Tekelec
For What's Next

WHITE PAPER

Machine-to-Machine Services

This paper examines how operators can capitalize on the growing M2M opportunity.

Tekelec Global Headquarters
+1.919.460.5500
888.628.5527
sales@tekelec.com

EMEA +44.1784.437000
APAC +65.6796.2288
CALA +1.919.460.5500

Tekelec has more than 25 office worldwide serving customers in more than 100 countries. Addresses, phone and fax numbers are listed on the Tekelec website at **www.tekelec.com/offices**.

Table of Contents

Introduction

Several technology trends are rapidly converging to accelerate the growth of the machine-to-machine (M2M) services market. First, most network-centric organizations, including telecommunications service providers, utility companies and enterprises, are migrating to all-IP infrastructures, a technology transition which encompasses home networks as well. Next, devices that access the network, particularly smartphones, are becoming ever more intelligent, expanding M2M's emphasis on both commercial and consumer markets. Finally, increased computing software and systems intelligence is enabling device data management, policy, and billing innovations which, in turn, support time-of-day discount pricing for M2M services. For mobile operators, the emerging M2M services sector presents tremendous opportunities to create new revenue streams, expand the customer base and strengthen margins.

However, M2M services, with a "hyper-connected" environment characterized by huge numbers of devices and applications, unpredictable traffic patterns and the "bursty" nature of billions of connections, also present several challenges for operators. These include network scalability; interoperability of 2G, 3G and LTE technologies; real-time quality of service requirements; and the security of both the network and the devices attached to it. If operators hope to accelerate M2M services revenues, they need networking solutions designed to help them overcome these challenges. Such solutions must be based on a thorough understanding of M2M services and what it takes to deliver those services efficiently, securely and cost-effectively.

A Market Sector Taking Off

While observers vary in their predictions as to how much the M2M services marketplace will grow in the next decade or two, everyone agrees that it is going to grow-and grow exponentially. For example, a study by U.K.-based research firm Analysys Mason expects the number of global M2M connections will rise from 62 million devices in 2010 to 2.1 billion devices in 2020. The GSM Association (GSMA) predicts the number of connected devices will be significantly higher than that-50 billion by 2025. ABI Research, in a study that focuses specifically on the mobile M2M market, forecasts that segment alone will grow from approximately 71 million global connections in 2009 to about 225 million connections by 2014. Viewing the market from another perspective, research firm iSuppli sees global revenues from sales of wireless modules for M2M systems rising nearly sevenfold between 2010 to 2014, from US$1billion in 2010 to US$6.5 billion by 2014. Regardless of how the potential of M2M services is evaluated, this sector clearly is poised, as the Analysys Mason study puts it, "to be one of the fastest-growing connectivity sectors in the next decade."

Major Impacts of M2M Services on the Mobile Network

A look at the ongoing evolution of the mobile network reveals three distinct stages. Prior to 2005, operators built their networks to handle mostly human-to-human voice and SMS traffic and to predict statistically the demand for network resources, including peak periods. About three years ago, operators began to optimize their networks to accommodate rising volumes of data traffic; while humans continue to be the main network subscribers, smartphones began driving the demand for data services, and operators started relying increasingly on policy control to manage network resources. Now, operators are expanding the network-optimization emphasis yet again, this time to "the Internet of things," meaning, most communications will occur between and among machines, rather than humans, and will be "bursty" in nature. Further, operators will be hard-pressed to predict the volume of traffic that will be generated by huge numbers of devices and applications. As mentioned earlier, the rapidly-evolving M2M services environment presents significant scalability, quality of service, radio-access interoperability and security challenges for operators.

M2M versus Human Mobility Services

To ensure that they select and deploy solutions that can tackle these challenges, leading operators are first taking a hard look at the differences between how people and machines use the network and for what purposes. The biggest difference lies with "how." For example, each of 10 million people likely makes two calls per peak hour on the network, while each of 30 million machines may send one message per week. Further, while a person may re-dial a call or call the mobile operator for assistance, none of the machines does that.

