

Uptime Institute®

Differentiating Your Datacentre in the Networked Future

John Duffin

Managing Director, South Asia. Uptime Institute

July 2017

© 2017 Uptime Institute, LLC

The Global Datacentre Authority







Uptime Institute® network



© 2017 Uptime Institute, LLC

The “Next Generation” Datacentre

- The Datacentre in 2020 and Beyond - (https://journal.uptimeinstitute.com/data-center-2020-beyond)
- The One “Datacentre” Model No Longer Exists
- Multiple Models Now Driven by
  - New Technology
  - Industrialisation
  - Standardisation
  - Automation
- Increasingly Workload Will Drive Design
- Resilience and Reliability Still Key

© 2017 Uptime Institute, LLC

"Next Generation" Technologies

- 10 Disruptive Technologies… Defined in 2013

TECHNOLOGY	HOW BIG	HOW FAST	HOW LIKELY	DISRUPTIVE RATING
Low-Power Servers	3.41	4.00	3.19	3.53
On-Site Clean Power Generation	3.13	2.31	2.42	2.62
Advanced DCIM	3.47	4.25	3.43	3.72
Cloud-Level Resiliency	3.95	4.06	3.36	3.79
Silicon Photonics	3.84	3.64	3.18	3.55
Chiller-Free Datacenters	3.56	3.54	3.35	3.48
Power-Proportional Computing	2.87	3.85	3.49	3.40
Flash Storage	3.66	4.59	4.09	4.11
Prefabricated Modular Datacenters	3.45	4.06	3.34	3.62
Memristors	3.88	3.31	2.83	3.34

© 2017 Uptime Institute, LLC

"Next Generation" Datacentres

- Types
  - Hyperscale – Owner Occupier and Colocation
  - Colocation – Traditional, Modular, Containers
  - Micromodular – Modular, Container
  - In Building – Rack, Multiple Racks
- Hybrid Infrastructures
  - Combinations of the above
- "Usage Class"
  - Core
  - Near Edge
  - Edge

© 2017 Uptime Institute, LLC

The Edge

- Data
  - Where data is initially generated – including by sensors and other devices (‘things’) – consumed and stored.
- Connections
  - Carrier-neutral or carrier-specific connectivity services with cross connects, cloud exchanges and/or direct connects (including dark fibre to other datacentres).
- Telecommunication Gateways
  - Entry or access points to local networks/WANs, including fibre and cellular, as well as to cloud computing and other IT service environments on networks.

© 2017 Uptime Institute, LLC

Where Does the Edge Live?

Figure 3: Key Characteristics of Micromodular Datacenters

Source: 451 Research, 2017

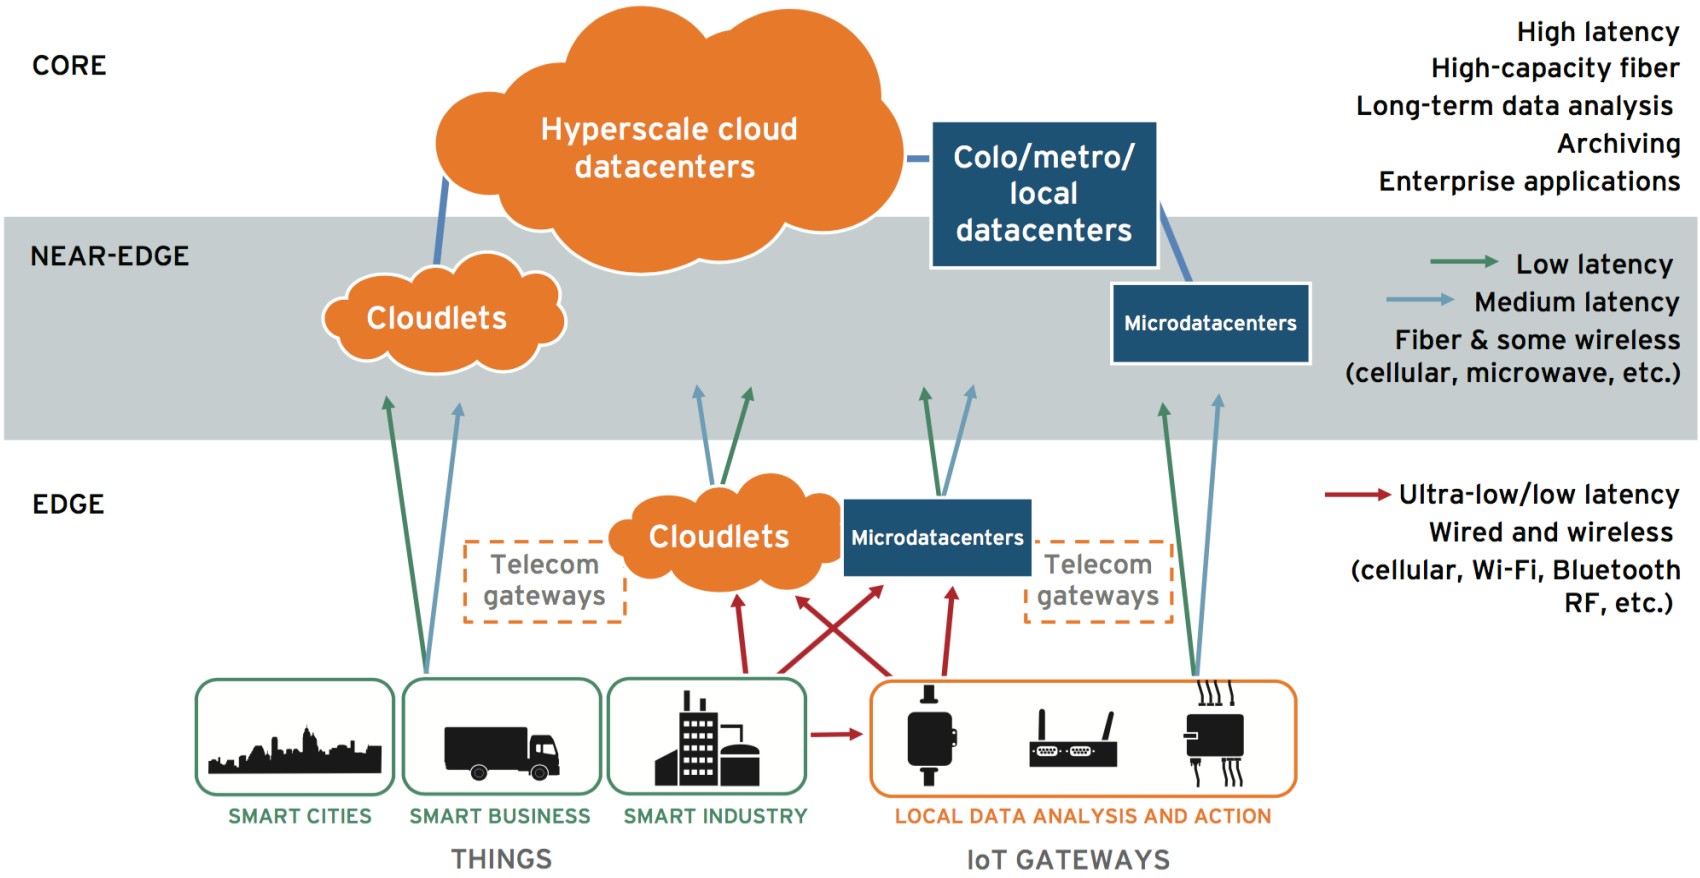
KEY CHARACTERISTICS	
Location	Near end user or last-mile communications.
Power Consumption	20kW-200 kW
Resiliency	Some redundancy and use of backup; unlikely to have concurrent maintainability. Some may have software failover.
Power Source	Grid, fuel cell, battery (Lead, Li-On or some other form), renewable energy source.
Power Conditioning	UPS and/or battery.
Power Backup	Usually (fuel cell, generator, battery).
Cooling	Indirect free air (fans plus evaporative) plus minimal mechanical (most likely); closed loop liquid (rarely).
IT Power	Usually <20-40 servers.
IT Configuration	Racks, rack scale or custom.
Security	Physical hard shell plus logical.
Fire Suppression	Yes.
Local Storage	Flash/SSD (usually).
Connectivity	One or two fiber and/or high-speed wireless (may be multiple).
Management	Remote – IT management and datacenter infrastructure management (DCIM).

