IOT-BASED BUILDING MANAGEMENT

SYSTEM RESILIENCY

Describing the inherent resilience of IoT-based Building Management Systems

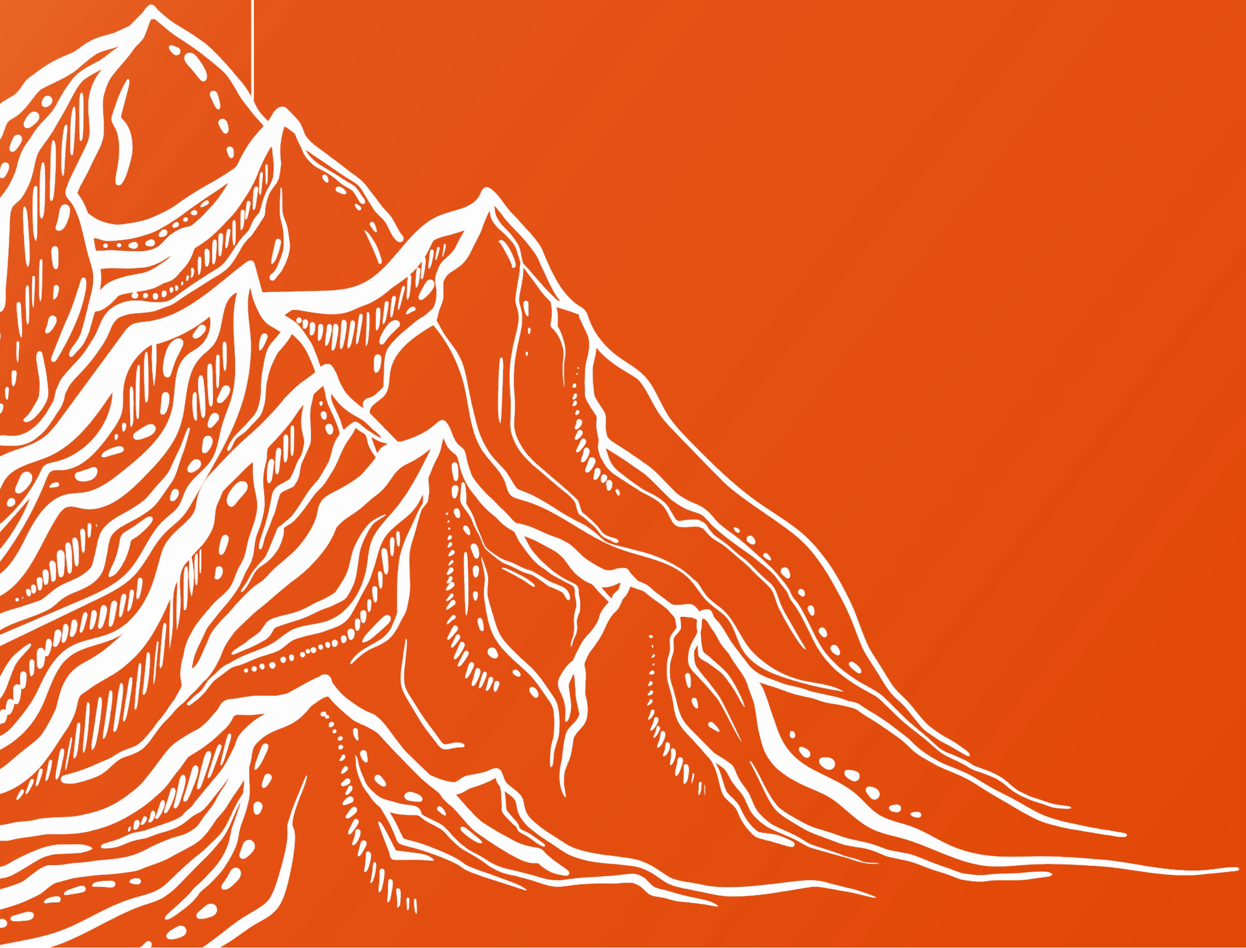

75F®

INTRODUCTION

The advancement of cloud technology and its ability to adapt to specific industry needs has revolutionized the way many businesses operate. From Google Nest in residential buildings to smart watches, cloud technology is well developed and stably deployed across many industries and countless B2B and B2C products. Yet, the commercial Building Management System (BMS) industry sticks out for its laggard adoption of cloud-native technology.

The BMS has been around in one form or another for decades — and the technological foundation of those systems in today's buildings are largely the same as they were in the 1980's and '90s. Without the flexibility of the Internet of Things (IoT), these traditional systems are rigid, highly customized, intensive, and expensive. While the traditional BMS is increasingly included in the project scope for many new, large developments and retrofitted into some commercial buildings for better efficiency and comfort, their functionality is largely offline and more complicated than necessary considering the power of modern technology.

Solutions using cloud computing, IoT-connected smart sensors, remote monitoring, meaningful reporting and analytics, and AI-tuned controlling are major tools for today's building management business leaders, who strive to keep costs down while demonstrating responsible and healthy management of their spaces to occupants and investors. With an IoT-based BMS, building operators have all these tools and more at their fingertips.

However beneficial these systems are, it's common for those unfamiliar with the IoT-based BMS to wonder how stable a system can be if it relies on an Internet connection to function. In short: The IoT-based BMS does not require the Internet to function at its core.

This white paper is the long answer to this question. What factors contribute to the stability of an IoT-based BMS? In the following chapters, we'll address:

- Inherent benefits of cloud-native building management compared a traditional approach
- The cloud environment: Why is it better, and is it secure?
- The system architecture of an IoT-based BMS
- System redundancy in each layer of the IoT-based BMS
- Disaster recovery should any layer fail to function
- Denial of Service (DoS) or Distributed Denial of Service (DDS) attacks and system response
- Basic security protocols at each stage of the product lifecycle

With the advent of cloud computing, access to more computationally intensive software, data storage, and data access is easily achieved. This allows devices to deliver functionality that would normally only be possible with high-end hardware and analysing aggregated datasets.

In an IoT-based BMS, Data Inputs From Sensors, equipment, and occupants move wirelessly through an Onsite Gateway to the cloud, where built-in sequences tailored to the building's equipment then make informed control decisions. These decisions travel down back to the building in the form of output decisions or micro adjustments seamlessly. Because the system is based in the cloud, both monitoring and updates can be carried out remotely.

Data that streams from an IoT-native system is clean from the ground up, as well, providing vast amounts of clear, readable and actionable data from our commercial buildings for the first time. This powerful dataset can automatically update the Digital Twin and gives building operators unprecedented access to data analytics that can be easily understood and securely shared with internal and external parties.

With automatic and industry-standard tagging, the opportunity for integration is endless. When combining the digital twin with deep learning approaches, complex detection, prediction and optimization functions can be achieved. This is very difficult to do with a traditional BMS.

