



# The Internet of Things–Ready Infrastructure

The world of smart devices talking to each other—and to us—is well underway and here to stay. To connect to the Internet of Things opportunity, it's key to design and build networking infrastructures that can handle massive amounts of new data.



## Introduction

The lines between our physical and digital worlds are blurring. Our daily experiences are—and will continue to be—impacted by the prevalence of Internet-connected devices and anytime, anywhere access to information. And there is no sign of things slowing down. Each year will see exponential growth in devices connected to the Internet. In fact, Gartner predicts there will be 25 billion connected “things” by 2020.<sup>1</sup>

These “things” will come in all shapes and sizes, from three-ton automobiles to clothing to under-the-skin blood sugar monitors and even our entire home. The world of smart devices talking to each other—and to us—is well underway. But reaping the business rewards will depend on the ability to design and build a networking infrastructure that successfully manages the flood of data that comes from this new Internet, the Internet of Things (IoT).

Just as LTE and technologies like Network Functions Virtualization (NFV), software-defined network (SDN), and VoLTE are transformative forces, there will be no escaping the effects of the IoT. How service providers choose to respond to this growing phenomenon and the explosion in connected devices, applications, and data that it brings will determine who benefits most in a market that IDC forecasts will grow to \$7.1 trillion in the next five years.<sup>2</sup>

Much like “bring your own device” transformed the workplace and enterprise mobility—not to mention subscriber expectations for faster, more reliable access to apps—the IoT will impact almost every aspect of our daily lives. Enhanced customer service and improved use of field assets have already been realized by early adapters across multiple industries. Moving forward, new business models will blossom and services will become more important than simple devices. With smartphones being the epicenter of people’s daily lives—and a crucial hub for IoT data traffic—finding a way to bring all this together is key for keeping pace in a constantly evolving landscape. And it must be done while ensuring network efficiency, interoperability, and security within your enterprise and optimal performance and security for the services you deliver to subscribers.



## WHITE PAPER

The Internet of Things–Ready Infrastructure

### More devices. More threats.

The IoT does not come without its challenges. Threats to data security, physical security, security of devices, regulations, privacy, encryption, authentication and a host of other issues all need to be addressed before the IoT can really become commonplace. These themes sound eerily similar to the ones surrounding the cloud only a couple of years ago. Now, consumer devices are the focus, and service providers will work to find new ways of driving greater operational efficiency and better management of infrastructure—for themselves and their customers. The challenge in harnessing this powerful force isn't limited to managing the sheer volume of data created. It's also making sense of that data to prioritize traffic and optimizing the application architecture itself.

Expert predictions tell us the IoT market size and the growing number of new devices will surpass anything we've seen. The wider implications for the underlying network infrastructure used to manage, monitor, and monetize these devices, though less obvious, will require considerable attention. The consequence of these devices and their supporting ecosystem failing could vary from a simple annoyance—device apps stuck on software updates—to something significantly worse, like a security breach targeting personally identifiable information. Vulnerabilities in devices will exist, devices will connect looking for updates, and patches will need to be pushed out. How will the infrastructure survive the onslaught?

The Internet of Things is not just about the things. It is really about the applications and services that enable them. F5 Synthesis™—the next generation model for supporting applications—can help apps (and the “things” they power) perform better with greater security and ensure that they are always available. A key component of the Synthesis model, F5@ Software-Defined Application Services™ (SDAS™) allows application delivery to evolve in alignment with new IoT innovations and device demands.

### What are these “things”?

The Internet of Things refers to the set of devices and systems that interconnect real-world sensors and actuators to the Internet.<sup>3</sup> This includes many different types of systems, such as:

- Mobile devices
- Smart meters and objects
- Wearable devices including clothing, health care implants, smart watches, and fitness devices
- Internet-connected automobiles
- Home automation systems, including thermostats, lighting, and home security
- Other measuring sensors for weather, traffic, ocean tides, road signals, and more



## WHITE PAPER


### The Internet of Things–Ready Infrastructure

These systems connect to the Internet or gateway in a variety of ways, such as long-range WiFi/Ethernet using IP protocols (TCP/UDP, including cellular); short-range Bluetooth low energy; short-range Near Field Communication; and other types of medium-range radio networks. Point-to-point radio links and serial lines are also used. Some devices/sensors connect directly to the Internet via an IP protocol and others with specific IoT protocols, such as Message Queue Telemetry Transport (MQTT), Constrained Application Protocol, and others. These may need specialized IoT networking hardware to “talk” to the data center.