© 2017 Uptime Institute, LLC

The Datacentre, Unfolded

Figure 2: Datacenters for IoT and Edge Computing

Source: 451 Research, 2017

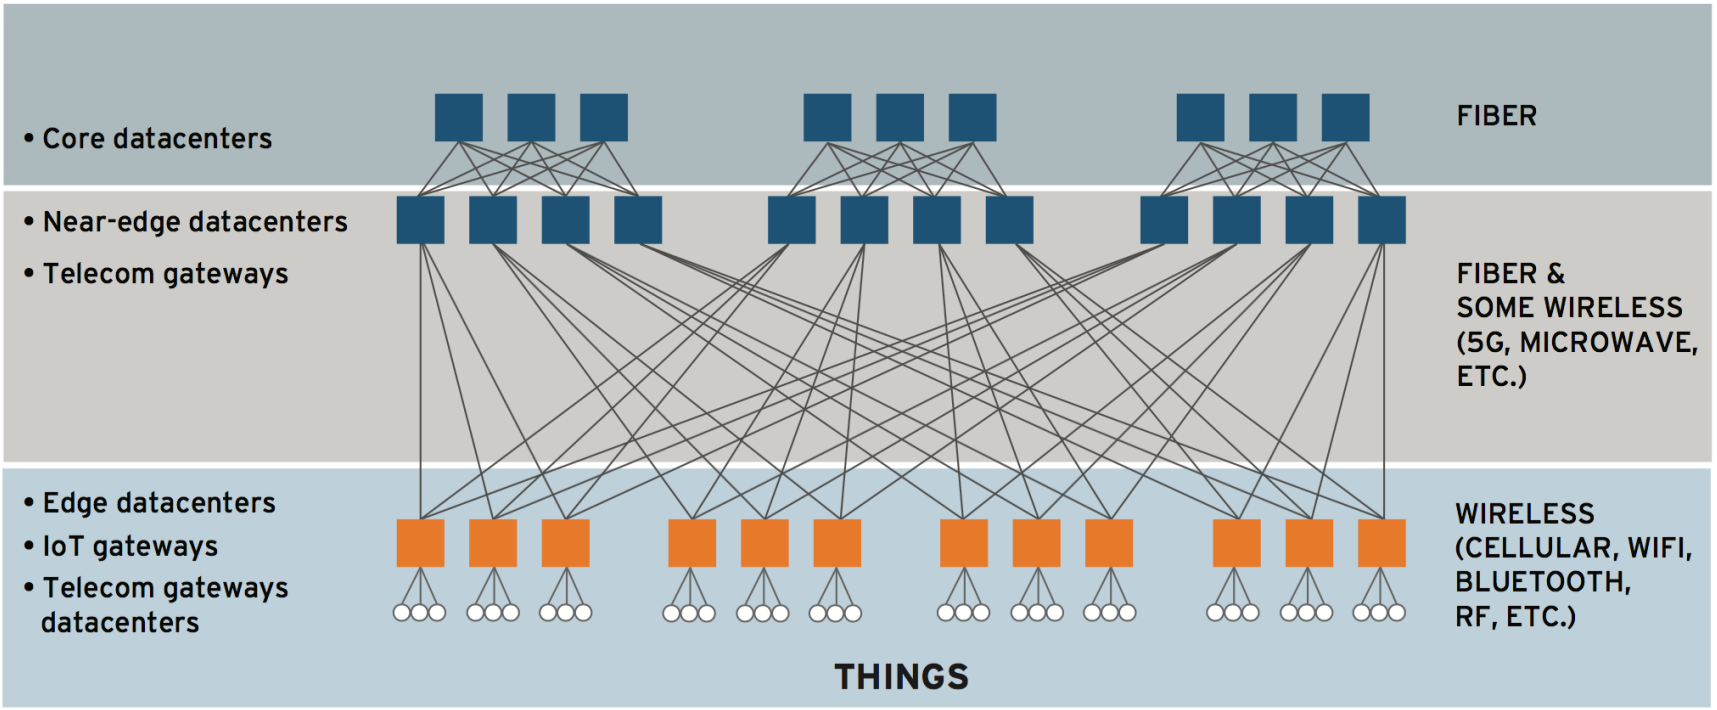


© 2017 Uptime Institute, LLC

Networks of the Future...?

Figure 1: Potential Networking and Datacenter Topologies at the Edge

Source: 451 Research, 2017



© 2017 Uptime Institute, LLC

Edge Use Cases…

Figure 4: Demand Characteristics for Major Edge Datacenter Use Cases

Source: 451 Research, 2017

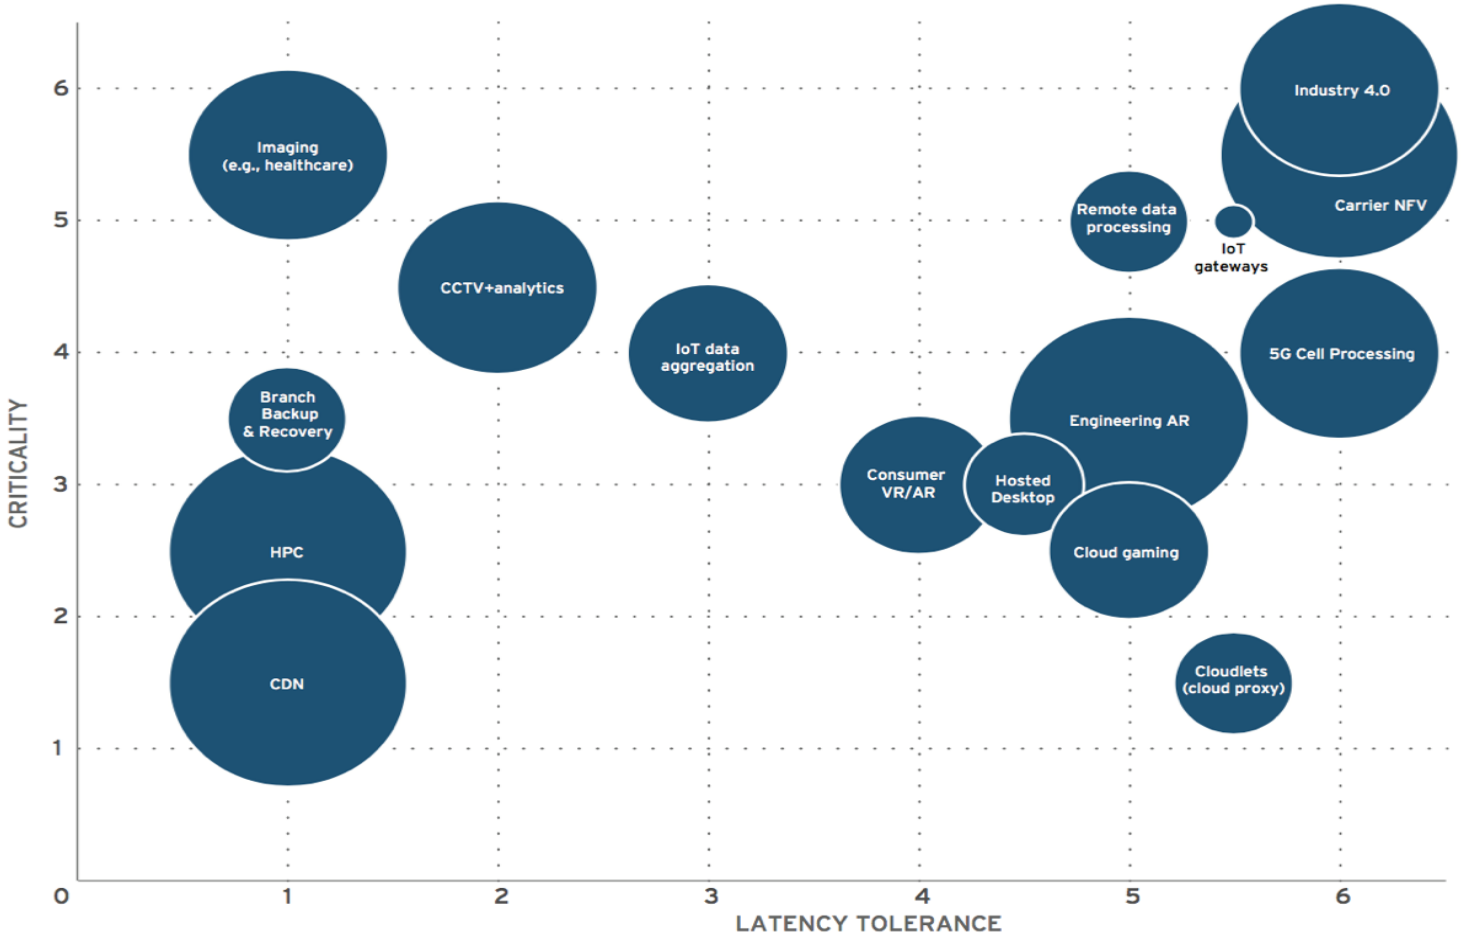
EDGE DATACENTER USE CASE	PROBABLE DEMAND CHARACTERISTICS	OPPORTUNITY FOR SUPPLIERS AND SERVICE PROVIDERS
CDNs	High and established demand that will continue to grow.	Existing growth opportunity.
Local Processing and Storage	Largely established demand that will continue to grow.	Upgrade opportunity: from server closets/rooms to microdatacenters/ near edge.
Off-Cloud Processing and Storage	Relatively low and established demand that will remain niche.	Limited, existing opportunity.
IoT	Demand is limited today, but rapid growth is expected.	New high-growth opportunity.
Next-Generation Networks and Distributed Cloud	Virtually no demand today but high growth is expected. The speed of deployment/adoption is unclear.	Future growth opportunity; could be tied to the buildout of local network capacity.

