NENA NG9-1-1 Operational Diversity and Redundancy Information Document

Abstract: Provide guidelines and best practices to implement and maintain high-availability NG9-1-1 systems using diversity and redundancy.

NENA NG9-1-1 Operational Diversity and Redundancy Information Document

NENA-INF-047.1-2025
DSC Approval: March 4, 2025
PRC Approval: March 19, 2025
NENA Board of Directors Approval: April 28, 2025
Next Scheduled Review Date: April 28, 2028

Prepared by:
National Emergency Number Association (NENA) Systems Security & Resiliency Committee, NG9-1-1 Operational Diversity & Redundancy Working Group

1 Executive Overview

Purpose

The purpose of this information document is to provide service providers and 9-1-1 authorities with guidance on how to meet the high-availability requirements necessary to support mission critical emergency response operations using diversity and redundancy best practices in the design, implementation, and operation of 9-1-1 systems.

This document provides the following information:

- Explains how system availability is calculated
- Identifies methods and procedures that can be used to achieve diversity in 9-1-1 system routes, suppliers, software, and PSAP building access
- Outlines how to manage system changes without negatively impacting availability
- Identifies considerations that should be addressed in system Service Level Agreements (SLAs)
- Discusses expectations, processes used, and considerations when 9-1-1 system outages occur
  Describes the process and suggested best practices for vendor management of subcontractors.

Scope

This document is applicable to all 9-1-1 stakeholders including, but not limited to the following:

- 9-1-1 Authorities
- Public Safety Answering Points (PSAPs) and Emergency Communication Centers (ECCs)
- NG9-1-1 Emergency Service Internet Protocol (IP) network (ESInet) providers
- Next Generation Core Service (NGCS) Functional Element (FE) providers
- Telecommunications Service Providers (TSP) (inclusive of all access, carrier, service, technology, or media types used)
- 9-1-1 and NG9-1-1 vendors, solution providers, application providers, integrator, equipment hardware and software providers, etc.

Table of Contents

NENA
INFORMATION DOCUMENT
NOTICE

This Information Document (INF) is published by the National Emergency Number Association (NENA) as an information source for 9-1-1 System Service Providers, network interface vendors, system vendors, telecommunication service providers, and 9-1-1 Authorities. It is not intended to provide complete design or operation specifications or parameters or to assure the quality of performance for systems that process such equipment or services.

NENA reserves the right to revise this Information Document for any reason including, but not limited to:

- Conformity with criteria or standards promulgated by various agencies,
- Utilization of advances in the state of the technical arts,
- Reflecting changes in the design of equipment, network interfaces, or services described herein.

This document is an information source for the voluntary use of communication centers. It is not intended to be a complete operational directive.

It is possible that certain advances in technology or changes in governmental regulations will precede these revisions. All NENA documents are subject to change as technology or other influencing factors change. Therefore, this NENA document should not be the only source of information used. NENA recommends that readers contact their 9-1-1 System Service Provider (9-1-1 SSP) representative to ensure compatibility with the 9-1-1 network, and their legal counsel to ensure compliance with current regulations.

Patents may cover the specifications, techniques, or network interface/system characteristics disclosed herein. No license expressed or implied is hereby granted. This document shall not be construed as a suggestion to any manufacturer to modify or change any of its products, nor does this document represent any commitment by NENA or any affiliate thereof to purchase any product whether or not it provides the described characteristics.

By using this document, the user agrees that NENA will have no liability for any consequential, incidental, special, or punitive damages arising from use of the document.

NENA's Committees have developed this document. Recommendations for changes to this document may be submitted to:

National Emergency Number Association
1700 Diagonal Rd, Suite 500
Alexandria, VA 22314
202.466.4911

or commleadership@nena.org

2 Document Conventions

NENA: The 9-1-1 Association improves 9-1-1 through research, standards development, training, education, outreach, and advocacy. Our vision is a public made safer and more secure through universally-available state-of-the-art 9-1-1 systems and better-trained 9-1-1 professionals. Learn more at https://www.nena.org.

2.1 NENA Intellectual Property Rights (IPR) and Antitrust Policy

NOTE – The user's attention is called to the possibility that compliance with this document may require use of an invention covered by patent rights. By publication of this document, NENA takes no position with respect to the validity of any such claim(s) or of any patent rights in connection therewith. If a patent holder has filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license, then details may be obtained from NENA by contacting the Committee Resource Manager identified on NENA's website at https://www.nena.org/ipr.

Consistent with the NENA IPR and Antitrust Policy, available at https://www.nena.org/ipr, NENA invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement any standards referenced by this document or to implement or follow any recommended best practices, procedures, or architectures contained herein.

Please address the information to:

National Emergency Number Association
1700 Diagonal Rd, Suite 500
Alexandria, VA 22314
202.466.4911

or commleadership@nena.org

2.2 Reason for Issue

NENA reserves the right to modify this document. Upon revision, the reason(s) will be provided in the table below.

Document Number	Approval Date	Reason For Issue/Reissue
NENA-INF-047.1-2025	April 28, 2025	Initial Document

3 Availability

Availability refers to the probability that a system will be available. In the context of 9-1-1, this means the probability the 9-1-1 system will meet its Service Level Agreements (SLAs) for availability. Typically, this is measured in percentage and is the basis for “5-9s” terminology. A 5-9’s system will be available 99.999% of the time, which means less than five minutes total down time in a year.

Highly reliable systems such as NG9-1-1 systems require availability that cannot be achieved with a single instance of a function or service. These systems can achieve the desired availability (e.g., 99.999% uptime, a.k.a. as 5-9s) by deploying multiple instances of a single function or service even if an individual instance has lower than the required availability, given that the addition of redundant elements provide the desired level of availability.

3.1 Calculating Availability

Availability = Uptime / (Uptime + Downtime)

Availability can be measured, using actual uptime and downtime, but to use measurements as the basis for predicting it, many years of data would have to be collected to get statistically meaningful data. Consequently, availability is typically predicted. The prediction relies on two quantities: Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR). Like availability, MTBF and MTTR can be measured or predicted, and like availability, one needs years of observations to get meaningful measured data. Consequently, they are typically predicted. In terms of MTBF and MTTR:

Availability = MTBF / (MTBF + MTTR)

3.2 Mean Time Between Failures

Complex systems have many subsystems in series with one another, where all of them must work for the system to work. For example, an NG9-1-1 Next Generation Core Services implementation may have an access router, a firewall, a session border controller, two or three computers, a media server, an egress router, and three network switches, all of which must work for the NGCS to meet its availability SLA.

For “n” subsystems,

Series MTBF = 1 / (1/MTBF1 + ... + 1/MTBFn)

MTBF, and therefore availability, goes down as the number of systems goes up, but redundancy can ameliorate this effect, as discussed below.

In turn, the MTBF of each subsystem is calculated using the MTBF of each component in the subsystem with the above formula. Typically, there is enough data on the components, or very similar components, to get statistically significant data. The MTBF of all the components are aggregated in that formula to yield the predicted MTBF of a server, router, session border controller, etc. The manufacturer of the device does this calculation, and all quality manufacturers will make the predicted MTBF value available to their customers on request. Very often, the information is protected with a non-disclosure agreement, where the customer of the manufacturer is not permitted to reveal the MTBF to its customers. For example, in 9-1-1 systems, this means that the state 9-1-1 authority may not be able to get MTBF numbers from its NGCS vendor. The vendor can get them from its supplier. Any manufacturer unable to provide MTBF values should not be considered for use in a 9-1-1 system regardless of the other qualities of its products. While it is also possible, it is usually not practical to use actual real-world experience with a particular model of a product that has been operated in a similar environment for a sufficiently long period with no changes in design to obtain a statistically significant MTBF.

When analyzing a system, it is essential that every component that needs to work for the system to meet its availability SLA is included in the calculation. In the above example, in addition to each subsystem including the network components, the MTBF of the Heating, Ventilation, and Air Conditioning (HVAC) system, the electrical power system, and even the physical building (which might, for example, be breached by a hurricane) must be determined and used in the series MTBF calculation above.

3.3 Mean Time To Repair

System design determines MTBF. Operations determine MTTR.

MTTR is the mean time from the time a failure occurs until the time the instance is made available again. With a parallel redundant active/active system, the MTTR refers to the repair time for a failed instance, while another instance maintains service. While it's possible to determine MTTR for each subsystem separately, most systems use a worst case MTTR for all subsystems. When calculating MTTR, one must consider the following:

- When (worst case) a technician will be available.
- How long it takes (worst case) to arrive on-site.
- How long it takes to diagnose a failure and determine what spare must be used to bring the instance back to operational.
- Where spares are kept, and if not at the failed site, how long it takes to retrieve the spare and get it to the site.
- How long it takes to replace the failed component.
- How long it takes to reconfigure the instance with the replaced component.
- How long it takes to determine that the problem is resolved.

It is essential to realize that it is worst case that matters, not average. Many assumptions commonly made in commercial systems don't work if 5-9s are the availability requirement. For example, the MTTR must assume worst case sparing: if the least-likely-to-fail component fails in three out of four redundant units, what is the MTTR for the third unit? That would be the MTTR used for the availability of that component.

The same is true of technician availability: if three of four units are down, what is the worst case for the technicians? If one technician must replace two units, then it takes much more time than if another technician arrives within the assumed arrival time to assist, assuming of course that the spare for that third unit is available at the same time as the first or second failure.

A common problem in operating highly available systems is the tendency to believe that because there are four (or five or six) instances of a given piece of equipment, that it's okay for it to take a long time to repair it. Technicians might be busy on other calls, or someone forgot to order a replacement spare. Since there are three (or more) working systems, the tendency is to not be in a hurry to repair the broken unit.

