November 2018

Context-Aware Identity Management Framework

Abstract

This report assesses how the emerging market for context-aware information can be applied to solve future identity management needs. It proposes a framework for context-aware identity management that builds upon existing and future authentication and authorization infrastructure to deliver a more robust set of solutions for network providers, enterprises, IT operations and consumer-based services. It also addresses the role of context-aware identity management in supporting the need of federated authentication and authorization processes when the context management functions reside in a different domain.

Foreword

As a leading technology and solutions development organization, the Alliance for Telecommunications Industry Solutions (ATIS) brings together the top global ICT companies to advance the industry's business priorities. ATIS' 150-member companies are currently working to address 5G, cybersecurity, robocall mitigation, IoT, artificial intelligence (AI)-enabled networks, the all-IP transition, network functions virtualization, smart cities, emergency services, network evolution, quality of service, billing support, operations and much more. These priorities follow a fast-track development lifecycle: from design and innovation through standards, specifications, requirements, business use cases, software toolkits, open source solutions and interoperability testing.

ATIS is accredited by the American National Standards Institute (ANSI). ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a founding Partner of the oneM2M global initiative, a member of the International Telecommunication Union (ITU) and a member of the Inter-American Telecommunication Commission (CITEL). For more information, visit www.atis.org.

Notice of Disclaimer and Limitation of Liability

The information provided in this document is directed solely to professionals who have the appropriate degree of experience to understand and interpret its contents in accordance with generally accepted engineering or other professional standards and applicable regulations. No recommendation as to products or vendors is made or should be implied.

NO REPRESENTATION OR WARRANTY IS MADE THAT THE INFORMATION IS TECHNICALLY ACCURATE OR SUFFICIENT OR CONFORMS TO ANY STATUTE, GOVERNMENTAL RULE OR REGULATION, AND FURTHER, NO REPRESENTATION OR WARRANTY IS MADE OFMERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR AGAINST INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. ATIS SHALL NOT BE LIABLE, BEYOND THE AMOUNT OF ANY SUM RECEIVED IN PAYMENT BY ATIS FOR THIS DOCUMENT, AND IN NO EVENT SHALL ATIS BE LIABLE FOR LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES. ATIS EXPRESSLY ADVISES THAT ANY AND ALL USE OF OR RELIANCE UPON THE INFORMATION PROVIDED IN THIS DOCUMENT IS AT THE RISK OF THE USER.

NOTE - The user’s attention is called to the possibility that compliance with this standard may require use of an invention covered by patent rights. By publication of this standard, no position is taken with respect to whether use of an invention covered by patent rights will be required, and if any such use is required no position is taken regarding the validity of this claim or any patent rights in connection therewith. Please refer to [http://www.atis.org/legal/patentinfo.asp] to determine if any statement has been filed by a patent holder indicating a willingness to grant a license either without compensation or on reasonable and non-discriminatory terms and conditions to applicants desiring to obtain a license.

Copyright Information

ATIS-I-0000070

Alliance for Telecommunications Industry Solutions
1200 G Street, NW, Suite 500
Washington, DC 20005

Contents

1. Introduction – The Future Context-aware Ecosystem

One of the most significant industry trends is the evolving marketplace for contextual data. Technology advancements in Internet of Things (IoT) sensors and devices, edge computing and storage, and a massive number of connected devices, are creating the abundance of contextual information available to networks. The growing need to effectively collect, manage and analyze this information is driving the development of new solutions that can effectively discover and apply context-based information to meet future market requirements.

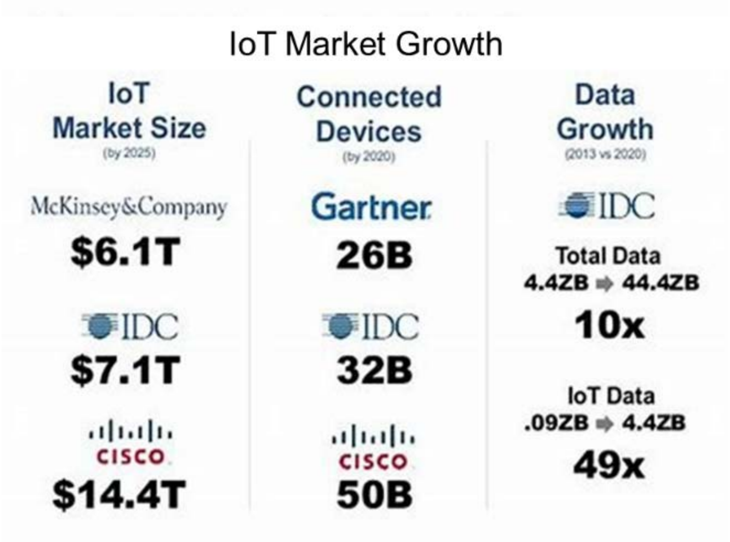

IoT Market Growth

IoT Market Size (by 2025)	Connected Devices (by 2020)	Data Growth (2013 vs 2020)
McKinsey&Company $6.1T	Gartner 26B	IDC Total Data 4.4ZB ➡ 44.4ZB 10x
IDC $7.1T	IDC 32B	IoT Data .09ZB ➡ 4.4ZB 49x
CISCO $14.4T	CISCO 50B	

The growth of IoT-based contextual data is creating new opportunities for context-enabled services and applications. It is the combination of IoT device growth and the expanding market for connected devices that is accelerating the availability of contextual information. This is further illustrated by the exponentially higher growth rate of IoT data with respect to data growth in general.

In most applications, context produces a higher intrinsic value for information versus raw data. Context-awareness infers that either the data, or the associated metadata, delivers some additional information about the environment surrounding the user, device or object. In this way, context-awareness will typically offer a greater level of reliability and usefulness surrounding the source and application of the data.

"The context of an entity is a **collection** of **measured** and **inferred knowledge** that describe the *state* and *environment* in which an entity exists or has existed."[1]

Extending these concepts to the way systems apply context awareness and situational awareness yields the following definitions:

- **Context Awareness** – The ability of a system to gather information about its environment and adapt behavior accordingly.

- **Situational Awareness** – The perception of the environment with respect to time or space, the comprehension of the environment, the projection of those data onto the task goals and the prediction of future events and state based on that understanding.

The remainder of this document will focus on the application of these principles to meet the future challenges of managing user and device access to a broad array of systems, services, applications and resources.

[1] Agoulmine, N. (Ed.). (2010). Autonomic network management principles: From concepts to applications. Academic Press.

2. Context-aware Identity Management

The Evolution of Identity Management

Over the last decade, identity management solutions have been in a continuous state of improvement, responding to new vulnerabilities and meeting the needs of new applications across industries. Enterprises, IT operations, network operators and vertical industries must continually apply more robust solutions to meet the challenges of this changing environment.

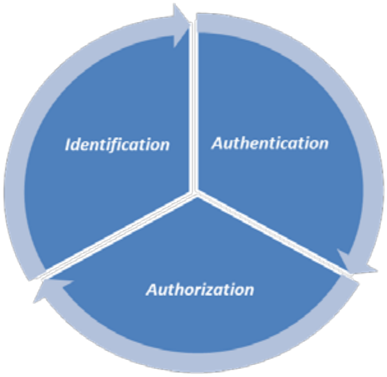

While many identity management models exist, there are some common elements that define this environment.

Identification is based on the recognition of unique identifiers associated with a user, device or object. In today's market, an identifier could be a name, email address, account number, IP address, MAC address, web domain, social media identity or any number of other attributes that uniquely define an entity.

