AUSTRALIAN COMMISSION
ON SAFETY AND QUALITY IN HEALTH CARE

Seventh clinical safety review of the My Health Record system

Review 7.3: Assessing downtime management best practices for clinical safety in health IT systems

October 2016

Contents

Background

The Australian Commission on Safety and Quality in Health Care (the Commission) has undertaken a clinical safety program for the My Health Record system since the system's implementation in 2012. In July 2015, the Australian Government Department of Health appointed the Commission to conduct the seventh clinical safety review of the system, with the oversight of the Commission's Clinical Safety Oversight Committee (CSOC).

The aim of the Commission's clinical safety reviews is to proactively identify potential clinical safety risks to, and arising from, the My Health Record system and to recommend suggested mitigation strategies. This will improve the overall safety and quality of the system over time.

Copies of the Commission's completed clinical safety reviews and the System Operator status reports against review recommendations to date are available on the Commission website.

The seventh clinical safety review of the My Health Record system was conducted by the Commission in 2016. Review 7, comprises three distinct review reports:

- Review 7.1: assessing the impact and safety of the use of the My Health Record system in emergency departments (the hospital emergency department review)
- Review 7.2: assessing the presentation to healthcare providers of the My Health Record system 'medications views' (the medications view review)
- Review 7.3: assessing downtime management best practices for clinical safety in health IT systems (the downtime management review).

This report presents the findings of clinical safety review 7.3. This review component assessed downtime management for the My Health Record system against current best practice for contingency planning to maintain clinical safety of digital health records.

Review objectives and scope

Downtime is a period of time that an IT system is not available, or only partially available, due to planned maintenance or an unplanned incident. In health care, the ability of downtime to disrupt service delivery and pose risks to patient safety increases as digitisation of electronic records, formularies, order entry, results reporting, decisions support and other elements of clinical care and documentation increases.

The aim of this review was to assess policies and procedures for managing the availability of the My Health Record system for healthcare providers and consumers.

Methodology

Policies and procedures for managing planned and unplanned downtime of the My Health Record system were compared with current best practice for contingency planning to maintain clinical safety of digital health records. The Commission also assessed actual availability of the My Health Record system using reports about downtime that were captured by the System Operator during a 12-month period.

This was a ‘desktop’ review. Assessment of gaps between actual practices and the documented procedures for handling My Health Record system downtime was out of scope.

The review assumed that all actual practices had been documented. Detailed examination of the system design, architecture and network infrastructure were also beyond the scope of this review. The System Operator was unable to provide some of the documents requested, such as disaster recovery plans, because of system security considerations.

No direct observations of patient care were undertaken as part of the review. The review was undertaken consistent with the *Privacy Act 1988* (Cwlth) and the *My Health Records Act 2012* (Cwlth).

Findings

The review of current best practice for managing healthcare IT system downtime identified that strategies to minimise disruption due to downtime should be integrated into system design and operation in three ways – high availability, business continuity and disaster recovery:

- Digital health record systems need to be designed to be continuously operational for a desirably long time (*high availability*).
- When systems are in routine use, the focus shifts to ensuring that critical functions are available to users during significant disruptive events (business continuity), and to recovering and protecting system information (disaster recovery).

A range of best practice strategies were identified and used to inform the findings and recommendations (listed below) for managing My Health Record system downtime. The desktop review of My Health Record system policies and procedures for managing downtime resulting in a total of 16 findings and 17 recommendations. Each finding has a risk rating and a related recommendation. The risk rating guide used for this review is in Appendix A.

No findings were assessed as critical or high risk. One finding was classified as major, ten findings were assessed as moderate, and five findings were classified as a minor risk to the My Health Record system.

Finding 1: Most maintenance and updates occurred outside of business hours on weekends.