This is a fallacy. The availability calculation is a function of MTBF and MTTR. If you assume a 4-hour MTTR, and you take two weeks to fix the broken unit, that will affect your MTTR. Unless you have a very large number of similar systems in the field, your MTTR will go up substantially when you add that 336-hour repair time to the calculation. Since most 5-9s systems do not have a lot of "cushion" in their availability calculation, the availability may not meet your Service Level Agreement. Every fault must be fixed within the MTTR used for the availability estimate. Of course, it is a mean, so it's not dominated by a single instance. But vendors must treat each failure as a critical failure and meet the stated MTTR. Customers must have visibility into the actual repair time so that they can have confidence that the contractual availability will be met. Vendor staff must treat each failure as if it was the penultimate failure that would cause the system to go down.

Another aspect of this fallacy is the effect of scheduled maintenance on availability. If you have four servers, and you used that number in the availability calculation, and you take one down for maintenance, then you only have three servers, and the probability that you can maintain a 5-9s availability goes down: you only need three failures to take the system down. You cannot expect to meet a 5-9s availability if you don't always have the stated number of parallel systems. There are two options: one is that the 5-9s SLA is reduced during scheduled maintenance. The other is that there are more redundant systems, such that when one is removed for maintenance, the remaining number is large enough to assure a 5-9s availability.

Any large vendor knows that maintenance activity quite often engenders failure, and while it's most often in the system that is directly undergoing the maintenance activity, it is not at all uncommon for other systems to be affected. Fat fingers, software bugs, and other effects often magnify the effect of what was thought to be a simple maintenance activity. Therefore, it is not advisable to squint and claim since maintenance occupies a small percentage of the measurement period, the effect of removing one of the redundant elements is not significant.

3.4 Increasing Availability through Redundancy

To achieve availability, we have “n” multiple systems in parallel and:

Parallel MTBF = MTBF x (1 + 1/2 + ... + 1/n)

There are two kinds of parallel redundancy: Active/Standby and Active/Active.

With Active/Standby, one instance is online and providing service, while redundant instances are on Standby. When the primary fails, one of the Standby instances takes over and provides service. Active/Standby affects MTTR and thus availability because the system is down until the failure is recognized and a standby instance is brought online.

Because of this, active/standby is rarely used in 5-9's systems but might be appropriate when only 3-9's availability is needed.

Active/Active systems employ multiple instances where all instances are online and available to provide service, where any instance can serve any transaction, call, etc. Active/Active is most used in 5-9s systems.

Since everything that can affect a system becomes part of the MTBF, the physical building, and its electrical, HVAC, and mechanical systems are part of the analysis. Weather is often a factor in a site staying available, and this means we almost always see multiple sites. It is important to make sure that the same severe weather event or other natural event (earthquakes for example) cannot affect all sites.

We often see deployments of two sites with two redundant elements of each type in a site, which claim to have 5-9s availability. This design originates in older carrier services where a single element (for example a frame relay switch) was providing a service, and getting the required MTBF and MTTR was possible in a carrier facility that had DC power, where the equipment worked over an extended temperature range and had sufficient shock resistance all of which was verified in special test labs. NG9-1-1 designs don't have such ability, although it is still possible to build more complex 5-9s systems using available technology. It is much more realistic to achieve 5-9s in three sites or other more hardware intensive configurations.

It is reasonable for a 9-1-1 authority to request a vendor "show their work" and provide details on how the stated availability will be achieved. However, purchasers should know that manufacturers often tightly restrict who can access their product's calculated MTBF, and the 9-1-1 Authority's vendor may not be allowed to reveal the MTBF numbers for some of the components they use. All quality vendors do MTBF calculations for their products, and refusal to supply MTBF is an indication that you have the wrong vendor. 9-1-1 Authorities must insist that calculations of availability of the entire path (typically from CO egress to PSAP ingress) be done and reviewed if any changes are made to the network. Where temporary changes are made, especially those which reduce diversity, recalculation of the system availability is required to assess the vulnerability of the system. This is especially important when planning conversions between systems, for example when transitioning from one NGCS operator to another or when upgrading a system or site. It may be acceptable to temporarily reduce availability for such changes, but they must only be permitted when calculations are performed that show the predicted availability during the conversion, and only with the 9-1-1 authority's knowledge of the lower availability.

3.5 Software Influence on Availability

Since nearly all the systems used in 9-1-1 systems are software based, and serious errors in the software can cause a system to not meet its availability SLA, it is very reasonable to ask how software figures into the availability calculation discussed above. While in some systems, an estimate of software defect occurrence is included in the MTBF numbers, since the MTTR of a software failure nearly always exceeds the allowable downtime, software is usually not included in the availability calculation, but software failures do count against the system availability metrics. Several sections of this document deal with how to minimize software failures.

3.6 Span of Availability Measurement

When evaluating availability, it is, of course, important to define the scope of the system to which the required availability is measured. In 9-1-1, the scope of the system to which we place a 5-9s availability requirement is from the egress of the telephone switch, or its equivalent to the ingress of the PSAP. Systems that rely on a single that rely on a single path or element (e.g., router or switch) within the Originating Service Provider (OSP) network or a single PSAP cannot achieve a 5-9s availability. For a legacy switch (e.g., tandem office switch or end office switch) in a NG9-1-1 system, the Legacy Network Gateway (LNG), or at least the protocol conversion portion of the LNG (the "PIF"), if connected locally to the switch can be considered part of the switch. However, if there is a network connecting the switch, an aggregation switch or multiplexer, another network, and the LNG (an "aggregation mechanism"), then these networks are part of the 9-1-1 system, and must be considered in the calculation of the availability of the entire NG9-1-1 system. The NGCS operator may be able to consider the characteristics of the aggregation mechanism in the calculation to achieve an overall 5-9s availability. Unless the entire 9-1-1 system including any such an aggregation mechanism can meet the overall required 5-9s, then such a mechanism may not be acceptable to the 9-1-1 Authorities.

4 Network Diversity and Redundancy

Networks used for 9-1-1 must be redundant and diverse. In terms of network design, there are four forms of diversity that must be considered: route diversity, supplier diversity, software diversity, and entrance diversity.

4.1 Route Diversity

IP networks that underlie NG9-1-1 systems have an important characteristic that differentiates them from most other network mechanisms; they automatically discover routes and use discovered routes based on efficiency metrics that the network learns. This is very different from other static network mechanisms that rely on pre-engineered backup paths. Because of this, the typical network designer does not have to anticipate all possible failures and design alternate paths for such failures. The network discovers paths itself. The routers on which IP networks are built determine any possible path that can get to the desired destination. The path may end up traveling through many routers, but if there is a path between two points, the network will discover it.

Failures are automatically accounted for, as an example, a path that used to exist, but no longer does, is automatically dropped. New links are automatically used. Once the routers learn what paths from that peer router exist, they are added to its local table of viable routes. As soon as packets are addressed to reachable entities on the new peer router, the local router will use that new link to reach them. Paths are continuously evaluated for how "good" they are based on packet performance, and the routers will automatically prefer "better" paths given a variety of choices. Every packet is routed independently of every other packet so that two packets in a row between two entities may travel different routes, especially if the metrics change.

IP networks are designed with graphs, where the routers are the nodes in the graph, and the links connect routers. There is almost never (with notable exceptions) a recognizable pattern in the graph. Connections are based on available circuits (e.g., fiber, wavelengths, Multiprotocol Label Switching [MPLS] Label Switched Paths [LSPs]), with different bandwidths and different endpoints. The goal is to create many paths between any two points, in the hope that in a failure, at least one path between two entities will survive, regardless of how many links or how circuitous the path is.

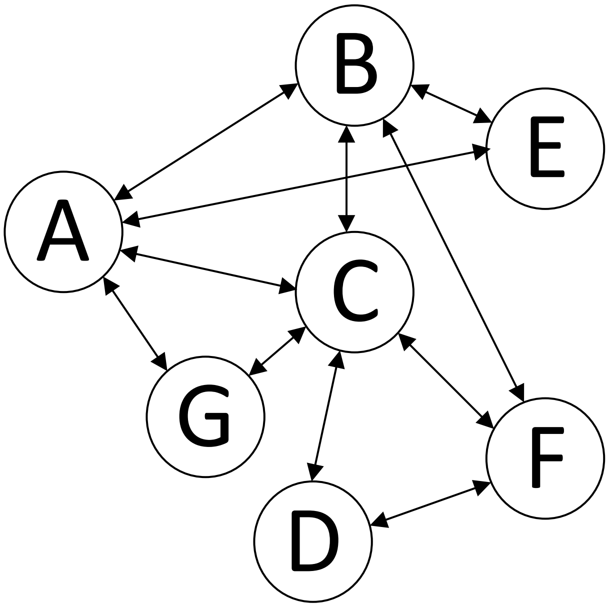

Figure 4-1 Example Network Graph

Figure 4-1 attempts to show nodes and links/interconnections within an IP network that do not seem to conform to any specific geometric pattern. The example illustrates the "as built," or real-world, arrangement of network elements and the links between them. This diagram is meant to show how networks should be documented with their geographic implications. Link costs may or may not be significantly affected by distance. The diagram shows how networks are built based on determining what is available and affordable in real life. The circles shown are meant to represent routers within an IP network and should include routers at a PSAP. The lines of interconnection represent links between routers, including those that are both logical and/or physical.

Figure 4-1 above shows a small network with between two to five links per node. A more realistic network would have many more nodes and ideally many more links per node. A minimum of three to seven links per node is recommended.

Some networks are split into "access/ingress," "core," and "termination/egress" networks. The core network is a graph, but the ingress and egress networks have a more recognizable pattern: a relatively small number of access routers are in the "core" router graph, but OSPs/PSAPs are directly connected to at least two access routers. This typically looks like a "star" connection with the access router as the hub and the OSPs/PSAPs as the spokes.

