BURNS McDONNELL

WHITE PAPER

Monitoring and Detection of Anomalies on Power Network Equipment Via SNMP and Syslog Can Pay Dividends

By David Eckert, Roger Luechtefeld, PE and Dustin Williams, PE

Electrical distribution system equipment can be at risk of failure without the monitoring, detection and notification properties offered by simple network management protocol (SNMP) and system logging protocol (syslog). Unfortunately, a lack of monitoring has resulted in system failures for some utilities.

A proactive approach for utilities is to monitor and detect potential substation equipment failures. Simple network management protocol (SNMP) is an internet-standard protocol that enables the collection, organization and modification of information about devices on networks. System logging protocol (syslog), the second key part of this monitoring and protection setup, sends event data logs to a central location for analysis and reporting. Both are proven monitoring tools that can provide the early warning needed to avert outages, equipment damage or injury to personnel.

Without the monitoring capabilities provided by SNMP and syslog, critical equipment could fail with little to no notice to network operations. Without these monitoring capabilities, temperature increases that could cause fires in substations might only be detected through supervisory control and data acquisition (SCADA) systems, often far too late for local emergency responders to arrive to put out a catastrophic fire.

Deploying SNMP can be complex and requires many skill sets to be implemented successfully on a network. This document will review and expand on the SNMP protocol by covering the aspects of object identifiers, trap collectors, SNMP versions and planning needs surrounding the real-world deployment with Ameren. SNMP can provide environmental notifications through device monitoring, security monitoring, device configuration, equipment management, and notifications of issues with the network device.

Ameren Protection Project

Ameren, a holding company with Missouri and Illinois utility operations along with electrical transmission operations in the Midcontinent Independent System Operator (MISO) region, worked with Burns & McDonnell to perform an SNMP rollout on corporate and operations equipment. Ameren has network management systems and is implementing SNMP for early warning and prevention of network conditions that often start as small, nearly undetectable events but can quickly escalate into widespread network outages if left unchecked.

The goal of this project was first to implement SNMP on devices and then work with the event management team to classify SNMP traps, organize patterns for trending alarms, and present data in a manner that enables the digital command center (DCC) to identify a network issue and work to resolve that issue before it becomes a severe issue. This project implemented SNMP on Cisco routers and switches, Nokia routers, Nokia dense wavelength division multiplexing (DWDM) equipment, and F5 load balancers.

Object Identifiers

The first step to deploying SNMP is to work with the vendors to get the management information database (MIB) files. An MIB is a structured database that includes objects that hold information about the different parameters and settings of network devices. These objects — called object identifiers (OID) — use a sequence of numbers to differentiate each OID. For example, the first part of an OID is an assigned number, such as 1.3.6.1.4.1.9 is an OID for Cisco and represents the Cisco organization. This OID is assigned by organizations such as the Internet Assigned Numbers Authority (IANA). The IANA registers this OID as a private organization. Using the OID 1.3.6.1.4.1.9 as an example, the different sections of the OID represent the information, as shown in Figure 1.

OIDs are assigned more values separated with periods to represent different data points. OIDs are hierarchical and organized in a tree structure for ease of navigation. Each tree level corresponds to a specific group of objects ranging from system settings to the state of protocols or ports. OIDs have different types, such as bits, integers or strings. OIDs are primarily used to monitor and retrieve information from network devices by using the SNMP GET and GETNEXT operations to retrieve current information about the device. Typically, OIDs are read-only, but the SNMP SET operation can configure a device by sending a new value for an OID, and the system will change the setting. However, proper access rights need to be configured to work correctly.

The MIB file for each particular model of equipment is different so it is important to have the vendor MIB for each. It is also important to know that multiple MIB files can be used for one device.

Typically, MIB files are split into different functions. For instance, one MIB file may be for functions shared by the same device class, such as routers with the same features, typical system settings, or hostnames. There are also unique MIB files with a set of OIDs for a particular feature, such as a GPS receiver only available on a specific device.