Risk rating: Minor

Analysis of temporal patterns showed that downtime was unevenly distributed over the week, with 46 per cent of downtime occurring on weekends. The eight planned events and four emergency extensions occurred before 9 am or after 8 pm; they were mostly restricted to weekends. Emergency extensions to scheduled maintenance occurred on Mondays, Thursdays and Sundays.

Recommendation 1: Identify the least disruptive windows for planned downtime based on actual patterns of system use by consumers and providers.

Finding 2: Reliance on the myGov portal was the biggest cause of unplanned downtime, which directly affects consumer access

Risk rating: Minor

An examination of system components revealed that software releases by the National Infrastructure Operator (NIO) accounted for 86 per cent of planned downtime. For unplanned events, the myGov service was responsible for 64 per cent of downtime, which mostly occurred between 6 am and midnight during the week (88 per cent, n = 23 events). These events directly affected consumer access to the system.

Although myGov can be considered outside the scope of the My Health Record System Operator, the reliability and availability of the myGov portal needs to be improved, which would in turn improve access to the My Health Record by consumers. This will be particularly important for the upcoming move to an opt-out model. The Commission was advised that improvements have been made to the reliability of the myGov service. The Commission was informed that, since November 2015, the Australian Government Department of Human Services (DHS) has gradually upgraded and transitioned myGov to Converged Infrastructure, following which there has been a reported decrease in unplanned downtime for myGov.

Provider access to the system was affected by downtime affecting the Healthcare Identifiers (HI) Service and NIO components and infrastructure, which accounted for 37 per cent of downtime.

When any planned change is performed on a My Health Record component, a window is selected following consultation between the NIO and the System Operator to ensure the time selected is suitable and during a low-usage period to reduce impact to users. The exception to this process involves a critical incident that requires an immediate solution to be applied.

Planned changes to DHS infrastructure are completed during low-traffic periods where existing systems provide 50 per cent of capacity during the change period. DHS has multimode systems in place to mitigate downtime.

The System Operator has confirmed that a read-only environment is now made available during major releases, whereby consumers and providers can view My Health Record information through the Consumer Portal (consumers) and business-to-business (B2B) gateway (providers), but they cannot make changes or upload new documents to the system.

During a major release, a message is posted on both the Consumer and Provider portals and the My Health Record website advising users that the system is temporarily unavailable. Similarly, users accessing the system via B2B transactions will also receive a message advising that the system is temporarily unavailable during a major release.

From a clinical safety perspective, it is important to note that the conformance requirements for the system also allow for queueing of clinical documents, should they be unable to be uploaded during downtime. The queueing process reduces the risk of clinical documents being lost during downtime.

Major system releases currently involve updates in a second environment (Disaster Recovery environment). These are then deployed into the live environment during a planned outage, at which time the system will be put in Read Only mode. Post-release, the redundant environment will be upgraded, at which point it will become the new Disaster Recovery environment. In the longer term, the system should move to a fault-tolerant architecture so that major system releases can be deployed in a redundant system and moved to the live environment when the update is successful, with no downtime. This will allow providers and consumers full access during planned system releases.

It is acknowledged that such a move would require significant investment and planning. The Commission has been advised that such zero downtime functionality was reviewed by the System Operator as part of the Read Only introduction and was found to have significant overhead costs. The Commission proposes that the possibility of using fault-tolerant architecture be reconsidered should the switch to opt-out at a national level be confirmed following the evaluation of the trial period.

Recommendation 2: That the Department of Human Services continues to improve the reliability of the myGov portal, thereby improving system availability for consumers. This is particularly important as use of the system grows over time.

Finding 3: My Health Record system downtimes were evenly split between planned and unplanned downtimes in the period examined.

Risk rating: Minor

Total downtime was evenly balanced between unplanned incidents (62 hours, 50 per cent) and planned events, which included release deployments (53 hours, 42 per cent) and emergency extensions (10 hours, 8 per cent)

Recommendation 3: That the system moves to a fault- tolerant architecture so that software can be updated on a redundant system and moved to the live environment when the update is successful, with no downtime.