© 2017 Uptime Institute, LLC

Workload Drives Design?

Figure 6: Classifying Select Workloads for Their Edge Affinity

Source: 451 Research,2017



© 2017 Uptime Institute, LLC

Latency Drives Resilience?

SCORE VALUE	LATENCY	VOLUME / DAY / SITE	CRITICALITY (AVAILABILITY OF DATA OR SITE)
6	sub-1ms	100TB+ (10Gbps+)	0 downtime/extreme business risk
5	1-10ms	10-100TB (1-10Gbps)	Few seconds/ high risk
4	10-100ms	1TB-10TB (0.1-1Gbps)	Few minutes / moderate risk
3	100-1000ms	100-1000GB (10-100Mbps)	Few hours / low risk
2	1-2s	10-100GB (1-10Mbps)	Few days / little risk
1	2s+	10GB (1Mbps)	Irrelevant / no risk

© 2017 Uptime Institute, LLC



Uptime Institute®

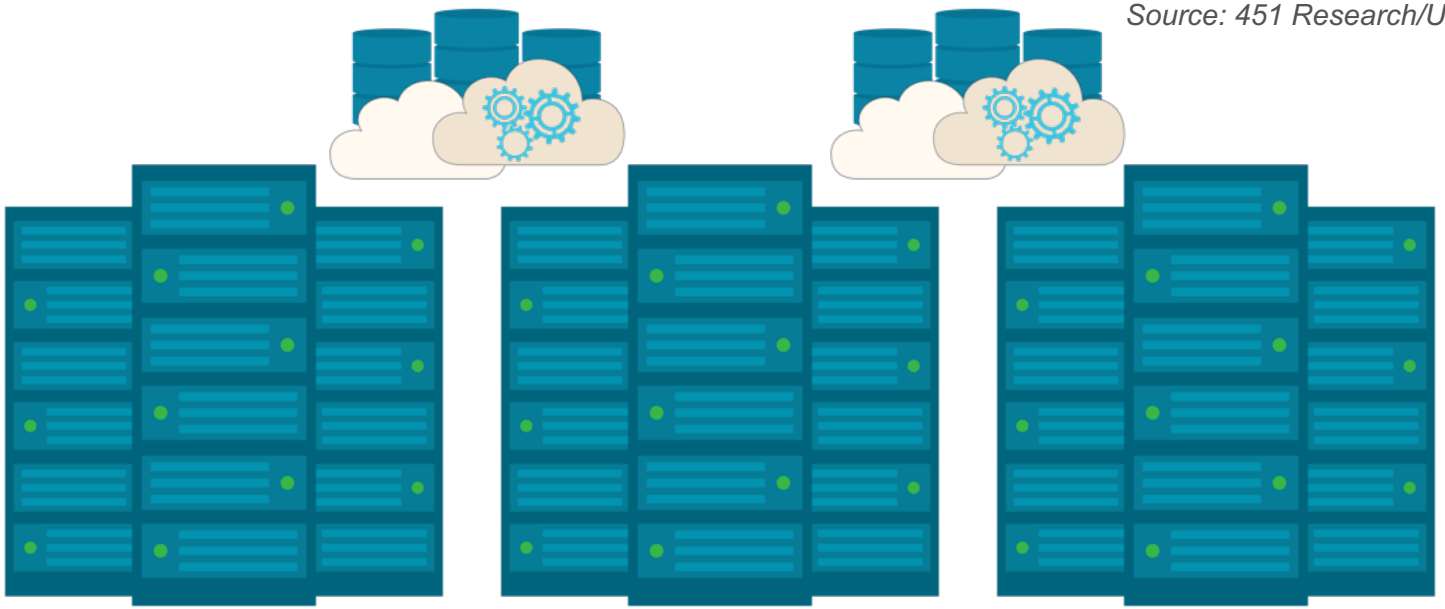
Resiliency in the Age of Cloud

© 2017 Uptime Institute, LLC

Definition of Resiliency

Resiliency - the **extent to which a system**, digital infrastructure, or application architecture is able to **maintain its intended service levels**, with **minimal or no impact on the users or business objectives,** in spite of planned and unplanned disruptions. It also describes the ability of a system, infrastructure or application to **recover full business operations after a disruption or disaster has occurred.**

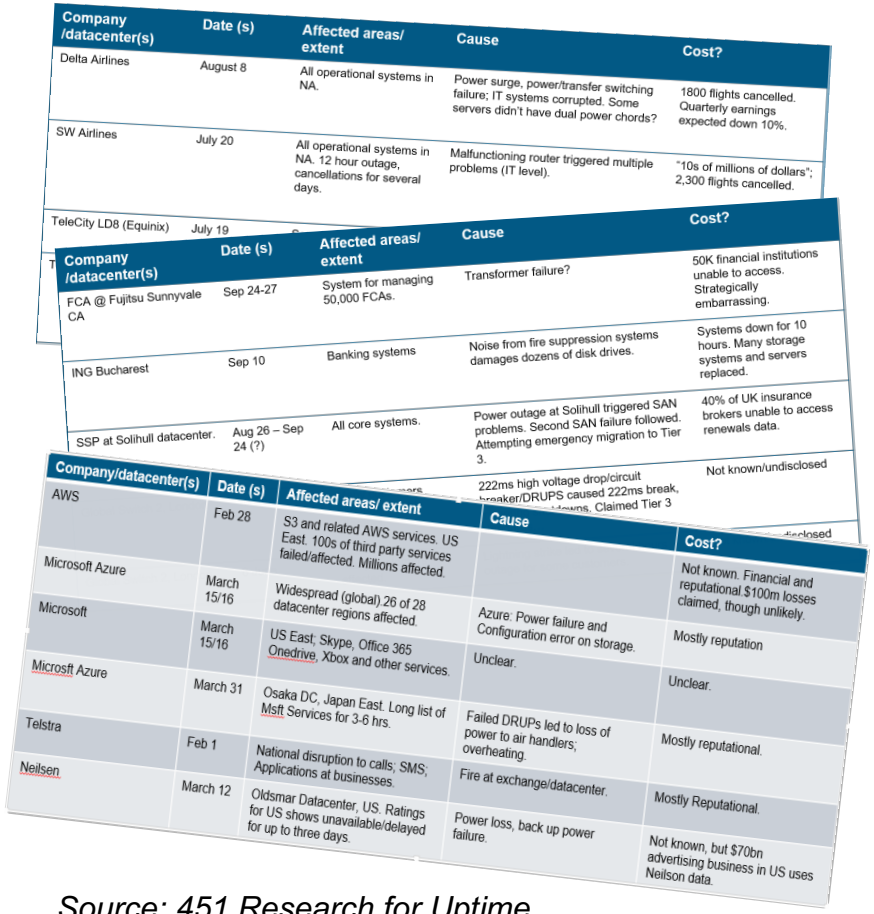
Source: 451 Research/Uptime Institute



© 2017 Uptime Institute, LLC

Big Outages - More Causes, More Effects

- 2016 and 2017 Outages Hit Revenues, Service and Reputation.
- Causes Stretch Across IT, Networks and Facilities.
- Complexity and Interdependencies Complicate Diagnosis and Recovery.