In most buildings with a traditional BMS, the system runs with a number of control devices and a PC-based server or proprietary software suite in a back office. Just as emails were once run through a company's own on-site server, this style of BMS is increasingly obsolete as it struggles to keep up with growing demands for meaningful analytics, solutions attainable to a wider market, and more.

Instead, like so many other aspects of our technological lives, the future of building management lies in the cloud considering all its advantages. With practically unlimited storage, a cloud-based BMS can handle vast quantities of data, improving the customisations and scalability that it brings with it.

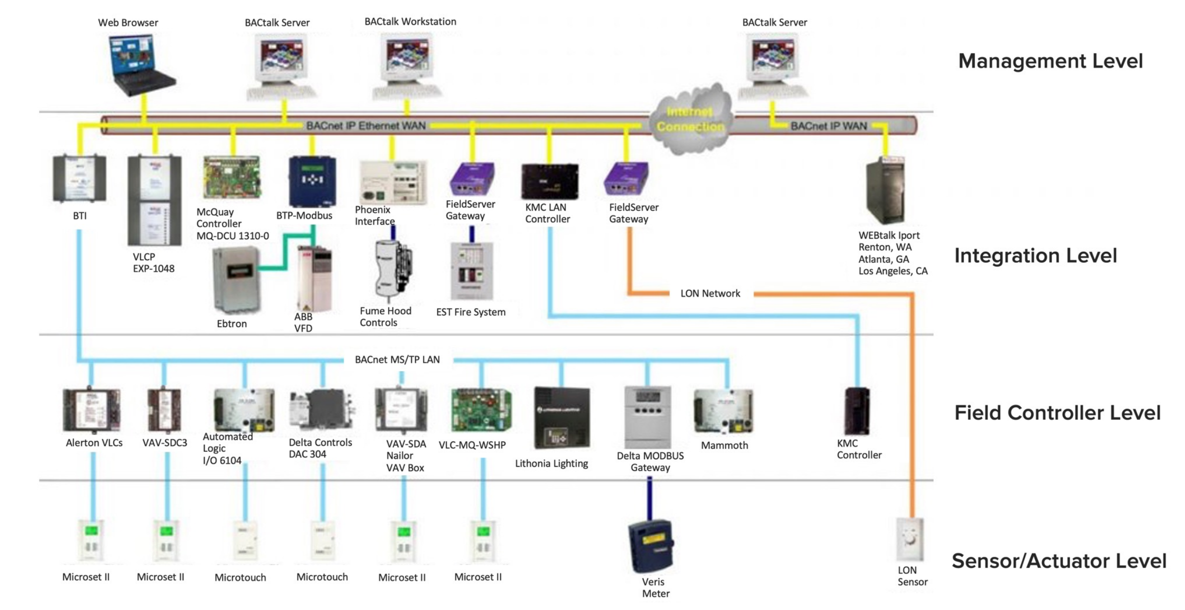

An example of the layers in a traditional BMS system architecture.

ADVANTAGES OF THE IOT-BASED BMS

Centralized Control Over Multiple Sites

At present, the majority of BMS systems are in owner-occupied buildings. Updating to a cloud-based BMS is one of the most meaningful changes a building manager overseeing multiple properties can make. While the traditional BMS system may only be accessed on site, moving to the cloud allows building managers to access multiple sites simultaneously and from any location. Building managers can securely access their buildings from anywhere and via any device. Most cloud services include streamlined features for easy client control and energy management, allowing building owners and building managers alike to monitor and adjust as needed. Client-focused dashboards are easily created, displaying the system information by client preference and in a uniform way across the portfolio.

Ultimate Efficiency with the Cloud

The superior data allowance in an IoT-based BMS system not only allows building managers to expand their IoT capabilities, but also improves flexibility of access and analytical power.

Cloud-based BMS interfaces allow the user to access the building's IoT from any internet-enabled device securely and from any location. With all data obtainable digitally, building owners and building managers can monitor and control things like metering, lighting, and HVAC from practically anywhere. The 75F platform also allows third-party API integration so other energy apps and building systems can be fully integrated — for example, making one single cohesive dashboard with all systems integrated. No more depending on phone calls to local facility managers to figure out how things are working (or not working). No more making site trips to get to the bottom of those comfort complaints. Facility managers already know what the problems are and how to fix them.

It's also possible to see data in real time, revealing significant metrics and trends as they happen. And, with expanded storage capacity, analytics can be run on huge collections of data, identifying new ways to use energy more efficiently and reduce costs.

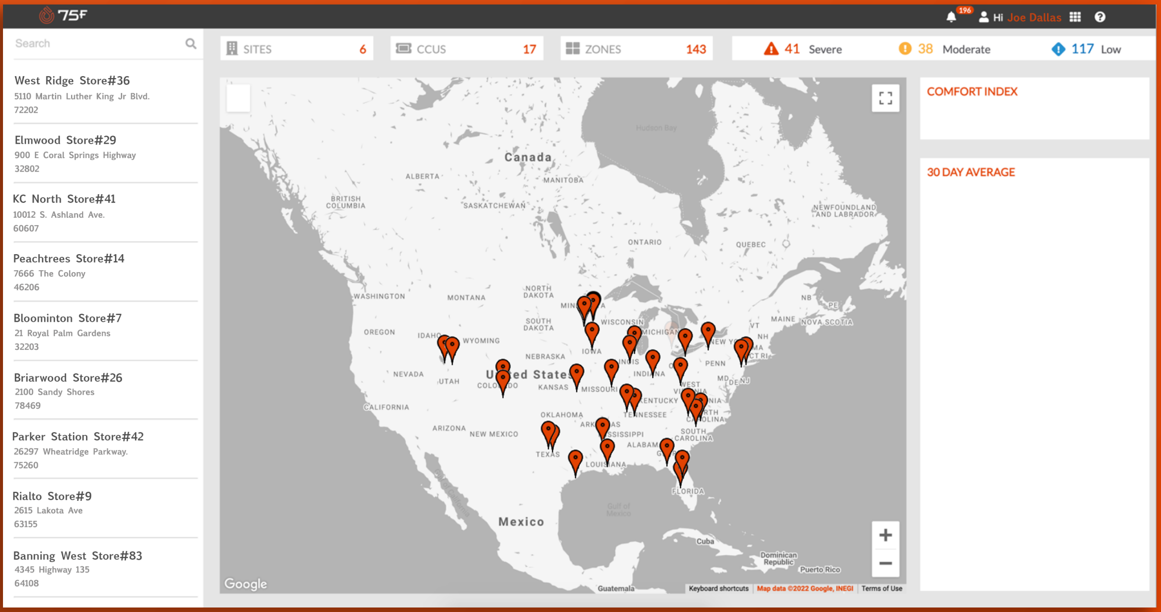

West Ridge Store#36
5110 Martin Luther King Jr Blvd.
72202

Elmwood Store#29
900 E Coral Springs Highway
32802

KC North Store#41
10012 S. Ashland Ave.
60607

Peachtrees Store#14
7666 The Colony
46206

Bloominton Store#7
21 Royal Palm Gardens
32203

Briarwood Store#26
2100 Sandy Shores
78469

Parker Station Store#42
26297 Wheatridge Parkway.
75260

Rialto Store#9
2615 Lakota Ave
63155

Banning West Store#83
4345 Highway 135
64108

An example dashboard in 75F's user portal, Facilisight.