[The recently implemented read- only environment upgrade strategy partially addresses this finding.]

Finding 4: There is no published documentation of availability management standards or current service level agreements.

Risk rating: Minor

The performance of an IT system is typically examined in relation to a service level agreement, which sets out the level of service expected by users from an operator, the metrics by which that service is measured, and the remedies or penalties (if any) should the agreed-upon levels not be achieved.

The review did not identify any references to availability management standards being followed (e.g. information technology infrastructure library [ITIL] practices). The review team was subsequently advised of the details of the current service level agreements, and the System Operator confirmed that ITIL standards are adhered to for incident management

Recommendation 4: That best practice availability management standards be documented, and availability metrics for each system component be published.

Recommendation 5: That system availability be set based on actual patterns of use by consumers and providers.

Finding 5: It is unclear how often the service level agreements (SLA) between the System Operator/ National Infrastructure Operator (NIO) and partner organisations have been reviewed.

Risk rating: Minor

The service level agreement between the System Operator and the NIO for My Health Record availability, including the national consumer portal, national provider portal and the national administration portal, is 99.5 per cent, excluding scheduled outages. The NIO also has service level agreements in place with partner organisations (e.g. Telstra) who are responsible for components of the system infrastructure that reflect the NIO's My Health Record availability service level agreements. The Commission has been advised that these agreements also document system monitoring processes for each relevant component of the system.

The DHS service level agreement (with the System Operator) for the HI Service and Medicare repositories is also 99.5 per cent per month, excluding scheduled outages. The

Commission has been advised that these agreements also document system monitoring processes for each relevant component of the system. The service level agreement for myGov availability is 99.5 per cent.

The Commission does not know if these SLAs have been reviewed and/or amended since the system went live in July 2012. Regular review of SLAs would be beneficial, particularly as system use grows.

Recommendation 6: Service level agreements should be reviewed to ensure they are providing the level of availability required, as usage of the system continues to increase.

Finding 6: Infrastructure and component failures contributed to unplanned downtime.

Risk rating: Moderate

The current level of redundancy for both infrastructure and application components could be improved. Overall, infrastructure and component failures equally contributed to unplanned downtime (8 per cent each). Although these figures are relatively low, these are largely avoidable if adequate redundancy and component monitoring are in place.

To ensure high availability through system design, single points of failure should be eliminated by adding redundancy so that the entire system does not fail when there is a component failure. The Safety Assurance Factors for Electronic Health Record Resilience (SAFER) guides recommend that data and application configurations are backed up and hardware systems are redundant. For example, mission-critical hardware systems (e.g. database servers, network routers, internet connections) should be duplicated.

The My Health Record system is supported by a disaster recovery environment. Full details of how replication to the disaster recovery environment occurs after an unplanned system release were not provided due to system security considerations. Although the documentation reviewed contained no details of the My Health Record system's disaster recovery architecture and processes, the System Operator has assured the Commission that these processes are documented.

DHS has advised that additional redundancy for the HI Service has been put in place by adding additional Java Virtual Machines during a scheduled upgrade.

Recommendation 7: That the System Operator improves reliability of both infrastructure and application components by eliminating single points of failure.

Finding 7: Details of the health checks for the business- to-business (B2B) gateway and the reporting portal have not been documented in the NIO Production Health Check Monitoring Approach.

Risk rating: Moderate

High availability is contingent on detecting failures when they occur. My Health Record system monitoring involves individual components (component-level health checks), as well as the system as a whole (end-to-end health checks). The Commission sought to assess the adequacy of this monitoring and the immediacy of generated alerts.

End-to-end monitoring seeks to mimic the interactions of end users with the system. In addition to the consumer, provider and administration portals, the B2B interface is critical to

My Health Record system operations because healthcare providers rely on this interface to post, replace and access documents in the My Health Record system from their clinical systems.