The reliability of the access network is much lower than the reliability of the core network. This is a design decision that prohibits the subscribers from being in the core network. Because there are a small number of access networks, and subscribers are only connected to a small number of them, a failure of a few access routers or links isolates the subscribers from the network.

NG9-1-1 networks need the highest possible reliability and should not be connected using a star configuration. Rather, the routers at the PSAPs, and ideally OSPs, should be treated as part of the core network and should have many connections to as many other routers in the network as possible. This includes links between routers located at a PSAP/OSP and routers at other PSAPs/OSP. The resulting network should look like an irregular graph, with 3-7 connections at every router, including PSAP/OSP routers. The connections should employ a variety of technologies, including, but not limited to, such technologies as MPLS, Digital Subscriber Line (DSL), wireless, cable modem, and satellite. If two PSAPs share a central office, there should be a connection between the PSAPs. For example, if they are on the same Synchronous Optical Networking (SONET) network, there should be a connection between them using that SONET network. If two PSAPs are on the same cable system, they should have connectivity via that cable network to each other. Rather than relying on a small number of core routers at data centers only, there should be additional core routers at concentrated interconnection facilities. These additional core routers are interconnected with the routers in the data centers and with each other. Anywhere that there are multiple paths to one or more nodes in the network, there should be a router with many connections.

Some of the connections, for example, cable modems and mobile network connections do not provide strictly point-to-point connectivity: they are public Internet connections. Virtual Private Network (VPN) connections can be used on such networks to improve the security of such connections. Multiple (3-5) VPNs should be provisioned across these public network connections so that even if the destination of the VPN experiences problems, other VPNs can function if traffic is passing through that public network.

4.2 Supplier Diversity

NG9-1-1 is more complex than legacy 9-1-1, which increases the difficulty of achieving reliability, but because it uses IP, more options have become available to improve diversity than were possible in E9-1-1 systems. Diversity may be difficult to achieve despite the availability of more options such as wireless communications.

PSAPs should have as many possible connections that are independent as possible. They should connect to every mobile network that is within range of the PSAP. If there is another central office or independent fiber network within a few miles of a PSAP, some kind of connection (e.g., microwave, short haul radio) should be made to a building that has connectivity to the other central office/fiber network.

Core routers must not have all connections from the same service provider. At least three service providers should provide service (if there are less than three, then perhaps the site for the router is the wrong place). A mix of technologies (e.g., DWDM, MPLS, SONET) should be used.

Caution must be taken when provisioning connections from multiple providers, especially in the "last mile"[1] to ensure they do not share some physical infrastructure. Service providers routinely use competitors' fiber or ducts to serve a wider area without the expense of trenching and cable pulling. For example, ordering IP circuits to get from Point A to Point B from two suppliers does not guarantee the circuits are not on the same strand of fiber. Supplier 1 may buy connectivity from Supplier 2.

It can be difficult to determine that such sharing of physical infrastructure occurs because the documentation of a path is hard to acquire and service providers are reluctant to share detailed engineering information on how their networks are constructed. All service providers have documentation on shared facilities, but it may be difficult to consolidate it into a meaningful description. It is also an issue when multiple providers use the same vendors since a software bug could theoretically affect all networks with the same software. Relying entirely or primarily on a single type of connection, such as all MPLS, should also be avoided, noting that in many service providers, a service like metro ethernet may in fact be provisioned on an underlying MPLS network.

4.3 Vendor Software Diversity

Currently 9-1-1 systems are typically not using diverse routers with diverse system software however future diversity configurations may include network components that might be a more robust design.

Networks should use a mix of at least two suppliers for routers: software failures in one brand of routers should not cause any loss of connectivity. Any given site should have at least two brands with the same function. Considerations related to maintenance and compatibility, testing, version control will increase both complexity and cost of such a design.

The term 'vendor software diversity' as used in this document refers to an approach that uses a combination of manufacturer, model, and software application versions of components that are unique in performing the same functions. For example, every NGCS

[1] The term "last mile" refers to the local connection from the service provider to the destination facility and is not intended to imply the actual length of such a connection.

and ESInet requires, minimally to operate, the use of at least one IP router within the core. A design with redundancy has at least two routers, both of which operate simultaneously (active-active). If both routers are the made by the same manufacturer, are the same model, and use the same version of the application, a defect in either the hardware or the application software may negatively impact both routers at the same time. If the concept of 'vendor software diversity' was applied to these routers, a defect would not create a single point of failure within the ESInet or NGCS core.

4.4 Entrance Diversity

The NG9-1-1 network should ideally be designed to remain operational during natural disasters or human-caused failures, including diverse entrances into a PSAP/ECC building.

There is a lot of emphasis on “no single point of failure” in NG9-1-1, and while redundant physical circuits are sometimes ordered, most PSAPs/ECCs do not have dual entrance facilities. So, the last mile, including the entrance to the PSAP/ECC, is almost always in the same conduit/trench. Backhoe related interruptions are a common cause of outages in the physical layer, but the cost of construction for dual entrances from the serving office/carrier facility may be prohibitive, or a physical alternate is not available. Entrance diversity may be achieved via a non-physical (wireless, satellite, microwave) method.

When contracting for NG9-1-1 ESInet Services, be aware that most providers leverage underground cabling and fiber facilities in newer construction to enter buildings. Whether your provider is reselling the existing access to the building, or they have built their own, most of the time they are using the same conduit from the street to the entrance to the building. If this entrance is compromised, both network connections can suffer the same fate.

Entrance diversity is achieved when physical facilities (e.g., circuits) are delivered through different areas of the building’s physical infrastructure. For instance, if the building has a circuit entering the west end of the building, another circuit would be ordered to enter through a different side of the building (e.g., south, north, or east). Often referred to as 'dual entry,' diverse entrance means that a secondary connection comes into the building in a different area than the first.

Planning for entrance connectivity should include different physical technology such as wireless, satellite, and/or cable providers with above-ground options for entrance diversity. Dual underground entrances and one above-ground represent a minimum recommended configuration. An ideal configuration would include multiple paths of in-ground and above-ground facilities (e.g., seven diverse paths and providers would be optimal). For two or more in-ground paths, each path should be separated by a distance that, given proper awareness and notification to authorities, reduces the risk that a single excavation adversely affects any other entrance facility.

4.5 Diversity Audits

Network Diversity Audits have been widely employed in traditional communications networks to help determine the perceived and actual engineered reliability of a communications network. When circuits and physical electromechanical connections, wires, pipes, conduit, and the like could be physically identified and followed, these audits helped to ensure the communication paths were physically separate and geographically diverse. They would help to assure network managers and owners that they were as secure as possible and removed from single points of failure that can be caused by a communications path sharing the same connectivity or conduit as other network components.

The traditional wired network diversity audit would demonstrate the disparate paths and connectivity that would provide confidence that a failure in one segment of the network, would not negatively impact the ongoing operations of other parts of the network and would help to assure delivery of the 9-1-1 call.

Diversity audits for legacy 9-1-1 networks were generally conducted by the 9-1-1 System Service Provider or the Incumbent Local Exchange Carrier (ILEC) at the request of the entity leasing or purchasing the network components such as the 9-1-1 Authority.

The NGCS operator may or may not be the provider of the network components in a NG9-1-1 system. The NGCS operator is responsible for diversity but may not be able to directly assess it. The underlying network provider may be the only entity who can conduct an adequate diversity assessment.

A diversity audit for legacy 9-1-1 networks would generally be conducted annually, and the results reviewed both internally by the service provider engineers and system administrators and shared with the network owner or lessee such as the state or county paying for the 9-1-1 service. Sometimes these audits were part of system service standards and/or contracts with the service provider. Reasons for performing these kinds of audits in legacy 9-1-1 networks include the requirement to ensure that diversity has not been negated due to network facility "grooming." Network grooming, which involves facility rearrangements, is typically done by the service provider as a regular part of their network management to seek the most efficient path for connections.

Diversity audits must be conducted on all paths. Just because a network facility is shared with other services, that does not relieve the NGCS operator from diversity audit responsibilities. For availability calculation, if a diversity audit cannot be conducted on a path, it should be assumed that the path is not diverse.

While IP networks ideally have many paths between two points, it is not unreasonable to determine that at least a minimum number of paths (including the most likely paths under normal circumstances) are redundant.

4.6 Service Provider Considerations

Despite challenges, there will always be components and network elements that can and should be part of a network diversity plan and therefore auditable to some extent (e.g., the "last mile").

Service provider(s) should endeavor to meet the needs of their clients by employing configurations that create robust and survivable networks. Many providers utilize networks from several providers to create networks that use a variety of connections from multiple sources. That does not mean, however, that the client can be assured they are purchasing a service that is truly diverse and adequately redundant. Network connectivity is not often disclosed, or easily discoverable, even if it were possible to provide this information to the client. In most cases in the NG9-1-1 environment, network designs do not deal with physical connections. The same technology allows for diversity in the number and types of providers and helps to ensure the most robust and reliable 9-1-1 networks. It also makes it difficult to trace a single communication path from one end to the other.

It is recommended that service providers demonstrate that they have a policy that governs their goals for diversity, processes that achieve those goals, and periodic audits that verify diversity conformance.

Audits may be performed solely as part of a service provider's policy or may be required by contract or SLA with the client.

4.7 PSAP Authority Considerations

PSAP Authorities, whether the contracting authority or acting as the system manager of their own system, should ensure their requirements stipulate the desired state of network reliability and survivability as well as response times, status reporting expectations, and after-action reports to include mitigation, if appropriate of future outage or service disruption situations.

The RFP for a system should clearly define the level of service the client demands as it relates to the performance of the network.

It is not enough to include the desired state of performance in the RFP. Additionally, at the contract development stage, the PSAP authority or their agent should ensure that SLAs are part of the contract and that they clearly define and outline the service level expectations. Another consideration that should be addressed is Telecommunications Service Priority (TSP). The PSAP authority should directly request the carrier for TSP on identified transport facilities that deliver 9-1-1 services.

