In the commercial world, workflow engines and collaborative “groupware” have long been part of transactional applications such as insurance claims processing. Why not apply those principles from the transactional world – commercial or military – to analytical environments?

Outbound Information Flows

Not to “pile on” the shortcomings of traditional BI systems, but here’s another one: the feeding of information from sources into a data warehouse or equivalent analytical environment is often an end-point: that information that has undergone such rigorous quality assurance and been synthesized together is available for the users of that environment...but nobody else!

In contrast, consider a missile warning system that synthesizes data from numerous sensors and, even when nothing is happening, launch-wise (thankfully!), regularly sends a status report to other critical links in the national defense system indicating such facts as the up/down state of each sensor, number of protocol errors in the past 24 hours, and similar facts. Essentially, a system that integrates data from many different sources (sounds like a data warehouse, right?) turns around and acts as a source system to some other downstream consolidator of information.

Most everyone can think of dozens of examples in his or her organization where there is a tremendous desire – an unfulfilled desire – to automatically feed selected outputs from a particular BI system to one or more environments elsewhere in the enterprise, to the benefit of everyone. Look for these opportunities; you won’t be sorry.

Fail-Safe Functionality Replication

One of the primary mantras of data warehousing has long been the quest for “a single version of the truth.” A laudable goal, no doubt; but I would argue that the typical data warehouse is more likely to be “a single version of an answer, and that answer might – or might not – be correct.” A better approach, I’ve always felt, is to have planned redundancy of selected key data managed by several components in various source systems and analytical environments throughout the enterprise, with codified, formalized, ever-present QA procedures running in the infrastructure that compare totals that should be same, do spot checks of data values that should be identical, and so on.

Planned redundancy is a long-standing concept in the C³I/ C⁴I world: multiple missile warning systems, multiple air defense systems, multiple intelligence systems. The secret, however, is that the redundancy must not only be planned, but actively (and correctly) managed. Having multiple intelligence systems with overlapping “watch lists” that have conflicting data is extremely detrimental and potentially catastrophic.

Remember this theme – “controlled replication is good; uncontrolled duplication is bad” – and you won’t go wrong.

Well-Defined Roles and Responsibilities

So far, the good news/bad news points made have dealt in matters of technology and systems architecture. Beyond these obviously critical pieces of the puzzle, one must also look at the personnel side of mission-critical BI. Unfortunately, many BI environments are structured in the same manner as a 1970s or 1980s-era class registration day at a 30,000-student public university...one big free-for-all in which lots of people line up for certain classes but other equally important classes are all but ignored.

OK, maybe it's not the best analogy, but the point is this: without some order and discipline as to which users should work with which data, a very real risk exists that certain segments of the BI environment will be over-served as a result of duplicate efforts by users in different organizations – or perhaps even the same organization – while other critical insights available go untapped because of the “role gap” situation.

As noted earlier, mission-critical military and intelligence systems are usually hallmarked by users working together collaboratively. However, it's more than just users working together: it's how they work together. The same principle should apply to BI environments: well-defined, precisely orchestrated roles and responsibilities for all involved to help ensure that the maximum value and insight is obtained from the synthesized data.

Planning for Personnel Succession

One final point to make also deals with the user community of a BI environment. Think about how many times a person disappears from an organization because of resignation, involuntary termination, job transfer, promotion, or some other reason...and the information that person used on a daily basis as part of his or her job suddenly is “lost” because that person's replacement doesn't work with the data the same way, or doesn't understand the nuances of the data and what it all means.

In contrast, military missions are fundamentally grounded in a never-ending turnover of personnel. True, some roles are served by civilian employees who may stay in their roles for ten, fifteen, or twenty years, but for the most part, critical operations roles are filled by enlisted personnel and officers for periods of one to three years, with succession an absolute certainty.

Succession planning involves extensive training in all aspects of a job: the work processes, the information, the interrelationships, even what to do in various crisis situations. These facets are, in the military world, drilled into the very fabric of newcomers through rigid training and even formal certification.

Why not apply these same principles to mission-critical BI environments? Newcomers should receive more than a user ID, a password, and a quick tour through the environment's online help facility. Planners and executives should assume that people coming and going will occur, and this must be factored into the environment to avoid

compromising the mission. Again, it takes more than just data and how it's delivered to users to create a viable mission-critical BI environment.

Conclusions

At the time of this writing (late 2004), we're about fifteen years into the data warehousing/business intelligence "modern era" that began in the 1989-1990 timeframe. On the one hand, these disciplines have certainly proven themselves to be more than passing fads like so many now-faded technology "movements." Still, as noted at the outset of this chapter, too many of these environments do little other than produce a lot of reports but very little timely, actionable, high-value insight.

One can look towards the mission-critical systems run by the military and intelligence communities for some critical hints and assistance on how to take the next step in the BI realm to enable the same degree of "mission-criticalness" (to coin a phrase) in how we use our terabytes of data for much more than most of us do today.