

Service Level Agreement

This **Service Level Agreement** ("SLA") sets forth the specific terms and conditions under which Lightedge Solutions, Inc. ("Lightedge") shall supply all Managed Services to Customer. The Master Agreement entered into between Lightedge and Customer fully incorporates the terms herein and provides that this Service Level Agreement, and Customer's execution of the Master Agreement constitutes acceptance of the terms and conditions stated herein. Capitalized terms used but not defined herein shall have the meanings set forth in the Master Agreement.

1. Definitions.

1.1 General Terms

"Maintenance Window" is the timeframe within which Service Changes are performed.

"Major Outages" are severe Service Outages affecting multiple Customers that require immediate corrective action.

"Service Change" is any configuration or maintenance change made to Service or Service Platform, whether Customer initiated or Lightedge initiated.

"Service Platform" refers to all physical gear used or required to deliver Service.

"Service Portal" refers to the web-based, Customer accessible portal by which Services can be managed and Tickets against Service Incidents and Service Requests can be opened. Service Portal can be accessed at <https://my.lightedge.com>.

"Service Requests" are general inquiries related to fulfilling standard changes, responding to requests for information, and fulfilling requests for access to Services. Service Requests are initiated by the Customer opening a Ticket with Lightedge Operations via phone, email, or the Service Portal.

"Ticket" is a documented record in the Lightedge Operations support database of an interaction with Customer. Tickets can be initiated by Customer or Lightedge.

1.2 Service Classification Terms

"Service Incident" is an unplanned event affecting a single Customer and causing a reduction in service quality but not yet causing a Service Outage. Although Lightedge may detect and proactively open a Ticket, Service Incidents are typically reported to Lightedge by the Customer.

"Service Problem" is a Service Incident affecting many Customers, or an event, or the potential for an event, which could cause an unplanned service interruption or degradation affecting many Customers.

"Service Outage" refers to a single Service Incident or Service Problem resulting in complete loss of Service or period(s) of total unavailability of Service.

"Service Failure" refers to any failure to meet any Service Level commitments as described in Sections five (5) and six (6) of this Service Level Agreement.

2. Claims and Exclusions.

2.1 SLA Credit Request Process.

Service Level Agreements (or SLAs) define availability, performance and other requirements of Lightedge in providing the Service(s). In the event Customer's Service fails to meet the applicable commitments outlined in this SLA, Customer shall be entitled to a credit adjustment to its account in accordance with the SLA Credits defined within this Service Level Agreement.

Customer must request any SLA Credit within 90 days of the event giving rise to the request by contacting Lightedge Accounting and requesting an "SLA Credit". Customer request for SLA Credit must be accompanied by a Service Ticket related to the Service Failure. SLA Credits will appear on Customer's bill within two (2) billing cycles.

2.2 SLA Credit Calculations.

"Affected Service" is portion of Service(s) affected or impacted by any Service Failure.

For the Affected Service, a percentage of that Service's monthly fees will be credited to Customer based on one of the following methods:

Percentage	Percentage SLA Credits are calculated as: Affected Service Monthly Recurring Charges x SLA Credit Percentage <i>For example, if monthly charges for affected Service total \$10k/month and SLA Credit is 20%, the Customer will receive credit of \$2k for the month during which outage occurred.</i>
Hourly	Hourly SLA Credits are calculated as: Affected Service Monthly Recurring Charges x (Total Hours SLA Credit / 720) <i>For example, if monthly charges for affected Service total \$10k/month and SLA Credit is 3.5 hours, the Customer will receive credit of \$48.61 (\$10k * (3.5 hrs / 720 total hours in month) for the month during which outage occurred.</i>

SLA Credits are based on the Monthly Recurring Charges normally paid to Lightedge for the affected Service and are exclusive of any applicable taxes, pass-through costs, or one-time charges.

In any calendar year, Customer's aggregated SLA Credits may not exceed, for any Service, four (4) months' worth of the monthly Service fees for the affected Service. In any billing month SLA Credits may not exceed, for any Service, fifty (50) percent of the monthly Service fees for the affected Service.

SLA Credits are made pro rata against the portion of Service affected by Service Failure:

Colocation Services	For any Service Failure of Colocation Services, SLA Credits are applied to the pro rata Affected Service Monthly Recurring Charges for the affected Colocation service(s) and any associated electrical service(s). <i>For example, if Customer is contracted for 50 kW of electricity delivered across many branch circuits and one branch circuit capable of supplying 10kW electricity is lost, the Customer will receive an SLA Credit will be calculated relative to the 10kW electricity (1/5th of electrical fees) and any impacted racks supplied by that branch circuit, not the entire 50kW of electricity and all colocation space.</i> <i>For example, if monthly charges for affected Service total \$10k/month and SLA Credit is 20%, the Customer will receive credit of \$2k for the month during which outage occurred.</i>
Cloud Services	For any Service Failure of Cloud Services, SLA Credits are applied to the pro rata Affected Service Monthly Recurring Charges for the affected Cloud Service(s). <i>For example, if Customer is contracted for a resource pool of 100 vCPU, 200GB of RAM, and 500GB of disk space, and one (1) VM of 1 vCPU, 2GB of RAM and 50GB of disk space fails, then SLA Credit will be calculated relative to 1/100th of the Monthly Recurring Charges for the entire cloud service.</i>
Infrastructure Services	For any Service Failure of Infrastructure Services, SLA Credits are applied to the pro rata Affected Service Monthly Recurring Charges for the affected Infrastructure Service(s). <i>For example, if Customer is contracted for ten (10) Bare Metal blades and one (1) blade is affected by a Service Failure, the SLA Credit will be calculated relative to the monthly contract value for that affected one (1) Service component, not the nine (9) unaffected Service components.</i>
Security Services	For any Service Failure of Security Services, SLA Credits are applied pro rata to the Monthly Recurring Charges for the affected Security Services and any add-on Services unavailable due to the Service Failure. <i>For example, if Customer is contracted for a Managed Firewall and a Service Failure causes that firewall to be unavailable, any SLA Credit due is applied against the affected firewall and not any Services being provided behind the firewall that may be unavailable due to a Service Failure of the firewall.</i>
Business Continuity Services	For any Service Failure of Business Continuity Services, SLA Credits are applied to the pro rata Affected Service Monthly Recurring Charges for the affected Business Continuity Service(s). <i>For example, if Customer is contracted for 10TB of backup on ten (10) servers and the backups fail on one (1) server consuming 1TB, then SLA Credit will be calculated relative to 1/10th of the Monthly Recurring Charges for the entire backup service.</i>
Connectivity Services	For any Service Failure, SLA Credits are applied to the Affected Service Monthly Recurring Charges for the affected Connectivity Service(s). <i>For example, if Customer is contracted for 1 gigabit of Internet with Standard DDoS Protection delivered via 2 cross-connects and a Service Problem or DDoS attack renders the Internet unreachable by Customer, then SLA Credit will be calculated relative to the Monthly Recurring Charges for all affected Connectivity Services.</i>

SLA Credits based on Downtime or Impact begin immediately. *For example, an SLA credit of "5% for each 30-minute period" will be calculated as follows: 1 second to 30 minutes Downtime = 5% SLA Credit; 30.1 to 60 minutes Downtime = 10% SLA Credit; 60.1 to 90 minutes Downtime = 15% SLA Credit. This continues until Service Credit maximums have been reached.*

2.3 SLA Exclusions.

SLAs do not apply and Lightedge is not responsible for failure to meet an SLA resulting from:

- failure of Customer to comply with other Lightedge agreement terms including the Master Agreement and Acceptable Use Policy which substantially attributes to Service Failure.
- failure of Customer to reasonably cooperate with Lightedge during testing, installation, maintenance or troubleshooting activities, and such cooperation is not provided after Lightedge notifies Customer that it is not cooperating, which substantially attributes to Service Failure or substantially extends the duration of Service Failure.
- Service Failures solely caused by Customer or agent of Customer.
- an inability to utilize Service due to a network issue outside of Lightedge control. *For example, if Customer is unable to access their equipment collocated at a Lightedge data center because the Internet at their office went down, that does not constitute a Service Failure of the Lightedge Colocation Service.*
- failure to adhere to Lightedge recommended configurations as documented in Service Guides and communicated to Customer via Ticket prior to such Service Failure.