Authentication encompasses processes related to access control and validating the identity of a user or device. More recently, it has relied on concepts such as stronger passwords, multi-factor processes and biometric authentication solutions. It is important to note that the industry has made significant investments in improving authentication techniques because this function acts as the front end to many other downstream processes.

Authorization involves the granting of privileges and assignment of roles once an identity has been authenticated in the system. In many cases today, privileges are pre-assigned to a specific identity once that user is authenticated within an environment. These privileges permit an identity to access specific resources based on their role, function, level or other attribute.

The evolution of authentication and authorization techniques represent a progression from *what a person knows* to *something a person has*, and more recently, to multi-factor approaches including biometric and behavioral factors. In the first case, "knowledge that a person has" is the basis for passwords and traditional security questions to help validate an individual user. Improvements in password strength or more extensive security profile questions have conversely created diminished user experiences when attempting to access applications across many security domains. Over the past decade, many enterprises and IT operations have relied increasingly on "something the user has" to validate identity, offering second-stage authentication or recovery from lost/forgotten passwords. Examples include security tokens, access cards or other credential-based media that an individual would use to access resources.

Multi-factor authentication (MFA) incorporates two or more elements of identity, which may include something you know, something you have or something you are (e.g., biometric features such as fingerprints, facial identity or eye recognition). MFA can be further optimized by delivery of these factors across different identity communications paths, such as using an SMS-delivered code to recover or reset an identity.

Two important developments are laying the groundwork for a more robust identity management strategy. First, applications are emerging that cross many industry sectors or domains and require authentication and authorization processes beyond single-sign-on capabilities. Second, the IoT market is beginning to radically impact the industry with a wealth of new data, including real-time contextual information. These developments, coupled with a desire to develop identity management approaches that significantly improve user experiences, have led to a disruptive opportunity to advance context-aware identity management.

Context-Aware Identity Management

Contextual information can be acquired from many sources. It can include data produced by IoT sensors, including environmental data, location, proximity, presence and sensory data. It can also include user preferences, profiles, behavior or other characteristics. When this information is collected, aggregated and analyzed, with appropriate privacy and consent controls, it can act to unify contextual information with identity management processes, creating the potential for a very robust solution. The fact that much of this information is real-time and dynamic enables a more powerful approach to managing authentication and authorization in the future.

Given these technology and market developments, a context-aware identity management solution applies the following elements:

"*Situational data* that enhances the identification, authentication or authorization of a user, device or object, and is *relevant* to a specific application or process."

This report's objective is to provide a detailed assessment of context-aware identity management approaches and to describe the framework and functional elements that will promote adoption and interworking in the future.

3. Market and Technology Drivers

Market Factors

Identity management opportunities related to the expanding IoT market are expected to have a major impact over the next few years. Rapid innovation across many IoT applications is driving diverse authentication and authorization approaches across industry verticals. Industrial IoT, transportation, energy, medical, home automation and other IoT-focused sectors are each demanding a high level of security robustness. While these approaches will offer specific benefits, there is the potential that diverse approaches will negatively impact user experiences for applications that intersect with multiple vertical markets.

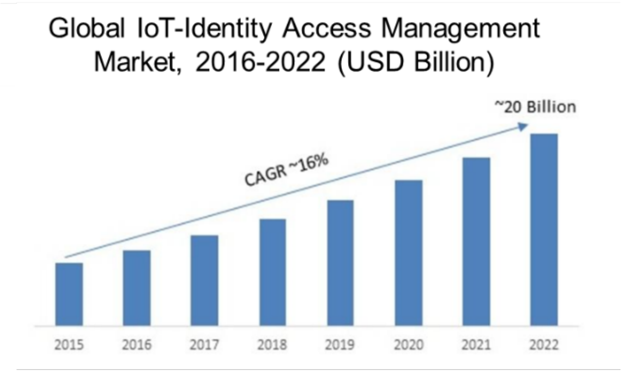
Global IoT-Identity Access Management Market, 2016-2022 (USD Billion)

Source: Market Research Future

Technology Drivers

There are several inherent benefits associated with identity management solutions derived from contextual information, including:

- The robustness of any identity management solution is defined by the degree to which the factors are independent of other factors. For example, if each factor of authentication is built on some level of personal information, then exposing personal data creates an overall vulnerability.

- Contextual information is dynamic (in most cases), so it reduces the need for the user to manage and reset identifying information.

- Data derived from context-aware sources is more difficult to manipulate by bad actors.

- Context-based authorization offers major enhancements over existing processes:

    - Currently, privileges are most often statically assigned or domain-based.

- Knowledge of contextual environment promotes dynamic treatment of privileges.
- Opportunities to integrate context with single-sign-on aspects.

Use Cases

There are several important use cases that span identity management applications across enterprise, IT operations, consumer and industry verticals. These use cases offer significant opportunities for increasing the robustness level and the user experience associated with future services.

The use cases listed below represent a cross-section of expected authentication and authorization scenarios that can be applied to the development of a functional framework for context-aware identity management.

Use Case 1 – Context-Enhanced Authorizations for Administrators/Developers

Description – Based on context, authorization levels will automatically be altered for privileged users such as system and network administrators and developers.

Current Approach – Typically, authorization is based on role-based access control (RBAC), and the RBAC does not changed based on external factors. Membership in predefined groups is the entire basis for access.

Future Context-Sensitive Response Options:

- **Time of Day/Day of Week**:
  - *Established normal time/day* – Fully authorized access based on role.
  - *Abnormal day/time* – User is segmented into a network with basic office functionality and is required to provide additional level of authorization prior to sensitive data access. Once identity is established, login can be re-initiated

and treated as a known source or dynamically reallocated to a known source network segment.

- **Location** – Source of login will trigger different authorization levels:
    - *Internal corporate office* – Fully authorized access based on role.
    - *Known source locations (home office)* – Fully authorized access based on role. Additional auditing enabled on any data movements.
    - *First/new access source* – User is segmented into a network with basic office functionality and is required to provide additional level of authorization prior to sensitive data access. Once identity is established, login can be re-initiated and treated as a known source or dynamically reallocated to a known source network segment.
    - *Known untrusted source* – Account is either disabled or moved to an isolated segment for monitoring where no live data is available if forensic evidence is required.
- **Resource Request** – Change response based on atypical resource request (e.g., servers, code repo, data stores).
    - *Typical resource request* – Requests within normal behavior patterns. Fully authorized access based on role.
    - *Direct resource support* – Access to a system or resource that is directly supportive of typical resource requests (e.g., normally manages firewalls, attempts to access directly connected routers or switches). Trigger additional auditing and potentially a re-authorization request.
    - *Unrelated resource request* – Examples include a mail admin attempting to access a database, or a firewall admin attempting to access a data center's top-of-rack gear. Require re-authorization and alert SIEM, or possibly deny connection.
- **Operational State** – Change restrictions based on operational state of the network:

  - Normal operating conditions – Normal authorization process and access permissions occur.

  - Detected attack – Restrict all access to only established norms for location, time/day, resource requests, etc.

- **Fail State** – Relax operational restrictions to allow for maximum support to return services to a normal operating state. This is an emergency state where it is all hands-on deck due to a service outage and it is necessary to grant access, such as a break-glass account process.

Use Case 2- Context-Enhanced Network Flow Authorizations

Description – Based on context, the network automatically alters the flow of network traffic.

Current Approach – Typically, the flow of network traffic is based on network discovery and routing protocols ACLs on network devices. Context for current operating conditions, with the exception of a "failure," do not alter network behavior.