Company /datacenter(s)	Date (s)	Affected areas/ extent	Cause	Cost?
Delta Airlines	August 8	All operational systems in NA.	Power surge, power/transfer switching failure; IT systems corrupted. Some servers didn't have dual power chords?	1800 flights cancelled. Quarterly earnings expected down 10%.
SW Airlines	July 20	All operational systems in NA. 12 hour outage, cancellations for several days.	Malfunctioning router triggered multiple problems (IT level).	"10s of millions of dollars"; 2,300 flights cancelled.
TeleCity LD8 (Equinix)	July 19			
FCA @ Fujitsu Sunnyvale CA	Sep 24-27	System for managing 50,000 FCAs.	Transformer failure?	50K financial institutions unable to access. Strategically embarrassing.
ING Bucharest	Sep 10	Banking systems	Noise from fire suppression systems damages dozens of disk drives.	Systems down for 10 hours. Many storage systems and servers replaced.
SSP at Solihull datacenter.	Aug 26 – Sep 24 (?)	All core systems.	Power outage at Solihull triggered SAN problems. Second SAN failure followed. Attempting emergency migration to Tier 3.	40% of UK insurance brokers unable to access renewals data.
			222ms high voltage drop/circuit breaker/DRUPS caused 222ms break, Claimed Tier 3	Not known/undisclosed
Company/datacenter(s)	Date (s)	Affected areas/ extent	Cause	Cost?
AWS	Feb 28	S3 and related AWS services. US East. 100s of third party services failed/affected. Millions affected.		Not known. Financial and reputational $100m losses claimed, though unlikely.
Microsoft Azure	March 15/16	Widespread (global) 26 of 28 datacenter regions affected.	Azure: Power failure and Configuration error on storage.	Mostly reputation
Microsoft	March 15/16	US East: Skype, Office 365 Onedrive, Xbox and other services.	Unclear.	Unclear.
Microsft Azure	March 31	Osaka DC, Japan East. Long list of Msft Services for 3-6 hrs.	Failed DRUPs led to loss of power to air handlers; overheating.	Mostly reputational.
Telstra	Feb 1	National disruption to calls; SMS; Applications at businesses.	Fire at exchange/datacenter.	Mostly Reputational.
Neilsen	March 12	Oldsmar Datacenter, US. Ratings for US shows unavailable/delayed for up to three days.	Power loss, back up power failure.	Not known, but $70bn advertising business in US uses Neilson data.

Source: 451 Research for Uptime Institute Network members

© 2017 Uptime Institute, LLC

Resiliency Trends

- More Sharing
- More Replication
- More in Software
- More Distributed
- More is Active



"…It is all about data: speed of access, availability and protection, both in terms of security and for regulatory requirements…"

Source : 451 Research and Uptime Institute

© 2017 Uptime Institute, LLC

Architectural Swing…

Scale up

- Transactional
- High Integrity
- Interdependent stack
- Single/mirrored sites
- DR/Back up
- Resilient facilities





© 2017 Uptime Institute, LLC

Architectural Swing…

Scale up

- Transactional
- High Integrity
- Interdependent stack
- Single/mirrored sites
- DR/Back up
- Resilient facilities



Scale out

- Distributed components
- Horizontally layered
- Highly virtualized
- Highly replicated
- Less transactional
- Available but integrity issues
- Less redundant facilities



Datacentre 1 Datacentre 2 Datacentre 3 Datacentre 4

© 2017 Uptime Institute, LLC

Datacentre Resiliency Types

•**Traditional Single-Site Availability**. This is the traditional setup, with high levels of redundancy at the infrastructure level, including facilities and basic IT.

Multi-site Resiliency

•**Linked-Site Resiliency.** This describes two or more lower-tier datacentres connected within a campus, region or zone using a dedicated network to achieve a higher level of availability than any individual site.

•**Distributed-Site Resiliency** – This is the term used to describe two or more independent sites using a shared internet/VPN network to provide resiliency through multiple asynchronously connected instances.

•**Cloud-Based Resiliency:** This is the term used to describe resiliency provided by distributing, virtualized applications, instances or containers, across multiple datacenters, using middleware, orchestration and distributed databases.

Source : 451 Research and Uptime Institute

© 2017 Uptime Institute, LLC

Resiliency Architecture Capabilities

	Support for rapid recovery	Supports continuous service through Planned Maintenance across all sites	Supports continuous service through Unplanned Major component loss across all sites	Supports continuous service through catastrophic single site loss	Data Integrity
Single Site (Tier III)	No	Yes	No	No	ACID
Single Site (Tier IV)	No	Yes	Yes	No	ACID
Linked Site – Asynchronous	Yes	If one site is Tier III or IV	If one site is Tier IV	Yes	BASE
Linked Site (asynchronous)	Yes	Yes	Yes	Yes	ACID
Distributed Site (asynchronous)	Yes	If one Site is Tier III or IV	If One Site is Tier IV	Yes	BASE
Cloud Based Resiliency (>2 Active DCs)	Yes	Yes	Yes	Yes	ACID

© 2017 Uptime Institute, LLC

Data Integrity….

ACID	BASE
***A**tomicity, **C**onsistency, **I**solation, **D**urability*	***B**asically **A**vailable, **S**oft State, **E**ventual consistency*
Strong consistency	Weak consistency (stale data?)
Isolation	Availability above consistency
Focus on “commit”	Best effort
Available/Consistent	Available/Partition tolerant
Transaction oriented	Data read/programmer managed
Robust Database/simpler code	Simpler database/harder code

© 2017 Uptime Institute, LLC

Architecture vs Data Integrity

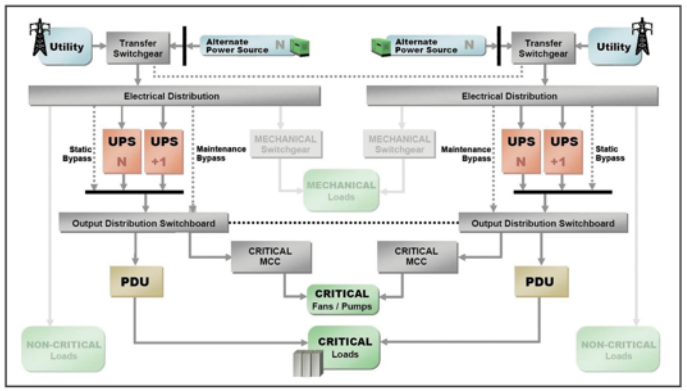
	Support for rapid recovery	Supports continuous service through Planned Maintenance across all sites	Supports continuous service through Unplanned Major component loss across all sites	Supports continuous service through catastrophic single site loss	Data Integrity
Single Site (Tier III)	No	Yes	No	No	**ACID**
Single Site (Tier IV)	No	Yes	Yes	No	**ACID**
Linked Site (asynchronous)	Yes	If one site is Tier III or IV	If one site is Tier IV	Yes	BASE
Linked Site – Synchronous	Yes	Yes	Yes	Yes	**ACID**
Distributed Site (asynchronous)	Yes	If one Site is Tier III or IV	If One Site is Tier IV	Yes	BASE
Cloud Based Resiliency (>2 Active DCs)	Yes	Yes	Yes	Yes	**ACID**

© 2017 Uptime Institute, LLC

Traditional Single-Site Resiliency

This is the traditional setup, with high levels of redundancy at the infrastructure level, including facilities and basic IT.

