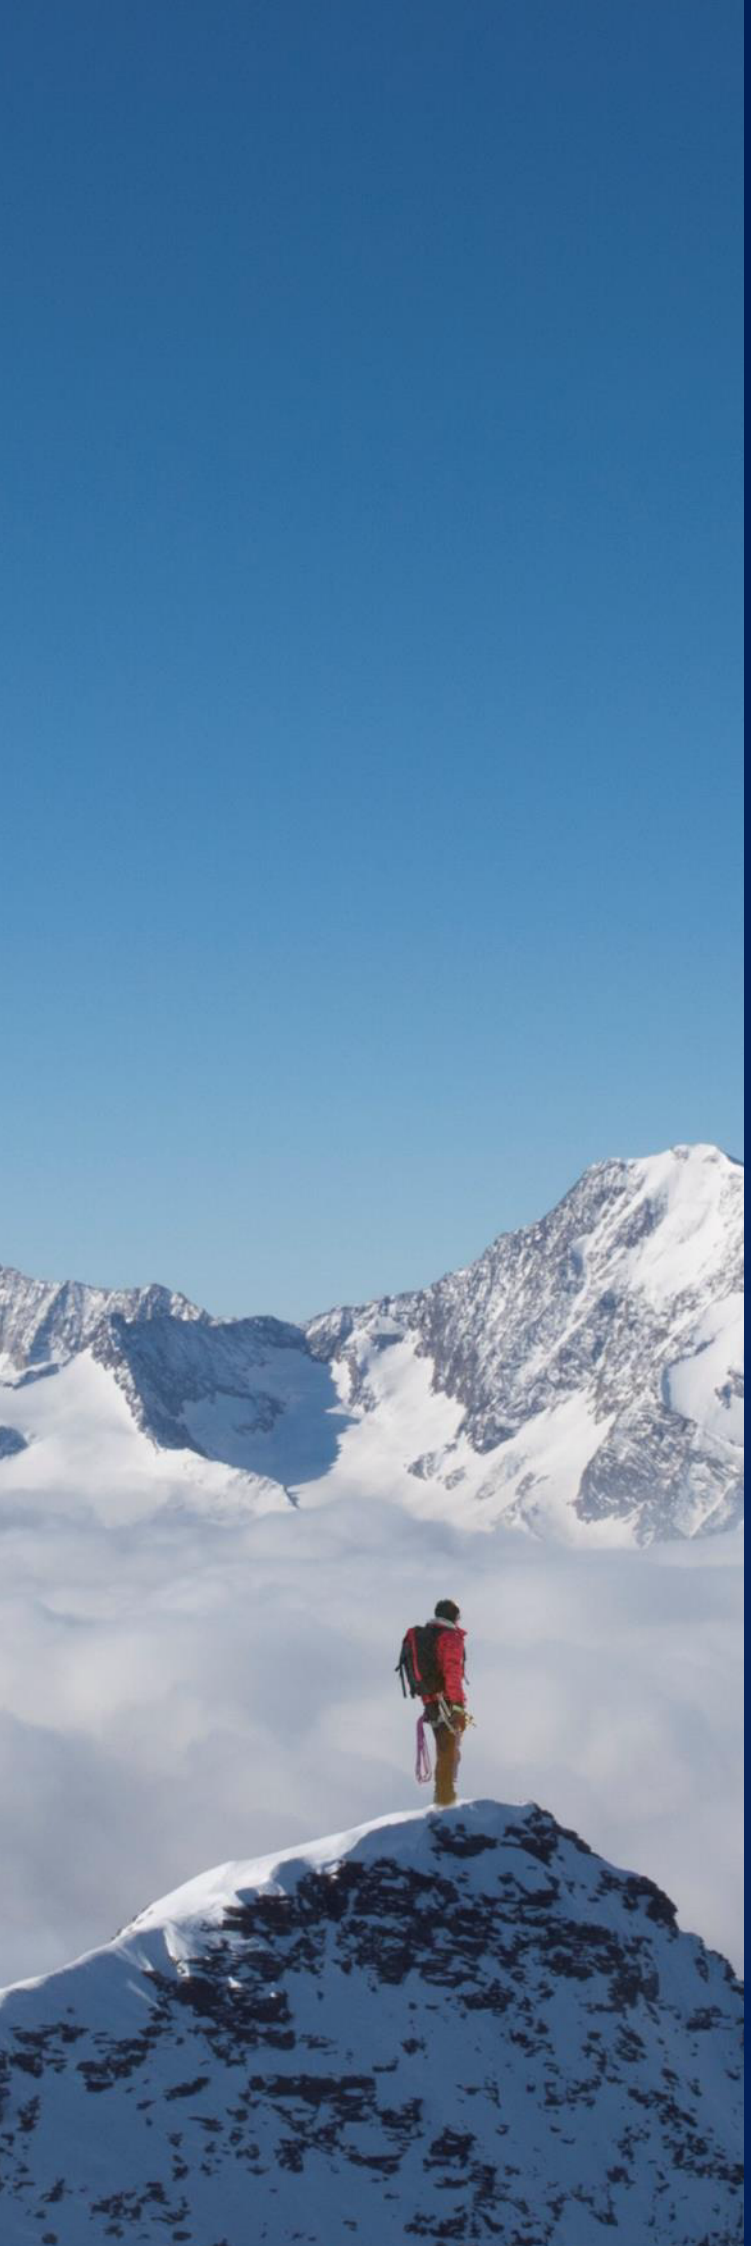

DREW & NAPIER

DREWTECH SERIES

CHAPTER 19

WHEN THE LIGHTS GO OFF – A REVIEW OF KEY LEGAL CONSIDERATIONS FOR HIGH-AVAILABILITY SYSTEMS

3 July 2026

LEGAL UPDATE

In this Update

An ever-increasing reliance on information technology systems, big data, and AI has as a corollary a need for systems that stay up and keep on going, even when the infrastructure around them breaks down. These are known as high-availability systems. This article examines key and oft-overlooked considerations for any organisation working with high-availability systems.

03
INTRODUCTION

03
UPTIME

04
DESIGN IMPLICATIONS

05
REGULATORY COMPLIANCE LANDSCAPE

06
CONCLUSION

INTRODUCTION

Technology has permeated every aspect of commerce. That much is well established. What is relatively newer is the degree to which almost every aspect of daily life now requires stable, resilient, and fault-tolerant technology infrastructure. Trivial examples include internet-connected smart beds which were stuck in inclined positions and heated up uncontrollably during a 2025 server outage. In a much more sobering example, two hospitals had to move to pen and paper during a server outage as a result of an intense heatwave. There is simply no alternative – organisations must ensure that their technology remains up, at the risk of both financial and human loss.

The legal edifice supporting high availability is built primarily in contract. Any competently drafted agreement should define the uptime commitment and specify the consequences of breach. It should deal with allocation of risk, which needs to be calibrated based on design choices made in the development of the infrastructure. It should deal with any regulatory requirements that may be imposed on your organisation.

UPTIME

What then should you look out for when negotiating an SLA? A key concept is the percentage of assured uptime. This is typically expressed as “nines” of uptime: 99.9% (three nines) permits roughly 8.7 hours of downtime per year; 99.99% (four nines) allows only 52 minutes; and 99.999% (five nines) permits a mere 5.26 minutes. The contractual machinery matters as much as the headline percentage. Scheduled maintenance, force majeure, customer-caused outages, degraded performance, regional outages and the method of measurement may materially affect whether downtime has occurred.

The decision as to the level of uptime required depends on the use case that the technology is supporting. Customers may tolerate being unable to doomscroll for the duration of a lunch break. Where seconds matter and lives are at stake, higher standards may be required.