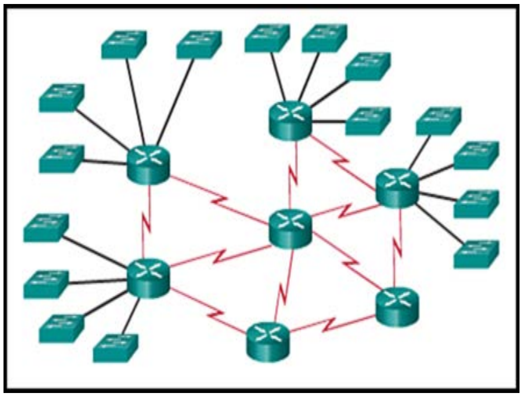

Future Context-Sensitive Response Options:

These options all presume a way to label data flows, data objects and threat models and an ability for the network infrastructure to react to those pieces of metadata.

- **Unknown Source/Identity Pairing** – Similar to a role access decision. When a user identity and source location are abnormal or first observed, the network will alter the flow through a device that executes full packet capture and an intercepting proxy to inspect the traffic. If the identity and source pairing is confirmed good, flow can be moved back to a normal segment, reducing latency and improving performance for the user and reducing security monitoring that the SOC would have to perform.

- **Abnormal Source/Destination Device Pairing** – When an abnormal device-to-device flow is detected, additional mitigations are automatically implemented. For example, if a print device is observed connecting to a database, the system will alert and shut down the flow automatically. The network could then restrict the access for that print device to interact with any other device on the network other than inbound print jobs from an authorized print server.

- **Time Sensitivity** – When data is considered time sensitive, the network will transfer the data via the highest bandwidth, lowest latency connections regardless of ownership devices or path. The system may alter routing and priority to further reduce overhead and accelerate content delivery.

- **Data Sensitivity** – When data is considered highly confidential, the network will transfer the data via a predetermined route through trusted infrastructure. This route may not be the shortest path and may move through slower links, thus delaying the delivery of the data, but the risk of compromise is greatly reduced. In some cases, when delivery cannot be assured because no path is found, the data may be queued for re-classification or alternate delivery, whichever is necessary to protect the data.

Use Case 3 -Context-Enhanced Authentication for File Exchanges

Description – Enabling context-based authentication for file exchanges helping to restrict transfer of IP to end points that pose enhanced risk.

Current Approach – There are solutions on the market today that use this concept but focus on user membership in security groups or the end point's IP address. These solutions tend to focus on the individual document through DRM-like capabilities but fail to add context to enhance the authentication or authorization processes.

Future Context-Sensitive Response Options:

When a file transfer is initiated, the system collects additional information to determine if the end points require enhanced authentication and to determine authorization to transfer.

These factors may need to leverage AI and machine learning (ML) capabilities to adjudicate the decision in near real time. Rule-based solutions may not be sufficient.

Context-enhancing examples include:

- User/machine entity making the request.
- Source and destination IP address.
- Anomalous pairing of user/IP.
- Agent strings for end user device.
- Source IP is known to be in a low-security, high-risk area (e.g., coffee shops, airports, foreign countries).
- Unauthorized attached devices.

Use Case 4 - Context-Controlled Access to Vehicle Data

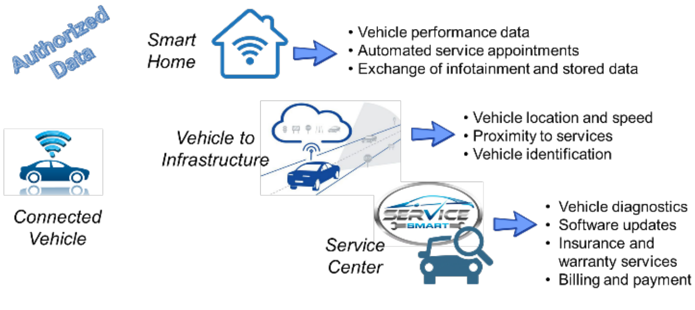
Authorized Data
Smart Home
• Vehicle performance data
• Automated service appointments
• Exchange of infotainment and stored data
Connected Vehicle
Vehicle to Infrastructure
• Vehicle location and speed
• Proximity to services
• Vehicle identification
SERVICE SMART
Service Center
• Vehicle diagnostics
• Software updates
• Insurance and warranty services
• Billing and payment

[5]

Description – Sharing of connected vehicle data based on location, proximity and overall contextual environment.

Current Approach – With the increasing availability of connected vehicles, information is shared via dedicated communication links to car manufacturers, subscribed security and emergency service providers, satellite-based navigation services and Internet-based applications. Each of these services generally requires a separate authentication process. Access to detailed vehicle diagnostics is typically provided via an on-board diagnostics port or some remote diagnostics reporting link.

Future Context-Sensitive Response Options

Connected vehicles are becoming more prevalent in the market, and vehicles are taking on a larger amount of computing and data storage capabilities. Vehicle data could be shared more efficiently across various environments based on contextual information coupled with appropriate authentication and authorization mechanisms. The type of data would be authorized by the user based on profiles, behaviors, preferences and other factors.

Contextual data associated with one application class could assist in authorizing the same user or vehicle across multiple applications and domains. The following are examples of vehicle-associated domains that could utilize contextual information to authenticate identities and authorize a prescribed set of applications:

- **Smart Home** – A vehicle returns home and gains authorized access to the home Wi-Fi network.
    - Vehicle automatically downloads performance data (e.g., mileage, MPG, notifications) to the driver’s home network.
    - Vehicle is synched with driver’s smart home applications to share appointments, maintenance reminders, etc.
    - Driver can download new infotainment files (e.g., audio, videos, travel information) to the vehicle.
    - Driver selectively downloads stored video files collected from the vehicle’s cameras.
- **Vehicle to Infrastructure** – A vehicle communicates with smart transportation infrastructure to access traffic and road conditions, navigational services, toll payments and other in-transit services.
    - Vehicle’s location and speed information is collected as anonymized data in support of smart transportation services within a municipality.
    - Basic vehicle data (vehicle capabilities) and destination can be used to advise the driver (on a customized basis) about the fastest route based on traffic and weather conditions.
    - Authorized location-based data is collected to advise the driver of nearby services (e.g., hotels, fuel or charging stations, restaurants).

- On-board authorized devices (e.g., transponders) are used to collect vehicle telemetry data at toll facilities.

- **Service Center** – A vehicle communicates with a service center to schedule appointments based on service notifications for immediate attention or normal maintenance needs.
  - An engine light notification generates an immediate appointment request to a driver's nearest or preferred service center location and communicates diagnostic information in advance of the appointment.
  - Within the perimeter of the service center, detailed vehicle information can be shared to advise customer as to severity and expected level of repair time. The latest software can be downloaded to the vehicle while at the service center.
  - Communications of diagnostics and repair costs can be automatically shared with the driver's car warranty services, insurance company, etc.
  - Billing and payment can be authorized to alleviate wait times and processing.

Use Case 5- Smart Health: Context-Protected Sharing of Patient Data

Description – Enabling context-based protection and authorized exchange of patient data based on environment.

Home Patient Monitoring

Physician's Office

Hospital

Smart Health and Fitness

Sharing of Patient's Health Data

Monitoring High Risk Patient's Movements

Current Approach – Smart health services and telemedicine are rapidly gaining acceptance. But sharing of patient data still requires complex patient authorization and action due to the need for privacy and restricted access to specific medical professionals.

Future Context-Sensitive Response Options:

The broad range of emerging patient healthcare and fitness options will drive a greater need for sharing of patient data. Knowledge of the patient's environment could greatly enhance this process by authorizing specific health information to be shared within a given set of contextually identified locations and situations.