At the IT level, resilience is further assured by internal replication and clustering.
Data/applications/configurations backed up off-site.



Pros:

- Extremely reliable/good performance for single sites.
- Under management control
- Hard wired designs reduce failure/unpredictability
- Prevents most serious and common failures (power)

Cons:

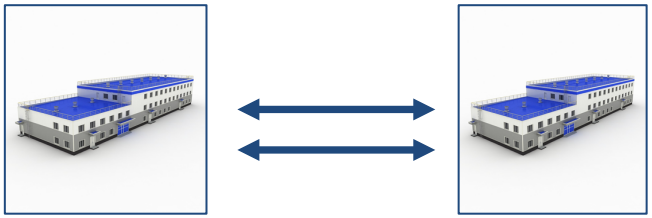
- Only partial solution for distributed loads
- Vulnerable to site specific issues
- High costs linear with scale

© 2017 Uptime Institute, LLC

Linked-Site Resiliency

• Two or more datacenters connected within a campus, region or zone using a dedicated network to achieve a higher level of availability than any individual site.

*At the IT level: C*an be used to support either synchronous (fault tolerant automated failover to the second site) or asynchronous replication (second site picks up the load).



Pros:

- Provides added resiliency to single sites
- Under management control
- Hard wired designs reduce failure/unpredictability
- Prevents most serious and common failures (power)
- Can be engineered after single site build

Cons:

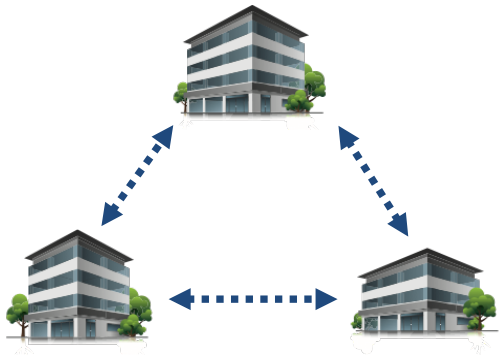
- Only partial solution for distributed loads
- Vulnerable to zonal incidents – requires out of region DR
- High costs linear with scale

© 2017 Uptime Institute, LLC

Distributed-Site Resiliency

Two or more independent sites using a shared internet/VPN network to provide resiliency through multiple asynchronously connected instances.

At the IT level: Underpins most disaster recovery services, especially the modern cloud iteration; Can support BASE architecture.



Pros:

- Can support extremely high availability/maintainability
- Can reduce/eliminate vulnerability to local/regional issues
- Can eliminate need for Disaster Recovery
- Can enable reduced investment in physical redundancy
- Very scalable – suited to Cloud native/scale out IT

Cons:

- Introduces IT complexity, expense
- Requires compromises – integrity v availability
- Requires scale or close collaboration

© 2017 Uptime Institute, LLC

Distributed Site Resiliency – How it Works

No of DCs	Server configuration at (sites)	Total # of servers	IT Capacity	Example Availability	Servers available if:		
					Site A goes off line	Site A + 1 off line	Site A + 2 off line
1	8	8	100%	99%	0	0	0
2	8 + 8	16	200%	99.99	8	0	0
3	4 + 4 + 4	12	150%	99.9999	8	4	0
4	4 + 4 + 2 + 2	12	150%	99.999999	8	4	2
5	2 + 2 + 2 + 2 + 2	10	125%	>99.99999	8	6	4

Availability and utilization can increase as datacenters are added, increasing efficiency.

© 2017 Uptime Institute, LLC

Cloud-Based Resiliency

Resiliency is provided by distributing, virtualized applications, instances or containers, across multiple datacenters, using middleware, orchestration and distributed databases.

Cloud based Resiliency moves resiliency up *to the IT level.* Any facility resilience achieved through redundancy provides added security, but may not prove essential.



Pros:

- Extremely high availability and performance
- Management control over wide area/globally
- May solve integrity v availability trade offs
- Eliminates need for DR
- Application programming simplicity

Cons:

- Requires control/ownership of datacenters/networks
- Complex, leading edge IT, some proprietary
- Requires scale - unsuited for enterprise use

© 2017 Uptime Institute, LLC

Resiliency Directions…

- Failure is no Longer On/Off
- More Requirement to Trust Others
- More Complexity
- Mixed Availability and Hybrid Strategies Emerging

In order to assess the business value of distributed resiliency, a multi-disciplined business and technical analysis is needed.

