

Findings Report on the July 8th Telstra Outage

By

Technology Audit Partners



August 24, 2026

1 Approach

Technology Audit Partners (TAP) has been engaged by Telstra to conduct an external expert review and analysis of the 8th of July 2026 network Outage, being the outage events from the beginning of Wednesday, July 8th through to Saturday, July 11th and also including follow-on impacts after the outage which involved post-Outage resolution by Telstra. The intention of this review is to provide a clear understanding of the events that took place, any events and decisions in the past leading up to and contributing to the Outage, how Telstra responded to recover from the Outage and our recommendations to address the technical, process and execution causes of the Outage, its scale and its duration. Where there are other matters related to this Outage that are being addressed directly by Telstra, they will not be included within this analysis unless it specifically connects to the scope of this external expert review.

Throughout this investigation, Telstra has provided all of the information and documentation that was requested by TAP in a prompt manner. TAP has also been provided access to every person we requested for an interview and our assessment is that all interview participants were open and transparent in their responses to our questions.

The Outage analysis of the events and causes was performed from three distinct perspectives:

- Detailed analysis of the events and circumstances and network conditions leading up to and including the Outage triggered when the incorrect date was sent to the network.
- Detailed analysis of the Outage response time focused on time to detect the Outage, time to know the cause of the incorrect date and time to respond and mitigate the source of incorrect date.
- Detailed analysis of the propagation of incorrect date across the network, the impacts and the associated remediation and recovery steps.

Based on detailed analysis from these three perspectives, the key findings will be elaborated and the associated recommendations for correction and/or improvement will be provided for the findings. In our findings when we mention concerns with Problem Management, Change Management and so forth, we are speaking of the commonly applied functions and processes and not the teams in Telstra that might carry the same names as those functions. These improvement recommendations are not targeted at or limited to individual teams.

2 Analysis

While the Outage first began to be widely visible to Telstra at 4:20am when an NTP (Network Time Protocol) alarm storm was reported and then at 4:29am when Incident Management Operations (IMO) systems detected that customer reports exceeded thresholds based on customer web searches for Telstra outages, the very beginning of the Outage started at the completion of a planned NTP (Network Time Protocol) timing chassis change at 2:50am when the chassis was powered up and back on line. This chassis was located in Melbourne and it was part of the mobile core timing architecture and was designed to provide stratum 2 and stratum 3 NTP services. When this server went back on line (at 2:50am), the GPS card in the chassis had reset during the change and began sending the incorrect date from 2006 into the network. The GPS card was actively used in that chassis and was missing a critical firmware update that was required to prevent this shift in time due to a known GPS rollover that was published in bulletins by the supplier. While the planned change of the NTP server chassis was the trigger point to begin the Outage, the steps leading to that point go back many years, as far back as 2020.

Telstra's NTP server network evolution over time

In 2010, Telstra designed a mobile core timing architecture that was fit for purpose. Three stratum 1 sources from the National Measurement Institute (NMI) were used to provide timing to two stratum 2 NTP servers (in Sydney and Melbourne). These provided time to three stratum 3 NTP servers (in Sydney, Melbourne and Perth). These then provided time to the various clients in Telstra's mobile core network. They were all configured in a Client-Server configuration with peering only allowed across the same stratum levels. This architecture provided a highly reliable and authoritative reference time source from NMI and enforced stable time sources between the layers of the mobile core timing architecture. It should be noted that this architecture and the architecture changes described below are no longer present in the Telstra network.

In 2020, as part of an upgrade to mobile core timing, several changes were made to the configuration of the mobile core timing architecture, introducing a replacement NTP timing chassis to replace the original aging hardware.

- 1) Due to a limitation in the NTP timing chassis, a stratum 2 NTP server could not provide time to a stratum 3 NTP server in the same chassis. Based on this limitation, Sydney stratum 3 NTP server could only obtain stratum 2 time from Melbourne stratum 2 NTP server and also Melbourne stratum 3 NTP server could only obtain stratum 2 time from Sydney stratum 2 NTP server. The Perth stratum 3 server in a different chassis, at a different location, continued to have stratum 2 redundancy.
- 2) A GPS card was introduced along with the NTP timing chassis. The architecture indicated that the GPS card was available to be used as an alternative failover source for time of day, but it was never configured for automatic failover and required manual reconfiguration to take effect.
- 3) The mobile core timing architecture was changed from client/server to peering, meaning that the NTP server in the NTP timing chassis could acquire any NTP source that it could find and select to use any of these sources that was providing the most stable and accurate time. This introduced uncertainty in the time sources being used in mobile core timing hierarchy.