Another difference between human-to-human and M2M services lies in the nature of the session. While a person communicates with the network for any one of several possible reasons, a machine does so only because of an event, for example, to provide a scheduled status report. Nearly always, this process, compared to a human-to-network session, is much more automated, systematic and monitored. Consequently, operators need M2M network solutions that are tailored to machines rather than people.

To ensure that they select and deploy solutions that can tackle these challenges, operators are first taking a hard look at the differences between how people and machines use the network and for what purposes.

Applications for Machines, for People or for Both?

Although some industry players classify M2M applications according to category or type, it may be more useful for an operator to look at them from a different perspective: applications that are not at all related to person-to-person communications, typically involving independent devices such as industrial meters or vehicle fleet-management devices, versus applications that relate in some manner to mobile applications used by people. Examples of the latter include:

- Security - vehicle security and anti-theft, as well as vehicle emergency calls
- Transport and logistics - navigation information
- Metering - consumption of electricity, gas and water
- Health - monitoring vital signs and remote medical diagnostics
- Smart living/entertainment - remote controls, synchronization and smart appliances

M2M applications that can be linked inside the network to people's existing mobile subscriptions offer operators enormous potential for strategic differentiation in the competitive marketplace. If an operator can create an ecosystem of devices around an individual subscriber, all connected and interrelated, then that operator can come up with some very innovative services which can translate into strategic differentiators in the marketplace.

Two Powerful Applications Illustrate the Security Challenges

The network implications of M2M services are evident in two initial applications, one of which is widely viewed as a positive development and the other as a destructive situation. The former is the smart grid, which is based on intelligent monitoring systems that are designed to make energy consumption more efficient and cost-effective.

In addition to the issue of how to ensure its security, the smart grid raises another important question for the network: who owns the device? Is it the user? the mobile-network operator? the utility company? The reasonable answer is all three, which means the device must have several identities within the network: one for whoever is managing it, one for whoever is accessing it, and another related to the type of operation handled by the device.

The destructive potential of some M2M applications became obvious in June 2010 with the discovery of Stuxnet, a Windows-based computer worm designed to spy on and damage industrial systems. The first such worm to include a programmable logic controller (PLC) rootkit, Stuxnet not only can re-program the PLCs but also hide those changes. According to news reports, Stuxnet may have damaged some of Iran's nuclear facilities. The implications for the nascent M2M services marketplace are enormous,

notably in terms of security. Will the "Internet of Things" become a prime target for virus writers and hackers and, if so, how do operators and users secure fleets of M2M devices against such attacks?

For mobile operators and other M2M service providers, as well as M2M device vendors, Stuxnet brings the issues of prevention and disaster recovery front and center. Although there is no silver bullet for preventing such attacks, one very effective strategy is to have comprehensive, real-time knowledge of all the devices deployed in the network. With M2M services, it is likely that a device may be out there in the network for 10 years, with no changes to its subscriber identity module (SIM) or to its hardware. That translates into the need for a database that contains not only that device's firmware information, serial number, and other identifying data but also information about its location and state. With all the relevant information residing in one location, mobile operators and other players in the M2M services chain can react quickly if an attack occurs.

Another important tactic in the overall attack-prevention strategy is guarding against theft. Unlike the situation with a person and a mobile device, no one is going to call the mobile operator if an M2M device is stolen. That creates the need for equipment identity register (EIR) capabilities within the database, along with device-tracking tools.

A third prevention tactic is to separate high-ARPU users from large-scale devices, at least on the home location register (HLR). A more fundamental question that mobile operators might ask themselves is whether they should even manage M2M devices on the same core network they use to manage human voice- and data-communications services.