- for any Services where Service Platform infrastructure must be deployed outside of a Lightedge data center, when Customer fails to provide suitable secure environment for infrastructure including but not limited to securely mounting/racking infrastructure, appropriate cooling and air handling of environment or securing infrastructure from theft.

In addition, SLAs do not apply:

- where Customer reports an SLA failure, but Lightedge does not find any SLA failure using commercially reasonable efforts to do so.
- when a Service Failure could have reasonably been prevented or mitigated by Customer. *For example, if Lightedge provides redundant Ethernet connections to Customer but Customer only utilizes one connection, and Service Incident or Problem would not have impacted Customer had they been properly utilizing the redundant connection provided by Lightedge.*

3. Operations Center and Support.

Lightedge Operations Center is staffed 24x7x365 and is responsible for the management and support of all Lightedge Services and associated Service Level Agreements. Communication with the Lightedge Operations Center will be in the English language.

3.1 Contact Information

<u>Lightedge Operations Center (24x7x365)</u> Phone: 1.877.589.3654 Ticketing: Opened via https://my.lightedge.com support portal
--

3.2 Planned Maintenance.

Service Changes may be required in order to keep Service in good operating order and will be referred to as "Planned Maintenance". Planned Maintenance that is expected to impact Service Delivery will be limited to Maintenance Windows as defined under Service Level Agreements and will not occur without Customer notification as detailed in accordance with Section 6 below. Except where otherwise noted, Service SLAs will not apply during Planned Maintenance.

3.3 Emergency Maintenance.

Major Outages or security events may require immediate or prompt maintenance. Any such maintenance will be considered "Emergency Maintenance". Emergency Maintenance may occur outside Planned Maintenance windows. Lightedge will make a reasonable effort to notify the Customer if feasible under the circumstances prior to Emergency Maintenance. Except where otherwise noted, Service SLAs will apply during Emergency Maintenance.

3.4 Service Monitoring and Notifications.

At Customer's request, Lightedge will notify Customer of a Service Outage via e-mail. Lightedge will provide 24x7x365 response to Customer or Lightedge Operations Center initiated alarms for Service availability issues.

3.5 Service Reporting and Cooperation.

Lightedge and Customer shall each use commercially reasonable efforts to keep the other informed regarding Service Incidents, Service Problems and Service Outages and shall share any information in their possession that would be reasonably available and necessary to determine the nature and scope of any Service Incidents, Service Problems and Service Outages. In addition, after any Service Outage, Lightedge will conduct a review to determine root causes of the Service Outage and actions that may be taken to prevent future Service Outages.

3.6 Helpdesk Response SLAs.

Lightedge maintains an Operations Center with enough staffing to promptly respond to Service Requests and Service Incidents.

Helpdesk Response SLAs are the Service Levels that Lightedge maintains for response to Service Requests. Any SLA Credits due for Helpdesk Response SLAs are applied to Monthly Recurring Charges for the affected service(s).

Helpdesk Response SLAs are provided during periods of normal operations and are exempt during periods of Force Majeure as defined in the Master Agreement.

Incident Response Times measures the duration between Ticket being opened by Customer and initial response by Lightedge to that Ticket.

Incident Update Times measures the duration between Ticket updates, excluding any time spent waiting for Customer.

Incident Severity	Incident Examples	Incident Response Times	Incident Update Times	Incident Resolution Times	SLA Credit
Critical	System or application is down and not responsive.	< 15 minutes	Every hour	All Service Incidents and	1 hour for each missed event

High	System or application is performing below normal levels, or there is an indication that something needs to be done to prevent an impending outage.	< 30 minutes	Every 2 hours	<i>Service Outages are resolved as quickly as possible.</i> <i>Remedy for not quickly restoring Services is already provided by Service SLA Credits.</i>	0.5 hour for each missed event
Moderate	Loss of redundancy but system is performing as normal, some service or task is needed by a certain time.	< 2 hours	Every 12 hours		0.25 hour for each missed event
Low	Some service or task is needed but not time sensitive, request for information.	< 24 hours	Best effort updates		0.1 hour for each missed event
Planning	Non-critical notes or informational communication.	Best effort response	Best effort updates		None

THIS SPACE INTENTIONALLY LEFT BLANK

COLOCATION SERVICES SLAs

4. Colocation Services.

"Colocation Services" are Private Suite, Cage, Rack or Shared Colocation Services residing in a Lightedge data center.

4.1 Colocation Planned Maintenance

Planned Maintenance for Colocation Services will be limited as follows:

Schedule: only performed between 12AM - 4AM Central Time
Notification: minimum of three (3) week advance notification to Customer

4.2 Electrical Power SLAs.

Service Level commitment for electrical power delivery in a Lightedge data center is:

● PRIVATE SUITE COLOCATION ● CAGE COLOCATION ● RACK COLOCATION ● SHARED COLOCATION									
Service Level	Service Level Agreement		SLA Credit						
REDUNDANT ELECTRICAL POWER 100%	DESCRIPTION: Lightedge delivers electrical power associated with Colocation Services in a redundant manner. Lightedge electrical supply infrastructure is designed and maintained for mission critical use. All key infrastructure components are maintained to N+1 redundancy. Power is delivered to Customer in redundant 2N configuration of A & B circuits. Power is always available (meaning electricity will always be available on EITHER the A OR B circuit even during Planned Maintenance. For redundant power, Customer shall not exceed 80% of the rated capacity of a single circuit of the combined A & B circuit pair or Service availability during maintenance or outage may be compromised. <i>For example, if the A circuit is loaded at 80%, the B circuit should not have any load so that the B circuit remains fully available should the A circuit fail. If the A circuit is loaded at 20%, the B circuit load should not exceed 60%.</i> Utility power is protected by Uninterruptible Power Supply ("UPS") systems or battery systems and backed up by generator power, with neither the A nor B feed sharing any critical components, and each feed being protected by a dedicated UPS system and generator. Generator power is exercised at least weekly, and a full load failover testing both the UPS systems and generator power is performed at least annually. Generator fuel sufficient for a minimum of one day runtime is maintained onsite. Priority fuel delivery agreements are maintained for each data center. Redundant electrical power is monitored in real-time by Lightedge via remote monitoring system that monitors critical Service components within the data center such as generator health, UPS health and available runtime, environmental controls, and electrical switch gear status and power delivery. Customer is responsible for attaching and maintaining their equipment in a manner that can utilize the redundant electrical power provided by Lightedge in an industry acceptable manner such as redundant power supplies connecting to the A & B circuits provided by Lightedge in a 2N manner. MEASURED BY: Service Failure occurs when usable and redundant electrical power is not supplied to Customer as follows:		10% for any Service Failure up to 30 minutes 5% for each additional 30 minutes of Service Failure						
	<table><tr><td>CAGE OR PRIVATE SUITE</td><td>RACK with CUSTOMER PROVIDED POWER STRIPS</td><td>SHARED COLO or RACK w/ LIGHTEDGE PROVIDED POWER STRIPS</td></tr><tr><td>Measured at output side of distribution panels (PDUs) feeding Customer space</td><td>Measured at input to Customer rack power strips (CDUs)</td><td>Measured at output of Lightedge supplied rack power strips (CDUs)</td></tr></table>			CAGE OR PRIVATE SUITE	RACK with CUSTOMER PROVIDED POWER STRIPS	SHARED COLO or RACK w/ LIGHTEDGE PROVIDED POWER STRIPS	Measured at output side of distribution panels (PDUs) feeding Customer space	Measured at input to Customer rack power strips (CDUs)	Measured at output of Lightedge supplied rack power strips (CDUs)
	CAGE OR PRIVATE SUITE	RACK with CUSTOMER PROVIDED POWER STRIPS		SHARED COLO or RACK w/ LIGHTEDGE PROVIDED POWER STRIPS					
	Measured at output side of distribution panels (PDUs) feeding Customer space	Measured at input to Customer rack power strips (CDUs)		Measured at output of Lightedge supplied rack power strips (CDUs)					
Service Failure begins:									