The following are examples of healthcare environments and the type of patient data that may be authorized in each case:

- **Home Patient Monitoring** – Requires a highly secure link between patient location and healthcare management to exchange health vitals, images, statistics, etc. The patient may pre-authorize additional level of surveillance under certain medial conditions.

- **Physician's Office** – Patient history and reports from other healthcare specialists.

- **Hospitals** – The patient's status and location within the hospital could be automatically exchanged with primary points of contact authorized by patient.

- **Smart Health and Fitness** – Fitness routines and basic health data could be shared with the user's primary physician or fitness expert within a given location or environment.

- **Monitoring a High-Risk Patient's Movements** – Patients with complex medical or cognitive disorders, location and vitals could be monitored when they move outside a certain environment, such as a hospital or assisted-living facility.

- **Sharing of Patient's Health Data** – Access via patient portals could be simplified with additional knowledge of patient's location, environment, proximity, etc.

4. Related Industry Activities

In assessing current industry activities related to the development of a context-based identity management framework, ATIS has taken on a unique role in defining an architecture that incorporates these principles. The following is a list of collaborative activities across the industry that are focused on different objectives, but include some related elements of work:

- **GSMA**: https://www.gsma.com/
    - Mobile Connect: https://www.gsma.com/identity/mobile-connect
    - Identity and Access Management (Link is to a liaison describing this work item for information): http://www.3gpp.org/ftp/Inbox/LSs_from_external_bodies/GSMA_PSMC/
- **MEF**: https://www.mef.net/
    - IoT Activity Group: https://www.mef.net/IMTC-resources-directory/uc/iot-activity-group/index.html
- **W3C**: http://www.w3.org/
    - W3C Credentials Community Group: https://www.w3.org/community/credentials/

5. Principles

As a precursor to the development of a context-aware identity-management framework, it is important to consider a list of principles and underlying assumptions that will serve as the basis for further architectural development.

Contextual Information

Contextual data may include some aspects of personal information or user/device profiles used in current authentication and authorization processes. For example, user-known information and user preferences are common sources of access management today. However, it is expected that an unbounded amount of contextual information will be generated by new IoT sources and connected devices.

The following are examples of contextual sources of information across users, devices and objects. In some cases, the contextual attributes may exist across multiple categories (e.g., proximity could be related to user and/or device).

Users

Personal information
Characteristics
Profiles
Preferences
Groups
Role

Devices

Location
Proximity
Presence
Environmental
Sensory
Categorization
Access type

Objects

Relationships
Associations
Movement

Visualizing this environment further, multiple contextual attributes can be collected and associated with a specific authentication or authorization request and collectively create a more robust and secure solution. The following diagram represents the intersection of multiple factors that can successfully authenticate a user or device or authorize a set of privileges.

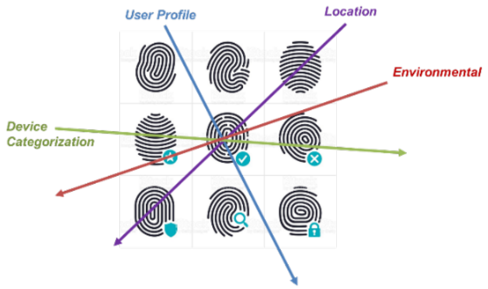

Identities and Identifiers

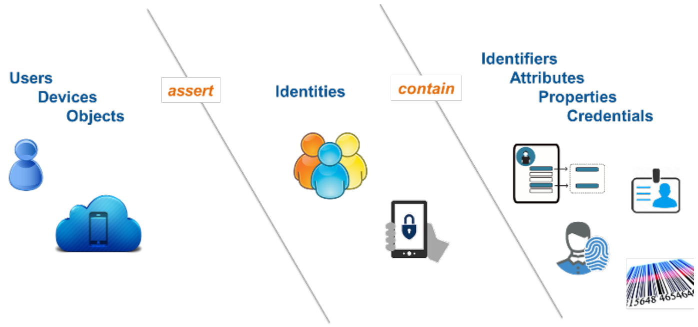

One of the important principles regarding identity management is that a user or device may assert one or more identities, which may then contain multiple identifiers, attributes, properties, etc. In this way, a single user or device may assert different identities depending on the application, domain or intended use. A simple example of this principle is that a single user may assert hundreds of different online identities for different online accounts, social media outlets, personal or business domains, etc. Likewise, a device may contain multiple intrinsic identities, such as in the case of those that support multiple SIM cards.

A context-aware identity management system is significantly more powerful if it can associate multiple identities with the same user. This enables the identity management system to apply knowledge of current or prior user behavior in order to make better decisions regarding authentication, authorization, policy and routing for a current interaction. However, with this additional capability, serious privacy concerns must be considered and resolved.

In a system capable of associating multiple identities containing various types of identifiers, the following is required from an architectural perspective:

- A database of known identities associated with each user/device, each potentially with its own set of policies (from the network point of view). In addition, there could be multiple policies associated with each identity, dependent on context.

- A mechanism for determining identity based on some combination of *direct user input* and/or *action being attempted*.

- A mechanism for applying policy appropriate to the identity and context, which will determine and/or modify network behavior.

Impact of Identity and Context on Network Behavior

From a network perspective, the intersection of identity and context could drive networks to take certain related actions:

- **Network Selection/Access** – A user or device may have multiple identities, each of which would drive the use of a particular network.
    - Example: The user subscribes to a public network and is also an employee of an enterprise. Depending on which identity is being asserted, it may be appropriate to connect to one network or another. (Enterprise or public network, assuming both are visible). If the user is asserting an enterprise identity and the enterprise network is unavailable, then connect via the public network, but automatically set up VPN and limit behavior (see next bullet).
- **Network Connection/Policy** (At the network level, not the application) – Depending on identity and context, the network makes only certain endpoints, applications and services available to the user/device at that time.
- **Priority/QoS/Route Selection/Security** – Depending on identity and context, a particular set of flows might be given different routing priority, utilize different routes and be given different levels of security.

6. User/Device Contextual History and Session State

Two types of memory can enhance the power of a context-aware identity management framework. The first is user/device history, which has to do with the ability of the identity management system to use historical user/device context information, or knowledge of the user/device's context during prior interactions with the system, in order to enhance policy decisions with respect to the current interaction. The second, session state, has to do with the ability to leverage awareness of potentially policy-impacting changes in the user/device environment during the course of an ongoing session.

Contextual History

The history of a user's/device's environment and its prior interactions with the system can be useful for driving policy decisions regarding the current interaction. For example, the confidence level for a device's system request can be increased if the association of a user and device can be ascertained based on past behavior. On the other hand, historical information can be used to flag anomalies that need further action to resolve. For example, if a user attempts to access a resource from San Francisco, and the historical data shows that he accessed the system from New York twenty minutes prior, the identity management system would need to take action to resolve this impossible change in location.

Contextual history may be maintained for each specific resource for which a particular identity is asserted. It also may be maintained more centrally, applying to multiple resources and asserted identities associated with the same user/device.
In the former case, the resource provider (application/network/service) would store and use contextual information about prior interactions of that user/device with that resource only. This mechanism is commonly applied by applications today.

In the latter case, a resource provider would be able to leverage information about the contextual history of the user/device over multiple asserted identities and resource providers. This historical information would be accessible from a centralized entity that had knowledge of a user's/device's multiple identities and resource providers. While this capability could greatly enhance the power of an identity management system, it also raises potentially serious privacy concerns that would need to be addressed and resolved.

Finally, contextual history may be maintained by the sources of contextual information themselves. For example, if a mobile network is viewed as a source of contextual information for resource providers, that network may maintain contextual information about the user/device that spans multiple asserted identities associated with that user/device (e.g., location history). This case raises privacy concerns similar to those raised in the previous paragraph and increases the importance of addressing privacy and anonymity in the system.