Co-Existence and Interoperability of 2G, 3G and LTE

LTE technology certainly is one factor driving the growth of the M2M market, specifically for high-bandwidth applications such as connected vehicles and video sharing. Yet LTE-based service will exist for some time in coverage islands that are located in a larger 2G/3G environment. From a subscriber data management (SDM) perspective, operators obviously cannot separate 2G, 3G and LTE devices into discrete silos, so they will need solutions that ensure seamless management of shared subscriber and device context and data between and among the different networks. By delivering a consolidated view of the subscriber and device across all three domains, such solutions ensure top-notch system performance and database integrity.

Four Network Areas Need Solutions for M2M Services

When it comes to network operations, M2M services have a major impact on four areas: SDM; messaging services; the policy and charging rules function (PCRF); and performance management. The following is a look at each of these four areas and how operators can optimize them for M2M services.

SDM

SDM is usually the first network area to be affected by M2M services, primarily because of the sheer scale of devices and the need to incorporate related information within associated databases, including the Home Location Register (HLR). There are several opportunities to optimize and improve legacy HLR platforms, which are designed for voice-centric human-to-human communications, so they can accommodate M2M services (see Figure 1).

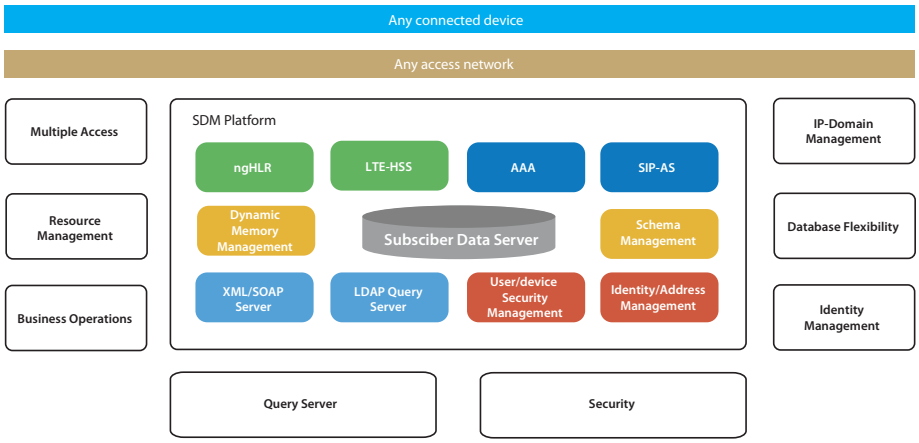

Figure 1: M2M-Optimized Subscriber Data Management

Multiple Access Domains - Operators need next-generation SDM platforms that are more than just next-generation HLRs. These M2M-optimized platforms must be able to track and manage devices across multiple access domains, from 2G and 3G to LTE, Wi-Fi and WiMax, each of which uses its own authentication and authorization functions. Further, such platforms not only have to maintain all those domains but also need to be able to select among them to ensure they can terminate the message to the appropriate domain. They also must tackle all the IP-domain functions, such as dynamic/static allocation of IP addresses; SIP-registration tracking; and network-initiated packet data protocol (PDP) context set-up.

If these functions co-exist in the same database and data server, within the same run time and application framework, operators can build and maintain smarter network capabilities, such as terminating SMS messages to the IP domain when the device is in,

for example, the home Wi-Fi environment, or establishing network-initiated connections. The ability to create such scenarios translates into the ability to conserve important network resources.

Scalability and Flexibility - Another critical set of capabilities for M2M-optimized SDMs includes scalability and flexibility. The required scalability comes from appropriate resource management which allows for the independent scaling of databases and applications. In addition, dynamic, i.e., intelligent, management of the memory for both active and dormant devices enables the system to scale to millions of devices in the most resource-efficient way possible.

The flexibility of the database and data model is extremely important for keeping critical information, such as device serial numbers and firmware information, close to information about the device's location and network state. The proximity of these information sets enables operators to create custom fields about specific devices and via such fields for all devices, obtain an instantaneous view of what is happening in the network right now.

