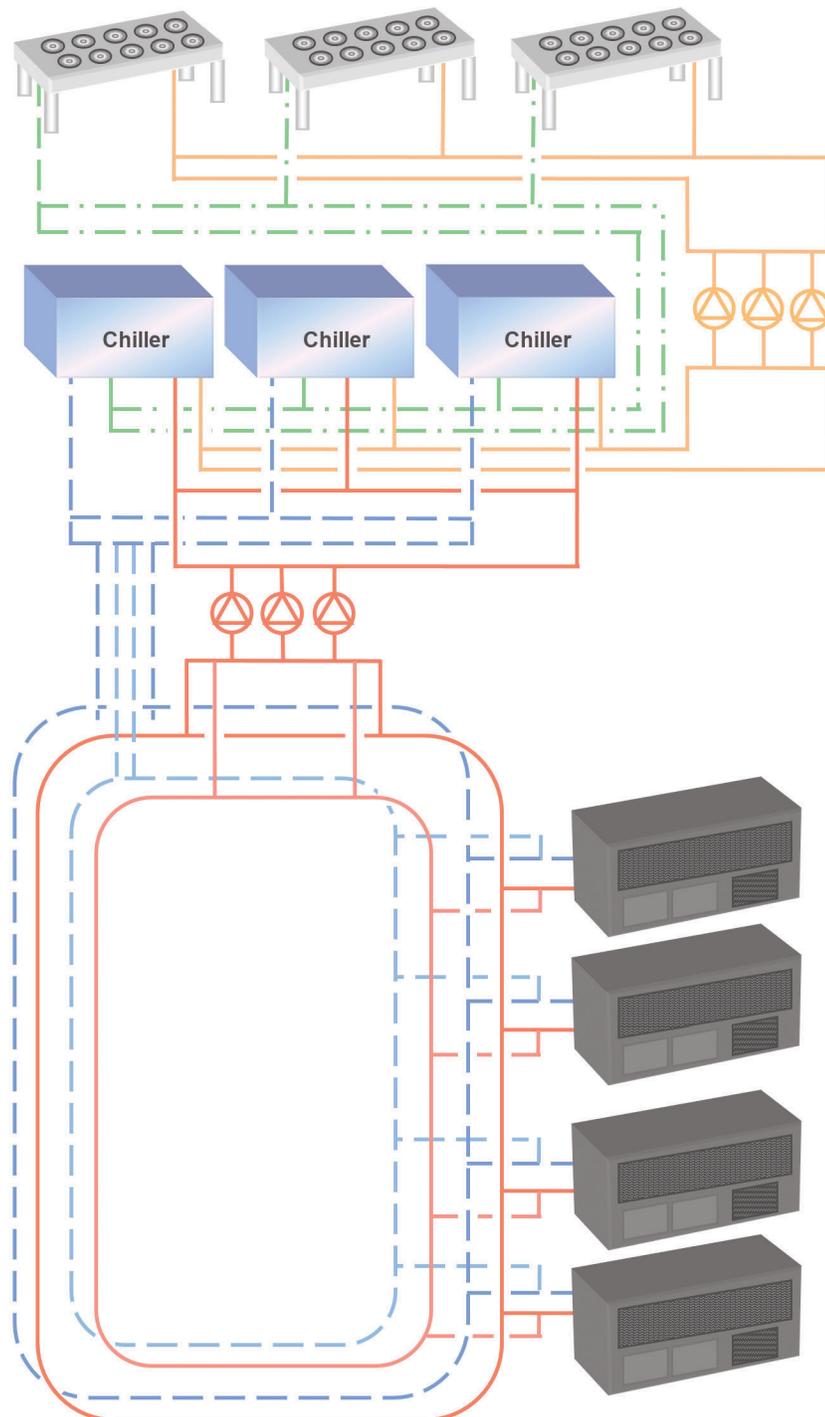


An Overview of the ANSI/BICSI 002-2019 Data Center Availability Class Methodology



The Start of Every Great Data Center is its Design

ANSI/BICSI 002-2019, *Data Center Design and Implementation Best Practices*

The foundation standard for data center design around the world, ANSI/BICSI 002-2019 continues its mission to provide requirements, guidelines and best practices applicable to any data center, from traditional, modular, and to the edge, hyperscale and beyond.

With 550 pages of content, the 2019 edition of BICSI 002 covers all aspects of data center design, including:

- Design Concepts and Planning
- Site Selection and Space Configuration
- Building Shell and Architecture
- Core Systems (e.g., Electrical, Mechanical, Network and Cable Plant)
- Facility and Building systems
- Security
- Commissioning

Written for international application by experts from all across the globe and representing all the all disciplines within data center design and implementation, BICSI 002 incorporates standards from ISO, TIA, CENELEC, ASHRAE and the latest in design concepts, such as the Open Compute Project®, enabling improved data center designs and operational results regardless of site location.