Remote Support for Stress-Free Building Management

By utilising a cloud-based BMS system, there is no longer a need for a dedicated computer on each site, nor is there a need for callouts in case of technical problems or issues in reporting. Instead, BMS experts are given permission to access your system remotely and support you with any problems as they arise.

With remote support, users can not only manage their building from any location, but also find, troubleshoot, and resolve problems from a laptop or mobile device.

Predictive Maintenance & Algorithm Tuning

Predictive maintenance is an invaluable feature the IoT-based BMS accomplishes with constant equipment monitoring and real-time data streams. IoT-based systems can predict a potential equipment issue, raise alerts, and provide insights into building performance over time. With this tool, facility managers can schedule maintenance before a critical equipment failure even occurs, saving valuable time, money, and comfort disasters. The predictive nature of IoT-based systems extends to everyday comfort, too. Predictive control sequence tuning enables the system to proactively control how equipment operate, eliminating hot and cold spots before they appear.

Summary of Key Advantages

- **Continuous Improvement** — Ability to peek into individual or overall building operations to evaluate energy costs and compare performance across sites.
- **Proactive Maintenance** – Proactive system monitoring with centralized dashboards across an entire portfolio, allowing comprehensive alerts and preemptive equipment service.
- **Occupant Comfort** – Ability to keep everyone, everywhere happy at all times. An IoT-based BMS uses predictive control based on weather and building conditions to ensure optimal comfort without human intervention.
- **Real-Time Alerts –** Unexpected events and out-of-spec conditions can trigger automatic alerts that give facility managers the time needed to take swift action and solve problems.
- **Improved Services** – Arming the right people with the right information. Overall building services are improved when decision makers, from local maintenance to corporate planners, are kept in the loop with valuable information sharing.
- **Customized Reporting** — Different Data Logical Views of the same data. A facility manager might want to see which sites are generating the most alarms and complaints, and on the other hand, the energy manager might want to compare energy usage across all sites to see which ones are the biggest energy hogs.
- **Crowd Sourcing** – The cloud makes it easier to share information and enable "bottom-up" solutions. Be it a control sequence or an analytics dashboard, easy sharing ensures all stakeholders are involved.

Buildings managed by an automated system needs to operate reliably irrespective of the option that is chosen for its data access and storage. Cloud data management on a BMS can be in n a public cloud environment — such as AWS EC2, Azure, or Google Cloud Platform — or any private cloud hosted.

There are several approaches to protecting a cloud BMS from downtime and disasters, but determining which approach is best depends upon multiple considerations. For any real-time distributed BMS, understanding the below three factors clarifies which path is best:

Using Fault-Tolerant (FT) & High-Availability (HA) Solutions

FT vendors will guarantee the applications they protect will have three-nines availability. Three-nines availability — 99.9% — allows 8 hours and 46 minutes of downtime per year.

Disaster Recovery (DR) Solutions

DR solutions are a related but separate approach to operational and management continuity. A DR solution is designed to ensure that operational data and infrastructure can survive a literal disaster like an earthquake or flood. A DR solution would replicate the infrastructure for the BMS in a geographically distant region that is unlikely to be affected by whatever disaster takes down the primary BMS infrastructure. With a DR solution, facility managers could run the BMS from the remote region or restore it when the local infrastructure is online again.

Cloud Considerations

From an application protection perspective, the cloud offers certain distinct advantages over an on-premises deployment. Building operators don't need to worry about acquiring the hardware, data center, real estate, or support personnel to manage and administer the infrastructure. The cloud service provider takes care of all of that on your behalf, and their operations are optimized to perform those tasks much more cost-effectively than anyone else could.

FT infrastructure in the cloud will still cost far more than HA infrastructure in the cloud, and it's only available through private cloud offerings sponsored by FT hardware providers. Either way, it's the service provider's responsibility to spin up the systems that will support the BMS solution. Particularly if you're using standard system configurations to design an HA infrastructure on AWS, Azure, or GCP, those services can spin up virtual machines to your specification in moments — and resize them instantly if your future needs demand a more powerful configuration.

Since cloud service providers operate their own data centers, it's easy to configure an HA failover cluster with VMs in separate cloud availability zones. You also gain 24/7 administrative services at a far lower cost because you don't have to hire your own teams of system administrators who can work 24/7. The teams supporting the cloud data centers are working 24/7, but because the cloud service providers can spread the cost of those teams across hundreds or thousands of clients, they can provide the same level service at a far lower cost to each client.