Identity Management - Management of identities in M2M services is another requirement for next-generation SDMs, and it is a particularly important one because of the predicted deployment of billions of devices in conjunction with the coming shortage of expensive identifier resources such as phone numbers and mobile station international subscriber directory numbers (MSISDNs). One way in which an SDM solution could tackle this challenge is to pool these identifier resources and dynamically assign them on an as-needed basis. Another approach is for the SDM solution to make other, more plentiful types of identifiers equivalent to phone numbers. For example, a device could use a SIP uniform resource identifier (URI), which would allow the termination of SMS and other types of sessions.

Security - Securing M2M services is an essential aspect of next-generation SDM platforms. Device authentication must include multiple simultaneous algorithms and a single sign-on function which can follow the device across any type of network. In addition, an EIR is essential, so the operator can identify and block stolen devices immediately.

Device authentication must include multiple simultaneous algorithms and a single sign-on function which can follow the device across any type of network.

Support for Business Operations - A strategic differentiator in M2M services is the ability to tailor operations to individual customer needs-in other words, to provide:

- an application programming interface (API) that is oriented to Web services;
- a framework for extensible markup language (XML) event-based notifications; and
- provisioning capabilities based on automated templates

For example, an electricity company or a vehicle fleet-management company wants to find a given device and check its health. This type of M2M service command requires a rich business API which, in turn, requires a flexible framework to support these simple-object-access-protocol (SOAP) operations.

Finally, a next-generation SDM should incorporate a frontend query server which features a high-performance, open, lightweight directory access protocol (LDAP) API; offline access; and reporting/integration with back office operations. Such a solution is essential to support database queries, analyses and reports.

M2M Network Control Center - Any successful M2M solution must work across network domains and protocols, i.e., the SS7 domain which manages the 2G and 3G worlds and the Diameter/SIP router for the LTE and IMS infrastructure. To control these M2M devices and address the scalability and security issues, operators can dynamically build a subscriber-aware point in between the two. This M2M Network Control Center (NCC) can receive copies of very specific information from the network and store these data alongside other important M2M information. As the NCC grows, the operator can use it for multiple purposes, among them:

- push information to the policy server to enrich the policy rules;
- analyze and monitor the network and, using that base of information, make the appropriate routing decisions; and
- use that same information to block access to the network for devices with specific firmware-in other words, control the floodgates when a security issue arises or an overload is imminent.

In addition, the NCC can be a big help in making transitions, either by enabling 2G/3G/LTE mobility or, for authentication purposes, simply functioning as a proxy between the different domains. Further, the NCC can play a major role when it comes to steering and protection. Many operators are concerned about who controls the schedule on which M2M devices communicate. Generally, the operators do not want the devices themselves to be in control of the communications schedule; given the unpredictability and the chances for a glitch, the devices could all connect simultaneously and bring down the network. Understandably, operators want a means to schedule device communications from a central point, and the M2M NCC, residing between the two signaling networks, is the ideal central scheduling point.

Finally, the NCC is an effective disaster-recovery tool. Acting as an HLR backup, the NCC dynamically stores all the necessary information; if an HLR problem occurs, the operator can keep the network up and running while solving the problem.

Messaging Services

Accommodating SMS - The second network area on which M2M communications has a significant impact is messaging, first because of the sheer volume of messages that M2M services must handle. Secondly, many M2M devices still leverage SMS, which has a global reach. Next, although SMS remains perfectly suitable for monitoring and controlling small-message traffic, its use of the traditional store-and-forward approach makes delivery times uncertain. Operators can resolve this issue by using SMS routers and first delivery attempt (FDA) technology to forward the message instantaneously.

By handling SMS traffic cost-effectively, SMS routers enable operators to lower the cost per SMS message. Basically, SMS routers free up legacy SMSC capacity by using FDA for mobile-originated (MO)/application-originated (AO) messages. That means SMSCs have to handle only about 15 percent of the message traffic that truly needs to be stored until delivery can be completed. This approach extends the life of capacity-stretched SMSCs.