© 2017 Uptime Institute, LLC



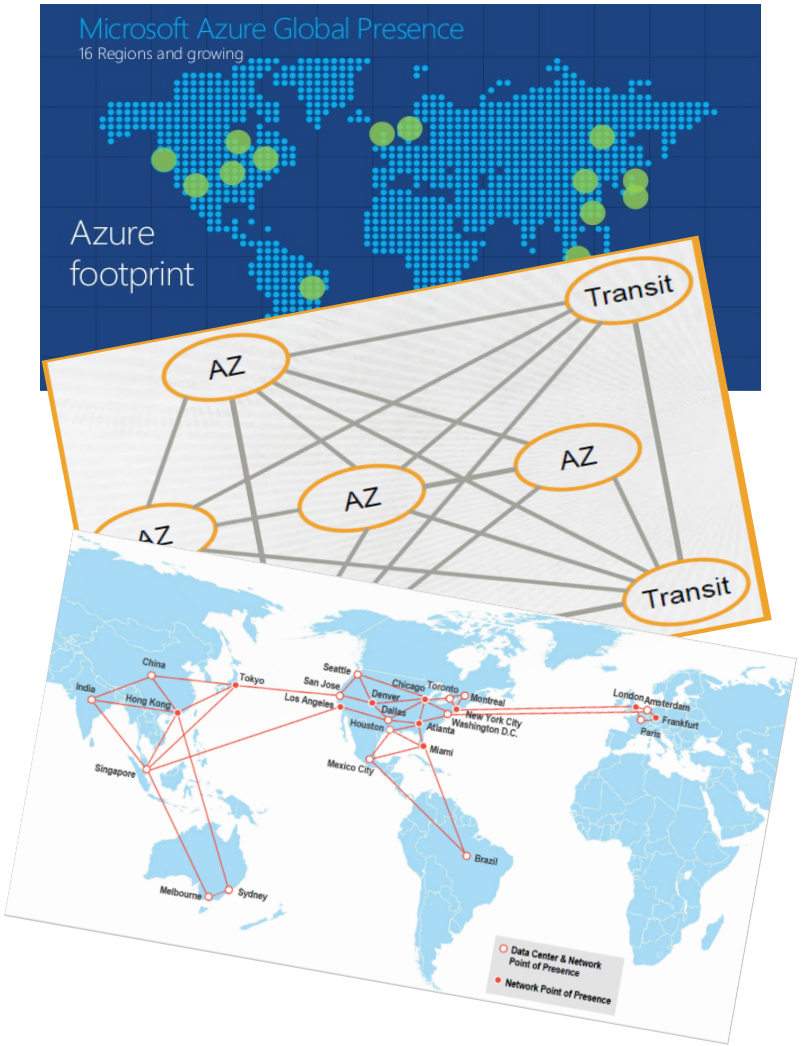
Uptime Institute®

Multi-Datacentre Infrastructure - Resiliency Assurance

© 2017 Uptime Institute, LLC

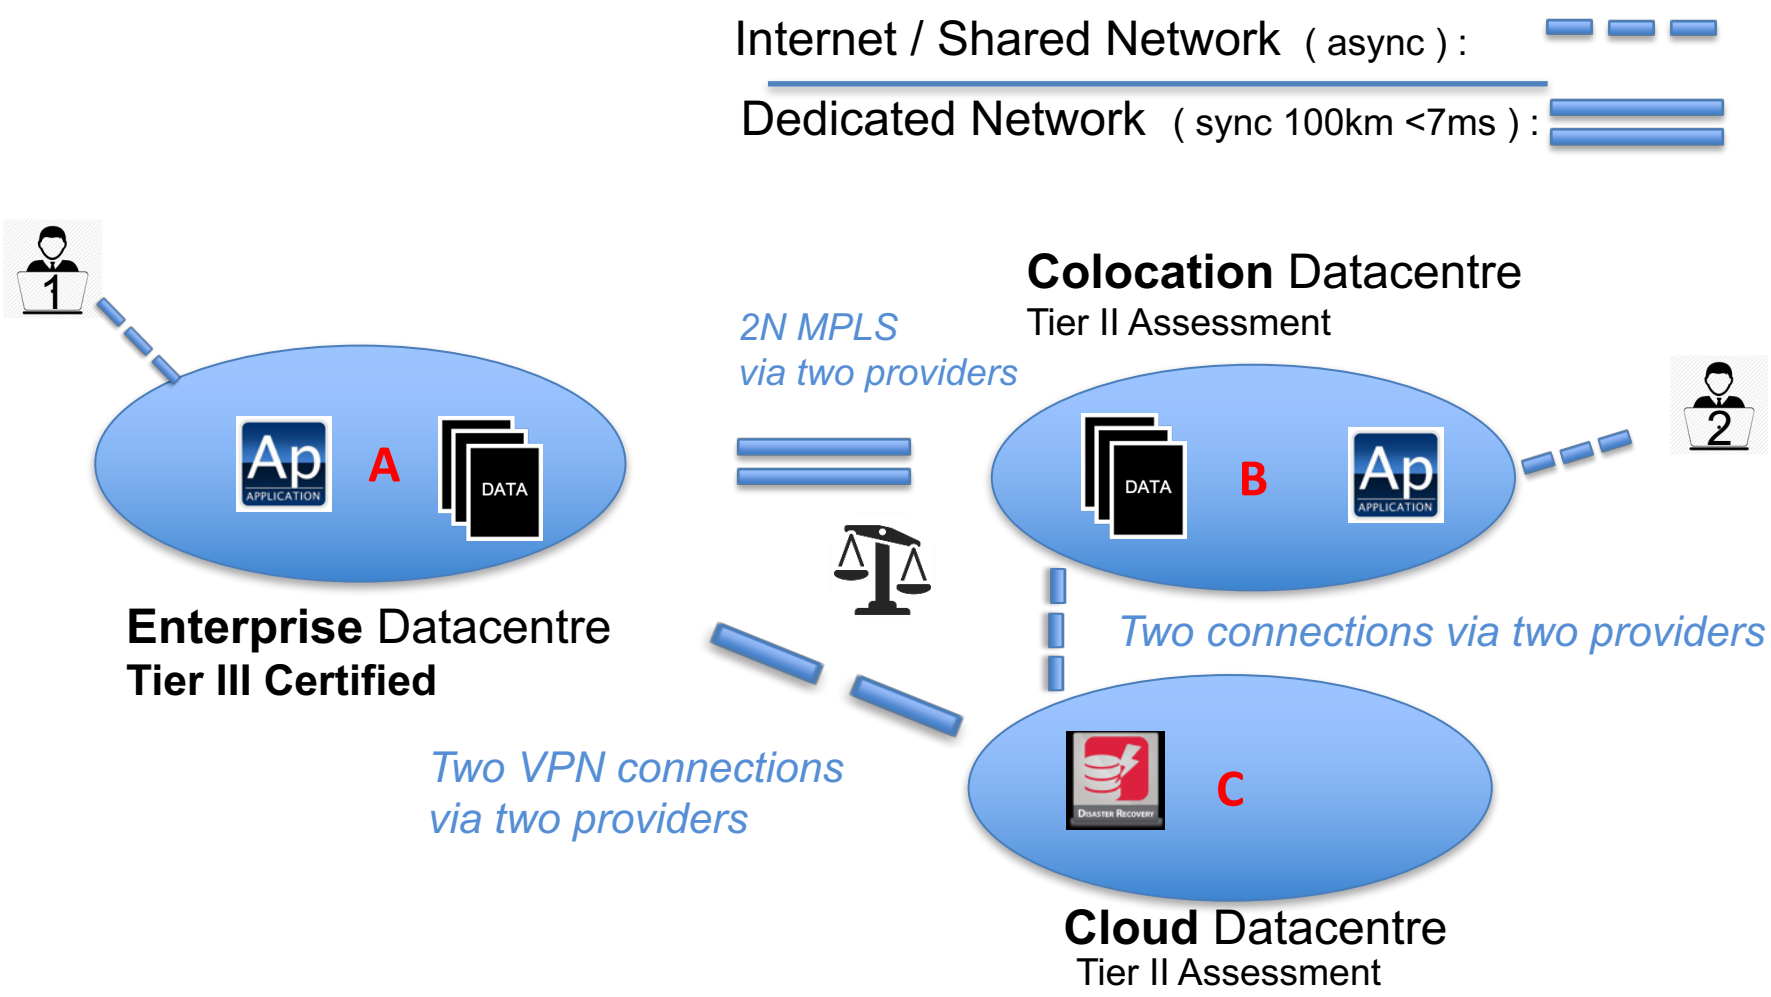
Benefits from Distributed Resiliency

- Can Offer Very High Availability
- Can Offer Very High Efficiency (Utilization of IT)
- Is (quite) Designed to Scale
- Can Deal Well With Single Site, Zonal and Regional Incidents
- Can Allow Rapid Deployment with Light Infrastructure



