IT Assessment: Council Email Publication and AI-Assisted Review

1. Cybersecurity Concerns with Proactive Publication

Information harvesting and aggregation

Publishing correspondence on an open website makes names, email addresses, phone numbers, signatures, organizations, and contextual information easier to collect at scale. Automated crawlers, data aggregators, and malicious actors can use exposed information to support spam, phishing, impersonation, social engineering, or other targeted activity.

Persistent hyperlink risk

Council correspondence may contain hyperlinks that are legitimate when an email is initially reviewed and published. The destination can change later through compromise, expiration, changes in ownership, or other events. A link that was safe at publication can therefore become malicious in the future while remaining accessible from City-hosted correspondence.

2. AI-Assisted Review and Redaction

IT believes an AI-assisted workflow is technically possible. The workflow could perform an initial screening and flag or propose redactions for predefined categories such as personally identifiable information, financial or medical information, threatening or inappropriate language, and other content selected by Legal or leadership.

Potential capability	Important limitation
Detect common PII and sensitive data patterns	Context-dependent information may be missed or incorrectly flagged.
Flag potentially threatening or inappropriate language	Sarcasm, coded language, context, and ambiguous wording can reduce reliability.
Propose redactions for human approval	A proposed redaction is not a legal determination that publication is appropriate.
Standardize an initial screening step	Automation can create false confidence if staff treat a 'no issues found' result as assurance.

3. Reliability and Human Review

The central limitation is reliability. Until a workflow is tested against representative Council correspondence, IT cannot determine how accurately it will identify all information that requires redaction or review. False positives would create unnecessary staff work; false negatives are more consequential because sensitive or otherwise problematic content could be published.

For that reason, AI should be treated as an assistive screening tool rather than the final publication authority. Regardless of the technology selected, a meaningful human review process would still be required to address legal and contextual judgments and to validate proposed redactions.

4. Development Feasibility

Based on the initial technical review, IT believes a prototype/beta workflow could likely be developed in roughly two to three weeks.

5. Decision Point

The policy choice is whether the transparency benefit of proactive Internet publication warrants the additional legal review, cybersecurity exposure, technology effort, and ongoing human oversight. IT can technically support either direction, but an AI solution should not be viewed as eliminating the risks or staff responsibilities identified by Legal and IT.

6. IT Recommendation

This assessment aligns with the City Attorney's staff recommendation to discontinue the current proactive posting practice. From a cybersecurity and operational standpoint, retaining Council correspondence as public records and providing responsive records through the established public-records-request (PRR) process would reduce the additional exposure, hyperlink risk, technology effort, and ongoing human-review burden associated with maintaining a proactive public web archive.