


National BILT Meeting Minutes
“AI Literacy” Vote and Discussion

MEETING DATE: Tuesday, August 25, 2026	MEETING TIME: 11:30am-1:00pm Eastern	MEETING PLACE: Zoom
RECORDER: Mark Dempsey	RECORDING: Available upon request	PREVIOUS MEETING: Soft skills vote meeting – June 8, 2026

MEMBERS PRESENT

BILT:		
Scott Andersen	Corey Kirkendoll, 5K Technical Services	YASH Technologies
College of the Desert	Durga Krishnamoorthy, Cognizant Technology	Kim Yohannan, MongoDB
Bob Hitchins, Fidelity National Financial	Lee Rosenfeld, McGraw Hill	Teresa Younkin, Mosaic Life Tech
Hye Tech Network & Security Solutions	Snohomish County	
Jess Irwin	Srinivas Sandiri	
EDUCATORS: College of Southern Nevada, Columbus State Community College, Delta College, Forsyth Technical, Frederick Community College, Lansing Community College, Merritt College, Polk State College, St. Petersburg College		
NITIC STAFF: Mark Dempsey, Ann Beheler, Andie Bonkowske, Diane Meza, Alie Hernandez, Larry McWherter, Stephanie Schuler, Christina Titus, Debbie Hecht, Grayson McKeown		

Agenda items	Discussion
NITIC and BILT overview	Larry welcomed the group, previewed the agenda, provided some “rules of engagement” for the Zoom call, then gave a quick overview of the NITIC grant. He explained that the goal of today’s meeting is to prioritize AI literacy skills for entry-level IT technicians in the next 12-36 months. Ann provided a quick overview of the BILT (Business and Industry Leadership Teams) because all of the BILT members on today’s call have attended past meetings. She explained that today’s discussion is on general AI literacy for a generic IT major, not for someone pursuing an AI specialist job. Ann also explained some of the work she’s been doing with colleges across the country on AI skills. She outlined also where this list of knowledge and skills came from – ▪ KSAs for the 5 foundational areas of the DOL AI Literacy Framework (see below) ▪ Other KSAs that other colleges have used for AI literacy for non-AI professionals ▪ A select group of AI literacy knowledge and skills for those majoring in an IT field



Ann next explained the 1-4 voting scale. Each item will receive its own individual vote.

- * “4” means absolutely must be in the curriculum
- * “3” means should be in curriculum
- * “2” means it’s a “nice to have”
- * “1” essentially means delete

Using a QR code (and URL in the chatbox), the SMEs voted electronically.

She reminded the group again that the BILT is considering literacy for the non-AI specialist.

Anything scoring 2.6 or less will be discussed because, typically, items with scores that low are left out of the curriculum. But the BILT is welcome to discuss any item.

Knowledge

KS - Knowledge of the roles of training data, algorithms, goals, parameters, and human design choices in shaping AI behavior, 2.38 average

One employer liked the idea of “training data” but that may be too advanced. He thought that concept could maybe be added to K4. Another employer asked if the idea was here was to set up an AI course. Ann said NITIC isn’t trying to tell community college how to handle this – some might create a course; others might create a workshop. She noted that many schools are pushing for all students to be AI literate. Anyone can use an LLM, so it’s important to teach students how to use LLMs responsibly. That employer replied that while a lot of the KS list should be required to know, he’s not sure all of it needs to be in a course. He’s noticed hiring managers are asking applicants to demonstrate how they’ve used AI to improve performance and productivity. Soon, everyone’s job is going to have AI integrated into it. He argued that AI shouldn’t be an independent course but threaded through every course. When a student graduates, he/she will know how to use AI to get what they need in any situation or workflow.

Another employer agreed. Before, you had to read and write to get a degree but now you must demonstrate the fundamentals of AI – not how to train it, but how to use it. He noted two categories to the KS list: the specialist who has to implement AI and the employee who must use it. He tried to emphasize that in his votes.

Ann explained the origin of the KS items: the first few are for the user, then next ones are more focused on the specialist. It can be a fine line.