	a) when power to BOTH A and B power circuits as delivered to Customer are unavailable at any time including during Planned Maintenance resulting in a complete interruption in electrical power supplied to Customer; OR b) when contracted power is unavailable to EITHER A or B power circuits as delivered to Customer power strip (CDU) except for periods of Planned or Emergency Maintenance, or during periods when the Data Center is actively experiencing an emergency which causes power to be diverted to either the A or B power circuits. <i>In other words, each of the A and B circuits by itself must be able to provide the contracted power, without the other circuit.</i> Service Failure ends when usable and redundant electrical power is supplied to Customer.	
NON-REDUNDANT ELECTRICAL POWER 99.99%	DESCRIPTION: Non-redundant power is an uncommon delivery method and would only be applicable if Customer specifically requested it. Lightedge recommends all Customers utilize the default (redundant) power delivery model defined above. In the case Customer requests non-redundant power, power circuits are delivered to Customer in non-redundant N configuration. For non-redundant power, Customer shall not exceed 80% of the rated capacity of any power circuit. Power may not be available on EITHER the A or B circuit during Planned Maintenance, allowing for maintenance of electrical supply gear. Non-redundant electrical power is monitored in real-time by Lightedge via remote monitoring system that monitors critical Service components within the data center such as generator health, UPS health and available runtime, environmental controls, and electrical switch gear status and power delivery. MEASURED BY: Service Failure occurs when usable electrical power is not supplied to Customer in accordance with Redundant Electrical Power SLA above. Service Failure begins when power to EITHER A or B power circuit delivered to Customer is unavailable outside of Planned Maintenance, resulting in a complete interruption in electrical power to Customer. Service Failure ends when usable electrical power is supplied to Customer.	5% for each 1 hour of Service Failure in excess of SLA

4.3 Cooling & Humidification SLAs.

Service Level commitment for air temperature and humidity in a Lightedge data center is:

● PRIVATE SUITE COLOCATION ● CAGE COLOCATION ● RACK COLOCATION ● SHARED COLOCATION		
Service Level	Service Level Agreement	SLA Credit
AIR TEMPERATURE 64.4°F - 80.6°F 100%	DESCRIPTION: Lightedge data center cooling is maintained in a manner suitable for use by high density compute, storage and networking infrastructure. Cooling infrastructure is designed and maintained with N+1 redundancy and complies with the ASHRAE Recommended Range for all Class A Data Centers. Cooling is designed for mission critical use and remains always available even during Planned Maintenance. Lightedge maintains the cooling setpoint at 72°F. MEASURED BY: Temperature is measured every 15 minutes via remote probes installed by Lightedge within the data center (for shared colo or racks) or dedicated Customer space (for cage or private suite). Remote probes are positioned in the cold supply aisles at top of aisle and located in the centerline of each cold row. Remote probes are placed a minimum of eight (8) feet apart from each other. Service Failure occurs when data center temperature falls outside target cooling range. Service Failure begins when minimum of two (2) remote probes falls outside target cooling range during same measurement period. Service Failure ends when all probes have returned to target cooling range.	> 80.6°F = 5% for each hour of Service Failure > 89.6°F = 15% for each hour of Service Failure
HUMIDITY 30-60% relative humidity 100%	DESCRIPTION: Lightedge data center air humidity is maintained in a manner suitable for use by high density compute, storage and networking infrastructure. Humidification infrastructure is designed and maintained with N+1 redundancy and complies with the ASHRAE Recommended Range for all Class A Data Centers. Humidification is designed for mission critical use and remains always available even during Planned Maintenance. Lightedge maintains the humidification setpoint at 45% relative humidity.	5% for each 6 hours of Service Failure

	MEASURED BY: Humidity is measured every 15 minutes via remote probes used for temperature monitoring and described within the Air Temperature SLA. Service Failure occurs when data center humidity falls outside target cooling range. Service Failure begins when minimum of two (2) remote probes falls outside target humidification range during same measurement period. Service Failure ends when all probes have returned to target humidification range.	
--	--	--

4.4 Physical Security SLAs.

Service Level commitment for physical security in a Lightedge data center is:

● PRIVATE SUITE COLOCATION ● CAGE COLOCATION ● RACK COLOCATION ● SHARED COLOCATION		
Service Level	Service Level Agreement	SLA Credit
COLOCATION SECURITY AVAILABILITY Mitigation: 15 minutes Repair: 1 day	DESCRIPTION: Secured colocation space is provided to Customer as security caging or colocation rack(s) for purpose of storing and operating Customer electronic equipment in a Lightedge data center. Colocation racks are typically four post racks with combination lockable front and rear doors and locked side panels. Side panels on adjacent colocation racks will not be removed. Lightedge will maintain the colocation space and all security components such as cage entry points, doors, side panels and door locks, in good working order. In the event any colocation rack(s) needs repair or replacement of the security components, Lightedge will provide additional controls until the colocation rack(s) are returned to normal service. <i>For example, if a door lock were the be inoperable, Lightedge might require escorted access into that data center quadrant until the door lock was repaired.</i> MEASURED BY: Failures of customer secured colocation space must be reported to Lightedge by Customer. Service Failure occurs when any infrastructure related to critical physical security controls fails and is left unmitigated or unrepaired for an extended period, resulting in a breach of physical security of the Customer collocated equipment installed within the rack. Service Failure begins when Customer reports failure of security component of their secured colocation space, as documented by a Service Ticket. Service Failure for mitigation SLA ends when equivalent security control has been implemented and noted by timestamp in Service Ticket. Service Failure for repair SLA ends when failed infrastructure has been restored back into normal service. The physical security control repair SLA is provided during periods of normal operations and is exempt during periods of Force Majeure as defined in the Master Agreement.	> 15 minutes for mitigation = 1% credit for each Service Failure > 24 hours for repair = 5% credit for each Service Failure
DATA CENTER SECURITY RESPONSE 5 minutes	DESCRIPTION: All entry and egress points into the data center are secured against unauthorized access. Each Customer entry point into the data center is secured by card and biometric access controls and is limited to approved Customers or Lightedge personnel. Critical operational areas such as AHU/HVAC rooms, UPS/Battery rooms and electrical rooms are secured by card access control and are limited to Lightedge personnel and approved contractors. Doors designated as emergency exits are equipped with delayed egress hardware and secured from external entry. All secured doors are alarmed for door forced and door ajar conditions. Surveillance cameras are maintained and monitored by Lightedge that cover a) all building perimeter entrances and exits, b) all access points to Customer colocation areas and c) any critical infrastructure necessary to support data center operations such as generator yards and electrical rooms. Lightedge will monitor physical security of the data center and will promptly respond to any physical security issues and alarms. Lightedge will provide fire alarm, fire suppression and emergency notification systems covering Customer colocation areas and critical infrastructure rooms. These fire systems will comply with	> 5 minutes = 1% credit for each Service Failure