See NENA Request for Proposal Considerations Information Document, NENA-INF-021 [2].

PSAP Authorities should review the provider's periodic audits to determine if the stipulated requirements are being met. This review would be done in accordance with the contract agreements for making audit information available. The contract should stipulate how proposed changes to diversity maintenance policies will be negotiated.

5 System Software/Hardware Release Procedures

It's not uncommon for a failure in any complex system to occur soon after the release of new software or a change in hardware. Any change to a stable system can cause issues. New software may have new defects, or procedures for installing new hardware or software may be lacking, may be wrong, or the common "fat finger" effect may cause problems. Controlling releases tightly in highly redundant systems can mitigate this kind of failure if procedures are established and followed closely that minimize the probability of a system failure or impairment. The Release Procedure includes a Release Process and the application of that Release Process.

5.1 Release Process

There must be a well-documented, strictly followed, release process. All the items in the following sections should be included in the release process. The release process must be written, reviewed, and widely disseminated. All participants in any aspect of the release process must be trained. Adherence to the release process must be mandated by senior management. The release process must include steps to handle unusual events so that they do not require abandoning the process when they occur. Finally, auditing of the release process must be completed regularly (e.g., quarterly).

5.2 Change Control Board (CCB)

There should be an organization that approves releases prior to implementing the updates. The CCB should be comprised of knowledgeable individuals primarily not directly part of the organization that does the release (so that its decisions cannot be overridden by common management). The CCB should review documentation submitted by the releasing organization that demonstrates that the process has been followed, such as testing, practice, roll back, and timing. CCB approval must be obtained before the first instance of a redundant system is updated.

5.3 Testing

Before any release, testing must be performed. Testing should be automated to the greatest extent possible to reduce the risk of unintended outcomes. Therefore, test suites must include test cases for normal and abnormal cases. If any failures occurred in the past, test cases for that failure and similar failures must be added to the test suite. It is recommended that each release include a substantial number of new test cases, with a goal of continuously improving the test suite.

For software, it is common to have "unit" tests (which test one module or file of code) and system tests (which test the entire system). Complex systems often have subsystem or service tests that test a functional element or similar subsystem. There are metrics available that assess the coverage of a test suite. For example, "code coverage" metrics measure how much of the code in a module is exercised by the test suite. The test process must specify what metrics are evaluated for each release, and what values for those metrics are acceptable. Minimum metric values should be evaluated at least yearly, and improvements mandated regularly as engineering skill levels and test suite size increases.

The software development process may sometimes be communicated to customers, but often the details of testing are not appropriate for customer release. A customer may reasonably request basic information from any vendor that allows them to assess whether the test processes used by a potential vendor are adequate for the task. Vendors should be willing to disclose basic information such as their basic test process, what metrics they track, what minimum standards they have for such metrics (e.g., code coverage), and how they approach regression testing (e.g., making sure known problems don't reoccur).

5.4 Continuous Integration/Continuous Deployment

"Continuous Integration," where automated code builds and test processes are performed for every code "check-in,"[2] should be part of every highly available system development process.

"Continuous Deployment" (CD), which is characterized by a highly automated, extensive test suite which allows releases to occur very frequently, even multiple times per day, is employed by many large software development teams. When code is "checked in" by a developer, the entire test suite is run automatically, including assessment of test coverage. If the test suite passes, the code is automatically deployed. The quality of code in many CD processes is excellent, and very good results can be expected if the test suite is comprehensive enough and trusted by the entire team and management. However, at the present time, CD is probably not appropriate for 9-1-1 systems, which need more stability, and more notice of pending changes than most systems do. Nevertheless, it is conceivable that given a sufficient test suite, an appropriate deployment strategy, and a cooperative customer base, CD may be appropriate for some systems in the future.

5.5 Method of Procedure (MOP)

A release requires a written Method of Procedure document that describes, in detail, how the release is accomplished. Each step undertaken by any staff member must be included

[2] In code writing practice an engineer would check out a code module to update it and would check it in when the update is completed.

in the MOP. It is essential that the MOP include every action required to accomplish the release, and frequent confirmation processes that confirm that a step was completed successfully should be included. The MOP must be comprehensive but can be broken into subsections that contain a subset of the complete MOP that a given staff member completes.

5.6 Practice

The MOP and the entire release process must be tested in a test environment that closely resembles the production release environment prior to the actual release. Where the system is deployed in more than one configuration, the release must be tested in each of them. Any deviation in expected behavior, any missed steps in the MOP, or any other problems must be documented, the MOP updated, and the practice repeated until the MOP is known to be good.

5.7 Roll-back

The MOP must include a roll-back process that returns the system to the same state it was in before the release was initiated. It is not acceptable for there not to be a roll-back for any aspect of a release. The roll-back must be practiced in the practice run in the test environment(s).

5.8 Maintenance Window

It is helpful to know in advance, the time when releases are to be performed. This is a "Maintenance Window," and typically occurs during expected low utilization of the system. The maintenance window should be defined as part of the contract, with changes subject to approval by the customer, although occasionally the end time of the window may need to be extended unilaterally by the vendor if a release process takes longer than the declared maintenance window. Emergency Maintenance (see below) is not confined to the maintenance window. Non-emergency releases should be limited to the maintenance window.

5.9 Notice

It is essential that customers and all support staff be aware of new releases so that any unusual behavior can be assessed in light of the new release. All on-duty staff in these organizations must be informed of the update. There should be a minimum of 30-day notice for non-emergency maintenance occurring outside the maintenance window (which should seldom occur).

5.10 Soak

Since it's known that systems often are unstable immediately after a change, highly redundant systems should be extremely cautious in rolling out a release to the redundant elements of a system. Initially, only one instance of a redundant system is updated, and it should be allowed to "soak" (i.e., observe how it works without changing anything else in that instance before updating another instance). Again, a soak period with no changes to that instance (or the initial instance), should be observed. Highly available systems typically are built with multiple instances per site and multiple sites. It is typical to update all instances at a site before updating any instances at another site. The soak period between instances at a site are typically short (a day perhaps for the first, and a few hours for subsequent instances at the site), while the soak period between the update of the last instance at a site to the first instance in another site is typically long.

This document recommends the initial soak period between sites be two to four weeks, the second be one week and the third and subsequent site soak periods be 24 hours. When observing these extended, between site soak periods, no other changes to the site that was updated should be permitted. Some overlap may be permitted, so that an update to a second site for a release on one system can occur during the soak period after updating the first site for a release on a different system although such overlaps should be avoided if practical.

Emergency Maintenance procedures may have dramatically lower soak times depending on the criticality of the update.

Soak periods are defined in engineering release procedure documents and are typically not subject to customer input.

5.11 Simultaneous Subsystem Version Differences

Software, libraries, and subsystems used by a highly available system must be able to operate with the above soak periods. That means that the software, library, or subsystem must work correctly when some instances are running different versions. This is most often a problem in code that synchronizes state between redundant instances. For example, a database must maintain a consistent state (content) in each of multiple instances. Call state replication mechanisms are another example. The software, libraries, and subsystems must be capable of handling the differences, including the back-out procedure.

Similar problems occur with schemas used in databases, interfaces, and the like: changes in the schemas need to be introduced gradually with the proper soak periods which means that the code must work with multiple versions of the schemas.

It is essential that the requirement to operate with different versions be maintained. Waiving the soak period should not be allowed under any circumstance. This often requires software to be written to expressly handle the differences. Generally, it should not be acceptable to have a "flag day" where some code is only executed when all instances are updated: the system should gracefully allow both versions to operate simultaneously, recognizing that new features may not be visible to users until all instances are updated.

There is some tension here in that "features may be not visible to users until all instances are updated" and "code only executed when all instances are updated" conflict. System design must endeavor to maximize the ability to exercise new code as soon as possible while dealing with the reality that a new feature cannot really be used until the soak is complete and all instances can support it.

5.12 Emergency Maintenance

When problems arise that are deemed critical to repair immediately, which includes critical security patches, releases may need to occur outside the maintenance window and with less than 30-day notice, and with greatly reduced soak periods. Emergency maintenance relieves the organization of the notice requirement and the restriction on occurring only in the Maintenance Window but does not relieve the organization of any other release procedures. In particular, the test requirements, MOP, practice, and roll-back requirements must be maintained. Emergency Maintenance procedures must be included in the process documents. Emergency maintenance procedures must be tested in advance and practiced regularly to ensure they can be completed efficiently when needed. Emergency Maintenance must not be used when the organization desires to avoid notice, timing, or soak period requirements. In all cases, Emergency Maintenance must be limited to unforeseeable, true emergencies. The Change Control Board must approve Emergency Maintenance the same as it does regular releases. The Change Control Board should be the arbiter of whether a release should be able to use the Emergency Maintenance procedures.

Testing of Emergency Maintenance Procedures may be accomplished, for example, by choosing a very low impact change to the code, such as a new feature that only affects a non-critical part of the system, and completing an emergency maintenance release of that code, with no prior notice to the internal teams involved. While it is important that all staff and customers are made aware that the test is indeed a test, testing should be completed following the same procedures that would be used in a real emergency. An "after-action" review of the release should consider how well the procedures worked, what changes might be needed, and what further training of staff may be required.

Emergency Maintenance is often necessitated by a failure and is incorporated in the Root Cause Analysis Plan (see Section 7.6 Root Cause Analysis).

5.13 Information for Purchasers

Most of the best practices described in this section are internal to a development organization. In some circumstances, purchasers, including organizations that bundle software in a larger system and end user customers, get some visibility into the procedures and practice of the development organization. It should be possible for a purchaser to obtain the current process documentation that describes how the development organization controls its software development process, and how discrepancies are handled. It is very reasonable when a customer has not received the availability it contracts for (SLA failure), that it receives a customer appropriate version of the root cause analysis. If follow up actions are needed, customers should receive commitments on when updates will be completed and should get timely reports on progress against those commitments.