Assurances are one side of the coin. The necessary concomitant is the remedies for breach. Service credits are the most common remedy for breaches of service level agreements. They operate as a contractually agreed credit or adjustment, for example as a percentage of monthly fees credited back to the customer for each hour or fraction of downtime below the uptime target. However, a critical drafting question is whether service credits constitute the exclusive remedy for downtime or merely the minimum remedy, preserving the customer’s right to claim general damages. From the vendor’s perspective,

exclusivity provides commercial certainty; from the customer's perspective, general damages may be necessary where losses vastly exceed a few months' fees.

DESIGN IMPLICATIONS

An interesting point in developing and operating a high-availability system is that it often involves significant physical infrastructure considerations not seen in more straightforward software builds. One such consideration is the power source being used. One individual does not need to deeply consider the electrical supply they receive from mains power. A large data centre, running operations for multiple organisations, needs to carefully consider and design for this.

To manage the need for a conditioned, regular flow of power, data centres will almost invariably have onsite backup power supply to deal with a widespread power outage. This is more complicated than simply starting an onsite diesel generator – as mentioned above, certain operating conditions may require downtime of not more than about 5 minutes a year, shorter than it could take for a generator to come online. Fortunately, there are established bridging solutions. One option is a "static uninterruptible power supply" – essentially, a set of chemical batteries that will provide power during the brief interval between loss of mains power and the backup generator start-up process. The other is a "rotary uninterruptible power supply" – a heavy flywheel that spins at dizzying speeds to store kinetic energy which can be converted back into electrical energy when needed.

The choice of technology is important. Battery installations engage chemical safety regulations applicable to battery storage (including lithium-ion fire risks that have prompted regulatory attention globally), hazardous substance disposal, and occupational health obligations for maintenance personnel. Rotary uninterruptible power supplies involve heavy metal moving at high speed, and may require additional professional and structural engineering considerations.