Further, by using advanced load balancing and throughput control techniques, SMS routers eliminate SMS bottlenecks and maximize SMSC capacity utilization. Peak-traffic control prevents legacy SMSC overload by directing high-volume traffic generated by applications, like tele-voting, directly to the application. Regular traffic is delivered to the destination handset, a high-performance SMS store solution or an existing legacy SMSC. Finally, SMS routers enable operators to establish priorities for SMS traffic and thereby offer improved quality of service (QoS) to premium customers.

Application Gateways - Operators also can use a short message peer to peer (SMPP) application gateway which helps to distribute traffic from multiple M2M applications toward the SMS infrastructure. Via a flexible, rules-based engine, an application gateway intelligently distributes application traffic directly to SMS routers, thus bypassing SMSCs and guaranteeing message delivery.

In addition, M2M communications service providers must open an API to application providers. Otherwise, an application such as fleet management, which needs to communicate with all the vehicles that are out there, would be forced to use the standard network-entry point, i.e., the SMSC type of interfaces. Via a messaging platform, operators can create application gateways-specific entry points for all these applications-toward the devices.

Security - To reduce the potential for malicious SMS messages disrupting M2M services, securing the connection between the device and their designated M2M application for

transportation of SMS messages is critical. Operators can use a SMS Security solution that filters incoming messages, using user definable screening capabilities and content filtering, to protect networks and devices from attack.

Policy and Charging Rules Function (PCRF)

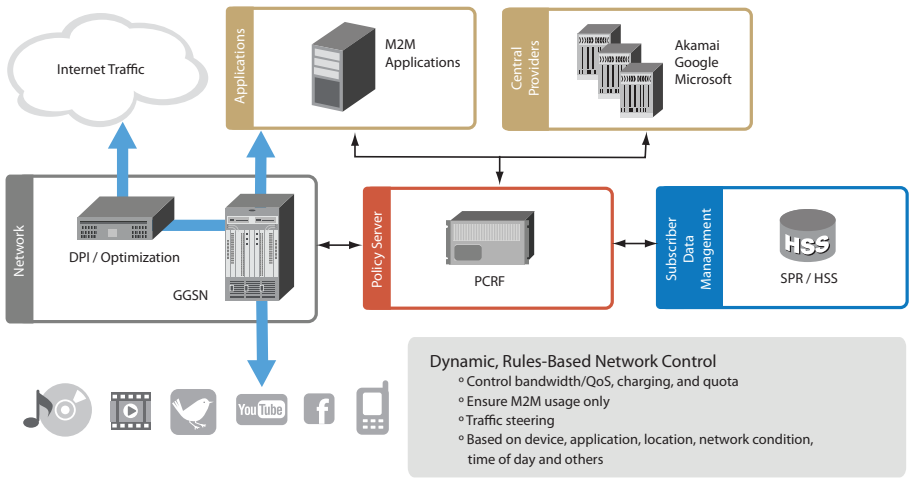

Figure 2: The Role of Policy in M2M

Differentiate Among Devices - As policy control becomes more tightly integrated with the SDM infrastructure, the emergence of M2M services creates critical requirements for the third network area on which M2M services will have a direct impact-the PCRF. These requirements include the ability to distinguish these M2M devices from devices used by subscribers and, using M2M profile, application, location, network condition and time of day, to differentiate among various device classes. Depending on the type of device in question and its network status, the policy server also must be able to offload traffic to different networks (see Figure 2).

Manage Relevant Resources - In addition, the PCRF must be able to manage and, in some cases, to guarantee bandwidth. Some M2M applications, for example, electricity metering, consume very little bandwidth, while others, such as surveillance cameras, need a great deal of guaranteed bandwidth to function properly. Clearly, an operator does not want to provide the same levels of quality of service (QoS), uplink and downlink bandwidth, latency, jitter, etc. to the two device types.