OID Section	OID Example Value	OID Section Number Meaning
1	1	ISO (International Organization for Standardization) assigned
2	3	Is an ORG (Organization)
3	6	Usually assigned to the DOD (Department of Defense)
4	1	Represents the internet
5	4	Designated for private enterprise use
6	1	Other identifiers
7	9	Private enterprise number (9 is Cisco)

Figure 1: Object identifiers (OID) use numerical classifications for various network device settings and parameters.

Making the decision on which type of OID to enable without investigation can result in enabling OIDs that are for features that a utility may not use or miss enabling OIDs that are critical features for network monitoring. Once the MIB files are obtained, OIDs are selected, and they must be classified as enabled and given a severity level.

MIB files can range from just a few OIDs to hundreds. Enabling all alarms could be counterproductive, especially since someone would have to enter them into a trap collector, discussed in more detail below. The next step is to assign each alarm a severity level. A common approach is to classify the alarms as critical, major, or minor. For example, critical alarms are conditions that can cause a communications outage, such as a routing protocol in the "down" state on a router. A loss of redundancy is an example of a major error, assuming the device has a redundant power supply, and there is time to replace the failed unit with no outage. A minor error is an event that needs to be addressed, but will not likely affect the network's operation.

Syslog servers receive information about unauthorized access attempts and information about changes on a system. This information can be critical for finding the root cause of network issues when a network change could be the culprit. Syslog servers receive and store log messages from devices on the network. Log messages contain a time

stamp, the host name of the device that sent the message, the severity level, facility and message content. The messages are sent sequentially and can be used as a log for future processing if an event requires security analysis or an audit trail. Real-time processing of syslog messages would offer real-time visibility into the network through a security information and management system; however, these systems may require specific integrations with different devices and are typically not used for condition alarming like SNMP.

Trap Collectors

Network devices use SNMP by sending traps to centralized servers called trap collectors, trap receivers or network management systems, depending on the system's capabilities. A trap collector sends queries to the network device to read specific data points (SNMP polling). Traps are asynchronous notifications, known as unsolicited messages, sent as soon as a change in status is detected. In some cases, the traps must be enabled on the device, as some settings default to not sending a trap on an OID change.

There are two types of traps: generic and specific. Generic traps are general types of traps that relay information about an event, while specific traps include detailed information about the events. Traps are sent encoded in ASN.1 (Abstract Syntax Notation One), a standard interface description language that defines data structures for use in networking.

The trap collector receives and processes SNMP traps sent by network devices. When selecting a trap collector, it is important to understand its capabilities, what SNMP versions are supported, and the difference between SNMP versions.

SNMP Versions

The earliest version of SNMP — SNMPv1 — has since been replaced by SNMPv2 and SNMPv3. SNMPv1 fulfilled the requirements for its time, but as the world of networking and the internet matured, security concerns became the driving force behind further modifications that led to SNMPv2.

SNMPv1 does not provide strong authentication or encryption, and this is the main security concern. SNMPv1 uses a community string sent in clear text. This community string is the string of text required for authentication and authorization to read the SNMP trap and write to the device. Since the community string is in clear text, each User Datagram Protocol (UDP) is susceptible to interception. Once intercepted, any packet decoding tool can read the password and gain the ability to change settings on devices using SNMPv1 with the exact community string.

SNMPv2 aimed to improve security features, but these features were not significant enough to solve the issue SNMPv1 possessed, the community strings. The most important problem with SNMPv2 was the lack of strong authentication and encryption. Even with the security concerns, SNMPv2 still managed to make SNMPv1 obsolete. The security enhancements include implementing community-based security models; other enhancements include bulk operations and improved error handling. The community-based security models introduced read-only and read-write communities, allowing better control over the protocol operations. This approach contrasts with SNMPv1, where only one community string provides limited access control. The primary security concern is still community strings since the password is sent in a plain text field, making it vulnerable to being easily intercepted and read with any packet decoder.