Based on documented procedures, there are gaps in the end-to-end checks. The health check for the provider portal is under development, but health checks for the B2B gateway, the reporting portal and the mobile app have not been documented. The System Operator has confirmed that the latest available revision of the NIO Health Check Monitoring document does not accurately reflect the current state of the system health checks.

Current tolerances for the end-to-end health checks could be too lenient once the expected transition to an opt-out model takes place. Currently, end-to-end health checks are conducted every 10 minutes, and a summary email is sent to the NIO Operations team mailbox. An SMS and email alert is sent to the NIO Operations on-call staff member only in the event of either three consecutive failures or if 30 per cent of the health checks fail within an hour.

Recommendation 8: Document and make available for routine review major end-to-end health checks for the provider portal, reporting portal, B2B gateway, business-to-mobile (B2M) gateway and the mobile app.

Finding 8: Only three component- level health checks are documented in the NIO Production Health Check Monitoring Approach.

Risk rating: Moderate

Component-level health checks seek to monitor individual system components that are local to the My Health Record system national infrastructure, as well as third-party applications linked to the My Health Record system. Only three component-level health checks are documented – the RLS-getDocumentList (internal), the Healthcare Identifiers Service (external) and the National Prescription and Dispense Repository (NPDR) (external) – from a total of around 30 system components. Of the total components, 19 are external to the My Health Record system national infrastructure.

The documentation states that not every component can be tested because of the nature of service or use of the component. However, it is not clear which components are not being tested because the document specifically states that additional health checks will be added. Plans to resolve individual component failure have not been documented.

The System Operator indicated that component-level health checks are in place for all components supporting transactions performed by consumers and healthcare providers, with two exceptions. These are for the component that manages connectivity with *myGov* (which is indirectly checked as part of the Consumer Portal end-to-end health check) and the one that manages access to the Provider Portal.

There is no documented evidence of any checks being performed on key system services, including:

- My Health Record view services (e.g. health record overview, Medicare information view and NPDR view)
- My Health Record ‘get document’ service (provides access to all the Clinical Document Architecture documents stored within the My Health Record)
- Medicare Repository’s ‘get document’ service.

Although these services may be covered in the end-to-end checks, such critical services should be checked at the component level to ensure faster identification of any failure.

Recommendation 9: Identify and implement health checks for critical system components, including the My Health Record view services, the My Health Record get document service and the Medicare Repository get document service.

Finding 9: Tolerances for end-to-end and component-level health checks could become too lenient with the transition to an opt-out model and as use of and reliance on My Health Record increases.

Risk rating: Moderate

With the expected transition to an opt-out model, the current tolerances for the component-level health checks could be too lenient. Component-level health checks are conducted every 5 minutes, and a summary email is sent to the NIO Operations team mailbox. An email and SMS alert is sent to the NIO Operations on-call staff member only in the event of either three consecutive failures or if 30 per cent of the health checks fail within an hour

Recommendation 10: Review current thresholds to determine whether lower thresholds for alerting, based on end-to-end and component health checks, are required.

Finding 10: There are no documented details in the NIO Production Health Check Monitoring Approach on how the system infrastructure is monitored.

Risk rating: Moderate

Other gaps in documentation relate to monitoring of data integrity across components. Continuous monitoring of data integrity is important to maintain system quality and safety. For example, the SAFER guides recommend that organisations monitor the total number of database transaction errors daily, and the percentage of data transaction errors that have been investigated and fixed.[1] The System Operator confirmed that the NIO monitors the total number of database transactions identified, investigated and fixed, and processes are in place to communicate issues to the vendor community.

The Commission was not provided with documented details on how system infrastructure is monitored. Critical infrastructure includes Central Processing Unit (CPU) and memory usage, network traffic and latency, available disk space, and storage area network read and write speed. By closely monitoring infrastructure, it is often possible to identify when an outage is imminent. This gives operators the opportunity to take proactive action and reduce the incidence of unplanned outages. The System Operator confirmed that the NIO's infrastructure providers monitor the system infrastructure on behalf of the System Operator in accordance with the terms of their service contracts.