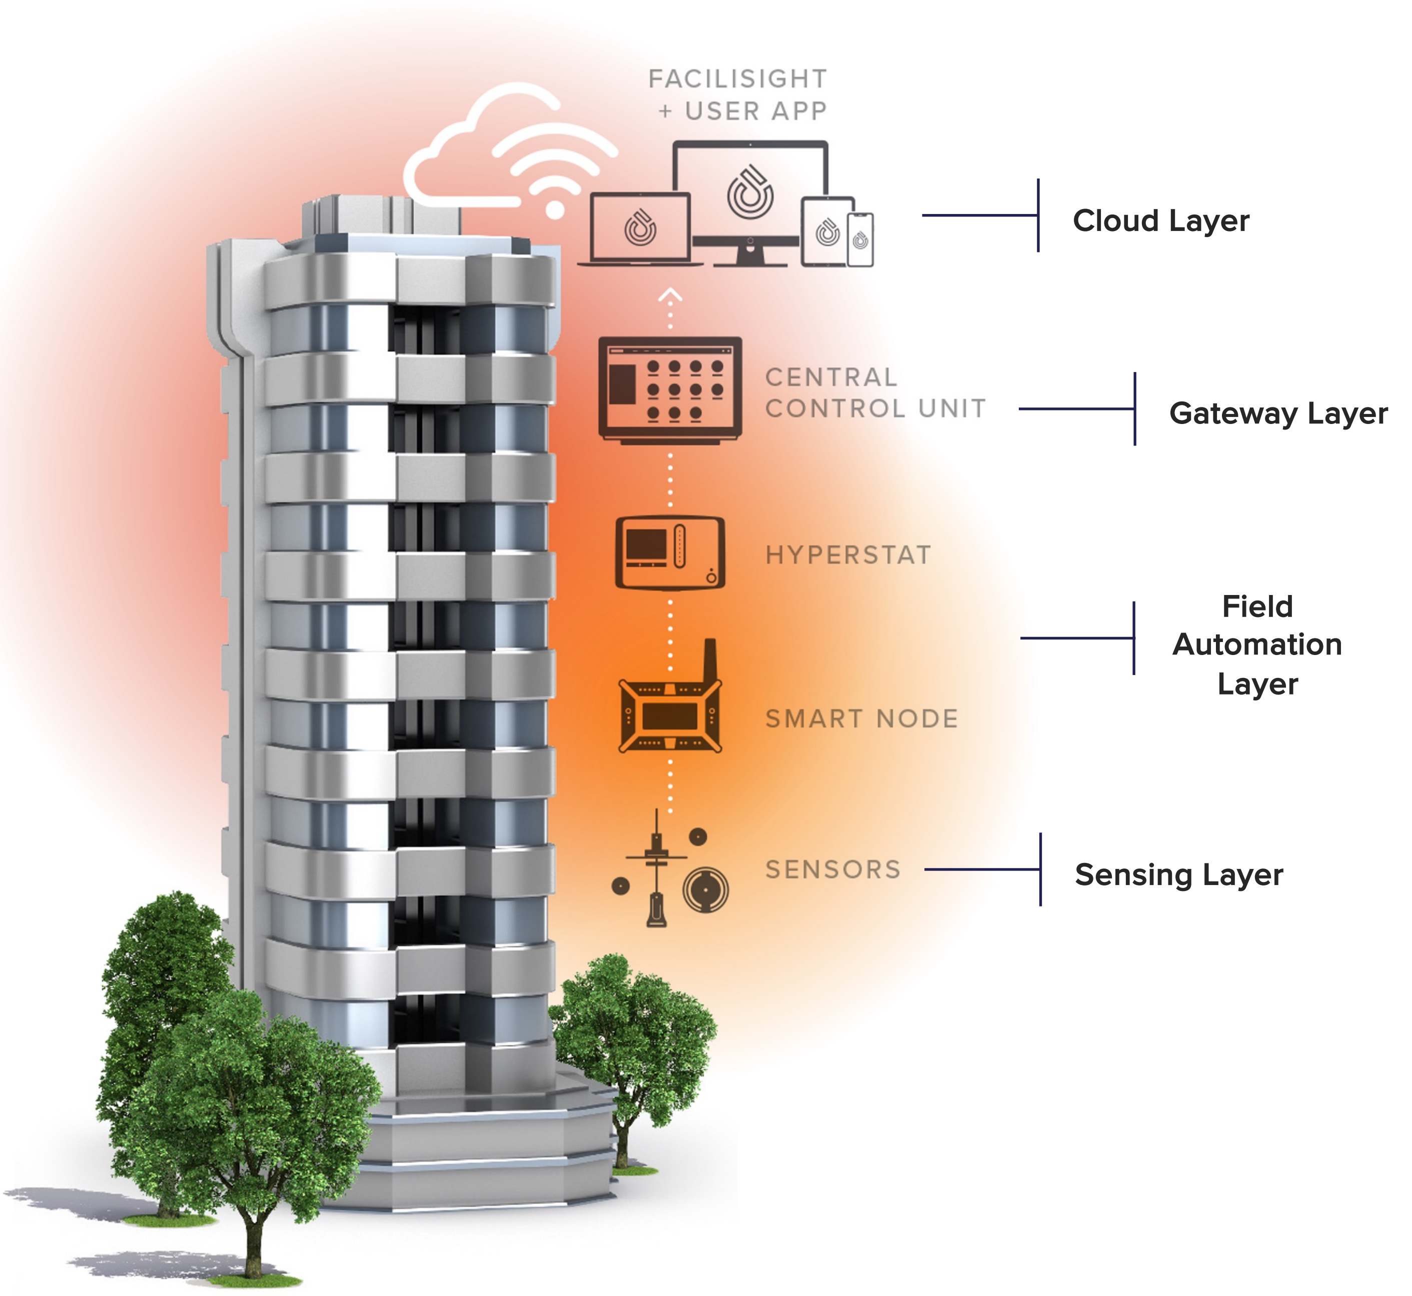

IoT-Based BMS System Architecture

75F’s IoT-based BMS encompasses all the advantages mentioned above with a highly scalable, fault tolerant, and always-available platform. In this section, we’ll describe how the system architecture is structured.

75F is a complete platform built on a Distributed Solution framework and leverages a flexible, microservices-based architecture to address how the system handles its various data layers and its redundancy handling mechanisms. This framework consists of four subsystems: device sensing layer, field automation layer, gateway layer, and cloud layer.

Device Sensing Layer

Data acquisition depends on local end sensing hardware capabilities. Typical data types collected from the BMS include all sorts of IAQ points including CO_2, volatile organic compounds (VOCs), occupancy, and more. The BMS will also collect current, voltage, time, location, ambient temperature, cell temperature, and the communication address of a cell or module. 75F's smart sensors use a proprietary 900 MHz wireless mesh network to communicate to the gateway layer.

Field Automation Layer

Field Controlling layer has terminal devices, which communicate and interact with the room control equipment and gateway systems. In most control path cases, monitoring and control functions are time critical, and when compounded by bandwidth limitations and redundancy support, edge computing functionality is the most computationally efficient structure to reduce mesh traffic in the 75F platform and cloud load. The 75F platform's edge computing framework consists of the Smart Node, HyperStat, and Helio Node, which all work as equipment controllers. These devices also use the wireless mesh network to communicate with the gateway layer.

Management Gateway Layer

All terminal devices deliver sensor and status data to this supervisory layer, which then facilitates the data's connection to the cloud. This layer runs smart algorithms on the aggregated data in order to control the equipment in the building centrally. The 75F's device on this layer is called the Central Control Unit, a wall-mounted tablet running Android. It also contains a Control Mote that provides inputs and outputs for connecting to central plant equipment.

Cloud Layer

75F's cloud layer is a distributed micro serviced system. The cloud system decomposes large data processing services into numerous small micro-services. 75F's cloud is easy to maintain and scale, with configurable storage, load balancing, and auto-scaling processing capacity. It enables the system to realize the value of real-time data analysis using IoT devices. Compared with gateway and edge computing, cloud computing provides more durable data storage and more powerful computing resources.

75F Cloud support comes in two flavors: shared and private. Shared clouds are the common cloud instance and share resources among clients, though the option to use a private stack is available. Private clouds are dedicated resources for individual clients. While a private cloud costs more than its public counterpart, it gives clients more control over the technology and security systems used, making it easier to standardize security processes and maintain accountability per the client's needs.

Apart from the advantages mentioned above, the 75F's cloud layer also combines collected data from 75F devices and similar external data sources with advanced deep learning algorithms to improve core control functions . The data will then be available to the user interface for visualization and will be backed-up for remote disaster recovery.

However, each layer in this framework works at being resilient to any faults that can occur and are reductant so that each layer can support itself. Each of the following sections will speak about how the 75F architecture across its platform supports redundancy at each layer, disaster recovery plans, the company's response strategy to DoS or DDoS attacks, and the security procedures in place for the whole platform.

75F's redundancy policy is about having backup components in the system that can take over if any layer in the platform fails. Any component that is critical to the 75F platform's operation and equipment control paths and could be a single point of failure with potential to stop the whole system's functionality has been made redundant. Each hardware or software layer is redundant to support the robust fault tolerant architecture.