Further, those required levels likely will vary even for one device, say, a wirelessly connected security system in the home. The system's periodic stay-alive and health checks will require a certain priority level but not necessarily significant bandwidth. However, if the alarm system detects movement and triggers the camera to function and send a recording feed to a central server, a different QoS level becomes necessary. In other words, based on an ongoing M2M communications session, the ability to change

policy rules, even for one device, is essential-not for the lifetime of the device but only when circumstances dictate.

It is the job of the PCRF to handle all these requirements. Located at the center of the policy architecture, the PCRF basically links the applications on top, the devices on one side and the subscriber and device data-residing in the subscription profile repository/ home subscriber server (SPR/HSS)-on the other side. Although it is the same policy architecture in current 3G and LTE networks, it must have some M2M service-specific rules to be an efficient, effective solution.

Performance Management and Network Planning

Finally, operators competing in the M2M marketplace need solutions that enable them to manage the network's performance, so they can comply with service level agreements (SLAs); detect and troubleshoot problems; and plan/optimize the network. For operators, the M2M customer is not each individual device but an enterprise such as a utility or a fleet-management company. Such customers require different, more stringent SLAs-specifying the level of network support and service the operator is to deliver-than operators provide to consumers. In the case of a utility, for example, an SLA might specify that a penalty will apply if for whatever reason an electricity meter is not able to submit its reading at the end of the month.

Or, a customer may request an SLA that guarantees the ability to refresh the firmware of that customer's devices a set number of times per hour, using over-the-air provisioning. Because such a capability will be critical for such M2M customers, operators must be able to track the relevant information to comply with these SLAs.

Because ARPU in M2M services likely will be much lower than for consumer services, the PCRF is essential for network planning and optimization as well. For example, using Tekelec's Performance Intelligence Center (PIC) products, an operator can obtain statistics on PDP context activations within a certain period of time for an M2M service, compare those with the number of PDP context activations in that same time period for a consumer service and then optimize M2M-related network resources accordingly.

Since many M2M services are new, it can be difficult for operators to determine the cost and expected revenues, for example, of adding 20 million meters and to determine the impact on the network of adding those meters. The answer to both questions will vary from operator to operator. Consequently, the ability to obtain accurate statistics related to M2M services is an operator's first step in determining whether it makes sense to sign that 20-million-device contract, and how much it is going to cost in new radio bandwidth and other network resources.

The ultimate success of the PCRF in ensuring operators can manage the performance of their networks at the required levels of granularity and optimize them for M2M services depends on establishing a feedback loop.

For example, Tekelec's solution uses a real-time feedback loop that enables operators to examine and adjust network resources continuously to deliver a consistent quality of experience (QoE) tailored to each customer. Using this flow of real-time information, operators can adapt quickly to the unpredictable demands of the data network. As shown in Figure 3, the continuous feedback loop performs three actions: analyze, decide and act.