These modifications degraded the mobile core timing architecture with a known degradation of lost redundancy of stratum 2 sources at Sydney and Melbourne and the unknown introduction of looping, a behaviour that is explained in section 4.1. We have no evidence that the personnel involved with the mobile core timing architecture changes which were made in 2020 were aware of or investigated this risk in the architecture.

In October 2025, personnel were working on the Melbourne NTP server due to an incident report that the NTP stratum 3 clock was losing connectivity multiple times per day to its stratum 2 source clock at Sydney, resulting in NTP alarms. The incident report also noted that the Melbourne NTP had become a stratum 5 source meaning that it was receiving time from a stratum 4 source in the mobile core timing architecture. Remember, Melbourne had only one stratum 2 source and when that source was lost, it used peering to find another source which happened to be a stratum 4 client of Sydney or Perth. To resolve this issue the GPS card was activated and connected to the stratum 3 NTP server in Melbourne as a means of replacing the lost Sydney stratum 2 source. Since there is no evidence of a prior design or configuration change review, the personnel working on the incident appear to have chosen this approach as the best means to address the incident. The GPS card is known as a reference source (aka stratum 0) and so this connection made Melbourne a stratum 1 NTP source, effectively the same level as the NMI stratum 1 sources in the mobile core timing architecture. This impacted the clients of Melbourne, making them stratum 2. In peering mode, these stratum 2 clients can be time servers at

the same priority as Sydney and Melbourne stratum 2 who were the intended stratum 2 sources. When Melbourne NTP was connected to the GPS card a 'Network at Risk' ticket should have been raised to find the real root cause of the frequent loss of the Sydney source and restore the NTP server back the intended design but this did not happen.

At the time of the October 2025 incident, the personnel involved did not recognize the meaning of Melbourne NTP at stratum 5; that looping was already occurring in mobile core timing when an intended NTP stratum 3 server was able to become a client of one (or more) of the clients and that a circular loop was established when the stratum 2 source was lost. This anomaly was observed but not noted and no problem ticket was created for further investigation.

As stated at the beginning, on 8 July 2026, during the chassis change to fix a power supply, the GPS card in use at Melbourne without the critical firmware update reset, and the date changed to 2006. Since Melbourne was operating at stratum 1 and peering was enabled and Melbourne stratum 2 had been turned off by the same chassis replacement, the incorrect date propagated throughout the mobile core timing network.

What Happened after the date changed to 2006

As the incorrect date propagated throughout mobile core timing, due to looping, all of the available NTP time sources began to agree that the date was 2006 and when the majority of available time sources agreed that the date was 2006, the clients - the network nodes and applications - accepted the incorrect date which began to cause failures with different impacts depending on the failure mode of the impacted node or application. Not all nodes and applications were impacted. As the incorrect date propagated the network impacts also propagated resulting in increasing impact to users relying on the Telstra mobile network, including Telstra customers and relevant MVNO customers.

Starting at around 4am July 8th, 2026, when Mobility Management Entity (MME) instances began to reset, end users of Telstra's network experienced varying levels of impact in their ability to communicate on the network, however not all users were affected. These impacts precluded some users from being able to make/receive voice calls, including Emergency Calls (per the Emergency Call Service Determination:2019, also referred to as Triple Zero calls), send/receive messages, or access mobile data services. Those without data services were unable to do things such as access the web, stream content and most mobile applications were unusable as well. This affected all forms of devices, including personal communication devices such as smartphones as well as IoT devices. Impacts extended to a variety of access technologies including 4G, 5G, fixed and Voice-over-WiFi.

Initially users were impacted by a loss of network capacity (due to failure of MME instances) which caused some users to be disconnected from the network. The remaining capacity was insufficient to enable those users to all reconnect deterministically nor was it capable of servicing the nominal network traffic load. So essentially the network was overloaded. As the network recovery progressed, the MME capacity was restored but other problems created longer lasting issues for some users. By 12:12pm on July 8th 2026 the majority of postpaid customer services were restored while a variety of residual issues continued to be investigated and resolved through July 11th 2026.

Network

The IP Multimedia Subsystem (IMS) Core was significantly impacted. It is responsible for signalling management between the user and the services they are trying to use. Until the correct network date was restored, information for approximately 30,000 IP address prefixes were incrementally corrupted (written with the wrong date) by the Diameter Signalling Controller (DSC) in such a way that access attempts using those addresses were blocked from using the network regardless of available capacity. Telstra, with help from their equipment vendor, applied a workaround during the Outage and then at a later time performed repairs to remove the bad IP address prefix entries.