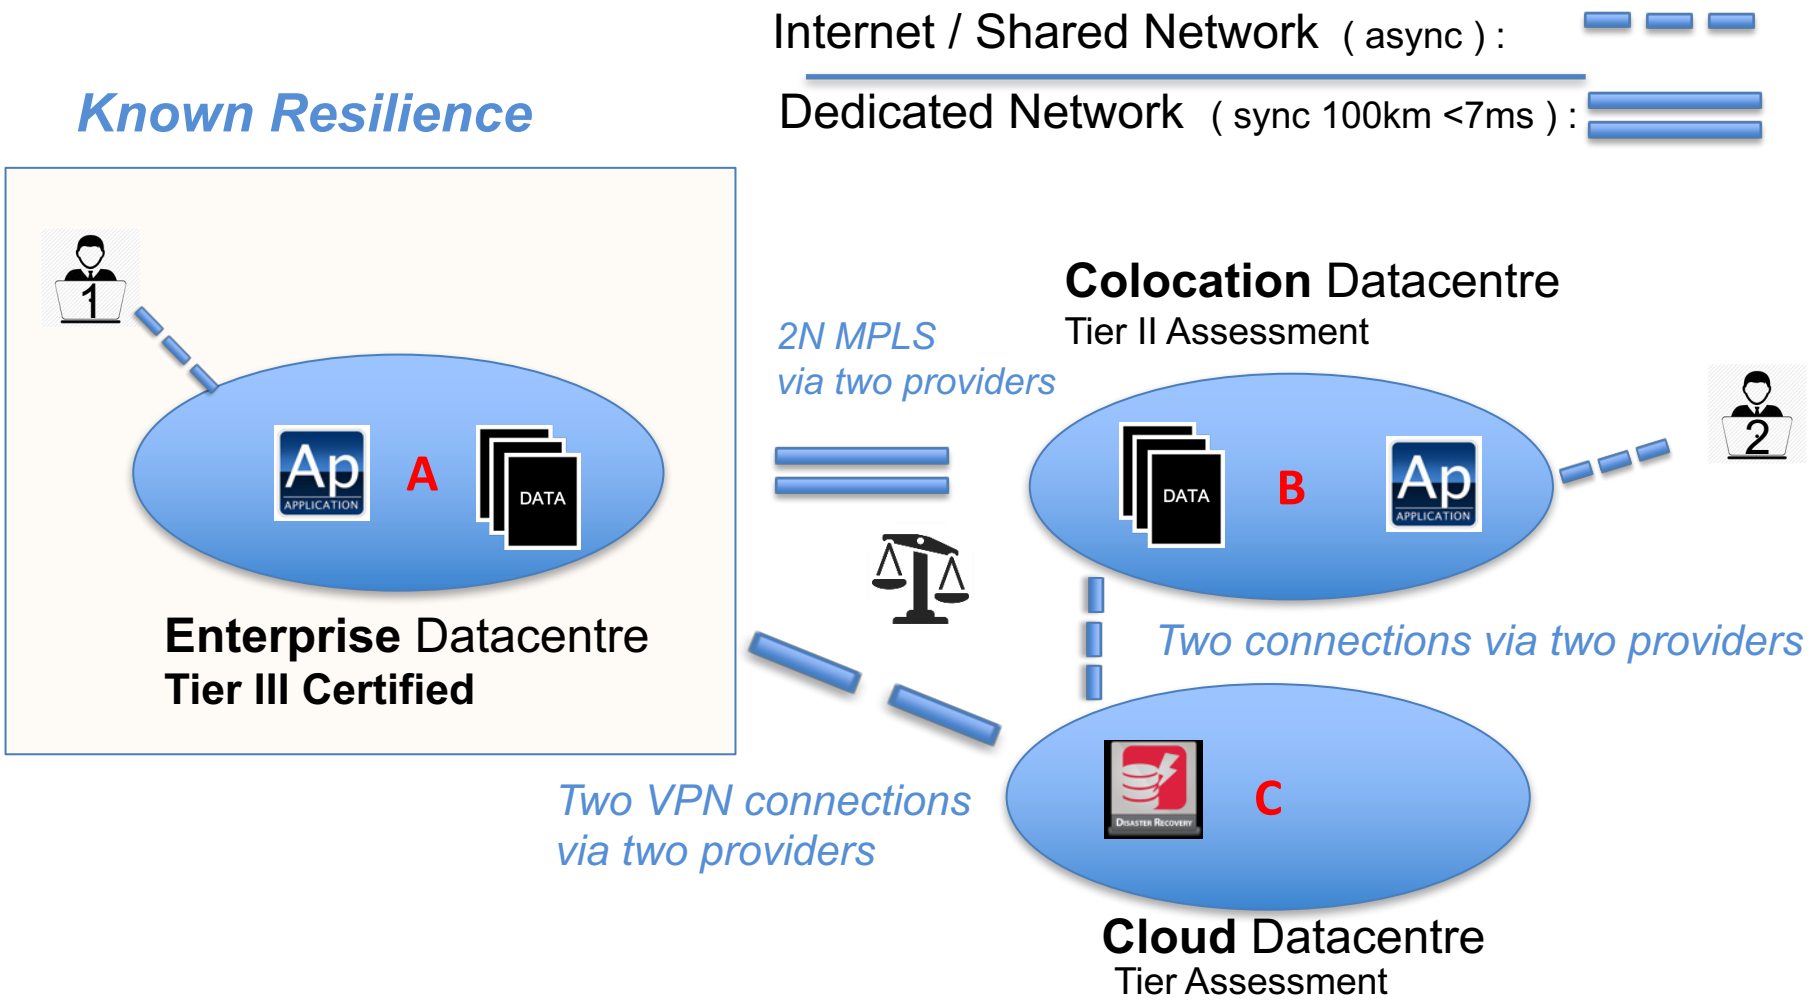
© 2017 Uptime Institute, LLC

Multi-Data Center Use Case



© 2017 Uptime Institute, LLC

Multi-Datacentre Use Case



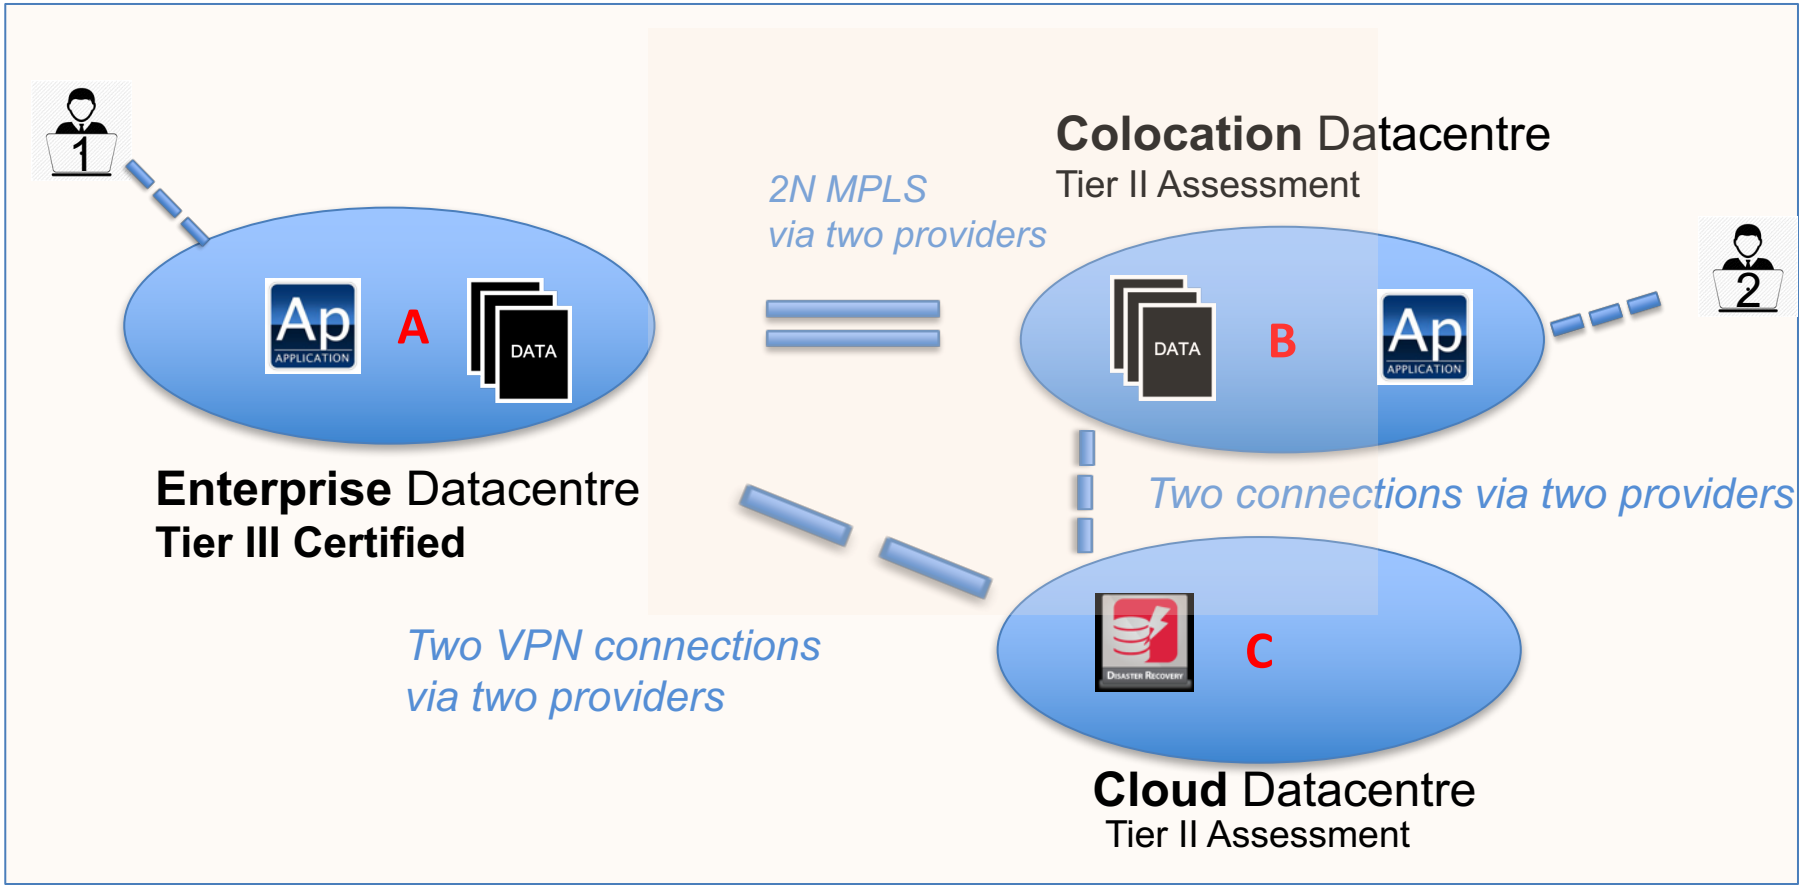
© 2017 Uptime Institute, LLC

Multi-Datacentre Use Case

Internet / Shared Network (async) :