Given the comprehensive nature of BICSI 002-2019, whether you are a data center designer, operator or involved in the planning, design, construction, implementation of a data center, there is something for everyone and more within BICSI 002.



Learn more about BICSI 002 and view the first five sections for free at www.bicsi.org/002

People have come to expect ready access to information 24 hours a day, every day. The Internet as well as more traditional enterprises—both business and governmental—operate 7 days a week, 24 hours a day.

With the increased reliance on 24/7 availability of information and data processing support, the reliance on mission-critical data processing facilities has also increased.

Mission-critical data centers have not traditionally been high-profile projects, yet their design issues are increasingly complex and critical. With an emerging design terminology and vocabulary, their rapid evolution calls for an exceptional degree of building and IT systems coordination and integration. These data centers are not merely warehouses for servers; instead, they rival medical operating rooms or semiconductor plants, with their precise environmental controls and power requirements.

To increase the likelihood of success of a mission-critical facility, required performance levels of availability and reliability should be defined, prior to the start or formalization of the design, procurement, and maintenance requirements and processes. Failure to define performance and availability levels prior to the project start often yields higher construction, implementation, and operational costs as well as inconsistent and unpredictable performance.

What is Availability?

Availability is the probability that a component or system is in a condition to perform its intended function. While similar to *reliability*, availability is affected by more events than a failure requiring repair or replacement of a component or system.

While there are different formulae to calculate availability for calculations involving systems, availability, in its simplest form, is the ratio of uptime observed during a specified interval

over the total time of that interval and can be expressed as the following equation:

$$\text{Availability} = \frac{\text{Uptime within Observation Interval}}{\text{Total Time of Observation Interval}}$$

While the previous equation can generate the availability of a system, the result does not provide information which can be used to improve the observed availability value. By splitting total time into its two primary elements (uptime and downtime), the equation changes to the following:

$$\text{Availability} = \frac{\text{Uptime}}{\text{Uptime} + \text{Downtime}}$$

Going one step further, downtime itself can be split into two types: scheduled and unscheduled. When the two types of downtime are inserted into the availability equation, it results in the following:

$$\text{Availability} = \frac{\text{Uptime}}{\text{Uptime} + \text{Scheduled Downtime} + \text{Unscheduled Downtime}}$$

Thus, availability can be increased by reductions in one or both types of downtime. Examples of common scheduled and unscheduled events are shown in Table 1.

Table 1 Common Downtime Events

<i>Scheduled Downtime</i>	<i>Unscheduled Downtime</i>
Preventive maintenance System and equipment setup and upgrades System testing/optimization Scheduled facilities related events Remedial maintenance	Repairs due to failure Maintenance delay Facility-related failures/outages

Risk Analysis

It is impossible to eliminate the risk of downtime, but risk reduction is an important planning element. In an increasingly competitive world, it is imperative to address downtime in business decisions. The design of systems supporting critical IT functions depends on the interaction between the criticality of the function and its operational profile.

Criticality is defined as the relative importance of a function or process as measured by the consequences of its failure or inability to function. The operational profile expresses the time intervals over which the function or process must operate.

To provide optimal design solutions for a mission-critical data center, consider several key factors. NFPA 75 identifies seven considerations for protection of the environment, equipment, function, programming, records, and supplies in a data center. These include:

- What are the life safety aspects of the function? For example, if the system failed unexpectedly, would lives be put at risk? Examples of such applications include automated safety systems, air traffic control, and emergency call centers.
- What is the threat to occupants or exposed property from natural, man-made, or technology-caused catastrophic events?
- What would be the economic loss to the organization from the loss of function or loss of records?
- What is the access to redundant off-site processing systems (e.g., “high performance computing”, massively paralleled systems, cloud service provider, disaster recovery site, backup data center)?

- What would be the economic loss to the organization from damaged or destroyed equipment?
- What would be the impact of disrupted service to the organization’s reputation? For example, would subscribers switch to a competitors’ service?
- What would be the regulatory or contractual impact, if any? For example, if unplanned downtime resulted in loss of telephone service or electrical service to the community, would there be penalties from the government?

Data Center Availability Classes

To a great degree, design decisions are guided by the identified Availability Class. Therefore, it is essential to understand the meaning of each Availability Class before determining an Availability Class for a specific data center

Availability Class 0

The objective of Class 0 is to support the basic requirements of the IT functions without supplementary equipment. Capital cost avoidance is the major driver. There is a high risk of downtime because of planned and unplanned events.

Availability Class 1

The objective of Class 1 is to support the basic requirements of the IT functions. There is a high risk of downtime because of planned and unplanned events. However, in Class 1 data centers, remedial maintenance can be performed during nonscheduled hours.