Another employer agreed that she rated the more technical KSs on the lower side where employees are using AI for job tasks and don't need to understand the technology. Ann agreed that basic AI literacy is likely important across all disciplines, not just IT disciplines. Something else the BILT needs to consider is how often these lists should be updated, which impacts how often AI courses and workshops should be updated. This is a challenge in academia where revisions and updates can take time.

K16 - Knowledge of the importance of corroborating AI-generated information with authoritative, current, and independent sources, 3.46 average.

One employer noted that most people struggle with this skill set so it's hard to expect a college student to do well with K16. Students should definitely learn how to look for independent sources, but he dislikes the word "authoritative" because AI tools and skills will change. Another employer agreed it's hard to find truly authoritative sources. Ann wondered about "human in the loop" – do humans still need to be involved? The BILT seemed to agree that AI is a useful tool, and humans need to stay engaged to use the tool correctly. Ann wonders if humans will always be in the loop, but that discussion can come later.

ACTION: remove the word "authoritative" from K16.

K20 - Knowledge of the importance of institutional, course, workplace, legal, and tool-specific policies governing AI use, 3.77 average.

One employer noted that when it comes to AI policies, students need to understand that in the absence of federal rules, a lot of states are stepping in to create legislation around AI use. Another employer agreed that K20 needs "federal" added. Ann wondered about adding "global" so we're not focusing just on American policies. Another employer agreed that students should be aware of regulations at all levels, including business-specific rules. The policies will be different depending on where you are and what you're trying to do. He wondered about combining this one with K16.

ACTION: clarify that K20 includes not just organization-specific rules but also state, federal, and global regulations.

K25 - Knowledge of the need to update AI knowledge as tools, capabilities, policies, and occupational applications change, 3.08 average.

One employer suggested that it be rewritten – it's not "AI knowledge" you're updating, it's "knowledge of AI." He noted that his company updates AI tools and services monthly because there are so many changes happening. He wondered about making it "general knowledge of AI." The BILT did not object to those changes. Ann noted again with the technology changing so rapidly, it is hard to create and maintain curriculum.

ACTION: consider rewriting K25 to read "Knowledge of the need to update the general knowledge of AI."

K27 - Knowledge of why digitization, usable data, and baseline digital literacy are prerequisites for effective AI use, 2.62 average.

One employer had a concern with "usable data." What is "usable data"? That seemed vague. He was okay with the "baseline digital literacy" phrase but disagreed with the statement that

these are all “prerequisites” for effective use. The more models you use, the better you will get. You don’t need to know anything to start. Another employer noted that he would expect entry-level workers to know about “enterprise digitization, maturity, or data readiness.” He would only expect them to know that “AI depends on good digital data.” Poor data or manual processes can limit AI effectiveness. He preferred focusing just on general literacy.

ACTION: consider removing from K27 the phrase about these being prerequisites.

K28 - Knowledge of how to distinguish rule-based automation, generative AI assistance, and agentic systems that can plan or act toward a goal, 2.62 average.

One employer thought this was too broad. He understands the need to distinguish between all of that, but it may go too deep. He suggested removing mention of agentic systems and sticking to just understanding automation and generative AI assistance. Another employer found the definition of these terms unclear. People may think they’re using an agent when they’re really using generative AI. Students need clear, accurate definitions of these terms and concepts, though she admits most employers won’t care so long as they get the answers and analysis they need. This may be more an “awareness” item and less a “knowledge” item. Another employer agreed K28 is better for an IT administrator to know which system provides which service. It’s often not a choice the end-user makes. Ann reminded the BILT that the focus of this discussion is on AI literacy for all IT majors, not necessarily just IT admins. That employer wondered if there should be a voting list for AI experts and one for AI users.

Another employer wondered about making the agentic portion of K28 into its own new K item. It should be captured. But he noted for now it’s “different enough” that many entry-level IT people may not know about it, nor would they know how or where to find it.

ACTION: consider emphasizing that K28 is more awareness than knowledge.