Points of Failure Evaluation

75F's goal with redundancy is to eliminate single points of failure and provide reliable uptime of the platform. Team engineers review the platform as a whole and consider the potential consequences of any one layer failing regularly. Some primary consequences of unplanned downtime are measured in lost energy savings, building controls equipment performance, and occupant comfort, among others.

The higher the business costs of these factors, the more likely a layer has been made redundant — the overall goal is to avoid any price of failure.

Another factor that affects redundancy requirements is how long it takes to restart and restore the platform layer's component if it stops. Backup and restore is covered in more detail in the next segment.

The 75F system architecture allows for redundancy and resiliency to be easily maintained with a focus on software frontend and backend, WiFi/Internet communication, hardware architecture, and wireless mesh network communication.

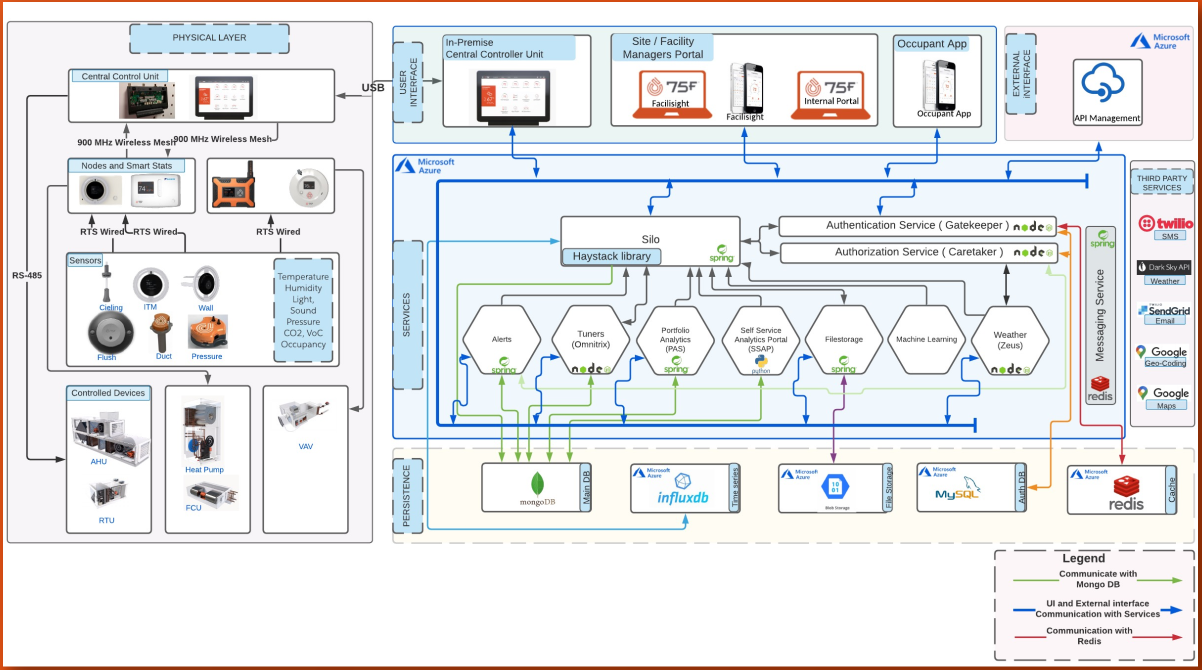

An example of a distributed microservice-based system architecture.

Software Frontend Redundancy

All front-end hosted web services within the 75F platform have geo-replication enabled so 75F can initiate a geo-failover to a geo-secondary in a different Azure cloud region. The geo-location replication support will make sure all web services are available all the time without any downtime.

Software Backend Redundancy

Cloud-hosted microservices in distributed architecture easily provide alternate paths for data to travel along in case any of the layers are affected, and most importantly, ensure the BMS runs smoothly even if any of the services and their communications are broken. Microservices enable only one smaller feature association with one microservice, so an unavailable microservice will affect only one small system capability. Even if one microservice is down, load balancing and auto scaling features of cloud platforms like Microsoft Azure will ensure there is an auto scaled instance that is always available.

Wi-Fi / Internet Communication Redundancy

The goal of consistent Internet cloud access is to prevent loss of cloud updates to enable remote control and continuous monitoring. In the case of a lost Internet connection on site, the CCU will continue to execute core algorithms so the building will remain unaffected.

Hardware Architecture Redundancy

In the physical layer, zone controllers like the Smart Node, Helio Node, and HyperStat have control algorithms coded into them. These take effect if communication to the CCU is lost so building equipment continues function as needed. This ensures building control and operation is continuous regardless of layer failures.

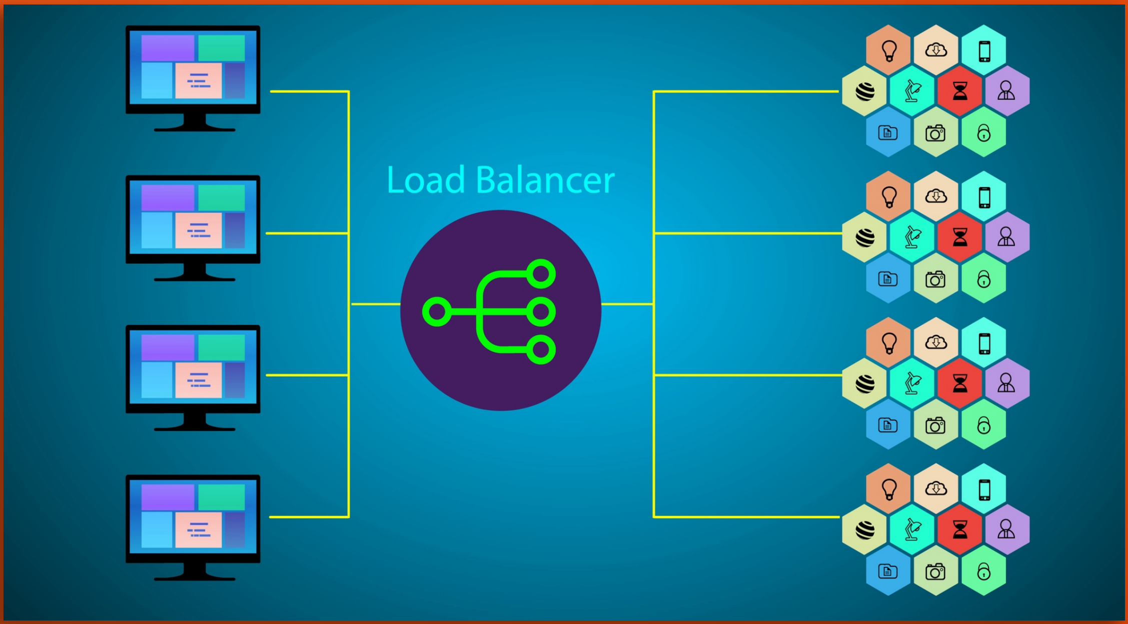

An example of load balancing in microservices architecture.