Availability Class 2

The objective of Class 2 is to provide a level of reliability higher than that defined in Class 1 to reduce the risk of downtime because of component failure. In Class 2 data centers, there is a moderate risk of downtime as a result of planned and unplanned events. Maintenance activities can typically be performed during unscheduled hours.

Availability Class 3

The objective of Class 3 is to provide additional reliability and maintainability to reduce the risk of downtime because of natural disasters, human-driven disasters, planned maintenance, and repair activities. Maintenance and repair activities will typically need to be performed during full production time with no opportunity for curtailed operations.

Availability Class 4

The objective of Class 4 is to eliminate downtime through the application of all tactics to provide continuous operation regardless of planned or unplanned activities. All recognizable single points of failure are eliminated.

Primary Concerns that Define the Availability Class

Each Availability Class is defined in terms of four areas of concern:

- 1) **Component redundancy** increases reliability by providing redundancy for critical high-risk, low-reliability components within systems.
- 2) **System redundancy** increases reliability even more by providing redundancy at the system level.
- 3) **Quality** ensures that high quality is designed and implemented in the data center, thereby reducing the risk of downtime due to failure during initial installation and/or premature wear.
- 4) **Survivability** refers to reducing the risk of downtime by protecting against external events such as physical forces, security breaches, and natural disasters.

Table 2 summarizes how each of these four factors is defined for each of the five Availability Classes.

Table 2 Summary of Areas of Concern for Availability Class

Class	Component Redundancy	System Redundancy	Quality Control	Survivability
Class 0/1	None	None	Standard commercial quality	None
Class 2	Redundancy is provided for critical components	None	Premium quality for critical components	Moderate hardening for physical security and structural integrity
Class 3	Redundancy is required for critical and noncritical components, except where the component is part of a redundant system.	System redundancy is required where component redundancy does not exist	Premium quality for all components	Significant hardening for physical security and structural integrity
Class 4	Redundancy is provided for all critical components and to increase maintainability; also provided for noncritical components.	System redundancy is provided with component redundancy so that overall reliability is maintained even during maintenance activities.	Premium quality for all components. Recommended to use different lots, model, and manufacturer to avoid common fault or recall.	All systems are self-supporting in any event and are protected against the highest levels of natural forces.

Availability Class Sub Groups

The data center is not just a facility or building, but it is a collection of services that supports the critical business processes. The data center services Availability Class model can be used to guide design and operational decisions for the following critical services:

- Facility: The facility systems (e.g., power, cooling, controls) can be categorized into one of the sub group Class F0 through Class F4.
- Cable Plant: The network cable plant topology can be categorized into one of the sub group Class C0 through Class C4.
- Network Infrastructure: The network architecture and topology can be categorized into one of the sub group Class N0 through Class N4.
- Data Processing and Storage Systems: The computer processing and storage systems can be categorized into one of the sub group Class S0 through Class S4.
- Applications: The applications can be categorized into one of the sub group Class A0 through Class A4.

Further information, including requirements and recommendations for each of these sub-groups can be found in ANSI/BICSI 002-2014.

Application Examples

Application examples for a Class F2, F3, and F4 include:

Class F2

The critical power, cooling, and network systems would need redundancy in those parts of the system that are most likely to fail. These would include any products that have a high parts count or moving parts, such as UPS, controls, air conditioning, generators, ATS or systems that are outside the control of the data center management such as network access carrier services. In addition, it may be appropriate to specify premium quality devices that provide longer life or better reliability.

Class F3

The critical power, cooling, and network systems must provide for reliable, continuous operations even when major components (or, where necessary, major subsystems) are out of service for repair or maintenance. To protect against unplanned downtime, the power, cooling, and network systems must be able to sustain operations while a dependent component or subsystem is out of service.

Class F4

The critical power, cooling, and network systems in a Class F4 facility must provide for reliable, continuous operations even when major components (or, where necessary, major subsystems) are out of service for repair or maintenance. To protect against unplanned downtime, systems must be able to sustain operations while a dependent component or subsystem is out of service.

Determining the Data Center Availability Class

While there are innumerable factors that can be evaluated in a mission-critical data center, there are three factors that can quickly be quantified.

These factors are:

- Operational requirements,
- Operational availability, and
- Impact of downtime,

and their interaction, as shown in Figure 1, provides an easy to apply method of determining an Availability Class.