MQTT is a “subscribe and publish” messaging protocol designed for lightweight machine-to-machine communications.<sup>4</sup> Originally developed by IBM, it is now an open standard. It needs a gateway or receiver (broker) to communicate. However, its primary purpose is to allow a device to send a very short message one hop to an MQ broker and to receive commands from that broker. Every message is published to a location, called a topic. Clients (the sensors) subscribe to various topics and when a message is published to the topic, the client/sensor gets it.

For example, imagine a simple network with three clients and a central broker.

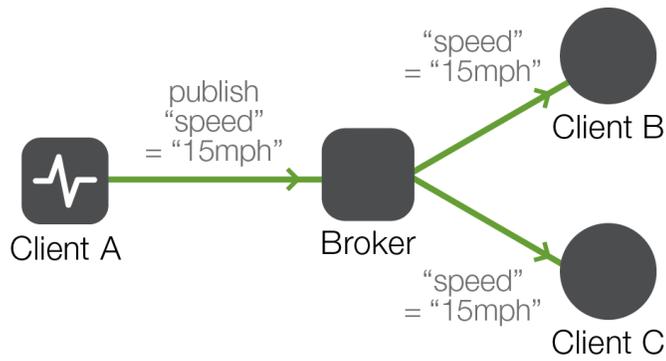
All three clients open TCP connections with the broker. Clients B and C subscribe to the topic speed.



At a later time, Client A publishes a value of 15mph for topic speed. The broker forwards the message to all subscribed clients.

## WHITE PAPER

### The Internet of Things–Ready Infrastructure



The publisher subscriber model allows MQTT clients to communicate one-to-one, one-to-many, and many-to-one.

The systems themselves typically fall into a few categories. The smallest devices have 8-bit embedded “system on a chip” controllers but no operating system. Others have a limited 32-bit architecture, like a home router, with or without a base OS. The most capable systems have either full 32-bit or 64-bit operating platforms, such as mobile phones. Customers could use their mobile phone to send the data, via the Internet, from the IoT device to the destined application.

Not only are we interacting with these devices, they are interacting with other machines to send specific information, which is called machine-to-machine (M2M) technology. The M2M fabric works in conjunction with the various systems that support wearables, home networks, and the widely deployed sensors that are part of them.

According to Gartner, the Internet of Things is not a single technology but a concept with embedded sensors driving the trend, real-time support and learning having a social impact and it allows businesses to make situational decisions based on the sensor’s information.<sup>5</sup> **No single architecture can address all the potential IoT device areas and the requirements of each. But a scalable architecture that can add or subtract resources to support a wide variety of scenarios can prepare service providers for the impact the IoT will have.**



## The IoT effect on applications

Many of today's traditional architectures will buckle under the increasing demand of all the connected devices. According to IDC, the rate at which applications double in the enterprise is once every four years.<sup>6</sup> This time frame is likely to be cut in half as more IoT devices need applications supporting them and service providers need to be ready for the increased demand.

As more applications are needed to run "things," traditional infrastructure concerns such as scale and reliability are becoming paramount. Additional challenges with identity and access, improving the end-user or subscriber experience, and the need for faster provisioning of services could overwhelm IT departments. A robust, scalable, and intelligent infrastructure will be necessary to handle the massive traffic growth.

The Domain Name System is the most likely method for connected devices to use to locate needed services, and it's potentially the means by which people will locate the devices themselves. There might be other schemas being planned, but those would require the adoption of a new technology naming standard, which would be costly, slow, and highly unlikely.