Wireless Mesh Network Communication Redundancy

Having a stable network connection for data transmission is very important for effective real-time data transmission between the gateway layer and the field automation/controller layer. 75F using wireless mesh layer in its platform allows for mesh network clusters to be formed, in which nodes form connection with as many other nodes as possible in a dynamic and non-hierarchical form to efficiently transfer data. In order to improve the reliability of the whole system, the functions at each Node/Stat in the operation process should be run locally. Then, data transmission, local update, data synchronization and local computing functions can be all realized in the edge Nodes and Stats. Even in the case that mesh communication is lost to the Control Mote board, the Node/Stat already has all the code to run the algorithm locally for the fail-safe mode to take effect.

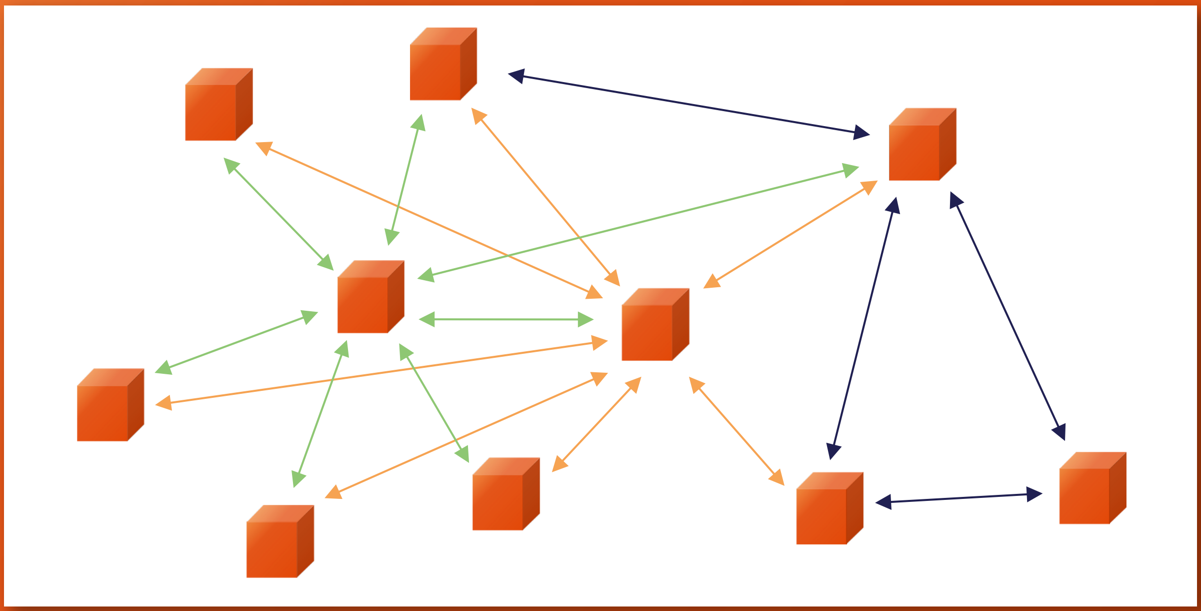

An example of a wireless mesh network topology.

75F Disaster recovery protocols are put in place to specify the way in which we resume regular operations after a disaster. This is typically accomplished through the resumption of all essential activities across our layers and the processes and systems used to support them.

For example, part of most disaster recovery plans involves regaining access to data, software, hardware, equipment, connectivity, and power. 75F's monitoring dashboards and alerting systems enable for easy monitoring of these data across a building and ensure that all backed up data is available for any replacing of controllers.

Disaster recovery must go according to a detailed, documented set of procedures designed to minimize the amount of time it takes for the platform and its supporting infrastructure to completely recover and be functional again.

How Does Disaster Recovery Work?

Disaster recovery depends on replicating data and all our product services in an area that will not be impacted by the disasters in question. In the event cloud service goes down due to a natural disaster, 75F protocols makes sure it can recover any lost data at a secondary location where the data has been backed up. A disaster recovery plan should account for disasters that are both geographically dependent and those that occur regardless of physical location.

Disaster Recovery vs. Building Management Continuity

Disaster Recovery is part of Business Continuity. 75F Business Continuity processes are a proactive effort to mitigate risks and plan for a building's operation to continue regardless of the type of interruption. 75F Disaster Recovery plans focus on the infrastructure and systems that need to resume operation after an interruption occurs.

Disaster Recovery at 75F

Disaster recovery involves delving into a number of methodologies and technologies. However, every effective disaster recovery strategy involves the following four elements:

1. Identification of business-critical assets
2. Backup mechanisms put into place for all function-critical resources
3. Evaluation of risks for services and data that need to be backed up periodically
4. Recovery mechanism put into place and verified

Active geo-replication is designed as a business continuity solution that enables quick disaster recovery of individual databases in case of a regional disaster or a large-scale outage. With geo-replication, we initiate a geo-failover to a geo-secondary in a different Azure region. The following section intends to provide a complete strategy for backing up and restoring devices and infrastructure that store critical data. This also has taken care of recovery of the failed systems to a functional state with minimal loss of content and with the quickest turnaround time.

DISASTER RECOVERY: BACKUP & RESTORE PROCESSES ACROSS THE PLATFORM

CCU-Level Building Data

In case any device is to be replaced, building-level data can be restored to its prior state with backed-up data.

1. Backup Frequency – System-level data for each CCU is backed up every 24 hours or customized time period.
2. Backup Retention — Two days
3. Backup Location — Azure storage
4. Backup Restore — A new CCU can be commissioned with available files to its backed-up state

Backup of Cloud Infrastructure & Configuration

Currently the configurations for productions services and servers are documented or stored in repositories. 75F has multiple Azure App Service, each with its own set of configurations. If the app services are accidentally removed, the configuration should be readily available to restore the app back to service.

Backup & Restore of Code & Wiki

1. Source code for the applications developed are stored in repositories
2. Backup Frequency – 12 hours
3. Backup Retention – Three days
4. Backup Location: Another azure account and geographical regions.
5. Backup Restore: A new repository can be created with available files

Cloud Databases Storing 75F Platform Data

All cloud databases storing 75F platform data for each building are backed up as mentioned below:

1. MongoDB
   - MongoDB is configured with Auto backups by default
   - Backup Frequency – Hourly Snapshot: Six Hours
   - Backup Frequency – Weekly Snapshot: Saturday
   - Backup Retention – Hourly Backup: Seven days
   - Backup Retention – Weekly Backup: Four weeks
   - Monthly Snapshot is also supported and retained for one year
2. InfluxDB
   - Backup Frequency – Full Backup: 24 hours
   - Backup Frequency – Incremental Backup: Two hours
   - Backup Retention: Two days
   - Backup Location: Azure Blob Storage
3. MySQL
   - Geo redundant backup
   - Backup Frequency – Full Backup: 24 Hours
   - Backup Frequency – Incremental Backup: Not required (Needs verification)
   - Backup Retention: Seven days
   - Backup Location: Azure Internal Storage. Cannot access the backup files or export the backed-up file to a storage.