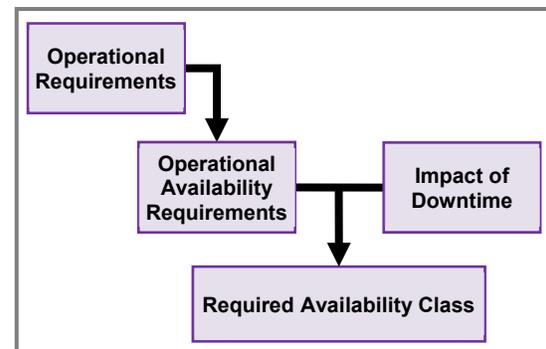


Figure 1 Interaction of Factors in Data Center Availability Class Determination

The following provides the four steps required to determine a data center's availability class.

Step 1: Identify Operational Requirements

The first step in determining the Availability Class associated with mission-critical data center services is to define the data center’s intended operational requirements. Sufficient resources must be available to achieve an acceptable level of quality over a given time period. IT functions that have a high-quality expectation over a longer time period are by definition more critical than those requiring less resources, lower quality, and/or are needed over a shorter time period.

While there are many factors in operations, the key factor to be assessed in this step is the amount of time to be provided for testing and maintenance activities that disrupt normal operation. This is often known as *planned maintenance shutdown*. Once the time for planned maintenance shutdowns is known, this value can be used within Table 3 to determine an Operational Level. The value of time used should not include projections for unplanned repairs or events.

The indicated Operational Level is then used in the next step.

Step 2: Quantify and Rank Operational Availability Requirements

The second step in determining the Availability Class is to identify the data center’s operational availability requirements, specifically the total uptime that the data center services must support without disruption. The term *availability* includes that ITE is operational and able to perform its function; it does not solely refer to operation of the supporting infrastructure.

Operational availability refers only to scheduled uptime—that is, the time during which the IT functions are actually expected to run.

These operational availability requirements are reflected by the determination of an Operational Availability rating. By using the Operational level determined in Step 1 and indexing that value with the allowed maximum annual downtime shown in Table 4, an Operational Availability Rating is indicated. The Operational Availability Rating will be used in Step 4 in conjunction with information derived in Step 3

Table 3 Identifying Operational Requirements: Time Available For Planned Maintenance Shutdown

Operational Level	Annual Hours Available for Planned Maintenance Shutdown	Description
0	> 400	Functions are operational less than 24 hours a day and less than 7 days a week. Scheduled maintenance down time is available during working hours and off-hours.
1	100-400	Functions are operational less than 24 hours a day and less than 7 days a week. Scheduled maintenance down time is available during working hours and off-hours.
2	50-99	Functions are operational up to 24 hours a day, up to 7 days a week, and up to 50 weeks per year; scheduled maintenance down time is available during working hours and off hours.
3	0-49	Functions are operational 24 hours a day, 7 days a week for 50 weeks or more. No scheduled maintenance down time is available during working hours.
4	0	Functions are operational 24 hours a day, 7 days a week for 52 weeks each year. No scheduled maintenance down time is available.

NOTE: The term *shutdown* means that operation has ceased; the equipment is not able to perform its mission during that time. Shutdown does *not* refer to the loss of system components if they do not disrupt the ability of the system to continue its mission.

Table 4 Identifying Operational Availability Rating: Maximum Annual Downtime (Availability %)

Operational Level (from Table 3)	Allowable Maximum Annual Downtime (minutes) Availability as % (Nines of Availability)				
	$x > 5000$ $x < 99\%$ (2-9's)	$5000 \geq x > 500$ $99\% \leq x < 99.9\%$ (3-9's)	$500 \geq x > 50$ $99.9\% \leq x < 99.99\%$ (4-9's)	$50 \geq x > 5$ $99.99\% \leq x < 99.999\%$ (5-9's)	$5 \geq x$ $99.999\% \leq x$ (6-9's)
Level 0	0	0	1	2	2
Level 1	0	1	2	2	2
Level 2	1	2	2	2	3
Level 3	2	2	2	3	4
Level 4	3	3	3	4	4

Of note, the cost of downtime must be weighed against the cost of mitigating risks in achieving high availability. Even an event such as a less than a second of power interruption or a few minutes of cooling interruption can result in hours of recovery time. Thus, the objective is to identify the intersection between the allowed maximum annual downtime and the intended operational level. A function or process that has a high availability requirement with a low operational profile has less risk associated with it than a similar function with a higher operational profile.

Step 3: Determine Impact of Downtime

The third step in determining the Availability Class is to identify the impact or consequences of downtime. This is an essential component of risk management because not all downtime has the same impact on mission-critical data center services. Identifying the impact of downtime on mission-

critical functions helps determine the tactics that will be deployed to mitigate downtime risk. As shown in Table 5, there are five impact classifications, each associated with a specific impact scope.

Step 4: Identify the Data Center Availability Class

The final step in determining the data center Availability Class is to combine the previously identified factors to arrive at a usable expression of availability. Since operational level is subsumed within the availability ranking, the task is to matrix the availability ranking against the impact of downtime to arrive at an appropriate Availability Class. Table 6 shows the intersection of these two values, and the resultant Data Center Availability Class.