## Security at every layer

Clearly, security must also be present since the IoT has the potential to weave vulnerabilities throughout the system—the ubiquity of connected devices presents a gold mine for attackers. Outpacing attackers in our current threat landscape will require more resources in order to minimize risk. Service providers will need to continue to harden their own infrastructures and look to cloud services like DDoS mitigation to lessen the effects of attacks.


At the same time, the explosion of embedded devices will further drive mainstream IPv6 adoption. There are several advantages to IPv6 such as a large namespace, address self-configuration, and the potential to remove Network Address Translation problems. The networking infrastructure will require some planning to embrace this shift. Components such as routers, firewalls, and application delivery controllers will need to be IPv6–ready and capable of understanding the protocols and data that devices will use to communicate.

To ensure security, intelligent routing, and analytics, networking layers will need to be fluent in the language that devices use. Understanding these protocols within the network will allow traffic to be secured, prioritized, and routed accordingly. Recognizing and prioritizing these messages will enable better scalability and manageability for the onslaught of device traffic and data. Intelligence will also be needed to categorize what data needs attention (like a health monitor alert) and what doesn't (like temperature is good).

## WHITE PAPER

### The Internet of Things–Ready Infrastructure

According to TechTarget, ensuring high availability of the IoT services will rely on boosting traffic management and monitoring.<sup>7</sup> This will both mitigate business continuity risks, and prevent potential losses. From a project planning standpoint, organizations need to do capacity planning and watch the growth rate of the network so that the increased demand for the required bandwidth can be met.



**Figure 1:** Intelligent application delivery for the IoT

## The IoT-ready infrastructure

The F5 SDAS fabric provides a foundation for building elastic IoT application services. F5 ScaleN™ technology enables on-demand application and operational scalability at the platform layer. This means the fabric can be deployed on a combination of hardware, software, and virtual form factors—as well as beyond the data center boundary into cloud computing environments.

### Avoiding a single point of failure

The IoT applications will come in all shapes and sizes. But no matter the size, availability is paramount to support both subscribers and the business. The most basic high-availability architecture is the typical three-tier design. A pair of Application Delivery Controllers (ADCs) in the DMZ terminates the connection. They, in turn, intelligently distribute the client request to a pool (multiple) of IoT application servers, which then query the database servers for the appropriate content. Each tier has redundant servers so in the event of a server outage, the others take the load and the system stays available.<sup>8</sup>



## WHITE PAPER

### The Internet of Things–Ready Infrastructure

This is a tried-and-true design for most operations and provides resilient application availability, IoT or not, within a typical data center. But fault tolerance between two data centers is even more reliable than multiple servers in a single location, simply because that one data center is a single point of failure.

## Cloud: The IoT enabler

The cloud has become one of the primary enablers for IoT. Within the next five years, more than 90 percent of all IoT data will be hosted on service provider platforms. That’s because cloud computing reduces the complexity of supporting IoT “Data Blending.”<sup>9</sup>

In order to achieve or even maintain continuous IoT application availability and keep up with the pace of new IoT application rollouts, service providers must explore expanding their data center options. Having access to cloud resources provides service providers with the agility and flexibility to quickly provision IoT services. The cloud offers organizations a way to manage IoT services, rather than boxes along with just-in-time provisioning. Cloud enables IT as a Service, just as IoT is a service, along with the flexibility to scale when needed.

Integrating cloud-based IoT resources into the architecture requires only a couple of pieces: connectivity, along with awareness of how those resources are being used.

The connectivity between a data center and the cloud is generally referred to as a cloud bridge. The cloud bridge connects the two data center worlds securely and provides a network compatibility layer that bridges the two networks. This provides a transparency that allows resources in either environment to communicate without concern for the underlying network topology.

Once a connection is established and network bridging capabilities are in place, resources provisioned in the cloud can be non-disruptively added to the data center–hosted pools. From there, load is distributed per the F5 BIG-IP® platform’s configuration for the resource, such as an IoT application.