Another question to consider is the amount of history to maintain. The relevance of past information will vary depending on the specific resource being requested. In the case where each resource provider maintains history about its own prior interactions with a particular asserted identity, the resource provider can determine how much history to store. But if history is to be maintained centrally, a mechanism for determining which information should be stored, for how long and subject to which privacy/sharing policies must be developed.

Session State

A user's/device's environment may change during the course of an interaction with a resource provider, and the new context might require modifying current policies. For example, a resource provider's policy might be to allow access to a secure file system if the user/device is in a secure location, but to deny that access if the user/device is connected via a public hotspot. A user/device might initiate that interaction with the secure file system while in a secure location but move to an unsecured location during the course of the session. It would be useful for the resource provider to have a mechanism for becoming aware of this change in context, and to be able to adjust access policies in response to it.

It is worthwhile for a resource provider to have a means for knowing when potentially policy-impacting changes in a user's/device's environment occur mid-session. Several mechanisms could apply, ranging from direct polling of contextual sources to variations on a subscribe/notify model. Each of these alternative mechanisms for tracking session state have implications in terms of complexity and load, and the mechanism adopted would need to reflect a more rigorous analysis of trade-offs.

7. Context-Aware Identity Management Framework

Introduction

The purpose of this functional framework is to promote a consistent and interoperable approach across industry for a context-aware identity management architecture. This section includes a visualization of the context-aware environment, a proposed architectural framework, functional description of the context manager and its related elements, and a high-level flow model. It is expected that this framework can help guide future development of solutions that adopt the basic principles of the ATIS context-aware identity management architecture to produce a common set of transactions that occur between the context-aware environment and the identity management infrastructure.

Visualization of Context-Aware Identity Management Physical Environment

From a high-level, the context-aware identity management environment functions as an intermediary between the users/devices seeking to access a given resource and the application/service/network that would provide that resource. As demonstrated below, this can be best understood by considering a contextual environment that surrounds the user or device and contains specific context-level information.

Visualization of CaIdM environment

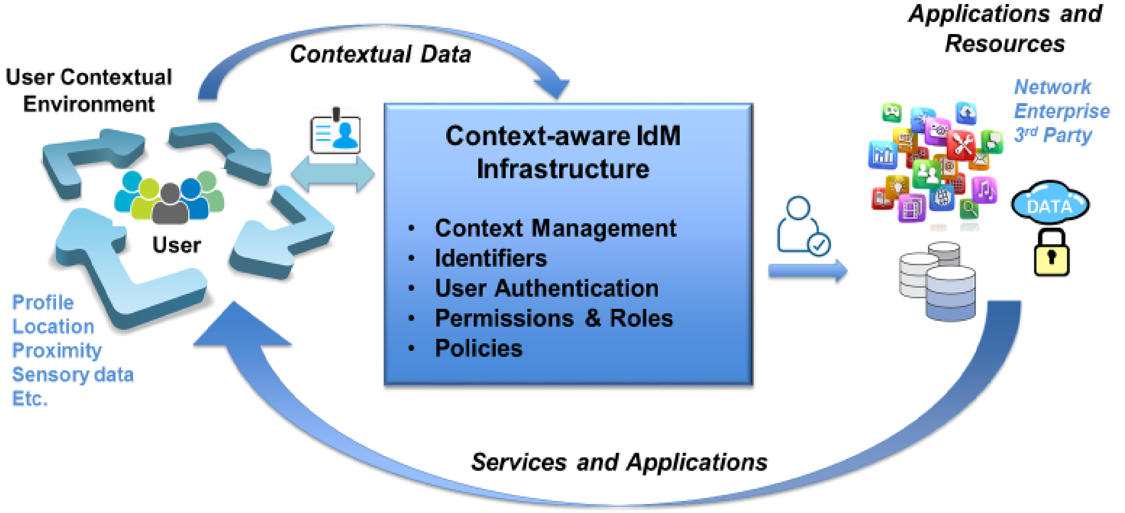

The context-aware identity management infrastructure illustrated in the middle of this diagram contains both the context management and the related policy enforcement functions, and the underlying identity management infrastructure that supports the primary functions of authentication and authorization. From a practical standpoint, it is understood that the user contextual environment and the identity management infrastructure environment may exist in the same ownership domain (e.g., network operator) or different domains.

In this visualization, the identity management infrastructure uses contextual data to supplement other identity management-related factors, such as passwords or smart tokens. Services and resources are allocated only after the application and/or device being used is properly authenticated and privileges are appropriately granted. The identity management infrastructure decides on the degree of contextual "assistance" that is required to authenticate a user/device or authorize a service or application.

Context-Aware Identity Management Framework

One of the underlying principles of this proposed architecture is the view that the context-aware domain acts as a producer of contextual data and delivers information that, after appropriate processing, forms a *contextual picture* to the consumer domain, containing the underlying identity management infrastructure. The identity management infrastructure could exist within the same ownership domain as the context manager, or it could act in an inter-domain arrangement. While it may be possible to integrate the context manager and identity management environments, the framework (as depicted below) is built upon a series of transactions that exchange contextual information in a request/response manner between the context-producer and context-consumer domains.

In many cases, the consumer domain will exist in a different ownership domain than the context manager. Therefore, a key attribute of this architecture is the exchange of contextual data between the context manager and federated, third-party or over-the-top (OTT) entities. In this way, the business environment surrounding the context manager can be viewed as delivering a set of context-aware services (context-as-a-service) for the purpose of authenticating or authorizing a user or device within a federated-domain. In addition, the context manager may provide other context-aware services beyond the scope of identity management.

Context-Aware Identity Management Functional Framework

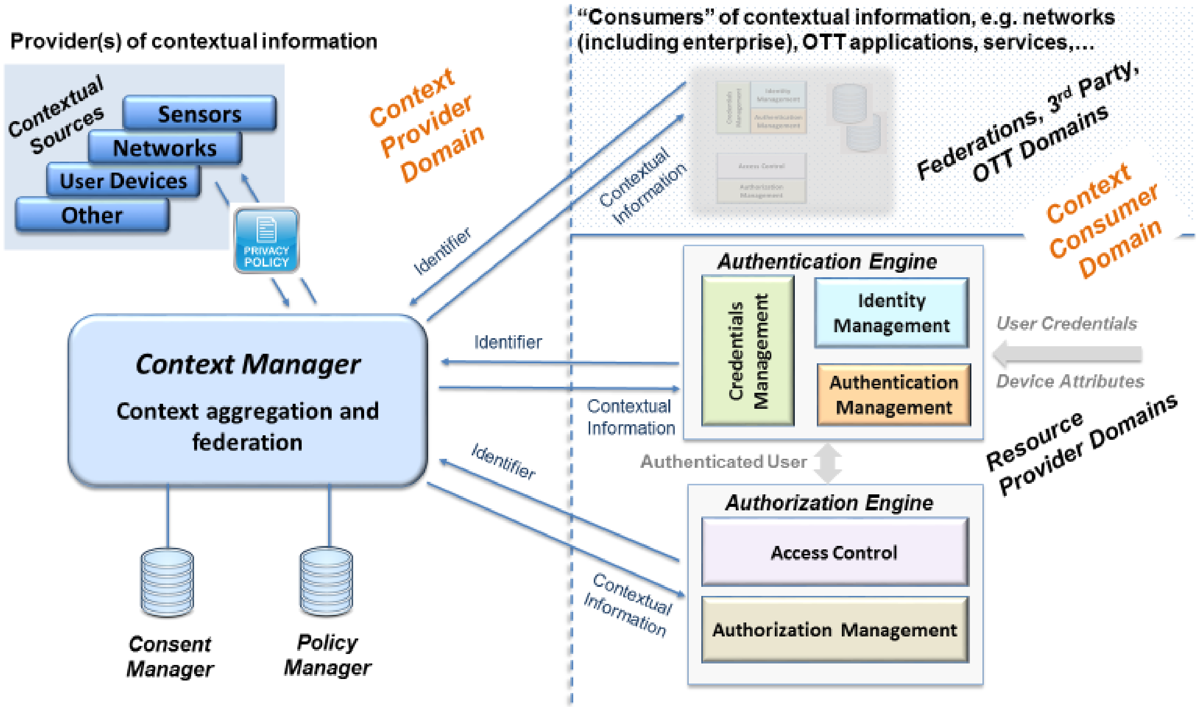

As an illustration of the flow of information through this model, the following example assumes that the context manager and identity management infrastructure reside in the same domain:

1. A user requests a resource and asserts an identity to the *resource provider*. The identity engine may exist within the same ownership domain as the context manager or it may be part of a federated, third party or OTT domain.

2. The resource provider's *authentication engine* determines, based on its policy, that contextual data is required to authenticate the user and requests contextual data (associated with a specific identifier/use/device) from the *context manager*.

3. The *context manager* associates the identifier with specific sources of contextual data and requests this information from the *contextual environment* surrounding the user.

4. Contextual data is returned to *context manager*, aggregated, formatted and returned to the authentication engine, consistent with consent and privacy policies.

5. The resource provider's *authentication engine* determines, based on its policy, if the user should be authenticated within the domain and either passes the authenticated user to the *authorization engine* or denies access.

6. The *authorization engine* may optionally request additional contextual information from the *context manager* for granting of specific privileges, roles or delegations.

7. The *context manager* performs similar actions for any requests from the *authorization engine*.

Note: The above example does not imply that the API between the authentication/authorization engine, and the context manager must follow the described data flow. The detailed API description is not part of the scope for this document.

8. Functional Elements

Context Manager

The context manager represents a key element of the context-aware identity management achitecture because it plays the central role in the context producer domain. As discussed earlier in this section, the context manager processes transactions in a request/response flow model because it receives a request for a specific identifier, retrieves contextual data from the approparaiate contextual source(s) and delivers this information to the requesting authentication or authorization engine within the context consumer domain.

Given the need to associate context sources with a target identifier, the context manager may acquire this knowledge in three ways:

1. The application, service or network making the request tells the context manager which sources to consult. This implies that the requesting consumer domain already has knowledge about the user's contextual environment.

2. The context manager has knowledge about an identifier's contextual environment and has the means to ensure that the source and the information received is trustworthy.

3. The context manager may store historical conextual data for identifiers based on policies.

The policy and consent managers interact with the context manager and the contextual sources to determine which information may be shared with the underlying identity management infrastructure or federation requesting the contextual data.

At a basic level, the context manager would process transactions and aggregate contextual data based on exchanges between context consumers and context sources, consistent with security, consent and privacy policies.

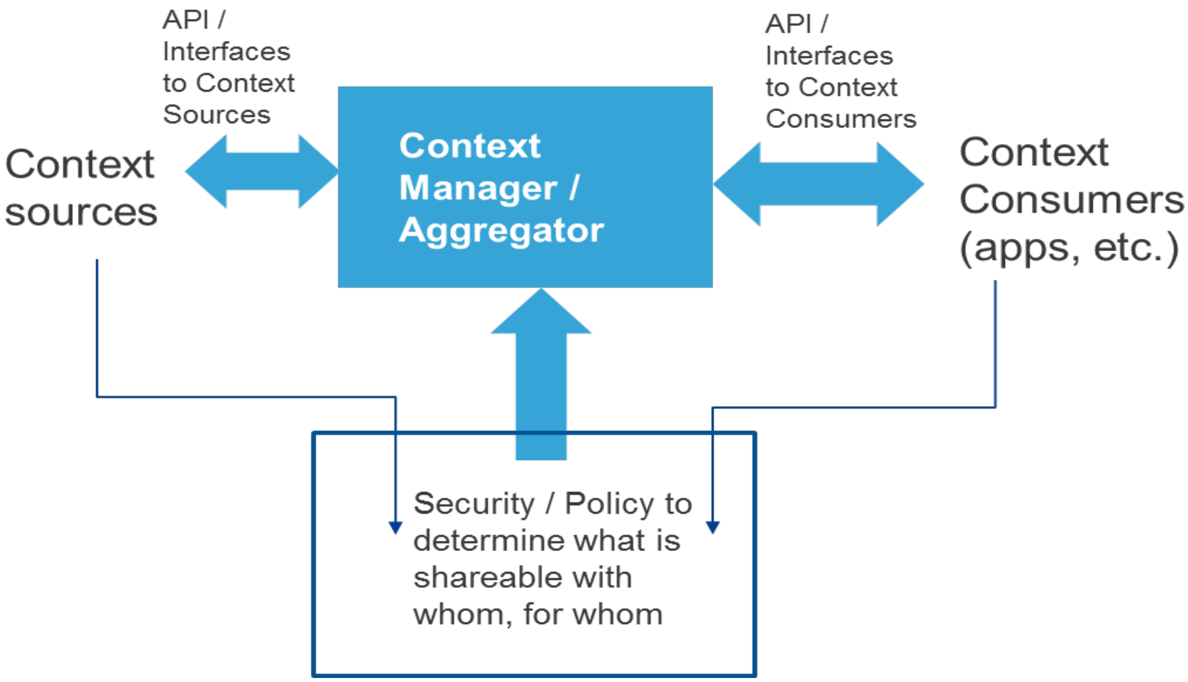
API / Interfaces to Context Sources
API / Interfaces to Context Consumers
Context sources
Context Manager / Aggregator
Context Consumers (apps, etc.)
Security / Policy to determine what is shareable with whom, for whom

This above figure represents the following set of API-based interactions between the context manager and the context sources and context consumers:

1. User/device requests service from consumer application.

2. The application determines which context is required in order to determine parameters of session (authentication, authorization, internal policies, etc.)

3. The application makes a request to the context manager for the information it needs, including asserted user identity and other information used by the context aggregator to determine:

    a. Does the application have the right to make this request for this user?

    b. What information is this application allowed to access?

4. The context manager requests information from appropriate context sources, including user ID, app ID and whatever else is necessary to determine whether the source will comply with the request.

5. Sources return the requested information, subject to their own configurations with respect to this user and this app.

6. The context manager returns the aggregated contextual picture to the requesting application, which then uses that information to do what it needs to do.

From a transactional standpoint, the following figure is an illustration of the role of the context manager in relation to the larger identity management ecosystem.