75F Platform layers are designed to be resilient in the face of any DoS and DDoS attacks. These attacks can be intended or unintended — for example, unintended attacks might include a rogue program that starts sending repeated requests.

Hardware Layer

Even the firmware on the hardware infrastructure is susceptible to several cyber-attacks due to the endpoint devices' restrictions in computation, storage, and communication capacity. This layer has been built with encrypted and proprietary data packets that ensure the mitigations steps are already put in place for any external data packets coming into the layer. Additionally, the right address of each Node/Stat allows for appropriate filtering when data packets are exchanged with these devices or even if external packets of data is flooded into the layer.

Software Cloud Layer

75F's cloud layer has its front end, platform services, and externally-exposed APIs for integrations, thus making the platform vulnerable to DoS/DDoS attacks. The two key considerations for mitigating large scale volumetric Dos or DDoS attacks are bandwidth (or transit) capacity and server capacity to absorb and mitigate attacks:

- **Transit capacity —** When architecting services and applications, our hosting provider offers ample redundant Internet connectivity that allows us to handle large volumes of traffic. Because the ultimate objective of a DDoS attack is to affect the availability of your resources/applications, 75F platform monitoring and alarming systems are equipped with the capability to take care of services availability by accurate infrastructure management and balancing options even during high volumes of traffic.

- **Server capacity —** Most DDoS attacks are volumetric attacks that use a lot of resources. 75F can quickly scale up or down on customer computation resources by auto-scaling the capability of our microservices and frontend hosted services. Additionally, we use load balancers to continually monitor and shift loads between resources to prevent overloading any one resource.
- **Proactive visibility and monitoring** — 75F's team proactively monitors for suspicious application traffic patterns and raises any needed alerts for further investigation.

The following section is an introduction to general security protocols at 75F. To read the entire procedure, visit the resources page on the 75F website.

INTRODUCTION TO 75F SECURITY PROTOCOLS

Security and privacy are major concerns when it comes to enterprise applications. The 75F platform incorporates protection in all stages of its lifecycle.

75F maintains transparency on all the data it collects from buildings and portfolios by creating, storing, using, and deleting data — and keeping all stakeholders informed on exactly what happens to this data and where it resides. From information protection, data management, and knowledge sharing to secure collaboration, 75F makes the most of your building information in a secure and user-friendly environment.

75F provides industry-leading security and privacy standards to its customers using a secure, enterprise-grade platform design for its whole product suite and cloud-based managed data. 75F alleviates security and privacy risks by isolating and protecting enterprise data sources and networks from client applications running on multi-platform devices and their networks.

The company's security protocol is a collection of security measures designed to protect cloud-based infrastructure, applications, and data. These measures ensure user and device authentication, data, and resource access control, and data privacy protection. This paper describes the security protocol and its features from three perspectives:

APPLICATIONS & COMMUNICATIONS — Communication across all layers aims to protect software application code & application data against any security threats

SECURITY OF PLATFORM SERVICES — Relies on a unified bundle of hardware and software that protect both infrastructure and software-defined hardware, storage, and network components along with the operating systems and applications that reside on those platforms

UNDERLYING CLOUD INFRASTRUCTURE — Procedures and technology that secure cloud computing environments against both external and internal security threats

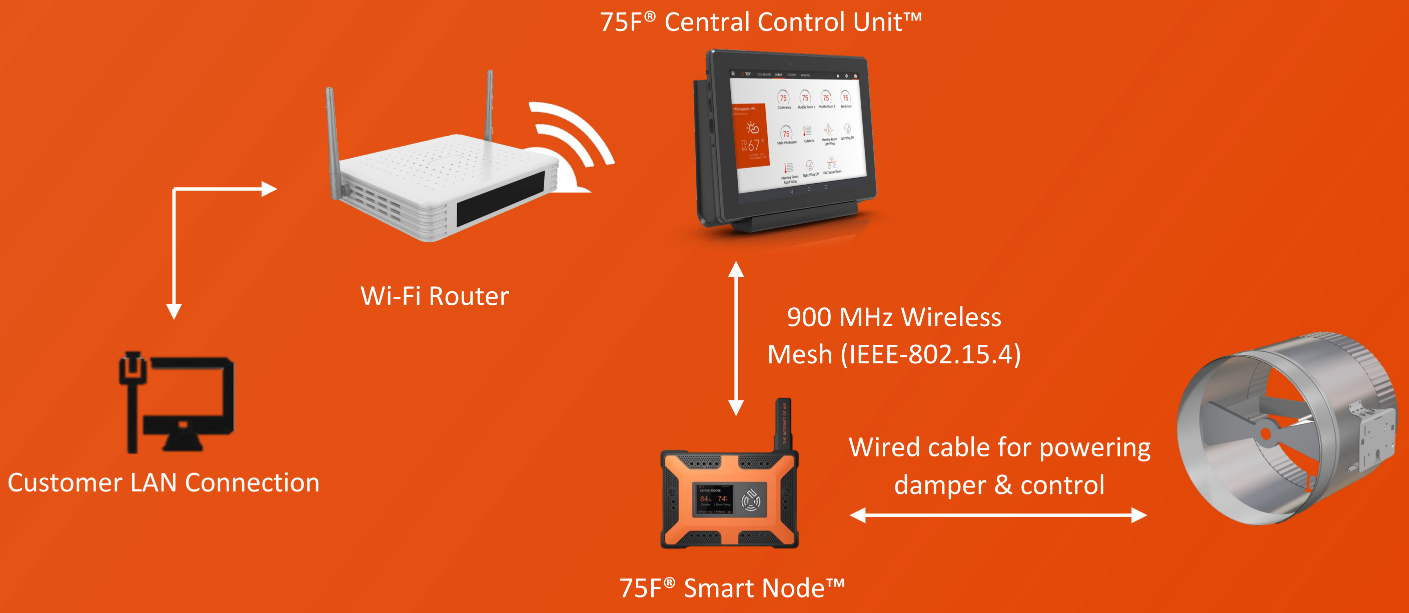

75F® Central Control Unit™

Wi-Fi Router

900 MHz Wireless Mesh (IEEE-802.15.4)

Customer LAN Connection

Wired cable for powering damper & control

75F® Smart Node™

SECURING THE SOFTWARE PLATFORM

The foundation for comprehensive platform and cloud security rests on four pillars: visibility and compliance; compute-based security; network protections; and identity security.

Securing the application and communication has a two-pronged approach:

1. User Management – Authentication and Authorization
2. Managing access to application functionality, API and data based on user management.

MICROSOFT AZURE AS 75F CLOUD PROVIDER

With Microsoft Azure acting as the 75F cloud provider, all the security aspects provided in Azure have been utilized and integrated into the 75F platform.