By integrating the enterprise data center with external clouds, the cloud becomes a secure extension of the enterprise’s IoT network. This enterprise-to-cloud network connection is encrypted and optimized for performance and bandwidth, thereby reducing the risks and lowering the effort involved in migrating IoT workloads to the cloud.



## WHITE PAPER

The Internet of Things–Ready Infrastructure

### Maintain seamless delivery

This hybrid infrastructure approach, including cloud resources, for IoT deployments not only allows service providers to distribute their IoT applications and services when it makes sense, but also provides global fault tolerance to the overall system. Depending on how the disaster recovery infrastructure is designed, this can be an active site, a hot standby, a leased hosting space, a cloud provider, or some other contained compute location. As soon as that IoT server, application, or even location starts to have trouble, a service provider can seamlessly maneuver around the issue and continue to deliver its services to the devices.

### Advantages for a range of industries

The various combinations of hybrid infrastructure types can be as diverse as the IoT situations that use them.

Service providers probably already have some level of hybrid, even if it is a mix of owned space plus SaaS. In most cases, the preference is to keep sensitive assets in-house, but this approach has started to shift, with increased migration of workloads to hybrid data centers. It all depends on the industry. For example, financial industries have different requirements than retail. Retail will certainly need a boost to its infrastructure as more customers will want to test IoT devices in-store. Whatever the case may be, service providers are in a unique situation as they will ultimately be touched by all this increased data traffic.

The service provider industry is well on its way to building out the IoT–ready infrastructures and services. One major service provider is deploying BIG-IP virtual editions to provide the ADC functionality needed for the scale and flexibility of its connected-car project. Virtualized solutions are required for NFV to enable the agility and elasticity necessary to support the IoT infrastructure demands.

### The IoT–ready platform

Because F5 application services share a common control plane—the F5 platform—the process of deploying and optimizing IoT application delivery services is simplified. With the elastic power of SDAS, organizations can rapidly provision IoT application services across the data center and into cloud computing environments, reducing the time and costs associated with deploying new applications and architectures.

The beauty of SDAS is that it can provide the global services to direct the IoT devices to the most appropriate data center or hybrid cloud depending on the request, context, and application health. Subscribers, employees, and the IoT devices themselves receive the most secure and fastest experience possible.



## WHITE PAPER

### The Internet of Things–Ready Infrastructure

The F5 high-performance services fabric supports traditional and emerging underlay networks. It can be deployed on top of a traditional IP and VLAN-based networks, works with SDN overlay networks using NVGRE or VXLAN (as well as a variety of less well-known overlay protocols), and integrates with SDN fabrics, such as those from Cisco/Insieme, Arista, and BigSwitch, among others.

## Hardware, software—or cloud

The services fabric model enables consolidation of services onto a common platform that can be deployed on hardware, on software, or in the cloud. This reduces operational overhead by standardizing management as well as enabling deployment processes that support continuous delivery efforts. By sharing service resources and leveraging fine-grained multi-tenancy, the cost of individual services is dramatically reduced, enabling all IoT applications—regardless of size—to take advantage of services that are beneficial to their security, reliability, and performance.

The F5 platform:

- Provides the network security to protect against inbound attacks
- Offloads SSL to improve the performance of the application servers
- Not only understands the application but also knows when it is having problems
- Ensures not only the best end-user experience but also quick and efficient data replication

F5 cloud solutions can automate and orchestrate the deployment of IoT application delivery services across both traditional and cloud infrastructures while also managing the dynamic redirection of workloads to the most suitable location. These application delivery services ensure predictable IoT experiences, a replicated security policy, and workload agility.

F5 BIG-IQ™ Cloud can federate management of BIG-IP solutions across both traditional and cloud infrastructures, helping organizations deploy and manage IoT delivery services in a fast, consistent, and repeatable manner, regardless of the underlying infrastructure. In addition, BIG-IQ Cloud integrates or interfaces with existing cloud orchestration engines such as VMware vCloud Orchestrator to streamline the overall process of deploying applications.