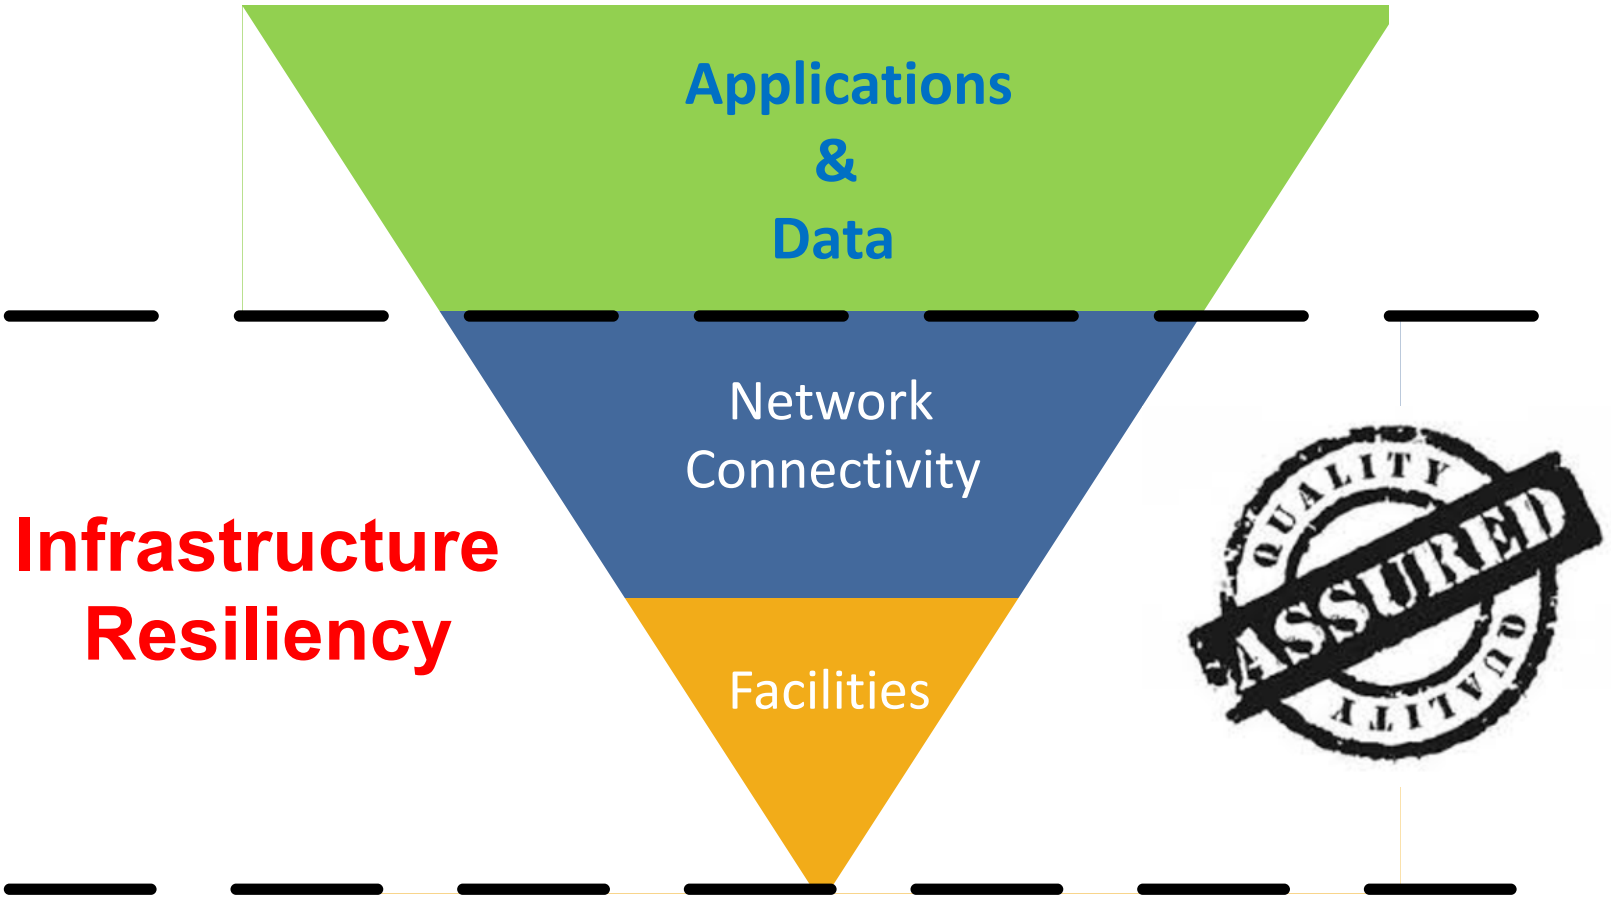
Dedicated Network (sync 100km <7ms) :

Unknown Resiliency



© 2017 Uptime Institute, LLC

Multi-Datacentre Resiliency Principle



© 2017 Uptime Institute, LLC

Multi-Site Resiliency Overview

- **Performance Based Network Resiliency Assurance**
  - Utilizes Uptime Institute's Networking Resiliency Protocol
    - Based on globally recognized Tier performance based classification system
    - Stand alone or multi-site approach
  - Physical and Logical Component Analysis
    - Network equipment resiliency
    - Network path resiliency
    - Throughput and Failover capabilities
    - Design and Operational evaluation
- **Datacentre Resiliency Assurance**
  - Uptime Tier Reliability Analysis to Determine Site Capabilities

© 2017 Uptime Institute, LLC

Multi-Datacentre Network Resiliency Matrix

	Connected	Redundant	Concurrent	Fault Tolerant
Physical Path diversity	No	No	No	Yes
Carrier Diversity	Yes	Yes	Yes	No
Recovery Path	Act/Pas	Act/Act	Act/Act(+1)	(2)Act/Act
Path Quality (Primary)	Shared	QoS	Dedicated capacity	Dedicated capacity
Path Quality (Recovery)	Shared	Shared	Dedicated capacity	Dedicated capacity
Path Control	None	SLA	SLA w/Dedicated capacity	Private
Power and Cooling Diversity	No	No	Yes	Yes
Network Equipment Physical Redundancy	N	N	N+1	2N
Network Equipment Logical Resiliency	No	<50%	>50%	Yes
Transport Resiliency	N	N	N+1	2N
Network Operational Processes	No	Yes	Yes	Yes
Physical Security	Yes	Yes	Yes	Yes
Logical security	Yes	Yes	Yes	Yes

© 2017 Uptime Institute, LLC

Multi-Datacentre Infrastructure Resiliency

		N Configuration				N+1 Configuration				Concurrent Maint				Fault Tolerant			
Enterprise Data Center	- Tier Rating	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4
Network	- Resiliency Certification	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
Colocation or Cloud Data Center	- Tier Assessment	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4
Infrastructure Resiliency Assurance Level		1	1	2	2	2	3	4	4	2	3	4	4	1	4	4	4

Performance Based Application & Data Availability = Multi-Datacentre Infrastructure Assurance

Infrastructure Resiliency Level Characteristics :

- = Infrastructure cannot be counted on to be resilient
- = Infrastructure provides resiliency during scheduled maintenance
- = Infrastructure provides fault tolerant resiliency

Application Design Characteristics :

- = Application must be loosely coupled to be resilient
- = Application may utilize infrastructure to avoid planned outages
- = Application may utilize infrastructure to be resilient against unplanned outages

© 2017 Uptime Institute, LLC

Uptime Institute Resiliance Directions…

- Assure Datacentre Availability through Tier Certification

- Provide Infrastructure Resiliency Assurance Qualification for Multi-Datacentre Infrastructures

Resiliency Assurance Developed with Industry Clients and Uptime Members to :

- Understand and articulate the value and benefits of new distributed models of IT resiliency

- Identify and understand the critical vulnerabilities, weaknesses and costs of distributed resiliency

- Evaluate methodologies and assurance processes for evaluating and rating resiliency, and the value of applying such methods across the industry

© 2017 Uptime Institute, LLC