Another improvement on SNMPv2 was the Get-Bulk operation. The Get-Bulk operation allowed a trap collector to request many OIDs in a single request instead of sending Get-Next for each query and creating multiple requests. With the Get-Next operation, security concerns arise because it gives access to man-in-the-middle attacks that can hijack sessions and eavesdrop to gain information about systems. It can also provide information to assist intruders to use enumeration attacks to gain sensitive information about the systems or deploy resource exhaustion attacks that could overload a device.

Error-handling improvements enabled better troubleshooting by identifying the root issue faster through more detailed error codes. SNMPv2c is a simplified community-based version of SNMPv2. The main difference is the removal of some protocol data units (PDUs) and a change in message format to better suit less-complex networks.

SNMPv3 is the only version with security features and should be considered over previous versions. It offers more security features than any SNMP version and introduces the User-Based Security Model (USM). The USM provides authentication and access control to reduce the possibility of a session hijack or eavesdropping and provides the creation of individual user accounts. Authentication in SNMPv3, specifically through the hash-based message authentication code (HMAC),

supports multiple cryptographic hash functions, such as the message-digest algorithm (MD5) and the secure hash algorithm (SHA). Additionally, SNMPv3 introduced DES (Data Encryption Standard) and AES (Advanced Encryption Standard) encryption capabilities to reduce or eliminate the risk of data compromise. SNMPv3 is also backward compatible with all previous versions of SNMP. Backward compatibility enables networks to transition to more current versions of SNMP without compromising network loss of services.

Proper Planning is Key

Implementing SNMP on large networks can pose challenges, with security being the foremost concern. SNMPv3 offers the most secure implementation of SNMP, addressing security issues related to reading or writing MIB tables. Unauthorized writing to network devices can lead to behaviors that impair critical infrastructure or disable security controls and access to substations. Modern transport devices enhance security by allowing configuration to direct SNMP traffic to the intended trap collector. Only trusted IP addresses can communicate via SNMP under this configuration. This effectively acts as an access control list (ACL) and prevents unauthorized IP addresses from sending or receiving SNMP traffic.

Performance impacts should also be considered during a wide-scale implementation. When planning for network traffic control, marking packets for quality of service (QoS) is essential. Where the packets are prioritized in the classification table of the QoS design is dependent on other applications. Applications such as relaying, SCADA, timing, phasor measurement units, and security are among the most critical ones. Depending on the bandwidth available to the transport network, prioritizing applications may pose no threat of dropping SNMP traffic. If the network is congested, it could pose a threat of dropped SNMP traffic.

Polling frequency of the trap collector determines how often the network nodes are polled for information. The frequency will directly impact both additional traffic on the network and compute required to process more messages. Similarly, the amount of data requested means more data to send and process. Routing protocol information, as one example, can carry significantly more data than something simple like port status or port traffic counters. Polling once per day to pull the entire MIB is a very different load on the network.

Properly sizing a trap collector with adequate interface capacity, storage, and compute requires an understanding of several factors. The system's sizing (capacity) needs to account for the number of nodes operating in the network, how often they send messages, and the size of each message. A 10-node network, polling devices once per minute, pulling a few specific MIB entries with SNMPv2 will be drastically less load on the network than 10,000 nodes, polling every 10 seconds, pulling the whole MIB and using SNMPv3. This is only part of the picture, as SNMP traps sent to the collector when an event occurs will be asynchronous from polling and require additional computing in the trap collector to process.

Real-World Example

To facilitate secure and efficient network management, Ameren has implemented SNMPv3 for its devices. Ameren's devices send traps to an OpenText collector, which classifies them on receipt. The following figures show three views that Ameren uses to monitor network health. These were developed to show information quickly so each responsible department could identify the information needed for its function.