Table 5 Classifying the Impact of Downtime on the Organization

Classification	Description – Impact of Downtime
Isolated	Local in scope, affecting only a single function or operation, resulting in a minor disruption or delay in achieving non-critical organizational objectives.
Minor	Local in scope, affecting only a single site, or resulting in a minor disruption or delay in achieving key organizational objectives.
Major	Regional in scope, affecting a portion of the enterprise (although not in its entirety) or resulting in a moderate disruption or delay in achieving key organizational objectives.
Severe	Multiregional in scope, affecting a major portion of the enterprise (although not in its entirety) or resulting in a major disruption or delay in achieving key organizational objectives.
Catastrophic	Affecting the quality of service delivery across the entire enterprise or resulting in a significant disruption or delay in achieving key organizational objectives.

Table 6 Determining Data Center Services Availability Class

Impact of Downtime (from Table 5)	Operational Availability Rating (from Table 4)				
	0	1	2	3	4
Isolated	Class 0	Class 0	Class 1	Class 3	Class 3
Minor	Class 0	Class 1	Class 2	Class 3	Class 3
Major	Class 1	Class 2	Class 2	Class 3	Class 3
Severe	Class 1	Class 2	Class 3	Class 3	Class 4
Catastrophic	Class 1	Class 2	Class 3	Class 4	Class 4

It is unlikely that a single data center would have all the applications, data processing, and storage platform systems aligned within a single reliability classification no matter what the targeted base data center reliability classification is.

Prior to virtualization, location transparent applications and cloud services, the optimal data center services configuration consisted of an alignment of the reliability classes across all the data center service layers. This provided the minimum required level of reliability and redundancy without over building any one of the data center service layers.

System designs with clustered systems having nodes spread across two or more Class 3 data centers can meet or exceed the uptime of a system in a single Class 4 data center. In such a design, the first failover is to the local node (synchronous), the second failover is to a nearby data center (~16 km [10 miles], and still synchronous), and the third is to a remote data center (but asynchronous).

Such a design does increase the facility's overhead and therefore, the cost. However, it offers a way for designers to avoid many of the costs associated with Class 4 data centers, whether owned, leased or collocated.

One of the values of the BICSI data center services reliability framework model is it can be used to:

- Identify the minimum reliability targets.
- Provide a structured methodical approach to guide decisions on how to adjust lower layer services to compensate for higher layer services reliability inadequacies.
- Guide discussions regarding the possible technical and cost benefits of increasing the reliability of the network architecture and higher layers above the targeted reliability class across multiple data centers so that cost savings can be realized by building each of the data centers facilities to a lower Class than the targeted reliability classification.

On the following pages, three examples are provided to illustrate how the framework provides for multi-data center architecture.

High Availability In-House Multi-Data Center Architecture Example

In this example, a customer has identified Class 3 as the targeted data center services reliability level. The customer has multiple facilities that can support critical data center functions. By provisioning the applications with high-availability configuration across two data center facilities, the customer will be able to achieve the targeted reliability and availability objectives.

It is important that any man-made or natural event common mode risks that may exist within the geographical region that is common between the two data centers be identified and evaluated. The communications between the two data centers can be synchronous or asynchronous, depending on the recovery point objective (RPO) and recovery time objective (RTO) of the disaster recovery/business continuity requirements and the physical distance limitations between the two data centers.

There are times when there are man-made or natural event common mode risks to both data centers that have been deemed an acceptable risk to the organization. An example would be multi-regional events, such as multi-State power outages, that an organization deems acceptable. There would be no loss of data within the data center (running on backup power sources); however, the users would not have access to the applications or data as their networks and systems would be off-line throughout the multi-state region. The organization might determine that the users would not have an expectation of accessing the data in this scenario, and there would be no loss of revenue or business reputation as a result. Therefore, the costs associated with building out multiple data centers across a wider geographical area (possibly outside synchronous communication capabilities) may not be justified.