ACTION: consider splitting the “agentic systems” element off K28 and creating a new K item.

K33 - Knowledge of how to distinguish artificial intelligence, machine learning, deep learning, generative AI, narrow AI, and major learning approaches at a conceptual level, 2.46 average.

K34 - Knowledge of how digitization, structured and unstructured data, data quality, and data availability enable or constrain AI-enabled IT services, 2.69 average.

K35 - Knowledge of, at a systems level, how models are trained, validated, and used for inference, including generative, multimodal, batch, real-time, cloud, and edge use, 2.15 average.

One employer found K33 and K35 were very similar, while K34 may go deeper than needed. He proposed consolidating all of them into a single high-level, generic K item that focused on “demonstrated understanding” of AI. No one on the BILT objected.

ACTION: consider combining K33, K34, and K35 into a single, higher-level K item.

K37 Knowledge of criteria for selecting AI tools or services, including task fit, performance, interpretability, privacy, security, licensing, cost, support, integration, and ethical implications, 2.46 average.

	<u>K38 Knowledge of how to differentiate no-code, low-code, and code-based AI implementation approaches and identify when each is appropriate, 2.23 average.</u> <u>K39 Knowledge of how to differentiate no-code, low-code, and code-based AI implementation approaches and identify when each is appropriate, 2.08 average.</u> <u>K40 Knowledge of how to distinguish AI-enhanced from AI-native systems and explain implications for availability, support, monitoring, vendor dependence, and risk, 2.23 average.</u> One employer suggested these four (K37, K38, K39, K40) also be combined into a single high-level K item. They're all "flavors within the concept of AI." To him, this is all still general knowledge, not a deep dive. No one on the BILT disagreed with this proposed change. ACTION: consider combining K36, K38, K39, and K40 into a single, higher-level K item. <u>K43 Knowledge of reliability, robustness, transparency, and explainability and their importance to testing, troubleshooting, auditing, and user trust, 2.46 average.</u> One employer would have voted higher if K43 was "knowledge of the definitions" rather than "knowledge of." None of the BILT members disagreed with that suggestion. ACTION: consider rewriting K43 to read "Knowledge of the definition of reliability et al."
Skills	Ann moved on to the skills. She noted that the skills got a lot of high votes – most of the skills items are in the "green" with higher vote averages than the knowledge areas. One employer pointed out that most of these feel like "core IT skills." <u>S28 - Skill in saving, labeling, versioning, and appropriately sharing effective prompts and relevant settings so AI-assisted work can be reproduced and improved, 2.77 average.</u> One employer didn't like the way this was written. Rather than "AI-assisted" work, it should be "human-assisted by AI." Another employer suggested that this is essential. Students need to learn the skills to leverage AI and make their job more productive. This is not just one course; it's a "longitudinal" thing across all of IT. Another employer agreed that this is a good concept. He wondered if S25 (Skill in monitoring changes in AI tools and policies, evaluate new capabilities, and pursue relevant next-step learning) should be combined into S28 because the tools are updating and changing so rapidly. The updated releases aren't coming just once a year anymore. This goes back to K20 about knowing the policies. We all need to be "policy works" when it comes to AI. <u>S-31 - Skill in defining the problem, users, constraints, desired outcomes, success criteria, and appropriate human role before selecting an AI-enabled solution, 3.08 average.</u> <u>S32 - Skill in assessing data for relevance, accuracy, completeness, cleanliness, integrity, provenance, authenticity, representativeness, and appropriateness for use, 3.08 average.</u> One employer suggested S31 and S32 were overarching "overall skills" for everybody and taught in every class. Students need to understand where AI make mistakes. There are a few others on the list that similarly be considered "overall" for everyone with "IT-specific" underneath. Ann offered to send this employer the worksheet to provide his notes after the