This first pane in Figure 2 is what the DCC views 24/7. These are the alarms for every device on the network. Along the top section of the pane are alarm classifications that give a value of active alarms. These alarm groups help prioritize actions taken by the DCC. The middle section lists individual alarms, including the time the alarm was last received, hostname of the device that created the alarm, the title that describes the alarm, the object type and the severity The bottom section ties alarms to incident tickets. The built-in automation system creates a ServiceNow ticket and pushes it to the system when specific alarms are generated, allowing the DCC to act on the alarm. This automation has improved reaction times and reduced outages by getting on top of issues as they occur.

As shown in Figure 3, the page is nicknamed the Donut Shop. The Donut Shop sorts each business technology category and allows the technology lead to determine the state of the devices using the list of alarm types below the donut. Each donut gives a quick view of the health and the different slice sizes, representing the number of each specific alarm. Technology leads can see the state of their technology and know where to focus their efforts to work toward an error-free environment.

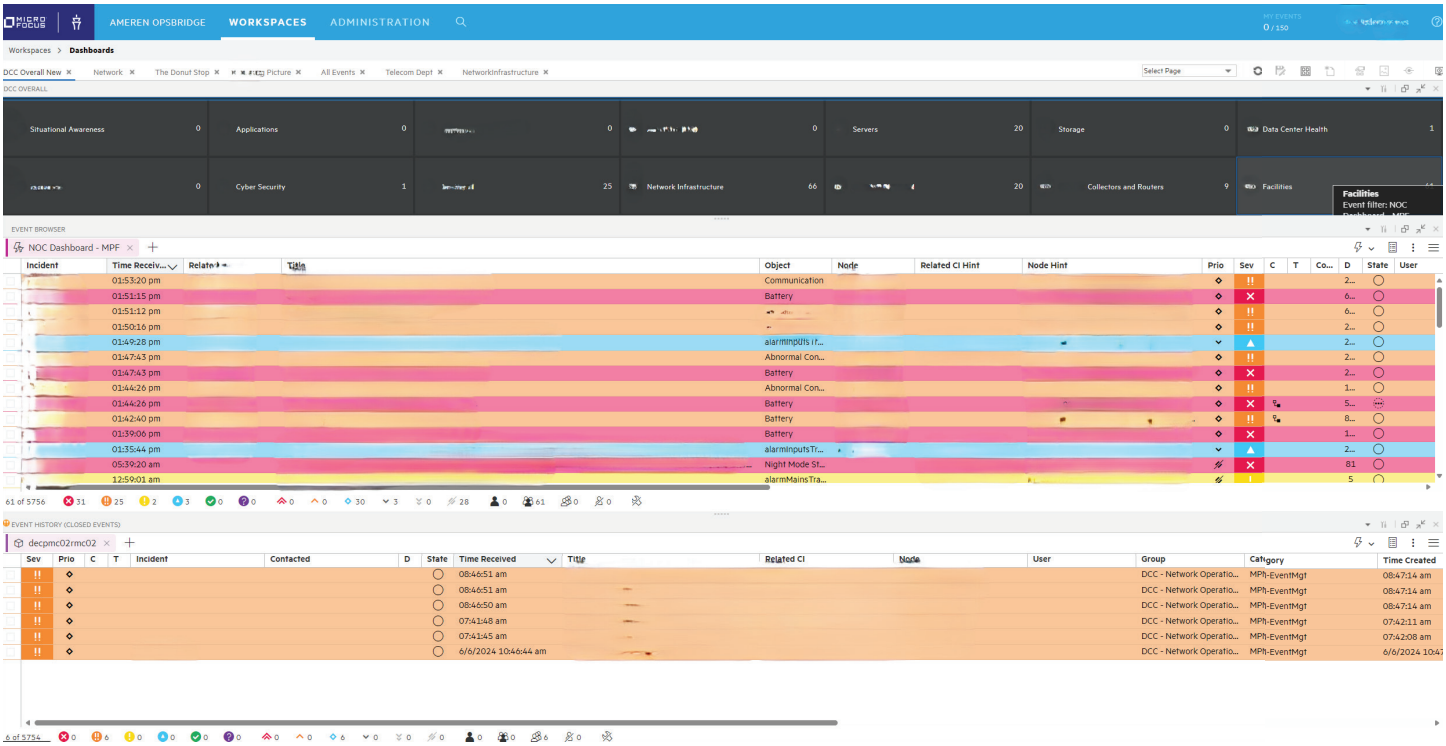

Figure 2: An example of a pane (with redaction) viewed by Ameren operators indicating network health.