Other applications and systems in the Telstra network were also impacted. These include billing, prepaid, performance metrics and other systems. While Emergency Calling does not use all of these services it was also impacted, however, Telstra worked effectively with their vendor to create a workaround to restore Emergency Call handling by the network.

The Outage recovery process worked well once the root cause of the problem was uncovered and the full MIM (Major Incident Management) bridge took over problem detection, isolation and recovery. It took several hours to determine the root cause of the Outage from its origin at 2:50am. There were several reasons for this primarily due to insufficient clarity of responsibility regarding NTP ownership and support, lack of visibility to the planned event to replace the Melbourne chassis that night, as well as a lack of understanding of the alarms that were the earliest indicators of a problem in the network.

Summary of Findings

Overall, we found 4 overarching findings as an outcome of this investigation:

- Unclear and insufficient NTP ownership and lack of recognition of its importance as a Sovereign Function (as discussed in section 3.1 below) which led to weakness in this area and not treating it as a high risk function of the network for the purpose of design, operations, resources and investment prioritisation.
- No end-to-end architecture for mobile core timing, including client design and behaviour, and no architecture and design oversight that provides the 2nd line support to challenge and steer design.
- Insufficient technical expertise for NTP and the need for stronger competency and more curiosity to investigate when something looks wrong.
- Increase attention to the operation and management process discipline in all areas. There were several instances of insufficient execution discipline for NTP including change management, configuration control and incident management including elements of attention to detail, record keeping and documentation and follow through.

These four summary findings are elaborated across the eleven finding areas where the discussion below will identify gaps from expected practice based on this analysis and provide recommendations for improvement.

3 Findings and Recommendations

This section is a list of findings and recommendations which are based on the TAP investigation into the Outage of 8 July 2026. It is broken down into process functional groupings to assist with the organization and planning of improvement effort. As noted above, these groupings are not Telstra teams or organization names.

A number of common findings have been identified which span one or more of the detailed recommendations. Before going into each of them individually, we further observe that underlying most of them are two key drivers which impact – both positively and negatively – the performance levels across many of the specific findings and carry additional recommendations as findings in their own right.

Common Finding Areas	NTP as a Sovereign Function	NTP Ownership Definition and Implementation
Architecture and Design	X	X
Configuration Management	X	X
Change Management	X	X
Problem Management	-	X
Incident Management	-	X
Risk Management	X	X
Resource & capability management	X	X
Investment / Lifecycle Management	X	-
Observability	X	-

Table 1 - Common Findings and Key Drivers

As shown in this table, each of these common finding areas, apply to either lack of treatment of NTP as a Sovereign Function, insufficient ownership or both as shown by the X in each column. Both of these key drivers are elaborated as additional macro level findings before analysing the individual common finding areas noted in the table above. It should be noted that the common findings are based on our in-depth review of mobile core timing associated with this Outage but it may be necessary to evaluate these findings in other areas to ensure that there are no gaps and weaknesses elsewhere.

3.1 NTP not identified as a Sovereign Function

Within modern networks, there are common functions that are highly leveraged across entire networks to provide key services. These functions are variously known as Common Failure Mode functions or Sovereign Functions. These functions create outsized risks to the network and must be managed with the utmost care. NTP was not designated by Telstra as one of these functions. Our recommendation is to establish a defined list of Sovereign Functions for Telstra and ensure NTP is listed as one of them. It is also recommended to review, and update as needed, the systematic handling of all Sovereign Functions to ensure sufficient care levels.

3.2 NTP Ownership Definition and Implementation

Not identifying NTP as a Sovereign Function (above) missed the opportunity for elevated performance in these common finding areas, where undefined or insufficient ownership appears to have degraded performance in many of these areas as noted in the table above. We observed instances where other areas and elements of the network and organizations were performing at better levels than NTP. Our recommendation is that in addition to treating NTP as Sovereign Function that NTP product ownership and responsibilities be reviewed and strengthened. Based on this Outage experience, it is also

recommended that defined roles and culture of Product Ownership be reviewed and updated based on learnings.

3.3 Architecture and Design

Architecture and design issues were substantially in play in setting the stage for this Outage going back as early as the change in 2020. There are also indications of insufficient architecture expertise in the areas involving NTP, and also the lack of an end-to-end architecture which integrates time serving/delivery design with time consumption/client design. We also observed that failure to understand the architecture and dependencies caused the instability of stratum imbalance in a network of mixed stratum sources to be overlooked. Telstra must establish a systematic end to end architecture for NTP timing and ensure the architects (internal and/or external) have expertise in the technology, risks and best practices. Additionally we recommend that Telstra bolsters independent review of new architectures and architecture change to ensure analysis of wider network impacts and possible failure modes.