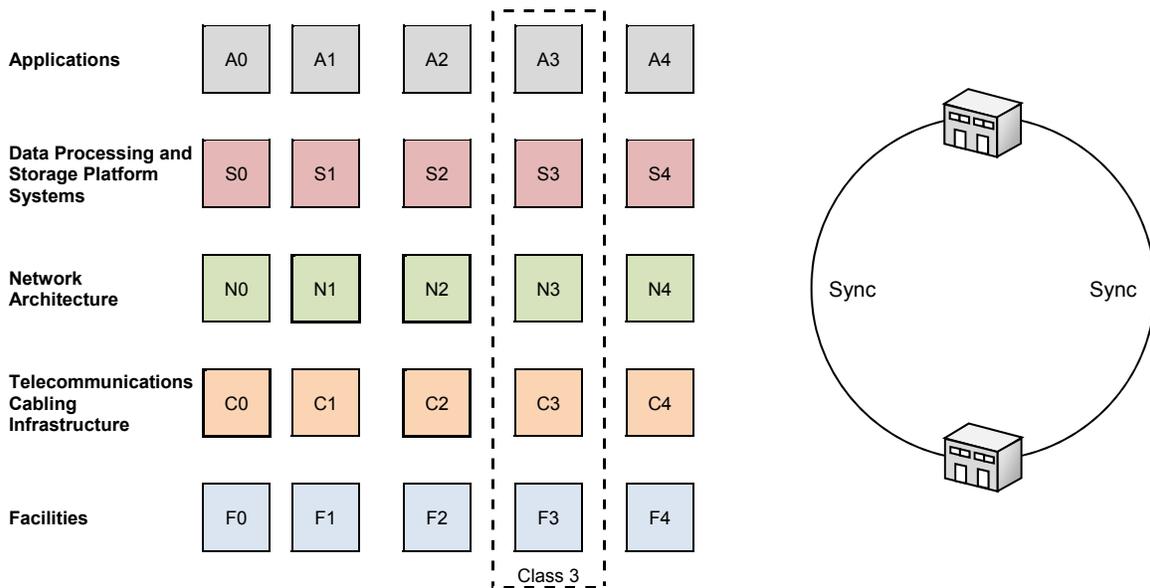


Figure 2 Multi-Data Center Class 3 Example

Private Cloud Multi-Data Center Architecture Examples

Private cloud services are implemented in customer-owned data centers. Private cloud applications are developed to improve scalability, speed of deployment, and reliability with the abstraction on the reliance on the lower layer data center services. Private cloud applications may enable the customer to implement highly reliable applications without requiring highly reliable lower layer data center services.

Private Cloud Multi-Data Center Architecture – Class 3 Solution/Three Class 2 Facilities

The first example is a customer that has identified at least two Class 3 data centers as the targeted data center services reliability level. The private cloud applications would be implemented across diverse geographical

regions. By provisioning the private cloud applications across three Class 2 data center facilities, the customer may be able to achieve similar reliability and availability objectives. The applications can move around each of the data center facilities with the loss of any one facility having little or no impact on the enterprise.

The two data centers connected via synchronous communications would be located within a common region. The data center that is connected via asynchronous communications would be located outside the region, ensuring no natural or man-made event represents a common mode of failure.

This example is not provided as a solution that will always equate to two Class 3 data centers; rather, it is provided to show how the data center services reliability framework can be used to evaluate various options.

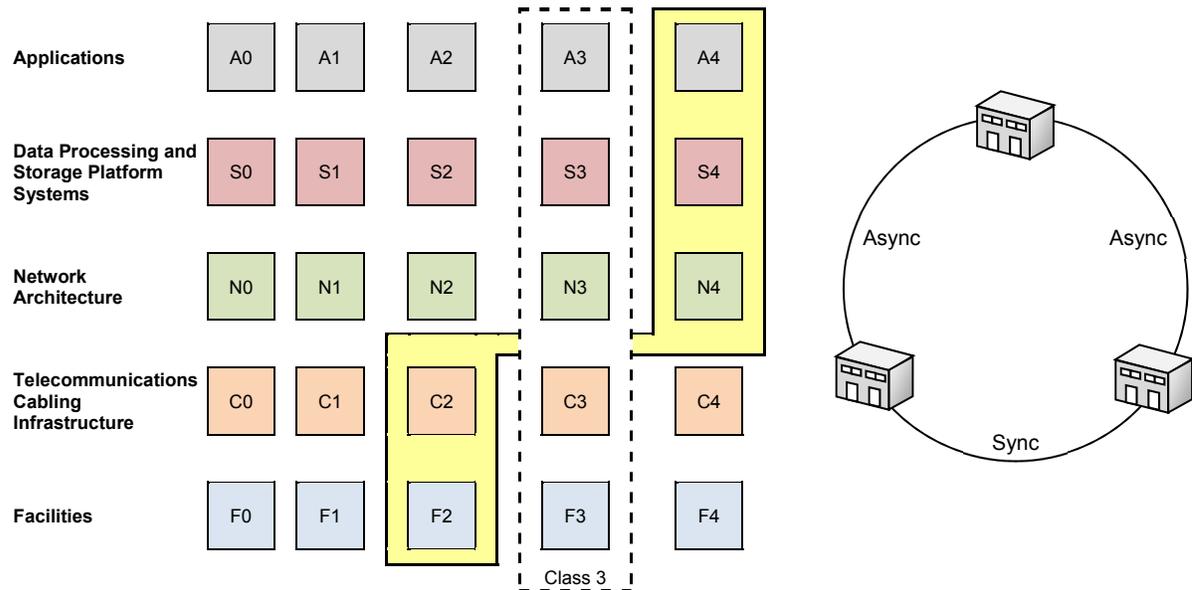


Figure 3 Multi-Data Center Class 3 Example With Three Class 2 Facilities

Private Cloud Multi-Data Center Architecture – Class 4 Solution/Four Class 2 Facilities