Of particular concern in NG9-1-1 systems where the prime contractor does not develop portions of the delivered system, is that end user customers, through the prime contractor, should be able to obtain the process documentation of the organization that developed the system and receive root cause analysis of issues that affect its purchase. This should be addressed in the contract with the prime contractor (e.g., during the RFP process).

6 Service Level Agreement (SLA)

An SLA related to system availability, performance, and restoration requirements must be clear, specific, and measurable. Both the service provider and the customer should be able to clearly understand what the SLA requirements are, and whether they've been met.

It is recommended that service providers demonstrate that they have a practice that governs the contracted goals for diversity, and processes that achieve those goals. Periodic audits that verify diversity conformance should be conducted, results documented, and modifications made to achieve diversity goals.

SLAs vary based on expectations and therefore are not standardized in their content. At a minimum, SLAs should consider including the following:

- Detailed network designs that include clear demarcation points
- Expectations of the parties that define roles and responsibilities during and after an outage
- Established service provider/vendor and PSAP authority communication plans
- Defined notification process and responsibilities of the parties when an outage is discovered/experienced
- Documented processes to perform a timely review of an outage report, a Root Cause Analysis (RCA), and distribution of the RCA results
- Maintenance window process and schedule including reporting expectations
- Process for reporting and resolving operational events that cause impairment of a critical service or application which may result in a service outage

Detailed explanation of RFP requirements and SLAs can be found in the NENA Request for Proposal (RFP) Considerations Information Document NENA-INF-021 [2]. In the absence of specific SLAs regarding outages, the following FCC definition [3] applies:

- **47 CFR § 4.5(a)** that defines an "outage" as "a significant degradation in the ability of an end user to establish and maintain a channel of communications as a result of failure or degradation in the performance of a communications provider's network".

> See also 47 CFR § 4.5(e) that defines an outage that potentially affects 9-1-1 as an outage that meets at least one of four identified criteria.

Details of an outage in an Originating Service Provider's (OSP) network that prevents their customer's from completing a 9-1-1 call are not addressed in this document.

PSAPs and other Public Safety entities can report outages, file requests, ask questions, or check the status of requests through the FCC's Public Safety Support Center at https://www.fcc.gov/general/public-safety-support-center.

7 When Outages Occur

Outages will occur even in high-availability systems. Proper management of the processes used to identify problems, restore service, and mitigate risks can minimize the impact of outages on the ability to provide reliable emergency services.

7.1 System Failure Terminology

For the purposes of this document, Outage is defined as a lack of service continuity or unexpected suspension of operation.

For the purposes of this document, Impairment is defined as a partial lack of service, including a deterioration, diminished capacity, or weakening of operations.

All outages are impairments but not all impairments are outages.

7.1.1 Service Outage

A service outage can occur within any area of a network and is generally defined within 9-1-1 as the emergency calling service that is not able to deliver calls (e.g., signaling, media, or data) over an emergency services network.

A single emergency call that does not complete may not necessarily constitute an outage, but an ongoing accumulation of emergency calls that don't get delivered across a geographic area is categorized as a service outage.

Outages that exceed the FCC rules threshold (e.g., 900,000 user minutes) are required to be reported to the FCC and the affected agencies. The criteria apply to wireless, Voice over Internet Protocol (VoIP), and wireline equally as discussed in FCC 911 Reliability, Network Outages, and Related Reporting Obligations [4].

Outages that don't meet the FCC requirements for reporting should be used as examples for an internal service improvement process and may or may not be made evident to other stakeholders based on requirements within contractual agreements, such as an SLA.

7.1.2 Service Impairment

Service impairments might only affect a percentage of emergency calls, may be of one or more specific types of calls (e.g., wireless, VoIP, wireline), or may be constrained to a specific geographic calling area or call answering PSAP. Other circumstances in which a service is still able to complete emergency calls, though does so with degraded performance during or after a system failover, lacks complete call information such as missing call back number, incomplete, corrupted, or missing location data.

Systems that continue to function, but that may have some redundant elements that are non-operational or are operating in a degraded state, are considered impaired while operating at a higher risk with reduced resiliency and are likely to suffer additional operational impacts until the conditions causing the impairment are resolved.

7.2 Notification

During service outages or impairments, internal notification procedures result in established processes to restore service. In some cases, notification of service disruption is required to external parties and must be made to relevant communications providers and 9-1-1 public safety entities in a timely manner and based on FCC requirements for specific thresholds of service. The notification requirements contained in FCC rules are established to ensure that affected providers of the service(s) notify stakeholders in a timely manner and provide relevant information about 9-1-1 outages that could potentially impact the PSAP and those making a 9-1-1 call.

7.2.1 Initial Outage Notification

FCC Mandated

Qualifying communication providers (wireline, cable, satellite, wireless, interconnected VoIP, and Signaling System 7 providers) are required to report network outages that last at least 30 minutes and satisfy other specific thresholds in the Commission’s Network Outage Reporting System (NORS) [5]. Data submitted to NORS is presumed confidential.

Initial outage notification requirement thresholds and processes are detailed by the type of service and are based on amendments to Part 4 of the Commission’s rules concerning disruptions to communications and improving 9-1-1 reliability [6].

Initial outage notification requirement thresholds and processes are detailed by the type of service (e.g., wireless, wireline, VoIP) and are based on FCC rules that generally adhere to the following (e.g., for wireless), “All wireless service providers shall submit electronically a Notification to the Commission within 120 minutes of discovering that they have experienced on any facilities that they own, operate, lease, or otherwise utilize, an outage of at least 30 minutes duration.”

Check the FCC for rules applicable to other services such as wireless, wireline, cable, VoIP, and SS7.

Company Specific Process-Contract

Initial outage reporting processes should follow the requirements per FCC 47 CFR § 4.9 Outage reporting requirements—threshold criteria [7], and should be specified within the service provider/customer contract.

7.2.2 Initial/Final Outage Reports

FCC Mandated

An Initial Communications Outage Report includes a submission to the FCC that follows requirements per FCC 47 CFR § 4.9 Outage reporting requirements—threshold criteria [7], which may vary by the type of service provided and should be specified within the service provider/customer contract. For example, in the case of wireless, FCC requirements specify, "Not later than 72 hours after discovering the outage, the provider shall submit electronically an Initial Communications Outage Report to the Commission." Not later than 30 days after discovering the outage, the provider shall submit electronically a Final Communications Outage Report to the Commission.

Company Specific Process-Contract

Outage report processes follow FCC requirements that may also include some customer notification based on conditions within the service provider/customer contract.

There is currently an ongoing Notice of Proposed Rulemaking that may change Outage and Reporting Rules.

7.3 Conference Bridge Management

It is common for an internal conference bridge to be opened by a vendor(s) when an outage or serious impairment is in progress. To be most effective, attendees must be controlled, and possible causes or fixes identified.

7.3.1 Attendees

The golden rule for who is allowed/required to be on the bridge is: if you cannot fix a potential cause of the outage/impairment or diagnose problems (or provide the data that directly leads to a diagnosis), you should not be on the bridge. There are two classes of interested parties whom experience has shown are not generally constructive contributors: people who feel they must have up-to-date status, and managers of people who do need to be on the bridge (unless they too can diagnose or fix potential issues). Priority attendance should be restricted to those who can repair the issue and their management team. Attendees who cannot diagnose the problem, fix the problem, or allocate resources to resolve the problem should be discouraged from participating on the call. Also, there is an exception: most organizations have some people who are good at asking probing questions and understand complex systems. Sometimes it's helpful to have one or two such people who are not immediately part of the team that can diagnose or fix problems, but whose questions might spark the path to the solution. A separate meeting may be used for other stakeholders (e.g., legal, regulatory) to update status.

7.3.2 Leadership

It is recommended that the leader of the team that manages the NOC/SOC or equivalent for the product be the bridge manager, who should make sure the "right" people are on the call (and the "wrong" people politely excused from the call), and make sure that the conversation is focused, respectful, and moving forward. It is imperative to have another attendee designated as a scribe who keeps notes and a third attendee who is assigned to provide frequent status updates to people not on the bridge. It is essential that the status be updated frequently and completely to other members of the vendor's team and customers so that attendance can be limited to those who positively contribute to fixing the problem.

7.3.3 Cadence

The cadence of a successful event bridge is (after attendees are brought up to date on status and known symptoms):

1. Someone suggests a possible cause or requests data that might lead to identifying the cause.
2. A short discussion is held to see if the suggestion is plausible/would provide useful information (and can be obtained quickly).
3. Seek a volunteer(s), or assign resources, to look at the code or other source information to see if that could potentially be a cause (note not "THE" cause) and what fix or workaround might address the problem, or to get the requested data.
4. When a response from work assigned above is complete, results are presented to bridge attendees.

Often suggestions may seem highly implausible to some attendees, but each idea must have proper consideration.

Based on follow-up (item 3), the scribe notes who is responsible for the investigation, and when the response is expected to be brought back to the bridge, the scribe notes the results.

7.3.4 Extended Bridge Functions

It sometimes happens that an outage/impairment lasts more than a few hours, and not infrequently more than a shift, or even a day. In that case, people join and leave the conference call with new people replacing others in the roles discussed above. It is important for all attendees to review the scribe's notes to stay abreast of the discussion.

7.4 Temporary Fixes vs. Permanent Fixes

When outages occur, the primary goal is to restore the system to an operational status. In pursuit of such goals, it is not uncommon or unreasonable to deploy a temporary fix (i.e., workaround) that does not actually address the root cause, but rather avoids the symptoms or provides a way to quickly recover. These workarounds are often a necessary and appropriate initial response. However, with a workaround in place and the pressure off, it is often the case that the workaround is left by default as the permanent fix (i.e., long-term solution) to the problem, rather than doing the additional work to address the actual root cause.