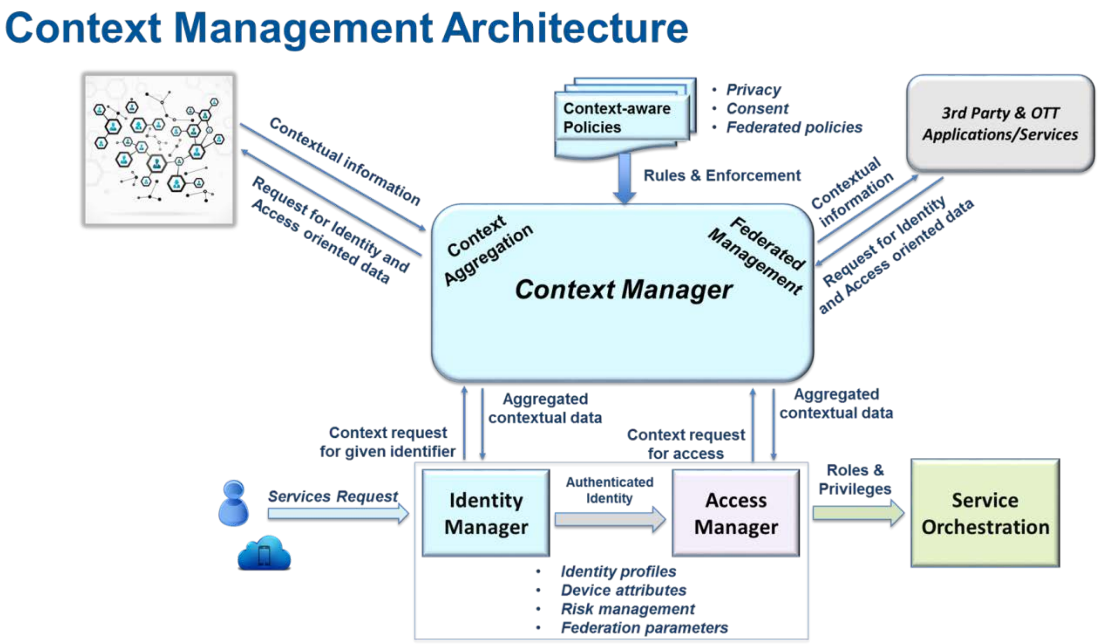

This transactional view of the context manager highlights the role of the context manager with respect in the underlying identity management infrastructure, whether it is part of the same ownership domain or is associated with a federated set of services, third party/OTT application. Although the identity management-related transactions, user/device service request and the service orchestration are beyond the scope of this architecture, they are illustrated above to show the transaction flow in the larger identity management ecosystem.

Authentication for a particular identity may not require additional context. However, context could drive the permissions for the authenticated identity incorporating place, time and environment. Optionally it could determine/modify how the network treats a request initiated by that identity given current context. For example, an authenticated first responder may only be allowed to upload video or update a database about a situation if they are located on the scene. They may also be given preferential QoS, low-latency/robust routing, additional security, if on the scene.

Contextual elements (e.g., location, proximity to other users, networks visible, etc.) may only become relevant after the user/device has been authenticated. The required subset of all possible contextual elements available might be determined only as a function of what the user/device is trying to do at that time.

Sources of Contextual Information

Contextual sources of data receive requests from the context manager and respond with data consistent with privacy and consent policies enforced by the policy manager.

Due to potential different organizations/companies generating the contextual data and the aggregator/federator, consent is needed from the user for every different source. Examples include location from the mobile network operator (MNO) or presence status from the messenger.

Policy Manager

The policy manager has overall responsibility for enforcing polices, consent and privacy across the context-aware identity management ecosystem. This includes privacy policies at the ingress to the context manager (from context sources) and at the egress of the context manager (after aggregation and formatting of data) to the identity manager of federated entity.

This architectural view of the policy manager assumes two levels of policy management and enforcement within the context domain:

1. A general "domain"-level policy that includes factors, such as required level of aggregation or anonymity of shared contextual data between context and a given consumer domain, or actions to be undertaken in the event that user-specific information is inadvertently discovered.

2. A specific user policy based on a user consent policy that is based on which type of information may be shared with a given consumer domain or application.

Consent Manager

The consent manager is responsible for informing the context manager (i.e., source of aggregated data) of user consent on a per-application basis (e.g., data may be used by MNO but not by web application X). User consent is needed to authorize which type of contextual data can be used by the application (e.g., location can be shared with financial institution, but not social media applications).

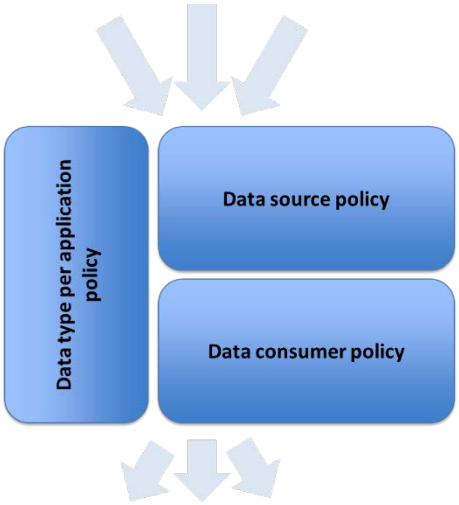

In most cases, the user consent policy will be stored in a consent repository. User consent is required to allow contextual information to be exchanged with the context manager and delivered to the consumer domain.

Given the increased focus on user and data privacy, it is recommended that the consent repository and the identity management in general engage with leading privacy advocacy groups and ensures compliance with current privacy practices. The White House published the *National Strategy for Trusted Identities in Cyberspace*[2] in April 2011, setting a national security priority focused on privacy and civil liberties with regard to identity. This led to the creation of the Trusted Identities Group within NIST[3] and the creation of *Special Publication 800-63, Digital Identity Guidelines*. In 2012, the release of the *Consumer Privacy Bill of Rights* outlined what users should expect in their service providers. In June 2016, the White House released the *National Privacy Research Strategy*,[4] which highlights critical goals and objectives for information security systems regarding privacy. Any context-aware identity management should be cognizant of the previous work and the challenges and to the greatest degree possible address and

[2] https://obamawhitehouse.archives.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf

[3] https://www.nist.gov/itl/tig

[4] https://www.nitrd.gov/PUBS/NationalPrivacyResearchStrategy.pdf

comply with the intent of these policies and similar policies, including the European Union’s General Data Protection Regulation (GDPR).

The value of anonymity in this solution space is an important consideration. Individuals wishing to remain anonymous have privacy rights, and a key element of any identity system is the awareness and degree to which it supports anonymous activity. It is reasonable to expect an identity attributable to a specific individual for critical services, however when anonymity is not possible the end user must be made aware of this restriction. Consequently, a user who requests anonymity must have those wishes respected but also necessarily be restricted from accessing services that require assured identities.

Third-Party & OTT Applications/Services

Federated services and applications that reside in a separate domain from the context manager will request contextual data associated with a specific identifier. The returned data will be used for identity and access control operations that exist within the third-party or OTT domain. The policy manager should enforce privacy and consent policies in accordance with user opt-in and federation requirements.

9. The Future Role of Context-Aware Identity Management

Market Opportunities

As the sources of network-accessible contextual data increase, it is expected that identity management architectures will take advantage of this information to enhance risk assessment techniques and create more robust security mechanisms for access to resources and services.

Over the next few years, it is envisioned that four primary applications will emerge around the intersection of context-awareness and identity services:

1. Context-Aware Authentication

One of the primary challenges being faced by multi-factor authentication (MFA) approaches today is the balance between security robustness and user acceptability. Each additional level of MFA adds another layer of complexity from the user perspective, in terms of information that a user may know (such as passwords or profile data) or something the user may have (such as secure tokens, smart cards, wireless keycards). In fact, the collision of increasingly more robust passwords and lack of user appetite for more complex authentication patterns has led to the emergence of passive authentication methods, which rely on factors such as user keystroke behaviors or other biometric factors.

MARKETSANDMARKETS EXPECTS THE GLOBAL PASSIVE AUTHENTICATION MARKET TO GROW FROM USD 440.0 MILLION IN 2017 TO USD 1,535.0 MILLION BY 2023, AT A COMPOUND ANNUAL GROWTH RATE (CAGR) OF 25.5% DURING THE FORECAST PERIOD.