The second example is a customer that has identified two Class 4 data centers as the targeted data center services reliability level. By provisioning the private cloud applications across four Class 2 data center facilities, both within a common region and outside common regions, the customer may be able to achieve similar reliability and availability objectives. The applications can move around each of the data center facilities with the loss of any one facility or facilities within a region having little or no impact on the enterprise.

Two of the data centers are connected via synchronous communications located within a common region. The pair of data centers located within each common region are connected via asynchronous communications. The pair of data centers would be located outside each other's region, ensuring no natural or man-made event represents a common mode of failure.

This example is not provided as a solution that will always equate to two Class 4 data centers, but it is provided to show how the data center services reliability framework can be used to evaluate various options.

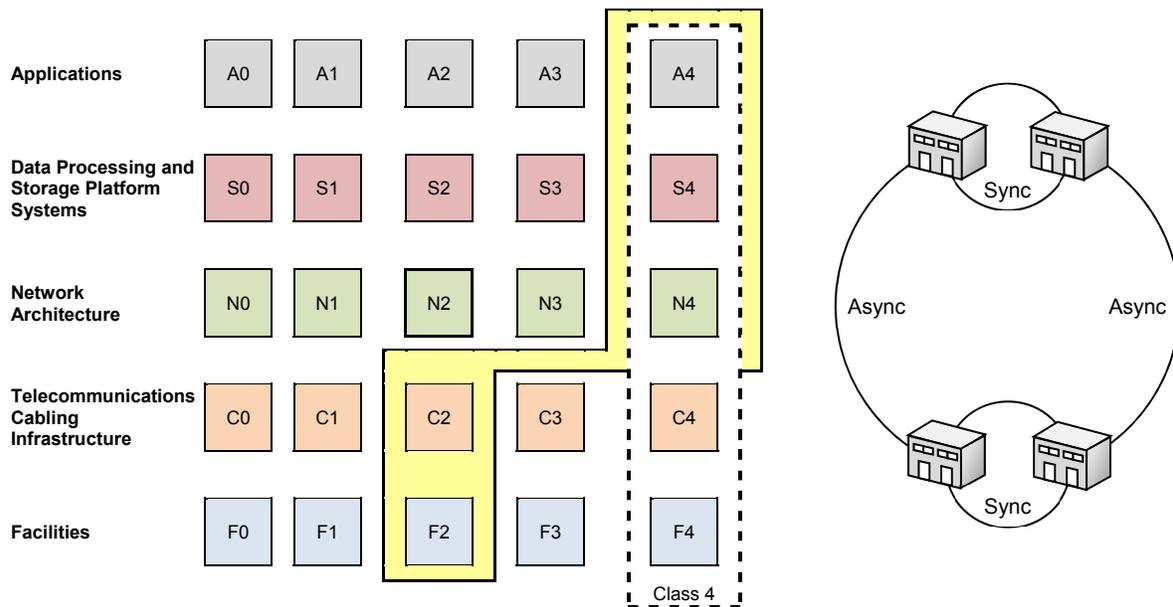


Figure 4 Multi-Data Center Class 4 Example with Four Class 2 Facilities

Reliability Planning Worksheet

Use the following planning guide to determine the data center availability class.

Project name: _____
Project number: _____
Project description: _____

Project location: _____

STEP 1: Determine Operational Requirements

- 1) How many hours of operation must be supported during a production week? _____
- 2) How many scheduled production weeks are there? (if production occurs every week enter 52.14)

- 3) Multiply line 1 by line 2, and enter here. This is annual production hours: _____
- 4) Subtract line 3 from 8,760, and enter the result here: _____
- 5) Are there additional available days or weekends each year for scheduled downtime that have not been accounted for in lines 2 or 3? Enter the total annual available hours: _____
- 6) Add lines 4 and 5 and enter the result (allowable annual maintenance hours) here: _____
- 7) If line 6 is greater than 400, the Operational Level is 0; otherwise, proceed to the next line.
- 8) If line 6 is greater 100, the Operational Level is 1; otherwise, proceed to the next line.
- 9) If line 6 is between 50 and 99, the Operational Level is 2; otherwise, proceed to the next line.
- 10) If line 6 is between 1 and 49, the Operational Level is 3; otherwise, the Operational Level is 4.