Recommendation 11: Ensure that the process for monitoring overall system performance and partial system availability is documented.

Finding 11: There are no documented details on the monitoring of data integrity across system components.

Risk rating: Moderate

The Commission was provided with the PCEHRSPEC-087 Reconciliation of PCEHR Data v3.0 document that outlines a number of reports available to extract key record data and metrics from both My Health Record and the Medicare Repository. However, the document does not cover how the available reports are used in practice to reconcile data across the systems.

Recommendation 12: Ensure monitoring of data integrity across all My Health Record system components is occurring and is documented.

Finding 12: There is considerable delay in communication of system status, particularly outside of business hours (i.e. 8.30 am to 5 pm).

Risk rating: Major

For planned events, the review team considers the current procedure in which consumers are directed to the helpline for further assistance as appropriate. However, there is a considerable lag with reporting unplanned events, and there are gaps in communication of system status, particularly outside of business hours (i.e. 8.30 am to 5 pm).

For unplanned incidents, the system status is updated on the service availability page only when incidents occur during business hours and if they are likely to continue for longer than 60 minutes. Outside of business hours, implementation of redirections and updates to the service availability page and myGov are at the discretion of designated incident managers.

Of the 26 unplanned incidents examined, 13 were shorter than 60 minutes (50 per cent). All but one of the remaining 13 occurred outside of business hours. Thus, it is highly likely that consumers were notified about only one incident over the entire 12-month period examined.

Recommendation 13: Notify system participants about unplanned interruptions as soon as possible.

Finding 13: There are gaps in communication of the system status – users are notified about downtime but not service degradation and partial availability.

Risk rating: Moderate

Users do not appear to be informed if the system is running slowly or is only partially available. This issue is likely to create frustration amongst users as the reliance of the My Health Record system increases over time.

Recommendation 14: Adopt proactive strategies (e.g. SMS, email, or through clinical systems) to notify providers about My Health Record system status, including service degradation and partial availability.

Finding 14: Procedures to communicate downtime to healthcare providers appear to be heavily reliant on vendors and the My Health Record website.

Risk rating: Moderate

In addition, procedures to communicate downtime events do not appear to extend beyond the consumer portal, the My Health Record and myGov websites and to software vendors whose products interact with the system. This includes notifications from DHS to vendors during periods of HI service downtime and subsequent service restoration. Consideration needs to be given as to whether relying on vendors to communicate downtime is augmented with additional notifications direct to registered healthcare organisations. The System Operator has advised the Commission that the incident management process follows ITIL standards, as per the relevant service contracts. The Commission also notes that there is no mention of the business-to-mobile gateway in the list of system components in the My Health Record Incident Management Framework.

The process for planned events is robust. Although source code and configuration management were not covered in the documentation reviewed, further discussions with the System Operator have confirmed that all source code and configuration items are managed using Accenture's proprietary delivery tools and methods (Accenture is the National Infrastructure Operator).

Recommendation 15: Develop capacity to communicate downtime through all available system interfaces, and consider direct notifications to registered healthcare organisations.

Finding 15: There are no identified measures to ensure data integrity across My Health Record system components immediately following a downtime event.

Risk rating: Moderate

The reviewer team considered the potential for errors in My Health Record system information relating to documents that are incorrect, absent, only partially present or delayed due to a downtime event), because these types of errors have been shown to pose risks to patient safety.[2] Examination of the three main system components B2B interface, NPDR and Medicare Repository did not identify measures to protect information in transit.

For example, posting a document via the B2B interface involves multiple transactions. Many transactions also extend to components outside the infrastructure of the National Infrastructure Operator. It is not clear what happens to documents and information in transit (i.e. partially uploaded) if such processes are disrupted by downtime.