## Extend, scale—and secure

F5 cloud solutions offer a rapid Application Delivery Network provisioning solution, drastically reducing the lead times for expanding IoT delivery capabilities across data centers, whether they are private or public. As a result, service providers can efficiently:

- Extend data centers to the cloud to support IoT deployments
- Scale IoT applications beyond the data center when required
- Secure and accelerate IoT connections to the cloud



## WHITE PAPER

### The Internet of Things–Ready Infrastructure

For maintenance situations, service providers no longer need to manually redirect traffic by configuring applications. Instead, IoT applications are proactively redirected to an alternate data center prior to maintenance.

The BIG-IP platform is application and location agnostic, meaning the type of application or where the application lives really does not matter. As long as you tell the BIG-IP platform where to find the IoT application, the BIG-IP platform will deliver it.

Bringing it all together, F5 Synthesis gives cloud and application providers, as well as mobile network operators, the architectural framework necessary to ensure the performance, reliability, and security of IoT applications.

## Conclusion

Connected devices are here to stay—forcing us to move forward into this brave new world where almost everything generates data traffic. While there's much to consider, proactively addressing these challenges and adopting new approaches for enabling an IoT–ready network will help service providers chart a clearer course toward success.

An IoT–ready environment can enable service providers to begin taking advantage of this societal shift without a wholesale rip-and-replace of existing technology. It also provides the breathing room needed to ensure that the coming rush of connected devices does not cripple the network infrastructure. This process ensures benefits will be realized without compromising on the operational governance required to ensure availability and security of the IoT network, data, and application resources. It also means service providers can manage IoT services instead of boxes.

However an IoT–ready infrastructure is constructed, it is a transformational journey. It is not something that should be taken lightly or without a long-term strategy in place. When done properly, an F5-powered IoT–ready infrastructure can bring significant benefits to service providers, their employees, and their subscribers.

---

F5 Silverline™ DDoS Protection is a service delivered via the Silverline DDoS Protection cloud-based platform that provides continuous DDoS protection, detection, and mitigation to stop even the largest volumetric DDoS attacks from reaching IoT networks.

---

---

<sup>1</sup> <http://www.gartner.com/newsroom/id/2905717>

<sup>2</sup> <http://www.forbes.com/sites/gilpress/2014/08/22/internet-of-things-by-the-numbers-market-estimates-and-forecasts/>

<sup>3</sup> <http://wso2.com/whitepapers/a-reference-architecture-for-the-internet-of-things/>

<sup>4</sup> [http://eclipse.org/community/eclipse\\_newsletter/2014/february/article2.php](http://eclipse.org/community/eclipse_newsletter/2014/february/article2.php)

<sup>5</sup> <http://www.gartner.com/webinar/2862818>

<sup>6</sup> IDC Directions, “Battle for the Future of the Datacenter: The Role of Disaggregated Systems,” Mar 2014



## WHITE PAPER

### The Internet of Things–Ready Infrastructure

<sup>7</sup> <http://searchsecurity.techtarget.com/tip/Internet-of-Things-IOT-Seven-enterprise-risks-to-consider>

<sup>8</sup> [http://support.rightscale.com/12-Guides/Designers\\_Guide/Cloud\\_Solution\\_Architectures/Designing\\_and\\_Deploying\\_High-Availability\\_Websites](http://support.rightscale.com/12-Guides/Designers_Guide/Cloud_Solution_Architectures/Designing_and_Deploying_High-Availability_Websites)

<sup>9</sup> IDC FutureScape Internet of Things:  
<http://www.machinetomachinemagazine.com/2014/12/04/idc-report-worldwide-iot-predictions-for-2015/>

#### F5 Networks, Inc.

401 Elliott Avenue West, Seattle, WA 98119  
888-882-4447 [www.f5.com](http://www.f5.com)

Americas  
[info@f5.com](mailto:info@f5.com)

Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

Japan  
[f5j-info@f5.com](mailto:f5j-info@f5.com)