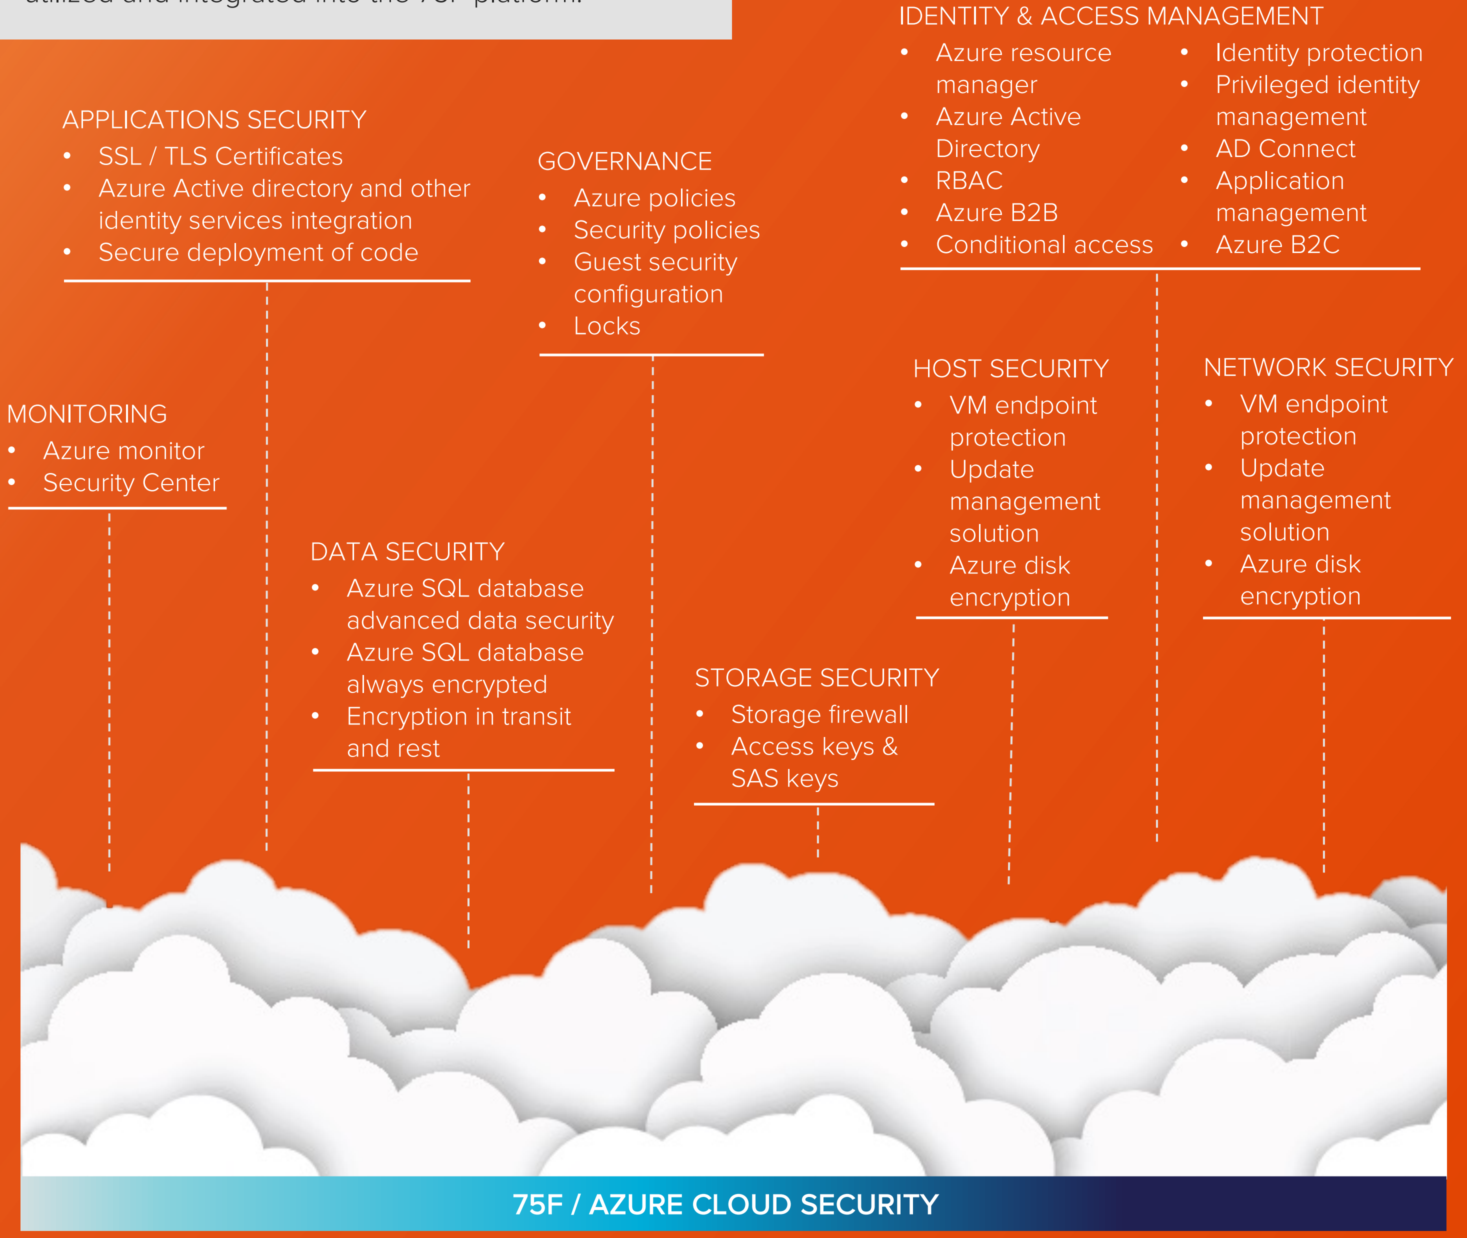

IDENTITY & ACCESS MANAGEMENT

- Azure resource manager
- Azure Active Directory
- RBAC
- Azure B2B
- Conditional access
- Identity protection
- Privileged identity management
- AD Connect
- Application management
- Azure B2C

APPLICATIONS SECURITY

- SSL / TLS Certificates
- Azure Active directory and other identity services integration
- Secure deployment of code

GOVERNANCE

- Azure policies
- Security policies
- Guest security configuration
- Locks

HOST SECURITY

- VM endpoint protection
- Update management solution
- Azure disk encryption

NETWORK SECURITY

- VM endpoint protection
- Update management solution
- Azure disk encryption

MONITORING

- Azure monitor
- Security Center

DATA SECURITY

- Azure SQL database advanced data security
- Azure SQL database always encrypted
- Encryption in transit and rest

STORAGE SECURITY

- Storage firewall
- Access keys & SAS keys

75F / AZURE CLOUD SECURITY