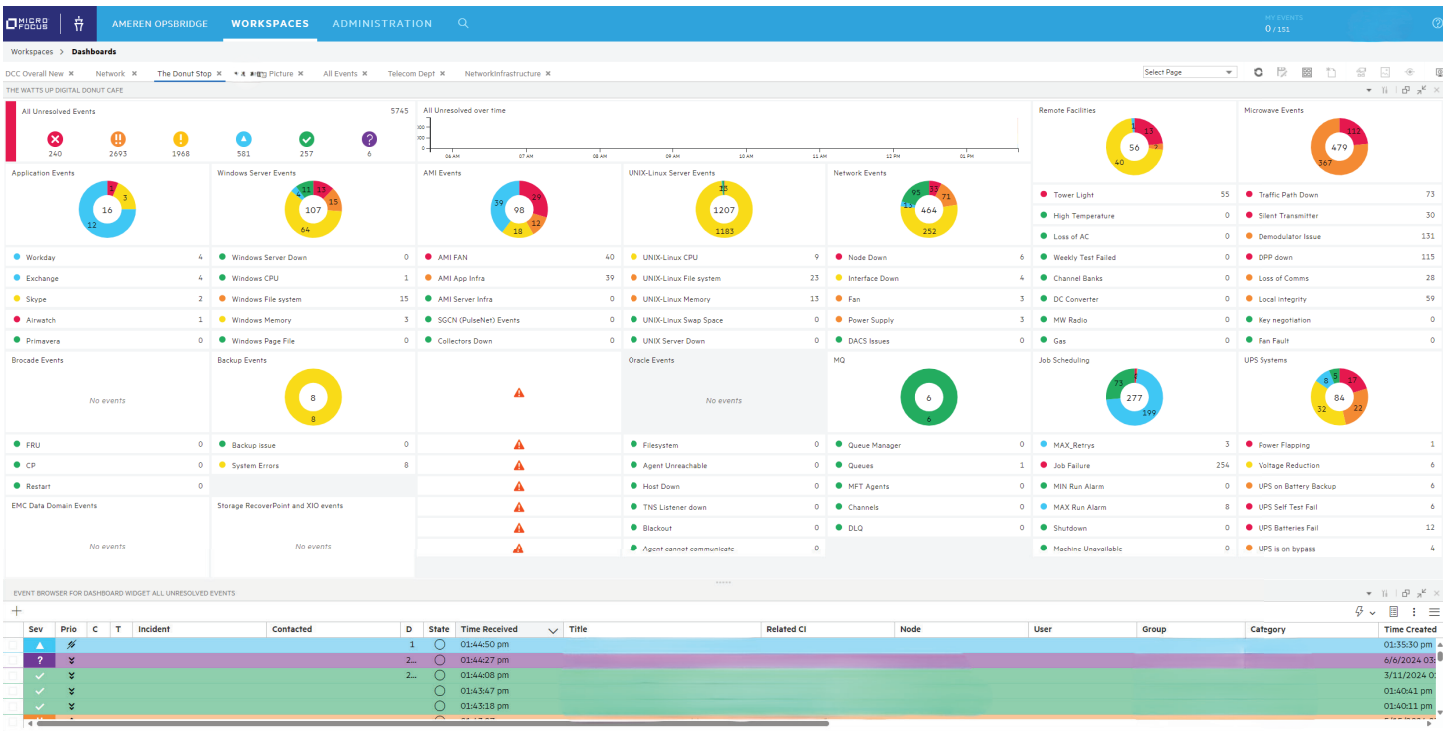

Figure 3: An example of the Donut Shop pane visually representing alarm types.