Any workaround must be recognized as a temporary way to avoid outages, and it is essential that the actual underlying problem gets fixed in a timely manner. This is very often complicated by the fact that a permanent fix is costly and/or time consuming and often the tendency is to leave the temporary fix "just a little longer" while some feature or new capability is addressed. It is also often the case that other tasks take priority and delay a permanent fix. When a problem causes an outage, and especially when the outage violates a Service Level Agreement, the permanent fix must be developed and deployed promptly. It is sometimes helpful to have a non-negotiable deadline for fixing problems that cause outages, one that prohibits prolonged delays to develop and deploy the permanent fix.

7.5 Wrap Up

When the system is returned to an acceptable level of functionality, a wrap up session should be held before attendees leave the bridge. Wrap up should document who is responsible for follow-up activities including:

1. Writing the incident report
2. Conducting a Root Cause Analysis (see Section 7.6 Root Cause Analysis)
3. Reviewing short term engineering and/or operations tasks related to the incident
4. Sharing a version of the incident report with the customer

7.6 Root Cause Analysis

When problems within high availability systems occur, especially a service failure that causes an SLA requirement to be missed, one of the most important follow-up actions, taken well after the event is over, is a Root Cause Analysis (RCA). An RCA is a process that seeks to identify the actual underlying cause of the incident so that it can be remedied in the near term and will mitigate the likelihood of recurrence in the long term. Most often, RCA investigation results are presented and discussed in a designated review meeting. A final report of the findings, actions, schedules, and commitments should be prepared and distributed. It is imperative that at the conclusion of the RCA process, the customer understands what went wrong and what mitigation steps were taken to prevent a future outage from recurring.

7.6.1 RCA Meeting Attendees

To ensure more objective analyses of the identified event, the group that is responsible for the event must be included in the RCA. However, the majority of attendees at this type of investigative RCA meeting should be from the parts of the organization outside the specific group which appears to be initially implicated by the outage in order to prevent internal bias during discovery. Consideration should be given to having the person assigned to run the meeting be from a different part organization to encourage independence in the investigation. Representatives of the engineering or support organization should be included in the meeting to provide background and technical knowledge.

In some circumstances the service provider will invite the customer, or a technical representative on behalf of the customer, to attend the RCA meeting and be copied on any documentation shared at the RCA meeting.

7.6.2 Targeting the Problem(s)

One of the hardest aspects of RCA is to avoid blame while working through the process to fully understand the event(s) that may have occurred. Especially in high availability systems, the root cause is almost never a failure of an individual act (such as a programmer introducing a bug into code, or a "fat-finger" data entry operation). Rather, failure in high availability systems is almost always a systemic failure such as inadequate backups, inadequate code verification, and absence of failure limitation processes. The purpose of the meeting is not to assign blame, but rather to reveal the problem or the process, or lack of process, that allowed a minor event to result in a major outage.

7.6.3 Failure Identification-The "Five Why's"

One of the most important tools in conducting a thorough RCA is to repeatedly ask the question of "Why." Usually, the RCA starts with a superficial symptom, for example, a queue became filled with no more calls allowed in. It is important to know why the queue filled. If the answer to that was that dequeuing failed, then ask, why did it fail? There is a saying that goes: unless the question "Why" has been asked at least five times, the true root cause has not yet been revealed.

7.6.4 Which Incidents Get RCA

In 9-1-1 systems, an RCA should be required whenever an SLA requirement is not met due to an unplanned outage. If the event was unusual in any way, conduct an RCA.

7.6.5 RCA Timing

An informal (verbal) Reason For Outage (RFO) along with immediate steps taken to restore service should be expected within 48 hours of a service outage. An RFO is not a substitute for the more in-depth, formal RCA process described below.

For the RCA, an initial written explanation of the problem and the planned resolution should be expected from the provider(s) within 10 business days. Typically, a final RCA meeting is expected to be held a few weeks after an incident is over, when the responsible organization has had enough time to thoroughly investigate what happened, and so the appropriate parties are at the RCA meeting.

7.6.6 Reporting RCA Results

The approved RCA finding is distributed internally and externally as needed. RCAs are typically for internal company use only. An RCA report may be distributed outside of the company, and any approval to do so must be per the direction of the company's executive management and included in the contract with the customer. This may involve a separate meeting.

Examples of outside recipients might include a customer or a regulatory entity.

7.6.7 Failure Occurrences

Failure occurrences will occur even in high-availability systems. Proper management of the processes used to identify problems, restore service, and mitigate risks can minimize the impact of outages on the ability to provide reliable emergency services.

7.6.7.1 Software fault

Software faults are responsible for many large-scale outages in highly available systems. The specific fault that triggers the individual failure is rarely the root cause. Highly available systems must be tolerant of such failures. Usually, there is a systematic failure that underlies the fault. The software change process might be deficient (e.g., automated test coverage was set too low) or the code may have insufficient defensive capability (e.g., not implementing rate limiting). Often, prior events may have occurred without proper action taken to avoid future problems (e.g., a prior failure showed that one process could starve some other process, but no resultant effort was taken to limit the subsequent impact, only that the original fault that caused the initial impact was addressed).

7.6.7.2 Hardware fault

Highly Available systems rely on fault tolerant, redundant system and network design but it is often the case that the source (i.e., manufacturer) of hardware used in the design is not sufficiently diverse, allowing for a vendor specific hardware or software failure. Issues that can cause hardware faults include: aging component, environmental impacts to components, (i.e. dust, overheating, improper A/C capacity), lack of properly conditioned power, or excessive moisture/humidity.

7.6.7.3 Nature related fault

Natural occurrences such as storms, floods, earthquakes, tornadoes, hurricanes, or tsunami(s) are all examples of unpredictable events. The goal of highly available systems design is to mitigate any of these potentially anticipated events. Hardened facilities, comprehensive and adequate analysis of events in nature, and the estimated likelihood of these events occurring should inform the network design. Design elements that include geo-diversity, consideration for diverse climate domains for data center placement, appropriate site selection, and general risk reduction will help to reduce these types of failures.

7.6.7.4 Human related fault

Human related faults can be caused by many factors. The most common faults are:

- Inattention to detail, including not carefully following processes and protocols
- Lack of training or improper training
- Improper supervision and lack of checks and balances
- Data input errors (e.g.,” fat fingering” data)
- Lack of system restoration and verification procedures
- Missing, inadequate, or untested Integration and Back Out Plan (IBOP)
- Uncalibrated, inoperable, or missing (e.g., theft) crucial equipment
- Incomplete or inadequate maintenance documentation
- Insufficient design documentation
- Unmarked or mismarked cables, ports, or other facilities
- Lack of notification surrounding planned or unplanned system or network impairment, including advanced, current, or post notification(s) that work is scheduled, is ongoing, or has been completed
- Documentation changes that are not properly controlled (e.g., not adequately noted, updated, or disseminated)
- Insufficient or missing test environment (e.g., test facility is not currently available, does not exist, or does not adequately represent the required production environment)

7.6.7.5 Power fault

These types of failures might include errors due to unplanned electrical service supply circuit failure, changes or modifications to a facility's power distribution systems including failed grounding systems, failures in Uninterruptible Power Supplies (UPS), adjunct maintenance equipment such as a missing, lapsed, overcommitted, or exhausted multi-tiered certified/authorized back-up fuel sources, generators, or battery exercisers. Peripheral causes could include stale fuel sources, low quality supply resources such as batteries, and aging system components.

7.6.7.6 Cybersecurity attack

Cybersecurity vulnerabilities and insufficient monitoring/mitigation of cybersecurity threats can be a source of failures. Without an adequate understanding of how these threats occur and how to mitigate them once they do occur, systems can be particularly vulnerable. Related causes can include inadequate controls and training, inadequate or lack of mitigation plans, insufficient awareness and detection training, inattentive monitoring, misconfigured firewalls and intrusion detection devices, and a lack of cybersecurity test plan preparation/exercise.

7.6.7.7 Other sources

Accidental, unplanned, or miscommunicated service impacts, damage that disrupts communication facilities (i.e. damage to ESInet circuits, damage to data centers, fire, or natural gas leak).

7.6.8 RCA in High-Availability Systems

Root causes in high availability systems is almost always a process problem. The root cause of an outage in a high availability system is usually not the triggering event, but it is the fundamental reason why the entire system failed. Highly available systems are (supposed to be) designed to be resilient in the face of multiple issues that degrade performance. Fixing the immediate failure may prevent that specific event from happening again but may not prevent similar events related to the same root cause from occurring in the future.

Even though a troubleshooting process may identify an initial cause of failure, more comprehensive analysis should be conducted to ensure the root cause has been defined and addressed. The RCA should identify how the failure occurred, what caused the failure, and what steps have been taken to ensure that the failure does not occur again. The specific instance, or category of instances, that may be a symptom, is not necessarily the root cause. There might be more than one process that failed. For example, the software development process might be a root cause, but it also may be the engineering response plan that failed to bring engineering resources into the incident handling process early enough that caused the SLA requirement to be missed.

7.6.9 Follow Up

During the RCA process, every identified problem must be documented. Each time "Why" is asked, an answer is sought, and a problem is identified that must be addressed. Often there is more than one problem identified for each "Why." Each of these problems must be recorded in a problem register or list.

For each contributing fault, the problem register must include:

- responsible party(ies) for the underlying cause,
- responsible party(ies) for the repair,
- a commitment on a repair schedule, and
- an action plan, which should be updated as new items are identified. This should include:
    - prioritization of the action plan item,
    - status meeting schedule, and
    - test plans and schedule to demonstrate that the fault is corrected.

NOTE: The action plan may be developed after the initial RCA meeting and may be modified as additional information is discovered.

All "process" documentation that needs to be updated as a result of the RCA should be completed as soon as possible.