3.4 Configuration Control

We did not observe any centralized and systematic repository of configuration information and records associated with NTP. There was no golden configuration, no documented configuration record for that matter other than the device itself against which to compare the device in operation. Pre-checks and impact assessment cannot be performed adequately without such records. Telstra should establish stronger configurations control for NTP, including documented configuration records and audits of the records themselves for completeness and quality and also audit documented configurations against actual deployment. We also advise a review of configuration control more broadly to identify if there are other areas with gaps.

3.5 Change Management

In the case of this Outage, we observed that there were no prechecks of equipment and software actually in place leading up to the change, inadequate technical review of the MOP (Method of Procedure) and localized approval – this despite the fact that the change was categorized as level 5 (the highest risk level). After our review of the MOP, it is our opinion that the MOP that was chosen was insufficient for this change and lacked critical detail and depth for the impacts of this change. In addition, after completion, the change was not tested to confirm that the key function of the chassis was working correctly (i.e., serving the correct time). Our recommendation is that Sovereign Functions (see above in section 3.1) receive the highest levels of change risk categorization, review and treatment including pre-checks. Furthermore, all MOPs for change should verify restoration of service function, in this case correct date. Practices for change start notification and change awareness across support personnel should be reviewed and updated as needed.

3.6 Problem Management

The industry standard would be logging intermittent issues and observed anomalies as a problem ticket and gathering evidence of the ongoing nature of the issue and possible solutions explored. There is no record of a problem ticket being opened or different solutions being explored for intermittent issues and anomalies for NTP we found from our analysis. Our recommendation is that stronger Problem Management practices should be applied in conjunction with addressing NTP ownership and capability.

3.7 Incident Management

Telstra appears to have strong Incident Management processes and people, in particular, once the full Outage MIM (Major Incident Management) process fully kicked in at 4:47am the process was organised

and effective in one of the most complex recoveries we have ever seen. However, NTP appears to have fallen into gap of coverage that impacted early outage detection and determination of cause. At the time of the Outage, support personnel could not find information about change that night or information about the people to contact or records of the servers themselves which delayed the determination of the source of the problem, how to get help and how to isolate the server. Our recommendation is that NTP support model is upgraded to be consistent with other groups and that NTP is audited to be included fully in standard Incident Management tools and databased used for support.

3.8 Risk Management

There was a failure to identify and manage this specific NTP risk despite evidence and concerns that were noted around the organization. Failure to observe the pattern of risk associated with lost or unreliable sync issues and then act on them proactively was a contributor to the conditions which allowed the Outage to occur, given that the GPS workaround was applied to address the stability issues. Our recommendation is to review bottom-up risk tracking for NTP to review constraints and risks are escalated and prioritised. Additionally, we recommend a review of risk management practices and procedures for Sovereign Functions generally.

3.9 Resource & capability management

There are too few engineering staff and insufficient expertise supporting the NTP platforms. On the night of the Outage the two primary NTP engineers who were involved in the planned change on the chassis were on mandatory standdown by the time the impact of that change became known. Overall, there is not sufficient knowledgeable staff to do all the necessary work and peer review for NTP to fulfill all product ownership and support responsibilities. Our recommendation is to review the NTP staffing to ensure it is sufficient and bolster capability depth for NTP and timing overall.

3.10 Investment / lifecycle management

The Transport & Fixed Access Engineering organization, which is responsible for the timing and synchronization services, has been unable to sufficiently modernize and maintain adequate support for NTP in part due to focus and priority as outlined in section 3.1 and 3.2 however it has been reported to also have been constrained in their choices by budgetary prioritisation. While these reports are not based on documentation and evidence, it is our opinion that the circumstances we found in this Outage analysis carry signs of these reports based on what has happened. Our recommendation is to ensure that investment in existing infrastructure, especially NTP and other Sovereign Functions, is balanced against growth and strategic objectives based on the increased risk posed by Sovereign Functions.

3.11 Observability

On the night of the Outage the NTP platform alarms were not surfaced in any of the standard support alarm monitoring tools and instead the NTP platform alarms were checked during business hours in the element management systems by a limited number of individuals. The client platform alarms did not have the correct severity or detail to aid troubleshooting or drive immediate action commensurate with the risk. Our recommendation is that NTP alarm monitoring is incorporated into standard support monitoring tools and 24x7 support and that alarm severity and description is reviewed and updated.

4 Background

4.1 NTP stratum levels and implications for network time configuration & management

The NTP protocol uses strata to indicate the expected accuracy of system clocks based on the number of levels from a reference clock. The device that has the most accuracy is the 'reference' clock—the clock from which the other devices on the system, synchronise their time. While the actual accuracy of a clock source is not determined solely by its stratum level, the stratum level is applied in the selection of final primary clock source when several good sources are presented.