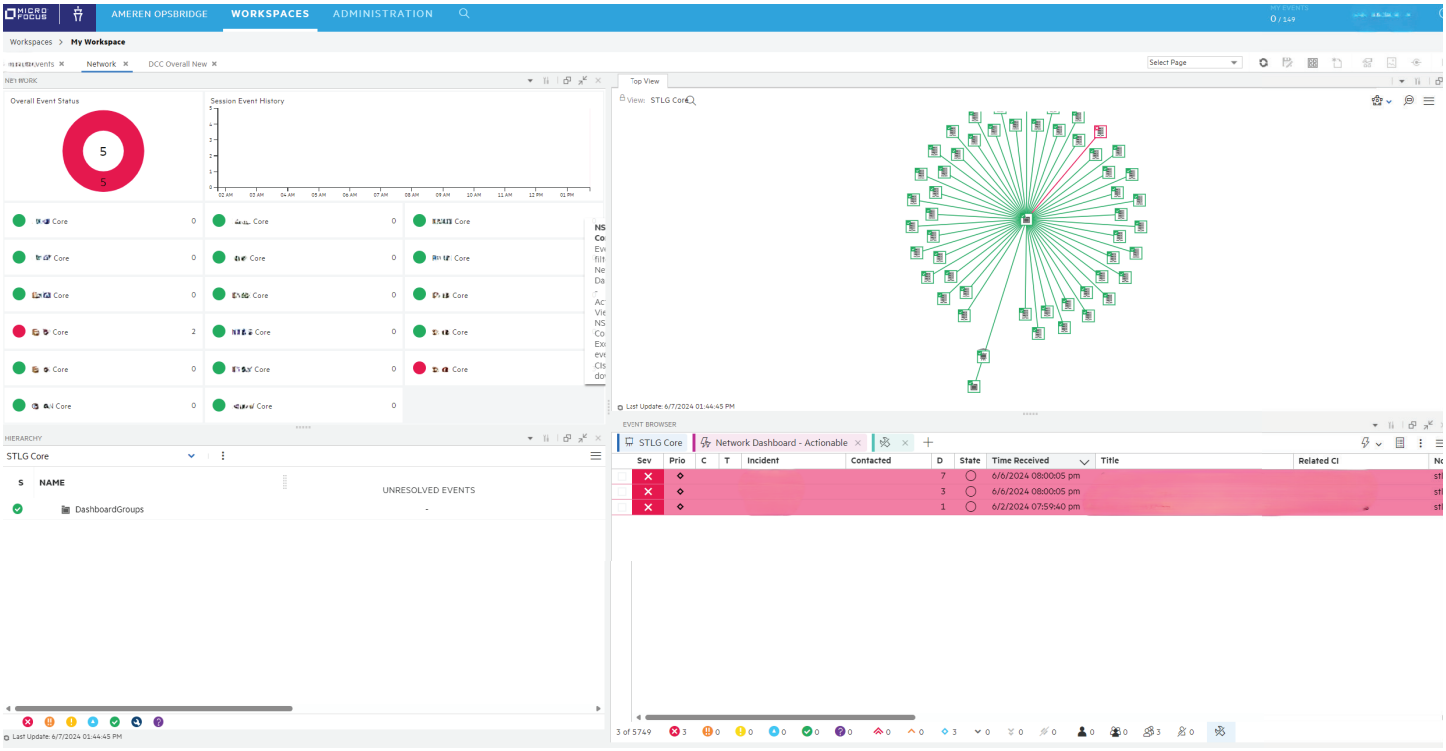

***Figure 4:** An example of the pane (with redaction) providing an executive view of the status of core network sites.*

This last pane in Figure 4 is what executive leadership uses to gain oversight of the state of each core location at a glance. This view provides a high-level view of the core sites with status: Red means there are active issues at the site, and green means they are operating as expected.