STEP 2: Determine Operational Availability Rank.

- 1) Based on the operational level from Step 1 above:
 - Level 0; Proceed to line 2.
 - Level 1; Proceed to line 3.
 - Level 2; Proceed to line 4.
 - Level 3; Proceed to line 5.
 - Level 4; Proceed to line 6.
- 2) Operational Level 0: If the maximum annual downtime is:
 - 500 minutes or greater, then the availability requirement is Operational Availability Rank 0.
 - Between 50 and 500 minutes, then the availability requirement is Operational Availability Rank 1.
 - Less than 50 minutes, then the availability requirement is Operational Availability Rank 2.Proceed to Step 3.
- 3) Operational Level 1: If the maximum annual downtime is:
 - 5000 minutes or greater, then the availability requirement is Operational Availability Rank 0.
 - Between 500 and 5000 minutes, then the availability requirement is Operational Availability Rank 1.
 - Less than 500 minutes, then the availability requirement is Operational Availability Rank 2.Proceed to Step 3.

<Worksheet continues on the next page>

- 4) Operational Level 2: If the maximum annual downtime is:
 - 5000 minutes or greater, then the availability requirement is Operational Availability Rank 1.
 - Between 5 and 5000 minutes, then the availability requirement is Operational Availability Rank 2.
 - Less than 5 minutes, then the availability requirement is Operational Availability Rank 3.
 Proceed to Step 3.

- 5) Operational Level 3: If the maximum annual downtime is:
 - 50 minutes or greater, then the availability requirement is Operational Availability Rank 2.
 - Between 5 and 50 minutes, then the availability requirement is Operational Availability Rank 3.
 - Less than 5 minutes, then the availability requirement is Operational Availability Rank 4.
 Proceed to Step 3.

- 6) Operational Level 4: If the maximum annual downtime is:
 - 50 minutes or greater, then the availability requirement is Operational Availability Rank 3.
 - Less than 50 minutes, then the availability requirement is Operational Availability Rank 4.
 Proceed to Step 3.

STEP 3: Define Mission-Critical Risk Level

Downtime will reduce or negatively impact operations (select one):

- Catastrophic (e.g., across the entire enterprise) _____
- Severe (e.g., across a wide portion of the enterprise) _____
- Major (e.g., across a single region or department) _____
- Minor (e.g., at a single location) _____
- Isolated (e.g., a single non-critical function) _____

STEP 4: Determine the Availability Class (using the following table)

- 1) Select the column from the Operational Availability Rank in Step 2.
- 2) Select the row from the Risk Level in Step 3.
- 3) Your Availability Class is where the two intersect: _____

Data Center Services Availability Class

Impact of Downtime	Operational Availability Rating				
	0	1	2	3	4
Isolated	Class 0	Class 0	Class 1	Class 3	Class 3
Minor	Class 0	Class 1	Class 2	Class 3	Class 3
Major	Class 1	Class 2	Class 2	Class 3	Class 3
Severe	Class 1	Class 2	Class 3	Class 3	Class 4
Catastrophic	Class 1	Class 2	Class 3	Class 4	Class 4

Further Your Data Center Career and Knowledge

BICSI Data Center Design Consultant (DCDC)

Data centers are more than just cabling and servers. Electrical and mechanical systems, as well as the physical building and building systems all interact to meet a data center's operational requirements and mission.

Be recognized for not only having the knowledge, but also the ability to apply your knowledge and skills in data center design. Those who have been awarded the DCDC come from different facets of the data center environment, be it data center design, construction or operations.

Relatively new to data centers? Qualifying to take the DCDC exam only requires a few years of approved and verifiable experience in the design or construction of data centers within the last six years. Hesitant because of today's data center work pace? The DCDC exam can be taken at any of 4,500 Pearson VUE Authorized Testing Centers worldwide, avoiding the need to take extended time away from the site.

Essentials of Data Center Projects, 1st Edition



Surrounding the design process is the actual project that brings a data center from idea to actuality. The *Essentials of Data Center Projects*, developed from years of data center professional experience, provides a holistic view of the overall data center project, from conceptualization and the design process, and through construction, commissioning, and closeout.

The *Essentials of Data Center Projects*, in conjunction with ANSI/BICSI 002-2019, serve as the foundation texts for the DCDC.

Learn more about all that BICSI has to offer for data centers and the DCDC at www.bicsi.org/dcdc



8610 Hidden River Parkway, Tampa, FL 33637-1000 USA

Tel: +1 813.979.1991 or 800.242.7405 (USA & Canada toll-free)

Fax: +1 813.971.4311; Email: bicsi@bicsi.org; Web: www.bicsi.org

© Copyright BICSI 2015, 2019. All Rights Reserved. BICSI is a registered trademark of BICSI, Inc.