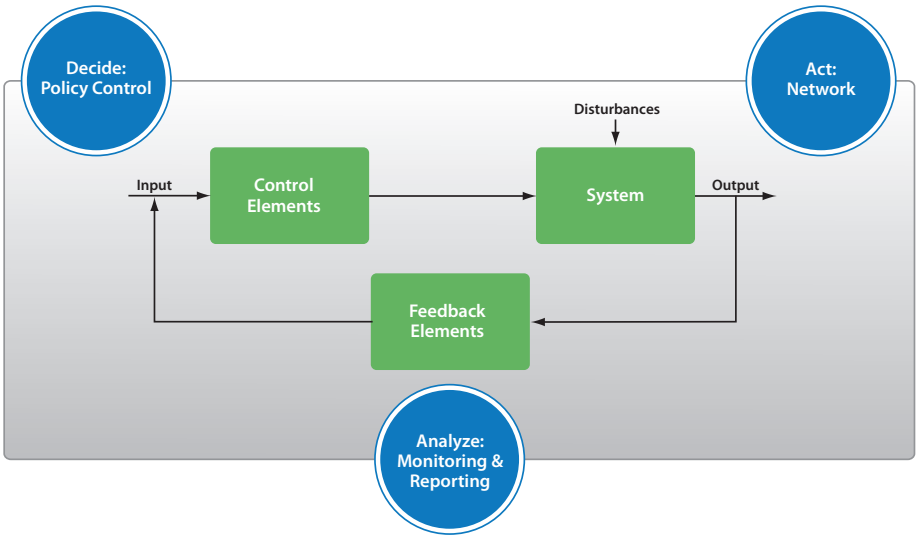

Figure 3: M2M Feedback Loop

- **Analyze** the activity on the network and QoE with Tekelec's Performance Intelligence Center (PIC) to ensure optimal use of resources and appropriate QoE for M2M services. Real-time network metrics identify issues and provide the information necessary for unparalleled levels of network decision-making.
- **Decide** where and how to handle traffic by using real-time data from the network and device information from the SDM system. Equipped with this context-aware information, operators can enhance the intelligence of their networks, making them device-aware as well as radio access network (RAN)-aware. With a sophisticated PCRF, operators can adjust policies, apply business rules and optimize traffic through routing and load balancing to improve network performance.
- **Act** on the decision to tune or balance traffic flows dynamically to ensure the delivery of high-priority traffic or latency-sensitive traffic with a guaranteed QoS.

The concept of a feedback loop is really about looking at what's going on in the network and making sure that M2M services do not flood the RAN. This is the future of policy — establishing this loop and controlling what radio and core network resources

are available or not in a given session and for a given device. By allowing the operator to monitor those interfaces and evaluate the relevant information, the feedback loop enables the operator to apply rules to the system, see the actual impact of those applied rules and, if necessary, modify them.

Summary

As mentioned at the outset, the opportunity for mobile operators in the emerging M2M services market is enormous, no matter how one measures it. Granted, some challenges confront operators that want to seize that opportunity, namely, scalability, 2G/3G/LTE interoperability, quality of service, and security. The current lack of M2M standards is another bump in the road, but many organizations, including Tekelec, are actively working to resolve that issue. For example, an initiative is underway within the European Telecommunications Standards Institute (ETSI), and the Global Standards Collaboration-15 task force, meeting recently in China, appointed the TIA TR-50 organization to coordinate various standardization efforts.

With M2M standardization underway, leading operators that want to capitalize on the huge M2M opportunity already are looking to vendors for help in tackling the other sets of challenges. They recognize the solutions they need to succeed in this emerging marketplace will not come from vendors that design their platforms around mobile consumers and mobile voice communications. Rather, the solutions that will help them capture market share, boost ARPU and enhance long-term profitability will come from vendors that truly understand M2M services and its unique networking requirements.

About Tekelec

Tekelec enables billions of people and devices to talk, text and access the Web. Our portfolio delivers a unique layer of intelligence allowing service providers to both manage and monetize the exponential growth in data traffic and applications. Tekelec has more than 25 offices around the world serving customers in more than 100 countries. For more information, please visit **www.tekelec.com**.

Appendix: Acronyms Used in This Document

AAA	Authentication, Authorization, and Accounting
API	Application Programming Interface
ARPU	Average Revenue Per User
EIR	Equipment Identity Register
FDA	First Delivery Attempt
GSM	GSM Association
HLR	Home Location Register
HSS	Home Subscriber Server
IP	Internet Protocol
LTE	Long Term Evolution
M2M	Machine to Machine
MSISDN	Mobile Station International Subscriber Directory Number
PCRF	Policy and Charging Rules Function
PDP	Packet Data Protocol
PLC	Programmable Logic Controller
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
SDM	Subscriber Data Management
SIM	Subscriber identity Module
SLA	Service Level Agreement
SMPP	Short message Peer to Peer Protocol
SMS	Short message Service
SOAP	Simple Object access Protocol
SS7	Signaling System 7
URI	Uniform Resource Identifier:
XML	Extensible Markup Language