A machine's time can be synchronised directly with the 'reference' clock, or via another machine (an 'NTP Server') which either synchronises its own clock directly with the 'reference' clock, or with the clock of another NTP server. Each level of synchronization is assigned a different stratum number. The lower the stratum number or level, the closer it is to the reference clock which makes it a preferred source, all other things being equal.

- Stratum 0 refers to the 'reference' clock itself—typically an atomic clock, GPS clock, or radio clock.
- Stratum 1 refers to any machine that synchronises its system clock directly with the 'reference' clock that resides at stratum 0. For example, a server that has a GPS unit plugged into one of its serial ports.
- Stratum 2 refers to any machine that synchronises its system clock with the clock on a stratum 1 server.
- Stratum 3 refers to any machine that synchronises its system clock with the clock on a stratum 2 server, and so on.

The furthest possible level is stratum 16. A machine at this level is deemed to not have a synchronised time source.

With machines that are configured to synchronise their system clocks via one of several NTP sources, the stratum level can vary depending which source is used. Such machines synchronise with whichever NTP time source is chosen by an algorithm that considers a number of factors - chiefly the agreement between sources and the precision of each, as well as stratum. Therefore, as NTP time sources are added, removed and changed, or as the connectivity network is changed, the stratum level can change and therefore impact the NTP server source selections of their clients which change their stratum level, thereby impacting other machines to change stratum level and on down through the network until a new stable time source hierarchy is established. Movement to a lower stratum will increase priority and utilization and movement to a higher stratum will decrease priority and utilization.

It is critical to have at least two independent reference sources. A timing loop is when sources that appear independent are in fact validating one another - directly, or by tracing back to a common reference. To avoid timing loops, you must clearly define the server hierarchy and document all relationships between servers. A consistent stratum architecture across your network simplifies troubleshooting and promotes predictable behaviour.

We observed that time source stratum changes and inconsistencies prior to October 2025 were present in Telstra's mobile core timing architecture.

4.2 GPS week number rollover behaviour

The Global Positioning System (GPS) broadcasts the number of weeks since the start of an epoch (epoch 0 started January 5, 1980). The main civil GPS signal (C/A code) broadcasts the GPS week number using

a 10-bit code with a maximum value of 1,023 weeks. This means every 19.7 years, the GPS week number in the C/A code rolls over to zero. This occurred for the first time in August 1999, and the second in April 2019. We are currently in the 2nd GPS epoch. The GPS receiver is responsible for determining the epoch and adding the correct multiple of 1024 to the broadcast week number. To avoid GPS device problems related to week number rollover, users must keep device software / firmware up to date. The phenomenon of GPS week rollover and its impacts is widely known across the industry.

When Telstra replaced the Melbourne NTP timing chassis on July 8, 2026, upon power up, the GPS card, without the critical firmware update installed, reset and used the previous epoch to calculate a date 1024 weeks in the past (November 2006) as the current date. This was then adopted by the Melbourne NTP server using that GPS card as its reference source.

4.3 Sovereign Functions (aka Common Failure Mode Functions)

Where Telstra applies the term Common Failure Mode functions, TAP uses the term Sovereign Functions¹ within its methodology. Sovereign Functions which provide critical network wide services also carry the risk of generating network wide outages. Critical functions include, but are not limited to, Border Gateway Protocol (BGP) for inter/intra-domain routing, Open Shortest Path First (OSPF) for intra-domain routing, Directory Name Service (DNS) for domain-to-IP resolution, Dynamic Host Configuration Protocol (DHCP) for Internet Protocol (IP) address assignment, Multiprotocol Label Switching (MPLS) for transport-independent traffic switching, Licensing, Firewalls and NTP/PTP (Precision Time Protocol) for time synchronization. All of these Sovereign Functions have been identified as sources of major network outages. Sovereign Functions will frequently ride across the network, both hardware and software, such that redundancy alone is not effective and neither are simple geography-based blast zones. Special care must be applied to network architecture and design with respect to Sovereign Functions and all practices for operations and change management must be done at the highest levels of review, approval and care. Personnel responsible for Sovereign Functions must be experienced specialists in their areas with current training on best practices.

¹ R. Owen, A. Bryant, L. Finch, D. Franklin, M. Abdollahi and M. Abolhasan, "Failures and Resilience in the IP Era: Navigating the Fragility of Modern Telecommunications Networks—The Sovereign Functions," in IEEE Access, vol. 13, pp. 155759-155777, 2025, doi: 10.1109/ACCESS.2025.3602054.