	requirements of the National Fire Protection Association (NFPA). Appropriate local and regional authorities will be notified upon activation of fire alarm or fire suppression systems. MEASURED BY: Physical security in each Lightedge data center is monitored for critical alerts by Lightedge 24x7x365 in real-time using a building management system (monitoring and management tool). Service Failure occurs when any alerts to critical security controls are not responded to promptly. Service Failure begins when a critical alert is generated from building management system related to physical security of the data center. Service Failure ends when Lightedge has documented response to critical alert as a Ticket and have begun investigation into the critical alert.	
DATA CENTER SECURITY AVAILABILITY Mitigation: 15 minutes Repair: 1 day	DESCRIPTION: Lightedge will maintain all equipment related to physical security such as doors, locks, surveillance cameras, and alarms, in good working order. In the event any security equipment needs repair or replacement, Lightedge will provide additional controls until such security equipment is returned to normal service. <i>For example, if a door alarm were to fail and require replacement, Lightedge might lock that door and require escorted access through the door with the failed alarm, until that alarm was repaired.</i> MEASURED BY: Physical security infrastructure in each Lightedge data center is monitored for failure by staff 24x7x365 in real-time using a building management system (monitoring and management tool). Service Failure occurs when any infrastructure related to critical physical security controls fails and is left unmitigated or unrepaired for an extended period, resulting in a breach of physical security of the data center. Service Failure begins when any infrastructure related to physical security in the data center fails. Service Failure for mitigation SLA ends when equivalent security control has been implemented and noted by timestamp in Service Ticket. Service Failure for repair SLA ends when failed infrastructure has been restored back into normal service. The physical security control repair SLA is provided during periods of normal operations and is exempt during periods of Force Majeure as defined in the Master Agreement.	> 15 minutes for mitigation = 1% credit for each Service Failure > 24 hours for repair = 5% credit for each Service Failure
DATA CENTER SECURITY RETENTION Surveillance video: 90-day retention Access logs: 1-year retention	DESCRIPTION: Physical security access logs and camera surveillance footage will be maintained by Lightedge for a minimum period-of-time. MEASURED BY: Physical security infrastructure in each Lightedge data center is monitored for failure by staff 24x7x365 in real-time using a building management system (monitoring and management tool). Service Failure occurs for each Service Request where Lightedge is unable to provide Customer with physical access logs or physical surveillance video falling within the retention period commitment.	5% credit for each Service Failure

DATA CENTER ACCESS REQUESTS 1 day	DESCRIPTION: Access requests grant or revoke physical access to Customer employees or agents to the Customer's Colocation Service. Access requests are initiated by Customer via Ticket. MEASURED BY: Service Failure occurs when any Customer access requests such as additions, deletions, revocations or changes are not fulfilled promptly by Lightedge. Service Failure begins when an access request is first requested by Customer, as documented by a Service Ticket. Service Failure ends when Lightedge has fulfilled the access request. Time elapsed between request for access change or revocation requests as documented by Ticket, and completion of change or revocation request.	> 1 day = 1% credit for each Service Failure
--	---	--

4.5 Cross Connect SLAs.

Service Level commitment for availability of cross connects is:

● CROSS CONNECTS		
Service Level	Service Level Agreement	SLA Credit
CROSS CONNECT AVAILABILITY: 99.99%	DESCRIPTION: Lightedge will maintain physical cross connects in good working order. This includes the structured cabling provided by Lightedge and used by Customer to transmit data. MEASURED BY: Cross Connect availability is not monitored by Lightedge. Lightedge recommends Customer uses industry standard monitoring tools to monitor network circuits traversing the Cross Connect. Service Failure occurs when Customer is unable to pass traffic across the cross connect provided by Lightedge due to a failure or damage to the physical cabling, except for periods of Planned or Emergency Maintenance. Service Failure begins when Customer reports issue with the cross connect by opening a Service Request. Service Failure ends when Lightedge restores the cross connect to working order as measured by Lightedge test gear.	5% for each 1 hour of Service Failure in excess of SLA

THIS SPACE INTENTIONALLY LEFT BLANK

CLOUD SERVICES SLAs

5. Cloud Services.

"Cloud Services" include Virtual Private Cloud, Dedicated Private Cloud, Public Cloud, Power Cloud and other compute services primarily managed by Lightedge.

5.1 Colocation Planned Maintenance

Planned Maintenance for Cloud Services will be limited as follows:

Schedule: only performed between 12AM - 4AM Central Time
Notification: minimum of one (1) week advance notification to Customer

5.2 Cloud Availability SLAs.

Service Level commitment for availability and performance of Lightedge Cloud Services:

● DEDICATED PRIVATE CLOUD ● VIRTUAL PRIVATE CLOUD ● IBMi/ POWER CLOUD ● WORKPLACE CLOUD ● HOSTING SERVICES – VIRTUAL DATA CENTER (Legacy) ● EDGE CLOUD		
Service Level	Service Level Agreement	SLA Credit
CLOUD SERVICE AVAILABILITY 100%	DESCRIPTION: Lightedge maintains Cloud Service Platforms with a variety of redundancies such as electrical, network, compute failover and storage protection to ensure uninterrupted service to Customer. Lightedge Cloud Services are designed and maintained for mission critical use. All key infrastructure components such as electrical and network connectivity are maintained to N+1 redundancy. Storage is protected using RAID or similar technologies. MEASURED BY: Cloud availability is measured every 5 minutes by Lightedge via remote monitoring system that monitors hundreds of Service components such as temperature, failed hardware components, network availability, service logs and software faults on all Cloud Service components. Service Failure occurs when Customer is unable to utilize Lightedge Cloud Services due to a problem with the Cloud Service infrastructure or software provided by Lightedge. <i>For example, a failure of a Lightedge compute host resulting in Customer VMs crashing would constitute a Service Failure. A Windows Server VM crashing due to an application fault would not constitute a Service Failure.</i> Service Failure begins when the monitoring system detects a Critical or Major event that results in an inability by Customer to utilize Lightedge Cloud Services, or such event results in the loss or complete disruption of Customer workload running on Lightedge Cloud Services. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 30 minutes of Service Failure
CLOUD SERVICE PERFORMANCE 100%	DESCRIPTION: Lightedge maintains the Cloud Service Platform in good working order and guarantees not to significantly oversubscribe or overload infrastructure, causing performance impact to Customer workload. MEASURED BY: Cloud performance is measured every 5 minutes via remote monitoring system that measures a variety of Service components such as host level CPU, disk performance, and network performance statistics on all Cloud Services components. Service Failure occurs when Cloud Service infrastructure is available but is adversely and substantially degraded such that Customer experiences an impact to their Service. <i>For example, an overloaded Lightedge compute host resulting in Customer VMs not having access to all of the memory they purchased, resulting in Customer VMs being available but having their performance adversely affected, would constitute a Service Failure. A Windows Server VM running out of memory due to Customer's application utilization all available memory would not constitute a Service Failure.</i> Service Failure begins when the monitoring system detects an abnormal or rise above normal thresholds that results in a substantial and adverse impact to Customer workload running on Lightedge Cloud Services. Service Failure ends when the monitoring system detects a return to normal conditions.	1% for each 1 hour of Service Failure