	meeting. The BILT liked this idea – NITIC will share the sheet with all BILT members for additional feedback. Another employer noted that a lot of the skills in the 20s seem “governance-focused” but there is not a lot of skills related to actual governance. Students need to know how organizational policies can impact business decisions. This may be more knowledge than skill: the employer can see “AI governance committees” driving most of this. Everyone else will be reacting to the policies that came out of the governance decisions. One employer discussed ITAR (International Traffic in Arms Regulations, which controls the handling of defense technology). Employees know that ITAR restricts certain information and demands specific handling. He sees AI in a similar way – you’re going to have to apply ITAR-style rules when working with AI. You need to know what you can and cannot do. <u>S34 - Skill in comparing AI tools or services using task fit, quality, security, privacy, integration, cost, latency, accessibility, supportability, and organizational requirements, 2.54 average.</u> One employer found this to be too broad. Those who disagreed based on the vote split did not speak up. <u>S38 - Skill in monitoring AI-enabled services and use logs, tests, user reports, and documented baselines to identify errors, drift, misuse, availability issues, or unexpected behavior, 2.54 average.</u> One employer noted this likely isn’t something an entry-level person will be doing. It’s too “deep” on the security side. Another employer agreed that even a senior-level infrastructure or coding technicians likely shouldn’t be doing this. S38 is more for security technicians. The first employer wondered about moving S38 up to the knowledge items and make it more about “knowledge of who to ask.” He noted that people have a lot of misconceptions about AI engines and LLMs and use them inappropriately. Mistakenly providing proprietary data into public LLMs is an ongoing problem. <u>S42 - Skill in evaluating changes in AI tools, services, threats, standards, and vendor capabilities and update technical practices accordingly, 2.62 average.</u> One employer thought this had already been covered elsewhere, which is why he voted it low. Maybe this one could be combined with some of the others. [Indeed, there may be overlap with S25 (Skill in monitoring changes in AI tools and policies, evaluating new capabilities, and pursuing relevant next-step learning).]
Voting for next time	Ann next asked the BILT to think about what might be missing from the list. One employer suggested “how to use AI to make my job better.” He is not sure he saw that threaded through the list. He saw definitions of AI and what to watch for, but nothing about how to leverage AI for their daily tasks. Another employer agreed – he refers to this as “AI proofing” your job. How do you make yourself more valuable to company by using AI tools? Ann agreed that if you don’t do this, you’ll be replaced by someone who can use AI. Another employer replied that some of this did appear on the list (e.g., K12 - Knowledge of how prompt clarity, specificity, context, audience, tone, constraints, and desired format influence AI outputs and S6 - Skill in writing clear prompts that specify the task, context,

	audience, goal, constraints, and desired output format), but reiterated the need for students to learn how to write specific, effective prompts. He referenced CATS – context, angle, task, style. “If you can provide those four pieces of information to even a modest AI, you'll get a reasonable answer.” Another employer agreed this should be a knowledge area. Once you have the knowledge, the skill will improve over time. Another employer mentioned again K28 and the need to have a “working knowledge of the current regulatory framework” that’s applicable to that specific job and company. It will vary based on location. One employer added that it’s up to the employer to provide the rules to employees. His company has AI rules, including not providing AI-generated results to anyone until you review and approve them because AI makes mistakes; that is, be sure you agree with the AI results to the best of your ability. Another employer agreed that a company needs to inform employees of policies, but his point about K28 is more about the regulatory framework above the company policy. A company could have written rules eight months ago, but maybe the government regulations have changed. If you just follow the eight-month old policies without knowing the larger government changes, you could be breaking the law. Ann wondered about a smaller company without large legal departments who can stay current with changing laws. The BILT seemed to agree that those smaller companies may need outside help to ensure compliance. One employer suggested that new hires should ask their company about their rules (and also local government regulations) for using AI. What am I allowed to do? He sees a lot of PII and HIPAA-protected data getting leaked. In this case, ask for permission and not forgiveness.
Conclusion	Ann thanked everyone for attending. She mentioned the next BILT meetings -- October 8, which will be a multi-disciplinary IT trends meeting, and December 8, when the BILT will vote and discuss technical support knowledge and skills.
Next Meeting: Thursday, October 8 (9:30am-10:30am Central/10:30am-11:30am Eastern) – IT trends	