While these will naturally be of interest to anyone developing and operating a data centre itself, it is also relevant to consumers. When entering into any agreement for supply of high-availability infrastructure, it is important to understand how your vendor will be ensuring the uptimes that it has contractually agreed to, so that you can satisfy yourself that there is a technical basis for the legal obligations that they have agreed to. While service credits may be all well and good, it can be cold comfort when faced with questions from your own downstream customers who cannot access your service.

REGULATORY COMPLIANCE LANDSCAPE

The foregoing is of general concern. However, organisations should additionally be aware of any additional requirements or obligations imposed on them which they should take into account when operating a high-availability system. A selection of possible concerns is set out below.

The Cybersecurity Act 2018 designates Critical Information Infrastructure (CII) across specified sectors (e.g., energy, healthcare, banking, transport, and others) whose continuous operation is necessary for national security, public safety, or the delivery of essential services. CII owners are subject to obligations including compliance with applicable codes of practice or standards of performance, regular cybersecurity audits and risk assessments, and reporting of prescribed cybersecurity incidents within prescribed timeframes. If any organisation is required to comply with the Cybersecurity Act and at the same time relies on a third-party vendor to provide technical support, it would be well advised to consider whether and how it can ensure that it can receive from its vendor the level of support it requires.

Singapore has also moved towards regulating foundational digital infrastructure more directly. The policy concern is reflected in the Government's study of a Digital Infrastructure Act to address broader resilience risks beyond cybersecurity (such as misconfigurations, physical hazards such as fires, and cooling system failures) for digital infrastructure providers that may not provide essential services but are heavily relied upon by the economy. The anticipated scope for the Digital Infrastructure Act signals that high-availability obligations will become increasingly statutory rather than purely contractual.

Organisations should note the interaction between these frameworks and sector-specific requirements such as the MAS technology risk management guidelines and notices. The regulatory landscape is converging: what was once sector-specific guidance for financial institutions is becoming an economy-wide obligation. In practice, failures by third-party infrastructure providers may expose regulated entities to direct regulatory liability, even where the root cause lies outside their control.

CONCLUSION

High availability is therefore not merely a technical aspiration, but a legal and operational discipline. Organisations procuring or providing high-availability systems should ensure that contractual commitments on uptime are matched by a realistic understanding of the infrastructure required to deliver them, the remedies available if they are not met, and the regulatory obligations that may apply where outages affect critical or widely relied-upon services. As dependence on digital infrastructure deepens, resilience will increasingly need to be built into contracts, systems, and compliance frameworks from the outset – before the lights go off.

UPDATES IN DREWTECH SERIES

1. Chapter 1: The Importance of an Exit Strategy in Technology Contracts <6 March 2019>
2. Chapter 2: Employees, Technology and A Legal Hangover – Bring Your Own Problems? <4 June 2019>
3. Chapter 3: I host, you post, I get sued? <24 September 2019>
4. Chapter 4: Diabolus ex machina <18 February 2020>
5. Chapter 5: Bringing hygiene online – the MAS notice on cyber hygiene <28 April 2020>
6. Chapter 6: Signing without signing – contactless contracts <16 July 2020>
7. Chapter 7: My Kingdom for a Horse – When your Systems are Held to Ransom <22 January 2021>
8. Chapter 8: New risks in new skins – Updates to the Guidelines on Risk Management Practices – Technology Risk <3 March 2021>
9. Chapter 9: Of blockchains and stumbling blocks <21 July 2021>
10. Chapter 10: Service by airdrop – no parachutes required <7 July 2022>
11. Chapter 11: Large Language Models and Larger Legal Minefields <4 April 2023>
12. Chapter 12: Beset on all sides – liability for data breaches <19 July 2023>
13. Chapter 13: Pitfalls of user-generated content <18 October 2023>
14. Chapter 14: Red queen races – vulnerability disclosure programs <1 August 2024>
15. Chapter 15: Looking at the man in the middle (in a cyber breach) – allocation of risk <5 March 2025>
16. Chapter 16: Speak, Friend, and Enter - Access Controls and Authorised Users <28 May 2025>

17. Chapter 17: I love the smell of AI in the morning – a review of an AI-generated letter of demand <12 November 2025>
18. Chapter 18: Schrodinger's Safe – Quantum Computing and the Law <22 April 2026>
19. Chapter 19: When the lights go off – a review of key legal considerations for high-availability systems <3 July 2026>

If you have any questions or comments on this article, please contact:

Rakesh Kirpalani
Director, Dispute Resolution & Information Technology
Chief Technology Officer
T: +65 6531 2521
E: rakesh.kirpalani@drewnapier.com

Ang Wei Shuen
Associate Director, Dispute Resolution & Information Technology
T: +65 6531 2797
E: weishuen.ang@drewnapier.com

DREW & NAPIER

Drew & Napier LLC
10 Collyer Quay
#10-01 Ocean Financial Centre
Singapore 049315

www.drewnapier.com

T : +65 6535 0733
T : +65 9726 0573 (After Hours)
F : +65 6535 4906