A Proven Protocol

SNMP provides network operators with quick access to automated alarming without the need to deploy additional devices. Understanding how SNMP works and the different versions is an important step in gaining the needed oversight for protecting critical equipment. Deployment can be complex, requiring a variety of skill sets to implement it correctly, so it is important to have an intentional design and plan.

Once SNMP is deployed on the network, the systems that receive traps can work in conjunction with other systems to automate responses, such as creating incident tickets for personnel to address. When properly configured, the SNMP provides notifications through monitoring security, device configuration, equipment status, and notifications of issues from the network. This provides a new level of network visibility that will be vitally important as utilities continue to navigate through ways to increase grid reliability. Without the monitoring capabilities provided by SNMP and syslog, critical equipment may run into failure instead of scheduling a timely repair.

Definitions of Terms

ACL — access control list

A list of rules used to control access to network resources by specifying which users or system processes are granted access and which are denied.

AES — Advanced Encryption Standard

A symmetric-key encryption algorithm adopted by the U.S. government to replace DES.

ASN.1 — Abstract Syntax Notation One

A standard notation used for describing data structures and encoding rules used in telecommunications and computer networking protocols.

DCC — digital command center

A centralized location from which IT professionals monitor, manage and maintain a company's network infrastructure. The DCC is responsible for maintaining network uptime, troubleshooting network issues, performing routine maintenance tasks and coordinating responses to network emergencies or outages. Typically staffed 24/7, DCC personnel use various monitoring tools and technologies to proactively identify and address network problems, to minimize disruptions and optimize network performance.

DES — Data Encryption Standard
A symmetric-key algorithm for the encryption of electronic data.

DWDM — dense wavelength division multiplexing
A technology used in fiber-optic communications to increase bandwidth by multiplexing multiple optical carrier signals onto a single optical fiber.

EGP — exterior gateway protocol
A protocol used to exchange routing information between different autonomous systems on the Internet.

HMAC — hash-based message authentication code
A mechanism for calculating a message authentication code (MAC) involving a cryptographic hash function in combination with a secret cryptographic key.

IANA — Internet Assigned Numbers Authority
A department of ICANN responsible for coordinating the global internet's systems of unique identifiers.

IGP — interior gateway protocol
A routing protocol used within an autonomous system (AS) to exchange routing information between routers.

MD5 — message digest algorithm 5
A widely used cryptographic hash function that produces a 128-bit (16-byte) hash value.

MIB — management information base
A database used for managing the devices in a communications network.

OID — object identifier
A unique identifier used in SNMP to identify managed objects in the MIB.

PDU — protocol data unit
A unit of data that is transmitted between network entities in a layered network architecture.

QoS — quality of service
A set of technologies and mechanisms used in computer networks to manage traffic and ensure a certain level of performance or service quality.

SCADA — supervisory control and data acquisition
A control system architecture used in industrial automation to control and monitor processes, infrastructure, and facilities.

SHA — secure hash algorithm
A family of cryptographic hash functions designed by the National Security Agency (NSA) and published by the National Institute of Standards and Technology (NIST).

SNMP — simple network management protocol
A protocol used for network management and monitoring of network-attached devices.

Syslog — system logging
A standard for computer message logging that allows different devices and applications to generate, forward and process log messages.

UDP — User Datagram Protocol
A connectionless transport protocol used in computer networking for transmitting data without the need for establishing a connection.

USM — user-based security model
A security model used in SNMPv3 to provide authentication and privacy services for SNMP messages.

About Burns & McDonnell

Burns & McDonnell is a family of companies bringing together an unmatched team of engineers, construction and craft professionals, architects, and more to design and build our critical infrastructure. With an integrated construction and design mindset, we offer full-service capabilities. Founded in 1898 and working from dozens of offices globally, Burns & McDonnell is 100% employee-owned. For more information, visit **burnsmcd.com**.

1004350-0824