The repair schedule and commitment to timelines should be widely communicated to all responsible parties. Affected customers should be notified of RCA conclusions, action plan, testing schedule, and summary results when testing is completed. Customers should be notified that all actions that were identified in the RCA have been completed.

8 Managing subcontractors

The prime vendor is responsible for meeting the applicable SLAs and conforming to the best practices in this document. The prime vendor must also ensure their subcontractors conform to the best practices in this document. When selecting a subcontractor for the deployment of an NG9-1-1 NGCS/ESInet, the prime vendor should view subcontractors as an integral part of their company's responsibilities and seek a coordinated approach to provide a consistently positive customer experience and successful outcome.

The selected subcontractor should have established management processes and procedures to meet customer expectations and prime vendor quality levels that are consistent with contract requirements.

8.1 Prime Vendor Management Procedures

The prime vendor should have a documented subcontractor management process with, at a minimum, two key deliverables:

1. A Statement of Work (SOW) that
   - establishes assigned responsibilities for providing services,
   - clearly defines lines of communication and authorities for an effective channel of communication, and
   - specifies deliverables and communicates due date(s) and schedule.
2. A main project plan which integrates the subcontractor(s) including an understanding of where their efforts fit into the overall project.
   - Establish, define, and coordinate subcontractor procedures with the selected prime vendor methods.
   - Apply management controls to coordinate work and outputs.
   - Establish performance measures and assess subcontractor performance and quality.
   - Apply quality standards and processes consistently across metrics, to measure variances, and adjust processes or provide proper feedback to the subcontractor.
   - Provide periodic performance feedback.

8.2 Subcontractor Management Procedures

Subcontractors should have a legally binding, written contract with the prime vendor. The subcontractor organization will designate a single point of contact for coordinating contractual matters.

9 Best Practices and Standards Related to Redundancy and Diversity

Standards typically outline requirements that a system must conform to, while Best Practices incorporate approaches to implementation that should be followed. The 9-1-1 industry relies on both published standards and recommended best practices to help establish consistent service offerings and methods of operation. This standardization is necessary to successfully interconnect systems and networks for a seamless and effective communication network. These standards are developed by Standards Development Organizations (SDOs) and are created under a stringent set of processes, inclusive of contributions by industry participants, including collaborative and open discussion, adjudication of comments or changes, and voting guidelines. There are several organizations whose standards and/or best practices have guided 9-1-1 in matters of system redundancy and diversity.[3]

The following Best Practices apply to redundancy and diversity (and can be found on the ATIS Industry Best Practices website [8][2]:

[3] Examples include ATIS, IETF, NENA, APCO.

- 13-9-0532

  Network Operators and Public Safety should periodically audit the physical and logical diversity called for by network design of their network segment(s) and take appropriate measures as needed.

- 13-10-5267

  Network Operators, Service Providers, Equipment Suppliers, Property Managers, and Public Safety should ensure that operating procedures are clearly defined and followed by personnel during emergency situations in order to avoid degradation of cyber and physical security due to a diversion.

- 13-12-0731

  Network Operators, Public Safety, and Service Providers should provide physical diversity on critical inter-office and wireless backhaul routes when justified by a risk or value analysis.

- 13-12-0762

  Network Operators should engineer networks supporting VoIP applications including access to NG9-1-1 Next Generation Core Services (NGCS) to provide redundant and highly available application layer services.

- 13-12-3224

  Network Operators, Service Providers, and Public Safety should use dedicated and diverse Signaling System No. 7 (SS7) or Multi-Frequency (MF) controlled trunk groups as feasible and commercially reasonable as possible for the normal routing of 9-1-1 calls from originating switching entities to 9-1-1 Selective Routers (SRs) or Legacy Network Gateway (for NG9-1-1) rather than using shared Public Switched Telephone Network (PSTN) trunk arrangements and where appropriate and necessary supported by service level agreements. Network Operators, Service Providers, and NG9-1-1 PSAPs should use dedicated, geo-diverse, and redundant IP connection points when feasible & commercially available.

- 13-12-3258

  Network Operators, Service Providers, and Public Safety should design Emergency Services IP Networks (ESInets), where technically and financially viable, with redundant interconnectivity to PSAPs using the characteristics of IP routing to maintain connectivity in the face of extensive disaster damage. Public Safety ESInets may use diverse private facilities or their functional equivalent (e.g., MPLS, Generic Routing Encapsulation [GRE] tunneling, VPN, or equally secure industry protocols) and where appropriate and supported by service level agreements.

- 13-12-3259

  Network Operators and Service Providers should design networks with redundant interconnectivity to Public Safety Emergency Services IP Networks (ESInets) using the characteristics of IP routing to maintain connectivity in the face of extensive disaster damage. OSPs may use diverse private facilities or their functional equivalent (e.g., MPLS, VPN, or equally secure industry protocols) and where appropriate and supported by service level agreements.
- 13-12-3276

  Network Operators, Service Providers, and Public Safety should where feasible, provide both physical and logical diversity of critical facilities links.
- 13-12-3277

  Network Operators, Service Providers, and Public Safety should identify and manage critical network elements and architecture that are essential for network connectivity and subscriber services considering security, functional redundancy, and geographical diversity.
- 13-12-5249

  Network Operators, Service Providers, and Public Safety should consider geographic separation of network redundancy during restoration, and address losses of redundancy and geographic separation following restoration.
- 13-12-8728

  Network Operators and Public Safety should consider industry guidelines for logical diversity (e.g., multi-homing) and perform network diversification validation on a scheduled basis (e.g., twice a year). Processes and procedures should exist for tracking discrepancies and maintaining a historical record. This applies to Public Safety only in an NG9-1-1 environment.

10 Abbreviations, Terms, and Definitions

See the NENA Knowledge Base (NENAkb) [1] for a Glossary of terms and abbreviations used in NENA documents. Abbreviations and terms used in this document are listed below with their definitions.

Term or Abbreviation (Expansion)	**Definition / Description**
9-1-1 SSP (9-1-1 System Service Provider)	An entity that provides systems and support necessary to enable 9-1-1 calling for one or more Public Safety Answering Points (PSAPs) in a specific geographic area. A 9-1-1 SSP may provide the systems and support for either E9-1-1 or NG9-1-1. In the context of E9-1-1, it is typically, but not always, an Incumbent Local Exchange Carrier (ILEC). This includes: • A method of interconnection for all telecommunications providers including but not limited to the wireline, wireless, and VoIP carriers • A method and mechanism for routing a 9-1-1 call to the Public Safety Answering Point (PSAP) with no degradation in service regardless of the technology used to originate the call • A method to provide accurate location information for an emergency caller to a PSAP and if required, to other emergency response agencies • Installation of PSAP call handling equipment and training of PSAP personnel when contracted to do so Coordinating with PSAP authorities and other telecommunications entities for troubleshooting and on issues involving contingency planning, disaster mitigation and recovery
A/C (Air Conditioning)	A system that keeps air cool and dry.
APCO (Association of Public Safety Communications Officials)	The world’s oldest and largest not-for-profit professional organization dedicated to the enhancement of public safety communications.
ATIS (Alliance for Telecommunications Industry Solutions)	A U.S.-based organization that is committed to rapidly developing and promoting technical and operations standards for the communications and related information technologies industry worldwide using a pragmatic, flexible, and open approach.
CCB (Change Control Board)	A committee that consists of Subject Matter Experts and Managers who decide whether to implement proposed changes to a system or project.

Term or Abbreviation (Expansion)	Definition / Description
CD (Continuous Deployment)	An approach to software and product development that includes continuous integration, testing, delivery, and deployment. This combination of processes allows the team to iterate and automate much of software development.
DSL (Digital Subscriber Line)	A "last mile" solution that uses existing telephony infrastructure to deliver high speed broadband access. DSL standards are administered by the Broadband Forum.
DWDM (Dense Wavelength Division Multiplexing)	An optical fiber multiplexing technology that is used to increase the bandwidth of existing fiber networks. It combines data signals from different sources over a single pair of optical fiber, while maintaining complete separation of the data streams.
ECC (Emergency Communication Center)	A facility designated to receive and process requests for emergency assistance, which may include 9-1-1 calls, determine the appropriate emergency response based on available resources, and coordinate the emergency response according to a specific operational policy.
ESInet (Emergency Service IP network)	A managed IP network that is used for emergency services communications, and which can be shared by all public safety agencies. It provides the IP transport infrastructure upon which independent application platforms and core services can be deployed, including, but not restricted to, those necessary for providing NG9-1-1 services. ESInets may be constructed from a mix of dedicated and shared facilities. ESInets may be interconnected at local, regional, state, federal, national, and international levels to form an IP-based internetwork (network of networks). The term ESInet designates the network, not the services that ride on the network. See NG9-1-1 Core Services.
FCC (Federal Communications Commission)	An independent U.S. government agency overseen by Congress, the commission is the United States' primary authority for communications law, regulation, and technological innovation. The FCC regulates interstate and international communications by radio, television, wire, satellite, and cable in all 50 states, the District of Columbia, and U.S. territories.
FE (Functional Element)	An abstract building block that consists of a set of interfaces and operations on those interfaces to accomplish a task. Mapping

Term or Abbreviation (Expansion)	Definition / Description
	between functional elements and physical implementations may be one-to-one, one-to-many, or many-to-one.
HVAC (Heating, Ventilation, and Air Conditioning)	The use of various technologies to control the temperature, humidity, and purity of the air in an enclosed space.
IBOP (Integration and Back Out Plan)	A strategy designed to implement system changes with the ability to reverse those changes already made in case of failure or undesired results.
IETF (Internet Engineering Task Force)	The lead standard-setting authority for Internet protocols.
INF (NENA Information Document)	A document published to distribute information on a particular subject to the public safety community. Information documents may contain background information, best practices, checklists, and other material representing the collective knowledge and experiences of the NENA community. These documents do not contain normative statements and are not intended to be used to establish conformance requirements in procurement or development activities. The NENA INF Template may be downloaded from the Administrative Procedures & Templates Documents Page on NENA Workspace.
IP (Internet Protocol)	The method by which data is sent from one computer to another on the Internet or other networks.
IPR (Intellectual Property Rights)	A category of legal rights that includes patents, published and unpublished patent applications, copyrights, trademarks, and trade secret rights, as well as any intellectual property right resembling a member of the foregoing list as such right may exist in a particular jurisdiction.
ISP (Internet Service Provider)	A company that provides Internet access to other companies and individuals.
LSP (Label Switched Path)	A unidirectional path through the MPLS network. You can set up an LSP using any of the signaling protocols such as LDP, RSVP, or BGP.
MF (Multi-Frequency)	A type of in-band signaling used on analog interoffice and 9-1-1 trunks.