5.3 Public Cloud

● AWS		
Service Level	Service Level Agreement	SLA Credit
SECURE GUARANTEE 100%	Lightedge will guarantee that if an AWS EC2 Instance is managed by Lightedge (a "Managed Instance"), (a) such Managed Instance will not get hacked or compromised, and (b) if such Managed Instance gets hacked or compromised, Lightedge will refund to Customer 5% of the monthly 24/7 AWS Management Fees paid to Lightedge for the hacked or compromised Managed Instance for each 30 minutes of downtime associated with correcting the security issue (up to 50% of your monthly fees for the affected Managed Instance). Downtime shall be measured from the time Customer opens a support ticket and such time as the security issue is contained, isolated, corrected or resolved. This guarantee shall only apply so long as Customer or any entity engaged, authorized or allowed by Customer to access the Managed Instance does not disrupt or disable any of Lightedge's security features, nor take any action or allow any action to be taken which compromises the security of the Managed Instance.	
UPTIME GUARANTEE 100%	Lightedge will guarantee that if an AWS EC2 & RDS instance is managed by Lightedge (a "Managed Instance") and Customer and Lightedge have agreed to implement and operate the Managed Instance according to Amazon's Best Practices for Fault- Tolerance and High-Availability (the "HA Managed Instance") as found at: https://docs.aws.amazon.com/whitepapers/latest/real-time-communication-on-aws/high-availability-and-scalability-on-aws.html such that the Managed Instance is replicated into another Amazon Web Services Availability Zone ("AWS AZ"), Lightedge shall guarantee that the HA Managed Instance shall be available 100% of the time. Should Customer not experience these uptime levels, we will refund Customer 5% of the monthly managed service fees (charged by Lightedge for our specific Managed Services and not any Amazon fees) for each 30 minutes of downtime (up to 50% of your monthly service fees for the affected HA Managed Instance. Downtime exists when your HA Managed Instance is unavailable and is measured from the time the HA Managed Instance goes down until the time it is back online, excluding any scheduled downtime or time required to failover to the alternate AWS AZ. These guarantees do not apply to: i.) any software or operating system related issues; or ii.) Managed Instances that are not HA Managed Instances, this guarantee does not apply should Amazon experience an outage that affects multiple Amazon Web Services Availability Zones or any other Amazon Web Services issue outside of the control of Lightedge.	
DATA BACKUP GUARANTEE 100%	We will guarantee that if an instance of yours, and its EBS data, is protected by our Managed Services and being backed up by us, (a) any data backups we take will be restorable to the condition that it was backed up, and (b) if such instance's data backup is not restorable to what was backed up, we will refund to you 100% of the monthly Lightedge Managed Services fees (for that month) for the instance whose data backup was not restorable to its backed up condition. This guarantee shall only apply so long as you do not disrupt or disable any of our data backup capabilities, nor take any action or allow any action to be taken which compromises our ability to perform data backups. Should Lightedge restore backed up data to its prior state, which happened to be corrupted or unusable prior to us backing it up, then this guarantee shall not apply.	

● Azure		
Service Level	Service Level Agreement	SLA Credit
SECURE GUARANTEE 100%	Lightedge will guarantee that if an Azure Virtual Machine ("Azure VM") is managed by Lightedge (a "Managed VM"), (a) such Managed VM will not get hacked or compromised, and (b) if such Managed VM gets hacked or compromised, Lightedge will refund to Customer 5% of the monthly 24/7 Azure Management Fees paid to Lightedge for the hacked or compromised Managed VM for each 30 minutes of downtime associated with correcting the security issue (up to 50% of your monthly fees for the affected Managed VM). Downtime shall be measured from the time Customer opens a support ticket and such time as the security issue is contained, isolated, corrected or resolved. This guarantee shall only apply so long as Customer or any entity engaged,	

	authorized or allowed by Customer to access the Managed VM does not disrupt or disable any of Lightedge's security features, nor take any action or allow any action to be taken which compromises the security of the Managed VM.
UPTIME GUARANTEE 100%	Lightedge will guarantee that if a VM is managed by Lightedge (a "Managed VM") and Customer and Lightedge have agreed to implement and operate the Managed VM according to Azure's Best Practices for Fault- Tolerance and High-Availability (the "HA Managed VM") as found at: https://azure.microsoft.com/en-us/resources/resilience-in-azure-whitepaper/ that the Managed VM is replicated into another Azure Region, Lightedge shall guarantee that the HA Managed VM shall be available 100% of the time. Should Customer not experience these uptime levels, we will refund Customer 5% of the monthly managed service fees (charged by Lightedge for our specific Managed Services and not any Azure fees) for each 30 minutes of downtime (up to 50% of your monthly service fees for the affected HA Managed VM. Downtime exists when your HA Managed VM is unavailable and is measured from the time the HA Managed VM goes down until the time it is back online, excluding any scheduled downtime or time required to failover to the alternate Azure Region. These guarantees do not apply to: i.) any software or operating system related issues; or ii.) Managed VM's that are not HA Managed VM's, this guarantee does not apply should Azure experience an outage that affects multiple Azure Regions or any other Azure issue outside of the control of Lightedge.
DATA BACKUP GUARANTEE 100%	We will guarantee that if a VM of yours, and its managed disk, is protected by our Managed Services and being backed up by us, (a) any data backups we take will be restorable to the condition that it was backed up, and (b) if such VM's data backup is not restorable to what was backed up, we will refund to you 100% of the monthly Lightedge Managed Services fees (for that month) for the VM whose data backup was not restorable to its backed up condition. This guarantee shall only apply so long as you do not disrupt or disable any of our data backup capabilities, nor take any action or allow any action to be taken which compromises our ability to perform data backups. Should Lightedge restore backed up data to its prior state, which happened to be corrupted or unusable prior to us backing it up, then this guarantee shall not apply.

5.4 Cloud Management SLAs.

Service Level commitment for helpdesk response times of any Cloud Services with Cloud Management:

● CLOUD MANAGEMENT (FULL ONLY)		
Service Level	Service Level Agreement	SLA Credit
As per Helpdesk SLA (Section 5.0)	Lightedge offers an optional Cloud Management Service where Lightedge takes on additional responsibilities such as monitoring and management of the operating system running on the Customer's Cloud Services. Response times for Service Incidents related to Cloud Services with Cloud Management are the same as defined under Section 5 Helpdesk SLAs, however Lightedge offers an increased SLA remedy for such Service Failures. Response is measured as described in Section 5.0 Helpdesk SLAs.	2x Helpdesk remedies

5.5 Storage Availability SLAs.

Service Level commitment for availability and performance of Lightedge Storage Services:

● SHARED SAN ● DEDICATED SAN ● OBJECT STORAGE		
Service Level	Service Level Agreement	SLA Credit
STORAGE SERVICE AVAILABILITY 100%	DESCRIPTION: Lightedge maintains Storage Service Platforms with a variety of redundancies such as electrical, network, redundant or scale-out controllers and RAID or erasure coding storage protections to ensure uninterrupted service to Customer. Lightedge Storage Services are designed and maintained for mission critical use. All key infrastructure components such as electrical and network connectivity are maintained to N+1 redundancy. MEASURED BY: Storage Service is measured every 5 minutes by Lightedge via remote monitoring system that monitors hundreds of Service components such as temperature, failed hardware	5% for each 30 minutes of Service Failure

	components, network availability, service logs and software faults on all Storage Service components. Service Failure occurs when Customer is unable to utilize Lightedge Storage Service due to a problem with the Storage Service infrastructure or software provided by Lightedge. <i>For example, a failure of a Lightedge storage node resulting in Customer being unable to access data stored on this Service would constitute a Service Failure. A network issue on Customer network resulting in Customer being unable to access this Service would not constitute a Service Failure.</i> Service Failure begins when the monitoring system detects a Critical or Major event that results in an inability by Customer to utilize Lightedge Storage Service. Service Failure ends when the monitoring system detects a return to normal conditions.	
STORAGE SERVICE DURABILITY 100%	DESCRIPTION: Lightedge maintains Storage Service Platforms with data storage redundancies to ensure data is not corrupted or lost while stored on Service Platform. Lightedge Storage Services are designed and maintained for mission critical use. Storage is protected using RAID or erasure coding. MEASURED BY: Data loss due to issues, corruption or bit rot due caused by the physical infrastructure. Service Platform is monitored in real-time for issues such as drive loss, SMART errors and other issues that could cause data loss. Serious infrastructure failures are logged by system when they occur. Service Failure occurs when Customer data is corrupted due to Service Platform issues while data resides on Service Platform resulting in data un-recoverability by Customer. <i>For example, corruption of data while being stored on the Service Platform would constitute a Service Failure. Corruption of data due to network or Customer issues such as ransomware would not constitute a Service Failure.</i> Service Failure occurs when the monitoring system detects, or when Lightedge Operations has verified, a data corruption event that results in unrecoverable data by Customer.	1% for each Service Failure
STORAGE SERVICE PERFORMANCE 100%	DESCRIPTION: Lightedge maintains Storage Service Platforms in good working order and guarantees not to significantly oversubscribe or overload infrastructure, causing performance impact to Customer use of the Service. MEASURED BY: Storage performance is measured every 5 minutes via remote monitoring system that measures a variety of Service components such as storage controllers, disk performance, and network performance statistics on all Storage Service components. Service Failure occurs when Storage Service infrastructure is available but is adversely and substantially degraded such that Customer experiences an impact to their Service. <i>For example, an oversubscribed Lightedge storage array resulting in Customer not having reliable access to their data at reasonable speeds would constitute a Service Failure. A Customer exhausting their purchased space would not constitute a Service Failure.</i> Service Failure begins when the monitoring system detects an abnormal or rise above normal thresholds that results in a substantial and adverse impact to Customer access to data stored on Lightedge Storage Services. Service Failure ends when the monitoring system detects a return to normal conditions.	1% for each 1 hour of Service Failure