The Commission conducted an incident review (PCEHRIII-68) into prescription documents not being uploaded into the My Health Record system due to intermittent Healthcare Identifiers Service connection issues in 2015. Following that review, the Commission recommended that the solution architecture for capturing prescription and dispense documents be reconfigured, so that the likelihood of a similar incident arising would be reduced. The recommendation made through the incident review is still valid, as the Commission understands the proposed changes are not yet in operation.

Recommendation 16: Refine processes and tools to review transactions and realign data when critical system transactions are disrupted by downtime.

Finding 16: Current documentation does not include the B2B gateway and procedures to operationalise existing tests.

Risk rating: Moderate

Residual effects are the effects of downtime that may persist until system processes return to a steady state.[1]

The review found no specific procedures to monitor the My Health Record system immediately after a downtime. Procedures to return the My Health Record system to normal operation are restricted to a business verification test plan, which includes a number of tests on certain aspects of the system infrastructure. The B2B gateway is absent from this list, and the procedure to operationalise these tests after a downtime is not documented

Recommendation 17: Document procedures to operationalise My Health Record business verification tests, including the B2B gateway, immediately after downtime to return the system to normal operation.

Conclusion

This review examined policies and procedures for managing the availability of the My Health Record system in the context of actual system availability over a 12-month period. The review team did not identify any critical issues, but there are several major and moderate issues with the policies and procedures that should be addressed to ensure adequate provider and consumer participation, particularly under future opt-out arrangements. Because downtime results in a decline in service quality on a large scale, affecting all consumers and providers, it directly affects the safe use of the My Health Record system and can profoundly influence its reputation. Growth of the system will result in providers and consumers increasingly relying on the system to support clinical decisions, meaning that downtime events will lead to increasing risk unless they are managed appropriately.

Appendix A Clinical safety review risk rating matrix

Review findings have been assigned one of five risk ratings – critical, major, moderate, minor and minimum, consistent with the review's clinical safety risk rating matrix (Table A1).

These categories have been confirmed by the Commission's Clinical Safety Oversight Committee and the My Health Record System Operator during the review process.

Table A1 Clinical safety review risk rating matrix

Risk rating	Reputation and public confidence of My Health Record / quality of service	Clinical safety harm	Control
Critical	Profound influence on the My Health Record system's reputation, resulting in a profound loss of public and healthcare provider participation Profound sustained degradation of service value and quality	A clinical incident resulting in patient death	Basic, supervisory and/or monitoring controls are inadequate and require urgent management attention A critical patient safety incident has occurred
Major	Significant influence on the My Health Record system's reputation, resulting in significant loss of public and healthcare provider participation Decline in service value and quality is recognised by a majority of patients or health service providers	A clinical incident resulting in major permanent loss of function	Basic, supervisory and/or monitoring controls are inadequate and require prompt management attention A major clinical safety incident has occurred
Moderate	Loss of reputation affecting participation in the My Health Record system Decline in service value and quality is recognised by a moderate number of patients and health service providers	A clinical incident resulting in permanent reduction in function	Basic, supervisory and/or monitoring controls are partly inadequate and require management attention High potential for a clinical safety incident
Minor	Mild damage to reputation of the My Health Record system Decline in service value and quality is recognised by the System Operator and My Health Record partners	A clinical incident resulting in increased level of care/intervention	Basic, supervisory and/or monitoring controls are operating as intended, recommendation for improvement to strengthen control
Minimum	Minimal impact on the My Health Record system's reputation Minimal effect on service value and quality	A clinical incident resulting in no injury	Basic, supervisory and/or monitoring controls are operating effectively, a process improvement opportunity exists

References

1. Sittig DF, Gonzalez D, Singh H. Contingency planning for electronic health record-based care continuity: a survey of recommended practices. Int J Med Inform 2014;83(11):797–804.

2. Coiera E, Magrabi F. Information system safety: guide to health informatics. Boca Raton, USA: CRC Press, Taylor & Francis Group, 2015:195–220.