Term or Abbreviation (Expansion)	Definition / Description
MOP (Method of Procedure)	A documented set of step-by-step instructions that outlines the specific actions and sequence of tasks required to complete a particular process or operation.
MPLS (Multiprotocol Label Switching)	A type of data-carrying technique for high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table. A mechanism that allows network administrators to perform a measure of traffic engineering within their networks.
MTBF (Mean Time Between Failures)	The predicted elapsed time between inherent failures of a mechanical or electronic system during normal system operation. MTBF can be calculated as the arithmetic mean (average) time between failures of a system.
MTTR (Mean Time To Repair)	The average amount of time it takes to repair or recover from an issue or failure in a system, equipment, or process.
NENA (National Emergency Number Association)	NENA is referred to as The 9-1-1 Association, which is fully dedicated to the continued improvement and modernization of the 9-1-1 emergency communication system. NENA's approach includes research, standards development, training, education, certification, outreach, and advocacy through communication with stakeholders. As an ANSI-accredited Standards Developer, NENA works with 9-1-1 professionals, public policy leaders, emergency services and telecommunications industry partners, like-minded public safety associations, and more. Current NENA activities center on awareness, documentation, and implementation for Next Generation 9-1-1 (NG9-1-1) and international three-digit emergency communication systems. NENA's worldwide members join with the emergency response community in striving to protect human life, preserve property, and maintain the security of all communities.
NGCS (Next Generation Core Service)	The set of services needed to process a 9-1-1 call on an ESInet. It includes, but is not limited to, the ESRP, ECRF, LVF, BCF, Bridge, Policy Store, Logging Services, and typical IP services such as DNS and DHCP. The term NG9-1-1 Core Services includes the services and not the network on which they operate. See ESInet (Emergency Services IP Network).

Term or Abbreviation (Expansion)	Definition / Description
NOC (Network Operations Center)	An in-house or outsourced team of IT professionals responsible for maintaining a service provider system's technical infrastructure
OSP (Originating Service Provider)	A communications provider that allows its users or subscribers to originate 9-1-1 voice or nonvoice messages from the public to public safety answering points, including but not limited to wireline, wireless, and voice over internet protocol services.
PSAP (Public Safety Answering Point)	A physical or virtual entity where 9-1-1 calls are delivered by the 9-1-1 Service Provider.
PSTN (Public Switched Telephone Network)	The network of equipment, lines, and controls assembled to establish communication paths between calling and called parties in North America.
RCA (Root Cause Analysis)	The process of finding and analyzing the causes of a problem or an event impacting the value delivery of an application or an organization. RCA can find one or more root causes underlying a problem needing to be addressed to solve and prevent the problem from recurring.
RFO (Reason For Outage)	A verbal or written explanation of the cause that prevented a service from working normally.
RFP (Request For Proposal)	A standardized document for requesting negotiated bids. In more general terms, it is an announcement from a customer or funding source that is seeking proposals for a specific program, project, or work effort.
SDO (Standards Development Organization)	An entity whose primary activities are developing, coordinating, promulgating, revising, amending, reissuing, interpreting, or otherwise maintaining standards that address the interests of a wide base of users outside the standards development organization.
SLA (Service Level Agreement)	A contract between a service provider (either internal or external) and the end user that defines the level of service expected from the service provider. SLAs are output-based in that their purpose is specifically to define what the customer will receive.
SOC (Security Operation Center)	An in-house or outsourced team of IT security professionals that monitors an organization's entire IT infrastructure, 24/7, to

Term or Abbreviation (Expansion)	Definition / Description
	detect cybersecurity events in real time and address them as quickly and effectively as possible.
SONET (Synchronous Optical Networking)	A standardized digital communication protocol that is used to transmit a large volume of data over relatively long distances using a fiber optic medium. With SONET, multiple digital data streams are transferred at the same time over optical fiber using LEDs and laser beams.
SOW (Statement Of Work)	A detailed document that specifies the objectives and deliverables for a particular project or service contract. It lists the goals of the project as well as all the activities, deliverables, and potential schedule.
SR (Selective Router)	The Central Office that provides the tandem switching of 9-1-1 calls. It controls delivery of the voice call with ANI to the PSAP and provides Selective Routing, Speed Calling, Selective Transfer, Fixed Transfer, and certain maintenance functions for each PSAP.
SS7 (Signaling System No. 7)	An out-of-band signaling system used to provide basic routing information, call set-up and other call termination functions. Signaling is removed from the voice channel itself and put on a separate data network.
TSP (Telecommunications Service Provider)	A business that provides voice or data transmission services. These services are provided over a telecommunications network that transmits any combination of voice, video and/or data between users. A TSP could be, but is not limited to, a Local Exchange Carrier (LEC), a wireless telecommunications provider, a Commercial Mobile Radio Service provider, or a PBX service provider.
UPS (Uninterruptible Power Supply)	Provides backup power when the regular power source fails or voltage drops to an unacceptable level. Also known as a battery backup.
VoIP (Voice over Internet Protocol)	A technology that permits delivery of voice calls and other real-time multimedia sessions over IP networks.
VPN (Virtual Private Network)	A network implemented on top of another network, and private from it, providing transparent services between networks or

Term or Abbreviation (Expansion)	Definition / Description
	devices and networks. VPNs often use some form of cryptographic security to provide this separation.

11 References

[1] National Emergency Number Association. "NENA Knowledge Base Glossary." Updated August 31, 2023. https://kb.nena.org/wiki/Category:Glossary.

[2] National Emergency Number Association. *NENA Request for Proposal (RFP) Considerations Information Document*. NENA-INF-021.1-2020. Alexandria, VA: NENA, approved January 19, 2020.

[3] Federal Communications Commission. *Definitions of outage, special offices and facilities, 911 special facilities, and 988 special facilities*. 47 CFR 4.5.

[4] Federal Communications Commission. "911 Reliability, Network Outages, and Related Reporting Obligations." Updated December 20, 2022. https://www.fcc.gov/enforcement/areas/911-outages-reporting.

[5] Federal Communications Commission. "Network Outage Reporting System (NORS) and Reporting Requirements." Updated November 30, 2023. https://www.fcc.gov/network-outage-reporting-system-nors.

[6] Federal Communications Commission. *Second Report and Order In the Matter of Amendments to Part 4 of the Commission's Rules Concerning Disruptions to Communications (PS Docket No, 15-80), Improving 911 Reliability (PS Docket No, 13-75), and New Part 4 of Commission's Rules Concerning Disruptions to Communications (ET Docket No, 04-05)*. FCC 22-88. Adopted November 17, 2022.

[7] Federal Communications Commission. *Outage reporting requirements—threshold criteria*. 47 CFR 4.9.

[8] Alliance for Telecommunications Industry Solutions. "Industry Best Practices." Accessed May 5, 2024. https://bp.atis.org.

ACKNOWLEDGEMENTS

The National Emergency Number Association (NENA) NG9-1-1 Operational Diversity & Redundancy Working Group developed this document.

NENA recognizes the following industry experts and their employers for their contributions to the development of this document.

Members	Employer
Dan Mongrain, Eng., Systems Security & Resiliency Committee Co-Chair	Motorola Solutions, Inc.

Members	Employer
Raymond Paddock, Systems Security & Resiliency Committee Co-Chair	Synergem Technologies, Inc.
Nancy Pollock, ENP, Co-Chair	911 Policy Consulting, LLC.
Steve Deloach, Co-Chair	Lumen Technologies, Inc.
Rick Blackwell, ENP, Notes Manager	Director, Marietta, SC, Retired
Tom Breen, ENP, Document Editor	SecuLore, an Exacom Company
Sarah Cook	T-Mobile, US, Inc.
Greg Denton	9-1-1 System Manager, State of Arizona
Pete Eggimann, ENP	Eggimann Technology Services, LLC
Pierre Foucault	Quebec 9-1-1 Agency, QC, CA
John Harding	Palm Beach County Sheriff's Office
Roger Marshall	Comtech
Fran Moore	City of Greenville, SC
Heather Musolff, ENP	Brevard County, FL
Theresa Reese	Ericsson, Inc.
Sarah Rollins	Consultant
Wendy Rooney, ENP	Spartanburg County SC
Brian Rosen	Brian Rosen Technologies, LLC.
Jeanne Sentz	Palm Beach County Sheriff's Office
Jason Sweney, ENP	City of Xenia, OH

Special Acknowledgements:

Sandy Dyre, ENP, Committee Resource Manager, and Cheryl Benjamin have facilitated the production of this document through the prescribed editing and approval process.

The Systems Security & Resiliency Committee, NG9-1-1 Operational Diversity & Redundancy Working Group is part of the NENA Development Group that is led by:

- Wendi Rooney, ENP, and Lisa Dodson, ENP, Development Steering Council Co-Chairs
- Brandon Abley, ENP, Chief Technology Officer
- April Heinze, ENP, VP, Chief of 9-1-1 Operations