THIS SPACE INTENTIONALLY LEFT BLANK

INFRASTRUCTURE SERVICES SLAs

6. Infrastructure Services.

"Infrastructure Services" include Bare Metal and Flex Cloud compute services which are primarily managed by Customer:

6.1 Colocation Planned Maintenance

Planned Maintenance for Infrastructure Services will be limited as follows:

Schedule: only performed between 12AM - 4AM Central Time
Notification: minimum of one (1) week advance notification to Customer

6.2 Infrastructure Replacement SLAs.

Service Level commitment for availability of Lightedge Infrastructure Services:

● BARE METAL		● FLEX CLOUD
Service Level	Service Level Agreement	SLA Credit
BARE METAL INFRASTRUCTURE REPLACEMENT 4 hours	DESCRIPTION: Lightedge maintains Infrastructure Service Platform hardware in good working order and guarantees replacement hardware will be readily available in the event of a hardware failure. Infrastructure replacement measures the availability of replacement hardware. Customer is responsible for managing the availability of any workload running on Infrastructure Services and restoration of such workload is not included in this SLA. <i>For example, if Customer is contracted for 3 Bare Metal servers, and 1 of them failed resulting in disruption to Customer workload, SLA measures time to replace the failed server and return the Customer to the 3 servers contracted for. It is expected that Customer has built enough redundancies into their design to minimize the impact of hardware failures to their workload.</i> MEASURED BY: Service Failure occurs when infrastructure provided by Lightedge to Customer develops a hardware fault or fails, substantially and adversely affecting Infrastructure Service. Service Failure begins when Lightedge is first made aware of hardware failure, either by Lightedge monitoring (for Managed Compute hosts) or by Customer Ticket (for Unmanaged Compute hosts). Service Failure ends when repairs are made to infrastructure or replacement infrastructure is made available to Customer for use. The infrastructure replacement SLA is provided during periods of normal operations and is exempt during periods of Force Majeure as defined in the Master Agreement.	> 4 hours = 5% credit > 12 hours = 10% credit > 24 hours = 25% credit

THIS SPACE INTENTIONALLY LEFT BLANK

CONNECTIVITY SERVICES SLAs

7. Connectivity Services.

"Connectivity Services" include any networking services managed by Lightedge including Internet, Inter-Data Center Backbone, Inter-Data Center Connects and Ports, Cloud Ports, Metro Connects and Data Center Ports.

7.1 Connectivity Planned Maintenance

Planned Maintenance for Connectivity Services will be limited as follows:

Schedule: only performed between 12AM - 4AM Central Time
Notification: minimum of one (1) week advance notification to Customer

7.2 Internet SLAs.

Service Level commitment for availability of Lightedge Internet Service:

● INTERNET		
Service Level	Service Level Agreement	SLA Credit
INTERNET SERVICE AVAILABILITY 100%	DESCRIPTION: Lightedge maintains multiple Internet peerings with a variety of Internet transit providers across a highly available, regionally resilient, inter data center network backbone. MEASURED BY: Internet availability is measured every 5 minutes via remote monitoring system that monitors network path between each Lightedge core router and each Internet peering. Service Failure occurs when Customer is unable to utilize Lightedge Internet Service to route packets to one of Lightedge's upstream transit providers due to a problem with the Lightedge core routing infrastructure or Internet peerings. Measurement is to any available transit provider peering and does not cover availability of entire Internet. <i>For example, a failure of the Lightedge peering with Level 3 that results in Customer being unable to get to any location on the Internet would constitute a Service Failure. A failure within the Internet, outside of Lightedge direct control, that results in Customer being able to get to most destinations on the Internet but not able to get to a particular destination, would not constitute a Service Failure of Lightedge's Internet service.</i> Service Failure begins when the monitoring system detects a complete loss of ICMP traffic or loss of Internet peering across any Lightedge peering that results in an inability by Customer to substantially utilize Lightedge Internet service. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 1 hour of Service Failure
DDoS SERVICE AVAILABILITY Mitigation: 2 hours	DESCRIPTION: Lightedge will maintain DDoS mitigation equipment at each Internet transit peering to mitigate volumetric, distributed denial-of-service attacks for Customers purchasing the optional DDoS Protection Service with Lightedge provided Internet Service. The DDoS mitigation equipment typically detects attacks in real-time. MEASURED BY: Availability of Customer's Internet after DDoS scrubbing is measured every 5 minutes via remote monitoring system that monitors the DDoS mitigation equipment. This monitoring shows the effect of scrubbing performed by the mitigation equipment. Service Failure occurs when DDoS scrubbing is not functioning while Customer is under DDoS attack, resulting in the loss of Internet Service Availability by the Customer. Service Failure begins when DDoS attack on Customer begins as noted by Service Ticket. Service Failure for mitigation SLA ends when equivalent security control or DDoS "scrubbing" has been implemented and noted by timestamp in Service Ticket. Service Failure does NOT occur when network attack directed at Customer could not have been reasonably mitigated using volumetric-based protection equipment and should have been mitigated by Customer security protection equipment such as a firewall or web application firewall (WAF).	5% for each 1 hour of Service Failure

THIS SPACE INTENTIONALLY LEFT BLANK

7.3 Inter-Data Center Network SLAs.

Service Level commitment for availability and performance of the Lightedge inter-data center backbone network:

● INTER-DC CONNECTS ● INTER-DC PORTS ● METRO CONNECTS																
Service Level	Service Level Agreement	SLA Credit														
BACKBONE NETWORK AVAILABILITY 100%	DESCRIPTION: Lightedge maintains a highly resilient, meshed network between all data center locations. Customers can utilize this Inter-Data Center backbone network for communication between Lightedge data centers by purchasing Inter-Data Center Connects or Ports to connect environments between multiple Lightedge data centers. In addition, Lightedge maintains connectivity referred to as Metro Connects to other local provider buildings in some markets. MEASURED BY: Backbone network availability is measured every 5 minutes via remote monitoring system that monitors network path between data centers from one Lightedge core router to another via ICMP probe. Service Failure occurs when Customer is unable to utilize Lightedge backbone network to pass traffic between Lightedge locations due to a problem with the Lightedge core routing infrastructure or backbone circuits managed by Lightedge. <i>For example, a failure of the Lightedge backbone network resulting in Customer VMs being unable to communicate with other VMs in a different Lightedge data center would constitute a Service Failure. A customer VM being unable to communicate to a destination across a 3rd party circuit provided by Customer would not constitute a Service Failure of Lightedge's backbone network.</i> Service Failure begins when the monitoring system detects a complete loss of ICMP traffic or LDP peering across any backbone network path that results in an inability by Customer to utilize Lightedge backbone network for communication between data centers. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 30 minutes of Service Failure														
PACKET DELIVERY 100%	DESCRIPTION: Packet delivery measures the reliability and potential oversubscription of a network path. Lightedge maintains oversubscription of the inter-data center backbone network to ensure traffic can be reliably across the Lightedge backbone network. Packet delivery is measured between Lightedge core routers and does not include Customer equipment. MEASURED BY: Packet delivery is measured every 5 minutes via remote monitoring system that monitors network path between data centers from one Lightedge core router to another via ICMP probe. Service Failure occurs when Customer use of the Lightedge backbone network is substantially degraded due to a problem with the Lightedge core routing infrastructure or backbone circuits managed by Lightedge. Service Failure begins when the monitoring system detects packet loss across any backbone network path that results in impact to Customer utilization of the Lightedge backbone network for communication between data centers. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 30 minutes of Service Failure														
PACKET LATENCY as per chart	DESCRIPTION: Packet latency measures the potential oversubscription of a network path by measuring the length of time for data packets to get from one location to another. Packet latency is measured between Lightedge core routers and does not include latency added by Customer equipment. <table><tr><td>Altoona, IA to Austin, TX</td><td>< 25 msecs</td></tr><tr><td>Altoona, IA to Kansas City, MO</td><td>< 10 msecs</td></tr><tr><td>Altoona, IA to Omaha, NE</td><td>< 5 msecs</td></tr><tr><td>Austin, TX to Kansas City, MO</td><td>< 15 msecs</td></tr><tr><td>Austin, TX to Omaha, NE</td><td>< 25 msecs</td></tr><tr><td>Kansas City, MO to Omaha, NE</td><td>< 5 msecs</td></tr><tr><td>Metro Connects (same metro)</td><td>< 5 msecs</td></tr></table> MEASURED BY: Packet latency is measured every 5 minutes via remote monitoring system that monitors network path between data centers from one Lightedge core router to another via ICMP probe.	Altoona, IA to Austin, TX	< 25 msecs	Altoona, IA to Kansas City, MO	< 10 msecs	Altoona, IA to Omaha, NE	< 5 msecs	Austin, TX to Kansas City, MO	< 15 msecs	Austin, TX to Omaha, NE	< 25 msecs	Kansas City, MO to Omaha, NE	< 5 msecs	Metro Connects (same metro)	< 5 msecs	5% for each 30 minutes of Service Failure
Altoona, IA to Austin, TX	< 25 msecs															
Altoona, IA to Kansas City, MO	< 10 msecs															
Altoona, IA to Omaha, NE	< 5 msecs															
Austin, TX to Kansas City, MO	< 15 msecs															
Austin, TX to Omaha, NE	< 25 msecs															
Kansas City, MO to Omaha, NE	< 5 msecs															
Metro Connects (same metro)	< 5 msecs															

	Service Failure occurs when Customer use of the Lightedge backbone network is substantially degraded due to a problem with the Lightedge core routing infrastructure or backbone circuits managed by Lightedge. Service Failure begins when the monitoring system detects high packet latency across any backbone network path that results in substantial impact to Customer utilization of the Lightedge backbone network for communication between data centers. Service Failure ends when the monitoring system detects a return to normal conditions.	
--	---	--

7.4 Data Center Network SLAs.

Service Level commitment for availability of the Lightedge data center switching network:

● DATA CENTER CONNECTS		
Service Level	Service Level Agreement	SLA Credit
DATA CENTER NETWORK AVAILABILITY 100%	DESCRIPTION: Lightedge maintains a highly resilient, redundant data center switching network in each data center. Customers can utilize this data center switching network for communication between their equipment by purchasing Data Center Ports instead of using their own Ethernet switch. For Data Center Ports that are redundant, ethernet handoff to Customer is delivered as a group of Ethernet cables from a multiple Lightedge data center Ethernet switches in a redundant port channel configuration. Network is always available (meaning connectivity will always be available on EITHER member of the network port group even during Planned Maintenance). MEASURED BY: Service availability is measured every 5 minutes via remote monitoring system that monitors network path between core switches and edge switches via ICMP probe. Service Failure occurs when Customer is unable to pass network traffic between any two points within a Lightedge data center across Lightedge's data center switching network. <i>For example, a failure of the Lightedge switching network resulting in two co-located Customer servers being unable to communicate with each other would constitute a Service Failure. A customer server being unable to communicate to another server due to the failure of Customer switching would not constitute a Service Failure of Lightedge's data center network.</i> Service Failure begins when the monitoring system detects a complete loss of ICMP traffic or peering to any data center switch that resulting in an inability by Customer to utilize Lightedge data center network for communication between their equipment. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 30 minutes of Service Failure

7.5 Cloud Port SLAs.

Service Level commitment for availability of the Lightedge Cloud Port Service:

● CLOUD PORTS		
Service Level	Service Level Agreement	SLA Credit
CLOUD PORT AVAILABILITY 99.99%	DESCRIPTION: Lightedge maintains peering routers which connect to a variety of private network connectivity providers. Cloud Ports are a Lightedge managed connection of these network circuits into Customer collocation or cloud services from Lightedge. MEASURED BY: Cloud Port availability is measured every 5 minutes via remote monitoring system that monitors network path between each Lightedge backbone network and Customer Cloud Port. Service Failure occurs when Customer is unable to utilize Lightedge Cloud Port to pass network traffic from the Customer environment at Lightedge to the Lightedge side of the Cloud Port. Measurement is to the specific Cloud Port used by Customer but does not cover availability of the network providers network. <i>For example, a failure of the Lightedge backbone network that results in Customer being unable to pass traffic across their Cloud Port connected to a Level 3 MPLS network would constitute a Service Failure. A failure of Level 3, within their own network and outside of Lightedge direct control, would not constitute a Service Failure of Lightedge's Cloud Port service.</i> Service Failure begins when the monitoring system detects a complete loss of ICMP traffic or loss of network traffic to Cloud Port that results in an inability by Customer to substantially utilize Cloud Port service. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 30 minutes of Service Failure in excess of SLA

THIS SPACE INTENTIONALLY LEFT BLANK

SECURITY SERVICE SLAs

8. Security Services.

"Security Services" include managed security infrastructure services such as Managed Firewalls, Enterprise Load Balancing and Private Application Delivery. Lightedge manages the infrastructure availability of these services and Customer or partner manages the security information and events (SIEM).

8.1 Security Planned Maintenance

Planned Maintenance for Security Services will be limited as follows:

Schedule: only performed between 12AM - 4AM Central Time
Notification: minimum of one (1) week advance notification to Customer

8.2 Security Availability SLAs.

Service Level commitment for availability of Lightedge Security services:

● MANAGED FIREWALLS ● ENTERPRISE LOAD BALANCING ● PRIVATE APPLICATION DELIVERY		
Service Level	Service Level Agreement	SLA Credit
SECURITY INFRASTRUCTURE AVAILABILITY 100%	DESCRIPTION: Lightedge maintains Security Services hardware in good working order and guarantees replacement hardware will be readily available in the event of a hardware failure. MEASURED BY: Security infrastructure availability is measured every 5 minutes via remote monitoring system that monitors hundreds of Service components such as temperature, failed hardware components, network availability, service logs and software faults on all Security Service components. Service Failure occurs when Customer is unable to utilize Lightedge Security Services due to a failure of the security infrastructure provided by Lightedge. <i>For example, a failure of a Lightedge firewall resulting in Customer being unable to send traffic across the firewall would constitute a Service Failure. A network issue resulting in firewall being unreachable would not constitute a Service Failure.</i> Service Failure begins when the monitoring system detects a Critical or Major event that results in an inability by Customer to utilize Lightedge Security Services, or such event results in the loss or complete disruption of network traffic traversing Lightedge Security Service infrastructure. Service Failure ends when the monitoring system detects a return to normal conditions.	5% for each 30 minutes of Service Failure
CUSTOMER PREMISE HARDWARE REPLACEMENT Next business day	DESCRIPTION: Lightedge maintains Security Services hardware in good working order and guarantees replacement hardware will be readily available in the event of a hardware failure for any hardware deployed outside a Lightedge data center. MEASURED BY: Security infrastructure replacement measures the availability of replacement hardware for Security devices deployed in a non-redundant manner. Whenever possible, Lightedge recommends customer deploy Security devices in a redundant manner. Customer is responsible for managing the availability of any workload protected by Security Services and restoration of such workload is not included in this SLA. <i>For example, if Customer is contracted for a single firewall. It is expected that Customer has built enough redundancies into their design to minimize the impact of such hardware failures to their workload.</i> Service Failure occurs when infrastructure provided by Lightedge to Customer develops a hardware fault or fails, substantially and adversely affecting Security Service. Service Failure begins when Lightedge is first made aware of hardware failure, either by monitoring or by Customer Ticket. Service Failure ends when repairs are made to infrastructure or replacement infrastructure is made available to Customer for use. The infrastructure replacement SLA is provided during periods of normal operations and is exempt during periods of Force Majeure as defined in the Master Agreement.	> 1 day = 10% credit > 2 days = 25% credit > 3 days = 50% credit

8.3 Security Reseller SLAs.

Lightedge resells Security solutions provided by partners. For these resold solutions, any SLA is provided directly by the partner to the Customer. Lightedge will assist in the management of Reseller SLAs but the supplier defines and manages the Service Level Agreement.