Context awareness will act as a powerful enabler, supplementing existing and future identity management infrastructure. Context can be gathered in the background, supplied to the identity manager and can create a contextual picture and corresponding risk factor. If the risk score exceeds a specific threshold, a user or device may be prompted for an additional level of MFA information. Similarly, if the contextual picture validates the user or device, the user experience can be greatly improved over conventional MFA approaches.

2. Context-Enhanced Authorization

Authentication
Who you are
Authorization
What you can do
CONTEXT

Access control and authorization approaches have evolved using relatively static data, as authenticated users and devices are often granted specific privileges and permissions based on stored profiles and roles. In some cases, privileged access to specific resources may be allowed or denied based solely on the domain of the requesting user or device.

Context-enhanced authorization opens a new realm of dynamic capabilities, as privileges and permissions can be altered as the contextual environment around the requesting entity shifts during a session. For example, authorized services may be granted based on location of the user, proximity of the device to other devices, behaviors of the user within a given environment or physical factors surrounding the user. All of these context-aware functions may be acted upon by the authorization manager without the direct intervention of the user. In this way, an authorized user achieves a better experience, and the identity domain is enhanced through a more robust solution.

3. Context-Aware Services

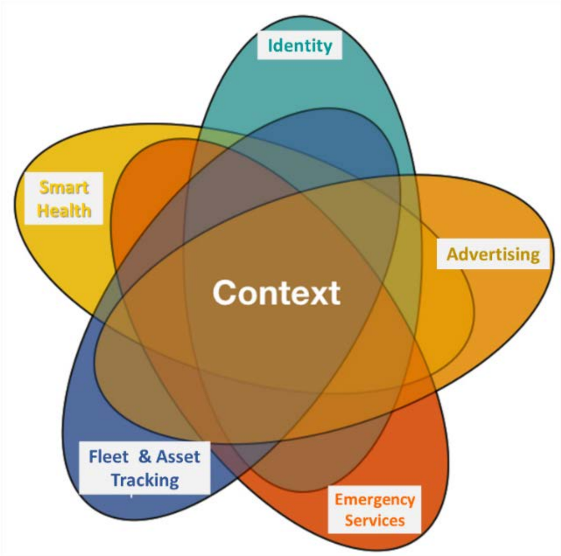

The context manager plays a key role in the emergence of context-aware identity management solutions. As described earlier in this document, the context manager may communicate with the underlying identity infrastructure within the same domain or may provide context as a service to third party or OTT entities. Although the context manager has been described within the scope of identity management in this report, there is a growing market for context-based services to support a broader set of applications. Therefore, it is expected that context manager investments can be leveraged by a broad array of services beyond identity management. These context-based services may include applications such as healthcare, transportation, marketing, advertising and public safety. While these applications are outside the scope of this report, they do provide reinforcement around the viability of a context manager platform.

Federated access to context-aware information is a key element of the framework described in this report. It is expected that many federated entities (third party, OTT, etc.) will gain value in acquiring contextual data from another entity (e.g., network provider) for the purpose of authenticating users and/or devices, or authorizing access to federated resources, without the need to invest in a context manager platform. As described earlier in this report, privacy and consent will need to be enforced across the federated and content manager domains as part of the policy management function.

4. Applications Across Adjoining Verticals

One of the major benefits of context-aware identity management is the role in authenticating or authorizing roles or privileges for applications across communications networks and verticals or between vertical markets.

User experience for identity management across adjoining industries will greatly impact adoption rates. For example, user experience across connected vehicle applications could offer seamless support of vehicle diagnostics, vehicle history, warranty services, infotainment, communications networks, etc. Applications that span multiple vertical markets offer the opportunity to gain access to a much greater level of contextual information. Identity management applications can leverage this data and create a much-improved user experience (UX) environment.

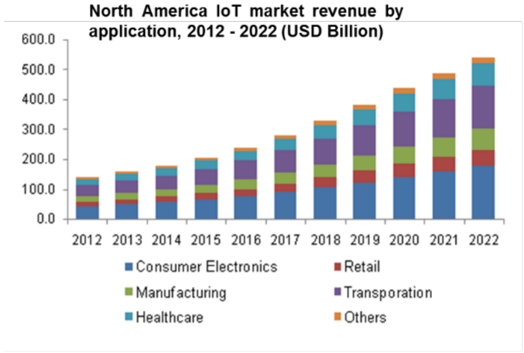

Source: Grand View Research

10. Findings and Next Steps

This report provides a detailed assessment of context-aware identity management approaches that can provide additional robustness to existing authentication and authorization infrastructures used by network operators, enterprises and third party/OTT entities. A list of significant findings is included below covering the business implications, framework and policy and privacy considerations associated with this architecture:

Future Business Landscape for Context-Aware Identity Management

- Security and privacy frameworks to authenticate users and authorize access to services and applications represent major investments by the industry.

THE VALUE OF CONTEXT-AWARE IDENTITY MANAGEMENT WILL GROW AS THE AVAILABILITY OF CONTEXTUAL INFORMATION EXPANDS.

- Network operators, verticals and application providers are already assessing more robust solutions to meet future identity management and access control needs.

- Given the fact that contextual-based solutions are already emerging at the device/application level, network-sourced contextual data can greatly expand the solution space, resulting in new business opportunities.

- The context manager can provide value-added capabilities on top of existing identity management solutions and can function in a producer/consumer model with identity management needs.

- Future delivery of contextual data can be viewed in a “context-as-a-service” sense to enhance identity management functionality and may be part of a wider array of contextual services offered by a context manager entity.

Architectural Framework for a New Context Domain

- This report defines the *context domain* to include the context manager, the related consent manager and policy manager, and the contextual information sources.

- Similarly, this report defines the *consumer domain* to include the underlying IdM infrastructure associated with a network provider, third party, federation, OTT app, etc.

- This report focuses on a request/response model between the context manager and the context sources.

- The context-aware identity management framework depends on a user or device asserting an identify for a given application and the identity management infrastructure conveying an" identifier" to the context domain.

THE CONTEXT MANAGEMENT ARCHITECTURE CAN BE OVERLAID ON TOP OF EXISTING AND FUTURE IDENTITY MANAGEMENT INFRASTRUCTURE AND ACT AS A NEW CONTEXT DOMAIN.

- The request to the context domain may optionally consist of either a request for specific contextual information or a request to verify user/device contextual information provided to the context manager.

- The content domain returns a contextual picture to the identity management consumer domain based on acquisition and aggregation of relevant contextual information.

Privacy and Policy Implications

- The policy manager exists within the context domain to enforce user/app level policies and domain level policies related to the requesting identity management consumer domain.

PRIVACY AND SECURITY SHOULD BE AT THE FOUNDATION OF ANY CONTEXT-AWARE IDENTITY MANAGEMENT SOLUTION.

- The consent manager acts as a repository of user/device/app level opt-in policies.

- The context source must include a set of security/privacy settings for the user and the requesting application.

- Given the diversity of IoT context sources and complexity, it is understood that the consent and policy managers will play a pivotal role in assuring enforcement of applicable security, privacy and opt-in decisions.

- Consent must be administered down to the application level in order to support multiple user-level opt-in policies associated with a given user or device identity.

Next Steps

Following the publication of this report, ATIS will:

1. Promote this report with the industry and with other standards development organizations engaged in identity management solutions.

2. Integrate the concepts of context-aware identity management into ATIS' larger cybersecurity agenda and next-generation network efforts.

3. Position this work as a valuable input to future requirements and study area development efforts associated with 3GPP.

4. Explore opportunities with industry to collaborate on concept or prototype testing and validation of these concepts.