Service Level commitment for resold Security Reseller solutions:

● VIRTUAL SOC BY IBM	● VIRTUAL SOC BY QRADAR
Service	Service Level Agreement
Virtual SOC by IBM	Refer to the IBM's VSOC SLA
Virtual SOC by QRadar	Refer to the CarbonHelix Statement of Work (SOW) for SLA

BUSINESS CONTINUITY SLAs

9. Business Continuity Services.

"Business Continuity Services" include backup and replication services completely managed by Lightedge such as Managed Backup & Recovery; VPC, DPC & Power Data Protection; VPC, DPC & Power Disaster Recovery. It also includes Lightedge offsite storage solutions used in conjunction with Customer managed backups such as Cloud Data Vault and offsite facilities reservations for use by Customer during disaster such as Workplace Recovery.

9.1 Business Continuity Planned Maintenance

Planned Maintenance for Business Continuity Services will be limited as follows:

MANAGED BACKUP & RECOVERY VPC, DPC & POWER DATA PROTECTION CLOUD DATA VAULT	VPC, DPC & POWER DISASTER RECOVERY	WORKPLACE RECOVERY
Schedule: only performed between 1PM – 5PM CDT Notification: minimum of one (1) week advance notification to Customer	Schedule: only performed between 12AM - 4AM CDT Notification: minimum of one (1) week advance notification to Customer	Schedule: performed anytime outside of occupancy Notification: minimum of one (1) week advance notification to Customer

9.2 Data Protection and Disaster Recovery Availability SLAs.

Service Level commitment for availability of all Lightedge Backup, Data Protection and Disaster Recovery Services:

● MANAGED BACKUP & RECOVERY		● VPC, DPC & POWER DATA PROTECTION	
● VPC, DPC & POWER DISASTER RECOVERY		● CLOUD DATA VAULT	
Service Level	Service Level Agreement		SLA Credit
BUSINESS CONTINUITY SERVICE AVAILABILITY 100%	DESCRIPTION: Lightedge maintains Business Continuity Services with a variety of redundancies such as electrical, network, compute failover and storage protection to ensure uninterrupted service to Customer. All key infrastructure components such as electrical and network connectivity are maintained to N+1 redundancy. Storage is protected using RAID or similar technologies.		
	Lightedge regularly monitors status of backup and replication jobs to ensure continuous protection of Customer data.		
	MEASURED BY: Service availability is measured by monitoring scheduled or continuous backup or replication jobs for successful completion on a regular basis.		5% for each Service Failure of Data Protection Services
	Service Failure occurs when Business Continuity Services do not reliably protect Customer data.		10% for each Service Failure of Disaster Recovery Services
	<i>For example, a failure of a Lightedge Service Platform component that causes a regularly scheduled backup or replication to fail, be missed, or be corrupted, would constitute a Service Failure. A failure on Customer VM causing a backup to fail would not constitute a Service Failure.</i>		
	Service Failure occurs:		
	a) when scheduled backups, scheduled replication or continuous replication has failed for greater than one ("1") business day, due to a Service Issue or Problem;		
	OR		

	b) for each backup or replication job which cannot be restored due to data corruption or other Lightedge fault within the control of Lightedge. Service Failure does NOT occur: a) when files or workload cannot be successfully protected due to an issue on Customer side such as a locked file or network issues outside of Lightedge control; b) when protected data that has been offloaded to external media by request of Customer and leaves the possession of Lightedge cannot be restored.	
--	--	--

9.3 Disaster Recovery RTO & RPO SLAs.

Service Level commitment for Recovery Time (RTO) and Recovery Point (RPO) of Lightedge Disaster Recovery Services:

	VPC, DPC & POWER STANDARD DR	VPC, DPC & POWER PREMIUM DR
RECOVERY TIME DESCRIPTION: RTO (or Recovery Time Objective) refers to the ability to promptly restore Customer data or workload from a Lightedge Disaster Recovery Service when it is needed by Customer during a Declared Event. MEASURED BY: Recovery Time SLAs are measured by the duration of time elapsing between when Service Incident or Declared Event is created (e.g. the timestamp of the Ticket request via the Service Portal) and when the data or recovered workload is made available for use. Recovery Time SLAs do not include time taken by Customer to create a Declared Event, or time spent executing runbook or recovery steps unrelated to the Lightedge Disaster Recovery Service. Service Failure occurs when Lightedge is not able to restore Customer data or workload within the timeframe defined by this SLA. Lightedge Recovery Cloud Service must be contractually reserved by Customer and adjacent the data being protected. Recovery Time SLAs will only apply for workload that can be recovered within the pre-committed Recovery Cloud.	2 hours	15 minutes
RECOVERY POINT (RPO) DESCRIPTION: RPO (or Recovery Point Objective) refers to the age of the Customer's data and the ability to restore a recent or up-to-date version of that protected data whenever needed by Customer during a Declared Event. <i>For example, a failure of a Lightedge backup platform or service platform component that causes replications to fail or becoming corrupted, necessitating restoration of older data, would constitute a Service Failure. A failure on Customer VM causing replication to fail would not constitute a Service Failure.</i> MEASURED BY: Recovery Point SLAs are measured by the age of the most recent backup or replication job containing restorable data and do not include time taken by Customer to create a Declared Event or to execute runbook steps unrelated to this Service. Standard Disaster Recovery Service level is limited by the type of disk being used on the production side as defined within this SLA. For lower Recovery Point requirements, faster disk or upgrading to Premium Disaster Recovery may be required. Service Failure occurs when Lightedge is not able to restore recent versions of Customer data as defined by this SLA.	<i>Based on storage tier being protected</i> Bulk: 24 hours Capacity: 12 hours Standard: 6 hours Performance: 4 hours	15 minutes
SLA CREDITS	2.5% for each Service Failure	5% for each Service Failure of Data Protection Services

9.4 Workplace Recovery SLAs.

Service Level commitment for Lightedge Workplace Recovery seat availability:

● WORKPLACE RECOVERY		
Service Level	Service Level Agreement	SLA Credit
Dedicated Seats: 100% available within 2 hours	DESCRIPTION: Workplace Recovery is a reservation of office space for use by Customer during disasters. Workplace Recovery environment will be maintained to ASHRAE standards for Class 3 office space. MEASURED BY: Service availability is measured by the availability of the recovery space for use by Customer after Customer declares intent to use space. Service Failure occurs when seats are not available for Customer use in a timely manner. <i>For example, a failure of a Lightedge to make disaster recovery space available in a timely fashion, or to maintain the disaster recovery space to occupancy standards, would constitute a Service Failure.</i>	10% for each 6 hours of Service Unavailability in excess of SLA
Shared Seats: First come, first serve, available within 4 hours		None

	<i>Customer being unable to use the disaster recovery space due to their own cause would not constitute a Service Failure.</i> Service Failure ends when Workplace Recovery environment is made available for Customer use. Workplace Recovery SLAs are provided during periods of normal operations and are exempt during periods of Force Majeure as defined in the Master Agreement.	
--	--	--